HPC USAGE ANALYTICS

Supercomputer Education & Research Centre

Akhila Prabhakaran

OVERVIEW: BATCH COMPUTE SERVERS

Dell Cluster : Batch cluster consists of 3 Nodes of Two Intel Quad Core X5570 Xeon CPUs.

Delta Cluster: 128 Core cluster with 8 Intel(R) Xeon(R) CPU E5-2670 Nodes of 2 processors (16 cores per node) with Intel Omni-path based Infiniband interconnection for MPI communication.

OVERVIEW: BATCH COMPUTE SERVERS

Fermi Cluster: 5 Node heterogeneous cluster dedicated to GPU jobs.

Tyrone Cluster: 800 core cluster composed of 17 Nodes (1 head and 16 execution nodes). This cluster is a heterogenous cluster composed of two types of nodes, 9 nodes with 32-cores each and 8-nodes with 64-cores each.

OVERVIEW: SAHASRAT

- **CPU only cluster:**
  - Intel Haswell 2.5 GHz based CPU cluster with 1376 nodes;
  - Each node has 2 CPU sockets with 12 cores each, 128GB RAM and connected using Cray Aries interconnect.
  - Total CPU compute capacity of 33024 cores and 176128 GB RAM.
- **Accelerator based clusters:** Two accelerator clusters
  - Nvidia GPU cards (44 nodes) [Tesla K40 card with 2880 cores and 12GB device memory]
  - Intel Xeon-Phi cards (24 nodes) Xeon-Phi Processor 7210 has 64 cores with 16GB accelerator memory and 96GB DDR4 memory per node.

OVERVIEW: SAHASRAT

- **High Speed Storage:**
  - 2 PB usable spaces provided by high speed **DDN storage unit** supporting Cray's parallel Lustre filesystem.
- **Software environment:**
  - Cray's customized Linux OS, called Cray Linux Environment,
  - Supports Intel and open-source based Gnu compilers
  - Parallel libraries like OpenMP, MPI, CUDA and Intel Cluster software
  - Extensive range of parallel Scientific and Mathematical libraries like BLAS, LAPACK, Scalapack, fftw, hdf5 etc.
  - DDT parallel debugger and profiler

OVERVIEW: WORKLOAD MANAGER

Portable Batch System (**PBS**) is a computer software that performs job scheduling.

Its primary task is to allocate computational tasks, i.e., batch jobs, among the available computing resources.

- **OpenPBS** — original open source version released in 1998 (not actively developed)
- **TORQUE** — a fork of OpenPBS.
- **PBS Professional** (PBS Pro) — the version of PBS offered by Altair Engineering that is dual licensed under an open source and a commercial license.

NEED FOR ANALYTICS

- Monitoring the availability of resources.
- Monitoring real-time usage of resources.
- Understanding the nature of HPC workloads.
- How are the allocated resources being used?
- Research domains being serviced by in-house HPC resources.
- Insight into user domain, application domain and usage patterns.
- Understand Queue Waiting Times & Job Efficiency.
- It is necessary that the data be presented in intuitive and visually understandable way so that quick inferences can be made.

THIRD PARTY ANALYTICS TOOLS

Ganglia: http://10.16.4.87/ganglia/

Real-time monitoring and execution environment and is used by hundreds of universities, private and government laboratories and commercial cluster implementers.

Ganglia in SERC:

Installed on the basic HPCs, i.e., on tesla, fermi, dell and tyrone clusters. Ganglia ties up these systems as a single pool.

Data regarding availability of nodes and cores and their average use is shown.

Who can access Ganglia?

SERC ANALYTICS PLATFORM FOR BATCH COMPUTE SERVERS

URL : www.serc.iisc.in/~hpc/WebPages/OTHERCLUSTERS/ClusterTyrone.html

Access:

Limited access to SERC administrators

Behind the scenes:

- PBS logs, NIS user information, Faculty and Department Information.
- Automated monthly scripts.
- Scripts in R to consolidate & generate HTML.
- Static HTML pages. (Linux, Apache)

SERC ANALYTICS PLATFORM FOR BATCH COMPUTE SERVERS

Basic HPC Clusters

Get information from PBS Monthly logs

Identify and associate jobs with research groups and departments.

Monthly utilization charts for each cluster.

Insights

Utilization of these clusters.

Department-wise & Research group wise utilization.

Potential inputs for capacity planning and new procurements.

SAHASRAT - QUEUES

Queue	Min & Max Cores	Wall Time Limit	Priority
idqueue	24-256	2 Hours	100
small	24-1032	24 Hours	300
small72	24-1032	24 Hours	300
medium	1033-8208	24 Hours	350
large	8209-22800	24 Hours	400
gpu queues	Core counts ranging from 1-12 and one GPU on a node	24 Hours	100

```
#!/bin/sh
#This job should be redirected to small
queue
#PBS -N jobname
#PBS -l select=100:ncpus=24
#PBS -l walltime=24:00:00
#PBS -l place=scatter
#PBS -l accelerator_type="None"
#PBS-S /bin/sh@sdb -V .
/opt/modules/default/init/sh cd <exepath>
aprun -j 1 -n 2400 -N 24 ./a.out
```

