Telemetry-enabled Customer Support using the Cray System Snapshot Analyzer (SSA)

Jeremy Duckworth, Jay Blakeborough, Kevin Coryell, Scott McLeod
Global Technical Service
Cray Inc.
Seattle, Washington, USA
{jeremyd,jblake,kcoryell,smcleod}@cray.com

***Abstract*— SSA is a Cray customer service application designed to support product issue diagnosis and reduce time to resolution. SSA is focused on the submission of product telemetry from a customer system, to Cray. SSA provides value to the support process by automating (1) the collection, submission and analysis of product diagnostic information (2) the collection, submission and analysis of product health information and (3) key aspects of the customer support process. The focal topic for this paper is a discourse on the benefits of SSA use. For background, we will provide a general overview of SSA and references for further reading. Next, we will provide an update on SSA product release and operations history. Finally, we will discuss the anticipated product roadmap for SSA.**

Keywords: Customer Support, Telemetry, Remote Product Monitoring

I. INTRODUCTION

A. Cray Customer Service

Cray is a global leader in supercomputing. Cray understands it is not just the products we make, but the customers we service, that ensure our success. That is why we are dedicated to meeting the needs of those who depend on our products by providing outstanding aftermarket support. To complement our highly-skilled and experienced global service teams, Cray continues to invest in and innovate around technologies designed to improve the support experience. Initially announced at the Cray User Group (CUG) Conference in 2015, the Cray System Snapshot Analyzer (SSA) is a relative newcomer to Cray's service technology framework – with a focus of providing telemetry-enabled support for our products.

B. Technology

SSA is a Cray customer service application designed to support product issue diagnosis and reduce time to resolution. SSA is focused on the submission of product telemetry from a customer system, to Cray.

SSA provides value to the support process by automating (1) the collection, submission and analysis of product diagnostic information (2) the collection, submission and analysis of product health information and (3) key aspects of the customer support process. This automation is enabled by a simple architecture that is based on making it easy to understand what information is being collected, how the information is transmitted to Cray and how Cray will then use the information. SSA can be used after a problem with a product is identified or, in some cases, SSA can proactively identify a problem. Through the collection of product configuration data on a routine basis, SSA can also track historical changes in a product's configuration – offering another source of diagnostic information to aide in problem resolution.

SSA has been designed to minimize the customer installation and configuration burden – making it easy to install and configure. Through integration with Cray's customer portal (CrayPort) [1] – we have also worked to ensure activation and management of the account information for SSA is not complicated. SSA is currently available free of charge to customers with a Cray service contract, on supported platforms.

C. Paper Organization

Section II of this paper provides background information, including terms and definitions, our motivation for pursuing this work, a brief history of SSA and references for further reading. Section III provides a discourse on the benefits of SSA use, followed by a presentation of select SSA operational metrics in section IV. Section V discusses future work, and the conclusion follows in section VI.

II. BACKGROUND

A. Terms and Definitions

1) Telemetry

The Merriam-Webster Online Dictionary [2] defines *telemetry* (in part) as, "*the science or process of telemetering data*" and "*data transmitted by telemetry*". They define *telemetering* [3] (in part) as, "*an electrical apparatus for measuring a quantity (such as pressure, speed, or temperature) and transmitting the result especially by radio to a distant station*."

For the purposes of this paper, we emphasize the remote, regular (rhythmic) and automated connotations

associated with telemetry – as a conceptual platform. Canonically, SSA is focused on the routine capture of information from Cray products in a customer environment and transmission of the resulting information securely back to Cray. Then, by analyzing and acting on the data, SSA is designed to deliver customer value. We overload the term (telemetry), as not all information captured and transmitted by SSA is automated by design and not all data captured can be strictly considered quantitative. We also believe there is residual value in the use of SSA when the telemetry apparatus is not active.

