



Technical Report

StorageGRID Webscale and Interica PARS

Oil and Gas Project Archiving to Object Storage

Chuck Smith and Steven Pruchniewski, NetApp
Chris Bearce and Kent Halfast, Interica

February 2017 | TR-4575

In partnership with:



Abstract

The Interica Project Archive and Retrieval System (PARS) archives and restores structured, semistructured, and unstructured data from file systems, databases, and applications. Customers who implement a private cloud solution can use NetApp® StorageGRID® Webscale to provide a resilient cloud storage target for Interica PARS. This combination gives customers the freedom to manage their data securely with policies to control locality and redundancy.

TABLE OF CONTENTS

1	Introduction	3
2	Solution Overview	3
2.1	PARS Overview	3
2.2	StorageGRID Overview	3
3	Solution Architecture	4
4	PARS Configuration and Best Practices	5
4.1	Users	5
4.2	Catalogs	5
4.3	Archive Sets	5
4.4	Destinations	6
4.5	Users and Groups	7
4.6	Schedules	7
4.7	PRM and PARS Application Connectors	8
4.8	PRM Project Discovery and Statistics Scans	9
5	StorageGRID Configuration	10
5.1	Configure NetApp StorageGRID for HTTPS	10
5.2	Configure StorageGRID for PARS	10
5.3	Data-Protection Options	11
5.4	ILM Configuration	11
5.5	Archive to External Amazon S3 Destination (Amazon Web Services)	12
6	Sizing Considerations	12
6.1	PARS Sizing	12
6.2	StorageGRID Sizing and Performance	13
7	For More Information	14

LIST OF TABLES

Table 1)	Metadata to add manually	9
----------	--------------------------	---

LIST OF FIGURES

Figure 1)	Configuration of PARS and StorageGRID Webscale tested in-house	4
-----------	--	---

1 Introduction

Oil and gas exploration companies are under the pressures of both time and cost. On one hand, the volume of project data that is generated by an increasingly diverse set of applications continues to grow while IT budgets continue to shrink. On the other hand, the desire to access even older projects nearly instantly is becoming a critical business concern, because reworking older data is more cost-effective than acquiring new data.

Multiple applications and client devices are proliferating petabytes of data across the globe. Geoscientists at each location increasingly demand 24/7 access from every device, and still IT must ensure the integrity and security of data.

As projects age, become inactive, or are put on hold, organizations must move data to lower-cost storage while protecting the results. And it is not just data that must be archived; metadata must be archived as well. Metadata includes information such as who worked the project, when it was last modified, the associated application version, where the project data is on a map, and so on.

Second- or third-tier storage for project archives must be more cost-effective than the production storage tier. It also requires metadata management, data searchability, retrievability, integrity, security, data movement, replication, and data migration over long periods. Compliance requirements mandate that the data be stored for years beyond the lifetime of the current physical storage media. To keep storage management operating costs down, data migration must become automatic and policy driven.

Most energy companies have used tape solutions for storing this dormant data, but retrieval of a sudden “hot again” project can be troublesome and lengthy and can often rely on “tribal” knowledge. And although tape storage is cost-effective, managing a tape infrastructure is not.

Organizations are trying to use cloud service providers to meet this archival storage need. They hope to leverage the cloud’s promise of durability, availability, performance, and cost-effectiveness, but questions concerning access costs, data governance, and vendor lock-in are emerging as use of the cloud evolves.

Interica and NetApp have partnered to develop and certify a solution. Their joint solution meets the requirements for the data from these oil and gas projects to be archived to (and retrieved from) a global, enterprise-wide object storage repository.

2 Solution Overview

2.1 PARS Overview

The Project Archive and Retrieval System (PARS) archives and restores structured, semistructured, and unstructured data from file systems, databases, and applications.

Specifically constructed to manage large datasets, the system has been further tailored by Interica to accommodate the particular needs of the oil and gas industry during its 20+ years of use in this vertical. PARS supports archiving complete projects from well-known oil and gas applications such as Schlumberger Petrel, Landmark DecisionSpace, IHS Markit Kingdom, Paradigm Epos, and many others.

