HEP Data-Intensive Distributed Cloud Computing Use Cases Document

CANARIE NEP-101 Project
University of Victoria HEP Computing Group

December 18, 2013

Version 0.2 Draft

Revision History

Date	Reason for Changes	Version
24/10/2013	Document skeleton.	0.1
18/12/2013	First draft.	0.2

Contents

List of Figures

1 Introduction

This document describes the requirements of the NEP-101 system in terms of use cases. A use cases document allows the reader to quickly identify the functionality that the system will offer to its users. The system is represented as a sequence of simple steps that describe the interaction between one or more actors (see below). In this document, these steps are described both textually and graphically in use case diagrams.

1.1 Definitions, acronyms and abbreviations

APF	The ATLAS AutoPyFactory generates batch jobs in response to workloads within the ATLAS PanDA queue. These APF batch jobs are payloadless, that is they do not contain the application or data to perform ATLAS analysis or simulation, but they are used to secure computing resources in a distributed grid or cloud environment. Having secured the required computing resources, an APF job selects and pulls a unit of work from the PanDA queue for execution.
ATLAS	A HEP experiment conducted at CERN laboratory in Geneva, Switzerland by more than 100 collaborating countries.
Cloud	A collection of computing resources and software that provide on demand through Web Services, Infrastructure As A Service (IaaS) Virtual Machines (see below).
DCCM	Distributed Cloud Computing Model as developed by NEP-52 and enhanced by NEP-101.
HEP	High Energy Physics, sometimes referred to as Particle Physics or Nuclear Physics
Image Repository	A storage capability employed by the NEP-101 system to manage virtual machine images. Depending on its' type, an image repository may be dedicated to a single cloud or shared by multiple clouds.
NEP-52	HEP Legacy Data project (Oct 2009 - Mar 2012), created the DCCM for high throughput, modest data input/output serial processing.
NEP-101	HEP Data-Intensive Distributed Cloud Computing project (Oct 2013 - Dec 2014), will extend the DCCM for data intensive ATLAS applications.
Virtual Machine (VM)	A complete computer system represented in software

2 Actors

An actor is someone or something outside the system that either acts on the system – a primary actor – or is acted on by the system – a secondary actor. An actor may be a person, a device, another system or sub-system, or time. Actors represent the different roles that something outside has in its relationship with the system whose functional requirements are

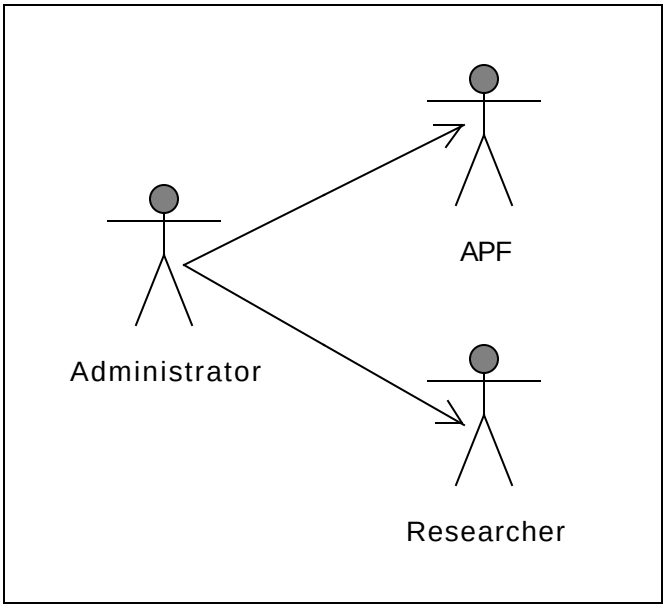

Figure 1: NEP-101 Primary Actors

being specified. An individual in the real world can be represented by several actors if they have several different roles and goals in regards to a system. They interact with system and do some action on that. [1]

2.1 Primary actors

The primary actors of the NEP-101 system and their relationship to each other are shown in Figure 1.

