Lancaster University

sBGP: A hybrid SDN approach to interdomain routing

uknof35, 8th September 2016

Nicholas Hart, Lancaster University

Lancaster University

Introduction

Lancaster University

The French Inquisition?

Where to does academia fit in networking development?

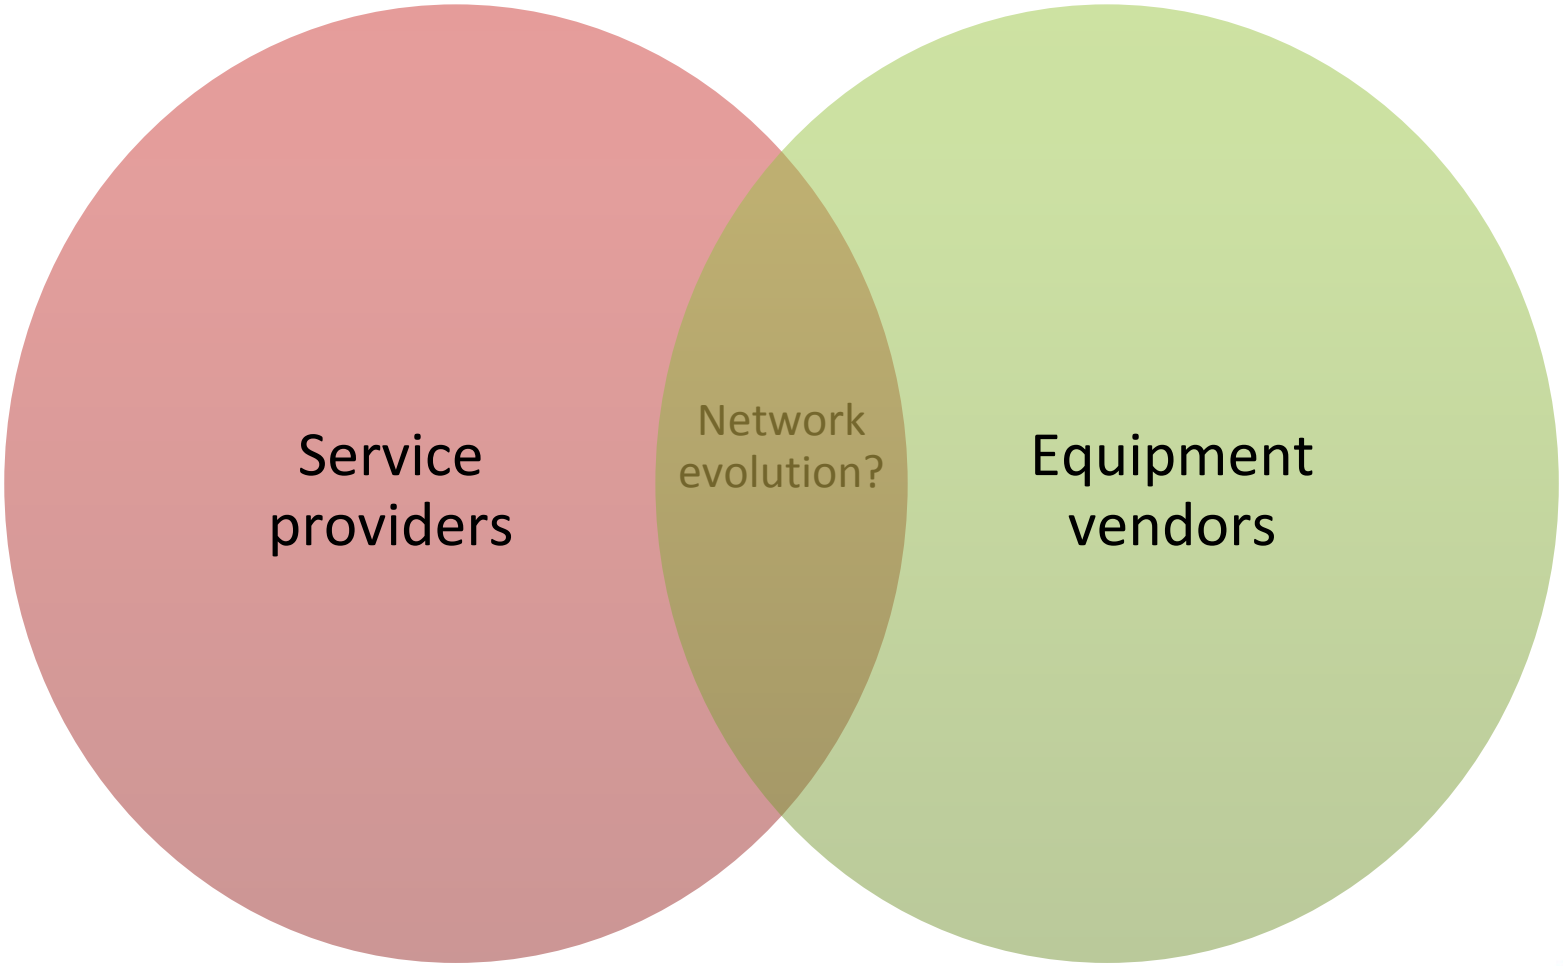

Where to does academia fit in networking development? (II)

