GENI Network Stitching

Proposed Architecture

GEC10

March 16, 2010

Tom Lehman (USC/ISI)
Xi Yang (USC/ISI)

Network Stitching

- Proposed Stitching Architecture Documents
  - http://geni.maxgigapop.net/twiki/bin/view/GENI/NetworkStitching
- GENI Network Stitching Architecture-Overview
- GENI Network Stitching Architecture-Example
- GENI Network Stitching Component Design-Example
- GENI Network Stitching Schema Example

Network Stitching

- The topic here is mostly focused on the mechanics of "network stitching"
- There is an associated topic which is network stitching related user and resource authentication/ authorization/policy application.
  - The network resources associated with stitching can be viewed as just another resource to be federated via integration into SFA.
  - user credentials, slice (resource) credentials: still many issues revolving around policy – definition/application/ enforcement
- These are important issues, but this talk is focused on what is required to accomplish the actual network stitching functions

Aggregate Inter-connections

- The inter-connections between Aggregates can be "viewed as a topology"

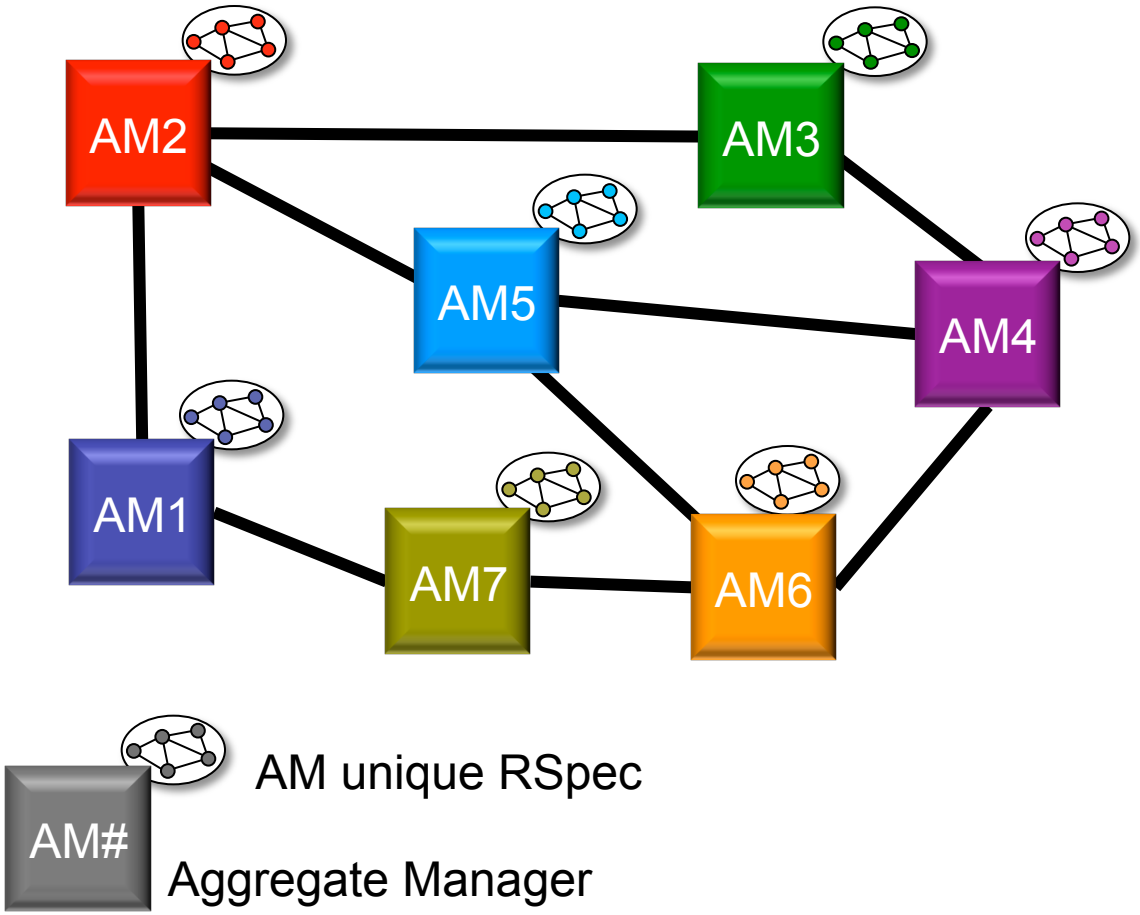

- This is a view of the AM inter-connection, not a specific slice instantiation
- represents resources potentially available for network stitching

Real World Considerations

- There are other networks in-between the AM inter-connections: sometimes dynamic, sometimes static

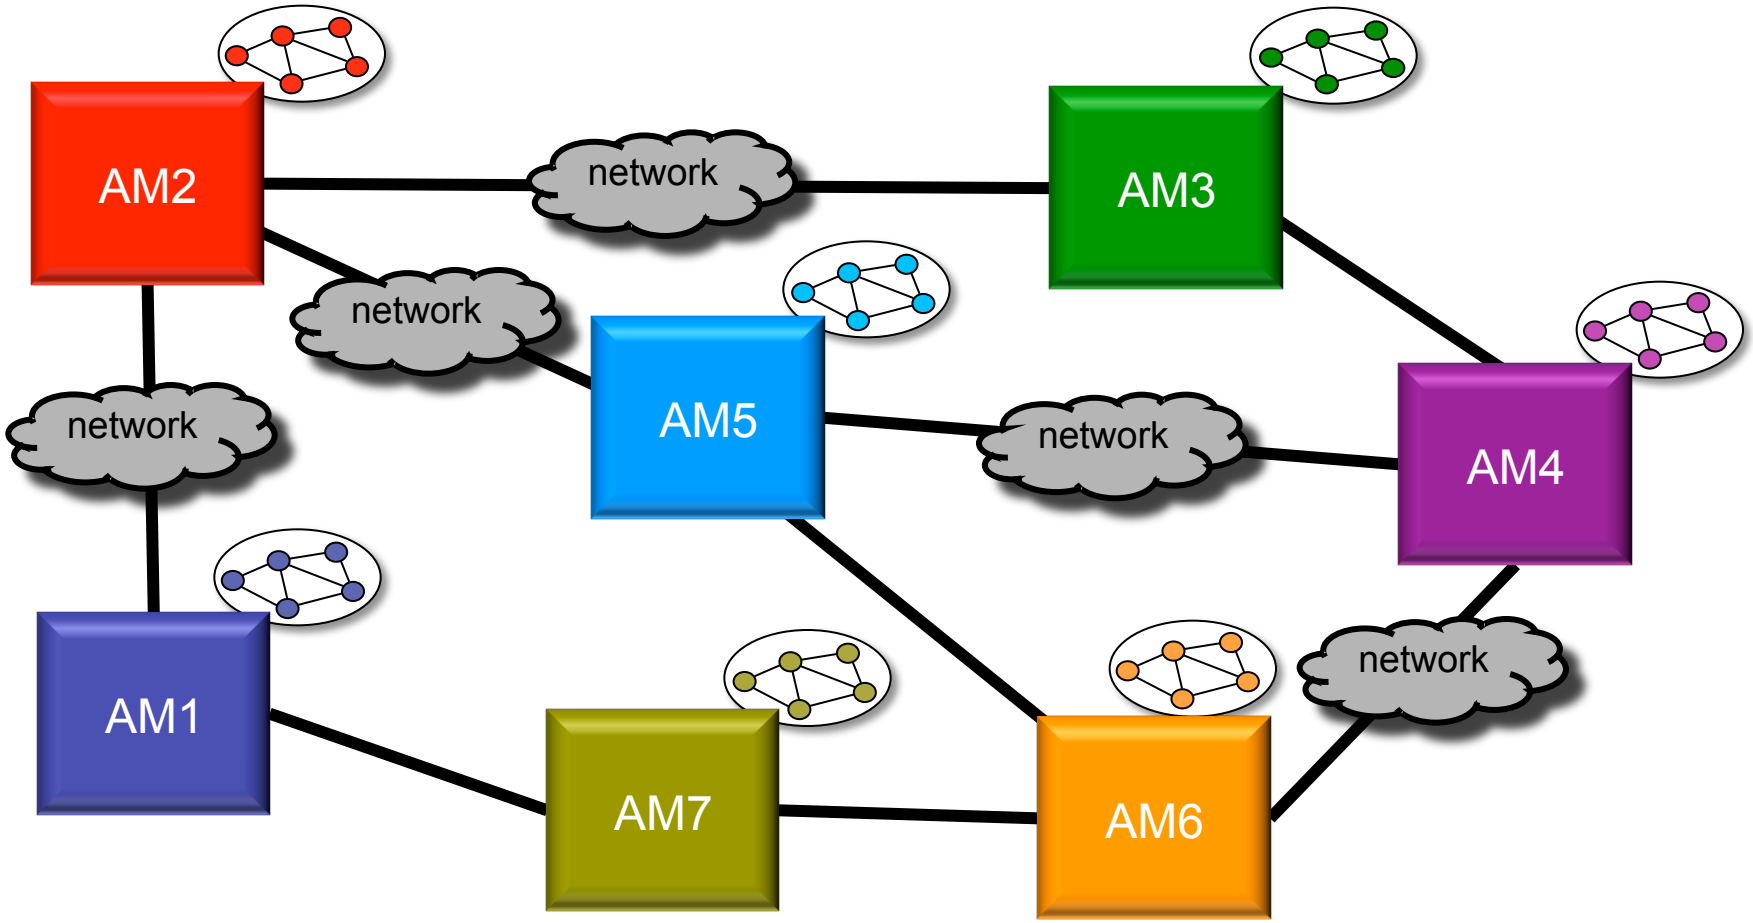

Network Stitching - Capabilities

- Create inter-aggregate network connections
- Typically Layer 2 (Ethernet VLANs)
  - but should also accommodate other layers/ technologies
- External Network navigation
  - dynamically provisioned and/or
  - constraints based on static configurations
- Able to locate and provision slice specific stitching resources (VLAN tag coordination, technology types, bw, others)

