Rethinking Path Validation

Russ White

Reality Check

- Right now there is no US Government mandate to do anything
- A mandate in the origin authentication area is probably immanent
- A mandate in the path validation space will probably happen eventually
- Are we happy with the options we have?

Origin Authentication

- AS65002 authorized to originate 2001:db8:0:1::/64
- AS65002 creates an RC signed with a private key and any additional parameters
- AS65002 places this in the RPKI database

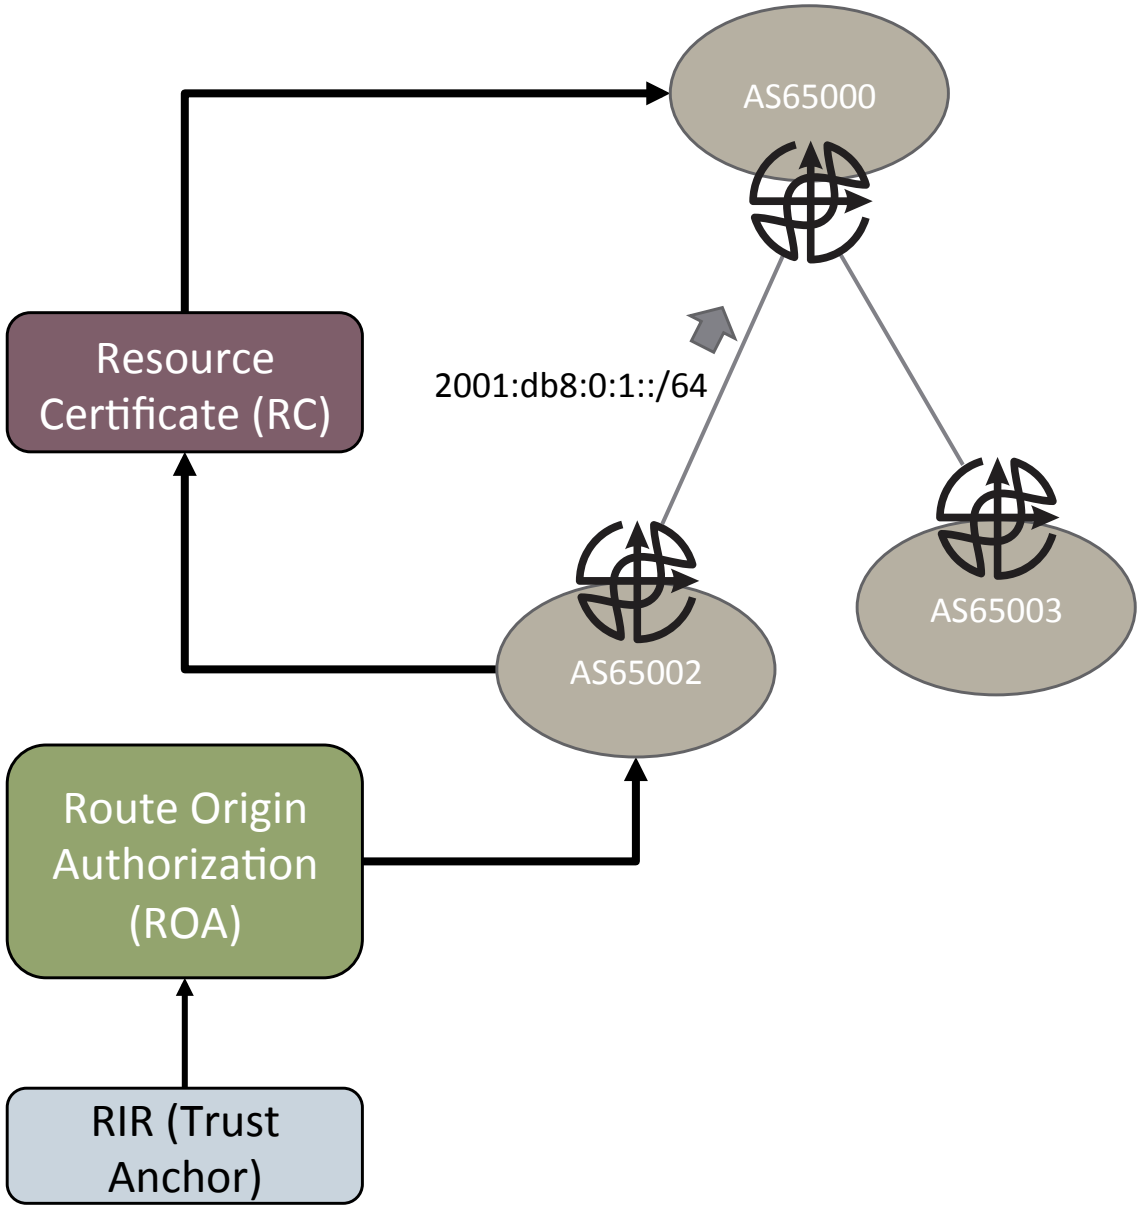

Origin Authentication

- AS65000 uses AS65002's public key to validate the ROA
- AS65000 can check the original authorization using the trust anchor's public key

