
Stackspin

Release 0.8.1

Greenhost

May 13, 2022

INSTALLATION

1	Installation overview	3
1.1	Preparing the provisioning machine	3
1.2	Create a python virtual environment (Optional)	4
1.3	Install requirements	4
1.4	Stackspin CLI client usage	4
2	Create and configure a kubernetes cluster	5
3	Install Stackspin	9
3.1	Step 1: Flux configuration	9
3.2	Step 2: Install core applications	10
3.3	Step 3: Install additional applications	11
3.4	Step 4: Validate setup	12
3.5	Step 5: Let us know!	12
4	Test your installation	13
4.1	Stackspin installation	13
4.2	Command line tests	13
4.3	Testing Instructions for all apps	13
4.4	Testing Instructions for specific apps	13
4.5	Providing feedback	15
5	Manage Stackspin	17
5.1	Dashboard	17
5.2	Users	17
6	Applications	19
6.1	Core applications	19
6.2	Optional applications	19
7	Logging	23
7.1	Central log aggregation	23
8	Monitoring	27
8.1	Grafana	27
8.2	Prometheus	27
8.3	Alertmanager	28
8.4	Email alerts	28
9	Maintenance	29
9.1	Backup	29

9.2	Restore	29
9.3	Change the IP of your cluster	31
9.4	Delete evicted pods	31
9.5	Apply changes to flux variables	31
10	Upgrade guide	33
10.1	Upgrading to 0.8	33
10.2	Upgrading to 0.7.0	35
10.3	Upgrading to 0.6.0	38
10.4	Upgrading from 0.4.0 to 0.5.0	38
10.5	Upgrading from 0.3.0 to 0.4.0	38
10.6	Upgrading to 0.3.0	39
11	Customizing	41
11.1	Prerequisites	41
11.2	Customize Stackspin applications	41
11.3	Adding custom apps to the cluster	42
12	Troubleshooting	43
12.1	Known issues	43
12.2	Run the CLI tests	43
12.3	SSH access	46
12.4	Using kubectl to debug your cluster	46
12.5	HTTPS certificates	47
12.6	Application installation or upgrade failures	49
12.7	Purge Stackspin and install from scratch	50
13	Security	51
13.1	Access control	51
14	Stackspin Design	53
14.1	Application build pipeline	53
14.2	Configuration	55
14.3	Application containers	56
14.4	Persistent data	56
14.5	Automatic updates	56
14.6	APIs	57
15	Instructions	59
15.1	Building a new version of Zulip	59
16	Reference documentation	61
16.1	File structure	61
17	Comparable projects	63
17.1	What makes Stackspin unique	63
17.2	Platforms	63

Stackspin is a platform that offers self-managed, click-and-play provisioning of online applications for Civil Society Organisations (CSOs). Users will be able to easily set up a self-hosted instance of Stackspin, so they can keep control over the data that gets entered into these applications.

Stackspin is:

- Open Source
- Self updating
- Easy to deploy
- Integrated

For more information, go to [the Stackspin website](#).

INSTALLATION OVERVIEW

Warning:

- Stackspin is still under heavy development and is not ready for production use! We anticipate major changes and do not guarantee a data-preserving upgrade path from current installations. However, we encourage you to try Stackspin and ask you to [report all issues you encounter](#).
- When you install Stackspin on a server, the installation process will make some substantial changes to the server's configuration, so please do not use a server that functions as anything other than a testing ground.

This page will guide you through the installation of the “Stackspin CLI”. The CLI can be used to set up Stackspin on a different machine that is connected to the internet. It also functions as a starting point for running tests.

1.1 Preparing the provisioning machine

You should have a trusted machine to run the installer on, this can be your laptop or pc. In this guide, we call it the *provisioning machine*. All commands in these installation instructions need to be run the provisioning machine, except when a command starts with `cluster$`

The following software setup is required **on your provisioning machine**:

- You need Python 3 with its development files, Pip and Git installed. On Debian/Ubuntu: `apt install python3-pip python3-dev git`
- We recommend using a [python virtualenv](#) to make sure Stackspin does not change any of your other python projects. Install virtualenv by running `pip3 install --user venv` or `apt install python3-venv`.
- You need a recent version of Pip. If you run into problems during `pip install` you might need to run `pip install --upgrade pip`.
- `kubectl` ([installation instructions](#))
- flux version 0.25.3 ([Download flux_0.25.3_linux_amd64.tar.gz](#) or use the official installation script: `curl -s https://fluxcd.io/install.sh | sudo FLUX_VERSION=0.25.3 bash`)

Then you can clone the Stackspin git repository and checkout the latest release branch (currently `v0.8`):

```
$ git clone -b v0.8 https://open.greenhost.net/stackspin/stackspin.git
$ cd stackspin
```

Apart from the previously installed software, Stackspin will only modify in the `clusters` subdirectory of the repository. So multiple provisioning machines can administer the same instance so long as that subdirectory is synced between them.

1.2 Create a python virtual environment (Optional)

Create a python virtual environment called “env” that uses python 3. This makes sure we do not change any of your other python projects. The second command “activates” the virtualenv.

Note: Activating the virtualenv means you will use that environment to install and run python programs instead of your global environment. If you close your terminal or open a new one, you need to activate the virtualenv again.

```
$ python3 -m venv env
$ . env/bin/activate
```

1.3 Install requirements

Next, install the Stackspin CLI client by running the following commands:

```
$ pip3 install -r requirements.txt
```

1.4 Stackspin CLI client usage

Now you can run the Stackspin CLI as follows:

```
$ python -m stackspin CLUSTER_NAME <command>
```

The CLI *always* needs a CLUSTER_NAME argument. Even for getting subcommand help messages. Be sure to run this command in the root directory of the git repository. In this tutorial, we’re using `stackspin.example.org` as the cluster name. Try it out by running:

```
$ python -m stackspin stackspin.example.org --help
```

The next chapter, *Create and configure a kubernetes cluster* will explain setting up and/or configuring a Kubernetes cluster to install Stackspin on.

CREATE AND CONFIGURE A KUBERNETES CLUSTER

Below, choose whether you want to install Stackspin on a VPS or an existing Kubernetes cluster. We recommend manually installing it on a VPS, this is the most tested solution.

A. Install on a VPS

We recommend manually installing Stackspin on a Greenhost VPS as explained below. However, if you plan on installing many Stackspin instances, click the tab “*ADVANCED: Use the Greenhost API*” to learn how to automatically create VPSs and DNS records and install Stackspin on them.

This document describes how you can use Stackspin’s CLI to install `k3s`, a lightweight Kubernetes distribution, on a virtual private server.

For Stackspin we will set up a “single-node” kubernetes cluster. This means everything runs on the same VPS. Support for “multi-node” clusters (a Kubernetes cluster on more than one VPS) will come in the future.

Manually install on a VPS

Prerequisites

These instructions assume you have a VPS or other kind of server ready that meets the following requirements:

- Debian 11 (bullseye) installed
- A public IP address
- The ability to create DNS records for this IP
- 6 cores and 12 GB of RAM
- At least 25GB of disk space for installation, plus more for application data. We recommend starting with at least 30GB.
- Root ssh access

You’ll need the machine’s *IP address*.

This guide assumes the VPS’s IP address is `1.2.3.4` and we install on the domain `stackspin.example.org`. Substitute these two variables with your IP address and the (sub)domain you want Stackpin to be available on.

Step 1: Create cluster configuration

To create a config for your cluster, use the `create` subcommand of the Stackspin CLI. First, choose a name (we chose `stackspin.example.org`) for your cluster. Then run the following command to get information about the `create` subcommand:

```
$ python -m stackspin stackspin.example.org create --help
```

Create the initial Stackspin configuration for your VPS by running the following command:

```
$ python -m stackspin stackspin.example.org create \  
stackspin.example.org \  
--ip-address 1.2.3.4
```

This configures your cluster under the fully qualified domain name (FQDN) `stackspin.example.org`. To break down the command:

- the first, positional argument `stackspin.example.org` tells the cluster the domain it will be hosted on. This should be a (subdomain of a) domain you own.
- `--ip-address 1.2.3.4` tells the script the IP address of your VPS. This will be used to find the VPS during the installation procedure.

The configuration has now been written to the `clusters/stackspin.example.org` on your provisioning machine.

ADVANCED: Use the Greenhost API

Prerequisites

Before you can start, you need the following things:

1. An API key with Customer rights.

Make sure the Stackspin CLI can use your token by storing it in a `COSMOS_API_TOKEN` variable.

```
$ export COSMOS_API_TOKEN=<paste your API key here>
```

In *the same terminal*, you can now use the `create` subcommand

2. An SSH key ID from the Cosmos Service Centre’s “VPS Cloud” section. This SSH key will get Root access on the VPS.
 - You can find your SSH key ID by going to VPS Cloud -> SSH keys and checking the link under “Show key”. The numerical part is your SSH key ID.
 - You can also use the API to list ssh keys and find it there. Read the [Greenhost API documentation](#) for more information*
3. The ability to create DNS records that point to the VPS that gets created

Step 1: Create VPS and cluster configuration

We’ll use the `create` subcommand of the Stackspin CLI to create a VPS and a corresponding configuration file on the provisioning machine.

