Pathway to Multi-Cloud Security Architecture

vARMOUR

PATHWAY TO MULTI-CLOUD SECURITY ARCHITECTURE

EXECUTIVE SUMMARY

Digitization of business has changed the way IT supports the needs of the organization. As a result, IT infrastructures have changed, moving more and more toward virtual, cloud and multi-cloud. Workloads with different performance, cost, and capability needs will benefit from being deployed on different types of cloud infrastructures.

Cloud-based infrastructures are designed to be agile, programmable, and developer-focused. However, these open and flexible characteristics increase attack surfaces, therefore exposing organizations to unforeseen security risk. Additionally, attackers are becoming more sophisticated in their adoption of these agile models of IT, developing cyber threat tools in no time and syndicating them across a brokered attacker ecosystem. From script-kiddies to well-funded nation states, this advancement further increases the risk profile for multi-cloud environments. The trend towards multi-cloud adoption and increased attacker sophistication creates an opportunity for IT and security leaders to build a security architecture that goes beyond existing perimeter-centric models, developed before the cloud era.

The following paper presents both a pathway to secure multi-clouds today and into the future. It is intended for use as a framework to help IT and security leaders build a more secure IT stack by proposing a vendor-neutral architecture that relies on interoperable, API-driven components.

INTRODUCTION:

MULTI-CLOUDS AND THE CHANGING ATTACKER

Digitization continues to disrupt industry after industry, making traditional organizational and business models obsolete and forcing adoption of new modes of IT. IT leaders are transforming their data center architectures to move from a reactive, inflexible organization to being a more proactive, agile part of the business that can respond quickly to changing business requirements. Over the next few years, Gartner estimates that 50% of organizations will be operating in these hybrid environments and leveraging multiple cloud providers on-premises and off. In addition, many organizations are using multi-clouds to operate bimodal IT - two distinct, separate modes of IT with different goals, processes and technologies, depending on factors such as reliability and agility.

Simultaneously, cyber threats are escalating with attackers embracing the same agile models of IT and co-developing software and control capabilities at a rapid rate - from yesterday's script kiddies to today's nation state attackers. The attackers can form and execute attacks more rapidly, and they only need to be successful once. As a result, a single breach of the perimeter can lead to situations where an attacker can begin amassing internal assets at their leisure undetected over a long period of time - up to 205 days undetected[1].

1. http://investors.fireeye.com/releasedetail.cfm?ReleaseID=897918

The typical kill chain model has the following levels:

EXTERNAL RECON ➜ INITIAL COMPROMISE ➜ INTERNAL RECON ➜ LATERAL MOVEMENT ➜ DATA STAGING ➜ EXFILTRATION.

The majority of enterprises today rely on stopping the initial compromise through deploying additional perimeter-based controls. This untenable approach relies on stopping attackers 100% of the time. From a defender's perspective, the best areas to focus your efforts are on Internal Recon, Lateral Movement, and Data Staging because 1) these are the longest stages (and therefore provide the most opportunity for identifying the activities), 2) the attacker's activities are more custom/improvised based on the environment, and 3) it's the chance to identify the attack before any damage has been done. The ability to create an economic disparity in an attack cycle through requiring greater effort than reward is the best-case outcome for multi-cloud security.

The trend towards multi-cloud infrastructure and increased attacker sophistication creates an opportunity for organizations to build a security architecture that goes beyond existing perimeter-centric models that were developed before the cloud era. A framework, defined as the Multi-Cloud Security Architecture (MCSA), enables organizations to move through levels based on their security maturity. The MCSA includes key attributes to support modern cloud environments, including:

- **Distributed security:** Placing deep security controls close to the assets it's protecting
- **API-driven extensibility:** Ability to provide ease of integration and extensibility from various vendors and across multi-clouds
- **Infrastructure independence:** Remain independent from workloads being protected to provide security integrity to maintain the appropriate security state as the workload travels throughout the clouds
- **Deep and actionable analytics:** Visualize and detect threats then be able to remediate or stop attacks

