



Managing Networks in a Software-Defined Future

Jeff Gehlbach

Southern California Linux Expo 2015 (SCaLE13x)

February 23, 2015

- ▶ Speaker vitals
- ▶ Elements of old-school networks
- ▶ Elements of software-defined networks
- ▶ A case study
- ▶ Conclusions
- ▶ Questions!

Jeff Gehlbach

Ten fingers, ten toes, some industry experience



- ▶ NASA NISN → Management of large IP networks
- ▶ Empire / Concord → Making and consulting on NMS
- ▶ BellSouth → Cranium formed into Bell shape
- ▶ OpenNMS Group → Making and consulting on *free* NMS!

Elements of old-school networks

A large, light green, curved decorative element that starts from the bottom left and arcs across the middle of the slide, ending on the right side.

Elements of old-school networks

Switches (oversimplified)



- ▶ Functions:
 - ▶ Switching L2 frames
 - ▶ Running STP
- ▶ Many physical ports, often modular
- ▶ High-throughput data plane
- ▶ Control plane driven by local config (!)

Elements of old-school networks


Routers (oversimplified)



- ▶ Functions:
 - ▶ Forwarding L3 packets
 - ▶ Running OSPF, BGP, *et al*
- ▶ Relatively few physical ports, often modular
- ▶ Medium- to high-throughput data plane
- ▶ Control plane driven by local config (!!)

Elements of old-school networks

Firewalls (oversimplified)



- ▶ Functions:
 - ▶ Forwarding L3 packets subject to a ruleset
 - ▶ Taking blame when anything breaks
- ▶ Relatively few physical ports, sometimes modular
- ▶ Low- to medium-throughput data plane
- ▶ Control plane driven by local config (!!!)

Elements of old-school networks

Inventory and configuration management



- ▶ Functions:
 - ▶ Making the right configs run on the right devices
 - ▶ Accounting for hardware elements in the network
 - ▶ Eating time and / or money
- ▶ Two separate problems really, each pretty hard
- ▶ Typically no understanding of configurations (control plane)
- ▶ Are all your nodes in your inventory?

Image: graemefazakerley / DeviantArt / CC BY-SA 3.0

Elements of old-school networks



Network management system (NMS)

- ▶ Functions – OSI FCAPS model:
 - ▶ **F**ault management*
 - ▶ **C**onfiguration management
 - ▶ **A**ccounting management
 - ▶ **P**erformance management*
 - ▶ **S**ecurity management
- ▶ OpenNMS adheres roughly to FCAPS
- ▶ Focus on fault (FM) and performance (PM)

Elements of old-school networks

FM and PM as implemented in OpenNMS

- ▶ **Provisioning** – how can we get nodes, interfaces, services into the system?
- ▶ **Service assurance** – how can we know whether important network entities are responsive?
- ▶ **Fault management** – how can a network element tell us it has a problem?
- ▶ **Performance management** – how can we quantify utilization of a network interface or a CPU?

Elements of old-school networks

Simple Network Management Protocol (SNMP)




- ▶ Functions:
 - ▶ NMS-to-managed-node data queries (GET / GET-BULK)
 - ▶ Managed-node-to-NMS unsolicited messages (TRAP)
- ▶ Routers, switches, *et al* are where the action happens
- ▶ The NMS talks to the SNMP agent on the managed node
- ▶ Data gathered: interface traffic, BGP statistics, environmentals...
- ▶ Extensible via Management Information Base (MIB)

Image: tedeytan / Wikimedia Commons / CC SA-2.0 Generic

Elements of old-school networks


In summary



- ▶ Many sovereign nodes with local configs driving control plane
- ▶ When we're lucky, traffic flows as intended
- ▶ Impossible to simulate accurately
- ▶ Clearly not designed by hackers

Image: D J Shin - My Toy Museum / Wikimedia Commons / CC BY-SA 3.0 Unported

Elements of software-defined networks


A large, light green, curved decorative element that starts from the bottom left and arcs across the middle of the slide, partially overlapping the text.

Separation of planes is key!

If you take away just one SDN fact:

SDN is about **separation** of
control plane from **data**
plane; and **programmability**.

Data plane



- ▶ Functions:
 - ▶ Moving frames or packets around
 - ▶ According to rules gotten from controller (“control plane”)
- ▶ Comparatively generic hardware
- ▶ Sometimes virtual
- ▶ Called “switches” regardless of role

Control plane



- ▶ Functions:
 - ▶ Control behavior of switches (“data plane”)
 - ▶ According to centrally-managed rules (eases config)
 - ▶ Across registered nodes (eases inventory)
 - ▶ Expose inventory, configuration, etc. via open APIs
 - ▶ Scripting hooks for network programmability
- ▶ Controller is just a general-purpose computer
- ▶ May have a bridge or flower tattooed on it
- ▶ May be virtual

In summary



- ▶ Relatively dumb switches
- ▶ Switch inventory, configurations centrally managed
- ▶ Programmability enables awesome wackiness, agility
- ▶ When we're skilled, traffic flows as intended
- ▶ Might even be unit-testable
- ▶ This is how hackers would build a network!

