

Scalability Testing with NFS File-IO

Goal: Analyze memory usage as the number of NFS clients is scaled up, and also when NFS endpoint settings are tweaked.

Hopefully by the end of this cookbook you'll have an idea how to assess how many File-IO endpoints and associated settings your system can support. This cookbook connects a LANforge system to a NFS file server. The file server in this example will be 10.17.1.1. It will be sharing 10.17.1.1:/mnt/tmpfs1. We will create 100, 250, then 500 readers and writers and compare memory usage as we also tweak the RW Size setting.

Note: For this example, we are using a CT523 with 8GB of memory. It is physically connected to the NFS file server.

1. First, **500** MAC-VLANs will need to be created.
 - A. In the **Port-Mgr** tab select port **eth1** and click **Create**.

- A. Select **MAC-VLAN**.
 - B. Set **Quantity** to **500**.
 - C. The starting IP address is 10.17.1.2/16 for this test. If your file server is on a different network, change the IP here accordingly. Make sure not to use the file server's IP address.
 - D. Click **Apply** and close the Port Create window once all ports are configured. Make sure all MAC-VLANs get IPs, this may take some time.

2. Create **250** NFS writers and **250** NFS readers. Batch create **50** at a time (49 for first batch of reader/writers). More information on creating File-IO endpoints can be found here: [LANforge File-IO with CIFS and NFS](#)

A. Any non-default settings for the **NFS writers** are listed below.

The screenshot shows the 'Create/Modify File Endpoint' dialog box for an NFS writer. The settings are as follows:

- Name: **nfs-writer001**
- FS-Type: **NFS**
- Port: **2 (eth1#0)**
- Min-RW-Size: **4k (4 KB)**
- Max-RW-Size: **4k (4 KB)**
- Min File Size: **large (1 MB)**
- Max File Size: **large (1 MB)**
- Min Read Rate: **T1 (1.544 Mbps)**
- Max Read Rate: **T1 (1.544 Mbps)**
- Min Write Rate: **100M (100 Mbps)**
- Max Write Rate: **100M (100 Mbps)**
- Quiesce After: **Forever (0)**
- Server: **10.17.1.1:/mnt/tmpfs1**
- Read/Write: **Write**
- Quiesce: **3 (3 sec)**
- Pattern: **increasing**
- Prefix: **AUTO**
- Use O_DIRECT**

- A. Name is **nfs-writer001**.
- B. FS-Type: **NFS**.
- C. Port is **eth1#0**.
- D. Min-RW and Max-RW sizes are **4k**.
- E. Min/Max File sizes are **1MB**.
- F. Min/Max Write rates are **100Mbps**.
- G. Quiesce After is **Forever**.
- H. Server is **10.17.1.1:/mnt/tmpfs1** (use your own file server settings here).
- I. Make sure Read/Write is set to **Write**.
- J. Turn on the **Use O_DIRECT** checkbox.

B. Any non-default settings for the **NFS readers** are listed below.

The screenshot shows the 'Create/Modify File Endpoint' dialog box for an NFS reader. The settings are as follows:

- Name: **nfs-reader001**
- FS-Type: **NFS**
- Port: **252 (eth1#250)**
- Min-RW-Size: **4k (4 KB)**
- Max-RW-Size: **4k (4 KB)**
- Min File Size: **large (1 MB)**
- Max File Size: **large (1 MB)**
- Min Read Rate: **T1 (1.544 Mbps)**
- Max Read Rate: **T1 (1.544 Mbps)**
- Min Write Rate: **100M (100 Mbps)**
- Max Write Rate: **100M (100 Mbps)**
- Quiesce After: **Forever (0)**
- Server: **10.17.1.1:/mnt/tmpfs1**
- Read/Write: **Read**
- Quiesce: **3 (3 sec)**
- Pattern: **increasing**
- Prefix: **nfs-writer001**
- Use O_DIRECT**

- A. Name is **nfs-reader001**.
- B. FS-Type: **NFS**.
- C. Port is **eth1#250**.
- D. Min-RW and Max-RW sizes are **4k**.
- E. Min/Max File sizes are **1MB**.
- F. Min/Max Write rates are **100Mbps**.
- G. Quiesce After is **Forever**.
- H. Server is **10.17.1.1:/mnt/tmpfs1** (use your own file server settings here).
- I. Set Read/Write to **Read**.
- J. Set Prefix to **nfs-writer001**.
- K. Turn on the **Use O_DIRECT** checkbox.

3. Start the first **50 writers** then the first **50 readers**. It's recommended to slowly start File-IO endpoints, for example, starting 25 writers then 25 readers then the same again until you have the desired amount running.