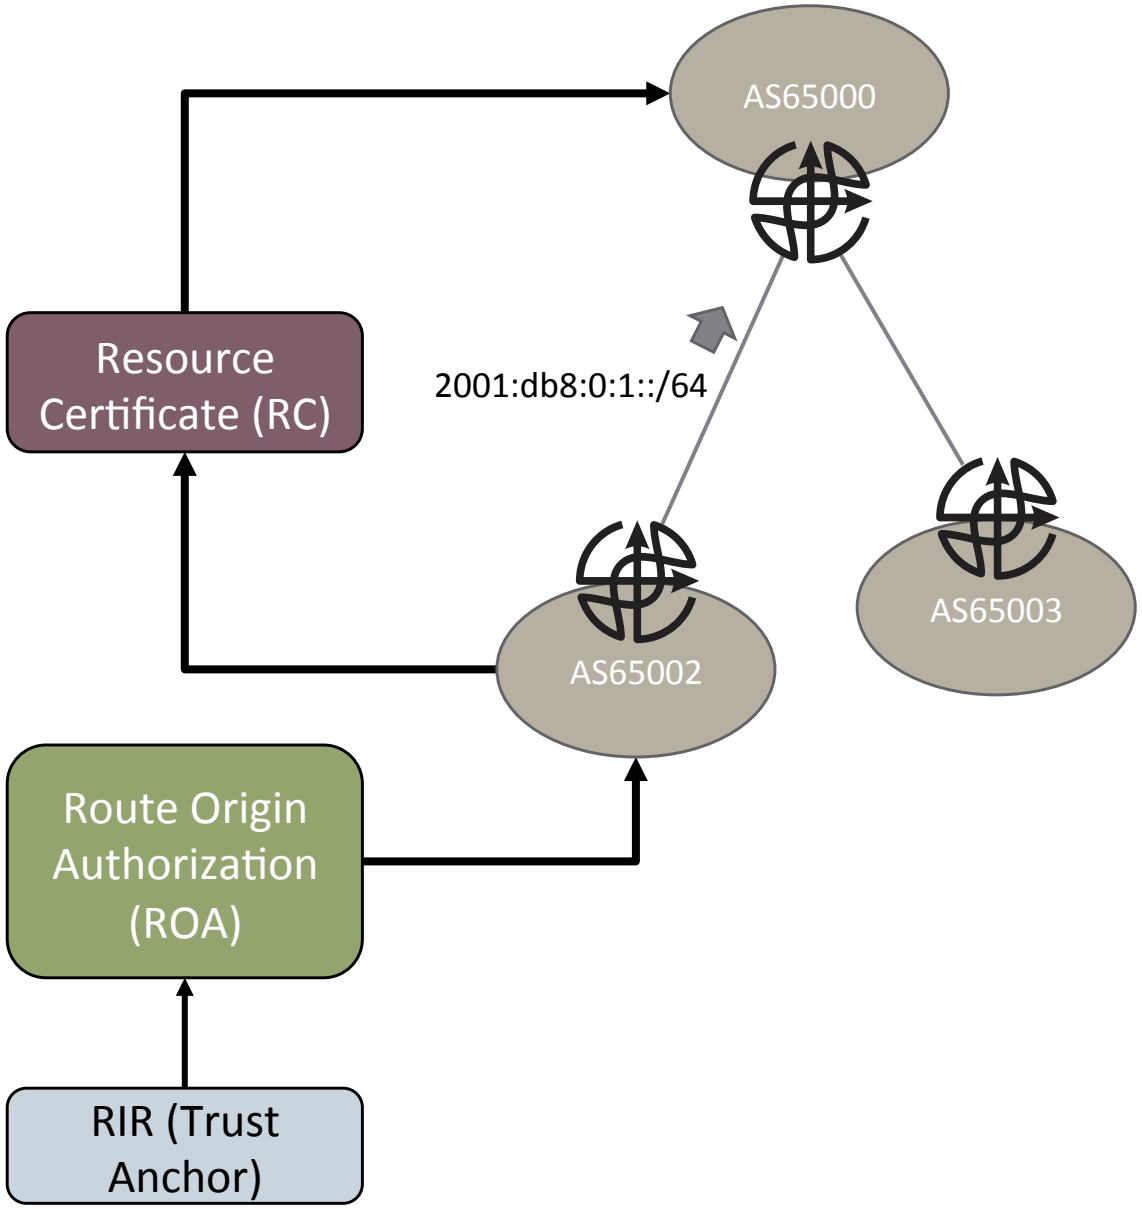

Origin Authentication

- AS65003 can advertise 192.0.2.0/24 with the AS Path [65002,65003]
- AS65000 will be none the wiser…
- To resolve this, *path validation* of some sort is needed

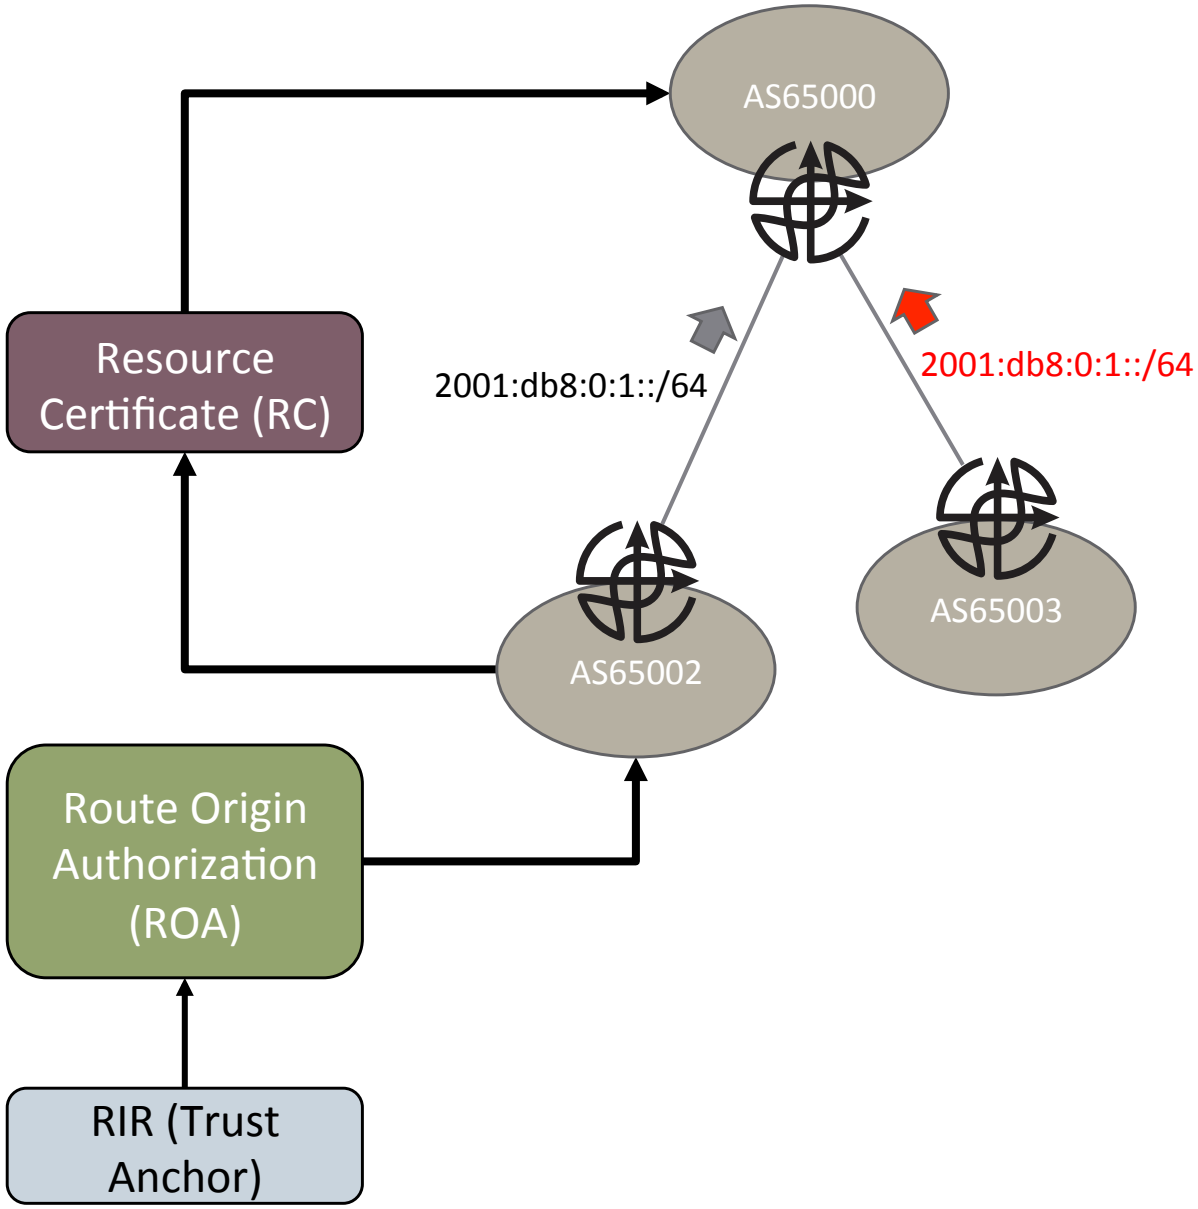

Rethinking Requirements

- Reuse BGP – trusted and understood
  - Address family (AF) or new message
  - No reason to reuse current bestpath for this application
- Reuse existing policy mechanisms if possible
- Don’t mess with origin authentication (in general)
  - *Allow* replacing rsync with BGP transport

