Monitoring HPC Security at LLNL

4th NIST HPC Security Workshop

Ian Lee
HPC Security Architect

2024-05-20

LLNL-PRES-864588

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

User Centric View of LC

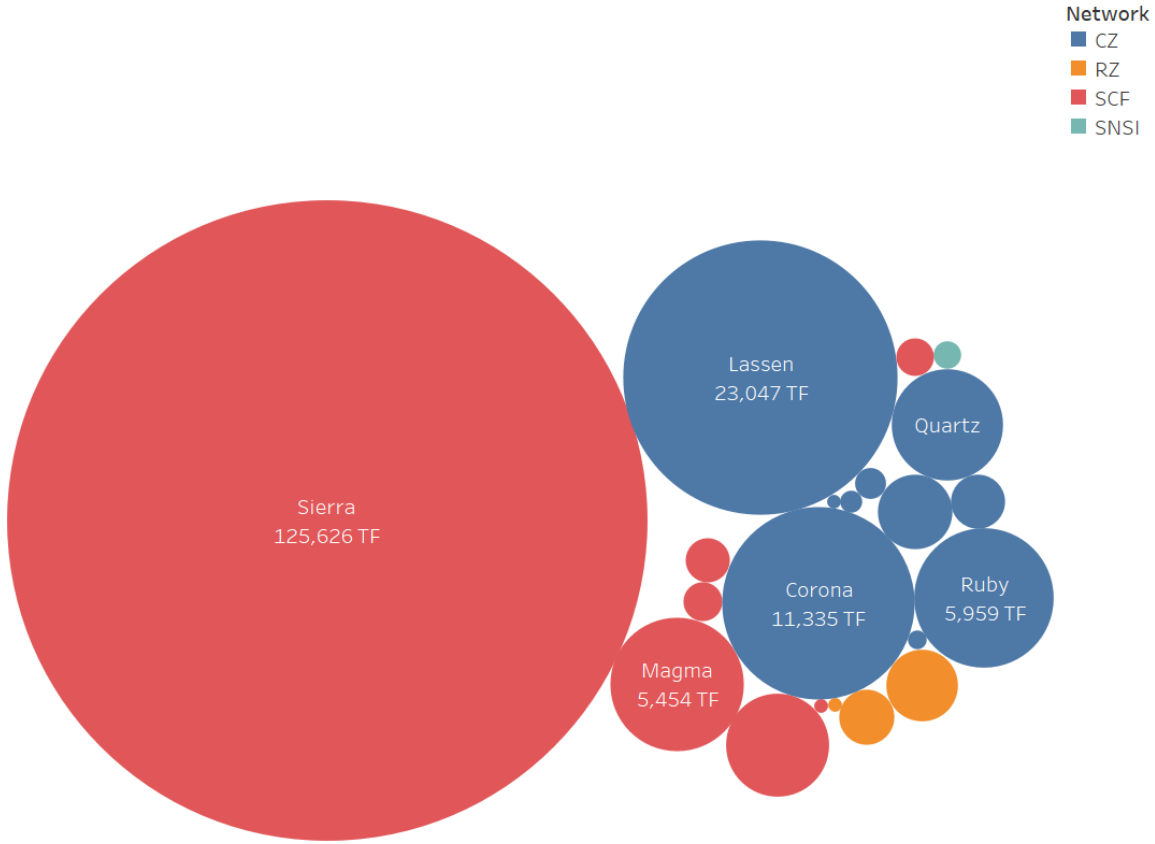

- Each bubble represents a cluster
- Size reflects **theoretical peak performance**, color indicates **computing zone**
- Only large user-facing compute clusters are represented

