LEVERAGING SPLUNK FOR CONTROL SYSTEM MONITORING AND MANAGEMENT

16th International Conference on Accelerator & Large Experimental Physics Control Systems (ICALEPCS)

October 8-13, 2017

Mikhail Fedorov
National Ignition Facility (NIF) Integrated Computer Control System (ICCS)

LLNL-PRES-739417
This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under contract DE-AC52-07NA27344. Lawrence Livermore National Security, LLC

Lawrence Livermore National Laboratory

Splunk at the National Ignition Facility (NIF)

- **Splunk Capabilities**
  - Indexing of unstructured log files
  - Powerful data search and analysis
  - Web visualizations and dashboards
- **Deployment at NIF**
  - Started in 2013
  - Well received by the controls team
  - Extended to one year of all logs
  - 20-50 GB/day of controls logs
  - Connected to other databases

Splunk became the primary tool for system monitoring, off-normals response and predictive trend analysis. Connected to enterprise data systems, Splunk is also used for data-driven project management.

Monitoring system “vitals”: CPU/RAM/Swap utilization

- **Resource utilization**
  - **2600 processes**
  - **Linux/Solaris/Windows**
  - **Load Patterns**
- **Process segmentation**
  - **Custom multi-panel dashboard**
  - **Processes are grouped by function**
  - **Database connection to Oracle Enterprise Manager (OEM)**

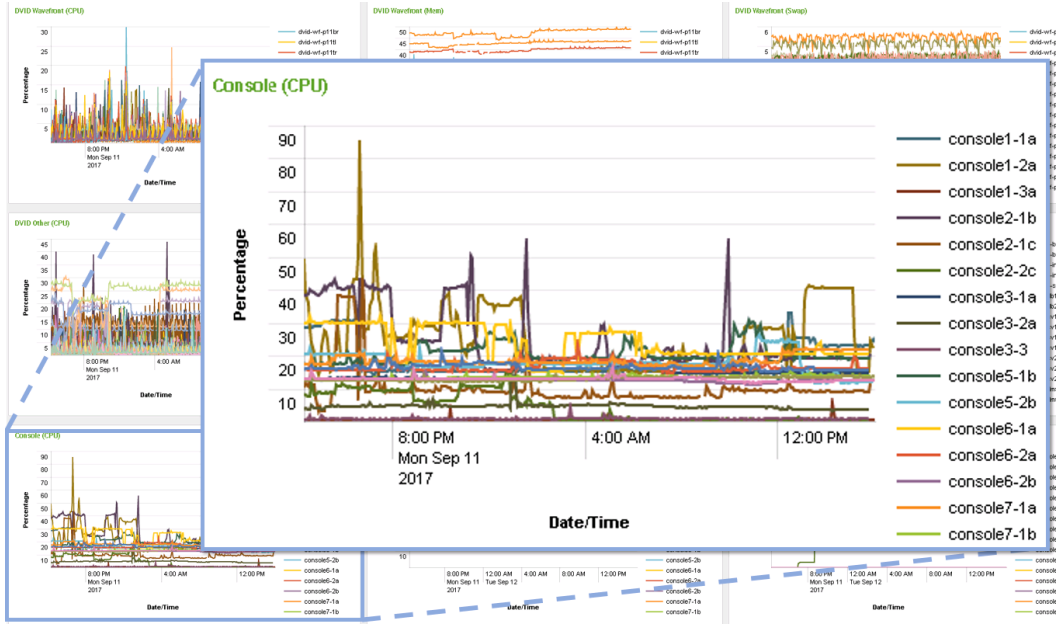

Custom dashboard displays similar processes together:
normal utilization patterns emerge, and outliers become clearly visible

Long-term trend monitoring

- **Monitoring NIF capacitor "health"**
  - 192 capacitor bank modules
  - New "health" concern
  - Long "incubation" time
- **One-year trend dashboard**
  - Schema-less data extraction
  - Looking back at one year of logs
  - Compact visualization of long trends for 192 capacitor banks