Image: Andreas Trepte / Wikimedia Commons / CC SA 2.5 Generic

A case study

A large, light green, curved decorative element that starts from the left side of the slide and curves upwards and to the right, passing behind the main title.

Controller: Project Floodlight

Controller: Project Floodlight

- ▶ Implements OpenFlow 1.0 – 1.4
- ▶ Apache-licensed
- ▶ Maintained by Big Switch Networks
- ▶ projectfloodlight.org

Switches: Open vSwitch / Fedora 21

Switches: Open vSwitch / Fedora 21

- ▶ Implements OpenFlow 1.3
- ▶ Apache-licensed
- ▶ Distributed maintainership
- ▶ Kernelspace implementation in Linux, FreeBSD
- ▶ Userspace implementation in NetBSD
- ▶ openvswitch.org

Provisioning / Inventory

A large, light green, curved swoosh graphic that starts from the bottom left, curves upwards and to the right, passing behind the main title, and then curves downwards towards the right edge of the slide.

SDN for provisioning

Provisioning: SDN controller as a source of truth

- ▶ Option 1: Push inventory from controller to OpenNMS API
 - ▶ Leans on SDN controller's internal programmability
- ▶ Option 2: Pull inventory from controller's API
 - ▶ Leans on SDN controller's API

Option 1: Push-mode

Floodlight → OpenNMS

- ▶ Floodlight features pluggable notification managers

```
1 public interface INotificationManager {
2     /**
3      * Post a notification. Depending on the underline implementation, it
4      * may write the notes to log file or send an SNMP notification/trap.
5      *
6      * @param notes    string message to be sent to receiver
7      */
8     public void postNotification(String notes);
9 }
```

```
1 private static class NotificationSwitchListener implements IOFSwitchListener {
2     // ...
3     @Override
4     public void switchAdded(DatapathId switchId) {
5         notifier.postNotification("Switch " + switchId + " connected.");
6     }
7     // ...
8 }
```


SDN for provisioning

Floodlight → OpenNMS



- ▶ Default implementation just squawks to syslog
 - ▶ Write a new one that POSTs to OpenNMS requisition ReST endpoint
 - ▶ Or just watch logs from outside, do the POST from there
- ▶ Doesn't seem the cleanest approach, but should be effective
- ▶ Floodlight / other SDN controller hackers, comments?

Option 2: Pull-mode

- ▶ Query Floodlight's core/controller/switches endpoint

```
1 // Output of http://mal:8080/wm/core/controller/switches/json
2 [
3   {
4     // Switch "wash"
5     "inetAddress": "/10.0.0.138:45261",
6     "connectedSince": 1424451598399,
7     "switchDPID": "00:00:26:09:6a:ae:e3:49"
8   },
9   {
10    // Switch "zoe"
11    "inetAddress": "/10.0.0.57:35907",
12    "connectedSince": 1424453016500,
13    "switchDPID": "00:00:d2:0b:68:3a:d2:49"
14  }
15 ]
```

- ▶ Query Floodlight's core/switch/<DPID> endpoint

```
1 // Output of http://mal:8080/wm/core/switch/00:00:d2:0b:68:3a:d2:49/desc/json
2 // This is "zoe"
3 {
4   "desc": {
5     "version": "OF_13",
6     "manufacturerDescription": "Nicira, Inc.",
7     "hardwareDescription": "Open vSwitch",
8     "softwareDescription": "2.3.1-git3282e51",
9     "serialNumber": "None",
10    "datapathDescription": "None"
11  }
12 }
```

- ▶ A bit short on details, but that's on Open vSwitch
- ▶ Anybody with Nexus, Arista, etc. gear see better data?

