Hyper-Resilient, Survivable Cloud Security Orchestration

Ed Amoroso, TAG Cyber
Marc Woolward, vArmour

Abstract

Several existing and new architectural methods are explained that increase the resilience and survivability of distributed cloud security orchestration as found in the vArmour solution set.

Introduction

Every cyber security architect knows that a single point of failure on an enterprise network can be targeted for degradation by attackers armed with botnets. Distributed denial of service (DDOS) attacks aimed at corporate gateways, for example, have been a source of concern for chief information security officers (CISOs) concerned with availability. In theory, a malicious actor could take out any exported service delivered through a gateway by simply flooding the egress capability of the link – and several prominent instances of this attack have been observed in recent years [1].

One technical solution to such single points of failure is to distribute the service delivery. This can be accomplished physically by distributing the data center entry points and links, not to mention increasing the bandwidth capacity of these links. The risk can also be reduced through more logical distribution using methods such as content distribution networks (CDNs), which create a distributed ingress/egress mesh that is much harder to attack with a DDOS flood from a botnet [2].

To visually represent the problem and its potential solutions, we can draw an enterprise perimeter LAN with a single narrow entry point, which represents the target for traditional DDOS attacks from botnets. One solution option is to physically distribute and expand physical access capability to offer services more robustly through multiple, wider entry points. Alternatively, one can introduce CDN front-end capabilities to reduce access survivability risk at the perimeter. These security methods are shown in the figure below.

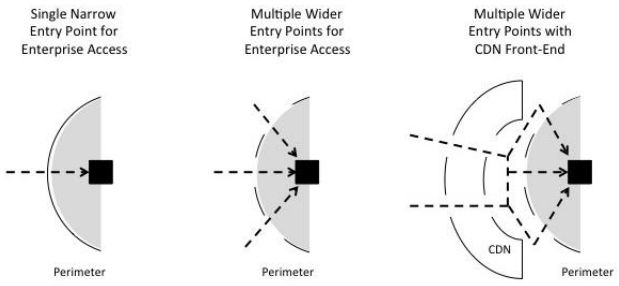

Figure 1. Increased Access Survivability Through Physical and CDN Distribution

While such solutions reduce the single point of failure risk for traditional service delivery on both Internet and mobile broadband infrastructure, enterprise LANs are now disappearing in favor of more complex, hybrid clouds. The resulting hybrid infrastructure includes distributed endpoints, workloads, hosting centers, and the like – so one would be inclined to suggest that hybrid networks, and the security systems deployed upon them, will be less

vulnerable to cyber attacks. Upon closer examination, however, one can observe new cyber risks that require more resilient orchestration of cloud infrastructure to deal with DDOS and other destructive attacks.

Adding Attack Resilience to Cloud Infrastructure

Many cloud orchestration solutions will introduce a single means for managing workload policy from early prototype to final test implementation of a given enterprise architecture. In conventional terms, this is like moving the security operations center (SOC) from the perimeter-protected LAN into the cloud. That is, the cloud architecture includes a single *Controller* (analogous to a SOC for security operations), which could be attacked or which could fail for many reasons. Within modern computer systems, we call this centralized function the *control plane*.

Any cyber offensive actor could thus paralyze the enterprise architecture by cutting off the orchestration capability centered in the new cloud-resident Controller. This might seem a trivial exploit, but it has high implications in environments that haven't properly considered the availability implications. For enterprise security teams with good reason to be concerned with the consequences of this resiliency and survivability scenario, we offer below seven complementary architectural methods for consideration:

Method 1: Resilient Endpoint Base Policy. By ensuring that distributed endpoint workloads have a stored local security policy management to support continued operation in the event of a Controller failure, the enterprise ensures a degree of failsafe operation. The vArmour solution works in this manner and has been a welcome resiliency feature for many enterprise security teams. The problem with relying on this approach as a sole means for survivability, however, is that if an urgent policy change is needed during a period in which the Controller is unavailable, the existing resilient base rules cannot be updated until the Controller is recovered. This is an existing problem with perimeter-protected LANs, so the scenario is not new; but it does highlight the need for additional resiliency improvement.

