



Advanced Monitoring Techniques for a Large Scale Data Processing Network



Brian Martin

CERN

On behalf of

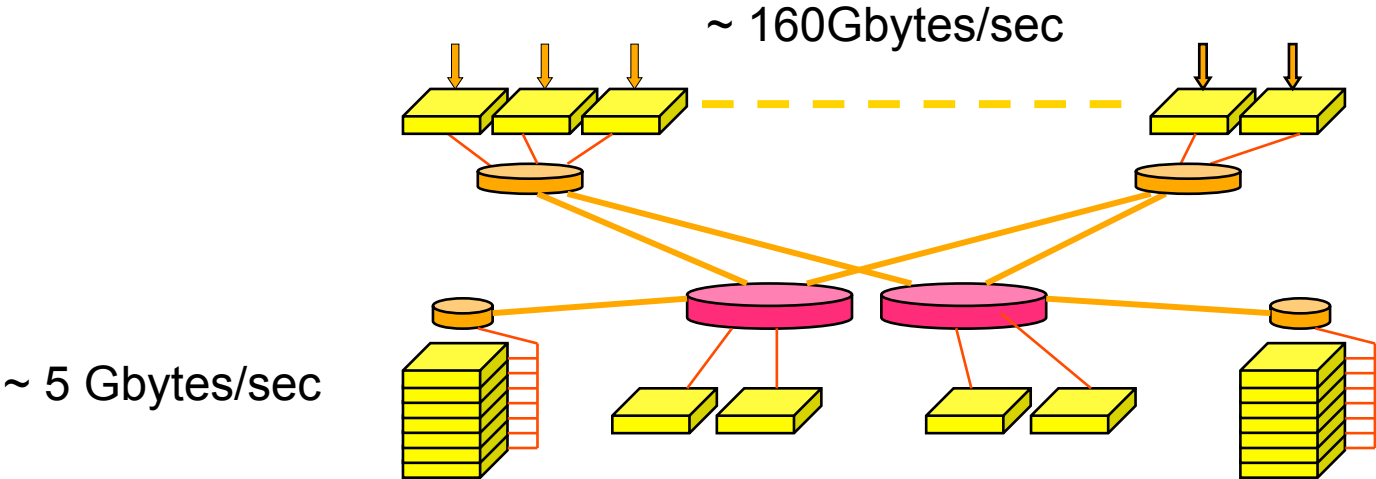
A. Al-Shabibi, S.M. Batraneanu,
D. Savu, R. Sjøen, S.N. Stancu
of the ATLAS collaboration

The ATLAS Experiment




- 3500 People
- 37 countries
- 7000 ton detector
- 80 Terabytes/s data output


Network Infrastructure




Network Infrastructure



Network Infrastructure



Network Infrastructure



8500 ports

Demanding customers!

- Network dimensioned to meet 'requirements'
- Maximum average link occupancy <60%
- **Should** mean peace of mind for Network Support
- **Actually** seen as a challenge by physicists

60% occupancy means 40% for free!

– Turn it up until something breaks

Monitor EVERYTHING

What is out there?

CA SPECTRUM

Great for tracking component failures

Much less great for accessing polled data

Report gateway limits

5 min polling of 120 switches and 2k hosts

No support for < 5 minute polling intervals

Polling drops

Deadlines to meet

Wrote own polling engine

APOLL

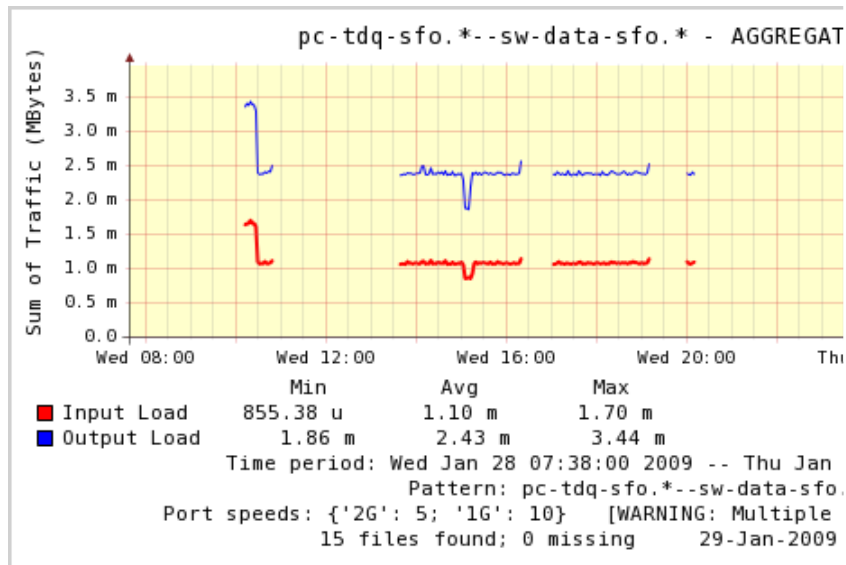
Fast and reliable

Writes to in-ram db

and RRD files

Solved the bug

Reconsider our options



What else is out there?


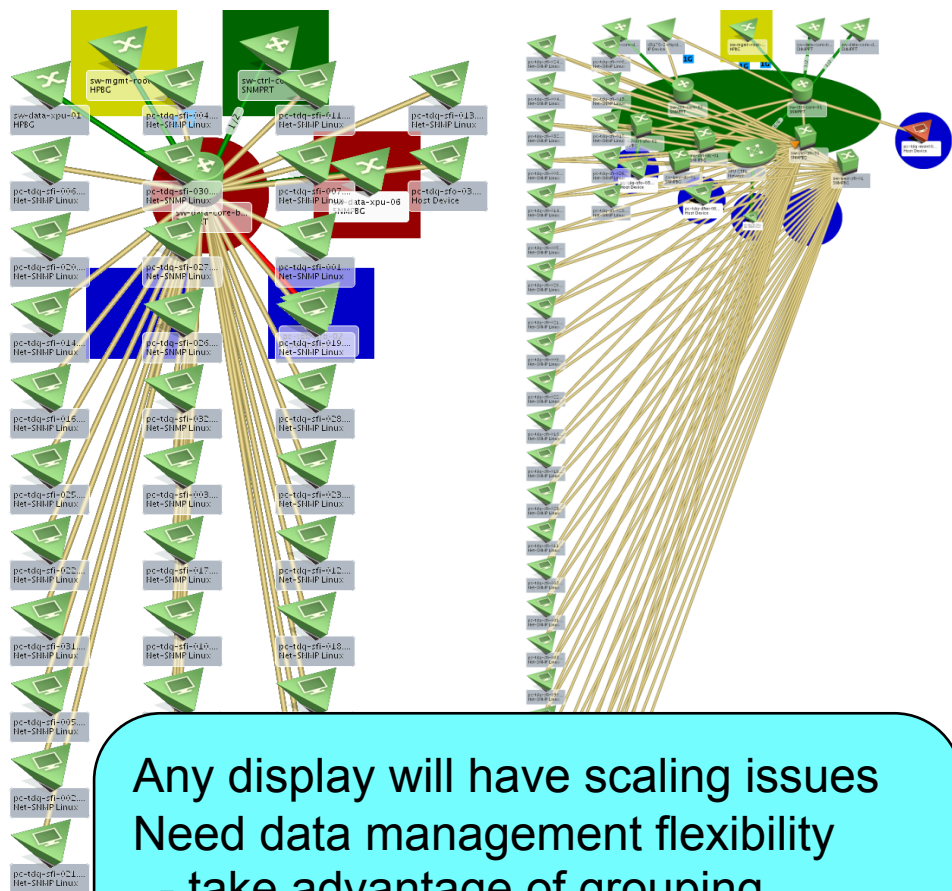
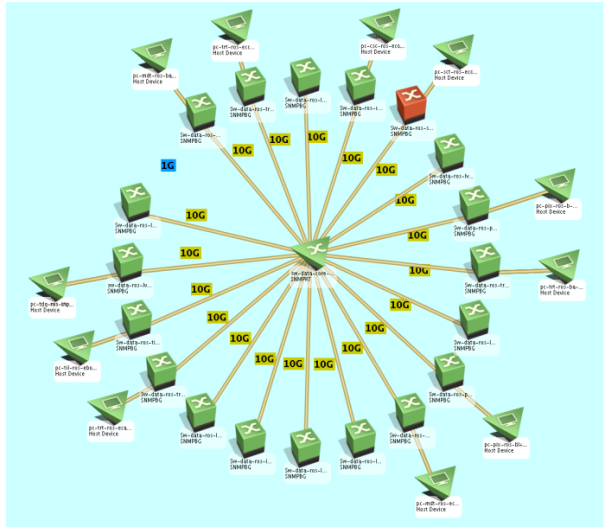
Spectrum	Best for failure alarms, solid archiving
Cacti	Good for visualising single plot RRD data No plot aggregation, Can't handle non RRD, or sFlow Can't use with external discovery
SPINE	Part of Cacti: Stable and fast poller
Also going to need	
Nagios	Great for collecting host CPU stats Poor plot visualisation
sFlow Collector	Our own development
PVSS	For environmental stats (Inherited)

How to pull it all together??