Notes

- Current bestpath in this context means current metrics, like MED, Local Pref, etc.
  - These don’t seem to apply to carrying certificates
  - A new AF can define its own metrics and bestpath algorithm
- Existing policy mechanisms primarily means communities in this context
  - Provide a common context for reachability and path security information
  - Provide a common policy that ties reachability and path security information
- There are concerns about the long term viability of rsync in this application
  - If we design the AF correctly, we can carry the current ROAs as well
  - Optional, but potentially useful; leave open for further discussion in the community

Rethinking Requirements

- Solve 80% of the problem space in a deployable way
  - Assume to be used in parallel with other mechanisms
  - Stateful inspection/IDS pair (separate baskets)
  - Don’t make the edge do crypto
- Persistence in the face of DDoS

Notes

- Any single mechanism probably isn’t going to solve every problem
  - If every problem can and should actually be solved at all
  - Think of a stateful packet filter (firewall/SPF) combined with an Intrusion Detection System (IDS)
  - The SPF doesn’t really catch every possible attack
  - Instead, we put in different systems to solve different parts of the problem
  - Given this, we should focus on solving “80%” of the problems
  - For instance, data analytics used across the table in near real time, in combination with DNS and traffic flow analysis, can probably catch some attacks or security problems more easily than a purely BGP based path validation system of any kind

Rethinking Requirements

- Hide things that aren't otherwise available
- Control where information is advertised
- Optionally attach peering types and other policy to specific relationships

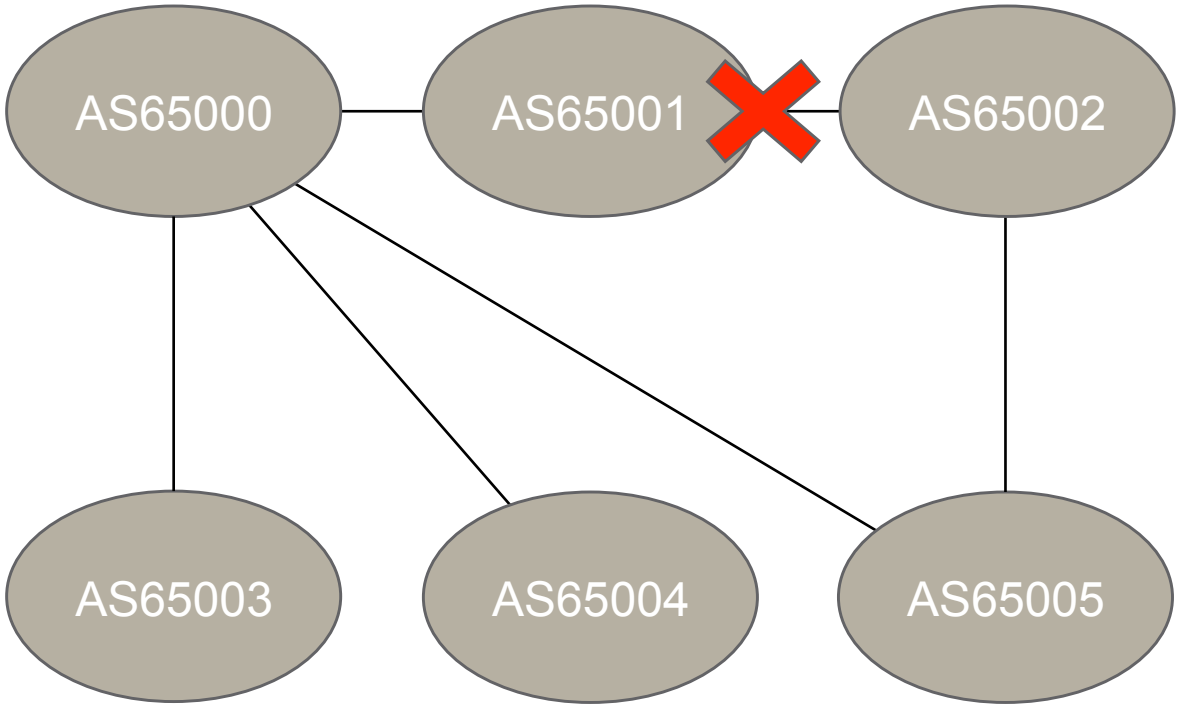

Notes

- AS65000 doesn’t want to advertise it’s connection to AS65003 unless the routes are being advertised
  - Backup routes, etc.
- AS65000 only wants its connection to AS65004 advertised to its peers, and not to their peers
  - Regional routing information, partnering relationships, etc.
- AS65000 wants to make certain other AS’ know that AS65005 is not a transit customer
  - So other AS’ should not see routes AS65000 advertises to AS65005 readvertised

Rethinking Requirements

- Overlay carrying new information
- Incremental deployment should add value incrementally

