Threat Model for BGP Path Security
draft-ietf-sidr-bgpsec-threats-07

Abstract

   This document describes a threat model for the context in which
   (E)BGP path security mechanisms will be developed. The threat model
   includes an analysis of the RPKI, and focuses on the ability of an AS
   to verify the authenticity of the AS path info received in a BGP
   update. We use the term PATHSEC to refer to any BGP path security
   technology that makes use of the RPKI. PATHSEC will secure BGP
   [RFC4271], consistent with the inter-AS security focus of the RPKI
   [RFC6480].

   The document characterizes classes of potential adversaries that are
   considered to be threats, and examines classes of attacks that might
   be launched against PATHSEC. It does not revisit attacks against
   unprotected BGP, as that topic has already been addressed in
   [RFC4271]. It concludes with brief discussion of residual
   vulnerabilities.

Status of This Memo

Copyright Notice

Table of Contents

1. Introduction

This document describes the security context in which PATHSEC is
intended to operate. (The term "PATHSEC" is employed in this
document to refer to any design used to achieve the path security
goal described in the SIDR WG charter. The charter focuses on
mechanisms that will enable an AS to determine if the AS_PATH
represented in a route represents the path via which the Network
Layer Reachability Information traveled. Other SIDR documents use
the term "BGPSEC" to refer to a specific design.) It discusses
classes of potential adversaries that are considered to be threats,
and classes of attacks that might be launched against PATHSEC.
Because PATHSEC will rely on the Resource Public Key Infrastructure
(RPKI) [RFC6480], threats and attacks against the RPKI are included.
This model also takes into consideration classes of attacks that are
enabled by the use of PATHSEC (based on the current PATHSEC design.)

The motivation for developing PATHSEC, i.e., residual security
concerns for BGP, is well described in several documents, including
"BGP Security Vulnerabilities Analysis" [RFC4272] and "Design and
Analysis of the Secure Border Gateway Protocol (S-BGP)" [Kent2000].
All of these documents note that BGP does not include mechanisms that
allow an Autonomous System (AS) to verify the legitimacy and
authenticity of BGP route advertisements. (BGP now mandates support
for mechanisms to secure peer-peer communication, i.e., for the links
that connect BGP routers. There are several secure protocol options
to addresses this security concern, e.g., IPsec [RFC4301] and TCP-AO
[RFC5925]. This document briefly notes the need to address this
aspect of BGP security, but focuses on application layer BGP security
issues that must be addressed by PATHSEC.)

RFC 4272 [RFC4272] succinctly notes:

   "BGP speakers themselves can inject bogus routing information,
   either by masquerading as any other legitimate BGP speaker, or by
   distributing unauthorized routing information as themselves.
   Historically, misconfigured and faulty routers have been
   responsible for widespread disruptions in the Internet. The
   legitimate BGP peers have the context and information to produce
   believable, yet bogus, routing information, and therefore have the
   opportunity to cause great damage. The cryptographic protections
   of [TCPMD5] and operational protections cannot exclude the bogus
   information arising from a legitimate peer. The risk of
   disruptions caused by legitimate BGP speakers is real and cannot
   be ignored."

PATHSEC is intended to address the concerns cited above, to provide
significantly improved path security, building upon the route
origination validation capability offered by use of the RPKI
[RFC6810]. Specifically, the RPKI enables relying parties (RPs) to
determine if the origin AS for a path was authorized to advertise the
prefix contained in a BGP update message. This security feature is
enabled by the use of two types of digitally signed data: a PKI
[RFC6487] that associates one or more prefixes with the public key(s)
of an address space holder, and Route Origination Authorizations
(ROAs) [RFC6482] that allows a prefix holder to specify the AS(es)
that are authorized to originate routes for a prefix.

The security model adopted for PATHSEC does not assume an "oracle"
that can see all of the BGP inputs and outputs associated with every
AS or every BGP router. Instead, the model is based on a local
notion of what constitutes legitimate, authorized behavior by the BGP
routers associated with an AS. This is an AS-centric model of secure
operation, consistent with the AS-centric model that BGP employs for
routing. This model forms the basis for the discussion that follows.

This document begins with a brief set of definitions relevant to the
subsequent sections. It then discusses classes of adversaries that
are perceived as viable threats against routing in the public
Internet. It continues to explore a range of attacks that might be
effected by these adversaries, against both path security and the
infrastructure upon which PATHSEC relies. It concludes with a brief
review of residual vulnerabilities, i.e., vulnerabilities that are
not addressed by use of the RPKI and that appear likely to be outside
the scope of PATHSEC mechanisms.

