

Securing the Border Gateway Protocol: A Status Update

Stephen T. Kent

BBN Technologies, 10 Moulton Street, Cambridge, MA, U.S. 02138
kent@bbn.com

Abstract. The Border Gateway Protocol (BGP) is a critical component of the Internet routing infrastructure, used to distribute routing information between autonomous systems (ASes). It is highly vulnerable to a variety of malicious attacks and benign operator errors. Under DARPA sponsorship, BBN has developed a secure version of BGP (S-BGP) that addresses most of BGP's architectural security problems. This paper reviews BGP vulnerabilities and their implications, derives security requirements based on the semantics of the protocol, and describes the S-BGP architecture. Refinements to the original S-BGP design, based on interactions with ISP operations personnel and further experience with a prototype implementation are presented, including a heuristic for significantly improving performance. The paper concludes with a comparison of S-BGP to other proposed approaches.

Problem Description

Routing in the public Internet is based on a distributed system composed of many routers, grouped into management domains called Autonomous Systems (ASes). Routing information is exchanged between ASes using the Border Gateway Protocol (BGP) [1], via UPDATE messages. BGP is highly vulnerable to a variety of attacks [2], due to the lack of a secure means of verifying the authorization of BGP control traffic. In April 1997, we began work on the security architecture described in this paper. We begin by reviewing the problem—a model for correct operation of BGP, BGP vulnerabilities and a threat model, and the goals, constraints and assumptions that underlie S-BGP. The reader is assumed to be familiar with the fundamentals of BGP.

BGP is used in two different contexts. External use of BGP (eBGP) propagates routes between ISPs, or between ISPs and subscriber networks that are connected to more than one ISP, i.e., multi-homed subscribers. BGP also is used internally, within an AS, to propagate routes acquired from other ASes. This use is referred to as internal BGP (iBGP). eBGP is the primary focus of this work, because failures of eBGP adversely affect subscribers outside the administrative boundary of the source of the failure. Nonetheless, some ISPs have expressed interest in using S-BGP to protect the

distribution of routes within an ISP. If route servers are employed for iBGP (see section 5.5), or if the number of iBGP peers is small, S-BGP may be a viable approach for iBGP security. We use “BGP” to refer to eBGP, unless otherwise noted.

A route is defined as an address prefix and a set of path attributes, one of which is an AS path. The AS path specifies the sequence of ASes that subscriber traffic will 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.

Each BGP router maintains a full routing table, 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, which limits 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

As we noted in [2], security for BGP is 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’s processing, storing, and distribution of this information in accordance with both the BGP specification and with the local routing policies. Many statements could be made in an effort to characterize correct operation, but they rest on two simple assumptions:

- Communications between peer BGP routers is authenticity & integrity secure
- BGP routers execute the route selection algorithm correctly and communicate the results

The first assumption is easily realized through the use of a suitable, point-to-point security protocol, e.g., IPsec. The second assumption is divisible into two cases: processing received UPDATES, and generation and transmission of UPDATES. From the perspective of an AS trying to protect itself against external attacks, correct operation of its own BGP routers is an implementation, not an architectural, security issue. However, an AS ought not rely on other ASes to operate properly, since such reliance leads to cascade failures. It is desirable for a BGP router to be able to verify that each UPDATE it receives from a peer is valid and timely. Validity of an UPDATE encompasses four primary criteria.

1. The router that sent the UPDATE was authorized to act on behalf of the AS it claims to represent (by virtue of placing that AS number in the AS path).
2. The AS from which the UPDATE emanates was authorized by the preceding AS in the AS path to advertise the prefixes contained within the UPDATE.
3. The first AS in the AS path was authorized, by the “owner” of the set of prefixes (NLRI), to advertise those prefixes.

4. If the UPDATE withdraws one or more routes, then the sender must have advertised the route(s) prior to withdrawing it (them).

There are limitations to the ability of any security mechanism to detect attacks. The local policy feature of BGP allows considerable latitude in UPDATE processing, so S-BGP cannot detect erroneous behavior that could 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 generate an UPDATE based on authentic, but old, information, e.g., withdrawing or reasserting a route based on outdated information. Thus the temporal accuracy of UPDATES, in the face of Byzantine failures, is enforced only very coarsely by these countermeasures.

