

[close window](#)

e-Newsletter Exclusive

[Print](#) 

Tools You Can Use: ncp and nweb

November 2008 | by [Jaqui Lynch](#)

Many of us are familiar with the nmon performance tool available for AIX*. This tool is primarily used for performance analysis and as a basis for tuning. It has also been used as a basis for some capacity planning. However, Nigel Griffiths, IBM Power Systems Technical Support, Europe, has another free tool that collects and displays workload-based information for System p* servers: ncp. The information collected and displayed by this free tool can be invaluable when determining which servers are over- or under-utilized.

Before using the tool, it's important to understand the tool is designed for capacity planning and consolidation studies. The "ncp" tool is designed to gather data suitable for capacity planning, and can then be used to analyze and visualize the data. The data is gathered as an hourly average each day—which is ideal for planning purposes, but not suitable for performance-tuning studies. Performance studies require much more granular data, such as that provided by nmon or other tools.

There are two components to the tool: ncp and nweb. The ncp tool collects the data and nweb is a simple Web server used to display the data. If the machine already has a web server installed, then you can use that rather than nweb. You can also use NFS or some other mechanism to export the data daily to a Web server on a different system. In my example, I installed nweb on the system I was monitoring to see what the tool provided. Version 2.1 (released in March 2004) reports either number of cores for dedicated LPARs, or number of virtual processors for those in the shared processor pool. I had ent=0.50 and VP=2 with simultaneous multithreading (SMT) on, and the configuration details showed this as two processors. So this is a great tool to use on dedicated servers and LPARs with dedicated cores, when you are trying to do consolidation studies. However, it probably isn't suitable for LPARs already in the shared processor pool.

The nweb tool is used by ncp as a small and safe Web server to deliver its graphs and data. Griffiths supplies nweb via "sample C code." It delivers static pages only to illustrate how a Web server works. It has a fixed number of file types and logs errors. It aims to be secure by virtue of only being 200 lines of C code that have been thoroughly reviewed and are publicly available. If you are interested in the nweb code, design and systems programming find more at: <http://www.ibm.com/developerworks/systems/library/es-nweb.html>

Specific information the tool provides includes machine configuration, average filesystem use, CPU utilization (user + system) broken out by day of the week and time of day, and disk I/O broken out by day of the week and time of day. The graphs are color-coded and you can click on any individual day and time to get more historic data. For example, you can look at overall CPU average for each day in the data; you can look at each day of the week and what the busy hours are for those days; and you can look at CPU averages for each day, split out by day of the week. Additionally, you can click on a highlighted box and get specific data. I clicked 2 p.m. Tuesday at and got the trend graph for 2 p.m. each Tuesday in the data set.

Viewing a graph like this over a long period of time allows you to see if there are trends, such as an increase over time on certain days. The color coding is also extremely useful; it makes it very easy to pick out periods worthy of deeper analysis. Disk I/O can also be analyzed in a similar manner to CPU utilization. Examples of a typical capacity planning session including graphs can be viewed at:

http://www.ibm.com/developerworks/eserver/articles/free_tools/index.html#N1021D

The tool itself is very easy to install and use. I chose the simplest of the options as follows:

1. Create a filesystem called /usr/local/perf/ncp. I always use separate filesystems when I am gathering data over time, as I do not want the data to go into /var or /usr—where it might fill up the filesystem and cause problems. By putting it into its own filesystem, the only thing affected is that one filesystem.
2. Create a user and group, both called ncp, and set them to own the filesystem—which becomes the home filesystem for userid ncp.
3. Download the zip file, unzip it and upload the tar file to the /usr/local/perf/ncp directory—where it was then untarred and ownership of all files was changed to the ncp userid and group.
4. Run ./ncp -? to check if the software is functioning.

