LINX News

Mike Hughes

mike@linx.net

UKNOF 4

Traffic nearly @ 100G

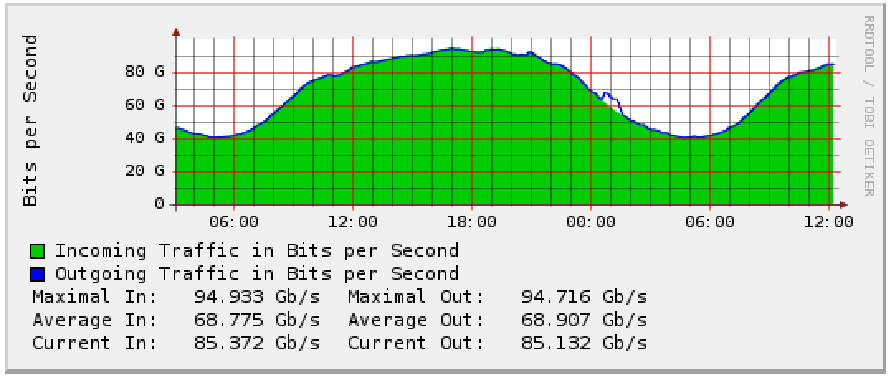

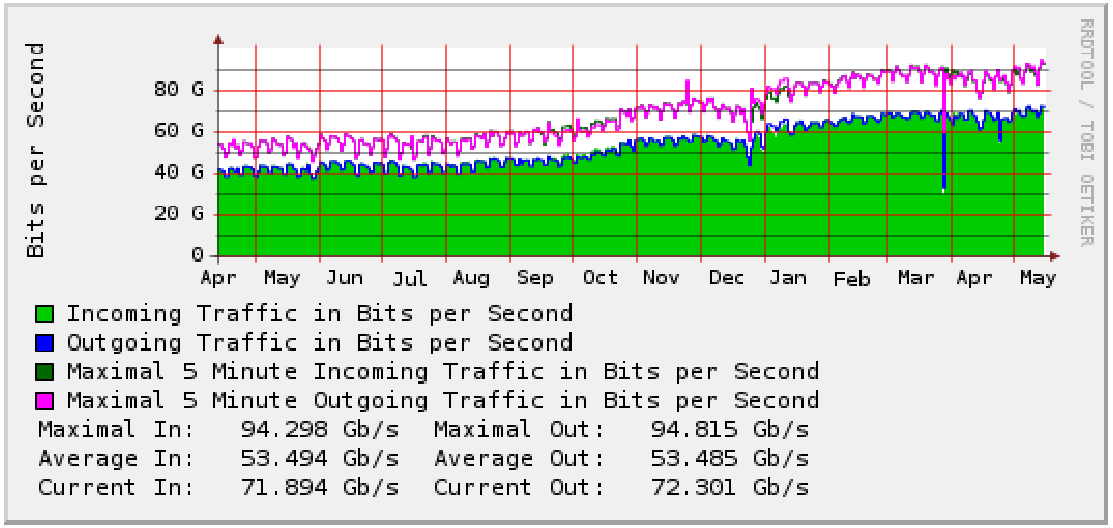

Membership

- Continues to grow – over 200 members
- Our first member from the Gulf
  - Etisalat (the UAE incumbent)
- No coverage of South America yet
  - Cost of transmission?
- Still a strong UK contingent

Network

- Continues to develop in response to member needs
  - Increasing number of 10G ports
- Upgrades are continuing
  - Completing deployment of Extreme upgrade
  - Deploying new Foundry hardware as well
  - Increasing metro backbone in Docklands
    - ISL Widening and multi-ring design

