OriginAI

Leveraging a comprehensive reference architecture for AI that significantly reduces time-to-insight

Featuring Scyld ClusterWare® and Red Hat® OpenShift®

Red Hat

PENGUIN COMPUTING

Contents

Solution-at-a-Glance

Features

- NVIDIA DGX A100 Systems
- NVIDIA Mellanox HDR Infiniband Networking
- Penguin Computing ActiveData™ Solution with WekaIO WekaFS
- Red Hat® Ceph® Storage
- Scyld ClusterWare®
- Red Hat® OpenShift®

Benefits

- Reduce time-to-insight.
- Maximize the performance and utility of high-value AI systems.
- Ensure increased productivity and highest ROI.
- Support scalability and flexibility.
- Ensure data security, resilience, and governance without compromising performance.
- Meet the demanding requirements of AI and analytics models.

The Challenge

Increasingly, organizations are recognizing the business potential artificial intelligence (AI), machine learning (ML), and deep learning (DL) represent; business derived from AI is expected to reach $3.9 trillion by 2022 (Gartner, 2018). As organizations move to adopt these technologies, CIOs are faced with the challenge of implementing complex new systems that are performant, bespoke, and secure.

In this highly specialized endeavor, the odds for success are stacked against CIOs: it is predicted that 85% of AI projects will ultimately not deliver for their organizations (Gartner, 2018). In order to be successful, every component of the AI infrastructure must be expertly tuned to support an organization's unique AI workload. Challenged by critical gaps in expertise, organizations are frequently unable to optimize, resulting in data path bottlenecks that render expensive resources idle, leading to lost productivity and increased time to insight. Adding to the complexity, the speed at which AI evolves demands flexible, scalable architectures that can keep pace as technology advances.

For organizations that are new to high performance computing (HPC) and high performance data analytics (HPDA) systems, developing AI infrastructure can be costly and time consuming, taking an estimated two to three years of internal research to bring AI-based products to market. In fact, 78% of AI projects are stalled even before they can be deployed (Dimensional Research, 2019).

Realizing a need for a comprehensive AI solution that encompasses architectural design, hardware and software services, hosting, and deployment, CIOs are increasingly seeking the expertise of single solution providers.

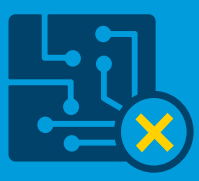

78% of AI projects are stalled even before they can be deployed.

Penguin Computing OriginAI

Penguin Computing partnered with WekaIO, Inc., under the Weka Within™ Certification Program, to create OriginAI™, a comprehensive, end-to-end solution for datacenter AI that breaks down software, server/switch hardware, data storage and governance, and infrastructure barriers that cause AI projects to fail or stall.

A full-service consultancy, Penguin Computing's Analytics Practice acts as a single point of reference for hardware, software, architectural design, hosting, and management, enabling organizations to focus their time and resources on the business and human challenges in bringing AI projects to production. By providing organizations with a comprehensive solution, Penguin Computing radically reduces the time to insight from years to months.

