

DataNexus with Vcinity UltimateX

Data access without boundaries...



Contents

- DataNexus™ with Vcinity Ultimate X Solution-at-a-Glance 4**
- Introduction 4**
- Penguin Computing DataNexus with Vcinity Ultimate X 5**
 - Data Technologies 6
 - Data Access Without Boundaries 6
 - Rapid Data Transfer 7
 - DataNexus with Vcinity Ultimate X Features 8
 - DataNexus with Vcinity Ultimate X Use Cases 9
 - Software Technologies 12
 - Access X 10
 - Command X 11
 - Compute Technologies 12
 - High Performance Computing 12
 - AI/ML and Analytics 12
 - Compute Virtualization 12
 - Penguin Computing Services 15
 - Design Services 16
 - Professional Services 16
 - Managed Services 16
 - Hosting Services 16
- Conclusion 17**
- Contact Us 17**

DataNexus™ with Vcinity™ UltimateX™

Solution-at-a-Glance

Features

- **Industry-proven Technologies** — Stage and pressurize data across a WAN to reduce packet loss, reduce network latency, and increase network bandwidth utilization.
- **Connected Data Sets** — Federated data can be virtually consolidated over one or more WAN connections to create a world wide data set.
- **Management Interface** — Command X provides an easy-to-use, web-based GUI to monitor performance and configure multiple nodes.
- **Flexible File Management System** — Access X provides easy-to-use, web-based comprehensive file management and synchronization software.

Benefits

- **Unbounded Access to Data** — Provide remote access to data anywhere, any time from any location, in a hybrid environment.
- **Hybrid Cloud Integration** — Seamlessly connect on-premises to cloud while retaining full control of your data.
- **Rapid Data Transfer** — Move terabytes to petabytes of data predictably, economically, and securely at unprecedented speeds.
- **Improved Efficiency** — Maximize available infrastructure, speed up workloads and shorten application run times.
- **Lower TCO** — Reduce “copy sprawl” for lower storage costs and stronger data security.
- **Improve Productivity** — Streamline distributed workflows and enhance real-time collaboration.

Introduction

Enterprises will waste billions of dollars by 2025 moving files around as applications and data have to be collocated.

Explosive data growth and hybrid, multi-cloud infrastructures exacerbate the problem by copying files everywhere and moving racks and trucks full of data around. The pervasive move-copy paradigm also affects productivity due to the delays in the arrival of data.

Existing approaches require the transfer and replication of the data, leading to copy sprawl and compromised data security. Even with data transfer, data arrives at the destination in an unpredictable time and performance varies with data type/size and application.

Utilizing existing technologies and applications is so cumbersome that even physical transportation of media or data is considered an acceptable solution.

*What if the distance
between all of your
locations didn't matter?*

*What if you could work on a file
in Mumbai and from New York
or London or Sydney without
having to download, copy or
transfer it in any way?*

*What if you had a
unified view of your data across
the globe, and the power to access,
edit or move it on demand
from anywhere at record speed?*

