# Software Defined Perimeter (SDP) and Zero Trust

cloud security alliance®

The permanent and official location for Software Defined Perimeter Working Group is
https://cloudsecurityalliance.org/software-defined-perimeter/

© 2020 Cloud Security Alliance – All Rights Reserved. You may download, store, display on your computer, view, print, and link to the Cloud Security Alliance at https://cloudsecurityalliance.org subject to the following: (a) the draft may be used solely for your personal, informational, non-commercial use; (b) the draft may not be modified or altered in any way; (c) the draft may not be redistributed; and (d) the trademark, copyright or other notices may not be removed. You may quote portions of the draft as permitted by the Fair Use provisions of the United States Copyright Act, provided that you attribute the portions to the Cloud Security Alliance.

# Acknowledgments


## Lead Authors:

Juanita Koilpillai
Nya Alison Murray

## Contributors:

Michael Roza
Matt Conran
Junaid Islam
Aditya Bhelke
Eitan Bremier
Tino Hirschmann
Steve Swift
Sam Heuchert
John Markh
Roupe Sahans
Oscar Monge Espana
Gerardo Di Giacomo
Vladimir Klasnya
J. Lam
Clara Andress
Dan Mountstephan
Manoj Sharma

## CSA Analysts:

Shamun Mahmud

## CSA Global Staff:

AnnMarie Ulskey (Design)


© Copyright 2020, Cloud Security Alliance. All rights reserved.

# Table of Contents


Acknowledgments ..... 3
Introduction ..... 5
- Goals ..... 7
- Audience ..... 8

Zero Trust Networking (ZTN) and SDP ..... 9
- Why Zero Trust ..... 9
- What Zero Trust Addresses ..... 10
- Implementing a Zero Trust Strategy ..... 12
- Benefits of a SDP Zero Trust Solution ..... 15
  - Security Benefits ..... 15
  - Business Benefits ..... 16
- SDP Zero Trust Strategic Approach and Proof of Concept ..... 17
- Technology Components and Infrastructure ..... 20
- Technology Risks and Issues ..... 21
- Assumptions ..... 21
- Technology Analysis ..... 21
- Required Resources ..... 22
- Key Industry Developments ..... 22
- Delivery Activities ..... 23
- Situation Analysis ..... 23
- Timeframes and Stakeholder Engagement ..... 23

References ..... 24


© Copyright 2020, Cloud Security Alliance. All rights reserved.

# Introduction

Software Defined Perimeter (SDP) is a network security architecture that is implemented to provide security at Layers 1-7 of the OSI network stack. An SDP implementation hides assets and uses a single packet to establish trust via a separate control and data plane prior to allowing connections to hidden assets. A Zero Trust implementation using Software Defined Perimeter (SDP) enables organizations to defend new variations of old attack methods that are constantly surfacing in existing network and infrastructure perimeter-centric networking models. Implementing SDP improves the security posture of businesses that face the challenge of continuously adapting to expanding attack surfaces that are increasingly more complex.

Originally, Zero Trust Network (ZTN) concepts were developed by the US Department of Defense (DoD) in the early 2000s while defining Global Information Grid (GIG) Network Operations (NetOps) Black Core routing and addressing architecture, part of the DoD's Netcentric Service Strategy. Over time, this concept evolved within the DoD intelligence and security communities into the current ZTN/SDP framework and test lab[1]. Around the same time, Forrester, a market research company that provides advice on technology began promoting ZTN as a worthwhile consideration for enterprise security teams. Today, Zero Trust has grown widely in adoption, as well as scope.

In the report entitled "Zero-Trust-eXtended-ZTX-Ecosystem," Forrester analysts observe that the changing nature of the network perimeter means that the historical context of Zero Trust architecture is transforming rapidly from "segmenting and securing the network across locations and hosting models." Forrester asserts that the current model, which supports the need to challenge and eliminate the inherent trust assumptions in current security strategies, suggests that a variety of new adaptive software-based approaches should also be considered. However, it does not identify a new direction for the "extended ecosystem framework."[2]

Essentially, Zero Trust is a network security concept centered on the belief that organizations should not automatically trust anything inside or outside traditional perimeters and aims to defend enterprise assets. Implementing Zero Trust requires the verification of anything and everything that tries to connect to assets before granting access and the continued evaluation of sessions during the entire duration of the connection. This is illustrated in Figure 1 where the National Institute of Standards and Technology (NIST) describes using 'trust boundaries.'

[1] https://www.secureworldexpo.com/industry-news/pentagon-zero-trust-security-framework
[2] https://www.em360tech.com/wp-content/uploads/2019/04/The-Forrester-Wave%E2%84%A2_-Zero-Trust-eXtended-ZTX-Ecosystem-Providers-Q4-2018-1-1.pdf

© Copyright 2020, Cloud Security Alliance. All rights reserved.

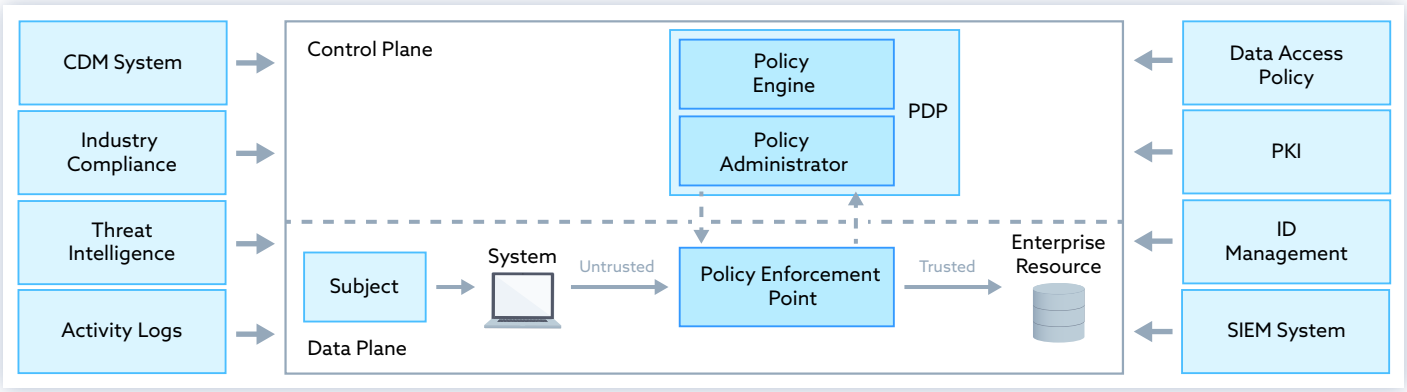


*Figure 1: Source: NIST, 800-207, Zero Trust Architecture 2nd Draft https://csrc.nist.gov/publications/detail/sp/800-207/draft*

