IEEE Hot Interconnects
August 18, 2005

ESnet On-Demand Secure Circuits and Advance Reservation System (OSCARS)

Chin Guok (chin@es.net)

Office of Science
U.S. DEPARTMENT OF ENERGY

QoS – New Network Service

- New network services are critical for ESnet to meet the needs of large-scale science like the LHC.
- Most important new network service is *dynamically provisioned virtual circuits* that provide:
  - Traffic Isolation
    - Will enable the use of high-performance, non-standard transport mechanisms that cannot co-exist with commodity TCP based transport. (see, e.g., Tom Dunigan's compendium http://www.csm.ornl.gov/~dunigan/netperf/netlinks.html)
  - Guaranteed Bandwidth
    - The only way that we have currently to address deadline scheduling – e.g. where fixed amounts of data have to reach sites on a fixed schedule in order that the processing does not fall behind far enough so that it could never catch up – very important for experiment data analysis.

Deployment Issues

- Establishing service offerings in the network.
  - Quality of Service (QoS) facilitates transport predictability (e.g. latency and jitter).
  - Enable traffic engineering mechanism:
    - Multi-Protocol Label Switching (MPLS) for packet switching.
    - RSVP-TE for signaling.
    - OSPF/ISIS-TE for routing.
- Accommodate networks that are shared resources.
  - Multiple QoS paths.
  - Guaranteed minimum level of service for best effort traffic.
  - Allocation Management
    - There will be hundreds of contenders with different science priorities.

End-to-End Issues (1/2)

- Virtual circuits must be set up end-to-end across multiple domains.
  - There are many issues that are poorly understood.
    - Ensuring consistent QoS characteristics end-to-end.
    - Inter-AS end-to-end MPLS LSP
      - RSVP packet filtering is almost non-existent. Extending RSVP into a neighboring AS would require implicit trust that their security and usage policies (and implementation) will prevent abuse of RSVP bandwidth requests and allocations.
      - MPLS packet filtering is also very immature compared to IP filtering.
      - MPLS, RSVP, and OSPF-TE were designed for intra-domain use.
    - Intra-AS MPLS LSP “stitched” together at AS boundaries.
      - Path between adjacent ASes must follow “normal” routing (e.g. route announcements between peering ASes).
      - Special case paths can be configured, but it must not alter “normal” routing (e.g. create a static route and isolate it within a routing instance). This is usually manually done and requires bilateral agreements to bypass normal routing paths and filter (e.g. RPF checks).

End-to-End Issues (2/2)

- To ensure compatibility, the work is a collaboration with the other major science R&E networks.
  - Code is being jointly developed with Internet2's Bandwidth Reservation for User Work (BRUW) project – part of the Abilene HOPI (Hybrid Optical-Packet Infrastructure) project.
  - Close cooperation with the GEANT virtual circuit project (Lightpaths – Joint Research Activity 3 project).

