



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

LiveData with MemVerge Memory Machine

Support application data needs with powerful,
flexible big memory



Contents

- Solution-at-a-Glance 3**
- Introduction 3**
- Penguin Computing LiveData 4**
 - Data Technologies 6
 - Enterprise Engineered Big Memory Computing Solution 7
 - Memory Machine Standard Edition 7
 - Memory Machine Advanced Edition. 7
 - Key Features & Benefits. 8
 - Software Technologies 9
 - Software Defined Memory Architecture 9
 - Compute Technologies 10
 - High Performance Workloads 10
 - Big Memory Computing Requirements 10
 - LiveData Solution Use Cases. 11
 - Increase Availability 11
 - Instantly Transmit & Store Massive Amounts of Real-time Data 11
 - Improve Memory Capacity & Provide Plug-Compatible Persistent
 - Memory Access for ML Applications 12
 - Scale Memory Capacity & Availability VMs 12
- Data Center Infrastructure. 12
 - Power 13
 - Cooling 13
- 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

- Tier Persistent Memory and DRAM for optimum performance.
- Low-latency memory replication.
- Virtualize to form a platform for enterprise-class data services.
- Plug-compatible access with existing applications.
- Recover hundreds of GB in seconds with ZeroIO™ memory snapshots.
- Clone databases in seconds.

Benefits


- Innovate and speed time-to-market.
- Scale memory capacity and improve system performance.
- Maintain availability - recover in seconds, not hours.
- Enable agility with efficient clone deployment, and fast application rollbacks.
- Avoid application disruption or rewrite.
- Reduce latency.

Introduction

Data has become one of the greatest drivers of our economy, our businesses, and our IT infrastructures. As data-intensive workloads scale, it's critical to implement data-driven, software-defined architectures that meet the demands of large data sets. These architectures are categorized into tiers based on characteristics such as performance, capacity, connectivity, and cost. They are 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 latency-sensitive workloads, a second fast and scalable flash tier for I/O-intensive workloads, and a capacity-optimized tier for long-term storage.

According to IDC, real-time data was less than 5% of all data in 2015 but is projected to comprise almost 30% of all data by 2024. They project that by 2021, 60-70% of the Global 2000 will have at least one mission-critical real-time workload. According to IDC, by 2021, 60-70% of the Global 2000 organizations will have at least one mission-critical real-time workload.

This Big Bang of real-time data is driving the expansion of real-time analytics and AI/ML applications into the mainstream. Another result is real-time applications are outpacing the capacity, performance, and availability capabilities of in-memory infrastructure. This can lead to congestion, I/O bottlenecks, storage outages, and cost overruns for data-intensive HPC and AI/ML workloads.



By 2021, 60-70% of the Global 2000 organizations will have at least one mission-critical, real-time workload.

These data bottlenecks, in turn, can lead to a situation known as Data Greater than Memory (DGM), when the data can no longer fit in memory. When this happens, data traditionally needs to overflow to hard drives or SSDs. This leads to a dramatic performance drop -- 1000 times slower, 5,000 times slower. It is then no longer a bottleneck, it's a roadblock.

Leveraging a powerful, software defined architecture to aggregate the performance and capacity of DRAM and persistent memory eliminates these roadblocks and enables innovative workflows. The right technical partner can provide an optimized platform with the architecture, integration, support, and managed services to ensure your success.


Penguin Computing LiveData

Penguin Computing has partnered with MemVerge to bring you Penguin Computing LiveData with MemVerge Memory Machine™. Penguin Computing LiveData addresses the DGM roadblock by providing a memory virtualization software layer that delivers software-defined memory services to the applications without application changes. This allows the thousands of applications running in the data center today to take advantage of higher capacity of memory at a lower cost.

MemVerge Memory machine allows you to massively scale out DRAM and persistent memory. The results are big memory pools where all applications and data can live. To support all application data needs, MemVerge has invented rich big memory data services such as snapshot, replication, and tiering that for the first time enable lightning fast recovery from in-memory application crashes. Existing tier-1 applications can run safely and transparently on big memory without application rewrites.

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

Worldwide Real-Time Data and Share, 2015-2024



Source: IDC® Global Datasphere, 2020

LiveData includes:




Data Technologies



Software Technologies



Compute Technologies




Data Center Infrastructure



Penguin Computing Services

Penguin Computing LiveData Components



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.

