Centre de Calcul de l'Institut National de Physique Nucleaire et de Physique des Particules

TYPICAL SYSLOG-NG USE-CASES

CC-IN2P3

FABIEN WERNLI

Outline

```
talk {
  infrastructure();
  altertative { ... };
  channel {
    storing(); alerting(); enriching();
  };
  channel {
    misc (); automation(); monitoring();
  };
  channel {
    appendix { ... };
    flags(if-time-permits);
  };
};
```

ARCHITECTURE

Architecture

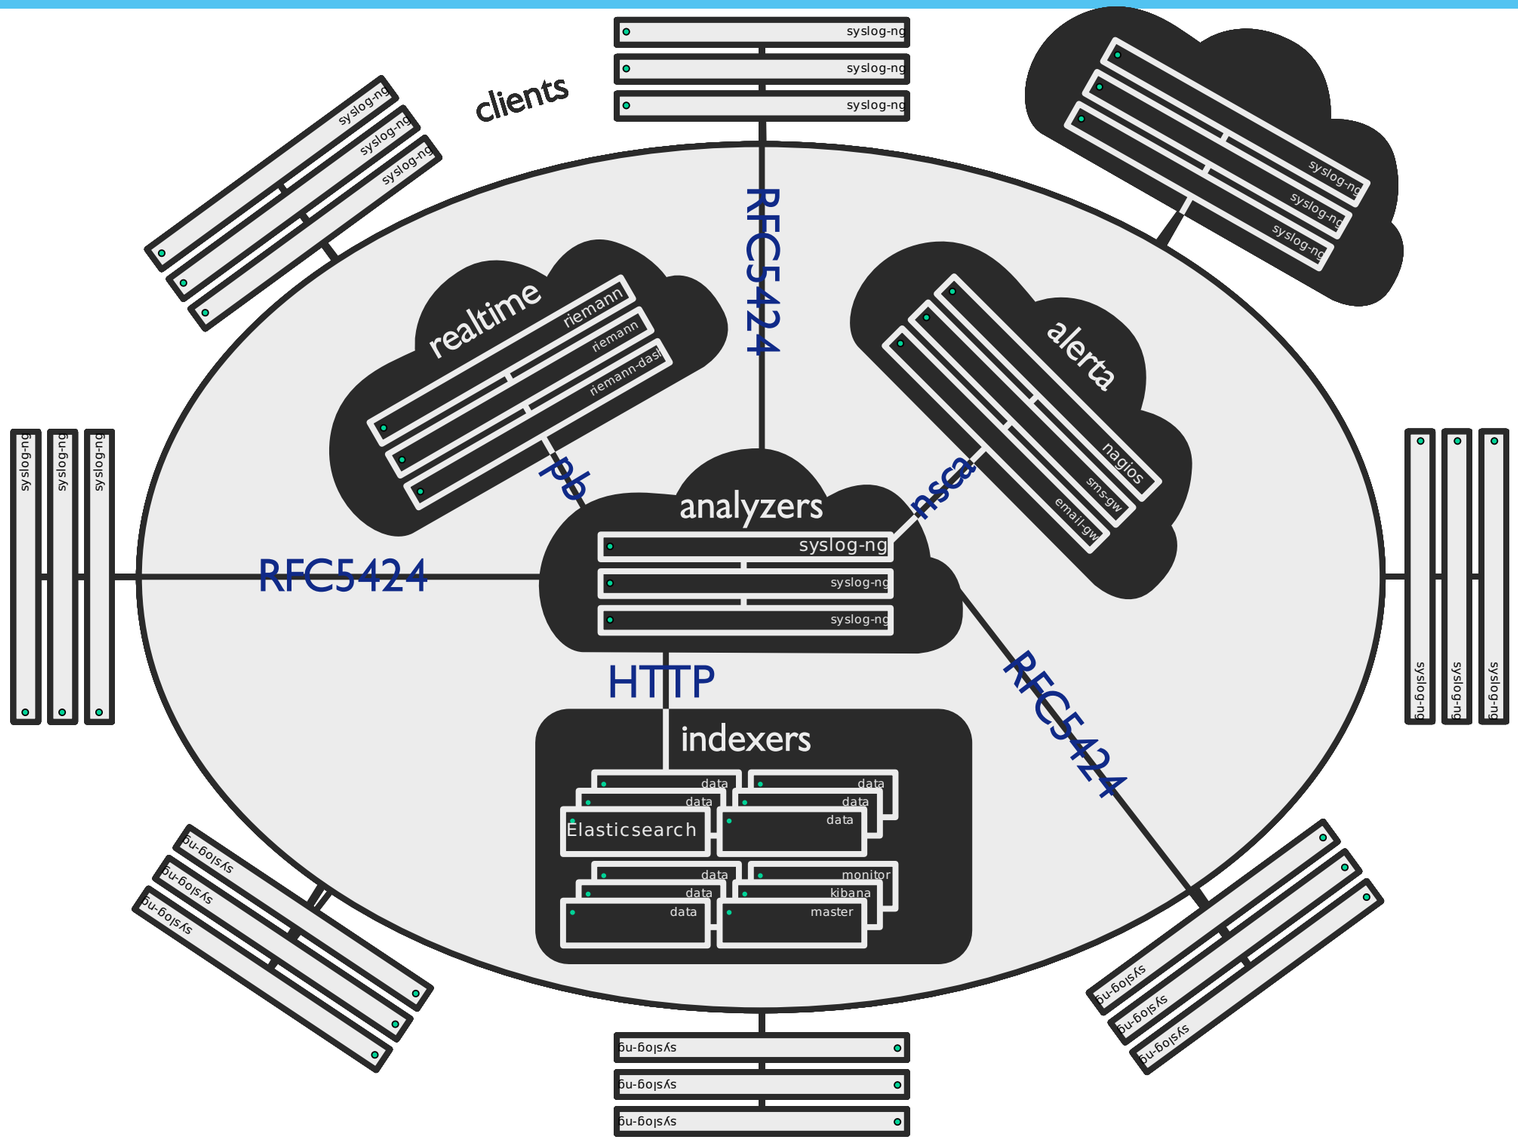

Infrastructure

Role	Platform	Max Usage
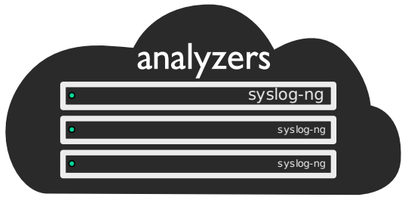	3 VMs 8GB/8CPU	1GB/2CPU
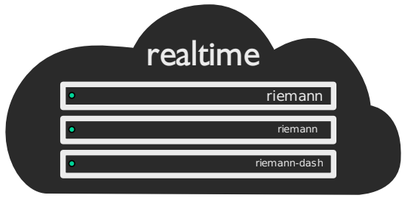	5 VMs 8GB/8CPU	6GB/4CPU
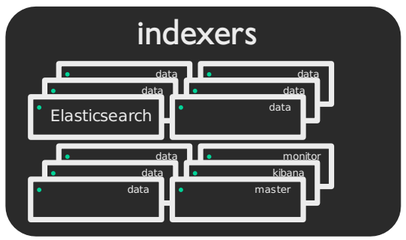	9 BMs 48GB/12CPU	32GB/10CPU

WHY SYSLOG-NG?

- Flexible
- Portable
- Fast
- Low resource footprint
- Friendly (ML, issues, PRs, IRC/gitter)

Alternatives

Alternative	Pros	Cons
rsyslog	infiltration, performance	config, documentation
logstash	community, flexible	speed, footprint
Elastic Beats	lightweight, portable	new, unflexible

STORING LOGS

ELASTICSEARCH DESTINATION

- uses JNI (libjvm.so)
- supported protocols: http(s), transport, and node