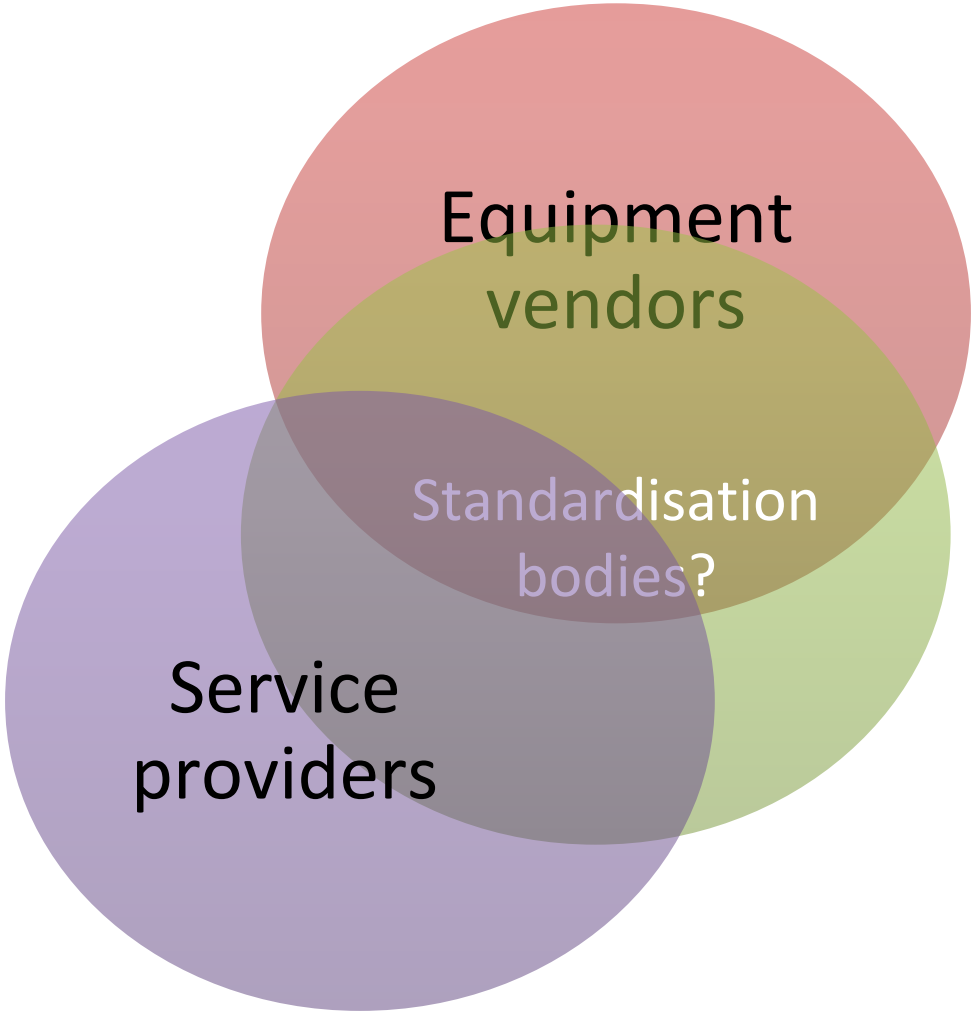

Where to does academia fit in networking development? (III)