Would like: One stop shopping

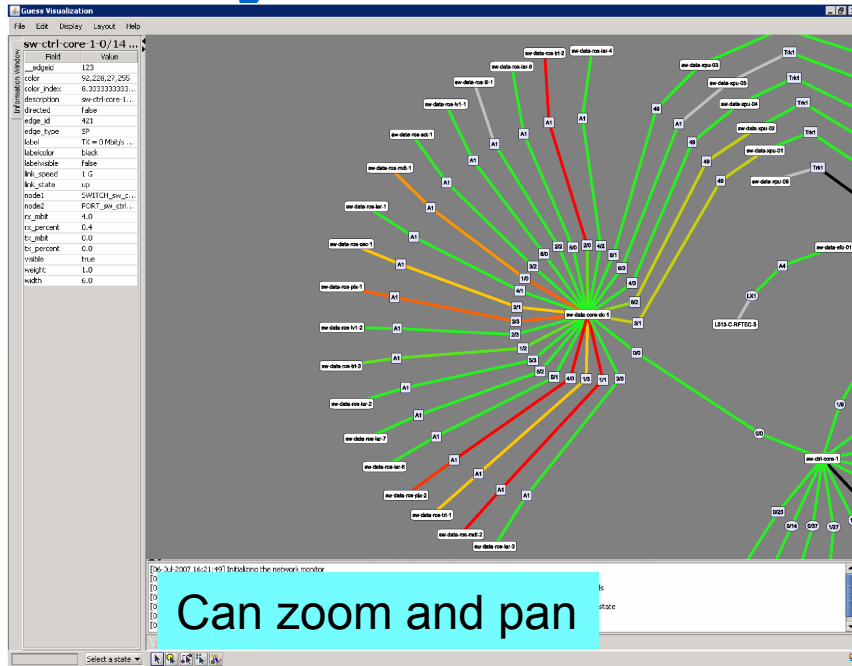
Don't Want: Multiple interfaces to learn and train

System Visualisation 1: Spectrum

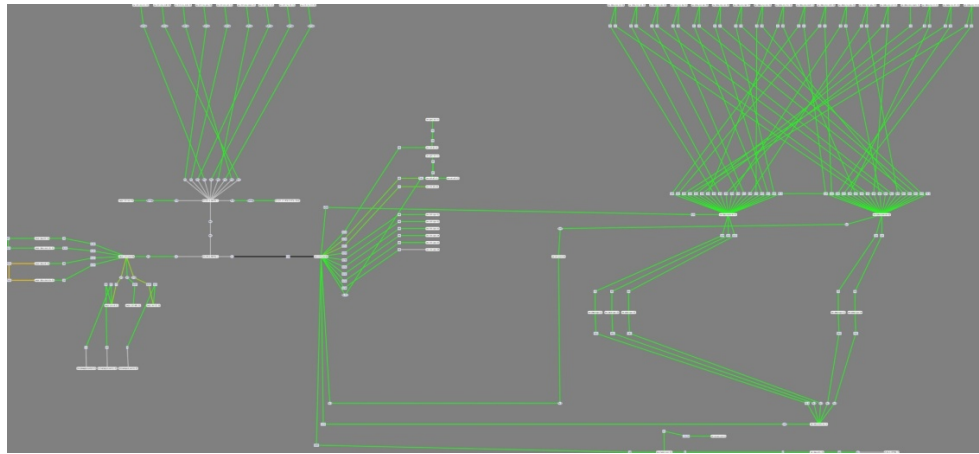


Any display will have scaling issues
 Need data management flexibility
 - take advantage of grouping
 - by affiliation
 - by network
 Need traffic data to be displayed


System Visualisation 2: GUESS




Traffic overlaid on the connectivity.
No hosts, only switches
Autoplacement different for each discovery



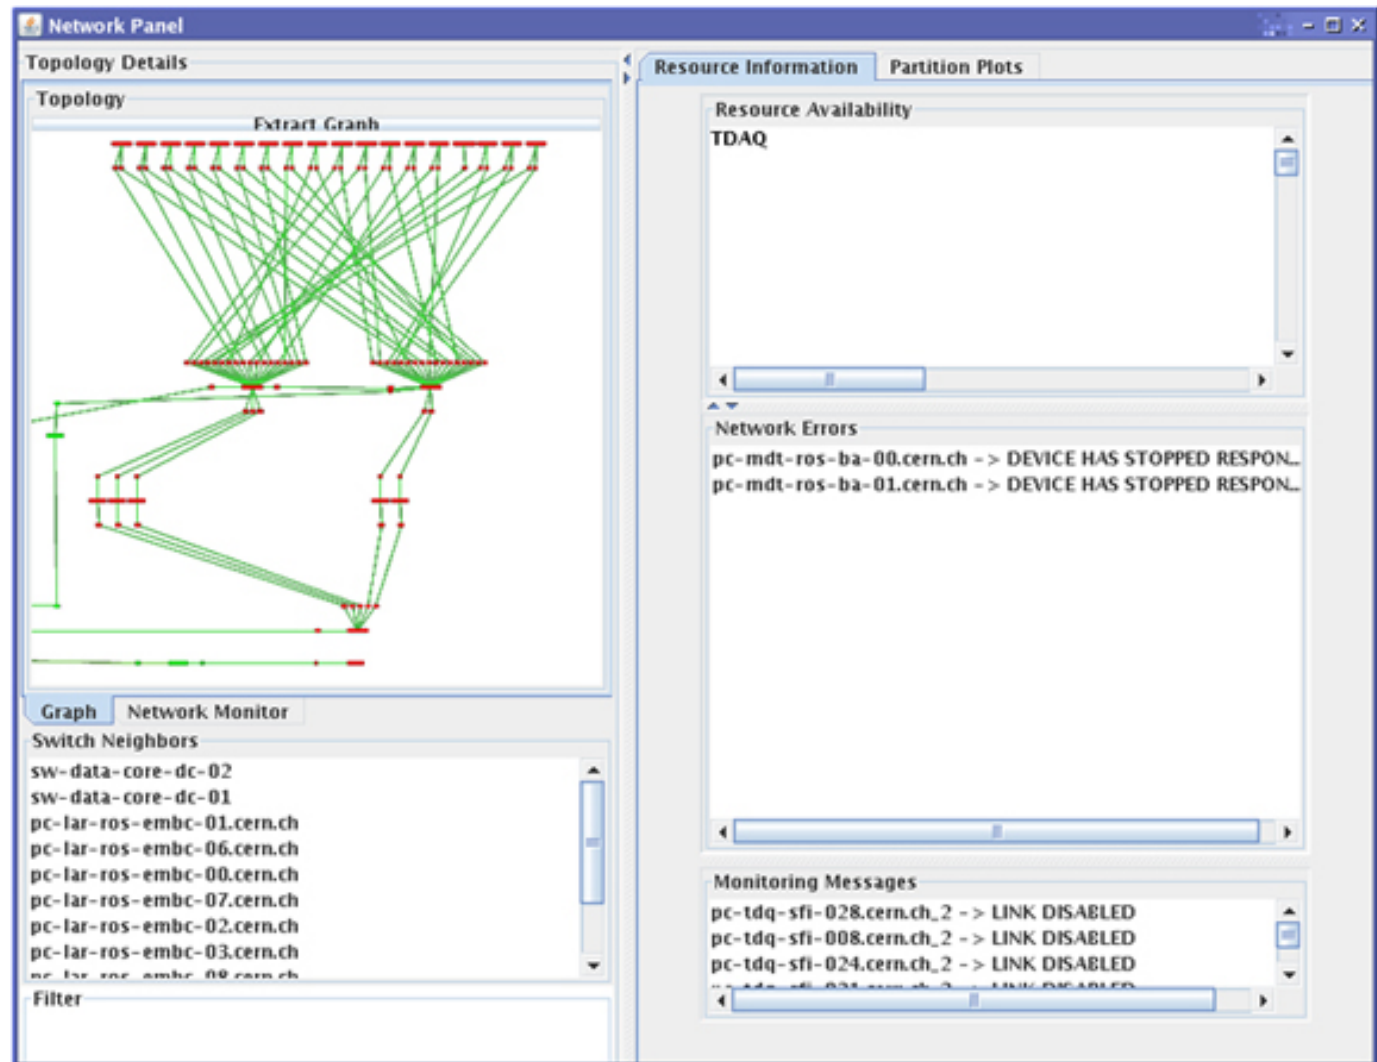
Added semi automatic placement
Sets and keeps the architecture
See the whole picture

