

Wisconsin Site Report

Tapas Sarangi



WISCONSIN
UNIVERSITY OF WISCONSIN-MADISON

T2_US_Wisconsin

Personnel:

- Sridhara Dasu (PI)
- Ajit Mohapatra, Tapas Sarangi (Tier-2)
- Carl Vuosalo (AAA)
- Help from department computing (Dan Bradley, others)

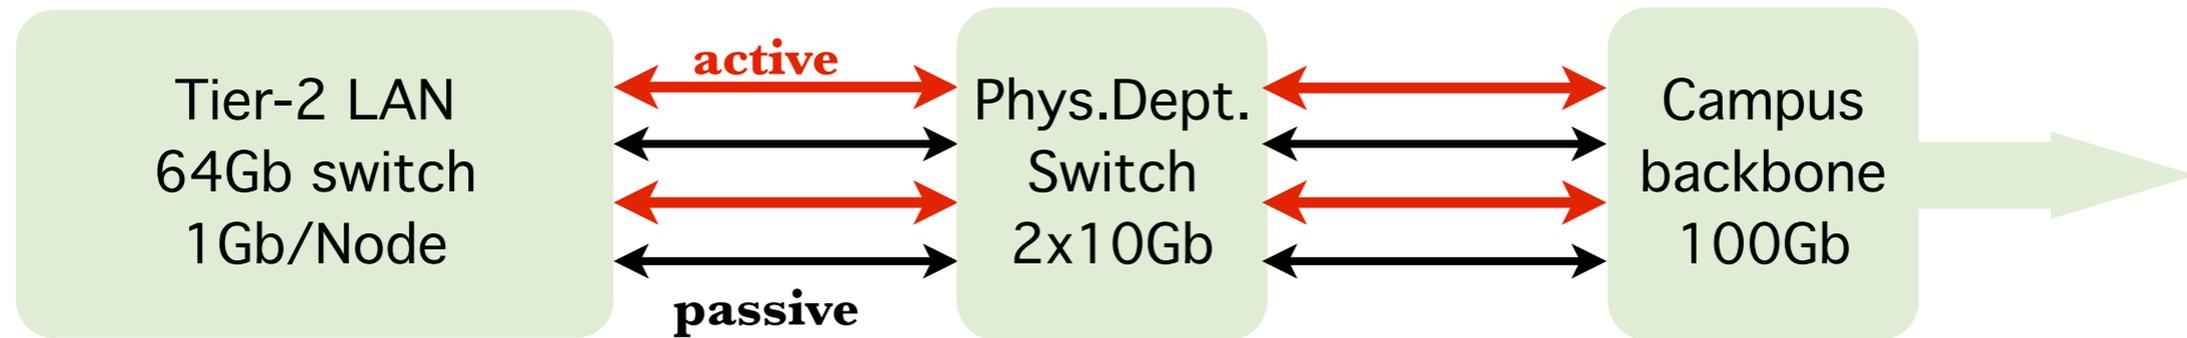
Gen	CPU Class	Slots*	HS06/slot**	Total HS06	Storage (TB)
g12	8 x 2.66 GHz Xeon X5355	224	8.31	1861	66
g14	8 x 3.00 GHz Xeon E5450	184	9.79	1801	88
g16	8 x 3.00 GHz Xeon E5472	432	10.17	4393	–
g18	16 x 2.4 GHz Opteron 6136	496	9.30	4612	189
g19	24 x 2.2 GHz Opteron 6174	696	8.51	5926	184
g20	16 x 2.67 GHz Xeon E5640	128	8.35	1069	167
g22	24 x 2.3 GHz Opteron 6176	408	8.65	3529	528
g23	24 x 2.3 GHz Opteron 6176	384	8.77	3368	–
g24	24 x 2.6GHz Opteron 6238	240	8.30	1993	310
g25	24 x 2.6GHz Opteron 6344	24	9.25	222	29
g26	32 x 2.2GHz Xeon E5-2660	960	9.64	9254	1056
g27	40 x 2.2GHz Xeon E5-2660	1200	9.64	11568	–
s15	8 x 2.60 GHz Xeon	–	–	–	247
s17	8 x 2.60 GHz Xeon	–	–	–	398
s21	8 x 2.40 GHz Xeon	–	–	–	607
Total	–	5376	9.0	53596	3867

