NASA EOSDIS Network Monitoring

New Active, Passive and Real-time Monitoring Approaches

Current Activities and Plans

JET Roadmap Workshop

April 13, 2004

EOSDIS Activities and Plans Agenda

- Current Activities:
  - Flow Graphs
  - Live Monitoring Tools
  - Flow Analysis Tools
  - Multi-node Active Measurements
  - Integrated Active Measurements
  - NetFlow Load Study
  - Sampling Study
- Planned Activities

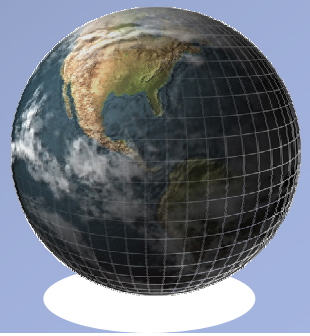

FlowGraphs – Flow Rate Graphs

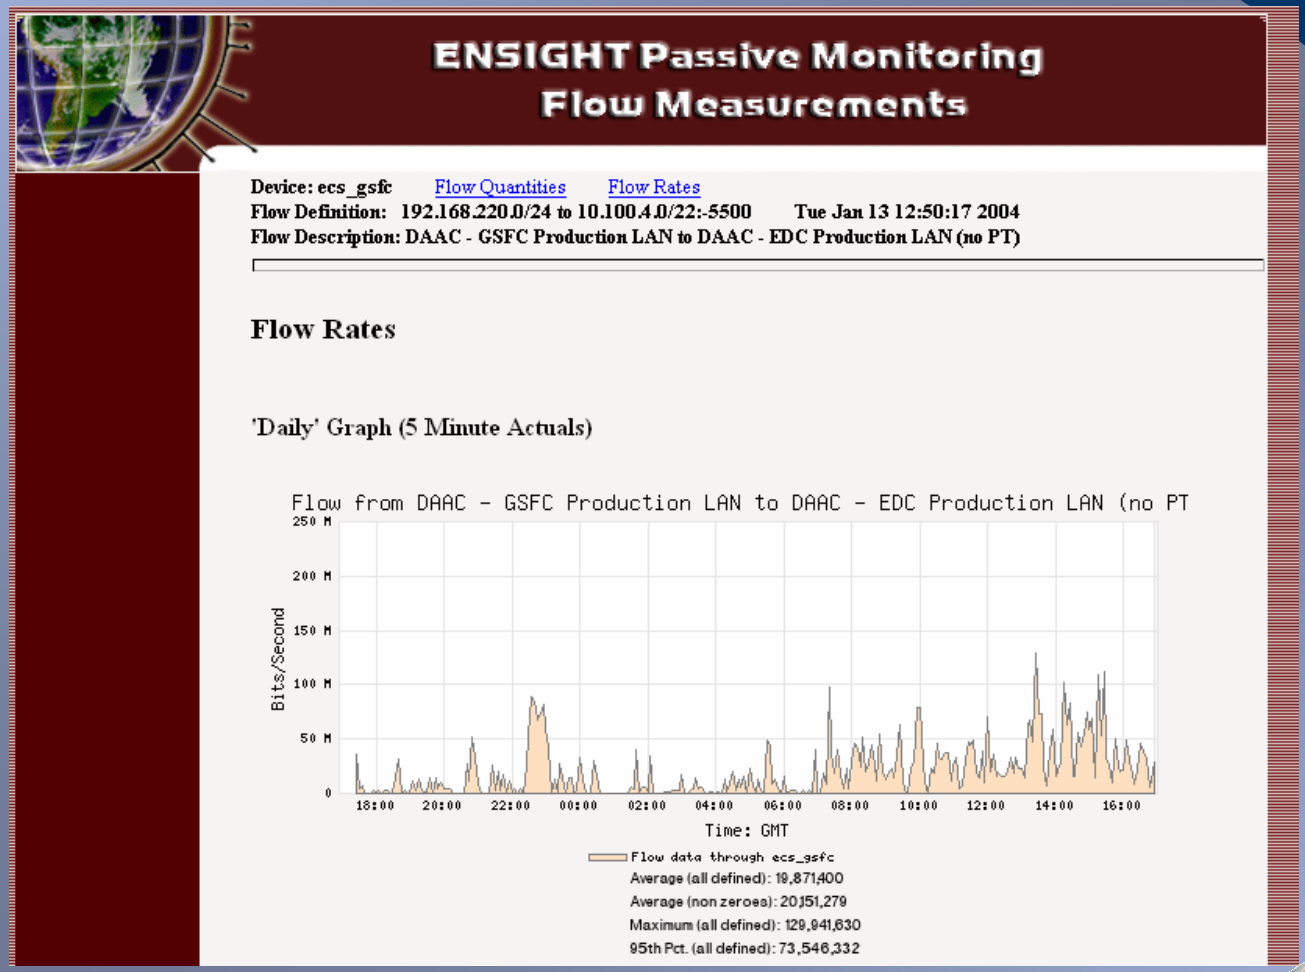

- **MRTG-like graphs for Flows instead of Interfaces**
- **Rates and Quantities are tracked**
- **This flow tracks science product delivery to archive**

FlowGraphs – Flow Quantity Graphs

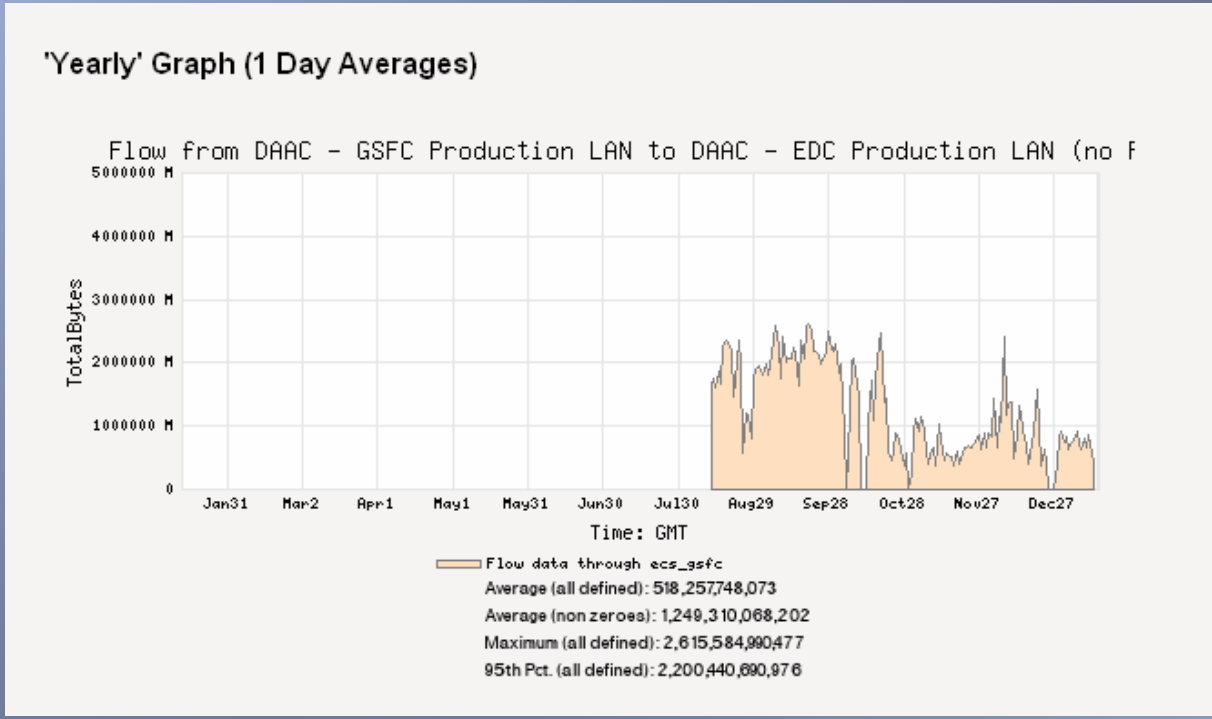

- **Shows ‘quantity’ of data for a particular flow**
- **Tracks delivery of science products to archive in Sioux Falls, SD**
- **Average 1.2 Terabytes/day; Peak: 2.6 Terabytes/day**
- **Useful for determining future resource requirements**