2. Terminology

The following security and routing terminology definitions are
employed in this document.

Adversary - An adversary is an entity (e.g., a person or an
organization) perceived as malicious, relative to the security policy
of a system. The decision to characterize an entity as an adversary
is made by those responsible for the security of a system. Often one
describes classes of adversaries with similar capabilities or
motivations, rather than specific individuals or organizations.

Attack - An attack is an action that attempts to violate the security
policy of a system, e.g., by exploiting a vulnerability. There is
often a many to one mapping of attacks to vulnerabilities, because
many different attacks may be used to exploit a vulnerability.

Autonomous System (AS) - An AS is a set of one or more IP networks
operated by a single administrative entity.

AS Number (ASN) - An ASN is a 2 or 4 byte number issued by a registry
to identify an AS in BGP.

Certification Authority (CA) - An entity that issues digital
certificates (e.g., X.509 certificates) and vouches for the binding
between the data items in a certificate.

Countermeasure - A countermeasure is a procedure or technique that
thwarts an attack, preventing it from being successful. Often
countermeasures are specific to attacks or classes of attacks.

Border Gateway Protocol (BGP) - A path vector protocol used to convey
"reachability" information among autonomous systems, in support of
inter-domain routing.

False (Route) Origination - If a network operator originates a route
for a prefix that the operator does not hold (and that it has not
been authorized to originate by the prefix holder, this is termed
false route origination.

Internet Service Provider (ISP) - An organization managing (and,
typically, selling,) Internet services to other organizations or
individuals.

Internet Number Resources (INRs) - IPv4 or IPv6 address space and
ASNs

Internet Registry - An organization that manages the allocation or
distribution of INRs. This encompasses the Internet Assigned Number
Authority (IANA), Regional Internet Registries (RIRs), National
Internet Registries (NIRs), and Local Internet Registries (LIRs,
network operators).

Man in the Middle (MITM) - A MITM is an entity that is able to
examine and modify traffic between two (or more) parties on a
communication path.

Network Operator - An entity that manages an AS and thus emits (E)BGP
updates, e.g., an ISP.

NOC (Network Operations Center) - A network operator employs a set
equipment and a staff to manage a network, typically on a 24/7 basis.
The equipment and staff are often referred to as the NOC for the
network.

Prefix - A prefix is an IP address and a mask used to specify a set
of addresses that are grouped together for purposes of routing.

Public Key Infrastructure (PKI) - A PKI is a collection of hardware,
software, people, policies, and procedures used to create, manage,
distribute, store, and revoke digital certificates.

Relying Parties (RPs) - An RP is an entity that makes use of signed
products from a PKI, i.e., relies on signed data that is verified
using certificates and Certificate Revocation Lists (CRLs) from a
PKI.

RPKI Repository System - The RPKI repository system consists of a
distributed set of loosely synchronized databases.

Resource PKI (RPKI) - A PKI operated by the entities that manage
INRs, and that issues X.509 certificates (and CRLs) that attest to
the holdings of INRs.

RPKI Signed Object - An RPKI signed object is a Cryptographic Message
Syntax (CMS)-encapsulated data object complying with the format and
semantics defined in [RFC6488].

Route - In the Internet, a route is a prefix and an associated
sequence of ASNs that indicates a path via which traffic destined for
the prefix can be directed. (The route includes the origin AS.)

Route leak - A route leak is said to occur when AS-A advertises
routes that it has received from an AS-B to AS-A's neighbors, but
AS-A is not viewed as a transit provider for the prefixes in the
route.

Threat - A threat is a motivated, capable adversary. An adversary
that is not motivated to launch an attack is not a threat. An
adversary that is motivated but not capable of launching an attack
also is not a threat.

Vulnerability - A vulnerability is a flaw or weakness in a system's
design, implementation, or operation and management that could be
exploited to violate the security policy of a system.

3. Threat Characterization

As noted in Section 2 above, a threat is defined as a motivated,
capable, adversary. The following classes of threats represent
classes of adversaries viewed as relevant to this environment.

