SAND2014-3230C

Routing is a Risky Business

Jeff Boote
Sandia National Labs
Livermore, CA
jwboote@sandia.gov

Carrie Gates
Dell Research
Round Rock, TX
carrie_gates@dell.com

Tom Kroeger
Sandia National Labs
Livermore, CA
tkroeger@sandia.gov

William Stout
Sandia National Labs
Albuquerque, NM
wmstout@sandia.gov

ABSTRACT

The Internet's routing infrastructure is based on BGP for route selection. This protocol is known to be insecure, which has been addressed by protocols such as S-BGP, which have not seen wide adoption due to the underlying economic model. One of the more difficult security issues is that network traffic can be hijacked without the end points realizing that their traffic has been rerouted through a malicious node — and such activities have been observed. In this paper we present a risk-based overlay network that allows end nodes to estimate the risk involved in using a particular first-hop node and therefore likely route. The new paradigm lies in adding security to performance calculations at the edge nodes of networks in a manner that takes into consideration the incomplete information and lack of true route control.

1. INTRODUCTION

The Internet is an essential underpinning for our nation's critical infrastructure. It is a primary communication channel for industry, utilities and government at all levels. Security for this distributedly owned (and mostly private) infrastructure is challenging. At the lowest levels of the network infrastructure the design would seem to sacrifice confidentiality and integrity to achieve very high levels of availability. Higher levels of the infrastructure attempt to add in confidentiality and integrity through the use of cryptography, but fundamentals such as attribution are challenging without support at the lower infrastructure levels. ISPs consider risk primarily in terms of connectivity and redundancy, while applications may have a more rich set of concerns such as not having others see their proprietary information.

A prime example of this trade-off can be seen at the routing layer of networking. Network administrators are concerned with maintaining the security of the infrastructure itself and of ensuring traffic is not interrupted. But from a transit perspective, administrators have very little involvement in preserving the confidentiality or integrity of the data transiting their network.

The Border Gateway Protocol (BGP) is the foundation for connectivity in the Internet. BGP defines how autonomous system (AS) networks are interconnected; however, BGP as a system is limited because it primarily depends on a transited trust arrangement where each peer network shares connectivity information with all direct peers. Specifically, BGP has the following limitations:

- Low resilience against malicious actors; route interception is becoming common.
- No notion of national borders.
- No mechanism for selected-path verification.
- Routing decisions are hidden from end applications, making risk management challenging for end users.

Protocols such as S-BGP secure the communication between connected peers but do not mitigate the risk from malicious actors within the community. Additionally, the costs associated with implementing S-BGP are not borne by the same entities that benefit the most from it, making deployment problematic.

The notion of using risk to determine if a transaction should be allowed or denied has been gaining traction, particularly within industry settings. A simple example of this is how credit card companies attempt to prevent credit card fraud — they calculate the risk of allowing a particular purchase to be made based on a number of factors, including geographic location, purchase history and known patterns for malicious actors. Similar approaches are being employed by Facebook and Google to determine if a user should be allowed to log into the account, even if the password is known. As such, risk measures are becoming part of access control decisions. And yet, the notion of risk has not been realistically applied to network control (at least to the best knowledge of the authors).

This paper presents a new paradigm of traffic routing that includes risk based on overall *security* needs (rather than

simply availability) in route determine. It examines the risks of interception by malicious actors and provides risk measures to aid edge networks in mitigating those risks.

The rest of this paper is organized as follows: we describe motivation, background and challenges in Sections 2, 3 and 4. We then go on to describe the proposed risk system in Section 5, and how this overlay might be implemented in Section 6. The approach is validated through three real-world use cases in Section 7, and an approach to testing this overlay (without first requiring wide-scale deployment!) is presented in Section 8. We then provide a description of requirements and suggestions to promote adoption, along with the limitations of this approach. In Section 11 we present comparisons to related work. We conclude in Section 12.

2. MOTIVATION

The Internet is an essential underpinning for our nation's critical infrastructure. It is a primary communication channel for industry, utilities and government at all levels.

At the lowest levels of the network infrastructure the design would seem to sacrifice confidentiality and integrity to achieve very high levels of availability. Higher levels of the infrastructure attempt to add in confidentiality and integrity through the use of cryptography, but fundamentals such as attribution are challenging without support at the lower infrastructure levels.

Recent revelations about the NSA and OpenSSL show that depending solely on a cryptographic solution is a brittle approach that can quickly loose all protections [9].

Cryptographic limitations and concerns of internet surveillance at U.S. based ISPs has countries such as Germany investigating modifications to the existing internet routing model. [14] The publicized modification would require data sourced in Germany, and destined for Germany, to never leave Germany. This kind of fine grained physical geography control is not currently available in the dominate inter-AS protocols used to control routing on the Internet such as the Border Gateway Protocol (BGP).

While the larger news story for the past year has been about surveillance at large telcos, the instances of Internet route hijacking has also increased [4]. In these kinds of attacks, an adversary will attempt to subvert the routing protocols to influence traffic to flow through an AS where they have a presence.