Here is an example of a complete creation command:

```
$ python -m stackspin stackspin.example.org create \  
--create-droplet \  
--create-hostname stackspin.example.org \  
--ssh-key-id 112 \  
--create-domain-records \  
--subdomain stackspin \  
example.org
```

Let’s break down the arguments:

- `--create-droplet`: Use the Greenhost API to create a new VPS
- `--create-hostname stackspin.example.org`: Create a VPS with hostname `stackspin.example.org`
- `--ssh-key-id 112`: Use SSH key ID 112 (you can find your SSH key ID in the [Cosmos Service Centre](#) under *VPS Cloud -> Installation SSH Keys*. Hover over a button there to see the ID in the URL it uses.
- `--create-domain-records`: Use the Greenhost API to create DNS records. If you do this, you can skip *Step 2: DNS Configuration*. The following records are automatically created:
 - An A record `stackspin.example.org` pointing to the VPS's IP address
 - A CNAME record `*.stackspin.example.org` pointing to `stackspin.example.org`.
- `--subdomain stackspin`: Only needed when you use `--create-domain-records` so the Greenhost API can find your domain. If you use the `--subdomain` argument, the positional argument **must not** include the subdomain (i.e., in the example command the positional argument is `example.org` and not `subdomain.example.org`).

If you want to host Stackspin on a root domain, you can omit the subdomain argument.

Step 2: DNS Configuration

Stackspin needs two DNS records that point to your cluster. It is important to create these DNS records before you start the installation, so Stackspin can secure your connections with https certificates from [Let's Encrypt](#).

Create these two DNS records:

- An A record `stackspin.example.org` pointing to the cluster's IP address,
- A CNAME record `*.stackspin.example.org` pointing to `stackspin.example.org`.

Note: It is also possible to host Stackspin on a domain without a dedicated subdomain. That does imply that the included WordPress site will be hosted on your root domain (i.e. `https://example.org`). In that case, make these DNS records instead:

- An A record `example.org` pointing to the cluster's IP address,
- A CNAME record `*.example.org` pointing to `example.org`.

Step 3: Create Kubernetes cluster

Before you continue, make sure your DNS configuration is propagated. You can do this by making sure ping shows your VPS's IP address:

```
$ ping stackspin.example.org
# Example output of the command:
PING stackspin.example.org (1.2.3.4) 56(84) bytes of data.
64 bytes from 1.2.3.4 (1.2.3.4): icmp_seq=1 ttl=62 time=9.43 ms
64 bytes from 1.2.3.4 (1.2.3.4): icmp_seq=2 ttl=62 time=52.5 ms
...
```

If you've never logged in to this machine before, you might need to accept the SSH hostkey before you can start the installation process. To be sure, log in to your VPS once before you continue:

```
$ ssh root@stackspin.example.org
```

If necessary, check your hostkey fingerprint and type `yes` to continue. You should get a shell on your VPS. Type `exit` to quit the shell and continue on your provisioning machine.

Run the `install` command with the CLI to completely configure your VPS for Stackspin.

```
$ python -m stackspin stackspin.example.org install
```

This will take a few minutes. It tries to connect to your VPS as the root user, then installs [k3s](#). In addition it installs command line tools [kubectl](#) (Kubernetes cli tool), [flux](#) (used for automated updates) and [velero](#) (Kubernetes resources and persistent volumes backup) so you can manage the Kubernetes cluster from a shell on the VPS.

Note: It is possible to re-run the `install` command with a newer version of the installation script. Refer to the [Upgrade guide](#) for details.

B. Install on existing cluster

Below, we describe how you can install Stackspin on an existing Kubernetes cluster. This process can be followed if you have a “managed Kubernetes” setup, or if you want to use a Kubernetes version other than K3s.

Note:

- Installation on an existing Kubernetes cluster is still experimental and might not work depending on the setup of your cluster/cloud provider. We are happy to receive feedback from your experiences but at the moment we cannot make any guarantees!
-

Prerequisites

- A single-node Kubernetes cluster (multi-node support has not been tested by the Stackspin team)
- A `kube_config.yml` file for API access

Configure DNS

Stackspin needs two DNS records that point to your cluster. It is important to create these DNS records before you start the installation, so Stackspin can secure your connections with https certificates from [Let's Encrypt](#).

Create these two DNS records:

- An A record `stackspin.example.org` pointing to the cluster's IP address,
- A CNAME record `*.stackspin.example.org` pointing to `stackspin.example.org`.

Note: It is also possible to host Stackspin on a domain without a dedicated subdomain. That does imply that the included WordPress site will be hosted on your root domain (i.e. `https://example.org`). In that case, make these DNS records instead:

- An A record `example.org` pointing to the cluster's IP address,
 - A CNAME record `*.example.org` pointing to `example.org`.
-

Setup

- Create a directory containing your cluster configuration, i.e. `mkdir -p clusters/stackspin.example.org`
- Copy your `kube_config.yml` file inside your cluster config directory and rename it to `kube_config_cluster.yml`: `cp kube_config.yml clusters/stackspin.example.org/kube_config_cluster.yml`

Now that you have a Kubernetes cluster running you can continue to [Install Stackspin](#).

INSTALL STACKSPIN

This guide explains how to install Stackspin on the Kubernetes cluster you have configured in the previous step.

3.1 Step 1: Flux configuration

Flux will run inside your Stackspin cluster to install and upgrade applications. It needs to be configured once, using the `flux` command line tool and scripts provided by us in the Stackspin repository.

3.1.1 Configuration

Copy the file `install/.flux.env.example` to your cluster dir `clusters/stackspin.example.org/.flux.env`. This file contains the last bit of information you need to configure. Make sure not to put any quotes around your values, because that can make the resulting yaml files invalid. Also, using a dollar sign `$` or double quote `"` may lead to problems, so please avoid using those characters in the values.

Cluster information

```
# The IP address of your cluster
ip_address=1.2.3.4
# The FQDN of your cluster
domain=stackspin.example.org
# The system administrator's email address. Alerts are sent to this address
admin_email=admin@example.org
```

Outgoing email

Stackspin uses SMTP to send emails. This is essential for finishing account setups with password recovery links. Additionally, apps like Nextcloud, Zulip and Wordpress will be able to send email notifications from the email address configured here. You also may receive alert notification emails from Stackspin's monitoring system. See [Email alerts](#) for more information about those alerts, especially during installation.

Because Stackspin does not include an email server, you need to search your (external) email provider's helpdesk for SMTP configuration details.

Edit the `clusters/stackspin.example.org/.flux.env` file as follows:

```
# Enable sending mails
outgoing_mail_enabled=true
# Email address that the cluster sends mails from. Needs to be an existing SMTP
# login
outgoing_mail_from_address=admin@example.org
# Same outgoing mail address, but only the part before the '@'
outgoing_mail_from_local_part=admin
# Same outgoing mail address, but only the part after the '@'
outgoing_mail_domain=example.org
# SMTP password for the outgoing mail address
outgoing_mail_smtp_password=CHANGEME
# SMTP username, often the same as the outgoing email address
outgoing_mail_smtp_user=admin@example.org
# SMTP login data.
outgoing_mail_smtp_host=smtp.greenhost.nl
outgoing_mail_smtp_authtype=LOGIN
outgoing_mail_smtp_port=587
```

Backups with Velero (Optional)

You can enable [Velero](#), a program that runs on your cluster and uploads backups of your cluster and user data to an S3 storage service of your choice.

If enabled, Velero will create a backup of your cluster once every night and upload it to the S3 storage you configure. This includes:

- your cluster state. Technically speaking, it will back up all Kubernetes resources in your cluster; this includes things like which applications are installed, including their version number and installation-time settings;
- persistent data of all applications: for example, single sign-on users that you created, Nextcloud files and meta-data, WordPress site data and comments, Zulip chat history, etc. A single exception to this is Prometheus data (statistics of system properties), which takes up a lot of space and we consider not valuable enough to back up.

It does not include anything on the VPS that you may have set up but is not part of Stackspin, like programs installed via `apt`, or data added to the VPS disk not through Stackspin.

To configure Velero, edit the file `clusters/stackspin.example.org/.flux.env`, and configure the settings with the `backup_s3_` prefix.

Then continue with the installation procedure as described below. At the end of the installation procedure, you have to install the `velero` application.

For information on how to use Velero with Stackspin, please see [Backup](#).

3.2 Step 2: Install core applications

Before you can start, you need to execute a few commands from the installation directory **on your provisioning machine**. Don't forget to replace `stackspin.example.org` with your domain.

We will use this variable in the following commands, set it to your cluster directory.

```
export CLUSTER_DIR=$PWD/clusters/stackspin.example.org
```

Make sure your `virtualenv` is activated.

```
. env/bin/activate
```

Copy the installation kustomization to your cluster directory.

```
cp install/kustomization.yaml $CLUSTER_DIR/
```

Tell kubectl to use your cluster's kube_config.

```
export KUBECONFIG=$CLUSTER_DIR/kube_config_cluster.yaml
```

Ensure flux-system namespace is created.

```
kubectl get namespace flux-system 2>/dev/null || kubectl create namespace flux-system
```

This inserts the configuration from .flux.env into your cluster as a “secret”.

```
kubectl apply -k $CLUSTER_DIR
```

After you have executed that code, your terminal should show:

```
secret/stackspin-cluster-variables created
```

Next, run:

```
./install/install-stackspin.sh
```

This installs the *core* of Stackspin into your cluster. To see what's included, check the `flux2/infrastructure` and the `flux2/core` folders in the [git repository](#).

