Syslog

Robert Petkus
NSLS-II Controls Group
2010 Fall EPICS Collaboration Meeting
October 12, 2010

Outline

- Syslog
- Rsyslog
- Syslog-ng
- Splunk
- LogZilla
- Test Bed

Syslog

Syslog

- is the standard logging solution on UNIX/LINUX systems and network routers/switches
- has evolved over time with several implementations => syslog, rsyslog, syslog-ng
- employs a layered architecture – separation of message content from transport
- reads and logs messages to log files, a console, and/or other systems
- supports output to named pipes (FIFOs) and remote logging (traditionally UDP/514)
- generates messages composed of (5) parts: Time Stamp, Program name, Facility, Priority, Log message

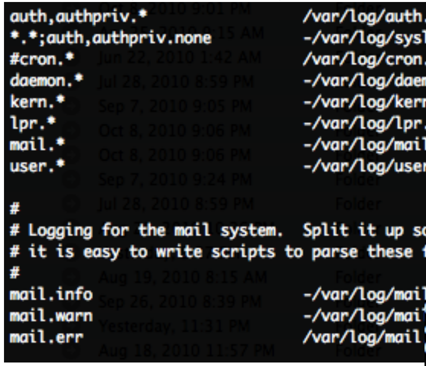

```
auth,authpriv.*                 /var/log/auth.log
*.*;auth,authpriv.none          -/var/log/syslog
#cron.*                         /var/log/cron.log
daemon.*                        -/var/log/daemon.log
kern.*                          -/var/log/kern.log
lpr.*                           -/var/log/lpr.log
mail.*                          -/var/log/mail.log
user.*                          -/var/log/user.log

#
# Logging for the mail system.  Split it up so that
# it is easy to write scripts to parse these files.
#
mail.info                       -/var/log/mail.info
mail.warn                       -/var/log/mai
mail.err                        /var/log/mail
```

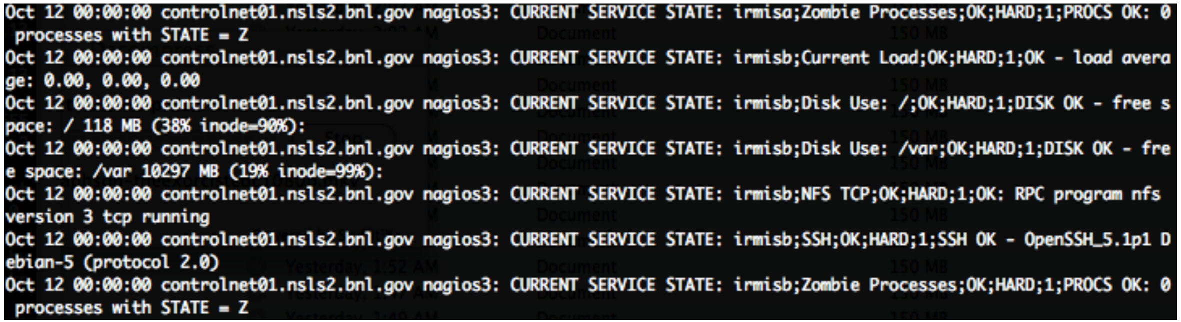

```
Oct 12 00:00:00 controlnet01.nsls2.bnl.gov nagios3: CURRENT SERVICE STATE: irmisa;Zombie Processes;OK;HARD;1;PROCS OK: 0
 processes with STATE = Z
Oct 12 00:00:00 controlnet01.nsls2.bnl.gov nagios3: CURRENT SERVICE STATE: irmisb;Current Load;OK;HARD;1;OK - load avera
ge: 0.00, 0.00, 0.00
Oct 12 00:00:00 controlnet01.nsls2.bnl.gov nagios3: CURRENT SERVICE STATE: irmisb;Disk Use: /;OK;HARD;1;DISK OK - free s
pace: / 118 MB (38% inode=90%):
Oct 12 00:00:00 controlnet01.nsls2.bnl.gov nagios3: CURRENT SERVICE STATE: irmisb;Disk Use: /var;OK;HARD;1;DISK OK - fre
e space: /var 10297 MB (19% inode=99%):
Oct 12 00:00:00 controlnet01.nsls2.bnl.gov nagios3: CURRENT SERVICE STATE: irmisb;NFS TCP;OK;HARD;1;OK: RPC program nfs
version 3 tcp running
Oct 12 00:00:00 controlnet01.nsls2.bnl.gov nagios3: CURRENT SERVICE STATE: irmisb;SSH;OK;HARD;1;SSH OK - OpenSSH_5.1p1 D
ebian-5 (protocol 2.0)
Oct 12 00:00:00 controlnet01.nsls2.bnl.gov nagios3: CURRENT SERVICE STATE: irmisb;Zombie Processes;OK;HARD;1;PROCS OK: 0
 processes with STATE = Z
```