Researcher This actor represents the main user of the NEP-101 system. Researchers have user accounts on the NEP-101 system and will use the system to do some research by processing some ATLAS data files and analyzing the output.

Administrator An Administrator's role is to configure and manage the NEP-101 system. These users are able to do all the actions the Researchers can do, plus some extra administration tasks on the NEP-101 system such as user account management, cloud configuration and software deployment.

VM-Instance This actor represents an instance of a VM image running on a cloud resource. VM-Instances interact with the system in order to fetch input data and store job output results.

[1] http://en.wikipedia.org/wiki/Use_case

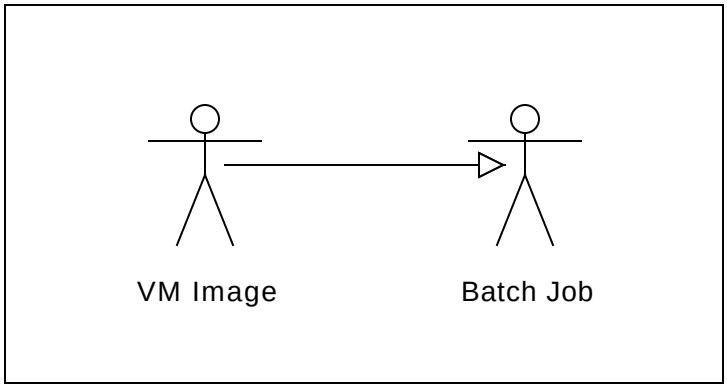

Figure 2: NEP-101 Secondary Actors

2.2 Secondary actors

The secondary actors of the NEP-101 system and their relationship to each other are shown in Figure 2.

VM-Image This actor is a VM image which may be manipulated by administrators and researchers alike. VM images may be created, modified, listed, renamed, copied, instantiated, shared and deleted. VM images may be generated and stored in a variety of different formats and may need conversion from one format to another for use on any particular cloud.

Storage Element This actor represents a data storage facility containing data the Researcher wishes to process using the resources provided by the VM Instance.

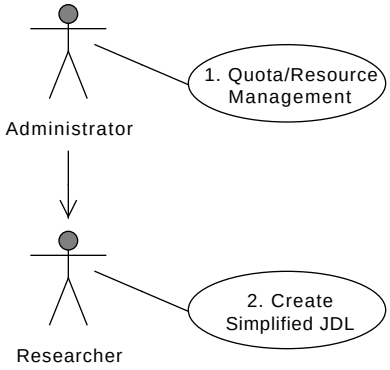

Figure 3: Batch Scheduling Use Cases

3 Batch Scheduling

This section contains detailed description of the Batch Scheduling use cases. See Figure 3 for a UML diagram of these use cases.

3.1 Use Case: Batch Job Creation using Simplified JDL

Description

Employing a text editor, the user creates or modifies Job Description Language (JDL) files; files that describe the environment, executable, and parameters for batch jobs. Previously, various specific parameters were required which were peculiar to a cloud or to a cloud type. New cloud/type agnostic parameters have rendered these specific parameters unnecessary, allowing a unified job description for all clouds.

Actors

- Administrator, APF (ATLAS Auto-Pilot-Factory), Researcher

Preconditions

- Consistent environment "flavors" across the clouds to be used.
- The availability of the required VM image on the clouds to be used.
- The availability of input data on the clouds to be used.

Basic flow

1. User creates a JDL file using a text editor.
2. User specifies the hardware (flavor) and software (image name) environment, the data requirements, the executable and job parameters (executable arguments).
3. User submits the JDL file to the HTCondor job queue.
4. System provides various commands to monitor the job progress and to retrieve the job output.

Trigger

A researcher or the APF needs to run ATLAS production jobs on the cloud.

Postconditions

- The required output is generated, retrieved, and stored.

