NANOG 80 Hackathon Recap

Michael Costello (just one of a cast of tens)

Hackathon stats

- 11th Hackathon (first virtual)
- Participants worked from 13:00 EDT Saturday to 17:00 EDT Sunday
- Participants from the US, Canada, and South Africa



Hackathon overview

- In support of NANOG's educational mission
- Organized by the Program Committee
- Weekend before the general conference
- Participants work individually or self-organize into groups
- Participants choose their own projects



Projects included

- A from-scratch BGP daemon
- Performance monitoring and notifications using home control systems
- Simulate traffic and visualize streaming telemetry data
- gnmi-gateway Kafka exporter and Netbox target loader
- A modern NNTP client library in Rust



Projects you'll see here today

- RIB compliance using Model Driven Telemetry
- A fully automated data center Proof of Concept
- A simple, automated, and easy-to-deploy SDN solution



Hackathon for NANOG 81

- We're looking for participants, mentors, and sponsors
- We'll be reviewing survey results ahead of planning
- See you in February







Lawrence BirdYordan Sutantohttps://github.com/TheBirdsNest/https://github.com/yordangit12/https://www.linkedin.com/in/lawrenceabird/https://www.linkedin.com/in/yordan-sutanto-5ba4487/

BGP RIB COMPLIANCE USING MDT

https://github.com/petermoorey/NANOG-8o-Hackathon/

NANOG-80 Hackathon



Peter Moorey https://github.com/petermoorey/ https://www.linkedin.com/in/pmoorey/



Vladimir Yakovlev https://github.com/VladimirGHC/ https://www.linkedin.com/in/vladimir-yakovlev-398ba5ao/

GOAL

- Evaluate various technology stacks
- Evaluate network telemetry capabilities
- Develop a network capable of showcasing various routing scenarios
- Provide real-time evaluation of BGP routes to detect:
 - Poorly configured route-maps
 - Incorrect provider policies
 - Route hijacking

TECHNOLOGY STACK



- CSR 1000V
- CML

- Service Provisioning
- Service Chaining

- InfluxDB
- Telegraf
- Chronograph

- Policy Processing
- Notification Handling
- Websocket Manager

LAB SETUP

- Challenges:
 - Are the regional route preferences being honored for the default route?
 - Is traffic to/from our providers being hijacked?



DATA COLLECTION

templates > 开 mdt_xml.jinja2

(contrig)
<mdt-config-data xmlns="http://cisco.com/ns/yang/Cisco-IOS-XE-mdt-cfg"></mdt-config-data>
<pre>{% for id, sub in data["subscriptions"].items() %}</pre>
<mdt-subscription></mdt-subscription>
<subscription-id>{{ id }}</subscription-id>
 kase>
<stream>yang-push</stream>
<source-vrf>{{ sub["sourcevrf"] }}</source-vrf>
<encoding>encode-kvgpb</encoding>
{% if sub["period"] == "on_change" %}
<no-synch-on-start>false</no-synch-on-start>
{% else %}
<period>{{ sub["period"] }}</period>
{% endif %}
<pre><xpath>{{ sub["xpath"] }}</xpath></pre>
<mdt-receivers></mdt-receivers>
<address>{{ sub["rx"]["ip"] }}</address>
<port>{{ sub["rx"]["tcp_port"] }}</port>
<protocol>grpc-tcp</protocol>
{% endfor %}

def main(grpc_host):

```
hosts = []
```

```
with open("mdt_subscriptions.yml", "r") as f:
   mdt_sub = safe_load(f)
```

```
with open(INVENTORY_FILE, "r") as f:
   inventory = safe_load(f)
```

```
for router in inventory['routers']:
   if inventory['routers'][router]['mdt_bgp']:
       hosts.append(inventory['routers'][router]['ip'])
```

```
for subscription in mdt_sub['subscriptions']:
   mdt_sub['subscriptions'][subscription]['rx']['ip'] = grpc_host
j2_env = Environment(Loader=FileSystemLoader("."), trim_blocks=True, autoescape=True)
template = j2_env.get_template("templates/mdt_xml.jinja2")
new_config = template.render(data=mdt_sub)
```

```
for host in hosts:
```

}

```
connect params = {
    "host": host,
    "username": USER,
    "password": PWD,
   "hostkey_verify": False,
    "allow_agent": False,
   "look_for_keys": False,
    "device_params": {"name": "csr"},
```

with manager.connect(**connect_params) as conn: print(f"NETCONF session connected: {host}")

```
# Perform the update, and if success, print a message
```

```
config_resp = conn.edit_config(target="running", config=new_config)
if config_resp.ok:
```

print(f"Added ({len(mdt_sub['subscriptions'])}) subscriptions")

print(f"NETCONF session disconnected: {host}")

