



Overview of UvA-SNE activities for international connectivity

Ralph Koning, Jeroen van der Ham
University of Amsterdam
SNE Group

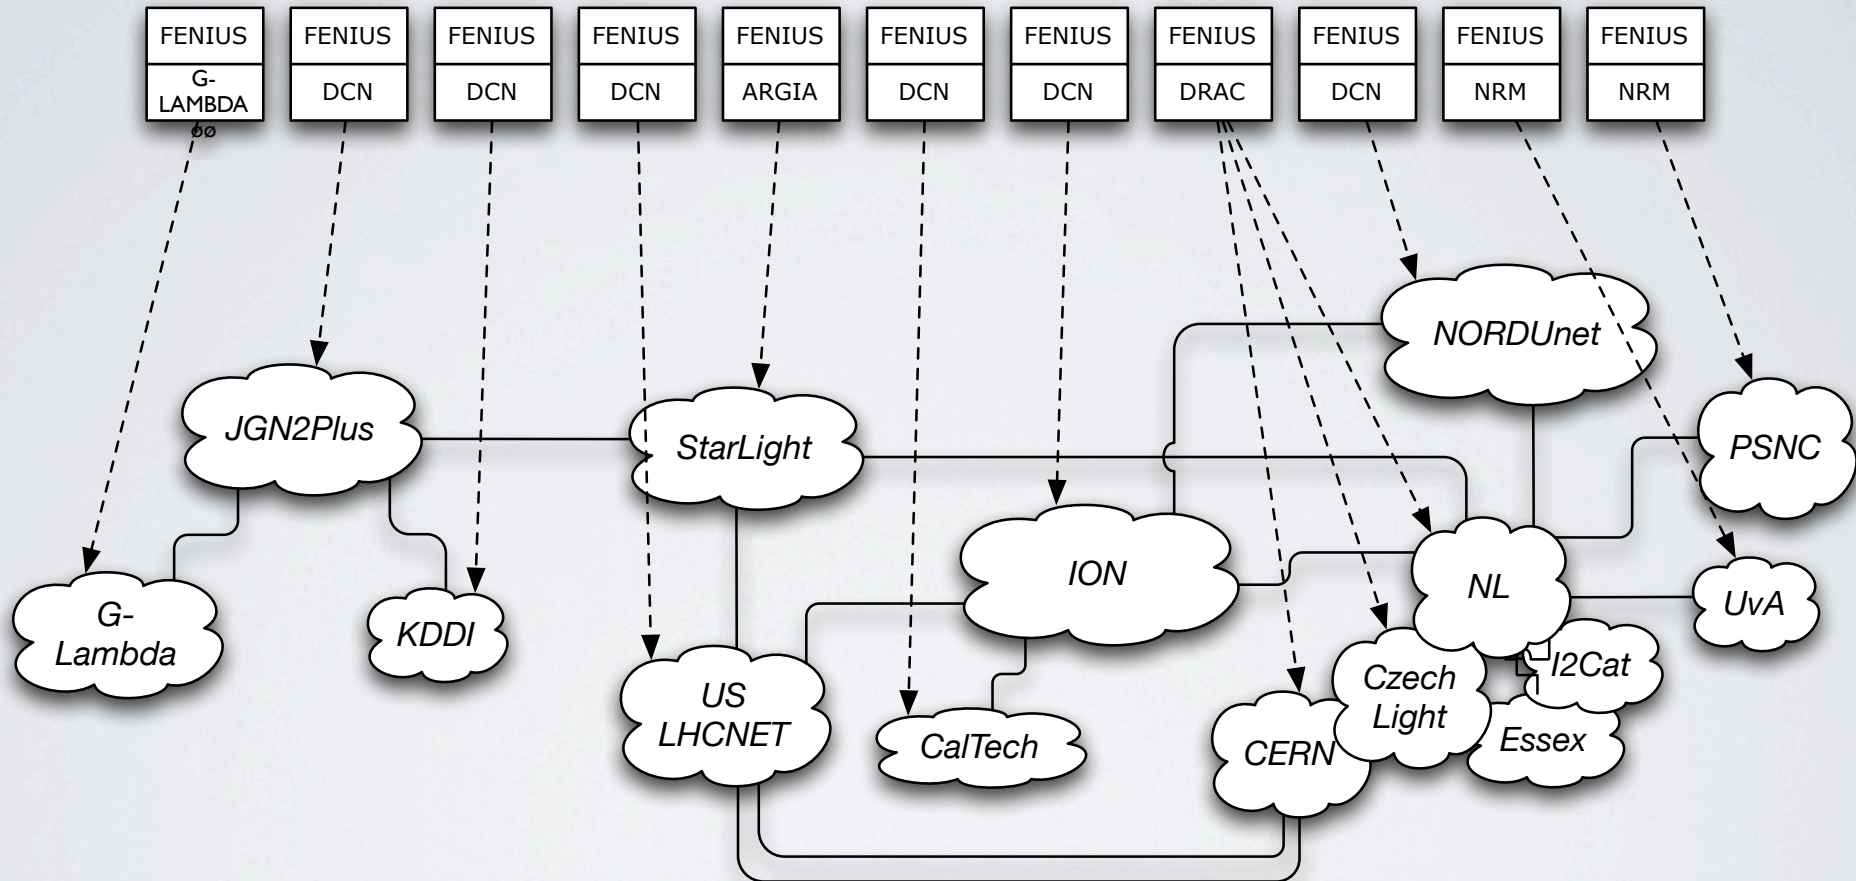
With contributions from: Paola Grosso, Cosmin Dumitru, Adianto
Wibisono and Cees de Laat