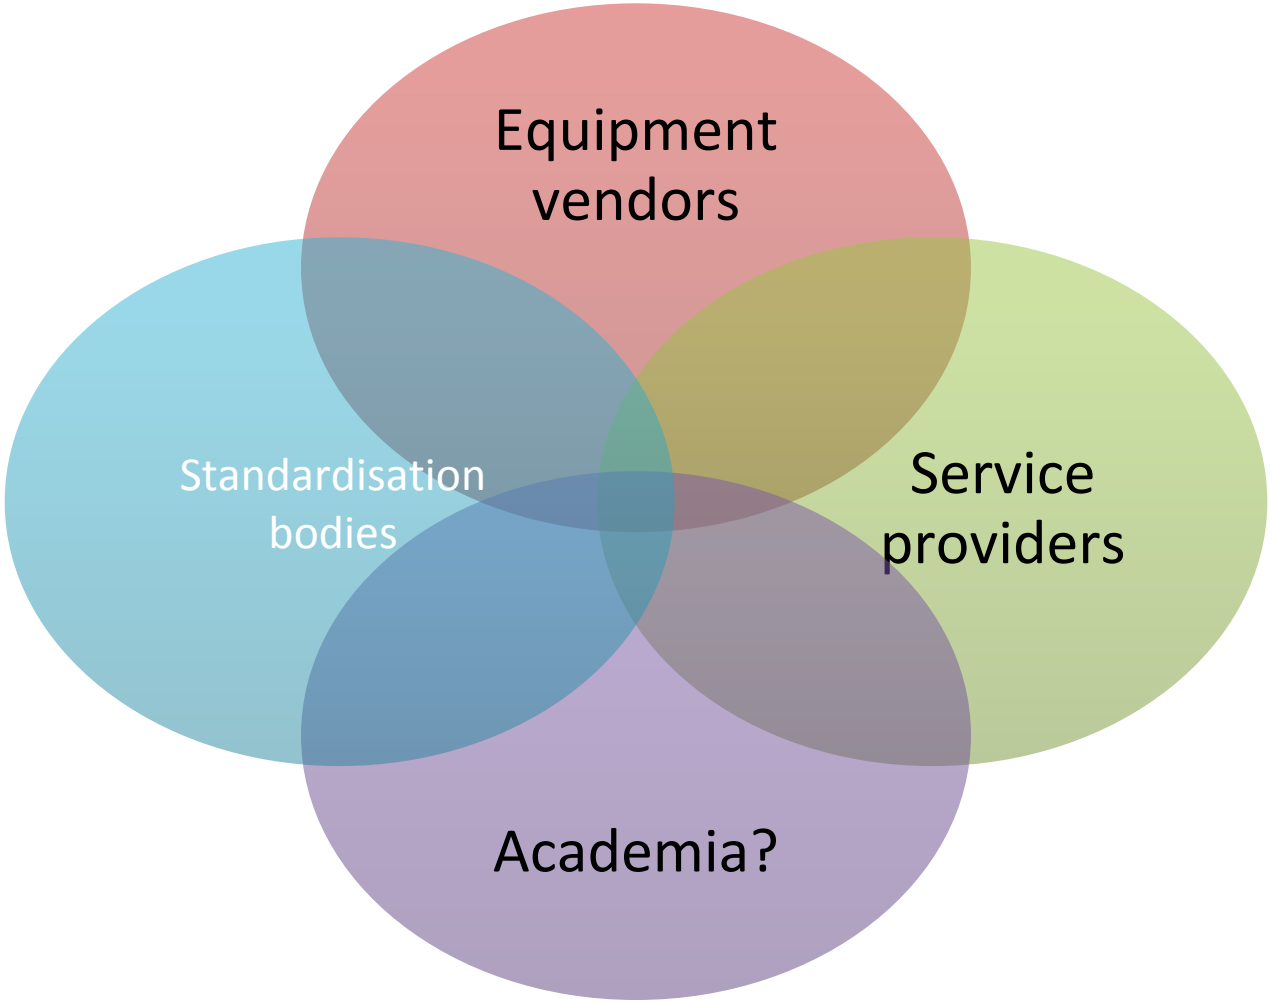

SDN pioneers - a reality check

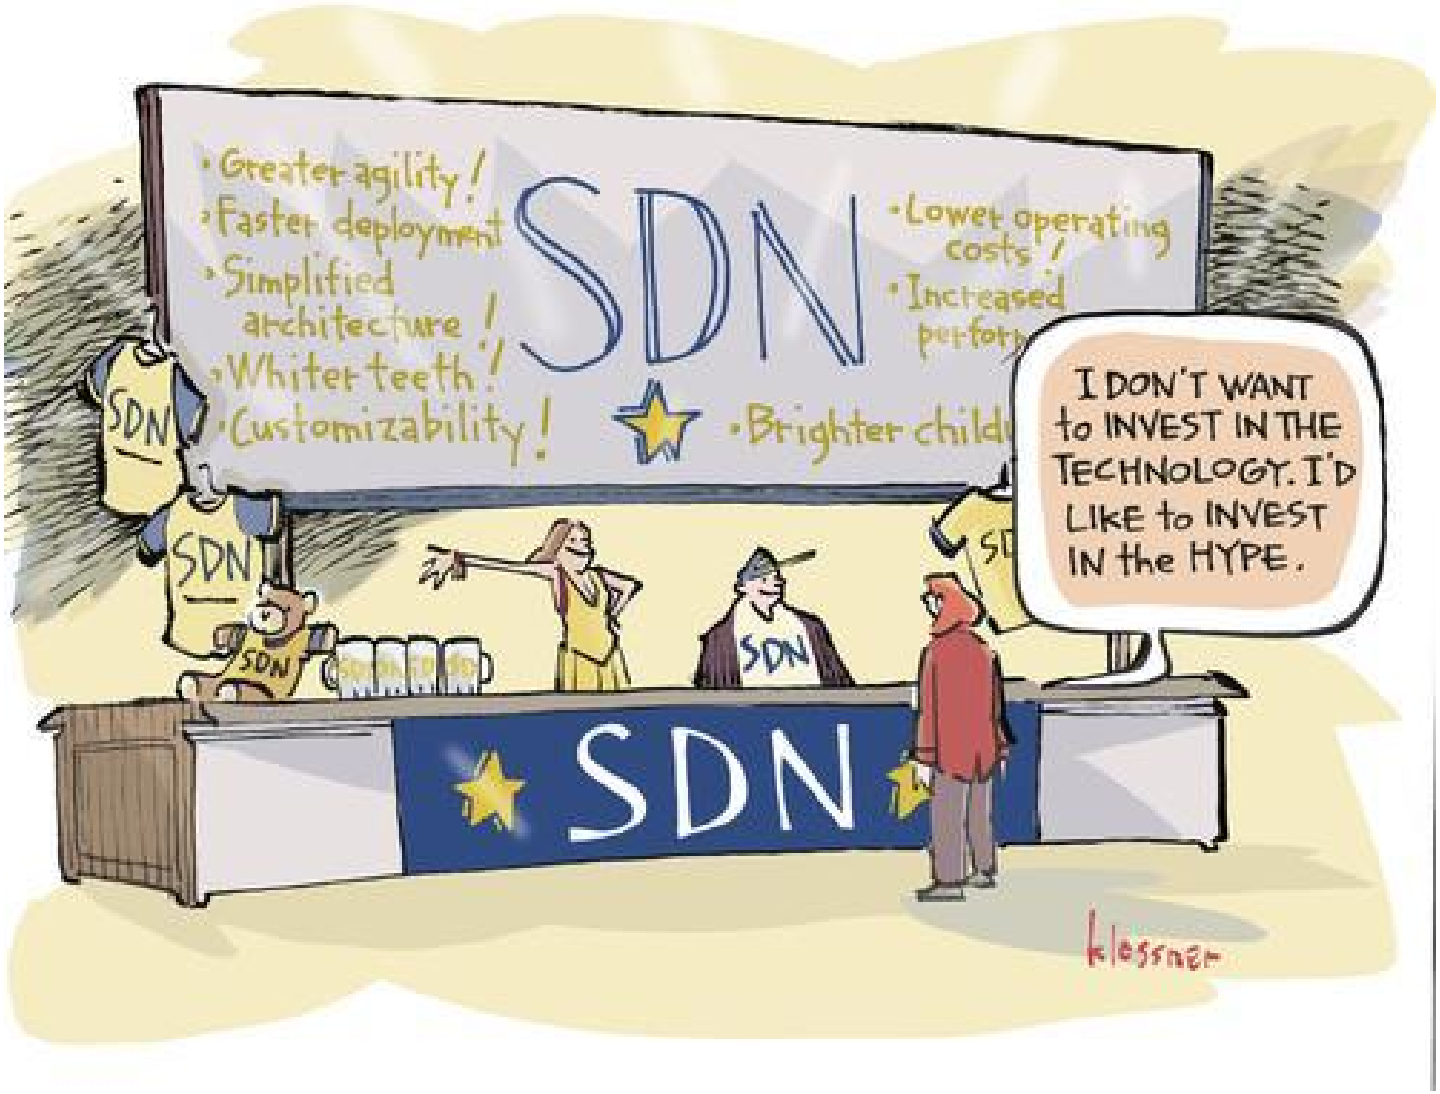

Lancaster University

Academia

The Register®
Biting the hand that feeds IT

DATA CENTRE SOFTWARE NETWORKS SECURITY INFRASTRUCTURE DEVOPS BUSINESS HARDWARE SCIENCE BOOTNOTES FORUMS

Data Centre ▸ Data Networking

OpenFlow controller design killing SDN, say network boffins

15

Packet serialisation means processors waste up 20 per cent of their time

22 Aug 2016 at 00:03, Richard Chirgwin 187

OpenFlow's architecture is inefficient, and caps performance while sucking unnecessary power.

That's the conclusion of a bunch of Comp. Sci boffins from researchers at Australian brain box Data61 and Sydney University, who assessed four major OpenFlow controllers – NOX, Maestro, Floodlight and Beacon. Their paper is at Arxiv.

Poor old OpenDaylight was also tested but not reported: "the performance [was] too low to provide any insightful comparison".

To cut to the chase: none of the controllers tested got anywhere close to line speed, whether running on

More like this

Sdn

Whitepaper Downloads

Agile operations and the three ways

Delivering business continuity for vital applications

Financing the IBM FlashSystem family

Six steps to control the uncontrollable

Solidifying protection.

Most read

Industry forums?
(The BGP router that squawked)

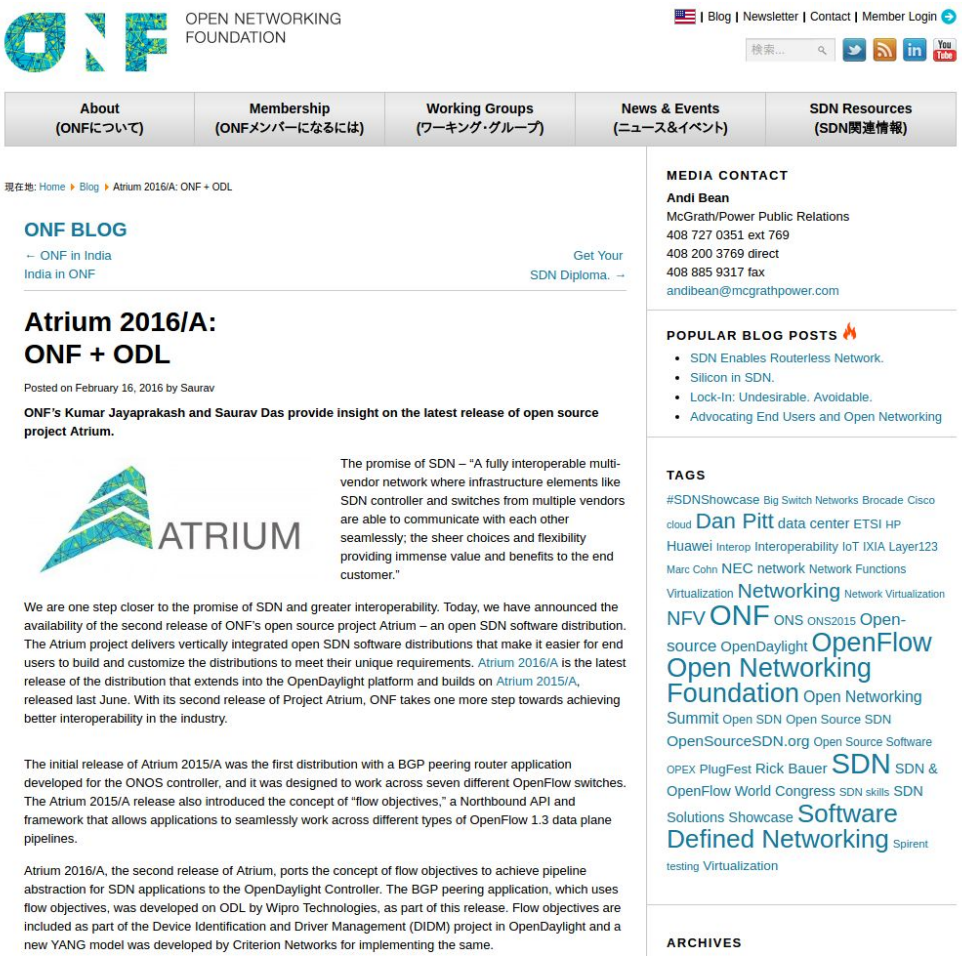

OPEN NETWORKING FOUNDATION

| Blog | Newsletter | Contact | Member Login

About (ONFについて) | Membership (ONFメンバーになるには) | Working Groups (ワーキング・グループ) | News & Events (ニュース&イベント) | SDN Resources (SDN関連情報)