PHILOSOPHICAL APPROACH FOR SECURING MULTI-CLOUDS

Security processes are cyclical in nature and multi-cloud security is no exception to this. Newly disclosed vulnerabilities, changes in threat intelligence, changes in IT infrastructure, and investigated incidents all provide inputs into this cycle that can be broken into four distinct phases: **Prevent, Detect, Respond, and Predict.**

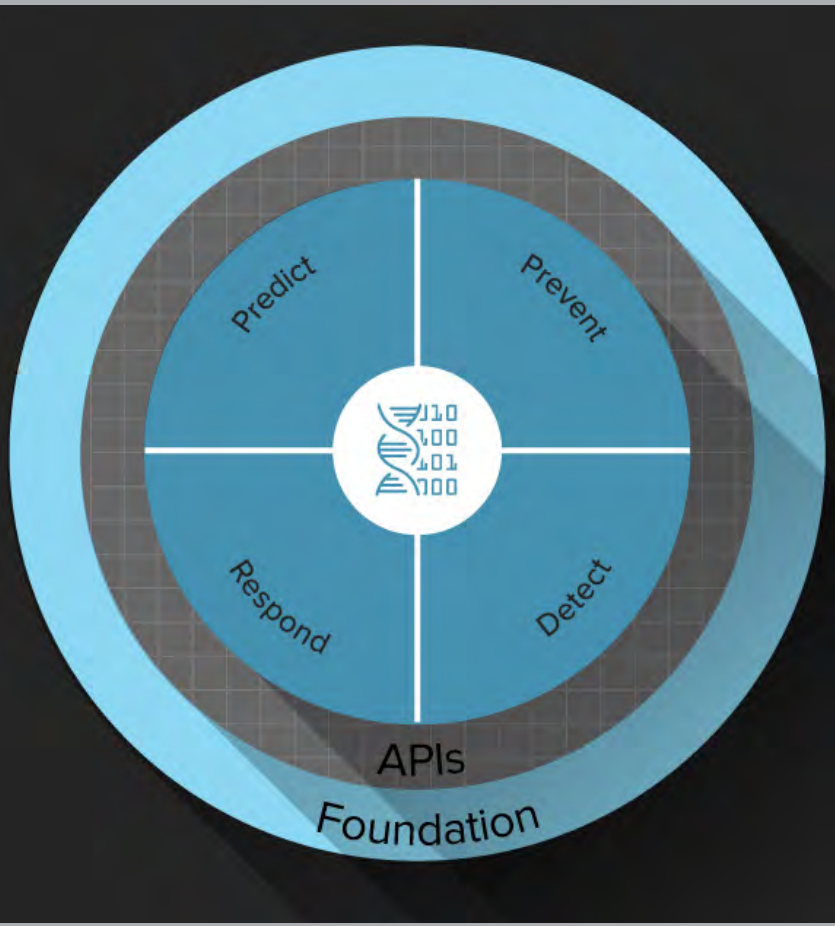

PREVENT

The tools and processes put in place across an organization to prevent an attack belong to this phase. Firewall policies, IPS, user segmentation, patch management, and infrastructure design are all aspects of prevention that are designed to reduce the potential attack surface that can be exploited for lateral movement to higher value assets. More sophisticated security programs may also begin deploying deception capabilities to further obfuscate the remaining attack surfaces.

DETECT

The set of technologies and their usage designed to identify attackers who have been able to bypass the controls deployed in Prevention fall into the phase of Detect. In addition to the tools often deployed to identify in-process attacks (i.e. IDS, WAF, anomaly detection, etc.), the processes used to consume the resulting alerts also fall into this phase. For example, aggregation and de-duplication of alerts, collection and evaluation of indicators, and promotion of alerts to incidents are all part of the Detect phase. Deception techniques also play a role in detection if attackers fall for the false attack surface deployed in the Prevent phase.

RESPOND

Often a largely process-driven phase, the Respond phase includes investigation, root cause analysis, remediation, documentation, and communication of incidents to stakeholders. More advanced security programs will employ significant automation to accelerate investigation and remediation timelines, enabling the human-driven steps in the process to be as efficient as possible.