Network Operators - A network operator may be a threat. An operator
may be motivated to cause BGP routers it controls to emit update
messages with inaccurate routing info, e.g., to cause traffic to flow
via paths that are economically advantageous for the operator. Such
updates might cause traffic to flow via paths that would otherwise be
rejected as less advantageous by other network operators. Because an
operator controls the BGP routers in its network, it is in a position
to modify their operation in arbitrary ways. Routers managed by a
network operator are vehicles for mounting MITM attacks on both
control and data plane traffic. If an operator participates in the
RPKI, it will have at least one CA resource certificate and may be
able to generate an arbitrary number of subordinate CA certificates
and ROAs. It will be authorized to populate (and may even host) its
own repository publication point. If it implements PATHSEC, and if
PATHSEC makes use of certificates associated with routers or ASes, it
will have the ability to issue such certificates for itself. If
PATHSEC digitally signs updates, it will be able to do so in a
fashion that will be accepted by PATHSEC-enabled neighbors.

Hackers - Hackers are considered a threat. A hacker might assume
control of network management computers and routers controlled by
operators, including operators that implement PATHSEC. In such
cases, hackers would be able to act as rogue network operators (see
above). It is assumed that hackers generally do not have the
capability to effect MITM attacks on most links between networks
(links used to transmit BGP and subscriber traffic). A hacker might
be recruited, without his/her knowledge, by criminals or by nations,
to act on their behalf. Hackers may be motivated by a desire for
"bragging rights" or for profit or to express support for a cause
("hacktivists" [Sam04]).

Criminals - Criminals may be a threat. Criminals might persuade (via
threats or extortion) a network operator to act as a rogue operator
(see above), and thus be able to effect a wide range of attacks.
Criminals might persuade the staff of a telecommunications provider
to enable MITM attacks on links between routers. Motivations for
criminals may include the ability to extort money from network
operators or network operator clients, e.g., by adversely affecting
routing for these network operators or their clients. Criminals also
may wish to manipulate routing to conceal the sources of spam, DoS
attacks, or other criminal activities.

Registries - Any registry in the RPKI could be a threat. Staff at
the registry are capable of manipulating repository content or
mismanaging the RPKI certificates that they issue. These actions
could adversely affect a network operator or a client of a network
operator. The staff could be motivated to do this based on political
pressure from the nation in which the registry operates (see below)
or due to criminal influence (see above).

Nations - A nation may be a threat. A nation may control one or more
network operators that operate in the nation, and thus can cause them
to act as rogue network operators. A nation may have a technical
active wiretapping capability (e.g., within its territory) that
enables it to effect MITM attacks on inter-network traffic. (This
capability may be facilitated by control or influence over a
telecommunications provider operating within the nation.) It may
have an ability to attack and take control of routers or management
network computers of network operators in other countries. A nation
may control a registry (e.g., an RIR) that operates within its
territory, and might force that registry to act in a rogue capacity.
National threat motivations include the desire to control the flow of
traffic to/from the nation or to divert traffic destined for other
nations (for passive or active wiretapping, including DoS).

4. Attack Characterization

This section describes classes of attacks that may be effected
against Internet routing (relative to the context described in
Section 1). Attacks are classified based on the target of the
attack, as an element of the routing system, or the routing security
infrastructure on which PATHSEC relies. In general, attacks of
interest are ones that attempt to violate the integrity or
authenticity of BGP traffic, or which violate the authorizations
associated with entities participating in the RPKI. Attacks that
violate the implied confidentiality of routing traffic, e.g., passive
wiretapping attacks, are not considered a requirement for BGP
security (see [RFC4272]).

4.1. Active wiretapping of sessions between routers

An adversary may attack the BGP (TCP) session that connects a pair of
BGP speakers. An active attack against a BGP (TCP) session can be
effected by directing traffic to a BGP speaker from some remote
point, or by being positioned as a MITM on the link that carries BGP
session traffic. Remote attacks can be effected by any adversary. A
MITM attack requires access to the link. Modern transport networks
may be as complex as the packet networks that utilize them for inter-
AS links. Thus these transport networks may present significant
attack surfaces. Nonetheless, only some classes of adversaries are
assumed to be capable of MITM attacks against a BGP session. MITM
attacks may be directed against BGP, PATHSEC-protected BGP, or
against TCP or IP. Such attacks include replay of selected BGP
messages, selective modification of BGP messages, and DoS attacks
against BGP routers. [RFC4272] describes several countermeasures for
such attacks, and thus this document does not further address such
attacks.

4.2. Attacks on a BGP router

An adversary may attack a BGP router, whether it implements PATHSEC
or not. Any adversary that controls routers legitimately, or that
can assume control of a router, is assumed to be able to effect the
types of attacks described below. Note that any router behavior that
can be ascribed to a local routing policy decision is not considered
to be an attack. This is because such behavior could be explained as
a result of local policy settings, and thus is beyond the scope of
what PATHSEC can detect as unauthorized behavior. Thus, for example,
a router may fail to propagate some or all route withdrawals or
effect "route leaks". (These behaviors are not precluded by the
specification for BGP, and might be the result of a local policy that
is not publicly disclosed. As a result, they are not considered
attacks. See Section 5 for additional discussion.)

