

Large-Scale Monitoring

Making data-driven decisions possible to address current and future goals

Kathleen Shoga, M. LeGendre, A. Scott, I. Lee (LLNL)

We need to understand the whole picture of performance across computing architectures and different workloads in LC to make informed decisions throughout the computing center. The increase in usage and integration of AI and quantum with HPC makes understanding even more important to optimize for the best combination of resources to complete mission goals quickly and efficiently. Funding: ~1–2 FTEs per year.

Goals

- Understand the effects resources have on each other.
- Utilize this information to create models and feedback loops.

Models/Metrics



Job Scheduler

Compute Resources

Example: Schedule jobs based on availability and resources available; e.g. AI or HPC specialized clusters, potential quantum resources. Provide feedback from resources to better schedule jobs.

Models/Metrics



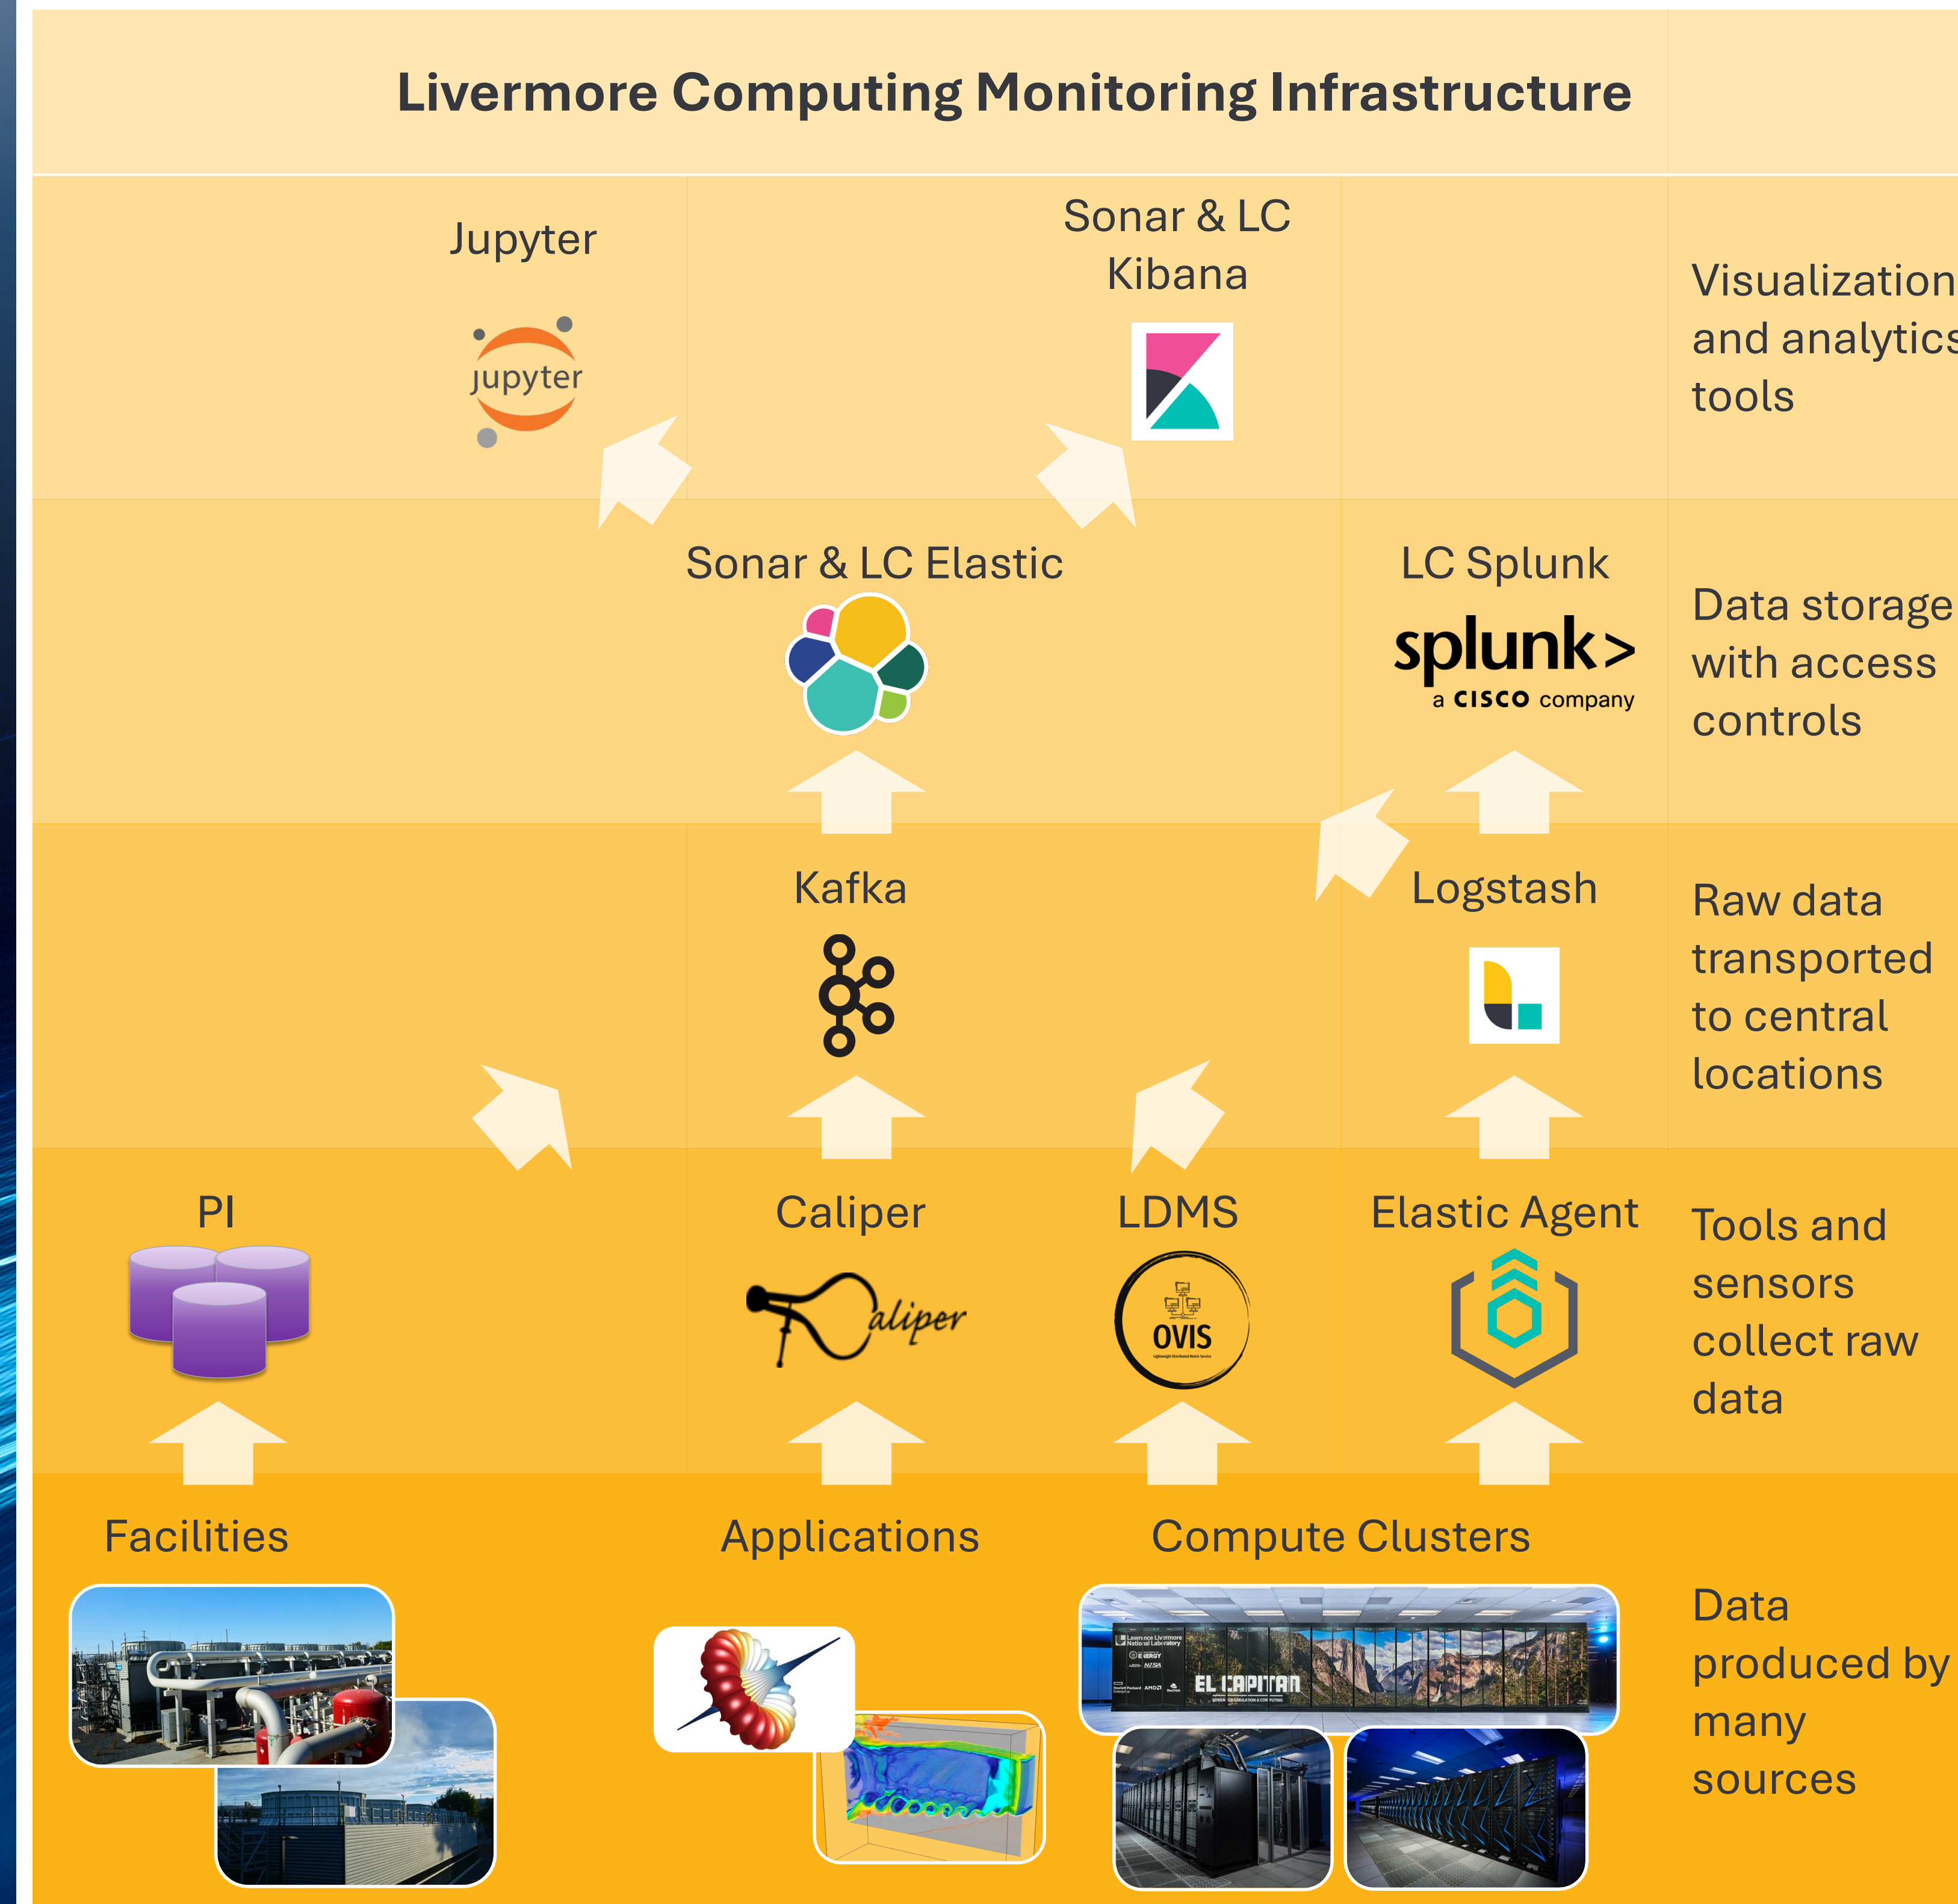
Facilities

Compute Clusters


Example: Perform appropriate response in facilities for cooling based on the workloads being done on compute clusters. Adjust workloads when downtimes are needed in facilities.

Combined Monitoring Infrastructure


- LC collects performance, metric, and log data from many sources, including compute clusters, facilities, and applications.
- Collected data is centralized with access controls for system administrators, researchers, application users, managers, and for system procurement decisions.



El Capitan Dashboards



Kibana visualization of system logs per hour (including El Capitan and its filesystem).



Kibana visualization of average GPU utilization across El Capitan.

Conclusions

- Large-scale monitoring is important to understand the performance and usage of hardware, workloads, and facilities.
- LLNL's combined expertise in monitoring metrics, logs, application performance, and facilities enable better informed procurement decisions, performance and other research, optimization of current systems, etc.

Collaborators



Center-wide monitoring enables data-driven procurements, research, security, and operations