SAHASRAT ANALYTICS PLATFORM DATA SOURCES

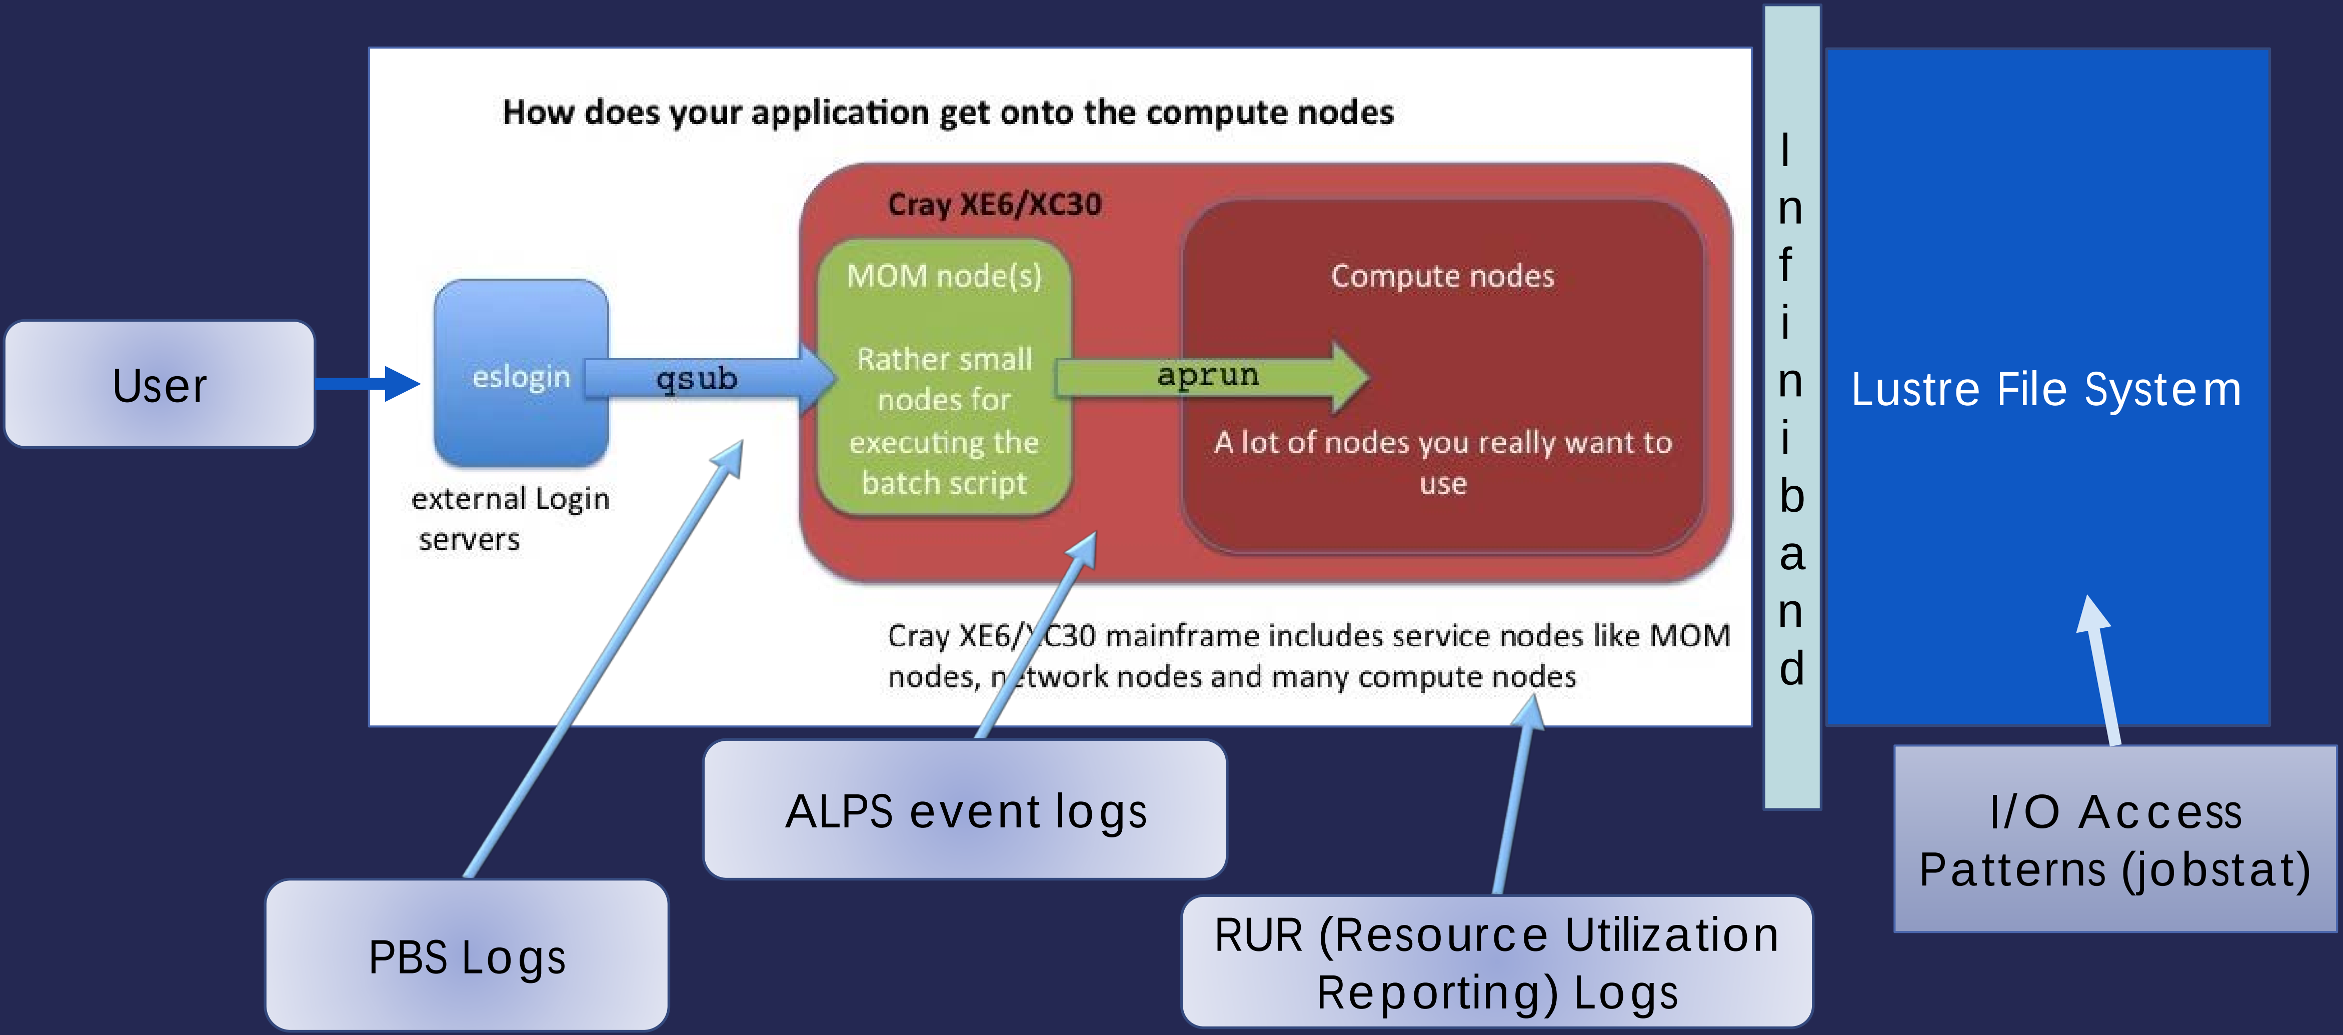

How does your application get onto the compute nodes

Cray XE6/XC30

MOM node(s)

Rather small nodes for executing the batch script

Compute nodes

A lot of nodes you really want to use

eslogin

external Login servers

qsub

aprun

Cray XE6/XC30 mainframe includes service nodes like MOM nodes, network nodes and many compute nodes

User

Infiniband

Lustre File System

PBS Logs

ALPS event logs

RUR (Resource Utilization Reporting) Logs

I/O Access Patterns (jobstat)

SERC ANALYTICS PLATFORM: DATA COLLECTION

PBSPro Accounting Logs – Job specific allocations and Queue Wait Times

ALPS event logs – Resource allocation in terms on nodes list, apstart, append times, cores used by each aprun statement.

RUR (Resource Utilization Reporting) Logs

- RUR is a tool for gathering statistics on how system resources are being used by applications.
- Resource utilization statistics are gathered from compute nodes.
- RUR runs primarily before the job has started and after it ends, ensuring minimal impact on performance.
- Details for each JOBID, APID (aprun) is specified in these logs
- Job Start Time, Job End Time, details for each aprun, apstart, append, utime, stime, Nodes allocated, gpu utilization, all aprun statements for the job etc.

