Report on Routing Resiliency Measurements Workshop

November 2-3, 2012
Atlanta, GA, USA

Introduction

This workshop was next in the series of events organized by the Internet Society in our effort to foster improvements in the resiliency of the Internet routing system and facilitate adoption of common solutions and best practices in this area. Two most recent events preceding this workshop were operator roundtable meeting in December 2011 [1] and May 2012 [2]. Participants included representatives of the leading network operators.

Discussions with network operators at these roundtable meetings indicated a need for a better understanding of the history and current state of challenge. The issue of routing security has been an object of constant attention by the industry for more than a decade. BGP vulnerabilities were identified as early as 1988 and many of them still exist today. Despite seeming fragility, global inter-domain routing of the Internet demonstrated exceptional stability and robustness. Indeed, apart from a few high-profile cases when intentional or unintentional configuration mistakes affected global routing for a limited time, threats from the routing system rarely endanger service providers' commitments.

Long-term data and analysis on frequency and types of BGP attacks are an important element in informing operator's risk assessment and selection of adequate tools and approaches. What level of attack has there been in the past -- to what extent do security incidents happen, but go unnoticed, or get dealt with inside a single network, possibly introducing collateral damage? Are the number and impact of service disruptions and malicious activity stable, increasing or decreasing? Can we understand why, and track it collectively?

The workshop was intended to address these and other questions related to the issue of routing resiliency.

The workshop was divided into three main sections:

- **Measurement methodology and frameworks**
  The focus of this session was on different methodologies, their limitations and available data sets used for the analysis of suspicious events related to inter-domain routing in the Internet.
- **Research analysis and operational data**
  In this session participants looked at the data related to routing resilience coming both from the research as well from the operational experience. Participants explored how these data related to risks, vulnerabilities and threats.

- **Metrics and long-term monitoring**
  In this session participants discussed what metrics could be a useful representation of routing resiliency in the Internet. How can these metrics be used in improving routing resiliency? What data is needed and what needs to be done to get these metrics collected? Is there a need for a long-term monitoring and trend analysis and what can be done in this area?

Measurement methodology and frameworks

There are several methodologies of BGP and traffic data analysis that might be complementary. Correlation of these data can yield deeper insight.

The workshop began with a session where participants discussed different methodologies their strong points and limitations. Part of this discussion was looking at available data sets used for the analysis of suspicious events related to inter-domain routing in the Internet.

One of the approaches that was used in a research effort by NIST that used registered data (e.g., RIR, IRRs, RPKI) as well as historical BGP trace data (e.g., Routeviews, RIS) for the identification and detection of routing anomalies, was presented by Kotikalapudi Sriram.

There are existing BGP robustness algorithms which seek to generate reliable white lists of {prefix, origin AS} pairs from historical BGP trace data. These white lists are used to assist in BGP routing anomaly detection. Sriram discussed a methodology for characterizing and filtering unstable BGP announcements so that the quality of a white list can be improved making them more reliable.

The goal of BGP anomaly detection algorithms is to identify and catch serious anomalies (e.g., prefix hijacks) while minimizing the probability of a false alarm. The study compared three algorithms: (1) Registry-data driven, (2) History-data driven, and (3) Hybrid (i.e., combined registry and history data driven). The study showed that false alarm probability can be reduced by augmenting the history data with registry data. Hence, the hybrid algorithm provides a better performance for anomaly detection, especially when the registry data is reliable as in the case of the RIPE's RIR/IRR data.

Sriram observed that while the RPKI repositories are beginning to evolve at various RIRs, it may take a few years before those repositories are complete. Any available RPKI data (e.g., ROAs) can be combined with cleaned-up IRR registry data and stable historical data to enhance BGP robustness in the short term. More details about this work and approach are available in [3] and [4].

Another approach, focused on passive measurements of the data plane - looking at incomplete flows (one-way communication attempts - packets starting a communication to a remote end-host that do not receive a reply), was presented by Xenofontas Dimitropoulos from ETH Zurich. Such incomplete flows may be an indication of network failures, attacks and misconfigurations, and, therefore, their analysis can provide an interesting insight into network incidents.