Plus Opportunistic resources

Job Scheduling & Multicore

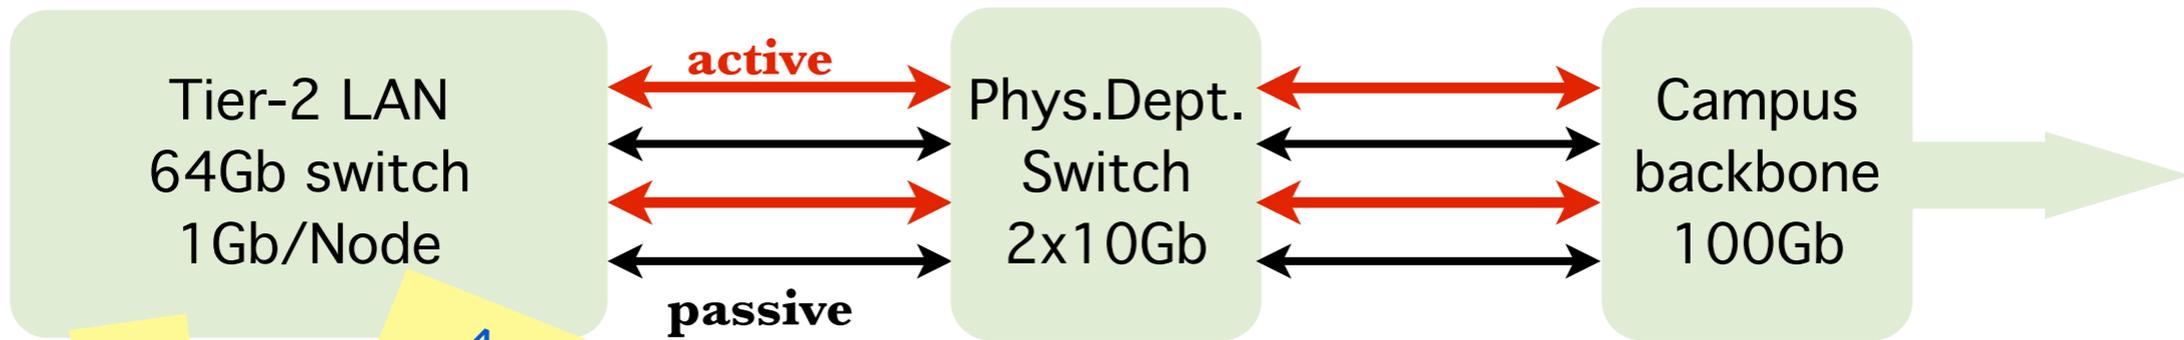
- Currently, ~5k cores plus opportunistic Campus HTCondor Cluster(CHTC) are provided for CMS jobs
- **HTCondor** is being used for job scheduling
- Use of **Xrootd** and **CVMFS** allows us to run jobs anywhere: on opportunistic resources and on OSG
- Opportunist jobs run only when there isn't enough CMS jobs to fill up the queues
- We have a very small fraction of machines that are configured with **Partitionable-Slots** to host multicore jobs
- As of now, we weren't required to convert significant amount of machines to host multicore jobs, but this can be achieved if needed...

