

# Cross-Aggregate Coordination

Robert Ricci

GEC5

July 22, 2009

# Challenge

- Some links cross aggregates
  - Two or more parties have to cooperate
- Cluster C has working AM support for:
  - VLANs *within* aggregate
  - Tunnels *between* aggregates
- Next challenge: VLANs *between* aggregates

# (Some) Tunnels Are Easy

- Current support in ProtoGENI: GRE
- Stateless
- Tunnel works when both ends are set up
- Very little coordination required
  - Just communicate endpoint IP addresses

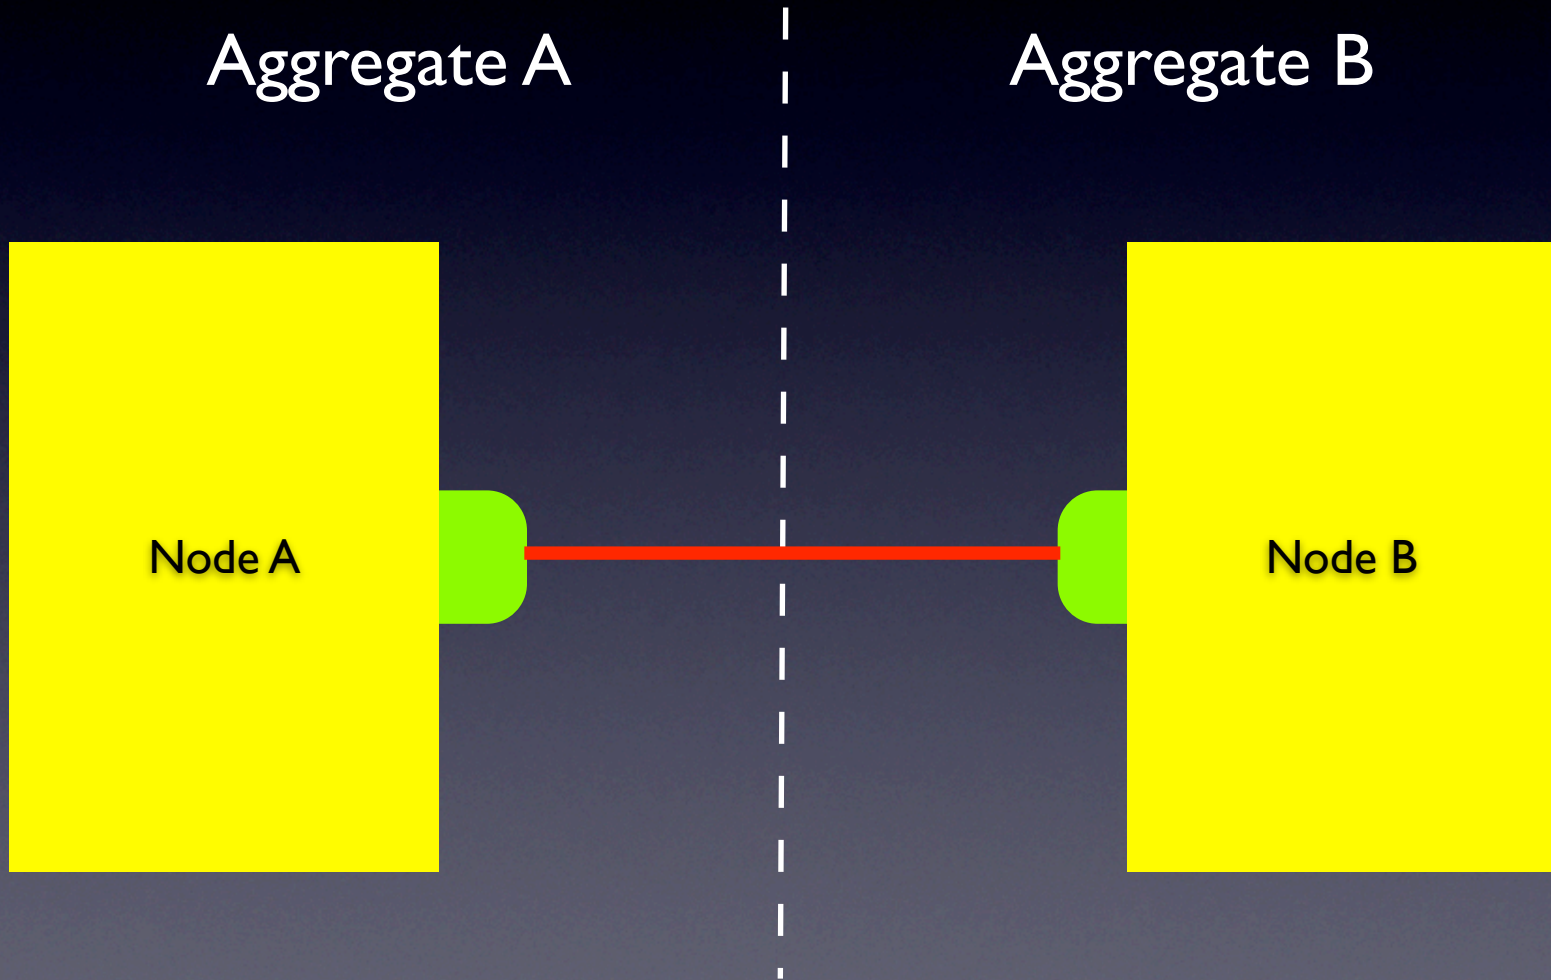
# Some Stitching is Harder

- Stateful setup (handshaking, ordering issues)
- Negotiation may be necessary
  - Encryption key establishment
  - Available identifiers
- Some are transitive problems
  - eg. VLAN #s (unless translation possible)

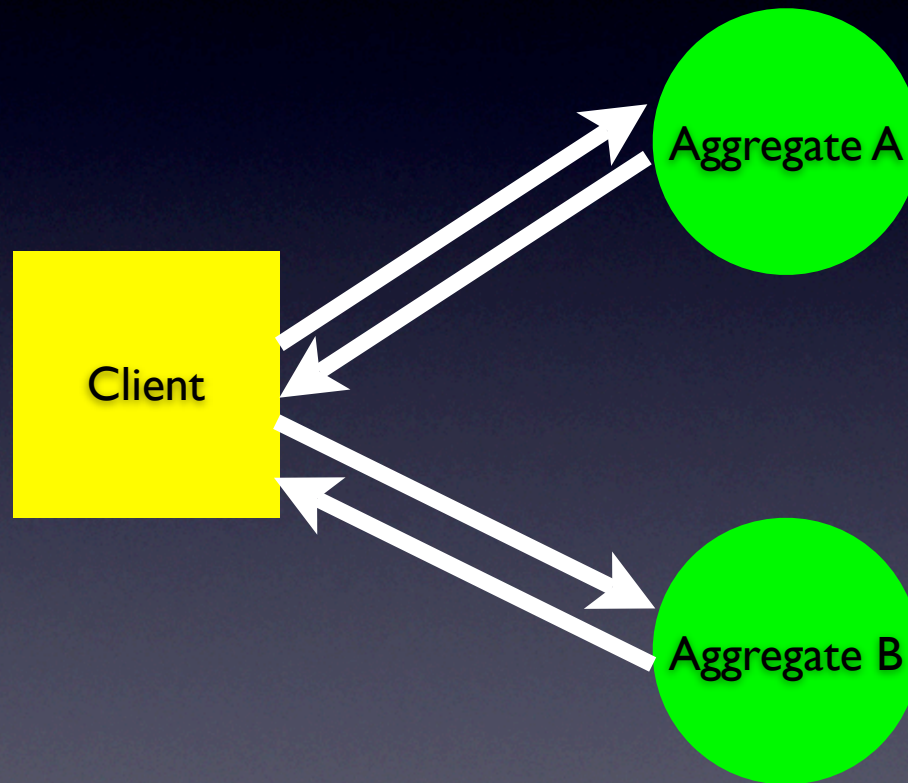
# Assumption

- Links may cross aggregates
  - Components don't
- Cross-aggregate links are established by endpoints within each aggregate