HPC Zones – User View

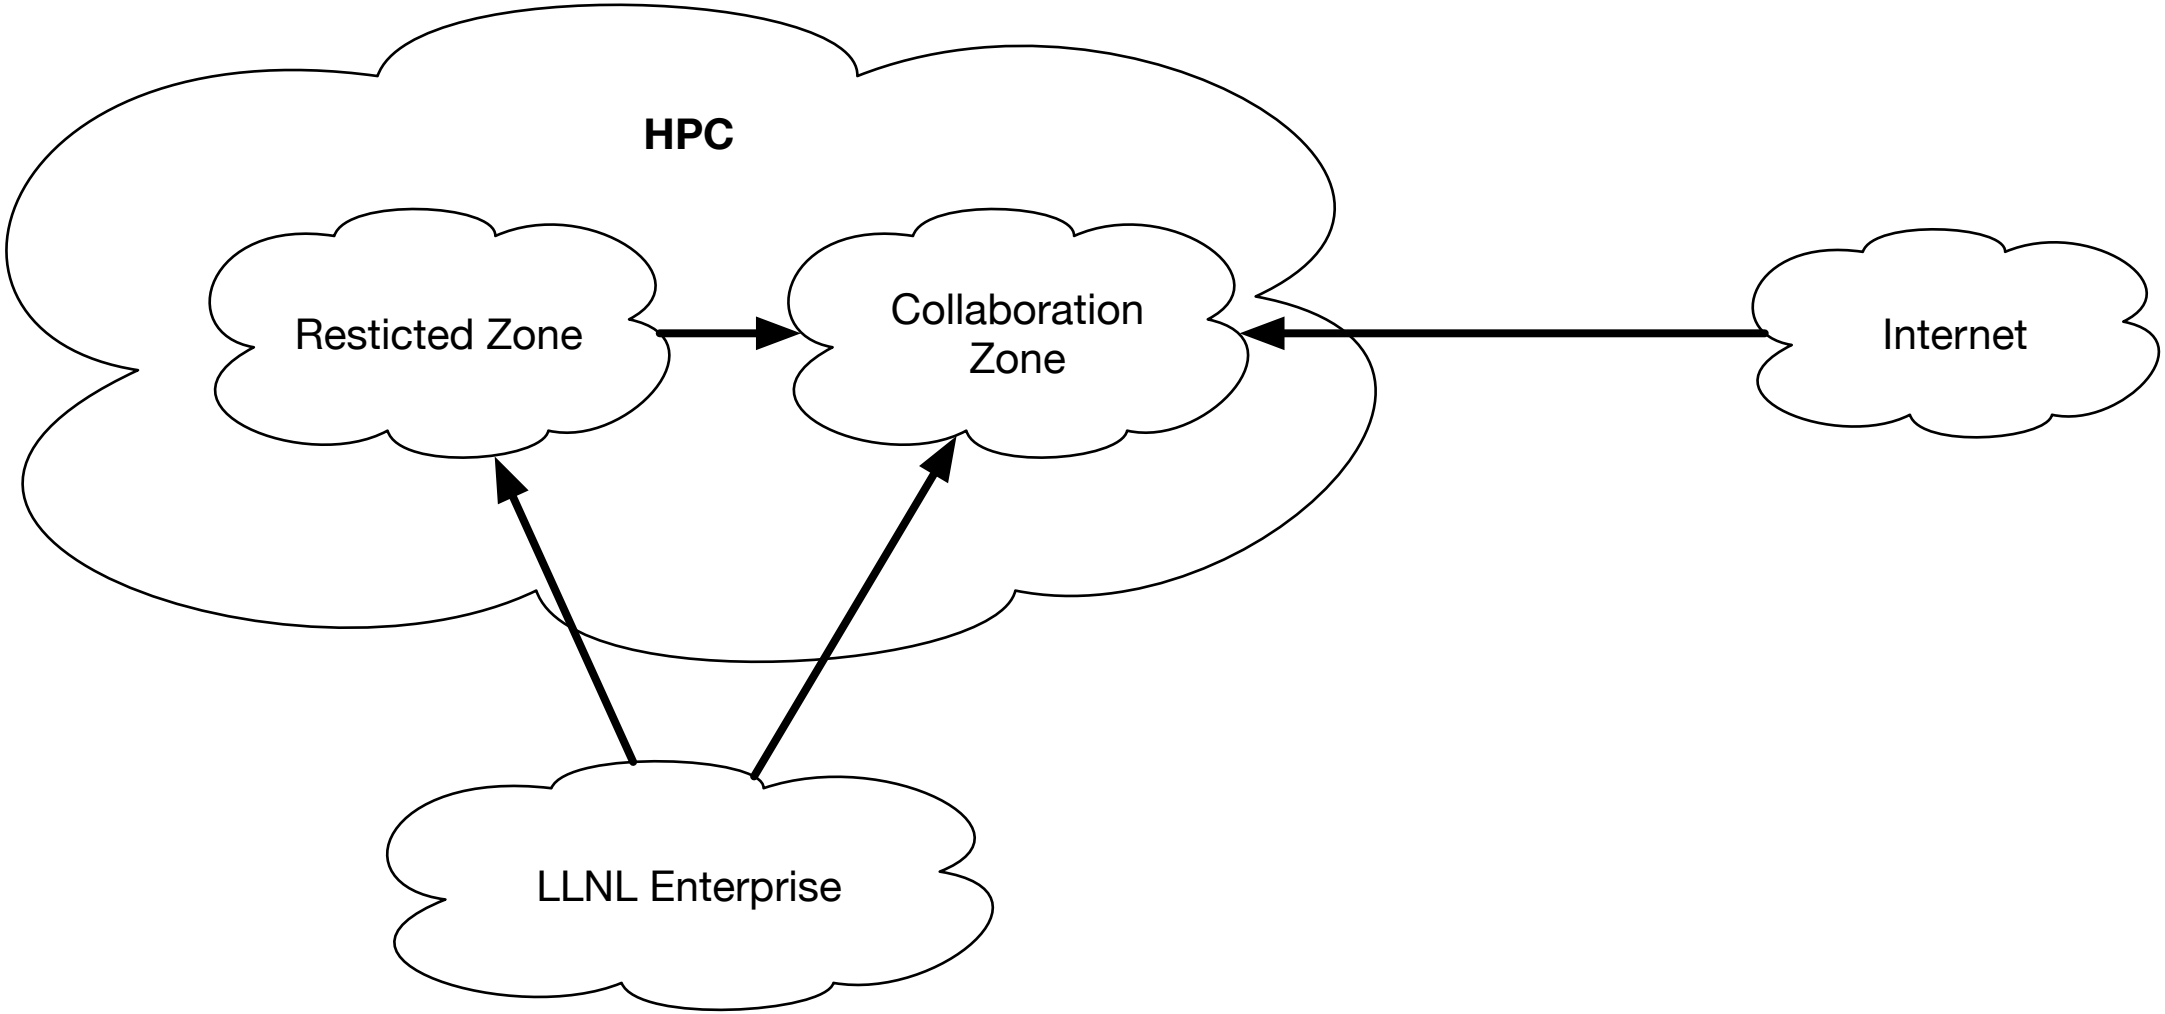

More Complete View of LC

RZ

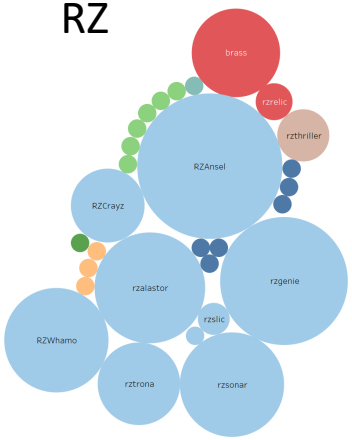

SCF

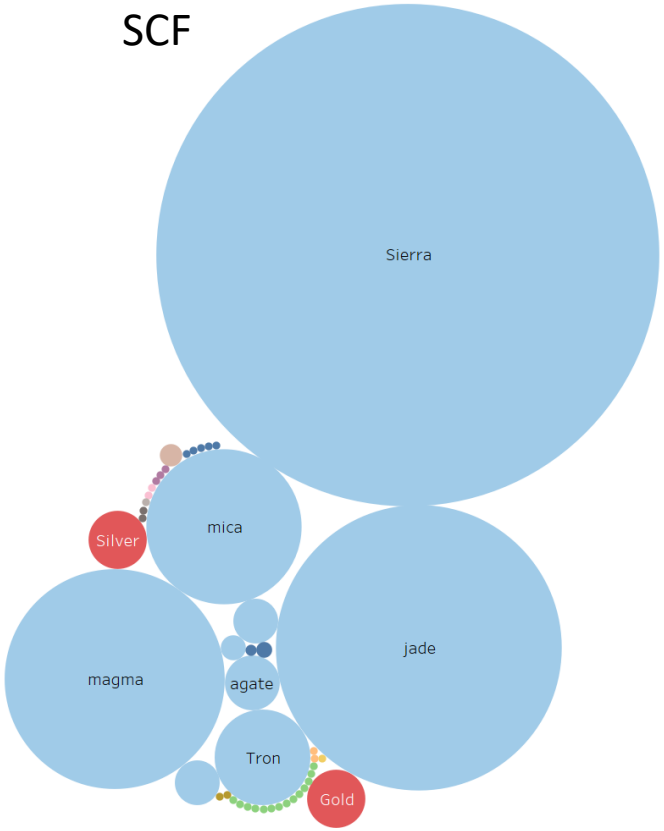

CZ

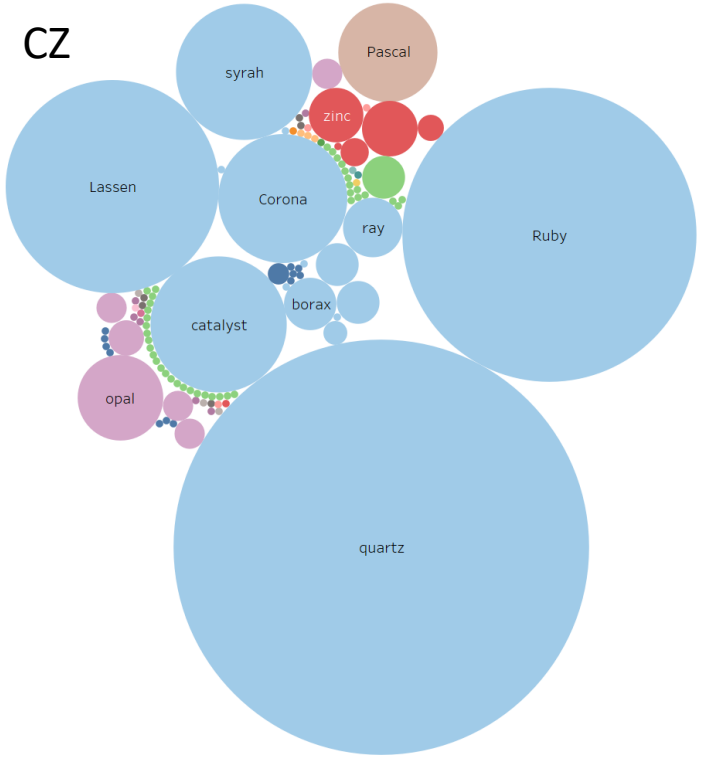

Type
- Center-Wide NFS
- Compute
- Enclave Server
- File Interchange
- File System Servers OCF
- Infrastructure
- Kerberos/LDAP
- LC Desktops
- License Servers
- Lustre Storage Clusters
- NAS Mgmt
- Network Management
- Network Monitoring
- Secure Web
- Servers
- Storage
- Testbed
- Visualization
- Web

Green

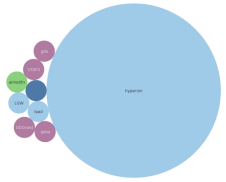

SNSI

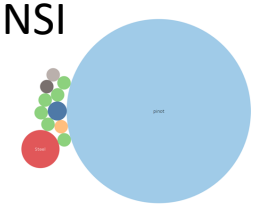

- Each bubble represents a managed cluster or system
- Size reflects **number of nodes**, color indicates **type** of system
- All production systems are shown, including non-user-facing systems

HPC Zones – Wider View

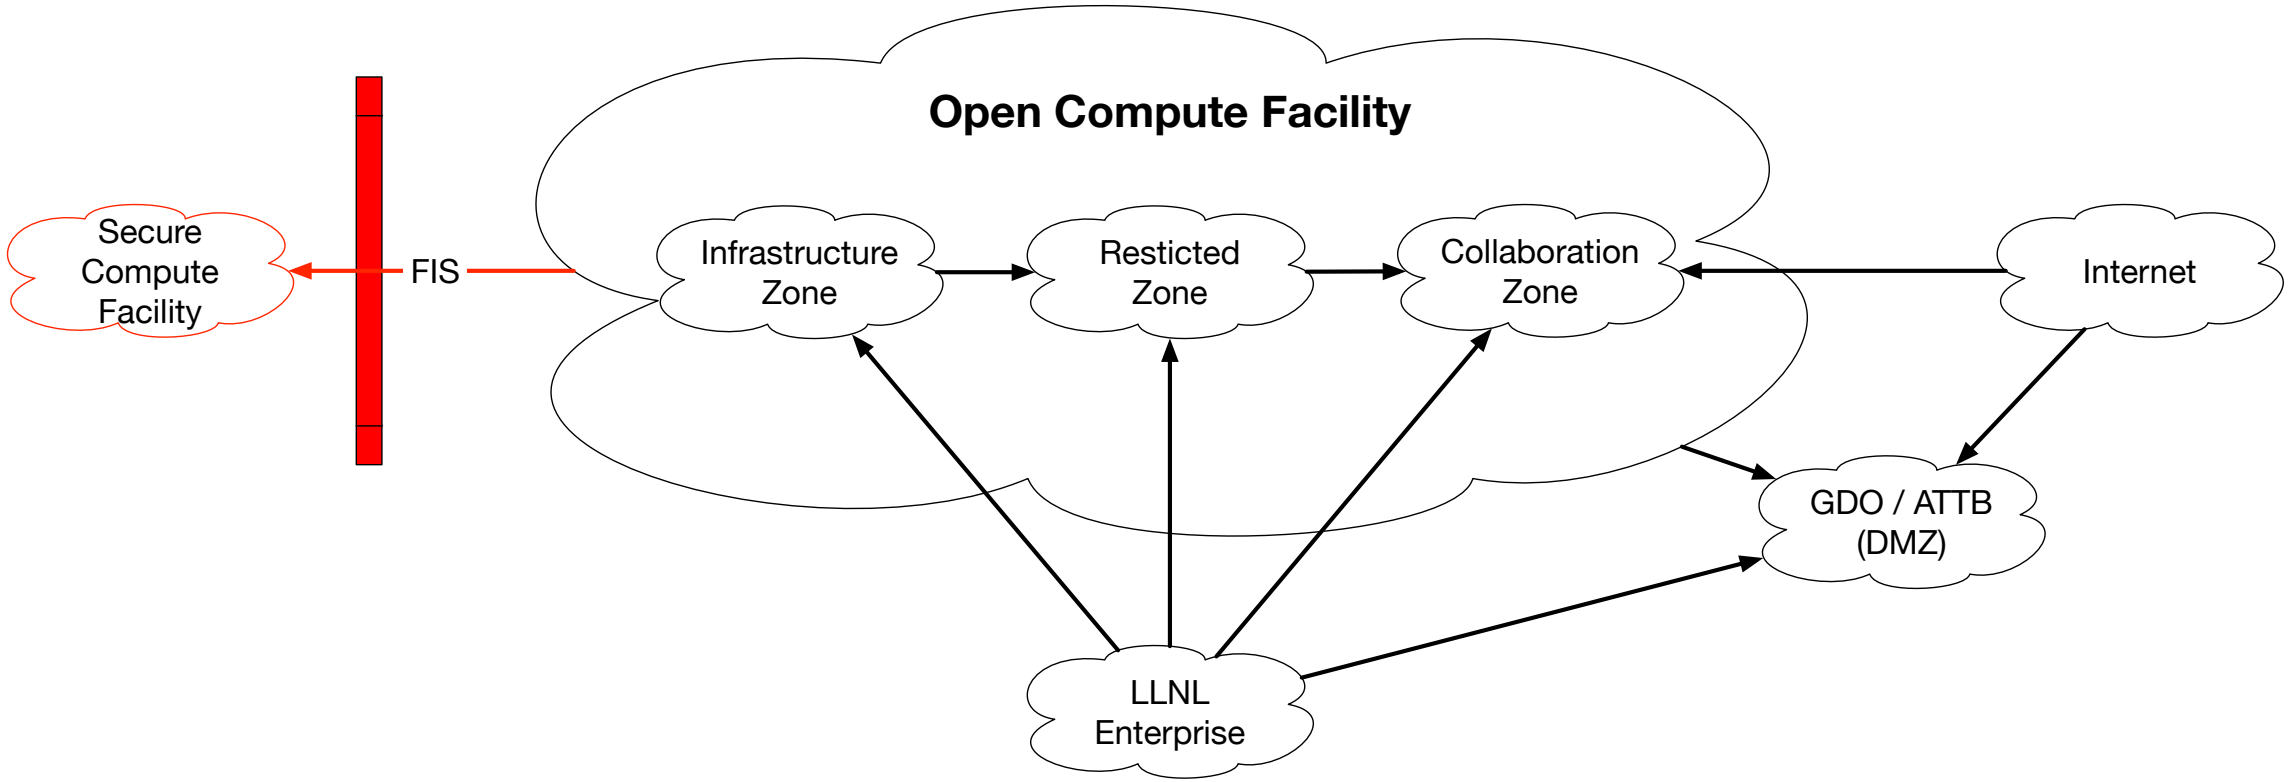

Where We're Heading

El Capitan

El Capitan vs the Rest

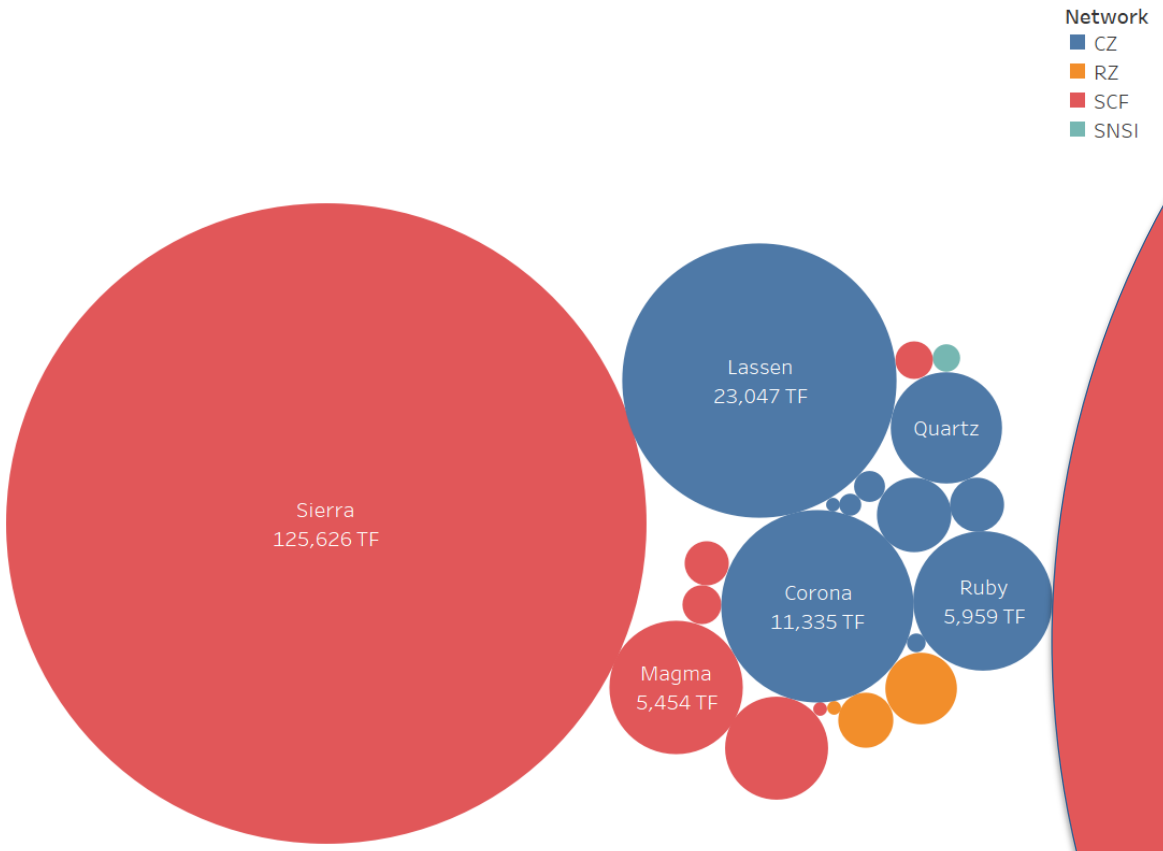

- Each bubble represents a cluster
- Size reflects **theoretical peak performance**, color indicates **compu**
- Only large user-facing compute clusters are represented