The analysis included identification and classification of the main causes of one-way traffic in the Internet based on passive measurements of 7.41 petabytes of traffic collected over 8 (2004-2011) years from an academic backbone network in Switzerland. The research shows that one-way traffic makes between 34% and 67% of the total number of flows, although it only accounts for only 3.4% and 0.79% of the total number of packets and bytes, respectively. The main sources of one-way traffic are: 1) scanning, which accounts for 83.5% of the one-way flows and 62.6% of the one-way packets; 2) peer-to-peer applications, which account for 6.7% of the one-way flows

Figure 1. A 15-min internal outage in ETH Zurich as identified by one-way flows

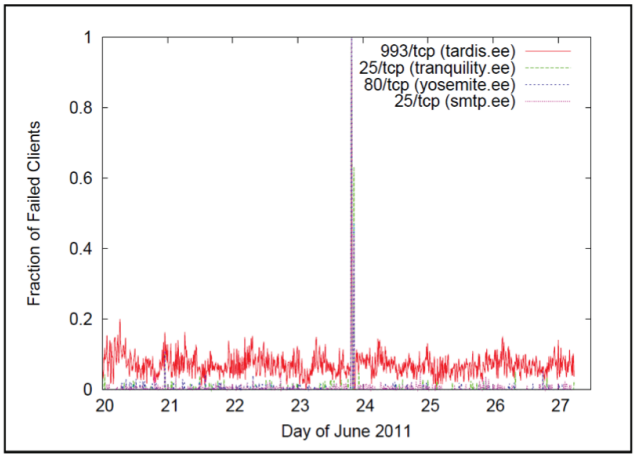

and 13.0% of the one-way packets; and 3) unreachable services, which account for 4.8% of the one-way flows and 10.1% of the one-way packets. One-way traffic measurements provide new possibilities for service availability monitoring enabling in particular to passively detect outages and assess their impact. Using the network of ETH Zurich as a case study, it was demonstrated how one-way traffic is useful for detecting network outages and for assessing their impact [5]. In one case, in particular, we identified a 15-minute outage, which affected the entire campus network of ETH Zurich (see fig.1). During this outage, 287,583 unique clients failed to access target services! More details about this work on detecting outages in local network services and in remote destinations can be found in [5], [6] and [7].

The Institute for Complex Additive Systems Analysis (ICASA) at the New Mexico Institute of Mining and Technology (NMT) has developed **capabilities that quickly provide information about Internet disruptions with global impact.** Specifically, these capabilities detect such disruptions, estimate the likely cause, and locate the effects in the topology of the Internet. This information is determined from open-source BGP data, using a combination of algorithms based on protocol-intrinsic data and behavioral models that capture more subtle information about the system. Dietrich Bachman presented this work.

In this approach ([8], [9]) reachability is defined on established sets of prefixes, categorized by any criteria desired, such as business entities, owner type (commercial vs. government vs. educational/NGO), or by country or registry. The first step is to determine for each prefix in the group whether it is routable at a given time. All of the answers to this question are aggregated to provide a simple 0-to-100% rating of the reachability of the whole group. If desired, multiple such samples can be calculated to build a plot over time. These plots can be used to observe massive reachability-degrading events, such as the government-ordered Internet shutoff in Egypt in 2011. Hidden Markov Models are employed to capture more subtle anomalies in the Internet's data plane, whose presence is sometimes not traceable in individual BGP updates, as with hijacks and reachability. These anomalies are visible in tracking the holistic behavior of the routing system, which HMMs accomplish by understanding (through training) a notion of "normal" behaviors of the system, as well as a host of previously observed abnormal behaviors. By employing multiple behavioral models of this kind, it is possible to both know when the system seems to be in an anomalous state, and to wager a guess as to what brought it there. Figure 2 shows the effect of invalid prefix origination by China Telecom in April 2010, where each data line represents the assessment of a historically-trained behavioral model to real-time incoming BGP data.