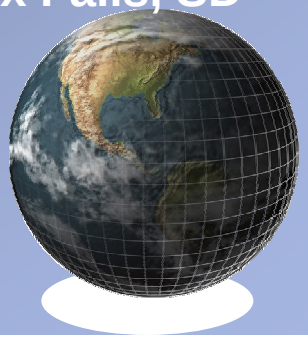

FlowGraphs – Behind the Scenes

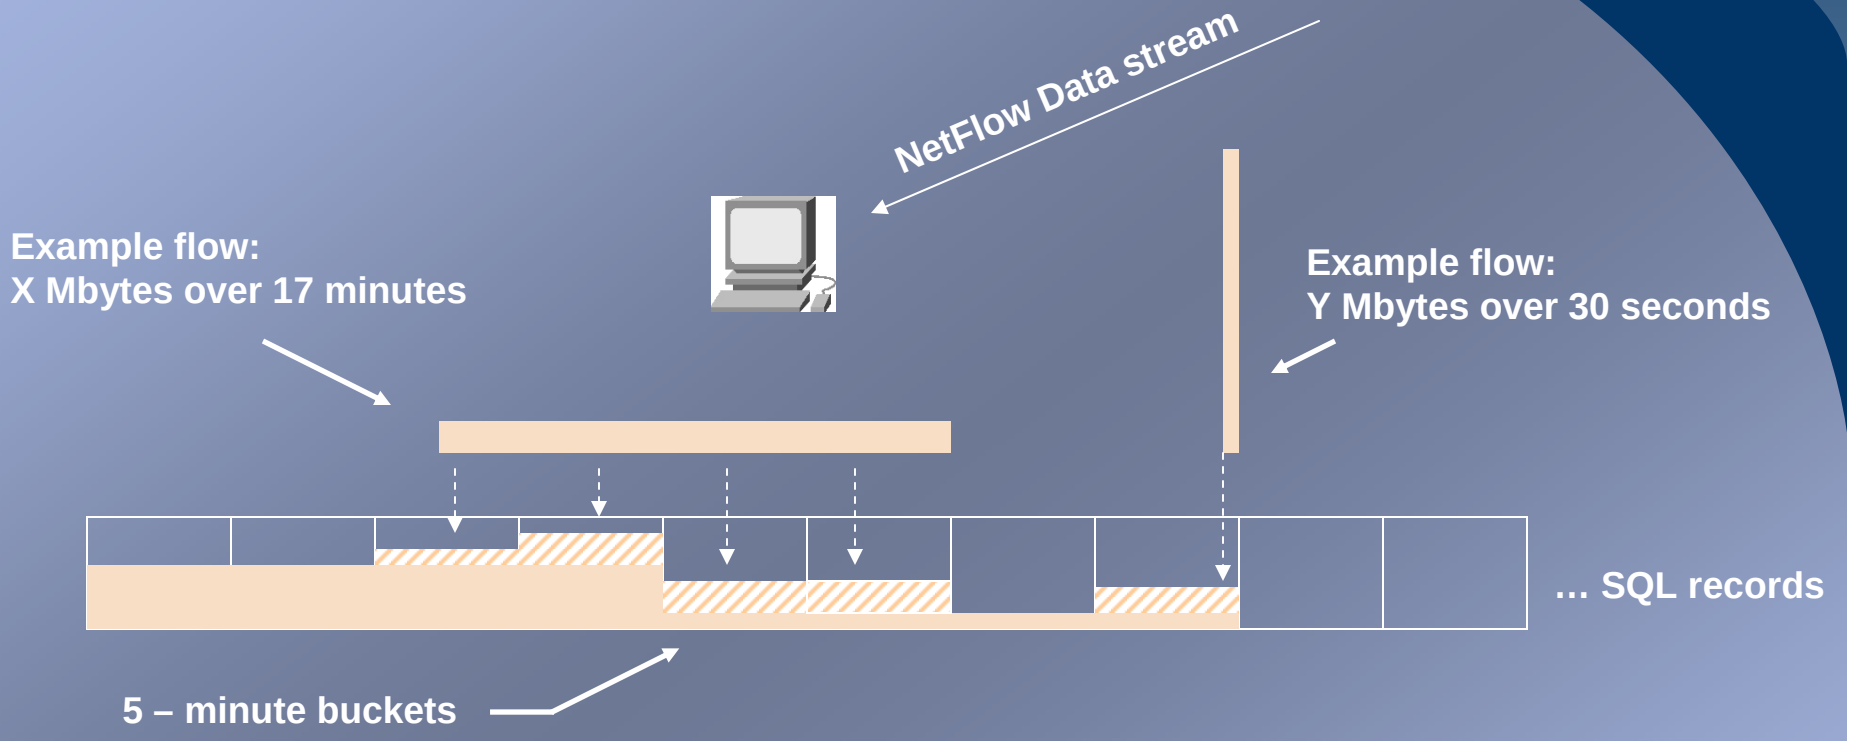

- For each defined Flow, all received NetFlow data is examined
- For defined Flows: byte count is pro-rated across 5-minute buckets, stored
- Flowgraph data accumulates over time, producing smooth graph

Live Monitoring of an Interface

Live Object Monitoring

Collection Period: 7 minutes Device: zzz_gsfc
Querying Interval: 5 seconds Object: 1.3.6.1.2.1.31.1.1.1.6.13
Graphing Interval: 5 seconds Description: ge800in:

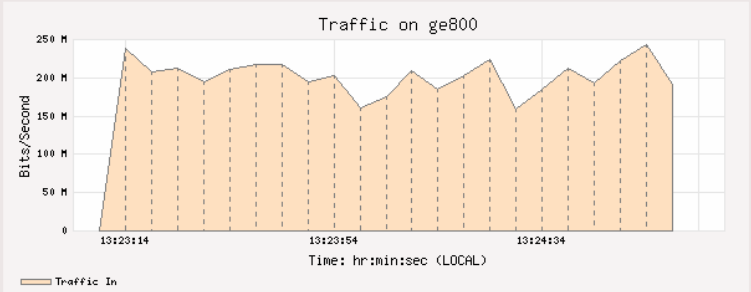

This web-page will continue to be updated with flow data for an additional 2 minutes after the collection period noted above. The page will be updated every 15 seconds. What am I looking at?

Host Summaries

Source	Destination	Bits	Rate
XXacXX1u.zzz.nasa.gov	xxxxx.xxxxxx.nasa.gov	11,338,247,280	103,074,975
XXfwi09.gsfcb.zzz.nasa.gov	samantha.xxxx.nasa.gov	1,870,807,856	17,007,344
XXisi08u.zzz.nasa.gov	miracle.rsmas.miami.edu	653,100,920	5,937,281
XXdps01u.zzz.nasa.gov	192.168.63.97	597,716,016	5,433,781
XXisi08u.zzz.nasa.gov	omisips2.xxxxxx.nasa.gov	563,851,168	5,125,919
XXisi08u.zzz.nasa.gov	192.168.69.20	468,005,368	4,254,594
XXisi08u.zzz.nasa.gov	aku.sci.yyyy.nasa.gov	308,348,496	2,803,168
XXfwi09.gsfcb.zzz.nasa.gov	l0dus02.xxxxb.zzz.nasa.gov	230,216,736	2,092,879
(Totals include values for all 143 host-pairs collected)		16,236,703,128	147,606,392

Individual Flows

Source	Port	Destination	Port	Bits	Flow Rate	Overall Rate
XXacXX1u.zzz.nasa.gov	33263	xxxxx.xxxxxx.nasa.gov	42167	5,316,804,424	117,972,938	48,334,585
XXacXX1u.zzz.nasa.gov	32795	xxxxx.xxxxxx.nasa.gov	41640	5,316,801,000	107,236,809	48,334,554
XXacXX1u.zzz.nasa.gov	65466	xxxxx.xxxxxx.nasa.gov	41298	700,442,512	125,123,705	6,367,659
XXdps01u.zzz.nasa.gov	65449	130.85.163.97	52900	597,713,184	9,271,902	5,433,756
XXfwi09.gsfcb.zzz.nasa.gov	47523	samantha.xxxx.nasa.gov	20	333,840,424	8,161,559	3,034,912
XXisi08u.zzz.nasa.gov	36949	omisips2.xxxxxx.nasa.gov	20	318,196,856	7,710,311	2,892,698
XXfwi09.gsfcb.zzz.nasa.gov	50672	samantha.xxxx.nasa.gov	20	310,324,832	8,595,303	2,821,134
XXfwi09.gsfcb.zzz.nasa.gov	50482	samantha.xxxx.nasa.gov	20	275,791,360	7,565,047	2,507,194
XXisi08u.zzz.nasa.gov	49049	192.225.69.20	20	265,682,920	12,249,097	2,415,299
XXisi08u.zzz.nasa.gov	49804	miracle.rsmas.miami.edu	20	259,392,872	5,347,755	2,358,117
XXfwi09.gsfcb.zzz.nasa.gov	48067	samantha.xxxx.nasa.gov	20	241,152,464	9,927,237	2,192,295
XXfwi09.gsfcb.zzz.nasa.gov	48423	l0dus02.xxxxb.zzz.nasa.gov	58294	229,620,232	6,931,931	2,087,456
XXfwi09.gsfcb.zzz.nasa.gov	48437	samantha.xxxx.nasa.gov	20	213,175,304	7,785,234	1,937,957
XXfwi09.gsfcb.zzz.nasa.gov	48903	samantha.xxxx.nasa.gov	20	206,842,752	9,771,483	1,880,388

This tool monitors an interface near in real-time.

Interface Utilization

The *Interface Utilization* graph is updated every 5 seconds. The interface was observed for 145 seconds at this point.

Host to Host

Live NetFlow data is collected and reduced using Mark Fullmer's Flow Tools suite. Totals are maintained for the most significant source-destination pairs. The report includes total bits and bits/second.

Host:Port to Host:Port

Data is also provided for host:port pairs. These connections make up the total contribution to the aggregate utilization shown in the graph. Note multiple parallel transactions contributing to large data transfers.

The *Flow Rate* is the rate of the flow itself. The *Overall Rate* is calculated from total bits per flow divided by elapsed collection time. The sum of this column should match the rate shown in the graphic.