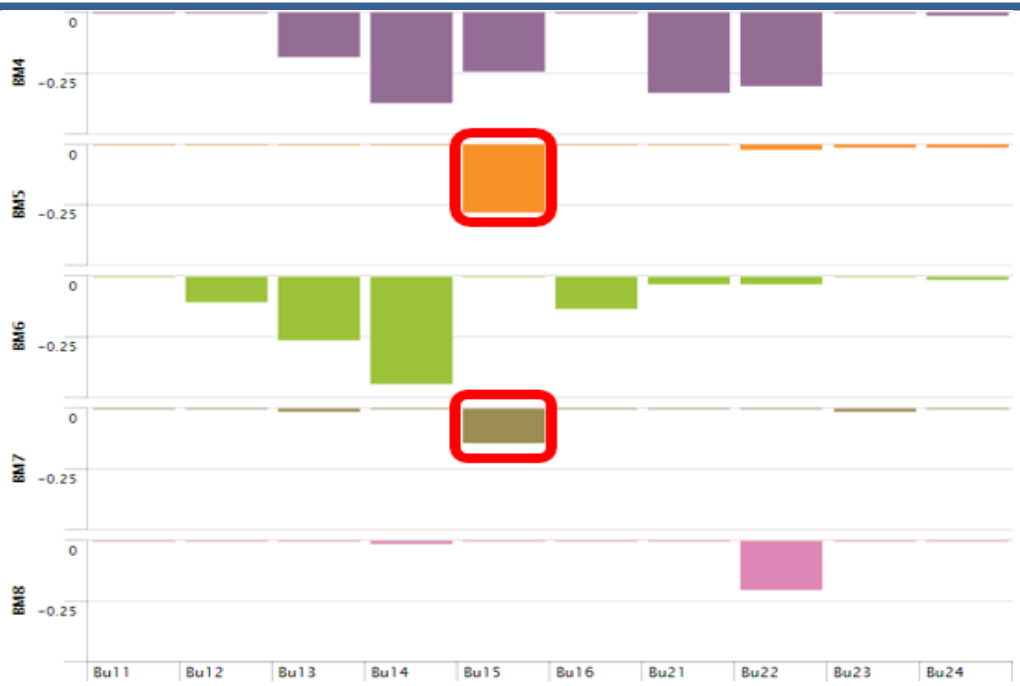

Schema-less data extraction turned one year of unstructured logs into a goldmine of long trend data, predicting need for capacitor bank "health" inspection

Gantt chart visualization of NIF shot cycle

- **Basic, frequently asked questions**
  - When was the laser shot fired?
  - How long was the preparation?
  - What was happening at a given time?

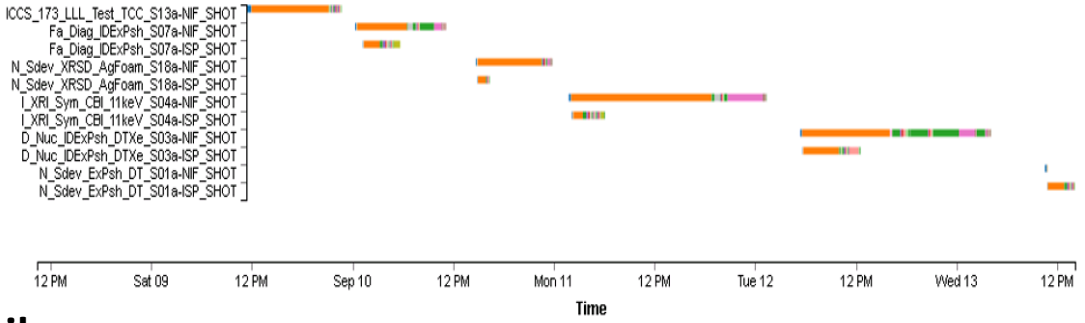

- **Navigating through layers of details**
  - Tens of thousands state transitions
  - Time scale from 1 ms to 1 week
  - Hierarchy of state transitions

Gantt chart visualization with custom drill-downs allows efficient navigation across time scales and numerous subsystems

Gantt chart zoom-in to “macrostep” level

- **Splunk drill-down feature**
  - Open new chart
  - Navigate to an external link
  - Trigger interactive change
- **NIF “macrostep” level of details**
  - Important transitions
  - Multiple NIF subsystems
  - Durations and parallelization