Activities in GigaPort

- Automated GOLE (Ralph, Jeroen)
 - experiment with next-generation switching equipment
 - develop of control-plane components
 - study of inter-GOLE provisioning efficiency
 - continue providing input in NSI/NML
- 40G/100G Ethernet (Ralph, SARA)
 - To pave the way for building very high end eScience infrastructures
- Aggregation and Multi-Layer Pathfinding Algorithms (Jeroen)

GLIF Automated GOLE Logical Topology





UvA GOLE Testbed SC10

sw* = Dell Powerconnect 5324

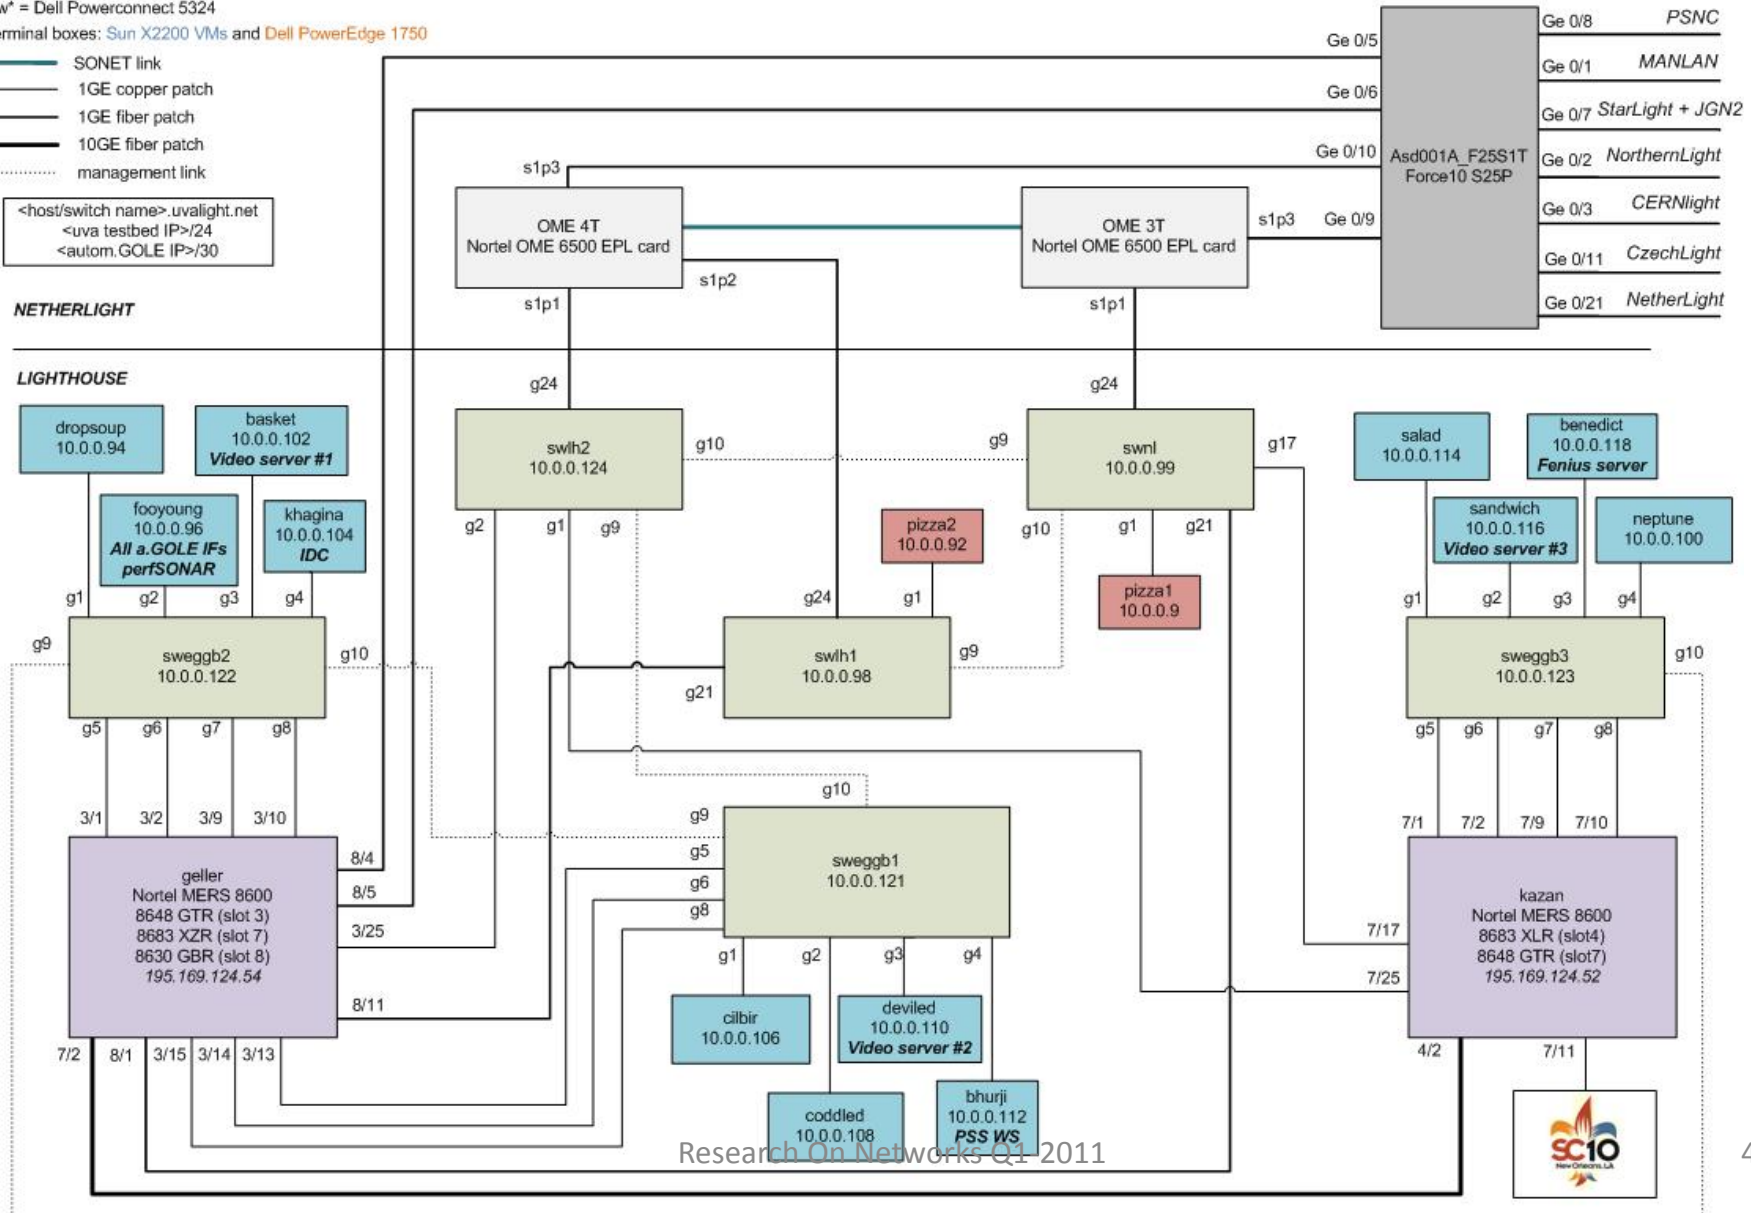
terminal boxes: Sun X2200 VMs and Dell PowerEdge 1750