- searchguard and https implemented by CC-IN2P3

Elasticsearch

```
destination d_elasticsearch {
  elasticsearch2(
    client-lib-dir("/usr/share/elasticsearch/plugins/search-guard-5/*.jar:/usr/share/elasticsearch/lib/")
    client-mode("https")
    concurrent-requests("16")
    disk-buffer(
      dir("/var/lib/syslog-ng-disq/")
      disk-buf-size(53687091200)
      mem-buf-size(1073741824)
    )
    flush-limit('1024')
    index("${__es_index:-syslog}-${YEAR}.${MONTH}.${DAY}")
    port('9200')
    server("node01 node02 node03 node04 node05")
    java_keystore_filepath("/etc/syslog-ng/coloss-analyzer-keystore.jks")
    java_keystore_password("terces")
    java_truststore_filepath("/etc/elasticsearch/coloss/truststore.jks")
    java_truststore_password("terces")
    http_auth_type("clientcert")
    skip-cluster-health-check("yes")
    template("$(format-json -s all-nv-pairs --rekey .SDATA.* --shift 7)")
    time-zone("UTC")
    type("${__es_type:-syslog}")
  );
};
```

SUGGESTIONS

- know your libjvm.so

```
Error initializing message pipeline;
```

- use large disq buffer (cluster restarts)
- bulk flush-limit('1024')
- multithread concurrent-requests("16")