Attacks on a router are equivalent to active wiretapping attacks (in
the most general sense) that manipulate (forge, tamper with, or
suppress) data contained in BGP updates. The list below illustrates
attacks of this type.

   AS Insertion: A router might insert one or more ASNs, other than
   its own ASN, into an update message. This violates the BGP spec
   and thus is considered an attack.

   False (Route) Origination: A router might originate a route for a
   prefix, when the AS that the router represents is not authorized
   to originate routes for that prefix. This is an attack, but it is
   addressed by the use of the RPKI [RFC6480].

   Secure Path Downgrade: A router might remove AS_PATH data from a
   PATHSEC-protected update that it receives, when forwarding this
   update to a PATHSEC-enabled neighbor. This behavior violates the
   PATHSEC security goals and thus is considered an attack.

   Invalid AS_PATH Data Insertion: A router might emit a PATHSEC-
   protected update with "bad" data (such as a signature), i.e.,
   PATHSEC data that cannot be validated by other PATHSEC routers.
   Such behavior is assumed to violate the PATHSEC goals and thus is
   considered an attack.

   Stale Path Announcement: If PATHSEC-secured announcements can
   expire, such an announcement may be propagated with PATHSEC data
   that is "expired". This behavior would violate the PATHSEC goals
   and is considered a type of replay attack.

   Premature Path Announcement Expiration: If a PATHSEC-secured
   announcement has an associated expiration time, a router might
   emit a PATHSEC-secured announcement with an expiry time that is
   very short. Unless the PATHSEC protocol specification mandates a
   minimum expiry time, this is not an attack. However, if such a
   time is mandated, this behavior becomes an attack. BGP speakers
   along a path generally cannot determine if an expiry time is
   "suspiciously short" since they cannot know how long a route may
   have been held by an earlier AS, prior to being released.

   MITM Attack: A cryptographic key used for point-to-point security
   (e.g., TCP-AO, TLS, or IPsec) between two BGP routers might be
   compromised (e.g., by extraction from a router). This would
   enable an adversary to effect MITM attacks on the link(s) where
   the key is used. Use of specific security mechanisms to protect
   inter-router links between ASes is outside the scope of PATHSEC.

Compromised Router Private Key: If PATHSEC mechanisms employ
public key cryptography, e.g., to digitally sign data in an
update, then a private key associated with a router or an AS might
be compromised by an attack against the router. An adversary with
access to this key would be able to generate updates that appear
to have passed through the AS that this router represents. Such
updates might be in injected on a link between the compromised
router and its neighbors, if that link is accessible to the
adversary. If the adversary controls another network, it could
use this key to forge signatures that appear to come from the AS
or router(s) in question, with some constraints. So, for example,
an adversary that controls another AS could use a compromised
router/AS key to issue PATHSEC-signed data that include the
targeted AS/router. (Neighbors of the adversary's AS ought not
accept a route that purports to emanate directly from the targeted
AS. So, an adversary could take a legitimate, protected route
that passes through the compromised AS, add itself as the next
hop, and then forward the resulting route to neighbors.)

Withdrawal Suppression Attack: A PATHSEC-protected update may be
signed and announced, and later withdrawn. An adversary
controlling intermediate routers could fail to propagate the
withdrawal. BGP is already vulnerable to behavior of this sort,
so withdrawal suppression is not characterized as an attack, under
the assumptions upon which this mode is based (i.e., no oracle).

4.3. Attacks on network operator management computers (non-CA computers)

An adversary may choose to attack computers used by a network
operator to manage its network, especially its routers. Such attacks
might be effected by an adversary who has compromised the security of
these computers. This might be effected via remote attacks,
extortion of network operations staff, etc. If an adversary
compromises NOC computers, he can execute any management function
that authorized network operations staff would have performed. Thus
the adversary could modify local routing policy to change
preferences, to black-hole certain routes, etc. This type of
behavior cannot be externally detected as an attack. Externally,
this appears as a form of rogue operator behavior. (Such behavior
might be perceived as accidental or malicious by other operators.)

If a network operator participates in the RPKI, an adversary could
manipulate the RP tools that extract data from the RPKI, causing the
output of these tools to be corrupted in various ways. For example,
an attack of this sort could cause the operator to view valid routes
as not validated, which could alter its routing behavior.