This paper presents a new paradigm of traffic routing that mitigates the risks of interception by malicious actors that are accepted members of the community. We apply risk to networks, and develop a risk-based overlay on-top of standard BGP routing. In particular, we propose the following mitigations:

- Use of secondary control-plane communication channels applied to topological representations of the full routing table
- The use of geo-location for networks to help determine if a route is considered risky
- Real-time performance feedback as compared to expected performance.

The end goal for this work is the design of a risk scoring system that allows an organization to configure their routing policies such that routing decisions can be partially based on the risk perceived in the route being suggested, in addition to shortest path, least cost and other traditional policy metrics. While we will focus on risks of interception and modification of packets, these mitigation techniques may have potential use for other security concerns such as denial of services.

3. BACKGROUND

The routing mechanisms we will discuss are primarily related to BGP as the most dominant inter-AS routing protocol. BGP at its core is a policy negotiation model, where each AS notifies its neighbors about the networks it can reach based on the local networks it controls and the announcements of its neighbors.

3.1 Secure BGP

Secure BGP (S-BGP, BGPSEC, soBGP, etc...) adds mitigations that if employed would make routing more secure. Using these technologies would allow an AS to determine if a given route originated from an AS with authority to announce a given network IP prefix but only if all ASs on the path are using it. Partial deployments may even make security worse for some.[12] Additionally, Secure BGP does nothing to expose any measure of risk to peers or end users.

3.2 ALTO

Application Layer Traffic Optimization (ALTO) is an architecture and protocol that was developed to provide P2P, overlay networks, and other application layer entities the ability to make policy related requests of the network infrastructure in ways that preserve the ability of the infrastructure to hide topology. [15] The kinds of questions the overlay networks are asking of the underlying network infrastructure are very similar to the ones needed to ascertain risk of interception. However, ALTO explicitly hides the topological information making geo-referenced risk measures infeasible with ALTO alone.

4. CHALLENGES

Security enhancements at the Internet routing level have been very difficult to advance. Each AS is interdependent of the other ASs resulting in a tragedy of the commons situation. The only security enhancements that are likely to succeed are ones that adhere to selfish motivations. The network operators expending the resources need to see benefits from the results without depending on neighbors to do the same.

Additionally, network operators are traditionally very secretive about topology and specifically peering agreements. This is an area of competitive advantage. One might question the ability for this to be a free market given the lack of transparency.

4.1 Solution Space

These considerations mean that any successful solution will need to provide the appropriate incentives to work. Specifically, the networks that most benefit from risk scoring of paths are the edge networks, the ones that provide network access for data providers or data consumers. This is because they are the ones with real users as customers — the customers that want their data kept confidential. And, this will only happen if the edge networks can make a business case from the effort. The work must cost equal or less than they can charge customers for it.

The transit networks will not initially see benefit from this until their customers (the edge networks) request it of them.

This leads to the follow-on requirement that the solution must provide benefit even with a very small deployment footprint. Ideally, a single AS should benefit by deploying the solution, although it is likely the benefit will increase as the deployment footprint increases.

5. A RISK OVERLAY

The first step to creating a risk based routing paradigm is to create a risk overlay of the network. The risk overlay will be built upon a topological foundation. The topological foundation will provide a logical graph representation of the global network of ASs.

Regional Internet Registries (RIRs) and the Routing Assets Database (RADB) create available databases of existing inter-AS path information [5]. This data, along with ALTO services and available BGP Looking Glass servers, can be used to generate this topological representation of the global network [10]. The topology representation does not have to be perfect, but it should be possible to assign confidence levels to different aspects of the topology. These confidence levels can eventually be used to help quantify the risk associated with using that topology. For example, if a given peering is shown to exist through all the exploratory methods described, it is more likely to be true than if only seen by one.

We propose utilizing a graph database to host this topological representation of the full global internet. This database would be queried by BGP instances to allow them to compare the BGP UPDATE messages from their peers with the derived graph topology database. Paths with a high degree of correlation would have a more advantageous risk score. This functionality allows an AS to evaluate a path announcement based on the full path rather than simply the neighbor, the originating AS, and the path length as is typically done today.

The actual implementation of this graph database will depend upon scaling issues. For example, it is possible that the RIRs could build and provide a graph based database for the use of ASs in their region, or this could be provided by third party vendors or, if implementations can be made efficient enough, individual ASs could deploy their own.

The next step to enhance the effectiveness of the topology graph is to augment the AS information with geo-references. The LAT/LON and political boundary information could potentially be added in by AS administrators themselves, or heuristically determined. Sterbenz *et al.* [17] goes into some detail as to why getting physical topology is important for risk, and also why it is do difficult to get. But again, we do not expect to get perfect data and will simply allow better data to eventually provide more advantageous risk scores in the hope that ISPs eventually see this as a way to attract more traffic and therefore more revenue. These geo-references would allow individual paths to be compared not only based on how well the peer announcements match, but also based on the risk associated with the location of the physical path.

