Securing the Border Gateway Protocol

by Stephen T. Kent, BBN Technologies

Routing in the public Internet is based on a distributed system composed of many routers, grouped into management domains called *Autonomous Systems* (ASes). ASes are operated by *Internet Service Providers* (ISPs) and by multihomed subscribers. (Throughout the remainder of this article, for brevity, we will talk in terms of ISPs, usually omitting references to multihomed subscribers.) Routing information is exchanged between ASes using the *Border Gateway Protocol* (BGP)[1], via UPDATE messages.

BGP is used in two different contexts. *External BGP* (eBGP) propagates routes between ISPs. BGP also is used within an AS to propagate routes acquired from other ASes. This latter use is referred to as *internal BGP* (iBGP). eBGP is the primary focus of this article, because failures of eBGP can adversely affect large portions of the Internet, well beyond the administrative boundary of the source of the failure. Nonetheless, some ISPs have expressed interest in protecting the distribution of routes within an ISP. The security technology discussed in this article can be used to secure iBGP, but eBGP is the focus of this article. We use the term "BGP" to refer to eBGP throughout the article.

BGP is highly vulnerable to a variety of attacks[2]. In some cases, this vulnerability arises because of a lack of integrity and authentication for BGP messages. However, the more substantive and harder problem is the lack of a secure means of verifying that BGP traffic is authorized, a concept explored in more detail in this article. In April 1997, BBN began work on the security architecture described here, a system we refer to as *S-BGP,* to address the vulnerabilities of BGP. This article begins by reviewing the problem, discusses a model for correct operation of BGP, presents a threat model, and states the goals and assumptions that underlie our proposed security architecture.

Before we begin the discussion of BGP in more detail, a few definitions are in order. A *route* is defined as an *address prefix* and a set of *path attributes*. One of the path attributes is an AS path, and that is the primary focus of BGP security considerations. The AS path specifies the sequence of ASes that subscriber traffic should traverse if forwarded via this route. When propagating an UPDATE to a neighboring AS, the BGP router prepends its AS number to the sequence, and may update certain other path attributes. The first AS included in the path is referred to as the *origin AS*.

Each BGP router (other than at the edges of the Internet) maintains a complete routing table, capable of routing traffic to any reachable destination, and sends its best route for each prefix to each neighbor. In BGP, "best" is very locally defined. The BGP route selection algorithm has few criteria that are universal, thus limiting the extent to which any security mechanism can detect and reject "bad" routes emitted by a neighbor.

Each ISP makes use of local policies that it need not disclose, and this gives BGP route selection a "black box" flavor, which has significant adverse implications for security.

Correct Operation of BGP

Security for BGP should be defined as the correct operation of BGP routers. This definition is based on the observation that any successful attack against BGP will result in other than correct operation, presumably yielding degraded routing. Correct operation of BGP depends upon the integrity, authenticity, and timeliness of the routing information it distributes, as well as each BGP router processing, storing, and distributing this information in accordance with both the BGP specification and local routing policies. Many statements could be made in an effort to characterize correct operation, but they rest on two simple assumptions.

First, control (vs. subscriber traffic) communication between neighbor BGP routers must be authenticity and integrity secure. This is easily achieved through the use of a point-to-point security protocol capable of protecting BGP traffic; for example, *IP Security* (IPSec). Second, BGP routers must execute the route selection algorithm correctly and communicate the results. There are two parts to this assumption: processing received UPDATEs, and generation and transmission of UPDATEs. In terms of an AS trying to protect itself against external attacks, correct operation of its own BGP routers is mostly a local security issue, but not an Internet-wide security issue. However, an AS should not rely on other ASes to operate properly; such reliance permits a failure in one AS to propagate to others, a domino failure effect. Thus it is important for a BGP router to be able to verify that each UPDATE it receives from a peer is valid (authorized) and timely.

The validity of an UPDATE message is based on four primary criteria:

- The router that sent the UPDATE was authorized to act on behalf of the AS it claims to represent; that is, the AS at the front of the AS path.
- The AS from which the UPDATE emanates was authorized by the preceding AS in the AS path (in the UPDATE message) to advertise the prefixes in the UPDATE.
- The first AS in the AS path was authorized, by the owner of the set of prefixes that are represented in the UPDATE, to advertise those prefixes.
- If the UPDATE withdraws one or more routes (specified by the prefixes for the routes), then the sender must have advertised each route prior to withdrawing it.

