# EAGER:Experimental Deployment of the ARTEMIS BGP Hijacking Detection Prototype in Research and Educational Networks

## 1 Introduction, Motivation and Goals

ARTEMIS (Automatic and Real-Time dEtection and MItigation System)[**artemis-site**], is a research effort between the INSPIRE group, FORTH, Greece [**inspire-forth**] and the Center for Applied Internet Data Analysis (CAIDA), University of California San Diego, USA[9]. ARTEMIS studies detection and mitigation techniques for BGP prefix hijacking focusing on the following novel aspects: a) the real-time monitoring of the inter-domain routing control plane using recently introduced streaming services by RIPE RIS and other data providers; b) the accurate detection of regular and advanced types of BGP prefix hijacking attacks, *i.e.,* man-in-the middle attacks, which can be very stealthy to detect; c) the evaluation of mitigation techniques that can be used to automatically resolve hijacking incidents within few seconds.

Based on results of our research study, we recently proposed an automated detection and mitigation approach that we describe in detail in a scientific paper [**artemis-paper**]. The ARTEMIS approach is entirely based on analysis of BGP data (*i.e.,* no data-plane measurements are necessary) and allows an operator not having to rely on third party services (with consequent loss of privacy, lack of detection accuracy, and practical deployment issues) for the detection of BGP hijacking attacks. Instead, it allows an operator to autonomously instantiate defense mechanisms that are more reliable, accurate and fast. The ARTEMIS approach relies on two key observations: *(i)* today's public BGP monitoring infrastructure (such as RouteViews [21] and RIPE RIS [20]) is much more advanced than when previous solutions for BGP hijacking detection were proposed, making it a valuable resource for comprehensive, live monitoring of the Internet control plane that is available to anybody; *(ii)* shifting from a third-party perspective to a self-operated approach enables us to effectively address the long-standing and persistent issues undermining the state of the art in BGP hijacking defense.

In the past months we presented ARTEMIS in various popular network operator meetings hosted by RIPE, IETF, Internet2, etc.[**artemis-ripe**, **artemis-ietf**, **artemis-i2**], and we registered significant interest in our approach. The feedback we received helped us shaping a first prototype code that implements the main techniques of the ARTEMIS approach. However, operators feedback also highlighted challenges related to our approach that are secondary in terms of conceptual detection and mitigation methodologies but crucial when trying to deploy it in a real operational network: *e.g.,* how to perform automated configuration as the operator makes changes to its BGP policies and announced routes in different points of its network? How to perform route de-aggregation or alter BGP communities in already de-aggregated routes that are normally distributed only within the operator's routing domain. Finally, several operators from research and educational networks showed their interest in collaborating to test our approach and help us identify and tackle operational challenges.

In this project, CAIDA proposes to carry out a first experimental assisted deployment in few selected operator networks supporting education, research and other no-profit activities in the US, such as Internet2 [**internet2**], MERIT[**merit**], and Great Plains Network[**gpn**]. The main goals of such experimentation are: *(i)* to evaluate the ARTEMIS methods and their implementation in real operational scenarios; *(ii)* to learn and address real deployment technical issues; *(iii)* to train operators in the use of the tool and contribute to improve the security of their networks. The duration of the project will be one year.

## 2 Background and State of the Art

Autonomous Systems (ASes) use the Border Gateway Protocol (BGP) [12] to advertise their IP prefixes and establish inter-domain routes in the Internet. BGP is a distributed protocol, lacking authentication of routes. As a result, an AS is able to advertise illegitimate routes for IP prefixes it does not own. These illegitimate 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* can be caused by router misconfiguration [1, 2] or malicious attacks [3, 19, 24]. Events with significant impact are frequently observed [4, 5, 6, 24], highlighting – despite the severity of such Internet infrastructural vulnerability – the ineffectiveness of existing countermeasures. For example, on the 6th of November 2017, millions of Comcast subscribers in the US could not use the Internet for approximately 90 minutes because of a prefix hijacking incident.

Currently, networks rely on *practical reactive mechanisms* to try to defend against prefix hijacking, since proposed *proactive* mechanisms [15, 23, 17, 18, 14] (*e.g.,* RPKI) are fully efficient only when globally deployed, and operators are reluctant to deploy them due to associated technical and financial costs [**lychev-2013-sigcomm**, **cooper-2013-hotnets**, **matsumoto-2017-privsec**]. Defending against hijacking reactively consists of two steps: *detection* and *mitigation*. Detection is mainly provided by third-party services, *e.g.,* [7], that notify networks about suspicious events involving their prefixes, based on routing information (such as traceroutes [25] or BGP updates [7]). The affected networks then proceed to mitigate the event, *e.g.,* by announcing more specific prefixes, or contacting other ASes to filter announcements.

However, due to a mix of technological and practical deployability issues, current reactive approaches are largely inadequate. Specifically, the state of the art suffers from 4 main problems, which the ARTEMIS approach instead addresses effectively:

- **Evasion**. None of the detection approaches in literature is capable of detecting all attack configurations (nor can they be easily combined), thus allowing sophisticated attackers to evade them. We propose a modular taxonomy describing all variations of attack scenarios and we use it to carefully analyze detection comprehensiveness of related work. ARTEMIS significantly overcomes limitations of the state of the art by covering all attack configurations in the context of a classic threat model.

