SlipStream: Data Model

C. Loomis (CNRS/LAL & SixSq)

9 December 2014

Orsay, France

Data Model

- The SlipStream data model defines how images and deployments fit together and are organized.
- The core entities (“modules”) are:
  - Project: Container of other modules
  - Image: Description of VM image
  - Deployment: Group of VM instances (nodes)
- All modules are versioned and the full history is always kept.

sixsq

Project

- Allows modules to be organized hierarchically.

- All users share a common root and consequently it is recommended to create your own root project for your modules.

sixsq

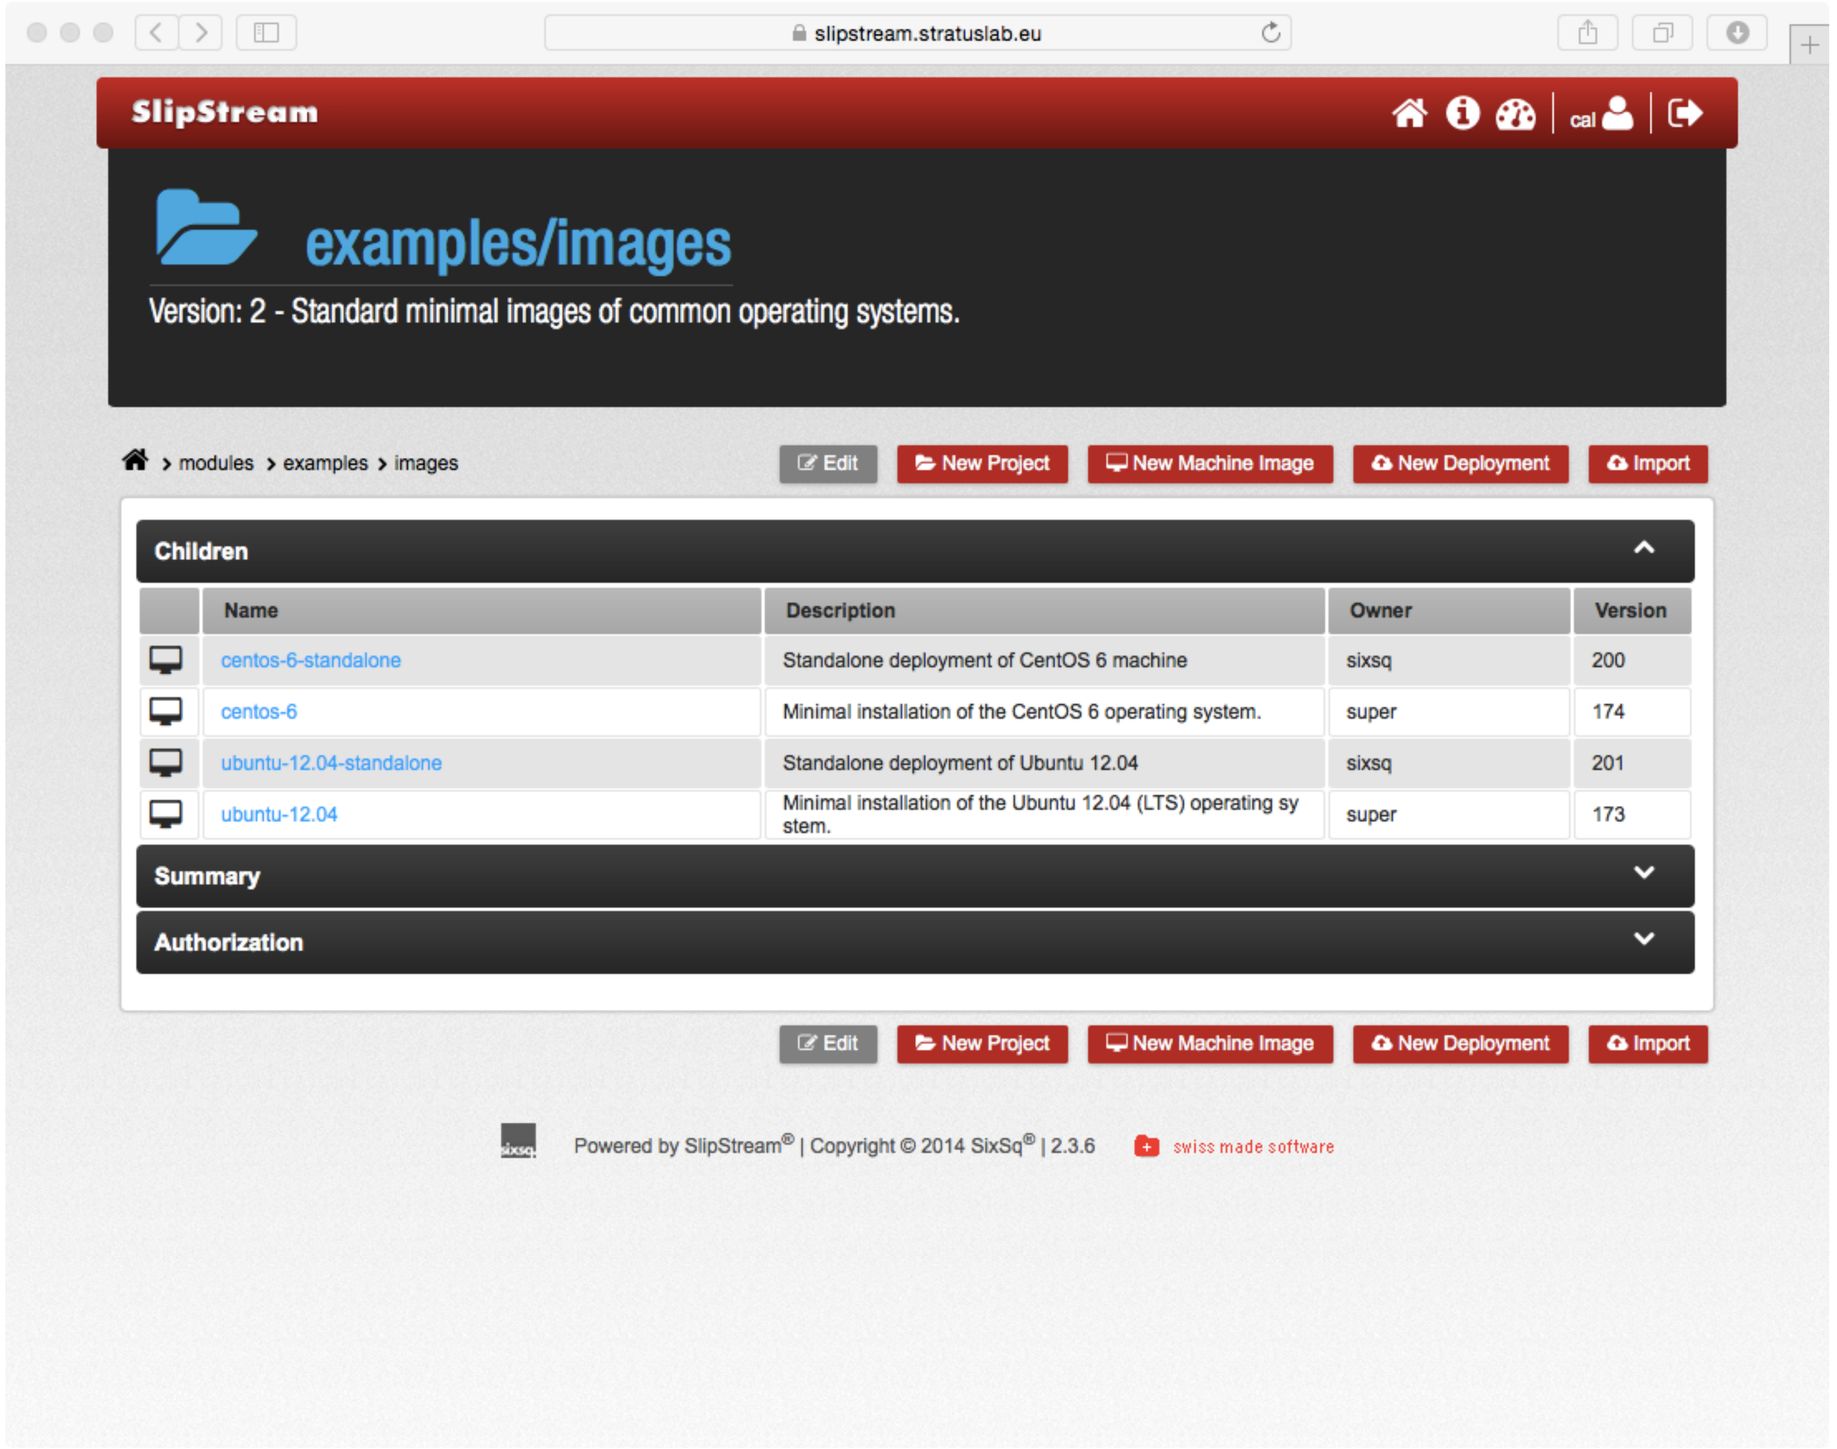
slipstream.stratuslab.eu
SlipStream
cal
examples/images
Version: 2 - Standard minimal images of common operating systems.
modules › examples › images
Edit
New Project
New Machine Image
New Deployment
Import
Children
Name
Description
Owner
Version
centos-6-standalone
Standalone deployment of CentOS 6 machine
sixsq
200
centos-6
Minimal installation of the CentOS 6 operating system.
super
174
ubuntu-12.04-standalone
Standalone deployment of Ubuntu 12.04
sixsq
201
ubuntu-12.04
Minimal installation of the Ubuntu 12.04 (LTS) operating sy stem.
super
173
Summary
Authorization
Edit
New Project
New Machine Image
New Deployment
Import
Powered by SlipStream® | Copyright © 2014 SixSq® | 2.3.6
swiss made software