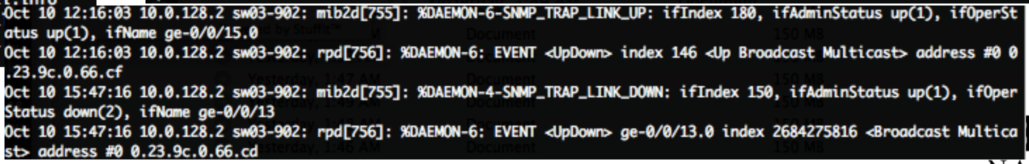

```
Oct 10 12:16:03 10.0.128.2 sw03-902: mib2d[755]: %DAEMON-6-SNMP_TRAP_LINK_UP: ifIndex 180, ifAdminStatus up(1), ifOperSt
atus up(1), ifName ge-0/0/15.0
Oct 10 12:16:03 10.0.128.2 sw03-902: rpd[756]: %DAEMON-6: EVENT <UpDown> index 146 <Up Broadcast Multicast> address #0 0
.23.9c.0.66.cf
Oct 10 15:47:16 10.0.128.2 sw03-902: mib2d[755]: %DAEMON-4-SNMP_TRAP_LINK_DOWN: ifIndex 150, ifAdminStatus up(1), ifOper
Status down(2), ifName ge-0/0/13
Oct 10 15:47:16 10.0.128.2 sw03-902: rpd[756]: %DAEMON-6: EVENT <UpDown> ge-0/0/13.0 index 2684275816 <Broadcast Multica
st> address #0 0.23.9c.0.66.cd
```