ALERTING

NAGIOS

```
destination d_nagios {
  program(
    /opt/bin/send_nsca
    flush-lines(1)
    template(
      "${nagios.host:-UNDEF}\t${nagios.service:-UNDEF}\t${nagios.status:-0}\t${nagios.messa
    )
  );
};
```

EMAIL

```
destination d_email {
  smtp(
    host("localhost"), port(25),
    from("syslog_ng" "noreply@cc.in2p3.fr"),
    to("${email.to}"),
    subject("[syslog_ng] ${PROGRAM} ${HOST} ${email.subject:-N/A}"),
    body("${email.body:-N/A}")
  );
};
```

RIEMANN

```
destination d_riemann {
  riemann(
    server("riemann.cc.in2p3.fr"),
    port(5555),
    type("tcp"),
    flush-lines(1),
    ttl("${ttl:-300}"),
    metric("$metric"),
    state("${state:-ok}"),
    attributes(
      scope(all-nv-pairs),
      key(".SDATA.*"
        rekey( shift(7) )
      ),
    ),
```

Alerting: routing

ROUTING

Alerting: routing

Pattern Matching (patterndb)

```
<rules>
 <rule provider='puppet' id='1311f61b-c2f5-4510-8b3b-6b263c9bd46e' class='system'>
  <patterns>
   <pattern>@ESTRING:: @@STRING::@@SET:: @@ESTRING:: @@ESTRING:: @@ES
  </patterns>
  <values>
   <value name='state'>warning</value>
   <value name='ttl'>7200</value>
  </values>
  <examples>
   <example>
    <test_message program='afs_fs'>Thu Jun  2 00:09:20 2016 Partition /vicepaa that conta
    <test_values>
     <test_value name='afs.partition'>vicepaa</test_value>
```

Alerting: routing

Filter

```
filter f_to_email {
  tags("f_to_email");
};
```

Log Path

```
log {
  source(s_system);
  source(s_network);
  ...
  parser(p_patterndb);
  ...
  log {
    filter(f_to_email);
    destination(d_email);
  };
  ...
}
```

EXAMPLES

GPFS

```
[W] allocLogBufs:no memory wait 5 seconds, 31 so far
[W] Inode space 41 in file system sps_hep is approaching the limit for the maximum number o
```

Node reboots

```
BIOS-e820: 0000000000100000 - 00000000cf379000 (usable)
Kernel command line: ro root=/dev/mapper/rootvg-root rd_NO_LUKS KEYBOARDTYPE=pc K
Initializing cgroup subsys blkio
```

FS

```
The number of I/O errors associated with a ZFS device exceeded
Buffer I/O error on device dm-6, logical block 64557071
Filesystem dm-6: xfs_log_force: error 5 returned.
```

ENRICHING LOGS

Enriching Logs

Puppet facts

```
rewrite r_sdata_facter {
  set("RedHat",
      value(".SDATA.facter.osfamily")
  );
  set("OpenStack Nova",
      value(".SDATA.facter.productname")
  );
};
```

CMDB

```
rewrite r_sdata_cmdb {
  set("workernode"
      value(".SDATA.cmdb.role")
  );
};
```

Enriching Logs

Send Structured Data using RFC5424

```
log {
  source(s_system);
  rewrite(r_sdata_facter);
  rewrite(r_sdata_cmdb);
  destination{
    network(
      "logs.cc.in2p3.fr"
      flags(syslog-protocol)
    );
  };
};
```