OSCARS: Guaranteed Bandwidth Service

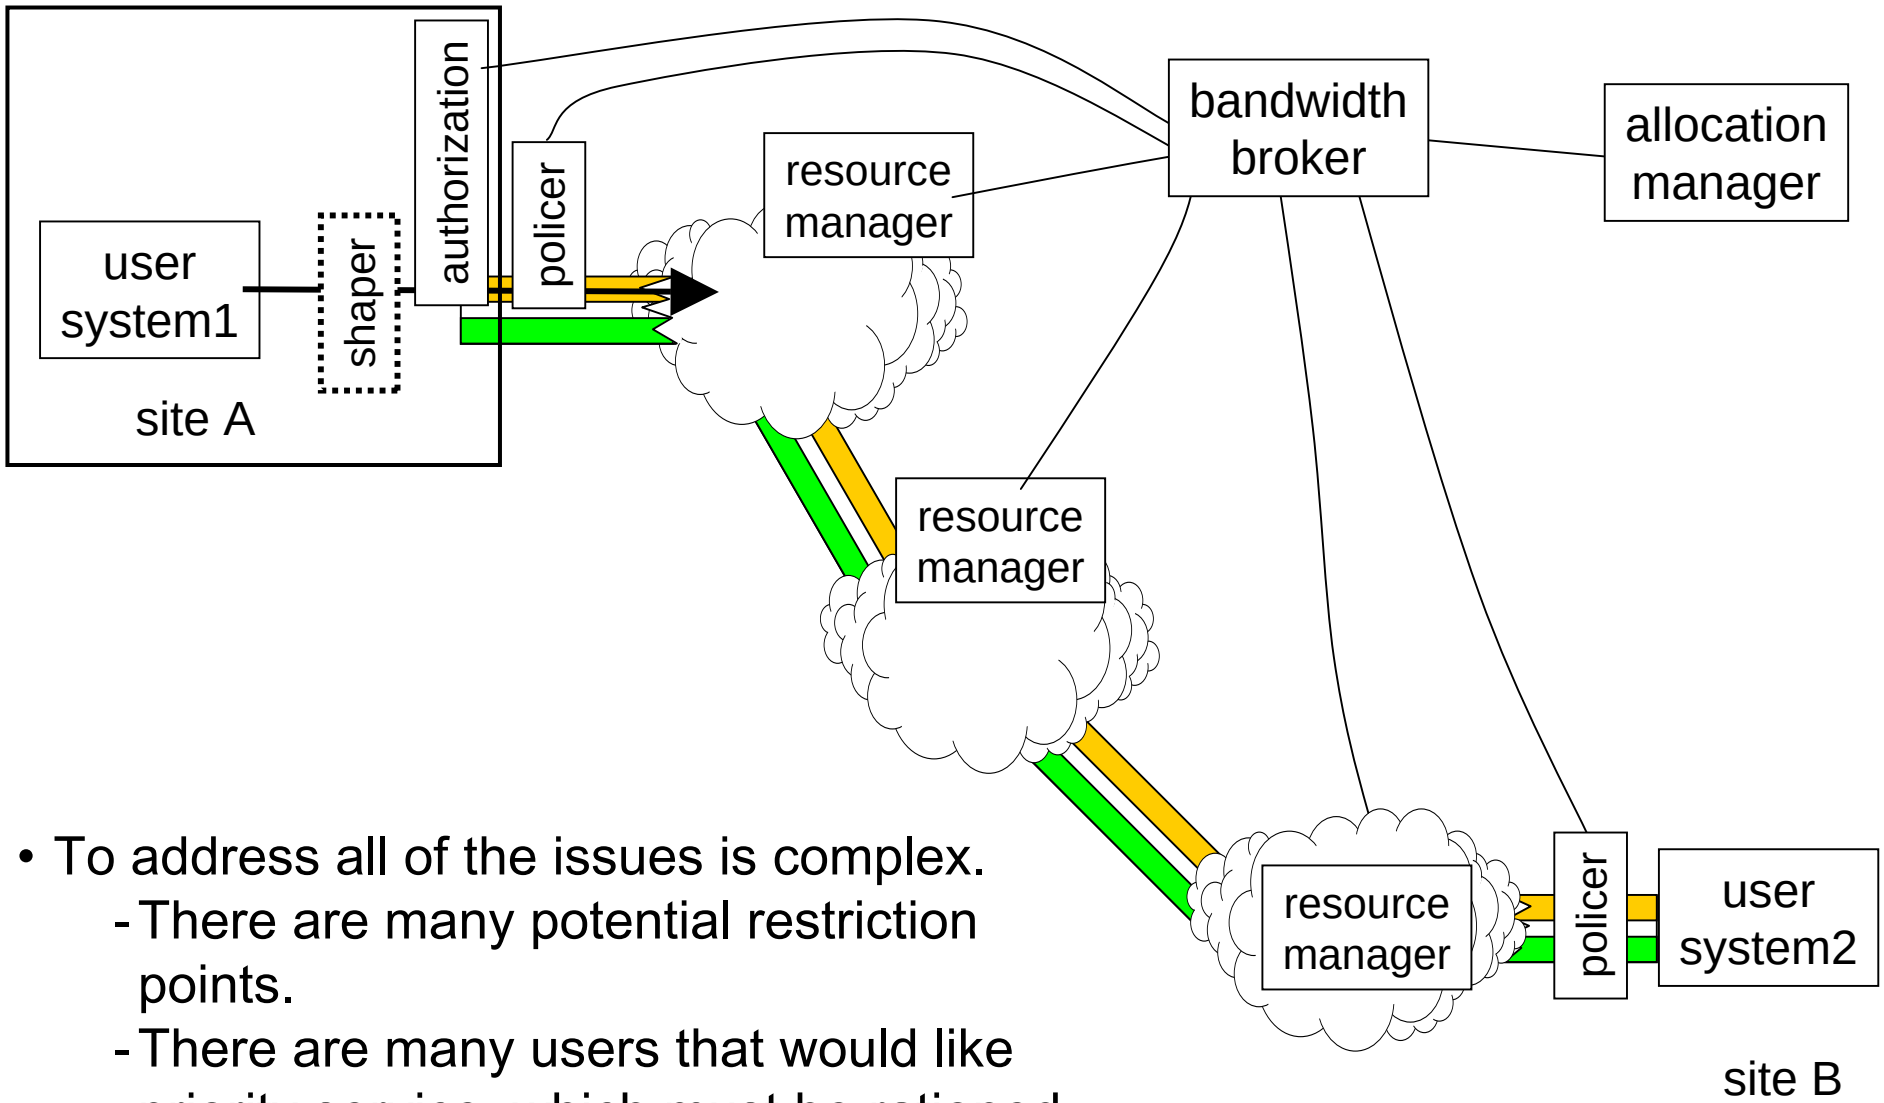

- To address all of the issues is complex.
  - There are many potential restriction points.
  - There are many users that would like priority service, which must be rationed.

OSCARS: Guaranteed Bandwidth Service

User's Perspective

- Making a reservation:
  - The user will make reservations through the Web-Based User Interface. (Subsequently, user applications may make reservations directly.)
  - Information which uniquely identifies the stream (e.g. source/destination IP addresses, source/destination port numbers, protocol), along with the duration and bandwidth requirement are entered.
  - A notice is returned if the reservation is permitted or denied.
- Claiming a reservation:
  - At the time when the reservation becomes active, the user simply sends traffic from the source to the destination.
  - The packets are filtered on the ESnet ingress router (using the information provided by the user) and injected into an LSP that was setup when the reservation became active.

OSCARS: Guaranteed Bandwidth Service

Contacts:

- Chin Guok (chin@es.net)
- David Robertson (dwrobertson@lbl.gov)

Website:

- http://www.es.net/oscars