



LambdaStation



Fermilab



Cal Tech

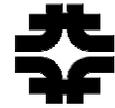
High-Performance Network
Research PI Meeting

BNL

Phil DeMar
September 29, 2005



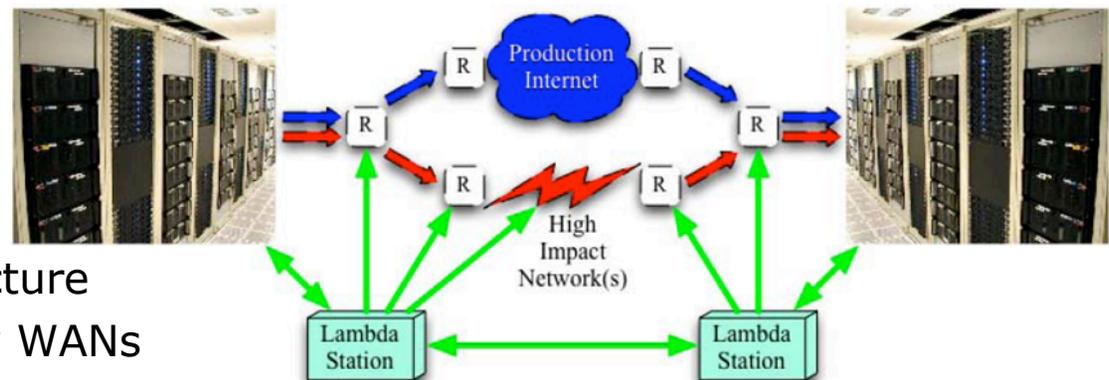
LambdaStation



- A network path forwarding service to interface production facilities with advanced research networks:
 - Goal is selective forwarding on a **per flow** basis
 - Alternate network paths for high impact data movement
 - Dynamic path modification, with graceful cutover & fallback
 - Current implementation based on policy-based routing & DSCP marking

- LambdaStation interacts with:

- Host applications / systems
- LAN infrastructure
- Site border infrastructure
- Advanced technology WANs
- Remote LambdaStations

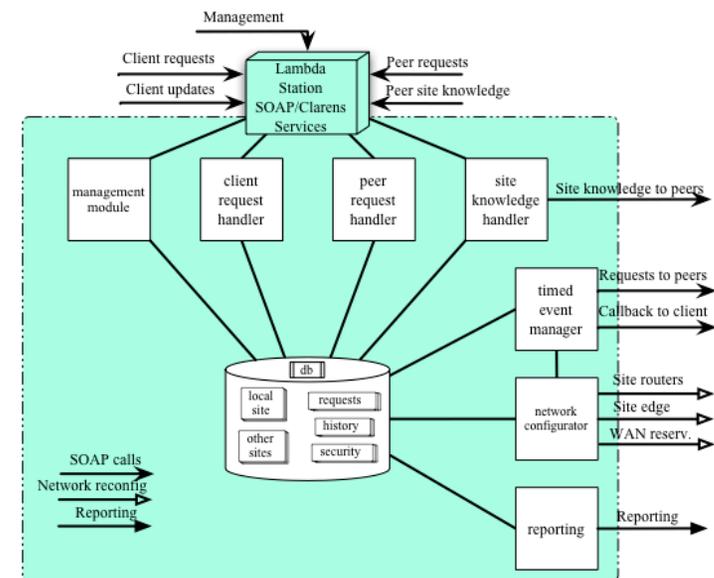




LambdaStation Building Blocks

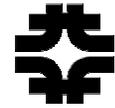


- Soap service (Clarens) supporting:
 - LambdaStation ↔ client communications
 - Inter-LambdaStation communications
 - Integrated scheduler function
- Network configuration module:
 - Internal path forwarding reconfiguration
 - WAN alternate path coordination
 - Site perimeter modification
- Local LambdaStation configuration:
 - Local site:
 - Alternate WAN path definitions
 - PBR client (internal alternate path definitions)
 - Local PBR clients (host systems/netblocks)
 - Requestors (identity strings for authentication)
 - At least one remote LambdaStation
 - Dynamic exchange of known LambdaStation information





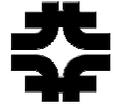
Development Status



- LambdaStation V1.0 implements full cycle functionality:
 - Functional Soap service:
 - LambdaStation interface protocol defined with 20 primitives
 - Ticket request process operational, including remote end feedback
 - _ Inter-LambdaStation coordination working...
 - SSL authentication (today...)
 - Dynamic network reconfiguration working:
 - PBR forwarding set & removed in both directions on multiple routers
 - Concurrent modification of interface ACLs to permit
 - LambdaStation-aware applications developed:
 - LSImperf
 - LSTraceroute
 - Monitoring capabilities developed:
 - Scripts developed for real time snmp interface displays
 - Quasi-real time analysis of flow data



