ISBN: 978-3-95450-193-9 doi:10.18429/JACoW-ICALEPCS2017-TUCPA02

LEVERAGING SPLUNK FOR CONTROL SYSTEM MONITORING AND MANAGEMENT

M. Fedorov, P. Adams, G. Brunton, B. Fishler, M. Flegel, K. Wilhelmsen, R. Wilson,
Lawrence Livermore National Laboratory, P.O. Box 808, Livermore, CA 94550, USA

Abstract

The National Ignition Facility (NIF) is the world's largest and most energetic laser experimental facility, with 192 beams capable of delivering 1.8 megajoules (MJ) and 500-terawatts of ultraviolet light to a target. To aid in NIF control system troubleshooting, the commercial product Splunk was introduced to collate and view system log files collected from 2,600 processes running on 1,800 servers, front-end processors, and embedded controllers [1]. We have since extended Splunk's access into current and historical control system configuration data, as well as experiment setup and results. Leveraging Splunk's built-in data visualization and analytical features, we have built custom tools to gain insight into the operation of the control system and to increase its reliability and integrity [3]. Use cases include predictive analytics for alerting on pending failures, analyzing shot operation critical paths to improve operational efficiency, performance monitoring, project management, and analyzing and monitoring system availability. This talk will cover the numerous ways we have leveraged Splunk to improve and maintain NIF's integrated control system.

INTRODUCTION

The National Ignition Facility at Lawrence Livermore National Laboratory (LLNL) is the world's most energetic laser system for experimental research in inertial confinement fusion (ICF) and high-energy-density (HED) physics. The NIF laser system consists of 192 laser beams, which are focused inside the 10-meter Target Chamber (TC), delivering up to 1.8 MJ of ultraviolet light onto the target.

Since 2013, the NIF Information Technology (IT) and Integrated Computer Control System (ICCS) organizations have been relying on Splunk for collecting, managing and analyzing computer log files [2]. Splunk is a commercial software system for processing unstructured log files into a centralized indexed database with powerful search, data processing, and visualization capabilities.

Based on the positive experience and value of the analytics insights, Splunk monitoring of the NIF control system has been extended and now includes all logs generated for one year. The production NIF control system generates 20-50 GB of logs per day, which constitutes 20-25% of total NIF Splunk daily volume. Control system log storage is currently 3.4 TB, while total NIF Splunk data size is 15.1TB.

Splunk's ability to process unstructured log files is of key importance for a specialized control system such as ICCS. Many log analytics systems are fine-tuned for a specific IT application: webserver logs, database logs, firewalls, etc. The majority of ICCS software is developed in-house and is unique to NIF – there is no commercial vendor or a market to develop analysis tools for our control system. With Splunk, there is a simple setup step of configuring log data sources: where they are coming from and the general format: timestamp, hostname and text body. Once sources are configured, the body of the log message is not constrained; any text will be imported, indexed and stored. There is no fixed data schema –- search, data extraction, analysis, and visualizations can be performed on any part(s) of the log messages.

In addition to its primary indexed log file storage, Splunk supports connectivity to external data sources and databases. For controls system applications, Splunk is connected to ICCS configuration and NIF Archive Oracle databases. For data-driven project management, Splunk is connected to NIF enterprise management and problem tracking systems, IT inventory, and monitoring databases and Jira issue-tracking software.

Since deployment, Splunk has become the primary tool for ICCS log analysis, and we have retired the previously used "snapshot" log capture system [2]. Splunk online training materials have helped to introduce ICCS developers to Splunk and its Search Processing Language (SPL). Many developers have progressed into advanced courses to achieve Splunk Power User certification [5]. ICCS software expertise is in the server-side and embedded control applications, not in the Web visualizations. With Splunk and SPL, our software engineers can create visualizations and dashboards without Web development skills and with minimal overhead. To encourage developers' adherence to best logging practices, we have placed ICCS development and test environments under Splunk monitoring. By making Splunk available early in the development cycle, we have assured that all interesting data are logged, developers practice their SPL skills, and Splunk dashboards are developed and tested well before production deployment.

CONTROL SYSTEM MONITORING

Performance Monitoring and Alerts

One of the traditional applications of the system monitoring tools is to observe "vitals" at the Operating System (OS) level: CPU load, memory utilization, and swap use. In a large distributed control system such as ICCS, visualization of this information coming from hundreds of computer hosts presents a challenge. If shown individually there are too many screens, and it is unclear what is normal and what is not. Combining all hosts into one screen creates a noisy, unusable chart.