3.3 Step 3: Install additional applications

After the script completes, you can install applications by running the other installation scripts in the `install` folder. At the moment, we have scripts to install:

- Monitoring stack (Prometheus, Grafana, Loki, Eventrouter) `./install/install-app.sh monitoring`
- Nextcloud and Onlyoffice with `./install/install-app.sh nextcloud`
- Zulip chat with `./install/install-app.sh zulip`
- Wekan with `./install/install-app.sh wekan`
- WordPress with `./install/install-app.sh wordpress`
- Velero with `./install/install-app.sh velero` (only if you have configured it in *Backups with Velero (Optional)*).

When the installation scripts complete, the application installation may still be running on the Stackspin cluster. You can monitor the progress by running `flux get kustomizations` (use `watch flux get kustomizations` to get updates). If all kustomizations have been applied correctly, you can monitor specific application releases by running `watch flux get helmreleases --all-namespaces`.

3.4 Step 4: Validate setup

Once the installation has been completed, you can log in on <https://dashboard.stackspin.example.org> (as always: replace *stackspin.example.org* with your domain). To get your login details, run:

```
$ python -m stackspin stackspin.example.org admin-credentials
```

Additionally, because Stackspin is still under development, we would like you to follow our *Test your installation* to make sure that the setup process went well.

3.5 Step 5: Let us know!

We would love to hear about your experience installing Stackspin. If you encountered any problems, please create an issue in our [issue tracker](#). If you didn't please still reach out as described on our [contact page](#) and tell us how you like Stackspin so far. We want to be in communication with our users, and we want to help you if you run into problems.

TEST YOUR INSTALLATION

Great that you want to take Stackspin for a test drive! This guide contains instructions to get you going, some pointers on what we think would be useful to test, and guesses at what results of those tests would be useful to write down. At any point please feel invited to test whatever functionality you come across, and reporting whatever you think is interesting. Our contact details are listed [here](#), and we'll describe how to give feedback via our issue tracker at the *end of these instructions*.

During these instructions, please replace *example.org* with your own domain name.

4.1 Stackspin installation

First we'd like you to setup a Stackspin cluster by yourself, following the *Installation overview* documentation and make sure you complete all steps.

4.2 Command line tests

Please *Run the CLI tests* which checks the overall functionality of your cluster and include the output in your feedback.

4.3 Testing Instructions for all apps

Please log in on <https://dashboard.stackspin.example.org> using single sign-on with your admin user email address and see if you have admin rights granted (usually there's an app specific admin panel available or admin functionality like configuring users). Afterwards logout and log in again as the non-admin single-sign-on user you created earlier in the Stackspin admin panel. You should *not* have any admin privileges now.

4.4 Testing Instructions for specific apps

4.4.1 Nextcloud

Please browse to Nextcloud using the link from your user dashboard app list and try to log in using single sign-on. Use the button labeled **Login with Stackspin**. Please try logging in with your admin account and configure the email settings as shown in the Usage doc. After that please log in with the user you created in the user panel.

Files & synchronization

Please try uploading, downloading, moving or copying files.

Calendar

Please test the basic functionality of the calendar, including advanced features like recurrence, custom notifications or inviting attendees.

Passwords

Please create and manage different passwords. In addition to that please try

Nextcloud sync applications

- Please try syncing with your smartphone or desktop using one of the apps mentioned in *Nextcloud*.

4.4.2 Onlyoffice

Creating a new office document

From the main Nextcloud webpage, please try to create a new office document, by clicking the round plus button near the top of the screen, then picking the Document type with the blue icon (third one from below on my screen), and enter a name for it. After that, please try some basic editing of the document, and save it. Maybe check you can open it again afterwards, and that it has the contents that you saved earlier.

Collaborating on an office document

This part of the test requires the cooperation of another person; feel free to skip it now if that's not convenient at this point.

- First, try to share your document with a different user.
- Then, try to open the shared document from a few different user accounts simultaneously, and let all participants edit the document mercilessly. There are also some collaboration features that you may want to try: on the left of the Onlyoffice screen there are buttons for chat and for text comments.

4.4.3 Wordpress

Using the link in the dashboard app, please try to log in to WordPress as the user you created earlier by pressing “Log in” and using the Login with OpenID Connect button.

Your user needs “Admin” privileges to edit the website.

4.5 Providing feedback

If you have not done so already, please create an account on <https://open.greenhost.net> (or log in with your existing github account) and [create a new issue](#) using the Feedback template.

Thanks a lot for your testing work! We'll use your input to try to improve Stackspin.

MANAGE STACKSPIN

After all the applications are installed, the first thing to do is log in to <https://dashboard.stackspin.example.org>. The Dashboard is the starting point for Stackspin administrators. Here you can create, edit and delete users, as well as see all the applications that run on your Stackspin cluster.

Your first login can happen in one of two ways:

1. Use the email address that was used during setup.

The admin email address has an account already. If you have access to its mailbox, press “Sign in” in the browser and then press “Forget password?” to start a password recovery procedure.

2. If you do not have access to the administrator email address, ask your system administrator to run the following command on their “provisioning machine” and share the admin credentials with you:

```
python3 -m stackspin stackspin.example.org admin-credentials
```

Note: these are the setup credentials, they will not be updated when you change the admin user’s password.

5.1 Dashboard

After logging in, you will see an overview of all the installed applications. The “Launch app” button on this screen will bring you to the application’s screen. If see a login screen, use the “log in with Stackspin” button to log in with your current user. You will not have to enter credentials, because you are already logged in via the Dashboard.

5.2 Users

The Dashboard also features a “Users” tab, where you can manage the Stackspin users. At the moment, all users will be able to log into all applications on your cluster with the “Log in with Stackspin” button. More fine grained control will be added in the future.

To create a new user, click the “Add new user” button. You need to specify the user’s (full) name under “Name” and a valid email address under “Email”.

Once you have created the user, they will not be automatically notified - please contact individual users and let them know they can finish their account setup by opening any Stackspin-related URL and clicking the “Forget password?” button on the “Login with Stackspin” screen. This will let them set up their password for the first time, and their account will be active from that moment.

Note: The Dashboard currently allows all users to manage the other users. In the (near) future, only users with “admin” privileges will be allowed to add, edit and remove other users.

Only add users for people that you trust.

To edit a user's name or email address, or delete a user, press the "Configure" button on the right side in the user table. If you want to delete several users, you can select them in the Users overview, using the select buttons on the left side. Then press the "Delete" button on the top right of the table.

Note: A note on deleting users

Deleting users from the Dashboard makes sure they can't log into applications anymore. However, the dashboard cannot control the user's data inside other applications.

To completely remove users from all the applications, you will need to remove them from inside the application's admin settings as well. Refer to the documentation of the specific application to see how that should be done.

If you only remove a user from the dashboard, but not from the application, it is possible that this user keeps access to the application through sessions inside application clients (for example the Zulip or Nextcloud desktop clients). These sessions can also typically be ended through the application's admin interface.

APPLICATIONS

6.1 Core applications

These applications are available after the installation is completed successfully:

6.1.1 Dashboard

The Dashboard is an application developed specifically for Stackspin. It allows you to manage Stackspin’s central user database. Soon, the dashboard will also allow you to configure, install and remove applications and set access rights for users.

Administrators can read how to use the Dashboard in *Manage Stackspin*

6.1.2 Grafana

[Grafana](#) is a dashboard application that shows you information about the status of your cluster collected by [Prometheus](#).

Single sign-on users

Users that have the “Admin” label turned on in the user panel, will have admin rights in Grafana. Other users are able to see graphs, but can not change anything.

6.2 Optional applications

6.2.1 Nextcloud

[Nextcloud](#) is a file sharing and communication platform. These Nectcloud apps can be accessed from the top navigation bar:

Files & synchronization

You can access your files with the Nextcloud Web interface and create, preview, edit, delete, share, and re-share files. See the [Files & synchronization user manual](#) for general usage and [Desktop and mobile synchronization](#) for setting up file sync on your mobile or desktop device.

Calendar

The Nextcloud Calendar app works similar to other calendar applications you can sync your Nextcloud calendars and events with. Follow the [Calendar user manual](#) for general usage and the *Nextcloud groupware docs* <https://docs.nextcloud.com/server/latest/user_manual/th/pim/index.html> for syncing your calendars with your mobile or desktop devices.

Passwords

A simple, yet feature rich password manager for Nextcloud. See [Password user handbook](#) for more details, including using the [Browser extensions](#).

