

# CIOQ

## **HPC User Forum**

Sept 2024

Only a few elite mechanics can make this car run at peak performance. That's who CIQ's PS Group is for you & your users with HPC/AI workloads.



# BioMed Case Study

- Tried cloud, wanted on prem system
  - 12 compute nodes plus head node
  - All work done remotely
- Mountain account with everything needed:
  - Rocky Linux
  - Warewulf
  - Apptainer, SLURM, etc.
- Seamless install:
  - Saved at least 8 hours of effort
  - Customer training was minimal since we replicated the way they worked in the cloud

# Meet the CIQ PS Group Leaders



**Nate Fuhriman**  
Technical Lead  
24 years in HPC  
Has fitted/put in production  
the 3<sup>rd</sup> largest system in the  
world on down



**Victor Wright**  
Project Management  
34 years in IT, 12 in HPC  
People-first mentality  
Honesty and trust are  
deep in company DNA



**Mark Wardrop**  
Director of Support  
30+ years in HPC  
Applications expert  
HPC analyst 2003-2022 on  
ConocoPhillips Seismic and  
Reservoir HPC team



**Devin Jensen**  
Business administrator  
30+ years in IT, 24 in HPC  
Champions customers  
Listener

# Your Trusted HPC Experts

## Software

install and configure Linux  
Set up provisioning system: (Warewulf, SMC, Bright)  
Set up workload managers (PBS, SLURM, & others) for end-to-end job submission  
Set up monitoring of all system resources  
Maintain CSP SAAS offering and/or custom software stack

## Configure and Maintain Clusters

Install applications  
Update software/BIOS/etc.  
Make resources available for end users  
Configure, manage, optimize, and have proactive support to resolve administrative issues  
Special projects as needed (IT, user, stack, etc.)

## Monitoring

Provide relevant reports and tracking to ensure system capabilities are functional as desired  
System utilization  
Storage metrics  
Testing and tuning  
Bottleneck identification  
Application benchmarking  
Memory and interconnect performance

## Client Hardware and Cloud Services

Maintain, install and or validate physical hardware (including GPUs)  
Help with POCs or benchmarks  
Auto-scaling  
Perfect cloud images  
Failover  
Backup  
Data recovery  
Setup and configuration of storage  
Adjust workload management for maximum efficiency and cost performance (including AI/ML, LLM, CFD, etc.)  
Firewalls and license servers  
Configure networking (IP, switches)  
Optimize high speed network (InfiniBand, Omnipath, 100g ethernet)  
Tune BIOS for HPC workloads

## Additional Custom Services

US-based standard support beyond normal business hours  
International support during normal business hours  
Various turnaround times (4, 12, 24 hours)  
L1 support (L2/L3 support is standard)  
Weekend, 24x5, and 24x7 support  
Various service levels depending on business needs and the number of users and applications  
Special projects beyond adding users, installing software, debugging system problems, etc.

# How We Build “Trusted Advisor” Status

- We create long-term relationships through a joint discovery process
  - Listen carefully to understand what your customers or prospects believe their issues are
  - Suggest possible solutions that we review and agree on together
  - Identify salient action items with follow up reporting discussions
  - Rinse and repeat as often as needed
- Key operating tenets:
  - Immediately support any issue deemed critical by client
  - Honestly admit when we don't know the answer or results disappoint
  - Say what we'll do and do what we said (return and report)
  - Selfless with best practices/happy to train
  - Under promise and over deliver
  - Treat others like we should be treated

# Milos in The Venetian, T-W 9am-5pm (OCW)



LAS VEGAS



The Venetian Las Vegas  
3355 Las Vegas Boulevard South  
Las Vegas, Nevada 89109

702-414-1270



Backup

# Summary from Callaway

A Topgolf Callaway Brands Corp. brand, Callaway Golf is the leading manufacturer of premium golf clubs, balls, performance gear, and accessories worldwide. Through an unwavering commitment to innovation, they push the limits of performance and create superior products designed to make every golfer a better golfer.

Needing to add more compute cores to the cloud, they chose Oracle Cloud Infrastructure for its ability to support their complex compute, storage, and networking requirements. With the help of CIQ's High Performance Computing (HPC) services, they expect a 40% improvement in performance on OCI compared with on-premises cores, ultimately helping them gain more market share within the golf industry.

# Callaway Golf Details

## Background

CIQ's PS group was instrumental in helping support Callaway Golf Company's transition to OCI in a recent HPC win. Callaway Golf was running out of room in their current data center and club designs were taking a year to create. Callaway Golf needed a platform to scale their workload intensive applications, which Nate advocated was best to be run on OCI. Nate is pivotal in ensuring HPC based projects and implementations go smoothly thereby ensuring that the consumption takes place within the customer's tenancy. CIQ's PS group has a unique position in the market helping customers make the most efficient use of their Linux based HPC systems after delivery. They provide the care and proactive support that these complex systems require, tailored specifically to customer's workloads.

## What did CIQ's PS group do?

Nate was actively involved in selecting OCI as the preferred Platform for Callaway. He knew that Oracle offered Callaway exactly what they were looking for: Bare Metal, Highly Performant Cloud built with RDMA based Networking, which emulated their on-prem cluster environment that he was managing for them. Nate upgraded Callaway's existing on-prem infrastructure to a modern OCI hybrid. Although the behind-the-scenes work was extensive, the migration process for Callaway's simulation engineers was smooth and invisible—they don't know where their jobs run (or care). Our work includes implementation, auto-scaling, firewalls, license servers, perfecting cloud images, and more.

## Customer Benefits

The OCI implementation worked through a sizable backlog within the first month of tenancy. Callaway considers their expansion key to increasing their market share in the golf industry and does not want their competition to know more than that. With OCI, their hybrid computing costs are more efficient for production, so the business benefit is keeping their market leadership position.

