appgate

Appgate SDP Reference Architectures

Type: Technical guide
Date: August 2023
Applies to: Appgate SDP v6.2 or later

Table of Contents

INTRODUCTION

Security and compliance are now mainstream enterprise requirements with visibility and support at the highest management levels. Previously, security was typically added after the fact and often seen as a barrier to business and innovation. With the software-defined perimeter (SDP) approach, the conversation has shifted to how fast and seamlessly SDP (Zero Trust Network Access (ZTNA)) can be deployed at scale into complex enterprise environments to achieve Zero Trust security maturity.

Often the main barrier to SDP or ZTNA adoption is how to make it work with legacy networks. Many large organizations are invested in significant legacy network infrastructures that cannot simply be taken down and replaced wholesale over a weekend. This means that new and old must live alongside one another for a period. During this time, old concepts like having three huge redundant points-of-access into the network, i.e., in Europe, Asia and the Americas, can gradually be dismantled and replaced with the fully distributed model which underpins SDP.

Appgate SDP, the industry's most comprehensive ZTNA solution, is a custom implementation of the SDP model. It implements the core architecture principles and has filled in the gaps with additional capabilities. The most unique aspects are that it operates at the network layer, providing direct connectivity from the user to multiple protected networks/resources while operating in real-time, responding to environmental changes as they happen.

Appgate SDP has some very specific advantages over traditional legacy infrastructures:

- ✓ Created for the hybrid cloud: Appgate SDP was designed to protect all enterprise and cloud resources, including transient workloads. It has a flexible, distributed deployment model to suit any architecture, automatically detects server instance creation, and leverages user and server attributes to determine access. Appgate SDP also bridges and integrates hybrid cloud infrastructure, controlling authenticated user access to appropriate cloud resources, wherever they may be.
- ✓ Seamless integrations: Appgate SDP can reduce cost, complexity and effort of configuring third-party access, privileged user access and cloud infrastructure security. It combines authorization, encryption and access control in one system, replacing many traditional solutions. Appgate SDP also integrates with identity management, multi-factor authentication and SIEM solutions, allowing enterprises to maintain their existing security infrastructure. This enforces strong authentication and enables better integration of security requirements into identity management life cycles.
- ✓ User-centric network security: Appgate SDP provides application and service-specific authentication and authorization to control network access inside and beyond the perimeter. It dynamically creates a secure, encrypted network segment of one that's tailored for each user session, based on user attributes. Network access rules aren't written once and saved forever but are created and enforced in real-time.
- ✓ Security where it's needed: Appgate SDP works on a distributed model. This allows for a topology that closely couples security controls with the hosts/apps themselves. This ensures traffic is encrypted over any network being used to access the hosts/apps and that the actual access controls are very close to the hosts/apps themselves.
- ✓ Compliance is key: Appgate SDP can help the enterprise reduce regulatory compliance costs by reducing scope and audit complexity. Cloud providers have some new tools that can help with a multitude of regulatory controls, but Appgate SDP can further enhance these. It can reduce the number systems that fall within audit scope which, in turn, might eliminate the need for some of the regulatory controls themselves. Robust logging provides all the evidence necessary to meet most audit requirements.

This document provides an overview of a typical reference architecture for Appgate SDP and some interim architectures that will allow a phased migration to an SDP ZTNA model.

If you're unfamiliar with the unique security, operational and ROI benefits of Appgate SDP direct-routed ZTNA versus the limitations of cloud-routed ZTNA solutions, see https://www.appgate.com/zero-trust-network-access.

WHAT IS THE SOFTWARE-DEFINED PERIMETER?

The software-defined perimeter is a very different approach to the problem of securing today's networks. Developed within the Cloud Security Alliance (CSA), its aim is to stop network attacks on application infrastructure, while ensuring user productivity and improving security operations efficiency. The CSA SDP working group developed a clean-sheet approach that combined on-device authentication, identity-based access and dynamically provisioned connectivity. While many of the security components in SDP are well-proven, the integration of these three elements is fairly novel.

More importantly, the SDP security model has been shown to stop network attacks including DoS, man-in-the-middle, server query (OWASP10) as well as advanced persistent threats (APT). SDP is not just another DMZ add-on to an existing set of security controls such proxies or VPNs. Rather, it provides a Zero Trust access alternative to outdated tools developed before hybrid cloud became all-pervasive.

It is important to understand that Appgate SDP has filled in some of the gaps in the SDP model as defined by CSA and extended that model with its own very distinctive flavor. However, the key concepts remain:

- ✓ The first principle is that traditional office perimeters no longer exist. Hybrid workforces and their devices now need to access enterprise networks from anywhere and hybrid workloads are scattered across on-prem, data centers and the cloud.
- ✓ The next is built on the Zero Trust security principle of least privilege, "authenticate first, connect second" approach. Unlike a traditional network that connects users in various roles or groups to a network segment and then relies on application-level permissions for authorization, Appgate SDP ZTNA creates individualized permissions; as a user's situation changes, the individualized permissions changes too. This allows for much more fine-grained access control.
- ✓ The third principle is that the access controls should be placed as close to the protected hosts as possible. When the user attempts to access a resource—for example by opening a web page on a protected server—the Appgate SDP Client redirects the request to the closest Gateway via a secure tunnel. This in turn applies additional policies in real time to control access based on the user's network location. This premise means that Appgate SDP Clients can make multiple connections to multiple gateways at the same time to meet the user's specific needs.

The "authenticate first, connect second" approach ensures that only authorized users can connect to network resources. Resources are rendered invisible to potentially dangerous reconnaissance which reduces the attack surface and significantly improves security.

Having multiple Gateways (access points) makes Appgate SDP ZTNA very suitable for hybrid environments, allowing consistent access policies to be applied to legacy network, data center and cloud environments simultaneously. Newly created Appgate SDP Sites are independent of one another and can be easily deployed without any long lead-times; they simply require internet connectivity.

THE NETWORK-CENTRIC MODEL

The network-centric model has endured for over a quarter of a century. It evolved when the Microsoft-powered PCs that ruled the world needed linking together by using token ring/ethernet-based LANs. At that time the internet was becoming established, so TCP/IP became the obvious way to link these PCs, then buildings, and eventually data centers.

But TCP/IP was never designed to address all the different use cases to which it has been applied. Maybe the biggest issue being its inherent lack of any upfront security when establishing new connections.

Today, networks struggle as IP-based devices are no longer tied to users' desks and connected via CAT5/6 cabling to a managed infrastructure within a defined boundary. Securing such traditional networks has become increasingly difficult, especially given the complex and disparate range of network security tools deployed to try to compensate for TCP/IP's shortcomings.

Redundancy is often provisioned by deploying multiple similar systems, themselves creating yet more network complexity. These systems all need to be configured to work with users on the inside as well as ones coming in via VPN from the outside.

The need for high availability (and users' different geographies) often requires sites to be linked to many other sites over some form of WAN. While these sites may appear somewhat separate from one another, often the rules governing the WAN are fairly open so lateral movement between sites is relatively easily achieved. In reality, they look like one big network thus making them very easy to explore and exploit once a bad actor has established a foothold.