Penguin Computing DataNexus with Vcinity UltimateX

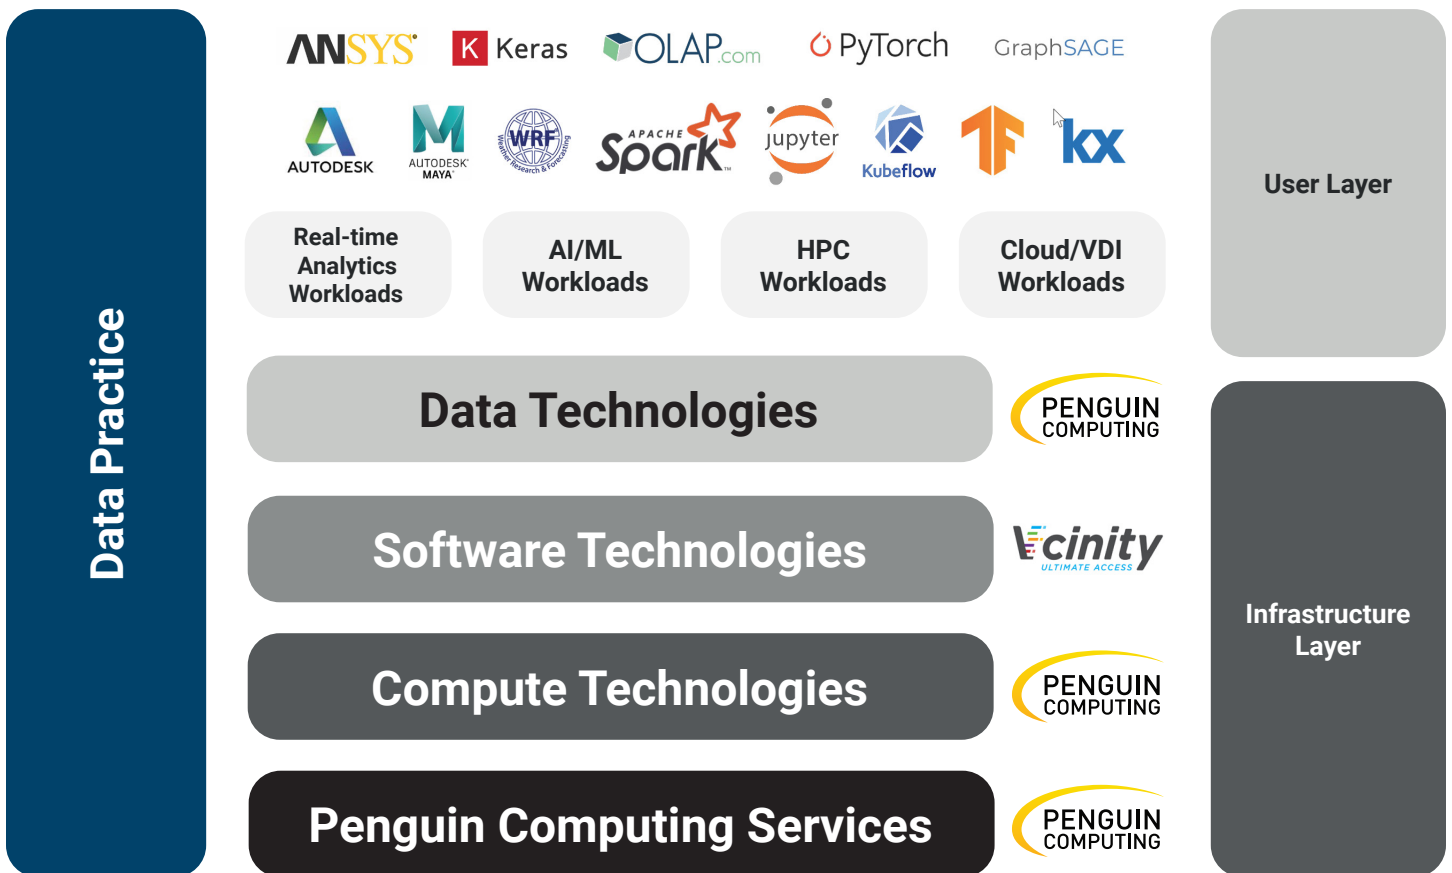
Penguin Computing collaborated with Vcinity to offer Penguin Computing DataNexus™ with UltimateX™, a data access solution that allows you to access and manipulate data remotely without the need to move it, or transfer it at unprecedented speeds when you do need to.


With the DataNexus no-copy, no-move paradigm, you don't need to collocate enterprise applications and data. You can collaborate on data sets without having to move it in or out of your on-premise, multi-cloud, or hybrid environments.

DataNexus transforms the WAN into a Global LAN, enabling instant, secure access to data anywhere without downloading, uploading, copying, pre-caching, or pre-staging. This allows for application- and data-agnostic operation to truly leverage enterprise data when and where it is created.

Penguin Computing DataNexus with Vcinity Ultimate X 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/Deep Learning, and Cloud to provide an end-to-end complete compute platform.

Penguin Computing DataNexus Components





Data Technologies

DataNexus with Vcinity Ultimate X connects geographically dispersed instances of compute and data. Where cloud delivered on the elasticity of infrastructure resources, DataNexus delivers on the flexibility of deploying these infrastructure resources. allowing you to work on anything from anywhere.

Data Access Without Boundaries

Using DataNexus with Vcinity Ultimate X, the distance between all of your locations no longer matters. You can work on a file in Mumbai and from New York or London or Sydney without having to download, copy or transfer it.

DataNexus Command X and Active X give you a unified view of your data across the globe, and the power to access, edit or transfer it on demand from anywhere. DataNexus gives you not just remote access, but also the capability to work on remote data, eliminating the distance and barriers across your distributed content. In real time, you can see edit and work on any file from anywhere, no matter where it says. No more downloading or replicating files.

