Evaluation of a new Grid Engine Monitoring and Reporting Setup

Thomas Finnern

HELMHOLTZ ASSOCIATION

DESY

Abstract for Conference

- Title
  - Evaluation of a new Grid Engine Monitoring and Reporting Setup
- Summary
  - Dashboards and Event Correlation for Grid Engine with Splunk
- Content
  - Splunk is a commercial software platform for collecting, searching, monitoring and analyzing machine data providing interactive real-time dashboards integrating multiple charts, reports and tables. We have been working on a grid engine setup based on the free branch supporting standard reporting and simple job and fairshare debugging with easy chart generation and smart event correlation features. On top of this we try understand the added value of the enterprise branch supporting integrated user authentication and role-based access controls. There is a plan to share our work in a public available Splunk grid engine app.

Outline

- Grid Engine Data and Splunk
  - Integration into Splunk
  - Job Data: submit, prolog, epilog plus Infos on Errors
  - SoGE (Son of Grid Engine) Data: Messages and Accounting
  - System Data: Resources and Usage (numbers, projects)
- Needs
  - Job Inspection (View 1)
  - Accounting and Reports (View 2)
  - Weekly Project Views (View 3)
  - Realtime System Data (View 4)
  - Role Based Access
  - Grid Engine App …
- Status and Conclusions

Integration into Splunk

- Splunk Indexing
  - Index: Separate Directory used by all SoGE Data
  - Index: Event indexing Time, Originating Host and Sourcetype
  - ASCII Store for Data Sets
  - Field Mapping on the Fly during Analysis
- Splunk Data Input
  - Extra Syslog Port mapped SoGE Index

    Submit, Jobstart, Jobend, System Data
  - Splunk Forwarder

    Running on Grid Engine Master
    Reliable Upload to Splunk Server
    Configured for Message File
    Configured for Accounting File