With this system, users can group these relevant data (and customizable metadata) from any source. They can also specify where PARS is to archive information, verify archived data against the original, search for archives based on the metadata, and more. A policy framework plus an Oracle Database allow the entry and extraction of metadata. The UI provides point-and-click access to management features, including archive, restore, delete, and replication.

2.2 StorageGRID Overview

NetApp StorageGRID Webscale is an enterprise-grade object store that supports widely adopted protocols such as Amazon Simple Storage Service (Amazon S3) and Swift. StorageGRID abstracts

storage resources across multiple logical and physical data centers. All resources are presented in a common namespace, enabling the service to continue uninterrupted through hardware and software upgrades and even through entire site outages.

NetApp StorageGRID Webscale can be deployed with various infrastructures. StorageGRID is a software-defined object-based storage platform that provides intelligent policy-driven data management. The ability to manage data while optimizing durability, protection, performance, and even physical placement across multiple geographies is key to meeting business requirements while reducing costs. Deploying NetApp StorageGRID Webscale with NetApp E-Series storage creates an enterprise-grade turnkey object storage appliance that is easy to deploy. Customers have the option to deploy Webscale nodes in any combination of virtual machines (VMs) on VMware or on OpenStack or as physical appliances.

StorageGRID Webscale:

- Delivers efficient, nonstop global operations on the premises and in the hybrid cloud
- Scales across multiple sites under a single global namespace
- Uses leading file and object-storage protocols for compatibility with your enterprise applications
- Lets you implement policy-driven, automated storage management
- Supports multitenancy and identity federation
- Provides a flexible deployment model, supporting both appliance and virtualized environments (software defined)
- Can reduce data storage costs compared with file-based storage systems

3 Solution Architecture

Figure 1) Configuration of PARS and StorageGRID Webscale tested in-house.

