A cloud for every workload

REIMAGINING MULTI-CLOUD STRATEGIES THROUGH PERSISTENT DATA CLASSIFICATION

Optimizing security

Most organizations — whether they're aware of it or not — are already using multiple clouds or are on the brink of doing so. However, multi-cloud usage is sometimes ad hoc and, likely out of timing or necessity, not always part of a carefully crafted enterprise multi-cloud strategy.

For organizations to truly optimize their security posture, it's critical to develop and implement strategies that leverage the optimal cloud — whether public, private or hybrid — for each of its data sets and applications.

The security that one cloud offers may be appropriate for one data set but excessive and wasteful for another. Building and executing this strategy is necessary to effectively safeguard the confidentiality, integrity and availability of the organization's data, while also maximizing the value of the organization's investment in cloud technology and related security controls.

The success of an organization's multi-cloud strategy is highly dependent on the organization leveraging data classification techniques to determine which types of clouds are appropriate for housing each of its data sets, and formalizing these decisions in policy. This policy must then be aggressively implemented by the organization to improve security while controlling costs.

It has become critically important for organizations to change the manner in which they store, access and use cloud-based data because of the poor security practices followed by most cloud service providers.

"For organizations to truly optimize their security posture, it's critical to develop and implement strategies that leverage the optimal cloud — whether public, private or hybrid — for each of its data sets and applications."

Understanding cloud architectures

Before we explore the various multi-cloud strategies and data classification requirements, a quick primer on typical cloud architectures is valuable. People typically think of three possible architectures for cloud deployments:

Public

Using a public cloud tends to be the most cost-effective cloud architecture in terms of operational costs, but it also tends to have the weakest security and the greatest threats.

Private

A private cloud tends to be somewhat more expensive than a public cloud because of the organization's acquisition of dedicated cloud infrastructure resources. However, a private cloud typically offers stronger security because of its relative isolation from outside entities.

Hybrid

A hybrid cloud is when an organization uses a public cloud and a private cloud together; an example is placing an application on a public cloud and verifying user identities via multifactor authentication for access to the organization's private cloud. Effectively, the public cloud has a direct connection with the private cloud. A hybrid architecture falls between a public cloud and a private cloud in terms of costs and security.

In addition to these architectures, a fourth architecture is available.

The virtual private cloud

This architecture combines the best attributes of public clouds and private clouds by establishing the logical equivalent of a private cloud within the physical infrastructure of a public cloud.

The goal of virtual private cloud architecture is to provide lower costs than a standard private cloud, while providing stronger security and management than a public cloud.

Public clouds usually have little or no security built in beyond the infrastructure layer, so customers in need of protection for their data and applications must either acquire add-on services through a third party, such as a managed security services overlay, or architect and implement their own solution, which requires a significant amount of time, cost and expertise.