El Capitan
~ 2 EF (~ 2,000,000 TF)

system.syslog per hour

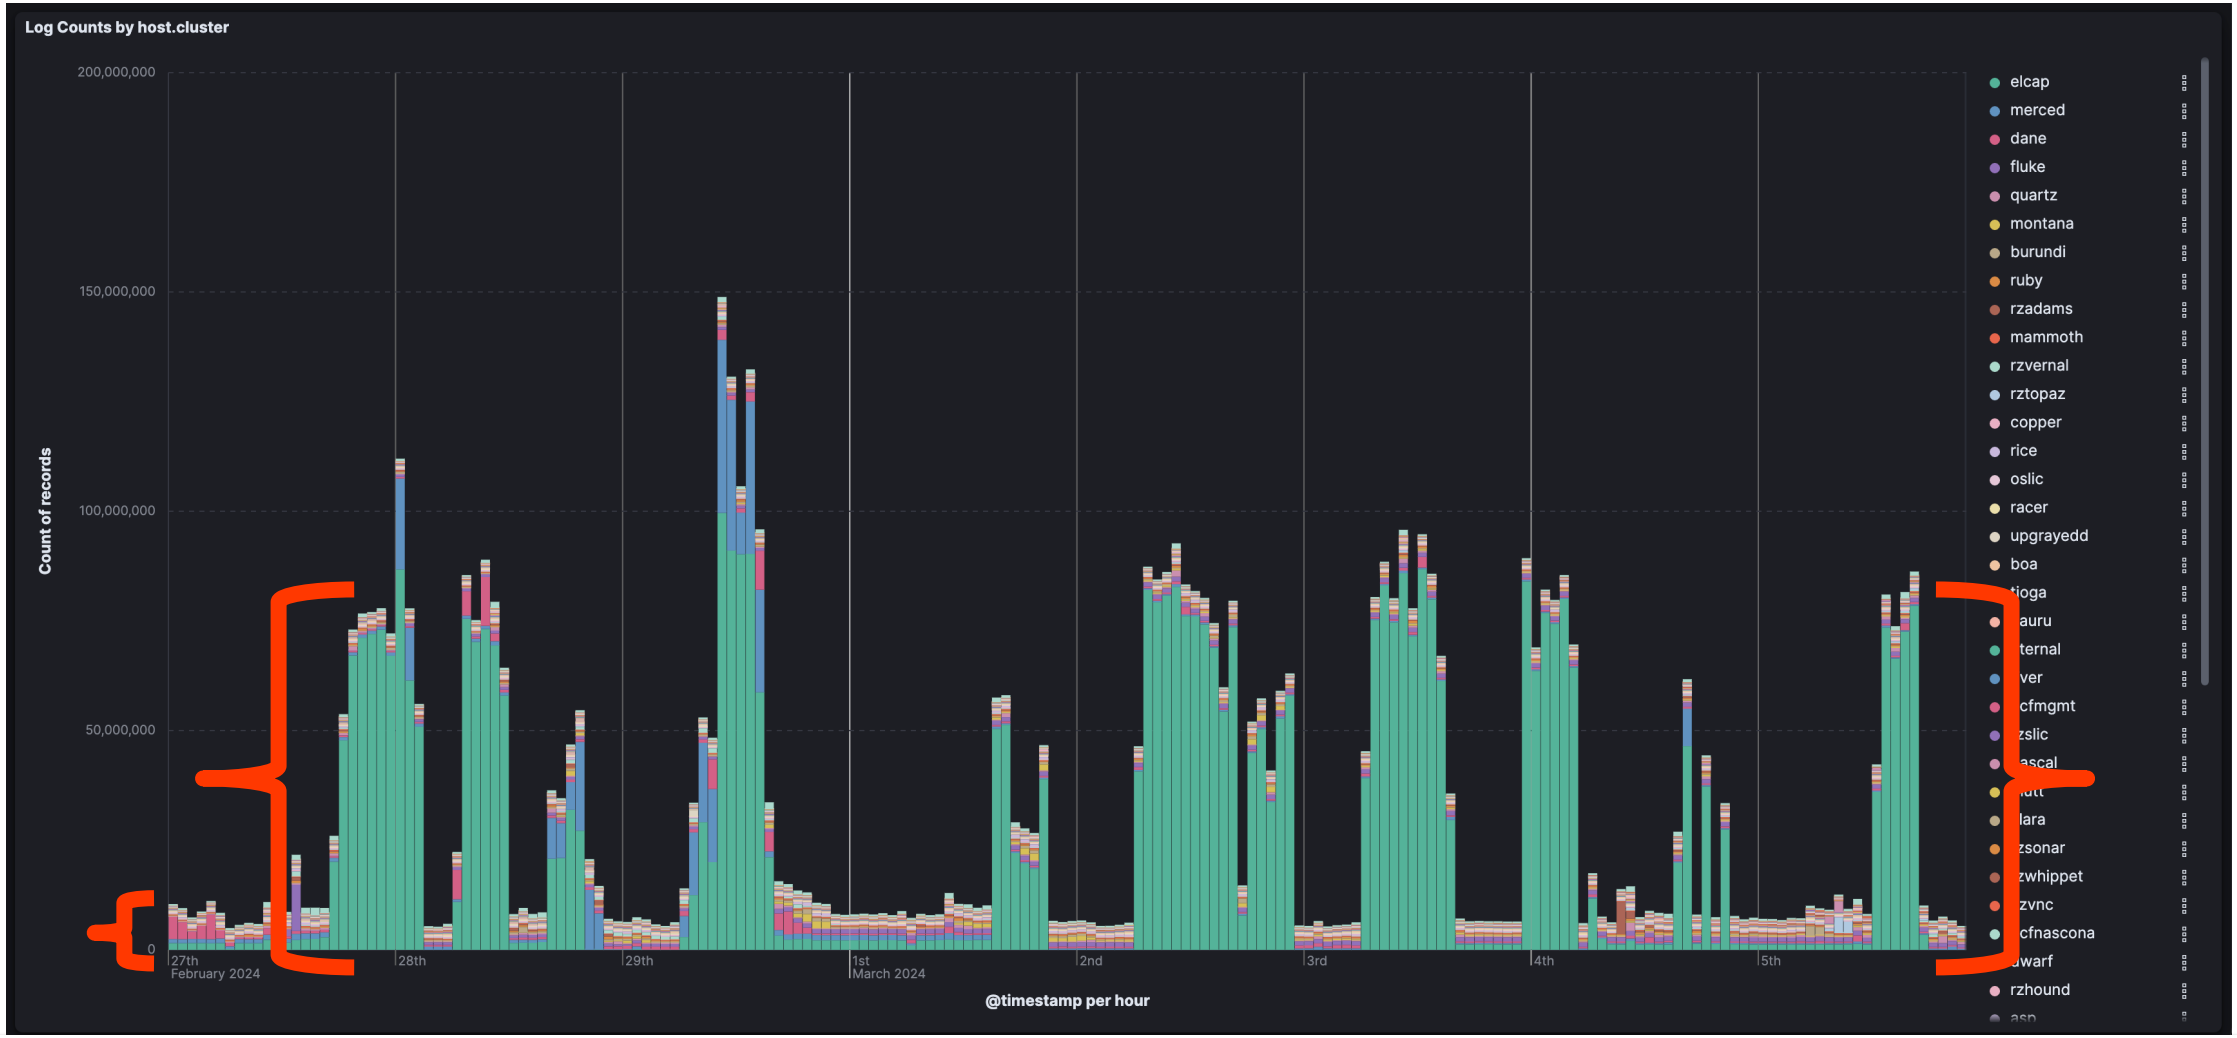

Logging Infrastructure

Logging Architecture

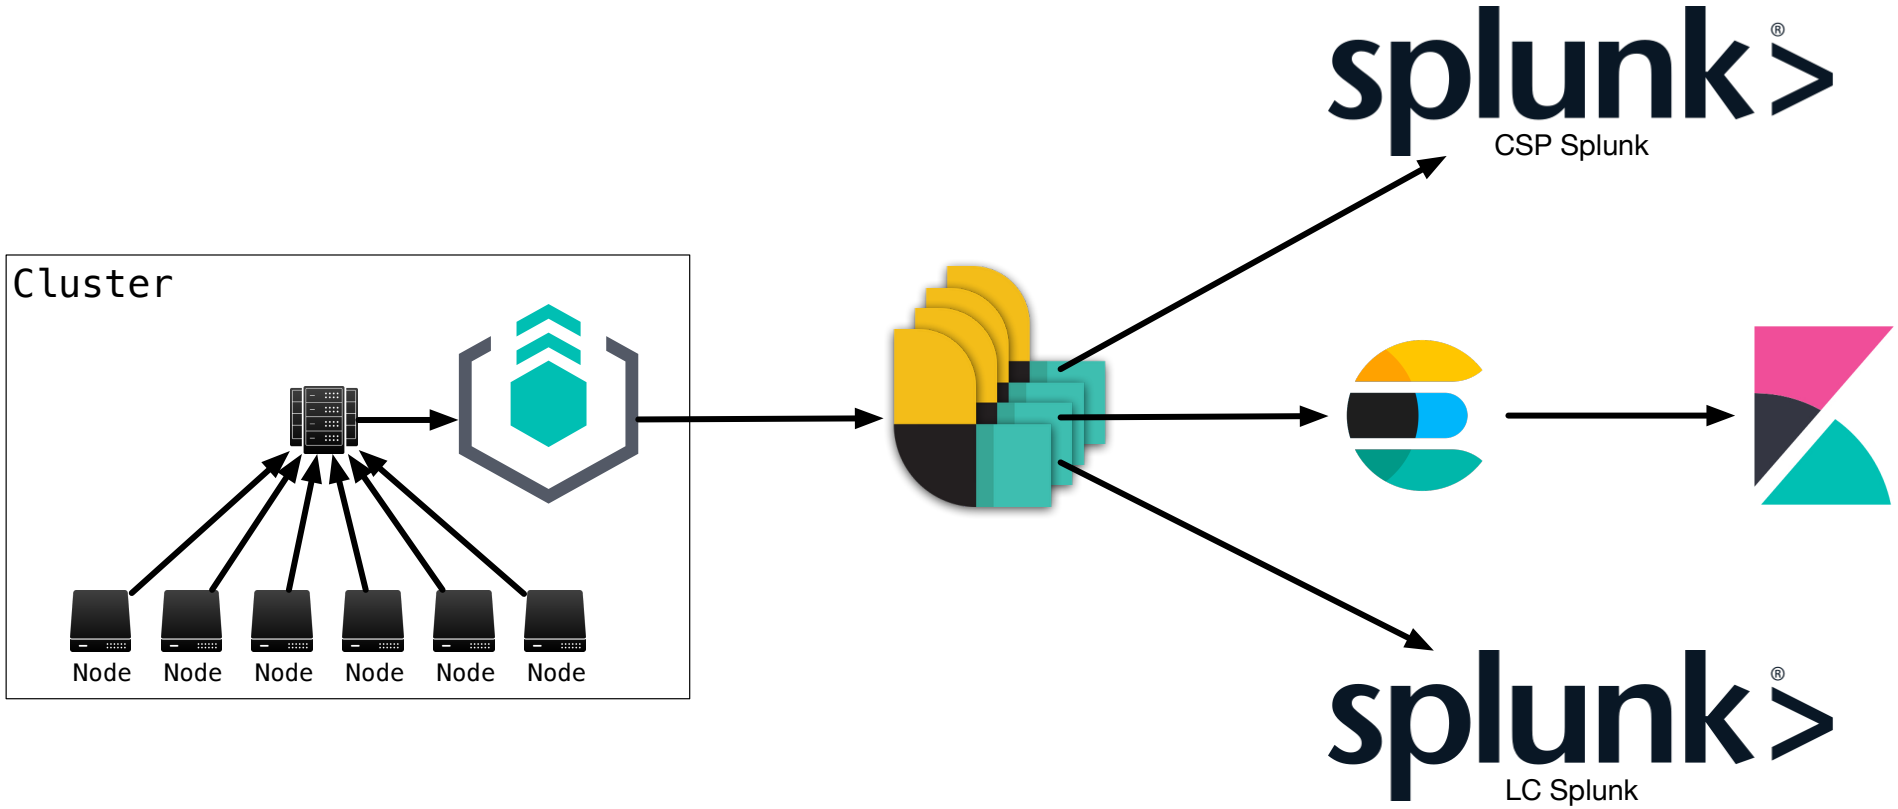

Current Service to Hardware Allocations