5. Run ls -l to ensure all files (especially the html ones) are present. The output should look like this:

```
total 208
-rw-r--r-- 1 ncp ncp 13545 Mar 14 2004 docs.html
-rwxr--r-x 1 ncp ncp 49 Mar 14 2004 dot.gif
-rw-r--r-- 1 ncp ncp 1655 Mar 14 2004 index.html
-rw-r--r-- 1 ncp ncp 5964 Apr 29 2004 install.html
-rw-r--r-- 1 ncp ncp 5694 Apr 29 2004 install.txt
-rwxr-xr-x 1 ncp ncp 48769 May 04 2004 ncp -rwxr-xr-x 1 ncp
ncp 8880 May 04 2004 nweb -rw-r--r-- 1 ncp ncp 136 May 04 2004
real.html
```

6. The next step is to create the subdirectories necessary by running ./ncp -x . (The . is necessary to create the files in the current directory.)

7. The directory list should now look like this:

```
total 504
-rw-r--r-- 1 root system 144 Nov 02 17:25 crontab.sample
-rw-r--r-- 1 root system 32390 Nov 02 17:25
database_sample.cpu
-rw-r--r-- 1 root system 22578 Nov 02 17:25 database_sample.df
-rw-r--r-- 1 root system 42450 Nov 02 17:25 database_sample.io
drwxr--r-x 2 root system 256 Nov 02 17:25 df
-rw-r--r-- 1 ncp ncp 13545 Mar 14 2004 docs.html
-rwxr--r-x 7 ncp ncp 49 Mar 14 2004 dot.gif
-rw-r--r-- 1 ncp ncp 1655 Mar 14 2004 index.html
-rw-r--r-- 1 ncp ncp 5964 Apr 29 2004 install.html
-rw-r--r-- 1 ncp ncp 5694 Apr 29 2004 install.txt
drwxr--r-x 2 root system 256 Nov 02 17:25 iostat
-rwxr-xr-x 1 ncp ncp 48769 May 04 2004 ncp
-rwxr-xr-x 1 ncp ncp 8880 May 04 2004 nweb
drwxr--r-x 4 root system 256 Nov 02 17:25 real
```

```
-rw-r--r-- 1 ncp ncp 136 May 04 2004 real.html
drwxr--r-x 4 root system 4096 Nov 02 17:25 sample
-rw-r--r-- 1 root system 40289 Nov 02 17:25 sample.html
drwxr--r-x 2 root system 256 Nov 02 17:25 vmstat
```

8. At this point, the programs are ready to start collecting data once the correct cron entries are set up. Using crontab -e (as userid ncp) add the following two entries:

```
1 0 * * * /usr/local/perf/ncp/ncp -c /usr/local/perf/ncp
0 1 * * * /usr/local/perf/ncp/ncp -g /usr/local/perf/ncp
```

The first entry tells ncp to collect the data and the second entry tells it to produce the graphs each night. After a few weeks, you should have enough data to be useful for trend analysis.

9. To view the data you have several choices. The simplest one is to use nweb. Since the data created was in /usr/local/perf/ncp, then nweb would be started as follows:

```
cd /usr/local/perf/ncp
./nweb 8081 /usr/local/perf/ncp
```

The above line starts the Web server on port 8081 but you can pick any unused port.

Now point your browser to: <http://systemname:8081/index.html>

You can click on the documentation, installation information or the sample graphs to get an idea of how it all works. To review your own data, click on "Capacity Planning from Collected Data." From here you can click on entries of interest and get some idea of how the LPAR or server is performing. The graphs provide a simple view that makes it easy to spot trends by day and time. It's not a prediction tool, but it can be used to do pick up basic trending by simply looking at the graphs.

More detailed information can be found in Nigel Griffith's article at Developerworks:

http://www.ibm.com/developerworks/eserver/articles/free_tools/index.html The tool can be downloaded from the link above after agreeing to the terms and conditions.

The ncp tool is a useful for looking at dedicated servers and LPARs that you are considering virtualizing or consolidating onto other servers. From the data produced, you can determine the average utilization by day of the week and hour of the day. This can be used to figure out what to combine and how much resource is needed. It's not as useful for LPARs already in the shared processor pool. Along with nmon, it's a good tool to have in your toolkit, but keep in mind it needs to gather data over time (in terms of weeks or months) in order to be fully useful.

IBM Systems Magazine is a trademark of International Business Machines Corporation. The editorial content of IBM Systems Magazine is placed on this website by MSP TechMedia under license from International Business Machines Corporation.

©2009 MSP Communications, Inc. All rights reserved.