appgate

Appgate SDP

– REFERENCE ARCHITECTURES

Type: Technical guide

Date: April 2020

Applies to: Appgate SDP v4.3 or newer

TABLE OF CONTENTS

1. INTRODUCTION

Security and compliance is now a mainstream requirement within the Enterprise and has both visibility and support at the highest levels. Previously, security was typically added after the fact and was often seen as a barrier to business & innovation. With the new Software-Defined Perimeter [SDP] approach, the conversation has shifted to how fast (and seamlessly) SDP can be deployed at scale into complex enterprise environments (to achieve Zero Trust).

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

Appgate SDP is Appgate's implementation of the SDP architecture. It implements the core architecture principles and has filled in the gaps with unique capabilities. The most unique aspects being 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 & cloud resources, including transient workloads. Appgate SDP 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 all elements of the enterprise hybrid cloud infrastructure, controlling access to authenticated users 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 point products. Appgate SDP also integrates with identity, multi-factor authentication and SIEM solutions, allowing enterprises to take advantage of existing security infrastructure. This enforces strong authentication and enables organizations to better integrate security requirements into their identity management life cycles.

- **User-Centric Network Security:** Appgate SDP provides application and service-specific authentication and authorization which controls network access both inside and beyond the perimeter. Appgate SDP 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 which closely couples the security controls with the hosts/apps themselves. This ensures traffic is encrypted over any networks 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. Appgate SDP 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 of the possible evidence necessary to meet most audit requirements.

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

If you're unfamiliar with SDP architecture that Appgate SDP implements, see https://www.appgate.com/software-defined-perimeter.

2. 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. Its aim was to solve the problem of stopping network attacks on application infrastructure, while ensuring user productivity and improving security operations efficiency. The SDP Workgroup 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 components 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 was not designed as another DMZ add-on to an existing set of security controls such Proxies or VPNs.

Rather, this new model provides an alternative to these outdated tools which were developed in a time before (hybrid) Cloud became all pervasive.

It is important to understand that Appgate SDP has filled in the gaps in the SDP model as defined by the Cloud Security Alliance — and extended the model with its own very distinctive flavor. However, the key principles remain:

- The first principle is that users and their devices live outside the new Software-Defined Perimeter. If you consider for a moment the data center model many businesses operate today; this forces users to operate outside of the building/network. So, this approach is not all that novel. Also, modern day working practices have moved on from the idea that everyone commutes to and works in the office.

- The next principle of a Software-Defined Perimeter is built on an "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, a Software-Defined Perimeter 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 Client redirects the request to the closest Gateway via a secure tunnel. This in turn applies additional policies in real time — for example, to control access based on the user's network location. This premise means that Clients can make multiple connections to multiple gateways at the same time to meet the user's specific needs.

With users staying outside the network — their devices are not being hand-carried back and forth across the network boundary, which effectively bypasses what is still many companies' main defensive line.

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 the SDP very suitable for hybrid environments — allowing consistent access policies to be applied to legacy network, data center and Cloud environments simultaneously. New Sites are very independent of one another and easily deployed with no long lead-times; they simply require Internet access.

3. TODAY'S NETWORK MODEL

While SDP is new, today's network model is already over twenty years old. It evolved when Microsoft-powered PCs and ethernet based LANs dominated the business world. At that time, the Internet was becoming well established, so TCP/IP seemed to be the logical way to connect up sites and later data centers. But TCP/IP was never designed to address all of these different use cases to which it has been applied. Maybe the biggest issue being its inherent lack of any up-front security when establishing new connections.

Today's network struggles as IP-based devices are no longer tied to users' desks and connected via CAT5 cabling to a managed infrastructure with a defined boundary. Securing today's networks has become increasingly difficult, especially given the increasingly complex and disparate range of (network) security systems deployed both within and around the edge of the network to try to compensate for TCP/IP's shortcomings.

Within these networks high availability has been provisioned by having multiple redundant systems with yet more network systems deployed to manage the traffic. These systems needing to be configured to work for users on the inside as well as ones coming in via VPN from the outside.

The need for redundancy and user's different geographies often require these network 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 exploit once a bad actor has established a foothold.