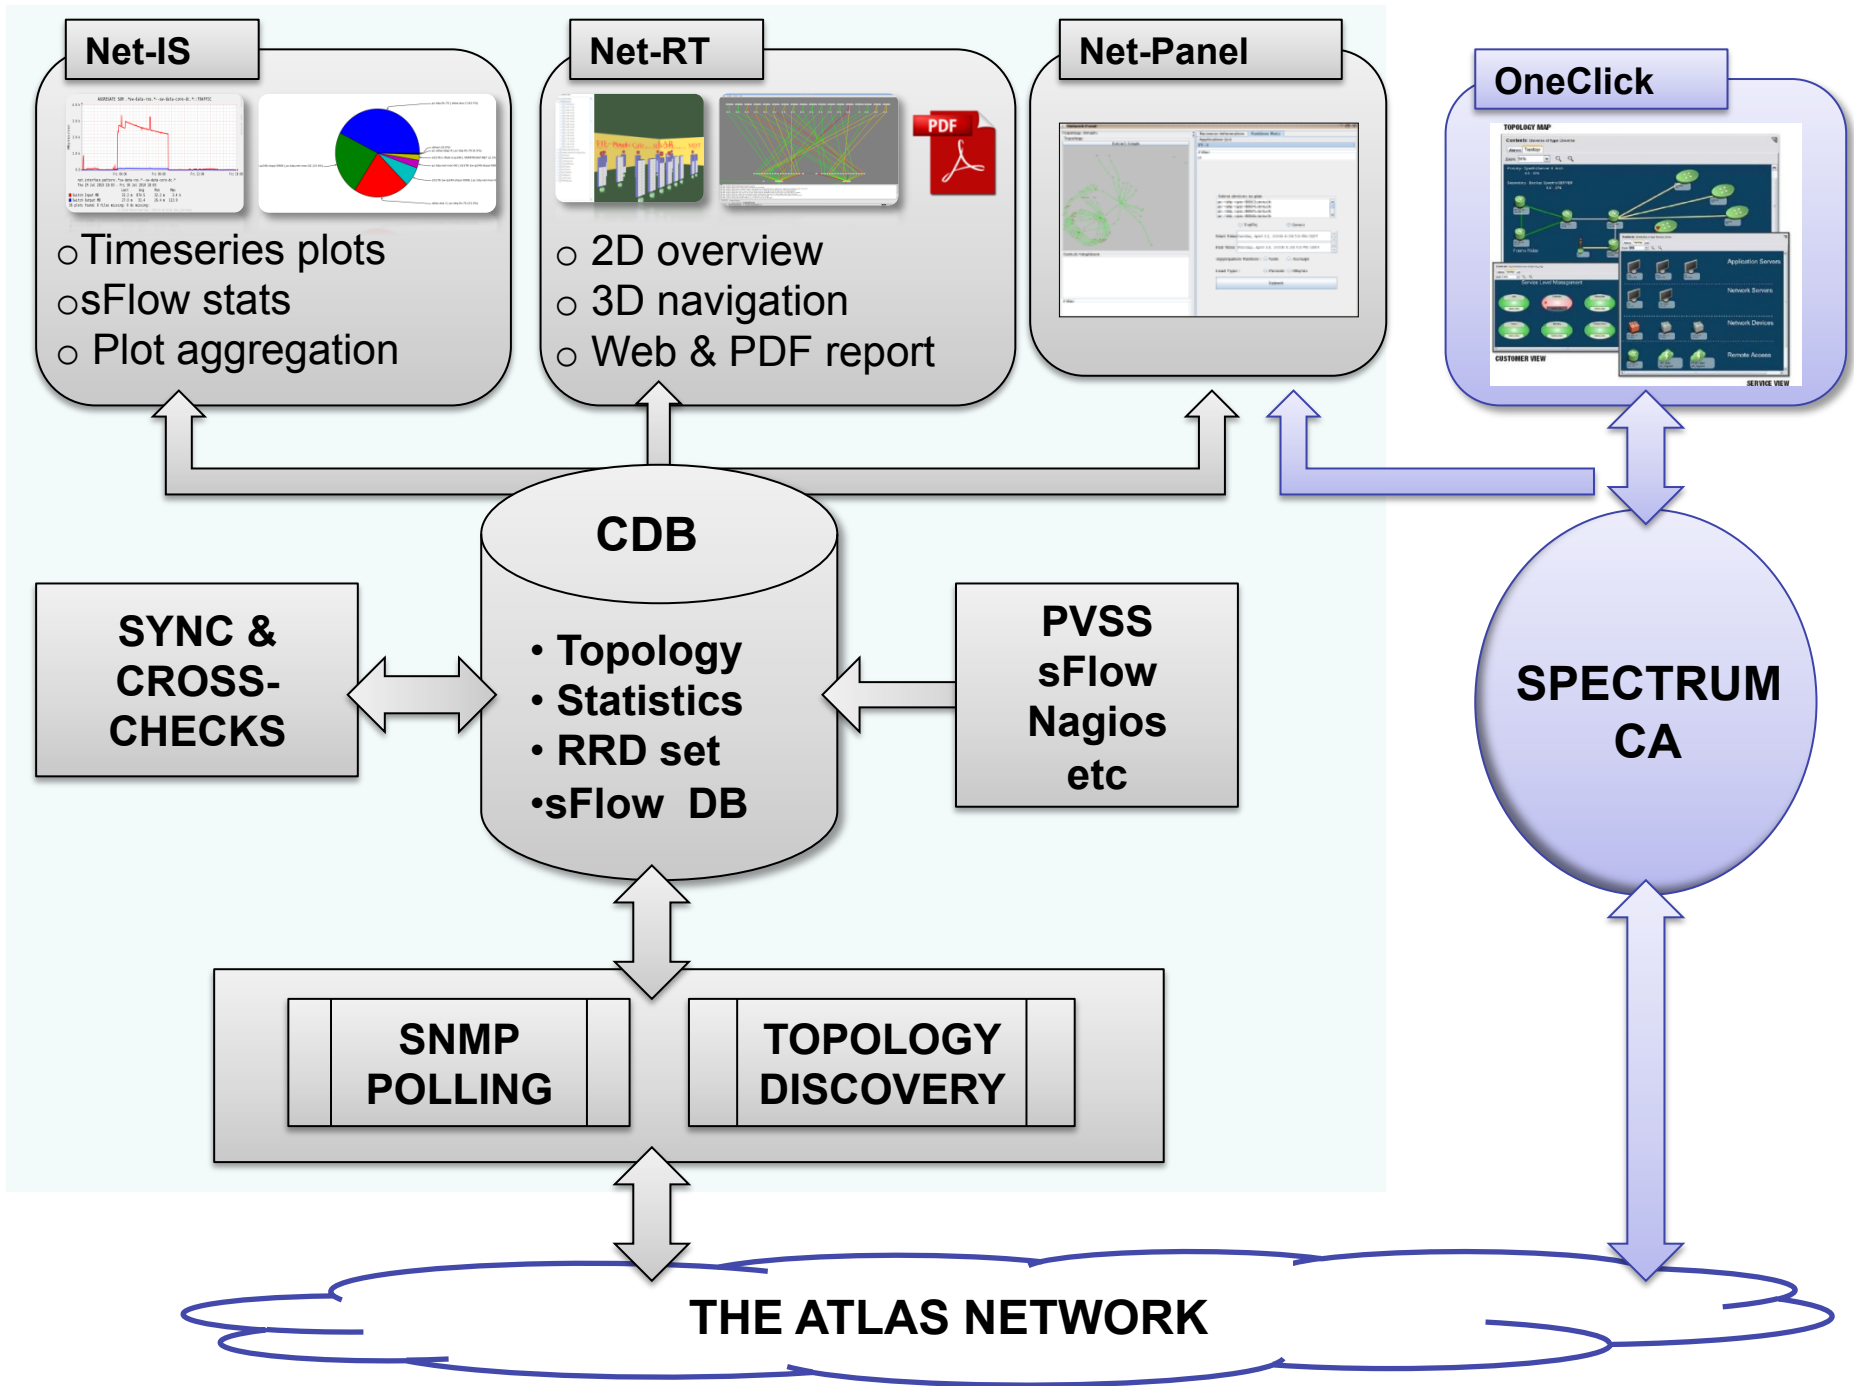




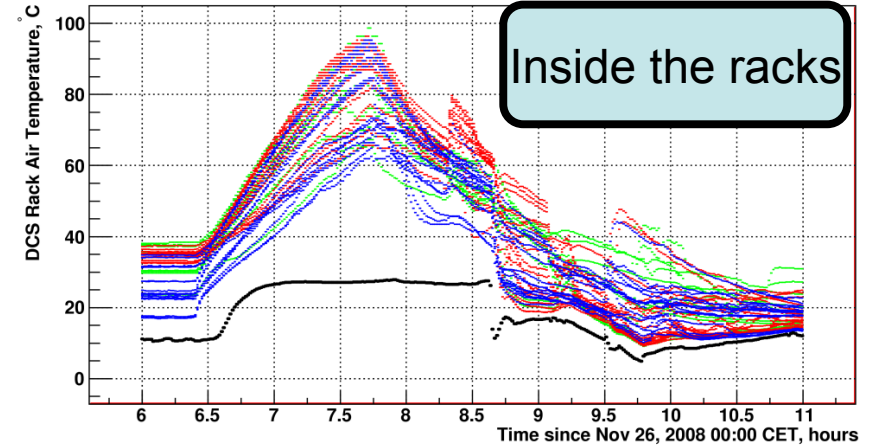
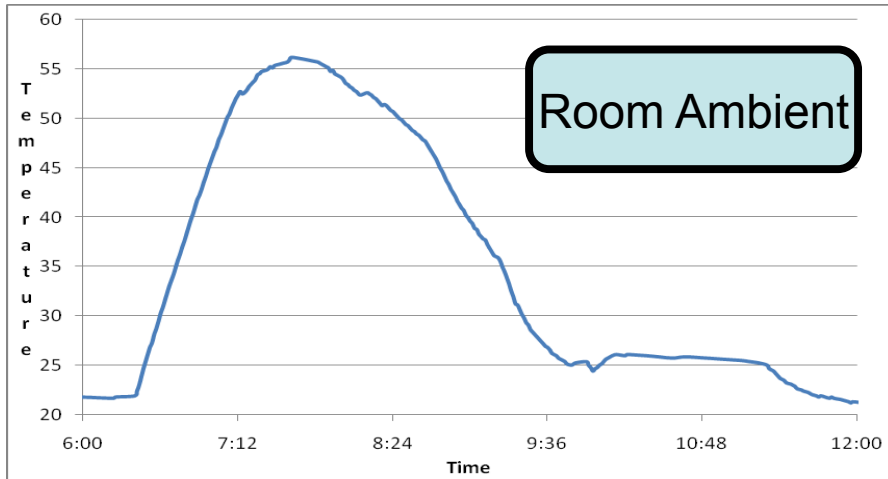
Atlas Users Network Panel

For operators we provide a summary of status per application set.





Wake Up Call



What to do when environmental's go wrong.

Have a TESTED fail-safe plan

Post -Mortem

Fiber connections they are the first to fail

Host Bios Cycle skipping / thermal limits suppressed

Host Bios Fan controls suppressed to protect discs

Routers PSU's/Fans went before auto protect

UPS Mission creep? Need uninterruptible cooling!

SMS Need meaningful short messages

Provided separate tap water supply

Scheduled progressive rack shut down: must be fast

Ensured environmentals go into monitoring.

Network Browser: System Summary

ATLAS Networking Team

- Home page
- GEN Confluence
- Control Room
 - Net-II
 - Net-TestTime
- External Apps
 - NAOS
 - NAOS-Graphs
 - Preser (DC)
 - sLog