現在地: Home ▸ Blog ▸ Atrium 2016/A: ONF + ODL

ONF BLOG

← ONF in India
India in ONF

Get Your
SDN Diploma. →

Atrium 2016/A: ONF + ODL

Posted on February 16, 2016 by Saurav

ONF's Kumar Jayaprakash and Saurav Das provide insight on the latest release of open source project Atrium.

ATRIUM

The promise of SDN – "A fully interoperable multi-vendor network where infrastructure elements like SDN controller and switches from multiple vendors are able to communicate with each other seamlessly; the sheer choices and flexibility providing immense value and benefits to the end customer."

We are one step closer to the promise of SDN and greater interoperability. Today, we have announced the availability of the second release of ONF's open source project Atrium – an open SDN software distribution. The Atrium project delivers vertically integrated open SDN software distributions that make it easier for end users to build and customize the distributions to meet their unique requirements. Atrium 2016/A is the latest release of the distribution that extends into the OpenDaylight platform and builds on Atrium 2015/A, released last June. With its second release of Project Atrium, ONF takes one more step towards achieving better interoperability in the industry.

The initial release of Atrium 2015/A was the first distribution with a BGP peering router application developed for the ONOS controller, and it was designed to work across seven different OpenFlow switches. The Atrium 2015/A release also introduced the concept of "flow objectives," a Northbound API and framework that allows applications to seamlessly work across different types of OpenFlow 1.3 data plane pipelines.

Atrium 2016/A, the second release of Atrium, ports the concept of flow objectives to achieve pipeline abstraction for SDN applications to the OpenDaylight Controller. The BGP peering application, which uses flow objectives, was developed on ODL by Wipro Technologies, as part of this release. Flow objectives are included as part of the Device Identification and Driver Management (DIDM) project in OpenDaylight and a new YANG model was developed by Criterion Networks for implementing the same.

MEDIA CONTACT

Andi Bean
McGrath/Power Public Relations
408 727 0351 ext 769
408 200 3769 direct
408 885 9317 fax
andibean@mcgrathpower.com

POPULAR BLOG POSTS

- SDN Enables Routerless Network.
- Silicon in SDN.
- Lock-In: Undesirable. Avoidable.
- Advocating End Users and Open Networking

TAGS

#SDNShowcase Big Switch Networks Brocade Cisco cloud Dan Pitt data center ETSI HP Huawei Interop Interoperability IoT IXIA Layer123 Marc Cohn NEC network Network Functions Virtualization Networking Network Virtualization NFV ONF ONS ONS2015 Open-source OpenDaylight OpenFlow Open Networking Foundation Open Networking Summit Open SDN Open Source SDN OpenSourceSDN.org Open Source Software OPEX PlugFest Rick Bauer SDN SDN & OpenFlow World Congress SDN skills SDN Solutions Showcase Software Defined Networking Spirent testing Virtualization

ARCHIVES

- PALO ALTO, Calif., June 9, 2015 - Open Networking Foundation Releases Atrium Open SDN Software Distribution "*Eases Entry to Open Source SDN Adoption; Solves Critical Integration Challenges Facing Today's Network Operators*"
- PALO ALTO, Calif., February 16, 2016 - Open Networking Foundation Releases Second Version of Atrium Open SDN Software Distribution "*Incorporates OpenDaylight, Improves ONOS Version, and Adds Leaf-Spine Fabric*"

Industry? (2014)

The network used to be programmed through what we call CLIs, or command-line interfaces. We're now changing that to create programmatic interfaces," Cisco Chief Strategy Officer Padmasree Warrior said at a press event earlier this year."

Industry? (2014)

PCWorld
FROM IDG

NEWS | REVIEWS | HOW-TO | VIDEO | BUSINESS | LAPTOPS | TABLETS | PHONES | HARDWA

Input Devices | Displays | Printers | Storage | **Networking** | Cameras

Home / Networking

Will software-defined networking doom the command line interface?

1 COMMENT

Stephen Lawson Aug 30, 2013 11:30 AM
IDG News Service

COMPUTERWORLD
FROM IDG

MORE

Cisco ONE is the company's strategy to make networks better understand applications. At its core is onePK (ONE Platform Kit), which will include 710 APIs (application programming interfaces) that developers can use to take advantage of features in Cisco's existing and future network equipment, according to Lloyd. Those APIs will let developers address the US$180 billion installed base of Cisco gear, he said.

Cisco ONE has been called Cisco's answer to SDN (software-defined networking), though the company says it's going beyond other SDN approaches, which focus on separating the control from the transport layer of the network.

"Our vision is much broader. We see the network as a platform," Warrior said. Cisco says its approach allows for more programmability.

The company's ASICs (application-specific integrated circuits) bring another element of programmability, Lloyd said, showing off chips for Cisco's Catalyst 3850 switch and ASR1000 Aggregation Services Routers. Developers will be able to gain access to the software that runs on those ASICs through Cisco ONE, Lloyd said.

"ASICs in the products, with software and services, I think, is going to allow Cisco to really, really shake this industry," Lloyd said.

Industry? (2016)

EVENTS | WHITEPAPERS | TOP 100 CIOS | RESEARCH | SMB SPOTLIGHT

computing

News | Big Data & Analytics | DevOps | Security | Internet of Things | OpenSource | Cloud & Infras

Networks

Cisco to lay off 5,500 staff as it 'pivots' towards software-defined networking

Profits up 20 per cent; jobs down seven per cent

Graeme Burton

@graemeburton

18 August 2016

0 Comments

Cisco has announced that it is cutting 5,500 jobs, or about seven per cent of its workforce - in a restructuring that is intended to shift the company's focus from hardware to software. Some 300 jobs will...