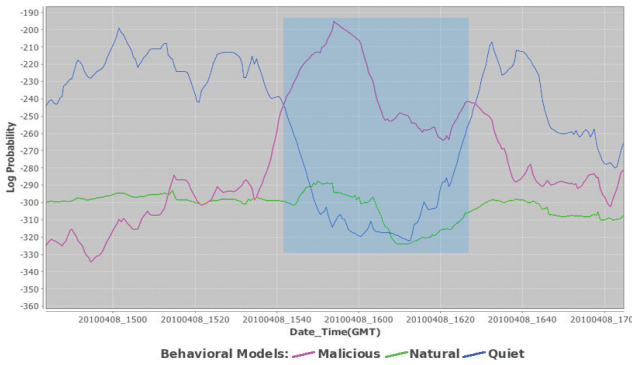

Figure 2. Classification HMM output during the event. Each model was trained on selections of BGP data concurrent with particular kinds of Internet stimuli (Malicious = worms, DoS, etc.; Natural = power outages, earthquakes, etc.; Quiet = periods of no known disruption).

Reduction of false positives remains an important objective of all presented efforts as the main requirement for their utility in operational tools. This task is challenging because of the dynamic nature of the BGP state globally as well as the lack of knowledge of the intent.

To reduce number of false positives researchers at the University of Arizona proposed a concept of **concurrent prefix hijacks** [10]. When a network originates prefixes of another network, it can be a hijack or due to an operational arrangement not known to the public. But when a network simultaneously originates prefixes of many other networks, it is highly likely to be a real hijack since operational arrangements with many different networks are unlikely to take effect at the same time. Based on this observation, they developed a scheme that detects concurrent prefix hijacks by correlating suspicious origin announcements and identifying networks that are offending many other networks simultaneously. The scheme doesn't require authoritative prefix ownership information. In order to facilitate real-time detection and reaction, scheme's parameters can be tuned to minimize false positives. As was demonstrated by some data in the next session and verified with network operators, this scheme can detect concurrent prefix hijacks with zero false positive.

Also, with regards to false positives tools like BGPmon [11], Cyclops [12] and EyeP [13] can provide a better picture, since the "intent" is documented in the tool itself by the users. Both BGPmon and Cyclops are network audit tools for service providers and enterprise networks, providing a mechanism to **compare the observed behavior of the network and its intended behavior.** Both are able to detect several forms of route hijack attacks, i.e. when Internet routes are maliciously diverted from their original state. Lixia Zhang from UCLA explained the capabilities of these tools.

In looking at routing incidents the main data source remains existing collections of BGP updates, mainly from the projects like RIS [14] and RouteViews [15]. Some methodologies combine it with other datasets, like flow information.

Discussion followed on what **other data sources** can be used to help reducing false positives – real routing incidents (attacks or simply misconfigurations) from normal changes in the state of the global routing system – changes of paths, origins and even reachability.

Various Internet Routing Registries (IRR [16]) provide dataset that can be used for this purpose. In the NIST research presented earlier ([3]) it was demonstrated how these data can further reduce false positives, despite the fact that these data are often claimed to be stale and incomplete.

While several presented approaches could effectively reduce false positives, from a research perspective the real unanswered question remains - how much goes unnoticed. How much confidence do we have that data coming from the analysis of limited data set, coming from limited number of vantage points with limited knowledge about the actual routing policies reflects the real state of affairs in this area?

Research analysis and operational data

Research data shows a relatively small number of routing incidents. This number goes up if a greater margin is allowed for false positives. It is difficult to say how much goes under the radar.

This corresponds well to the operator experience. It is quite possible that the actual number of incidents is higher, but they either go unnoticed, or get attributed to other causes. Operator's focus on routing security in a narrow sense is limited, more attention is on routing resilience in a wider sense.

Lixia Zhang presented some observations from the work on concurrent prefix hijacks [10]. Applying this detection scheme to RouteViews BGP data from 2003 through 2010, 5 to 20 concurrent prefix hijacks were detected each year (see fig. 3). They are typically hijacks of prefixes of a few tens of other networks, which last from a few minutes to a few hours, and pollute routes at most vantage points, meaning that the damage to data traffic could be widespread. Detected events from 2008, 2009, and 2010 were verified via email communication with network operators. All the 21 detected events were confirmed as real hijacks.