The network-centric model (cont.)

Today's typical company network might look something like this:

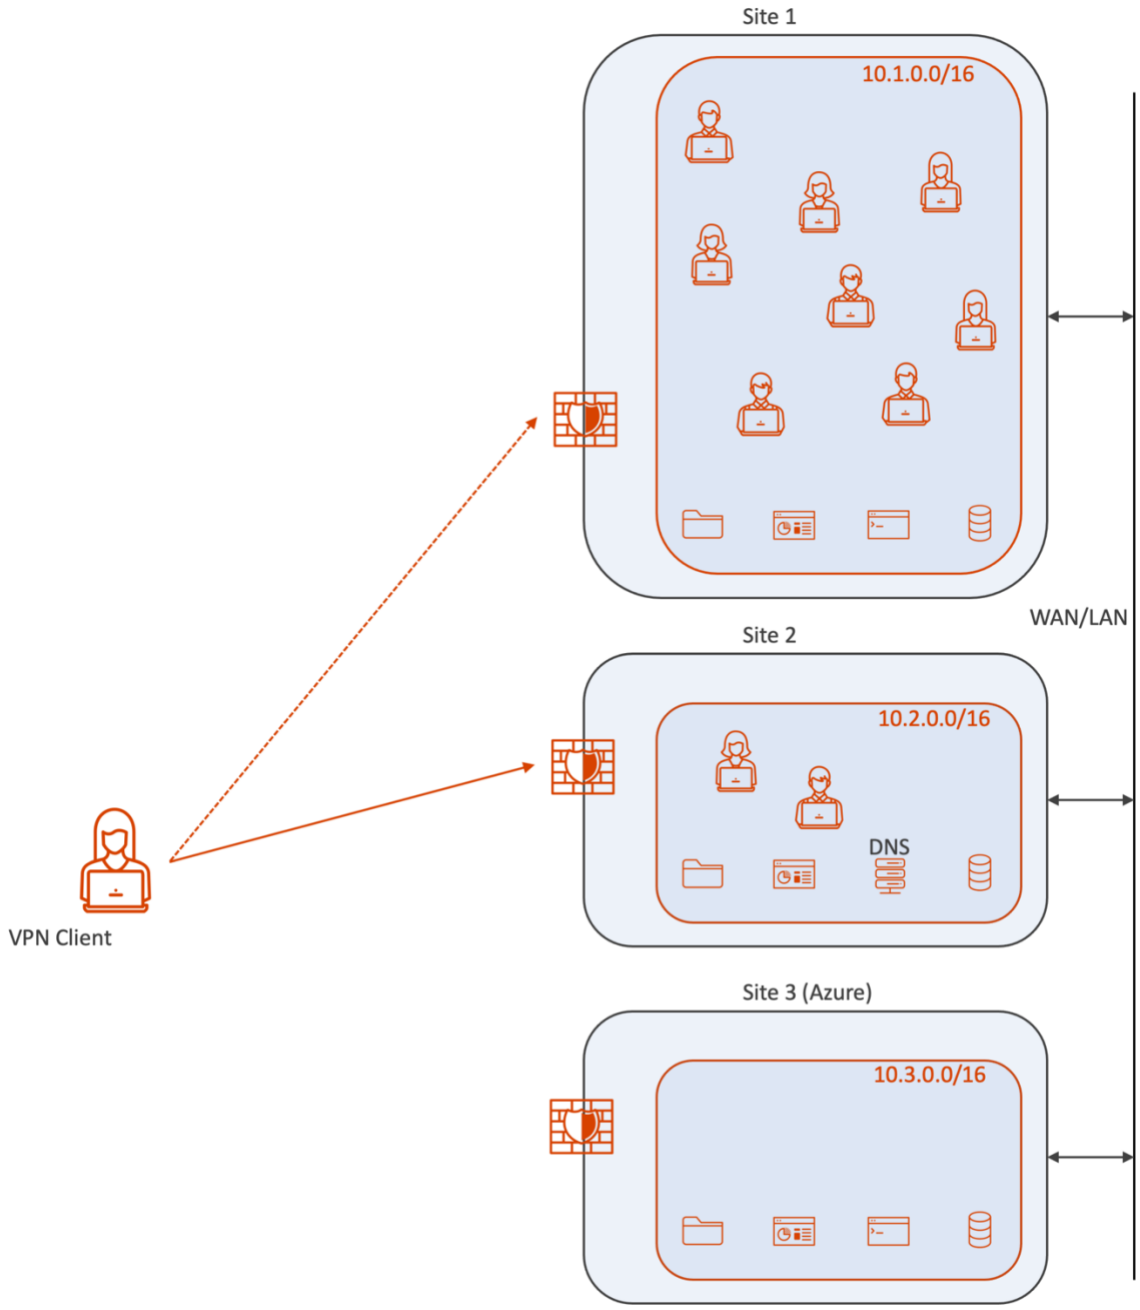

Figure 1

The network-centric model (cont.)

No longer fit for purpose

The Cloud: Cloud adoption is well established with almost all organizations using it to some degree. Often it is tacked on as shown in the diagram above (Site 3). But this model is not sustainable as many businesses will use hybrid cloud for the foreseeable future and cannot continue to backhaul all the connections to the cloud.

Complexity beyond comprehension: In some regards, networking is still seen as an "inside" thing that should provide some protection from the "outside." In reality IP-networks have changed little over the last quarter of a century except for the proliferation of point solutions that have filled most businesses DMZs in a vain attempt to secure this increasingly complex nightmare.

User behavior: Users have changed and expectations about where people are allowed to work are now very different and the office increasingly redundant. Many organizations now use contractors and third parties that are not employees as such, but still need secure access to corporate IT assets.

DevOps: Systems are no longer the static machines they used to be. The IT environment is now dynamic, and legacy network and access infrastructure is no longer fit for this purpose. Access needs to follow the IT environment in near real-time allowing users to be connected to their protected hosts even if they are only there for an hour or two.

External Threats: The speed at which new exploits are introduced is frightening. The attack surface of a traditional (multi-site) network is vast. The internal exploitation paths are countless. Taken together these make the job of defending these networks next to impossible unless, of course, there's a way to cloak all access and make the network effectively disappear.

HOW DOES APPGATE SDP WORK?