SDN strategy:

“”

Hilton Romanski
SVP and Chief Strategy Officer

Industry? (2016)

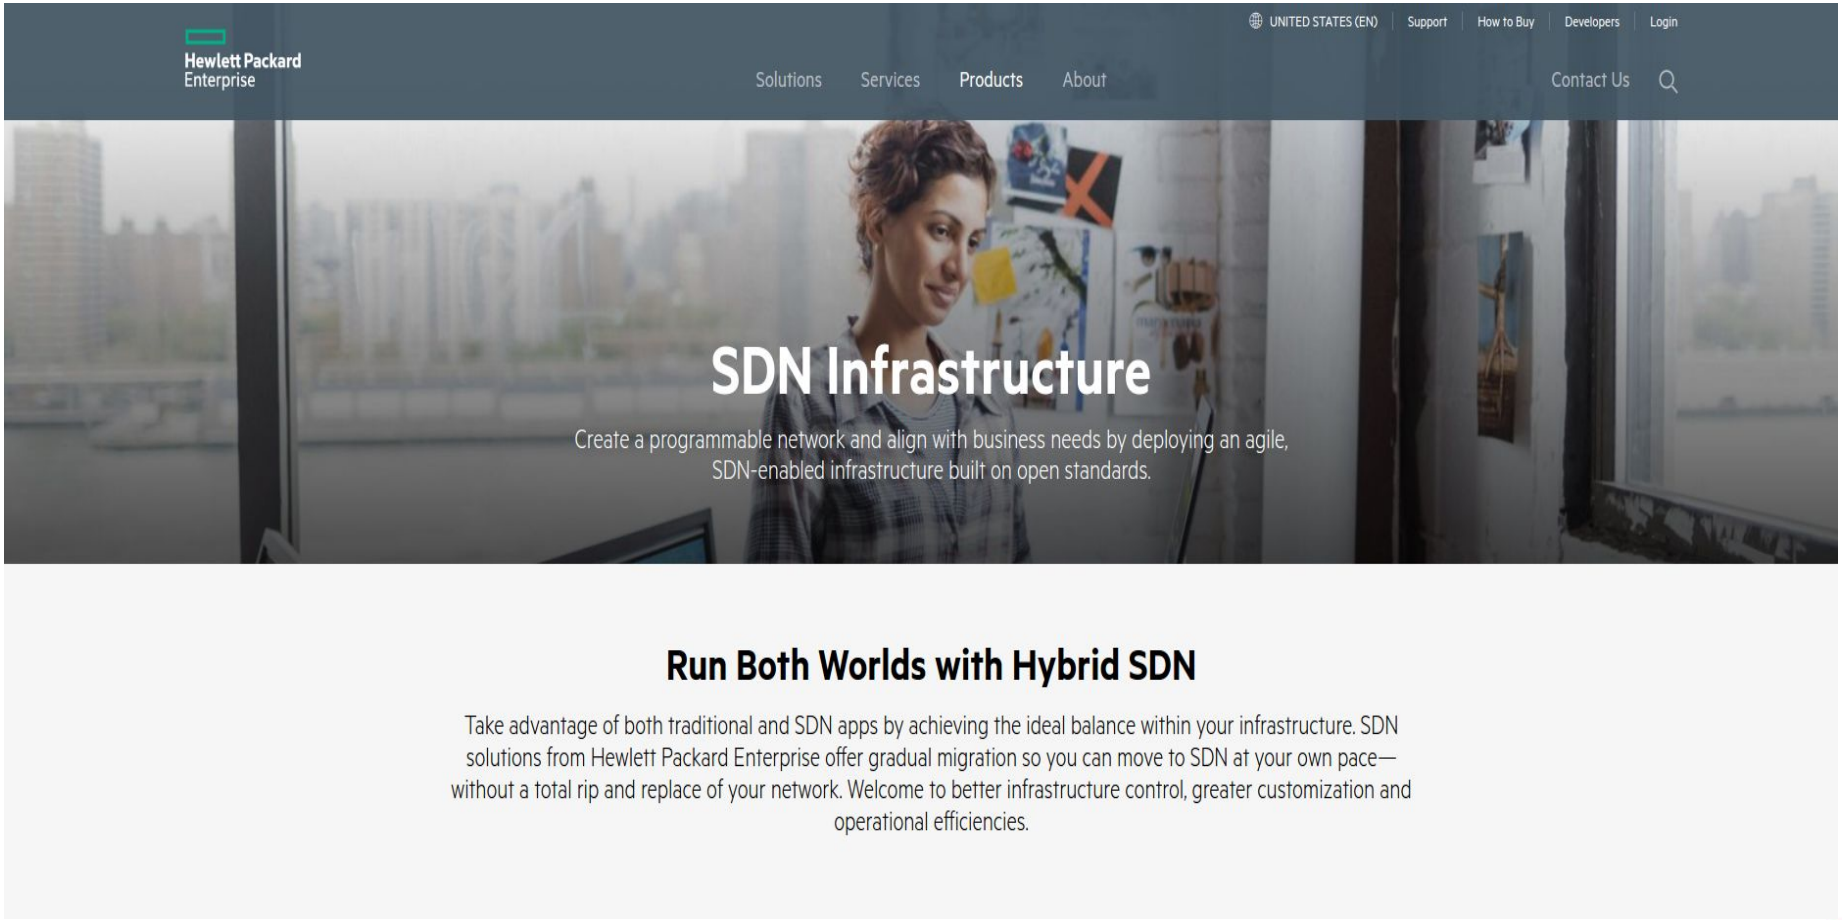

Industry+Academia (2016)

Incremental Deployment of SDN in Hybrid Enterprise and ISP Networks

David Ke Hong*, Yadi Ma†, Sujata Banerjee†, Z. Morley Mao*
* University of Michigan, † Hewlett Packard Labs
{kehong, zmao}@umich.edu, {yadi.ma, sujata.banerjee}@hpe.com

ABSTRACT

Introducing SDN into an existing network causes both deployment and operational issues. A systematic incremental deployment methodology as well as a hybrid operation model is needed. We present such a system for incremental deployment of hybrid SDN networks consisting of both legacy forwarding devices (i.e., traditional IP routers) and programmable SDN switches. We design the system on a production SDN controller to answer the following questions: which legacy devices to upgrade to SDN, and how legacy and SDN devices can interoperate in a hybrid environment to satisfy a variety of traffic engineering (TE) goals such as load balancing and fast failure recovery. Evalua-

ularly true in large scale legacy enterprise networks where it is common to have long lifecycles of network devices. Thus enterprise and ISP network operators resort to incrementally deploying SDN devices in their existing networks and tailor their SDN applications to work in a hybrid environment. For example, many enterprises have resorted to upgrading just their edge devices to SDN for QoS and security related applications. This upgrade strategy may not work well for TE and failure recovery applications.