Network at U. Wisconsin



Current network status for T2_Wisconsin cluster

Network at U. Wisconsin

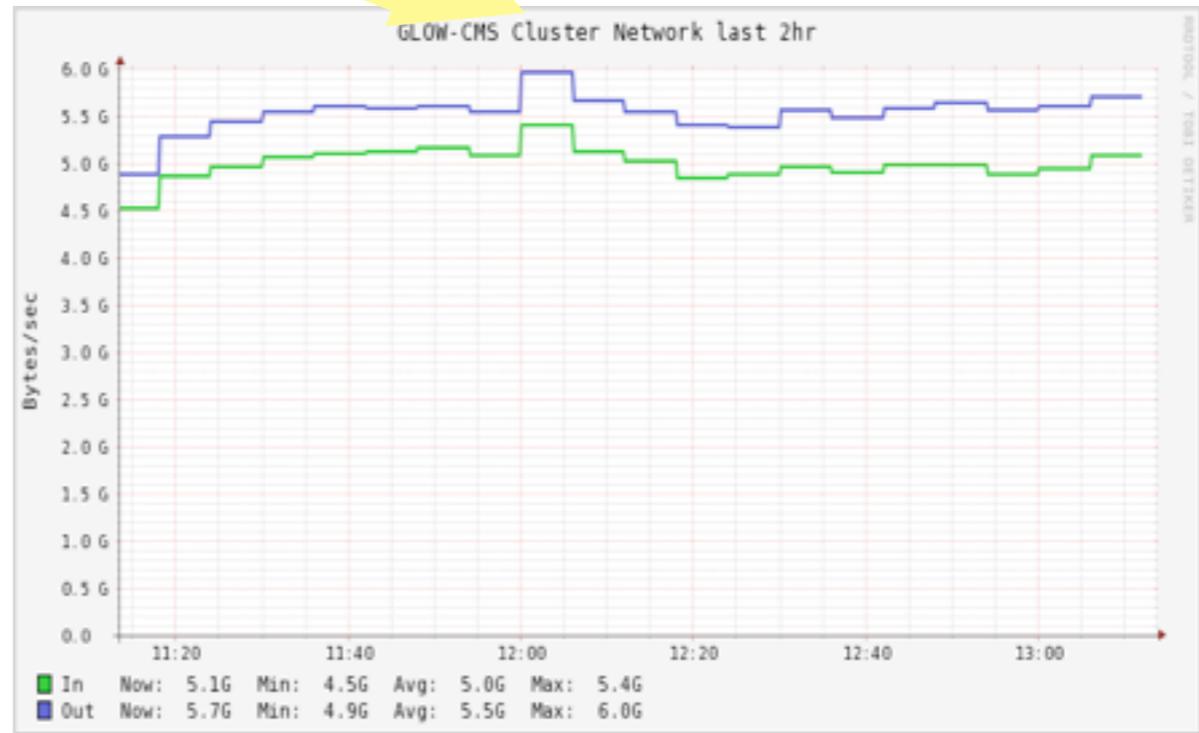


UPGRADE

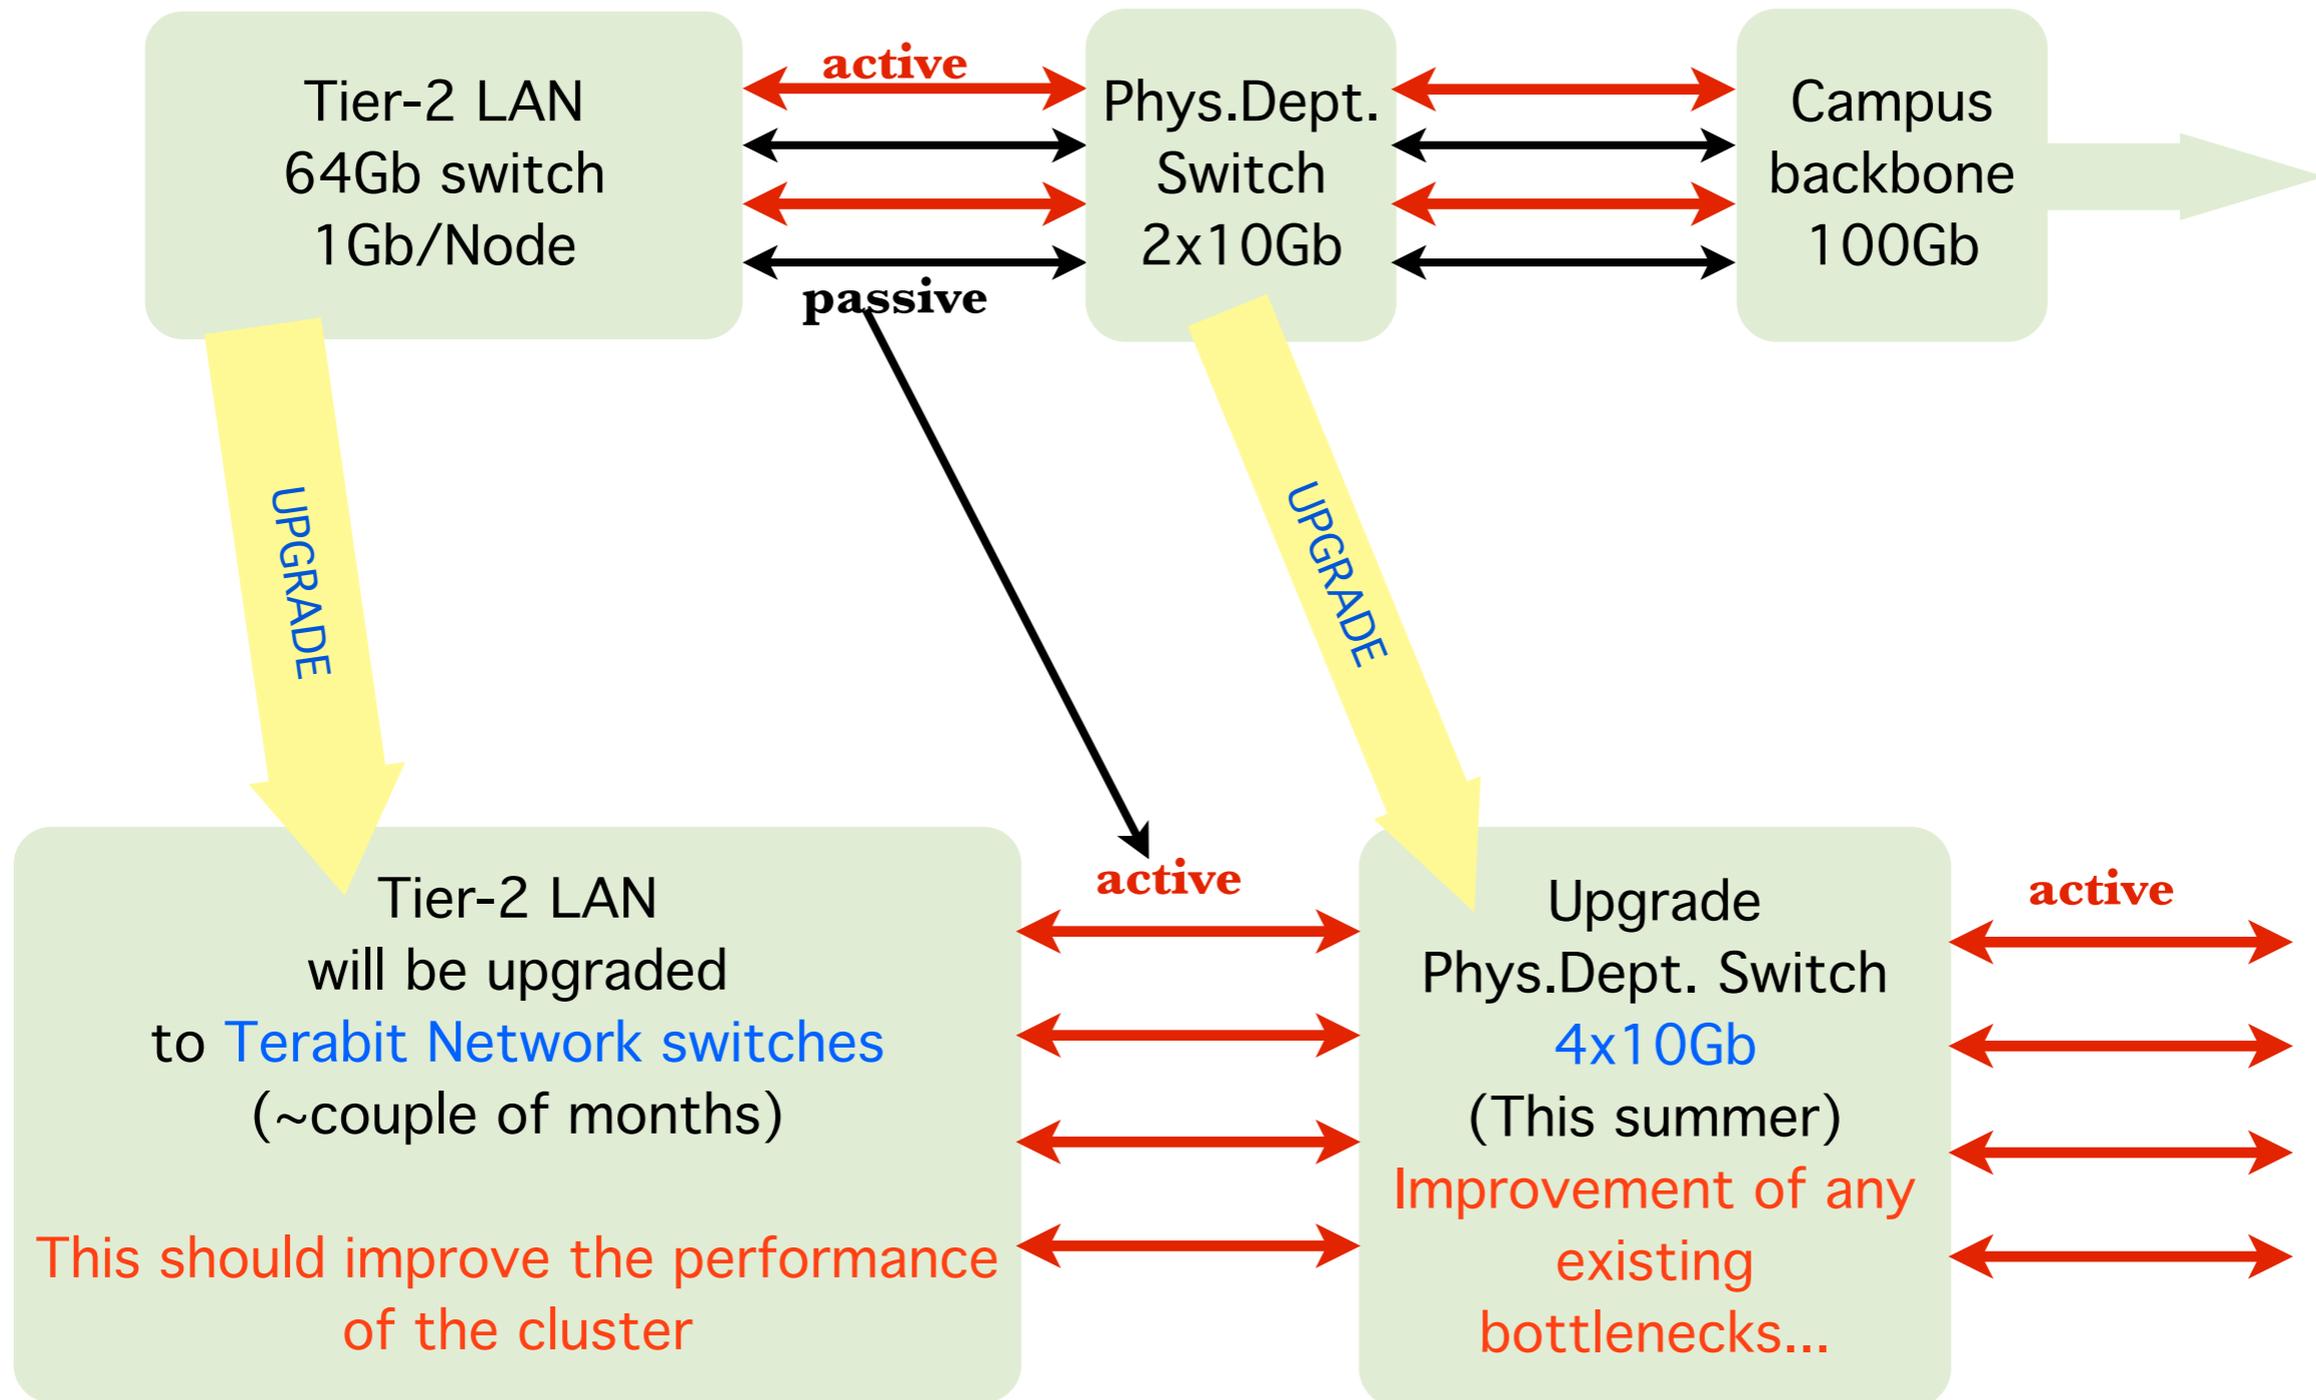
At times internal network is maxed out...

