COMMON:
Coordinated Multi-layer Multi-domain Optical Network Framework for Large-scale Science Applications (2010-2013)

PI: Dr. Vinod Vokkarane

Associate Professor, University of Massachusetts at Dartmouth

(Currently on Sabbatical: Visiting Scientist, MIT)

Contact: vvokkarane@ieee.org

Project Website: http://www.cis.umassd.edu/~vvokkarane/common/

Supported by DOE ASCR
under grant DE-SC0004909

Brookhaven National Laboratory (BNL) January 12-13, 2012

Outline

- Introduction and project objectives
- Year 1 Objectives:
  - Anycast Multi-domain Service
  - Multicast-Overlay Algorithms
- Year 2 and 3 Project Objectives
  - Multi/Manycast-Overlay Deployment
  - Survivable Connections
  - QoS support
- Multi-domain Anycast Demo

COMMON Project Team

UMass Team

- Dr. Vinod Vokkarane (PI)
- Dr. Arush Gadkar (Post-Doc)
- Dr. Joan Triay (Visiting Fulbright Scholar)
- Bharath Ramaprasad (MS)
- Mark Boddie (BS-MS)
- Tim Entel (BS-MS)
- Jeremy Plante (Ph.D.)
- Thilo Schoendienst (Ph.D.)

ESnet/LBNL Team

- Chin Guok
- Andrew Lake
- Eric Pouyoul
- Brian Tierney

Introduction

- To support large-scale science applications we need to provision network resources across multiple layers and multiple domains.
- The network needs to provision connections between clients efficiently.
  - Immediate reservation (IR): network provisioning “immediately” when the connection request arrives.
  - Advance reservation (AR): resources can be reserved at some point in the future.

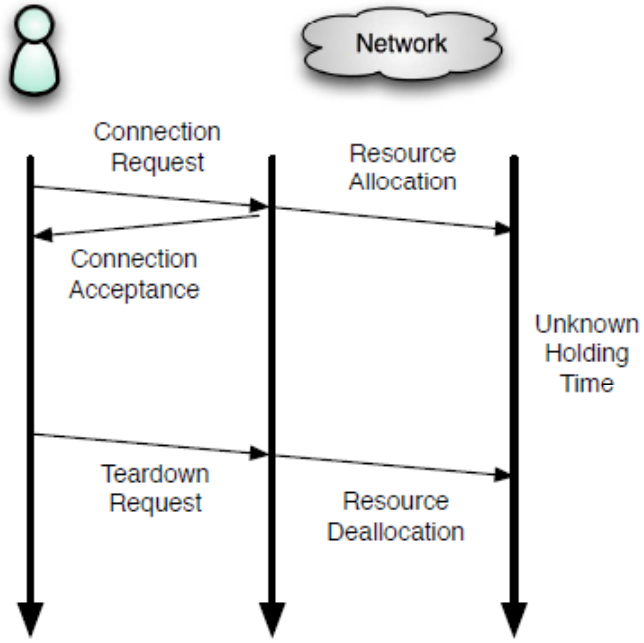

(a) Immediate Reservation

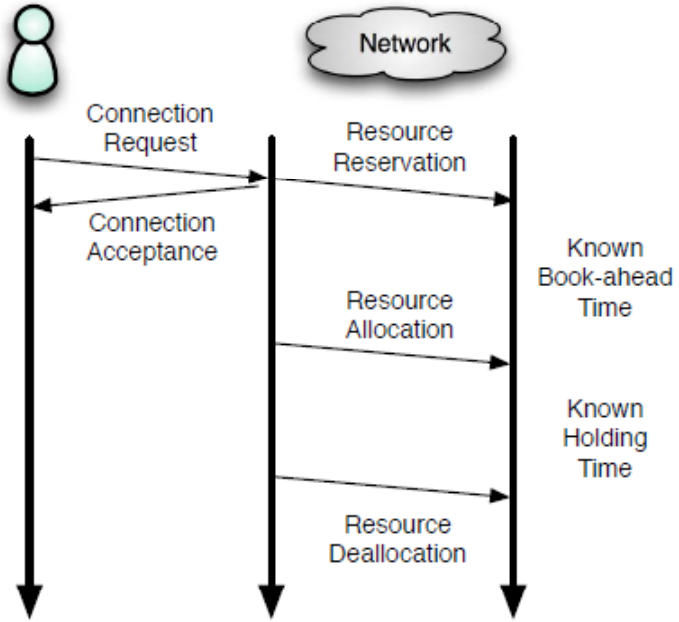

(b) Advance Reservation

Communication Services: Anycast

- Unicast Vs Anycast

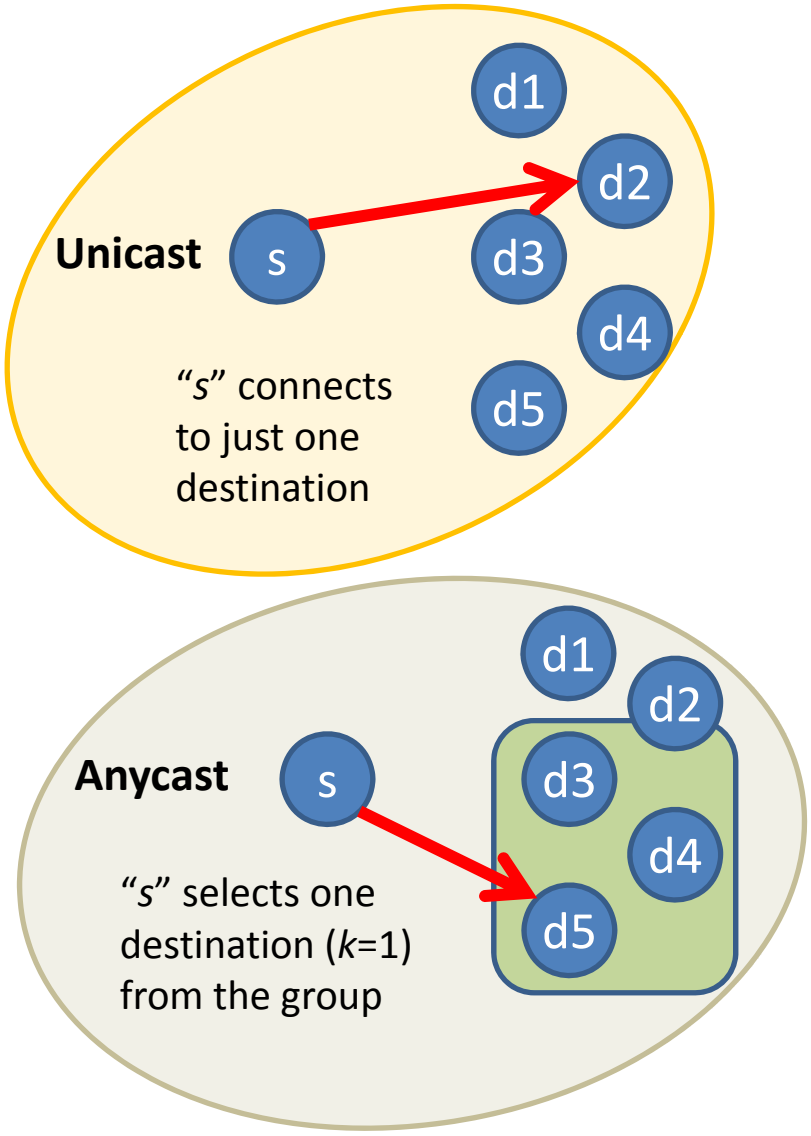

- Unicast request is represented by a tuple (*s,d*), where *s* is the source node and *d* is the destination node.
- Anycast request is represented by a tuple (*s,{D}*) where *s* is the source node and the {*D*} is the set of candidate destination nodes.
- In Anycast, the source node communicates with **any one** node from the set of candidate destination nodes.

Communication Services: Multicast/Manycast

- Multicast Vs Manycast

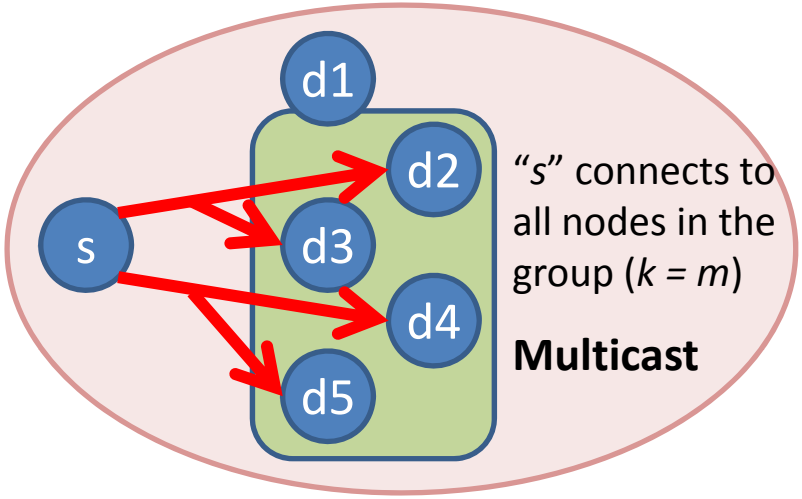

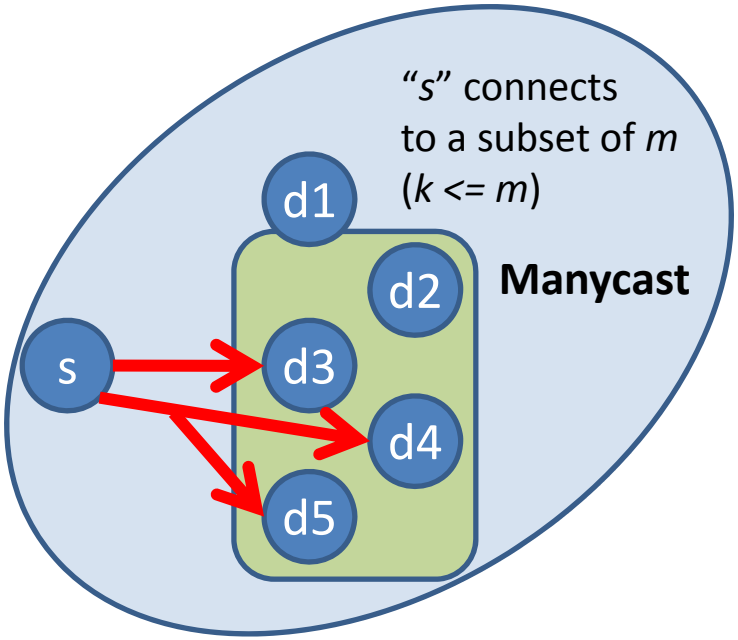

- Multicast request is represented by a tuple *(s,{D}),* where s is the source node and *D* is the set of destination nodes (d_1, d_2, …,d_m}.
- In Multicast, source node communicates with each destination node in {*D*}.
- Manycast request is represented by a tuple (*s,{D},k*) ,where *s* is the source node and the {*D*} is the set of candidate destination nodes. The source node communicates with any *k* nodes from the set {*D*}.

COMMON Project Objectives

- Design and Implement new services for advance reservation of network resources across multiple domains and multiple layers.

- COMMON Focus Areas:
  - Deploy Anycast AR algorithms on OSCARS (Year 1).
  - Develop Multi/Manycast Overlay models and deploy them on OSCARS (Y1-Y2).
  - Design and Deploy survivability techniques for multi-domain networks (Y2-Y3).
  - Design QoS mechanisms to support scientific applications on multi-domain networks and deploy them on OSCARS (Y2-Y3).
  - Extend QoS and Survivability mechanisms to multi-layer scenarios and deploy them on OSCARS (Y3).

Outline

- Introduction and project objectives
- Year 1 Deliverable: Anycast Multi-domain Service
- Year 2 and 3 Project Objectives
- Anycast Demo

Year 1: Deployment of Anycast Service on OSCARS

Objectives

- Design and implement a production-ready anycast service extension to existing OSCARS framework.
- Improve connection acceptance probability and user experience for anycast-aware services.

Impact

- Provide scientific community with ability to:
  - (a) Allow for destination-agnostic service hosting on large-scale networks.
  - (b) Increase service acceptance.

Design & Implementation (Complete)

- Designed anycast service as a PCE extension.
- Implementation of the PCE modules to find anycast connectivity, remove the unavailable resources, and select the best possible destination.
- Successfully completed Stress, Regression, and Integration testing of the anycast modules on OSCARS 0.6 (Q4, 2011).
- Hot deployment ready (PnP capable) anycast version of OSCARS 0.6 available at:

https://oscars.es.net/repos/oscars/branches/common-anycast/

- Plan to work with ESnet and ESG group to attach this service to a specific application.

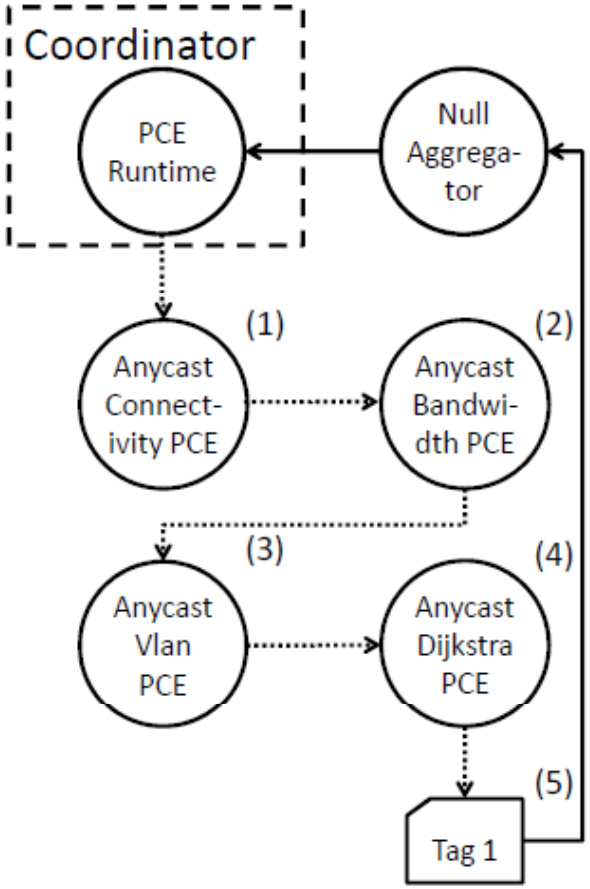

End-to-end Anycast Flowchart

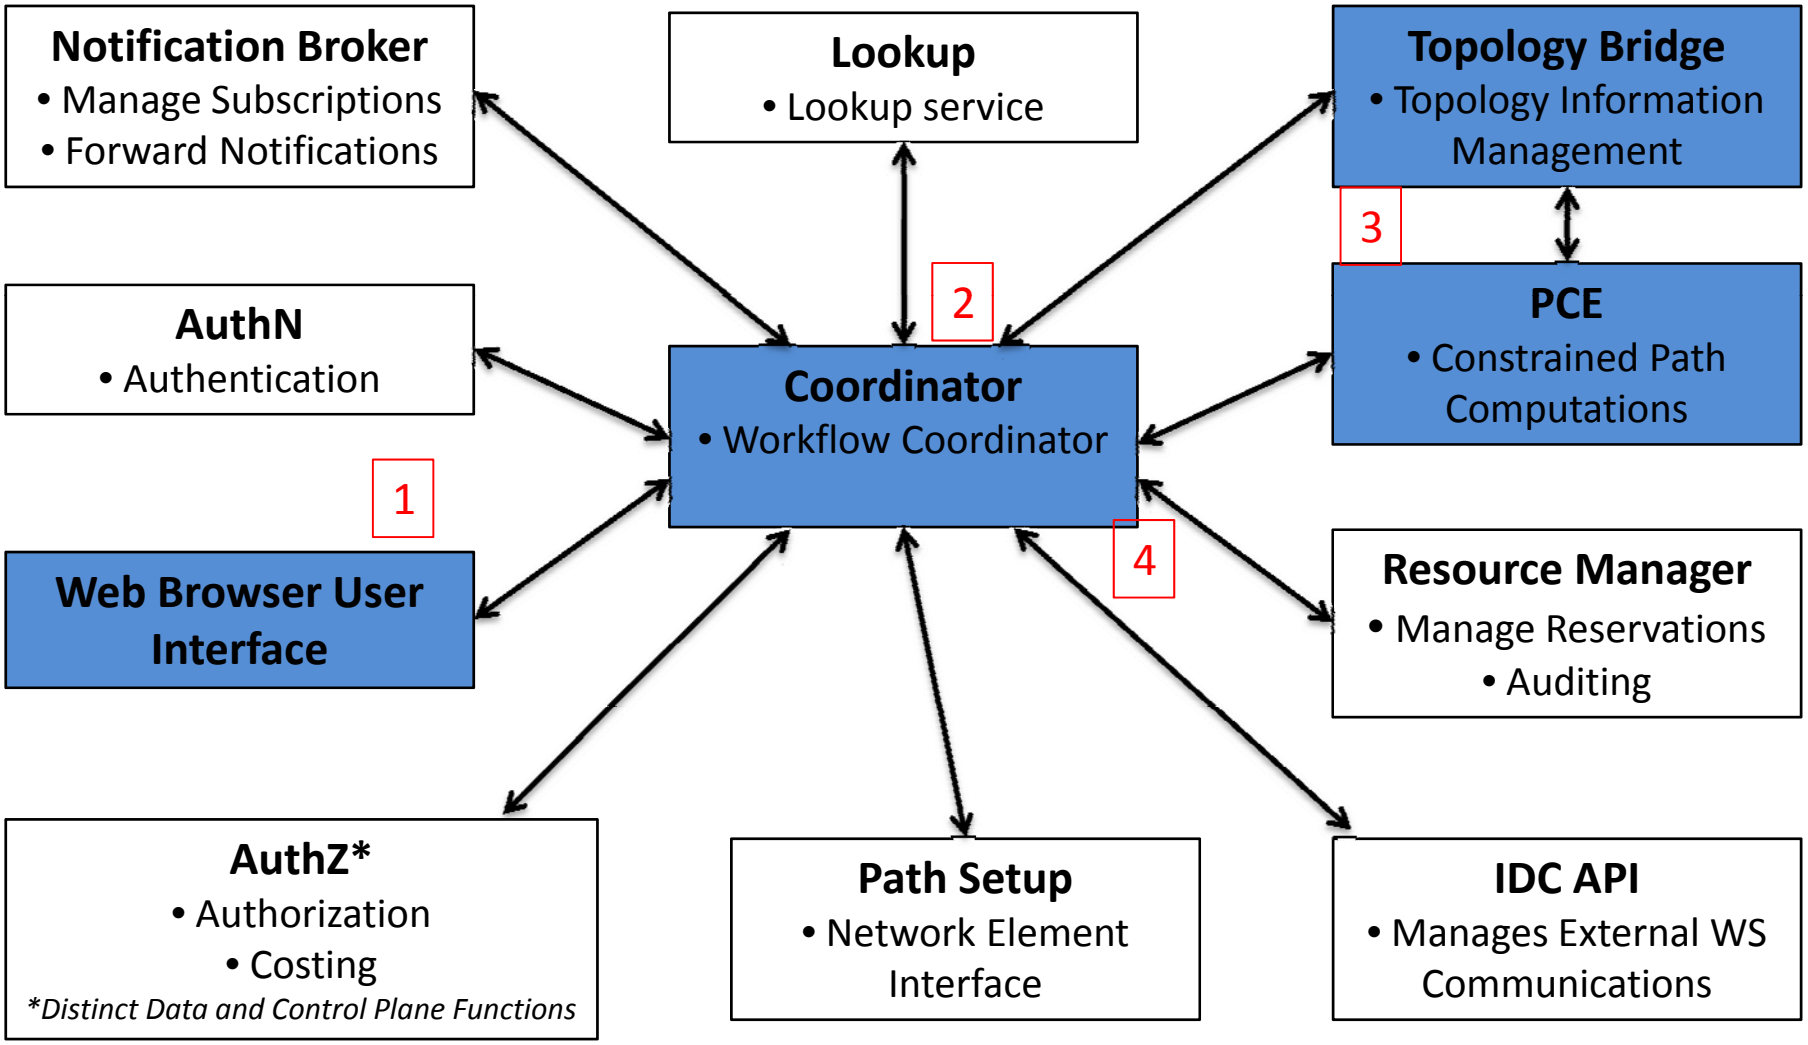

End-to-end Anycast Flowchart

- 1 The user interface servlets would process the anycast request as a unicast request with a big exception: the destination field will be a list of destination nodes (the anycast destination set).
  - An option is to encapsulate the anycast data as an OptionalConstraintType, in addition to the rest of parameters mapped into a UserRequestConstraintType. Both, UserRequestConstraintType and the OptionalConstraintType will be part of the ResCreateContent.
  - The ResCreateContent will be passed to the Coordinator to further process the anycast request.
- 2 The Coordinator, through the CreateReservationRequest, will get the ResCreateContent and map the user request constraints and optional constraints into a PCEData object.
  - The PCERuntime will handle the query process to the PCE.
- 3 The PCE (using the design proposed in the following slides) will make use of the OptionalConstraintType (which carries the list of destinations).
- 4 The result of the PCE will be the path from the source node to a single destination node, so, from the path reservation and PSS modules standpoint, the rest of the flowchart will work as a unicast request.

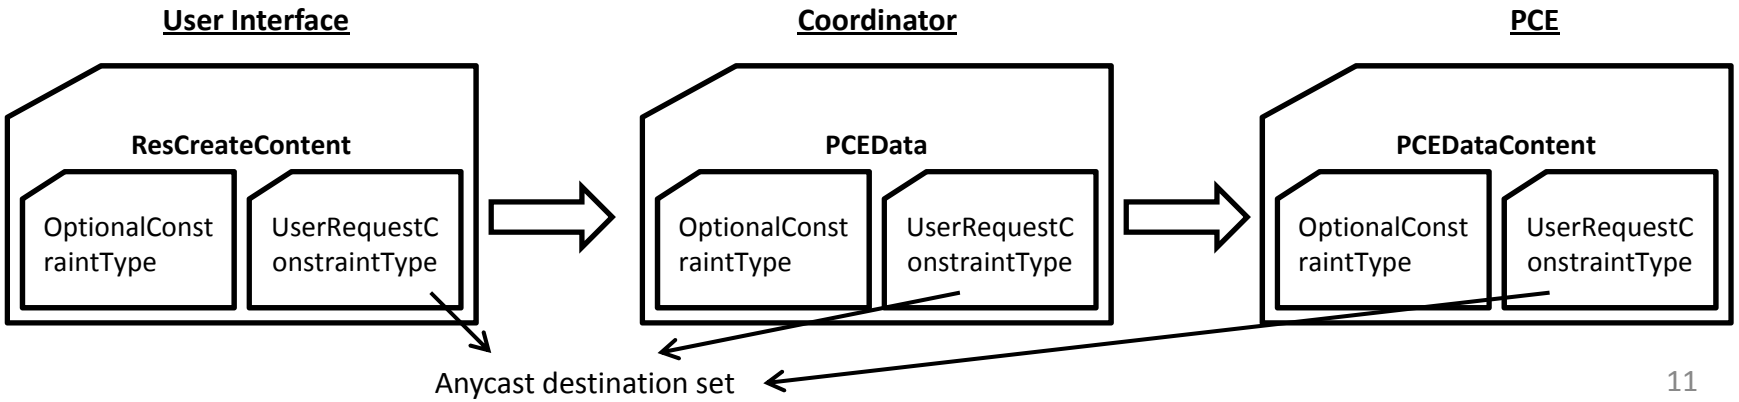

Anycast PCE – Design

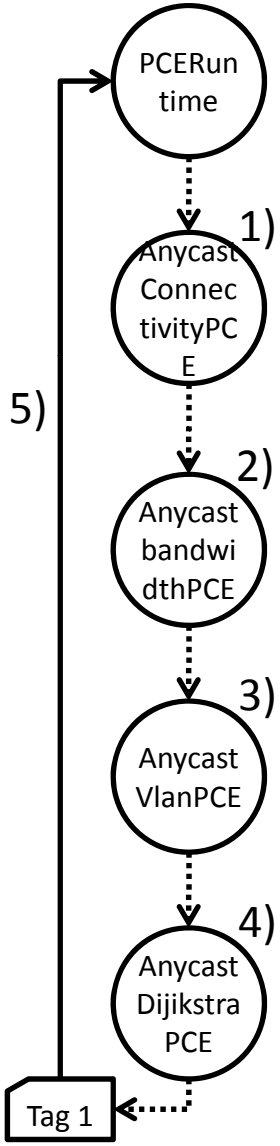

- In this design, we need to:
  - implement a new connectivityPCE and dijkstraPCE which are anycast-aware.