The reason to do this is because the system may become unresponsive if it can't handle the number of File-IO endpoints.

Note: If you notice the system slowing down, connections not starting, or connections stuck in WAITING state, or any weirdness in general, you should stop 25 or 50 writers/readers at a time until the system clears up.

LANforge Manager Version(5.3.4)

Control Reporting Tear-Off Info Plugins

Stop All Restart Manager Refresh HELP

Layer-4 Generic Test Mgr Test Group Resource Mgr Event Log Alerts Port Mgr vAP Stations Messages

Status Layer-3 L3 Endps VoIP/RTP VoIP/RTP Endps Armageddon WanLinks Attenuators File-IO

Rpt Timer: default (5 s) Go Test Manager all

Select All Start Stop Quiesce Clear

Create Modify Batch Modify Delete

Cross Connects for Selected Test Manager

Name	EID	Type	Status	Read-Bps	Rx-Bps-20s	Files-Read	Buf-RD	Bytes-RD	Write-Bps	Tx-Bps-20s	Files-
nfs-writer001	1.1.2.5...	NFS	Run	0	0	0	0	0	21,515,751	21,779,280	
nfs-writer002	1.1.3.5...	NFS	Run	0	0	0	0	0	21,391,168	21,654,267	
nfs-writer003	1.1.4.5...	NFS	Run	0	0	0	0	0	20,480,000	20,731,892	
nfs-writer004	1.1.5.5...	NFS	Run	0	0	0	0	0	22,649,273	22,860,735	
nfs-writer005	1.1.6.5...	NFS	Run	0	0	0	0	0	21,865,464	22,068,494	
nfs-writer006	1.1.7.5...	NFS	Run	0	0	0	0	0	21,523,062	21,721,817	
nfs-writer007	1.1.8.5...	NFS	Run	0	0	0	0	0	21,457,671	21,656,914	
nfs-writer008	1.1.9.5...	NFS	Run	0	0	0	0	0	21,386,693	21,475,817	
nfs-writer009	1.1.10....	NFS	Run	0	0	0	0	0	21,621,625	21,708,458	
nfs-writer010	1.1.11....	NFS	Run	0	0	0	0	0	21,219,629	21,309,128	
nfs-writer011	1.1.12....	NFS	Run	0	0	0	0	0	21,374,660	21,460,501	
nfs-writer012	1.1.13....	NFS	Run	0	0	0	0	0	21,704,260	21,794,708	
nfs-writer013	1.1.14....	NFS	Run	0	0	0	0	0	20,791,500	20,878,144	
nfs-writer014	1.1.15....	NFS	Run	0	0	0	0	0	21,189,858	21,278,162	
nfs-writer015	1.1.16....	NFS	Run	0	0	0	0	0	21,025,465	21,114,144	
nfs-writer016	1.1.17....	NFS	Run	0	0	0	0	0	20,492,175	20,577,572	
nfs-writer017	1.1.18....	NFS	Run	0	0	0	0	0	21,041,802	21,128,428	
nfs-writer018	1.1.19....	NFS	Run	0	0	0	0	0	20,593,996	20,676,703	
nfs-writer019	1.1.20....	NFS	Run	0	0	0	0	0	20,592,273	20,678,086	
nfs-writer020	1.1.21....	NFS	Run	0	0	0	0	0	20,903,773	20,985,617	
nfs-writer021	1.1.22....	NFS	Run	0	0	0	0	0	20,966,465	21,049,610	
nfs-writer022	1.1.23....	NFS	Run	0	0	0	0	0	20,978,302	21,064,668	
nfs-writer023	1.1.24....	NFS	Run	0	0	0	0	0	20,969,219	21,055,547	
nfs-writer024	1.1.25....	NFS	Run	0	0	0	0	0	21,343,792	21,430,585	
nfs-writer025	1.1.26....	NFS	Run	0	0	0	0	0	21,441,394	21,530,746	
nfs-writer026	1.1.27....	NFS	Stopped	0	0	0	0	0	0	0	
nfs-writer027	1.1.28....	NFS	Stopped	0	0	0	0	0	0	0	
nfs-writer028	1.1.29....	NFS	Stopped	0	0	0	0	0	0	0	