The final aspect of our risk overlay system is to incorporate real-time measurements of network traffic to validate the expectations of the topology graph. We propose using simple sampling techniques to grab small subsets of traffic transiting the network at many locations, and using the TTL and the TCP timestamp of packets to determine hops and latencies from senders of that traffic. IP geo-location and latency modeling of the internet (network tomography) can be used to infer latencies between router nodes and therefore provide baseline expectations for given AS paths [13]. This information can be populated into the topological graph database allowing end ASs to compare the expected number of hops and packet latencies with the actual values seen by transiting traffic. The more these values differ, the higher risk score would be associated. However, it is possible we will need to explore dampening techniques for this option to deal with the long-tail latency distributions seen on congested connections.

6. USING THE RISK OVERLAY

Now that we have built a Risk Overlay utilizing published RIR data and observational network information, we can employ this additional information to make more informed routing decisions and allow end-clients to participate and manage their own risk.

6.1 Risk Factors

We employ three specific risk factors that are enabled by this risk overlay. Each risk factor consists of a comparison between observed and expected values.

x_1**: Peer announced path correlates with RIR deduced graph:** In this case the node performing the risk calculation deduces a routing graph from the topology database (as described in Section 5). It then compares this graph to the routing information being supplied by its peers. That is, assume that a given node has three peers. It compares the information supplied by each of these peers against the graph that it has already deduced in order to determine if the peers match (or mostly match), or if there is a significant discrepancy. The actual value is calculated as one minus a percentage of match between the deduced database and the peer announcement. This provides redundancy for peer reported messages and makes it more difficult to subvert regions of the internet, effectively shrinking the attack surface to the RIRs. Thus the risk increases as $x_1 \rightarrow 1$.

x_2**: Selected path traverses (or not) specific geopolitical regions:** While the actual path between two nodes can not be specified, nor even known with certainty, the ex-

pected path can be inferred from the geo-location references in the topology database (as described in Section 5). This information can be used to determine if the nodes in the expected path (starting with a given first hop) are within a specific political or geographical boundary. In the simple case, the value for x_2 is simply one minus the percentage of nodes that are within the desired boundary. In the more complex case, an organization might want to specify that certain boundaries are "riskier" than others, or even have some locations be considered as completely inappropriate (in which case the value for $x_2 = 1$). Thus the risk increases as $x_2 \to 1$.

x_3: Network tomography correlates with expected path: An additional metric for determining if the path being followed is the expected path is to use actual data packet information to determine if we are seeing the delays we would expect for a given route. More specifically, routers can sample traffic and determine if TCP timestamps (correlated between send/recv pairs) match expected delays (s). The expected delays (e) to different ASs can be calculated *a priori* based on their estimated physical location and the speed of light. Given that the exact physical location and length of fiber will likely not be known exactly, "fuzziness", or error bounds, can be added to the expected value to account for uncertainty. Whenever the time to transit an AS exceeds this value, the actual route being used comes into question. The greater the deviation, the more suspect the route. More specifically, if $s < e$ then $x_3 = 0$ else $x_3 = \frac{s-e}{V}$ where V is the largest possible delay that is suspect but not known to be definitively an unacceptable route. If $x_3 > 1$ then it is rounded down to 1. Thus x_3 represents the risk that the route being taken is not the expected route based on expected route timings. If no traffic has yet been seen to a given route, $x_3 = 0.5$ to indicate neutral risk. As traffic is seen, a cache table will be populated with both expected and experienced delay times making this risk factor easy to calculate each time a BGP update takes place.

It should be noted that the internet does not generally allow a sender to explicitly specify the route of the traffic. This is true of every AS along a path. Therefore, once a packet has left an AS the eventual route may be very different from the expected route. This distributed control aspect of the internet is exactly what makes it effective for high-availability functionality; however, it does this at the expense of confidentiality and integrity. These two final risk factors help us mitigate this concern.

6.2 BGP Decisions

We have identified three specific risk factors to incorporate into a risk overlay, but note that the design is such that additional factors can be added, and that individual ASs will have the ability to determine how each of these factors effect their routing decisions.

We develop metrics to measure the effectiveness of the modified protocols from both a security and performance perspective, applied to the global system and each individual AS. If the value for any of $x_1 = 0$, $x_2 = 0$, or $x_3 = 0$, then $R = 0$. Likewise, if $\omega_1 + \omega_2 + \omega_3 = 0$, then $R = 0$ (and Risk measures are disabled). Otherwise the following equation, which defines an overall risk score, is used:

$$R = \frac{\omega_1 x_1 + \omega_2 x_2 + \omega_3 x_3}{\omega_1 + \omega_2 + \omega_3} \quad (1)$$

where ω represents a weight on each of the different values x (described in Section 6.1 above). These weights (ideally, but not necessarily, in the range [0..1]) are decided based on site policies that allow each site to determine their own priorities on each of the measured values. If a factor is not important to a site, then its weight can be set to 0. If each factor is equally important, than each weight can be set to 1. By dividing by the sum of the weights, we allow a site to set a weight (and hence a factor) to zero for the overall risk score (indicating that they don't care about that particular risk), and judge the risk entirely based on the remaining factors. The resulting risk value is in the range [0..1]. We have chosen a simple linear function rather than a more complex calculation in order to support the speed at which routers need to operate.