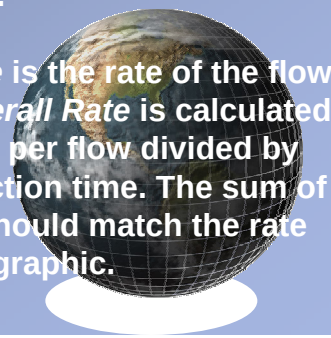

Creating a Custom Flowgraph

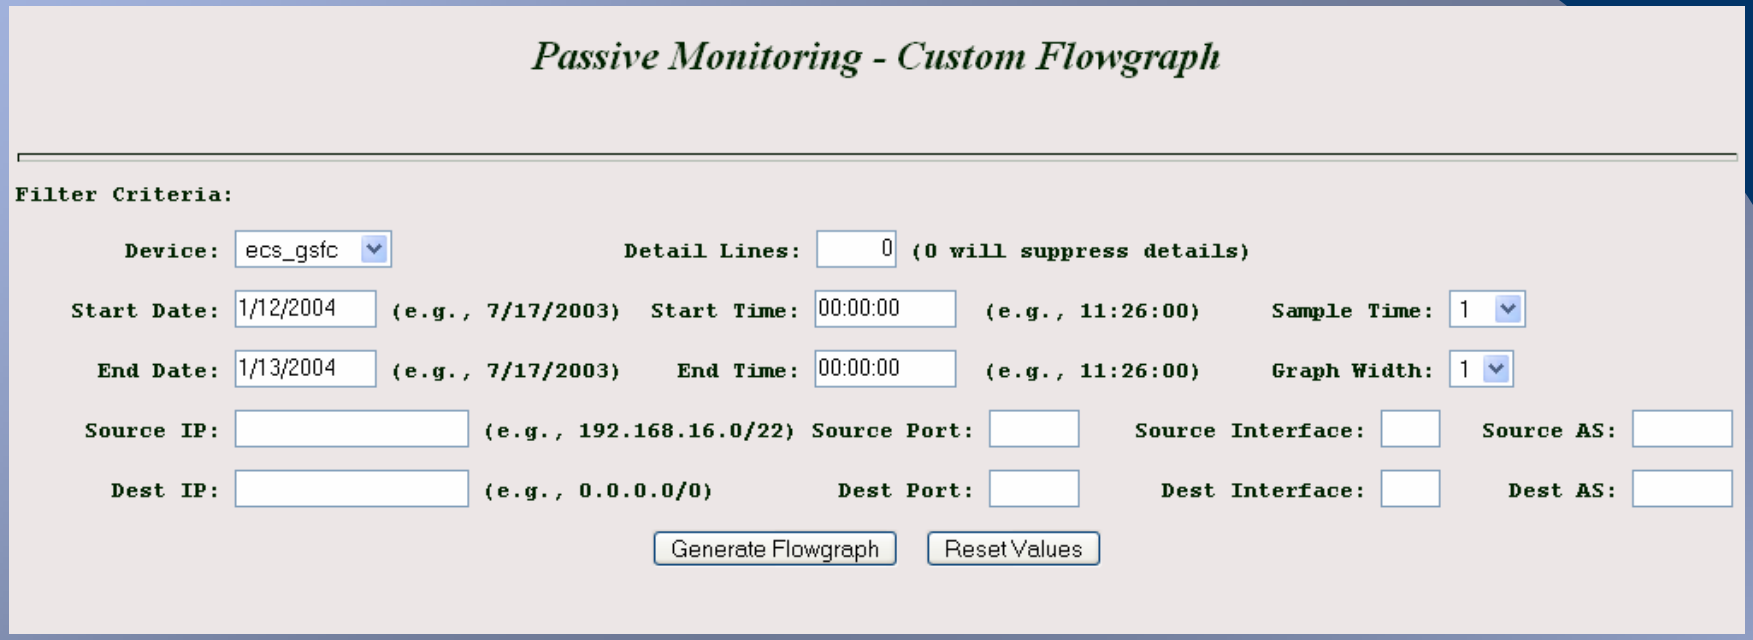

- Provides for graphical analysis of stored flow data
- Can look at time period, subnetworks, ports, interfaces, AS, etc.
- Graph Width option allows for unpacking dense graphs
- Detail lines permits accompanying textual breakout of Flowgraph

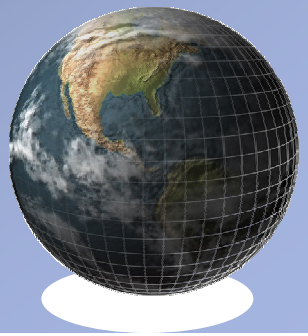

Custom Flow Graph

Custom Flowgraph

Report Parameters

Start: 1/12/2004 07:25:00
End: 1/12/2004 08:25:00
Sample Time: 1
Graph Width: 1

Source Address:
Source Port:
Source Interface:
Destination Address: 137.78.233.0/24
Destination Port: 20
Destination Interface:

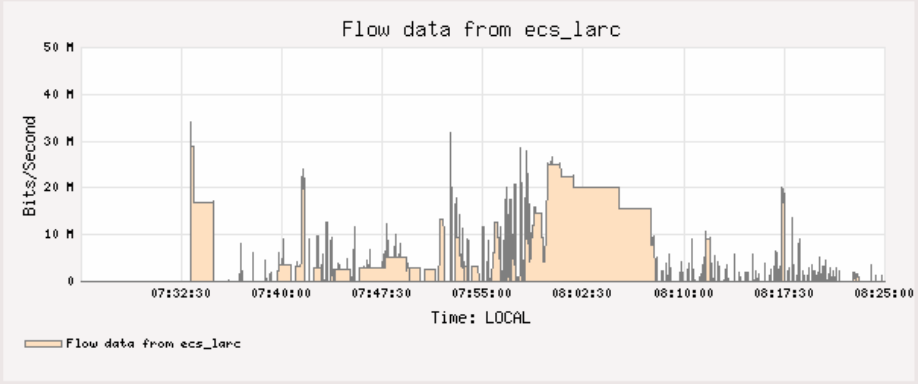

Start	End	Len	Source Host	Port	Destination Host	Port	Total Bytes	Mbps
07:33:06	07:33:18	12.1	129.165.196.2	44174	137.78.233.40	20	19,462,222	12.872
07:33:06	07:34:49	103.4	129.165.196.2	44172	137.78.233.40	20	219,407,519	16.982
07:39:43	07:40:37	54.5	129.165.196.2	51579	137.78.233.40	20	26,069,800	3.825
07:40:55	07:41:39	44.3	129.165.196.2	51579	137.78.233.40	20	17,379,792	3.137
07:41:24	07:41:38	14.2	129.165.196.2	63566	137.78.233.40	20	31,077,552	17.464
07:42:18	07:42:53	34.7	129.165.196.2	51579	137.78.233.40	20	13,034,844	3.002
07:43:57	07:45:04	66.9	129.165.196.2	51579	137.78.233.40	20	21,724,740	2.598
07:45:42	07:49:14	211.5	129.165.196.2	51579	137.78.233.40	20	73,864,116	2.794
07:47:39	07:49:25	106.4	129.165.196.2	56660	137.78.233.40	20	31,787,389	2.389
07:49:30	07:50:16	46.5	129.165.196.2	51579	137.78.233.40	20	17,379,792	2.993

- Examination of FTP transfers between two networks
- One hour time period examined
- Ex.: Used to troubleshoot slow FTPs, and exonerate network

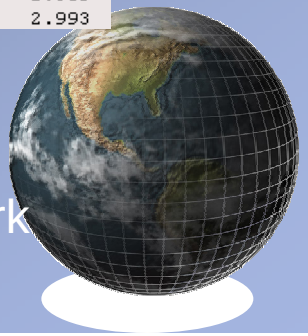

Creating a Custom Flow Report

Passive Monitoring - Custom Reports

Filter Criteria:

Device: ecs_gsfc

Start Date: 4/11/2004 (e.g., 7/17/2003) Start Time: 10:00:00 (e.g., 11:26:00)

End Date: 4/11/2004 (e.g., 7/17/2003) End Time: 14:00:00 (e.g., 11:26:00)

Source IP: (e.g., 192.168.16.0/22) Source Port: Source Interface: Source AS:

Dest IP: (e.g., 0.0.0.0/0) Dest Port: Dest Interface: Dest AS:

Note: a minus sign (-) will negate an entry in the fields above (e.g. -1774 for AS, would mean any AS but 1774)

Report Type:

Statistics: Statistics Reports Printed: Print Reports

Sort Field: 4 Cutoff Lines: 100 Resolve Addresses: Y

Generate Report Reset Values

- **Serves primarily as an HTML front-end to Mark Fullmer's *flow-tools***
- **Permits options of various flow tools reports and statistics**
- **Ability to resolve IP addresses to names**

Custom Flow Report

New Custom Report

Powered by Mark Fullmer's Flow Tools Suite!