Figure 3. Number of concurrent hijacks detected

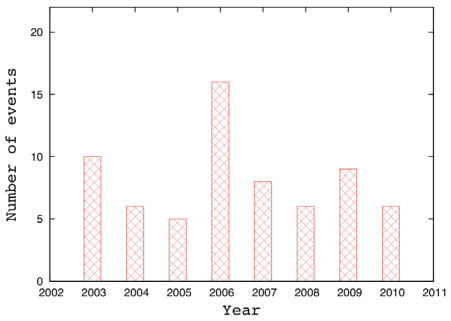

Interestingly most events are not mentioned in operator mailing lists such as NANOG or identified in research literature, implying that the network community in general is not aware of these hijacks. Furthermore, most operators of victim prefixes told us that they were unaware of their prefixes being hijacked. This result is a strong evidence of the urgency of the routing security problem and illustrates the importance of continuous measurement and monitoring.

When addressing routing security issues, as with any security-related activity, a network operator needs factual data and a good understanding of what's going on in the system to better inform the process of risk assessment and the selection of adequate tools and approaches. It is also important to measure the effect of such tools once they are deployed, and monitor the changing dynamics of the environment. Because the inter-domain routing system is globally interdependent, such monitoring and measurements should be long-term and be done on a global scale.

In this context, one important data set is **operational statistics of incidents related to routing security, as registered by a network operator.** The Internet Society attempted to collect some of these statistics by conducting a small informal survey among the operators.

Requested data included type, duration and frequency of registered incidents, as well as typical measures taken to resolve. A separate set of questions was dedicated to any other violations of policy detected. For instance, if an operator has prefix filters configured, how often BGP updates from a neighbor get filtered out.

Some preliminary results, although not being statistically representative, indicated that frequency of the incidents is very low, or they are not detected (whether they are occurring or not) at all. Given that many of routing incidents as was demonstrated in “Concurrent Prefix Hijacks: Occurrence and Impacts” [10] do not exceed 30 min, a case (a customer complaint, an alert from a monitoring system) might be already resolved before reaching NOC, or before closing a ticket. It was also suggested that because of low frequency operators may not have an explicit “closing code’ related to routing when closing a trouble ticket.

The most common problem experienced by the operators was number of prefixes received in a peering session exceeding the maximum prefix value. This typically happens 1-2 times a week and is a result of misconfiguration or simply growth of a peer. It appears that other policy violations were not tracked or logged dues to technical limitations or lack of tools (even if policy controls were in place).

The results also showed limited capabilities that operators have in place for collecting data and receiving alerts related to violations of the routing policy. That might be an indication that incidents happen infrequently, or that operators are simply not aware of, or not recognizing the incidents resulting from routing anomalies. But in any case, as was noted during the discussion, if one can’t get a rough probability of events and quantify the potential loss from risks being materialized then it is very difficult to convince operators to deploy preventive or mitigating system with non-zero incremental cost.

Another preliminary results were presented from the analysis of outage posts in NANOG (http://www.nanog.org/mailinglist/) and Outages (http://puck.nether.net/mailman/listinfo/outages) mailing lists, which are frequently used by network operators to report problems. Posts to these lists between Jan 2010 and Oct 2012 were manually examined looking for posts that clearly mentioned the words “outage”, “hijacking”, “route leak”, “unreachable” or “down”. Events that turned out to be internal network issues were omitted. Evolution of number of incidents over time is shown on fig.4.

Figure 4 Evolution of number of reported incidents on NANOG and Outages mailing lists

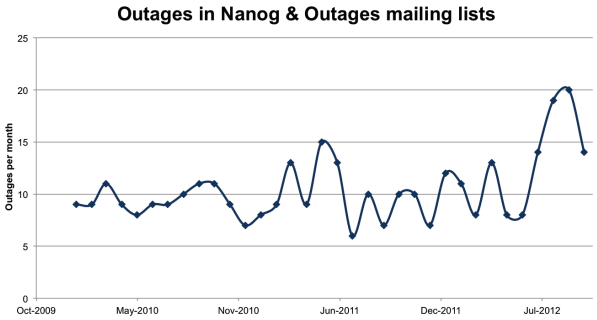

A total of 361 outages were identified and for 37.2% of them researchers were able to classify the causes. The distribution of outages by its cause is shown in the table:

Outage cause	Count	Percentage
DoS	6	1.7
Fiber cut	48	13.3
Hardware failure	17	4.7
Maintenance	4	1.1
Peering issue	4	1.1
Power outage	7	1.9
Congestion	4	1.1
Routing issue	18	5
Route hijack/leak	15	4.2
Unknown	226	62.8
Other	12	3.1
Total	**361**	**100**

This is work was conducted by ETH Zurich together with Mentari Djatmiko from NICTA, Australia. Related to this was a discussion around the question at what level of risks, or probability of events, proactive measures are more cost-efficient compared to reactive measures. Again, a better grasp on data could provide helpful answers in this area.

The discussion touched on another common violation of policy – so-called route leaks. A typical example of a route leak is a customer "leaking" routes learnt from its upstream provider to another (upstream provider). In this case the customer effectively becomes a transit provider for the leaked prefixes for the second provider, causing outages, decreased performance and collateral damage.

Route leaks are difficult to detect without knowing the intent (an operator's routing policy), but it was demonstrated that some heuristic techniques could be effectively applied under certain constraints. For instance, appearance in the AS-PATH of more than 2 "major networks" consequently, most probably indicates a route leak. Jared Mauch from NTT America demonstrated a tool (http://puck.nether.net/bgp/leakinfo.cgi) that he runs for more than 4 years, which analyses BGP updates collected by RouteViews looking for "suspicious" patterns in the AS-PATHs. Currently, 19 "major networks" are defined. There are known business relationships that are documented and excluded from the count.

Since the beginning of 2008 the system collected 9.6 million anomalies (BGP updates with a route leak). On some days the number of anomalous updates exceeded 180,000. Problems were manually reported to operators and in many cases got fixed with a tactical patch or policy changes. One of the prevention techniques is the implementation of filtering of customer updates (e.g. preventing "major networks" from being seen in the update).

Discussion of route leaks widened the context of routing resiliency. For example, as the number of transitive optional BGP attributes grows, the way they are handled, raises serious concerns among the operators. This opens an opportunity for launching a “packet bomb” into the infrastructure, resulting in nodes several hops away tearing down a session, causing surge of BGP updates and collateral damage.

A question emerged on whether **purely reactive measures**, informed by tools like Cyclops or BGPmon, are an effective approach for routing resiliency. Detect-and-React solutions have certain advantages and may be an effective routing infrastructure protection, because:

- They do not require changes to the routing Infrastructure
- They do not rely on any single component to function correctly
- They do not require participation of all networks
- And “given enough eyeballs, all faults will surface”.

At the same time some participants noted that purely reactive approach takes time and analysis and doesn’t always scale. Research and operational data showing the real level of threat from route mistakes and hijacks could allow operators to determine the optimal mix of reactive and proactive measures.

In this context it was emphasized that the importance of personal network among the operators and ability to easily find and approach a NOC contact cannot be overstated. Internet routing resiliency is not only based on technology, but to a great extent on expertise of technical staff and the ability to promptly take coordinated actions responding to a wide range of incidents affecting inter-domain routing.

At the end of the session an important point was made that even if individual operators experience may not warrant an action purely based on risk analysis and ROI, if we agree that level of “brokenness” globally demands improvements, they need to be facilitated in the spirit of collaborative responsibility, not only looking after operator’s own corner of the Internet.

Metrics and long-term monitoring

While the current state of routing resilience may look acceptable, long-term monitoring and trend analysis is necessary for better awareness and as an early warning system. Definition of useful and actionable metrics is crucial here.

The discussion started with a general question on the need of measurements and, more specifically, of what the metrics are, which that we can track over time and that will give operators the information they need to act, or not to act based on their threat model?

As was presented by Doug Montgomery from NIST, the need for measurements boils down to answering questions about the **problem space** (e.g. How “big” is the routing security problem in the existing / future system? What are the frequency and ramifications of historical events? What are the risks of known existing vulnerabilities?) and the **solution space** (How “effective” is/would-be a given approach in mitigating these problems? In reducing the frequency/likelihood and ramifications of events. What is the rate of false-positive/undesirable side effects? At what cost of deployment/operation?).

