June 2022

# OBTAINING VISIBILITY INTO A SECURE ACCESS SERVICES EDGE (SASE) NETWORK

Indermeet Gandhi

Robert Barton

Jerome Henry

# OBTAINING VISIBILITY INTO A SECURE ACCESS SERVICES EDGE (SASE) NETWORK

AUTHORS:
Indermeet Gandhi
Robert Barton
Jerome Henry

## ABSTRACT

Techniques are presented herein that provide visibility into Secure Access Services Edge (SASE) components for real user traffic as well as for synthetic traffic. Aspects of the presented techniques support the stacking of labels as a packet travels through a service chain and then the review of same for audit purposes, thus providing audit and assurance capabilities for a customer. Those capabilities may be used to reassure the customer that they are actually receiving the correct combination of SASE services. Aspect of the presented techniques may be used to reveal violations or exceptions in a service level agreement (SLA) where, for example, certain security services were not properly engaged. Such information can help in providing a customer with complete insight and troubleshooting support.

## DETAILED DESCRIPTION

The Secure Access Services Edge (SASE) space is a must-win market for network equipment vendors. SASE components may include software-defined wide-area networks (SD-WANs), a cloud access security broker (CASB), a cloud-delivered firewall (CDFW), remote browser inspection (RBI) capabilities, Domain Name System (DNS) security, and secure web gateways (SWGs).

Network equipment vendor SD-WAN solutions include different products and cloud security (encompassing a CDFW, SWGs, zero trust network access (ZTNA), a virtual private network (VPN), data loss prevention (DLP) capabilities, and a CASB) and may be provided by a cloud-based Secure Internet Gateway (SIG) platform.

In current deployments, cloud security may operate like a black box to customers. A customer may purchase a list of SASE services, but it can be very difficult to provide audit and assurance capabilities to the customer to reassure them that they are actually

receiving the correct combination of services that they require for each flow. In other words, how can a customer be assured that they are being provided with the correct security service chain as flows traverse the SASE cloud?

Figure 1, below, illustrates an exemplary arrangement that is reflective of the above discussion.

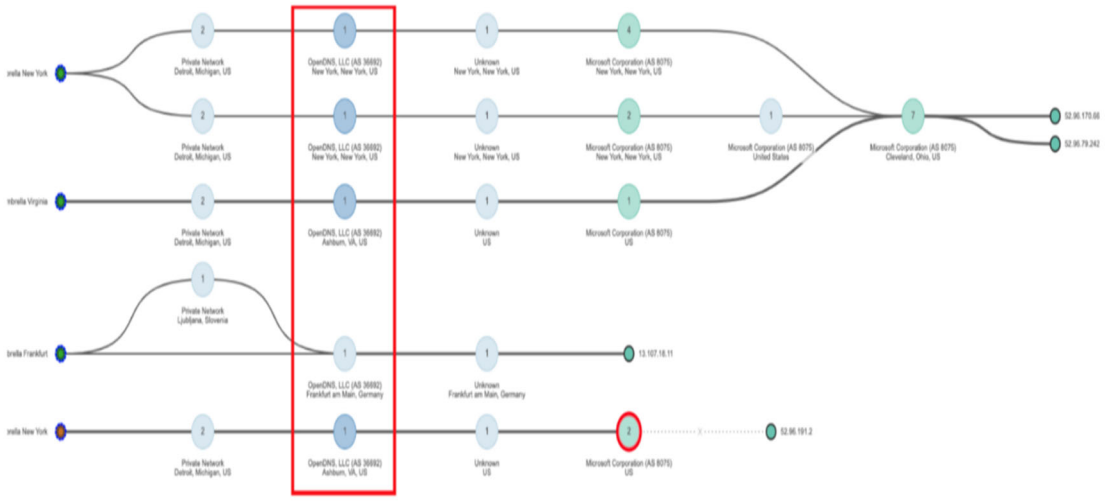

*Figure 1: Exemplary Arrangement*

As depicted in Figure 1, above, customers only see the services of a cloud-based SIG platform as a single hop (e.g., OpenDNS) without knowing which services are being applied (i.e., service assurance). Such a view limits a customer's ability to troubleshoot issues or provide an audit trail for compliance reasons. Further, if a customer has applied a policy of skipping decryption for a certain set of services (such as an online business productivity application), then the visual representation should provide a hop-by-hop view of decryption being skipped or decryption being applied for different applications.