Threat Model and BGP Vulnerabilities

BGP has many vulnerabilities that can be exploited to cause improper routing or non-delivery of subscriber traffic, network congestion, and traffic delays. Misrouting attacks facilitate passive and active wiretapping of subscriber traffic, and thus an attack against BGP may be part of a larger attack against subscriber computers.

Routers are vulnerable in both the architectural and implementation domains. Implementation vulnerabilities may allow an attacker to assume control of a router, to cause it to operate maliciously, or to cause the router to crash, and thus deny service. Architectural vulnerabilities permit various forms of attack, independent of implementation details, and thus are potentially more damaging, as they persist across all implementations. To make Internet routing robust, both forms of vulnerabilities must be addressed. S-BGP does not directly address implementation vulnerabilities, but it does limit the impact of such vulnerabilities. Use of S-BGP by an AS protects that AS against many attacks that result from security failures suffered by other ASes.

BGP can be attacked in many ways. Communication between BGP peers can be subjected to active and passive wiretapping. A router's BGP software, configuration information, or routing databases may be modified or replaced illicitly via unauthorized access to a router, or to a server from which router software is downloaded, or via an attacked distribution channel. Most of these attacks transform routers into hostile insiders. Effective security measures must address such Byzantine failures.

Many countermeasures could be employed in an attempt to address these vulnerabilities. Better physical and procedural security for network management facilities, and routers would help. Cryptographic protection of BGP traffic between routers and for network management traffic would also reduce some of these vulnerabilities. However, improved 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 proved to be a source of several significant Internet outages in the past and seems likely to persist. Any security approach that relies on ISPs to act properly, that relies on "trust" among ISPs, violates the "principle of least privilege" and leaves the Internet routing system vulnerable at its weakest link. In contrast, the security ap-

proach described here 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.

Goals, Constraints, and Assumptions

In order to create countermeasures that are both effective and practical, the S-BGP architecture is based on the following goals, constraints, and assumptions.

Any proposed security architecture must exhibit dynamics consistent with the existing system, e.g., responding automatically to topology changes, including the addition of new networks, routers and ASes. 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 multi-protocol support (e.g., MPLS).

The S-BGP architecture must be incrementally deployable; there cannot be a “flag day” when all BGP routers suddenly begin executing S-BGP.

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 provide a “firebreak” that prevents a compromised AS from propagating erroneous routing data to other (secured) ASes.

S-BGP Public Key Infrastructure (PKI)

S-BGP uses a PKI based on X.509 (v3) certificates to enable BGP routers to validate the authorization of BGP routers to represent ASes and prefixes. This PKI was described in [24] and the reader is referred to that paper for additional details. The S-BGP 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 many of the “trust” issues that often complicate the creation of a PKI.

This PKI is unusual in that it focuses on authorization, not authentication; the names used in most of the certificates in this PKI are not meaningful outside of S-BGP.

S-BGP calls for a certificate to be issued to each organization that is granted "right to use" of a portion of the IP address space. This certificate is issued through the same chain of entities that today is responsible for address allocation starting with the IANA¹. If an ISP or subscriber owns multiple prefixes, we issue a single certificate containing a list of prefixes, to minimize the number of certificates needed to validate an UPDATE.

This PKI represents the assignment of prefixes by binding prefixes to a public key belonging to the organization to which the prefixes have been assigned. These certificates are used to prove "ownership" of one or more prefixes². Each certificate in this PKI contains a private extension that specifies the set of prefixes that has been allocated to the organization. We use the Domain Component (RFC 2247) construct in the subject name in each certificate to represent a DNS-style name for an organization.

Certificates issued under this PKI also represent the binding between an organization and the AS numbers allocated to it. The PKI allows an organization to certify that a router represents the organization's AS(es). Here too, the PKI parallels existing "trust relationships," i.e., the IANA assigns AS numbers to RIRs, which in turn assign AS numbers to ISPs or subscribers that run BGP, and they certify their routers.

Attestations

An attestation is a digitally signed datum asserting that its target (an AS) is authorized by the signer (an organization) to advertise a path to the specified prefix(es). There are two types of attestations, address and route. They share a single format.

