IPv6 Networking in SDN

Eric Choi
Senior Product Manager, Product Management
Service Provider Group, APJ

Brocade in SDN Community

Active in all SDN Industry Community Efforts

- Open Network Foundation OpenFlow
  - ONF launched in March 2011 with Brocade as founding member
  - Curt Beckmann chairs FAWG, is a member of CAB
- OpenStack
  - Brocade is a member, has developed VDX and ADX plug-ins
- Network Functions Virtualisation, an ETSI Industry Specification Group
  - Brocade is a participant, hosted NFV Infrastructure meeting in September
  - Two key Vyatta participants: CTO Robert Bays and Distinguished Engr Mukhtiar Shaikh
- OpenDaylight Project
  - Brocade's SP CTO David Meyer chairs Technical Steering Committee

Things to learn (hint!)

- *OpenFlow is clearly an SDN protocol. What others are there?*
- *What makes newer versions of OpenFlow tricky on hardware?*
- *What version of OpenFlow has what you need for IPv6?*
- *What is the current "stable release" of OpenFlow?*

SDN protocols

- SDN is “Software Defined Networking”
  - Most accepted definition: SDN decouples control plane from data plane
  - Note: “Real” SDN should be fully interoperable (to deliver the full promise)
  - This implies a fully open standard. Hence ONF/OpenFlow
  - OpenFlow is the best known SDN control protocol
    - But it’s not the only, and arguably not the first. E.g. CAPWAP, many others
    - Still, OpenFlow is much broader protocol (a double edged notion!)
- Other control protocols with growing SDN relevance
  - Path Computation Element Protocol (PCEP), at IETF
  - Interface to Internet Routing System (I2RS), at IETF
- And data protocols
  - VXLAN, NVGRE, STT, which are part of NVo3, at IETF

SDN benefits

- What's so good about separating control and data?
  - It lets you put the brains outside the box
  - That sounds like we've decided to move the brains outside…
  - But really, there are new cases where brains are already outside
- At human scales and speeds, reactive networks are okay
- With orchestration (brains outside), networks need to be proactive
  - Amazon cloud, etc.
- Most SPs now looking to put network functions into VMs: NFV
  - Provide more services faster, etc
  - Reduce OpEx and errors, etc

Is SDN IPv6 Ready?

- IPv6 has been the next Y2K since, yeah, Y2K
  - Been getting pretty real for a couple of years
- SDN hype really started to crank up in March 2011
  - Now starting to get real here and there
- But SDN protocols are still developing
  - The broadly adopted version of OpenFlow doesn't support v6
  - A newer version of OpenFlow does support IPv6, but no deployed yet
  - I2RS and PCEP still in-process

OpenFlow Switch protocol

- OpenFlow Switch (OF-Switch or OFS) is the main OF protocol
  - There is also “OF-Config”, still relatively new and not yet broadly adopted
- OFS1.0, Dec 2009, almost 4 years ago
  - Simple (and limited) single flow table. Well-adopted. No IPv6
- OFS1.1, Feb 2011, almost 3 years ago
  - Just before ONF launched; a hurry-up job (no working code), no IPv6
  - Lots of new features, like powerful (and challenging) multiple flow tables
- OFS1.2, Dec 2011, a bit of IPv6
- OFS1.3, Mar 2012, solid IPv6. Chosen as a “focus release”
- OFS1.4, Jun 2013 (after deliberate gap)
  - Work has also begun on OFS1.5

Testing and Adoption

- OF Conformance Testing has started!
  - 2 labs certified; 1 switch has achieved certification for OFS1.0
- Despite wide support of OFS1.0, interest in IPv6, etc, only minor testing of OFS1.3 last June.
  - But last November a full OFS1.3 plugfest was held
  - Some good news, but no formal reports available at this point
    - My expectation: probably some difficulties with multiple table scenarios
- Why has OFS1.3 adoption taken so much longer?
  - Theories: Chicken-and-egg (switch/controller)… Specs changing too fast
    - And maybe tricky multiple tables are a factor (I'm convinced)
    - Meanwhile, tons of optional functionality make interop confusing too
  - Some firms may choose to do single-table OFS1.3 (it is quite legal)

