Routing Resiliency Survey

An Internet Society Pilot Project to Analyze Routing Incidents and their Impacts

Internet attacks or instability resulting from vulnerabilities in the global routing system have been the topic of interest and concern for many years. The Internet Society convened a study to better understand the scope and volume of these phenomena with the help of the network operator community.

Background

Improving the security and resilience of the global routing system is key to ensuring the Internet is a reliable platform for communication around the world. Border Gateway Protocol (BGP), outlined in RFC4271[1], is the routing protocol used in the global inter-domain system. As stated in the "Good Practices in Resilient Internet Interconnection[2]" ENISA report, "given its key importance to the Internet, BGP is surprisingly susceptible to malfunctions, and despite recent technical advances and significant attention devoted to its vulnerabilities, BGP is still considered by most providers and experts "the Achilles' heel of the Internet". Indeed, most of the recent interconnection-related incidents and Internet outages have revolved around BGP failures and vulnerabilities, and as BGP disorders tend to spread fast within this protocol monoculture, many incidents can have regional, national or even international impact."

Vulnerabilities of BGP are known and are well documented in RFC 4772 "BGP Security Vulnerabilities Analysis"[3]. One of the threats is the possibility to inject false (bogus) routing information in the global routing system by a BGP speaker (a router) that may get propagated by other networks.

From RFC 4772: "Bogus routing information can have many different effects on routing behavior. If the bogus information removes routing information for a particular network, that network can become unreachable for the portion of the Internet that accepts the bogus information. If the bogus information changes the route to a network, then packets destined for that network may be forwarded by a sub-optimal path, or by a path that does not follow the expected policy, or by a path that will not forward the traffic. Consequently, traffic to that network could be delayed by a path that is longer than necessary. The network could become unreachable from areas where the bogus information is accepted. Traffic might also be forwarded along a path that permits some adversary to view or modify the data."

One example of such bogus information is an announcement by a network of a prefix that does not belong to that network. This phenomenon is called "prefix hijacking" and is an occasional cause of outages and a tool for DoS attacks. In most cases, traffic destined to that prefix gets "black-holed" and is not delivered to the intended network. Such announcements may be caused by misconfiguration or by malicious intent to mount a DoS attack. Prefix hijacking was the main focus of this project.

Survey Approach

Operational data is an important foundation for monitoring developing trends and making rational decisions to address security issues related to routing. It is also important to measure the effect of routing security tools and technologies once they are deployed. Because the inter-domain routing

system is global, such monitoring and measurements should be long-term and be done on a global scale.

Various measurement and research initiatives exist in this area. In 2012, the Internet Society hosted a workshop to facilitate discussion about measurement frameworks and collected data. A workshop report can be found at http://www.internetsociety.org/doc/report-routing-resiliency-measurements-workshop. One of the conclusions was that there is low awareness at the individual operators' level of the risks and that better monitoring and data collection may calibrate operational experience, resulting in an increased awareness of routing incidents and better understanding of their operational and economic impact.

However, currently, there is no coordinated approach across network operators to collect or analyze this kind of data, especially from an impact point of view – a basis for risk assessment and global trend analysis – leaving network operators with only incomplete or anecdotal evidence for understanding Internet-wide routing security issues and making adequate decisions.

To address this gap, the Internet Society ran a Routing Resilience Survey pilot project[4] aimed at the collection of incident data related to routing resiliency. The project, undertaken in partnership with the Border Gateway Protocol (BGP) monitoring service BGPmon (http://www.bgpmon.net), aimed to analyze various aspects of routing security as well as provide a six-month snapshot of routing incidents and their impacts, as registered by network operators.

Questions we tried to answer with this survey:

- How frequently do incidents happen and what are their impacts?
- To what extent do network operators care about routing security?
    - Monitoring (reactive tools)
    - Routing controls (proactive tools)
- Are there any regional differences with regard to incidents?
- How does the picture from the outside differ from the operator's perception/impact?

We did not ask about how the incidents were resolved, nor did we investigate the cause of incidents.

Survey

Participants were asked to provide two kinds of information:

1. General network information (e.g. number of peers, clients, transit providers, etc.), and
2. Data related to routing security incidents via an automated monitoring effort.

General network information

At the beginning, participants were asked to complete a registration process, filling out a web form containing questions related to network type, connectivity, and practices used in mitigating routing security incidents.

Data related to routing security incidents via an automated monitoring effort

BGPmon generated alerts related to changes in originating Autonomous System Numbers (ASNs) for the participant prefixes. The list of prefixes that were monitored for a given AS was generated from the observed BGP announcements.

