


Deploying LoGS to Analyze Console Logs on an IBM JS20



James E. Prewett

<http://www.hpc.unm.edu/~download/LoGS>

OpenPGP key pub 1024D/31816D93 2003-01-06 James Prewett <download@hpc.unm.edu>

Key fingerprint = F618 949F 5FB8 918E 8378 8C1F BFFC DDC6 3181 6D93



Ristra

- 96 JS20 Blades
- Two 1.6 GHz PowerPC 970 processors / blade
- 4 GB RAM / blade
- Myrinet Interconnect
- X335 Management node
- xCAT cluster management software



The Problem

- Management node writing all console output to disk
- Problems with SOL system caused 150 messages/second/machine to be logged
- Heavy Disk I/O on management node

(LOGS)



Serial Over LAN (SOL) Failures

- SOL is unavailable on a single system
- SOL is unavailable on all of the systems in a particular chassis
- SOL has been disabled on a single system

(LOGS)



Responses to SOL Failures

- SOL is unavailable on a single system
 - **Reboot the System***
- SOL is unavailable on all of the systems in a particular chassis
 - **Reset the SOL subsystem on the Chassis**
- SOL has been disabled on a single system
 - **Enable SOL for that system**

(LOGS)



My Solution:

(LOGS)

(LOGS)



Quick intro to LoGS

- Rule Based Log Analysis Tool
- Dynamic Ruleset
- Efficient Ruleset Design
- Written entirely in Common Lisp
 - Lisp used for both configuration and extension
- (mostly) ANSI compliant code, works with many ANSI Common Lisps
- Hacker-friendly



LoGS Components

- Messages
- Rules
- Rulesets
- Contexts
- Actions



LoGS Performance

- Generally > 1% of CPU
- > 128 MB Memory under heavy load
- Generally > 10 MB Memory (including Lisp footprint)
- > 70,000 messages/second processed

(LoGS)



Strategy

- Replace log files with FIFOs and use LoGS to only write “interesting” messages to disk to reduce I/O load on “admin node”
- Filter out known uninteresting messages
- Respond to SOL failures
- Respond to certain other known failures

(LoGS)



Filtering

- Remove unwanted messages from the log stream
- good for reducing overall volume of log data

(LOGS)



Filtering Rule

- :: throw away all messages of the form
- :: “username: USERID”

(filter
(lambda (message)
(equal
“username: USERID”
(message message))))

(LOGS)



Responding to SOL Failures

- SOL is unavailable on a single system
 - Alert an Administrator
- SOL is unavailable on all of the systems in a particular chassis
 - Reset the SOL subsystem on the Chassis
- SOL has been disabled on a single system
 - Enable SOL for that system

(LOGS)



SOL unavailable on a single system

```
;; match a single "SOL is not ready" message
;; create a rule to ignore further messages from
;; this host for 5 minutes
(make-instance 'rule
  :match (lambda (message)
    (multiple-value-bind (match-start match-end)
      (cl-ppcre::scan "SOL is not ready" (message message))
      (when match-start
        (values t '((from-file ,(from-file message)))))))
  :continuep t
  :actions
  (list
    ;; note this in the "important" logfile
    (file-write
      (format () "/var/log/LoGS/important/logfile-~A"
        (car (last (CL-PPCRE:SPLIT "/" (from-file message))))))
    ;; ignore further messages from this host for 5 minutes
    (lambda (message)
      (rule-before (make-instance 'rule
        :timeout (+ *now* (* INTERNAL-TIME-UNITS-PER-SECOND 300))
        :environment env ;; grab environment from return of match function
        :match (lambda (message)
          (and (equal from-file (from-file message))
            (cl-ppcre::scan "SOL is not ready" (message message)))))))
```

(LOGS)