Data in Europe can be processed by computing systems in the US without the need to move the data.

Rapid Data Transfer

Penguin Computing and Vcinity teamed up to showcase the unprecedented ability of DataNexus with Vcinity Ultimate X to migrate 1 Petabyte (PB) of data across the U.S. in just over 23 hours. The data was transferred using a 100Gbps WAN link across approximately 2,800 miles. The demonstration integrated Vcinity's technologies with those of Penguin Computing using industry-standard protocols.

The test continued for just over 23 hours to transfer a total of 1.0556 PB of data. The integrated solution was not optimized specifically for file transfer applications, yet it achieves optimal, predictable, and repeatable performance leveraging existing resources.

The test configuration consisted of six pairs of Penguin Computing DataNexus servers running Vcinity's ULT X software suite and RAD X-1040 WAN Fabric Extenders connected to a 100G Ethernet switch. Each RAD X-1040 WAN connection was rate-shaped at a maximum bandwidth of 16.666Gbps feeding the aggregate 100Gbps WAN connection with 70 milliseconds of latency, typical of a dedicated circuit between the U.S. East and West Coasts. A DataNexus server connected via InfiniBand to a RAD X-1040 fabric extender sat on each side of the WAN as connection bookends. Using file synchronization capabilities of Access XTM, the basic test procedure repeatedly sent datasets of mixed files and sizes totaling ~12.6TB per run from the source server bank to the target server bank for a total of 13 runs. Then the last run on four servers used a reduced fileset (so as not to greatly exceed an aggregate of 1 PB).

DataNexus with Vcinity UltimateX Features

Unprecedented Performance

- Near-real-time access to remote data as if it were local
- Transfer of terabytes to petabytes of data at the speed of light over virtually any distance (for example, 1PB in less than a day over a 100Gbps link across the U.S.)
- ~95% end-to-end bandwidth utilization of terrestrial and satellite links independent of size of the network connection

Maximum Scalability

- Predictable low-latency performance over any distance and independent of IT, LAN, WAN and transport protocols
- Linear performance with the addition of ULT X nodes, links, and bandwidth
- Remote edge applications enabled by software-only solution (ULT X-1000v)
- Policy-based multi-tenancy support
- Comprehensive policy and security (LDAP_credentials)
- Global Namespace support

Integration with Existing Infrastructure

- Seamless integration into existing workflows with no application rewrite
- Industry standard NAS with NFSv3/4 and SMBv3 connectivity
- Expandable tiered storage with embedded drives or external iSCSI/FC/IB attached storage for performance levelling across the network (ULT X-1000/1000e/1100)
- Deployable in any virtual environment—KVM, VMware ESXi, AWS EC2 VM, AWS F1, Microsoft Hyper-V, Xen or bare-metal (ULT X1000v)
- Data movement APIs for seamless integration with existing automated workflows

Ease of Operation

- Data access through simple network-mapped storage
- Rack-mount or travel case options for AC or DC powered, low power solution (ULT X1000/1000e/1100)
- Ruggedized portable version (ULT X-1000e)
- Integration with network managers via SNMP
- RESTCONF/NETCONF APIs for seamless integration into existing orchestration platforms
- Comprehensive file management (copying and moving files/folders) and synchronization features with flexible scheduling modes (Manual, Scheduled and Continuous) and powerful synchronization modes

DataNexus with UltimateX Use Cases

DataNexus provides an enterprise platform built on proven technologies to create worldwide data sets from federated data that spans data center, edge, and cloud environments. It offers compelling benefits for a variety of customers in media & entertainment, oil & gas, financial services, healthcare, and government.

Run applications using data from anywhere

Typically, data needs to be moved to users and applications prior to its processing, which slows down business agility while the user waits for transfer of data. ULT X remotely accesses data without moving it enabling live editing or processing of data in-place and accelerates workflows.

Rapid Data Transfer

DataNexus Rapid Data Transfer enables the movement of data among high performance servers at record speeds. See Rapid Data Transfer for details.

AI/ML/Deep Learning

AI/ML platforms are most effective when they have access to large-source data sets, however, these data sets are often not co-located with the compute resources. This creates a need to move the data which can be very expensive. Most organizations get around this by moving and using a small subset of the data, reducing the overall efficacy of the AI/ML platform. DataNexus allows the analytics platform to leverage the complete source dataset, without having to move or reduce the dataset, unleashing the full value of the data.

