

## Sisyphus Log Analysis Toolkit

Date: start:  << < = > >>  
 First: 2006/04/01 00:00:00  
 stop:  Last: 2006/07/13 06:00:00  
 << < = > >>

Host:

Job ID:

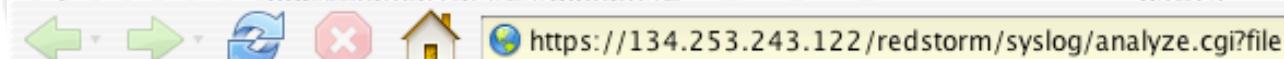
Word:

Display top N:  sorted by

List: [Documents](#) [?](#)

Displaying 100 of -1 docs. 0 selected.

<input type="checkbox"/>	<input checked="" type="checkbox"/>	<a href="#">YYYY/MM/DD/HH</a>	<a href="#">HOST</a>	<a href="#">bytes</a>
<input type="checkbox"/>	<input type="checkbox"/>	<a href="#">docs/2006/07/03/23/nid00832</a>		65275
<input type="checkbox"/>	<input type="checkbox"/>	<a href="#">docs/2006/07/04/nid00835</a>		83069
<input type="checkbox"/>	<input type="checkbox"/>	<a href="#">docs/2006/07/03/23/nid00543</a>		122268



<input type="checkbox"/>	3	24	150	100	kern err kernel: LustreError: * pi
<input type="checkbox"/>	4	1	0	0	kern err kernel: LustreError: * *
server code, just going back to sleep ns: mds-mds_15_UUID lock: 00000100397a8cc0/0x2a					
PLN flags: 4000 remote: 0x0 expref: -99 pid: 1031					
Jul 3 23:53:57 nid00832 kern err kernel: LustreError: 1513:0:(client.e:945:ptlrc_expire_one					
req@0000010039001600 x2377216/t0 08->ost00543_0_UUID@NID_543_UUID:6 lens 2					
Jul 3 23:53:57 nid00832 kern alert kernel: LustreError: dumping log to /tmp/lustre-log-nid00					

## Event logs are:

- **Ubiquitous** (nearly all types of com)
- **Informational** (nearly always the first)
- **Voluminous** (infeasible to read all of them)
- **Noisy** (many messages are completely useless)

• **Non-stationary** (characteristics)

## Interrogating (grep) or monitoring (log)

- Is common
- Can be easy (or NOT)
- Can be effective (or NOT)

Many hard questions cannot be easily answered

Under what conditions does string “x” appear in log  
^ e.g. nodes, times, locations

*Premise:*

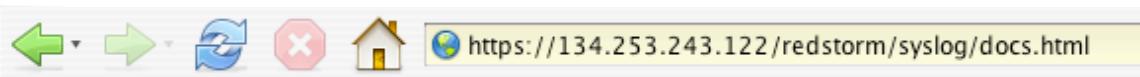
**Similar computers  
correctly execute  
the same program  
should produce  
similar log messages**

*Process:*

**Group** log messages into “document”  
according to time and host factors

**Encode** as a “term-doc” matrix:

**$X=GL$**  where  $G$  is a per-term  
information weight regarding  
some factor of interest  
(Host, Time, App, User, etc)



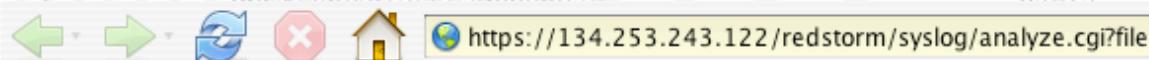
## Sisyphus Log Analysis Toolkit

(1) Date: start:  << < = > >>  
stop:  First: 2006/04/01 00:00:00  
Host:  Last: 2006/07/13 06:00:00  
Job ID:   
Word:   
Display top N:  sorted by

List:

Displaying 100 of -1 docs. 0 selected

<input type="checkbox"/>	<input checked="" type="checkbox"/>	YYYY/MM/DD HH HOST	bytes
<input type="checkbox"/>	<input checked="" type="checkbox"/>	<a href="#">docs/2006/07/03/23/nid00832</a>	65275
<input type="checkbox"/>	<input checked="" type="checkbox"/>	<a href="#">docs/2006/07/07/04/nid00835</a>	83069
<input type="checkbox"/>	<input checked="" type="checkbox"/>	<a href="#">docs/2006/07/03/23/nid00543</a>	122268



(2)

<input type="checkbox"/>	3	24	150	100	kern err kernel: LustreError: * p
<input type="checkbox"/>	4	1	0	0	kern err kernel: LustreError: * *

(4)

server code, just going back to sleep ns: mds-mds\_15\_UUID lock: 00000100397a8cc0/0x2a  
PLN flags: 4000 remote: 0x0 expref: -99 pid: 1031  
Jul 3 23:53:57 nid00832 kern err kernel: LustreError: 1513:0:(client.c:945:ptlrpc\_expire\_one\_request()) req@0000010039001600 x2377216/t0 o8->ost00543\_0\_UUID@NID\_543 UUID:6 lens 2  
Jul 3 23:53:57 nid00832 kern alert kernel: LustreError: dumping log to /tmp/lustre-log-nid00832

(5)

<input type="checkbox"/>	<a href="#">ID</a>	<a href="#">pos</a>	<a href="#">word</a>	<a href="#">host info</a>
<input checked="" type="checkbox"/>	195	4	[:obdclass:obd_dev+4272]	4.57
<input type="checkbox"/>	127	5	1513:0:(client.c:945:ptlrpc_expire_one_request())	4.29
<input type="checkbox"/>	318	25	mds-mds_15_UUID	3.91

(6)

## How effective is it?

Of the 33 **GL** Combinations Tested,  
**host information weighting** provides  
exceptionally good fault detection.