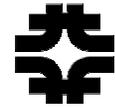
Deployment & Testing



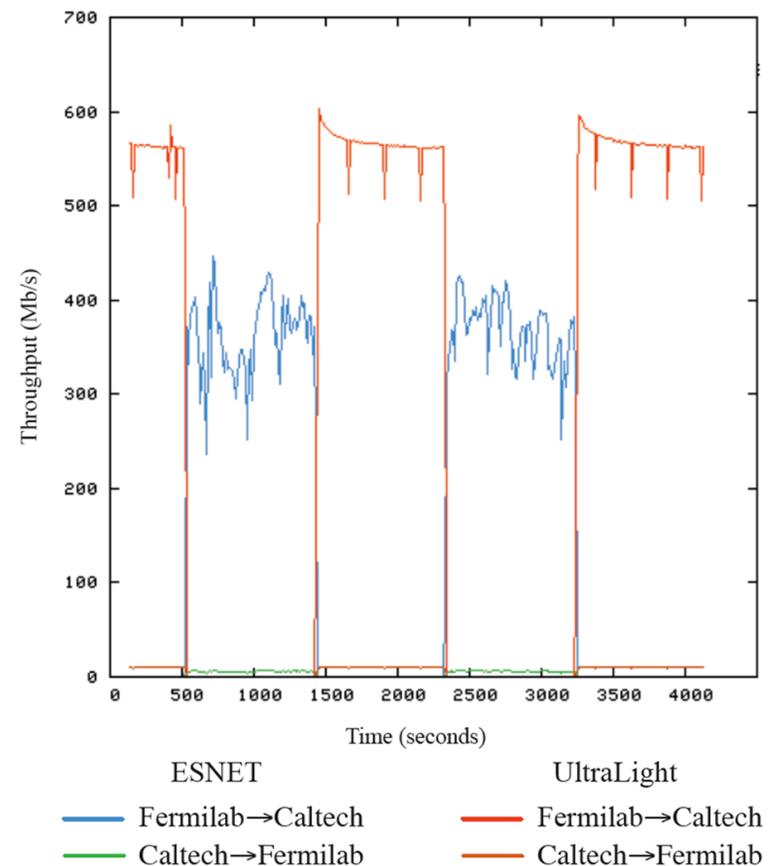
- LambdaStation servers deployed at FNAL (2) & Cal Tech
- Alternate WAN path testing with UltraScience Net & UltraLight conducted
 - Static scheduling & WAN path setup, at this stage...
- Test network environment developed:
 - Cisco & Force10 (soon...) routers for PBR manipulation
 - Four 10GE systems and 60+ node test farm
- Multiple simultaneous alternate path requests tested:
 - Three in each direction worked fine; scalable to order 10s
- Designed simulation of more advanced reservation system with resource scheduler:
 - Concerns about deadlock conditions & mutual exclusion
 - Four LambdaStations with three concurrent requests tested OK



Path Switching Experiment

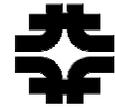


- Continuous GridFTP transfers:
 - Path toggled at 15 min intervals by LambdaStation servers
 - 1GE NIC, dual 750Mhz source
 - 18 TCP streams
 - Not optimized for performance
- Each path appears reasonably consistent:
 - “Choppiness” of blue (ESnet) path due to OC12 congestion
 - Red (UltraLight) spikes are periodic system resource issue
- Overall, impact of dynamic path changes were not significant