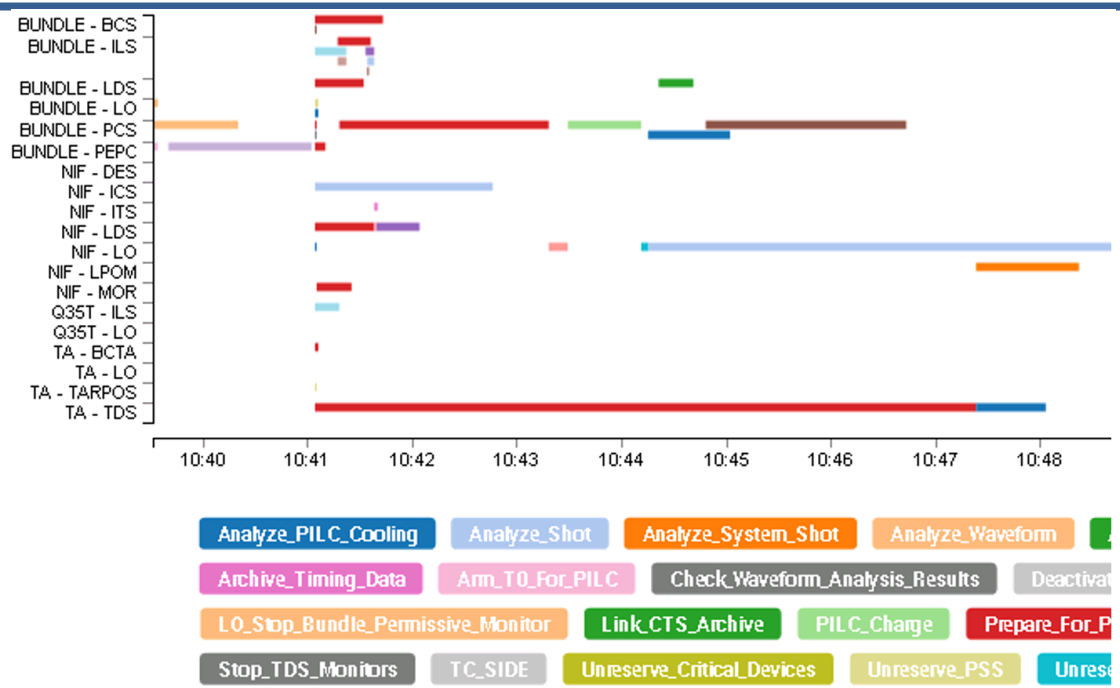

Navigating from top “shot” level down to “macrostep” level provides an insight into the duration of significant operations and attained parallelization of processes

Gantt chart zoom-in to “step” level

- NIF “step” level of details
  - One subsystem
  - Low-level operation
  - Multiple locations
- Drill-down to log entries
  - Full text of log messages
  - Filtered by the subsystem
  - Filtered by the location
  - Filtered by the “step” time interval