Enterprise Engineered Big Memory Computing Solution

MemVerge Memory Machine responds to the need for a modern in-memory computing model that addresses the DGM syndrome. The first in the new class of big memory software, Memory Machine virtualizes DRAM and persistent memory so that memory can be accessed without code changes, scale-out in clusters to provide the capacity needed by real-time analytics and AI/ML applications, and deliver enterprise-class data services for high-availability.

MemVerge Memory Machine, with its unique IP, memory allocation, software algorithms, and memory-peering algorithms, allows Penguin Computing LiveData to make application performance DRAM-like, and in some applications even exceed DRAM performance.

LiveData provides an engineered and optimized turnkey big memory computing solution catered to memory-centric workloads. LiveData can be right sized to your workloads based on standard and advanced features.

Memory Machine Standard Edition

Businesses that need to address the cost of their growing memory footprint can leverage Memory Machine Standard Edition to provide applications with transparent access to virtualized pools of persistent memory to accommodate larger jobs. Memory Machine also allows lower-cost persistent memory to perform like DRAM. Instead of modifying multiple applications, IT organizations can now efficiently deploy big memory across a data center with a single Memory Machine virtualization layer.

Memory Machine Advanced Edition

Memory Machine Advanced Edition includes all the capabilities of Memory Machine Standard Edition, plus transparent access to persistence via the industry's first enterprise-class data services for highly available memory.

IT organizations rely on data services for their disk and all-flash storage, especially snapshots and replication which enable crash recovery. Memory Machine's patented ZeroIO™ memory snapshot technology eliminates IO to storage. The result is the ability to snapshot and recover terabytes of data from persistent memory in a few seconds, instead of taking minutes to hours from storage. In addition, Memory Machine's Pub/Sub feature provides a lightning fast in-memory message bus that can achieve high availability through data replication.

Key Features & Benefits



Virtualize to form a platform for enterprise-class data services – Persistent Memory Machine virtualizes DRAM and persistent memory so that data in-memory can be accessed, tiered, scaled, and protected.



Plug-compatible access with existing applications – Memory Machine provides access to persistent memory without changes to applications. It looks like DRAM to your apps.



Tier Persistent Memory and DRAM for optimum performance – Memory Machine creates a DRAM “fast tier” and a PMEM “persistent tier”.



Low-latency memory replication – For apps like Pub/Sub, Memory Machine uses RDMA to publish high volumes of messages to remote subscribers with ultra-low latency.



Recover hundreds of GB in seconds with ZeroIO memory snapshots – To protect large memory blast zones, Memory Machine provides memory snapshot and replication for lightning-fast crash recovery that requires zero IO to storage.




Clone databases in seconds – Developers and the applications they create can use snapshots to clone terabyte database instances while sharing existing memory.



Software Technologies


Software Defined Memory Architecture

Memory Machine is a Linux-based software subscription deployed on a single server or in a cluster. Once DRAM and persistent memory are virtualized, the Memory Machine hypervisor can guarantee performance to specific applications with QoS, set-up DRAM as a “fast” tier for persistent memory and provide enterprise-class memory data services based on ZeroIO™ memory-to-memory snapshots.



No matter where you're operating from edge to core, workload portability is important in workflow optimization. Containerization allows you to build, deploy, run, and manage modern applications, and deliver them at scale faster and more reliably to the end user. The ability to support containers and traditional bare metal workloads is important to ensuring you can accelerate your workload and decrease your time-to-insight. The ability to accelerate your workloads from edge to core and in any environment is important to delivering workload portability.

LiveData with MemVerge Memory Machine allows you to do all of that even better. LiveData supports HPC schedulers, Docker, Kubernetes, and OpenShift.




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.



Big Memory Computing Requirements

Accelerated computing has enabled workloads to run on many fewer systems than previously possible. Workloads that once required 100 servers can now run on just 10. The trend to increase the computing capability of a single system has exposed the memory as a bottleneck in many workloads.

These bottlenecks force users to limit the size of their models to fit within a limited memory footprint based on DRAM technologies. These artificial restrictions reduce the accuracy of certain workloads, and cause others to take much longer to complete, because users have to run many jobs in several chunks before aggregating the outputs together. Using Intel Optane persistent memory, the total addressable memory footprint of a server can be increased between two to four times that of a system using DRAM only.



Designing a workload optimized memory system includes in-system memory-to-flash and memory-to-network optimizations. These design choices ensure that the memory footprint can run extremely large models, enabling accelerated-computing systems to compute optimally, increasing workload accuracy and improving completion time.