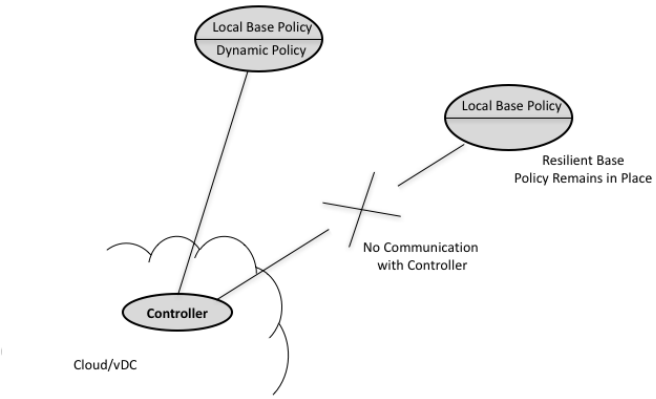

Figure 2. Resilient Endpoint Base Policy

A best practice for determining local base policy is that it should include the entire list of essential rules. If the Controller is in full communication, there should be no operational difference between local base rules and the dynamic updates being made in real-time. Obviously, once the Controller loses connectivity, any additional dynamic updates would be no longer possible until communication is re-established.

Method 2: Redundant Controller Orchestration. A second approach to improving resiliency and survivability in distributed cloud security orchestration involves redundant operation of the Controller, often as a live node and hot standby. This method roughly doubles the likelihood that an attack will be survived, but it is obviously vulnerable to any determined adversary targeting both the live and standby systems. While this risk already exists with live and standby SOCs inside perimeter LANs, our resiliency expectations for hybrid and public cloud architectures are typically much higher. Redundant operation is thus an acceptable solution for now, especially in conjunction with local base policy resiliency, but it will require additional improvements moving forward.

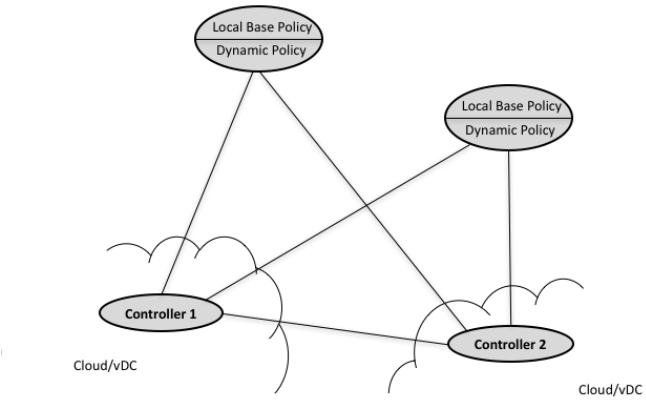

Figure 3. Redundant Controller Orchestration

One point worth adding for redundant Controller orchestration is that the failover cutover mechanism cannot fail if the live node degrades. For example, if the network links for the live Controller are saturated by an attack, then the cutover mechanism must also not fail. It is worth noting that redundancy can include two, three (common within clustered architectures), or more Controller nodes, with the provision that practical enterprise operational costs might grow accordingly.

Method 3: Intermediary Control Plane Nodes. A third approach to increased resiliency of cloud security orchestration, also implemented by vArmour, involves using intermediary nodes between cloud-hosted endpoint workloads and the highest-level Controller nodes. In this approach, a middle layer of nodes (often in an 8:1 to 32:1 ratio with Controller nodes) has the intermediate responsibility to absorb availability attacks, and to distribute the management workload across multiple nodes. Combining this intermediate Control Plane node method with redundant Controllers and resilient base policy represents the current state of the art in secure cloud orchestration survivability. Additional methods are presented below, but enterprise network security managers employing these first three methods are well ahead of the current norm.

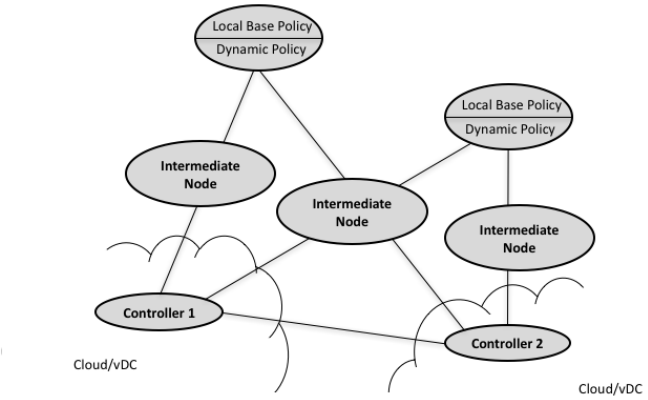

Figure 4. Intermediary Controller Nodes

A recommended enterprise architectural approach involves evolution of Controller and intermediate nodes toward common functionality utilizing cluster based approaches. This increases redundancy, while also preserving the intermediate task concept. For now, this solution is effective in distributing workloads and an adversary would need to invest significant effort to cause outages.

