This paper was presented as part of the 14th IEEE Global Internet Symposium (GI) 2011 at IEEE INFOCOM 2011

Exploiting BGP Scoping Services to Violate Internet Transit Policies

Pierre Francois
ICTEAM
Université catholique de Louvain (UCL)
Belgium
pierre.francois@uclouvain.be

Paolo Lucente
KPN International
The Netherlands
paolo.lucente@kpn.com

***Abstract*— Inter-domain path propagation limitation services are commonly implemented by IP transit providers, to offer flexible traffic engineering means to their customers. By combining such limitation services with the injection of more specific prefixes, operators of an Internet service provider can bring the global Internet routing system in a state which violates the policies of some other Internet service providers.**

In this paper, we describe the conditions for such policy violations to happen, and provide a specific analysis for the case of violations within the Tier-1 Internet service provider clique. We discuss the incentives for each actor of the violation and show that these can be beneficial to the parties causing them.

We present two approaches to tackle this problem. The first one is to instrument ways to detect the violations and react a posteriori, while the violation is happening. The second approach is to pro-actively defend against the violations, notably by black-holing policy violating traffic. Acknowledging its simplicity, we finally advocate for the former approach, as we observe that typically, the ISPs concerned by the violation have established business relationships and may hence prefer solutions which do not break the IP transit service.

I. Introduction

With the years, while the Internet profitability has moved on value-added services to end-users, providing wholesale and retail clean-pipe Internet access has become a very competitive and low margin business. In this context, sparing on IP transit costs thus became a tempting choice for many ISPs.

To be attractive or more efficient on typical customer support queries, many IP transit providers allow their customers to perform flexible inbound traffic engineering, by influencing the set of peers and transit providers to whom their paths are advertised. Such a flexible routing service is required to reach some desired transit market share distribution among the upstream providers of an ISP. This service is implemented by letting the BGP decision and filtering processes be influenced by path attribute setting by the customers, and has been found to be extensively used over the Internet [1].

In the data-plane of the routing system, however, the longest prefix match forwarding rule precedes the application of BGP policies. The existence of a prefix p that is more specific than a prefix P in the routing information base will let packets whose destination matches p be forwarded according to the best nexthop obtained for p. The forwarding behavior for packets destined to p thus disregards the policies applied in the control plane for the selection of the best nexthop for P. When combined with carefully tweaked advertisements of a more specific prefix, this data-plane property can lead to policy violations at another ISP. Such violations are obtained by letting the routing system converge to policy compliant states for both p and P, with a limited propagation of the routing information of p, that we refer to as “prefix scoping”. Due to this restriction of the scope of p, some data-plane paths will be made of the combination of forwarding states for P and p. In some cases, such combined data-plane paths do not fit with the policy of some ISPs.

ISPs which do not defend against these may find themselves offering “free transit”. More importantly, ISPs which do not give themselves the means of detecting policy violations may offer such free transit for long periods or large amounts of traffic. Also, ISPs offering prefix scoping services to their customer branch may be identified as the culprit of a policy violation happening at a competitor’s network, even though policy violation was not the goal underlying the provided service.

This document describes the routing services and the operational habits leading to such potential threats, in Section II. In Section III, we provide a detailed analysis of the conditions under which those policy violations can happen, depending on properties of the scoping services. We present a specific analysis for policy violations happening at the Tier-1 level, and discuss the incentives for an ISP to trigger such violations. In section IV, we discuss two families of solutions to the problem. The first one is to let an ISP detect the occurrence of such violations, and take non technical actions to solve the problem. The second family of solutions tries to anticipate the occurrence of the problem and not let those policy violations occur. We advocate for the first family of solutions, as we observe that it is practically hard to enforce the respect of such policies without risking to transiently black-hole traffic ultimately destined to customers, or render the scoping service useless.

II. More specific prefixes advertisement and scoping

In this section, we describe the BGP routing features which, when jointly used, can lead to policy violations. We emphasize

978-1-4244-9920-5/11/$26.00 ©2011 IEEE

their utility to achieve flexible routing, and observe how they can lead to policy violations.

