

## POD Summary

---

Penguin Computing's HPC Cloud, POD, allows organizations to utilize a high-performance, bare-metal, HPC computing environment in the cloud without having to invest in on-premises infrastructure. POD eliminates many of the performance, scalability, and security challenges associated with the shared infrastructure of multi-tenant, cloud environments.

Jobs are easy to submit and monitor from either a traditional Linux CLI environment, or through an intuitive and secure web portal. Choose from CPU or GPU clusters. POD's HPC clusters are ready-to-run with hundreds of pre-installed applications, eliminating much of the complexity of building, managing, and scaling high-performance computing environments. This efficiency and economy of scale saves both capital and operational costs while ensuring a clear pricing model.

## Resources for any workload

---

POD's bare-metal, on demand HPC compute clusters are ideal for organizations in manufacturing, biosciences, research, energy, design, and finance - or any organization with high-performance computing needs.

Penguin Computing has over two decades of experience in optimizing HPC environments and applications.

POD users enjoy the included support from our HPC experts who assist in optimizing applications, designing and troubleshooting workflows, and getting the best experience with POD.



## 3D Remote Visualization

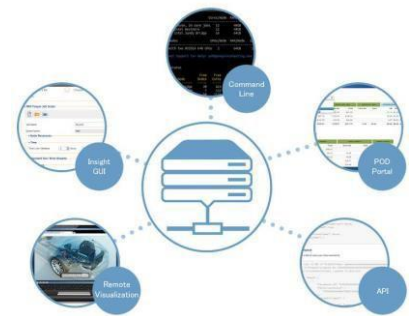
---

- On Demand 3D Remote Desktops
- Run real-time, interactive GUI workflows and 3D visualization
- Clientless remote desktop. No browser plugin or application necessary
- Designed for HPC desktop applications and post-processing tools
- High-end workstation performance
- NVIDIA® GRID technology provides H.264 quality with minimal latency
- Secure HTTPS access - no additional ports needed through your firewall
- Support of all major browsers (Chrome, Firefox, Safari, Internet Explorer)

## POD Features

---

- True HPC Computing - bare metal on non-virtualized nodes
- Free technical support from HPC experts
- Support staff with backgrounds in Science and Engineering
- No data transfer charges - in or out
- Predictable Billing - simple, pay-per-use
- No set-up charges



- Hosted in Tier III, US data centers
- SSAE 16 SOC1 Type II Audited data centers
- Redundant, secure Internet access through multiple Tier 1 network providers

## POD HPC Resources

---

### MT3 Cluster – AMD with GPUs

AMD EPYC 7642

- 48 Cores @2.40GHz
- 512GB of 3200MHz RAM
- 8 AMD Radeon Instinct MI50s per node
- 2TB local scratch
- Non-blocking Mellanox IB
- 1PB of Weka Storage
- 20 all flash storage servers
- Expandable and Tierable

### MT2 Cluster – Intel CPUs

Intel Skylake Nodes

- Dual Intel Xeon Gold 6148 @2.4GHz
- 40 cores per node
- 384 GB RAM per node

Intel Broadwell Nodes

- Dual Intel Xeon E5-2680 v4 @2.4GHz
- 28 cores per node
- 256 GB RAM per node / 9 GB per core

## High Speed Data Storage

---

- Billed for actual average monthly usage
- Private home volume per user, shared volumes can be configured
- IB Connected Lustre and WekaIO parallel file systems
- Local SSD Scratch available at no extra charge

## More Information

---

We are happy to answer any questions you may have about POD.

<https://pod.penguincomputing.com/contact>  
[podsales@penguincomputing.com](mailto:podsales@penguincomputing.com)

+1 (888)-PENGUIN (736-4846)