HARDENING SOFTWARE DEFINED NETWORKS

INDIANA UNIVERSITY

JULY 2014

FINAL TECHNICAL REPORT

STINFO COPY

AIR FORCE RESEARCH LABORATORY
INFORMATION DIRECTORATE

■ **AIR FORCE MATERIEL COMMAND**　　■　**UNITED STATES AIR FORCE**　　■　**ROME, NY 13441**

NOTICE AND SIGNATURE PAGE

AFRL-RI-RS-TR-2014-200 HAS BEEN REVIEWED AND IS APPROVED FOR PUBLICATION IN ACCORDANCE WITH ASSIGNED DISTRIBUTION STATEMENT.

FOR THE DIRECTOR:

/ S /
CARL R. THOMAS
Work Unit Manager

/ S /
MARK H. LINDERMAN
Technical Advisor, Computing &
Communications Division
Information Directorate

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. REPORT DATE *(DD-MM-YYYY)*	2. REPORT TYPE	3. DATES COVERED *(From - To)*
JULY 2014	FINAL TECHNICAL REPORT	DEC 2012 – JAN 2014

4. TITLE AND SUBTITLE

HARDENING SOFTWARE DEFINED NETWORKS

5a. CONTRACT NUMBER
FA8750-13-2-0023

5b. GRANT NUMBER
N/A

5c. PROGRAM ELEMENT NUMBER
62303E

6. AUTHOR(S)
Jean Camp, Ross Anderson, Ander Odlyzko, Zhi-Lang Zhang,
Chris Hall, Chris Small, Tim Kelley

5d. PROJECT NUMBER
PROC

5e. TASK NUMBER
ED

5f. WORK UNIT NUMBER
UI

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
Indiana University
205C Bryan Hall, 105 N. Indiana Ave
Bloomington, IN 47405

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)

Air Force Research Laboratory/RITA
525 Brooks Road
Rome NY 13441-4505

10. SPONSOR/MONITOR'S ACRONYM(S)
AFRL/RI

11. SPONSOR/MONITOR'S REPORT NUMBER

AFRL-RI-RS-TR-2014-200

12. DISTRIBUTION AVAILABILITY STATEMENT
Approved for Public Release; Distribution Unlimited. This report is the result of contracted fundamental research deemed exempt from public affairs security and policy review in accordance with SAF/AQR memorandum dated 10 Dec 08 and AFRL/CA policy clarification memorandum dated 16 Jan 09.

13. SUPPLEMENTARY NOTES

14. ABSTRACT
Software Defined Networking (SDN) presents an extremely rare point of inflection which offers the potential to leverage the economics of SDN to harden the network as a whole. Utilizing this inflection point requires security technologies that have two characteristics. First, security technologies must be incentive-aligned for initial adoption. Securing SDN requires designing technologies that provide immediate returns for the early adopters. Compare with BGPSEC, which helps only peers and not the investing organization. We have a demonstration providing risk-aware routing given the previous RIB. Second, the technologies must function without complete adoption. And of course, third, these must be resilient against attack. Compare with egress filtering, which works with ISP adoption. We offer a proof of concept showing herd immunity to classes of DoS attacks with partial adoption by second-tier ISP's. Failing to secure next-generation networks risks increasingly vulnerable cyber=physical systems, including homes and even individual persons as the internet of things is diffused to households and surgeries. We focused on six use cases: data centers, then large ISPs, an IXP case, two cyber-physical cases, and the case of the next generation battlefield. The two cyber-physical cases were international airports and industrial control systems.

15. SUBJECT TERMS
Software Defined Networks, Network Security, Security

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON **CARL THOMAS**
a. REPORT U	**b. ABSTRACT** U	**c. THIS PAGE** U	SAR	65	19b. TELEPHONE NUMBER *(Include area code)* **315-330-2600**

Standard Form 298 (Rev. 8-98)
Prescribed by ANSI Std. Z39.18

TABLE OF CONTENTS

LIST OF FIGURES

1   SUMMARY

Distributed traffic engineering requires mutual authentication of components. This is still not a solved problem in the context of client/server and router/router interactions. DNSSEC is being deployed, and BGPSEC is being standardized. However it's not at all clear if there are adequate deployment incentives for BGPSEC as it's currently conceived, or whether it can cope with all attacks. As Software Defined Networking (SDN) islands emerge, first in data centers, then at IXPs, then in corporate networks, and then in consumer ISPs, the solution of how to securely link them may give us a once-in-a-generation opportunity to revisit the assumptions underlying today's networking.

Confounding this are the concurrent changes in network transparency with the simultaneous (and completely uncoordinated) movements from IPv4 and BGP to a mixed networks with IPv4 and IPv6 both routed and switched by BGP and SDN. BGP is designed to be an information hiding protocol; it lets ASes conceal sensitive business information from competitors with whom they peer. IPv6 removes the information and device hiding that is an inherent result of NATs. Yet transparency is good for detecting malicious behavior. For example, observations of traffic across borders can provide an indication of the sources of malicious traffic. SDN will centralize control information and may thus force a rethink of transparency.

There is also a risk that SDN itself will be adopted in an insecure manner, repeating and exacerbating the vulnerabilities of today's network. Investment in security lags; for example, only one SDN switch supports even basic TLS. Issues of security network management applications, evaluation of complex interactions, and the inseparable economic and usability questions are unexplored.

This report enumerates both the potential of SDN and the challenges in meeting that potential. What are the challenges, threats, potential, and overall implications for Software Defined Networks (SDN) in terms of creating a resilient network? To answer this question, we have created clear threat models grounded in documented and realistic use cases; enumerated the resulting authentication requirements; modeled the next-generation network to evaluate authentication interactions; and finally constructed a prototype that demonstrates practical forward movement to meeting these challenges.

The research project associated with this illustrates substantive proof of progress in terms of securing SDN and leveraging SDN to increase BGP reliance, as well as moving forward network modeling to identify emergent concerns before these become practical problems.

In summary, this report describes the path and achievements under FA8750–13–2–0023. The initial approach to resilient SDN was domain-specific. In order to create a comprehensive view of future attacks and thus emergent security requirements, we began the project with a set of critical use cases (as noted above and included as appendices). As the critical use cases matured, the threat model reified, and the project moved directly to a vision of the security challenges as existing at distinct layers rather than being differentiated by domain. The result is a comprehensive description of both the threat models and the authentication requirements for reliable network management. A critical part of this threat model is the recognition of the need for a management plane to leverage all the potential of SDN for simultaneous simplification of network management and assurance of network reliance. Thus we identify threats on a data plane (isolated from but informing the control plane), the control plane, the management plane, and the human layer where final decision-making authority resides.

2 HARDENING THE SOFTWARE-DEFINED NETWORK

2.1 Introduction

Many of the vulnerabilities in the current Internet, especially those associated with BGP, can be attributed to the complexity of network management in the existing IP networks. Currently network management requires distributed management and manual configurations of individual devices, often using low-level commands or scripts.

The current coexistence of the control and data planes on the same network amplifies such vulnerabilities. For instance, the misconfigurations of BGP results in routing failures, as in the myriad examples offered below. The combination of data and control causes an invisibility of physical risks and choke points. Conversely, the necessary openness for configuration of individual devices results in high levels of network device information availability for attackers. (For readability, current networks which use BGP, eBGP, etc. are all referred to as "BGP" networks in this report.)

One solution to this complexity is Software Defined Networking (SDN). In Software Defined Networks, the control and data planes are isolated, making distributed traffic engineering far more feasible. This separation enables a level of abstraction that can resolve many current network vulnerabilities. SDN has the potential to resolve long-standing challenges in authentication in routing and traffic management (mitigating outages and denial of service). In traditional security terms, SDN can address authentication, integrity, and availability. The move to SDN presents an extremely rare point of inflection which offers the potential to leverage the economics of SDN to harden the network as a whole. The demonstration implemented under this contract proves SDN can be secured, and when secured in an incentive-aware manner, simultaneously improves reliance in the BGP network to which it connects.

The use cases we examined were the currently most common case of data centers, then large ISPs, based on this an IXP case, then two cyber-physical cases, and the case of the next generation battlefield. The two cyber-physical cases were international airports and industrial control systems (i.e., SCADA). These cases are included as appendices. These cases force us to consider SDN in the broadest possible future use, resulting in a layered approach to authentication and security requirements. This approach is embedded in the following final report, with the data plane, control plane, and the necessary but rarely examined management plane each considered. We address how to model these distinct interacting layers, such that the models address all layers from physical to human. In particular, modeling SDN as opposed to pure BGP requires models that can embed the ability of the layers to act upon each other in very distinct ways. Examining the literature, we selected bipartite and tripartite network models are those which can address the emerging properties in a SDN/BGP network structure. We drew particularly on biological models, where components are both tightly integrated and highly specialized.

To illustrate that functional forward movement is possible in this practically framed and theoretically well-grounded approach, we concluded with the construction of an open source network component embedding some of our findings. This component consists of two primary elements, the first a re-instantiation of Quagga to provide a route information base (RIB) constructed from BGP updates. The second component, named Bongo (to align with alcelaphine and equine naming traditions), translate the RIB to a flow information base. By generating a flow information base (FLIP) rather than responding to queries individually, Bongo provides rule compression preventing various forms of controller DoS attacks described below. In doing this, Bongo inherently implements no-cost egress filtering, addressing the problem of today's amplification attacks. Bongo is further designed to enable arbitrary management constraints on flow construction, for example, rejecting a route if a particular AS is to be avoided in transit.

The designs provided here address both first order incentives (for diffusion) and second-order implications, thus illustrating the potential with a network where each individual acting in its own best interest creates herd immunity to certain classes of attacks.

2.2   Methods, Assumptions and Procedures

Software Defined Networks (SDN) offer a fundamentally new approach to traffic engineering and routing, creating more static flows over relatively stable routes, enabling quality of service, and potentially offering superior network utilization. Yet the promises for superior traffic engineering and security are inherently interdependent. Dependable traffic engineering requires not just the resolution of conflicts between resource requests, but mutual authentication, trustworthy platforms, trustworthy information, and trustworthy data. A system without adequate trust engineering can be neither reliable nor resilient, much less optimally managed.

The control plane is and has changed. The hourglass concept, with a single orderly flow of packets through the universal single IP layer, is outdated. Multiple cooperative and competing layers, content distribution networks, and business-based routing would have changed this model without the introduction of SDN. The changes to the control plane require empirical evaluation, building on the emergent science of cybersecurity, including taxonomies and algorithms to identify that which is meaningful and/or predictable. Multidisciplinary and interdisciplinary methodologies are needed to understand the current evolution, and guide against the worst case. The emergence of SDN makes the timing of this research critical. SDN provides an inflection point, a moment of opportunity to reach the best case for emergent and legacy networks.

Software Defined Networks differ from current networks in two main ways. First, they separate the data and control planes, implementing out-of-band control for packet-switched networks. As such, SDN offer many potential advantages for network resilience, control, and management. Second, software defined networks move control (and thus risk) from the expensive optimized hardware of routers to software running on commodity hardware.

Figure 1 Definitions; ME – Management Element, CE – Control Element, RE – Routing Element, FE – Forwarding Element
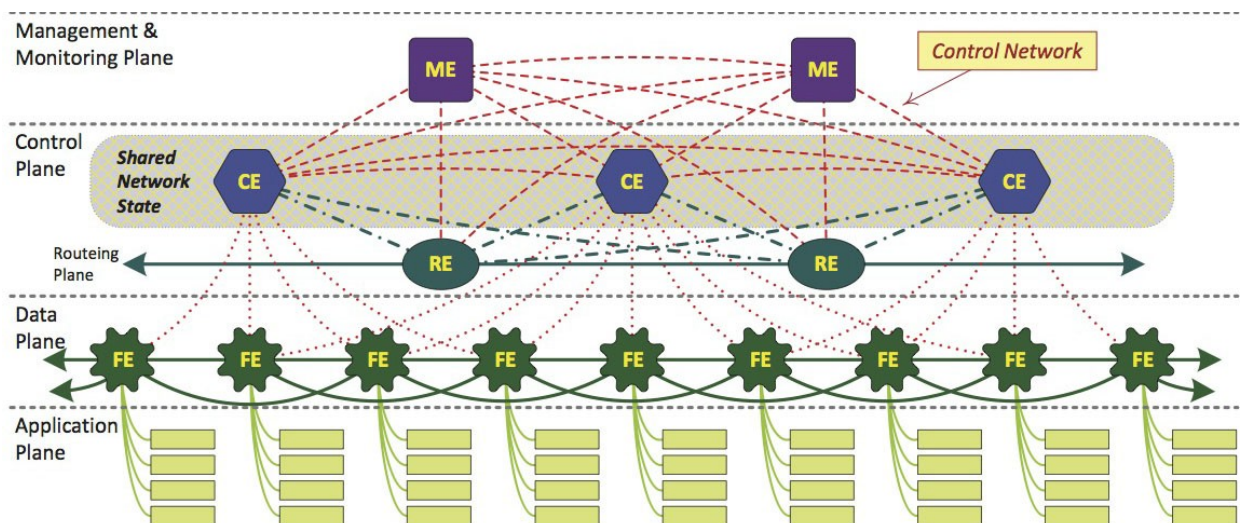

Figure 1 The Northbound Interface to the Management Plane and Southbound to Data Plane;

Yet almost every advantage can also be an opportunity for an attacker. For example, a single SDN controller greatly simplifies network management, but simultaneously risks a single point of failure. Changing the economics and mechanics of routing changes the economics and mechanics of security and network resilience. As SDN and IPv6 concurrently diffuse, new forms of unpredicted cascading failure may also arise.

Today, SDN is primarily adopted by data centers, each within a single organization and in clearly defined physical locations, connected with identified data links. The next step in SDN diffusion is likely to include Internet exchange points. These have a single organizational trust structure but are places where many network operators host their equipment: consumer ISPs, transit providers, CDNs, and even intelligence agencies. If SDN diffuses into inter-exchanges, it is likely to expand down the network hierarchy into homes and devices. While SDN may exist at multiple levels in the network (from home networks to Internet exchange points) the focus here is on the risks at the higher levels.

There are many flavors of SDN, but in this report we focus on one particular implementation, OpenFlow. OpenFlow instantiates paths through the network as deterministic flows; the path selected by a controller from a point in the network will be the same for all switches in the location. That path is determined by the switches' controller; and the packet forwarding is implemented by switches based on controller rules as opposed to being optimized by routers. Flows are not circuits; flows from one switch to another can be determined by a single controller (no negotiation between controllers is required) and there are none of the quality of service guarantees associated with virtual circuits. OpenFlow has succeeded in the marketplace due to its early availability on merchant silicon switches and its greater programmability for applications (such as traffic engineering).

SDN has succeeded because of price. The Amazon cloud networking from Cisco, using custom routers and BGP, would have cost almost a billion dollars. Yet the SDN alternative is only eleven million. Orders of magnitude differences in price are the core driver for SDN adoption; thus it will happen with or without security.

Yet security is also inherent in SDN, albeit as yet unmet potential. Filtering, resilience against control plane misinformation (or misconfiguration) attacks, and load balancing are at the intersection of traffic engineering and security. But if a controller cannot be secured, the traffic engineering promises cannot be kept.

The next section examines the current vulnerabilities in BGP. To illustrate that these can be mitigated or exacerbated with SDN, these are translated into the various vulnerabilities at the different layers at the new networking paradigm.

2.2.1   The Resilience of BGP

The Internet, as a network of networks, is also a network of trust, instantiated in the practical motto "send conservatively, accept liberally". A most clear example is the updating of router tables based on unsubstantiated announcements from other networks. Thanks to this trust, the control plane can be extremely responsive to failures, and recover quickly. The tragic attack on the World Trade Centers did lead to horrific, major loss of life. As an aside it also resulted in the loss of more than three million data and three hundred thousand voice (recoverable) circuits, resulting in a 6% loss of American connectivity. In addition to the physical destruction of switching locations, there were cascading failures from power loss [59]. Yet the Internet, through accepted updates, maintained connectivity. Much of this resilience was the result of engineers trusting each other, and executives willing to forego negotiation before connection. Yet as networks become more automated, that response is not necessarily

reproducible, as earlier failures have not been subject to complete analysis [9].

The very trust that enables resilience can lead to failures when principals lack competence or benevolence. The current BGP-based Internet has proven simultaneously resilient and fragile. It is resilient in its ability to resist, adapt to, and recover from attacks and failures. It is most fragile in the nature of these recoveries: recovery depends on coordinated ad-hoc human responses, as critical as knowing who to call to run a fiber across an exchange center or shut off a machine. The Internet control plane has failed to be resilient to some classes of attacks, such as route leaks, denial-of-service attacks, and botnets.

Some major failures have been caused by network configuration errors. China Telecom announced 15% of all IPv4 space in April of 2010, resulting in loss of traffic for 18 minutes. Some commentators thought this could have been a cyber-war exercise, rhetoric escalated to the 'testing a cybernuke' level of (potentially dangerous) discourse. However, most observers accepted China Telecom's explanation that it was an error. Given that the traffic did not reach its intended destination, it would have been a very clumsy attack. It would also have been erratic: Level 3 and AT&T were notably different in their response [38].

Another route leak that denied service at scale was the misconfiguration of a small Australian ISP in 2012, which took Australia down for hours after it announced all the routes from two larger ISPs to each other [67]. This is the most common routing failure: a straight-forward failure of human factors, which vendors blame on operator error and operators blame on poorly-designed and error-prone control interfaces. The ISP which suffered the outage can be blamed for not filtering route announcements appropriately. There are open economic questions around liability: do such outages cause users to switch ISPs? What's the optimal level of route filtering, for each ISP and for the Internet as a whole? As the outages that result from route leaks are immediately apparent and repaired within hours, such questions rumble in the background rather than becoming major industry issues. SDN may change this game.

Malice can also change the game. When an entity misrepresents its location in a path, rather than claiming to own a destination, the errors are less obvious. In this case, traffic continues to be delivered and such a routing configuration could remain stable for long periods. This occurred again between China Telecom and AT&T, this time for a period of some months [43]. In that particular case, Facebook traffic was routed through China. Note that while the login to Facebook is protected by TLS, no updates are encrypted. Thus a significant amount of global traffic was routed via a nation where Facebook adoption is remarkably low. Such attacks have become more common, for example, the report of massive rerouting through the Ukraine and the Scandinavian route subversions in 2013![26].

In addition to errors and odd incidents, there have also been a wide range of political attacks. The most famous is from Pakistan in 2008, where Pakistan objected to several YouTube videos sufficiently to block all of YouTube. An internal address for YouTube was intended to be announced within Pakistan Telecom; however, it was broadcast across northern Africa and Europe, leading to a service outage lasting several hours [39]. It may be argued that the intention to block within Pakistan was a sovereign political decision, and the leak itself was a human factors problem: a blunder rather than an attack. Yet Internet blocking incidents during the Arab Spring have often been seen as attacks by governments on their populations, as in the case of the rapid drop-offs for Egyptian and Libyan populations [26, 10]. In fact the earliest political Internet blocking may have been during the Serbian atrocities during the dissolution of Yugoslavia, and the subsequent war-crimes trials of Serbian leaders allow us to unambiguously describe such actions as 'attacks'; similar arguments can be made in the case of Libya and Egypt.