A. Definitions

More specific advertisement is defined as the propagation of paths towards a sub-block of an IP block within the BGP routing system, *in addition* to the propagation of the paths for the IP block itself. As an illustration, in Figure 1, MHS advertises a path towards $P = 10.0.0.0/20$, to its providers, A, B, and E. In addition, it advertises a path towards a more specific prefix, $p = 10.0.0.0/24$, only to B and E.

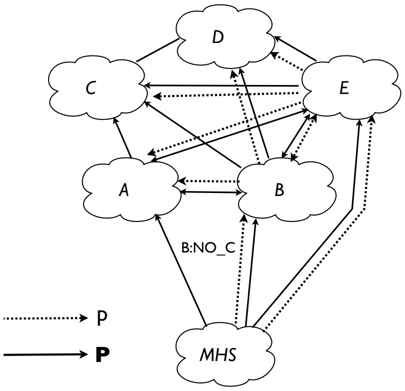

Fig. 1. Reference Environment

Prefix scoping is defined as the use of BGP configuration tweaks, which result in an incomplete dissemination of the paths towards a given prefix. That is, with prefix scoping, paths towards the prefix are only propagated over a subset of the inter-domain links over which it would have been propagated according to the commonly established routing policies of the ASes forming the routing system.

In the example of Figure 1, the selective advertisement of p to only ISP B and E is a form of prefix scoping. Prefix scoping is also typically implemented as a service offered by an ISP to its customers, in order to let them better control their incoming traffic. For example, the customers of the Sprint network can decide to which of Sprint's peers a path for a given prefix will be propagated further. This scoping is achieved through the use of BGP communities tagged to the paths upon their advertisement to Sprint's routers [2]. In Figure 1, upon the advertisement of its path for p to ISP B, MHS tags the path with a dedicated BGP community, "$B : NO\_C$", notifying ISP B that this path should not be propagated further to ISP C.

Depending on the provided service, the customer can use combinations of such communities to tweak path propagation with interesting expressivity, as discussed hereafter.

B. Effects of Scoping More Specifics

Restricting the propagation of a path towards a more specific prefix to some selected peers and providers of an AS can lead to traffic market share shifts. While these shifts remain compliant with individual ISP policies, and could not be achieved without the use of scoping services, their flexibility can also lead to a global routing system such that the transit service of other ISPs is abused.

1) In-policy Business Shifts: Let us analyze the routing system depicted in Figure 1. MHS originates P, which it advertises to A, B, and E. With the selective additional advertisement of p to B and E, MHS enforces the traffic for p to only flow in from B and E. ISP A does not receive a customer path for p, but only receives one from its peers B and E. A will thus reach p through B or E, while keeping using its direct link with MHS for destinations in P that are not covered by p. ISP A will not re-advertise the peer path for p to its other peers. The result of this selective advertisement of p by MHS is thus that A now only owns the part of the transit market shares for p which stem from its own customer base. By using the scoping service of ISP B, MHS further defines that it does not want B to be used by C to reach destinations in p. As a result of this additional scoping, C only knows a path for p via E, and hence will select it as best.

With this technique, B is thus given the market shares for reaching p from its own customer base, as well as potentially from the customer base of A and D. The market share for reaching p from the customer base of C is given to E. As compared to triggering AS path prepending upon advertisement of p to C by B, the scoping approach makes sure that C will not locally override its preference for a path through ISP B.

In this example, the resulting routing system for the destination prefix p is policy compliant, and could not be achieved by MHS without the local configuration tweaks provided by scoping services. Advertising a path for p to A would indeed prevent MHS from shifting its entire incoming traffic for p along its links with B and E. Letting B advertise a path for p to C would potentially leave traffic coming in from C along the link with B, instead of E.

We can observe that in this scenario, the scoping being performed does not prevent any AS from receiving a path towards p. All ASes in this network forward traffic in the data-plane according to existing routing entries for p, which have been created as a result of a policy-compliant path propagation process. For this reason, any BGP routing system, constrained by scoping services, which ensures that all networks know a path for a more specific prefix being advertised, will be policy compliant.

2) Out of Policy Business Shifts: The scoping being performed on a more specific prefix might no longer let routing information for that specific prefix be spread to all ASes of the routing system. In such cases, some ASes will route traffic falling into the range of the more specific prefix, p, according to the routing information obtained for the larger range covering it, P.