Tier-2 LAN will be upgraded to Terabit Network switches (~couple of months)

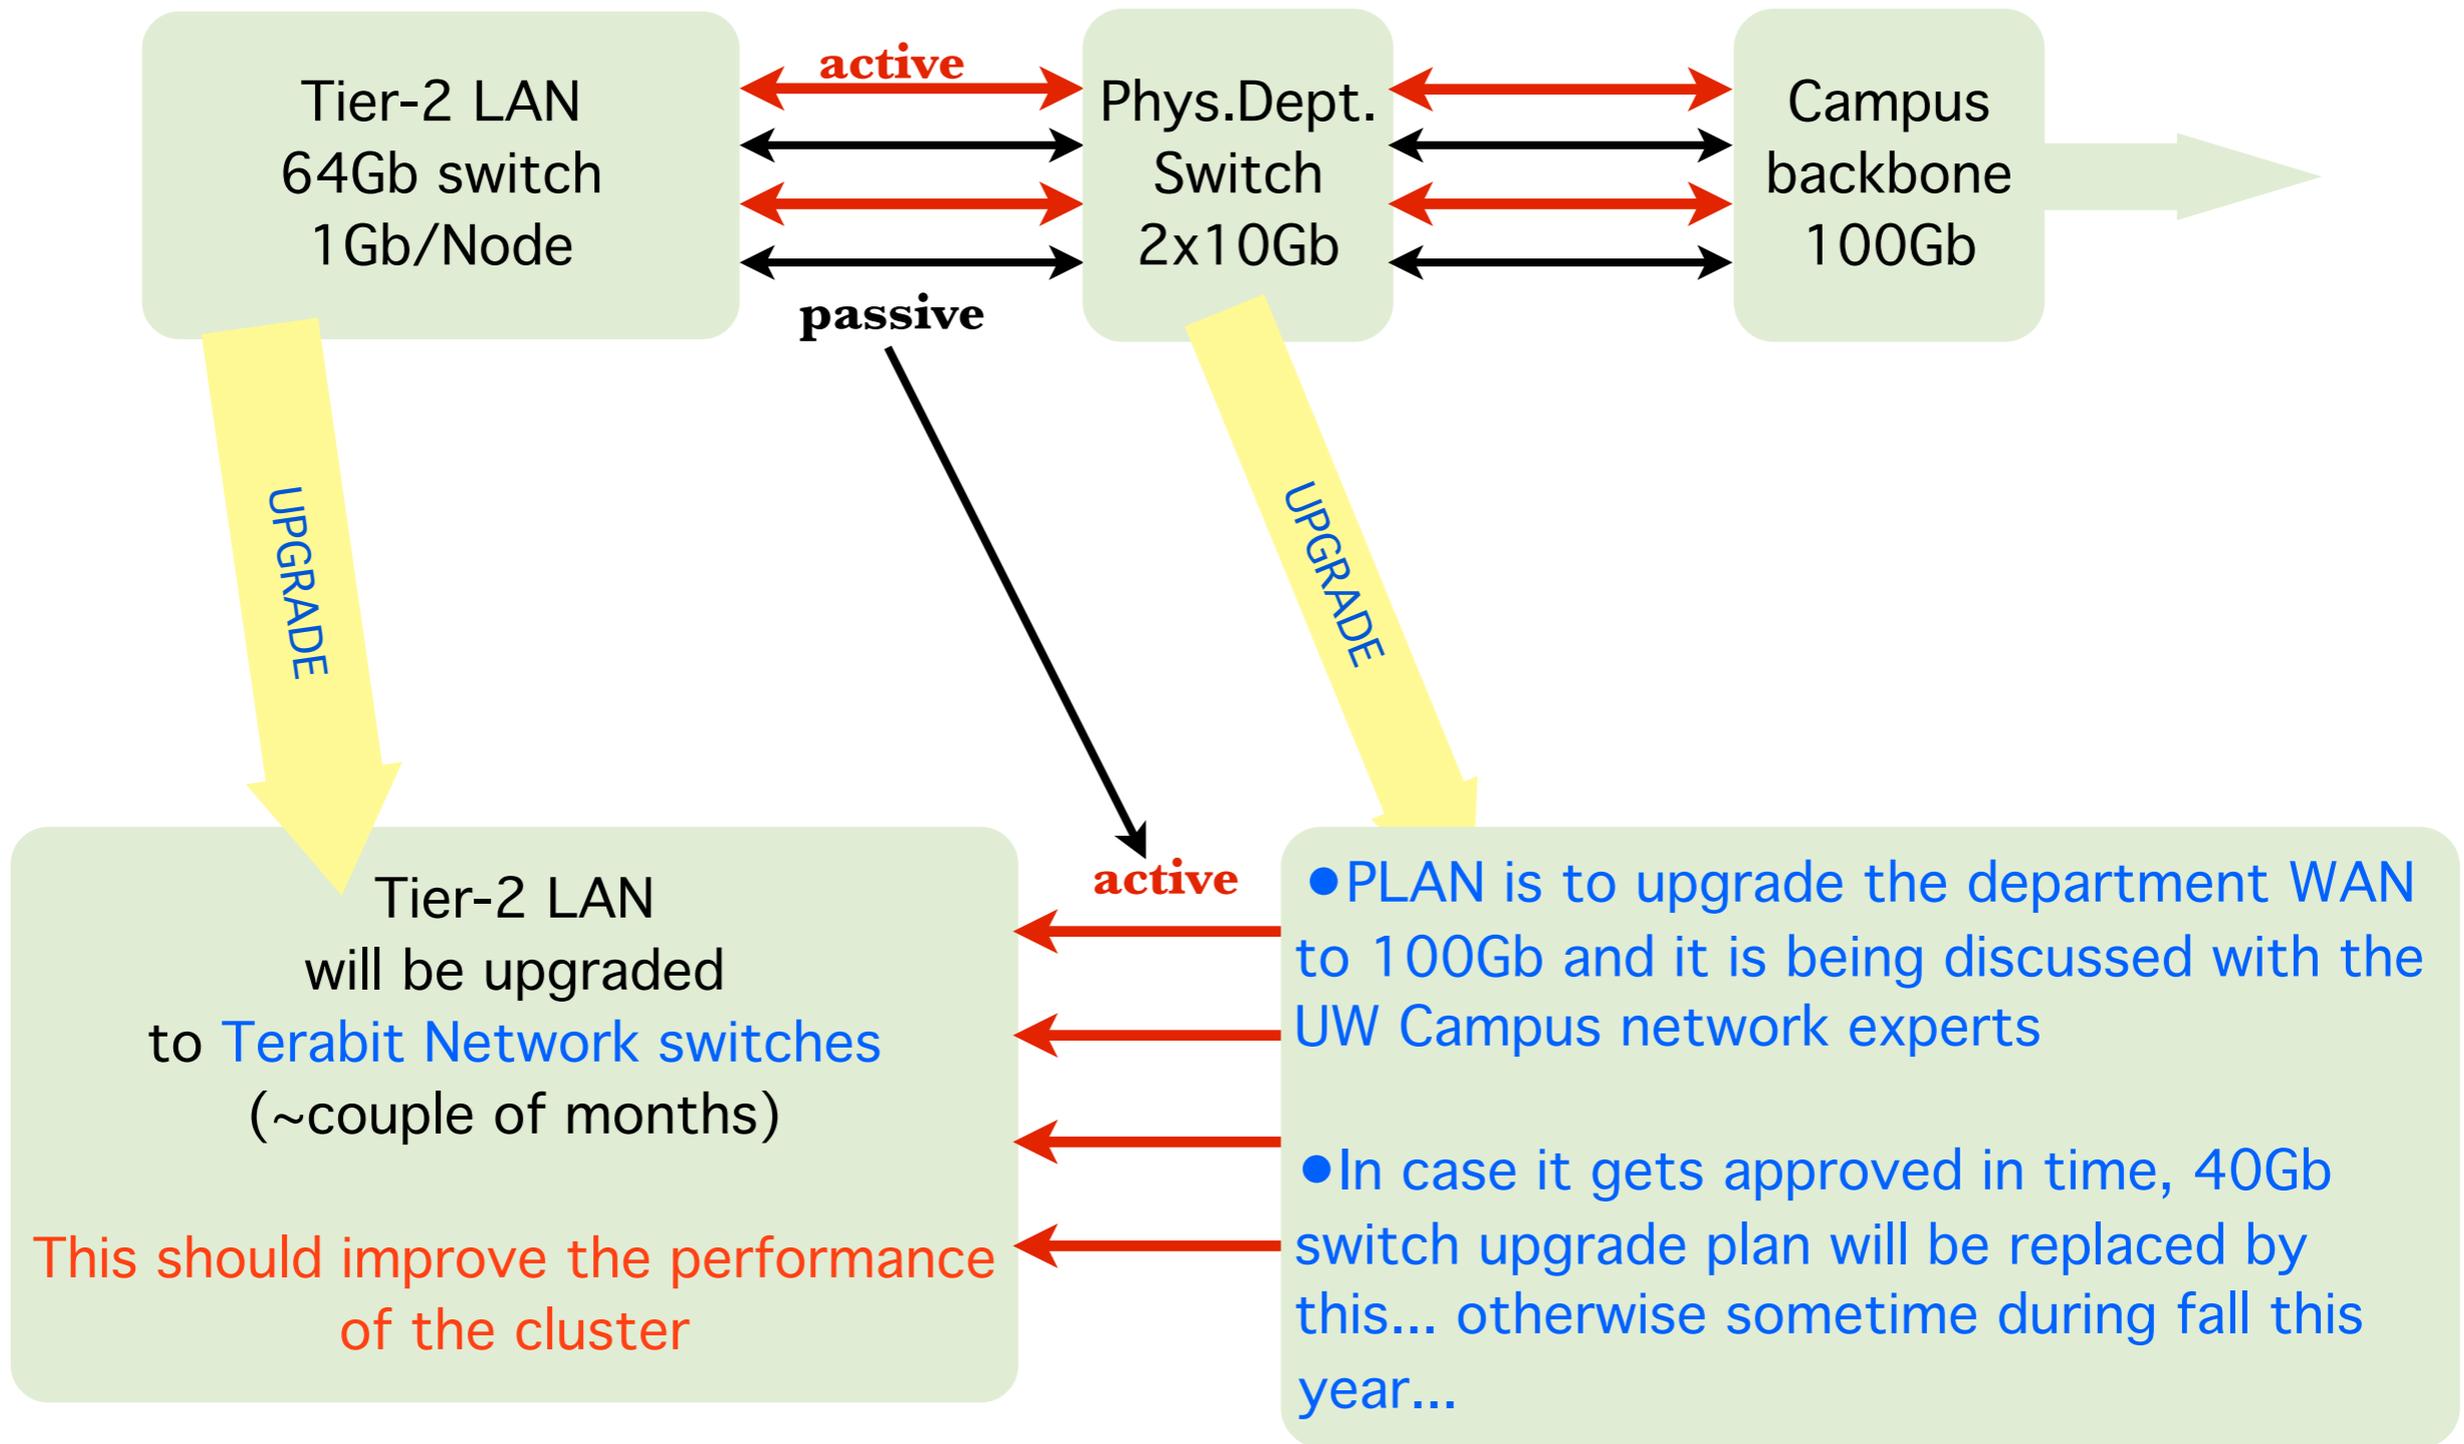
This should improve the performance of the cluster