Initially, each participant was presented with several past possible security incidents (if any severe incidents were detected over the previous year). After that, a weekly report of recent incidents was generated and made available to only that participant via the survey web portal. An email notification was also sent indicating that new events were available. A screenshot of the incident dashboard that was available to participants is provided on Figure 1 below.

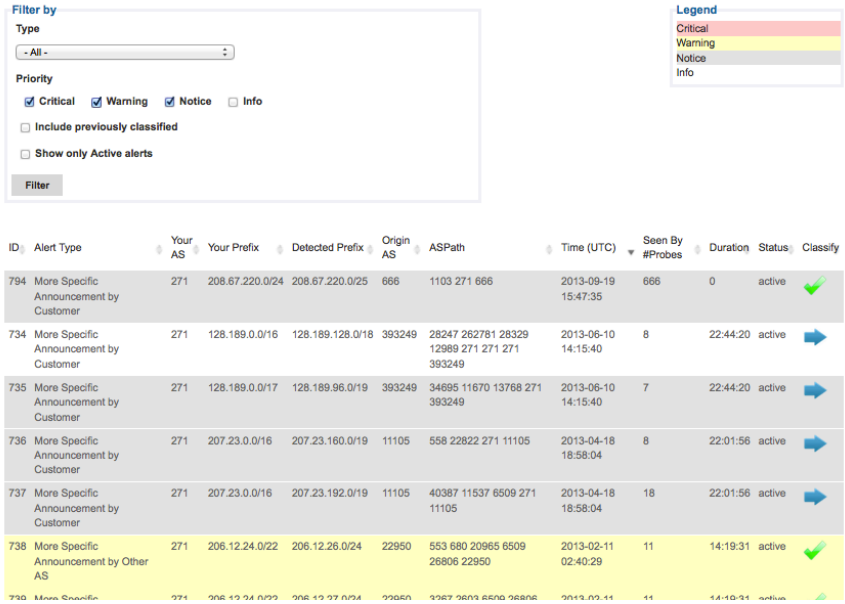

Figure 1. An example of the incident dashboard.

To validate the information in the overview, a participant needed to log into the portal and provide additional information for every incident listed. This additional information included:

Severity of the incident:

- Severe, caused prolonged service disruption and/or customer complaints
- Moderate
- Insignificant
- Cannot determine / not related
- Not an incident

How the participant learned about the incident:

- Network monitoring system alert
- Customer call
- This alert
- Not an incident

Privacy Concerns

The Internet Society respected the sensitivity of some of the data involved in this effort. Therefore, it committed to ensuring participant-specific information remained confidential. All data collected was stored on Internet Society servers. Any information or analyses shared beyond a specific network was fully anonymized.

Overview of Participating Networks

More than 30 operators signed up for the Survey, among which 27 operators were active in completing it. The participants represented a mixture of Tier 1, Tier 2/3, cloud and content delivery, enterprise, and other types of network operators from all around the world. Because participants were also allowed to classify events related to their customers, the Survey represented 239 Autonomous Systems in total.

Before participating in the incident classification effort, each participant indicated what routing security controls they deployed in their networks. The results of this network information survey are presented in Figure 2 below.

It is interesting to see that almost 90% of the participants impose a maximum amount of prefixes they allow to receive from a BGP neighbor. Looking at stricter tools like filtering, 80% of participants use prefix filtering for their customers and peers (almost 70% do this for their customers and about 30% for peers; 20% do both), more than 60% apply some sort of AS-PATH filtering and 80% filter "bogon" announcements, like reserved and private space. Less than half of the participating networks (44%) register routing information in an IRR that they subsequently use to generate filters. None of the networks used RPKI.