As a result, the forwarding paths being followed by the concerned traffic will be made of one part of a path towards P, followed by a path towards p. The AS being at the junction of these two parts may see its policy violated; traffic received over the ingress link, due to the advertisement of P, may be not supposed to be forwarded over the egress link selected for sending traffic destined to p.

The conditions for those violations to happen are as follows:

1) The victim ISP only receives paths towards the more specific prefix, p, from peers or transit providers.
2) The victim ISP has a customer path towards the covering prefix, P.
3) A subset of the victim ISP peers and providers did not receive a path for p, and have selected a path through the victim ISP as best for P.

When these conditions are met, at least one provider or peer of the victim ISP sends traffic destined to p towards the victim ISP, according to its routing entry for P. The victim ISP itself uses peers or transit providers to forward that traffic further on, resulting in a violation of its policies.

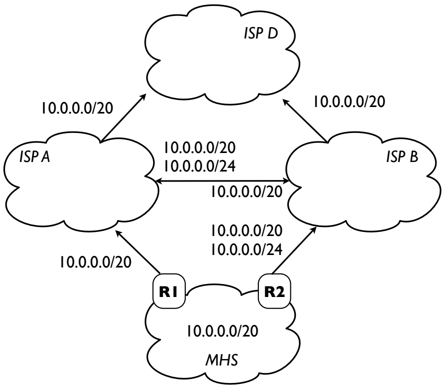

Fig. 2. Policy Violation at ISP A

Let us illustrate such a scenario, in Figure 2. MHS advertises $P = 10.0.0.0/20$ to both of its providers, ISP A and B. Let us assume that some routers in ISP D, a provider of A and B, prefer the path via ISP A for prefix P. MHS scopes the advertisement of a path to $p = 10.0.0.0/24$, so that it is only propagated to B, and uses the scoping service of B to let B only propagate that path to ISP A. From a control plane point of view, all the paths being propagated and selected as best are compliant with each individual ISP's policy. However, from a data-plane point of view, the policy of A is being violated. Traffic destined to p, originated by ISP D, will be forwarded from a provider link towards a peer link by A. Indeed, BGP routers at D are not aware of a path for p and will thus keep on forwarding any packet falling in the p range according to their routing entry for P, via A.

III. Policy violations, Properties and Incentives

In this section, we study properties of the scoping service provided by ISPs. We first observe that the networks playing a role in the policy violation commonly have direct business relationships. Next, we focus on the particular case of prefix scoping among the Tier-1 clique. We then discuss the potentiality of policy violation due to scoping services provided by lower tier networks, as a function of the flexibility of their offered service. Finally, based on these observations, we discuss the incentives for ISPs to perform inbound traffic engineering that will lead to policy violations.

A. Properties of Scoping Services

1) Relationships Between Affected ISPs: Scoping services are usually restricted by their providers. The set of ASes that can take advantage of the service is commonly limited to the set of customers of the ISP. Practically, this limitation is implemented by stripping off the communities aimed at using its scoping services, from the paths that it receives over its peering and provider links. The scoping provider thus usually belongs to the provider cone of the owner of P.

The set of neighboring ASes to which a prefix can be prevented from being propagated is typically strictly defined. In [2], the set of neighboring ASes to which scoping can be performed is listed, and customers of Sprint can perform path scoping on a per-peer basis. In [3], the customers of Cogent can define the scoping behavior of Cogent BGP speakers, on a per-region basis.

As the policy violation always occurs by having peers and providers of the victim forward traffic to it according to P, it is necessary that the victim's routing policy lets P be advertised to these. Hence, P must be a customer path for the victim. Consequently, the victim always lies in the provider cone of the owner of P.

As the policy violation takes place by having the scoping provider advertise p, a customer prefix, towards the victim and not to some other providers and peers, the scoping provider most often has a direct provider-customer or a peering relationship with the victim.

We thus observe that there always exists a direct or indirect business relationship between the originator of the concerned destination, the ISP providing the scoping service to the originator, and the victim of the policy violation. Such an observation plays an important role in the analysis of the incentives for violating policies, as discussed later in this section, as well as in the analysis of the solution space to handle these violations, as discussed in Section IV.