PREDICT

The incorporation of findings from incidents, external threat intelligence, Red Team exercises, gap analysis, and pen testing into the various systems and processes throughout the security cycle constitute the Predict phase. This phase is concerned with ensuring that the entire security cycle continues to become more robust and that inefficiencies are systematically minimized or eliminated.

EVOLVING THE SECURITY CYCLE

Organizations have an opportunity to evolve their capabilities across the four phases of the security cycle in response to changing IT infrastructures and a move toward multi-cloud environments. With this security evolution, they are able to not only maintain their security posture in the face of rapidly changing technologies, but also shift the odds in their favor against attackers. Several factors come to bear in this evolution, including the leveraging of “home court advantage” for deception techniques, the sharing of inter- and intra-industry threat intelligence, the use of automation and autonomics to optimize the efficacy of the organization’s security personnel, and others. What follows is an overview of the pathway toward multi-cloud security architecture, rooted in the evolution of organizations’ security capabilities across the cycle phases.

Organizations have an opportunity to assess their current security solutions and their ability to move through the cycle as quickly as needed to meet their business agility needs, balanced with risk. Once a baseline is set, organizations are able to move along the pathway to Multi-Cloud Security Architecture to the level that matches their business speed and risk requirements. At each level, organizations will add new security capabilities and experience increasing benefits that will allow them to pass through the security cycle more and more quickly and precisely to a desired level that fits their business.

LEVEL 1

DEFENSIVE SECURITY

LEVEL 1 : DEFENSIVE SECURITY

What is Defensive Security?

In a Defensive Security level, the primary focus is keeping attackers out through the use of perimeter defenses designed to stop externally-sourced threats. This provides security professionals with an entry-level of protection at the perimeters of their infrastructure. However, in this model, once past gateway-based protection and detection technologies, attackers are easily able to move through the flat pools of infrastructure commonly found in cloud data centers. This is primarily due to a lack of internal visibility and control tools that are not present in a Defensive Security model. This is an assumed minimum level of security present in a majority of enterprise security models today.

Key Capabilities and Use Cases

- **Perimeter-focused prevention and detection (NGFW, sandboxing, NGIPS, DLP, content security).**
- **Continuous monitoring of perimeter defenses (e.g. SIEM, sampled network traffic).**

BENEFITS OF DEFENSIVE SECURITY

Defensive Security gives security professionals visibility into basic threats that are performing external reconnaissance. An easy example of this would be mass phishing campaigns that are non-targeted in nature. Defensive Security can also be useful for providing basic controls related to preventing the leakage of sensitive data externally (e.g. sending documents through Gmail).

PREPARATION FOR THE NEXT LEVEL: DISTRIBUTED SECURITY

- Understand the multi-cloud infrastructure model your company has and is planning to deploy (e.g. OpenStack, AWS, hyper-converged, etc.).
- Engage with the cloud architecture team and understand what instance types are planning to be offered in the cloud environments (e.g. service catalog inventory).
- Understand the migration plan for existing applications into the new cloud environment. Assess the sensitivity of the data contained in those applications and opportunities to 'build in' security controls as part of migration.
- Understand the rollout plan for new applications into the cloud environment. Assess the sensitivity of the data required by those applications.
- Assess the scalability of existing logging and monitoring infrastructure (e.g. SIEM log data stores).
- Define success criteria based on existing number of security events detected and stopped using a Defensive Security model. Use this as a basis for quantitative comparison later against a Distributed Security model. More events detected in a Distributed Security model is desirable.
- (Optional) Calculate the cost to purchase and operate internal security appliances, such as data center firewalls.

