

Creating UNIX Network Usage Charts and Displaying Them on an Intranet

Stephen B Taubman, Federal Reserve Board, Washington, DC
David Kelly, Federal National Mortgage Corporation, Washington, DC

Abstract

Our UNIX Network consists of several sub-area networks (SANs) which are linked together via a common backbone. This paper will describe how we produce charts which display how the computing resources of each SAN are being used. We display the mean and 90th percentile for five measurements of computer usage. The charts are created as gif files and displayed on the Federal Reserve Board's local intranet using HTML programs and the Netscape Web Browser. All work was done on the Sun Solaris 2 operating system using SAS Release 6.11. BASE SAS and SAS/GRAPH as well as the SAS macro facility were used. The raw data is obtained from software supplied by Sun Microsystems. A novice SAS user should understand this paper.

The Raw Data

The raw data is obtained from software supplied by Sun Microsystems. Monthly raw data files consist of one observation taken every eight minutes for each network machine. Each observation contains 78 variables. A SAS dataset is created from the raw data. The variables in the SAS dataset consist of the machine name and the following 5 usage measurements of interest:

1. CPU usage as a percentage of capacity.
2. One-minute load average, which is the number of processes running in an average one-minute interval.
3. Total packets transferred, which measures the amount of network data transferred to and from each machine.
4. Percent of network collisions, which is the number of network packets that collide as a percentage of total network traffic for that machine.
5. Swap-outs, which in SUNOS is the amount of system swapping, generally between RAM and the swap file system in order to free RAM space.

Some inferences about network machines that can be made by observing these variables are:

* A machine with a high load average and low CPU usage might have a lot of dormant processes or processes waiting for input.

* However, a high load average together with high CPU usage may indicate an over-used machine.

* A machine reporting a large number of packets transferred either may have too many heavily used file systems installed or could be making too much use of remote file systems.

* A high percent of net collisions combined with a low value for packets transferred suggests a problem with hardware or cabling.

Bar Charts for an Individual SAN.

Our network consists of several sub-area networks (SANs) which are linked together via a common backbone. Each SAN consists one or more machines. For a given SAN we compare each machine in the SAN to the SAN average. We are interested in two statistics for each of the five measurements or variables, the mean and the 90th percentile. For example, if the 90th percentile for % CPU usage is 40 percent, then that machine operated at 40 percent or less of CPU capacity 90 percent of the time. It can also be said that the machine operated at 40 percent or more of CPU capacity only 10 percent of the time. The 90th percentile statistic provides a sense of the peak load whereas the mean provides the average load. PROC UNIVARIATE is used to compute the 90th percentile and PROC MEANS calculates the means. PROC GCHART, with the HBAR option, draws the graphs. A different graph is produced for each variable and each SAN. For comparison purposes, these individual SAN charts also display a vertical bar representing a network-wide mean for the given variable. An example of a Bar Chart for an Individual SAN is shown in Figure 1. The machines in the SAN are represented by M1 through M8.

Bar Charts Comparing SANs

The second type of bar chart shows the mean values for each SAN. One chart is produced for each variable. The bars are arranged with the highest values on the top of the page. The number of machines in each SAN is also displayed. An example of a bar chart comparing SANs is shown in Figure 2. The SAN names are S01 through S16.

Displaying the Network Usage Charts on the Local Internet

All charts are created as a graphics stream file containing color gif images. This is a new feature of UNIX SAS Release 6.11. Creating a gif file is described on page 33 of the manual 'SAS/GRAPH Image Extensions'. The GOPTIONS statement used was

```
goptions reset=global device=imggif gsfname=gsasfile cback=white ctext=black;
```

where *gsasfile* is the filename reference pointing to the UNIX file name of the created gif file.

The gif files are ultimately stored in directories and then referenced by an HTML file and displayed with the web browser.

The following sequence will display a network usage chart:

1. A user invokes the web browser Netscape on our local intranet, accesses a page containing UNIX network information, and selects the 'Network Usage Chart' button.

2. The next page has three buttons.

- individual SAN plots
- plots of all SANs
- how the plots are created

The HTML code for this page contains these records:

```
<A HREF="file://directory_indiv">individual SAN plots</A>  
<A HREF="file://directory_all"> plots of all SANs </A>  
<A HREF=info.html> how the plots are created </A>
```

where

directory_indiv is the name of a directory which contains one subdirectory for each SAN. These subdirectories are named after the SAN and each subdirectory contains 5 gif files (one for each variable).

directory_all is the name of the directory containing the 5 gif files (one for each variable) representing the plots comparing all the SANs.

info.html is the name of the file containing information about the plots,

3a. If the user selects the 'individual SAN plots' button, the next page will list the SAN names, each of which is a directory name in directory_indiv. The top of the page looks like

- s01
- s02
- s03

The user then selects one of these SAN name buttons. If s01 is chosen for example, the subsequent page lists the files in directory_indiv/s01. These files (identical names within each SAN subdirectory) are the 5 gif files, one for each variable of interest. They are:

- cpu.gif
- loadav.gif
- ncoll.gif
- swpout.gif
- totalp.gif

If any of these buttons is selected, the corresponding graph of all machines in SAN s01 will be displayed for that variable.

3b. If the user selects the 'plots of all sans' button, the next page lists the files in the directory_all directory. In this case it contains the 5 gif files and looks like:

- cpu.gif
- loadav.gif
- ncoll.gif
- swpout.gif
- totalp.gif

Each selection will display the corresponding graph in Netscape. An HTML program need not be written for this page. It will be generated by Netscape.

3c. If the users selects the 'how the plots are created' button, a page is displayed that lists the info.html file containing information about creating the plots.

Acknowledgement

The authors wish to thank Ray Board and Bruce Gilson for comments throughout the development of this paper. The authors can be contacted at :

Federal Reserve Board
Mail Stop 096
Washington, DC 20551
(202)452-2517
m1sbt00@frb.gov

FIGURE 1
PERCENT OF NET COLLISIONS SAN S13




FIGURE 2
ONE - MINUTE LOAD AVERAGE BY SAN