There are some limitations to the ability of any practical security mechanism to detect all BGP security failures. The local policy feature of BGP allows each ISP considerable latitude in how UPDATEs are processed, making it difficult for an external observer—for example, a router in a neighboring AS—to determine if a router is operating properly.

This is because such behavior might be attributed to local policies not visible outside an AS. To address such attacks, the semantics of BGP itself would have to change. Moreover, because UPDATEs do not carry sequence numbers, a BGP router can emit an UPDATE based on authentic, but old, information; for example, withdrawing or reasserting a route based on outdated information. Thus the temporal accuracy of UPDATEs, in the face of Byzantine failures, is hard to enforce, except in a very coarse fashion. (Simply speaking, a *Byzantine failure* is one in which a nominally trusted or authorized entity misbehaves.)

Threat Model and BGP Vulnerabilities

Routers exhibit both architectural and implementation vulnerabilities. Implementation vulnerabilities are the result of errors that arise in developing design details or coding; for example, translating the BGP specs into software. Architectural vulnerabilities permit various forms of attack, independent of implementation details, and thus are potentially more damaging, because they persist across all implementations. To make Internet routing robust, both forms of vulnerabilities must be addressed. BGP vulnerabilities can be exploited to cause improper routing or nondelivery of subscriber traffic, network congestion, and traffic delays. Misrouting attacks can be used to facilitate both passive and active wiretapping of subscriber traffic. Often an attack against BGP may be part of a larger attack against subscriber computers. For example, there have been BGP attacks that seek to misroute queries to *Domain Name System* (DNS) root servers, as part of an attack against subscriber systems.

BGP can be attacked in many ways. Communication between BGP peers can be subjected to active or passive wiretapping. The BGP software, configuration information, or routing databases of a router may be modified or replaced via unauthorized access to a router, or to a server or management workstation from which router software is downloaded. These latter attacks transform routers into hostile insiders, so security measures must address such Byzantine failures.

Improved physical and procedural security for network management facilities, and routers, and cryptographic security for BGP traffic between routers would help reduce some of these vulnerabilities. However, physical and procedural security is expensive and imperfect, and these countermeasures would not protect the Internet against accidental or malicious misconfiguration by operators, nor against attacks that mimic such errors. Misconfiguration of this sort has been a source of Internet outages in the past and seems likely to persist. Any security approach that relies on ISPs to act properly violates the "principle of least privilege" and leaves the Internet routing system vulnerable at its weakest link. In contrast, the security approach described in this article satisfies this principle, so that any attack on any component of the routing system is limited in its impact on the Internet as a whole.

Routers also are susceptible to resource exhaustion attacks based on delivery of large quantities of management traffic, BGP or otherwise. This vulnerability arises because these devices are designed with the not unreasonable model that management traffic is a very tiny percentage of all the traffic that arrives at a router. Router interfaces can deliver traffic to the management processor at very high rates, because they are designed to accommodate subscriber traffic flows. Solutions to this problem need to be generic, to accommodate all types of router management traffic, and thus are outside the scope of the BGP security measures discussed in this article.

Goals, Constraints, and Assumptions

Any proposed security architecture must exhibit dynamics consistent with the existing BGP system; for example, responding automatically to topology changes, including the addition of new networks, routers, and ASes. These actions take place on different time scales and have different scopes. For example, in the current BGP system, if an ISP replaces a failed router, the action can take place fairly quickly and has only local impact, because ISPs are not aware of the identity of routers in other, non-neighboring, ISPs. The issuance of new AS numbers, representing new nets, is not a fast process, nor is the allocation of new blocks of address space (new prefixes). But both of these actions are globally visible. Changes in routes also may have global impact, and they may occur very quickly.

Solutions also must scale in a manner consistent with the growth of the Internet. The countermeasures must be consistent with the BGP protocol standards and with the likely evolution of these standards. This includes packet size limits and features such as path aggregation, communities, and multiprotocol support (for example, *Multiprotocol Label Switching* [MPLS]). The security measures must be incrementally deployable; there cannot be a "flag day" when all BGP routers suddenly begin executing a new security protocol. It is desirable to not create new organizational entities that must be accepted as authorities by ISPs and subscribers, in order to make routing secure.