LEVEL 2

DISTRIBUTED SECURITY

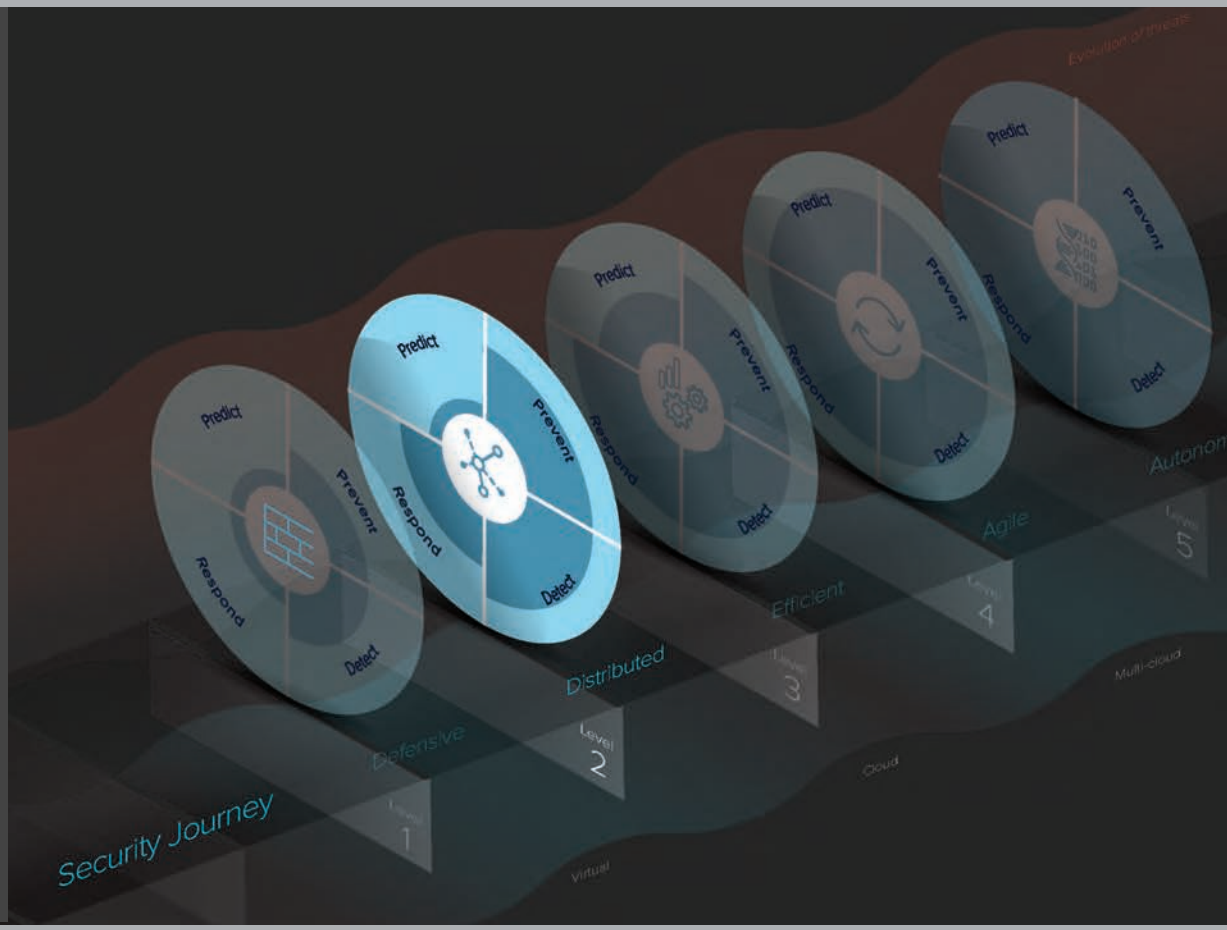

LEVEL 2 : DISTRIBUTED SECURITY

What is Distributed Security?

Security professionals who recognize the fundamental imbalance presented by a Defensive Security approach (i.e. “the attacker only has to be right once”) will benefit from adopting a Distributed Security model. A Distributed Security model focuses on controlling risk by gaining visibility and control internally at a intra-workload level by moving protection next to the asset being protected. Once security controls are closest to each workload, architects can then begin to apply these technologies to threat intelligence and the enrichment of log data. When preparing for future levels of multi-cloud security, care should be taken at the Distributed level to select controls that are programmable, scalable, and independent of the workload being protected.

