silver peak™

White Paper

From Zero Touch Provisioning to Secure Business Intent

Flexible Orchestration with Silver Peak's EdgeConnect SD-WAN Solution

White Paper

From Zero Touch Provisioning to Secure Business Intent

Flexible Orchestration with Silver Peak's EdgeConnect SD-WAN Solution

Executive Summary

Essential requirements for software-defined WANs (SD-WANs) include granular visibility into both legacy and cloud applications, as well as centrally assigned business intent policies to secure and control all WAN traffic. To be implemented at scale, these capabilities must be easily distributed from headquarters to branch offices.

The flexible orchestration included with the Unity EdgeConnect solution ensures rapid branch rollouts with zero touch provisioning (ZTP). As part of the configuration, administrators map local traffic classes into deployment profiles. These policies are then folded into discretely managed virtual topologies, referred to as "business intent overlays." This paper illustrates the ease of branch rollout and the network-wide implementation of business intent.

Using key tenets of software-defined networking (SDN) and virtualization, these virtual overlays ensure proper end-to-end handling of WAN traffic according to defined business intent. Silver Peak's SD-WAN architecture makes optimal routing, quality and optimization decisions regardless of the underlay network infrastructure and topology.

Finally, granular security is assured though a capability known as micro-segmentation, where individual workloads are mapped to underlying resources, and security controls are applied accordingly. By mapping global business policies into local office profiles, the Silver Peak EdgeConnect solution ensures a highly-visible and tightly-controlled high-performance enterprise WAN.

Rapid Installation in the Branch Office

A large-scale SD-WAN rollout requires a rapid installation process for individual offices, particularly when there are branch offices without dedicated IT staff.

Figure 1 shows a simplified picture of branch office installation. Elements that play a key role in this process include:

- Unity EdgeConnect instance deployed at a branch office
- Customer premise equipment (CPE) connecting to MPLS or broadband Internet (e.g. an edge router)
- Silver Peak's cloud portal, accessible over the public Internet
- Unity Orchestrator deployed at the enterprise headquarters

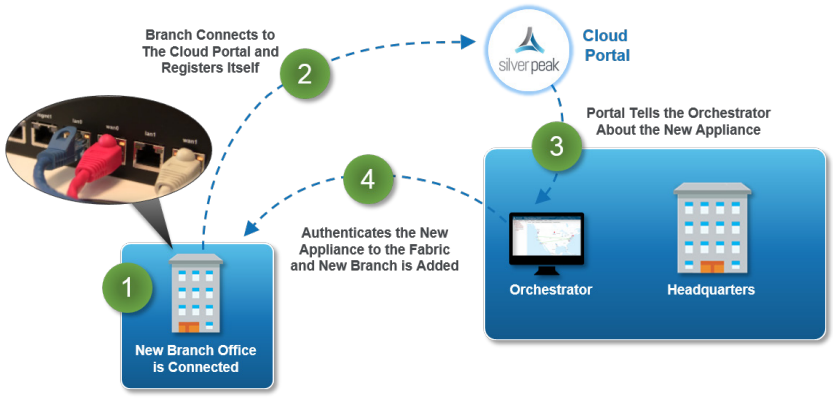

Figure 1: Simplified view of zero touch provisioning.

Silver Peak EdgeConnect zero touch provisioning (ZTP) only requires an Internet (or MPLS or other private WAN) connection at the branch. Provisioning begins by simply plugging in WAN Ethernet (your Internet connection), LAN Ethernet and power.

At that point, the EdgeConnect appliance connects to the Silver Peak cloud portal and registers itself. A registration handshake takes place between the EdgeConnect instance and the Orchestrator. If you deploy an EdgeConnect within or alongside an existing edge router, there is a DHCP request/response from the EdgeConnect software to the existing CPE, which then provides the EdgeConnect with an IP address. The EdgeConnect appliance then uses DNS to resolve the cloud portal's IP address.[2]

When the EdgeConnect appliance automatically registers itself with the Silver Peak cloud portal, it provides its license key, which may be a serial number (for a physical device) or an assigned customer code (for a virtual instance). The portal looks up the key in its database, and connects to the customer's Unity Orchestrator (typically at headquarters), providing both the key and the EdgeConnect's IP address.