When talking about the problem space the challenge is in the development of reliable metrics, especially if we want to measure the effectiveness of deployed solutions. Since addressing events that go unnoticed is the biggest problem, how can we begin to characterize the **level of uncertainty** in our measurements?

Talking about the solution space, some participants mentioned that reasonable estimates of risks and ramifications are important information for deciding an appropriate set of solutions. It was further discussed whether historical data can help in identifying the risks. It is very clear that vulnerabilities in the routing system exists and they have been exploited many times, as research and operational experience shows. It was also noted that in a security context sometimes potential threats are more important than the historical data.

While everyone agreed that the system is vulnerable to misconfigurations, the participants engaged in the discussion of whether the routing system is an attractive target for malicious attacks. It was noted that attacking control plane of the infrastructure might be less attractive than attacking the services that run on top of that infrastructure, and besides, for the latter one needs the routing infrastructure working. Acknowledging that such analysis is speculative, it was suggested that a long-term phenomenal approach may provide a useful insight into how the system evolves. For example, a shift in the attack vector (e.g. from services to the infrastructure) would be manifested by a change in the metrics.

In other words, if we can determine the state of health of the Internet, and the routing system specifically, - useful and actionable metrics, - then we can do trend analysis and have a better understanding of whether there is a problem that is growing or decreasing, and whether collaborative efforts are working to address actual problems. This doesn't necessarily help identify solutions directly -- but goes to what people care about (brokenness of the Internet).

Following up on the discussion of vulnerabilities and risks, Eric Osterweil from Verisign Labs presented the concept of "attack surface" as an approach to is to quantify the potential vulnerabilities a complex system may face. However, given the high degree of interdependencies among networked systems, it remains an open challenge whether it is possible to quantify the attack surface of such systems.

He presented a novel methodology that offers a repeatable way to quantify the systemic dependencies of real Internet systems. He presented two specific examples of how this methodology can be applied: the X.509 CA verification system and a newly standardized alternative approach called DANE, and early thoughts on the RPKI. This methodology might help to both model the systemic dependencies of actual Internet-scale systems, and to formally quantify the often elusive notion of a system's attack surface.

Looking forward

Development of metrics and consistent vocabulary

The workshop participants discussed a number of various methodologies that can be used in the analysis of routing security and resiliency. Good examples of such analysis were also presented and discussed. At the same time the participants noted lack of common set of metrics that can usefully and consistently describe the state of the routing resiliency. The development of such metrics could be very useful for long-term monitoring and trend analysis, and in fact can facilitate such efforts.

Part of the problem lies in the fact that there is also no vocabulary used to describe various incidents. This partly explains the challenge in conducting the operator survey that the Internet Society conducted informally in preparation for this workshop. The taxonomy of incidents varies from operator to operator and often doesn't include classification and codes specifically related to routing resilience. The development of common vocabulary is crucial for collection of consistent operational statistics.

More research into "false negatives" – how much is going unnoticed

From an operator's perspective and the utility of tools based on the analysis of the routing system the reduction of false positives in the main objective. However from a point of view of the overall understanding of the resiliency of the routing system it is important to estimate how much, or what types of the events got filtered out or have not been noticed at all.

Raising awareness of the real safety level

If an operator is not monitoring violations of their routing policy (route hijacks, leaks, etc.) they are not aware of real threats coming from the routing system. Subsequently, there is little motivation to deploy additional controls. Also, the effectiveness of the deployed measures is hard to estimate. Better monitoring and collection of operational data may calibrate operational experience – resulting in an increased awareness of routing incidents, better understanding of their operational and economic impact. It might also reveal the real origin of such incidents as related to routing security, otherwise resolved with unrelated "closing codes".

Focus on the low hanging fruit – existing best practices tailored to what needs to be most protected

Deployment of some of the best practices, including prefix filtering, limiting number of prefixes received from a neighbor, as well as improvements in the detection and mitigation techniques can be the most effective first step in improving routing resilience. The application of these practices is well understood and if coupled with operator's understanding of the critical elements of the infrastructure, can result in a substantial protection.