Key Capabilities and Use Cases

- **Broad data center visibility (intra-workload).**
- **Broad controls, regardless of workload placement (intra-workload).**
- **Programmable architecture (e.g. robust API's for visibility and control).**
- **High-efficacy threat intelligence integrated into analytics platforms for detection of low-sophistication attacks.**
- **Enrichment of log data with relevant security information.**
- **Discovery of mode-1 applications destined for cloud-migration.**
- **Security analytics to consume the data from broad visibility.**

BENEFITS OF DISTRIBUTED SECURITY

Distributed Security enables professionals to control risk and brand damage caused by low-sophistication attacks that are able to bypass perimeter-based defenses. This is because the internal attack surface area can be minimized by distributing security controls across the multi-cloud environment. An increase in attacker sophistication is presumed due to the compression of the time window for internal reconnaissance and lateral movement caused by the implementation of these internal controls. The distributed nature of this approach enables organizations to focus on contextual awareness and threat intelligence sharing that are key to support later levels.

Intra-workload controls can also provide a direct cost-containment strategy. This benefit can be calculated by allowing mixed-sensitivity workloads to reside in a multi-tenant environment on a common pool of infrastructure as opposed to building dedicated, siloed infrastructure. A specific example of this would be allowing PCI v3.x regulated workloads that store, process, or forward payment card information to reside on the same compute infrastructure as out-of-scope assets. Distributed Security empowers security and IT professionals to deliver business value by gaining even more utilization from their cloud infrastructure investments.

PREPARATION FOR THE NEXT LEVEL: EFFICIENT SECURITY

- Identify key sources of contextual information (e.g. identity stores, systems of record, cloud management platforms, etc.).
- Assess the fidelity and accuracy of the data contained in those systems.
- Inventory currently in-use threat intelligence feeds and their feasibility for inclusion into the Multi-Cloud Security Architecture.
- Evaluate the effectiveness of existing security information sharing programs with a focus on stopping advanced attacks.
- Create a decision support framework to be used in defining appropriate response controls once initial compromise has happened.
- Create a baseline measurement of the time required for security policies to be changed in Defensive and Distributed levels. Use this as a quantitative measure of the effectiveness of an Efficient Security level.
- Understand the organization's cloud-tenant responsibility model, including:
    - Will tenant owners be able to order services in an orchestrated fashion?
    - What interfaces will cloud-tenants be able to interact with?
- (Optional) Invest in cross training of level-1 threat analyst teams with skills to support later levels.
- (Optional) Calculate the operational cost associated with making security policy changes (e.g. number of FTE hours required to log a CMDB policy change ticket, engage with technology risk and application owner teams and alter policy). Build ROI model associated with self-service aspects of 'Efficient Security'.

LEVEL 3

EFFICIENT SECURITY

LEVEL 3 : EFFICIENT SECURITY

What is Efficient Security?

Efficient Security is about making the security system more productive through the integration of context and automation to improve day-to-day security operations of clouds. At this level, security operators can reduce the tedium of security tasks and create a more responsive system in the event of an attack. Security templates become integrated with service catalogues to reduce the friction of delivering security controls in a secure, dynamic and intent-driven manner. At this level, the variety of processing actions available expands, including the ability to execute deeper analytics as risk levels increase.

Key Capabilities and Use Cases

- **Workload metadata is put into template images for embedding security policy creation during IT-controlled workload lifecycle events from a central service catalog.**
- **Automated alert triage for high-volume security events.**
- **Creation of security policy for legacy applications that have been transitioned into the cloud.**
- **Advanced security analytics combined with deep visibility (e.g. intra-workload packet capture) to provide the visibility required to detect advanced adversaries.**

BENEFITS OF EFFICIENT SECURITY

In an Efficient Security level, security policy creation for IT-controlled service lifecycle workflows are now largely orchestrated automatically from the initial application owner's request. In scenarios where a business application owner requires a new system to be created (e.g. a virtual workload), there is a semi-automated workflow similar to a CMDB request. This request is handled by the instantiation of a template from a cloud service catalog. These templates now will contain the metadata necessary for creating base security policy around this run-time workload. By embedding security policy creation into the metadata contained in the cloud service catalog, security operations are able to enable faster service provisioning while constraining the internal attack surface.

