A Survey among Network Operators on BGP Prefix Hijacking

Pavlos Sermpezis[1], Vasileios Kotronis[1], Alberto Dainotti[2],
Xenofontas Dimitropoulos[1,3]

[1]FORTH-ICS [2]CAIDA, UC San Diego [3]University of Crete

ABSTRACT

BGP prefix hijacking is a threat to Internet operators and users. Several mechanisms or modifications to BGP that protect the Internet against it have been proposed. However, the reality is that most operators have not deployed them and are reluctant to do so in the near future. Instead, they rely on basic - and often inefficient - proactive defenses to reduce the impact of hijacking events, or on detection based on third party services and reactive approaches that might take up to several hours. In this work, we present the results of a survey we conducted among 75 network operators to study: (a) the operators' awareness of BGP prefix hijacking attacks, (b) presently used defenses (if any) against BGP prefix hijacking, (c) the willingness to adopt new defense mechanisms, and (d) reasons that may hinder the deployment of BGP prefix hijacking defenses. We expect the findings of this survey to increase the understanding of existing BGP hijacking defenses and the needs of network operators, as well as contribute towards designing new defense mechanisms that satisfy the requirements of the operators.

1 INTRODUCTION

BGP prefix hijacking 101. Autonomous Systems (ASes) use the Border Gateway Protocol (BGP) [15] to advertise address space (as IPv4/IPv6 network prefixes) and establish inter-domain routes in the Internet. BGP is a distributed protocol, lacking authentication of advertised routes. As a result, an AS is able to advertise illegitimate routes for IP prefixes it does not own. These advertisements propagate and "pollute" many ASes, or even the entire Internet, affecting service availability, integrity, and confidentiality of communications. This phenomenon, called *BGP prefix hijacking*, is frequently observed [25], and can be caused by router misconfigurations [1, 2] or malicious attacks [3, 22, 25].

Current defenses are not sufficient. Currently, networks rely on *practical reactive mechanisms* to defend against prefix hijacking, since *proactive mechanisms* such as RPKI [16–19, 24] are fully efficient only when globally deployed, and operators are reluctant to deploy them due to associated technical and financial costs [11,12,14,20,21]. Reactive mechanisms mainly operate in two stages: *detection* (*e.g.,* based on monitoring data) and *mitigation* (*e.g.,* based on local network actions, such as originating BGP advertisements) of the hijack. The speed of the reactive defenses is crucial; even short-lived events can have severe consequences [3]. However, the reality shows that, currently, hijacking events are not quickly mitigated. For instance, back in 2008, a hijacking event affected YouTube's prefixes and disrupted its services for 2 hours [9]. More recently, in Sep. 2016, BackConnect (AS203959) hijacked, at different times, several ASes; the events lasted for several hours [4]. In Jan. 2017, the Iranian state telecom TIC hijacked disparate pornographic websites for more than a day [5]. In Apr. 2017, financial services, like Visa and Mastercard, and security companies, like Symantec, were hijacked by a Russian company for seven minutes [6].

Survey motivation and contributions. To surpass existing shortcomings and achieve a swift and efficient resolution of hijacking events, new defense approaches that fit the needs and requirements of the operators are needed. To this end, we launched a survey [8] to increase the understanding of currently used BGP hijacking defenses, and to receive feedback directly from network operators about their needs. The main motivation for this survey was to propose and design such a defense approach that matches the community needs; in fact, we acquired valuable information from this survey that helped us design a defense system called ARTEMIS [23].

However, the findings of the survey are more general and can be beneficial for both researchers and operators. Researchers can evaluate the severity of the problem of BGP prefix hijacking as it is seen from the operator community, and investigate new defense mechanisms capitalizing on current operational practices. Operators can be informed about the trends in the BGP prefix hijacking issue and the employed defenses, provide valuable feedback to the network community themselves, and adjust accordingly the way they manage and protect their networks against hijacks.

