


Organizations are not protecting their BGP session. Take the time to ask the question Do we have our BGP ports protected? Are you:

1. Worked with your BGP peers to deploy the General Purpose TTL Mechanism (GTSM).
2. Deploy Control Plane ACLs on the Router to ONLY permit the BGP session from your peer (explicit permit algorithm).
3. Deploy Infrastructure ACLs (iACL) on the edge of to protect your network devices.
4. Configured the iACLs on the edge of your network to include anti-spoof. Internet “miscreants” can spoof your IPs and attack your devices (getting through ACLs).
5. Worked with your BGP peers to deploy TCP Authentication Option (TCP-AO).

If not, work with your peers to deploy an Infrastructure ACL (iACL) to cover all your network devices, deploy specific data plane ACLs on your routers/switches to protect them, work with your vendor to deploy their “Control Plane Protection,” work with your peers to configure General Purpose TTL Security Mechanism (GTSM) on all you BGP peering sessions, and set up appropriate monitoring so you know when people are trying to DDoS you port 179. This approach establishes a layered BGP session defense for each router/switch.



None of these tasks requires any CAPEX. Your network devices can do all of this TODAY. Here is a step-by-step guide to help your organization architect BGP Session Resiliency.

RECOMMENDATION! Read through RFC 7454 – BGP Operations and Security.

Step 1 – Set up an Infrastructure ACL (iACL) on your Network to protect your Routing Control Plane.

Infrastructure ACLs (iACLs) are essential to protect your network! These ACLs are Explicit Permit ACLs. That means EVERYTHING is DENIED by DEFAULT. You then explicitly permit the IPv4/IPv6 blocks, protocols, source IP/port, destination IP/port, and other fields.


Routes had problems with large ACLs that checked all packets in the past. This is no longer the case. Most of the large broadband providers in the US deploy Exploitable Port Filters that combine the iACL with port filtering to protect their customers (see [Filtering Exploitable Ports and Minimizing Risk from the Internet and from Your Customers](#))

BGP protection happens with an iACL rule that blocks BGP access to your router's control and management plane and EXPLICITLY PERMITTING the BGP speakers who need to "BGP Peer" with your router.

How would you configure this policy? It would be specific to your network and topology. There are many paths to success. Explore the options, talk to your peers, consult with your vendors, and WRITE DOWN THE POLICY!

Step 2 – Set up an Infrastructure ACL on your Router/Switch to Protect your Route Processor

Remember, DDoS Threat Actors can get inside your network. They can reflect off devices or use violated devices to attack you from within. "Defence in Depth" does not work in a world where you cannot trust anything. You need "Dipping Dot" security where everything is distrustful, and you put up protections. The network device's iACL would protect the router from a port reflection attack inside the network.



Yes, “zero trust” extends inside the network. The router’s iACL would permit all the data plane traffic but have explicit ACLs to protect the control and management plane of the router.

Infrastructure ACLs (iACL) Deployment Guidelines

There will be pushback from people saying, “My router cannot support ACLs on the Internet edge.” That is NO EXCUSE in today’s Internet. The attitude’s consequence opens your network undefended, with your routers open to DDoS attacks.

Given this, here are guidelines for deploying iACLs:

- iACL would be set up on the network's Autonomous System (ASN) edge.
- Anti-Spoof rules will be in the iACL. These will block any IP source that equals your IPv4 and IPv6 blocks. A miscreant on the Internet should not be able to send a packet whose source matches your IP blocks.
- Consider blocking destination IPs to network devices, interfaces, and core infrastructure. This would have a “first layer” of packet filtering protecting routers, switches, interfaces, and other critical network infrastructure.
- Explicit Permission for outside routers & services to connect to your network devices. BGP Peering would be one example.
- Consider [Filtering Exploitable Ports and Minimizing Risk from the Internet and from Your Customers](#). Filtering exploitable ports on the edge of the ASN has protected major operators like Cox, Comcast, Charter, Rogers, Bell Canada, and many other large Operators.

[Slide on Operators]

Step 3 – Control Plane Protection

Many routers have some type of control plan protection. Specialized “Control Plane” ACLs and rate limiters inside the router apply policy before the packet arrives on the BGP process. It would be worth reading [RFC 6192 Protecting the Router Control Plane](#). This is a peer-reviewed RFC providing collected practices for control plane protection.