Migrations/ On-boarding/ Syncing

Moving workloads in and out of colo, public or private cloud platforms is the largest barrier for adoption of hybrid IT strategies. It becomes very costly to migrate workloads with the additional risk of placing wrong workloads onto the platform. ULT X moves data at unprecedented speeds, which enables workloads to be dynamically rotated in hybrid environments based on business/mission needs resulting in better SLAs and less risk.

Breaking down enterprise data silos

NAS silos limit enterprise flexibility with restricted access to enterprise information and complex security management and protection. ULT X's federated and consolidated NAS enables a unified data fabric with clients accessing remote NAS resources via their local ULT X. There is no client software needed and file shares are presented with standard SMB/NFS interfaces.

Agile data access for disconnected/disadvantaged environments

Intermittent cloud-enterprise connectivity poses the primary challenge for cyber foraging use cases. Existing protocols support limited data transfer ability and utilization of the available connectivity resulting in reduced agility of the edge node-based architecture and posing challenges to mission success. ULT X empowers edge node-based architectures to transfer data from command center/ cloud to the edge node for data pre-staging and caching at the tactical edge. It also allows the user to reach in, to process data in-place immediately where it is stored or to transfer data between edge node and command center/cloud during mission operations (that is, in the event of loss of connectivity, mission completion or data uploading).

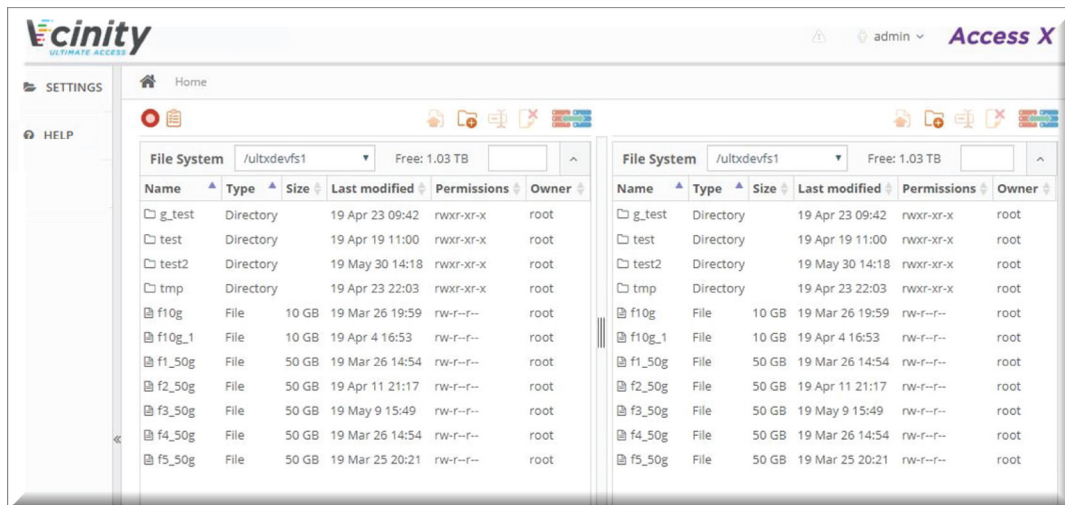


Software Technologies

DataNexus provides power software technologies for file management and synchronization and the monitoring and configuration of nodes:

- **Access X** – A tool for even more efficient, comprehensive file management and synchronization. Access X gives you unparalleled access and control and is ideal for enterprises that want to fully leverage the file sharing capabilities that DataNexus enables.
- **Command X** – A management system for an easy, intuitive user experience. It provides a web-based interface that enables monitoring and configuration of nodes across multiple sites.

Access X - Efficient File Management



Browse File Systems

The Access X Dashboard allows you to browse a file system by selecting a file system from the pull-down. Once a file system has been selected, browse the files and folders displayed in the pane. Select a folder to move down into the folder, and select the folder “.” at the top of the pane to move up out of the folder.

Copy or Move Files and Folders

You can drag files, folders, or collections between Source and Destination window panes, so either window may serve as the Source or the Destination window. Each drag and drop action is treated as a Task and you can view Task progress in the Status window.

You can also move files by using the Transfer function, which allows you to select pre-defined destinations, or new destinations to be defined. In addition, there is an option to select move or copy for the transfer.