CIQ's PS group provided the ability to scale up and down as needed to maximize investment into expensive 3rd party design licenses from Ansys.

OCI Competitive Differentiators:

- Data ingress and egress with FastConnect (We don't hold data hostage)
- Consistent regional pricing

- High caliber networking capabilities.
- OCI delivers a hybrid environment that works with Callaway's business practices so no additional training for engineers is required—Nate makes it so the engineers don't know (or care) where their jobs run
- Access to latest technologies (CPU and GPU) as soon as they are released on OCI. CPU/GPU increased performance also lowers 3rd party licenses costs.
- Over 40% performance gain compared to on-prem
- Ease of data mobility
- Unlimited scalability for all their workloads



Only a few elite mechanics can make this car run at peak performance. That's who CIQ's PS Group is for you & your users with HPC/AI workloads.

CIQ was started in 2019 by Greg Kurtzer, the founder of CentOS and Rocky Linux, to disrupt enterprise computing. This US-based, venture-backed company modernizes the software infrastructure for enterprise and government agencies.

We are creators and collaborators across these domains.



**ROCKY LINUX**

Secure, Stable, Trusted  
Enterprise Linux



**MOUNTAIN**

Complete Infrastructure  
Management & Security



**ASCENDER**

Infrastructure  
Automation



**FUZZBALL**

Cloud-Native Federated  
HPC 2.0 Platform



**APPTAINER**

Secure Application  
Containers



**WAREWULF**

Cluster Management  
and Provisioning

# The Rise of Rocky Linux


Inventing the successor to CentOS & changing the world of enterprise computing

**Dec 2020 -  
CentOS EOL /  
Rocky Begins**

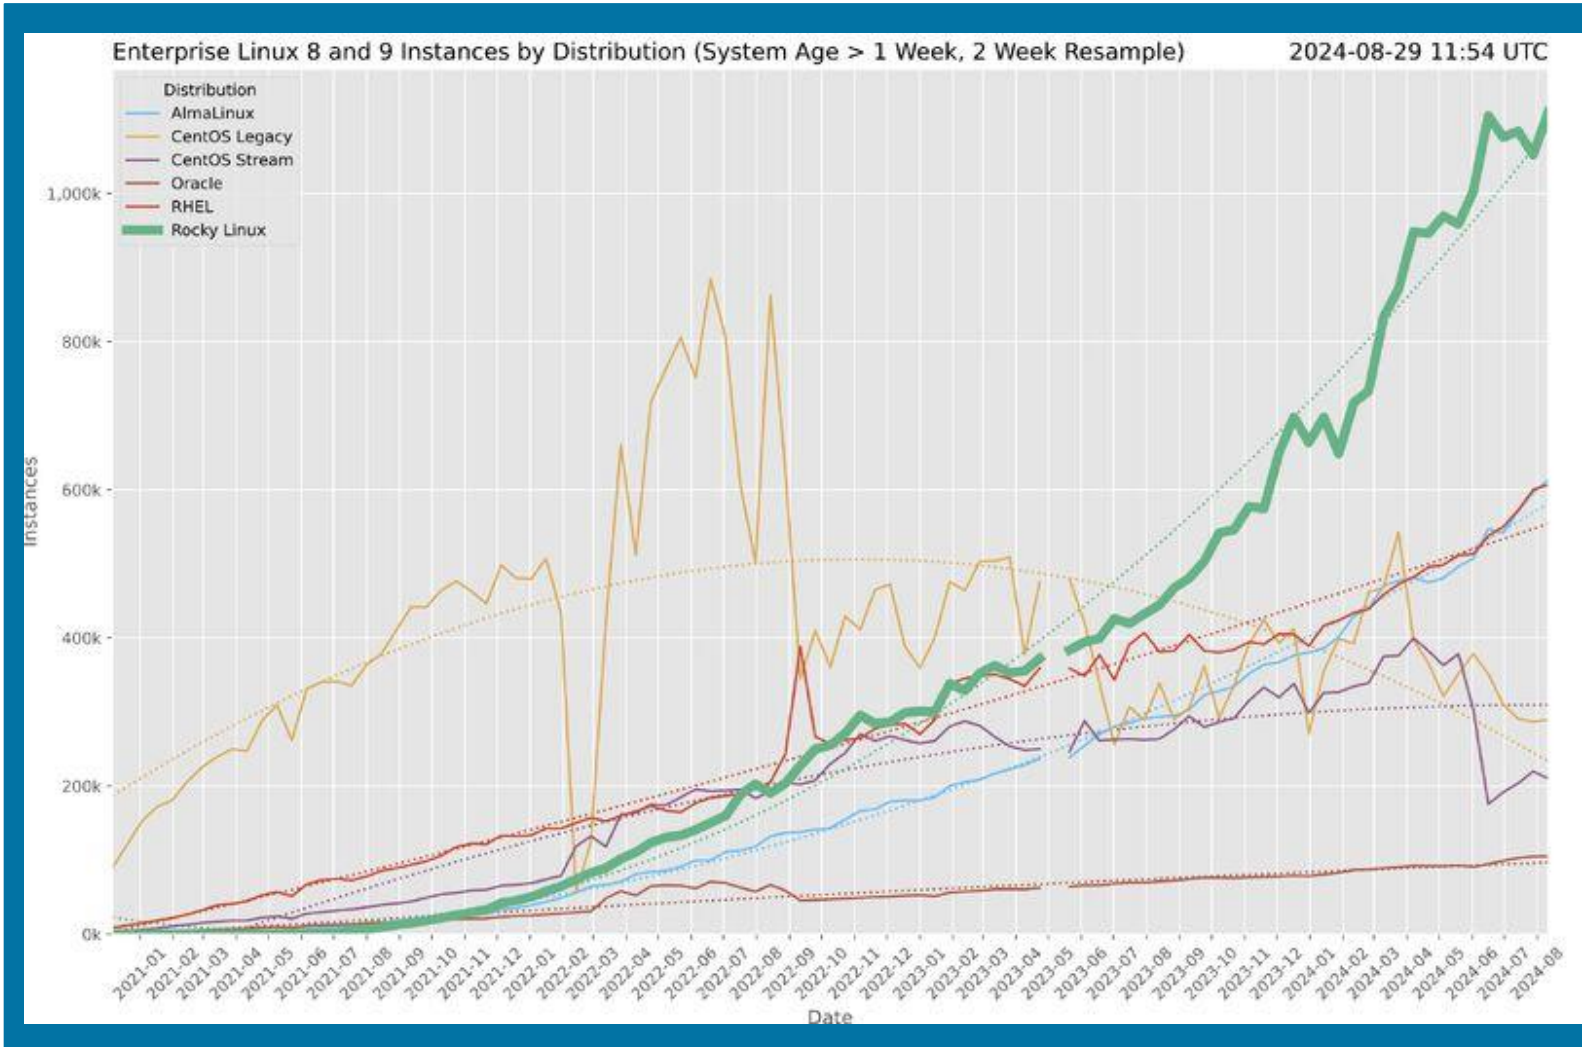
**Feb 2021 -  
RESF Founding**

**May 15, 2022 -  
8.6 LTS Release**

**2023 -  
Cloud Marketplace  
Expansion**



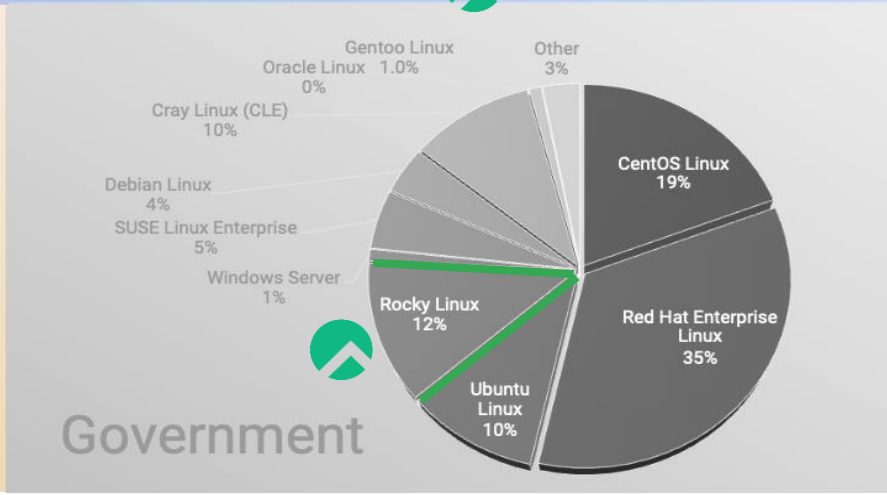



# Adoption Keeps Climbing



## Rocky = Rocket Ship

- **Growth:** #1 Fastest Growing Enterprise Linux
- **Testing:** #1 EPEL deployment, predictive of future growth
- **Cloud Deployments:** Over 1B+ instances deployed in 2023
- **Community Support:** Recent industry news has 10x user and developer support
- **Accelerated Adoption:** Major Enterprise and Public Sector deployments are heavily increasing due to licensing and subscription costs. Typical customer savings >50%

# Rocky Linux Usage Survey



- Active Deployments
- AL / ML / HPC Workloads
- Performance Intensive
- Stable
- Secure

# Product Overview

# CIQ Enterprise Linux Platform



## Rocky Linux:

- **Compatible:** Alternative for Red Hat Enterprise Linux & CentOS