OpenNMS ← Floodlight

- ▶ Query Floodlight's /core/switch/<DPID>/port-desc endpoint

```
1 // Output of http://mal:8080/wm/core/switch/00:00:d2:0b:68:3a:d2:49/port-desc/json
2 // This is switch "zoe"
3 {
4   "version": "OF_13",
5   "portDesc": [
6     {
7       "portNumber": "1",
8       "hardwareAddress": "06:4e:04:ca:b5:70",
9       "name": "eth1",
10      "config": "1",
11      "state": "1", // ...
12      "currSpeed": "1000000",
13      "maxSpeed": "10000000"
14    },
15    {
16      "portNumber": "local",
17      "hardwareAddress": "d2:0b:68:3a:d2:49",
18      "name": "br-int", // ...
19    }
20  ]
21 }
```

- ▶ Build a requisition (PRIS source plugin)
 - ▶ Foreign-ID = DPID

```
1 <?xml version="1.0"?>
2 <model-import foreign-source="floodlight-switches">
3   <node node-label="wash" foreign-id="00:00:26:09:6a:ae:e3:49">
4     <interface ip-addr="10.0.0.138" descr="" status="1" snmp-primary="P" />
5   </node>
6   <node node-label="zoe" foreign-id="00:00:d2:0b:68:3a:d2:49">
7     <interface ip-addr="10.0.0.57" descr="" status="1" snmp-primary="P" />
8   </node>
9 </model-import>
```

Service assurance

A large, light green, curved decorative element that starts from the bottom left, curves upwards and to the right, and then levels off towards the right edge of the slide.

SDN for network management

Service assurance (“are the switches up?”)

- ▶ Controller exposes presence / absence of switches
- ▶ Most other measures best done through synthetic transactions directly to switches
- ▶ Seems not much will change soon in this facet

A large, light green, curved graphic element that starts from the bottom left and arcs across the middle of the slide towards the right.

Fault

SDN for network management



Fault management (“ZOMG a switch broke!”)

- ▶ Controller able to send unsolicited messages to an NMS
- ▶ Similar in function to SNMP traps
- ▶ Examples
 - ▶ “Switch 00:00:00:00:de:ad:be:ef joined the controller”
 - ▶ “Switch 00:00:00:00:ca:fe:ca:fe left without saying goodbye”
- ▶ Not yet well-developed in main Floodlight code base
 - ▶ Downstream OEMs may provide their own NotificationManagers
- ▶ OpenNMS can reparent data onto the correct node (switch) via its Event Translator facility

Performance

A large, light green, curved decorative element that starts from the bottom left and arcs across the middle of the slide, ending on the right side.

SDN for network management



Performance management (“How busy is that switch interface?”)

- ▶ Floodlight exposes interface-level metrics and other stuff via ReST
- ▶ OpenNMS can collect performance data directly via ReST using XMLCollector with JSON handler
- ▶ Data trivially reparented onto the correct node (switch)

Performance management ("How busy is that switch interface?")

- ▶ Query Floodlight's /core/switch/<DPID>/port endpoint

```
1 // Output of http://mal:8080/wm/core/switch/00:00:d2:0b:68:3a:d2:49/port/json
2 // This is switch "zoe"
3 {
4   "version": "OF_13",
5   "port": [
6     {
7       "portNumber": "1",
8       "receivePackets": "5213610",
9       "transmitPackets": "2947725",
10      "receiveBytes": "2855576667",
11      "transmitBytes": "2354303692",
12      "receiveDropped": "0",
13      "transmitDropped": "0",
14      "receiveErrors": "0",
15      "transmitErrors": "0",
16      "receiveFrameErrors": "0",
17      "receiveOverrunErrors": "0",
18      "receiveCRCErrors": "0",
19      "collisions": "0", // ...
20    }, // ...
21  ]
22 }
```

SDN for network management

Performance management (“How busy is that switch interface?”)

A large, light green, curved graphic element that starts from the bottom left and arcs across the middle of the slide towards the right.

But...


SDN for network management



Performance management (“How busy is that switch interface?”)

- ▶ Scalability of ReST / JSON-based collection to huge networks is unproven
- ▶ Most SDN switches on the market also support SNMP
- ▶ Every NMS in the world groks SNMP already
 - ▶ Prediction: Gradual transition from SNMP to controller API
 - ▶ Consistency across controller APIs is key

Conclusions

A large, light green, curved decorative element that starts from the bottom left and arcs across the middle of the slide, ending on the right side.

Conclusions

How OpenNMS is coping



- ▶ It's still early days for SDN on the ground
 - ▶ Standards landscape frequently changing
 - ▶ Most deployments we see are hybrid
- ▶ We've had some practice with similar movements
- ▶ Work on SDN full time? Let's chat over a beer.

Is SNMP finally dead?

Predicted since late 1990s or earlier

- ▶ Not yet. Sorry.
- ▶ Problems? Sure.
 - ▶ Painful to implement
 - ▶ SMI struggles to model really complex relationships
 - ▶ Stateless nature increasingly problematic with larger data sets
- ▶ Still useful, though
- ▶ Entrenchment + utility = durability

Questions!

jeffg@opennms.org

IRC: Freenode jeffg / #opennms

Twitter: @jeffgdotorg

License: Creative Commons Attribution-ShareAlike 4.0 International