Single / Multiple Flow Tables

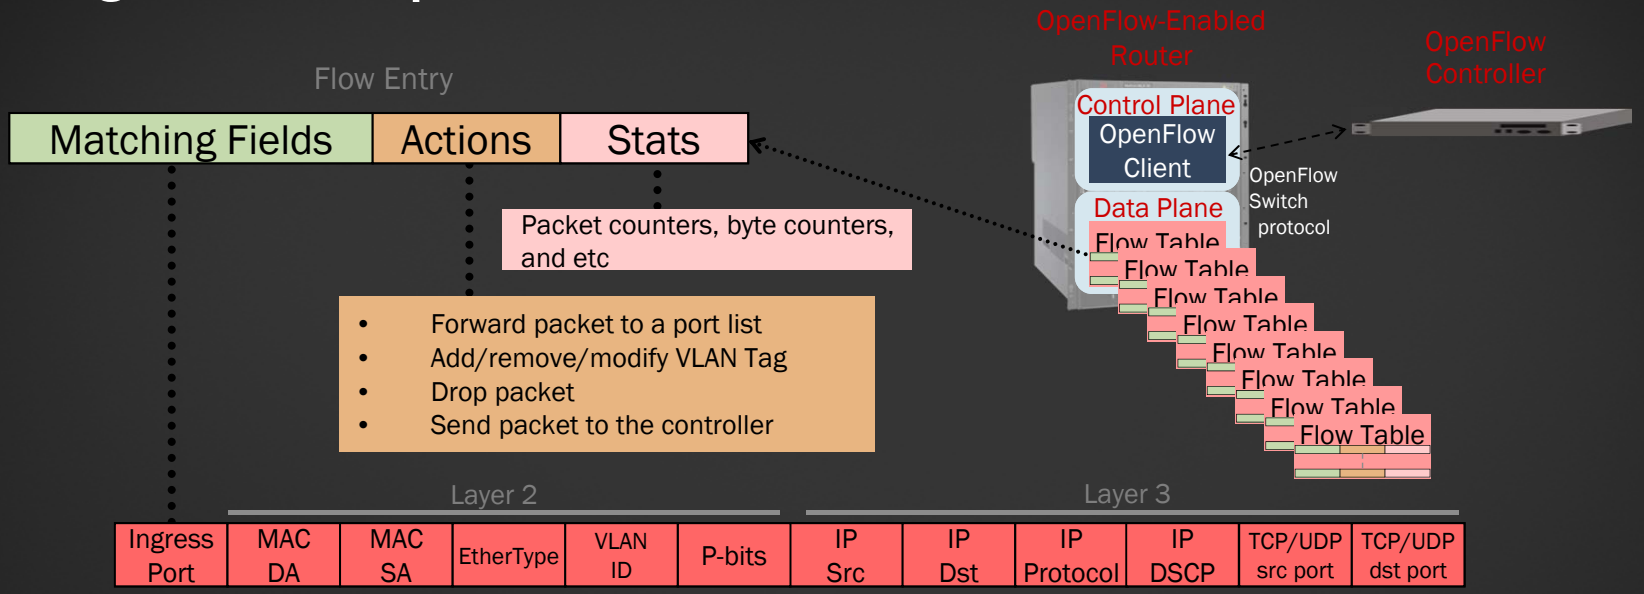

- Each flow table entry contains a set of rules to match (e.g., IP src) and an action list to be executed in case of a match (e.g., forward to port list)
- Single flow table: a message tells the switch the full story
- Single flow table: the switch processing “pipeline” is trivial (not even really a pipeline)
- Multiple flow tables: no message tells full story; “pipelines” no longer trivial, # of variables is huge