- **Stability:** Long-Term Supported images for 18 months
- **Performance:** Optimized and upstream kernel images
- **Secure:** FIPS 140-3 certification in progress
- **Choice:** Choose the version that is best for your use case



## Ascender Automation:

- **Increased Efficiency:** Save time, reduces manual errors, and free up time
- **Improved Consistency:** Standardize processes and configurations
- **Enhanced Scalability:** Manage larger and more complex infrastructure easily
- **Improved Security:** RBAC, credential management, and audit trails
- **Collaboration:** Enable multiple teams with centralized control



## Mountain:

- **Centralized Repository:** Secure access for Rocky Linux, Apptainer and Warewulf, container images, and optimized Rocky Linux versions
- **Fine-grained Access Control:** Define precise access permissions for users and teams based on roles (RBAC) and subscriptions
- **Secure Delivery Pipeline:** Integrate with CI/CD tools to automate artifact creation, testing, and secure distribution

# Choose the ideal version of Rocky Linux & Support SLA

## 8.x or 9.x Point Release

- Latest version
- Updates every 6 months

## 8.x or 9.x Long-Term Support

- Built for stability
- Aligns to even-numbered point release
- Supported for 18 months after the point release changes

**Standard**  
**8am-6pm local time**

**Advanced**  
**24/7**

# Rocky Linux LTS Timeline

| Rocky Linux 8 Release Timeline |      |      |      |      |      |      |      |      |      | *Planned EOL: May 31, 2029 |
|--------------------------------|------|------|------|------|------|------|------|------|------|----------------------------|
| OS Release                     | 2021 | 2022 | 2023 | 2024 | 2025 | 2026 | 2027 | 2028 | 2029 |                            |
| *8.10 (Final)                  |      |      |      |      |      |      |      |      | *    |                            |
| 8.9                            |      |      |      |      |      |      |      |      |      |                            |
| 8.8                            |      |      |      |      |      |      |      |      |      |                            |
| 8.7                            |      |      |      |      |      |      |      |      |      |                            |
| 8.6                            |      |      |      |      |      |      |      |      |      |                            |
| 8.5                            |      |      |      |      |      |      |      |      |      |                            |
| 8.4                            |      |      |      |      |      |      |      |      |      |                            |