Structure. In Section 2 we present the questions of the survey, and in Section 3 we discuss the main findings and their

implications. The detailed results are presented in Figures 1, 2, and 3.

2 SURVEY PROFILE AND QUESTIONS

We launched a survey [8] on network operators' mailing lists, such as NANOG and RIPE. The survey is anonymous and comprises 21 questions studying *(a)* the operators' awareness of BGP prefix hijacking attacks, *(b)* presently used defenses against BGP prefix hijacking, *(c)* the willingness to adopt new defense mechanisms, and *(d)* reasons that may hinder the deployment of BGP prefix hijacking defenses. We received answers from 75 participants operating a broad variety of networks (Fig. 1(a)) all over the world (Figs. 1(b) and 1(c)), working at different positions (Fig. 1(d)).

The survey/questionnaire is composed of three parts.

(1) Information about the participants and their organizations (4 questions). In the first part we ask the participants to provide information about the type (*e.g.,* ISP, CDN, IXP) and location of their organization, as well as their work position in the organization. The questions and results are presented in Fig. 1.

(2) Knowledge and Experience with BGP Prefix Hijacking (6 questions). The second part consists of questions related to the participants' awareness and concern about BGP prefix hijacking, including their experience with past hijacking events on their networks. The questions and results are presented in Fig. 2.

(3) Defenses against BGP Prefix Hijacking (11 questions). The last part asks the participants about (i) the defenses they use (if any) against BGP prefix hijacking, such as RPKI, (ii) how they detect and mitigate a hijacking event affecting their prefixes, and (iii) the characteristics they consider desirable (or not) in a future defense (detection/mitigation) system. The questions and results are presented in Fig. 3.

3 SURVEY RESULTS

We classify the survey findings in 4 categories, which we present in the following sections: (i) evaluation of impact of hijacks (Section 3.1), (ii) general information about current defense mechanisms employed against hijacks (Section 3.2), (iii) specific information on the detection and mitigation stages in today's operations (Section 3.3), and (iv) requirements posed on new mitigation mechanisms (*e.g.,* involving outsourcing defense functionality to third parties), as well as the willingness of operators to adopt them (Section 3.4).

3.1 Impact of Hijacks

BGP prefix hijacking is a real threat and concerns network operators. More than 40% of the operators reported that their organization has been a victim of a hijack in the past (Fig. 2(e)). However, the vast majority is concerned about BGP prefix hijacking in the Internet (Fig. 2(b)) and its potential impact on their own networks (Fig. 2(c)). Almost all operators are knowledgeable on the issue of hijacks and the involved mechanisms (Fig. 2(b)).

Hijacks have a severe and lasting impact. Operators evaluate the impact of a potential hijack targeting their network (in terms of *duration* and *number of disrupted services*) as shown in Fig. 2(d). The vast majority (76%) expects the impact of a hijack to last for a long time (few hours or more), while opinions are divided on whether the hijack will affect a few or many of their services/clients, indicating that there are concerns both for extended (*e.g.,* route leaks) and limited/targeted (*e.g.,* malicious attacks) hijacks. Moreover, their past experience (Fig. 2(f)) shows that most hijacks indeed lasted long: more than 57% of hijacks lasted more than an hour, while 25% lasted *more than a day*; around 28% are short-term hijacks, lasting a few minutes (14.3%) or seconds (14.3%).

3.2 Defenses against Hijacks

RPKI deployment is limited. In accordance with previous studies [13], most of the network operators (71%) answered that they have not deployed RPKI as a proactive defense mechanism in their networks (Fig. 3(a)); very few (12%) use the full functionality of RPKI (Route Origin Authorisation - ROA and Route Origin Validation - ROV). There are various reasons for this, as shown in Fig. 3(b); deployment lags mainly due to RPKI's *limited adoption* and *little security benefits*, but also due to the increased *CAPEX and OPEX costs*, and increased *complexity* and *processing overhead* associated with the protocol mechanisms. Therefore, about 60% of the operators (Fig. 3(c)) resort to other mechanisms and practical defenses to protect their networks against BGP hijacks.

