IMPERVA®

MIGRATION GUIDE

Cloud Migration Guide

Contents

About this Guide

The Cloud Migration Guide provides direction as you begin your journey of migrating applications and data to the cloud.

Migrating to the cloud is not a simple matter. There are many possible paths and options to consider. And there are a variety of security concerns.

This guide focuses on application and data security. It shows you how to migrate your security policies from on-premises to the cloud, depending on the deployment you select. It delineates the various options available to you and helps you identify a path that makes the most sense for your site-specific needs. It also looks at how your business model may change when you make the move to the cloud and how a flexible licensing model can be central to that change.

The Cloud Migration Guide is organized into four parts:

Cloud Migration Approaches: We review the various cloud migration strategies.

Application and Data Security Concerns: We describe the "shared responsibility" model that is central to understanding each "actor's" responsibility for the security of applications and data in the cloud.

Policy Migration: We discuss migrating IT policies from an on-premises architecture to a cloud-based architecture.

Business and Technical Considerations: We discuss flexible application security requirements and various licensing models for cloud-based security solutions, including those offered by Imperva.

Cloud Migration Approaches

We are in the middle of a computing revolution. At its core is the rapid development and deployment of data-centric applications. This revolution is based on the adoption of Agile development methodologies and the formation of DevOps teams. Concurrent with this shift toward rapid development and deployment is the move from traditional on-premises IT infrastructure toward computing in the cloud and a microservices-based architecture. All signs point to the rapid adoption of cloud computing as a means to facilitate not just changes in development methodologies, but changes to the very nature of IT architecture and maintenance.

As with all disruptive technologies, any cloud migration initiative must be approached with caution, and based on a strategy that includes big-picture thinking and exacting attention to security issues. The advantages offered by a cloud-based environment are well-documented. Nevertheless, there are numerous choices to be made that can transform the complexities of the migration process into a relatively smooth transition, especially with regard to application and data security.

This chapter describes the options available to you when contemplating a migration of application and data resources to the cloud.

Migration Strategy Fundamentals

As with Agile methodology, you must always have your eye on business objectives when developing your migration strategy. Your strategy will require conformance to your organization's specific business mission, but there are fundamental components that are essential to all cloud migration initiatives:

- Define an end-to-end strategy that takes into consideration your business objectives as well as the impact of cloud migration on IT operations.
- Take the opportunity to discover and evaluate your enterprise application portfolio to see where inefficiencies exist, and where they can be remediated and optimized with available cloud services.
- Re-design your business applications to integrate effectively with the specific service models offered in the cloud.
- Understand the model of shared responsibility as it relates to security policy and risk mitigation, and develop policies and controls accordingly.

After you have a thorough and accurate picture of your application portfolio, you can begin thinking about where to start your application migration. There will be "low-hanging fruit" that will be the easiest to migrate, along with other applications that present complexities requiring additional time and attention.

Develop a plan that can be used as a framework for each application that you migrate to the cloud and the order in which they are to be migrated, but keep in mind that your plan will likely need to be amended and modified as you proceed through your application portfolio.

Service Models

There are three commonly recognized service models, sometimes referred to as the SPI (Software, Platform and Infrastructure) tiers (SaaS, PaaS, and IaaS), that describe the foundational categories of cloud services:

- **Software as a Service (SaaS)** can be compared to a one-stop shop that provides everything you need to run an application. Typically, SaaS providers build applications on top of Platform as a Service (PaaS) and (IaaS) Infrastructure as a Service to take advantage of all the inherent economic benefits of the IaaS and PaaS service models.

  SaaS provides all the usual components of an on-premises application, including an application/logic layer, data storage, and access via API calls. Typically, there are multiple presentation layers available to the end user, including mobile apps, web browsers, and public API access.

  SaaS examples: Google Apps, Salesforce, Workday, Concur, Citrix GoToMeeting, Cisco WebEx.

- **Platform as a Service (PaaS)** provides a platform that offers management of servers, networks, and other system components. Cloud users only see the platform, not the underlying infrastructure.