Penguin Computing OriginAI Rack Layout

Penguin Computing OriginAI Components

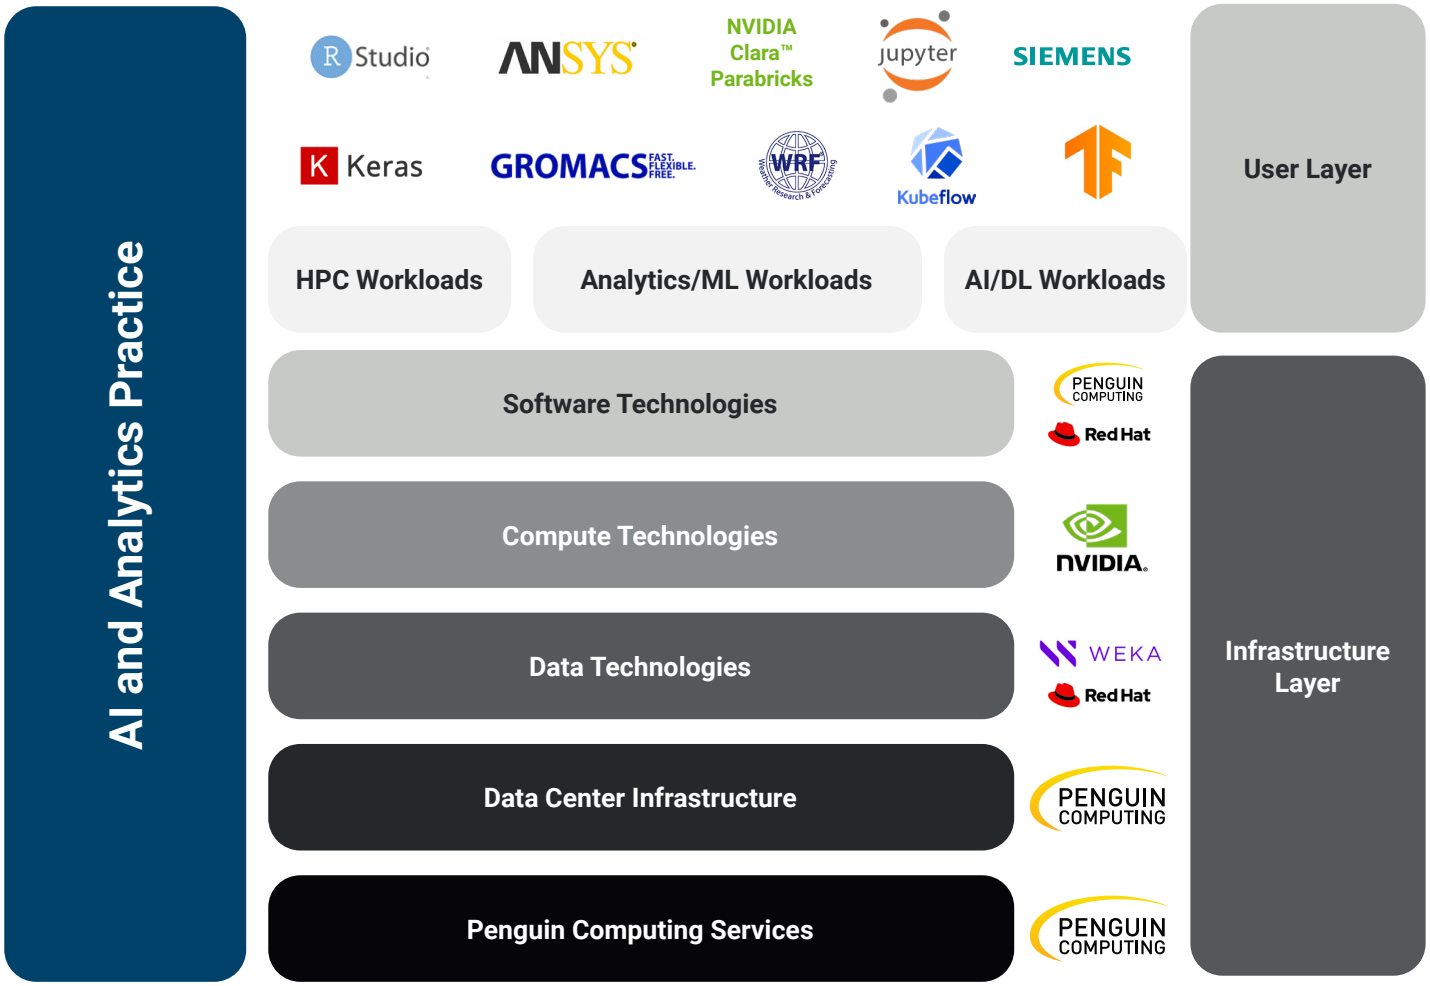

Software Technologies

Scyld ClusterWare

OriginAI leverages Penguin Computing's Scyld ClusterWare cluster orchestration software. Scyld ClusterWare provides a complete HPC environment that supports Slurm, OpenPBS, and TORQUE to handle the scheduling and queueing of HPC jobs. Scyld ClusterWare also provisions the hardware to operate as a single, unified cluster by booting compute nodes using PXE, establishing IP address using DHCP, monitoring node health, and collecting metric data across the cluster.