Report Parameters:

```
Report: 132 Columns                    Sort Field: n/a        Start: January 13, 2004 12:00:00
Device: ecs_gsfc                     Lines Cutoff: 100          End: January 13, 2004 12:00:05
Source Port:      Destination Port: -5500
Source I/F:      Destination I/F: 12
```

Start	End	Sif	SrcIPaddress	SrcP	DIf	DstIPaddress	DstP	P	Fl	Pkts	Octets
0113.12:00:00.458	0113.12:00:00.574	13	192.168.254.126	54737	12	172.16.76.74	80	6	3	10	1480
0113.11:59:57.706	0113.12:00:03.286	13	192.168.254.126	54729	12	172.16.117.206	80	6	3	25	1773
0113.12:00:00.390	0113.12:00:00.390	13	192.168.254.126	54736	12	10.104.110.153	25	6	2	1	60
0113.12:00:00.458	0113.12:00:00.458	13	192.168.254.126	0	12	10.240.76.74	2048	1	0	1	1500
0113.12:00:00.622	0113.12:00:00.662	13	192.168.220.74	123	12	192.168.193.2	36536	17	0	4	304
0113.12:00:00.622	0113.12:00:00.622	13	192.168.220.2	0	12	192.168.193.2	2048	1	0	1	1500
0113.12:00:00.630	0113.12:00:00.666	13	192.168.220.74	123	12	192.168.193.130	57992	17	0	4	304
0113.12:00:00.818	0113.12:00:00.818	13	192.168.254.14	21	12	172.16.115.87	51934	6	0	1	40
0113.11:59:58.546	0113.12:00:02.294	13	192.168.254.126	49493	12	10.46.247.45	22	6	0	13	1380
0113.12:00:03.558	0113.12:00:03.558	13	192.168.254.14	64460	12	192.168.110.158	4693	6	0	1	41
0113.12:00:03.846	0113.12:00:03.846	13	192.168.220.2	0	12	172.16.200.238	2048	1	0	1	1500
0113.12:00:04.066	0113.12:00:04.066	13	192.168.254.14	35151	12	192.168.110.158	4719	6	0	1	41
0113.11:59:46.510	0113.12:00:04.574	13	192.168.254.126	21482	12	192.168.4.174	21	6	2	3	180

- **Ex.: Produces a '132 Column' report from Flow Tools**
- **This particular report looks at a 5 second period**
- **Background script concatenates files 2 hours beyond specified time**

Tracking NetFlow Impact on Network Resources

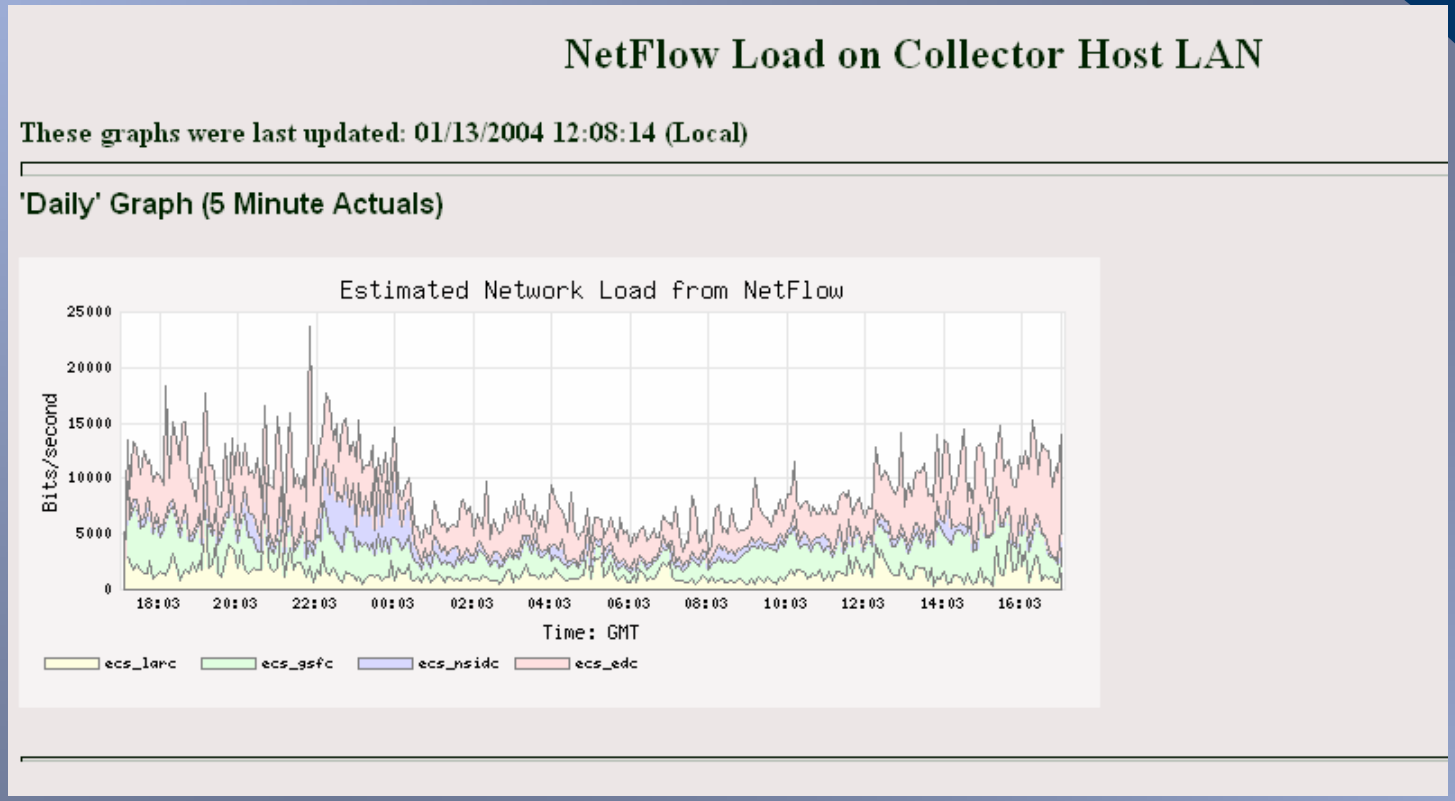

- **NetFlow load is tracked on local LAN**
- **TCPdump data is collected, extrapolated**
- **Graph indicates load on WAN, LAN caused by NetFlow**
- **Rarely more than 15K bits/ second for 4 routers**

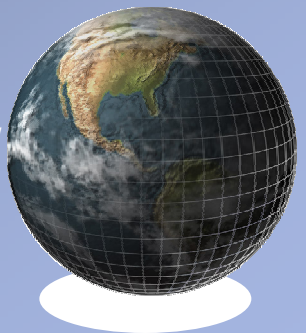

1 Week Round Trip Time Chart

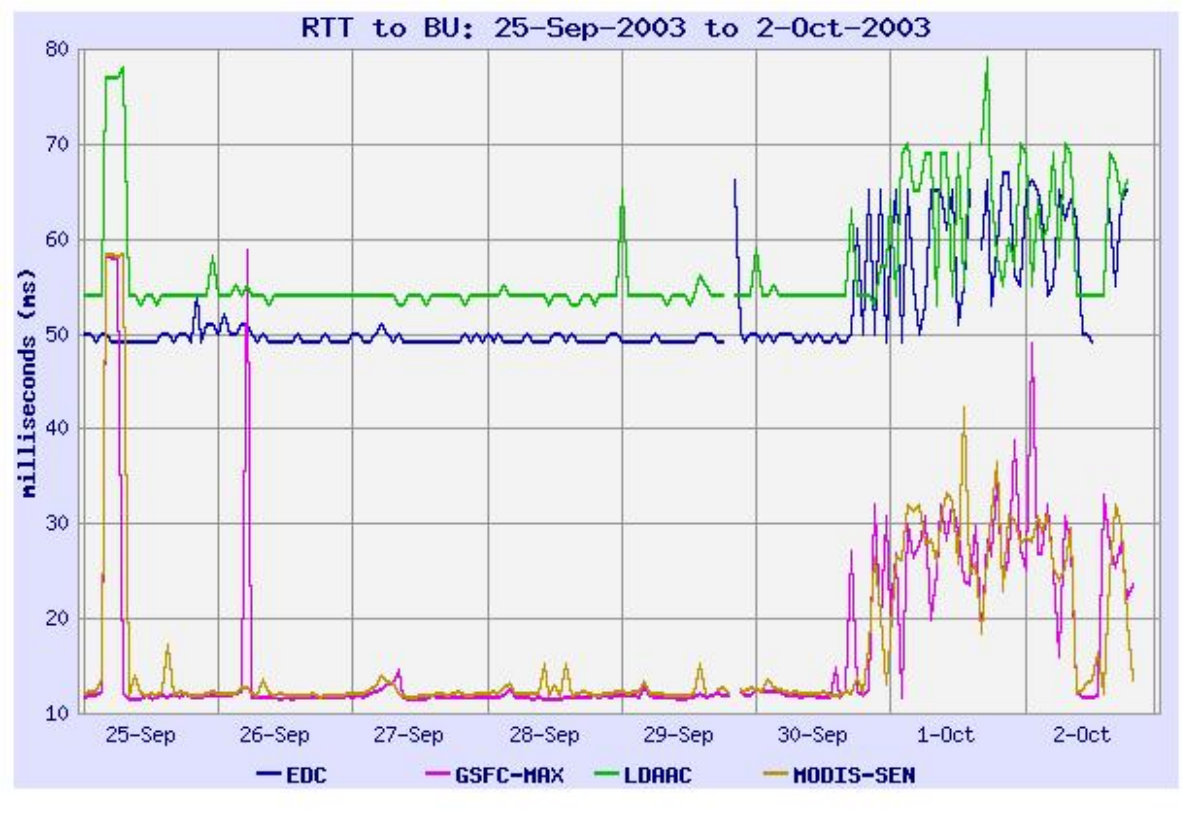

This **Round Trip Time (RTT)** chart shows RTTs between the same hosts and the performance test host at BU. It reveals the benefit of displaying multiple sites on the same chart. In this case, all source-destination pairs experienced significant RTT increase at the same time. This would place the problem closer to the BU end.