Appgate SDP Zero Trust Network Access utilizes user context to dynamically create a secure, encrypted network “segment of one” tailored for each user session. The Appgate SDP usage model is the same wherever it is deployed and can be described in a few simple steps. (Figure 2)

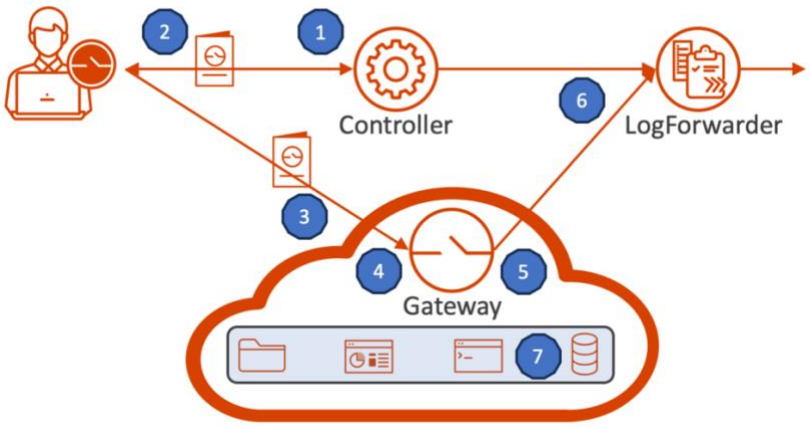

Figure 2

Step-by-Step

1. User authentication: The user authenticates to the Appgate SDP Controller, which is optionally connected to an enterprise’s IAM system (LDAP (AD), OAuth, etc.)
2. Controller applies required Policies: The Controller applies Policies assigned to the user based on user attributes, groups, and context, and then issues signed Entitlement tokens listing the resources the user is allowed to access.
3. User accesses resources: The authenticated user can now request access to protected resources behind a Gateway.
4. Gateway evaluates user Entitlements: The Gateway evaluates Entitlements in real-time, ensuring that all Conditions are met. For example: network location, time of day, device health, or service metadata such as security groups. Users may be prompted for additional information, such as a one-time password.
5. Gateway opens connection to resource: If the Gateway determines that required Conditions are successfully met, it opens a connection to the protected resources specified for the user.
6. All steps in the process are logged: Throughout the authentication and authorization process, all the steps pertaining to the user, resources, and decisions made by the Controller and Gateways are logged. These logs are normally sent to an enterprise SIEM using the LogForwarder for additional analysis and action. There is an integrated LogServer which can be used initially which is outside the scope of this Reference Architecture document.
7. Detects new services: The Gateway constantly monitors for the creation of new hosts/services, and based on this metadata and the user’s Entitlements, adjusts user access as necessary.

The Components

The entire Appgate SDP system is designed to be distributed and to offer high availability. One key element that underpins this is the use of single packet authorization (SPA) to hide TCP ports on servers/appliances until a very specific wake-up packet is received from a connecting device. This allows the appliances to be cloaked from hostile actors thus enhancing availability for authorized users.

Let's look at how the individual components contribute to the highly available operation of the overall Collective.

Controller

For full high availability operation, Clients are designed to handle multiple DNS A records to obtain the list of all available Controllers they can try. Otherwise, a load-balancer can be used to distribute the traffic based on performance, weighting, geography, etc.

The Controllers utilize multi-master database synchronization for high availability based on an eventual consistency concept. The communication between Controllers happens over bi- directional TCP port 443 (mutual TLS connectivity). This includes the database synchronization which is mainly used by the Controllers to store the static configuration—Policies and Entitlements, etc.—and the occasional changes to the user's IP addresses. Most database operations are therefore read operations used to generate the tokens which contain the session information. The use of eventual consistency synchronization means there is no real-time syncing required for the system to operate and no waiting until all databases have processed a record update before handling the next transaction.

The Controller is normally connected to an identity provider [IdP], which serves to validate user authentication and act as the source of user attributes and group memberships. The IdP may be located anywhere, as long as the Controller can access it. Appgate SDP supports AD, LDAP, OAUTH, SAML and Radius-based identity systems. Multiple IdPs can be specified in the Controller that will be tried sequentially. If no IdPs respond, then the Client will try the next Controller that it is aware of.

Client

The Client must use a specific pre-shared SPA key to open the TCP (UDP) connection with Appgate SDP appliances.

The Client connects on port 443 (TLS) and passes only system control data. At first connection it checks that the Controllers certificate is valid as it will be used in this and subsequent communications to verify the authenticity of the Controller.

The Client asks for authentication credentials and optionally the device needs to go through an on-boarding procedure which might include the use of use a one-time password (using the built-in or an external Radius service configured by the customer). The device will receive an on-boarding cookie that glues user credentials with device ID, thus marking the device as friendly. When on-boarding is disabled, a valid user (username and password) will only be able to use friendly devices.

The Client generates its own private/public key pair and based on this receives a Client Certificate signed by the CA that will be used for all subsequent mutual (D)TLS connections between the Client and Gateways. This traffic from Clients to Gateway comprises both system control data and application data.

LogServer/LogForwarder

As already mentioned, the Collective comes with a choice of functions that allow audit logs to be collected and distributed. This ensures that there is an accurate record of all accesses to the protected environment. The deployment of these functions are not covered in this paper.

Portal/Connector

There are two other functions that allow clientless access to the Collective. The Portal is provided for user access while the Connector is aimed more towards unattended devices such as IoT devices. However, the Connector can also be used to connect small branch offices to the Collective. These functions are not covered in this paper.

Metrics Aggregator

The Metrics Aggregator function provides a means of exporting Prometheus metrics from the Collective. As with most enterprise access solutions, the use of monitoring tools is vital to ensure the health of the infrastructure. This is not covered in this paper.

Gateway

After successful authentication, the Client receives a number of Site-based Entitlement tokens signed by one of the Controllers which contains the list of all available Gateways for that Site and the list of descriptive network Entitlement Actions for that user on that Site. (A Site is a special Appgate SDP term denoting a logical grouping of protected hosts. There might be two different Sites defined at one physical location or indeed one Site might span across several physical locations.)

The Client connects to one of the Gateways in the Site based on a pre-set weighting. No load balancers are required in front of the Gateways. The Gateway will verify the token's signature before using any defined Actions to create a micro private firewall for this specific user session.

Gateways will receive mutual (D)TLS traffic on port 443 from the Clients and will send a specific pre-shared SPA key to establish a mutual TLS connection on port 443 to Controllers, LogForwarders, etc. Connections to other appliances do not need to be in place for the Client to be able to use a Gateway. There is no communication between Gateways.

Gateways can be deployed in a number of different ways. Before we go on to look at the broader network architecture options available, it is worth just focusing on the Gateways themselves to understand their internal deployment options.

Component deployment

Below are shown some simplistic deployment options for the Gateway:

Single interface Site

This is a typical deployment scenario where the Gateway sits in some sort of DMZ or Security zone. Rules can allow inbound traffic to the single NIC and allow traffic from the NIC to the protected resources. The performance of the NIC is restricted as it must handle traffic in both directions.