If an adversary invoked the tool used to manage the repository
publication point for this operator, it could delete any objects
stored there (certificates, CRLs, manifests, ROAs, or subordinate CA
certificates). This could affect the routing status of entities that
have allocations/assignments from this network operator (e.g., by
deleting their CA certificates).

An adversary could invoke the tool used to request certificate
revocation, causing router certificates, ROAs, or subordinate CA
certificates to be revoked. An attack of this sort could affect not
only this operator, but also any operators that receive allocations/
assignments from it, e.g., because their CA certificates were
revoked.

If an operator is PATHSEC-enabled, an attack of this sort could cause
the affected operator to be viewed as not PATHSEC-enabled, possibly
making routes it emits be less preferred by other operators.

If an adversary invoked a tool used to request ROAs, it could
effectively re-allocate some of the prefixes allocated/assigned to
the network operator (e.g., by modifying the origin AS in ROAs).
This might cause other PATHSEC-enabled networks to view the affected
network as no longer originating routes for these prefixes. Multi-
homed subscribers of this operator who received an allocation from
the operator might find their traffic was now routed via other
connections.

If the network operator is PATHSEC-enabled, and make use of
certificates associated with routers/ASes, an adversary could invoke
a tool used to request such certificates. The adversary could then
replace valid certificates for routers/ASes with ones that might be
rejected by PATHSEC-enabled neighbors.

4.4. Attacks on a repository publication point

A critical element of the RPKI is the repository system. An
adversary might attack a repository, or a publication point within a
repository, to adversely affect routing.

This section considers only those attacks that can be launched by any
adversary who controls a computer hosting one or more repository
publication points, without access to the cryptographic keys needed
to generate valid RPKI signed products. Such attacks might be
effected by an insider or an external threat. Because all repository
objects are digitally signed, attacks of this sort translate into DoS
attacks against the RPKI RPs. There are a few distinct forms of such
attacks, as described below.

Note first that the RPKI calls for RPs to cache the data they acquire and verify from the repository system [RFC6480][RFC6481]. Attacks that delete signed products, that insert products with "bad" signatures, that tamper with object signatures, or that replace newer objects with older (valid) ones, can be detected by RPs (with a few exceptions). RPs are expected to make use of local caches. If repository publication points are unavailable or the retrieved data is corrupted, an RP can revert to using the cached data. This behavior helps insulate RPs from the immediate effects of DoS attacks on publication points.

Each RPKI data object has an associated date at which it expires, or is considered stale. (Certificates expire, CRLs become stale.) When an RP uses cached data it is a local decision how to deal with stale or expired data. It is common in PKIs to make use of stale certificate revocation status data, when fresher data is not available. Use of expired certificates is less common, although not unknown. Each RP will decide, locally, whether to continue to make use of or ignore cached RPKI objects that are stale or expired.

If an adversary inserts an object into a publication point, and the object has a "bad" signature, the object will not be accepted and used by RPs.

If an adversary modifies any signed product at a publication point, the signature on the product will fail, causing RPs to not accept it. This is equivalent to deleting the object, in many respects.

If an adversary deletes one or more CA certificates, ROAs or the CRL for a publication point, the manifest for that publication point will allow an RP to detect this attack. An RP can continue to use the last valid instance of the deleted object (as a local policy option), thus minimizing the impact of such an attack.

If an adversary deletes a manifest (and does not replace it with an older instance), that is detectable by RPs. Such behavior should result in the CA (or publication point maintainer) being notified of the problem. An RP can continue to use the last valid instance of the deleted manifest (a local policy option), thus minimizing the impact of such an attack.

If an adversary deletes newly added CA certificates or ROAs, and replaces the current manifest with the previous manifest, the manifest (and the CRL that it matches) will be "stale" (see [RFC6486]). This alerts an RP that there may be a problem. The RP should use the information from a Ghostbuster record [RFC6493] to contact the entity responsible for the publication point, requesting that entity to remedy the problem (e.g., republish the missing CA

certificates and/or ROAs). An RP cannot know the content of the new
certificates or ROAs that are not present, but it can continue to use
what it has cached. An attack of this sort will, at least
temporarily, cause RPs to be unaware of the newly published objects.
INRs associated with these objects will be treated as
unauthenticated.