POLICY EVALUATION

1	# PCD PTP Compliance Poli	cy file for dome lab
	# BGP KIB Compliance Poli	cy file for demo lab
	neet enumers Cicco TOE VE	han energing state date/han neute unfs/han neute unf
	noot_source: Cisco-103-Ac	-bgp-oper.bgp-state-data/bgp-route-vrts/bgp-route-vit
	pagions:	gp_rouce_filter/bgp_rouce_encry/bgp_pach_encry/
	omen:	
		# Bonn
		# Bomo
		# Rolle
		# Boston
	- 10.200.49.5	# Duite
	- 10.200.49.0	# Quito
	external:	
		# Americas DC
	- 10.200.49.8	# Europe DC
		# Match default roule
	match: explicit	# Match type is explicit. Nib prefix must match 0.0.0.0/0 exactly
	region: americas	# Region to monitor for KIB updates on
	attributes:	# List of attributes to evaluate and expected values
	community: 10	With the defender works
	0.0.0/0	# Match default route
	match: explicit	# Match type is explicit. RIB prefix must match 0.0.0.0/0 exactly
	region: emea	# Region to monitor for RIB updates on
	attributes	# List of attributes to evaluate and expected values
	community: 10	0:2 # Community value expected in RIB update
	192.168.0.0/16:	
	match: any	# any match will include any network prefix that is within the sub
	region: external	
	attributes	
	origin: 100	
	# community:	
	as_path: 139/	9 13979
	150.0.0.0/16:	
	match: any	
	region: external	
	attributes:	
	origin: 150	

	<pre>is_in_policy(self, prefix: str) -> str:</pre>
	try:
	<pre>target_network = IPNetwork(prefix)</pre>
	<pre>for prefix, policy in self.policies.items():</pre>
	parent_network = IPNetwork(prefix)
	<pre>if policy['match'] == 'any':</pre>
	if target network in parent network:
	<pre>print(f"{target network} in {parent network}")</pre>
	return str(parent network)
	<pre>elif policy['match'] == 'explicit':</pre>
	<pre>if target_network == parent_network:</pre>
	<pre>print(f"{target_network} is {parent_network}")</pre>
	return prefix
	else:
	raise NotImplementedError
	except KeyError:
	<pre>print("Malformed Policy")</pre>
	return None
еŗ	evaluate(set) -> bool:
	query = + SELECT + PROM {SELF.SOURCE} WHERE CLIME > HOW() - 155
	pinic(query)
	for result in results get noints(self source).
	<pre>prefix nath = 'bop route af/bop route filter/bop route entry/prefi</pre>
	prefix = result[prefix path]
	<pre>print(f"Prefix is {prefix} on source {result['source']}")</pre>
	if prefix:
	<pre>compliance = {}</pre>
	<pre>policy_prefix = self.is_in_policy(prefix)</pre>
	<pre>print(f"Policy prefix: {policy prefix}")</pre>

print(f"evaluating prefix {prefix} in policy")
if policy_prefix:
 policy = self.policies[policy_prefix]

t, or the subnet itself

VISUALIZATION



NANOG 80 HACKATHON

- Aakash Rawal
- Hast Patel
- Mukesh Jaiswal
- Nikhil Gadre
- Swati Niture





PROPOSAL

Proof of Concept for a fully automated Tier 2 ISP Data Center – Deployment and Monitoring

What did we achieve ?

- Centralized control and monitoring of the network
- Cost-effectiveness (reduces OpEx)
- One-click configuration and deployment (saves time)
- Reduction in human errors



HIGH-LEVEL DIAGRAM













NETWORK MANAGEMENT AND AUTOMATION SYSTEM (NMAS)

Web App

- Generate Configuration Files
- Push Configuration to Devices
- Monitor Individual Devices and Network
- Compare Running Config to Golden Config
- Backup Configuration to Cloud



GENERATING CONFIG FILES

- Python
- Jinja2 Templates
- Modules
 - Netmiko
 - NAPALM
 - Flask
- Multi-threading



