NANOG-BCOP

NANOG 63 – 2 February 2015 – San Antonio, TX

What’s a BCOP?

Best Current Operational Practice

- A living document describing the best operational practices currently agreed on by subject matter experts
  - Vetted and periodically reviewed by the global network engineering community (GNEC)

The Problem

- Operational knowledge tends to be "tribal"
  - Presentations, hallway conversations, internal documents, in someone's head…
  - Technology, tools, and practices change over time…
- There are hundreds of operational forums globally
  - Archives stored in different formats, some searchable, rarely have speech text or video, no vetting, and state unknown.
- How do I find up-to-date, relevant information when I need it?

The BCOP Solution

- Open, Transparent, Bottom-up, and Community led
  - Community driven, community written, community vetted Best Current Operational Practices from an open forum, list, and publicly searchable site.
  - Community written and approved Development Process for BCOPs
  - Everyone is welcome to participate
- 80/20 model

Today’s Agenda

- Global BCOP Update
- Ethernet OAM BCOP
- DDoS/DoS Attack BCOP
- eBGP Configuration BCOP
- Public Peering Exchange Participant BCOP
- Open Mic

BCOP activity around the world:

http://www.internetsociety.org/deploy360/about/bcop/

- Africa region: A BCOP group was started under AfNOG, lead by Douglas Onyango
- Asia: We expect the first BoF to be held at APRICOT 2015!
- Europe: RIPE BCOP Task Force created, co-chaired by Benno Overeider and Jan Žorž
- Latin America: A BCOP Task Force was started under LACNOG, lead by Luis Balbinot and Pedro R Torres Jr.
- North America: NANOG BCOP Committee established, co-chaired by Aaron Hughes and Chris Grundemann

AfNOG BCOP

- First introduced in May of 2013
- Held a BoF in Abidjan at AfriNIC19
- Most recent BoF at AIS/AfriNIC 20 in Djibouti (June 2014).
- Current focus:
- Put in place a mailing list
  - Using http://www.afnog.org/mailinglist.php for now
- Create an online BCOP document repository
- Development of two or more drafts
- A session at AFRINIC 21 in Mauritius (Nov 2014).

AfNOG BCOP documents in the works:

“IPv6 questions/answers cheat sheet specific to Africa”

Contributors: Alfred Arouna

- Aims to consolidate common questions and best answers in a kind of IPv6 questions/answers cheat sheet specific to Africa.

RIPE BCOP

- RIPE BCOP Task Force charter page:
- http://www.ripe.net/ripe/groups/tf/best-current-operational-practices-task-force
- Mailing-list:
- https://www.ripe.net/mailman/listinfo/bcop

RIPE BCOP documents in the works:

“IPv6 troubleshooting for residential helpdesks”

Contributors: Lee Howard, John Jason Brzozowski, David Freedman, Jason Fesler, Tim Chown, Sander Steffann, Chris Grundemann, Jen Linkova, Chris Tuska, Daniel Breuer, Jan Žorž

- Starting point for technical support staff at ISPs or enterprise IT helpdesks
- Addresses the “fear of the unknown” problem at many organizations
- Provides a solid first step for front-line support personnel.

RIPE BCOP documents in the works:

“DNSSEC operational practices for authoritative name servers”

Contributors: Matthijs Mekking

- Protocol default values

- + Cryptographical considerations?

+ ZSK/KSK split or CSK?
+ When to rollover?
+ Values for signature validities, re-sign, refresh, …
+ NSEC or NSEC3?
+ If NSEC3, when to resalt?

- Key management

- + Generation: Number of participants?

+ Delivery: Integrity checks? Audit trail?
+ Storage: Online or offline? HSM or not?
+ Usage: Who can use? How to (de)activate?

Available software

+ Standalone solutions: OpenDNSSEC, BIND, Knot, …
+ Combinations: ldnsutils + NSD, …
+ Closed source: Microsoft DNS, Nominum, ...