3.2 Use Care: Quota Management and Simplified Resource Specification

Description

Using the available API query and manage the resource quota based on information available from the service. Previously such values were specified in a configuration file for the cloud resource

Actors

- Administrator

Preconditions

- Information regarding quotas is available through the service API

Basic flow

1. Administrator specifies basic information for a cloud service in the configuration file
2. The program contacts the service and queries for quota information
3. The program uses the quota information to determine scheduling in place the previously configured values

Trigger

The Administrator adds a new cloud resource and tells the program to reload the configuration

Postconditions

- The quota for the cloud resource is now accurately tracked by the service instead of externally

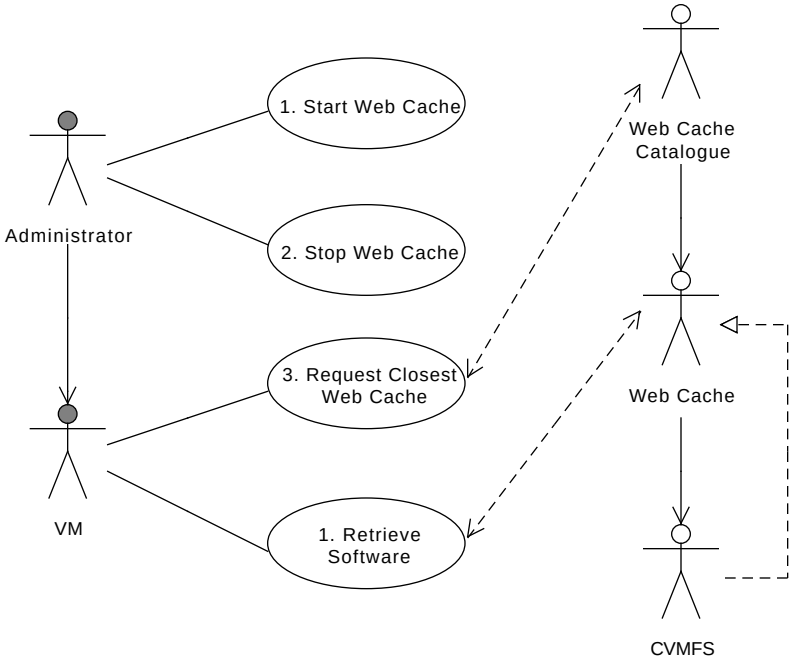

Figure 4: Software Distribution Use Cases

4 Software Distribution

This section contains detailed description of the Software Distribution use cases. See Figure 4 for a UML diagram of these use cases.

4.1 Start a web cache

Description

To optimization network utilization, a DCCM administrator decides that an additional web cache is required.

Actors

- Administrator, web cache VM appliance.

Preconditions

- No suitable (eg.. too remote, too busy, etc) web cache is available.

Basic flow

1. DCCM administrator identifies opportunity and location for network optimization.
2. DCCM administrator uses there cloud credentials to authenticate with the identified hosting cloud.
3. DCCM administrator manually instantiates a web cache VM appliance using an image in the cloud's repository.

Postconditions

- Hosting cloud has a web cache running locally.

4.2 Stop a web cache

Description

To minimize resource utilization, a DCCM administrator decides that a running web cache is no longer required.

Actors

- Administrator, web cache VM appliance.

Preconditions

- A running web cache is consuming resources but has no clients.

Basic flow

1. DCCM administrator identifies the redundant web cache.
2. DCCM administrator uses there cloud credentials to authenticate with the identified hosting cloud.
3. DCCM administrator manually terminates the web cache.

Postconditions

- Hosting cloud is no longer running the redundant web cache.

4.3 Request closest web cache

Description

VM-Instance needs a web cache to effectively access software used by the researcher

Actors

- VM-Instance, CVMFS web cache, cache catalogue

Preconditions

- CVMFS web cache advertises to cache catalogue

Basic flow

1. VM Instance requests nearest CVMFS web cache from cache catalogue
2. Cache catalogue matches VM Instance to CVMFS web cache by geographical location
3. VM Instance configures CVMFS client to use CVMFS web cache

Postconditions

- CVMFS web cache provides software to VM Instance

4.4 Software access using CVMFS

Description

Actors

- Researcher, Batch Job, VM-Instance, CVMFS client, CVMFS server, CVMFS web cache