Network at U. Wisconsin

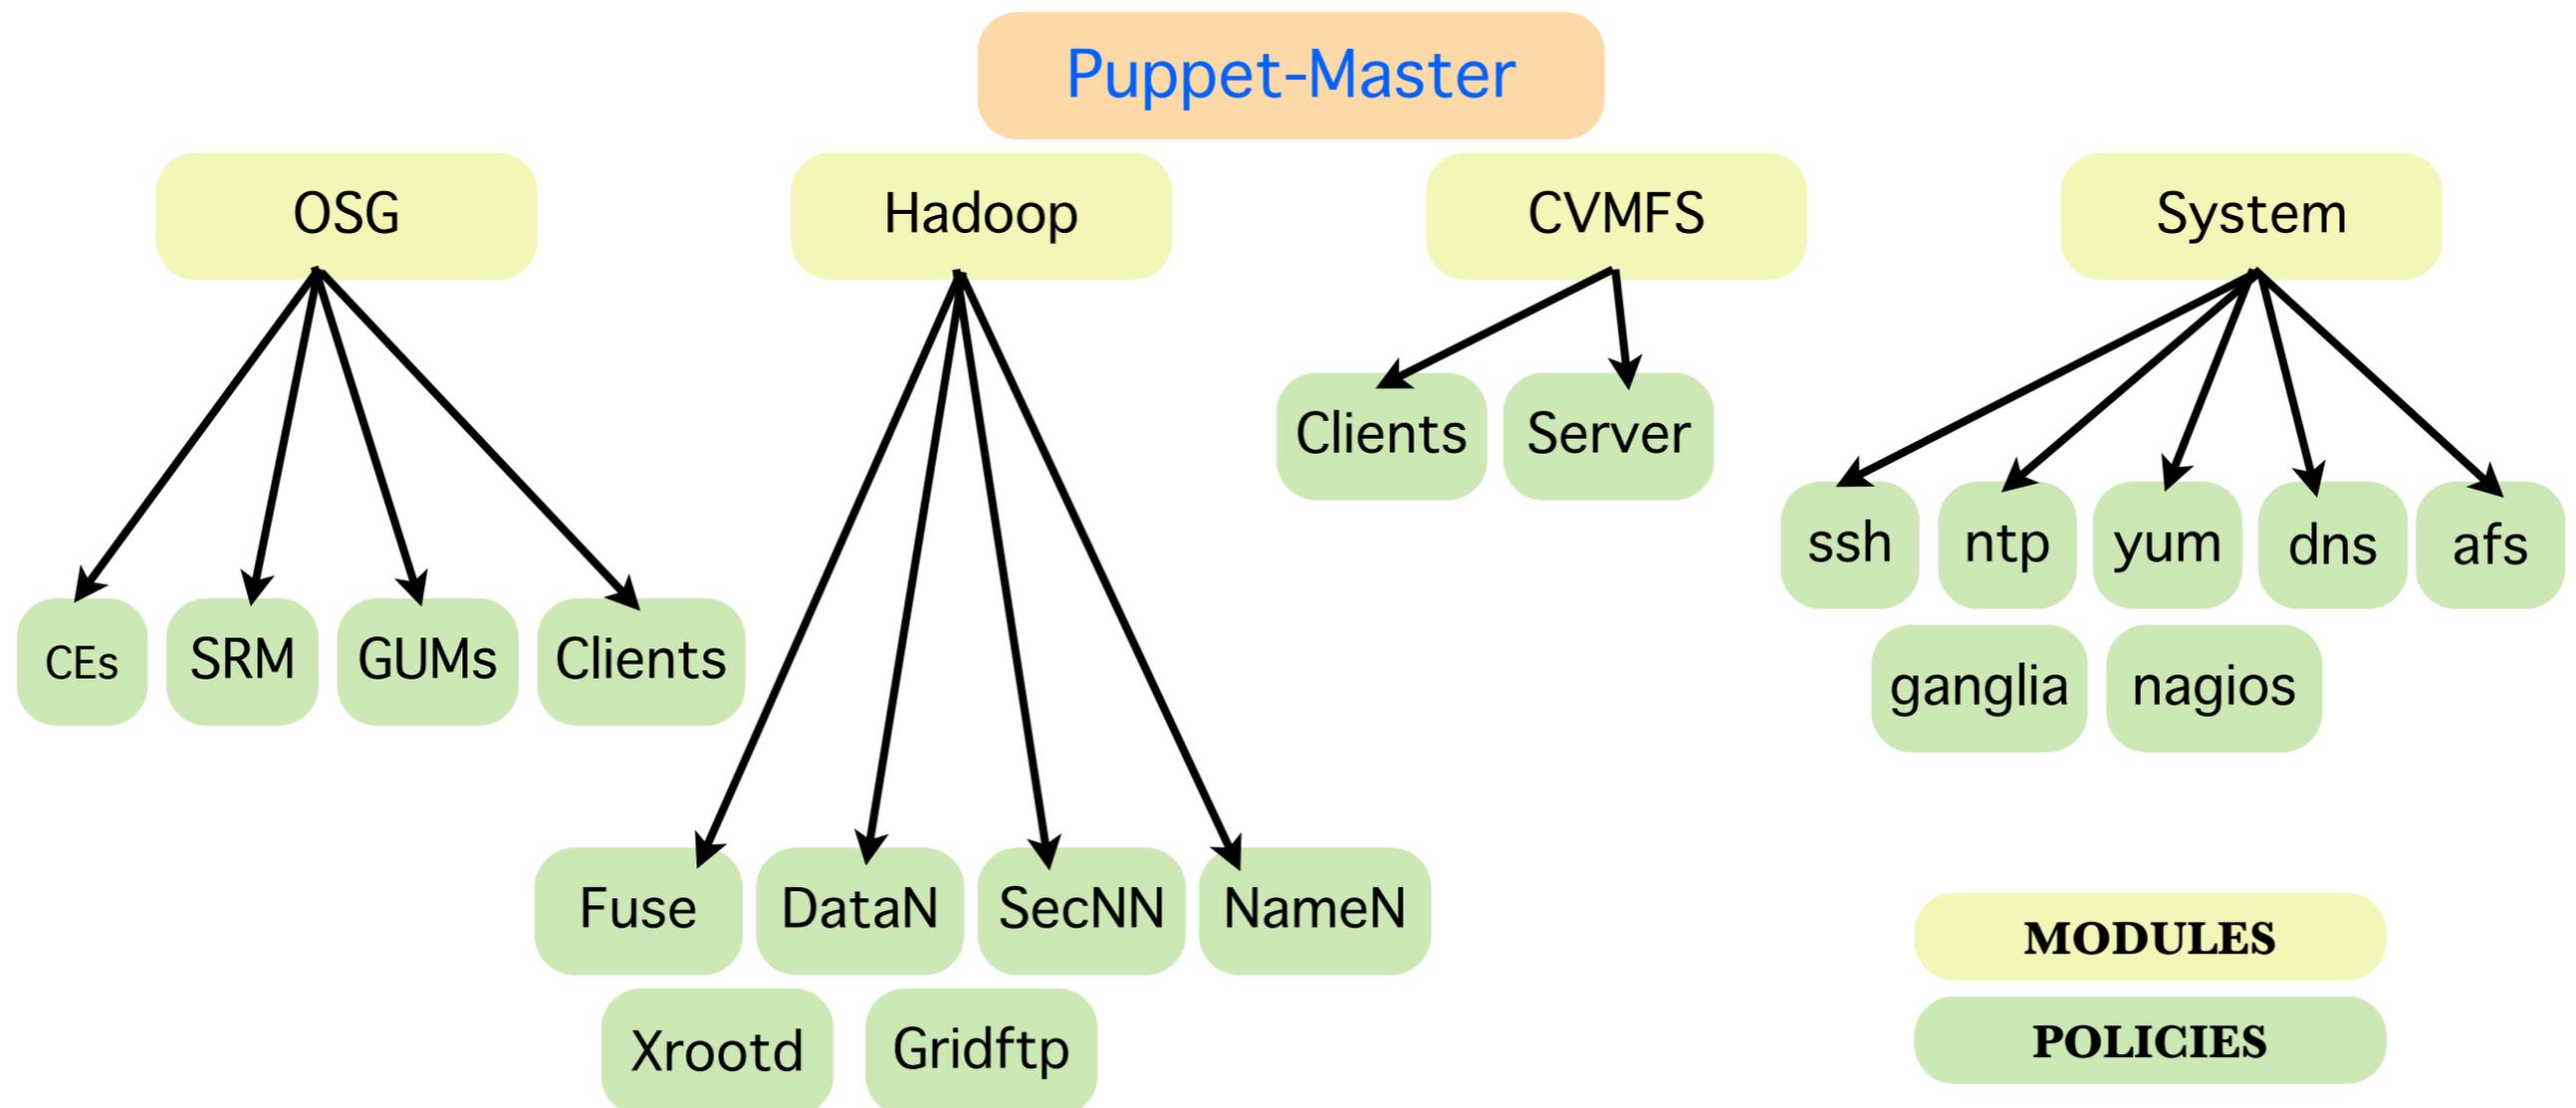


Network at U. Wisconsin



Cluster management, Automation

- Most of the important services (srm, gums, secondaryNN, condor...) are put into virtual machines
 - We use **ganeti** to manage/backup virtual machine cluster
- Cluster is managed by **Puppet (master)**
 - Backend with **passenger** (apache-2 ext. module) that let you run Rail apps and provides better performance for larger number of clients



Others

- We are using IPv4 connections, but IPv6 is available at the switch level
 - Implementation on the OS/application level will happen sometime soon...
- Changes related to FTS3 (debug instance) has been implemented
 - PhedEx transfer hasn't shown signs of trouble so far...
- perfSonar tests are now running and publishing information correctly
- Recent issue with DashBoard monitoring showing much more running jobs at Wisconsin than actual (**Fixed**):
 - Timeout bug at central monitoring side
 - On the T2 side, some jobs never reported back to DB
- After major upgrade last month ([osg-32](#), [hadoop-2](#), [condor-8](#), [SL6](#) for [CEs](#), [GUMS](#), [SRMs](#)), no surprises, T2 operation is smooth without any problem...