  This service model is most interesting to developers who need an application platform (for example, a place to run code), or to developers of database applications. Server management is handled in the cloud and is based on utilization, freeing the developer from managing individual servers, patches, and other underlying infrastructure.

  PaaS examples: Salesforce Heroku, AWS Elastic Beanstalk, Microsoft Azure, Engine Yard, and Apprenda.

- **Infrastructure as a Service (IaaS)** provides a shared pool of computer, network, and storage resources. Cloud providers use the technique of "abstraction," typically through virtualization, to create a pool of resources. The abstracted resources are then "orchestrated" by a set of connectivity and delivery tools. Orchestration is the process that coordinates the abstracted resources, creates the individual pools, and automates the delivery of the pools to end users.

  IaaS examples: DigitalOcean, Linode, Rackspace, Amazon Web Services (AWS), Cisco Metapod, Microsoft Azure, Google Compute Engine (GCE), Joyent.

  Most, if not all communications between cloud components are typically handled by application programming interfaces (APIs). A set of APIs is often made available to the cloud user so they can manage resources and configuration. These APIs use representational state transfer (REST) or RESTful web services as a way of providing interoperability between computer systems on the internet.

Comparison of Shared Responsibility

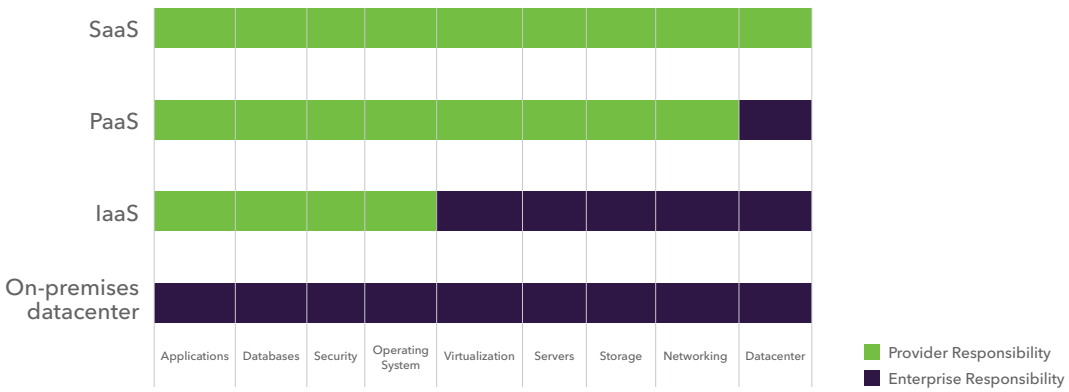

Deployment Models

Cloud deployment models apply across the entire range of service models. They describe how cloud technologies are implemented and consumed:

- **Public Cloud:** Owned and operated by a cloud service provider and made available to the public or a major industrial sector on a pay-as-you-go basis.
- **Private Cloud:** Operated solely for a single organization and managed by the organization or by a third party. A private cloud may be located on- or off-premises.
- **Community Cloud:** Shared by several organizations with common concerns. These concerns can include security requirements or governance considerations. The deployment can be managed by the community or by a third party. A community cloud can be located on- or off-premises.
- **Hybrid Cloud:** Typically comprises two or more clouds (private, community, or public) and possibly an on-premises infrastructure as well. The different cloud deployments maintain their discrete identities but are able to interoperate via standard or proprietary technologies. An example of interoperability is "cloud bursting," which enables an application to run in a private cloud and burst into a public cloud during increased demands for computing capacity.