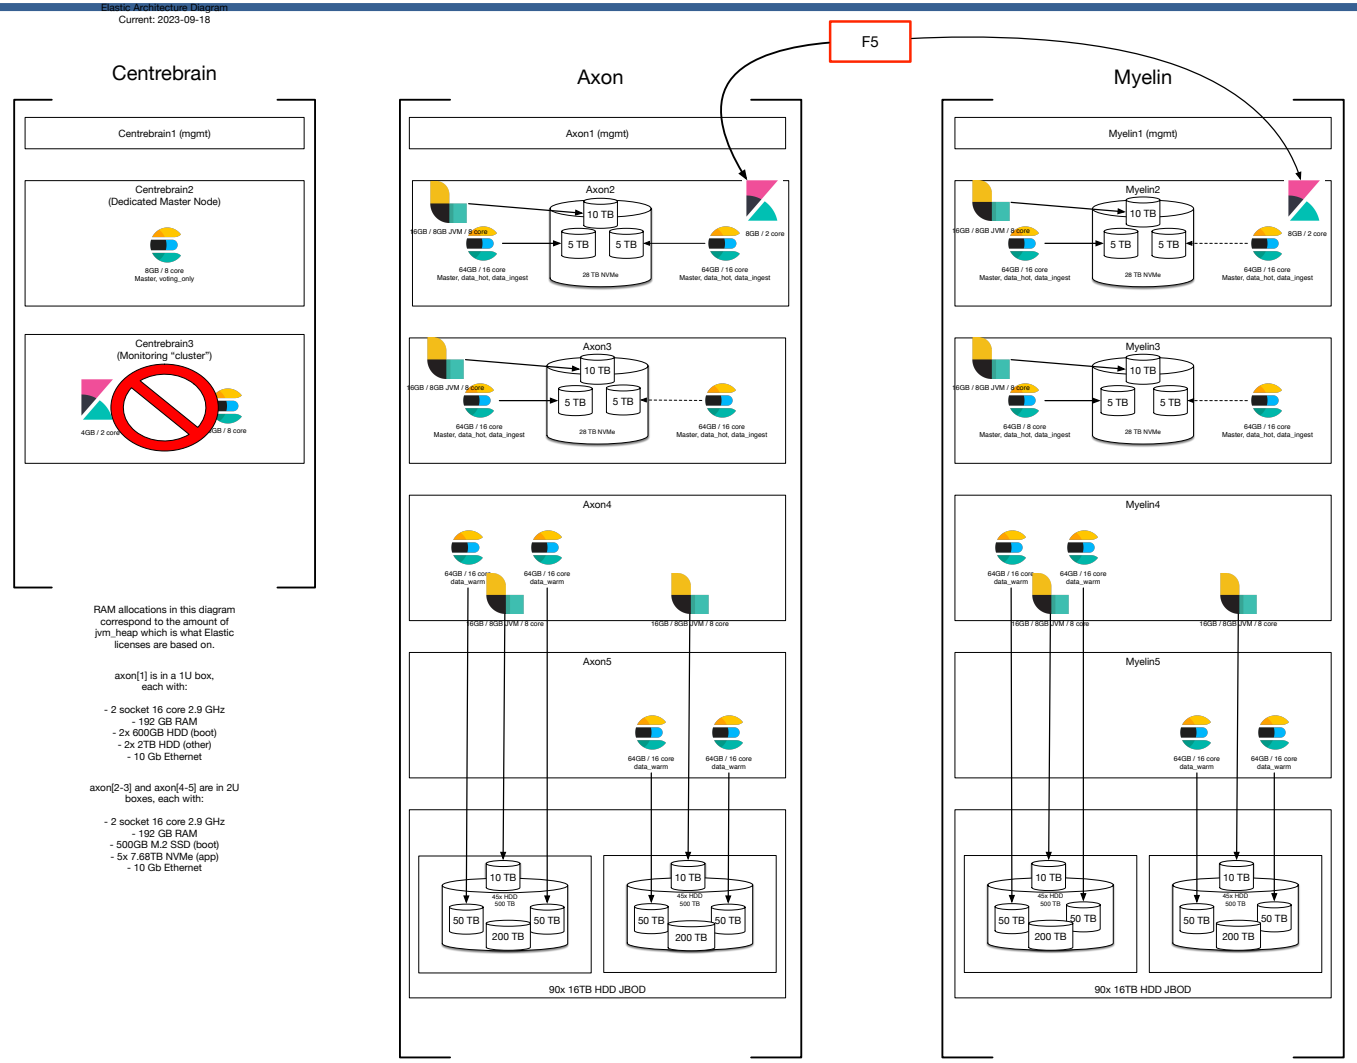

~ 112TB NVMe (total)
~ 2PB HDD (total)

Cluster Stats Today

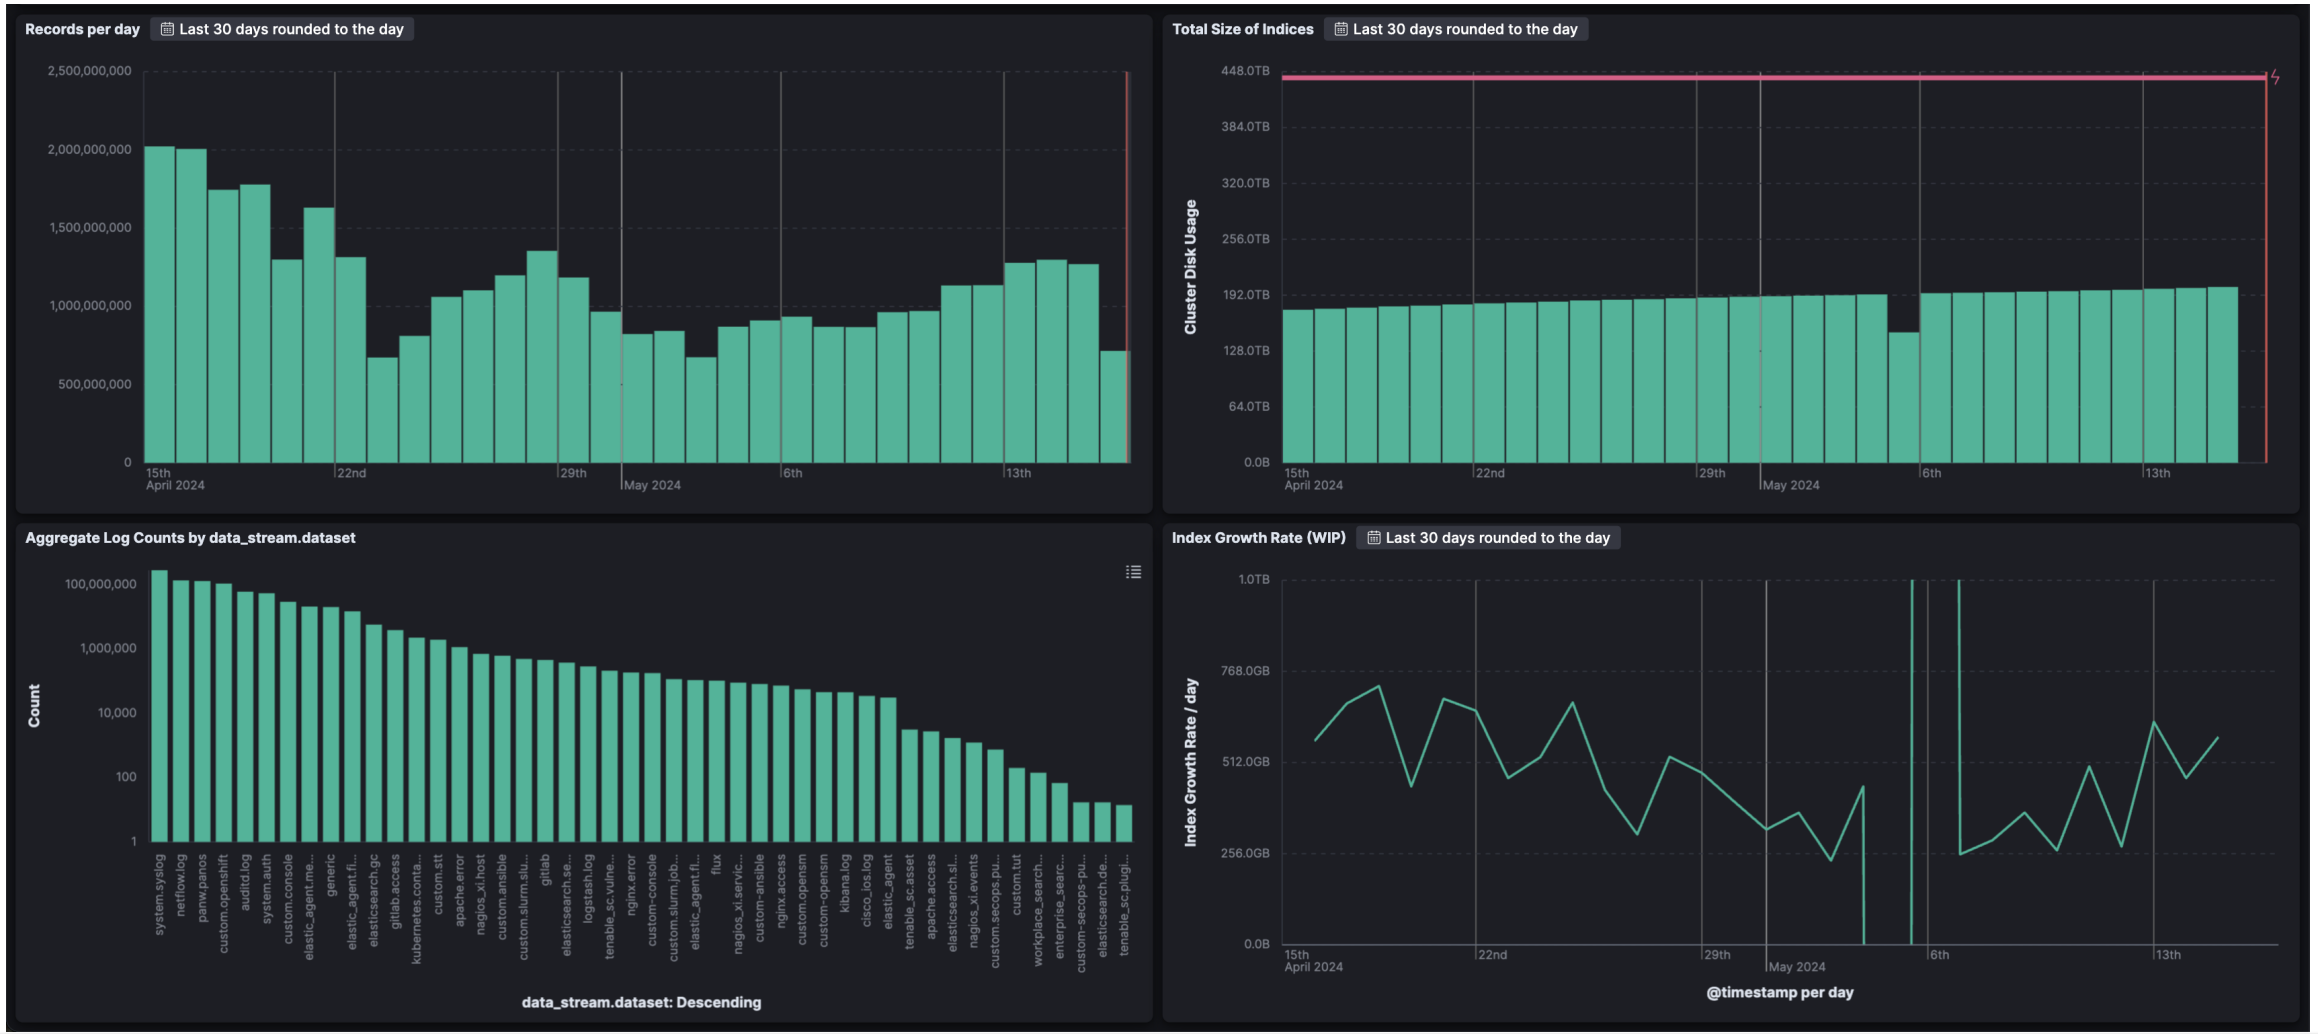

Log Sources Breakdown

Example: Auditd log issue

Jan 1 – Jan 2 (1.25 seconds)