S-BGP Architecture

S-BGP consists of four major elements:

- A *Public Key Infrastructure* (PKI) that represents the ownership and delegation of address prefixes and AS numbers
- *Address Attestations* that the owner of a prefix uses to authorize an AS to originate routes to the prefix
- *Route Attestations* that an AS creates to authorize a neighbor to advertise prefixes
- *IPSec* for point-to-point security of BGP traffic transmitted between routers

These elements are used by an S-BGP router to secure communication with neighbors, and to generate and validate UPDATE messages relative to the authorization model represented by the PKI and address attestations. Together, the combination of these security mechanisms prevents a compromised AS from propagating erroneous routing data to other, secured ASes. Each element is described in more detail in the following section.

S-BGP Public Key Infrastructure

S-BGP uses a PKI based on X.509 (v3) certificates to enable routers to validate the authorization of other routers to represent ASes (ISPs). The PKI also allows routers to verify the authorization of each ISP as the owner of one or more prefixes (contiguous blocks of address space). This PKI was described in[14], and the reader is referred to that paper for additional details. The PKI parallels the existing IP address and AS number assignment delegation system and takes advantage of this infrastructure. Because the PKI mirrors existing infrastructure, it avoids most of the "trust" issues that often complicate the creation of a PKI. This PKI is unusual in that it emphasizes authorization, not authentication. The names used in the certificates in this PKI are not employed to determine whether a given ISP or router is authorized to do anything, and the names are not even meaningful outside of S-BGP.

S-BGP calls for a certificate to be issued to each ISP (or subscriber) that owns (more properly, has a right to use) a portion of the IP address space. This certificate is issued through the same procedures employed for address allocation, starting with the *Internet Assigned Numbers Authority* (IANA) and continuing through a *Regional Internet Registry* (RIR), and, if applicable, an ISP. If an ISP owns multiple prefixes, we issue a single certificate containing a list of prefixes, to minimize the number of certificates in the system. The PKI represents address-space ownership by binding prefixes to a public key belonging to the ISP to which the prefixes have been assigned. Each certificate contains a private extension that specifies the set of prefixes that has been allocated to the ISP. Certificates issued under this PKI also represent the binding between an ISP and the AS numbers allocated to it. The PKI allows each ISP to issue certificates to its routers, certifying that these routers represent the ISP and hence, the ASes owned by the ISP. Here too, the PKI parallels the existing AS allocation system; that is, the IANA allocates AS numbers to RIRs, which in turn assign AS numbers to ISPs that run S-BGP.

Attestations

An *attestation* is a digitally signed datum asserting that its target (an AS) is authorized by the signer (an ISP) to advertise a path to one or more specified prefixes. There are two types of attestations, address and route, which share a common format. For an *Address Attestation* (AA), the signer is the ISP or subscriber that controls the prefixes in the AA, and the target is a set of ASes that the ISP/subscriber authorizes to originate a route to the prefixes. AAs are relatively static data items, because relationships between address-space owners and ISPs change relatively slowly.

For a *Route Attestation* (RA), the signer is an S-BGP router (operating on behalf of an ISP), and the target is an AS or set of ASes, representing the neighbors to which the UPDATE containing the RA will be sent. RAs, unlike AAs, are very dynamic, possibly changing for each transmitted UPDATE.

UPDATE Validation

Attestations and certificates are used by S-BGP routers to validate routes asserted in UPDATE messages; that is, to verify that the first AS in the route has been authorized to advertise the prefixes by the prefix owner(s), and that each subsequent AS has been authorized to advertise the route for the prefixes by the preceding AS in the route. To validate a route received from AS_n, AS_{n+1} requires:

- An AA for each organization owning a prefix represented in the UPDATE (not for prefixes in the UPDATE that represent routes being withdrawn)
- A certified public key for each organization owning a prefix in the UPDATE
- An RA corresponding to each AS along the path (AS_n to AS_1), where the RA generated and signed by the router in AS_n encompasses the *Network Layer Reachability Information* (NLRI) and the path from AS_{n+1} through AS_1
- A certified public key for each S-BGP router that signed an RA along the path (AS_n to AS_1), to check the signatures on the corresponding RAs

An S-BGP router verifies that the advertised prefixes and the origin AS are consistent with AA information. The router verifies the signature on each RA and verifies the correspondence between the signer of the RA and the authorization to represent the AS in question. There also must be a correspondence between each AS in the path and an appropriate RA. If all of these checks pass, the UPDATE is valid.