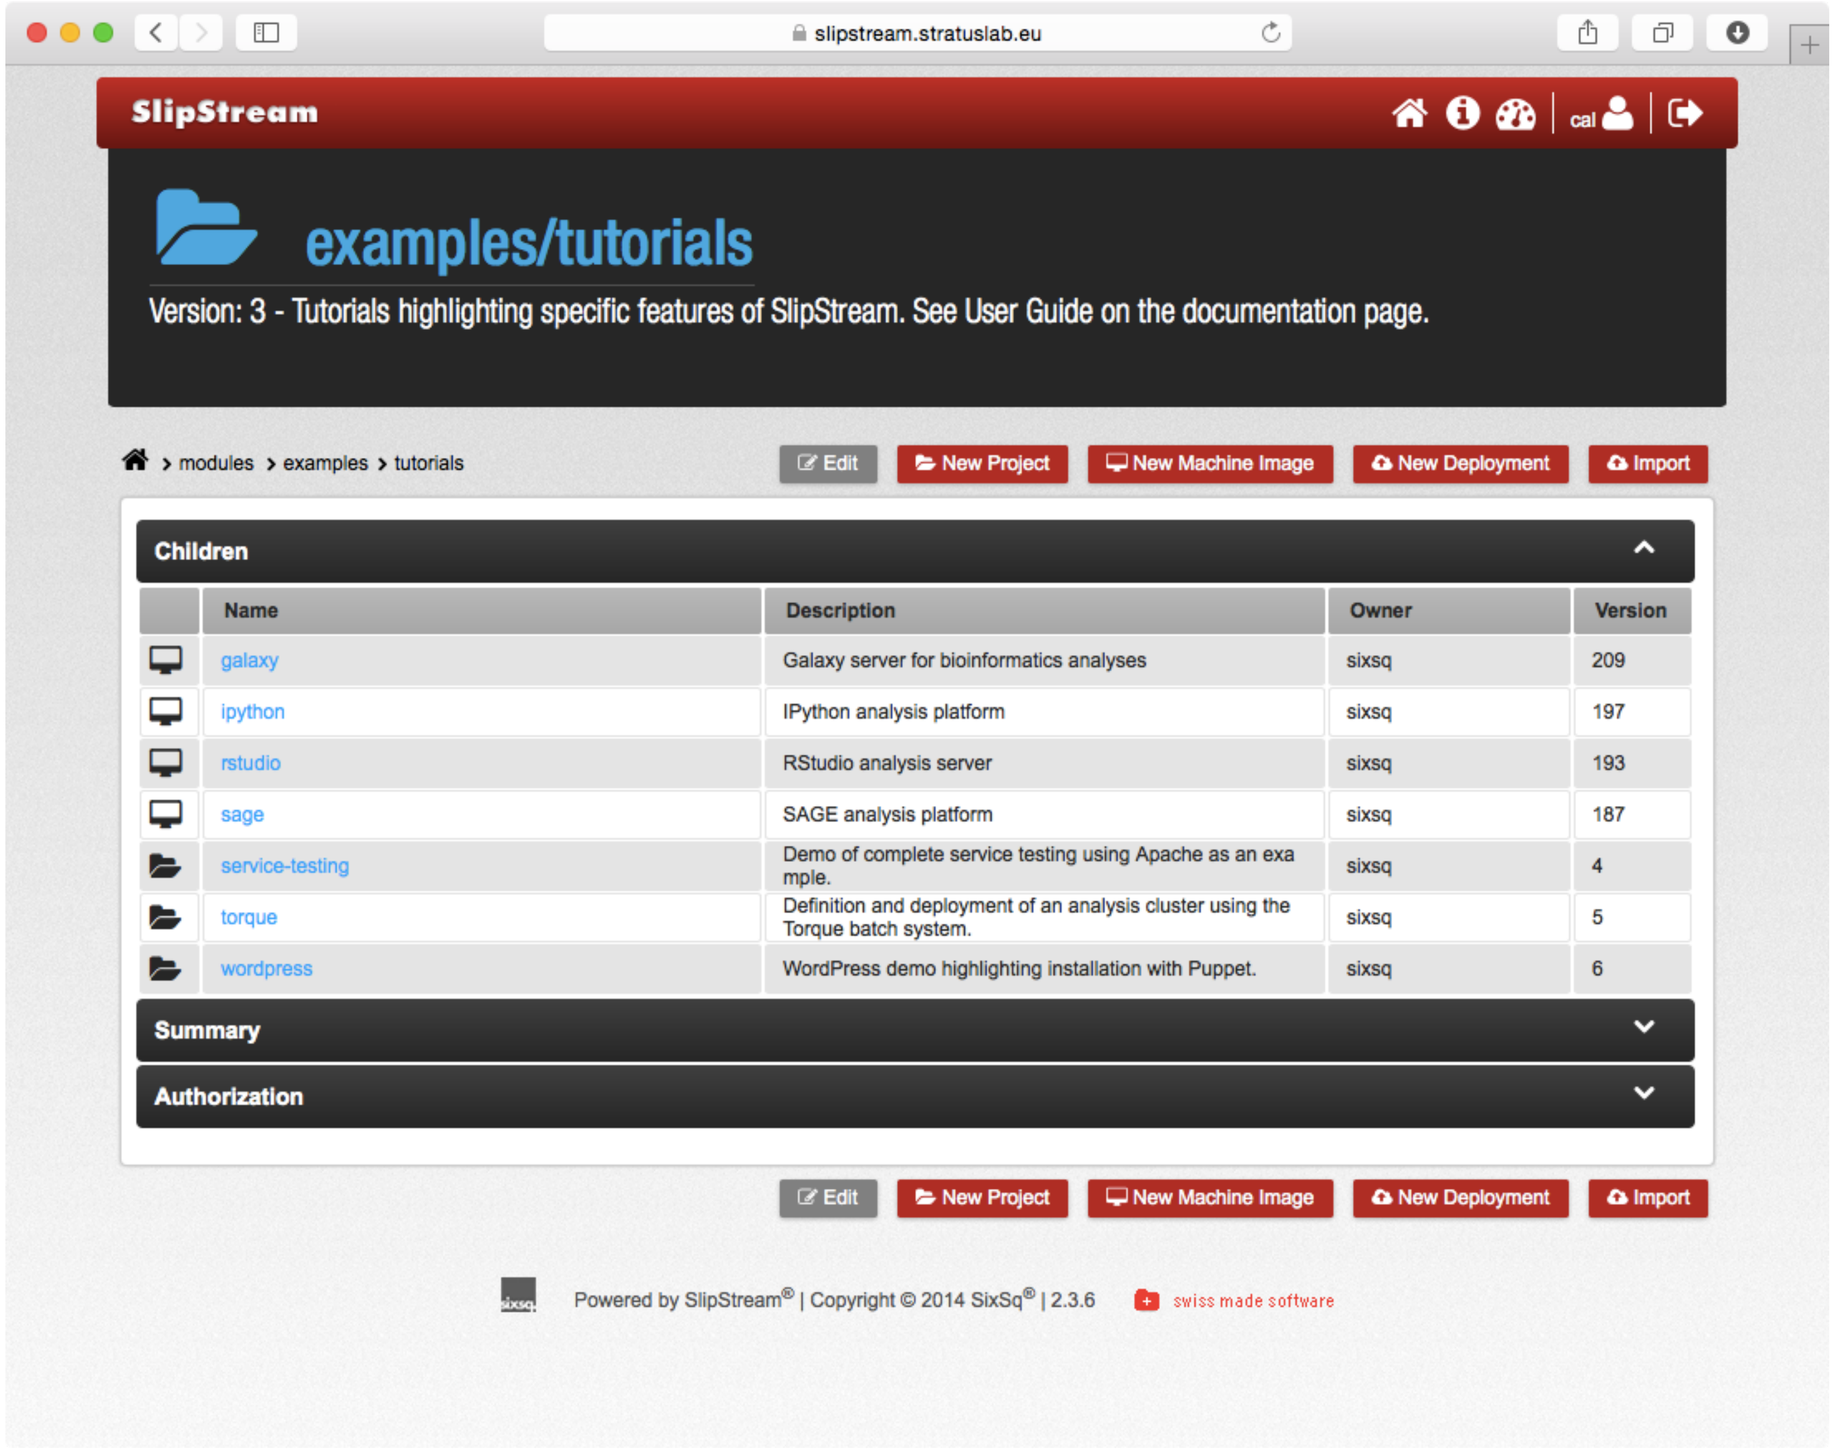
slipstream.stratuslab.eu
SlipStream
cal
examples/tutorials
Version: 3 - Tutorials highlighting specific features of SlipStream. See User Guide on the documentation page.
modules › examples › tutorials
Edit
New Project
New Machine Image
New Deployment
Import
Children
Name
Description
Owner
Version
galaxy
Galaxy server for bioinformatics analyses
sixsq
209
ipython
IPython analysis platform
sixsq
197
rstudio
RStudio analysis server
sixsq
193
sage
SAGE analysis platform
sixsq
187
service-testing
Demo of complete service testing using Apache as an exa mple.
sixsq
4
torque
Definition and deployment of an analysis cluster using the Torque batch system.
sixsq
5
wordpress
WordPress demo highlighting installation with Puppet.
sixsq
6
Summary
Authorization
Edit
New Project
New Machine Image
New Deployment
Import
Powered by SlipStream® | Copyright © 2014 SixSq® | 2.3.6
swiss made software