- **Accuracy**. Legitimate changes in the routing policies of a network (*e.g.,* announcing a sub-prefix for traffic engineering or establishing a new peering connection), could be considered suspicious events by the majority of third-party detection systems [10, 16, 25, 22, 13]. To avoid this, operators would need to timely inform third parties about every routing decision they make and share private information. On the other hand, adopting a less strict policy to compensate for the lack of updated information and reduce false positives (FP), incurs the danger of neglecting real hijacking events (false negatives – FN). We designed ARTEMIS detection to be run directly by the network operator without relying on a third party, thus leveraging fully and constantly (and potentially automatically) updated information that enables 0% FP and FN for most of the attack scenarios and a configurable FP–FN trade-off otherwise.

- **Speed**. A side effect of the inaccuracy of third-party approaches is the need for manual verification of alerts, which inevitably causes slow mitigation of malicious events (*e.g.,* hours or days). Few minutes of diverted traffic can cause large financial losses due to service unavailability or security breaches. On the contrary, the ARTEMIS approach is a fully automated solution integrating detection and mitigation, allowing an AS to quickly neutralize attacks.

- **Privacy and Flexibility**. One of the issues that impedes the adoption of third-party detection is privacy, *e.g.,* ISPs usually do not disclose their peering policies. Similarly, operators are sometimes reluctant to adopt mitigation services requiring other organizations to announce their prefixes or tunnel their traffic. ARTEMIS offers full privacy for detection and the option to achieve self-operated mitigation. Another factor affecting willingness to externalize mitigation is cost. Trade-offs between cost, privacy, and risk may be evaluated differently by the same organization for distinct prefixes they own. Leveraging the availability of local private information and its fully automated approach, ARTEMIS offers the flexibility to customize mitigation (*e.g.,* self-operated or third-party-assisted) per prefix and per attack class.

## 3 Proposed Tasks

We have already identified few candidate networks that support research, education and no-profit activities in the US and we will select at least two participants to collaborate with in the course of the project. Operators that have already expressed their interest and commitment to collaborate include Internet2 [**internet2**], Great Plains Networks [**gpn**] and MERIT Networks [**merit**]. To achieve the goals highlighted in Section 1, we will organize our collaboration with each operator into a series of tasks that might be iterated at different stages as well as carried out in partial overlap, depending on the lessons learned during the project execution, the challenges that arise, and the potential solutions we identify.

First of all, we will work with each operator to provision the suitable hardware, or discuss solutions based on virtual machines running in cloud environments, where to deploy the ARTEMIS prototype software. We will start with deploying a default "test" configuration with a manually-created configuration file to monitor only a sample set of prefixes of the operator.

The next step is to collaborate with the operator to discuss their current internal procedures to maintain the inventory of announced address blocks, originating Autonomous Systems, BGP peering sessions (and therefore neighboring ASes), etc. We received preliminary feedback from different network administrators about the potential lack of complete or up-to-date BGP routes inventory in large networks. Working jointly with the operator, we will identify strategies to extract and merge information from their routers, route servers and route reflectors. We will then proceed to develop proof-of-concept software (*e.g.,* based on exabgp[11] or CAIDA BGPStream [8] through the BGP Monitoring Protocol [**bmp**]) to automatically and continuously extract such information and synthetize it into the ARTEMIS configuration file format. In parallel, we will discuss and implement strategies to merge configuration data automatically generated with the pre-existing configuration and potentially with additional input from a human operator. Another approach that we will evaluate, is performing the discovery of the visible (*i.e.,* based on the current status of the network) inventory of BGP routes, origins and neighbors by using data from RouteViews and RIPE RIS collectors and similar types of vantage points in order to assist the operator and identify approaches to auto-generate a basic ARTEMIS configuration. We will discuss the pros and cons of all the possible solutions with the operator and experiment with several of them. This step might encourage or require the operator to introduce improvements to their network management procedures and policies for what concerns BGP operations.

Once we will have deployed full monitoring of the operator's prefixes, we will initially configure the ARTEMIS prototype to alert for any type and confidence level of suspicious events and we will start collecting alert data (subsequently we will re-iterate this task tuning the sensitivity of our detection algorithms). This is one of the crucial activities in this project, since with the operator we will jointly investigate and discuss the alerts generated by ARTEMIS and how its configuration can be tailored to satisfy their expectations and requirements. We will learn about potential false positives and false negatives, about different levels of priority and preference of alerting and mitigation of the operator based on both the type of event and the alert information generated by ARTEMIS. We will also investigate these events in detail to ascertain their nature. For example, we might manually execute (or automatically trigger) traceroute measurements to examine routes at the data plane.