These are the mobile apps that you can use to access your passwords from your smartphone:

- [NC passwords app for Android (by joleaf)](<https://gitlab.com/joleaf/nc-passwords-app>)
- [Nextcloud password app for Android (by daper)](<https://github.com/daper/nextcloud-passwords-app>)
- [Nextcloud passwords for iOS](<https://github.com/johannes-schliephake/nextcloud-passwords-ios>)

You'll find how to configure file or calendar sync with your smartphone or desktop in the [Nextcloud Groupware documentation](#).

6.2.2 Onlyoffice

[Onlyoffice](#) is an online document editing suite. You can open documents in Onlyoffice by clicking them in Nextcloud. You can open new documents by clicking the “Plus” button in Nextcloud and selecting Document, Spreadsheet or Presentation.

6.2.3 Zulip

[Zulip](#) is a team chat application.

Unfortunately, Zulip does not yet listen to the Stackspin roles, so once you have installed Zulip, you need to make one of your users the owner. That allows this user to manage Zulip's settings from inside the application. This owner can also edit other users' roles.

For this, you have to replace <your_email> in the the following command and run it on your cluster:

```
` kubectl exec -n stackspin-apps zulip-0 -c zulip -- sudo -u zulip /home/zulip/
deployments/current/manage.py change_user_role <your_email> owner --realm 2 `
```

6.2.4 Wekan

Wekan is a Kanban board application.

6.2.5 WordPress

WordPress is a website content management system.

Known issues

If you want admin privileges inside WordPress, you need to set this in the Dashboard. Then, you need to log in to WordPress again.

If this is your first time logging in to WordPress, you have to log out and in again for your admin privileges to be applied.

For more information and technical details, refer to [the issue in our issue tracker](#).

LOGGING

Logs from pods and containers can be read in different ways:

- In the cluster filesystem at `/var/log/pods/` or `/var/logs/containers/`.
- Using `kubectl logs`
- Querying aggregated logs with Grafana, see below.

7.1 Central log aggregation

We use `Promtail`, `Loki` and `Grafana` for easy access of aggregated logs. The [Loki documentation](#) is a good starting point how this setup works. There are two ways of viewing aggregated logs:

- Via the Grafana web interface
- Using the `logcli` command line tool

7.1.1 Viewing logs in Grafana

The [Using Loki in Grafana](#) gets you started with querying your cluster logs with Grafana. You will find the Loki Grafana integration on your cluster at <https://grafana.stackspin.example.org/explore> together with some generic query examples.

Please follow [LogQL query examples](#) for more LogQL query examples.

7.1.2 Query logs with logcli

Please refer to `logcli` for installing `logcli` on your Laptop. The create a port-forwarding to your cluster using the `kubectl` tool:

```
kubectl -n stackspin port-forward pod/loki-0 3100
```

In another terminal you can now use `logcli` to query loki like this:

```
logcli query '{app=~".+"}'
```

Please follow [LogQL query examples](#) for more LogQL query examples.

Search older messages (in this case the last week and limit the output to 2000 lines):

```
logcli query --since=168h --limit=2000 --forward '{app="helm-controller"}'
```

7.1.3 LogQL query examples

Please also refer to the [LogQL documentation](#) and the [log queries documentation](#) .

Query all aggregated logs (unfortunately we can't find a better way of doing this since LogQL always expects a stream label to get queried):

```
{app=~".+"}
```

Query all logs for a keyword:

```
{app=~".+"} |= "error"
```

Query all k8s apps for errors using a regular expression:

```
{app=~".+"} |~ `(error|fail|exception|fatal)`
```

Flux

Flux is responsible for installing applications. It uses four controllers:

- `source-controller` that tracks Helm and Git repositories like <https://open.greenhost.net/stackspin/stackspin> for updates.
- `kustomize-controller` to deploy kustomizations that often install helmreleases.
- `helm-controller` to deploy the helmreleases.
- `notification-controller` that is responsible for inbound and outbound flux messages

Query all messages from the `source-controller`:

```
{app="source-controller"}
```

Query all messages from `flux` and `helm-controller`:

```
{app=~"(source-controller|helm-controller)"}
```

`helm-controller` messages containing `wordpress`:

```
'{app = "helm-controller"} |= "wordpress"'
```

`helm-controller` messages containing `wordpress` without `unchanged` events (to only show the installation messages):

```
'{app = "helm-controller"} |= "wordpress" != "unchanged"'
```

Filter out redundant `helm-controller` messages:

```
'{app="helm-controller"} !~↵  
↵ `(unchanged|event=refreshed|method=Sync|component=checkpoint)`'
```

Cert-manager

Cert manager is responsible for requesting Let's Encrypt TLS certificates.

Query cert-manager messages containing chat:

```
'{app="cert-manager"} |= "chat"'
```

Hydra

Hydra is the single sign-on system.

Show only warnings and errors from hydra:

```
{container_name="hydra"} != "level=info"
```

Debug oauth2 single sign-on with zulip:

```
{container_name=~"(hydra|zulip)"}
```

Etc

Query kubernetes events processed by the eventrouter app containing warning:

```
'{app="eventrouter"} |~ "warning"'
```


MONITORING

For monitoring your Stackspin cluster we included the [kube-prometheus-stack](#) helm chart, which bundles the applications [Grafana](#), [Prometheus](#) and [Alertmanager](#), and also includes pre-configured Prometheus alerts and Grafana dashboards.

8.1 Grafana

Grafana can be accessed by clicking on the **Monitoring** icon in the **Utilities** Section of the dashboard. Use Stackspin single sign-on to login.

8.1.1 Dashboards

Browse through the pre-configured dashboards to explore metrics of your Stackspin cluster. Describing every dashboard would be too much here, reach out for us if you don't find what you are looking for.

8.1.2 Browse aggregated logs in Grafana

See *Viewing logs in Grafana* how to do this.

8.2 Prometheus

Prometheus can be reached by adding `prometheus.` in front of your cluster domain, i.e. `https://prometheus.stackspin.example.org`. Until we [configure single sign-on for prometheus](#) you need to login using basic auth. The user name is `admin`, the password can get retrieved by running

```
python -m stackspin CLUSTERNAME secrets | grep prometheus-basic-auth
```

8.3 Alertmanager

Alertmanager can be reached by adding `alertmanager.` in front of your cluster domain, i.e. `https://alertmanager.stackspin.example.org`. Until we [configure single sign-on for prometheus](#) you need to login using basic auth. The user name is `admin`, the password can get retrieved by running

```
python -m stackspin CLUSTERNAME secrets | grep alertmanager-basic-auth
```

8.4 Email alerts

From time to time you might get email alerts sent by [Alertmanager](#) to the email address you have set in the cluster configuration. Common alerts include (listed by the `alername` references in the email body):

- **KubeJobCompletion:** A job did not complete successfully. Often happens during initial setup phase. If the alert persists use i.e. `kubectl -n stackspin-apps get jobs` to see all jobs in the `stackspin-apps` namespace and delete the failed job to silence the alert with i.e. `kubectl -n stackspin-apps delete job nc-nextcloud-cron-27444460`.
- **ReconciliationFailure:** A `flux helmRelease` could not get reconciled successfully. This also happen often during initial setup phase. It can have different root causes though. Use `flux -n stackspin-apps get helmreleases` to view the current state of all `helmReleases` in the `stackspin-apps` namespace. In case the `helmRelease` in question is stuck in a `install retries exhausted` or `upgrade retries exhausted` state you can force a reconciliation with

```
flux -n stackspin-apps suspend helmrelease zulip  
flux -n stackspin-apps resume helmrelease zulip
```

Depending on the underlying cause this will fix the `helmRelease` state or not. For more information on this issue see [helmrelease upgrade retries exhausted regression](#)

MAINTENANCE

9.1 Backup

9.1.1 On your provisioning machine

During the installation process, a cluster config directory is created on your provisioning machine, located in the top-level sub-directory `clusters` in your clone of the `stackspin` git repository. Although these files are not essential for your Stackspin cluster to continue functioning, you may want to back this folder up because it allows easy access to your cluster.

9.1.2 On your cluster

Stackspin supports using the program `Velero` to make backups of your Stackspin instance to external storage via the S3 API. See *Backups with Velero (Optional)* in the installation instructions for setup details.

For the maintenance operations described below – in particular, restoring backups – you need the `velero` client program installed, typically on your provisioning machine although you can also run it on the VPS if preferred. You may find it at [Velero's github release page](#).

By default `Velero` will make nightly backups of the entire cluster (minus Prometheus data). To make a manual backup, run

```
cluster$ velero create backup BACKUP_NAME --exclude-namespaces velero --wait
```

from your VPS. See `velero --help` for other commands, and [Velero's documentation](#) for more information.

Note: in case you want to make an (additional) backup of application data via alternate means, all persistent volume data of the cluster are stored in directories under `/var/lib/Stackspin/local-storage`.

9.2 Restore

Restoring from backups is a process that for now has to be done via the command line. We intend to allow doing this from the Stackspin dashboard instead in the near future.

These instructions explain how to restore the persistent data of an individual app (such as Nextcloud, or Zulip) to a previous point in time, from a backup to S3-compatible storage made using `velero`, on a Stackspin cluster that is in a healthy state. Using backups to recover from more severe problems, like a broken or completely destroyed Stackspin cluster, is also possible, by reinstalling the cluster from scratch and restoring individual app data on top of that. However, that procedure is not so streamlined and not documented here. If you are in that situation, please [reach out to us](#) for advice or assistance.