1 Week Packet Loss Chart

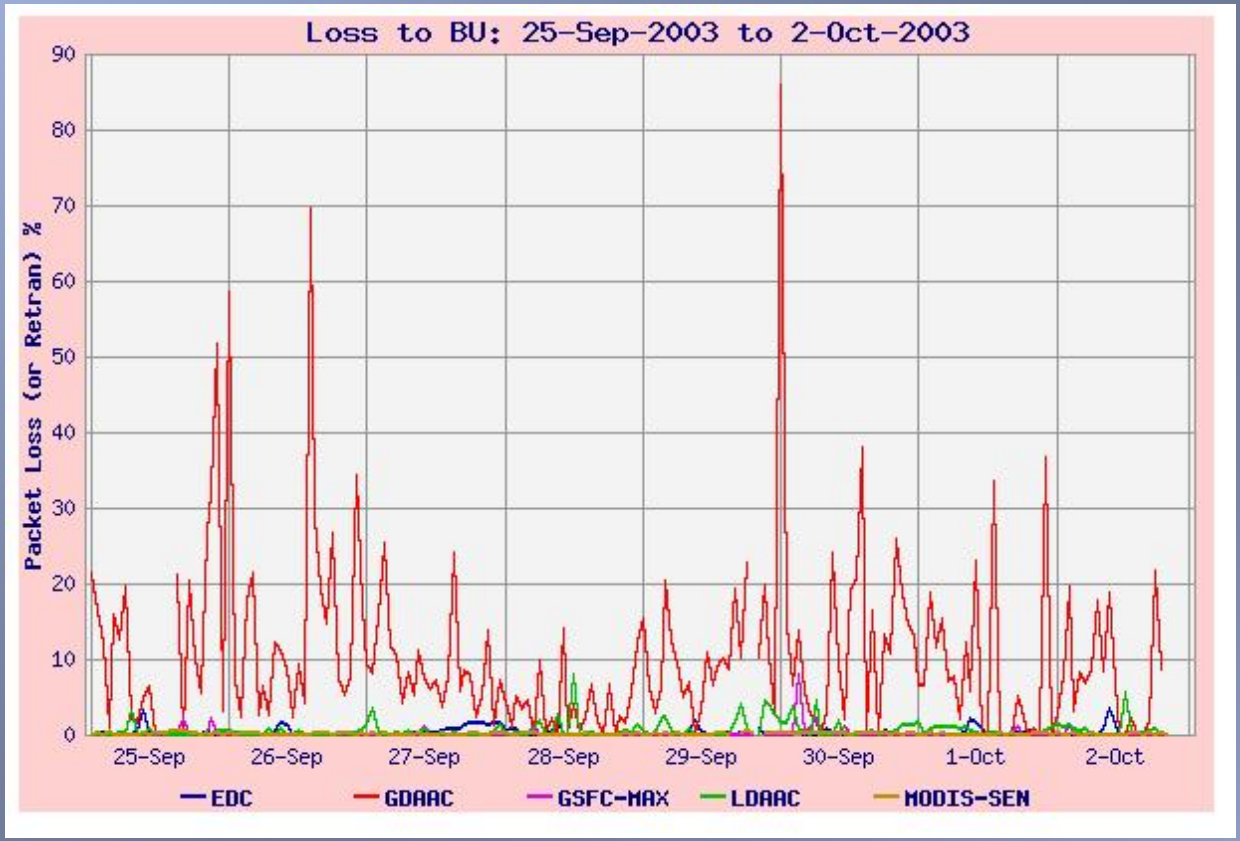

On the other hand, this **Packet-Loss** chart shows low packet-loss to BU from most of the source nodes, but high losses from one source. This would indicate a problem closer to the source node.

Integrated Charts

- The problem: Neither iperf nor MRTG alone is sufficient to characterize the performance of a circuit
  - MRTG will be low if users are idle
  - But Iperf results will appear low if competing with active user flows
- Solution: Add the iperf and MRTG measurements together.
  - But beware: there are some difficulties:
    - Hard to get data for the same time period
      - iperf runs 30 seconds; MRTG taken at 5 mins
  - The measurements are at different protocol levels (layer 2 vs. TCP)
    - Could take a "discount" from the MRTG to account for overhead
- Improved Solution: Add the iperf and applicable Flow data
  - Flow data can be obtained for small time periods
  - But still susceptible to interference

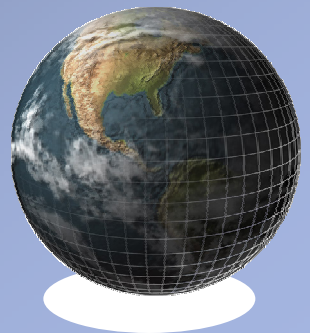

An "Integrated" Chart

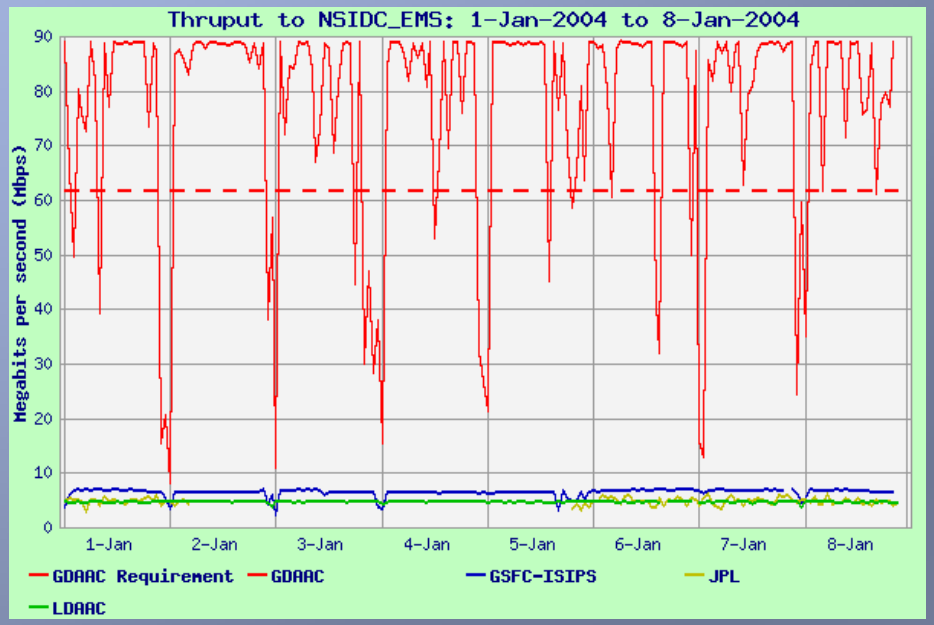

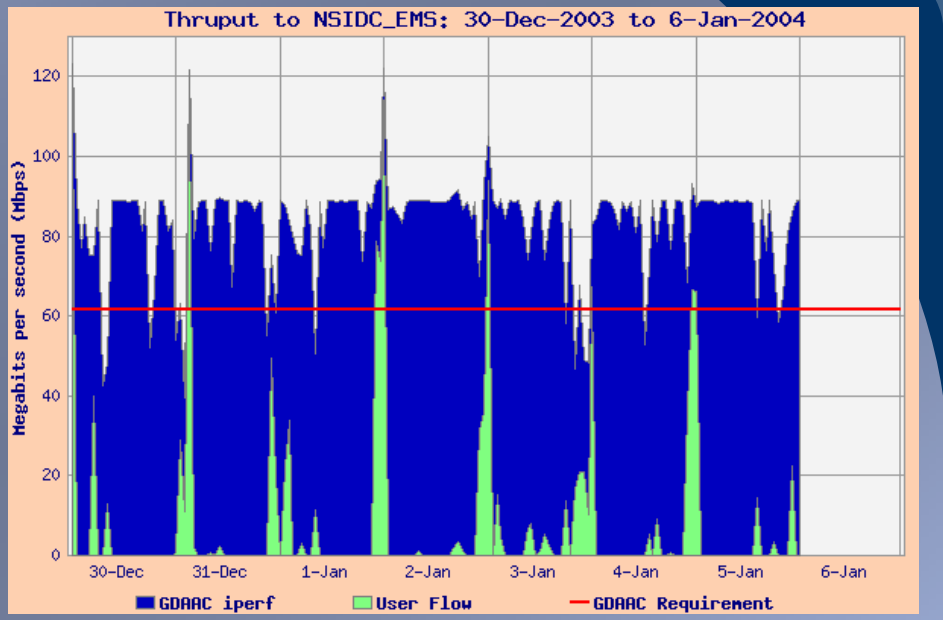

The thruput chart (on the left) shows deep periodic iperf drops (red line), while the integrated chart shows that these low iperf results correspond to high user flows.

