

## pixstor Technical Specifications & Features

### Scalability and Performance

- Guaranteed Performance up to 99% capacity utilisation
- Maximum Capacity:  $2^{99}$  Bytes
- Maximum Number of Files and Directories:  $2^{64}$
- Max Throughput Performance Per Storage Server: 20GB/s each Max
- Throughput Performance Per Gateway Node: 10GB/s each Max IOPS
- Performance Per Storage Server: 250,000 4K Random
- Maximum Number of Storage Arrays Per 4 Storage Servers: 8, plus expansions
- Native Clients available for: **Linux | Windows**
- Max Performance per Native Client: **Linux:** 20GB/s | **Windows:** 4GB/s
- Max Performance Per NAS Client: **NFS (\*nix) | 10Gb:** 1GB/s | **25Gb:** 1.5Gb/s
- SMB (Windows): **10Gb:** 1GB/s | **25Gb:** 1.5GB/s
- SMB (OSX): **10Gb:** 1GB/s | **40Gb:** 1.25GB/s
- Support Disked Types: NL-SAS, SAS, SSD, Flash, NVMe
- Local Read Only Cache: SSD/NVMe drives in Linux Native Clients and Gateway nodes that significantly reduce access latency for applications such as Houdini. Up to 10X increase in simultaneous renders with small investment
- RDMA over Converged Ethernet Support (Linux): Ultra Low latency connectivity for line rate bandwidth, minimal CPU overhead and a small number of application threads
- NVMeOF Support: Remotely hosted NVMe drives delivered over commodity 40, 50 and 100Gb Ethernet with no loss in performance delivered direct to clients (Excelero NVMeMesh)
- Memory Caching: Storage and Gateway node supports up to 1TB of memory-based cache
- No separate metadata services required - each storage node acts as a metadata server
- Node-based Quality of Service for limiting the impact of either management or user-based activity on the system
- Network-level bandwidth shaping for preventing over subscription from ingest, data movement and render based tasks
- Large block size support (up to 8MB) without sacrificing capacity or small file performance

### Data Management & Protection

- Up to 7 pools of storage in every pixstor namespace (for a maximum total capacity of  $2^{99}$  bytes), each with their own unique capacity and performance characteristics
- Automated data placement, migration, replication, extended attribute (tags) and deletion policies, with no-tree-walk candidate selection
- Synchronous Replication, both at a granular level and for the entire namespace
- Asynchronous Replication and Disk-Based Backups with pixstor Sync Fileset functionality, providing “filesystem within a filesystem” granularity for data placement, snapshots and quotas
  - Up to 256,000 “Copy-on-Write” Snapshots
  - Automated snapshot creation and deletion
  - Snapshots tightly integrated with Windows Previous Versions for self- service user restores across SMB
  - File clones for writeable snapshots
  - User, Group and Fileset Directory Quotas
- Immutability and WORM support
- Automated Compression Policies with transparent decompression
- Duplicate File Identification for identifying wasted capacity
- Transparent migration and recall of data to third party cloud, object, tape and NAS resources with pixstor and ngenea
  - Capable of migrating data to multiple external targets simultaneously
- Data migration from legacy services with ngenea Reverse Stubbing technology

## Scale-Out NAS

- N+Many Active NFSv3, NFSv4, SMB2, SMB3, FTP, S3 and SFTP services across servers and network interfaces
- Industry Leading OSX/MacOS support for SMB (AAPL Samba Extensions)
- Scales out with pixstor Gateway nodes
- Granular configuration down to the individual client for workload-specific tuning without compromise
- Multiple Subnet and Tagged VLAN support
- Unified NAS and Object support

## Systems Management & Programmability

- Web-Based Administration, Performance and Health Dashboards
- CLI-based administration, performance and health scripts
- Email and Slack-based Alerts
- Event/Response engine for automatically performing tasks based on file system events (node loss, file system mount, disk failure, etc)
- Web-Based Capacity Analytics for determining who, what and where capacity is being consumed, including data that has been migrated by ngenea
- InfluxDB hosted metrics for easy access from 3rd party monitoring and graphing tools
- Customisable performance dashboards
- Salt-based configuration management, deployment and upgrade automation
- Python API for majority of filesystem management operations, allowing tight integration of the pipeline/workflow with the storage
- REST API for common filesystem, snapshot and share management activities
- Python-wrapped C API for performance intensive scan and update tasks