9.2.1 Select backup

To show a list of available backups, perform the following command on your VPS:

```
kubectl get backup -A
```

Once you have chosen a backup to restore from, record its name as written in the `kubectl` output.

Note: Please be aware that for technical reasons the restore operation will restore not only the persistent data from this backup, but also the app’s software version that was running at that time. Although the auto-update mechanism should in turn update the app to a recent version, and the recent app version should be able to automatically perform any necessary data format migrations on the old data, this operation has not been well tested for older backups, so please proceed carefully. As an example of what could go wrong, Nextcloud requires upgrades to be done in a serial fashion, never skipping a major version upgrade, so if your backup is from two or more major Nextcloud versions ago, some manual intervention is required. If you have any doubts, please [reach out to us](#).

9.2.2 Restore app data

Warning: Please note that restoring data is a destructive operation! It will replace the app’s data as they are now. There is no way to undo a restore operation, unless you have a copy of the current app data, in the form of a current Stackspin backup or an app-specific data export. For that reason, we recommend making another backup right before beginning a restore operation.

To restore the data of app `$app` from the backup named `$backup`, perform the following commands:

```
flux suspend kustomization $app
flux suspend helmrelease -n stackspin-apps $app
kubectl delete all -n stackspin-apps -l stackspin.net/backupSet=$app
kubectl delete pvc -n stackspin-apps -l stackspin.net/backupSet=$app
velero restore create arbitrary-name-of-restore-operation --from-backup=$backup -l_
↪stackspin.net/backupSet=$app
# At this point, please first wait for the restore operation to finish, see
# text below.
flux resume helmrelease -n stackspin-apps $app
flux resume kustomization $app
```

Note: Specifically for Nextcloud, the `kubectl delete pvc ...` command might hang due to a Kubernetes job that references that PVC. To solve that, look for such jobs using `kubectl get job -n stackspin-apps` and delete any finished ones using `kubectl delete job` That should let the `kubectl delete pvc ...` command finish; if it was already terminated, run it again.

The `velero restore create ...` command initiates the restore operation, but it doesn’t wait until the operation is complete. You may use the commands suggested in the terminal output to check on the status of the operation. Additionally, once the restore operation is finished, it may take some more time for the various app components to be fully started and for the app to be operational again.

9.3 Change the IP of your cluster

In case your cluster needs to migrate to another IP, make sure to update the IP address in `/etc/rancher/k3s/k3s.yaml` and, if applicable, your local kube config and `inventory.yml` in the cluster directory `clusters/stackspin.example.org`.

9.4 Delete evicted pods

In case your cluster disk is full, kubernetes [taints](#) the node with `DiskPressure`. Then it tries to evict pods, which is pointless in a single node setup but can still happen. We have experienced hundreds of pods in evicted state that still showed up after `DiskPressure` had recovered. See also the [out of resource handling with kubelet](#) documentation.

You can delete all evicted pods with this command:

```
kubectl get pods --all-namespaces -ojson | jq -r '.items[] | select(.status.reason!
↳=null) | select(.status.reason | contains("Evicted")) | .metadata.name + " " + .
↳metadata.namespace' | xargs -n2 -l bash -c 'kubectl delete pods $0 --namespace=$1'
```

9.5 Apply changes to flux variables

Before installing, you configured cluster variables in your cluster directory in `.flux.env`. If you change any of these variables after installation you can apply the changes by following the [Step 2: Install core applications](#) instructions until the step `kubectl apply -k $CLUSTER_DIR`. Then, use the following command that will apply the changes to all installed kustomizations:

```
flux get -A kustomizations --no-header | awk -F' ' '{system("flux reconcile -n " $1 "
↳kustomization " $2)}'
```


UPGRADE GUIDE

10.1 Upgrading to 0.8

Note: 0.8 introduce many breaking changes. We gave our best to make the upgrade smooth but this will require a lot of manual intervention. Please reach out to us for help if needed !

When upgrading to version 0.8 OpenAppStack will be renamed to its final name: *Stackspin*. This comes with many changes, some of which need to be applied manually.

We have written a script to automate a lot of the preparations for the upgrade. However, afterwards you might need to get your hands dirty to get all your applications to work again. **Read this whole upgrade guide carefully, before you get started!**

```
# Log in to your Stackspin server
ssh <server>
# Download our upgrade script
wget https://open.greenhost.net/stackspin/stackspin/-/raw/main/upgrade-scripts/to-0.8.0/
↪rename-to-stackspin.sh
chmod +x rename-to-stackspin.sh
```

First of all, if you have any `-override` configmaps or secrets, you'll want to move them from the `oas` namespace to the `stackspin` namespace, and from `oas-apps` to `stackspin-apps` (you also need to make these namespaces first). You also need to rename them from `oas-X` to `stackspin-X`. You can use a command like this to rename the cm and move it to the right namespace.

```
kubectl get cm -n oas-apps oas-$APP-override -o json | jq '.metadata.name="stackspin-
↪$APP-override"' | jq '.metadata.namespace="stackspin-apps"' | kubectl apply -f -
```

This script will cause serious down time and it will not do everything for you. Rather, it will prepare your cluster for the upgrade.

The script does the following:

1. Install `jq`
2. Shut down the cluster, make a back-up of the data, and bring the cluster back up
3. Copy all relevant `oas-*` secrets to `stackspin-*`
4. Move all `PersistentVolumeClaims` to the `stackspin` and `stackspin-apps` namespaces and sets the `PersistentVolumes ReclaimPolicy` to "Retain" so your data is not accidentally deleted.
5. Delete all OAS `flux` kustomizations

6. Delete the `oas` and `oas-apps` namespace
7. Create the new `stackspin` source and customization

Because there are not many Stackspin users yet, the script can need some manual adjustments. It was written for clusters on which all applications are installed. If you have *not* installed some of the applications, please remove these applications from the script manually.

```
# Execute the upgrade preparation script
./rename-to-stackspin.sh
```

After this, you need to update secrets and Flux in the cluster by running `install/install-stackspin.sh`. Then re-install applications by running `install/install-app.sh <app>` from the Stackspin repository. See the application specific upgrade guides below.

After all your applications work again, you can clean up the old secrets and reset the Persistent Volume ReclaimPolicy to Delete

```
wget https://open.greenhost.net/stackspin/stackspin/-/raw/main/upgrade-scripts/to-0.8.0/
↪cleanup.sh
chmod +x cleanup.sh
./cleanup.sh
```

10.1.1 Nextcloud

Your SSO users will have new usernames, because the OIDC provider has been renamed from `oas` to `stackspin` and because the new SSO system uses UUIDs to uniquely identify users.

You can choose from these options:

1. Manually re-upload and re-share your files after logging in to your new user for the first time.
2. It is possible to transfer files from your previous user to the new user. To do so, find your new username. It is visible in Settings -> Sharing behind “Your Federated Cloud ID” after you’ve logged out and in to Nextcloud with the new SSO (the part *before* the @).

```
# Exec into the Nextcloud container
kubectl exec -n stackspin-apps nc-nextcloud-xxx-xxx -it -- /bin/bash
# Change to the www-data user
su www-data -s /bin/bash
# Repeat this command for each username
php occ files:transfer-ownership oas-<old username> <new user ID>
# Note: the files are transferred to a subfolder in the new user's
# directory
```

Depending on when you first installed Nextcloud, the `setup-apps` job may fail during the upgrade. If that happens, execute these commands in order to update the failing apps to their newest version, and to remove old files that can cause problems.

```
kubectl exec -n stackspin-apps deployment/nc-nextcloud -- rm -r /var/www/html/custom_
↪apps/onlyoffice
kubectl exec -n stackspin-apps deployment/nc-nextcloud -- rm -r /var/www/html/custom_
↪apps/sociallogin
flux suspend hr -n stackspin-apps nextcloud && flux resume hr -n stackspin-apps nextcloud
```

10.1.2 Rocket.Chat

We replaced Rocket.Chat with [Zulip](#) in this release. If you want to migrate your Rocket.Chat data to your new [Zulip](#) installation please refer to [Import from Rocket.Chat](#).

10.1.3 Monitoring

The monitoring stack will work after the upgrade, but monitoring data from the previous version will not be available.

10.1.4 Wekan

In our testing we didn't need to change anything for Wekan to work.

10.1.5 WordPress

In our testing we didn't need to change anything for WordPress to work.

10.2 Upgrading to 0.7.0

Because of [problems with Helm and secret management](#) we had to move away from using a helm chart for application secrets, and now use scripts that run during installation to manage secrets. Because we have removed the `oas-secrets` helm chart, Flux will remove the secrets that it has generated. **It is important that you back up these secrets before switching from v0.6 to v0.7!**

Note: Before you start, please ensure that you have the right `yq` tool installed, because you will need it later. There are two very different versions of `yq`. The one you need is the go based [yq from Mike Farah](#), which installs the same binary name as the `python-yq` one, while both have different command sets. The `yq` needed here can be installed by running `sudo snap install yq`, `brew install yq` or with other methods from the [yq installation instructions](#).