The Orchestrator alerts the IT (or network) administrator about the presence of the new EdgeConnect with an email containing a URL, which directs to a page to accept or deny the new device. At this point, the EdgeConnect device must be "approved" by the IT (or network) administrator through the Orchestrator's user interface in order to be added to the network (Figure 2).

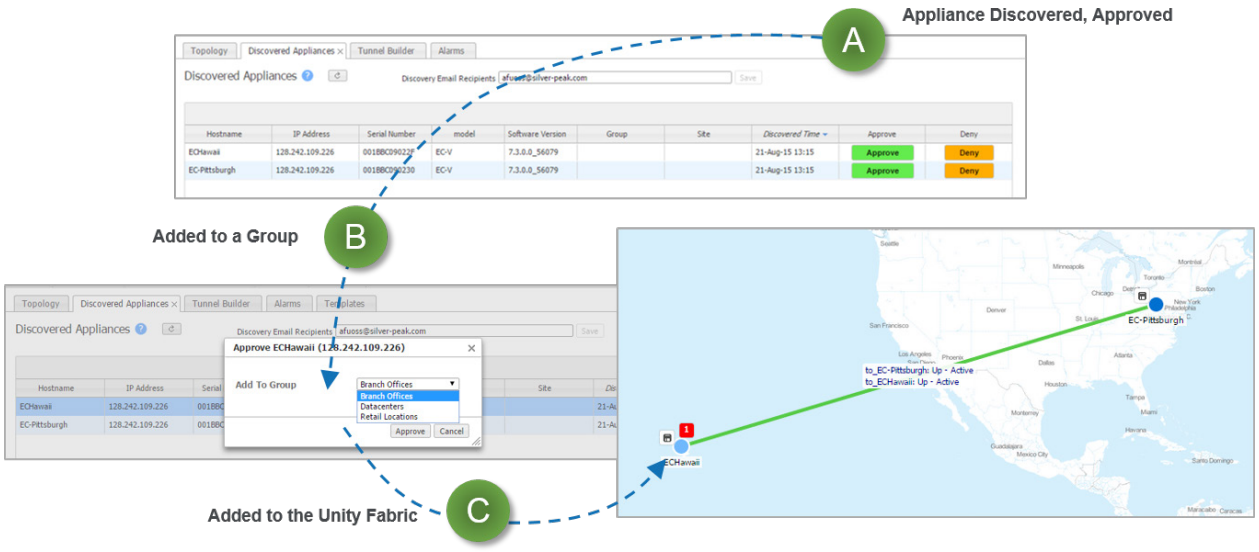

Figure 2: Centrally authenticating EdgeConnect.

Figure 2 illustrates what happens after the appliance is discovered. The administrator at Headquarters recognizes the device and approves it. Upon the approval, the appliance is added to a group of all "Branch Office" appliances, and is then added (policies are pushed) to the Unity fabric. The new device then appears in the Orchestrator topology map.

At this point, Orchestrator builds a secure tunnel to the EdgeConnect device, establishing a management plane. EdgeConnect is then ready to accept local profiles that have been pre-established at Headquarters.

[2] *In a smaller (thin) branch, there may not be an edge router, and EdgeConnect serves as the next hop gateway for branch hosts.*

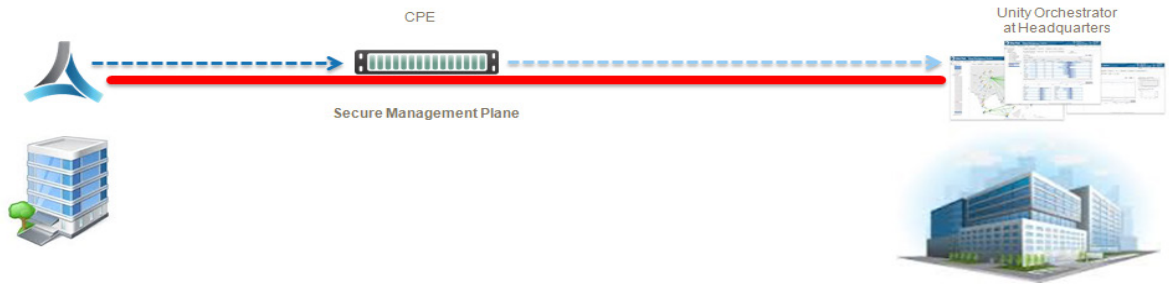

Figure 3: Secure management plane between EdgeConnect and Orchestrator.