In keeping with the theme of efficient security operations, it is also important to acknowledge the operational cost savings by automating out the tedium required to deal with the large numbers of security events that are expected to be introduced with Distributed Security in multi-cloud infrastructure. Once this tedium is removed, it frees up senior threat response personnel to focus on advanced attempts to move laterally through the environment. The ability to use federated threat intelligence with deep visibility (for one specific example - intra-workload packet capture) further compresses the ability for attackers to exploit internal data center weaknesses.

Furthermore, operator error is substantially reduced by automation improvements to both policy definition and threat event processing functions. Risk associated with human error executing complex security related functions is substantially mitigated at this level.

Lastly, by using the intra-workload visibility capabilities introduced in the Distributed Security level, it is possible to create a declarative security policy semantic based on the intent of the operator. By taking this approach of using broad visibility, it will enable the migration of legacy applications that have been identified as candidates for cloud migration. Threat operations personnel are also able to better control risk through the visual linkage of intent-driven policy combined with deep visibility.

PREPARATION FOR THE NEXT LEVEL: AGILE SECURITY

- Understand the Continuous Integration and Continuous Deployment (CI/DC) tool chain in use by your development team (e.g. Chef, Puppet, Jenkins, etc.).
- Ensure internal inventories provide accurate and sufficiently granular information to enable further abstractions of policy model.
- Understand the most common workflows used by developers either via CI/CD interfaces or through Cloud Management Platforms (OpenStack Horizon, AWS CloudFormation, RedHat CloudForms).
- Identify planned PaaS deployments and desired vendors (e.g. Mesos, Pivotal CloudFoundry, RedHat OpenShift, etc.).
- Develop a decision support framework based upon risk that will identify assets that will benefit from enhanced detection and response techniques, such as attacker deception.
- Create a baseline measurement of the time required for security policies to be changed in an Efficient Security level. Use this as a quantitative measure of the effectiveness of an Agile Security level.
- Create a set of standard governance policies that will supersede developer requested policies when appropriate.
- Develop a Continuous Integration / Continuous Deployment approach to introducing new tooling in the Multi-Cloud Security Architecture. Because of the API-driven extensibility of the system, it will be possible to integrate new security tools as and when necessary into the architecture. It is recommended to build operational unit tests to validate that new tools as well as upgraded components have been well integrated into the architecture on a continual basis.

LEVEL 4

AGILE SECURITY

LEVEL 4 : AGILE SECURITY

What is Agile Security?

Agile Security increases the efficiency of service creation by allowing a business application owner (i.e. application developer) to request security policies in a common language of risk. This means that a cloud tenant is able to request services based upon the needs of an application's security profile. This federation of policy responsibility requires that a governance model is introduced in tandem, so that business application owners are responsible for their policy requests.

Conditional policy structures allow for the appropriate insertion of deep visibility and control technologies based upon contextual knowledge of the workload being protected. As an example, this means that it will be possible to understand that a sensitive workload is undergoing internal reconnaissance and based upon the risk of compromise to that asset either respond with advanced attacker deception or enforce more stringent deep controls, such as limiting the number of API calls that are able to be used.

Key Capabilities and Use Cases

- **Simplified tenant-managed security policy model.**
- **Cloud governance to provide guard rails for tenant-managed policy.**
- **Advanced detection and response capabilities (E-W and N-S) in the form of attacker deception techniques and deep packet processing capabilities.**
- **Conditional policy structure allowing for the insertion of appropriate security measures based upon the risk and value of the asset being attacked.**
- **Increased control depth beyond application-awareness.**
- **The perimeter security controls have been largely re-architected as services across multi-cloud.**

BENEFITS OF AGILE SECURITY

There are three primary benefits from deploying an Agile Security model. First, operational costs, time to provision, and operator-induced errors are further removed from data center operations through tenant-based policy management. Second, responsible application owners are able to further reduce the attack surface of their applications to levels that are able to stop all but the most advanced of attacks. In the case where application owners are overzealous in their requests, a proper governance model will help minimize risk (with the ability to reconcile requests against predefined standards) of misuse or misconfiguration. Lastly, the ability to provide variable response actions based upon contextual knowledge of the environment, reduces the overall risks of the system.