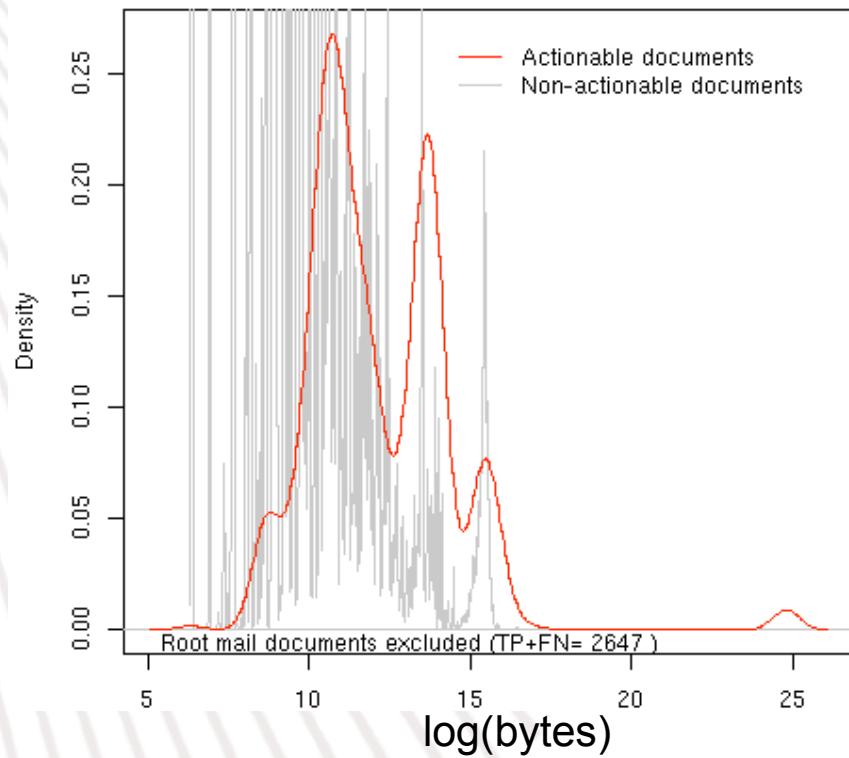
On a 512 Node Linux Cluster at SNL,  
60% of faults are detected with an  
alarm precision rate of at least 80%.  
This corresponds to an excellent false  
positive rate of 0.07%.

Results for Red Storm, Blue Gene,  
Thunderbird, and Liberty are pending...

# Effectiveness

## Bytes

VS





## Sisyphus Log Analysis Toolkit

Date: start:  << < = > >>  
First: 2006/04/01 00:00:00  
stop:  Last: 2006/07/13 06:00:00  
<< < = > >>

Host: (not

Job ID:

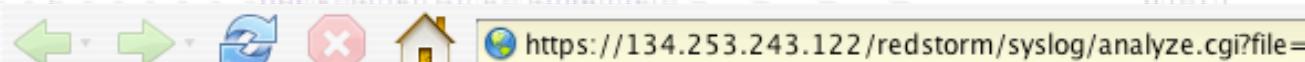
Word:

Display top N:  sorted by

List: [Documents](#) [?](#)

Displaying 100 of -1 docs. 0 selected. (

<input type="checkbox"/>	<input checked="" type="checkbox"/>	YYYY/MM/DD/HH	HOST	bytes
		<a href="#">docs/2006/07/03/23/nid00832</a>		65275
		<a href="#">docs/2006/07/07/04/nid00835</a>		83069
		<a href="#">docs/2006/07/03/23/nid00543</a>		122268



<input type="checkbox"/>	3	24	150	100	kern err kernel: LustreError: * pre
<input type="checkbox"/>	4	1	0	0	kern err kernel: LustreError: * * *

server code, just going back to sleep ns: mds-mds\_15\_UUID lock: 00000100397a8cc0/0x2a7  
PLN flags: 4000 remote: 0x0 expref: -99 pid: 1031

Jul 3 23:53:57 nid00832 kern err kernel: LustreError: 1513:0:(client.c:945:ptlRpc\_expire\_one\_req@0000010039001600 x2377216/t0 o8->ost00543\_0\_UUID@NID\_543\_UUID:6 lens 24

Jul 3 23:53:57 nid00832 kern alert kernel: LustreError: dumping log to /tmp/lustre-log-nid008