

Tier2 Management At Manchester

Colin Morey
(colin.morey@manchester.ac.uk)

Maximum Efficiency with Minimum Effort



Copyright Colin Morey

The Nodes

- 2x2.8Ghz Xeons
- 2x250GB hdd
- 4GB ram
- 2x1GigE



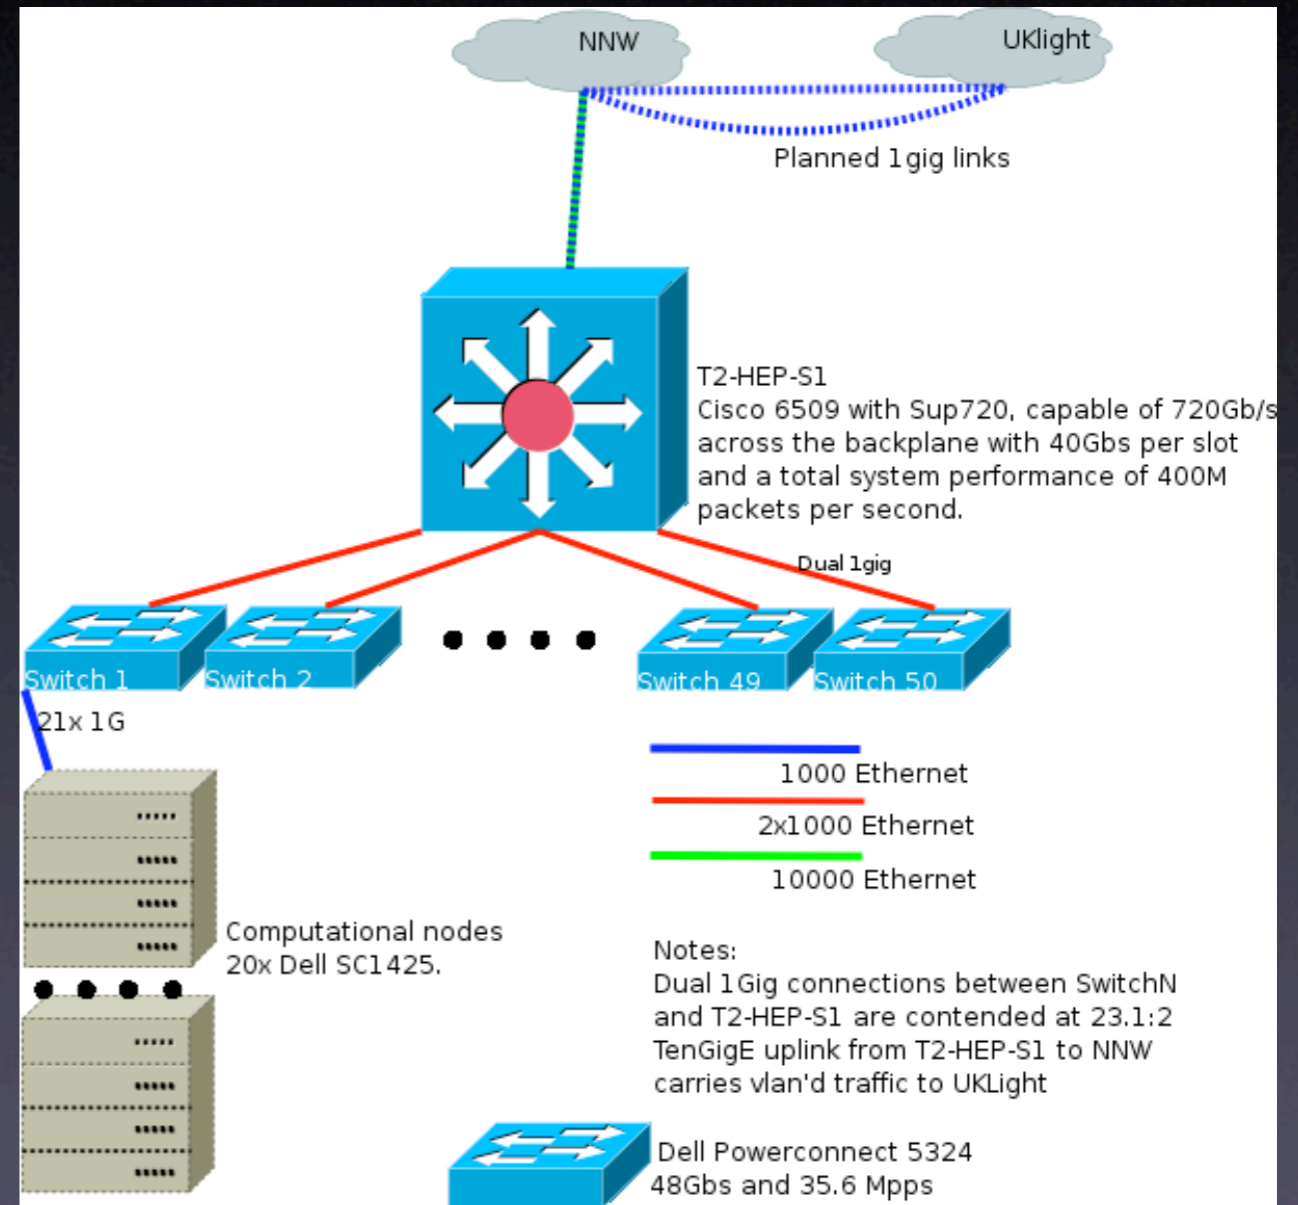
The Farm

- 1000 Nodes
- 2000 x2.8Ghz Xeons
- 1/2 Petabyte storage
- 4TB ram



The Network

- 51 Switches
- 50 Dell 5324
- 1 Cisco 6509
- 1250 Gbit links



Automation

- Build/install - Kickstart
- Configuration - Cfengine
- Administration - ssh
- Updates - yum

