



ActiveData with WekaFS

Prevent data bottlenecks and speed time-to-market
with the world's fastest shared file system



Contents

- Solution-at-a-Glance 4**
- Introduction 4**
- Penguin Computing ActiveData 5**
 - Data Technologies 6
 - WekaIO WekaFS 6
 - Accelerate Performance. 7
 - Multi-Tier Single Namespace 7
 - Space Without Limits 8
 - Enterprise Engineered High Performance Storage. 8
 - Software Technologies 9
 - Software Defined Architecture for High Performance Storage 10
 - Persistence Storage Volumes with Kubernetes CSI Plugin. 10
 - Compute Technologies 11
 - High Performance Workloads 11
 - High Performance Data Requirements 11
 - Data Center Infrastructure. 12
 - Power 12
 - Cooling 12
 - Penguin Computing Services 13
 - Design Services 14
 - Professional Services 14
 - Managed Services. 14
 - Hosting Services. 14
- Conclusion 15**
- Contact Us. 15**

Solution-at-a-Glance

Features

- Balanced high-performance flash storage for predictable performance that scales.
- Workload-optimized storage with high bandwidth and low latency.
- Flexibility for dedicated or converged deployment models.
- Supports bare-metal, containerized, virtual, cloud and hybrid deployments.

Benefits

- Reduce operating costs.
- Speed time-to-market.
- Dramatically reduce TCO with industry standard servers in a software-defined architecture.
- Improve employee productivity.
- Scale from as small as 30TB to hundreds 100's of petabytes in a single namespace.
- Optimize scale and cost with flash and disk combinations.
- Improve performance even in mixed small and large file workloads.
- Keep data safe from threat or rogue actors with encryption and authentication.
- Run seamlessly on premises, in the cloud, or both.
- Push backups straight to cloud for long-term retention.

Introduction

Data is being generated in larger volumes and at faster rates than ever before. The amount of data created over the next three years will be more than the data created over the past 30 years, and the world will create more than three times the data over the next five years than it did in the previous five.

[IDC 2020](#) IDC also predicts that the Global Datasphere will grow from 33 Zettabytes in 2018 to 175 Zettabytes by 2025. [IDC 2025](#)

This explosive growth in data creation can lead to congestion, I/O bottlenecks, storage outages, and cost overruns for High Performance Computing (HPC), High Performance Data Analytics (HPDA), Artificial Intelligence (AI), and Machine Learning (ML) workloads. The data analysis performance gap has widened with the continued densification of compute to support today's workloads. While network and compute can be virtualized to operate at scale very effectively, storage has remained isolated in silos that impede modern workflows.

As a result of these challenges, important business insights remain out of reach of decision makers. In order to extract these valuable insights, it is critical to implement innovative software-defined architectures and tailored data pipelines that meet workload demands. These architectures are categorized into tiers based on characteristics such as performance, capacity, connectivity, and cost. They can be combined in a variety of data tiering strategies to efficiently optimize a complete data pipeline for the unique requirements of an organization's workloads.

