



## NEMo User Guide

© 2007 Gear6

## Introduction

Document version 1.2, last updated 1/22/07

### About NEMo

NEMo is a software-based tool that enables data center and storage managers to identify I/O bottlenecks and assess the benefits of storage acceleration. Using NEMo results, Gear6 can gauge how well the CACHEfx appliance will accelerate bottlenecked I/O traffic. NEMo is focused on NFS v3 environments.

### How it works

NEMo consists of two components: a network data (NFS data) capture component (collector.pl script) and an offline analysis component of the captured log data.

### Capture

NEMo uses tcpdump, a utility included in Linux and Unix OS distributions to capture network activity for analysis. NEMo operates in a two step process. First, it initiates the tcpdump command through a PERL script with a set of configurable parameters. Next, it retains only file attributes and strips out all file data to ensure security and privacy. NEMo also uses nfsstat, another standard Linux/Unix tool.

### Analysis

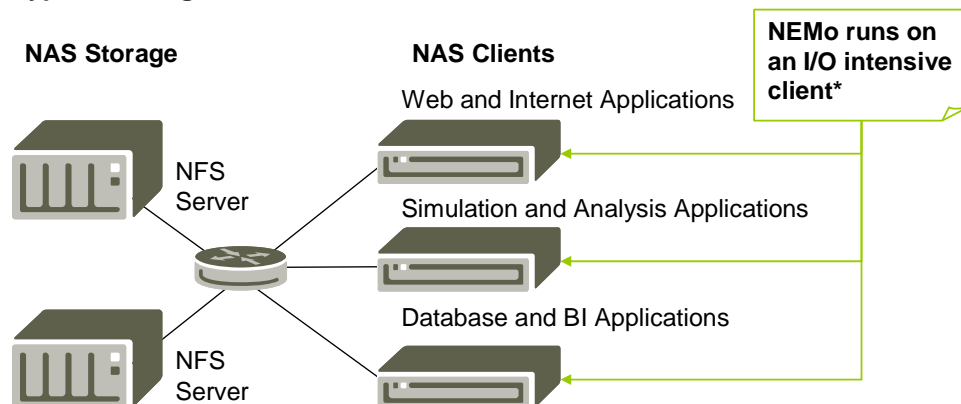
The log data collected in step one is manually uploaded to Gear6 for off-line analysis. The resulting report succinctly describes the observed I/O pattern and predicts the benefits from the CACHEfx appliance.

### Security and Privacy

NEMo collects file attribute information from NFS commands. NEMo DOES NOT RETAIN OR TRANSMIT SPECIFIC INFORMATION FROM WITHIN THE FILES. All file data other than top level attributes is systematically removed to ensure absolute security and privacy.

Additionally, the NEMo script Collector.pl is in clear text and can be examined to determine that it does not violate any security requirements.

### Typical Configuration



## Running NEMo and submitting log data

### Obtaining NEMo

NEMo is available at:  
[www.gear6.com/NEMo](http://www.gear6.com/NEMo)

### Running NEMo

Run NEMo on a Linux or UNIX server machine that hosts an I/O intensive application and uses NAS/NFS based networked storage.

1. As root user, make a directory in a location of your choice and unpack NEMo in it
2. Again as root user, create a directory on a local partition (cannot be an NFS mounted partition) that has at least as much storage as you intend to capture (see NEMo usage in Appendix 2 below), typically, 100MB-200MB.
3. Invoke collector.pl script as described in Appendix 1

### Log submission

1. Go to the directory containing the log data
2. Use gzip to pack/compress the log file(s) into one compressed archive file
3. FTP the log file using the FTP information provided by your Gear6 contact.  
Be sure to notify Gear6 that the file has been uploaded.

### Analysis Results and Report

1. Following your log submission, a Gear6 representative will contact you regarding your report.

### NEMo Support

If you have questions or issues with installing or using NEMo, please email your questions to [support@gear6.com](mailto:support@gear6.com) with Subject line: "NEMo Support". In the body please include details on:

- Contact information
- How you obtained NEMo (via email from gear6 or web download from gear6.com)
- NEMo version info
- Relevant environment details such as hardware architecture and OS version, tcpdump version, NFS version and NAS storage being used (NetApp filers, EMC Celerra, etc)

## APPENDIX 1: NEMo Usage and tcpdump Util Details

### NEMo Usage

```
[root@homer Nemo]$ ./nemo-collect.pl --help
Usage
./nemo-collect.pl
    --ifc NetworkInterface
    --ip FilerIPAddr
    --tcpdump tcpdump Path
    --nfsstat NfsstatPath
    [--dir OutputDir]
    [--duration Duration(Min)]
    [--maxdata MaxDataSize(MB)]
    [--help]
```

**--ifc:** Interface to listen on. Ex. eth0

**--ip:** Filer IP address.

**--tcpdump:** Pathname of local tcpdump utility. Ex. /usr/sbin/tcpdump

**--nfsstat:** Pathname of local nfsstat utility. Ex. /usr/sbin/nfsstat

**--dir:** Directory to which all output files are sent. Default is /tmp

**--duration:** Max number of minutes to run collector. Default is 60 min

**--maxdata:** Maximum quantity of data gathered by collector. Default is 100MB

### Check version of tcpdump on your Linux/UNIX OS

```
[root@homer]# /usr/sbin/tcpdump --version
tcpdump version 3.8
libpcap version 0.8.3
Usage: tcpdump [-aAdDeflLnNOPqRStuUvX] [-c count] [ -C file_size ]
[ -E algo:secret ] [ -F file ] [ -i interface ] [ -M secret ]
[ -r file ] [ -s snaplen ] [ -T type ] [ -w file ] [ -W filecount ]
[ -y datalinktype ] [ -Z user ]
[ expression ]
```