A better mechanism is needed for showing the hop-by-hop visual flows of how customer data traffic is flowing across the SASE service (e.g., from a CDFW to a SWG to a CASB, etc.). Such a depiction would also expose the shortcomings of some network equipment vendors who do not have a full portfolio of cloud security services.

To address the challenge that was described above, techniques are presented herein that provide visibility into the SASE components for real user traffic as well as for synthetic traffic. While SASE, inbound Operation, Administration, and Maintenance (iOAM)

capabilities, and tagging mechanisms for Internet Protocol (IP) packets are well-known, aspects of the presented techniques support the stacking of labels as a packet travels through a service chain and then the review of same for audit purposes.

Figure 2, below, depicts elements of an illustrative arrangement according to aspects of the techniques presented herein and reflective of the above discussion.

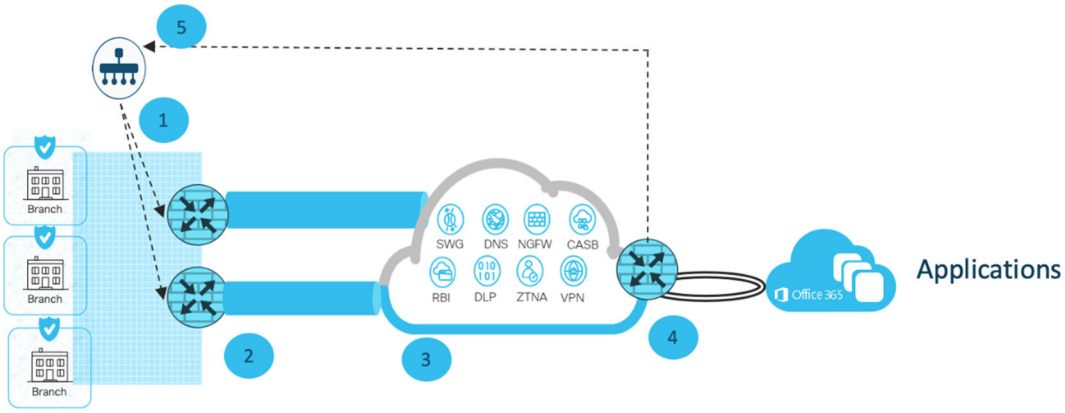


*Figure 2: Illustrative Arrangement*

Figure 2, above, depicts a series of steps. Those steps, which are labeled 1 through 5 in the figure, will be described below.

In Step 1, a network controller programs an SD-WAN edge router by marking the tunnels to the SASE service. A connection to the SASE cloud may also be initiated by an edge client or agent.

During Step 2, when a connection is made from an originating device, a service flag or label is written into the IP packet header (i.e., into a metadata field that serves as an audit flag) that will be used to alert the SASE provider that a service chain audit trail must be created for the instant flow. A timestamp may also be appended to the flag. Thus, when the source traffic leaves either the user machine or the SD-WAN edge, a services flag is added to each packet of the flow. Such an addition may be imposed on a per-application basis, or it may be imposed for an entire device or user. A services flag as described above may be added as part of the IP options field for IP version 4 (IPv4) traffic, or it may be an IP version 6 (IPv6) extension header, or it may be some other custom flag.

