panzura

Panzura White Paper

Hybrid Cloud NAS for On-Premise and In-Cloud File Services with Panzura and Google Cloud Storage

By: Rich Weber, Product Management at Panzura

This article describes how Panzura works with Google Cloud Storage

Panzura Freedom™, an intelligent hybrid cloud storage solution, consolidates storage, replaces legacy NAS, and enables real-time, global collaboration without limits. With Panzura Freedom, enterprises gain the economics, scalability, and durability of the cloud while achieving local and in-cloud data center performance.

Panzura's integration with cloud storage provides enterprises in media and entertainment, genomics, life sciences, healthcare, oil and gas, financial services, and other industries the ability to deploy local and performant file services on premises at each location as well as in the compute cloud that consolidates all your data in Google Cloud Storage. Data is cached locally in each Panzura filer while storing the authoritative data in Google Cloud Storage. Panzura CloudFS™, the first enterprise file system purpose built for the cloud, provides you with a global namespace and real-time file locking services under the same UNC path for both NFS and SMB file services.

The following diagram shows how Panzura Filers integrate with on-premises resources, Google Cloud Storage, and Google Compute instances running on the Google Cloud Platform.

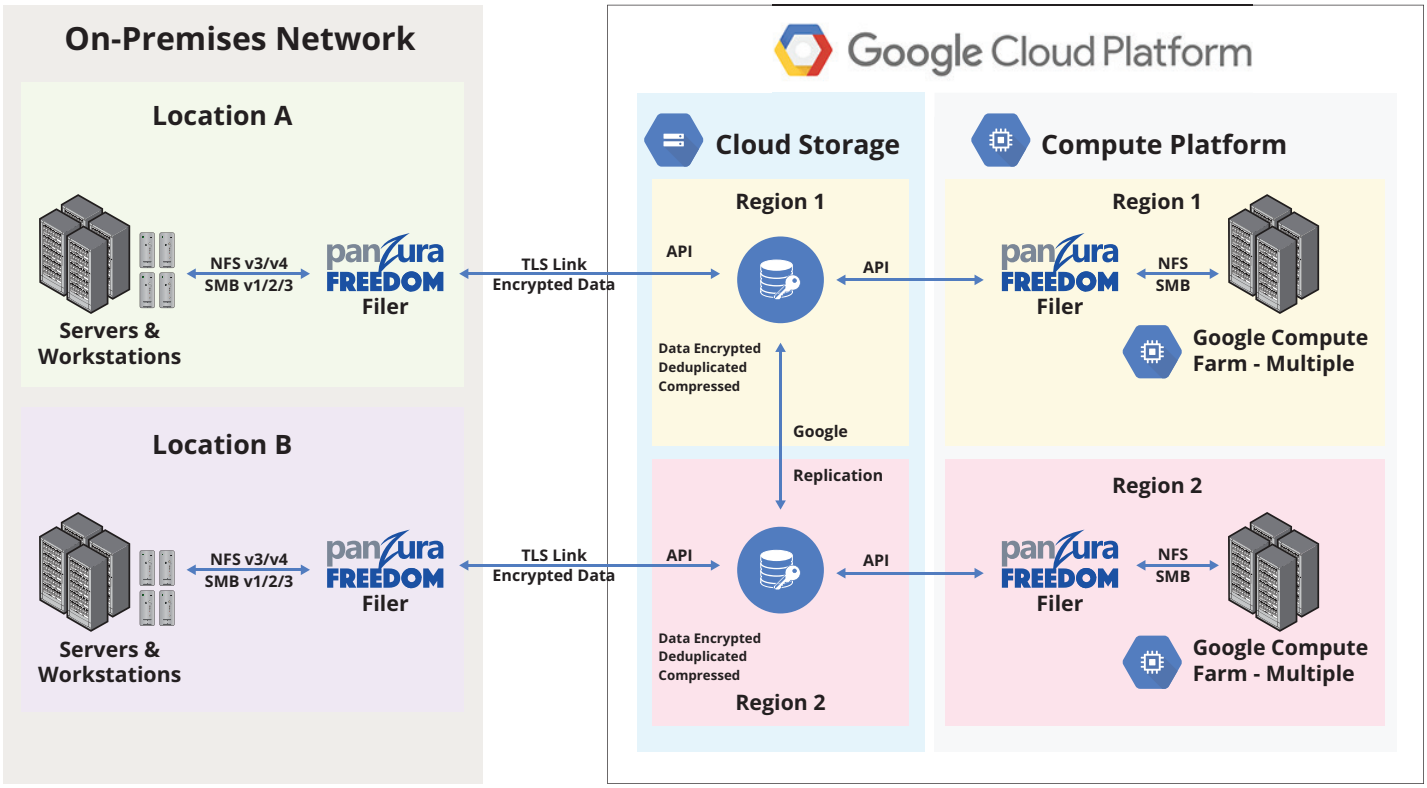

Panzura Freedom offers enterprise customers the ability to integrate their native file services in their on-premises data centers as well as inside the Google Compute environment with Google Cloud Storage. This provides a single cohesive file data storage solution with local file performance to a centralized storage solution. Each Panzura Filer provides local caching services where data is automatically and intelligently cached, pulling data directly from the Google Storage platform. Data is globally deduplicated, compressed, and encrypted using the customers crypto keys. The Panzura solution provides immediate file consistency with global file locking services as needed.

The Panzura Freedom Filer can be deployed locally as either a virtual machine or on Panzura provided hardware. Deployment of the Filer can be launched as a Google Compute instance for in-cloud direct access to the same data available on-premise as part of the overall Panzura CloudFS global file system. With the integrated Panzura hybrid cloud NAS solution integrated with GCP you see the following benefits:

- Maximum scale, performance, and resource utilization
- Real-time, global file and byte-range locking
- File services to power data center workloads and mobile clients
- Freedom-based business and deployment models
- Automation and centralized management
- Complete cloud data protection and high-availability with no single point of failure

Panzura allows you to manage your file services as a single solution for both on-premise and in-cloud compute workloads. Panzura CloudFS puts the data near the applications and users where it is needed, regardless of where those services are. Some of the features included as part of Panzura Freedom include;

Protocol Support: Panzura supports the NFS v3/4 and SMB v1/2/3 file protocols for connection to compute farm servers and user workstations.

Storage Efficiencies: By splitting the metadata and data, and storing them in Google's storage cloud, each system has access to the entire data footprint in a single namespace. This allows the Panzura solution to perform live global deduplication and compression automatically reducing the overall storage footprint.