- Address attestation (AA)— the signer of an AA is the organization that "owns" the prefix(es) in the AA, and the target is a set of ASes that the organization authorizes to originate a route to the prefix(es), i.e., the ISPs with which the issuer has a traffic carriage arrangement. AAs are relatively static data items, since relationships between address space owners and ISPs change relatively slowly.
- Route attestation (RA)— the signer of an RA is an S-BGP router in an ISP. The target is a set of ASes, representing the neighbors to which UPDATEs containing the RA will be sent. Note that the router signing an RA might sign a separate RA for each neighbor, or it may sign a single RA directed to all of its neighbors. The latter option permits a router to reduce its digital signature burden, so long as the same parameters appear in the UPDATEs sent to each neighbor. RAs, unlike AAs, are very dynamic, possibly changing for each transmitted UPDATE.

¹ In [24] we proposed cross-certification among the RIRs to create an analogue to a single root PKI, in case politics prevented adoption of a single root.

²One could use X.509 attribute certificates to represent this authorization, but they offer little benefit in this context and would increase the certificate processing burden.

UPDATE Validation

Attestations and certificates are used by BGP routers to validate routes asserted in UPDATE messages, i.e., 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 NLRI portion of the UPDATE
- a valid public key for each organization owning a prefix in the NLRI
- an RA corresponding to each AS along the path (AS_n to AS_1), where the RA generated and signed by router in AS_n encompasses the 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 an 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.

Address attestations 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 use of 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, which usually means keys for every other S-BGP router. Each router also needs access to all AAs, to verify that the origin AS is authorized to originate a route to the prefix(es) in the UPDATE.

S-BGP does not distribute certificates, CRLs, or AAs via UPDATES. Transmission of these items via UPDATES would be very wasteful of bandwidth, as each BGP router would receive many redundant copies from its peers. Also, an UPDATE is limited to 4096 bytes and thus generally could not carry this data. 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 (processed against CRLs) and reduced to a more compact, “extracted file” format³ 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.

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

³ Only the public key, subject name, and selected extensions need be retained.

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, ISPs periodically (e.g., daily), upload and download new/changed certificates, CRLs, and AAs. The repositories periodically transfer new data to one another to maintain loose synchronization. ISPs process the repository information to create “extracted files” and transfer them to their routers.

Since certificates, AAs, and CRLs are signed and carry validity interval information, they require minimal additional security. Nonetheless, S-BGP employs SSL, 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.

Distribution of Route Attestations

Route attestations (RAs) are distributed with BGP UPDATEs in a newly defined, optional, transitive path attribute. Because RAs may change quickly, it is important that they accompany the UPDATEs that are validated using the RAs. When an S-BGP router opens a BGP session with a peer, transmitting a portion of its routing information database via UPDATEs, relevant RAs are sent with each UPDATE, and with subsequent UPDATEs sent in response to route changes. These attestations employ a compact encoding scheme to help ensure that they fit within the BGP packet size limits, even when route or address aggregation is employed. (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 Loc-RIB (if the route is selected). As noted below in Section 4, the bandwidth required to support in-band distribution of route attestations is negligible (compared to user traffic).

Although the RA mechanism was designed to protect AS path data, it can also accommodate other new path attributes, e.g., communities [13] and confederations [14]. 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.

IPsec and Router Authentication

S-BGP uses IPsec [8,9,10], specifically the Encapsulating Security Payload (ESP) protocol, to provide authentication, data integrity, and anti-replay for all BGP traffic between neighboring routers. The Internet Key Exchange protocol (IKE) [11,12] is used for key management services in support of ESP. The PKI established for S-BGP includes certificates for IKE, separate from those used for RA processing.

How S-BGP Addresses BGP Vulnerabilities

Together, the S-BGP PKI and AAs support validation of router assertions about:

- the ASes the router is authorized to represent
- the prefixes an AS is authorized to originate
- the prefixes an AS has been authorized to advertise by other ASes

The UPDATE validation procedure described earlier ensures that every AS along the path has been authorized by the preceding AS to advertise the prefixes in the UPDATE, and that the origin AS was authorized by the prefix user.

AAs allow a router to detect any attempt by an AS to advertise itself as an origin for a prefix unless the prefix owner has authorized the AS to do so. The use of RAs in UPDATES allows an S-BGP router to detect any tampering with a path by any intermediate router. This includes attempts to add to the set of prefixes in the NLRI, to add or remove AS numbers from the AS path, or to synthesize a bogus UPDATE.

The use of IPsec protects all S-BGP traffic between routers against active wiretap attacks. It is necessary to prevent a wiretapper from sending UPDATES that only withdraw routes (and thus would not contain any RAs to be validated) and to prevent such an attacker from replaying valid UPDATES. IPsec also protects the router against various TCP-based attacks, including SYN flooding and spoofed RSTs (resets).