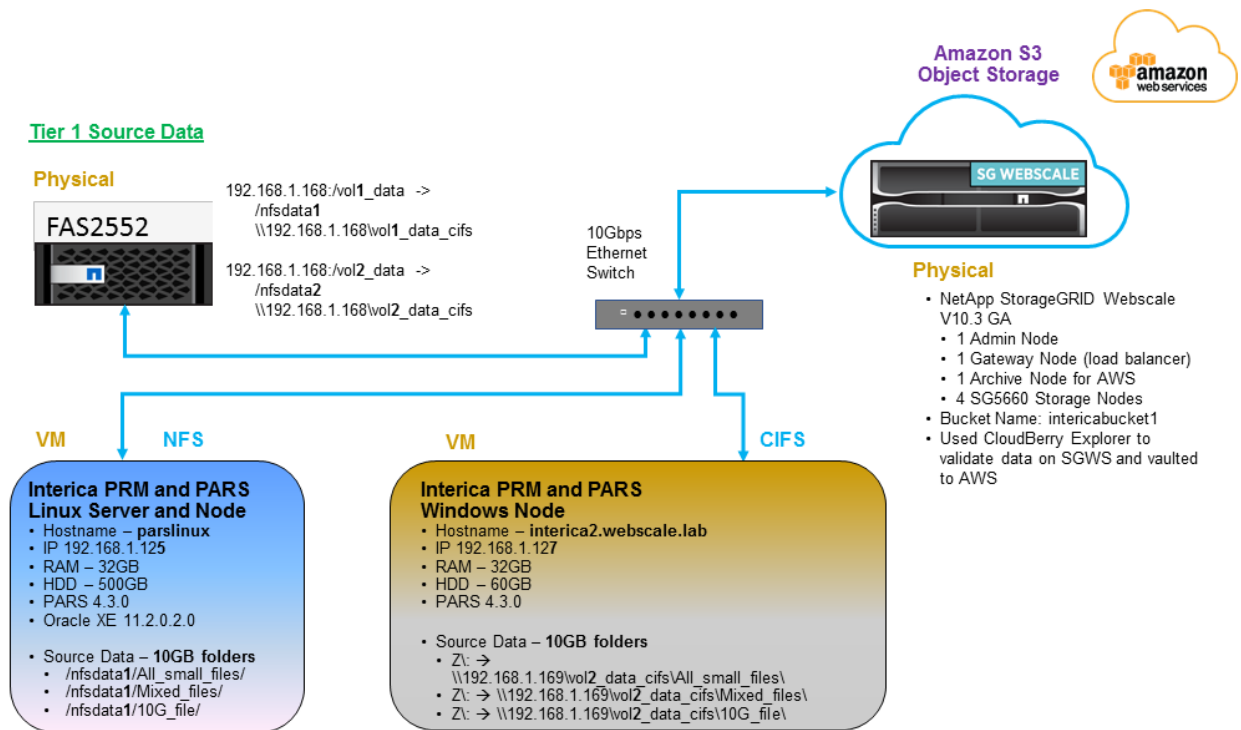


Figure 1 shows a configuration that was tested in-house to certify PARS with NetApp StorageGRID Webscale. These components are the minimum required for a successful installation, except for the four storage nodes that are used. The minimum number of storage nodes required for a StorageGRID installation is three.

The source data was on a FAS2552 storage controller over 10GbE network connectivity. The input data represented what a typical oil and gas interpretation project would contain: seismic, well, interpretation data, and other project files. Through the PARS application, the project data was discovered and archived from the NetApp FAS2552 storage over both the NFS and CIFS/SMB protocols. These protocols are the most common for accessing projects that are used by applications in the upstream geo-computing environment.

Depending on the customer's requirements, this configuration can be expanded to multiple geographic sites. The StorageGRID configuration can be extended, with the global namespace, to up to 16 sites. Although each site must have a minimum of three storage nodes, sites do not need to be symmetrical in number or capacity of storage nodes per site.

Note: StorageGRID Webscale can scale to up to 16 sites (data centers), 100 billion objects, and 70PB in size.

4 PARS Configuration and Best Practices

There are subtle variations in setting up an Interica Project Resource Manager (PRM) and PARS system so that it is simple, secure, efficient, and easy to work with. The configuration shown here gives one such example. Various system architectures and unchangeable legacy data management processes might require changes to this configuration. For those scenarios, a consultant is generally recommended. Against a purist philosophy, the list does indeed mix policies and practices, but it is intended to be a short guide and yet have meaningful information that can be readily put to use.

4.1 Users

To set up users:

1. For parity, create at least two users for each role or role combination.
2. Give the most senior data managers a role combination of server administrator and module administrator so that they can fully manage the PRM and PARS system. A PARS audit log tracks actions by users.
3. It is not recommended to create more than three users with server administrator and module administrator privileges unless the users communicate well and frequently about changes to PARS and the surrounding geotechnical environment.

4.2 Catalogs

Create catalogs based on business units. A catalog is the highest categorization level of archives.

So that test archives are not mixed in with production archives, always have a `TEST` catalog with archive sets that have a `TEST*` archive destination.

4.3 Archive Sets

Create archive sets based on geographic areas. Some companies align their application project area names (for example, R5000 districts) with their PARS archive sets. This approach is the best practice because everything is logical, well thought out, and aligned for simplicity. Note that archive sets are associated with an archive destination. An archive set is the second-highest categorization level of archives.

4.4 Destinations

Create one or two production destinations with meaningful names that represent where the archive will exist. Examples include `ARCHIVE_ %Bucket_Name%` for archives that go to an Amazon S3 archive destination, `ARCHIVE_DISK` for archives that go to a disk file system, and `ARCHIVE_LTO5` for archives that go to an Interica IDS storage set called `ARCHIVE_LTO5`. Create a `TEST_ %Bucket_Name%`, `TEST_TAPE`, and a `TEST_DISK` destination for testing archives. Fewer destinations mean fewer items to manage. PARS and IDS (for tape) handle where the archives exist, and it is easy to find out exactly where they are—down to the object storage bucket, file system, or tape media if necessary.

Cloud Object Storage Destinations

A little more information is needed for this type of destination:

1. **Create an Amazon S3 endpoint first.** Following is an example of the PARS endpoint configuration dialog box. The information that you need from your object storage provider is in the first five fields and in the last checkbox in the dialog box.

The screenshot shows a dialog box titled "Add S3 Endpoints" with a close button (X) in the top right corner. The dialog contains the following fields and controls:

- Name:
- Access Key:
- Secret Key:
- Endpoint URL:
- Protocol:
- Segment Size(MB):
- Write Part Size(MB):
- Write Part Concurrency:
- Read Part Size(MB):
- Read Part Concurrency:
- Use Path Style Access:

At the bottom of the dialog are four buttons: "Test S3 Endpoint", "OK", "Cancel", and a help icon (?).

