The Inherent Vulnerability of the Border Gateway Protocol (BGP)

Denis Richard
Denis.richard@tufts.edu
Ming Chow

Abstract

At DEFCON 16, Anton Kapela and Alex Pilosov demonstrated a technique to stealthily steal Internet traffic along the lines of a MITM attack. The attack utilizes an inherent feature of the Border Gateway Protocol (BGP), fooling routers into redirecting traffic to an eavesdropper's network. Because it does not exploit a bug in the protocol, but exploits the natural way BGP works, such attacks are hard to detect and would require an overhaul of the protocol to prevent them. Secure BGP (S-BGP) is the name of such an overhaul and is designed to verify an AS's authenticity through a hierarchical process of authentication by the succeeding ASes in a route's path. This paper outlines the inherent flaw of the BGP and analyzes the problems faced by the implementation of the revised S-BGP protocol.

To the community

It is clear that the implementation of the Internet's currently used exterior routing dates back to a time when the number of autonomous networks was substantially smaller and could rely on a system of trust more realistically. While this assumption allowed the use inferior routing infrastructure, less bandwidth, and more flexibility, the increasing growth of networks composing the Internet and the benefits of the latter's exploitation forces us to revise this assumption in order to guarantee the safety and reliability of its traffic control. For years, a group of individuals emphasized that BGP is inherently vulnerable to the propagation of illegitimate routes that are hard to authenticate, but the lack of concrete examples of such exploitations allowed the issue of this critical revision to go by largely unaddressed. However, recent examples of black holes, large-scale traffic redirects, and interceptions, stress the importance of understanding the implications of using this ubiquitous protocol, and furthermore the gravity of revising it in order to enable a secure means of authenticating the routes it propagates. The only thing that can force Autonomous Systems to fix this issue is for their customers to demand security solutions. This paper attempts to increase the awareness about the severity of this issue on the part of the latter.

Introduction

Networks are largely composed of intermediate or end systems, referred to as routers or hosts. Internet routing is based on a distributed system composed of many networks that are grouped into individual management domains called Autonomous Systems (ASes). There is no central "core" to the Internet and in order to communicate routing information between these disparate parts that may be using different internal routing protocols, ASes make use of the Border Gateway Protocol (BGP), which is currently the most widely used exterior routing protocol because it allows for fully decentralized routing. In such a design, the interior protocol provides optimum paths through the individual ASes, while BGP provides edge-to edge routing across them.

The Internet consists of tens of thousands of ASes that use BGP to exchange information on how to reach blocks of IP addresses. An AS itself can be seen as a collection of IP routing prefixes under the control of one or more network operators that present a common, clearly defined routing policy to the Internet. Each AS is allocated a unique Autonomous System Number (ASN). The root of this allocation is the Internet Assigned Numbers Authority (IANA), whose duties are currently performed by the Internet Corporation for Assigned Names and Numbers (ICANN). IANA then assigns large address blocks and ASNs to Regional Internet Registries (RIRs), which subdivide these and further delegate them to Local Internet Registries (LIRs). Similar subdividing may then be repeated several times at lower levels of delegation.

The address space allocation follows the concept of cluster addressing and its notation defined by Classless Inter-Domain Routing (CIDR): a syntax used for the

specification of IP addresses and their routing prefix, which appends a slash character to the address followed by the decimal number of leading bits to the routing prefix. The number of addresses of a subnet can thus be calculated as $2^{(adressSize - prefixSize)}$, where the address size is 128 for IPv6 and 32 for IPv4.

The Internet does not run without BGP, as it is the glue that holds its various parts together. BGP might not carry end user traffic, but when an end user wants to communicate to an IP address, its ISP consults a BGP table - which contains a list of known routers - for the best route to the destination. This table is built up through BGP exchanges across ASes, which declare the range of IP addresses, or prefixes, to which they are able to deliver traffic. A cost metric is then associated with the path to each router so that the best available route is chosen. BGP's main function is thus to exchange network reachability information with others in order to track the networks through which the traffic would have to pass to reach the final destination.

Among routing protocols, BGP is unique in using TCP as its transport protocol. A TCP session is manually configured between neighbors on port 179 in order to conduct the exchange. 19-byte keep-alive messages are then sent every 30 seconds to maintain this connection.