Delete Files and Folders

The Access X Dashboard allows you to browse a file system by selecting a file system from the pull-down. Once a file system has been selected, browse the files and folders displayed in the pane. Select a folder to move down into the folder, and select the folder “..” at the top of the pane to move up out of the folder.

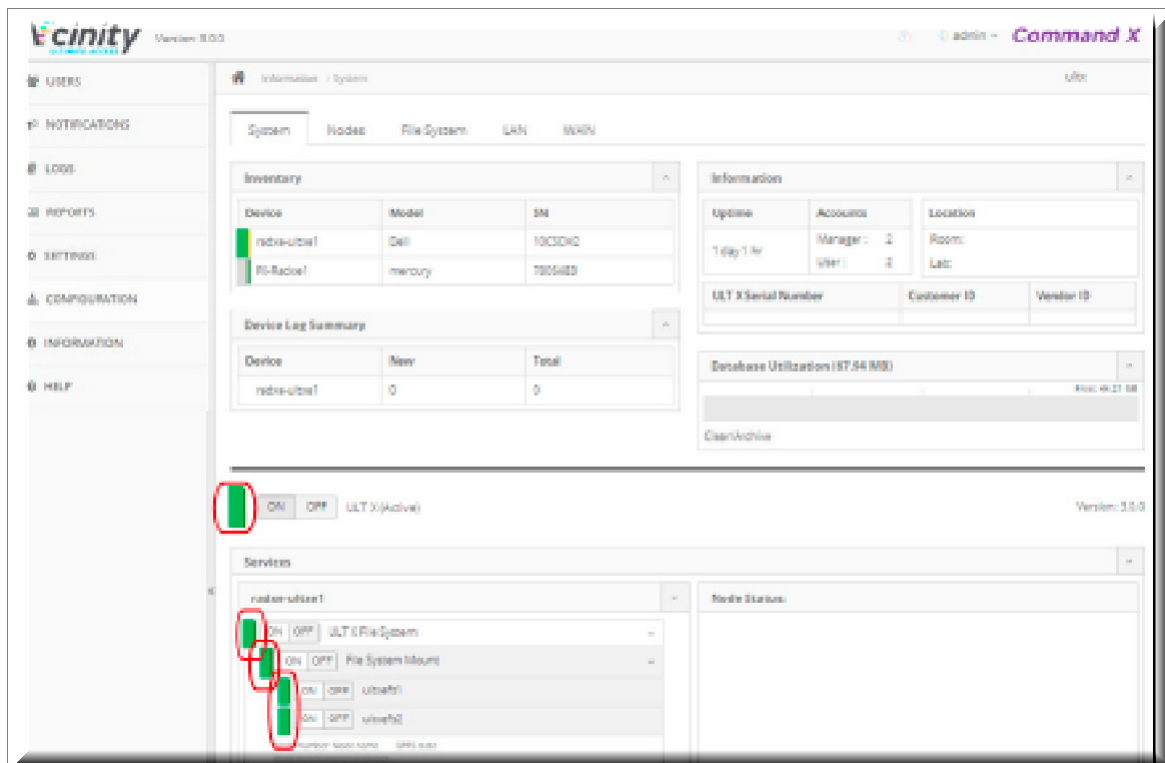
View Task Status


The bottom half of the Dashboard serves as the Status window. The Status window allows you to view task status.

Command X - Node Monitoring & Configuration

Command X is a robust Graphical User Interface (GUI) management tool for DataNexus with Vcinity Ultimate X. Command X™ configures ULT X nodes across multiple sites within the converged environment. It provides HTML GUIs for easy management and monitoring of infrastructure resources and data-level operations enabled by the solution. This results in reduced staff training and learning curves for an intuitive, seamless end-user experience.

Command X displays key UTL X-related information such as performance monitoring, error notifications, logging and account management. Unlike file transfer applications that require software installation at each client, UTL X is installed at each site. This plug and play deployment is further enhanced by Command X's easy-to-use management interface.





Compute Technologies

Enterprise workloads often require high core count, high clock speed, high memory bandwidth, low latency communication, and accelerated computing using GPUs, FPGAs, and ASICs. Powerful solutions leverage industry-leading technologies from Intel, AMD, NVIDIA, and other technology providers to enable a complete technology ecosystem supporting many different workloads.

DataNexus with Vcinity Ultimate X can be combined with Penguin's TrueHPC, OriginAI, and CloudBase solutions to provide heterogeneous computing environments using workload-optimized server building blocks for many types of enterprise workloads.

