BGP Testing:
Why Be So Negative?

Brent Imhoff
Scott Poretsky

NANOG 30
February 10, 2004

True in 1996. True Today.

North American Network Operators Group

Re: "Basic BGP configuration problem"

From: Henry Kilmer

Date: Tue Oct 01 15:51:51 ***1996***

No one is immune to bugs in code. -Hank

North American Network Operators Group

Re: "Basic BGP configuration problem"

From: Perry E. Metzger

Date: Tue Oct 01 16:13:55 ***1996***

Henry Kilmer writes:

>Yes. And it wasn't the configs that were wrong. It was a BGP related
{Vendor name removed} bug.

That's the best way to find these sorts of things -- in test...

Perry

Negative BGP Conditions in Networks

- BGP Update message errors – Software crash may occur if improperly handled. NLRIs may be incorrectly advertised.
- BGP Route Explosion – Out-of-Memory condition may occur, which could force router reboot
- BGP Processing – High CPU Utilization may persist on router, which could cause BGP sessions to flap and lead to more network instability
- Incorrect BGP Path Selection Process – Could produce routing loop or route oscillation
- Stuck Routes – Could produce routing loop

Negative Testing

- Negative Testing=

  Tests designed to verify that the router under test correctly responds to error conditions in the network

- Negative Testing different from functionality, conformance, and interoperability testing
  - These verify correct operation with known expected behavior

- Negative Testing more difficult to define because number of error conditions is boundless

Great Question!

North American Network Operators Group

BGP testing?

From: Timothy Brown

Date: Fri Nov 17 19:20:01 2000

Hey folks, Does anyone have a script or a series thereof to do large-scale BGP testing? I'm looking for scripts that will generate and nail down several hundred networks of varying sizes, and/or fake peering relationships with a similar purpose, and/or do things that don't meet the BGP protocol standards, etc. Thanks for any responses. Tim

IETF BMWG Routing Benchmarking

- BMWG Benchmarks single device Performance, not Conformance and not Negative Testing
- Current Routing Benchmarks cover FIB Scaling, Forwarding Performance, and Convergence that are fundamentals of Negative Testing:
  - Terminology for Forwarding Information Base (FIB) based Router Performance (RFC 3222)
  - Terminology for Benchmarking BGP Device Convergence in the Control Plane (draft-ietf-bmwg-conterm-05.txt)
  - Benchmarking Terminology and Methodology for IGP Data Plane Route Convergence (draft-ietf-bmwg-igp-dataplane-conv-term [and meth]-02.txt)
- http://www.ietf.org/html.charters/bmwg-charter.html

Negative Testing of BGP UPDATE Errors

- AS Path List
  - Routing Loop
  - Maximum AS Path Length
- Missing/Incorrect/Errored Attributes
  - EBGP, IBGP, RR, Confeds
- NLRI with incorrect next hop
  - Receiving router is next hop
  - NLRI is next-hop
  - Unreachable Next-Hop
- Interoperability negotiating capabilities
  - UPDATE message includes parameter that wasn't agreed upon

Negative Testing of BGP Resources

- BGP Route Explosion –
  - Force Out-of-Memory condition
  - Advertise NLRIs until no memory available
  - Advertise additional NLRIs and observed router behavior
- BGP Processing –
  - High CPU Utilization
  - Remove EBGP Peering Session from which most FIB routes learned
  - Force Convergence Event during a Convergence Event
- Redistribute BGP into IGP

BGP Convergence Tests

Test Setup

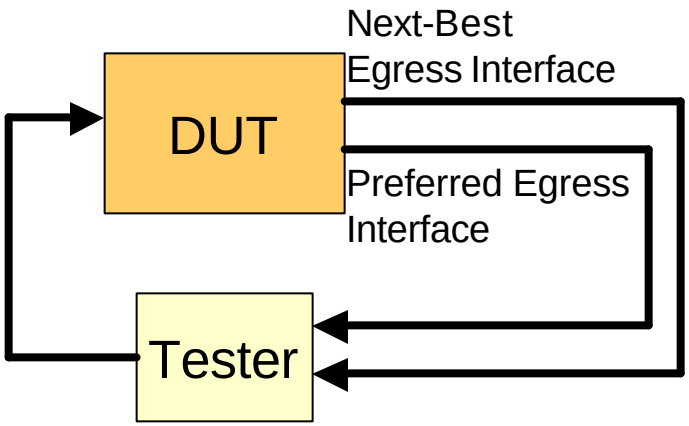

Convergence Events

Link Failure
- -Local Interface Failure
- -Neighbor Interface Failure
- -Remote Interface Failure

Layer 2 Failure (PPP, GigE)

IGP Adjacency Failure

Route Withdrawal

Cost Change

Test Procedure

- DUT has two paths (via Link 1 and Link 2) to reach destinations
- Tester sends traffic to DUT to all destinations in FIB
- DUT by default prefers lower cost path via Link 1
- Convergence Event reroutes traffic to Link 2
- Observe recovery to forwarding at line rate and Calculate Convergence Time

Negative Testing of BGP Functionality

- Path Decision Process
  - advertise same NLRI from multiple neighbors
  - cause each step of the decision process to be used
  - Use different IGP next-hops
  - force execution of entire Process
- Invalid peering
  - MD-5 Authentication with invalid password/keys
  - Incorrect AS number

Negative Testing of BGP Configuration

- BGP stability/scaling testing can be impacted by the configuration of the following:
- Hello/Keepalive Timer
  - Some routers exhibit degraded behavior when using a setting of 30 seconds / 90 seconds instead of the default 60 seconds / 180 seconds
- Update Rate
  - Some routers exhibit degraded behavior when increasing to 5K NLRIs/second or higher
  - Routers tend to be stable at 2K
- Peer-Groups
  - Use of Peer-Groups can improve memory utilization to increase the number of peers and routes

Test Tools for BGP Negative Testing

- Routing Protocol emulation from commercial router test equipment has become very advanced in past year:
  - Ability to emulate iBGP or eBGP
  - Ability to emulate IGP on same test port
  - User Configuration of Update Rate and Timers
  - Feature to load external route table for building Update messages
  - These tests may be scripted
  - Test Vendors: Need canned script for Selection Process!
- Freeware emulators can be extended via scripts to generate negative conditions:
  - Python Routing Toolkit from Sprint labs
  - bgpsim, which is part of MRT from Merit, is available at http://www.sourceforge.com

Summary

- Negative Testing is a critical component of router evaluation prior to deployment.
- Some BGP implementations have caused network instability because of lack of Negative Testing
- BGP Negative Testing should drive High Memory and CPU Utilization
- IETF BMWG is addressing FIB Scaling and Convergence
- Additional BGP Negative Testing should include:
  - BGP Update Messages
  - Route Convergence
  - Path Selection Process
  - Peering
  - Configuration
- Sophisticated Test Tools are commercially available today perform this negative testing

Comments?