# Cross Aggregate Link



# Alternative A

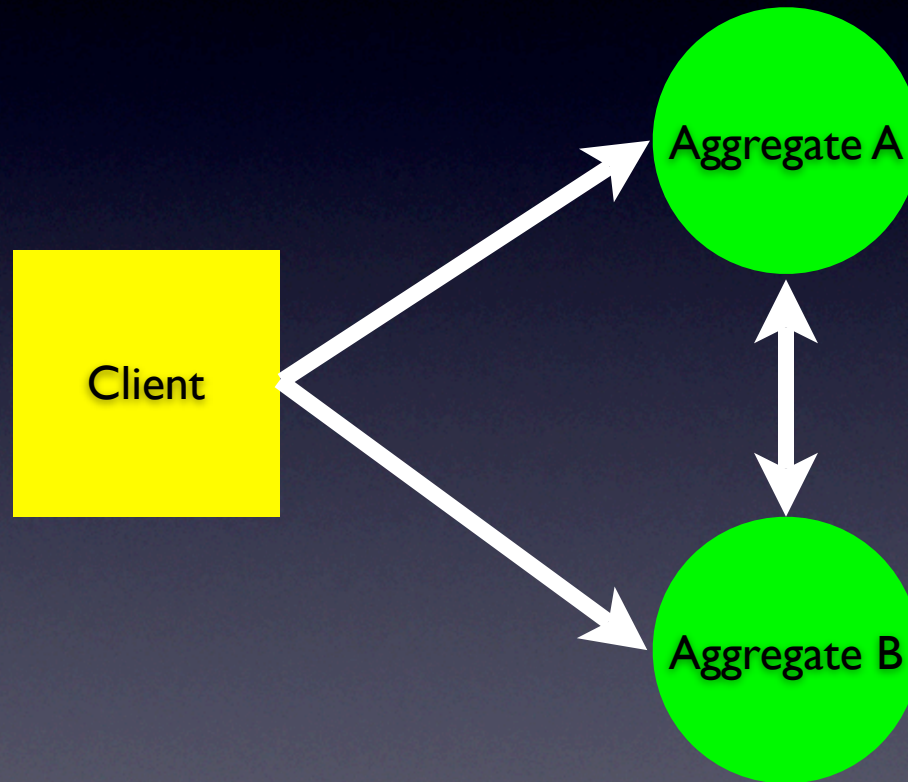


# Alternative A

- Client negotiates with each AM (through RSpec)
- Good
  - Doesn't require new AM API calls
- Bad
  - Client must know a lot about coordination
  - Complicated data path