Local EdgeConnect Deployment Profiles

Orchestrator allows branch offices to be provisioned quickly and consistently—configuration drift is easily avoided in an EdgeConnect rollout. Of course, planning takes place before this process to ensure that EdgeConnect profiles are defined before they are deployed into the branches, as all profiles obey the global business intent policies of the enterprise. Figure 4 shows a sample profile for a small branch office that supports voice, data-intensive applications, and database replication.

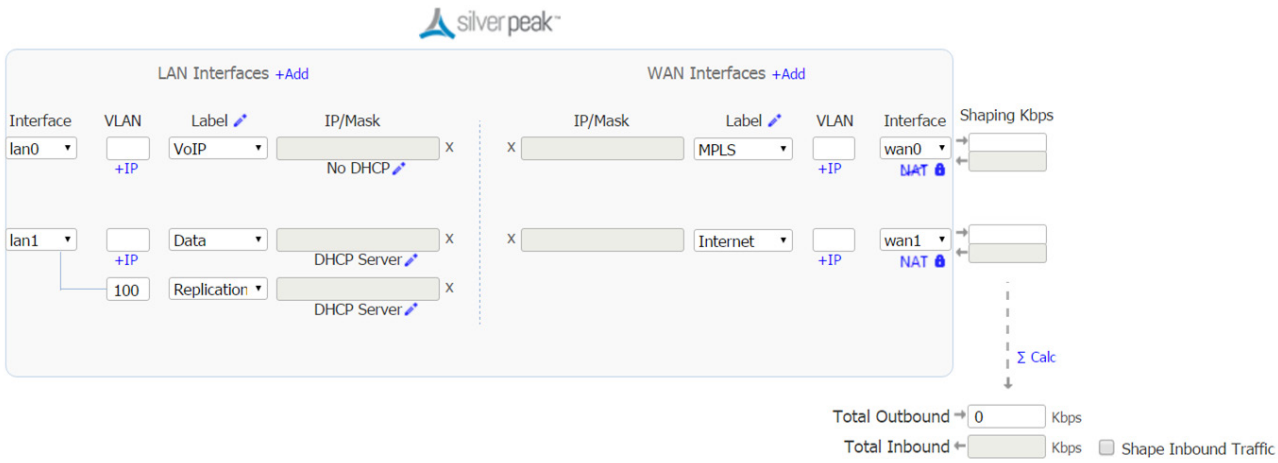

Figure 4: EdgeConnect profiles deployed into branch offices.

Profiles provide "personality" to the EdgeConnect instances. For example, they may describe small, medium, or large-sized branch offices. Figure 4 shows the LAN ports and the WAN uplinks that are available in a small branch profile.

LAN ports have interface labels assigned to them that describe the traffic types and applications that will use them. For instance, voice traffic is assigned to lan0, and the different VLANs on lan1 carry Data and Replication traffic.

The interfaces are also labeled on the WAN side. Here, wan0 connects to an MPLS service, and wan1 connects to the public Internet. Both WAN interfaces are encrypted and the Internet uplink is configured for Network Address Translation (NAT). Combining these port configurations into a single profile simplifies the installation process and prevents configuration errors.

Most importantly, network interfaces are assigned to labels rather than to individual port names. This facilitates end-to-end consistency and the ability to define the behavior of these interfaces.

Some local overrides may be made to the profile when applying it to the branch. For instance, in Figure 5, the administrator can change the WAN links (swapping wan0 and wan1) to match how the EdgeConnect appliance is physically connected.