B. Motivation

Systems with similar properties to SSA are not novel to the computing industry. However, SSA is a relatively new platform for Cray. Through concerted collaboration within Cray and with our customers, we are positioning the (SSA) platform to *assist* in gradually realizing a simple goal – *maximize Cray product resources that are available for customer use*. Abstractly, we can approach this ideal goal intuitively through optimization problems. These are:

(1) Minimize Product Downtime

Example methods and practices include:

- Ensure product operators are trained and properly equipped, have contextual product information available when needed, and apply methodical product management practices
- Ensure product is properly configured
- Ensure product operating environment is within manufacturer specification(s)
- Undertake proactive maintenance actions to address a continuum of product failure modes, as early as possible to avoid or abate impact
- Minimize the amount of time an unanticipated product failure takes to resolve
- Ensure product updates are applied in a timely and consistent manner

(2) Maximize Product Performance

Example methods and practices include:

- Identify when products are over-subscribed for customer workload(s)
- Identify product use instances or patterns that yield poor or sub-optimal performance relative to the product's design
- Recommend product configuration change(s) that yield better resource availability and/or higher performance resources as a dynamic to customer workload(s)

The two problems above should be considered interdependent. While perhaps well known to experienced technical or high performance computing (HPC) professionals, finding practical solutions to these problems is often very challenging – notably as the size, complexity and required availability of a solution increases.

Today with SSA, our focus is on delivering capabilities designed to minimize product downtime. We are also using SSA to drive efficiency within the customer support process. Our motivation for this paper is to publish our position on how SSA can provide value to Cray customers – so that customers can make an informed decision on their participation in SSA.

C. Cray System Snapshot Analzyer (SSA)

1) General Architecture

The SSA architecture can be distilled into three areas, (1) cross-product standard client integrated into product-specific interfaces, (2) the data upload interface accessible at `https://ssa.cray.com` and (3) SSA back-end analytics and automation. Data flows from the client, to the upload service. Data is uploaded either on a recurring, automated schedule or on-demand, depending on the specific use of SSA. No content is downloaded to the customer system. Fig. 1 illustrates the high-level SSA Architecture. For additional detail on the SSA Architecture, please consult the SSA Whitepaper [4].

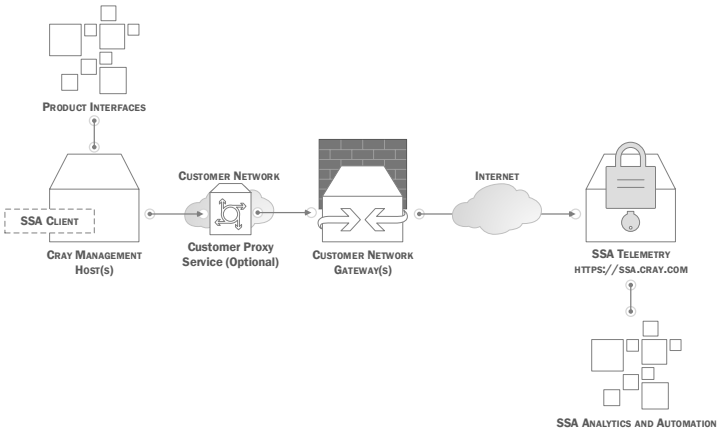

Figure 1. High-level SSA Architecture

2) Program Management

The SSA program is managed by the director of Cray's Global Technical Service Organization (formally Continuation Engineering). The SSA Development Team is currently comprised of four, full-time software engineers. The SSA Development Team is responsible for a quarterly SSA client release cadence on qualified products, in addition to the service infrastructure that supports the

`https://ssa.cray.com` interfaces and SSA automation system. The SSA Team works closely with many Cray internal organizations for alignment towards driving value for Cray customers and Cray.

3) *Brief History*

The SSA program started internally at Cray in 2012, followed by a customer early access release in 2014 and an initial general availability (GA) release in June of 2015. We currently support the Cray® XE6™, XK6™, XK7™, XC30™ and XC40™ systems running the SMW 7.2- and CLE 5.2-based software environments and the Cray® Sonexion 900™, 1600™, 2000™ and 3000™ systems running the 1.2, 1.3, 1.4, 1.5, 2.0 and 2.1 software environments. There have been nine client releases in SSA history, the most recent being the 1.4.0 release – made available in March 2017. SSA clients are normally released on a quarterly basis.

