Open Science Grid

Syslog-NG and Centralized Logging

Suchandra Thapa

Computation Institute

University of Chicago

HPDC 2007

Monterey, California - June 25, 2007

Three topics of interest

- Benefits of centralized logging
- Syslog-ng
- Common log formats

Why centralized logging?

- Allows logs for a cluster to be checked on a single system
- No need to log into multiple systems to figure out what is happening
- Allows for more powerful data mining on log files (more on this later)

Applications?

- What does centralizing system logs give you?
- Why go through the trouble of implementing and maintaining a centralized logging setup?

Troubleshooting

- Log messages from several systems are collated on a single system
- Related messages from different systems can be correlated and queried on a single location
- Access to multiple machines are not needed, just access to the central logging host
- Can aggregate information into a database to allow for easy searching

Troubleshooting before

- Identify log events on worker node, ce, or client
- Check other machines (gatekeepers, compute element, clients) in order to get other events that might be relevant
- May need to involve other sysadmins and people in order to access log files

Troubleshooting After

- Identify log events on compute element, gatekeeper, or client machine
- Search web interface to logging database for events that might be related in time or by content
- Can potentially be done with just access to the logging database (no need to involve multiple people to track down trivial problems)

Security Applications

- Centralization of logs allows auditing to be done more effectively
- Suspicious patterns can be more effectively picked up
- Postmortem assesments of breaches can be done more easily