Syslog Configuration Rules

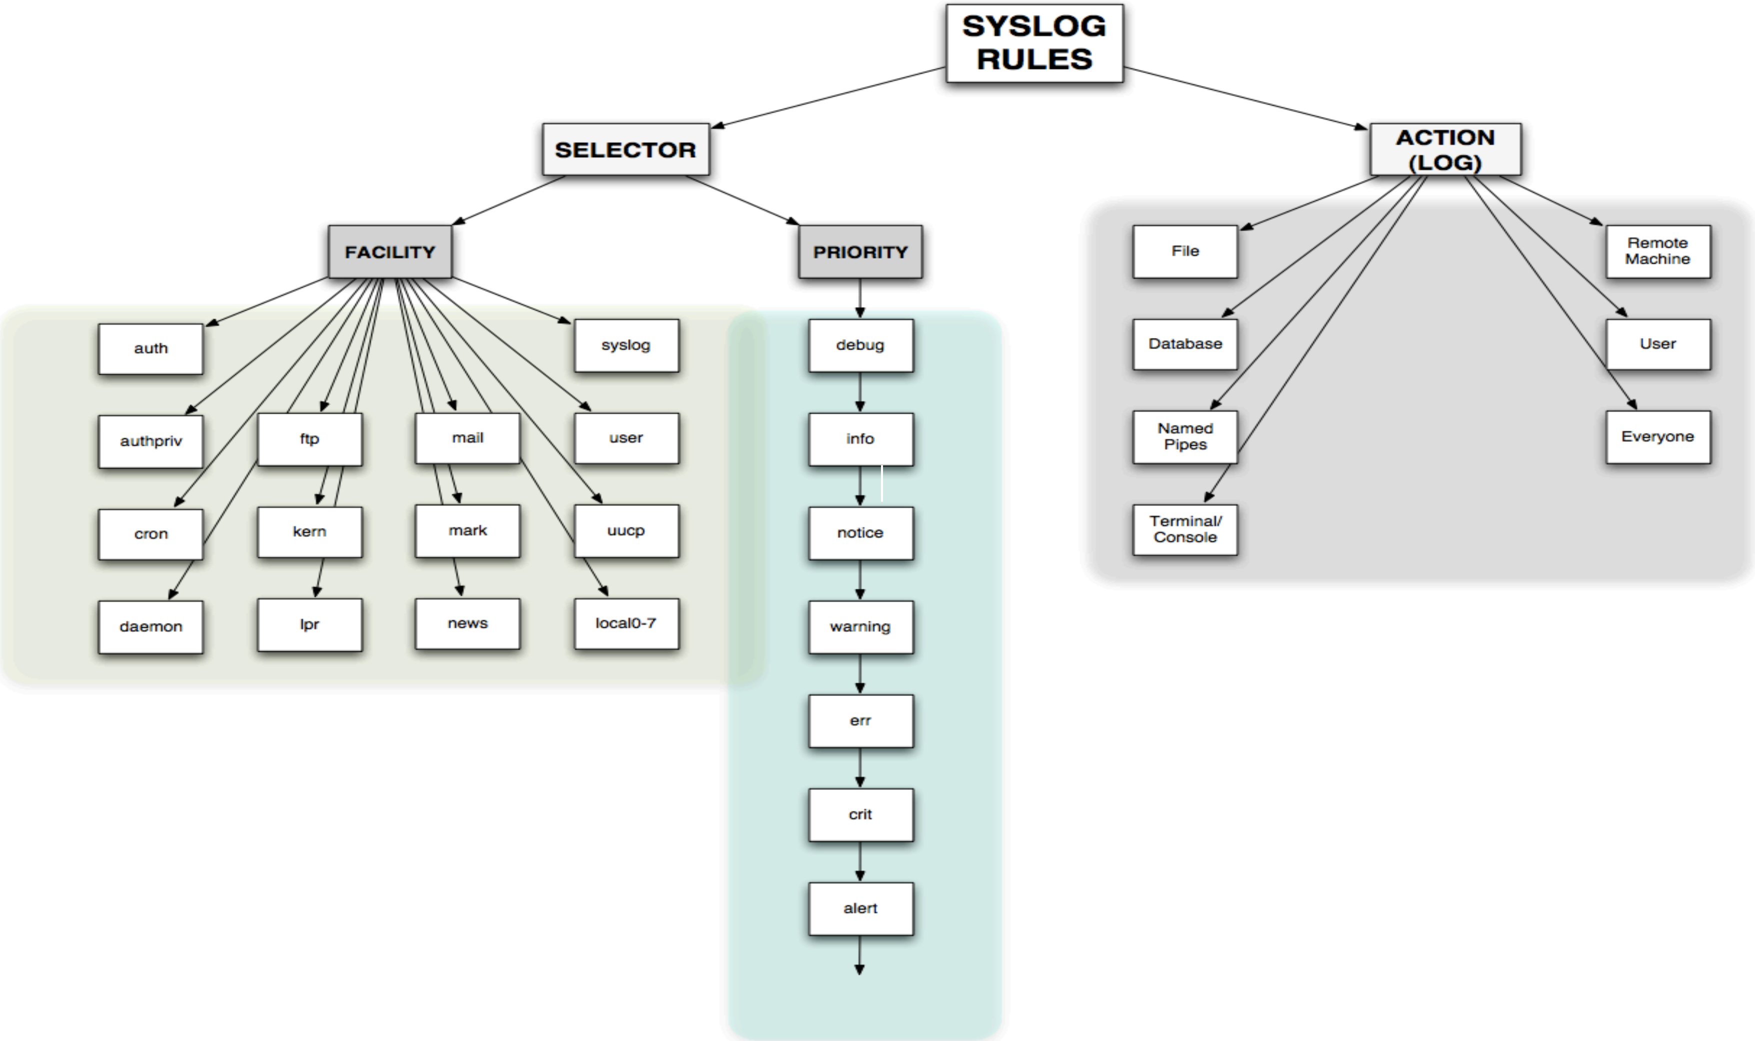

U.S. DEPARTMENT OF ENERGY

BROOKHAVEN NATIONAL LABORATORY

RSyslog & Syslog-ng

Rsyslog improves upon syslog with

- native support to write logs to a database => MySQL, Postgres, OpenTDS, SQLLite, libdbi
- the ability to send email based on a trigger
- support for TCP (improved reliability over UDP) and RELP (improved reliability over TCP)
- Encryption (SSL/TLS)
- filters supporting regular expressions
- data compression (zlib) on the fly (send & receive)
- On-demand disk spooling for both scheduled log processing and data buffering