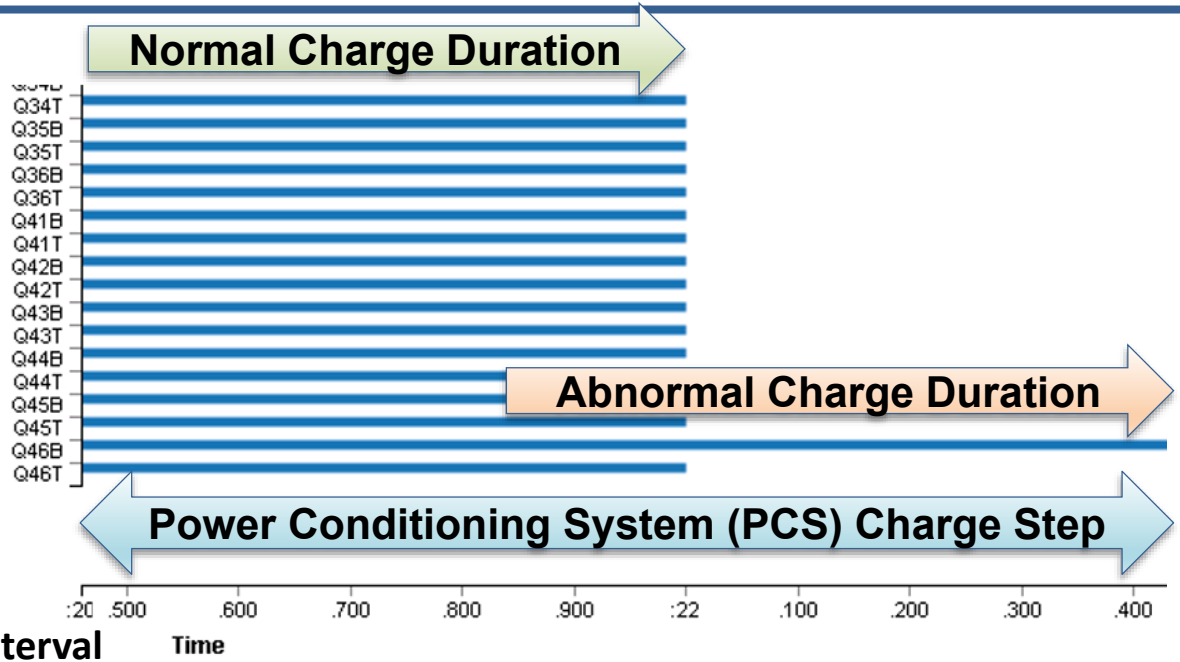

Diving into the third level of individual shot “steps” allows to compare durations across multiple system locations. Further drill-down leads to a filtered view of text log messages.

Finding laggards

- **NIF hardware symmetry**
  - 192 laser beams
  - 48 quads (x4) or 24 bundles (x8)
  - 4 clusters
  - 2 laser bays
- **Expected similarity in state patterns**
  - Majority establishes the reference
  - Outliers usually point to a problem

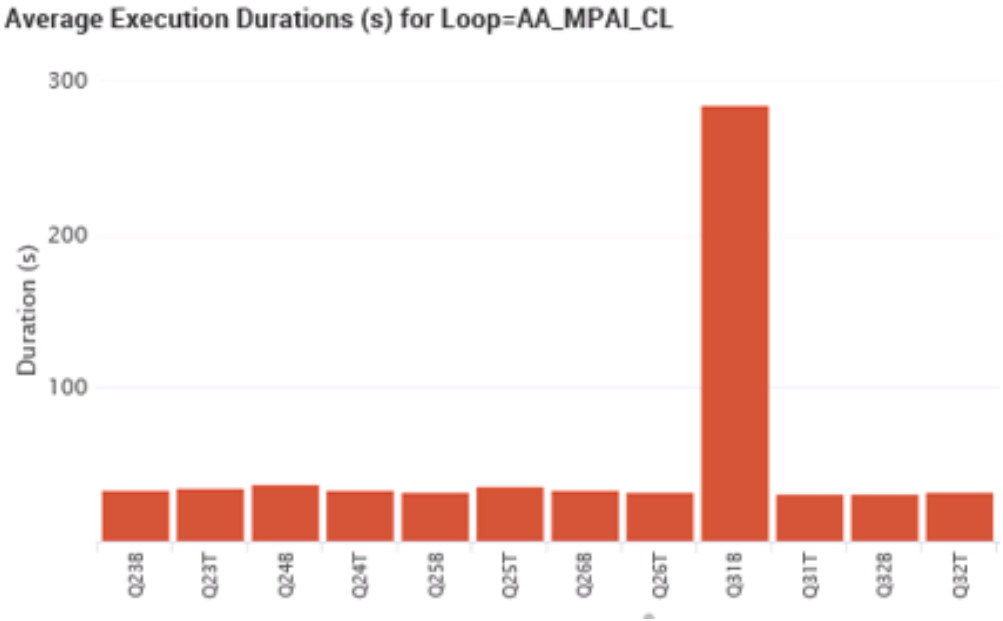

Discovering insightful performance metrics and collecting them into effective dashboard visualizations helped to identify failing hardware, outdated reference data and incorrect control system configuration

Data-Driven Project Management: Jira Developer Load Planning

- **Visualizing planned load from Jira**
  - **Splunk connected to Jira database**
  - **Custom status dashboards**
  - **Powerful data analysis**
  - **Drill-downs for details**
  - **Links to Jira**