The readability and usefulness of ICCS performance monitors greatly improved after we configured a custom Splunk dashboard which segments the "vitals" into several groups of comparable hosts: framework servers (Solaris), supervisory applications (Linux), front-end processors (FEPs) (Linux, Windows), and operator consoles (Windows) (Fig. 1).

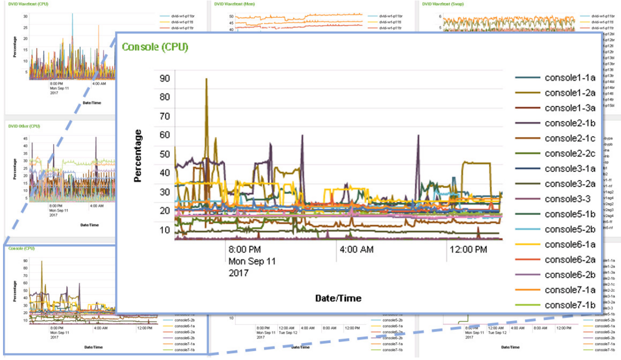

Figure 1: Operating System "vitals" segmented by host type.

Resource utilization patterns vary from group to group; some FEPs are continuously acquiring data, while framework servers peak at the time of a NIF shot. But within these groups, the utilization patterns are similar, so the well-behaving majority can serve as a visual reference for the group. After we have gathered the similarly functioning hosts into the groups, normal utilization patterns have clearly emerged on the charts, and abnormal outliers became easily identifiable.

Interestingly, at that time we did not even have Splunk collecting "vitals" raw data. While Splunk forwarders can be deployed and configured to capture OS load metrics, we did not have Splunk forwarders installed on every control system host. However, we previously had Oracle Enterprise Manager (OEM) deployed and already collecting OS load metrics. By configuring a Splunk data source pointing to the OEM database, we have obtained access to all historical and new streaming OEM data. Using SPL, we have quickly prototyped data segmentation and visualization and come up with an effective dashboard, optimized for our system usage patterns.

In addition to the visual dashboards, Splunk alerts were configured to email responsible individuals (RIs) whenever unhealthy levels of OS loads are recorded. After being set and almost forgotten, these alerts have repeatedly proven their value by sending email notifications when a software or system reconfiguration has produced an unexpected, surprising effect on one of the systems. For example, a new framework tool is deployed across all video front-end processors (FEP). As expected, and proven by the pre-deployment tests, 99% behave correctly. However, three x-ray imaging systems got pushed into high random access memory (RAM) utilization due to the limitations of their hardware platform. The problem was noted and preemptive mitigation actions have been initiated based on the Splunk email alert.

In addition to OS "vitals," similar monitoring is configured at the Java Virtual Machine (JVM) level, since the majority of the control system applications are Java-based. Heap size and Garbage Collection (GC) times are proven to be the most useful for producing relevant actionable alerts.

At the lowest hardware interfacing level, Splunk is setup to monitor and alert when VxWorks motion control software detects an unexpected VME input-output (I/O) board state. Since such situations require an immediate repair action, these Splunk alerts are routed to the NIF hardware technician paging system.

Long-Term Trend Monitoring

Unlike traditional database or archive systems, Splunk does not have a schema. There is no need to define data values of interest up front. Instead, any part of the logging output, or a combination of such parts, or their timestamps, duration, or patterns, can be defined as a data value in SPL expression and used for extraction, analysis, and visualization. The flexibility of going back and defining "schema" in the past is especially helpful when a new problem or concern comes to attention. Months of bulk logs captured by Splunk then become a data mine for insights supported by quantitative data.

The NIF laser system utilizes large, state of the art capacitors to pulse its flash-lamps during the shot. The capacitors age and eventually require replacement to prevent failures. The health of the capacitor can be assessed during an inspection, which requires operational downtime. We have been asked to help pre-select at-risk capacitors for inspection prioritization. By reviewing logs produced by the NIF Power Conditioning FEP, we have identified a log entry relevant to capacitor health. With the logs collected in Splunk storage, we immediately gained access to a one-year time-series for each of NIF's 192 capacitor bank modules. By comparing trends for good and failing capacitors, we have established criteria for the at-risk inspection.

