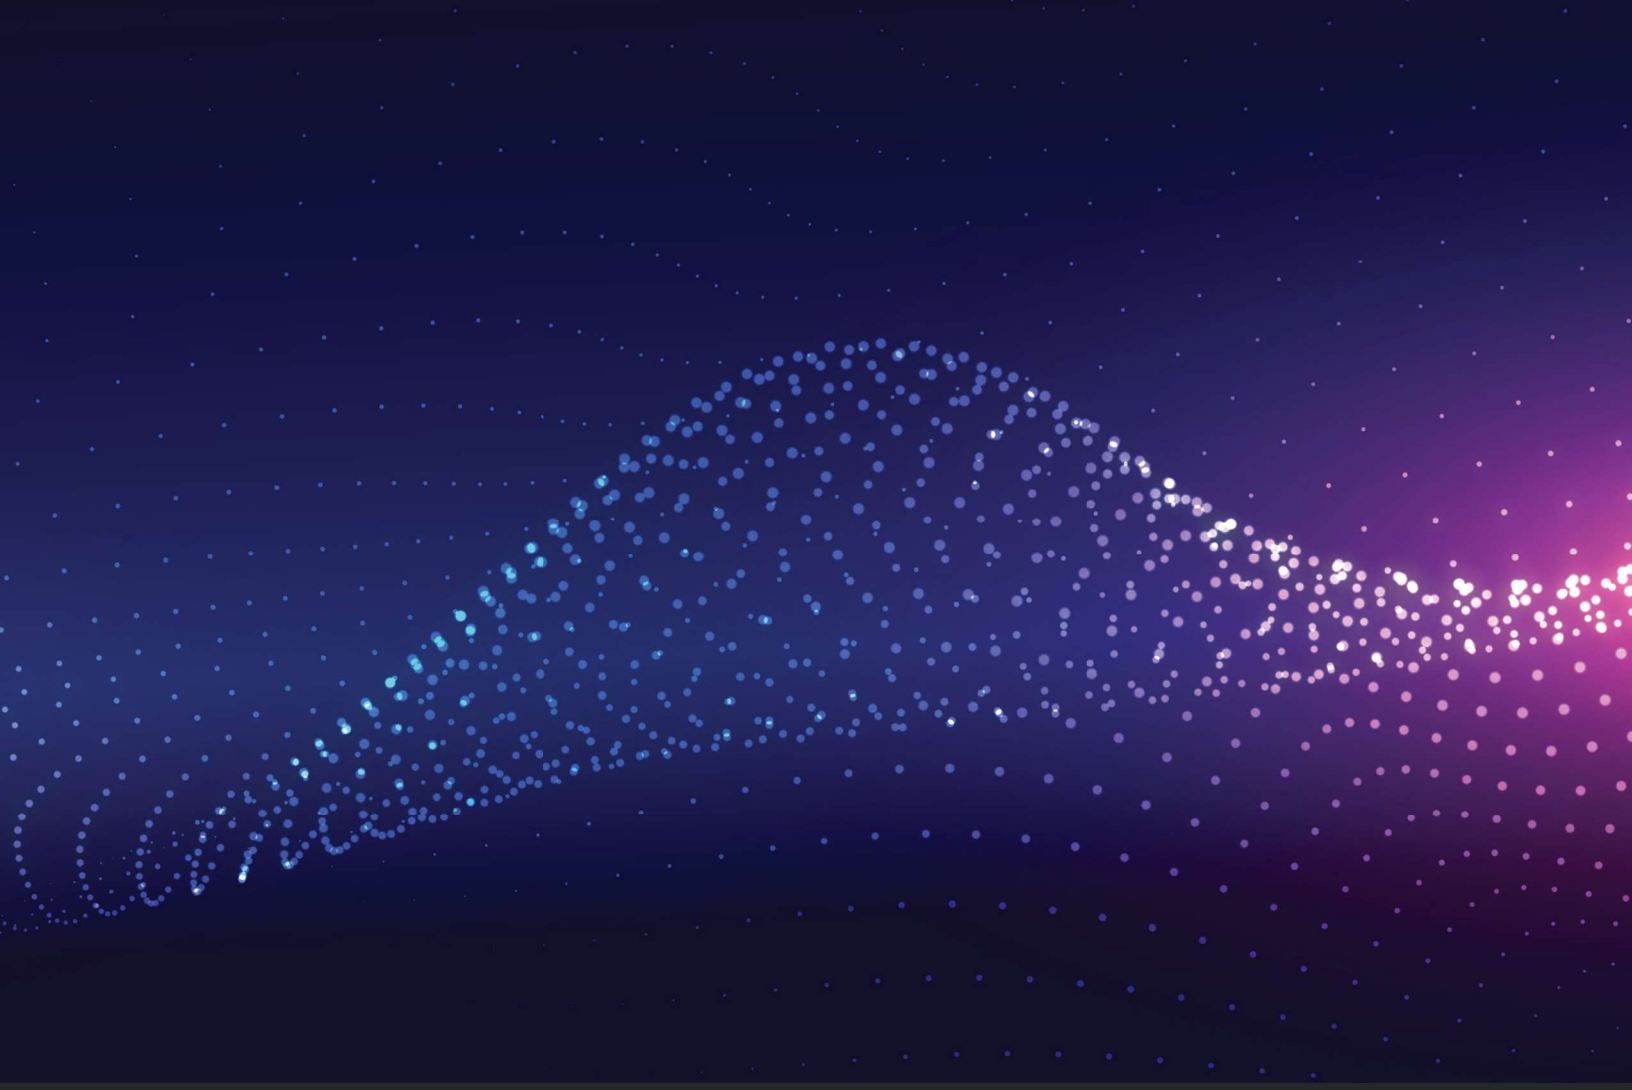

TrueHPC with Scyld ClusterWare

Driving innovation with ready-to-run HPC

Featuring Relion® servers powered by Intel® cores.