 - Log Manager
 - ITAO Analytics
- Internal Apps
 - Switch Logs
 - Net-Consistency
 - OneClick
 - GENES
- Consistency
 - GEN Equipment
 - Yield
- Last Resort
 - LR-Net-II
 - LR-Net-RT
- Topology
 - Net-Resort
 - Decompose
 - Preseries
- Development
 - Net-II Admin
 - Net-TEST
- Links
 - Extra (net team only)
 - Obsolete/ SuScript
 - Obsolete/ sFlow

Erzaed / Colson

View type: Aggregates | Interval: 2010-07-29 23:00 to 2010-07-30 08:05 (3h | 24h) | Engine: Default | Filter: | Refresh | Stop | Search

NETWORK: Traffic Aggregates | Data/ Ctrl Flow Maps: Data Flow Map

Data Flow Map (brief view)

Plot details

AGGREGATE AVG . *513.*

Mbytes/sec

900
800
700
600
500
400
300
200
100
0

23:00 00:00 01:00 02:00 03:00

net.interface.pattern: *513.*.sw-data-sfo.*
Thu 29 Jul 2010 23:00 - Fri 30 Jul 2010 08:00

	Last	Avg
Forward flow MB	562.8	322.4
Reverse flow MB	29.8	16.3

3 plots found; 0 files missing; 0 ds missing;

