



CRAY



SSA, ClusterStor Call-home Service Actions, and an Introduction to the Cray Central Telemetry and Triage Services (C2TS)

CUG 2018

Jeremy Duckworth, Cray Inc.



Overview

- **Target Audience**
 - System Owners, Architects, Analysts, Administrators, and Operators
- **Touch on the Benefit of Cray's Call-home Platform**
 - Improve customer system availability through proactive support
 - Lower customer and Cray operational costs through support process automation
- **Outline**
 - Provide a Cray SSA operational update
 - Provide a brief history of Cray ClusterStor and SSA Integration
 - Describe and define ClusterStor Reliability, Availability, and Serviceability (RAS) telemetry types
 - Preview Cray Call-home Service Action Features for ClusterStor
 - Discuss Support Data Capture on ClusterStor, future use of SSA
 - Preview Next Generation Call-home Architecture: Cray Central Telemetry and Triage Services (C2TS)
- **Summary**
- **Q&A**

Cray SSA Operational Update

- **Call-home production use started June 2015**
- **Active at 21 customer sites, 78 systems reporting**
- **Processed ~ 80K call-home snapshots, 48K product health events**
- **Logged 414 proactive and triage case association events**
- **Field Notice (FN) 6122 contains client release details**

Brief History: ClusterStor and SSA Integration



- **SSA was qualified for Sonexion in September 2016**
- **In June 2017, SSA became the official RAS call-home transport on Sonexion**
- **Cray acquired the ClusterStor Product in 2017**
- **In March 2018, the first unified SSA Client released for Sonexion and ClusterStor**
- **Cray recommends all ClusterStor and Sonexion customers use SSA for call-home enablement**

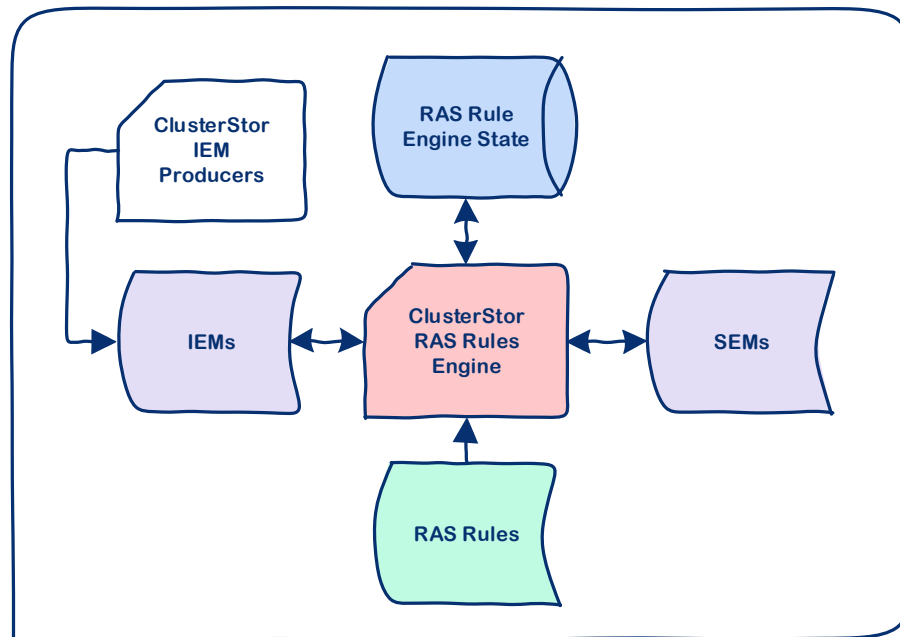
ClusterStor RAS Telemetry Types

- **Interesting Event Messages (IEMs)**
 - Structured, highly contextual diagnostic events
 - Cover a broad range of topics
 - Lustre, disk and enclosure services, systems management, ...
- **Service Event Messages (SEMs)**
 - Signal when ‘service is needed’, clear when service is completed
 - Generated through analysis of IEMs and related context
 - Translate to service event e-mails and alerts local to product
- **Machine Reportable Product Data (MRPD)**
 - Infrequent (daily) snapshots of system inventory, status, and, diagnostic information

ClusterStor Standalone RAS Rules Engine

● Overview

- Analyzes stream of RAS telemetry events
- Orchestrates Response Actions
 - (internal) Updates rule engine state and context
 - Creates SEMs and thus local service alerts