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

AppGate SDP Reference Architecture – Today's model

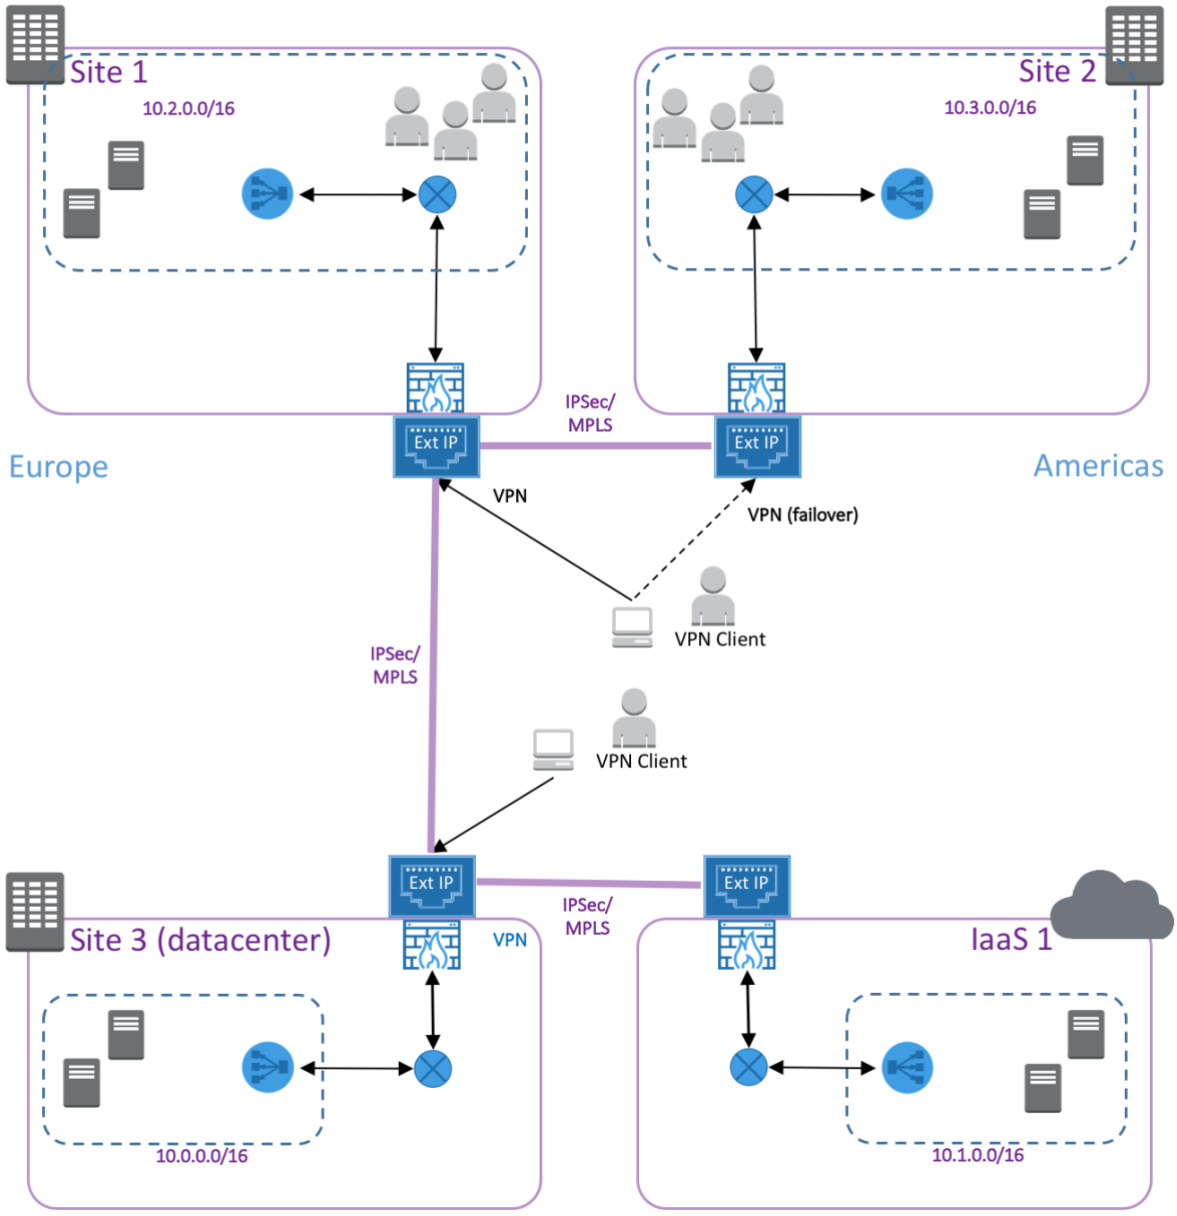

Before examining the Appgate SDP model, let's remind ourselves of the four main changes driving many organizations to realize that the time has at last arrived when we should be retiring these legacy networks and replacing them with the Software-Defined Perimeter:

Mobile working & personal devices: People's lives are increasingly always-on and there is an expectation that they can work when and where they like, often on a device of their choosing. People no longer have devices tied to the network; indeed, their devices may cross the network boundary many times per day or live on the outside permanently.

The Cloud: Cloud adoption is now well established with almost all businesses using it to some degree. Often it is tacked on as shown above but this model is not sustainable as many businesses will be hybrid Cloud users for a long time to come and cannot continue to backhaul all the connections to the Cloud.

Complexity beyond comprehension: Networking is still seen as an 'inside' thing which should provide some protection from the 'outside'. In reality IP-networks have changed little over the last 20 years except for the proliferation of point solutions that have filled most businesses DMZs in a vain attempt to secure this increasingly complex nightmare.

External Threats: The speed with which new exploits are utilized is frightening. The attack surface of a traditional (multi-site) network is vast. These together are making the job of defending these networks next to impossible unless you can cloak all access and make the network effectively disappear.

4. HOW DOES APPGATE SDP WORK?

Appgate SDP 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:

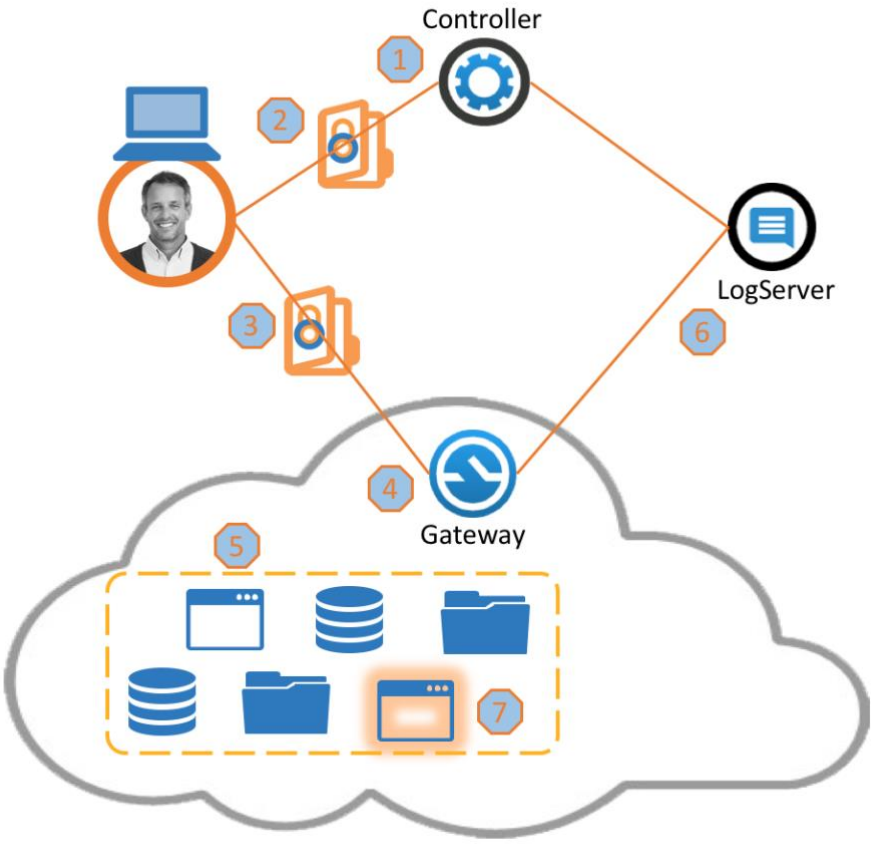

Step-by-Step

1. User authentication: The user authenticates to the Appgate SDP Controller, which is optionally connected to an enterprise's IAM system (AD, LDAP, Radius or SAML)
2. Controller applies access policies: The Controller applies policies assigned to the user based on user attributes, roles, and context, and then issues signed Entitlement tokens listing the resources the user is allowed access to.
3. User accesses resources: The authenticated user can now access 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 have been successfully met, it will open a connection to the protected resource specified by 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. There is an integral LogServer which can be used for this or logs can be sent to the enterprise SIEM for additional action. The LogServer 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* users thus enhancing availability for *authorized* users.

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

Controllers

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 on bi-directional TCP port 444 (mutual TLS connectivity) for control channel messages and bi-directional TCP port 5432 for encrypted mutual database synchronization. The database is mainly used by the Controllers to store the configuration, Policies and Entitlements 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, Radius and SAML-based identity systems. Multiple IdPs can be specified in the Controller which will be tried sequentially. If no IdPs are available then the next Controller is tried (if the Client is aware of any others).

Clients

The Client has to 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 will check 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 will ask 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.

Gateways

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 term denoting a logical grouping of protected hosts which might span more than one physical location.)

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 and receive mutual TLS traffic on port 444 with Controllers and LogServer. There is no communication between Gateways whatsoever.

Gateways can be deployed in a number of different ways. Before we go on to look at the different network architectures that can be used it is worth just focusing on the Gateways themselves to understand their immediate deployment options.

Gateway deployment options

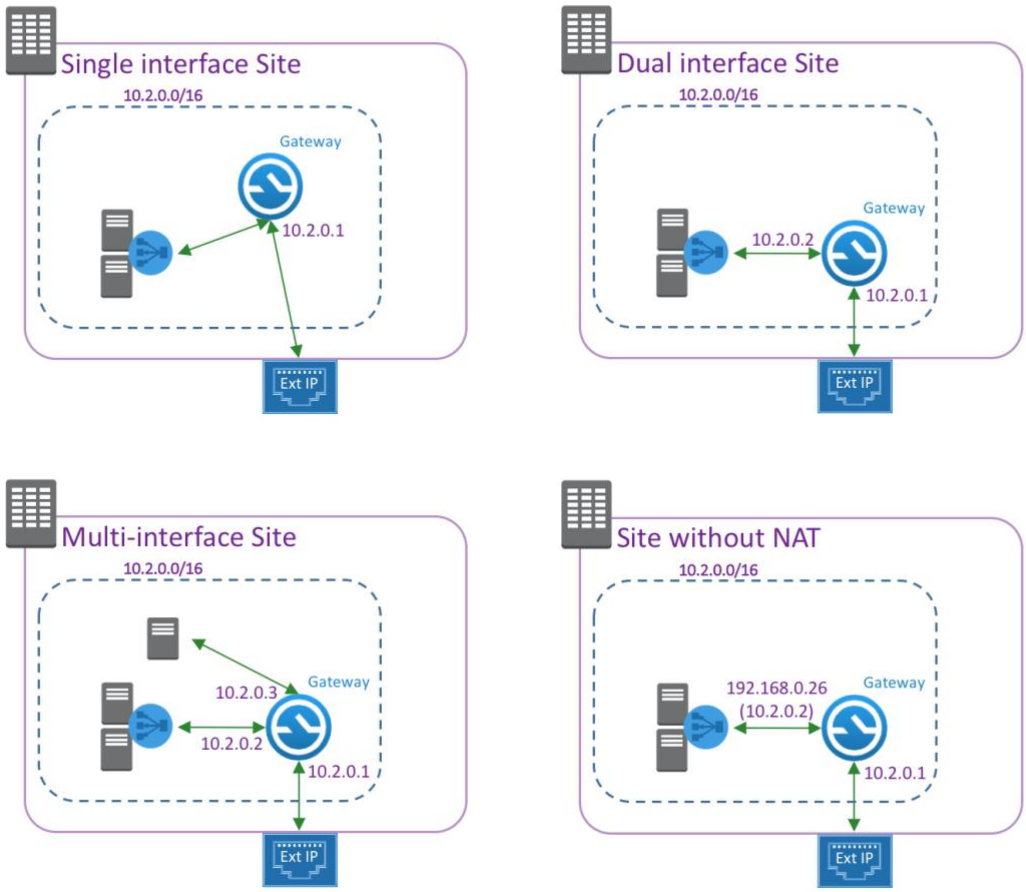

Single interface Site

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 also allow traffic from the NIC to the protected resources. The performance of the NIC is restricted as it has to handle traffic in both directions.

Dual interface Site

Also a quite typical deployment scenario where the Gateway sits inline. It may still reside in a DMZ or Security zone but can often be positioned to sit alongside a traditional firewall as opposed to in series with it.

Multi-interface Site

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

Site without NAT

Without NAT, the user's tunnel IP address is presented on the network instead of the Gateway's IP address. This may require some routing changes on the network. This mode is required when stateful failover between Gateways is a requirement.

Textbook SDP Model

In a textbook Appgate SDP deployment the company network might be very different from what is typically in place today.

The Appgate SDP collective is fully distributed across all Sites where there are resources that require protection. Controllers are deployed at different locations for HA/resilience. The user traffic to Controllers is minimal and infrequent so location/geography is only of secondary importance.

All peer-to-peer traffic can be sent across the Internet because it uses mutually validated TLS. Access to peers should be restricted using firewall rules.

User access is via SPA protected ports, so the SDP system is cloaked for all unauthorized users coming from the Internet.