| Rocky Linux 9 Release Timeline |      |      |      |      |      |      |      |      |      |      |      |      | *Planned EOL: May 31, 2032 |
|--------------------------------|------|------|------|------|------|------|------|------|------|------|------|------|----------------------------|
| OS Release                     | 2021 | 2022 | 2023 | 2024 | 2025 | 2026 | 2027 | 2028 | 2029 | 2030 | 2031 | 2032 |                            |
| *9.10 (Final)                  |      |      |      |      |      |      |      |      |      |      |      | *    |                            |
| 9.9                            |      |      |      |      |      |      |      |      |      |      |      |      |                            |
| 9.8                            |      |      |      |      |      |      |      |      |      |      |      |      |                            |
| 9.7                            |      |      |      |      |      |      |      |      |      |      |      |      |                            |
| 9.6                            |      |      |      |      |      |      |      |      |      |      |      |      |                            |
| 9.5                            |      |      |      |      |      |      |      |      |      |      |      |      |                            |
| 9.4                            |      |      |      |      |      |      |      |      |      |      |      |      |                            |
| 9.3                            |      |      |      |      |      |      |      |      |      |      |      |      |                            |
| 9.2                            |      |      |      |      |      |      |      |      |      |      |      |      |                            |
| 9.1                            |      |      |      |      |      |      |      |      |      |      |      |      |                            |
| 9.0                            |      |      |      |      |      |      |      |      |      |      |      |      |                            |

|      |   |
|------|---|
| Key: |   |
|      | Rocky Linux Minor Release                               |
|      | CIQ Standard Long Term Support                          |
|      | Extended CIQ Long Term Support                          |
|      | CIQ Certified Rocky Compliant Releases (DISA/STIG/FIPS) |

“LTS” is given only to non-zero even-numbered releases, like: 8.6, 8.8, 9.2, 9.4, etc. for 18 months past release expiration

“Extended LTS” versions will continue to receive updates beyond the standard LTS product. To align with CIQ’s FIPS product CIQ will extend LTS for Rocky 8.6, 9.2, and 9.6

Rocky Linux versions 8.10 and 9.10 will receive package updates for 5 years


## Security Vulnerability Policy

- Security related backports are done for packages that are marked with a CVSS score of 8.0 or higher.
- Packages with security issues marked as Important or Critical from upstream may be eligible for backports
- Packages with a CVSS score lower than 8.0 may receive backports, at CIQ’s discretion
- Backported packages may receive needed bug fixes, at CIQ’s discretion (not security related)

# For Everyday Production and Development Environments



Rocky Linux: Latest  
Rocky Linux LTS: 8.6, 8.8, 9.2, 9.4  
Rocky Linux FIPS LTS 8.6  
Container Images: Rocky Linx 8 & 9