ClusterStor RAS Telemetry: Examples

IEM

SEM

```
{
  "local_time": "Thu, 03 May 2018 14:25:01 CST",
  "message_type": "iem",
  "messages": [
    {
      "data": {
        "confirmed_time": 0.0,
        "creation_time": 1525375281.394192,
        "event_code": "001001001",
        "event_data": {
          "dcs_timestamp": 1525375281,
          "enclosure_serial_number": "REDACTED-ENCL-SERIAL",
          "fru_serial_number": "REDACTED-DRIVE-SERIAL",
          "index": "13",
          "status": "failed",
          "type": "disk"
        },
        "id": "REDACTED-ENCL-SERIAL:disk:13",
        "resolved_time": 0.0,
        "state": "EVALUATION",
        "version": 2
      }
    }
  ]
}
```

Disk drive failure under evaluation

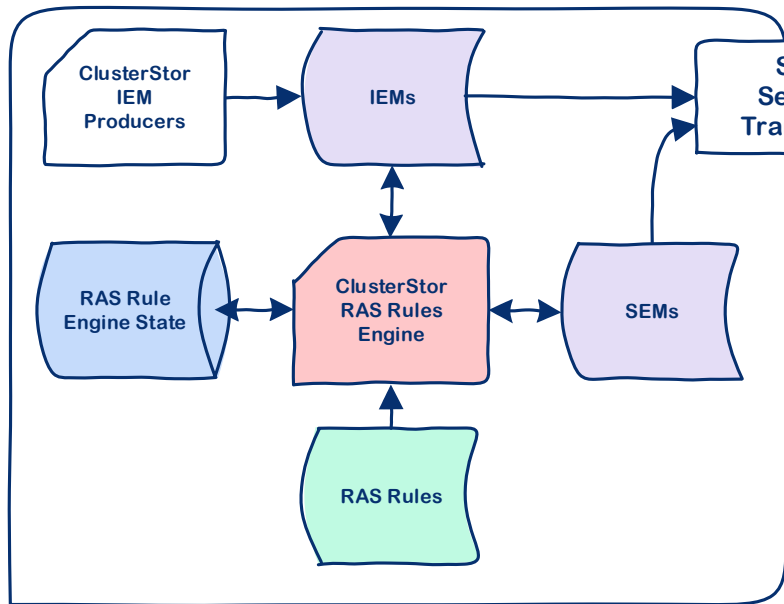
```
{
  "local_time": "Thu, 03 May 2018 14:39:45 CST",
  "message_type": "sem",
  "messages": [
    {
      "completion_time": 0.0,
      "confirmed_time": 1525376181.398169,
      "creation_time": 1525376181.3996589,
      "dcs_timestamp": 1525375281,
      "event_code": "002005001",
      "event_description": "Disk drive needs replacement",
      "fru": {
        "disk_installed": "1",
        "dm_report_t10": "11110111100",
        "firmware": "EOG5",
        "manufacturer": "SEAGATE",
        "part_number": "ST6000NM0034",
        "sect": "512",
        "serial_number": "REDACTED-DRIVE-SERIAL",
        "status": "Failed",
        "t10_enabled": true
      },
      "location": {
        "enclosure_location": "16U",
        "enclosure_model": "5U84",
        "index": "13",
        "rack": "R1C4"
      }
    }
  ]
}
```