Despite the extensive security offered by S-BGP, there exist architectural vulnerabilities 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 re-advertised. The router also may suppress UPDATES, including ones that withdraw routes. These vulnerabilities exist because BGP UPDATES do not carry sequence numbers or timestamps that could be used to determine their timeliness. However, RAs do carry an expiration date & 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 paper, we updated our data, utilizing a variety of sources, e.g., 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.

Some BGP and S-BGP Parameters

The backbone routers of the major ISPs have a route to every reachable IP address. As of 2003, the routing information databases (Loc-RIBs) in these routers contain about

125,000 IPv4 address prefixes. Each route contains an average of about 3.7 ASes, and typically there would be one route attestation per AS, which provides a basis for calculating how much space is devoted to RAs in UPDATE messages and in RIBs.

Over a 24 hour period, a typical BGP router receives an average of about one UPDATE per minute per peer. Thus a router at an Internet exchange with 30 peers, receives about .5 UPDATES per second, on average. This rate is affected somewhat by Internet growth, but it is primarily a function of link, component, or congestion failures and recoveries.

We originally estimated the peak, per-minute rate for UPDATES at about 10 times the average. However, more recent data suggests that, in times of extreme stress, the peak UPDATE rate might be as much as 200 times the average. Analysis shows that about 50% of all UPDATES are sent as a result of route "flaps," i.e., transient communication failures that, when remedied, result in a return to the former route. This sort of routing behavior has long been characteristic of the Internet⁴ [3].

The X.509 certificates used in S-BGP are about 600 bytes long. The certificate database will grow each year as more prefixes, ASes, and S-BGP routers are added. We estimate the current database size at about 75-85 Mbytes. The CRL database associated with these certificates adds to this total, but since most of these certificates are issued to organizations and devices (vs. people) and the expected revocation rate should be relatively low and thus CRLs ought not grow large.

S-BGP Processing

The computation burden for signature generation and validation in S-BGP has attracted considerable attention, as well it should. After all, routers today do not process digital signatures and this new burden must be considered carefully. Under normal conditions, UPDATES processing represents a minimal burden for most BGP routers⁵. However, when routes are changing rapidly, the BGP processing load can rise dramatically, and when a BGP router reboots, it receives complete routing tables (via UPDATES) from each of its neighbors. The time required by BGP to process all of these UPDATES represents a significant processing burden. Better algorithms and heuristics are needed to allow routers to better cope with UPDATE surges. Such algorithms should be developed irrespective of the use of S-BGP, but S-BGP would allow these algorithms to operate with confidence about the source and integrity of UPDATES.

In previous analysis, we assumed that each received UPDATE would contain about 3.6 RAs (now updated to 3.7), and would result in transmission of an UPDATE with one new signature. This was an over simplification; a router generates and transmits an UPDATE only if the newly received route is "better" than the current best route, or

⁴ In a discussion with David Mills, an architect of the NSFNET, he confirmed that route flapping has been a characteristic of the Internet since the mid-80's.

⁵ Most subscriber traffic traverses a router via a "fast path" which often uses hardware for path selection. Management traffic, such as BGP, is directed to a general purpose processor and associated memory, which processes the traffic and executes routing algorithms.

if that route is withdrawn by the UPDATE. When a router has many peers, most of the UPDATES it receives will not trigger a change in its view of the best route.

On the other hand, when a router does select a new route, an UPDATE may be constructed and sent to each neighbor, requiring a 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, and thus UPDATES, larger.

This observation 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, i.e., 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 will detect such fraudulent UPDATES and could choose to drop its connection to a peer that behaved this way.

Our initial analysis yielded a peak signature verification rate of about 9/second, for a router with 30 peers, and taking advantage of a depth 1 cache. Given the more thoughtful analysis above, and the more realistic surge UPDATE rates, it is no longer clear what constitutes a good estimate for typical & surge signature validation/generation rates. One could argue for use of a crypto processor to accommodate worst case (200-fold surge) UPDATE rates at a router with many peers. One also could argue that deferring validation unless a received UPDATE would trigger transmission of an UPDATE would reduce the crypto burden to a level that is well within the capabilities of modern, general purpose CPUs. We have not constructed a new analytic model or simulation to evaluate the heuristic.