PREPARATION FOR THE NEXT LEVEL: AUTONOMIC SECURITY

- Re-validate contextual metadata from prior levels. Specifically pay attention to asset criticality and business impact values.
- Identify security workflows that required the intervention of senior threat response personnel.
- Identify mission-critical assets that cannot be taken offline, even in the event of system compromise.
- Review tenant-requests that triggered governance-based rules.
- Test algorithm effectiveness and continually refine models for efficacy and safety.
- Retire legacy security features currently stranded in perimeter DMZs in readiness to establish ‘borderless’ multi-cloud architectures.

LEVEL 5

AUTONOMIC SECURITY

LEVEL 5 : AUTONOMIC SECURITY

What is Autonomic Security?

Autonomic Security minimizes the interaction required for the creation of policies and response to threats - evolving policy definition from declarative creation of relationships by application owners, to the automatic permissioning of relationships based upon organizational and technical policies (approved application, protocol, and organizational ‘patterns’ and toxic ‘anti-patterns or ‘guide rails’). Machine learning techniques are used to converge policy definitions around observed models and to recognize significant deviations or violations to those models over time, such as those present in a compromised system.

At an Autonomic Security level, risk of deploying workloads and applications to various venues within the multi-cloud is automatically managed by risk brokering system, which evaluates prevailing threat conditions associated to each location, application risk context, and makes deployment decisions accordingly.

Key Capabilities and Use Cases

- **Deploy machine-learning to further reduce the threat attack surface by adapting to previous attacks and 'in scope' changes to application behavior (policy autonomics).**
- **Predictive analytics of policy definitions based upon prior knowledge of service lifecycles and boundaries defined according to corporate 'acceptable behavior' models.**
- **Risk brokerage model informing workload placement decisions in multi-clouds.**
- **The perimeter has been largely dissolved apart from coarse-grained controls oriented towards the Internet. Security is baked in as part of cloud solution delivery.**

BENEFITS OF AUTONOMIC SECURITY

Reducing necessary operator interaction and minimizing exposed complexity utilize machine-learning algorithms can identify significant anomalies and behavioral changes at an Autonomic level. Security is now managed by exception - saving security operators' time. In addition, risk-brokering systems reduce the inherent risk associated with exposing security policy and governance to end users that can act as internal "weak links." Organization at this level will build a self-healing, self-optimizing security system that progresses through the security cycle with the greatest speed and efficacy for multi-cloud infrastructure - reducing risk without slowing down the business.

CONCLUSION

Digitization is happening now... but what could be next? The attacker landscape will continue to evolve with new technology trends being adopted by organizations. In this way, security must be architected for today's cyber threats, but adapt for what is to come in the future. This requires a completely new architectural approach to security that is API-driven and independent of underlying infrastructure, while providing the deep visibility and control that is needed in multi-cloud environments.

However, organizations will not get there overnight. As organizations have evolved IT to meet their business needs, security must have a similar transformation to match the level of risk and agility tolerated by each unique organization. Adoption of the Multi-Cloud Security Architecture in partnership with vArmour and the surrounding ecosystem can give customers the opportunity to improve the speed and precision of their security cycle. In this way, organizations can safely embrace the power of multi-cloud for their business now and in the future.

NEXT STEPS

For more information on the pathway to multi-cloud security, visit our website at www.varmour.com for our CTO Marc Woolward's perspective in his webinar and blog series.