1) The AnycastConnectivityPCE computes the connectivity topology from the source to all destinations in the anycast set.
2) The constraint set is sent to the AnycastBandwidthPCE to check the bandwidth availability on the anycast topology.
3) The AnycastVlanPCE receives the pruned topology to the selected destination and checks for the vlanPCE.
4) The AnycastDijkstraPCE gets the constraints and pruned input topology from the bandwidthPCE and checks the shortest path for each destination in the anycast set. It also selects the final destination.
5) It then returns the final reply to the PCERuntime

Multi-domain Anycast

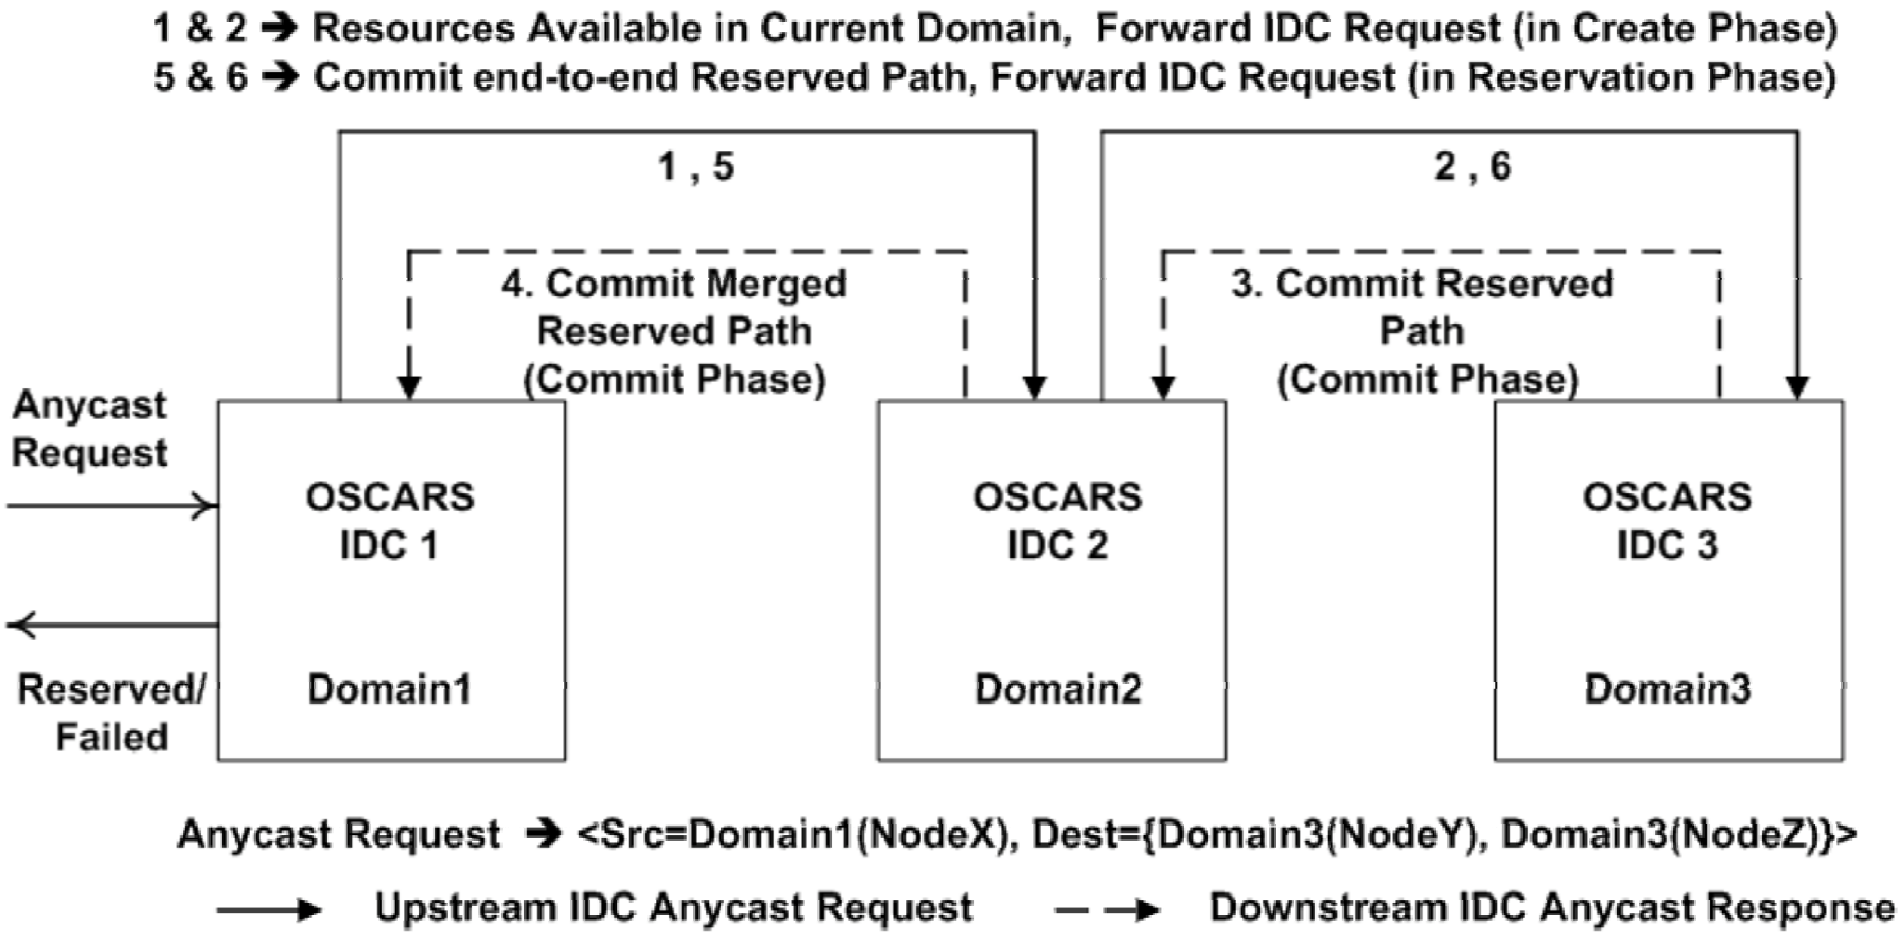

Performance of Anycast Service for OSCARS

Results for single domain

- We simulated 30 unique sets of 100 AR requests (and present the average values).
- All links are bi-directional and are assumed to have 1 Gb/s bandwidth.
- For each request, the source node and destination node(s) are uniformly distributed.
- Request bandwidth demands are uniformly distributed in the range [100 Mb/s, 500 Mb/s], in increments of 100 Mb/s.
- All requests are scheduled to reserve, transmit, and release network resources within two hours such that we stress test the network by increased traffic loads in this time frame.
- The *correlation factor* corresponds to the probability that requests overlap during that two-hour window.

Correlation	m=2	m=3	m=4	m=5
0.1	48.11	50.94	49.06	51.89
0.3	29.59	29.34	32.65	33.93
0.5	24.81	29.92	33.08	35.04
0.7	14.30	21.57	22.02	22.25
0.9	14.44	17.91	20.77	21.21
1.0	11.43	15.76	16.16	18.20

Percentage blocking reduction of anycast m/1 over unicast.

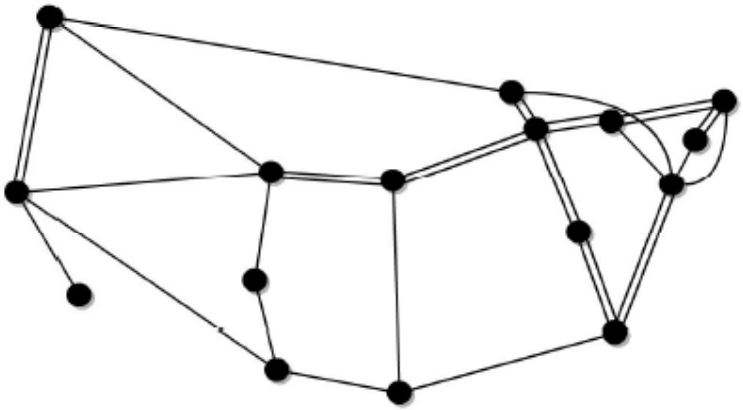

16-node ESnet SDN core network topology used in obtaining results.

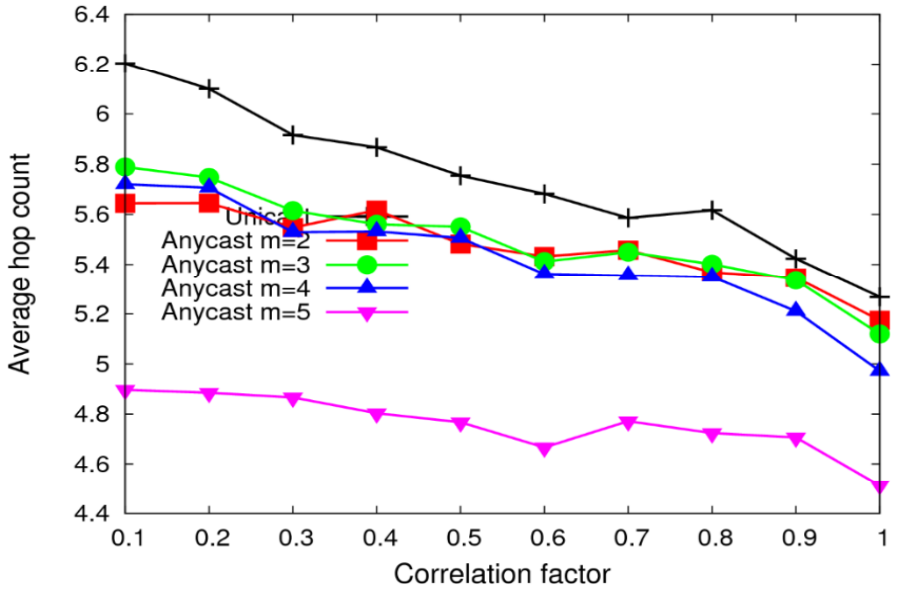

Average hop-count of successfully provisioned requests: unicast vs. anycast m/1 .

Performance of Anycast Service for OSCARS

Results for single domain

- We simulated 30 unique sets of 100 AR requests (and present the average values).
- All links are bi-directional and are assumed to have 10 Gb/s bandwidth.
- For each request, the source node and destination node(s) are uniformly distributed.
- Request bandwidth demands are uniformly distributed in the range [1 Gb/s, 5 Gb/s], in increments of 1 Gb/s.
- Anycast significantly outperforms unicast at all loads.

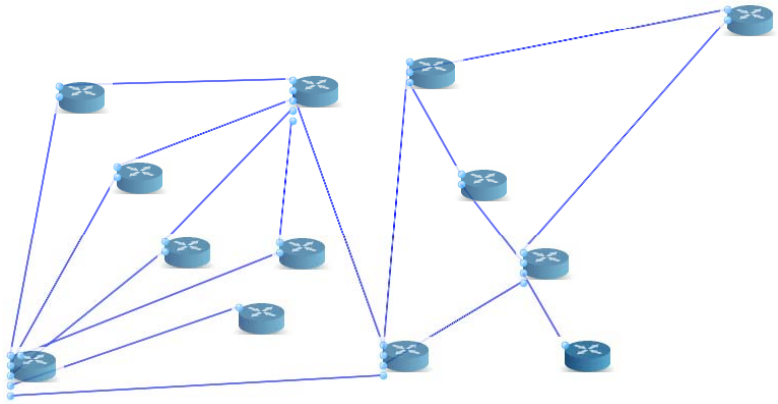

13-node GÉANT network topology used in obtaining results.

Blocking probability of various services (unicast vs. anycast *m*/1)

Correlation	*m*=2	*m*=3	*m*=4	*m*=5
0.1	39.18	42.54	42.91	42.91
0.3	17.35	21.02	23.26	24.70
0.5	11.00	15.70	17.38	18.98
0.7	10.44	12.33	12.92	15.07
0.9	4.056	6.82	7.11	8.23
1.0	3.924	5.50	7.08	6.87

Percentage blocking reduction of anycast *m*/1 over unicast.

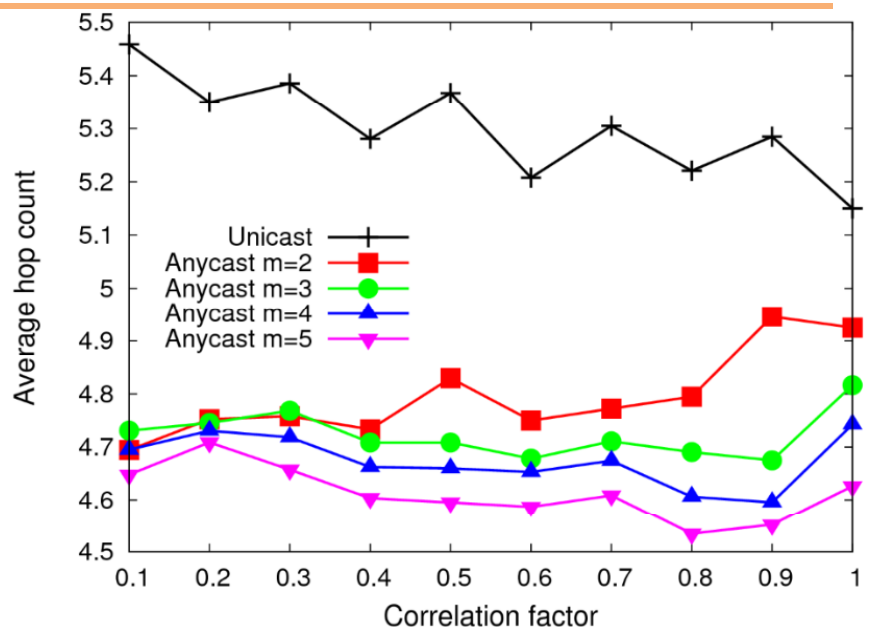

Average hop-count of successfully provisioned requests: unicast vs. anycast *m*/1 .

Performance of Anycast Service for OSCARS

Results for multi-domain

- We simulated 5 unique sets of 50 AR requests (and present the average values).
- All links are bi-directional and are assumed to have 10 Gb/s bandwidth.
- Each request, has source node in ESnet and destination node(s) in GEANT.
- Request bandwidth demands are uniformly distributed in the range [1000 Mb/s, 5000 Mb/s], with step granularity of 1000 Mb/s.
- 2 inter-domain links between ESnet and GEANT.
- Remaining assumptions similar to single domain.

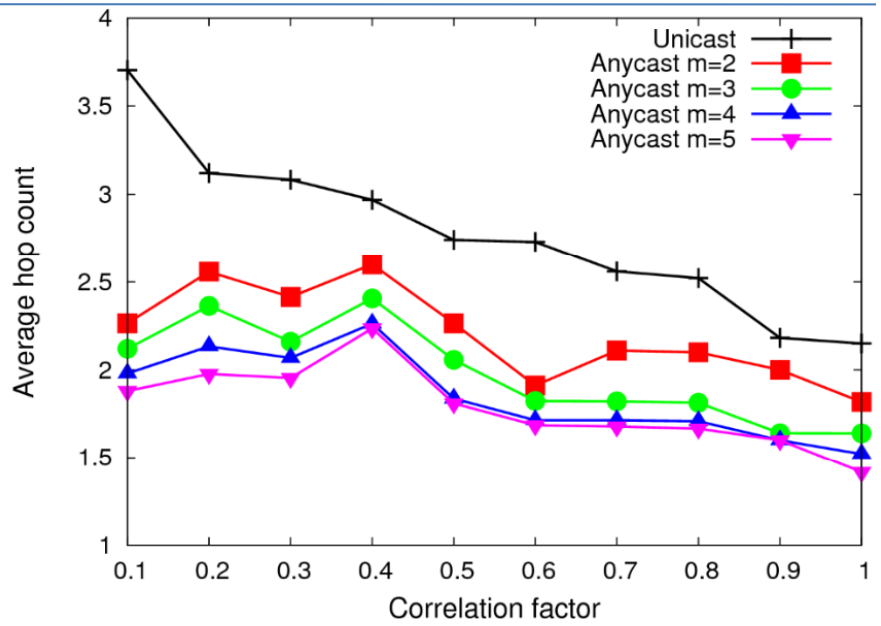

Average hop-count of successfully provisioned requests: unicast vs. anycast *m*/1 .

Correlation	*m*=2	*m*=3	*m*=4	*m*=5
0.1	66.67	100.0	100.0	100.0
0.3	35.71	50.00	57.14	57.14
0.5	15.79	21.05	31.58	31.58
0.7	4.00	12.00	12.00	12.00
0.9	3.57	10.71	10.71	10.71
1.0	6.67	10.00	10.00	13.33

Percentage blocking reduction of anycast *m*/1 over unicast.

Outline

- Introduction and project objectives
- Year 1 Deliverable: Anycast Multi-domain Service
- Year 2 and 3 Project Objectives
- Anycast Demo

Year 1-2: Deployment of Multi/Manycast Overlay on OSCARS

Objectives

- To support point-to-multipoint connections.
- To develop an overlay model to support Multicast/Manycast communication paradigms over point-to-point unicast connections in OSCARS.

Impact

- Allow scientific community the ability to:
  - (a) Use a multicast service and increase the service acceptance.
  - (b) Provide different connection setup choices with different quality of service (QoS) to the scientists.

Progress

- Proposed two overlay models: Drop at member node (DAMN) and Drop at any node (DAAN).
- Compared the performance of above models to the naïve method of supporting multicast (MVWU).
- Conducted simulations on 14-node ESnet.
- Blocking performance results show significant improvement due to DAMN and DAAN algorithms
- Both these overlay models will be tested and integrated into the OSCARS system (Year 2).

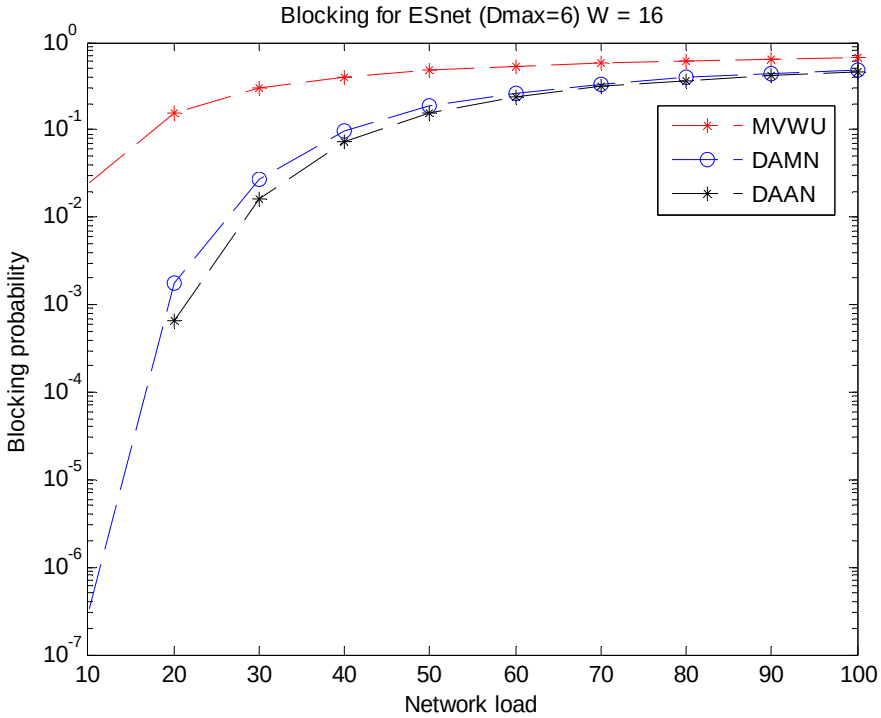

Year 1-2: Deployment of Multi/Manycast Overlay on OSCARS

- Two Multicast/Manycast overlay approaches proposed to provide point-to-multipoint (P2MP) communication over unicast-only optical/VLAN layer.
  - **MVWU – single logical-hop overlay**
    - Source of the multicast/manycast request establishes an independent lightpath (VC) to each destination.
    - It is possible that these lightpaths overlap, thus making inefficient use of available bandwidth.
    - This can lead to unnecessarily high connection blocking.
  - **DAMN/DAAN – multiple logical-hop overlay**
    - Source of the request establishes a lightpath (or VC) to one destination.
    - The next destination can be reached by a lightpath (or VC) from the source or from the first destination.
    - Creates a Steiner tree routing scheme wherein each lightpath (or VC) may be viewed as a hop in the logical overlay layer.
- We plan to implement both overlay mechanisms on OSCARS.

Year 1-2 development: MVWU vs. DAMN:
MVWU : 3 Wavelengths DAMN: 1 Wavelength

Consider 2 Multicast requests: R1 → 1{2, 5, 6} and R2 → 4{2, 3, 5}

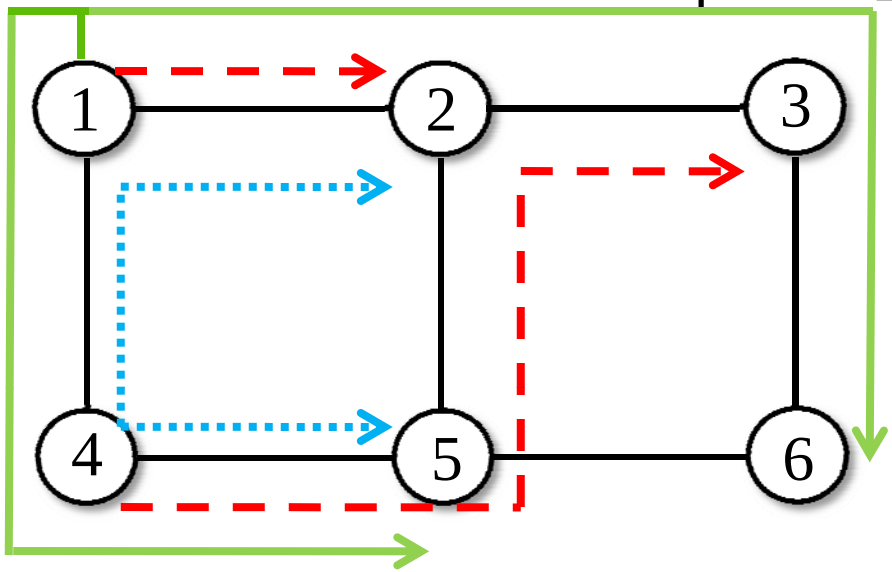

MVWU: Routing in Physical Topology

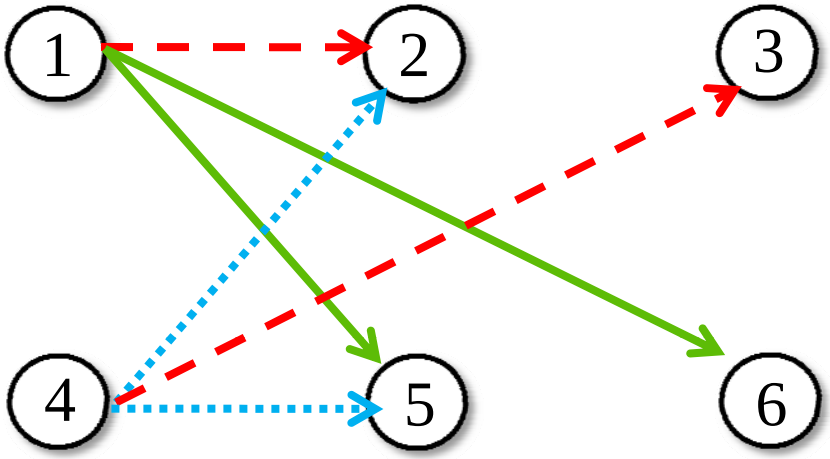

MVWU: Logical Topology

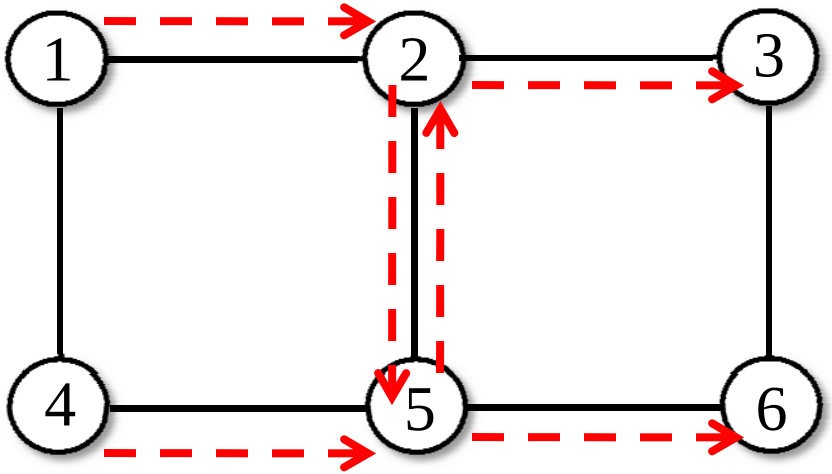

DAMN: Routing in Physical Topology

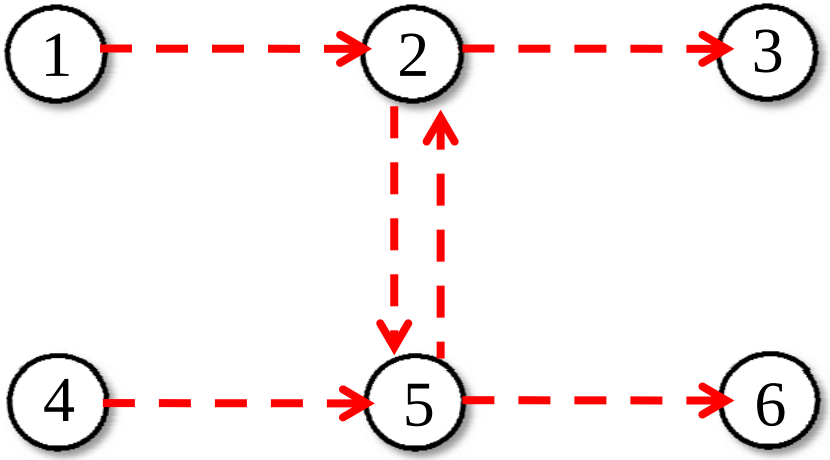

DAMN: Logical Topology

Year 2-3: Design Survivability Techniques for use with OSCARS

Objectives

- Design and implement survivability techniques on OSCARS using path protection and destination relocation.
- Extend this feature to a multi-domain and multi-layer network.

Impact

- Additional resources are reserved for each request
- Users are protected from single-link failures and destination failures.