A site can determine its tolerance to risk, and set an appropriate threshold based on that tolerance. In order to determine an appropriate tolerance, we recommend calculating risk values for every connection for some given period of time (e.g., a day), and examining the values obtained before deciding on an appropriate threshold. A site might also use simulations (see Section 8) based on this actual collected data before deciding on a threshold.

In time, this calculation can be modified to take into account the history of a given route, so that a route that has consistently provided a low risk score is given preference over other routes. Conversely, should a high(er), or anomalous, risk score be encountered for this route, the anomalousness of the score might cause the risk for this route to be elevated even further, such that the route is not selected even if the risk score is with acceptable limits. While we note this as a possible algorithm enhancement, we do not address this idea any further in this paper.

6.3 Network Deployment

One way to incorporate this information into existing deployments is to use each of these factors to provide a combined risk score that can be used within the LOCAL_PREF mechanism of BGP. In general, the path with the highest LOCAL_PREF score will be selected. (Other tie-breaking methods come into play as well.) LOCAL_PREF is a good choice for risk score inclusion because that is the part of BGP expected to indicate the cost/benefit to the local AS and risk management is part of that measure. This is in contrast to the boolean choices available with route filtering where an AS can simply accept or not accept a route based on security controls. The following equation describes how a risk-inclusive LOCAL_PREF would be computed:

$$LOCAL\_PREF' = LOCAL\_PREF * (1.0 - \omega * R) \quad (2)$$

Where ω represents the overall weight a local AS wants to afford risk measures relative to their baseline LOCAL_PREF and $LOCAL\_PREF'$ would be used in place of the statically defined LOCAL_PREF within the BGP route selection for

inclusion in the Forwarding Information Base (FIB) of the routers.

6.3.1 Too Risky To Route

The above methodology allows an AS to prioritize a less risky route over more risky routes. But, what if a given route is deemed too risky to use, and it is the only route available? This question is anathema to network operators because dropped traffic is dropped business (unless this is the policy the customer wants).

However, a policy could be implemented to deal with traffic that matches routes with a $LOCAL\_PREF'$ below a specified threshold with two rules.

1. If there is a less specific prefix that matches the given network, then remove the route from the table and use the less specific network route.
2. Otherwise, replace the route with a special null route. Traffic that matches this route should be dropped and the router should notify the sender that it was unable to forward the traffic by sending an "ICMP Administratively Prohibited" message (Type 3, Code 9) [3].

Additionally, if a route is determined to be too risky by an AS, it should not be shared with its peers. This will ensure traffic destined to the remote network that is not sourced by the current AS is not sent to it, and it will eventually not have to drop further transit traffic. This propagates the administrative decision closer to the sender of the traffic.

7. USE CASES

In order to explain how this risk overlay would work if deployed, we examine three use cases based on real world events.

7.1 YouTube Blocked by Pakistan

In 2008, Pakistan Telecom blocked access to YouTube based on an order from the Pakistani government. In doing so, they started advertising a route that was a subset of YouTube (a /24 rather than the full /22). This advertisement propagated across the internet, with most traffic to YouTube selecting this new route (to Pakistan) due to the more specific routing advertisement. The net result was that the majority of the internet was unable to access YouTube for nearly two hours [6].

In this case, our risk overlay would detect the change by noticing that the peer announced path did not correlate with the RIR deduced graph, and so the value for x_1 would approach 1, thus driving the risk calculation to approach 1. If this change was not detected by x_1, it might be detected by x_2, where it would be noted that the selected path potentially traverses a non-desirable geo-political location. Finally, the lack of ACKs would drive the observed time to route to infinity, thus driving x_3 to 1.

7.2 German Traffic Only Please

In the wake of the Snowden revelations, particularly of accusations of the US spying on German Chancellor Angela Merkel, Deutsche Telekom (a German communications company) has announced that it intends to develop a routing system that will keep German traffic within Germany. The goal is to prevent any traffic that originates in Germany, and that is destined for a location within Germany, from crossing any international border. The broader debate calls for a European Union level initiative that would keep EU traffic within the EU [1].

We take such desires into consideration with the second factor of our risk metric (x_2). In this case, we use the RIR-deduced graph to infer the path that will be selected, and determine what routes that path traverses, using this to determine ultimately if we accept the risk of our traffic transiting certain countries. Thus we have defined a risk metric that is sensitive to national boundaries.