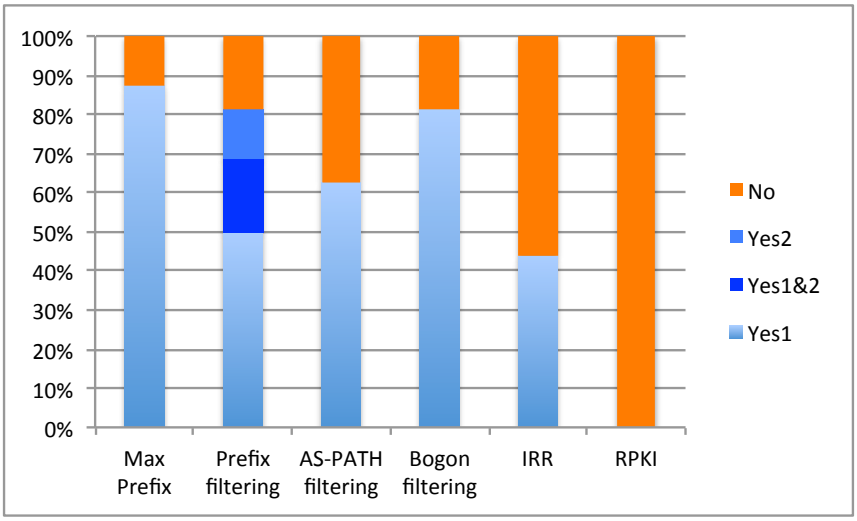

Figure 2. Use of routing security controls among the participating networks. For prefix filtering, Yes1=applied to customer announcements, Yes2=applied to peer announcements, Yes1&2=applied to both.

These statistics may not be representative for the Internet at large, since the network operators involved in this Survey were interested in participating in the incident classification effort, which could be an indication that they were aware of routing security issues and took them seriously.

Incident Data Analysis

BGPmon generated events based on observed changes in BGP announcements for participating networks. Not all changes were reported as an event, though.

BGPmon reported only based on historical data. So if an AS began announcing a prefix that hadn't been announced before, such event would not be reported. Also, incidents involving manipulation of AS-PATH attribute were not reported. For example, if an AS that is not authorized to announce a particular prefix prepends an AS number of the legitimate AS, such incident would not be detected and reported.

BGP updates when an existing prefix, overlapping prefix (less specific) or a more specific prefix were announced by another AS were reported as events.

Impact Severity

The distribution of registered incidents and their impact is shown in Figure 3 below. The chart only shows incidents that happened during the duration of the project (November 2013 – June 2014) and excludes historical events. We also removed the “unknown” and “not an incident” events from the graph, to focus on incidents that had some impact on network operations. And even with our relatively small set of surveyed networks, one can notice moderate and some severe incidents.

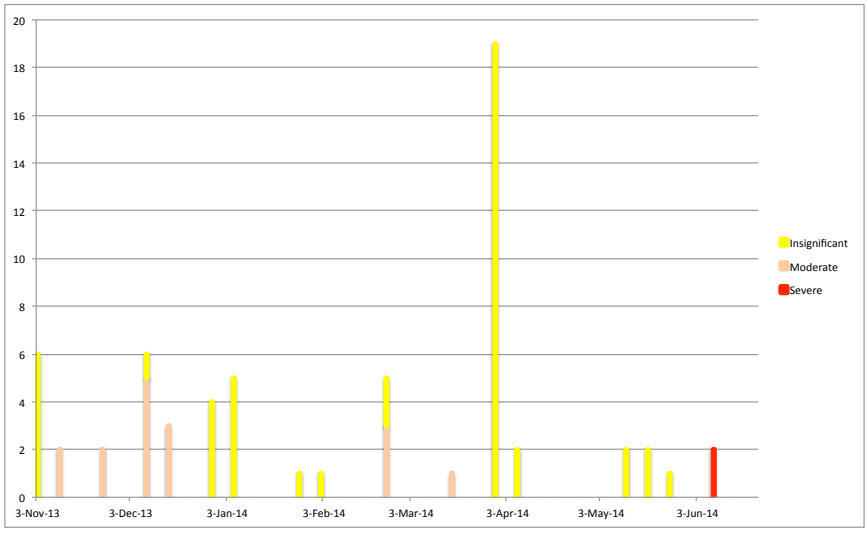

Figure 3. Distribution of classified incidents (events that were marked as Severe, Moderate, or Insignificant) for all participants excluding historical events.

To give an idea of the distribution of events we included also “not an incident” and “unclassified” events in Figure 4 below. It is clear that false positives dominate the generated events, which may indicate one of the reasons why monitoring tools are not widely used.