SAMPLE LOG

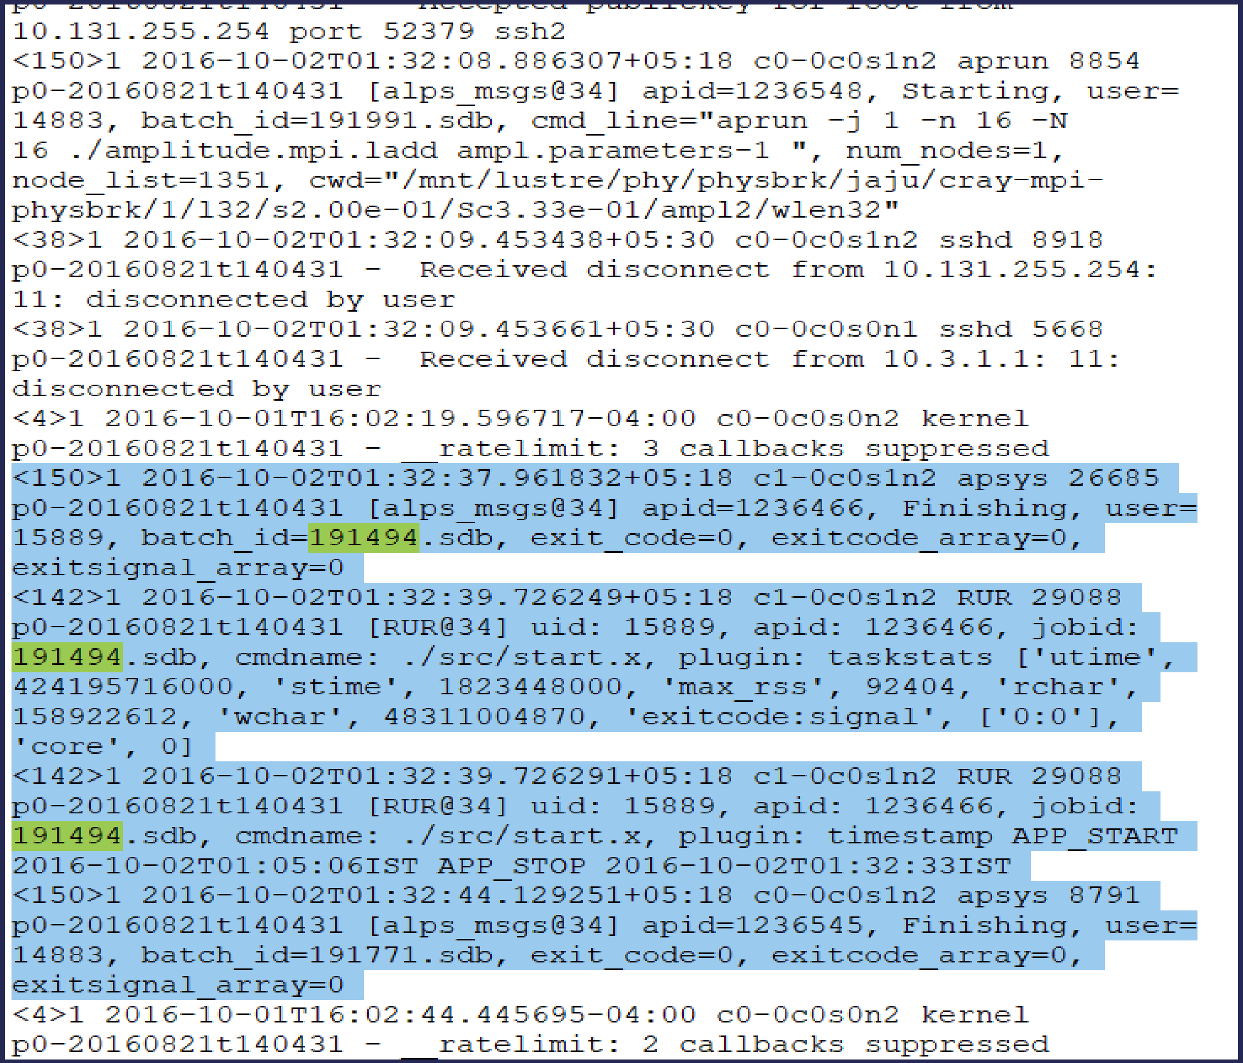

```
p0-20160821t140431 - Accepted publickey for root from
10.131.255.254 port 52379 ssh2
<150>1 2016-10-02T01:32:08.886307+05:18 c0-0c0s1n2 aprun 8854
p0-20160821t140431 [alps_msgs@34] apid=1236548, Starting, user=
14883, batch_id=191991.sdb, cmd_line="aprun -j 1 -n 16 -N
16 ./amplitude.mpi.ladd ampl.parameters-1 ", num_nodes=1,
node_list=1351, cwd="/mnt/lustre/phy/physbrk/jaju/cray-mpi-
physbrk/1/132/s2.00e-01/Sc3.33e-01/ampl2/wlen32"
<38>1 2016-10-02T01:32:09.453438+05:30 c0-0c0s1n2 sshd 8918
p0-20160821t140431 -  Received disconnect from 10.131.255.254:
11: disconnected by user
<38>1 2016-10-02T01:32:09.453661+05:30 c0-0c0s0n1 sshd 5668
p0-20160821t140431 -  Received disconnect from 10.3.1.1: 11:
disconnected by user
<4>1 2016-10-01T16:02:19.596717-04:00 c0-0c0s0n2 kernel
p0-20160821t140431 - __ratelimit: 3 callbacks suppressed
<150>1 2016-10-02T01:32:37.961832+05:18 c1-0c0s1n2 apsys 26685
p0-20160821t140431 [alps_msgs@34] apid=1236466, Finishing, user=
15889, batch_id=191494.sdb, exit_code=0, exitcode_array=0,
exitsignal_array=0
<142>1 2016-10-02T01:32:39.726249+05:18 c1-0c0s1n2 RUR 29088
p0-20160821t140431 [RUR@34] uid: 15889, apid: 1236466, jobid:
191494.sdb, cmdname: ./src/start.x, plugin: taskstats ['utime',
424195716000, 'stime', 1823448000, 'max_rss', 92404, 'rchar',
158922612, 'wchar', 48311004870, 'exitcode:signal', ['0:0'],
'core', 0]
<142>1 2016-10-02T01:32:39.726291+05:18 c1-0c0s1n2 RUR 29088
p0-20160821t140431 [RUR@34] uid: 15889, apid: 1236466, jobid:
191494.sdb, cmdname: ./src/start.x, plugin: timestamp APP_START
2016-10-02T01:05:06IST APP_STOP 2016-10-02T01:32:33IST
<150>1 2016-10-02T01:32:44.129251+05:18 c0-0c0s1n2 apsys 8791
p0-20160821t140431 [alps_msgs@34] apid=1236545, Finishing, user=
14883, batch_id=191771.sdb, exit_code=0, exitcode_array=0,
exitsignal_array=0
<4>1 2016-10-01T16:02:44.445695-04:00 c0-0c0s0n2 kernel
p0-20160821t140431 - __ratelimit: 2 callbacks suppressed
```