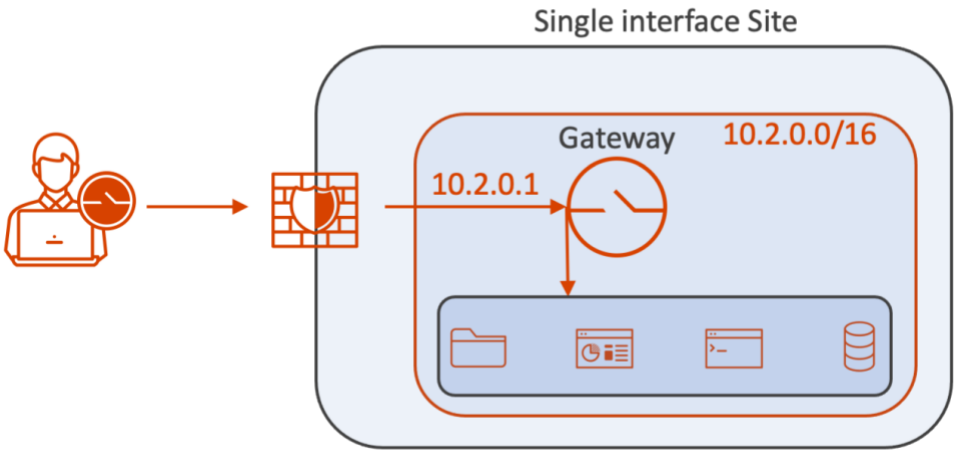

Figure 3

Dual-interface Site

A quite typical deployment scenario where the Gateway sits inline. It may still reside in a DMZ or security zone behind a firewall; however it can often be positioned to sit alongside a firewall which might be handling the outbound traffic.

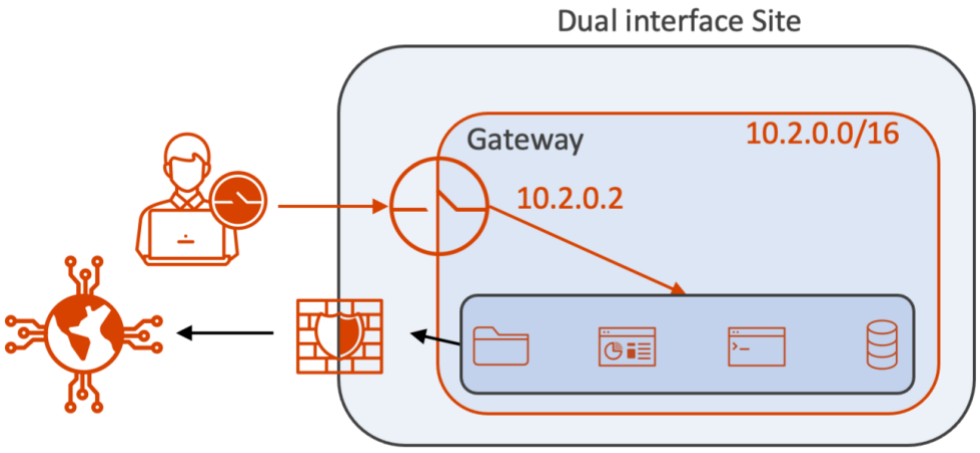

Figure 4

Multi-home Site

This deployment scenario is useful where the Gateway may need to access different subnets isolated from one another. Maybe one requires a higher security level such as hosting PCI-related services.

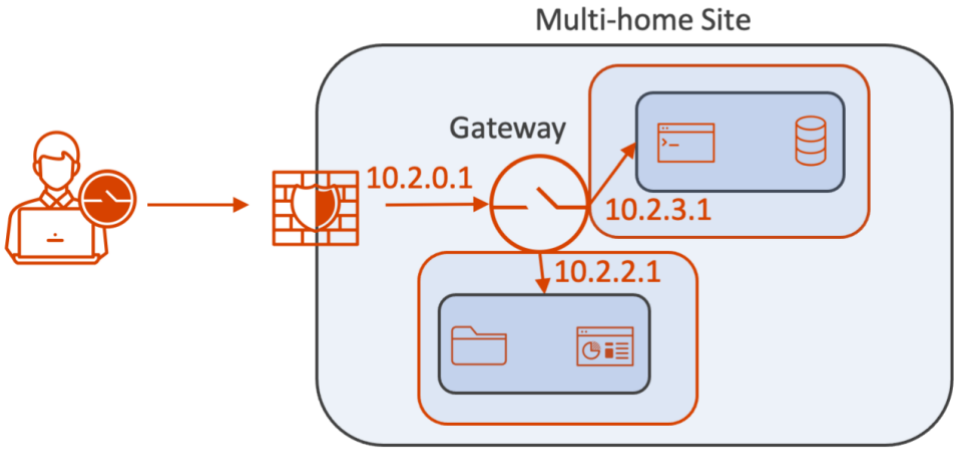

Figure 5

Site Without NAT

By default, Gateways use SNAT. Without SNAT, the user's tunnel IP address is presented on the network instead of the Gateway's IP address. This mode is required when attempting to make reverse connections (towards the user) or when stateful failover between Gateways is a requirement. This usually requires some routing changes on the network to ensure traffic towards the user is routed correctly.

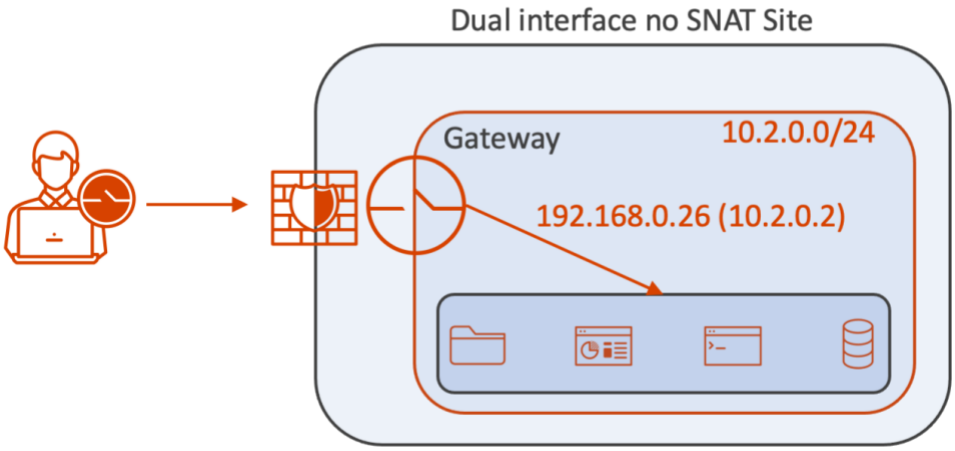

Figure 6

THE SDP JOURNEY

Most organizations have existing (typically complex and messy) networks, hosting numerous users and production applications. SDP can still be deployed in these transitional environments to very good effect. So let's look first at the textbook Appgate SDP ZTNA scenario and then at three different scenarios that might form part of any implementation journey.

- Legacy network with users on the inside
- WAN with single tunnel (from external users)
- WAN with multiple tunnels (from external users)

Many organizations have a heavy investment in legacy network infrastructure with many mature systems that may be operationally fragile. Their related access policies are not well documented or understood and are often enforced at different points on multiple different systems. There are also sizable installed user bases making any change harder to plan and implement.

Why make the journey?

Many of today's network components already can implement access rules based on the user's identity, but these are not widely implemented because the cost is high because there is no centralized way to manage policy and the benefit is low due to the topology of today's networks.