2. Set the Segment Size to 1GB or 1024MB. Testing was performed with 2GB segment sizes, also, and no appreciable performance gains were shown. Segment Size is used to segment the archives where a particular Amazon S3 destination does not support multipart reads. It therefore allows PARS to support the restoration of selected smaller segments of the archive for a subset or partial restore. Segmentation also confirms that no archive in the cloud exceeds the 5TB limit that is currently in effect.
3. Set the Write Part Size and the Read Part Size to 5MB each. This setting gave the best archive and restore performance during PARS certification testing against NetApp StorageGRID Webscale cloud object storage.

4. Set the Write Part Concurrency and the Read Part Concurrency (number of I/O threads) to 30 each. This setting gave the best archive and restore performance during PARS certification testing against NetApp StorageGRID Webscale cloud object storage.
5. Select the Use Path Style Access box if you want to access Amazon S3 buckets in the following way:
<http://endpointurl/bucketname>

If this field is not selected, Amazon S3 buckets are accessed as follows:

<http://bucketname.endpointurl>

In most cases, Use Path Style Access should be selected. However, sometimes (if you're accessing an endpoint by way of a proxy server, for example) it might be possible to access a bucket only by using "nonpath style."

4.5 Users and Groups

To set up users and groups:

1. Try to group into fewer than 20 the areas that you work and create the archive sets with the same names as those areas. Start small and grow the areas and archive only after careful thought.
Capture more detailed area information in the actual archive metadata. A good list to start with is:
 - North America
 - Gulf of Mexico (optional)
 - South America
 - Alaska (optional)
 - North Sea (optional)
 - Europe
 - Middle East
 - Asia
 - Africa
 - India
2. If only the PARS module administrators will be archiving, then disregard the following setup of other users and groups (items 3 through 7).
3. Create PARS groups based on the areas being worked.
4. Create PARS users and assign them to the PARS groups that match the areas in which they work.
5. Give PARS groups view access to the catalog that is associated with the business unit within which they work.
6. Give PARS groups view-only access to the archive sets that are associated with the area within which they work.
7. Give PARS groups view-only access to the application connectors if you want them to build archive jobs and submit them to the pending queue. That way, the archive jobs can later be run by a senior data manager after the manager has inspected them. This approach helps confirm that all metadata has been input correctly and that the number of jobs that run on the system at any one time is balanced.

4.6 Schedules

Create predefined schedules that allow the project list and archive jobs to be performed when the network and NAS systems are least active. Work with IT to understand the IT backup windows and when the system is most active. Following is a good set of example schedules:

- Monday–Friday at 6 p.m.

- Monday–Friday at 9 p.m.
- Monday–Friday at midnight
- Monday–Friday at 3 a.m.
- Saturday at midnight
- Saturday at 3 a.m.
- Saturday at 6 a.m.

4.7 PRM and PARS Application Connectors

To set up application connectors:

1. Under the General tab, name the application connector the same as the application under which projects will be archived.
2. The Source Nodes tab is for file browse–type application connectors. Add only a list of the file systems in which the data that you want is expected to be found. Do not put just /, a.k.a. the root file system, as a path here. If you do, you might get archives with data that you don't want because **all** file systems live under /, and any of them could be selected.
3. Under the Parameters tab:
 - a. Filename Format—For content clarity when searching log files, a best practice is to use the macro `%op_%appname_%project_%date_%time` to name the archive files.
 - b. Job Log Directory—Always put the application connector name here. Log files are divided by this directory (grouped by application). This approach avoids one directory having all the log files and makes it easier to find them.
 - c. Archive Verify: User Selectable—Select On. (Verifying that the archive is good by reading it back from the destination media and by performing a bit-by-bit comparison with the source data files is a best practice in most cases.)
 - d. Follow File Links—Select Always On. (If file links are encountered, follow them and archive the linked file.)
 - e. Follow Directory Links—Select Always On. (If directory links are encountered, follow them and archive the data in the linked directory.)
 - f. Delete After Archive—Select Always Off. (You can use this option only for simple application connectors, and it is a best practice to have the user or data manager manually delete the source project data after archiving.)
 - g. Restore Verify: User Selectable—Select Off. (Generally, you don't need to verify that the data restored was actually the same data on the archive media.)
 - h. Overwrite—Select Always Off. (This approach prevents accidental overwriting of data during a restore.)
 - i. Logging Options—Select Show Full Path. (This selection provides a good balance between presenting file name only and a detailed list of all files in the log.)
4. For Metadata, a check across all application connectors for consistency is recommended. Defining consistent and meaningful metadata across all applications that provide a similar function, for example, seismic interpretation, is recommended. This approach adds value to the archive asset for the company, and it helps the future restorer understand what is in the archive even if tribal knowledge about the archived data is long gone.