7.3 Targeted Misdirection

Renesys noted in 2013 that they had observed several man-in-the-middle (MITM) attacks against approximately 1500 IP blocks, with each attack lasting from minutes to days [8]. Unlike classic MITM attacks, which required getting physical access to the routing infrastructure, these attacks were carried out via route hijacking using the BGP protocol. Specifically, the attacks caused network traffic to be redirected to a given point (where, presumably, the traffic was captured) and then routed to its final destination. For example, in February 2013, traffic from several different countries was routed through Belarus (e.g., [8] gives an example of traffic from Guadalajara, Mexico, to Washington, DC, that routes through Belarus due to the BGP peering announcements). They also cite a second example where traffic from one site in Denver to another site *in Denver* was routed via Iceland. Given that the delays from this are minimal (from an end user perspective), and that the traffic actually reaches its destination, most organizations will never notice such an attack. Renesys notes also in its report that it has been unable to determine any attribution, or even if these were actual malicious attacks or software errors.

Were our risk overlay deployed by an end node (e.g., Guadalajara), then the first few packets or connections would likely be delivered via the hijacked route. However, the network tomography (x_3) would soon indicate that the TCP timestamps being observed on the return traffic (e.g., any ACKs received) would be significantly different from the expected values, forcing $x_3 = 1$ and thus $R = 1$, causing the traffic to either be rerouted or dropped with ICMP messages sent to the originator.

8. TESTING APPROACH

Thus far, we have outlined some of the deficiencies inherent in BGP regarding the capability for routers to select paths that are deemed "safe" for select traffic (either transit traffic, or edge-device traffic). Our response to these deficiencies is addressed by the development of a risk-overlay solution that produces a metric which routers may use to influence route selections (using risk factors). The Use Cases described in Section 7 show how our metric could have been used to avert (or aid) events that transpired. This section specifies how we might implement our solution and test it in a controlled environment.

8.1 Route Selection

When a router receives traffic to forward, it employs prefix-matching against its BGP table to determine which port (and path) to forward the traffic. The primary BGP route table is often populated by the shortest routes (default). However, other routes may also be present in the general route table. For our risk overlay system, as described in Section 6, we exploit the LOCAL_PREF attribute in order to force route selection in a router (and the AS itself, for that matter); a route selection is influenced by the following risk factors:

1. Route comparison, based on locally generated topology
2. Avoidance of untrusted ASs
3. Delay/latency analysis

For (1), x_1, a global view of the routed networked is required in order to algorithmically determine shortest, valid path(s) between endpoints. When transmission is required, routes to the destination network are first generated from the network data. This can be done by formation of an adjacency matrix M from the collected data (RIR, Looking Glass, etc) and use of a shortest path algorithm (Dijkstra, $A*$, etc). Shortest, valid routes are then compared with the routes in the BGP table. Those routes whose paths closely match those generated from M will tend to 0 (less risk).

For (2), x_2, a predetermined list of untrusted ASs are collected, based local domain preferences. For those routes selected in (1), additional risks metrics are applied to each based on the ASs that comprise the path. This risk score is then calculated based on a percentage of untrusted ASs in said path (with 0 having no untrusted ASs in the path).

For (3), x_3, an estimated latency value is calculated for the path. This value is dependent on the acquisistion of latency values between ASs and source-destinations pairs; the latencies can be captured passively by monitoring TCP timestamps and TTLs. For a previously traveled path, this value can be stored and used for subsequent communications. For new or aged-out paths, subsets of the selected path previously traveled may be accumulated to deduce the delay for a total path. If this is not possible, then the risk value is set to 0.5 (neutral risk) until latency data can be acquired. Risk is determined based on the comparison between the estimated (e) and received latencies (s), given a threshold V. If $s \leq e$ then the risk is 0. This is a actively applied metric, in that risk should vary as the latency times fluctuate.

Thus, for each route that is received at a router in a given AS (via BGP updates), the LOCAL_PREF attribute is reassigned (Eq 6.3), based on the calcuation of the risk metric R (Eq 6.2). Hence, augmentation of the Quagga BGP source-code shall be drawn from the following algorithms:

```
1) Event: Received BGP Update U
     w': [0,1] (Weight of using risk-overlay)
     w{1,2,3}: [0,1] (Weights on x{1,2,3})

     for each path p in U
       Apply local policy (e.g., set LOCAL_PREF_p for p)
       Calculate x1, x2, x3 for p and set R
       set LOCAL_PREF_p = LOCAL_PREF_p*(1 - w'*R)

2) Event: Received traffic on external interface
     s = source prefix
     Find path p in BGP table T
       where s is destination network
     Retrieve e
     Test e,s against threshold V
     Reassign x3 for p
```

There is no requirement to impose additional algorithms on ingress traffic from the local AS bound for external ASs. The modification of the LOCAL_PREF attribute alleviates this need by influencing route selection as the routes are entered into the BGP table. It should be noted that this BGP risk overlay is not limited to just edge ASs. Intermediate ASs (routers) may establish trusted routes to other intermediate and endpoint ASs through the course of traffic passing.

With the necessary modifications specified for our AS routers, we now address our methodology to test our risk overlay system.

