Silver Peak Security and Optimization Assessment

May 4, 2021

For more info, contact:
Wingspan Assessments | All-CCIE Team
16755 Von Karman Ave, #200, Irvine, CA 92602
info@wingspanassessments.com | 1(949) 445-6006

Client Questionnaire Form

Manufacturer	Assessment Type	Qty of Nodes
☐ Amazon Web Services	☐ Audit & Compliance	☐ # Cloud ____
☐ Arista	☑ Best Practice Architecture	☐ # Compute/Storage ___
☐ Azure	☐ Diagram & Map	☐ # Firewalls ____
☐ Cisco	☐ Feature Enhancements	☐ # Routing Instances _____
☐ F5	☐ Legal	☐ # Switches ____
☐ Fortinet	☐ Operational Efficiency	☑ # SDWAN Node 50
☐ Google Cloud	☐ Redundancy	☐ # Virtualization Nodes __
☐ Palo Alto Networks	☐ Scalability & Growth	☐ # Voice Nodes _____
☑ Silver Peak	☑ Security - Design/Architecture	☐ # Other _____
☐ VMware	☑ Security - CVEs	
☐ Other _____	☐ Speed	
	☐ Other _____	

For more info, contact:
Wingspan Assessments | All-CCIE Team
16755 Von Karman Ave, #200, Irvine, CA 92602
info@wingspanassessments.com | 1(949) 445-6006

Table of Contents

For more info, contact:
Wingspan Assessments | All-CCIE Team
16755 Von Karman Ave, #200, Irvine, CA 92602
info@wingspanassessments.com | 1(949) 445-6006

Executive Summary

Anonymous Corporation engaged WingSpan, a Silver Pea Authorized Deployment Partner, for an assessment of the existing Silver Peak environment. The goal of the Assessment was to assess the security posture and functional best practices of a Client's Silver Peak network, integrated with a Cisco network. The concluding result found that the Silver Peak fabric is well architected, and no major issues were found from a functional or security standpoint. In the course of analysis, WingSpan did identify some recommendations that could benefit Anonymous Corporation.

The WingSpan process taken to ultimately finalize these results are as follows:

1) First, the Assessment engagement began with a full-scale network discovery of LAN/WAN/SAN and Cloud networks to better understand the current state of the overall architecture at Anonymous Corporation. The discovery was conducted by our certified experts utilizing the industry's best network discovery software.
2) Second, the WingSpan engineers honed in on the following goals:
    - Assess Silver-Peak for security vulnerabilities and CVEs.
    - Assess Silver-Peak for functional best practices
    - Assess Silver-Peak for future enhancements and new features
    - Assess Silver-Peak best design practices for an optimized—high throughput and redundant SDWAN network connecting on-premise, cloud, and disparate datacenters.
3) Finally, the findings were analyzed holistically and presented in this document. In addition, remediation or improvement items are presented in this document.

For more info, contact:
Wingspan Assessments | All-CCIE Team
16755 Von Karman Ave, #200, Irvine, CA 92602
info@wingspanassessments.com | 1(949) 445-6006

1. Findings – Existing Environment

In our analysis, we highlight and provide a perspective of some of the most significant elements from the SD-WAN fabric overall. Given that some components, such as route policies, QoS policies, optimization policies, etc., are neither playing a critical role nor customized, we did not necessarily focus on them.

1.1 Business Intent Overlays

Business Intent Overlays provide the most comprehensive set of policies on the fabric and dictate the treatment of the traffic in terms of SD-WAN routing, topology, QoS, WAN utilization, boost, etc. The list of overlays are:

- **LocalBreakout.** This overlay is applied to all appliances in the fabric and is utilized to locally send Internet traffic to specific targets, without the need to backhaul the traffic to the datacenter.

- **Voice_Video.** This overlay handles the traffic related to collaboration/multimedia applications. It is also applied to all appliances, and uses high availability as its link bonding policy, therefore sends exact copies of the same packet across each of the available WAN circuits. It is important that this overlay remains set up with a mesh topology, as it facilitates direct communication between remote sites.

- **BOOST.** It is used for a very specific subset of source-destination subnets in the environment. Its traffic class, for QoS purposes, is second in order of priority after Voice_Video. This overlay is applied to 30 sites, from which 24 have actual boost license configured.

- **CriticalApps.** The CriticalApps overlay is meant to serve the most critical enterprise applications. Aside from matching default applications, this overlay has been configured to include VDI, banking-specific applications, ERPs, etc. It also has traffic class 2 and backhauls all Internet traffic matching this overlay through the datacenter. The link bonding policy here is high quality, which sends all the traffic

through the best path, and puts Forward Error Correction (parity) packets on the second-best path.