Practical defenses include route filtering, extensive peering, and de-aggregation. The responses to the (optional) question "what other defense mechanisms are used by networks" are shown in Fig. 3(d). The majority of the participants, *i.e.,* 17 networks (among those who provided answers for this optional question), use *route filtering* as a proactive defense to protect their own and their customers' prefixes from being hijacked. *Route filtering* is implemented in various ways (based on their answers) including for example: prefix origin (e.g., from IRR records) or AS-path filtering; filtering at edge routers (with customers/peers) or route servers (at IXPs). Less popular approaches are *anycast* (2 answers) and *prefix de-aggregation* (4 answers). Finally, 5 operators (from CDNs or tier-1 networks) mention that they peer with many other networks extensively; this helps them protect their networks from hijacking events (*i.e.,* by reducing their impact).

3.3 Detection and Mitigation of Hijacks.

Hijack detection mainly relies on third parties. The majority of networks (61.3%) use a third party detection service, which notifies them about hijacking incidents against their prefixes (Fig. 3(e)). BGPmon [7] is the most popular detection service, according to Fig 3(f). The satisfaction of operators from third parties generally varies a lot; some use them because they are satisfied and others because there are no alternatives (*e.g.,* it is not possible to develop their own detection service)[1]. Moreover, 17.3% of networks also practically rely on third parties, since they expect to get notified about a hijack by receiving notification from colleagues, clients, mailing lists, etc. In total, 78.6% of the networks rely on third parties for the detection of hijacks against their prefixes. About one third of the networks have deployed a local hijack detection mechanism (*e.g.,* by monitoring the disruption of their services)[2]. Finally, a non-negligible percentage of 8% would probably not learn about a hijack.

Mitigating through de-aggregation and contacting other networks. Asking operators what would be the countermeasures they would take to mitigate a prefix hijacking[3], the majority (62.7%; Fig. 3(g)) responded that they would announce more specific prefixes (de-aggregation) and contact the offending network (i.e., the hijacker) or its providers. 5.3% would follow *only* the former approach (de-aggregation) and 25.3% *only* the latter (contacting other operators). This indicates that although de-aggregation is not widely used currently (see Fig. 3(d)), operators still find it a good solution and are willing to proceed to similar actions after a hijacking event–affecting their–networks has taken place.

3.4 New Mitigation Mechanisms

The survey results show that the main practices that networks currently use for hijack mitigation comprise prefix de-aggregation and contacting other networks (Figs. 3(d) and 3(g)). Since these approaches have some important shortcomings, *e.g.,* de-aggregation is not efficient when a /24 prefix is hijacked (due to upstream filtering), and contacting network operators is usually done manually and thus adds significant delay to the mitigation process, we ask the network operators about their willingness to deploy new mitigation mechanisms, as well as what desired characteristics these mechanisms should possess.

We first ask them about their willingness to outsource functions related to the detection and mitigation of hijacks to a third party, in order to enhance their defenses. 61% of the operators are not willing to proceed to such outsourcing practices (Fig. 3(h)). This shows that a potential mechanism should not be entirely based on outsourcing, since this would not be acceptable by many networks. Flexible approaches that could be operated in two modes, *i.e.,* self-operated and outsourced, could be promising, since a significant percentage of 39% does not reject the possibility to outsource such functions.

The reasons for the operators' reluctance to outsource are given in Fig. 3(i), where the associated (high) cost and the need to share private information about their network are the main factors. Administrative and technical overhead may also prevent outsourcing. This is a first indication about the characteristics of a potential defense system: low cost, privacy-preserving, and easy to operate and manage.

More specific results about what would be the information/control that they would *not* be willing to share/allow with an outsourcing organization, are given in Fig. 3(j). As it can be seen, most of them are willing to share information about their prefixes and AS-neighbors (95%), as well as their routing policies (80%). A smaller percentage would allow BGP announcements to be controlled or implemented by the outsourcing organization.