High Performance Computing

Penguin Computing TrueHPC™ provides a complete solution built on workload-optimized server building blocks and Scyld ClusterWare 11 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. DataNexus can be paired with TrueHPC to support HPC workloads in need of high capacity, scale-out storage.

AI/ML and Analytics

The Penguin Computing OriginAI™ solution uses AI software to provide a high-performance Deep Learning training environment for large-scale, multi-user AI software development teams. It includes the operating system (RHEL), cluster management (Scyld ClusterWare® 11), orchestration tools and workload schedulers (Red Hat OpenShift), AI libraries and frameworks, and optimized containers. DataNexus can be paired with OriginAI to support AI workloads with object storage.

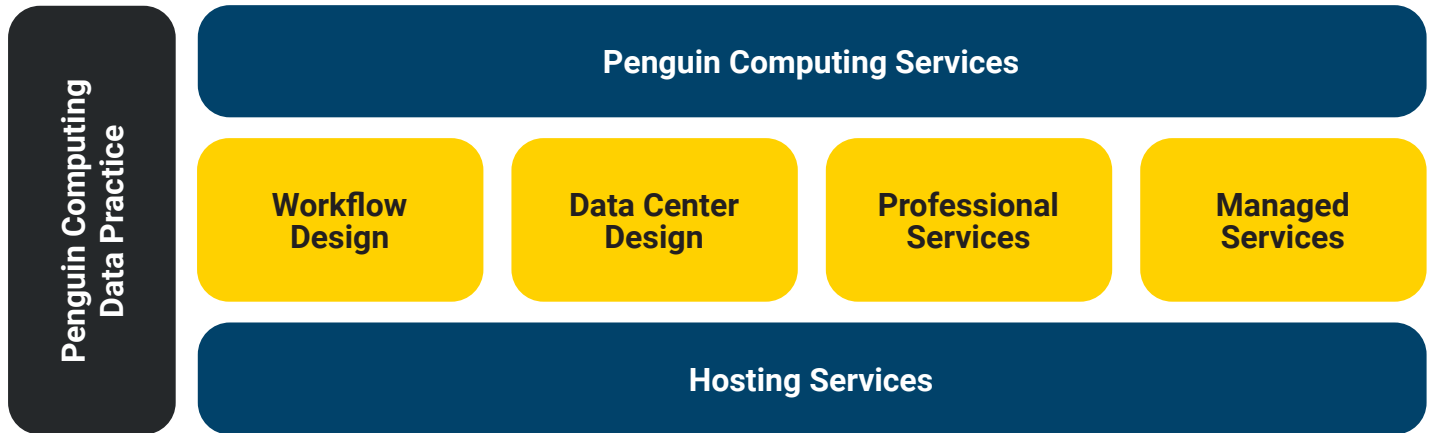
Compute Virtualization

Penguin Computing CloudBase™ with Red Hat OpenStack and/or Red Hat OpenShift is built on optimized server building blocks and cloud-native technologies to orchestrate an on-premises or hybrid cloud environment that enables workload portability and infrastructure flexibility for containerized and virtualized applications. DataNexus can be paired with CloudBase to support cloud-native environments with block and object storage.



Penguin Computing Services

Penguin Computing DataNexus with Vcinity Ultimate X is a comprehensive, end-to-end solution that organizations can leverage to jump-start their 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.

DataNexus 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 DataNexus 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 DataNexus with Vcinity Ultimate X allows you to solve your large data challenges by delivering a proven and tested turn-key enterprise solution that gives you the ability to provide unbounded and secure access to data in either a cloud or hybrid cloud environment. Seamlessly integrate and present as standard storage types without application changes.

If you have large amounts of data that you need to move, transfer, consolidate across multiple, disparate silos, DataNexus can help you do that as well, and at speeds up to 10X faster than other solutions.

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 turnkey data extender solution that can:

- Provide remote access to data anywhere, any time from any location, in a hybrid environment.
- Move terabytes to petabytes of data predictably, economically, and securely at unprecedented speeds.
- Seamlessly connect on-premises to cloud while retaining full control of your data.
- Maximize available infrastructure, speeds workloads and shortens application run times.
- Reduce “copy sprawl” for lower storage costs and stronger data security.
- Streamline distributed workflows and enhance real-time collaboration.



**PENGUIN
COMPUTING**

Expanding the world's vision of what is possible