So what is Zero Trust? According to Forrester, there are three main concepts of Zero Trust:

- Introducing the concept of trust to the network, so that it becomes natural to ensure that all resources are securely accessed no matter who creates the traffic or from where it originates, regardless of location or hosting model, cloud, on-premises or collocated resources.
- Adopting a least privilege strategy (LPS) that enforces access control to eliminate the human temptation to access restricted resources.
- Continuously logging and analyzing user traffic inspection for signs of suspicious activity.

What is SDP? Software Defined Perimeter (SDP) is the most advanced implementation of a Zero Trust strategy. Cloud Security Alliance has adopted and is advocating the following constructs applied to network connectivity:

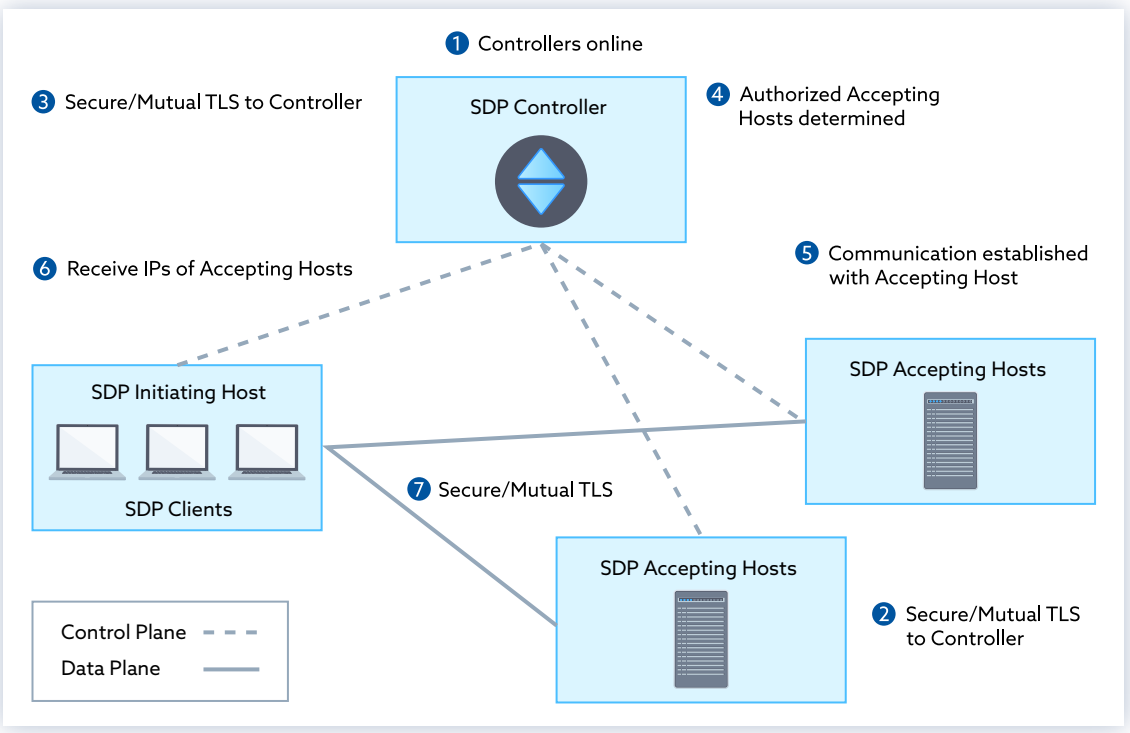


*Figure 2: SDP Architecture (previously published by CSA in SDP Specification 1.0)*

© Copyright 2020, Cloud Security Alliance. All rights reserved.

- Separating the control plane where trust is established from the data plane where actual data is transferred.
- Hiding the infrastructure (e.g. blackening the servers) using a dynamic deny-all firewall (not deny-all, allow exceptions) - the point where all unauthorized packets are dropped for logging and analyzing traffic.
- Using single packet authorization to authenticate and authorize users and validate devices for access to protected services - least privilege is implicit in this protocol.

Since SDP is agnostic of the underlying IP-based infrastructure and hones in on securing all connections using said infrastructure, it is the best architecture for adopting a Zero Trust strategy, as it can be applied at the OSI/TCP/IP network layer before the transport layer protocols, and prior to the application of the session layer. This is important as the transport layer, that provides host-to-host communication services for applications, and the session layer, the mechanism for opening, closing and managing a session between end-user application processes, both have known and undiscovered weaknesses, for example Transport Layer Security (TLS) vulnerabilities and TCP/IP SYN-ACK attacks on session establishment.

The following table relates the ISO Open Systems Interconnection (OSI) model to the Internet Engineering Task Force (IETF) TCP/IP protocol.

| # | OSI Layer | TCP/IP Layer | Protocol Data Unit | Description |
|---|---|---|---|---|
| 7 | Application | Application | Data | Network process to application |
| 6 | Presentation | | Data | Data representation and encryption |
| 5 | Session | | Data | Interhost communication |
| 4 | Transport | Transport | Segments | End-to-end connections and reliability |
| 3 | Network | Internet | Packets | Path determination addressing |
| 2 | Datalinks | Network Access | Frames | Physical addressing |
| 1 | Physical | | Bits | Media, signal and binary transmission |

*Figure 3: Source: https://www.iso.org/ics/35.100/x/ and https://tools.ietf.org/html/rfc1180*

# Goals

This paper will show how SDP can be used to implement ZTNs and why SDP is applied to network connectivity and is the most advanced ZTN implementation.