If you're unsure which `yq` you have installed, look at the output of `yq --help` and make sure `eval` shows up under `Available Commands:`.

To back-up your secrets, run the following script:

```
bash
#!/usr/bin/env bash

mkdir secrets-backup

kubectrl get secret -o yaml -n flux-system oas-cluster-variables > secrets-backup/oas-
↪cluster-variables.yaml
kubectrl get secret -o yaml -n flux-system oas-wordpress-variables > secrets-backup/oas-
↪wordpress-variables.yaml
kubectrl get secret -o yaml -n flux-system oas-wekan-variables > secrets-backup/oas-
↪wekan-variables.yaml
kubectrl get secret -o yaml -n flux-system oas-single-sign-on-variables > secrets-backup/
↪oas-single-sign-on-variables.yaml
kubectrl get secret -o yaml -n flux-system oas-rocketchat-variables > secrets-backup/oas-
↪rocketchat-variables.yaml
```

(continues on next page)

(continued from previous page)

```
kubectl get secret -o yaml -n flux-system oas-kube-prometheus-stack-variables > secrets-
↳ backup/oas-kube-prometheus-stack-variables.yaml
kubectl get secret -o yaml -n oas oas-prometheus-basic-auth > secrets-backup/
↳ oas-prometheus-basic-auth.yaml
kubectl get secret -o yaml -n oas oas-alertmanager-basic-auth > secrets-backup/
↳ oas-alertmanager-basic-auth.yaml
kubectl get secret -o yaml -n flux-system oas-oauth-variables > secrets-backup/oas-
↳ oauth-variables.yaml
kubectl get secret -o yaml -n flux-system oas-nextcloud-variables > secrets-backup/oas-
↳ nextcloud-variables.yaml
```

This script assumes you have all applications enabled. You might get an error like:

```
Error from server (NotFound): secrets "oas-wekan-variables" not found
```

This is not a problem, but it *does* mean you need to add an oauth secret for Wekan to the file `secrets-backup/oas-oauth-variables.yaml`. Copy one of the lines under “data:”, rename the field to `wekan_oauth_client_secret` and enter a different random password. Make sure to base64 encode it (`echo "<your random password>" | base64`).

This script creates a directory called `secrets-backup` and places the secrets that have been generated by Helm in it as `yaml` files.

Now you can upgrade your cluster by running `kubectl -n flux-system patch gitrepository openappstack --type merge -p '{"spec":{"ref":{"branch":"v0.7"}}}'` or by editing the `gitrepository` object manually with `kubectl -n flux-system edit gitrepository openappstack` and setting `spec.ref.branch` to `v0.7`.

Flux will now start updating your cluster to version 0.7. This process will fail, because it will remove the secrets that you just backed up. Make sure that the `oas-secrets helmrelease` has been removed by running `flux get hr -A`. You might also see that some `helmreleases` start failing to be installed because important secrets do not exist anymore.

As soon as the `oas-secrets helmrelease` does not exist anymore, you can run the following code:

```
#!/usr/bin/env bash

# Again: make sure you use https://github.com/mikefarah/yq -- install with `snap install`
↳ yq`
yq eval 'del(.metadata.annotations,.metadata.labels,.metadata.creationTimestamp,.
↳ metadata.resourceVersion,.metadata.uid)' secrets-backup/oas-wordpress-variables.yaml |
↳ kubectl apply -f -
yq eval 'del(.metadata.annotations,.metadata.labels,.metadata.creationTimestamp,.
↳ metadata.resourceVersion,.metadata.uid)' secrets-backup/oas-wekan-variables.yaml |
↳ kubectl apply -f -
yq eval 'del(.metadata.annotations,.metadata.labels,.metadata.creationTimestamp,.
↳ metadata.resourceVersion,.metadata.uid)' secrets-backup/oas-single-sign-on-variables.
↳ yaml | kubectl apply -f -
yq eval 'del(.metadata.annotations,.metadata.labels,.metadata.creationTimestamp,.
↳ metadata.resourceVersion,.metadata.uid)' secrets-backup/oas-rocketchat-variables.yaml
↳ | kubectl apply -f -
yq eval 'del(.metadata.annotations,.metadata.labels,.metadata.creationTimestamp,.
↳ metadata.resourceVersion,.metadata.uid)' secrets-backup/oas-kube-prometheus-stack-
↳ variables.yaml | kubectl apply -f -
yq eval 'del(.metadata.annotations,.metadata.labels,.metadata.creationTimestamp,.
↳ metadata.resourceVersion,.metadata.uid)' secrets-backup/oas-prometheus-basic-auth.yaml
↳ | kubectl apply -f -
```

(continues on next page)

(continued from previous page)

```

yq eval 'del(.metadata.annotations,.metadata.labels,.metadata.creationTimestamp,.
↳metadata.resourceVersion,.metadata.uid)' secrets-backup/oas-alertmanager-basic-auth.
↳yaml | kubectl apply -f -
yq eval 'del(.metadata.annotations,.metadata.labels,.metadata.creationTimestamp,.
↳metadata.resourceVersion,.metadata.uid)' secrets-backup/oas-oauth-variables.yaml |
↳kubectl apply -f -
yq eval 'del(.metadata.annotations,.metadata.labels,.metadata.creationTimestamp,.
↳metadata.resourceVersion,.metadata.uid)' secrets-backup/oas-nextcloud-variables.yaml |
↳kubectl apply -f -

```

Again this script assumes you have all applications installed. If you get the following error, you can ignore it:

```

error: error validating "STDIN": error validating data: [apiVersion not set, kind not
↳set]; if you choose to ignore these errors, turn validation off with --validate=false

```

Now Flux should succeed in finishing the update. Some helmreleases or kustomizations might have already failed because the secrets did not exist. Once failed, you can retrigger reconciliation of a kustomization using the commands `flux reconcile kustomization ...` or `flux reconcile helmrelease ...`. This can take quite a while (over an hour some times), because Flux waits for some long timeouts before giving up and re-starting a reconciliation.

10.2.1 Potential upgrade issues

Some errors we've seen during our own upgrade process, and how to solve them:

SSO helm upgrade failed

```

oas          single-sign-on          False Helm upgrade failed: template: single-sign-on/
↳templates/secret-oauth2-clients.yaml:9:55: executing "single-sign-on/templates/secret-
↳oauth2-clients.yaml" at <b64enc>: invalid value; expected string 0.2.2      False

```

This means that the `single-sign-on` helmrelease was created with empty oauth secrets. The secrets will get a value once the core *kustomization* is reconciled: `flux reconcile ks core` should solve the problem.

If that does not solve the problem, you should check if the secret contains a value for all the apps:

```

# kubectl get secret -n flux-system oas-oauth-variables -o yaml
apiVersion: v1
data:
  grafana_oauth_client_secret: <redacted>
  nextcloud_oauth_client_secret: <redacted>
  rocketchat_oauth_client_secret: <redacted>
  userpanel_oauth_client_secret: <redacted>
  wekan_oauth_client_secret: <redacted>
  wordpress_oauth_client_secret: <redacted>
  ...

```

If your secret lacks one of these variables, use `kubectl edit` to add them. You can use any password generator to generate a password for it. Make sure to base64 encode the data before you enter it in the secret.

Loki upgrade retries exhausted

While running `flux get helmrelease -A`, you'll see:

```
oas          loki          False  upgrade retries exhausted  2.5.2  ─
└─False
```

This happens sometimes because Loki takes a long time to upgrade. Usually it is solved by running `flux reconcile hr loki -n oas` again.

10.3 Upgrading to 0.6.0

A few things are important when upgrading to 0.6.0:

- We now use Flux 2 and the installation procedure has been overhauled. For this reason we advice you to set up a completely new cluster.
- Copy your configuration details from `settings.yaml` to a new `.flux.env`. See `install/.flux.env.example` and the *Installation overview* instructions for more information.

Please [reach out to us](#) if you are using, or plan to use OAS in production.

10.4 Upgrading from 0.4.0 to 0.5.0

Unfortunately we can't ensure a smooth upgrade for this version neither. Please read the section below on how to do an upgrade by installing a the new OAS version from scratch after backing up your data.

10.5 Upgrading from 0.3.0 to 0.4.0

There is no easy upgrade path from version 0.3.0 to version 0.4.0. As far as we know, nobody was running OpenAppStack apart from the developers, so we assume this is not a problem.

If you do need to upgrade, this is how you can migrate your data. Backup all the data available under `/var/lib/OpenAppStack/local-storage`, create a new cluster using the installation instructions, and putting back the data. This migration procedure might not work perfectly.

Use `kubectl get pvc -A` on your old cluster to get a mapping of all the PVC uuids (and thus their folder names in `/var/lib/OpenAppStack/local-storage`) to the pods they are bound to.

Then, delete your old OpenAppStack, and install a new one with version number 0.4.0 or higher. You can upload your backed up data into `/var/lib/OpenAppStack/local-storage`. All PVCs will have new unique IDs (and thus different folder names). You have to manually match the folders from your backup with the new folders.