Image

- Describes the content of a virtual machine image.

- Two varieties:
  - Native Image
  - Derived Image

sixsq

Native Image

- Reference to a pre-existing image in a particular cloud infrastructure.

- These are created and maintained outside of SlipStream, but are used as the basis for derived images.
- Generally assume that minimal Ubuntu 12.04 and CentOS 6.x images exist for all clouds.

sixsq

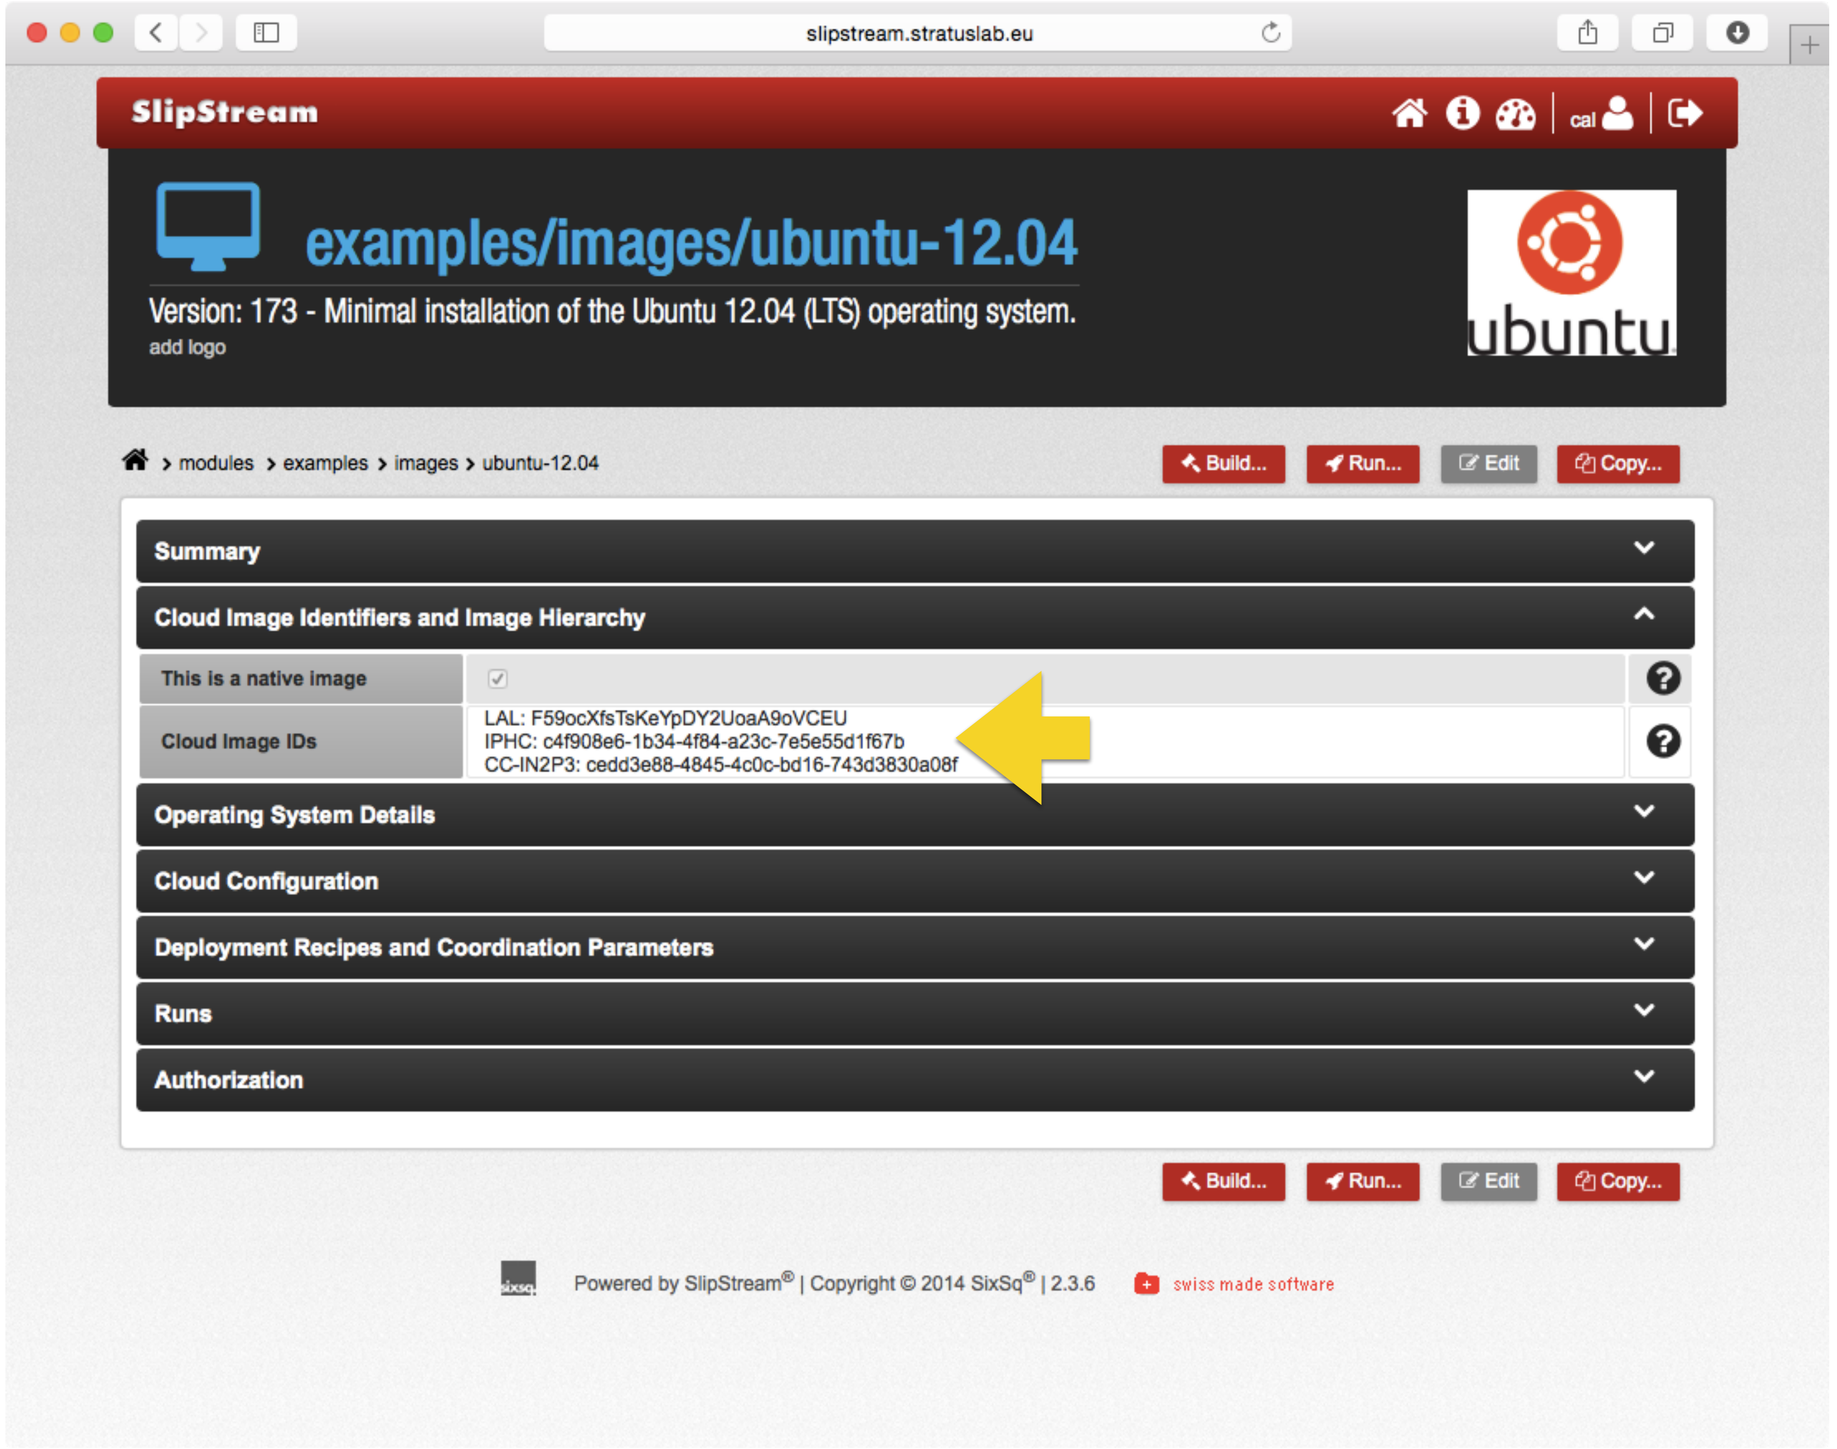
slipstream.stratuslab.eu
SlipStream
cal
examples/images/ubuntu-12.04
Version: 173 - Minimal installation of the Ubuntu 12.04 (LTS) operating system.
add logo
ubuntu
modules › examples › images › ubuntu-12.04
Build...
Run...
Edit
Copy...
Summary
Cloud Image Identifiers and Image Hierarchy
This is a native image
Cloud Image IDs
LAL: F59ocXfsTsKeYpDY2UoaA9oVCEU
IPHC: c4f908e6-1b34-4f84-a23c-7e5e55d1f67b
CC-IN2P3: cedd3e88-4845-4c0c-bd16-743d3830a08f
Operating System Details
Cloud Configuration
Deployment Recipes and Coordination Parameters
Runs
Authorization
Build...
Run...
Edit
Copy...
Powered by SlipStream® | Copyright © 2014 SixSq® | 2.3.6
swiss made software

Derived Image

- Inherits from a native image or another derived image, optionally adding “recipes” for enhancing the image.