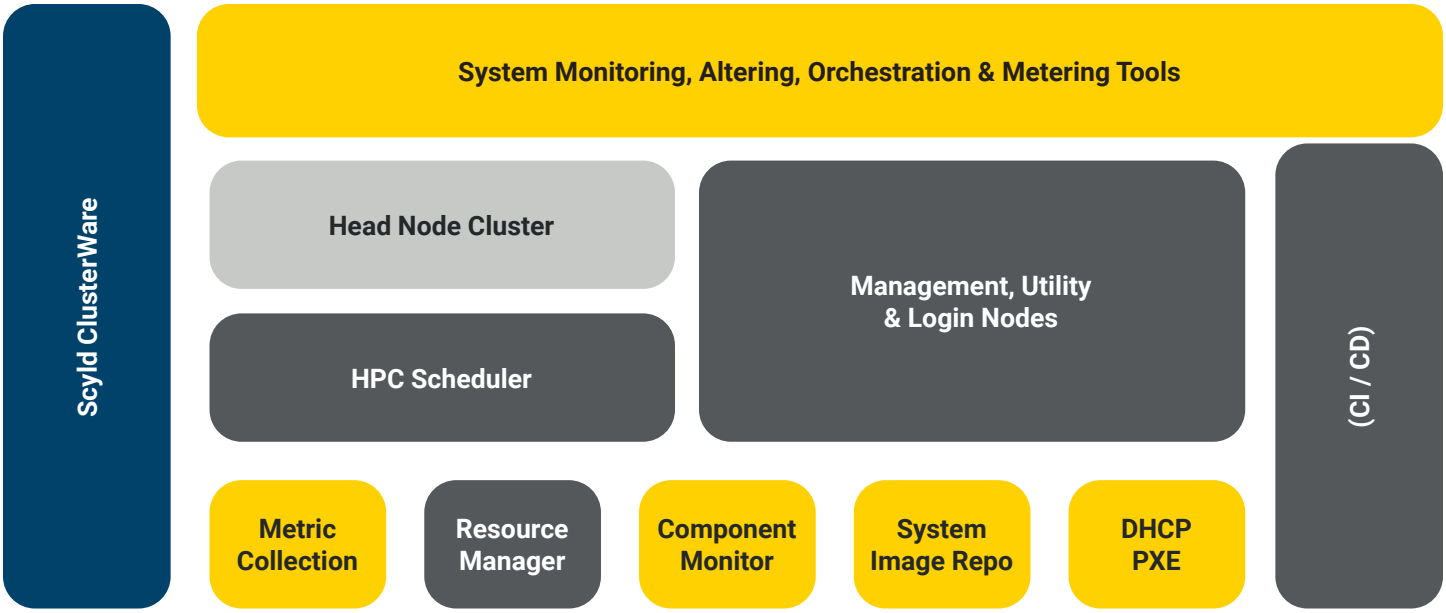

Alerting and Monitoring

Scyld ClusterWare supports alerting features that can integrate into enterprise communication tools, such as email, Slack, PagerDuty, and more to send out important alerts to different groups within an organization regarding the current status of an OriginAI implementation. This information is also centralized to the head node, providing system administrators with a single pane of glass displaying the status of every node in the cluster.

Cluster Customization

Scyld ClusterWare provides a robust HPC orchestration capability for system administrators to provision and manage an HPC environment. This capability has been designed to provide a familiar, intuitive, and documented method for adjusting the cluster and providing change management. Scyld ClusterWare enables experienced Linux administrators to use familiar tools, methodologies, and best practices to customize their OriginAI solution.

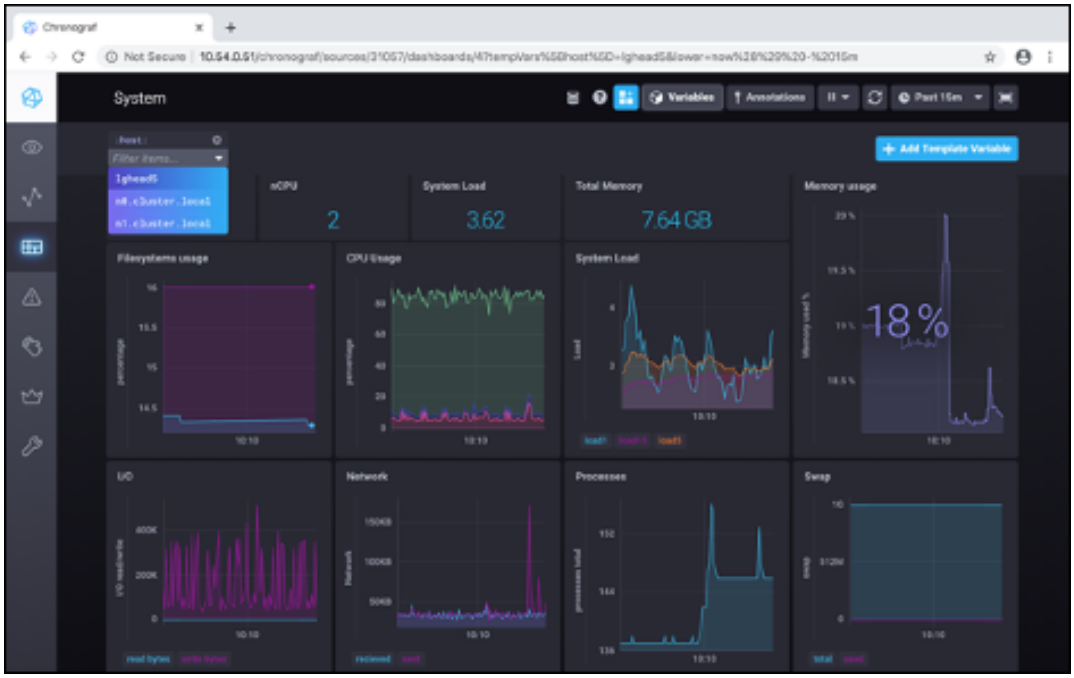

System Image by Application

Scyld ClusterWare also supports custom system image deployment. Users can save a compute node image into a repository managed by the Scyld ClusterWare head node. System images can be completely different from the operating system that the head node uses. The head node could be running RHEL 8, while the compute nodes are running RHEL 6, RHEL 7, RHEL 8, Ubuntu 16.04, Ubuntu 18.04, or some combination of operating systems and system versions across the cluster.

Eliminating Compute Silos

Scyld ClusterWare's ability to dynamically deploy system images per job allows system administrators to centralize their compute infrastructure and eliminate compute silos that would otherwise require a dedicated workstation or compute cluster to support very specific operating systems that some number of user applications would require. The ability to centralize compute resources into one dynamic environment relieves administrative headaches and allows IT teams to focus their efforts on supporting a centralized environment where updates, patches, fixes, and upgrades have a greater impact across their organization.

High Availability

Scyld ClusterWare provides a robust high availability HPC cluster, that allows the system to remain operational when unexpected failures occur, and during scheduled maintenance windows. The HA feature of Scyld ClusterWare enables system administrators to upgrade a kernel or reboot a head node without disrupting the functional use of the environment to end users.