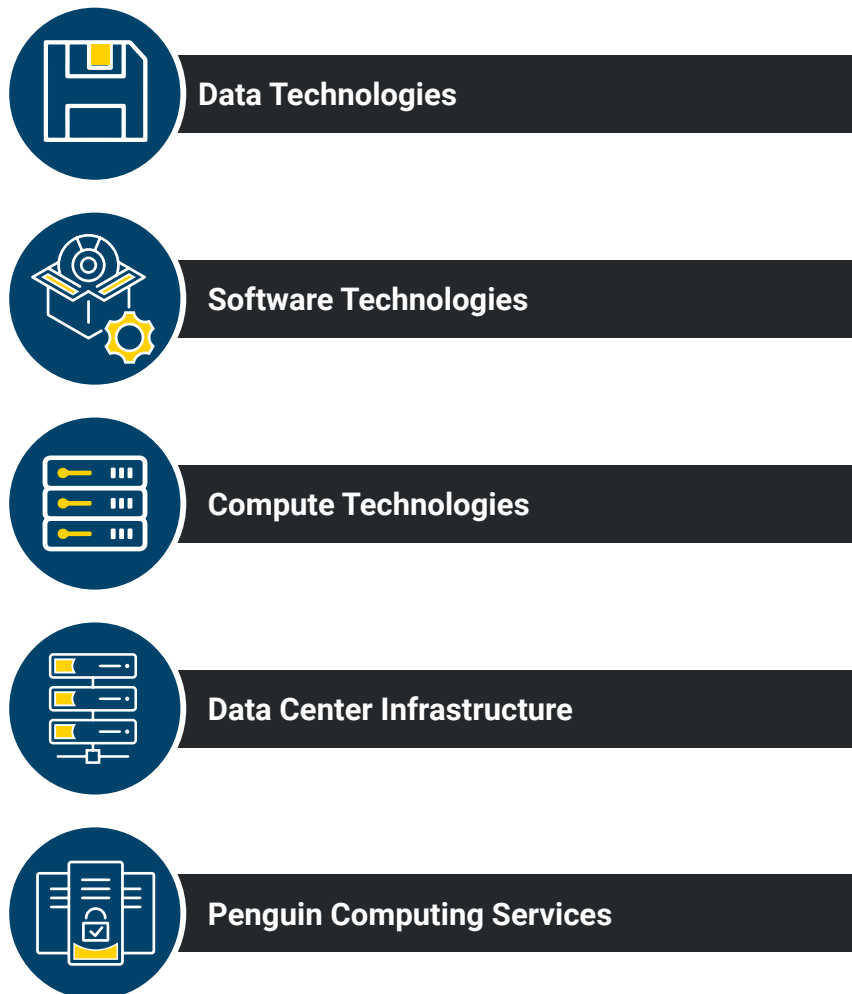
A typical data-intensive pipeline will benefit from a fast, in-memory tier for the most demanding workloads, a second fast and scalable NVMe tier for staging and at-scale operations, and a capacity optimized tier for long-term storage.

