

# Log Analysis and Iltop LUG 2011

John Hammond  
TACC

The University of Texas at Austin

[jhammond@tacc.utexas.edu](mailto:jhammond@tacc.utexas.edu)

# Collaborators

- Tommy Minyard, TACC
- James C. Browne, U.T. C.S.
- Edward Chuah, Howie Kuo, Gary Khoon,  
William Tjhi, IHPC Singapore

# Part 1: Log Analysis

*Log analysis is for the birds.*

-Kent B.

*System logs are a kind of cognitive black hole. It's easy to pour vast amounts of time and energy in to staring at them, but then have nothing ever seem to come out.*

-A.

# Why bother with logs?

- Good trace (`printk( )`) better than debugger, SystemTap, bad trace, ...
- They're there: easy to collect with no added infrastructure.

# Rational message encoding

In the kernel, modify `printk()` by inserting call to a new function, `rat_printk()`.

Every call to `printk()`,

```
printk(KERN_ERR "device failed: Error %d: %s\n", errno, errmsg);
```

would produce the normal message:

```
Mar 10 15:09:30 host1 kernel: device failed: Error 157: Not a toaster
```

along with an unformatted version of the same (line breaks inserted for clarity):

```
time:1268255370\0
host:host1\0
progname:kernel\0
0:<3>Print failed: Error %d: %s\n\0
1:157\0
2:Not a toaster\0\0
```

# Transport parallel to `printk()`

- Store in ring buffer in kernel.
- Read from `/proc/rat-kmsg`  
(message oriented)
- Run `rat-syslogd` to add timestamp, hostname, store, forward, ...

# Comparison 1

```
Jul  8 12:27:52 i115-312 kernel: test_sine.x[27673]: segfault \  
at 00002aaa9e094c88 rip 00002b92854db19e rsp 00007fff25bc0b40 \  
error 4
```

```
time:1278610053\  
host:i115-312\  
jobid:1461336\  
prog:kernel\  
0:%s%s[%d]: segfault at %016lx rip %016lx rsp %016lx error %lx\  
1:<6>  
2:test_sine.x\  
3:27673\  
4:46912284216456\  
5:47908301681054\  
6:140733826468672\  
7:4
```

# Comparison 2

```
LustreError: 11-0: an error occurred while communicating with \
129.114.97.41@o2ib. The ost_write operation failed with -107
```

```
0:%sLustre: %x-%x: an error occurred while communicating with \
%s. The %s operation failed with %d\n\0
1:<3>\0
2:17\0
3:0\0
4:129.114.97.41@o2ib\0
5:ost_write\0
6:-107\0
```