2) The Tier-1 Case: Tier-1 networks usually provide scoping services, letting their customers prevent their paths from being propagated to some of the other Tier-1s of the clique. Here, we capture the properties of Tier-1 level scoping service which lead to policy violations at some other Tier-1 of the clique. We use Figure 3 to illustrate these properties. In this example, All ISPs but MHS are forming a Tier-1 clique.

Let us define $\mathcal{T}$, the set of Tier-1 networks. They form a clique of established peering links and exchange their customer routes over these. $\mathcal{T}_c^P$ denotes the set of Tier-1 networks which receive a path for P from their customer base. In our example, this set is made of A, B, and E. $\mathcal{T}_c^p$ denotes the set of Tier-1 networks which receive a path for p from their customers. In our example, this set is made of B and E.

$\mathcal{V} = \mathcal{T}_c^P \backslash \mathcal{T}_c^p$ forms the set of potential victims of the scoping. This set is non empty when scoping is performed

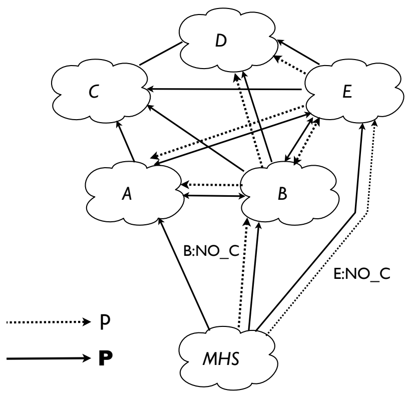

Fig. 3. Tier-1 Policy Violations

by the customer, resulting in p being selectively announced to a strict subset of his T1 providers. In our example, A belongs to $\mathcal{V}$. $\mathcal{S} = \mathcal{T}_c^P \bigcap \mathcal{T}_c^p$ forms the set of Tier-1 whose scoping services might violate the policies of the set of ISPs belonging to $\mathcal{V}$. When their scoping service is used to prevent paths from being propagated to some Tier-1 which neither belong to $\mathcal{A}$ nor V, then a violation can occur at ASes belonging to $\mathcal{V}$.

In our example, $\mathcal{S}$ is made of B and E. C receives the path for p from neither B nor E, and hence may forward traffic destined to p towards A according to its routing entry for P, resulting in a policy violation at ISP A.

From this analysis, it can be concluded that policy violations at the Tier-1 level are very likely to happen as soon as one is using scoping services for a more specific prefix, while letting the prefix covering it advertised without scoping. For an ISP that is dual-homed to two Tier-1 networks, using the scoping service of one of its providers for an advertised more specific prefix, preventing a third Tier-1 from receiving a path for that prefix, leads to policy violation at the other Tier-1 provider. That is, using more specific prefix scoping at the Tier-1 level is very likely to trigger policy violations.

3) Properties of Scoping Services Provided by Non T1s: If the scoping service offered by a non Tier-1 ISP enforces the propagation of its customer paths to its providers, policy violations will commonly not occur.

Indeed, for condition 3 in Section II-B.2 to be met in this case, it would be necessary to not have a policy compliant sequence of ISPs from the providers and peers of the victim towards the scoping provider. This does not happen in a routing system that is resilient against the failure of the links between the victim and the scoping provider, thanks to transit connectivity. Such lack of resilience is unlikely to take place w.r.t. the Tier-1 connectivity.

As soon as the scoping service of a lower tier ISP can be used to not let it propagate paths to its providers, policy violations may occur. Such complete services are commonly offered by lower Tier ISPs, e.g. by [4], and hence policy violations should not be considered as a problem that only Tier-1 would face.

B. Incentives for Policy Violations

Let us discuss the incentives for individual or a group of ISPs to abuse scoping of more specifics. Typically, the victim, ISP A, is a global Tier-1 network which has better reach than ISP B, a regional challenger. MHS is connected to both and pays advantageous tariffs to ISP B. MHS also reckons that the reach of ISP B to some networks out of the region is poor, hence puts the violation in place striving to get advantage of ISP A global scope at the ISP B rates.