Mitigation of a BGP hijacking attack is the final goal of a complete defense approach. There are two classes of problems that we will discuss with operators to learn about their needs and their requirements and preferences in terms of privacy, speed of execution vs accuracy, etc. The first one is related to how they envision in an ideal case a suitable strategy of mitigation for different categories of address blocks (and therefore BGP prefixes) based on their relevance and role within the organization, and consequently discuss the trade-offs involved in realistic scenarios. The second class of problems is related to technical aspects and deployment issues: *e.g.,* are certain prefixes already de-aggregated within the operator's routing domain but announced externally in aggregated form because of BGP community tags associated with them? What are possible technical means - that are compatible with the organization security and management policies - to interface with their routers and force prefix de-aggregation?

While experimenting with the detection of events on actual operational prefixes is an activity that does not alter the behavior of the operator's network in terms of BGP announcement, mitigation (*i.e.,* prefix de-aggregation) can have a visible impact on BGP operations. We will limit our collaboration with an operator with respect to BGP hijacking mitigation to the discussion and analysis of all the technical, administrative, and management aspects mentioned above for operators that prefer to postpone experimenting with automated measures (*e.g.,* because of concerns with their legal department). With all the others we will collaborate to experiment with such capabilities at various levels. As a first step, we will experiment with only one non-operational prefix made available by the operator for test purposes. We will then use external research infrastructure deployed at different locations in the Internet topology and worldwide, such as the PEERING testbed [**peering**] to hijack the monitored prefix and first ensure that monitoring works correctly and then testing automated de-aggregation based on the preliminary analysis and agreements with the operator (*e.g.,* removing/altering BGP communities, or de-aggregate prefixes at the border routers). Besides testing and tuning the correct execution of mitigation, these set of experiments will provide precious data informing us on how effective ARTEMIS can be in neutralizing attacks for the specific network under exam. We will analyze these results in collaboration with the operator and architect more evaluations and potential code updates. We will then consider, in agreement with the operator, to extend automated mitigation to operational prefixes announced by their AS and continue experimentation based on real anomalies observed in the wild.

Finally, we will work closely with the operator's technical staff to teach them effective use of our prototype and to collect feedback that will influence new code iterations. We will discuss functionalities for visual interfaces and usability requirements and we will evaluate the potential integration with existing platforms used by the operator.

## 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] https://arstechnica.com/security/2017/04/russian-controlled-telecom-hijacks-financial-services-internet-traffic/.

[5] http://dyn.com/blog/iran-leaks-censorship-via-bgp-hijacks/.

[6] http://seclists.org/nanog/2016/Sep/122.

[7] *BGPmon (commercial)*. http://www.bgpmon.net.

[8] *BGPStream*. https://bgpstream.caida.org/.

[9] CAIDA. http://www.caida.org/.

[10] Ying-Ju Chi, Ricardo Oliveira, and Lixia Zhang. "Cyclops: the AS-level connectivity observatory". In: *ACM SIGCOMM Computer Communication Review* 38.5 (2008), pp. 5–16.

[11] Exa-Networks. *exabgp: The BGP swiss army knife of networking*. https://github.com/Exa-Networks/exabgp.

[12] Susan Hares, Yakov Rekhter, and Tony Li. *A border gateway protocol 4 (BGP-4)*. https://tools.ietf.org/html/rfc4271. 2006.

[13] Xin Hu and Z Morley Mao. "Accurate real-time identification of IP prefix hijacking". In: *IEEE Symposium on Security and Privacy*. 2007, pp. 3–17.

[14] Josh Karlin, Stephanie Forrest, and Jennifer Rexford. "Pretty good BGP: Improving BGP by cautiously adopting routes". In: *Proc. IEEE ICNP*. 2006.

[15] Stephen Kent, Charles Lynn, and Karen Seo. "Secure border gateway protocol (S-BGP)". In: *IEEE Journal on Selected Areas in Communications* 18.4 (2000), pp. 582–592.

[16] Mohit Lad et al. "PHAS: A Prefix Hijack Alert System." In: *Usenix Security*. 2006.

[17] Matt Lepinski. "BGPSEC protocol specification". In: (2015).

[18] Matt Lepinski, Richard Barnes, and Stephen Kent. "An infrastructure to support secure internet routing". In: (2012).

[19] Anirudh Ramachandran and Nick Feamster. "Understanding the network-level behavior of spammers". In: *ACM SIGCOMM Computer Communication Review* 36.4 (2006), pp. 291–302.

[20] *RIPE RIS - Streaming Service*. https://labs.ripe.net/Members/colin_petrie/updates-to-the-ripe-ncc-routing-information-service.

[21] *University of Oregon Route Views Project*. http://www.routeviews.org/.

[22] Xingang Shi et al. "Detecting prefix hijackings in the Internet with Argus". In: *Proc. ACM IMC*. 2012.

[23] Lakshminarayanan Subramanian et al. "Listen and whisper: Security mechanisms for BGP". In: *Proc. NSDI*. 2004.

[24] Pierre-Antoine Vervier, Olivier Thonnard, and Marc Dacier. "Mind Your Blocks: On the Stealthiness of Malicious BGP Hijacks." In: *Proc. NDSS*. 2015.

[25] Changxi Zheng et al. “A light-weight distributed scheme for detecting IP prefix hijacks in real-time”. In: *ACM SIGCOMM Computer Communication Review*. Vol. 37. 4. 2007, pp. 277–288.