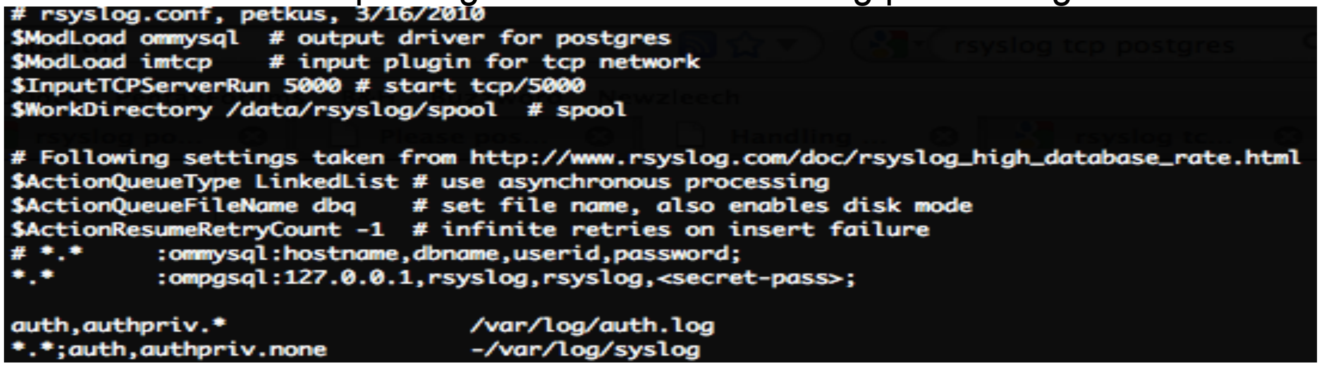

```
# rsyslog.conf, petkus, 3/16/2010
$ModLoad ommysql  # output driver for postgres
$ModLoad imtcp    # input plugin for tcp network
$InputTCPServerRun 5000 # start tcp/5000
$WorkDirectory /data/rsyslog/spool  # spool

# Following settings taken from http://www.rsyslog.com/doc/rsyslog_high_database_rate.html
$ActionQueueType LinkedList # use asynchronous processing
$ActionQueueFileName dbq    # set file name, also enables disk mode
$ActionResumeRetryCount -1  # infinite retries on insert failure
# *.*     :ommysql:hostname,dbname,userid,password;
*.*       :ompgsql:127.0.0.1,rsyslog,rsyslog,<secret-pass>;

auth,authpriv.*                 /var/log/auth.log
*.*;auth,authpriv.none          -/var/log/syslog
```

Syslog-ng

Syslog-ng competes with Rsyslog and offers

- direct database access (MSSQL, MySQL, Oracle, Postgres, SQLite3)
- high performance => 75k messages/s real time and >24GB raw logs/hour
- robust TCP / encryption
- advanced configurability => message sorting, parsing, rewriting, classification in real time
- human readable pattern matching (and regex)
- precision time-stamping => millisecond resolution

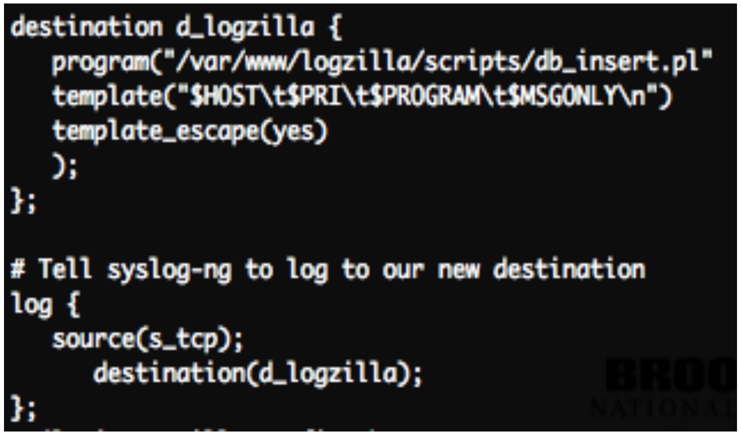

```
destination d_logzilla {
   program("/var/www/logzilla/scripts/db_insert.pl"
   template("$HOST\t$PRI\t$PROGRAM\t$MSGONLY\n")
   template_escape(yes)
   );
};

# Tell syslog-ng to log to our new destination
log {
   source(s_tcp);
      destination(d_logzilla);
};
```