Practical load balancing issues may also push the MHS to trigger the violation. Given the wider customer and peering cone of ISP A, MHS may struggle to effectively load-balance inbound traffic among its upstream providers. It then resorts to violating ISP A via ISP B with more specific routes for a subset of its total address space, still keeping the transit reach of the victim under use.

The abuse favors ISP B for peering traffic ratio accounting, as the policy violating state will increase the amount of traffic that ISP A sends over its peering link with ISP B.

Once the MHS and the scoping service provider share business interests, the incentives become trivial. MHS indeed moves transit costs payed to a third party to the "internal" cost structure of the MHS and ISP B, while ISP B saves on transit costs.

IV. Defending against out of policy cases

An ISP may want to detect two concerning situations. The first, most obvious, one is the case when the ISP is being exploited with the use of scoping. The second case of interest is when the ISP is unwantedly responsible of policy violations at its peers and providers network.

A. Defending Against Policy Violations

We analyze two different approaches for the operator of an ISP to defend against policy violations that could happen in his network. The first approach simply consists in giving oneself the ability to detect the occurrence of a policy violation, and react a posteriori to have the problem be solved by the culprit of the violation, through non technical means. The second approach that we investigate consists in enforcing the respect of policies by technical means. Here, we discuss both approaches and advocate for the use of the former.

1) Detecting violations: Detecting policy violations in its own network can be done by looking at internal BGP data to see whether there exists in the RIB a prefix p more specific than P such that the nexthop for p is through a peer (or a provider) while P is routed through a customer. Only running this check is unfortunately tainted with false positives. For example, in Figure 1, at ISP A, this check would trigger a warning upon a selective advertisement of p by MHS to B, even if B would not scope the propagation of p. Although A loses traffic shares in this situation, it does not have its policies violated, so that the check would wrongly raise a warning.

These checks miss the third condition for a policy violation to happen, as described in Section II-B. That is, it does not capture whether providers and peers of the ISP are actually using A to reach destinations within p (by using their route for P). In order to make the detection accurate, looking glasses around the providers and peers of the ISP can be consulted to see if some ASes are propagating a path towards P through A, and no paths towards p, to their own customers, peers, or providers. This would mean that ASes in the surrounding area of the current AS are forwarding packets based on the routing entry for the less specific prefix only, across A.

Relying on the availability of looking glasses in the surrounding area of the ISP is however not always practical for an ISP. An alternative approach is to use telemetry export protocols (NetFlow, IPFIX [5], [6]) to check whether traffic to destinations covered by p actually flows in along provider and peering links of ISP A. This would be done for all p such that there exists a couple (p, P) satisfying the conditions expressed above.

While this would work, separation between telemetry and routing data make the approach rather expensive from a computational point of view: requirement is in fact to save micro-flows to persistent storage for a-posteriori lookup against an external BGP RIB, and leveraging expensive IPCs for the task. In this sense, a computational savvy method to detect such violations is to integrate telemetry and routing data in a single memory footprint, using a dedicated tool.

We suggest to use *pmacct* [7], as it allows for spatial as well as temporal aggregation, and grouping of micro-flows on persistent storage. It leverages direct memory access and lightweight inter-thread IPCs. *pmacct* is already used by many ISPs to perform essential network management tasks of IP accounting such as billing, graphing network resources usage, analyzing live or historical traffic trends, steering BGP peerings, triggering real-time alerts, and monitoring some types of SLAs.

To provide such features, *pmacct* integrates NetFlow and sFlow data collection features as well as a BGP daemon aimed at collecting control plane information. For operators who have already *pmacct* deployed, detecting policy violations will thus require very few additional configuration or scripting.

We suggest to use local preference information found in BGP paths received by its BGP daemon [8], or communities tagged to these paths by the operator, to identify to which type of egress link a BGP path leads. *pmacct* also allows to tag each ingress link, either physical or logical, with the type of peering relationship it serves. It thus becomes simple to check against telemetry aggregated data to see whether any flow is transiting from any peer or provider to any other peer or provider. A tutorial on how to use *pmacct* to detect policy violations can be found at [9].