Finally, according to operators, the importance of different characteristics that a hijack defense system should have, is shown in Fig. 3(l) (ranked from the highest to the lowest importance). A graphical representation of the importance of these characteristics for the network operators is given in Fig. 3(k), where the rightmost characteristics are considered of the highest importance. The speed and effectiveness of the mitigation stage, as well as the self-operability and low cost and management overhead, are the highest-ranked characteristics. Moreover, the detection stage is required to generate few false positives, which indicates the need for high levels of detection accuracy.

4 CONCLUSION

In this work, to increase community understanding of existing BGP hijacking defenses and the needs of network operators, we presented the results of a survey of 75 network operators around the world.

Through the survey, we verified our intuition that BGP prefix hijacking is a real threat and concerns the vast majority of network operators; in fact, hijacks can have a severe and lasting impact on their own networks. In the context of

[1]This variation can be observed in the following examples from the detailed answers in our survey:
"pretty happy with it", "no issues so far", "it works fine [..], but is relatively limited", "I hate it", "It's ok", " it seems to work quite well the few times I have needed it", "Better than nothing, but a lot of false alerts", "It rules!", "It is very noisy because it does not know a damn thing about IXP route servers", "Not great".

[2]Among the "other" answers, a high percentage of answers relates to the observation of -or, reception of complaints about- disruption in their services.

[3]Note that for this question, we provided the choices, based on the answers received in a preliminary version of our survey; operators could have answered in a different way if this was a completely open question.

combatting such hijacks, operators can use proactive or reactive techniques. On the one hand, proactive mechanisms, such as RPKI, have gained extremely little traction for multiple reasons, including limited adoption and high cost and complexity of deployment. On the other hand, practical reactive defenses such as contacting other networks, route filtering, extensive peering and prefix de-aggregation are usually preferred methods to mitigate hijacks; however, each has its own significant limitations, ranging from very slow mitigation speeds (*e.g.,* contacting other operators) to inefficient mitigation (*e.g.,* de-aggregation for /24 prefixes).

In terms of detection, we observe that operators mainly rely on third parties, such as BGPmon. However, the level of satisfaction varies wildly across operators. Moreover, most of them are reluctant to perform similar outsourcing for the mitigation of the hijacks themselves; in fact, there are mixed feelings about the kind and amount of information they would be willing to disclose to the third party, as well as the involved costs and technical and administrative overhead. The speed and effectiveness of the mitigation stage, as well as the self-operability and low cost and management overhead, are of paramount importance; moreover, the detection stage is required to generate few false positives, mandating high levels of detection accuracy. The findings of this survey could inform the design and implementation of new concepts and methodologies, such as ARTEMIS [10, 23], as well as more secure inter-domain routing protocols in general.

ACKNOWLEDGEMENTS

This work was supported by the European Research Council grant agreement no. 338402, the National Science Foundation grant CNS-1423659, and the Department of Homeland Security (DHS) Science and Technology Directorate, Cyber Security Division (DHS S&T/CSD) via contract number HHSP233201600012C.