Grid Engine Data and Splunk

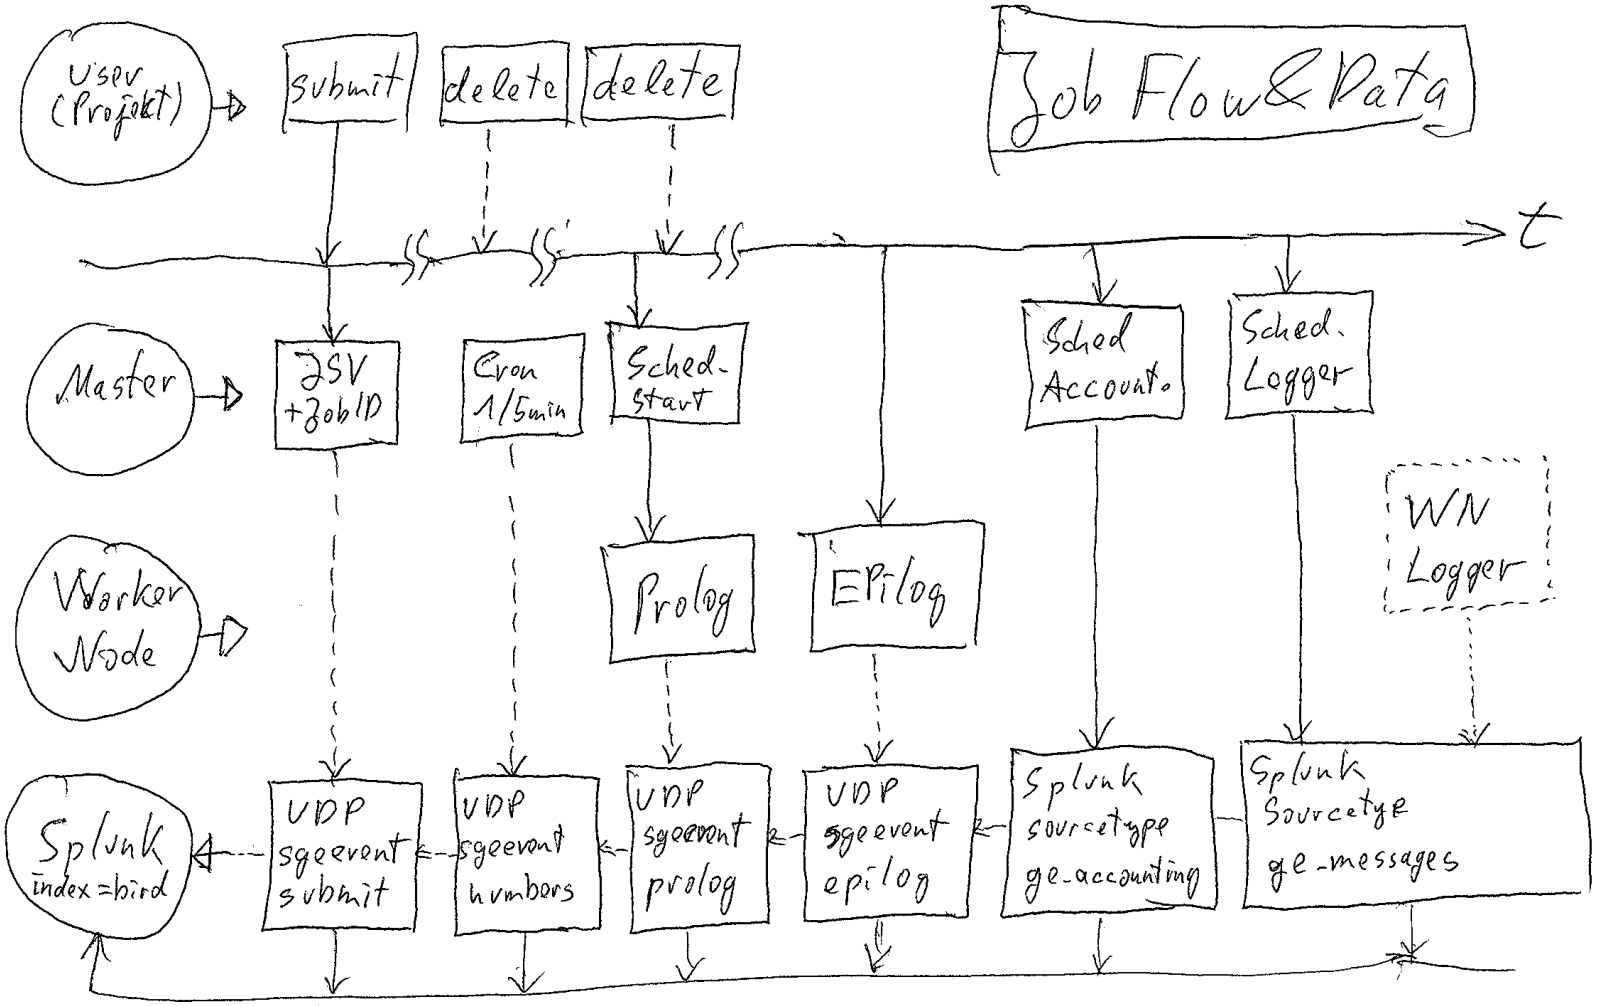

Job Data via UDP Syslog

- Setup
  - Perl Script in JSV(Job Submit Verification on Server), Prolog and Epilog
- Submit
  - eventtype=“sgelog” sgeevent=“submit"
  - sgeuser sgejobid sgeroot sgecell sgehost
- Prolog, Epilog
  - eventtype=sgelog sgeevent={prolog|epilog}
  - sgeuser sgejobid sgeroot sgecell sgehost sgequeue sgeslots sgetaskid sgearch sgehosts sgepe sgesubmithost