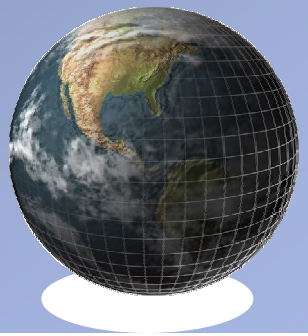

Hourly "Integrated" Chart

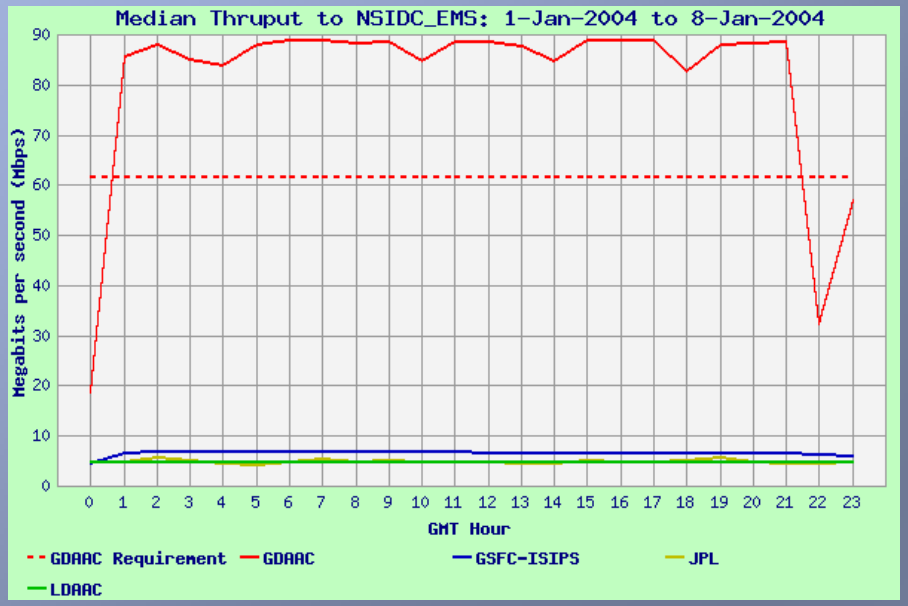

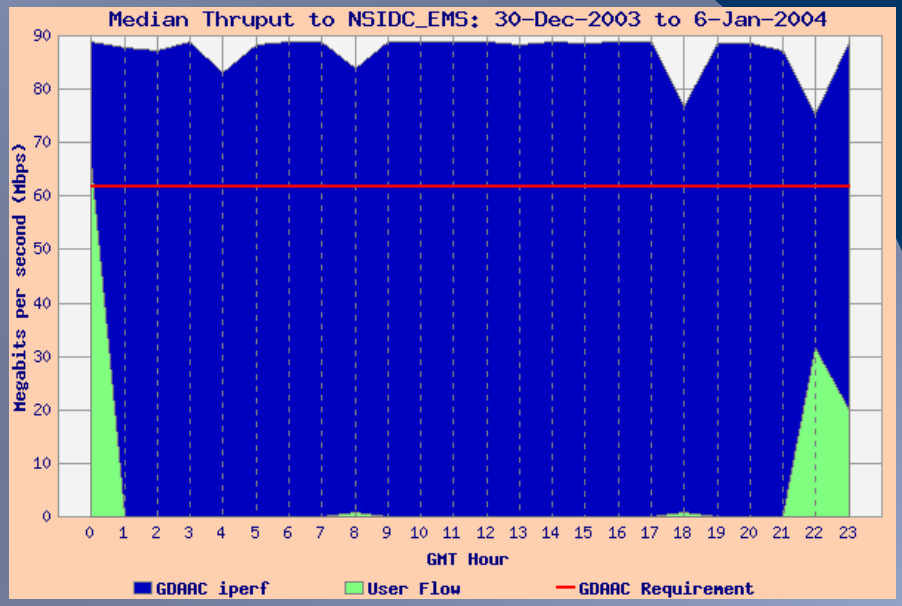

The thruput chart (on the left) shows that iperf drops for 3 hours around midnight (red line), while the integrated chart shows that these hours correspond closely to high user flows.

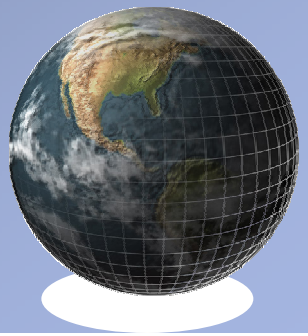

Appendix 1: NetFlow storage requirements

First four routers

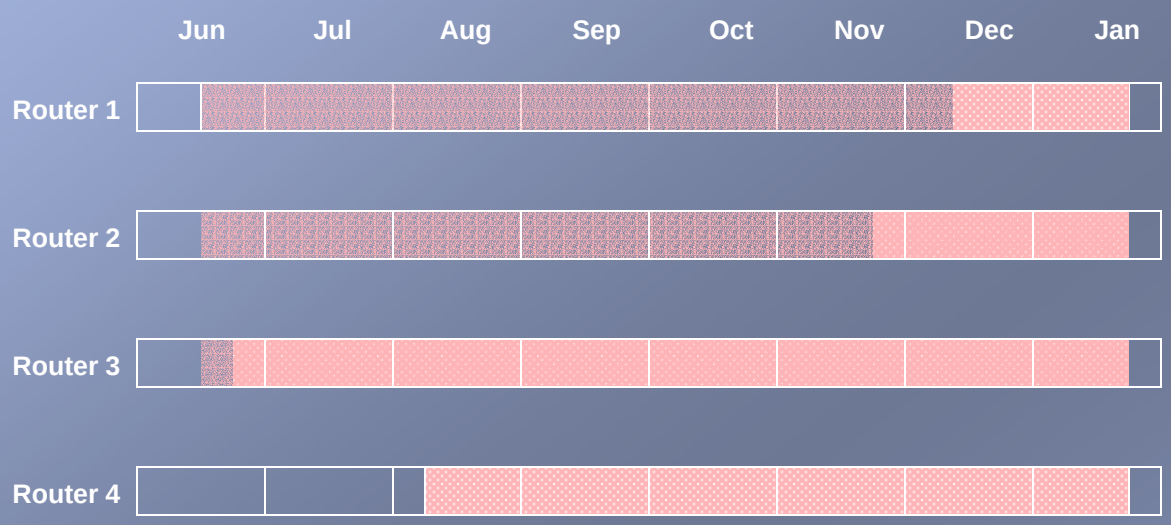

	Tot. Size*	Num. Files**	Avg. File
Router 1	634M	5870	111K
Router 2	641M	8401	78K
Router 3	651M	18727	36K
Router 4	487M	16111	32K

* blocksize = 1048576, includes dirs.
** 1 file = 1 day

Flow tools collection for each router provided with 600M bytes storage. Flow tools will expire older data as necessary to preserve total maximum allocation of 600M bytes.

Need to review interfaces being collected for overlap and reallocate various per router allocations so that time period of retrievable data is roughly equal between all routers.

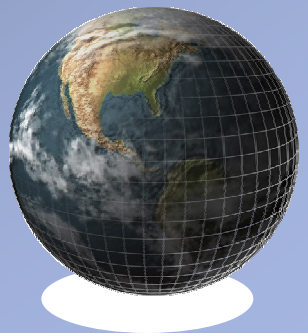

Recent measurement activities: Load Study

Problem: Large science data sets transferred out of GSFC to other NASA sites were arriving corrupted requiring re-sends.

Interesting characteristic: Analysis revealed short portions of data set had had zero's and one's flipped. This was getting past the TCP checksum!

Resolution: removal of intermediate router that was being over-tasked (high CPU utilization)

Byproduct: Netflow generation was not the culprit; very small CPU impact.