LiveData Solution Use Cases

MemVerge’s prediction in the era of big memory computing is that all applications will eventually run in memory that is abundant, persistent, and highly available. As 451 Research reports, “The abundance and persistence will result from the higher physical density and lower cost of storage-class memory (SCM) compared with conventional DRAM memory, as well as their persistence and non-volatility. The high availability actually refers to the availability of applications, which will be boosted by the persistence of SCMs, which will allow for significantly faster application recoveries after server crashes or restarts. According to MemVerge, availability will also be boosted by the use of the high-speed memory-based snapshot and replication functions of its Memory Machine software.”¹ Since its inception in 2017, MemVerge has been helping customers achieve their system performance and business goals in a variety of data-intensive applications.

Increase Availability

“Slow” is the state of the art for IMDB crash recovery. Typically, a crashed system can be restored from disk or SSD replaying the transaction log to catch up to the last saved state. For one financial services customer of MemVerge, this process takes 3 hours to recover 500GB.

As the blast zone for IMDB crashes grows wider and deeper, Intel and MemVerge respond with big memory. With persistent memory and Memory Machine software, memory can safely scale-out to petabytes because it’s now possible to recover from crashes in seconds. The state-of-the-art for IMDB crash-recovery has been changed from “slow” to “fast”.

Instantly Transmit & Store Massive Amounts of Real-time Data

Applications ranging from news, to risk analysis, to stock trading rely on streaming real-time data from a network of exchanges. As the number of nodes and messages grows, instability emerges, limiting the scalability of a publish/subscribe network.

1 | https://www.memverge.com/web/wp-content/uploads/2020/08/451_Reprint_MemVerge_03AUG2020.pdf

The Memory Machine software publish/subscribe messaging API provides the tools to create a pub/sub system to instantly transmit and store massive amounts of time sequence data such as stock exchange trading ticks.

Improve Memory Capacity & Provide Plug-Compatible Persistent Memory Access for ML Applications


Inference from a pretrained model often requires large databases that must reside in system memory to achieve sufficient transactional performance. When resources become tight, nominally less used memory items are written to disk to make more space, only to be swapped back later. This back-and-forth has severe impacts, drastically reducing the performance.

Deep Learning Recommendation Models (DLRMs), image recognition, and the open source Stanford GraphSAGE are three ML applications that Memory Machine software has powered to use large persistent-memory configurations to eliminate the performance penalty of memory swapping.

Scale Memory Capacity & Availability in VMs

Host is not responsive. Virtual machines cannot be started. Applications in virtual machines run out of memory. These are all caused by the DGM problem that we described earlier. The capacity, availability, and crash recovery challenge is even greater when large data sets and real-time analytics are involved.

Memory Machine addresses a VM's need for more memory with support for higher-density / lower-cost persistent memory that allows more memory to be configured in a single physical server. Memory Machine software then takes capacity to the next level with the ability to scale-out memory in a virtual server cluster to form massive memory lakes. Memory Machine protects all that data with memory data services including snapshot, replication and lightning-fast recovery.



Data Center Infrastructure

LiveData 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

LiveData 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

LiveData can be air cooled with traditional HVAC equipment. Penguin Computing recommends using a combination of air cooling and liquid cooling when deploying LiveData 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. LiveData 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

LiveData is a comprehensive, end-to-end solution that organizations can leverage to jump-start their big memory computing 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 the environment, maximum rack power consumption, power phase balance, efficient heat removal, and the optimal networking topologies when using low-latency, high throughput interconnects.

LiveData 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 LiveData 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

LiveData with MemVerge Memory Machine provides a complete, end-to-end big memory Computing solution for memory-centric real - time workloads that includes a flexible, scalable, workload-optimized data infrastructure and management tools that allow you to not only monitor your big memory computing resources, but easily manage and deploy big memory computing workloads..

LiveData frees organizations from having to focus valuable time and human resources on creating an software defined architecture and deploying memory-centric infrastructure 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 big memory 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 use LiveData with MemVerge Memory Machine to open the door to big memory, a world of abundance, persistence, and high availability where you can:

- Innovate and speed time-to-market.
- Scale memory capacity and improve system performance.
- Maintain availability - recover in seconds, not hours.
- Improve agility with efficient clone deployment, and fast application rollbacks.
- Avoid application disruption or rewrite.
- Reduce latency.



**PENGUIN
COMPUTING**

Expanding the world's vision of what is possible