Network Stitching - Capabilities

- Be able to accomplish the above in context of the heterogeneous AM and interconnect topology
- Fit into the framework defined by the GENI Aggregate Manager API and RSpec model

Network Stitching – Functions

- Ability for a specific AM to understand where its external connection touch points are located and to represent these in its RSpec (then available to external functions)
- Ability for the "stitching function" to develop a global view of the interconnections between Aggregates based on the AM provided RSpec information, and a representation of the external network topologies between them.

Network Stitching – Functions

- Ability for a "stitching function" to use this global AM inter-connection view to initiate specific resource provisioning actions on AMs and external network resources.
- Ability to accomplish the above in an environment where there may be multiple resource provisioning actions (slice instantiations which include network stitching) simultaneously.

Architecture Components

Six Network Stitching Architecture Components are defined

- **Stitching Resource Element**
- **Common Stitching Topology Schema**
- **Stitching Topology Service**
- **Stitching Path Computation Function**
- **Stitching Workflow Function**
- **GENI AM API Stitching Extensions**

Architecture Components

- **Stitching Resource Element** – RSPEC information which defines an aggregate's external touch points to other aggregates and external networks
  - needs to be formatted in accordance with a "Common Stitching Topology Schema"
  - In an environment where all aggregates have unique RSPEC schemas, this is really a standalone "Stitching RSPEC" which defines external connections
    - if there was a common RSPEC format for external communications, then stitching information could be included in that

Architecture Components

- **Common Stitching Topology Schema** – in order to plan and conduct stitching operations across multiple aggregates we need to have a common schema format for how all aggregate describe their external touch points
  - Stitching Resource Element (i.e., Stitching RSPEC to be formatted in accordance with this)

Architecture Components

- **Stitching Topology Service**
  - collects Stitching RSPECs from all GENI AMs
  - Allows the building of larger views (global if desired) of inter-connected GENI AM space
- **Stitching Path Computation Function**
  - has the ability to obtain the global network stitching topology view from the Stitching Topology Service, perform slice instantiation specific path computations

Architecture Components

- **Stitching Workflow Function**
  - has the ability to take the path computations results and execute the workflow steps required to accomplish network stitching. This includes interacting with the specific AM APIs and external network resources
- **GENI AM API Stitching Extensions**
  - extensions to the GENI AM API to support robust and globally scalable network stitching functions.

Client Desired Slice Topology

AM A — AM C

Figure 3a
Tree model - Client Tools Initiated

The following has already happened but is not shown in this diagram:

- Client Application has retrieved RSpecs from all Aggregates of interest (ListResources for AM A and AM C)
- Client Application has formulated RSpecs to be included in CreateSliver Requests