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

AppGate SDP Reference Architecture – SDP model

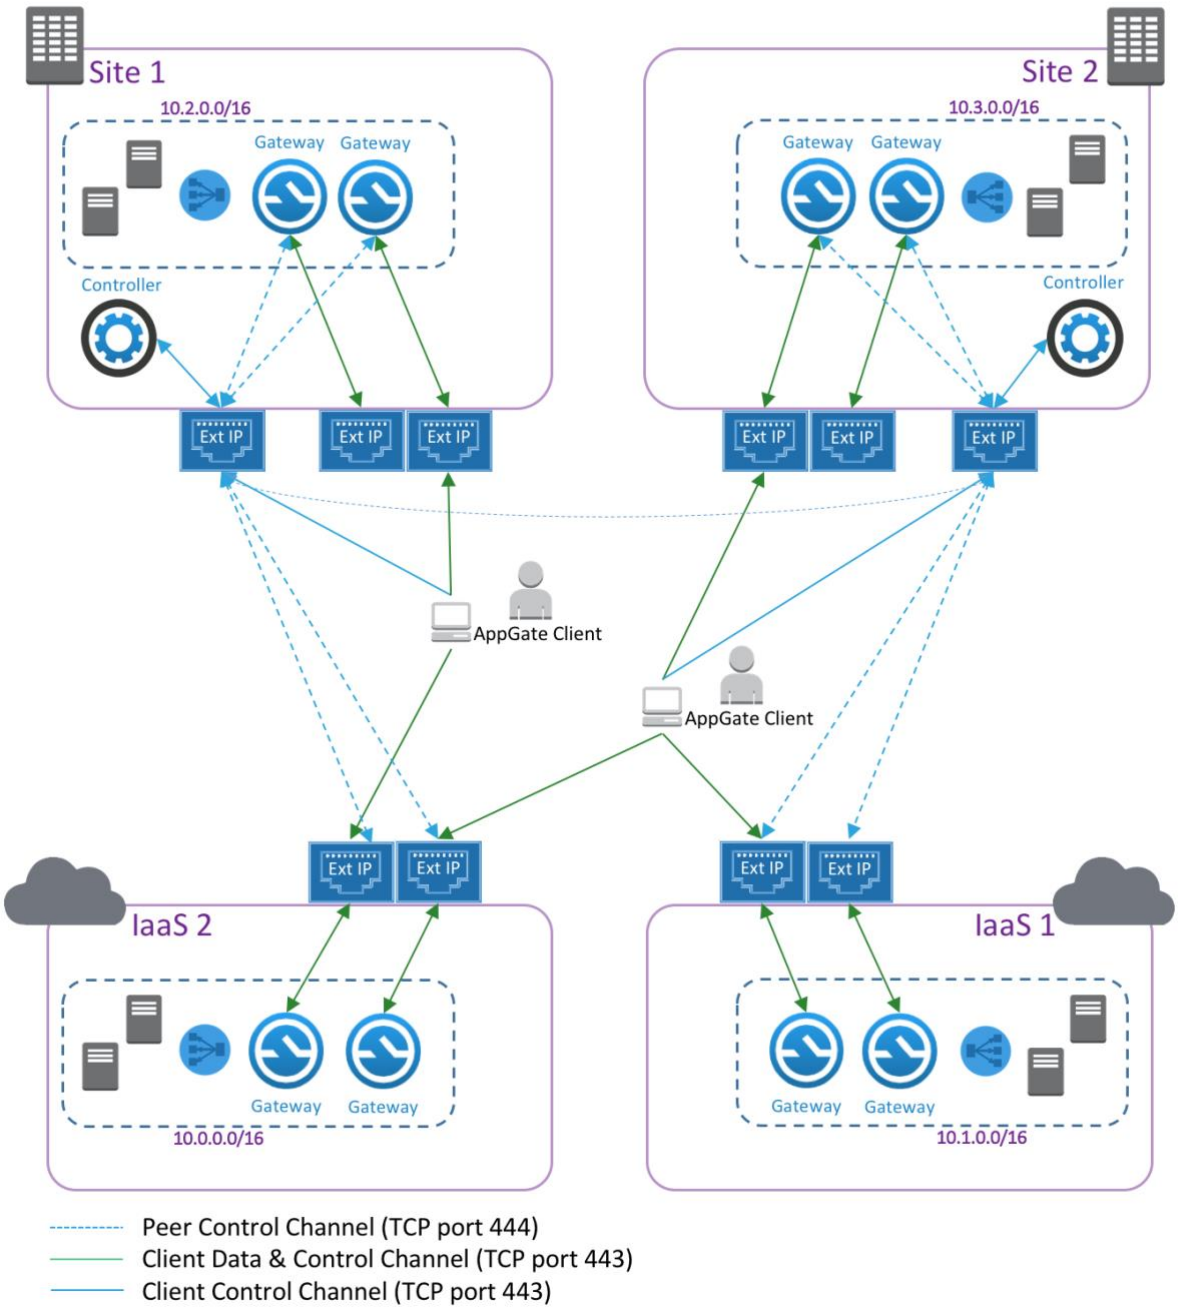

In this model there are 2 Controllers at different locations but with the same DNS name (and 2 DNS A records) to provide high availability.

The Controllers talk with all the gateways, the *local* Gateways using their external DNS names. This avoids any potential issues with users moving from the inside to the outside and allows for easier relocation of appliances during any phased SDP deployment.

Clients will connect to one Gateway on every Site where they have been 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 are a minimum of 2 Gateways per site to provide HA (and load balancing). In the event one fails then the Clients automatically try any others that are listed for the Site. These Gateways could provide access to the protected applications directly; however for end-to-end HA, application access should be via some form of load balancer using virtual IPs.

Management network

Not shown in above is any reference to management networks which are common place in today's data centers. However all Appgate SDP appliances come with the ability to support multiple NICs. It is therefore very easy to link all the appliances to any management network while maintaining full separation from the user traffic. The (delegated) administration policies can also be tailored so that only Clients connecting from the IP range of the management network would be granted admin UI access.