8.2 Testing Methodology

To test our extension to BGP, we intend to deploy a virtualized environment with emulated routers, switches and hosts. The environment will not only provide the platform to emulate the experiments, but also the mechanisms to collect emulytic data for analysis. For the emulated router, we propose to augment the existing open-source BGP daemon Quagga with our risk-based path selection algorithms.

Building representative models of the Internet's routing infrastructure is not as simple. Often, a domain (or AS) in the Internet may have a unique view of the global connectivity mesh due to business relations between ASs (sharing/not sharing routes); routing through some ASs may be likened to black boxes, where actual paths taken may be not made public. Thus, to scope our experiment, we will focus on one or more geographical areas encompassing a sufficiently large number of ASs (using public services such as Looking Glass to build the topology). To maintain some fidelity with the hierarchical nature of the routed Internet, this subset shall contain Tier1/Tier2 ISPs and edge-networks. To further reduce complexity in routing, each AS will be represented by a single Quagga router with multiple interfaces for external AS connctivity.

We will also introduce emulated malicious actors in the environment, to influence the path generation processes in the global routing table via route-hijacking, MITM, and black-holing. These attacks will be carried out in both the classic and enhanced BGP versions, with the latter using our risk factors to determine how they sway route paths. Finally, we will tag a subset of the ASs in the topology as untrusted, where routes that contain them should be avoided. This topology is displayed in Figure 1.

To test our system, we will use a factorial experimental design (Figure 2). Our System Under Test is the system of interconnected routers running the Quagga process. Our three primary factors, or input parameters to the SUT will be: (1) the size of the network; (2) ASs to avoid (based on

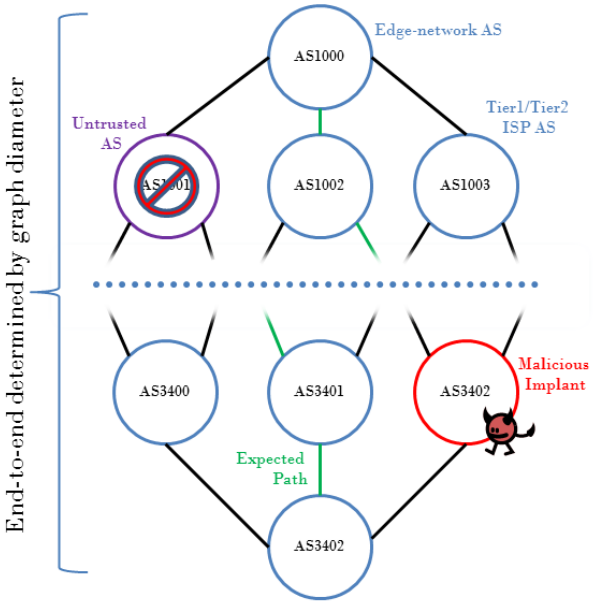

Figure 1: Experimental Topology

Figure 2: System Under Test

geolocation); and (3) malicious actors in the network. Two edge-network nodes in the system with the greatest diameter will be selected for the end-to-end communications. The output from the SUT will be the risk-scoring metrics used to determine the route (x_1, x_2, x_3), as well as the selected route itself. Based on the route selection, a binary metric will be used to capture whether the selected path is optimal for the risk-score (i.e., it does not traverse untrusted ASs and avoids the malicious implants).

The levels for each of the three primary factors are outlined in Table 1. Tests will be run on each of the network sizes of 10, 1K, and 10K ASs. For each of these network sizes, some percentage of ASs will be untrusted and some percentage malicous (randomized, with overlap allowed). The secondary factor, as mentioned above, will be experiments run in classic BGP, and those run in our enhanced BGP. Since the experiment is full-factorial, measurements will be gathered for every combination of the factors. Hence, the total number of base experiments is given by $3^3 * 2^1 = 54$. Additionally, to minimize variability and approach a normal distribution for the collections, each experiment will be executed 20 times between pair-wise edge networks, bringing the total number of experiments to $54 * 20 = 1080$.

It is our expectation that our analysis will show how well our system performs against classic BGP w.r.t. untrusted zones and malicious implants. Futhermore, we will also have

Table 1: Factorial Experiment Factor Levels

Parameter	1	2	3
Graph Density	100	1K	10K
Avoided ASs	10%	25%	50%
Implanted Hijacks	5%	10%	25%

the ability to statistically describe how well the risk overlay performs as the network, untrusted zones and malicious actors scale.

9. ADOPTION STRATEGIES

Practical solutions could eventually be advanced to the Internet Engineering Task Force (IETF).

Modifying a working, deployed system the scale of the Internet is a difficult enterprise since backwards compatibility must be maintained. Incentives for upgrades must be aligned to benefit individual networks for deployments to happen. S-BGP, DNSSEC and IPv6 are all poster-children for what happens when incentives are not aligned.

One advantage of using an overlay technique for the risk, is the risk metrics could be used by end clients and applications before it is actually used by ISPs and backbone networks for routing decisions. ISPs that don't provide the information would get a larger risk scoring than those that did provide the information (assuming the information was consistent with the performance metrics). This would steer traffic towards ISPs that did implement the protocol overlay, and would therefore increase the use of those networks and their ability to charge for service.