BGP is an incremental protocol. If new routes become available or existing ones are withdrawn, an UPDATE message is sent to all peers. UPDATE messages consist of a list of address prefixes that have become reachable or unavailable. BGP is a path-vector protocol, so in addition the lists, the message also contains the current inter-AS hop and characteristics of the cumulative path that can be used to reach these prefixes. UPDATE messages are then stored and processed in the router's Routing Information Bases (RIBs), which are separated into three conceptual parts. The Adj-RIB-In stores the input received from other BGP speakers, which is then processed

into the Loc-RIB according to local policies, where the routes actually used by the local speaker are stored. The Adj-RIB-Out stores the routes and information a speaker wants to advertise to its peers. In summary, the unprocessed routing information sits in the Adj-RIB-In, the Loc-RIB contains the routes selected by the local BGP speaker's decision process, and the Adj-RIB-Out table then organizes these for advertisement to specific peers.

The conceptual separation of these tables implies the importance of filtering in BGP exchanges, which takes place for a number of economic, strategic and security reasons that are at the discretion of the individual ASes. The input information stored in the Loc RIB-In is sufficient for constructing an AS connectivity graph. In accordance to an AS's administrative implementation of the filtering process, learned routes can then be dropped, modified and redistributed via route-map mechanisms. Typically, these rules will differ for peers that are interior or exterior to the AS, but if two systems deliver traffic for a given address, the one with the narrower prefix is usually preferred. In addition, the decision process will almost always prefer the route with the shortest AS-path, which is the set of ASes that must be traversed to reach the given destination.

The current version of BGP makes the implementation of a filtering process even more important due to a number of vulnerabilities that can cause problems such as mis- or non-delivery of traffic, network resource misuse and congestion, which are most often unintentional. BGP's correct operation depends upon the "integrity, authenticity, and timeliness of the routing information it distributes as well as each BGP speaker's processing, storing and distribution of this information in accordance with both the BGP specification and with the local routing policies of the BGP speaker's AS" (Kent). In order to comply, each UPDATE must be sent and received

by the intended peers, the address space advertisement must be authorized on behalf of the sender's parent organization, and remain unmodified en route.

BGP was largely developed in the 1970s with trustworthy assumptions on behalf of the then-nascent network and its "inadequacies […] have been obvious since a time shortly after being drawn up on a napkin" (Dugan). Thus, the limited guarantees provided by BGP now contribute to serious instabilities and outages because the implementations of its necessary procedures remain at the discretion of the ASes' router administration, which is vulnerable to the human element and physical access. Even if one disregards malicious behavior on behalf of ASN-issued networks, service providers working with IP networks are very clear that BGP is the most complex and difficult to configure Internet protocol, which creates problems for the execution of standard procedures by itself.

The infrastructure built on top of BGP is thus highly vulnerable to a variety of malicious attacks because of the lack of secure means of authenticating the legitimacy of BGP control traffic, as well as the underlying TCP session it is built on, which has its own set of vulnerabilities. In such a framework, AS origins, routes and advertised address spaces can be maliciously misconfigured by adversaries in order to affect the routing process, effectively redirecting traffic, which enables black holes, wiretaps and server masquerades. While such instances can be mitigated through the implementation of route filtering and route registries, it is not a replacement for a strong security architecture. In addition to the heuristic nature of the filtering process, limited computational power and the unwillingness of ISPs to disclose too much information about its customer base have a severely dampening effect on the possibility of maintaining prefix registries, and thus on the effectiveness of filtering in general. Therefore, there is no way to guarantee that a BGP-speaking router uses the

AS number it was allocated or that it holds the address space it advertises. Most of BGP's security problems thus stem from an uncertainty about the relationship between IP prefixes and ASNs, the use of TCP as its underlying transport protocol and the potential to tweak route advertisements in order to subvert routing policies.

Action Items

In 1998, Peiter "Mudge" Zatko, a noted computer security expert, testified in Congress that he could bring down the Internet in 30 minutes using a BGP attack and further demonstrated to government agents how BGP could also be used to eavesdrop (Zetter 2008) on the traffic it coordinates. Using BGP, Anyone with control over a BGP router could advertise any address block using a narrower range of IP addresses than any other AS. In this case, the announcement could propagate across the worldwide network of BGP speakers once neighbors select this route and start advertising it to their peers. In a couple of minutes, "data that should have headed to the intended networks would begin arriving to the eavesdropper's network instead" (Zetter 2013).