- **Transactional**. This overlay takes care of data traffic more tolerant to quality variations. It is set up with traffic class 3, and load balances between the available WAN paths for increased throughput, still adding Forward Error Correction. Backhauling to datacenter is also configured on this overlay.

- **Backup_VMotion.** Mainly oriented to replication and backup traffic. This overlay uses high efficiency link bonding, which also load balances, but in this case, there is no protection from Forward Error Correction, which exposes the traffic to potential loss in case of quality issues with the circuits. This traffic is placed on a lower traffic class (5).

- **DefaultOverlay.** This overlay acts as a "catch-all", the packets that did not match an overlay will be placed on this one. Also uses high throughput and as its link bonding policy and backhauls traffic through the datacenter. It is allocated with traffic class priority 4.

1.2 Security Policies

The way the security policies are laid out is a rather simple but effective means to protect the perimeter of the appliances. The zones have been standardized across the SD-WAN in the following manner:

- **Zone Inside.** Corresponds to the LAN-facing interface of the appliances for each location. Allows interface traffic coming from the MPLS network, as well as the inside.

- **Zone Default**. Assigned to the interface connected to the MPLS circuit. As a standard, it is located on the "01" appliance from the HA pair. Interface traffic from the same zone and the LAN is allowed.

- **Zone Outside.** The outside zone applies to all Internet-facing interfaces. The Internet connections are physically connected to the "02" device from the HA pair. Allows packets coming from the MPLS and inside interfaces matching the access-list ACL_Region, which includes specific types of traffic like ICMP, Webex,

Zoom, SkypeForBusiness, and several banking services. The table below shows the matrix view of the standard security policy template.

----table omitted for Confidentiality------

Table 1. Silver Peak security policy matrix.

In addition to the security zones, there is an inline appliance, either a Symantec or Checkpoint, between the site core switches and the LAN interfaces of each Silver Peak device. This device is performing an IPS role, monitoring/capturing data and enforcing action based on threats, botnets, and high severity events.

1.3 Templates

We consider the use of templates across SD-WAN fabric as one of the key components of this particular deployment. The different attributes are allocated into domestic, regional and site specific according to their applicability. This strategy greatly facilitates scalability and expediency as it relates to rollout of new sites. The breakdown of template attributes is as follows:

- **Default Template Group**. Fabric wide. Used in all appliances. Its attributes are:
  - Authentication mechanism.
  - SNMP configuration.
  - DNS information.
  - Banners.
  - Administrative distance.
  - Shaper.
  - Access lists.
  - Security policies.

  - CLI commands.
  - Session management values.
- **Region - <Central/East/West/Corp >.** Applicable according a specific area covering multiple sites. Used in remote sites only.
  - Netflow configuration.
  - Peer priority.

- **Site - <Sitename>.** Site-specific template.
  - Syslog server.
  - NTP information.

1.4 Boost

As mentioned previously, boost, protocol acceleration and data compression, is currently enabled on25 devices. The current amount of boost bandwidth licensed is 1000Mbps, to be allocated across those devices. There is also only one Business Intent Overlay with boost enabled. Below is the current boost allocation for the sites consuming this resource.

Appliance Name	Boost Allocated
CORP-SP-01	100.0 Mbps
AMAZON-SP-01	100.0 Mbps
ATL-SP-01	80.0 Mbps
MIA-SP-01	50.0 Mbps
IND-SP-01	50.0 Mbps
ARK-SP-01	50.0 Mbps
BAK-SP-01	40.0 Mbps
LAS-SP-01	40.0 Mbps

SEA-SP-01	20.0 Mbps
HOU-SP-01	20.0 Mbps
HUN-SP-01	20.0 Mbps
SAN-SP-01	20.0 Mbps
NV-SP-01	10.0 Mbps
KILM-01	5.0 Mbps

Table 2. Boost allocation per appliance.

1.5 High Availability

The vast majority of sites in the SD-WAN environment are deployed using high availability. WAN connections are split across the two HA members, where Internet circuits are connected to the 02 appliance, and the MPLS link is connected to the 01 HA participant device. As far as preference, the Silver Peak device named "-02" is set to be primary from a core traffic forwarding perspective, by standard.