The most straightforward way to limit network access is destruction of the infrastructure, as shown in cable cuts in January of 2008. These attacks on the physical infrastructure near Egypt effected large parts of the world in terms of reliable service, but reachability was generally maintained during the ten days of repairs [56]. The fragility of the information infrastructure in terms of the physical infrastructure was also highlighted by the disruption of network traffic to Armenia [58]. An elderly woman, surviving by scavenging scrap metal, discovered what could have been her finest meal when she found a large copper cable close to the surface. That so much of the nation's connectivity depended on this one cable was not apparent until she sliced out a few meters.

It is a general problem that the opacity of BGP makes it difficult to understand the redundancy, or lack of redundancy, in a network. Another example was the Buncefield incident in the UK where an explosion at a fuel storage depot destroyed a number of fibers leading to surprising network outages where the primary and secondary network connections of firms and hospitals had been routed through adjacent fibers without anyone being aware. There are so many layers of subcontracting and outsourcing in the telecoms business that tracking the physical infrastructure on which a system relies is both difficult and expensive. The visibility provided by the abstractions in SDN may make such weaknesses more visible. Besides failures of data and failures of updates there is one notable failure of software which points to another potential vulnerability. For the better part of an hour in August 2010, BGP routes read an attribute of a RIPE-announced address such that Cisco routers interpreted this as a command to drop the route [62]. Details of many of these failures, and on BGP resilience generally, can be found at [15].

2.2.2 SDN Resiliency and Security: Modeling And Analytics

Modeling emerging networks offers guidance to network designers in terms of resilience and potentially offers the ability to identify emergent threats before these emerge. Abstracted network models are used in a wide range of domains including biology, cognitive science, physics, and infrastructure to identify characteristics of multilayered networks. Bipartite and tripartite models are potentially most promising (and somewhat underutilized) in the modeling of emergent networks. Specifically, these models can offer insight into the robustness of networks under arbitrary failures. Bipartite and tripartite models are particularly well-suited to a confluence of traditional networks and software defined networks where SDN components are instantiated on shared hardware. There are preliminary results on a simple topology showing the ability to model cascading failures. A straight-forward extension of the bipartite model into a tripartite network representation of a simple software-defined network (showing controllers, switches and data as separate) is shown. This enables first order modeling of cascading failures when there exist virtual topology with interdependency in the physical network (e.g., multiple switches on one physical box).

One of the key characteristics of a network structure is how gracefully it fails under different conditions. Current network architecture implementations have been constructed and evolved to a network that has been operationally robust against failures and errors, but less resilient against concentrated attacks. Indeed some argue that networks fault tolerance has resulted in a vulnerability to targeted attacks [18, 19]. However, there is skepticism about the applicability of these results, as they are based on topological models and fail to capture many of the real-world failure and recovery dynamics on these networks [68, 49].

Static measures of failure are captured in attack and fault tolerance metrics [5]. In these regards, the physical structure of the Internet and the virtual network on top of its structure, the WWW, as currently defined, have similar topological properties [23]. There has been considerable effort in mapping and understanding the topological structure of the Internet and the World-Wide Web (WWW)

in order to understand its resiliency in the face of attacks and random errors [33]. The topological properties of the Internet and WWW had made them robust against random faults, but potentially susceptible to targeted attacks [19, 46].

Moreover, early analysis of resiliency failed to consider the dynamics of failure. One way of measuring the dynamic properties is cascading failure. A cascading failure takes into account additional network failures that arise when a small initial failure triggers complex interactions [52, 22].

Analysis of the effects of cascading failures has been used to examine structure in metabolic networks [2, 35] and electrical systems [28, 37, 25]. Effects of cascading have also been modeled in the study of communication networks such as the AS-Level topology of the Internet [24]. Studies of cascading failure should incorporate both the topological and dynamic properties of the system. If done well, these models suggest possible methods of mitigating risks of cascades [54]. Ideally, modeling cascading studies can make these failures less likely, and recovery less difficult.

Only recently has the study of network resilience, both static and dynamic measures, begun to examine the effects of interdependent networks. Much of the work in cascades in networks of networks has been done in the physics community [12, 31, 44].

There is also work in communication [69], power infrastructure [11], and travel networks [14]. The work in interdependent and multi-mode networks shows that analysis of aggregated unipartite or projected unipartite networks does not correspond to the more refined method of multinetwork analysis [14, 36].

Few of these works, however, incorporate specific attributes of the physical systems being modeled. While ignoring specific qualities can be useful for simplified, abstract statistics and dynamics of a network topology, topological data alone has been found to be inadequate at predicting real-world cascades in power systems [37]. Both recent and past models focus on different aspects of network robustness. For example, work on abstract networks tends to focus on node failures [29, 69], while work on physical systems looks primarily at link failures [11, 14, 44]. Ideally, for robustness analysis of software-defined networks, we would want to incorporate both node and edge failures in our model.

The separate data and control network structures in a SDN create a network of interdependent networks, which must be modeled as such. In addition there is interdependency with the physical network infrastructure [70]. For example, if one physical machine fails many virtual forwarding elements will fail simultaneously. Data forwarding elements determine where to send data based on information from the control elements. This procedure looks much like a metabolic network where metabolite nodes are connected to reaction nodes. If the metabolite is a reactant, it has an out-degree toward a reaction node, but if it is a product of a reaction, it has an in-degree coming from the reaction, allowing a node to be described by its in and out degrees [2].

There are some differences, as data forwarding elements can only transmit data to other elements they are physically connected to. Also, unlike metabolic networks, the control elements are in SDN share information about network structure. Furthermore, most connections in a metabolic network are one-way, while in SDNs connections can be defined in numerous ways. Still, the analogy is useful; removing a control element means the forwarding element loses the ability to send data to the correct location. Removing a critical forwarding element means that even when control units are aware of where the data should go, they are unable to produce a connecting element, and the techniques for analyzing bipartite cascading failures should still be viable, with additional modifications.

2.2.3 Security Engineering & Economics

Although we are constantly reminded, often through vivid examples, of the pervasive insecurities of all IT systems, the world is doing quite well, for the most part. What is the root of the tolerance of all those catastrophes? It appears that the main explanation lies at the intersection of economics, sociology, and psychology (including a large dose of human factors contributions, of the kind discussed below). Technology is extremely important, but abundant evidence illustrates that absolute security with zero risk is unattainable in practice. As connectivity with perfect security is impossible, economics in its basic form can inform decisions on how much to invest in security and in what forms of security, in balancing costs and benefits.

The BGP network, in particular, has been remarkably successful. Even though it has many widely acknowledged deficiencies, it has been the workhorse of the Internet for several decades, as this network grew several orders of magnitude in terms of either ASes, number of devices attached to it, or traffic. The usual explanation for this success is that it is due to the small community of network engineers, where peer pressure and personal contacts suffice to preserve a network that functions smoothly. That appears to be correct, but the next question is, why would this not continue? The usual response is that the traditional approaches will not scale, as the Internet grows and becomes more heterogeneous. That is also likely largely correct, and we are beginning to see phenomena on the network, such as the recent large scale man-in-the-middle traffic diversions that appear not to be errors, but carefully constructed attacks. But the basic principle of assuring cooperation from network operators could likely be applied in the future, by providing the right tools (such as unalterable and therefore trustworthy logs, together with analysis and dissemination programs) to provide traceability, which is key to holding players accountable for their actions, and thereby inducing them to stick to community norms of behavior. When combined with the right incentives, this approach could induce desirable cooperation from operators.

The other fundamental conjunction of traffic engineering and security economics is transparency of the network. Incentives to provide support for malicious actions will decrease when such actions are more transparent. Current efforts are ad-hoc, and require considerable cooperation. Spamhaus is a canonical example: it is without state authority or enforcement power, but its contacts and reputation enabled this cooperative venture to scale into an effective mechanism for policing smaller sites. Yet it took years for the notorious spam-producing ISP McColo to be subject to refusal to peer [17].

2.2.4 Human Factors

In applications from air traffic control [47] to radiotherapy machines [48], poor usability engineering has proven deadly. Ignoring consistent but irrational human behavior in the design of security mechanisms is itself irrational [53].

An understanding of technology alone has not been adequate for defeating worms, viruses, and malware [7] and this will not change in the management plane. Human incentives have been illustrated to be important, (e.g. [6]), as has usability [21]. However, incentives assume that individuals are rational [61, 71], implementing a calculus of risk [20]. Usability assumes that the individuals will have an exogenous awareness of and desire to engage with risk reducing technology. The bounded rationality of end-users is being addressed through nudging interfaces [4]. Simultaneously, user awareness is being addressed through security education [63]. Integration of the scholarship on perceived risk is a necessary complement to these approaches.

In the physical realm, individuals can use visual, geographical, and tactile information to evaluate the authenticity and trustworthiness of a service provider or peer [55, 8]. In the virtual realm,

transactions are characterized by spatial, temporal, and social separation [34, 13]. This separation simplifies masquerade attacks in part by decreasing the cost of constructing a false business facade. While there exists a range of security protocols that are testament to their creators, Internet-based confidence scams continue to increase in profitability and sophistication. Medical identity theft, for example, is an increasing activity, and an increasing concern [45, 50]. Such fraud is made more simple, ironically, on the information infrastructure.

Human-factors research results sometimes lead to simple solutions, such as checklists for pilots; while others require significant redesign. Usability engineering and human-centered design offer half a century of science and experience that can be the foundation of solutions to challenges of human interaction with secure networks.

From a dependability point of view, 3% of failures are a result of classic failures in code, or bugs. Most failures are a result of misunderstanding the requirements or operator errors. Indeed operator failures can be conceived of as a misunderstanding of the human requirements of operators [1]. A systems level end-to-end view requires integrating requirements as a design basis and critiquing the design from a risk perspective. It is necessary to examine both potential failures of the system, and the risks inherent in success. A true end-to-end perspective includes the complete integrated system: people, sensors, community, etc. It may be necessary to be explicit in designs when dealing with group dynamics by using forward secrecy and backward secrecy: changing group authentications when existing member(s) leave to prevent the departing members from decrypting the future messages, and changing the authentication (e.g., a group key) when new members join can prevent the joining members from decrypting the previous information (even if they are intercepted and stored). The user-centric approach to our research suggests the use of distributed key agreement schemes. However, it is also necessary to add information about the relative security of potential entrants in the secure group to allow users to implement in technologically strong manners their own social judgments on the value of authenticating a device and a person.

In contrast, consider SDN in a data center configuration. There is a single organization. The connections are very long-lived and high bandwidth. The contents of a X.509 certificate are inappropriate. For example, domain name is meaningless. Appropriate fields may include organizational unit, or the supervisor of the person instantiating the switch. An appropriate configuration may have an organization creating its own elliptic curves with extremely long-lived sessions keys, further constrained so that the flows must follow a specific fiber. However, this would be completely unrealistic for a home network with your average person setting up authentication between devices. X.509 and ubiquitously trusted certificate authorities are fundamentally ill-suited to both these contexts.

Currently proposed cross-domain trust mechanisms seek to minimize computational costs and management overhead. For example, commerce systems minimize key generation by linking all attributes and rights to a single commerce-enabling certificate. These keys are validated by a single root. This creates a single point of failure for the entire system (the root) as well as a single point of failure for the consumer (the key). The only similar system in the United States is the currency system, where the failure of the US Treasury would yield complete collapse. In family systems, individual businesses, and even religions there are multiple levels and power points. In physical security, any key is a part of a key ring, so that the failure of the validity of one key does not destroy the strength of all locks.

An examination of the certificates currently used by commerce sites illustrates some rather extreme problems [30]. Twenty-two commerce sites share the same certificate. Effectively, these websites were set up with the certificate as shipped and never reconfigured. Thus any organization that uses that domain certificate to enable commerce is, in terms of the core cryptographic authenticating infrastructure of the Internet, indistinguishable from these commerce sites.

The previously noted modeling enables the identification of useful points of leverage, rather than pursuing security across all the networked homes of America (and the world). Of those points of leverage, could the users recognize them as risky? In what cases do users knowingly accept the risks for some benefit or perceived benefit, and in what cases are the risks simply invisible? Individual decisions cannot be predicted, but aggregate behaviors can be observed and understood, and that understanding can be used to improve the design of technical mechanisms. Similarly, risk behaviors may change in individual cases, but people in the aggregate will continue to see risks as framed by possible benefits and so, in the aggregate, people will behave in systematic ways [65, 32].

2.3 Results and Discussion

2.3.1 Secure Integration Of BGP And Software Defined Networks

Current networks are vulnerable not only to malicious attacks but also to severe disruption by simple misconfigurations. The attacks experienced so far have been local; there has been no global attack. Yet if a capable nation state or substate group were to try to take down the Internet by disrupting the interconnection system, how much impact might they have? Concerns like this motivate the drive to introduce BGPSEC, but neither has it been widely adopted nor has it been proven in smaller-scale operation. Such issues could be addressed as SDN diffuses, creating a combined network that is more resilient. Yet, currently, not even the issues of mutual authentication within SDN have been addressed.

Even more so, the understanding of the possible interactions of SDN and BGP networks is inchoate. The demonstration implemented under this contract illustrates one possible mode of interaction at the single-controller scale. Yet this is a proof of concept, showing progress towards a more resilient integrated network, not a declaration that this is the sole possible mode of operation. DNSSEC, BGP, BGPSEC will all interact with the control plane and are part of the answers to the questions above. But what is needed, is still undefined, and requires a multi-year commitment to the technology and resources for a trust ecosystem. Creating a resilient infrastructure requires reasoning about risks at a network scale.

As software defined networks diffuse, they will be integrated with BGP-centric networks at every level. The separation of the physical and logical components of the network may arguably be more valuable in cyber-physical systems and networks managed by less skilled operators than in the larger infrastructure.

A good example of cyber-physical systems are SCADA or, in general, industrial control systems. There, millions of devices that were designed with the assumption of physical security and isolated networks are now connected over the Internet. This trend that will not only accelerate but also include OpenFlow controllers. Legacy control system components on the electrical grid, natural gas pipelines, pumping stations, and networked controllers in the extraction industries are increasingly connected. These devices are fundamentally different in terms of vulnerabilities than the current PC or mobile device model. There is no simple method to update the software when vulnerabilities are found. Update cycles can be years or even decades rather than weeks. The level of expertise required to move from known vulnerability to subverted system varies widely; currently no simple automation of code suitable for ego-driven amateurs is available (i.e., no script kiddies) but this may be only temporary. The view of the small community of SDN security experts is that the only practical medium-term solution is reparameterization, and this can be done much more cheaply and effectively using SDN. However, SDN not only has the potential to provide better isolation for legacy cyber-physical systems, but also (if wrongly deployed) to provide ease of malicious mapping and simplified, more scalable attacks. Effective mutual authentication of cyber-physical systems will require effective architectures, trustworthy code,

and correct administration of SDN. All the challenges noted below are compounded in cyber-physical systems; yet the core authentication requirements remain.

Similarly SDN can be used to provide home network control and isolation. Our current traffic-engineering and network management technologies are not up to the task of preventing malicious distributed hosts from self-organizing into massive botnets. Should SDN be unable to improve on current methods to mitigate the chronic insecurity of networks owned by naive users, it may not offer much overall improvement to the network as a whole.

2.3.2 Management Plane

An SDN controller has two sets of interfaces, commonly considered north and south. The southbound interface views the switches under the controller's domain. The northbound interface provides a view of the state of the exterior network; interacting with multiple protocols and responding to changes.

The risks of the southbound interface are primarily denial of service and disclosure of information. The channel between controllers and switches is typically low-bandwidth, as only control information is needed, and denial of service is correspondingly simplified.

The risks of the northbound interface are from interaction with the network and interaction with what can be termed the management plane itself. The management plane will provide information about the state of the network to which controllers will respond. This is key to the dependability and performance of SDN but is not yet standardized.

The management plane is likely to be instantiated as a series of applications engaging with different protocols (e.g., eBGP) in order to take full advantage of SDN traffic engineering capabilities. Indirect manipulation through incomplete or malicious monitoring data could be bidirectional, with insecure controllers being leveraged for surveillance, stolen network resources, or for amplifying attacks from the SDN network into the BGP network. The management plane consists of software which may come from diverse sources. Monitoring software will be marketed, installed, and updated in the way that applications are today. Monitoring information, manipulation of the networks, and direct alteration of controllers are risks from malicious code. While an SDN app store is eminently foreseeable, the security questions have yet to be asked: what would the permissions look like for an SDN app? There have been no proposals.

Of the traditional security goals of confidentiality, integrity and availability, traffic engineering in general (and SDN in particular) mainly delivers availability. Confidentiality and integrity are not entirely excluded: they can always be provided by encryption at higher layers, but they can be greatly helped by dependable separation, and separation may be the most economical way to protect legacy networks such as industrial control systems where retrofitting cryptography may be impractical.

Most denial of service issues demand traffic engineering as a component to any solution. The exclusion of unwanted traffic is made possible by OpenFlow. Consider that new and unknown traffic can be given a priority, as with FRESCO, or filtered out, as with the demonstration project developed under this proposal. Filtering out traffic not intended for legitimate routes may function effectively in normal operation. But determining whether such rules would be effective under different failure conditions is not at all trivial. The Internet's resilience in the face of natural disasters and acts of war has been remarkable. Strict filtering in network exchange points might limit that resilience in emergencies. There are complex open questions around how to manage filtering at scale.

This leads to the next category of traffic management: allocating resources. The two options of

prioritizing versus filtering are extreme cases. In reality, operations will require a more subtle approach to resource allocation, in particular, resource allocation from remote requests. Currently SDN installations are sufficiently organizationally centralized or built in conjunction with organizational trust where reputation is embedded such that remote and unknown entities are not a problem. Similarly, the past two decades of Internet growth have depended on human trust, as key technical staff at large ASes know and work with each other to resolve problems. This system is starting to fray, and it's not clear how it can scale to networks that are partially trusted.

The minimal research challenges to meet in order to create SDN that is more resilient and secure at the management plane are

• the security of management applications,

• the reliable interoperability of applications from a potentially wide range of sources, and

• documentation and understanding of failure modes of different management approaches.

In addition, all of these examinations must be grounded in an understanding of the authentication requirements at the different layers.

2.3.3 Control Plane Challenges

The control plane is a defined component in standard SDN designs and instantiations.

There have been several papers on creating secure networks using OpenFlow. FRESCO [64] operates on top of the NOX controller and provides a programming framework to execute and link together security-related applications. Jafarian et. al [40] developed a system using OpenFlow that makes the IP addresses of internal hosts appear to change frequently to external networks to make network reconnaissance and attacks difficult. NICE [16] uses OpenFlow to build a DDoS detection and mitigation system for large infrastructure-as-a-service cloud providers. However, these works are focused on providing security to networks controlled with OpenFlow and rely on the assumption that the underlying OpenFlow network is secure.

FortNOX [60] was developed as an extension to the NOX OpenFlow controller to deal with conflicting and possibly malicious OpenFlow applications by adding role-based authorization and constraints to the permitted rules that an OpenFlow application can send to switches. In this case, an "application" could be anything that wants to modify, record, duplicate, or block network traffic using OpenFlow (e.g. firewalls, intrusion prevention systems, traffic logging, etc.). FortNOX addresses some part of the authentication requirements between the control and management planes by constraining access from the management plane into the control plane.