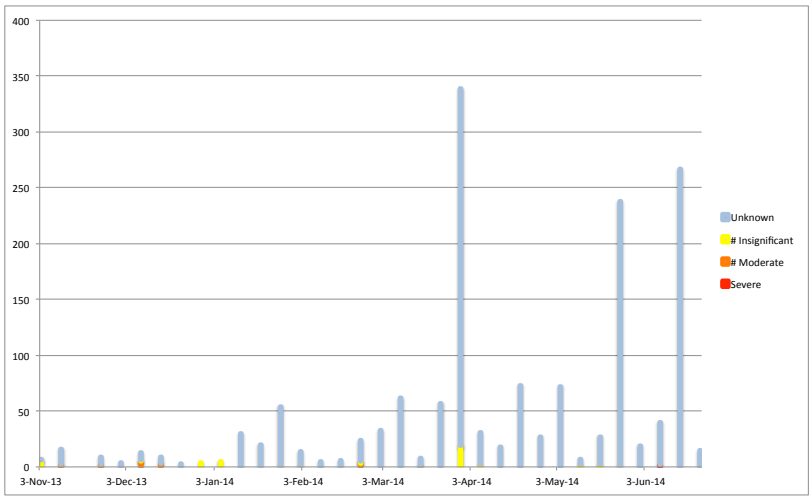

Figure 4. All events for all participants, including the false positives: “not an incident” and “unknown” (excluding historical events).

It is interesting to look at the difference in perception about the severity of an event from an operator's perspective and a third party's observing changes in the global BGP state. For the latter we used BGPmon's classification of events based on the type of the event (e.g. a new prefix announced by the same customer, or a new AS) and its global visibility (i.e. number of peers that saw this event). This chart is presented as Figure 5 below.

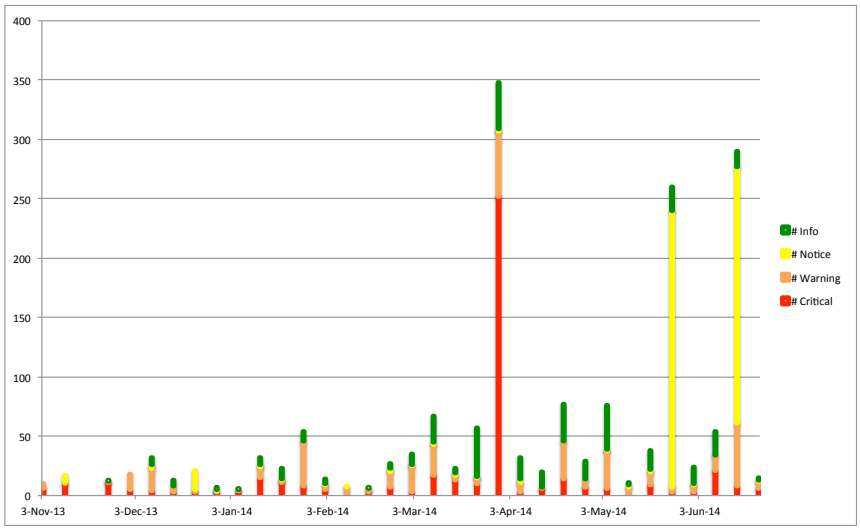

Figure 5. Distribution of the events by "criticality" estimated by the BGPmon service based on heuristics.

It is clear that the impact from an outsider's point of view seems to be higher than the impact that the operators experienced themselves. Part of this might be related to the fact that the BGPmon heuristics allow for a certain level of false positives, partly because some of the participants were not able to classify the events due to lack of resources, but also because some of the incidents went unnoticed and were difficult to track back (we asked to classify events retrospectively, based on weekly reports).

If we look at the impact from a slightly different angle, see Figure 6 below, we see that false positives constitute at least 18% of all events. Real incidents sum up to about 4% of all events. These numbers may be higher, depending on what is in the "unknown" category – events that were not classified.

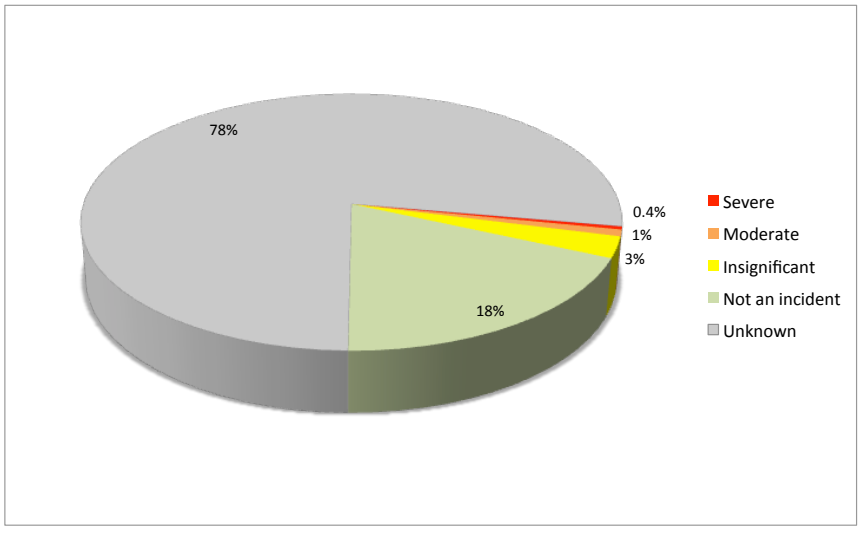

Figure 6. Distribution of incidents by severity of impact

Again, if we compare this distribution with the one from a BGPmon point of view the difference is significant (Figure 7):