RIPE BCOP documents in the works:

“BGP Best Current Operational Practices”

Contributors: Pierre Lorinquer, Observatory Team (G. Valadon, M. Feuillet, F. Contat) and operators Association Kazar, France-IX, Jaguar Network, Neo Telecoms, Orange, RENATER, SFR

- Definitions:

Interconnection types

- Direct interconnection
- IXP Peering
- IXP Route-server
- Multihop

- *AS relationships*
- Transit / Customer (leaf)
- Transit / Small transit
- Peering

- Recommendations:

AS relationship dependent

- TCP-Authentication
- AS-PATH filtering
- Prefixes filtering (route objects)
- Max-prefix
- Private AS removing

- *General recommendations*
- Martians filtering
- Bogons filtering
- Default route filtering
- Log
- Graceful restart

LACNOG BCOP

- The group has asked for a webpage under the LACNOG umbrella.

- Mailing List: https://mail.lacnic.net/mailman/listinfo/bcop

- The group still has to decide on primary language of the produced documents (Spanish/Portuguese/English).

- They recently held a BoF at LACNOG 2014 / LACNIC 22 in Santiago, Chile (October 2014)

LACNOG BCOP documents in the works:

“LacNOG BCOP Development Process document”

Contributors: Pedro R. Torres Jr., Luis Balbinot

- A development process is important for capture the Best Current Operational Practices in documentation format that is uniform and easy to read.
- LacNOG BCOP TF decided to set the format and procedure first and then start capturing the Best Current Operational Practices into documents.

Ethernet OAM BCOP

Mark Calkins, NANOG63, San Antonio, 2.2.2015

BCOP Summary

- This BCOP aims to provide insight into how Ethernet OAM is best deployed within todays service provider networks.
- This BCOP will try to capture current and emerging best practices for uses of Ethernet OAM technologies.
- The primary focus of this BCOP is to de-mystify EOAM protocols and practices.

BCOP Background

- This BCOP is needed because Ethernet OAM is not widely understood outside of the service provider community.

Participants

- Shepherd: Mark Calkins
- Current SMEs:
  - Voitek Kozak
  - Jean-François Lévesque
  - Mark Calkins

Draft BCOP Outline

- General EOAM BCOPs
  - High level, What, When, Where
- Link Level EOAM BCOPs
  - How LFM functions, why it is good
  - Best practices for LFM's link monitoring
- Service Layer EOAM BCOPs
  - Service Layer OAM orientation
  - How CFM functions, why it is good
  - Fault Management
  - Performance Management

BCOP Content for Review

- All content is up for review

Next Steps

- Last Call for feedback
- Compile feedback received and modify document as required
- Wiki-tize document
- PDF document

Thank You

- Thank you to my SMEs
  - Voitek Kozak
  - Jean-François Lévesque
- Thank you to all who have offered input and feedback so far.

LAST CALL

- Last Call for feedback

- Get involved!
  - Flexible time commitment
  - Contribute as much or as little as you can
  - The more voices we can include the better
  - Email mark.calkins@gmail.com to be included

DDoS/DoS Attack BCOP

Yardiel Fuentes, NANOG 63 - San Antonio, TX

Participants

- Shepherd: Yardiel Fuentes
- Current SMEs:
  - Rich Compton
  - Prabhu Gurumurthy
  - Damon Fortune
  - John W
  - Yardiel Fuentes
- Other contributions from:
  - Link King

BCOP Summary

- This BCOP appeal shares practices which have performed in production environments as a guide on what to do before, during, and after a DDoS/DoS attack.
- This BCOP document focuses on providing, in a vendor-agnostic framework, guidelines at the different stage of dealing with DDoS/DoS attacks
- Check out http://bcop.nanog.org/index.php/BCOP_Drafts

BCOP Background

- This BCOP is needed because of the increase number and intensity of DDoS/DoS attacks NANOG engineers have to willing or unwillingly deal with.
- The need for practical info obtained from defending production networks
- Effort was triggered by multiple requests for information, the NANOG email list "frequent topics" and track recommendations

Draft BCOP Outline

- What to do prior to a DDoS/DoS attack:
  - Customized packet filtering (Firewall Filters, ACLs)
  - Disable open recursion for internal DNS with explicit external trusted DNS servers.
  - Identify and mitigate open relays (open resolver/ Shadowserver) including in your customer's networks
  - Consider DDoS/DoS mitigation subscription services
  - Be familiar with your network's normal session traffic patterns
  - Adhere to BCP38, BCP84 and RFC2827 tips like RPF