© Copyright 2020, Cloud Security Alliance. All rights reserved.

# Audience

Security professionals, CIOs, CISOs and other corporate executives who are embracing Zero Trust as the breakthrough for effectively protecting against large-scale breaches are the intended audience for this paper.

© Copyright 2020, Cloud Security Alliance. All rights reserved.

# Zero Trust Networking (ZTN) and SDP

The security industry acknowledges that existing defense mechanisms are only partially successful. The execution of SDP can be applied before TCP/IP and TLS, which reduces the likelihood of these and other vulnerable protocols being used as attack vectors by threat actors. Software Defined Perimeter implementations compliant with the CSA SDP version 1 specification create zero trust implementations that prevent common methods of attack such as DDoS, credential theft, and the notorious top ten threats published by the Open Web Application Security Project (OWASP). SDP renders assets invisible and prevents access until the associated identity is successfully authenticated and authorized for access to these assets for a proven zero trust implementation.

In practical terms, "Zero Trust" is the philosophy behind the SDP architecture. SDP's basic tenets are ABCD: "Assume nothing, Believe nobody, Check everything, Defeat threats." While SDP ZT is meant to be applied at the Network Layer 3 of the International Standards Organization (ISO) OSI mode, in view of common architecture patterns such as applications accessing hybrid cloud services, care must be taken to apply ZTN as close to a domain perimeter as possible, to ensure optimal performance and prevent unnecessary service latency

## Why Zero Trust

Today's network security implementations can be compared to the analogy of building walls and doors, allowing for criminals to attempt to pick the locks of the doors. Organizations today rely on their security 'door locks' and then heavily monitor the locks to ensure that criminals don't break in. It is much better to ring-fence digital assets and then rely on vetting the threat to keep out unauthorized users. We may want to see who is knocking but definitely must prevent malicious acts by denying the opportunity to pick the locks. This is the essence of why there is a pressing need for effective Zero Trust deployments. Furthermore, it is well known that threat actors' primary goal is to penetrate networks and then gain lateral movement in order to access systems with increased privilege credentials. Zero Trust can prevent unauthorized users from hiding their activities, limiting access to authorized users.

The following issues require a rapid change in the way network security is implemented.]

a) The Changing Perimeter

The past paradigm of a fixed network perimeter, with trusted internal network segments protected by network appliances such as load balancers and firewalls, has been superseded by virtualized networks and the realization that the network protocols of the past are not secure-by-design. In fact, many current network protocols, such as IPSec and SSL VPN have known vulnerabilities[3]. In addition, the plethora of mobile and IoT devices challenges the essence of the traditional fixed network perimeter network.

With the introduction of cloud, the environment has changed. Add to cloud BYOD requirements, machine to machine connections, the rise in remote access and the rise in

[3] https://crypto.stanford.edu/cs155old/cs155-spring11/lectures/08-tcp-dns.pdf

© Copyright 2020, Cloud Security Alliance. All rights reserved.

phishing attacks, the legacy approach is constantly challenged. There are many internal devices and a variety of users. A common use case is on-site contractors must access network resources both on premise and in the cloud. Also trending are hybrid architectures where corporate workstations are using the cloud to facilitate co-located facilities with end-users moving off-site to customer and partner locations. Moreover, in these scenarios, domain perimeters are being redefined with site-to-site connections including interconnectivity to third parties.

b) The IP Address Challenge

Everything today relies on IP addresses for trust at layers 1-4 of the OSI stack, but this presents a problem: IP addresses lack any type of user knowledge to validate the device request integrity. There is simply no way for an IP address to have user context. IP addresses simply provide connectivity information but provide no indication of trustworthiness of the endpoint or the user. TCP is a bidirectional protocol at layer 4 of the OSI network stack, so internal trusted hosts communicating with external untrusted hosts can receive untrusted messages.

Any changes to IP addresses can mean complex configuration, allowing errors to creep into network security groups and network access control lists. Forgotten internal hosts can provide an entry point to hackers by providing default responses to past protocols such as ICMP network support. Finally, IP addresses should not be used as anchors for network locations because IP addresses change, for example with dynamic allocation, or when a user moves from one location to another.

c) The Challenge of Implementation of Integrated Controls

Visibility and transparency of network connections is problematic for network security and security tools implementation. Currently integration of controls is performed by gathering data in multiple logs forwarded to Security Information & Event Management (SIEM) or Security Orchestration Automation Response (SOAR) technologies for analysis.

A single point of trust for network connections is difficult to implement. Integrating identity management prior to allowing access through a firewall is a resource intensive task. In addition, for most development/operations/network teams, use of secure coding practices, application layer firewalls and anti DDoS protections is very much an afterthought.

Providing individual applications with the ability to control security posture is currently a huge challenge. Retrofitting security into application and container platforms requires integration of access controls, identity management, token management, firewall management, code, script, pipeline and image scanning, as well as orchestration the integrated whole. This is proving very difficult for most teams.

# What Zero Trust Addresses

The following common weaknesses are inherent in how networks are architected today, giving rise to the need for a new way of designing networks for security.

© Copyright 2020, Cloud Security Alliance. All rights reserved.

a) **Connect first and then authenticate** - In most network installations, access is allowed prior to authentication. Since there are no foolproof gatekeepers to challenge identity claims, access control mechanisms can be bypassed. Encrypted or not, authentication, authorization, and token-based access control systems may have multiple flaws.