Indeed, performing a hijacking attack in BGP can be a relatively simple task, but it is also BGP's most dominant security issue, and can be exploited in a number of ways. Furthermore, malicious attempts are often hidden in the number of times hijacking takes place unintentionally, often as the result of a typo in a routing announcement or some other mistake in the router's configuration. When this does occur, it generally results in a black hole and a following outage, as the routed traffic never reaches its destination. A major instance of this took place in 1997, when a small ISP in Florida misconfigured one of its routers to advertise optimal connectivity

to all Internet destinations. Because these statements were not validated by the upstream, they were widely accepted and as a result, most traffic was routed to this small ISP, which quickly became overwhelmed and effectively crippled the Internet for almost two hours (Butler).

Of course, hijacking can be executed purposefully as well. A greater danger lies in the possibility of interception when, after passing through a malicious AS, the traffic is then forwarded back to the intended recipient. The communication between the two parties thus remains uninterrupted and can be inspected and modified without the parties being immediately aware of this. BGP eavesdropping has long been a known weakness (Zetter 2013), but no one was known to have intentionally exploited it until two security researchers demonstrated it at DefCon 9, where they intercepted traffic bound for the conference and channeled it to a system in New York before redirecting it back to DefCon in Las Vegas. This vulnerability is so severe that it could allow any AS with a mischievous agenda to intercept massive amounts of data, and even modify it en route. According to one of the two researches, they weren't even "doing anything out of the ordinary. […] The problem arises from the level of interconnectivity that's needed to maintain this mess, to keep it all working" (Zetter 2008). Rather than attacking a bug or flaw in BGP, the technique simply takes advantage of the protocol's trust-based architecture.

The demonstrated MITM technique is not trivial, since it requires steps to be taken in order to avoid the redirected traffic to boomerang right back to the malicious AS. To do this, the technique demonstrated at DefCon embeds a loop in an AS_PATH and sends this announcement to select routers, which will then reject the route from their Adj-RIB, enabling the stolen traffic to be redirected through these to its rightful recipients. Earlier this year, researchers at Renesys uncovered such a

MITM attack that targeted traffic heading to government agencies, corporate offices and other recipients in the US and elsewhere. For months, this traffic was mysteriously redirected to Belarus and Iceland before being redirected back on its way to its legitimate destinations (Zetter 2013). According to Renesys, analysts had never seen anything that looked like an intentional attack before, and these were recurring on an almost daily basis. This means that MITM BGP route hijacking has now moved from a theoretical concern whose importance was emphasized for years by a small group, to something that may be happening very frequently without anyone noticing it. However, one cannot carry this kind of attack without leaving a permanent footprint in global routing tables, and according to Renesys, the reason why attackers still proceed with this technique is because they believe that nobody is looking (Cowie).

BGP uses TCP as its transport protocol, which is also vulnerable to tampering, eavesdropping and DoS attacks. Another possibility is for an adversary to redirect traffic and present the user with an impersonation of the intended destination, enabling the possibility of phishing attacks. One must keep in mind that these techniques are made possible due to the nature of route updates, which can change quite frequently due to a variety of natural reasons such as company mergers, natural disasters, server fallouts etc… On good days, routing paths can remain fairly static, but "when the Internet has a bad hair day, the rate of BGP path updates goes up by a factor of 200 or 400" (Zetter 2008). The lack of secure means to authenticate these changes makes the network built on top of BGP extremely vulnerable to malicious attacks.

Defenses

A potential solution to these problems has long been seen as the establishment of a functional Internet Routing Registry (IRR). However doing this at a large scale causes substantial problems, as filtering and route authentication become more and more cumbersome. Kapela, one of the researchers who demonstrated the MITM attack at DefCon, says that eavesdropping could be thwarted if ISP's filtered aggressively, but that this process is labor intensive and requires absolutely all ISPs to participate. Furthermore, such a registry would require ISPs to disclose their address space for all customers, which is something they tend to want to keep private.

To really address the problem, a secure inter-domain protocol must display Byzantine robustness, meaning that a given malicious or faulty behavior should be able to be determined by all other hosts in a finite time period. Finding such a solution is an active area of research (Butler), at the forefront of which is S-BGP, which has been developed by the BBN. S-BGP makes use of a Public Key Infrastructure (PKI) that would store signed certificates issued to ISPs attesting to their address space and AS numbers. Propagated route advertisements would thus be signed, easily verified, and transmitted to the next authorized hop. This means that "nobody could put themselves into the chain, into the path, unless they had been authorized to do so by the preceding AS router in the path" (Zetter 2008). Such a PKI parallels the IP address and ASN assignment, which takes advantage of the existing, hierarchical infrastructure stretching from organizations to ISPs to RIRs all the way up to ICANN, the ultimate authority for address allocation. A new transitive path is employed to carry these digital signatures within a BGP UPDATE. To simplify, each AS signs the route attestation as it traverses the network, making it possible for the router to

validate not only the path, but also that the ASes were traversed in the indicated order and that no ASes were added or removed by an adversary.