REFERENCES

[1] https://www.ripe.net/publications/news/industry-developments/youtube-hijacking-a-ripe-ncc-ris-case-study.
[2] http://www.bgpmon.net/chinese-isp-hijacked-10-of-the-internet/.
[3] https://www.wired.com/2014/08/isp-bitcoin-theft/.
[4] http://seclists.org/nanog/2016/Sep/122.
[5] http://dyn.com/blog/iran-leaks-censorship-via-bgp-hijacks/.
[6] https://arstechnica.com/security/2017/04/russian-controlled-telecom-hijacks-financial-services-internet-traffic/.
[7] BGPmon (commercial). http://www.bgpmon.net.
[8] Survey on BGP prefix hijacking. http://tinyurl.com/hijack-survey.
[9] YouTube Hijacking: A RIPE NCC RIS case study. http://www.ripe.net/internet-coordination/news/industry-developments/youtube-hijacking-a-ripe-ncc-ris-case-study, March 2008.
[10] G. Chaviaras, P. Gigis, P. Sermpezis, and X. Dimitropoulos. ARTEMIS: Real-time detection and automatic mitigation for bgp prefix hijacking. In *Proc. ACM SIGCOMM Demo*, 2016.
[11] D. Cooper, E. Heilman, K. Brogle, L. Reyzin, and S. Goldberg. On the Risk of Misbehaving RPKI Authorities. In *Proc. of ACM Workshop on Hot Topics in Networks (HotNets-XII)*, 2013.
[12] W. George. Adventures in RPKI (non) Deployment. https://www.nanog.org/sites/default/files/wednesday_george_adventuresinrpki_62.9.pdf, 2014. NANOG presentation.
[13] Y. Gilad, A. Cohen, A. Herzberg, M. Schapira, and H. Shulman. Are we there yet? on RPKI's deployment and security. NDSS 2017, to appear, http://eprint.iacr.org/2016/1010.pdf, 2016.
[14] S. Goldberg. Why is it taking so long to secure internet routing? *Communications of the ACM*, 57(10):56–63, 2014.
[15] S. Hares, Y. Rekhter, and T. Li. A border gateway protocol 4 (bgp-4). https://tools.ietf.org/html/rfc4271, 2006.
[16] J. Karlin, S. Forrest, and J. Rexford. Pretty good bgp: Improving bgp by cautiously adopting routes. In *Proc. IEEE ICNP*, 2006.
[17] S. Kent, C. Lynn, and K. Seo. Secure border gateway protocol (s-bgp). *IEEE Journal on Selected Areas in Communications*, 18(4):582–592, 2000.
[18] M. Lepinski. Bgpsec protocol specification. 2015.
[19] M. Lepinski, R. Barnes, and S. Kent. An infrastructure to support secure internet routing. 2012.
[20] R. Lychev, S. Goldberg, and M. Schapira. BGP Security in Partial Deployment: Is the Juice Worth the Squeeze? In *Proc. of ACM SIGCOMM*, 2013.
[21] S. Matsumoto, R. M. Reischuk, P. Szalachowski, T. H.-J. Kim, and A. Perrig. Authentication Challenges in a Global Environment. *ACM Trans. Priv. Secur.*, 20:1:1–1:34, 2017.
[22] A. Ramachandran and N. Feamster. Understanding the network-level behavior of spammers. *ACM SIGCOMM Computer Communication Review*, 36(4):291–302, 2006.
[23] P. Sermpezis, V. Kotronis, P. Gigis, X. Dimitropoulos, J. H. Park, D. Cicalese, A. King, and A. Dainotti. ARTEMIS: Neutralizing BGP Hijacking within a Minute. arxiv.org/abs/1801.01085, 2017. arXiv 1801.01085.
[24] L. Subramanian, V. Roth, I. Stoica, S. Shenker, and R. Katz. Listen and whisper: Security mechanisms for bgp. In *Proc. NSDI*, 2004.
[25] P.-A. Vervier, O. Thonnard, and M. Dacier. Mind your blocks: On the stealthiness of malicious bgp hijacks. In *Proc. NDSS*, 2015.

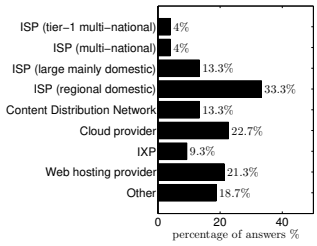

(a) Q1: Which term(s) would best characterize your organization?

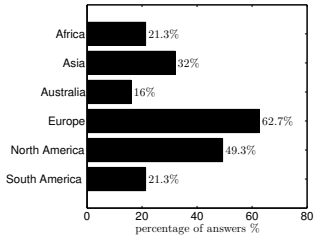

(b) Q2: In which continent(s) does your company operate?