Automate daily tasks  
Validate and report with Ledger



# CIQ Performance Computing Platform

CIQ Enterprise Linux Platform



## Apptainer (former Singularity)

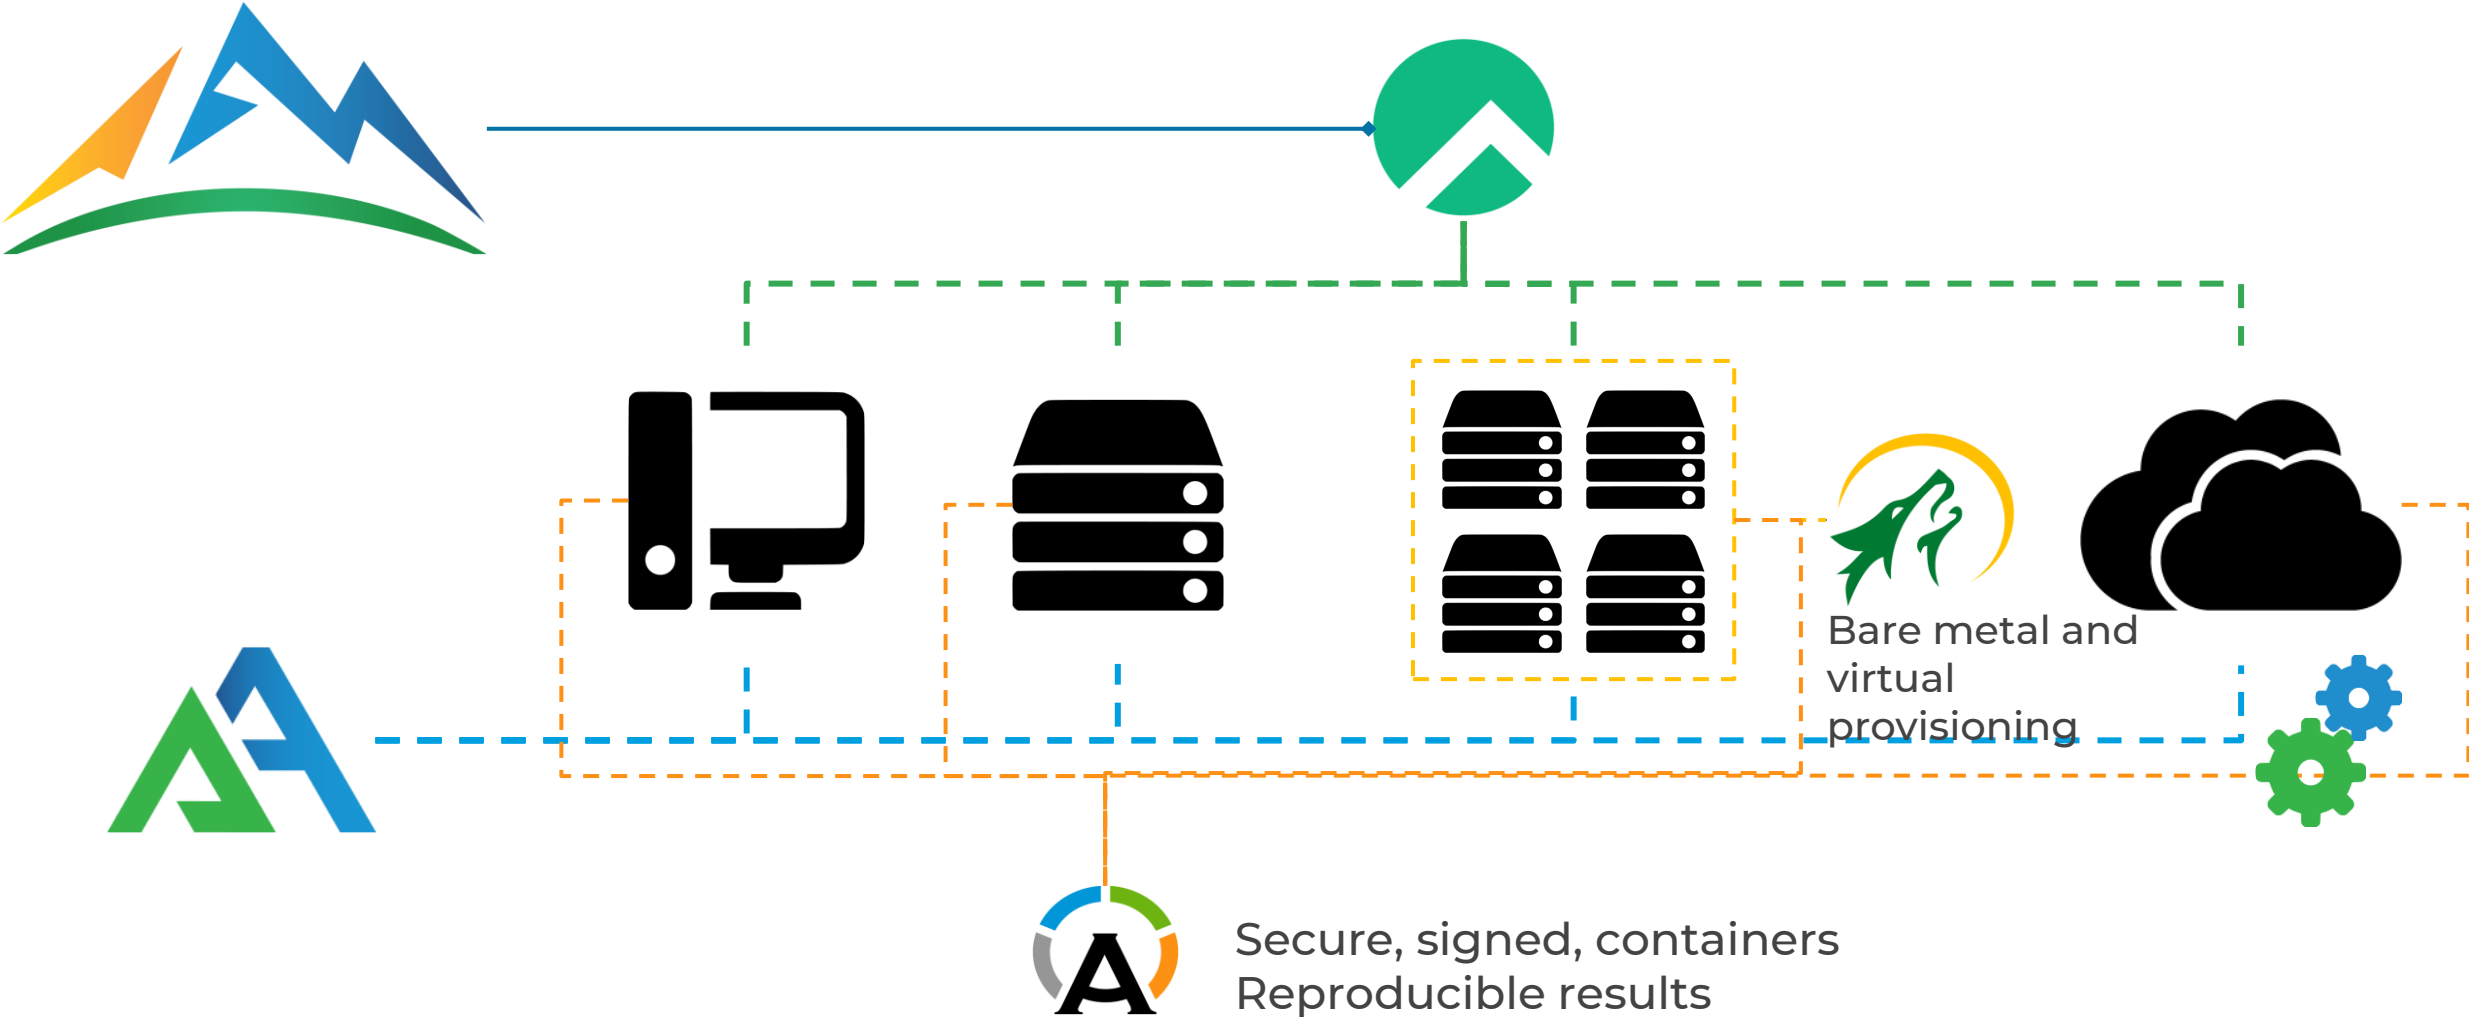
- **Verifiable Security:** Reproducible results, using cryptographic signatures, an immutable container image format, and in-memory decryption
- **Easy Integration:** Make use of GPUs, high speed networks, parallel file systems on a cluster or server by default
- **Mobility:** The single file SIF container format is easy to transport and share



## Warewulf

- **Lightweight & Simple:** no underlying system dependencies or changes to the provisioned cluster node operating systems
- **Flexible:** Address needs from workstations to Supercomputing Centers
- **Agnostic:** Deploy on your preferred Enterprise Linux distribution and underlying hardware
- **Secure:** Stateless provisioning system that supports SELinux out of the box