However, S-BGP still faces significant barriers to its adoption. The new protocol adds additional complexity, infrastructure and bandwidth requirements. Its implementation requires the participation of several distinct organizations, since individuals cannot justify the expense of investing in this technology unless others have done so, a classic chicken-and-egg problem. In addition, there is resistance to a global PKI with a single root of trust. However, algorithmic research has clearly demonstrated a sharp difference between BGP and S-BGP. While usual attacks on the former are solvable in polynomial time, they are NP-HARD for the latter (Chiesa).

Conclusion

Since its inception in the 70s, the inherent flaws regarding the authenticity of routes advertised by BGP-speakers have been given little attention due to the major overhaul necessary to overcome these problems along with a persistent, yet consistently crumbling notion of inter-domain trust that might have mitigated, but surely helped ignore the possibility of their exploitation. Over time however, the fact that individual ASes can redirect traffic along unintended routes has given rise to an increasing number of examples that showcase how the safety and reliability of the Internet can easily be subverted both intentionally and unintentionally.

In response to this inherent problem facing the Internet's correct, safe and reliable functioning, numerous revisions and patches have been proposed. Out of these, S-BGP stands out with its means to securely authenticate the legitimacy of a propagated path using a Public Key Infrastructure. While the efficacy of its deterrence

is proven algorithmically, its implementation still entails higher costs in the form of hardware, oversight and bandwidth. Thus, its implementation still faces hurdles in the form of capital investment and an initial, substantial move on the part of numerous organizations. But given the severity of potential repercussions and threats presented by the current protocol, these are unlikely to subsist for long and a new protocol and its entailing infrastructure, will soon have to be implemented.

Bibliography

Butler, Kevin, and Patrick McDaniel. "A Survey of BGP Security Issues and Solutions." *Uoregon.edu*. University of Oregon, 1 Jan. 2010. Web. 15 Dec. 2013.

Chiesa, Marco, Giuseppe Di Battista, Thomas Erlebach, and Maurizio Parignani. *Computational Complexity of Traffic Hijacking under BGP and S-BGP*. *Arxiv.org*. Dept. of Computer Science and Automation, Roma Tre University, 21 May 2012. Web. 15 Dec. 2013.

Cowie, Jim. "The New Threat: Targeted Internet Traffic Misdirection." *Renesys The New Threat Targeted Internet Traffic Misdirection Comments*. N.p., 19 Nov. 2013. Web. 15 Dec. 2013.

Dugan, Stephen. "$tea£ing with BGP." Lecture. Blackhat. *Blackhat.com*. Blackhat, 2002. Web. 15 Dec. 2013.

Goldberg, Sharon. "How Secure Are Secure BGP Protocols." Speech. NANOG 49. San Francisco. *Nanog.org*. 15 June 2010. Web. 15 Dec. 2013.

Kent, Stephen, Charles Lynn, and Karen Seo. "Secure Border Gateway Protocol (S-BGP)." *Cs.jhu.edu*. N.p., 4 Apr. 2000. Web. 15 Dec. 2013.

Pilosov, Alex, and Tony Kapela. "Stealing the Internet." Speech. DefCon 9. Las Vegas. *Defcon.com*. 10 Aug. 2008. Web.

Rekhter, Y., T. Li, and S. Hare. "RFC 4271 - A Border Gateway Protocol 4 (BGP-4)." *RFC 4271 - A Border Gateway Protocol 4 (BGP-4)*. N.p., Jan. 2006. Web. 15 Dec. 2013.

"Secure BGP Project (S-BGP)." *Secure BGP*. BBN, 2000. Web. 15 Dec. 2013.

"Spam and Fraud Activity Trends." *Endpoint, Cloud, Mobile & Virtual Security Solutions*. N.p., n.d. Web. 15 Dec. 2013.

Zetter, Kim. "Revealed: The Internet's Biggest Security Hole." *Wired.com*. Conde Nast Digital, 26 Aug. 2008. Web. 15 Dec. 2013.