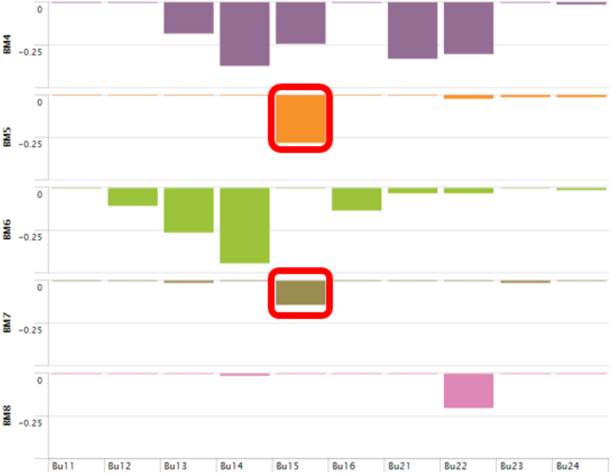

Figure 2: Capacitor health dashboard (fragment).

Using Splunk rapid development tools, a dashboard was developed for the NIF Power Conditioning group which guided prioritization of the inspections and replacement processes (Fig.2).

Gantt Chart Visualization of NIF Shot Cycle

Experiments at NIF are structured as "shots" – sequences of steps in which the entire facility is orchestrated to get set up, aligned, and then driven into main laser firing by the ICCS Shot Director software [4].

Most of the facility troubleshooting or performance analysis starts with a time-sequence question – when has a given activity happened, or what was the facility doing at a certain time? The Shot Director software does log when and what step has been executed; however, the volume of the data is overwhelming.

To help navigate through the NIF shot sequences, a hierarchical dashboard was developed, relying on Splunk's "drill-down" feature – a custom action can be assigned to an element on the chart, such as expanding that element into the next level of detail. At the top level of the dashboard the multi-day overview is presented, with several NIF experiments ("shots"). For each shot on the graph, only key state transitions are shown, such as "Begin Shot," "Implement Plan," "Ready," "Rod," and "System Countdown." (Fig.3)

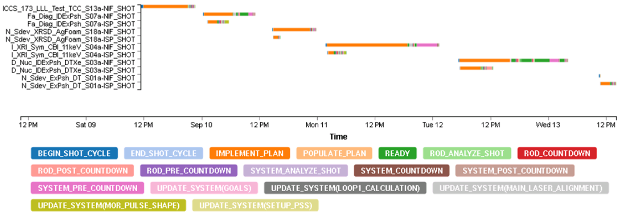

Figure 3: Gantt chart visualization of several NIF shots

The state bars are clickable, so the user can narrow in on one of states of a specific shot and see more details (Fig.4).

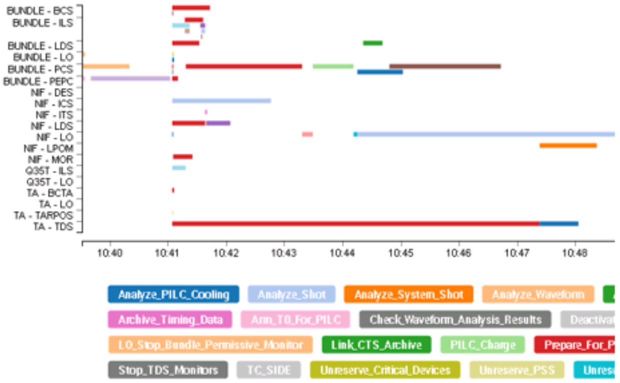

Figure 4: Drill-down into shot countdown (fragment)

The steps on this level also support Splunk "drill down." For example, if a critical-path step is being investigated, the next screen may reveal that the Quad 46B Power Conditioning unit took an abnormally long time to charge, and therefore is responsible for the critical path delay (Fig.5)

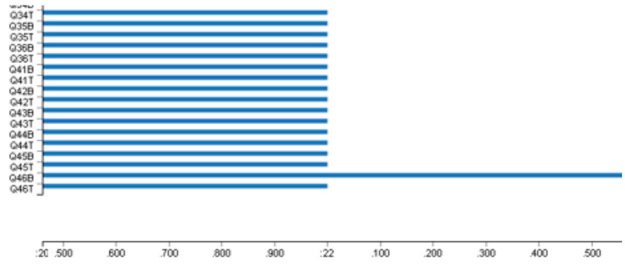

Figure 5: Abnormal charge time in one laser quad.

Finding Laggards

The NIF laser has 192 laser beams organized into 48 quads or 24 bundles, further grouped into 4 clusters in 2 laser bays. The underlying symmetry in NIF hardware design leads to the expectation of similar performance across the locations. As demonstrated by the Fig.5 example in the previous subsection, a drastically different performance from one location is likely to indicate a hardware or system configuration issue.