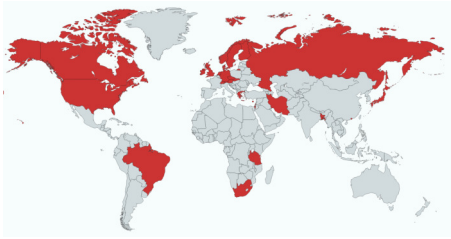

(c) Q3: In which country(-ies) does your company operate? (optional)

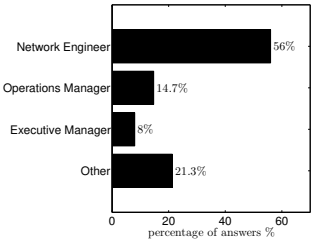

(d) Q4: What is your position in your company?

Figure 1: Survey results – Information about the participants and their organizations

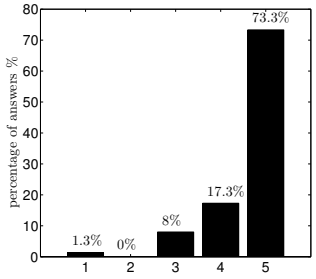

(a) Q5: Do you know what BGP prefix hijacking is and how it can happen?

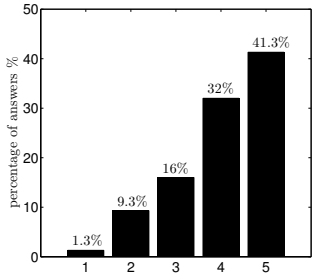

(b) Q6: How concerned are you about BGP prefix hijacking incidents on the Internet?

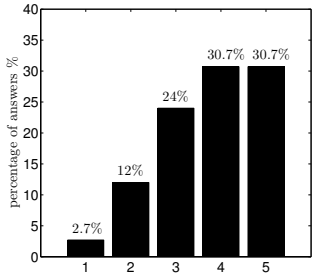

(c) Q7: Are you concerned that your network may be a victim of a BGP prefix hijacking incident in the future?

	no impact	~min.	~hours
few services/ clients	0%	9.3%	28.0%
many services/ clients	0%	9.3%	48.0%

(d) Q8: How severe do you consider the potential impact of a BGP prefix hijacking against your network?

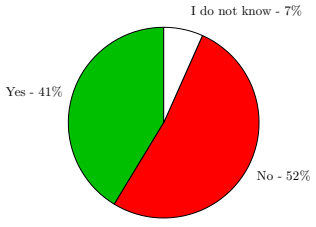

(e) Q9: Has your organization been a victim of a BGP prefix hijacking incident in the past?

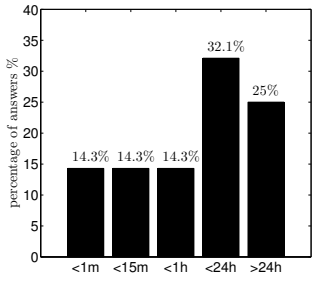

(f) Q10: If your organization was a victim of a BGP prefix hijacking incident, for how long was your network affected? (optional)

Figure 2: Survey results – Knowledge and Experience with BGP Prefix Hijacking

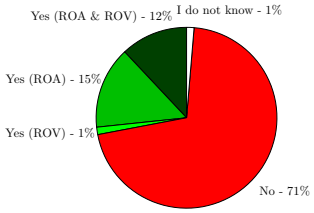

(a) Q11: Do you use RPKI in your network?

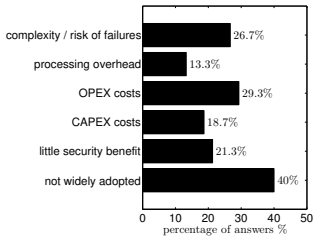

(b) Q12: If no (in Q11), what are the main reasons for not using RPKI? (optional)

(c) Q13: Do you use in your network any other defense mechanisms (other than RPKI) that protect your/others' prefixes from BGP prefix hijacking?