1	{% for k in natinfo["Interfaceout"] %}	HPOLAAA		1	{% for k in ospfinfo["int"] %}	Witherson and a second
2	<pre>interface {{ k }}</pre>	10000000000000000000000000000000000000		2	<pre>int {{ k }}</pre>	
	ip nat outside				<pre>ip address {{ ospfinfo["int"][k] }}</pre>	
	{% endfor %}				no shut	
			Т		exit	
	{% for k in natinfo["Interfacein"] %}				{% endfor %}	
	int {{ k }}				{% for j in range((ospfinfo["network"]) length) %}	
	ip nat inside				router ospf {{ ospfinfo["pid"] }}	
	{% endfor %}				<pre>network {{ ospfinfo["network"][j] }} area {{ ospfinfo["area"]</pre>	}
10				10	exit	
11	{% for k in natinfo["accesslist1"] %}			11	{% endfor %}	
12	access-list {{ k }}			12		
13	{% endfor %}			13		
14						
15	{% set count=namespace(value=1) %}					
16	{% if natinfo["accesslist1"] length == 1 %}					
17	<pre>ip nat inside source list 1 interface {{natinfo["Interfaceout"]</pre>					
18	{% else %}					
19	{% for k in natinfo["Interfaceout"] %}					
20	<pre>ip nat inside source list {{count.value}} interface {{k}} overl</pre>					
21	{% set count.value=count.value +1 %}					
22	{% endfor %}					
23	{% endif %}					



🗐 nat_info.txt - Notepad —		R6_natconf.txt - Notepad	- 🗆	\times
File Edit Format View Help		<u>F</u> ile <u>E</u> dit F <u>o</u> rmat <u>V</u> iew <u>H</u> elp		
K	^	interface fa1/1		1
		ip nat outside		
"R1":		interface fa2/0		
{"Interfaceout":["fa4/1"], "accesslist1": ["1 permit 198.51.0.0		ip nat outside		
0.0.255.255"], "Interfacein":["fa0/1","fa1/0","f0/0"]},		interface fa0/0		
		ip nat outside		
"R2":				
{"Interfaceout":["fa4/1"], "accesslist1": ["1 permit 198.51.0.0		int fa4/1		
0.0.255.255"], "Interfacein":["fa0/1","fa1/0","f0/0"]},		ip nat inside		
"DC".				
KD : ("Tutter General Hard No. 1/1" "General Here (0.0) "General Here (0.0)		access-list 1 permit 205.0.0.0 0.0.255.255		
{ Interfaceout : [fai/1 , faz/0 , fa0/0], accessiisti : [1 per	L.WTC	access-fist 2 permit 205.0.0.0 0.0.255.255		
205.0.0.0 0.0.255.255 , 2 permit 205.0.0 0.0.255.255 , 5 permit 205.0.0.0 0.0.255.255"], "Interfacein":["fa4/1"]},	L	access-iist 5 permit 205.0.0.0 0.0.255.255		
		ip nat inside source list 1 interface fa1/1 overload		
"R6":		ip nat inside source list 2 interface fa2/0 overload		
<pre>{"Interfaceout":["fa1/1","fa2/0","fa0/0"], "accesslist1": ["1 per</pre>	rmit	ip nat inside source list 3 interface fa0/0 overload		
205.0.0.0 0.0.255.255","2 permit 205.0.0.0 0.0.255.255","3 permit	t			
205.0.0.0 0.0.255.255"], "Interfacein":["fa4/1"]}				
}				



GGG Boulder Network Management and Automation Station (NMAS) Version 1.0
OSPF config files are successfully created, the files are:
R1_ospfconf.txt
R2_ospfconf.txt
R3_ospfconf.txt
R4_ospfconf.txt
R5_ospfconf.txt
R6_ospfconf.txt
press <u>here</u> to return to home page.

A COLUMN TO A COLUMN

1. ST 1. 1

States and Provide and

and the second



PUSH CONFIGS TO DEVICES

Netmiko, Napalm

- Interface Configs
- OSPF
- NAT
- BGP
- IPSec
- 6to4 Tunneling













MONITOR INDIVIDUAL DEVICES

- Device Interfaces
 - IP / Status
- OSPF Neighborship
- BGP Neighborship
- View Running Config
- Show Commands



Boulder
Network Management and Automation Station (NMAS)
Version 1.0
Interfaces IPs
BGP data
OSPF data
Running configs
press <u>here</u> to return to home page.