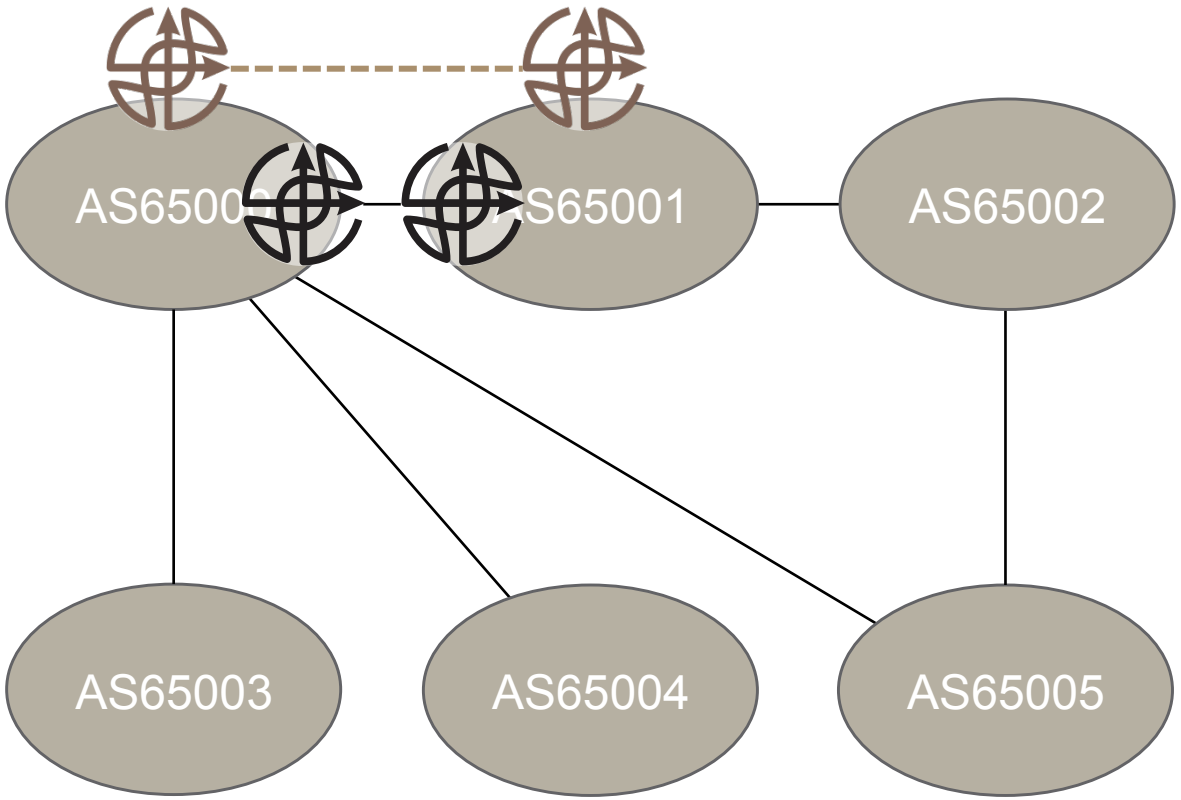

Conceptually

- AS Level semantics
- Only AS level changes are reflected in the base advertisements
- More detail *may* be included

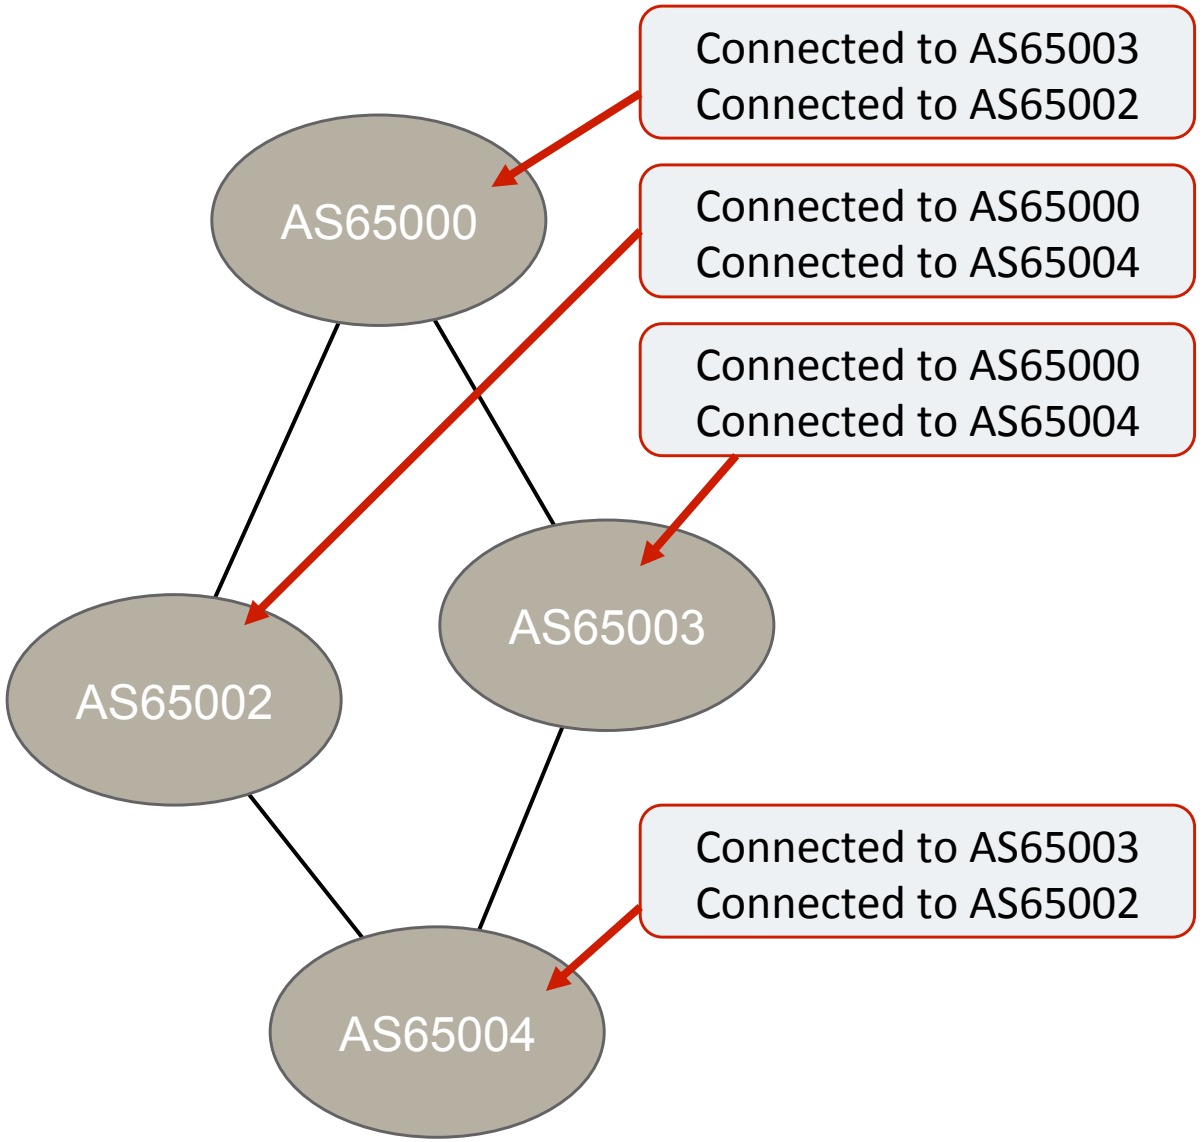

Conceptually
(simpler version)

- Build a set of path pairs
- Each path pair can contain policy
- These can be used as a set of path filters at the AS edge