This comparative approach has been utilized by NIF Automatic Alignment to identify laser beams and quads which throttle overall performance. Laser alignments are performed automatically and concurrently, so the slowest performing location determines the total duration of the operation.

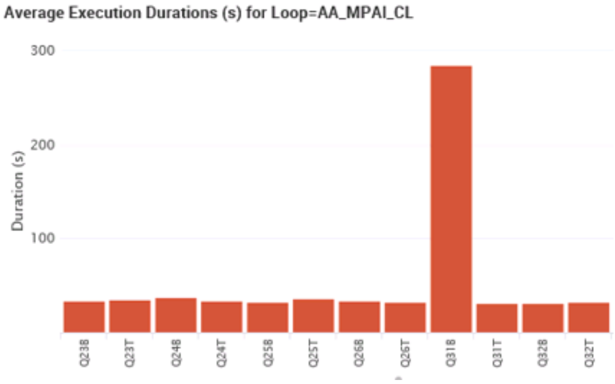

Figure 6: Duration of alignments across quads (fragment).

The alignment duration dashboard tool identifies alignment operations and laser locations with abnormally slow performance compared to average durations (Fig.6). The tool helped to identify failing hardware, incorrect reference data, or control system configuration. As a result, the underlying issues were resolved and performance was improved, supporting a higher shot rate for the entire facility [3].

SPLUNK FOR PROJECT MANAGEMENT

Data-Driven Project Management

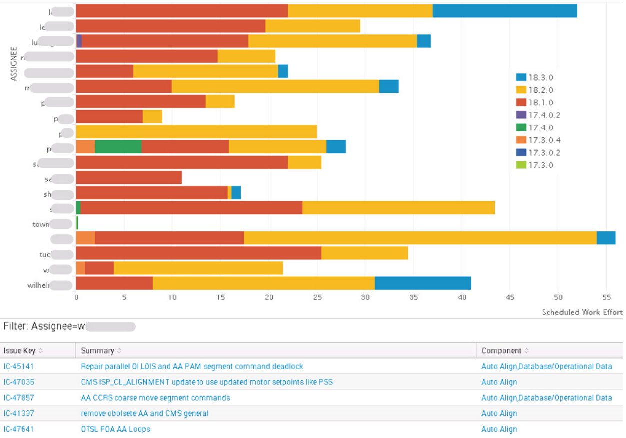

Filter: Assignee=wi

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

Figure 7: Developer load planning, Jira data connection.

Once the ICCS team became familiar with Splunk for analysis and visualization of the control system events and metrics, it became natural and efficient to use same tool for building dashboards for other tasks such as project management.

ICCS is relying on Atlassian Jira [6] for tracking software issues and development tasks. Jira comes with its own dashboards, but we wanted the full power of data manipulation, custom visualizations, and drill-downs. Using the Splunk DB Connect plugin [7], we have exposed Jira data to Splunk and developed dashboards for monitoring and planning team workloads (Fig. 7).

Visualizations to Support Engineering Processes

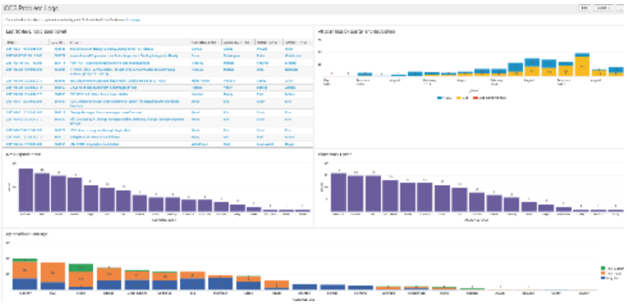

Figure 8: Issue Disposition Dashboard (overview)

We have found that a clear, relevant dashboard makes an effective tool to encourage or enforce new management processes across the team.

For example, when we needed to focus developers' attention on prompt analysis and disposition of the incoming problem logs, we established a new policy with timelines for the initial assessment and resolution of the issues. Introduction of the policy was supported by development of a dashboard which reflects developers' progress toward satisfying the established goals. The dashboard is projected during the status meetings and immediately indicates how well the new policy is being implemented and where attention needs to be focused (Fig. 8, top-level overview)

Universal access to all data with DB Connect and ease of rapid dashboard development with SPL were key enabling factors, since long development time or a significant overhead would make this approach unfeasible.

