
UForge AppCenter User Documentation

Release 3.8

FUJITSU

Jun 10, 2019

1	Getting Started	3
1.1	Signing In to UForge Portal	3
1.2	Keyboard Actions	5
1.3	Basic Concepts	6
2	Managing Your Accounts	9
2.1	Modifying Your User Profile	9
2.2	Viewing Your Statistics	10
2.3	Changing Your Password	11
3	Managing Your Credentials	13
3.1	Managing Cloud Accounts	13
3.2	Managing Your Artifact Accounts	15
3.3	Managing API Keys	16
3.4	Managing SSH Keys	17
4	Managing Appliance Templates	19
4.1	Supported Operating Systems	19
4.2	Creating an Appliance Template	20
4.3	Modifying an Appliance Template	24
4.4	Managing the OS Profile	25
4.5	Updating the Install Profile	26
4.6	Managing Configuration	33
4.7	Adding Custom Software Components	35
4.8	Tracking OS Package Updates	42
4.9	Listing Appliance Templates	44
4.10	Cloning an Appliance Template	45
4.11	Importing and Exporting Templates	45
5	Generating and Publishing Machine Images	51
5.1	Supported Machine Image Types	51
5.2	Generating a Machine Image	53
5.3	Publishing a Machine Image	57
6	Deploying a Machine Image	63
6.1	Deploying a Machine Image	63
6.2	Listing Deployed Machine Images	66

6.3	Launching a PXE Image	67
7	Blueprinting	71
7.1	Creating a Blueprint	71
7.2	Deploying a Blueprint	74
8	Migrating Live Workloads	77
8.1	Migration Pre-requisites	77
8.2	Migration Process Overview	79
8.3	Migrating Using Lift & Shift	81
8.4	Migrating Using Re-platform	85
8.5	Scanning a Linux Machine	86
8.6	Scanning a Microsoft Windows Machine	88
8.7	Viewing and Using a Scan	91
8.8	Differences between Source and Target Instances	102
9	Using Workspaces	107
9.1	Creating a Workspace	107
9.2	The Activity Stream	108
9.3	Managing Workspace Members	108
9.4	Sharing an Appliance Template in a Workspace	109
9.5	Adding a Comment to a Shared Appliance Template	109
10	Using the REST API	111
10.1	Response & Error Codes	111
10.2	Sending a Request	113
10.3	Response Body Types	114
10.4	Using the API Keys	114
10.5	Query Parameters	115
10.6	REST API Examples	115
11	Using the Java SDK	125
11.1	Download and Installing the SDK	125
11.2	Communicating with UForge	125
11.3	Creating an Appliance Template	126
11.4	Adding an OS Profile	127
11.5	Generating a Machine Image	128
11.6	Publishing an Image	128
11.7	Adding a Project from the Project Catalog	129
11.8	Uploading a Software Component	130
11.9	Adding a Boot Script	130
12	Using the Python SDK	133
12.1	Download and Installing the SDK	133
12.2	Communicating with UForge	135
12.3	Creating an Appliance Template	135
12.4	Creating My Software	137
12.5	Creating a Project for a Specific OS	139
13	Hammr Command Line Tool	143
13.1	Getting the Source Code	143
13.2	Further Reading	144
14	Using UForge CLI	145
14.1	Installing UForge CLI	145

14.2	Launching UForge CLI	145
15	3rd Party Components	147
15.1	Trademarks	147
15.2	Copyright FUJITSU LIMITED 2019	147
15.3	3rd Party Components Used	147
16	Changelog	153
16.1	3.8.fp13	153
16.2	3.8.fp12	154
16.3	3.8.fp11	155
16.4	3.8.fp10	156
16.5	3.8.fp9	157
16.6	3.8.fp8	158
16.7	3.8.fp7	160
16.8	3.8.fp6	161
16.9	3.8.fp5	162
16.10	3.8.fp4	163
16.11	3.8.fp3	164
16.12	3.8.fp2	165
16.13	3.8.fp1	167
16.14	3.8	169
16.15	3.7.fp8	170
16.16	3.7.fp7	172
16.17	3.7.fp6	173
16.18	3.7.fp5	174
16.19	3.7.fp4	175
16.20	3.7.fp3	176
16.21	3.7.fp2	178
16.22	3.7.fp1	180
16.23	3.7	181
16.24	3.6.fp2	182
16.25	3.6.fp1	184
17	Trademarks	187
18	Copyright FUJITSU LIMITED 2019	189
19	High Risk Activity	191
20	Export Restrictions	193

Note: There are multiple options for reading this documentation - click on the link at the lower left hand corner for these options.

Contents:

1.1 Signing In to UForge Portal


To sign in, go to the UForge Portal sign in page:

<https://your-uforge-server-hostname/uforge/>

The UForge Portal has the following pages, accessible from the left-hand sidebar:

- **Dashboard:** It shows statistics on your UForge usage.
- **Apps:** This is where your appliances are created and listed. You also go to this page to add custom software, update packages in appliances, and create images, among other things.
- **Blueprints:** This page allows you to create and deploy multi-node applications.
- **Deployments:** This area lists published images that have at least one instance running on Amazon, Open-Stack, Cloudstack and Microsoft Azure. This menu is only visible if you have the correct entitlements. Contact your UForge administrator for more information.
- **Collaboration:** This is a private area where you can share appliances with other users who are part of your workspace. These users must be invited and join your workspace. They can be part of your organization or part of another organization.
- **Migrations:** This is where you can launch a migration or scan.
- **Credentials:** This is where you manage your cloud account information, SSH keys and API keys.
- **Profile:** This is your UForge account information.
- **Administration:** (only for administrators of platform). Provides administration tasks including operating system and formats management.

Warning: Depending upon your access rights one or more of these tabs may not be visible.



1.1.1 Supported Browsers for UI Access

The following browser versions are officially supported when using the user interface:

- Firefox v35 or later
- Chrome v29 or later
- Safari v9 or later
- Internet Explorer 11 or later
- Opera v15 or later
- Edge v42 or later

Note: Javascript is required when using the UI. The following error will appear if you have disabled Javascript (check your browser configuration, plugins or security settings).



1.1.2 Language Selection for the Graphical User Interface

The UForge UI is accessible in different languages. You can select the language using the the drop-down menu at the top right-hand of the UI, next to the sign-in button.



Note: The language choice is saved at the level of the browser and not of the user account.

1.2 Keyboard Actions

1.2.1 Escape Key

The UForge UI uses the Escape (ESC) key as a hotkey to cancel or close a popup window without saving or executing any actions. However, in some cases, using the hotkey ESC will redirect you to a different page from the one from which the popup was opened.

1.3 Basic Concepts

1.3.1 Organization

UForge AppCenter is a multi-tenant platform which can serve multiple users. All the resources of the platform are held within an Organization. The organization contains:


- One or more users
- One or more operating systems
- A project catalog containing software components that can be used by its users
- One or more formats available to generate images



1.3.2 User

Each user on the platform has:

- an **Appliance Library** containing all the appliance templates created by the user
- a **Software Library** (My Software) containing all the custom software uploaded by the user, which can be used in one or more appliance templates
- a list of one or more **Cloud Accounts** to allow the user to publish and register generated machine images to various cloud and virtualization platforms



1.3.3 Appliance Templates

An Appliance Template is meta-data describing a software stack. It consists of five layers, namely:

- an `Install Profile` (mandatory) - specific information for the first time the image boots
- an `OS Profile` (mandatory) - a list of operating system packages. Each operating system within the organization provides one or more standard OS profiles to choose from when creating the OS Profile of the appliance template. It is also possible to create custom OS profiles.
- `Projects` (optional) - a list of Project software components chosen from the Organization's Project Catalog
- `My Software` (optional) - a list of software components chosen from the User's private "Software Library"
- `Configuration` (optional) - configuration information including boot scripts and/or other software components to manage the image after provisioning

Depending on the user's roles and privileges, the user may only have access to a restricted number of operating systems, projects and image formats the organization has to offer.

Using an appliance template, the user can generate machine images in different formats. For some formats, the user can publish and register machine images to a target cloud or virtualization platform. Each appliance template stores meta-data regarding each machine image generated and published.

1.3.4 Workspaces

Each user can also create and join **workspaces**. A workspace is an area for members to collaborate and share appliances. A workspace is created and maintained by users. The user can invite members to be part of a workspace. When the user invites a member that is not part of the UForge database, an email is sent to the new member to invite them to register on UForge.

The workspaces are listed under the `Collaboration` tab. Each workspace has:

- an activity stream, which lists the members' comments, the activities, such as invites and appliances shared
- a templates page, which lists all the templates shared with the people that are part of the workspace
- a members page where the user who created the workspace (the workspace administrator) can invite new members, delete members or change the role of a workspace member.

Members of a workspace are either:

- **Administrator.** This is generally the user who has created the workspace. There can be several administrators in a workspace. The administrator can invite or delete members and is able to delete a workspace. The administrator has all the same basic rights as the collaborator.
- **Collaborator.** The collaborator has the same basic rights as the Guest, but can also share templates.
- **Guest.** They can read and post to the activity stream, and import templates into their private appliance library.

Managing Your Accounts

The `Profile` tab on the left-hand sidebar allows you to modify a large number of your personal details, including your password.

The following sub-sections go into detail on how to manage your account information:

2.1 Modifying Your User Profile

You can edit your user profile by going to the `Profile` page. The only mandatory information is your email address.

The screenshot shows a user profile editing interface. On the left, there's a sidebar with a hamburger menu icon and a 'save' button. The main content area is divided into three sections: 'Account Summary', 'Address', and 'Company Information'. Each section has a set of input fields for user details. The 'Account Summary' section includes fields for First Name, Surname, Email, and Website. The 'Address' section includes fields for City/Town, State, Country, Home Phone, and Mobile Phone. The 'Company Information' section includes fields for Company, Job Title, Company Website, and Office Phone. A 'save' button is located in the top right corner of the form.

You can add or modify your name, company information and address and click `save`.

To add a photo to your profile or modify the existing one:

1. Go to the `My Account` page.
2. Click on the `Profile` tab.
3. Click on the round photo icon on the left hand side.
4. Navigate to the desired image (preferably a .jpeg or .png).

5. Click Open.

2.2 Viewing Your Statistics

2.2.1 Your Usage

When you log in to the UForge Portal, the first page you will see is the `Dashboard > Usage` sub tab. This view displays your statistics on the platform.

The first part of the page displays your current usage and quota information for appliance templates, software components, machine image generations and scans created.

The `Appliances` summary shows the number of appliance templates you currently have in your Appliance Library. Your quota usage and limit is displayed under the summary information. If you have reached your quota limit, you can delete existing appliance templates to free up space to create new ones. This information is not shown if you do not have access to creating and managing appliance templates.


The `Software Components` summary is the number of software components (also known as software bundles) you currently have in your Software Library. Your quota usage and limit is displayed under the summary information. If you have reached your quota limit, you can delete existing software components to free up space to create new ones. This information is not shown if you do not have access to creating and managing appliance templates (as these software components are added to appliance templates). This information is not shown if you do not have access to creating and managing appliance templates.

The `Generations` summary is the number of machine images you have currently stored on the platform. These machine images are generated from appliance templates or scans. Your quota usage and limit is displayed under the summary information. The quota information is measured on the number of successful machine images you have requested over a period of time; and not necessarily the number of machine images you currently have listed on your account. Even if you delete an existing machine image, the quota usage will not be changed. Your quota usage will be reset if you have a quota that is associated with a frequency (e.g. 5 generations per month). If you cancel a generation or if there is an error during the generation, the quota value is not updated.

The `Scans` summary displays the current number of successful live machine scans stored on the platform. Your quota usage and limit is displayed under the summary information. The quota information is measured on the number of successful scans you have requested over a period of time. This number does not necessarily reflect the number of scans you currently have in your account. If you delete a scan from your account the quota value is not changed. Your quota usage will be reset if you have a quota that is associated with a frequency (e.g. 5 scans per month). If you cancel a scan or if there is an error during a scan, the quota usage is not updated. This information is not shown if you do not have access to migration features.


The rest of the `Usage` page provides more detailed activity information including:

- `Your Disk Usage`. This shows the disk space consumed by you for storing software components, machine images and scan data. This includes any software that has been imported. If you have reached your quota limit, you can delete software components, machine images or scans to save space.
- `An Appliance Activity` displaying a timeline of the number of appliance templates and machine image generations created. It also provides a breakdown of operating systems used in the appliances and the types of formats of the generated machine images.
- `A Scan Activity` displaying a timeline of the number scans and machine image generations created. It also provides a breakdown of operating systems found for each scan and the types of formats of the generated machine images.



2.2.2 Generations

The Dashboard > Generations view provides you a current list of machine images stored in your account along with any machine images you have published to one or more cloud environments.



2.3 Changing Your Password

To modify your password:

1. Go to the **Profile** tab on the left-hand sidebar.
2. In the top right, click on the key icon.

change password

Change Your Password

Current Password

New Password

Verify New Password

Town

State

Country

Phone

Phone

3. Enter your current password and your new password.
4. Click ok.

Managing Your Credentials


The `Credentials` tab on the left-hand sidebar allows you to modify your credentials such as: Cloud Account information, API Keys and SSH keys.

The following sub-sections go into detail on how to manage your credentials:

3.1 Managing Cloud Accounts

In order to publish an image to a cloud using UForge AppCenter, you will need to add your cloud credentials to UForge. You must have a cloud account prior to setting up your credentials on the platform. Have all the information for your cloud account on hand before starting.

1. Go to the `Credentials` tab in the left-hand sidebar.
2. Go to `Cloud Account`.
3. Select the type of account you want to create from the `Add account` drop-down menu at the top right.



4. Scroll over any given field for more information on the mandatory information to provide for a given cloud account.

Note: In order to set up the credentials for a Microsoft Azure account, there are a number of steps you must complete first. For more information refer to [Setting up a Microsoft Azure Account](#).

5. Click `Create` to complete.

3.1.1 Setting up a Microsoft Azure Account


If you want to publish images to Microsoft Azure, you must first set up the correct Microsoft Azure credentials in UForge. To do so, first, from your Microsoft Azure account you will need to:

1. Create an Azure Active Directory application (as described in the [official Microsoft Azure documentation related to Azure Active Directory creation](#)).
2. Get the subscription ID from your Microsoft Azure account.
3. Get the application ID and authentication key (as described in [official Microsoft Azure documentation related to application ID and authentication key](#)).
4. Retrieve the tenant ID which you will need to enter in UForge credentials (as described in [official Microsoft Azure documentation related to tenant ID retrieval](#)).
5. Assign the `contributor` role to the application (as described in [official Microsoft Azure documentation related to application role assignment](#)).

Note: The above steps are also required when setting up a Microsoft Azure account for Azure Stack.

Then you will need to fill in your credential information in UForge. To do so:

1. Go to the `Credentials` tab in the left-hand sidebar.
2. Go to `Cloud Account`.
3. Select `Microsoft Azure` from the `Add account` drop-down menu at the top right.



Note: When creating a cloud account for Azure Stack, it is also necessary to turn on `Use private endpoint`, and in `Endpoint URL` input the Azure Stack's Resource Manager Endpoint. For example: `https://management.local.azurestack.external`.

3.1.2 Setting up an AWS Account to Publish Windows Images

Note: Before publishing a Windows image to AWS, `vmimport` role must be properly assigned to an AWS account. An access right to S3 bucket is also required for an account. Refer to the procedure in [VM Import Service Role section of the AWS documentation](#). Because the name of a S3 bucket changes for each image publish, you should give access to any buckets using a wildcard `*`.

The following is an example of `role-policy.json`:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "s3:GetBucketLocation",
        "s3:GetObject",
        "s3:ListBucket"
      ],
      "Resource": [
        "arn:aws:s3:::*"
      ]
    },
    {
      "Effect": "Allow",
      "Action": [
        "ec2:ModifySnapshotAttribute",
        "ec2:CopySnapshot",
        "ec2:RegisterImage",
        "ec2:Describe*"
      ],
      "Resource": "*"
    }
  ]
}
```

3.2 Managing Your Artifact Accounts

In order to download files for My Software, you can create an artifact account in UForge.

1. Go to the `Credentials` tab in the left-hand sidebar.
2. Go to `Artifact Account`.
3. Click on `create`.

4. Enter a name for the account.
5. Select the type from the drop-down menu.
6. Enter the host name and port.
7. Click `create` to complete.

3.3 Managing API Keys

API Keys are used to communicate with UForge AppCenter more securely when using the platform APIs (rather than using basic authentication).

If you have the rights, you will be able to manage your API keys. See your UForge Administrator to be given these rights, if needed.

The number of API key pairs that a user can create is set by the administrator. If you have reached your quota of key pairs, contact your administrator.

To create a key pair:

1. Go to the `Credentials` on the left-hand sidebar.
2. Go to `API Keys`.
3. Click `create`.
4. A pop-up window will appear. If you want to add a tag or description, click on the toggle. Enter the tag and or description (optional).

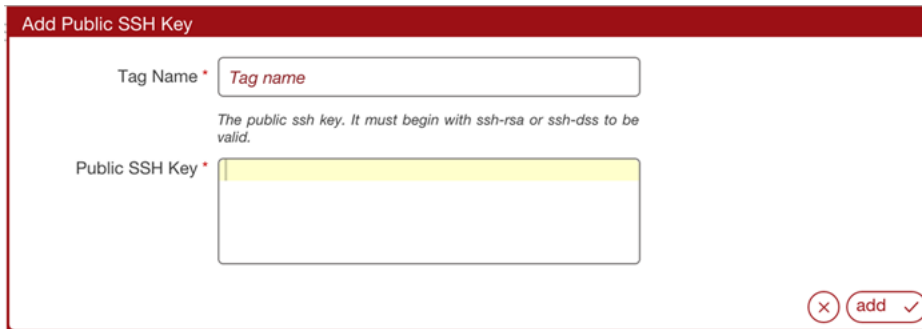
5. Click `ok`.

3.4 Managing SSH Keys

You can manage one or more SSH keys that can be added to an appliance template.

To add an SSH key:

1. Go to the `Credentials` tab on the left-hand sidebar.
2. Go to `SSH Keys`.
3. Click `create`.
4. Enter the key name and the key details.



The screenshot shows a form titled "Add Public SSH Key" with a red header bar. It contains two input fields: "Tag Name" with a red asterisk and a placeholder "Tag name", and "Public SSH Key" with a red asterisk and a yellow highlight. A red error message is visible below the "Tag Name" field: "The public ssh key. It must begin with ssh-rsa or ssh-dss to be valid." At the bottom right, there are two buttons: a red "x" button and a red "add" button with a checkmark.

5. Click `add`.

Managing Appliance Templates

An appliance template is a meta-data description of an entire software stack. The following sub-sections go into detail on how to create and manage your appliance templates:

4.1 Supported Operating Systems

The following is a list of supported OSes that users can use as the guest operating system when creating their appliance templates.

OS	Factory	Migration
CentOS	5.2+ (32bit and 64bit), 6 (32bit 64bit), 7 (64bit)	5.2+ (32bit and 64bit), 6 (32bit
Debian	6 (Squeeze), 7 (Wheezy), 8 (Jessie), 9 (Stretch) (32bit and 64bit)	6 (Squeeze), 7 (Wheezy), 8 (Jessie), 9 (Stretch) (32bit and 64bit)
Fedora	8 to 22	8 to 22
SUSE Linux Enterprise Server	11.3, 11.4 (32bit and 64bit) 12.1, 12.2, 12.3 (64bit)	11.3, 11.4 (32bit and 64bit) 12.1, 12.2, 12.3 (64bit)
Open SUSE	11.3, 11.4 (32bit and 64bit) 12.1, 12.2, 12.3 (64bit) 42.1 42.2 42.3 (64 bit)	11.3, 11.4 (32bit and 64bit) 12.1, 12.2, 12.3 (64bit) 42.1 42.2 42.3 (64 bit)
Red Hat Enterprise Linux*	5.2+ (32bit and 64bit), 6 (32bit and 64bit), 7 (64bit)	5.2+ (32bit and 64bit), 6 (32bit and 64bit), 7 (64bit)
Oracle Linux	5.2+ (32bit and 64bit), 6 (32bit and 64bit), 7 (64bit)	5.2+ (32bit and 64bit), 6 (32bit and 64bit), 7 (64bit)
Scientific Linux	5.2+, 6 (32bit and 64bit), 7 (64bit)	5.2+, 6 (32bit and 64bit), 7 (64bit)
Ubuntu LTS	10.04 (Lucid), 12.04 (Precise), 14.04 (Trusty), 16.04 (Xenial), 18.04 (Bionic) (32bit and 64bit)	10.04 (Lucid), 12.04 (Precise), 14.04 (Trusty), 16.04 (Xenial), 18.04 (Bionic) (32bit and 64bit)
Microsoft Windows Server	2008R2, 2012, 2012R2, 2016	2008R2, 2012, 2012R2, 2016

Note: For Red Hat Enterprise Linux you must provide the ISO images or access to a repository.

Warning: If you intend to generate machine images for cloud environments, ensure that the operating system you are using in the machine image is correctly supported by the cloud environment. For example, Microsoft Azure supports the following [operating systems](#).

4.1.1 Notes on Licensing

When using UForge, you have to comply with the license agreement of OSes and software which UForge handles, in particular:

- **Publishing OS image of RHEL (Red Hat Enterprise Linux) subscription to public cloud** Cloud provider has to be CCSP (Certified Cloud & Service Provider) and you must register to Red Hat Cloud Access. For more details, please confirm with cloud provider.

- **Scanning server** You have to check whether the licenses of OS and software which the source machine contains allow you to use them on the destination server which you are migrating to.

If the source machine contains rpm packages which Red Hat provides, check with the administrator whether UForge repository contains these packages, because UForge automatically regenerates rpm packages not contained in the repository, and regenerated packages are NOT supported by Red Hat.

On UForge Portal, you can see the list of rpm packages which the source machine contains and header `In Repo` indicates whether or not the package comes from the repository (refer to [Viewing and Using a Scan](#)). Once migration is done, you can see where the package comes from by rpm command on the destination server. If regenerated, `Build Host` is overwritten as `uforge`.

- **Handling Microsoft Windows** UForge users must acquire Windows license in order to handle Windows OSes in UForge. When publishing Windows OS image or scanning Windows server, you have to confirm usage conditions of cloud provider and virtualization software which you publish to or scan.

4.2 Creating an Appliance Template

You can create either Linux-based or Windows-based appliance templates. The steps differ slightly. Please refer to the appropriate section below.

4.2.1 Creating a Linux-based Appliance

To create a new appliance in your private workspace:

1. Select `Apps` icon on the left.
2. On the `Appliance Library` page, click on `create` in the top right.
3. Enter the `Name` and `Version` of the appliance.

4. From the drop-down menus, select the operating system (distribution, release and architecture).
5. Click the `create` button. This creates a skeleton of an appliance template in the platform which you can now customize with operating system packages, middleware and application software.

Note: If SELinux is installed (ie the file `/etc/selinux/config` exists), the filesystem will be relabeled on the first boot of a UForge generated machine image in order to add the SELinux context in the all system files extended attribute. At boot time, `init.rc` checks for the existence of `/.autorelabel`. If this file exists, SELinux performs a complete file system relabel (using the `/sbin/fixfiles -f -F relabel` command), and then deletes `/.autorelabel`.)

6. You should now see the appliance overview page. You can add a description to your appliance (optional) and a logo (optional). The logo format must be `.jpg`, `.jpeg` or `.png`.
7. An OS profile is mandatory. See [Adding a New OS Profile](#). However, you can leave the appliance at this point and edit it later.
8. If you have made any modifications, click the checkmark to save.

Note: When you create an appliance, the packages are stored locally in the UForge cache repository. This ensures that the packages will always be available.

Note: Due to a [known missing dependency issue](#) with packages `nss-softokn` and `nss-softokn-freebl` for CentOS and Red Hat Enterprise Linux versions 6.5 and 6.6, if your OS profile includes one of these packages you have to:

- Either add the missing dependency manually
 - Go to Stack > OS Profile
 - Make sure that you have both `nss-softokn-freebl` and `nss-softokn 3.14.3`. If one of the packages is missing, you must add it to the OS profile.
- Or update your OS to version 6.7 (or above)
 - Go to Updates tab
 - Click on 6.7 milestone (or above)
 - Click `simulate` then `update`

Otherwise, image generation will fail with a dependency checking error.

4.2.2 Creating a Windows-based Appliance

To create a Windows Appliance:


1. From the Apps tab, select `create`.
2. Enter the appliance name and version.
3. Choose `Windows` from the OS Distribution drop-down menu.

4. Select the Release and Architecture from the drop-down menus.
5. Click `create`.
6. From the `Stack` page, select the OS Profile from the drop-down list under `Profile name`. Click `save`.

Note: Once you have chosen the OS Profile, you cannot add any packages or run updates. The OS Profile is static. Once created, if you select OS Profile, you will only be able to view the details of the profile you selected.


7. Set the Install profile and click `Save`. For more information, refer to [Updating a Windows-based Install Profile](#).

Note: Unlike Linux, the following cannot be set for Windows appliances: Root User, Users and Groups, Security, Kernel, Keyboards, Licences.



8. Optionally you can add partitions.

- Click on Partitioning and select Advanced Partitioning.
- Click on the + sign at the top right.
- You can modify the name and partition type
- Select the filesystem to `ntfs` and mount point to `D:` (for example).
- Enter the size. The install disk should be 12 Gb for core versions and not less than 32Gb for the full version.
- Check the box in the `Grow` column if you want the partition to be growable.
- Click save.



9. Add software bundles from the Projects or MySoftware pages.

Warning: Software bundles included in `MySoftware` and `Projects` will be put on the image disk but the UForge generation tool WILL NOT install them even if these are executable/installers files (.exe, .msi, etc.). It is up to the end user to manually complete the installation of the software bundles.

For Windows, .exe or .msi files can be given extra parameters. The parameters depend on the .exe or .msi file, and can be used for example for silent installation, providing extra configuration values, etc.

Note: For Windows, with Software bundle whose name begins with `UForgeWinDrivers`, you can specify drivers to be installed in a generated machine image.

Note: A binary called `uforge-install-config` is embedded at generation time, which helps the final user of the Appliance do the last-mile configuration.


Note: A Windows appliance created using a golden image will not list the applications and services under `Stack > OS Profile`.

4.3 Modifying an Appliance Template

You can modify and update appliances that are in your library, including ones that have been imported.

To modify an appliance template:

1. Click on the appliance template to modify.
2. From the `Overview` page you can add or change the logo, modify the name, version or description.
3. On the `Stack` page, you will notice the appliance toolbox on the left-hand side.



The toolbox allows you to define the five key elements of an appliance, namely:

- `Install Profile` – to customize the questions asked when the image is booted for the very first time (or during installation for a physical image format). It also allows you to customize the disk size and partitioning. For more information see [Updating the Install Profile](#).
- `OS Profile` – (mandatory) to choose the operating system packages that are to be used for the appliance. For more information, see [Managing the OS Profile](#). Note that the OS Profile cannot be modified for Windows-based appliances. Refer to [Windows-based Appliance Updates](#).
- `Projects` – to access the UForge Project Catalog. This catalog provides a set of commonly used 3rd party software components when building appliances. The Project Catalog is maintained by the UForge administrator. To add software from the Project Catalog to an application, see [Adding Software from the Project Catalog](#)
- `My Software` – to add to the appliance any software components that you have uploaded. This is also where you can use the Overlay features to manage where the files are installed during generation, if UForge should unzip archives as part of the generation, and set if UForge should on the contrary not install native OS packages. For more information, see [Adding Software from Your Software Library](#)
- `Configuration` – to add boot scripts to configure the appliance after provisioning. For more information, see [Managing Configuration](#)

4.3.1 Adding a Logo

You can add or modify the appliance logo as follows:

1. Click on the appliance to modify.
2. From the `Overview` page click on the square and plus (+) image on the left hand-side.
3. Select the image you want to use as the logo. The format must be .jpg, .jpeg or .png. The image is saved automatically.

4.4 Managing the OS Profile


4.4.1 Adding a New OS Profile

Every appliance must have an OS profile, which contains all the operating system packages for the appliance. UForge allows you to easily create an OS profile from a set of standard profiles. You can then add specific operating system packages.

The goal is to only include the operating system packages you require to run your application. This process is known as JeOS (Just Enough Operating System). By only using the operating system packages you need, not only do you reduce the footprint of the resulting machine image, you also make the machine image easier to maintain (as there are less updates) and (hypothetically) more secure as there will be less unwanted services started.

To add an OS profile to your appliance:

1. Double click on the appliance you want to edit.
2. Go to the `Stack` page.
3. Click on `OS Profile` in the toolbox.
4. From the drop-down menu, select the `OS profile` template you wish to use. The operating system packages are added automatically and your appliance revision number is increased.



4.4.2 Adding Packages to the OS Profile

You may want to add packages that are provided as part of the operating system distribution. To get a list of all the packages that correspond to your search criteria:

1. Select the appliance to modify and go to the `Stack` page.
2. Click on `OS Profile` in the toolbox.

3. Enter your search string. For example, add `php` as a search string to get a list of all the PHP packages provided by CentOS.
4. Click the search icon.

Search for Packages

x | 🔍

☐ Show only 32-bit packages

5. Optionally you can filter the results by selecting `Show only 32-bit packages` (this displays only 32bit packages available) or by entering text in the `filter` box.
6. Select the packages you want to add.
7. Click the `save` button to add the packages to the OS profile.

Note: When you create an appliance, the packages are stored locally in the UForge cache repository. This ensures that the packages will always be available. However, UForge tracks all available updates.

For more information on package updates, see [Tracking OS Package Updates](#). If you want to make sure you always have a specific version of a package, read [Pinning Packages](#).

4.5 Updating the Install Profile

The `Install Profile` on the `Stack` page allows you to customize the questions asked when the image is booted for the very first time (or during installation for a physical image format). The install profile is mandatory. However, the options you set will differ depending on if your appliance is Linux-based or Windows-based. Please refer to the appropriate section below.

You can also refer to the following sub-sections:

4.5.1 Adding Users and Groups

You can add operating system users and groups to a Linux-based appliance Install Profile. These will be integrated to the appliance template.

Warning: The users and groups created in UForge are not linked. If you create a user and list it as part of a specific group, you must then also create the group; otherwise the image generation will fail.

To add a user to an appliance:

1. Select the appliance template you want to modify.
2. From the `Stack` page, click on `Install Profile` in the toolbox.
3. Select `Users and Groups`.
4. Click the plus symbol next to the `Users` table. The `Create User` page will be displayed.

Create user

User *

Full name *

System User ☐ Mark this user as being a system user

User id ☒ Set the user id automatically

1000

Shell * /bin/bash

☐ No console login

Home Directory

Primary Group ☒ Automatically create primary group for user

Secondary Groups Secondary groups for the user (e.g. grp1, grp2, grp3, etc.)

Password

Verify Password

create +

5. Enter a user name.
6. If you want to manually enter the user ID, deselect set the user id automatically and enter the ID number.
7. If you want to manually create the primary group the user is part of, deselect automatically create primary group for user and enter the group name.
8. Click create.

Warning: If you create a user and list it as part of a specific group, you must then also create the group; otherwise the image generation will fail.

To add a group to a Linux-based appliance install profile:

1. Select the appliance template you want to modify.
2. From the Stack page, click on Install Profile in the toolbox.
3. Select Users and Groups.
4. Click the plus symbol next to the Groups table. The Create Group page will be displayed.

Create Group

Name

System Group ☐ Mark this group as being a system group

Group id ☒ Set the group id automatically

1000

create +

5. Enter a group name.
6. Check if you want this group to be a system group.
7. If you want to manually enter the group ID, deselect Set the group id automatically and enter the group ID number.
8. Click create.

4.5.2 Configuring Multiple Network Internet Cards

You can configure multiple network internet cards (multi-NIC) as part of your appliance template in the `Install Profile`. This is only available for Linux-based appliances.

To set multi-NICs for an appliance template:

1. Select the appliance you want to modify.
2. From the `Stack` page, click on `Install Profile` in the toolbox.
3. Select `Network`.
4. Click `create`. By default it is set to `Automatic`.

5. Select `Manual`. The following page will appear.

6. Enter the name and IP information.
7. Click `save`.
8. Optionally you can change the network card order. Click on the symbol on the left of the pencil icon, on the main `Network` page.



4.5.3 Configuring Advanced Partitioning

You can configure advanced partitioning as part of your appliance template in the `Install Profile`. The elements you can configure will depend if your template is Linux or Windows based.

Advanced Partitioning for Linux

The following example assumes that you want to build the following partitions, with a virtual hard drive of 20 GB.