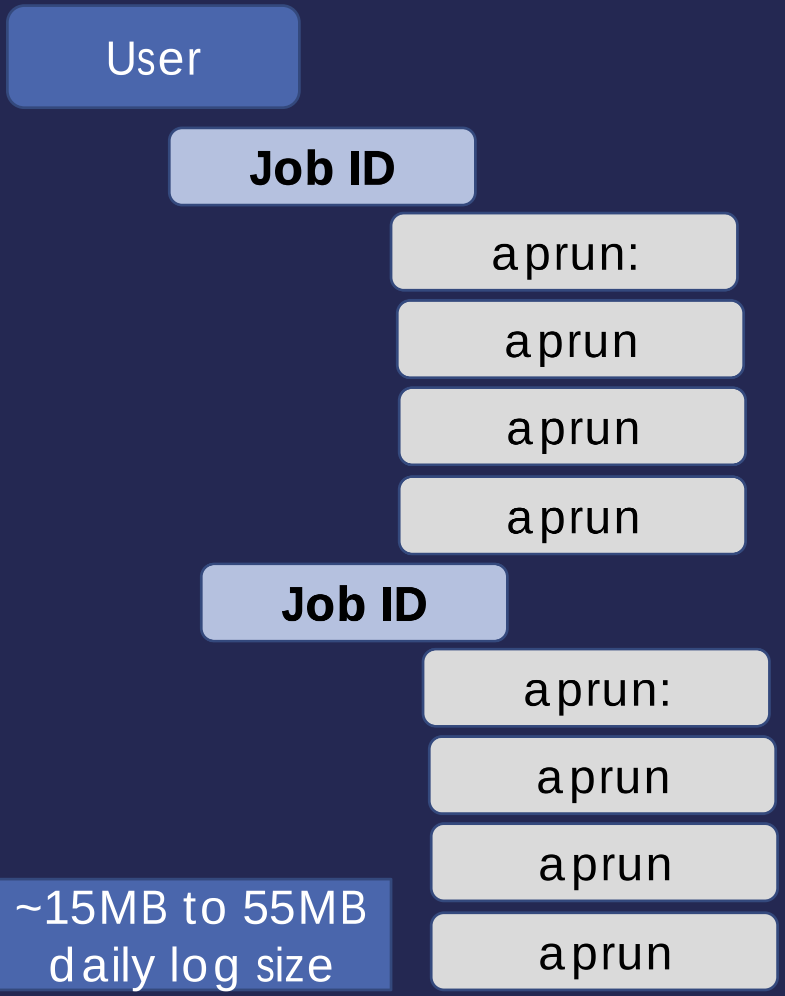

User

Job ID

aprun:

aprun

aprun

aprun

Job ID

aprun:

aprun

aprun

aprun

~15MB to 55MB daily log size

HPC USAGE ANALYTICS PLATFORM: BUILDING BLOCKS

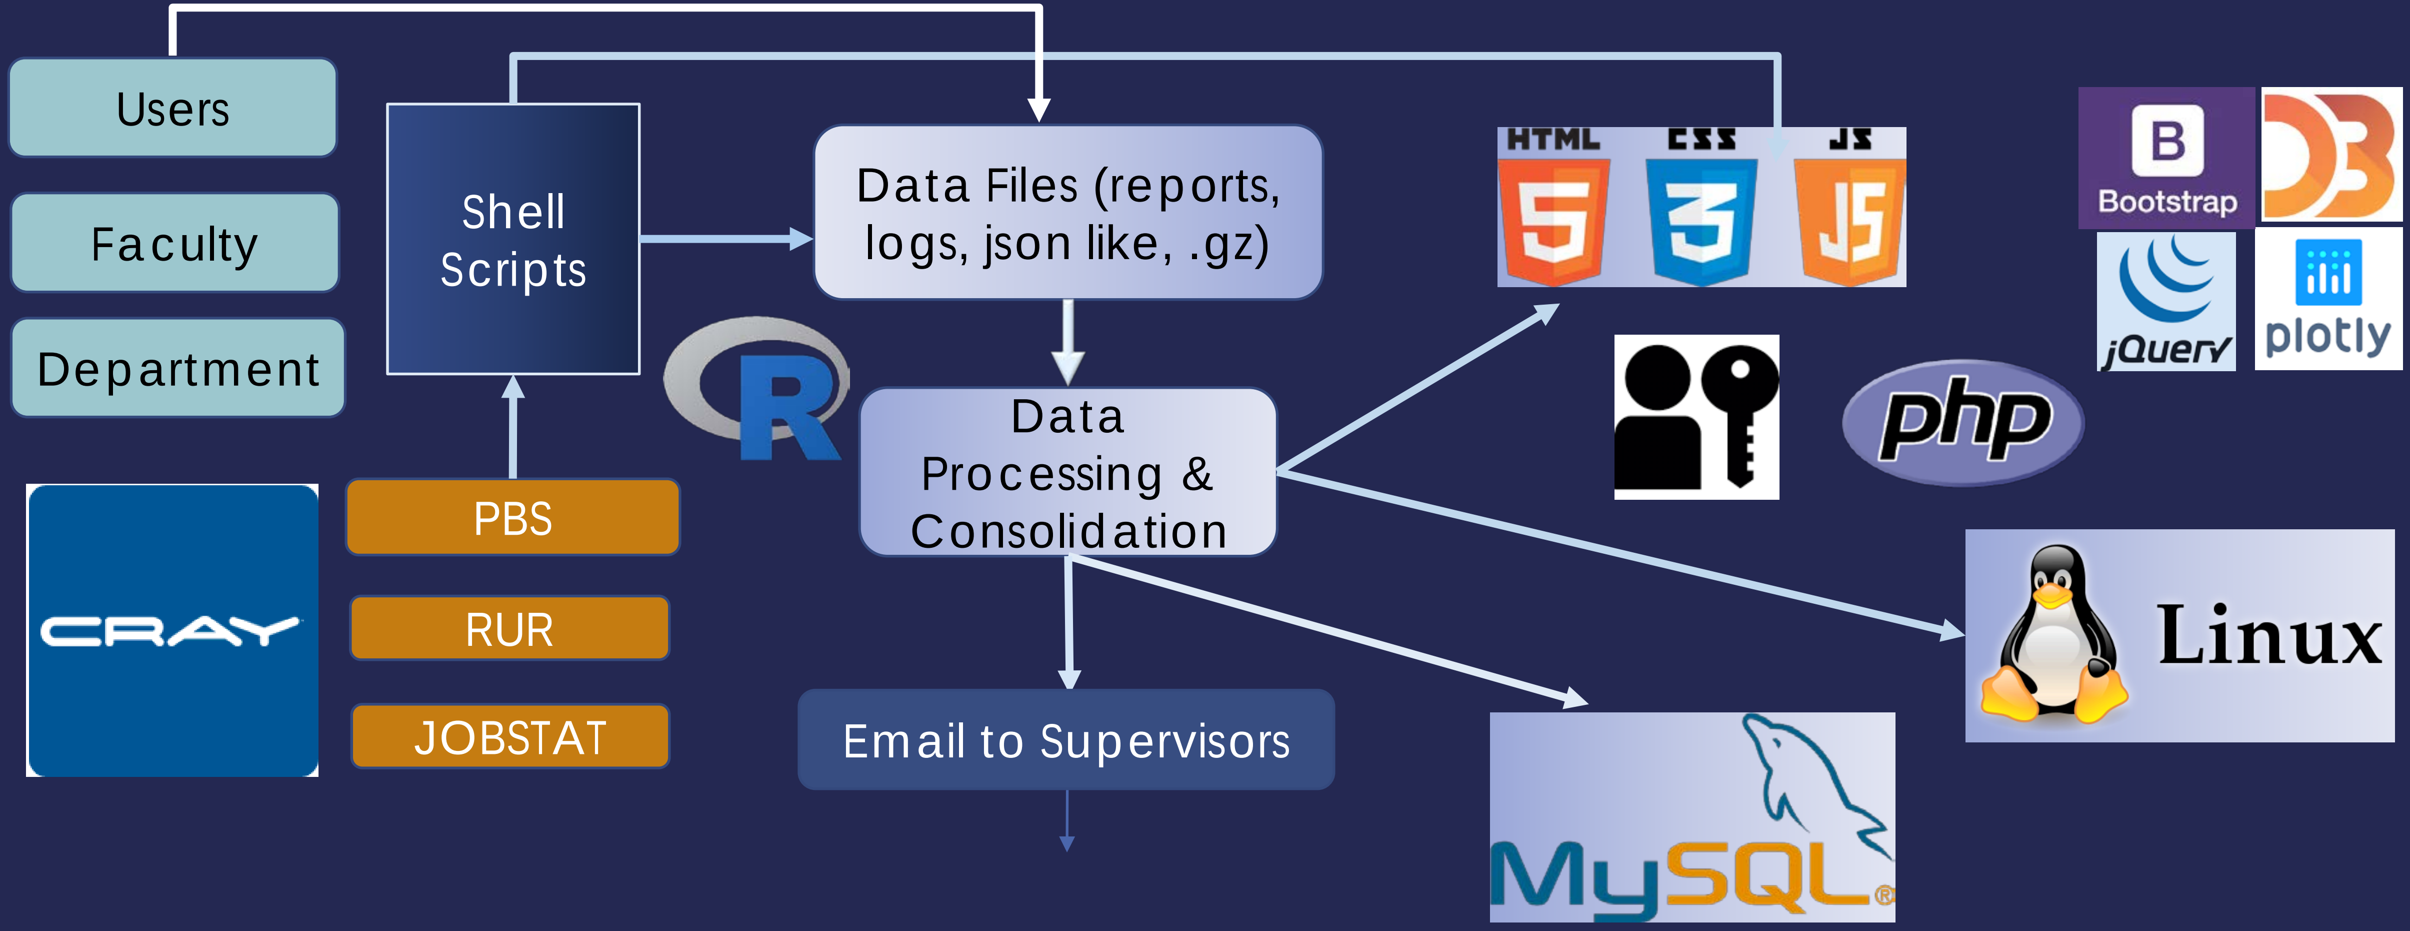

Users

Faculty

Department

Shell Scripts

Data Files (reports, logs, json like, .gz)

Data Processing & Consolidation

Email to Supervisors

PBS

RUR

JOBSTAT

SERC ANALYTICS PLATFORM

URL : www.serc.iisc.in/~hpc/WebPages/UserDashboard/RunningJobs.html

Access : By User Role (Admins, Supervisors, Users)
Authentication: Computational ID

Realtime Insights:

- ✓Realtime SahasraT Availability & Usage
- ✓Realtime User Jobs on Queue with Estimated Queue WaitTime
- ✓Realtime Running Jobs Status & Distribution of Utilization
- ✓Realtime view of IOSTATS (aggregated across all OSTs)

ESTIMATED JOB START TIME

✓Estimates the start time for all jobs on Queue using a PBS Server Attribute.

