Cloud Adoption Strategy for the Planetary Data System Geosciences Node. L. E. Arvidson, D.V. Politte, D. Scholes, McDonnell Center for the Space Sciences, Department of Earth and Planetary Sciences, Washington University in Saint Louis, 1 Brookings Drive, Campus Box 1169, Saint Louis, Mo, 63130, lars@wunder.wustl.edu; politte@wunder.wustl.edu; scholes@wunder.wustl.edu

Introduction: The Geosciences Node (https://pds-geosciences.wustl.edu) of NASA's Planetary Data System (PDS) stores and distributes ~250 terabytes of archive data. In addition to storing PDS archive files, hundreds of terabytes of disk space are used for archive preparation tasks, databases, virtual server images, etc. Managing such a large amount of data requires an advanced IT infrastructure and the node is always looking for ways to improve the efficiency and functionality of the system. By migrating subsets of our systems to the cloud, we have eased the burden of managing time consuming IT-related tasks, while positioning the node to begin investigating more advanced cloud processing techniques.

Current Cloud Use: The PDS Geosciences Node identified several IT operations to migrate to the cloud. The decision to migrate onsite systems are made based on their required effort to manage, cost to operate, and the possibility of improved functionality by moving to the cloud.

"Cold" Deep Archive Backup to Offline Cloud. For disaster recovery purposes, and to adhere to PDS data integrity policy, a tertiary offline copy of the PDS archive and supporting system files are maintained offsite.

Historically, the disaster recovery backup was copied to tape, physically transported, and stored at an offsite facility. This process was time consuming and expensive. A more efficient process was developed in 2019 using AWS Glacier Deep Archive [2], which is intended to be a long-term offline cloud storage platform and serves a similar function to tape backups. The disaster recovery backups were migrated from magnetic tape to Glacier Deep Archive to address a number of IT challenges including effort required to manage backup jobs and perform hardware maintenance, high hardware support costs, and the requirement of an offsite storage location.

Existing data protection software was used to redirect backups from the tape library to the cloud. The scheduled backup jobs scan and index the local storage cluster and writes incremental changes to Glacier. The node is currently backing up over 400 terabytes of PDS archive data, databases, virtual server images, etc. Once requested, data stored in Glacier are made available to download in less than 12 hours. The recovery time using this processes is comparable to tape backups.

Close to a 50% cost saving was observed by migrating disaster recovery backups to the cloud. Along with the low cost per terabyte for the Glacier tier, IT administrators are no longer required to manage tapes and transport them to remote facilities.

"Warm" Backup to Online Cloud. In an effort to reduce costs, the PDS Geosciences Node migrated its warm (secondary) copy of PDS archive data to Azure's Blob tier of cloud storage instead of replacing onsite storage systems that were reaching the end of service life.

The secondary copy is an online mirror of the primary on-premises storage system. Files can be accessed at the secondary location in the case of the loss or corruption of files on the primary copy.

Azure Blob cloud storage [3] was selected in 2020 for this project due to contracts between Washington University and Microsoft that provide competitive rates for cloud offerings and egress fees. Blob storage maintains files in an online, accessible state and can be accessed through various Microsoft tools and APIs. Files are stored as objects rather than through a common directory structure. Metadata values, including file checksums, are stored with the file objects.

The PDS Geosciences Node uses a combination of an in-house file catalog and the Microsoft Azcopy tool to maintain synchronization of the secondary cloud archive copy with the local primary archive copy. The secondary copy is updated after the PDS releases new data. The high bandwidth between on-premises systems and Azure is sufficient to adequately keep both copies of archive data synced within a few hours.

Storing the secondary copy of PDS archive data in the Azure cloud not only saves on hardware and maintenance costs, but also positions the node to begin developing more advanced methods to process and analyze data in the cloud.

Pilot of CRISM Processing in the Cloud. The PDS Geosciences Node is also exploring in-cloud data analysis using the secondary warm archive backup. This is meant to allow users of the data to take advantage of cost benefits by performing data analysis in virtual machines co-located in the same zone as the cloud data. The lack of cloud egress when a user accesses data from such a VM reduces costs.