Strategy

- We plan to implement basic survivability techniques to protect against link and node failures.
  - **Path protection** enables survivable connections by reserving a link-disjoint backup path.
  - OSCARS will provision both primary and link-disjoint backup paths for each connection request using a new ProtectionPCE stack.
- We plan to extend our anycast PCE design to account for single-link/node failure via anycast path protection and anycast protection using destination relocation.
  - **Anycast path protection** allows the ProtectionPCE to pick the anycast destination which has the least-cost link-disjoint path pair from the source.
  - **Anycast relocation** allows for the backup path to be routed to an alternate destination in the anycast set. Example: R(S, {D_1,D_2}) ➔ Primary: S-> D_1; Backup: S->D_2.

Year 2-3: Design Survivability Techniques for use with OSCARS

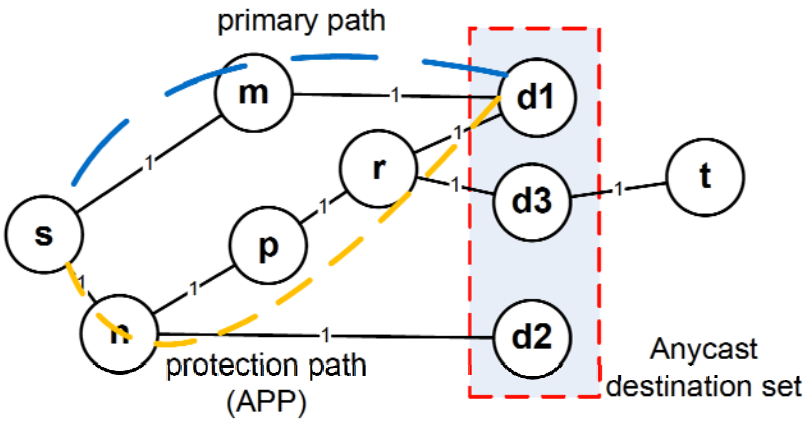

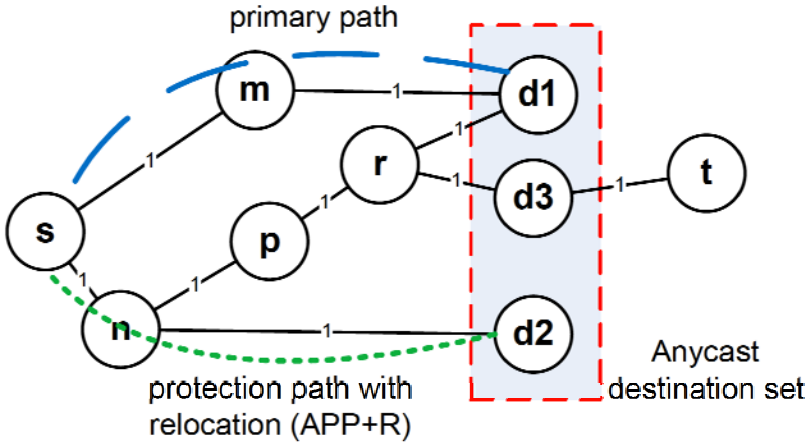

Fig A: Anycast Path Protection

Fig B: Anycast Protection with Destination Relocation

Fig A: the backup path in yellow is routed to the same anycast destination.

Fig B: shows how relocation can reduce resource consumption and reduce potential blocking by allowing more flexibility in the choice of the backup path.

We will look in to implementation of dynamic **restoration** mechanisms for all the above survivability techniques.

Year 2-3: Supporting Quality of Service (QoS) in OSCARS

Objectives

- To provide for user-profile based access to network resources like domains, nodes, ports, and links while provisioning connection requests.
- To support multi-layer and multi-constraint restrictions on requests based on user privileges.

Impact

- Classification of user requests based on service requirements.
- Preferential treatment for high-priority requests originating from data/resource-intensive scientific applications.

Progress

- Designing QoS user profile database extensions based on existing authentication & authorization databases.
- Research and Development on traffic policing and shaping at a request provisioning level driven by user profiles. This can be done using several metrics that balance link bandwidth and VLAN availability using policies like least-used, most-used, most recently used, least recently used, and random.

Year 2-3: What-if Driven Multi-Constrained/Layered OSCARS

Objectives

- Multi-domain offline/inline negotiation protocol for querying and/or reserving the best possible circuit by providing different viable reservation solutions which best suits the user, based on the user profile.
- Rank the various viable connection paths using key performance indicators (KPIs) to best suit the application requirements.

Impact

- Solution matches to user/application requirements as several viable reservation solutions are ranked.
- Re-attempts of reservation for failed reservation requests are minimized.

Progress (in collaboration with ARCHSTONE)

- Implementation in progress of the What-if engine and What-if driven GUI for OSCARS.
- Implementation of a separate offline and inline query + reservation workflow for what-if for multi-domain also in progress.

Year 2: What-if Driven Multi-Constrained/Layered OSCARS

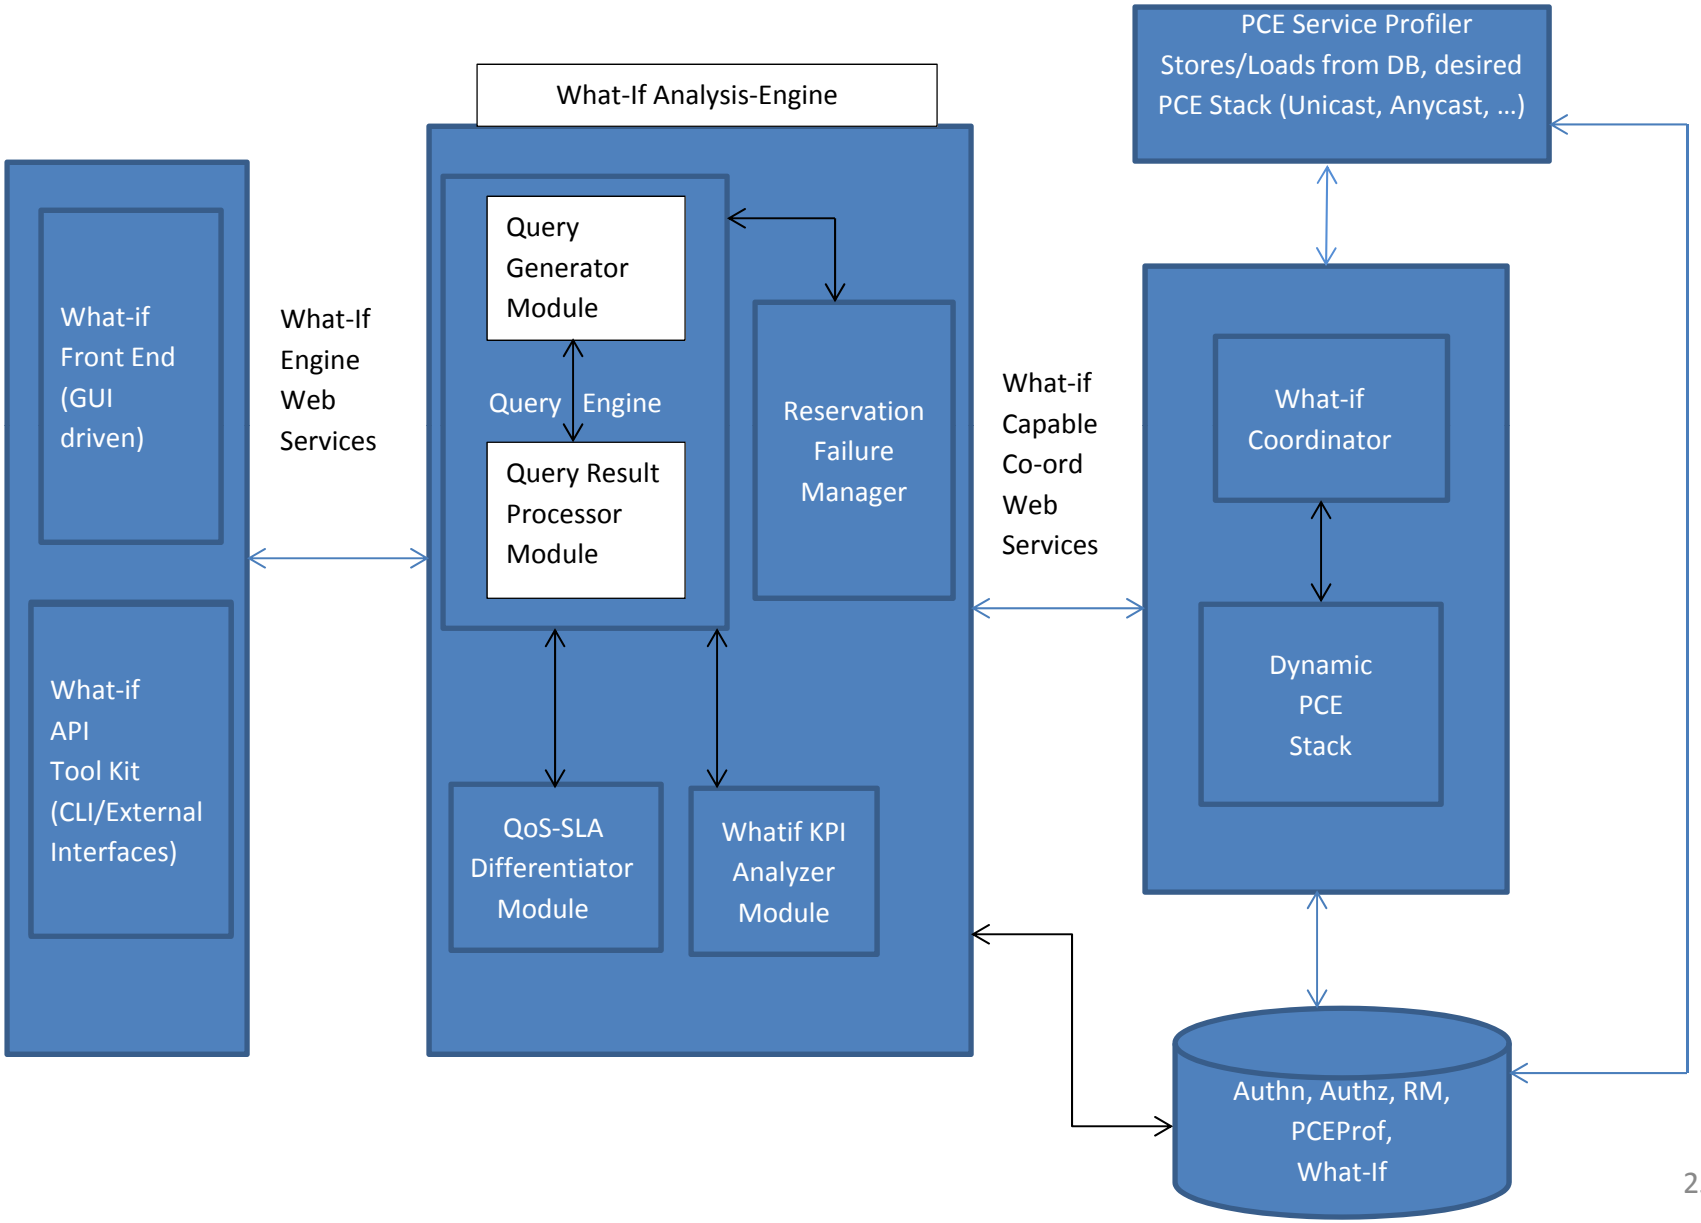

Research Papers and Journals relevant to COMMON progress

Year 1: Anycast and Advance Reservation

[**1**] Mark Boddie, Timothy Entel, Chin Guok, Andrew Lake, Jeremy Plante, Eric Pouyoul, Bharath H. Ramaprasad, Brian Tierney, Joan Triay, and Vinod M. Vokkarane, "On Extending ESnets OSCARS with a Multi-Domain Anycast Service", *submitted to ONDM 2012*.

[**2**] Bharath H. Ramaprasad, Arush Gadkar, and Vinod M. Vokkarane, "Dynamic anycasting over wavelength routed networks with lightpath switching," IEEE *High Performance Switching and Routing (HPSR), 2011,* July 2011.

[**3**] Neal Charbonneau, Arush G. Gadkar, Bharath H. Ramaprasad, and Vinod M. Vokkarane, "Dynamic Circuit Provisioning in All-Optical WDM Networks Using Lightpath Switching," *accepted, Elsevier Optical Switching and Networking, Special Issue on IEEE ANTS 2010*, Nov. 2011.

[**4**] Neal Charbonneau, Chin Guok, Inder Monga, and Vinod M. Vokkarane, "Advance Reservation Frameworks in Hybrid IP-WDM Networks," IEEE Communications Magazine, Special Issue on Hybrid Networking: Evolution Towards Combined IP Services and Dynamic Circuit Network Capabilities, May 2011.

Year 2: Multicast/Manycast Overlay & QoS in OSCARS

[**5**] Arush Gadkar, Jeremy Plante, and Vinod Vokkarane, "Static Multicast Overlay in WDM Unicast Networks for Large-Scale Scientific Applications," *Proceedings, IEEE ICCCN 2011*, Maui, Hawaii, August 2011.

[**6**] Arush Gadkar and Jeremy Plante, "Dynamic Multicasting in WDM Optical Unicast Networks for Bandwidth-Intensive Applications," *Proceedings, IEEE Globecom 2011*, Houston, Texas, December 2011.

[**7**] Arush Gadkar, Jeremy Plante, and Vinod Vokkarane, "Manycasting: Energy-Efficient Multicasting in WDM Optical Unicast Networks," *Proceedings, IEEE Globecom 2011*, Houston, Texas, December 2011.

[**8**] Jeremy Plante, Arush Gadkar, and Vinod Vokkarane, "Dynamic Manycasting in Optical Split-Incapable WDM Networks for Supporting High-Bandwidth Applications," *to appear, Proceedings, IEEE International Conference on Computing, Networking and Communications (ICNC)*, Maui, Hawaii, February 2012.

[**9**] Jeremy Plante, Arush Gadkar, and Vinod Vokkarane, "Multicast Overlay for High-Bandwidth Applications", *submitted to the IEEE Journal of Optical Communications and Networking (JOCN)*.

[**10**] J. Triay, C. Cervell´o-Pastor, and V. M. Vokkarane, "Analytical Model for Hybrid Immediate and Advance Reservation in Optical WDM Networks*," in Proc. of IEEE GLOBECOM 2011.*

[**11**] J. Triay, C. Cervell´o-Pastor, and V. M. Vokkarane, "Computing approximate blocking probabilities for hybrid immediate and advance reservation in optical WDM networks," *submitted to IEEE/ACM Transactions on Networking, 2011.*

Outline

- Introduction and project objectives
- Year 1 Deliverable: Anycast Multi-domain Service
- Year 2 and 3 Project Objectives
- Multi-Domain Anycast Demo

Multi-domain Anycast Communication for OSCARS 0.6

- Scenario 1) Connection Blocking
- Scenario 2) Resource Consumption (Hop Count)
- Evaluate blocking, resource consumption, and request provisioning time.

Scenario 1 – Connection Blocking

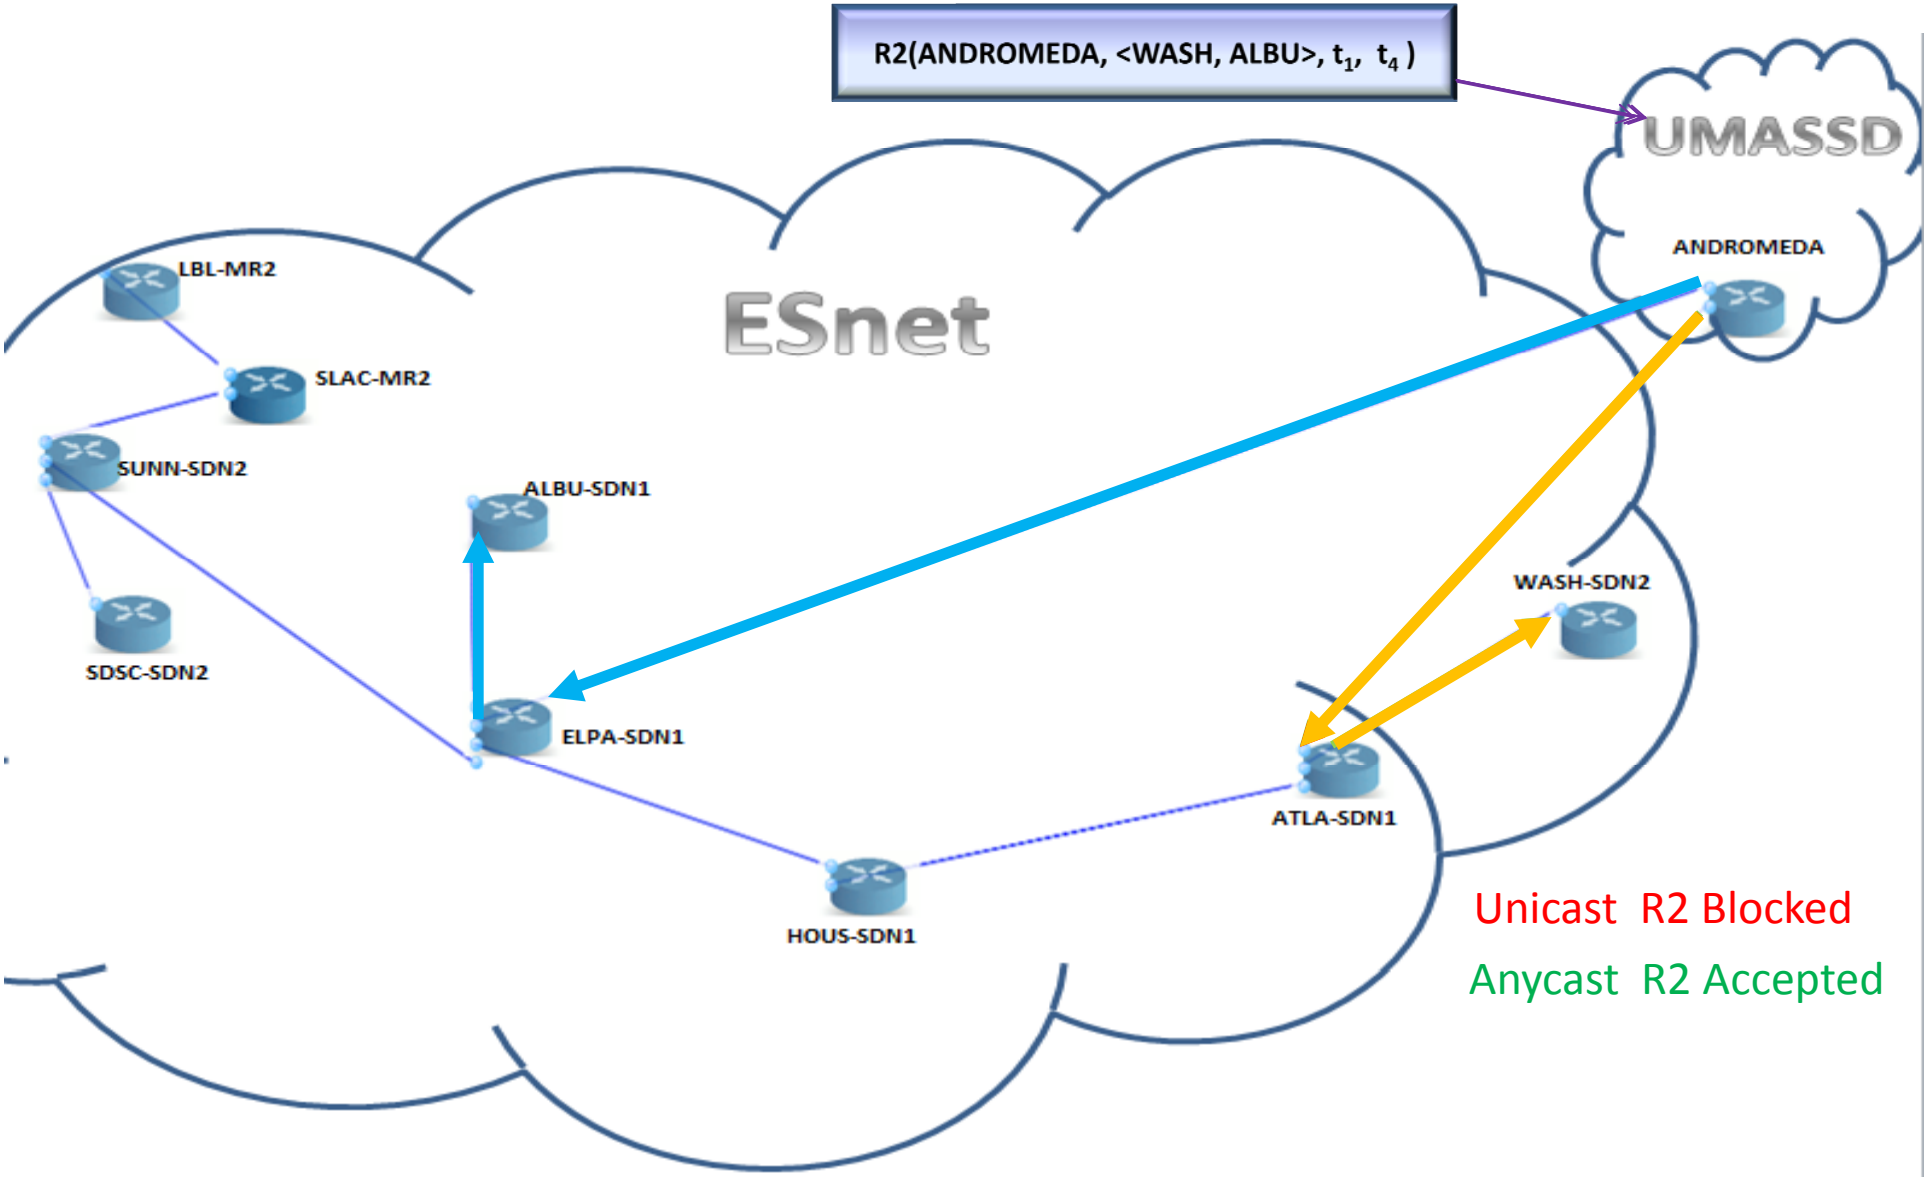

Unicast Requests experience high blocking

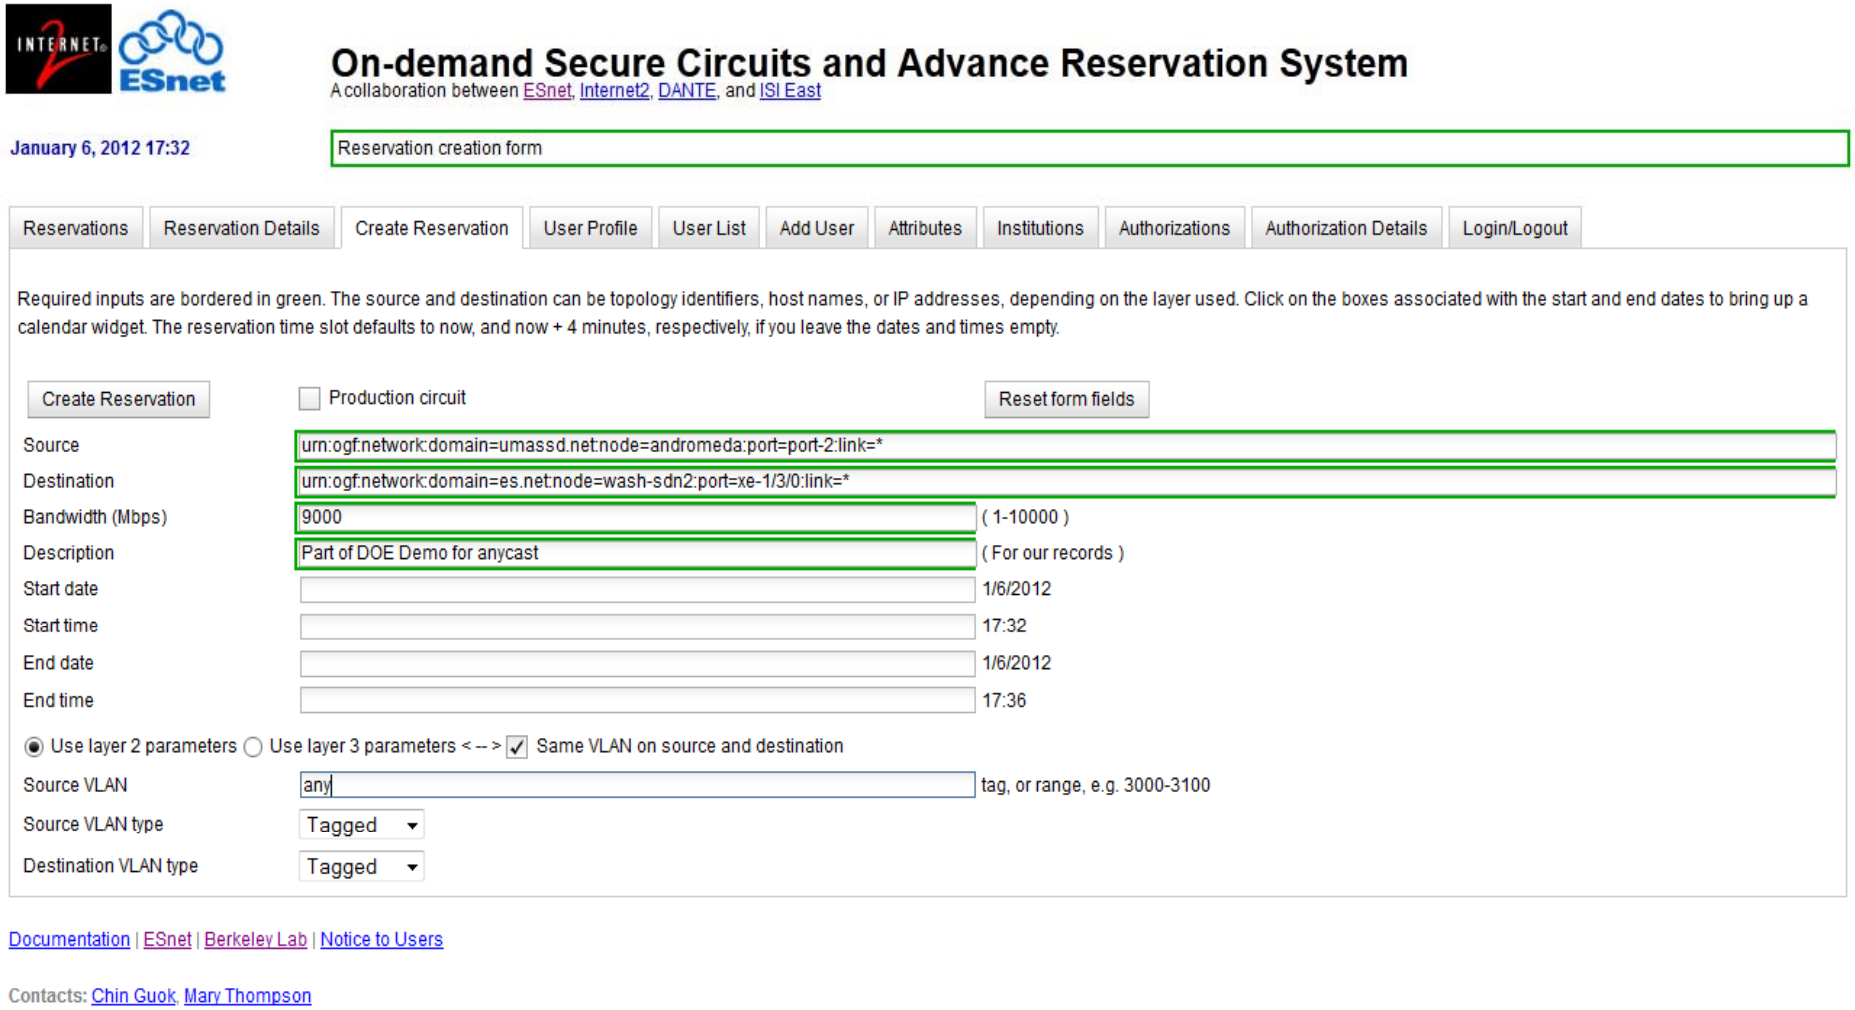

On-demand Secure Circuits and Advance Reservation System

A collaboration between ESnet, Internet2, DANTE, and ISI East

January 6, 2012 17:32

Reservation creation form

Reservations | Reservation Details | Create Reservation | User Profile | User List | Add User | Attributes | Institutions | Authorizations | Authorization Details | Login/Logout

Required inputs are bordered in green. The source and destination can be topology identifiers, host names, or IP addresses, depending on the layer used. Click on the boxes associated with the start and end dates to bring up a calendar widget. The reservation time slot defaults to now, and now + 4 minutes, respectively, if you leave the dates and times empty.