Support Interdisciplinary Communications

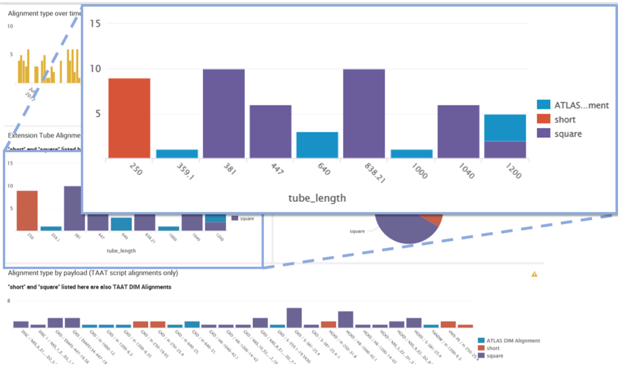

Figure 9: Analysis of alignment tool utilization patterns.

With firm adoption of Splunk within our ICCS computer controls team, we have started to use dashboards when communicating with non-computer groups across the larger NIF organization. Clear, accurate, real-time dashboards with visualizations of the key performance indicators (KPIs) help to avoid misunderstandings in multidisciplinary projects.

In the diagnostic alignment area, the ICCS team was responsible for the introduction of two new alignment efficiency tools: The Target Area Alignment Tool (TAAT) [8] and the Advanced Tracking Laser Alignment System (ATLAS) [9]. Both ICCS software and alignment teams were interested in knowing how often the new tools were used and how successful they were. A Splunk dashboard was developed to monitor tools utilization patterns over time as well as by different diagnostic categories. The dashboard has helped both teams to track problems and direct efforts during introduction of these new technologies (Fig. 9).

THE NIF SPLUNK PLATFORM

With the development of many new dashboards, and finding new uses for Splunk analysis within the control system and across the entire NIF organization, the load on our Splunk indexing and web search engines has increased. To assure availability and performance, Splunk "health" also required monitoring.

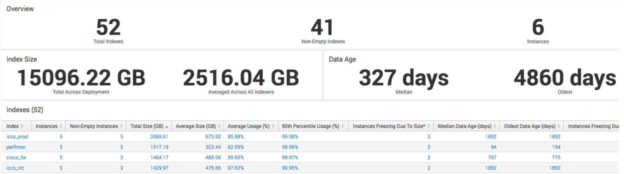

Overview

52 Total Indexes

41 Non-Empty Indexes

6 Instances

Index Size

15096.22 GB Total Across Deployment

2516.04 GB Averaged Across All Indexers

Data Age

327 days Median

4860 days Oldest

Indexes (52)

Index	Instances	Non-Empty Instances	Total Size (GB)	Average Size (GB)	Average Usage (%)	90th Percentile Usage (%)	Instances Freezing Due To Size*	Median Data Age (days)	Oldest Data Age (days)	Instances Freezing Du
iccs_prod	5	5	3369.61	673.92	85.98%	99.98%	3	1892	1892	
perfmon	5	5	1517.19	303.44	62.09%	99.96%	3	94	154	
cisco_fw	5	3	1464.17	488.06	99.95%	99.97%	3	767	775	
iccs_int	5	3	1429.97	476.66	97.62%	99.96%	2	1892	1892	

Figure 10: Splunk index summary dashboard

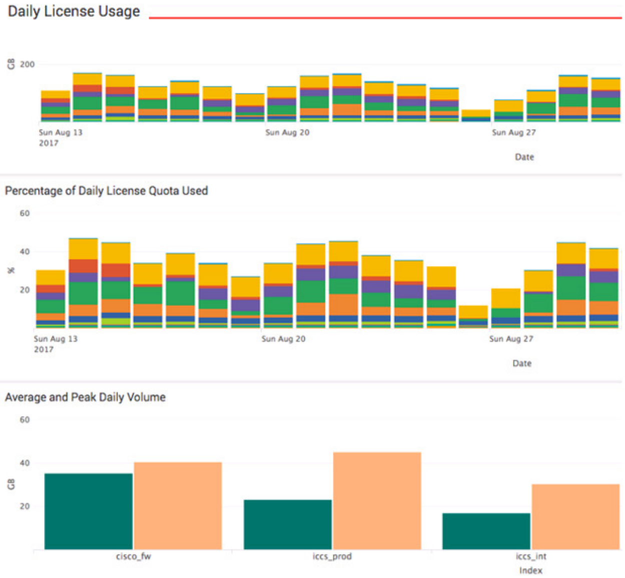

Figure 11: Daily usage visualization.