Preconditions

- Researcher places software on CVMFS server
- VM-Instance runs CVMFS client
- CVMFS client is connected to CVMFS server and CVMFS web proxy

Basic flow

1. batch job access software in directory tree supplied by CVMFS server
2. VM Instance retrieves software on demand from CVMFS server using the CVMFS web cache

Trigger

Batch job requests software from CVMFS

Postconditions

- Software is available on the VM Instance
- Batch job uses software

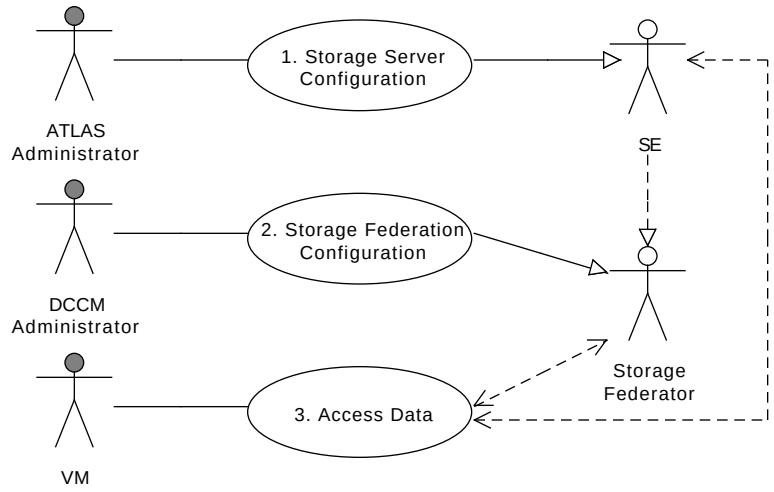

Figure 5: Storage Federation Use Cases

5 Storage Federation

This section contains detailed description of the Storage Federation use cases. See Figure 5 for a UML diagram of these use cases.

5.1 Use Case: Storage Server Configuration

Description

ATLAS batch production processing requires access to data using standard protocols and in specific formats. Servers meeting these requirements are referred to as Storage Elements (SE) and consist of hardware and software together with large amounts (multiple terabytes) of data copied from the ATLAS collaboration. To provide service to the local ATLAS (Tier 2 and 3) processing centres and provide data integrity and redundancy, many SEs containing overlapping and duplicate ATLAS data are already distributed around the globe. These SEs are built by local administrators using local resources.

Actors

- ATLAS administrator.
- ATLAS data plus hardware and software resources.

Preconditions

- ATLAS is available for distribution.
- Sufficient resources are available to build the SE.

Basic flow

1. Administrator configures a server with appropriate hardware and software resources.
2. Administrator copies the required ATLAS data to the new SE.
3. Administrator enables access to authorized users.

Trigger

Researchers require access to ATLAS data.

Postconditions

- Researchers can access ATLAS data locally.
- Optionally, data can be accessed remotely.

5.2 Use Case: Storage Federation Configuration

Description

In DCCM, ATLAS batch processing is distributed around the globe. Each job, which can run anywhere within the DCCM environment, will need to access data from one or more of the many distributed SEs. The storage federator provides a catalogue and single point of access for all jobs, redirecting them to the nearest SE for any particular data element.

Actors

- DCCM administrator.
- Hardware and software resources.

Preconditions

- One or more ATLAS SEs available for federation.
- Sufficient resources are available to build the federation server.

Basic flow

1. DCCM administrator configures a server with appropriate hardware and software resources.
2. DCCM administrator obtains and configures access to all SEs to be federated.
3. DCCM administrator starts the federation service.

Trigger

ATLAS jobs require single point of access to ATLAS data.

Postconditions

- ATLAS jobs running within DCCM can access ATLAS data for the nearest data repository.

5.3 Use Case: Batch Input from Storage Federation

Description

- VM-Instances running on a distributed cloud access data provided by one or more storage elements.
- VM-Instance uses optimal storage elements determined by relative location.
- Performance requirements: ATLAS pile-up jobs - down: 800 Mb/h/job, up: 20 Mb/h/job.