CASTOR

Seeing the hosts: Aggregation

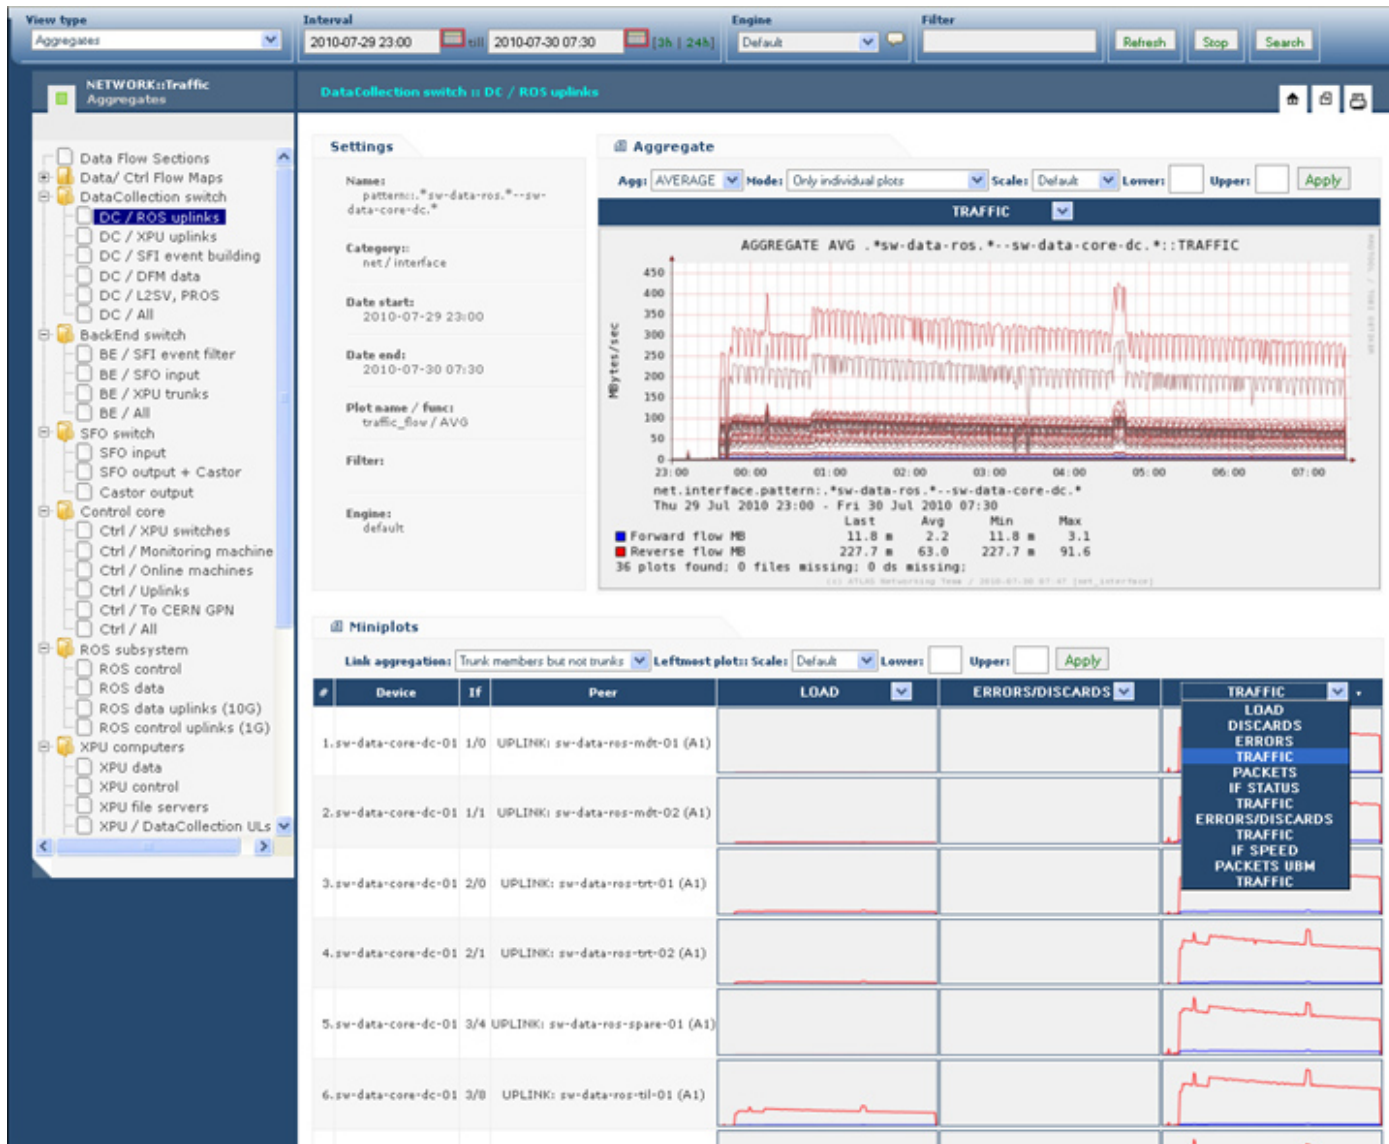
Every tool can produce a single plot. Usually auto-scaled. We are interested in traffic of course,

but also load as % of capacity and of course discards

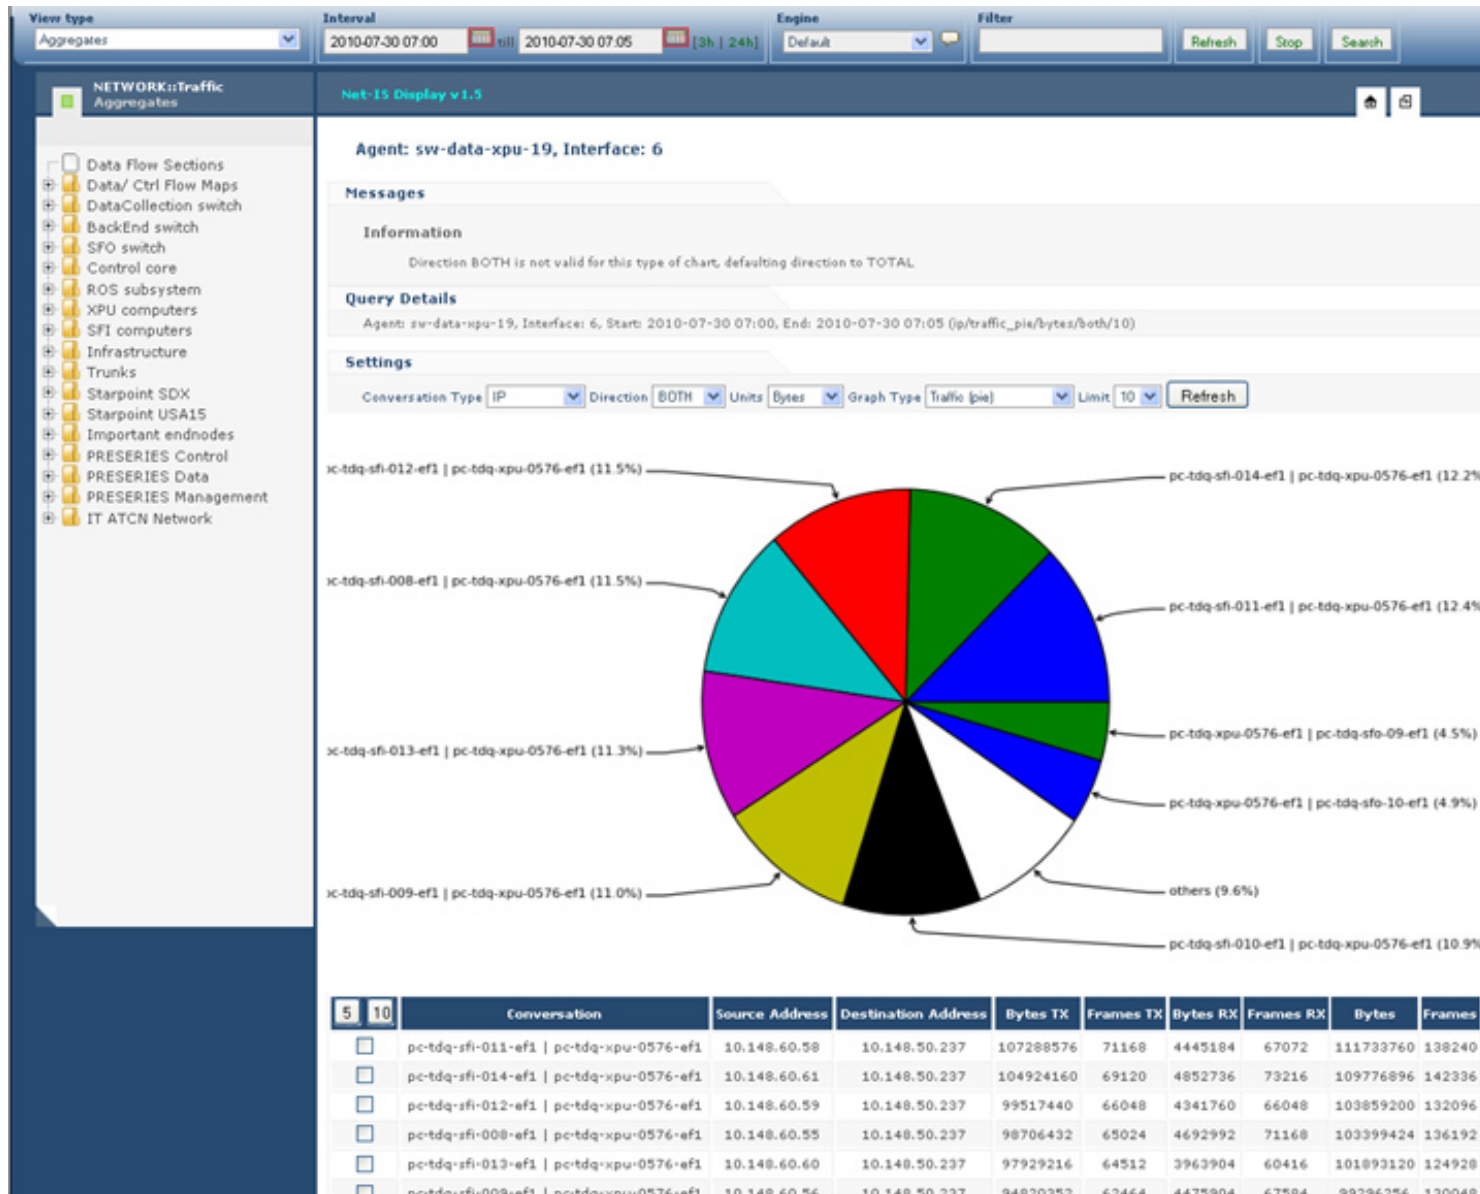


And would love to see any of them for a whole set of peers

Network Browser: Link Aggregates



Network Browser: sFlow



Network Browser: NAGIOS

View type: Category / Device | Interval: 2010-07-29 08:04 till 2010-07-30 08:05 [3h | 24h] | Engine: Default | Filter: xpu-01 | Refresh | Stop | Search