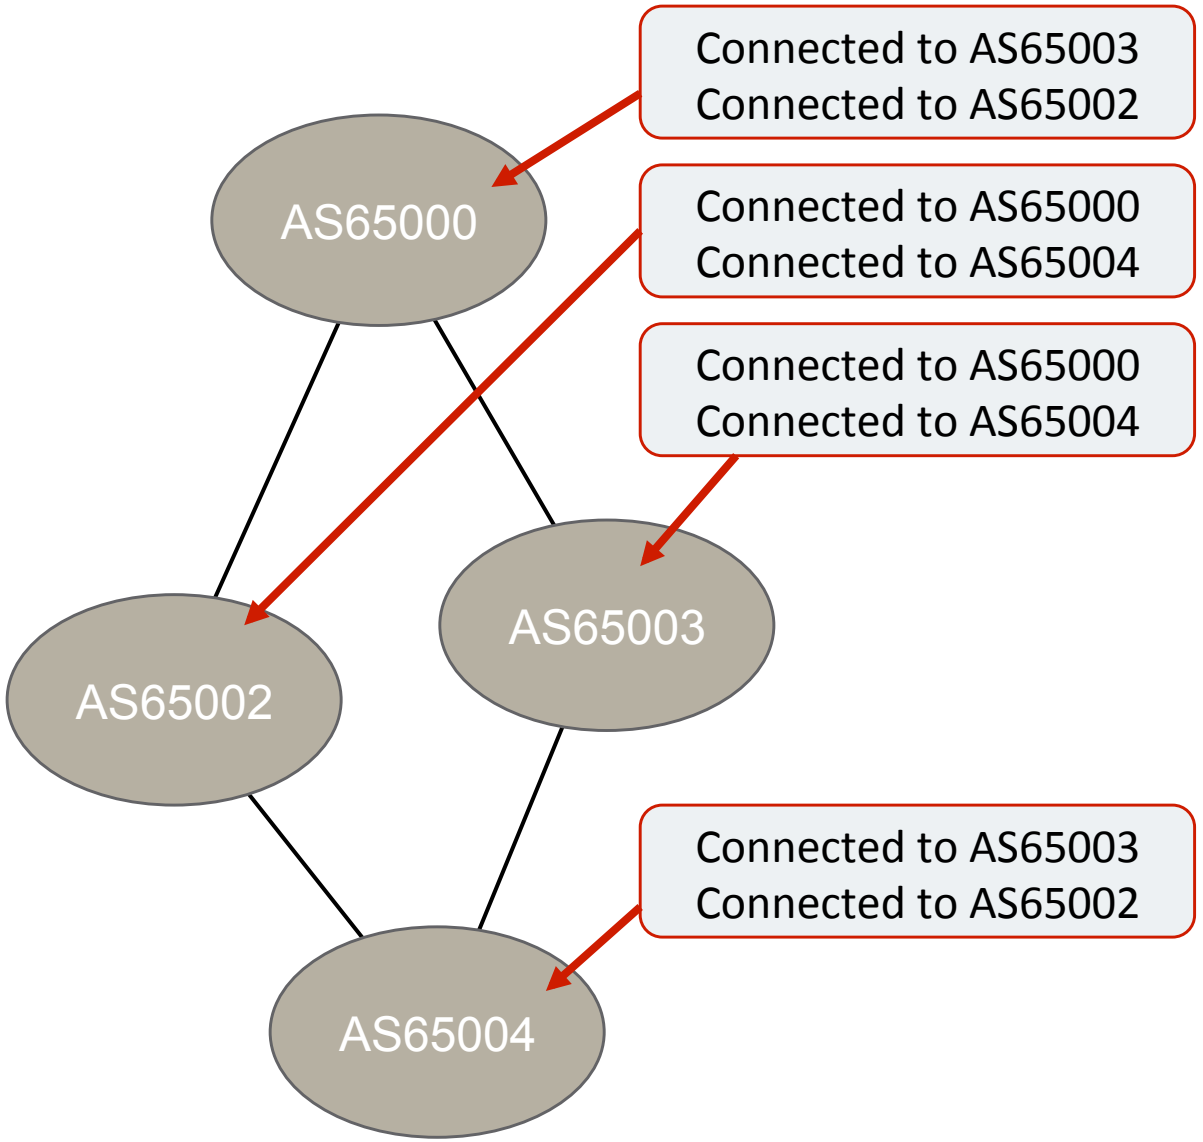

Conceptually

(simpler version)

- For instance, if an advertisement is received with the AS path [65004,65003] at AS65000
  - Is AS65004 connected to AS65003? Yes
  - Is there any policy along the path that says I shouldn't be receiving this route? No
  - Am I connected to AS65000? Yes
  - 80%+ certain this is a good route
- *Leave it to reactive/future systems to resolve the rest*

Conceptually
(more complex version)

- Tree Based DAG
  - AS' are edges
  - Connections are nodes
- Policy hangs off nodes
- *Path State Vector*