Enriching Logs

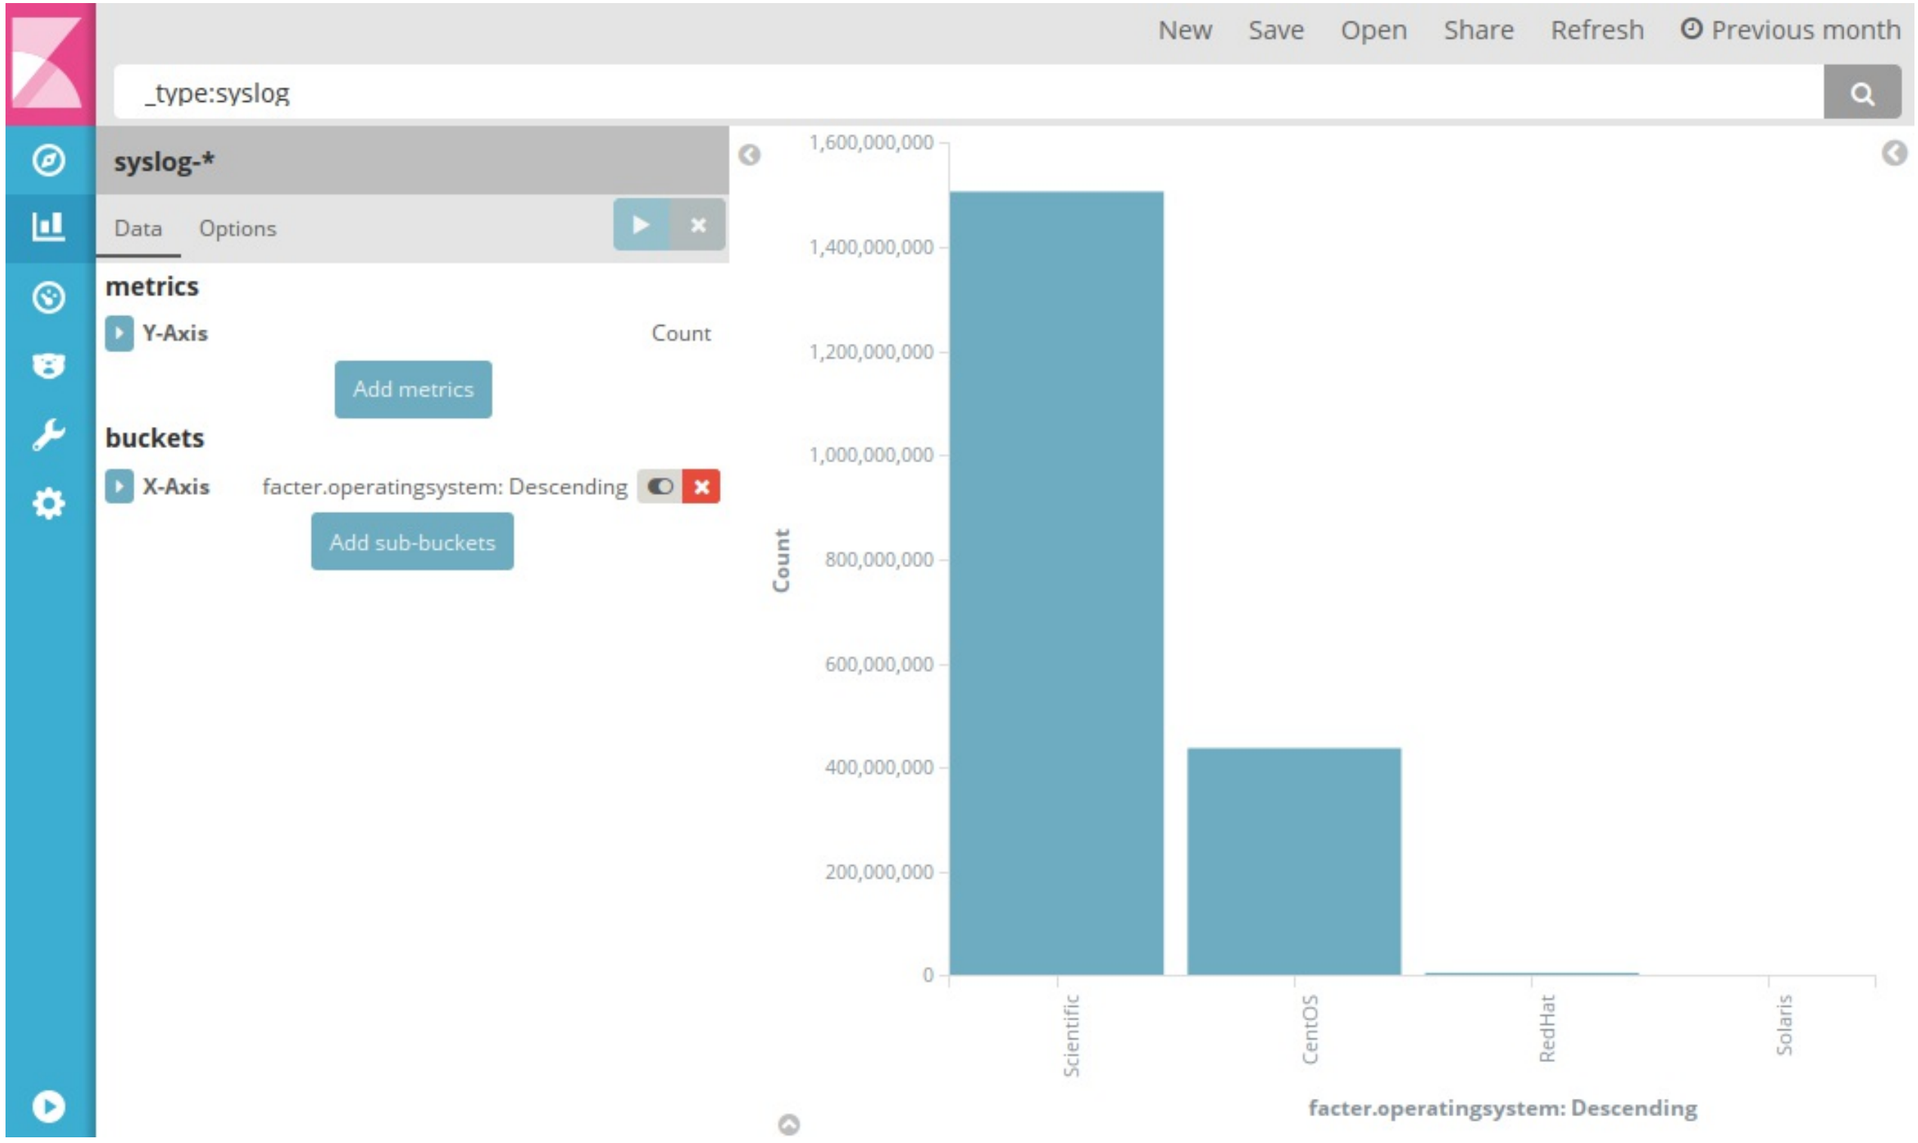

Enriching Logs

Or use an external file (available since 3.8.1)

```
parser p_uppet_facts {
  add-contextual-data(
    selector("$HOST"),
    database("/path/to/puppet-facts.csv"),
  );
};
```

MISC

Mailbox source

```
source s_mbox {
  channel {
    source {
      mbox("/var/spool/mail/syslog_ng");
    };
    parser {
      json-parser(
        template("$(python mbox)")
      );
    };
    junction {
      channel {
        date-parser(
          format("%a, %d %b %Y %T %z")
          template("${Date}")
```

Mailbox source - but why?

SERIOUSLY?

- sure, for syslog-unfriendly tools
- and quick'n'dirty solutions
- "You Know, for Search"
- Ex: appliances, electrical equipment, ...
- Ex: yum, (ana)cron, ...