Actors

- Administrator, APF, Researcher

Preconditions

- ATLAS Storage elements running WebDAV.
- A Storage Federation service running on each cloud.

Basic flow

1. User submits a batch job.
2. Batch job contacts local federator service to ascertain the nearest copy of the required input data.
3. The federator supplies either the URI for the required input data or an error response (eq. invalid request, data unavailable, etc.)
4. Storage element permits data access based on credentials supplied by the Researcher?
5. Batch job either processes the data or fails.

Trigger

An ATLAS production job was placed in the batch queue.

Postconditions

- ATLAS production output or failure messages generated.
- Output to a SE via Federation? Does the Federation need to become a SE for Panda?

6 VM Image Distribution

This section contains detailed description of the VM Image Distribution use cases. The uses cases are broken down into the following sub categories: user account management, user repository management, and image management.

6.1 User Account Management Use Cases

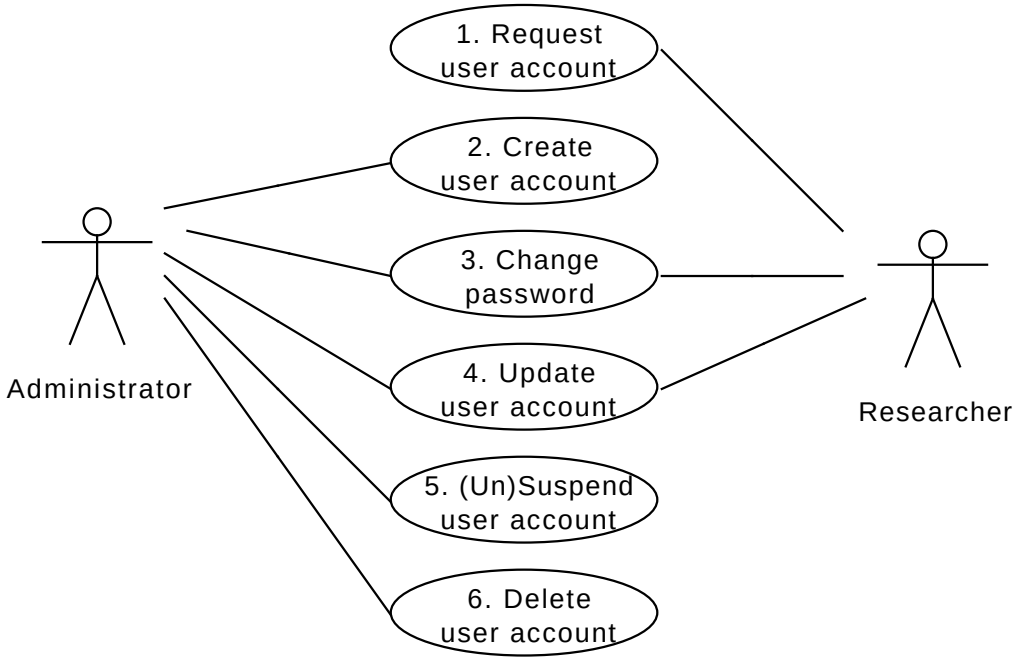

Figure 6: User Account Management Use Cases

6.1.1 Use Case: Request user account

Description

User requests a VM distribution account from the DCCM administrator.

Actors

- Researcher.

Preconditions

- The user is entitled to the services and can provide the necessary identity credentials.

Basic flow

1. User sends an account request to the DCCM administrator, specifying the preferred account name and required resources.
2. Administrator processes the account request (see below).

Trigger

Researcher needs to use the DCCM and propagate VM images to the target clouds.

Postconditions

- Researcher has an account and log in to the VM distribution service.

6.1.2 Use Case: Create user account

Description

The administrator receives a request for a VM Image Distribution account from a new user. The request includes the user's credentials, eg. x.509 distinguished name, ssl public key, ldap, keystone ID, etc. What kind of authentication do we want for the service (x.509, or keystone, ldap, domain defined, other)? The account is created and the user receives notification together with instructions on how to access it.