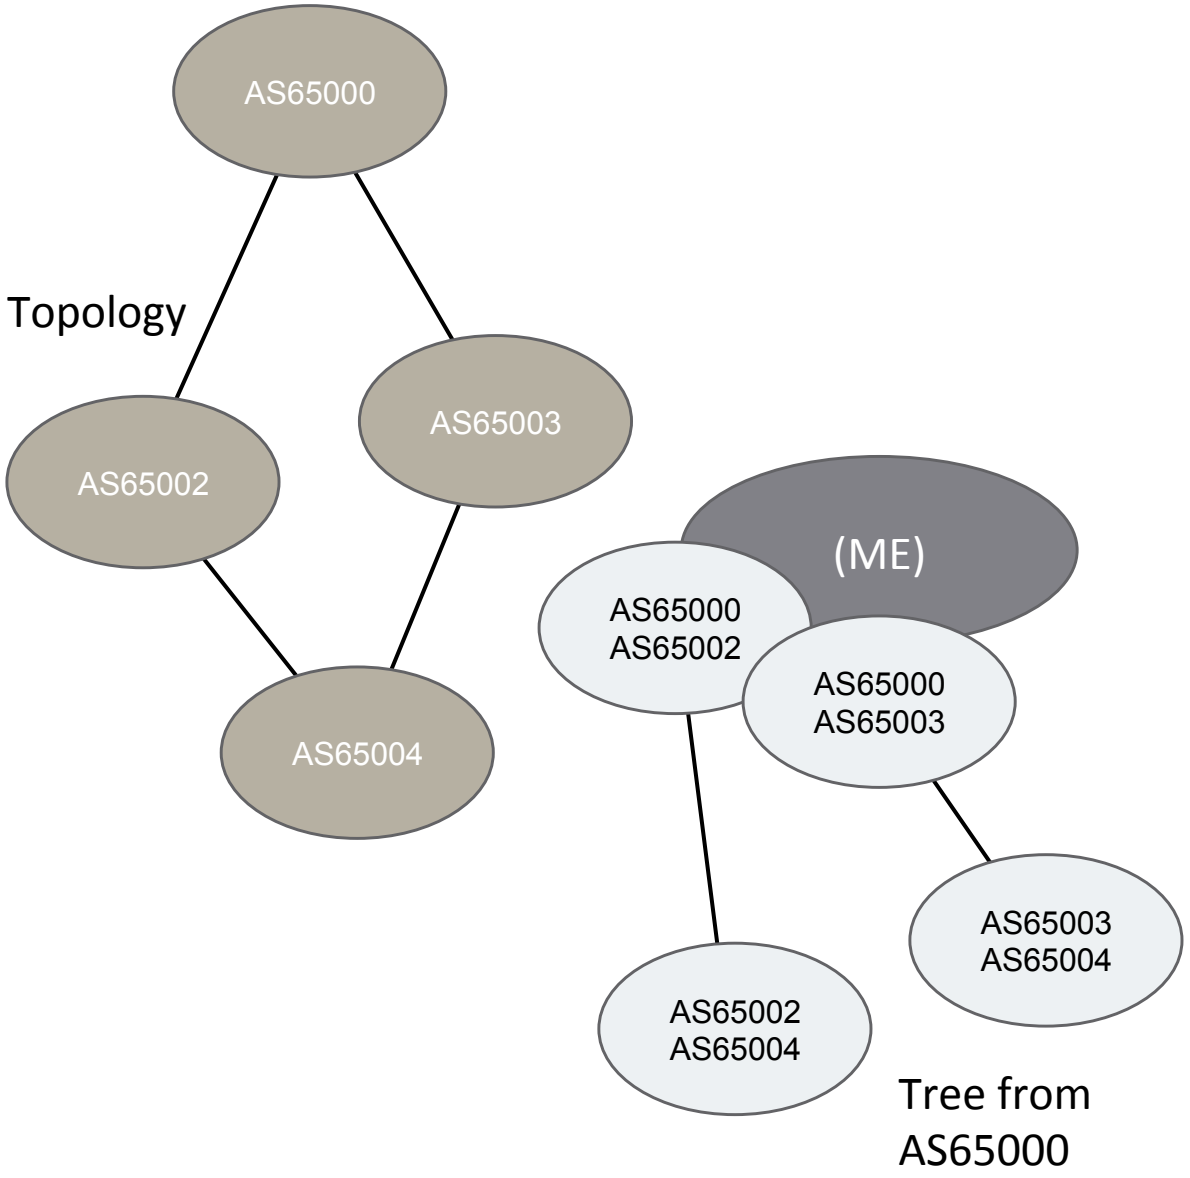

Conceptually

(more complex version)

- DAG: directed acyclic graph
  - Like an SPF, only containing all possible paths, rather than just the best path
  - Contains loops, which is okay for this application
- For any advertisement received, start with the origin and "walk" the DAG
- If I can reach myself without encountering policy problems, the route is valid
- *Leave it to reactive/future systems to resolve the rest*

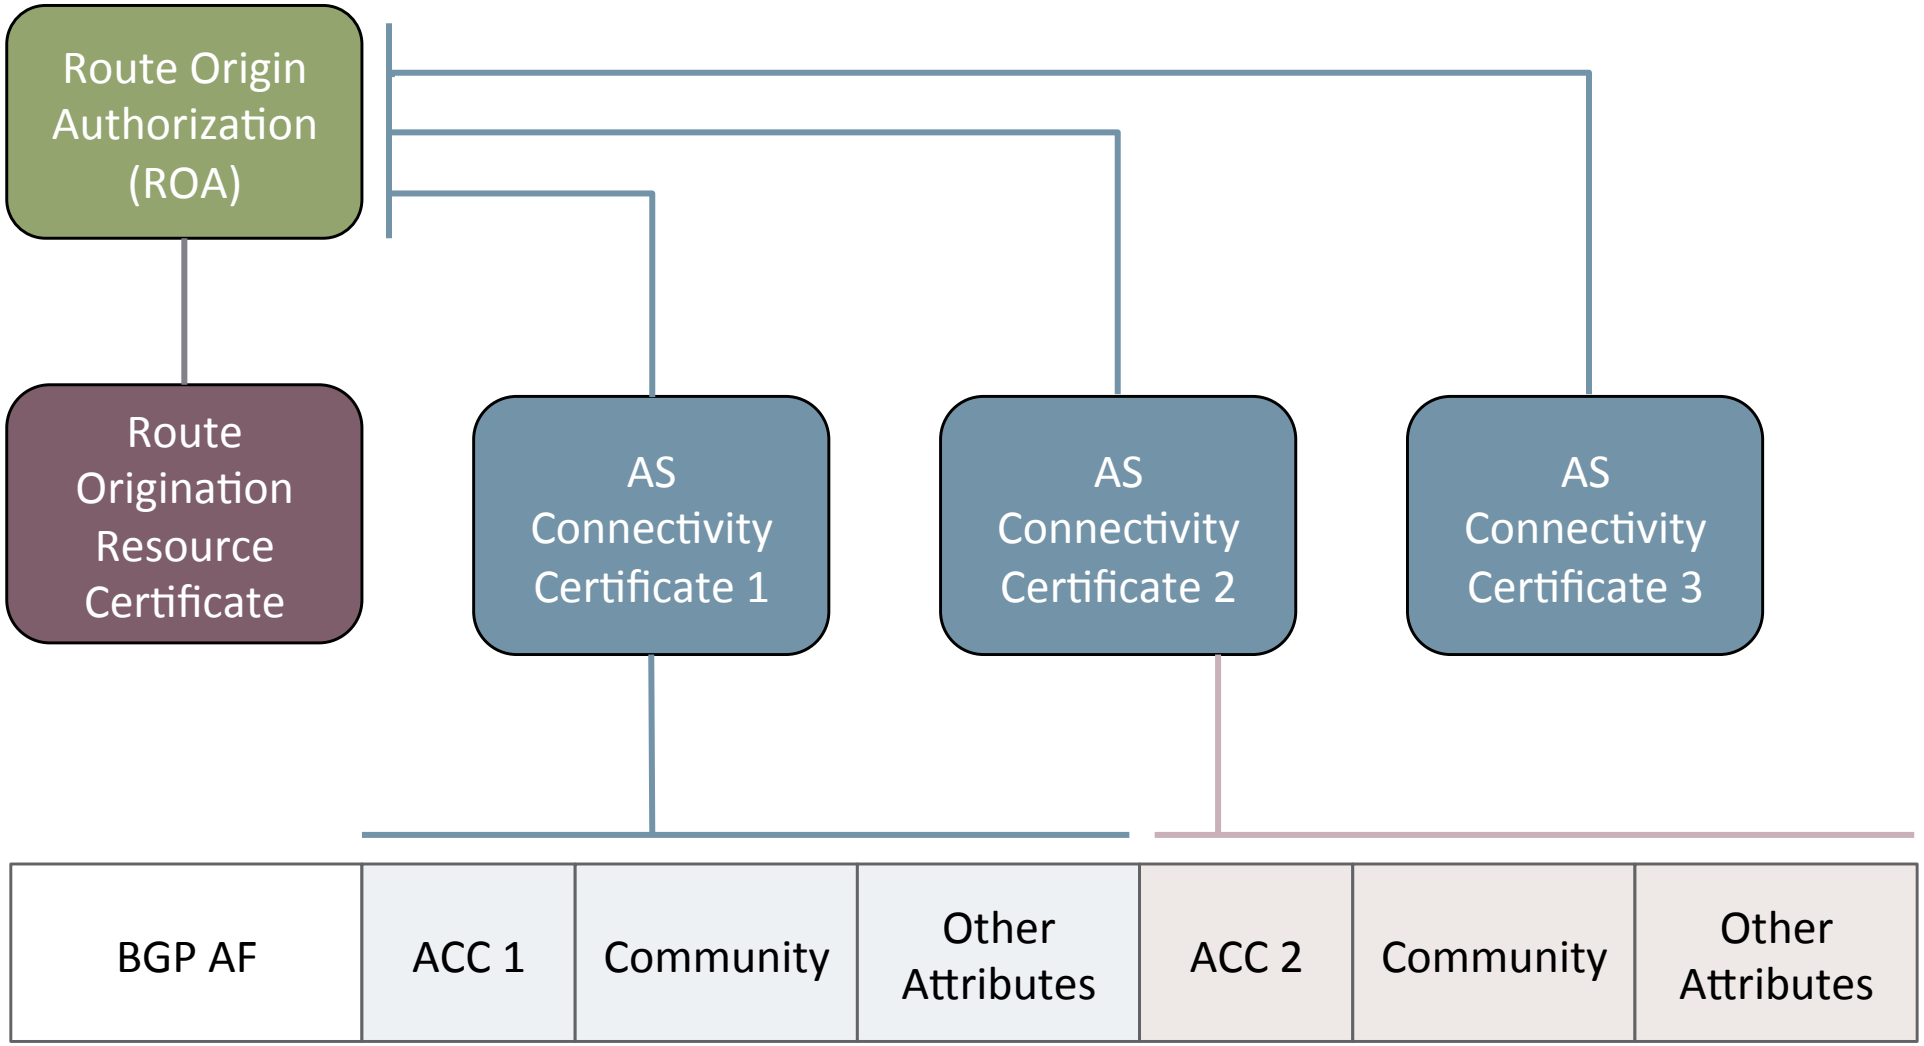
Route Origin Authorization (ROA)
Route Origination Resource Certificate
AS Connectivity Certificate 1
AS Connectivity Certificate 2
AS Connectivity Certificate 3
BGP AF
ACC 1
Community
Other Attributes
ACC 2
Community
Other Attributes