Mailbox source - Examples

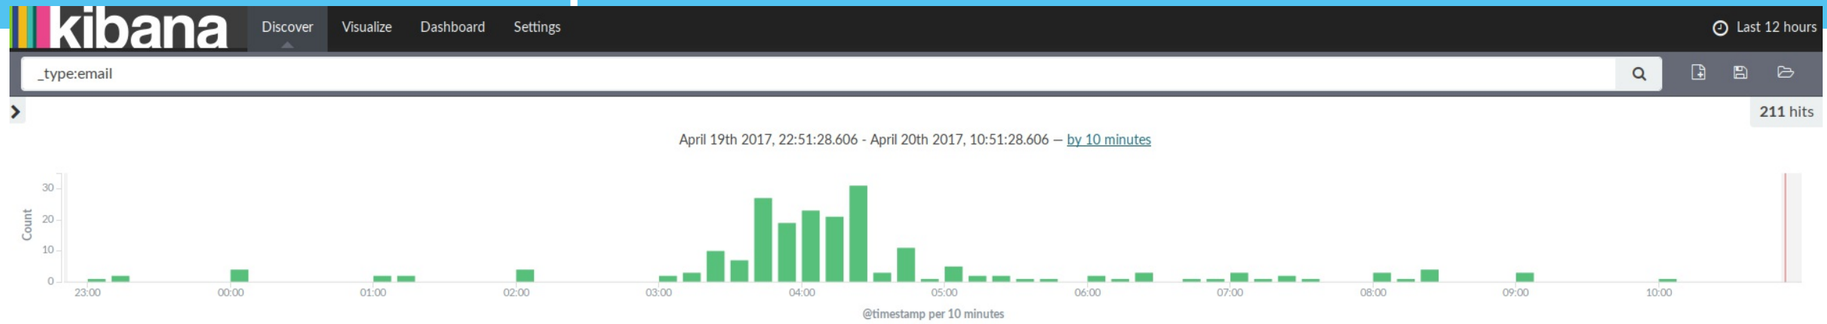

Time	From	To	Subject	message
April 20th 2017, 10:07:04.000	"(Cron Daemon)" <root@ccosvm0870.in2p3.fr>	root@ccosvm0870.in2p3.fr	Cron <root@ccosvm0870> run-parts /etc/cron.hourly	/etc/cron.hourly/0yum-hourly.cron: Repository mongodb-org-3.2 is listed more than once in the configuration
April 20th 2017, 09:07:02.000	"(Cron Daemon)" <root@ccosvm0870.in2p3.fr>	root@ccosvm0870.in2p3.fr	Cron <root@ccosvm0870> run-parts /etc/cron.hourly	/etc/cron.hourly/0yum-hourly.cron: Repository mongodb-org-3.2 is listed more than once in the configuration
April 20th 2017, 09:07:02.000	"(Cron Daemon)" <root@cctest30.in2p3.fr>	root@cctest30.in2p3.fr	Cron <root@ccosvm0964> run-parts /etc/cron.hourly	/etc/cron.hourly/0yum-hourly.cron: Repository updates is listed more than once in the configuration Repository extras is listed more than once in the configuration Repository centosplus is listed more than once in the configuration
April 20th 2017, 09:07:02.000	"(Cron Daemon)" <root@ccsquidli02.in2p3.fr>	root@ccsquidli02.in2p3.fr	Cron <root@ccsquidli02> run-parts /etc/cron.hourly	/etc/cron.hourly/logrotate: error: error opening /var/log/squid/squid-log.3.json.1: No such file or directory
April 20th 2017, 08:28:03.000	Anacron <root@cctest30.in2p3.fr>	root@cctest30.in2p3.fr	Anacron job 'cron.daily' on ccosvm0863	/etc/cron.daily/0yum-daily.cron: warning: /etc/cloud/cloud.cfg created as /etc/cloud/cloud.cfg.rpmnew warning: /etc/syslog-ng/syslog-ng.conf created as /etc/syslog-ng/syslog-ng.conf.rpmnew

```
libtool-ltdl              x86_64  2.4.2-22.el7_3         centos-updates   49 k
lvm2                      x86_64  7:2.02.166-1.el7_3.4   updates         1.1 M
-    - 12/14692: sysadmin@cc.in2p3.fr   Yum: Updates installed on ccosvm0863 -- (53%)
```

HPSS Correlation

```
Record type=EVENT, Event time=2017/04/21 09:22:35 CEST, Severity=NONE
Subsystem=MPSR, Message#=63, Error code=0 Desc
name=MPS1,
Routine=mps_DiskMigr ( line 6715 ) PID=52853,
Node=xxxx.in2p3.fr, User=
Type=OPERATION INITIATION, Object Class=37, Request Id=0
Disk migration start (SClassId 14, SubSysId 1).
```

```
Record type=EVENT, Event time=2017/04/21 09:48:05 CEST, Severity=NONE
Subsystem=MPSR, Message#=64, Error code=0
Desc name=MPS1,
Routine=mps_RecordStats ( line 2321 )
PID=52853,
Node=xxxx.in2p3.fr, User=
Type=OPERATION COMPLETION, Object Class=37, Request Id=0
Disk migration end (SClassId 14, SubSysId 1, Files 37, Bytes 578924559613, Errors 0).
```

HPSS Correlation

TABLE

0 to 100 of 500 available for paging

@timestamp	hpss.msg_id	hpss.message	hpss.sclass_id	hpss.subsys_id	hpss.stat.migratedbytes	hpss.stat.migration_rate	hpss.stat.purgedbytes
2017-04-27T07:30:39.000+02:00	MPSR0066	Disk purge end	14	3			5320122302464
2017-04-27T07:24:40.000+02:00	MPSR0066	Disk purge end	12	1			1674634592256
2017-04-27T07:22:05.000+02:00	MPSR0066	Disk purge end	14	5			4300336005120
2017-04-27T06:29:00.000+02:00	MPSR0064	Disk migration end	14	5	921968717959	1092	
2017-04-27T06:25:46.000+02:00	MPSR0064	Disk migration end	10	5	1922680326	13	
2017-04-27T06:12:00.000+02:00	MPSR0064	Disk migration end	12	4	168497876523	354	
2017-04-27T06:10:40.000+02:00	MPSR0066	Disk purge end	14	3			5345623670784
2017-04-27T05:55:16.000+02:00	MPSR0064	Disk migration end	14	4	11650244952	68	
2017-04-27T05:29:17.000+02:00	MPSR0066	Disk purge end	14	1			3752727674880
2017-04-27T05:05:24.000+02:00	MPSR0064	Disk migration end	12	5	265135502343	458	

Correlating events using group-by()