- SONET link
- 1GE copper patch
- 1GE fiber patch
- 10GE fiber patch
- management link

```
<host/switch name>.uvalight.net  
<uva testbed IP>/24  
<autom.GOLE IP>/30
```

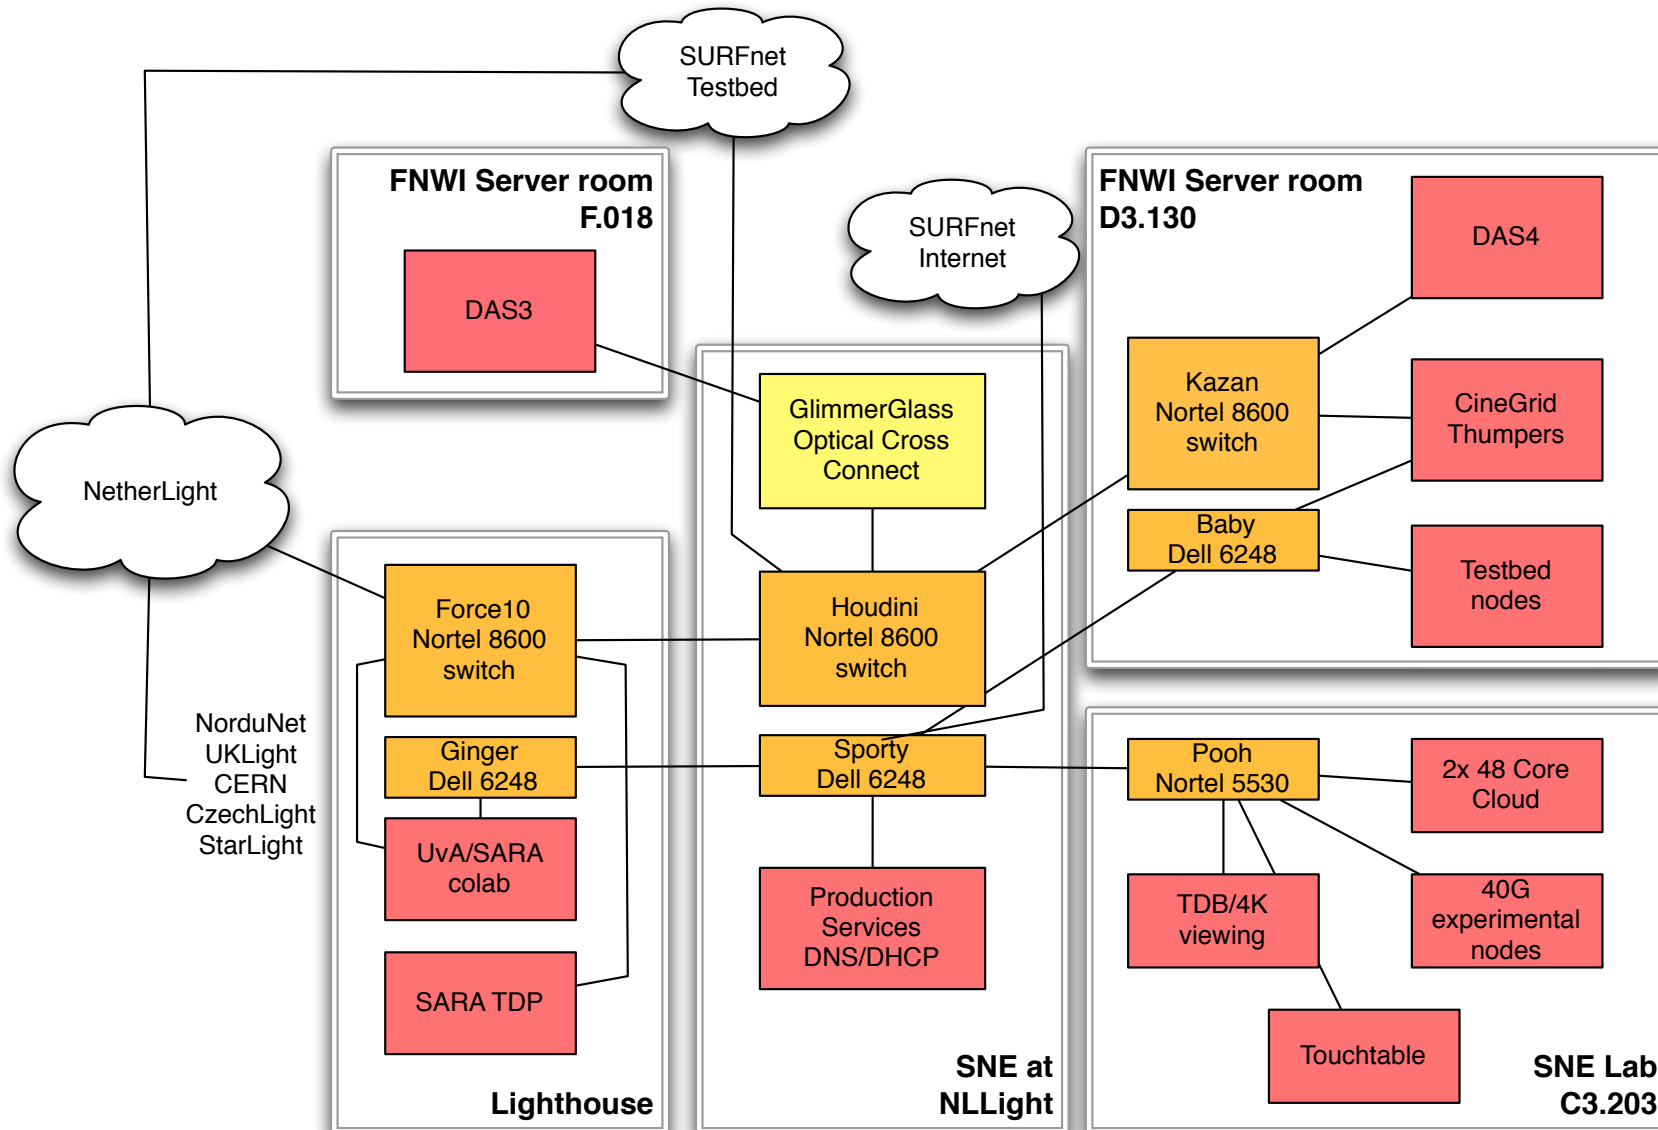
NETHERLIGHT

LIGHTHOUSE





'new' UvA Testbed





Resources for automated GOLE

- Available hardware
 - 2x MERS 8600
 - 1x Force10 Etherscale E600
 - FTOS version outdated (2005)
 - Dell 5324
 - 3 Sun x2200 nodes or virtual machines
- Connectivity
 - 2 links to SURFnet F10
 - 3 links to Testbed OME's