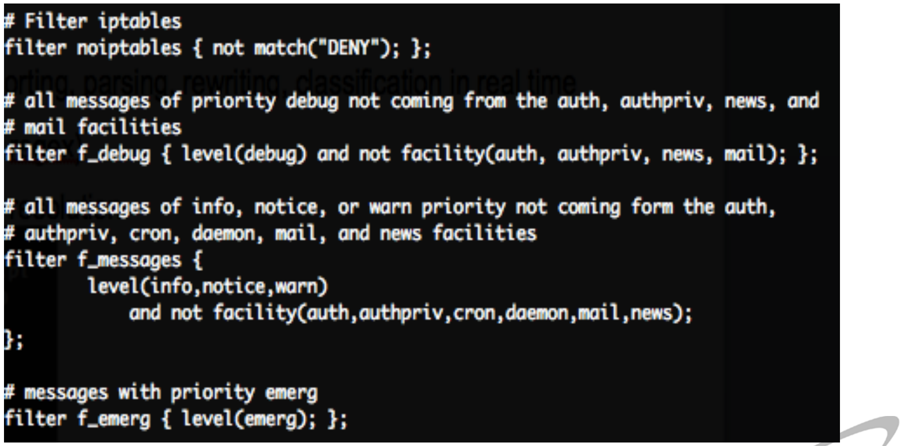

```
# Filter iptables
filter noiptables { not match("DENY"); };

# all messages of priority debug not coming from the auth, authpriv, news, and
# mail facilities
filter f_debug { level(debug) and not facility(auth, authpriv, news, mail); };

# all messages of info, notice, or warn priority not coming form the auth,
# authpriv, cron, daemon, mail, and news facilities
filter f_messages {
        level(info,notice,warn)
            and not facility(auth,authpriv,cron,daemon,mail,news);
};

# messages with priority emerg
filter f_emerg { level(emerg); };
```

Log Analysis => Splunk

What is Splunk? A system administrator search engine

- Search and analyze data from servers, apps, network appliances indexed in real time
- Generate reports, audits, sign data
- Data sources can be logs, alerts, scripts, archive files, SNMP trap data, etc.
- Configure alerts to send emails/daily reports/SNMP messages and trigger scripts
- Ability to forward data from one/many Splunk instance(s) to another (forwarder – receiver)
  - Data centralization, load-balancing, data cloning, data routing, distributed search
  - (2) flavors: Regular (forwards raw or parsed data) & Light (raw or unparsed)
- Timestamp modification/manipulation; Train to recognize new Timestamp formats
- Creation of tags to cluster groups of hosts, fields, sourcetypes, etc.
- LDAP authentication

Splunk Indexing

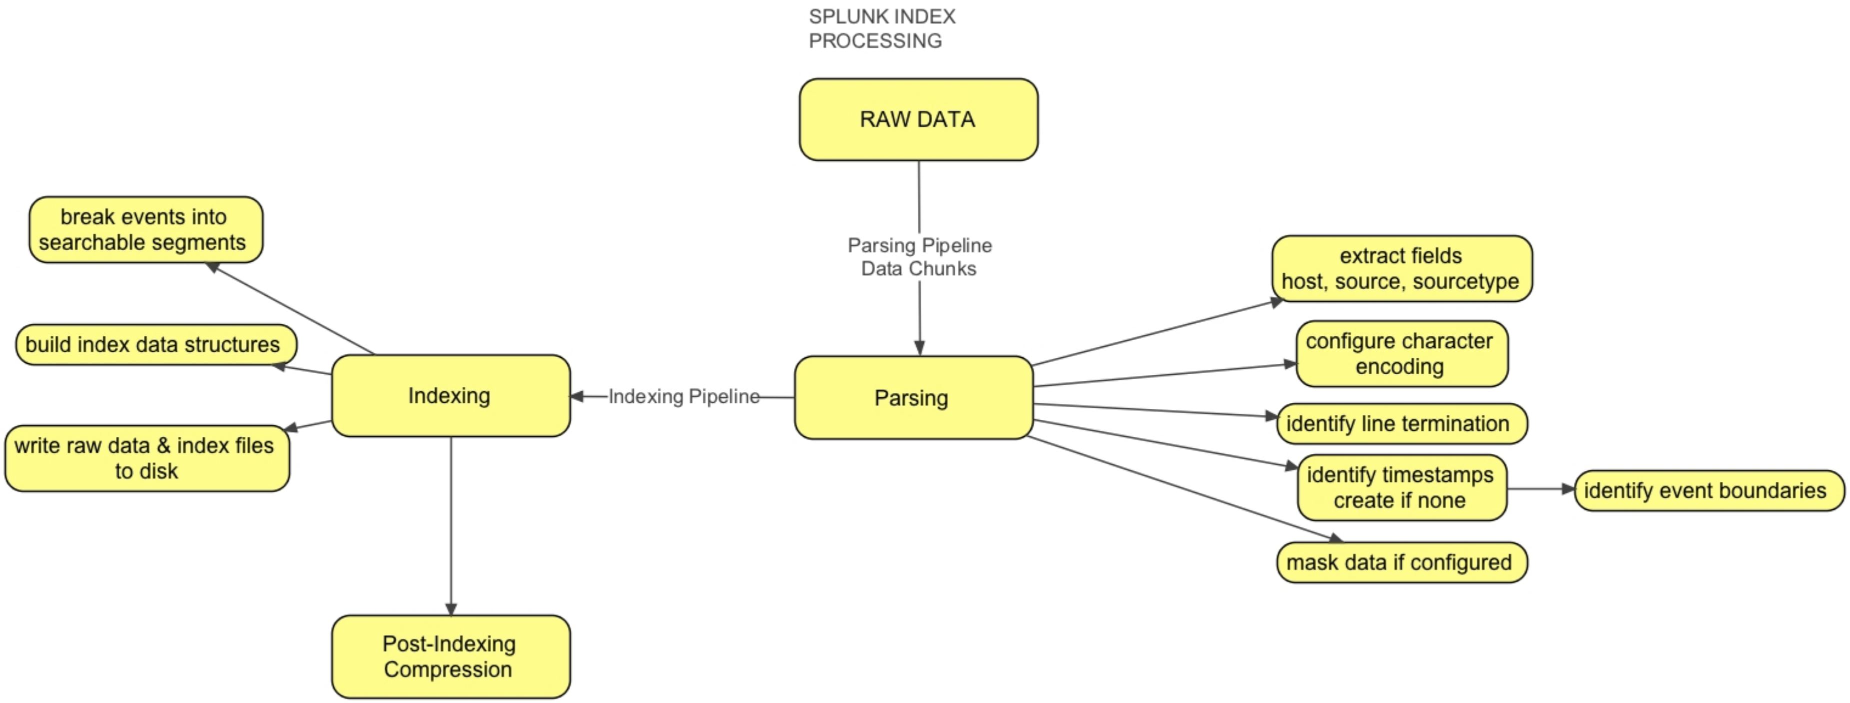
SPLUNK INDEX PROCESSING
RAW DATA
Parsing Pipeline Data Chunks
Parsing
extract fields host, source, sourcetype
configure character encoding
identify line termination
identify timestamps create if none
identify event boundaries
mask data if configured
Indexing Pipeline
Indexing
break events into searchable segments
build index data structures
write raw data & index files to disk
Post-Indexing Compression