Create Reservation | Production circuit | Reset form fields

Source	urn:ogf:network:domain=umassd.net:node=andromeda:port=port-2:link=*	
Destination	urn:ogf:network:domain=es.net:node=wash-sdn2:port=xe-1/3/0:link=*	
Bandwidth (Mbps)	9000	(1-10000)
Description	Part of DOE Demo for anycast	(For our records)
Start date		1/6/2012
Start time		17:32
End date		1/6/2012
End time		17:36

Use layer 2 parameters | Use layer 3 parameters < -- > | Same VLAN on source and destination

Source VLAN	any	tag, or range, e.g. 3000-3100
Source VLAN type	Tagged	
Destination VLAN type	Tagged	

Documentation | ESnet | Berkeley Lab | Notice to Users

Contacts: Chin Guok, Mary Thompson

Unicast Requests experience high blocking

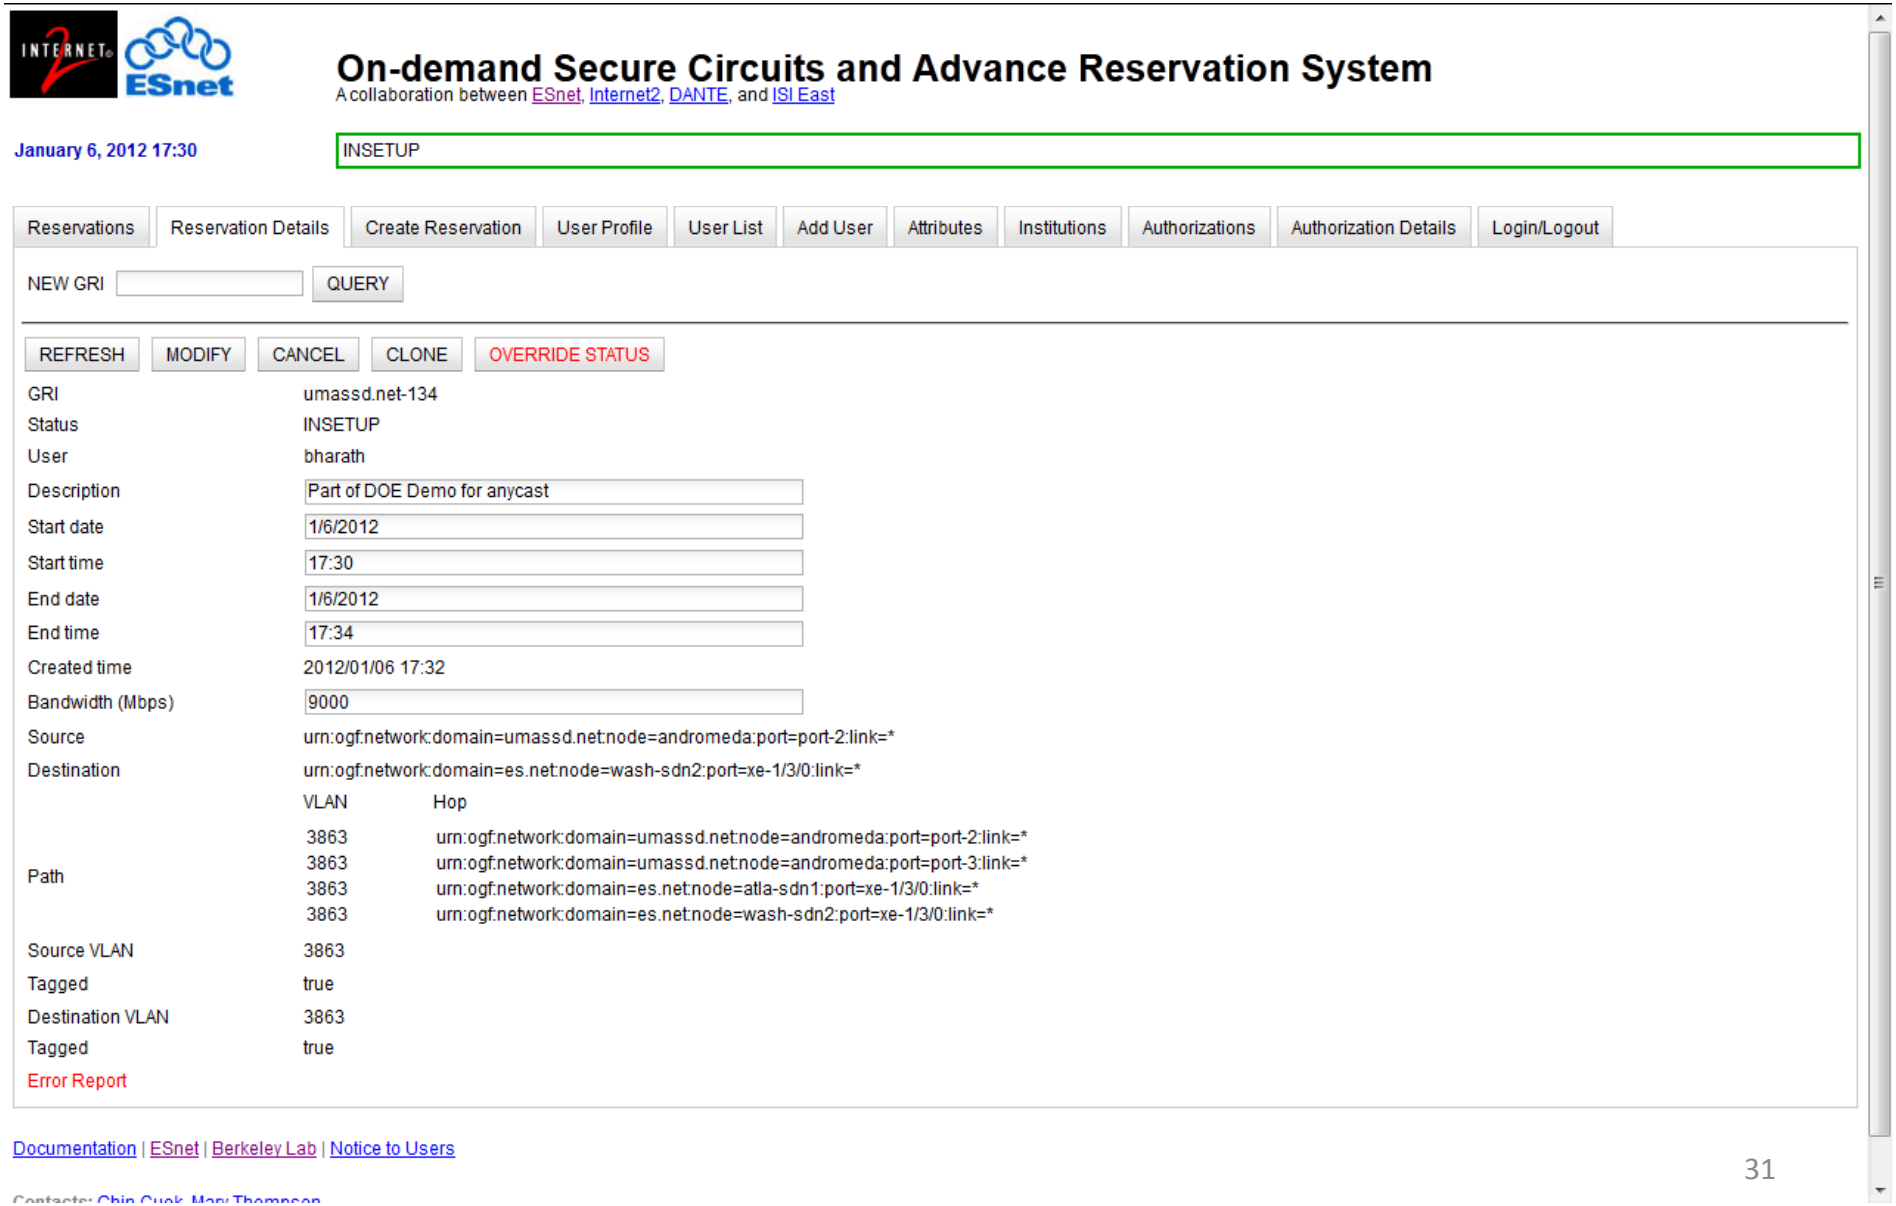

On-demand Secure Circuits and Advance Reservation System

A collaboration between ESnet, Internet2, DANTE, and ISI East

January 6, 2012 17:30 INSETUP

Reservations | Reservation Details | Create Reservation | User Profile | User List | Add User | Attributes | Institutions | Authorizations | Authorization Details | Login/Logout

NEW GRI [ ] QUERY

REFRESH | MODIFY | CANCEL | CLONE | OVERRIDE STATUS

GRI	umassd.net-134
Status	INSETUP
User	bharath
Description	Part of DOE Demo for anycast
Start date	1/6/2012
Start time	17:30
End date	1/6/2012
End time	17:34
Created time	2012/01/06 17:32
Bandwidth (Mbps)	9000
Source	urn:ogf:network:domain=umassd.net:node=andromeda:port=port-2:link=*
Destination	urn:ogf:network:domain=es.net:node=wash-sdn2:port=xe-1/3/0:link=*

Path	VLAN	Hop
	3863	urn:ogf:network:domain=umassd.net:node=andromeda:port=port-2:link=*
	3863	urn:ogf:network:domain=umassd.net:node=andromeda:port=port-3:link=*
	3863	urn:ogf:network:domain=es.net:node=atla-sdn1:port=xe-1/3/0:link=*
	3863	urn:ogf:network:domain=es.net:node=wash-sdn2:port=xe-1/3/0:link=*

Source VLAN	3863
Tagged	true
Destination VLAN	3863
Tagged	true

Error Report

Documentation | ESnet | Berkeley Lab | Notice to Users

Contacts: Chin Guok, Mary Thompson

Unicast Requests experience high blocking

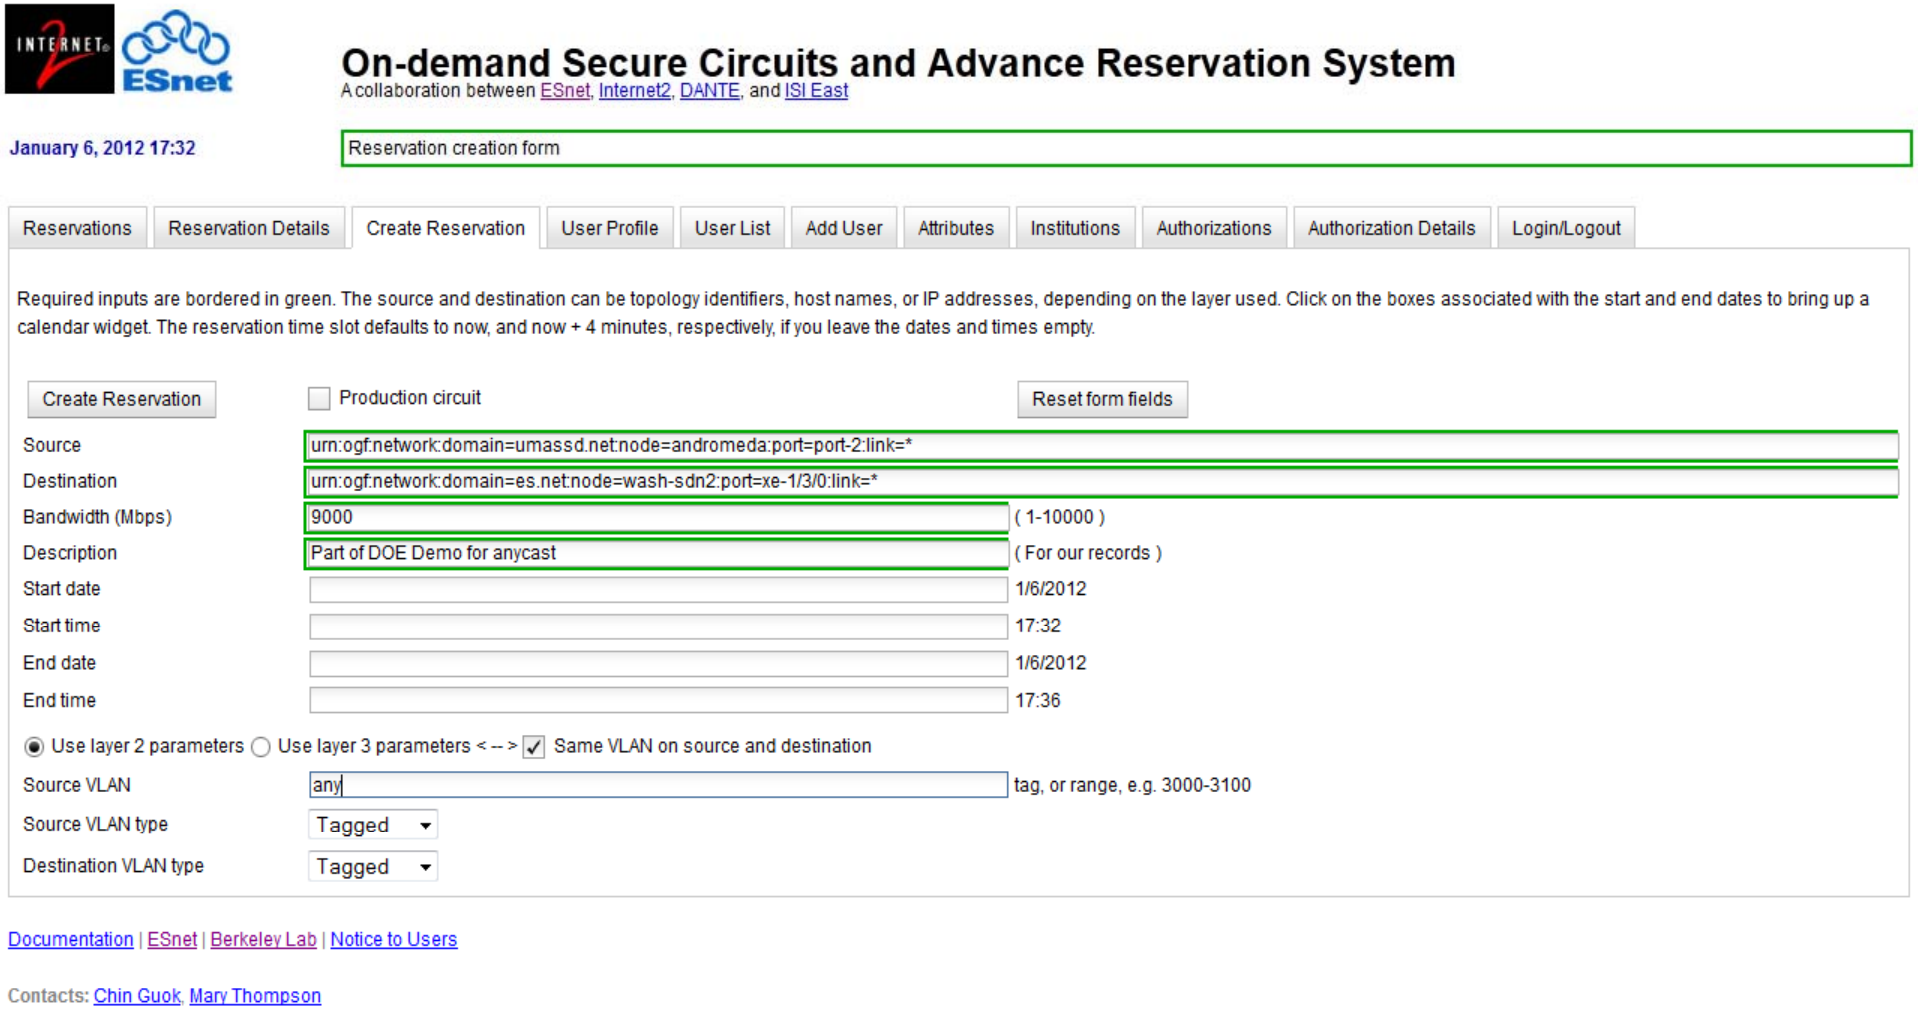

On-demand Secure Circuits and Advance Reservation System

A collaboration between ESnet, Internet2, DANTE, and ISI East

January 6, 2012 17:32

Reservation creation form

Reservations | Reservation Details | Create Reservation | User Profile | User List | Add User | Attributes | Institutions | Authorizations | Authorization Details | Login/Logout

Required inputs are bordered in green. The source and destination can be topology identifiers, host names, or IP addresses, depending on the layer used. Click on the boxes associated with the start and end dates to bring up a calendar widget. The reservation time slot defaults to now, and now + 4 minutes, respectively, if you leave the dates and times empty.

Create Reservation | Production circuit | Reset form fields

Source	urn:ogf:network:domain=umassd.net:node=andromeda:port=port-2:link=*	
Destination	urn:ogf:network:domain=es.net:node=wash-sdn2:port=xe-1/3/0:link=*	
Bandwidth (Mbps)	9000	(1-10000)
Description	Part of DOE Demo for anycast	(For our records)
Start date		1/6/2012
Start time		17:32
End date		1/6/2012
End time		17:36

Use layer 2 parameters | Use layer 3 parameters < -- > | Same VLAN on source and destination

Source VLAN	any	tag, or range, e.g. 3000-3100
Source VLAN type	Tagged	
Destination VLAN type	Tagged	

Documentation | ESnet | Berkeley Lab | Notice to Users

Contacts: Chin Guok, Mary Thompson

Unicast Requests experience high blocking

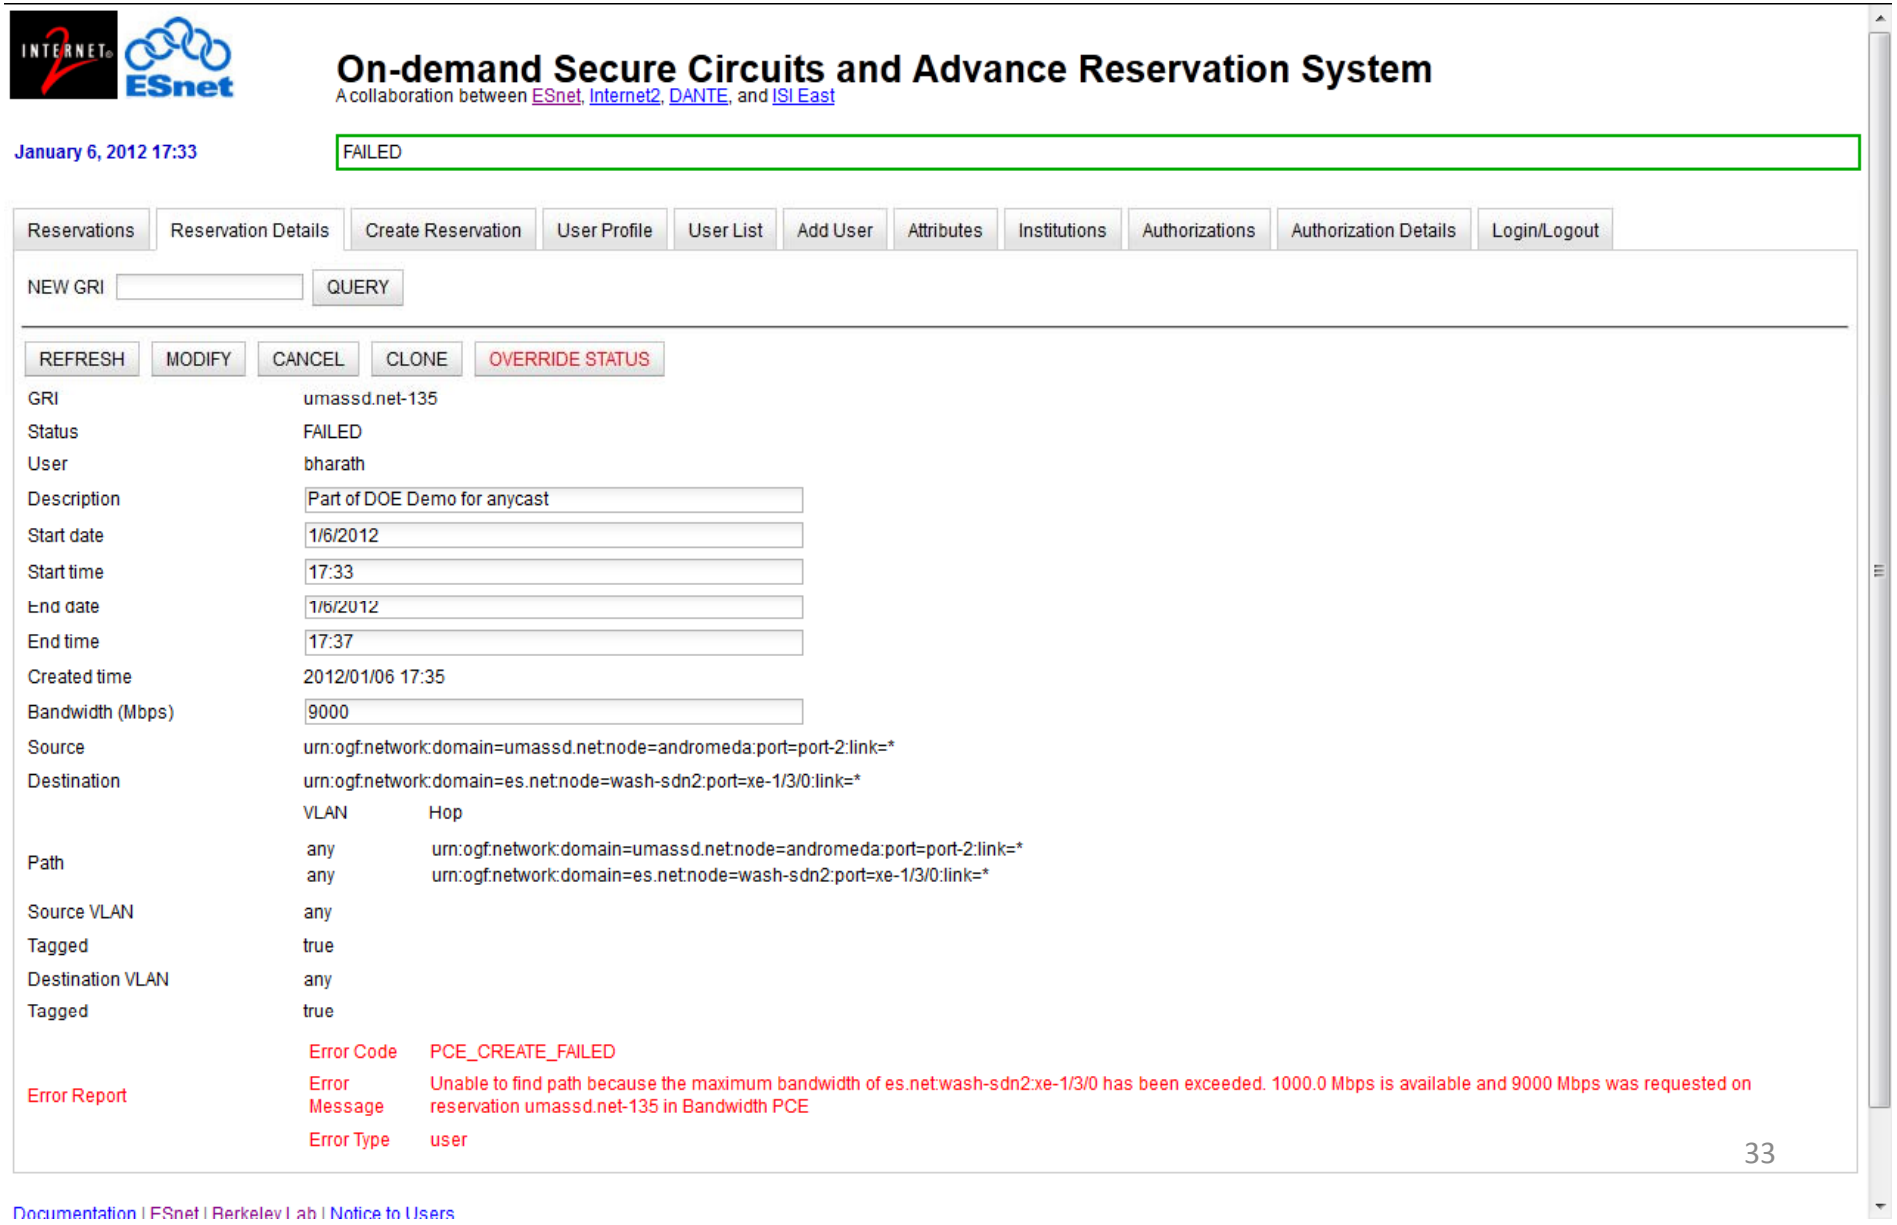

On-demand Secure Circuits and Advance Reservation System
A collaboration between ESnet, Internet2, DANTE, and ISI East