Penguin Computing ActiveData

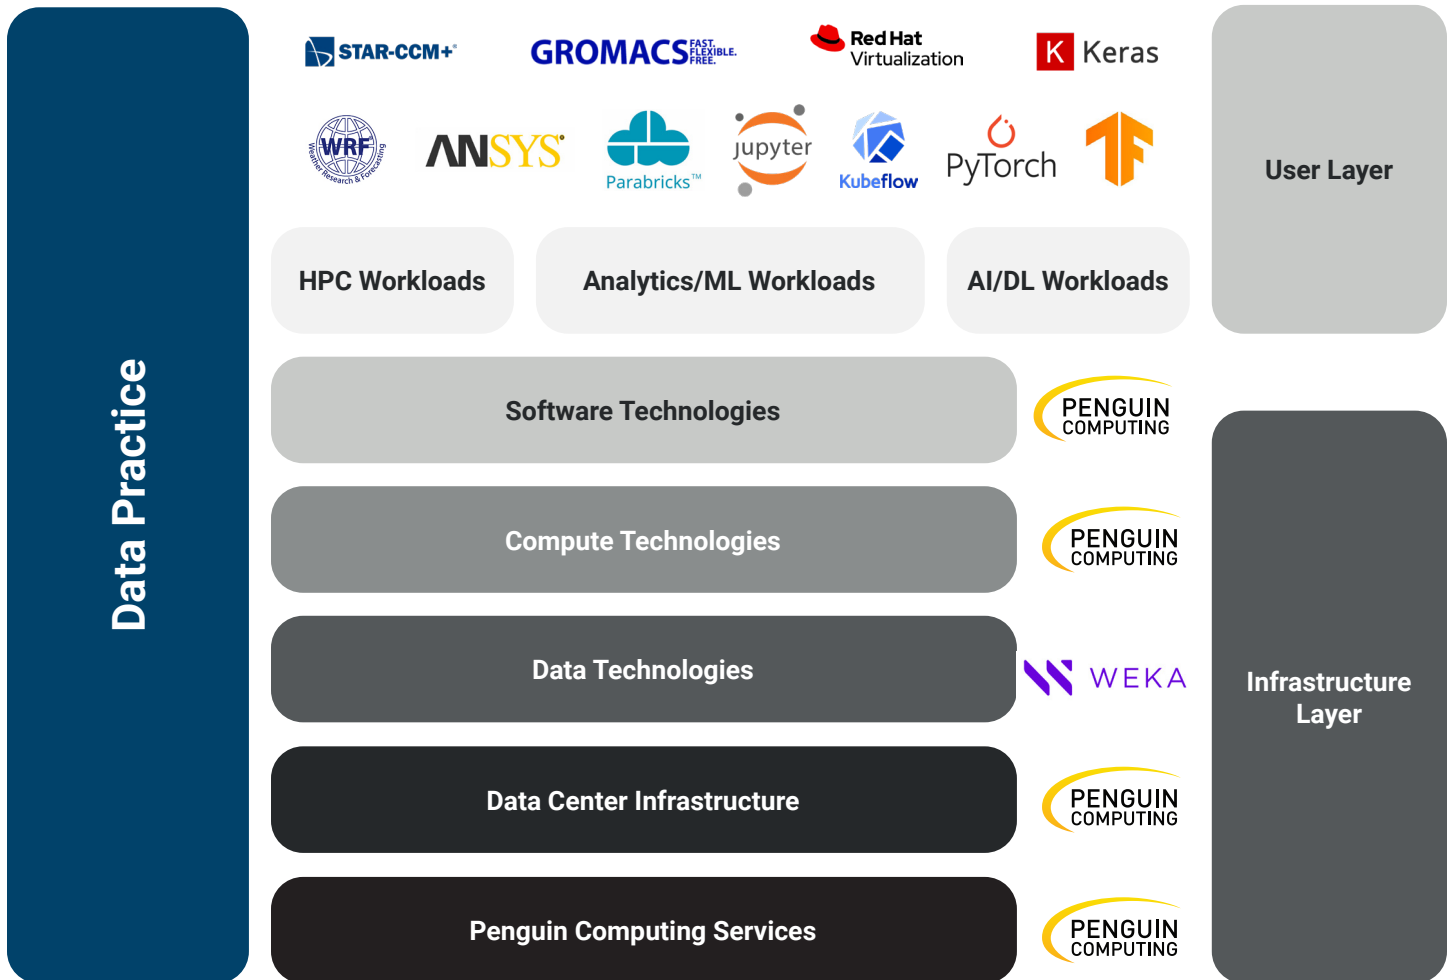
Penguin Computing has partnered with WekaIO, Inc., to offer Penguin Computing ActiveData™ with WekaFS, a storage solution that solves data acceleration challenges. ActiveData with WekaFS is a complete turnkey high-performance file-based storage solution that is highly scalable and easy to deploy, configure, manage, and expand. The design philosophy behind the WekaFS is to create a single storage architecture with the performance of all-flash arrays, the simplicity of network-attached storage (NAS), and the scalability of the cloud.

ActiveData with WekaFS can be integrated into your existing bare-metal, containerized, virtual, or cloud environments. It can be implemented alone as a solution or in combination with other Penguin Computing solutions for Data, HPC, AI/Analytics, and Cloud to provide an end-to-end complete compute platform.

ActiveData includes:



Penguin Computing ActiveData Components



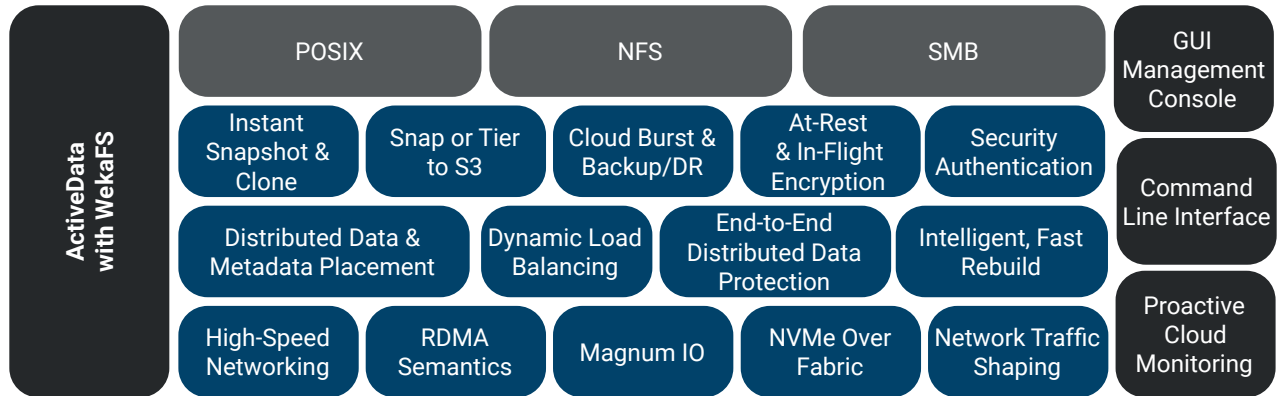
Data Technologies

ActiveData with WekaFS addresses the unique challenges of high IO storage workloads with the scalability needed to tame ever-growing data sets. ActiveData with WekaFS delivers massive bandwidth for ingest and training, ultra-low latency for improved inferencing, and storage features to manage data workflows.

WekaIO WekaFS

WekaFS is the world's fastest shared parallel file system and it delivers unmatched performance at any scale while offering the same enterprise features and benefits of traditional storage appliances. WekaFS is built for those who solve big problems, and runs on Penguin Computing engineered and optimized hardware and cloud infrastructure, whether on-premises, in the cloud, or shared across both. It gives you control of hardware choices and achieves dramatic cost savings for IT infrastructure.

WekaFS is an enterprise grade modern file system supporting multiple workloads using multiple protocols that is highly scalable and easy to manage. Flexible management through GUI or CLI along with proactive monitoring provides operational efficiency, allowing a storage admin to manage storage clusters that scale from terabytes to petabytes to exabytes or more.




Accelerate Performance

WekaFS eliminates any CPU or GPU performance bottlenecks and scales without limits or increased latency. It delivers 10 times the performance of legacy network attached storage (NAS) and 3 times the performance of local server storage. WekaFS can achieve 10s of millions of IOPS at sub-250 microsecond latency.

Multi-Tier Single Namespace

WekaFS can provide a single namespace for a local HPC cluster and an on-premise or cloud-native S3 accessible data lake. WekaFS can automatically tier between a flash-based high performance storage system and a high capacity spinning disk-based storage system to simplify the storage environment. Applications can mount WekaFS and access petabytes of flash or exabytes of spinning disk without ever knowing the difference. The ability to wrap multiple tiers of storage into a single namespace allows organizations to centralize their storage and eliminate data silos.



Features

- Balanced high-performance flash storage for predictable performance that scales.
- Workload-optimized storage with high bandwidth and low latency.
- Flexibility for dedicated or converged deployment models.
- Supports bare-metal, containerized, virtual, cloud and hybrid deployments.
- Penguin Computing provides a single point of support for quick problem resolution.

Scale Without Limits

With an integrated cloud tiering system and simple, intuitive management, WekaFS can scale to hundreds of petabytes in a single namespace – billions of directories and trillions of files. Flexible deployment across public and private cloud ensures quick response to changing needs.

Enterprise Engineered High Performance Storage

ActiveData with WekaFS is optimized to leverage the speed and low latency of NVMe technology in optimized server building blocks and supports both small and large file access – either randomly or sequentially – at the lowest latency. ActiveData is engineered for balanced internal and external IO performance.



ActiveData with WekaFS Server

NVMe drives have full bandwidth to high speed network devices in the system across the same CPU complex to ensure the data path is optimized for the lowest latency and highest throughput possible. ActiveData with WekaFS is ideal for performance-intensive workloads such as artificial intelligence, deep learning, data analytics, life sciences research, financial modeling, engineering development, government/university research, and more.