In a similar context, FlowVisor [3] acts as a mediator between controllers and switches to apply limitations to the rules created by controllers. It does this by rewriting the rules generated by the controllers to restrict their effect to a specific "slice" of the network. These "slices" can be defined by physical ports, or by packet headers (e.g. only traffic on TCP port 80). Thus FlowVisor could function to mitigate the impact from a subverted controller on the data plane, just as FortNOX offers to mitigate the impact on a subverted controller from the management plane.

Another difference between FortNOX and FlowVisor is that FlowVisor runs separately from the controller (normally on a different host), where FortNOX is a single controller that executes many concurrent applications. Both of these applications focus on restricting untrusted controllers or applications running on controllers; however, there has not been any work published examining the vulnerabilities that emerge from OpenFlow network designs or vulnerabilities with the protocol.

2.3.4 Data Plane Challenges

The data plane functionality implements flows as instructed by controllers, and in turn provides information to controllers and thus input to data management. Data plane operations are implemented by switches, which are slaves to controllers, and query controllers when they encounter a packet which does not correspond to a flow.

The corresponding vulnerabilities include denial of service by overloading the low-bandwidth channel to the controller, by sending a steady flow of queries. A more subtle attack is creating unlikely or impossible situations and sending this false data to the controllers in order to attack processing ability. If the controller is configured to create a new rule and share it across switches, this attack could cause an overload of the flow rule space for an entire set of controllers.

And of course switches themselves are vulnerable. The lack of mutual authentication means that a malicious entity can implant additional flow rules, for example duplicating flows to obtain content. The lack of switch-to-switch coordinating can create loops, which could be forced by timing attacks. For example, switch one could believe the path to network one is through switch two. By altering switch two to believe the path to network one is to switch one, a loop could be created by an attacker, thus denying service. This reifies the need for controller/switch authentication noted previously. Query management, query priority, rule management, and state validation are solutions to these potential vulnerabilities, but none of these are solved problems.

As these attacks and requirements illustrate, the data plane, the control plane, and the management plane are not entirely discrete. Information flows across these barriers. Currently, in theory there are three security architectures for OpenFlow in the marketplace, but only the first exists in practice.

The first relies on physical security. Installations with a single controller or very few controllers use physical walls and guards to ensure secure connections between controllers and switches; and thus data plane security. Remote access is disabled, and authentication consists of being adjacent to the machine. This is currently deployed in many installations, and is likely to be the norm for personal area networks if SDN is adopted in home networking [51].

The second approach assumes that passwords are adequate. Passwords are generally still sent in the clear, as TLS isn't provided on any switch other than the NEC IP8800. This approach is particularly problematic if listening mode is enabled. Using passwords, data controllers can be bound to specific controllers using an approach that harkens back to rhosts.

The third possible approach, and a natural step beyond this second approach is TLS within a single trust domain. At the time of this report, this was currently under consideration but not implemented. (As of this report, Google has not yet documented the changes made to its SDN infrastructure in winter of 2013, so there may be an extant implementation.) With this, authentication between switches and controllers can be as complex as desired by the operator. Current approaches to network modeling cannot address the layers of SDN particularly the potential failures in authentication or isolation.

2.3.5 Sample Results of SDN/BGP Modeling

In this section we demonstrate a modification of the bipartite network robustness model to a tripartite SDN network topology using a single-node, random attack (rather than the cascading failures above). The topology here is a uniform random topology with data nodes having on average connections to 5 switches, and switches have on average 3 connections to control elements. In the actual use cases, the topology will be different and may be time variant, but this sample will give a basic overview of the technique.

This example only uses random node failure, rather than any sort of targeting. This can be the same as a given device failing, or an attack on a random piece of infrastructure due to lack of information. In this case, Node number 31–a control element–fails, and is removed (Figure 3a). This leads to the destruction of links in the Switch-Layer isolating Node 28. The isolation of Node 28 impacts the ability for Node 16 to connect to the rest of the network (Fig. 3b).

After this first cascade finishes, there are 4 of 5 control elements still functioning, 9 of 10 switches still functioning, and 19 of 20 data elements still have the possibility of connecting via some path through the network. However, these are simple metrics of robustness. More complete statistics have been developed and can be used to evaluate the different projections at the end of the cascade [5, 46, 66]. In practice, we would want to run many samples of random failures of a range of given sizes and compare the number of initial failures with our measure of resiliency.

The modeling, investigating, and thus understanding of the emerging fragmented network structure can be improved by leveraging models from disciplines beyond information networking. To prove this a simple model has been constructed showing the failure of connectivity under the given topology.

(a)Tripartite Representation of SDN

The model and thus sample results represent failure at a very abstract level in a very simple network. However, even at this simple level the model illustrated the trade-offs in a single versus multiple SDN controllers.

(b) Projection of Switch Layer of SDN, through Controllers

(c)   Projection   Data   of   Layer through Switches

Figure 2: A Random, Restricted Tripartite Topology

Figure 2 shows an initial view of a random, restricted tripartite topology with data elements represented in green, switch elements in blue, and control elements in red Fig. 2a. Fig. 2b (lower left) is the projection of the network of switches mediated through the control elements. Fig. 2c (lower right) is the projection of the network of data elements mediated through the switch elements.

Of course, much is to be done to generate more realistic models. For example, it would be useful to add distance between nodes, as well as connection capacity. In particular, expanded models can examine the propagation of routing information within a pure SDN network and within a combined BGP/SDN network.

Thus, we have illustrated that there are extant modeling techniques from other disciplines appropriate for SDN networks, or combined BGP/SDN networks.

(a) Control Node 31 fails an is removed, isolating Switch Node Node16.

(b) Switch Node 28 fails and is removed, isolating Data Node16

(c) Data node 16 fails and is removed
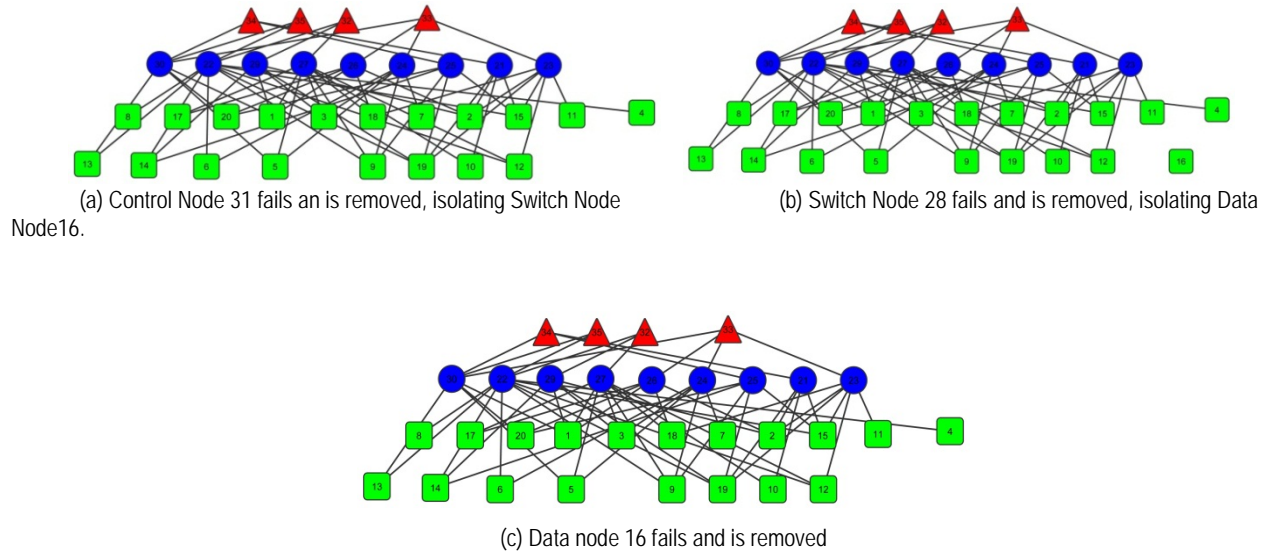
Figure 3 Cascading Failure Due to Loss of Random Control Node

(a) Switches

(b) Data

Figure 4 Final projections of the switch network (Fig. 4a) and data network (Fig. 4b)

Figures 3 and 4 shows modeling appropriate for the types of attacks on the trustworthy, untrustworthy, or partially trustworthy components in pure SDN or SDN/BGP networks. The simple model above is included as a proof of concept, to illustrate the potential of cross-disciplinary modeling for identifying potential challenges is reachability or cascading in SDN.

2.3.6 Proof of Concept: Quagga and Bongo

Having defined the scope of the challenge to leveraging the point of inflection that is SDN for a more resilient network, can progress be made? The challenges are both severe and subtle, and the solutions require a wide range of expertise. The purpose of Bongo, combined with the rewrite of Quagga, is to illustrate that SDN provides an opportunity to strengthen the entire BGP infrastructure by transparently optimizing and hardening SDN islands. Essentially, Bongo knows the status and topology of a subset of the network (everything southbound) and leverages this for improved security and performance across SDN. Quagga is redesigned to give Bongo the status of the network northbound to inform the decisions flowing downward.

The long-term design of Bongo is to do the following:

1. at all times, costlessly implement ingress and egress filtering of data flows;

2. accept route updates at the control plane;

3. process and update these route updates, given a known SDN state; and

4. update the controller state so the controller installs updated rules in switches.

As a result there's limited ability for attackers to leverage devices southbound of Bongo to deny traffic that violates BGP policies at ingress. A more profound result is that changes in the larger network may be delayed, or routed around by Bongo. This ability to decrease the rate of change in the network based on arbitrary policies is profound. For example, updates that demand traffic flow through an untrusted AS can be rejected. (However, the potential implications of this for the larger BGP network are potentially immense in terms of router storms, and not well understood, so such decisions should not be taken lightly.)

What are perceived now as the necessary tasks for secure BGP (i.e., S-BGP, route authentication) are handled by Quagga. However, higher level trust architecture can be added by Bongo using a trust API, on a logically distinct COTS processor. The demonstrated prototype system shows a simple but previously infeasible analysis of the RIB. The system will detect paths distributed with multiple hops within one AS; which is sometimes done for economic reasons. The current instantiation of the reputation system will also identify loops, thus offering the possibility of identifying content hijacking. Even if the only result is identification of current path hijacking, because SDN does not give one AS any power to determine the routing of another, this is a fundamental change in today's network where path hijacking is identified by expert examination often after significant time has elapsed. Thus, hijacking incidents occur frequently (presumably usually in error) but may persist for months.

The demonstration illustrates in one case economic alignment of security and operations in the merged networks (BGP, SDN). It offers the promise of a more resilient BGP though merging this with islands of SDN.

Quagga has been rewritten for the interface between Bongo and BGP. Quagga accepts routing updates to create a distinct routing information base. Quagga translates BGP updates to create a

customized Routing information Base (RIB) for each routing context; for example, for a large ISP with different locations or a content distribution network. Quagga integrates discrete routing decisions to form groups; it is intended for clusters of devices that are at the forwarding layer. At its most basic, the project instantiation of Quagga provides an additional layer of abstraction between BGP and actual packet transit.

Quagga can make different routing contexts for each of these discrete contexts by extracting the handling of the next hops. Thus each separate route context can have its own distinct metrics without requiring either per-router metrics or metrics across an entire AS. In practice, there will be created a single BGP instance which accepts all the BGP feeds from all the peers across the organization, (distributed data center, content distribution, etc.) and then similarly sends out comprehensive updates to all peers. Yet on the south side of Quagga, different SDN islands can receive the different information from Quagga, but the network as a whole views all the islands as a single contiguous whole.

There are serious empirical questions to be asked about routing. One goal of working with Quagga was to provide a platform to answer those questions.

Current Quagga instantiation provided by Hall and Anderson upon request, was planned to be transfered to New Zealand January 2014. Current demonstration code illustrating filtering and the creation of the FLIP for Bongo is available upon request.

2.3.6.1 Architecture

At a minimum, Bongo communicates with two external entities: an SDN controller to affect the flow of traffic on a network, and a BGP API to react to route announcements. Depending on the desired goals, Bongo could additionally communicate with other data sources to provide additional information such as reputation data and reliability metrics.

The current Bongo instantiation is designed to interact with an instantiation of Quagga designed to simplify network management as described above.

Bongo utilizes the northbound API of an SDN controller to install flows into a network to permit or deny any given block of IP addresses. In the initial version it communicates with the POX controller; however, it should easily be adaptable to any other SDN controllers that allows basic permit and deny rules for IP blocks.

In order to retrieve information about current BGP announcements, it will leverage the API of a BGP speaker to obtain access to the entire routing information base about its neighbors. Access to the routing information base instead of the forwarding information base is necessary because one of the purposes of Bongo is to determine which announced routes have preference based on any organizationally selected basis.

As shown in the figure below, RIB is converted into a series of flow rules (which we dubbed a Flow Information Base, FLIB) that is shared across the controllers from Bongo. The controllers then redirect the data on the newly-instantiated flow rules, received from Bongo as a FLIB. The figure also shows the information returning up the stack. The layers of abstraction and decision-making are abstracted from the data flow, meaning that the rate of change of the network is marginally slower as well as subject to high degrees of local customization. This introduces timing problems. Thus the information from the data layer is needed for two functions. First, to ensure that the flow information base changes have been, in fact, instantiated and second to provide information about congestion, delays, and loads to inform future Bongo decisions.

At the most basic level, Bongo allows costless filtering. However, Bongo can also examine route

updates to address known weaknesses in the current routing plane.

2.3.6.2 Filtering

Bongo allows developers to customize the way network traffic is forwarded based on information received from BGP neighbors. One practical use case is filtering out traffic sourced from addresses not advertised by a BGP neighbor. In addition to improving the security of the overall Internet, dropping spoofed packets lowers the burden on an ISPs network. This provides an economic incentive to implement this filtering.

This is similar to the ingress filtering suggested by BCP38 (urlhttp://tools.ietf.org/html/bcp38), which states that an ISP should only allow packets from a customer if the source IP addresses match the ones allocated to that customer. If every ISP followed this practice, source-spoofing on the Internet would be a solved problem. However, enough ISPs fail to implement source verification that DNS amplification attacks have resulted in DDoS attacks up to 120 Gbps in size.
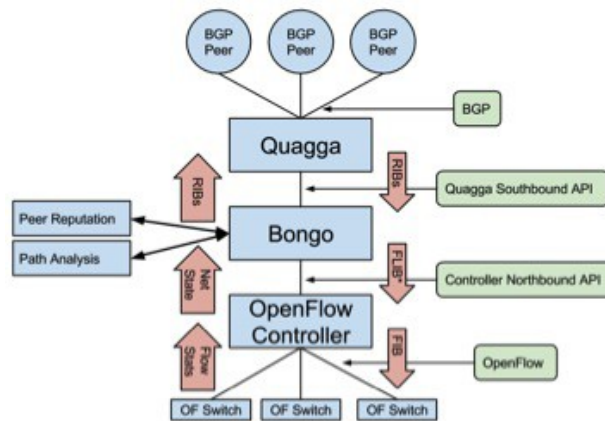
Figure 5 Information Flow North and South of Bongo and Redesigned Quagga

To accomplish this, Bongo uses the BGP announcements received from a given BGP neighbor to generate a set of flows that only allow traffic sourced from the networks in those announcements. These flows are then installed into an OpenFlow switch at the peering point with the neighbor so filtering is performed on the edge of the network at line rate.

2.3.6.3 Route Acceptance

Bongo is responsible for constructing the routing table (flow information base) from the routing information base it receives from each neighbor via the BGP speaker API (e.g. Quagga). This permits custom logic based on arbitrary metrics and/or external data sources to affect how routes from neighbors are handled.

For example, Bongo could reference a list of AS reputations and ignore announcements that result in traffic being directed towards low-reputation autonomous systems as long as the previous path remains. This would mitigate situations where an ISP attempts context exfiltration via route hijacking. Normally, this is implemented by announcing a shorter path to the target. As long as the previous route still exists, Bongo would ignore the update and build the RIB based on the longer trusted path and stop the propagation of the shorter path. An alert could then be generated so an administrator could approve the route if he/she determines it to be valid.

Consider the current threat of information acquisition via route-hijacking. Shorter paths are the quickest and most simple way of hijacking a route. Claiming to be a customer, or being a customer, makes a route more attractive by default. Distance measures and changes in AS route behavior are also sources that can be used to refuse updates, delay updates, or simply generate alters. As the number of backbone providers is limited, organizations with distrust or even enmity share the same providers over some hopes. Yet ASes that have never been peers or transit points suddenly announcing themselves as such is not an unusual component of route hijacking.

2.3.6.4 Testbed Architecture

The current test network is composed of a single VM, two HP OpenFlow switches, and a single Spirent test chassis. The VM is running Quagga, Bongo (not right now obviously), and POX. Bongo interfaces with both Quagga and POX via separate IPC APIs. Each OpenFlow switch runs two virtual OpenFlow instances, and connects a total of eight virtual routers simulated by the Spirent Test Chassis to the Quagga route server. The test chassis is also responsible for traffic generation matching the advertised routes of the simulated routers. In addition, the test chassis is capable of producing erroneous traffic for verification of implemented routes across the network data plane.

2.3.7 Impact

There were four sources of impact from this project: technology transfer, publications, participation, and code. The previous section provided information about the code which was developed under this one-year project. The other components are addressed briefly in this closing section.

2.3.7.1 Technology Transfer

The most ambitious and significant technology transfer will continue after the close of the project. This is the transfer of the modified Quagga discussed in Section 2.3.6 to Citylink (New Zealand's Internet Exchange Points). Citilink will evaluate the Bongo/Quagga demonstration code in spring of 2014.

A second form of technology transfer was in the form of the personnel supported with the year-long project. This included Kevin Benton spending one month in a cooperative visit (as part of the management plan) at the University of Cambridge. Following this one-month visit, he then spent the remaining summer months working full time at Big Switch. The work he completed at Big Switch was presented under their auspices the first week of November at the OpenStack Summit in Hong Kong. The second significant placement was Dongting Yu from Cambridge to SRI International for the summer of 2013. The third significant student placement, in terms of technology transfer, was Zheng Dong of Indiana to Microsoft Research for the summer of 2013.

2.3.8 Publications

Publications are the easiest to identify components of knowledge transfer. The following list the nearly two dozen accepted publications. It does not include publications in preparation or currently under review.

1. Andrew Odlyzko: Papers on Communication Networks and Related Topics. Will smart pricing finally take off? To appear in Smart Data Pricing, S. Sen, C. Joe-Wong, S. Ha, and M. Chiang, eds., Wiley, 2014.

2. Dongting Yu, Andrew W. Moore, Chris Hall, Ross Anderson, Authentication for Resilience: the Case of SDN. In Proceedings of the Security Protocols Workshop 2013 (LNCS vol. 8263).