Zetter, Kim. "Someone's Been Siphoning Data Through a Huge Security Hole in the Internet." *Wired.com*. Conde Nast Digital, 03 Dec. 0013. Web. 15 Dec. 2013.

BGP Hijack

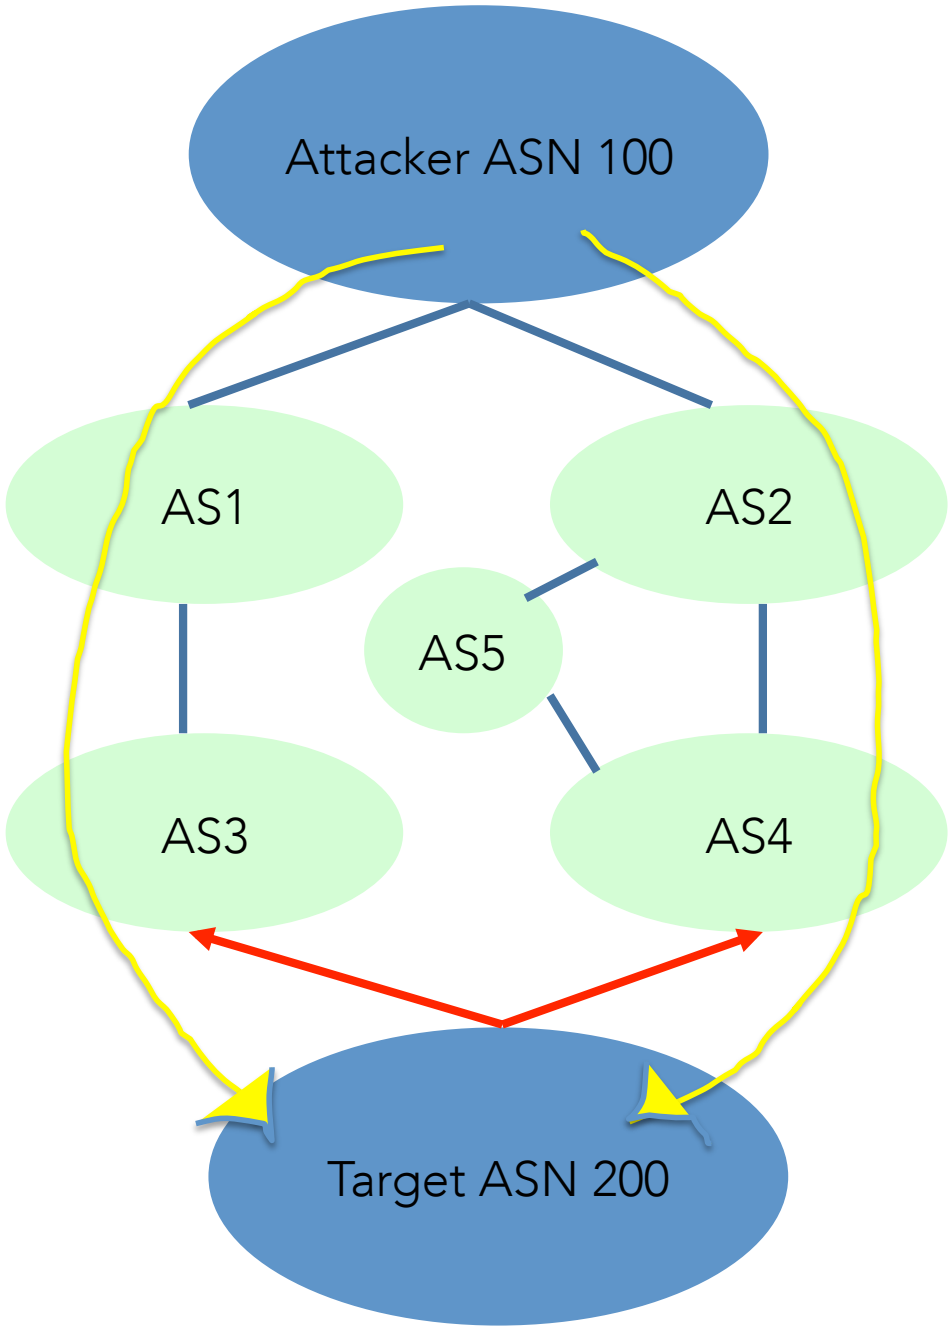

ASN 200 advertises x/15 address block