If a CA revokes a CA certificate or a ROA (via deleting the
corresponding EE certificate), and the adversary tries to reinstate
that CA certificate or ROA, the adversary would have to rollback the
CRL and the manifest to undo this action by the CA. As above, this
would make the CRL and manifest stale, and this is detectable by RPs.
An RP cannot know which CA certificates or ROAs were deleted.
Depending on local policy, the RP might use the cached instances of
the affected objects, and thus be tricked into making decisions based
on these revoked objects. Here too the goal is that the CA will be
notified of the problem (by RPs) and will remedy the error.

In the attack scenarios above, when a CRL or manifest is described as
stale, this means that the next issue date for the CRL or manifest
has passed. Until the next issue date, an RP will not detect the
attack. Thus it behooves CAs to select CRL/manifest lifetimes (the
two are linked) that represent an acceptable trade-off between risk
and operational burdens.

Attacks effected by adversaries that are legitimate managers of
publication points can have much greater effects, and are discussed
below under attacks on or by CAs.

4.5. Attacks on an RPKI CA

Every entity to which INRs have been allocated/assigned is a CA in
the RPKI. Each CA is nominally responsible for managing the
repository publication point for the set of signed products that it
generates. (An INR holder may choose to outsource the operation of
the RPKI CA function, and the associated publication point. In such
cases, the organization operating on behalf of the INR holder becomes
the CA, from an operational and security perspective. The following
discussion does not distinguish such outsourced CA operations.)

Note that attacks attributable to a CA may be the result of malice by
the CA (i.e., the CA is the adversary) or they may result from a
compromise of the CA.

All of adversaries listed in Section 2 are presumed to be capable of
launching attacks against the computers used to perform CA functions.
Some adversaries might effect an attack on a CA by violating
personnel or physical security controls as well. The distinction

between CA as adversary vs. CA as an attack victim is important.
Only in the latter case should one expect the CA to remedy problems
caused by a attack once the attack has been detected. (If a CA does
not take such action, the effects are the same as if the CA is an
adversary.)

Note that most of the attacks described below do not require
disclosure of a CA's private key to an adversary. If the adversary
can gain control of the computer used to issue certificates, it can
effect these attacks, even though the private key for the CA remains
"secure" (i.e., not disclosed to unauthorized parties). However, if
the CA is not the adversary, and if the CA's private key is not
compromised, then recovery from these attacks is much easier. This
motivates use of hardware security modules to protect CA keys, at
least for higher tiers in the RPKI.

An attack by a CA can result in revocation or replacement of any of
the certificates that the CA has issued. Revocation of a certificate
should cause RPs to delete the (formerly) valid certificate (and
associated signed object, in the case of a revoked EE certificate)
that they have cached. This would cause repository objects (e.g., CA
certificates and ROAs) that are verified under that certificate to be
considered invalid, transitively. As a result, RPs would not
consider as valid any ROAs or PATHSEC-protected updates based on
these certificates, which would make routes dependent on them to be
less preferred. Because a CA that revokes a certificate is
authorized to do so, this sort of attack cannot be detected,
intrinsically, by most RPs. However, the entities affected by the
revocation or replacement of CA certificates can be expected to
detect the attack and contact the CA to effect remediation. If the
CA was not the adversary, it should be able to issue new certificates
and restore the publication point.

An adversary that controls the CA for a publication point can publish
signed products that create more subtle types of DoS attacks against
RPs. For example, such an attacker could create subordinate CA
certificates with Subject Information Access (SIA) pointers that lead
RPs on a "wild goose chase" looking for additional publication points
and signed products. An attacker could publish certificates with
very brief validity intervals, or CRLs and manifests that become
"stale" very quickly. This sort of attack would cause RPs to access
repositories more frequently, and that might interfere with
legitimate accesses by other RPs.

An attacker with this capability could create very large numbers of
ROAs to be processed (with prefixes that are consistent with the
allocation for the CA), and correspondingly large manifests. An
attacker could create very deep subtrees with many ROAs per

publication point, etc. All of these types of DoS attacks against
RPs are feasible within the syntactic and semantic constraints
established for RPKI certificates, CRLs, and signed objects.

An attack that results in revocation and replacement (e.g., key
rollover or certificate renewal) of a CA certificate would cause RPs
to replace the old, valid certificate with the new one. This new
certificate might contain a public key that does not correspond to
the private key held by the certificate subject. That would cause
objects signed by that subject to be rejected as invalid, and prevent
the affected subject from being able to sign new objects. As above,
RPs would not consider as valid any ROAs issued under the affected CA
certificate, and updates based on router certificates issued by the
affected CA would be rejected. This would make routes dependent on
these signed products to be less preferred. However, the constraints
imposed by the use of RFC 3779 [RFC3779] extensions do prevent a
compromised CA from issuing (valid) certificates with INRs outside
the scope of the CA, thus limiting the impact of the attack.