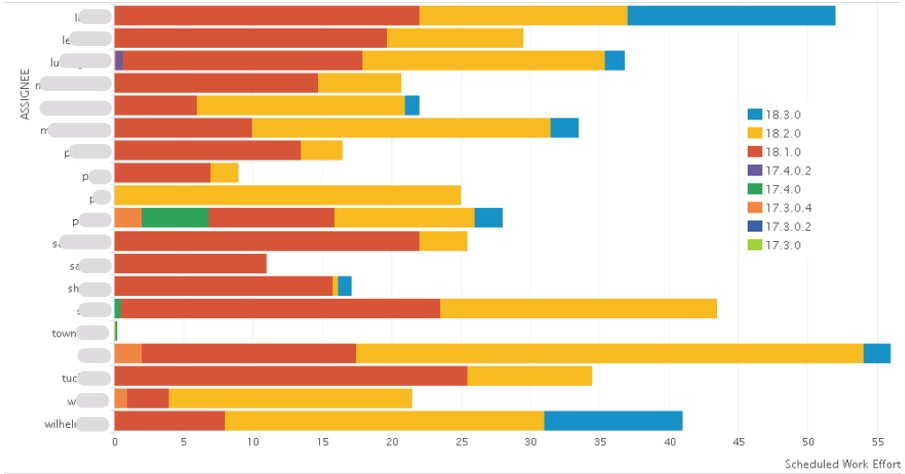

Filter: Assignee=w

Issue Key	Summary	Component
IC-45141	Repair parallel OI LOIS and AA PAM segment command deadlock	Auto Align,Database/Operational Data
IC-47035	CMS ISP_CL_ALIGNMENT update to use updated motor setpoints like PSS	Auto Align
IC-47857	AA CCRS coarse move segment commands	Auto Align,Database/Operational Data
IC-41337	remove obolsete AA and CMS general	Auto Align
IC-47641	OTSL FOA AA Loops	Auto Align

Leveraging Splunk dashboard skills and experience with Search Processing Language (SPL) to apply powerful analysis and visualization techniques to software team management

Data-Driven Project Management: Support Engineering Process

- **Support introduction of a new engineering process or a policy with a status dashboard**
- **Immediate feedback to monitor acceptance progress and focus attention**
- **Fast, low effort development**
- **Access to enterprise data systems**

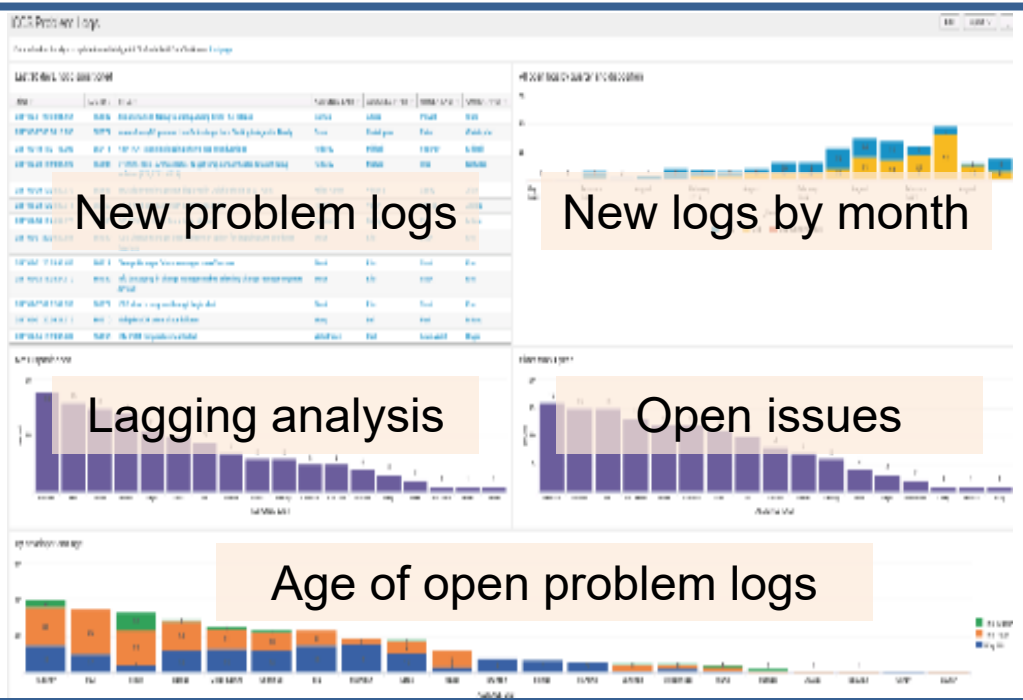

Clear, real-time, online visualization of a new policy rollout provides immediate feedback to the team and supports rational acceptance of new process across organization

Data-Driven Project Management: Interdisciplinary Communications

- **Preventing misunderstandings in multidisciplinary projects**
- **Define “what matters”, metrics which can be measured**
  - **Clear success criteria**
  - **Supported by a quick visualization**
  - **Monitor progress**
  - **Focus attention**
  - **Complete and finalize**