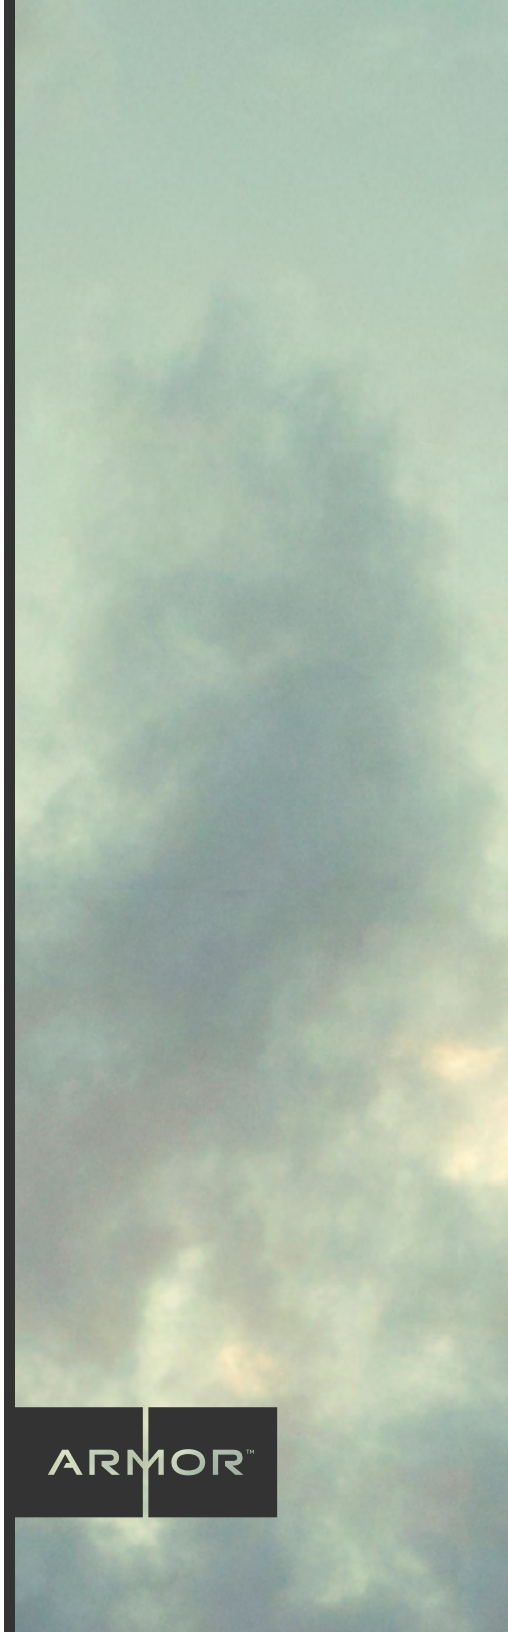

Clouds of every type, every use

When people comment that an organization is using "the cloud," odds are good that the organization is really using multiple clouds.

This may indicate the use of multiple cloud architectures — such as both public and private clouds — or it may indicate the use of multiple cloud services with the same architecture, such as two or more public clouds.

In a recent survey by Dimensional Research,[1] more than 40 percent of surveyed organizations in North America were already using multiple clouds, and another 38 percent were planning to use multiple clouds in the future.

Increase functionality

At an enterprise level, it is often advantageous to use multiple clouds instead of a single cloud. One possible reason is functionality; an organization may need to use multiple clouds because some clouds offer necessary functionality (e.g., particular software-as-a-service (SaaS) applications) that other clouds do not.

Be resilient

Another possible reason is resilience. Even though clouds generally provide high-availability, there are risks in putting critical assets in a single cloud. Should there be a failure or compromise of that cloud, for example, it could result in a major outage or security incident for the entire organization. From a deployment standpoint, it typically involves little additional effort to use multiple clouds instead of a single cloud. This is because it's already customary — and usually necessary — to divide an organization's data and applications/services into numerous cloud workloads.

Managing security & scalability

However, the security and management of multiple cloud environments can be expensive, time-consuming and labor-intensive. These challenges are exacerbated by the lack of a single dashboard or interface for all environments.

Taking advantage of the scalability and flexibility of cloud architecture necessitates taking on this challenge, as does sheer practicality; there is a limit to the resources any single workload can use.

Having multiple workloads also supports security because it better isolates sensitive data from unauthorized users. This approach improves management by allowing different administrators to each control and monitor only the necessary portions of the enterprise's data and applications.

" ... more than 40 percent of surveyed organizations in North America were already using multiple clouds ... "

DIMENSIONAL RESEARCH

1 Cloud Adoption Study: Global Survey of IT Professionals (North America vs. Global)," http://www.equinix.com/resources/analyst-reports/cloud-adoption-study/

ARMOR™

A factor of three

Organizations are increasingly dividing their workloads across multiple clouds. In addition to the aforementioned functionality and resilience motivations, another trio of factors must be considered: security, cost and performance.

These three factors are so closely related to each other that they must be considered together.

Moreover, they are three particularly important factors; in the recent Dimensional Research study, these three factors were identified as the biggest challenges to organizations wanting to expand their cloud deployments, with security by far the leading concern.

There are three primary considerations for evaluating a cloud environment.

1. Security
2. Cost
3. Peformance

Security is critical, but customizable

Different clouds offer different levels of security. For example, some workloads contain more sensitive data than others, so they may require a high degree of protection, such as that provided by a virtual private cloud from a secure cloud services provider.

Other workloads contain less sensitive data or do not contain any sensitive data at all, so they may require a lower degree of protection — for example, the security offered by a managed security overlay on top of a public cloud.

Sliding the scale between cost & performance

Achieving a high degree of protection for the most sensitive data workloads may involve somewhat increased direct costs and reduced performance, particularly if the security controls have not been carefully selected, implemented, integrated and managed.

But this protection may be required — for example, by regulatory compliance efforts — and the costs of an incident caused by having insufficient protection may greatly outweigh the additional cost of using the protection in the first place.

On the other hand, non-sensitive data workloads will not materially benefit from having more stringent security controls in place, so a less rigorous (and potentially lower-cost) security solution is appropriate for these workloads.

In short, organizations should match each workload with a cloud that provides the appropriate level of security, instead of trying to secure all workloads the same way. This helps an organization to achieve security as efficiently as possible.