Local Caching Services: The Panzura Freedom Filer's automated intelligent local caching services maximize system responsiveness to user/machine requests and provide local data center performance with LAN-speed access to files in a global cloud storage system. The intelligent caching system dynamically and automatically manages where data is cached for maximum performance. Data caching ensures LAN-speed access to specific files regardless of their age or usage patterns.

Global Namespace: Metadata for the entire file system is automatically synchronized across all Filers configured in CloudFS, through objects stored in Google Cloud Storage so that each mounted client has access to the same data through a common UNC file path regardless of that Filer's geographic location.

Global File Locking: File locking services are present locally for read/write access at each Filer regardless of the protocol used, with remote Filers in read-only at a minimum. Distributed global read/write locking services is available for SMB protocol based workflows as needed.

SmartSync: SmartSync provides real-time locking services with immediate data consistency, even if that data hasn't been updated in the cloud through peer-to-peer services that exchange not only locking operations, but transfer changed block data or pointers between Filers enabling real-time collaboration for SMB workflows.

Data Encryption: With customer owned and managed keys, Panzura Freedom will automatically encrypt all data before it is loaded on Google Storage allowing customers to secure their data not only in transit but at rest in the cloud with the Panzura FIPS 140-2 compliant solution.

Global Snapshots: Nearly limitless snapshots of your file system are supported, providing you with the ability to retain hourly, daily, weekly, monthly, and yearly snapshots over long periods as required by your internal policies or external regulatory requirements.

Mobile Data Access: Panzura Mobile provides mobile access to the same global data on Android, IOS, Windows, Mac, and Web-based clients.

Support for Google Storage Tiers: Full support for use of Google Near Line and Cold Line Cloud Storage tiers providing NFS and SMB file services on the front end with all tiers of Google object storage in the cloud.

RESTful API: Panzura Freedom provides a fully integrated API allowing you to programmatically perform operational and configuration operations as needed.

Use Cases

Deploying workloads on Panzura Freedom gives you the ability to access your data both on-premises and in the Google Compute Cloud while consolidating your data in Google Cloud Storage. Many industries are embracing the use of cloud technologies to improve productivity, reduce cost, and utilize cloud compute resources for both bursting and expansion of their own local capabilities.

Rendering for Media and Content Creation for Media and Entertainment

- Provide a centralized cloud data model with local access making rendering workloads possible both in-cloud and on-premise without a transport event.
- Utilize in-cloud compute resources for rendering workloads as needed.
- Reduce cost of the overall storage solution through consolidation and elimination of replicated copies of data.
- Boost productivity and meet deadlines by utilizing in-cloud compute resources working on the same data you have locally.
- Automatically encrypt all data using customer loaded crypto keys.
- Reduce network overhead consumption by eliminating duplicated data replication between data centers and/or in-cloud resources.
- Remove infrastructure headaches and let creative talent concentrate on content.