Notes

- This is one option for encoding this type of information
- Treats the certificate as essentially opaque to BGP
  - BGP is just transporting this stuff
- Communities and other attributes can be added on to supply common inter and intra AS policy
- Sequence number is included for freshness of information
- Packet formats in flux at this point

Operationally

- AS65000 advertises three *connectivity sets*
- [65000,65003]
  - Community bound
  - Only advertised when routes from AS65003 are advertised

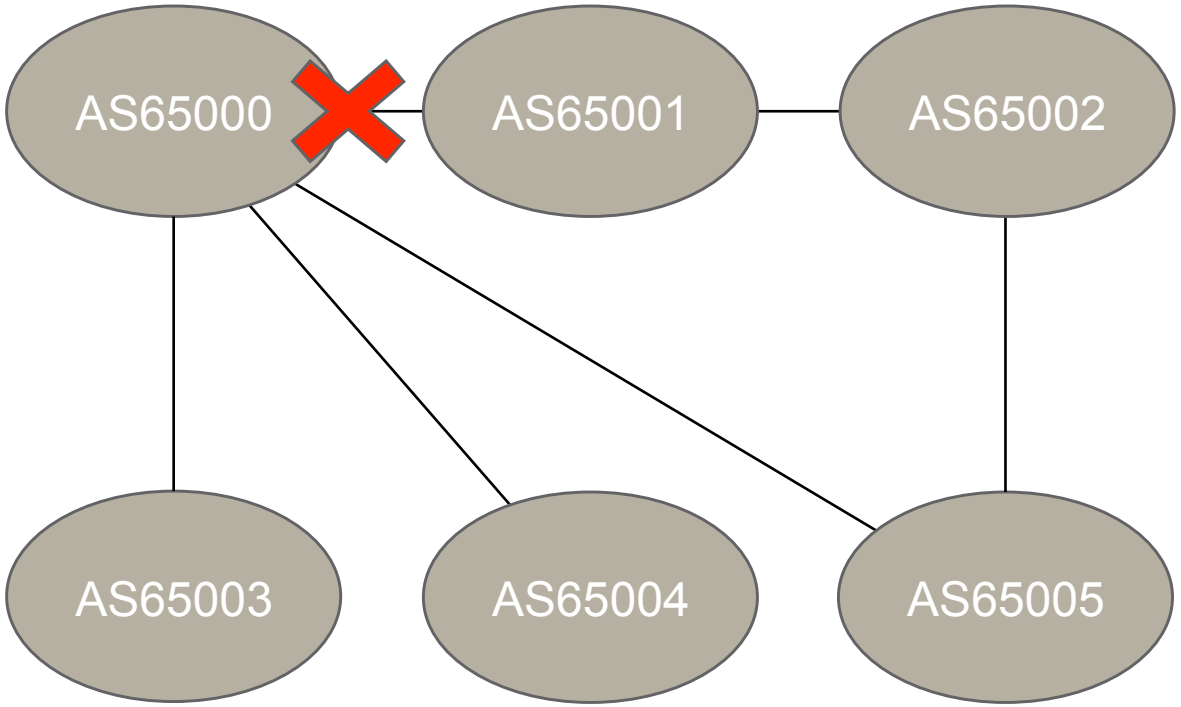

Operationally

- [65000,65004]
  - Community bound to be blocked at the AS65001=>AS65002 edge
- [65000,65005]
  - Marked as non-transit peering relationship

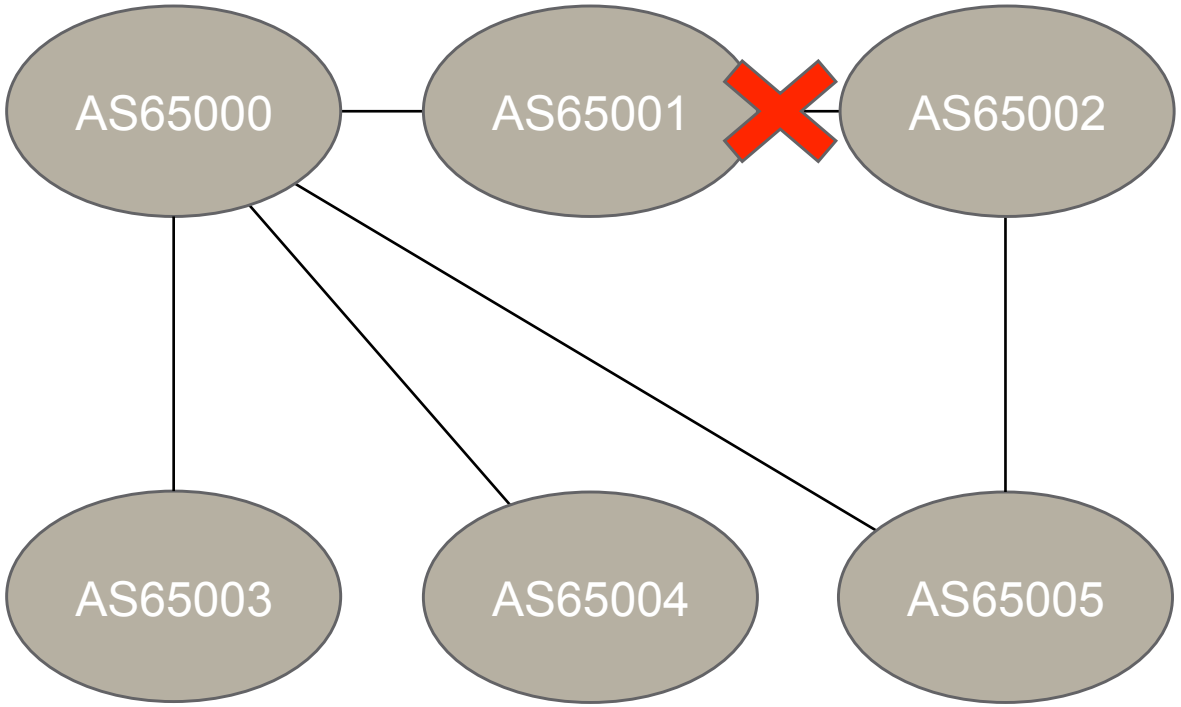

Attacks Resolved

- AS65005 advertises with a path of [65001,65005]
- AS65001 is not advertising a connection to AS65005
- AS65004 can reject the route