An adversary that controls a CA could issue CA certificates with
overlapping INRs to different entities, when no transfer of INRs is
intended. This could cause confusion for RPs as conflicting ROAs
could be issued by the distinct (subordinate) CAs.

An adversary could replace a CA certificate, use the corresponding
private key to issue new signed products, and then publish them at a
publication point controlled by the attacker. This would effectively
transfer the affected INRs to the adversary, or to a third party of
his choosing. The result would be to cause RPs to view the entity
that controls the private key in question as the legitimate INR
holder. Again the constraints imposed by the use of RFC 3779
extensions prevent a compromised CA from issuing (valid) certificates
with INRs outside the scope of the CA, thus limiting the impact of
the attack.

Finally, an entity that manages a repository publication point can
inadvertently act as an attacker (an example of Walt Kelly's most
famous "Pogo" quote [Kelly70]). For example, a CA might fail to
replace its own certificate in a timely fashion (well before it
expires). If might fail to issue its CRL and manifest prior to
expiration, creating stale instances of these products that cause
concern for RPs. A CA with many subordinate CAs (e.g., an RIR or
NIR) might fail to distribute the expiration times for the CA
certificates that it issues. A network with many ROAs might do the
same for the EE certificates associated with the ROAs it generates.
A CA could rollover its key, but fail to reissue subordinate CA
certificates under its new key. Poor planning with regard to rekey
intervals for managed CAs could impose undue burdens for RPs, despite

a lack of malicious intent. All of these example of mismanagement
could adversely affect RPs, despite the absence of malicious intent.

5. Residual Vulnerabilities

The RPKI, upon which PATHSEC relies, has several residual
vulnerabilities that were discussed in the preceding text
(Section 4.4 and Section 4.5). These vulnerabilities are of two
principle forms:

o the RPKI repository system may be attacked in ways that make its
   contents unavailable, not current, or inconsistent. The principle
   defense against most forms of DoS attacks is the use of a local
   cache by each RP. The local cache ensures availability of
   previously-acquired RPKI data, in the event that a repository is
   inaccessible or if repository contents are deleted (maliciously).
   Nonetheless, the system cannot ensure that every RP will always
   have access to up-to-date RPKI data. An RP, when it detects a
   problem with acquired repository data has two options:

   1. The RP may choose to make use of its local cache, employing
       local configuration settings that tolerate expired or stale
       objects. (Such behavior is, nominally, always within the
       purview of an RP in PKI.) Using cached, expired or stale data
       subjects the RP to attacks that take advantage of the RP's
       ignorance of changes to this data.

   2. The RP may chose to purge expired objects. Purging expired
       objects removes the security info associated with the real
       world INRs to which the objects refer. This is equivalent to
       the affected INRs not having been afforded protection via the
       RPKI. Since use of the RPKI (and PATHSEC) is voluntary, there
       may always be set of INRs that are not protected by these
       mechanisms. Thus purging moves the affected INRs to the set
       of non-participating INR holders. This more conservative
       response enables an attacker to move INRs from the protected
       to the unprotected set.

o any CA in the RPKI may misbehave within the bounds of the INRs
   allocated to it, e.g., it may issue certificates with duplicate
   resource allocations or revoke certificates inappropriately. This
   vulnerability is intrinsic in any PKI, but its impact is limited
   in the RPKI because of the use of RFC 3779 extensions. It is
   anticipated that RPs will deal with such misbehavior through
   administrative means, once it is detected.

PATHSEC has a separate set of residual vulnerabilities:

o It has been stated that "route leaks" are viewed as a routing
   security problem by many operators. However, BGP itself does not
   include semantics that preclude what many perceive as route leaks,
   and there is no definition of the term in any RFC. This makes it
   inappropriate to address route leaks in this document.
   Additionally, route leaks are outside the scope of PATHSEC, based
   on the SIDR charter. That charter focuses on the integrity and
   authenticity of the data contained in the AS_path. If, at a later
   time, the SIDR charter is amended to include route leaks, and an
   appropriate definition exists, this document should be revised.