Job Inspection (View 1): Find active users, jobs and hosts

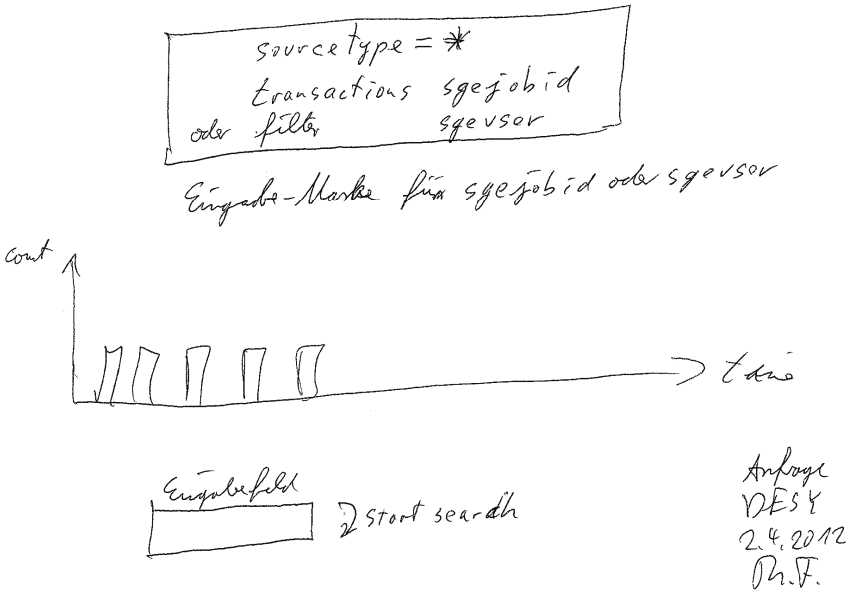

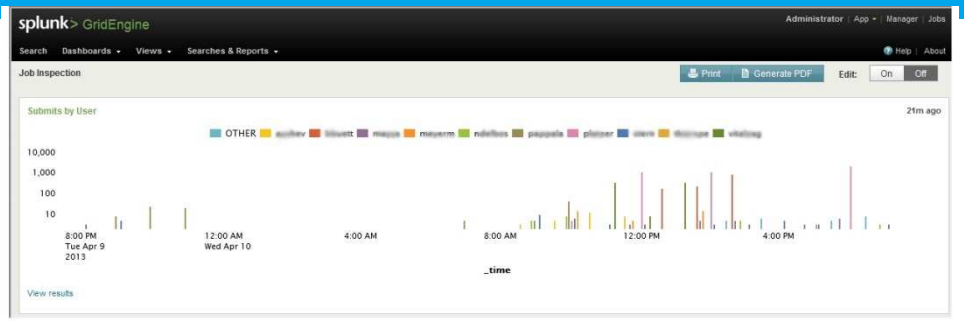

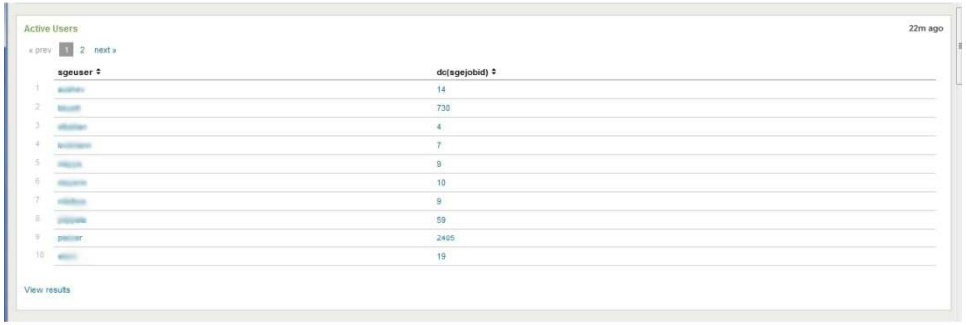

- index=bird | transaction sgejobid startswith=submit span=<time> | search sgejobid=* sgeuser=* | chart values(sgeproject) values(sgeuser) values(sgejobid) values(sgehost) by sgeuser

Job Data via TCP Splunk Forwarder

- Setup
  - Splunk Forwarder (same RPM as on Splunk Master)
  - Reliable Data Collection over Splunk Protocol for accounting and messages files
- Grid Engine Messages
  - Running on qmaster, optionally on workernode
  - sourcetype=ge_messages sgeevent=„{I,W,E}“ source sgejobid sgettaskid sgemessage sgescope
- Grid Engine Accounting
  - Running on qmaster
  - sourcetype=ge_accounting source sgejobid sgedistro sgeuser sgear_submission_time sgearid sgecpu sgedepartment sgeend_time sgeexit_status sgefailed sgegranted_pe sgegroup sgehost sgeio sgeiow sgejobname sgemaxvmem sgemem sgepe_taskid sgepriority sgeproject sgequeue sgeru_idrss sgeru_inblock sgeru_ismrss sgeru_isrss sgeru_ixrss sgeru_majflt sgeru_maxrss sgeru_minflt sgeru_msgrcv sgeru_msgsnd sgeru_nivcsw sgeru_nsignals sgeru_nswap sgeru_nvcsw sgeru_oublock sgeru_stime sgeru_utime sgeru_wallclock sgeslots sgestart_time sgesubmission_time sgetask_number