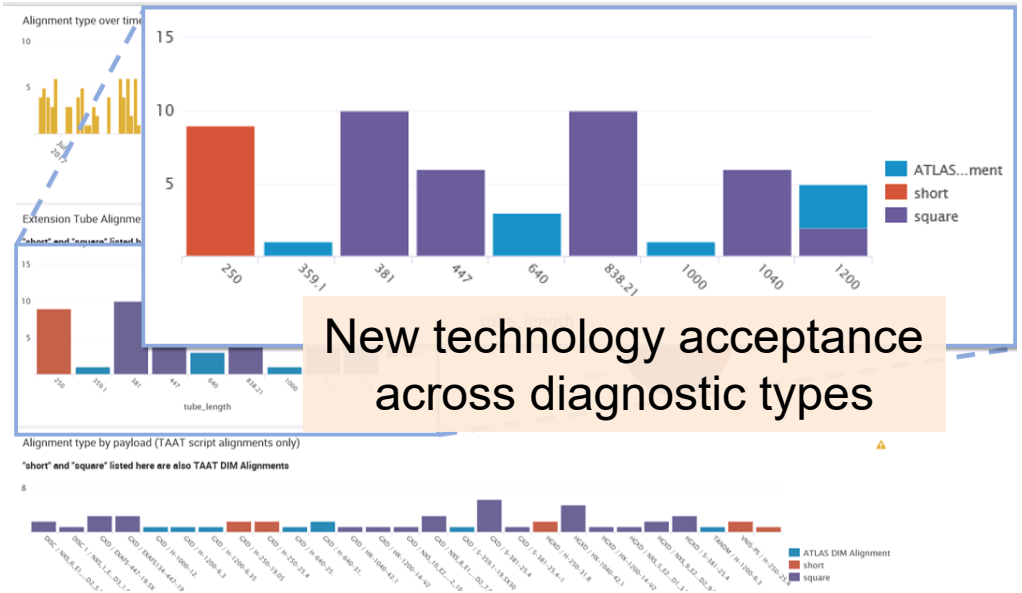

Agreeing upfront on “what matters”, and developing dashboards of key performance indicators (KPIs) helps to avoid misunderstandings in multidisciplinary projects.

Monitoring Splunk performance

- **Splunk performance indicators**
  - **Data source for built-in and custom dashboards**
  - **Daily indexing volume, storage size, index age**
  - **Indexing and search performance metrics**
  - **User access audits**

