IPv6 first hop security in our cloud environment

David Freedman – Claranet – IPv6 Security Workshop – July 2017

claranet

hosting | applications | networks

Some background - 2011

- In 2011 we did launched a cloud computing product.
  - We called it VDC (Virtual Datacentre).
- At the time, customers traditionally hosted services on their own hardware.
  - Either on their own premises or in our Datacentres.
- Customer hardware usually enterprise names / brands.
  - We had to build trust in a nascent market.
  - This meant anything we built had to be based on the same names / brands.
- However, the flesh was willing, but the spirit was weak
  - No decent orchestration or management software, portals etc..
  - Big vendors full of ideas, but no solutions to the problems.
  - Eventually, we found a supplier to work with for this software.
  - But we largely had to develop the networking ourselves.

IPv4 model, SIAs and DIAs

- **VM configuration and provisioning workflow:**
  - **Pick an flavour of VM / Configuration.**
  - **Add some networking to it, in the form of a vNIC.**
  - **Three flavours of vNIC available:**
    - Public (which we internally call SIA)
    - External (which we internally call DIA)
    - Private
  - **Public vNICs share a broadcast domain.**
  - **External and Private vNICs have a dedicated routing domain.**
    - Private is completely private to the customer living entirely inside the virtualisation domain.
    - DIA exists on the physical network and can be joined to other things and services.
    - SIA has a shared routing domain.
  - **Machines encouraged to request address via long lived DHCP**
    - Only mandatory in SIA.
  - **Custom DHCP server serves state from provisioning DB, does not use leases.**

What is SIA?

- **SIA enables you to obtain an address quickly.**
  - Pick from a pool, your VM can have a public address directly attached.
  - No NAT (unless your VM is a NAT box of course).
- **SIA is a shared routing domain.**
  - It is also a shared broadcast domain (in theory).
  - We don't segment customers any more than we would filter them from each-other.
  - We do however have an FHS security model.
  - Flows which do not meet the security criteria are dropped.

SIA IPv4 FHS Security Model
(non-exhaustive)

- Only permit IPv4 and ARP EtherTypes.
- Only permit source MACs you own.
- Only permit destination MACs in-domain.
- Only authorised DHCP servers on the LAN.
- Only permit ARP replies for your DHCP address.
- Only permit IP source address you are assigned.

SIA IPv4 FHS Security Model

(non-exhaustive)

- **Only permit IPv4 and ARP EtherTypes.**
- **Only permit source MACs you own.**
- **Only permit destination MACs in-domain.**
- **Only authorised DHCP servers on the LAN.**
- **Only permit ARP replies for your DHCP address.**
- **Only permit IP source address you are assigned.**

Dynamic MAC filtering (first three items)

Vendor Feature : DHCP Snooping (DHCP servers)

Vendor Feature : Authorized ARP (ARP replies)

Vendor Feature : IP Source Guard (IP source address)

SIA IPv6 Implementation

- **Same concept required – shared broadcast domain.**
- **IPv6 SIA a /64 with stateful DHCPv6 service + delegation.**
- **Threat model is therefore:**
  - **Attacks on neighbor discovery (control plane)**
    - Unauthorised neighborships / poisoning.
    - ND cache exhaustion..
  - **Attacks on router advertisement (control plane)**
    - Unauthorised router advertisements.
  - **Spoofing (forwarding plane).**
    - Unauthorised source addresses and prefixes.

Start with DHCPv6

- **Host steered toward DHCPv6 service via RA managed config.**
  - DHCPv6 has does appear on-link (though, really not – conceived pre-RFC6939)
  - Set M flag, clear A flag (important).
  - DHCPv6 “lease” reflects their provisioned address.
  - Delegation made if the host is to be a router.
- **Record of address and delegation added to bindings DB.**
  - Vendor calls this ‘Glean’
  - Glean can be applied to RA/ND (stateless) and/or RA/DHCPv6 (stateful)
  - Configured in stateful mode.
- **Prevent any rogue DHCPv6 servers.**
  - Vendor calls this ‘DHCPv6 Guard’
  - Just like IPv4 counterpart, blocks unauthorised DHCPv6 replies.
  - Susceptible to evasion scenarios (need additional mitigation).

Neighbor Discovery / RA

- **Bindings are used to validate further ND**
  - **Vendor calls this “ND Inspection”**
  - **Invalid ND packets are dropped before doing anything else.**
  - **NA assertions validated against bindings DB**
    - Validates neighbor address, bound MAC and source MAC.
    - This mitigates against NA spoofing and poisoning.
- **Router Advertisements also validated**
  - **Industry & Vendor call this “RA Guard” (RFC6105)**
  - **Block router advertisements from unauthorised sources.**
  - **Attempt to mitigate evasion scenarios listed in RFC7113.**
- **ND rate limited**
  - **Vendor calls these “ND Cache Interface Limit” & “ND Resolution Rate Limit ”**
  - **Queued requests and interface cache size limited.**

Source and Destination Validated

- **IPv6 Global sources validated against bindings DB**
  - Vendor calls this both“IPv6 Source Guard” and “IPv6 Prefix Guard”
  - Link local allowed, but global auto-configured address not.
  - This mitigates against global source spoofing.
  - uRPF also enabled upstream for protection.
- **IPv6 Global destinations validated against bindings DB**
  - Vendor calls this “IPv6 Destination Guard”
  - This mitigates against destination spraying / cache exhaustion.

Problems

- **Bugs – Lots of bugs**
  - Features not working.
  - Memory leaks.
  - Overzealous defaults.
- **Keeping state**
  - ND Bindings needs to be backed up.
  - If you have multiple units, you have to be ‘creative’
- **Emergencies**
  - Loss of bindings or corruptions of state, again, you have to be ’creative’
  - Amnesty scripts.

Was it all worth it?

- **Mostly vendor features, but some already standardised.**
  - Most of them didn't appear in mature code until recently.
- **Alternative was to forward traffic via hypervisor.**
  - Potentially not having shared broadcast domain.
  - However, needs of home-grown code, supporting forwarding security.
- **Another alternative, don't offer SIA**
  - Public cloud providers are using similar model, why would we opt out?
  - Customers require this level of flexibility.

Any questions?