Method 4: Data and Control Plane Separation: A fourth approach involves separating data and control plane functions into and out of the nodes in a distributed, hybrid cloud configuration. Many current architectures require certain types of network traffic, such as packets with special setting or traffic not matching a local cache or policy, to be processed by the centralized control system. This allows an attacker to overwhelm the distributed system by injecting large volumes of traffic. To protect a distributed system from data plane attacks, one can create an air gap between data plane elements processing network traffic and control functions. To achieve this capability, the vArmour architecture ensures that each data plane element can operate as a standalone function from the perspective of traffic processing. In other words, there are no situations in which an attacker can craft data plane activity which will be processed by the control function within the system.

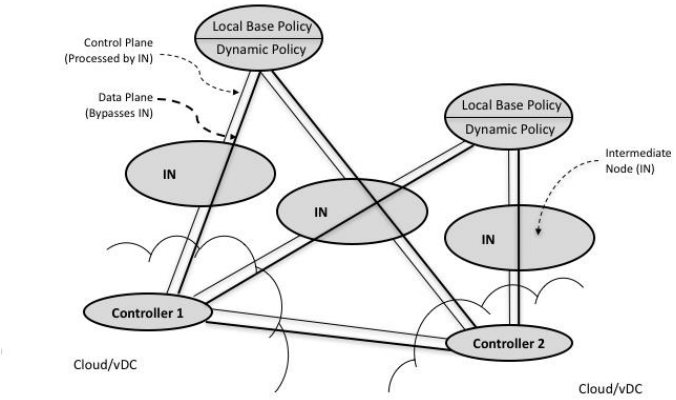

Figure 5. Control and Data Plane Separation

This architectural approach has been broadly adopted in basic network systems for many years. It is commonly referenced as a move from *flow* to *topology*-based forwarding models.

However, within a security system, the requirements to manage dynamic state generated by traffic activity, including TCP sequence numbers and application level protocol state, alongside the ability to recover quickly from component failure, present a more difficult problem. The vArmour team has addressed the requirement to separate control and data planes in a highly stateful system by introducing a mid-tier state recovery layer. The separation of data plane and control functions is fundamental in avoiding cascading failures caused by relatively simple denial of service (DOS) attacks, regardless of subsequent methods (methods 1-3) used to create hyper-resilience.

Method 5: Hyper-Resilient Survivability Protocols. As critical infrastructure companies begin to move truly consequential infrastructure to cloud, attention to the highest possible levels of resiliency becomes essential. One potential approach borrows from the botnet community, which uses fast flux DNS techniques to ensure that bot-infected nodes can continue to communicate in the presence of command and control (C&C) failures. A provocative idea thus emerges for cloud orchestration in which each node can play the role of high-level Controller or even intermediary node. In such cases, the possibility for hyper-resilience would grow proportionally with the size of the network. Thus, if a Controller node is providing policy orchestration and control for a set of endpoint nodes and that Controller node is destroyed by cyber attackers, then in this hyper-resiliency scheme, control would automatically shift to another node. This might be another Controller (as in the redundant scheme), but if that is unavailable, then a protocol would determine which node might be selected to perform that control function. Micro-services architectures that decouple functions from executing nodes could enable this form of anti-fragile paradigm. Despite the general nature of this approach, the argument can certainly be made that extending Controller functions to endpoints might be going a bit too far in any environment except those requiring extreme levels of hyper-resiliency such as warfighting scenarios.