Unfortunately, we find a lack of strategies from existing work that can systematically compute the best upgrade options for the maximum return of network benefits, as well as operate the hybrid infrastructure in an optimal way. Our

Academia? (2016)

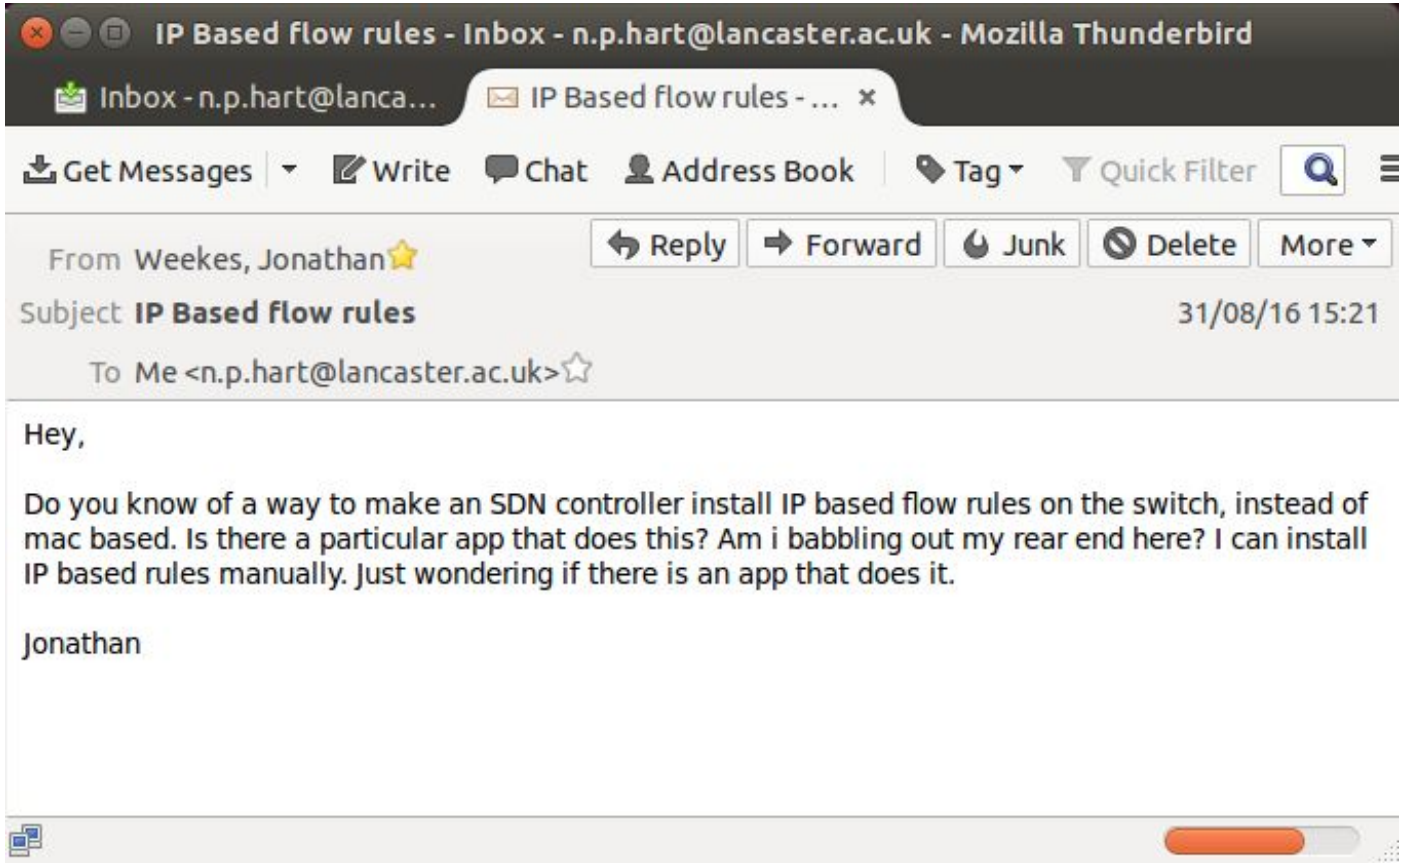

Academia? (2016)