## Search

- Web-based UI for searching through assets, using 'human' search terms
- Directed search filters for targeted, granular searches
- Automated indexing of data with deep metadata inspection
- Automated proxy generation
- Plugin Based Architecture for rapid generation of new metadata inspection, proxy generation and front-end functionality
- REST API
- Desktop integration for opening results directly on the desktop (Chrome support)
- Filelist exports based on search results for importation into third party applications (BIN generation, etc)
- Integration with popular AI tools such as Machine Box and Microsoft AI
- Scale out architecture for large deployments and ingests

## Security & User Management

- Windows, NFSv4 and POSIX
- ACL Support
- Active Directory, LDAP and Local User Authentication
- Pluggable ID Mapping: AD/LDAP Source
- Repeatable Cross-Platform RID-based ID mapping
- Autogenerated Mapping
- Multi-tenancy and share isolation via Container Support (Control Plane/Data Plane Separation)
- Server-side Firewall Protection
- Sudo-wrapped filesystem communication
- Optional SSL-Encrypted transfers for cross-cluster communication
- SSL encrypted data transfers to supported ngenea object targets
- Kerberized NFSv4, SMB signing and encryption
- Oauth2 Compliance for all Web-based and REST services
- "Air-gapped" Upgrade Capabilities

## Current Support For

- Amazon S3
- Azure Blob
- Google GCS
- Spectra Logic Black Pearl
- Scalify
- Caringo Swarm
- IBM Cloud Object Store
- NetApp Storage Grid
- EMC ECS
- HGST Active Scale
- SwiftStack
- Swift (S3)
- Ceph (S3 GW)
- Isilon
- Nexenta
- Qumulo
- StorNext

## Collaboration & Data Sharing

- Remote caching functionality for data sharing across large geographical distances
- “Reverse stubbing” for publishing data to remote pixstor installations from both pixstor and 3rd NAS solutions such as Isilon and Qumulo - both on prem and in the cloud with ngenea
- Pre-migration with ngenea for pushing data to remote cloud, object and NAS solutions, including geo-replication (AWS, GCS)
- Cloud bursting support
- Multi-cluster support provides secure, remote mounting of file systems between multiple pixstor clusters

## Cloud

- Full push-button deployable in the Cloud
- Data can be cached, ingested and shared with on-premise resources
- Marketplace availability for **AWS** and **GCP**
- Data-tiering from cloud-based block storage to GCS or S3 to reduce capacity costs

## Data Management & Protection

- Up to 7 pools of storage in every pixstor namespace (for a maximum total capacity of  $2^{99}$  bytes), each with their own unique capacity and performance characteristics
- Automated data placement, migration, replication, extended attribute (tags) and deletion policies, with no-tree-walk candidate selection
- Synchronous Replication, both at a granular level and for the entire namespace
- Asynchronous Replication and Disk-Based Backups with pixstor Sync Fileset functionality, providing “filesystem within a filesystem” granularity for data placement, snapshots and quotas
  - Up to 256,000 “Copy-on-Write” Snapshots
  - Automated snapshot creation and deletion
  - Snapshots tightly integrated with Windows Previous Versions for self- service user restores across SMB
  - File clones for writeable snapshots
  - User, Group and Fileset Directory Quotas
- Immutability and WORM support
- Automated Compression Policies with transparent decompression Duplicate File Identification for identifying wasted capacity
- Transparent migration and recall of data to third party cloud, object, tape and NAS resources with pixstor and ngenea
  - Capable of migrating data to multiple external targets simultaneously
- Data migration from legacy services with technology
- ngenea Reverse Stubbing

## About arcastream

arcastream equips teams with smarter, better orchestrated free-flowing data to power faster discovery, solving and innovation.

Our purpose-built, software-defined storage and data solutions work to complement and improve your existing infrastructure and provide a sustainable way to manage rapidly growing infrastructure.

Customer success is at the heart of our business. We offer solutions that fit your objectives and technical infrastructure, working with all major hardware and public cloud providers to do so.

arcastream is privately-owned and headquartered in the UK, with offices in the USA and Germany.

---

### For more information on **arcastream** solutions:

UK Sales: +44 (0) 345 058 2374

e: [info@arcastream.com](mailto:info@arcastream.com)

w: [www.arcastream.com](http://www.arcastream.com)