Additionally, if you want to re-use your old `settings.yaml` file, this data needs to be added to it:

```
backup:
  s3:
    # Disabled by default. To enable, change to `true` and configure the
    # settings below. You'll also want to add "velero" to the enabled
    # applications a bit further in this file.
    # Finally, you'll also need to provide access credentials as
    # secrets; see the documentation:
```

(continues on next page)

(continued from previous page)

```

# https://docs.openappstack.net/en/latest/installation_instructions.html#step-2-
↪optional-cluster-backups-using-velero
enabled: false
# URL of S3 service. Please use the principal domain name here, without the
# bucket name.
url: "https://store.greenhost.net"
# Region of S3 service that's used for backups.
# For some on-premise providers this may be irrelevant, but the S3
# apparently requires it at some point.
region: "ceph"
# Name of the S3 bucket that backups will be stored in.
# This has to exist already: Velero will not create it for you.
bucket: "openappstack-backup"
# Prefix that's added to backup filenames.
prefix: "test-instance"

# A whitelist of applications that will be enabled.
enabled_applications:
# System components, necessary for the system to function.
- 'cert-manager'
- 'letsencrypt-production'
- 'letsencrypt-staging'
- 'ingress'
- 'local-path-provisioner'
- 'single-sign-on'
# The backup system Velero is disabled by default, see settings under `backup` above.
# - 'velero'
# Applications.
- 'grafana'
- 'loki'
- 'promtail'
- 'nextcloud'
- 'prometheus'
- 'rocketchat'
- 'wordpress'

```

10.6 Upgrading to 0.3.0

Upgrading from versions earlier than 0.3.0 requires manual intervention.

- Move your local `settings.yml` file to a different location:

```

cd CLUSTER_DIR
mkdir -p ./group_vars/all/
mv settings.yml ./group_vars/all/

```

- Flux is now used to install and update applications. For that reason, we need you to remove all helm charts (WARNING: You will lose your data!):

```

helm delete --purge oas-test-cert-manager oas-test-local-storage \
  oas-test-prometheus oas-test-proxy oas-test-files`

```

- After removing all helm charts, you probably also want to remove all the pvcs that are left behind. Flux will not re-use the database PVCs created for these applications. Find all the pvcs by running `kubect1 get pvc --namespace oas-apps` and `kubect1 get pvc --namespace oas`

CUSTOMIZING

In this guide we show how to customize your cluster installation, i.e. if you want to install additional applications, or change the configuration of existing apps installed by Stackspin this is the right place. Customizing other parts of your cluster is possible but not yet covered by this guide. This guide is written for users with advanced knowledge of the tools behind Stackspin, most importantly: Kubernetes, Helm, Ansible and Flux 2.

Warning: Customizing your Stackspin cluster could break your cluster in a way that it's not easy to recover. Please be aware of the potential risk when proceeding.

11.1 Prerequisites

- A functional Stackspin cluster installed following the [Stackspin installation instructions](#)

11.2 Customize Stackspin applications

Apps deployed by Stackspin are configured using helm values from templates in `flux2/apps/<application>/release.yaml`. It is possible to override values from the helmrelease by adding a custom ConfigMap or Secret to the cluster. The secret or configmap name is specified in the `valuesFrom` section of the `release.yaml` file. Read more in the [Flux documentation](#)

11.2.1 Example: Customize Nextcloud to work with staging certificates

Our CI pipeline works with staging certificates from Let's Encrypt, for that reason we need to allow insecure connections for the integration with ONLYOFFICE. You can find the file at `install/overrides/stackspin-nextcloud-override.yaml`.

To apply it, run the following commands:

```
# If you want to run this on your provisioning machine, tell kubectl to use
# your cluster:
export KUBECONFIG=$PWD/clusters/stackspin.example.org/kube_config_cluster.yml
# Check the current state of the helmrelease you want to modify:
flux get helmrelease -A
# If all is OK, make sure to apply your override configmap or secret in the
# same namespace as your helmrelease with the '-n' argument
kubectl apply \
```

(continues on next page)

(continued from previous page)

```
-n stackspin-apps \  
-f ./install/overrides/stackspin-nextcloud-override.yaml
```

11.3 Adding custom apps to the cluster

Stackspin uses Flux 2 to install and auto-update applications. If you want to install extra applications or other things into the Kubernetes cluster, our advice would be to set up your own GitRepository and add it to the Flux system.

When you do this, you are fully responsible for keeping those applications secure and updated. If any of those applications is insecure, that can also invalidate the security of your Stackspin applications, because they are part of the same cluster and VPS.

Refer to the [Flux 2 documentation](#) for more information.

TROUBLESHOOTING

If you run into problems, there are a few things you can do to research the problem. This document describes what you can do.

Note: `cluster$` indicates that the commands should be run as root on your Stackspin machine. All other commands need to be run on your *provisioning machine*.

We would love to hear from you! If you have problems, please create an issue in our [issue tracker](#) or reach out as described on our [contact page](#). We want to be in communication with our users, and we want to help you if you run into problems.

12.1 Known issues

If you run into a problem, please check our [issue tracker](#) to see if others have run into the same problem. We might have suggested a workaround or temporary solution in one of our issues. If your problem is not described in an issue, please open a new one so we can solve the problems you encounter.

12.2 Run the CLI tests

To get an overall status of your cluster, run the tests from the command line. There are two types of tests: `testinfra` tests verify the setup, `Taiko` tests check the webservice availability.

12.2.1 Testinfra tests

Testinfra tests are split into two groups, let's call them *blackbox* and *clearbox* tests. The blackbox tests run on your provisioning machine and test the Stackspin cluster from the outside. For example, the certificate check will check if the Stackspin returns valid certificates for the provided services. The clearbox tests run on the Stackspin host and check i.e. if docker is installed in the right version etc. Our testinfra tests are a combination of blackbox and clearbox tests.

First, enter the `test` directory in the Git repository **on your provisioning machine**.

```
cd test
```

To run the test against your cluster, first export the `CLUSTER_DIR` environment variable with the location of your cluster config directory (replace `stackspin.example.org` with your cluster name): Then export the `KUBECONFIG` variable that lets `kubectl` talk to your cluster. Thirdly, export a variable that contains the password for the Prometheus HTTP endpoint.

```
export CLUSTER_DIR="${PWD}/../clusters/stackspin.example.org"
export KUBECONFIG=${CLUSTER_DIR}/kube_config_cluster.yml
export BASIC_AUTH_PW=$(kubectl get secret -n stackspin stackspin-prometheus-basic-auth --
↳template '{{ .data.pass }}' | base64 -d)
```

Run all tests

```
py.test -s --ansible-inventory=${CLUSTER_DIR}/inventory.yml --hosts='ansible://*'
```

Test all applications

This will check for:

- The applications return proper certificates
- All helm releases are successfully installed
- All app pods are running and healthy (this test includes all optional applications)

These tests includes all optional applications and will fail for optional applications that are not installed.

```
pytest -s -m 'app' --connection=ansible --ansible-inventory=${CLUSTER_DIR}/inventory.yml
↳--hosts='ansible://*'
```

Tests a specific application

```
pytest -s -m 'app' --app="wordpress" --connection=ansible --ansible-inventory=${CLUSTER_
↳DIR}/inventory.yml --hosts='ansible://*'
```

Known Issues

The default ssh backend for testinfra tests is *paramiko*, which does not work out of the box. It fails to connect to the host because the *ed25519* hostkey is not verified. Therefore we need to force plain *ssh://* with either `connection=ssh` or `--hosts=ssh://...`

12.2.2 Taiko tests

Taiko tests run in a browser to test if all the interfaces are up and running and correctly connected to each other. They are integrated in the `stackspin` CLI command suite.

Prerequisites

Install **Taiko** on your provisioning machine:

```
npm install -g taiko
```

Run Taiko tests

To run all Taiko tests, run the following command in this repository:

```
python -m stackspin CLUSTERNAME test
```

To learn more about the `test` subcommand, run:

```
python -m stackspin CLUSTERNAME test --help
```

You can also only run a Taiko test for a specific application, i.e.:

```
python -m stackspin CLUSTERNAME test --taiko-tags nextcloud
```

12.2.3 Advanced usage

Testinfra tests

Specify host manually:

```
py.test -s --hosts='ssh://root@example.stackspin.net'
```

Run only tests tagged with *prometheus*:

```
py.test -s --ansible-inventory=${CLUSTER_DIR}/inventory.yml --hosts='ansible://*' -m
↳prometheus
```

Run cert test manually using the ansible inventory file:

```
py.test -s --ansible-inventory=${CLUSTER_DIR}/inventory.yml --hosts='ansible://*' -m
↳certs
```

Run cert test manually against a different cluster, not configured in any ansible inventory file, either by using `pytest`:

```
FQDN='example.stackspin.net' py.test -sv -m 'certs'
```

or directly:

```
FQDN='example.stackspin.net' pytest/test_certs.py
```

Running Testinfra tests with local gitlab-runner docker executor

Export the following environment variables like this:

```
export CI_REGISTRY_IMAGE='open.greenhost.net:4567/stackspin/stackspin'
export SSH_PRIVATE_KEY="$(cat ~/.ssh/id_ed25519_stackspin_ci)"
export COSMOS_API_TOKEN='...'
```

then:

```
gitlab-runner exec docker --env CI_REGISTRY_IMAGE="$CI_REGISTRY_IMAGE" --env SSH_PRIVATE_
↪KEY="$SSH_PRIVATE_KEY" --env COSMOS_API_TOKEN="$COSMOS_API_TOKEN" bootstrap
```

Advanced Taiko tests

If you want to use Taiko without invoking the stackspin CLI, go to the `test/taiko` directory and run:

For nextcloud & onlyoffice tests:

```
export DOMAIN='stackspin.example.net'
export SSO_USERNAME='user1'
export SSO_USER_PW='...'
export TAIKO_TESTS='nextcloud'
taiko --observe taiko-tests.js
```

You can replace `nextcloud` with `grafana` or `wordpress` to test the other applications, or with `all` to test all applications.