3. Dongting Yu, SDN Security and Resilience. 6 Aug 2013 at ESnet/Lawrence Berkeley National Laboratory

4. Dongting Yu, SDN Security and Resilience. 7 Aug 2013 at Open Networking Lab (ON.LAB)

5. Dongting Yu, Security: a killer app for SDN?. 19 Sept 2013 at SICSA Software-Defined Networking workshop

6. L. Jean Camp, Designing for Trust. RSA Invited Speaker (Bedford, MA) 5 November 2013.

7. L. Jean Camp, Security Designers are the Weakest Link. Communications Futures Program (CFP) Speaker Series CSAIL MIT, (Cambridge, MA) 22 October 2013.

8. L. Jean Camp, Efficient Methods to Guard Against Online Risk. Executive Office of the President, National Security Staff, (Washington, DC) 27 September 2013.

9. L. Jean Camp, Security as a Common Pools Good. 2013 Blouin Creative Leadership Summit, Metropolitan Club (New York, NY) 25 September 2013.

10. L. Jean Camp, Security, Usability, and Why We Have Neither. HotSec '13: 2013 USENIX Summit on Hot Topics in Security (Washington, DC) 13 August 2013.

11. L. Jean Camp, Economics of Cybersecurity. National Grid Cyber-security Research Centre, University of Aberdeen, (Aberdeen, UK) 15 March 2013.

12. L. Jean Camp, Enabling Cybersecurity. Horizon Digital Economy, U. of Nottingham (Nottingham, UK) 13 March 2013.

13. L. Jean Camp, Risk Communication for Cybersecurity. Computer Science and Telecommunications Board, National Academy of Sciences, (Washington, DC) 12 March 2013.

14. V. Garg, and L. Jean Camp, Heuristics and Biases: Implications for Security Design, IEEE Technology & Society, Mar. 2013.

15. Z. Dong, A. Kapadia and L Jean Camp, Pinning and Binning: Building Whitelists and Blacklists Using Machine Learning. ACSAC Extended Abstracts, (New Orleans, LA) 3-7 December 2013.

16. Vaibhav Garg and L Jean Camp, Spare the Rod Spoil the Security?. TPRC, (Arlington, VA) 26-30 September 2013.

17. Kevin Benton, L Jean Camp & Chris Small, OpenFlow Vulnerability Assessment. HotSDN, August 2013, (Hong Kong) (extended abstract)

18. Shaddi Hasan,Yahel Ben-David,Colin Scott,Eric A. Brewer,Scott Shenker: Enhancing rural connectivity with software defined networks. ACM DEV 2013: 49

19. Sam Whitlock,Colin Scott,Scott Shenker:Brief announcement: techniques for programmatically troubleshooting distributed systems. PODC 2013: 134-136

20. Brandon Heller,Colin Scott,Nick McKeown,Scott Shenker,Andreas Wundsam,Hongyi Zeng,Sam Whitlock,Vimalkumar Jeyakumar,Nikhil Handigol,James McCauley,Kyriakos Zarifis,Peyman Kazemian:Leveraging SDN layering to systematically troubleshoot networks. HotSDN 2013: 37-42

21. Aurojit Panda,Colin Scott,Ali Ghodsi,Teemu Koponen,Scott Shenker:CAP for networks. HotSDN 2013:91-96

22. Seyed Kaveh Fayazbakhsh,Yin Lin,Amin Tootoonchian,Ali Ghodsi,Teemu Koponen,Bruce M. Maggs,K. C. Ng,Vyas Sekar,Scott Shenker: Less pain, most of the gain: incrementally deployable

23. Sangjin Han (U.C.Berkeley), Norbert Egi (Huawei Corp.), Aurojit Panda, Sylvia Ratnasamy (U.C.Berkeley), Guangyu Shi (Huawei Corp.), Scott Shenker (U.C.Berkeley and ICSI). Network Support for Resource Disaggregation in Next-Generation Data Centers. Hotnets 2013

2.3.9 Participation

In addition to formal outreach, there were presentations explaining our work beyond Big Switch, Microsoft, and SRI International. Specifically, the example of Bongo as an exemplar use of OpenFlow was presented at SDN World Congress, Bath Homburg (18 Oct 2013) by Uwe Dahlmann of Indiana. Kevin Benton participated in the Cyber Defence Information Exchange of the Allied Command Transformation (which is a component of the Strategic Command of NATO) 15-19 April of 2013. Andrew Moore participating in the Open Networking Summit, a highly restricted event of SDN customers, 4 February 2013. Hall was invited to attend the DIMACS Software Defined Network Workshop in Rutgers on 3-4 December 2012. He also traveled from London to Santa Clara to participate in the Open Networking Summit in April 2014. While in the San Francisco Bay area, he also visited the project participants at Berkeley, as well as handling project outreach and knowledge transfer with the Open Source Routing Forum and ON.Lab.

Project leads Anderson, Camp, Small, and Odlyzko combined the in-person project team meeting with the Workshop on the Economics of Information Security in May of 2013. Professors Camp and Anderson also participated in the invitation-only event Security and Human Behavior in California 2-3 June 2013. Dongting Yu was invited to the SICSA workshop in Edinburgh on 19 October 2013.

Other participation and outreach is embedded within the Critical Use Cases (which follow as appendices). In particular, the Battlefield Use Case included participants from Crane Surface Warfare Center. After the end of funding, Crane will continue to fund a doctoral student at Indiana University who will work on white box switches. Cambridge University used extensive social contacts for the development of the ISP and Cyber-Physical Systems Use Cases, too numerous to list here.

2.4   Conclusions

The international and multidisciplinary team has completed a one-year project and defined the requirements to leverage the emergence of SDN as a force in high-level networking. Jean Camp was the team lead, from Indiana University. The United Kingdom component of the team consisted of Ross Anderson and Andrew Moore from the University of Cambridge, working with Chris Hall. Andrew Odlyzko, and Zhi-Li Zhang from the University of Minnesota provided modeling expertise. Scott Shenker from University of California Berkeley closed out the roster.

Understanding the challenges in an integrated SDN/BGP network requires high level threat modeling; not only formal network modeling to determine efficacy of historical attacks but also integration of security requirements, including economic and game-theoretic techniques that deal with dynamic systems, incomplete information and other forms of uncertainty, and interactions between rational self-optimizing agents.

To address the myriad risk, we began our analysis with a set of critical use cases (included here as appendices). From these cases, we developed authentication and operational requirements for each layer, including the core contribution of a clear abstraction of the management plane. Based on these we simultaneously modeled the network using these assumptions and created a prototype that illustrates progress against the challenges identified.

2.4.1 Management Plane

The management challenge has two future trends: traffic engineering application and interoperability at scale. Management planes are currently isolated, with little research in scale or diffusion of failures and attacks. Applications are one-off research proposals.

Recall that the results are that the resilience and security research challenges include the security (in terms of both errors and malicious failure); interoperability of applications and authentication modes; and an understanding of diffusion of failures.

The core challenges that must be addressed at the management plane are the security of application code, application interoperability, and implications for reliability. The domain is new. While the modeling section addresses some of the differences, there are a range of tools that may be applied to the SDN management plane challenges. Approaches in other domains include the Apple centralized verification approach, the reputation systems and sandboxing approach taken by Android, static code auditing, source authentication, as well as test beds, and test suites run by third parties.

In addition, all of these examinations must be grounded in an understanding of the authentication requirements at the different layers. Certificate architectures as currently constructed are ill-suited for a controller/switch architecture, or even for verification of management applications. No current certificate architecture addresses the requirements of SDN.

2.4.2 Control Plane

Isolation of the control plane will not of itself resolve the vulnerabilities and questions with respect to transparency and mutual authentication in networks. In fact, there is serious risk that SDN will exacerbate the current difficulties if adopted in its lowest-cost, insecure configuration.

Securing and protecting the SDN control plane is vital to the reliability and resilience of SDN. This includes both securing and protecting the SDN controllers, but also securing and protecting the "dissemination channels" through which control decisions are delivered to individual forwarding elements.

The research challenges to meet in order to create SDN that is more resilient and secure include

• SDN authentication infrastructure for remote components;

• isolation from BGP network failures; and in normal operation,

• secure integration and appropriate isolation of legacy sub-networks, BGP networks, and SDN networks.

The interaction of OpenFlow with extant networks has not been examined. OpenFlow instantiations have been assumed to be isolated from BGP networks. Controllers must either interact directly with BGP, or must have interactions with BGP negotiated through the management plane. Both paths (communications north and south of the controller) are illustrated in Figure 1.

2.4.3 Data Plane Challenges

For SDN's potential to be realized, the problem of how to authenticate mutually suspicious principals is urgent. Without this we cannot join up SDN islands in different ASes, or for that matter, different compartments of the same corporate network cannot be reliably and securely joined.

Multi-controller architectures are not yet documented, and make the challenges noted above more complex. As networks of controllers emerge and the number of switches controlled moves into the thousands, naive TLS authentication mechanisms will not scale. Multicontroller and interexchange

architectures will have to deal with untrustworthy control information not only from BGP but also from potentially hostile or incompetent external controllers.

2.4.4 Future Multicontroller Interactions

Security engineering decisions cannot be made with the assumptions that attacks will continue as usual. Consider a DoS attack. Once ingress and egress filtering are made affordable and easy to implement via SDN, the attacks will evolve to leverage SDN just as such attacks now leverage DNS resolvers. For example, consider if traffic is to be sent from Chicago to Indiana University. Should this traffic go through Purdue or IUPUI on the academic networks? Is the link to IUPUI congested? Or is there a rogue controller trying to swamp Purdue by announcing that all other routes are congested and that only Purdue has excess capacity? If an attacker can use misinformation to misdirect traffic engineering decisions, and at the same time implement a denial-of-service attack against the controller / switch channel, he could potentially get much more leverage than with DNS.

SDN networks will initially be newer, and less mature code is likely to have more bugs [57] as compared to more mature code. The response to the telephony incident included improved testing, but testing of each system component is not adequate to understand the resilience of the larger system. The modeling section below begins to tackle issues of the entire system and its resilience, and better transparency will improve network awareness. However, there is a clear need for standards and a testing suite, if not testing facilities, for SDN apps. The appropriate model for this is an open question.

In terms of certification provision, the current model of certificate authorities is known to be problematic. The certificate issuance and validation infrastructure used for the web will inform the trust challenges with SDN, but only with great serendipity would they correspond to the needs of a cross-domain SDN world. The likelihood that this promiscuous model that encourages industrial-scale generation of certificates for unverified requesters will match SDN is small. The SDN world will have fewer, bigger islands, with clusters of controllers taking more important decisions.

2.4.5 Cascading Failure and Robustness in SDNs

In the description of methods above, we identify a range of possible approaches to network modeling from other disciplines. In this section, we conclude the bioinformatics offers a particularly suitable set of tools, and provide some sample results building on metabolic models. Viewing the structure of the control elements along with forwarding elements as similar to a metabolic network results in each forwarding element's interaction with a control element generating "production" of the next forwarding element. This analogy allows us to draw on the deep literature within Systems Biology analyzing the effects of "knocking out" given elements in metabolic networks.

A promising technique, and the one reported here, is proposed by Smart et al. where they view the network as a percolation process [2]. They define a viable metabolite as one that participates in at least one generating and one consuming reaction. Generating reaction create a resource that is consumed in a reaction to produce a second resource. For example, two generating reactions producing proteins enables a consuming reaction to then produce a single protein complex. In Biology this is known as the Topological Flux Balance (TFB) criterion. Expressed in network terms, a node i is viable if and only if $k_{i,in}$ and $k_{i,out} \geq 1$, where $k_{i,in}$ and $k_{i,out}$ are the in and out degrees of node i respectively. External nodes, those that provide either only input (nutrients) or only output (end-products), are exempt from the TFB criterion, but must maintain at least one of $k_{i,in}$ or $k_{i,out} \geq 1$ to be considered viable [2]. Obviously, these correspond to network sources, destinations, and flows.

Under this model, a cascading failure is measured by applying a recursive algorithm where a random node is removed, then the network is fed back into the TFB criterion, and all non-viable nodes are removed. Then the network is passed back to the algorithm until there is no change in the network [2]. This metric gives a good measure of the effects of the removal of a single node, but it does not give a good comparison to other networks, nor the significance of the given cascade. For this we turn to the method developed by Gu¨ ell et al [35].

Gu¨ ell et al. expand on the methods developed by Smart et al. to examine the effects of knockout in comparison with a null model of a random network with the same number of edges and molecular mass conservation [35]. Molecular mass conservation in our case relates to the ability of the resulting forwarding elements to handle the load balance of another (failed) forwarding element. In addition to the study of individual knockout, their technique includes the study of failures in pairs of reactions. The modeling of failure in reaction pairs can be used to model the interdependence of virtual network components in the physical layer.

2.4.6   Security Engineering & Economics

The demonstration project by Indiana University and Cambridge University illustrates that SDN can change the economics of data filtering: an SDN-based exchange point can bring direct cost savings to its operator. This simple change in the economics of filtering will also change the economics of denial-of-service attacks. Currently, it is simple to harness large botnets in such an attack, as shown clearly by the attack on Spamhaus. Making traffic management and failures visible can vastly improve global network management. No doubt attackers will respond by changing the nature of DoS attacks, but in the face of better coordinated defense, attacks should require much larger contributions by individual bots and the networks which host them. But this will also make bot detection easier.

One goal of the modeling in the previous section was to model potentially complex interactions to identify potential deployment incentives. Understanding the source of local pain, and resolving that, creates an incentive to adopt a security technology in pursuit of local resilience, reachability, and prevented cascading failures.

One goal of the following section is to explore the necessary human components in incentive alignment, for example, usability engineering makes systems easier to use and thus decreases the cost of adoption.

2.4.7   Human Factors

What systematic behaviors are needed at the end points to create a more resilient network? What authentication structures are needed to support these behaviors?

Previous work has found that interface design [42], group affiliation [27], and communication [41] affect the extension of trust and willingness to accept risk when interacting with the network. This work in the social sciences proceeded and to some degree predicted the work on social navigation. While early studies focused on the effect of computer mediation on the extension of trust, they do not address the core issue of the trustworthiness of the underlying computer technology with which individuals interact.

The tightknit group of engineers who came of age with the Internet are unlikely to be present at or after the first SDN storms. The emergence of an apps model further means the number of potential programmers goes from network engineers to everyone who believes that they can write an app; with the corresponding loss of trust and diversity of expertise.

Code maturity parallels the issues of operator experience. Code maturity and interoperability are

both economic and human factors concerns. A smaller base of trained operators is expected with SDN, as the complexity is moved to the code. However, the loss of tacit knowledge combined with potential interoperability failures has not been considered in system design.

SDN will extend beyond the trained operators into the home, as it diffuses across the network. What are the reasonable expectations for home users? What technical design decisions can affect how network operators and individual users are connected in terms of expectations, organization, and risk? These are profoundly different questions but the same methods apply. Network engineers and operators will be the first category. To begin to design for such a group, with specific expertise and culture, it is critical to first listen to them. Clearly emergent threats and next-generation vulnerabilities will not be foremost on their minds. However, basic understanding of how the people interoperate with the system is critical to understanding how to provide them information and support.

Examples of human challenges in the current control plane include misconfiguration due to inexperience as well as allegedly 'fat' fingers. (Perhaps this euphemism has led to acceptance of what could be a solvable human factors problem, where the methods of usability, safety engineering, and human factors could be applied.) In addition, there are more subtle challenges. One of these is tacit knowledge in network operations. There is a generation of networking engineers who came of age professionally with the networks, with their experience and knowledge base expanding concurrently with the network itself. The small base of trained operators risks being overwhelmed by the massive base of programmers who believe themselves capable of programming apps for any possible problem.

2.4.8 Demonstrations

In this work we have shown the promise of software defined networking (SDN) in terms of the creation of herd immunity against types of attacks that have proven intractable in current BGP networks. We used biological models to show the capacity to characterize, and thus engineer herd immunity on the network against certain classes of attacks. A core component of the larger argument, as well as the examples, is the selection of participants so that there is both individual incentive to participate and the potential for emergent herd immunity.

This work documented both a demonstration and a proof of concept of using the biological approach to designing herd immunity. Both designs integrate incentive-aware engineering, such that initial investments provide immediate benefits for the adopter. The herd immunity is illustrated with the biological models, and shows the potential levering of networks effects with less than ubiquitous adoption. The first is an illustration of how one AS can protect itself against bad routes by identifying these and delaying them, encrypting payloads, or choosing alternative services in response. Theoretically there is a tipping point at which level the adoption of route filtering is sufficient to protect the entire network and enable effective (in practice) route rejections utilizing the additional layer of isolated control in SDN. The second is an example of ingress filtering to use in Tier 2 networks against Tier 3 networks that enables them to drop spoofed traffic. Again, there is a level of adoption that creates population or herd immunity for the entire network. Both of these approaches depend on historical analysis of routes and peers.

Bongo and the new instantiation of Quagga are grounded in incentive-aligned design. While each of the technologies has secondary benefits to the Internet as a whole, both have primary design goals of providing value to the entity that must adopt them. Bongo's approach allows ISPs to implement source verification on inter-ISP links. So when an ISP fails to prevent spoofing, spoofed packets will still be filtered at the BGP peering points with other ISPs. This brings the second, point Bongo is also designed with a goal to be effective across the network with widespread but not complete adoption by Tier 2 providers. Any technology that requires total adoption implicitly requires that the malicious cooperate

in their own defeat.   Similar to the vaccination's herd immunity, only a subset of ISPs have to implement Bongo to mitigate spoofing. Spoofing and denial of service attacks would still exist within the domain of a single ISP

With Bongo, organizations can choose to reject routes or even send traffic through one provider as opposed to another based on their different route acceptances. For example, if one organization is accepting routes through a hostile jurisdiction, then some organizations may switch traffic to another provider. Senders (e.g., corporate or public sector networks) can choose to delay information (such as remote back-ups). Alternatively, originators of traffic can alter their own patterns based on the received RIB, deliberately avoiding alternative routes. This decreases the rate of change which results in a potential decrease in responsiveness when there are genuine disasters or failures. But this also has the potential to make route-hijacking more detectable, and less likely to cause widespread difficulties. In addition, it makes route-hijacking attacks (usually for content exfiltration) more difficult to hide.

2.5   Recommendations

In the rush to get SDN products and companies to market, critical aspects of security are being neglected, from authentication through denial-of-service to the design of an access control architecture for SDN apps.

We recommend a significant interdisciplinary investment in securing SDN. Having demonstrated designs that provide individual incentives and offer the promise of herd immunity, we illustrated the potential for utilizing this point of leverage for a more diverse and a more secure control plane. In addition to networking expertise, the following areas are needed for optimal designs:

- authentication and security,

- economics and security,

- human factors, and

- modeling expertise from biological and physical sciences.

2.5.1   Authentication and Resilience in the Devices

Recall that a SDN controller has two interfaces: southbound views the switches it controls while northbound interacts with the management plane. Both can fail. A typical risk to the southbound interface is that overloads can cause cache misses and performance degradation. For example, in an agile application such as a load balancer, a DDoS attack that attempts to open a large number of new flows, or a naturally occurring phenomenon such as a ping sweep can downgrade performance.