Under Step 3, as the user data enters the SASE provider the provider recognizes the services flag and copies it into a Generic Network Virtualization Encapsulation (Geneve) option header. At ingress to the SASE cloud a timestamp may be appended for assurance purposes. Aspects of the techniques presented herein extend the Geneve option header to include an audit trail for the SASE service chain (e.g., a CASB, SWGs, RBI, encryption, DLP facilities, etc.). As the user traffic flows across the SASE provider, each service it passes through in the service chain may be enhanced to add their own identifying service flag. According to aspects of the techniques presented herein, those service flags may be stacked, one-by-one, to form an audit trail in the Geneve option header. Such an audit trail provides evidence that the traffic flow actually passed through each intended service. In brief, the stack of labels may be added to the packet at each service element (along with a timestamp) to provide assurance that the service chain was correctly executed, within performance requirements, as the traffic flow moved through the SASE cloud.

During Step 4, at the egress point of the SASE cloud provider, a router may strip off the Geneve option header, map the stack of services flags to the list of the services that the flow traversed, and report the same to a database. The label stack and timestamps may then be used for analysis and assurance purposes.

Under Step 5, the stack of labels may be analyzed and the results may be compared to a service level agreement (SLA) to both ensure the type of security services that were executed as well as the performance metrics that were realized. Such an audit trail may also be provided to a customer. A visual interface may be provided to the customer that displays a hop-by-hop assurance in a topological view, indicating the exact service chain for different users, applications, devices, etc. Such a service chain audit may also be used to reveal violations or exceptions in an SLA where certain security services were not properly engaged, thus providing the customer with complete insight and troubleshooting support.

According to further aspects of the techniques presented herein, synthetic traffic may also be generated to test various applications and devices to ensure, through the production of an audit trail for that traffic, that those entities are compliant with a policy. Additionally, iOAM traffic may be generated by a customer to continuously validate the assurance of a security service chain.

Figure 3, below, presents elements of a hop-by-hop view of an illustrative cloud-based SIG platform according to aspects of the techniques presented herein.

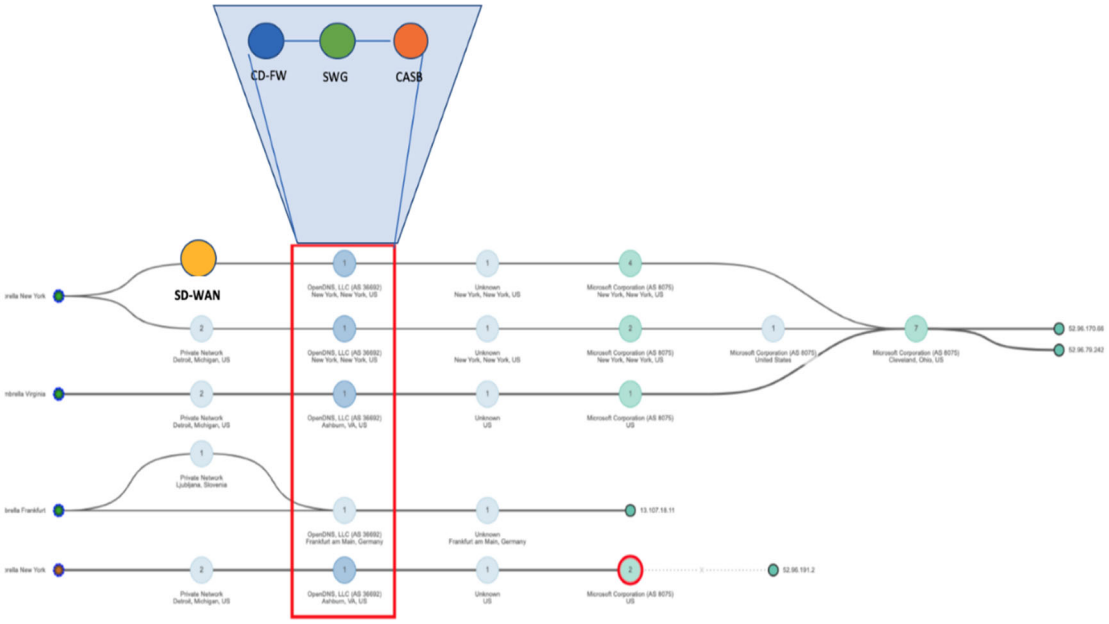


*Figure 3: Illustrative Hop-by-Hop View*

It is important to note that while the concept of header tags is well known, the techniques presented herein completely invert that concept. Under aspects of the presented techniques, each service element in the SASE cloud may impose its own validation tag as a flow progresses. Such an approach provides a secure audit trail and a chain of trust that proves that the service was delivered as expected or as required by an SLA.