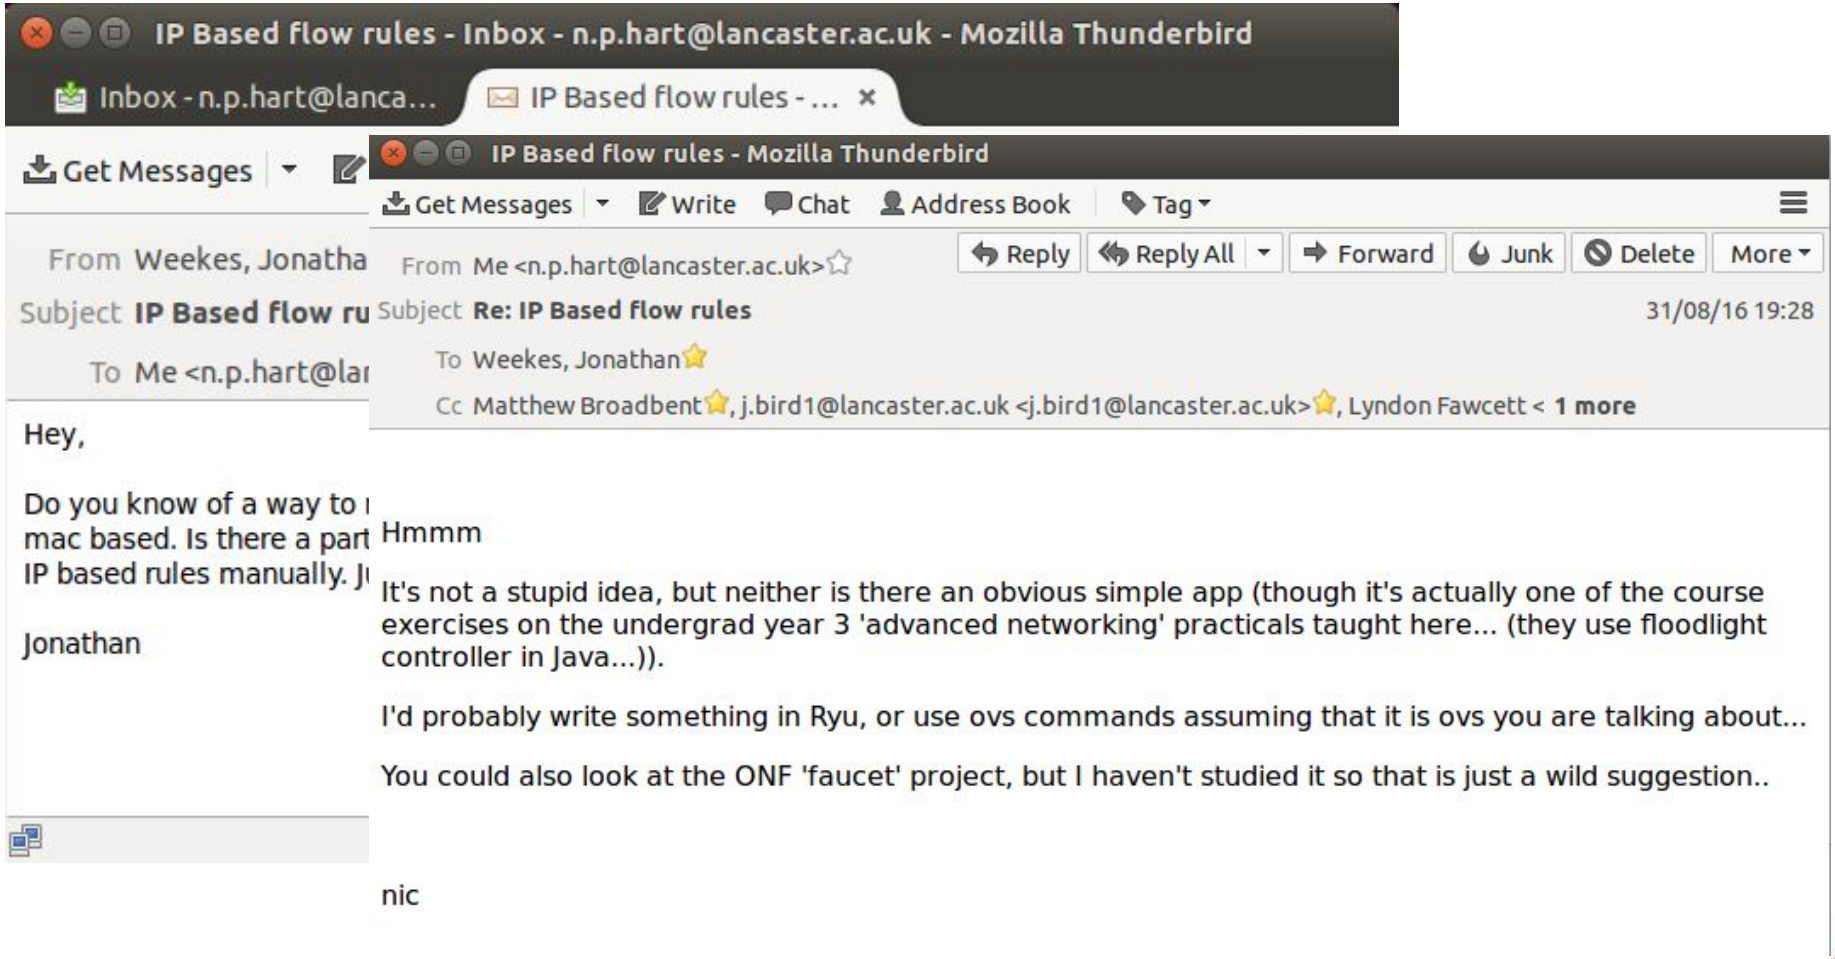

Lancaster University

sBGP: A hybrid SDN approach to interdomain routing

uknof35, 8th September 2016

Nicholas Hart, Lancaster University

Research Context: Current internet architectures

- Terminology: ASes, iBGP, eBGP, peering
- What does a transit ISP do?
- Learn routes (BGP inbound)
- Advertise routes (BGP outbound)
- Simple architecture
  - Edge routers process offered peer routes (eBGP), accept/reject, and propagate to the rest of their peers (iBGP).
  - Edge and internal routers learn and use these routes (iBGP), other edge routers re-advertise the routes learned from the
  - Source edge router
    - route selection - edge router applies 'policy' to decide which routes to accept
    - route selection is complex, AS path is central to the decision
  - Exit route selection - edge routers must choose whether and which routes to re-advertise - this can be based on commercial or operational considerations, not just technical

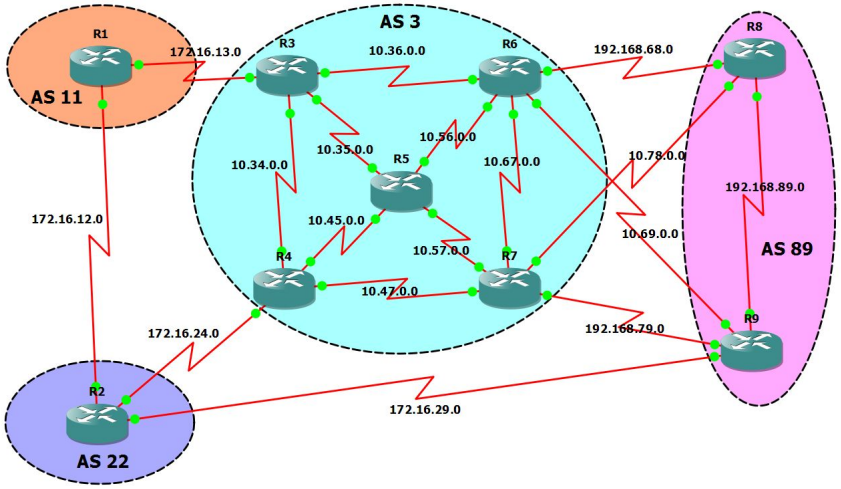

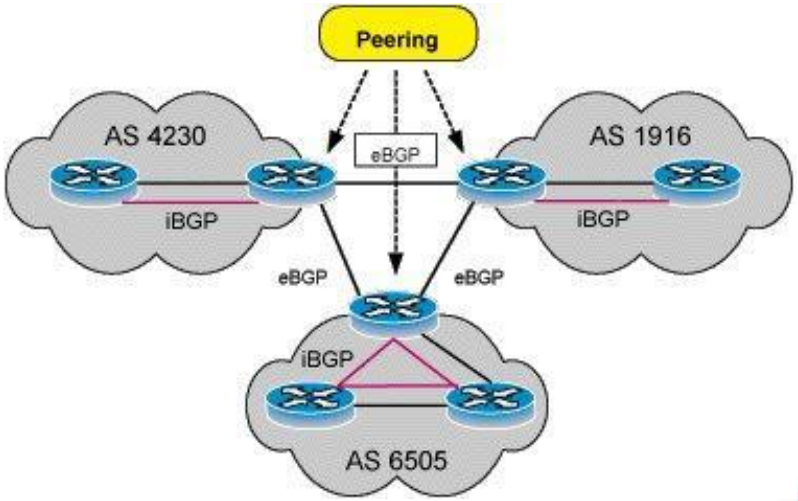

Research Context: SDN introduction

- SDN ancestry
- SDN principles
- SDN architectures
- Controllers and forwarding agents
- Control plane and data plane
- OpenFlow and other ‘south-bound’ protocols