January 6, 2012 17:33 — FAILED

Reservations | Reservation Details | Create Reservation | User Profile | User List | Add User | Attributes | Institutions | Authorizations | Authorization Details | Login/Logout

NEW GRI [ ] QUERY

REFRESH | MODIFY | CANCEL | CLONE | OVERRIDE STATUS

GRI	umassd.net-135
Status	FAILED
User	bharath
Description	Part of DOE Demo for anycast
Start date	1/6/2012
Start time	17:33
End date	1/6/2012
End time	17:37
Created time	2012/01/06 17:35
Bandwidth (Mbps)	9000
Source	urn:ogf:network:domain=umassd.net:node=andromeda:port=port-2:link=*
Destination	urn:ogf:network:domain=es.net:node=wash-sdn2:port=xe-1/3/0:link=*

Path	VLAN	Hop
	any	urn:ogf:network:domain=umassd.net:node=andromeda:port=port-2:link=*
	any	urn:ogf:network:domain=es.net:node=wash-sdn2:port=xe-1/3/0:link=*

Source VLAN	any
Tagged	true
Destination VLAN	any
Tagged	true

Error Report		
	Error Code	PCE_CREATE_FAILED
	Error Message	Unable to find path because the maximum bandwidth of es.net:wash-sdn2:xe-1/3/0 has been exceeded. 1000.0 Mbps is available and 9000 Mbps was requested on reservation umassd.net-135 in Bandwidth PCE
	Error Type	user

Documentation | ESnet | Berkeley Lab | Notice to Users

Anycast Eliminates or Reduces blocking

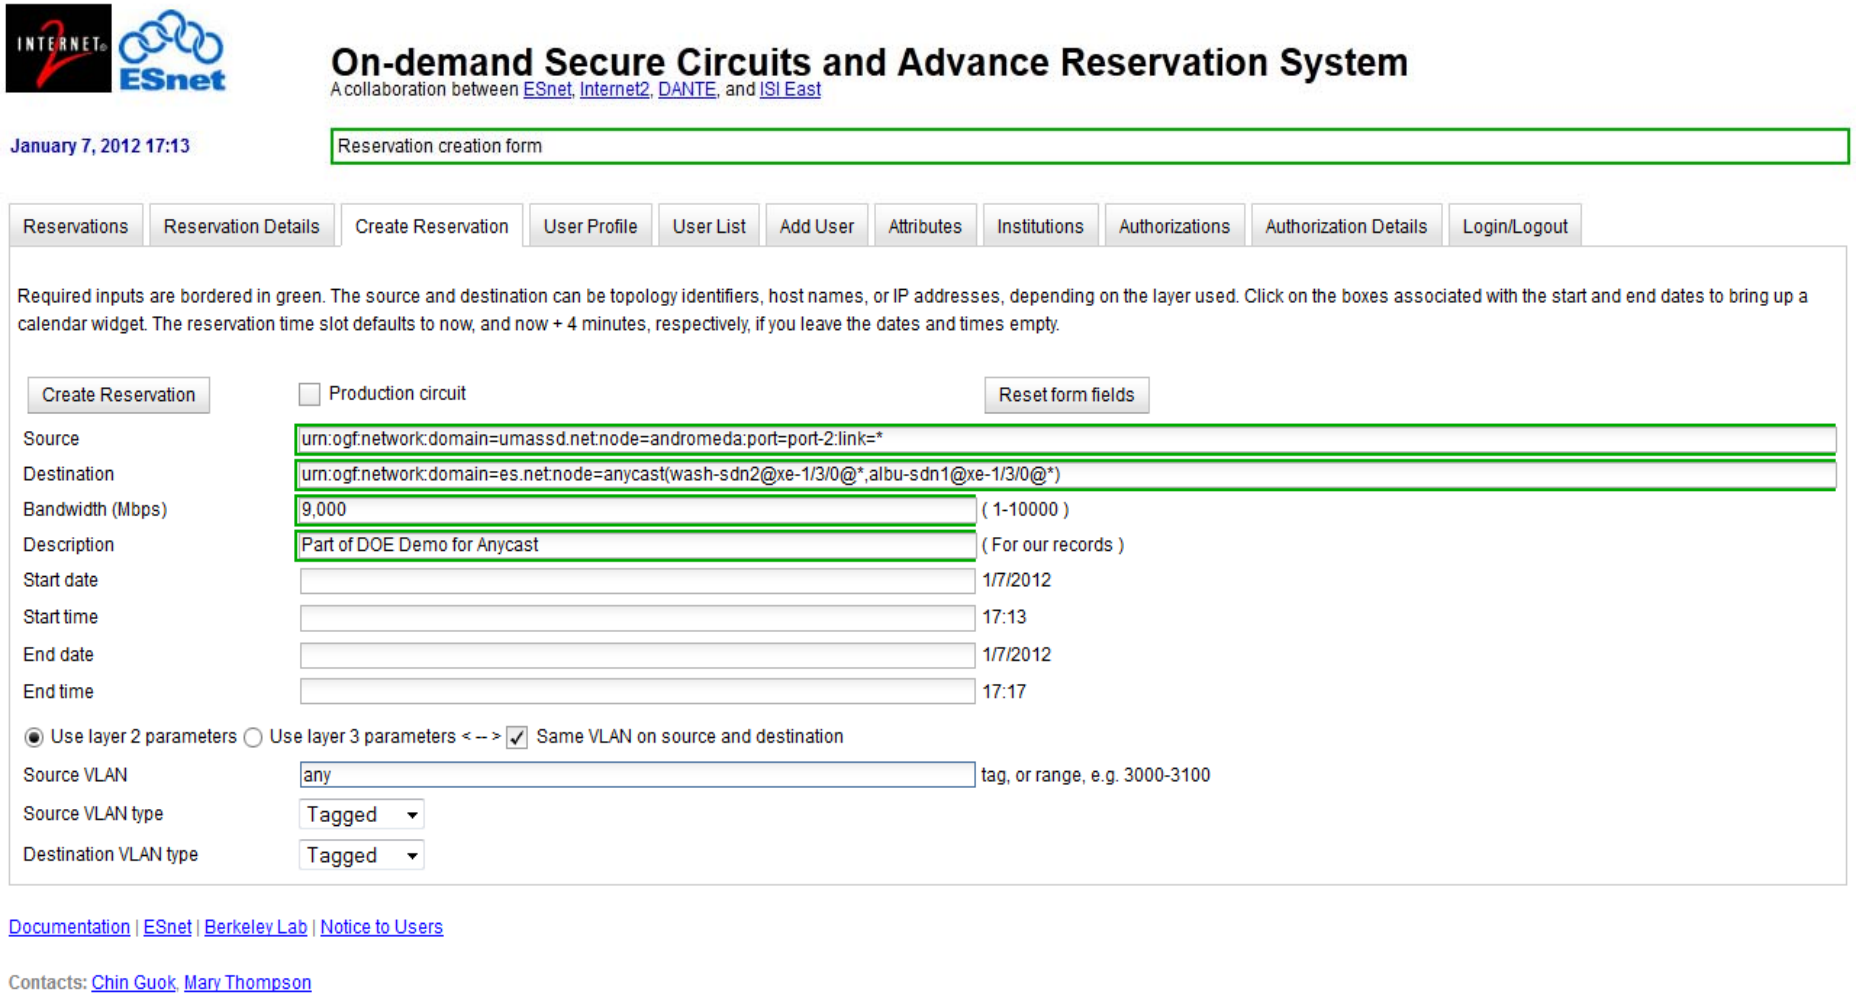

On-demand Secure Circuits and Advance Reservation System

A collaboration between ESnet, Internet2, DANTE, and ISI East

January 7, 2012 17:13

Reservation creation form

Reservations | Reservation Details | Create Reservation | User Profile | User List | Add User | Attributes | Institutions | Authorizations | Authorization Details | Login/Logout

Required inputs are bordered in green. The source and destination can be topology identifiers, host names, or IP addresses, depending on the layer used. Click on the boxes associated with the start and end dates to bring up a calendar widget. The reservation time slot defaults to now, and now + 4 minutes, respectively, if you leave the dates and times empty.

Create Reservation | Production circuit | Reset form fields

Source	urn:ogf:network:domain=umassd.net:node=andromeda:port=port-2:link=*	
Destination	urn:ogf:network:domain=es.net:node=anycast(wash-sdn2@xe-1/3/0@*,albu-sdn1@xe-1/3/0@*)	
Bandwidth (Mbps)	9,000	(1-10000)
Description	Part of DOE Demo for Anycast	(For our records)
Start date		1/7/2012
Start time		17:13
End date		1/7/2012
End time		17:17

Use layer 2 parameters | Use layer 3 parameters < -- > | Same VLAN on source and destination

Source VLAN	any	tag, or range, e.g. 3000-3100
Source VLAN type	Tagged	
Destination VLAN type	Tagged	

Documentation | ESnet | Berkeley Lab | Notice to Users

Contacts: Chin Guok, Mary Thompson

Anycast Eliminates or Reduces blocking

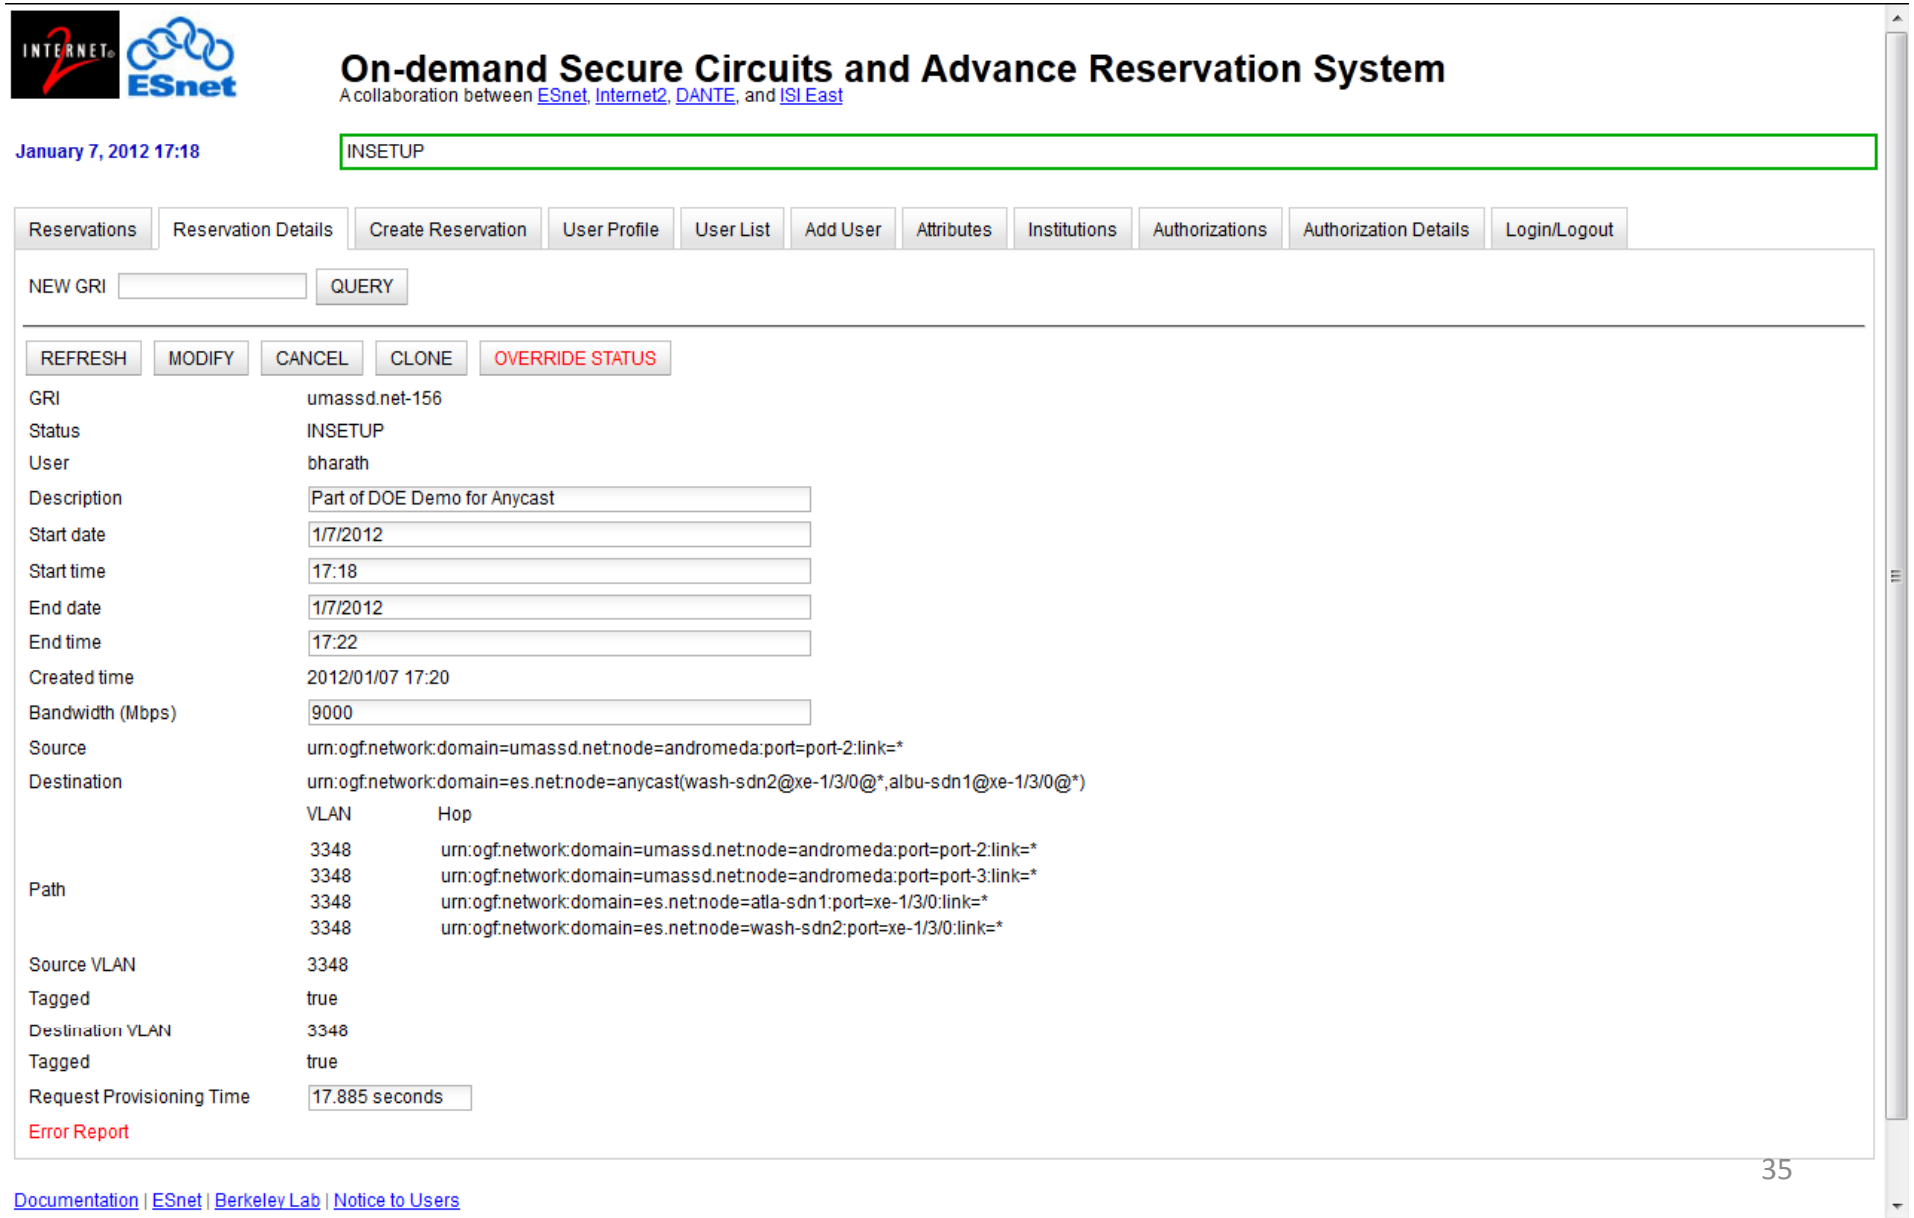

On-demand Secure Circuits and Advance Reservation System

A collaboration between ESnet, Internet2, DANTE, and ISI East

January 7, 2012 17:18 INSETUP

Reservations | Reservation Details | Create Reservation | User Profile | User List | Add User | Attributes | Institutions | Authorizations | Authorization Details | Login/Logout

NEW GRI [ ] QUERY

REFRESH | MODIFY | CANCEL | CLONE | OVERRIDE STATUS

GRI	umassd.net-156
Status	INSETUP
User	bharath
Description	Part of DOE Demo for Anycast
Start date	1/7/2012
Start time	17:18
End date	1/7/2012
End time	17:22
Created time	2012/01/07 17:20
Bandwidth (Mbps)	9000
Source	urn:ogf:network:domain=umassd.net:node=andromeda:port=port-2:link=*
Destination	urn:ogf:network:domain=es.net:node=anycast(wash-sdn2@xe-1/3/0@*,albu-sdn1@xe-1/3/0@*)

Path	VLAN	Hop
	3348	urn:ogf:network:domain=umassd.net:node=andromeda:port=port-2:link=*
	3348	urn:ogf:network:domain=umassd.net:node=andromeda:port=port-3:link=*
	3348	urn:ogf:network:domain=es.net:node=atla-sdn1:port=xe-1/3/0:link=*
	3348	urn:ogf:network:domain=es.net:node=wash-sdn2:port=xe-1/3/0:link=*

Source VLAN	3348
Tagged	true
Destination VLAN	3348
Tagged	true
Request Provisioning Time	17.885 seconds

Error Report

Documentation | ESnet | Berkeley Lab | Notice to Users

Anycast Eliminates or Reduces blocking

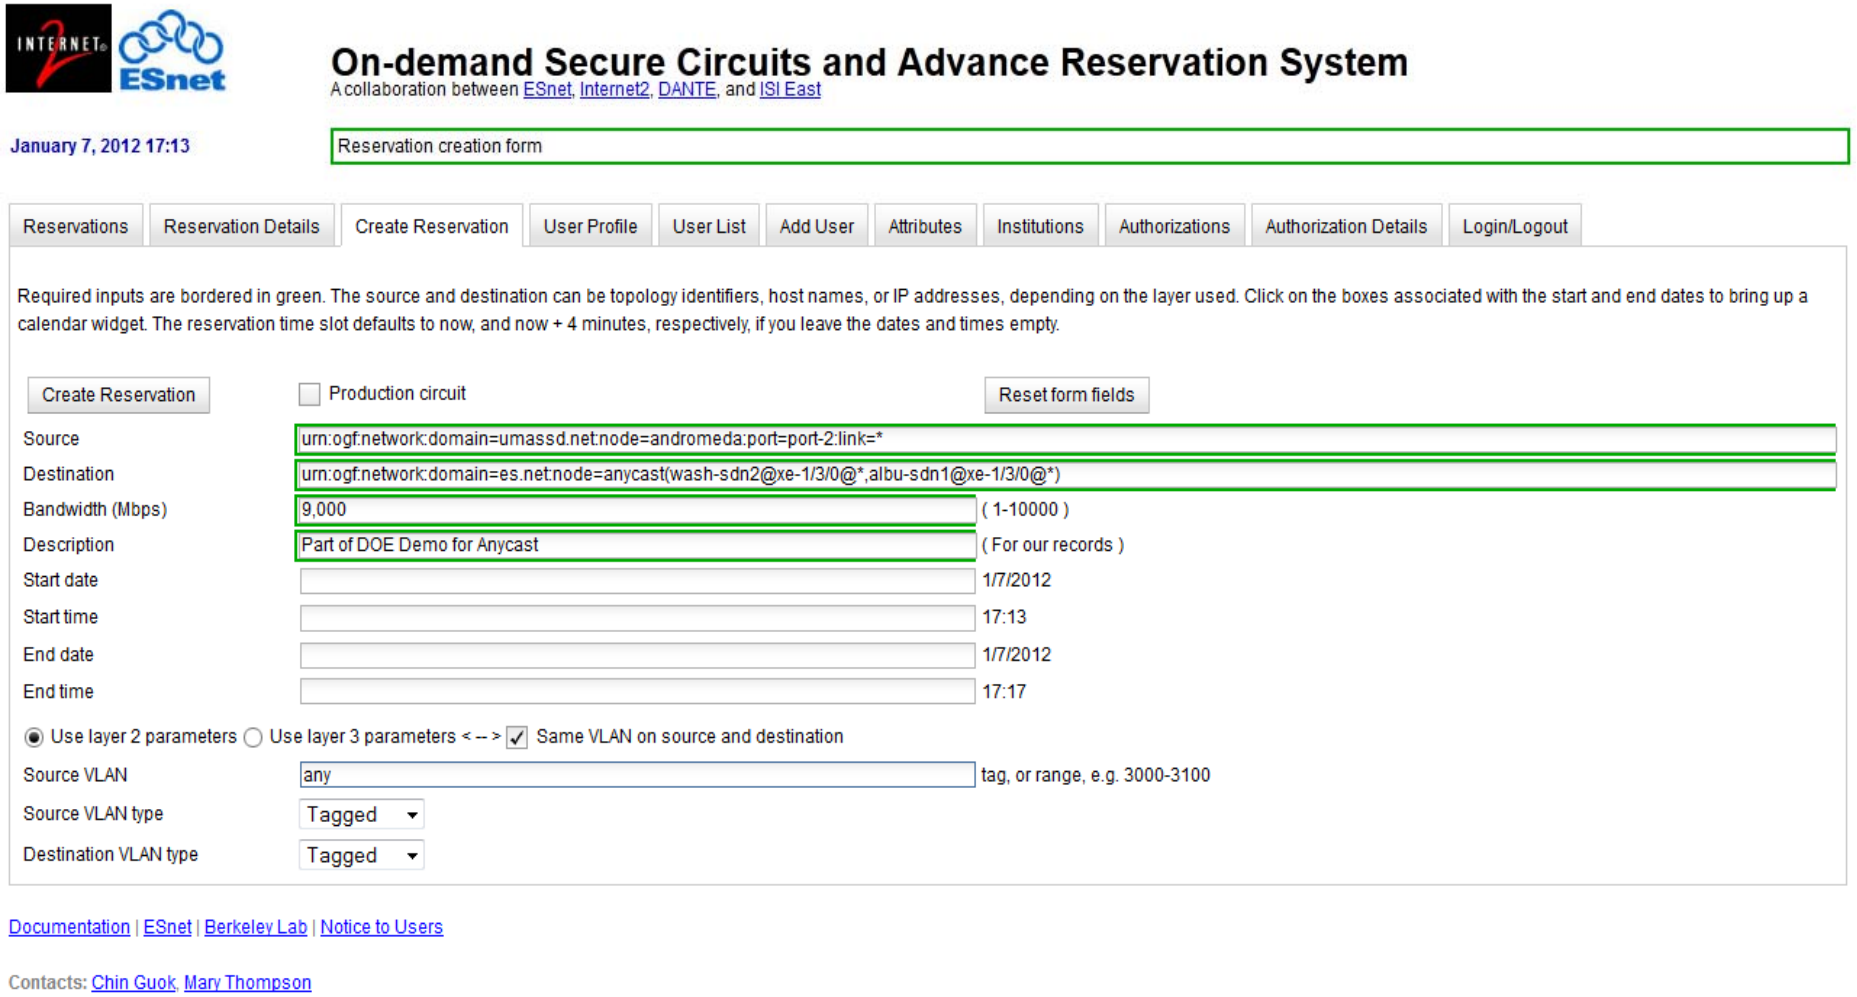

On-demand Secure Circuits and Advance Reservation System

A collaboration between ESnet, Internet2, DANTE, and ISI East

January 7, 2012 17:13

Reservation creation form

Reservations | Reservation Details | Create Reservation | User Profile | User List | Add User | Attributes | Institutions | Authorizations | Authorization Details | Login/Logout

Required inputs are bordered in green. The source and destination can be topology identifiers, host names, or IP addresses, depending on the layer used. Click on the boxes associated with the start and end dates to bring up a calendar widget. The reservation time slot defaults to now, and now + 4 minutes, respectively, if you leave the dates and times empty.