Logged in to: brent-523:4002 as: Admin

4. Now that 50 File-IO endpoints are running, the memory will be analyzed via htop.

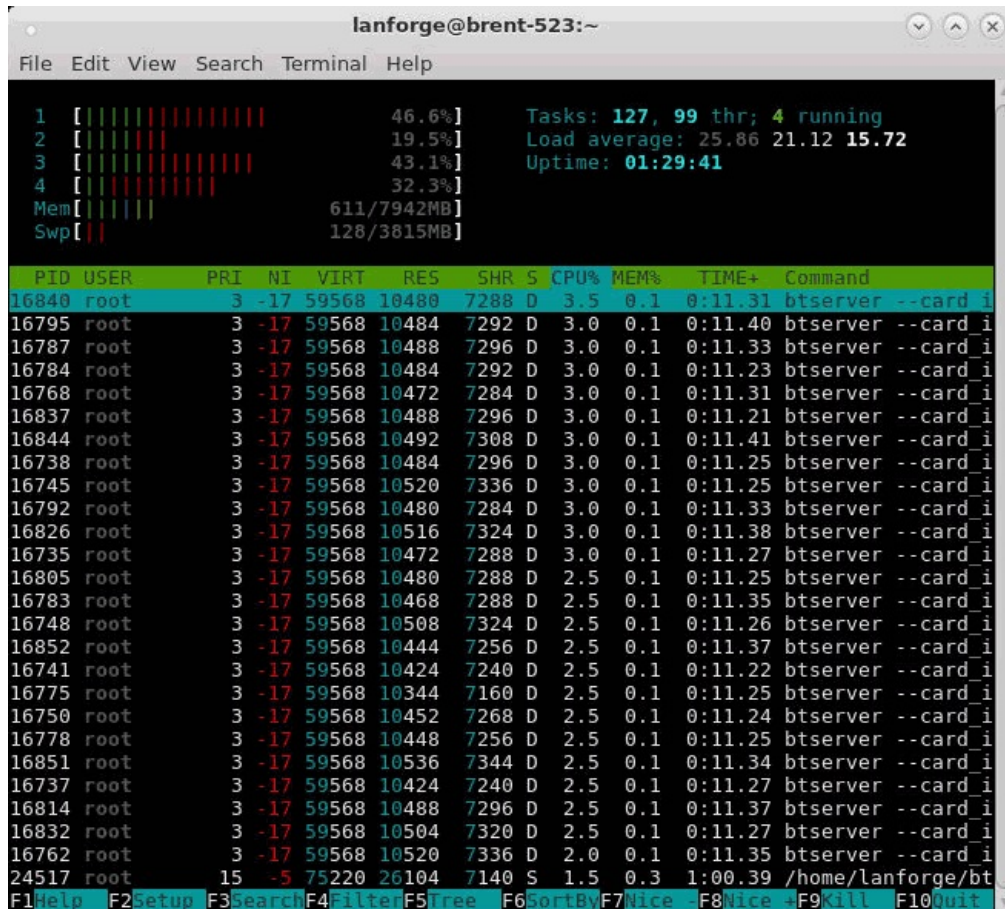
A. Open a terminal on the LANforge system using a preferred method (either directly or remotely through ssh, desktop, or vncviewer).

```
lanforge@brent-523:~
File Edit View Search Terminal Help
[lanforge@brent-523 ~]$
```

B. Show a list of btserver processes (LANforge processes) by running the command:

```
htop -p `pgrep btserver | tr "\n" "," | sed 's/,,$//`
```

If htop isn't installed run: `sudo yum install htop`



The screenshot shows a terminal window titled 'lanforge@brent-523:~'. At the top, system statistics are displayed: CPU usage at 46.6%, memory at 611/7942MB, and swap at 128/3815MB. Below this, a table of running processes is shown. The table has columns for PID, USER, PRI, NI, VIRT, RES, SHR, S, CPU%, MEM%, TIME+, and Command. The 'RES' column shows values around 10,400 KiB for most btserver processes. The command column shows 'btserver --card i'. At the bottom of the terminal, function key shortcuts are listed: F1 Help, F2 Setup, F3 Search, F4 Filter, F5 Tree, F6 SortBy, F7 Nice, F8 Nice, F9 Kill, F10 Quit.