✓PBS allows setting of the **est_start_time_freq** for retrieving estimated start times for all jobs.

✓By setting the **est_start_time_freq = 0**, we have been able to retrieve the estimated start time for each job on Queue.

✓These estimates are retrieved using the **qstat -T** command.

✓Using a cron-job that runs every 10 minutes, we get this information for all Queued Jobs.

✓The estimated start time for each job is computed by PBS on every scheduler cycle and is likely to change based on backfilling.

JOB CPU EFFICIENCY

✓CPU Efficiency = $\dfrac{\text{Utime + Stime}}{\text{Cores X Wall Time}}$

✓Computed for each aprun statement

✓Averaged for a Job

✓Number of cores equivalent to hyperthreaded cores for jobs that have implicitly or explicitly used -j2 on the aprun command.

```
2013-11-02T11:09:49.457770-05:00 c0-0c1s1n2 RUR 2417 p0-20131101t153028 [RUR@34] uid:
12345, apid: 86989, jobid: 23460, cmdname: /lus/tmp/rur01.2338/./CPU01-2338 plugin:
taskstats ['utime', 10000000, 'stime', 0, 'max_rss', 940, 'rchar', 107480, 'wchar', 90,
'exitcode:signal', ['0:0'], 'core', 0]
```

JOB I/O STATISTICS: LUSTRE FILE SYSTEM

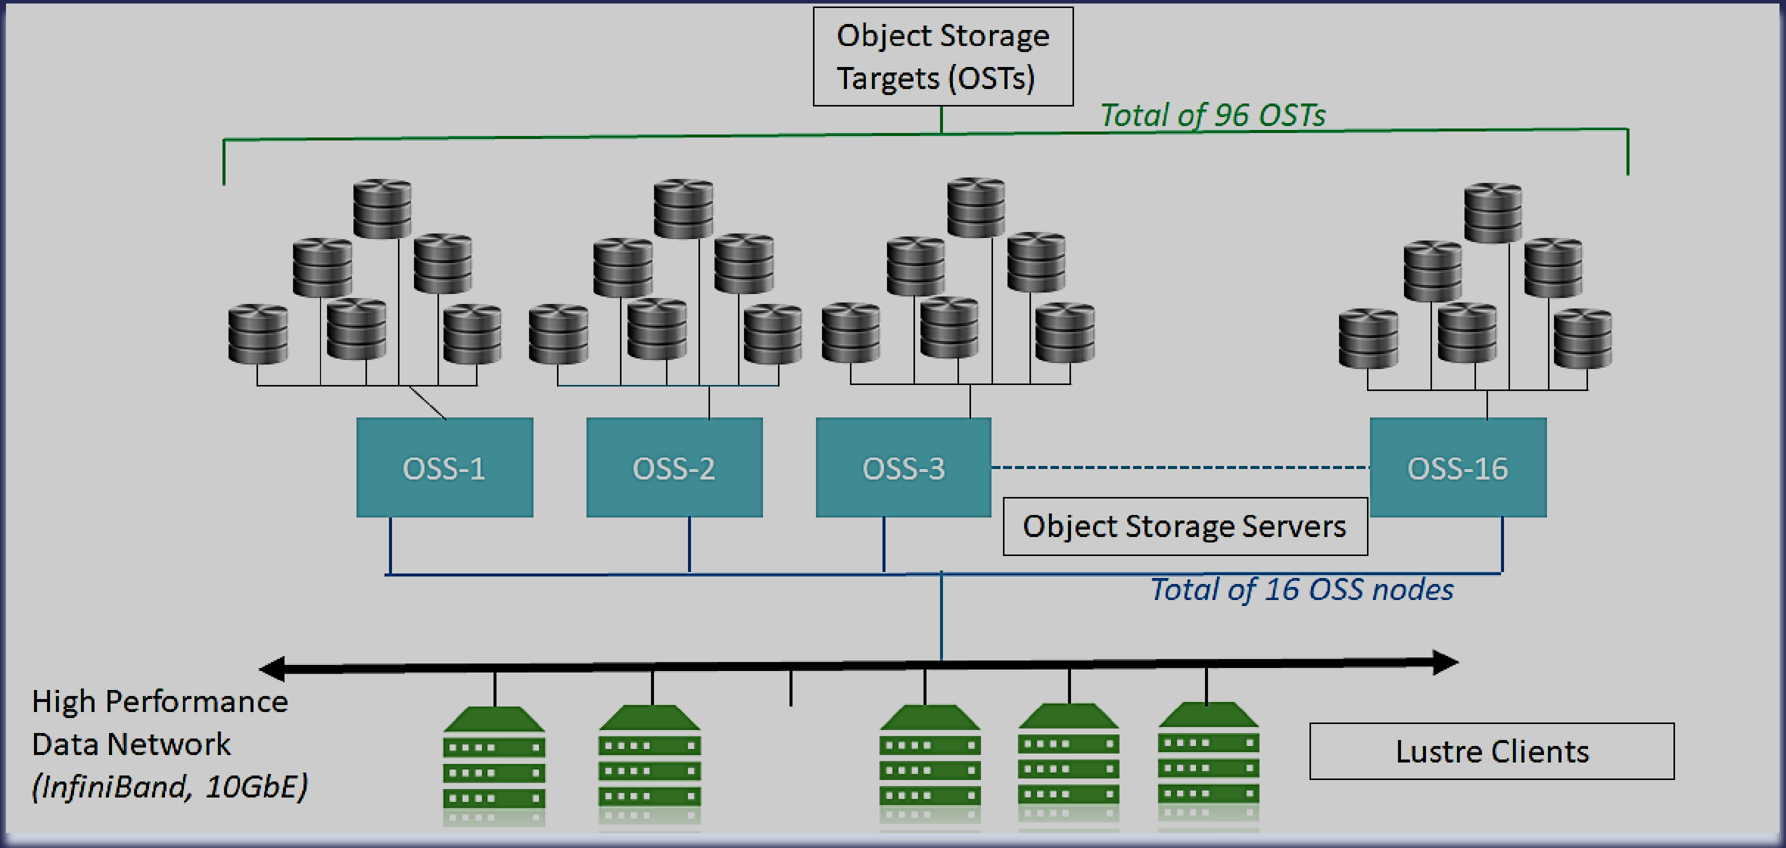

JOB I/O STATISTICS

- ✓A script runs every 8 minutes
- ✓SSH to the 16 OSS nodes of Cray
- ✓Gather the ***job_stats*** data for the 96 OSTs
- ✓Each OSS will provide stats for the 6 OSTs its responsible for
- ✓The script takes 5-6 seconds to complete.
  At the end of each job run, we get a single output file.

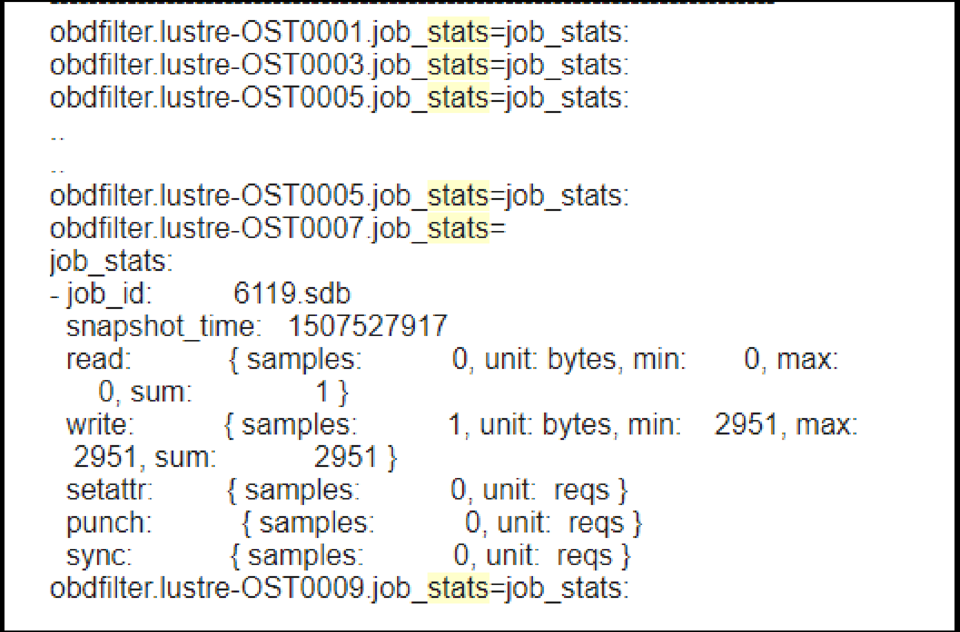

```
obdfilter.lustre-OST0001.job_stats=job_stats:
obdfilter.lustre-OST0003.job_stats=job_stats:
obdfilter.lustre-OST0005.job_stats=job_stats:
..
..