intel

PENGUIN COMPUTING

Contents

Solution-at-a-Glance

Features

- A comprehensive, proven OCP-compliant compute architecture for HPC.
- Workload-optimized compute servers.
- Rapid image-based provisioning of bare metal servers.
- Flexible infrastructure to support heterogeneous compute.
- Innovative, configurable cooling and power management.
- Support for Slurm, OpenPBS, and TORQUE HPC job schedulers.
- Centralized cluster reporting and diagnostics in a single pane of glass.
- Enhanced Support for secure environments (SELinux with MLS, FIPS 140-2).
- Pre-bundled HPC software platform, backed by commercial support.
- Professional support by HPC and cluster management experts.

Benefits

- Jump start HPC initiatives with a ready-to-run enterprise HPC cluster.
- Accelerate time-to-innovation.
- Accelerate adoption of rapidly changing technologies.
- Efficiently manage change.
- Lower total cost of ownership.
- Reduce risk with high availability (HA) features.
- Free up administrators so they can focus on tasks specific to your business.
- Customize and manage the HPC environment with intuitive, supported tools.

The HPC Challenges

The HPC landscape has changed dramatically over the past two decades. New entrants to the market are leveraging new strategies to push the boundaries of the discipline. At the same time, the underlying technology is evolving at an increasingly rapid pace, enhancing the capabilities of organizations that can adapt and improving their competitiveness. Whether you are a new user of HPC trying to harness the power of this technology for the first time or a seasoned practitioner looking for a performance edge, consider the following when adopting these technological advances:

Platform Complexity: Designing the right HPC platform for targeted workloads is a complex task. Ensuring that the compute, storage, and networking subsystems are well designed individually and function in a balanced manner together is critical. There are many degrees of freedom in the process, and new variables are introduced with each technological advance in these subsystems. A poor design choice at any point in the process can negatively impact performance and reliability, and can significantly reduce the value of any HPC investment.

Risk Mitigation: The complexity of HPC platforms can increase points of failure. Sometimes these failures can range from subtle performance degradation that is difficult to identify to catastrophic failures of key elements of the architecture that result in measurable loss of productivity. Fortunately, elements of high availability can be designed into HPC architectures in order to avoid the worst failures.

Hardware Abstraction: HPC users are moving closer to their workloads and farther from the underlying hardware that runs them. This very effective approach to HPC leverages cloud-native technologies to accelerate innovation by allowing domain experts to do what they do best: deliver research results. However, the fact that engineers are less concerned with hardware — combined with the fact that those hardware platforms are increasingly more complex — places an additional burden on research infrastructure teams to bridge that gap. Adopting HPC for the first time, augmenting an existing HPC platform with emerging technologies, or repatriating cloud-based HPC resources all drive a need for new tools and skill sets.

Cluster Management, Control, & Customization: The productive output of an HPC platform relies primarily on its architecture and secondarily on the effectiveness of the management layer and how it is customized to an organization's workflows. Tools that automate the basic management functions of an HPC environment in a manner that is conducive to the workflows of the user base are critical. Proper selection, implementation, and ongoing management of this layer boosts productivity for users and administrators alike and frees both to excel the more complex aspects of their work.

Exponential Data Growth: While HPC users have grappled with the challenges of processing, moving and extracting insight from data for decades, those challenges have been pushed to the fore as our global data production increases. These challenges are further compounded with the move to cloud and now the emergence of powerful hybrid and multi-cloud architectures. Grappling with issues such as data silos, expensive bandwidth, the drive to real-time insight, and a geographically distributed workforce can be daunting without the powerful new tools and innovative strategies required to succeed.

Data Center Infrastructure: HPC platforms have always tested the infrastructure limits of contemporary data centers. That has never been more true than today as we push the envelope with new demands on connectivity, space, cooling, and power densities that reach up to 100kW per rack. The symbiosis between the compute platforms and the data centers that house them call for a unified approach to the complete HPC architecture from the platform outward. Even if a new data center is not in your plans, there are many power and cooling technologies that can enable existing data centers to accommodate the new standards in HPC.

Speed your time to results with HPC infrastructure that is engineered by a partner who knows HPC and cluster orchestration tools that were created specifically for HPC administrators by engineers who manage HPC environments.

Penguin Computing TrueHPC

For more than two decades, Penguin Computing has been helping organizations build HPC architectures to meet the requirements of the most demanding workloads.

Penguin Computing TrueHPC™ provides a complete solution built on our workload-optimized server building blocks and our Scyld ClusterWare® HPC orchestration technologies.