4) *Getting Started with SSA*

SSA is currently available free of charge to customers with a Cray service contract on supported platforms. CrayPort Knowledge Article ID 4546, "*Getting Started with the Cray System Snapshot Analyzer (SSA)*" [5] contains the necessary steps to download the SSA client, accept the end user license agreement (EULA) [6] and manage activation of the upload service account. Release notes for each release and other release-specific documentation are available for download. The release notes contain hyperlinks to the SSA User Guide associated with the release. For additional SSA resources, please visit the SSA section of Cray's website [7].

III. SSA VALUE PROPOSITION

In this section, we present the current and prospective capabilities of SSA. We also present a value proposition for each capability. While we do expect to eventually develop the capabilities discussed as prospective, their realization is not guaranteed in whole or in the form suggested.

Results of analysis generated after SSA data is uploaded to Cray are currently only available to Cray employees. We are actively investigating the introduction of interfaces that would allow customer consumption of these results. Data SSA collects can be reviewed locally if desired, see the SSA Whitepaper [4] for additional information.

Operationally, SSA is evolving based on feedback from customer support cases, Cray engineers and Cray customers. The resulting feedback model is the engine that drives iterative improvement in SSA.

A. *Product Inventory and Configuration Baselining*

1) *Current Capabilities*

A daily scheduled data collection and upload process takes place on products with SSA installed and configured for upload. The data collected consists of system hardware and software inventories, configuration and other sources of product data. A single collection can consist of hundreds to thousands of data points. Over time, a repository of data collections is stored, each representing a point-in-time snapshot of the source product.

As part of these snapshots being uploaded to Cray, a change-analytics system analyzes each data element in a snapshot (e.g., files, command outputs and other data sources) to identify when changes occur – relative to other snapshots captured against the product.

2) *Customer Value*

a) Cray has to ask fewer questions of the customer when addressing a support concern

With very few exceptions, Cray requests, and our customers are required to provide, details on the inventory or configuration of their system. Delays in satisfying these requests can result in adding time to the resolution of a support request. This effect can be compounded if multiple parties are involved on either end of the dialogue.

Cray's ability to query a database of information for the data they would normally request can reduce the amount of time required to resolve a support request.

b) Cray can proactively search the SSA database to answer questions that previously required a field survey

On a frequent basis, a survey of the field for information on characteristics of installed products is necessary. Sometimes the questions are simple, such as, *what is the current distribution of programming environments and versions in use*? Sometimes the questions are more complicated and contextual, such as, *how many of a specific processor and memory combination are running in a computing element with firmware version x, driver version y, …*? *How many combinations are there*?

Cray's ability to query a database of information for the data we would normally have to survey for should reduce the amount of time required of customers and Cray field engineering. Automating the collection and analysis of the information also improves the freshness and quality of the data, leading to better insights.

c) Cray can leverage historical data from a reporting product to understand 'what a system looked like' at a previous point in time, and what may have changed

Product configuration and state change often. Change is introduced from many sources, including system failure, conscious management decisions or opaquely by system updates. Managing conscious configuration

change is difficult – and unauthorized or poorly socialized changes are not uncommon. In systems where performance is a first-class citizen, the performance of the system is very often associated with concerted configuration tuning against a known baseline. When these tuning parameters are changed (or not established on system reboot/reset cycle) or the baseline changes, it is common that the performance characteristics of the system degrade.

There are many tightly coupled change management solutions available on the market. These solutions typically require additional commitment or investment – notably deep system expertise, customer staff requirements and rigid systems management guidelines. The solutions are typically policy-based, where change is accounted for against a policy and policy violations are expressed as exceptions. From a vendor (Cray) perspective, this class of application is not suitable for use diagnostically (to help Cray resolve a customer support concern).

Deriving value from change analysis typically requires, a) separating noise from signal b) understanding the change context and application(s) wherein the change occurred and c) causal correlation of the observed behavior with the change. Ideally, in practice, the use of change analysis to diagnose or troubleshoot a problem is iterative and informed by a rough timeline and set of symptoms or observations.