The risks of the northbound interface are much more diverse. If the management plane is instantiated as a series of SDN applications then controllers exposed not just to the benefits of software in terms of flexibility and innovation, but also to all the issues of software security, coupled with the problems of networks too. That the more-regulated and more mature telephony network was brought down by cascading software failures in February 1998 indicates that network maturity does not harden against software failures. (A software failure disrupted telephony on the East Coast in February 1998 http://www.wtonline.com/vol13_no18/special_report/271.html.)

To put this in the most plain language the questions that need timely answers include the following:

- What decisions are made under SDN? (network action)

- What information do you need to make those decisions? (data authentication)

- From what parties can you get that information? (role attestations)

- What do you need to know about those parties? (attribute attestations)

- Where will the decisions be taken? (access control architecture) These raise issues of certificate provision, reputation, and compliance.

For compliance, there are no open forums for SDN compliance standards or evaluations. Standards organizations for capital-intensive industries without transparency or openness have no history of investing in resilience or security absent external forcing functions. This vacuum has resulted in the current inadequate state of authentication and access control, with no empirical or historical reason to believe that continuing on the current path will fill that space.

Flow service applications will face a virtual network, making decisions based on data from ever-increasing range of sources. The questions of trusting the data were identified in the DoS issue above. What of trusting the code itself? A single authority (which is the Apple store model) is unlikely to be feasible. The problems of code authentication in a multi-provider model (i.e., the Android model) are demonstrably unresolved. There is clearly a need for testing, testbed(s), and compliance verification beyond vendors and tight knit groups of engineers. Software engineering offers predictive models for vulnerabilities based on age of code and other variables; analysis of attack surfaces is another example of a domain with offerings yet underutilized.

2.5.2   Modeling and Implications on the Network

Individual security may be adequate for stand-alone systems but security in the control plane has unique requirements beyond incentive-align designed. Second order effects can begin to be understood by bringing in modeling techniques from other disciplines. Using modeling techniques that are standard in bio-informatics we have illustrated the potential for incentive aligned design to secure not only endpoints but also to mitigate attacks that have proven intractable with current individual approaches.

The questions that need to be asked revolve around the resiliency and robustness of software defined networks in different conditions. Additionally, there are issues of control plane injection attacks and cross-application interactions. In BGP networks, due to the unification of data and control elements, network analysis can be done using well defined techniques. SDNs, by contrast, must be represented in a different manner due the separation of data and control elements.

Modeling also offers insight into the long term economic implications of design. Economic considerations of costs, scale, and diffusion are critical to create solutions with real world impact.

2.5.3   Economics Factors

Incentives are particularly critical for the design of network systems because of the strong network externalities. BGPSEC is an example of a protocol that was operationally and economically misaligned to the problem it was intended to solve. BGPSEC makes both requirements and benefits that discourage incentive alignment and its externalities make it hard for early adopters to reap much benefit. Local, incremental, and early benefits are the key to rapid adoption.

One way to implement incentive alignment is to move specific components of security away from

the weakest-link or public-good model by creating designs that change the underlying economics. Creating systems where exclusion is possible alters the nature of security from public good to club goods; creating systems which are rivalrous creates security that can be adopted as a club good. The fundamentals of economics offer the potential to build products where benefits accrue quickly to adopters. In such cases there can be external benefits (i.e., potential herd immunity) or only local benefit.

Club goods theories provide a strong theoretical foundation for determining the importance and risks of scaling. What is the requirement for the scale of a social network for a particular application to be reliable and functional? When the application is scaled up, what is the effect on the risk? Is the risk increased in terms of disclosure while it is decreased in terms of denial of service? There are inflection points which minimize all risks in the design and use of each application, and these can be informed with theories from club goods as human incentives are included in design.

To summarize, a core engineering economics requirement in security technologies for SDN is incentive-aligned design. Such designs must address coordination problems, meaning that the incentives cannot apply only after there is widespread adoption. First and second order effects are components of design for diffusion; this requires partial deployment incentives.

2.5.4 Human Factors

Studies of usability, risk communication, economics, and human factors are critical for a resilient next-generation network. Research done by computer scientists and network engineers alone will not enable SDN to reach its full potential or have the optimal effect in terms of a more resilient network. Testing must include engineers and operators with different levels of expertise, and explore the different cultural factors in ASes large and small, from the private and public sectors.

Even the most exact expert information on computer security is unlikely to result in human beings applying a calculus of risk. Not surprisingly, experts are humans too. Interaction with experts can focus on identifying patterns, and enabling greater depth of investigation. In contrast, for the home network, non-experts cannot be expected to identify patterns or express nuanced trust policies. Humans have consistent practices when they estimate individual risks.

In no research domain are human factors for SDN implementations being studied. Despite a hard-won understanding of the importance of human factors from access control to XML, and the fact that most real service outages come from route leaks caused by BGP misconfiguration, and that these in turn are caused by the existing vendors' command-line interfaces which date from the 1980s and differ from each other in confusing ways, human factors is not an active area of research in SDN. (Nor is human factors expertise visible in the SDN teams from the major vendors.)

Rolling out SDN without concern for usability would be irresponsible – especially as usability issues with legacy routers underlie most route leaks to date. SDN holds out the promise of providing network information through a layer of logical abstraction and thereby helping to solve many hard challenges in router configuration; but that can only happen if designers get usability right, as that's where the real failures are happening today.

2.5.5 Closing

The current approach to SDN security is inadequate. During the year of the project, both its importance and timeliness were reified by market events, illustrating the expansion of SDN envisioned in the proposal was prescient. The expansion of SDN as described in the proposal was also verified when the weaknesses of SDN were identified by data loss by Google, and by the continued diffusion of SDN

products with limited or no security. However, the team has had some impact on the market even in the single year, as illustrated in the technology transfer section.

The economics of SDN are a powerful driver, and 2013 was the opening of a window of opportunity to effect the implications of SDN adoption and diffusion into the wider network. There will be some point at which SDN is so widely deployed that the window will close.

3   REFERENCES

[1] Summary of a Workshop on Software Certification and Dependability. 2004.

[2] Cascading failure and robustness in metabolic networks. Proceedings of the National Academy of Sciences of the United States of America, 105(36):13223–8, Sept. 2008.

[3] Opennetworkinglab/flowvisor github. https://github.com/OPENNETWORKINGLAB/ flowvisor, Mar 2013.

[4] A. Acquisti. Nudging privacy: The behavioral economics of personal information. Security & Privacy, IEEE, 7(6):82–85, 2009.

[5] R. Albert, H. Jeong, and A. Barabasi. Error and attack tolerance of complex networks. Nature, 406(6794):378–82, July 2000.

[6] R. Anderson. Why information security is hard-an economic perspective. In ACSAC '01: Proc. of the 17th Annual Computer Security Applications Conference, page 358, Washington, DC, USA, 2001. IEEE Computer Society.

[7] R. Banham. The enemy within the key to a secure computer system is an honourable and educated workforce. CFO, 20(13):81–86, 2004.

[8] J. Blythe, L. J. Camp, and V. Garg. Targeted risk communication for computer security. In International Conference on Intelligent User Interfaces, 2 2011.

[9] V. J. Bono. Explanation and Apology. Merit NANOG Mail Archives, 5 1977.

[10] W. Bowman and L. J. Camp. Protecting the internet from dictators: Technical and policy solutions to ensure online freedoms. The Innovation Journal, 2013.

[11] C. D. Brummitt, R. M. D'Souza, and E. a. Leicht. Suppressing cascades of load in interdependent networks. Proceedings of the National Academy of Sciences of the United States of America, 109(12):E680–9, Mar. 2012.

[12] S. V. Buldyrev, R. Parshani, G. Paul, H. E. Stanley, and S. Havlin. Catastrophic cascade of failures in interdependent networks. Nature, 464(7291):1025–8, Apr. 2010.

[13] L. J. Camp. Reconceptualizing the role of security user. Daedalus, 140(4):93–107, 2011.

[14] A. Cardillo, M. Zanin, J. Go´ mez-Garden˜ es, M. Romance, A. J. Garc´ıa del Amo, and S. Boccaletti. Modeling the multi-layer nature of the European Air Transport Network: Resilience and passengers re-scheduling under random failures. The European Physical Journal Special Topics, 215(1):23–33, Jan. 2013.

[15] R. C. Chris Hall and R. Anderson. Resilience of the Internet Interconnection Ecosystem. European Network and Information Security Agency, 2010.

[16]   C. Chung, P. Khatkar, T. Xing, J. Lee, and D. Huang. Nice: Network intrusion detection and countermeasure selection in virtual network systems. 2013.

[17] R. Clayton. How much did shutting down mccolo help. Proc. of 6th CEAS, 2009.

[18] R. Cohen, K. Erez, D. Ben-Avraham, and S. Havlin. Resilience of the internet to random breakdowns. Physical review letters, 85(21):4626–8, Nov. 2000.

[19]   R. Cohen, K. Erez, D. Ben-Avraham, and S. Havlin. Breakdown of the Internet under Intentional Attack. Physical Review Letters, 86(16):3682–3685, Apr. 2001.

[20] D. Cornish and R. Clarke.   The reasoning criminal: Rational choice perspectives on offending. Springer-Verlag, New York, NY, 1986.

[21]   L. Cranor and S. Garfinkel. Security and usability: Designing secure systems that people can use. O'Reilly Media, Inc., Sebastopol, CA, 2005.

[22] P. Crucitti, V. Latora, and M. Marchiori. Model for cascading failures in complex networks. Physical Review E, 69(4):045104, Apr. 2004.

[23] P. Crucitti, V. Latora, M. Marchiori, and A. Rapisarda. Efficiency of scale-free networks: error and attack tolerance. Physica A, 320:622–642, 2003.

[24]   P. Crucitti, V. Latora, M. Marchiori, and A. Rapisarda. Error and attack tolerance of complex networks. Physica A: Statistical Mechanics and its Applications, 340(1-3):388–394, Sept. 2004.

[25]   V. Cupac, J. T. Lizier, and M. Prokopenko. Comparing dynamics of cascading failures between network-centric and power flow models. International Journal of Electrical Power & Energy Systems, 49:369–379, July 2013.

[26]   A. Dainotti, E. Aben, A. King, K. Benson, Y. Hyun, and K. Claffy. Monitoring large-scale internet outages. BGPMon, jun 2013.

[27] R. M. Dawes, J. McTavish, and H. Shaklee. Behavior, communication, and assumptions about other people's behavior in a commons dilemma situation. Journal of Personality and Social Psychology, 35(1):1, 1977.

[28] I. Dobson, B. a. Carreras, V. E. Lynch, and D. E. Newman. Complex systems analysis of series of blackouts: cascading failure, critical points, and self-organization. Chaos (Woodbury, N.Y.), 17(2):026103, June 2007.

[29]   G. Dong, L. Tian, D. Zhou, R. Du, J. Xiao, and H. E. Stanley. Robustness of n interdependent networks with partial support-dependence relationship. EPL (Europhysics Letters), 102(6):68004, June 2013.

[30]   Z. Dong, J. Camp, and J. Blythe. Beyond the lock icon: Inferring website categories from ssl certificates, 2013.

[31] J. Gao, S. V. Buldyrev, S. Havlin, and H. E. Stanley. Robustness of a Network of Networks. Physical Review Letters, 107(19):195701, Nov. 2011.

[32] V. Garg and J. Camp. Heuristics and biases: Implications for security design. Technology and Society Magazine, IEEE, 32(1):73–79, 2013.

[33] R. Govindan and H. Tangmunarunkit. Heuristics for Internet map discovery. In Proceedings IEEE INFOCOM 2000. Conference on Computer Communications. Nineteenth Annual Joint Conference of the IEEE Computer and Communications Societies (Cat. No.00CH37064), volume 3, pages 1371–1380. IEEE.

[34] S. Grabner-Kraeuter. The role of consumers' trust in online-shopping. Journal of Business Ethics, 39, August 2002.

[35] O. Gu¨ ell, F. Sague´s, G. Basler, Z. Nikoloski, and M. A. Serrano. Assessing the significance of knockout cascades in metabolic networks. pages 1–15, Oct. 2012.

[36] J.-L. Guillaume and M. Latapy. Bipartite graphs as models of complex networks. Physica A: Statistical Mechanics and its Applications, 371(2):795–813, Nov. 2006.

[37] P. Hines, E. Cotilla-Sanchez, and S. Blumsack. Do topological models provide good information about electricity infrastructure vulnerability? Chaos (Woodbury, N.Y.), 20(3):033122, Sept. 2010.

[38] R. Hiran, N. Carlsson, and P. Gill. Characterizing large-scale routing anomalies: a case study of the china telecom incident. In Passive and Active Measurement, pages 229–238. Springer, 2013.

[39] P. Hunter. Pakistan youtube block exposes fundamental internet security weakness: Concern that pakistani action affected youtube access elsewhere in world. Computer Fraud & Security, 2008(4):10–11, 2008.

[40] J. H. Jafarian, E. Al-Shaer, and Q. Duan. Openflow random host mutation: transparent moving target defense using software defined networking. In Proceedings of the first workshop on Hot topics in software defined networks, HotSDN '12, pages 127–132, New York, NY, USA, 2012. ACM.

[41] N. L. Kerr and M. KAUFMAN-GILLILAND. Communication, commitment, and cooperation in social dilemmas. Journal of Personality and Social Psychology, 66(3):513–529, 1994.

[42] S. Kiesler, L. Sproull, K. Waters, et al. A prisoner's dilemma experiment on cooperation with people and human-like computers. Journal of personality and social psychology, 70:47–65, 1996.

[43] J. Kirk. At&t facebook traffic takes a loop through china. Computerworld, (4), 2011.

[44] M. Kurant, P. Thiran, and P. Hagmann. Error and attack tolerance of layered complex networks. Physical Review E, 76(2):026103, Aug. 2007.

[45] L. Lafferty. Medical identity theft: the future threat of health care fraud is now. Journal of Health Care Compliance, 9(1):11–20, 2007.

[46] V. Latora and M. Marchiori. Efficient Behavior of Small-World Networks. Physical Review Letters, 87(19):198701, Oct. 2001.

[47] N. Leveson. Engineering a safer world: Systems thinking applied to safety. MIT Press, 2011.

[48] N. G. Leveson and C. S. Turner. An investigation of the therac-25 accidents. Computer, 26(7):18–41, 1993.

[49] L. Li, D. Alderson, W. Willinger, and J. Doyle. A first-principles approach to understanding the internet's router-level topology. ACM SIGCOMM Computer Communication Review, 34(4):3, Oct. 2004.

[50] T. Moore, R. Clayton, and R. Anderson. The economics of online crime. The Journal of Economic Perspectives, 23(3):3–20, 2009.

[51] R. a. Mortier. Control and understanding: Owning your home network. In Communication Systems and Networks (COMSNETS), 2012 Fourth International Conference on, pages 1–10, 2012.

[52] A. Motter and Y.-C. Lai. Cascade-based attacks on complex networks. Physical Review E, 66(6):065102, Dec. 2002.

[53] A. Newell and H. Simon. Human Problem Solving. Prentice-Hall, 1972.

[54] D. Newth and J. Ash. Evolving cascading failure resilience in complex networks. Complexity International, 11(2005):125–136, 2005.

[55] H. Nissenbaum. Securing trust online: Wisdom or oxymoron. Boston University Law Review, 81(3):635–664, June 2001.

[56] M. Omer, R. Nilchiani, and A. Mostashari. Measuring the resilience of the global internet infrastructure system. In Systems Conference, 2009 3rd Annual IEEE, pages 156–162. IEEE, 2009.

[57] A. Ozment. Bug auctions: Vulnerability markets reconsidered. In Third Workshop on the Economics of Information Security, Minneapolis, MN, USA, June 2004.

[58] T. Parfitt. Georgian woman cuts off web access to whole of Armenia. The Guardian, apr 2011.

[59] C. Partridge. The Internet under Crisis Conditions: Learning from September 11. The National Academies Press, February.

[60] P. Porras, S. Shin, V. Yegneswaran, M. Fong, M. Tyson, and G. Gu. A security enforcement kernel for openflow networks. In Proceedings of the first workshop on Hot topics in software defined networks, HotSDN '12, pages 121–126, New York, NY, USA, 2012. ACM.

[61] R. Posner. Rational choice, behavioral economics, and the law. Stanford Law Review, 50(5):1551–1575, 1998.

[62] E. Romijn. Ripe ncc and duke university bgp experiment, 2010.

[63] S. Sheng, B. Magnien, P. Kumaraguru, A. Acquisti, L. F. Cranor, J. Hong, and E. Nunge. Anti-phishing phil: the design and evaluation of a game that teaches people not to fall for phish. In Proceedings of the 3rd symposium on Usable privacy and security, SOUPS '07, pages 88–99, New York, NY, USA, 2007. ACM.

[64] S. Shin, P. Porras, V. Yegneswaran, M. Fong, G. Gu, and M. Tyson. Fresco: Modular composable security services for software-defined networks. Internet Society NDSS (Feb. 2013). To appear, 2013.

[65] P. Slovic. Perceptions of risk. Science, pages 280–285, 1987.

[66] H. Tangmunarunkit, R. Govindan, S. Jamin, S. Shenker, and W. Willinger. Network topology generators: degree-based vs. structural. In Proceedings of the 2002 conference on Applications, technologies, architectures, and protocols for computer communications SIGCOMM '02, page 147, New York, New York, USA, 2002. ACM Press.

[67] A. Toonk. How the internet in Australia went down under. BGPMon, 2 2012.

[68] D. Towsley. On distinguishing between Internet power law topology generators. In Proceedings.Twenty-First Annual Joint Conference of the IEEE Computer and Communications Societies, volume 2, pages 638–647. IEEE, 2002.

[69] O. Yagan and D. Cochran. Optimal Allocation of Interconnecting Links in Cyber-Physical Systems: Interdependence, Cascading Failures, and Robustness. IEEE Transactions on Parallel and Distributed Systems, 23(9):1708–1720, Sept. 2012.

[70] C. H. Yeung and D. Saad. Networking - A statistical physics perspective. Journal of Physics A: Mathematical and Theoretical, 46(10):103001, Mar. 2013.

[71] R. Zeckhauser. Comments: behavioral versus rational economics: what you see is what you conquer. The Journal of Business, 59(4):435–449, 1986.

APPENDIX: CRITICAL USE CASES

The end result of this research was a layer model for security threats and potential in SDN. The starting point was a series of use cases. The use cases are included in the following Appendices.

In order to scope the problem of protecting future Software Defined Networks, We have selected four families of critical use cases. There is some commonality between them, but their differences are likely to lead to diversity in threat models and architectures. The most important dimensions that are highly variable are one of trust of the larger network, the type of failure cases that are intolerable, and the incentives which different stakeholders have to contribute to the protection, and more generally the dependability, of the switching fabric.

Our first use case is the data center, which is also the focus of initial deployments of SDN technology. Here the system is under the control of a single actor with well-known topology and performance requirements; the case for SDN deployment rests heavily on the ability to save costs, both in terms of moving from proprietary to commodity hardware and also in terms of providing higher level abstractions to enable large fleets of similar machines to be managed in an efficient and centralized way. Operators will use SDN tools to deliver not just costs savings but also high levels of availability using virtualization and redundancy. Although they are in complete control of their platforms, they may have tenants who compete with each other, and must therefore provide strong separation between mutually suspicious customers.

