February 2020

ASIC-FRIENDLY SRV6-BASED SD-WAN SERVICE THEFT PREVENTION MECHANISM

Clarence Filsfils

Pablo Camarillo

ASIC-FRIENDLY SRV6-BASED SD-WAN SERVICE THEFT PREVENTION MECHANISM

AUTHORS:

Clarence Filsfils
Pablo Camarillo

ABSTRACT

Presented herein is a security mechanism that prevents service theft with software-defined networking in a wide area network (SD-WAN) services provided with Segment Routing over IPv6 dataplane (SRv6). The security mechanism described herein is valid in SRv6 deployments, but also provides line-rate security at PEs that do not have dedicated crypto-hardware.

DETAILED DESCRIPTION

Segment Routing over IPv6 dataplane (SRv6) technology provides virtual private networks (VPNs) with underlay service-level agreement (SLA) requirements. This applies to both Over-The-Top VPNs and Software-Defined WANs (SD-WANs). SD-WANS are particularly interesting for operators, who are seeing a constant decline in revenues due to enterprise-based SD-WAN solutions. Providing an SRv6-based SD-WAN service allows operators to bring-back part of this lost revenue.

In case that a customer-provided equipment (CPE) is directly connected to the operator's network, there is no security risk. However, in case that the CPE is remotely connected, then there is a need to provide a security mechanism that prevents service theft. This mechanism must ensure that an SD-WAN customer is only able to access the services that have been instantiated for him (i.e., customer 1 cannot access the services of customer 2).

For example, assume there are two customers (customer 1 and customer 2), and two providers (provider A and provider B). Customer 1 hires an SD-WAN service from provider A and provider B. Each provider respectively assign anycast SIDs A:88:1:: and

B:88:1:: to this SD-WAN service. Both A:88:1:: and B:88:1:: are accessible in the open Internet, but customer 2 must be unable to access the service behind both as he is unauthorized to use those SD-WAN services.

In one solution, an operator might provide the SD-WAN service over several PEs that are sharing an anycast address. Similarly, the customer might have more than a single CPE. A scalable security solution is needed that does not require security association for each [Source,Destination] pair.

It is noted that neither encryption nor data integrity are part of the requirements and do not provide any benefit (since the two SD-WAN CPEs will always have an underlying end-to-end IPSec session). However, it is extremely interesting to provide DDoS and anti-replay protection to the SD-WAN service. It is also noted that the SD-WAN service is provided from a PE router on an operator network that, conventionally, does not have any crypto hardware installed thereon. Therefore, any solution that relies on the use of crypto-hardware (e.g., IPSec AH) is not considered acceptable.

The security mechanisms presented herein leverage a common application-specific integrated circuit (ASIC) construct, known as the forwarding information base (FIB). Through the use of a nested FIB lookup mechanism, two solutions with different levels of protection are presented herein. In the simplest solution, referred to below as "Solution A," the mechanism provides the same level of security as HMAC without the requirement on crypto-hardware. In the more advanced solution, referred to below as "Solution B," the techniques provide loose anti-replay protection based on synchronized time across devices at the expense of CPU utilization for regular FIB updates (e.g., additional 10% CPU overhead). These two mechanism are likely sufficient to prevent service theft in SRv6 SD-WAN deployments, where encryption and data integrity are orthogonal to these mechanisms.

Solution A: FIB-based Security Mechanism

- The PE nodes advertise an SD-WAN service accessible through the anycast SID A:88:1::/64 (SID argument size 64b).
- The SID is associated with a unique shared secret key.
- On the PE router:

  - The main FIB contains an entry for A:88:1::AAAA:BBBB:CCC:DDDD, where AAAA:BBBB:CCCC:DDDD is the result of the hash of [IPv6_CPE_SA, SID, key] -trimmed to the most-significant 64b-.
  - This FIB entry (/128) is associated with a particular VRF.
    - The VRF will contain the IP addresses of the customer CPEs. The uRPF check will be performed on this particular VRF.
    - Additionally, the VRF is populated with the globally routable SIDs and all the destinations that are known to be reachable once that SD-WAN SID has been visited.
- On the CPE:
  - Receives the advertisement of the anycast SID A:88:1::/64.
  - Computes the full IPv6 DA A:88:1::AAAA:BBBB:CCCC:DDDD a single time based on the hash of [IPv6_SA, SD-WAN Service SID, key].
  - All traffic is sent using that IPv6 DA.
- Upon traffic reception:
  - CPE receives customer traffic, encapsulates it and adds as outer IPv6 DA the address A:88:1::AAAA:BBBB:CCCC:DDDD as computed previously.
  - Packet is routed up to a PE router, that upon reception, checks whether there is a valid FIB entry for that SID.
  - That FIB entry points to a particular VRF. Within this VRF the uRPF check is performed for the CPE Source Address.
  - The SRv6 BSID is expanded into the actual SID list. The lookup to resolve the first SID is done on that VRF as well.