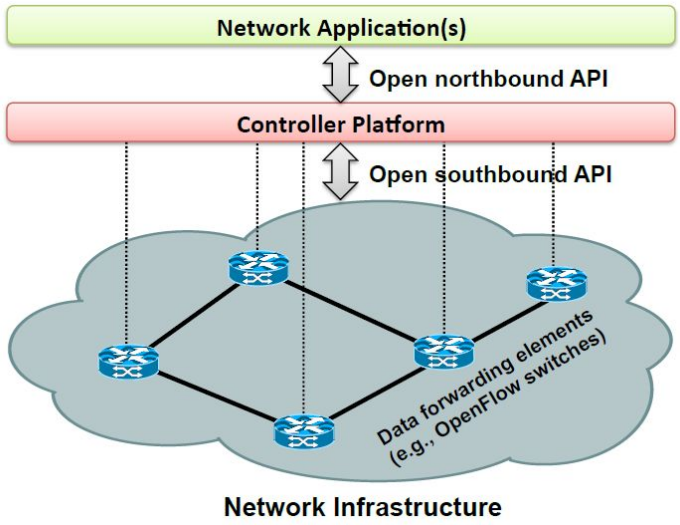

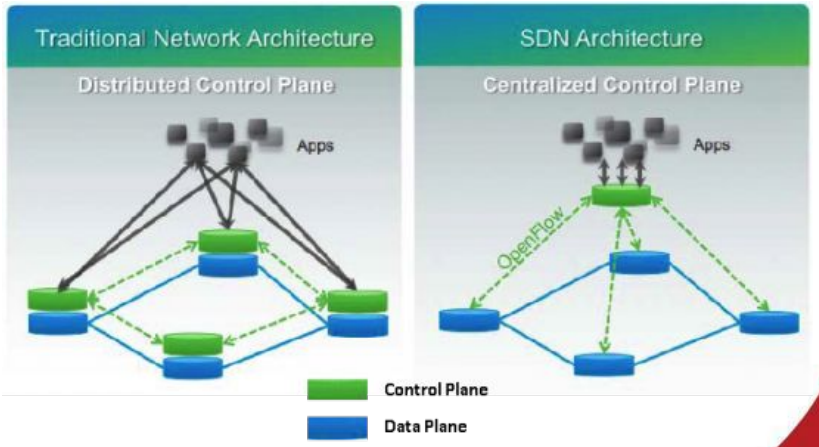

Motivation: IDR and SDN today

- Core internet routing (IDR) isn't benefitting from SDN
  - SDN devices don't have the capacity or performance
  - OpenFlow doesn't have the semantics to define complex behaviours
  - OpenFlow doesn't have the performance required
  - Core routers already have the forwarding capabilities, OpenFlow offers few advantages*
- Core internet routing (IDR) really needs SDN
  - Centralised policy and policy database management needed
  - Policy complexity too high for configuration mechanism
  - Security challenges not addressed
  - Migration to vendor neutral solutions blocked by distributed routing policy implementations - routers as white boxes as a migration strategy
- SDN applications limited by south-bound protocol choices
  - OpenFlow is 'the only game in town'
  - The fast recovery problem, and OAM
  - BGP has already been extended to support many other applications - it is a 'meta' protocol

The sBGP architecture - ‘hybrid SDN’

- What is different? - explicitly - all of the BGP connections between routers have been removed, replaced with BGP connections to a routing control system.
- Logically: the route selection process has been relocated from the edge routers, to the routing control system. BGP is retained for the internal control, but with reduced functionality. It could be replaced in future, e.g. by OpenFlow.
- Consequence: the responsibility for policy is now under software control.

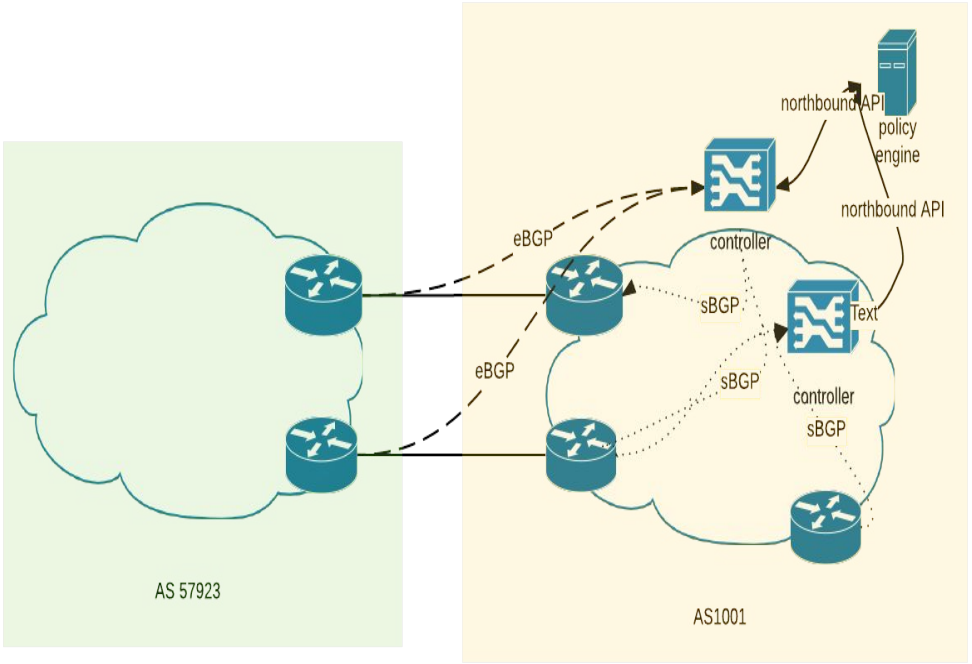

Lancaster University

Research Questions

- Hybrid SDN
  - Does it solve the problems we set? Does it perform, scale? Is it secure, resilient?
  - What is the role of the SDN controller for sBGP? Can we use it in existing SDN/OpenFlow environments?
  - Can it handle today's IDR stress points? (route flaps, erroneous routing data,..)
- Can sBGP match OpenFlow semantics for security applications (BGP extensions? Hybrid OF/sBGP?)
  - Source address and port matching, rate limiting, packet duplication
- Matching evolution of intra-domain transport to centralised paradigm
  - MPLS, PCE, LISP, segment routing

Lancaster University

Industry impact

- Industry challenges
  - BGP configuration is hard, rarely automated, and often mistakes are made
  - Few ISPs implement the current full set of recommendations
  - Routing configuration is used to resolve (d)DOS attacks, usually by hand
  - There is little scope for ‘sense checking’ new routes based on history or central data
  - Attacks on routing infrastructure are increasing in volume and sophistication
- Vision
  - Next generation internet architectures are like nuclear fusion (always 20 years away)
  - My goal is demonstrating near term, real world, feasibility and value

Historical Context - RCP (Routing Control Platform)

- To what extent was RCP successful/adopted/evolved?
- Where it didn't, why didn't it?
- Are the challenges which led to the RCP papers still present?
- Are there other (routing policy) challenges which weren't considered, and would an RCP approach address them?
- Categorise the design space of separated forwarding and routing architectures: are there qualitatively different approaches than RCP?
- Are there any hybrid architectures? do they represent an evolution/migration path to full separation
- Does the evolution of new technologies for inter-domain transport, e.g. PCE, MPLS-xx, LISP, segment routing, change the landscape for separated routing and forwarding?
- Does it make sense to decompose horizontally the IDR problem - i.e. centralise some functions (policy data, policy representation), but distribute implementation?
- Considering problems like security vulnerabilities and mitigation strategies, or optimising traffic, or protecting against unintentional disruption, can the case be made the existing architectures are fundamentally incapable of addressing them?

- Design and implementation of a routing control platform. 2nd USENIX NSDI, May 2005.(M. Caesar, N. Feamster, J. Rexford, A. Shaikh, J. van der Merwe.)
- The case for separating routing from routers. ACM SIGCOMM Workshop on Future Directions in Network Architecture, Sept. 2004. (N. Feamster, H. Balakrishnan, J. Rexford, A. Shaikh, K. van der Merwe.)

Contact

n.p.hart@lancaster.ac.uk

Contact - and an appeal!

n.p.hart@lancaster.ac.uk

I’m looking for industry partners to collaborate with.
From

- Feedback - suggestions, criticisms, questions...
- Simple information gathering (how do you automate BGP configuration - what is the level of effort required - if you could do more, cheaply, safely, what would it be?)
- Sharing traffic data - understanding the profile of routing change, and interaction between ‘policy’ and traffic - analysing critical network ‘events’.
- Evaluating approaches to automated policy implementation

(Obviously...) - confidentiality strictly observed.....