ActiveData with WekaFS can be deployed as a turnkey solution or in combination with other Penguin Computing solutions for Data, HPC, Cloud, and AI/Analytics. The solution also supports a hybrid cloud model, allowing enterprises to leverage on-demand public compute resources for cloud-bursting, remote backup, and disaster recovery.



Software Technologies

Software-Defined Architecture for High Performance Storage

ActiveData with WekaFS integrates and works with your workload-optimized software infrastructure. Software-defined architectures are key to delivering a flexible data center that allows organizations to deliver accelerated data-intensive workloads, decreasing time to insight and discovery.

Weka's parallel file system is designed to provide a cloud-like experience, whether you run your applications on-premises or move them to the cloud. WekaFS provides seamless operation in today's hybrid and multi-cloud architectures.

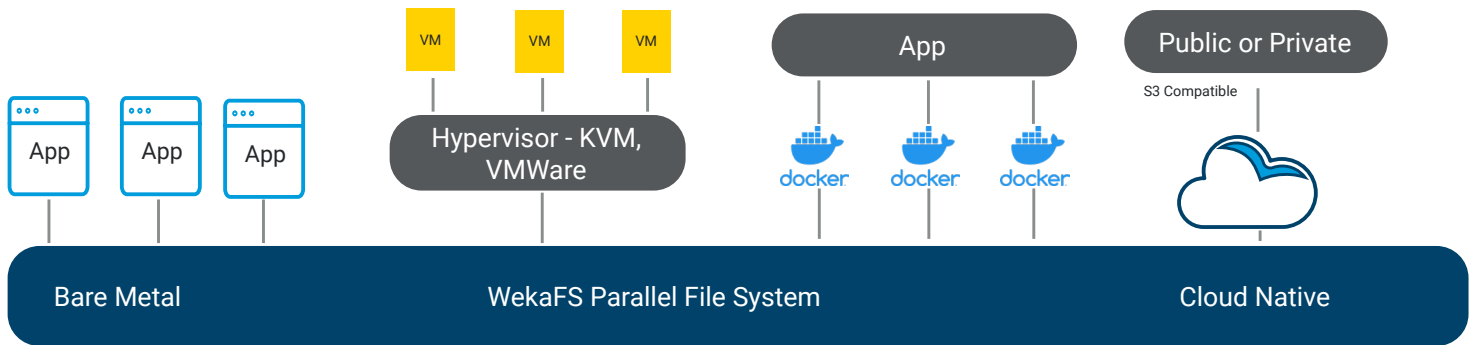
Most legacy parallel file systems overlay file management software on top of block storage, creating a layered architecture that limits performance. WekaFS is a distributed, parallel file system that eliminates the traditional block-volume layer managing underlying storage resources. This vertically integrated architecture does not suffer from the limitations of other shared storage solutions and delivers both scalability and efficiency.

ActiveData with WekaFS supports all major Linux distributions and leverages virtualization and low-level Linux container techniques to run its own RTOS (Real-Time Operating System) in the user space, alongside the original Linux kernel. Weka manages its assigned resources (CPU cores, memory regions, network interface cards, and SSDs) to provide process scheduling and memory management, and to control the I/O and networking stacks. By not relying on the Linux kernel, WekaFS effectively utilizes a zero-copy architecture with much more predictable latencies.

WekaFS functionality running in its RTOS is comprised of the following software components:

- **File Services (Front End)** – Manages multi-protocol connectivity.
- **File System Clustering (Back End)** – Manages data distribution, data protection, and file system.
- **SSD Access Agent** – Transforms the SSD into an efficient networked device.
- **Management Node** – Manages events, CLI, statistics, and call-home capability.
- **Object Connector** – Read and write to the object store.

Bypassing the kernel means that Weka’s software stack is not only faster with lower latency, but is also portable across different bare-metal, virtualized, containerized, and cloud-instanced environments.



Supported Protocols

Clients with the appropriate credentials and privileges can create, modify, and read data using one of the following protocols:

- POSIX
- NFS (Network File System) v3
- SMB (Server Message Block) v2 and v3
- S3 via gateway
- NVIDIA® GPUDirect® Storage (GDS)⁴


Persistence Storage Volumes with Kubernetes CSI Plugin

The combination of containers, Kubernetes orchestration, and software-defined scale-out storage is a winning match. Enabling these elements achieves both application and data portability across the edge to core to the cloud.