A key benefit of the Appgate SDP model is that access rules can be a user-based segment of one. By using attributes (group memberships) and context (location), it is very easy to implement Policies that initially allow a group of existing users access to a whole subnet. Because the Policy objects behave independently, each can then be fine-tuned when the time is right until the ultimate goal of a segment of one is achieved for all users.

In a textbook software-defined [perimeter deployment the company's network might look very different from what is considered a traditional network topology.

A single Appgate SDP Collective will span all locations where there are resources that require protection. Controllers are deployed in the Cloud and/or at different locations for HA/resilience. The user traffic to Controllers is minimal and infrequent so location/geography is only of secondary importance. The Controllers do not really feature in the different scenarios discussed later in this paper.

Because appliance-to-appliance traffic uses SPA protected mutually validated TLS, it can all be sent across the internet. Where appropriate, access to appliances can be further restricted using firewall rules.

User access is also via SPA-protected, mutually validated TLS, so the SDP system is cloaked for all unauthorized users coming from the Internet.

Appgate SDP Sites are largely independent of one another but where protected hosts require some form of server-to-server communication between Sites, enterprises can still use Appgate SDP for this. (Note, this is not covered in this document.)

Deployment scenario A: Textbook SDP implementation

In this scenario (see Figure 7 on page 15):

- Appgate SDP is deployed on every Site
- Users do not co-habit inside the network
- There are two Controllers – usually in different locations/zones
- Sites are independent (the WAN may no longer be used)

The Controllers could be deployed using Appgate ZTP into Appgate's hosted environment. They share the same DNS name (and two DNS A records) to provide high availability.

Once seeded, the Gateways can communicate with the Controllers. The local Gateways use their external DNS names thus avoiding any potential issues with users moving from the inside to the outside; this also allows for easier relocation of appliances during any phased SDP deployment.

Clients connect to one Gateway on every Site where granted an Entitlement. One of the Sites will host the internal DNS server used by the Client, so a DNS Entitlement must be included for that Site.

There is a minimum of two Gateways per Site to provide HA and load balancing. In the event one fails, then the Clients automatically try any others listed for the Site. These Gateways can provide access to the protected applications directly; however, for end-to-end HA, application access can be via some form of load balancer using virtual IPs.

Management network

Not discussed above is any reference to management networks that are commonplace in today's data centers. All Appgate SDP appliances come with the ability to support multiple NICs and to configure one as a dedicated management NIC. It is therefore very easy to join all the appliances to a management network for all the admin types of traffic such as SSH.

Delegated administration Policies can also be tailored so that only connections from the IP range of the management network would be granted admin UI access.

Deployment scenario A (cont.)

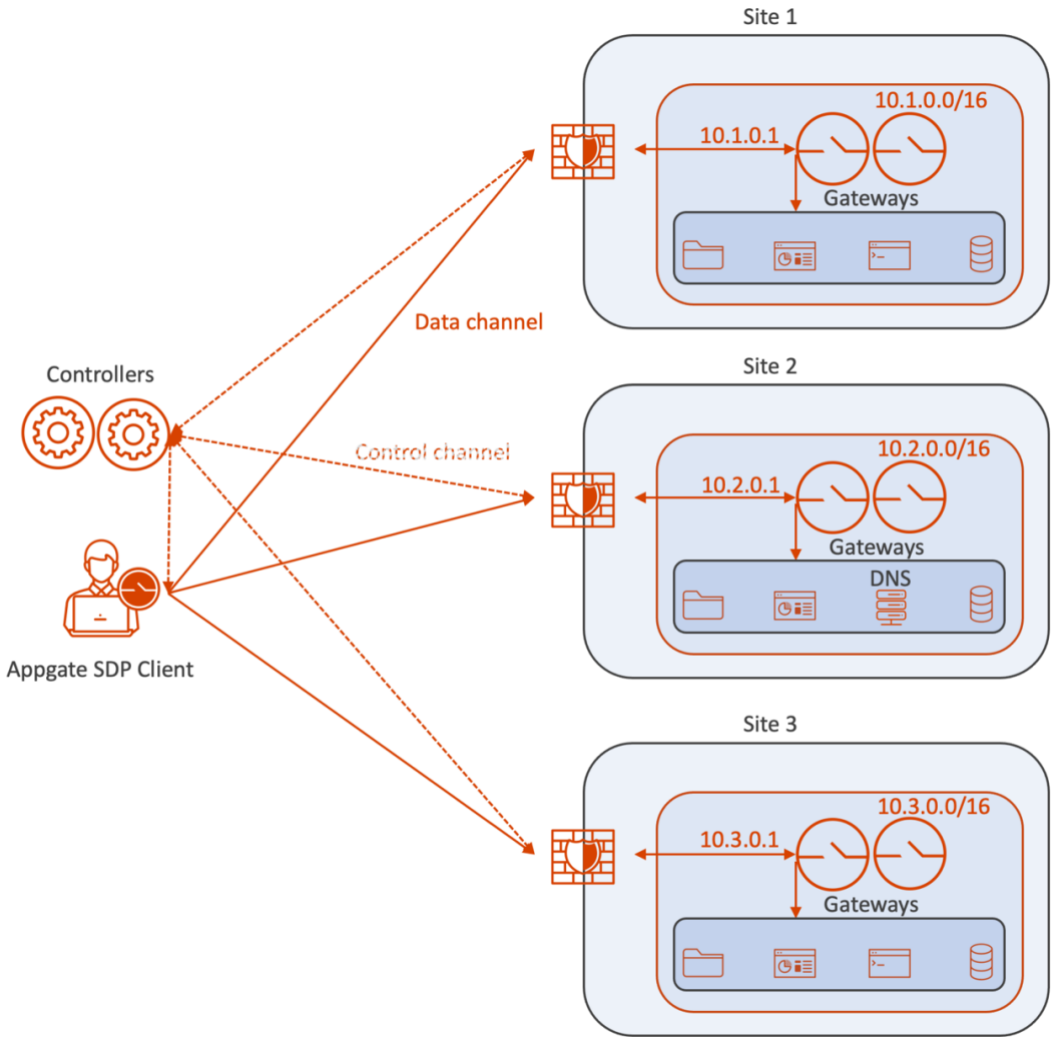

Figure 7

Benefits ✓ Encrypted connections (no VPN required) ✓ Network resource cloaking (SPA) ✓ Device on-boarding ✓ User-based rules (no NAC required) ✓ Multi-Gateway access (no load balancing) ✓ Users live outside the protected network
Issues ✓ Backwards compatibility with today's topology
Mitigation ✓ Use the intermediate topologies suggested below

Deployment scenario B: SDP with users on the inside

In this scenario (see Figure 8 on page 17):

- The Appgate SDP system is deployed on only one Site
- Users co-habit inside the network
- The protected resources are segmented from the users

The first key milestone on the journey to SDP is to split users away from the network where the (protected) hosts reside. Organizations sometimes look to NAC solutions to solve this problem. Originally, NAC was driven by the need to enforce access policies for Windows PCs then it grew to controlling access from personally owned devices connecting to network environments. NAC's application footprint is quite limited as it does not extend well beyond the network boundary or into the Cloud. It also has the problem that users and protected hosts are still sharing the same network and are often relying on VLANs or crude permission based controls (LDAP) for segmentation once access is granted.