Additionally, given the desire of many nations to enforce national network boundaries, providing this information in existing networks would make national networks unneeded. Therefore government policy could make this happen, or industry could make it happen to forestall regulatory action.

10. LIMITATIONS

Special care will need to be taken when looking at making routing decisions based on feedback from performance metrics. Routing is not a closed system and modifying the routing will change the performance seen across those paths. History shows us that performance based routing decisions need to be tempered to avoid negative consequences [11].

Distributed control of the internet has allowed it to scale and be incredibly resilient in the face of network partitions (e.g. backhoes or bombs). Adding risk to the routing and sending decision making by definition ensures some of that traffic will not get through. The question is who should be making the trade-off decisions — and at what timescale. A general goal of this work is to make the routing fabric more transparent so all stake-holders have the ability to evaluate the risks.

Because it will take time to evaluate new paths as they change, and especially to incorporate performance feedback, risk scoring will be be delayed.

11. COMPARISON TO RELATED WORK

Risk has been addressed lately in the area of communications networks. For example, Cholda *et al.* [7] describe how risk-awareness would benefit communications networks. In particular, they delineate the risks associated with a loss of availability (e.g., network failures) and describe how to determine the business risk associated with network downtime. Thus their risk model focuses on economic consequences based on probabilities of different events happening that affect network availability and dependability. In contrast, in this paper we focus on *security* risks that might occur from using specific routes, where availability might not actually be impacted. Further, Cholda *et al.* develop an off-line risk model for organizations and suggest mitigations such as having multiple carriers and appropriate SLAs. In contrast, we have designed an overlay network that makes risk calculations in real-time (or near real time) in order to provide a judgement call on every connection regarding the security of the route.

There have also been efforts in the security space at applying risk models. For example, Teo *et al.* [18] have looked at creating a risk-aware network access management system. In their architecture, risk is used as an input in order to determine if network access rules should be dynamically modified (e.g., if there is a security threat). Some of the risk factors they use include malformed packets, malicious/abnormal packet content and anomalous behavior. This work focuses on determining if network access should be granted or denied at some border, based on the threat model of malicious traffic trying to enter a network. Other authors, such as Ahmed *et al.* [2], have looked at using security risk factors (such as known existing vulnerabilities and vulnerability trends) in order to determine appropriate security configuration for network services. In contrast to the referenced (and similar) work, our system is based on the threat model of *outgoing* traffic being intercepted *after* it has left our network, and therefore judging the risk of such actions based on the (inferred) route being taken.

Of closer relationship to our approach is a paper by Snigurov and Chakryan [16] that combines Quality of Service (QoS) and information security into a single metric. They present a risk score based on the probability of realizing a threat to confidentiality, integrity or availability, where the risk score is adjusted based on the information being transmitted, and then used as a weight on the particular route. Unlike our work, they do not state how the probabilities can be calculated to a practical level, nor do they consider that the exact route might not be known.

Also closely related is research on security within ad-hoc networks, where it is recognized that a node in the network can be a malicious actor. Rather than determining the risk of communicating with a given node, the research addresses the related concept of trust. We consider trust to be related to risk in this context because it is defined as being "based on the expectation that the other entity will perform a particular action important to the trustor, irrespective of the ability to monitor or control that other entity" [19] — a definition that is similar to our operational environment. However, Yan *et al.* calculate trust before there are any network communications, rather than determining trust (risk) on the fly (as we do in our model). Further, they allow for the calculation of a secure route through the network, while we operate under the condition that the route cannot be controlled.

Similarly, Yi *et al.* [20] have proposed Security-Aware ad hoc Routing (SAR), which uses security metrics as part of their route determination. Similar to our usage of risk metrics, the authors argue that applications must be able to specify the quality of protection or security provided by a route. As in the paper by Yan *et al.*, however, their approach assumes that the route can be determined ahead of time, and that its security level can therefore be ascertained, while our operating environment is more limited than this.

12. CONCLUSIONS

Risk analysis for network security has traditionally been investigated from one of two perspectives: either how do we keep malicious traffic from entering a network or how do you determine the riskiness of a particular network route where the network route is known *a priori*. Given the deployment of BGP as the core internet routing protocol, organizations in reality have no control over what route is chosen for their network traffic. The result has ranged from instances of services being knocked off the internet due to misconfigurations to suspected malicious hijacking of network traffic. There have also been increasing calls for assurances that network traffic will stay within specified (geopolitical) borders due to political reasons.

In order to address these threats, we have presented an approach to risk-based routing that accepts that an organization (AS) will never have control over a given route, and instead provides measures that it can use to determine the risk associated with sending the traffic on to the next hop based on the expected routing path. In addition to a description of the protocol, we have addressed issues including strategies to promote adoption of the protocol and realistic restrictions regarding routing speeds and requirements in the core network. We provide a testing methodology to determine the capabilities of the protocol, along with a discussion of three different real-world use cases.