However, the metadata added to each application connector should not result in a list that is so onerous to add that it will never be manually entered. When possible, make PARS pick lists of the metadata. Table 1 is a reasonable list of manual metadata to add to each application connector. Note that PARS has much more metadata that is added to the archive automatically by default.

Table 1) Metadata to add manually.

Metadata Name	Mode	Help Comment
Project Type	User Selectable/Optional	Standalone, Master, Shared, Subproject, Interpretation Project, Survey, 2D, 3D, etc.
Requestor	User Selectable/Mandatory	Person or group requesting that the archive be performed (archive initiator).
Archive Type	User Selectable/Mandatory	Why the archive is being performed; e.g., milestone, inactive project, etc.
Country	User Selectable/Mandatory	Country with the most jurisdiction over the project extent.
Area	User Selectable/Optional	Main geographical or business areas of the project.
Field	User Selectable/Optional	Fields covered by the project.
Basin	User Selectable/Optional	Basins covered by the project.
Business Dept/Unit	User Selectable/Optional	Business department and or unit.
Application Version	Automatic/Manual	The version of the application to which the project belongs. Make a manual pick list if not automatic for the application connector.
Project Extent*	Automatic/Manual	Minimum and maximum latitude and longitude bounding the project. Can also be a single point or polygon set of latitudes and longitudes.
Display Projection	Automatic/Manual	Display or project CRS (Cartographic Reference System) + datum; e.g., map projection such as UTM15, state plane, etc.
Storage Projection	Automatic/Manual	Master CRS (Cartographic Reference System) + datum; e.g., original or stored master projection such as UTM15, state plane, etc.; loaded into master project.
Datum	Automatic/Manual	Geodetic datum if not included with preceding projections; e.g., World Geodetic System (WGS) 1984.
Attached File**	Automatic/Manual	HTML links to one or more files (reports, maps, images, etc.) associated with application project. Files to be placed in default path/name to be picked up during archiving.

* Might not always be automatic; might have to be drawn or pasted into decimal degree fields where not supported yet.

** Might not always be automatic; if not automatic, then place attached files in an `Attached_Files` folder in a top-level project directory.

- a. Catalogs—Check the catalog box to have an application connector show up in the appropriate catalog during the archive process.
- b. User Access—Assign privileges to users and groups based on previously mentioned steps.

4.8 PRM Project Discovery and Statistics Scans

To set up these project scans:

1. Create project discovery scans (generates the project lists) with multiple applications and file systems selected so that they run more efficiently. PARS uses multithreaded scanning technology that excludes directories in which projects have been found. This feature results in faster scan times than if separate scans were set up for each application, covering overlapping file systems.

2. Name the project discovery scans following this example: %Application1 %Application2 %Application3... Project List.
3. You do not have to create project discovery scans for applications that have a system such as a central database from which a project list can be queried. For example, don't create project discovery scans for EPOS4, R5000, or GeoFrame projects.
4. Schedule project list scans to run nightly so that the project list for PARS archives and PRM statistics scans is never more than 24 hours old.
5. Create separate PRM project statistics scans per application for smaller, easier-to-read log files that are focused on the one application. These scans generally run longer than the project discovery scans because they are gathering the metadata.
6. Name the PRM project statistics scan following this example: %Application Statistics.
7. Run the first PRM project statistics scan manually after the system has been installed and validated for production. To allow time for analysis and action, schedule subsequent PRM project statistics scans to run monthly.
8. If the PRM project statistics scans are running longer than anticipated, set the Max Duration hours to a suitable window, such as 8 hours. The scan then runs only when scheduled and for the Max Duration hours that you set. The next scan starts where the previous scan left off. Multiple scans result in obtaining statistics for all the projects that are discovered in the discovery scans.