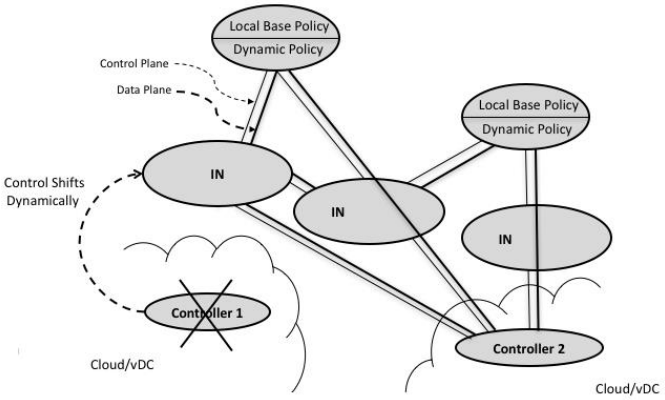

Figure 6. Hyper Resiliency Protocol for Cloud Orchestration

Note that this method continues the evolution of intermediate nodes to Controller nodes by assigning consistent design and implementation to every node, regardless of function. This is not a current practice by solution designers, even advanced ones such as vArmour, but it is worthy of consideration by the research community. One design issue with this method is that the failover protocol might be required to utilize a prioritization scheme, which would start with Controller node cutover in the presence of failure, followed by intermediate node cutover next, and finally to endpoint workload cutover if necessary. It is hard to imagine an easy way to corrupt such a scheme from an availability perspective. The only reasonable

offensive strategy would be to target the cutover scheme, which implies that it would require clean, well-tested code.

Method 6: Partitioning failure domains: To achieve the highest levels of resiliency and survivability, large scale distributed systems and networks can implement scaling mechanisms known as failure domains, which place partitions between domains sharing state. Within our examples so far, we have discussed distributed systems which share both operational state and base policy via some sort of shared state model or database. Even single systems with hyper-resilient properties can be susceptible to persistent attacks which cascade through multiple Controller transitions. By localizing operational functions to individual failure domains and logically centralizing the policy computation function for a security system, it is possible to provide greater robustness to the overall architecture. A logically centralized mechanism called a *policy compute engine* (PCE) can be used to deploy policy changes to each failure domain across a discrete API boundary, thereby preventing attacks on an individual failure domain cascading back towards the policy root.

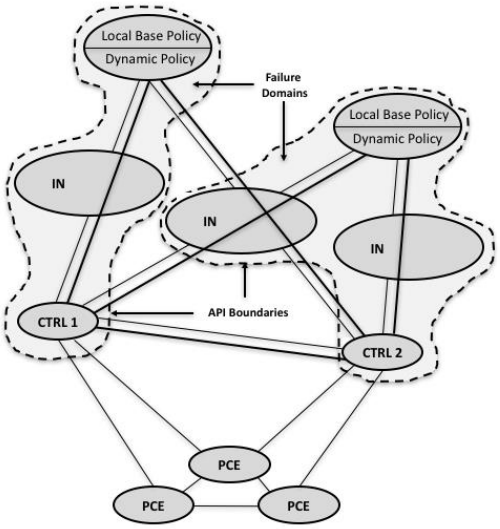

Figure 7. Implementing Failure Domains in a Distributed Hybrid Cloud

Malicious attacks upon the PCE root can be further mitigated by deploying some of the earlier resiliency methods described in this paper, as well as by allowing security policy to be directly applied to an active Controller within each security domain in a so-called *break-glass* scenario disabling the PCE. Furthermore, since a network boundary is applied between PCE and security failure domains, it becomes much simpler to harden the boundary between attackers and critical control infrastructure. The vArmour team is currently demonstrating this advanced capability of Hyper-Resilience within their PCE component, Policy Architect.

Concluding Remarks

Cyber security vendors such as vArmour have been considering many of these methods for some time as part of their commercial deployments. This follows from the observation that the benefits of cyber security are consistent with the goals of overall resilience and operational robustness. That such features are embedded in commercial tools is testament

to the growing awareness in the cyber security community that cloud solutions are no longer just innovation demos, but have evolved to robust practical deployments that now require attention to practical, operational concerns such as hyper-resiliency. The world's best cyber offensive actors will soon begin targeting cloud deployments as more critical infrastructure moves in that direction. Those of us in the cyber defensive community thus need to operate at the top of our games.

References

[1] http://www.nytimes.com/2012/10/01/business/cyberattacks-on-6-american-banks-frustrate-customers.html

[2] https://www.akamai.com/us/en/resources/protect-against-ddos-attacks.jsp