Enterprise Security

Scyld ClusterWare supports deploying compute nodes with SELinux enforced, and with an MLS policy enabled. This allows Scyld ClusterWare to be deployed in a manner that supports FIPS 140-2 requirements. OriginAI is tested and delivered as an integrated solution with security and support as paramount components.

Scyld ClusterWare Advantages

- Rapid provisioning for technical computing environments
- Designed to manage optimized HPC clusters and their coupled enterprise services
- Single source for tested HPC middleware (MPI implementations, HPC schedulers)
- Image based node management facilitates simplified change management
- Flexible provisioning options (for example, diskless, diskfull, network mounted)
- Robust high availability architecture prevents downtime when unexpected failures occur
- Supports SELinux in MLS mode and FIPS 140-2 implementations
- Monitoring GUI for visualizing system telemetry and building custom dashboards
- Notification and alerting integration with email, Slack, and PagerDuty

Red Hat OpenShift

Red Hat® OpenShift® provides a container-based application development platform, built on massively scalable cloud native infrastructure, all managed through a common management framework. Customers can move existing workloads to scale-out cloud infrastructure and accelerate new cloud-based services for private cloud and application development. With OriginAI with Red Hat OpenShift, an operations team can deliver public cloud-like services on premises to developers and business units while maintaining control and visibility. Red Hat OpenShift enables a unified management framework to truly enable workload portability edge-to-core strategies to deliver the same user experiences on-premises or in public cloud environments.

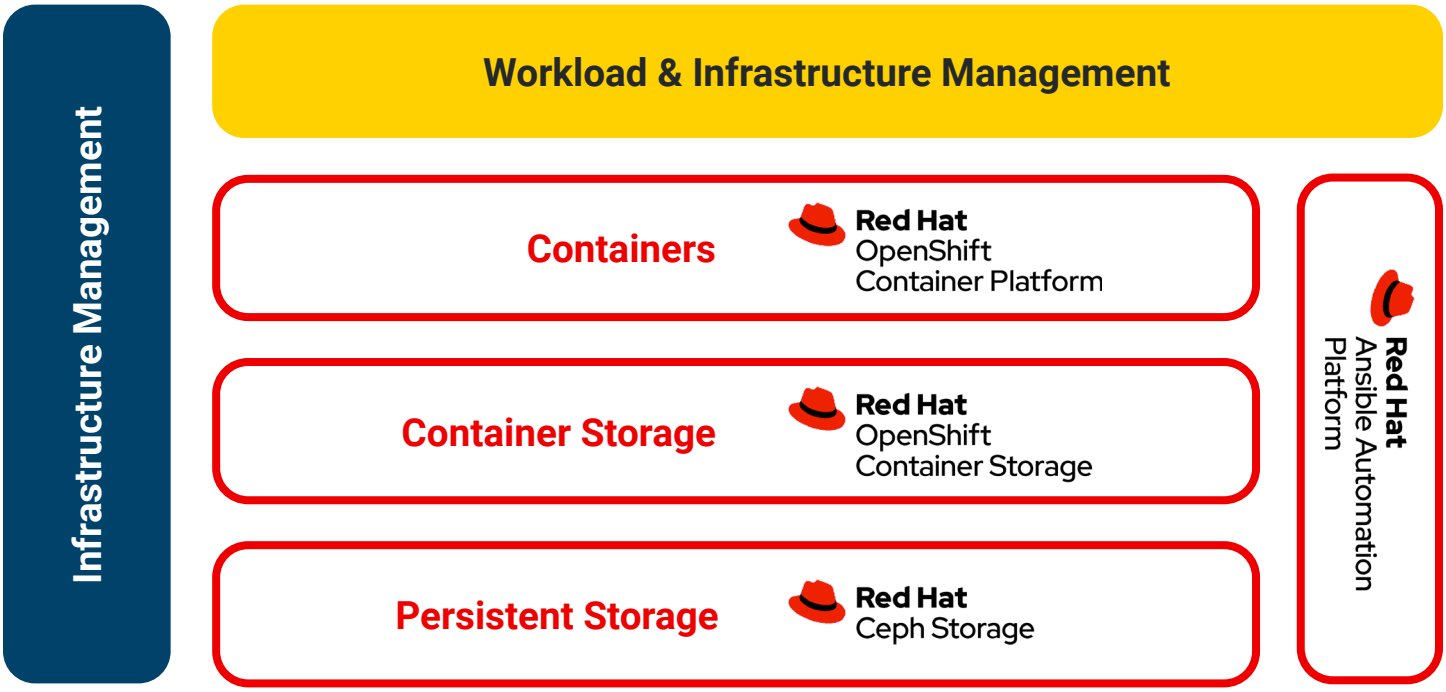

At its infrastructure foundation, OriginAI leverages Red Hat OpenShift to build a private, cloud-based Red Hat Enterprise Virtualization based on high-performance virtualization. This provides a secure, scalable foundation for hosting the Openshift Container Platform. OpenShift automates the development and administration of container-based applications.