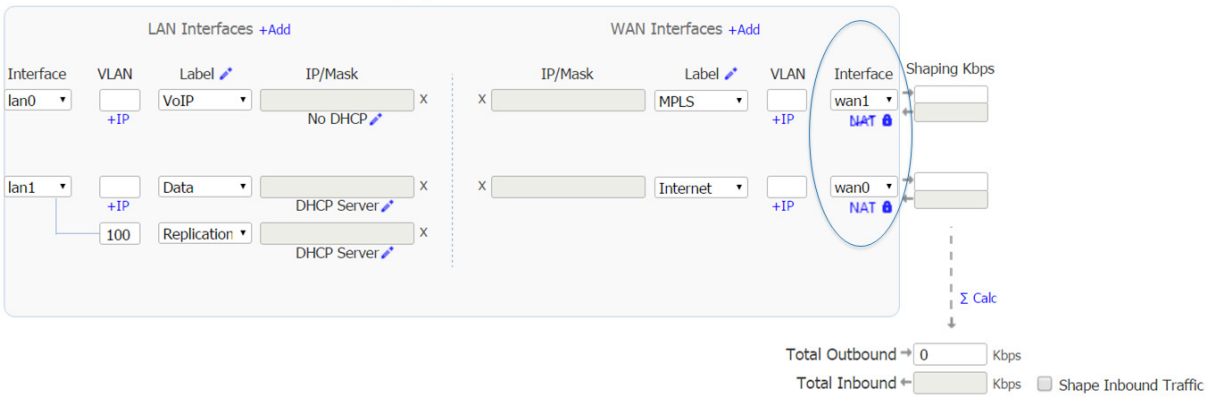

Figure 5: Modifying the profile to match the individual EdgeConnect.

In Figure 6, the Seattle EdgeConnect has been added to a mesh configuration and connected to other Silver Peak EdgeConnect instances, in this case, Chicago, Dallas, Denver, and Los Angeles.

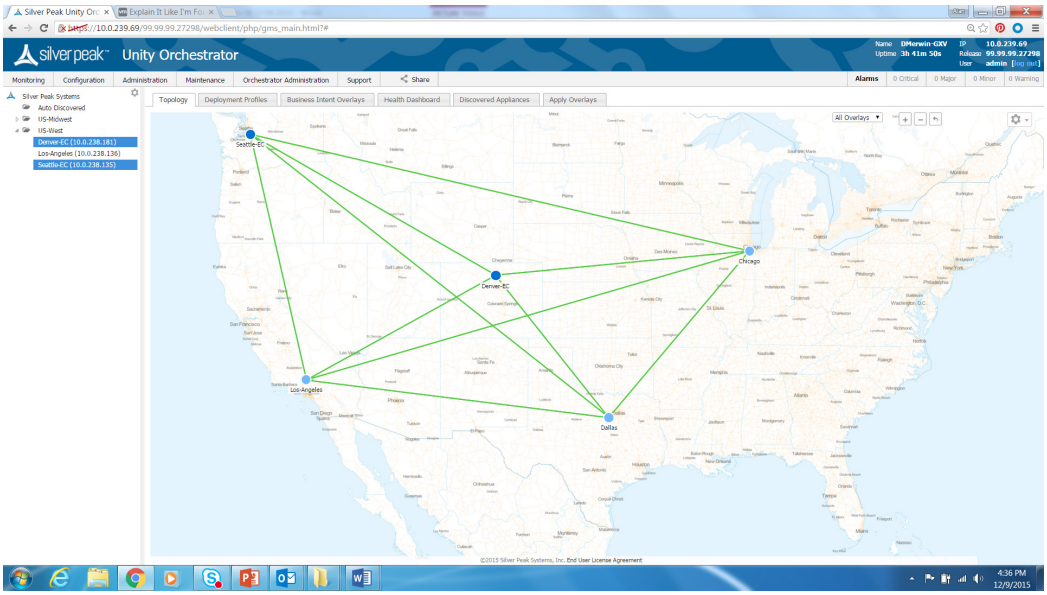

Figure 6: The topology view after the registered EdgeConnect has an assigned profile.

Global Business Intent Overlays

SD-WAN overlay networks are created independently from the physical network, and from each other. Each overlay has a unique tunneling mechanism that keeps its behavior—including topology, security, and forwarding rules—separate from the others.

This design principle allows for high-scale and secure application segmentation. Each overlay scales automatically as EdgeConnect endpoints are added to the SD-WAN fabric, and configuration integrity is maintained as each site maps the local profile into a global overlay.

Figure 7 indicates a common infrastructure (physical MPLS, Internet and cellular networks) serving individual needs of separate overlays. In this case, there are three separate overlays for Guest Wi-Fi, Enterprise Applications, and VoIP. The overlays apply globally to the enterprise network, and the local profiles for each Silver Peak EdgeConnect instance map the overlays into the LAN interfaces at the branch.