Processing Genomics for Life Sciences

- Provide a centralized cloud data model with local access of large research datasets both in-cloud and on-premises without a transport event.
- Provide high performance local access both on-premises and in-cloud to research datasets for both staff and automated workloads.
- Reduce cost of overall storage solution through consolidation and elimination of replicated copies of data.
- Automatically encrypt sensitive research datasets using customer loaded crypto keys.
- Expand existing research data workflows in the cloud while accessing the same data.

Running Simulations and Archiving Data for Financial Services

- Provide a centralized cloud data model, with local access for financial analytic and simulation workloads, possible both in-cloud and on-premise without a transport event.
- Ingest long-term data retention on-premise and automatically store in Google Cloud Storage taking advantage of Near Line and Cold Line storage tiers for additional cost savings.
- Expand your analytics compute as much as required without the headache of data transport or replica data management.
- Make the results of analytical and simulation data processing immediately available to on-premise locations through a unified file storage solution.

Seismic Data Processing in the Cloud for Oil and Gas

- Provide a centralized cloud data model, with local access for seismic data sets for processing analytic workloads, possible both in-cloud and on-premise without a transport event.
- Long term data retention or archival of seismic and other data in Google Cloud Storage, taking advantage of Near Line and Cold Line storage tiers for additional cost savings.
- Expand your seismic data processing compute capacity as much as required without the headache of data transport or replicated data management.
- Make the results of seismic data processing and analytics immediately available to on-premises locations through a unified file storage solution.

Medical Imaging Storage in the Cloud for Healthcare

- Fully HIPAA and FIPS 140-2 compliant solution.
- Long term retention for medical images and other related unstructured data in Google Cloud Storage, taking advantage of Near Line and Cold Line storage tiers for additional cost savings.
- Fully encrypted data both in transit and at rest in Google Cloud Storage using customer owned and managed crypto keys for regulatory compliance.
- Access to the same global data sets from any location or in-cloud compute environments.

File Shares and Home Directories

- Panzura Freedom hybrid cloud NAS consolidates all of your file shares and user data in Google Cloud Storage, eliminating replicated data and providing a single global version of the truth without having to manage multiple versions of your data.
- Full cloud based disaster recovery options. With Panzura Freedom, all your data is automatically backed up in Google's resilient Cloud Storage. Data is recoverable from any location through either HA options or a full cloud restoration.
- With all of your data automatically consolidated and stored in Google Cloud Storage and virtually limitless snapshots, the need for onsite backups is all but eliminated.
- With global file locking and a global namespace users can easily collaborate on the same data even across multiple geographically dispersed sites.
- Data is cached locally for LAN access performance speed. Automated caching policies make it easy with advanced caching policy manipulation for the most demanding workflows.
- Mobile client support that easily puts your data in the hands of traveling employees or 3rd party collaborators when needed.

Database and VM Backups in the cloud

- Panzura can be your backup target for SQL dumps and VM backups that automatically store your data in Google Cloud Storage using either standard, near line or cold line tiers at an affordable cost over long retention periods.
- Fully encrypted data both in transit and at rest meeting your regulatory compliance requirements with FIPS 140-2 and HIPAA compliant Panzura solution.
- Use Veeam to backup your VMs on-premises and restore them in the cloud to migrate or expand your virtual environment into the cloud compute space.

What's Next

Read more about the Panzura Freedom Filer solution;

- Panzura and the Google Cloud Platform
- Panzura Freedom Filer Solution
- More Panzura resources and documentation

About Panzura

Panzura is the leader in enterprise hybrid cloud NAS and cloud data management. Panzura Freedom, an intelligent hybrid cloud storage platform, represents a breakthrough in managing explosive growth in unstructured data, delivering data center performance with the economics, scalability, and durability of the cloud. Panzura CloudFS underpins Panzura Freedom and incorporates intelligent file services backed by 27 patents. Using Panzura, enterprises shift from legacy storage as they consolidate and simplify their storage into the cloud. Organizations like Milwaukee Tool, American College of Radiology, and Chevron use Panzura for active archive, hybrid cloud NAS, and cross-site collaboration.

Please visit www.panzura.com for further information.

Panzura, Inc. | 695 Campbell Technology Pkwy #225, Campbell, CA, USA | 855-PANZURA | www.panzura.com