BGP Security Requirements

An Overview of the Current Work in the IETF

BGP Security Requirements

- Scope
- Deployment Requirements
- Trust Models
- The AS-Path and NLRIs
- Address Allocation and Advertisement
- Logging/Tracking
- Transport Protection
- Current Proposed Systems

Scope

- **The RPsec WG within the IETF is currently documenting BGP security requirements.**
- **In Scope:**

  Originating False Data, Routing Database Integrity, Peering Integrity
- **Out of Scope:**

  Any attack where BGP isn't directly manipulated, Data Packet Delivery

Scope

- **Inform any BGP security mechanisms designed and proposed.**
- **Provide a set of baseline requirements any proposed security system can be judged against.**
- **Don’t try to define perfect security, but rather balance deployment and security.**
- **Consider what works in the real world.**

Deployment Requirements

- **SHOULD NOT Slow Down BGP Convergence**
- **This may be an “impossible” requirement, but it’s a goal the community *should* strive for.**
- **It’s a SHOULD, rather than a MUST.**
- **Routing convergence speed is critical**
- **BGP on-the-wire optimizations, such as packing, SHOULD NOT be impacted by security mechanisms. Verification of routing information SHOULD be real time, but MAY be periodic.**

Deployment Requirements

- **MUST Be Incrementally Deployable**

  Elements

  - MUST be able to handle secure and unsecure routes in the same way.
  - MUST allow BGP speakers running in an unsecure environment to peer with speakers running in a secure environment.
  - MUST use backward compatible message formatting, etc.
  - SHOULD allow a BGP speaker to tell the difference between an altered secure route and an unsecure route.

Deployment Requirements

- **MUST Provide Local Trust Decision Point**

  Trust model MUST provide local policy implementation for routing information authenticity.

  Follows the current model of local policy about information received and acted on from external peers.

  Leaves final decision on how to act in reference to learned security information up to the local AS.

  Allows for variability in different environments and internetworks in reference to security levels and policies.

The Trust Model

- **MUST Support A Distributed Trust Model**

  The optimal trust model may vary.

  Military/government may find a strictly hierarchical trust model more appropriate.

  Internetworks under a single overall entity may find a strictly hierarchical model more appropriate.

  Large scale internetworks built on a contractual basis may find a distributed trust model more appropriate.

The Trust Model (Cont.)

- A strict hierarchy is a subset of a distributed trust model.
  - This single trust model encompasses several modes of operation.
  - Any proposed solution should be able to support a number of deployments if it supports a distributed trust model.

Routing Information Validation

- **MUST Verify Origin AS' Authorization to Advertise**

  The trust model is critical in this area.

  In some environments, only a strict hierarchy of address allocation will be acceptable, in others, a web of trust, or distributed trust, may work better.

Routing Information Validation

- **MUST verify the AS path corresponds to a valid path in the Internetwork**

  Does {4, 3, 2, 1} exist and did (or could) the advertisement traverse it? Note a relation to non-repudiation here.

- **MUST ensure the first element of the AS path is the same as the transmitting peer's AS**

  A must make certain 2 is the first AS on the AS Path when accepting an update from B

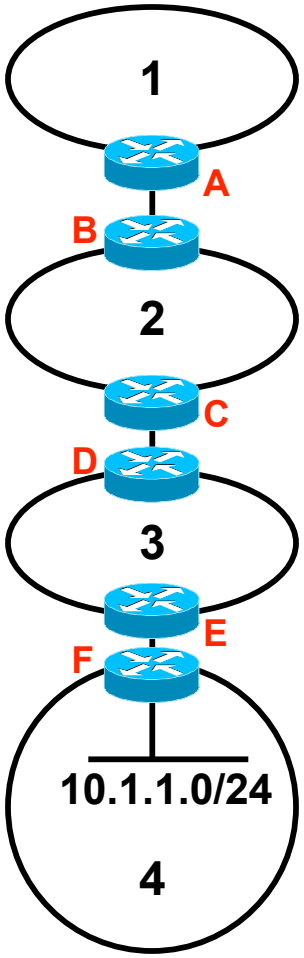

Logging/Tracking

- **SHOULD Provide Non-Repudiation of Updates**

  The receiver should be able to verify who originated a specific update, and track the update through the internetwork.

- **MUST Provide for Logging**

  Must be in a “standard format.” And is an essential part of any security mechanism

  Whether logging should be directly specified in the proposed security mechanism is currently under discussion.

Transport Protection

- **MUST Include Transport Protection Between BGP Speakers**

  MAY reference systems already in existence to solve this problem.

  SHOULD use the same keys for security throughout the proposed security solution (one key per AS is preferred).

Questions

- **How Should the IETF Handle Proposed Security Mechanisms?**

  Form a design team to build a system using existing proposals as references, and filtering it through the requirements?

  Advance any proposed system meeting the requirements laid out in the requirements draft to experimental, and revisit based on actual deployment in several years?

  Other options?