SYSADMIN::PC Sensors
Category / Device

Filter ON

- PC Atlas
- PUB
- CR
- PC LAR
- PC MDT
- PC PIX
- PC Preseries
- PC RPC
- PC SCT
- PC TDQ
 - AMS
 - DFM
 - MON
 - LFS
 - ONL
 - pROS
 - ROS
 - SCR
 - SFI
 - SFO
 - XPU**
 - pc-tdq-xpu-0001
 - pc-tdq-xpu-0002
 - pc-tdq-xpu-0003
 - pc-tdq-xpu-0004
 - pc-tdq-xpu-0005
 - pc-tdq-xpu-0006
 - pc-tdq-xpu-0007
 - pc-tdq-xpu-0008
 - pc-tdq-xpu-0009
 - pc-tdq-xpu-0010
 - pc-tdq-xpu-0011
 - pc-tdq-xpu-0012
 - pc-tdq-xpu-0013
 - pc-tdq-xpu-0014
 - pc-tdq-xpu-0015

PC TDQ :: XPU

Settings

Names: category:~pc-tdq-xpu-
Category: sys / device
Date starts: 2010-07-29 08:04
Date ends: 2010-07-30 08:05
Plot name / func: cpu1_temperature / AVG
Filter: xpu-01
Engine: default

Aggregate

Agg: AVERAGE | Mode: Individual and aggregate plots | Scale: Default | Lower: | Upper: | Apply

CPU1 Temperature

AGGREGATE AVG ~pc-tdq-xpu-::CPU1 Temperature

sys.device.category:~pc-tdq-xpu-
filter:xpu-01
Thu 29 Jul 2010 08:04 - Fri 30 Jul 2010 08:05

	Last	Avg	Min	Max
■ Current (C)	34.0	31.7	29.4	41.1
■ Max (C)	65.0	65.0	65.0	65.0
■ Min (C)	0.0	0.0	0.0	0.0

72 plots found; 2 files missing; 0 ds missing;

Miniplots

Leftmost plots: Scale: Default | Lower: | Upper: | Apply

#	Device	CPU Load	CPU1 Temperature	Ping
1.	pc-tdq-xpu-0100			<ul style="list-style-type: none"> CPU1 Temperature System Temperature CPU2 Temperature Ping Power Supply Voltages Memory Usage CPU Load
2.	pc-tdq-xpu-0101			
3.	pc-tdq-xpu-0102			
4.	pc-tdq-xpu-0103			
5.	pc-tdq-xpu-0104			

Network Browser: Environmental

View type: Room/Rack | Interval: 2010-07-29 08:04 to 2010-07-30 08:05 [2h | 24h] | Engine: Default | Filter: [] [Refresh] [Stop] [Search]

DCS::Environment
Rooms / Rack

- SDX Level 2
 - Y.03-02.D2
 - Y.03-04.D2
 - Y.03-06.D2
 - Y.04-02.D2
 - Y.04-04.D2
 - Y.04-06.D2
 - Y.05-02.D2
 - Y.05-04.D2
 - Y.06-02.D2
 - Y.06-04.D2
 - Y.07-02.D2
 - Y.07-04.D2
 - Y.08-02.D2
 - Y.08-04.D2
 - Y.08-06.D2
 - Y.09-02.D2
 - Y.09-04.D2
 - Y.09-06.D2
 - Y.10-02.D2
 - Y.10-04.D2
 - Y.10-06.D2
 - Y.11-02.D2
 - Y.11-04.D2
 - Y.11-06.D2
 - Y.12-02.D2
 - Y.12-04.D2
 - Y.12-06.D2
 - Y.13-02.D2
 - Y.13-04.D2
 - Y.13-06.D2
 - Y.14-02.D2
 - Y.14-04.D2
 - Y.14-06.D2
 - Y.15-02.D2
 - Y.15-04.D2

Settings

Name: rroom:SDX Level 2
Category: net/rack
Date starts: 2010-07-29 08:04
Date end: 2010-07-30 08:05
Plot name / func: air_temperature / AVG
Filters:
Engine: default

Aggregate

Agg: AVERAGE | Mode: Only individual plots | Scale: Default | Lower: [] | Upper: [] | Apply

Air temperature

AGGREGATE AVG SDX Level 2::Air temperature

net.rack.room:SDX Level 2
Thu 29 Jul 2010 08:04 - Fri 30 Jul 2010 08:05

	Last	Avg	Min	Max
Temperature (C)	33.7	34.6	33.6	35.9

46 plots found; 0 files missing; 0 ds missing

Miniplots

Leftmost plots: Scale: Default | Lower: [] | Upper: [] | Apply

#	Device	Room	Air temperature	Water Temperature	Power Status (Breaker 1)
1.	Y.16-06.D2	SDX Level 2	NO DATA	NO DATA	Power Status (Breaker 1) Fan Status Water Temperature Power Status (Breaker 2) Air temperature
2.	Y.12-02.D2	SDX Level 2	[Line]	[Line]	
3.	Y.07-02.D2	SDX Level 2	[Line]	[Line]	
4.	Y.16-02.D2	SDX Level 2	[Line]	[Line]	
5.	Y.11-02.D2	SDX Level 2	[Line]	[Line]	
6.	Y.10-02.D2	SDX Level 2	[Line]	[Line]	

Dynamic pages with real time traffic

ATLAS Networking
Network Interfaces :: Discovered but down

Network Interfaces::Discovered but down

#	Device	Interface	Peer	STATUS	LastInfo
1	D3125-1V-IP16-SHP3L-2919	42	PS-RPC-H1 (landb)	Down	11.2 min ago
2	D3125-1V-IP16-SHP3L-2919	43	PS-RPC-L2 (landb)	Down	11.2 min ago
3	D3125-1V-IP16-SHP3L-2919	44	PS-RPC-L1 (landb)	Down	11.2 min ago
4	D3125-1V-IP16-SHP3L-2919	29	ncatmdns3.ctrl0	Down	11.2 min ago

ATLAS Networking
Network Interfaces :: Not operating at full speed

Network Interfaces::Not operating at full speed

#	Device	Interface	Peer	LastInfo	speed	normal_speed
1	sw-data-core-dc-01	Po 1	UPLINK: sw-data-core-dc-02 (Po 1)	9.3 min ago	30000.0	40000
2	sw-data-core-dc-02	Po 1	UPLINK: sw-data-core-dc-01 (Po 1)	9.3 min ago	30000.0	40000
3	swps-ctrl-core-01	25	sbc-preseries-roib-00.ctrl0	9.3 min ago	100.0	1000
4	sw-ctrl-xpu-18	11	pc-tdq-xpu-0548.ctrl0	9.3 min ago	100.0	1000
5	sw-ipmi-sfi-01	30	pc-tdq-sfi-018.mgmt	9.3 min ago	100.0	1000
6	sw-data-core-dc-02	11/13	pc-tdq-sfi-018.data_dc2	9.3 min ago	100.0	1000
7	sw-ctrl-sfi-01	34	pc-tdq-sfi-018.ctrl0	9.3 min ago	100.0	1000
8	sw-data-xpu-23	17	pc-tdq-lfs-xpu-23.data0 (default)	9.3 min ago	10.0	1000
9	sw-ctrl-xpu-03	17	pc-tdq-lfs-19.mgmt	9.3 min ago	10.0	1000
10	swps-ctrl-core-01	24	pc-preseries-l1src-01.ctrl0	9.3 min ago	100.0	1000