Benefits	Issues
• 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	• Backwards compatibility with today's topology **Mitigation** • Use the intermediate topologies suggested below

5. JOURNEY FROM TODAY'S NETWORK TO SDP

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 at 3 different models each using Appgate SDP as its basis, 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. The related access policies are not well documented/understood and are enforced at different points across multiple different systems. Finally, there are often sizable user bases involved making change hard to plan and implement. So how do you get from today's networks to the SDP model of tomorrow?

Many of today's network components already have the ability to implement access rules based on the user's identity but they are not widely implemented because the cost is high (there is no centralized way to manage policy) and the benefit low (limited by flawed topology of today's networks).

One of the key benefits of Appgate SDP model is the 'segment of one' based on user context. By using attributes, such as group memberships, it is quite easy to implement some simple departmental Policies inside Appgate SDP, which might for instance allow a group of existing users access to a whole subnet.

Because these policy objects behave quite independently, each can be fine-tuned when the time is right until the ultimate 'segment of one' is achieved for all user groups.

SDP with Users on the Inside

The first key milestone on the journey from today's networks to SDP is to remove users 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 is quite limited as it does not extend 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) once access has been 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. The similarity is that both manage the process that allows users to be granted access to protected hosts. In using SDP to solve a problem where a traditional NAC solution might be considered, the fundamental prerequisite is that users MUST first be segmented away from the protected hosts. This may 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 then 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 is done effectively—any lingering NAC requirement should vanish.

The user network would allow access to the Appgate SDP system (but not the protected hosts) and the Client would establish a secure connection to the Gateway. Policy would allow the users access to the protected hosts – possibly using a subnet wide access rule if the specific access rules were not known at the time. Best practice suggests setting the internal DNS to resolve to the external Gateway IP address to avoid any issues when users move from the user network to VPN (outside); and after all this is only a transitional state on the journey to full SDP.

The local Gateway (to the user) could still resolve hosts on other sites according to the existing networks DNS servers and the traffic forwarded over the WAN to any such hosts.

Having said that, either of the next two scenarios could also be used, in which case it might be better to deploy Gateways to all Sites and allow these users to connect to multiple Gateways if hosts on other sites were included in any access Policy.

AppGate SDP Reference Architecture – SDP applied to today's model – user on the inside