AAs are not used to check withdrawn routes in an UPDATE. Use of IPSec to secure communication between each pair of S-BGP routers, plus the fact that BGP uses a separate *Adjacency Routing Information Base* (Adj-RIB-In) for each neighbor, ensures that only the advertiser of a route can withdraw it.

Distribution of S-BGP Data

Each S-BGP router must have the public keys required to validate the RAs in UPDATEs, a scenario that translates into securely distributed keys for every router that implements S-BGP (and that is reachable via an S-BGP path). Each router also needs access to all AA information, to verify that the origin AS is authorized to originate a route to the prefixes in the UPDATE. S-BGP does not distribute certificates, *Certificate Revocation Lists* (CRLs), or AAs via UPDATE messages; transmission of these items via UPDATEs would be very wasteful of bandwidth, because each BGP router would receive many redundant copies from its neighbors.

Also, an UPDATE is limited to 4096 bytes and thus generally could not carry all of this data for the route represented by the UPDATE. Instead, S-BGP distributes this data to routers via out-of-band means. The data is relatively static and thus is a good candidate for caching and incremental update. Moreover, the certificates and AAs can be validated and reduced to a more compact format by ISP operation centers prior to distribution to routers. This avoids the need for each router to perform this processing, saving both bandwidth and storage space. It also means that routers do not need to be able to parse X.509 certificates and validate certificate paths for S-BGP purposes, although some capability in this area may be required for IPSec key management.

S-BGP uses *repositories* for distribution of this data. We initially described a model in which a few replicated, loosely synchronized repositories were operated by the RIRs. Discussions with ISPs suggest a model in which major ISPs and Internet exchanges operate repositories, and smaller ISPs and subscribers make use of these repositories. In either model, each ISP periodically, for example daily, uploads new/changed certificates, its current CRL, and AAs. Each ISP also downloads all of this data for all other ISPs that are running S-BGP. The repositories periodically transfer new data to one another to maintain loose synchronization. ISPs process the repository information to create more compact files that contain the AA data and the public keys and prefix and AS data from the certificates, but none of the certificate management information or CRLs. These resulting "extracted" files are transferred to the routers executing S-BGP under the control of the ISP.

Because certificates, AAs, and CRLs are signed and carry validity interval information, they require minimal additional security while in transit to or from a repository or while stored on a repository. Nonetheless, S-BGP employs the *Secure Sockets Layer* (SSL) protocol, with both client and server certificates, to protect access to the repositories, as a countermeasure to denial-of-service attacks. The simple, hierarchic structure of the PKI allows repositories to automatically effect access control checks on the uploaded data, for example, to prevent one ISP from accidentally or maliciously overwriting the certificates, CRLs, and AAs from another ISP.

Distribution of Route Attestations

S-BGP distributes RAs with BGP UPDATEs in a newly defined, optional, *transitive path attribute*. Because routes may change quickly, it is important that RAs accompany the UPDATEs that are validated using them. If any other means of distribution is employed for this data, there is a likelihood that the UPDATEs and the data will be out of synch, creating a conundrum for a router; that is, what should the router do when the UPDATE and the security data differ? RAs employ a compact encoding scheme to help ensure that they fit within the BGP packet size limits, even when route or address aggregation occurs. (S-BGP accommodates aggregation by explicitly including signed attribute data that otherwise would be lost when aggregation occurs.) An S-BGP router receiving an UPDATE from a peer caches the RAs with the route in the Adj-RIB for the peer, and in the *Local Routing Information Base* [Loc-RIB] (if the route is selected).

As noted in the following discussion, the bandwidth required to support in-band distribution of route attestations is negligible (compared to subscriber traffic).