2) Preventing violations: An operator can technically ensure that the traffic destined to a given prefix will be forwarded from an entry point of the AS, only on the basis of the set of paths that have been advertised over that entry point [10]. Under such "Neighbor-Specific BGP" deployment, an ISP can no longer find itself at the junction between a prefix of a data-plane path for P, followed by a data-plane path for p. The traffic destined to p, subject to the violation, would thus not be forwarded to ISP B by ISP A, but directly to MHS, according to the routing state of P.

Such a transit network deployment can however be considered as radically different from usual Internet transit service deployment approaches. Operators deploying Neighbor-Specific BGP for the flexibility of the services that it provides will defend against policy violations as a side gain, but it is unlikely that such deployment would be performed for the sake of policy violation avoidance.

An alternative approach is to trigger a similar behavior with control plane actions. The more specific prefixes that are subject to policy violation can indeed be filtered out at the peering session over which they are received. In the example of Figure 2, A would filter out $10.0.0.0/24$ in its eBGP in-filter associated with the eBGP session with B. As a result, the traffic destined to that /24 would be forwarded by A along its link with MHS, despite the actions performed by MHS to have this traffic coming in through its link with B.

The operator relying on such counter-measures will still have to discover the set of prefixes to be filtered out at its borders. A complementary tool aimed at detecting these violations is thus still necessary for this solution. Moreover, the maintenance of the prefix-list to be filtered out might be considered overwhelming by some operators, and considered less attractive even if a degree of automation is possible.

A third approach is to automate BGP router configurations so as to have them dynamically install an access-list made of the prefixes towards which the forwarding of traffic from that interface would lead to a policy violation. In the example of Figure 2, ISP A would install an access-list denying packets matching 10.0.0.0/24, for the interfaces connecting its providers and peers. As a result, the traffic destined to that /24 would be dropped, despite the existence of a non policy-violating route towards 10.0.0.0/20.

Note that these techniques either let policy violating packets be dropped, or forwarded according to a biased routing state that does not follow the routing tweaks performed by customers of the ISP putting them in place. As these may be not aware of the policy violation that they are triggering, automating such a behavior might be considered as a too aggressive measure.

3) Detection or prevention ?: As shown in Section III, the actors of the policy violation usually have established business relationships. Depending on the business relevance of the amount of violating traffic, operators will thus typically react to such events in a smooth way. It should thus be recommended not to deploy automated solutions that would black-hole traffic deemed at violating the policies of the ISP. In case a policy violation takes place, the scoping service provider and the MHS will be advised of the issue and asked to take action. In case the violation causes urgent issues such as link saturation, the victim will typically try to filter specific routes without breaking connectivity. Upon reiteration of the issue or when not observing co-operation by the involved parties, another possible scenario is the victim resorting to black-holing or requesting compensation for the violation. The latter being justified by the fact that the victim has been acting as a provider of the scoping service provider for the duration of the violation.

B. Preventing oneself to be the culprit

In order to provide a prefix scoping service that will not lead to policy violations in distant ASes, an ISP should make sure that it does not allow for its scoping service to lead to cases satisfying the three conditions defined in Section II, for some distant AS.

Preventing conditions 1 and 2 from being satisfied at the same time at a distant AS would basically render the scoping service useless. Indeed, many such cases reflect perfectly valid, policy-compliant routing states, that customers are actually looking for through the use of that scoping service. The scoping service provider must thus take care, when selectively propagating a path for p to a given peer or provider, that all the peers and the provider branch of that ISP will also receive a path for p, so that condition 3 is not satisfied at the same time as conditions 1 and 2.

Constraining the scoping service to some selected combinations of neighboring ASes is however not sufficient to ensure such a property. For example, in Figure 3, we can observe that ISP B is getting benefits of the policy violation at ISP A. However, the use of the scoping service of B by MHS is in itself perfectly legitimate. The policy violation indeed takes place because MHS also used the scoping service of ISP E in order for C to not be advertised of the path towards p. Such a behavior is not under the control of ISP B.

As explained in Section III, a non Tier-1 network which enforces the propagation of its customer paths to all its providers will typically not be the culprit of a violation. However, it renders the scoping service less efficient for its customers. Note that if the ISP does not strip off the communities aimed at scoping services at its own providers, a policy violation can take place higher in the routing hierarchy, independent of the restrictions it sets on the use of its own scoping service.