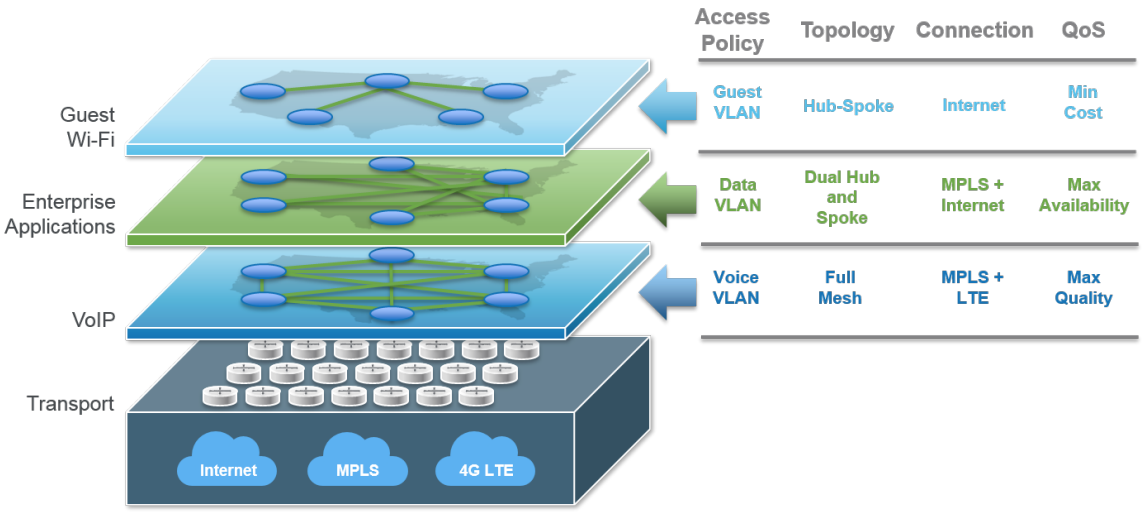

Figure 7: Sample business intent overlays.

At a logical level, different topologies and behaviors define the requirements of these applications. For instance, you may not want the same topology for CRM applications as for voice-over-IP (VoIP) applications. A dual hub and spoke configuration, with each of two data centers acting as hubs, would make sense for CRM, whereas VoIP typically would be configured in a full mesh to each destination.

Business intent overlays follow, and benefit from, the operational models of compute virtualization:

- Independence from the physical layer is maintained because the overlay decouples the delivery of business intent and applications from the vendor choice and hardware refresh cycle of the underlay (physical network of switches and routers)
- Secure physical, virtual, and control networks are isolated because each overlay describes a logical network for the application that can have a different topology—including addressing and access control—from the physical network
- High availability (HA) and ample bandwidth are facilitated via integration with route policies including dynamic path control (DPC) techniques that emphasize HA, maximum throughput or load balancing; applications are segmented according to required service-level guarantees such as minimum bandwidth or Quality of Service (QoS)
- Application visibility provides full knowledge and control of all applications crossing your enterprise WAN with real-time graphs at the Layer 7 application level, including web services over HTTP(s)

The single-screen user interface for business intent overlays is shown in Figure 8.

For a particular overlay, you select the topology, including participating hubs, which may be connected in a mesh, partial mesh, or hub-and-spoke model. You also select the access policy, brownout thresholds, underlying transports, and a bonding policy that may be optimized for either real-time (high availability) or data intensive (high throughput) applications.