# For Performance Intensive Compute Environments



# CIQ Bridge Solution

CIQ will solve this problem with CIQ Bridge to address -

- Support for CVSS Scores 7+ CVE 9-7
- Critical security issues fixed as outlined, but no active development or backporting
- No guarantee of non-security related updates (bug fixes)
- Proprietary drivers not supported; troubleshooting requires an additional charge
- No Desktop Support
- Functionality fixes provided under a PS-style engagement

To address common vulnerabilities with a score lower than 7, CIQ will offer pre-purchased Engineering Hours that can be used in the following flow-

1. Customer requests CVE lower than a score of 7 to be addressed via CIQ Support Ticket
2. CIQ reviews the CVE requests and begins assessing the feasibility of addressing the CVE
  - a. Cost = 5 hours per CVE request
3. If CIQ can address the requested the CVE, CIQ will come back with a Statement of Work that include the cost against the committed Engineering hours and a timeline to complete

# Enterprise-Ready Solutions

## CIQ Enables-

- Software Infrastructure
- HPC / AI / ML
- Automation
- Customized and Optimized Images
- Stability
- Security
- Compliance
- HUGE Cost Savings

## Product License Models:



### Per Node on-premise:

Each deployment of Rocky Linux

- Node
- HPC Cluster Nodes: Head and Compute
- Ideal for new deployments

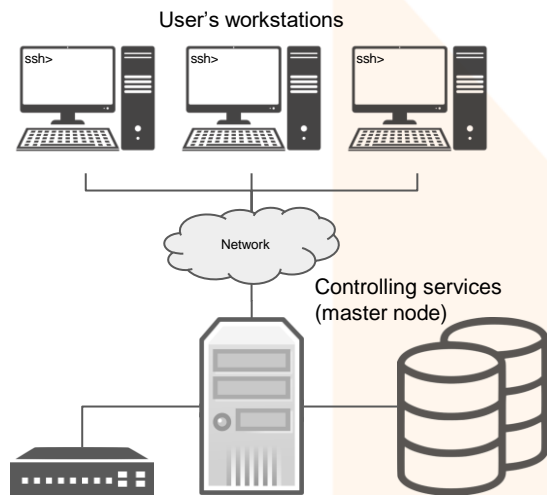
# CIQ Support Levels

| CIQ Support Offerings | Standard  | Premium   |
|-----------------------|---|---|
| Hours of Coverage     | Standard Business Day 8am - 6pm in your local timezone <sup>3</sup> | 24x7 FTS Support <sup>1</sup>   |
| Support Channels      | Web & Email   | Web & Email<br>Named Customer Success Manager<br>Slack Access to CSM <sup>2</sup> |

| Response Times      | Initial and Ongoing Response | Initial and Ongoing Response  |
|---------------------|------------------------------|---|
| Severity 1 (Urgent) | 1 business hour              | Initial Response: 30 minutes<br>Ongoing Response: 1 hour or as agreed |
| Severity 2 (High)   | 4 business hours             | 2 hours or as agreed  |
| Severity 3 (Normal) | 1 business day               | 4 business hours or as agreed   |
| Severity 4 (Low)    | 2 business days              | 1 business days or as agreed  |

# Traditional HPC


## Nearly Every HPC System Today is Based on a Legacy Architecture



The standard HPC cluster architecture today is a near 30 year old design (the “Beowulf”) generally consisting of:

- Controlling services: batch scheduling, login/interactivity, ssh, etc.
- Shared storage for: data, applications, output, etc.
- Private cluster network(s) to connect all compute and storage
- Compute nodes