Auditing Project Architecture

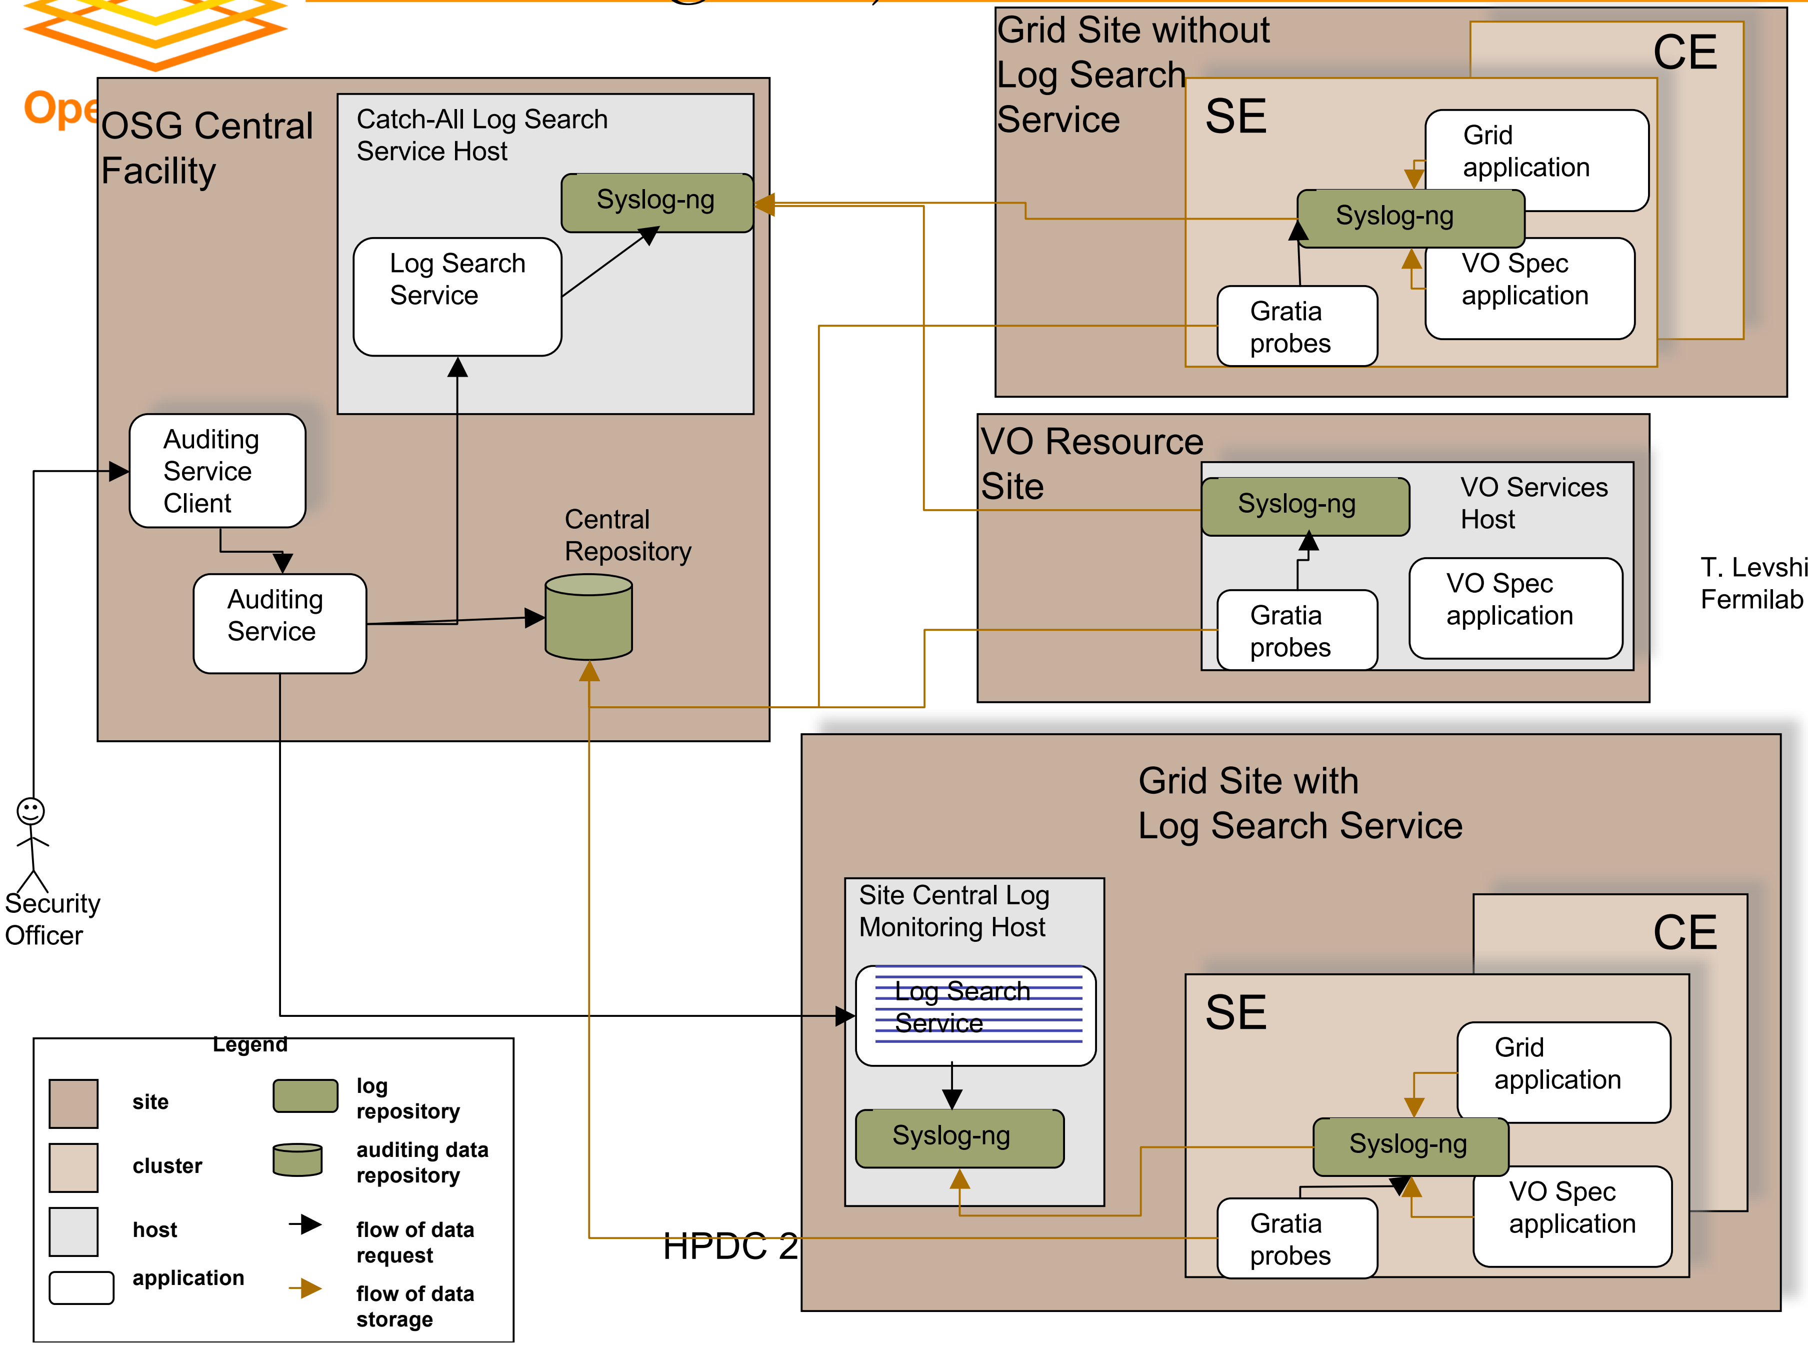

T. Levshina
Fermilab

Why Syslog-ng?