Containerized, Cloud-Native Workload Portability

As cloud-native technologies like containers and Kubernetes mature rapidly, they are quickly becoming the preferred way to build new software experiences and modernize existing applications, workloads, and workflows at scale and across on premises to public clouds and multi-cloud. Creating value depends on the ability to deliver applications and workloads faster. This is being driven by the explosion of data-driven workloads in AI/ML, Analytics, IoT, and other emerging technologies. Cloud-native technologies are driving this innovation culture. Enterprise customers now seek container development platforms that accelerate and simplify the development and operations (DevOps) of cloud-native apps wherever and however firms build and deploy them. OriginAI not only provides comprehensive container infrastructure lifecycle operations from the data center to the cloud to the edge, it also helps developers modernize apps and innovate workloads with integrated service catalogs and microservices, service mesh, and serverless features.

OriginAI powered by the Red Hat OpenShift Container Platform delivers a balanced blend of development and operations features that:

- **Simplify cloud-native app development with rich development services** — OriginAI draws developers in and jump-starts both development and app modernization with microservices frameworks, serverless support, continuous integration and delivery (CI/CD) integrations, dependency management, and app lifecycle management features like code quality checks and vulnerability scanning. OriginAI helps developers focus on business logic with comprehensive service catalogs and prebuilt DevOps automations and integration.
- **Enable distributed infrastructure operations from data center to cloud to edge** — Data-driven workloads are increasingly distributed and hybrid. OriginAI offers model-driven configuration, monitoring, security, and cluster lifecycle features for unified multi-cloud cluster operations. OriginAI extends operational control to the edge and supports thousands of clusters.

Customer Benefits

OriginAI with Red Hat OpenShift provides a common interface and technology stack for operations, IT administration, development, and lines of business. OriginAI includes a management framework across application development and infrastructure layers, along with complete operation and lifecycle management as well as proactive risk mitigation.

Developers use OriginAI to gain rapid access to compute power and to continuously integrate and deploy applications across a broad range of languages and frameworks. Operators can simultaneously monitor and govern these services and applications across a hybrid infrastructure, from development to production. CIOs can now align better with business requirements, meeting market needs and customer requirements.

Red Hat OpenShift Advantages

- **Integrated components** — Tightly integrated, fully supported components that act in concert to provide an open hybrid cloud.
- **Unified management** — A single management framework across infrastructure and application development layers, plus complete operation and lifecycle management with proactive risk mitigation.
- **Full-featured application development and containers** — Flexible container-based application development abilities via OpenShift Enterprise.
- **Open and interoperable** — Leveraging Red Hat's open API exposure, OriginAI allows customers to enhance or replace existing components with their choice of existing technologies, enabling a true open technology approach.
- **Massive ecosystem** — It's easy to add networking, storage, and other cloud-native solutions, including Penguin Computing, Red Hat, and third-party products. In addition, leveraging Penguin engineered and optimized hardware to deliver a complete solution, OriginAI integrates into your on-premises environment and monitors activity in public clouds.

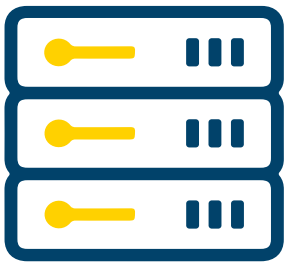

Compute Technologies

Penguin Computing partners with industry leaders such as WekaIO to design and build workload-specific computing platforms that accelerate time to insight. The design process includes in-system device-to-device bandwidth and resource optimizations and system-to-system communication optimizations. These design choices ensure that the underlying computing and networking infrastructure are optimized for the workloads that customers will run on them.

The NVIDIA DGX A100 system is the universal system for all AI infrastructure, from analytics to training to inference. Combining the DGX A100 system with the NVIDIA Mellanox HDR Infiniband Network enables in-network computing through the Co-Design Scalable Hierarchical Aggregation and Reduction Protocol (SHARP)™ technology results in an order of magnitude of application performance improvements.

NVIDIA DGX A100 System

- Universal system for all AI infrastructure, from analytics to training to inference
- 5 petaFLOPS of performance in a 6U form factor, setting a new bar for compute density
- Eight integrated A100 GPUs for unprecedented acceleration
- Optimized for NVIDIA CUDA-X™ software

NVIDIA Mellanox QM8700 HDR Infiniband

- 16Tb/s of non-blocking bandwidth with sub 130ns port-to-port latency
- Forty 200Gb/s full bi-directional bandwidth ports
- Scalable Hierarchical Aggregation and Reduction Protocol (SHARP) technology enables in-network computing to accelerate communications frameworks, resulting in order of magnitude application performance improvements

Data Technologies

As data-intensive workloads scale, it's critical to implement data-driven, software-defined architectures that meet the demands of large data sets. Optimized, accelerated data platforms promise an immediate and tangible solution for delivering discovery and insight from machine-generated data. These platforms combined with accelerated compute and the right software create a new storage category. This approach provides a data store that delivers an enterprise-ready, unified data platform that performs across the entire environment, while also providing essential data security, resilience, and governance. This type of platform is a requirement for data management in the era of AI.