Solution A, defined above, has several benefits. For example, this solution utilizes standard/regular router operations (e.g., plain VRF and plain FIB lookup) and standard router behavior. This solution does not require crypto hardware now no hashing in the dataplane, but provides the same level of security as HMAC. HMAC additionally checks the full segment list, but in the SDWAN context this does not provide any benefit as there is only a single segment and the SRH is not present. If, in the above example, customer 2 finds out the SID and the shared secret of the SD-WAN service of provider A, this information is insufficient to steal the service. This solution also protects from DDoS attacks. In particular, if the IPv6 DA does not contain the hash, then packet will be dropped in the FIB lookup. If the packet source is incorrect, then the packet will fail the uRPF check on the VRF. As such, there is no bottleneck on hash computation and DDoS traffic will be dropped at the FIB.

However, if the source address is spoofed, this Solution A does not protect Customer 2 from sending traffic to the services of Customer 1 (e.g., Customer 1 would be billed incorrectly). In addition, anti-replay protection is not provided.

Solution B: FIB-based Security Mechanism with Anti-Replay Protection

As per CAIDA studies, 12% of IPv6 blocks are still spoofable (e.g., due to lack of implementation of BCP38 or RFC6092). In such cases, the mechanism described in Solution A still has a threat to possible spoofing attacks, or alternatively replay attacks that might cause incorrect overbilling to Customer 1 in the above example. As such, Solution B extends the security mechanism with a "loose" anti-replay protection mechanism. This mechanism leverages the tight timing synchronization requirements that have been introduced in the past few years due to 5G, and it its capable of providing anti-replay protection on a one-to-any basis.

In particular, Solution B modifies the above Solution A as follows:

- All the PEs and CPEs have their time synchronized (sync-margin within micro-seconds).
- The CPE will have an associated /64 prefix.
    - Every 10 seconds, a new IPv6 SA will be generated.
    - This address is generated by concatenating the IPv6 prefix with the hash on [IPv6_SA_Prefix, SD-WAN Service SID, key, time t] trimmed to the most-significant 64b.
    - The egress packets will be sent towards the anycast address of the SD-WAN service using as IPv6 Source Address the address computed in that period of time t.

- Every PE node of provider A will have a FIB entry for the SD-WAN service instantiated on the anycast SRv6 SID A:88:1::AAAA:BBBB:CCCC:DDDD, where:
    - AAAA:BBBB:CCCC:DDDD is the result of the hash of [IPv6_SA_Prefix, SID, key, time *t*] trimmed to 64b.
    - Every 10 seconds, the FIB entry is updated replacing the old address A:88:1::AAAA:BBBB:CCCC:DDDD for A:88:1::FFFF:GGGG:HHHH:IIII, whereas FFFF:GGGG:HHHH:IIII is the result of a new hash computation with the actual time (*t+10*).
    - This FIB entry is linked with a particular VRF.
        - This VRF is used for the uRPF check, and it contains the IPv6 Source Address for each possible CPE of Customer 1. Note that the address is computed at the PE node independently and it is the source address that the CPE is allowed use on that particular timeslot t.
- Upon setting up the service:
    - The SDWAN controller creates a single security association in between the all the PEs and CPEs that establish the base-time and the frequency (e.g. every 10 seconds) of IPv6 address rollover
    - Every 10 seconds, both the CPEs and the PEs recompute the IPv6 CPE SA and SD-WAN service SID addressing based on the new time-slot
    - Upon reception of customer traffic, CPE steers the traffic into the updated SDWAN SID and forwards
    - Additionally, both the CPEs and PEs maintain an additional FIB entry for the previous time-slot during a fraction of the rollover time to avoid dropping packets during transition periods.

Solution B, defined above, has several benefits. For example, this solution provides a secured SD-WAN service with protection against over-billing of service and a “loose" anti-replay protection mechanism that neither relies on counters, nor has state per (S, D) pairs. In addition, the dataplane does not require hashing on a per-packet basis and no crypto hardware required.

With this solution, complete replay protection is not afforded. However, the state is traded off at PEs and CPEs for a small window in which packets may be replayed. If the rollover period is made sufficiently short, attackers will be forced to expend a non-trivial amount of processing in order to perform a successful attack. For example, the hackers would need to deploy an automatic tool to capture packets, and then resend them to the group as if they were from the originator. It is also noted that, if it is determined that one particular SD-WAN service (SID) is being targeted, it is possible to increase the rollover time for that single SID until it is made small enough (under 0.5s) to make the attack unaffordable.