- New system logging utility
- Can replace regular syslog daemon or can be used in parallel
- More powerful facilities for filtering, formatting, and redirecting log messages
- Open source license -- can be redistributed and modified if needed

Syslog-ng logging facilities

- Can filter log messages based on log level, system host, facility, ip address or regular expressions on the message
- Can reformat and modify messages using template facilities
- Inputs can be files or sockets
- Outputs can be other hosts, files, or sockets

Open Science Grid

More syslog-ng capabilities

- Arbitary fan in and fan out for forwarding logs
- Messages can be sent to different destinations based on originating host or other filter criteria (incoming connection details, message tags, regex, etc.)
- Connections to remote hosts can be encrypted if needed using stunnel
- Messages can be altered in flight allowing hostnames to be added to ease in categorizing messages

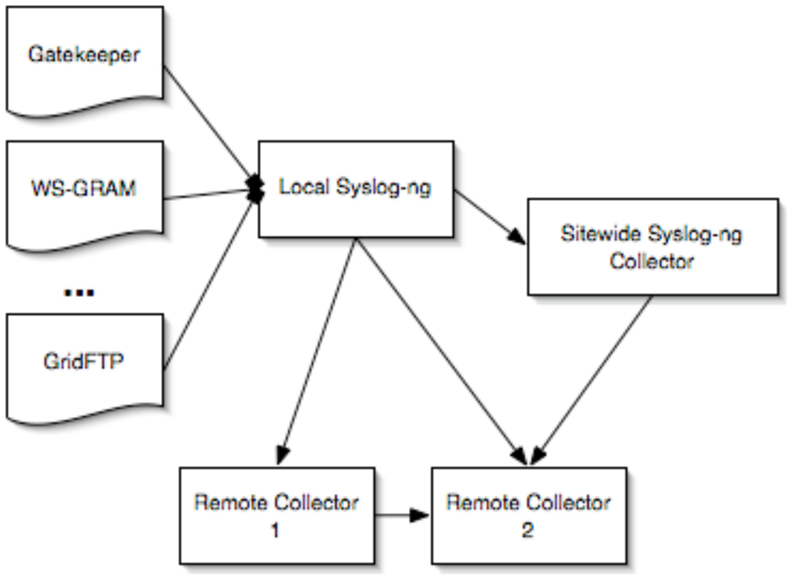

Syslog-ng for OSG

- Currently available in VDT 1.7.x releases
- Will be available in VDT 1.8.0, the next stable VDT release
- Efforts are underway to incorporate it into the OSG 0.8.0 release
- Currently used in parallel with native syslog

Current usage in OSG

- All log messages from compute elements in the validation testbed (VTB) sent to central server
- Log messages redirected from this server to 2 other servers (server hosting splunk at U of C and CEDPS system at LBNL)
- Log messages also formatted and archived in a database

DQ2 for ATLAS

- All messages get collected on central host
- 6 of 8 us tier2 sites currently sending logs to central host
- Will eventually incorporate splunk and/or php syslog-ng for log analysis and searches over the web