The Switch Pipeline

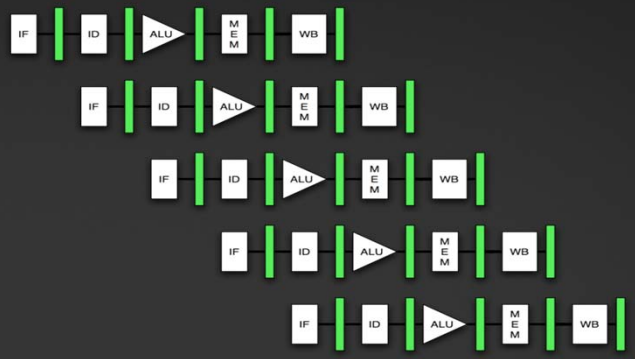

http://commons.wikimedia.org/wiki/File:Pipeline-2.png

- "A sequence of processing logic
  - Wide variety in today's networking ASICs
- When OpenFlow had a trivial pipeline (one table), easy to code
  - Switch software could translate individual messages to ASIC pipeline
- But with non-trivial pipelines and incomplete messages
  - Only works when switch pipeline matches OpenFlow pipeline (CPU, NPU)
    - ASIC, FPGA pipelines don't match; "pipeline mapping" is needed
    - No mapping mechanism planned for
      - Switch developers try to figure out mapping using OpenFlow Switch control messages
    - Interoperable (no "tricks") multi-table support is arguably impossible on ASICs
      - Some inadequate messages defined, but no best practices, etc

Forwarding Abstractions WG

- FAWG goal: Make multi-table OpenFlow work on today's ASICs
  - Make it practical to develop, and reliable interoperable
- Not a protocol problem… we need to change the framework
  - Help with that "pipeline mapping" problem, don't need to do it at run-time
- Network operators won't run new, risky stuff in production
  - Everything will have been tested before (by vendors at least, probably lab)
  - If it's all done before, then mapping solved…why do it again at run time?
  - So, new picture: let's sort out the mapping before hand
  - First, create a detailed description of the pipeline
  - Then "compile it" (map it) onto ASICs

FAWG uses TTPs (= NDMs)

- The “detailed description of a pipeline” is a model of switch behavior; a “datapath model”.
- The 1st gen datapath model is for today’s OpenFlow
  - Today’s OpenFlow is all about Tables.
  - So we called our 1st gen datapath models “Table Type Patterns”…TTPs
- FAWG and OF-Config WG codeveloped a way for controllers and switches to dynamically select between datapath models
  - We want this selection process to work for future generations also
  - We call the umbrella term, “Negotiable Datapath Models”… NDMs

More on TTPs / NDMs

- FAWG's new framework assumes there will be several (10? 100?) common TTPs, each targeting a particular market segment.
  - E.g. "cloud data center" or "MPLS WAN virtualization" or "campus"
- But FAWG doesn't want to be a bottleneck by defining all TTPs
  - So instead, plan is to enable SDN community to define TTPs they need
  - Which also allows for TTPs under NDA (at least until after achieve GA)
- TTP framework is in "beta", starting "PoC" code. Done Q1?
- Along the way, we found unexpected benefits for TTPs
  - Useful for test profiles, even for products that are not TTP aware
  - Useful for product data sheets, to make interoperability more visible
  - Also useful for soft switches (NPUs, CPUs) to run faster, scale farther

The controller story

- The SDN architecture includes a logically centralized (though likely clustered or otherwise distributed) SDN controller.
  - In some versions, there may be multiple SDN clusters
- There are numerous Open Source controllers
  - Each has its champion; a vendor that effectively controls the development
  - Really, there were too many… better to have a de facto dominant one
- OpenDaylight consortium
  - First big release Dec 9!
  - Still needs to prove itself
  - But chances are pretty good

Summary

- OFS1.3 is a stable release that provides good IPv6
- OFS1.3 being broadly developed, well-attended plugfest
- OFS conformance testing labs are making progress
  - OFS1.0 certification has been done
  - OFS1.3 conformance testing is in development
  - Good news, but many optional functions… Conformance means what?
- Multi-table OpenFlow (like OFS1.3) has benefits and challenges
  - Vendors not required to offer multiple tables
  - FAWG (my WG) has new framework (“TTPs”) in beta, likely done Q1 ’14
- Single table IPv6 work can begin as soon as product ship
  - Single table can be great way to start, and may also work long term
- Multiple table will likely be ready late 2014

THANK YOU

Eric Choi
Senior Product Manager, Product Management
Service Provider Group, APJ

email: echoi@brocade.com