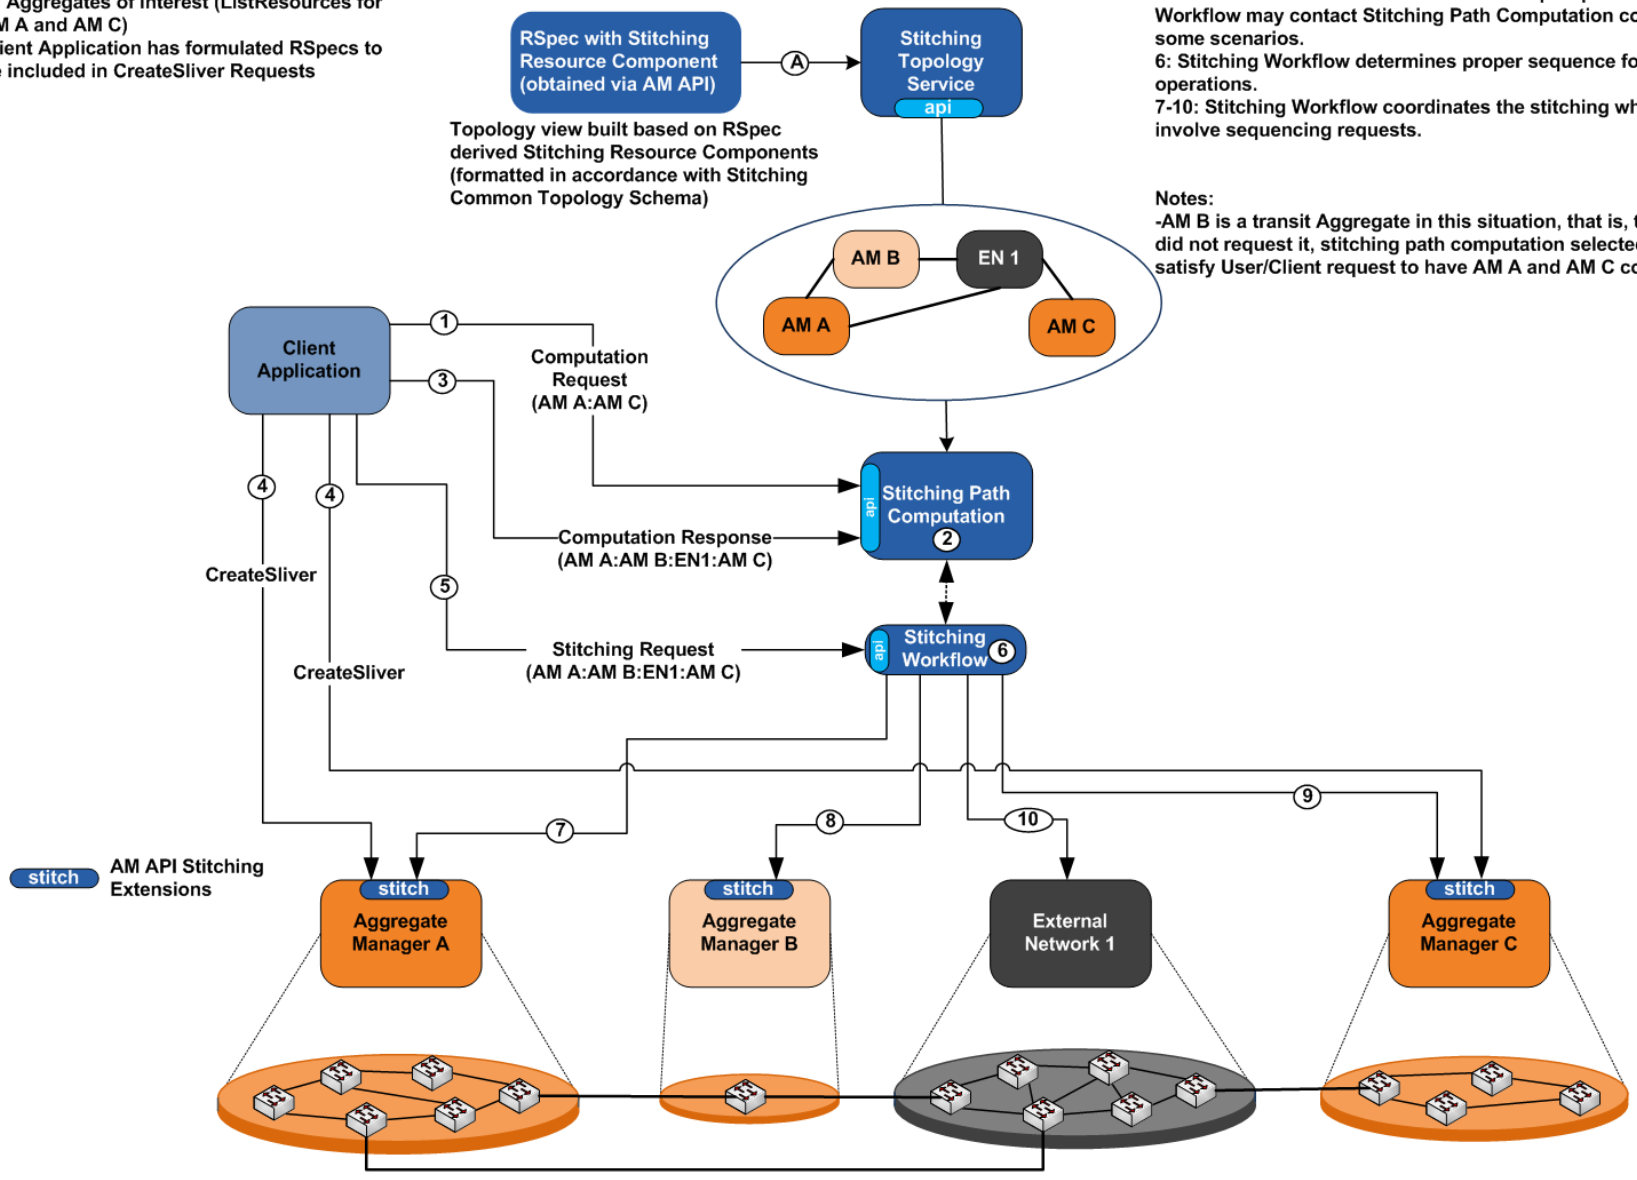

Steps:
1: Client Application asks for a stitching path computation for the desired slice of AM A: AM C
2: Path Computation Function retrieves topology and computes path
3: Client Application gets back response of AM A:AM B:EN 1: AM C, which includes the needed stitching aggregates and external networks.
4: Client uses CreateSliver to make request from AM A and AM C, which are the AMs of interest for this slice.
5: Client Application requests stitching operations from Stitching Workflow. This could be loose or strict hop request. Stitching Workflow may contact Stitching Path Computation component in some scenarios.
6: Stitching Workflow determines proper sequence for stitching operations.
7-10: Stitching Workflow coordinates the stitching which may involve sequencing requests.

Notes:
-AM B is a transit Aggregate in this situation, that is, the Client/User did not request it, stitching path computation selected as a way to satisfy User/Client request to have AM A and AM C connected.

v5

Client Desired Slice Topology

AM A — AM C

Figure 3b
Tree model - Aggregate Manager Initiated

The following has already happened but is not shown in this diagram:

- Client Application has retrieved RSpecs from all Aggregates of interest (ListResources for AM A and AM C)
- Client Application has formulated RSpecs to be included in CreateSliver Requests

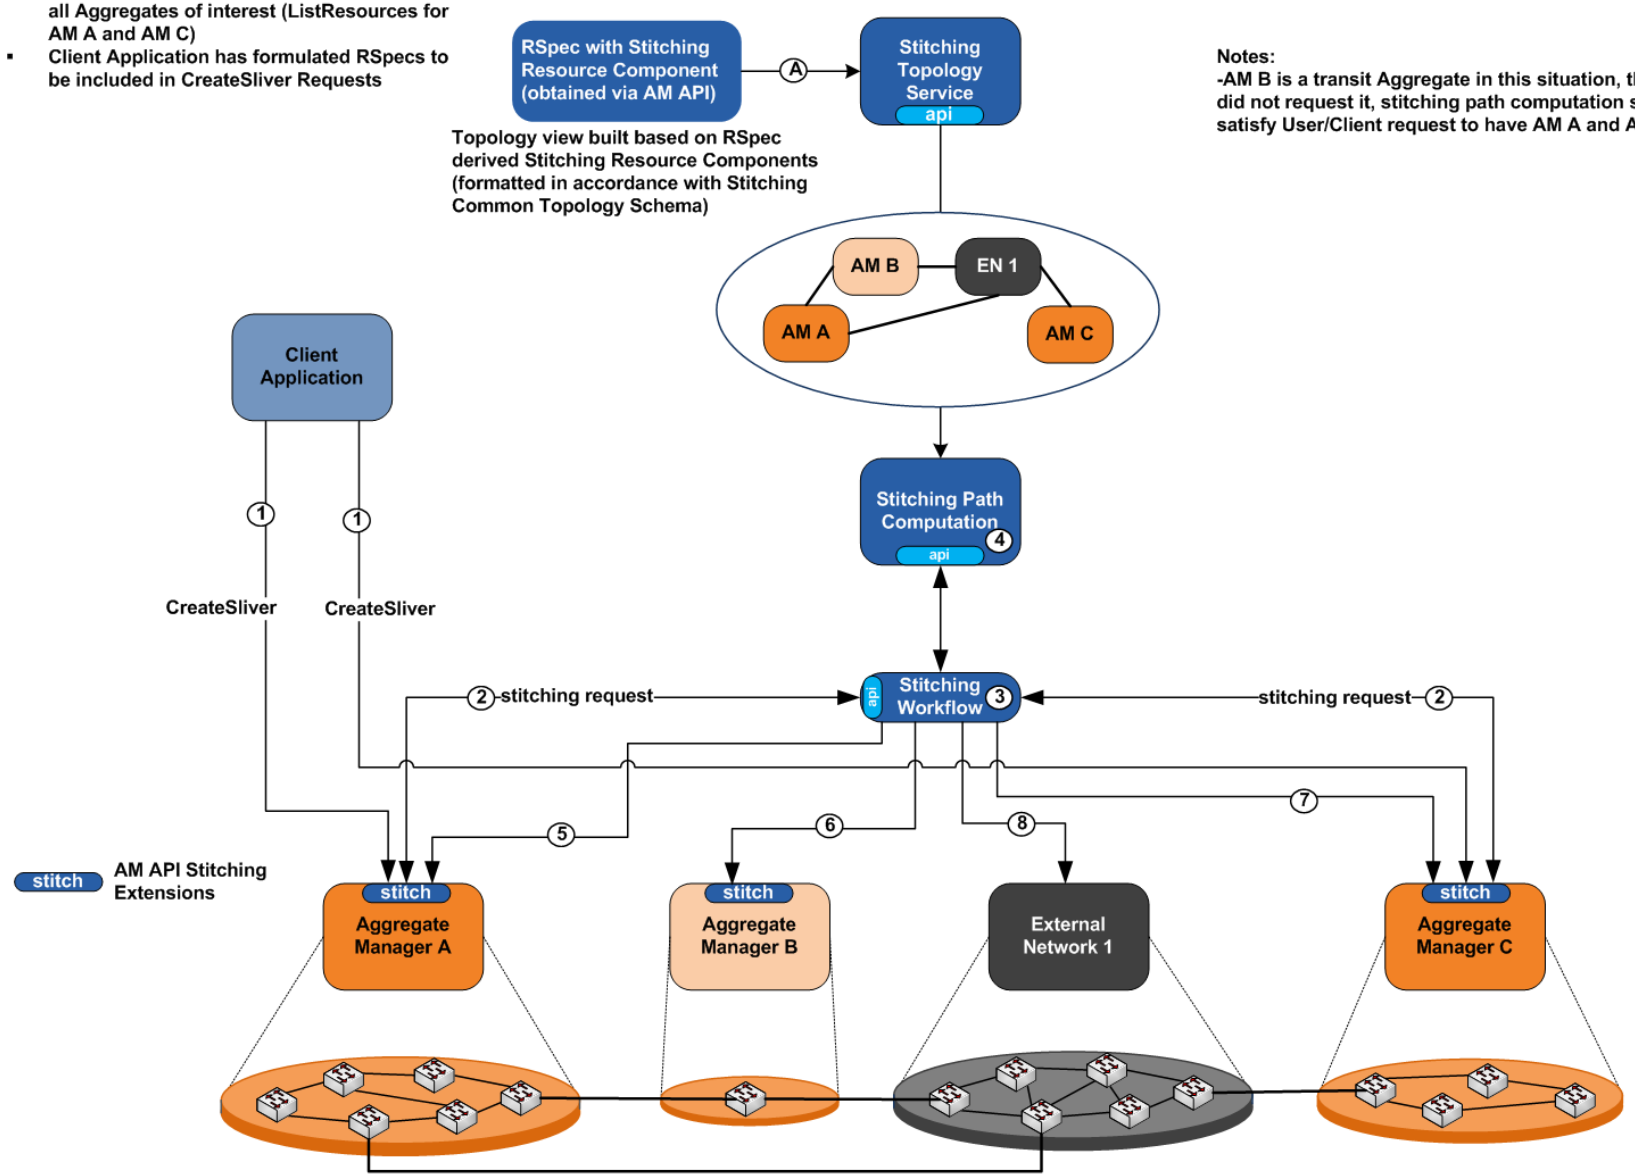

Steps:
1: Client uses CreateSliver to make request from AM A and AM C, which are the AMs of interest for this slice.
2: AM A and AM C realize they are not directly connected to each other, so ask for help from Stitching Workflow process
3-4: Stitching Workflow contacts Stitching Path Computation and gets a result that says use AM A:AM B:EN 1:AM C
5-8: Stitching Workflow coordinates the stitching which may involve sequencing requests.