### Tcpdump invocation details within Collector.pl:

As you can see below, we only invoke tcpdump with a restricted set of options that does not give the NEMo PERL script collector.pl any visibility into the actual payload data of the tcp packets:

```
# tcpdump -vvv -n -tt -u -s 1500 -i eth0 port 2049 src or dst 10.1.1.23
```

The tcpdump options string used above in collector.pl are explained below:

**-vvv:** Detailed NFS trace output

**-n:** Host addresses rather than host names (this is for ease of parsing. The host addresses are subsequently stripped from the collector log output entries. The filer IP address submitted as an argument is included in the log file name. This helps to identify and distinguish between datasets collected for different file servers)

- tt:** Un-formatted timestamps per record entry
- u:** Un-decoded file handles
- s 1500:** captures whole IP packets (note: under the options we use for tcpdump, the file data is never seen by the nemo-collect.pl script)
- i eth0:** Specifies the Ethernet network interface on the NFS client machine for tcpdump to listen on
- port 2049:** Specifies that tcpdump should listen on port 2049 (NFS traffic)
- src or dst 10.1.1.23:** Specifies the source or destination IP address of the NFS file server (ex: NetApp filer)

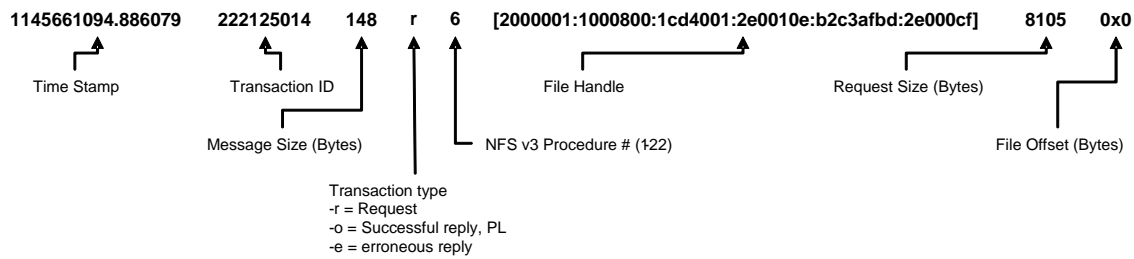
**Sample output of tcpdump and collector.pl log file:**

Output of “tcpdump -vvv -n -tt -u -s 1500 -i eth0 port 2049 src or dst 10.1.1.23”:

**Output of NEMO-invoked tcpdump**

```
1145661094.886079 IP (tos 0x0, ttl 64, id 17444, offset 0, flags [DF], proto 6, length: 200) 10.1.2.242. 222125014 >
10.1.2.200.2049: 148 read fh[2000001:1000800:1cd4001:2e0010e:b2c3afbd:2e000cf] 8105 bytes @ 0xf000
```

**Output of collector.pl to log file**



## **APPENDIX 2: NEMo Frequently Asked Questions**

### **Where do I run NEMo?**

NEMo is run on any application server that has an I/O intensive application accessing an NFS filer.

### **Will NEMo disrupt my applications?**

No, it does not disrupt your applications. tcpdump runs in the background with minimal impact on any running applications.

### **Does NEMo affect the performance of my running applications?**

NEMo has little impact to running applications. Overhead ranges from 5-10%.

### **What specific attributes does NEMo capture?**

NEMo captures NFS operations such as READ, WRITE, ACCESS, and GETATTR, and file specific attributes such as File Handle, File Size, File Type, and Access and Modification Times.

### **Where does NEMo put the log files?**

NEMo will place the log files in a user-specified directory.

### **How big are the log files? Are there more than one?**

Depending on the duration of the tcpdump sampling, and the amount of data, NEMo may create multiple log files. Each log file has a maximum size of approximately 100MB.

### **Why should I run NEMo?**

NEMo will allow Gear6 to assess and share the specific benefits of using the CACHEfx appliance in your existing storage environment. We will also provide you with a detailed report outlining a recommended configuration and identifying the expected return on investment.

### **Can I run NEMo for my Fibre Channel or iSCSI Storage?**

Not yet. NEMo currently supports NFS-based storage.

### **Can I run NEMo on multiple NFS clients simultaneously?**

Not yet. NEMo runs on a single client. If the client chose is representative of the entire client group, NEMo results can be extrapolated accordingly.

### **How do I get started?**

Download NEMo from the Gear6 website and run it on any I/O intensive application server. If you would like assistance, contact Gear6 and we will be happy to have a representative assist you.

### **Who do I contact if I run into issues with NEMo?**

Please email your questions to support@gear6.com with Subject line: "NEMo Support". In the body please include the following details:

- Contact information
- How you obtained NEMo (via email from gear6 or web download from gear6.com)
- NEMo version info
- Relevant environment details such as hardware architecture and OS version, tcpdump version, NFS version and NAS storage being used (NetApp filers, EMC Celerra, etc)