The predominant network protocol used today for connectivity is the Transport Control Protocol (TCP). Applications operate with a Connect First, Authenticate Second model when they use this protocol for connectivity. When a client wants to communicate and have access to an application, clients first need to set up a connection. Once the connection is established, then clients authenticate. Once the client has authenticated, data can be exchanged.

In this model, clients are allowed to connect to the network first, and this allows unauthorized user ingress. Clients are then authenticated but only after connection is allowed. This means connected and unauthorized users are now in the network and can perform malicious activity. As there is no awareness of who the legitimate clients are until authentication happens, these users typically bypass the authentication methods when their identity claims are not challenged.

In essence, devices are given IP addresses to connect to the Internet, which forces organizations to do three things:

- Deny the bad actors who are attempting to connect, relying on threat intelligence to provide the identification of these parties.
- Lock down the machine so it is airtight, i.e. with vulnerability, patch and configuration management. This has proved impossible.
- Implement a network layer firewall device with no user context. These firewalls are vulnerable to internal attacks, out-of-date static configurations. (NOTE: Next Generation Firewalls (NGFs) do address user context, application context and session context into consideration but are still IP-based, with uncertain results because of application layer vulnerabilities. See SDP Architecture document for details.)

SDP Perspective: **None of these techniques are effective at preventing attacks. A Zero Trust implementation requires immunity from all layers of attack on network, hosting and application platform infrastructure.**

b) **Monitoring endpoints is compute, network and human resource intensive -** Endpoint monitoring using AI cannot yet correctly detect or prevent unauthorized access. Virtual variations on isolation of protected resources can be compromised over time by capturing identity details, understanding authorization mechanisms and spoofing authentication credentials of people, roles and applications.

Today, artificial intelligence models are currently simple behavioral models, for the most part based on multiple linear regression analysis and/or expert systems, or neural networks which trained to detect patterns. AI security detection models can be extended to time-based events, providing there is sufficient time series data. These models are for a non-evolving system, mostly detecting patterns of incursion after the fact. While AI is on the path to rapid development, at present, skilled security professionals are required to provide the analysis to detect and prevent new and evolving threats. Large volumes of data combined with well-trained models may be able to detect well-known attack

© Copyright 2020, Cloud Security Alliance. All rights reserved.

vectors. However, to detect new avenues of ingress with fraudulent intent, that have never been seen before, requires a combination of performance monitoring, pattern analysis of transaction data and analysis provided by security specialists. Relying on endpoint monitoring alone still leaves enterprises vulnerable to undetectable attacks.

SDP Perspective: **For highly confidential data, the best method of security is to prevent attacks before they occur. An SDP Zero Trust deployment can deny risky transactions based on a single packet analysis revealing a lack of positive identification.**

c) **Packet inspection has no user context** - Network packet inspection has its limitations in that packet 'analysis' happens at the application layer, so incursions can happen prior to detection.

Network single packet inspection to identify connections are innovative and successful within bounds. These methods are only as secure as TCP/IP and TLS protocols and application code.

Traditionally packet inspection happens on or close to the firewall with an Intrusion Detection System (IDS) and/or on strategically identified areas important for monitoring. Traditional firewalls typically control access to network resources based on source IP addresses. The fundamental challenge with inspecting packets is the problem of identifying the user from the source IP. The tools for inspection are based on IP addresses. While some attacks such as DDoS and malware may be detected using existing techniques, the vast majority of attacks such as code injections and credential theft require a context to detect, as they are performed at the application layer.

SDP Perspective: **On the contrary, SDP does have packet inspection end user context. With an SDP Zero Trust deployment, dropped packets gathered at SDP gateways can be forwarded for out-of-band inspection and analysis. Combined with network data, a risk profile can be detected before ingress.**

# Implementing a Zero Trust Strategy

Zero Trust is a philosophy for designing network security architecture in a way that withholds access until a user, device or even an individual packet has been thoroughly inspected and authenticated and authorized. Even then, only the least amount of necessary access is granted based on authorization to access. The following constructs are required to adopt a Zero Trust strategy.

a) Authentication before access

Using VPNs and Firewalls to establish Zero Trust allows the user to connect to services (e.g. a mail server). Firewalls can be set up to blacklist IPs, and services can be set up to determine which IP addresses to allow or deny. VPNs can be configured to only allow the users on the network who have the authorized VPN client and the appropriate keys, suggesting a Zero Trust has been implemented. However, unauthorized users who clone the VPN client and steal the keys can also access the mail server and then guess other user names and passwords and perform malicious acts such as DDoS, credential theft, etc. The VPN allows a user to log into the network and deny other services (e.g. SharePoint) not on the mail server network

© Copyright 2020, Cloud Security Alliance. All rights reserved.

segment. If unauthorized users are already in the network, lateral access to a SharePoint server is a common occurrence. Access before authentication allows users more access to services than is intended by access rights.

In order to ensure authentication before access, there is an implied requirement: a control plane for authentication that is separate from the data plane. To ensure acceptable response times, a mechanism for immediate authentication is also required.

b) Capability to limit network connectivity and exposure

Public/private clouds do configure perimeter network security. They provide a layered approach to security, stream logs to monitoring tools and provide insight and hybrid service control policies. However, these features do not address the problem of challenging authentication prior to access.

Strong supporting measures for cloud native platform and application services include inbound/outbound security configuration and corporate network policy configuration. An industry standard practice for strong authentication and authorization is mutual TLS (two-way SSL) certificates. A better approach is to require authentication before access, the drop or forward packets at the network layer with traffic management provided by an SDN controller interfacing to an SDP Zero Trust Deployment. With this architecture, the SDN infrastructure can drop network connections if authentication fails.

c) Granular trust authentication mechanism

Network Layer VPNs and firewalls and application layer firewalls and SSL VPNs do not have explicit fine-grained access control. A Zero Trust deployment implicitly requires not only policy-based authorization but also identity authentication in the context of network micro segmentation and distributed service connectivity and interconnectivity across hybrid private/public multi cloud scenarios.