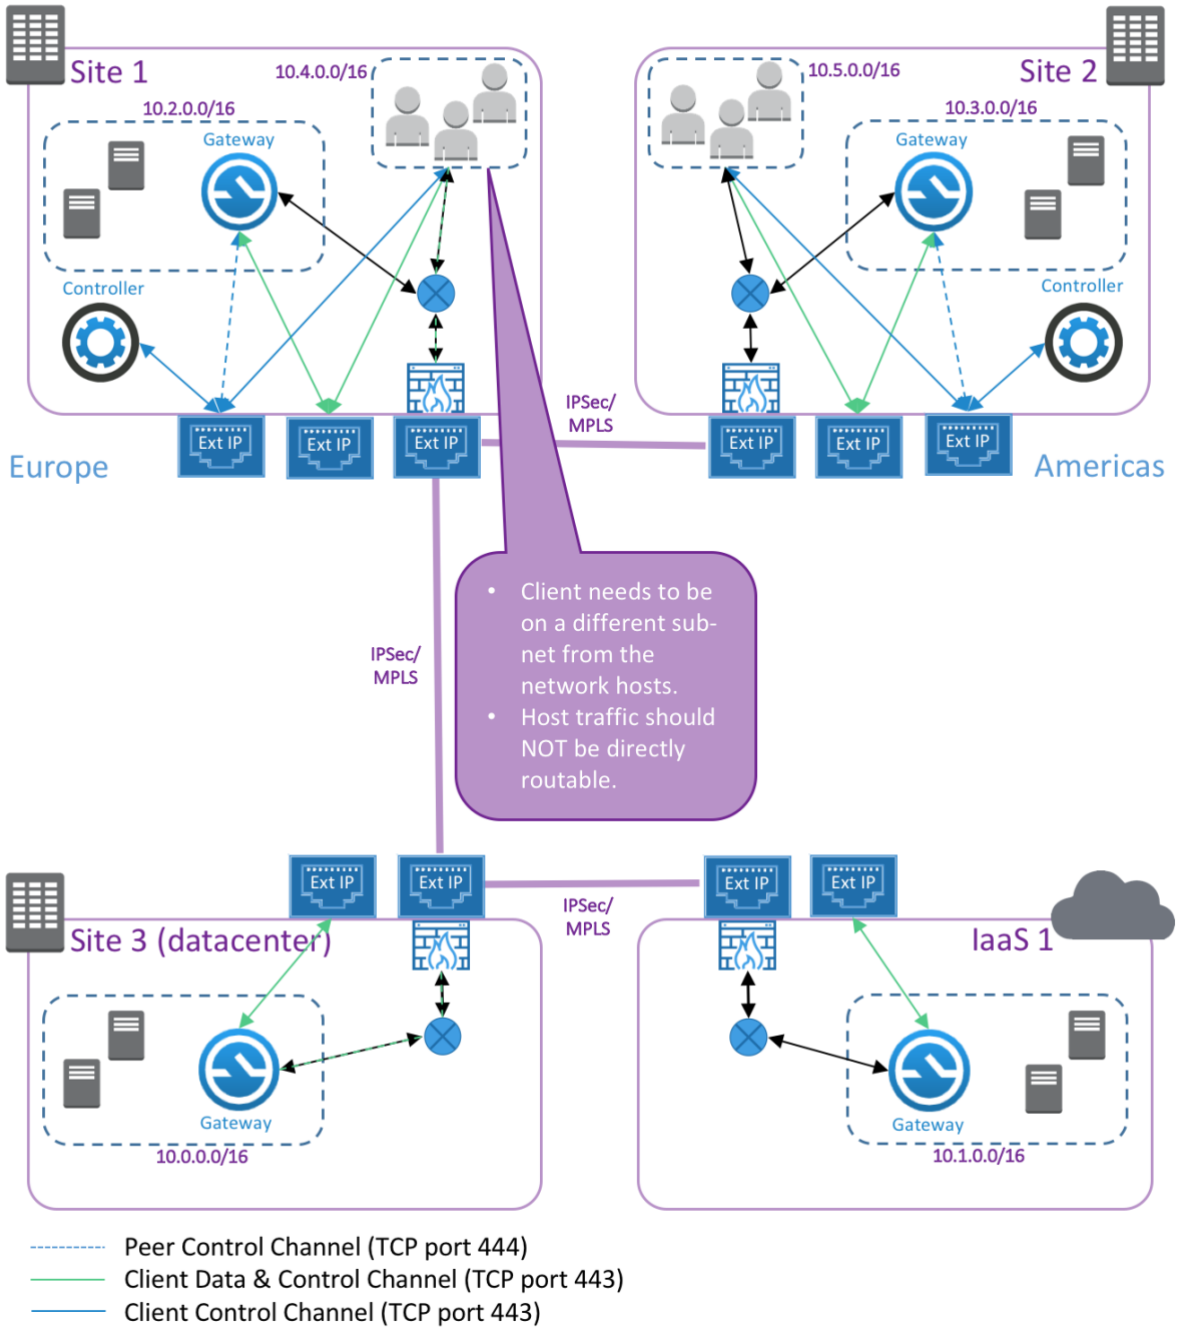

In this scenario the external capabilities of the Appgate SDP system need not necessarily be used. So all the traffic (TLS) could be routed over internal networks and the WAN. The Gateway prevents direct access to the protected hosts and any network routing is updated to reflect this. The Clients would build routing tables based on the configured Entitlements.

These routes would point to the Appgate SDP tunnels for access to all the protected applications. So when a user tries to access a protected host the traffic is directed over the LAN to the local Gateway. If some of these protected hosts were in the data center then the Gateway there would handle this, the tunnel being established over the existing WAN.

Benefits	Issues
• Encrypted connections on user network • Device On-boarding • Protected hosts are not visible o the LAN/WAN • User based rules (no NAC required) • Users live outside the protected network	1. Need to keep users away from hosts 2. When used outside the enterprise network, internal/external DNS can be challenging **Mitigation** 1. Need to do some network segmentation 2. Follow the DNS guide lines in the manual

SDP with Single Tunnel (VPN Alternative)

In the single tunnel scenario the Appgate SDP system deployment is conceived such that the system can be extended/grown over time towards the SDP ideal model. To begin with this model can be thought of as a simple VPN replacement. This scenario might be required because there is no Internet access provision to some Sites (within the WAN) at the current time.

The Controllers can be geographically distributed using the ideal SDP model. At least one Gateway is deployed; usually on the site where the existing VPN infrastructure is located. If two points of presence are used for VPN access; for instance, one in Europe and one in the Americas, then 2 Gateways can be deployed. Geo-location claims can be used to direct the users to the appropriate Gateway. The network continues to provide all the backhaul routes across the WAN.

This might be a good point to explain how the system would failover European users to the Americas if the European site went down. Controllers are not an issue, as was explained earlier DNS round-robin would always allow users to find the other Controller. But how do you fail one Site over to the other? In this example the simplest solution would be to have Entitlements configured for apps on both Sites. Normally DNS for crm.mycompany would return crm.europe.mycompany which would be used in the European Sites Entitlements. However when the European site was down then instead DNS for crm.mycompany would return crm.americas.mycompany which would be used in the Americas Sites Entitlements.

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 and with users connecting to only one Gateway, access to the protected hosts not on the Gateway's home Site would still be over the WAN. So effectively we have only one (Appgate) Site!

To define all the Entitlements on the one Site at this point would make migration harder when the time comes spliting the Site into smaller parts. If 2 points of presence were used (2 Sites) then the problem is even worse as all the Entitlements need to be defined twice!

To support this interim stage on the journey, it is possible to specify a ‘Deployment Site’ in the Policy form; this will effectively override the normal SDP situation where the Entitlements relate to their specified Site. All Entitlements will now be forced to the ‘Deployment Site’, but this means the system can be configured for multiple Sites even though there is only one operational.

As Sites are populated with Gateways and moved away from being part of the company WAN to being stand-alone SDP sites; migration is a simple as deselecting the ‘Deployment Site’ in each Policy!

AppGate SDP Reference Architecture – SDP applied to today's model – single tunnel