Php Syslog-ng Query Interface

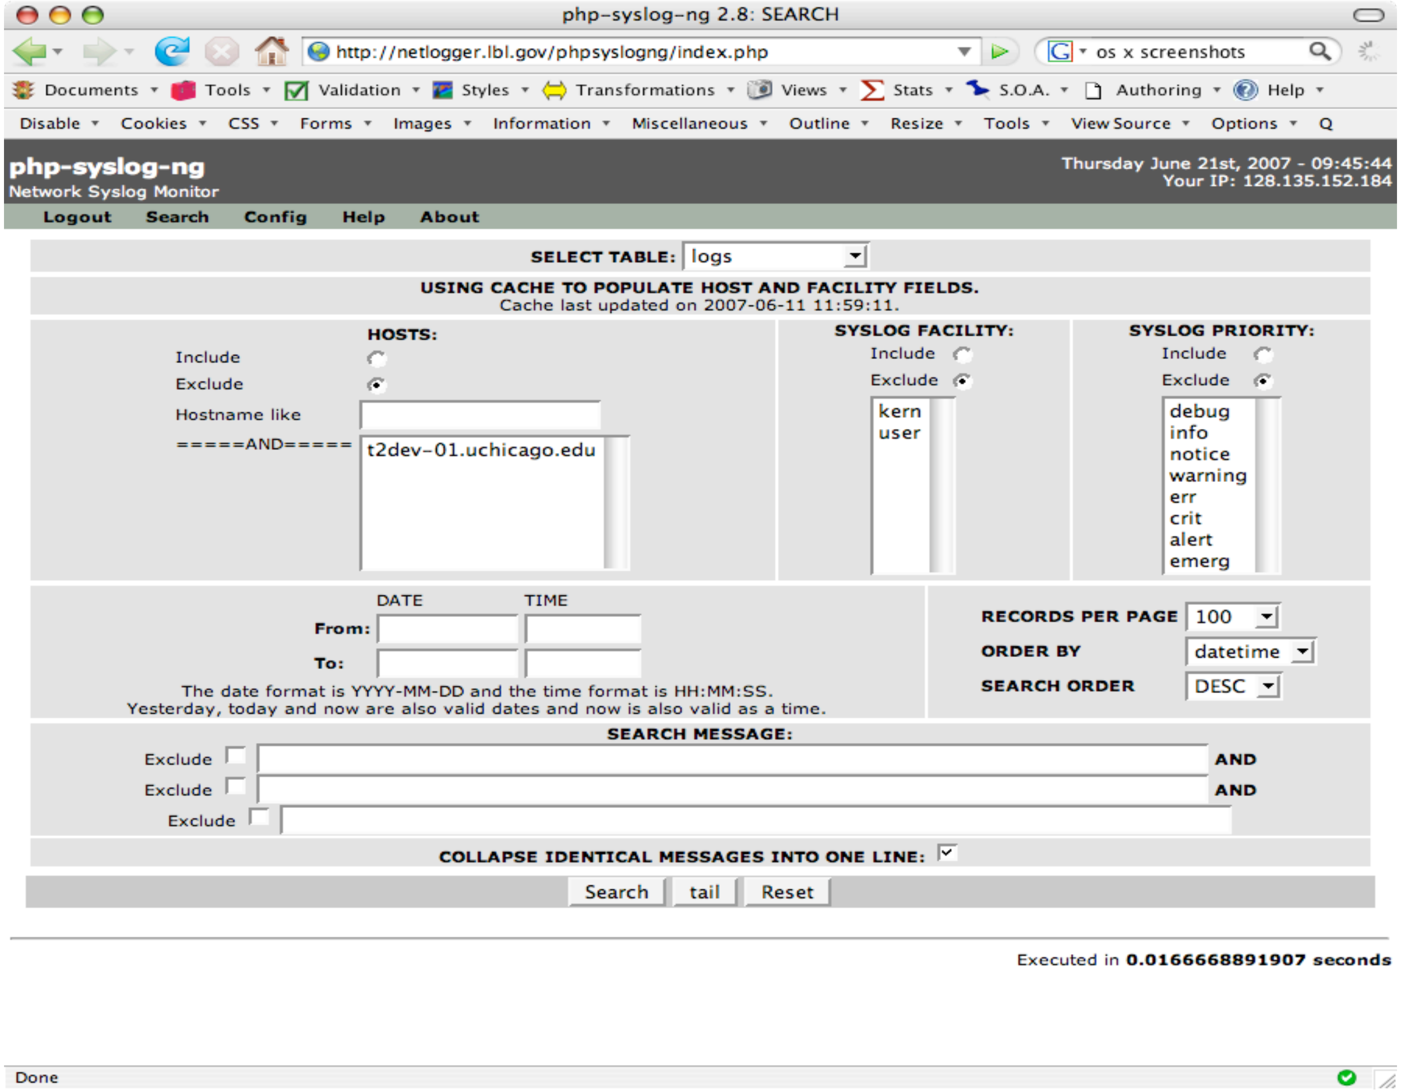
php-syslog-ng 2.8: SEARCH
http://netlogger.lbl.gov/phpsyslogng/index.php
os x screenshots
Documents ▾ Tools ▾ Validation ▾ Styles ▾ Transformations ▾ Views ▾ Stats ▾ S.O.A. ▾ Authoring ▾ Help ▾
Disable ▾ Cookies ▾ CSS ▾ Forms ▾ Images ▾ Information ▾ Miscellaneous ▾ Outline ▾ Resize ▾ Tools ▾ View Source ▾ Options ▾ Q
Logout Search Config Help About
Search tail Reset
Done

php-syslog-ng
Network Syslog Monitor

Thursday June 21st, 2007 - 09:45:44
Your IP: 128.135.152.184

SELECT TABLE: logs

USING CACHE TO POPULATE HOST AND FACILITY FIELDS.
Cache last updated on 2007-06-11 11:59:11.