PID	USER	PRI	NI	VIRT	RES	SHR	S	CPU%	MEM%	TIME+	Command
16840	root	3	-17	59568	10480	7288	D	3.5	0.1	0:11.31	btserver --card i
16795	root	3	-17	59568	10484	7292	D	3.0	0.1	0:11.40	btserver --card i
16787	root	3	-17	59568	10488	7296	D	3.0	0.1	0:11.33	btserver --card i
16784	root	3	-17	59568	10484	7292	D	3.0	0.1	0:11.23	btserver --card i
16768	root	3	-17	59568	10472	7284	D	3.0	0.1	0:11.31	btserver --card i
16837	root	3	-17	59568	10488	7296	D	3.0	0.1	0:11.21	btserver --card i
16844	root	3	-17	59568	10492	7308	D	3.0	0.1	0:11.41	btserver --card i
16738	root	3	-17	59568	10484	7296	D	3.0	0.1	0:11.25	btserver --card i
16745	root	3	-17	59568	10520	7336	D	3.0	0.1	0:11.25	btserver --card i
16792	root	3	-17	59568	10480	7284	D	3.0	0.1	0:11.33	btserver --card i
16826	root	3	-17	59568	10516	7324	D	3.0	0.1	0:11.38	btserver --card i
16735	root	3	-17	59568	10472	7288	D	3.0	0.1	0:11.27	btserver --card i
16805	root	3	-17	59568	10480	7288	D	2.5	0.1	0:11.25	btserver --card i
16783	root	3	-17	59568	10468	7288	D	2.5	0.1	0:11.35	btserver --card i
16748	root	3	-17	59568	10508	7324	D	2.5	0.1	0:11.26	btserver --card i
16852	root	3	-17	59568	10444	7256	D	2.5	0.1	0:11.37	btserver --card i
16741	root	3	-17	59568	10424	7240	D	2.5	0.1	0:11.22	btserver --card i
16775	root	3	-17	59568	10344	7160	D	2.5	0.1	0:11.25	btserver --card i
16750	root	3	-17	59568	10452	7268	D	2.5	0.1	0:11.24	btserver --card i
16778	root	3	-17	59568	10448	7256	D	2.5	0.1	0:11.25	btserver --card i
16851	root	3	-17	59568	10536	7344	D	2.5	0.1	0:11.34	btserver --card i
16737	root	3	-17	59568	10424	7240	D	2.5	0.1	0:11.27	btserver --card i
16814	root	3	-17	59568	10488	7296	D	2.5	0.1	0:11.37	btserver --card i
16832	root	3	-17	59568	10504	7320	D	2.5	0.1	0:11.27	btserver --card i
16762	root	3	-17	59568	10520	7336	D	2.0	0.1	0:11.35	btserver --card i
24517	root	15	-5	75220	26104	7140	S	1.5	0.3	1:00.39	/home/lanforge/bt

A. The majority of btserver processes with the same or similar values here represent the running File-IO endpoints. The **RES** (resident value) column represents the memory used by these File-IO endpoints. The current test uses an average of about 10,400 KiB or **10.6 MB** per btserver process. These values should be monitored as the File-IO settings are adjusted.

NOTE: You will need to restart htop if the endpoints are stopped/restarted. Just press **q** to quit out and run the above htop command again.

C. Total up the RES memory used by btserver. This is useful to compare memory usage between a different number of connections. For example you could compare how much memory 100 endpoints use compared to 500 endpoints.

The image shows two terminal windows from a user named 'lanforge' on a machine named 'brent-523'. The top terminal window shows the command `ps aux | grep btserver | awk '{s+=$6} END {print s}'` being executed, which outputs the number `1130916`. The bottom terminal window shows a `while` loop command: `while sleep 1; do ps aux | grep btserver | awk '{s+=$6} END {print s}'; done`. This loop outputs a list of memory usage values for `btserver` processes every second, including `1130916`, `1130984`, `1130936`, `1130920`, `1130912`, `1130972`, `1130984`, `1130976`, `1130880`, `1130916`, `1130940`, `1130880`, `1130908`, `1130892`, and `1130880`.

- A. To show the total RES memory used, run the command:
`ps aux | grep btserver | awk '{s+=$6} END {print s}'`
To repeat it every second, use the below command. Press **Ctrl+C** to stop.
`while sleep 1; do ps aux | grep btserver | awk '{s+=$6} END {print s}'; done`
- B. The total memory for 50 NFS File-IO endpoints with a **Min/Max RW rate of 4k** is about 1,130,000 KiB which converts to around **1,157 MB**.

5. While the File-IO endpoints are running, change the **Min-RW** and **Max-RW** settings.

A. Select all File-IO endpoints and click **Batch Modify**.

Min Read Rate:	NA	Max Read Rate:	NA
Min Write Rate:	NA	Max Write Rate:	NA
Min-RW-Size:	1MB (1 MB)	Max-RW-Size:	1MB (1 MB)
Min File Size:	NA	Max File Size:	NA
Pattern:	NA	Prefix:	NA
Server:	NA	Options:	NA
File #:	NA	Flags:	NA
Do-CRC:	NA	Retry-Timer:	NA
Read/Write:	NA	FS-Type:	NA
Quiesce After	NA (-1)		