References

[1] Routing Security: Report on 2nd Internet Society Operator Roundtable, 13 - 14 December 2011. http://www.internetsociety.org/routing-security-report-2nd-internet-society-operator-roundtable

[2] From Routing Security to Secure Routing, Report on 3rd Operator Roundtable on Routing Security, 31 May 2012, http://www.internetsociety.org/routing-security-report-3rd-internet-society-operator-roundtable.

[3] K. Sriram, O. Borchert, O. Kim, and P. Gleichmann, and D. Montgomery, "A Comparative Analysis of BGP Anomaly Detection and Robustness Algorithms," Proceedings of the Cybersecurity Applications and Technology Conference for Homeland Security (CATCH), Washington D.C., March 3-4, 2009, pp. 25-38. http://www.nist.gov/itl/antd/upload/NIST_BGP_Robustness-2.pdf

[4] O. Kim, K. Sriram, O. Borchert, P. Gleichmann, and D. Montgomery, "An Analysis of ARIN NetHandles with OriginAS Data and Analysis of RIR/IRR Registry Data", Presented at ARIN XXIII, San Antonio, TX, April 26-29, 2009. https://www.arin.net/participate/meetings/reports/ARIN_XXIII/pdf/monday/nethandles.pdf

[5] E. Glatz and X. Dimitropoulos. Classifying Internet One-way Traffic. ACM SIGCOMM IMC Internet Measurement Conference, Nov. 2012.

[6] Classifying One-way Traffic in the Internet. http://www.ow-class.ethz.ch/

[7] D. Schatzmann, S. Leinen, J. Kögel and W. Mühlbauer. FACT: Flow-based Approach for Connectivity Tracking. Passive and Active Measurement conference, Mar. 2011.

[8] K. Glass, R. Colbaugh, M. Planck. "Automatically Identifying the Sources of Large Internet Events", http://www.icasa.nmt.edu/Content/publication/LogicalLocalizationInternet.pdf

[9] M. Planck, K. Glass, I. Lyman and R. Colbaugh "A Framework for Near Real-Time Event Characterization Within the Internet", http://www.icasa.nmt.edu/Content/publication/DisruptiveInternetEvents.pdf

[10] V. Khare, Q. Jun, B. Zhang, “Concurrent Prefix Hijacks: Occurrence and Impacts,” ACM/USENIX Internet Measurement Conference (IMC), 2012, www.cs.arizona.edu/~bzhang/paper/12-imc-hijack.pdf

[11] BGP Monitoring System (BGPmon), Colorado State Universityб http://bgpmon.netsec.colostate.edu/

[12] Cyclops, UCLA, http://cyclops.cs.ucla.edu/

[13] EyeP, UCLA, http://eyep.cs.ucla.edu/

[14] Routing Information Service (RIS), RIPE NCC, http://www.ripe.net/data-tools/stats/ris/routing-information-service

[15] Route Views Project, University of Oregon, http://www.routeviews.org/

[16] Internet Routing Registries, http://www.irr.net/docs/list.html

List of participants

- Shane Amante, Level 3 Communications, Inc.
- Dietrich Bachman, NMT ICASA
- Leslie Daigle, Internet Society
- Alberto Dainotti, CAIDA, UCSD
- Xenofontas Dimitropoulos, ETH Zurich
- Mat Ford, Internet Society
- Phillipa Gill, The Citizen Lab
- Jared Mauch, NTT America
- Danny McPherson, Verisign
- Douglas Montgomery, USA NIST
- Andy T. Ogielski, Renesys
- Eric Osterweil, Verisign Labs
- Andrei Robachevsky, Internet Society
- Phil Roberts, Internet Society
- Kotikalapudi Sriram, USA NIST
- Lixia Zhang, UCLA

Internet Society
Galerie Jean-Malbuisson, 15
CH-1204 Geneva
Switzerland
Tel: +41 22 807 1444
Fax: +41 22 807 1445
www.internetsociety.org

1775 Wiehle Ave.
Suite 201
Reston, VA 20190
USA
Tel: +1 703 439 2120
Fax: +1 703 326 9881
Email: info@isoc.org