MSN 2016

Nicholas Hart
Coseners, July 2016
Supervisors: Nicholas Race, David Hutchison

‘sBGP: A hybrid SDN approach to interdomain routing’

PhD Studentship Application

Nicholas Hart
June 2016
Supervisors: Nicholas Race, David Hutchison

‘sBGP: A hybrid SDN approach to interdomain routing’

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

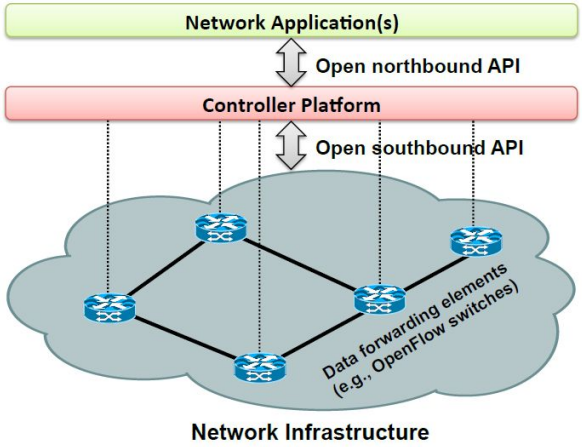

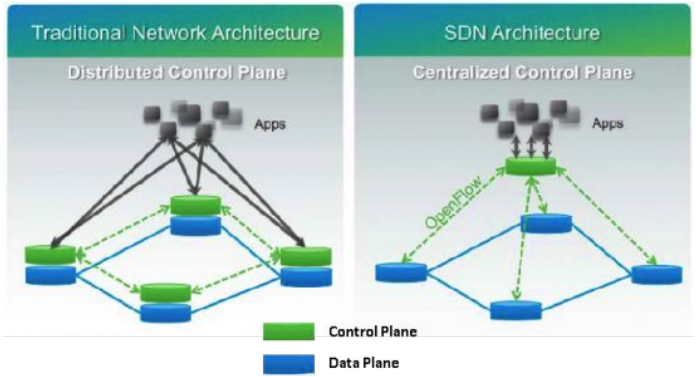

Motivation: IDR and SDN today

- 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.

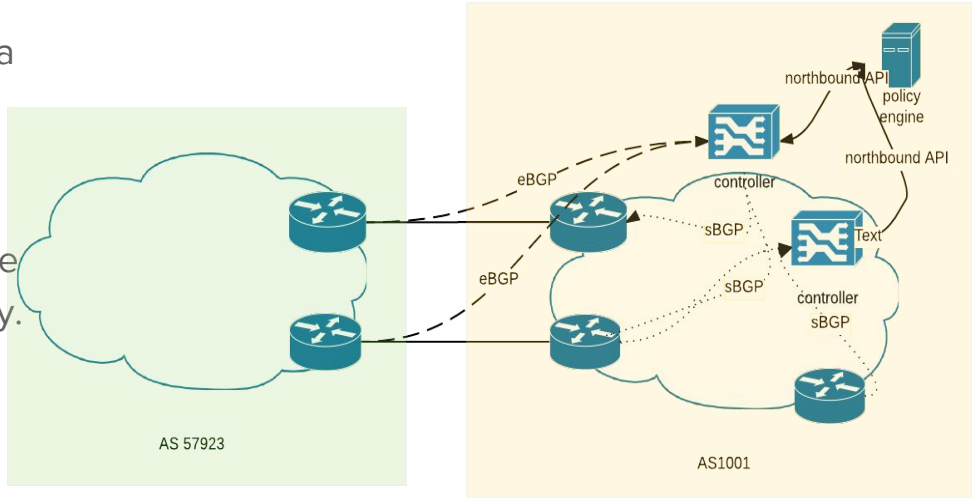

Research questions:

- 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

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

Work plan

- Papers and venues
  - Coseners (July ‘16)(accepted), UKNOF 35 (Sept ‘16) (accepted)
  - Hybrid SDN, sBGP - workshop paper (sigcomm,conext,infocom)
  - Survey - Core Internet - challenges and opportunities - journals - IEEE....
  - DSLs for routing policy - network management venues....
- Architecture, Design, Implementation & Evaluation
  - BGP testbed
    - Already partially complete -Juniper M320, Cisco 7404 , SDN based, Quagga
  - Functional implementations
    - Extending existing controller (ONOS) to support sBGP
    - Implementation of client side sBGP (Quagga/netlink)
    - Implementation of policy DSL and northbound SDN API for route selection
  - SP collaborations: JANET, TNP, B4RN, a medium to large transit ISP
    - (already initiated, except the last)
- Write-up

Evaluation

- Implementation based
  - Quantitative
    - Demonstrating performance parity with existing technology for e.g. convergence, re-routing
  - Qualitative
    - Functional equivalence, with simplified/centralised configuration
    - Additional functionality, e.g. automated DDOS mitigation, intelligent route selection, route testing
    - Evaluating richness/completeness of policy DSL for expressing real-world policy based on industry partners inputs
    - LIke for like studies using BGP and OpenFlow
- Modeling based
  - Model implementation based real world data
    - Validate correctness
    - Evaluate effectiveness of better policy routing

(supplementary slide) The Research Topic:

‘sBGP: A hybrid SDN approach to interdomain routing’ - An SDN based architecture for Internet routing, based on BGP as the principal southbound interface protocol.’

Deconstructing the thesis topic:

SDN - (re-)programmable forwarding devices, open/standardised protocol - OpenFlow

hybrid SDN - keeping the best of the classical network switch/router capabilities

interdomain routing - internet glue - basis of the Internet

BGP - *the* interdomain routing protocol

southbound interface protocol - the core protocol in any SDN system

(supplementary slide) 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

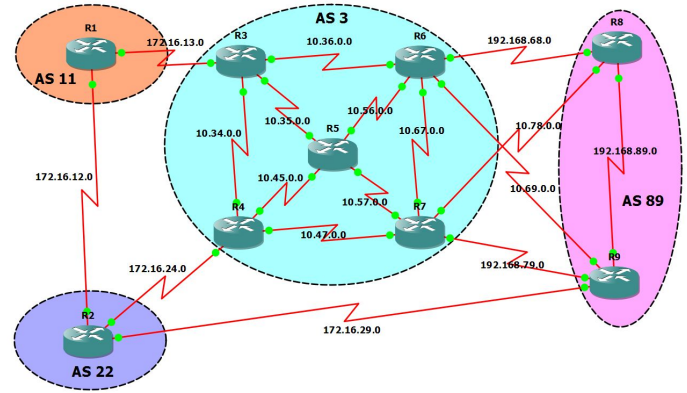

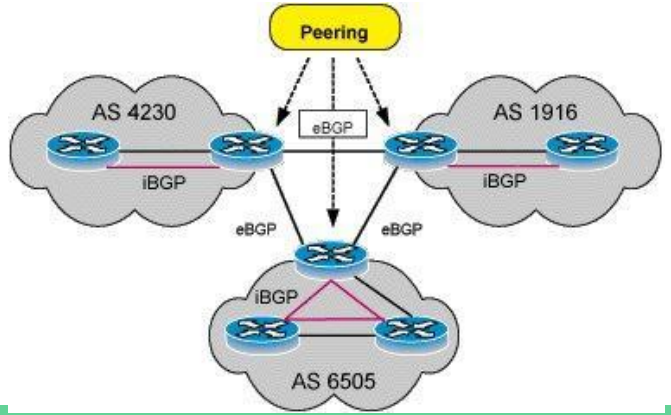

(supplementary slide) RCP topics:

RCP = Routing Control Platform, see e.g.

- 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.)

- 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?

(supplementary slide) About Myself

- Industry career
- OU and FP
- Lancaster
  - Papers
  - SDN and cloud work
  - TOUCAN
  - Talks
  - TA work
  - Contacts & endorsements