Notes:
-AM B is a transit Aggregate in this situation, that is, the Client/User did not request it, stitching path computation selected it as a way to satisfy User/Client request to have AM A and AM C connected.

v5

Client Desired Slice Topology

AM A — AM C

Figure 3c
Chain model

The following has already happened but is not shown in this diagram:

- Client Application has retrieved RSpecs from all Aggregates of interest (ListResources for AM A and AM C)
- Client Application has formulated RSpecs to be included in CreateSliver Requests

Steps:
1: Client presents RSpec for Slice which includes AM A and AM C
2: Stitching path computation finds a path that says AM A:AM B:EN 1:AM C
3: Client gets a response: AM A:AM B:EN 1:AM C
4: Client goes thru CreateSliver exchange with AM A, AM B, AM C
5-7: The CreateSliver request to AM A includes the path computation results. As a result, the AM A knows that it is first in chain (and knows that path is AM A, AM B, EN1 AM C) and starts the chain stitching process. AM A sends a range of vlans (with suggested vlan) to AM B, AM B adjusts range and sends to EN 1, EN1 adjusts range and sends to AM C, AM C randomly picks a vlan from range and sends info back down chain toward AM A. This may require a Proxy AM to cover the EN1.

Notes:
-AM B is a transit Aggregate in this situation, that is, the Client/User did not request it, stitching path computation selected it as a way to satisfy User/Client request to have AM A and AM C connected.
-The Stitching Workflow component is shown here, but not used in this scenario, because we used the chain method

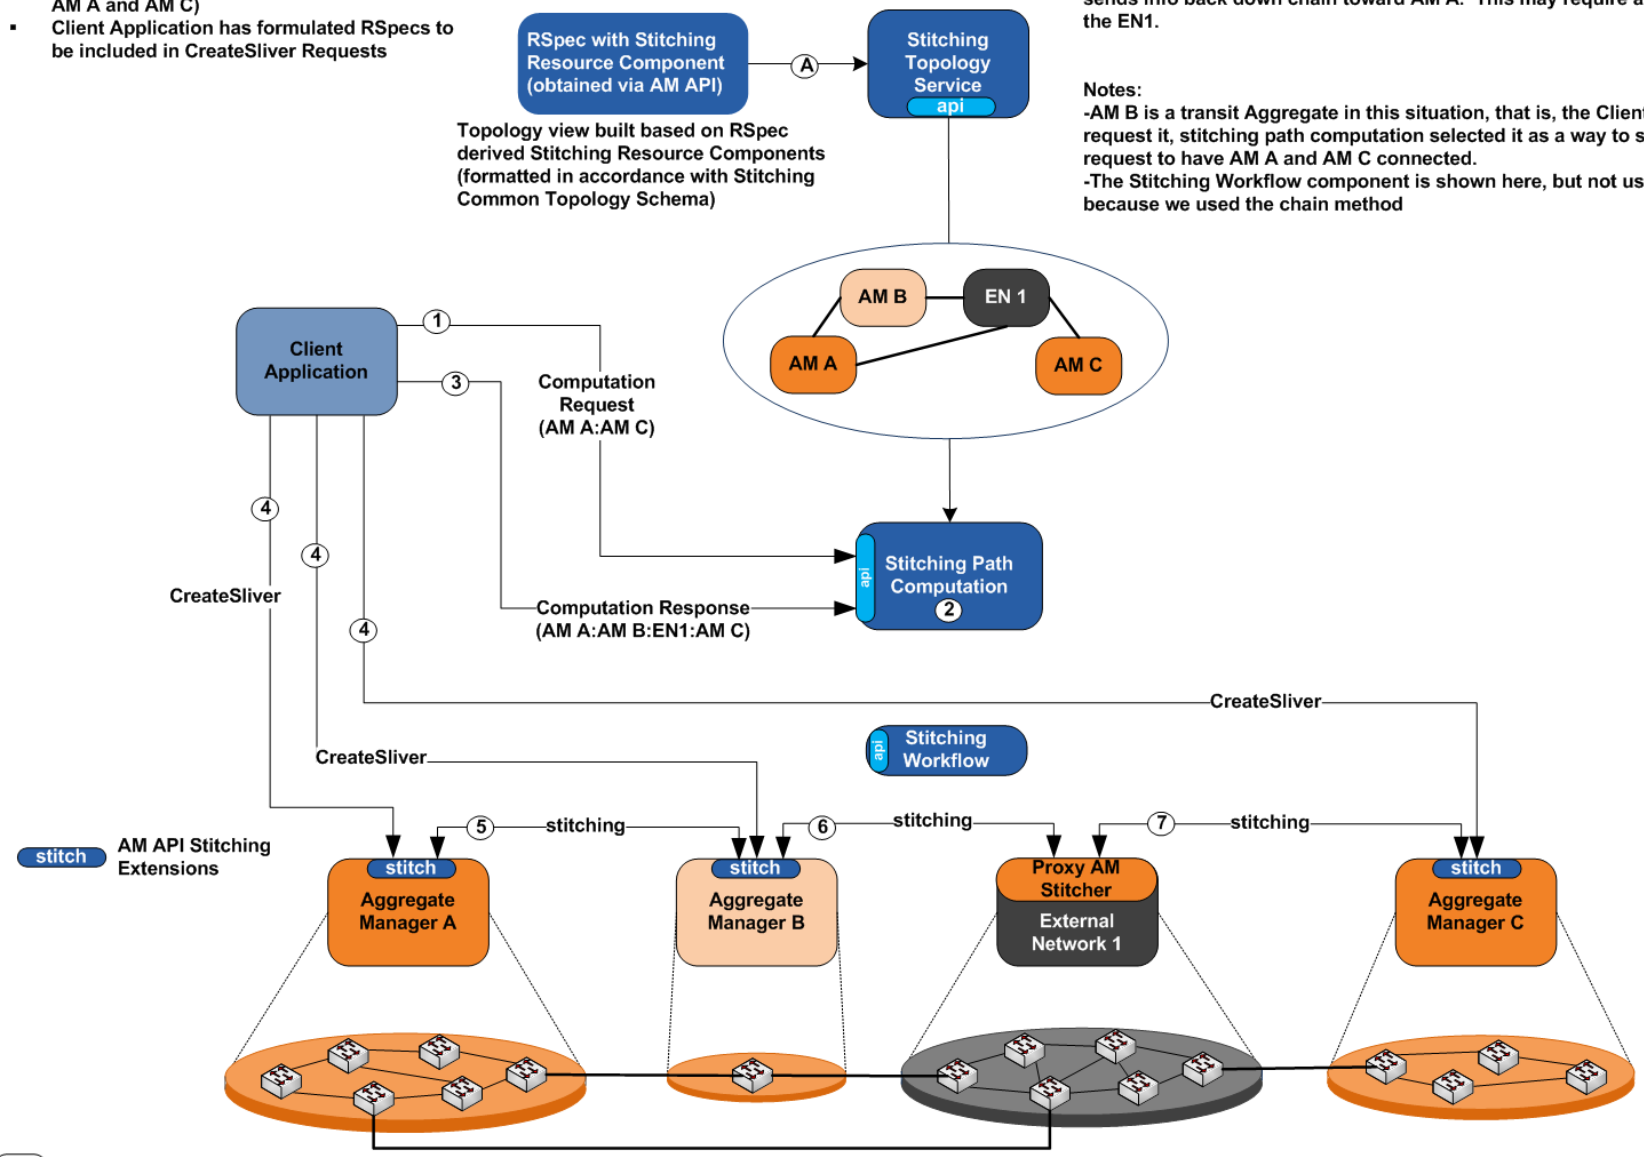

v5

Client Desired Slice Topology

AM A — AM C

Figure 3d
Hybrid Stitching Example
Tree and Chain

The following has already happened but not shown in this diagram:

- Client Application has retrieved RSpecs from all Aggregates of interest (ListResources for AM A and AM C)
- Client Application has formulated RSpecs to be included in GetTicket/CreateSliver Requests