Although the RA mechanism was designed to protect AS path data, it can also accommodate other new path attributes; for example, communities[11] and confederations[12]. Specifically, there is a provision to indicate what data, in addition to the AS path, is covered by the digital signature that is part of the RA.

Putting It All Together

Figure 1 illustrates how the major elements of S-BGP interact, using a simplified example. The figure shows two ISPs, each with a *Network Operations Center* (NOC), a repository, and three routers. A third ISP is represented by a single (S-BGP-enabled) router. Each ISP interacts with an RIR to acquire a certificate representing the prefixes and AS numbers assigned to the ISP. Each NOC interacts with a repository to upload data (certificates, CRLs, and AAs) from that ISP, and to download the same data acquired from all other ISPs. The repositories interact with one another to exchange uploaded ISP data, to make that data available to all other ISPs. Within an ISP, the NOC pushes a copy of the extracted certificate and AA data, produced from the downloads acquired from a repository, to each router. Routers exchange UPDATE messages, containing RAs, that enable validation of each received UPDATE.

Figure 1: S-BGP Element Interactions

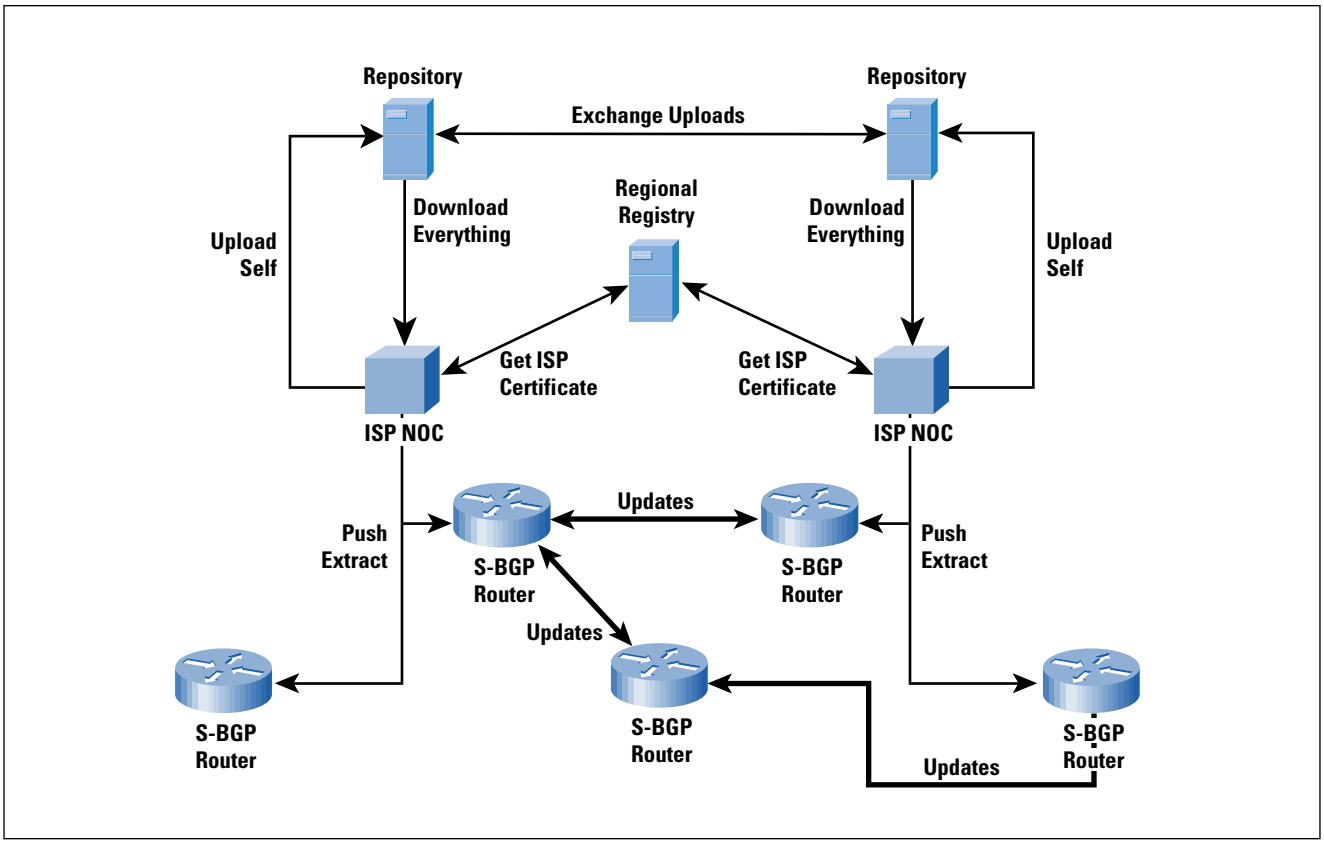

IPSec and Router Authentication

S-BGP uses IPSec[6,7,8], specifically the *Encapsulating Security Payload* (ESP) protocol, to provide authentication, data integrity, and antireplay for all BGP traffic between neighboring routers. The *Internet Key Exchange* (IKE) protocol[9,10] is used for key management services in support of ESP. The S-BGP PKI includes certificates for IKE, separate from those used for RA processing.

The use of IPSec is preferable to the current option of the *Message Digest Algorithm 5* (MD5) TCP checksum option[15], in several respects. IPSec uses keyed hash functions in a way that is cryptographically more secure that the MD5 checksum option, and IKE provides automated key management, a feature sorely lacking in the option. Protecting BGP traffic at the IP layer, vs. the TCP layer, counters more vulnerabilities, because the TCP implementation is protected as well, for example, including SYN flooding and spoofed RSTs (resets), are rejected.

Residual Vulnerabilities in S-BGP

Despite the extensive security offered by S-BGP, architectural vulnerabilities exist that are not eliminated by its use. For example, an S-BGP router may reassert a route that was withdrawn earlier, even if the route has not been readvertised. The router also may suppress UPDATEs, including ones that withdraw routes. These vulnerabilities exist because BGP UPDATEs do not carry sequence numbers or time stamps that could be used to determine their timeliness. However, RAs do carry an expiration date and time, so there is a limit on how long an attestation can be misused this way. S-BGP restricts malicious behavior to the set of actions for which a router or AS is authorized, based on externally verifiable, authoritative constraints.