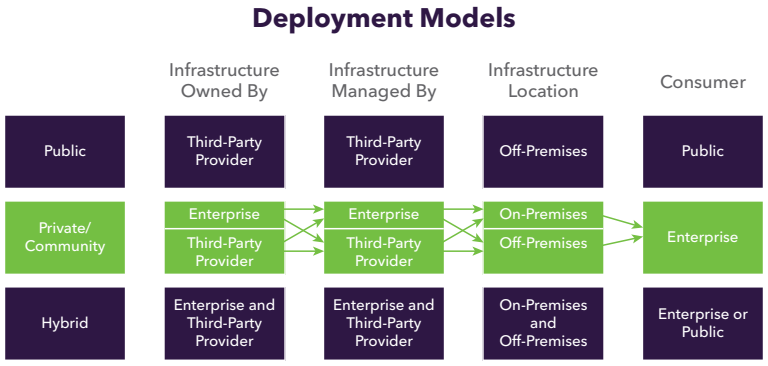

Five Application Migration Strategies

Thankfully, the process of transitioning to the cloud offers a great deal of flexibility conducive to a cautious approach. Transitioning is not an "all or nothing" proposition. It is certainly possible, and indeed in many cases desirable, to leave some applications running in a local, traditional datacenter while others are moved to the cloud. This "hybrid model" makes it possible for companies to move their applications to the cloud at their own pace.

The complexity of migrating existing applications varies, depending on the architecture and existing licensing arrangements. We strongly recommend starting with something on the low-complexity end of the spectrum for the obvious reason that it will be easier to complete—which will give you some immediate positive reinforcement or "quick wins" as you learn. If you think about the universe of applications to migrate on a spectrum of complexity, a virtualized, microservices-oriented architecture would appear on the low-complexity end of the spectrum, and a monolithic mainframe at the high-complexity end of the spectrum.

The 5 most common application migration strategies are:

1. Rehosting

Sometimes referred to as "lift and shift," rehosting simply entails redeploying applications to a cloud-based hardware environment and making the appropriate changes to the application's host configuration. This type of migration can provide a quick and easy cloud migration solution. There are trade-offs to this strategy, as the IaaS-based benefits of elasticity and scalability are not available with a rehosting deployment.

However, the solution is made even more appealing by the availability of automated tools such as Amazon Web Services VM Import/Export. Nevertheless, some customers prefer to "learn by doing," and opt to deploy the rehosting process manually. In either case, once you have applications actually running in the cloud they tend to be easier to optimize and re-architect.

Rehosting is particularly effective in a large-scale enterprise migration. Some organizations have realized a cost-savings of as much as 30 percent, without having to implement any cloud-specific optimizations.

2. Replatforming

This strategy entails running applications on the cloud provider's infrastructure. You might make a few cloud-related optimizations to achieve some tangible benefit with relative ease, but you aren't spending developer cycles to change an application's core architecture.

Advantages of replatforming include its "backward compatibility" that allows developers to reuse familiar resources, including legacy programming languages, development frameworks, and existing caches of an organization's vital code.

An unfortunate downside to this strategy is the nascent state of the PaaS market, which doesn't always provide some of the familiar capabilities offered to developers by existing platforms.

3. Repurchasing

This solution most often means discarding a legacy application or application platform, and deploying commercially available software delivered as a service. The solution reduces the need for a development team when requirements for a business function change quickly. The repurchasing option often manifests as a move to a SaaS platform such as Salesforce.com or Drupal. Disadvantages can include inconsistent naming conventions, interoperability issues, and vendor lock-in.

4. Refactoring / Re-architecting

This solution involves re-imagining how an application is architected and developed, typically using the cloud-native features of PaaS. This is usually driven by a strong business need to add features, scale, or performance that would otherwise be difficult to achieve in the application's existing environment.

Unfortunately, this means the loss of legacy code and familiar development frameworks. On the flip side, it also means access to world-class developer's tools available via the provider's platform. Examples of productivity tools provided by PaaS providers include application templates and data models that can be customized, and developer communities that supply pre-built components.

The primary disadvantage to a PaaS arrangement is that the customer becomes extremely dependent on the provider. The fallout from a disengagement with the vendor over policies or pricing can be quite disruptive. A switch in vendors can mean abandoning most if not all of a customer's re-architected applications.

5. Retiring

As stated previously, the initial step in the cloud migration process is the discovery of your entire IT portfolio. Often, this discovery process entails application metering to determine the actual usage of deployed applications. It's not unusual to find that anywhere between 10 to 20 percent of an enterprise IT estate is no longer being used. Retiring these unused applications can have a positive impact on the company's bottom line; not just with the cost savings realized by no longer maintaining the applications, but by allowing IT resources to be devoted elsewhere, and by minimizing security concerns for the obsolete applications.