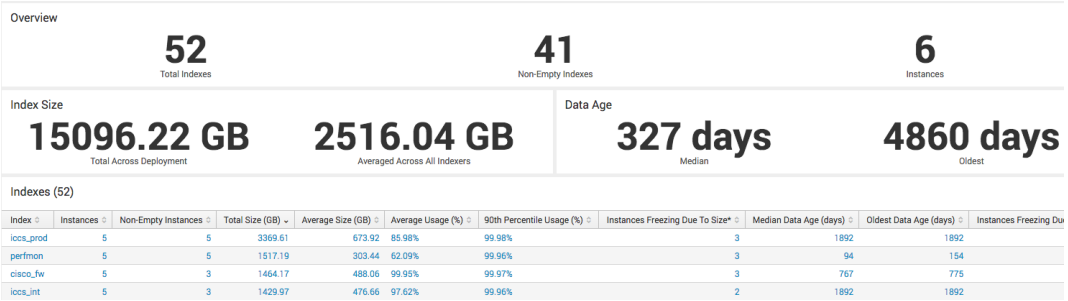

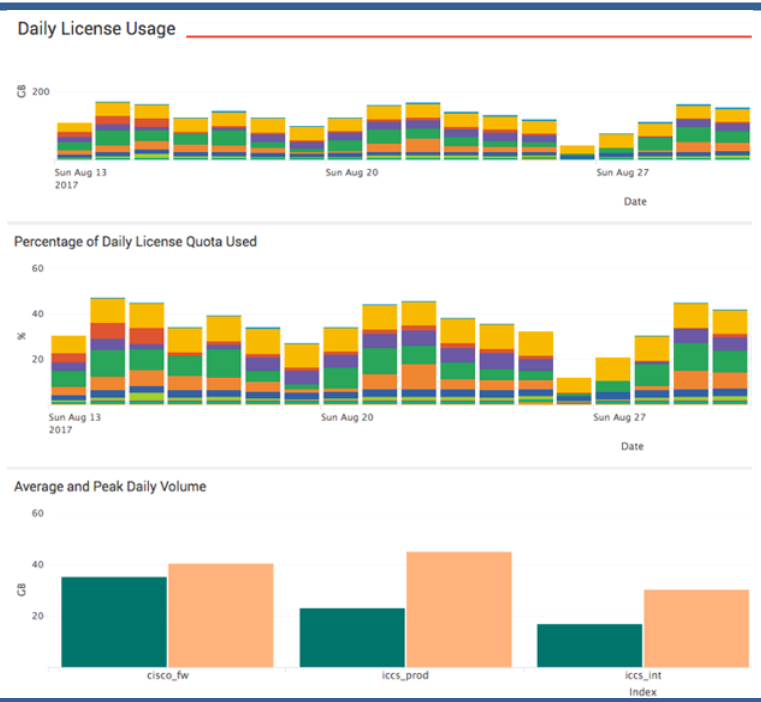

Consistent with its universal data model approach, Splunk exposes its own performance data to SPL queries and visualizations so the metrics can be monitored with either built-in or custom dashboards

Splunk performance: lessons learned

- **Search and Indexer require a cluster of performant bare metal hardware**
- **A deployment server is required to manage a production system**
- **Sizing and quantity of indexer "buckets" are critically important for Splunk performance, it is the best to set them correctly from the beginning**
- **Virtual machines are sufficiently performant for forwarders and auxiliary functions**

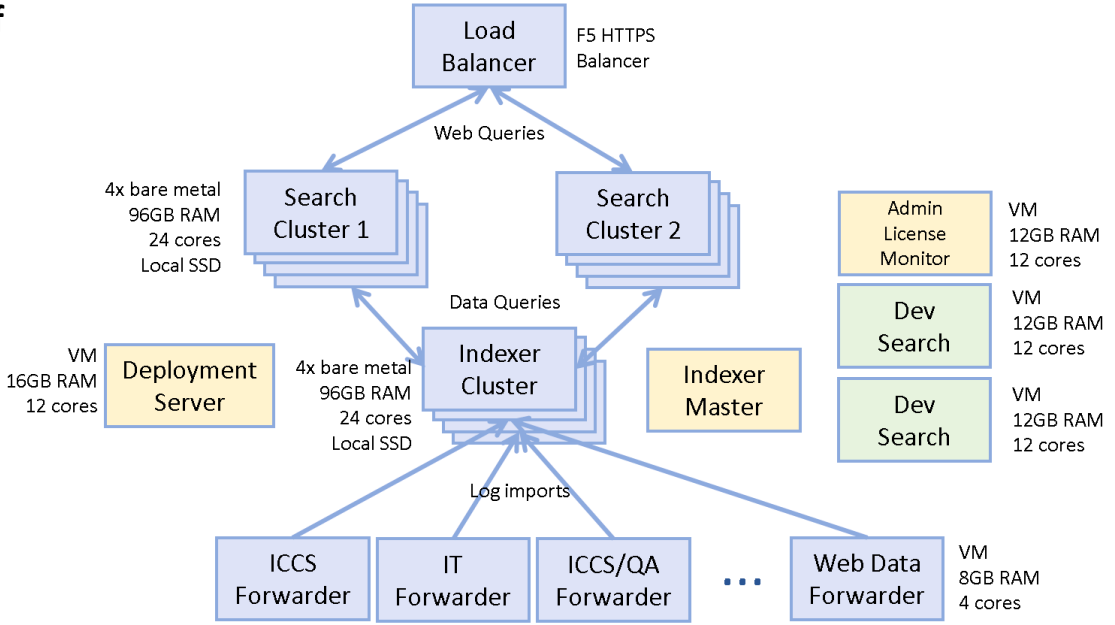

With our increased reliance on Splunk for monitoring and analysis, the indexer and search loads have grown. NIF Splunk system configuration has been adjusted to maintain performance.

Benefits of using Splunk for monitoring and management of NIF control system

- Universal data analysis and visualization tool
- Efficient schema-less indexing of unstructured log files
- Rapid “one-line” data analysis with SPL
- Ease of creating effective web visualizations
- Access to training, support, online community


Lawrence Livermore
National Laboratory