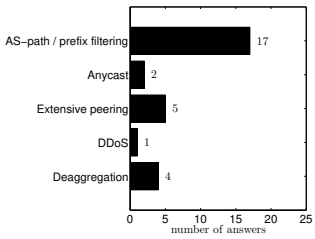

(d) Q14: If yes (in Q13), what mechanisms do you use? Could you provide a brief description? (open question/answers - optional)

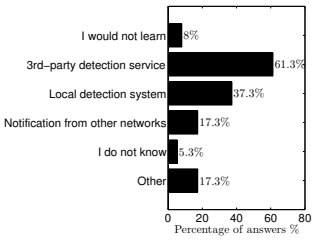

(e) Q15: In your network, how would you learn about a hijacking incident against your prefix(es)?

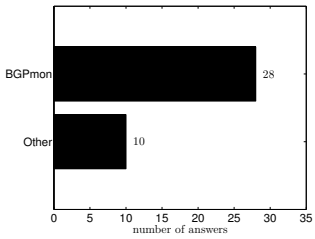

(f) Q16: If you use a local or third-party detection service or system, could you please give us more details about it? (open question/answers - optional)

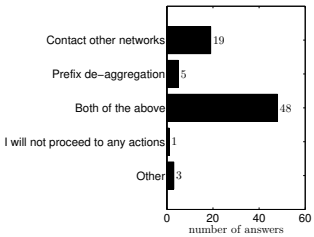

(g) Q17: How would you mitigate a hijack against your prefixes if you were notified about an on-going event?

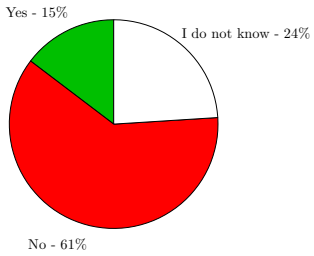

(h) Q18: Would you outsource functions relating to the detection and mitigation of prefix hijacking incidents to a third-party, if this helps your organization reduce its risks?

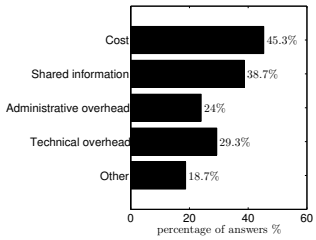

(i) Q19: If no (in Q18), what are the main factors that would affect your decision not to outsource prefix hijacking mitigation? (max 2 answers - optional)

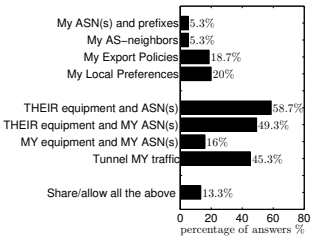

(j) Q20: Assuming you fully trust an outsourcing organization for prefix hijacking mitigation, what is the information/control (if any) you are still NOT willing to share/allow?

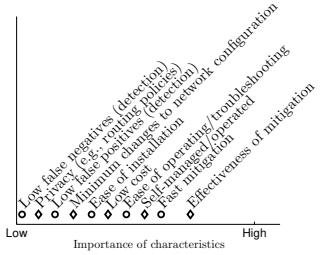

(k) Q21: How important do you consider the following characteristics for the deployment of a new defense system in your network?

Importance (0: Low ... 3:High)	0	1	2	3
Effectiveness of mitigation	0 %	2.7%	49.3%	48.0%
Fast mitigation	2.7%	14.7%	38.7%	44.0%
Self-managed/operated	1.3%	18.7%	38.7%	41.3%
Ease of operating/troubleshooting	0 %	21.3%	40.0%	38.7%
Low cost	6.7%	17.3%	52.0%	24.0%
Ease of installation	1.3%	30.7%	53.3%	14.7%
Minimum changes to network configuration	5.3%	26.7%	49.3%	18.7%
Low false positives (detection)	2.7%	12.0%	29.3%	41.3%
Privacy (e.g., routing policies)	8.0%	32.0%	32.0%	28.0%
Low false negatives (detection)	1.3%	13.3%	36.0%	34.7%

(l) Detailed answers for Q21

Figure 3: Survey results – Defenses against BGP Prefix Hijacking