Checklist: Essential Characteristics of Cloud Computing

When contemplating cloud migration, it is of the utmost importance that you ensure certain features are included by your cloud provider. Listed here are all the characteristics essential to cloud computing. Your provider's cloud environment must include these characteristics. If it lacks any of these characteristics, it is likely not a cloud, or it doesn't provide the necessary characteristics of a cloud.

- ❑ **Resource pooling:** Resource pooling is the most fundamental characteristic. Resource pooling describes a situation in which providers serve multiple clients, or "tenants" with provisional and scalable services. These services can be adjusted to suit each client's needs without any changes being apparent to the client or end user.
- ❑ **Self-service:** Customers can change their levels of service at will without being subject to any of the limitations of physical or virtual resources. Customers provision resources from the pool using on-demand self-service. They manage their resources themselves, without actually contacting a service provider.

- ❑ **Centralized networking resources:** Customers have access to networking resources from a centralized third-party provider using wide area networking (WAN) or internet-based access technologies. The network can be shared as well as the computing resources. All resources are available over the network, without any need for direct physical access.
- ❑ **Elasticity:** The cloud environment must be "elastic." It must be able to adapt to workload changes by provisioning and de-provisioning resources in an autonomic manner, such that at each point in time the available resources match the current demand as closely as possible. This allows customers to more closely match resource consumption with demand.
- ❑ **Pay-per-use:** Pay-per-use plans are where customers are billed based on the resources they use. Specific instrumentation must be included in the application to support metering for billing. This plan has the advantage that it relates the customers' costs to their usage. This is where the term utility computing comes from, since computing resources can now be consumed like water and electricity, with the client only paying for what they use.

Application and Data Security for SPI

From a high-altitude viewpoint, cloud security is based on a model of "shared responsibility" in which the concern for security maps to the degree of control any given actor has over the architecture stack. Thus, the most important security consideration is knowing exactly who is responsible for what in any given cloud project:

- **Software as a Service:** The cloud provider is responsible for the bulk of security concerns. Fortunately, cloud vendors can provide a high level of security for cloud services, delivering secure application development and operation with features such as application code scanning, application security management and vulnerability detection. Bear in mind, cloud consumers must always ensure the security of their application data and the endpoints that are used to access cloud services.
- **Platform as a Service:** The PaaS cloud provider is typically responsible for the security of the physical infrastructure. The consumer is responsible for everything they implement on the platform, including how they configure any offered security features. Thus, at the application layer and the account and access control layer, the consumer must tend to the same security concerns as with an on-premises installation.
- **Infrastructure as a Service:** IaaS is essentially a data center in the cloud. The cloud provider has primary responsibility for the physical security of the servers and the data vulnerability of the network itself. The cloud user is responsible for the security of everything they build on the infrastructure.

The cloud user should be aware that data is particularly vulnerable to a breach when it is being migrated from onsite servers to an offsite cloud. Physical servers are only as secure as the level of hardening applied to them, and unsecured equipment can easily become the source of a data breach.

In summary, identity and access management is essentially the responsibility of the cloud consumer in the IaaS model, since the provider only operates the physical or virtual infrastructure. More of a shared responsibility exists with PaaS and SaaS. In those models, access management is the domain of the user, the provider is responsible for API security and auditing. Identity management, including privileged user management, is also a shared responsibility between cloud provider and consumer.

A Simple Cloud Security Process Model

The development of a comprehensive cloud security process must take into account a wide range of implementation details such as design models and reference architectures. The following high-level process model for managing cloud security contains only the most essential items that must appear in a cloud security process model:

- Identify enterprise governance, risk, and compliance requirements, and legacy mitigation controls.
- Evaluate and select a cloud provider, a service model, and a deployment model.
- Select your cloud provider, service, and deployment models.

- Define the architecture of your deployment.
- Assess the security controls and identify control gaps.
- Design and implement controls to fill the gaps.
- Develop and implement a migration strategy.
- Modify your implementation as necessary.