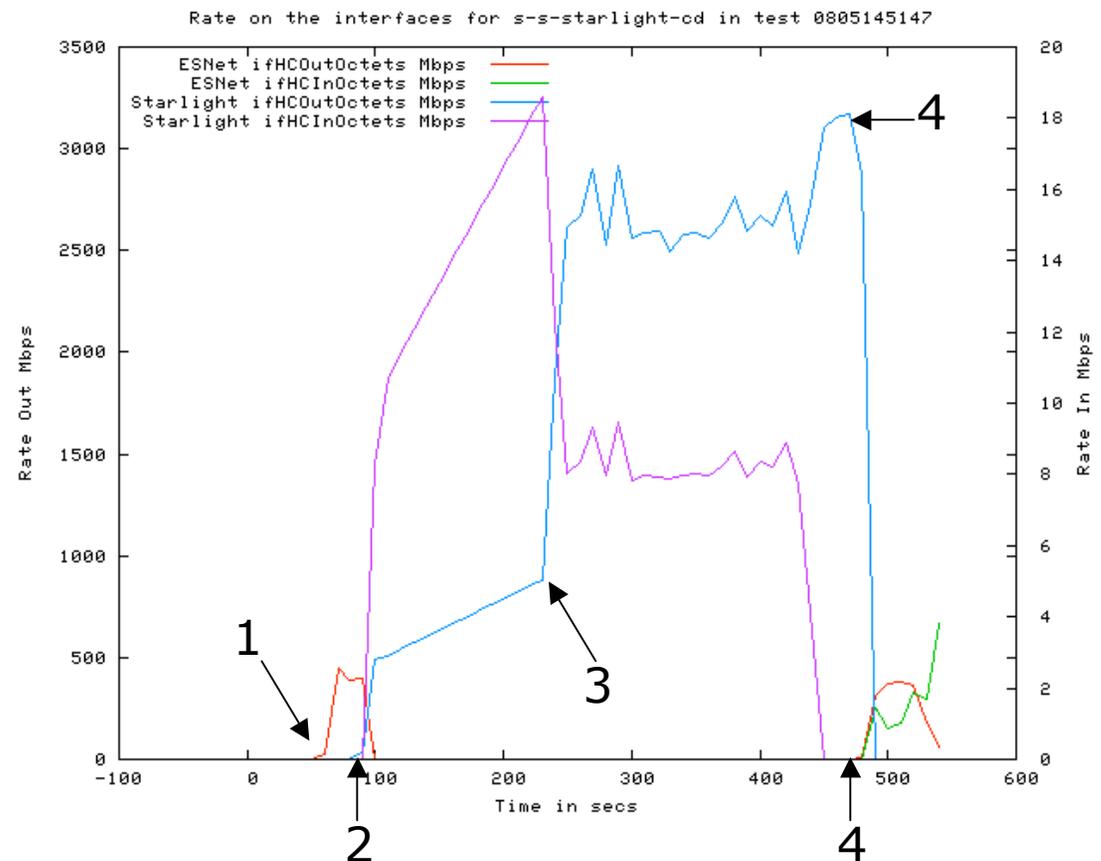


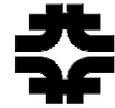


LSiperf End-to-End Test



1. Data transfer started:
 - 10GE host; 5 tcp streams
 - Network path is via ESnet
 - OC12 bottleneck...
 - Path MTU is 1500B
 - LambdaStation service ticket is opened
2. LambdaStation changes network path to USN
3. Host path MTUD check detects larger path MTU
4. LambdaStation service ticket expires:
 - Network path changed back to ESnet



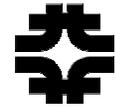


Future Directions

- Development Plans:
 - LambdaStation awareness in Storage Resource Manager:
 - GridFTP DSCP tagging
 - Control channel dialog with GridFTP server
 - Distributable LambdaStation kit
 - Automated coordination/setup with WANs
 - Authentication based on GRID certificates
 - Integral path switching monitoring capability
- Testing & Deployment:
 - Deploy LambdaStations at additional sites
 - CERN, others?
 - Test across different advanced technology WANs
 - MPLS tunnels (OSCARs)
 - Test with different (non-Cisco...) LAN hardware
 - Instigate extensive failure mode testing



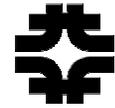
Closing Comments



- LambdaStation Project progress:
 - In line with our project plan, and about what was expected
 - Outside interest remains very high
 - LambdaStation deals with a problem space that's still largely avoided
 - Other uses for LambdaStation have been suggested:
 - _ Dynamic bypass of perimeter firewall for high impact traffic
- Issues & Concerns Remain:
 - Complexity
 - Level of customization:
 - May result in high level of local effort to implement
 - Uncertainty about adapting to less traditional (Cisco...) environments
 - Application-awareness is a 'stubborn' issue
 - Avoiding mission-creep:
 - LambdaStation is an alternate path forwarding service, not a bandwidth reservation service



Supplemental Slide 1: LS Interface Primitives



- Informative Group:

whoami	return identity of local or remote LS
sayHello	a probe message to test a remote LS
ip2client	return site and PBR-client identity based for IP address if exists
getClientInfo	return description of PBR-client
getKnownClients	return the list of known clients (site,client)
getKnownLambdas	return list of all known lambdastations
getStationParameters	return parameters of lambdatstion
NetConfigMode	return current mode of network configuring (static or dynamic)

- Service Group:

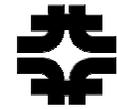
openSvcTicket	place request to establish an alternative path
cancelTicket	abnormal completion of ticket
completeTicket	normal complete of ticket before it expires otherwise expiration process will be applied
getTicket (getTicketStatus)	return ticket's parameters, e.g. its current status
getFlowSpec	return the full description of flows associated with ticket
updateFlowSpec (updateDSCP)	update flow specification
updateTicket	update ticket parameters
getMyRemoteID	get ID of remote ticket associated with localID

- Internal:

getNextDSCP	select the next free DSCP code (in case site use dynamic mode for DSCP assignment)
assignDSCP	assign DSCP to ticket
updateMyLambdaDefs	check for updates (new, remove) PBR clients
localRequest/remoteRequester	determine the authenticated requester of ticket

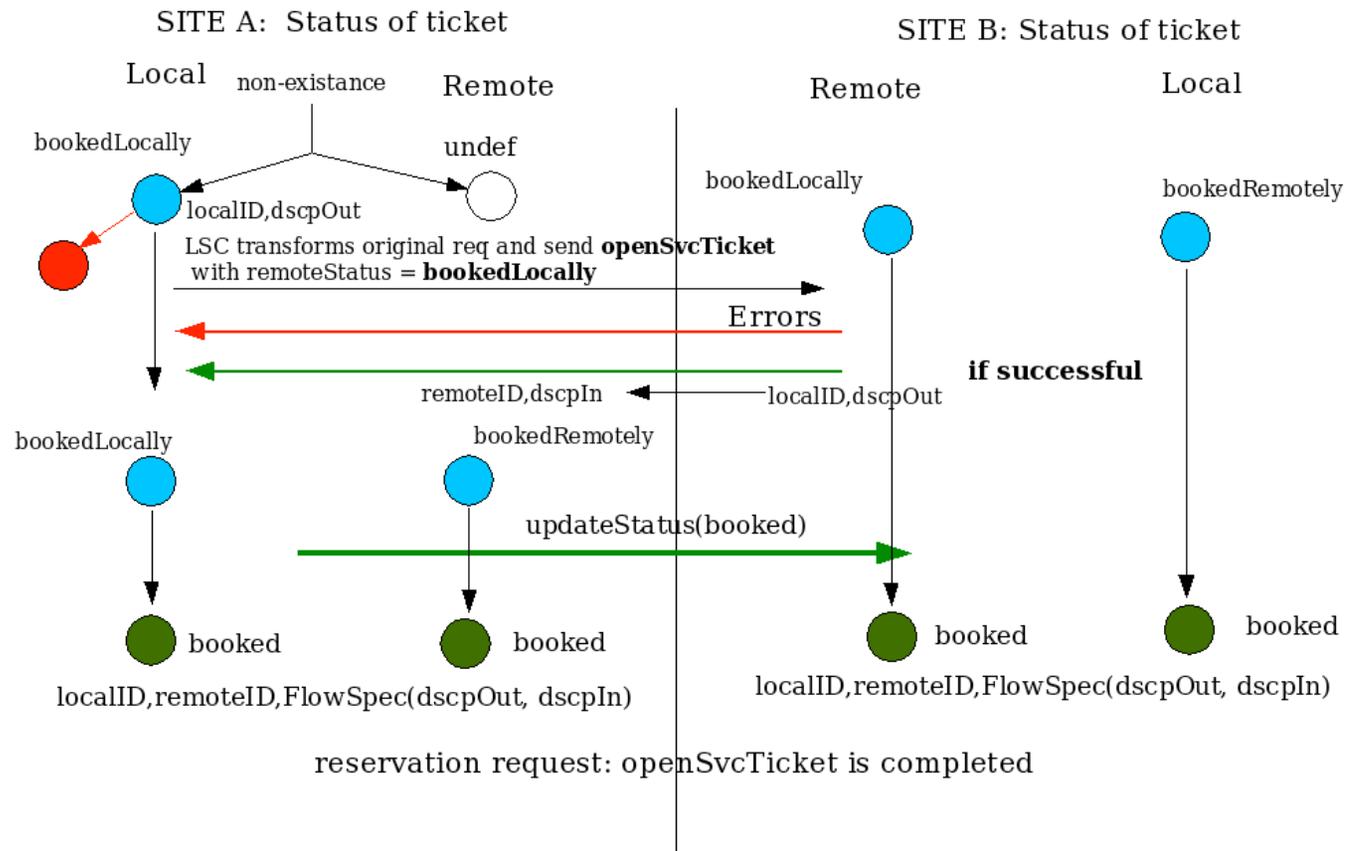


Supplemental Slide 2: Open Service Ticket



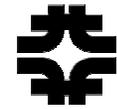
A.Bobyshev, 04/27/2005

openSvcTicket – open service ticket





Supplemental Slide 3: Ticket State Diagram



The diagram of ticket's states