Artificial Intelligence Data Life Cycle

WekaIO WekaFS

Addressing storage I/O requirements at each stage in the AI data pipeline, the Penguin Computing ActiveData™ solution with WekaFS delivers massive bandwidth for ingest and training, ultra-low latency for improved inferencing, and storage features to manage data workflows.

WekaFS is the world's fastest shared parallel file system. As the only shared storage solution that provides end-to-end data management for data-intensive AI workloads, WekaFS enables companies to scale performance across the GPU cluster.

- Delivers unmatched performance at any scale while offering the same enterprise features and benefits of traditional storage
- Prevents GPU starvation, easily meeting the I/O requirements of the most demanding AI and analytics models
- Delivers 10x the performance of network attached storage (NAS) systems and 3x the performance of local server storage

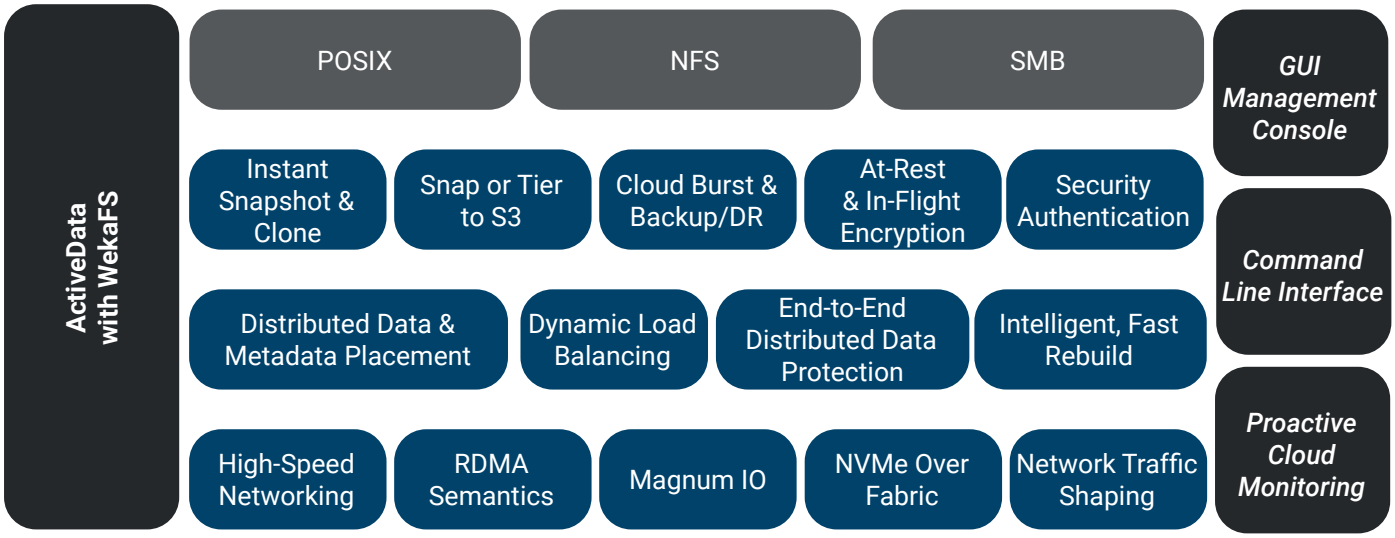

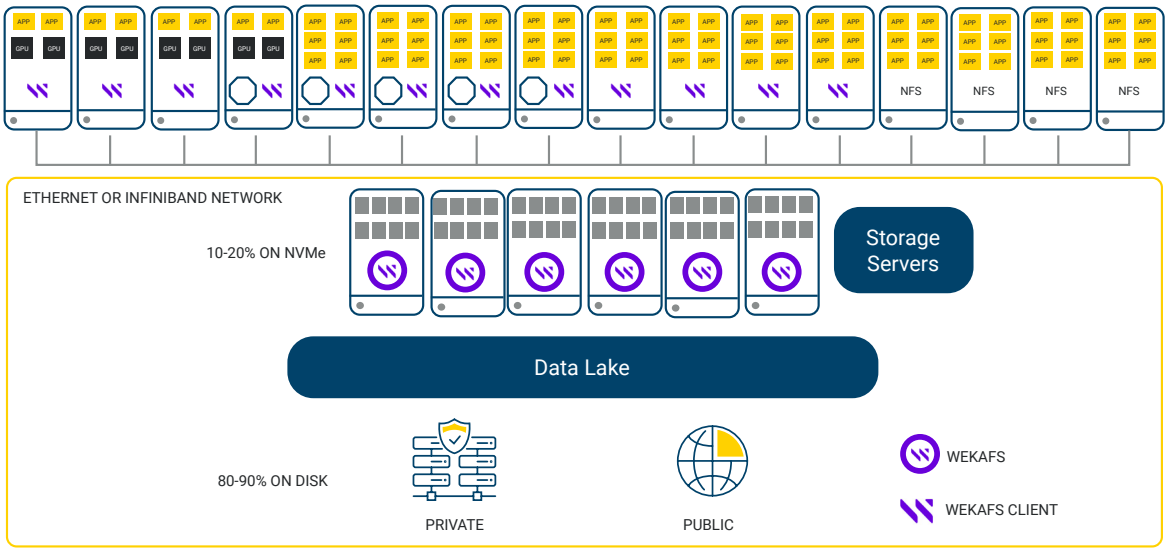

WekaFS Advantages

- Balanced flash storage and high-performance networking for predictable performance that scales
- Multi-workflow optimized storage with massive bandwidth for ingest and training, ultra-low latency for improved inferencing, and storage features to manage data workflows
- Can be deployed as a dedicated storage server (appliance model) or integrated into the application cluster (converged)
- Supports a hybrid cloud model, allowing enterprises to leverage on-demand public compute resources for cloud-bursting, remote backup, and disaster recovery
- Penguin Computing provides a single point of support for quick problem resolution