Performance and Operational Issues

In developing the S-BGP architecture, we paid close attention to the performance and operational impact of the proposed countermeasures, and reported our analysis in earlier papers. In preparing this article, we updated our data, utilizing a variety of sources; for example, the *Route Views* project. Although much data about BGP and associated infrastructure is available, other data is difficult to acquire in a fashion that is representative of a "typical" BGP router. This is because each AS in the Internet embodies a slightly different view of connectivity, as a result of local policy filters applied by other ASes.

It is important that the transmission, storage, and processing requirements imposed by S-BGP not be so great as to overwhelm routers. Each of these requirements must be analyzed separately.

The transmission of RAs in UPDATEs does significantly increase the size of these messages, by about 800 percent. However, because the volume of this traffic is minuscule relative to subscriber traffic, the increase is negligent. The set of files containing certificates, AAs, and CRLs would be about 75–85 MB. Daily transmission of these files between ISPs and repositories would not represent a significant increase in traffic volume for the Internet.

Although the transmission overhead is not a concern, storage of the RAs in each Adj-RIB and the Loc-RIB is a problem. The additional space required to hold these RAs is estimated at about 30–35 MB per peer, if S-BGP were fully deployed today. This is a modest amount of memory for a typical router with a few peers, but a significant amount of storage for routers at Internet exchanges, where a router may have tens or even hundreds of peers.

Thus the management CPU in a router might need a gigabyte or more of RAM under these conditions. (When a large ISP peers with many other ISPs at an exchange, the peering is not symmetric; that is, the large ISP accepts only a few routes from each of the smaller ISPs, filtering out the rest. Thus the amount of additional memory required for RAs in Adj-RIBs for each of these small ISP peers may be considerably less than for symmetric peer relationships.) This requisite memory seems modest by current workstation standards, but most deployed routers cannot be configured with this much memory.