Many organizations start their container journey with stateless workloads that are easier to transition to a cloud-native microservices architecture. However, the majority of HPC, Data Analytics and AI/ML applications today are stateful workloads that live throughout the enterprise. Organizations seek to modernize and containerize these applications without significant refactoring – while ensuring production-grade security and persistent data storage.

The Weka CSI Plugin provides support for static and dynamic provisioning of persistent container storage for Kubernetes orchestration, by implementing all three services: Identity, Controller, and Node.

ActiveData with WekaFS can be combined with other Penguin Computing HPC solutions, such as TrueHPC™ and InsightHPC™ that provide job management and other cluster management capabilities.

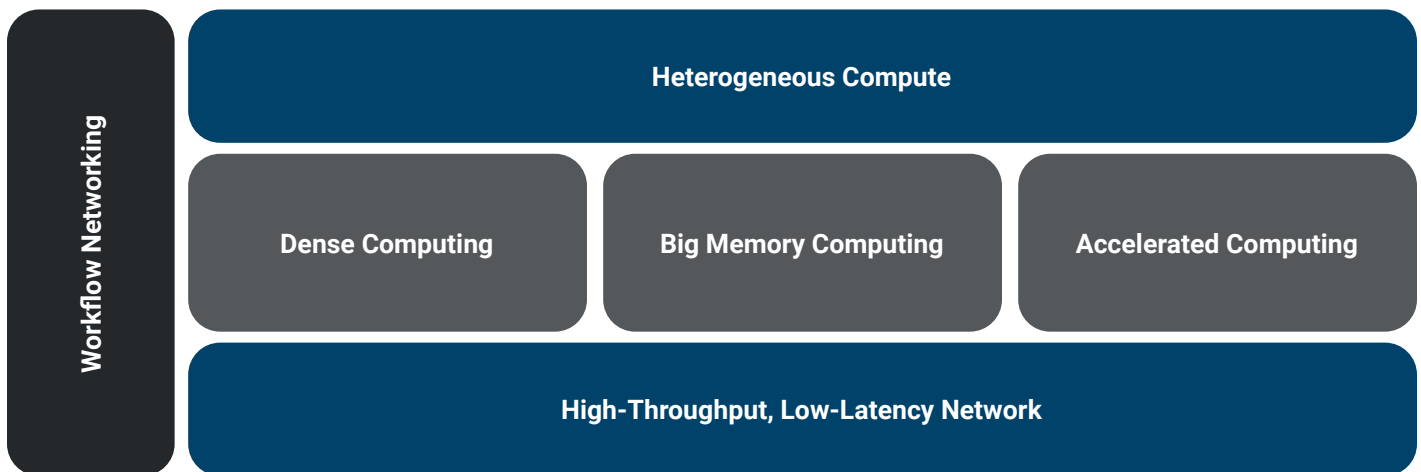


Compute Technologies

High Performance Workloads

High performance workloads require a large number of cores and high core clock speeds to achieve the best performance possible. These workloads also require high performance interconnects because many workloads span multiple servers, requiring constant node-to-node communication that benefits from high-throughput and low-latency network technologies.

Memory-centric workloads call for additional server memory resources to support applications that require extreme read and write performance and extremely low latency. Accelerated-computing workloads require enterprise accelerators such as GPUs, FPGAs, and ASICs to drastically improve the performance of certain applications.




High Performance Data Requirements

Accelerating computing has enabled workloads to run on many fewer systems than previously possible. Workloads that once require 100 servers can now run on just 10. This trend to increase the computing capability of a single system has increased the I/O demand per server.

Data sets are growing exponentially, creating storage bottlenecks in application workflows. These bottlenecks cause accelerated compute systems to idle, reducing the ROI of those systems. In order to maximize the ROI of accelerated compute infrastructure, organizations need to optimize their storage for the ingest needs of their accelerating computing.