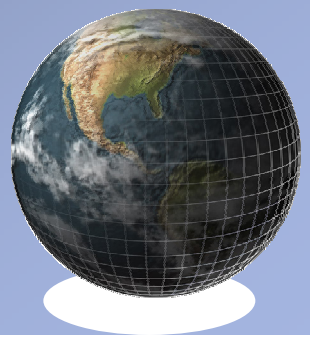

Recent measurement activities: Load Study

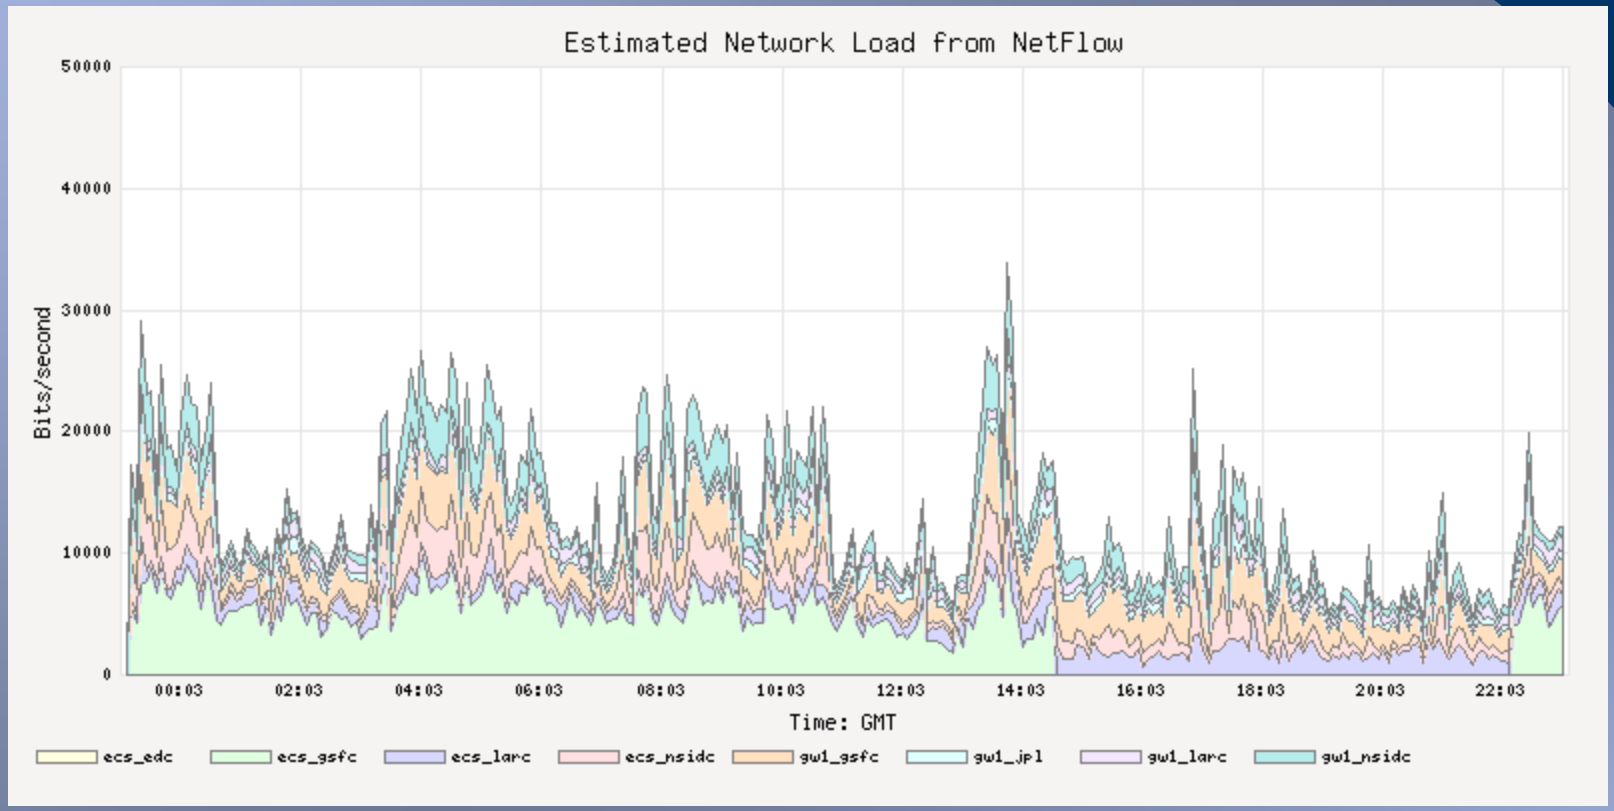

For a test, NetFlow was turned off from 2/12/04 9:30 AM to 2/12/04 5:00 PM
Router: ecs_gsfc Chart shows no data received during period
All times are GMT (-5 EST)

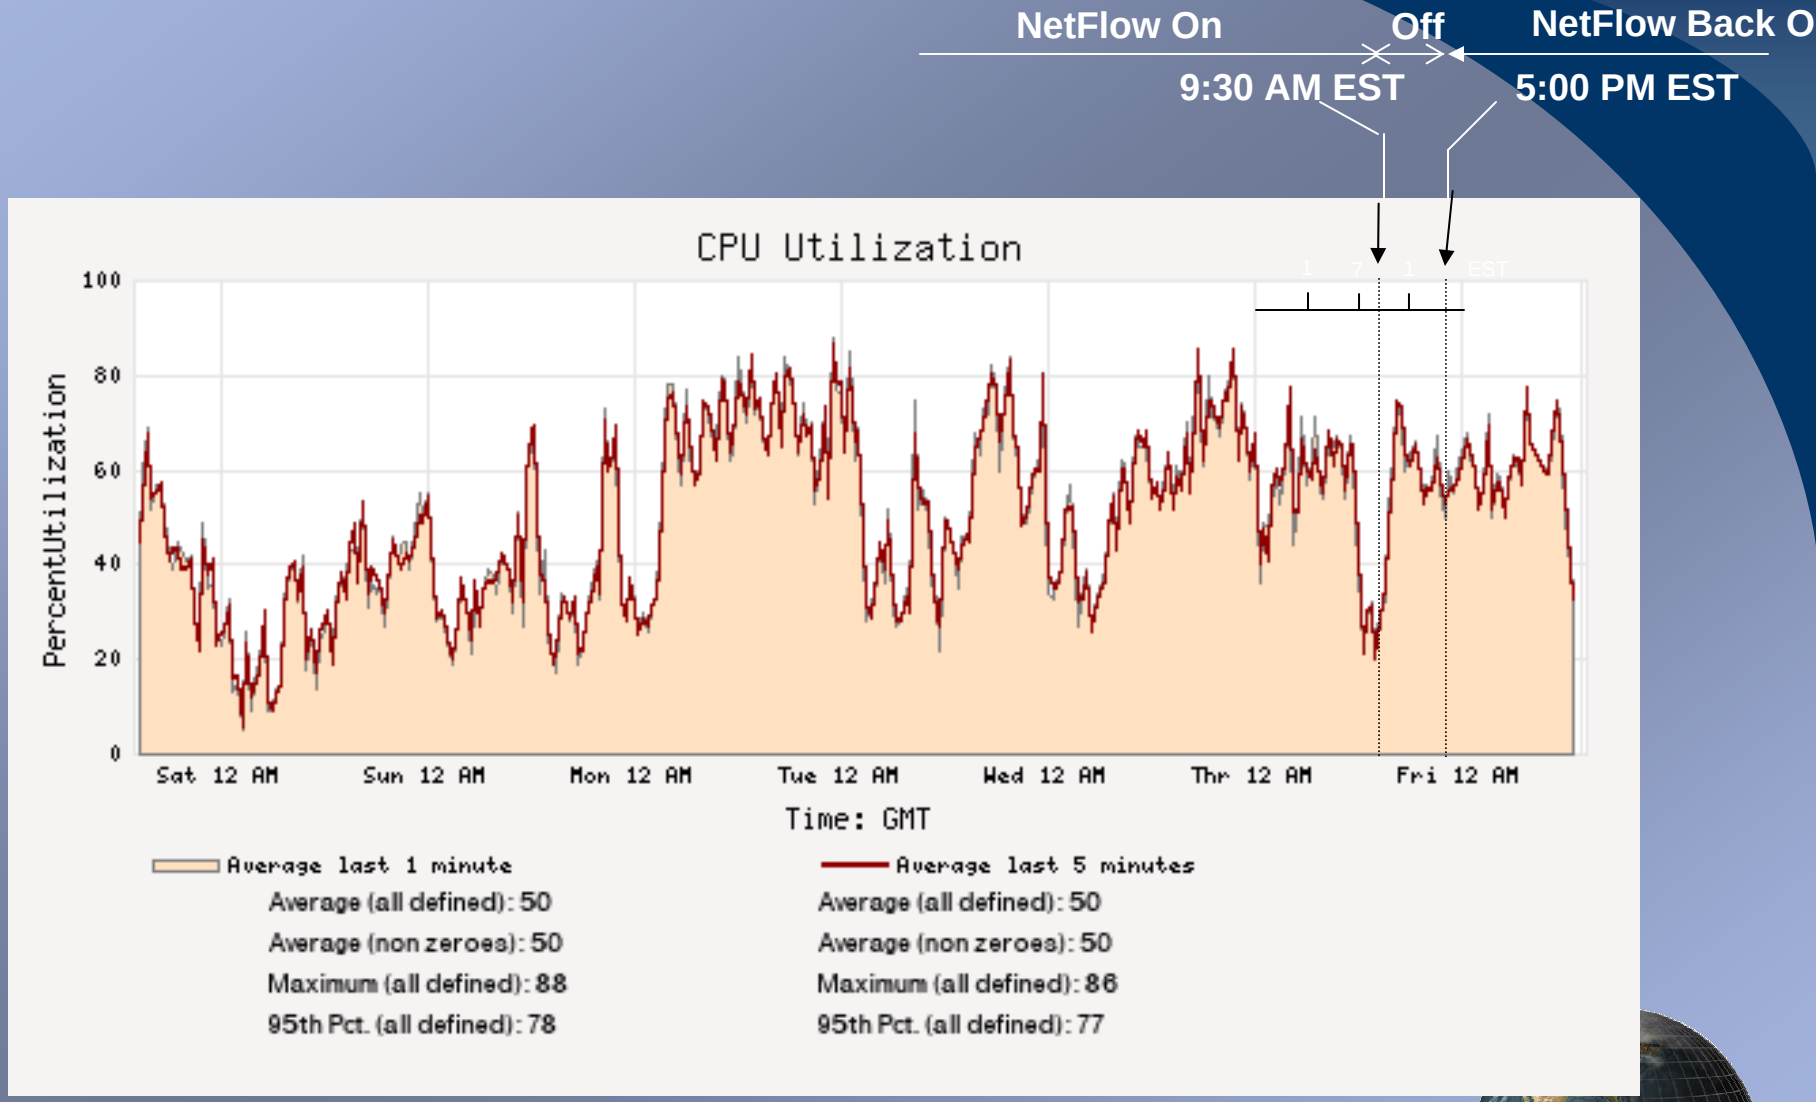

For a test, NetFlow was turned off from 2/12/04 9:30 AM to 2/12/04 5:00 PM Router: ecs_gsfc, 30 Minutes averages