Foundry Network Plan

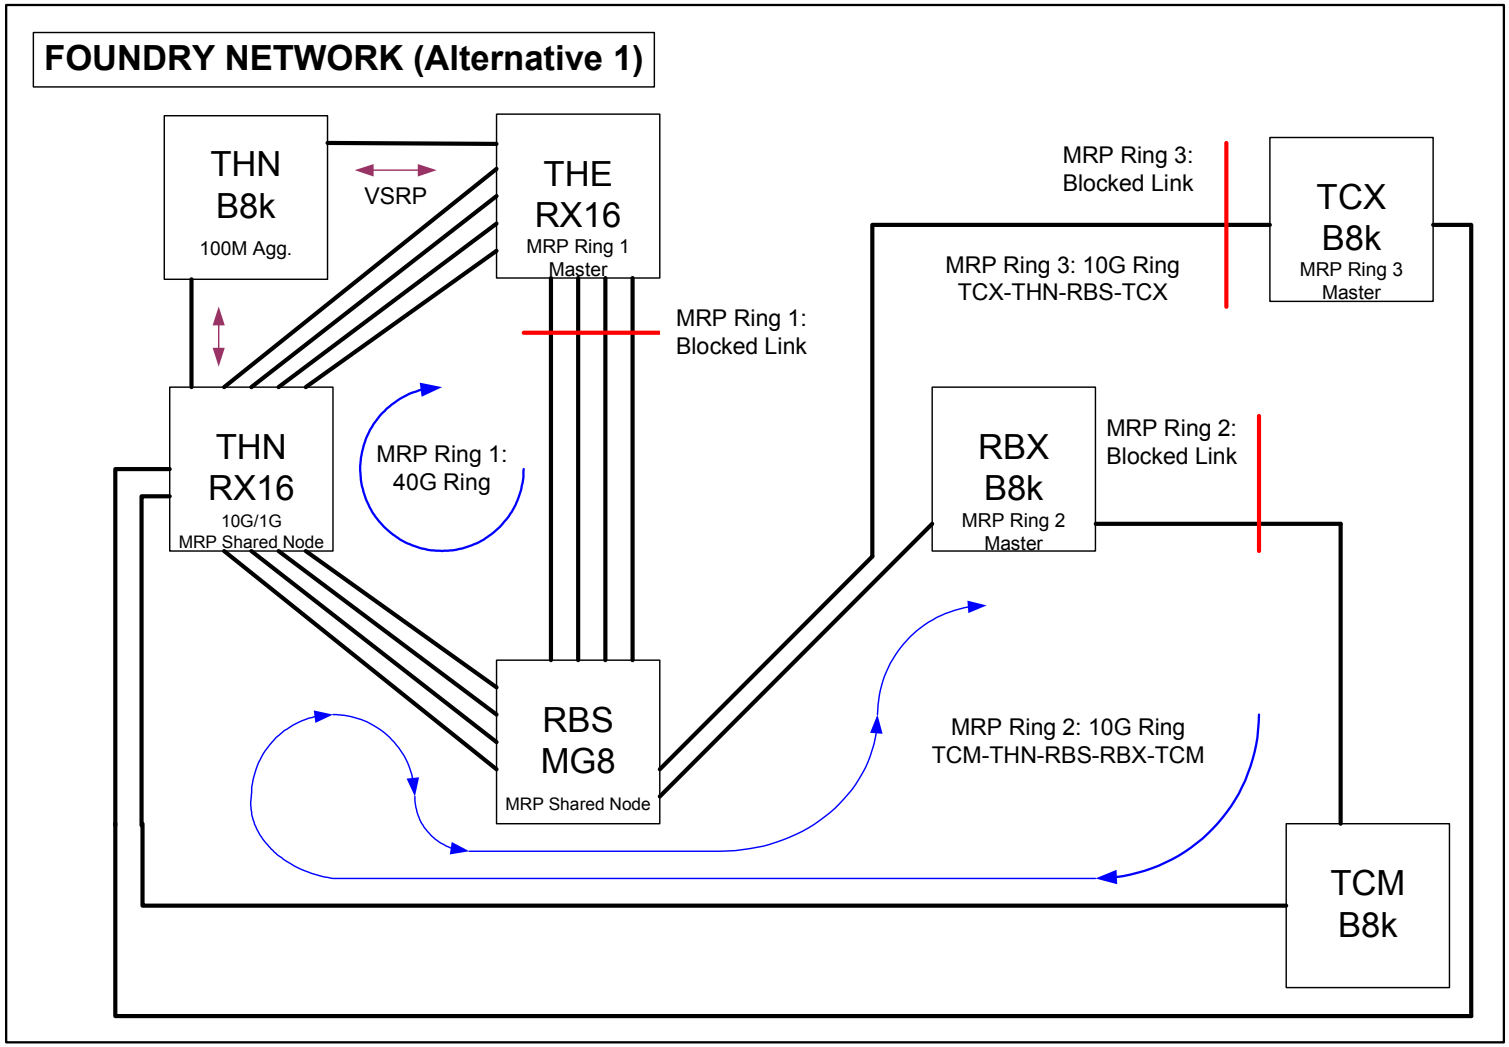

Extreme Network Plan

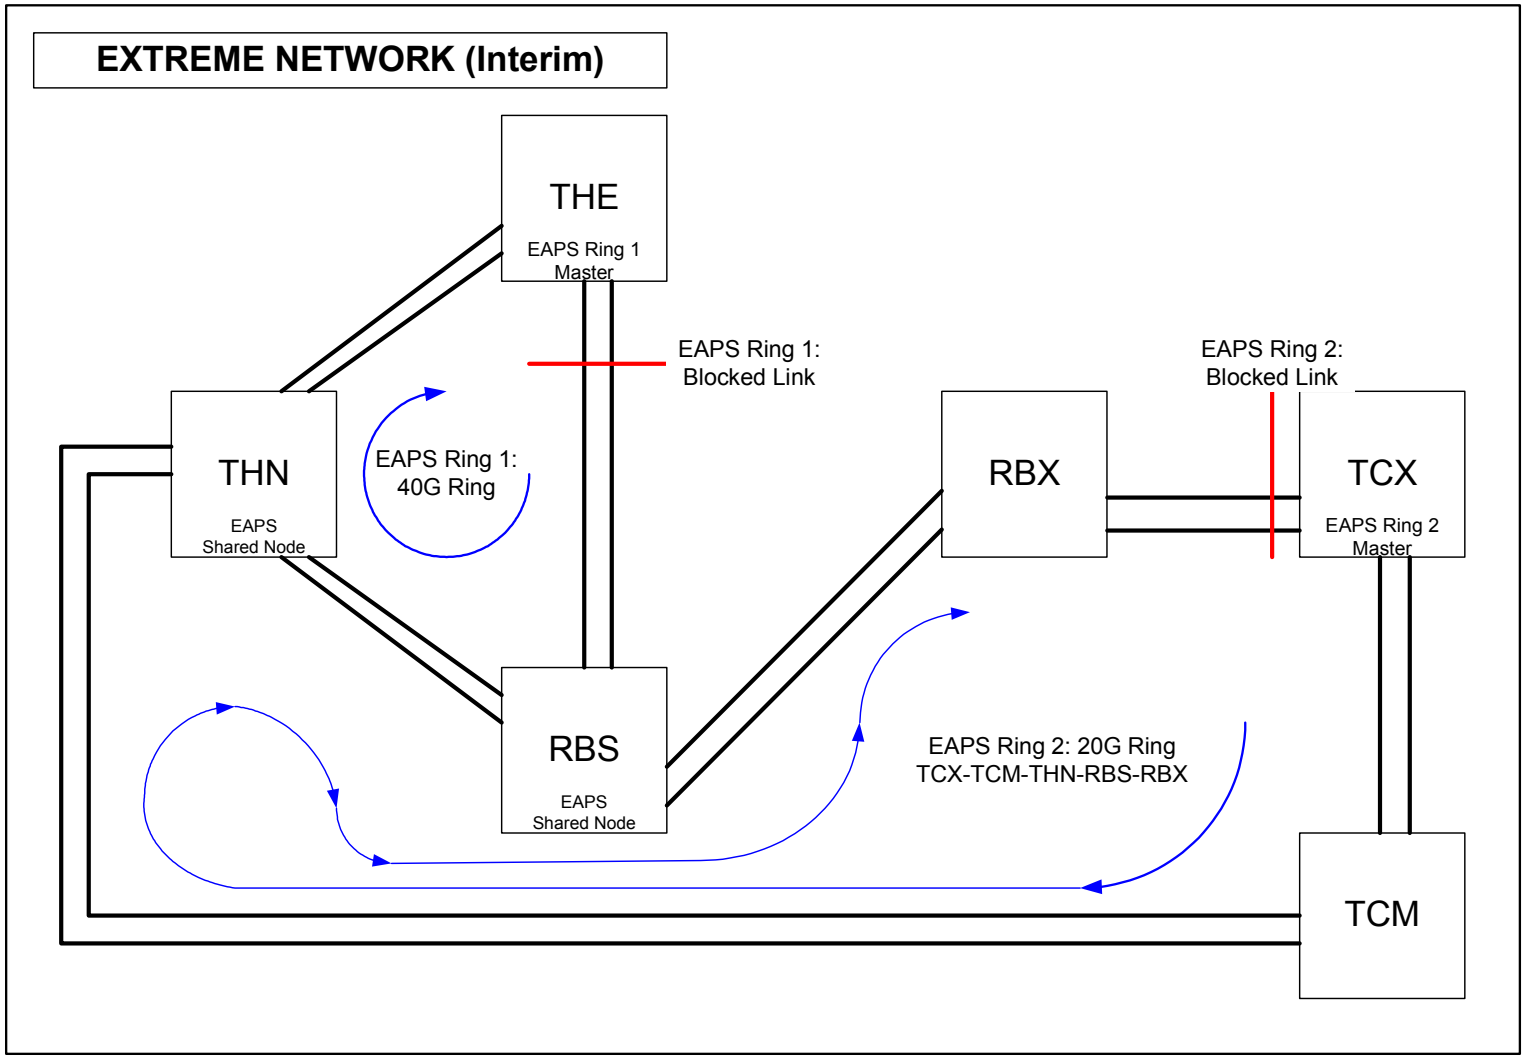

Sflow Tools

- We're enabling Sflow data collection on the network
- Takes samples of traffic from each port (1 in 2048 frames)
  - To get a "flavour" of the traffic on each port
- This is then propagated via management network to a collector server
  - Which is a challenge in itself

Sflow Server

- 2 x 3.2GHz Xeon CPU
- 4GB RAM
- Disks
  - Mirrored 500GB OS disk (+ for data overflow)
  - 2TB RAID5 (data)
  - online hot spare (usable to extend data volume)
  - All same disk size – KISS principle
- Separate NICs for sFlow data

Sflow: Issues

- Initial tests/coding workable
  - Data from a couple of switches
  - ~10,000 mac:mac pairs
- Scaling problems when full foundry-net sFlow-enabled
  - ~35000 mac:mac pairs
  - pmacct memory-tables working fine
  - Writing to ~35000 RRD files: time-related issues

Sflow: Disk I/O Issues

- System was constrained by disk I/O
- Synchronous writes to filesystem
- RAID5 used – faster writing than RAID0 (mirror) [obviously]
- Filesystem testing (xfs chosen)
- Filesystem tuning
- Less granular sampling (temporarily)

Sflow: Other Considerations

- Writing to RAM (tmpfs) & periodic sync'ing to disk
  - Faster
  - (but) large potential for data-loss
- Writing to database rather than individual RRD files
  - fewer random disk writes, read from RAM
  - but requires non-RRD graphing tools
- Extreme LAN (226-net) will add further overhead
- Since been offered access to other existing code