5 StorageGRID Configuration

5.1 Configure NetApp StorageGRID for HTTPS

A best practice is to configure StorageGRID to use a Secure Sockets Layer (SSL) certificate that is provided by a certificate authority (a commercial SSL or one provided by a customer's IT).

You must configure StorageGRID to use an SSL key. For more information, see "Configuring a Custom Server Certificate" in the "StorageGRID Webscale 10.3 Simple Storage Service Implementation Guide."

By default, the StorageGRID gateway node (load balancer) uses port 8082 for HTTPS communication. You can configure a redirect on the gateway node to use port 443. You can also use third-party or open-source load balancers if you prefer.

5.2 Configure StorageGRID for PARS

After StorageGRID has been configured, an endpoint (Amazon S3) can be configured for ingest from the PARS application:

- The StorageGRID administrator creates an Amazon S3 account to be used by PARS and provides the following information:
 - Access Key ID
 - Secret Access Key
 - Endpoint, typically a DNS entry pointing to a StorageGRID gateway node IP and port 8082

Save Keys

⚠ The Access Key and Secret Access Key will not be accessible once this dialog is closed. Please retain a copy of these keys by using the Download option or copying them directly.

Access Key ID:

Secret Access Key:

Download

Finish

- Create an information lifecycle management (ILM) policy to protect and distribute PARS data.

5.3 Data-Protection Options

Before you deploy this solution, consider the customer's high-level IT infrastructure and access patterns:

- To how many sites will the solution be deployed?
- From where will the data be accessed?
- Is a disaster recovery plan in place?
- What level of failure does the customer want to protect against?
 - Data center or regional disaster
 - Single node or server failure

StorageGRID is designed to be a multisite, eventually consistent object store. Customers can choose to implement a single site or up to 16 sites. Sites are logical. They can be geographically separated or within a physical data center, depending on the data-protection level that the customer requires.

Data ingested into the StorageGRID solution is protected at multiple levels. With object storage, the concept of primary and disaster recovery data centers no longer applies. Data is placed as per the policy; data can be stored in a single site or can be replicated across all sites. Every site is active for ingest and retrieval, greatly simplifying business continuity planning. A single namespace is presented to the application, Interica PARS in this example. PARS must know only the endpoint address and bucket.

PARS objects (archive segments) are 1GB in size as long as the segment size is set to 1024MB per the PARS Endpoint Configuration dialog box. This makes them an excellent candidate for erasure coding data protection. To PARS, an archive on an Amazon S3 destination appears as a single file even though it could be split into segments. Therefore, if it is replicated to a destination where segmentation cannot occur, it appears as a single file again.

For details about creating a data-protection policy, see the StorageGRID Webscale Administrator Guide, https://library.netapp.com/ecm/ecm_download_file/ECMLP2411995.

5.4 ILM Configuration

The StorageGRID administrator can choose to configure a specific ILM rule for PARS data or simply use ILM rules that are already in place.

Data placement and protection within the grid are driven by ILM policy. This policy defines the type of replica that is stored and how many and where in the grid the objects are stored. ILM policy is driven by system or user metadata. Customers can choose any quantity and combination to meet data-protection and service-level agreements. Also, any combination of replica types and moves from one type to another over time, based on access or other metadata, can be configured.

Supported replica types are:

- Full copies—one or many
- Erasure coded
- Tape
- Cloud tier (requires an archive node, see section 5.5)

For details about constructing and managing ILM policies, see the StorageGRID Webscale Administrator Guide, https://library.netapp.com/ecm/ecm_download_file/ECMLP2411995.

5.5 Archive to External Amazon S3 Destination (Amazon Web Services)

The archive node, set up through the StorageGRID administrator GUI, provides an interface through which object data is sent to an external archival storage system for long-term storage. You can configure an archive node to target any external archival storage system that can interface with the StorageGRID Webscale system through the Amazon S3 API, such Amazon Web Services (AWS).