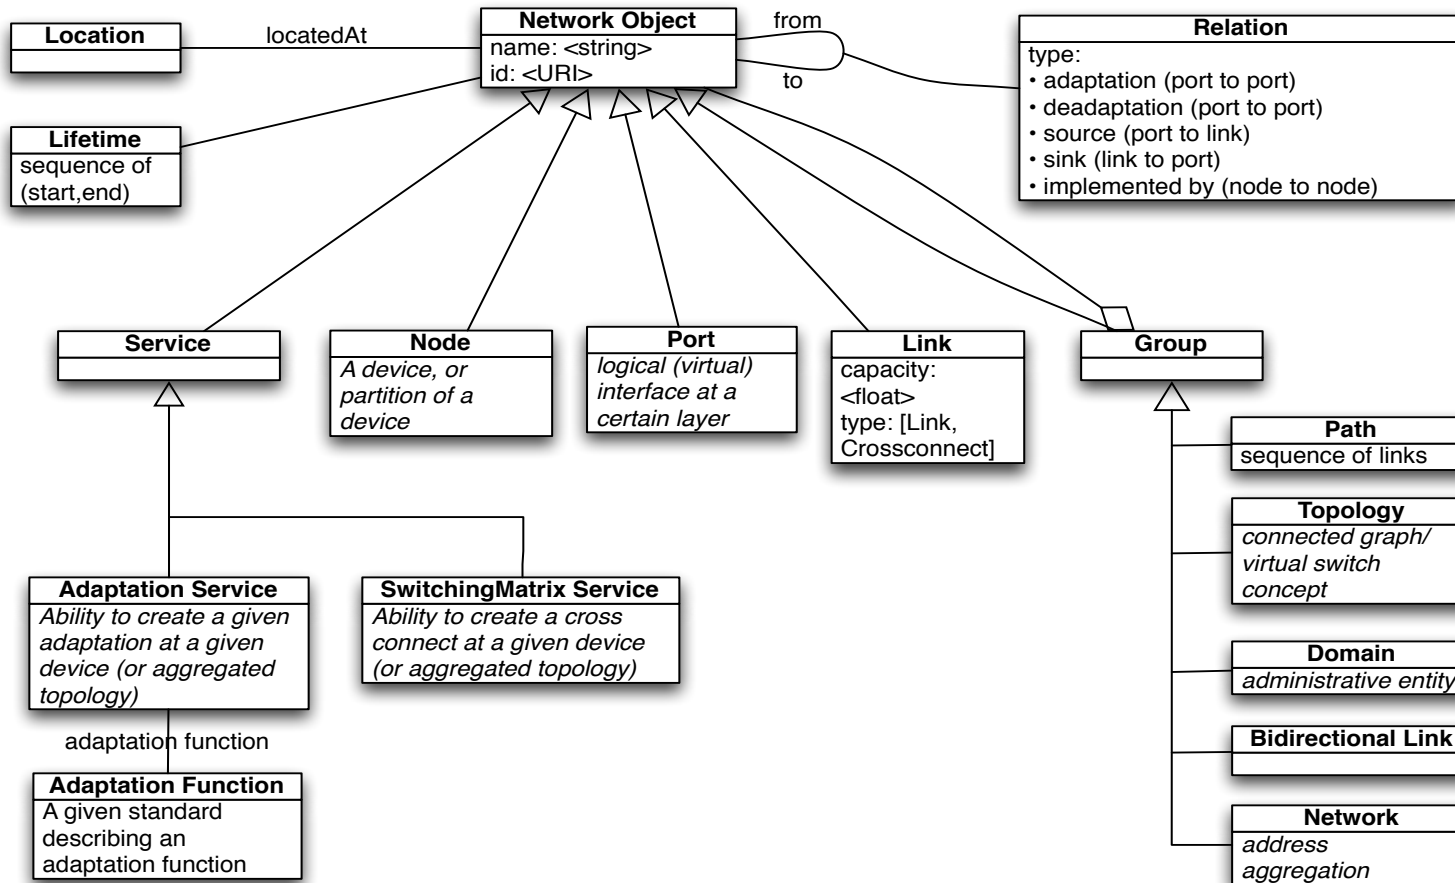


Future Testbed

- Future UvA testbed
 - See if we can upgrade FTOS on our Force10
 - MERS 8600 support for OpenDRAC
 - End of Life
 - Buy something new
- Mac-in-Mac
 - MERS 8600 supports PBB/802.1ah
 - Can we support this with newer FTOS load?



NML Schema



- End goal is to use NML to describe the Topology
 - But this a lot of work



Automated GOLE Topology description

- Simple RDF/OWL based topology description
 - Just enough information to make it usable
 - Connection points
 - Connections
 - Endpoints
 - Lowers threshold allowing new users participate more easily
 - Should be easy to map to NML/NDL



Ontology Editor

SNE Editor SNE's Network Topology OWL Editor for Cinegrid

New Load Save Upload Delete Help

Graph View OWL View

Cinegrid OWL

- Content
 - Item
 - AudioClip
 - VideoClip
 - List
- Document
 - Element
- Device
 - Host
 - Projector
 - Screen
- Group
 - Cluster
- Exchange
 - Node
- Service
 - Authenticator
 - Indexer
 - XpIndexer
 - Storage
 - iRODSStore
 - LocalStore
 - NFSStore
 - Streamer
 - NTTStreamer
 - SAGEStreamer
 - ZaxelStreamer
 - Transcoder
 - Visualizer
 - NTTVisualizer
 - SAGEVisualizer