A DO TO	Boulder
A CONTRACTOR	Network Management and Automation Station (NMAS)
	Version 1.0
	Enter show command
1 w/00	show ip route
	Enter host
	192.168.100.1
	Enter username
5	test
	Enter password
1	test
	Enter device-type i.e cisco_ios
	cisco_ios Submit



COMPARE RUNNING CONFIG TO GOLDEN CONFIG

- Backup Running Config
- Compare
 - Latest Backed Up Running Config
 - Golden Config
- Show Difference





['Difference found for R1', 'Difference found for R2', 'Difference found for R3', 'Difference not found for R4', 'Difference found for R5', 'Difference found for R1', 'Difference foun





BACKUP CONFIG FILES TO CLOUD

Actions ✓ US West (N. California) US West (N. California) US West (N. California) US West (N. California) US West (N. California) Viewing 1 to 31 Viewing 1 to 31 V	Overview
Actions ~ US West (N. California) C Actions ~ Viewing 1 to 31 Size ~ Storage class ~ Last modified ~ Size ~ Storage class ~ Size ~ Storage class ~ Oct 18, 2020 10:44:20 AM GMT 0 B Standard O O Standard Oct 18, 2020 12:45:41 PM GMT 2.3 KB Standard O O O Standard O O Standard O </td <td>Q Type a prefix</td>	Q Type a prefix
Viewing 1 to 31 Last modified ▼ Size ▼ Storage class ▼ Oct 18, 2020 10:44:20 AM GMT- 0600 0 B Standard Oct 18, 2020 12:45:41 PM GMT- 0600 2.3 KB Standard Oct 18, 2020 12:45:41 PM GMT- 0600 2.5 KB Standard Oct 18, 2020 10:44:18 AM GMT- 0600 2.5 KB Standard	1 Upload
Last modified ▼Size ▼Storage class ▼Oct 18, 2020 10:44:20 AM GMT- 06000 BStandardOct 18, 2020 12:45:41 PM GMT- 06002.3 KBStandardOct 18, 2020 10:44:18 AM GMT- 06002.5 KBStandardOct 18, 2020 10:44:18 AM GMT- 06002.5 KBStandard	
Oct 18, 2020 10:44:20 AM GMT- 0600 0 B Standard Oct 18, 2020 12:45:41 PM GMT- 0600 2.3 KB Standard Oct 18, 2020 10:44:18 AM GMT- 0600 2.5 KB Standard Oct 18, 2020 10:44:18 AM GMT- 0600 2.5 KB Standard	Name ▼
Oct 18, 2020 12:45:41 PM GMT- 0600 2.3 KB Standard Oct 18, 2020 10:44:18 AM GMT- 0600 2.5 KB Standard Oct 18, 2020 12:45:40 PM GMT- 0600 2.5 KB Standard	□ 🖹 R12_20
Oct 18, 2020 10:44:18 AM GMT- 0600 2.5 KB Standard Oct 18, 2020 12:45:40 PM GMT- 0600 2.5 KB Standard	☐ R12_20
Oct 18, 2020 12:45:40 PM GMT- 0600 2.5 KB Standard	□ 🖹 R1_202
0.140,0000 (0.4440 MM ONT	□ 🖹 R1_202
Oct 18 2020 10:44:18 AM GMT-	≝ R1202



FUTURE ENHANCEMENTS

- Zero Touch Provisioning (ZTP) Management Network
- Automated Device Backups Cron Job
- Security
 - Firewall
 - Access-lists
- Redundancy
 - For Default Gateway



THANK YOU

A Simple, Automated and Easy to Deploy SDN Solution

> Roni Mukherjee MS, Telecommunications 2019-21, CU Boulder

Why SDN?

- Reduced CapEx: Simple device move control away from devices
- Reduced OpEx: Simplified network monitoring automation
- Build your solution
- Faster troubleshooting

Hybrid Network

- Combination of SDN and Traditional Routing and Switching Networks
- No need to change traditional network
- SDN network is spun up to join the network
- Based on the implementation, SDN Controller can be anywhere in the network
 - Only, IP connectivity is required to establish a TCP/TLS connection between the SDN switch (s) and the controller (s)

Hybrid Network – Proof of Concept



Automated SDN Deployment

 Mininet boots up, gets DHCP address, routing enabled, controller configured – Automated

PS C:\MyCanvas\NANOG80_hackathon\submission> python .\sdn_network_spinup.py ******** The IP leased to Mininet VM is: 192.168.100.2

******* OpenFlow version changed to --> 1.3

******* Controller configured!

******** Controller connected!

PS C:\MyCanvas\NANOG80_hackathon\submission>

Advantages of the proposed solution

- Easy to deploy Automated script (boot up/routing/monitoring)
- SDN Centralized control, global view, reduced CapEx and OpEx
- API Use of APIs make it faster and human readable
- No need to change existing network
- Reuse SDN multipurpose simple devices
- GUI: Greatly increased troubleshooting efficiency

