Trusted Domain SRv6

draft-raviolli-intarea-trusted-domain-srv6-02

INTAREA WG meeting, IETF 118

Prague, 8th November 2023

Why does this exist?

- SRv6 has known security vulnerabilities should an attacker be able to insert packets into the SRv6 domain
- RFC8402 section 8 clearly states that SRv6 that leaks beyond the boundaries of the trusted domain creates a security violation
- RFC8754 also states, in section 5, that to mitigate attacks filtering into the domain is required
- LPM filtering on every “external” facing port on a broad scale is neither scalable nor practical in an operational sense
  - Competently filtering SRv6 frames leaked from/between downstream transit customers – where they have IPv6 transit – is currently impossible

Attack Vectors

- Some of the attack vectors are documented at:
  - https://mailarchive.ietf.org/arch/msg/v6ops/GbWiie-bjQ_Bp1JKB1PlDh_fPdc/
  - https://datatracker.ietf.org/doc/draft-li-spring-srv6-security-consideration/

What the draft aims to accomplish

- Provide operators the ability to protect themselves and/or their customers, by creating a deterministic means by which to distinguish SRv6 from IPv6
- Avoid any re-work of SRv6 itself; all functionality needs to be maintained, while providing the ability to operate an SRv6 network in a “fail closed” manner
- Avoid creating any additional work that would require changes to silicon
- Enhance the deployment appeal of SRv6 by allowing operators who will not deploy SRv6 because of the stated security concerns, a solution to mitigate those shortcomings.
- Maintain the ability to run SRv6 without this mechanism for those who choose to run without the imposition of an EtherType.

How we will accomplish this

- A global knob on a device that tells the device to run in “Trusted Domain SRv6” mode
    - If enabled, no SID processing will occur on packets that do not contain an SRv6 Trusted Domain EtherType
- A per-interface knob to enable processing of the SRv6 Trusted Domain EtherType
    - This is disabled by default, and unless enabled, packets containing this EtherType will be dropped on ingress
- The per interface method replicates the same method used to secure MPLS – on almost all devices MPLS requires explicit enablement on MPLS capable ports

Another note

- It may be possible to impose/enforce the SRv6 Trusted Domain EtherType on the ”border” interfaces, and then revert to standard IPv6 EtherTypes on the inside of the network
- While not documented in the draft, this approach has its own set of security concerns and is not an approach we would recommend, thus choosing to explicitly rule designs of this nature as being out of scope of this document.

Changes to the draft since -01

- Language updates for clarity, following feedback from IETF 117
- The addition of an ‘Applicability Considerations’ section
  - Describes the concerns of the authors in the context of an underlay network carrying multiple IPv6 customer networks transited via overlays
  - Better represents within the draft the security concerns of service providers considering SRv6

WG adoption

- We are grateful of the feedback provided by the community during the last INTAREA WG meeting
- It is the belief of the authors that the draft is in a good shape to be formally adopted by the WG, and would like to request such adoption

Thanks for listening.

Questions?