Actors

- Administrator, Researcher

Preconditions

- The user is entitled to the services and can provide the necessary identity credentials.

Basic flow

1. User requests via email the creation of a personal account and provides the necessary credentials, eg. x.509 distinguished name, ssl public key, etc.
2. Administrator creates the account and sends notification and instructions to the user.
3. User logs into their account and performs any first time configuration, eg. reset password, etc.

Trigger

A Researcher needs to manage VM images across multiple cloud repositories.

Postconditions

- The user is able to use the VM Image Distribution service.

6.1.3 Use Case: Change user password

Description

The user (or an administrator) can change the password on an account.

Actors

- Administrator, Researcher

Preconditions

- The user has access to a valid account.

Basic flow

1. User logs into the account.
2. User enters into the "Change Password" dialog and selects, verifies, and saves a new password.

Trigger

A user changes the password for security reasons or convenience.

Postconditions

- The user is able to use access their account with the new password.

6.1.4 Use Case: Update user account

Description

The user (or an administrator) can change account information.

Actors

- Administrator, Researcher

Preconditions

- The user has access to a valid account.

Basic flow

1. User logs into the account.
2. User enters into the "Update Account" dialog, updates and saves the account information.

Trigger

A user needs to rectify incorrect account information.

Postconditions

- The user account correctly reflects the user status.

6.1.5 Use Case: (Un)Suspend user account

Description

The user's privileges can be temporarily revoked or restored.

Actors

- Administrator

Preconditions

- A valid user account needs to enabled or disabled.

Basic flow

1. Administrator logs in.
2. Administrator enters into the "Account Suspend/Unsuspend" dialog, changes and saves the status.

Trigger

An administrator needs to temporarily revoke or restore a user privileges.

Postconditions

- The user's account privileges are changed.

6.1.6 Use Case: Delete user account

Description
Redundant user accounts can be deleted.

Actors

- Administrator

Preconditions

- A redundant user account needs to be removed.

Basic flow

1. Administrator logs in.
2. Administrator enters into the "Account Delete" dialog and deletes the redundant account..

Trigger
An administrator needs to unnecessary accounts.

Postconditions

- The user's account is removed.

6.2 Repository Management Use Cases

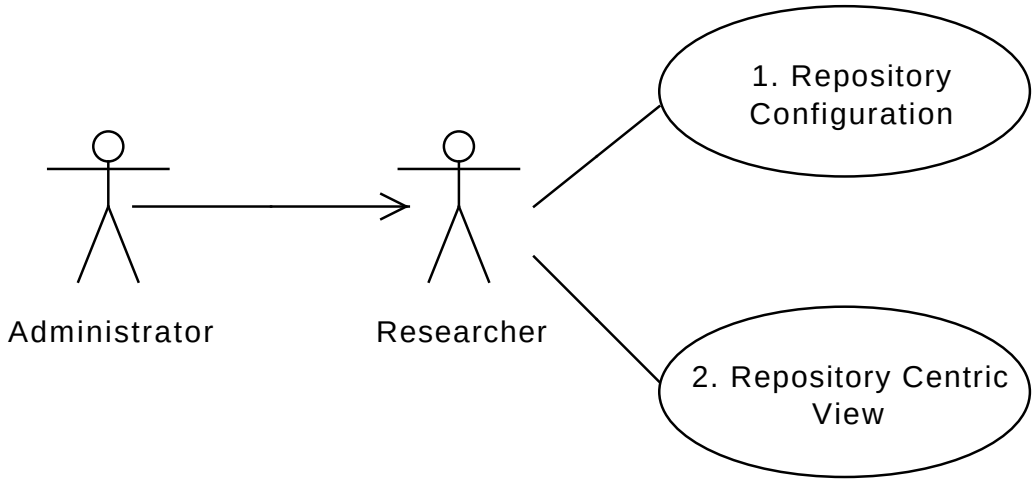

Figure 7: Repository Management Use Cases

6.2.1 Use Case: Repository Configuration (User)

Description

Before a user can manage the images on a repository, the need to identify the repository and their credentials on it. By defining multiple repositories, a user can then replicate images across repositories as dictated by their processing requirements. The system allows for both manual, automatic and scripted replication.