Each migration process should be evaluated based on its own set of configurations and technologies, even when these projects are based on a single provider. The security controls for an application deployed on pure IaaS in one provider may look very different than a similar project that instead uses more PaaS from that same provider.

The key is to identify security requirements, define the architecture, and determine the control gaps based on the existing security features of the cloud platform. It's essential that you know your cloud provider's security measures and underlying architecture before you start translating your security requirements into cloud-based controls.

Checklist: Application and Data Security for SPI

Recommendations:

- ❑ Cloud users must understand the differences between cloud computing and traditional infrastructure or virtualization, and how abstraction and orchestration impact security.
- ❑ Cloud users should evaluate their cloud provider's internal security controls and customer security features so the cloud user can make an informed decision.
- ❑ Cloud users should, for any given cloud project, build a responsibilities matrix to document who is implementing which controls and how. This should also align with any necessary compliance standards.
- ❑ Cloud users should become familiar with the NIST model for cloud computing and the CSA reference architecture.
- ❑ Cloud users should use available tools and questionnaires to evaluate and compare cloud providers.
- ❑ Cloud users should use available tools to assess and document cloud project security and compliance requirements and controls, as well as who is responsible for each.
- ❑ Cloud users should use a cloud security process model to select providers, design architectures, identify control gaps, and implement security and compliance controls.
- ❑ Cloud users must establish security measures, such as a web application firewall (WAF), that allow only authorized web traffic to enter their cloud-based data center.

Security Policy Migration

Migrating applications to the cloud can, if not proactively addressed, result in incomplete or inconsistent security policies that conflict with your on-premises security policies. For example:

- On-premises application, data, device, and web controls will be disabled or not available in the cloud, unless specifically enabled or configured for the cloud-based applications.
- Some on-premises services, such as Active Directory, don't directly map to cloud versions. Instead, you can synchronize, but not migrate, on-premises directories, computer accounts, group policies, etc.
- Lack of a unified management platform for both on-premises and cloud application security will make it difficult to implement consistent security policies across the application portfolio and harder to investigate security incidents.

What this means is that you need to examine your on-premises policies (user access and authorization, network traffic, system and application configuration, event logging and monitoring, and so on) and map to cloud-based policies. There are multiple ways to address this:

- Many solutions support both on-premises cloud-based deployment. In this case, by deploying in the cloud, policy migration is straightforward.
- Some clouds offer their own security services. In these cases, it may be possible to create policies that accomplish the same thing, but it is rare to be able to automatically translate a policy from one security solution to another.
- Look for solutions that provide a single management interface for cloud and on-premises application security to simplify and streamline policy management.

Checklist: Policy Migration

Recommendations:

❑ Identify your customized policies.

❑ Identify the default policy settings provided by your cloud host.

❑ Determine which policies you need to migrate to the cloud and what type of policies will be managed by your cloud host.

❑ Determine if, and where, you have the option to deploy the same security solution both on-premises and across cloud deployments.

❑ Determine if the same management console for on-premises applications can be used for migrated, cloud-based applications. Additionally, determine whether the same policies for on-premises applications can be used in the cloud.

❑ Determine whether there are any automated processes for migrating policies and if not, what the guidelines for manually migrating policies are.

❑ Determine compatibility between your policies and your host's. If not compatible, you'll need to either deploy different applications or find a different host.

Technical and Business Considerations

If you're like many businesses, you're moving applications into public and private cloud infrastructure because you've seen how the cloud's agility, resiliency and scalability drives business growth. Fortunately, rolling out new apps in the cloud is easy with containers, micro-services, and DevOps supporting you. But what's not always as easy to figure out is application security—especially if you're in the midst of migrating and need to keep apps secure both on-premises and in the cloud.

Make no mistake: your apps will be attacked. According to the 2017 Verizon Data Breach Investigations Report, web app attacks are by far the number one cause of data breaches and denial of service incidents are more common than any other security incident.

The good news? You can secure your apps as easily as you can roll them out when you have a flexible, scalable security solution in place.