Network Layer Firewalls merit specific consideration. They are static, so user groups are used to provide granular trust. It is not unusual to have a group of users from a variety of departments with different roles needing access to the same service with the same IP address. Firewalls rules are static and rely only on network information. They do not dynamically change based on context, i.e. the level of trust required for a given device from a given network. A frequent use case is where a user requests access through riskier network such as an internet café. If local firewall or antivirus software has been turned off by malware or by accident, this will not be detected by traditional firewall.
A case can be made for IPSec VPNs, which do not access identity attributes for authentication prior to allowing access. Instead IPSec VPNs are reliant on tokens and credentials that may have been intercepted. SSL VPNs have known vulnerabilities.

In light of these limitations, a network perimeter Zero Trust approach is more secure with a granular trust authentication mechanism and policy-based authorization.

© Copyright 2020, Cloud Security Alliance. All rights reserved.

d) Monitoring suspicious activity

Consider when authentication of identity attributes fails. The capability to forward suspicious activity based on packet inspection to endpoint logging and monitoring services provides really useful inputs to the security orchestration, automation and response (SOAR) technologies that enable organizations to take inputs from a variety of sources (mostly from security information and event management (SIEM) systems, see SDP Architecture Guide for details). Automation refers to the workflow processes that are initiated to gather data, to be integrated and orchestrated, providing operational intelligence and visualization graphs and dashboards. Zero Trust implementations can forward useful intelligence for input to SOAR AI models and the proper monitoring of suspicious activity.

© Copyright 2020, Cloud Security Alliance. All rights reserved.

# Benefits of an SDP Zero Trust Solution

A Software Defined Perimeter Zero Trust Solution can provide the following security and Business benefits as defined in the CSA SDP Architecture Specification.

## SECURITY BENEFITS

| Benefit | Description |
|---|---|
| Reduced attack surface | Protects critical assets and infrastructure by separating access control and data planes to render each of them "black," thereby blocking potential network-based attacks |
| Protects critical assets and infrastructure | Enhances protection for cloud applications by hiding them:<br>• Gives more centralized control to business/system owners<br>• Provides visibility to all authorized connections from whom, where, when, to what<br>• Enables instant monitoring because controls are integrated |
| Ability to hide assets by denying connectivity | Enables deny-all gateway until users/devices are authenticated and authorized to access the assets |
| Reduced cost of ownership | Reduces costs for endpoint threat prevention/detection Reduces cost for incident response Reduces complexity for integrating controls |
| Provides connection-based security architecture | Provides connection-based security architecture instead of IP-based alternatives, because today's explosion of IPs and loss of perimeter in cloud environments render IP-based securities weak |
| Provides an integrated security architecture | Provides an integrated security architecture that is otherwise hard to achieve with existing security point products, such as NAC or anti-malware SDP integrates the following discrete architectural elements:<br>• User-aware applications<br>• Client-aware devices<br>• Network-aware firewalls/gateways |
| Using Single Packet Authorization | Determines connections and enables integrated controls for authentication and authorization |
| Requires pre-vetting of connections | Allows control of all connections based on pre-vetting of who can connect, from which devices, to what services, infrastructure and other parameters |
| Authenticates BEFORE allowing access to resources | Implements a separate control and data channel Enables validation prior to TLS/TCP handshake Provides fine-grained access control that is implicit in the design Allows enforcement of two-way mutually encrypted communications |
| Open specification | Allows vetting community Facilitates participation in hackathons |

© Copyright 2020, Cloud Security Alliance. All rights reserved.

## BUSINESS BENEFITS

| Benefit | Description |
|---|---|
| Cost and labor savings | Reduces licensing and support costs since traditional network security components are replaced with SDP. Reduces operational complexity and reliance on traditional security tools due to implementation and enforcement of security policies using SDP Reduces costs by minimizing or replacing MPLS or leased line utilization. Organizations can reduce or eliminate the use of private backbone networks Brings efficiency and simplicity to organizations, which can ultimately help reduce labor needs |
| Increased agility of IT operations | IT processes can act as a drag on business processes. SDP implementations, on the other hand, can be driven automatically by IT or IAM events. These benefits accelerate IT, making it more responsive to business and security demands. |
| GRC benefits | Delivers reduced risk compared to traditional approaches SDP suppresses threats and reduces attack surfaces, preventing network-based attacks and the exploitation of application vulnerabilities. SDP can feed into and respond to GRC systems (such as when integrating with SIEM) to streamline compliance activities for systems and applications. |
| Compliance benefits | Compliance data collection, reporting, and auditing processes can be improved by SDP through the centralized control of connections from users on registered devices to specific applications/services. SDP can provide additional traceability of connectivity for online businesses. The network micro-segmentation provided by SDP is frequently used to reduce compliance scope, which can have a significant impact on compliance reporting efforts. |
| Secure cloud computing adoption | Can help enterprises rapidly, confidently, and securely adopt cloud architectures by reducing the costs and complexity of the required security architecture to support applications in the public-cloud, private-cloud, data-center, and mixed environments. New applications can be deployed faster with equivalent or better security than other options. |
| Business agility and innovation | Enables businesses to implement their priorities quickly and securely. Examples include:<br>• Enables transition from on-premises call-center agents to home-based agents<br>• Enables the outsourcing of non-core business functions to specialized third parties<br>• Enables customer-facing kiosks on remote third-party networks and locations<br>• Enables deployment of company assets onto customer sites, creating stronger integration with customers and generating new revenue |

© Copyright 2020, Cloud Security Alliance. All rights reserved.

| | |
|---|---|
| Facilitates business transformation | Facilitates IoT adoption via segmentation and permissions Allows access to transformation engineers without compromising existing services Creates next gen secure systems combining IoT with private permission Blockchain |

# SDP Zero Trust Strategic Approach and Proof of Concept

In light of large-scale enterprise breaches, separating sensitive information resources into high-security networks for essential services and protecting data privacy, important measures to take. Recent analysis by the CSA 2019 Cloud Security Threat Report indicates that risky human behavior continues to be responsible for a significant proportion of data breaches, along with cloud malware injection and DDOS incidents.