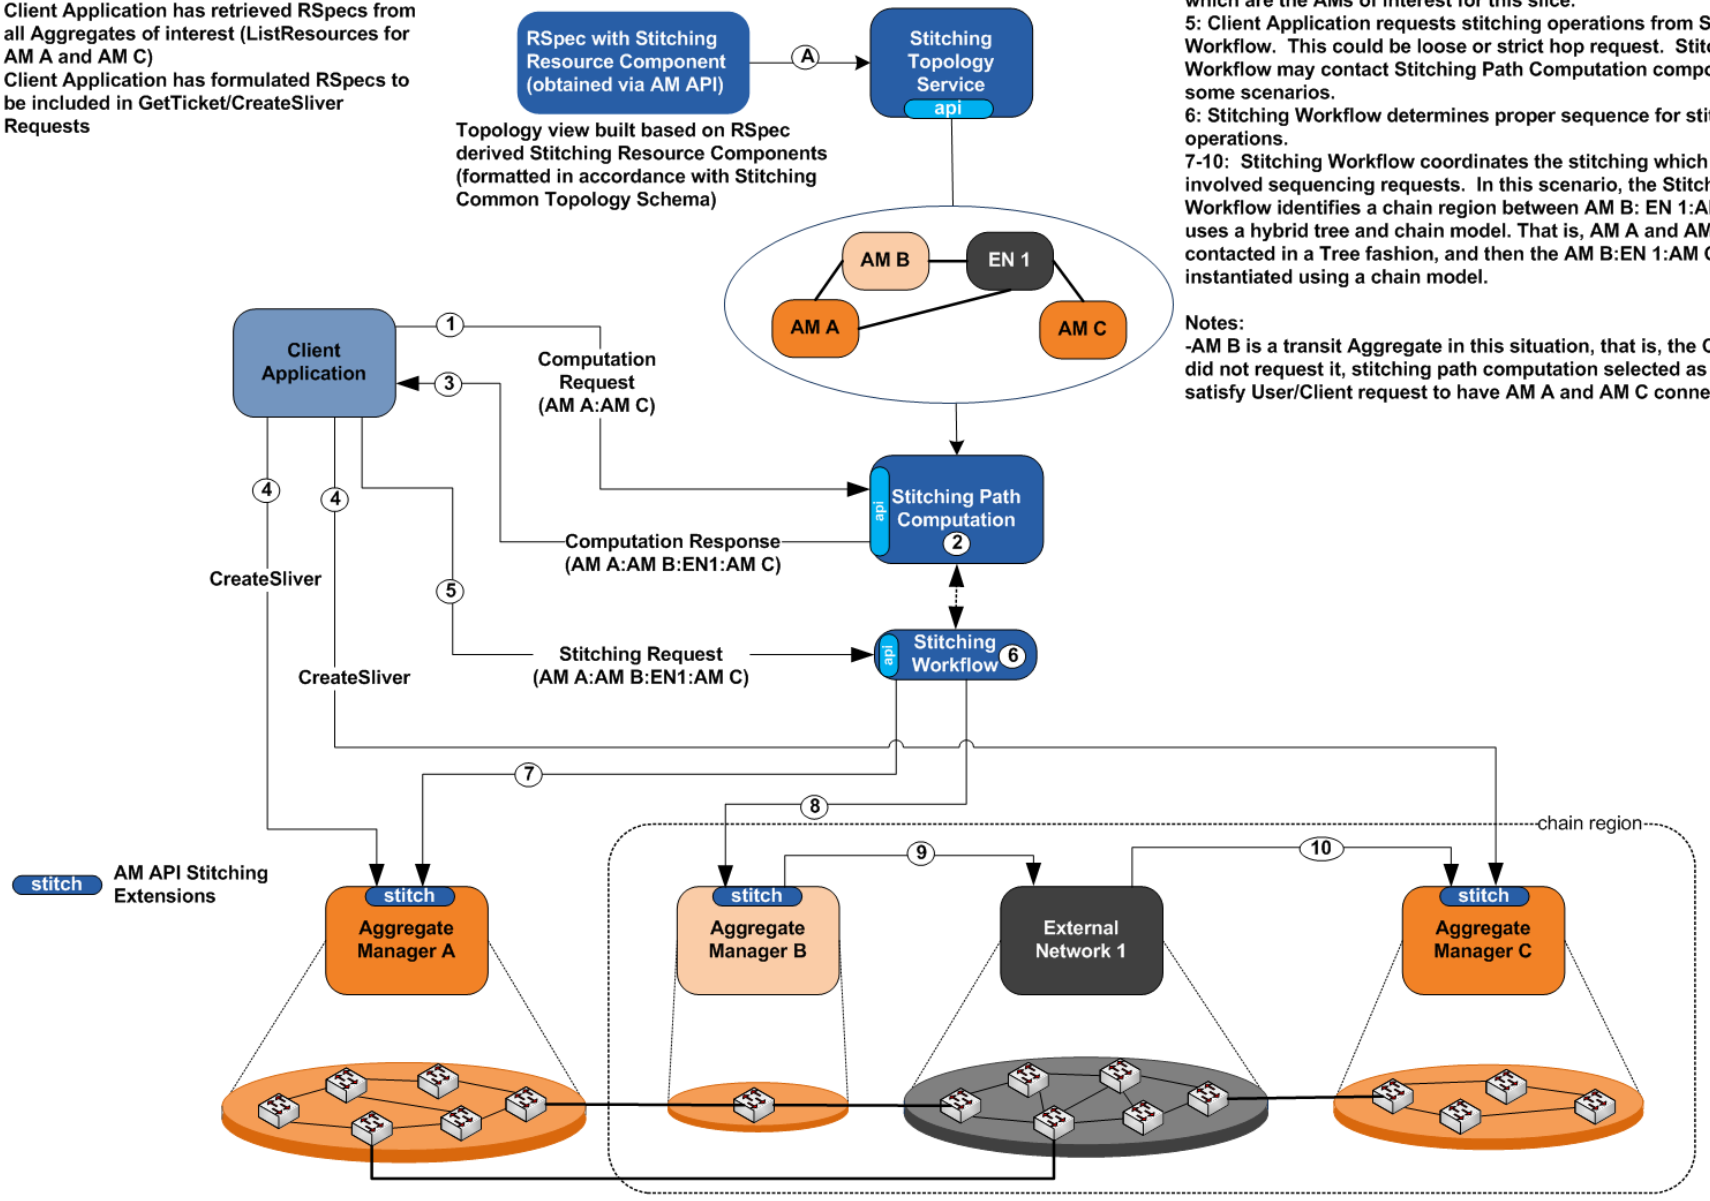

Steps:
1: Client Application asks for a stitching path computation for the desired slice of AM A: AM C
2: Path Computation Function retrieves topology and computes path
3: Client Application gets back response of AM A:AM B:EN 1: AM C, which includes the needed stitching aggregates and external networks.
4: Client use CreateSliver to make request from AM A and AM C, which are the AMs of interest for this slice.
5: Client Application requests stitching operations from Stitching Workflow. This could be loose or strict hop request. Stitching Workflow may contact Stitching Path Computation component in some scenarios.
6: Stitching Workflow determines proper sequence for stitching operations.
7-10: Stitching Workflow coordinates the stitching which may involved sequencing requests. In this scenario, the Stitching Workflow identifies a chain region between AM B: EN 1:AM C and uses a hybrid tree and chain model. That is, AM A and AM B are contacted in a Tree fashion, and then the AM B:EN 1:AM C path is instantiated using a chain model.

Notes:
-AM B is a transit Aggregate in this situation, that is, the Client/User did not request it, stitching path computation selected as a way to satisfy User/Client request to have AM A and AM C connected.

Stitching RSPEC and Topology Service

- The Topology Service is populated with information from individual AM Stitching RSPECs
- This information is not intended to be dynamic. An external link may be identified as a 10Gbps link, with a VLAN range of 1000-2000 available for dynamic aggregate stitching.

Stitching RSPEC and Topology Service

- However, the fact that 3 Gbps and VLAN 1051 may be already reserved at any one instant in time is generally not maintained in the Topology Service.
- The Topology Service is generally utilized to seed the Stitching Path Computation and Workflow Functions. It is through these real-time processes that the slice specific multi-aggregate resource identification and provisioning actions occur.

Component Design Example

- GENI Network Stitching Component Design-Example
  - Stitching Resource Element Definition
  - Stitching Topology Service API
  - Stitching Path Computation API
  - Stitching Workflow API
  - AM API Network Stitching Extensions
  - Common Stitching Topology Schema
- provides an example of a specific design
  - intent is to show an example, not necessarily to define the final design details

Stitching RSPEC Example

- Leverages IETF GMPLS information model for describing network elements, links, and interconnects.
- Stitching web site has the following:
  - geni-stitching.xsd – schema leveraging IETF GMPLS information model
  - geni-stitching-example.xml – example of how MAX Aggregate would advertise an external connection to ProtoGENI Aggregate
  - stitching-protogeni-rspec-v2-style.xml – same example as above, but formatted using ProtoGENI RSPEC v2 schema

