PURESTORAGE®

RADIOLOGICS

SOLUTION BRIEF

Accelerate Medical Imaging Research with Radiologics on Pure Storage®

Accelerate workflows with a purpose-built high-performance computing architecture.

Medical imaging modalities such as x-rays, MRIs, CTs, and digital pathology produce close to 450 petabytes of data per year.[1] Clinical researchers visually interpret images to better understand cancer and neurological conditions and are increasingly using machine learning and artificial intelligence (ML and AI) to augment their capabilities.

Consequently, imaging puts intense pressure on storage and compute. Many teams deploy XNAT, an open-source imaging informatics platform that supports ML and annotation workflows, especially through the Radiologics application suite. However, traditional high-performance computing (HPC) filesystems can't handle these workloads efficiently because imaging datasets contain multiple small and large files with high metadata. To help solve this, researchers and IT leaders at life-sciences organizations are reconsidering their data infrastructure to get the most value from their imaging data.

Boosting Medical Imaging

Imaging is the dominant tool used to diagnose human health and disease. To augment physician and researcher capabilities, imaging is becoming increasingly quantitative and highly dependent on advanced computational processes. Unfortunately, supporting infrastructure has been slow to follow.

Today's radiology is powered by picture archival and communication systems (PACS) which often don't have advanced image computing capabilities. In addition, medical imaging data produces file size variability that traditional HPC file systems are not optimized to handle: The average file size for MRIs and CTs, for example, is less than 50KB but can go as high as 900MB for digital pathology images.

Speed Up with Agility

- Greater than 10x faster data ingestion from PACS into XNAT
- Self-tuning for diverse file workloads

Rapid Scaling

- Scalable from 66TB to petabytes without downtime

Single Platform

- Store PACS, VNA, XNAT, training, and models data on one platform

XNAT easily ingests this varied data from a clinical PACS and makes it available for research. But moving that data from clinical PACS into XNAT can be very slow on legacy HPC storage because it needs to be tuned for each different data type. The lack of simplicity creates deployment and runtime bottlenecks.

Pure Storage has partnered with the Radiologics XNAT Computational Radiology Platform, which sits side-by-side with enterprise PACS and EMR systems to integrate patient analytics and enable the entire pipeline from XNAT to downstream AI applications and model deployment for researchers and clinicians. Together, Pure and Radiologics bring a scalable turnkey solution for clinical research in biopharma and academic medical centers.

XNAT Works Better on Pure

Pure Storage FlashBlade® supports modern cloud-like architectures such as XNAT, is purpose-built for HPC workloads, and can self-tune for different I/O types, such as small or large files and metadata. FlashBlade performs significantly faster than traditional HPC environments by automatically adjusting to variable file types. In fact, XNAT pipelines see more than 10 times improvement on ingest performance on FlashBlade. In addition, researchers do not need to re-engineer the entire solution every time they change file types. Instead, they can seamlessly move between and study different imaging modalities, and focus on research and clinical outcomes.

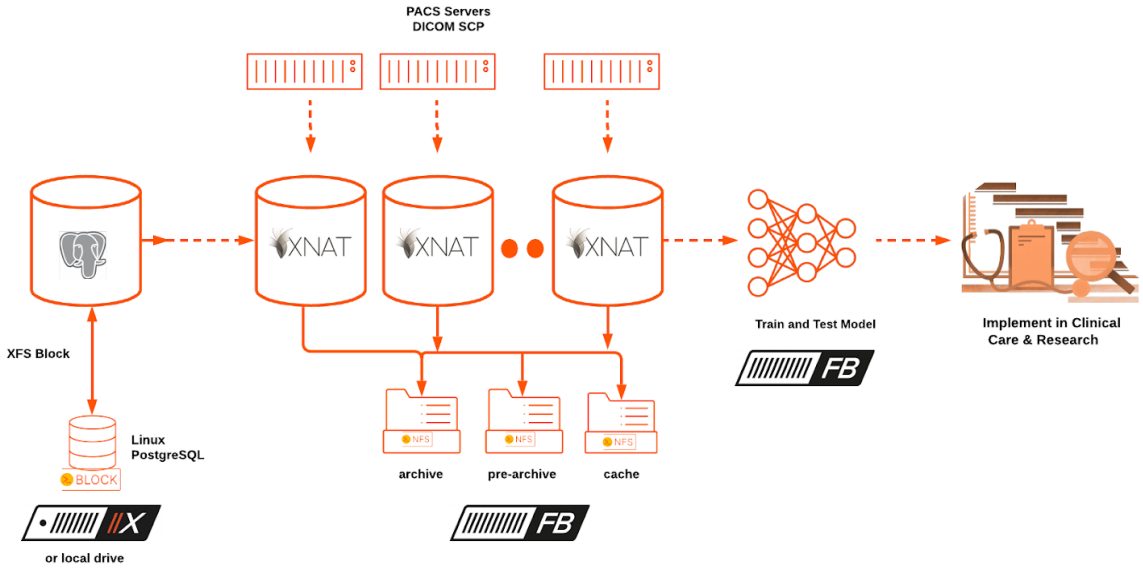

Figure 1. Pure FlashBlade solves the IO bottlenecks in XNAT pipelines.

Because FlashBlade scales well and allows horizontal scaling, you can use multiple XNAT servers to improve pipeline throughput. Aside from ingestion, FlashBlade can also accelerate the training and testing part of the image computation workflow. In addition, teams can use FlashBlade to power more than just XNAT; keeping PACS and VNA data on FlashBlade can ultimately speed up all downstream research processes.

FlashBlade includes many features to help optimize imaging analytics, including:

- Self-tuning and optimizing for small files, large files, and metadata
- Support for NFS, SMB, and S3 protocols
- Cloud-enabled
- Performance that scales linearly
- Simple deployment and management

Pure, NVIDIA, and Radiologics: A Turnkey Solution for Imaging Analytics

Power your medical imaging AI workflows with AIRI®, a jointly architected solution from Pure Storage and NVIDIA that extends the power of NVIDIA DGX A100 systems with the massively parallel, all-flash performance of Pure FlashBlade. AIRI linearly scales training and inference performance with GPU count for various models, including ResNet-50, while FlashBlade can minimize data bottlenecks in the flow. Using the Radiologics software suite on top of AIRI can be ideal for turnkey imaging informatics pipeline deployment.

Modern Data Protection

As imaging data becomes increasingly massive, there's a huge need to protect patient information, research findings, and data backups in the event of an attack. The Pure1 Meta® analytics platform helps data teams synthesize intelligence from thousands of devices and quickly identify attempted intrusions. Combining FlashBlade with Splunk or Elasticsearch creates a powerful data analytics and security platform. Pure SafeMode™ Snapshots are designed to keep an attacker or rogue insider from deleting backups, even if administrator credentials have been compromised, and helps protect data if an attack occurs. At the same time, FlashBlade delivers rapid restore capabilities of up to 270TB per hour.

Additional Resources

- Learn more about enterprise imaging solutions from Pure.
- Explore Pure solutions for life sciences and healthcare.

[1] IDC Health Insights, 2017.

PS2082-01 06/2021