The second use case is the ISP or IXP. Larger scale network service providers also face severe cost pressures (the industry has recently been shaken out and consolidated by price competition) and may have customers who are mutually distrustful; consider for example an ISP whose customers include several dozen banks, whether in Manhattan or London or Tokyo. Their operations are on a much larger scale than a data center with a wide range of competence and experience of personnel across the globe. The majority of network failures are a result of error and misconfiguration, not hostility. The promise of SDN is not just to save costs but to improve reliability by making networks more manageable, less liable to configuration errors, and easier to repair when things go wrong.

The third use case is the large cyber-physical system. Here we consider two examples: a large airport (think JFK or Chicago or Heathrow) and a large industrial control system (think a large petrochemical complex or a future regional smart grid). For our example, the generic International Airport "IAirport" has 180,000 staff working for 3000 companies, including mutually hostile carriers from countries at war or with deep-rooted aggression. Airports are potential targets for terrorist activity, or for strategic attacks by a hostile state in times of tension; the same holds for energy supply. The scale and complexity of the interactions along with the risk provide a microcosm of the challenges to be met in a safety-critical organization that combines private-sector and public-sector operations.

The fourth use case is the military and intelligence community. The next generation war-fighter, in contrast, has decision-making even more acutely grounded in life and death and is likely to face capable motivated opponents in times of tension. The rapid deployment of large forces is likely to depend not just on dedicated military communications equipment, but increasingly on the ability to acquire virtual networks from civilian providers as needed, and to configure and defend them appropriately. These four cases are intended to provide concrete discussion points in the large space of possible future SDN deployment. In each case the studies begin with the operational requirements including descriptions of the likely and conceivable failure modes and attacks. Stakeholders are described, including bystanders and possible adversaries. From this, the use cases move to evaluations of the attack surfaces with a

particular focus on the conditions under which attacks are immediately detectable and those under which subversion may remain unknown for some time (as with Stuxnet). The use cases close with architectural options for mitigating the use and abuse cases, constrained by engineering limits and informed by first-order deployability implications.

The use cases all include SDN federated with BGP/BGPSec either within the operation (e.g., Airport, ISPs), at the boundaries (i.e., Data Centers), or dynamically as the location and infrastructure availability on the ground changes. The cases assume FlowVisor, to provide continuity, and lightweight TCP. Authentication requirements will evolve from each case, and not be determined a priority.

A1 Battlefield Use Case

The Armed Forces has widely varying network contexts. There are the established networks, from Pentagon to Pacific, that are either classified or unclassified. Classified networks ideally are implemented with an air gap. Unclassified networks are in a state of cyberwar. It is the state of assault by nation-states and actors with equivalent resources that distinguish this use case. Economics of attack and defense are arguably inapplicable when the adversary has effectively unlimited funds and is not seeking monetization.

A second, similar category of networks are ones that are established, but mobile. These networks include the ones found on naval ships. Unlike the first category of networks, these have to operate under very strict limitations. Equipment failures cannot be fixed by simple replacements if a vessel is thousands of miles from the nearest friendly port. Therefore, equipment must undergo rigorous certification procedures, making tasks as simple as firmware upgrades long, arduous procedures. Additionally, the rigid requirements of the network make enforcing network compliance a difficult task when housing personnel that bring their own devices.

The final category of networks of interest are those that must be immediately deployed, often in domains with little preexisting infrastructure. The preexisting infrastructure may be putatively under the control of allies; however, even in this case the insider threat is so extreme as to make these network components effectively untrustworthy (e.g. vehicles in the front lines of combat). In this case operators themselves are untrustworthy without considering the more fundamental question of SDN as allowing trusted operations on untrusted hardware.

A1.1 Operators and Stakeholders

At the most fundamental level, the operators and stakeholders in the battlefield or disaster preparedness case are the people in the boots on the ground. Too formal a set of requirements for the devices software as a service will result in subversion. The capacity to isolate the networks that are created by personnel bringing their own devices combines these. For example, requiring that no personnel bring any gaming devices to extended deployments may appear a reasonable policy, particularly to senior officers whose childhood included more Parcheesi than Princess Peach. However, such constraints are likely to be subverted in practice. Individuals in high-stress and high-risk situations are more likely to subvert policy for their own requirements for communication and stress release. Currently, the vast majority of security policy violations are in the theater. By creating the possibility of true isolation between bring-your-own-device networks and the operation networks, SDN offers the ability to set appropriate multi-level requirements. Recognizing individuals in the field with their own devices as legitimate stakeholders can be integrated into a resilient SDN.

A second set of stakeholders are allies who may include less trusted insiders who are working with military personnel. The tragic increase of blue on green violence in 2012 is testament to the limitations of political alliances to ensure allegiances of individuals. Because of the complex interaction of political

and military, proximity authentication and requirements are quite distinct from the commercial domain.

At a higher level, stakeholders include the operators of classified and unclassified networks. For military networks this includes identity providers and trusted certificate providers. Included here are those who test, maintain, and upgrade the networks. The short product lifecycle characteristic to information technology exacerbates the conflict of interest between network operator and those responsible for assurance. As new products and methods of communication become available to individuals in the field, those responsible for authentication find the laboratory assumptions no longer hold. A firewall that is secure today may see the creation of a tunnel by a new consumer innovation brought by an airman; which then allows an attacker through the network into the less secured interior. An action as simple as upgrading the browser may be necessary for security or may instead introduce a range of vulnerabilities.

Testing and upgrade components are unique and critical to the defense space. Formal auditing expense and difficulty is a critical barrier to adoption and upgrading of military systems. SDN offers the ability, in the production of few core hardware components, to implement a wide range of network functionality while keeping the complexity of the firmware on network hardware relatively simple and static. For example, a new routing protocol can be implemented in the controller and applied to the entire network without any firmware changes to the network hardware.

Battlefield networks have a distinction that may allow for more effective use of SDN: a hierarchical network of distributed authentication tokens or 'cat cards'.

A1.2 Technical Operations

The three components of threats above match to a higher-level conceptualization of SDN.

The closing air gap between classified and unclassified networks is a case of secure internetworking. The essential functional requirements are isolation and ideally even invisibility.

Compartmentalization is a requirement for multiple SDN applications, such as cyber-physical systems. In the Next Generation Battlefield there is a requirement for dynamic compartmentalization. As components are added or become unreachable the system must retain the original security policy. The system not only will maintain its integrity but also be able to introduce new elements quickly. In a chaotic environment being unable to use the network may be life threatening in the battlefield. Users unable to use the network may also use insecure channels if the system is unreliable or inflexible. The Next Generation Battlefield requires Battlefield and other defense networks with a greatly increased capacity and number of devices. For example high volumes of visual data may need to be transmitted and analyzed to deal with a new physical threat.

Currently, DoD networks require extensive change management procedures and require significant analysis of any change made to a network to determine the impact on security. SDN technologies may allow for the rapid reconfiguration of the network to add additional circuits and network hardware. The logically centralized controller topology allows new components to be added without having to verify the entire system. Conformance can be broken down into multiple separate tasks of ensuring the security policy is being maintained by the software; ensuring secure, reliable connectivity; and verifying the network device properly enacts the rules specified by the software. The core value of SDN, that is networking logic that is hardware agnostic, inherently aligns with the testing and upgrading requirements for configuration management. Hardware upgrades would no longer require extensive testing of various network protocol implementations on the hardware; and, conversely, network

software changes can be made without having to upgrade and test the hardware. SDN specifications provide a clearly defined common denominator that hardware manufactures and software engineers can target to guarantee interoperability between the two.

The air gap between classified and non-classified networks is closing. Deployed tactical systems seek isolation and invisibility from other networks, as opposed to more classic requirements for assurance. OpenFlow offers the exact specification of what can be connected. Imagine connecting an external subsystem that only needs to talk to one or two devices or systems. In a SDN network it would be possible to leverage the existing ethernet infrastructure to allow the new external subsystem to address only those that are desired, and indeed to find only those for which connections are desired. The closure of the air gap is a challenge of internetworking resilience.

Physical reconfiguration is a unique challenge on the battlefield because it occurs at high stress moments during physical attack. Orchestration and automatic reconfiguration provide ways to mitigate the issues of network reconfiguration under battlefield conditions. SDN controllers and other orchestration systems could react to the changing conditions and automatically reconfigure the network and other resources. SDN networks can resist degradation such as radio interference of the communication channel by using redundant transmission of the message on different frequencies or type of transmission technology.

Weaponry and armor will be networked, as exemplified by the goals of the Joint Tactical Radio System. The ideal of every soldier being a data command point with the ability to send and receive data-rich images and to provide different ways to communicate. The increased use of large data flows from many different sources and receivers makes the networking more difficult. Not only does the network have to handle increasingly large data flows, the direction of the information has changed. Information gathering and retrieval can be done at any location giving each soldier a greater awareness of the battlefield conditions. However reliability, security, and performance of the network transmitting this information needs to be assured. Multipath transmissions may allow greater bandwidth utilizing all possible paths to transmit high priority information. Critical transmissions can receive guaranteed bandwidth or utilize quality of service parameters configured in a SDN protocol.

A1.3 Vulnerabilities

While it's the case that network service providers and data centers require isolation, these are much less likely to face advance, persistent threats of the nation-state level. It's certainly the case that physical attacks are a unique concern.

The military network is characterized by hardware tokens that identify individuals and devices. The unique vulnerability is that these can be captured and not identified as being held by hostiles. Captures of individuals, devices, and tokens are serious issues. In the battlefield, possession of devices can change suddenly and not be detected for hours.

Information sharing requirements in the military do not map directly to industry requirements. In the case of isolation it may be acceptable for a line employee to be isolated, and to wait for technical assistance. However, such delays are not tolerable in the battlefield or emergency environment. Conversely, overcommitting to resources is a standard approach to deployed network configuration. Thus the capacity to isolate quickly upon risk of subversion may be more acceptable in a military domain than in an industrial application. Perversely, more than any other domain, respect for the restraints on attention span are most critical in a battlefield.

- The Black Box Full characterization and qualification of a component (e.g. NIC, switch, router, etc) is made difficult by elements of the component that cannot be subject to full inspection and classification. These potentially unknowable elements regularly include firmware and programmable logic devices. Mitigations to the black box problem often include high-coverage testing and guard elements which increase system cost and development time. These mitigations increase cost, and must be repeated for every new device connected to the network. As different generations of hardware connect to the network, it becomes more challenging to examine every possible interaction.
- Configuration management. Even very minor changes to a systems hardware or software configuration, can bring about very significant changes in behavior. For this reason, configuration management in critical real-time systems is a high priority. Commercial-off-the-shelf (COTS) products often have no requirement or incentive to maintain configuration management. In fact, the pressures of the commercial domain (cost priorities and lean operating principles) often creates an environment that is not conducive to constraints on configuration.
- Coordinated (and restricted) interfaces. Large systems often categorized as systems of systems often have independently developed subsystems that rely on highly-coordinated interfaces. While many commercial standards (IEEE 802.3 is just one example) meet the performance requirements, use of such powerful interfaces presents a large threat surface and an enormous testing challenge. Accordingly, such broad interfaces must be profiled to assess the necessary subset of features that are required and subject to validation. Similarly, there is interest in managing interfaces such that subsystems cannot be easily spoofed. There has been interesting work in this area that includes domain specific languages (DSLs) (see Meredith Patterson 28C3 talk).

This is exacerbated but the challenge that military authentication does not happen at the network level. It happens with certificates on the machine, using certificates across the network, and in informal (often policy-violating) interactions in the field.

The interaction of survivability and isolation must be function-dependent. The application isolation of SDN may angle a more refined definition of the value of operational connectivity, as some applications are better isolated while others should allow more risky connections. Consider a metaphorical analog in the physical realm. For example, Humvees and tanks have push start buttons because denial of service is a larger threat than having the equipment overrun. If it's necessary to move, moving takes precedent over security concerns. However, access to a tank does not necessarily provide authenticated access to the weapons systems. Similarly, networks carry a wide range of applications, and the ability of SDN to provide varying levels of service and survivability to different applications is a potential game-changer. If a truck is fully networked, and the back and top are blown off the truck, the communication and functionality of the truck must remain. If a truck falls into enemy hands, there is no feasible way to prevent its use. However, if the operators of the truck are unable to provide technical credentials then denial of service to weaponry may be the correct operational choice.

A1.4 Organization Requirements.

The reality of military configuration is that preparation may have high levels of resources, capacity for testing, and well-controlled authentication before configuration. However, many of these will be highly compartmentalized in order to prevent an internal adversary or inadvertent leakage to provide

information on scale or participants of an action. Once an action has begun, and at any moment after initiation, moments of reconfiguration are extremely likely to be associated with cognitive overload and data uncertainty.

• The Black Box SDN, or more specifically OpenFlow, implementation could allow a standardized method of interfacing with the hardware component, eliminating unique programming interfaces for different devices. Use of SDN could allow for standardized test cases which could be used across multiple products with minimal change driving commonality during evaluation and qualification.

• Configuration management. By constraining actions a priori, the behaviors of a device can be constrained before and during deployment, and should a device exhibit behaviors beyond its constraints then it can be disabled or removed. This allows an organization to place integrity or confidentiality above survivability; and to develop a flexible policy for these trade-offs in different situations.

• Coordinated (and restricted) interfaces. The hardware devices and pre-configuration can provide necessary isolation in daily operations. However, it is famously true that no operational plan survives actual engagement. Adding real time level of authentication can enable the network to provide more fluid permissions. Combining the hardware with social engagement in defined operational situations can enable robust responses to changes in operation. For example, expanding x590 to include some human-readable but network-authenticated challenge-respond to enable the level of human authentication that is necessary can limit misuse of automated recognition of captured devices, but create a level of vulnerability to social engineering. Examination of trust mechanisms beyond the isolated technical certificate is necessary to address the integration of cultures (e.g., different regional first responders or allied forces) and organizations.

A2 Cyber-Physical Use Case 1: International Airports

The third class of use case is the cyber-physical system. We give two examples; an airport and a large industrial control system.

In a large international airport, the network is behind the scenes but still visible in interacting with a large number of diverse users. It is a complex physical environment relying on a robust network to ensure its daily operations. We could discuss this in the context of any large international airport (JFK, O'Hare, LAX, Hong Kong) but for simplicity we will refer to this setting as "IAirport".

A2.1 Operators and Stakeholders

IAirport has 180,000 staff (in the sense of people who have been issued with an access badge); they work for some 3,000 organizations, ranging from the IAirport management office through the major contractors and the airlines down to small franchise operators, like the concourse restaurant operations. Some organizations are arms of the Host State (FAA, TSA, BAA, Customs, etc..) and deal with classified information, others are arms of other states, such as a foreign nations National Airlines (British Airways, Air Canada, Air China, Air Koryo, etc..). The only common background check performed on IAirport staff is that they are not on various blacklists of terrorist suspects or convicted of serious crimes; in the event of the foreign employees of foreign organizations, real checks are not possible.

All of these organizations must use a single network; it would not be acceptable for 3,000 firms to have their own cabling and other infrastructure. The network must also serve the general public, via wi-fi hotspots where passengers can pay for service via dozens of providers and roaming agreements; it also serves emergency services by supporting mobile radio of various kinds. There are more specialized shared services, such as a broadcast data channel which announces which plane is at which gate and a communications system that relays data to aircraft. Many of these shared services are critical; without

the aircraft communications system, for example, airlines would take an extra 15–45 minutes to turn round aircraft as flight plans would have to be filed manually.

These numerous users and stakeholders need the network for different purposes and thus have different requirements. We identify four main groups of users: infrastructure-owner, main customers, auxiliary customers, and general public.

The owner, or its lead contractor, lays out the network infrastructure. This is crucial as many things must be got right to give users the dependability guarantees they need. The infrastructure sets the playground for the others to join. The owner is the key decision-maker on network requirements, design, installation, and maintenance. Without proper planning, a failure at this layer will affect the business of the whole operation. In the case of IAirport, the host government also has a vested interest as continued operations are of national importance. An attack that significantly disrupted airplane or passenger movement could have macroeconomic effects even if it did not lead to casualties.

The second group is the 'anchor' tenants such as the major airlines, the police, and border and customers services. These customers may have similar needs, but the airlines (for example) are in competition with each other. In more extreme cases two state-owned airlines may represent two countries at war – so the case for secure separation between users is strong. These customers not only need to connect their machines and devices together at the location and to servers offsite, but also to take over shared facilities on a temporary basis (as when an airline uses a dozen check-in desks for a few hours before one of its planes departs). They may want the network to roll out special technologies to support their business. Both airlines and law-enforcement agencies need the network to run their core business: if an airline cannot get through to its booking system (typically run by a third-party contractor) then it does not know which passengers may board which plane and its operations cease (or at best continue in highly-degraded form, relying on paper documents that more and more customers simply don't carry).

Auxiliary customers such as restaurants, shops, and third-party wireless Internet providers are not critical to the airport's operations in the sense that passengers could board and planes could fly without them. However the rent they pay is a critical part of the airport's operating income, and if their network service were interrupted for a substantial period there would be consequences. For example, high-value stores could not operate profitably without the ability to do online payment card transactions. So these users' availability requirement may not be very high, but it is not zero. Furthermore, as auxiliary networks are used by many staff who have had only cursory background checks, they can be an entry point for malicious actions.

The general population is the largest group of users by headcount. Many participants within this space need to interact with the network: passengers, hotel guests, shoppers, and service staff. There are also criminals from baggage thieves to people who set up bogus wi-fi connections to conduct phishing and pharming attacks. So the general public access may be another entry point for tactical or strategic attacks.

A2.2 Technical Operations

A cyber-physical network needs to provide several properties: low latency, high bandwidth, resilience, virtualization, and flexibility. An application may require one or more of these properties to function.

Latency is most important in infrastructure-type functions. Timely feedback of individual sensors throughout the building, such as CCTV cameras and fire alarms, is essential to continue operations. Emergency radio equipment is also critical. Latency is achieved by both having a good network design

and the ability to prioritize traffic from critical applications.

The large number of staff, firms, and passengers also use a lot of bandwidth. Advertising, CCTV, flight information, and business transactions are joined by passenger wi-fi use to soak up the megabits.

Resilience, too, is important as an airport depends on many critical applications, from obvious ones such as check-in and baggage handling to the less visible such as the RFID locks on thousands of doors, the communications systems used to send and receive data from aircraft, and the classified government systems that connect border agency staff with visa and blacklist systems.

Virtualization is important for individual airlines, shops and other firms operating in the airport as they need to connect to their corporate networks over the network operated by IAirport. Virtualization must also support separation, to prevent virtual networks interfering with each other and denying service. It is not desirable for any of the airport's tenants to be able to deny service to critical flight-operations or national security systems.

Lastly, networks need to accommodate the many new devices and protocols that will emerge over the lifespan of a business such as an airport (which has been in operation for over 60 years). The deployment of new networks and services must be possible without interfering with legacy communications and applications.

In order to give some idea of the scale, one single terminal at IAirport has

• 1000 fixed and 500 mobile video cameras (10Gb/s) • 500 displays (10Gb/s)

• Biometric scanners (10Gb/s)

• Private and Public Fixed and Wireless LAN (20 Gb/s)

• Cellular services (10 Gb/s)

• Mission Critical Voice/Data and private radio (0.5 Gb/s)

• Passive RFID (0.2 Gb/s)

• Active locatable RFID (5 Gb/s)

The assumed aggregate peak rate is in the tens of Gb/s.