Additionally, the Internet Engineering Task Force (IETF) Request for Comments (8979) discusses the form and function of variable-length context headers that carry into a service chain information about a subscriber and performance requirements. However, that RFC does not address auditing or validating service elements in any way, let alone using a progressive label stack as packets transit the service chain for audit and assurance purposes, as do the techniques presented herein.

While the ability to tag specific packets with actions for troubleshooting purposes is well known, the techniques presented herein support more than just imposing a new tag for troubleshooting. Aspects of the presented techniques validate an end-to-end (E2E) and

hop-by-hop assurance through a SASE service chain – i.e., both that the required services were executed and that the desired performance targets were met.

As described previously, aspects of the presented techniques support not just a tag for interesting traffic, but rather a stack of labels at each service element along with a timestamp. Beginning at ingress, then at each hop, a new label may be added to the stack. At egress, all of the labels may be removed and then used as a service chain assurance mechanism. Not only does such an approach provide evidence to a customer that their flows are being given the security services that they signed up for, but use of the timestamps allows total latency to be measured and performance to be expressed on a per-flow basis. All of that information may then be provided back to the customer.

Unlike on-premise deployments, where services physically reside in customer networks under the customer's ownership and control, cloud security services follow a different (e.g., an as-a-service) paradigm. When a customer purchases a list of SASE services it can be very difficult to provide audit and assurance capabilities to the customer to reassure the customer that they actually receive the correct combination of services that they require for each flow. In other words, how can the customer be assured that they are being provided with the correct security service chain as flows traverse the service chain in a SASE cloud (where the service is indeed a black box and the question becomes one of verifying the outcome)? Service chain audits, according to aspects of the techniques presented herein and as described above, address that challenge. Service chain audits may also be used to reveal violations or exceptions in an SLA where certain security services were not properly engaged. Such information can help in providing a customer with complete insight and troubleshooting support.

As described in the above narrative, the techniques presented herein comprise a number of capabilities. For example, techniques herein provide for the ability to impose a label in a packet header at a point of ingress to an SASE cloud to identify the time of ingress, an entry point, and a customer identifier (ID). Such a label may be used at the point of egress to determine an E2E assurance through the service chain. Further, techniques herein may provide for imposing a unique identifier or label in the packet, along with a timestamp, at each element of a service chain. Additionally, techniques herein may provide for

stacking the identifying labels (as described above) as the packet progresses through a service chain.

Still further, techniques herein may provide for removing all of the labels at an egress, exporting the same, and storing the same in an analysis engine. Moreover, techniques herein may provide for the analysis engine comparing the executed service chain (per-flow) with an SLA and providing evidence, back to a customer, that the SLA is being adhered to (i.e., that all of the service elements were executed). Finally, techniques herein may provide for using the labels and the timestamps (as described above) to produce a performance analysis for a service chain for each individual flow indicating a hop-by-hop and E2E performance through the service chain.

In summary, techniques have been presented herein that provide visibility into the SASE components for real user traffic as well as for synthetic traffic. Aspects of the presented techniques support the stacking of labels as a packet travels through a service chain and then the review of same for audit purposes, thus providing audit and assurance capabilities for a customer. Those capabilities may be used to reassure the customer that they are actually receiving the correct combination of SASE services. Aspect of the presented techniques may be used to reveal violations or exceptions in an SLA where, for example, certain security services were not properly engaged. Such information can help in providing a customer with complete insight and troubleshooting support.