Example Schema
Information Model Close to UNIS

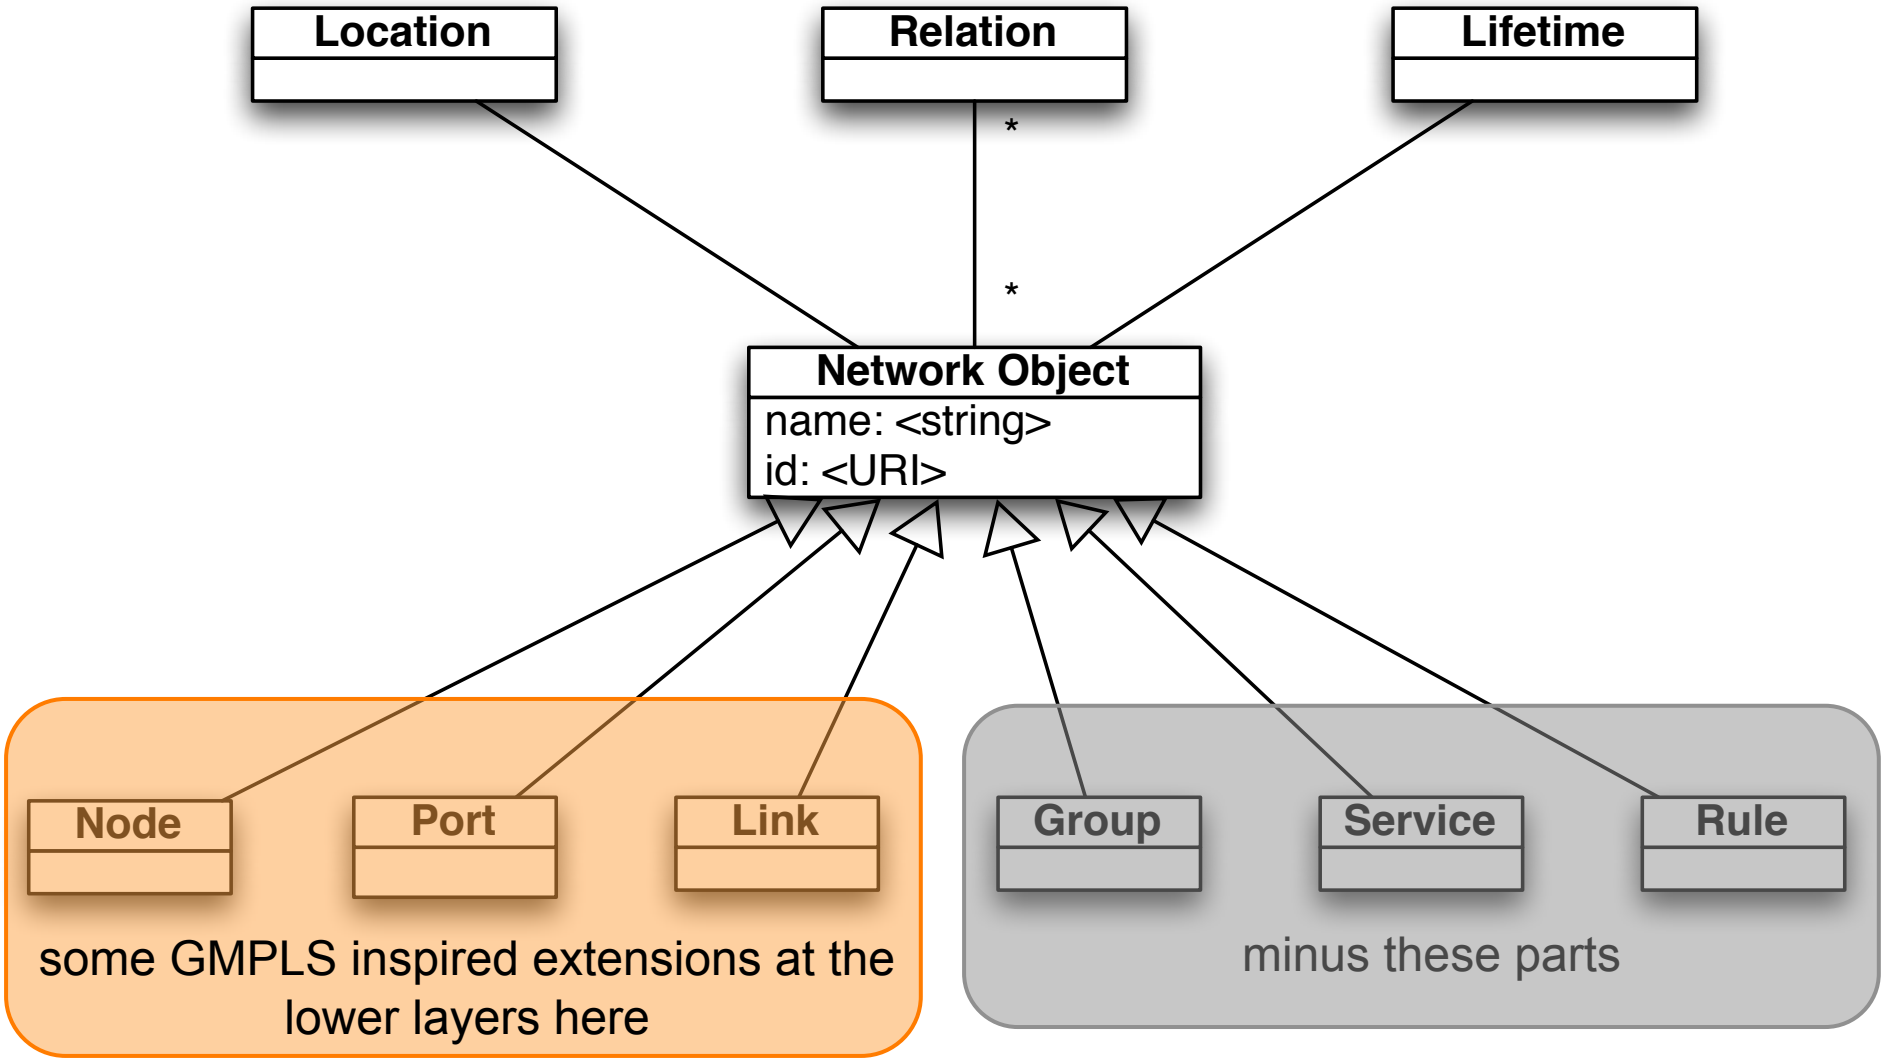

Stitching Architecture Notes

- The intent is that a single architecture and suite of components described will support multiple workflows and deployment configurations (tree, chain, hybrid)
- The key items to enable this are
  - agreement on a Common Stitching Topology Schema
  - all aggregates to make available their external connection information in this format (Stitching Resource Element)

Stitching Implementation Notes

- There are multiple implementation options for each of the architecture components
- One option is that the following could be made available as general GENI Infrastructure Services:
  - Stitching Topology Service
  - Stitching Path Computation
  - Stitching Workflow Function
- Projects could still develop their own versions of the above if expanded or project unique capabilities desired

What is Needed Next

- Stitching RSPEC agreement:
  - information model and content
- Common Topology Schema Format
- API definitions for the other components
  - Stitching Topology Service API
  - Stitching Path Computation API
  - Stitching Workflow API
  - AM API Network Stitching Extensions
- Implementation plan

Summary

- A GENI Stitching Architecture has been proposed (six architecture components)
- An example design for each of the components has been developed to provide something specific to comment against
- Community discussions have begun and will continue
- The objective is to develop a consensus and move to implementation and testing
- Please participate and help us accomplish these goals

The End

Thank-you!

extras

Stitching Example

- Desired GENI Slice
  - slivers on ProtoGENI, MAX, BBN, ORCA Aggregates
  - Three "Stitching Segments" required