```
part /boot -fstype=ext4 -size=500 -ondisk=sda
part pv.1 -grow -size=1 -ondisk=sda
volgroup ROOTVG -pesize=4096 pv.1
logvol / -fstype=ext4 -name=LogVolROOT -vgname=ROOTVG -size=3072
logvol swap -name=LogVolSWAP -vgname=ROOTVG -size=1024
logvol /usr -fstype=ext4 -name=LogVolUSR -vgname=ROOTVG -size=5120
logvol /var -fstype=ext4 -name=LogVolVAR -vgname=ROOTVG -size=1024
logvol /home -fstype=ext4 -name=LogVolHOME -vgname=ROOTVG -size=5120
logvol /tmp -fstype=ext4 -name=LogVolTMP -vgname=ROOTVG -size=1024
logvol /opt -fstype=ext4 -name=LogVolOPT -vgname=ROOTVG -size=1024
```

To set advanced partitioning for an appliance template:

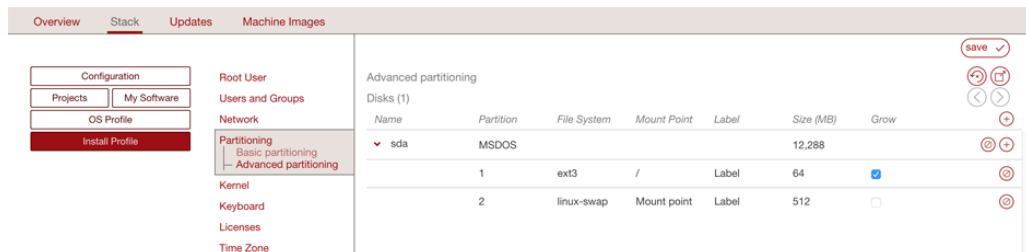
1. Select the appliance you want to modify.
2. From the `Stack` page, click on `Install Profile` in the toolbox.
3. Select `Partitioning`, then `Advanced Partitioning`.

Note: In order to view the information more easily, you can click on the enlarge button in the top right. This opens a separate window where you will see all the advanced partitioning.

4. Advanced partitioning works sequentially, from top to bottom. The three sections offered by the UI, which are `Disks`, `Logical Groups`, and `Logical Volumes` should be filled in order, sequentially:
 - disks with partitions including the total virtual disk size required
 - logical group(s), assigning the associated physical extent (partition)
 - specify the logical volume specification for the created logical group(s)



5. Click on the arrow in Disks next to sda. You will see the default disks.



6. Delete the default `linux swap` partition by clicking the x at the end of the line with `linux-swap`. You must not set the swap size to 0.
7. Click on the size of the MSDOS partition to set it to 20000.
8. Click on the partition 1 information to modify the file system to `ext4` and the mount point to `/boot`.
9. Click on the + sign to create a new partition with type `lvm2` and size set to 18000 MB.
10. Unselect `Grow` and set the size of the `/boot` disk to 500.
11. In the Logical Groups section, click on the + sign and set the name of the logical group. For this example: `ROOTVG`.

Note: Image generation will fail when migrating if the volume group name set in the Partitioning Table is the same as the name of LVM volume group in UForge server.

12. Next to the newly created volume group, click on the + sign to create a new volume extent. A pop-up window will appear proposing a `sda/2` physical extent with size automatically set to 18000 MB. Click `ok`.

The screenshot shows the 'Edit Advanced Partitioning Table' window with three sections: Disks (1), Groups (1), and Volumes(2).

Disks (1)

Name	Partition	File System	Mount Point	Label	Size (MB)	Grow
sda	MSDOS				12,288	
	1	ext3	/	Label	64	☒
	2	linux-swap	Mount point	Label	512	☐
	3	lvm2	Mount point	Label	4,096	☐

Groups (1)

Name	Physical Extend	Size (MB)
vg0		4,096
	Partition: sda/3	4,096

Volumes(2)

Name	File System	Mount Point	Size (MB)	Grow
vg0			4,096	
lvol0	unformatted	Mount point	0	☐
lvol1	unformatted	Mount point	0	☐

13. Create the logical volumes one by one, or create them all at once and then edit the respective specifications. For each logical volume to create, click on the + sign in the Logical Volumes section. For our example, you will need seven logical volumes.


- LogVolROOT ext4 / 3072
- LogVolSWAP linux-swap 1024
- LogVolUSR ext4 /usr 5120
- LogVolVAR ext4 /var 1024
- LogVolHOME ext4 /home 5120
- LogVolTMP ext4 /tmp 1024
- LogVolOPT ext4 /opt 1024

14. Click Save.

Advanced Partitioning for Windows

You can set an advanced partitioning table for a Windows-based appliance template. To set advanced partitioning:

1. Click on Partitioning and select Advanced Partitioning
2. Click on the green + sign at the top.
3. You can modify the name and partitions type
4. Select the filesystem to ntfs and mount point, for example: D : .
5. Enter the size. The install disk should be at least 14 Gb for core versions and 20Gb for full versions
6. Check the box in the Grow column if you want the partition to be growable.
7. Click save.



4.5.4 Updating a Linux-based Install Profile

You can define the following as part of a Linux-based appliance install profile:

- **Root User:** The root user password by default is prompted during the first boot of the machine image i.e. ask during installation. However, you can pre-set a root password. You can enter an SSH key to allow users to login as root. If you select `Disable root password login via SSH`, root will still be able to login from the console.
- **Users and Groups (optional):** you can add operating system users and groups. See [Adding Users and Groups](#) for more information
- **Network:** You can set the internet settings. The default is set automatically. See [Configuring Multiple Network Internet Cards](#)
- **Security:** You can activate or deactivate the firewall present in the filesystem when launching the appliance (regardless of whether the firewall is iptables or other). Firewall is set to `Off` by default. You can also set the SELinux configuration here.
- **Partitioning:** You can modify the disk and swap size for the automatic set up, select ask during install, or set up Advanced Partitioning (for several disks). For more information see [Configuring Advanced Partitioning](#).

Note: If you want to generate a K5 image or an OpenStack image, the appliance partition should contain only 1 disk.


- **Kernel:** You can add kernel parameters by clicking plus, entering data and click save.
- **Keyboard:** default is ask during installation. You can choose set automatically and select the keyboard from the drop down menu.
- **Licenses:** default is accept licenses during installation.
- **Time Zone:** default is set automatically to London.
- **Welcome Message:** You can enter a welcome message.

Note: For basic partitioning disk size, you must ensure that the disk is large enough to store all the binaries and files for the appliance template. For windows based operating systems, it is advised to have a disk size of at least 14GB for core versions, and at least 20GB for full versions.

4.5.5 Updating a Windows-based Install Profile

You can define the following as part of a Windows-based appliance install profile:

- **Sysprep:** Allows you to indicate that sysprep should be run on first boot. If you do not select to run sysprep, you cannot set an administrator password as Administrator option will not be visible.
- **Activation Key:** You can set the Windows Activation Key. If it is not entered in the Install Profile, the default key will be used for a 30-day trial period once the appliance is booted.
- **Administrator:** To set the administrator password. Can be one of Ask during installation, None or Set automatically. Note that this option is visible only if you have selected to run sysprep under Sysprep option.
- **Network:** You can set the internet settings. The default is set automatically. See [Configuring Multiple Network Internet Cards](#)
- **Partitioning:** You can modify the disk and swap size for the automatic set up, select ask during install, or set up Advanced Partitioning (for several disks). For more information see [Advanced Partitioning for Windows](#).
- **Time Zone:** default is set automatically to London.
- **Welcome Message:** You can enter a welcome message.



4.6 Managing Configuration

When you create your appliance, you can include boot scripts to configure the appliance after provisioning or to configure the appliance to communicate with a DevOps platform for further communication steps.

4.6.1 Adding a Boot Script

You can add boot scripts that will be run either the first time the appliance is booted or every time the machine or virtual machine is started. The boot scripts will be run once all the software and appliance packages are installed, prior to launching the machine. The scripts are run in numeric then alphabetical order. The boot scripts will be executed as root.


Note: The firstboots are run before everyboot scripts. Bootscripts named 1_script will be run before a_script, which will be run before script_1, which is run before script_a.

Warning: Only .bat files will be executed for Windows. If you want to upload a Powershell script then you should upload it to My Software and call the execution of the Powershell script from a .bat configuration bootscript.

If you want to install software or packages as part of the installation, you can use My Software to upload overlay files (e.g. /etc/profile.d/xxx.sh). For more information, refer to [Adding Software from Your Software Library](#).


To add a boot script to your appliance:

1. Select the appliance you want to modify.
2. From the Stack page, click on Configuration in the toolbox.
3. Select Boot Scripts.
4. Click on new or upload button on the right hand side.



If you are creating a new boot script:

1. Enter the name.
2. Select the type: If you select first boot, then the boot script will be launched once, the first time the instance is launched. If you select every boot, then the boot script will be launched every time the instance is rebooted.
3. Select the boot order.
4. Enter the contents of the boot script.
5. Click save.



If you are uploading an existing boot script:

1. Select the type: If you select `first boot`, then the boot script will be launched once, the first time the instance is launched. If you select `every boot`, then the boot script will be launched every time the instance is rebooted.
2. Select the boot order.
3. Click `choose` to locate your file and click `open`.
4. Click `save`.

4.7 Adding Custom Software Components


There are two ways to add 3rd party software components or your own software to an appliance template.

- Using the `Project Catalog`. This catalog is public to all the users on UForge and is maintained by the privileged users and administrators.
- If the catalog does not contain the software component you are looking for, you can upload the software into the `My Software` library. This is your own private software library, allowing you to upload any software into UForge and be able to add it to any of your appliance templates.

4.7.1 Adding Software from the Project Catalog

The following describes how to add MySQL and Apache to an appliance.

1. Select the appliance to modify and go to the `Stack` page.
2. Click on the `Projects` in the toolbox. This displays the Project Catalog. The top table lists which projects have already been added to the appliance.



3. Add the projects to your appliance either by clicking their `add` button or by selecting them and clicking the `add selected` button. You can scroll through the available projects or enter a search string to filter the list. The changes are automatically saved.

Overview

Stack

Updates

Machine Images

Summary

Configuration

Projects

My Software

OS Profile

Install Profile

Attached Projects

Filter

x

Q

Project	Category	Version	Size	Tag	License	Format dependent
Apache HTTP server	Web/Application Servers	2.4.6	9 MB	Distribution	Apache License	remove
MySQL Server	Database	1.0	15 KB	Distribution	GPLv2	remove

Available Projects

+ add selected

Filter

x

Q

Project	Category	Version	Size	Tag	License	Format dependent
☒ Asterisk	Communications	1.0	15 KB	Distribution	GPLv2	+ add
☐ Bugzilla	Software Development	1.0	25 KB	Distribution	MPL	+ add
☒ ClamAV	Security	1.0	15 KB	Distribution	GPLv2	+ add
☐ Drupal	Content Management	1.0	15 KB	Distribution	GPLv2	+ add
☐ Eclipse	Software Development	1.0	16 KB	Distribution	EPL	+ add
☐ Exim	Communications	1.0	15 KB	Distribution	GPLv2	+ add
☐ Jetty	Web/Application Servers	1.0	10 KB	Distribution	Apache License	+ add

4.7.2 Adding Software from Your Software Library

You can add your own custom software either using the `Software Library` or by including boot scripts. My Software overlay files (for example `/etc/profile.d/xxx.sh`) will be run before bootscripts when the machine is booted.

The following is a list of supported file formats:

- **Linux only**, note by default `.rpm` and `.deb` files will be installed at generation. This can be modified when you upload the files (see procedure below):
 - `.rpm` (“`rpm i`” will be executed)
 - `.deb` (“`dpkg -i`” will be executed)
- **Windows only:**
 - `.msi`
 - `.exe`
- **Linux and Windows compatible:**
 - `.tar.bz2`
 - `.bz2`
 - `.tar.gz`
 - `.tgz.bz2`
 - `.tgz`
 - `.gz`
 - `.tar`
 - `.zip`
 - `.tar.zip`


To add custom software components to an appliance:

1. From the `Apps` tab, click on `Software Library` in the top left hand side of the UI. This opens your private `Software Library`.
2. Click on the `add software` in the top right hand side.

- You are now prompted for the name, version and maintainer of the software component you would like to add. You can also set a tag and category. Click `create` to complete.

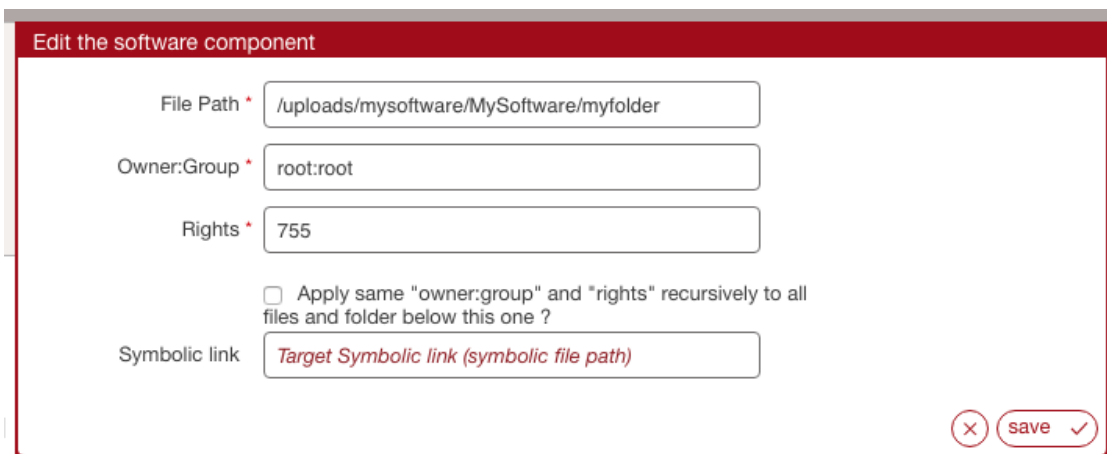
Note: If you want to create specific software bundle with the VirtIO drivers which will allow you to publish Windows images to Red Hat, then the name must begin with `UForgeWinDrivers`.

- The software Overview page will now open. You can modify the name and version, and add a description.
- To upload the files, go to the Files page.



Note: If you want to group a set of files, you can create a folder by clicking `create folder`. Then enter a folder name and click `create folder`. Now if you want to put files in this folder, click on the upload icon next to the folder name.

If you create a folder, you can then click on the pencil icon to edit the file properties. Here you can select to apply the same owner and rights to all the files added to the folder.



Warning: Folders are not supported if you are uploading VirtIO drivers for publishing Windows images to Red Hat.

You can now upload the files as follows:

1. Click `upload` to select the files you want to add and click open.
2. Click `fetch` to set an archive location where the files should be retrieved. When using the `fetch` option you can indicate a remote URL or artifact account (for information on artifact accounts, refer to [Managing Your Artifact Accounts](#)). You can also select to have the files uploaded to the UForge by checking `use cache`.

Warning: If you are uploading VirtIO drivers for publishing Windows images to Red Hat, archive files such as zip or tar are not supported.

6. Optionally, you can modify the default `install path` that will be used. If you have modified the name of your appliance, it may not be reflected in the install path.
7. By default `.deb` and `.rpm` files will be installed at generation time. Optionally you can edit the settings for those files
 - Click on the pencil on the right hand side of your uploaded file.
 - Un-select `Install this native package as part of the generation`.
 - Click the check mark to save your changes.

In this case, the `.deb` or `.rpm` archive file will be in the directory but will not be installed at generation time.


8. Optionally, you can select to unzip archives as part of the generation. To do so:

- Click on the pencil on the right hand side of your uploaded file.

- You can then edit the install path and select if it should be extracted.
- Click the check mark to save your changes.



- Optionally, you can restrict the distribution or target format that the software applies to, from the [Restrictions](#) page. The restriction rule is set as a regular expression. For more information, refer to [Restricting Software Components for OSes and Formats](#).
- From the [Repository Packages](#) page, you can set the packages with which your software is compatible. This page will only be visible if the restriction rule matches only one distribution. You can search for packages. Select and click the down arrow to add them. Click [save](#).
- From the [Boot Scripts](#) page, upload any boot scripts you want to add to this software. You can either:**
 - Upload an existing boot script file by using the [upload](#) icon.
 - Create a new boot script by clicking [new](#).



In both cases, you must select the type. If you select `first boot`, then the boot script will be launched once, the first time the instance is launched. If you select `every boot`, then the boot script will be launched every time the instance is rebooted. You must also set the boot order.

- From the [Licenses](#) page, select a standard licence, or select [Custom](#) and then click [upload](#) to upload any licenses you want to add to this software. Changes will be automatically saved.
- Add the uploaded software component to the appliance. Click on the [Appliance Library](#) to view your Appliance Library. Double-click on the appliance template you want to add the software to.
- Go to the [Stack](#) page and click on the [My Software](#) button in the toolbox.

15. Select the software components you want to add.

Note: You can create a Docker image directly from a software using the UForge UI. Refer to *Generating a Docker Image from Software Components*.

4.7.3 Restricting Software Components for OSES and Formats

Under *My Software* you can restrict the usage of a software bundle based on a distribution name, family, architecture or for a specific machine image format.


To set restriction rules:

1. Under the Apps go to the Software Library page.
2. Select the software you want to modify.
3. **Go to the Restrictions tab. Enter the restriction rule. The restriction rule is represented by a logical expression with f**
 - object is either `Distribution` or `TargetFormat`
 - for `Distribution` field is `family`, `pkgType`, `name`, `version` or `arch`. The version must be a major version.
 - for `TargetFormat` field is `name` or `type`
 - value is the value you want to match with the fields. For example, `CentOS` for `Distribution` name, `linux` for `Distribution` family, `x86_64` for `Distribution` arch, `VirtualBox` for `TargetFormat` name, `cloud` for `TargetFormat` type.
 - logical operator is `||` for OR and `&&` for AND
 - carriage return is not authorized

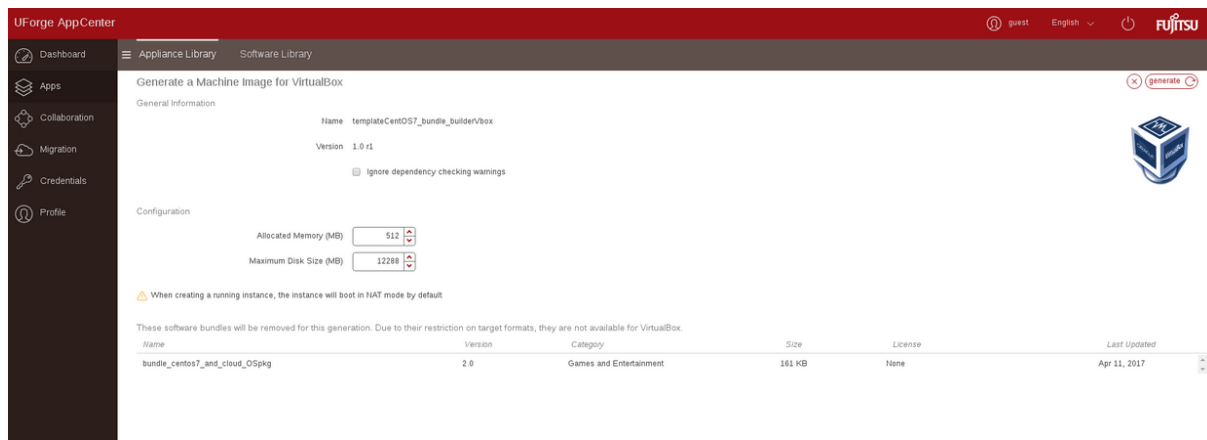
For example, if the software bundle is designed only for distributions CentOS 7 x86_64 or Debian 8 x86_64, or for image format with type virtual, then you would note the Restriction rule as follows:

```
(Distribution#arch=x86_64 && ((Distribution#name=CentOS && Distribution#version=7) ||
↪ (Distribution#name=Debian && Distribution#version=8))) || TargetFormat#type=virtual
```

Once you have set a restriction rule, you will see a cube next to the software component, under the column `Format dependent`.



Warning: If your software bundle is limited to a certain target format and you generate an image in another format, your appliance will be generated but the software bundle will not be part of the final image. A note indicating this will appear when you select to generate the machine image, as in the following image.



Restriction Rule Values

By using the command-line tool `hammr`, you are able to retrieve available distributions:

```
$ hammr os list --url https://uforge.example.com/api -u username -p password
```

Where **Distribution** fields possible values are:

- `name`: the value listed in the *Name* column
- `version`: the value listed in the *Version* column
- `arch`: the value listed in the *Architecture* column
- `family`: can be one of the following `linux`, `unix` or `windows`
- `pkgType`: can be either `application/x-rpm` or `application/x-debian-package`

By using the command-line tool `hammr`, you are able to retrieve available machine image formats:

```
$ hammr format list --url https://uforge.example.com/api -u username -p password
```

Where **TargetFormat** fields possible values are:

- `name`: the value listed in the *Builder Type* column

- `type`: can be one of the following `cloud`, `container`, `physical` or `virtual`

4.8 Tracking OS Package Updates

All the OS packages added to the `OS Profile` section of the appliance templates are tracked for any updates by UForge AppCenter. Based on a timestamp stored in the appliance template, UForge AppCenter can detect any OS package updates that are available. Updates are displayed in the user interface for each appliance template.

If you are in grid view:



If you are in table view, it is listed in the `Updates` column:

≡

Appliance Library

Software Library

create +

Created Appliances

Imported Appliances

Shared Appliances

Appliances (4)

Filter

x

🔍

☐	Name	Version	Revision	OS	Last Updated	Updates
☐	CentOS Basic	 1.0	5	CentOS 7 x86_64	Aug 1, 2016	6 ↺ ↻
☐	nginx	 1.0	1	CentOS 7.1 x86_64	Aug 1, 2016	↺ ↻

You can then easily update the packages using the UForge GUI. Using this tool, you can also roll-back to previous versions of OS packages.

Note: This feature is only available for appliance templates using Linux operating systems.

To update the OS packages:

1. Select your appliance from your appliance library under the `Apps` tab.
2. Go to the `Updates` page.
3. You can see from the graphic when updates were made and how many are available.



4. You can select current updates (if any) or select to return to a previous version of the OS by moving the cursor on the graph.
5. Click `simulate` to see the changes that will be applied. The changes will be listed at the bottom of the screen. Scroll to view the results.
6. If you want to apply the changes listed, then click `update`.

Note: The triangle indicates a Milestone. For CentOS, this is the versions (6.1, 6.2 etc). Milestones are customized by the UForge administrator.

4.8.1 Windows-based Appliance Updates

For Windows-based appliances UForge will indicate the number of updates available, however you cannot use this procedure to update the packages for an existing Windows appliance.

In order to benefit from a newer version of Windows, you will have to:

1. Create a new Golden Image, either from scan or manually. See your administrator.
2. Create a new appliance.
3. You can re-use the MySoftware components contained in the current appliance.
4. You can download from the current template the boot scripts and save them on your local hard drive. You can then upload them to the new appliance.
5. You must re-produce the configuration (Install Profile, Configuration).

4.8.2 Pinning Packages

UForge allows you to “pin” certain packages in your appliance (previously referred to a “sticky” package). This means that during image generation, the package version you have pinned is chosen regardless of the current appliance template timestamp for calculating package versions. All the package dependencies of this package are also calculated.

To pin a specific package:

1. Select the appliance you want to modify.
2. Go to the Stack page.
3. From the OS profile, click on the pinned bar in the right hand side of the package info. In the following image, the first package has been pinned.

Overview

Stack

Updates

Machine Images

Summary

Configuration

Projects

My Software

OS Profile

Install Profile

OS Profile Packages (34)


delete os profile

Search for package(s) in distribution

Filter

☐	Package name	Version	Release	Architecture	Repository	Pinned
☐	bash	4.2.46	20.el7_2	x86_64	CentOS 7.2.1511 x86_64 updates	
☐	bash-completion	2.1	6.el7	noarch	CentOS 7.0.1406 x86_64 os	
☐	bzip2	1.0.6	13.el7	x86_64	CentOS 7.2.1511 x86_64 os	
☐	centos-release	7	2.1511.el7.centos.2.10	x86_64	CentOS 7.2.1511 x86_64 os	
☐	chkconfig	1.3.61	5.el7_2.1	x86_64	CentOS 7.2.1511 x86_64 updates	

4. A pop-up window will list all of the versions of the package available, allowing you to select the version you want to pin.




5. Choose the version of the package you want to pin, then click `save`.

4.9 Listing Appliance Templates

To view the appliance templates you have access, go to the `Apps` tab in the user interface. By default the appliance templates are listed by latest first.

Under the `Appliance Library`, the appliances are organized as follows:

- `Created Appliances` are appliance templates you created
- `Imported Appliances` have been imported from a `Collaboration workspace`
- `Shared Appliances` are appliance templates that you have shared to a `Collaboration workspace`



You can view the details of an appliance template, generate images from them and comment on them.


4.9.1 Searching Appliance Templates

To find a particular appliance template you can:

1. Go to the `Apps` tab on the left sidebar.
2. Use the search engine. The search runs on the appliance name or the OS name. To use the search engine, enter the text in the filter field.



- If you are in the grid view you can sort the appliances. From the drop-down menu select how you want to sort the appliances in your library.



4.10 Cloning an Appliance Template

Once you have created an appliance, you can clone it. This clones all the meta-data of the original appliance template.

Note: Your clone will not include any machine image information that was generated or published for the original.

You can clone an appliance template as follows:

- Select the appliance you want to clone.
- Click `clone` in the top right hand toolbar.



- Enter the appliance name and version of the clone.
- Click `clone`.

4.11 Importing and Exporting Templates

You can import and export appliance templates. When exporting, an archive is created of the appliance template. This archive includes a meta-data file describing the appliance template (based on the [hammr](#) specification) as well as any bundled software that was initially uploaded as part of the template creation.


Likewise, an archive can be imported to the UForge platform, creating a new appliance template in your Appliance Library under the Imported Appliances section.

Warning: If you want to export an appliance template with software bundles in UForge 3.7 and import it in UForge 3.8 or more, you will need to modify your template to make it compatible. Refer to [Updating a 3.7 Appliance Template Before Exporting For 3.8](#).


4.11.1 Exporting

To export an existing appliance:


1. Go to your Appliance Library.
2. Click on the `export` icon on the right hand side of the appliance template in question to export.




3. A pop-up will appear, allowing you to select the format you would like (yaml or json).



4. Click export. This will start the export process.



5. Once the export is complete, you will be prompted to download the archive file.



The equivalent export feature is available when editing an appliance template.

Note: An exported Windows-based template does not contain the “Activation Key” to ensure this data remains confidential.

4.11.2 Importing

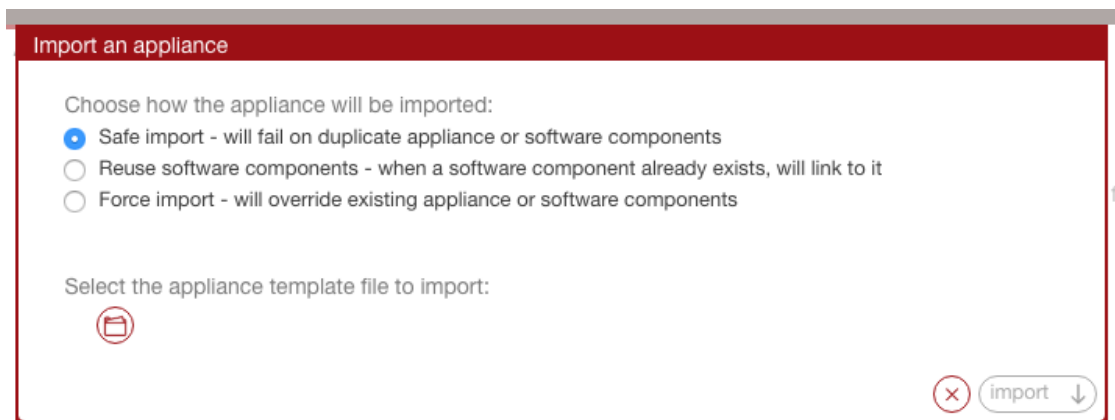
Warning: You will only be able to import a Windows appliance to a UForge that has the same golden image as the one used by the appliance. If the golden does not exist on the target UForge you must export the appliance without OS Profile or remove some OS section fields in template file. Refer to [Updating a Windows Appliance Before Import](#).

To import an archive:


1. Go to your `Appliance Library`.
2. Click on the `import` button at the top right hand side of the view.



3. A number of options are proposed. Select the method of import and click the folder icon to select the archive to import.



4. Click `import`. This will start the import process. The import progress will be shown.



5. Once the import is complete, a new appliance template can be found in the `Imported Appliances` sub-section of your `Appliance Library`.

4.11.3 Updating a Windows Appliance Before Import

When you import a Windows appliance, UForge checks if the golden image used to create the appliance exists on the UForge server. If it does not exist, the import will fail. In order to import the template, you must export the appliance and remove the OS profile.

For example, in order to remove some OS section fields before importing your Windows appliance:

1. Go to your `Appliance Library`.
2. Click on the `export` icon on the right hand side of the appliance template to export.
3. Once the export is complete, you will be prompted to download the archive file.
4. Open the archive file. The OS section should look something like:

```
os:
  name: "Windows"
  version: "Server2012R2"
  arch: "x86_64"
  profile: "Win2K12R2 scan Scan #1"
  windowsEdition: "standard"
  windowsType: "full"
  windowsLanguage: "English"
```

5. Update the OS section to remove the lines `profile`, `windowsEdition`, `windowsType`, `windowsLanguage`. For example:

```
os:
  name: "Windows"
  version: "Server2012R2"
  arch: "x86_64"
```

6. Save the file.
7. Go to your `Appliance Library`.
8. Click on the `import` button at the top right hand side of the view.
9. A number of options are proposed. Select the method of import and click the folder icon to select the archive to import.
10. Click `import`. This will start the import process.

You can also remove the OS Profile from the UI by going to the `Stack` page of the appliance before exporting.

4.11.4 Updating a 3.7 Appliance Template Before Exporting For 3.8

When you export an appliance in UForge 3.7 that contains software bundles, and import it to UForge 3.8, you will encounter compatibility issues. In order to import the template, you must modify the file containing the description of your appliance (`template.yml` or `template.json`).

In this file, there is a section `bundles`, that contains a subsection `oses`. You need to remove the entire content of the subsection `oses`, and replace it by a section `restrictionRule`. Refer to [Restricting Software Components for OSES and Formats](#) for detailed documentation about restriction rules.

For example, if your `oses` subsection is like this:

```
- name: "CentOS"
  version: "7"
  arch: "x86_64"
- name: "Debian"
  version: "8"
  arch: "x86_64"
```

Then you need to write this restriction rule instead:

```
(Distribution#name=CentOS && Distribution#version=7 && Distribution#arch=x86_64) ||
↪ (Distribution#name=Debian && Distribution#version=8 && Distribution#arch=x86_64)
```

When you export a Windows appliance in UForge 3.7, you will also encounter compatibility issues. In order to import the template, you must modify the file containing the description of your appliance (`template.yml` or `template.json`).

In this file, there is a section `os` that contains various information and in particular the profile name. The new 3.8 template now has the following additional fields: `windowsEdition`, `windowsType` and `windowsLanguage`. You will need to add these fields to be able to import the template.

For example, if your `os` section is like this for 3.7:

```
name: "Windows"
version: "Server2012R2"
arch: "x86_64"
updateTo: "12-19-2016 10:32:07 GMT"
profile: "Standard Full English"
```

Then you should update it with the new fields, and it should look something like:

```
name: "Windows"
version: "Server2012R2"
arch: "x86_64"
updateTo: "12-19-2016 10:32:07 GMT"
profile: "Standard Full English"
windowsEdition: "standard"
windowsType: "full"
windowsLanguage: "English"
```

The profile should exist on the target UForge platform, otherwise you must delete the `profile` field and import the template like this:

```
name: "Windows"
version: "Server2012R2"
arch: "x86_64"
updateTo: "12-19-2016 10:32:07 GMT"
```

Thus, your appliance will be imported and you will be able to setup the OSProfile manually in the UI.

Generating and Publishing Machine Images

You can generate a machine image from an existing appliance template in your library. Once the machine image is generated, it can be published to a cloud platform. The following sub-sections go into detail on how to generate and publish a machine image:

5.1 Supported Machine Image Types

With UForge you can create machine images in the following formats.

Physical Format	Additional information
ISO	not supported for a Windows-based appliance
PXE	Only for CentOS systems

Virtual	Additional information
Hyper-V	not supported for Hyper-V on Windows 2016
KVM	none
OVF	Supported hardware types: 4 (ESXi >=3.x) 7 (ESXi >=4.x) 9 (ESXi >=5.x) 11 (ESXi>=6.x)
QCOW2	none
Raw	none
tar.gz	none
Vagrant Base Box	none
VHD	none
Virtual Box	none
VMware vCenter	Supported hardware types: 4 (ESXi >=3.x) 7 (ESXi >=4.x) 9 (ESXi >=5.x) 11 (ESXi>=6.x) For Windows, ensure your HW version is the most up to date for the ESXi host
VMware Server	Supported hardware types: 4 (ESXi >=3.x) 7 (ESXi >=4.x) 9 (ESXi >=5.x) 11 (ESXi>=6.x)
Vagrant	none
Xen	none
Citrix XenServer	none

Container	Additional information
Docker	not supported for a Windows-based appliance
LXC	not supported for a Windows-based appliance

Cloud	Additional information
Amazon AWS	none
Microsoft Azure	Ubuntu <= 12.04 not supported
Cloudstack	Target formats: CloudStack VMWare (OVA) CloudStack Citrix Xen (VHD) CloudStack KVM (QCOW2)
Fujitsu K5	32 bit OSes are not supported For Windows, Evaluation Editions are not supported
Google Compute Engine	none
OpenShift	not supported for a Windows-based appliance
OpenStack	Target formats: OpenStack QCOW2 OpenStack VMDK OpenStack VDI OpenStack VHD
Oracle Cloud	- For OS versions supported by Oracle Cloud, refer to FAQ at Oracle Cloud - Debian 8 is not supported by UForge for publication to Oracle Cloud. - Supported subscriptions are Metered Service Offerings. Nonmetered Service Offerings are not supported by UForge.
Outscale	none
SUSE Cloud	none
VMware vCloud Director	Supported hardware types: 4 (ESXi >=3.x) 7 (ESXi >=4.x) 9 (ESXi >=5.x)

5.2 Generating a Machine Image

Once an appliance template has been created, you can generate a machine image that packages the stack to run on a particular virtual, cloud, container or physical environment. For a complete list of supported machine image formats, refer to *Supported Machine Image Types*.

Note: If you want to generate a K5 image or an OpenStack image, the appliance partition should contain only 1 disk.

Note: If generating CentOS 6.5/6.6 image, refer to *the following note* if your OS profile includes either of the following packages: *nss-softoken* or *nss-softoken-freebl*.

To generate a machine image:

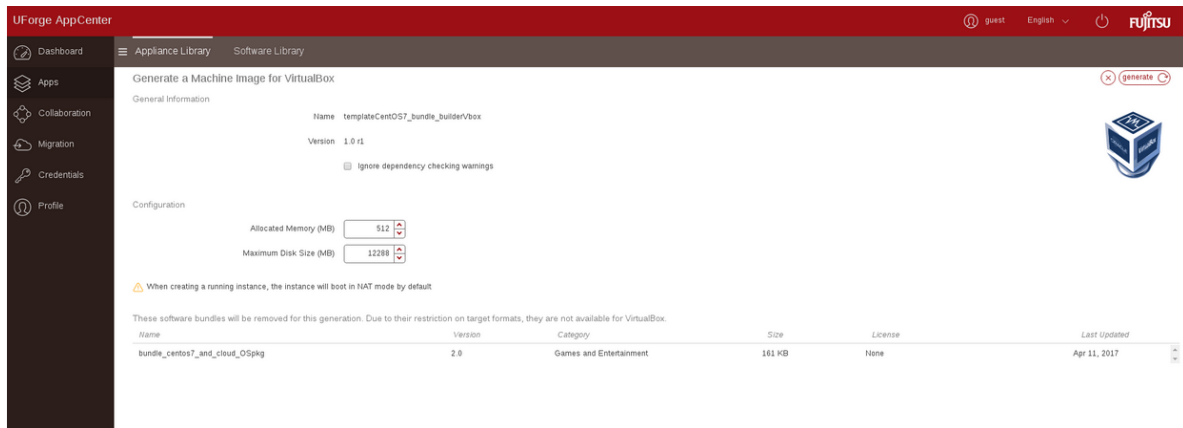
1. Go to the Apps tab.
2. Select the appliance from the Appliance Library.
3. Click on the `generate` button at the top right to display all possible image formats which can be generated. The formats are organized by type: Cloud, Container, Virtual, Physical.
4. Choose the image format you would like to generate. For a Docker image, refer to *Generating a Docker Image from an Appliance*.

Note: If you are generating an image for VMware vCenter make sure that your generated image's hardware version is the most up to date for the ESXi host that you are using for publish. For more information please refer to *VMWare recommendations*

5. You will see a recap of the image you are about to generate.

Note: If you want to ignore dependency checking during image generation (for example because you have knowingly removed a package dependency that is not required in your environment), then you check the option `ignore dependency checking warnings`. Note that this is an advanced option and should not be used systematically.

Warning: If your software bundle is limited to a certain target format and you generate an image in another format, your appliance will be generated but the software bundle will not be part of the final image. A note indicating this will appear when you select to generate the machine image, as in the following image.



6. You can set the disk size, then click the `generate` button to launch a generation in UForge for this appliance template.

Note: Depending on the packages you install and the size of your software, make sure that the disk size is large enough to accommodate the software to be installed. For Windows-based operating systems, it is advised to have a disk size of 32GB, as recommended by Microsoft. Moreover, depending on the OS profile used, and Project or MySoftware components attached to the template, the disk size may have to be larger.

The generation will take a few minutes to complete (depending on the number of packages in the appliance template and the disk size chosen). The generation progress is shown.

Note: Some services are disabled or enabled depending on the target machine image being created (refer to [Service State](#)).

Once the generation is complete, you can download the image locally, or for certain cloud formats register the machine image directly to the target environment using your cloud credentials.

- For PXE, you cannot download the image but can launch it through an iPXE shell. For more information, refer to [Launching a PXE Image](#).
- For KVM, `launch.sh` script is included in the archive and you can use it to start the instance. However, before using it you need to modify it according to your environment.

You will note that a package `uforge-install-config` is injected in the generated image. This file is responsible for:

1. launching the dialog for the install profile configurations which are not automatic (keyboard, root password, licenses, time zone, static IP)
2. executing the installation bootscript of the template

If the template is configured to be fully automatic in the `Install Profile` and has no bootscript that is supposed to run at every boot, the `uforge-install-config` package and associated `/etc/UShareSoft` directory can be removed safely.

However, it is preferable to leave this file.

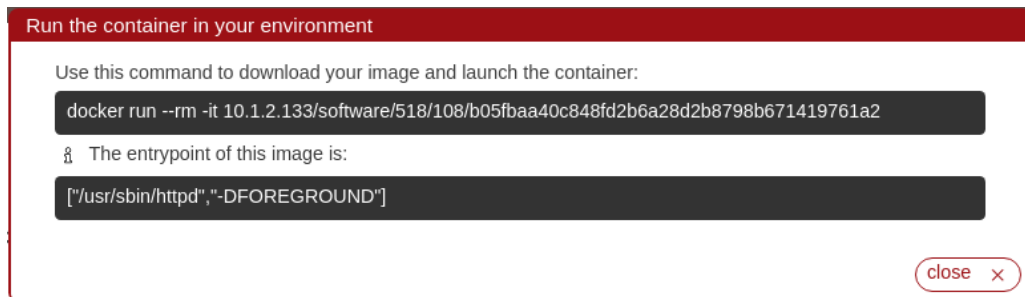
5.2.1 Generating a Docker Image from an Appliance

Note: You cannot generate a Docker image from a Windows template or a Windows scan.

Note: If you generate a Docker image from an appliance that does not have an install profile set as automatic, then you will not be able to launch the Docker image.

To generate a machine image from an appliance:

1. Go to the Apps tab.
2. Select the appliance from the Appliance Library.
3. Click on the generate button at the top right to display all possible image formats which can be generated.
4. Choose Container, then Docker image format.
5. Enter the Docker entrypoint in exec form. For example `["/usr/sbin/httpd", "-DFOREGROUND"]`.
6. Click the generate button to launch a generation in UForge for this appliance template.
7. Once the generation is complete, the Docker image will be listed in your library under Machine Images.
8. You can (optionally) click on download. A pop-up similar to the following will be displayed. If you copy and run the command in the pop-up, this will download, import and run the Docker image from your Docker client.



5.2.2 Generating a Docker Image from Software Components

Note: You cannot generate a Docker image if any of the restriction rules that have been set for the software exclude Docker publication or generation.


Note: When generating a Docker image from software components, a reduced set of packages will be included. For example, the kernel will not be included. This will result in a lighter-weight image for the following OSes:

- CentOS 7
- Red Hat Enterprise Linux 7
- Debian 9
- Ubuntu 18.04

- SUSE Enterprise Linux Server 12
 - openSUSE 42
-

To generate a machine image from a software component:

1. Go to the `Apps` tab.
2. Select the software from the `Software Library`.
3. Click on the `generate` button at the top right.



4. Select the distribution and click `generate`.
5. The possible image formats which can be generated are displayed. Select the format you want to use.
6. Enter the `Docker` entrypoint in `exec` form. For example `["/usr/sbin/httpd", "-DFOREGROUND"]`.
7. Click the `generate` button to launch a generation in UForge for this software component.
8. Once the generation is complete, the `Docker` image will be listed on the `Dashboard` under `Generations`. You can now view the packages, publish or download your image.

5.2.3 Generating an OpenShift Image

Note: You cannot generate an OpenShift image from a Windows template or a Windows scan.

To generate a machine image:

1. Go to the `Apps` tab.
2. Select the appliance from the `Appliance Library`.
3. Click on the `generate` button at the top right to display all possible image formats which can be generated.
4. Choose `Cloud`, then `OpenShift` image format.
5. Fill the entrypoint in `exec` form. For example `["/usr/sbin/httpd", "-DFOREGROUND"]`.
6. Click the `generate` button to launch a generation in UForge for this appliance template.

7. Once the generation is complete, the OpenShift image will be listed in your library under Machine Images.

Note: It is not possible to publish from a Docker image to OpenShift.

5.3 Publishing a Machine Image

Note: When publishing an image, you have to comply with the license agreement of OSes and software which UForge handles, in particular:

- **Publishing OS image of RHEL subscription to public cloud** Cloud provider has to be CCSP (Certified Cloud & Service Provider) and be registered to Red Hat Cloud Access. For more details, please confirm with cloud provider.
- **Publishing Windows OS image** You must acquire Windows licenses in order to handle Windows OSes in UForge and confirm usage conditions of cloud provider and virtualization software which you publish to.


In order to publish a machine image to a cloud environment or container, you must already have credentials to access that environment.

Note: If you want to publish a Windows image to KVM on Red Hat Linux you need to inject specific VirtIO drivers. See the specific procedure documented in [Publishing a Windows Image to KVM on Red Hat Linux](#).

Note: To boot up a Windows image on KVM, you should specify 'VirtIO' as a disk bus type. You should not specify 'IDE' or other types.

Note: Before publishing a Windows image to AWS, `vmimport` role must be properly assigned to an AWS account. An access right to S3 bucket is also required for an account. Refer to the procedure in [VM Import Service Role section of the AWS documentation](#). Because the name of a S3 bucket changes for each image publish, you should give access to any buckets using a wildcard *. For an example, refer to :ref: `windows-aws-s3`.

1. If not already done, create an account for the target environment. For more information, see [Managing Cloud Accounts](#).
2. Go to the appliance and click the Machine Images page. If you have not generated a machine image, you will need to do so as described in [Generating a Machine Image](#).
3. Click on the arrow pointing to the right to publish your image.



- Following the instructions in the tooltips, choose the account and fill in any additional information required.

Note: When publishing to Microsoft Azure you should note the following restrictions:

- The storage account must exist in your Microsoft account
- The container is mandatory. If the value you set does not match an existing container it will be created with the new name you have set.
- The blob and cloud image name must be unique. If they already exist, this will overwrite an existing one with the same name.
- The resource group is optional. If you enter a value here, it must already exist in your Microsoft account.

Overview

Stack

Updates

Machine Images

Summary

Publish and Register Machine Image to Microsoft Azure

ⓧ publish →

Account *

UShareSoft Azure Resource Manager

+

Storage Account *

Microsoft Azure storage account

Container *

Microsoft Azure container

Blob *

ldp-219-guest-CentOS7.vhd

Cloud Image Name *

ldp-219-guest-CentOS7

Resource Group

Microsoft Azure resource group

-
- Click publish.

Note: Publishing an image to Amazon and Outscale will be billed to the user account. Trial Amazon accounts are not supported for publishing images from UForge. Only full Amazon accounts can be used.

Note: Currently, publications to Docker are only supported for Docker Registry v2.

Note: Currently, publishing a Windows Server 2008 R2 machine image to Azure Stack is not supported.

- The publication will take a few minutes to complete (depending on the size of the image and the network connectivity between UForge and the target environment). The publication progress is shown. At the end of the publication, the machine image has been published by UForge to your target environment. The published image can be found in the target cloud environment or container.

Note: If your publication to AWS or Outscale fails with a message `Unable to connect to proxy`, this may be due to your security settings on your AWS or Outscale account. Your default security group in your AWS or Outscale account must allow SSH from UForge AppCenter compute nodes. You may need to contact your administrator to obtain the IP addresses of the compute nodes in the case of a multi-node deployment of AppCenter.

UForge does not launch instances in the target cloud environment. If you wish to launch an instance from this machine image, you should go to your target cloud environment console for further actions.

Note: When you publish an image to Amazon or OpenStack, and you have the correct UForge entitlements, you will be able to launch your machine image directly from UForge using the `Deploy` option, without having to connect to your cloud platform. Refer to [Deploying a Machine Image](#).

Note: When you publish an image to VMware, the result will be a template and not a virtual machine.

Note: When you publish an image to Azure Stack, the result will be a VHD file on the specified BLOB. Virtual Machines can then be created from these VHD files either manually or by using scripts.

Warning: If your Fujitsu K5 publish failed, there may be data published to cloud, incurring costs, even if not visible on your cloud account. You should run a cleanup manually. Refer to Chapter 2.7 Object storage of the [FUJITSU Cloud Service K5 IaaS API Reference \(Foundation Service\) guide](#)

5.3.1 Publishing a Windows Image to KVM on Red Hat Linux

If you want to publish a Windows image to KVM on Red Hat Linux you need to inject specific VirtIO drivers.

1. In order to add the specific drivers, the Red Hat VirtIO drivers can be either added to AppCenter as a Project if it has been added to your AppCenter by the administrator, or using MySoftware.

Note: If you create a specific VirtIO driver (refer to [Adding Software from Your Software Library](#) for more information), the software bundle MUST have a prefixed name `UForgeWinDrivers`.

The following files should be included in your software bundle:

- viostor.sys
- viostor.inf
- viostor.cat
- vioscsi.sys
- vioscsi.inf
- vioscsi.cat
- netkvm.sys
- netkvm.inf
- netkvm.cat

1. When creating your appliance template, add the Project or MySoftware bundle that contains the VirtIO drivers.
2. Generate the machine image.
3. Publish the machine image.

5.3.2 Publishing a Windows Image to Azure Using Migration

Note: In order to install Azure Virtual Machine Agent, .NET Framework 4.0 or later is required and should be installed in the source machine or in the golden image in advance.

When publishing a machine image to Azure through the migration workflow, there are some prior steps to be performed. Please refer to Microsoft's guide on [How to prepare a VHD image for upload](#), more specifically the following sections:

- Set Windows configurations for Azure
- Check the Windows services
- Update Remote Desktop registry settings
- Configure Windows Firewall rules
- Verify VM is healthy, secure, and accessible with RDP
- Install Windows Updates

When publishing a Windows machine image to Azure, depending on the generation method of the machine image, the published image can be listed on Azure portal under either `Images` or `Disks`, according to the cases listed below:

1. The published image will be listed under `Disks` when:

- performing a blackbox migration of a Windows machine;
- performing a whitebox migration of a Windows machine with the appliance configured as to not run sysprep automatically on its first boot.

Note: In these cases, the Azure Virtual Machine Agent should be installed manually before the migration process.

2. The published image will be listed under `Images` when:

- performing a whitebox migration of a Windows machine with the appliance configured to run sysprep automatically on its first boot;

Note: If `Run Sysprep` is enabled in the appliance configuration, any password settings defined under the `Administrator` section of the appliance will not be reflected in the published image. Password settings should be configured on the Azure portal instead.

- the Windows machine image was generated using a manually created Golden Image.

Warning: When publishing to Azure an image generated using a manually created Golden Image, the `Install Profile` default settings of the appliance should not be modified. If you need to change the disk size, you can change it on the image generation page.

Note: In these cases, the Azure Virtual Machine Agent will be installed automatically as part of Azure's internal processing to register the machine image.

Deploying a Machine Image

UForge allows you to automatically deploy a published image to a cloud account previously set up in your UForge credentials. The following sub-sections go into detail on how to deploy a machine image:

6.1 Deploying a Machine Image

Note: Currently, Linux machine images published on Amazon, OpenStack, Cloudstack and Microsoft Azure can be deployed from UForge. You can deploy a Windows based appliance only to Amazon and Azure. Deploy for Azure Stack is not supported.

You can only deploy a machine image if you have AMP installed. For more information on installing AMP, refer to the official [Cloudsoft AMP documentation](#)

The credentials to deploy a machine image are those used to publish it. For Amazon, only full accounts can be used, trial accounts are not supported for publishing images from UForge.

6.1.1 Pre-requisites

Check the following before starting your deployment.

For Windows:

- Check that there are no pending Microsoft updates when you scan the source system for deployment.
- For Windows Server 2008R2, you need to upgrade WinRM to 3.0 on the source machine before scanning.
- For Windows deployments to Azure, do the following on the source machine before scanning:
 - Edit Windows Firewall rules as follows:
 1. Enable the Windows Remote Management (HTTP-In) rule.

2. Define `local` address as “Any”.

– Launch command prompt as administrator and run the following WinRM commands:


```
winrm qc
winrm set winrm/config/service/auth @{Basic="true"}
winrm set winrm/config/service @{AllowUnencrypted="true"}
```

- For Amazon, a Windows image with `Run Sysprep` enabled in its template configuration is not supported for deployment. You should uncheck `Run Sysprep` option in `Install Profile`, when you generate the image to deploy.
- For Azure, a Windows image with `Run Sysprep` disabled in its template configuration is not supported for deployment. You should check `Run Sysprep` option in `Install Profile`, when you generate the image to deploy.

Note: For all deployed Windows virtual machines, if WinRM service is not enabled by default, then it will be activated during the deployment process.


6.1.2 Deployment Procedure

1. If not already done, create an account for the target environment. For more information, see [Managing Cloud Accounts](#).
2. Go to the appliance and click the `Machine Images` page. If you have not generated a machine image, you will need to do so as described in [Generating a Machine Image](#).
3. Click on the arrow pointing to the right to publish your image.



4. Following the instructions, choose the account and fill in any additional information required.

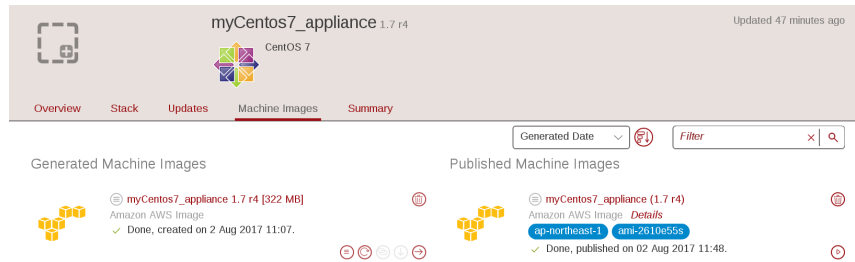
Warning: If your appliance has any elements such as passwords that must be set up by the user or licenses accepted at installation, you will see a warning when you publish the image which indicates `This image requires additional user interactions at startup` (as shown below). In this case, you will not be able to use the `deploy` feature.



5. Click publish.

Note: Publishing an image will be billed to the user account.

- The publication will take a few minutes to complete (depending on the size of the image and the network connectivity between UForge and the target environment). The publication progress is shown. At the end of the publication, the machine image has been published by UForge to your target environment.
- Under Published Machine Images, next to your appliance, you will now see the deploy (play) icon. In order to deploy your machine image, click on this icon.



- In the deploy view, enter the name of the deployment. Click on deploy. This will connect to cloud platform and launch your machine image instance. Once deployed, a green button will appear under the Status column.

For Linux images, you will need to input the Instance Name, Minimal number of cores and Minimal RAM values.

Deploy an instance to Amazon



⚠ Deploying an instance on this cloud platform may result in external charges.

Instance Name *

Minimal number of cores *

Minimal RAM (in MB) *

For Windows images, you will need to supply an additional User name and User password for an administrator account to be used for WinRM.

Deploy an instance to Amazon



⚠ Deploying an instance on this cloud platform may result in external charges.

Instance Name *

Minimal number of cores *

Minimal RAM (in MB) *

User name *

User password *

Note: When deploying a Windows machine image to Azure, you cannot use Administrator or Guest as the User name.

Warning: Windows Evaluation Editions are not supported

9. From the `Deployments` page you can see the instances deployed and their status. If the status is green it is up and running. If it is red, there has been an issue during deployment and/or it is stopped.

Deployments (3)

Cloud	Application	Locations	Sources Used	State
	AWS instance	us-northeast-1	myCentos7_appliance	Starting 
	Instance from scan	3116338284062466	Deployment_blackbox Scan #1	Running 
	OpenStack instance	3116338284062466	myCentos6_appliance	Running 

Note: Using this procedure, all deployed Linux virtual machines will have ssh port (22) open in their security group. All deployed Windows virtual machines will have RDP port (3389) and WinRM ports (5985, 5986) open in their security group.

Note: It may take a long time for a Windows virtual machine deployed to Azure to be provisioned.

Note: Currently, deployment of a Windows virtual machine from a published `Disk` is not supported. For more information on publishing of Windows machine images, see [Publishing a Windows Image to Azure Using Migration](#).

10. To stop a deployed instance, click on the delete (garbage) icon. This will remove it from your cloud.

Warning: Terminating an OpenStack deployment may fail due to a [known issue in AMP](#). Click again on the delete (garbage) icon to work around the issue.

6.2 Listing Deployed Machine Images

In order to view a list of the deployed machine images, go to the `Deployments` tab. A window similar to the following will appear.

Deployments				
Deployments (4)				
Cloud	Application	Locations	Sources Used	State
	My Instance (SSH)	westeurope	 Centos7	 Running 
	My Instance	westeurope	 Centos7	 Running 
	My Instance	eu-west-1	 Centos7	 Running 
	My Instance	3116338284062466	 Centos7	 Running 

From this window you can see the cloud platform the published machine image was deployed to, the name of the deployed instance, the location of the machine image published, the source machine image (Appliance or Scan) and the state of the deployed instance on the cloud platform.

Note: If you click on the delete (garbage) icon, this will stop your deployed instance and remove it from your cloud.

If you click on the application name you will see the details of the instance that has been deployed, including the name of the instance and the Hostname or IP address of the machine and the user name and the SSH key (if any). An example code will also be given to connect to your instance.

ldp-centos-simple_v1_r24

LD

MyAppService

Service

Status: running

Name: MyAppService

Type: com.ubersoft.brocklyn.type.LinuxAppliance

Appliance

Name: ldp-centos7

Version: 1

UUID: a2d83ec3-d883-4b78-85c8-853d6d10904

Machine Characteristics

Hostname: ec2-54-194-139-17.eu-west-1.compute.amazonaws.com

Instance Type: m3.medium

Memory Size: 3840 MB

Cores Number: 1

Number of Disks: 1

Total Disk Size: 2 GB

Location

Provider: aws-ec2

Region: eu-west-1

Cloud Image Id: ami-0bc2804fb3b6ef5ce

Connection

Connection User: root

UForge SSH Key Name: • LuubP

Example: ssh -root@ec2-54-194-139-17.eu-west-1.compute.amazonaws.com

6.3 Launching a PXE Image

The PXE format can be used to install an OS through the network, based on a PXE image created from your UForge. The only difference is that the PXE format is not downloaded with the UI of the UForge, but through an iPXE shell. By default the connection will be via https but if it fails it will fall back to http.

Note: The PXE format is only available for CentOS-based images.

You can execute iPXE script provided by UForge directly from iPXE shell. To do so, first make sure that your CentOS-based image is created in PXE format.

Note: You must be in an iPXE shell. Most virtual machines already implement iPXE. Make sure that you have the latest version of iPXE installed.

1. Start your machine and boot it on the iPXE shell.
2. Make sure the network interface is configured by executing “dhcp” in the iPXE shell.
3. Run the following command:

```
chain http://YOUR_FORGE_ADDRESS/resources/ipxe-init
```


For example:

```
chain http://10.1.2.197/resources/ipxe-init
```

or:

```
chain https://uforge.yourdomain.com/resources/ipxe-init
```

4. The following screen should appear and the boot process can be started.



```

Please complete the information below
Protocol: https
Your UForge IP address: <Select to change value>
Your root context: "/api"
Your credentials: <Select to change value>
Exit to shell
```

Enter your UForge IP, root context and credentials. Then select Connect.

5. You can then select the PXE image you want to launch.

Warning: It is possible to have a kernel panic error during the installation. This is caused by a bug in the Linux kernel (fixed in 3.16). In that case, just restart the installation.

Warning: When using Virtual Box, depending on the version, the installation can freeze. You may want to upgrade VirtualBox (version >=5.1.26).

6.3.1 Initializing the Boot Process Using a Customized ipxe.iso File

Most virtual machines already implement iPXE. You can customize to ease the use of iPXE in uForge.

The process of booting through iPXE can be started by using an `ipxe.iso` file. This file is built using the [iPXE open source project](#).

To do so, clone or download the ipxe open source project and build it. However, there are some parameters that can be used and

- You can embed a script in `ipxe.iso` so that it automatically executes some iPXE commands when booting on it. We already provide a script that allows to select the UForge from which you want to download PXE images at: `PATH_TO_THE_SCRIPT`. For example:

```
make EMBED=/home/user/scripts/myScript.ipxe
```


- You can activate the HTTPS protocol to download files. First, check in `src/config/general.h` that there is the line `define DOWNLOAD_PROTO_HTTPS` and not `undef DOWNLOAD_PROTO_HTTPS`. Then build the `ipxe.iso` by adding the argument: `TRUST=/etc/ssl/certs/CertificateAuthority.pem`. To build iPXE with HTTPS enabled run:

```
make TRUST=/home/uforge.yourdomain.com.crt
```

More information about how to build your project can be found on ipxe.org.

The project must be build in `/src` and the resulting output is found at `/src/bin/ipxe.iso`. This ISO can then be used through `qemu`, `VirtualBox`, `VMware vcenter` or on a USB stick to boot using iPXE protocol.

Once you are booting on `ipxe.iso`, the following screen should appear and the boot process can be started.




Enter your UForge IP, root context and credentials. Then select `Connect`. You can then select the PXE image you want to launch.

Warning: It is possible to have a kernel panic error during the installation. This is caused by a bug in the Linux kernel (fixed in 3.16). In that case, just restart the installation.

Warning: When using Virtual Box, depending on the version, the installation can freeze. You may want to upgrade VirtualBox (version `>=5.1.26`).

6.3.2 Using a DHCP Server that Redirects to an iPXE Script

The options of a DHCP server can be modified so that it redirects to an iPXE script (the one given on a UForge for example). To do so, the field named `next-server` must be associated with the filename to be executed. Using this, it is possible to automatically execute the script given on a UForge (at http://YOUR_FORGE_ADRESS/resources/pxe). Then in an iPXE shell, execute the command `dhcp`. The following screen should appear and the boot process can be started.



Enter your UForge IP, root context and credentials. Then select `Connect`. You can then select the PXE image you want to launch.

Warning: It is possible to have a kernel panic error during the installation. This is caused by a bug in the Linux kernel (fixed in 3.16). In that case, just restart the installation.

Warning: When using Virtual Box, depending on the version, the installation can freeze. You may want to upgrade VirtualBox (version $\geq 5.1.26$).

UForge allows you to create a multi-node application, called a blueprint, based on existing appliances and MySoftware in your UForge platform. The following sub-sections go into detail on how to deploy create and deploy a blueprint:

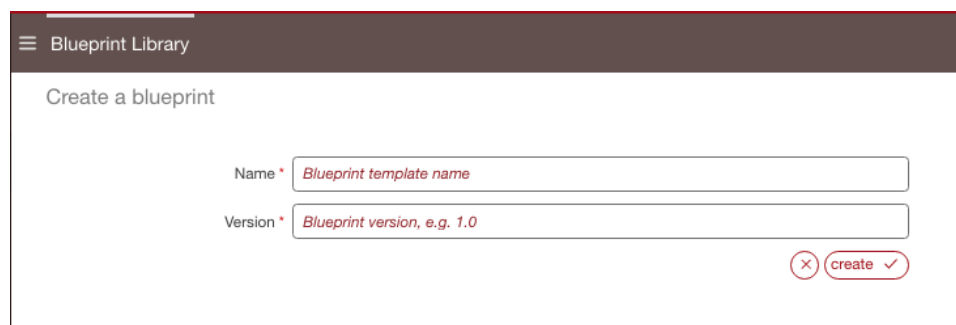
7.1 Creating a Blueprint

A blueprint is a multi-node application, based on Linux. It contains UForge appliances and MySoftware components. It configures dependencies between services, network topology, lifecycle and policies. Once your blueprint is created, it can be deployed on multiple cloud platforms (refer to [Deploying a Blueprint](#))

Note: Before you can create a blueprint, you must already have appliances created on your UForge. Refer to [Creating an Appliance Template](#).


To create a new blueprint:

1. Select Blueprints icon on the left.
2. On the Blueprint Library page, click on create in the top right.
3. Enter the Name and Version of the appliance.



The screenshot shows the 'Blueprint Library' interface. At the top, there is a dark header bar with a hamburger menu icon and the text 'Blueprint Library'. Below this, the main content area is titled 'Create a blueprint'. It contains two input fields: 'Name *' with a placeholder 'Blueprint template name' and 'Version *' with a placeholder 'Blueprint version, e.g. 1.0'. At the bottom right of the form, there is a red 'X' icon, a button labeled 'create', and a green checkmark icon.

4. Click the create button. The following window will be displayed.



You can choose to create and configure your blueprint either from:


- Graphical Designer, as described in this example
- YAML Editor

By default, the application does not have a name. To modify the name, select `New application` and in the right-hand configuration panel, enter the name. You can also enter an ID number.

Note: The blueprint application name and the name of the blueprint in UForge Blueprint Library can be different.

5. Select the appliances you want to add to the blueprint and drag them into the tree view. Your first element should be dragged onto the `New application` element. If you add an appliance while one is selected (red circle around appliance), it will be added as a child element of that appliance. You can drag and drop the elements as you wish. To add applications you can also click on an application from the left panel and click `add to...`

The following is a sample blueprint.



You can modify the name and configuration info of all the appliances in the right-hand configuration panel. To close this configuration panel, click X. To reopen the panel, click on the square in the top right.

6. Once you have created your blueprint, click on `save`. It will now be visible in the Blueprint Library.

Note: If you update any of the appliances used in the blueprint once the blueprint is created, this will be updated in the blueprint also. However, if your blueprint is already deployed, the changes will not be in the deployed blueprint. For more information about deploying blueprints, refer to [Deploying a Blueprint](#).

Note: If any of your appliances has a warning sign next to it in the graphical designer, this indicates the appliance is not ready to be used in a blueprint. Click the warning icon to get information on what actions are needed to make it ready.

7.1.1 Using the YAML Editor to Create or Modify a Blueprint

If you choose to create or edit a blueprint with the YAML editor, you should note the following:

The section must be written with YAML blueprinting language. This language is defined by the open-source Apache Brooklyn project. For more information, you can refer to [Apache Brooklyn documentation](#).

The `id` attribute can be used in the YAML configuration to retrieve information from an item in the services list. For example you can set `myService` as a service ID and then use it in the `shell.env` configuration of another service using the following Brooklyn command:


In the `customize.command` attribute you can execute commands using the shell language. A customized command is not stored in the UForge appliance nor UForge images in cloud catalog. These commands will be executed with the user `uforge` during the deployment in the deployed instances; therefore, you will be allowed to update your scripts without regenerating UForge images.

If you need root access for your commands, use the `sudo`. Also, ensure that packages `sudo` is included inside your appliance's OS profile.

You can use your own variables by declaring them in the `shell.env` attribute. Then you can use them with this syntax: `${MY_VARIABLE}`.

7.1.2 Modifying the Blueprint Configuration

To modify the blueprint configuration, you can select any of the elements of the blueprint and modify the configuration elements in the right-hand panel.



You can either create a new custom configuration or use a standard one from the list.

Note: If you update any of the appliances used in the blueprint once the blueprint is created, this will be updated in the blueprint also. However, if your blueprint is already deployed, the changes will not be in the deployed blueprint.

7.2 Deploying a Blueprint

Once you have created a blueprint you can deploy it to a cloud platform. For details about creating a blueprint, refer to *Creating a Blueprint*.

To deploy a blueprint:

1. From the **Blueprints** page, select the blueprint you want to deploy.
2. Click on **deploy** in the top right.
3. Enter the Name of the deployment.
4. **Select the target platform you want to deploy to.**

Note: Currently only AWS and Azure are supported.

5. Select the cloud account you want to use to deploy the blueprint.
6. From the drop-down menu, select the **Region** in which the blueprint should be deployed and fill in any additional information required.
7. On the bottom of the page, there is a list of the appliances used by your blueprint. For each of these appliances, you can see their current status. All the appliances need to be ready (an appliance is ready when its last version

is published to the selected region). If an appliance is not in `Ready` state, you will not be able to publish. You can prepare it for deployment directly from this page by clicking `prepare`.

Deploy MediaWiki

Deploying a blueprint may result in external charges

Name *

Target Platform *

Account *

Region *

Storage Account *

Appliances status

Name	Version	Revision	Status
mariadb	1	3	Not generated for this revision
mediawiki	1	3	Ready

The appliances that were not in `Ready` state will go to state `Generation ongoing` then to `Publish ongoing`.

Note: Using the `prepare` button will automate the generation and publication of all the appliances that are not in ready state. However, you could also generate and publish each appliance separately.

8. Once all the appliances are ready, click on `deploy` in the top right. Once you launch the deployment, you will be redirected to the `Deployments` page. From here you can see the status of the blueprint. The elements will be flashing green throughout the deployment process. Then the status of each element is color coded:

- green is ok
- orange is stopped
- red is error

The following is an example of a deployment.

Deployments

Graphical runtime

mediawiki_v1_r6

mariadb

mediaWiki

Service

Status: running

Name: mediaWiki

Type: com.usharesoft.brooklyn.type.LinuxAppliance

Appliance

Name: kdp-ck-centos7

Version: 1

UUID: a2d63ec3-c683-4b78-b5c8-853c6a810904

Machine Characteristics

Hostname: ec2-54-194-139-17.eu-west-1.compute.amazonaws.com

Instance Type: m3.medium

Memory Size: 3840 MB

Cores Number: 1


Number of Disks: 1

Total Disk Size: 2 GB

Location

Provider: aws-ec2

If there is an issue or error on one of the appliances, it will appear in red and the error will be shown with a warning symbol, as shown below.



You can delete your deployment by clicking on `terminate` in the top right. This will stop and delete the deployment not only from UForge but also from your cloud account.

9. If you click on an application name, you can view the details of the deployment. For more information please visit [Apache Brooklyn Official Documentation](#)

Migrating Live Workloads

UForge AppCenter offers the capability to migrate a live system using two different methods:

- using `Lift & Shift`. This is an automated process which covers three steps: scan, generate and publish.
- using `Re-platform`. In this case the steps are executed separately by the user: scan, update, generate and publish. To update, you must import as appliance and modify it prior to the final migration.

In both cases, UForge will “deep scan” a live system and report back the meta-data of every file and package that makes up the running workload.

The goal of `Lift & Shift` is to produce a near identical copy of the currently running workload. However, there will always be small differences between the two workloads after migration is complete, notably some services are disabled or enabled depending on the target machine image being created (refer to *Service State*). For more information on the differences after migration, refer to *Differences between Source and Target Instances*.

The goal of `Re-platform` is to analyze an existing system and allow the user to make the necessary changes (potentially add or remove software, upgrade, re-configure etc).

The following sub-sections describe in detail the different steps, requirements and limitations of migrating your system, for each of the methods possible:

8.1 Migration Pre-requisites

You must have root access on the target system in order to complete a migration, as you will need to copy and run a binary file on the target system.

Warning: UForge AppCenter migration does not support multi-kernels. When scanning a machine with more than one kernel, only the kernel running will be scanned and imported.

In order to migrate a system, it must meet the following conditions:

- The operating system must be in one of the supported formats (refer to *Supported Operating Systems*).

Note: Your scan will take longer if not all minor versions of a distribution are installed in your UForge AppCenter. For example, if you scan a CentOS 6.8 machine, but your AppCenter has only been populated with packages up to CentOS 6.7, then the AppCenter will use the machine's yum repo to download the missing packages. As a result, the scan will take longer to complete.

- The image formats must be in supported formats (refer to *Supported Machine Image Types*)
- Any pre-install or post-install scripts on the system you are about to scan should only use ASCII characters. Otherwise UForge AppCenter will return a scan error: DB Error - invalid characters.
- Custom packages on the live system to be scanned should not contain references to package dependencies as relative paths. They should be expressed as absolute paths.
- If custom packages are installed using `--nodeps` flag, the scan process will not detect these packages. When running `Re-platform`, UForge AppCenter will check for these dependencies. If they are custom packages that are not on the live system, the generation will fail. Therefore, it is recommended to provide a custom repository with all the necessary custom packages. Otherwise, they can be added after the scan to the appliance template in `My software`.
- For Windows, the partition format must be NTFS
- For Windows, GPT is not supported. Scanned Windows machine has to have MBR.

Warning: Currently, UForge is not able to migrate the Yum repository GPG keys except the ones located in the `/etc/pki/rpm-gpg/` directory. This means that the user will have to accept the repository GPG key when the user installs or updates a package. The user will have to do this only once per repository.

8.1.1 Migrating Windows to K5

If you plan to migrate a Windows instance onto [K5 Fujitsu Public Cloud](#), you must uninstall CloudBase-Init (if installed) before scanning.

For more detailed information, refer to [official Fujitsu K5 IaaS Documentation](#).

8.1.2 Migrating Windows

If you plan to migrate a Windows machine to any cloud, you need to ensure that you have the following RDP update: KB4103723.

8.1.3 Migrating Linux to VCenter

If you plan to migrate a Linux instance from a platform using cloud-init (such as OpenStack) to VCenter you must follow the following steps:

- Uninstall `cloud-init` (if installed) before scanning
- Use `Re-platform`. Once the appliance is imported, remove `cloud-init` package before generating.

8.1.4 Migrating Linux to Microsoft Azure

Currently, to publish to Microsoft Azure platform <https://azure.microsoft.com/en-us/> you must do the following before scanning:

1. Unless you are migrating from CentOS7+ or Red Hat Enterprise Linux 7+, uninstall NetworkManager (if installed).
2. Uninstall the Microsoft Azure agent, i.e. WALinuxAgent and waagent packages (if installed).

8.1.5 Migrating Windows to Microsoft Azure


You must install the Azure Virtual Machine Agent on the source machine before scanning if you plan to migrate using Lift & Shift or using Re-platform and the appliance is configured as to not run sysprep automatically on its first boot.

8.2 Migration Process Overview

The five main steps of migration are as follows:

1. The live system is “deep scanned”, detecting all the files and configuration information.
2. The scan report and overlay is sent to UForge AppCenter.
3. A machine image is generated from the scan. With Re-platform you can import the scan prior to machine image generation to change the content.
4. The machine image is uploaded and registered to the target environment.
5. The registered machine image (also known as a template in cloud terminology) can be used to provision one or more instances. These instances have near identical content to the original live system.

Lift and Shift automatically executes steps 1, 2, 3 and 4. With Re-platform, each of these steps are done manually, allowing the user to make the necessary changes after step 2.



For both migration using Lift & Shift and Re-platform, UForge will scan the system you want to migrate and detect the following information:


- all the files and packages on the system (including configuration information). If you have selected a scan without overlay, then extra files and specific configuration information will be detected but will not be included in the report.

- root and user password (encrypted)
- timezone. As part of the generation, UForge will use the same timezone as the original system. UForge creates a symbolic link for `/etc/localtime` that point to the timezone set in `/usr/share/zoneinfo`.
- keyboard settings
- kernel parameters
- users and groups
- filesystem layout (partitioning)
- **SELinux settings are added to the `Install Profile` for CentOS, Red Hat and Scientific Linux only. Note the following**
 - if SELinux is not activated on the scanned machine, it will not be activated on a VM generated from this scan.
 - if SELinux is activated, it will be activated with the same rules, and the file system will be relabeled on the first boot because of `.autorelabel` file (and this file is deleted by SELinux once the relabel is done). Parameters from the `/etc/selinux/config` file other than the SELinux mode are lost during the appliance import. If you want to modify some of these parameters you will need to add a MySoftware overlay with a `/etc/selinux/config` file after the scan. Refer to [Adding Software from Your Software Library](#).

Warning: All other information found in the scan is not used (reset) and the `Install Profile` default information is used. This includes:

- Networking information (IP address). In the case of Windows systems, or with NIC cards managed with NetworkManager, only the first NIC is taken into consideration. Note that if the first card is static, it will be changed to DHCP.
- SSH keys. You will need to manually add the ssh keys to the install profile.

Warning: Currently, UForge is not able to migrate the Yum repository GPG keys. This means that the user will have to accept the repository GPG key when the user installs or updates a package. The user will have to do this only once per repository.



8.2.1 Scan vs Scan with Overlay

You can select two types of scan: scan or scan with overlay. A scan will scan the target environment, but will only return the system packages and configuration known to UForge. A scan with overlay will include an overlay report which details all the extra files, packages or specific configuration. The overlay is all the things that are missing compared to a known state (a previous scan of a machine or the operating system native packages). This overlay is a standard tar archive. Once created, it is uploaded via HTTPS to the UForge AppCenter. The overlay is not built on the scan target but it is stream uploaded (faster and does not need any space on the scan target machine).

Note: Scan with overlay is only supported for Linux-based machines.

Warning: The following file types are not included in the scan overlay file:

- character device
- block device
- FIFO (named pipe)
- socket

8.3 Migrating Using Lift & Shift

You can use the `Lift & Shift` tab on the `Migrations` page to complete a migration automatically. When you run an automatic migration, you cannot modify any of the elements of the system.

Note: Refer to [Migrating Using Re-platform](#) if you want to modify the machine before migration.

The automatic migration process includes the following steps:

1. Scan the source machine.

2. Generate a machine image from the scan result.
3. Publish the generated machine image to the selected target environment.

8.3.1 Launching an Automatic Migration

All pre-requisites listed in the section *Migration Pre-requisites* are also applicable to Lift & Shift migration.


In order to run a migration:

1. Go to the Lift & Shift tab on the Migrations page.
2. Click on the migrate button on the top right.



3. Enter the name of your migration.

4. From the drop-down menu select the target platform and type. Enter all other required fields, based on the target you have selected.
5. Click next.
6. Follow the instructions on the UForge AppCenter pop-up window.



7. Download binary locally by clicking Download.

Note: Older versions of curl command may not support a secure connection with UForge server. If you encounter a *SSL connect error*, you can use a web browser to download the binary or download the binary on another machine and copy it to the scanned machine.


8. Copy the binary on the source environment you want to migrate.
9. Open a terminal window and login to the source environment.
10. Launch the migrate command on the running source environment to start the migration.

Note: For Linux migration, you can exclude non-native files and directories using `-e` parameter. The `-o` parameter will include the overlay (this option is used by default). The overlay is all the things that are missing compared to a known state (a previous scan of a machine or the operating system native packages). This overlay is a standard tar archive.

Note: The uforge migrate tool self extracts in `/tmp` directory before execution. If the machine is configured with `noexec` option on `/tmp` partition a `Permission denied` error will occur. In that case, you can specify the environment variable `TMPDIR` to use a different directory. For example:

```
TMPDIR=/root/ ./uforge-migrate.bin -U https://<uforge-ip>/api -u
↪<username> \
    -n "Test_migration" -o
```

11. Click next to view the migration status.
12. You can return to the Lift & Shift page to view the status of your migrations.




Note: You should note that if you cancel a migration, this will stop the process of migration at the step currently being done. Therefore, for example, if the image generation is complete, it will be visible on your platform. Also, if publication has started, this may incur costs. You should check your cloud account. For more details on the constraints of publishing an image, refer to [Publishing a Machine Image](#).

From the status page you can view more detailed information about the scanned instance. Refer to [Viewing Details of an Automatic Migration](#).

8.3.2 Viewing Details of an Automatic Migration

Once your automatic migration has completed, you can view the details of the status and details of the machine migrated from the `Lift & Shift` tab on the `Migrations` page.



From the status page you can view more detailed information about the appliance scanned by clicking:

- [Details](#) for more information about the published appliance. This will include the cloud ID, cloud account used for publication, if the publication succeeded.
- [View Scan Details](#) which will bring you to the details of the scan appliance on the `Re-platform` page. From here you can see all the packages and files that are part of the appliance.
- [View Generations Details](#) which will bring you to the `Generations` page. This page lists ALL generations (not just from this migration). From here you can launch a publish.
- [View Publication Details](#) which will bring you to the `Generations` page.

8.4 Migrating Using Re-platform

The entire migration process using `Re-platform` (previously referred to as `scan`) has 5 main steps:

1. Scan the source machine. You can select two types of scan: `scan` or `scan with overlay`. A `scan` will scan the target environment, but will only return the system packages and configuration known to UForge. A `scan with overlay` will include an overlay report which details all the extra files, packages or specific configuration.
-
- Note:** Scan with overlay is only supported for Linux-based machines.
-
2. Report the scan results to UForge AppCenter, where the platform analyzes the report. The scan differentiates between known packages and extra files. This report is transferred via HTTPS to UForge AppCenter.
 3. The results are used to build an archive. The extra files are only included in the archive if the user is performing a scan with overlay. The archive is sent back to the platform.
 4. The archive is uncompressed, and the information is stored in the database as a `Scan`.
 5. You can use the scan to generate machine images or imported to create a new appliance template. The generated machine image can then be published to the target environment and instances can be provisioned.

Note: When generating a machine image from an appliance template, package dependencies are calculated using the list of packages in the OS Profile. Any missing packages from the OS Profile are added prior to generating the machine image.

Warning: There may be cases where packages that are present on the scanned server have been removed or replaced in the newer version of the OS. Such packages will obviously not be present in UForge AppCenter’s package repositories, and will therefore be added to a `My Software` component, instead of being added to the `OS Profile`. Therefore, if you update the OS version when performing a migration using `re-platform`, and if there are such packages, you should first manually remove or replace them in the template, as instructed by the distribution vendor, before generating the machine image.

For detailed procedures on migration using `Re-platform` refer to the following sections:

- *Scanning a Linux Machine*
- *Scanning a Microsoft Windows Machine*

8.5 Scanning a Linux Machine

Warning: When scanning a Linux machine, you have to check whether the licenses of OS and software which the source machine contains allow you to use them on the destination server which you are migrating to. For more detail, refer to *Notes on Licensing*.

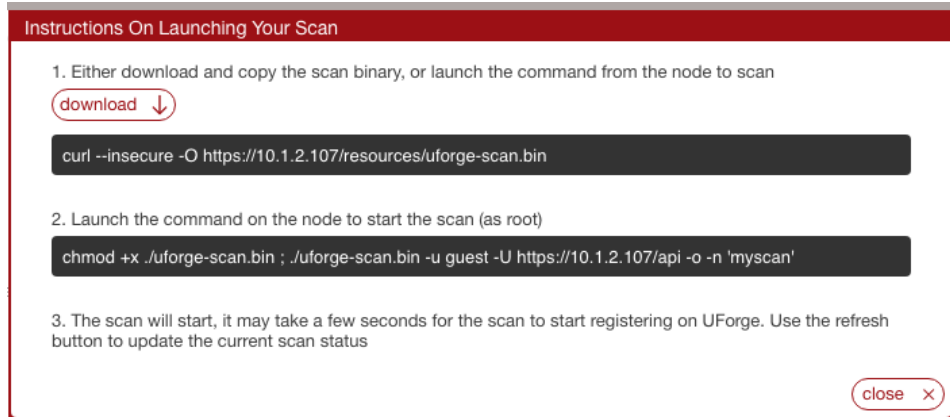
To carry out a scan, go to `Re-platform` tab on the `Migrations` page:

1. Click on `scan` in the top right.
2. Enter a name for the scan of the target system you want to migrate.

3. Select `Linux` from the drop-down menu.
4. Leave `Include overlay` selected if you want to run a scan with overlay. This will return all the packages, files and configuration information of the source machine.

Note: If you do not include the overlay, all files needed to rebuild packages are still transferred to UForge, including the package configuration files. However, all “extra_files” are excluded from the scan.

5. If you want to exclude certain directories or files from the scan then click **add** and enter the directory path or full pathname of the file.
6. Click **Next**. Follow the instructions on the UForge AppCenter GUI.



7. Download binary locally by clicking **Download**.

Note: Older versions of curl command may not support a secure connection with UForge server. If you encounter a *SSL connect error*, you can use a web browser to download the binary or download the binary on another machine and copy it to the scanned machine.

8. Copy the binary on the target environment you want to migrate.
9. Open a terminal window and login to the target environment.
10. Run the scan command on the running target environment to start a scan of the system you want to migrate. Note the `-o` option in the scan command indicates that you will launch a scan with overlay. The binary identifies the packages, files and custom files on the system.

By default the scan data will be saved in `/tmp`. You can modify the directory where the data will be saved using the `-t` option in order to ensure that there is enough space to save the scan data.

Use the `-e` option if you want to exclude certain files or directories from the scan.

You can also use API keys to run the command. In this case, in the command you copied, remove the password and enter the API keys using `-a` option for the public key and `-s` option for the secret key. For example:

```
./uforge-scan.[bin/exe] -u <username> -a <public-key> -s <secret-key> -U http://ip:port/ufws -n 'Test_scan'
```

Note: The `-n` option to indicate the scan name is mandatory. If a scan with this name already exists, the scan will be grouped under this name.

Note: The uforge scan tool self extracts in `/tmp` directory before execution. If the machine is configured with `noexec` option on `/tmp` partition a `Permission denied` error will occur. In that case, you can specify the environment variable `TMPDIR` to use a different directory. For example:

```
TMPDIR=/root/ ./uforge-scan.bin -u <username> -a <public-key> -s  
↩<secret-key> \  
-U https://<uforge-ip>/api -n 'Test_scan'
```

11. A report is sent to UForge AppCenter which can be used for migration. To view the progress, go back to the `Re-platform` page.

Note: The duration of the scan depends on:

- the power of the machine in the target environment,
- the complexity of the target environment OS (number of packages installed),
- the network bandwidth between the target environment and UForge.

Scans of typical simple target environments can last about 5 to 15 minutes. In the case of larger and more complex target environments, together with poorer bandwidth, one can experience durations of up to one hour.

12. To view the details of a scan, click on the scan and refer to [Viewing and Using a Scan](#).
13. To complete your migration, you will need to generate a machine image and publish it. Refer to [Generate a Machine Image from Scan](#).

Note: Image generation will fail when migrating if the source server has the same LVM volume group name as the UForge server's one. It will fail also if the volume group name set in the Partitioning Table is the same as the name of LVM volume group in UForge server.

8.6 Scanning a Microsoft Windows Machine

Note: When scanning a Microsoft Windows machine, you must acquire Windows licenses in order to handle Windows OS in UForge and confirm usage conditions of cloud provider and virtualization software which you scan and migrate to.

Restrictions

- It is not possible to scan a Windows “Core” system for all versions: 2008R2, 2012, 2012R2 and 2016.
- Windows Evaluation Editions are not supported for scanning.

Warning: If the target filesystem is NTFS, the scan is optimized by extracting only “used space” from target disks on the source system. Some applications might hold their data on “free space,” which is not used by the operating system. If the source system of the scan has such applications installed, these applications may not work correctly on a machine image generated from the scan. With the other filesystems, such as FAT, ReFS and so on, all the space including “free space” on target disks will be copied by the scan.

Pre-requisites


For Windows Server 2008R2, you will need to run a command in PowerShell on the source machine before scanning `winrm set winrm/config/service '{@AllowUnencrypted="true"}'`. This command will

disable WinRM encryption over HTTP which is not supported by AMP. More information can be found on the Cloudsoft AMP documentation, in section [WinRM Connectivity Diagnostics](#).

If you plan to migrate a Windows instance onto [K5 Fujitsu Public Cloud](#), you must uninstall CloudBase-Init (if installed) before scanning. For more detailed information, please refer to [official Fujitsu K5 IaaS Documentation](#).

To carry out a scan, go to the `Re-platform` tab on the `Migrations` page:

1. Click on `scan` in the top right.
2. Enter a name for the scan of the target system you want to migrate.
3. Select `Windows` from the drop-down menu and click next.
4. Select if you want to use the graphical user interface or the command line. Follow the instructions on the pop-up to download the scan binary.



When you run the `uforge-scan.exe` command the `-n` option to indicate the scan name is mandatory. If a scan with this name already exists, the scan will be grouped under this name.

By default the scan data will be saved in `/tmp`. You can modify the directory where the data will be saved using the `-t` option in order to ensure that there is enough space to save the scan data.

Use the `-e` option if you want to exclude certain files or directories from the scan. Use `-E` option if you want to exclude a harddisk.

5. You can then launch the scan. The following image illustrates a Windows scan with all of the possible options checked. They are described in the following steps.

UForge Scan

Connection

UForge URL:

Use API keys authentication: ☒

Username:

Public key:

Secret key:

Proxy authentication required: ☒

Proxy User:

Proxy Password:

Scan

Scanned Instance Name:

Use Local Storage: ☒

Temporary Storage Dir:

Disks or Partitions to exclude:

- ☐ \Device\Harddisk0 { C: }
- ☐ \Device\Harddisk0\Partition1
- ☐ \Device\Harddisk0\Partition2 (C:)
- ☐ \Device\Harddisk1 { E: F: G: }
- ☐ \Device\Harddisk1\Partition1 (E:)
- ☐ \Device\Harddisk1\Partition2 (F:)
- ☐ \Device\Harddisk1\Partition3 (G:)

Cancel Scan

6. Optionally you can use API keys. In this case, check `Use API keys authentication` and enter the public and secret key information.
7. If you want to set up a proxy, check `Proxy authentication required` and enter the proxy information.
8. Fill in the `Scanned Instance Name`.

Note: Scan name can be mix of alphanumeric characters, spaces and the following special characters `._-`. All other characters are not supported at this time.

9. Optionally you can select `Use local storage`. This means that the scan will be not be done in streaming but in 2 phases. First the data will be stored on a temporary storage drive during the scan process. This temporary storage can be a local directory or a virtual space on the network. It can also be on the same partition that is being scanned (provided there is enough space available). It must be at least half the size of the machine you want to scan.

Note: If you are using local storage, UForge AppCenter will generate a script for you named

`upload_raw_partition_images.bat` that you will have to launch at the end, to upload the archive to UForge AppCenter once the scan is complete.

10. If you want to exclude certain partitions from the scan then check the boxes accordingly in the section `Disks` or `Partitions` to exclude.
11. Click `Scan` to launch the scan. A report is sent to UForge AppCenter which can be used for migration. To view the progress, go back to the `Scans` page.

Note: `uforge-scan.exe` uses Microsoft Windows snapshot mechanism, which requires free disk space on each scanned partition. If not enough free disk space is available, this may lead to a scan failure.

12. To view the details of a scan, click on the scan and refer to [Viewing and Using a Scan](#).

Note: The log files of `uforge-scan.exe` on the source Microsoft Windows machine are stored under the following directory: `%SystemRoot%\System32\LogFiles\UForge`.

13. To complete your migration, you will need to generate a machine image and publish it. Refer to [Generate a Machine Image from Scan](#).

8.7 Viewing and Using a Scan

Once you have run a scan, you can view the scan details.

To view the scans, go to the `Re-platform` tab on the `Migrations` page. All the scans are listed by name. If you have run the same scan with the same name, it will appear as `Scan <number>`.

Lift & Shift		Re-platform		
Scanned Nodes				
The scan analyses the source system and stores the results as meta-data. You can generate machine image or create an appliance template from this scan.				
			Filter	x Q
Name	Status	Distribution		
migration-from-CentOS-to-V Center-scan (1)		CentOS 7 x86_64		
Scan #1	Finished (less than a minute ago)			
migration-from-CentOS-to- K5-scan (1)		CentOS 7 x86_64		
Scan #1	Finished (less than a minute ago)			

From this page you can use the scans to do the following:

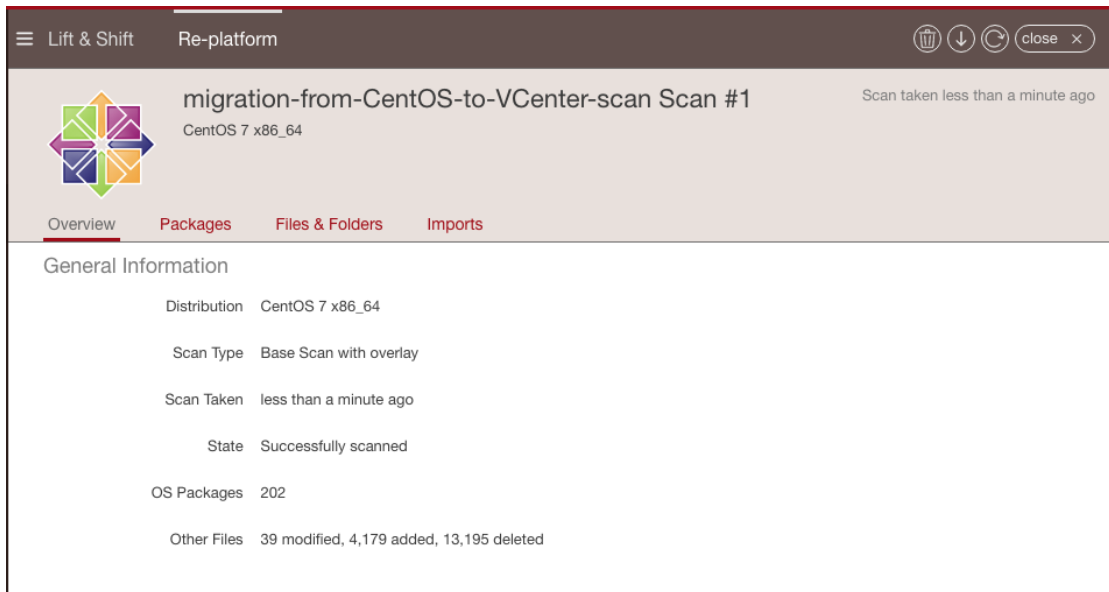
8.7.1 Viewing the Details of Scan

You can view the details of the scan report that will display the packages that are present as well as the filesystem detected on the scanned source instance. The information available depends on whether the scanned instance is Linux-based or Windows-based.

Linux-based

To view the details of a scanned Linux-based instance:


1. Go to the `Re-platform` tab on the `Migrations` page.
2. Click on the scan to view the details.



The screenshot shows the 'Re-platform' tab in UForge AppCenter. The selected scan is 'migration-from-CentOS-to-VCenter-scan Scan #1' for 'CentOS 7 x86_64'. The 'Overview' tab is active, displaying the following information:

- Distribution:** CentOS 7 x86_64
- Scan Type:** Base Scan with overlay
- Scan Taken:** less than a minute ago
- State:** Successfully scanned
- OS Packages:** 202
- Other Files:** 39 modified, 4,179 added, 13,195 deleted

3. All of the packages and non-native files are listed under the `Packages` tab.



The screenshot shows the 'Packages' tab for the same scan. A checkbox 'Only show packages with changes' is checked. A search filter is present. The following table lists the packages:

Name	Version	Release	Architecture	In Repo
acl	2.2.51	12.el7	x86_64	✓
audit-libs	2.4.1	5.el7	x86_64	✓
basesystem	10.0	7.el7.centos	noarch	✓
bash	4.2.46	20.el7_2	x86_64	✓
bind-libs-lite	9.9.4	29.el7_2.4	x86_64	✓
bind-license	9.9.4	29.el7_2.4	noarch	✓
binutils	2.23.52.0.1	55.el7	x86_64	✓

4. You can also filter the packages that have been modified (UForge AppCenter compares the packages scanned with its repo) by checking `Only show the packages with changes`.
5. To view the details of a package, click on the package name and then the arrow.

Note: The number of packages between your scanned system and the one in UForge AppCenter will differ for several reasons. First, if you had more than 1 kernel only 1 is imported into UForge AppCenter.

Also, UForge AppCenter adds files for install configuration and install profile.

6. From this page you can now import as appliance (*Create an Appliance from a Scan*) or generate an image (*Generate a Machine Image from Scan*).


Windows-based

Note: The details of the scan are for information purposes only. They cannot be modified. The following applications in your scanned system are listed in the Applications tab.

- Applications displayed in Control Panel > Programs and Features page
- Windows Store applications

To view the details of a scanned Windows-based instance:

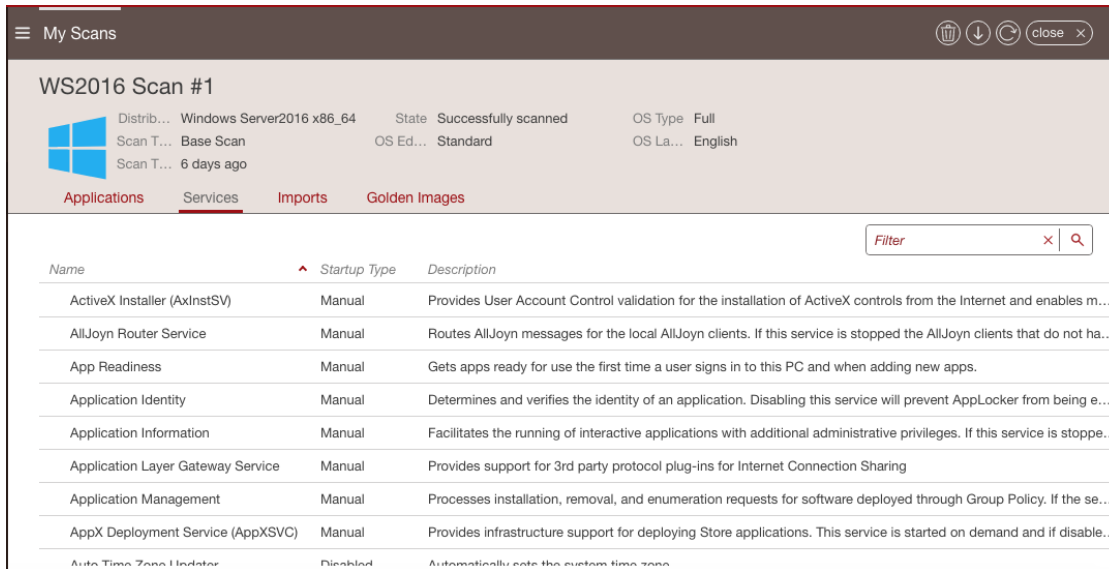
1. Go to the Re-platform tab on the Migrations page.
2. Click on the scan to view the details.
3. To view the Windows applications, go to the Applications tab.



The screenshot shows the 'My Scans' interface. At the top, there's a header with 'My Scans' and a close button. Below it, the scan details for 'WS2016 Scan #1' are displayed. The details include: Distrib... Windows Server2016 x86_64, State Successfully scanned, OS Type Full, OS Ed... Standard, OS La... English, Scan T... Base Scan, and Scan T... 6 days ago. Below the details, there are four tabs: Applications (selected), Services, Imports, and Golden Images. The Applications tab shows a table of installed applications.

Name	Version	Path	Size	Architecture	Windows store
Microsoft Visual C++ 2008 Redistrib...	9.0.30729.6...		1 MB	x64	
Microsoft Visual C++ 2008 Redistrib...	9.0.30729.4...		872 KB	x86	
Microsoft.AAD.BrokerPlugin	1000.14393....		N/A	neutral	✓
Microsoft.AccountsControl	10.0.14393....		N/A	neutral	✓
Microsoft.BioEnrollment	10.0.14393.0		N/A	neutral	✓
Microsoft.LockApp	10.0.14393.0		N/A	neutral	✓
Microsoft.Windows.Appprep.ChxApp	1000.14393....		N/A	neutral	✓
Microsoft.Windows.AssignedAccess...	1000.14393....		N/A	neutral	✓

4. To view the Windows services, go to the Services tab.



- From this page you can now import as appliance (*Create an Appliance from a Scan*) or generate an image (*Generate a Machine Image from Scan*).


8.7.2 Generate a Machine Image from Scan

Once you have scanned a source system, you can generate a machine image directly from the scan report. When you generate a machine image from a scan, all the information included in the scan report is used in constructing the new machine image (except extra files if you performed a scan without overlay). However, prior to the generation starting, you will be prompted to indicate if you want to change some basic settings of the filesystem, namely the overall disk size and the swap size. You cannot set the swap size to 0. If you want to delete the swap partition, you must do this in Advanced Partitioning (refer to *Configuring Advanced Partitioning*.)

Note: When generating a machine image, some services are disabled or enabled depending on the target machine image being created (refer to *Service State*).

To create an image from a scan:

- Go to the Re-platform tab on the Migrations page.
- Click on the scan to view details.
- Click on the generate icon in the top right.



- Select the image format you want.
- Click Generate. UForge AppCenter generates the machine image:

- Dependency checking is SKIPPED. This is done intentionally so that UForge does not alter the package list manifest detected during the scan process.
- Creates the disk ready for installation (using the disk size and partitioning by the user if they have requested a change)
- Installs the native os packages from the scan report
- Applies the overlay file from the scan report, if you performed a scan with overlay
- Applies the low configuration information detected in the scan report (passwords, time-zone, keyboard, etc)
- Applies any specific libraries or configuration depending on the machine image format chosen (e.g for AWS UForge adds the required AWS libraries)

Note: The networking information is treated differently depending on whether the IP address of the workload being migrated is using a static IP address or DHCP.

- Static IP Addresses: The current information detected during the scan is kept. During the generation phase, this networking information is also kept. Consequently, the new machine instance has the same static IP address set.
 - Dynamic (DHCP) IP addresses: In this case, the networking information, is reset namely the IP address information is removed during the generation process, and is setup as DHCP. When the new machine instance is provisioned, the instance sends a request to the local cloud DHCP service to get a new IP address.
-

6. You can view the progress and the completed generation on the `Dashboard` tab under `Generations`. You can provision one or more instances from the machine image. Each instance being a near identical workload from the original. For information on the differences after migration, refer to *Differences between Source and Target Instances*.

Note: In the case of a migration from a para-virtualized platform to a non para-virtualized platform, UForge AppCenter injects everything that is needed to make the machine work (the kernel and its tools). Based on the packages discovered on the scan target and on the underlying operating system, UForge AppCenter calculates the most accurate kernel version to inject for your machine.


8.7.3 Create an Appliance from a Scan

You can create an appliance template from a scan, using `import`. This allows you to change and modify the contents of the machine that has been scanned. Once you create an appliance template, all packages and files that are known will be listed under `OS profile`, while “unknown” packages and files will be listed under `MySoftware`. When you import a scan as an appliance template, the import process basically transforms the meta-data of the scan report to an appliance template.

To create an appliance template from a scan, from the `Re-platform` tab on the `Migrations` page:

1. Click on the scan from which you want to generate an image.
2. Select the `import` button (downward arrow) at the top right to create an appliance template from the scan. The following example is for Linux, but you can also import a Windows scan.


Note: If you have administrator rights, for a Windows scan you can choose to import as appliance or import scan as a golden image. For more information, refer to the Admin guide.



3. Enter the appliance name and version.
4. Click import. The progress will be shown under the Imports tab. Once complete you can find the appliance also under the Apps tab, under Imported Appliances.

You can now generate a machine image and share it, as you would any other appliance template. It is also possible to push it to a remote environment. The image is then ready to be launched in the new environment and the migration is finished.

Note: When you create a Windows appliance from scan, the applications and services will be visible under OS Profile on the Stack page.



More importantly you can now change the contents of the original scanned system. If you go to the Apps tab, the new appliance template will be listed in the Imported Appliances section. Double-click on it to view the details or modify it.

Warning: Once you have created a Windows appliance using import from a scan, you will only be able to import this appliance to another UForge platform by removing some OS section fields. Refer to [Updating a Windows Appliance Before Import](#).

Changing a Configuration with “No-Console” Features

This is typically the case when migrating to Azure. Azure does not have any console facility today during first boot of the instance. When you migrate a workload using `Lift & Shift`, no install type questions are asked on first boot. This is due to:

- root password is copied
- SSH keys are copied
- partitioning table is preserved
- keyboard is preserved
- timezone is preserved

Access to the machine would typically be done via SSH.

Migration using `Re-platform` provides more flexibility. As soon as you import a scan, which effectively creates an appliance template, you will have access to the Install Profile. This allows you to reset and change many of the “installation”/first boot parameters including prompting the end user to provide the information (for example: ask the end user to set the root password). Any prompt to the end user is normally displayed in the console. However if the user logs into the machine for the first time via SSH, these prompts are displayed in the SSH terminal and not the console.

Warning: If you decide to prompt the user for the root password, then an SSH key mechanism must already be determined (private key owned by the user and public key set in the Install Profile). Otherwise the user will be locked out of the system with no way to SSH into the machine as no password has been set yet.

Changing Configuration Information

If you want to modify configuration information of a scanned system, there are several solutions:

1. Using pre-install and post-install scripts
2. Using boot scripts
3. Integrate with a Configuration Management platform

In each case, you must import the scan as an appliance template.

Solution #1: Using pre-install and post-install scripts in the package mechanism: RPM and DEB package mechanism allows you to register scripts that are executed at various moments during the installation of the package. By packaging your middleware or application binaries as a native package you can register such scripts. These scripts are automatically taken into account as part of the machine image generation process. These packages can be added into the appliance template – either in a custom repository known by UForge AppCenter (in this case the packages are displayed as part of the repository and are added in OS Profile) or as part of a Project or My Software component.

Solution #2: Using boot scripts. UForge AppCenter allows you to add boot scripts in the appliance template. These boot scripts are executed the first time the migrated instance is provisioned. Boot scripts can also be registered to be run every time the instance is rebooted.

Solution #3: Integrate with a Configuration Management platform: There are many 3rd party platforms including Puppet, Chef, Ansible and Saltstack that can be used to configure middleware and application layers. Once a system has been migrated or a machine image generated from an appliance template, such configuration management platforms can be used for package update and configuration. You may need to include a bootstrapping mechanism to register the instance to the configuration management platform of your choice. This bootstrapping can be done using boot scripts (see solution #2).

Adding Security Patches

When a scan is imported as an appliance template all the native packages detected from the scan are compared with the UForge AppCenter package repositories. UForge AppCenter will immediately inform you whether new package versions are available for your scan report. Using the “appliance update” feature, a graph is displayed showing you all the available updates and allows the user to update the appliance template to the latest available packages. Once the appliance template has been updated, the user can then generate the machine image and register the machine image on the target environment. The migrated instance will have the latest package updates.

Of course this is not the only way to update a migrated system. The administrator can update the live system using the standard operating system update mechanism. Depending on the operating system this will be yum, apt, yast etc. The administrator can run this update manually, or add a boot script in the appliance template that carries out the update during first boot.

This allows the administrator to decide to use other configuration management platforms (Puppet, Chef, Ansible, Saltstack, BMC Bladelogic to name a few) to manage such updates. For some of these configuration platforms though, you will need to add a boot script as part of the appliance template to bootstrap the running instance with the configuration management platform.

Changing the OS Version of Middleware

Native packages, middleware and application software can be changed or swapped out; and the user can use the appliance update mechanism to determine if any package updates are available that can be applied prior to generating and migrating the workload.

Note: Changing the operating system for example from CentOS to Ubuntu is not supported.

For a list of supported OSes for Migration, see the table in [Supported Operating Systems](#).

Major OS versions, for example upgrading from CentOS 5.0 to CentOS 6.0 is not supported automatically, though as we have the complete list of operating system packages from the scan, a new appliance template can be constructed with the new operating system version.

This process can further be automated by using the command-line tool hammr (see hammr.io). Hammr allows you to create identical machine images from a single configuration file (in JSON or YAML). For more information refer to the [hammr documentation](#).

Qualification of any middleware and application software is strongly recommended.

Modifying the Scan Overlay

Note: This section only applies if you performed a scan with overlay.

When you import a scan as an appliance template, the overlay created as part of the scan process is registered as a My Software component. This My Software component is added to the appliance template.

The My Software component created from the overlay contains two archives. The first includes all the native package meta-data changes (permissions, ownership changes) and data changes (due to configuration modifications through the lifetime of the live machine). The second archive includes all files that are not part of any native package.

To modify a file in this overlay, you need to download, extract, modify and re-upload it to MySoftware once the changes have been made.

8.7.4 Comparing Scans

As scans are meta-data, you can compare two scans to determine their differences. This can be used to detect the differences between two live machines (for example between staging and production) or to detect changes of the live machine over time.


Warning: This is only supported for Linux based source instances.

There are two methods to compare scans:

- clicking the `compare` button (balance icon) at the top right hand side of the `Re-platform` page, then selecting the scans and clicking `compare`.



- selecting a specific scan and going to the `Delta` tab. By default the selected scan will be compared to a base machine. You can then modify the scan to compare to, from the `Compared to` section, at the top of the page.




Warning: The source and target scan must be of the same type. For example, if the source scan is a scan with overlay, then it must be compared to a scan with overlay. Similarly, a scan without overlay can only be compared to a scan without overlay.

UForge lists all the differences between the two systems on the `Delta` tab. The changes are color coded:

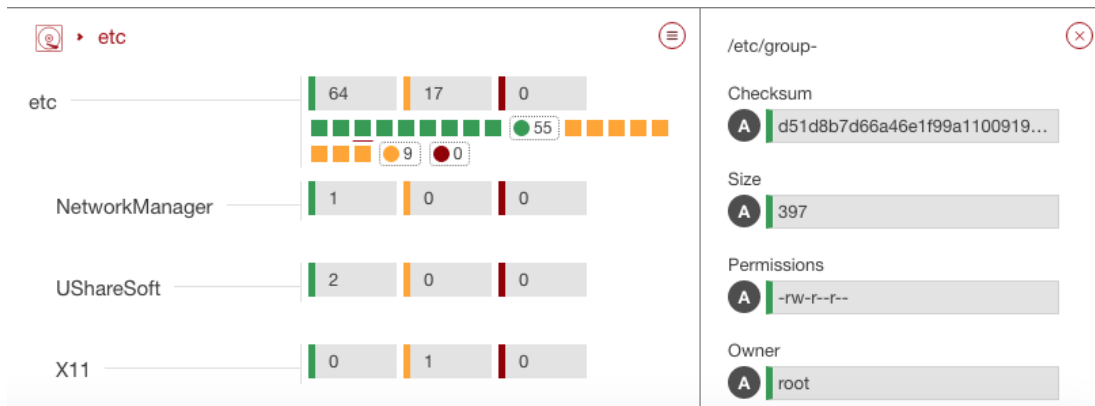
- green for added
- orange for modified
- red for deleted

The number shown corresponds to the number of files that are different between the two scans. When you drill down into a folder (as shown below), you can see the number of files modified depicted as clickable squares, while the number of files modified in sub-folders are indicated by a number next to the colored dot.

In the example below, there are a total of 64 added files in the `/etc` folder, with 55 of these new files being added to sub-folders of `/etc`.



If you click on any of the squares, you will see the details of the file that differs from the base scan.



You can also see the delta as a list.


The screenshot shows the delta view as a list of files. The table is as follows:

Name	Size	Modified	Owner:Group
.pwd.lock	0	Dec 3, 2018	root:root
.updated	163	Dec 3, 2018	root:root
group-	397	Dec 3, 2018	root:root
gshadow-	312	Dec 3, 2018	root:root
nsswitch.conf.bak	1735	Dec 3, 2018	root:root

Searching for Files in a Scan

The delta tab of the Re-platform page allows you to search and filter the scan compare results.

If you click on one of the categories of scan results (Added, Modified or Deleted), this will help you sort the results. In the example below, the Modified and Deleted categories have been clicked on to deselect, allowing you to see more quickly where the added files are.



If you are looking for a particular file name or type, then you can enter this in the search field and click on the search button to view the results.

8.7.5 Synchronize Target Environment with Scanned Machine

At the end of the migration process, you can synchronize the running migrated instance with the source machine data.

Warning: This feature is only available for scans without overlay for the following Operating Systems:

- CentOS 6, 7
- Red Hat Enterprise Linux 6, 7
- Debian 6, 7, 8
- Ubuntu 16.04, 14.04, 12.04
- openSUSE 42

From the scan detail view, you can get the instructions to synchronize your data by clicking the Synchronize button.

Note: `uforge-sync` is based on `rsync`, which does not handle BTRFS subvolume synchronization well. If you plan you synchronize `snapshots` directories from a BTRFS partition, it is better to exclude these directories from the scan and use another tool to handle synchronization, based on `btrfs send` and `btrfs receive`.

Download the uforge-sync Tool

You need the `uforge-sync` binary in the source machine, available at `/resources/uforge-sync.bin`, for example:

```
$ curl --insecure -O https://10.1.2.206/resources/uforge-sync.bin
```

Note: Older versions of `curl` command may not support a secure connection with UForge server. If you encounter a `SSL connect error`, you can use a web browser to download the binary or download the binary on another machine and

copy it to your machine.

Launch the uforge-sync Tool

Once the `uforge-sync` binary has been copied to the source machine, you can launch it:

```
$ chmod +x ./uforge-sync.bin ; ./uforge-sync.bin -U http://10.1.2.206/api -u guest \
-n users/guest/scannedinstances/1/scans/1 -d 11.12.13.14 -i ~/.ssh/id_rsa
```

The following are mandatory parameters:

- `-U`: the UForge server URL
- `-u`: the UForge user
- `-n`: the scan URI
- `-d`: the migrated instance address

Note: The usage of an SSH key pair is strongly recommended to connect to the migrated instance. You must specify the private key path using the `-i` parameter, and have the public key in the migrated instance's `~/.ssh/authorized_keys` file.

If you encounter an error message indicating the SSL certificate subject name does not match the target host name, you can add the `-k` option to your command line to prevent `uforge-sync.bin` from strictly checking the SSL certificate.

Note: The `uforge-sync` tool may overwrite existing configuration files of running services. Please restart services to load new configuration, if necessary.

Note: The `uforge-sync` tool self extracts in `/tmp` directory before execution. If the machine source or destination is configured with `noexec` option on `/tmp` partition a `Permission denied` error will occur. In that case, you can specify the environment variable `TMPDIR` on the source machine to use a different directory. For example:

```
TMPDIR=/root/ ./uforge-sync.bin -U https://<uforge-ip>/api -u <username> \
-n users/guest/scannedinstances/1/scans/1 -d <target-ip> -i ~/.ssh/id_rsa
```

8.8 Differences between Source and Target Instances

The goal of migration is to create an “identical” copy of the source system and run it on a new target environment. The word identical though is slightly misleading, as there are small differences that must be carried out to ensure the migrated system correctly boots and runs on the target environment. This section highlights these differences.

8.8.1 IP Address and MAC Addresses

When a live machine is scanned the networking is treated differently depending on whether the IP address is static or DHCP.

- **Static IP Addresses:** UForge keeps the information detected during the scan. During the generation phase, this networking information is also kept. Consequently, the new machine instance has the same static IP address set. If you need to change the static IP address, then you should refer to Advanced Migration procedure. By importing the scan report into an appliance template, you will have access to the `Install Profile` allowing you to change NIC settings prior to the generation process.
- **Dynamic (DHCP) IP Addresses:** In this case, the operating system resets the networking information at boot time. When the new machine instance is provisioned, the instance will send a request to the local cloud DHCP service to get a new IP address.

Consequently networking configuration files including `/etc/hosts` may be different between the source and target systems.

8.8.2 Network Device Name

Currently, the network device name is renamed after a migration using `Lift & Shift`. The new device name will be `eth0`, regardless of the name prior to the migration.

8.8.3 New Libraries Added

Depending on the source and target environments, some packages or libraries are added to ensure that the system boots correctly on the target environment:

- For migrations using `Re-platform` (see *Migrating Using Re-platform*): UForge injects `uforge-install-config` and `uforge-install-profile` packages. These packages provides features provided by templating to prompt users for some additional information during the initial booting of the system (root password, keyboard, timezone etc based on the `Install Profile` settings) and the execution of any boot scripts saved as part of the appliance template.
- `Para-virtualized Source -> Full Virtualized Target (Linux)`: UForge injects the `kernel` package and its dependent libraries to boot the system.
- `Source -> Azure (Linux)`: UForge injects the Windows Azure Linux agent `walinuxagent`
- `Source -> Azure (Windows)`: UForge injects `Windows VM Agent.msi`
- `Source -> OpenStack (Linux)`: UForge injects `cloud-init` package

8.8.4 System Clock `/etc/adjtime`

For Linux systems, the `/etc/adjtime` might be different between source and target systems. The Hardware Clock is usually not very accurate. However, much of its inaccuracy is completely predictable - it gains or loses the same amount of time every day. This is called systematic drift. The `/etc/adjtime` file keeps historical information on the clock's drift. Changing hardware environments may change the contents of this file when hardware clock is adjusted (in many cases the kernel will automatically adjust the hardware clock periodically).

8.8.5 Contents of `/etc/fstab` File

UForge regenerates the `/etc/fstab` file during machine image generation. Hardware and device mappings may differ between source and target environments. For example, `/dev/xvda` could be translated to `/dev/sda`.

File system options (general, dump and check options) are not supported by UForge. Consequently during a migration these values will be reset to their system defaults.

Furthermore, a `cdrom` mount point is added by UForge to permit some cloud platform tools to work correctly.

Some differences may arise between the scanned server and the generated template in the contents of the `/etc/fstab` file. In particular, lines pertaining to NFS (Network FileSystem) mounts are not kept during a migration, because the machine images resulting from the migration would have to be instantiated on the same network as the source machine, with similar network parameters, which cannot always be the case. This also avoids the situations where the inability to mount an NFS share would prevent the (generated) server from booting.

8.8.6 Contents of `/etc/shadow` File

In Linux, the `/etc/shadow` file is used to store user passwords (along with other related information) in an encrypted format, in a place that is only accessible to root and privileged users, for an enhanced level of security compared to the previously used `/etc/passwd` file.

UForge regenerates the `/etc/shadow` file during machine image generation. Differences may occur in the `/etc/shadow` file, between the source machine and the generated template, especially regarding the 3rd to 7th fields of each line. For complete transparency, here is the role of these fields. Such differences, should they occur, do not incur any security risk.

- 3rd field: number of days since Jan 1, 1970 when password was last changed.
- 4th field: the minimum number of days required between password changes.
- 5th field: the maximum number of days the password is valid (after which users are forced to change their password).
- 6th field: the number of days before password is to expire that users are warned that their password must be changed.
- 7th field: the number of days after password expires that account is disabled.

8.8.7 Service State

When generating a machine image from a scan, certain services are disabled or enabled depending on the target machine image being created. The following changes are common to all formats:

- `libvirtd` disabled
- `sshd` enabled
- rewrite grub configuration and `initramfs/initrd`

If present, `NetworkManager` is enabled, otherwise `network` is enabled. If the user requested a firewall, services `iptables` or `ip6tables` will be enabled. Otherwise, `iptables` or `ip6tables` are always disabled.

Note: `NetworkManager` is the name used by some operating systems which is equivalent to `network-manager`. The name `network` is used by some operating systems which is equivalent to `networking`.

EC2 AMI Image

- `hal` disabled
- `haddaemon` disabled
- `network` enabled
- `ip6tables` disabled
- `iptables` disabled

OpenStack Image

- hal disabled
- haddaemon disabled
- network enabled
- ip6tables disabled
- iptables disabled

Microsoft Azure Image

- network enabled

8.8.8 Red Hat vs CentOS Packages

In addition to the above, there are other minor differences between source and target instances after migration. If CentOS packages have been installed on a Red Hat server, with the same functionality and the same version number as equivalent existing Red Hat packages, then after migration, these CentOS packages will be replaced by their Red Hat counterparts, because the vendor information present in RPM packages is not accounted for during migration: only the package name and version number are.

Warning: Installing non-Red Hat packages on a Red Hat server, or replacing Red Hat packages by CentOS packages, is strongly discouraged, because it voids the warranty. Red Hat can refuse support requests on a modified machine.

8.8.9 GRUB2

GRUB2 is always re-installed when migrating a Linux system. Therefore, some files in the folder `/boot/grub2` might be different between source and target. This is especially true for `/boot/grub2/i386-pc/core.img` since UForge installs GRUB2, regardless of the source machine configuration, with the following modules: `biosdisk linux boot part_msdos lvm ext2 gettext xfs`.

Using Workspaces

UForge AppCenter allows users to create workspaces. Only users that have been given the right to create workspaces will be able to create and manage workspaces. If this is not the case, contact your UForge Administrator.

Under the `Collaboration` tab, users can create a number of workspaces, invite other users and share appliance templates with the members of their workspace. A collaboration workspace allows you to restrict the users that can see and use your appliance templates by inviting them to join your workspace. Only users who accept the invitation will be able to view your appliances.

The following sub-sections go into detail on how to manage workspaces:

9.1 Creating a Workspace

To create a workspace:

1. Go to the `Collaboration` tab in the left-hand sidebar.
2. Click on `create workspace` in the top right hand corner.
3. Enter the workspace name.
4. Select the organization from the drop-down menu.
5. Click `create workspace`.



The screenshot shows a web interface for creating a new workspace. At the top, there is a dark header bar with a hamburger menu icon on the left and a button labeled 'create workspace' with a red 'x' icon on the right. Below the header, the main content area is titled 'Create a New Workspace'. It contains two input fields: a text field labeled 'Name' with the value 'Dev Team' and a dropdown menu labeled 'Org' with the value 'UShareSoft'.

You can now invite users, share appliance templates, and post comments in the activity stream.

9.2 The Activity Stream

Each workspace has an Activity Stream. This is a log of the current actions being carried out by the members of the workspace (for example adding a new appliance template).

Members can also add comments to the workspace for other members of the workspace to see. For example, to ask questions or share information. Your comments will be posted in chronological order in the activity stream.

To add a comment in a workspace, do the following:

1. Go to the `Collaboration` tab.
2. If you have several workspaces, select the workspace you want from the drop-down menu at the top right.
3. Enter your comment in the square.
4. Click `post`.

You can also post comments on a specific appliance that is shared in the workspace. For more information, see [Adding a Comment to a Shared Appliance Template](#).


9.3 Managing Workspace Members

9.3.1 Inviting Members to Your Workspace

Once you have created a workspace, you can invite users to it. By inviting users, they will be able to view appliance templates you have added to the workspace, and depending upon their access rights share their appliance templates.

To invite users to your workspace:

1. Go to the `Members` page of the `Collaboration` tab.
2. Click on `invite` in the top right.



3. Enter the email addresses of the people you want to invite to your workspace. If the email you specify is not recognized by UForge AppCenter, you will be prompted to invite them to join the platform.
4. You can modify the invitation message. This message will be included in the email inviting members to join your workspace.
5. Click `invite`.

9.3.2 Managing Member Access Rights

Once a user has joined your workspace, you (or another workspace administrator) can modify their status. By default, when users accept your invitation and join your workspace they will be `collaborators`.

Members of a workspace are either:

- **Guest.** A guest can read and post to the activity stream, and import appliances into their private appliance library.
- **Collaborator.** The collaborator has the same basic rights as the Guest, but can also share appliances.
- **Administrator.** This is generally the user who has created the workspace. There must be at least one administrator in a workspace, though there can be more. The administrator can invite or delete members and is able to delete a workspace. The administrator has all the same basic rights as the collaborator.

9.3.3 Deleting a Member

If you are a workspace administrator, you can delete any member of the workspace simply by going to the `Members` page and clicking the small symbol to the right of the member's name and info.

9.4 Sharing an Appliance Template in a Workspace

The main purpose of the workspace is to share appliance templates with a group of users.

To share an appliance template:

1. Go to the `Collaboration` tab.
2. If you have several workspaces, select the workspace you want to add your appliance template to from the drop-down menu on the top right.
3. Go to the `Shared Templates` page.
4. Click on `share` in the top right. This lists your private appliance templates.
5. Select an appliance template from your list and click next arrow.
6. Add a description.
7. Click `share` to push the appliance template to the workspace.

Once an appliance template has been shared, other users in the workspace can import it into their private appliance library. They can then modify, use and share the appliance as they wish.

Warning: Any changes you make to your original appliance will NOT be propagated to the shared copy in the collaboration workspace. You will need to share the new version again for members of your workspace to see the changes. The older version will be overwritten.

9.5 Adding a Comment to a Shared Appliance Template

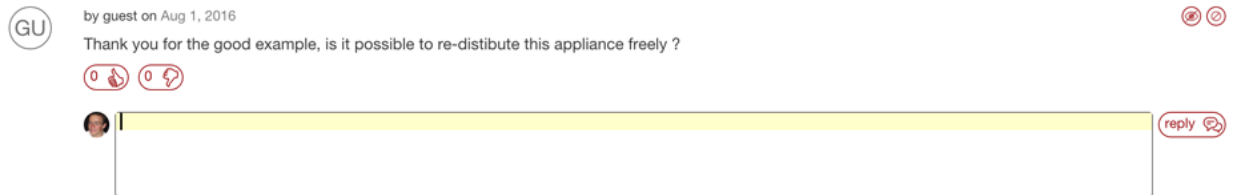
To add a comment to a specific appliance template in a workspace, do the following:

1. Go to the `Collaboration` tab.
2. If you have several workspaces, select the workspace you want from the drop-down menu on the top right.
3. Go to the `Shared Templates` page.
4. Select the appliance template you want to comment on.
5. Click on `Comments`.

6. Add your text.
7. Click post.



If you want to reply to a comment, enter your comment below the initial comment and click `reply`.



Using the REST API

UForge API is a RESTful resource.

The UForge API follows the design principles of Representational State Transfer (REST). UForge platform provides a set of resources (the API), each of which is referenced with a global identifier (URI). In order to manipulate these resources, clients communicate via the standard HTTP(S) protocol and exchange representations of these resources in a specific format. The documentation notes which formats are available for each method. The API presently supports XML and JSON. To get the results in the format of your choice, the `Accept-Type` header attribute is used in the request.

To make the UForge API easier to use, UForge has a Java SDK and Python SDK. You can create the API for other languages.

Communication with UForge is done via HTTP(S). For security reasons it is recommended to use HTTPS, however you may submit HTTP requests for debugging purposes.

- GET requests retrieve data
- POST requests create data
- PUT requests modify existing data
- DELETE request destroy existing data


API methods that require a particular HTTP method will return an error if you do not make your request with the correct one. All HTTP methods return codified response codes.

For a complete list of all the REST interface, refer to the [REST APIs Reference documentation](#).

10.1 Response & Error Codes

The UForge API returns typical HTTP status codes for every request received. HTTP status codes in the 200 range mean that the request was successful, while the 400 and 500 range indicates an error. The status codes in the 300 range are reserved for redirection.

Some of the error codes are returned by the web server container, while all the other response codes are returned by the UForge REST web service. Errors returned by the UForge REST web service may include a more detailed error message indicating why the request failed.



10.1.1 Success Codes

The following codes indicate a successful connection. The response may also include a body section providing the requested information. This information is in a MIME format specified as acceptable in the request.

- **200 OK:** Success. The request was fulfilled.
- **201 Created:** Following a POST command, a new resource has been created. The new resource URI is included in the response.
- **204 No Response:** Server has received the request but there is no information to send back. This is usually the case in a DELETE command, where a resource has been deleted
- **304 Not Modified:** Used when a client does a Conditional GET Request. If the document has not been modified since the date and time specified in If-Modified-Since field, the server responds with a 304 status code and does not send the document body to the client. The purpose of this feature is to allow efficient updates of local cache information (including relevant meta-information) without requiring the overhead of multiple HTTP requests (e.g. a HEAD followed by a GET) and minimizing the transmission of information already known by the requesting client (usually a caching proxy).

10.1.2 Error Codes

The 4xx codes are intended for cases where the client seems to have erred, and the 5xx codes for cases where the server identifies that the server has erred. It is impossible to distinguish these cases in general, so the difference is only informational. The UForge platform will attempt to provide a detailed error message to help the client diagnose the problem.

- **400 Bad Request:** The request has bad syntax or was inherently impossible to satisfy.
- **401 Unauthorized:** The request did not provide an acceptable authorization parameter. The client should retry the request with a suitable Authorization header.
- **403 Forbidden:** The client does not have the privileges to access this resource. Authorization will not help.
- **404 Not found:** The server did not find anything matching the resource provided in the request.
- **409 Conflict:** Following a POST command, if the resource being created already exists.

- 415 Unsupported Media Type: The server refuses to service the request because the entity of the request is in a format not supported by the requested resource for the requested method.
- 500 Internal Error: The server encountered an unexpected condition which prevented it from fulfilling the request.
- 502 Bad Gateway: UForge is down or being upgraded.
- 503 Service Unavailable: UForge is overloaded with requests. Try again later.

10.2 Sending a Request

The UForge Platform Services are all RESTful services, where clients communicate via the standard HTTP(S) protocol. That means you can easily construct request URLs that will work on the command line and in your code.

All UForge requests (with some exceptions) require authentication information as part of the request.

The UForge REST API uses a public and secret API key pair for authenticating each request. The public key is inserted as a query in the request URI. The secret key is then used to encode the entire URI to create a signature using HMAC_SHA1. This signature is then added to the end of the request URI.

Note that you can use Basic Authentication by adding an extra HTTP header `Authorization:Basic username:password`. However, this is less secure. We recommend this only be used on local area networks for instance.

All request URLs start with the hostname of where UForge is running, the port where UForge is listening for incoming requests, the service name and version number. This is known as the BASE URL. Such request URLs resemble the following sample:

```
https://myuforge.example.com:443/api
```

Even though UForge accepts HTTP requests, it is highly recommended for security reasons that HTTPS requests be used. HTTP requests should only be used for debugging purposes. Sensitive information will be exposed using HTTP.

10.2.1 The Request Headers

UForge expects certain headers containing authentication information be present as part of the URL request. UForge also accepts other header information, for example, to specify response content type and caching.

10.2.2 Request Example

The following is an example of a request sent to a UForge platform with hostname `10.0.0.20` using `cURL` to get the user `myUser`. Note that the response body (the user information) has been omitted here for clarity:

```
$ curl 'http://10.1.2.206/api/users/myUser?apiKey=XX8Bs2prKPdFrKH_
↪i4rsW7WR0f4FQ05IO7A8vuQUoNDino-7513mmEDecIAzpeMwWXZvnyZ6W0bJTKBwvc&
↪signature=3qDloxLwOI321BJlpDZ6Dzmqbac%3D' -H "Accept: application/xml" -v

* Trying 10.1.2.206...
* Connected to 10.1.2.206 (10.1.2.206) port 80 (#0)
> GET /api/users/guest HTTP/1.1
> Host: 10.1.2.206
> User-Agent: curl/7.42.1
> Accept: application/xml
```

(continues on next page)

(continued from previous page)

```
< HTTP/1.1 200 OK
< Date: Mon, 29 May 2017 14:28:19 GMT
< Server: Apache
< Last-Modified: Mon, 29 May 2017 08:49:24 GMT
< ETag: "837201f6b809de2aecedca4814e7a85e5"
< Content-Language: en
< version: 3.7.4-SNAPSHOT
< Content-Type: application/json
< Set-Cookie: JSESSIONID=708921B1F0C2AFA55262119F5E321FAF; Path=/ufws/; HttpOnly;
↪HttpOnly;Secure
< Connection: close
< Transfer-Encoding: chunked
```

The example illustrates the following:

- a GET request is sent (cURL by default uses GET) on the resource: `/users/myUser`
- an API key is used in this case for authorization
- the `Accept` header is being used to request that the response be sent in XML.
- the response header includes `ETag` and `Last-Modified` allowing cache validation and a conditional GET request.

10.3 Response Body Types

The API response types supported are XML or JSON. The `Accept` header is used in the request to determine which response type you would like.

- For JSON, use: `-H "Accept: application/json"`
- For XML, use: `-H "Accept: application/xml"`

If no accept header is used, then XML is returned by default.

For example, to retrieve the user information in JSON, the following request is used:

```
$ curl "http://10.1.2.206/api/users/guest" -X GET -u "guest:guest_password" -H
↪ "Accept: application/json"
```

Note: This example uses basic authentication (which is not advised). Furthermore, the response body i.e. the user information has been omitted here for clarity.

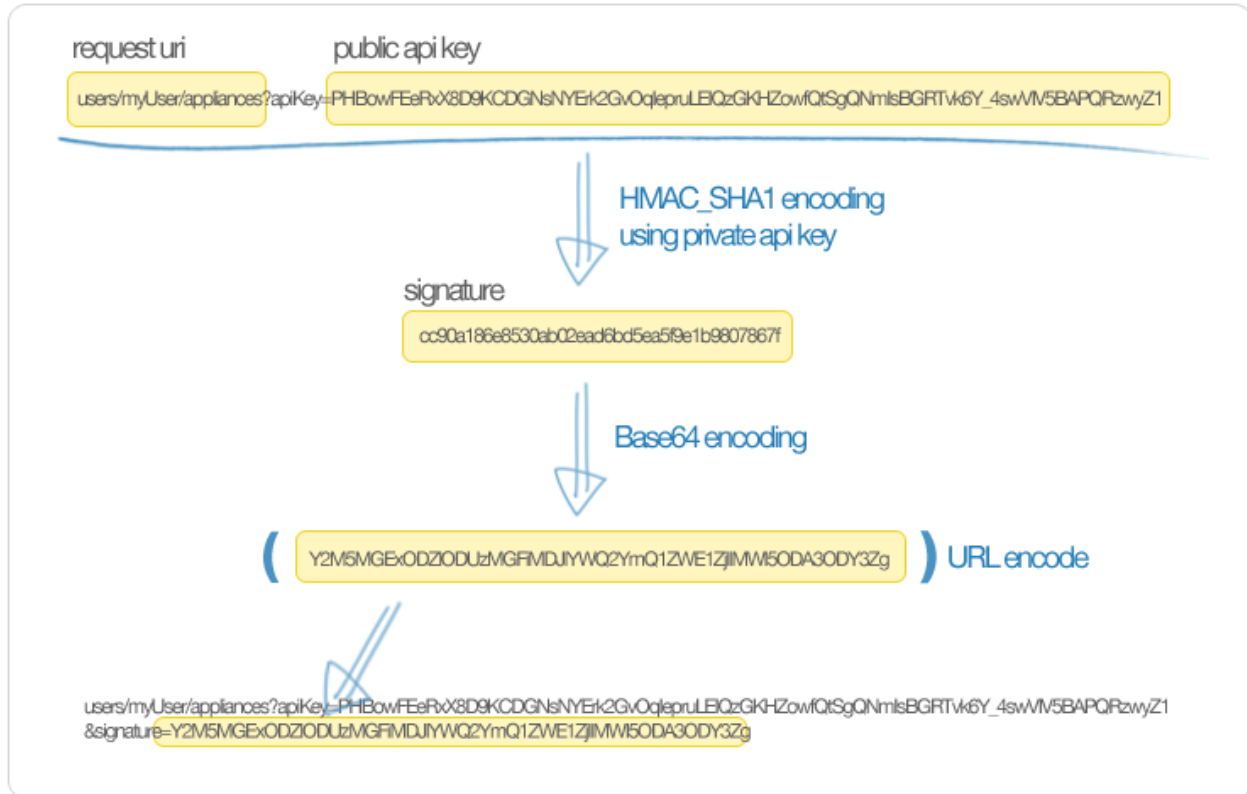
10.4 Using the API Keys

To use the UForge AppCenter APIs, it is recommended to use a public and secret API key as part of the request. This allows UForge AppCenter to correctly authenticate and authorize the request. API key pairs are managed on the `My Accounts` page under `API Key`. If you cannot see this tab, then you do not have the right to access UForge via the APIs. Contact your administrator for an initial API key pair.

The API keys are used inside and to sign each request URI to the UForge platform. The creation of a properly signed request URI is done in 5 steps:

1. Add the public API key to the end of the request URI with the query parameter `apiKey`.

2. Encrypt the request URI using HMAC_SHA1 with your secret API key to create a signature string.
3. Encode the signature string using Base64.
4. URL encode the signature string.
5. Add the signature to the end of the request URI created in step 1 with the query parameter `signature`.



10.5 Query Parameters

Certain resources within the UForge AppCenter API allow query parameters (or query strings) as part of the request URI. This allows you to pass extra parameters during search requests or to restrict the response data.

To pass a query parameter in a URL, the question mark symbol (?) is used as a separator. For example:

```
http://server/uripath?query_string
```

The query string is composed of one or more field-value pairs, each separated by the equals symbol (=). The series of field pairs is separated by the ampersand symbol (&). For example:

```
http://server/uripath?field1=value1&field2=value2&field3=value3
```

10.6 REST API Examples

For clarity, the following examples use basic authentication to communicate with UForge. As this is an insecure request, this is not recommended for production use. All response information is also omitted.

10.6.1 Retrieving a User

To retrieve a user's profile information use the API resource:

GET /users/{uid}

- uid: is the login of the user

Example:

```
$ curl "http://10.1.2.206/api/users/guest" -X GET -H "Authorization: Basic_
↪guest:guest_password" -H "Accept: application/xml" | tidy -xml -indent -quiet
```

References:

- API reference: [user_get](#)

10.6.2 Adding a Cloud Account

A cloud account is used to register machine images that have been generated from an appliance template. To create a cloud account use the API resource:

POST /users/{uid}/accounts

- uid: is the login of the user
- credAccount: CredAccount object you wish to create in the request body

Example:

```
$ curl "http://10.1.2.206/api/users/guest/accounts" -X POST -H "Authorization: Basic_
↪guest:guest_password" -H "Content-Type: application/xml" -H "Accept: application/xml
↪" --data-binary "@representation.xml" | tidy -xml -indent -quiet
```

The `representation.xml` content (the request body):

```
<ns0:credAccount
  xmlns:ns0="http://www.usharesoft.com/uforge"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:type="ns0:CredAccountOpenStack">
  <name>OpenStack John</name>
  <targetPlatform>
    <name>OpenStack</name>
    <type>openstack</type>
  </targetPlatform>
  <glanceUrl>http://ip:9292</glanceUrl>
  <keystoneUrl>http://ip:5000</keystoneUrl>
  <login>username</login>
  <password>password</password>
  <keystoneVersion>v3</keystoneVersion>
</ns0:credAccount>
```

References:

- API reference: [cloudAccount_create](#)
- [CredAccount](#) object

10.6.3 Get User Appliance Templates

To retrieve all the appliance templates for a particular user, use the following API resource:

GET /users/{uid}/appliances

- uid: is the login of the user

Example:

```
$ curl "http://10.1.2.206/api/users/guest/appliances" -X GET -H "Authorization: Basic_
↳ guest:guest_password" -H "Accept: application/xml" | tidy -xml -indent -quiet
```

References:

- API resource: [appliance_getAll](#)

10.6.4 Create an Appliance Template

An appliance template contains the model of the software stack. The model includes all the operating system packages, middleware and application software for generating an image that can be provisioned on a virtual or cloud platform. To create an appliance template, you need to decide which operating system to build the template from, as well as the name and version.

To create an appliance template, the following API resource is used:

POST /users/{uid}/appliances

- uid: is the login of the user
- appliance: Appliance template object you wish to create in the request body

Example:

```
$ curl "http://10.1.2.206/api/users/guest/appliances" -X POST -H "Authorization:_
↳ Basic guest:guest_password" -H "Content-Type: application/xml" -H "Accept:_
↳ application/xml" --data-binary "@representation.xml" | tidy -xml -indent -quiet
```

The `representation.xml` content (the request body):

```
<ns0:appliance
  xmlns:ns0="http://www.usharesoft.com/uforge">
  <name>My Appliance</name>
  <version>1.0</version>
  <description>Appliance created with UForge API</description>
  <distributionUri>distributions/1</distributionUri>
  <orgUri>orgs/1</orgUri>
</ns0:appliance>
```

References:

- API resource: [appliance_create](#)
- [Appliance object](#)

10.6.5 Retrieve Available OS Profiles for an Operating System

Each appliance template can contain an OS profile. This lists the operating system packages for the appliance template. To help users to create OS profiles for an appliance template, each Operating System registered in the UForge platform

has one or more base OS profiles (also known as OS templates) the user can choose from to get started. To retrieve the list of available OS profiles (or templates) for a particular operating system, use the following API resource:

GET /distributions/{id}/profiles

- id: is the id of the Operating System

Example:

```
$ curl "http://10.1.2.206/api/distributions/1/profiles" -X GET -H "Authorization:
↪Basic guest:guest_password" -H "Accept: application/xml" | tidy -xml -indent -quiet
```

References:

- API resource `osTemplate_getAll`
- OS Profile object

10.6.6 Add an OS Profile to an Appliance Template

To add an OS profile (template) to an appliance template, first list all the OS profiles for the operating system (see *Retrieve Available OS Profiles for an Operating System*) and note the name and URI of the profile you would like to add. Then use the following API resource to add this OS profile:

POST /users/{uid}/appliances/{aid}/osprofile

- uid: is the login of the user
- aid: the id of the appliance template where the os profile will be added
- OS Profile: OS profile to add to the appliance template (in the request body)

Example (Linux):

```
$ curl "http://10.1.2.206/api/appliances/1616/osprofile" -X POST -H "Authorization:
↪Basic guest:guest_password" -H "Content-Type: application/xml" -H "Accept:
↪application/xml" --data-binary "@distribprofile.xml" | tidy -xml -indent -quiet
```

The `distribprofile.xml` content (the request body):

```
<ns0:linuxProfile
  xmlns:ns0="http://www.usharesoft.com/uforge">
  <name>Minimal</name>
  <standardProfileUri>distribution/1/profiles/4</standardProfileUri>
</ns0:linuxProfile>
```

Example (Windows):

```
$ curl "http://10.1.2.206/api/appliances/1616/osprofile" -X POST -H "Authorization:
↪Basic guest:guest_password" -H "Content-Type: application/xml" -H "Accept:
↪application/xml" --data-binary "@distribprofile.xml" | tidy -xml -indent -quiet
```

The `distribprofile.xml` content (the request body):

```
<ns0:windowsProfile
  xmlns:ns0="http://www.usharesoft.com/uforge">
  <name>Standard 2012 English with agent</name>
  <standardProfileUri>distribution/9/profiles/27</standardProfileUri>
</ns0:windowsProfile>
```

References:

- API resource `applianceOSProfile_create`
- Appliance object
- Linux OS Profile object
- Windows OS Profile object

10.6.7 Searching for Operating System Packages (Linux Only)

Registered Operating Systems in UForge have package repositories attached to them (Linux only). You can search for packages in these repositories. These packages can then be added to an appliance's OS profile. To search for packages, the following API resource is used:

GET `/distributions/{id}/pkgs`

- `id`: is the id of the Operating System

This would retrieve all the packages for this operating system. To carry out a search, a `search_criteria` is used. This is based on FIQL, allowing you to search for packages based on different attribute values (including names, dates etc).

Example (retrieving all packages with the name `nginx`):

```
$ curl "http://10.1.2.206/api/distributions/1/pkgs&query=(name=='*nginx*'" -X GET -H
↪ "Authorization: Basic guest:guest_password" -H "Accept: application/xml" | tidy -
↪ xml -indent -quiet
```

Note: Wildcards can be used in FIQL queries.

References:

- API resource `osPkg_getAll`
- Package object

10.6.8 Adding a Package to an Appliance's OS Profile (Linux Only)

Once an Appliance Template has an OS profile, you can add or remove operating system packages to it. To add one or more packages, you will need the URI of the packages(s) you wish to add. You can [search for packages](#) to retrieve this information.

Use the following API resource to add or remove these packages to the OS profile of an appliance template:

PUT `/users/{uid}/appliances/{aid}/osprofile/{osid}/pkgs`

- `uid`: is the login of the user
- `aid`: the id of the appliance template
- `osid`: the id of the os profile to which the packages should be added

Example:

```
$ curl "http://10.1.2.206/api/appliances/1616/osprofile/2040/pkgs" -X PUT -H
↪ "Authorization: Basic guest:guest_password" -H "Accept: application/xml" --data-
↪ binary "@pkg.xml | tidy -xml -indent -quiet
```

The `pkg.xml` content (the request body):

```
<ns0:packages xmlns:ns0="http://www.usharesoft.com/uforge">
  <addedPkgUri>
    <uri>distributions/1/pkgs/631993</uri>
  </addedPkgUri>
</ns0:packages>
```

Removing the same package, the `pkg.xml` content would be (the request body):

```
<ns0:packages xmlns:ns0="http://www.usharesoft.com/uforge">
  <deletedPkgUri>
    <uri>distributions/1/pkgs/631993</uri>
  </deletedPkgUri>
</ns0:packages>
```

References:

- API resource `applianceOSProfilePkg_updateAll`
- Package object

10.6.9 Adding Custom Software to an Appliance Template

Software components can also be added to an appliance template. This is done in four stages:

1. A software component container is created. This is registered into your software library.
2. Create a software artifact container. This is the meta-data container for an uploaded file.
3. Upload the software files into this software artifact container.
4. Add the software component to an appliance template. Note that this software component can be added to multiple appliance templates.

To create the software component container, use the following API resource:

POST /users/{uid}/mysoftware

- `uid`: is the login of the user

Example:

```
$ curl "http://10.1.2.206/api/users/guest" -X POST -H "Authorization: Basic_
↪guest:guest_password" -H "Accept: application/xml" --data-binary "@software.xml |_
↪tidy -xml -indent -quiet
```

The `software.xml` content (the request body):

```
<ns0:mySoftware xmlns:ns0="http://www.usharesoft.com/uforge">
  <name>Zabbix</name>
  <version>3.0.1</version>
</ns0:mySoftware>
```

Once created, note down the `artifactsUri` of this software component. This is the URI we need to use to register one or more artifact objects.

Now we can create an artifact container. To do this use the `uri` of the software component. The resource API is:

POST /users/{uid}/mysoftware/{msid}/artifacts

- `uid`: is the login of the user
- `msid`: the id of the software component created

Example (uploading a RPM, but this can be any file type):

```
$ curl "http://10.1.2.206/api/users/guest/mysoftware/918/artifacts" -X POST -H
↪ "Authorization: Basic guest:guest_password" -H "Accept: application/xml" --data-
↪ binary "@artifact.xml | tidy -xml -indent -quiet
```

The artifact.xml content (the request body):

```
<ns0:softwareFile
  xmlns:ns0="http://www.usharesoft.com/uforge">
  <name>
    zabbix-release-3.0-1.el6.noarch.rpm
  </name>
  <fullName>
    zabbix-release-3.0-1.el6.noarch.rpm
  </fullName>
  <origName>
    zabbix-release-3.0-1.el6.noarch.rpm
  </origName>
  <subSoftwareArtifacts/>
</ns0:softwareFile>
```

Now you can upload the binary. Note down the `binaryUri` of the newly created artifact object. This is the resource uri you use to upload the file:

POST /users/{uid}/mysoftware/{msid}/artifacts/{said}/bin/{fileName}

- uid: is the login of the user
- msid: the id of the software component created
- said: the id of the software artifact
- fileName (optional): The filename to upload

Example (uploading a RPM, but this can be any file type):

```
$ curl "http://10.1.2.206/api/users/guest/mysoftware/918/artifacts/1078/bin/" -X POST
↪ -H "Authorization: Basic guest:guest_password" -H "Accept: application/xml" --data-
↪ binary "/path/to/file/zabbix-release-3.0-1.el6.noarch.rpm" | tidy -xml -indent -
↪ quiet
```

Finally you can now add this software component to an appliance template. The following resource API is used:

PUT /users/{uid}/appliances/{aid}

- uid: is the login of the user
- aid: the id of the appliance template

Example:

```
$ curl "http://10.1.2.206/api/users/guest/api/appliances/1616" -X PUT -H
↪ "Authorization: Basic guest:guest_password" -H "Accept: application/xml" --data-
↪ binary "@appliancesoftware.xml" | tidy -xml -indent -quiet
```

The appliancesoftware.xml content (the request body):

```
<ns0:appliance xmlns:ns0="http://www.usharesoft.com/uforge">
  <name>My Appliance</name>
  <version>1.0</version>
```

(continues on next page)

(continued from previous page)

```

<distributionUri>distributions/1</distributionUri>
<orgUri>orgs/1</orgUri>
<uri>users/guest/api/appliances/1616</uri>
<mySoftwareList>
  <mySoftware>
    <name>Zabbix</name>
    <version>3.0.1</version>
    <uri>users/guest/mysoftware/918</uri>
  </mySoftware>
</mySoftwareList>
</ns0:appliance>

```

References:

- Create software component resource `mySoftware_create`
- Create software artifact resource `mySoftwareArtifact_add`
- Upload a binary file `mySoftwareArtifact_upload`
- Add software component to an appliance template `appliance_update`
- Software Component object
- Software Artifact object
- Appliance object

10.6.10 Generate a Machine Image

Machine images can be generated from appliance templates by using the following API resource:

POST /users/{uid}/appliances/{aid}/images

- uid: is the login of the user
- aid: the id of the appliance template

Example (generating an OpenStack QCOW2 image):

```

$ curl "http://10.1.2.206/api/users/guest/api/appliances/1616/images" -X POST -H
↪ "Authorization: Basic guest:guest_password" -H "Accept: application/xml" --data-
↪ binary "@generateimage.xml" | tidy -xml -indent -quiet

```

The `generateimage.xml` content (the request body):

```

<ns0:image xmlns:ns0="http://www.usharesoft.com/uforge">
<compress>false</compress>
<targetFormat>
  <name>OpenStack QCOW2</name>
</targetFormat>
<installProfile>
  <memorySize>512</memorySize>
  <diskSize>2048</diskSize>
</installProfile>
</ns0:image>

```

References:

- Generate a machine image `machineImage_generate`

- [Machine Image object](#)

10.6.11 Publish/Register a Machine Image to a Cloud Environment

Once a machine image has been generated, for certain formats, this machine image can be published (also known as registered) to a corresponding Cloud environment.

To publish a machine image, the following API resource is used:

POST /users/{uid}/appliances/{aid}/images/{itid}/pimages

- uid: is the login of the user
- aid: the id of the appliance template
- itid: the id of the generated machine image

Example (publishing to OpenStack):

```
$ curl "http://10.1.2.206/api/users/guest/api/appliances/1616/images/346/pimages" -X_
↪POST -H "Authorization: Basic guest:guest_password" -H "Accept: application/xml" --
↪data-binary "@publishimage.xml" | tidy -xml -indent -quiet
```

The `publishimage.xml` content (the request body):

```
<ns0:publishImage xmlns:ns0="http://www.usharesoft.com/uforge"
xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xsi:type="ns0:PublishImageOpenStack">
  <applianceUri>users/guest/api/appliances/1616</applianceUri>
  <credAccount xsi:type="ns0:CredAccountOpenStack">
    <targetPlatform>
      <name>openstack</name>
    </targetPlatform>
    <glanceUrl>http://ip:9292</glanceUrl>
    <keystoneUrl>http://ip:5000</keystoneUrl>
    <login>username</login>
    <password>password</password>
    <keystoneVersion>v3</keystoneVersion>
  </credAccount>
  <imageUri>users/guest/api/appliances/1616/images/346</imageUri>
  <keystoneDomain>Keystone Domain Example</keystoneDomain>
  <keystoneProject>Keystone Project Example</keystoneProject>
  <displayName>Machine Image Name Example</displayName>
  <publicImage>false</publicImage>
</ns0:publishImage>
```

References:

- Publish a machine image [machineImage_publish](#)
- [Machine Image object](#)
- [Publish Machine Image object](#)

11.1 Download and Installing the SDK

To use the UForge AppCenter Java API, you need the jar files. You can either download the jar bundle and then add the dependency jar files to your classpath, or use maven to handle all the dependencies by adding the following to your `pom.xml` file:

```
<dependency>
  <groupId>com.usharesoft</groupId>
  <artifactId>uforge-client</artifactId>
  <version>3.6</version>
</dependency>

<repository>
  <id>maven.usharesoft.com</id>
  <name>UShareSoft Repository</name>
  <url>http://maven.usharesoft.com/repository/official/</url>
  <layout>default</layout>
</repository>
```

To use the API, you must have:

- The base URL of the UForge AppCenter platform you will be communicating with, for example: `https://factory.usharesoft.com:443`
- An account on the UForge AppCenter platform you will be communicating with
- An API key pair (secret and public) for the account (or use Basic Authentication)

11.2 Communicating with UForge

The `UForgeConnector` class provides all the lower level communication with UForge AppCenter by leveraging the `Jersey client API`. This class creates HTTP packets with the correct header information and constructs the request

URI to authenticate the request (using the secret and public keys). The response is parsed by JAXB to provide POJO Java classes of the response information.

Normally the first step is to get the user information of the account being used to authenticate. The response provides the URIs to the organizations, appliances and software this user has access to. The code below shows how to recuperate the user information. Note, as UForge is completely RESTful, when the method `login()` is used, no session is created between the client and UForge. Each request will reuse the authentication information stored in this `UForgeConnector` instance.

```
import com.usharesoft.client.common.connector.UForgeConnector;
import com.usharesoft.client.common.objects.User;
import org.apache.log4j.Logger;

private static String UFORGE_BASE_URI = "https://factory.usharesoft.com:443";
private static String USER_NAME = "myUser";
public static String SECRET_API_KEY = "b90240x-
↪N5jvPqzYw8IlfYvFRuSQv9sFuNM30gNAwZ4RqY5nOt2zPd8XyOS2hp3oKB09-XsYbNr_e4apR";

public static String PUBLIC_API_KEY =
↪"PHBowFEeRxX8D9KCDGNsNYErk2GvOqIepruLEIQzGKHZowfQtSgQNmIsBGRtVtk6Y_4swVIV5BAPQRzwyZ1
↪";

UForgeConnector connector = new UForgeConnector(UFORGE_BASE_URI, USER_NAME, PUBLIC_
↪API_KEY, SECRET_API_KEY);
User me = connector.login();
logger.info("Successfully connected to UForge. User: " + me.getLoginName());
```

11.3 Creating an Appliance Template

An Appliance Template contains the model of the software stack. The model includes all the operating system packages, middleware and application software for generating an image that can be provisioned on a virtual or cloud platform. To create an appliance template, you need to decide which operating system to build the template from, as well as the name and version.

When creating an appliance or choosing an operating system, you must choose the organization where to create the appliance or to search an operating system. The user must be a member of the organization to have authorization to search the organization resources. By default a user will be a member of at least one organization.

The following code provides an example of constructing an appliance template.

```
// require to have the organization information in the connector when getting
↪distributions
// and creating appliances
// get the first organization of the user
URI orgUri = me.getOrgUris().getUris().iterator().next();
OrgDelegate orgDelegate = new OrgDelegate(connector);
Org org = orgDelegate.get(orgUri);
connector.setOrg(org);

// Use the connector holding the user authentication information to get
// the operating system to use for constructing the appliance template
DistributionDelegate distributionDelegate = new DistributionDelegate(connector);
Distribution distribution = distributionDelegate.get(Distribution.CENTOS_5_6_I386);

// Create the appliance
ApplianceDelegate applianceDelegate = new ApplianceDelegate(connector);
```

(continues on next page)

(continued from previous page)

```
Appliance appliance = new Appliance(distribution, "WordPress", "3.2.1");
appliance = applianceDelegate.create(appliance);
```

11.4 Adding an OS Profile

An Appliance Template must contain an operating system profile. This profile contains a subset of operating system packages required by the middleware and application software to run correctly. Each operating system provided by UForge contains a set of standard operating system profile templates to choose from. These contain commonly used package bundles for the operating system to run, providing the basic operating system services.

The “Minimal” OS profile contains the minimum set of packages for the operating system to run properly and provide a minimum set of networking services and administration tools.

The following code shows how to create a new OS profile from a standard OS profile template and add it to an Appliance Template.

```
// Get the minimum template from the distribution
DistribProfileTemplate osProfileTemplate = distributionDelegate.
    ↪getProfile(distribution, DistribProfile.MINIMAL);

// Create the os profile that will be added to the appliance from the template
DistribProfile osProfile = new DistribProfile(osProfileTemplate);
osProfile.setDistribProfileTemplate(osProfileTemplate);

// Create this os profile in the appliance
ApplianceOSProfileDelegate aospDelegate = new ApplianceOSProfileDelegate(connector);
osProfile = aospDelegate.create(this.appliance, osProfile);
```

Extra packages can be added to the appliance template’s OS profile.

```
// Get the current list of packages in the os profile
PkgList pkgs = aospDelegate.getPkgs(osProfile);

// Add more packages
// php53
Pkg pkg = distributionDelegate.getPkg(distribution, "php53");
if (pkg == null) {
    logger.error("Unable to retrieve the package php53 for this distribution");
    return;
}
pkgs.addExtraPackage(pkg);

// php53-common
pkg = distributionDelegate.getPkg(distribution, "php53-common");
if (pkg == null) {
    logger.error("Unable to retrieve the package php53-common for this distribution");
    return;
}
pkgs.addExtraPackage(pkg);

// php53-cli
pkg = distributionDelegate.getPkg(distribution, "php53-cli");
if (pkg == null) {
    logger.error("Unable to retrieve the package php53-cli for this distribution");
```

(continues on next page)

(continued from previous page)

```

        return false;
    }
    pkgs.addExtraPackage(pkg);

    // php53-mysql
    pkg = distributionDelegate.getPkg(distribution, "php53-mysql");
    if (pkg == null) {
        logger.error("Unable to retrieve the package php53-mysql for this distribution");
        return false;
    }
    pkgs.addExtraPackage(pkg);

    // update the os profile with the new package list
    aospDelegate.updatePkgs(osProfile, pkgList);

```

11.5 Generating a Machine Image

Once you are happy with the contents of an appliance template, you can then generate a machine image to practically any hypervisor or cloud environment. The following code generates a CloudStack VHD image (for Xen hypervisor). For some image types you can select the disk size and the RAM of the virtual machine to be created. These can be updated once provisioned in the cloud environment. If you have set advanced partitioning in the installation profile, then this will be used instead for the disk size. The generation is done asynchronously; the generation status gives the progress of this generation.

```

ImageDelegate imageDelegate = new ImageDelegate(connector);
Image image = new Image(appliance, ImageFormat.CLOUDCOM_VHD_FORMAT);
image.setCompress(true); // create an archive (.gz)
image.setVmDiskSize(4096); // 4GB
image.setVmMemorySize(256); // 256MB

// Launch the generation
image = imageDelegate.generate(appliance, image);

// Check the generation status every 5 seconds
OpStatus status = imageDelegate.getStatus(image);
while ( status.isComplete() == false ) {
    try {
        Thread.sleep(5000);
    } catch (InterruptedException ex) {
        //Error in a thread while trying to get the status of the cloud.com image_
        ↪ generation
    }
    // Get the status
    status = imageDelegate.getStatus(image);
}

// Generation complete!

```

11.6 Publishing an Image

UForge has connectors to many of the popular cloud platforms including Amazon, Microsoft Azure, Google Compute Engine, OpenStack and CloudStack to name a few. Once an image has been generated you can either download

the image or publish directly to a cloud environment using your own cloud account credentials. Like generations, publishing images is asynchronous. You can get the progress of the publish from the publish status. The following code publishes a generated VHD image directly to the template library of a CloudStack environment.

```
// Create the credential information to communicate with the Cloud.com environment
CredAccount credAccount = new CredAccount();
credAccount.setType(InfraType.CLOUD_COM);
credAccount.setName("My Cloud.Com Account");
credAccount.setServerUrl(new URI("http://10.0.1.251:8080/client/api"));
credAccount.setPublicAPIKey("8pqgg0HV8ocpt6j8qYiCpDZ4cqzbtLaxCErIOpCD0r9VOjnILgahX85_
↪J2CFvC8863en3NGZEFiLY82PKbAfwQ");
credAccount.setSecretAPIKey("9Q-vVxokmMbI_
↪l4t7aAfbocTgoLBlnt4lXy6iLZUfC6PzAdXNy2rRegAWhBMF3mQ9jk4MtPaCVULDu7ORxX2ZQ");

// Add the zone on where the image should be published
credAccount.setZoneName("zone1");

// Provide information on the image being uploaded to the Cloud.com template library
credAccount.setDisplayName("WordPress image");
credAccount.setDescription("WordPress image for the cloud.com platform");
credAccount.setPasswordEnabled(true);
credAccount.setFeaturedEnabled(false);

// Allow this image to be accessed to all cloud.com users
credAccount.setPublicImage(true);

// Publish the image
PublishDelegate publishDelegate = new PublishDelegate(connector);
PublishImage publishImage = new PublishImage();
publishImage.setCredAccount(credAccount);
publishImage.setImage(image);
publishImage = publishDelegate.publish(appliance, publishImage);

// Get the status of the publish
OpStatus status = publishDelegate.getStatus(publishImage);
while (status.isComplete() == false) {
    try {
        Thread.sleep(5000);
    } catch (InterruptedException ex) {
        // Error in a thread while trying to get the status of the iso image_
↪generation
    }
    // get the status
    status = publishDelegate.getStatus(publishImage);
}

// Publish finished
// Can check for errors
if (status.isError()) {
    // error occurred during the publish
}
```

11.7 Adding a Project from the Project Catalog

Each UForge organization provides a Project Catalog where commonly used software components can be added to an appliance template. The Project Catalog must belong to the same organization as the appliance. The following code

adds some projects to an appliance template.

```
ProjectDelegate projectDelegate = new ProjectDelegate(connector);

// Add Apache HTTP Server
Project project = projectDelegate.get(distribution, "Apache HTTP server", "2.2.3",
↪null);
if (project == null) {
    logger.error("Unable to retrieve project");
    return;
}
appliance.addProject(project);

// Add MySQL
Project project = projectDelegate.get(distribution, "MySQL5 Server", "5.0.77", null);
if (project == null) {
    logger.error("Unable to retrieve project");
    return;
}
appliance.addProject(project);

// Save the updated appliance instance
applianceDelegate.update(appliance);
```

11.8 Uploading a Software Component

You can upload your own software components to a private software library (My Software Library). This software can then be added to any of your appliance templates. This provides a mechanism to compliment the Project Catalog. The following code shows how to upload files, attach a license and then add it to an appliance template.

```
MySoftwareDelegate mySoftwareDelegate = new MySoftwareDelegate(connector);
File wpf = new File("wordpress-3.2.1.zip");
File lf = new File("wp-license.html");

// Create a My Software component and upload the wordpress zip file
MySoftware mySoftware = mySoftwareDelegate.upload("WordPress", "3.2.1", wpf);

// Attach a license
mySoftware = mySoftwareDelegate.uploadLicense(mySoftware, lf);

// Add the software component to the appliance
appliance.addMySoftware(mySoftware);

// Save the changes to the appliance
applianceDelegate.update(appliance);
```

11.9 Adding a Boot Script

Boot scripts can be added to the appliance template allowing initial configuration to be executed either during the first time the image is started or during every boot of the image.

The following code shows how to upload a boot script to an appliance.


```
ApplianceConfigDelegate configDelegate = new ApplianceConfigDelegate(connector);

// upload a boot script to the appliance
File bsf = new File("myscript.sh");
BootScript bootscript = new BootScript();
// only execute this boot script once during first boot
bootscript.setBootType(BootScript.FIRST_BOOT);
bootscript = configDelegate.uploadBootScript(appliance, bootscript, bsf);
```


12.1 Download and Installing the SDK

The Python API is supported on all major operating systems: Linux, Mac-OS, and Windows. The easiest way to install the API is using `pip`, the widely used package management system for installing and managing software packages written in Python.

Note: This API is used by the open source Hammr project. You can find more examples on how to use the API in the source code of this project.

12.1.1 Installing pip

If you already have `pip` installed on your system, you can skip this section.

To install or upgrade `pip`, download this file:

```
https://raw.githubusercontent.com/pypa/pip/master/contrib/get-pip.py
```

Then run the command:

```
$ python get-pip.py
```

For more information on installing `pip`, refer to the official `pip` documentation: <http://www.pip-installer.org/en/latest/installing.html>

12.1.2 Installing UForge Python API

Once `pip` has been installed, you can now install the UForge Python API packages. You may have to run this command as `sudo` or administrator.

See below the instructions for installing the Python API for your target platform:

For Linux

First, you need to install extra packages on your system, and then install the SDK package.

Debian

```
$ apt-get install python-dev gcc
$ pip install uforge-python-sdk
```

Red Hat and CentOS

```
$ yum install gcc python-devel
$ pip install uforge-python-sdk
```

Upgrading the SDK

To upgrade an already installed SDK, run:

```
$ pip install --upgrade uforge-python-sdk
```

For Mac OS

As a pre-requisite, you need to have XCode installed.

To install the SDK run:

```
$ export ARCHFLAGS="-Wno-error=unused-command-line-argument-hard-error-in-future"
$ pip install uforge-python-sdk
```

Upgrading the SDK

To upgrade an already installed SDK, run:

```
$ pip install --upgrade uforge-python-sdk
```

For Microsoft Windows

Run the following command:

```
c:\Python27> .\Scripts\easy_install.exe uforge-python-sdk
```

Upgrading the SDK

To upgrade an already installed SDK, run:

```
c:\Python27> .\Scripts\easy_install.exe --upgrade uforge-python-sdk
```

12.2 Communicating with UForge

The Python API provides all the lower level communication with UForge AppCenter by creating HTTP request packets with the header information to authenticate correctly.

Normally the first step is to get the user information of the account being used to authenticate. The response provides the URIs to the organizations, appliances and software this user has access to. The code below shows how to recuperate the user information. Note, as UForge is completely RESTful, when the method `login()` is used, no session is created between the client and UForge. Each request will reuse the authentication information stored in this `api` instance.

```
# Import the Uforge python API
from uforge.application import Api

login='username'
passwd='password'

# Create the API object
api = Api(url = 'https://uforge.example.com/api',
          username = login, password = passwd,
          disable_ssl_certificate_validation = True)

# Send a request (getting the user object)
user = api.Users(login).Get()
```

Note: All UForge Python objects are in the file `objects.py`, contained in the Python SDK. You can find all attribute names for each object type.

12.3 Creating an Appliance Template

An Appliance Template contains the model of the software stack. The model includes all the operating system packages, middleware and application software for generating an image that can be provisioned on a virtual or cloud platform. To create an appliance template, you need to decide which operating system to construct the template from, as well as the name and version.

When creating an appliance or choosing an operating system, you must choose the organization where to create the appliance or to search an operating system. The user must be a member of the organization to search the organization resources.

The following code provides an example for constructing an appliance template. You will then need to add an OS profile to the template. Refer to [Adding an OS Profile to the Appliance](#).

```
# Import the Uforge python API
from uforge.application import Api
from uforge.objects.uforge import *

login='username'
passwd='password'
```

(continues on next page)

(continued from previous page)

```

# Create the API object
api = Api(url = 'https://uforge.example.com/api',
          username = login, password = passwd,
          disable_ssl_certificate_validation = True)

# all Orgs
newOrg = api.Users(login).Orgs().Getall()

# all Distributions
newDistribution = api.Users(login).Distros.Getall()

newAppliance = appliance()
newAppliance.name = "apptestpythonsdk"
newAppliance.version = "2.3"
# Let's use first organization
newAppliance.orgUri = newOrg.orgs.org[0].uri
# Let's use first distribution
newAppliance.distributionUri = newDistribution.distributions.distribution[0].uri

# create appliance
try:
    createdAppliance = api.Users(login).Appliances.Create(newAppliance)
except Exception as e:
    print str(e.args[0].statusCode) + " " + e.args[0].localizedErrorMsg.message
    sys.exit(1)
print "Created appliance: " + createdAppliance.uri

```

12.3.1 Adding an OS Profile to the Appliance

An Appliance Template must contain an operating system profile. This profile contains a subset of operating system packages required by the middleware an application software to run correctly. Each operating system provided by UForge contains a set of standard operating system profile templates to choose from. These contain commonly used package bundles for the operating system to run, providing the basic operating system services.

The “Minimal” OS profile contains the minimum set of packages for the operating system to run properly and provide a minimum set of networking services and administration tools.

The following code shows how to create a new OS profile from a standard OS profile template and add it to an Appliance Template. The Appliance template must already be created, as described in *Creating an Appliance Template*.

```

# List available osprofiles for this distribution to apply on appliance
newProfile = api.Distributions(newDistribution.distributions.distribution[0].dbId).
    ↳Profiles.Getall()
newProfile = newProfile.distribProfileTemplates.distribProfileTemplate
print "Will set profile: " + newProfile[0].name + " - " + newProfile[0].uri

newApplianceOSProfile = distribProfile()
newApplianceOSProfile.standardProfileUri = newProfile[0].uri

# apply os profile
try:
    api.Users(login).Appliances(createdAppliance.dbId).Osprofile().
    ↳Create(newApplianceOSProfile)
except Exception as e:

```

(continues on next page)

(continued from previous page)

```

    print str(e.args[0].statusCode)+" "+e.args[0].localizedErrorMsg.message
    sys.exit(1)
print "Profile applied."

```

12.3.2 Generating a Machine Image

Once you are happy with the contents of an appliance template, you can then generate a machine image to practically any hypervisor or cloud environment. The following code generates a CloudStack VHD image (for Xen hypervisor). For some image types you can select the disk size and the RAM of the virtual machine to be created. These can be updated once provisioned in the cloud environment. If you have set advanced partitioning in the installation profile, then this will be used instead for the disk size. The generation is done asynchronously.

```

# all target formats, search for VirtualBox
newTargetFormat = api.Users(login).Targetformats.Getall()
for i in newTargetFormat.targetFormats.targetFormat:
    if i.name == "VirtualBox":
        print "VirtualBox target format found for user " + login + "."
        break

if i.name != "VirtualBox":
    print "VirtualBox target format not found for user " + login + ", will not
    ↪generate."
    sys.exit(1)

newImage = image()
newImage.compress = "true"
newTargetFormat = targetFormat()
newTargetFormat.name = "VirtualBox"
newImage.targetFormat = newTargetFormat
newInstallProfile = installProfile()
newInstallProfile.memorySize = 512
newImage.installProfile = newInstallProfile

# generate target Virtualbox
try:
    api.Users(login).Appliances(createdAppliance.dbId).Images().Generate(newImage)
except Exception as e:
    print str(e.args[0].statusCode)+" "+e.args[0].localizedErrorMsg.message
    sys.exit(1)
print "Launched generation of appliance for VirtualBox target format."

```

12.4 Creating My Software

In addition to projects and OSes, you can add your own personal software to an appliance. In order to do this, you must create a My Software container, add the packages (refer to [Adding a Package to My Software](#)), and (optionally) add a license (refer to [Uploading a License to My Software](#)). Once you have created the My Software, you can then add it to an appliance.

```

# Import the Uforge python API
from uforge.application import Api
from uforge.objects.uforge import *

```

(continues on next page)

(continued from previous page)

```

import os

login='username'
passwd='password'

# Create the API object
api = Api(url = 'https://uforge.example.com/api',
          username = login, password = passwd,
          disable_ssl_certificate_validation = True)

newMySoftware = mySoftware()
newMySoftware.name = "newmstest"
newMySoftware.version = "5.6"
newMySoftware.description = "This is the description."

try:
    createdMySoftware = api.Users(login).Mysoftware().Create(newMySoftware)
except Exception as e:
    print str(e.args[0].statusCode)+" "+e.args[0].localizedErrorMsg.message
    sys.exit(1)
print "Created mysoftware: " + createdMySoftware.uri

```

12.4.1 Adding a Package to My Software

Once you have created the My Software container, you can add packages (files) using the following code. For a list of supported file formats, refer to the list provided in *Adding Software from Your Software Library*.

```

myNewPackage = package()
myNewPackage.origName = "/etc/redhat-release"

try:
    myNewPackage.size = os.stat(myNewPackage.origName).st_size
except Exception as e:
    print "Problem reading file " + myNewPackage.origName + ". Will not be added_
↳to mysoftware."
    sys.exit(1)

# add package file to mysoftware
try:
    api.Users(login).Mysoftware(createdMySoftware.dbId).Pkgs().Add(myNewPackage)
except Exception as e:
    print str(e.args[0].statusCode)+" "+e.args[0].localizedErrorMsg.message
    sys.exit(1)
print "Added package file " + myNewPackage.origName + " to mysoftware."

```

12.4.2 Uploading a License to My Software

If you want to add a license file to your software (optional), add the following code.

```

try:
    myLicenseFile = open("/etc/redhat-release")
except Exception as e:
    print "Problem reading file " + myLicenseFile.name + ". Will not add license_
↳to mysoftware."

```

(continues on next page)

(continued from previous page)

```

        sys.exit(1)

# add license file to mysoftware
try:
    api.Users(login).Mysoftware(createdMySoftware.dbId).
    ↳ Licenses(createdMySoftware.license.dbId).Uploadfile(myLicenseFile)
except Exception as e:
    print str(e.args[0].statusCode) + "    " + e.args[0].localizedErrorMsg.message
    myLicenseFile.close()
    sys.exit(1)
print "Added license file " + myLicenseFile.name + " to mysoftware."

myLicenseFile.close()

```

12.4.3 Listing All My Software

You can list all your private software for your account as follows:

```

print "Listing all mysoftware:"
# all MySoftware
newMySoftware = api.Users(login).Mysoftware().Getall()
for i in newMySoftware.mySoftwareList.mySoftware:
    print " ID " + str(i.dbId) + "    " + i.name + " v" + i.version + "    " + ("",
    ↳ "IMPORTED")[i.imported == True]

```

12.5 Creating a Project for a Specific OS

You can also add third-party software to an appliance using projects. The catalog of projects is public to all users on UForge and is maintained by the privileged users and administrators.

```

# Import the Uforge python API
from uforge.application import Api
from uforge.objects.uforge import *

login='username'
passwd='password'

# Create the API object
api = Api(url = 'https://uforge.example.com/api',
          username = login, password = passwd,
          disable_ssl_certificate_validation = True)

# all Distributions
newDistribution = api.Users(login).Distros.Getall()

# all Organisations
newOrg = api.Users(login).Orgs().Getall()

# create a project for 1st distribution and 1st organisation
newMyProject = project()
newMyProject.name = "testnewproject"
newMyProject.version = "1.9"

```

(continues on next page)

(continued from previous page)

```

newMyProject.release = "14"
newMyProject.category = "Blogging"
newMyProject.shortTag = "INTERNAL"
newMyProject.description = "test new project description"
newMyProject.company = company()
newMyProject.company.name = "MyCompany"
newMyProject.license = license()
newMyProject.license.type = "Custom"
newMyProject.distributionUri = newDistribution.distributions.distribution[0].uri

try:
    createdproject = api.Orgs(newOrg.orgs.org[0].dbId).Projects.
    ↪Create(newMyProject)
except Exception as e:
    print str(e.args[0].statusCode)+" "+e.args[0].localizedErrorMsg.message
    sys.exit(1)
print "Created project: " + createdproject.uri

```

12.5.1 Listing all Projects for a Specific OS

```

# listing all projects for same distribution
print "Listing projects for distribution " + newDistribution.distributions.
    ↪distribution[0].name + " " + \
        newDistribution.distributions.distribution[0].version + " " + \
        newDistribution.distributions.distribution[0].arch + ":"
print "-v-----"
newProjects = api.Distributions(newDistribution.distributions.distribution[0].dbId).
    ↪Projects.Getall()
for i in newProjects.projects.project:
    print "++ " + i.uri + " - " + "\033[1m\033[93m" + i.name + " v" + i.version
    ↪+ " r" + \
        str(i.release) + "\033[0m" + "\t(tag: " + i.shortTag + ") (size: +" +
    ↪str(i.size/1024/1024) + "MB)"
    print (" |", " ")[i == newProjects.projects.project[-1]] + "    path: " + i.
    ↪defaultInstallLocation
    print (" |", " ")[i == newProjects.projects.project[-1]] + "    category: " +
    ↪i.category + \
        "                               maintainer: " + i.company.name
    if i.description:
        print (" |", " ")[i == newProjects.projects.project[-1]] + "
    ↪description: " + i.description

```

12.5.2 Listing Target Formats and Target Platforms

```

# Import the Uforge python API
from uforge.application import Api
from uforge.objects.uforge import *

login='username'
passwd='password'

# Create the API object

```

(continues on next page)

(continued from previous page)

```

api = Api(url = 'https://uforge.example.com/api',
          username = login, password = passwd,
          disable_ssl_certificate_validation = True)

# all Organisations
newOrg = api.Users(login).Orgs().Getall()

targetformatswithtargetplatform = []
# all target platforms & target formats
allTargetPlatforms = api.Orgs(newOrg.orgs.org[0].dbId).Targetplatforms().Getall()
for i in allTargetPlatforms.targetPlatforms.targetPlatform:
    print i.uri + " \033[1m\033[93m" + i.name + ":\033[0m" + \
          " (type " + i.type + ") (" + \
          ("not active, ","active, ")[i.active] + \
          ("no access)","access")[i.access]

    allTargetFormats = api.Orgs(newOrg.orgs.org[0].dbId).Targetplatforms(i.dbId).
    ↪Targetformats().Getallformats()
    for u in allTargetFormats.targetFormats.targetFormat:
        targetformatswithtargetplatform.append(u.uri)
        print "          " + u.uri + " : \033[96m" + u.name + "\033[0m (type " + u.
    ↪type + ") (" + \
          ("not active, ","active, ")[u.active] + \
          ("no access)","access")[u.access]

print "\033[1m\033[93mNo target platform:\033[0m"
allTargetFormats = api.Orgs(newOrg.orgs.org[0].dbId).Targetformats().Getall()
for u in allTargetFormats.targetFormats.targetFormat:
    if u.uri not in targetformatswithtargetplatform:
        print "          " + u.uri + " : \033[96m" + u.name + "\033[0m (type " + u.
    ↪type + ") (" + \
          ("not active, ","active, ")[u.active] + \
          ("no access)","access")[u.access]

```


CHAPTER 13

Hammr Command Line Tool

UForge AppCenter provides an open source command-line tool called `hammr` to allow users to create machine images for different environments from a single configuration file.

Build consistent and repeatable machine images for :



Hammr is a lightweight client-side tool based on Python, and can be installed on all major operating systems.

Hammr can be used as part of your “DevOps tool chain” and in conjunction with other tools such as Jenkins, Chef, Puppet and SaltStack, allowing you to easily build your machine images and maintain your live running instances. Hammr also has migration capabilities, allowing you to scan a live system, generate a machine image for a different environment as well as export it back to a configuration file for sharing.

13.1 Getting the Source Code

All the source code is available on [GitHub](#).

13.2 Further Reading

All the documentation on installing and using hammr can be found on: <http://hammr.io>

You can manage UForge AppCenter using the UForge command-line interface (CLI).

14.1 Installing UForge CLI

The UForge CLI is available for CentOS 7 clients.

In order to install the version of the UForge CLI which corresponds to your UForge AppCenter, please ask your administrator to provide you with the `/etc/yum.repos.d/uforge-ee-centos.repo` file of the UForge AppCenter installation.

Then, run the following commands:

To install the UForge CLI run:

```
yum clean all
yum install uforge-cli
```

14.2 Launching UForge CLI

To launch UForge CLI locally on the machine, launch:

```
/opt/UShareSoft/uforge/cli/bin/uforge -u LOGIN_NAME -p PASSWORD -U http://IP-OF-
↪WEBSERVICE-NODE/api
```

The `-p` argument is optional. If the password is not entered in the command, the user will be prompted to enter their password.

Use `--help` to get a list of available commands.

Note: When accessing from outside the machine, use: <https://IP-OF-WEBSERVICE-NODE/api>

15.1 Trademarks

UForge is a registered trademark of UShareSoft, a Fujitsu company.

LINUX is a registered trademark of Linus Torvalds.

Microsoft and Windows are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.

Oracle, Java, and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle Corporation and/or its affiliates.

Apache Ant, Ant, and Apache are trademarks of The Apache Software Foundation.

Red Hat Enterprise Linux is a trademark of Red Hat.

UNIX is a registered trademark of the Open Group in the United States and in other countries.

Other company names and product names are trademarks or registered trademarks of their respective owners.

15.2 Copyright FUJITSU LIMITED 2019

All rights reserved, including those of translation into other languages. No part of this manual may be reproduced in any form whatsoever without the written permission of FUJITSU LIMITED

15.3 3rd Party Components Used

- Ace: <http://ace.c9.io/>
- Ace-gwt: <https://github.com/daveho/AceGWT>
- Angular: <https://angular.io/>

- Angular JS: <https://angularjs.org/>
- Angular Resize Event: <https://github.com/vdolek/angular-resize-event>
- Apache Brooklyn: <https://brooklyn.apache.org/>
- Apache Commons: <http://commons.apache.org/>
- Apache HttpComponents: <http://hc.apache.org/>
- Apache Syncope: <https://syncope.apache.org/>
- Apache Tomcat: <http://tomcat.apache.org/>
- Apt libraries: <https://packages.debian.org/wheezy/apt>
- argparse: <https://pypi.python.org/pypi/argparse/>
- AWS SDK: <http://aws.amazon.com/sdk-for-java/>
- Azure SDK: <https://github.com/Azure/azure-sdk-for-java>
- bzip2-libs: <http://www.bzip.org>
- CentOS: <https://www.centos.org/>
- cmd2: <https://pypi.python.org/pypi/cmd2/>
- cyrus-sasl-lib: <http://asg.web.cmu.edu/sasl/sasl-library.html>
- device-mapper-event-libs: <http://sources.redhat.com/lvm2>
- device-mapper-libs: <http://sources.redhat.com/lvm2>
- dnf: <http://dnf.baseurl.org/>
- docker-distribution: <https://github.com/docker/distribution>
- Dozer: <http://dozer.sourceforge.net/>
- D3: <http://d3js.org/>
- EclipseLink: <http://www.eclipse.org/eclipselink/>
- expat: <http://www.libexpat.org/>
- Filebeat: <https://www.elastic.co/products/beats/filebeat>
- fontconfig: <http://fontconfig.org>
- glibc: <http://www.gnu.org/software/libc/>
- GWT: <http://www.gwtproject.org/>
- GWTD3: <https://github.com/gwtd3/gwt-d3>
- Gwt-log: <http://code.google.com/p/gwt-log/>
- Gwt-reflector: <https://github.com/sykesd/gwt-reflector>
- Gwt-Storage: <https://github.com/seanchenxi/gwt-storage>
- Gwt-visualization: <https://code.google.com/p/gwt-google-apis/wiki/VisualizationGettingStarted>
- Hibernate: <http://hibernate.org/>
- Hibernate c3p0: <http://www.mchange.com/projects/c3p0/>
- hivex: <http://libguestfs.org/>
- HtmlCompressor: <http://code.google.com/p/htmlcompressor/>

- `httplib2`: <https://pypi.python.org/pypi/httplib2/>
- `hurry.filesize`: <https://pypi.python.org/pypi/hurry.filesize/>
- `ImageMagic`: <http://www.imagemagick.org/>
- `Jackson`: <https://github.com/FasterXML/jackson>
- `Jackson-dataformat-xml`: <https://github.com/FasterXML/jackson-dataformat-xml>
- `Java MariaDb Client Library (JDBC)`: <https://mariadb.com/products/connectors-plugins>
- `Jawr`: <https://jawr.java.net/>
- `jclouds`: <https://jclouds.apache.org/>
- `Jersey`: <https://jersey.github.io/>
- `Jersey server`: <https://jersey.java.net/>
- `Jiruka RSQL`: <https://github.com/jirutka/rsql-parser>
- `junit-xml`: <https://pypi.python.org/pypi/junit-xml/>
- `krb5-libs`: <http://web.mit.edu/kerberos/www/>
- `lcms2`: <http://www.littlecms.com/>
- `lesscss`: <http://lesscss.org/>
- `libblkid`: <ftp://ftp.kernel.org/pub/linux/utils/util-linux-ng>
- `libcom_err`: <http://e2fsprogs.sourceforge.net/>
- `libcurl`: <http://curl.haxx.se/>
- `libgcc`: <http://gcc.gnu.org/>
- `libgcrypt`
- `libgpg-error`: <ftp://ftp.gnupg.org/gcrypt/libgpg-error/>
- `libICE`: <http://www.x.org>
- `libidn`: <http://www.gnu.org/software/libidn/>
- `liblvm2cmd.so.2.02`
- `liblvm2app.so.2.2`
- `libmount lib`: <http://en.wikipedia.org/wiki/Util-linux>
- `libpcrc`: <http://www.pcre.org/>
- `libpng`: <http://www.libpng.org/pub/png/>
- `libselenium`: <http://oss.tresys.com/git/selinux.git>
- `libsepol`: <http://www.selinuxproject.org>
- `libSM`: <http://www.x.org>
- `libssh2`: <http://www.libssh2.org/>
- `libtar`: <http://www.feep.net/libtar/>
- `libtool-ltdl`: <http://www.gnu.org/software/libtool/>
- `libxcb`: <http://xcb.freedesktop.org/>
- `libXau`: <http://www.x.org>

- libuuid: <ftp://ftp.kernel.org/pub/linux/utils/>
- libXext: <http://www.x.org>
- libXt: <http://www.x.org>
- libX11: <http://www.x.org>
- lib magic 1.0.0
- lvm2-devel
- log4j: <http://logging.apache.org/log4j/1.2/>
- MariaDB: <https://mariadb.org/>
- MariaDB J-Connector: <https://mariadb.org/>
- Material Design Icon: <https://material.io/tools/icons>
- ncurses-libs: <http://invisible-island.net/ncurses/ncurses.html>
- nspr: <http://www.mozilla.org/projects/nspr/>
- nss: <http://www.mozilla.org/projects/security/pki/nss/>
- nss-util: <http://www.mozilla.org/projects/security/pki/nss/>
- nss-softokn-freebl: <http://www.mozilla.org/projects.security/pki/nss>
- OAR: <https://oar.imag.fr/>
- OpenDJ: <https://forgerock.github.io/opendj-community-edition/>
- openldap: <http://www.openldap.org/>
- openssl-libs: <http://www.openssl.org>
- paramiko: <https://pypi.python.org/pypi/paramiko/>
- progressbar: <https://pypi.python.org/pypi/progressbar/>
- pyparsing: <https://pypi.python.org/pypi/pyparsing/>
- pyreadline: <https://pypi.python.org/pypi/pyreadline/>
- RabbitMQ: <https://www.rabbitmq.com/>
- RabbitMQ client: <https://www.rabbitmq.com/clients.html>
- Reflections: <http://code.google.com/p/reflections/>
- Rhino: <https://developer.mozilla.org/en-US/docs/Mozilla/Projects/Rhino>
- Rpm libraries: <http://rpm.org/releases/>
- Saxon: <http://saxon.sourceforge.net/>
- slf4j: <http://www.slf4j.org/>
- Squid: <http://www.squid-cache.org/>
- systemd-libs: <http://www.freedesktop.org/wiki/Software/systemd>
- termcolor: <https://pypi.python.org/pypi/termcolor/>
- texttable: <https://pypi.python.org/pypi/texttable/>
- Tika: <http://tika.apache.org/>
- tinysql: <http://www.grinninglizard.com/tinysql>

- uforge_python_sdk: https://pypi.python.org/pypi/uforge_python_sdk
- URL Rewrite filter: <http://tuckey.org/urlrewrite/>
- util-linux-ng
- Vault: <https://www.vaultproject.io/>
- VMware vjjava: <http://vjjava.sourceforge.net/>
- VMware SDK: <https://developercenter.vmware.com/web/sdk/5.5.0/vcloud-java>
- Weld: <http://weld.cdi-spec.org/>
- xz-libs: <http://tukaani.org/xz>
- Xz libraries: <http://tukaani.org/xz/>
- zlib: <http://www.gzip.org/zlib>
- zypper: <https://en.opensuse.org/Portal:Libzypp>

16.1 3.8.fp13

Release Date: 2019-06-10

16.1.1 New Features

- Docker images now include entrypt. Previously generated Docker based images have to be regenerated to include it.
- **Lightweight container images. Docker images built by UForge AppCenter from Software Bundles now include a reduced set of base images.**
 - RedHat Enterprise Linux 7
 - Debian 9
 - Ubuntu 18.04
 - SUSE Enterprise Linux Server 12
 - openSUSE 42
- **Blueprint deployment has been improved. Users can now deploy blueprint on Microsoft Azure Public Cloud.**
 - Please note the following restriction: the browsers that support the Blueprint Runtime are the latest public releases of Firefox and Chrome.

16.1.2 Enhancements

- root ssh access is allowed for UForge Azure images

- When creating a blueprint, if an appliance is interactive, the “Deploy” button cannot be selected. A “warning” icon is positioned on the appliance. When the warning icon is selected, a popup opens with a button “Configure Install Profile”. Selecting the button saves the blueprint and redirects the user to the appliance profile allowing him to modify it.
- New deployment view: the new UI allows you to manage and monitor the evolution of your blueprint deployment. It provides clearly status for each instances deployed, indicating any errors as well as configuration and runtime information. Please note UForge RPM installation will update the Brooklyn plugin. If the AMP/Brooklyn endpoint is configured, the Brooklyn server must be restarted after installing UForge RPM.

16.1.3 Bug Fixes

- 1343 Repository refresh fails for repositories where file listing is disabled
- 12200 We can’t download some binaries if the root context is “/”
- 13203 Supervisor can not change supervised user
- 13411 Need to login with root/sudo user to VM before being able to use it even if no interaction required at installation
- 13538 Scan fails with API key authentication if user name includes ‘@’
- 13583 Remove call to another unexistant supervision icon
- 13654 Error when adding a new SSH key that ends with a blank line
- 13867 No clear error given when deploying an image that is deleted from Azure/AWS
- 13904 500 call failed in blueprint YAML editor

16.2 3.8.fp12

Release Date: 2019-04-29

16.2.1 New Features

- Users can publish Windows machine images to Azure Stack.
- Lightweight container images. Docker images built by UForge AppCenter for Software Bundles used to have unnecessary packages such as kernel. They now include a reduced set of packages, included in a base image, that allow users to build images on top of it. Only CentOS 7 is supported for now.

Warning: <i>Container</i> is now an OS Profile provided by UForge for all Linux distributions. Existing OS Profiles with this name must be renamed before the update.

16.2.2 Enhancements

- Ability to view the list of packages for images generated from a software component.
- Ability to generate a Docker image from a Linux appliance without the kernel.

16.2.3 Bug Fixes

- 13489 Fix incompatible version error when generating Oracle 7 image
- 13551 Make all Azure agent packages sticky
- 13350 UForge system update: `yum update docker` breaks docker command
- 10912 Empty folders in MySoftware are not present in generated machine
- 12518 Fix ssh connection for openstack
- 13291 Slash in MySoftware files not well handled
- 13302 Some appliances with an interactive profile do not have the warning icon when composing a blueprint
- 11945 Publishing to AWS stops at 39% when image size is large
- 12893 Only the bottom part of the appliance logo can be clicked
- 12472 Icon overlap in “Advanced Partitioning” and “Welcome Message” views
- 12877 Icons overlap after leaving edit network card without using ok or cancel
- 12239 Status of some ‘services’ like ufw is not detected by scan agent
- 11940 One volume remains after AWS publish cancelled

16.3 3.8.fp11

Release Date: 2019-03-18

16.3.1 New Features

- Develop and test your software applications faster. Thanks to the new local docker registry installed with UForge, you can docker run your containers directly.

Note: Previously generated Docker based images will no longer be downloadable. You must regenerate them to use this feature.

16.3.2 Enhancements

- Blueprint deployment has been improved. The user can now generate and publish all images of the blueprint appliances by clicking on the new `prepare` button.
- Blueprint Composer now raises a warning if one of the appliance component configuration requires user interaction at first boot
- UForge provides a PowerShell command to download `uforge-scan.exe` from a Windows Server

16.3.3 Bug Fixes

- 13028 Imported appliances from workspace break the blueprint appliance list
- 10553 Windows generation for some formats fails due to some registry keys

- 10854 'uforge template list' shows created time in local time while modified time in UTC
- 13077 Creating an Oracle-based target format using the CLI raises an error
- 13080 Can't connect to a CentOS 6.7 AWS instance with SSH from a "whitebox" migration
- 12854 'uforge pimages list' fails with AttributeError
- 9637 Creating a migration with unauthorized characters in name succeed through the API
- 6495 Highlighted element does not handle overflow correctly
- 12759 ImageDTO returned from API has always CentOS 7 as distributionName when image generated from software
- 12284 Windows on Azure generated from UForge-published image for Blackbox migration does not start correctly
- 12909 Mandatory fields are not highlighted in account summary of user
- 12116 Command to launch distribution population is wrong in administration part
- 11942 Sourced network interfaces are not detected by the scan agent
- 6316 Windows scan fails if second disk is full
- 12856 uforge-sync.bin overwrites PID file and makes a service behave strangely
- 12605 Files & Folders scan tab displays wrong information
- 10096 There are inconsistencies in the French translation
- 12806 Downloading an OpenShift image directly retrieves the archive without explanation
- 12588 Cannot delete multiple users in Workspace Members page
- 12734 Publishing to Azure ARM, K5 and VCenter terminate with error when it takes more than 8 hours
- 10068 Error generating Windows Server 2012r2 Standard Full English for Amazon AWS (VirtIO drivers issue)
- 9404 Error while generating an Azure image from a Windows Server 2012-R2 appliance
- 9247 Unable to generate Windows Server 2012R2

16.4 3.8.fp10

Release Date: 2019-02-04

16.4.1 New Features

- Added a new `Delta` tab in Migrations section to visually compare the difference two scan results (from one or more live instances). This allows users to visualize which files have been added, modified or deleted. Users can browse the filesystem, search and filter the results. This feature does not provide differences in the packages added, modified or deleted between the two scans at this time.
- **Introduce the new user interface Blueprint Composer which allows users to create and deploy multi-node applications. PL**
 - This release currently supports Linux and Amazon AWS only.
 - For the deployment of the blueprint, UForge must be associated with the new AMP version 5.3

- You need to install the Brooklyn Plugin Uforge after the installation of AMP server by following the procedure described in Admin Guide->Further AppCenter Configuration->Configuring Cloudsoft AMP
- The Browsers that support the Blueprint Composer are the latest public releases of Firefox and Chrome.
- Only the English version of the Blueprint Composer view is available.
- Blueprints created with a version of UForge older than 3.8.FP10 will be deleted
- Note that this new release brings an API break for Blueprint DTO model.

16.4.2 Known Issues

- Package file deletions are not detected when comparing scans

16.4.3 Enhancements

- Disable Software Bundle generation page if no base OS available
- Inject firewall configuration package when “ask during installation” is selected on Centos 7
- Disable unsupported “ask during installation” firewall option on Debian and Ubuntu

16.4.4 Bug Fixes

- 8626 After scanning Windows 2016, one service, CDPUserSvc_XXXXXX, is missing on the list of scan results
- 8871 Changing partitioning from Advanced to Basic in a template imported from a Windows scan leads to generation error
- 10519 ISO image created by UForge does not respect UTC
- 11777 Password is displayed in clear text on Summary for an application
- 12265 When publishing an appliance from blueprint UI, the row is moved to the bottom of the table
- 12280 CentOS 5 scan fails without explicit message when duplicated GPG Pubkey packages are installed
- 12293 Updates counter on appliance view is placed too far to the right
- 12337 Copyright customized through config.xml is not shown
- 12433 User password displayed in clear when deploying Windows image on AWS
- 12474 Reset the password after an attempt of sign-in with a wrong password fails
- 12480 Root password is displayed in clear inside the generated machine image
- 12483 A previously assigned mountpoint can't be reused in the UI after a file system type change
- 12619 Password is displayed in clear text on Summary for a workspace

16.5 3.8.fp9

Release Date: 2018-12-21

16.5.1 New Features

- Ability to generate a software component to a Docker image. The user can select any Linux distribution as a “Base OS”.
- Support of hotkey Escape (ESC) to cancel/close a popup window.
- UForge logs can now be pushed to an ELK instance.

16.5.2 Enhancements

- UForge Microsoft Azure images now use the Azure agent version 2.2.14-1 for Debian 8 (Jessie)
- UForge Microsoft Azure images now use the Azure agent version 2.2.18-3 for Debian 9 (Stretch)
- Improved Outscale publish connector robustness
- Blueprint deploy view has been improved. The user can now see which blueprint appliances are not ready to be deployed.
- Blueprint deploy view has been improved. Required actions for deploying the blueprint can now be triggered from this view.

16.5.3 Bug Fixes

- 6596 Unknown server error when editing incremental scan after deleting the base scan
- 9171 Publication to Fujitsu K5 timeout with slow network
- 9419 The forgotten password email contains both the username and the password
- 11191 Firewall is not configured on an image generated from an appliance with firewall activated
- 11812 Google certificate is in clear in log when uploaded
- 11933 uforge-scan.bin fails to execute when /tmp has noexec mode
- 11970 Deleting pinned package looks ok in UI but package is not deleted
- 12193 Cannot publish to Outscale us-west-1
- 12349 uforge-scan.bin cannot be downloaded from the UForge UI
- 12352 SELinux configuration is not supported for Oracle Linux

16.6 3.8.fp8

Release Date: 2018-11-12

16.6.1 New Features

- Ubuntu 18.04 supported, except for synchronization feature in migration workflow.

16.6.2 Enhancements

- Using uforge-sync binary, users can now synchronize the target environment with scans (without overlay) of openSUSE 42 systems.
- Improve usability of the add and delete actions on Projects page.
- Blueprint deploy view has been improved. Appliances used in the blueprint, with their status, are now displayed.

16.6.3 API changes

- Google Cloud Engine authentication method has been updated to support the new format used to authenticate to the platform. As a result, credentials accounts have changed. The certificate is no longer a .p12 file but a .json file.

Note: Old Google Cloud Engine credential accounts will no longer be usable. You must replace them by new ones in the correct format.

16.6.4 Bug Fixes

- 11941 Get requests to vault fail in proxy environment
- 11863 Cannot login to migrated CentOS6 image on AWS with SSH key
- 11799 Outscale images built by UForge do not boot on Outscale
- 11637 Empty directories are not synchronized to the target machine with uforge-sync.bin
- 11608 Credentials secret keys are visible in clear for Outscale, Amazon and CloudStack
- 11548 “500 call failed” is shown on non-english summary tab when a new Windows template is created
- 11532 OpenShift installation is incomplete when upgrading to 3.8.fp6 with many users
- 11499 Vault in proxy environment does not work
- 11390 User email address should not be exposed to the other user
- 11354 Partitioning Table Volumes subtitle misses a white space
- 10870 Publishing to Google Compute creates unnecessary disks and images
- 10697 Cannot publish to GCE when AppCenter is behind a Proxy Server
- 10503 Mislabeled UForge on Update tab when creating a Windows appliance
- 10444 When user quota limit of appliance is set, “Quota used” increases by 2 when importing from scan
- 10443 Cannot create Azure VM from published VHD from AWS Ubuntu template - No NIC detected
- 8989 Some French translations are not accurate or missing
- 8897 Spelling mistakes in English i18n constants
- 5224 Typo: “Unformatted” in Install Profile -> Partitioning
- 1351 “Internal server error.” displayed when publishing a Google Compute Engine with wrong credentials

16.7 3.8.fp7

Release Date: 2018-10-01

16.7.1 New Features

- UForge administrator can now register and manage software repositories and operating systems from the user interface
- Users can deploy Windows instances from published machine images to Azure

16.7.2 Enhancements

- UForge Microsoft Azure images now use the Azure agent version 2.2.21 for Ubuntu 14.04 and 16.04
- UForge Microsoft Azure images now use the Azure agent version 2.2.18 for Red Hat Enterprise Linux 6 and 7
- Using `uforge-sync` binary, users can now synchronize the target environment with scans (without overlay) of Debian 6 systems
- New icon in the UI for blueprints

16.7.3 API changes

- Update the Repository DTO model: rename field `officiallySupported` to `coreRepository`

16.7.4 Deprecated Features

- UForge command line tool option `--officiallySupported` for command `org repo create` is deprecated. Please Use `--coreRepository` instead

16.7.5 Bug Fixes

- 8341 Fixed issue allowing japanese characters to be used in the description field for an OS Profile
- 8934 Fixed issue to allow a user to correctly delete a pinned package in the OS profile of a template
- 8936 Fixed validation tooltips where backslashes are actually not supported
- 8940 Fixed unclear error message when editing a software component bootscript used in an ongoing generation
- 10708 Fixed issue when generating an ISO image from a scan of a live system with CentOS installed
- 10822 Fixed issue generating Ubuntu 10.04 with the latest `debootstrap` package
- 11096 Fixed the unpinning of a package from failing in the UI
- 11201 Fixed arrow buttons from expanding in Distribution > OS Profile view in the UI
- 11226 Cannot log in to a CentOS AWS instance with SSH key pair set by AWS
- 11292 Replaced `vssadmin.exe` with `diskshadow.exe` to properly flush Windows registry during migration process
- 11349 Fix to display tenant name in the details of a published image for OpenStack

- 11375 Fixed refresh issue for the top navigation menu to display its children when the window is resized horizontally
- 11376 Fixed refresh issue for the main navigation menu to display its items when the window is resized vertically
- 11492 Fixed misalignment in the Updates notification information in the UI
- 11504 Fixed issue with the blueprint view filter being cut during loading

16.8 3.8.fp6

Release Date: 2018-08-20

16.8.1 New Features

- New option to migrate (Lift & Shift) an instance without transferring any overlay information (only keep the operating system information, remove software application and users data)
- Using uforge-sync binary, users can now synchronize the target environment with scans (without overlay) of Ubuntu 16.04, 14.04, 12.04 systems.
- Support generation and publication of Linux machine images for OpenShift. However, it is no longer possible to publish from a Docker image to OpenShift.

16.8.2 Enhancements

- UForge Microsoft Azure images now use the Azure agent version 2.2.18 for CentOS
- Improve display of error details for failed migrations
- Enhance Clone Appliance view to display the version and revision of the current appliance to be cloned
- Improve usability of the add and delete actions on MySoftware page
- Clicking on “UForge AppCenter” (top-left corner) now redirects the user to the dashboard
- Clicking on an Appliance now redirects to the Stack tab
- Ability to publish Windows 2016 appliances on Fujitsu K5 cloud

16.8.3 API changes

- Update the Image DTO model: rename field applianceUri to parentUri.

16.8.4 Bug Fixes

- 5175 UI returns 500 call failed when portal has changed and requires clearing cache and reloading
- 7195 /etc/sysconfig/system-config-firewall file created after migration though it is not supported in RHEL 5.2
- 8050 CentOS 5 scan fails with duplicate GPG Pubkey package installed without explicit message
- 8439 UForge version displayed in the portal is incorrect
- 8724 The file name of a cloned software is incorrect

- 9475 /etc/UShareSoft/uforge-install-config-CheckRootLogin.sh not found after CentOS blackbox migration to AWS
- 10246 CLI timeout following *subscription os add* with many users
- 10653 CLI command *template info --all* always displays 0
- 10811 Deployment of Linux images to Microsoft Azure does not take ssh key into account
- 10478 Deleting a publication raises errors
- 11045 Deleting two publications raises errors
- 11059 Migration does not launch generation in a multi-node UForge environment
- 11170 Impossible to delete a PublishImage that comes from a migration
- 11171 Deleting an Image from a Scan does not work
- 11318 Some dependencies of platform tools are not injected when generating from a scan
- 11343 uforge-install-config does not execute correctly for Ubuntu 14.04

16.9 3.8.fp5

Release Date: 2018-07-09

16.9.1 New Features

- Allow users to cancel running migrations
- API users can now publish to OpenShift from a Docker image (compatible with Hammr CLI), tested on OpenShift Online and OpenShift Origin
- BTRFS filesystem support
- New customizable opt-in message in Sign Up page.
- Support Entrypoint in Docker images
- Users can deploy Windows instances from published machine images to AWS.

16.9.2 Deprecated Features

- Remove support of following formats: Abiquo, Eucalyptus, Flexiant, Nimbula

16.9.3 Enhancements

- The uforge-migrate binary displays now the progress of each phase.
- User can specify the network bandwidth allowed for data transfer when scanning a Windows system.

16.9.4 Bug Fixes

- 10251 VirtualBox image of UFIAB fails to boot with initrd root filesystem dependency failure
- 10331 Generation stucked at 55% and nothing work anymore after it
- 10335 Failed to import OVA image into vCenter
- 10430 Publish to all regions of Outscale does not work in UForge (unsupported regions, ami id out of date, wrong user for connection)
- 10555 On AppCenter with many users, appliance GET for one user has performance issues
- 7617 In the publish views, some select lists are randomly sorted
- 8638 License cannot be changed in the clone Software
- 9847 [Documentation] File size of /boot/grub2/i386-pc/core.img is changed during blackbox/whitebox migration
- 10217 Floating point exception occurs on uforge-scan.bin
- 10325 OAR jobs logs show WELD “Exception in thread”
- 1415 Portal - Language selection menu truncated in firefox
- 9672 SLES 11 scan sticks during creating report
- 9676 When adding a certificate for the creation of a google compute engine, the spinner never stops
- 9782 Windows black box migration failed at publication with classCastException
- 9836 [Documentation] configuration to connect to AMP is incorrect for multinodes environnement
- 9956 Overlapping text when generate an K5 image from a Linux scan
- 9961 Migration is stuck in progress
- 10103 Import a bundle first as first action will block subsequent template imports
- 10105 Scrollbar is not well displayed on deployments and blueprints views
- 10588 Windows on Azure generated from UForge-published image for Whitebox migration does not start correctly.
- 10657 You cannot save a Windows template imported from a scan with an error.
- 10795 Generation of a Debian 8 Server OS profile image fails with server install profile

16.10 3.8.fp4

Release Date: 2018-05-01

16.10.1 New Features

- The automated migration process is now available for Windows
- Cloud credentials are now stored in a new secret manager (Vault) in order to improve security
- New option to exclude some files and directories when migrating a live machine
- Introduce the new Blueprint module which allows users to create and deploy multi-node applications. This release currently supports Linux and Amazon AWS only.

16.10.2 Enhancements

- Improve the migration details page in the user interface

16.10.3 Bug Fixes

- 3695 No space left on virtual disk is not caught as an error and generation returns “internal server error”
- 9013 Docker image format generated by UForge is incompatible with latest Docker executables
- 9044 Publication to AWS or Outscale: message when cannot connect to proxy is misleading
- 9406 Azure publishing feature transfers VM images via HTTP instead of HTTPS
- 9555 Subscription quota update command without a limit sets the limit to 0
- 9679 Google Compute Engine Regions are outdated
- 9733 Use public IP address instead of private one in the security group rule when publishing to AWS
- 9918 UI freeze in Google Chrome when entering the ‘Stack’ tab
- 9940 Cannot publish to azure with existing account through REST API
- 9941 Image generation for K5 failed while checking the dependencies
- 10099 Boot scripts cannot be found in the cloned MySoftware

16.11 3.8.fp3

Release Date: 2018-04-16

16.11.1 New Features

- Windows system with an extended partition is now supported for templating, image generation, and scanning.
- The automated migration process is now available for all supported platforms

16.11.2 Enhancements

- Visualise all the software (libraries, drivers or packages) that is automatically injected by UForge during the generation of a machine image for a specific target cloud environment.
- “Migrations” tab now contains both automated migration and scan features, available under “Lift & Shift” and “Re-platform” sub menus
- Allow users to delete multiple migrations
- 9057 Remove End-of-life Microsoft Azure Classic

16.11.3 Restrictions

- Currently the image generation of Linux system for K5 migration fails. This is due to a known issue during dependency checking. 9941 Image generation for K5 failed while checking the dependencies

16.11.4 Bug Fixes

- 9937 Segmentation fault when generating a machine image
- 9762 Default chunk size for publishing to K5 is too low and leads to K5 error
- 9411 When scanning a machine, UForge portal UI displays a big OS Logo. Fix the logo size in the header and allow to debug CSS from remote computers
- 9771 Docker image format should not be proposed for Windows migrations
- 9781 When selecting Azure platform, AWS or K5 for generation, an unknown server error is displayed
- 9424 When a migration is deleted, the error message of the migration tool is not relevant
- 9641 Delete account with a certificate, after a publish, fails
- 9639 Wrong URI for Scan installProfile
- 9657 Unable to publish Docker images to Azure Containers
- 9615 Generation dashboard is in error after a Migration Generation Stage has completed
- 8917 After Blackbox Migration, RHEL5.3 is updated to RHEL5.5.
- 9598 Provisioning on Azure fails due to No DVD device
- 9582 SLES generation for Azure format does not work (no platform tools injected)
- 9258 Unable to do a Debian 8 blackbox migration to Microsoft Azure
- 9074 Improve error handling in getCredAccountResources service
- 9188 Partitions in LVM logical groups/volumes appear in wrong order if a group's name is changed
- 9532 When importing an appliance, Firewall is set to Ask during installation
- 9502 NIC configuration disappears on install profile in an imported appliance from a CentOS5.7 scan
- 8682 If the root context is not modified in the deployment wizard, the page redirected to at the end is / which is forbidden
- 9081 "should contains" to be replaced with "should contain" in bootscript name in UI
- 6200 Fix scan installProfile URI
- 8973 Ubuntu 16.04 image generation requires debootstrap to be of version at least 1.0.85 and does not support "proposed" packages
- 9094 If a user's home directory is in a multiple levels folder hierarchy, the user creation fails
- 8951 Remove GoldenPath field from uforge.xsd
- 9176 Please remove obsolete Squid directives from /etc/squid/squid.conf
- 7937 Debian dependency checker returns too much detailed message when failed.
- 8371 Add SYSPREP setting in windows yml template
- 9317 uforge_update.sh fails with SQL error in db_modifs_180130-01.sql

16.12 3.8.fp2

Release Date: 2018-03-05

16.12.1 New Features

- The migration process (scan, generate and publish) can now be automated by using a simple and intuitive workflow in the user interface. The user simply creates a new migration through the interface, launches it from the server to migrate and follows the overall status.

Note:

- This new feature is available under the “Migrations” tab. This tab already existed. This previous tab has been renamed “Scans”.
 - This feature is currently available for the following platforms : Fujitsu K5, Microsoft Azure, VMware vCenter, Openstack and SUSE Cloud.
-

16.12.2 Enhancements

- 9004 As a user, I would like to have my appliance revision incremented when I modify an attached Software-Bundle
- 5994 Provide a way to offuscate and reveal passwords in the UI
- 9054 RHEL 7 and 6 should be supported in K5 format

16.12.3 Bug Fixes

- 9091 iptables rules have been changed - whitebox migration
- 8721 Scanning CentOS 6 generated with UForge results in an error
- 8648 The Deployment Wizard should block non-numeric inputs on the Proxy Port
- 8646 Debian 9 missing in the deployment wizard
- 8745 Disk usage increases when uploading twice the same file for limited quota user
- 9178 Overlapping text in the summary view of an appliance template
- 8664 systemd-tmpfiles-clean.timer clears /tmp thus removes the symlinks and breaks AppCenter
- 8613 [RHEL7/CentOS7] /boot/grub2/device.map is cleared during Blackbox/Whitebox Migration
- 8222 Software bundles are not extracted in the correct directory
- 8847 ComboList in portal appears empty after selection with firefox
- 8921 Add scan import to golden cancel webservice
- 9024 Outscale image generation limited to 10 GB disk size
- 8701 UForge deployment fails behind proxy because of unaccessible ntp server
- 8422 openssh bits in uforge template do not need to be sticky anymore
- 9198 Service mysql restart display FAILED message
- 8865 Despite deleting the golden images, the files of the golden image are not deleted in the file system
- 9053 user login and password are sent to the user in the same e-mail message
- 8842 Docker publishing cannot be canceled
- 9010 After a file upload error, the value of consumed diskusage increases when the quota is changed to unlimited

- 8412 Ubuntu 14.04 generation fails with stack overflow error
- 7493 A letter ‘&’ in comment field of /etc/passwd file changed to ‘&’ after migration
- 8873 CPU usage of Dozer thread sticks to over 99% and never ends
- 8995 script machine_infos.sh fails
- 9045 Shell injection, the user can execute command as tomcat user when calling publish api
- 9026 template imported from scan fails to generate
- 8899 openssh package version has been changed after a white box migration
- 9185 If user role is only Migrator, an error occurs in cloud account selection of publish image
- 9165 No need to call reset_eventcontroller.sh in the crontab anymore
- 9125 After Black box migration, Firewall setting changes to enable in Cent OS 6.
- 9257 Error occurs when migrating to Hyper-V of Windows Server 2012 R2
- 9108 [Server-side]Add the Timezone param to uforge-install-config.conf
- 9089 [Server-side]don’t write a firewall param if the template is Windows
- 8320 “user enable” uforge-cli command always resets password
- 8987 “Request timeout” is be shown during a generation and requires refreshing the screen
- 8444 Missing /opt/Tomcat/.bashrc file
- 9144 a logical group disappears after removing an LVM disk
- 9175 removed partitions from a volume group still remain in the group
- 9210 Scanning a CentOS 7 server with a CD in the drive causes an additional hard disk to be added

16.13 3.8.fp1

Release Date: 2018-02-01

16.13.1 New Features

- Users can deploy Linux instances from published machine images to Apache CloudStack.
- Fujitsu Cloud Service K5 jp-east-2 region is supported.

16.13.2 Enhancements

- The uforge-sync binary now requires the API endpoint, to improve usability.
- Enhance UI headers for Apps and Migration tab.
- Add French internationalization.

16.13.3 Bug Fixes

- 835 Refresh the generation page causes a 500 call failed
- 1060 Files permissions changes after blackbox migration
- 1064 Timezone is always reset to Europe/London after a blackbox migration.
- 1416 Portal - MySoftware - Files - package file path not restored
- 6769 The presence of a malformed filename in the source filesystem causes the scan to hang badly (segmentation fault)
- 7019 Filename vCneter.log spelled wrong. Should be vCenter.log.
- 7021 Backslash not properly escaped in credentials causes VMware vCenter publish to fail
- 7087 UTC and ARC settings in /etc/sysconfig/clock has been changed during blackbox migration
- 7112 In a blackbox migration /etc/USharesoft/ files are not deleted
- 7134 'org repo delete' fails frequently and the error message is confusing
- 7187 Parameters in /etc/fstab has been changed during blackbox generation
- 7193 /etc/gshadow has been changed during the blackbox migration
- 7196 /etc/shadow lock and 'no password' options not taken into account during migration
- 7214 Popup have an unexpected scroll bar
- 7275 Error message does not include any information when publish to AWS failed.
- 7284 Directories/files changed during blackbox migration
- 7416 Packages of custom repo still visible even after repo detach
- 7423 The UI view for searching and adding an OS package to a template shows too many versions
- 7659 Blackbox migration of CentOS7.2 on fresh forge failed "Detaching loop"
- 7680 /etc/sysconfig/clock file is added in Blackbox and Whitebox migration
- 7684 /etc/sysconfig/kernel is modified after Blackbox and Whitebox migration
- 7687 Hammr deploy OpenStack retrieval timed out
- 7712 Viewing bootscript of a cloned template raises a 404 error
- 7730 Uploading several files to a Software bundle randomly leads to 500 error
- 7738 Bad concatenation in kernel parameters after two blackbox of a debian appliance
- 7747 The field for disk size at generation for AWS should be in GiB
- 7748 Unsupported AWS region are displayed in the publish view (cn-north-1, us-gov-west-1)
- 7758 UForge cli takes minutes to manipulate repositories
- 7819 uforge-sync does not resolve fully qualified names for AppCenter endpoint
- 7869 Enabled Firewall becomes disabled after Scan and Import
- 7935 uforge-scan output is not proper english
- 7940 Version of uforge-scan is not consistent with version of UForge platform it has been downloaded from
- 8054 "UForge critical error" e-mail is sent after a successful scan import
- 8055 "C:fakepath" is displayed when selecting an appliance archive to import

- 8062 Publishing a compressed image failed on OpenStack