TrueHPC is a ready-to-run, enterprise-supported HPC solution that leverages high-performance, low-latency networking to deliver a workload-optimized high performance computing infrastructure that scales to meet your workloads.

TrueHPC can be implemented alone as a solution or in combination with other Penguin Computing solutions for Data, AI/Analytics, and Cloud to provide an end-to-end complete compute platform.

TrueHPC includes:

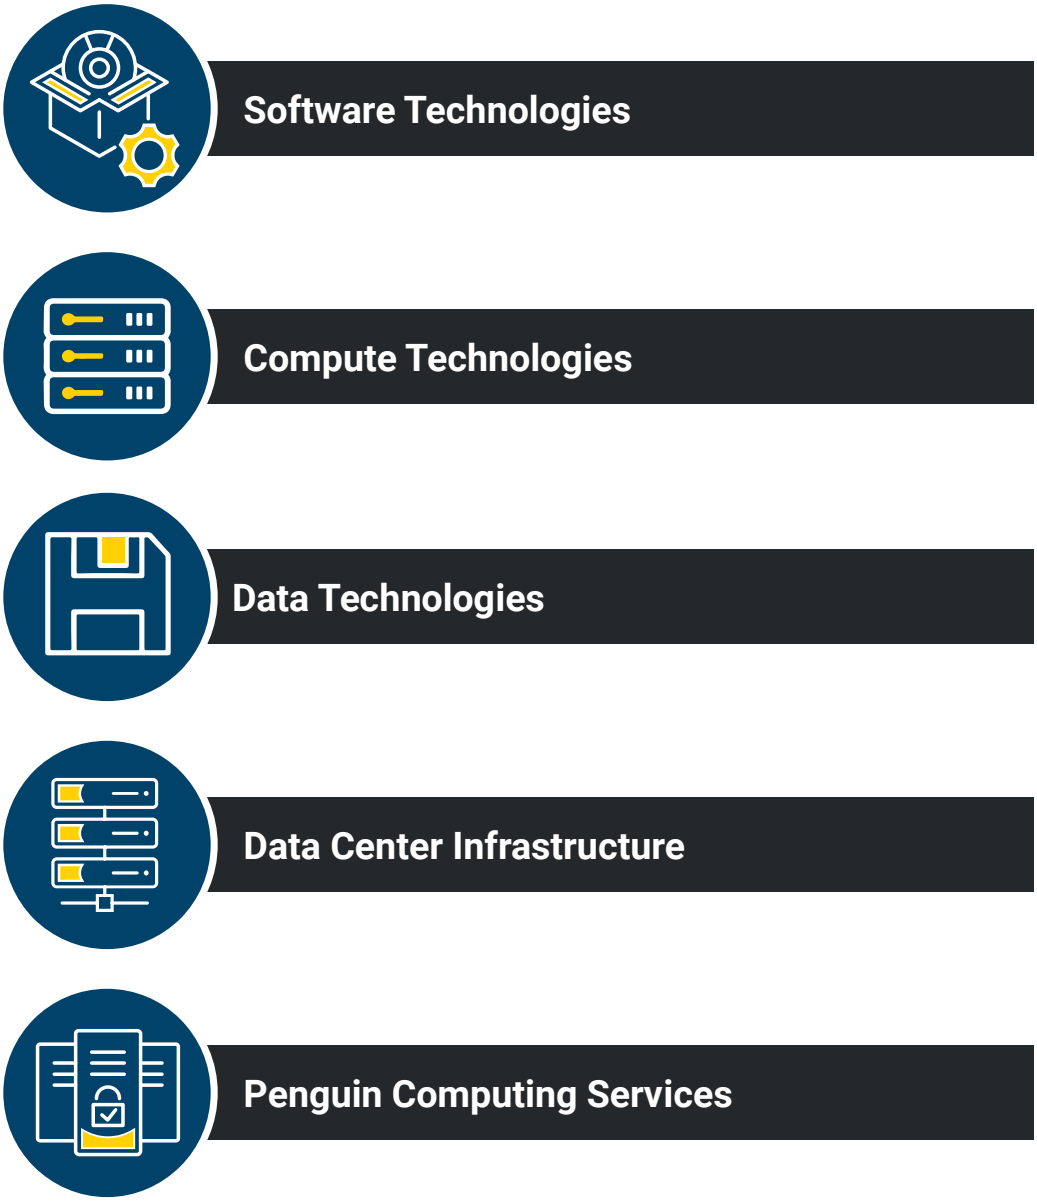

Penguin Computing TrueHPC Components

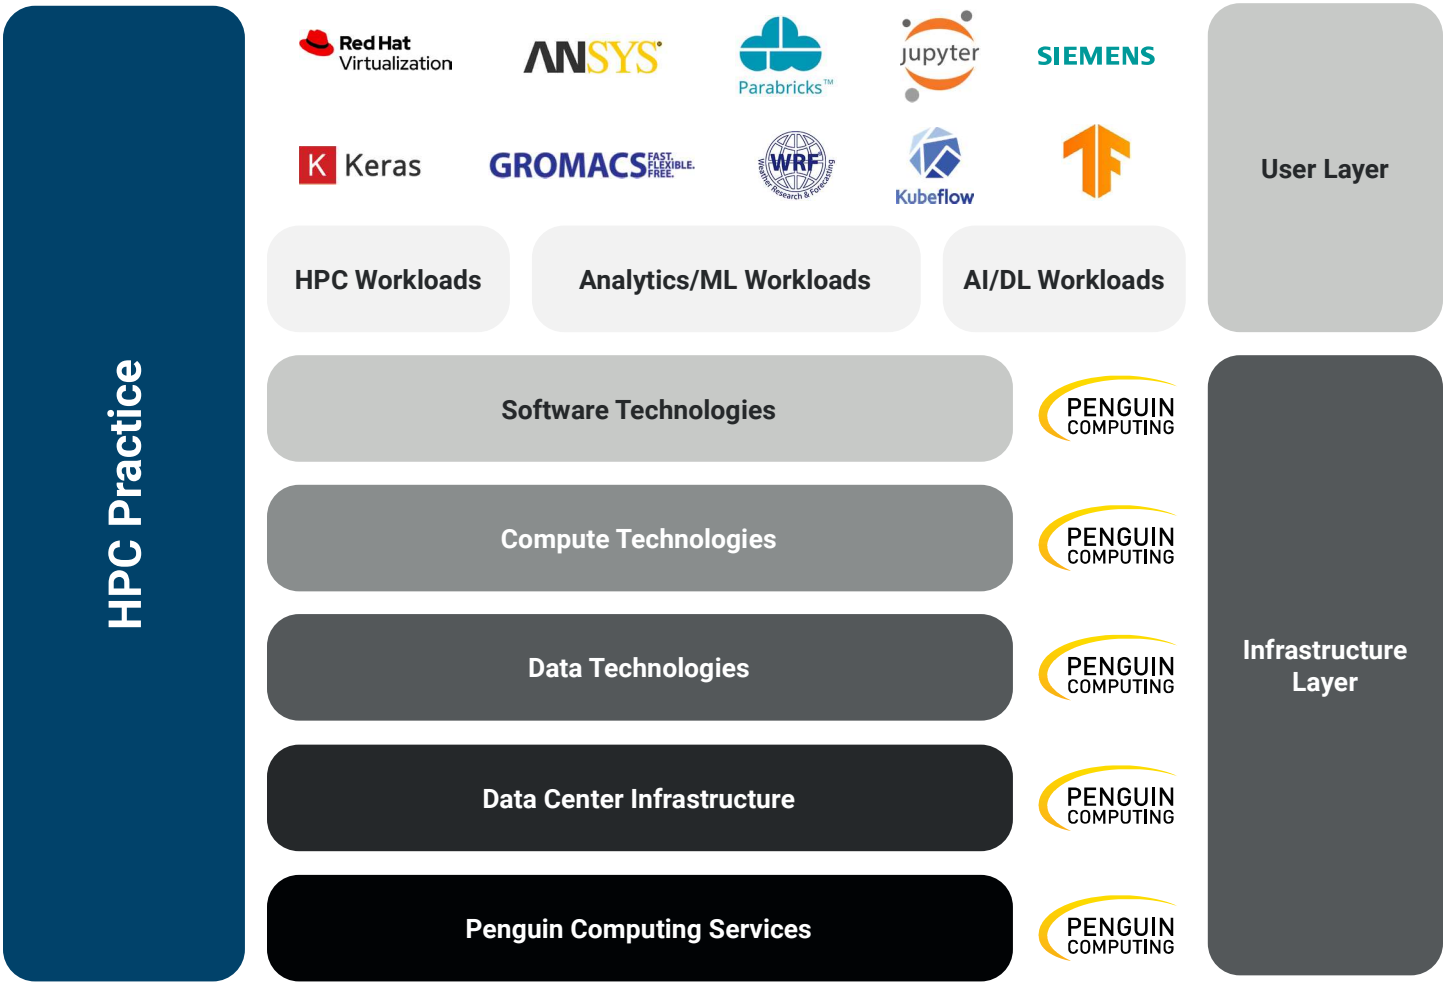

Software Technologies

HPC Cluster Orchestration

TrueHPC 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.