As previously stated, SSA routinely collects system configuration and state data and stores it as a series of time-centric snapshots – also cataloguing change as a feature of the Cray-centralized analytics system. This change is expressed in a form that identifies when an individual element is added, deleted or its attributes are modified. The attributes that are tracked vary based on the type of data element, but abstractly include all meta-data describing the object in addition to full data contents. We can easily determine things like a configuration file changed (including ownership, permissions, size, contents) and can also enumerate a full change comparison between two versions. The same is true for other data types, like command output. Each unique version of data element is tracked, resulting in a set whose cardinality is typically more *interesting* (less noisy) relative to the number of snapshots (samples) in the client's population.

Cray's ability to leverage product change information to compare a customer system with itself at an earlier point in time, or compare multiple customer systems to each other can reduce the amount of time required to resolve a support request.

3) Prospective Capabilities

Prospective capabilities include:

- Validating the configuration of a product against a standard vendor baseline (e.g., auditing the configuration to ensure it is valid and up to date)
- Making better use of existing 'gold' (known good) snapshots and tracking longer configuration baseline histories
- SSA integration in the Cray's manufacturing process, and tracking a product's configuration through integration and into production
- Synchronization of SSA-provided product inventories with a manufacturing baseline for use in repair part logistics or similar applications
- Application of more sophisticated change analysis techniques that incorporate machine learning and larger product populations to determine if change is *interesting*

B. Product Triage Data Collection

1) Current Capabilities

A customer has the option to invoke a 'triage snapshot' on one of their Cray products. This type of SSA client use is typically used on-demand at the request of Cray customer service. SSA in this case is designed to capture a "checklist" of information that Cray may routinely need to make progress on a customer support concern.

A customer can optionally associate the triage snapshot with a support case identifier. When a snapshot is uploaded with such an association, the case is automatically updated with a snapshot reference. With this reference, Cray engineers assisting with the case then access the diagnostic information via a standard process.

2) Customer Value

a) The data that is collected in response to a support concern approaches a stable prescription and the resulting triage information is easier to collect and manage

With very few exceptions, Cray requests, and our customers are required to provide, diagnostic information. Delays in satisfying these requests can result in adding time to the resolution of a support request. This effect can be compounded if multiple parties are involved on either end of the dialogue.

Providing our customers with a simple interface to capture and upload triage information and providing our engineers with adequate information to assist the customer can reduce the amount of time required to resolve a support request.

3) Prospective Capabilities

Prospective capabilities include:

- Automated triage snapshot invocation (in response to a detected product issue)
- Automated triage data analysis

C. Product Health Monitoring

1) Current Capabilities

Product health information is collected on a relatively frequent interval. Currently, the health sub-system expresses system health in a *stateful* way, producing discrete health checks. Each health check is associated with a status code. The status codes are rudimentary (e.g., pass, fail, …). Each check also has an extended list of attributes, including the severity of a check if it failed, how the check relates to source data collected from the system (for provenance), a series of identifiers, a title and a description. The analysis required to produce the health checks takes place on the product (e.g., by the SSA client). The health check severities are mapped to Cray Service Case 'Priority' as a suggestion of how urgently an issue should be addressed.

If the SSA client detects a change in the set of health checks being monitoring, it will trigger an upload of the health snapshot where changes were identified. Once uploaded, SSA analyzes the health snapshot and generates discrete change events (e.g., pass to fail, fail to pass, severity change, …) associated with individual health checks. The history of a unique health check is also maintained with complete detail.

Cray employees can currently subscribe to receive e-mail notification of health events when they occur. After receiving an e-mail, a subscriber has the opportunity through a simple user interface to convert the health check into a support case – an entry point to remedial or proactive activities. Cray engineers are currently using and providing feedback on this system.

2) Customer Value

a) All products qualified for SSA have at least a basic set of health checks available

Monitoring the health of complex products is challenging. Managing the influx of automated alerts is also a challenge – with a poorly performing system having significant human impact (e.g., notification fatigue, false positives). Cray customers use a myriad of approaches toward monitoring their systems, yielding vastly different results. Cray is leveraging SSA to provide a standard set of supported health checks for a product, that see iterative improvement as new releases are made available. We currently limit the alerts to Cray employees, in large part because we are still refining the technology. We want to ensure alerts have a comfortably low false-positive rate and that notifications are actionable. We further want to ensure we have an appropriate level of integration with other systems at Cray. Product health monitoring allows Cray to promptly respond to product issues, in some cases proactively – ideally leading to improved system availability.