NDL Domain

- NDL Topology
- NDL VPN
- WGS84 Position
- Geyser
- Composed

Object Properties

Owl Classes

Data properties

Properties

Title	CDL_Amsterdam
Description	Attempting to build http://cinegrid.uvalight.nl/owl/odi-amsterdam-2.0.owl within this editor

Outline

Info

SNE Cinegrid editor is designed to allow user to easily create network topology based on Network Description Language and Cinegrid Ontology. Extensions to other ontologies will be possible in the future.

Host

Data properties of Host	
Name	Amsterdam

Host

Data properties of Host	
Name	cgvideo
hostName	cgvideo.uvalight.nl
cpuThreads	8
osFlavor	Ubuntu
osName	Linux
osVersion	9.04
cpuFreq	2900
cpuArch	amd64
cpuCache	1024

Host

Data properties of Host	
Name	cgdaemon
hostName	cgdaemon.uva.netherlight
cpuThreads	4
osFlavor	Ubuntu
osName	Linux
osVersion	8.04
cpuFreq	1000
cpuArch	amd64
cpuCache	1024

Host

Data properties of Host	
Name	suitcees
hostName	suitcees.science.uva.nl
cpuThreads	8
osFlavor	Solaris
osName	SunOS
osVersion	5.11
cpuFreq	1386
cpuCache	9999

LocalStore

Data properties of LocalStore	
Name	LocalStore_cgvideo
assignedStorageSpace	9999
storageType	9999
usedStorageSpace	9999
availableStorageSpace	9999

LocalStore

Data properties of LocalStore	
Name	LocalStore_cgdaemon
assignedStorageSpace	9999
storageType	9999
usedStorageSpace	9999
availableStorageSpace	9999

SAGEVisualizer

Data properties of SAGEVisualizer	
Name	SAGEVisualizer_cgvideo
pixelsY	3840
pixelsX	2160
displayType	

SAGEStreamer

Data properties of SAGEStreamer	
Name	SAGEStreamer_cgdaemon
streamingType	

iRODSStore

Data properties of iRODSStore	
Name	iRODSStore_suitcees
assignedStorageSpace	9999
storageType	9999
usedStorageSpace	9999
availableStorageSpace	9999

LocalStore

Data properties of LocalStore	
Name	LocalStore_suitcees
assignedStorageSpace	9999
storageType	9999
usedStorageSpace	9999
availableStorageSpace	9999

Host

Data properties of Host	
Name	cgtrfm
hostName	cgtrfm.vlan400.uvalight.nl
cpuThreads	1

Host

Data properties of Host	
Name	cgdevil
hostName	cgdevil.uvalight.nl

Host

Data properties of Host	
Name	basketcees
hostName	basketcees.science.uva.nl
cpuThreads	8

LocalStore

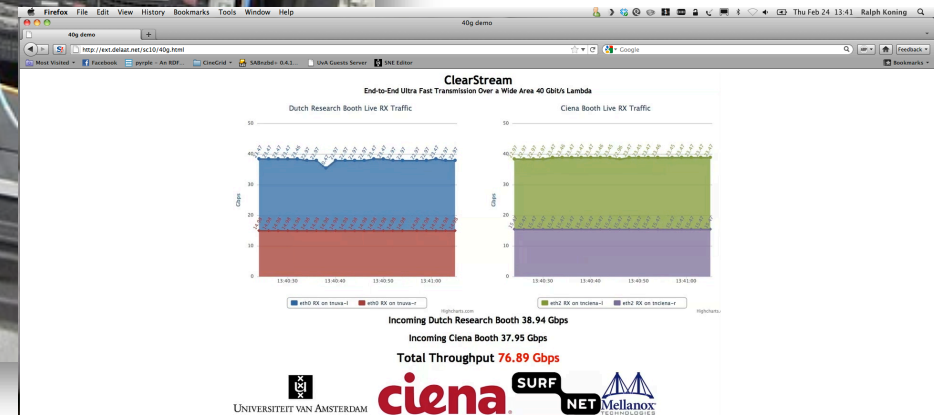
Data properties of LocalStore	
Name	LocalStore_suitcees
assignedStorageSpace	9999
storageType	9999
usedStorageSpace	9999
availableStorageSpace	9999

