



Contribution ID: 327

Type: **not specified**

# SYN Proxy at Scale with BPF

*Wednesday, 15 November 2023 16:30 (30 minutes)*

SYN Cookie is a technique used to protect servers from malicious connection requests. Under SYN flood, the Linux TCP stack encodes the client information into the initial sequence number (ISN) of SYN+ACK, which is called SYN Cookie, and decodes that from ACK of 3WHS so that the kernel can release resources for the connection and stays stateless during 3WHS.

For security reasons, SYN Cookie is calculated with some host-specific secrets, so can only the generator validate the cookie. Even with SYN Cookie, intermediate nodes between the client and the server must keep tracking the connection. SYN Cookie does consume resources and thus is NOT stateless in the network.

SYN Proxy reduces such unwanted resource allocation by handling 3WHS at the edge network. After 3WHS with a client, SYN Proxy generally restores the initial SYN packet from SYN Cookie and forwards it to the backend server to initiate 3WHS. However, the ISN in SYN+ACK is selected randomly by the server and does not match the SYN Cookie. To be transparent to the client and the server, SYN Proxy must keep the ISN mapping and fix the SEQ/ACK numbers in all packets. This solution also is not stateless and does not scale well for a service with high bandwidth.

This talk will cover

- what the kernel encodes into SYN Cookie
- our stateless SYN Proxy and kernel module
- ongoing effort to add a new kfunc to replace the module

**Primary author:** IWASHIMA, Kuniyuki (Amazon Web Services)

**Presenter:** IWASHIMA, Kuniyuki (Amazon Web Services)

**Session Classification:** eBPF & Networking

**Track Classification:** eBPF & Networking Track