3) Prospective Capabilities

Prospective capabilities include:

- Greater configurability of health sub-system on client
- Extended stateful health check attributes (better express complex systems)
- Product event processing and analysis (to complement stateful health checks)
- Secondary health analysis by SSA centralized analytics
- Automatically shipping/staging repair parts based on an identified failure

IV. SSA Program Metrics

In this section, we present select metrics describing the SSA program, including SSA account activations, SSA product population (uploading systems), snapshot population statistics, related service case statistics and SSA health event statistics.

A. Customer Adoption

As of April 11, 2017, there were 16 total activated and active customers using SSA (uploading data). Fig. 2 shows the rate of customer activations and cumulative activations since July 2015.

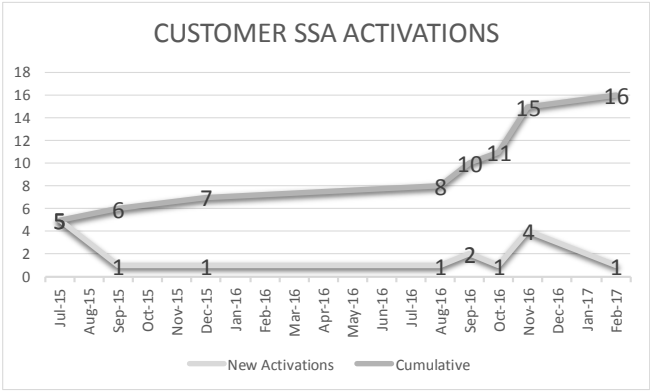

Figure 2. Customer SSA Account Activations

B. Reporting Product Population

As of April 14, 2017, there were 42 actively reporting customer systems. The Cray® XC™ product makes up the majority. Fig. 3 provides a breakdown by product family.

Please note that, in Fig. 3, Cray® XK™ systems are grouped with the Cray® XE™ Family.

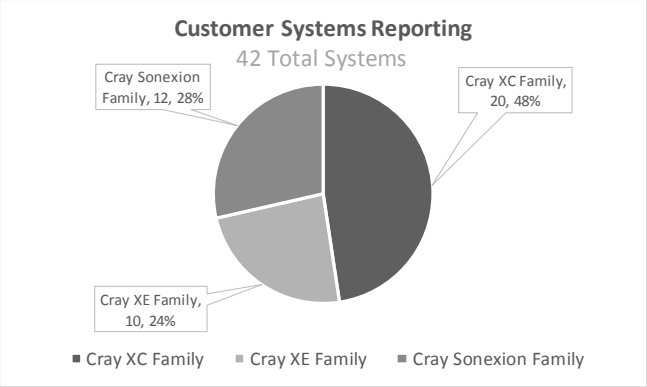

Figure 3. Products Reporting by Product Family

C. *Snapshot Statistics*

Over 21,000 snapshots have been processed by the SSA back-end system. There are three standard types of snapshots, or snapshot channels – daily (configuration baseline), health and triage. Configuration baseline snapshots are normally uploaded for each reporting product, daily (some systems report multiple snapshots for different logical aspects of the system). Health snapshots are only uploaded based on change detected during client processing. Triage snapshots are captured and uploaded on demand. For each reporting *device*, a maximum of 31 snapshots is retained based on current policy. After the window is exceeded, older snapshots *expire*. Some data is persisted outside of the snapshot system (e.g., health events and other analytics products).

Fig. 4 shows the population on April 12, 2017 of daily snapshots based on collection date for the previous six months (1,233 total). Fig. 5 and 6 show health (1,019 total) and triage (157 total) snapshots for the same period, respectively.