"Proactively planning which cloud to use based on each workload's characteristics can be called developing a multi-cloud strategy."

ARMOR™

Data classification assumes greater role

The key to developing a successful multi-cloud strategy is to focus on understanding and documenting how different data sets should be categorized, managed, stored and protected.

Some data sets require more rigorous security than others in every organization. It's important to define two or more classifications for data that encompass all of the organization's data that is or may become cloud-based. These classifications are necessarily organization-specific because each organization has a unique environment and set of requirements.

Class of two

Imagine the simplest definition possible: two classifications. One of these classifications could be for data that, if lost, would materially damage the organization. The other classification would then be for all other data.

Data in the "material damage" classification might include the organization's personally identifiable information (PII), financial records, electronic protected health information (ePHI), and other data that is subject to regulatory compliance initiatives, as well as the organization's most important intellectual property.

Based on the definition of this classification, the organization could restrict cloud-based deployments of this data to only the most secure options available, such as a virtual private cloud with strong built-in security tools and services, so that the data is as strongly protected as possible while still being accessible by authorized users.

Data in the other classification could be placed into a public cloud, with standard security precautions (e.g., managed security overlay) to provide basic protection for the confidentiality, integrity and availability of the data.

Take inventory, then assess

The first major step in assigning the classifications is to inventory the organization's data, if this hasn't recently been done. This includes characteristics of the data such as:

1. Where it is located
2. How it flows from one location to another
3. Where & how it is used

This inventory is often the most challenging step for organizations, particularly if divisions within the organization acquire their own IT services.

It may be necessary to acquire specialized tools or services to search for cloud-based "shadow IT" services being used to store sensitive data for the organization without the knowledge of the security staff.

The second major step is to review the inventory and group items with similar security, performance and availability needs together into levels. Generally it is best to come up with a small number of meaningful levels.

For example, the levels might be Low, Low-Medium, Medium, Medium-High and High. Such levels could be defined for security, performance and availability individually.