Objects that are stored in StorageGRID can be vaulted to an AWS endpoint upon initial ingest of the object or later. In both cases, information lifecycle management (ILM) policies are configured through the StorageGRID Webscale GUI where the archive node is selected as the target for content placement. AWS then stores and manages the object itself; however, the vaulted object's metadata is kept on the storage nodes within StorageGRID. This approach is taken for proper and complete retrieval if an ILM policy dictates it, or for the simple restore of the data from the PARS application, in this case.

If you want to restore an archive project that has been vaulted to AWS, you can simply use the PARS web UI to restore the project as normal. PARS knows that the project has been archived to StorageGRID and requests a `GET` for the data. StorageGRID knows that the project has been vaulted to AWS and issues another `GET` for the data. This process occurs seamlessly, from the PARS user's standpoint.

6 Sizing Considerations

6.1 PARS Sizing

Following are the requirements for hardware and OS on which to install PARS, for both Microsoft Windows and Linux.

Requirements for PARS Windows VM or physical server:

- OS—Windows Server 2012, 64 bit, joined to a Microsoft Windows Active Directory domain with DNS working between all servers and storage in the test architecture
- Two quad-core CPUs
- 32GB RAM
- 500GB disk space for application installation
- 10Gb Ethernet
- PARS 4.3 licenses to be supplied by Interica (**a service account with domain administrator privilege is necessary to install and configure PARS and for PARS service to run**)

Requirements for PARS Linux VM or physical server:

- Linux OS—Red Hat or CentOS 6.x, 64 bit, with DNS working between all servers and storage in the test architecture
- Two quad-core CPUs
- 32GB RAM
- 500GB disk space for application installation
- 10Gb Ethernet
- PARS 4.3 license to be supplied by Interica (**root or sudo access is necessary to install and configure PARS and for PARS service to run**)

Database requirements:

- Oracle 10 or 11 installed locally on any PARS server or on a remote server; client database administrator or Interica to create the schema at installation
- SQL Server 2012 installed locally on the Windows PARS server or on a remote Windows server; client database administrator or Interica to create the schema at installation

6.2 StorageGRID Sizing and Performance

NetApp StorageGRID sizing depends on several factors, including:

- Number of storage nodes in a grid
- Number of gateway nodes in a grid
- Types of storage in use (SSD, NL-SAS, archive tape and cloud, or third-party array)
- Number of objects in the grid
- Workload (read- or write-heavy; small or large objects, and so on)
- Number and type of replicas (erasure coded or full replication copy)
- Data center design (number of sites, number of nodes per site)

PARS typically is considered to be a large object workload, meaning that it includes a few large objects instead of many billions of small objects. For large object workloads, denser storage nodes such as the NetApp SG5660 appliance are recommended. A typical PARS archive of a project can range from 10GB to 100GB+.

Each StorageGRID storage node supports up to 375 million objects—with a minimum of three storage nodes per site, so the object count should not be a gating factor in archiving PARS objects. However, object count should still be considered, especially if StorageGRID will host storage for other applications (other than PARS).

A StorageGRID sizing calculator is available for properly designing options such as erasure coding, the number of sites, the number of storage nodes, replicas, and so on. For assistance, contact your local NetApp sales systems engineer.

Detailed performance characteristics of StorageGRID Webscale are documented in TR-6773: “StorageGRID Webscale Performance”, available upon request.

7 For More Information

About NetApp

Leading organizations worldwide count on NetApp for software, systems and services to manage and store their data. Customers value our teamwork, expertise and passion for helping them succeed now and into the future. www.netapp.com.

About Interica

Interica is the new force in E&P data management, delivering integrity, control, access and efficiency to data owners and users in Oil and Gas, across the globe. We provide a growing range of specialist software and service-based solutions that help you to manage your subsurface data more effectively, to improve decision making and hence generate value for your organization. Interica has been created by bringing together Enigma Data and InfoAsset. Established businesses with proven international success, securing the mission critical data of some of the world's largest organizations. <http://www.interica.com>

Refer to the [Interoperability Matrix Tool \(IMT\)](#) on the NetApp Support site to validate that the exact product and feature versions described in this document are supported for your specific environment. The NetApp IMT defines the product components and versions that can be used to construct configurations that are supported by NetApp. Specific results depend on each customer's installation in accordance with published specifications.

Copyright Information

Copyright © 1994–2017 NetApp, Inc. All rights reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

Trademark Information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.