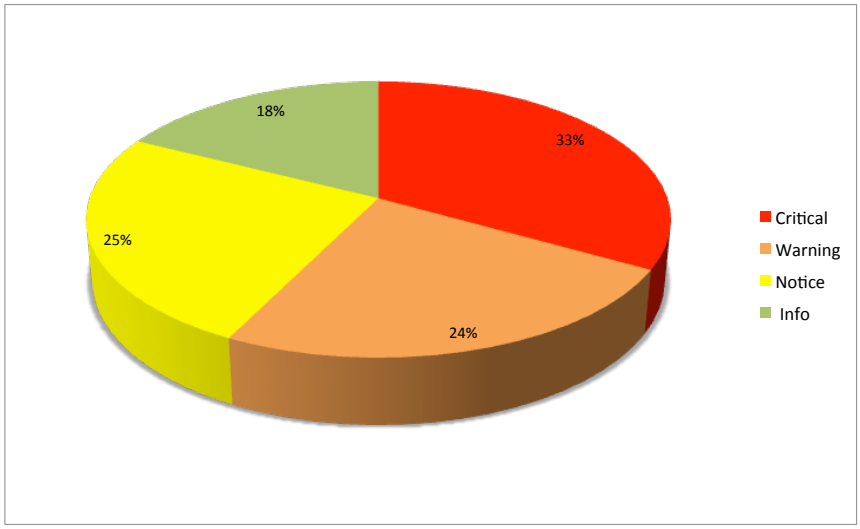

Figure 7. Distribution of events as classified by BGPmon

Learning About the Incidents

For each event, the participant was given four options to indicate how they learned about the incident. The results are presented in Figure 8 below.

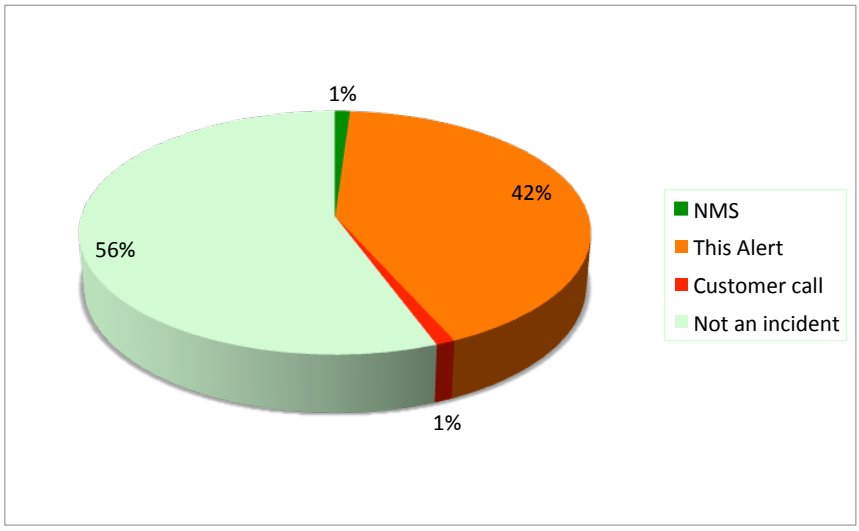

Figure 8. How participants learned about an incident

From this distribution, we see that more than half of all events that participants classified were false positives. Of the rest, the majority would have probably gone unnoticed if the participant hadn't received an alert from this Survey. Customer calls and deployed network monitoring systems (NMSes) contribute equally, each 1% of all classified events.

Frequency and Duration of the Incidents

As seen from the results of the data collection, incidents happen quite rarely. Of 239 monitored autonomous systems, only 10% of them registered an event at all during the six-month duration of

the Survey. From a NOC perspective (participants that monitored their networks as well as their customer networks), severe incidents happen on average once every three to six months, and insignificant, but still requiring reparations, events happen on average once or twice per month.