NAC wraps a mixed-use legacy network with (somewhat limited) access policies which selectively allow only friendly actors to join the mix. SDP's primary goal is to utilize powerful access policies to selectively allow friendly actors to connect to defined resources (or networks) and block all else. Confusion sometimes arises in that both implement mechanisms that allows users to be granted access to protected hosts.

When using SDP (to solve a problem where a traditional NAC solution might be considered), there is a prerequisite that users MUST first be segmented away from the protected hosts. This could be as simple as splitting the WIFI from the wired network or creating a couple of subnets and then positioning an Appgate SDP Gateway such that it can securely connect them back together again. Once the split is done, SDP does not care if the user network is at the same location (e.g. the WIFI network) or if users are working remotely (VPN style); the same controls can be used for both (although the Policies might be different). And if the split has been done effectively - any lingering NAC requirement should automatically

In this scenario the external capabilities of the Appgate SDP system need not necessarily be used. The user network would allow access to the Appgate SDP Gateways (but not the protected hosts) and the Client would establish a secure tunnel to the Gateways and on to the protected hosts – possibly using a subnet wide access rights if more specific access rules were not known at the time.

Directing user traffic to the Gateways

When an internal user wants to connect to the Gateways there is the option to set the internal DNS to resolve to the Gateway's external IP address. On Site 2 the user was given the external IP address of the Gateway (56.67.4.20) so their traffic was hairpinned on the firewall (blue line). This avoids any issues when users move from the internal network to remotely. Another option is to "Connect via local network" which can be enabled for a Site. This means the Controller will issue Entitlement tokens with the internal hostname/IP of the Gateways for the Site the user is currently on. On Site 1 the user was given the internal IP address of the Gateway (10.1.0.2) so their traffic went direct to the Gateway (blue line).

Routing application traffic

Once connected, the Clients build routing tables based on the Entitlements. These routes point to the Appgate SDP tunnels for access to all the protected hosts. The Gateways that control this access should ideally serve the protected hosts on each Site separately, in which case the Client traffic could be routed over the WAN/LAN (SDP model). Alternatively, the Gateways can serve all the protected hosts on both Sites so the traffic from the Gateways can be routed over the WAN/LAN. This also allows controlled access to locations that are not yet enabled with Gateways. (Figure 8)

Deployment scenario B (cont.)

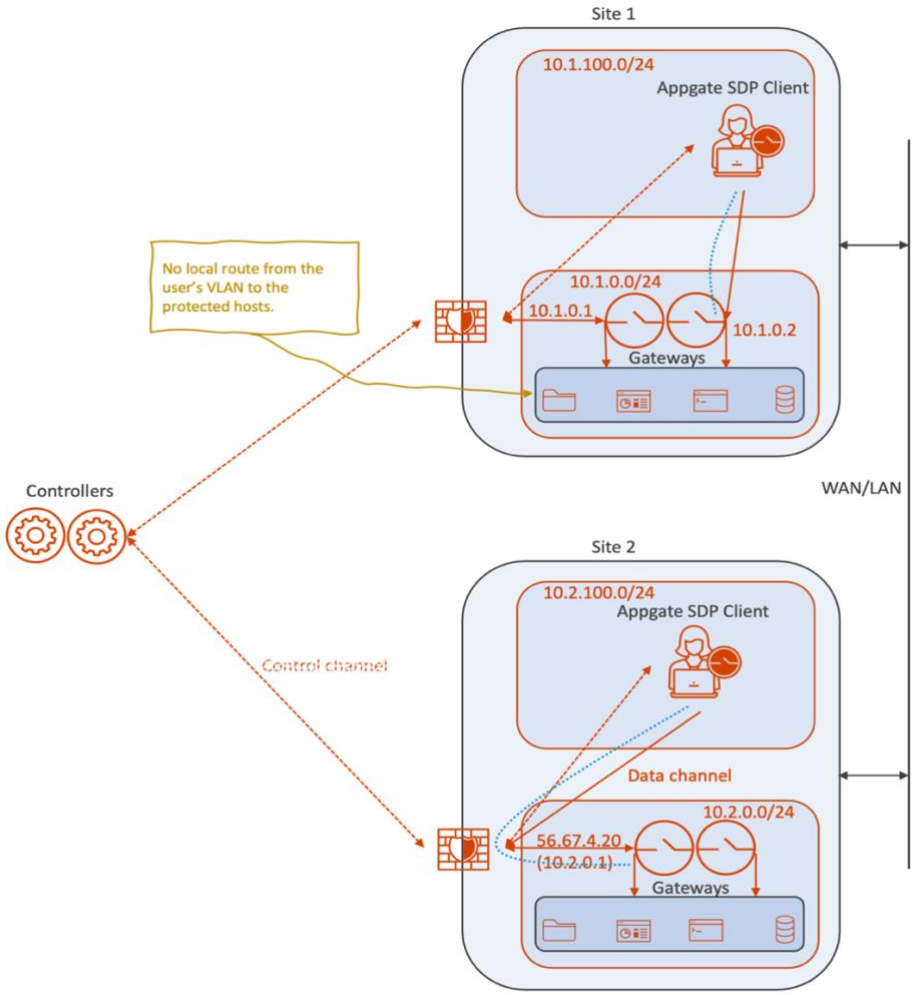

Figure 8

Benefits ✓ Encrypted connections on user network ✓ Device on-boarding ✓ Protected hosts are not visible on the WAN/LAN ✓ User based rules (no NAC required) ✓ Users live outside the protected network
Issues ✓ Need to keep users away from hosts ✓ When used outside the enterprise network, internal/external DNS can be challenging
Mitigation ✓ Need to do some network segmentation ✓ Follow the DNS guidelines in the manual

Deployment scenario C: SDP with single tunnel (VPN look-alike)

In this scenario (see Figure 9 on page 20):

- Operates as a simple VPN replacement.
- Only a single tunnel is used (by the Clients).
- Site failover is provided.

This scenario can be driven by the user environment such as rolling out new devices or be driven by the network environment such as there is no separate internet access at some Sites within the WAN. However, the Appgate SDP deployment model for this scenario is conceived such that the system can be extended/grown over time toward the more distributed SDP ZTNA model.

The Controllers should be cloud-based using the ideal SDP model; as explained earlier DNS round-robin will always allow users to find a Controllers if one was down. At least two Gateways are deployed, usually at the location of the existing VPN infrastructure. If two points of presence are used for VPN access—for instance, one in Europe and one in the Americas—then two Sites should be defined and Gateways deployed in both locations.

There are two things that now need to be considered before rushing to the configuration screens:

- How do I configure this now to make the SDP journey easy for administrators to avoid defining policies, entitlements and Sites multiple times)?
- How do I direct my users to one Site for now and multiple Sites later?

Future-proof configuration

At this point on the journey, the Sites are still linked by the company's WAN and with users connecting to Site 1 only; access to the protected hosts on Site 2 would still be over the WAN.