obdfilter.lustre-OST0005.job_stats=job_stats:
obdfilter.lustre-OST0007.job_stats=
job_stats:
- job_id:          6119.sdb
  snapshot_time:   1507527917
  read:            { samples:          0, unit: bytes, min:       0, max:
     0, sum:               1 }
  write:           { samples:          1, unit: bytes, min:    2951, max:
   2951, sum:           2951 }
  setattr:         { samples:          0, unit:  reqs }
  punch:             { samples:          0, unit:  reqs }
  sync:            { samples:          0, unit:  reqs }
obdfilter.lustre-OST0009.job_stats=job_stats:
```

Each log is 400 – 800KB,
10800 logs per day
~ few GB per day

SERC ANALYTICS PLATFORM

URL : www.serc.iisc.in/~hpc/WebPages/CRAY/StaticCoreUsageReport.html

Access : Restricted

Insights:

Overall System Utilization

Workloads by Queue and Core counts

Timeline of availability and utilization of resources

Queue Wait Times & Insight into jobs on special/reserved queues.

Job Completion Statistics

SERC ANALYTICS PLATFORM

Workload Insights:

- ✓Overall Availability & Utilization of Resources.
- ✓Monthly Utilization – by Queue, Department, Research Group & User
- ✓Job size characterization based on historical usage patterns.
- ✓Job success & exit code patterns based on historical usage.
- ✓Resource Allocated versus Resource Used
- ✓Job wise: Efficiency, Queue Wait Time and IO Stats
- ✓Monthly Summary of Usage emailed to research groups.
- ✓System Outage Calendar

SERC USER'S FORUM

URL : https://usersupport.serc.iisc.in/

Accessible using computational ID & password

Forum is broadly organized into 4 categories

- FAQ
- Systems Support
- Software Support
- HPC Discussion Forum

THANK YOU