SOL unavailable on a single system (cont.)

```
6 04:02:01 aloe syslogd 1.3-3: restart (remote reception).
6 04:02:01 I254 syslogd 1.4.1: restart.
6 04:02:01 I109 anacron[32216]: Updated timestamp for job `cron.daily' to 2004-06-01
6 04:02:01 aloe syslogd 1.3-3: restart (remote reception).
6 04:02:01 I254 anacron[32216]: Updated timestamp for job `cron.daily' to 2004-06-01
6 04:02:01 I109 syslogd 1.4.1: restart.
6 04:02:01 I254 syslogd 1.4.1: restart.
6 04:02:01 I107 CROND[32213]: (root) CMD (run-parts /etc/cron.daily)
;; match a single "SOL is not ready" message
6 04:02:01 I107 CROND[32213]: (root) CMD (run-parts /etc/cron.daily)
;; send email to admins alerting them to SOL error
(make-instance 'rule
:match (lambda (message)
(multiple-value-bind (match-start match-end)
(cl-ppcre::scan "SOL is not ready" (message message))
(when match-start
(values t '((from-file ,(from-file message)))))))
:continup t
:actions
(list
(lambda (message)
(email message "administrators@somedomain.edu"
(format () "SOL error on host:~A" from-file))))))
6 04:02:03 I208 kernel: nfs: server pinon not responding, still trying
6 04:02:05 I167 kernel: nfs: server pinon not responding, still trying
6 04:02:08 aloe PAM_pwdb[5857]: (su) session opened for user news by (uid=0)
6 04:02:14 I169 kernel: nfs: server pinon not responding, s
6 04:02:16 I167 kernel: nfs: server pinon OK
```

(LOGS)



SOL unavailable on entire chassis

```
6 04:02:01 aloe syslogd 1.3-3: restart (remote reception).
6 04:02:01 I254 syslogd 1.4.1: restart.
6 04:02:01 aloe syslogd 1.3-3: restart (remote reception).
6 04:02:01 I109 anacron[32216]: Updated timestamp for job `cron.daily' to 2004-06-01
6 04:02:01 aloe syslogd 1.3-3: restart (remote reception).
6 04:02:01 I254 anacron[32216]: Updated timestamp for job `cron.daily' to 2004-06-01
6 04:02:01 I109 syslogd 1.4.1: restart.
6 04:02:01 I254 anacron[32216]: Updated timestamp for job `cron.daily' to 2004-06-01
6 04:02:01 I109 syslogd 1.4.1: restart.
6 04:02:01 I204 CROND[22193]: (root) CMD (run-parts /etc/cron.daily)
6 04:02:01 I204 syslogd 1.4.1: restart.
6 04:02:03 I208 kernel: nfs: server pinon not responding, still trying
6 04:02:05 I167 kernel: nfs: server pinon not responding, still trying
6 04:02:08 aloe PAM_pwdb[5857]: (su) session opened for user news by (uid=0)
6 04:02:14 I169 kernel: nfs: server pinon not responding, s
6 04:02:16 I167 kernel: nfs: server pinon OK
```

- Maintain table of hosts found to have SOL problems
- When all blades in chassis have the problem:
 - call a shell command to reset SOL on the chassis (mpareset in xcat)
 - Log action to special log file

(LOGS)



SOL disabled on a single system

- Notice SOL unavailable
- Call shell command to re-enable SOL
- Custom expect script similar to scripts in xCAT

(LOGS)



Other things we can find

- Disk errors often don't show up in /var/log/ messages, but are generally logged to the console!
- Software version mis-matches (eg. kernel differences can be found on boot)

(LOGS)



Why LoGS?

- I wrote it (and know it well)
- Efficient ruleset design allows for quicker processing than flat rulesets
- Maintaining tables and other tricks are difficult at best with other tools
- Regular expression don't fulfill the needs of all of these examples!

(LoGS)



LoGS Info

- Website:
 - <http://www.hpc.unm.edu/~download/LoGS/>
- CVS Repository:
 - <http://savannah.nongnu.org/projects/LoGS/>
- Mailing list
 - <http://lists.nongnu.org/mailman/listinfo/logs-devel>

(LoGS)