Create Reservation | Production circuit | Reset form fields

Source	urn:ogf:network:domain=umassd.net:node=andromeda:port=port-2:link=*	
Destination	urn:ogf:network:domain=es.net:node=anycast(wash-sdn2@xe-1/3/0@*,albu-sdn1@xe-1/3/0@*)	
Bandwidth (Mbps)	9,000	(1-10000)
Description	Part of DOE Demo for Anycast	(For our records)
Start date		1/7/2012
Start time		17:13
End date		1/7/2012
End time		17:17

Use layer 2 parameters | Use layer 3 parameters < -- > | Same VLAN on source and destination

Source VLAN	any	tag, or range, e.g. 3000-3100
Source VLAN type	Tagged	
Destination VLAN type	Tagged	

Documentation | ESnet | Berkeley Lab | Notice to Users

Contacts: Chin Guok, Mary Thompson

Anycast Eliminates or Reduces blocking

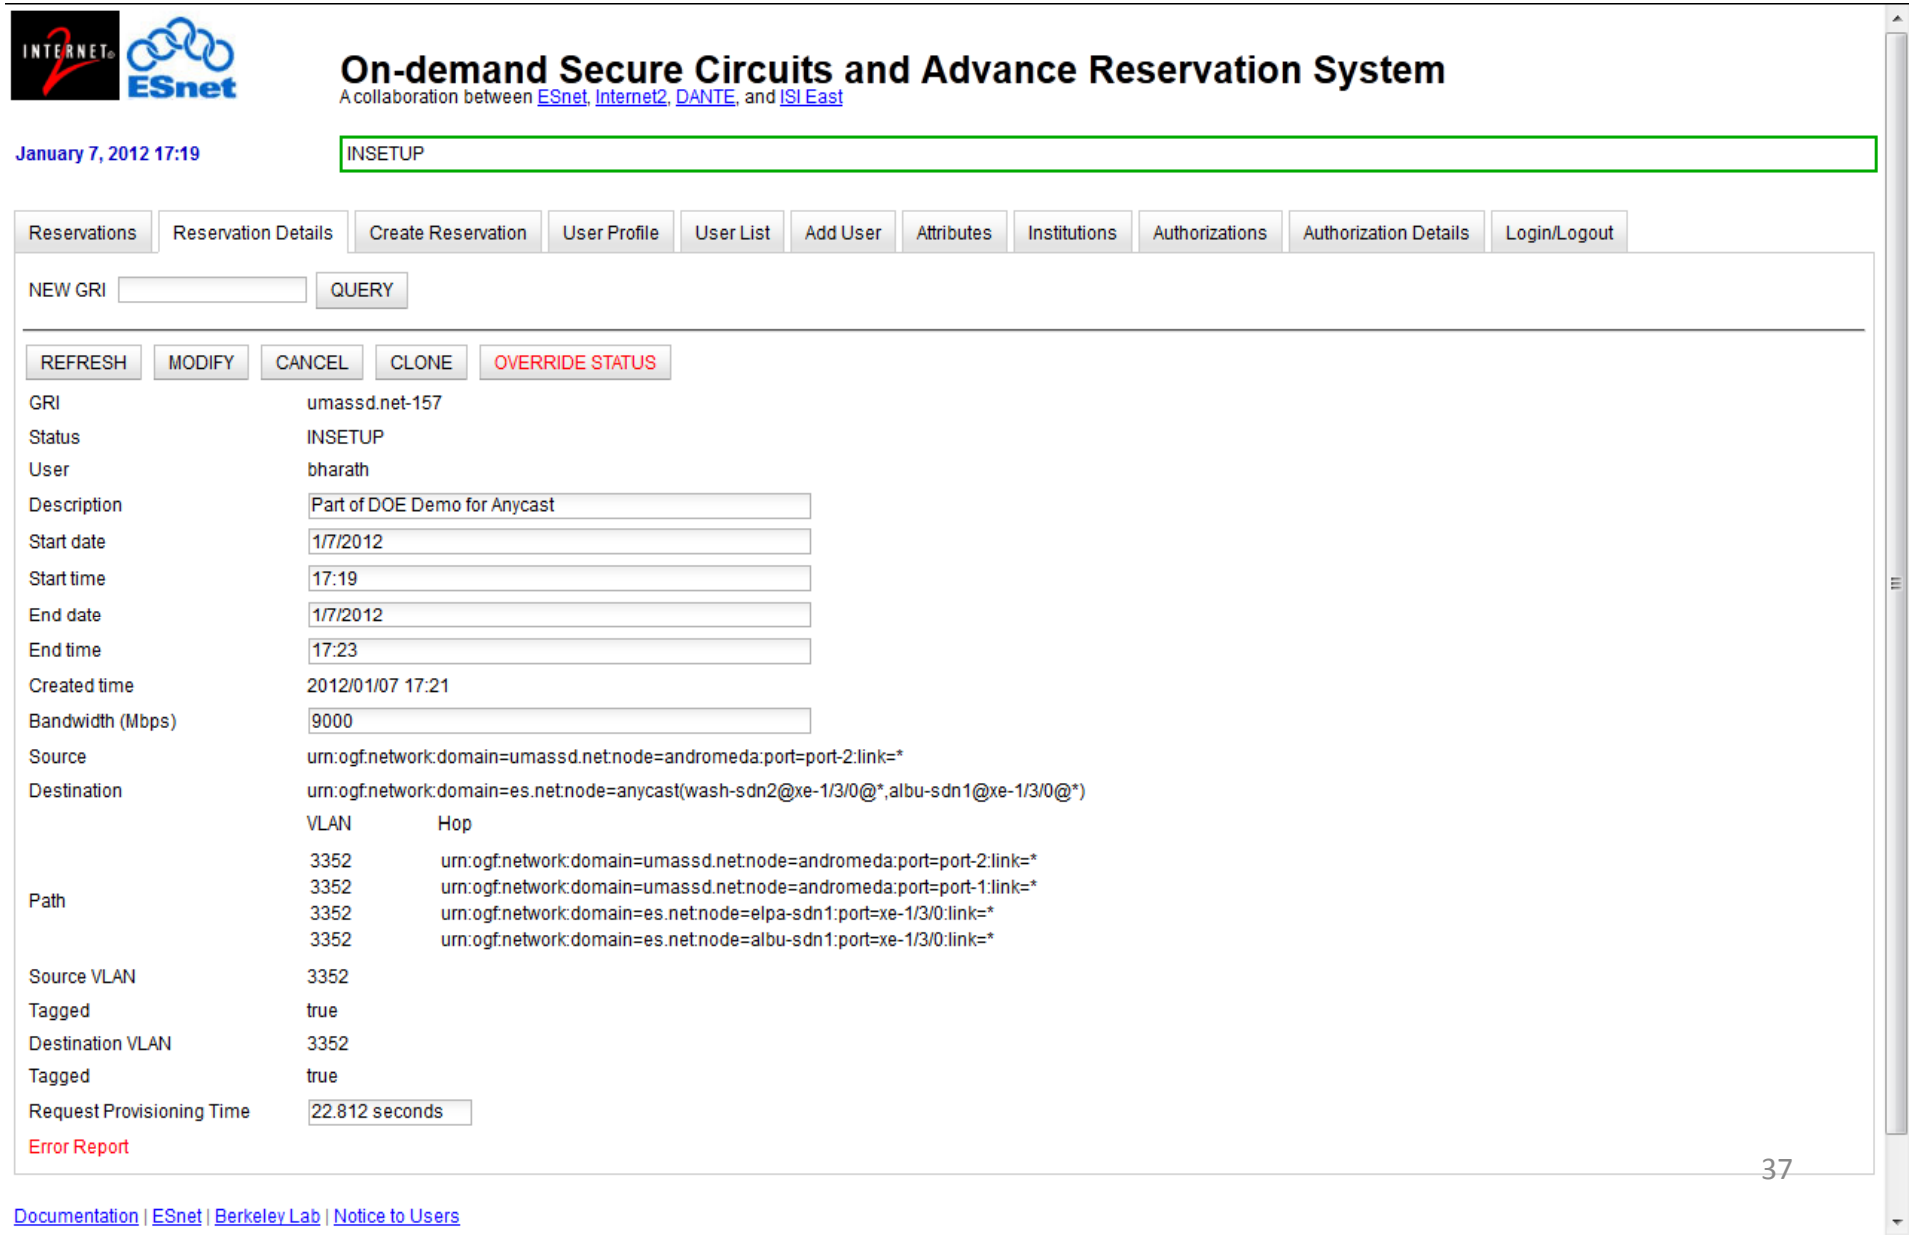

On-demand Secure Circuits and Advance Reservation System
A collaboration between ESnet, Internet2, DANTE, and ISI East

January 7, 2012 17:19 INSETUP

Reservations | Reservation Details | Create Reservation | User Profile | User List | Add User | Attributes | Institutions | Authorizations | Authorization Details | Login/Logout

NEW GRI [ ] QUERY

REFRESH | MODIFY | CANCEL | CLONE | OVERRIDE STATUS

GRI	umassd.net-157
Status	INSETUP
User	bharath
Description	Part of DOE Demo for Anycast
Start date	1/7/2012
Start time	17:19
End date	1/7/2012
End time	17:23
Created time	2012/01/07 17:21
Bandwidth (Mbps)	9000
Source	urn:ogf:network:domain=umassd.net:node=andromeda:port=port-2:link=*
Destination	urn:ogf:network:domain=es.net:node=anycast(wash-sdn2@xe-1/3/0@*,albu-sdn1@xe-1/3/0@*)

Path

VLAN	Hop
3352	urn:ogf:network:domain=umassd.net:node=andromeda:port=port-2:link=*
3352	urn:ogf:network:domain=umassd.net:node=andromeda:port=port-1:link=*
3352	urn:ogf:network:domain=es.net:node=elpa-sdn1:port=xe-1/3/0:link=*
3352	urn:ogf:network:domain=es.net:node=albu-sdn1:port=xe-1/3/0:link=*

Source VLAN	3352
Tagged	true
Destination VLAN	3352
Tagged	true
Request Provisioning Time	22.812 seconds

Error Report

Documentation | ESnet | Berkeley Lab | Notice to Users

Scenario 2 – Average Hop Count

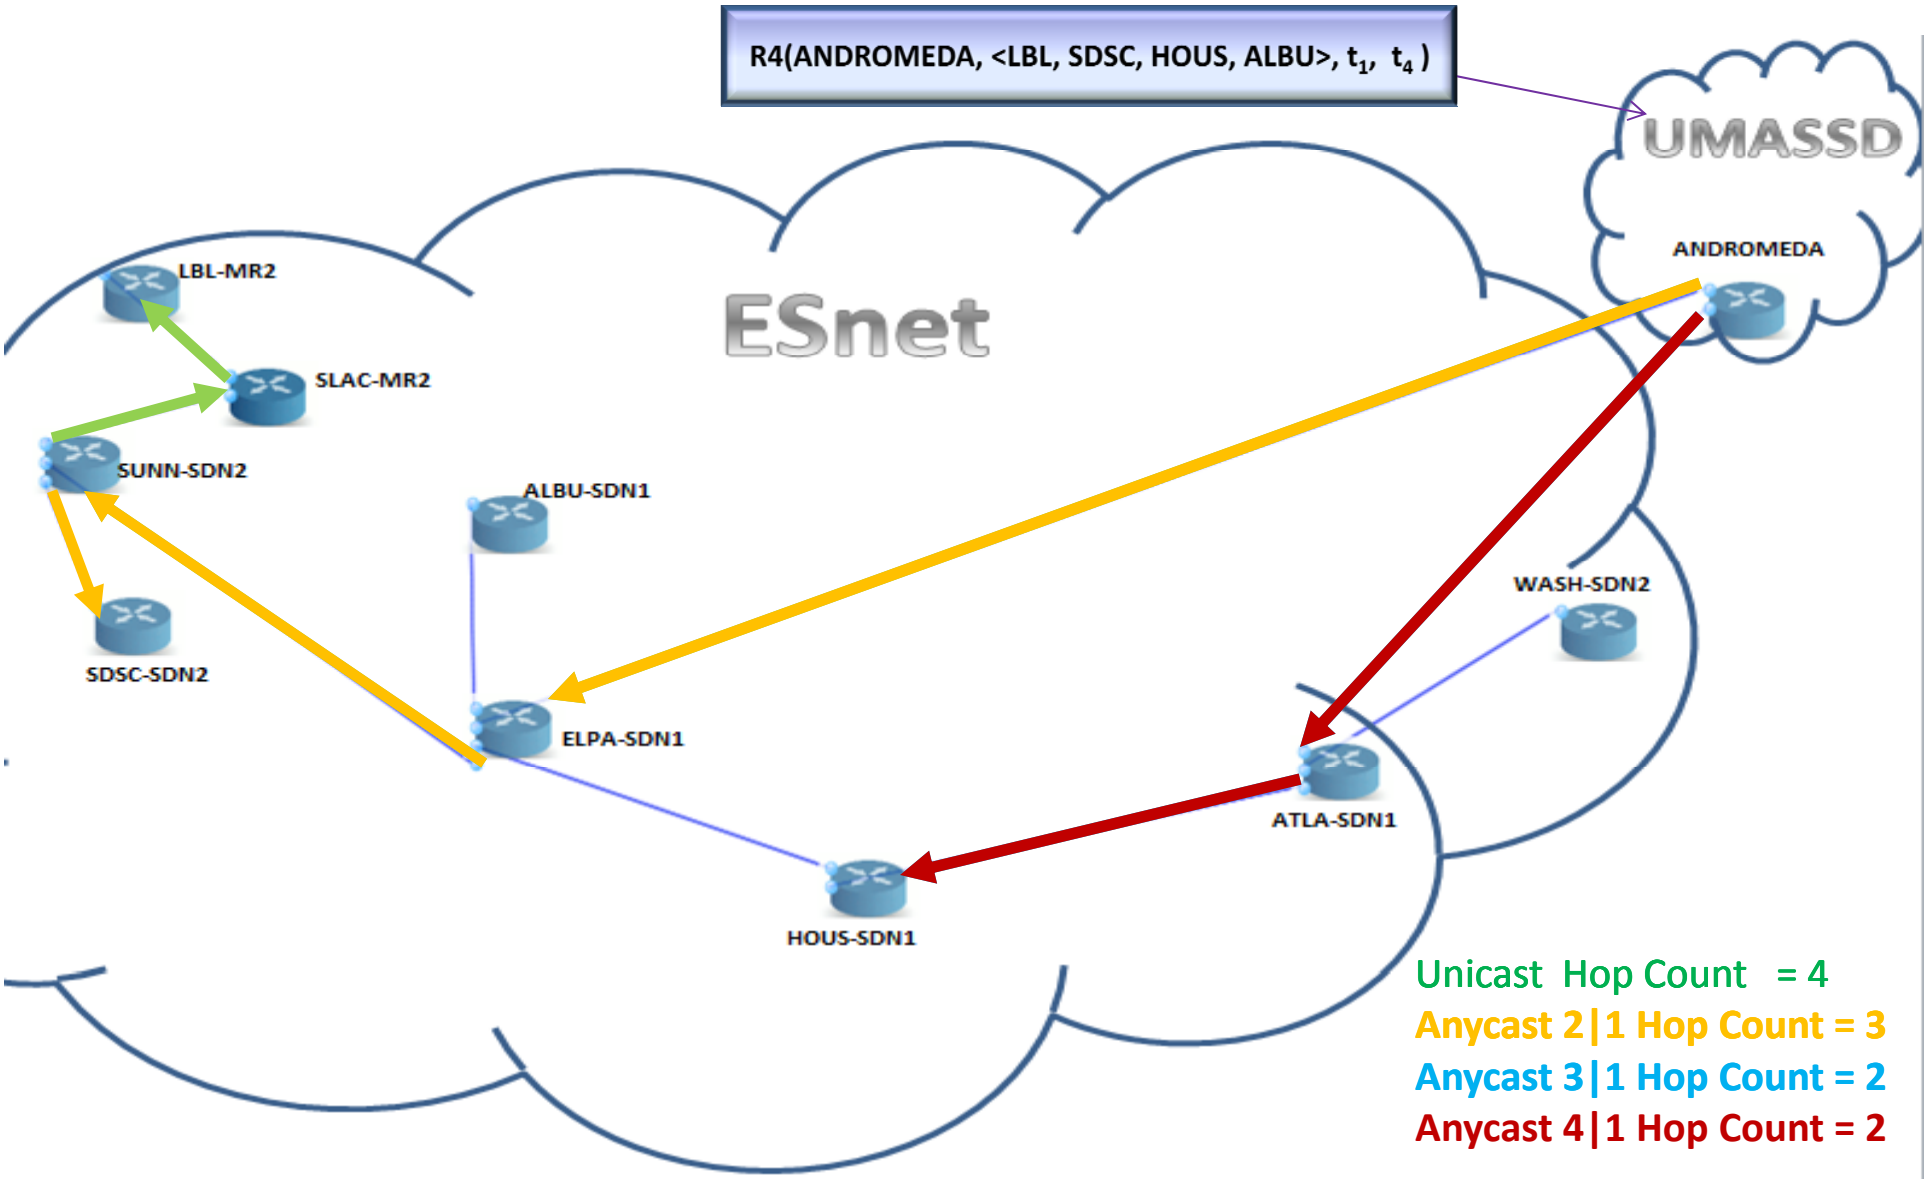

Hop Count Measure for Unicast

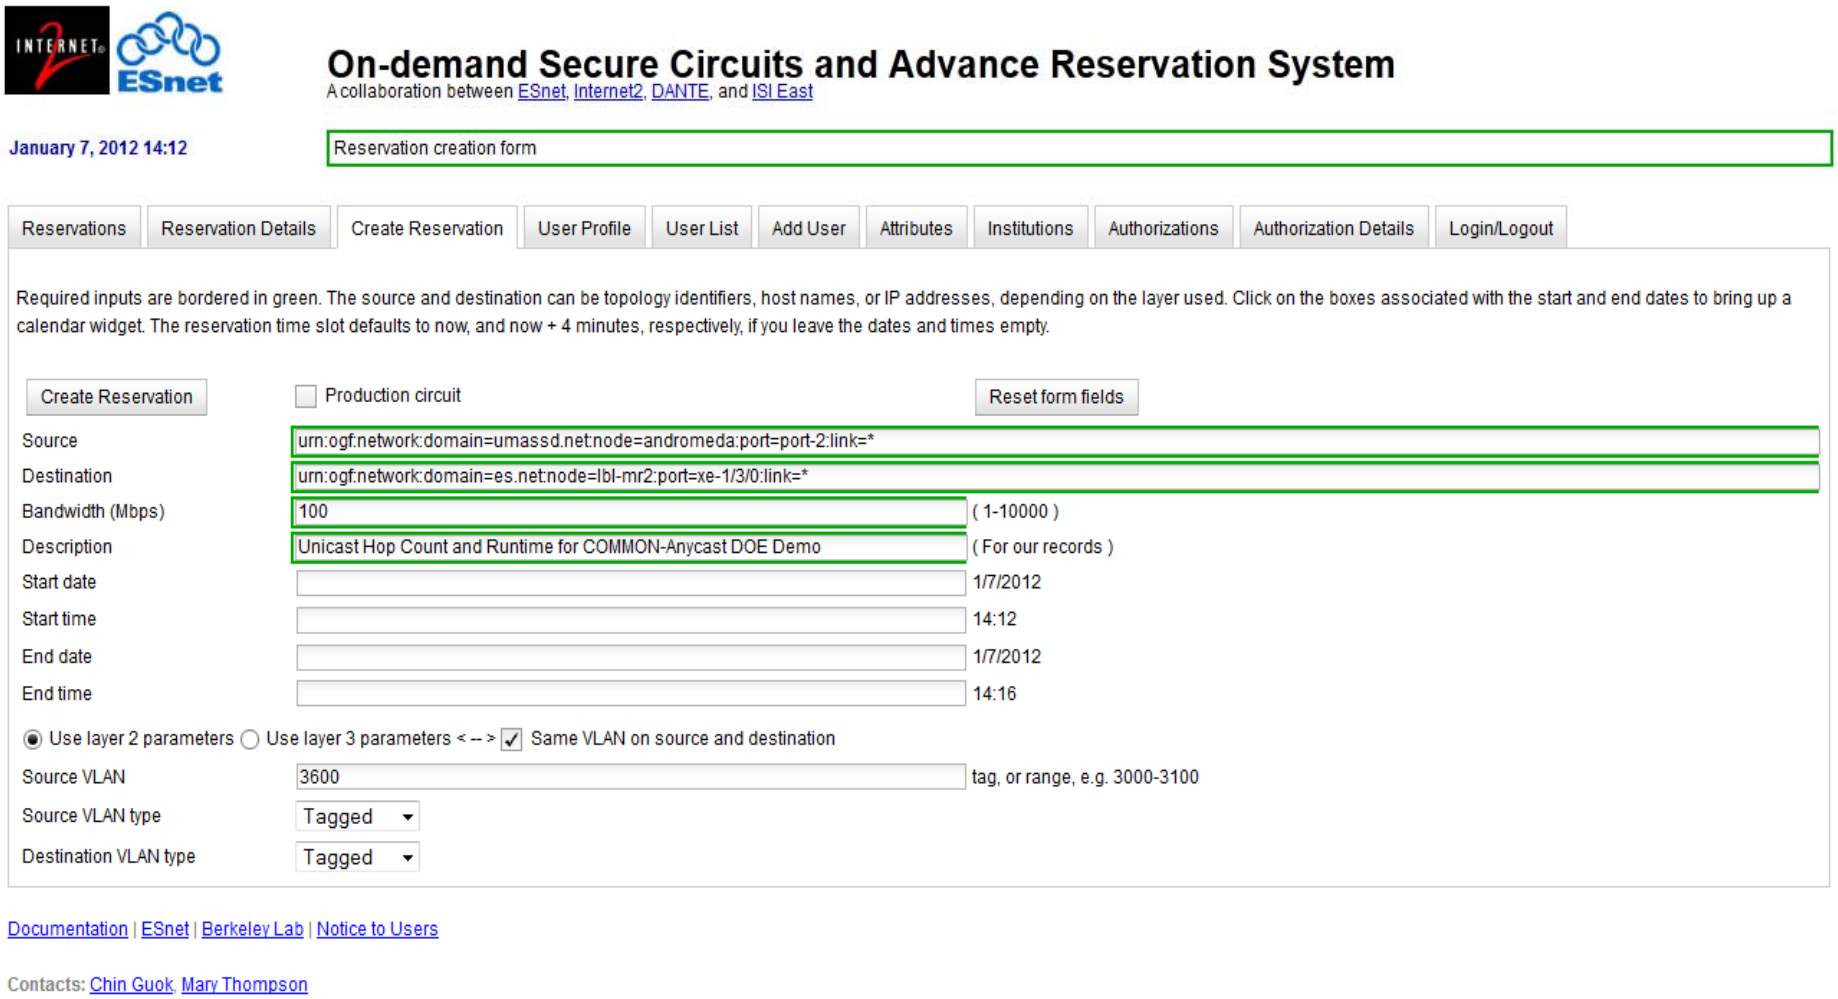

On-demand Secure Circuits and Advance Reservation System

A collaboration between ESnet, Internet2, DANTE, and ISI East

January 7, 2012 14:12

Reservation creation form

Reservations | Reservation Details | Create Reservation | User Profile | User List | Add User | Attributes | Institutions | Authorizations | Authorization Details | Login/Logout

Required inputs are bordered in green. The source and destination can be topology identifiers, host names, or IP addresses, depending on the layer used. Click on the boxes associated with the start and end dates to bring up a calendar widget. The reservation time slot defaults to now, and now + 4 minutes, respectively, if you leave the dates and times empty.