Splunk Search

Example 1: Keep only search results that have the specified "src" or "dst" values.

```
src="10.9.165.*" OR dst="10.9.165.8"
```

Example 2: Search for events with either codes 10 or 29, and a host that isn't "localhost" and an xqp that is greater than 5

```
(code=10 OR code=29) host!="localhost" xqp>5
```

Example 3: Search for events with "404" and from host "webserver1"

```
404 host="webserver1"
```

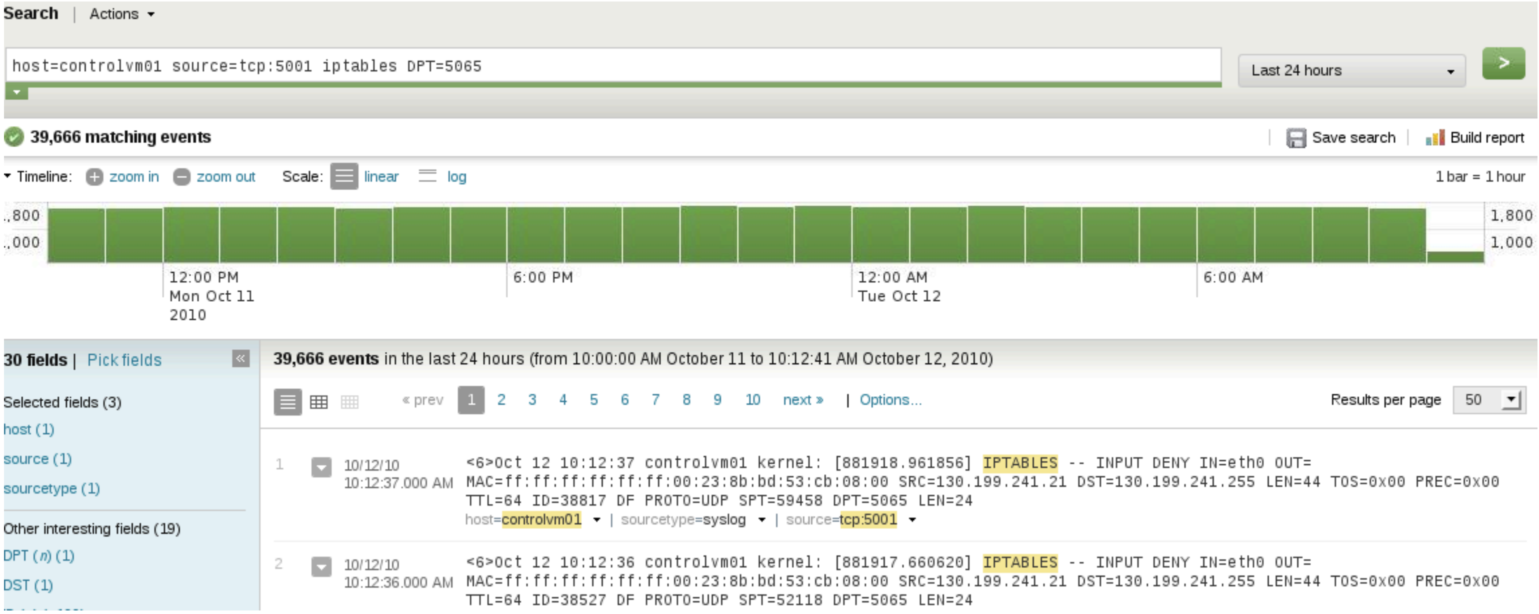

```
host=controlvm01 source=tcp:5001 iptables DPT=5065
```