AUTOMATION

ihrwein.syslog-ng

```
syslog_ng_client_destinations:
  - "log.cc.in2p3.fr":
      proto: udp
      port: 514
      filters:
        - f_syslog
        - f_all_but_debug
```

ccin2p3/patterndb

```
patterndb::simple::action:
  CPU_SOFT_LOCKUP:
    rule: 7ae85d0e-0e42-48d8-9992-2f125c9ae310
    rate: "1/300"
    condition: '"$(context-length)" >= "10"'
    message:
      inherit_properties: TRUE
      tags:
        - f_to_email
      values:
        email.to: p'tach@cc.in2p3.fr
        email.subject: soft lockups on CPU
        email.body: '${MESSAGE}'
        state: warning
        PROGRAM: 'soft_lockup'
```

Contributions welcome!

ccin2p3/syslog_ng

```
syslog_ng::config:
  version:
    content: "@version: %{syslog_ng_version}"
    order: '02'
  scl:
    content: '@include scl.conf'
    order: '03'
syslog_ng::filter:
  f_messages:
    params:
      - level: [ "info..emerg" ]
```

Contributions welcome!

MONITORING

syslog-ng-ctl

syslog-ng-ctl

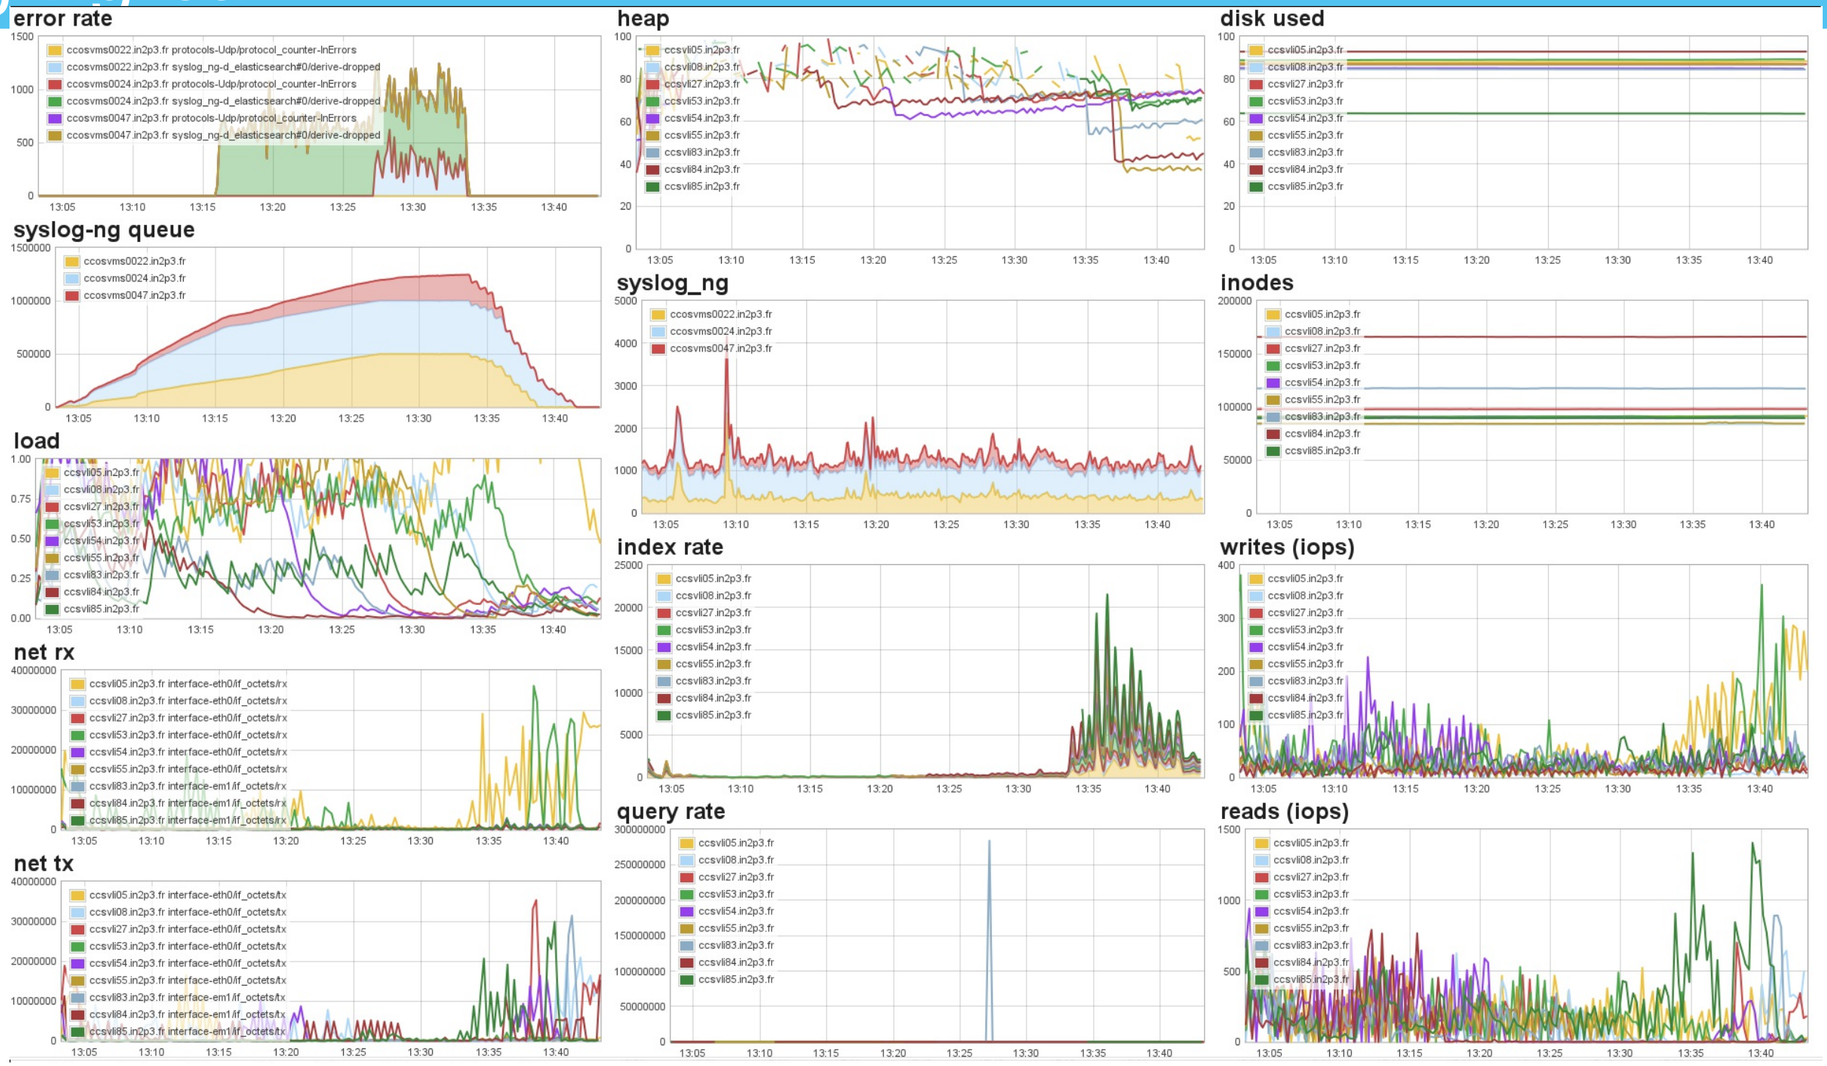

GOTCHAS/TRAPS

- EPEL version not supporting X
  - workaround: use unofficial packages or build from source
- EL7 hard RPM deps on libvirt, cloud-init
- /etc/logrotate.d/syslog conflict with rsyslog.rpm