A new network architecture paradigm, known as Cloud Security Alliance's (CSA) Software Defined Perimeter (SDP) protocol was initiated in 2013. It was designed to create an architecture for positive identification of network connections from single packet inspection prior to accessing sensitive data. Implicit in this architecture is the separation of the control plane where trust is established, from the data plane where actual data is transferred. This removes the vulnerabilities inherent in TCP and TLS termination, as well as the vagaries of network firewalling by IP Network Address Translation (NAT) tables.

SDP provides a simple means of preventing the negative consequences of people bypassing enterprise and legal security controls in the cloud.  Adopting an SDP implementation enforces the separation of establishing trust from data transfers. Network segmentation and the establishment of micro networks, so important for multi-cloud deployments, also benefit from adopting a software defined perimeter ZT architecture.

Combining Software Defined Perimeter with multi-factor authentication and improved access control/authorization mechanisms puts organizations on a strategic path to addressing security vulnerabilities and large-scale intrusions. Software defined perimeters enforce security policies at configuration and deployment in addition to runtime detection and response.

© Copyright 2020, Cloud Security Alliance. All rights reserved.

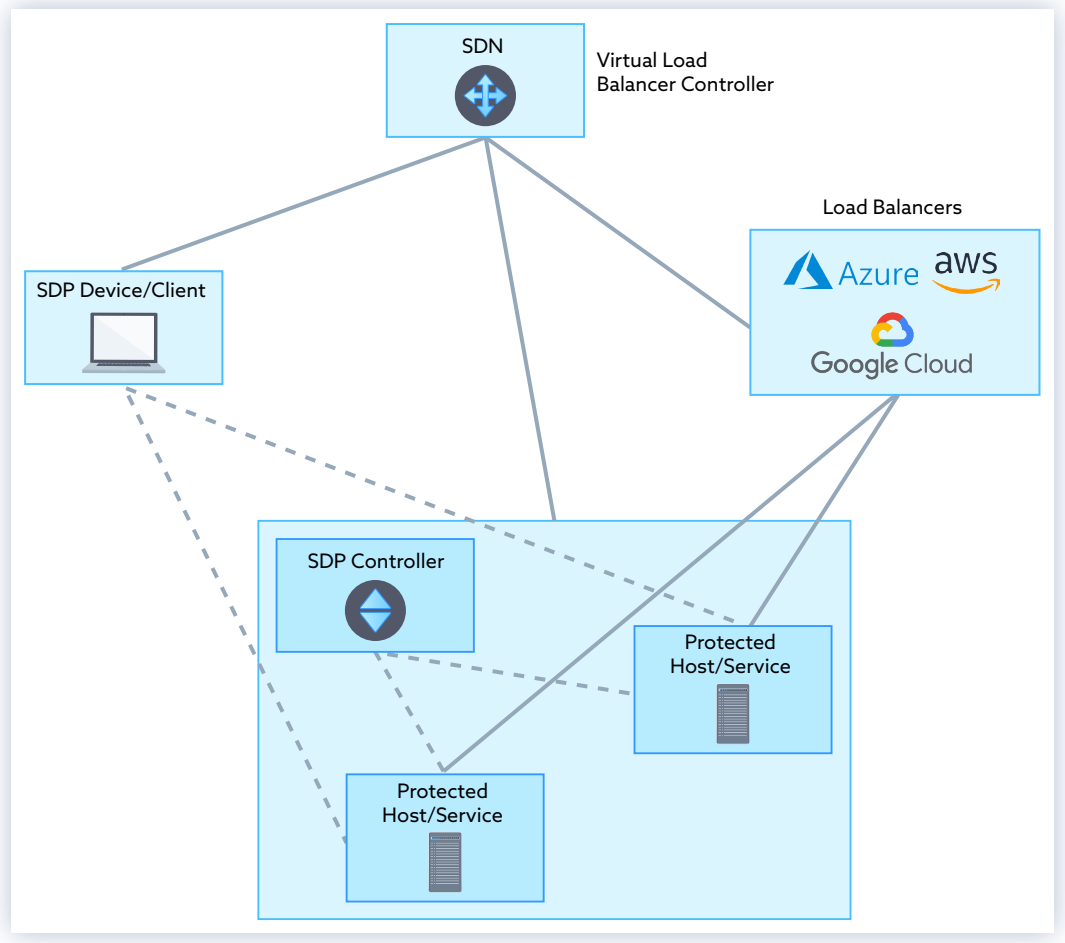


*Figure 4: Hybrid Cloud Environment*

An SDP architecture proof of concept (POC) can demonstrate how SDP addresses the challenges of application delivery in a hybrid multi-cloud environment. Specifically, the POC demonstrates the following:

- communications that are classified as highly sensitive can be secured over any type of network, even the internet, from one secure environment to another without having to run the gauntlet of network layer to application layer insecurities, using an SDP approach;
- advances in Software Defined Networking can support a Software Defined Perimeter by providing the support of a separate control and data plane as well as a deny-all/allow firewall implementation; and
- the SDP approach to network forwarding across a hybrid multi-cloud deployment is perfectly aligned with the principles of zero trust networking based on a single packet inspection.

The inference is that an SDP deployment applied at the network layer uses an interface from an SDN controller to route connections from the initiating host to an SDP controller. The preferable interface for configuring this routing is a REST interface enabling self-service configuration.

The rationale for providing this network layer SDP demonstration is to address the problems caused by applying Zero Trust at the Application Layer after TLS termination. Most of the existing "Zero

© Copyright 2020, Cloud Security Alliance. All rights reserved.

Trust" security measures are applied as authentication and "sometimes" authorization based on policy after the termination of TLS certificates. Certificate validation is a complex verification and validation process, and there are known possible vulnerabilities with TLS 1.2, TLS 1.3 and mutual TLS.