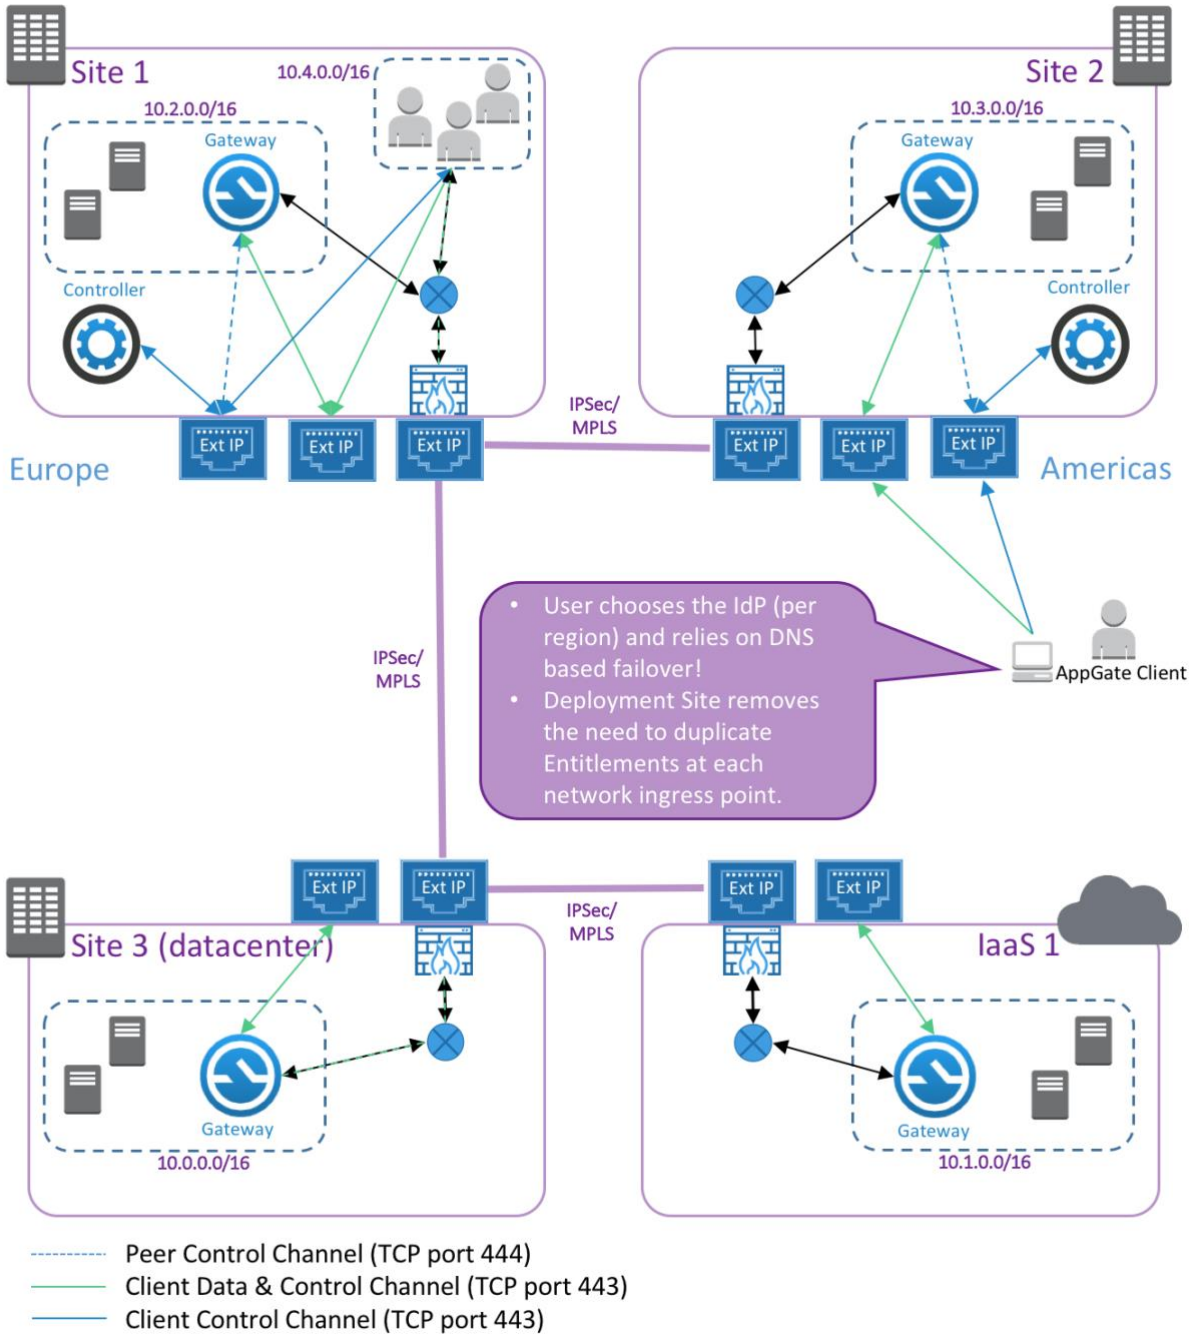

In this scenario the external capabilities of the Appgate SDP system are brought into play but the fully distributed nature of SDP is kept for another day.

Two points of presence are used one for Europe and one for the Americas. Gateways can still be deployed to other Sites but only the two Deployment Sites needs to be fully implemented. These Sites will handle all the traffic from users then the back-haul connections to the protected hosts over the existing WAN just as a VPN solution would do today.

Users only connect to one site but if one Policy was configured per (future) Site, then as soon as another Site is ready to become stand-alone, the ‘Deployment Site’ option is removed and the Clients would now build an additional tunnel to that Site and route the traffic there directly.

Benefits	Issues
• Encrypted connections • Device On-boarding • External services are ‘cloaked’ • Geographic based access for users	1. Must have duplicate Entitlements on all Sites 2. No real-time failover between regions Mitigations 1. Use Deployment Site option 2. Give users the choice of regional IdPs to use

SDP with Multiple Tunnels

In the multiple entry point scenario the Appgate SDP system deployment is conceived such that the system is part way towards the SDP ideal model. To begin with, this model is likely to be used by the large enterprise who may have several points of presence and users might use more than one as they travel between regions. It might also reflect a stopping point part way along the journey to the ideal SDP model where some sites are deployed (but not all).

The Controllers can be distributed using the ideal SDP 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 used to implement the WAN.

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. Where some sites remain to be provisioned fully for SDP a ‘Deployment Site’ can still be used to override the Site defined in the Entitlement.

If SNAT is not being used on the Gateways then the users tunnel IP address will be advertised on all the used Sites even ones still linked by the WAN. This is OK for traffic originating from the Client as it is stateful. However for traffic originating from a Site the user’s IP address may appear in several places at the same time causing routing issues.

To support this interim stage on the journey, in the Site form it is possible to specify ‘IP pool mapping’. This will effectively override the normal IP tunnel address that the Client uses when connecting to a given Site. So the default IP pool 10.100.0.0/24 addresses might be mapped to their equivalent in the 10.100.2.0/24 range on site 2 and to the 10.100.3.0/24 range on site 3. This solves the routing issue that can exist while Sites remain linked by the WAN.

In this scenario:

- The Appgate SDP system may be deployed distributed
- The protected resources sit behind gateways wherever required
- Users co-habit inside the network.