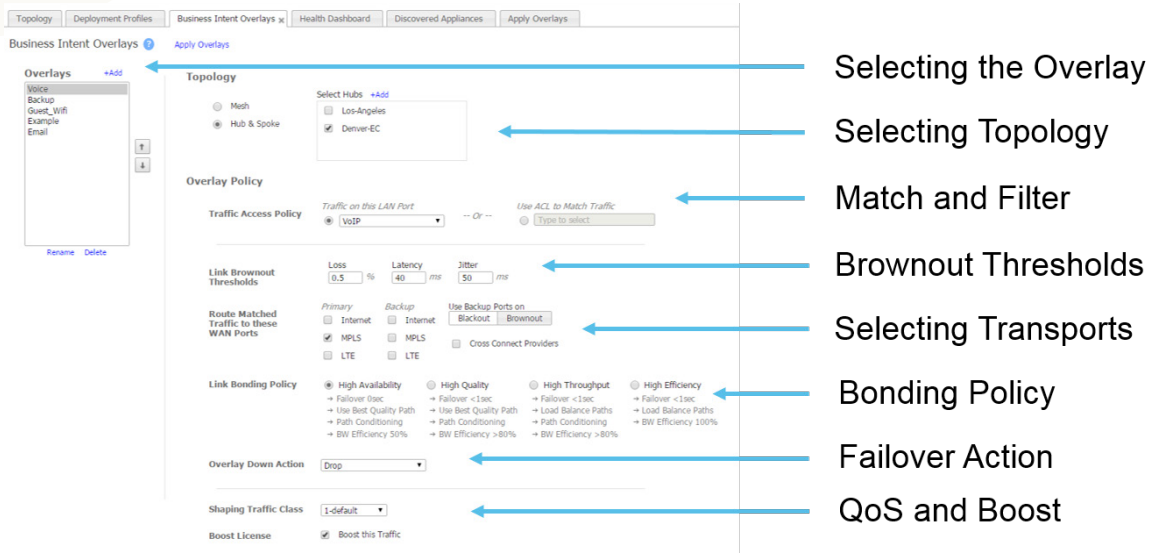

Figure 8: User interface for business intent overlays.

Security: Micro-segmentation and WAN Hardening

Why have multiple overlays? Why not handle all applications within a single overlay? The answer lies in the fact that with parallel virtual networks, each isolated from the others and from the underlying physical network, we can provide more-targeted policies that accurately match business logic to the underlying network.

A virtualized SD-WAN with micro-segmentation provides significant security benefits. In fact, this fine-grained segmentation approach produces a zero-trust architecture wherein security controls can be applied to a very small group of resources—for instance, a particular branch location's use of a CRM application.

Micro-segmentation is a best practice approach to security, but difficult to apply in WAN environments. The inherent security and automation capabilities of the EdgeConnect solution make micro-segmentation feasible across the enterprise WAN. In this model, fine-grained network controls enable security policies to be applied flexibly to a logical network interface (Figure 9).

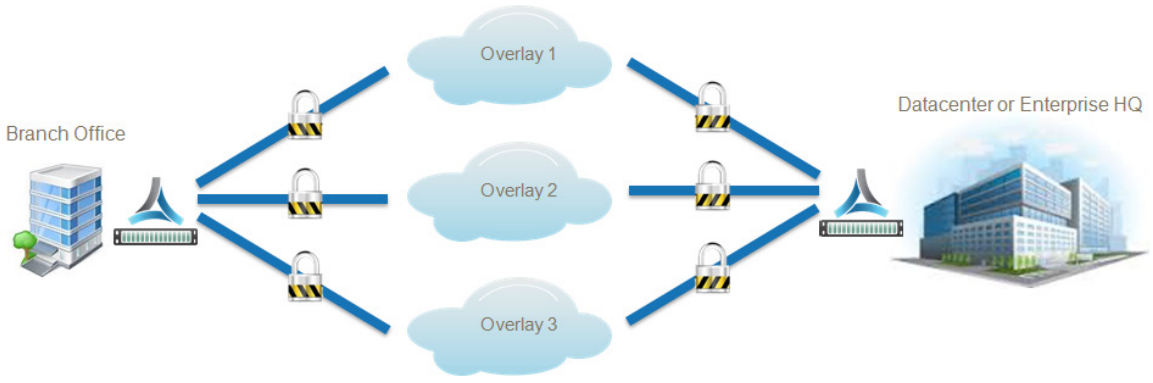

Figure 9: WAN hardening in business intent overlays.

With a virtualized network, it is much more straightforward to associate a security policy with a workload. Furthermore, each hardened interface is encrypted.

Conclusion

The integration of zero touch provisioning with EdgeConnect profiles, and the further integration of local profiles into the global business intent of Silver Peak's virtual overlays, fulfills the promise of SD-WAN with virtualized wide area networking. Enterprises are assured that their business needs are folded into the enterprise WAN, that the best forwarding decisions are made at any given time, ensuring that individual workloads are entirely secure in a zero trust model.