Initialization/reboot of a BGP router also results a surge in UPDATE processing, and the deferred processing heuristic is applicable here too, even though reboots are relatively infrequent. Saving RIBs in non-volatile storage also addresses this problem.

Transmission bandwidth

Transmission of RAs in UPDATES increases the average size of these messages to about 600 bytes. This is a significant percentage increase (over 800%), but UPDATES represent a very, very small amount of data vs. subscriber traffic. Downloading the certificate, CRL, and AA databases contributes an insignificant increment to this overhead. Full database download, from a repository to an ISP might entail a 75-85 Mbyte file transfer by each ISP. Even if performed more than once a day, these transfers would be swamped by subscriber traffic. Thus the impact on utilization of Internet bandwidth due to transmission of all of the countermeasures data is minimal.

RIB Size

UPDATEs received from neighbors are held in Adj-RIBs and in the Loc-RIB. The space required for RAs is estimated at about 30-35 Mbytes per peer 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 of peers. Thus the management CPU in a router might need up a gigabyte of RAM under some conditions, a modest amount by current workstation standards.

Unfortunately, most currently deployed BGP routers cannot be configured with more than 128M or 256M of RAM; additional RAM would be needed in these routers to support full deployment of S-BGP. Over time it is reasonable to assume that routers could be configured with enough RAM, but this analysis shows that full deployment is not feasible with the currently deployed router base. To add RAM, and possibly to add non-volatile 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.

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. IANA and RIRs must expand 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, subject to the constraint that only neighboring ASes will benefit directly 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 non-contiguous deployment of S-BGP are quite complex, hence our comment that only contiguous deployment is a viable strategy.

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, to avoid route loops within the AS.

Related Work

Any discussion of routing security must include a reference to the first significant treatment of the topic, Radia Perlman's thesis [21].

Several papers on routing security have been published over the last decade, but most deal with “toy” protocols, not with BGP specifically. A number of these papers made suggestions that the techniques they developed would be applicable to BGP, but the assertions proved to be incorrect. Fast signatures based on hash chains have been proposed for this purpose, most recently in [20], but these proposals also have failed to make a solid case for their applicability to BGP.

Some ISPs do make use of a keyed, MD5 integrity check with TCP for BGP transport [16]. This mechanism is less desirable than the use of IPsec in S-BGP, due to its lack of automated key management and operation above the IP layer.

It has been suggested [17] that one could use the DNS and DNSSEC [18] to distribute the information contained in AAs. This mechanism does not address route authorization, nor does the proposal describe in detail how this data would be distributed to BGP routers, and thus it is at best a part of a solution.

Several papers have proposed using Internet Routing Registries [19] or servers operated by ISPs [22] as a basis for distributing data for use in detecting unauthorized route advertisements. The proposals do not address how the accuracy of the information placed in these registries would be verified. The latter proposal suggests that servers operated by ISPs would communicate to verify routes, when routers detect suspicious UPDATES, but this merely creates another path for propagating erroneous data.

Any approach that relies on repositories to propagate routing (vs. origin AS authorization) data will be less dynamic than routing changes, creating problems when route authorizations change quickly, a not uncommon occurrence in response to major outages. Finally, the Internet routing registries (as opposed to RIRs) are “artificial” entities from an authorization perspective, which creates additional concerns.

The most recent proposal in the BGP security arena is soBGP, described in a set of individual Internet Drafts submitted by a team of engineers from Cisco. The name suggests that soBGP focuses on securing origin AS data, but the proposal has evolved to encompass security for AS paths (routes). At this stage, soBGP is not a security architecture for BGP. It is a “Chinese menu” set of components that cannot be analyzed as a system, because it allows a variety of options for various aspects of the protocol, and mandates no choices among these options. Absent such choices, interoperability cannot be assured among ASes, nor can the impact of the system be evaluated. For example, soBGP allows distribution of signed route data via repositories, or in-band (via new BGP protocol extensions). It allows the computation of authorized routes by routers, or by a NOC that distributes the results to the routers in its AS at some unspecified interval. One cannot meaningfully compare soBGP to S-BGP at this time, because the former does not yet reflect choices that permit such comparisons.

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 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 is currently the only complete proposal for addressing BGP security problems. It detects and rejects unauthorized UPDATE messages, irrespective of the means by which they arise, e.g., misconfiguration, active wiretapping, compromise of routers or management systems, etc. S-BGP addresses the timeliness of UPDATE messages only in a limited fashion. S-BGP also does not address an existing, significant problem for BGP routers, i.e., rapid demuxing of management traffic to avoid processor overload. The former problem is a side effect of the lack of such capabilities in BGP itself; the latter is a problem not unique to BGP.