Apply OK Cancel

A. Set the **Min/Max RW sizes** to **1MB**.

B. Click **OK**

6. Analyze how the Min/Max RW size setting change has affected the memory (see steps 4B and 4C).

```

lanforge@brent-523:~
File Edit View Search Terminal Help

1 [||||] 13.3% Tasks: 204, 102 thr; 1 running
2 [||||] 6.7% Load average: 46.19 54.05 54.10
3 [|||] 4.3% Uptime: 01:01:02
4 [||||] 11.0%
Mem[|||||] 906/7942MB
Swp[ ] 0/3815MB

PID USER PRI NI VIRT RES SHR S CPU% MEM% TIME+ Command
2376 root 3 -17 97840 48036 7508 S 1.0 0.6 1:24.07 /home/lanforge/bt
2394 root 15 -5 75256 25312 7388 S 0.5 0.3 0:50.77 /home/lanforge/bt
22182 root 3 -17 60600 11496 7284 S 0.5 0.1 0:29.35 btserver --card_i
22316 root 3 -17 60600 11544 7332 D 0.5 0.1 0:29.21 btserver --card_i
22210 root 3 -17 60600 11612 7396 S 0.5 0.1 0:29.12 btserver --card_i
22315 root 3 -17 60600 11548 7332 S 0.5 0.1 0:29.76 btserver --card_i
22263 root 3 -17 60600 11612 7396 D 0.5 0.1 0:29.17 btserver --card_i
22192 root 3 -17 60600 11636 7420 D 0.5 0.1 0:29.19 btserver --card_i
22153 root 3 -17 60600 11548 7336 D 0.5 0.1 0:29.17 btserver --card_i
22261 root 3 -17 60600 11620 7408 D 0.5 0.1 0:29.35 btserver --card_i
22313 root 3 -17 60600 11612 7396 D 0.5 0.1 0:29.08 btserver --card_i
22197 root 3 -17 60600 11520 7308 D 0.5 0.1 0:29.21 btserver --card_i
22169 root 3 -17 60600 11556 7340 D 0.5 0.1 0:29.05 btserver --card_i
22292 root 3 -17 60600 11628 7412 D 0.5 0.1 0:29.69 btserver --card_i
23304 root 3 -17 60600 11628 7420 S 0.5 0.1 0:07.70 btserver --card_i
22013 root 3 -17 60600 11624 7412 D 0.0 0.1 0:31.33 btserver --card_i
22140 root 3 -17 60600 11592 7376 D 0.0 0.1 0:29.00 btserver --card_i
22194 root 3 -17 60600 11616 7400 D 0.0 0.1 0:29.21 btserver --card_i
22230 root 3 -17 60600 11584 7368 D 0.0 0.1 0:29.15 btserver --card_i
22296 root 3 -17 60600 11632 7420 D 0.0 0.1 0:29.17 btserver --card_i
23130 root 3 -17 60600 11564 7352 S 0.0 0.1 0:07.62 btserver --card_i
23207 root 3 -17 60600 11604 7396 S 0.0 0.1 0:07.62 btserver --card_i

lanforge@brent-523:~
File Edit View Search Terminal Help

[lanforge@brent-523 ~]$ ps aux | grep btserver | awk '{s+=$6} END {print s}'
1233688
[lanforge@brent-523 ~]$


```

A. Changing the Min/Max RW sizes to 1MB increased the btserver memory use to about 11,600 KiB or **11.9 MB** from 10,400 KiB or **10.6 MB**. About a 1,200 KiB or **1.3 MB** difference.


The total memory increased to 1,233,000 KiB, **1,262 MB** from 1,130,000 KiB, **1,157 MB**. About a 103,000 KiB or **105 MB** difference.

7. Set the endpoints **Min/Max RW** size back to **4k**.

A. Stop all running file-IO endpoints by selecting them and clicking **Stop**.



B. Select all file-IO endpoints and click **Batch Modify**.



- A. Set the **Min/Max RW sizes** to **4k**.
- B. Click **OK**.

8. Repeat steps **3 through 7** for **125 writers/readers** and then again for **250 writers/readers**.

9. The results from the above tests are shown in the below tables. They show how RW size affects 100, 250, and 500 NFS File-IO endpoints.

Individual Process Memory Usage (MB)				
		Number of File-IO Endpoints		
		100	250	500
RW Size	4 KB	10.6	10.8	10.8
	1 MB	11.9	11.9	11.8

Total Memory Usage (MB)				
		Number of File-IO Endpoints		
		100	250	500
RW Size	4 KB	1,157	2,804	4,134
	1 MB	1,262	3,031	6,013

Candela Technologies, Inc., 2417 Main Street, Suite 201, Ferndale, WA 98248, USA
www.candelatech.com | sales@candelatech.com | +1.360.380.1618