Kickstart

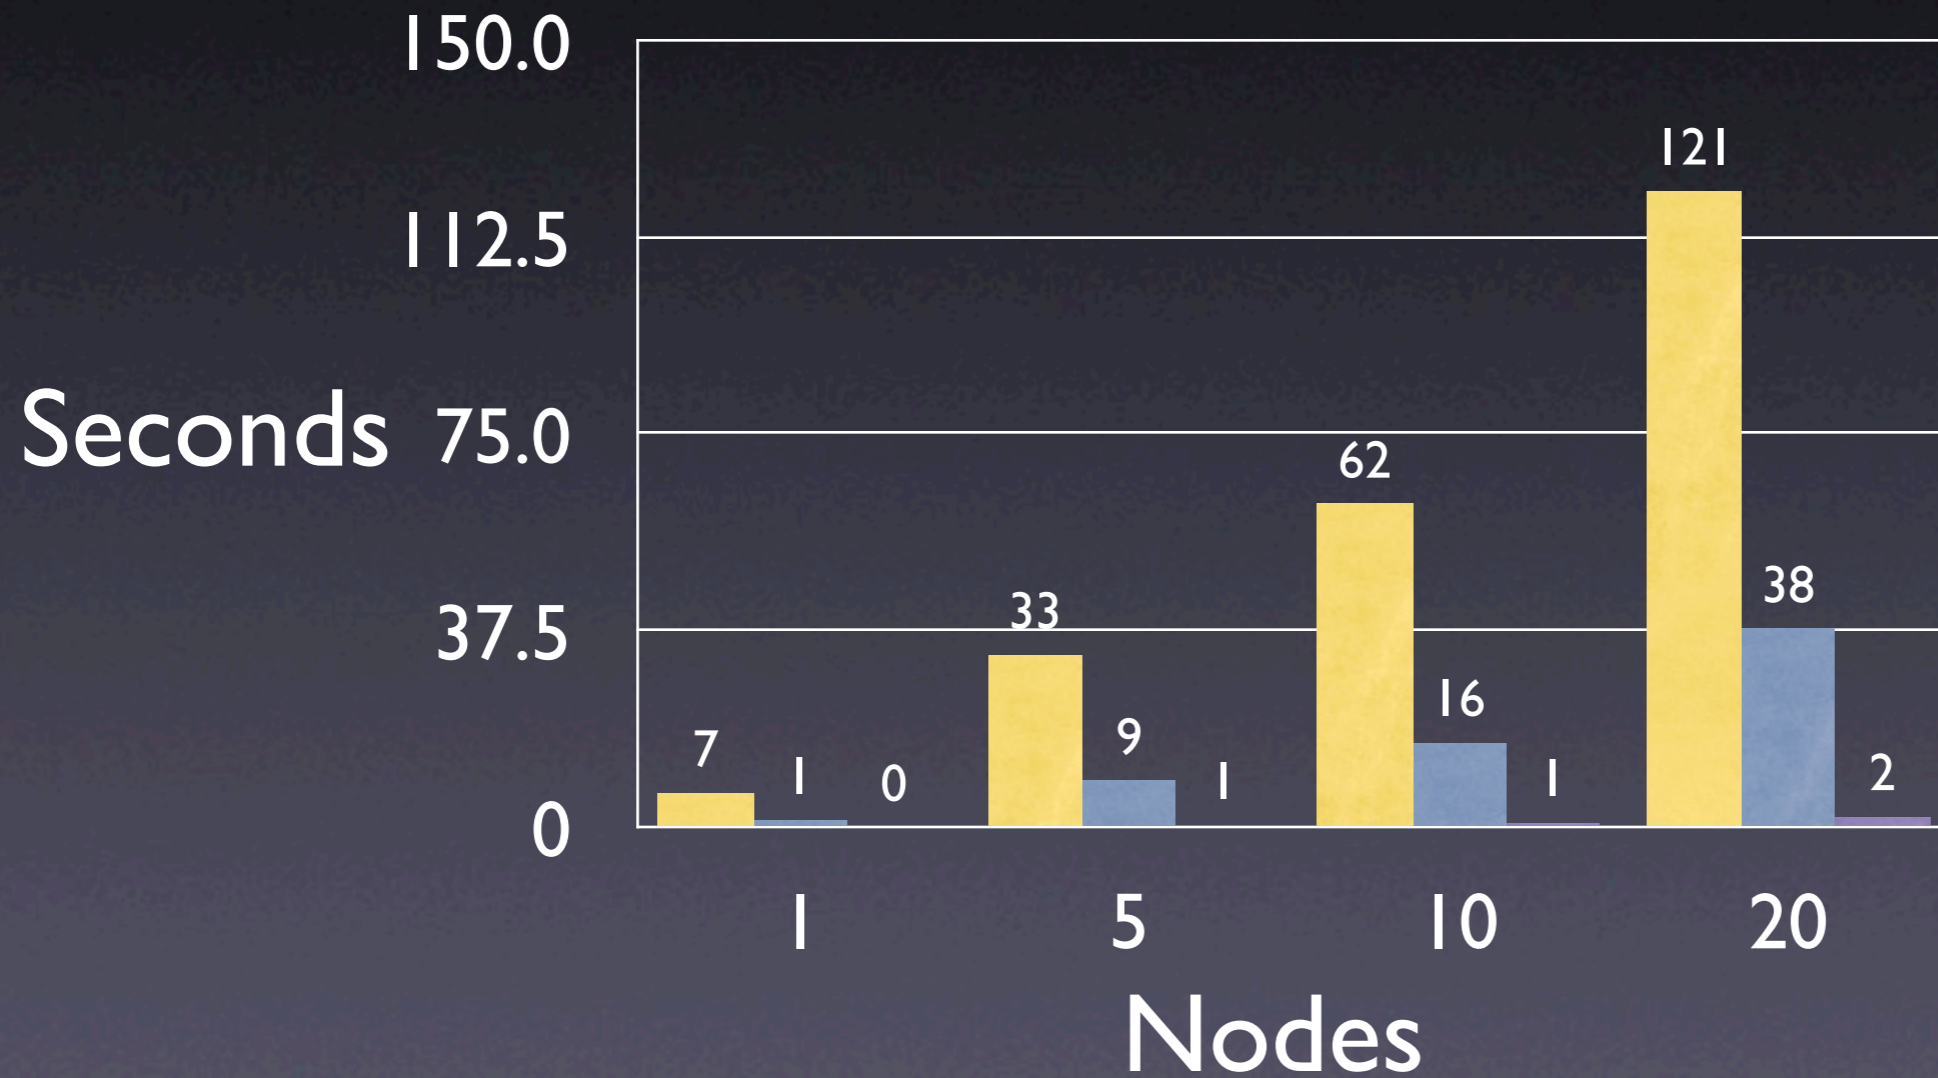
- 2 DHCP servers provide leases for the whole network, including the powerbars, and the LOM interfaces on the nodes.
- 4 Servers in total, connected via 2xGigE, All server are identical and currently kept in sync via rsync from a designated 'master',
- With 3 servers, we were able to supply data at an aggregate rate of 600MB/s

Cfengine

- Early Stages of deployment.
- Very, Very powerful.
- Integration with Kickstart.