- All derived images ultimately inherit from a native image within a particular cloud infrastructure.
- Start up time can be optimized by “creating” images on clouds that support it.

sixsq

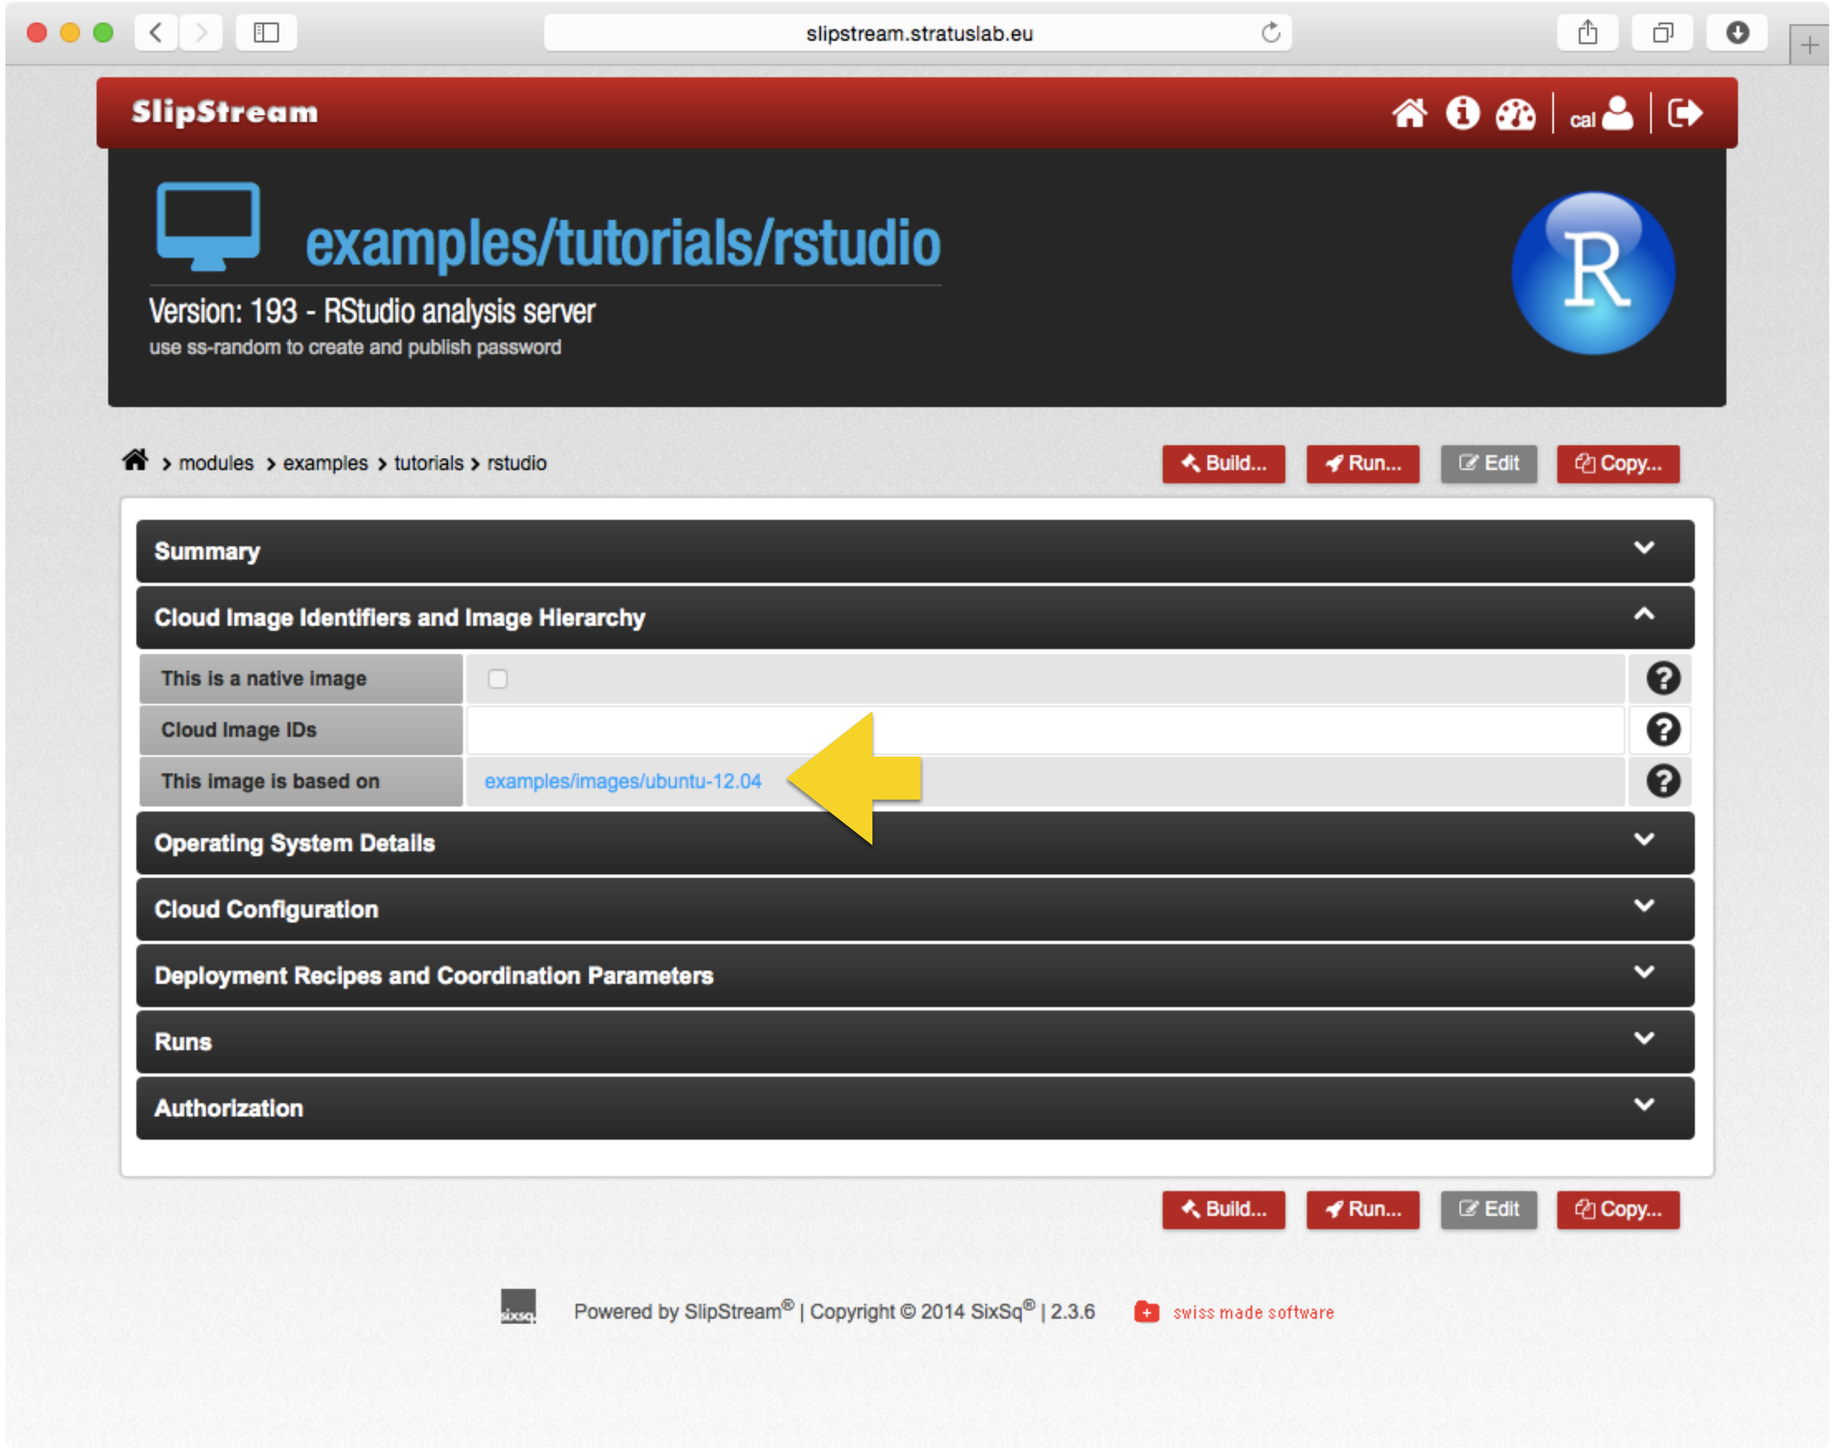
slipstream.stratuslab.eu
SlipStream
cal
examples/tutorials/rstudio
Version: 193 - RStudio analysis server
use ss-random to create and publish password
R
modules › examples › tutorials › rstudio
Build...
Run...
Edit
Copy...
Summary
Cloud Image Identifiers and Image Hierarchy
This is a native image
Cloud Image IDs
This image is based on
examples/images/ubuntu-12.04
Operating System Details
Cloud Configuration
Deployment Recipes and Coordination Parameters
Runs
Authorization
Build...
Run...
Edit
Copy...
Powered by SlipStream® | Copyright © 2014 SixSq® | 2.3.6
swiss made software