Red Hat Ceph Storage

Users of Origin AI with Red Hat® Ceph® Storage, can leverage the Ceph Storage open, massively scalable, highly flexible software-defined storage system that is now bundled with Red Hat OpenShift Platform. Red Hat Ceph Storage is expertly architected and integrated into complete turnkey solutions on optimized Penguin hardware. Red Hat Ceph Storage is overwhelmingly preferred by OpenShift users because of its seamless integration with OpenShift's architecture and storage components.

Cost-Effective, Petabyte-Scale Object and Block Storage

Red Hat Ceph Storage is an award-winning, production-ready implementation of Ceph that is optimized for large installations (typically a petabyte or greater) and offers both object and block storage, ideal for data lakes, elastic media storage, backup and recovery, and virtualization platforms. Red Hat Ceph Storage eliminates single points of failure and protects data through dynamic data distribution. The system is both self-healing and self-managing, and this reduces administrative overhead costs.

Red Hat Ceph Storage is designed to provide the data protection, reliability, and availability required for demanding, web-scale object storage and flexibility to support a wide variety of workloads and use cases.

Red Hat Ceph Storage Advantages

- **Flexibility** — Supports object and block storage with application programming interfaces (APIs) for Red Hat Ceph storage native object access protocols as well as OpenStack Swift and S3.
- **Elasticity —** Storage pools are abstracted from hardware, so each can be managed independently. Grow or shrink storage pools without down time and easily scale-out to exabyte capacities single management framework across infrastructure and application development layers, plus complete operation and lifecycle management with proactive risk mitigation.
- **Reliability and availability —** Striping, erasure coding, replication, and snapshots, along with self-healing and automatic rebalancing ensure data protection and resilience. lexible container-based application development abilities via OpenShift Enterprise.
- **Simplified administration —** Manage and monitor the entire storage cluster from an easy-to-use graphical console, command-line interface, or programmatically everaging Red Hat's open API exposure,
- **Performance —** Tune with storage policies that reflect service level agreements, performance requirements, and failure.
- **Security —** Integrate with authentication and authorization services, such as Active Directory, Lightweight Directory Access Protocol, and Keystone. Limit access at the pool, user, bucket, or data level. Implement cluster-level and at-rest encryption. une with storage policies that reflect service level agreements, performance requirements, and failure.

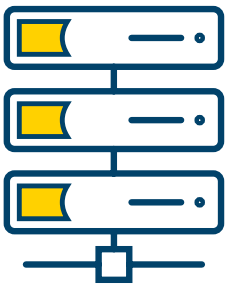

Data Center Infrastructure

OriginAI EIA racks can be deployed in 42U or 48U high x 750mm wide x 1200mm deep options. This enables quick and painless serviceability by ensuring there is sufficient space around and behind each system to install, remove, or cable them properly.

Penguin Computing has partnered with leading data center facility pioneers who can support the demanding characteristics of today's HPC platforms.

Power

In an Origin EIA rack, Power Distribution Units (PDUs) can be installed vertically at the rear of the rack and mounted with the LED display facing to the rear. This allows system administrators to quickly view the power status of each PDU at a glance. A fully populated EIA rack can require 30 to 60 kW of power, requiring two to four PDUs to ensure adequate rack-level power.

N+1 Power Redundancy: Power redundancy is provided at the system level. Every system uses multiple Power Supply Units (PSUs) to ensure there are N+1 PSUs per system, where N is the number of PSUs needed to fully power the system, and an additional PSU is installed to account for at least 1 PSU failure without impacting the uptime of the system.

Cooling

Removing the heat from an OriginAI rack can be a challenge for most data centers, especially if they are not set up to handle HPC environments. Most data centers use HVAC systems to pump low-temperature air into the data center to prevent the environment from overheating.

Rear Door Heat Exchanger: Although an OriginAI rack can be cooled using traditional air conditioning, we highly recommend using a cooling method that better targets the heat, such as rear door heat exchangers that can be mounted to the rear of each rack. Rear door heat exchangers work by capturing heat as it leaves the rack, ensuring the exhaust air temperature is net neutral to the ambient air temperature in the data center.

Penguin Computing Services

OriginAI s a comprehensive, end-to-end solution that organizations can leverage to jump-start their AI, ML, and DL initiatives. In some cases, the reference architecture will directly meet the needs of the organization, right out of the box. However, most often there will be additional design, deployment, integration, and hosting considerations that need to be addressed.