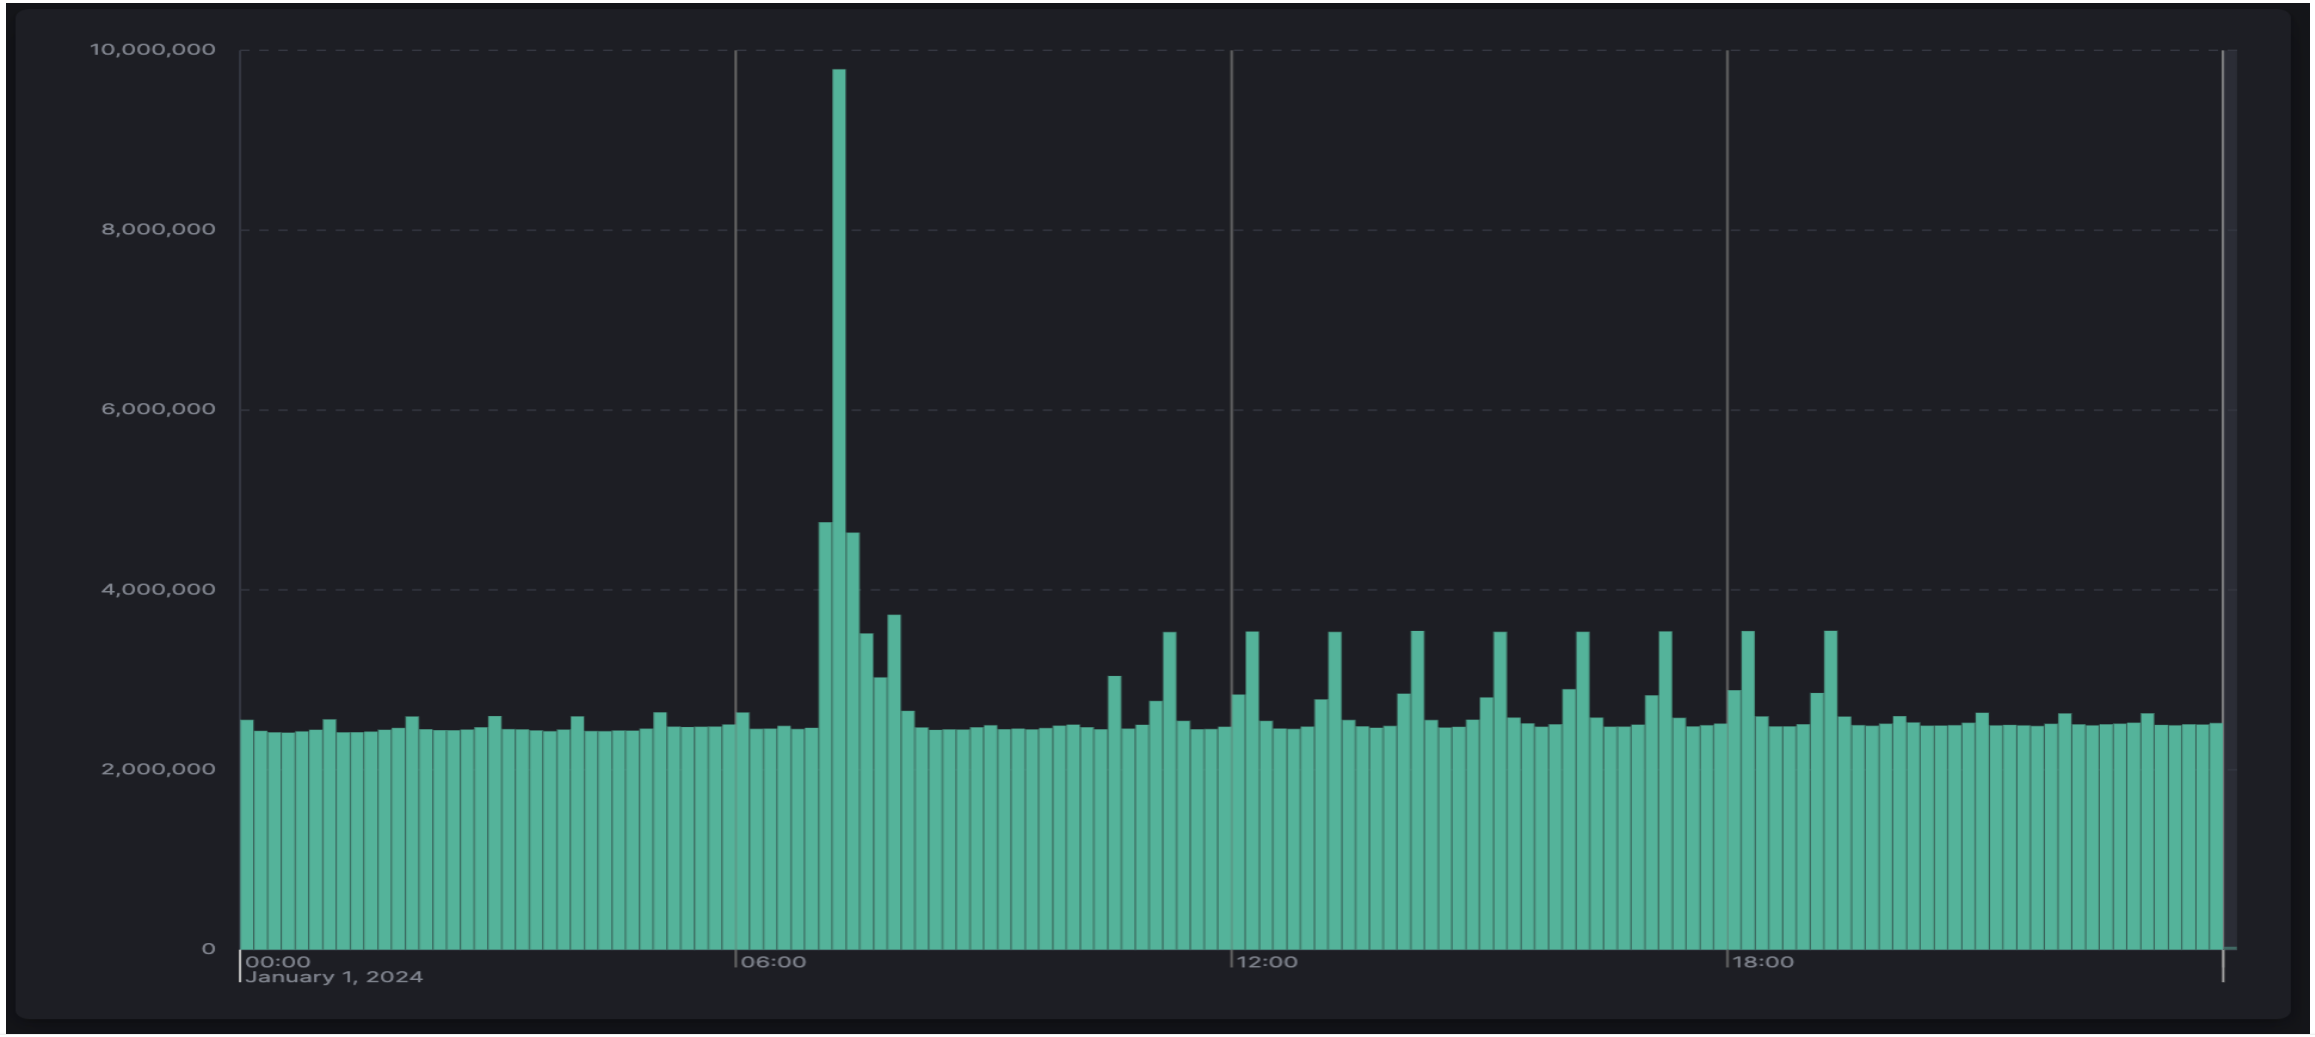

Jan 1 – Jan 8 (20 seconds)

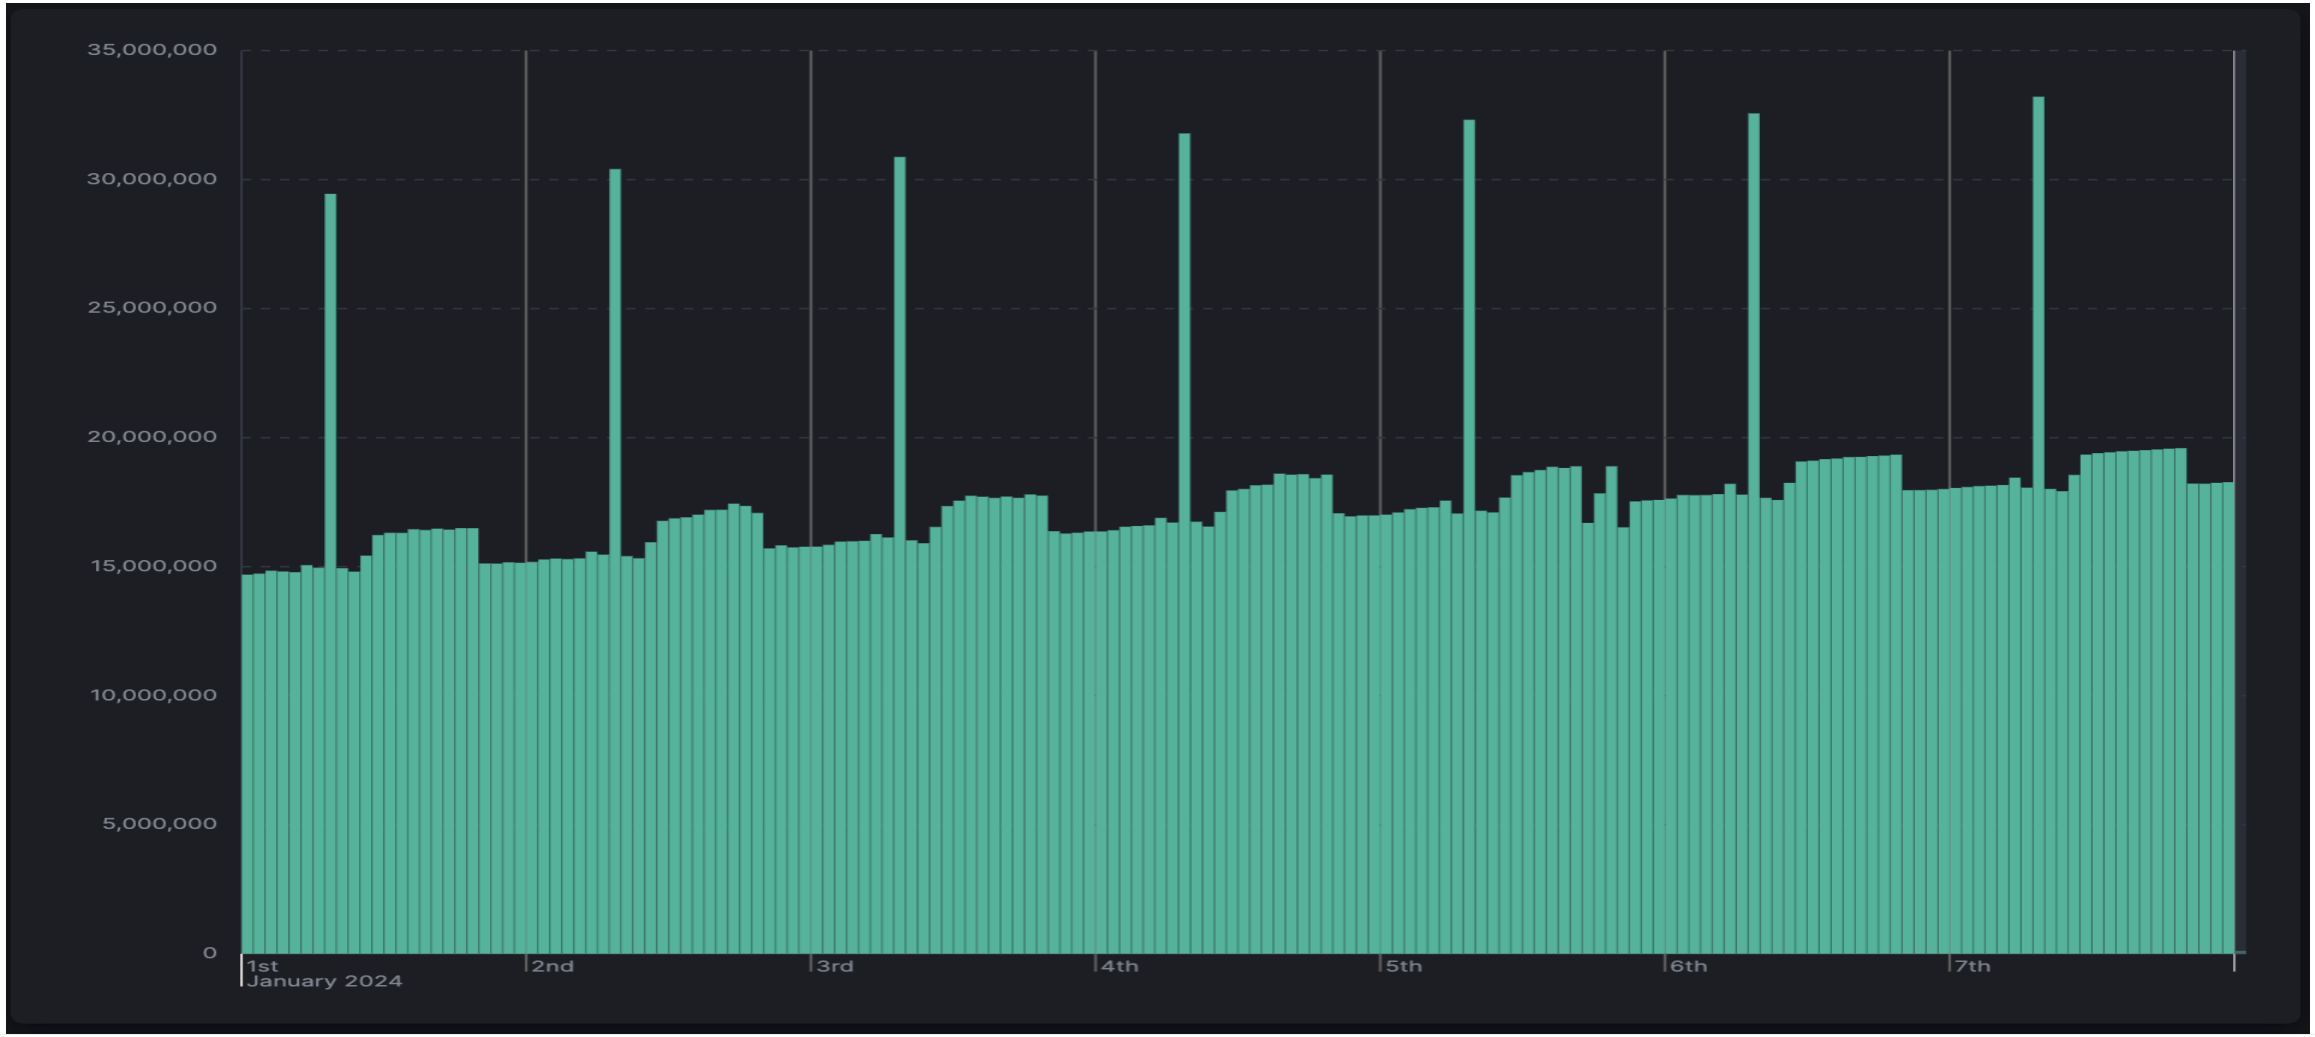

Nov 1 – Jan 8 (42 seconds)