- 8063 UForge update logs show WELD “Exception in thread”
- 8064 Software bundles are not extracted in the correct directory
- 8076 “500 call failed error” when uploading a boot script to a project catalog
- 8097 VMware vCenter publish fails in multiple vlan/vnic environment
- 8102 Display explicit error message when template has no partition
- 8146 Typo in Artifact account in the creation page
- 8180 Folders where VMware vCenter templates will be published are changed randomly
- 8192 OpenStack generation from scan fails with message Installed packages more than expected (240 > 237)
- 8214 When moving from Name to Version with the tab key, the Version box becomes red (error)
- 8306 Export, Import and Scan features does not work when UForge user login contains ‘@’
- 8322 Publishing a compressed image failed on VMware vCenter
- 8326 The order of NICs is changed by exporting/importing a template appliance
- 8476 uforge-cli command template info throws AttributeError: NoneType for Windows Appliances
- 8649 Only one architecture of an OS package is kept when there were multiple in the imported template
- 8889 uforge-cli command template info throws AttributeError for Windows Appliances
- 8898 Publish on OVH Openstack does not work

16.13.4 Known issues

- In some situation, deployment in CloudStack could fail if the CloudStack image is duplicated in different zones.

16.13.5 Compatibility issues

- The import / export of appliance templates from UForge 3.7 to UForge 3.8 may not work if the template contains software bundles. Please refer to the section Importing and Exporting Templates (Updating a 3.7 Appliance) to make your template compatible.

16.14 3.8

Release Date: 2018-02-01

16.14.1 New Features

- SLES 11 and 12 operating system supported for all features (templating and migration)
- OpenSUSE 42.x operating system supported for all features (templating and migration)
- Debian 9 (Stretch) operating system supported for all features (templating and migration)

For other features, please refer to 3.7.fp8 release notes

16.14.2 Migrating to 3.8

For specificities relating to migrating a 3.7 or 3.7.fpx to 3.8 please refer to the section Migrating UForge from 3.7 to 3.8 in the Admin guide.

16.14.3 Bug Fixes

- 8656 Estimated size of Windows templates is 0 B
- 8653 Generation fails for an imported Windows template built on a “Scan To Golden” profile
- 8578 Generation does not finish if there are volume groups though it was cancelled or got an error
- 8577 Image generation of a CentOS 6.7 scan from ISO fails in grub installation
- 8507 Update error message about RHEL not supported for K5 in UForge
- 8505 Publishing a Docker image fails in slow network environment
- 8501 Export, import and scan features do not work when UForge user login contains @
- 8486 Image generated from a CentOS 7.1 scan from ISO fails to boot showing the grub shell
- 8499 UNIX group ID is not taken into account when import a template
- 8437 Name and downloadId missing in the download URL for appliances generated from a template
- 8417 Windows image generation from a legacy golden fails without displaying the details if required disk is too small
- 8309 Windows automatically shuts down after being instantiated on AWS
- 8270 Scan of CentOS 7 fails with message `Unable to rebuild package dialog 1.2 x86_64 on 3.8`
- 8094 Whitebox image generation failure with `non encrypted password error`
- 8078 Add arch selection, in order to allow install of package with multiple architectures
- 7831 Scan on CentOS 7.4 with LVM fails silently and causes generation error

16.15 3.7.fpx8

Release Date: 2017-10-16

16.15.1 New Features

- Using `uforge-sync` binary, users can now synchronize the target environment with scans (without overlay) of CentOS 6, 7, Red Hat Enterprise Linux 6, 7 and Debian 7, 8 systems.
- Microsoft Azure connector has been updated. Previously with UForge the machine image was publish as a “vhd” blob file in the Azure cloud Account. Now an image will be accessible in the cloud console from this blob file. In order to support this additional information must be entered in `Credentials` for Microsoft Azure ARM connector.
- Support generation and registration of machine images for Oracle Cloud with the metered service subscriptions.
- Users can deploy Linux instances from published machine images to Microsoft Azure ARM.

Note: If you have an existing Microsoft Azure ARM account already setup in UForge, then you must update the credential information.

16.15.2 Enhancements

- Improved deploy button tooltip in Dashboard view
- Replace spinner by ProcessStepWidget for OpenStack
- UForge users can inject specific VirtIO drivers for Windows appliances
- Amazon AWS connector can now publish Windows images with multiple disks

16.15.3 Bug Fixes

- 1311 Error “WELD-ENV-002002: Weld SE container was already shut down” can be displayed in the portal when generation failed
- 6196 Image generation from a scan fails when the repository is updated by the spider simultaneously
- 6359 Scan comparison shows two packages with different versions instead of package’s target scan
- 6669 Installing UForge AppCenter in a root context other than /uforge breaks some features
- 6848 Disk order and partition number are not kept after migration
- 6862 All fields in deployment tables should be displayed entirely
- 6957 When scanning a RHEL machine, UForge portal UI displays a big RHEL Logo in IE
- 7004 /etc/sudoers is reinitialized after migration
- 7016 CentOS images from blackbox migration fail to start on Microsoft Azure: no WALinuxAgent installed
- 7076 Generation error when extracting overlay if size is bigger than / (root) partition
- 7109 Tooltip of source used on a deployments is wrong if come from a scan
- 7114 Protect Deploy Activity from incomplete publish image
- 7149 When scanning Windows 2012 R2 and blackboxing it to VirtualBox, Windows requires to change admin password at first boot
- 7150 Error when specifying a directory of more than depth 1 in mount points in install profile
- 7164 Blackbox migration of debian 7, 8 and ubuntu 14 does not boot on major clouds due to DHCP ipv6 activation
- 7184 NetworkManager package is present in “server” profile and the generation does not work with Azure
- 7194 CentOS 5.11 scan fails at phase 4/7 by segmentation fault
- 7253 Scan fails with SQL Error: 1205, SQLState: HY000 when running two scans concurrently
- 7408 CentOS whiteBox migration to Microsoft Azure: wrong version of WALinuxAgent selected
- 7510 CentOS 7.4 and Oracle Linux 7.4 fail to boot
- 7673 Generation of a migrated debian 8 fails randomly
- 7686 Whitebox migration : multinic method of second interface is disabled instead of static or manual
- 7697 File System type not set properly for logical partitions

- 7711 Outscale cloud: cannot see and publish in new regions

16.16 3.7.fp7

Release Date: 2017-09-04

16.16.1 New Features

- Fujitsu Cloud Service K5 US, Finland and Spain regions are supported.
- Amazon AWS Ohio, Mumbai, London and Canada regions are supported.
- Introduce a REST API for users to list files to synchronize to the target environment after a CentOS scan without overlay.
- Users can deploy instances from published machine images to OpenStack.
- Images can be created in PXE format for CentOS.

16.16.2 Enhancements

- Improve UI text and tooltip message for K5 Project ID

16.16.3 Bug Fixes

- 944 Scanning failed at Phase 6 (heap memory in eventcontroller)
- 960 Failed to delete together two or more artifact accounts with error
- 985 Error occurs in the UForge CLI images list if user has generations only from scan
- 1323 Using AWS zone ap-south-1 (Mumbai) with the API raises a publication error
- 1370 CLI command “uforge image list” results in SimpleFacetValueError when the keyboard is “jp”
- 1379 Artifact accounts display bug under IE
- 1419 Package kernel-PAE not recognized as a kernel when doing a black box migration
- 5740 Incomplete Japanese translation of the “Pull a remote file” dialog box
- 6103 Modifying a used software component restriction rule raises an internal server error
- 6108 Support /dev/cciss/cXdXpX disks that exist on HP server
- 6133 AWS connector logs are set to DEBUG and should be set to NORMAL
- 6157 Debian Scan: All the files are uploaded to UForge server when scanning with overlay
- 6162 Tooltip when generating from scan (blackbox) mentions install profile changes
- 6165 Comparing two scans, there are no strike-through on the delete files
- 6235 Uploading boot scripts or my software using IE causes an error if the local directory path is included
- 6261 UForge UI for Windows scan using the command line misses the -p parameter
- 6268 Viewing My Software from Imported Scan raises an unknown server error
- 6385 When a scan source CentOS 7 machine has “/boot/grub/grub.conf”, the generated image does not boot

- 6386 Messaging bus consumer breaks down if cloud-init is installed on UForge server
- 6444 Error in NIC API examples, request URI is wrong
- 6501 Deleted package files are recovered after migration
- 6507 Debian migration cannot detect missing info changes
- 6573 “Uploading Archive 0 %” and “Transfer in progress -1 %” are displayed alternately by template import
- 6611 Error message when trying to delete a UForge OS profile milestone not self-explanatory
- 6614 UForge web service response (401 unauthorized) is not RFC compliant
- 6683 Heap memory error when scanning VM with huge files number
- 6753 AWS publish failed in some circumstances
- 6760 uforge-install-config TUI does not appear when using K5 Console
- 6768 Unable to build a package using rpmgen with a file located in /
- 6853 Unable to rebuild RPM, using rpmgen, with hook scripts containing a commented spec file section tag like `##install`
- 6903 Outscale Publish final status never reached
- 6906 yum is injected during blackbox migration
- 6917 When cloning an appliance, the parentApplianceUri of the clone is null

16.17 3.7.fp6

Release Date: 2017-07-24

16.17.1 New Features

- Introduce a new deployment feature which allows users to deploy published machine images directly to Amazon, without having to connect to their Amazon account.
- The scan of Windows is optimized by extracting only “used space” from target disks on the source system. “Free space” on target disks will not be copied by the scan.
- Fujitsu Cloud Service K5 Germany region is supported.

16.17.2 Enhancements

- Improved the information displayed in the banner when administering OS Profiles. Now the date the OS profile was created is displayed (for Windows only), as well as the size and the associated distribution (for both Windows and Linux)
- Support Linux multiple disks publication to AWS

16.17.3 Bug Fixes

- 851 After K5 Black box migration, some packages were updated
- 884 uforge-scan.bin ignores option `-e " / "`

- 953 Cannot add a license in a MySoftware
- 972 Internal error happens when clicking directory name { | } in Files & Folders in Migration
- 982 debootstrap.log should be preserved
- 1001 Publish Outscale changes proxy instance ID configuration
- 1003 Image generation fails for a CentOS 6.1 imported from scan and upgraded to the 6.5 milestone
- 1061 /etc/ssh/sshd_config file is changed after the migration.
- 1304 All POST and PUT API examples in the documentation miss Content-Type parameter
- 1305 API doc cloudAccount_create needs correcting
- 1306 API doc for creating API key pair needs updating
- 1317 Dashboard quotas are used more than 100%
- 1353 Correct https_proxy that breaks perl cloud (openstack) connectors
- 1393 Remove uforge-anytermd and remove its pid file after uninstalling uforge-install-config
- 1402 Migrator Role does not allow to generate image from a Scan
- 5447 Scanning a UForge server raises an internal server error and a Mapping error
- 6084 Images list Cli command returns wrong OS name
- 6127 The latest perl-Compress-Raw-Zlib package is not used
- 6147 Generation fails when selinux packages are manually specified in a MySoftware
- 6148 /etc/ssh/sshd_config modifications are ignored after a blackbox migration
- 7431 Cannot generate debian 8 when /tmp partition is small and has large extra files
- 7806 Scan comparison raises a 500 call failed error when clicking on a package modification of the comparator
- 7842 Cloud account password is saved as plain text in text file
- 7915 cleanup_tickets.sh and cleanup_scans.sh do not delete Generated Images from Scans
- 7986 The UI in Stack u003e OS profile displays the latest version of the OS packages instead of the one used in the appliance template

16.18 3.7.fp5

Release Date: 2017-06-12

16.18.1 New Features

- Add a mechanism to restrict the usage of a project (for Administrator) or software bundle (in MySoftware for users) based on a Distribution name, family, architecture or for an output format
- Administrators can now create a golden image from the UI. Once a Windows scan is complete, the Administrator can import the scan as a golden image. This golden image will be available to users to create new Windows appliances.
- Publication to VMware vCenter improved. UForge now publishes templates to VMware vCenter, rather than instances. The datacenter information (ESXi hypervisor, datastore, network name, etc) is automatically retrieved by UForge and prefilled for publication to VMware VCenter.

16.18.2 Bug Fixes

- 7560 Oracle Linux is treated as RHEL at scan
- 7622 OpenSUSE generation Failed with default OS Profile due to package conflict.
- 7423 CentOS 7.3 VBox image never ends up booting if '/' partition is a logical volume
- 7429 Error message is always logged in oar error log "unary operator expected" when generating Linux image.
- 7361 Windows generation error when disk too small does not raise an understandable error message
- 7620 db_modifs and postupdate modifs applied to several versions of UForge are not handled properly by update_mechanism
- 7758 On the generation UI pages, there is a CSS issue between headers and content
- 7853 License of WS2008R2 is displayed on WS2012R2
- 7771 Hover message on items in scan list is not internationalized
- 7871 Updating the certificate of a google cred account generates a null pointer exception and "this should never happen, please updateTemplateInfo" in the UI
- 7682 Published image tag summary displayed wrong tooltip
- 7635 The type of the password input field of artifact accounts is inconsistent in the UI
- 7584 Applications and Services for Windows are not parsed correctly
- 7767 Missing timezone data on branch master
- 7897 500 error occurs in image generation using a template without a partition table
- 7669 The EventController service does not consume ScanAction event correctly
- 6285 AWS publication is not working behind an external http proxy
- 7630 Outscale publish connector is not working anymore
- 6789 When booting a migrated instance, haldaemon starts although autostart setting is off in the source machine
- 7298 Import/Export Software bundle fails with "Permission denied error"

16.19 3.7.fp4

Release Date: 2017-05-02

16.19.1 New Features

- New user dashboard providing usage statistics and quota information
- New option to scan a live machine without transferring any overlay information (allowing a light-weight audit of the instance)
- Support to create appliance templates based on Windows 2016 operating system
- Ability to scan and migrate Windows 2016 instances
- Application and services information now captured and displayed when scanning a Windows-based instance

16.19.2 Enhancements

- Ability to export an appliance template in either YAML or JSON format (default now YAML)
- Enhanced the information displayed after registering machine images to a cloud environment. machine ID and cloud location (region, zone etc depending upon the cloud target) now displayed in the UI

16.19.3 Bug Fixes

- 7553 A workload based on Scientific Linux cannot be scanned
- 7546 Scanning failed at Phase 6 (heap memory in eventcontroller to the even bus - message too large)
- 7534 Wrong values in /etc/fstab if the appliance has both partition '/' and partition '/boot'
- 7521 `hammr template import` fails for certain types of advanced partitioning tables
- 7500 K5 publication fails with message "Error creating publish command for K5"
- 7436 500 call fail when displaying the detailed information of a scan when `i18n` is Japanese
- 7403 Windows scan command displayed in the UI is wrong
- 7369 Error badly handled during appliance import if message contains ""
- 7360 Oracle Linux 7 and Scientific Linux 7 machine images do not boot if the appliance templates has logical volumes
- 7340 Scanning a server with a file larger than 40 GB fails
- 7314 Cannot generate a machine image for Fujitsu K5 format from a scan
- 7229 Registering a machine image to AWS fails with `Java PublishCommunicator` error
- 7157 The scan binary ignores option `-e "/"`
- 7153 Scan cannot treat files whose name includes `>`
- 7147 Docker publish does not work anymore
- 7092 When launching Service Management Tool from `run -> services.msc`, an error occurs
- 7071 The check box `Ignore dependency checking warnings` is displayed in the UI when a Windows image is created
- 7063 Inconsistent update of template revision
- 6960 Simultaneous scans of two CentOS 7.3 machines fails
- 6932 When cloning an imported appliance and exporting, the wrong page is displayed
- 6748 Unable to download a generated machine image via the UI twice
- 5977 When resetting password, the information message to indicate that an email has been sent is badly positioned
- 5907 When inviting a collaborator to a workspace, email textbox is case insensitive
- 5074 Bad vertical aligned text in expandable button

16.20 3.7.fp3

Release Date: 2017-03-21

16.20.1 New Features

- Users can now import a Windows based scan, creating an appliance template. This allows users to update the appliance template prior to migration.
- Users can specify to run `sysprep` as part of a machine image generation for Windows-based appliances that have been imported from a Scan. This allows users to provide a new administrator password as part of the install profile.
- Ability to trigger Repository updates manually via an API call.

16.20.2 Enhancements

- The UI updated to display the language, type and edition of Windows OS profile
- The UI can be customized to allow hyperlinks in the footer or header to either open in a new tab (default) or in the same tab (replacing the UI).
- API Keys now have optional name and description meta-data to help the user identify what API keys are used for.

16.20.3 Bug Fixes

- 7146 Scan cannot treat files whose name includes >
- 6995 The scan status is not updated to `error` when the error occurs during uploading
- 6993 A file or directory name whose include a line feed (LF) is not present in the scan result
- 7069 Upload a logo which is not `png` or `jpg` raise an error but erase the existing logo
- 7065 Incorrect warning message when appliance has multiple disks during generation of some formats
- 7061 Issue when adding PDF files as custom license to project (should not be allowed)
- 7035 `rpmgen -e` (exclude dir list) option is not working correctly
- 7074 MySoftware files are not copied on the filesystem when generating CentOS7 ISO images
- 7024 Windows scan of a machine with 2 disks, when user excludes 1 disk, UForge still creates 2 disks in the scan meta data
- 7067 `uforge org category delete` fails with two arguments
- 7029 Cannot create unformatted logical volumes
- 6939 My profile picture is not displayed on Activity Stream Workspace
- 7048 Search for packages does not take into account hour of the day
- 6873 Amazon publication - S3 bucket is not necessary anymore
- 7009 UForge root password can not be changed wrong message
- 7002 Spider do not cleanup all temporary dirs in `/tmp`
- 6948 Projects non-native files are ignored if my software has the same name.
- 7003 Windows generation is failing during OS check
- 6998 When exporting a linux appliance without OS Profile an internal server error is raised
- 6986 After delete a custom license in MySoftware or Project , the icon `done.svg` is still visible

- 6971 After K5 Black box migration, Firewall setting changes to enable in Cent OS 6.
- 6970 CentOS 6 scan and generation leads to an error
- 6884 Generation of AWS image for Windows Server 2012R2 fails with illegal seek exception
- 6834 After the migration from 3.5.1 to 3.6, created API keys no longer displayed in the UI
- 6964 Canceling the K5 publication finishes with ERROR message.
- 6961 Incoherence in template and mysoftware revisions when sharing to workspace
- 6963 Internal generation tools must generate the correct guestOS inside vmx when windows+vmware
- 6747 An image can be downloaded more than once by using the URL with same Download ID
- 6855 Cannot retrieve directory from remote site with http basic authentication in software library.
- 6794 Documentation mentions copyright in customisation but copyright is not displayed
- 6870 A generated CentOS 6.8 image kernel panics if it has a logical volume in the partitioning table
- 6815 Cannot pull files from FTP in MySoftware.
- 6875 When uploading a file for the second time the confirm popup has two handlers and so the action is carried out twice
- 6872 Success message for `org os add` is not correct
- 6800 Cannot download non-cached software artifact correctly if the remote file size has been changed.
- 6819 While scanning Windows OS, Scan progress is continued to copy on the clipboard.

16.21 3.7.fp2

Release Date: 2017-02-13

16.21.1 New Features

- Support registration of machine images for Azure ARM and Azure Enterprise Accounts
- Support for Ubuntu 16.04
- Ability to register docker images built in UForge to DockerHub. This includes managing credential information to authenticate against DockerHub.
- In `Projects` or `My Software` can now provide restrictions to determine if they are compatible with a particular OS family, type or version.

16.21.2 Enhancements

- Renamed `VM Builder` Tab in the UI to `Apps`.
- Better internal logging information when publishing/registering machine images to a target cloud environment.
- Better validation in the web service for information used in publishing/registering machine images.
- Better UX experience when managing and choosing `pinned` (or `sticky`) packages.
- UI now displaying the size of the generated machine images.

- Can now delete an invitation of a user to a Workspace if a user has invited someone to join a collaboration workspace, and the person is not responding, there is no way to cancel the invitation.
- Added an `Id` column for all UForge CLI commands that lists information (for better referencing in other commands).
- Added the ability to reset a user's password via the UForge CLI (`--resetPassword` option).

16.21.3 Compatibility Issues

Migrating to UForge 3.7-2 will have the following compatibility issues:

- any Windows golden image that use a non-standard Edition (for example `Windows K5` instead of the official `Standard`, `Enterprise`, `Webserver` or `Database`) will be changed to `Standard` edition. A warning will be added to the log files. If you would like to change the Edition of the golden image, you should re-register the golden image with `org golden create` command.

Warning: Fujitsu is not legally responsible for any damage or loss caused by the possible inconsistency between the assumed and the actual Editions.

The following API interface and calls have been modified:

- The object `DistribProfile` is now an abstract object and is implemented by either `linuxProfile-object` or `windowsProfile-object` (which are new object types).
- The deprecated object `DistribProfileTemplate` has now been deleted. The object `distribProfile-object` is now used. The attribute `standardProfileUri` is now deprecated and been set to `null`.

Due to the above object changes, the following API calls have been modified:

- `orgOSWindows-add`
- `orgOSWindows-delete`
- `osTemplate-getAll`
- `osTemplate-get`
- `orgOSWindows-getAll`

The following API calls have been added to enhance scanned Windows-based workloads:

- `workspaceTemplateOSApplications-get`
- `workspaceTemplateOSServices-get`
- `workspaceTemplateOSPartitionTable-get`

16.21.4 Bug Fixes

- 6853 While scanning Windows OS, Scan progress is continued to copy on the clipboard.
- 6821 Blob name must finish with `.vhd` and add some information in the publish popup.
- 6820 Issues in properties `i18n` file.
- 6809 OpenStack account turned into another type of cloud account after the migration from 3.5.1 to 3.6.
- 6706 Fix backward compatibility for golden edition with custom names.
- 5607 Even if the scan ends the UI continues to ask for information of the scan.

- 6737 Sub menu scrollable inside the Dashboard.
- 6734 Cannot delete template with software component from workspace.
- 6732 Unexpected scroll bar in My Software view.
- 6716 Cannot download rpms from yum repos whilst scanning a centos system.
- 6713 Error message containing typo for windows disk size.
- 6711 Golden location is retrieved from Pkgs table, it should be retrieved from WindowsProfile table.
- 6672 Scan does not read KEYBOARD in metadata.
- 6646 File conflicts against packages built with when installing centos distribution packages.
- 6639 Primary disk size is changed to the other disk size on UI when having multiple disks.
- 6627 Cannot export a template if the software component has rpm file in Repository Packages tab.
- 6614 Creating folder failed but displayed on UI.
- 6599 i18n properties breaking master build.
- 6596 Imported appliances from archive are not counted statistics in Dashboard.
- 6529 Image generation fails when a template includes rpm file with no cached.
- 6497 Can't display Projects as guest user.
- 6495 The `org golden xxx` command fails if edition name in db is not allowed.
- 6492 Badly formed error label for Credentials Microsoft Azure.
- 6480 Spelling mistake retrieving remote path and error message shown.
- 6478 Sharing a template in collaboration, including software that does not use the cache of the fetch, raises an Internal Server Error.
- 6460 Imported appliances are counted as created on statistics in Dashboard.

16.22 3.7.fp1

Release Date: 2017-01-09

16.22.1 New Features

- Multi-NIC support for Linux based appliance templates.
- Driver injection improvements (internal mechanism) for Windows-based appliance templates.

16.22.2 Bug Fixes

- 6326 Impossible to publish an OpenStack VDI image
- 6323 Cloud account name appears twice in the public informations in UI for all Cloud formats
- 6234 Sticky package of imported template is not shown in the UI
- 6141 User gets a 500 call failed if a custom target platform has been added but not enabled specifically for the user
- 6042 OS packages are not sortable in the Repository column

- 6237 Spelling mistakes in the API docs
- 6222 Format enabling/disabling not working when updating the UI config
- 6453 Impossible to generate image when install profile contains users
- 6199 Migration fails because the user ID taken from a scan and user ID that the package makes overlap.
- 6409 OE-lite can't fetch QT source file
- 6206 Filter inactive pkgs on `DistributionPackages.getAll()` method
- 6200 Scanning a disabled OS is possible
- 6190 Scanning an azure vm with advance partitioning : install profile partitioning not correct
- 6180 Errors outputted into `/oar/oar_scan_job*.stderr` when scanning CentOS 6
- 6154 Launching windows scan binary from command line with API key does not launch the scan
- 6134 Pkg overlay archive are built differently if a black box migration is done first or if it's a scan import to appliance
- 6309 Several concurrent generations could fail if there are uncached software bundles files in it
- 6211 Creating a two bootscripts with same name does not show an error message
- 6194 Japanese Characters are OK to use but encoded incorrectly for `Tag` and `Maintainer` fields of a software component
- 6193 Same rpm file can be uploaded without overwritten to a software component
- 6178 Errors outputted into `/oar/job_finalize.log` when generating CentOS image
- 6169 Total Disk Usage doesn't count the size of files uploaded to software components
- 6027 Exported template has lost some information on `MySoftware`
- 6346 WARP should skip to inject uforge agent in the specific condition
- 6327 Scripts are not imported when sharing a template in a `Workspace`
- 6057 Yum update error `uforge_update.sh: line 660: [: too many arguments`
- 6055 The volume shadow copy is not deleted after scan of Windows.
- 6007 Code in `distrotools/lib/str.[c|h]` in function `repl_str()` cannot compile for windows using `mingw c++`
- 6440 Can't display `Projects` as guest user
- 6453 Impossible to generate image when install profile contains users

16.23 3.7

Release Date: 2016-12-27

16.23.1 New Features

None (released based 3.6-fp2)

16.23.2 Bug Fixes

- 6537 Removed AMI format for AWS S3
- 6521 Launching windows scan binary from command line with API key does not launch the scan
- 6517 Impossible to know which publish image on UForge corresponds to which Image in K5 portal
- 6515 CentOS 6 images can be accessed with SSH on K5
- 6513 Validation for K5 publish view is not properly handled
- 6511 Launching uforge-scan.exe from command prompt still fails if the file path includes Japanese characters
- 6507 The `uforge-install-config` binary for windows does not start because `uforge-install-profile-1-1.noarch.zip` does not contain the correct directory structure.
- 6505 The `no_console` file is not created for Windows.
- 6504 Problem with OpenDJ port 4444 usage in several UForge config scripts
- 6503 The `uforge.conf.Orig` contains plain passwords with very weak permission
- 6502 AWS connector uses a fixed size 3.4 GB disk and publication fails for larger images
- 6422 Uploading an avatar image twice, the first image is still used
- 6410 Loading page empty during 5 seconds for the first time in `Software Library` view
- 5897 If a space is used in cloud accounts in openstack in the URL, then an internal error is observed
- 5849 Displaying the logo in view package details of a target format is not displayed
- 6488 Impossible to generate image when install profile contains users
- 6362 AWS resource connector no longer work due to credential changes
- 6064 The CLI command `org repo update` returns exception if `--type` param value is invalid.
- 5900 Generation sometimes fails if the second disk of the appliance is too small

16.24 3.6.fp2

Release Date: 2016-12-05

16.24.1 New Features

- Fujitsu K5 support. Can now register machine images generated on the platform to Fujitsu K5.

Note: The following operating systems are supported for the moment (others will be supported soon):

- CentOS 7.0
 - Ubuntu 14.04
-

- SELinux support when creating appliance templates and during migration
- Docker machine image generation support. This allows users to build docker base images.
- When scanning Windows machines, the scan report now includes the services detected.

Note: The platform does not support the comparison of windows-based scans at this time.

16.24.2 RFEs

- Better progress status when scanning Windows machines
- Less restrictive validation of website information in the MySoftware/Project Overview
- New icons for ‘pull’ and ‘upload’ for software/project files management
- Added directory icon when displaying all the files for software/project files view
- When deleting a folder, the confirm message should be more explicit (that all sub folders and files will also be deleted)
- Better explanation of the “cached” option for software/project files in the UI
- Managing licenses for software/project components; there is now an explicit delete button to remove an uploaded license file

16.24.3 Bug Fixes

- 6123 Publishing a generation from a scan results in 500 error in UI
- 6089 Member’s role on workspace couldn’t be changed if language is set as French or Japanese
- 6017 Canceling from Appliance Create no longer returns to previous page
- 5946 Publishing to CloudStack fails with the next error: vhd.gz: No such file or directory
- 5942 RHEL is added despite launching *org os add* for Oracle Linux or Scientific Linux with cli
- 5909 User ID and group ID of the install profile can be set 0
- 5906 UserResourcesAccessRights database mapping not proxied
- 5896 Deployment fails due to NIC settings
- 5892 Deployment fails when using eth1
- 5843 “org category delete” raises an error
- 5777 Launching uforge-scan.exe from command prompt fails with an error if the file path to the binary includes Japanese characters.
- 5762 Cannot register the third disk with a VirtualBox image
- 5756 New users, the default country is: Abkhazia
- 5754 opening the Dashboard > Generations page first shows progress bar for all publications
- 5752 Number of MySoftware components not properly refreshed in the UI
- 5750 Number of Appliance not properly refreshed in the UI
- 5748 The diskusage of “uforge user quota list” is displayed by byte
- 5684 Invite the same user in the collaboration members list does not show error message
- 5676 Duplicated variable in /etc/default/grub if distribution provides default values.
- 5647 Keyboard and kernel parameters are not taken into the scan report on CentOS 7 scan.

- 5635 Broken incremental scan for windows 2012R2
- 5627 Cancelling scan via ctrl+c is not correctly displayed in the UI
- 5625 uforge-scan does not respect bandwidth limit
- 5623 When the image of CentOS7 is generated, RPM-GPG-KEY-CentOS import read fails
- 5621 rpmsgen fails to build package if file path in %file includes space.
- 5570 Impossible to delete an incremental scan
- 5562 UForge CLI accesses to interactive mode even if the user or password are wrong
- 5560 The input value of the activation key is not saved in a Windows appliance
- 5342 Scan incremental with Ubuntu does not appear in UI
- 5265 No dialog box displayed after running an instance on Azure

16.25 3.6.fp1

Release Date: 2016-10-31

16.25.1 New Features

- Import/Export of appliance templates in the user interface
- **Software (MySoftware) and Project bundles now consolidated. New features added including:**
 - pulling files from remote locations (HTTP, FTP endpoints) so the user no longer requires to upload software components to the platform
 - pulling files can be cached for future generations or pulled on each generation
 - file permissions added for files and directories
 - can create directory structures in a software bundle
 - can add tagging information to a software bundle
 - can add native packages from OS repositories to a software bundle
 - can add boot scripts to a software bundle
 - identify the software bundle only being supported on a subset of operating systems
- API keys can be used for authentication when running a scan for migration.
- Scan messages and error messages cleaned up and more understandable
- Japanese language localization for the UI

16.25.2 Bug Fixes

- 5293 Image generation error: Windows image must have at least 512 MB of memory
- 5729 Issues with migration from 3.5.1. to 3.6
- 5465 Build fails due to unreachable rpm-4.11.2.tar.bz2
- 5740 Fix DB schema checks

- 5331 AWS publish no longer works
- 5637 Windows generation from scan fails at boot
- 5427 Unable to generate a virtual machine with LVM inside a MSDOS disk
- 5291 All combo boxes are empty when a value has been selected
- 5876 Logo broken on Dashboard
- 5444 Unable to populate Fedora/RHEL distributions
- 5420 When a template is removed from a workspace, a DELETE error appears in the log file
- 5527 Subscription info does not list the frequency of quotas
- 5494 Scan fails because of files of type c (character device file)
- 5483 The service command not found in a machine generated by UForge
- 5442 The file deletion of Project fails
- 5429 Root can disable root account in UForge CLI
- 5746 Timeout of 10 seconds for the UForge CLI is not usable
- 5563 Internal error in Migration tab
- 5558 500 Call Fail Error when generating an image from scan
- 5556 The targetformat of Amazon is not displayed when generating an image
- 5553 If a scan is deleted, the image generated from the same scan is not deleted
- 5551 Spelling mistake in UI when publishing to Flexiant
- 5549 The error of Keystone version is displayed in Keystone Server URL
- 5403 Scan fails when trying to rebuild a non repo package

CHAPTER 17

Trademarks

UForge is a registered trademark of UShareSoft, a Fujitsu company.

LINUX is a registered trademark of Linus Torvalds.

Microsoft and Windows are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.

Oracle, Java, and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle Corporation and/or its affiliates.

Apache Ant, Ant, and Apache are trademarks of The Apache Software Foundation.

UNIX is a registered trademark of the Open Group in the United States and in other countries.

Red Hat Enterprise Linux is a trademark of Red Hat.

Other company names and product names are trademarks or registered trademarks of their respective owners.

CHAPTER 18

Copyright FUJITSU LIMITED 2019

All rights reserved, including those of translation into other languages. No part of this manual may be reproduced in any form whatsoever without the written permission of FUJITSU LIMITED

CHAPTER 19

High Risk Activity

The Customer acknowledges and agrees that the Product is designed, developed and manufactured as contemplated for general use, including without limitation, general office use, personal use, household use, and ordinary industrial use, but is not designed, developed and manufactured as contemplated for use accompanying fatal risks or dangers that, unless extremely high safety is secured, could lead directly to death, personal injury, severe physical damage or other loss (hereinafter “High Safety Required Use”), including without limitation, nuclear reaction control in nuclear facility, aircraft flight control, air traffic control, mass transport control, medical life support system, missile launch control in weapon system. The Customer shall not use the Product without securing the sufficient safety required for the High Safety Required Use. In addition, FUJITSU (or other affiliate’s name) shall not be liable against the Customer and/or any third party for any claims or damages arising in connection with the High Safety Required Use of the Product.

CHAPTER 20

Export Restrictions

Exportation/release of this document may require necessary procedures in accordance with the regulations of your resident country and/or US export control laws.