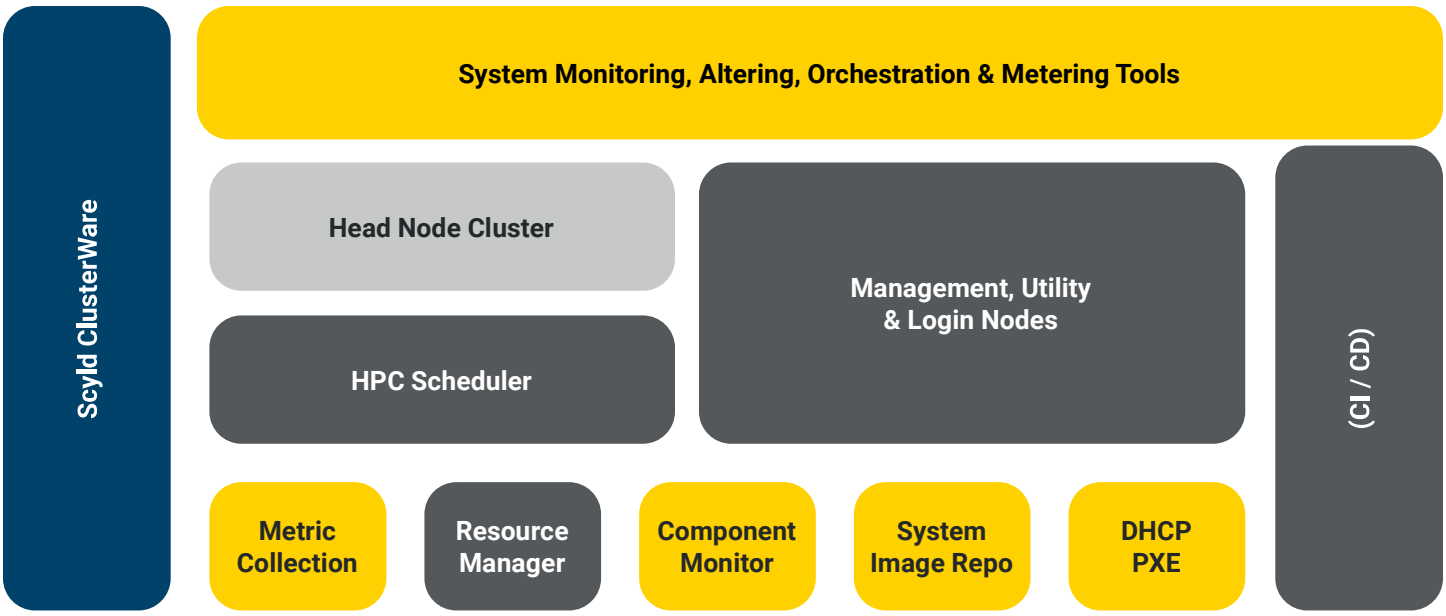

Alerting and Monitoring

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 a TrueHPC 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.

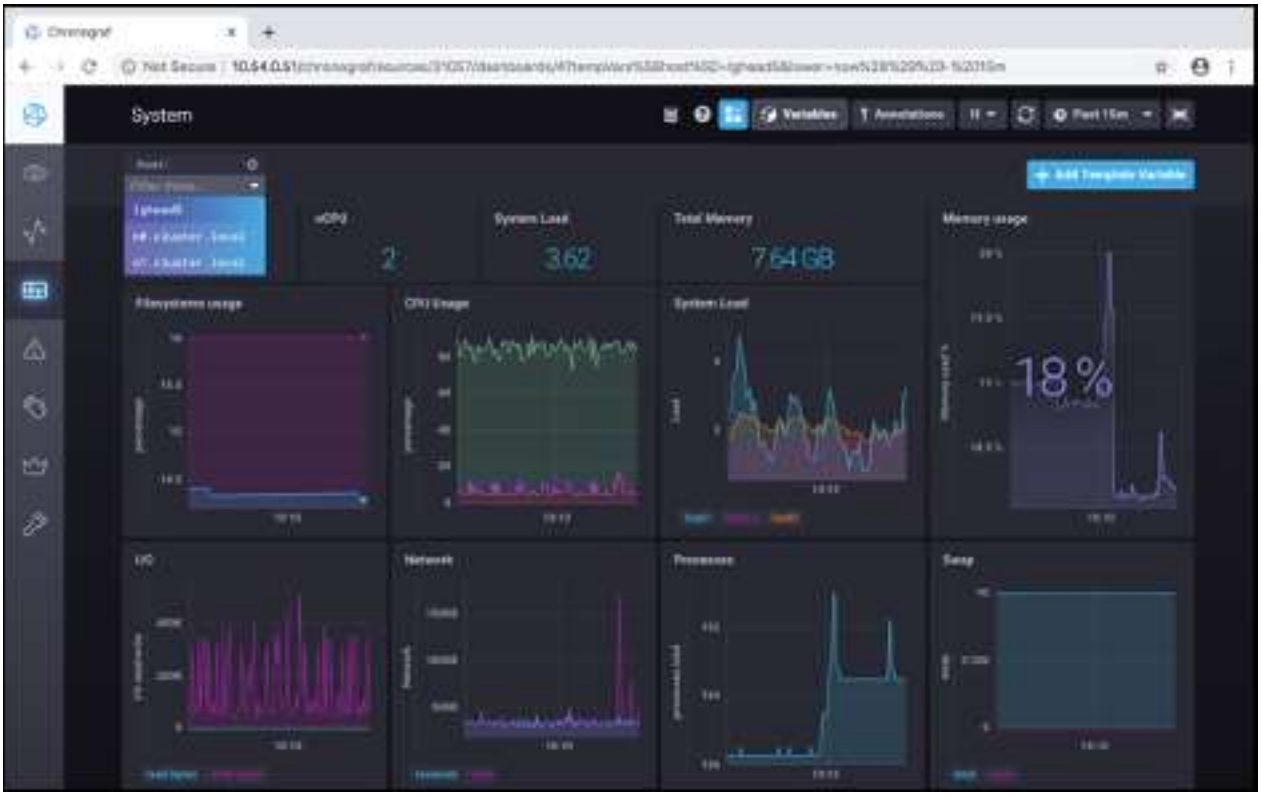

Cluster Customization

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 TrueHPC 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. The TrueHPC solution is tested and delivered as an integrated solution with security and support as paramount components.