These considerations illustrate that providing a scoping service that is not too restrictive, while not allowing policy violation at distant ISPs is not an easy engineering process. It thus appears to be not practical for an ISP to restrict its scoping service for the sake of defending its surrounding ISPs policies. In the first place, ISPs providing path scoping services should recommend their customers to not use path scoping as the first mean to meet their inbound traffic share distribution goals. Indeed, in many cases, triggering AS Path prepending instead of a "no export" behavior is sufficient to meet these requirements. These remotely triggered prepending actions cannot lead to limited visibility of the paths that they affect, and hence cannot lead to policy violations. Path scoping should thus be presented as a "last resort" operation, to be used when routing at the targeted ISPs is not influenced by the provided AS Path prepending service.

As for the detection of policy violations in its own network, an ISP could also follow a reactive approach, i.e., be able to detect policy violations consequent to the use of its scoping service, a posteriori. Using only locally available BGP control plane data (the set of paths received at the borders of its AS) is not sufficient for an ISP to *detect* such a violation. Using the same example of Figure 3, one can see that the set of paths received by ISP B from its neighboring ASes would be the same if MHS was not using the scoping service of ISP E and hence would not trigger the violation. The operator would thus need to rely on BGP looking glasses to analyze the effect of the use of its scoping service on distant ASes.

V. Conclusion

In this paper, we have shown that some of the BGP routing tweaks and services currently offered and used by ISPs can lead to policy violations, when used together. We have discussed the incentives for ISPs to actually trigger these violations, and have shown that simple prefix-scoping usage on more specific prefixes can already lead to policy violations, notably among the ISPs forming the Tier-1 clique.

We discussed the various approaches for dealing with such violations, and concluded that detection and a posteriori reaction is the approach to be followed. To detect such violations in a lightweight fashion, we recommend ISPs to rely on their already deployed IP management software, and use them to perform the appropriate checks on their transit behavior. We provided means to perform such a detection in $pmacct$, an open-source IP management tool.

The policy violations depicted in this paper are all taking place in the context of usual, well-known, ISP transit policies following "Customer-Provider" and Peering links. As a further work, we will extend our analysis to the cases where some ISPs deviate from this simple model.

Acknowledgements

Pierre Francois is partially funded by the "Fonds national de la recherche scientifique".

References

[1] B. Donnet and O. Bonaventure, "On BGP Communities," *ACM SIGCOMM Computer Communication Review*, vol. 38, no. 2, pp. 55–59, April 2008.

[2] "Sprintlink's BGP Policy." [Online]. Available: https://www.sprint.net/index.php?p=policy_bgp

[3] "Cogent Communications User Guide." [Online]. Available: http://www.cogentco.com/us/cs_faq.php

[4] "Easynet BGP Policy Description." [Online]. Available: http://www.onesc.net/communities/as4589/

[5] B. Claise, "Cisco Systems NetFlow Services Export Version 9," RFC 3954 (Informational), Internet Engineering Task Force, Oct. 2004. [Online]. Available: http://www.ietf.org/rfc/rfc3954.txt

[6] ——, "Specification of the IP Flow Information Export (IPFIX) Protocol for the Exchange of IP Traffic Flow Information," RFC 5101 (Proposed Standard), Internet Engineering Task Force, Jan. 2008. [Online]. Available: http://www.ietf.org/rfc/rfc5101.txt

[7] "The pmacct Project." [Online]. Available: http://www.pmacct.net/

[8] "Building traffic matrices to support peering decisions." [Online]. Available: www.pmacct.net/building_traffic_matrices_n49.pdf

[9] "Detecting routing policy violations." [Online]. Available: http://wiki.pmacct.net/DetectingRoutingViolations

[10] Y. Wang, M. Schapira, and J. Rexford, "Neighbor-specific BGP: more flexible routing policies while improving global stability," in *Proceedings of the eleventh international joint conference on Measurement and modeling of computer systems*, ser. SIGMETRICS '09. New York, NY, USA: ACM, 2009, pp. 217–228. [Online]. Available: http://doi.acm.org/10.1145/1555349.1555375