There are a number of initiatives to address Zero Trust approaches to content delivery from all major Cloud Service Providers. To date there are no applications of Zero Trust solely at the Network Layer, particularly involving hybrid multi-cloud deployments. In this instance, "hybrid cloud" refers to connectivity from private clouds to enterprise to data centers, while "multi-cloud" refers to network connections across different public and private clouds. Industry sources indicate that most enterprises now have, or are intending to have, a hybrid multi-cloud strategy.

The proof of concept takes advantage of the fact that virtualization of networks makes it possible for security-related actions to be performed in the control plane of SDP implementations. To put this in perspective, the widespread adoption and evolution of Software Defined Networking (SDN) has enabled the service providers to simplify network management. However, this wide adoption of SDNs is posing real challenges on how to provide proper authentication, access control, data privacy, and data integrity among others for the API-driven orchestration of network routing. Although SDN allows virtual networks provision on demand for both efficient data transport and fine-grained control services, current security practices were not designed to match the complexity and challenges that emerge from the integration of these software defined infrastructure. However, the Software Defined Network paradigm allows for an SDN controller to call a Software Defined Perimeter service that can orchestrate connections and perform an allow/deny action on a network connection based on the SDP identity and device verification of the request[4]. The SDP controller then instructs the SDN to either to route the connection to the accepting host or to drop the connection when the packet identification attributes do not pass the required checks.

[4] On the Security of SDN: A Completed Secure and Scalable Framework Using the Software-Defined Perimeter (http://sdpcenter.com/resources/research/)

© Copyright 2020, Cloud Security Alliance. All rights reserved.

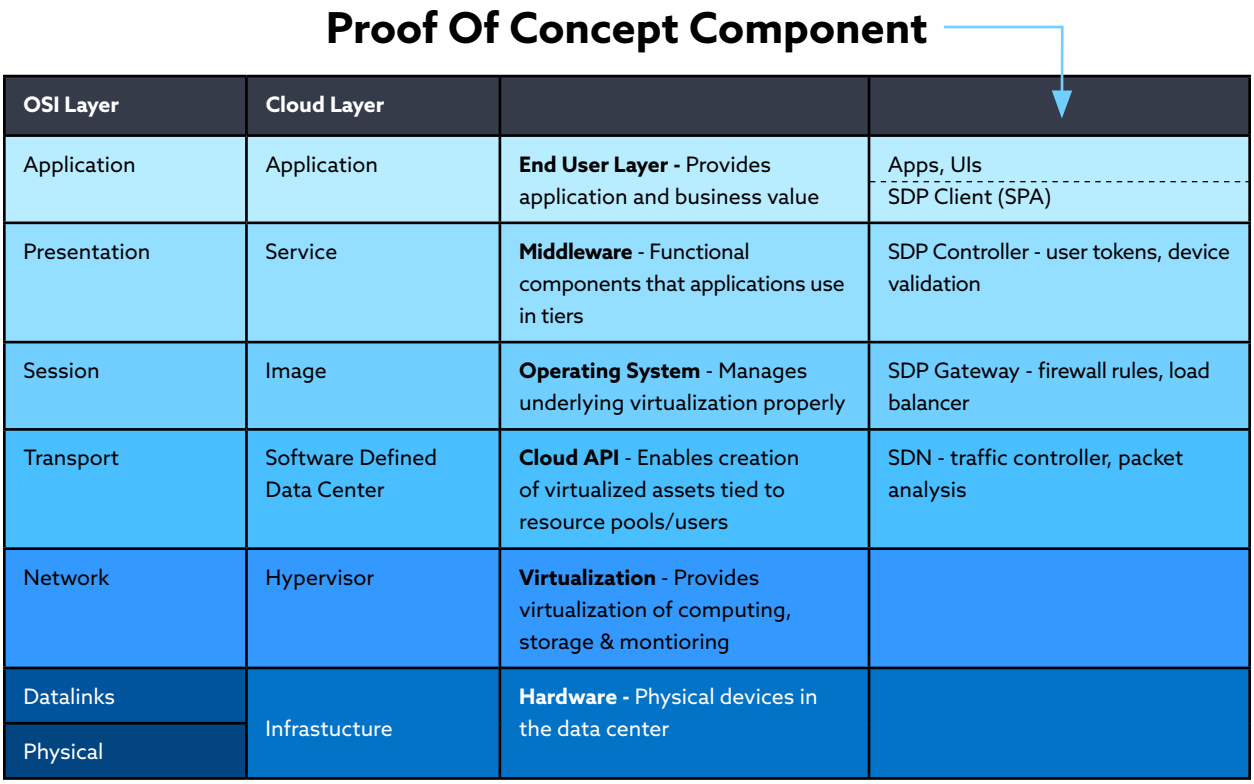


*Figure 5: Proof of Concept SDP Components*

# Technology components and Infrastructure

The following services are required to demonstrate that an SDP deployment can address vulnerabilities in enterprise information technology capabilities. This demonstration builds on the existing SDP open source deployment. The demonstration is to be made publicly available, and where open source is not possible, technology suppliers are to provide clear instructions on configuration and deployment of components.

**SDP control and data plane technology components**

1. SDP agent deployed on an SDP client
2. SDP controller hosted in a location that is accessible to the SDN Virtual Load Balancer (VLB)
3. SDP host endpoints that require a Zero Trust allow/deny security posture (For the purposes of this proof of concept these servers can be VMs deployed on a public cloud accessible by an external cloud load balancer.)
4. Network connectivity from the public internet

**Network Load Balancer Controller and public cloud technology components**

1. SDN Virtual Load Balancer capable of routing an SDP request to an identity access control microservice and determining an allow/deny response
2. Public cloud external load balancers for VLB to forward requests to SDP accepting hosts/services deployed to public cloud services

© Copyright 2020, Cloud Security Alliance. All rights reserved.