Scyld ClusterWare Features

- 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

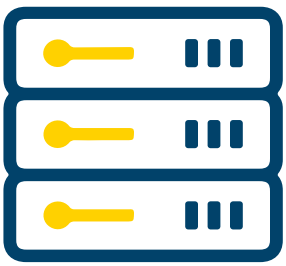

Compute Technologies

Heterogeneous Compute

TrueHPC leverages industry-leading technologies from Intel, AMD, NVIDIA, and other technology providers to enable a complete technology ecosystem supporting many different workloads. HPC workloads often require high core count, high clock speed, high memory bandwidth, low latency communication, and/or accelerated computing using GPUs, FPGAs, and ASICs. TrueHPC supports heterogeneous computing environments within a single architecture using workload-optimized server building blocks for many types of high performance workloads.

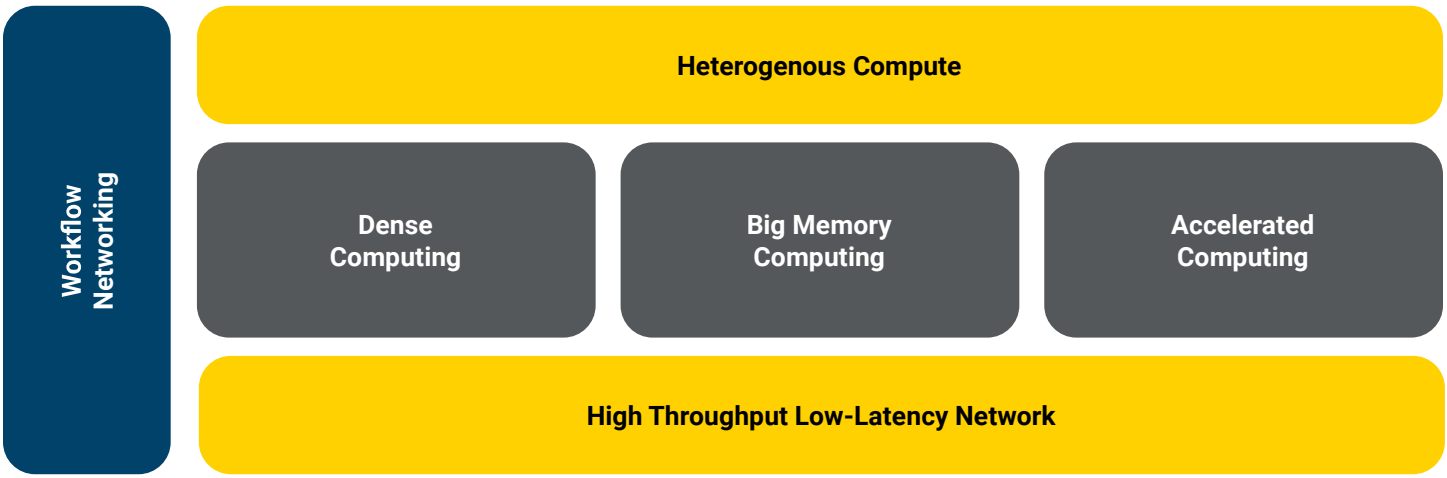

Workload-Optimized Servers

HPC workloads typically 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 HPC workloads span multiple servers, requiring constant node-to-node communication that benefits from high-throughput and low-latency network technologies. Optimized server building blocks for HPC workloads need to provide many cores with high clock speeds and low-latency, high-throughput interconnects to provide the best application performance possible.

Memory-centric workloads call for additional server memory resources to support applications that require extreme read and write performance and extremely low latency. Optimized server building blocks for big memory computing require as much memory bandwidth, capacity, and clock speed as possible to provide the best application performance.

Accelerated computing workloads require enterprise accelerators, such as GPUs, FPGAs, and ASICs to drastically improve the performance of certain applications. Optimized server building blocks for accelerated computing require in-system, device-to-device communication optimizations to ensure that accelerators can communicate with CPUs, SSDs, NICs, and other accelerators without communication bottlenecks in order to provide the best application performance possible.

TrueHPC supports several server building blocks optimized to improve the performance for certain workloads.

Relion XO1132g Server - General Purpose Computing

- Up to 56 Intel Scalable Cores
- Up to 3.9GHz Core Clock Speeds
- Up to 4TB DDR4 2933MHz RAM
- Up to 100Gb HDR InfiniBand

Relion XO1114GT Server - Big Memory and Accelerated Computing

- Up to 56 Intel Scalable Cores
- Up to 3.9GHz Core Clock Speeds
- Up to 6TB DDR4 2933MHz RAM
- Up to 100Gb HDR InfiniBand
- Up to 4 x PCIe Gen 3 GPGPU

Relion XO1114GTS Server - Big Memory and Accelerated Computing

- Up to 56 Intel Scalable Cores
- Up to 3.9GHz Core Clock Speeds
- Up to 6TB DDR4 2933MHz RAM
- Up to 100Gb HDR InfiniBand
- Supports 4 x SXM2 GPU

High Speed Low-Latency Interconnects

The TrueHPC solution supports the leading high-throughput, low-latency networking interconnects that help maximize the performance of an HPC cluster for certain workloads.