Disk drive failure confirmed, context added

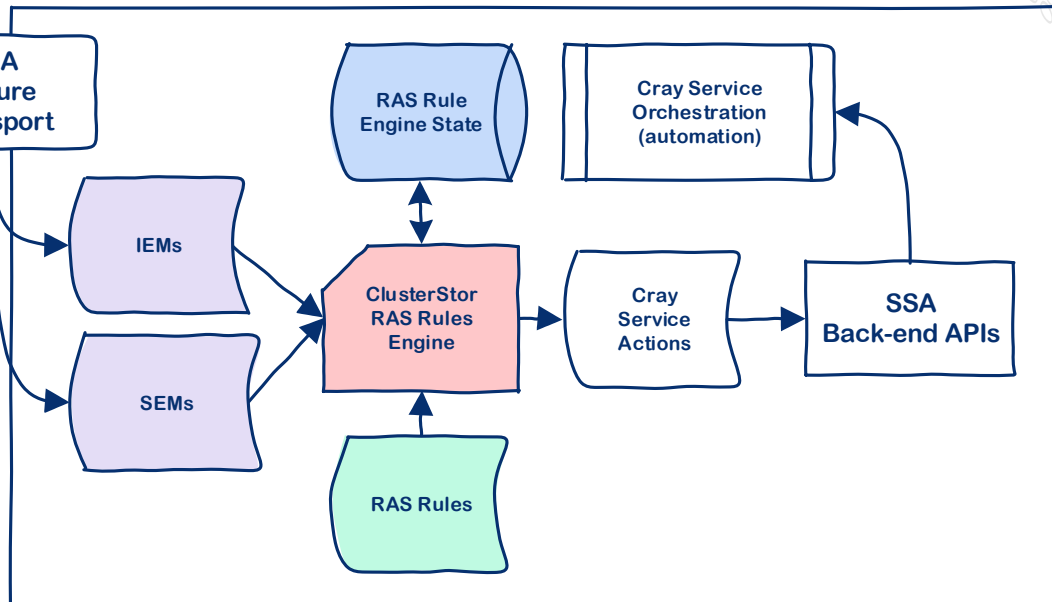
ClusterStor RAS Call-home



ClusterStor RAS Standalone



ClusterStor RAS Call-Home



COMPUTE

STORE

ANALYZE

ClusterStor Call-home Service Actions

- **Automated Field Replaceable Unit (FRU) Repair Part Orders**
 - On FRU Failure, create and populate a service case, create a corresponding repair part request, and ship part to the customer
- **Proactive Service Recommendations**
 - Utilizes up-to-date knowledge of product releases and defects
 - Rules can be vetted via call-home, updated in future standalone engine updates
 - Types of recommendations Cray is investigating
 - ‘you might want to upgrade’ or ‘the following issues exist for your release...’
 - Faster/better detection for subtle or complex product failures
 - For example, excessive errors on a data path, possible enclosure failures

ClusterStor Support Data Capture

- **ClusterStor currently has a Support Bundle Feature**
 - Collects system status, logs, ...
 - Data used by support and product teams for root cause analysis of Cray internal and customer issues
- **SSA currently supports ‘triage’ snapshots**
 - Goal is very similar to that of Support Bundle
- **Plan is to reach feature parity in SSA, deprecate Support Bundle over time**
 - Provides streamlined tooling, improved test, and standard interfaces across Cray Products
 - SSA triage snapshots are integrated into Cray customer service systems

- **Cray Central Telemetry and Triage Services (C2TS)**
 - Evolution of Cray's Call-home platform
 - Design informed from lessons learned after five years of experience with SSA
 - Requirements driven by
 - Parity-or-better SSA features
 - Cray's next generation supercomputing platform
 - Better alignment with ClusterStor, ClusterStor RAS