Work with your vendor to explore these options. Here are a couple of examples:

- Cisco: [Control Plane Policing](#)
- Juniper: [Configuring Control Plane DDoS Protection](#)
- NVIDIA: [Control Plane Policing](#)

Step 4 – General Purpose TTL Security Mechanism (GTSM)

Most BGP sessions are one hop away. If the session is one-hop way, then the Time To Live (TTL) of the packet would be between 254 – 250 (depending on the TCP checks). A simple security approach would be to check the TTL of the packet. If it is between 255 and 250, then the packet is safer. If the TTL of the packet is 180, then it is many hops away and most likely a bad packet.

GTSM is a technique configured on the routers to use this “predictable TTL” as another security tool.

GTSM by itself will not stop DDoS Flood from reaching your route processor. But, if you know the TTL Range of the BGP sessions, your vendor’s ACL might allow you to put the TTL range into the iACL or Control Plane ACL.

Step 5 – TCP AO (replacing BGP MD5)

TCP Authentication Option (TCP-AO) replaces the BGM MD5 feature. RFC 5925 ‘The TCP Authentication Option would now be used between BGP peers. TCP AO protects very low-volume TCP state attacks that would drop the session.

Most BGP session DDoS attacks are volumetric attacks that focus on saturating the control plane to the BGP session. However, BGP Session RST attacks have happened in

the past. GTSM + TCP-AO, along with layers of ACLs, increases the router and BGP session's resilience to these direct port 179 attacks.

NOTE: Have a conversation with your vendors. TCP-AO's code is not rolled out consistently with vendors. As people say, "TCP-AO is a feature now in any RFP/RFQs we send to the usual gang of \$vendors."

- Juniper: [TCP Authentication Option \(TCP-AO\)](#). Presentation: [The TCP Authentication Option - \(TCP-AO\)](#) By Melchior Aelmans
- Cisco: [BGP Support for TCP Authentication Option](#)
- Nokia: [TCP enhanced authentication option](#)
- Huawei: [TCP-AO support for BGP authentication configuration](#)
- Mikrotik: Sorry - there is ZERO support for TCP-AO on Mikrotik.
- FFRouting: [Feature Request: TCP-AO \(RFC 5925\) #7240](#)
- APNIC Blog: [It's time to replace MD5 with TCP-AO](#) by Melchior Aelmans on 28 Jul 2021
- [Production Deployment of TCP Authentication Option](#) by [Andrew Gallo](#) – 27 Sep 2022 Contributors: Ron Bonica, Melchior Aelmans

Take a moment to watch these videos:

- [SwiNOG#38 | TCP Authentication Option: How it works and where it can be used | Remi Locherer | Arista](#) to get an excellent TCP-AO update.
- Packet Pushers: [BGP Authentication? Use TCP/AO Instead of MD5! With Melchior Aelmans](#)

Finally – Monitor for Session Attacks

Monitor the layers of defense. You need to alarm and investigate if an attacker is trying to hit your BGP Session. BGP Session attacks indicate an experienced and sophisticated attacker who has done their homework. When you see this, investigate. BGP might have been protected, but the threat actor might be trying something else.

How to monitor for BGP Session attack attempts? The options depend on your network telemetry architecture. It could be as simple as an ACL log, Netflow/IPFIX, or one of your security tools.

Extra – Don't Use the Physical Interface as the BGP IPs

The most resilient networks will ensure their BGP Peering subnets are not advertised on protected physical links or architected in a way where they cannot be externally targeted. For example, putting your BGP session on a separate IP block from the connected interface is one way to make it harder for a threat actor to find what to target.

There are multiple approaches that all need to match your network topology, your peer's topology, and your vendor's capabilities. This approach works, but needs to be architected into your BGP peering.

What about

What about [BFD \(Bidirectional Forwarding Detection RFC5880\)](#), BGP Hold Timers, and [Route Flap Damping \(RFC2439\)](#)? These are BGP Resiliency Widgets. They are widgets, sprockets, and cogs that must be part of your BGP Resiliency architecture. These do have value. They are used. But, they are architected into a BGP resiliency architecture. Seek conversations with your peers, [the operations community](#), and other operators about what they do.