Draft BCOP Outline

- What to do during a DDoS/DoS attack:
  - Implement passive collection of data (Jflow, net flow, port mirroring, taps)
  - Establish baselines of normal traffic as reference for identifying DDoS/DoS flood traffic
  - Deploy DoS mitigation solutions and additional options (including Flowspec) for NLRI rate-limiting
  - Seek to identify the Type of DDoS attack (TCP win 0, Synfloods, single-source UDP/LOIC-based, botnet, etc)

Draft BCOP Outline

- What to do after a DDoS/DoS attack:
  - Identify the type of DDoS/DoS attack via collected data analysis. Port-mortem analysis is useful
  - Check if FFs/ACLs could be fine tuned to prevent re-occurrences
  - (controversial point) If decision is to notify authorities, the IC3 and FBI official stand is that they are interested in being informed about these Attacks — Particularly, after the Sony Attack and the FCC official interest on these.
  - Share your DDoS/DoS attack details with NANOG community (such as this BCOP)

“Nice to have” Content

- What sensible approaches or steps in addressing DoS/DDoS attacks are not as effective as initially considered ?

- Have you found legal or law enforcement agencies helpful at all ? — names/details welcome…

- Should an FAQ or brief educational session be included in this BCOP ?

- Comments Welcome — email me at yardiel@gmail.com

Join Us!

- Are you interested in DDoS/DoS Attacks?
- Do you have real-world experience with DoS Attacks?
- Are you interested in helping the NANOG community address DoS attacks ?
- Get involved!
  - Flexible time commitment (all interactions done via email)
  - Contribute as much or as little as you can
  - The more voices we can include the better
  - Email yardiel@gmail.com to join this BCOP team

eBGP Configuration BCOP

Bill Armstrong, NANOG 63, San Antonio, TX, 02-02-2015
The Operators cry, "Ki yippee yi!"

BCOP Summary

- This BCOP aims to provide a singular, consistent view of industry standard eBGP interconnection methodologies
- This BCOP will also document pre and post turn-up validation practices and IRR Etiquette
- The primary focus of this BCOP is eBGP KNOW-HOW

BCOP Background

- Although eBGP peering sessions are turned up everyday the one you turn up tomorrow could be the other guy's first. This BCOP is needed to make sure the other guy knows what to expect.
- The creation of this BCOP was prompted after reading through a sordid 6 day cut-over that played out on the NANOG List
  - *Despite best laid plans by the OP the remote Peer was unable to stay up*
  - *Common expectations between peers were not set*
  - *The final resolution was only a max-prefix adjustment away*
- Doing things inconsistently CONSUMES TIME
  - No peering session should take 6 days to come up
  - No one should have to play Russian roulette when it comes to something as fundamental as a Peering turn up.

Draft BCOP Outline

1 BCOP Summary (Appeal)
2 BCOP Background / History
3 BCOP
3.1 What is BGP
3.1.1 Who needs BGP
3.1.2 Internal BGP (ibgp)
3.1.3 External BGP (ebgp)
3.1.4 Route Advertisement in IBGP vs EBGP
3.1.4.1 IBGP
3.1.4.2 EBGP
3.1.5 Loop avoidance mechanism in IBGP vs EBGP
3.1.5.1 IBGP
3.1.5.2 EBGP
3.1.6 BGP best Path selection refresher
3.1.6.1 Juniper
3.1.6.2 Cisco
3.2 Pre-turn-up considerations
3.2.1 Relationship Types
3.2.1.1 Transit
3.2.1.2 Peering
3.2.2 Interconnection Types
3.2.2.1 Point to Point Interface peering
3.2.2.2 eBGP Multi-hop peering
3.2.2.3 Load sharing
3.3 Policy Considerations
3.3.1 Inbound policy classification approach.
3.3.2 Inbound policy definitions and examples.
??!?!??!!?WHAT ABOUT OUTBOUND!!?!??!?
3.3.2.1 Transit inbound filters both for IPv4 and IPv6
3.3.2.2 Transit inbound IPv4 filters
3.3.3 Transit inbound IPv6 filters
3.3.3.1 Communities
3.3.3.1.1 Downstream Communities
3.3.3.1.2 Transit Communities
3.3.4 IRR\PeeringDB
3.3.4.1 IRR basic building blocks
3.3.4.1.1 Maintainer Object
3.3.4.1.2 Route/Route6 Object
3.3.4.1.3 AS-set object
3.3.5 Getting Started
3.3.5.1 Choose a IRR database
3.3.5.1.1 Gather information about our network
3.3.5.1.2 Create maintainer object
3.3.5.1.3 Create route/route6 objects
3.3.5.1.4 Create a as-set object
3.3.5.1.5 Notify your peers
3.4 Turning up eBGP Peering
3.4.1 Testing and Validation - NEED MORE INFO/ Removed?
4 BCOP Conclusion