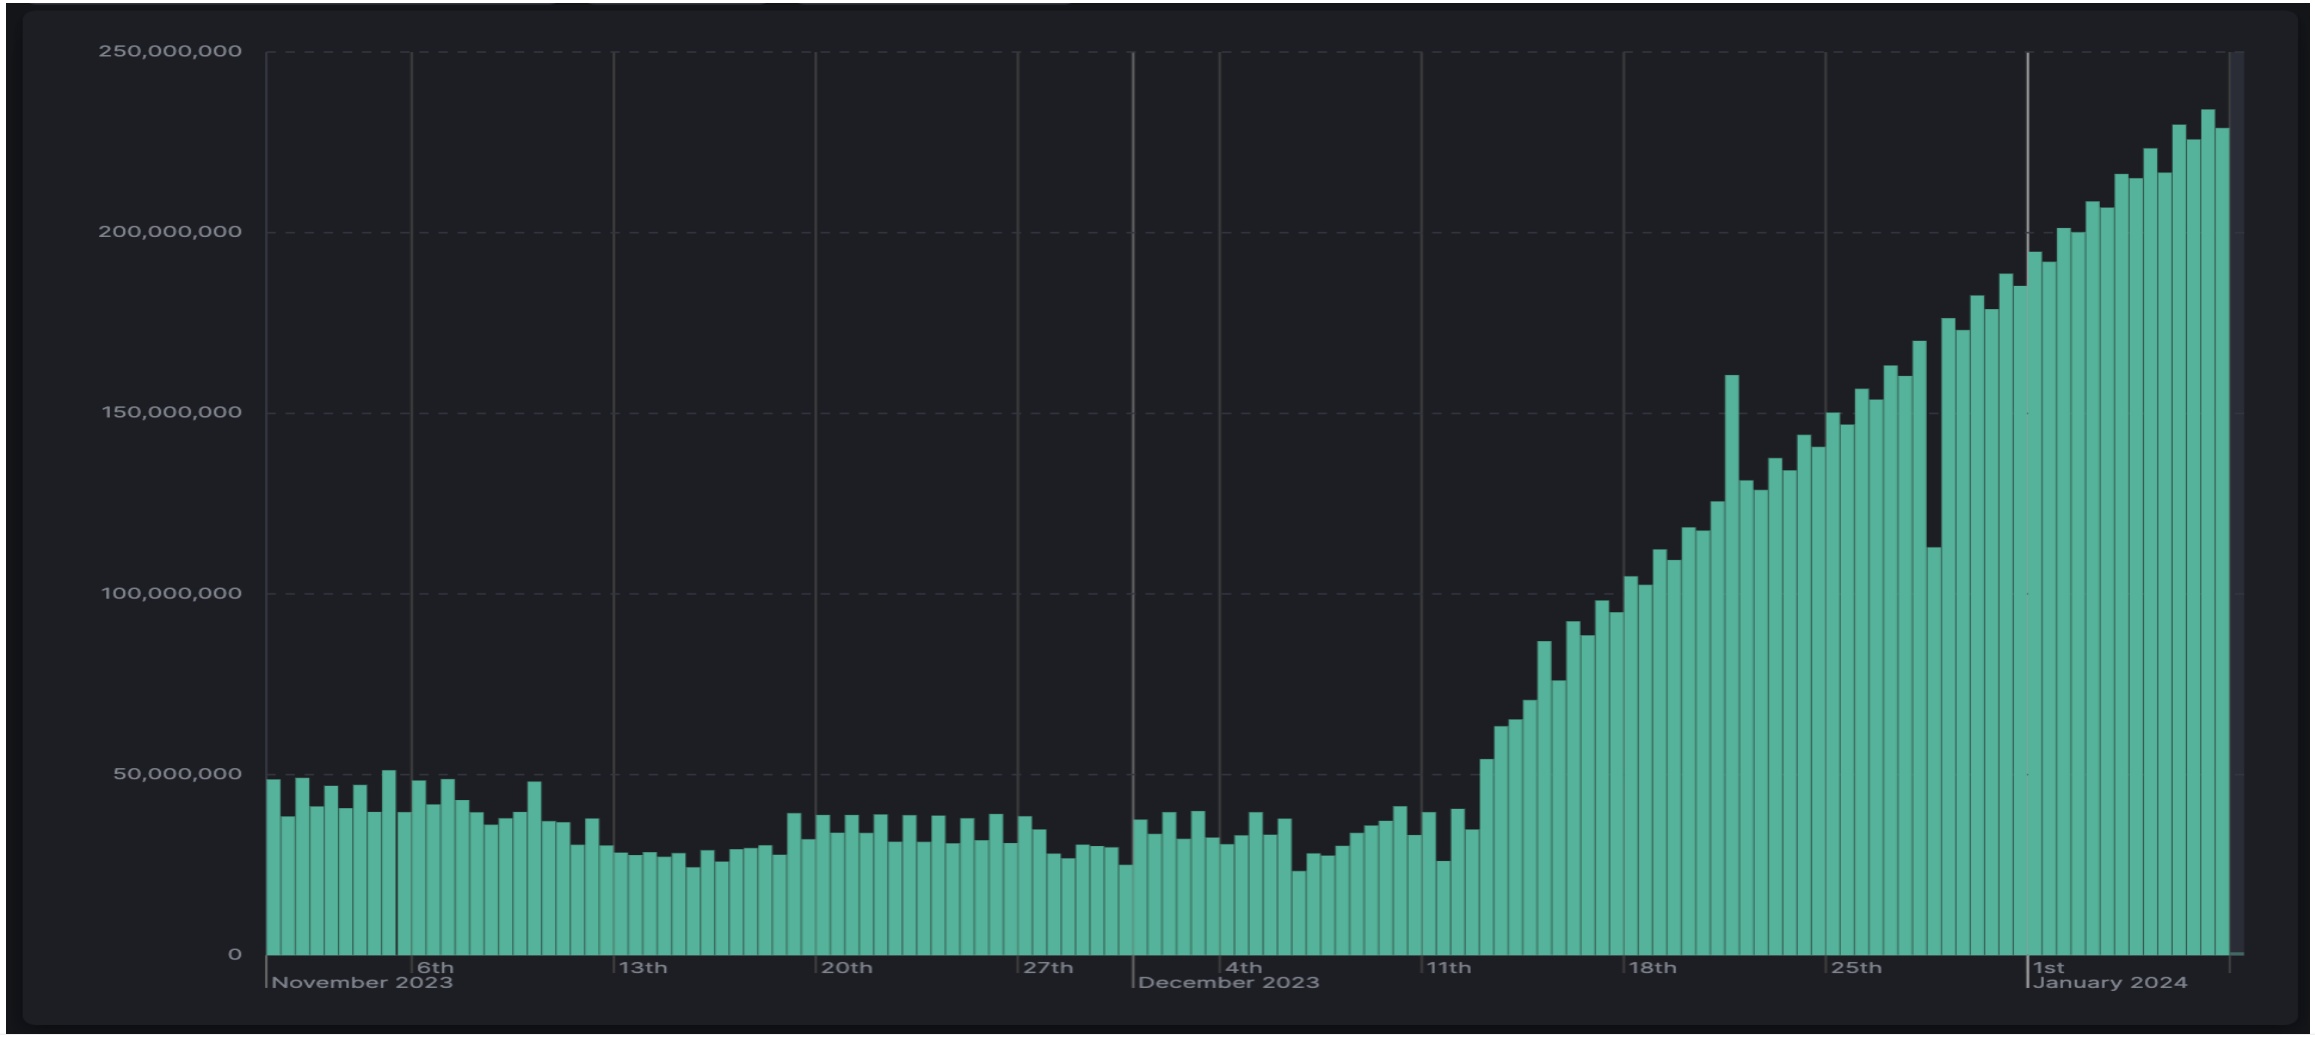

Monitoring

ServiceNow Incident Activity | RMA Dashboard | ZFS Panel | ZFS Panel Lite | IB Fabric Checker | Resolution Notes Review | WEG POG

[Operations] Hotline Events Today

Impacted systems	Event details	Start time	End time
all	Default desktop updates on all VNC servers. Only affects new sessions.	2024-03-05T07:30	2024-03-05T07:30
pascal	TOSS-4.7-4 rolling update on Pascal. See 'news TOSS_4.7-4_update' Autogenerated tickets are disabled for the 'Impacted Systems' between 0	2024-03-05T09:00	2024-03-05T09:00
poodle	TOSS-4.7-4 update on Poodle. See 'news TOSS_4.7-4_update' Autogenerated tickets are disabled for the 'Impacted Systems' between 07:00 a	2024-03-05T08:00	2024-03-05T10:00
rztopaz or rzwhippet	TOSS-4.7-4 rolling update on RZTopaz, RZWhippet. See 'news TOSS_4.7-4_update' Autogenerated tickets are disabled for the 'Impacted Syste	2024-03-05T08:00	2024-03-05T08:00

Clicking on the objects below should lead you to details for their respective panels. Give it a shot.

[Nagios] Alerts Today
313

[HPSS] Alarms Last 1 hour
8

[Operations] Unavailable file systems
0

[Operations] Center scanner offenders
0

[FIS] Files waiting
0

[Operations] MSINFO Actionables Last 15 minutes
1

[Operations] ServiceNow / OCF Unassigned Incidents Last 90 days rounded to the day

[Operations] ServiceNow / SCF Unassigned Incidents Last 90 days rounded to the day

short_description:** and assignment_group:("LC Operations" or "LC NAS" or "LC Data Storage" or "LC Production Systems" or "LC Lustre" or "LC Networks")

Last 10 years | Refresh

Incident # Any	Short description Any	Cluster Any	Created by Any
Assigned to Any	Assignment group Any	State Any	Close code Any
Shift / Assigned to Any	Shift / Opened by Any		

Use the search bar at the top for searching more broadly like 'short_description:*quartz*'
Click on the dots on the right side of the rows below to view Incident information.

Incident	Created on	Updated on	Short description	Recent Incidents	Assigned to	Assignment group	State	Close code
INC0406042	Mar 5, 2024 @ 21:50:52.000	Mar 5, 2024 @ 23:23:30.000	[redacted]	4	[redacted]	LC Operations	Active	-
INC0406041	Mar 5, 2024 @ 21:14:29.000	Mar 5, 2024 @ 21:29:13.000		0		LC Operations	Active	-
INC0406035	Mar 5, 2024 @ 19:17:53.000	Mar 5, 2024 @ 20:40:04.000		43		LC Operations	Resolved	Resolved Remotely
INC0406032	Mar 5, 2024 @ 19:13:34.000	Mar 5, 2024 @ 19:14:49.000		0		LC Production Systems	Active	-
INC0406011	Mar 5, 2024 @ 18:40:16.000	Mar 5, 2024 @ 19:34:27.000		2		LC Operations	Resolved	Cancelled
INC0405988	Mar 5, 2024 @ 17:48:56.000	Mar 5, 2024 @ 18:54:36.000		0		LC Operations	Resolved	Software - Bounced
INC0405946	Mar 5, 2024 @ 16:53:50.000	Mar 5, 2024 @ 19:46:53.000		0		LC Operations	Resolved	Cancelled
INC0405938	Mar 5, 2024 @ 16:40:58.000	Mar 5, 2024 @ 18:58:26.000		1		LC Operations	Resolved	Solved