ATLAS Networking
Network Interfaces :: Switch Uplinks

Network Interfaces::Switch Uplinks

#	Device	Interface	Peer	LoadInOut	LoadIn	LoadOut	TrafficIn_MB	TrafficOut_MB	LastInfo
1	D3178-2-RHP2L-1	K23	UPLINK: D2175-RS-IPX-SHP3M-1 (24)	51.57 %	25.65 %	25.91 %	30.58	30.89	11.9 min ago
2	D3125-2-RHP2L-1	L19	UPLINK: D3125-2V-IP19-SHP3M-0511 (24)	50.52 %	25.29 %	25.23 %	30.15	30.08	7.9 min ago
3	D3125-2V-IP19-SHP3M-0511	24	UPLINK: D3125-2-RHP2L-1 (L19)	50.34 %	25.07 %	25.28 %	29.88	30.13	7.9 min ago
4	D2175-RS-IPX-SHP3M-1	24	UPLINK: D3178-2-RHP2L-1 (K23)	49.89 %	25.13 %	24.77 %	29.96	29.52	7.9 min ago
5	D3178-2-RHP2L-1	A1	UPLINK: D3125-2-RHP2L-1 (A1)	11.84 %	6.19 %	5.64 %	73.84	67.28	11.9 min ago
6	D3125-2-RHP2L-1	L18	UPLINK: D3125-2V-IP18-SHP3M-2411 (24)	8.48 %	3.17 %	5.31 %	3.78	6.33	7.9 min ago
7	D3178-2-RHP2L-1	B1	UPLINK: sw-ctrl-core-01 (6/2)	6.06 %	5.25 %	0.81 %	62.54	9.65	11.9 min ago
8	sw-ctrl-ros-mdt-02	23	UPLINK: D3125-2-RHP2L-1 (K12)	6.07 %	1.02 %	5.05 %	1.22	6.02	11.9 min ago
9	D3125-2V-IP18-SHP3M-2411	24	UPLINK: D3125-2-RHP2L-1 (L18)	7.57 %	4.83 %	2.75 %	5.75	3.27	11.9 min ago
10	D3125-1V-IP15-SHP3M-2619	24	UPLINK: D3125-2-RHP2L-1 (L15)	4.94 %	4.76 %	0.18 %	5.68	0.22	7.9 min ago
11	D3125-2-RHP2L-1	K12	UPLINK: sw-ctrl-ros-mdt-02 (23)	4.87 %	4.71 %	0.16 %	5.62	0.19	7.9 min ago
12	D3125-2-RHP2L-1	L15	UPLINK: D3125-1V-IP15-SHP3M-2619 (24)	4.70 %	0.17 %	4.53 %	0.21	5.4	7.9 min ago
13	D3125-2-RHP2L-1	A1	UPLINK: D3178-2-RHP2L-1 (A1)	6.57 %	3.23 %	3.34 %	38.45	39.86	7.9 min ago

2D Display Limitations

GUESS display good for switch to switch traffic

Can't incorporate host details - - -overwhelming

The Browser gives all the host details

BUT I WANT:

- all the detail when something goes wrong
- to see the neighbours of a problem node
- a system wide view for visual correlation
- different detail depending on my viewpoint
- fly-through, and navigation and visible errors and pop ups and reports...

I want GOOGLE Earth for my network.

3D Display

Google Earth as inspiration
Variable detail as a function of
viewing distance

Variable viewing angles
Intuitive navigation

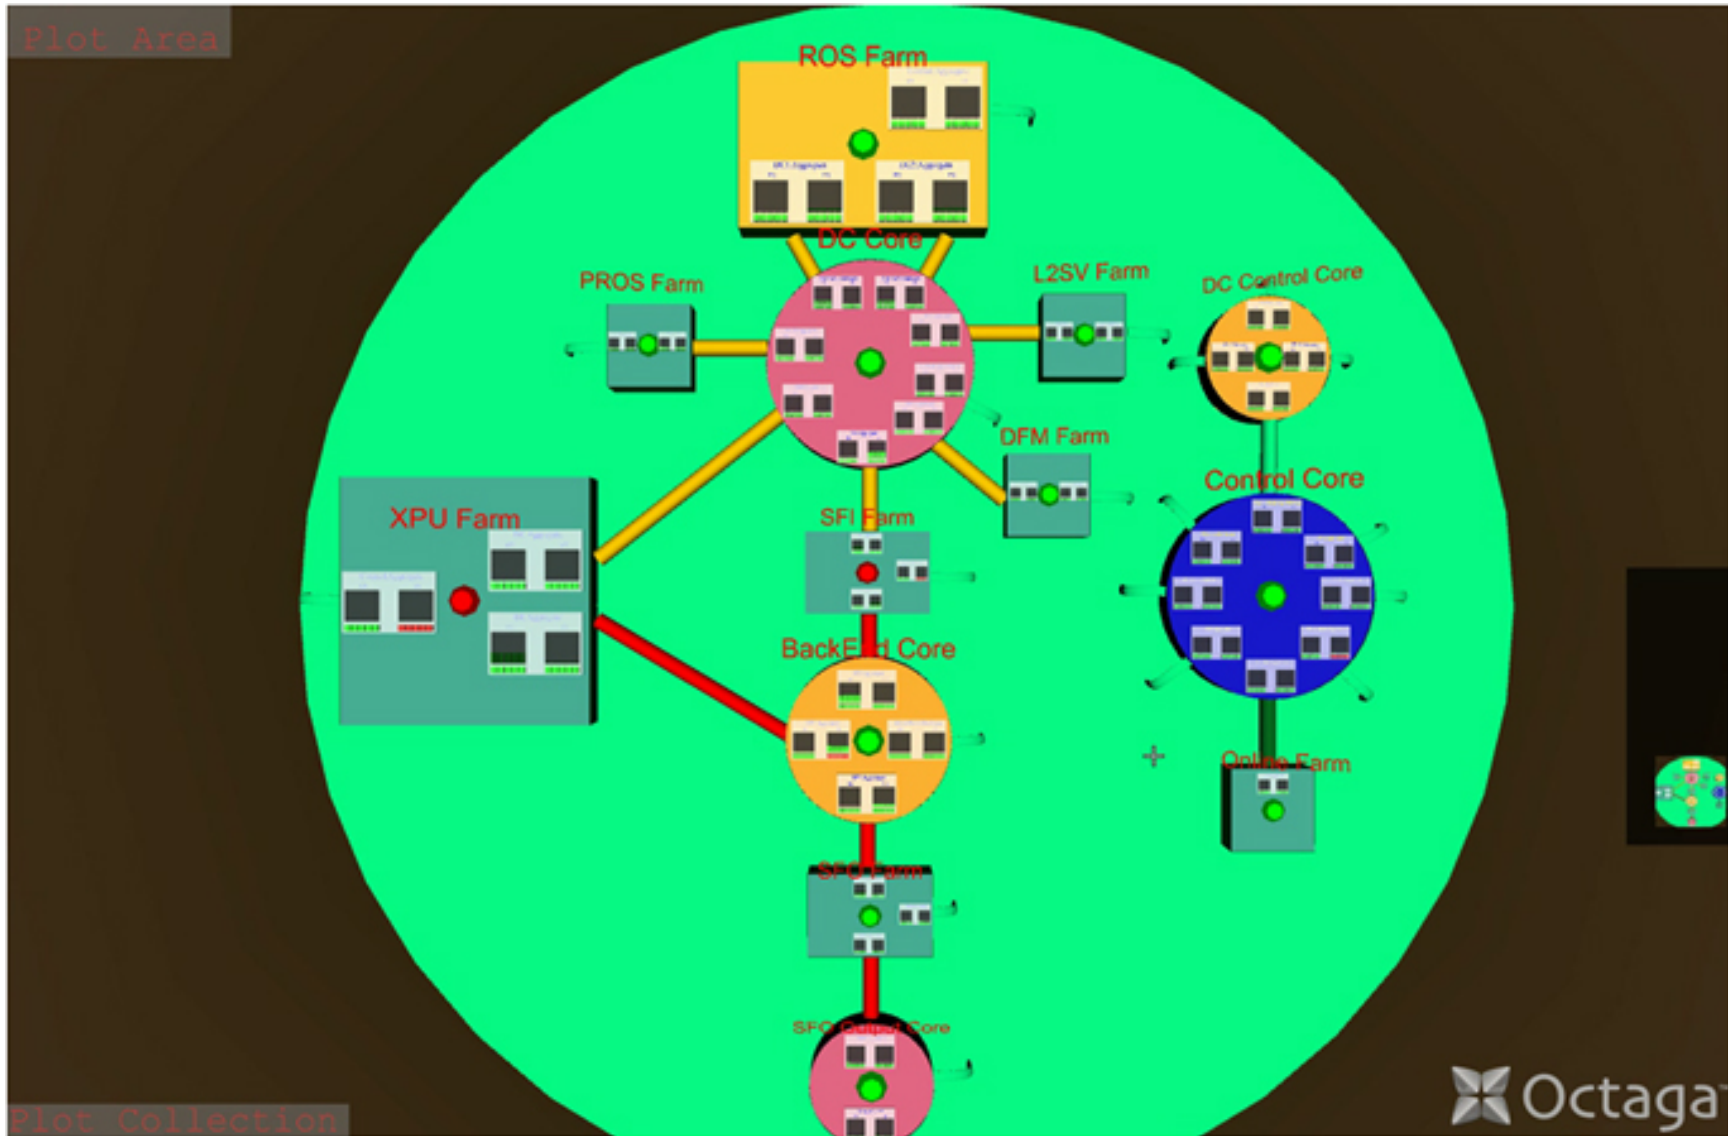
Unfortunately Google didn't
cope with our dynamic update
requirements