Designing a workload-optimized storage system includes in-system, device-to-device bandwidth and resource optimizations and system-to-system network communication optimizations. These design choices ensure that the underlying computing, storage, and networking infrastructure are optimized for the workloads that customers will run on them.

Data Center Infrastructure



ActiveData can be built using both a traditional 19" rack platform and a modern 21" OCP (Open Compute Project) platform. Traditional 19" rack infrastructures are supported in almost every data center worldwide and in a variety of dimensions. Modern 21" OCP rack infrastructures require data centers that can support the most demanding physical and power densities. Penguin Computing has partnered with leading data center facility pioneers who can support the demanding characteristics of today's HPC platforms.

Power

ActiveData supports three-phase 50A or 60A, 208V, 277V, or 480V power options as well as A+B fully redundant power, or N+1 redundant power. 21" OCP also supports 12V or 48V power delivered directly to the servers, which enable much higher power density per rack.

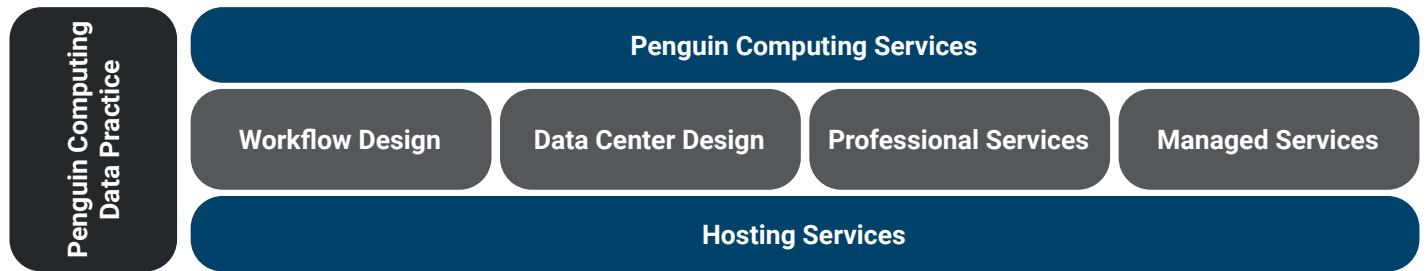
Cooling

ActiveData can be air cooled with traditional HVAC equipment. Penguin Computing recommends using a combination of air cooling and liquid cooling when deploying ActiveData into a data center not designed for high-power equipment. Rear Door Heat Exchangers capture hot air exhaust at the rear of the rack, and can be deployed on most 19" and 21" rack infrastructures. ActiveData is also designed to integrate Direct-To-Chip cooling options that capture heat directly from the CPU block. This cooling solution removes 85% of server heat before it's transferred into the air, and can be used in select 21" infrastructures.



Penguin Computing Services

Penguin Computing ActiveData is a comprehensive, end-to-end solution that organizations can leverage to jump-start their high performance data initiatives. In some cases, the solution 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, scale of the environment, maximum rack power consumption, power phase balance, efficient heat removal, and the optimal networking topologies when using low-latency, high throughput interconnects.

ActiveData is supported by Penguin Computing engineering services, which can include 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 ActiveData 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 ActiveData provides a single, secure, end-to-end solution for performance data requirements that includes a flexible, scalable, workload-optimized data infrastructure and management tools that allow you to not only monitor your storage, but easily manage and provision file systems..

Penguin Computing ActiveData frees organizations from having to focus valuable time and human resources on creating an architecture and software-defined storage solution from scratch, to lower TCO, reduce risk, and accelerate time-to-innovation.

Penguin Computing applies 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 data-driven computing initiatives to production.

Contact Us

Use this [form](#) or call Penguin Computing today at 1-888-736-4846 to find out how you can take advantage of a production-ready NVMe storage solution that can:

- Speed time-to-market
- Stay ahead of emerging technologies
- Reduce operating costs
- Improve employee productivity
- Improve performance across all file sizes



**PENGUIN
COMPUTING**

Expanding the world's vision of what is possible