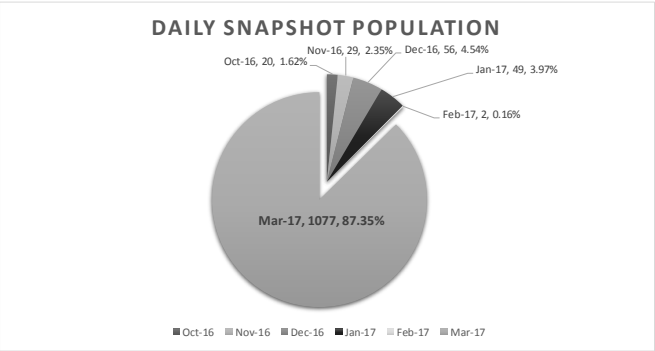

Figure 4. Customer reported daily snapshots as of April 12, 2017

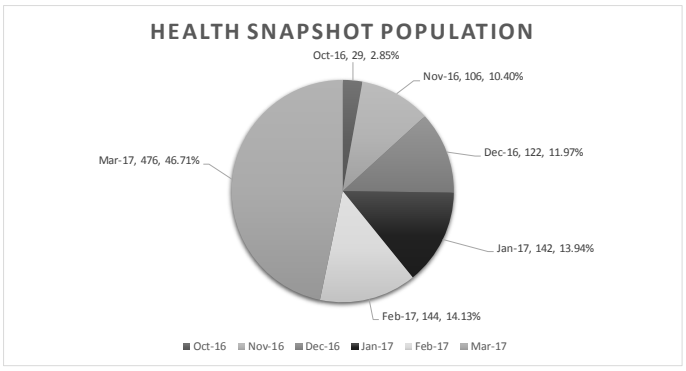

Figure 5. Customer reported health snapshots as of April 12, 2017

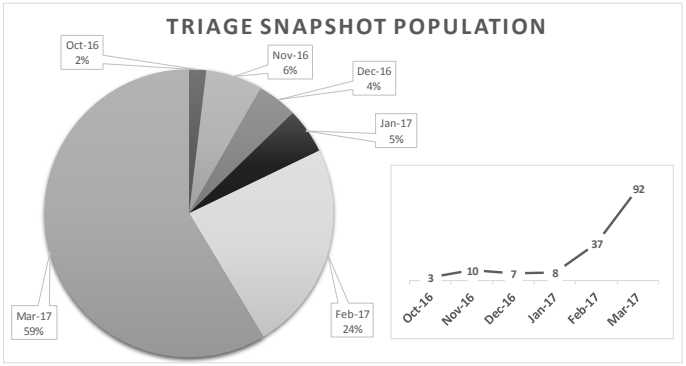

Figure 6. Customer reported triage snapshots as of April 12, 2017

D. *Service Case Interface Statistics*

Fig. 7 shows the number of customer support cases created directly from SSA health notifications, for the last six months. Fig. 8 shows the number of triage snapshots automatically associated with customer support cases.