SDN GUI – Reusable/Extendable Global View

- Purpose-built for the proposed network
- Web Application Console logs
- Web GUI
- Real-time graphical plots of the network (auto-refresh every 5 seconds)
- Greatly reduced troubleshooting steps match on 40 fields and create own flow

SDN Topology



SDN GUI – Console Logs

PS C:\MyCanvas\NANOG80_hackathon\submission> python .\sdn_gui.py
******* SDN GUI running
* Serving Flask app "sdn_gui" (lazy loading)
* Environment: production
Use a production WSGI server instead.
* Debug mode: on
* Restarting with stat
******* SDN GUI running
* Debugger is active!
* Debugger PIN: 123-884-149
* Running on http://127.0.0.1:5000/ (Press CTRL+C to quit)
Number of Packet_INs to controller in this iteration is: 3
Number of Rx packets of port1 of switch1 in this iteration is: 31
Number of Tx packets of port1 of switch1 in this iteration is: 31
Number of matched traffic pattern sourced from H1 and destined to H2 in this iteration is: 0
Number of violation count for the block traffic flow rule to H4 in this iteration is: 1
Number of HTTP packets at switch1 in this iteration is: 24
127.0.0.1 [17/Oct/2020 23:15:10] "GET / HTTP/1.1" 200 -
127.0.0.1 [17/Oct/2020 23:15:10] "GET /static/img/1.jpg HTTP/1.1" 304 -
Number of Packet_INs to controller in this iteration is: 4
Number of Rx packets of port1 of switch1 in this iteration is: 33
Number of Tx packets of port1 of switch1 in this iteration is: 35
Number of matched traffic pattern sourced from H1 and destined to H2 in this iteration is: 1
Number of violation count for the block traffic flow rule to H4 in this iteration is: 3
Number of HIIP packets at switchi in this iteration is: 24
127.0.0.1 [17/Oct/2020 23:15:22] "GET / HTTP/1.1" 200 -
127.0.0.1 [17/Oct/2020 23:15:22] "GET /static/img/1.jpg HTTP/1.1" 304 -

SDN GUI – Web GUI



SDN GUI – Web GUI



SDN Web GUI – Switch Tx/Rx Monitor

- Tx/Rx traffic has a similar pattern.
- If there is a significant difference between Tx/Rx traffic, captured on graph and an email is sent to the administrator for immediate action.
- Graph is plotted against realtime data and refreshed every 5 seconds.



SDN Web GUI – Firewall Monitor

- Custom firewall rule is created.
- If there is a significant increase in violation count, captured on graph and email is sent to administrator for immediate action.
- Graph is plotted against realtime data and refreshed every 5 seconds.



SDN Web GUI – Web Traffic Monitor

- One of the host is web server and serving HTTP requests.
- This is a critical service.
- If there is a significant change in traffic pattern, captured on graph and an email is sent to the administrator for immediate action.
- Graph is plotted against real-time data and refreshed every 5 seconds.



SDN Web GUI – Pattern Traffic Monitor

- Custom match is created for traffic pattern, e.g. traffic sourced from MAC address of H1 and destined to MAC address of H2.
- This is considered as critical service.
- If there is a significant change in traffic pattern, captured on graph and an email is sent to the administrator for immediate action.
- Graph is plotted against real-time data and refreshed every 5 seconds.



SDN Web GUI – OpenFlow Packet_In Monitor

- Traffic from switch to controller.
- This is the traffic that matches table-miss.
- Graph is plotted against real-time data and refreshed every 5 seconds.



SDN Web GUI – Violation email

Hey Admin > Inbox ×



9

-

-

ronismtp97@gmail.com HIGH VIOLATION COUNT!!

ronismtp97@gmail.com HIGH VIOLATION COUNT!!

ronismtp97@gmail.com

to 💌

HIGH VIOLATION COUNT!!

Tools used

- GNS3
- Cisco IOS routers
- Mininet
- Ryu SDN controller
- Ubuntu VMs
- VirtualBox

• Flask • HTML/CSS Netmiko • NETCONF • Python Wireshark

Conclusion

- Custom flow rules using any match criteria (from 40 different fields)
- Graphical representation significantly increased troubleshooting efficiency
- Hybrid Network
 - Reduced CapEx: Simple device, reusability
 - Reduced OpEx: Significantly improved network monitoring and troubleshooting through automation
- Automated deployment of the Hybrid-SDN solution

GitHub

https://github.com/rm-viable/NANOG80-hackathon