In this section, we'll discuss what you need to consider to securely migrate apps to the cloud and how a flexible licensing model can keep your applications secure wherever they live.

The Security Model in the Public Cloud

Leading cloud vendors introduced a shared responsibility model for security in the cloud. Amazon states that AWS has "responsibility for security of the cloud," while customers have "responsibility for security in the cloud." Microsoft Azure, Google Cloud and other vendors also adopted this model. What does it mean for you? Cloud vendors provide the tools and services to secure the infrastructure (such as networking and compute machines), while you are responsible for things like network traffic protection and application security. For example, cloud vendors help to restrict access to the compute instances (AWS EC2/Azure VM/Google CE) on which the web server is deployed (by using security groups/firewalls and other methods); they also deny web traffic from accessing restricted ports by setting only the needed HTTP or HTTPS listeners in the public endpoints (usually the load balancer).

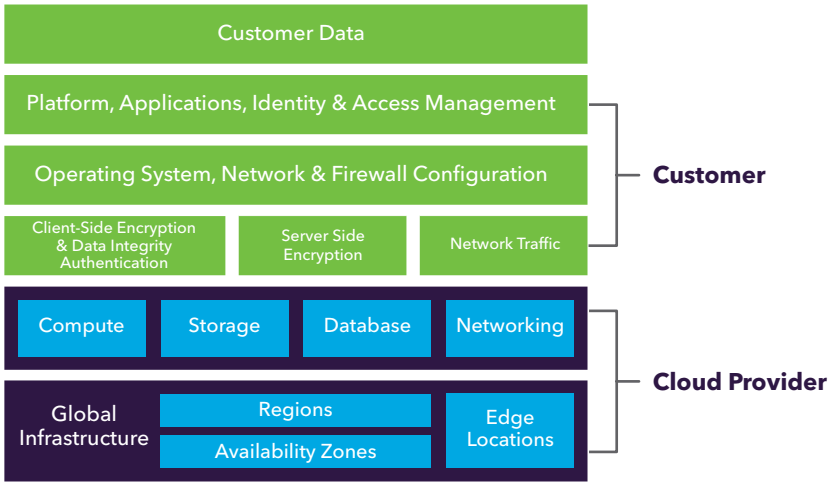

But public cloud vendors do not provide the necessary tools to fully protect against application attacks such as the OWASP Top 10 risks or automated attacks. It's your responsibility to establish security measures that allow only authorized web traffic to enter your cloud-based data center—just as with a physical data center. Securing web traffic in physical data centers is typically done by a web application firewall (WAF) and fortunately, WAF can be deployed in the public cloud as well.

Choose Flexible Application Security for the Cloud

When choosing solutions to mitigate different web application threats, it's important to make sure that they offer flexibility to select the tools you need. The first mitigation layer is usually common for all attackers, it denies access from badly-reputed sources ("malicious IPs") and blocks requests based on predefined signatures. This solution could be useful against generic types of attacks, like a botnet attack looking for known vulnerabilities. The more targeted the attack is though, the more fine-grained the tools required to mitigate it—and the higher the level of control your security team needs. When an attacker tries to launch an attack tailored to a specific web service, you need customizable tools to block it.

An ideal solution would offer both generic and customizable tools with the flexibility to be deployed within the private network and public cloud while giving your security administrator full control, including deployment topology and security configuration. An application security solution that is deployed in the public cloud should support several key attributes:

Burst capacity:

Scalability is highly important in cloud environments, so the solution must be able to scale to process more web traffic as the load grows. It should detect critical issues like high CPU or bandwidth utilization, and automatically spawn new security instances which then register with the existing cluster of gateways and synchronize security settings.

Automatically spawn new security instances which then register with the existing cluster of gateways.

Multi-cloud security:

A security solution should support all the major public cloud infrastructures (AWS, Azure or Google Cloud Platform) and your own datacenter so you can secure applications wherever they live—now and in the future.

DevOps ready:

Security should never stop your organization from being agile. A security solution should employ machine learning to automatically understand application behavior and tune policies so that you can quickly rollout applications using DevOps without changing configuration on your existing security deployments.