Operating System Versions

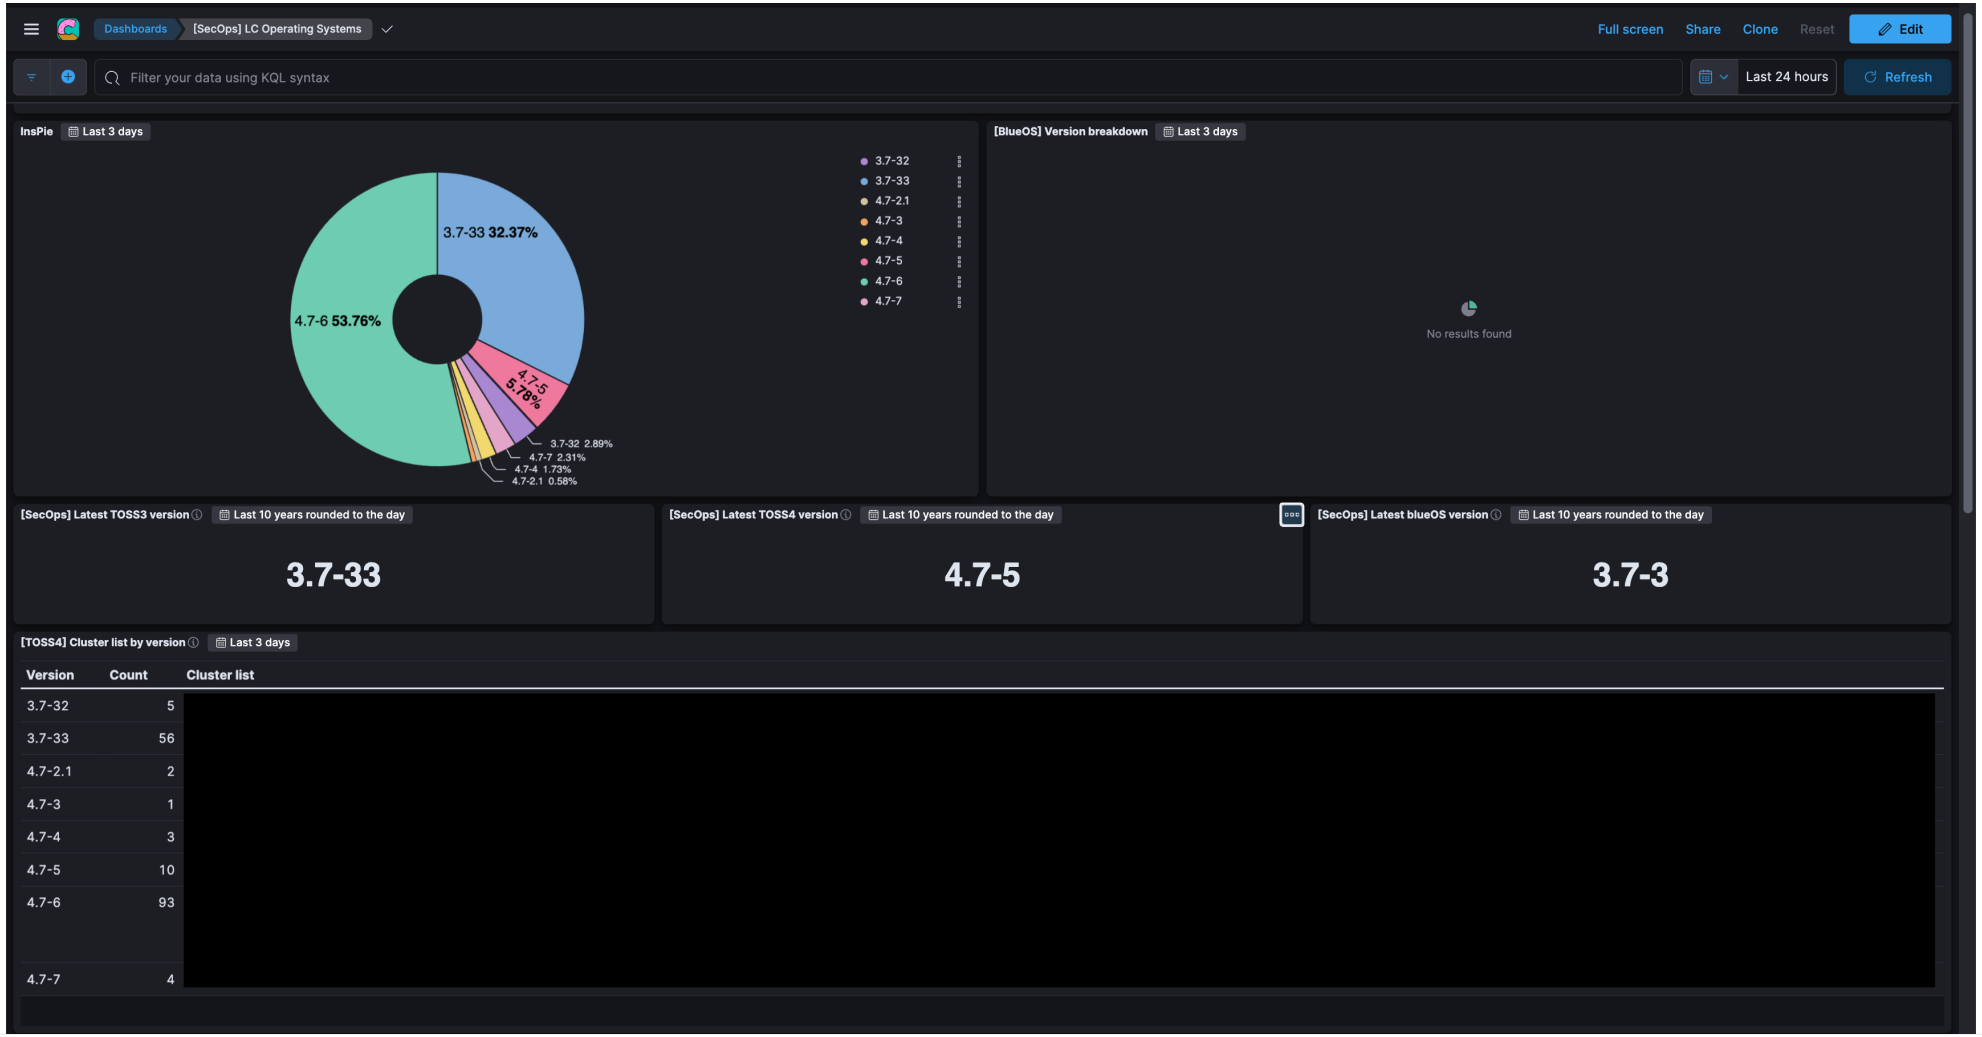

PAN Firewall Traffic Flows

SSH Authentications

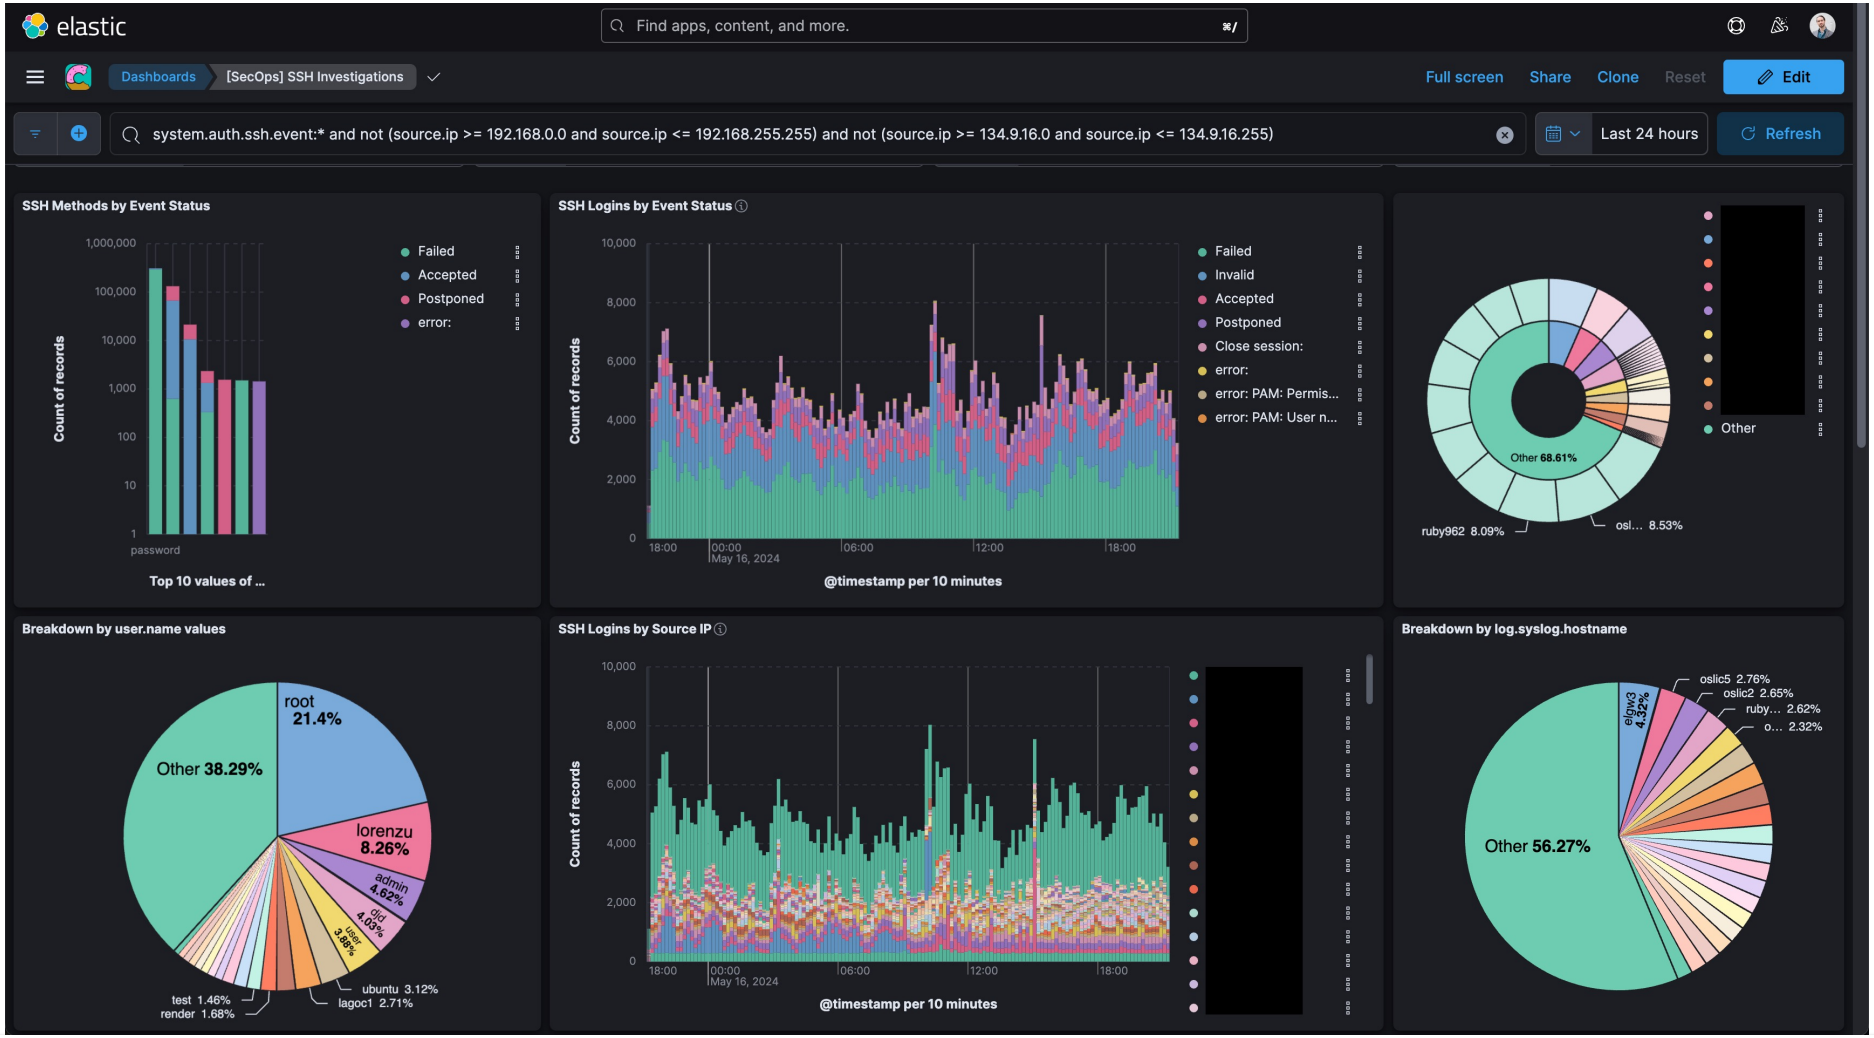

Kibana Security Dashboards

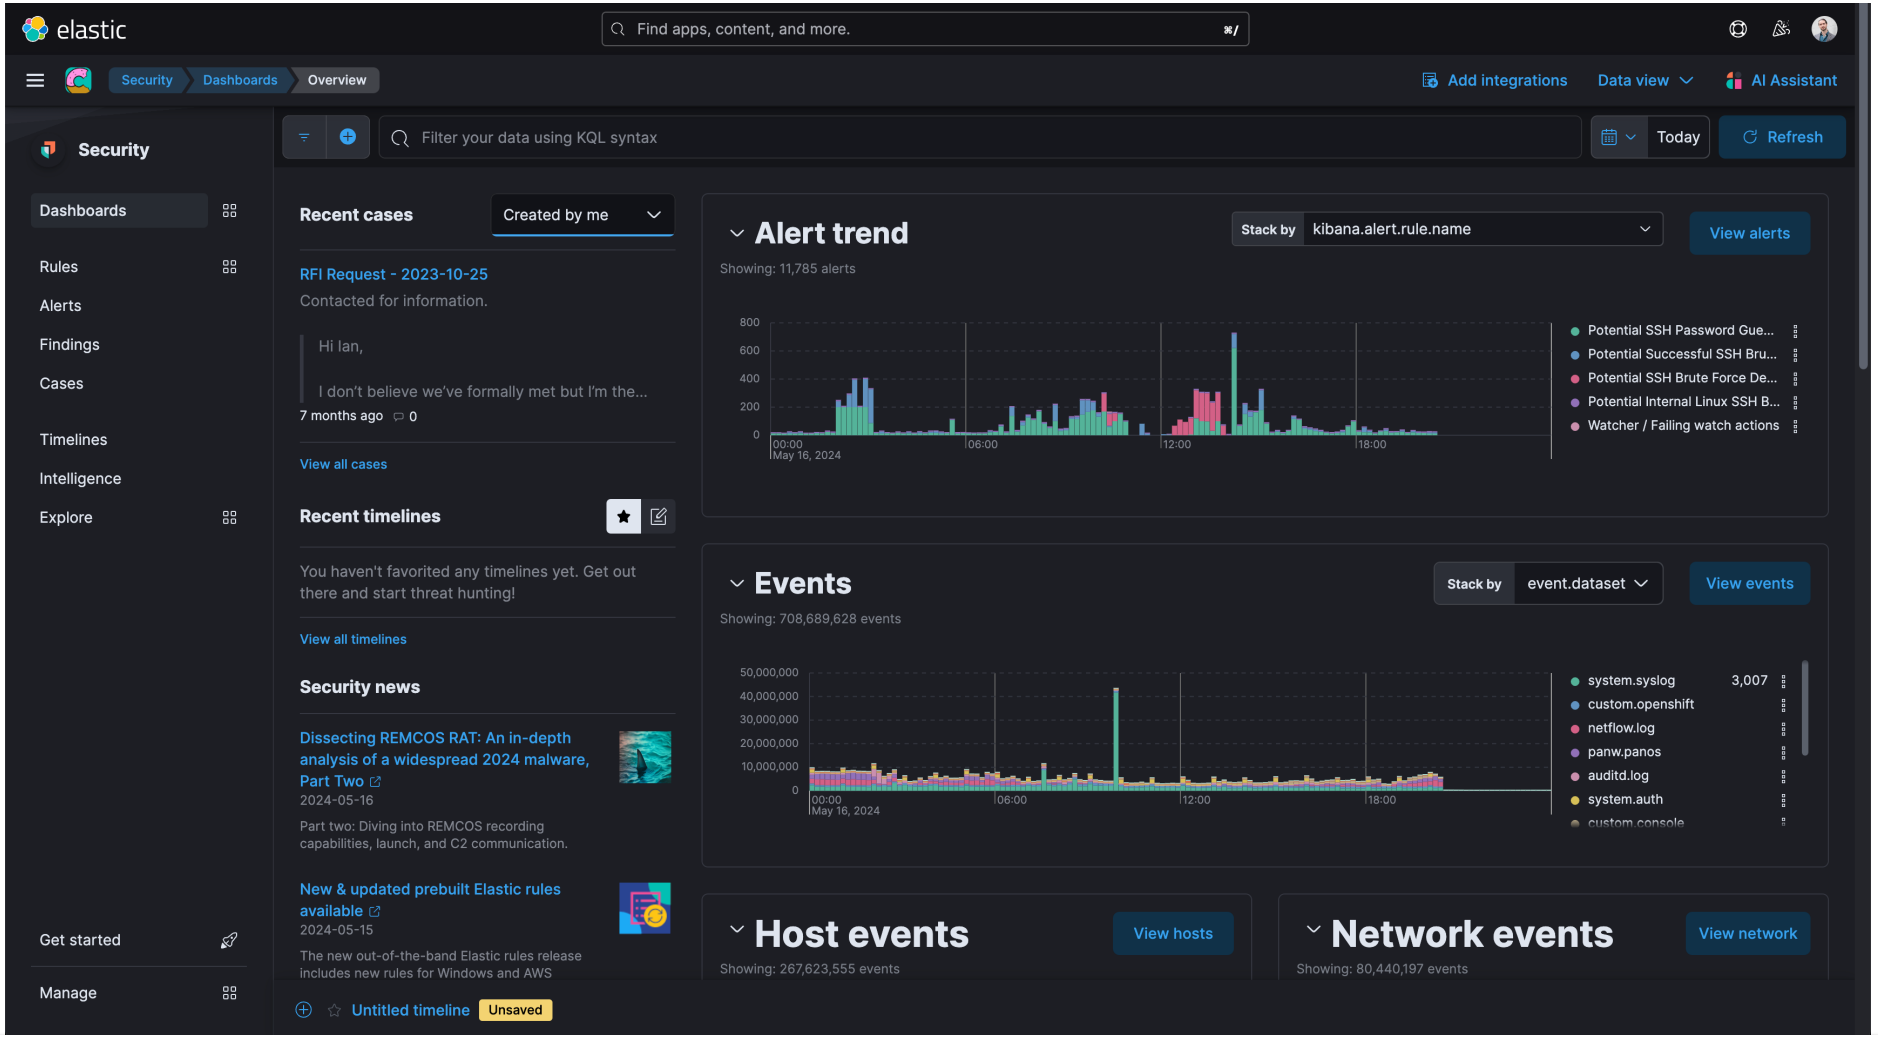

Monitoring Vision Going Forward

- Explore other offerings
  - ML / Anomaly Detection
  - Enterprise Search (unified search across web + confluence + gitlab)
  - Elastic Defend
- More automated alerts / processes

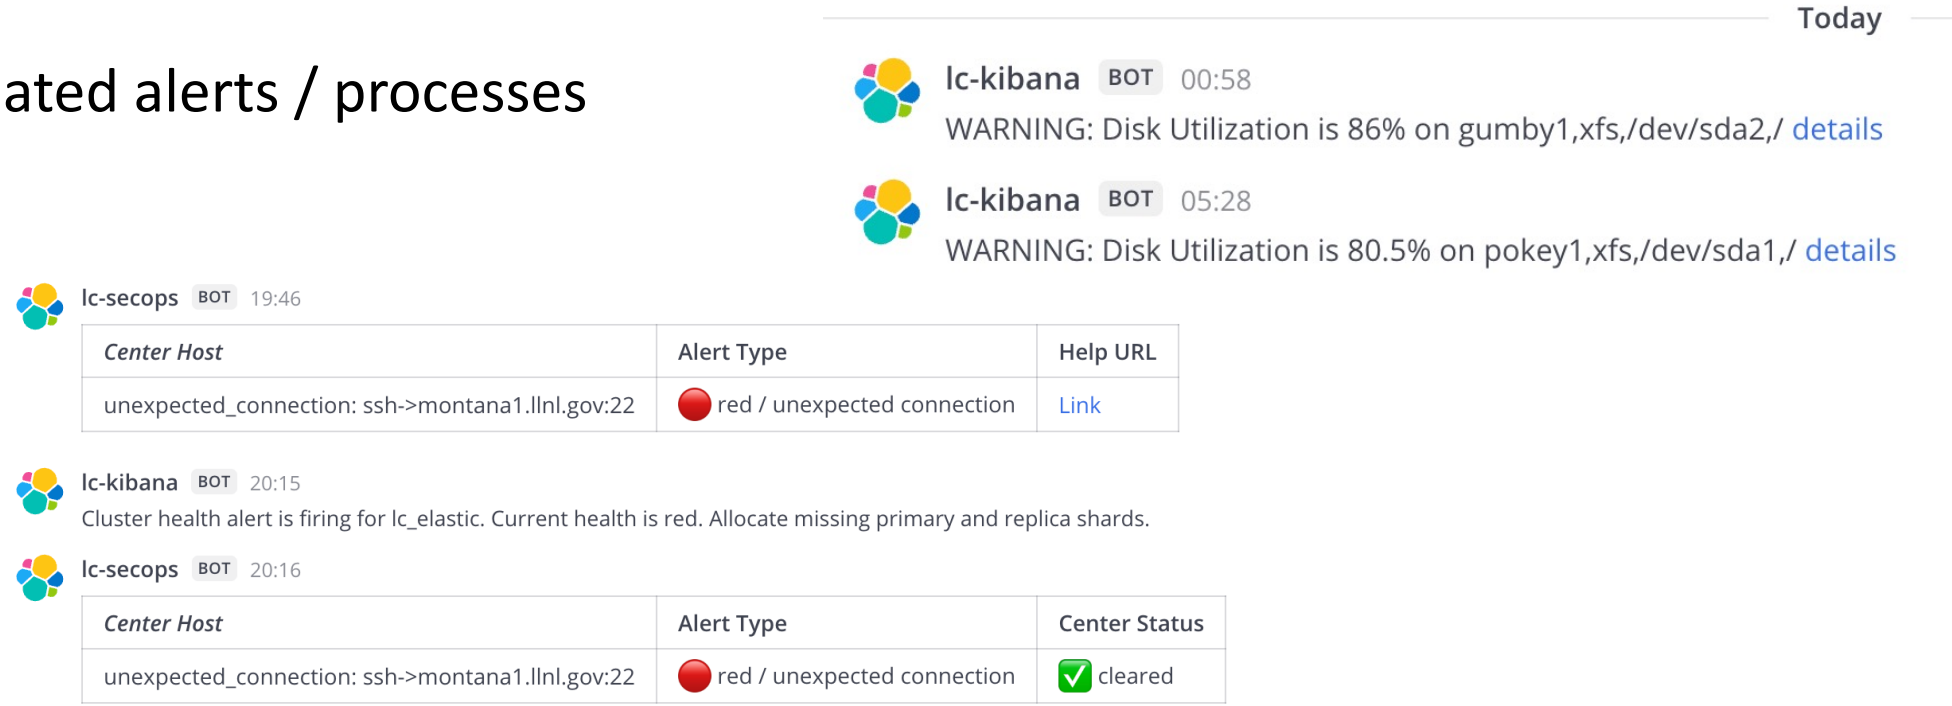

Security Baseline

- Security requirements often quite prescriptive
  - STIG > CIS Benchmark > Vendor Guideline > generic NIST 800-53 controls

- Developed a STIG for the TOSS operation system with DISA
  - Inspired by RHEL 8 STIG, which TOSS 4 is derived from
  - Small tweaks: adjust some DoD specific language to make compatible for other Gov agencies
  - Larger requests: no explicit allow-listing of software on TOSS, being a software development OS
  - HPC specific: RHEL STIG says 10 concurrent sessions for DOS reasons, TOSS STIG allows 256

- Need to regularly check and validation configuration
  - https://github.com/llnl/toss-stig

Community Work

- https://github.com/llnl/cmvl (WIP)
  - Repository of Elastic, Splunk, etc queries, dashboards, and visualizations
- https://github.com/LLNL/elastic-stacker
  - Export saved objects from Kibana for sharing
- https://github.com/LLNL/toss-configs (WIP)
  - Configuration files and scripts for setting up and maintaining TOSS HPC systems