HOSTS:
Include ○
Exclude ●
Hostname like
=====AND===== t2dev-01.uchicago.edu

SYSLOG FACILITY:
Include ○
Exclude ●
kern
user

SYSLOG PRIORITY:
Include ○
Exclude ●
debug
info
notice
warning
err
crit
alert
emerg

DATE TIME
From:
To:
The date format is YYYY-MM-DD and the time format is HH:MM:SS.
Yesterday, today and now are also valid dates and now is also valid as a time.

RECORDS PER PAGE 100
ORDER BY datetime
SEARCH ORDER DESC

SEARCH MESSAGE:
Exclude ☐ **AND**
Exclude ☐ **AND**
Exclude ☐

COLLAPSE IDENTICAL MESSAGES INTO ONE LINE: ☑

Executed in **0.0166668891907 seconds**

Php syslog-ng

- Php and mysql based addon to syslog-ng
- Log messages are formatted and placed in mysql database
- PHP based web interface to allow messages to be queried and displayed
- Open source license so can be modified to meet specific project needs

Php syslog-ng interface

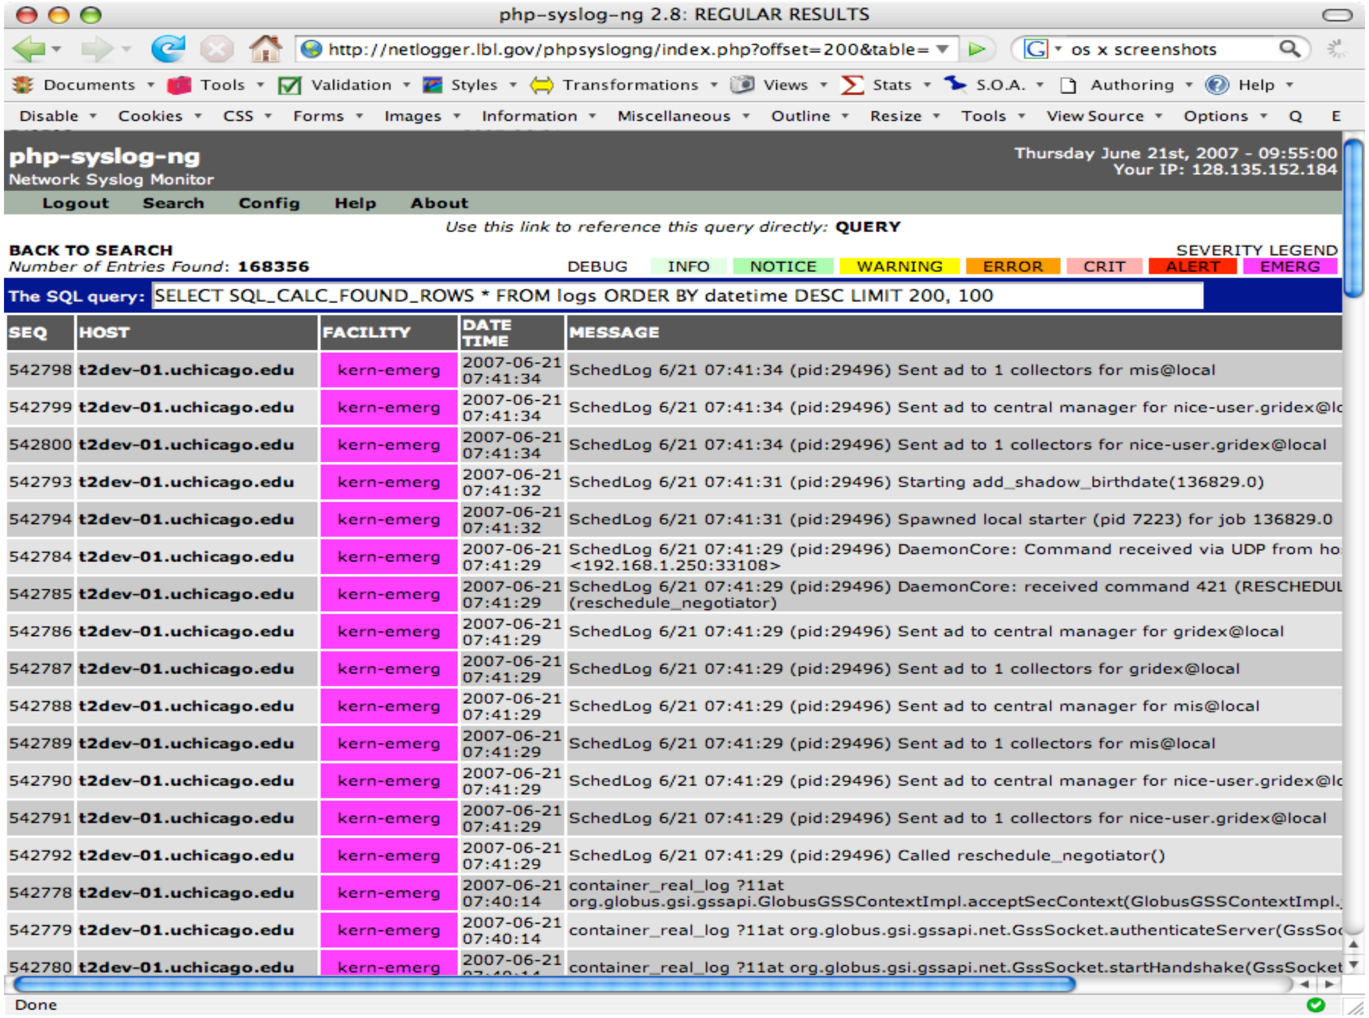

php-syslog-ng
Network Syslog Monitor

Thursday June 21st, 2007 - 09:55:00
Your IP: 128.135.152.184

Logout Search Config Help About

Use this link to reference this query directly: **QUERY**

BACK TO SEARCH
Number of Entries Found: **168356**

SEVERITY LEGEND
DEBUG INFO NOTICE WARNING ERROR CRIT ALERT EMERG

The SQL query: SELECT SQL_CALC_FOUND_ROWS * FROM logs ORDER BY datetime DESC LIMIT 200, 100

SEQ	HOST	FACILITY	DATE TIME	MESSAGE
542798	**t2dev-01.uchicago.edu**	kern-emerg	2007-06-21 07:41:34	SchedLog 6/21 07:41:34 (pid:29496) Sent ad to 1 collectors for mis@local
542799	**t2dev-01.uchicago.edu**	kern-emerg	2007-06-21 07:41:34	SchedLog 6/21 07:41:34 (pid:29496) Sent ad to central manager for nice-user.gridex@lc
542800	**t2dev-01.uchicago.edu**	kern-emerg	2007-06-21 07:41:34	SchedLog 6/21 07:41:34 (pid:29496) Sent ad to 1 collectors for nice-user.gridex@local
542793	**t2dev-01.uchicago.edu**	kern-emerg	2007-06-21 07:41:32	SchedLog 6/21 07:41:31 (pid:29496) Starting add_shadow_birthdate(136829.0)