Security solutions should employ machine learning to automatically understand application behavior.

Automation:

Dynamic cloud environments require automation to launch, scale, tune policies and handle maintenance operations. Your cloud application security solution needs to provide tools that facilitate this automation. You should be able to orchestrate automation with native services (such as CloudFormation in AWS and ARM in Azure), or by other third-party tools (Puppet or Chef, for example).

High availability:

Business continuity demands that your security solution be highly available. Usually this means that the solution needs to utilize the cloud-native tools for high availability (for example, span over availability zones in AWS or be part of availability set in Azure).

Centralized management for hybrid deployment:

Migration to the public cloud often happens over time and a hybrid deployment mode – in which some of the workload remains in the physical data center and some in the cloud – is very common. Operationally, this raises the question of how to apply the same organizational security policies in both the physical data center and the cloud. A security solution should have centralized management that can control deployments in both locations.

A security solution should have centralized management that can control deployments in both locations.

Pricing for Applications

Applications are moving to a more automated architecture and they're being developed and rolled out faster than ever. If any of the following apply to you, then you are in need of a flexible licensing solution for security:

- Moving to a microservices architecture
- Planning to use serverless computing such as AWS Lambda
- Deploying containers instead of traditional virtual machines
- Have a dedicated application DevOps team in your organization
- Concerned about your API security
- Moving your applications from on-premises to public cloud infrastructure like AWS, Azure or Google Cloud Platform
- Need to keep certain applications on-premises and need security for both cloud and on-premises

Imperva FlexProtect offers a single subscription with the flexibility to mix and match application security tools so you can secure applications wherever they live. FlexProtect security tools protect application portfolios in-the-cloud and on-premises, and keep your applications safe while you navigate the uncertainties of moving to a virtual, cloud-centric architecture.

With FlexProtect, you can choose from the following Imperva security solutions:

Imperva Incapsula combines security with performance optimization and load balancing to provide complete cloud-based protection for your business. Incapsula is a cloud-based service designed to stop Internet threats like massive DDoS attacks and bad bots. You can deploy it as a managed service or as a cloud service.

Imperva SecureSphere Web Application Firewall (WAF) protects business critical applications from sophisticated cyberattacks. SecureSphere blocks threats that target application vulnerabilities and illegitimate access to web services. You can deploy it as an on-demand service or as a virtual appliance in hosted cloud environments.

Imperva ThreatRadar is an advance-warning system that stops emerging threats before they impact your business. By collecting, comparing and analyzing attack data from a variety of trusted sources, ThreatRadar delivers early detection and an effective defense against constantly evolving threats.

Imperva application security solutions are available in a flexible, hybrid model that combines cloud-based services with virtual appliances to deliver application security and DDoS defense for the cloud.

Imperva Can Help Secure Your Applications Wherever They Live

Now is the time to act. It is widely expected that the majority of applications will move to cloud infrastructure by 2020. The agility with which new applications are developed and existing ones are modified in the cloud creates a greater chance of exposing vulnerabilities. So, it's imperative to provide the same or higher level of application security in the cloud as you do on premises.

Your organization needs a simple and flexible solution to facilitate a smooth transition from on-premises to the cloud. Imperva offers a solution that scales as the business does while allowing you to mix and match tools based on your application security requirements. With FlexProtect, Imperva removes the dilemma associated with cloud migration planning and future proofs application security investments.

Contact us today to find out how the FlexProtect licensing model can help you keep your apps safe wherever they live, now and in the future.

About Imperva

Imperva® (NASDAQ: IMPV) is a leading provider of cyber security solutions that protect business-critical data and applications. The company's SecureSphere, CounterBreach, and Incapsula product lines enable organizations to discover assets and risks, protect information wherever it lives—in the cloud and on-premises—and comply with regulations. The Imperva Defense Center, a research team comprised of some of the world's leading experts in data and application security, continually enhances Imperva products with up-to-the minute threat intelligence, and publishes reports that provide insight and guidance on the latest threats and how to mitigate them. Imperva is headquartered in Redwood Shores, California. www.imperva.com