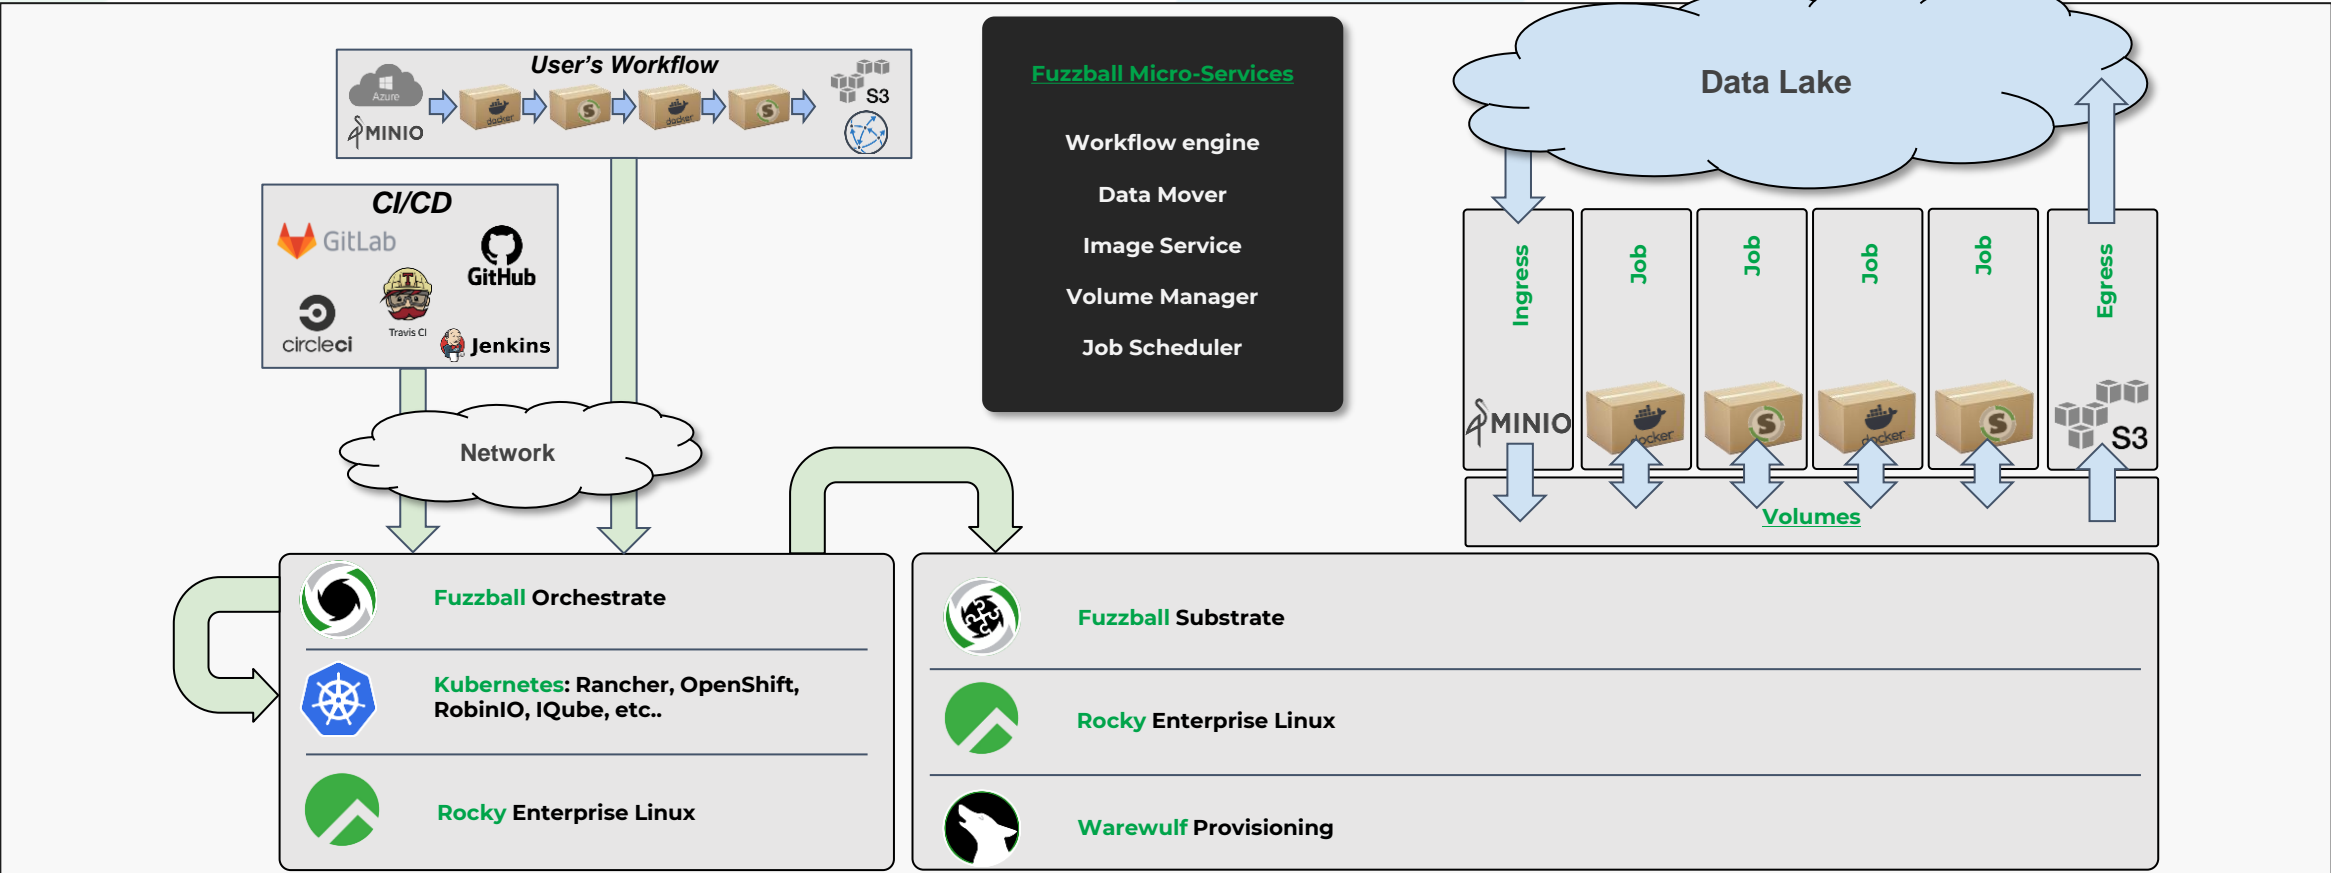
**We have learned to scale up this architecture massively to meet the needs from a simple system to approaching exascale**



Lots of compute nodes

# Fuzzball: Orchestrate

A **Fuzzball cluster** allows users to submit workflows to **Fuzzball Orchestrate**, which then manages the data, container images, and storage volumes while setting up the job graph and compute resources. The jobs are run via **Fuzzball Substrate** which manages those workflows on thousands of compute nodes!



*Management Cluster (1-10s of nodes)*

*Compute Cluster (10s-1000s of nodes)*

# Tested Use Cases / Industry Examples

## GenAI / Tech

- **Meta's Llama 2/3 (LLM), Stable Diffusion (image gen), GPT-like models (LLMs)**
- **Training/inference with major AI frameworks**
  - PyTorch
  - Tensorflow
  - JAX-based models
  - Hyperparameter tuning with W&B
- **Interactive Jupyter notebooks**
- **Salmon and similar Python-based bioscience tools**
- **Specialized Use Cases**
  - GasNET/Chapel language support
  - Exascale Computing Project apps
  - Remote GPU profiling via NVIDIA tools
  - Satellite signal processing submitted via Fuzzball SDK
  - VNCs into programming language debuggers/interfaces like R