NFSStore

Data properties of NFSStore	
Name	
assignedStorageSpace	
storageType	
usedStorageSpace	
availableStorageSpace	



40G Ethernet






40G Ethernet (2)

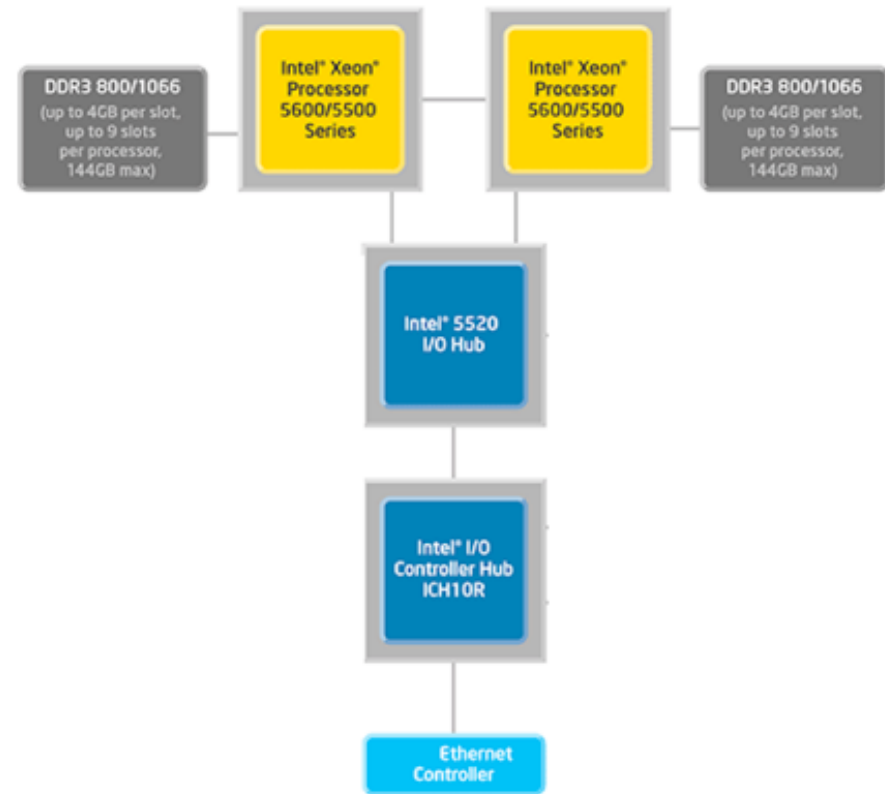
- Successful demos at:
 - GLIF 2010 (1650km Amsterdam – Geneva)
 - Supercomputing 2010 (End to end 40GE)
- Current server architecture not ready.
 - PCI-express gen2-8x 26 Gb/s limit
- Still looking for a killer app.
 - Single stream
 - No iperf



Server Architecture




DELL R815
4 x AMD Opteron 6100



Supermicro X8DTT-HIBQF
2 x Intel Xeon




Hardware Limitations - AMD






Hardware Limitations - Intel

Single stream TCP performance with iperf pinned on both nodes to specific cores - interrupts handled on the receiver node by core 7
300s average





Extending InfiniBand Over 40G Ethernet Wide Area Network Demonstration Connectivity Diagram





40GE Research Ideas

- Look how 40G IB performs over long distances/high latency
- SSD to SSD file replication over 40G
 - TCP/IP based on old constraints
 - Little memory
 - Slow diskdrives
 - Look into methods to improve this.



References

- Ontology Editor
 - <http://sneneteditor.appspot.com>
- Network Markup Language (NML)
 - http://www.ogf.org/gf/group_info/view.php?group=nml-wg
- 40G Ethernet
 - ClearStream: Prototyping 40 Gbps Transparent End-to-End Connectivity
 - OnVector 2011 Conference Proceedings