- https://github.com/llnl/toss-stig
  - Ansible implementation of the TOSS STIG

HPC Security Technical Exchange – August 2024

- August 5 – 8, Lawrence Livermore National Laboratory, CA
- Government focus; CUI up to TS//SCI
- Registration to open imminently, including Call For Topics / Prompts
  - Contact ian@llnl.gov for details / “wait list”

Thank you!

Happy to chat and answer questions!

ian@llnl.gov

@IanLee1521

What is Elastic?

- Elastic Stack is a collection of software for logging, reporting, and visualization of data
  - Formerly “ELK stack” or just “ELK”

- Consists of Elasticsearch, Logstash, Kibana, Beats, and more

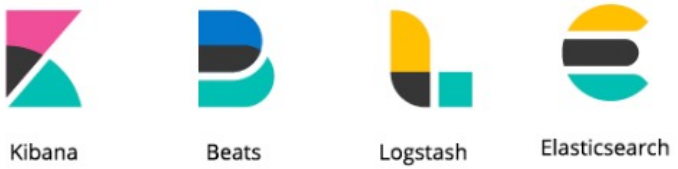

- Open source components, commercial support, similar idea to GitLab

Why Elastic?

- Were already using parts of the Elastic stack before (Logstash, *Beats)
- Better integration / extension support
  - Enterprise Search, Machine Learning tools, Elastic Integrations via Agent
- Performance claims (should be significantly faster searches compared to Splunk)
  - Reality has been a bit mixed here, and there is definitely room to continue tuning our deployment.
- Compared notes with ORNL folks who moved (at least partially) from Splunk to Elastic

Deployment

- GitLab CI + Ansible configuration (separate from the TOSS Ansible repo)
  - James Taliaferro did a talk at S3C at NLIT going into detail on this
- Very fast to destroy and rebuild the Elastic clusters
- Straightforward to scale up the service to meet demand