For reference, the following sites are the only ones currently not operating in high availability mode:

- ATL-SP-01
- HUN-SP-01
- SAN-SP-01
- CAN-SP-01
- ALA-SP-01

1.6 LAN Integration and Routing

The overall design involves each Silver Peak HA member connected to both distribution DC switches. There are five flavors as far as the methodology utilized to route traffic between the Silver Peak appliances and the LAN.

1.6.1 BGP. This is the method used by the datacenters. By design, the #1 device is the primary path for both the core switches and the Silver Peak appliances. The path preference mechanism on the core switches is driven by the MED attribute. From the Silver Peak standpoint, local preference is used to determine the preferred path. Figure 1 depicts the logical integration.

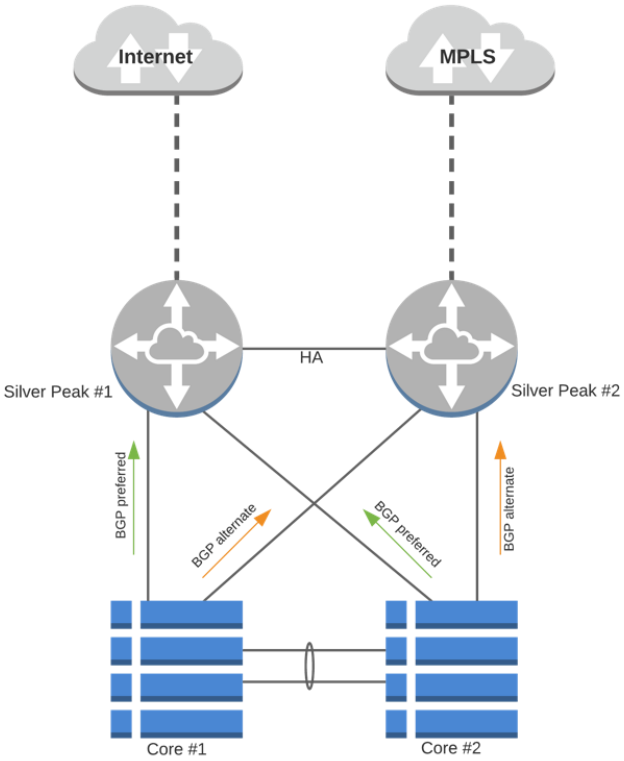

Figure 1. BGP integration setup.

1.6.2 VRRP. This FHRP is the most common approach and is utilized by the remote sites. The core switch leverages static routes pointing to the Silver Peak Virtual IP address. The LAN interfaces on both Silver Peak appliances share a common network segment with the core switch(es). The VRRP master is the Silver Peak "02" appliance by standard. The diagram below (figure 2) illustrates the scenario.

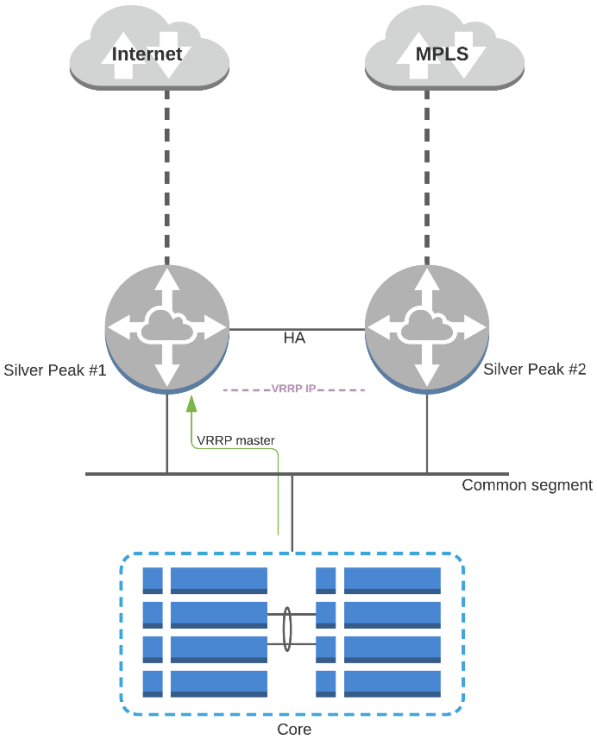

Figure 2. VRRP integration setup.

1.6.3 Standalone Silver Peak appliance with static routing. This method is currently utilized by five sites, and given that the setup is comprised by only one Silver Peak appliance, the routing mechanism is rather straightforward. Static routing pointing at the LAN0 interface of the Silver Peak as its next hop. Figure 3 displays this specific setup.

1.6.4 Standalone Silver Peak appliance with no L3 switch. This is the case for two locations, including the AWS instance. In this scenario, the Silver Peak acts as default gateway for the LAN, therefore no internal routing is required. Please see figure 4 for reference.