Participants

- Shepherd: Bill Armstrong
- Current SMEs:
  - Alex Saroyan
  - Brian Schleeper
  - Courtney Smith
  - Mannan Venkatesan
  - Raghav Bhargava
  - Russell Harrison
  - Scott Dalton
  - Scott Morris
  - Umair Arshad
- Other contributions from:
  - Karsten Thomann

Current Status

- The Draft is posted:
    - http://http://bcop.nanog.org/index.php/EBGP_Configuration_BCOP_v0.1
- We have a decent amount of information but we NEED SOME HELP!
    - Baltimore got some new participants BUT
        - We have only enriched what we have
        - The "Testing and Validation" section STILL has very little in there and is perhaps the MOST critical portion of the document
    - When the draft was posted questions surrounding PE to CE Outbound prefix filtering came up but we don't have any good answers…Yet? [LOOK AT AUDIENCE WANTINGLY]

Join Us!

- Are you an expert in eBGP Configuration and Testing?
- Do you have real-world experience with eBGP Policy ‘Stuff’?
- Are you interested in eBGP Configuration best Practices?
- Get involved!
  - Flexible time commitment
  - Contribute as much or as little as you can
  - The more voices we can include the better
  - PLEASE Email wrarmstrong@gmail.com to be included

Public Peering Exchange Participant BCOP

Shawn Hsiao, NANOG 63, San Antonio TX, 2.2.2015

BCOP Summary

- This BCOP aims to update current “Public Peering Exchange" BCOP
  - Add IXP prefix handling advice
  - Remove information pertaining to the operation of an exchange into a separate document, and re-focus the document toward exchange participants
  - Other updates as needed

BCOP Background

- From a discussion thread in 01-15-2014 regarding handling of IXP prefixes, there are several approaches discussed and different opinions raised. The update to BCOP aims to document and analyze these approaches, and make recommendations

- There are also other topics that would be beneficial for the participants
  - Also sharing and cross-reference contents from eBGP Configuration BCOP

Participants

- Shepherd: Shawn Hsiao
- Current SMEs:

Draft BCOP Outline

- Other considerations, e.g,
  - Check IXP policy on prefix distribution
  - Not accepting IXP routes from other AS
  - BGP Resiliency via BGP Timers or BFD
  - Traffic engineering and peering in multiple locations with a same peer

Needed BCOP Content

- Document and analyze approaches for handling of IXP prefixes, and make recommendations
- Sharing and cross-referencing contents from eBGP Configuration BCOP

Next Steps

- Looking for SME who wants to help!
- Document and analyze approaches for handling of IXP prefixes, and make recommendations
  - Collaborating with eBGP Configuration BCOP team

Join Us!

- Are you an expert in Public Exchange Peering?
- Do you have real-world experience with Public Exchange Peering as a participants?
- Are you interested in Public Exchange Peering?
- Get involved!
  - Flexible time commitment
  - Contribute as much or as little as you can
  - The more voices we can include the better
  - Email phsiao@tripadvisor.com to be included

Open Mic

- Global coordination?
  - ISOC?
  - IETF?
- Curate vs. Create
  - Inventing wheels…
- What’s on your mind?

Get Involved!

- BCOP Appeals (questions to answer): http://bcop.nanog.org/index.php/Appeals
- Draft BCOPs (in progress): http://bcop.nanog.org/index.php/BCOP_Drafts
- Mailing List: http://mailman.nanog.org/mailman/listinfo/bcop

Thank You

Let’s Continue the Conversation at the Social…