A2.3 Vulnerabilities

At a high level, we are concerned with both strategic attacks and tactical attacks. Strategic attacks involve an attempt to close down or cripple the airport, perhaps as a service-denial attack in wartime. In the IAirport case, apart from the most obvious case of cutting the power for a complete network outage, a targeted attack on the passenger booking system can result in total chaos as airline employees need to fall back to paper passenger lists, while a takedown of the communications system can delay plane turnaround time by forcing flight plans to be filed manually.

Tactical attacks have smaller goals, for example if Country-A Air tries to sabotage the operations of Country-B Air. This is just an example of what can happen when many mutually mistrusting parties operate in the same network. In fact, we find this mistrust everywhere: not only do participants of the same function not trust each other (as with airlines); different applications within the same network also distrust each other. The network delivers traffic for safety critical systems (e.g. emergency radio), operational critical systems communications system, and public wi-fi, all with the same physical infrastructure. How do we ensure, for example, that untrusted applications don't deny service to critical ones by hogging bandwidth?

Gaining physical access to devices can potentially allow an adversary to replace a device with one of his own. In this attack, the attacker can do much more than eavesdropping but can actively participate in the network. With an owned device, such as a switch, it is possible to cause havoc, such as by repeatedly advertising and tearing down bad routes, leading to a network collapse. More subtle attacks can involve tromboning – altering network configuration so that a competitor's traffic passes a machine controlled by the attacker, to facilitate eavesdropping, traffic analysis, or targeted service denial. Such an attack may be hard to detect, if carried out competently. SDN has the potential to support robust virtualization that would make such attacks harder.

Similarly, a software vulnerability exploitable over the network can potentially allow an attacker to take over one or more routers or other devices. Such exploits can potentially allow access to the device to be escalated into an attack on the network. Standard software vulnerability classes and attack surfaces apply here, as routers and switches are based increasingly on commodity software and standard attacks may be used in the time window between a vulnerability being disclosed in Linux or FreeBSD and its being patched in network hardware employing a version of such software.

A2.4 Requirements for a Resilience Architecture

The required resilience of the network is to ensure network connectivity as well as the quality of service for a number of critical services in decreasing priority order, starting with safety and emergency communications, down through flight operations, through airline systems to the services offered to the general public. Further virtual network separation is desirable to protect airlines against interfering with each other's service, whether accidentally or otherwise. The network managers must be able to measure and monitor changes and faults and modify network topology and behavior in response as appropriate.

A2.5 Organization Requirements

Apart from technical and operational requirements, there are three linked business requirements from the network operator: the need for abstractions, the need for automation, and the need to reduce costs.

The network in a setting such as an airport is complex and expensive. Current management tools are inadequate, being based on router command lines that differ from one vendor to another and which don't support the atomic, consistent, isolated, and durable transactions which network operators really need. The lack of appropriate abstractions entails a pervasive lack of contextual information which is not just inconvenient but can easily result in operator errors. Abstractions can hide the parts of the network that are not relevant to the task in hand; this is just elementary computer science, and applying it to networking is one of the big promises of SDN.

Abstraction will support another requirement from operators, namely better automation. Current network technologies leave too much room for human error. Some tasks, such as adding a new airline or shop, may be repeated many times and should become routine; dependencies with local implementation detail must be better hidden to make this simple and dependable.

The main driver in day-to-day operations is of course, cost. The air transport industry operates on tight margins. Investments in new technologies will only happen if there is an unavoidable regulatory mandate, or to save money.

A3 Cyber-Physical Use Case 2: Industrial Control Systems

Industrial control systems are also of critical concern because of a modern state's near-total dependency on civilian utility networks.

A3.1 Operators and Stakeholders

The most important utility from the critical-infrastructure viewpoint is electric power, without which almost everything else comes to a standstill; but there are other critical systems such as oil refining, railway signaling, and water treatment. It was realized fifteen years ago that such networks were becoming exposed to cyber-attack because of the rapid adoption of IP networking. The protocols most commonly used in industrial control systems evolved in a world of closed and dedicated networks with no need for authentication or encryption. The move to IP was driven by cost pressures but left operators vulnerable; anyone in the world who knew the IP address of a sensor could read it, and anyone who knew the address of an actuator could operate it. Since the alarm was sounded in 1998, and especially since 9/11, considerable efforts have been expended by both state and private-sector actors in protecting critical control systems. In what follows we will discuss the electricity industry; similar comments apply, mutatis mutandis, to petrochemicals, signaling, water-treatment, and indeed industrial production.

A3.2 Technical Operations

A small installation such as an electricity substation might have 100–200 programmable devices attached to a substation LAN, including transformers, circuit breakers, reclosers and meters. Traffic on the LAN is not encrypted or authenticated, as there are stringent latency requirements, so anyone with access to the wiring could disrupt operations. Anyone with physical access could do this anyway by operating manual override switches, so the issue is whether an attacker might get remote access to a device on the LAN. Security at present depends on a station controller, which is on the LAN, and a gateway which is attached to the controller and also to WAN communications (typically over the Internet to a network operations center, protected by TLS). It's critical that these devices not be vulnerable to remote software attack, and that they provide effective protection to internal devices.

In effect, the security architecture is one of re-perimeterisation. It is not generally feasible to retrofit authentication or other security mechanisms because of the variety of equipment whose service life is generally measured in decades rather than years. Work is underway to agree on new versions of control system protocols that do support authentication; perhaps within five years new equipment will support this. However it is likely to be decades before most systems are replaced.

The same applies to larger installations such as power stations and network control centers. Here however the re-perimeterisation is much more complex. A power station may have communications at five different safety integrity levels:

- the safety systems will typically be at SIL 3, and must not be vulnerable to interference or service denial attacks from any lower level. The safety systems prevent failures leading to loss of life or catastrophic damage to plant; for example, by closing down a nuclear reactor if the reaction exceeds specified limits.
- the control systems will typically by at SIL 2, and also must not be vulnerable to interference or attacks from below.
- monitoring systems will typically be at SIL 1. Although they cannot affect higher levels directly, the loss of monitoring systems can make control systems unusable leading to a precautionary

plant shutdown. So they too must be protected from problems at lower levels.

- Below this are the plant's executive information systems and business processes such as invoicing and payroll. Although these systems lie below the mandatory access control framework of the SILs, there may be a business case for further network segregation, for example to protect internal financial systems from Internet-facing web and mail servers.

As in the airport case, the applications at these levels have differing requirements for latency, bandwidth, resilience, virtualization, and flexibility. There is nothing like the diversity of organizations, but there is some: the vendors of various pieces of equipment will have maintenance access, as will control-systems contractors. Here the issue is not so much the management of a complex high-bandwidth network with some separation requirements, as the maintenance of high-quality separation between critical networks in a complex environment where separation can easily break down – as we shall now discuss.

A3.3 Vulnerabilities

Power stations, network control centers and large-scale substations are likely targets in the event of cyber-attack by a hostile state, or by a capable substate group such as militant environmentalists. A power station has been affected accidentally by malware when a flash worm spammed the monitoring network, which would have caused a safety shutdown had it been operational at the time. A more deliberate attack might follow the Stuxnet model, with targeted malware introduced via USB drives left abandoned for an operator to find and introduce to the network. An important line of defense is to prevent software making its way from open systems to the SIL1 and higher domains. The strict network separation that SDN networks can support is attractive here.

A typical system has network vulnerabilities that arise spontaneously. A search engine built to discover control systems found over a thousand of them accessible on the Internet. Traditional network management technologies make it hard to manage separation dependably; the combination of complexity and obscurity makes it hard for people to understand what's connected to what, and as people modify things to get their work done, paths open up to the wider Internet. The principal benefit of SDN lies in providing much better tools to enforce perimeters, providing high assurance that critical sensors and actuators never become accessible from outside.

A3.4 Requirements for a Resilience Architecture

As before, a resilient network will ensure connectivity and quality of service for critical services in decreasing priority order, starting with SIL3, then SIL2, then SIL1, then corporate communications. It will also use virtualization to protect each layer against lower layers.

A3.5 Organization Requirements

A power station is, like an airport, a major capital asset with a lifespan measured in decades. Even if an individual nuclear reactor is decommissioned after 40 years, it's common for new reactors to be built next to old ones, as the local communities accept them and value the jobs. Power station operators are regulated; the regulation is even fiercer for transmission and distribution operators who maintain the power grid. A big issue in the USA (and also in much of Europe) is that security expenditures are not part of the regulated cost base, and thus they come directly off the company's bottom line. As the operators are typically funded by debt as well as equity this creates major resistance to any security investment that cannot save money directly.

In the USA, attempts to make power station operators invest in information security backfired in 2009. FERC/NERC regulations stipulated that all critical assets had by then to meet minimum information security standards; in the case of generating plant, a 'critical' unit was one with black start capability – the means of coming online in the absence of grid power. Hydro plant does have black start capability; nuclear does not; and coal-fired stations generally only do if they have auxiliary diesel generators. Operators' response to the FERC/NERC regulations was to scrap diesel generators rather than installing firewalls, thus making the US power grids less resilient.

So, just as in the airport case, an attractive feature of SDN is that it has the potential to save money as well as providing more robust network virtualization. Its deployability in lean commercial environments is one of its strong points.

A4 Data Center Types

Modern data centers exhibit vast architectural diversity, however, the differences can be abstracted by two sets of models that capture relevant security aspects. The first one focuses on defining administrative domains, while the latter is more concerned with identifying and defining stakeholders and their incentives.

Boundaries between different administrative domains dictate how a specific security mechanism is implemented, while, in turn, the assumptions on trust (as well as other factors) dictate where that boundary is.

A4.1 Structural Models

• End-host virtualization allows an operator to retain control over the end. This is very important since it allows complexity to be moved out of the network. For instance, an isolation or a denial-of-service prevention mechanisms can be both implemented inside of a hypervisor (push-back filters) and the network (VLANs). Although functionally equivalent, they have varying effects on performance, scalability, cost and security.

• Network virtualization allows a tenant to gain control over its network. The meaning of control is somewhat ambiguous, since it can refer to multitude of things including resource control, routing control, and other things.

A4.2 Organizational Models

• Multi-Tenant data centers may have hundreds of thousands of customers all utilizing the same resources. Competitors may be using the same infrastructure posing difficult issues with separation. Denial of Service attacks are exacerbated as a significant 'insider' threat. As the infrastructure is available for any customer, an attacker may have privileged access if it can be collocated with their victim. Stakeholders not only include the tenants of data centers but also their customers. A security incident affecting a multi-tenant data center may not only affect the direct tenants but all of the data of the customers of applications running on the infrastructure.

• Private/Enterprise data centers may seem to have fewer stakeholders than a multi-tenant data center but there are still issues for their customers and the threat of data loss and security. Individual enterprises may have specific requirements dependent on their business (financial, medical, etc.). Many businesses, especially small or mid-sized enterprises, do not have experienced personnel capable of understanding the security implications of a network change. In short, enterprise data centers are very similar to the multi-tenant ones, however, more emphasis is put towards cost and/or security.

• Content Providers may have millions of customers using their services. What differentiates a

content provider is that is has no notion of tenant, and consequently, security is shifted towards availability and resource allocation. Any disruption to the service can have widely reaching effects. Content providers such as social networks and cloud storage deliver increasingly large volumes of data. The loss of a day's service by (say) a video rental firm can have a major business impact.

A4.3 Incentives

Initially, when the first data-center operators started building their infrastructure, they ran into an awkward problem of adopting the existing enterprise tools (such as VLANs and firewalls) and practices to an ill-suited environment. Consequently to this day we are still lacking proper primitives to describe desired security and isolation policies between tenants. Drivers for SDN deployment include:

• Network virtualization and isolation are major drivers of possible SDN deployments. It provides mechanisms to segregate traffic for both security and traffic engineering purposes. It also covers attempts to rectify the limitations of existing protocols.

• Fine-grain control over network resource allocation would provide the ability to enforce variety of policies between tenants, therefore, mitigating certain denial of service attacks. The choice of policy is left to the tenants or the application.

• Consolidation of middlebox and network in networks that contain large numbers of 'middleboxes' such as firewalls, load balancers, wan accelerators. These boxes not only complicate the network and increase costs; they also can provide performance bottlenecks. The behavior of traffic becomes much more complicated because of multiple different devices affecting data transmission in ways that are hard to analyze and can interact in subtle ways, sometimes causing serious failures or creating security vulnerabilities.

• Merging of L2 and L3 in a unified network fabric. An example is the use of SDN technologies to reduce the problems inherent with large broadcast domains. Large Layer 2 domains allow for transparent live migration of services from one physical host to another without requiring renumbering. The large number of broadcast messages and risks of broadcast storms however makes this impractical beyond a certain point due to issues of reliability and performance. SDN technologies could provide the means to maintain core connectivity for established virtual networks during such a storm; however, this is a function of configuration and not innate to the technology.

• Reduced management complexity and cost are core value propositions for SDN. Established routes can be established as flows, with less vulnerability to route churn or other influences from external parties. High value or highly sensitive routes can be isolated to the point of invisibility from other network participants.

• Open integration allows operators to rapidly develop and deploy new application services without having to wait for the lengthy process of standardization and adoption.

A4.4 Weaknesses

Increased control granularity and flexibility comes at a cost. Unless the system is designed in a scalable manner, two new vectors of attacks (resource exhaustion) become possible.

- Data-plane resource exhaustion: The total number of actions installed can potentially render the system unusable.
- Control-plane resource exhaustion: In a reactive model, a controller can become unresponsive in case of a high churn.
- Although not related to data-centers, it is worth mentioning the following two aspects that are

not present in the traditional networks.

- Failure and recovery: Inconsistent views due to out-of-band control.
- Software bugs: Now that the network control plane has been moved to software, it becomes theoretically possible for an attacker to take control of a controller.
- Control Plane Injection
- Control Plane DoS
- Cross Application interactions

A4.5 Technical Operations

Multi-Tenant Data Centers have high levels of redundancy and multiple availability zones where failures are expected not to be cross domains. However extremely large numbers of users use these services. When a service outage effects one data center it may affect thousands of applications and millions of users, even for one availability zone. Multi-tenant data centers could be high value targets for a strategic attacker because of the number of services that use them.

Compliance issues for enterprises may also be salient. A financial institution may not only have operational security goals but may also require mechanisms to audit and verify any network security properties that a bank relies on for compliance reasons, for example to enforce a Chinese Wall between retail and investment operations.

Other systems with significant compliance requirements may include medical records systems, and other systems holding sensitive personal information (in terms of EU data protection law and sector-specific US laws such as HIPAA). This may require that systems be separated so that only clinical personnel and approved support staff have access to them. Similar though generally more stringent provisions hold in respect of classified information held not just by government departments, but also defense contractors. Health and defense information may have geographical limitations, in that it may not be stored or transmitted outside a given country of alliance without special protection measures. Such protection properties must be capable of being audited.

As well as separation, transparency is becoming steadily more important. At present, customers of cloud enterprises have little or no visibility into their data storage and transfer. Just as a cloud service provider might face a business demand for a 'Switzerland only' virtual network for a bank or an 'EU only' network for personal health information, so there may be further demands from firms who do not want any corporate information stored in jurisdictions where it might be more vulnerable to subpoena or to governmental coercion.

A5 Large ISP Use Case

A5.1 Introduction

By large ISP we mean an ISP with a physically large network, serving many external customers and providing a broad range of network and hosting services. We have in mind ISPs with national up to global reach, including the so-called 'Tier 1' ISPs. Before considering what a Next Generation ISP might look like, we need a reasonably simple model of a current ISP where we have: a number of geographically separated "Points-of-Presence" (PoPs) connected by the ISP's "Core Network". The Core Network will use a variety of routing protocols—notably OSPF or IS-IS, and iBGP—which may be overlaid over MPLS and/or Layer 2 networks.

  • at any given PoP, there may be:

  – "border" connections to other ISPs: peers and (except at Tier 1) transit providers—either directly

or via Internet Exchange Points (IXPs). These connections all use eBGP.

a local Data Center, where there may be some quantity of:

∗ ISP internal services, including: internal management and monitoring systems (OSS), DNS, email servers, etc.

∗ customer equipment or equipment provided for customer use.

∗ local Content Distribution Network (CDN) equipment: content caching. (Connections to

external CDNs are, essentially, peering connections.)

− customer connections: either simple (Default Route) connections or Transit Customer connections (using eBGP). Transit customers may be other ISPs or multi-homed end-customers.
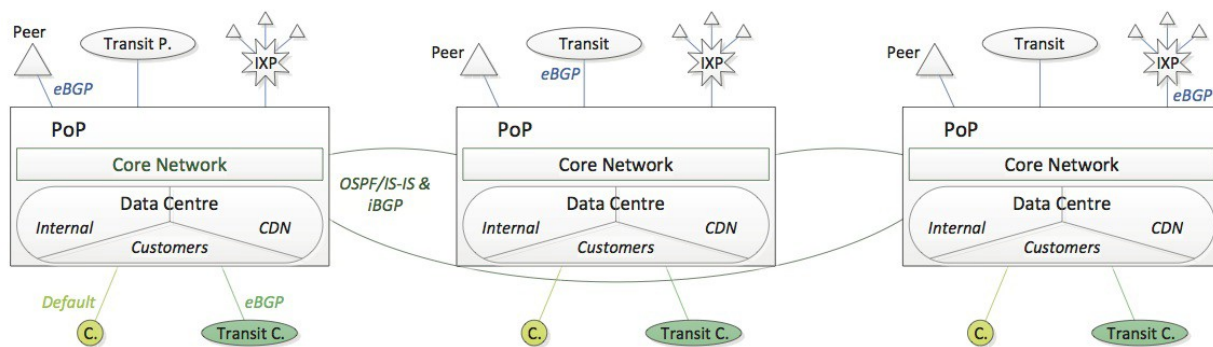

Figure 6 Model of Large ISP

Within the PoP the Site Network will connect things locally and to the Core Network. That Site Network will use a variety of routing protocols and lower layer networks.

Not explicit in this model of an ISP are:

- network management and monitoring: the equipment for these will be distributed across the PoPs and connected by some internal network within a PoP, and across the Core Network between PoPs and to one or more Network Operation Centers. There may be some entirely separate way of reaching some PoPs, for disaster recovery.
- network services: services such as VPNs will overlay the connections and networks shown.
    - Internet Access is also provided over the connections and networks shown.
    - the infrastructure for the PoPs: the buildings, their security, the reliable supply of electricity and cooling, etc.
    - the network infrastructure between PoPs: from the fiber upwards.

The Data Center component of the PoP will vary in size and complexity. For this component, this use case overlaps the Data Center (qv) and the IAirport (qv) use cases. What distinguishes this use case are:

- geography: the ISP's PoPs may be widely geographically spread, so the network between those PoPs may have significant latency and be less reliable than the network within a PoP.
- interconnection with other networks: which is a quite different from connections within a

network.