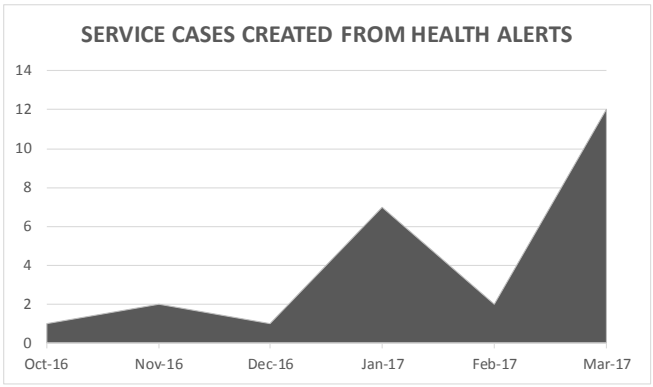

Figure 7. SSA Health Notifications Converted to Service Cases

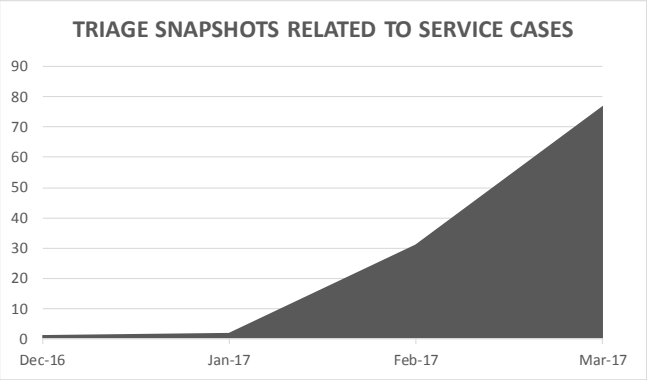

Figure 8. Triage Snapshots Associated with Service Cases

V. FUTURE WORK

There are many areas of opportunity for future work within the SSA program, some of which we covered in section III. In this section, we provide our 2017 roadmap and other focal areas for future work.

For our 2017 roadmap, we are tracking the following activity:

- Continued incremental bug fixes and features on the Cray® Sonexion™, XC™ and XE6™ (including XK) products
- Initial support for SMW 8 and CLE 6
- Initial support for the Cray® Urika GX™ product
- Initial integration of SSA inventory and status information into Cray customer Support asset tracking systems
- Introduction of a product event-analysis processing environment and health processing improvements

Strategically, we are collaborating with multiple teams at Cray – working toward a unified platform strategy that fully leverages our expertise and product portfolio.

Perhaps most importantly, we are firmly committed to working closely with the Cray customer support teams and Cray customers. The feedback and requirements provided from these communities is undeniably critical to the future of SSA. We will endeavor to align and refine our operational metrics as a measure of how well we are meeting the needs of our community.

VI. CONCLUSION

SSA has been generally available since June 2015. Working with our community, SSA has seen steady improvement over nine client releases – including increasing the number of supported platforms. We have also made considerable improvements to our back-end analytics and automation environment, with over 21,000 snapshots processed in production.

We believe the customer value proposition for SSA is firmly established. We have endeavored to introduce and support these beliefs in this paper. We are keenly interested in obtaining customer feedback to help guide the future of this capability.

SSA customer adoption is gaining momentum. Cray's investment in SSA is increasing. Our support teams are working closely with customers to help them get started with SSA. Adoption and use by our customers and support teams has led to several wins in the field. We are also working to ensure we are driving improvements based on missed opportunities. We would like to thank customers who are using and helping to improve SSA. For customers not yet using SSA, a reminder that SSA is currently free of charge to those with a Cray service contract and an SSA supported platform.

ACKNOWLEDGEMENT

We would like to thank the Cray User Group (CUG) for the opportunity to present this paper at the 2017 conference. We would also like to again thank customers that are using SSA – with a special thanks to those who have submitted requests for enhancement or bugs (we are listening!). We would also like to thank our colleagues at Cray for their support of SSA. Lastly, we would like to thank our technology partners – notably ISOdx Solutions [8] – for their continued support and dedication to the success of SSA.

REFERENCES

[1] Cray Customer Portal (CrayPort) (2017, April 11). [Online]. Available: https://crayport.cray.com/

[2] Definition of "telemetry". (2017, April 11) . [Online]. Available: https://www.merriam-webster.com/dictionary/telemetry

[3] Definition of "telemetering". (2017, April 11) . [Online]. Available: https://www.merriam-webster.com/dictionary/telemetering

[4] J. Duckworth, Cray System Snapshot Analyzer (SSA) White Paper. Seattle, WA: Cray, Inc, 2015. [Online]. Available: http://www.cray.com/sites/default/files/resources/Cray-SSA-White-Paper.pdf

[5] Cray Customer Portal (CrayPort) KB 4546 (2017, April 11). [Online]. Available: https://portal.cray.com/Support/apex/ka_how_to?Id=kA33300000008ToACAU

[6] Cray System Snapshot Analyzer (SSA) End User License Aggrement (EULA) (2017, April 11). [Online]. Available: http://www.cray.com/sites/default/files/resources/Cray-SSA-EULA.pdf

[7] Cray customer service: Cray System Snapshot Analyzer (SSA) (2017, April 11). [Online]. Available: http://www.cray.com/support/cray-system-snapshot-analyzer

[8] ISOdx Solutions (2017, April 13). [Online]. Available: http://www.isodxsolutions.com