1.6.5 Standalone Silver Peak with OSPF. This is the case for HUN. The Silver Peak appliance and the core switch share an OSPF neighbor relationship which allows them to exchange routes. This is referenced in figure 5.

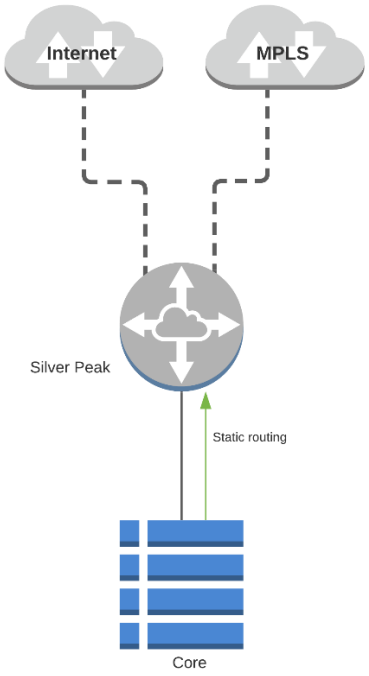

Figure 3. Static routing integration setup.

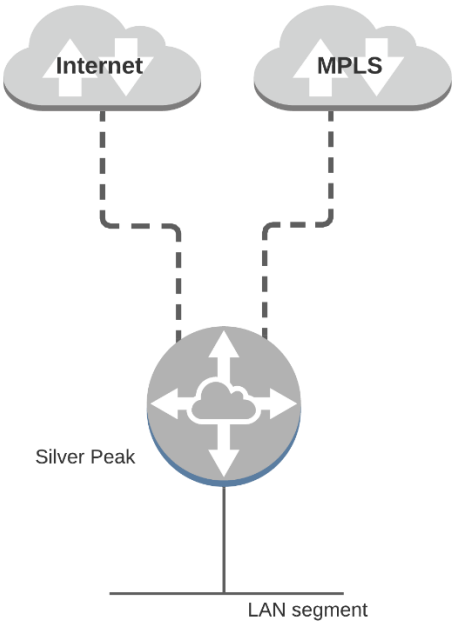

Figure 4. Standalone with no L3 switch integration.

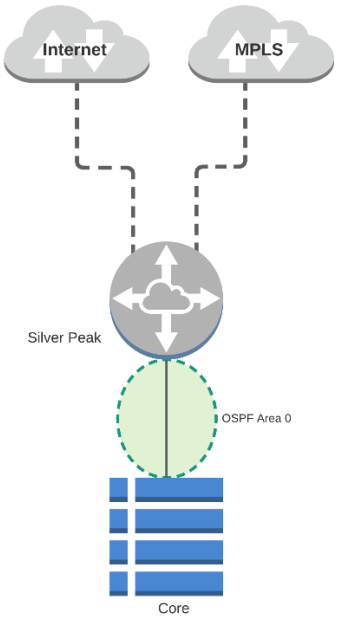

Figure 5. Standalone with OSPF integration setup.

2. Gaps and Recommendations

2.1 Architecture-wide Review

We were able to perform discovery of the Silver Peak fabric, including the orchestrator, datacenter sites, and remote sites. Also, for us to better understand the different traffic patterns of the enterprise network, we went ahead and analyzed a large number of network devices, including Cisco and Palo Alto.

From a global architecture and design standpoint, we have not found major gaps or vulnerabilities. The SD-WAN fabric is running within utilization thresholds, and there are good standards in place, mainly from the proper use of templates, which enables operational and deployment efficiency, and the avoidance of a complex routing design. As far as security, good measures are in place and ensuring the environment is tightly protected: not publicly exposing the orchestrator from incoming requests, using "IKEless" IPsec UDP tunnels, security zones, avoiding "Allow All" firewall modes to interfaces directly exposed to Internet. Also, secure management mechanisms like RADIUS, SNMPv3, TLSv1.2 (orchestrator) are in place, and legacy ones are excluded. In addition, the Orchestrator security vulnerability recently published by Silver Peak is addressed here by having Orchestrator version 8.10.13.

It is worth mentioning the potential initiative of placing a firewall in front of the Silver Peak appliance. By placing the firewall in front, the pass-through traffic will be subject to inspection. This can also be achieved enabling the inspection before it reaches the Silver Peak. There are different approaches, but from the experience, we often see from the banking industry the approach of placing the firewall in front, sometimes even placed both before and after the traffic traverses the Silver Peak appliance.