Actors

- Administrator, Researcher

Preconditions

- User has authorization and credentials to login to the VM distribution service and to access one or more clouds/cloud repositories.

Basic flow

1. User logs into the target cloud and retrieves their credentials.
2. User logs into the VM distribution service and enters the Repository Configuration dialog.
3. User completes the dialog for each target cloud, defining the name, location (URL), and credentials.

Trigger

User needs to propagate an image to one or more repositories in order to run VMs on the target clouds.

Postconditions

- User image is propagated to one or more repositories.

6.2.2 Use Case: Display Repository-centric View

Description

Show all repositories used/configured by the user and which images they contain.

Actors

- Administrator, Researcher

Preconditions

- User has authorization and credentials to login to the VM distribution service and to access one or more clouds/cloud repositories.

Basic flow

1. User logs into the VM distribution service.
2. User requests repository centric view of the repositories and images.

Trigger

Process is invoked at the user's discretion.

Postconditions

- User has a list of the repositories and images organized by repositories.

6.3 Image Management Use Cases

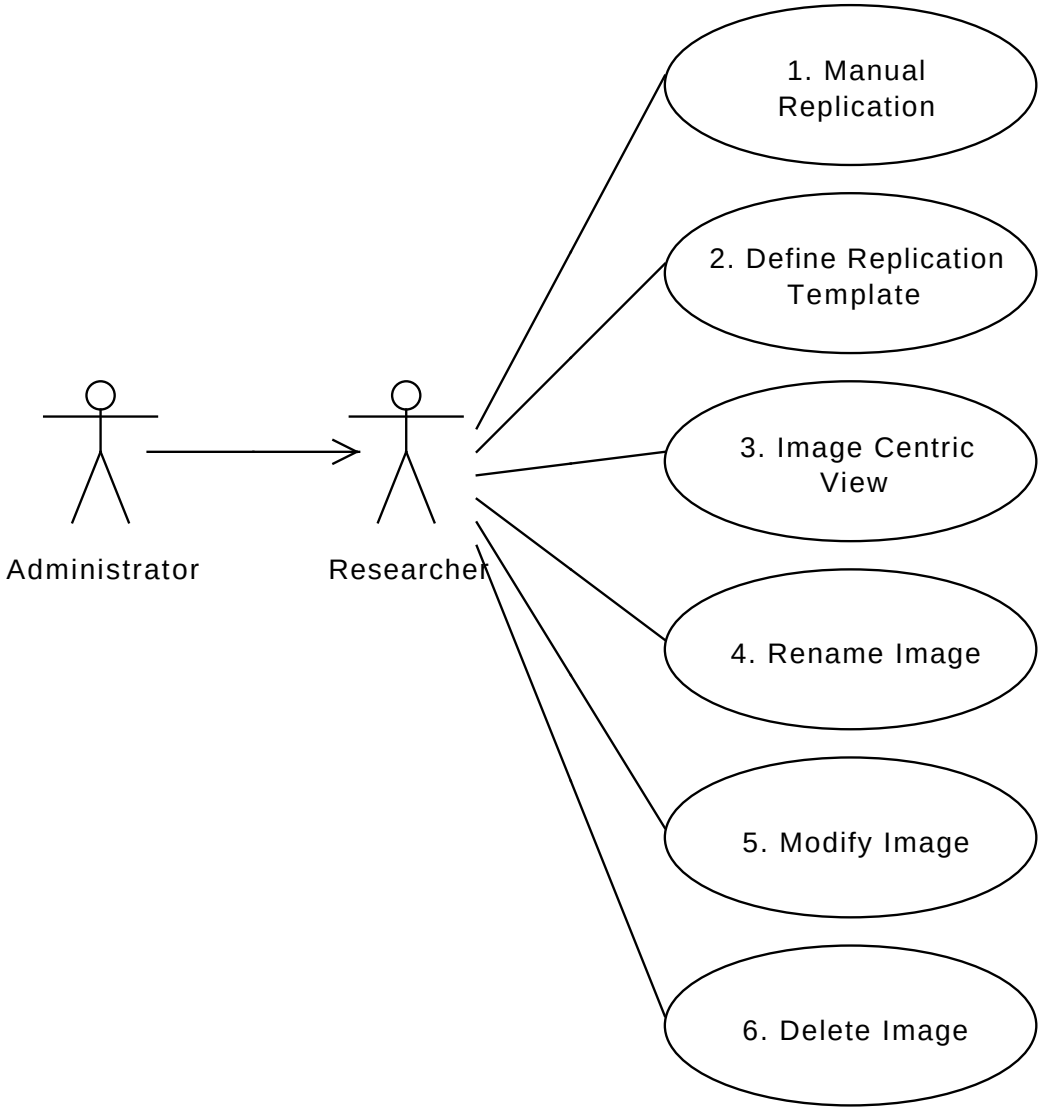

Figure 8: Image Management Use Cases

6.3.1 Use Case: Manual Replication

Description

Show all repositories used/configured by the user and which images they contain.

Actors

- Administrator, Researcher

Preconditions

- User has authorization and credentials to login to the VM distribution service and to access one or more clouds/cloud repositories.

Basic flow

1. User logs into the VM distribution service and enters the Image Propagation dialog.
2. User completes the dialog for each target cloud.

Trigger

Process is invoked at the user's discretion.

Postconditions

- Image is copied as requested to each target cloud.

6.3.2 Use Case: Define Replication Template

Description

To facilitate automatic replication, the user defines templates. A template will identify an image by name and indicate where copies of the image are to reside. It will also indicate under what circumstances copies of the image are to be refreshed.

Actors

- Administrator, Researcher

Preconditions

- User has authorization and credentials to login to the VM distribution service and to access one or more clouds/cloud repositories.

Basic flow

1. User logs into the VM distribution service and enters the Propagation Template dialog.