AppGate SDP Reference Architecture – SDP applied to today's model – multiple tunnels

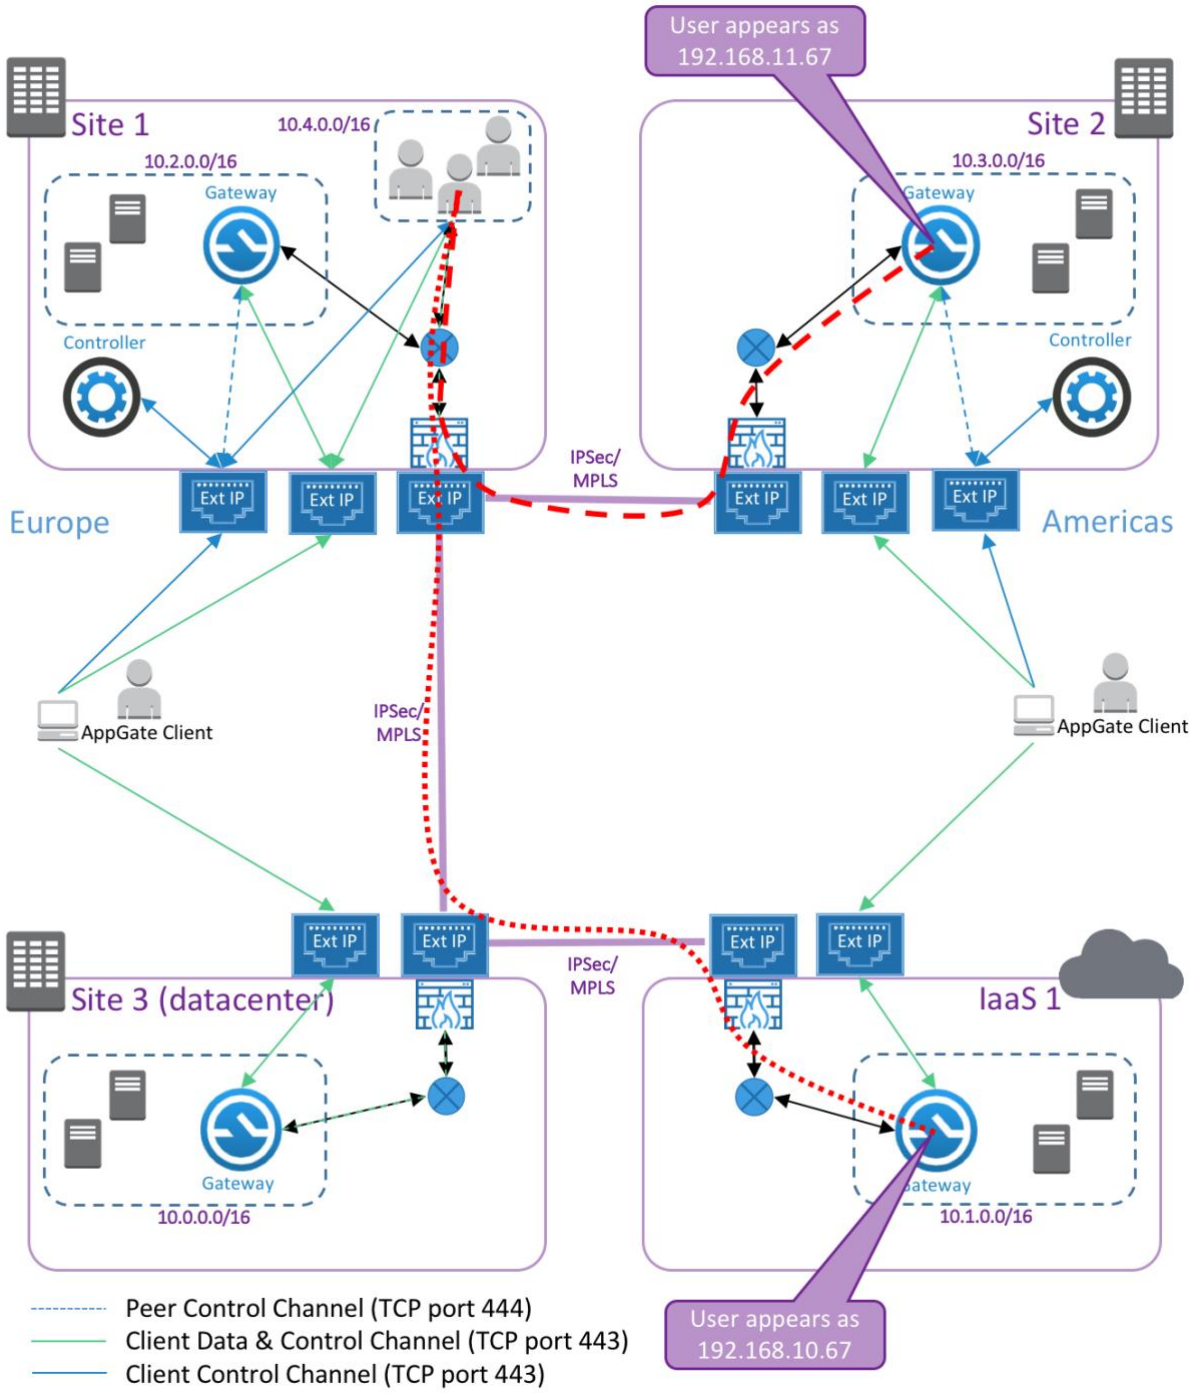

In this scenario the fully distributed nature of the Appgate SDP system is used. Users connect to multiple Sites according to their Entitlements just as if SDP was in full use. These Sites will handle all the traffic from users and the back haul connections to the protected hosts will be done locally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

With 'IP pool mapping' an Appgate SDP user now presents a different IP address on the Sites they connect to, so when an (internal) user wants to connect to the Appgate SDP user then the issue of them appearing twice on the network with the same IP address is resolved.

If the target IP address is 192.168.10.67, then only the dots route will be used as this is now unambiguous.

Mapped IP pools must be the same size as the default IP pool, so with many sites this can use up a lot in internal IP addresses.

Benefits	Issues
• Encrypted connections • Device On-boarding • External services are 'cloaked'	1. tunIP appears in multiple parts of the network 2. Can be some routing issues to resolve Mitigations 1. Use Mapped IPs per site 2. Ensure Entitlements are set per Site

6. RESOURCES

Further Appgate SDP product documentation is available here:

- Admin Guide: https://sdphelp.appgate.com/adminguide
- Client User Guide: https://sdphelp.appgate.com/userguide
- Web site resources: https://www.appgate.com/software-defined-perimeter

Thank you, and we hope you find Appgate SDP to be a valuable solution to your security challenges.