slipstream.stratuslab.eu

examples/tutorials/rstudio

Version: 193 - RStudio analysis server

use ss-random to create and publish password

modules › examples › tutorials › rstudio

Build... Run... Edit Copy...

Summary

Cloud Image Identifiers and Image Hierarchy

Operating System Details

Cloud Configuration

Deployment Recipes and Coordination Parameters

Execute | Parameters

Standard recipe (fires during Running phase)

```
16  #
17  # install RStudio
18  #
19  wget http://download2.rstudio.org/rstudio-server-0.97.551-amd64.deb
20  gdebi --non-interactive rstudio-server-0.97.551-amd64.deb
21
22  #
23  # put this on standard port
24  #
25  echo 'www-port=80' > /etc/rstudio/rserver.conf
26
27  #
28  # finish by upgrading entire system
29  #
30  apt-get -o DPkg::options::="--force-confdef" -o DPkg::options::="--force-confold" -y upgrade
31
32  #
33  # ==========
34  #
```

Runs

sixsq

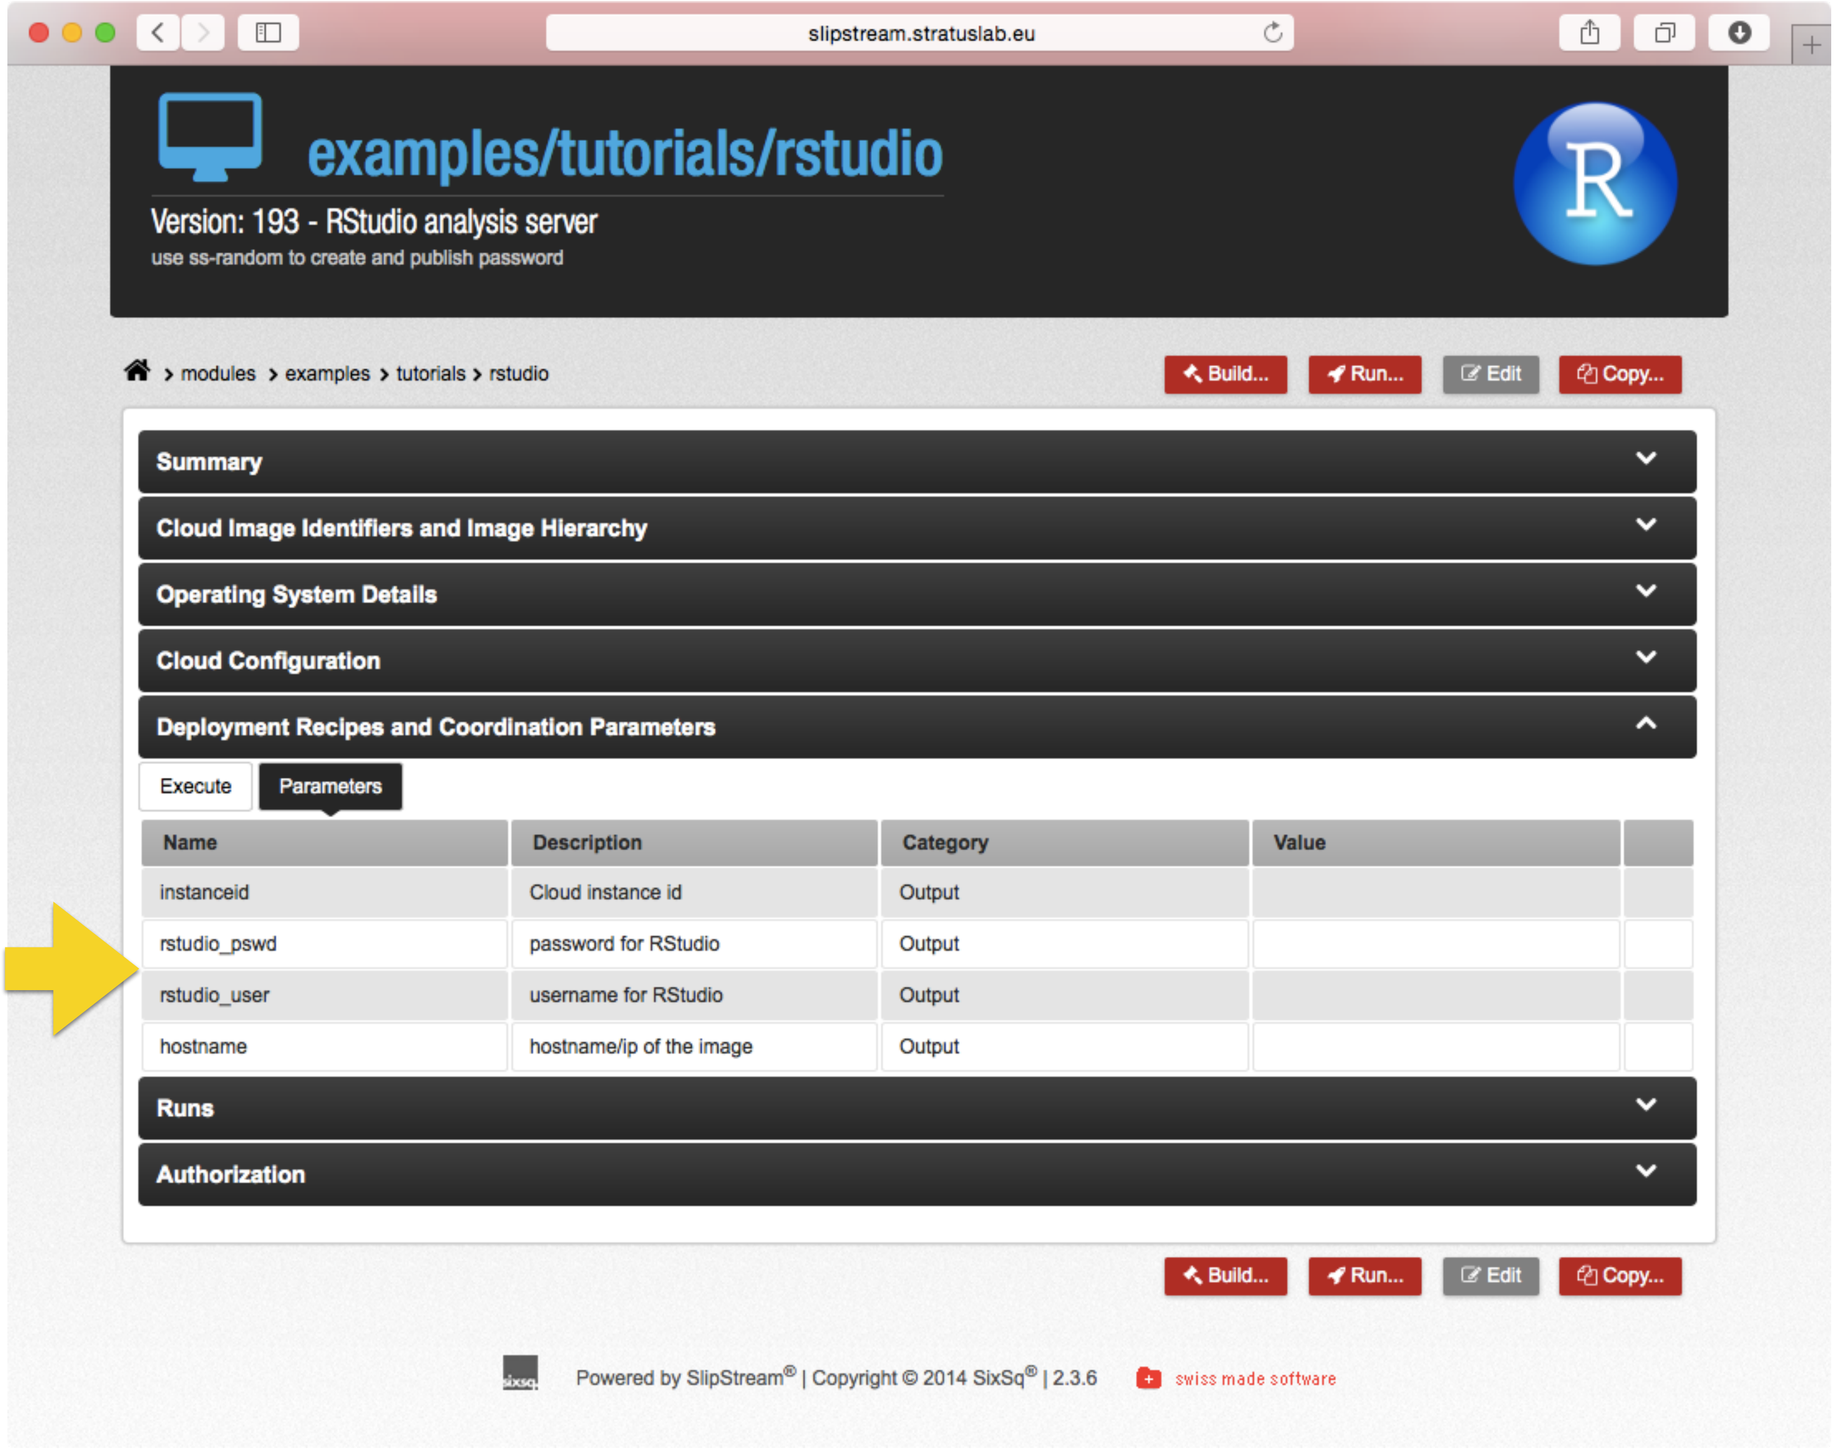
slipstream.stratuslab.eu
examples/tutorials/rstudio
Version: 193 - RStudio analysis server
use ss-random to create and publish password
modules › examples › tutorials › rstudio
Build...
Run...
Edit
Copy...
Summary
Cloud Image Identifiers and Image Hierarchy
Operating System Details
Cloud Configuration
Deployment Recipes and Coordination Parameters
Execute
Parameters
Name
Description
Category
Value
instanceid
Cloud instance id
Output
rstudio_pswd
password for RStudio
Output
rstudio_user
username for RStudio
Output
hostname
hostname/ip of the image
Output
Runs
Authorization
Build...
Run...
Edit
Copy...
Powered by SlipStream® | Copyright © 2014 SixSq® | 2.3.6
swiss made software