OpenSSH

- yum info - with key
- uptime - no key
- uptime - with key



Yum

- Easy to use package management
- Cryptographic checking of packages
- Supports multiple repositories
- Supports autonomous execution

Monitoring

- Network Usage - ntop, mrtg, rrdtool
- Environment - mrtg, rrdtool
- Service status - nagios

NTop

- Uses a concept of traffic flows
- Act as a netflow target
- Clear and concise display of detailed information

MRTG

Pros

- Reliable
- Well known
- Mature
- Capable
- Not restricted to SNMP targets

MRTG

Cons


- Scalability issues
- Integer only
- Unable to record samples at a resolution < 1 min
- Only 2 items per target

RRDTool

- Concerned only with the Storage and displaying of data.
- Backend storage and front end visualisation of data for MRTG.
- Multiple interfaces for inserting data into datasets.
- RPN can be confusing for the beginner, but it's capable of producing good results.

RRDTool

- With a bit of Perl and bash (and coffee), we can create graphs like these.



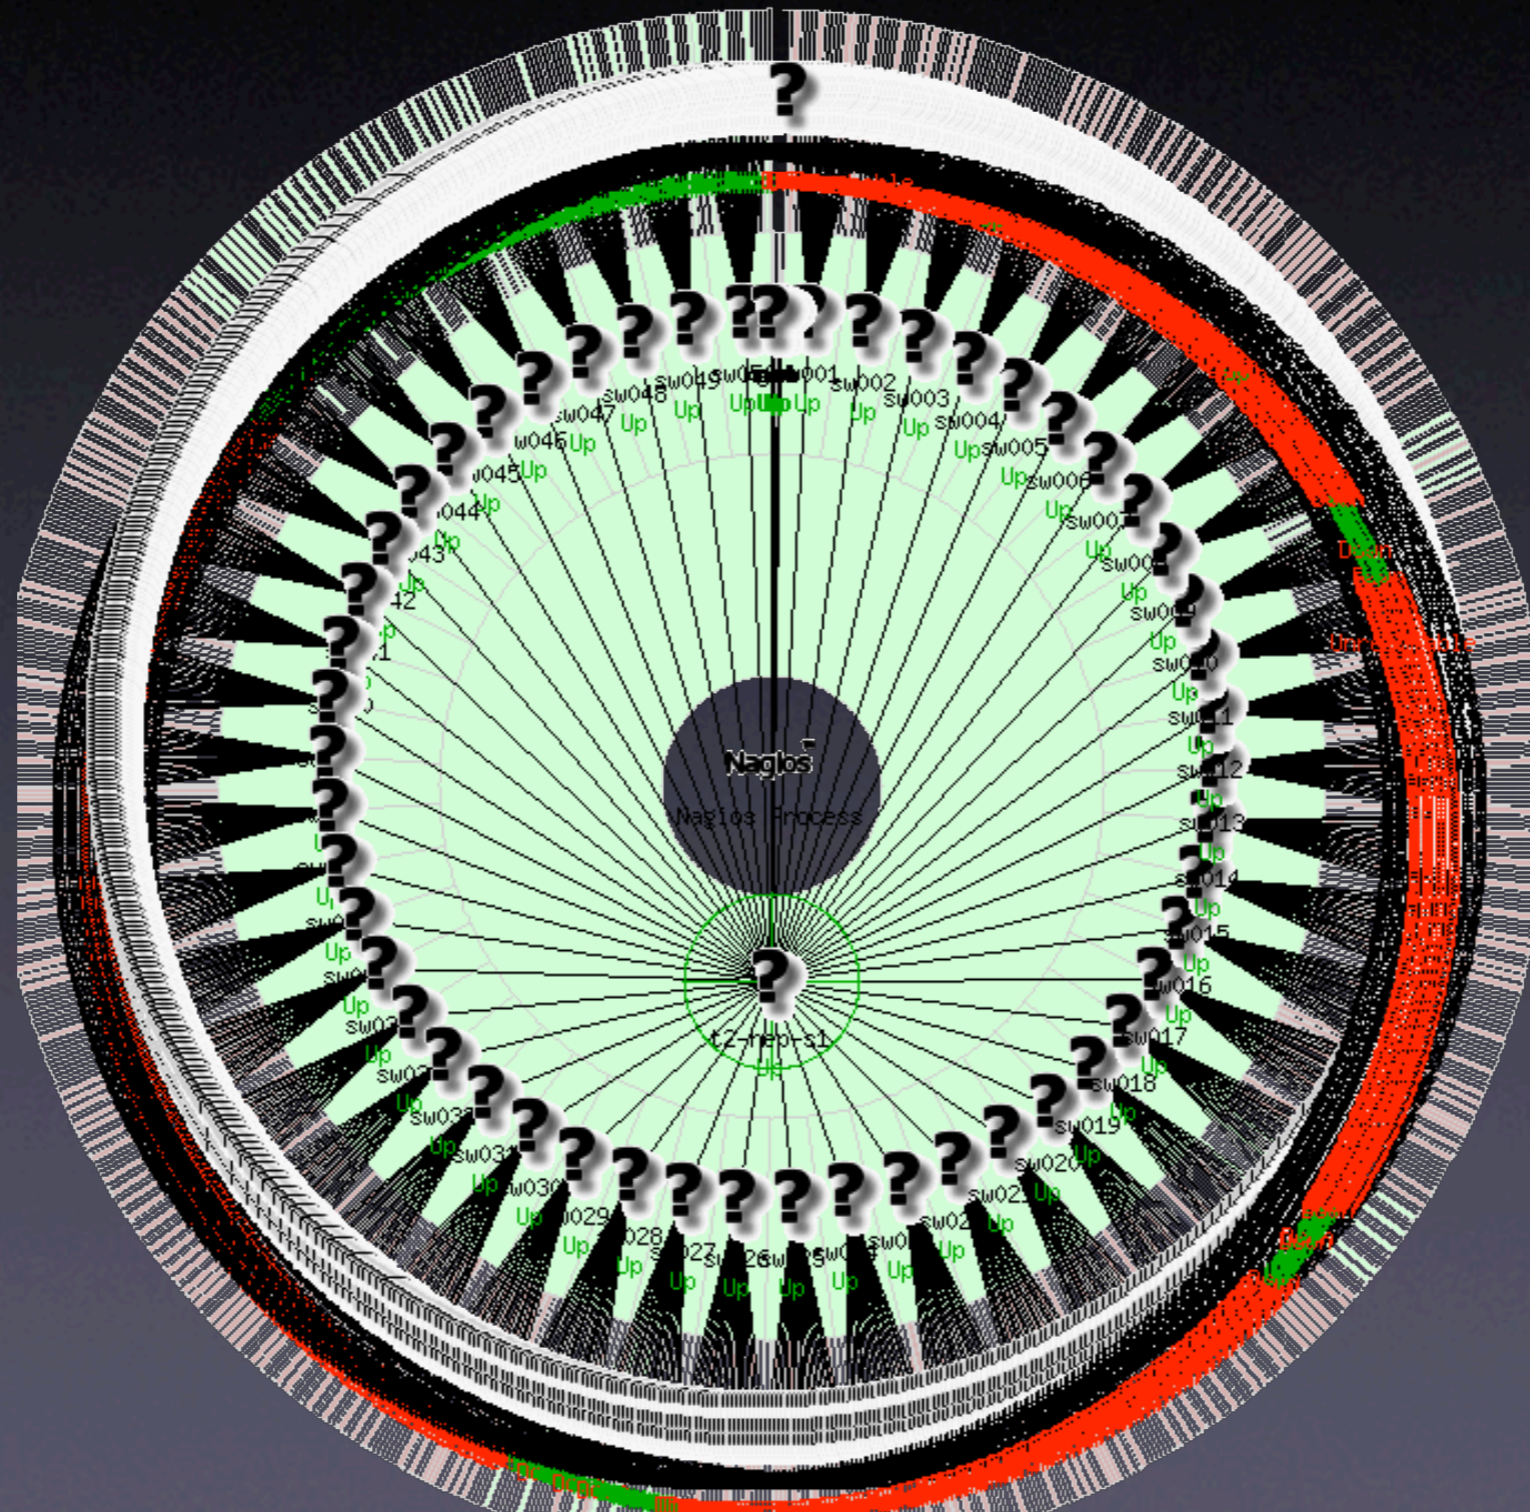
Nagios

- Service based monitoring,
- New revision (version 2 upwards) introduces classes and dependancies
- easily extensible