This workflow also has two main benefits for users. First, the user does not need to purchase and maintain on-premises hardware for high-performance workloads;

they might stand up a virtual machine that meets their computing requirements, run their programs, and then decommission the VM. This method minimizes costs for the user. Second, the user does not need to download all the data to be analyzed to their on-premises computer over the internet. The transfer from our secondary backup to their cloud VM uses the much faster intra-zone network, and only the final results need to be transmitted over the open internet back to their own computer (see Figure 1). The cloud analysis workflow saves data transfer time in the typical case, where the input data is large and the analysis results are small.

Our pilot project in the cloud involves the analysis of CRISM spectrometer data using JCAT (Java CRISM Analysis Tool) [1]. We created a Windows Server 2019 VM co-located with our secondary backup in Azure, connected to it via Remote Desktop Protocol (RDP), installed JCAT, and copied input data files from our cloud copy onto the VM using the Python API for Microsoft Azure Blob Storage. JCAT was used to analyze the data and the small result files were sent back to the on-premises computer via the RDP client.

This procedure worked correctly, though we note the need to reduce the inconvenience of several steps to make them accessible to data users. These include the need to install all non-OS software, including basic enabling software like the Java Runtime Environment, onto each user-facing VM. Methods will also be developed to provide to users direct download links for archive data stored on the secondary cloud copy.

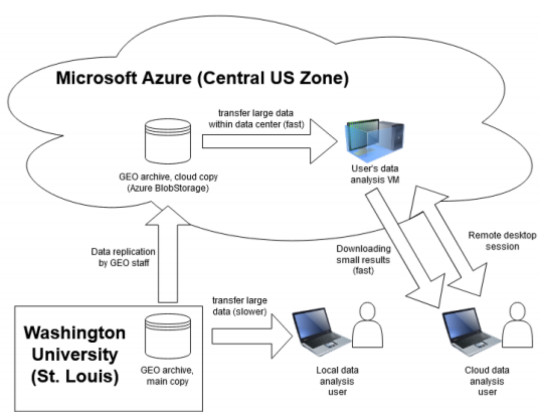

Figure 1. Diagram showing differences between local and cloud-based data processing of PDS Geosciences Node data.

Lessons Learned: Data replication to the cloud proved to be challenging because the on-premises and cloud storage systems use different metadata systems to track file modified and creation times. Traditional replication tools could not correctly mirror the two storage systems. The solution was to use Azure's Azcopy CLI sync tool in conjunction with in-house developed data validation software to detect differences.

Accessing files stored in the Azure Blob tier of cloud storage is not possible using traditional storage protocols such as SMB, NFS, or iSCSI. Custom APIs were written using the Azure SDK. There are options to use file servers in the cloud that use traditional storage protocols, but they were cost prohibitive.

There is a deep learning curve to understand cloud security, authentication, networking, etc. Sufficient time should be dedicated to researching the cloud to determine if your workload is a good candidate for migration.

Understanding how data egress is charged and how to mitigate these costs is important. One way to address egress fees is to require users to access your cloud data from within the same zone. There are also more advanced methods to reduce egress that governs the number of bytes that each user can transfer out of your zome.

Next Steps: Additional cloud implementations to add value to the PDS Geosciences' archive data are being considered.

In an effort to improve data availability for users, the node is considering methods in which PDS archives could be served directly from the cloud in the event of catastrophic or network failure to the primary on-premises copy.

Developing new methods in which users can use cloud computing techniques will be the primary focus for future cloud implementations. In addition to building on cloud pilots that have already been completed, the node plans to investigate cloud computing methods in which users are able to mount PDS archive data to their own compute instances or even "check out" a pre-built VM with data analytic tools already installed.

Acknowledgments: The Geosciences Node IT systems are developed and operated through funding provided by NASA.

References: [1] Politte D. V. et al. (2021) LPSC LII, Abstract #2396. [2] AWS Glacier Deep Archive https://aws.amazon.com/blogs/aws/new-amazon-s3-storage-class-glacier-deep-archive/. [3] Azure Blob https://azure.microsoft.com/en-us/services/storage/blobs/.