Deployments

- Group of virtual machines instances that are managed as a single entity.

- No need to create single-machine, deployments; images themselves can be run directly.
- In a multi-machine deployment, each machine has a multiplicity that can be changed at deployment time and optionally at runtime.

sixsq

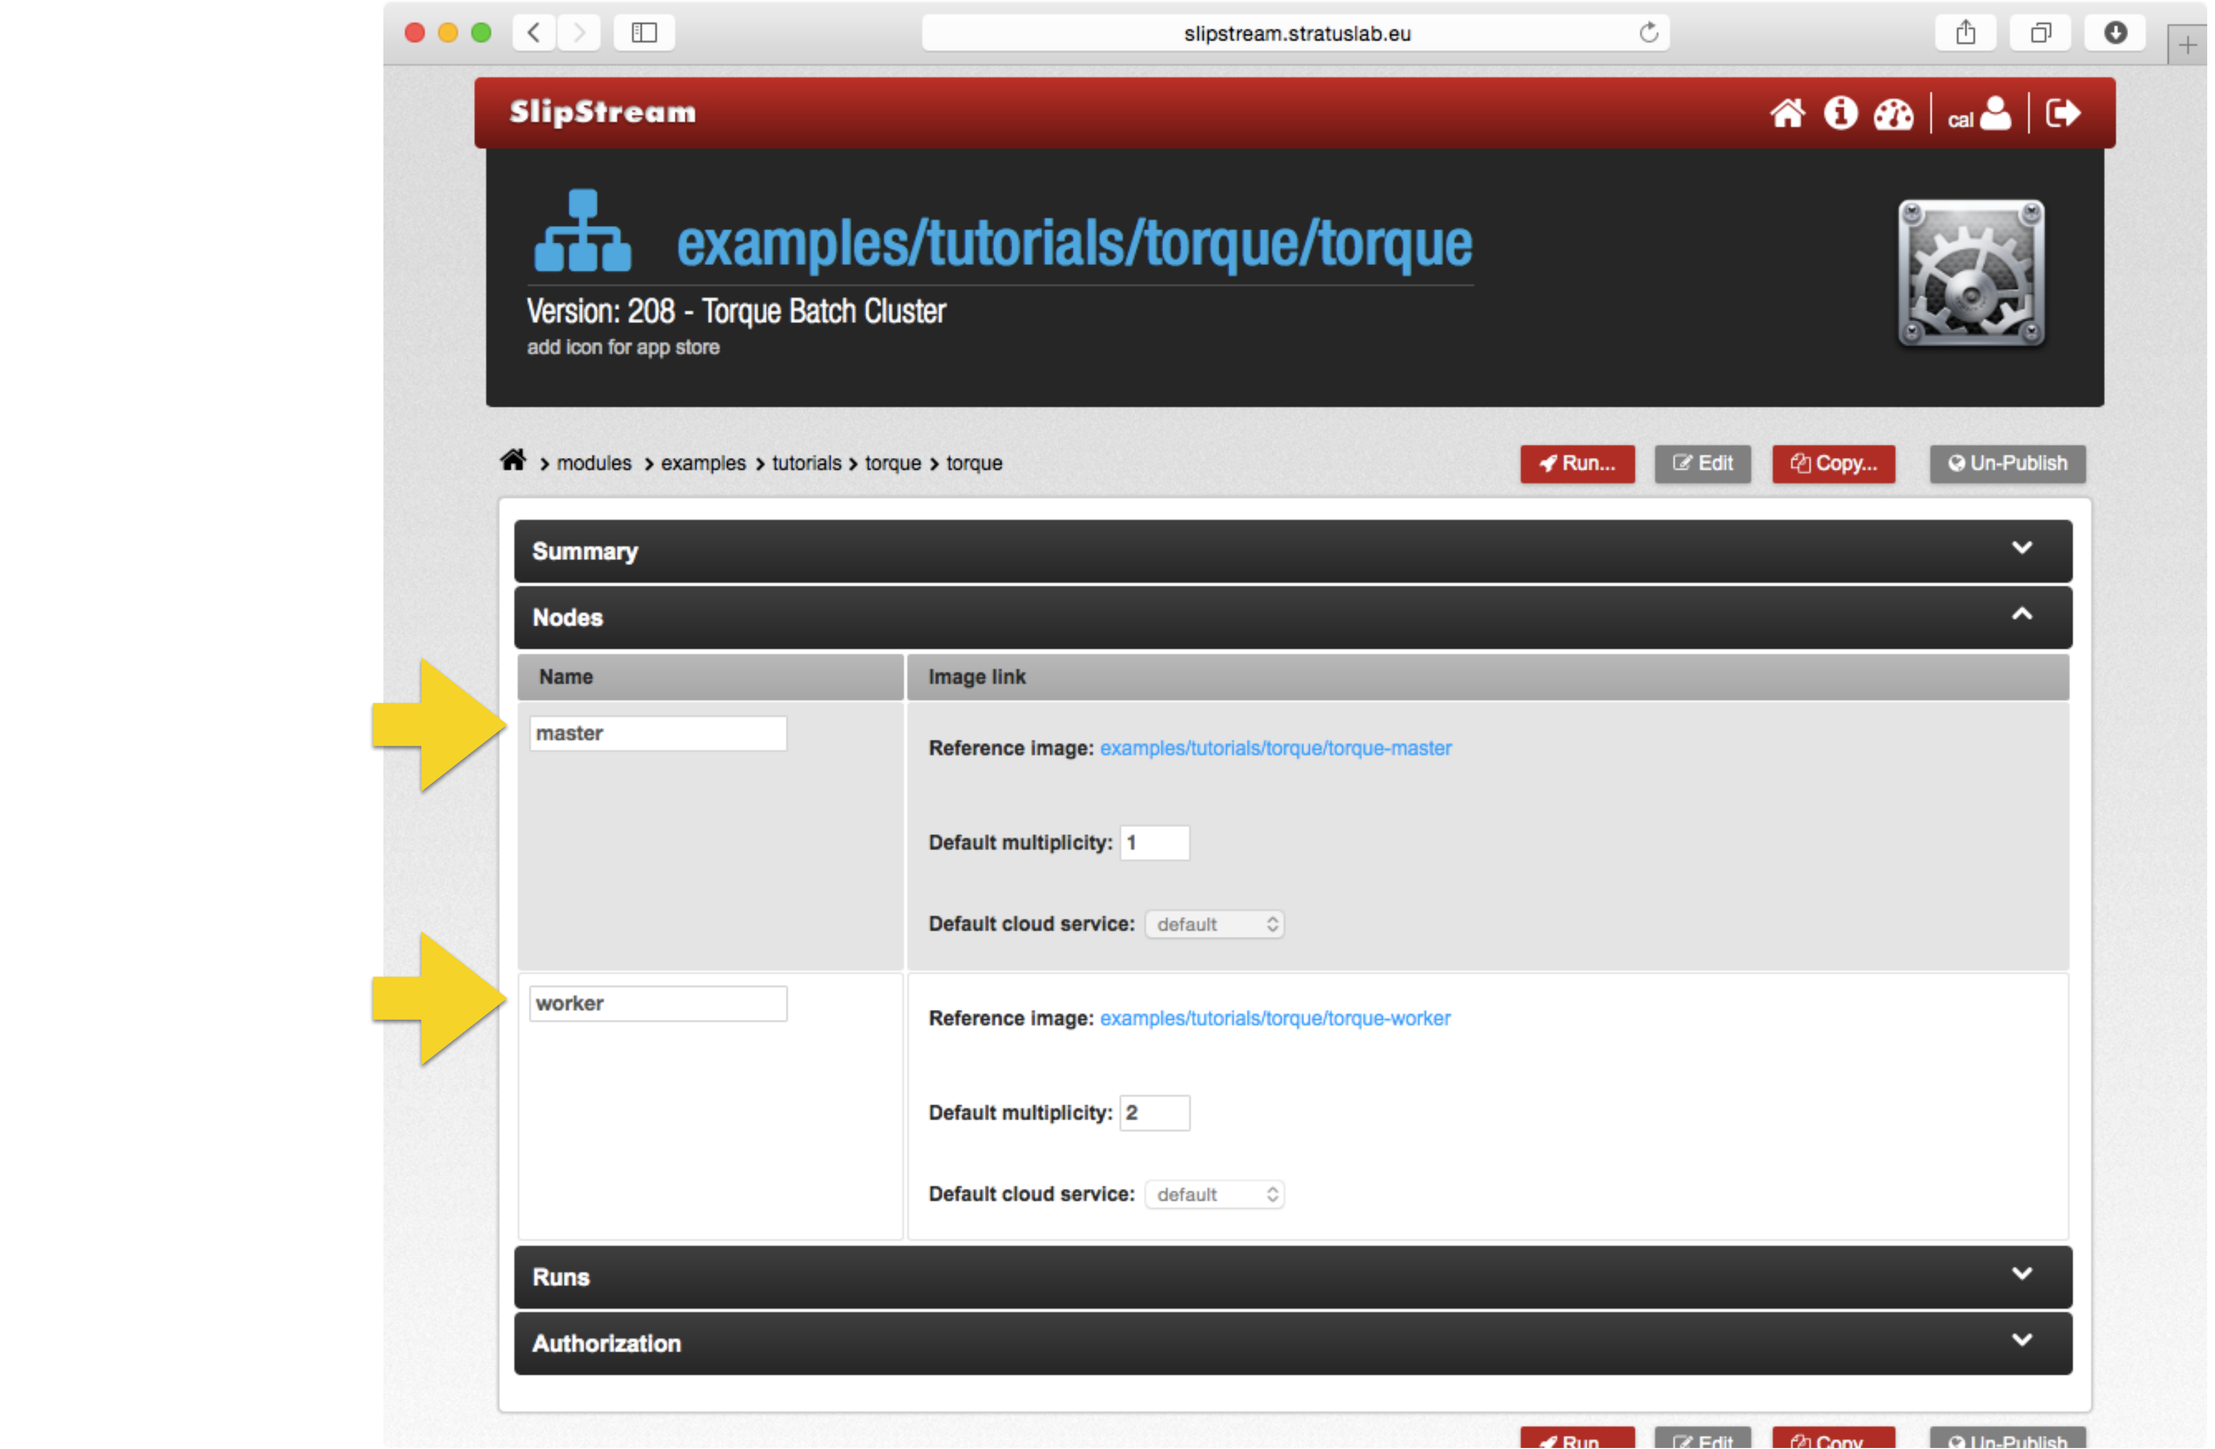
slipstream.stratuslab.eu
SlipStream
cal
examples/tutorials/torque/torque
Version: 208 - Torque Batch Cluster
add icon for app store
modules › examples › tutorials › torque › torque
Run...
Edit
Copy...
Un-Publish
Summary
Nodes
Name
Image link
master
Reference image: examples/tutorials/torque/torque-master
Default multiplicity: 1
Default cloud service: default
worker
Reference image: examples/tutorials/torque/torque-worker
Default multiplicity: 2
Default cloud service: default
Runs
Authorization

Runs

- Contains the state for a single deployment instance, includes the state of each virtual machine as well as the values of all defined parameters.

sixsq

Next Step

- Creating a web server deployment.

sixsq