

PanView Analytics

High-Performance Data Insight Software for PanFS

Maximizing the value of existing data is critical for high-performance computing (HPC), artificial intelligence, and machine learning (AI/ML) workloads. As unstructured data within organizations continues to scale to petabytes, understanding that data becomes just as important as storing it.

Data-driven organizations need comprehensive insight into their storage environments to understand their space consumption and asset availability. Compounding this challenge, multiple vendors are often used for storage, creating data silos and storage islands. Organizations that lack insight into their data stores risk losing their competitive advantage.

Panasas provides PanView™ Analytics, a comprehensive and easy-to-use data insight solution for HPC and AI/ML environments that helps organizations extract full value from their data. Combined with the included PanMove™ Advanced software, users can discover, search, and analyze their data files to mobilize them where and when they are needed. PanMove Advanced and PanView Analytics are the result of a collaboration between Panasas and Attempo, a leading provider of data protection and data management software.

PanView Analytics provides clear analytical information and powerful search capabilities that result in better decision making and reduced operating costs for infrastructure and personnel. Users access PanView Analytics through a single centralized dashboard that presents a global view of Panasas storage systems as well as backup and archive destinations. PanView Analytics provides a graphical representation of the data lifecycle and creates customized reports.

State-of-the-art HPC and AI/ML solutions require storage solutions that are the highest in reliability, scalability, and performance, and they must also be simple to manage. Whether it's pushing the boundaries in life sciences, manufacturing, energy, financial services, media and entertainment, or academic and government research, Panasas delivers the PanFS parallel file system with the PanFS software suite that includes PanView Analytics for high-performance data insights.

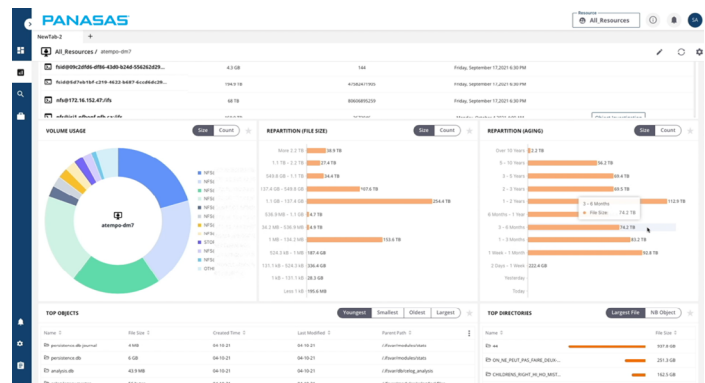


Figure 1. PanView Analytics Overview Page.

PanView Analytics Benefits

- Reduce total cost of ownership (TCO) by optimizing different storage tiers.
- Analyze data usage metrics.
- Identify data via metadata.
- Gain global understanding of storage and data.

Take Action

With PanView, customers are enabled to act on their infrastructure and data:

- File metadata visibility (date, size, type, etc.).
- Capacity planning and keeping storage costs under control.
- Consolidating storage from multiple silos.
- Identifying “cold” data to transfer to lower-cost storage.

View Reports

PanView is accessed through a single centralized dashboard:

- Global view of Panasas storage systems and backup and archive destinations.
- Graphical representation of data lifecycle (“hot” and “cold” data).
- Consolidated views and creation of customized reports.

ActiveStor Storage Nodes

PanView Analytics is supported on ActiveStor® Ultra, ActiveStor® Flash, ActiveStor® Ultra XL, and ActiveStor® Classic storage nodes. All nodes are visible through the same graphical interface.

Web Interface

PanView Analytics comes with an easy-to-use Web interface that includes a navigation pane providing access to the following tabs:

- Dashboard – Present information on the selected resource and customize the dashboard to show metrics and analytics.
- Analytics – Explore and organize metrics for the selected time period and resource types.
- Search – Perform advanced file search operations by path, file name and tag.
- Log out – Sign out from the application.

Data Visualization

Data visualization can be defined based on the selected resource, date range, and data exploration range. The metric types are selectable from a predefined list of categories that include:

- Child details – Child details are children of the resource selected.
- Object investigation – This button enables you to investigate the selected object and volumes.
- Volume usage – The volume usage on present child categories is displayed in a pie chart.
- Repartition (file size) – The files that are available for the selected data are categorized by file size in a horizontal bar chart (e.g., between the range of 5 GB–10 GB file size there could be a total of 20 GB of files).
- Repartition (aging) – The files that are available for the selected data are categorized by age in a horizontal bar chart (e.g., between the range of 3–5 years duration there could be a total of 10 GB of files).
- Top objects – Top objects consist of most-recent search operations performed on files and directories.
- Top directories – The largest files and directories are displayed with order selectable between the Both/Files/Directories categories to represent the largest files.

About Panasas

Panasas builds a portfolio of data solutions that deliver exceptional performance, unlimited scalability, and unparalleled reliability – all at the best total cost of ownership and lowest administrative overhead. The Panasas data engine accelerates AI and high performance applications in manufacturing, life sciences, energy, media, financial services, and government. The company's flagship PanFS® data engine and ActiveStor® storage solutions uniquely combine extreme performance, scalability, and security with the reliability and simplicity of a self-managed, self-healing architecture. The Panasas data engine solves the world's most challenging problems: curing diseases, designing the next jetliner, creating mind-blowing visual effects, and using AI to predict new possibilities.

Worldwide Office
1-888-PANASAS
info@panasas.com

Panasas Headquarters
San Jose, CA, USA
Panasas Research & Development
Pittsburgh, PA, USA

Panasas EMEA
Oxford, United Kingdom
emeainfo@panasas.com

Panasas APAC
Sydney, Australia
apacinfo@panasas.com

Panasas China
Shanghai, China
chinainfo@panasas.com

Search Query

The Web interface includes an object explorer that allows easy location and identification of files and supports simple and advanced search queries. In addition, filters can be converted into queries.

The search path includes the resource name within which you perform the search. To specify searches, the following fields are supported:

- Field name (e.g., file size).
- Operator (e.g., equal).
- Setting a value (e.g., 5).

Dashboard

PanView Analytics also comes with a dashboard that provides an overview and a customized dashboard section. The dashboard overview, shown in Figure 2 below, displays information on the snapshot taken for the selected resource including:

- Selected resource and its total, occupied and available space.
- Number of resources inside the selected resource.
- Number of directories and files within them.
- Number of platforms and volumes.



Figure 2. PanView Analytics Dashboard.

The customized dashboard can show metrics and analytics for a specified resource.

Timely, High-Quality Service and Support

Unlike open-source solutions and even commercial alternatives from broad portfolio vendors, Panasas offers timely, world-class support.

More Information

For more information about PanView Analytics, contact your local Panasas representative or visit panasas.com/products/panview.