Last 24 hours

39,666 matching events

30 fields

Selected fields (3)
host (1)
source (1)
sourcetype (1)

Other interesting fields (19)
DPT (*n*) (1)
DST (1)

39,666 events in the last 24 hours (from 10:00:00 AM October 11 to 10:12:41 AM October 12, 2010)

Results per page 50

```
1  10/12/10           <6>Oct 12 10:12:37 controlvm01 kernel: [881918.961856] IPTABLES -- INPUT DENY IN=eth0 OUT=
   10:12:37.000 AM    MAC=ff:ff:ff:ff:ff:ff:00:23:8b:bd:53:cb:08:00 SRC=130.199.241.21 DST=130.199.241.255 LEN=44 TOS=0x00 PREC=0x00
                      TTL=64 ID=38817 DF PROTO=UDP SPT=59458 DPT=5065 LEN=24
                      host=controlvm01 ▾ | sourcetype=syslog ▾ | source=tcp:5001 ▾

2  10/12/10           <6>Oct 12 10:12:36 controlvm01 kernel: [881917.660620] IPTABLES -- INPUT DENY IN=eth0 OUT=
   10:12:36.000 AM    MAC=ff:ff:ff:ff:ff:ff:00:23:8b:bd:53:cb:08:00 SRC=130.199.241.21 DST=130.199.241.255 LEN=44 TOS=0x00 PREC=0x00
                      TTL=64 ID=38527 DF PROTO=UDP SPT=52118 DPT=5065 LEN=24
```

Log Analysis => LogZilla, etc.

LogZilla

- Web front-end providing real-time access to syslog messages logged to MySQL
- Customized searches/report generation based on host, facility, priority, etc.
- Fast search via Sphinx => MySQL batch index and data search
    - 60+ MB/sec indexing performance
- Limited functionality compared to Splunk

LogZilla Web Interface

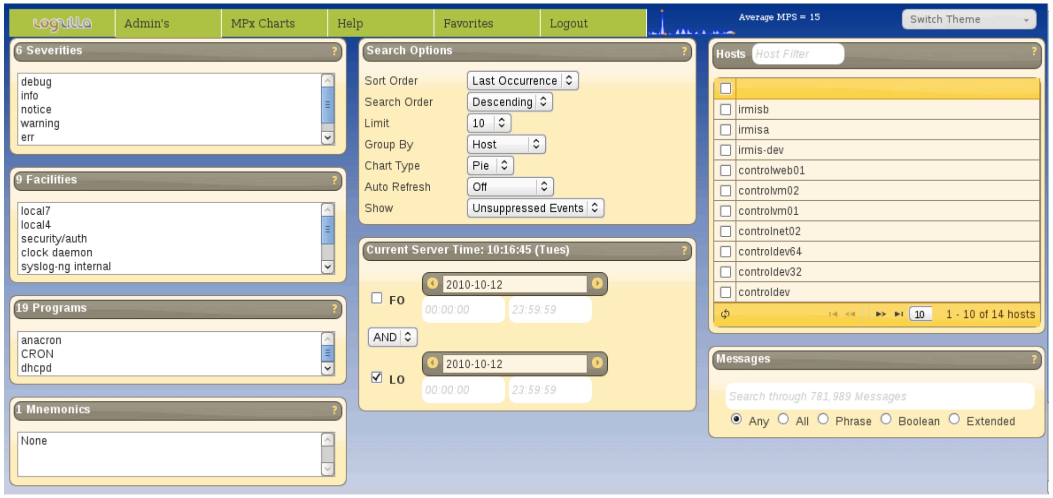

6 Severities

debug
info
notice
warning
err

9 Facilities

local7
local4
security/auth
clock daemon
syslog-ng internal

19 Programs

anacron
CRON
dhcpd

1 Mnemonics

None

Search Options

Sort Order	Last Occurrence
Search Order	Descending
Limit	10
Group By	Host
Chart Type	Pie
Auto Refresh	Off
Show	Unsuppressed Events