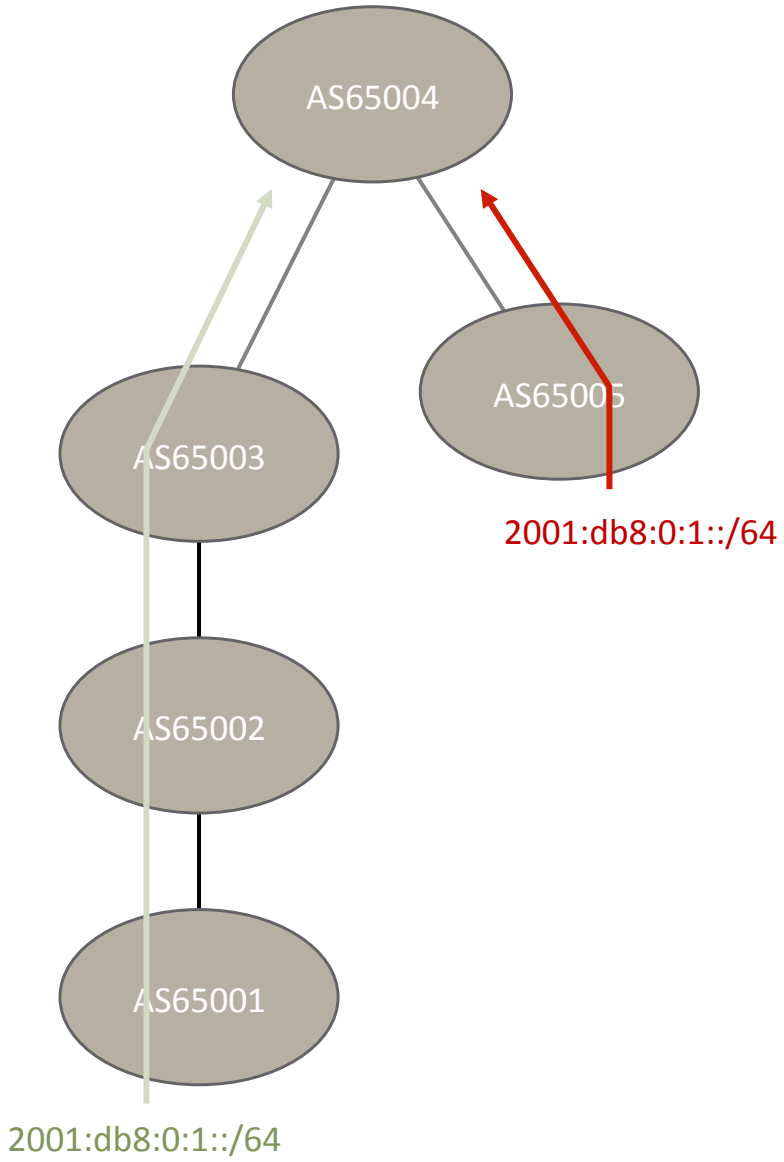

Attacks Resolved

- AS65001 is not transit
- AS65002 can mark AS65001 as not transit
- AS65005 can drop the route based on this
- *This is optional, but as more policy is exposed, more can be enforced*

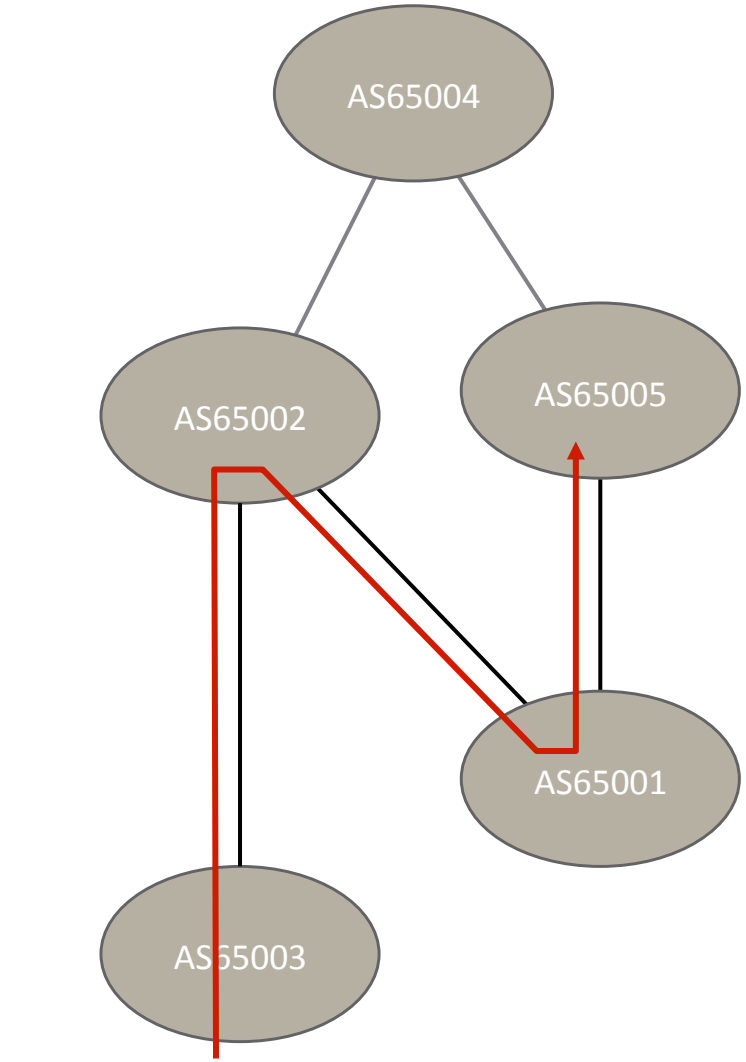

Thoughts on this solution

- Would meet the objectives of reasonably worded government mandate
- Would protect 80% or more of what needs to be protected
  - Works with existing origin validation to stop hijacking
  - Stops truly “out of path” man in the middle attacks
  - Provides a home for some policy *when desired*
  - Protects provider “private links,” etc.

Path Forward

- Small group formed to work on this
  - Increasing group size over time as folks are interested
  - Need to avoid boiling the ocean or building a camel if possible
- We need community support to build a deployable system that solves the set of problems we care about
- Eventually take this to the IETF
- If a mandate is forthcoming…
  - Hopefully we have a system in place that operators can live with

Questions?