The computational burden of router processing of RAs in UPDATEs is a function of the path length in each UPDATE and the rate at which UPDATEs arrive. The arrival rate is a function of the number of S-BGP peers the router sees, and the rate at which each peer sends UPDATEs. Our analysis suggests that the long-term (24-hour) UPDATE rate for a router with 30 peers is about 0.5 UPDATEs per second. On average, each UPDATE would contain about 3.7 RAs. We originally estimated the busy minute rate as about 10 times the average rate. At this rate, a router could probably perform the requisite signature verification in software (about 18 signature verifications per second). Recent evidence suggests a factor of 100–200 might be a better estimate, in light of experience with major worm attacks, and at that rate it would be hard for software to keep pace.

Heuristics are available to reduce this burden. Analysis shows that about 50 percent of all UPDATEs are sent as a result of route "flaps"; that is, transient communication failures that, when remedied, result in a return to the former route. Thus if a router maintained a depth-two cache for each Adj-RIB-In, it could avoid signature validation about 50 percent of the time. However, this would double the storage requirements for these RIBs, and that would exacerbate the storage problem cited previously.

Our previous analysis also assumed that receipt of each UPDATE would result in transmission of an UPDATE with one new signature. This was an oversimplification; a router generates and transmits an UPDATE only if the newly received route is "better" than the current best route (for the prefix), or if the best route for the prefix is withdrawn by the UPDATE. When a router has many peers, most of the UPDATEs it receives may not yield a better route, and thus will not trigger transmission of a new UPDATE.

On the other hand, when a router does select a new route, an UPDATE may be constructed and sent to each neighbor, requiring one signature per neighbor. This is because an RA specifies the AS number of the neighbor to which it is directed. It is possible to construct an RA that identifies the next hop as a set of AS numbers, corresponding to all the neighbors to which an UPDATE is authorized to be sent. The downside of this strategy is that it makes the RAs larger, contributing to the storage problem noted previously.

The observation made previously suggests a heuristic for UPDATE processing to mitigate signature validation costs. A router can defer validation of the RAs in any UPDATE that it receives, if the UPDATE would not represent a new best route. This optimization could be especially helpful for routers that receive the greatest number of UPDATEs; that is, routers with many neighbors. One might worry that this strategy allows an attacker to force processing, by sending what would be considered "very good" routes, but an S-BGP router could detect such fraudulent UPDATEs and could choose to drop its connection to a peer that behaved this way, in order to counter such an attack.

Initialization/reboot of a BGP router also results in a surge in UPDATE processing, and the deferred processing heuristic is applicable here too, even though reboots are relatively infrequent. Saving RIBs in nonvolatile storage addresses this problem. Most deployed routers do not have sufficient nonvolatile storage to adopt this strategy, but some do have hard drives that would easily accommodate the RIBs.

It is reasonable to assume that next-generation routers could be configured with enough RAM for the RIBs, but this analysis shows that full deployment is not feasible with the currently deployed router base. To add RAM, and possibly to add nonvolatile storage, router vendors will have to upgrade the processor boards where net management processing takes place. That suggests that addition of a crypto accelerator chip would be prudent as part of the board redesign process, for example, to deal with surge conditions noted previously.

Deployment and Transition Issues

Adoption of S-BGP requires cooperation among several groups. ISPs and subscribers running BGP must cooperate to generate and distribute AAs. Major ISPs must implement the S-BGP security mechanisms in order to offer significant benefit to the Internet community. The IANA and RIRs must enhance operational procedures to support generation of prefix and AS number allocation certificates. Router vendors need to offer additional storage in next-generation products, or offer ancillary devices for use with existing router products, and revise BGP software to support S-BGP.

There is some good news; S-BGP can be deployed incrementally. Only neighboring ASes receive full benefit from such deployment. Although we chose a transitive path attribute syntax to carry RAs, and thus it might be possible for non-neighbor ASes to exchange RAs, it seems likely that intervening ASes would not have sufficient storage for the RAs in their RIBs.

Also, the controls needed in routers to take advantage of noncontiguous deployment of S-BGP are quite complex, hence our suggestion that only contiguous deployment of S-BGP be attempted.

External routes received from S-BGP peers need to be redistributed within the AS, both to interior routers and to other border routers, in order to maintain a consistent and stable view of the exterior routes across the AS. Thus an AS must switch to using S-BGP for all its border routers at once, to avoid route loops within the AS.