542794	**t2dev-01.uchicago.edu**	kern-emerg	2007-06-21 07:41:32	SchedLog 6/21 07:41:31 (pid:29496) Spawned local starter (pid 7223) for job 136829.0
542784	**t2dev-01.uchicago.edu**	kern-emerg	2007-06-21 07:41:29	SchedLog 6/21 07:41:29 (pid:29496) DaemonCore: Command received via UDP from ho: <192.168.1.250:33108>
542785	**t2dev-01.uchicago.edu**	kern-emerg	2007-06-21 07:41:29	SchedLog 6/21 07:41:29 (pid:29496) DaemonCore: received command 421 (RESCHEDUL (reschedule_negotiator)
542786	**t2dev-01.uchicago.edu**	kern-emerg	2007-06-21 07:41:29	SchedLog 6/21 07:41:29 (pid:29496) Sent ad to central manager for gridex@local
542787	**t2dev-01.uchicago.edu**	kern-emerg	2007-06-21 07:41:29	SchedLog 6/21 07:41:29 (pid:29496) Sent ad to 1 collectors for gridex@local
542788	**t2dev-01.uchicago.edu**	kern-emerg	2007-06-21 07:41:29	SchedLog 6/21 07:41:29 (pid:29496) Sent ad to central manager for mis@local
542789	**t2dev-01.uchicago.edu**	kern-emerg	2007-06-21 07:41:29	SchedLog 6/21 07:41:29 (pid:29496) Sent ad to 1 collectors for mis@local
542790	**t2dev-01.uchicago.edu**	kern-emerg	2007-06-21 07:41:29	SchedLog 6/21 07:41:29 (pid:29496) Sent ad to central manager for nice-user.gridex@lc
542791	**t2dev-01.uchicago.edu**	kern-emerg	2007-06-21 07:41:29	SchedLog 6/21 07:41:29 (pid:29496) Sent ad to 1 collectors for nice-user.gridex@local
542792	**t2dev-01.uchicago.edu**	kern-emerg	2007-06-21 07:41:29	SchedLog 6/21 07:41:29 (pid:29496) Called reschedule_negotiator()
542778	**t2dev-01.uchicago.edu**	kern-emerg	2007-06-21 07:40:14	container_real_log ?11at org.globus.gsi.gssapi.GlobusGSSContextImpl.acceptSecContext(GlobusGSSContextImpl.j
542779	**t2dev-01.uchicago.edu**	kern-emerg	2007-06-21 07:40:14	container_real_log ?11at org.globus.gsi.gssapi.net.GssSocket.authenticateServer(GssSoc
542780	**t2dev-01.uchicago.edu**	kern-emerg	2007-06-21 07:40:14	container_real_log ?11at org.globus.gsi.gssapi.net.GssSocket.startHandshake(GssSocket