2. User completes the dialog.

Trigger

Process is invoked at the user's discretion.

Postconditions

- Images which are the subject of a propagation template will be propagated automatically.

6.3.3 Use Case: Display Image-centric View

Description

Show all images available/owned by the user and which repositories they are stored in.

Actors

- Administrator, Researcher

Preconditions

- User has authorization and credentials to login to the VM distribution service and to access one or more clouds/cloud repositories.

Basic flow

1. User logs into the VM distribution service.
2. User requests image centric view of the images and repositories .

Trigger

Process is invoked at the user's discretion.

Postconditions

- User has a list of the images and repositories organized by images.

6.3.4 Use Case: Image Rename

Description

Show all images available/owned by the user and which repositories they are stored in.

Actors

- Administrator, Researcher

Preconditions

- User has authorization and credentials to login to the VM distribution service and to access one or more clouds/cloud repositories.

Basic flow

1. User logs into the VM distribution service and enters the Image Rename dialog.
2. User completes the dialog.

Trigger

Process is invoked at the user's discretion.

Postconditions

- The image is renamed at all locations where it is resident.

6.3.5 Use Case: Image Modify

Description

Show all images available/owned by the user and which repositories they are stored in.

Actors

- Administrator, Researcher

Preconditions

- User has authorization and credentials to login to the VM distribution service and to access one or more clouds/cloud repositories.

Basic flow

1. User modifies an image's master copy.

Trigger

Process is invoked at the user's discretion.

Postconditions

- If managed by a template, the modified image is copied all other target locations automatically.

6.3.6 Use Case: Image Deletion

Description

Show all images available/owned by the user and which repositories they are stored in.

Actors

- Administrator, Researcher

Preconditions

- User has authorization and credentials to login to the VM distribution service and to access one or more clouds/cloud repositories.

Basic flow

1. User logs into the VM distribution service and enters the Image Delete dialog.
2. User completes the dialog.

Trigger

Process is invoked at the user's discretion.

Postconditions

- The image is removed from all target clouds.