Status

As of early 2003, an implementation of S-BGP has been developed and demonstrated on small numbers of workstations representing small numbers of ASes. We also developed software for a simple repository, and for NOC tools that support secure upload and download of certificates, CRLs, and AAs to and from repositories, and for certificate management for NOC personnel and routers. This suite of software, plus CA software from another *Defense Advanced Research Projects Agency* (DARPA) program, provide all of the elements needed to represent a full S-BGP system. All of this software is available in open source form.

Summary

S-BGP represents a comprehensive approach to addressing a wide range of security concerns associated with BGP. It detects and rejects unauthorized UPDATE messages, irrespective of the means by which they arise; for example, misconfiguration, active wiretapping, compromise of routers or management systems, etc. S-BGP is not perfect; it has a few residual vulnerabilities, but these pale in comparison to the security features S-BGP provides, and removal of these vulnerabilities would require more fundamental changes to BGP semantics.

The S-BGP design is based on a top-down security analysis, starting with the semantics of BGP and factoring in the wide range of attacks that have or could be launched against the existing infrastructure.

Acknowledgements

Many individuals contributed to the design and development of S-BGP, including Christine Jones, Charlie Lynn, Joanne Mikkelson, and Karen Seo.

References

[1] Y. Rekhter and T. Li, "A Border Gateway Protocol 4 (BGP-4)," RFC 1771, March 1995.

[2] S. Kent, C. Lynn, and K. Seo, "Secure Border Gateway Protocol (S-BGP)," *IEEE Journal on Selected Areas in Communications,* Vol. 18, No. 4, April 2000.

[3] C. Villamizar, R. Chandra, and R. Govindan, "BGP Route Flap Damping," RFC 2439, November 1998.

[4] B.R. Smith, and J.J. Garcia-Luna-Aceves, "Securing the Border Gateway Routing Protocol," Proceedings of Global Internet '96, November 1996.

[5] S. Murphy, panel presentation on "Security Architecture for the Internet Infrastructure," Symposium on Network and Distributed System Security, April 1995.

[6] S. Kent and R. Atkinson, "Security Architecture for the Internet Protocol," RFC 2401, November 1998.

[7] R. Glenn and S. Kent, "The NULL Encryption Algorithm and Its Use with IPsec," RFC 2410, November 1998.

[8] S. Kent and R. Atkinson, "IP Encapsulating Security Payload (ESP)," RFC 2406, November 1998.

[9] D. Maughan, M. Schertler, M. Schneider, and J. Turner, "Internet Security Association and Key Management Protocol (ISAKMP)," RFC 2408, November 1998.

[10] D. Harkins and D. Carrel, "The Internet Key Exchange (IKE)," RFC 2409, November 1998.

[11] R. Chandra, P. Traina, and T. Li, "BGP Communities Attribute," RFC 1997, August 1996.

[12] P. Traina, "Autonomous System Confederations for BGP," RFC 1965, June 1996.

[13] T. Bates, R. Chandra, D. Katz, and Y. Rekhter, "Multiprotocol Extensions for BGP-4," RFC 2283, February 1998.

[14] K. Seo, C. Lynn, and S. Kent, "Public-Key Infrastructure for the Secure Border Gateway Protocol (S-BGP)," DARPA Information Survivability Conference and Exposition, June 2001.

[15] A. Heffernan, "Protection of BGP Sessions via the TCP MD5 Signature Option," RFC 2385, August 1998.

STEPHEN KENT received the S.M., E.E., and Ph.D. degrees in computer science from MIT, and a B.S. in mathematics from Loyola University of New Orleans. He has worked at BBN for over 25 years, where he serves today as Chief Scientist–Information Security. He served on the IAB for over a decade, and chaired the Privacy & Security Research Group of the IRTF and the PEM WG in the IETF, where he currently co-chairs the PKIX WG. He has served on several committees for the National Research Council, and chairs a committee on authentication and privacy for the NRC. His current work focuses on PKI issues, BGP security, and very high speed IP encryption. He is a Fellow of the ACM, and a member of the Internet Society and Sigma Xi. His e-mail address is: **kent@bbn.com**