## Industry Examples

- **Life Sciences / Pharma**
  - GROMACS (drug discovery)
  - LAMMPS (drug discovery, NL nuclear apps)
  - STAR genomic alignment (genomics)
- **Simulation / Manufacturing**
  - MATLAB
  - Ansys tools
  - Altair OpenRadioss (FEM, structural simulation, etc.)
  - WRF (weather simulation for public power utilities)
  - OpenFOAM w/ Paraview visualization (CFD, aerodynamics, fluid flow)
  - QMCPACK (quantum dynamics)
  - Quantum ESPRESSO (quantum dynamics)
  - Arkouda - Exploratory Data Analysis

CIQ

SA

Workflows

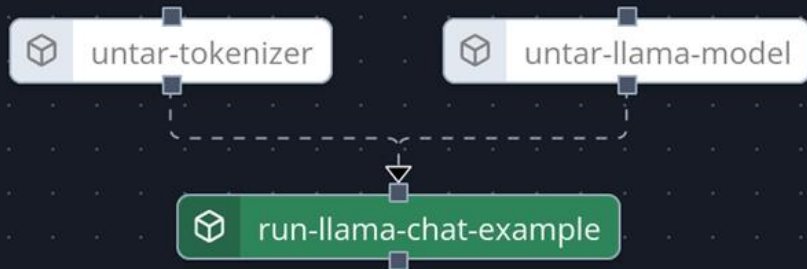
Workflow Editor

ACCOUNT CONFIGURATION

Definitions

Secrets

Users



### Edit run-llama-chat-example

Job Environment Resources

**Name**  
run-llama-chat-example

**Requires**  
untar-tokenizer  
untar-llama-model

**Add Required Job**  
Jobs that must be completed before this job may be started.

**Command**  
/bin/sh  
-c  
torchrun --nproc\_per\_node 1 ./example\_chat\_comple

**Add**  
Command and arguments to be executed in the Job container.

**Policy**  
Timeout Duration: e.g. 2m  
Retry Attempts: [dropdown]

**Job Type**  
Default Multinode Task Array

**Isolated** [toggle]

Force the use of a container network namespace

Delete Close

```
run-llama-chat-example:
  cwd: /llama-farm/llama
  image:
    uri: >-
      oras://us-west1-docker.pkg.dev/ciq-sa/containers/llama-2:apptainer-stable
  secrets:
    password: ${secret "CIQ_SA_RO" }}$
    username: _json_key_base64
  mounts:
    llama-barn:
      location: /llama-barn
  command:
    - /bin/sh
    - '-c'
    - >-
      torchrun --nproc_per_node 1 ./example_chat_completion.py --ckpt_dir
      /llama-barn/llama-2-7b-chat/ --tokenizer_path
      /llama-barn/tokenizer.model --max_seq_len 512 --max_batch_size 6
  requires:
    - untar-tokenizer
    - untar-llama-model
  resource:
    cpu:
      cores: 4
    memory:
      size: 58 GiB
    devices:
      nvidia.com/gpu: 1
```

```
volumes:
  llama-barn:
    type: EPHEMERAL
    ingress:
      - source:
          uri: s3://co-ciq-misc-support/fburt-ai/llama-2-7b-chat.tar.gz
          secrets:
            s3Region: us-east-2
            s3AccessKey: ${secret "FBURT_AWS_KEY" }}$
            s3AccessKeyId: ${secret "FBURT_AWS_KEY_ID" }}$
          destination:
            uri: file://llama-2-7b-chat.tar.gz
      - source:
          uri: s3://co-ciq-misc-support/fburt-ai/tokenizer.tar.gz
          secrets:
            s3Region: us-east-2
            s3AccessKey: ${secret "FBURT_AWS_KEY" }}$
            s3AccessKeyId: ${secret "FBURT_AWS_KEY_ID" }}$
          destination:
            uri: file://tokenizer.tar.gz
```

## Start Workflow?



Name

Optional display name for workflow

Workflow has:

 3 jobs

 1 volumes

Cancel

▶ Start Workflow

# Success!



Workflow **run-llama-chat-example** started successfully.

Done

Go to status →

ciq  
SA

Workflows

Workflow Editor

ACCOUNT CONFIGURATION

Definitions

Secrets

Users

f burt@ciq.co  
f burt@ciq.co

**WORKFLOW OVERVIEW**  
**run-llama-chat-example**  
 11 minutes ago  
 f burt@ciq.co

**CREATE VOLUME**  
**llama-barn**  
 16s

**PULL IMAGE**  
**oras://us-west1-docker.pkg.dev/ciq-sa/containers/llama-2:apptainer-stable**  
 17s

**TRANSFER FILE**  
 s3://co-ciq-misc-support/fburt-ai/llama-2-7b-chat.tar.gz  
 file://llama-2-7b-chat.tar.gz  
 9m 24s

**TRANSFER FILE**  
 s3://co-ciq-misc-support/fburt-ai/tokenizer.tar.gz  
 file://tokenizer.tar.gz  
 0s

**RUN JOB**  
**untar-llama-model**  
 oras://us-west1-docker.pkg.dev/ciq-sa/containers/llama-2:apptainer-stable  
 llama-barn

**RUN JOB**  
**untar-tokenizer**  
 oras://us-west1-docker.pkg.dev/ciq-sa/containers/llama-2:apptainer-stable  
 llama-barn

**RUN JOB**  
**run-llama-chat-example**  
 oras://us-west1-docker.pkg.dev/ciq-sa/containers/llama-2:apptainer-stable  
 llama-barn

Details Definition

Workflow Overview

| STATUS  | CREATED            | STARTED            | FINISHED | DURATION |
|---------|--------------------|--------------------|----------|----------|
| Started | 10/25/23, 12:41 PM | 10/25/23, 12:41 PM | --       | --       |

# Upcoming Events

- CIQ team at Oracle CloudWorld next week
- Dr. Dave Godlove speaking at HPC+AI on Wall Street
  - Sept. 17-19
- Forrest Burt speaking at OpenRadioss Day in Philadelphia
  - Sept. 26
- CIQ team at SC24 in Atlanta
  - Nov. 17-22

**Thank you**