Which propagates through the Ases

Fastest routes are then determined

BGP Hijack

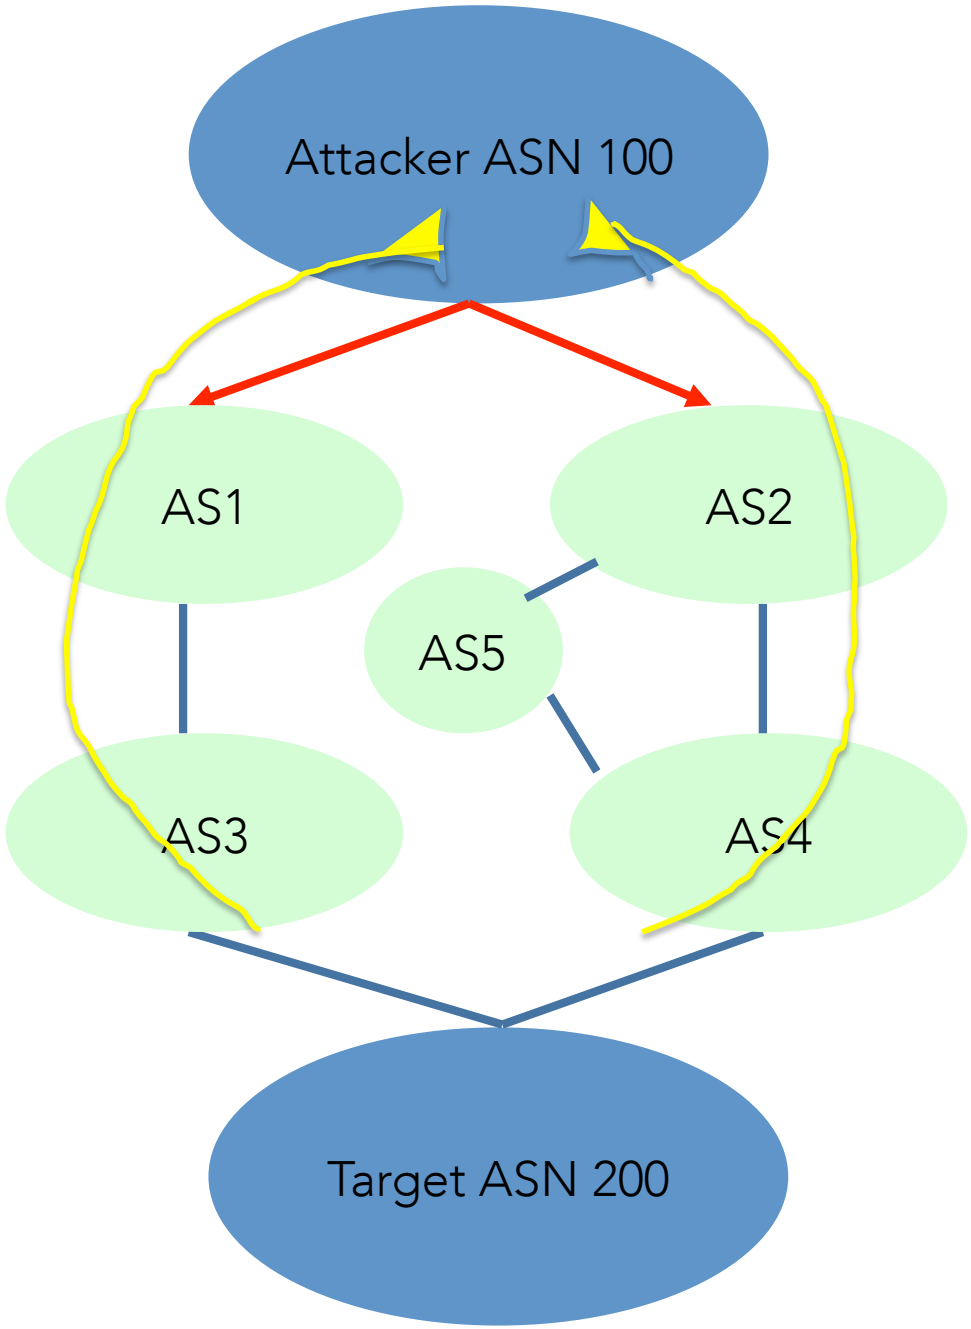

Attacker then advertises x/20 address block

Which propagates through the Ases

Preffered routes are redetermined, traffic effectively redirected