Create Reservation | Production circuit | Reset form fields

Source	urn:ogf:network:domain=umassd.net:node=andromeda:port=port-2:link=*	
Destination	urn:ogf:network:domain=es.net:node=lbl-mr2:port=xe-1/3/0:link=*	
Bandwidth (Mbps)	100	(1-10000)
Description	Unicast Hop Count and Runtime for COMMON-Anycast DOE Demo	(For our records)
Start date		1/7/2012
Start time		14:12
End date		1/7/2012
End time		14:16

Use layer 2 parameters | Use layer 3 parameters < -- > | Same VLAN on source and destination

Source VLAN	3600	tag, or range, e.g. 3000-3100
Source VLAN type	Tagged	
Destination VLAN type	Tagged	

Documentation | ESnet | Berkeley Lab | Notice to Users

Contacts: Chin Guok, Mary Thompson

Hop Count Measure for Unicast

Reservations | Reservation Details | Create Reservation | Login/Logout

NEW GRI [ ] QUERY

REFRESH | MODIFY | CANCEL | CLONE | CREATE PATH | TEAR DOWN PATH | OVERRIDE STATUS

GRI	umassd.net-151
Status	FINISHED
User	umassd.net
Description	Unicast Hop Count and Runtime for COMMON-Anycast DOE Demo
Start date	1/7/2012
Start time	14:13
End date	1/7/2012
End time	14:17
Created time	2012/01/07 14:13
Bandwidth (Mbps)	100
Source	urn:ogf:network:domain=umassd.net:node=andromeda:port=port-2:link=*
Destination	urn:ogf:network:domain=es.net:node=lbl-mr2:port=xe-1/3/0:link=*

	VLAN	Hop
Local path	3600	urn:ogf:network:domain=es.net:node=elpa-sdn1:port=xe-1/3/0:link=*
	0,2-4094	urn:ogf:network:domain=es.net:node=elpa-sdn1:port=xe-0/1/0:link=xe-0/1/0.0
	0,2-4094	urn:ogf:network:domain=es.net:node=sunn-sdn2:port=xe-3/2/0:link=xe-3/2/0.0
	0,2-4094	urn:ogf:network:domain=es.net:node=sunn-sdn2:port=xe-3/0/0:link=xe-3/0/0.0
	0,2-4094	urn:ogf:network:domain=es.net:node=slac-mr2:port=xe-1/1/0:link=xe-1/1/0.0
	0,2-4094	urn:ogf:network:domain=es.net:node=slac-mr2:port=xe-0/0/0:link=xe-0/0/0.0
	0,2-4094	urn:ogf:network:domain=es.net:node=lbl-mr2:port=xe-0/0/0:link=xe-0/0/0.0
	3600	urn:ogf:network:domain=es.net:node=lbl-mr2:port=xe-1/3/0:link=*
Interdomain path	3600	urn:ogf:network:domain=umassd.net:node=andromeda:port=port-2:link=*
	3600	urn:ogf:network:domain=umassd.net:node=andromeda:port=port-1:link=*
	3600	urn:ogf:network:domain=es.net:node=elpa-sdn1:port=xe-1/3/0:link=*
	3600	urn:ogf:network:domain=es.net:node=lbl-mr2:port=xe-1/3/0:link=*

Source VLAN	3600
Tagged	true
Destination VLAN	3600
Tagged	true

Documentation | ESnet | Berkeley Lab | Notice to Users

Reduced Hop Counts for Anycast 2|1 compared to Unicast

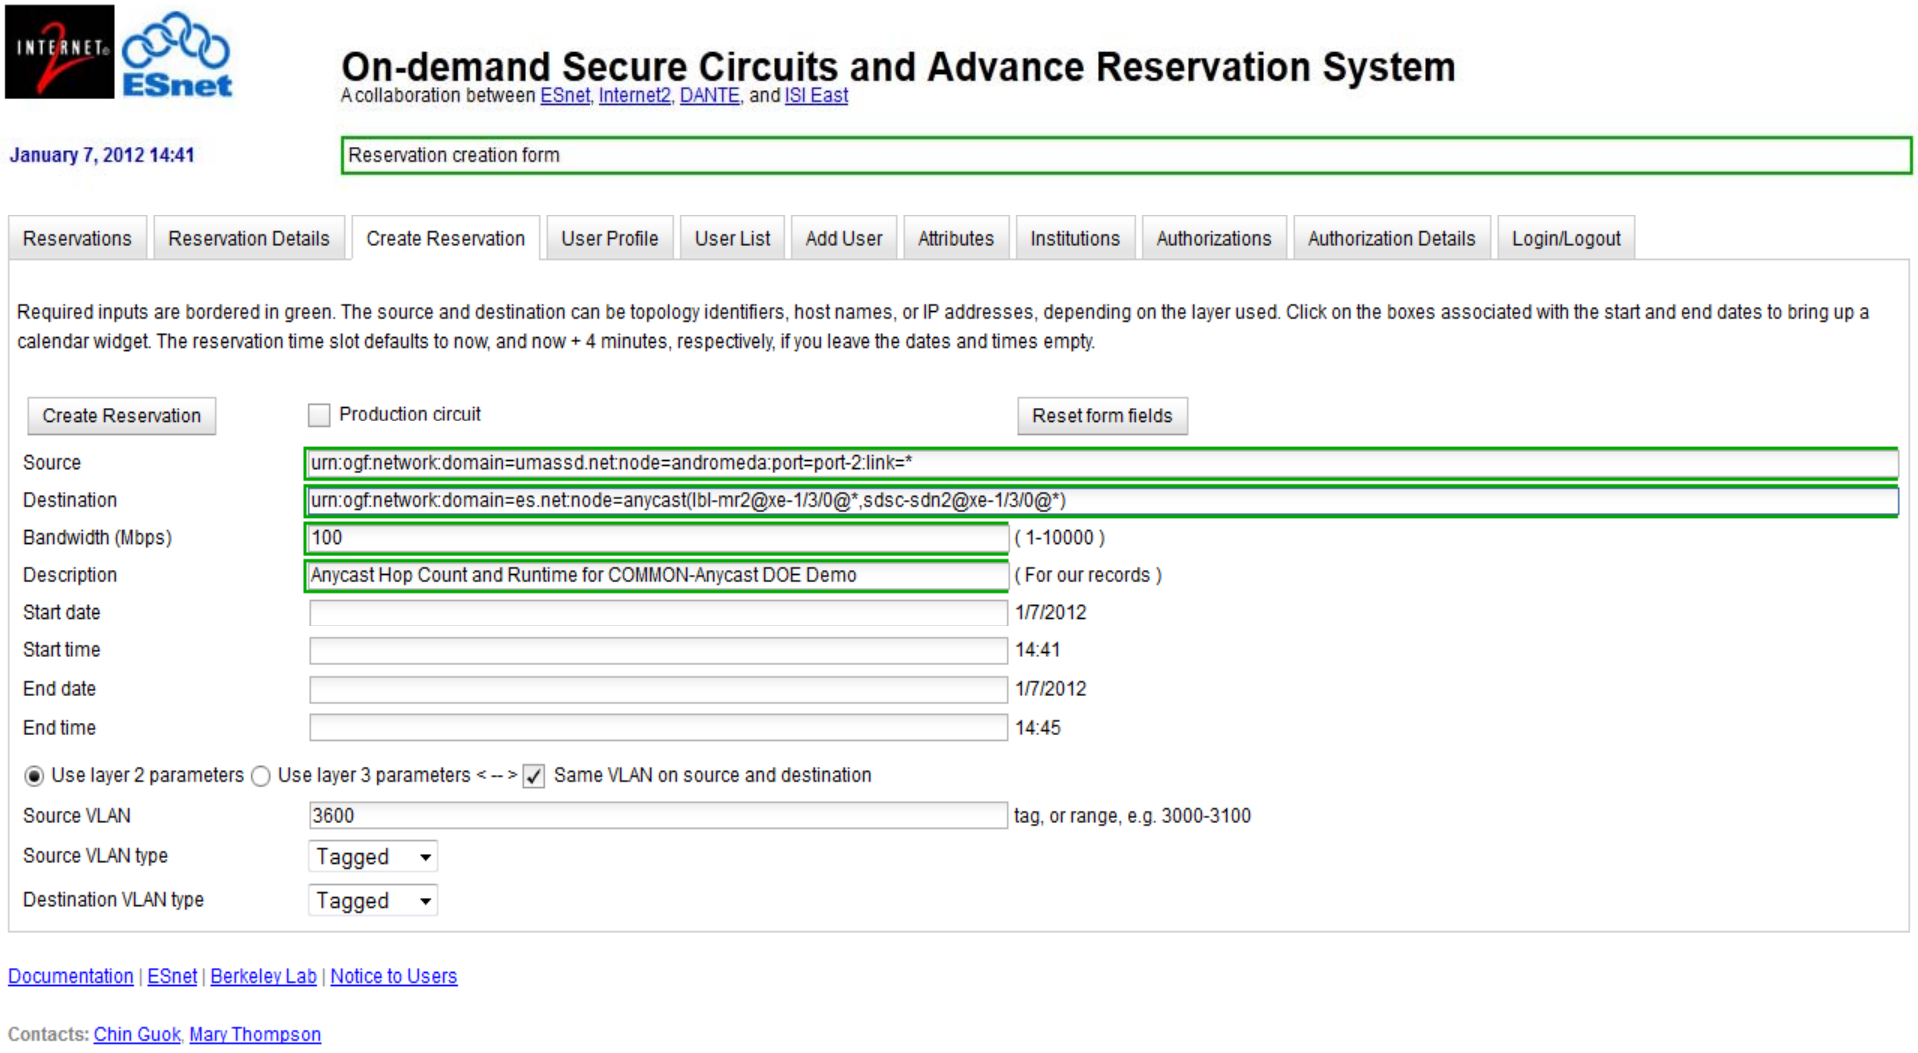

Reduced Hop Counts for Anycast 2|1 compared to Unicast

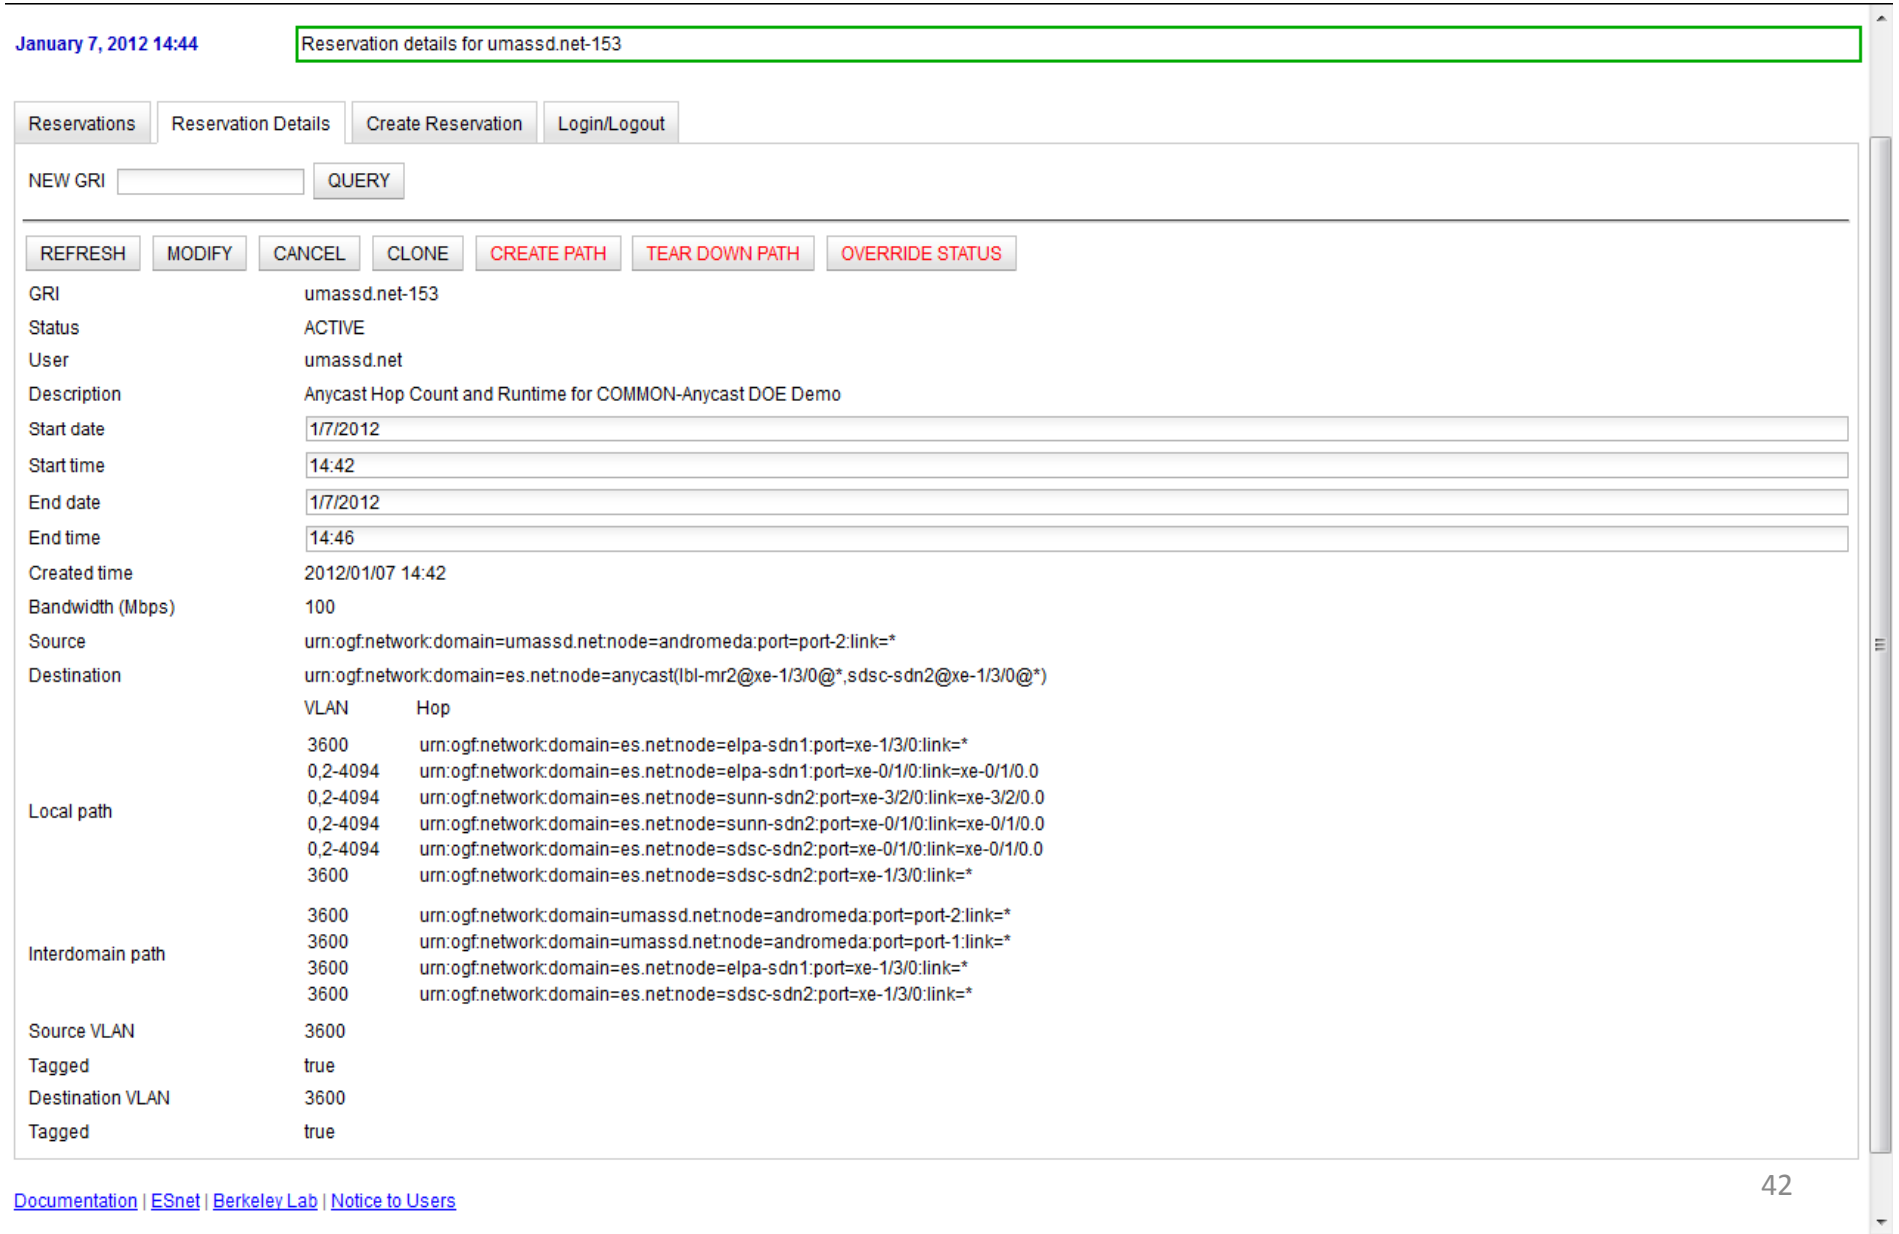

January 7, 2012 14:44

Reservation details for umassd.net-153

Reservations | Reservation Details | Create Reservation | Login/Logout

NEW GRI [ ] QUERY

REFRESH | MODIFY | CANCEL | CLONE | CREATE PATH | TEAR DOWN PATH | OVERRIDE STATUS

GRI	umassd.net-153
Status	ACTIVE
User	umassd.net
Description	Anycast Hop Count and Runtime for COMMON-Anycast DOE Demo
Start date	1/7/2012
Start time	14:42
End date	1/7/2012
End time	14:46
Created time	2012/01/07 14:42
Bandwidth (Mbps)	100
Source	urn:ogf:network:domain=umassd.net:node=andromeda:port=port-2:link=*
Destination	urn:ogf:network:domain=es.net:node=anycast(lbl-mr2@xe-1/3/0@*,sdsc-sdn2@xe-1/3/0@*)

	VLAN	Hop
Local path	3600	urn:ogf:network:domain=es.net:node=elpa-sdn1:port=xe-1/3/0:link=*
	0,2-4094	urn:ogf:network:domain=es.net:node=elpa-sdn1:port=xe-0/1/0:link=xe-0/1/0.0
	0,2-4094	urn:ogf:network:domain=es.net:node=sunn-sdn2:port=xe-3/2/0:link=xe-3/2/0.0
	0,2-4094	urn:ogf:network:domain=es.net:node=sunn-sdn2:port=xe-0/1/0:link=xe-0/1/0.0
	0,2-4094	urn:ogf:network:domain=es.net:node=sdsc-sdn2:port=xe-0/1/0:link=xe-0/1/0.0
	3600	urn:ogf:network:domain=es.net:node=sdsc-sdn2:port=xe-1/3/0:link=*
Interdomain path	3600	urn:ogf:network:domain=umassd.net:node=andromeda:port=port-2:link=*
	3600	urn:ogf:network:domain=umassd.net:node=andromeda:port=port-1:link=*
	3600	urn:ogf:network:domain=es.net:node=elpa-sdn1:port=xe-1/3/0:link=*
	3600	urn:ogf:network:domain=es.net:node=sdsc-sdn2:port=xe-1/3/0:link=*

