BGPSEC

Potential Optimizations for AS-PATH Prepending and Transparent Route Servers.

sidr wg / Québec City

2011.07.28

Doug Montgomery <dougm@nist.gov>
Randy Bush <randy@psg.com>

bgpsec-00 Level Set

- **Requirements**: Provide cryptographic assurance that:
  - Origin AS was authorized by IP holder to announce route.
  - Every AS in the AS_Path explicitly authorized the advertisement of the route to the subsequent AS in the AS_Path.
- **Semantics**:

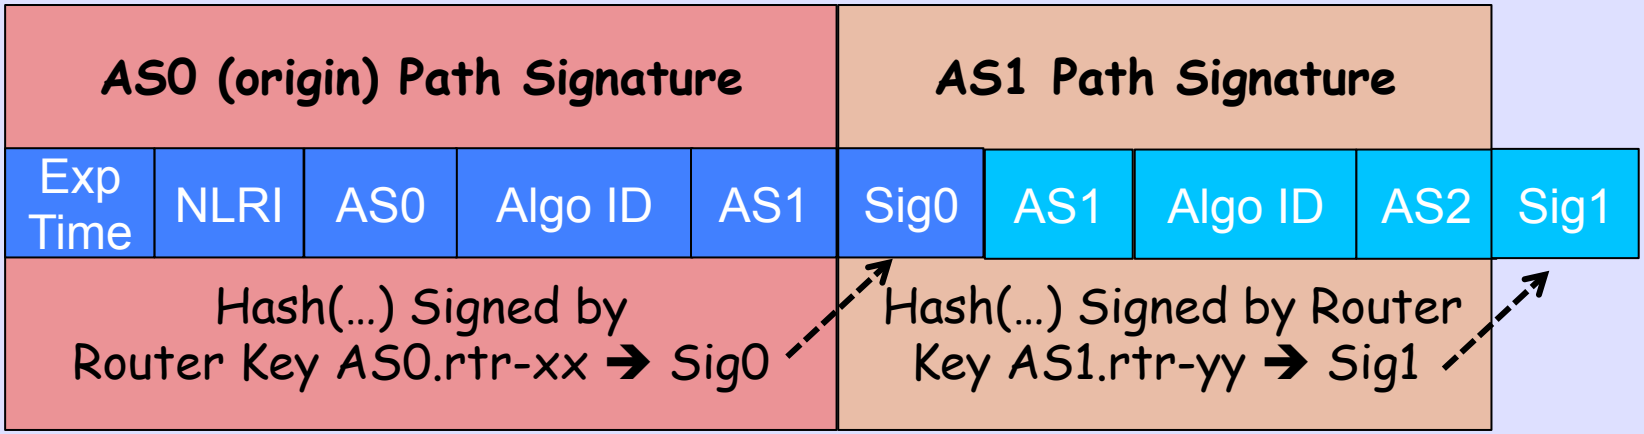

  - Each AS's Signature computed over Target AS for the update (forward chaining of path) and previous AS Sig (backward chaining of authorization).
  - Expire Time & beaconing subject of another presentation …..

bgpsec-00 Path_Sigs

- **Syntax**:
  - BGPSEC Path_Signatures attribute elements correspond 1-1 with AS_Path attribute elements.
  - AS_Path attribute umodified, nor is Path data (ASNs) replicated in Path_Sig.