- blocking syslog() syscalls
  - ldap workaround: owner(-1) group(-1)
  - dns workaround: /etc/hosts or use IPs

OTHER DESTINATIONS

- HTTP (libcurl)
- kafka (2 impl.)
- HDFS (java)
- SQL
- collectd (WIP)

SYSLOG-NG REFERENCES

- unofficial-ng packages
- documentation
- the libjvm.so problem
- ldap deadlock
- dns deadlock
- syslog-ng java drivers

OTHER REFERENCES

- elastic beats
- rsyslog
- logstash
- Search Guard
- ansible module
- puppet patterndb
- puppet syslog_ng

APPENDICES

Mailbox source - multiline

```
block source mbox(filename()) {
    file(
        "`filename`"
        log-msg-size(10000000)
        log-fetch-limit(1)
        flags(no-parse)
        multi-line-mode(prefix-suffix)
        multi-line-prefix('^From ')
      );
};
```

Mailbox source - python

```
def mbox(logmsg):
  lines = logmsg.MSG.splitlines()
  first_line = lines.pop(0)
  if not re.match(r'^From ', first_line):
    return json.dumps({"mbox.error": "doesn't look like an email"})

  first_line = first_line.split(None,2)
  out = {}
  out['Envelope'] = first_line[1]
  out['Isodate'] = first_line[2]
  line = "foo"
  while (len(line) > 0):
    line = lines.pop(0)
    if re.match(r'^.*: ',line):
      d = line.split(': ')
```