# Comparison 3

```
LustreError: 1943:0:(osc_request.c:1143:can_merge_pages()) is \
it ok to have flags 0xc20 and 0x420 in the same brw?
```

```
0:%sLustreError: %d:%d:(%s:%d:%s()) is it ok to have flags 0x%x\
and 0x%x in the same brw?\n
1:<3>
2:1943
3:0
4:osc_request.c
5:1143
6:can_merge_pages
7:3104
8:1056
```

```
Apr 12 01:07:04 i115-409 kernel: Lustre: 3570:(client.c:1476:\
ptlrpc_expire_one_request()) @@@ Request x136339563010 0532 sent \
from scratch-OST009a-osc-ffff810877104c00 to NID 129.114.97.67@...
```

```
$ rat-log-find -msg 'Apr 12 01:07:04 i115-409 kernel: Lustre: 3570'
time:1302588424
host:i115-409
jobid:1902176
prog:kernel
0:%sLustre: %d:(%s:%d:%s()) @@@ Request x%Lu sent from %s to NID ...
1:<4>
2:3570
3:client.c
4:1476
5:ptlrpc_expire_one_request
6:1363395630100532
7:scratch-OST009a-osc-ffff810877104c00
8:129.114.97.67@o2ib
...
40:0
```

```
Apr 12 01:07:04 i115-409 kernel: Lustre: 3570:(client.c:1476:\
ptlrpc_expire_one_request()) @@@ Request x136339563010 0532 \
sent from scratch-OST009a-osc-ffff810877104c00 to NID 129.114.97.67\
@o2ib 7s ago has timed out (7s prior to deadline).
Apr 12 01:07:04 i115-409 kernel: req@ffff8103ee3a7000 \
x1363395630100532/t0 o101->scratch-OST009a_UUID@\
129.114.97.67@o2ib:28/4 lens 576/544 e 0 to 1 dl 1302588424 \
ref 1 fl Rpc:/0/0 rc 0/0
```

```
$ rat-log-find -msg 'Apr 12 01:07:04 i115-409' -print 0
%sLustre: %d:(%s:%d:%s()) @@@ Request x%Lu sent from %s to NID %s\
%lus ago has %s (%lds prior to deadline).\n req@%p x%Lu/t%Ld \
o%d->%s@%s:%d/%d lens %d/%d e %d to %d dl %ld ref %d fl %s:%s%s%s%s\
%s%s%s%s%s/%x/%x rc %d/%d\n
```

```
$ rat-log-find -msg 'Apr 12 01:07:04 i115-409' -print 8
129.114.97.67@o2ib
```

```
$ rat-log-find -msg 'Apr 12 01:07:04 i115-409' -print 15
101
```

# Part 2: TACC Iltop

- Lustre load monitor with batch (SGE,...) scheduler integration
- Point-and-shoot operation.
- Helps decide where to send the goons.


# Real\* Screenshot

```
$ sudo tacc_lltop work
```

JOBID	WR_MB	RD_MB	REQS	OWNER	WORKDIR
1901914	248	0	641	nellie	/work/01251/nellie/XYZ20/Lu
1901919	180	0	727	nellie	/work/01251/nellie/XYZ20/Lu
1901959	143	0	622	nellie	/work/01251/nellie/XYZ20/Lu
ilogin3	90	33	2463		
1901923	87	0	524	nellie	/work/01251/nellie/XYZ20/Lu
1902767	86	0	276	alice	/work/01117/alice/CC_1p
1901957	80	0	1683	nellie	/work/01251/nellie/XYZ20/Lu
1902703	51	0	104	edward	/work/01682/edward/H20_flu
1901838	44	0	88	samuel	/work/00934/samuel/4_1
1902684	10	0	23	mary	/work/00491/mary/sims/dma
1901166	9	0	23	james	/work/01516/james/big
1900672	8	0	19	james	/work/01516/james/small
1902252	6	0	66	cora	/work/01333/cora/123/prod
1901738	6	0	41	annie	/work/00287/annie/fgh20
1901917	5	0	2340	nellie	/work/01251/nellie/XYZ20/Lu
1901305	4	0	29	minnie	/work/00914/minnie/MUT/prod
1902646	4	19	52	frank	/work/01640/frank/CORAL/R1
1900742	3	0	35	bertha	/work/01246/bertha/amber/cl

```
...
```

# lltop architecture



# Design goals.

- Single executable to install on servers.
- Invocation and transport by ssh.
- Hookable
  - LNET-address to hostname conversion.
  - Hookable hostname to jobid conversion.
  - FS name to server list.



<https://github.com/jhammond/lltop>



# Part 3: Lustre /proc

- Mostly very good.
- Needs some attention.
  - Some are just wrong.
  - Some are only partially wrong.
  - What is a reg?
- Force multiplier.
  - Enable toolmakers to add value without mucking about in Lustre source.

# Wish list 1.

- Make it easier for zero-conf tools to associate servers, /proc directories with mount points.
  - What if they mount two scratches?
  - Appending “\_UUID” does not make it so.
- Make “Bytes on the wire” stats available in one file.

# Wish list 2.

- Keep /proc.
- Text only.
- No special tools please.
- Strict “No ASCII art policy.”