So we went looking for display
software that does.

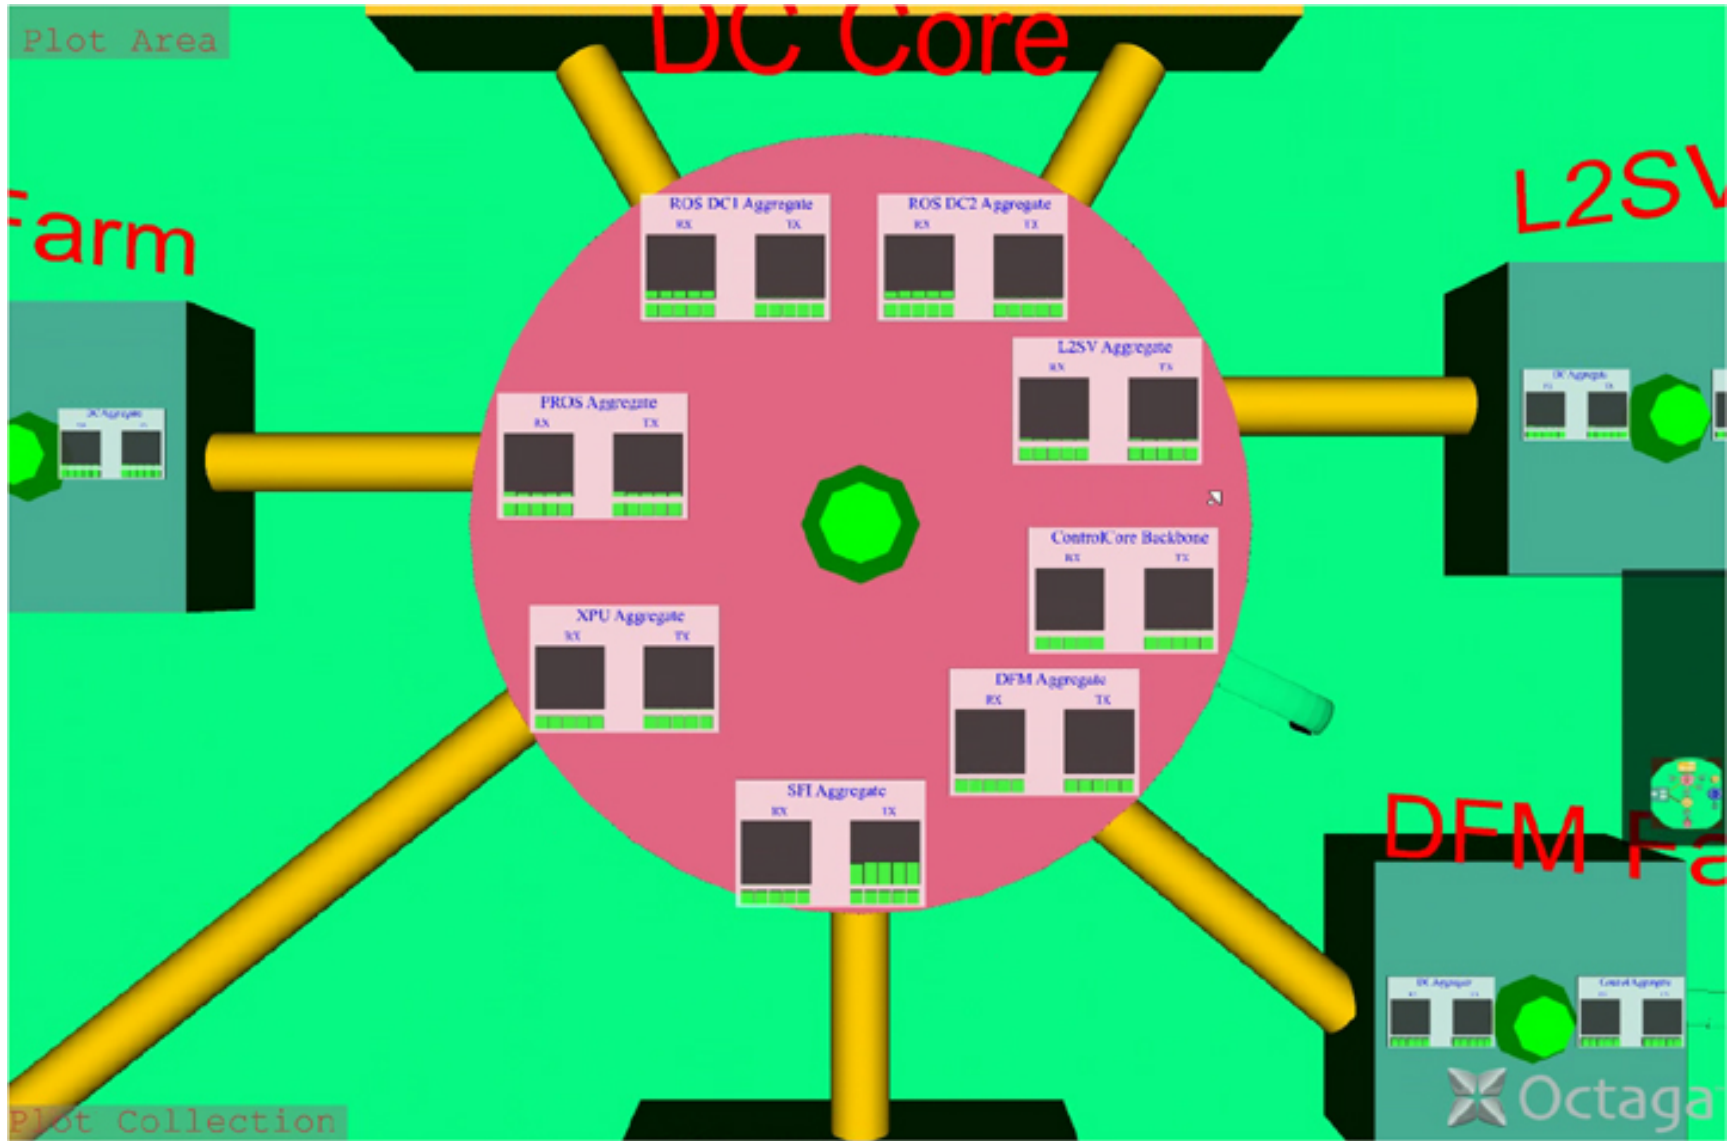
X3D (enhanced VRML)
Octaga Player



3D Top Level View



3D Top Level View



Detail of Statistics Windows



Distance Sensitive Detail



Conclusions

- Total system coverage for ~3K nodes with ~8K ports
- Advanced network browsing
- Affiliated and agglomerated plotting
- Full detail down to processor level
- Automated detection for system wide traffic / error thresholds
- Visual intuitive feedback for system diagnosis and monitoring
- 2 and 3 dimensional displays

Future work

- Rules based expert system
- Application centric view of network (not connection centric)

And to finish we have a fly-through of the 3D network display