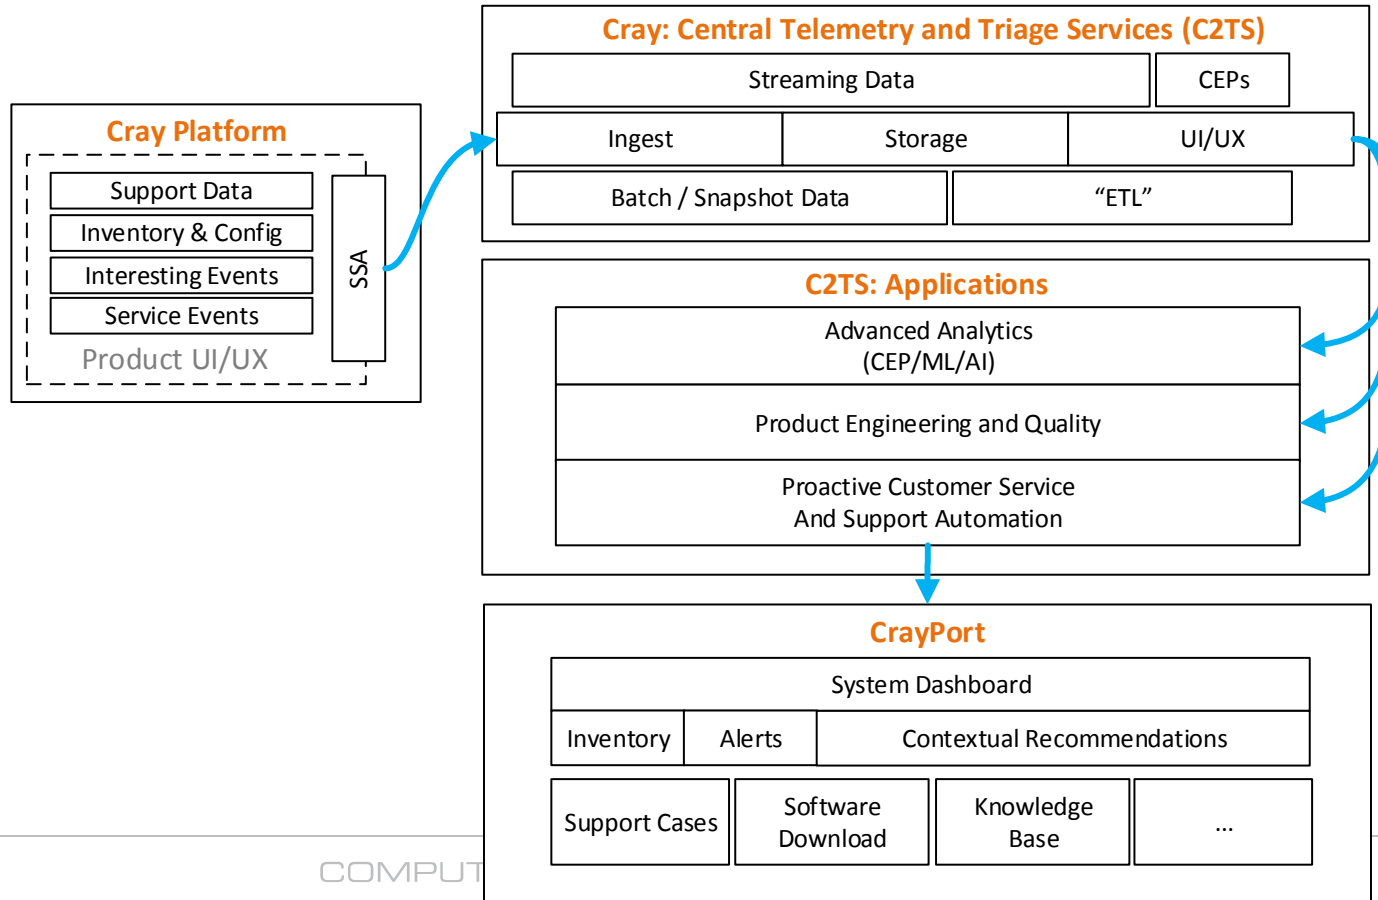

SSA's Relationship to C2TS

- **C2TS new 'back-end' for SSA and future call-home services**
- **Major usability enhancements in C2TS for SSA**
- **Adds cleaner delineation between data lake and applications**
 - Embraces the concept of other application stacks leveraging call-home data lake services

C2TS Motivations

- Availability
 - Cray supports customers 24x7x365
- Scalability
 - Number of systems, data size and data volume increasing
 - Time to solution requirements decreasing
- Usability
 - Cray product support and engineering teams, need to focus on data, not tooling
 - Streaming data is better processing model for some use cases
 - Get more call-home data in front of customers, versus behind-the-curtain
- Programmability
 - Less esoteric APIs, data schemas
- Serviceability
 - Easy to diagnose, fix issues, understand behavior

C2TS Reference Architecture



COMPUT

Summary

- **Cray ClusterStor is designed to enable proactive and automated product support via call-home telemetry**
- **SSA provides a secure call-home transport and a product support automation capability across Cray platforms**
- **Cray plans to streamline support data collection on ClusterStor, using SSA**
- **C2TS is Cray's next generation call-home platform, designed with Cray's next generation system platforms in mind and lessons learned from five years of experience operating SSA**

Legal Disclaimer

Information in this document is provided in connection with Cray Inc. products. No license, express or implied, to any intellectual property rights is granted by this document.

Cray Inc. may make changes to specifications and product descriptions at any time, without notice.

All products, dates and figures specified are preliminary based on current expectations, and are subject to change without notice.

Cray hardware and software products may contain design defects or errors known as errata, which may cause the product to deviate from published specifications. Current characterized errata are available on request.

Cray uses codenames internally to identify products that are in development and not yet publicly announced for release. Customers and other third parties are not authorized by Cray Inc. to use codenames in advertising, promotion or marketing and any use of Cray Inc. internal codenames is at the sole risk of the user.

Performance tests and ratings are measured using specific systems and/or components and reflect the approximate performance of Cray Inc. products as measured by those tests. Any difference in system hardware or software design or configuration may affect actual performance.

The following are trademarks of Cray Inc. and are registered in the United States and other countries: CRAY and design, SONEXION, URIKA and YARCDATA. The following are trademarks of Cray Inc.: CHAPEL, CLUSTER CONNECT, CLUSTERSTOR, CRAYDOC, CRAYPAT, CRAYPORT, DATAWARP, ECOPHLEX, LIBSCI, NODEKARE, REVEAL. The following system family marks, and associated model number marks, are trademarks of Cray Inc.: CS, CX, XC, XE, XK, XMT and XT. The registered trademark LINUX is used pursuant to a sublicense from LMI, the exclusive licensee of Linus Torvalds, owner of the mark on a worldwide basis. Other trademarks used on this website are the property of their respective owners.

Q&A

A scenic view of a European city, likely Copenhagen, with colorful buildings and a prominent church spire, reflected in the water. The sky is clear and blue, and the water is calm, showing clear reflections of the buildings and the sky.

Jeremy Duckworth
jeremyd@cray.com

Tim Morneau
tmorneau@cray.com