Consistent with the universal data access approach, Splunk exposes its own internal performance data to SPL queries and visualizations so the metrics can be monitored with either built-in or custom dashboards. (Figs.10, 11)

At one point, our increased Splunk usage has resulted in search and indexing delays. By learning about Splunk performance tuning and through the on-call and on-site customer support, we have identified deficiencies in our early architecture and deployment:

- Search and Indexer configurations require a cluster of performant bare metal hardware, not a virtual machine.
- A deployment server is required to manage a production system; maintaining configuration files manually is not sustainable and is inconsistent and error-prone.
- Sizing and quantity of indexer “buckets” have critical effects on the indexer performance; e.g., oversized “buckets” never transition to “cold” long-term storage and may have a blocking effect on new data imports.
- VMs are sufficiently performant for forwarders and auxiliary functions (deployment server, cluster masters, administration, license, and monitoring).

The NIF Splunk architecture has been adjusted to address these performance recommendations (Fig.12)

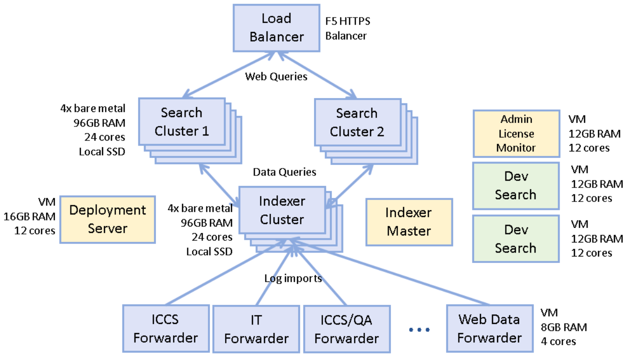

Figure 12: Optimized NIF Splunk platform

CONCLUSION

Our four years of experience using Splunk at NIF/ICCS has confirmed its value for control system monitoring and its uses have grown into numerous data analysis and visualization applications benefiting project management and facility operations.

The broad adoption of Splunk was facilitated by the following qualities of the system:

- Universal data analysis and visualization tool
- Efficient schema-less indexer of unstructured log files
- Connectors to external databases and data sources
- Rapid “one-liner” data analysis with SPL
- Ease of creating effective web visualizations
- Access to training, support, and an online community [10]

We have learned that Splunk system architecture and storage configuration may have a dramatic effect on the indexing performance and user experience. NIF/ICCS Splunk platform and configuration have been tuned and adjusted to accommodate growth of the indexed data, an expanded user base, and increased complexity of the data analysis and visualizations.

ACKNOWLEDGMENT

This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344. LLNL-CONF-739363.

REFERENCES

[1] G. Brunton, *et al.*, “Status of the National Ignition Facility (NIF) Integrated Control and Information Systems”, *16th International Conference on Accelerator and Large Experimental Physics Control Systems* (*ICALEPCS2017*), Barcelona, Spain, October 2017, MOAPL03.

[2] J. Fisher, *et al.*, “Monitoring of the National Ignition Facility Integrated Computer Control System”, *14th International Conference on Accelerator and Large Experimental Physics Control Systems* (*ICALEPCS2013*), San Francisco, CA, 2013.

[3] G. Brunton, *et al.*, “Shot Rate Improvement Strive for the National Ignition Facility (NIF),” *15th International Conference on Accelerator and Large Experimental Physics Control Systems* (*ICALEPCS2015*), Melbourne, Australia, MOD3O03

[4] D. Mathisen, “Orchestrating Shots for the National Ignition Facility”, IAEA 8th Technical Meeting, San Francisco, 2011.

[5] Splunk Certified User, `https://www.splunk.com/view/education/SPCAAAJEN`

[6] Atlassian Jira Issue and Project Tracking Software, `https://www.atlassian.com/software/jira`

[7] Splunk DB Connect, `https://splunkbase.splunk.com/app/2686/`

[8] M. Fedorov, et al., “New visual alignment sequencer tool improves efficiency of shot operations at the National Ignition Facility”, *16th International Conference on Accelerator and Large Experimental Physics Control Systems* (*ICALEPCS2017*), Barcelona, Spain, October 2017, TUMPA01.

[9] E. Wilson, *et al.*, “Experiences with Laser Survey Instrument Based Approach to National Ignition Facility Diagnostic Alignments”, presented at the *16th International Conference on Accelerator and Large Experimental Physics Control Systems* (*ICALEPCS2017*), Barcelona, Spain, October 2017, MOCPL02.

[10] Splunk Answers, `https://answers.splunk.com`