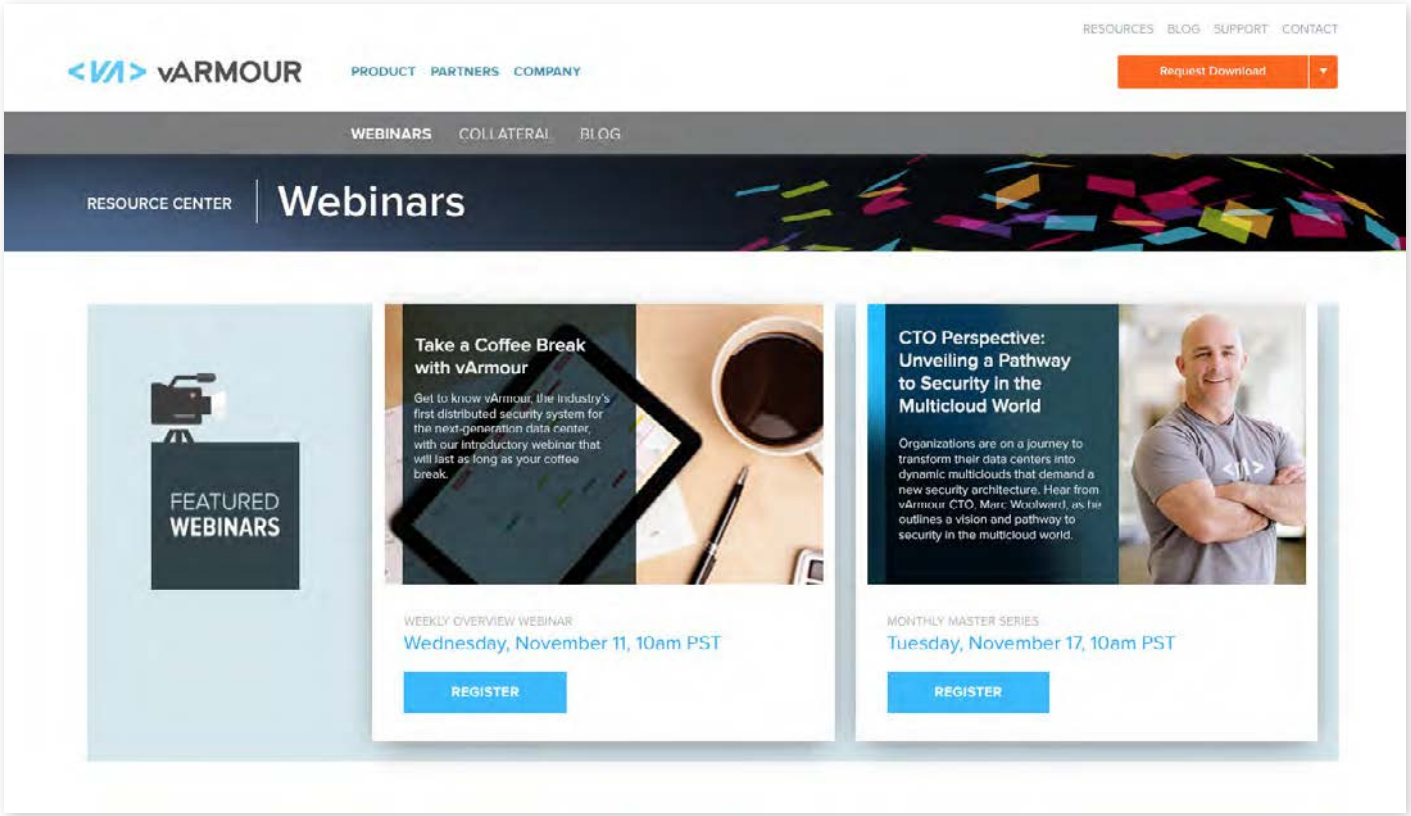

To learn where your organization is on its security pathway to multi-clouds and how vArmour and its partners can help, call us today at 650-564-5100 or email us at info@varmour.com.

ABOUT vARMOUR

vArmour, the data center and cloud security company, delivers software-based segmentation and micro-segmentation to protect critical applications and workloads with the industry's first distributed security system. Based in Mountain View, CA, the company was founded in 2011 and is backed by top investors including Highland Capital Partners, Menlo Ventures, Columbus Nova Technology Partners, Work-Bench Ventures, Allegis Capital, Redline Capital, and Telstra. The vArmour DSS Distributed Security System is deployed across the world's largest banks, telecom service providers, government agencies, healthcare providers, and retailers. Partnering with companies including AWS, Cisco and HPE, vArmour builds security into modern infrastructures with a simple and scalable approach that drives unparalleled agility and operational efficiency. Learn more at www.varmour.com