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. Initial funding was provided by NSA, in April of 1997, yielding a first cut design. DARPA provided continued funding, under Dr. Hilarie Orman and Dr. Douglas Maughan, that enabled us to refine, implement and test the design, and to create the current prototype. The author would also like to thank Christine Jones, Charlie Lynn, Joanne Mikkelson, and Karen Seo for their efforts on this project.

References

1. Y. Rekhter, T. Li, "A Border Gateway Protocol 4 (BGP-4)," RFC 1771, March 1995.
2. S. Kent, C. Lynn, and K. Seo, "Secure Boarder Gateway Protocol (S-BGP)," IEEE Journal on Selected Areas in Communications, vol. 18, no. 4, April 2000.
3. C. Villamizar, R. Chandra, R. Govindan., "BGP Route Flap Damping," RFC 2439, November 1998.
4. Smith, B.R. and Garcia-Luna-Aceves, J.J., "Securing the Border Gateway Routing Protocol," Proceedings of Global Internet '96, November 1996.
5. Smith, B.R, Murphy, S., and Garcia-Luna-Aceves, J.J., "Securing Distance-Vector Routing Protocols," Symposium on Network and Distributed System Security, February 1997.

6. Kumar, B., "Integration of Security in Network Routing Protocols," ACM SIGSAC Review, vol.11, no.2, Spring 1993.
7. Murphy, S., panel presentation on "Security Architecture for the Internet Infrastructure," Symposium on Network and Distributed System Security, April 1995.
8. S. Kent, R. Atkinson, "Security Architecture for the Internet Protocol", RFC 2401, November 1998.
9. R. Glenn & S. Kent, "The NULL Encryption Algorithm and its Use with IPsec," RFC 2410, November 1998.
10. S. Kent, R. Atkinson, "IP Encapsulating Security Payload (ESP)," RFC 2406, November 1998.
11. D. Maughan, M. Schertler, M. Schneider, J. Turner, "Internet Security Association and Key Management Protocol (ISAKMP)," RFC 2408, November 1998..
12. D. Harkins, D. Carrel, "The Internet Key Exchange (IKE)," RFC 2406, November 1998.
13. R. Chandra, P. Traina, T. Li, "BGP Communities Attribute", RFC 1997, August 1996.
14. P. Traina, "Autonomous System Confederations for BGP," RFC 1965, June 1996.
15. T. Bates, R. Chandra, D. Katz, Y. Rekhter, "Multiprotocol Extensions for BGP-4," RFC 2283, February 1998.
16. A. Heffernan, "Protection of BGP Sessions via the TCP MD5 Signature Option," RFC 2385, August 1998.
17. T. Bates, R. Bush, T. Li, Y. Rekhter, "DNS-based NLRI origin AS verification in BGP," presentation at NANOG 12, February 1998, <http://www.nanog.org/mtg-9802>.
18. D. Eastlake, 3rd, C. Kaufman, "Domain Name System Security Extensions," RFC 2065, January 1997.
19. C. Alaettinoglu, T. Bates, E. Gerich, D. Karrenberg, D. Meyer, M. Terpstra, C. Villamizar, "Routing Policy Specification Language (RPSL)," RFC 2280, January 1998.
20. Yih-Chun Hu, A. Perrig, and D. Johnson, "Efficient Security Mechanisms for Routing Protocols," Network and Distributed System Security Symposium, February, 2003.
21. R. Perlman, "Network Layer Protocols With Byzantine Robustness," MIT/LCS/TR-429, October, 1988.
22. G. Goodell, W. Aiello, T. Griffin, J. Ioannidis, P. McDaniel, and A. Rubin, "Working Around BGP: An Incremental Approach to Improving Security and Accuracy for Interdomain Routing," Network and Distributed System Security Symposium, February, 2003.
23. J. Ng, "Extensions to BGP to Support Secure Origin BGP (soBGP)," www.ietf.org/internet-drafts/draft-ng-sobgp-bgp-extensions-00.txt.
24. Seo, K., Lynn, C., Kent, S., "Public-Key Infrastructure for the Secure Border Gateway Protocol (S-BGP)," DARPA Information Survivability Conference and Exposition, June 2001.