Cornelis Omni-Path Fabric 100 48-Port Switch

- 48 x 100Gb Omni-Path Ports
- 9.6Tbps Aggregate Throughput
- 110ns Port-to-Port Latency

Data Technologies

HPC is moving toward data-driven workloads that consume and generate large amounts of data. This data growth drives the need for data solutions that can scale to exabyte capacities. HPC environments have data requirements that create data workflow and infrastructure challenges related to management and orchestration.

Data I/O requirements weigh heavily on the overall success of an HPC solution. I/O patterns and performance vary across different tiers of storage in the environment — from memory, to flash, to cold storage. TrueHPC integrates with the data solutions in Penguin Computing's Data Practice, which cover the entire spectrum of I/O — from memory, to flash, to cold storage — to support the entire data lifecycle.

Big Memory Computing

Some HPC workloads can require massive memory storage or high memory performance. By utilizing the LiveData™ solution, TrueHPC can support memory-centric workloads that require high memory capacity, persistent memory, and high memory tier performance.

High Performance Storage

HPC workloads require high performance storage that can fulfill the ingress and egress demands of high performance workloads. TrueHPC can be paired with the ActiveData™ solution to provide high performance storage for data-heavy computing workloads. Just as with TrueHPC, the ActiveData solution can leverage the industry-leading technologies that are tuned for specific customer

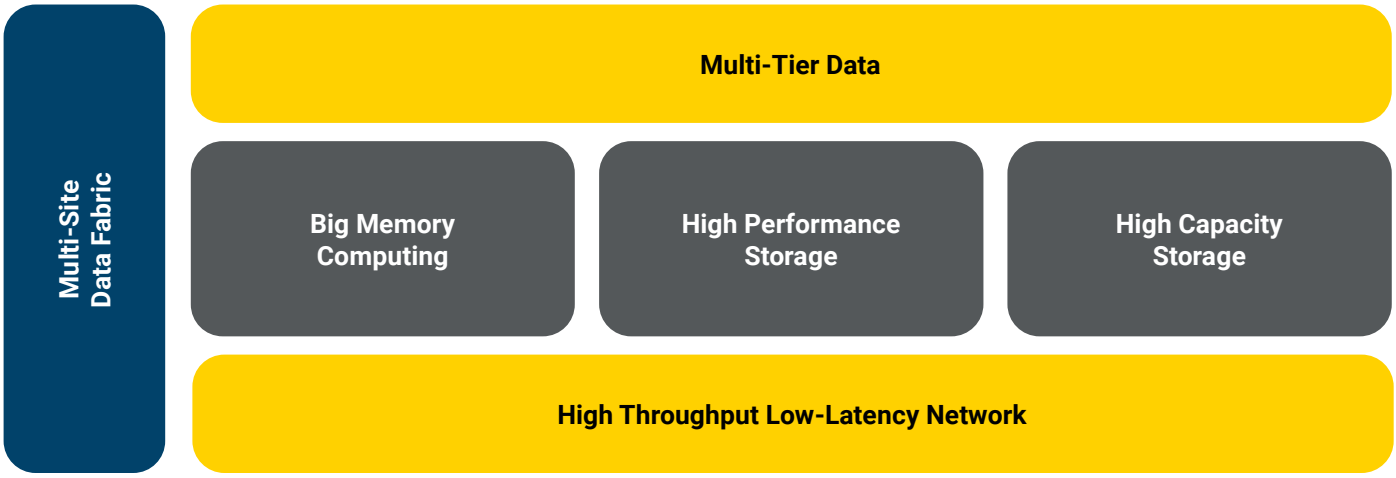

High Capacity Storage

HPC workloads often ingest or output massive amounts of data that must be kept in a general purpose storage environment when not being used for computing. TrueHPC can be paired with the DeepData™ solution to provide a scale out capacity optimized storage tier best suited for storing long-term data.

Multi-Site Data Fabric

Some HPC environments require connectivity to the cloud or another site. Some workloads might require cross-organizational collaboration on datasets that span multiple locations. TrueHPC can access data sets across the world as if they were local using the DataNexus™ solution.

Data Center Infrastructure

TrueHPC 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.

Tundra

Tundra racks are supported in many data centers designed for High Performance Computing (HPC) or Hyperscale environments, such as Penguin Computing data centers. Tundra racks can be deployed in 40OU or 44OU high x 762mm 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.

Open Unit: Each Open Unit (OU) is 1.889 inches tall - for comparison, a standard Rack Unit (RU) is 1.75 inches tall. Regarding height, a 40OU rack is equivelant to a 42U rack, and a 44OU rack is equivalent to a 48U rack.

Network Cable Serviceability: Network cabling is accessible from the front of a Tundra rack. This allows all systems to be serviced from the cold isle in an area of comfort, rather than from the rear of the rack where the hot air is exhausted. Network switches are mounted in the rack with their ports facing the front of the rack - requiring reversed airflow units.

Power

Tundra racks provide the same rack-level power aggregation, redundancy, and efficiency of OCP, and takes power delivery one step further by supporting up to 100kW of power per rack. These power optimizations make Tundra a robust platform for TrueHPC scale-out technical and accelerated computing environments.