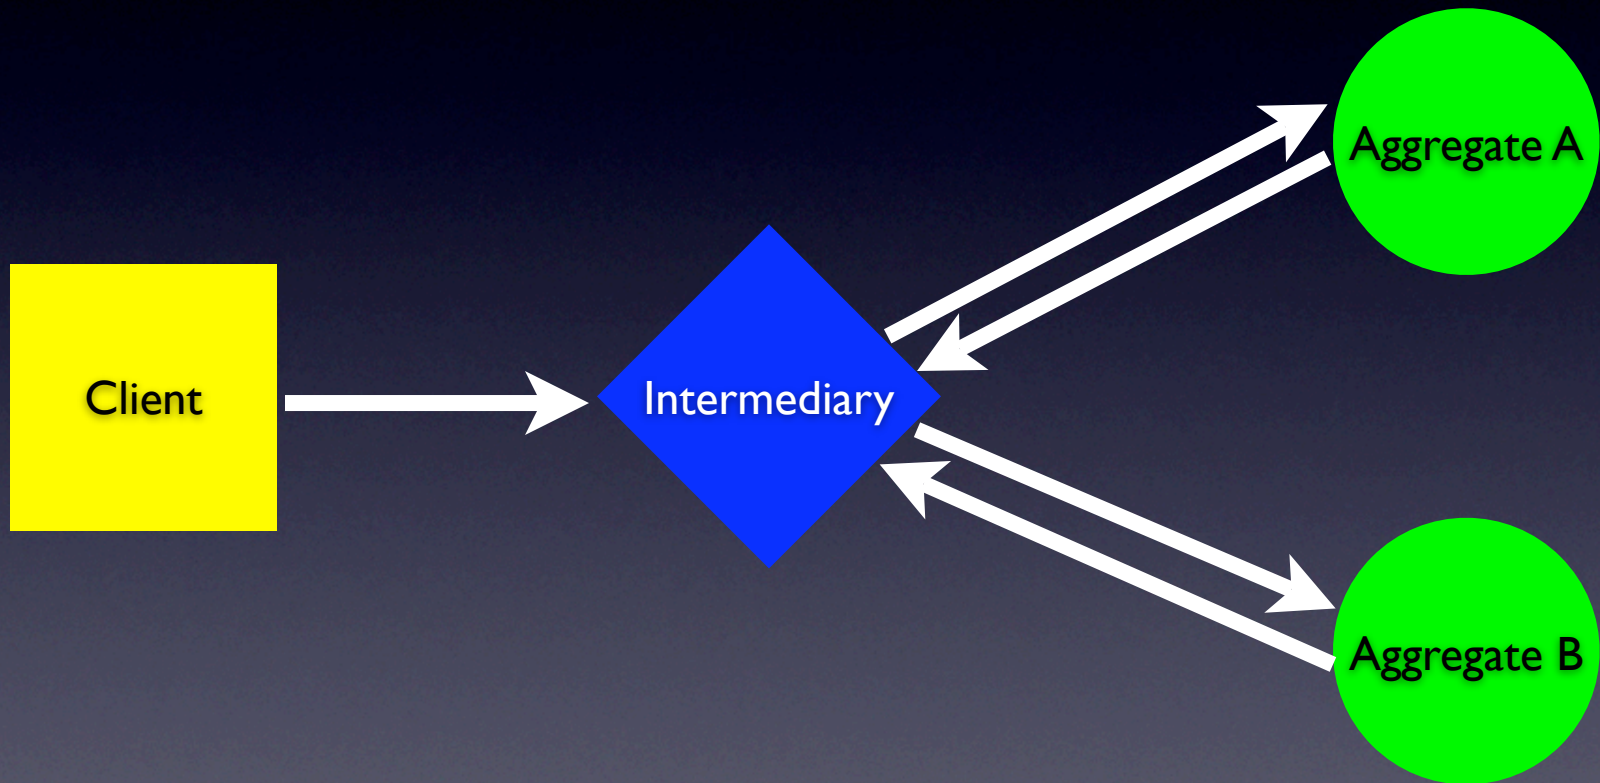
# Alternative B



# Alternative B

- AMs coordinate among themselves
  - Using a new standardized control plane API
  - RSpec *could* be medium
- Good: Client doesn't need to know details
- Bad: Hard to get global view
  - Hard to handle multi-party


# Alternative C



# Alternative C

- Untrusted intermediary negotiates for client
  - Intermediary acts on client's behalf
  - “Aggregate of Aggregates”
- Good: Easy to set up on the fly
- Bad: Just shifts problem around

# Alternative D



# Alternative D

- Same as C, but intermediary is trusted
  - Pre-established
- Good
  - Easy to have global, authoritative view
- Bad
  - Scales poorly
  - Must be set up in advance

# Alternatives

- A: Client does the work
- B: AMs talk amongst themselves
- C: Untrusted intermediary
- D: Trusted intermediary

# ProtoGENI Plan

- Hybrid of B (direct AM coordination) and D (trusted intermediary)
- New AM call, only for use by other AMs
- AMs handle two-party arrangements directly
- Trusted intermediary negotiates multi-party
  - Trusted authority picks VLAN #



# Advantages

- Client doesn't have to understand AM coordination
- Technology-specific knowledge lives in AMs
- Likely case: Backbone can be trusted party for assigning VLAN #s

# Disadvantages

- Both (or many) endpoints must be in RSpec
- Client must still contact all endpoints
  - Endpoints must wait until others have been contacted to finish setup

<http://protogeni.net>