Splunk

- Commercial softwre used to archive and query log messages
- Web interface allows log messages to be categorized and correlated
- Messages can be queried and sorted based on categorization and other parameters
- Used at fermilab as well for internal logging collection

Open Science Grid

Splunk Interface

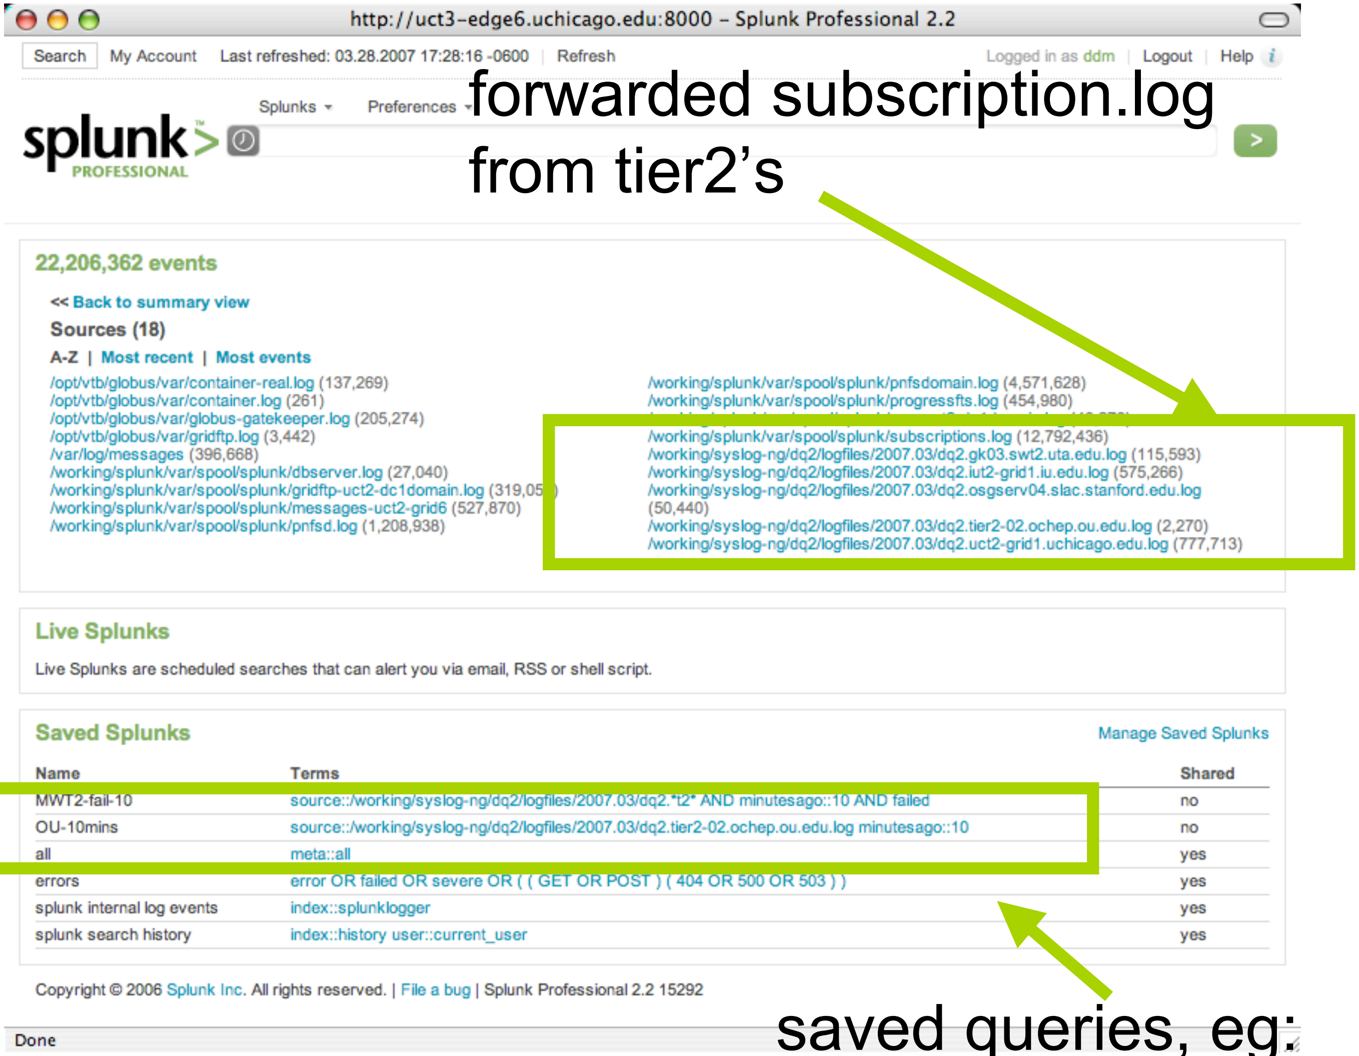
http://uct3-edge6.uchicago.edu:8000 – Splunk Professional 2.2
Search
My Account
Last refreshed: 03.28.2007 17:28:16 -0600
Refresh
Logged in as ddm
Logout
Help
Splunks
Preferences
splunk> PROFESSIONAL
>