By default, entitlements need to be assigned to a Site, and because this feels a bit like one Appgate SDP Site, it is tempting to create all the Entitlements for just one Site. But, doing this makes migration harder when the time comes to split the Sites apart. So, this is not an ideal solution as all Policies and Entitlements would need redefining for the split Sites later in the Journey.

To support this interim stage, Appgate SDP provides a means of overriding this default so that Entitlements specified for Site 2 could be temporarily added to Site 1, allowing users to connect to Site 1 and still get access to their all their Entitlements.

To do this, define the appropriate Entitlements for each Site (as if the SDP model was in use) and then create a Policy per Site and specify an 'Override Site' in the Policy for Site 2. This will override the normal SDP model and force the Entitlements in that Policy to use an 'Override Site', this case Site 1. This means the system can be configured for multiple Sites even though there is only one operational. As stand-alone Sites are enabled (and moved away from being part of the company WAN), migration is as simple as turning off the 'Override Site' in each Policy. The Clients would then build additional tunnels to the new Sites and route the traffic there directly.

Directing user traffic to the right Sites

If you have just one Site then this is easy, however Appgate SDP is geo-aware, so with more than one Site, it can easily direct users to their nearest Site:

- Assign just one Policy for geo-location use. In the Policy, override the entitlements using *Override with nearest Site* under Override Site. This will force the entitlements to be assigned to the nearest geo-located Site. Geo-located Sites must be enabled in Sites by checking *Use for nearest Site selection*.
- Use the geo-location claim to assign either a Policy for Site 1 or another for Site 2. Override the Entitlements in the latter Policy to Site 2 using *Override by selecting from list* under Override Site.

The network continues to provide all the backhaul routes across the WAN.

Failover

This is also a good time to explain how the system would failover users of Site 1 (Europe) to Site 2 (Americas) if it went down.

One obvious solution might be to allow users to have two different Client profiles set up to connect to the two different Sites. Both Sites would allow access to all the apps either locally or across the WAN. The DNS would be used to direct traffic, so for the app crm.mycompany - DNS would return crm.europe.mycompany (which is what would be used in the European Site's Entitlements). However, when the European Site was down, then instead DNS for crm.mycompany would return crm.americas.mycompany (which is what would be used in the Americas Site's Entitlements). Not an ideal solution as all Entitlements have to be defined twice, differently.

To help, Appgate SDP has a built-in fallback Site awareness so with more than one Site, users can easily be switched if their main Site goes down. Site 1 has a fallback Site specified, in this case Site 2. The Client has *use fallback Site* enabled in their Policy. The fallback process is controlled by the Client and after a minimum unresponsive period of 30s the Client will decide to switch to the fallback Site.

Fallback capability should be used carefully, for instance by the sysadmins responsible for getting Site 1 back up and running. By limiting the numbers of users who can fallback to Site 2 in the situation where there is a problem with Site 1, will prevent Site 2 becoming overloaded and potentially failing as well.

In this scenario the external capabilities of the Appgate SDP system are brought into play but the fully distributed nature of SDP is kept for another day.

Two points of presence are used as active Sites. These Sites handle all user traffic then the back-haul connections to the protected hosts over the existing WAN just as a VPN solution would do today.

Deployment scenario C (cont.)

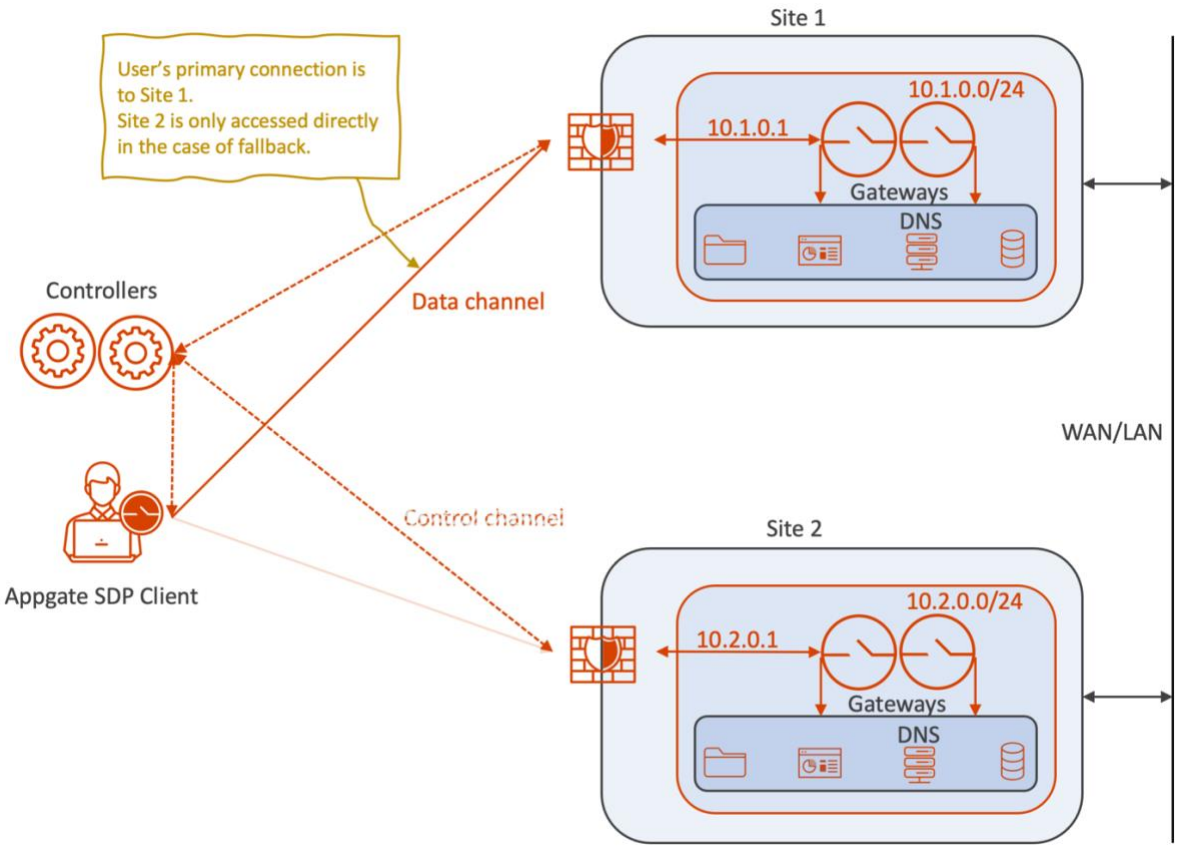

Figure 9

Benefits ✓ Encrypted connections ✓ Device On-boarding ✓ External services are 'cloaked' ✓ Geographic based access for users
Issues ✓ May require duplicate Entitlements on all Sites
Mitigations ✓ Use Override Site option ✓ Give users the choice of regional IdPs to use

Deployment scenario D: SDP with Multiple Tunnels

In this scenario (see Figure 10 on page 22):

- The Appgate SDP system is deployed in a distributed model
- The protected resources are only available from behind Gateways
- Some users might still co-habit inside the network
- The WAN has not yet been removed

In the multiple entry point scenario, the deployment model goes most of the way toward the ideal Zero Trust Network Access model. Two or more Gateways are deployed at the points of presence that are/were used for VPN access. Additional Gateways can be deployed on any new stand-alone Sites such as in the cloud. The network retains the (majority of) backhaul routes across the existing WAN.