Reports and Accounting (View 2): Grid Engine Accounting (30d)

- Pie „Sum CPU Seconds by Project“
- Timechart „ Sum CPU Seconds by Project“
- Queue Timing „Wait Times and Wall Times by Queue“
- Timechart „Jobs in Error by Project“
- Table „All Values“

Accounting Analysis

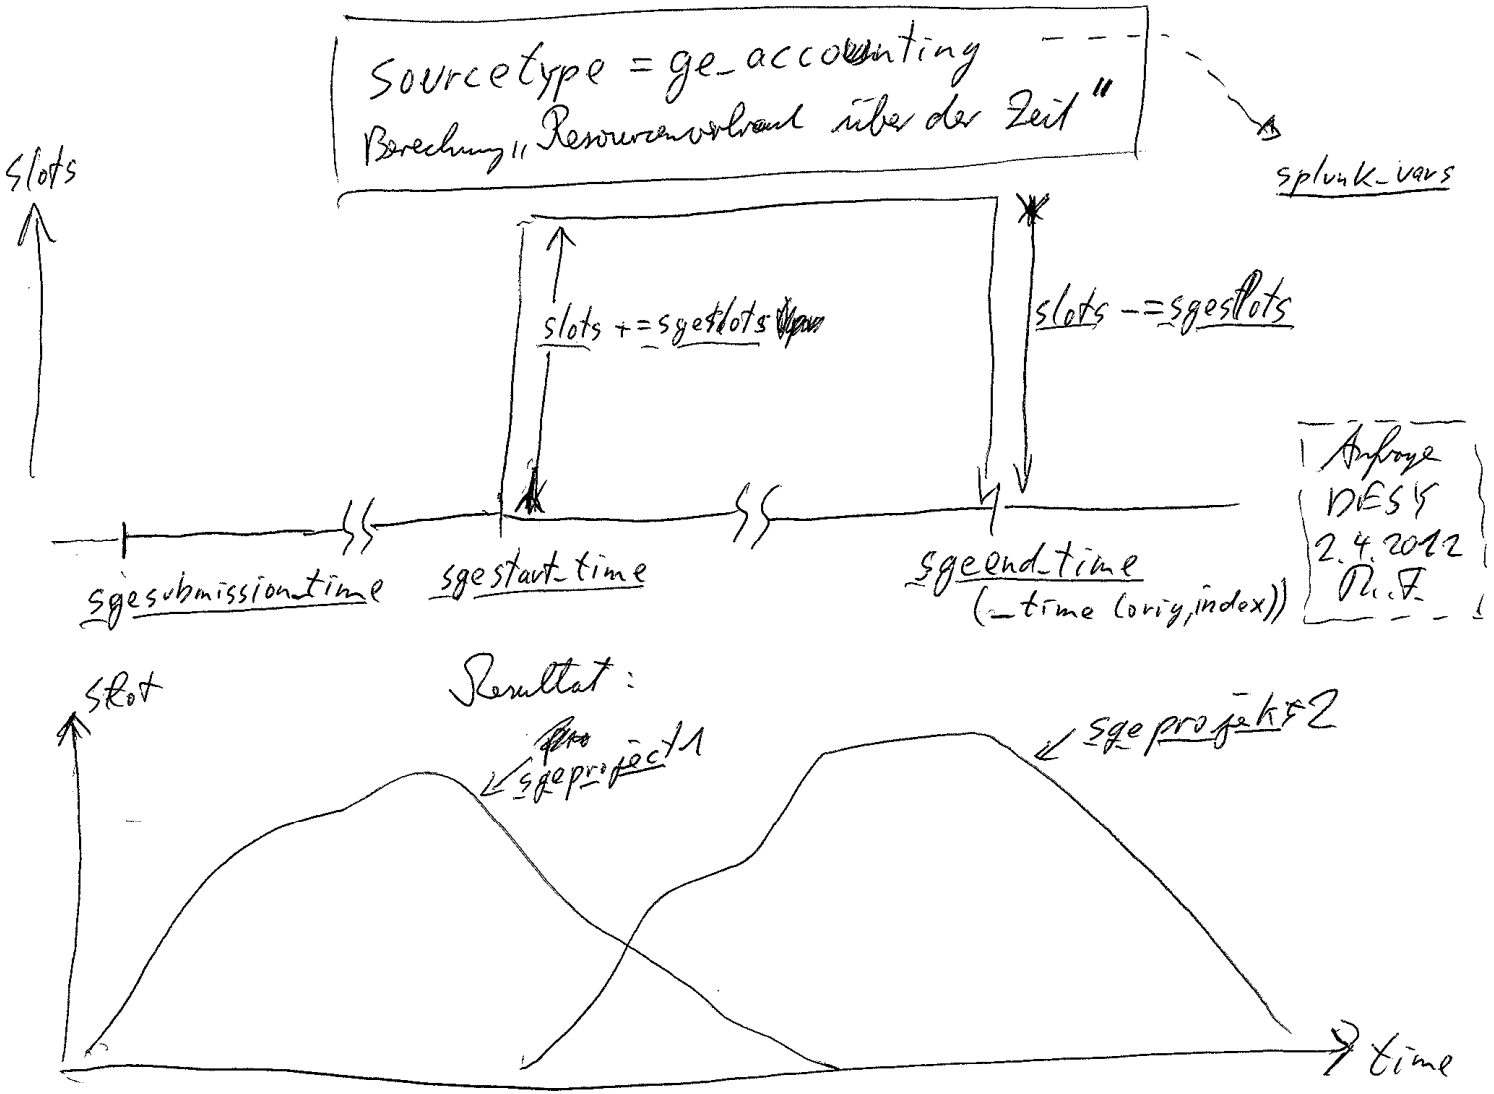

System Data via UDP Syslog

- Setup
  - Perlscript running Commands qhost and qstat in Cron Job on Master and Slave
  - qhost provides Worker Node Resources, qstat shows Project Data
- sgeevent=„numbers“
  - sgelog: sgehost="global" sgeShareTotal sgeSumProject-<ProjectName> sgeSumJobs-Error sgeSumJobs-Waiting sgeSumJobs-SlotRun sgeSumJobs-Running sgeSumShares-<ProjectName> sgeSumShare-Store sgeSumShare-Memory sgeSumShare-Cores sgeSumShare-Hosts sgeSumCores-Total sgeSumCores-Available sgeSumDistro-<Name> sgeSumQueue-l<QName> sgeSumQueue-total sgeSumStore-mem_used sgeSumStore-h_vmem sgeSumStore-h_ftotal sgeSumStore-mem_total sgeSumStore-h_fused
- sgeevent=„projects“
  - sgelog: sgehost="global" sgeproject sgeshare sgestckt sgeovrts SlotInError sgeotckt sgetckts sgeftckt JobsInError SlotRunning sgemem JobsRunning sgeio sgecpu JobsWaiting