Current Server Time: 10:16:45 (Tues)

□ FO 2010-10-12 00:00:00 23:59:59

AND

☑ LO 2010-10-12 00:00:00 23:59:59

Hosts Host Filter

□
□ irmisb
□ irmisa
□ irmis-dev
□ controlweb01
□ controlvm02
□ controlvm01
□ controlnet02
□ controldev64
□ controldev32
□ controldev

10 1 - 10 of 14 hosts

Messages

Search through 781,989 Messages

● Any ○ All ○ Phrase ○ Boolean ○ Extended

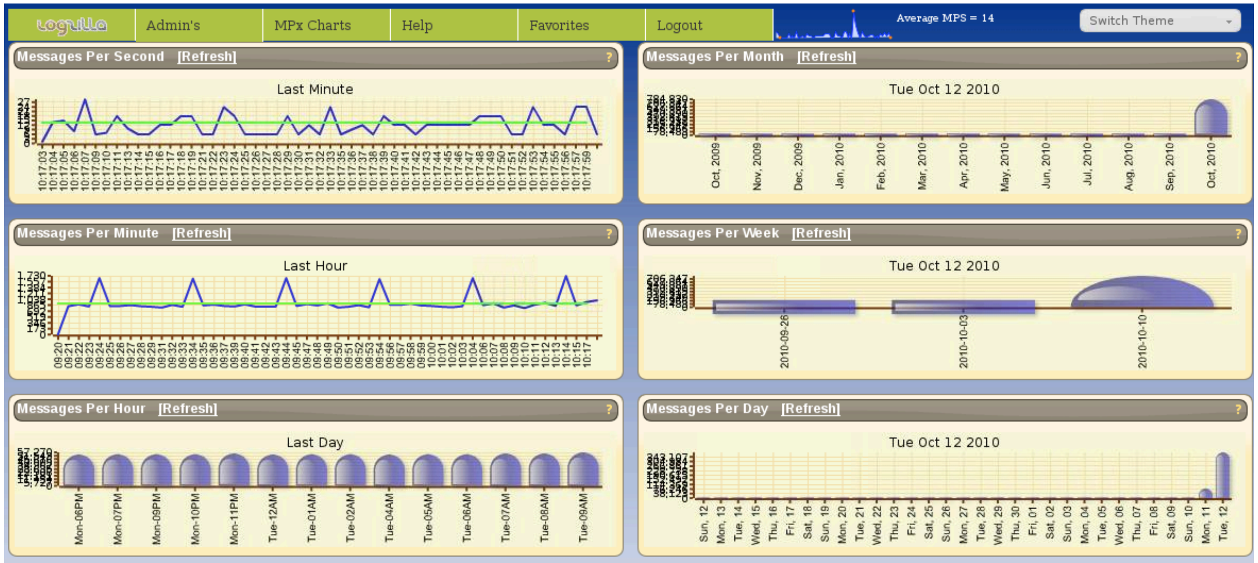

Prototype Environment at NSLS-II

In preparation of deploying server infrastructure at the production facility, we've

- Deployed a central log server (syslog-ng) collecting logs from all internal systems (~20)
  - (2) streams (to simultaneously run Splunk & LogZilla)
    - Stream A => TCP forked to both ASCII text and MySQL (LogZilla)
    - Stream B => TCP direct to Splunk DB
      - No performance bottlenecks (GbE, private net) but scale-out will require RAID array
  - Splunkd configured as a "collector"
- On client-side
  - Syslog-ng packages and configs pushed to clients via Puppet
  - Noisy logs (DHCP, Iptables, etc.) filtered-out locally but sent over wire to central log
  - Interesting clients with non-syslog app logs (NX, Virtualbox, conserver, Apache) run Splunk as a "light forwarder" to the Splunk collector on central log.

Resources

- Syslog Protocol Standard – RFC 5424 (http://tools.ietf.org/html/rfc5424)
- Rsyslog (http://www.rsyslog.com/)
- Syslog-ng (https://www.balabit.com/network-security/syslog-ng)
- Splunk (http://www.splunk.com)
- LogZilla (http://nms.gdd.net/index.php/LogZilla)
- Sphinx open-source SQL full-text search engine (http://sphinxsearch.com/)

Thanks

Questions – Comments ?

U.S. DEPARTMENT OF ENERGY

BROOKHAVEN
NATIONAL LABORATORY