A5.2 SDN/NGN and Intra-ISP Networks

The working model of an SDN described above has been successfully applied to Data Centers. In a Data Center there may be a very large number of switches/routers and an even larger number of devices in a network with disparate requirements for connecting some devices together, and for ensuring some devices remain separate. Devices in the network may include "middle boxes", such as firewalls, load-balancers, traffic-shapers, and so on. SDN brings software and processing power to bear on all this complexity.

While a large Data Center may be complicated by scale and diversity of connections, it also has properties which fit with the SDN approach:

• it is straightforward to separate the control network from the data plane. It may be possible to physically separate the control network, using separate switches and links between the ME, CE and FE. If not actually separate, the control network may be implemented as separate VLAN(s).

• the control network can be given as much bandwidth as it needs, and is physically relatively small, so throughput and latency to and between CEs should not be an issue.

• the control network is in a controlled and benign environment, so can be expected to be reliable, which all contribute to being able to maintain the required Shared Network State. Note that we do not, here, concern ourselves with exactly how the SDN maintains this state.

Turning to the application of SDN to the ISP, clearly within a PoP the Data Center part can follow the model above. For the ISP's Core Network, perhaps the SDN can be extended across all PoPs, as shown in Figure 7.


Figure 7 Sketch of SDN ISP Core Network

Here it is supposed that the Control Network that allows the ME and CE to coordinate the SDN is extended as an overlay or virtual network over the ISP core network. The issues here are that, unlike within the data center:

- the control network is strongly dependent on the data plane.
- the bandwidth available to the control network may be limited, and the network is physically large, so throughput and latency to and between CEs may be an issue.
- the control network is in an uncontrolled environment, so cannot be relied upon

Approved for Public Release; Distribution Unlimited.

On the other hand, it is only necessary for the extended shared network state to cover the core network, which will generally be relatively simple.

SDN for this application requires mechanisms which replace the routing protocols which currently make forwarding decisions based on the dynamic state of the network and distribute that state around the network. The Shared Network State abstraction allows for new and better ways to manage and make those forwarding decisions, but depends on being able to keep track of changes in the network in a timely fashion and being able to maintain stability—both issues which current routing protocols struggle with.

At present one can only speculate whether some form of SDN, along these lines or otherwise, will replace today's routed core networks. However, if so, we believe that the SDN will comprise a network of ME, CE, and FE devices—what the ME and the CE will do, exactly, remains an open question.

A5.3 Transitional or Hybrid ISP Networks

Assuming that ISP networks move to some form of SDN over time, there will need to be a means to incrementally replace existing routed networks. So, whatever form the SDN takes, it will need to interoperate with current routed networks. The essence of this is the separation of routing from forwarding—so that SDN parts of an ISP network speak routing protocols in order to exchange routing information with existing parts of the network, but make forwarding decisions and manage the data plane in their own way.
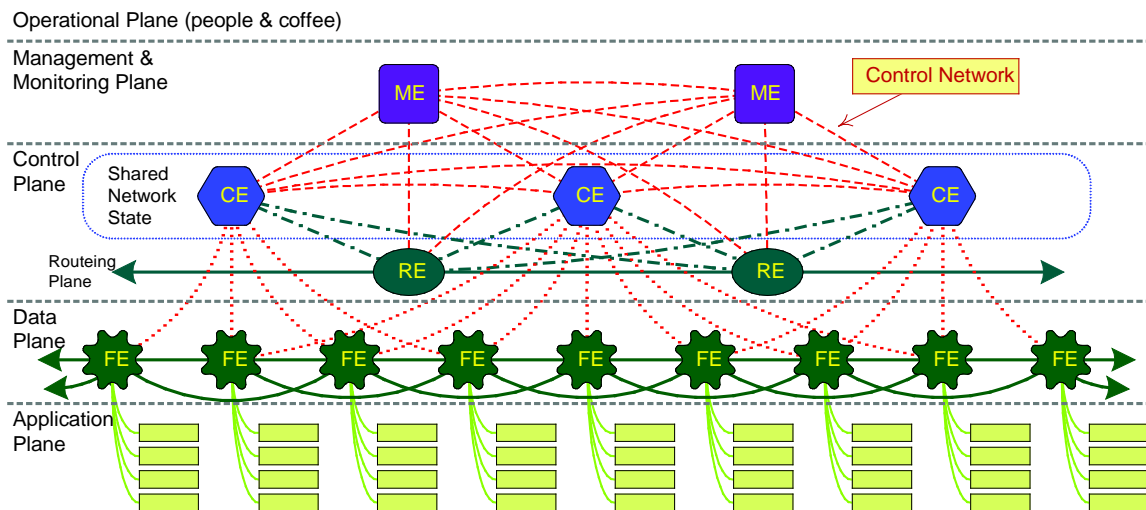

Figure 8: Transitional SDN ISP Network

Here the Routing Elements (RE) speak routing protocols to each other and to routers in other parts of the network. The RE are under the management of the ME, in the same way as the CE. The RE exchange information about the network state with the CE. Note that the RE are entirely separate from the Data Plane, except to the extent that their connections to each other and, especially, to existing routers may be implemented by the data plane.

This, essentially, completes the separation of routing from forwarding that exists in current integrated routers. It also separates not only the forwarding, but also the forwarding decision making from the distribution of routing information. The business of routing is divided in three parts:

1. the routing protocol, which specifies the information which is distributed across the network, and how that is achieved. This dictates what the router stores in its Routing-Information-Base (RIB).

2. the configuration of routing policies, such that the routers which comprise a network collectively deliver what the network operator wants.

3. deciding how to forward packets given the contents of the RIB and the routing policy—that is to say, deciding what to store in the Forwarding Information Base (FIB).

Separating routing from routers makes it possible to then decompose the business of routing, so that:

- the configuration of policy can be centralized, so that the operator can configure their network, not a collection of routers.
- more software and greater computing resources can be applied to making forwarding decisions: to improve traffic engineering, pre-calculate fail-over paths, improve network utilization, and so on.

This transitional organization applies equally to intra-ISP and inter-ISP routing. For inter-ISP routing the RE would speak eBGP to peers and transit providers. The RE would speak iBGP to existing routers and/or route reflectors. The opening up of the business of routing opens up the possibility of replacing iBGP with something less prone to spending tens of seconds or more "hunting" for new paths when things change.

There are a number of incentives for ISPs to move in this direction:

- improvements in network management and configuration—reducing cost and providing greater control and flexibility.
- with a centralized view of the network and central management of its configuration, it should be possible to model network behavior, and check configuration changes before they are applied to the network.
- opening the market for new suppliers of separate forwarding devices.
- innovation in network management, routing, and control software.

A5.4 SDN/NGN and Inter-ISP Networks

Thus far we have considered only the ISP's own networks, within and between its PoPs. Now we consider how inter-ISP—inter-AS, peering and transit connections—might change in an SDN/NGN world.

Currently, the essence of Inter-ISP networking is BGP. BGP is a simple protocol, carrying a relatively small amount of information about a relatively large number of routes. What makes BGP complicated is firstly the scale of the task, distributing routes for and across the entire Internet, and secondly all the policy bells and whistles intended to allow the ISP to manage their connections to the rest of the Internet.

The wonder of BGP is not so much that it is far from perfect, but that it works at all. Amongst the issues with BGP are:

1. the speed at which the BGP mesh can respond to changes is deliberately damped in order to maintain stability. This can lead to some routing losses measured in (small numbers of) minutes.

2. for eBGP, particularly, there is a strong, implicit link between the BGP session and the routes advertised in that session—the control and the data planes are bound together. This is because most eBGP sessions run over a point to point connection between a router in one AS and its correspondent in another. The routes advertised in the session will naturally have a next-hop which uses the same point to point connection. For resilience two ASes may establish two separate interconnections, with two separate eBGP sessions. The failure of one connection causes a ripple at the BGP level, where it would be preferable to manage the recovery at a lower, faster level.

3. the strictly limited support for traffic engineering, particularly beyond the network's borders. This is partly to do with limitations in the protocol, but a lot to do with the independence of every AS and the implicit "best-efforts only" nature of the wider Internet.

4. BGP carries only information about reachability. It carries no information about capacity.

5. the absence of any means to verify that a route arriving via BGP is kosher—saving the presence of BGPSEC.

6. the absence of any means to detect "route leaks", BGPSEC notwithstanding.

None of these issues are directly related to the SDN/NGN separation of control plane from data plane. Mostly these are deep issues with the structure of the Internet. And the Internet is of a size that any change will take time. So, it is hard to imagine that there will be a swift resolution.

However, the separation of routing from routers offers the best opportunity yet to extend or replace BGP, starting, perhaps with iBGP.

A5.5 Operators and Stakeholders

In general terms we expect an NGN ISP to have a network comprising a Data Plane, made up of a number of Forwarding Elements, under the control of a Control Plane made up of a number of Control Elements (which directly control Forwarding Elements) and a number of Management Elements (which manage the network) and (possibly) a number of Routing Elements (connecting to existing routed networks). It seems likely that the Forwarding Elements will be distinct devices, while the functions of the Control Plane may be combined and implemented as integrated or separate software systems and applications, spread across some number of actual devices.

Without attempting, at this early stage in the development of SDN/NGN ISPs, to define how each of the elements will, eventually, work, we can identify a general structure comprising:

• control devices: spread across the ISP's PoPs

• some network connecting the control devices—the "control network"

• forwarding devices

• some network connecting forwarding devices to their controllers—the "command network"

It seems likely that each forwarding device will be under the control of a small number of controllers (for resilience) and that those controllers will be local (within the PoP).

A "pure" forwarding device might be defined as one which does exactly as it is told by a single (or replicated) controller. A device which integrates different sets of commands from different controllers may be deemed to be a hybrid (low level) controller and forwarding device. A hybrid device would be (in our terms) connected to the "control network".

The NGN ISP may be expected to have a core network and various other networks interconnected

by that core network—as now—implemented by some hierarchy of control and forwarding layers.

So, to secure the control of the NGN ISP it is necessary to:

• secure the various elements: so that each one cannot be subverted or prevented from doing their intended job.

• authenticate connections in the control and command networks: so that each element only talks and listens to the elements it should talk and listen to.

• protect the connections in the control and command networks: so that data is not lost, modified, added or delayed.

The data exchanged is not strictly secret, but it would do no harm to:

• encrypt connections in the control and command networks.

To ensure that the control of the NGN ISP is reliable requires redundancy of elements and network. Implementing the control and command networks as physically separate networks has obvious advantages. Implementing them as separate virtual networks, with some priority to ensure the availability of bandwidth, is the next best thing. Where the control network extends between ISP PoPs, a virtual network layer is the best option.

Assuming that control of the NGN ISP is insulated from outside interference, we may worry about whether inside interference can be detected and dealt with—for example, misconfiguration arising from human error or from malice. Some consistency checking in the control plane is required to check that what the network is being told to do is valid and correct. By valid we mean that it is consistent with the network topology and capabilities. By correct we mean that it is consistent with what the operator wants the network to do. So, it is valid to tell the network to do something it can do, but it is only correct to do so if the network then does something the operator wants.

Checking for validity and correctness requires a specification of the network topology and policies which can be checked against. One of the advantages of a NGN ISP network is that the control plane may be driven by just such high level specifications. Checking lower level configuration against the higher level specification could detect errors in the generation of that lower level configuration. Closing the loop and checking the actual behavior of the network against the high level specifications may detect errors at any point in the process of telling the network what to do, and could detect interference which all other measures has failed to detect.

There are a number of parties involved:

1. the network operator, including:

    (a) the operator's own NOC.

    (b) subcontractors.

2. the network's customers, who may have:

    (a) virtual private networks.

    (b) virtual private data centers over which they may wish to have (at least virtual) control.

3. the network's peers and transit providers: where eBGP (or some future replacement) provides an arm's length connection between the networks, but which carries information which the ISP control plane must be able to depend on for its external routes. With BGPSEC there are other parties involved,

providing the data-bases which contain keys used to authenticate BGP messages and which attest to an AS's right to originate a set of prefixes.

For the large ISP its subcontractors may be a particular concern. When equipment is replaced or newly installed, it must be attached to the control network, started up and brought under the control of the ISP's NOC. If a third, fourth, fifth… party is doing the work on site in some remote PoP then the ISP's NOC needs some reliable way of remotely installing the keys required for the new device to authenticate itself to others, and of configuring it so that it will not disrupt the control network the moment it is connected. In a remote PoP the ISP may wish to consider the possibility of extra equipment being placed in their network, and of extra software and configuration being added to devices as they are installed.

For virtual private networks and data centers, the question is how deep into the control network the customer needs to be admitted. One approach may be to allow the customer access to their own high level specification of their virtual infrastructure, and nothing more. The ISP's management systems can then translate and check that specification, before allowing it to be reflected in the actual network. More complicated is the creation of vertical or horizontal slices of the ISP's infrastructure, and allowing the customer to reach in and manage those—clearly this requires some means to ensure that the customer can affect only their slice, and that what they do with that slice does not exceed capacity or other limits on the service provided.

A6 IXP Use Case

An Internet Exchange Point (IXP) is, essentially, a switch—that is, Layer 2 infrastructure to which many ISPs can connect, and over which those ISPs can establish Peering (or, possibly, transit) connections.

The larger IXPs operate 1+1 redundant switching infrastructure, in some cases supporting a virtual Layer 2 mesh so that failures of the underlying infrastructure are invisible. For large flows and connections, some IXPs provide direct Layer 1 connections between ISPs and some use Layer 1 switches to switch connections between redundant devices.

An IXP is fairly straightforward and should change only as connections are added or removed. For the largest IXPs the challenges are traffic volumes and reliability requirements. Nevertheless, centralizing the configuration and control of the switching layer could make the job of running an IXP easier. The virtualization of the switch infrastructure may be achieved more easily and less expensively with a new (SDN) control plane over an (OpenFlow) data plane.

At an IXP it is up to the individual ISPs to establish and maintain eBGP sessions, to exchange routes and traffic. The IXP itself is not involved in this process, though the IXP must provide adequate capacity—when the ISPs ask for it—and it is in the IXP's reputational and commercial interest to ensure that all its clients maintain adequate capacity.

However, most IXPs also provide a Route Server. To peer with others at the IXP, an ISP can connect to the Route Server, so that a single eBGP session between the ISP and the Route Server replaces many connections between the ISP and all the other ISPs at the exchange. This replaces a full mesh of individual eBGP sessions by a hub and spoke arrangement, where the Route Server is, effectively, a proxy for all its clients. The principal advantage of the Route Server is that a new ISP joining the exchange does not have to persuade all existing ISPs at the exchange to establish a new eBGP session— where the marginal cost to each of the existing ISPs can mean that the task lingers at the bottom of the list.

Many ISPs connect to the Route Server and will peer with any and all other ISPs. Some ISPs connect

to the Route Server but wish to pick and choose their peering partners. So, the Route Server must provide a "peering matrix" function, which provides, at a minimum, each ISP with ability to deny or permit announcements to other ISPs. Because the Route Server exchanges routes using BGP, where it has more than one route for a given prefix the Route Server must select the best one on behalf of each client. So, each client is delegating some policy to the Route Server—though current Route Servers do not allow the client much, if any, control over the route selection made on its behalf. This is an area where opening up the eBGP routing layer may allow the Route Server and its clients to cooperate more closely, or to do away with the Route Server altogether (either by automating the process of setting up a new peering connection, or by providing a form of broadcast for announcements).

Some IXPs configure their Route Servers to filter incoming route announcements (that is, announcements coming from each client) to allow only prefixes which are known to be valid. This filtering may be configured from routing policy published in an Regional Internet Registry (RIR). It is not, strictly speaking, the IXP's business what routes its clients announce to each other. However, some feel they should ensure that only valid routes are announced by their Route Servers.

In a "Software Defined IXP" we may see a full integration from the ISP policy down to the packet forwarding, so that:

- routes announced (broadcast) by the Route Server are checked for validity against the originating clients' policies.
- packets forwarded are checked for validity against the routes announced.

The effect of which is to ensure that nothing crosses the exchange which should not cross it.

The control plane of the IXP is entirely a matter for the IXP. So, this might be secured simply by ensuring that no third party has access to the control and command network(s). But the control plane would be more secure if it were secured as if third parties could access it. There is not much which is novel here.

The use of Route Servers at exchanges is an interesting example of what may be achieved by the SDN approach. First, providing a better and less expensive way to build and manage an IXP. Second, by allowing the control plane to be extended—once it is unbundled from the data plane, and given fine grained control over the forwarding plane—to improve the working of the network: in this case by ensuring the IXP only carries the routes and traffic it is intended to carry.

Further, given the ability of both the IXP and the clients to extend what the Route Server does and how clients interact with it, it would be possible for each client to apply its own policy to the routes available at the IXP. This implies some mechanism for each client to influence "its part" of the Route Server, in much the same way as a Virtual Private Network client needs to be able to manage "their part" of the host ISP's network.

LIST OF SYMBOLS, ABBREVIATION AND ACRONYMS

API: Application Program Interface

AS: Autonomous System

BAA: British Airport Authority

BCP38: Best Current Practices (BCP)38, another term for RFC 2827, is a best practice methodology around ingress traffic filtering

BGP: Border Gateway Protocol

BGPSEC: Border Gateway Protocol Security Extensions

CCTV: Closed Circuit Television

CDN: Content Distribution Network

CE: Control Element

COTS: Commercial off-the-shelf

DNS: Domain Name System

DNS: Domain Name Server

DNSSEC: Domain Name System Security Extensions

DoD: Department of Defense

DDoS: Distributed Denial of Service

DoS: Denial of Service

DSL: Domain Specific Language

eBGP: External Border Gateway Protocol

FAA: Federal Aviation Administration

FE: Forwarding Element

FERC: Federal Energy Regulatory Commission

FLIB: Flow Information Base

HIPAA: Health Insurance Portability and Accountability Act of 1996

iBGP: Internal Border Gateway Protocol

IP: Internet Protocol

IPC: Inter-process Communication

IPv4: Internet Protocol version 4

IPv6: Internet Protocol version 6

ISP: Internet Service Provider

IS-IS: Intermediate System to Intermediate System

IUPUI: Indiana University – Purdue University Indianapolis

IXP: Internet Exchange Point

JFK: John F Kennedy International Airport

LAN: Local Area Network

LAX: Los Angeles International Airport

ME: Management Element

MPLS: Multiprotocol Label Switching

NAT: Network Address Translation

NERC: Nuclear Energy Regulatory Commission

NGN: Next Generation Network

NOC: Network Operation Center

OSPF: Open Shortest Path First

OSS: Operational Support System

PoP: Point of Presence

POX: An open source development platform for Python-based Software Defined Networking (SDN) control applications, such as OpenFlow

RE: Routing Element

RFID: Radio Frequency Identification

RIB: Route Information Base

RIR: Regional Internet Registry

SCADA: Supervisory Control and Data Acquisition

SDN: Software Defined Networks

SIL: Safety Integrity Levels

TCP: Transmission Control Protocol

TFP: Topological Flux Balance

TLS: Transport Layer Security

VPM: Virtual Private Network

VM: Virtual Machine

WAN: Wide Area Network

WWW: World Wide Web

X.509: Public Key Infrastructure Certificate Specification

X590: Public Key Infrastructure Certificate

XML: Extensible Markup Language