Power Aggregation: Tundra racks use a rack-level power aggregation solution to deliver power efficiently to the systems in the rack. Three-phase 208V, 277V, or 408V power inputs are connected into Power Distribution Units (PDUs) that reside at the rear of the rack. These PDUs are redundantly connected to four Power Shelves that are installed in the rack. Power Shelves contain several Power Rectifiers, and are connected to three 12V copper Bus Bars that run vertically along the rear of the rack. The rear of a Tundra server has metal connectors that make contact with the copper Bus Bars when the server is mounted into the rack. The Tundra server receives power after the connection is made between the server and the Bus Bars.

N+1 Power Redundancy: A fully populated Tundra rack can require 30 to 100 kW of power. Power redundancy is provided at the rack level by using four Power Shelves, each having six Power Rectifiers. The Power Shelves are configured to provide enough power to satisfy a fully populated Tundra rack using N+1 rectifiers, where N is the number of rectifiers needed to fully power the rack, and an additional rectifier is installed to account for at least 1 rectifier failure without impacting the uptime of the systems.

Cooling

Removing the heat from a TrueHPC 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 that data center to prevent the environment from overheating.

Direct-to-Chip: Direct to chip liquid cooling involves bringing a liquid coolant via tubes directly to cooling blocks that are installed on the processor - replacing a standard air cooled heat sink. Liquid flows through the cooling block, carrying heat away from the chip and out of the server. Since this system cools processors directly, it's one of the most effective and efficient forms of data center heat removal. A Cooling Distribution Unit (CDU) aggregates the liquid from the systems at the rack level. The CDU can either transfer the heat to a local facility water supply, to the air in the data center, or to the air outside the facility. Direct to chip liquid cooling enables deploying high power dependent processor technologies that are impossible to deploy with air cooling alone.

Rear Door Heat Exchanger: Although a TrueHPC 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

TrueHPC is a comprehensive, end-to-end solution that organizations can leverage to jump-start their HPC 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, 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.

TrueHPC 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 TrueHPC needs - as a service.

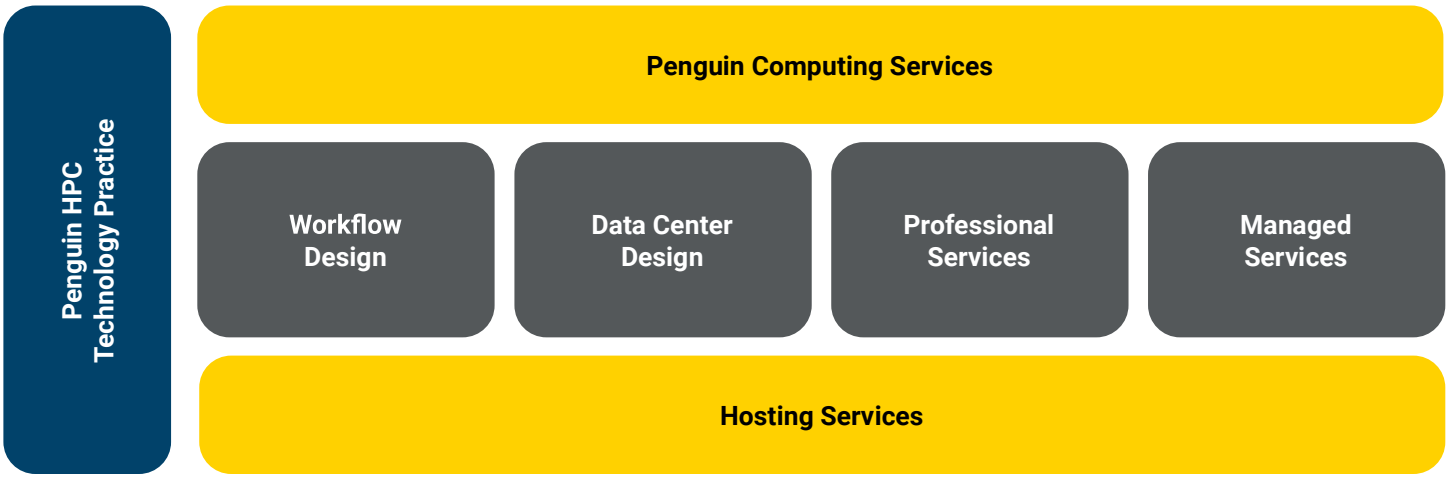

Design Services

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

TrueHPC provides a single, secure, end-to-end solution for HPC that includes a flexible, scalable, workload-optimized compute infrastructure. TrueHPC provides cluster management tools that allow you to not only quickly and easily provision technical compute environments but to efficiently monitor and manage them.

TrueHPC frees organizations from having to focus valuable time and human resources on creating a cluster architecture and management solution from scratch, allowing them 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 HPC initiatives 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 HPC initiatives with a ready-to-run solution that addresses compute infrastructure, software orchestration, data management, and infrastructure design to help you meet your project goals and drive innovation.

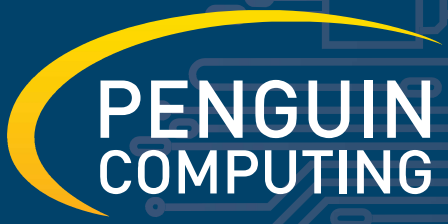

Expanding the world's vision of what is possible