To begin with, this model is likely to be used by the large enterprise that may have several points of presence and users might use more than one as they travel between regions.

An absolute Appgate SDP requirement is that the Entitlements (defines access to protected hosts) are defined by Site. At this point on the journey all the Sites are still linked by the company's WAN. Users are going to be connecting to multiple Sites to access all their allowed protected hosts.

Gateways normally use SNAT so users appear to be coming from the Gateways' IP addresses (Site 3). If SNAT is not being used, then the users' will appear to be coming from their tunnel IP address (Site 2). If SNAT is not being used on multiple Sites then their tunnel IP address will be advertised on more than one Sites at the same time across the WAN causing potential routing issues.

Avoiding routing issues across the WAN

To help in this situation, When you disable source NAT on Gateways, there is the option to specify an 'IP pool mapping' for the Site (Site 1). There is a further choice to *allocate* or *translate* this new range of IP addresses. This allows the default IP tunnel address that the Client uses to be overridden for a given Site. A 192.168.0.0/24 IP address might be allocated one from the mapped 10.1.100.0/24 range on Site 1.

By using 'IP pool mapping' an Appgate SDP user can now presents a unique IP address on each Site that they connect to, so when an (internal) user wants to connect to a remote Appgate SDP user (blue line), there is no issue of them having duplicate IP addresses on the WAN.

Translated IP pools must be the same size as the default IP pool, but allocated IP pools can be much smaller which is useful when sharing an existing IP range or when only a few users appear on certain Sites.

In this scenario the fully distributed nature of the Appgate SDP system is used. Users connect to multiple Sites according to their Entitlements just as if SDP was in full use. These Sites will handle all the traffic from users and the back haul connections to the protected hosts will be done locally by each Gateway.

Where some locations remain to be provisioned with SDP Gateways the a 'Override Site' can still be used to override the Site defined in the Entitlement and route the protected hosts' traffic over the WAN to the correct location.

Deployment scenario D (cont.)

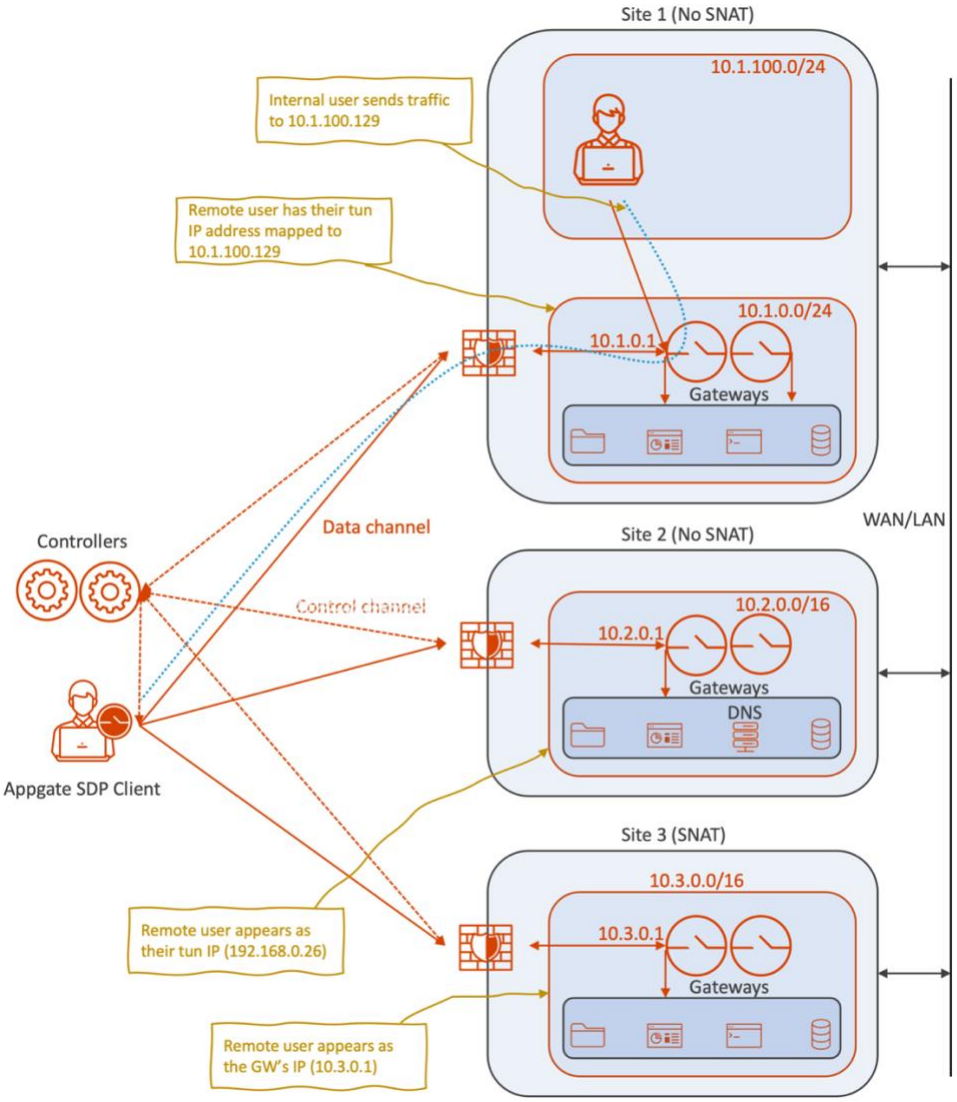

Figure 10

Benefits ✓ Encrypted connections ✓ Device On-boarding ✓ External services are cloaked ✓ Allows Entitlements to be set per Site (for full SDP working)
Issues ✓ Tun IP appears in multiple parts of the network ✓ Can be some routing issues to resolve
Mitigations ✓ Use Mapped IPs per Site

ABOUT APPGATE

Appgate is the secure access company. We empower how people work and connect by providing solutions purpose-built on Zero Trust security principles. This people-defined security approach enables fast, simple, and secure connections from any device and location to workloads across any IT infrastructure in cloud, on-premises and hybrid environments. Appgate helps organizations and government agencies worldwide start where they are, accelerate their Zero Trust journey and plan for their future. Learn more at www.appgate.com

RESOURCES

Further Appgate SDP product information is available here:

Admin Guide: https://sdphelp.appgate.com/adminguide

Client User Guide: https://sdphelp.appgate.com/userguide

More about Appgate SDP: www.appgate.com/ztna

More about ZTP: https://www.appgate.com/zero-trust-network-access/zero-trust-platform

We hope you find Appgate SDP to be a valuable solution to your security challenges.