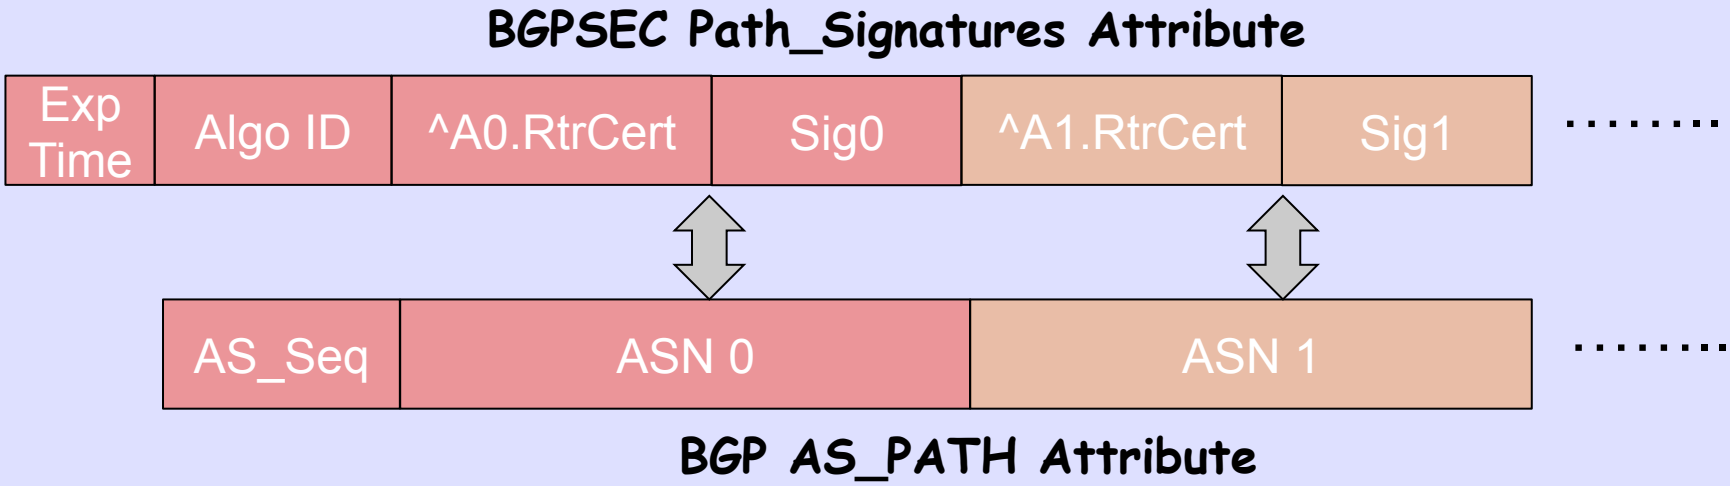

- **bgpsec-00 – focus on requirements and semantics – purposefully ignore syntax optimizations until we get the first two right.**
  - **See: draft-sriram-bgpsec-design-choices-00**

Optimizations / Enhancements

- **From recent WG discussions**:
  - **AS_Path Prepending** – current 1-1 correspondence of elements would require repeating signatures when using Path prepending.
  - **Transparent Route Servers**: - AS_Path does not reflect the actual sequence of AS's that the update traversed.
- **Going Forward**:
  - Important to separate the requirement, semantics and syntax discussions of how we address these and future enhancements.
    - **Requirement?** – ensure BGPSEC doesn't interfere with some current/proposed use case, or enhance BGPSEC to protect use case?
  - Depending upon which BGP services / capabilities / uses we are trying to preserve / protect there are numerous ways to spec solutions.
- **Strawman Approaches in response to recent discussions …**

Prepending Strawman

- **Requirements:**
  - Support current uses of prepending without incurring expense of repeating Path_Signature elements.
  - Use BGPSEC to protect prepended AS's from modification.
- **Semantics:**
  - Prepend Count (pCNT) included as input to ASx signature noting how many times ASx appears in the actual AS_PATH.
  - Normally pCNT=1.