Next, we will present specific items. A good number of them have been already discussed during status meetings.

2.2 Configuration/Operational Gaps

- The current orchestrator version (8.10.14.50069) allows for upgrade to a more recent appliance software version, in case it is being considered in the short term.
    - From the currently installed VXOA version (8.2.9.12_99) it is recommended to upgrade to the newer ECOS 8.4.0.9. In order to get to this version, it is necessary an intermediate step to 8.3.9.15.
    - It is important to keep in mind that a static route configured on an appliance after upgrading to version 8.3.x will not be advertised to other routing peers on remote EdgeConnect appliances running 8.2.9.x and earlier, therefore, it is recommended with appliance version 8.3.X to perform the upgrade of all EdgeConnect devices within the same timeframe/window, or as close as possible to avoid disruption for traffic using the static routes within that scenario. Another option is to upgrade all EdgeConnect appliances with BGP/OSPF peering on the LAN interfaces first.
    - The upgrade process migrates previous configurations by creating system generated route maps. It is very important to make sure that the route maps are properly created, inheriting all the routing information, and confirming no statements are blocking routes.

- (Default template group) Auth/Radius/TACACS+ template. Since it's only using local authentication, it could be removed from the template in case this is not something to be implemented in the near future. This template is using local authentication only and pertains only to appliances.

- Following the previous idea, it is suggested under certain troubleshooting scenarios to have a non-admin local user account on the appliances, in case SSH or LAN-access is required.

- The HUN Internet interface was found set to "default" zone security. This has already been addressed by the engineering team at Anonymous Corporation.

- The application INDY_BANK is configured in the overlay ACL under the overlay LocalBreakout. This application is currently allocated as a compound application,

most likely inherited from a legacy user-defined application. There is a bug where rules are being processed incorrectly if specified using both Range and Compound format, as in the overlay ACL. It might be possible for the site ATL to hit this bug, which has been addressed in the release 8.3.0.5.

- As far as the topology, it would be worth in the future considering a hub and spoke setup for the overlays forwarding data centralized in the datacenters. The Voice_Video overlay, given its nature or likelihood to facilitate connectivity between remote sites, makes more sense to be deployed as a mesh.

- In regards to the opportunity to regionalize traffic, given the reach of the MPLS network and Internet circuits, the geographic distribution of the sites in the SDWAN fabric, and the centralized nature of the traffic towards the datacenters, maintaining a global region is sufficient. As far as what we can determine within the scope of this assessment, the major reason to use regions within this fabric is if each datacenter serves most of the resources for its region, leaving a lower demand of inter-region traffic.

- The replication-related applications in use, present in the Transactional overlay can be consolidated into the Backup_VMotion overlay. There are also some applications of this type with minimal to no usage at all, like Dropbox, Github, Spoint, etc.

- Presently, the allocated boost does not highly exceed utilization thresholds that can lead to bandwidth throttling, in general. There are two sites with relatively frequent extra boost requirements, CORP and DAL. As the traffic demand grows, it is recommended to allocate more boost bandwidth to the sites/appliances consuming it, otherwise it can throttle the device throughput overall.

- The BGP peering with the MPLS provider at the SEA and LAS locations is set up as branch. It is recommended to change it to PE-router for consistency.

- It is suggested to make sure the MPLS provider is applying a QoS profile that honors the values from the Silver Peak shaper configuration, in order to avoid traffic starvation and drops. Given that the shaper is set from a template, the profile also needs to be the same at all sites.

- The label VER_MPLS can be cleaned up from the Business Intent Overlays, as it is not being used.

- There are different locations where some traffic is being forwarded directly to the firewall. The routes for these destinations don't exist in the Silver Peak fabric. It has been confirmed with the engineering team at Anonymous Corporation that this traffic is related to third party sites, HVAC, remote clinics not on SD-WAN, etc. This is a good measure in the sense that the firewall is directly protecting such third-party traffic, offloading this function from the SD-WAN.

- Due to the absence of a default route, the sites with no operational Internet connection are unable to reach the Silver Peak cloud portal. In order to address this, it would be applicable to either advertise the cloud portal subnets via the MPLS network, or create an exception via a route policy.

Appendix C. Version Log

Editor	Date	Comments	Revision
Oscar Ramirez, CCIE, SPSX	01/30/2020	Original version	1.0

For more info, contact:
Wingspan Assessments | All-CCIE Team
16755 Von Karman Ave, #200, Irvine, CA 92602
info@wingspanassessments.com | 1(949) 445-6006