o PATHSEC is not required to protect all attributes associated with
   an AS_PATH, even though some of these attributes may be employed
   as inputs to routing decisions. Thus attacks that modify (or
   strip) these other attributes are not prevented/detected by
   PATHSEC. The SIDR charter calls for protecting only the info
   needed to verify that a received route traversed the ASes in
   question, and that the NLRI in the route is what was advertised.
   (The AS_PATH data also may have traversed ASes within a
   confederation that are not represented. However, these ASes are
   not externally visible, and thus do not influence route selection,
   so their omission in this context is not a security concern.)
   Thus, protection of other attributes is outside the scope of the
   charter, at the time this document was prepared. If, at a later
   time, the SIDR charter is amended to include protection of
   additional BGP attributes, this document should be revised.

o PATHSEC cannot ensure that an AS will withdraw a route when the AS
   no longer has a route for a prefix, as noted in Section 4.2.
   PATHSEC may incorporate features to limit the lifetime of an
   advertisement. Such lifetime limits provide an upper bound on the
   time that the failure to withdraw a route will remain effective.

6. Security Considerations

A threat model is, by definition, a security-centric document.
Unlike a protocol description, a threat model does not create
security problems nor purport to address security problems. This
model postulates a set of threats (i.e., motivated, capable
adversaries) and examines classes of attacks that these threats are
capable of effecting, based on the motivations ascribed to the
threats. It describes the impact of these types of attacks on
PATHSEC, including on the RPKI on which PATHSEC relies. It describes
how the design of the RPKI (and the PATHSEC design goals) address
classes of attacks, where applicable. It also notes residual
vulnerabilities.

7. IANA Considerations

   [Note to IANA, to be removed prior to publication: there are no IANA
   considerations stated in this version of the document.]

8. Acknowledgements

   TBD

9. Informative References

   [Kelly70] Kelly, W., "'We Have Met the Enemy, and He is Us': Pogo
              Earth Day Poster", April 1970.

   [Kent2000]
              Kent, S., Lynn, C., and K. Seo, "Design and Analysis of
              the Secure Border Gateway Protocol (S-BGP)", IEEE DISCEX
              Conference, June 2000.

   [RFC3779] Lynn, C., Kent, S., and K. Seo, "X.509 Extensions for IP
              Addresses and AS Identifiers", RFC 3779, June 2004.

   [RFC4271] Rekhter, Y., Li, T., and S. Hares, "A Border Gateway
              Protocol 4 (BGP-4)", RFC 4271, January 2006.

   [RFC4272] Murphy, S., "BGP Security Vulnerabilities Analysis", RFC
              4272, January 2006.

   [RFC4301] Kent, S. and K. Seo, "Security Architecture for the
              Internet Protocol", RFC 4301, December 2005.

   [RFC5925] Touch, J., Mankin, A., and R. Bonica, "The TCP
              Authentication Option", RFC 5925, June 2010.

   [RFC6480] Lepinski, M. and S. Kent, "An Infrastructure to Support
              Secure Internet Routing", RFC 6480, February 2012.

   [RFC6481] Huston, G., Loomans, R., and G. Michaelson, "A Profile for
              Resource Certificate Repository Structure", RFC 6481,
              February 2012.

   [RFC6482] Lepinski, M., Kent, S., and D. Kong, "A Profile for Route
              Origin Authorizations (ROAs)", RFC 6482, February 2012.

   [RFC6486] Austein, R., Huston, G., Kent, S., and M. Lepinski,
              "Manifests for the Resource Public Key Infrastructure
              (RPKI)", RFC 6486, February 2012.

   [RFC6487]   Huston, G., Michaelson, G., and R. Loomans, "A Profile for
               X.509 PKIX Resource Certificates", RFC 6487, February
               2012.

   [RFC6488]   Lepinski, M., Chi, A., and S. Kent, "Signed Object
               Template for the Resource Public Key Infrastructure
               (RPKI)", RFC 6488, February 2012.

   [RFC6493]   Bush, R., "The Resource Public Key Infrastructure (RPKI)
               Ghostbusters Record", RFC 6493, February 2012.

   [RFC6810]   Bush, R. and R. Austein, "The Resource Public Key
               Infrastructure (RPKI) to Router Protocol", RFC 6810,
               January 2013.

   [Sam04]   Samuel, A., "Hacktivism and the Future of Political
               Participation", Ph.D. dissertation, Harvard University,
               August 2004.

Authors' Addresses

   Stephen Kent
   BBN Technologies
   10 Moulton St.
   Cambridge, MA 02138
   US

   Email: kent@bbn.com

   Andrew Chi
   University of North Carolina - Chapel Hill
   c/o Department of Computer Science
   CB 3175, Sitterson Hall
   Chapel Hill, NC 27599
   US

   Email: achi@cs.unc.edu