# Technology Risks and Issues

There are risks when deploying SDP/Zero Trust at the network layer along with software-defined networking and virtual load balancers. As SDP allow/deny is a binary choice for network connectivity, clearly this implementation insinuates a single point of failure. It is therefore critical that the integrated access control mechanism at the data plane has deep security. This ensures that the attribute used for identification is securely planned, built and run.

# Assumptions

1. An existing open source SDP deployment is to be used as the basis for the proof-of-concept demonstration.
2. A virtual load balancer is selected that is capable of calling a microservice and forwarding requests to common cloud and on-premises load balancers.
3. Supplier technology deployment of the proof-of-concept is to be made publicly available, with detailed implementation details not including proprietary or private capabilities.
4. Virtual Private Cloud networking micro deployments are to be made available for the purposes of the PoC.
5. A device or VM acting as an SDP initiating host/server is to be made available.
6. A test environment is to include test cases to cover "eligible" and "ineligible" connections that is, the identity attributes in the request packet are either matched in the identity service (connection allowed) or not (connection dropped).

# Technology Analysis

The technology components required to deploy a true network layer Zero Trust allow/deny connectivity posture requires access to a new connection prior to the application of network protocols at the accepting host.

The proof of concept requires a component to be deployed during traffic management, during routing, prior to TLS certificate termination and without exposing the actual accepting host prior to acknowledgement on the final TCP/IP endpoint destination. This has the obvious advantages of preventing incursion through certificate weaknesses, as well as preventing DDoS packets reaching the target.

The technology required is the deployment of a packet inspection service with access to identity attributes directly from the virtual load balancer. The connection can be routed via an SDP controller service that makes the decision to drop the connection or allow forwarding to the accepting host server based on the identity attributes service.

To facilitate current network environments where multiple environments may be connected for a single or multiple service, the technology required to interface with the SDP deployment is a virtual load balancer able to interface with Cloud Service Provider and enterprise load balancers.

© Copyright 2020, Cloud Security Alliance. All rights reserved.

This technology must also be able to connect to a VM deploying the SDP controller, as well as intercepting SDP relevant requests from an initiating client/server device or VM.

# Required Resources

Components required to deliver the SDP Zero Trust proof of concept demonstration are as follows:

1. Client/Server initiating a network connection
2. Internet network connectivity
3. Virtual load balancer capable of calling REST services prior to forwarding based on packet inspection
4. SDP controller deployment microservice
5. Client/Servers accepting the network connection
6. CSP/enterprise external load balancers to forward requests to accepting client/servers

Note: Client/server is a general term, with no particular inbound/outbound direction implied.

| Requirements | Components |
|---|---|
| Connections require authentication before access | Identity attributes verification service deployed at SDP controller |
| Capability to limit network connectivity to vetted connections and exposure | Virtual Load Balancer controller to drop connections that are not vetted by SDP controller |
| Identity and Access Management granular controls | Single Packet Inspection of each connection forwarded by VLB deployed at SDP controller to authenticate each connection at runtime |
| Forwarding of suspicious activity to monitoring system | VLB controller to forward information about suspicious connections to an endpoint monitoring SIEM service |

# Key Industry Developments

SDN advances, specifically virtual load balancer controllers configured to REST service APIs, and the capability to route network connections make this proof of concept viable.
Virtual load balancer control plane services are therefore capable of making intelligent decisions based on connection request network packets. This means that the deployment of network layer authentication of identity attributes can now be achieved by calling a perimeter-protected service over REST that can validate packet identity attributes.

© Copyright 2020, Cloud Security Alliance. All rights reserved.

# Delivery Activities

Implementation of an SDP Zero Trust proof-of-concept demonstration requires the following activities:

1. Set up of virtual private cloud networks and virtual machines
2. Establishment of internet connectivity between endpoints
3. Identity validation microservice
4. Set up of virtual load balancer, routing to CSP external load balancers public IP addresses
5. Packet inspection to determine SDP connections
6. Extraction of identity attributes from connection single packet inspection
7. REST service to identity validation microservice
8. Routing of network connections between VLB and SDP controllers

# Situation Analysis

Currently many suppliers and vendors are claiming 'Zero Trust' capability for their product and service offerings. While the following capabilities and activities support a Zero Trust network capability, they do not constitute Network Layer Zero Trust without being able to authenticate prior to authorized access to a network endpoint.

- Configuration of perimeter network security
- Streaming logs to insight monitoring tools
- Configuration of hybrid service control policies
- Inbound/Outbound firewall security configuration
- Corporate network policy configuration
- Authentication with mutual TLS certificates
- Authorization with Single Packet Authorization

Most authentication takes place after the TCP/IP protocol acknowledgement and after TLS certificate termination validation.

# Timeframes and Stakeholder Engagement

Suppliers of technology components who are interested in participating in an SDP Zero Trust demonstration are invited to contact the Cloud Security Alliance Software Defined Perimeter Working Group by email to smahmud@cloudsecurityalliance.org.

© Copyright 2020, Cloud Security Alliance. All rights reserved.

# References


Cloud Security Alliance Initiatives:

- SDP Architecture Guide published May 2019
  https://cloudsecurityalliance.org/artifacts/sdp-architecture-guide-v2/
- SDP as a DDoS Defense Mechanism published October 2019
  https://cloudsecurityalliance.org/artifacts/software-defined-perimeter-as-a-ddos-prevention-mechanism/
- Specification 2.0 in Jan 2020 - In progress

Market Awareness and Adoption Overview:

- Cloud Security Alliance The State of SDP Survey: A Summary

Open Source Reference Implementation (funded by DHS):

- http://sdpcenter.com/test-sdp/

Zero Trust Presentation to OMG (Object Management Group):

- https://cloudsecurityalliance.org/artifacts/sdp-the-most-advanced-zero-trust-architecture/

US Department of Defense Net-Centric Services Strategy:

- https://dodcio.defense.gov/Portals/0/documents/DoD_NetCentricServicesStrategy.pdf


© Copyright 2020, Cloud Security Alliance. All rights reserved.