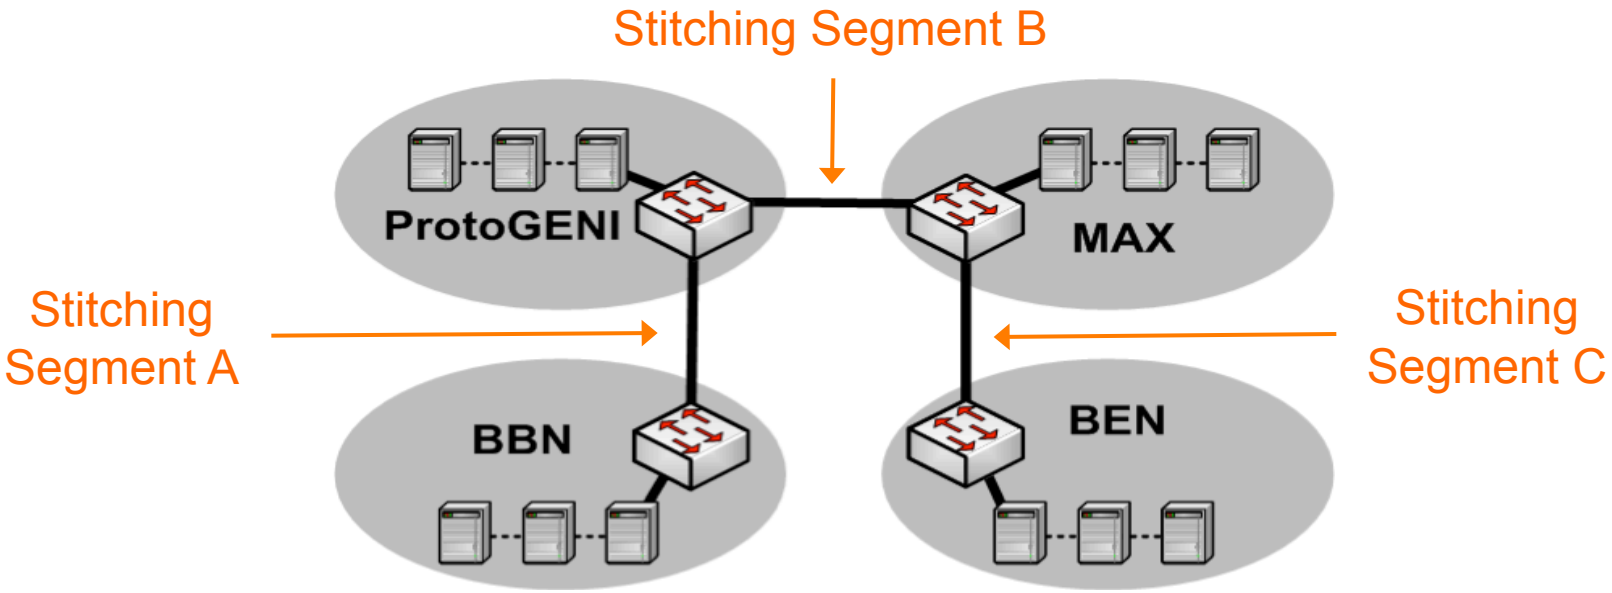

Figure 1 Use Case One Slice Topology

Stitching Example

- From the Topology Service (via the AM provided Stitching RSPEC) we can construct the below inter-connected aggregate view

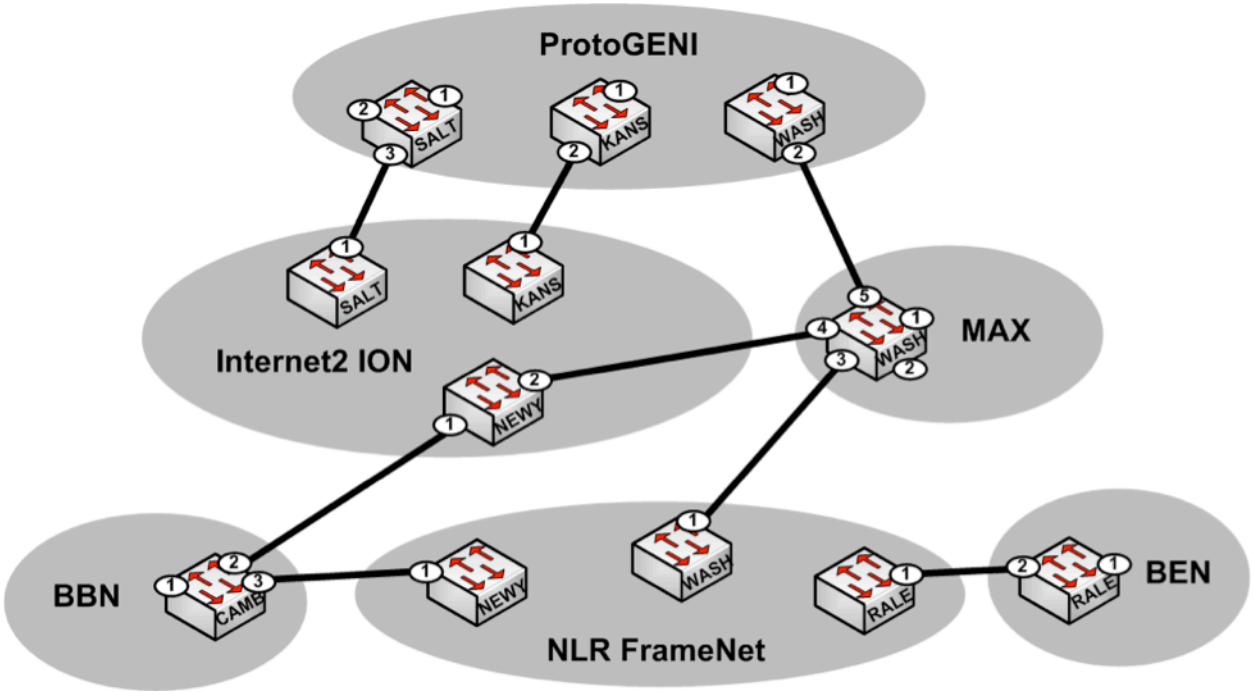

Figure 2 GENI Aggregate Topology

Stitching Example

- A client request may provide the following information:
  - Stitch A request:
    - Source: <bbn.com:CAMB:1>, VLAN=any
    - Destination: <protogeni.net:SALT:1>, VLAN=any
    - Bandwidth: 2 Gbps
  - Stitch B request:
    - Source: <geni.maxgigapop.net:wash:5>, VLAN=any
    - Destination: < protogeni.net:SALT:2>, VLAN=any
    - Bandwidth: 2 Gbps

Stitching Example

- A client request may be provide the following information (continued):
  - Stitch C request:
    - Source: <ben.renci.org:RALE:1>, VLAN=any
    - Destination: < geni.maxgigapop.net:wash:2>, VLAN=any
    - Bandwidth: 2 Gbps

Stitching Example

- The Path Computation Function may take the Client Request, and the topology information from the Topology service and construct this:

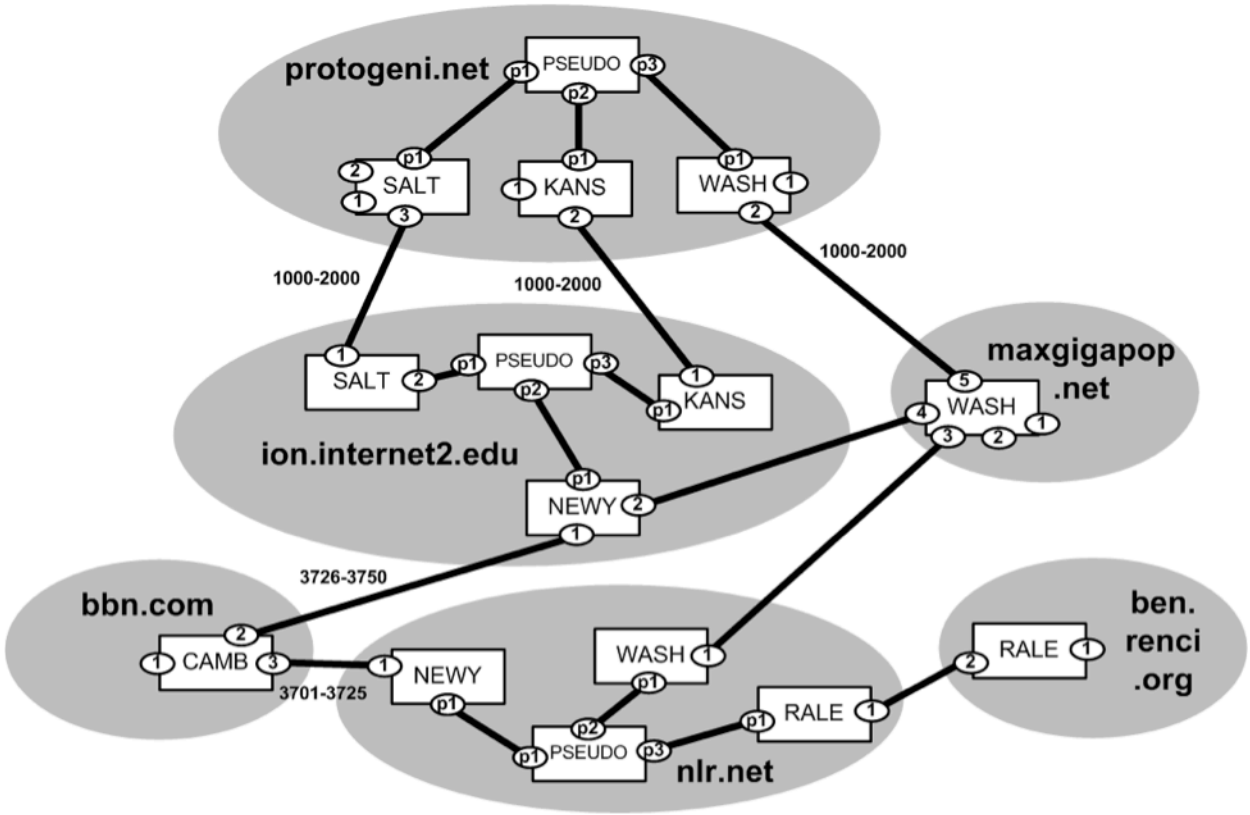

Figure 3 Aggregate Topology Converted for Path Computation

Stitching Example

- After path computation three results may be identified like this:

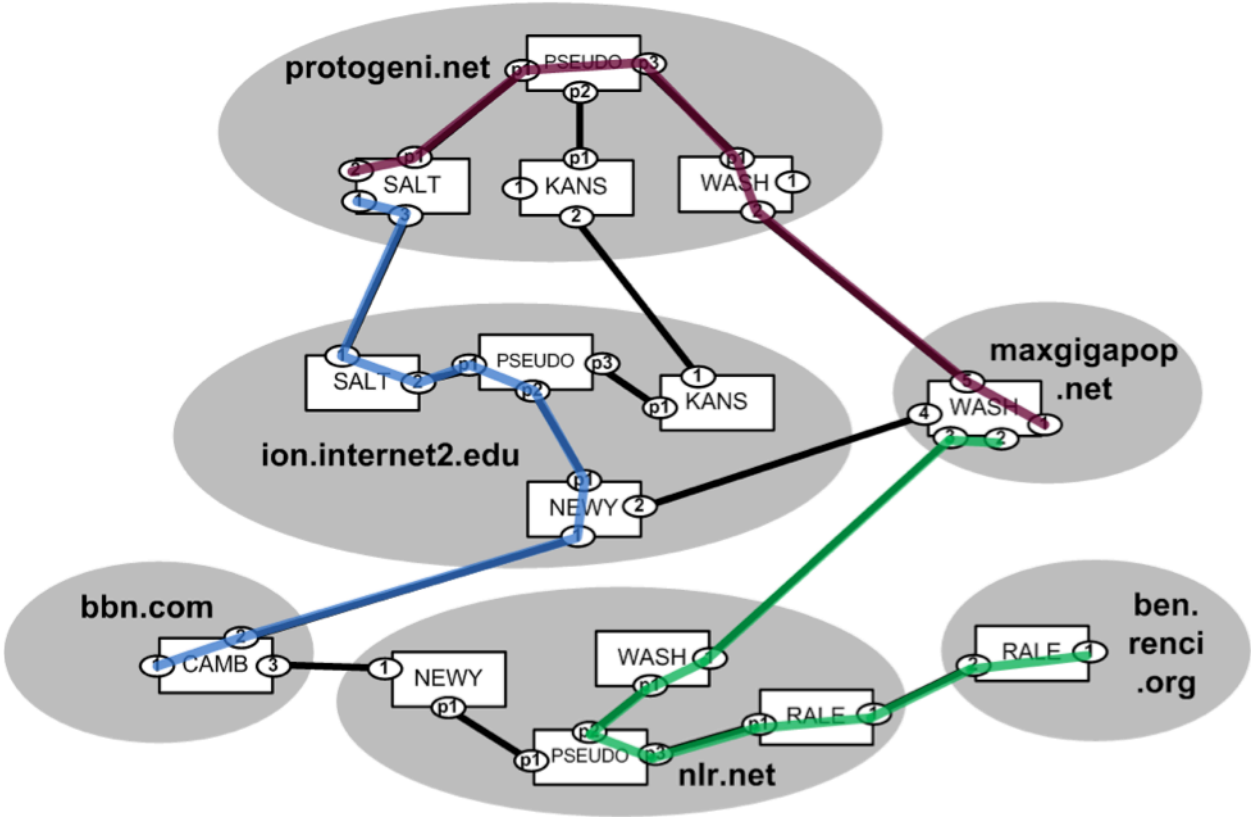

Figure 4 Path Computation Results

Stitching Example

- Stitch A path computation result:
  - Source: <bbn.com:CAMB:1>, VLAN=any
  - Destination: <protogeni.net:SALT:1>, VLAN=any
  - Bandwidth: 2 Gbps
  - Path Information:
    - <bbn.com:CAMB:1>
    - <bbn.com:CAMP:2>
    - <ion.internet2.edu:NEWY:1>
    - <ion.internet2.edu:SALT:1>
    - <protogeni.net:SALT:3>
    - <protogeni.net:SALT:1>

Stitching Example

- Stitch B path computation result:
  - Source: <geni.maxgigapop.net:wash:5>, VLAN=any
  - Destination: < protogeni.net:SALT:2>, VLAN=any
  - Bandwidth: 2 Gbps
  - Path Information:
    - <ben.renci.org:RALE:1>
    - <ben.renci.org:RALE:2>
    - <nlr.net:RALE:1>
    - <nlr.net:WASH:1>
    - <maxgigapop.net:WASH:3>
    - <maxgigapop.net:WASH:2>

Stitching Example

- Stitch C path computation result:
  - Source: <ben.renci.org:RALE:1>, VLAN=any
  - Destination: < geni.maxgigapop.net:wash:2>, VLAN=any
  - Bandwidth: 2 Gbps
  - Path Information:
    - <maxgigapop.net:WASH:1>
    - <maxgigapop.net:WASH:5>
    - <protogeni.net:WASH:2>
    - <protogeni.net:SALT:2>

Stitching Example

- stitching steps
- Stitch A workflow process shown
- Goes to ION last since it has VLAN translation

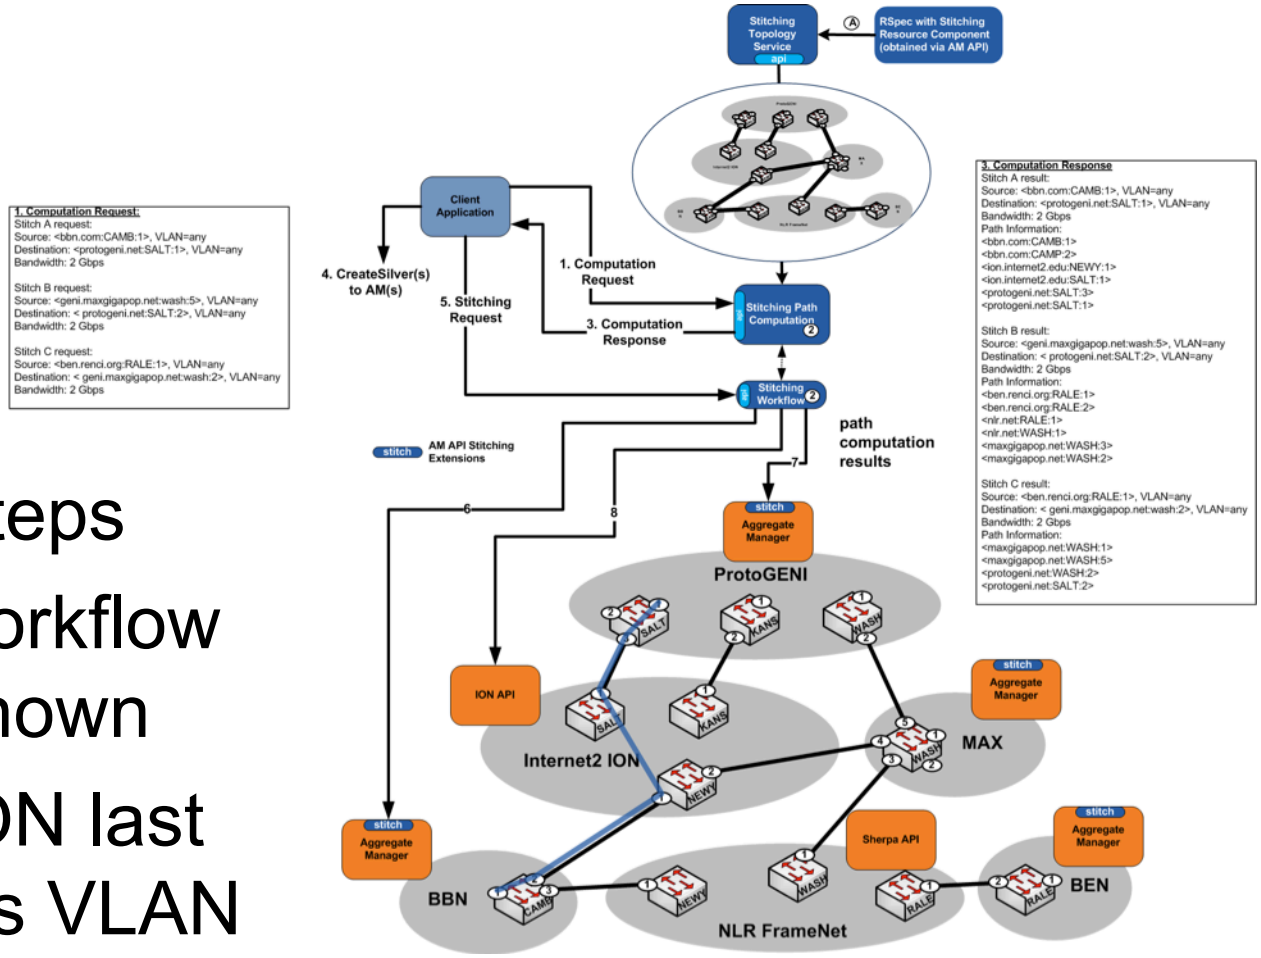

Figure 5 Use Case One Stitching Workflow