

Re: TCP RST attack

Source: <http://www.derkeiler.com/Mailing-Lists/FreeBSD-Security/2004-04/0104.html>

From: Matthew Dillon (dillon_at_apollo.backplane.com)

Date: 04/20/04

Date: Tue, 20 Apr 2004 13:45:20 -0700 (PDT)

To: Charles Swiger <cswiger@mac.com>

Well, the advisory certainly exaggerates some, but I think there is a real issue with BGP and I have to respectfully disagree with DES. Route flapping cannot be solved by reducing hysteresis, at least not unless you want your backbone provider to cut you off! Flapping is a major problem... less of one now than when I was doing BEST a decade ago, but still very serious. When a BGP session flaps it has to resynchronize, and resynchronization can take a significant period of time, bandwidth, and router resources to accomplish. You can't just reconnect and pick up where you left off (if you could it would be a non–problem).

On the other hand, BGP can be trivially protected. You don't need ingress or egress filtering at all (by which I mean IP block filtering), you simply disable the routing of any packet to or from port 179. 99.9% of all BGP links are direct connections (meaning that they terminate at a router rather than pass through one). No packet to or from port 179 has any business being routed from one network to another in virtually all BGP link setups so the fix is utterly trivial.

–Matt

freebsd-security@freebsd.org mailing list

<http://lists.freebsd.org/mailman/listinfo/freebsd-security>

To unsubscribe, send any mail to "freebsd-security-unsubscribe@freebsd.org"