Penguin Computing provides services that consider rack and floor space, how to scale the environment, maximum rack power consumption, power phase balance, efficient heat removal, and the optimal networking topologies when using low-latency, high throughput interconnects.

OriginAI is supported by Penguin Computing engineering services, including Design Services, Professional Services, Managed Services, and Hosting Services

Data center hosting services are offered through Penguin Computing's strong partnerships with data center service providers. Our partners can provide the space, power, and cooling that OriginAI needs - as a service.

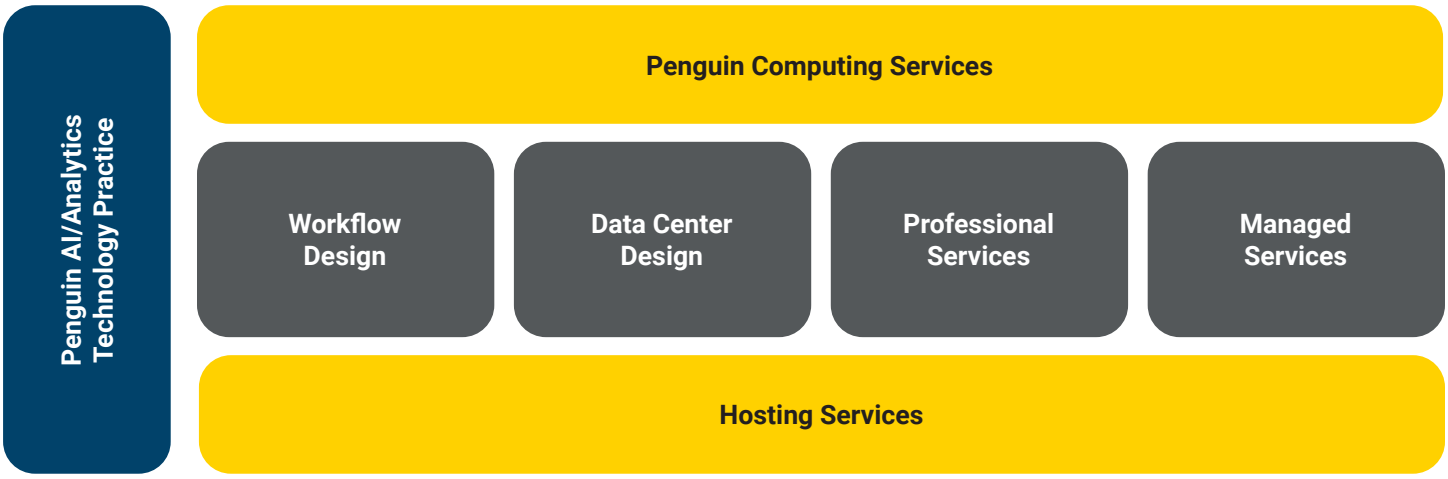

Design Services

Workflow Design

- Software Orchestration
- Compute Performance
- Multi-Node Communication
- Data Storage and Data Tiering
- Data Ingest and Egest
- Environment Sizing

Data Center Design

- Rack and Floor Space
- Environment Scalability
- Maximum Power Consumption
- Power Phase Balance
- Efficient Cooling and Heat Removal
- Optimal Networking Topologies

Professional Services

Stand Up and Initialization

- System Burn-In Testing
- Racking and Cabling
- Software Installation & Tuning
- On-Site Deployment and Integration

Hosting Services

Data Center Hosting

- Penguin Data Center
- Customer Data Center
- Power, Space, and Cooling Management
- Monthly or Annual Billing (As-A-Service)

Managed Services

System Administration:

- Complete Hands-Off Experience
- Augment Existing IT Capabilities
- Collaborate with Penguin Support
- Tens to Thousands of Servers
- Terabytes to Exabytes of Data
- Multi Data Center Support

Conclusion

Penguin Computing OriginAI provides a single, secure, end-to-end solution that includes architectural design, hardware and software services, hosting, and deployment. Penguin Computing OriginAI frees private and public sector organizations from having to focus valuable time and human resources on creating an architecture from scratch, delivering cost savings, decreasing risk, and accelerating time to insight.

Penguin Computing can apply our decades of experience to create quality, integrated solutions for our clients. We offer a wide range of professional and managed services that can quickly bring your artificial intelligence, machine learning, and deep learning initiatives products to production.

Contact Us

Use this **form** or call Penguin Computing today at 1-888-736-4846 to find out how you can jump start your artificial intelligence, machine learning, and deep learning initiatives with a reference architecture that addresses software orchestration, compute infrastructure, Data, and infrastructure design to:

- Accelerate time to insight from years to months
- Maximize the performance and utility of high-value AI systems
- Ensure increased productivity and highest ROI

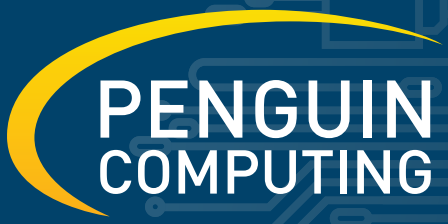

Expanding the world's vision of what is possible