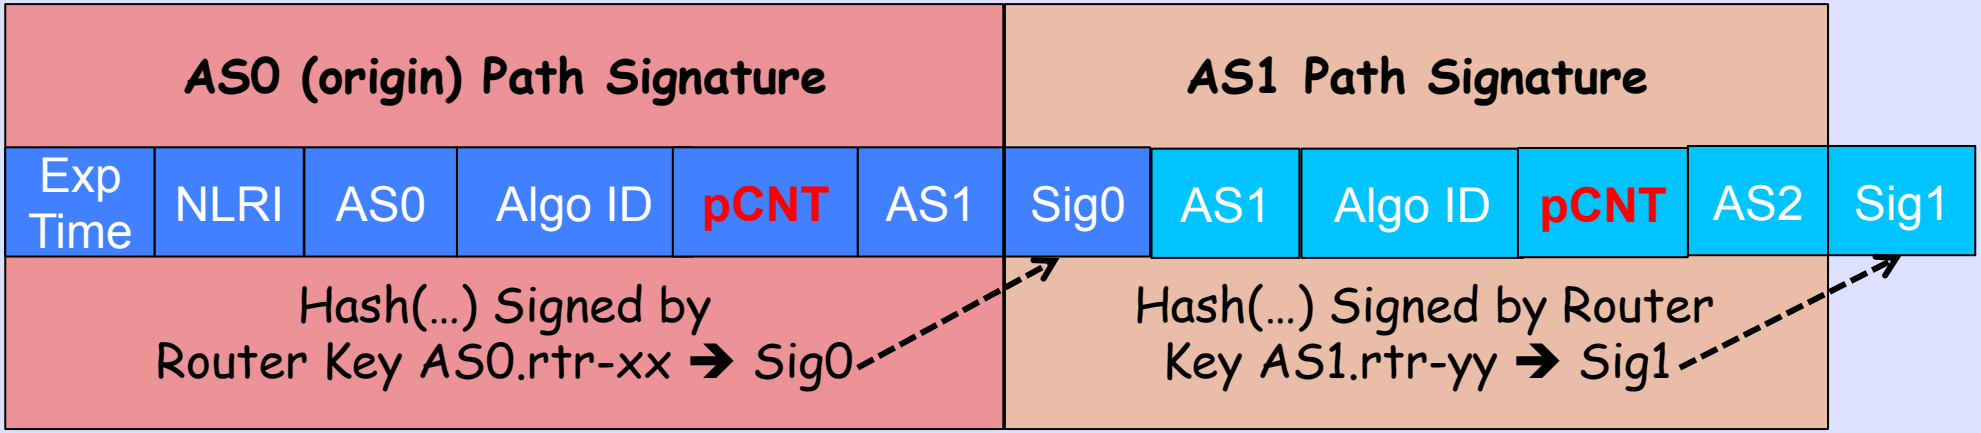

Transparent Route Servers

- **Use Case:**
  - Multi-Lateral Peering where sender does not know all the receivers.
  - Router Server's AS not included in AS_PATH so as to not contribute to the Path Length for purposes of downstream best path computation.

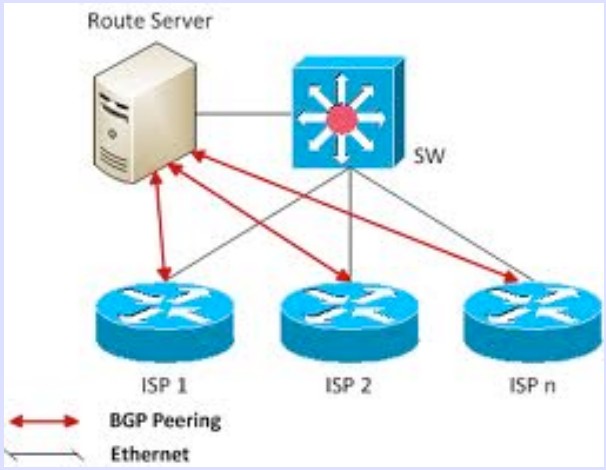

- **BGPSEC Issue:**
  - BGPSEC speaker can't forward sign to other RS customers – not peering with them directly, may not even explicitly know them, defeats the transmission efficiency of RS architectures.
  - Total transparency – violates the fundamental service of BGPSEC to provide a cryptographically verified sequence of route authorizations.

"Translucent" RS Strawman

- **Requirements:**
  - Support current business/use model of transparent RS's.
  - Use BGPSEC to protect / verify the complete AS_PATH (including RS AS) without impacting AS_Path_Length computations.
  - Be completely transparent when sending updates to non-BGPSEC peers.
- **Semantics**:
  - RS AS fully participates in BGPSEC to/from its customer AS's, including explicitly carrying the RS AS# in the update.
  - Use previously proposed Prepend Count (pCNT=0) to indicate that an AS is operating as a transparent RS.
  - BGPSEC validates complete AS_PATH, but pCNT=0 hops do not contribute to Path_Length.
  - When sending to non-BGPSEC peer, RS AS# is stripped.

Further Details …

- **Not important now … unless we agree on the requirements/semantics:**
  - … if we do agree on the basic approach, then we can consider further details.
  - Tradeoffs between optimization and minimizing BGPSEC changes to current BGP attributes / behaviors.
- **Strawman Syntax:**
  - Extend BGPSEC Path_Signatures Field to carry pCNT for each Sig.
    - Modify Sig generation / verification procedures to address pCNT.
    - Modify Path_Length computation on BGPSEC implementations to ignore pCNT=0 hops.
  - Add rules to strip RS AS when sending to non-BGPSEC peer.