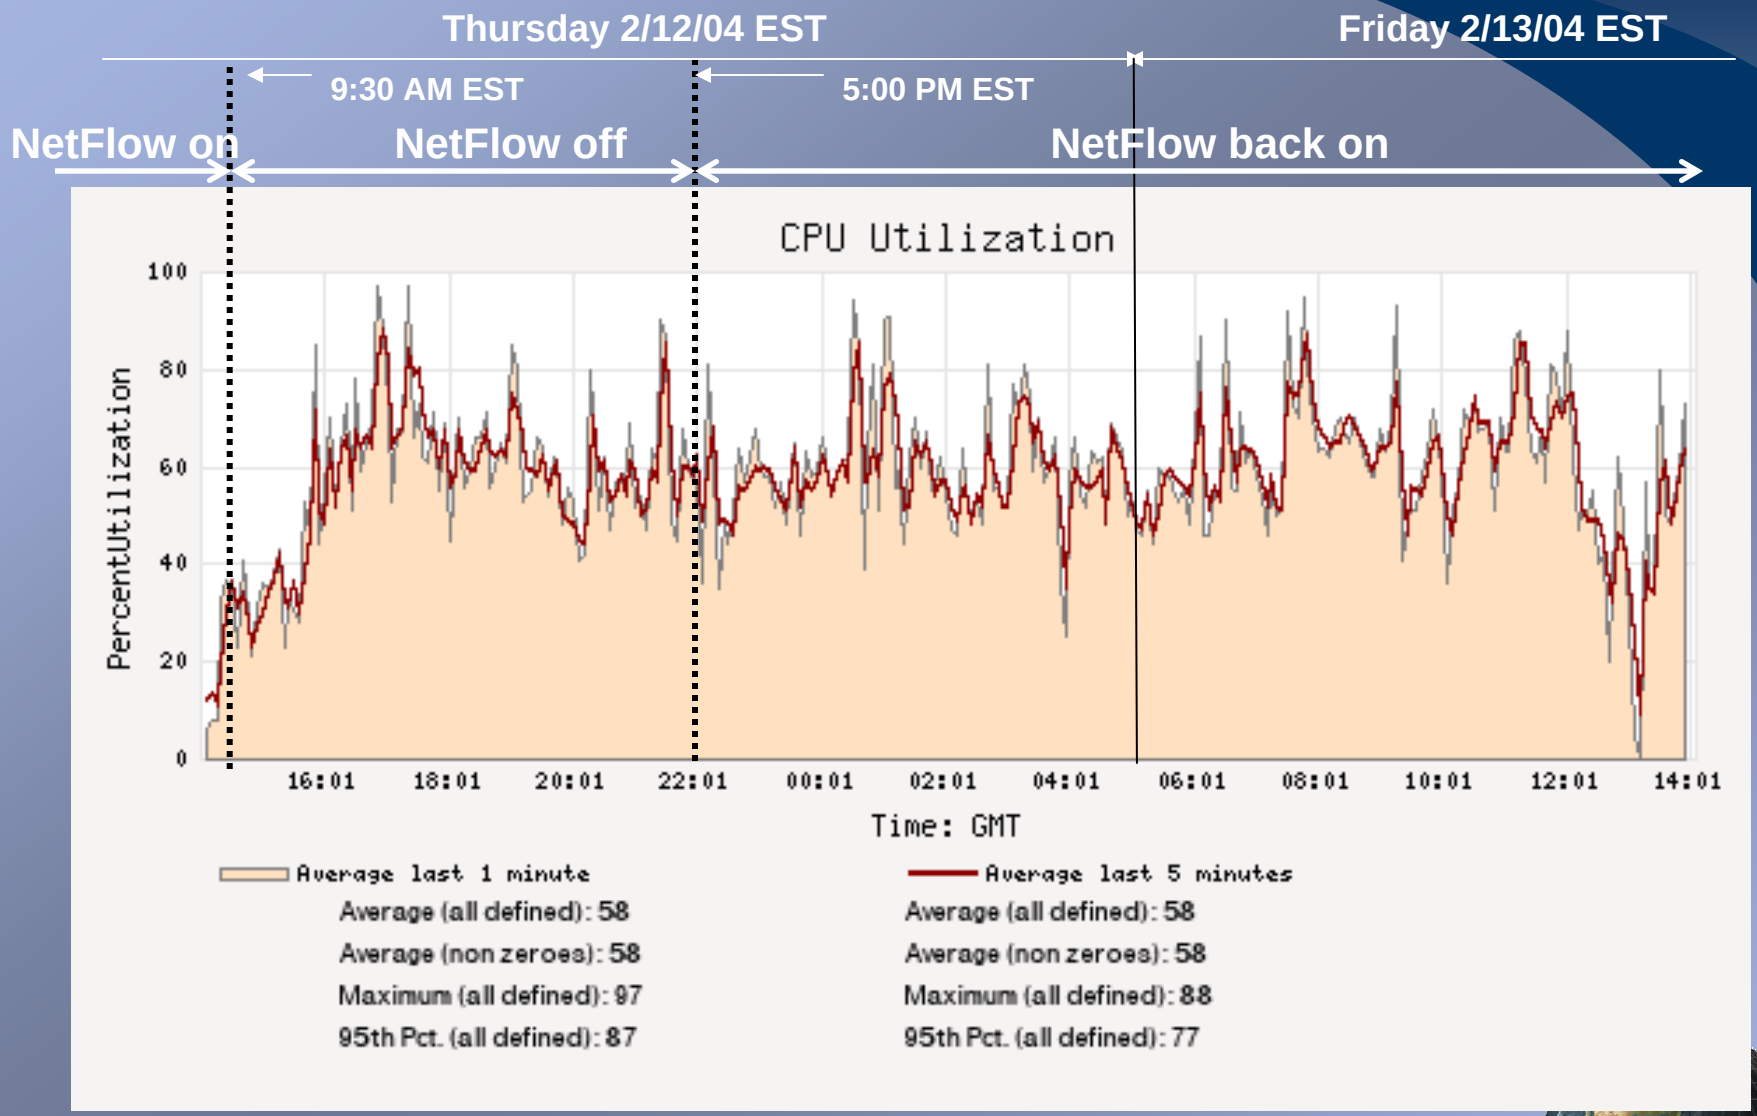

For a test, NetFlow was turned off from 2/12/04 9:30 AM to 2/12/04 5:00 PM
Router: ecs_gsfc 5 minute averages

Recent measurement activities: Sampling

EDC router sampling:	237,717,884,804
x 10:	2,377,178,848,040
GSFC full amount:	2,367,564,379,229
Delta:	9,614,468,811
Pct.	00.41%

Conditions: Relatively fixed traffic patterns, smaller number of flows, large flows.

Note: This activitiy involved the coordination of two separate organizations demonstrating the ease of sharing data (security permitting.)

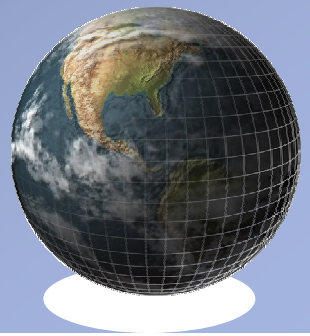

Recent measurement activities: Sampling

```
0330.04:49:50.445 33    129.165.220.71  36583 0     152.61.5.2      65444 6   3  16          18329
0330.04:49:50.235 33    129.165.220.71  36581 0     152.61.5.2      65443 6   3  14          16275
0330.04:49:53.045 34    152.61.5.2      65454 33    129.165.220.71  36593 6   3  26          1360
0330.04:49:53.075 33    129.165.220.71  36593 0     152.61.5.2      65454 6   3  209         306256
0330.04:49:53.145 33    129.165.220.71  21    0     152.61.5.2      65437 6   3  16          2513
0330.04:49:53.105 34    152.61.5.2      65437 33    129.165.220.71  21    6   3  18          1093
0330.04:49:55.785 34    152.61.5.2      38852 33    129.165.220.71  21    6   3  11          618
0330.04:49:55.825 33    129.165.220.71  21    0     152.61.5.2      38852 6   3  13          2146
0330.04:57:58.015 34    152.61.5.2      2331  33    129.165.220.71  21    6   3  0           0
0330.04:51:06.795 34    152.61.5.2      65063 33    129.165.220.71  21    6   3  0           0
0330.04:58:08.205 34    152.61.5.2      2412  33    129.165.220.71  38444 6   3  0           0
0330.04:51:29.925 33    129.165.220.71  21    0     152.61.5.2      65464 6   3  0           0
0330.04:59:49.665 34    152.61.5.2      2893  33    129.165.220.71  21    6   3  0           0
0330.04:51:00.975 34    152.61.5.2      65481 33    129.165.220.71  21    6   3  0           0
0330.05:02:08.575 33    166.61.8.117    2615  0     166.61.8.118    179   6   0  4           246
0330.04:57:58.055 33    129.165.220.71  21    0     152.61.5.2      2331  6   3  0           0
0330.05:02:03.525 34    152.61.5.2      3467  33    129.165.220.71  21    6   3  18          1097
0330.04:58:08.235 33    129.165.220.71  38444 0     152.61.5.2      2412  6   3  0           0
0330.04:51:29.845 33    129.165.220.71  36628 0     152.61.5.2      65502 6   3  0           0
0330.04:53:51.325 34    152.61.5.2      1520  33    129.165.220.71  21    6   3  0           0
0330.04:58:08.645 34    152.61.5.2      2335  33    129.165.220.71  21    6   3  0           0
0330.04:59:55.275 33    129.165.220.71  21    0     152.61.5.2      2896  6   3  0           0
0330.04:56:49.785 34    152.61.5.2      1981  33    129.165.220.71  21    6   2  0           0
0330.05:02:03.555 33    129.165.220.71  21    0     152.61.5.2      3467  6   3  16          2527
```

Removal of two (innocuous) Netflow configuration statements in router corrected problem

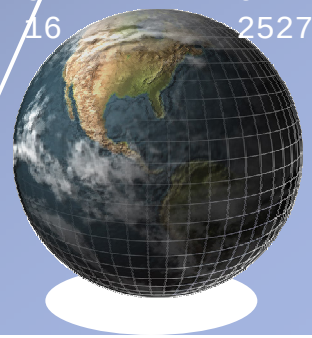

Planned measurement activities

- Recently completed majority of ENSIGHT system
- Continue to investigate improving Integrated active measurements
- Add FlowScan to web-site
- Consider an automated approach to transferring passive measurements into modeling tool
- Use performance tools to analyse EOS network for major transition