Source VLAN	3600
Tagged	true
Destination VLAN	3600
Tagged	true

Documentation | ESnet | Berkeley Lab | Notice to Users

Reduced Hop Counts for Anycast 3|1 compared to Anycast 2|1

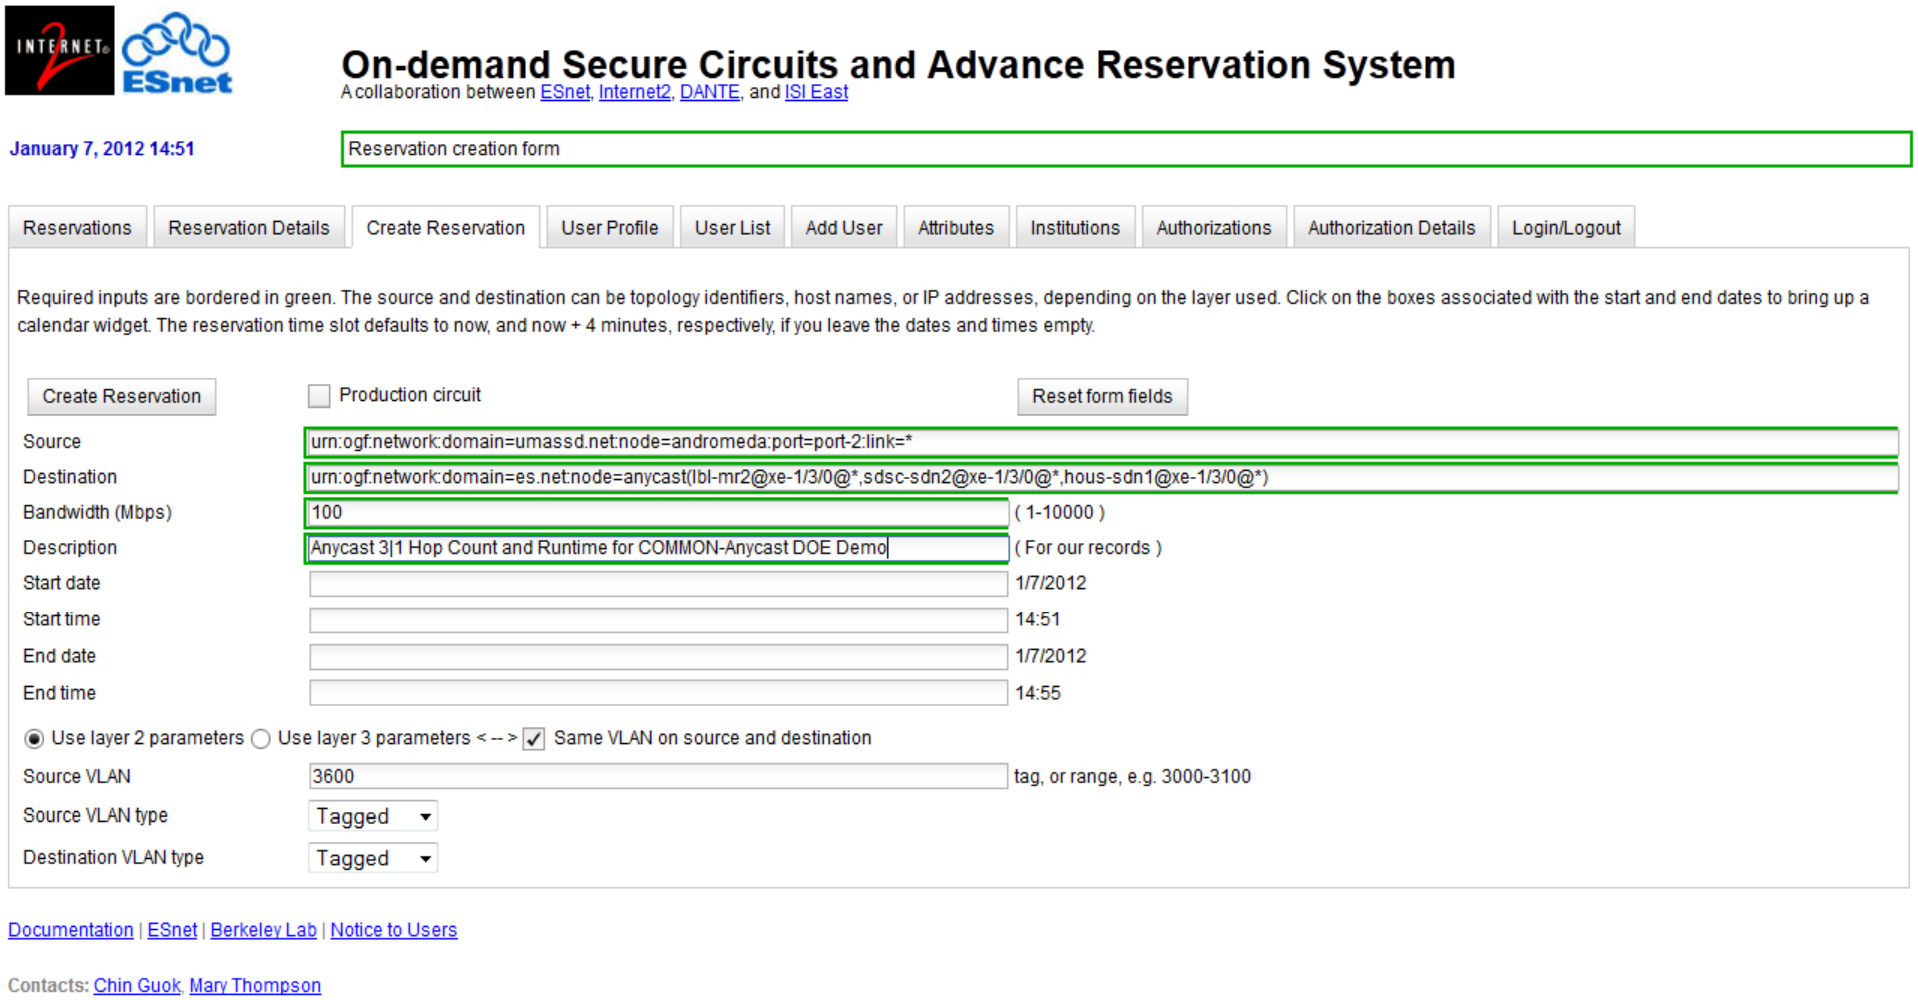

Reduced Hop Counts for Anycast 3|1 compared to Anycast 2|1

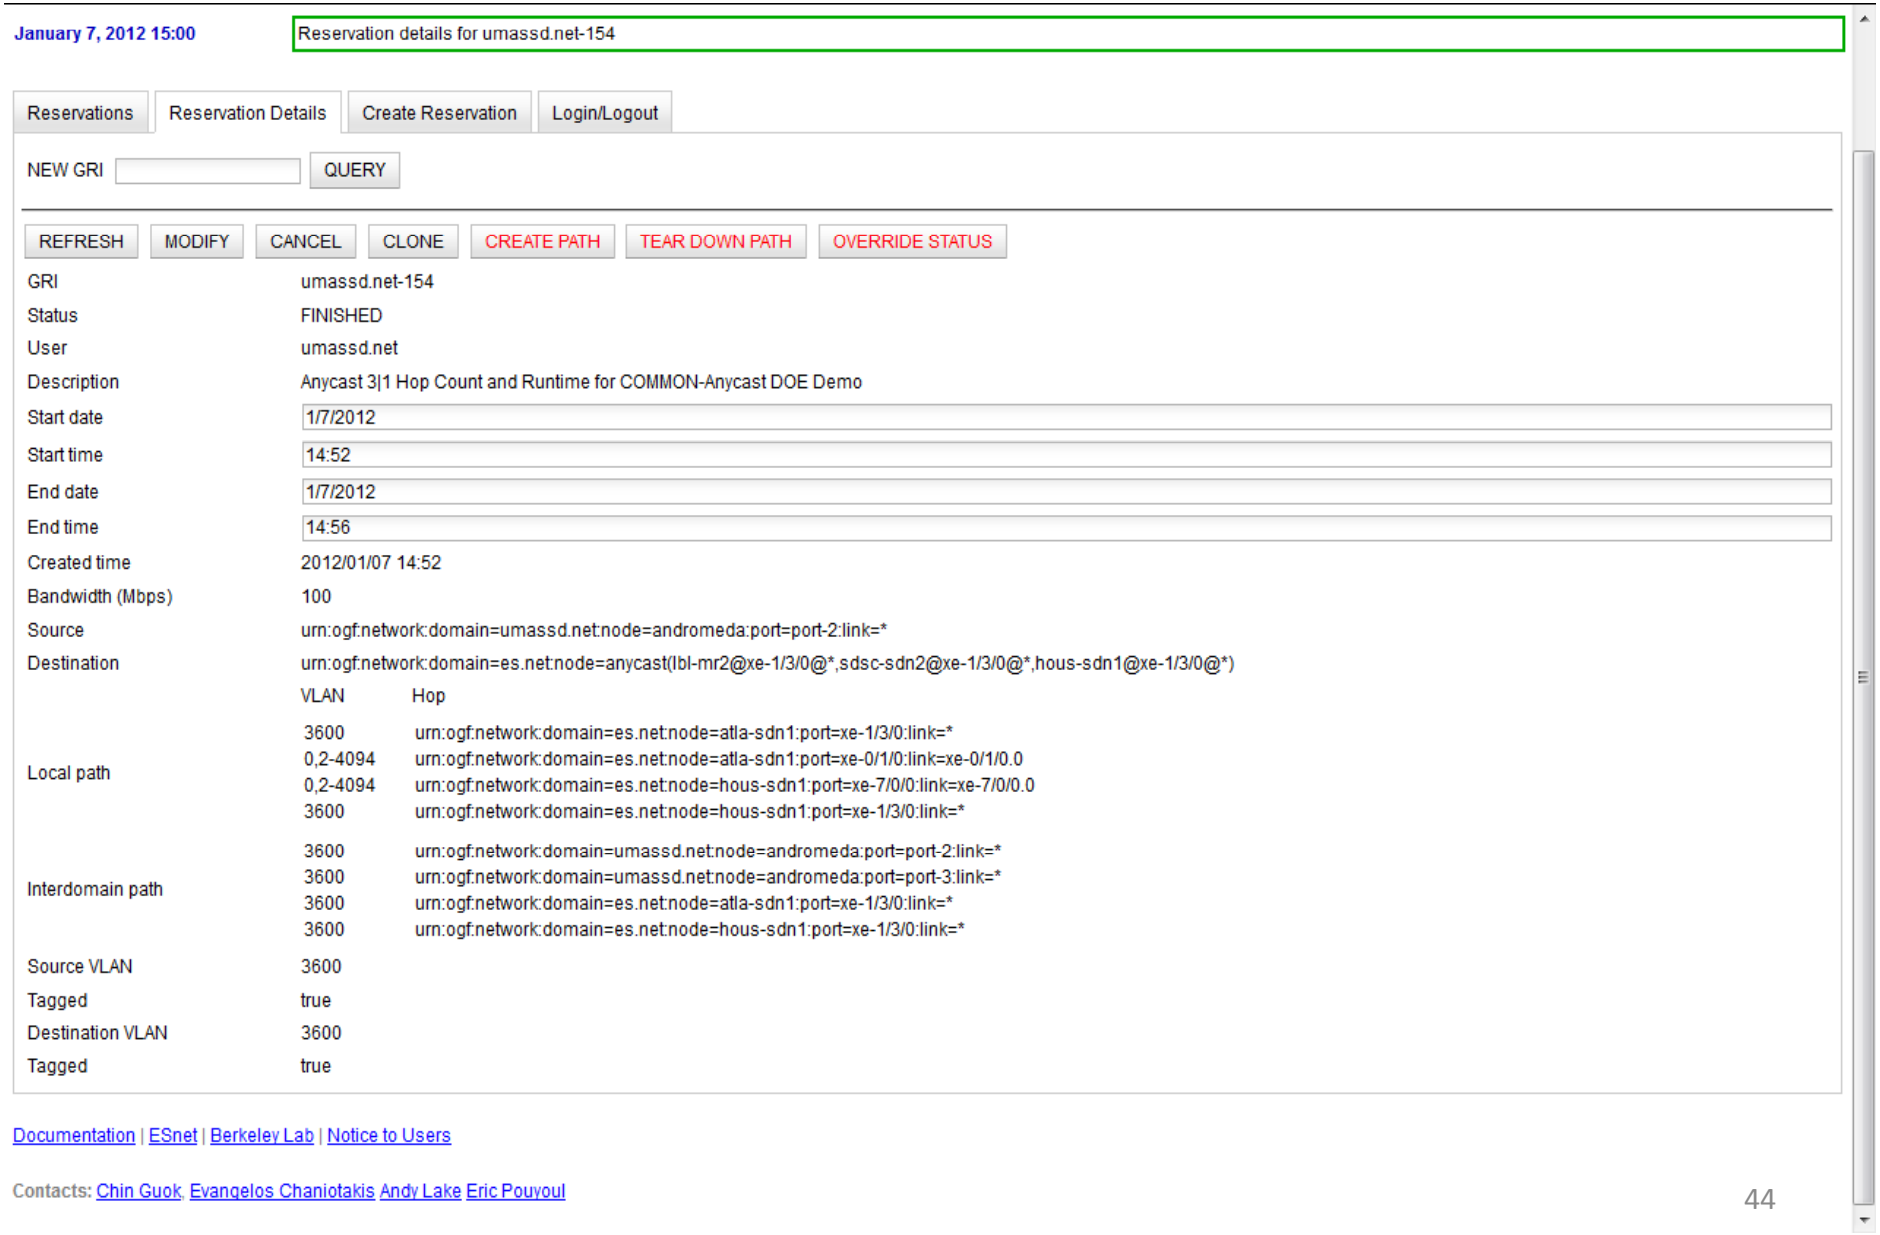

January 7, 2012 15:00

Reservation details for umassd.net-154

Reservations | Reservation Details | Create Reservation | Login/Logout

NEW GRI [ ] QUERY

REFRESH | MODIFY | CANCEL | CLONE | CREATE PATH | TEAR DOWN PATH | OVERRIDE STATUS

GRI	umassd.net-154
Status	FINISHED
User	umassd.net
Description	Anycast 3|1 Hop Count and Runtime for COMMON-Anycast DOE Demo
Start date	1/7/2012
Start time	14:52
End date	1/7/2012
End time	14:56
Created time	2012/01/07 14:52
Bandwidth (Mbps)	100
Source	urn:ogf:network:domain=umassd.net:node=andromeda:port=port-2:link=*
Destination	urn:ogf:network:domain=es.net:node=anycast(lbl-mr2@xe-1/3/0@*,sdsc-sdn2@xe-1/3/0@*,hous-sdn1@xe-1/3/0@*)

	VLAN	Hop
Local path	3600	urn:ogf:network:domain=es.net:node=atla-sdn1:port=xe-1/3/0:link=*
	0,2-4094	urn:ogf:network:domain=es.net:node=atla-sdn1:port=xe-0/1/0:link=xe-0/1/0.0
	0,2-4094	urn:ogf:network:domain=es.net:node=hous-sdn1:port=xe-7/0/0:link=xe-7/0/0.0
	3600	urn:ogf:network:domain=es.net:node=hous-sdn1:port=xe-1/3/0:link=*
Interdomain path	3600	urn:ogf:network:domain=umassd.net:node=andromeda:port=port-2:link=*
	3600	urn:ogf:network:domain=umassd.net:node=andromeda:port=port-3:link=*
	3600	urn:ogf:network:domain=es.net:node=atla-sdn1:port=xe-1/3/0:link=*
	3600	urn:ogf:network:domain=es.net:node=hous-sdn1:port=xe-1/3/0:link=*

Source VLAN	3600
Tagged	true
Destination VLAN	3600
Tagged	true

Documentation | ESnet | Berkeley Lab | Notice to Users

Contacts: Chin Guok, Evangelos Chaniotakis Andy Lake Eric Pouyoul

Reduced Hop Counts for Anycast 4|1 compared to Anycast 3|1

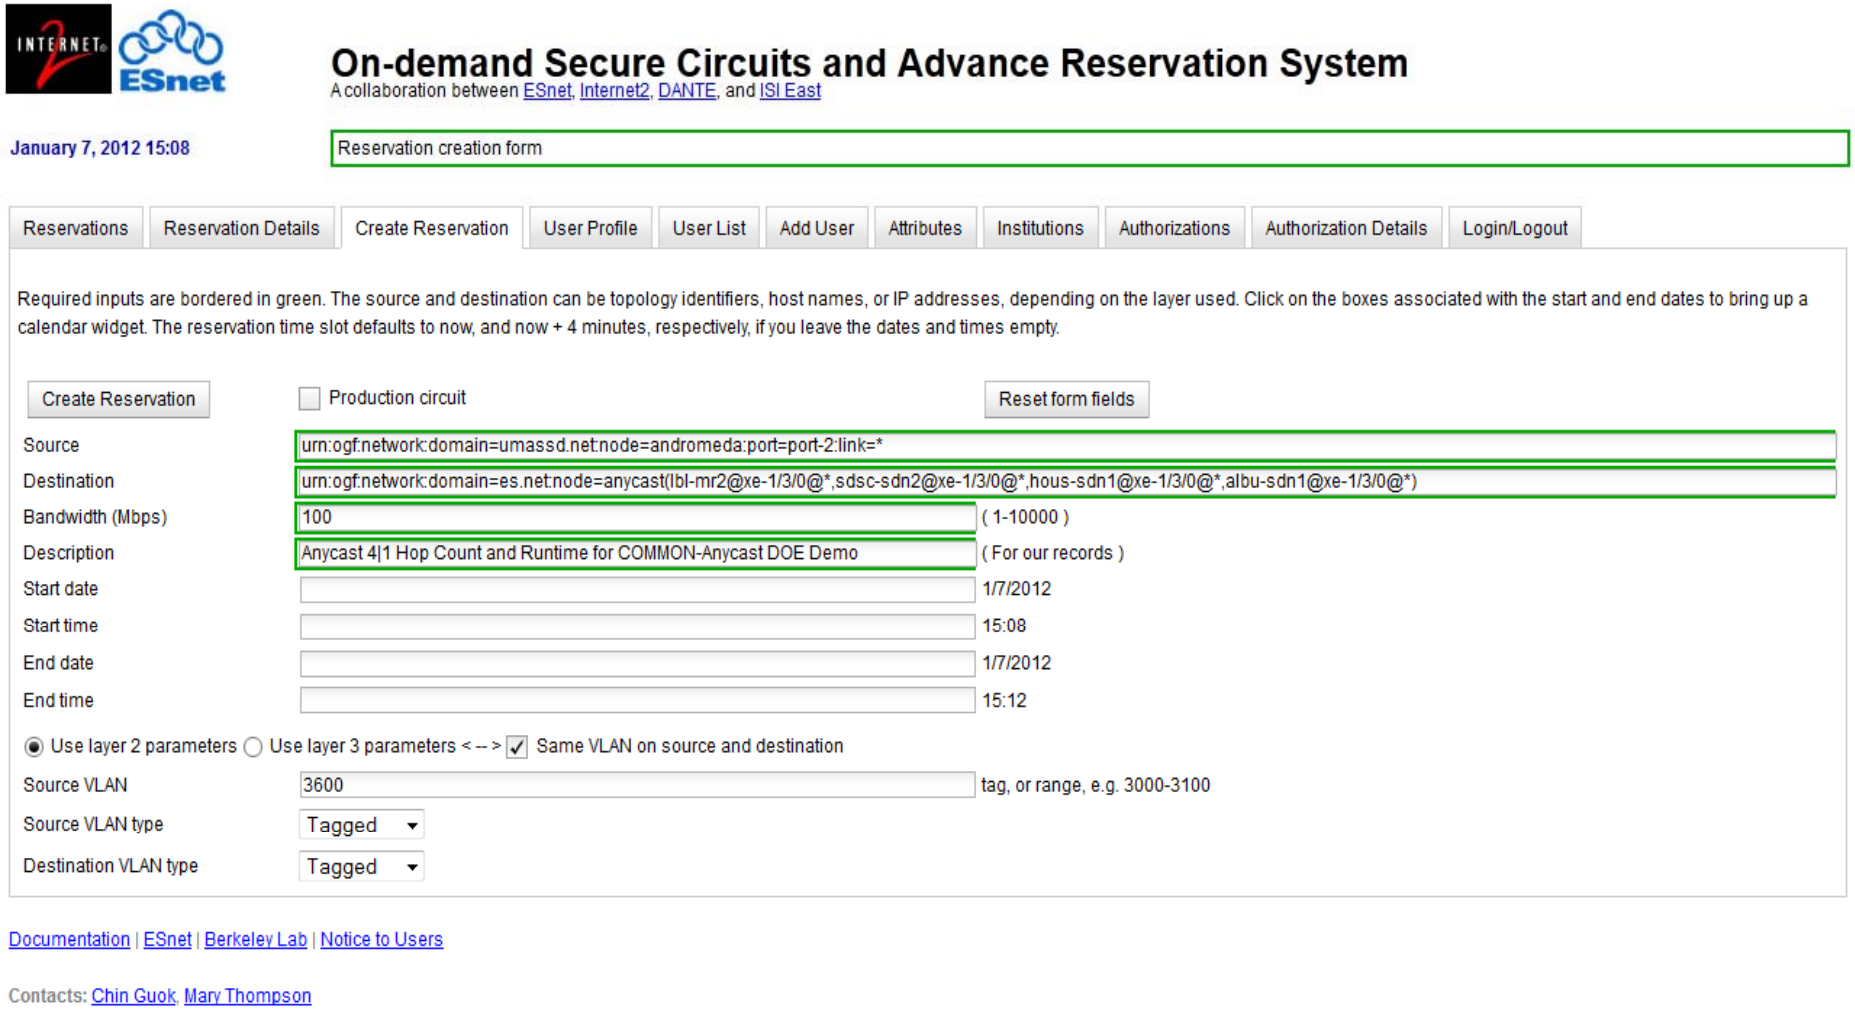

Reduced Hop Counts for Anycast 4|1 compared to Anycast 3|1

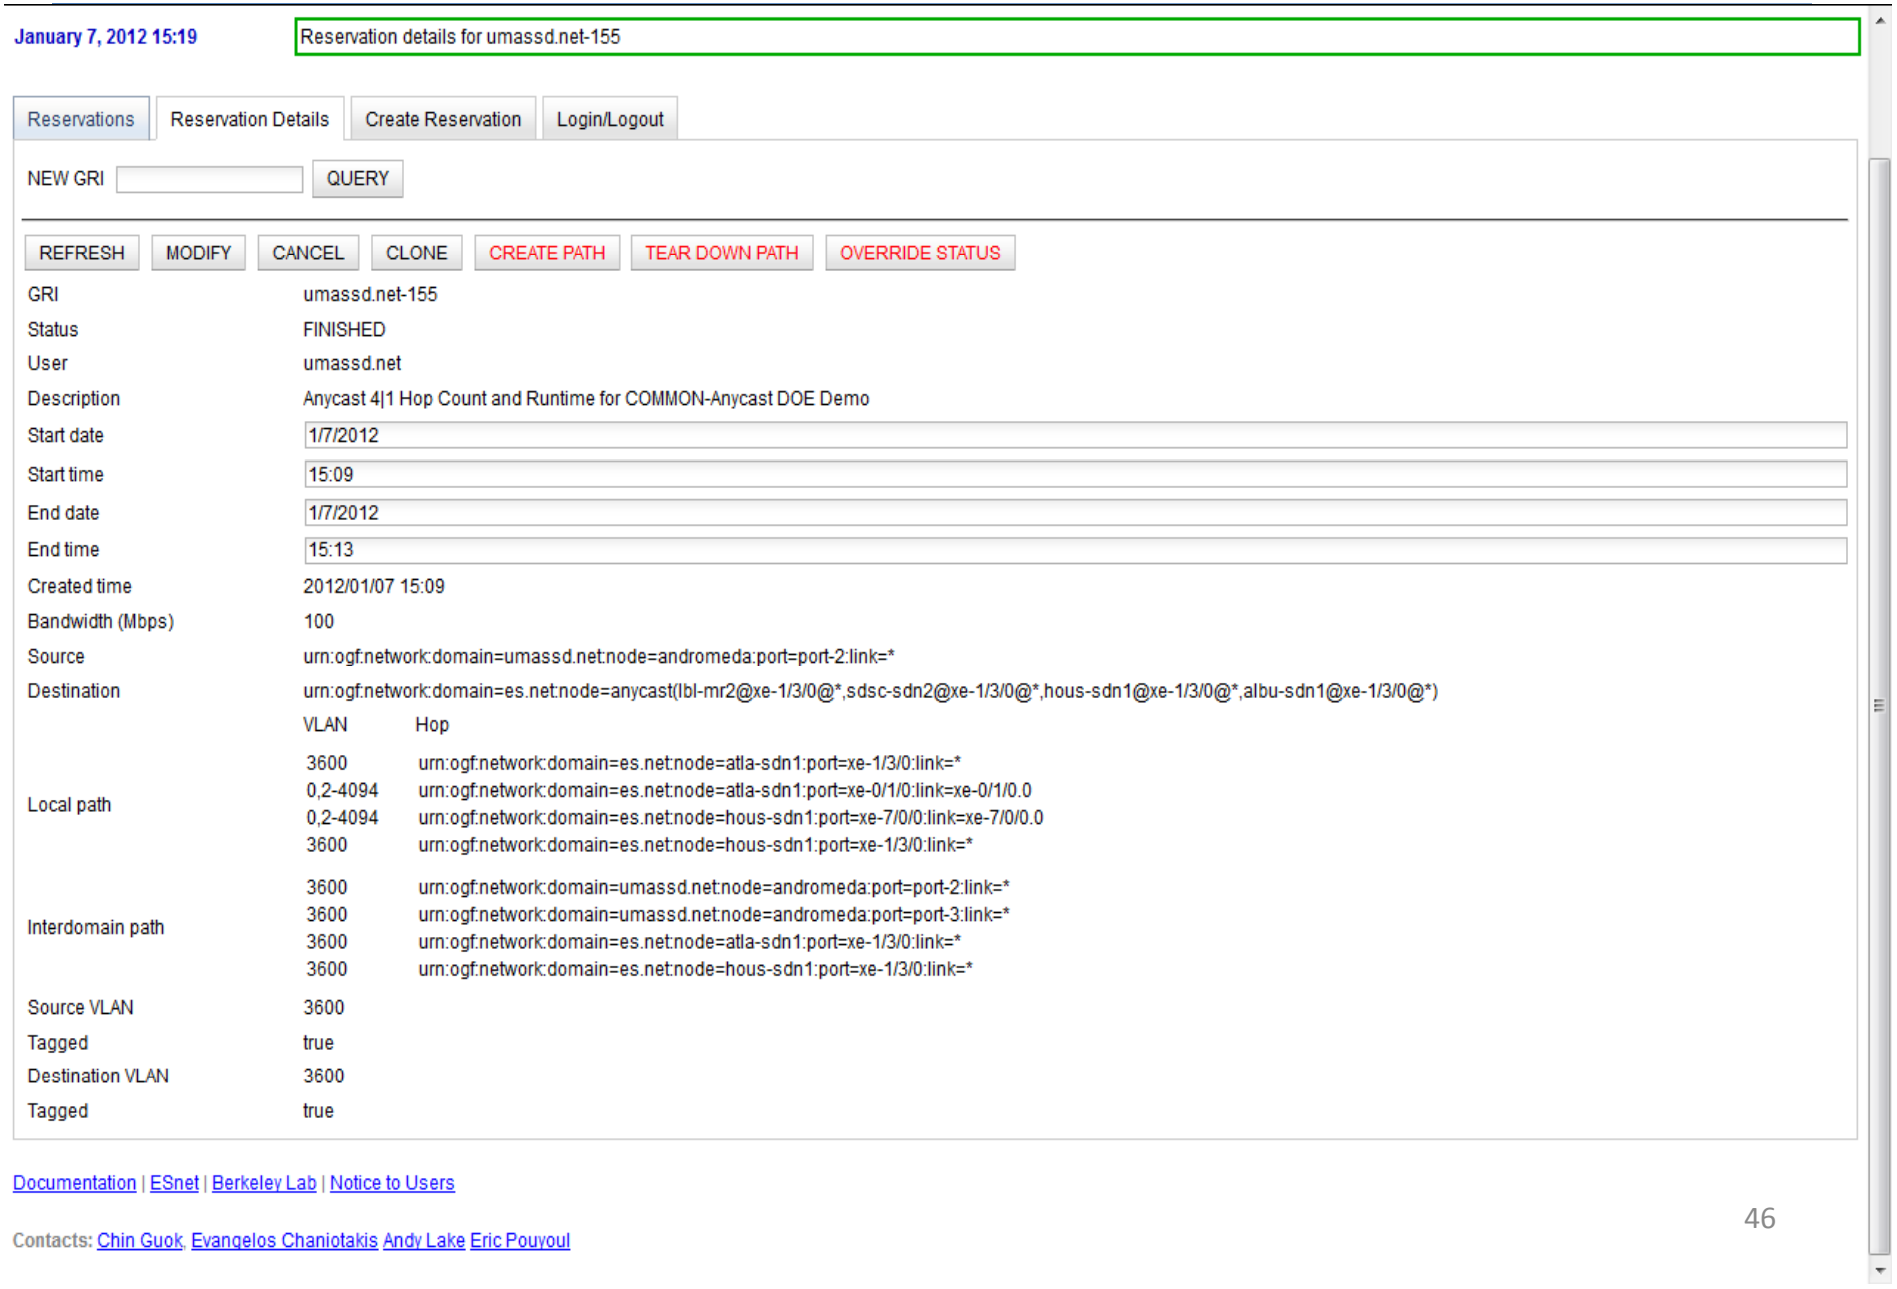

January 7, 2012 15:19

Reservation details for umassd.net-155

Reservations | Reservation Details | Create Reservation | Login/Logout

NEW GRI [ ] QUERY

REFRESH | MODIFY | CANCEL | CLONE | CREATE PATH | TEAR DOWN PATH | OVERRIDE STATUS

GRI	umassd.net-155
Status	FINISHED
User	umassd.net
Description	Anycast 4|1 Hop Count and Runtime for COMMON-Anycast DOE Demo
Start date	1/7/2012
Start time	15:09
End date	1/7/2012
End time	15:13
Created time	2012/01/07 15:09
Bandwidth (Mbps)	100
Source	urn:ogf:network:domain=umassd.net:node=andromeda:port=port-2:link=*
Destination	urn:ogf:network:domain=es.net:node=anycast(lbl-mr2@xe-1/3/0@*,sdsc-sdn2@xe-1/3/0@*,hous-sdn1@xe-1/3/0@*,albu-sdn1@xe-1/3/0@*)

	VLAN	Hop
Local path	3600	urn:ogf:network:domain=es.net:node=atla-sdn1:port=xe-1/3/0:link=*
	0,2-4094	urn:ogf:network:domain=es.net:node=atla-sdn1:port=xe-0/1/0:link=xe-0/1/0.0
	0,2-4094	urn:ogf:network:domain=es.net:node=hous-sdn1:port=xe-7/0/0:link=xe-7/0/0.0
	3600	urn:ogf:network:domain=es.net:node=hous-sdn1:port=xe-1/3/0:link=*
Interdomain path	3600	urn:ogf:network:domain=umassd.net:node=andromeda:port=port-2:link=*
	3600	urn:ogf:network:domain=umassd.net:node=andromeda:port=port-3:link=*
	3600	urn:ogf:network:domain=es.net:node=atla-sdn1:port=xe-1/3/0:link=*
	3600	urn:ogf:network:domain=es.net:node=hous-sdn1:port=xe-1/3/0:link=*

Source VLAN	3600
Tagged	true
Destination VLAN	3600
Tagged	true

Documentation | ESnet | Berkeley Lab | Notice to Users

Contacts: Chin Guok, Evangelos Chaniotakis Andy Lake Eric Pouyoul

Connection Request Provisioning Time

The following are the actual request provisioning time obtained:

- Unicast Request Provisioning Time = 16.975 seconds
- Anycast 2|1 Request Provisioning Time = 19.15 seconds
- Anycast 3|1 Request Provisioning Time = 21.316 seconds
- Anycast 4|1 Request Provisioning Time = 23.229 seconds

Benefits of Anycast over Unicast OSCARS on live deployment

In summary, during this demo we observed the following:

1. Anycast as a communication paradigm for OSCARS eliminates or reduces blocking significantly when compared to using unicast.

2. Anycast as a communication paradigm for OSCARS significantly reduces average Hop Counts required to establish circuits when compared to unicast, thereby reducing network signaling considerably as well as utilizing optimally least number of network resources .

3. Provisioning time (run-time complexity) for Anycast M|1 for $2 \leq M \leq 4$ is comparable to that of Unicast as there is only a cumulative 2 second increase in provisioning time for an unit increase in cardinality of the Anycast set when compared to unicast .