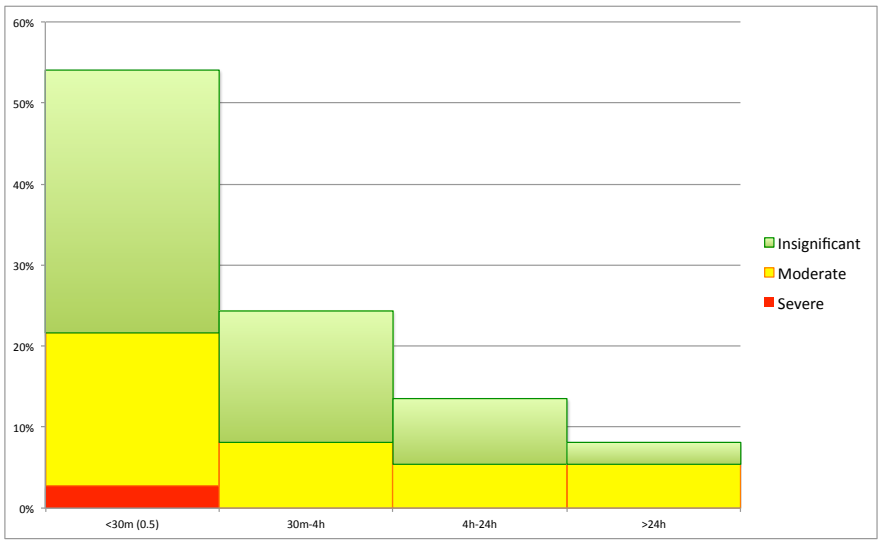

Figure 9. Distribution of the duration of incidents

The distribution of the duration of incidents is shown in Figure 9 above. More than half of all incidents are resolved within 30 minutes. All severe incidents registered during the Survey lasted less than 30 minutes. Most of the incidents were resolved within 24 hours.

Conclusions and Recommendations

Incidents with real impact are rare. That means that while network operators are aware of the vulnerabilities of the routing system, risks associated with them are perceived as low. In such circumstances, reactive measures seem to be more appropriate and proactive protection is deployed only if it has low operational costs associated with it.

High percentage of false positives. The difference between an outsider's view of suspicious routing events and an operator's own experience is very significant. Even for classified incidents, false positives constituted at least 56% of the total. In practice this number is even higher. This means that using external monitoring as raw input to the NOC may be challenging. At the same time, with some knowledge of intent, the number of false positives can be dramatically reduced. This can be done by the NOC itself (e.g. by filtering out alerts related to a connected new customer) or by publicizing the "intent" (e.g. by creating corresponding records in an IRR or RPKI repository – ROAs, or directly to a monitoring service).

Incidents are fixed quite quickly. Moderate and insignificant incidents do not last longer than a few hours, and severe incidents are mitigated even more quickly. This adds more justification for using proactive measures, but their effectiveness relies strongly on the effectiveness of global coordination among network operators.

Further steps

The pilot indicated some challenges. Overcoming these could have improved the results of the effort and its impact.

- Only half of all events were classified. That leaves some questions unanswered – were these events unimportant false positives participants did not bother classifying, or was it lack of participant resources, especially if problems were real and resources were needed to cope with them?
- Low number of participants made the results less statistically representative and did not allow for exploring other types of correlations, like geographical differences.
- We did not look at **how** incidents were mitigated. We also did not explore the **causes** of incidents, which may in some cases be more obvious to an external observer than to the affected party.

Possible further steps:

1. Better understand motivating factors for operators to participate in such efforts.
2. Focus on specific incidents that got external visibility and analysis (e.g. YouTube, IndoSat, etc.) continuously monitoring developing trends in terms of cause, nature and impact of these incidents. That will reduce the notification frequency, their focus and reduce number of false positives significantly.
3. While reducing the number of alerts and making them more targeted to the surveyed group, request more extensive information for each incident – how was it mitigated, what was the cause of it, etc.
4. Improve the user interface.
5. Produce timely reports, adding value to the participation in the Survey.

[1] http://datatracker.ietf.org/doc/rfc4271

[2] https://www.enisa.europa.eu/activities/Resilience-and-CIIP/critical-infrastructure-and-services/inter-x/resilience-of-interconnections/enisa-report-on-resilient-internet-interconnections/at_download/fullReport

[3] http://datatracker.ietf.org/doc/rfc4272

[4] https://www.internetsociety.org/rrs/

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