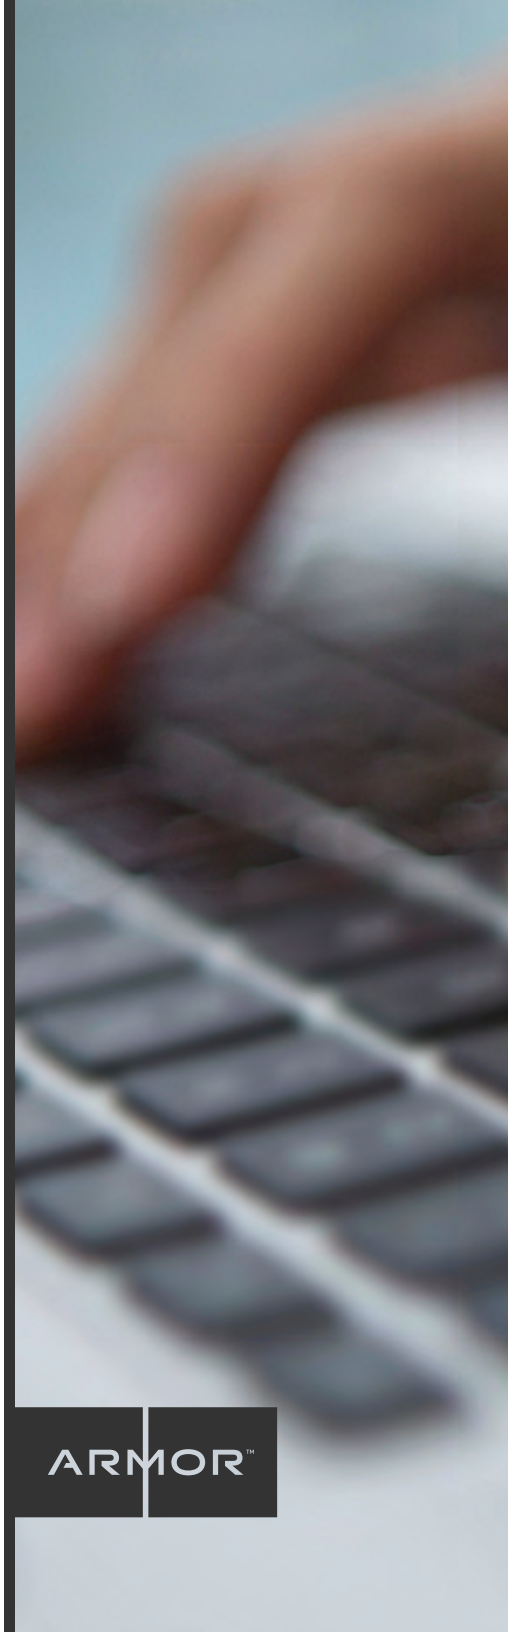

Right security, right workload

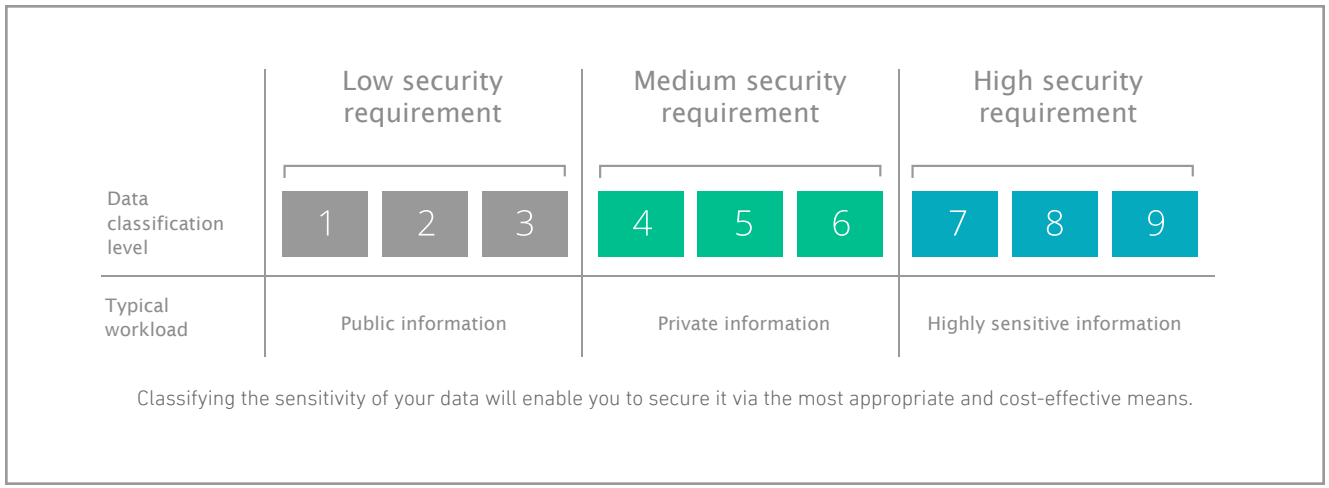

Classifying the sensitivity of your data will enable you to secure it via the most appropriate and cost-effective means.

Map your data path

The final major step is to map the established levels to the classifications. To continue the previous example, the organization might determine that for now, only data at security levels Medium-High and High should be classified as "material damage" data.

Two years from now, the organization may evaluate changes in risk and decide to reclassify the Medium security level as "material damage" data as well. This, in turn, may necessitate migrating the Medium-level data from standard public clouds to virtual private clouds that can provide stronger security protections for the data.

This three-step process is only illustrative of how an organization could implement and use data classifications for planning the security of cloud workloads. Each organization should define its own process that takes advantage of existing data classifications, data inventories and other such resources.

The important thing is to ensure that any data destined for the cloud is evaluated to determine what its security, performance and availability needs are, and that it is migrated to the appropriate cloud architecture that meets these needs as efficiently as possible.

"The key to developing a successful multi-cloud strategy is to focus on understanding and documenting how different data sets should be categorized, managed, stored and protected."

ARMOR™

Invest in your data

Determining which cloud an organization should use for deploying a particular data set and associated application can be a surprisingly complex process, with several factors to consider.

The key to simplifying and streamlining this decision-making process is understanding the nature of each of the organization's data sets and having a clear data classification policy defined for which data set characteristics require the use of certain cloud architectures.

For example, the organization's most sensitive data may necessitate the strongest cloud security available, such as a virtual private cloud with robust built-in security features.

Less sensitive data may be sufficiently protected by a managed security overlay placed on top of a public cloud. Even data that is not considered sensitive at all still needs security protection for its integrity and availability.

Because each organization has its own unique requirements and data sets, it is up to each organization to define its own data classification policy and ensure that it is enforced.

Without these controls in place, some of the organization's data is almost certainly going to be stored in clouds that do not provide adequate security, and a major data breach will be inevitable.

"Organizations must invest the necessary time and resources in multi-cloud strategy planning and enforcement to support the short- and long-term security, performance and availability of their valuable data."

US 2360 Campbell Creek Boulevard, Suite 525, Richardson, Texas 75082 | Phone: +1 877 262 3473
UK 268 Bath Road, Slough, Berkshire SL1 4AX | Phone: +44 800 500 3167