Sflow: Peering Stats Beta

Traffic flows from port

Port
AS Number 13037
IP Address 195.66.224.158
MAC Address 00:90:69:c8:74:5d

Thu, 18 May 2006 12:22:37 BST

	Day	Week
20547 195.66.224.72 00:08:02:ef:e9:bd 20/7/2 uksolutions	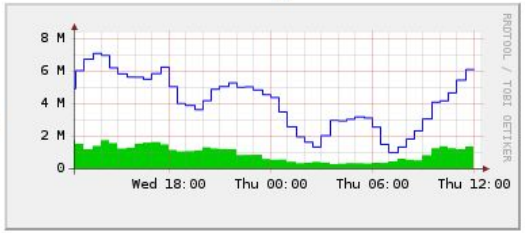	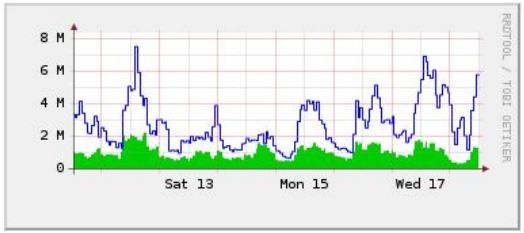
2529 195.66.224.12 00:0b:60:2c:c0:ab 2/3/39 demon	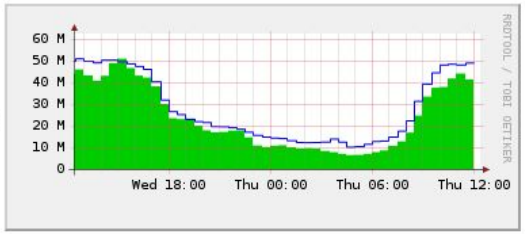	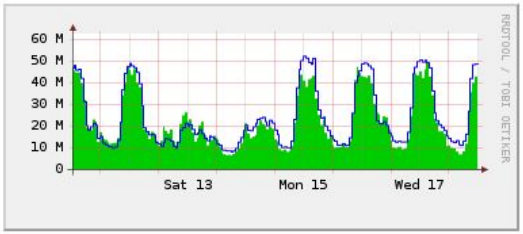
15766 195.66.224.221 00:13:19:fb:66:1b 3/7/11 domicil	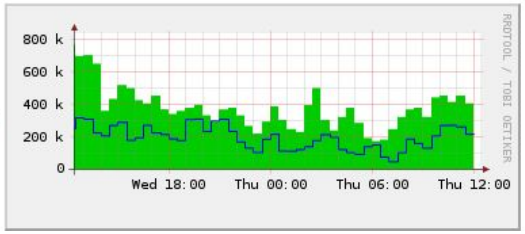	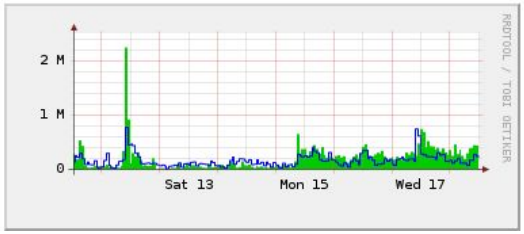

Other Tools

- Route Aggregation BCP checker tool
  - Checks your ASN's aggregation against the LINX BCP
  - Slightly more digestible reporting than CIDR Report, for example
- "Peering Professional" tool, in association with Renesys
  - Provides "what-if"-type analysis
  - Aimed at peering co-ordinators and capacity planners

Sflow: Other Applications

- Currently working with Richard Clayton (Cambridge University/THUS) on the SpamHINTS project
  - Providing Sflow data samples of TCP header info for mail traffic
    - No TCP payload, no actual SMTP headers or content
  - Characterising Spam by TCP “fingerprinting”
  - Plan to move to a “live” filtered stream

Governance

- LINX 53 General Meeting earlier this week
- Members voted to drop the “traffic-based” element of the fees
  - A return to port-based charging
- Board elections
  - Patrick Gilmore replaces Nigel Titley after 12 years’ service to LINX

Questions?

mike@linx.net

http://www.linx.net/