Nagios

- \leq 1 minute: 331 (15.3%)
- \leq 5 minutes: 1984 (91.4%)
- \leq 15 minutes: 2170 (100.0%)
- \leq 1 hour: 2170 (100.0%)
- Since program start: 2170 (100.0%)

Nagios



Node Management

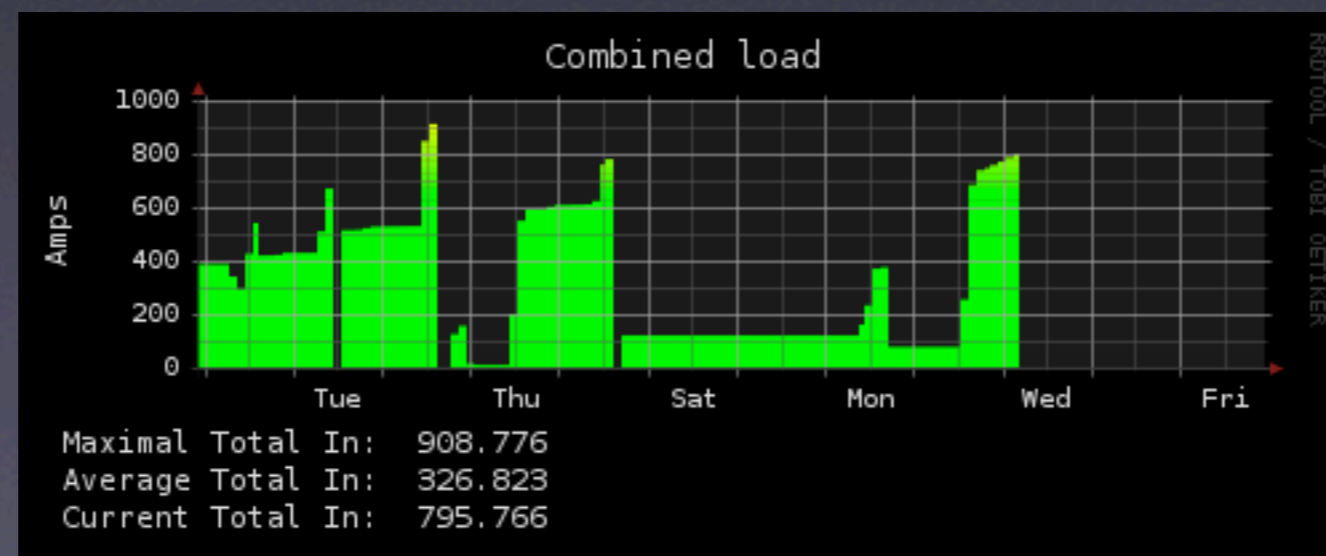
- All Nodes have a separate 'ipmi' address that provides LOM
- Remote Chassis Control, power on/off/cycle
- Reports Chassis status, temps, fan speed, onboard event log.

Facilities Management

- Power bars in each Node rack
- Exports current draw per bank of sockets via SNMP
- Provides remote control of Power to each node.

Facilities Management

- Power bars in each Node rack
- Exports current draw per bank of sockets via SNMP
- Provides remote control of Power to each node.



fin.



resources

- mrtg <http://www.mrtg.org/>
- rrdtool <http://www.rrdtool.org>
- cfengine <http://www.cfengine.org>
- ntop <http://www.ntop.org>
- nagios <http://www.nagios.org>
- OpenSSH <http://www.openssh.com>
- kickstart <http://www.redhat.com/docs/manuals/linux/RHL-9-Manual/custom-guide/part-install-info.html>
- yum <http://linux.duke.edu/projects/yum/>
- <http://linux.duke.edu/projects/yum/>