22,206,362 events

<< Back to summary view

Sources (18)

A-Z | Most recent | Most events

/opt/vtb/globus/var/container-real.log (137,269)
/opt/vtb/globus/var/container.log (261)
/opt/vtb/globus/var/globus-gatekeeper.log (205,274)
/opt/vtb/globus/var/gridftp.log (3,442)
/var/log/messages (396,668)
/working/splunk/var/spool/splunk/dbserver.log (27,040)
/working/splunk/var/spool/splunk/gridftp-uct2-dc1domain.log (319,05)
/working/splunk/var/spool/splunk/messages-uct2-grid6 (527,870)
/working/splunk/var/spool/splunk/pnfsd.log (1,208,938)

/working/splunk/var/spool/splunk/pnfsdomain.log (4,571,628)
/working/splunk/var/spool/splunk/progressfts.log (454,980)
/working/splunk/var/spool/splunk/subscriptions.log (12,792,436)
/working/syslog-ng/dq2/logfiles/2007.03/dq2.gk03.swt2.uta.edu.log (115,593)
/working/syslog-ng/dq2/logfiles/2007.03/dq2.iut2-grid1.iu.edu.log (575,266)
/working/syslog-ng/dq2/logfiles/2007.03/dq2.osgserv04.slac.stanford.edu.log (50,440)
/working/syslog-ng/dq2/logfiles/2007.03/dq2.tier2-02.ochep.ou.edu.log (2,270)
/working/syslog-ng/dq2/logfiles/2007.03/dq2.uct2-grid1.uchicago.edu.log (777,713)

Live Splunks

Live Splunks are scheduled searches that can alert you via email, RSS or shell script.

Saved Splunks

Manage Saved Splunks

Name	Terms	Shared
MWT2-fail-10	source::/working/syslog-ng/dq2/logfiles/2007.03/dq2.*t2* AND minutesago::10 AND failed	no
OU-10mins	source::/working/syslog-ng/dq2/logfiles/2007.03/dq2.tier2-02.ochep.ou.edu.log minutesago::10	no
all	meta::all	yes
errors	error OR failed OR severe OR ((GET OR POST) (404 OR 500 OR 503))	yes
splunk internal log events	index::splunklogger	yes
splunk search history	index::history user::current_user	yes

Copyright © 2006 Splunk Inc. All rights reserved. | File a bug | Splunk Professional 2.2 15292

Done

forwarded subscription.log from tier2's

saved queries, eg: "failure" in last 10 mins

Open Science Grid

Common Log Formats

- Why is this needed?

Current situation

- Different applications format log messages differently
- Difficult to extract information reliably
- Need to use regular expressions to obtain messages
- Not scalable

Work to rationalize log messages

- CEDPS logging best practices documentation
  - Defines standard layout
  - Defines location of information (e.g. timestamps, event names) and formatting of information (e.g. timestamps in UTC with date followed by time, etc)
  - Makes suggestions as to which events should be logged
- Currently working on syslog-ng approaches
- Have scripts that transform messages from a globus to common format

Conclusion

- Centralized logging provides benefits in managing clusters
- Troubleshooting and security analysis can be made easier by having a central repository for log files
- Even more benefits if log messages use a common format
- Implementation of centralized logging can done relatively quickly and without disturbing existing logging systems
- Questions?

Links

- Syslog-ng - http://www.balabit.com/products/syslog-ng/
- OSG - https://twiki.grid.iu.edu/twiki/bin/view/Integration/SysLogNg
- CEDS Logging Best Practices - http://www.cedps.net/wiki/index.php/LoggingBestPractices

Thank You

Special thanks to: Robert Gardner, Alain Roy, Brian Tierney