13. REFERENCES

[1] Leila Abbound and Peter Maushagen. Germany wants a german internet as spying scandal rankles. `http://www.reuters.com/article/2013/10/25/us-usa-spying-germany-idUSBRE99O09S20131025`, October 2013. Last Visited: April 16, 2014.

[2] M.S. Ahmed, E. Al-Shaer, M.M. Taibah, M. Abedin, and L. Khan. Towards autonomic risk-aware security configuration. In *Proceedings of the 2008 IEEE Network Operations and Management Symposium*, pages 722–725, 2008.

[3] F. Baker. Requirements for IP Version 4 Routers. RFC 1812 (Proposed Standard), June 1995. Updated by RFCs 2644, 6633.

[4] Hitesh Ballani, Paul Francis, and Xinyang Zhang. A study of prefix hijacking and interception in the internet. In *Proceedings of the 2007 Conference on Applications, Technologies, Architectures, and*

Protocols for Computer Communications, SIGCOMM '07, pages 265–276, New York, NY, USA, 2007. ACM.

[5] L. Blunk, J. Damas, F. Parent, and A. Robachevsky. Routing Policy Specification Language next generation (RPSLng). RFC 4012 (Proposed Standard), March 2005.

[6] Martin Brown. Pakistan hijacks youtube. `http://www.renesys.com/2008/02/pakistan-hijacks-youtube-1/`, February 2008. Last Visited: April 16, 2014.

[7] Piotr Cholda, Eirik L. Folstad, Bjarne E. Helvik, Pirkko Kuusela, Maurizio Naldi, and Ilkka Norros. Towards risk-aware communications networking. *Reliability Engineering and System Safety*, 109:160–174, 2013.

[8] Jim Cowie. The new threat: Targeted internet traffic misdirection. `http://www.renesys.com/2013/11/mitm-internet-hijacking/`, November 2013. Last Visited: April 16, 2014.

[9] Dan Kaminsky. Be still my breaking heart. `http://dankaminsky.com/2014/04/10/heartbleed/`, April 2014. Last Visited: April 17, 2014.

[10] Akmal Khan, Taekyoung Kwon, Hyun-chul Kim, and Yanghee Choi. As-level topology collection through looking glass servers. In *Proceedings of the 2013 Conference on Internet Measurement Conference*, IMC '13, pages 235–242, New York, NY, USA, 2013. ACM.

[11] D. Richard Kuhn, Kotikalapudi Sriram, and Douglas C. Montgomery. Sp 800-54. border gateway protocol security. Technical report, National Institute of Standards & Technology, Gaithersburg, MD, United States, 2007.

[12] Robert Lychev, Sharon Goldberg, and Michael Schapira. Bgp security in partial deployment: Is the juice worth the squeeze? In *Proceedings of the ACM SIGCOMM 2013 Conference on SIGCOMM*, SIGCOMM '13, pages 171–182, New York, NY, USA, 2013. ACM.

[13] Harsha V. Madhyastha, Thomas Anderson, Arvind Krishnamurthy, Neil Spring, and Arun Venkataramani. A structural approach to latency prediction. In *Proceedings of the 6th ACM SIGCOMM Conference on Internet Measurement*, IMC '06, pages 99–104, New York, NY, USA, 2006. ACM.

[14] Louisa Schaefer. Deutsche telekom: 'internet data made in germany should stay in germany'. `http://www.dw.de/deutsche-telekom-internet-data-made-in-germany-should-stay-in-germany/a-17165891`, October 2013. Last Visited: April 16, 2014.

[15] J. Seedorf and E. Burger. Application-Layer Traffic Optimization (ALTO) Problem Statement. RFC 5693 (Informational), October 2009.

[16] Arkadij Snigurov and Vadim Chakryan. Approach of routing metrics formation based on information security risk. In *Proceedings of the 2013 CADSM*, pages 339–340, 2013.

[17] JamesP.G. Sterbenz, EgemenK. etinkaya, MahmoodA. Hameed, Abdul Jabbar, Shi Qian, and JustinP. Rohrer. Evaluation of network resilience, survivability, and disruption tolerance: analysis, topology generation, simulation, and experimentation. *Telecommunication Systems*, 52(2):705–736, 2013.

[18] Lawrence Teo, Gail-Joon Ahn, and Yuliang Zheng. Dynamic and risk-aware network access management. In *Proceedings of the 2003 ACM Symposium on Access Control Models and Technologies*, pages 217–230, 2003.

[19] Zheng Yan, Peng Zhang, and Teemupekka Virtanen. Trust evaluation based security solution in ad hoc networks. In *Proceedings of the Seventh Nordic Workshop on Secure IT Systems*, 2003.

[20] Seung Yi, Prasad Naldurg, and Robin Kravets. Security-aware ad hoc routing for wireless networks. In *Proceedings of the 2nd ACM International Symposium on Mobile Ad Hoc Networking and Computing*, pages 299–302, 2001.