Weekly Project View (View 3): Project Infos plus System Infos

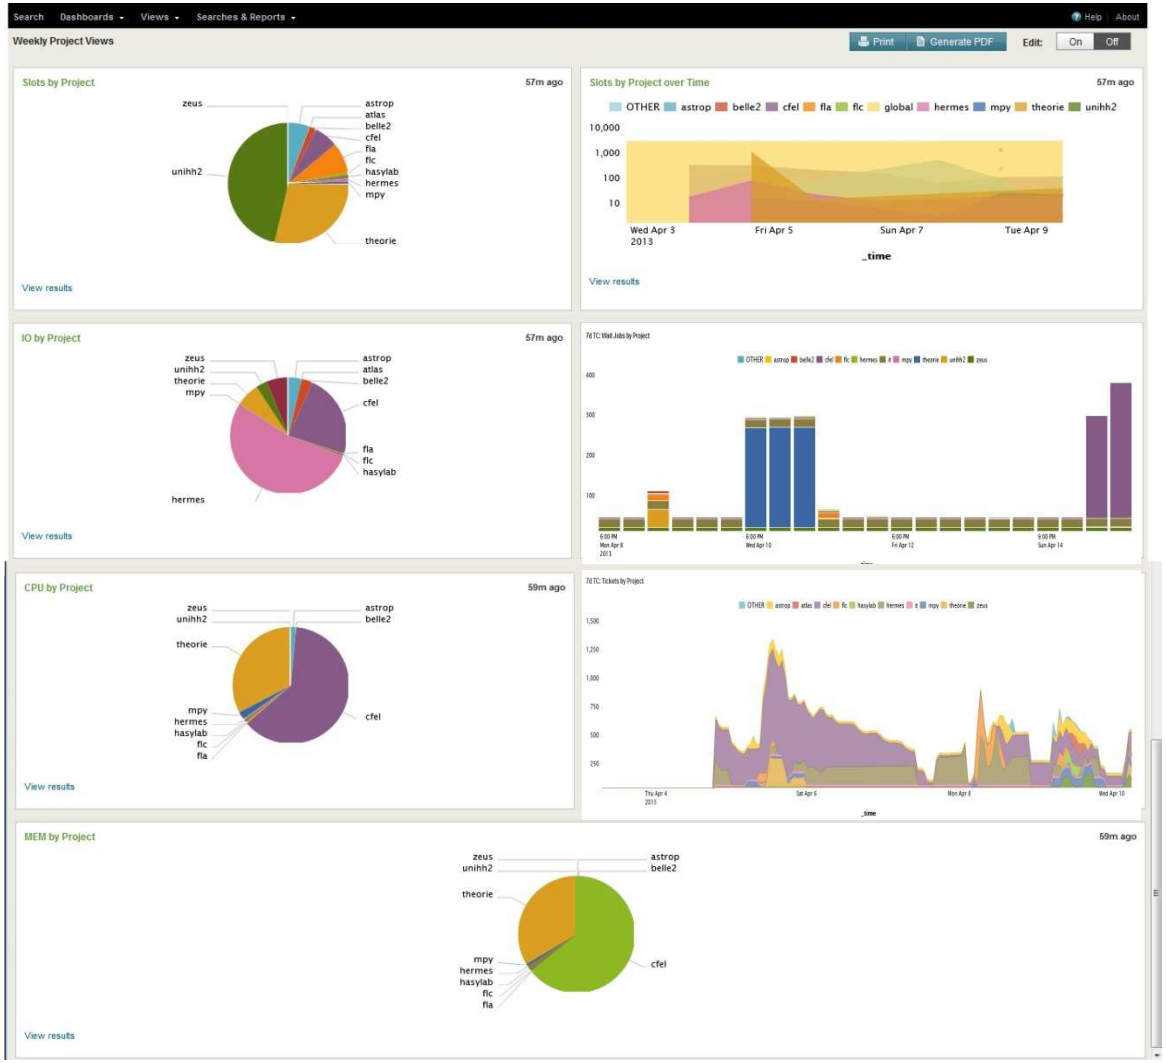

- Pies „Slots, CPU, IO and MEM by Project“
- Timechart „Slots, Waiting Jobs and Tickets by Project“

Realtime System Data (View 4): System Health and Trends

- Last 15 Minutes Jobs, Slots, Waits and Errors
- Timechart 24 Hours Jobs, Slots, Waits and Errors
- Timechart „Core Usage by Queue“
- Sparkling Lines: Trends for Unavailable cores, Distros, Memory Usage
- Table „All Summed Values“

Enterprise vs. Free Version

- Price per Data Volume …
- Limit / License Handling
- More Data (> 500 M)
- Role Based Data Access
- Auto Report
- HA

- Limit on Data Volume …
- Limit / License Handling
- 500 Mbyte
- Interactive Analysis
- Overall Correlations
- Easy Debugging
- Puppet Install

Status and Conclusion

> Status

- Work in Progress
- Finetuning Reports
- Checking Data Consistency
- Still Learning Splunk
- …

> Conclusions

- Would like to buy

  Role Based Data Access
  High Availibilty
  …