12.3 SSH access

You can SSH login to your VPS. Some programs that are available to the root user on the VPS:

- `kubect1`, the Kubernetes control program. The root user is connected to the cluster automatically.
- `helm` is the “Kubernetes package manager”. Use i.e. `helm ls --all-namespaces` to see what apps are installed in your cluster. You can also use it to perform manual upgrades; see `helm --help`.
- `flux` is the `flux` command line tool

12.4 Using kubect1 to debug your cluster

You can use `kubect1`, the Kubernetes control program, to find and manipulate your Kubernetes cluster. Once you have installed `kubect1`, to get access to your cluster with the Stackspin CLI:

```
$ python -m stackspin stackspin.example.org info
```

Look for these lines:

```
To use kubect1 with this cluster, copy-paste this in your terminal:
export KUBECONFIG=/home/you/projects/stackspin/clusters/stackspin.example.org/kube_
↪config_cluster.yml
```

Copy the whole export line into your terminal. **In the same terminal window**, `kubect1` will from now on connect to your cluster.

12.5 HTTPS certificates

Stackspin uses `cert-manager` to automatically fetch [Let's Encrypt](#) certificates for all deployed services. If you experience invalid SSL certificates, i.e. your browser warns you when visiting Zulip (<https://zulip.stackspin.example.org>), a useful resource for troubleshooting is the official `cert-manager` [Troubleshooting Issuing ACME Certificates](#) documentation. First, try this:

In this example we fix a failed certificate request for `https://chat.stackspin.example.org`. We will start by checking if `cert-manager` is set up correctly.

12.5.1 Is your cluster using the live ACME server?

```
$ kubectl get clusterissuers -o yaml | grep 'server:'
```

Should return `server: https://acme-v02.api.letsencrypt.org/directory` and not something with the word `staging` in it.

12.5.2 Are all cert-manager pods in the `stackspin` namespace in the `READY` state?

```
$ kubectl -n cert-manager get pods
```

`Cert-manager` uses a “custom resource” to keep track of your certificates, so you can also check the status of your certificates by running:

This returns all the certificates for all applications on your system. The command includes example output of healthy certificates.

```
$ kubectl get certificates -A
NAMESPACE   NAME                                     READY   SECRET                               AGE
stackspin   hydra-public.tls                       True    hydra-public.tls                    14d
stackspin   single-sign-on-userpanel.tls          True    single-sign-on-userpanel.tls        14d
stackspin-apps  stackspin-nextcloud-files             True    stackspin-nextcloud-               14d
stackspin-apps  stackspin-nextcloud-office            True    stackspin-nextcloud-               14d
stackspin   grafana-tls                            True    grafana-tls                         13d
stackspin   alertmanager-tls                      True    alertmanager-tls                   13d
stackspin   prometheus-tls                        True    prometheus-tls                     13d
```

If there are problems, you can check for the specific `certificaterequests`:

```
$ kubectl get certificaterquests -A
```

For even more information, inspect the logs of the `cert-manager` pod:

```
$ kubectl -n stackspin logs -l "app.kubernetes.io/name=cert-manager"
```

You can `grep` for your cluster domain or for any specific subdomain to narrow down results.

12.5.3 Example

Query for failed certificates, -requests, challenges or orders:

```
$ kubectl get --all-namespaces certificate,certificaterequest,challenge,order | grep -iE
↪ '(false|pending)'
```

stackspin-apps	certificate.cert-manager.io/stackspin-zulip	False	↪
↪	stackspin-zulip	15h	
stackspin-apps	certificaterequest.cert-manager.io/stackspin-zulip-2045852889		↪
↪	False	15h	
stackspin-apps	challenge.acme.cert-manager.io/stackspin-zulip-2045852889-1775447563-837515681	pending	↪
↪	chat.stackspin.example.org	15h	
stackspin-apps	order.acme.cert-manager.io/stackspin-zulip-2045852889-1775447563		↪
↪	pending	15h	

We see that the zulip certificate resources have been in a bad state for 15 hours.

Show certificate resource status message:

```
$ kubectl -n stackspin-apps get certificate stackspin-zulip -o jsonpath="{.status.
↪ conditions[*]['message']}"
Waiting for CertificateRequest "stackspin-zulip-2045852889" to complete
```

We see that the *certificate* is waiting for the *certificaterequest*, let's query its status message:

```
$ kubectl -n stackspin-apps get certificaterequest stackspin-zulip-2045852889 -o
↪ jsonpath="{.status.conditions[*]['message']}"
Waiting on certificate issuance from order stackspin-apps/stackspin-zulip-2045852889-
↪ 1775447563: "pending"
```

Show the related order resource and look at the status and events:

```
$ kubectl -n stackspin-apps describe order stackspin-zulip-2045852889-1775447563
```

Show the failed challenge resource reason:

```
$ kubectl -n stackspin-apps get challenge stackspin-zulip-2045852889-1775447563-
↪ 837515681 -o jsonpath='{.status.reason}'
Waiting for http-01 challenge propagation: wrong status code '503', expected '200'
```

In this example, deleting the challenge fixed the issue and a proper certificate could get fetched:

```
$ kubectl -n stackspin-apps delete challenges.acme.cert-manager.io stackspin-zulip-
↪ 2045852889-1775447563-837515681
```

12.6 Application installation or upgrade failures

Application installations and upgrades are managed by `flux`. Flux uses `helm-controller` to install and upgrade applications with `helm charts`.

An application installed with Flux consists of a *kustomization*. This resource defines where the information about the application is stored in our Git repository. The *kustomization* contains a *helmrelease*, which is an object that represents an installation of a Helm chart. Read more about the difference between *kustomizations* and *helmreleases* in the [flux documentation](#).

To find out if all *kustomizations* have been applied correctly, run the following flux command in your cluster or from the provisioning machine:

```
cluster$ flux get kustomizations -A
```

If there is an issue, use `kubectl` to inspect the respective service, for example `nginx`:

```
kubectl describe helmrelease -n stackspin nginx | less
```

If all your *kustomizations* are in a *Ready* state, take a look at your *helmreleases*:

```
cluster$ flux get helmreleases -A
```

Often, you can resolve complications with *kustomizations* or *helmreleases* by telling Flux to *reconcile* them:

```
cluster$ flux reconcile helmrelease nextcloud
```

This brings the Nextcloud *helmrelease* into a state that our Stackspin wants it to be in.

12.6.1 Common installation failures

var substitution failed

When you execute `flux get kustomization` and you see this error:

```
var substitution failed for 'kube-prometheus-stack': YAMLToJSON: yaml: line 32: found ↵
↵ character that cannot start any token
```

That can mean that one of your values contains a double quote (") or that you quoted a value in `.flux.env` during the *Step 1: Flux configuration*. Make sure that `.flux.env` does not contain any values that are quoted.

If you need to change `.flux.env`, run the following commands:

```
$ kubectl apply -k $CLUSTER_DIR
```

Afterwards, you can speed up the process that fixes your *kustomization*, by running `flux reconcile kustomization kube-prometheus-stack` (replace `kube-prometheus-stack` with the *kustomization* mentioned in the error message).

12.7 Purge Stackspin and install from scratch

Warning:

You will lose all your data! This completely destroys Stackspin and takes everything offline. If you chose to do this, you will need to re-install Stackspin and make sure that your data is stored somewhere other than the VPS that runs Stackspin.

If things ever fail beyond possible recovery, here is how to completely purge a Stackspin installation to start from scratch:

```
cluster$ /usr/local/bin/k3s-killall.sh
cluster$ systemctl disable k3s
cluster$ rm -rf /var/lib/{rancher,Stackspin,kubelet} /etc/rancher /var/log/{Stackspin,
↳ containers,pods} /tmp/k3s /etc/systemd/system/k3s.service
cluster$ systemctl reboot
```

13.1 Access control

By default, the applications on your Stackspin cluster will be exposed to the whole internet (although they are password protected).

If you like to limit who can access your cluster resources you can configure the Stackspin ingress (`ingress-nginx`) to only accept connections from a certain IP address or range.

Follow the *Customize Stackspin applications* instructions, but use the following secret as `install/overrides/stackspin-nginx-override.yml` and apply the secret in the `stackspin` namespace instead of `stackspin-apps`. Replace the source range with the IP address ranges you want to allow.

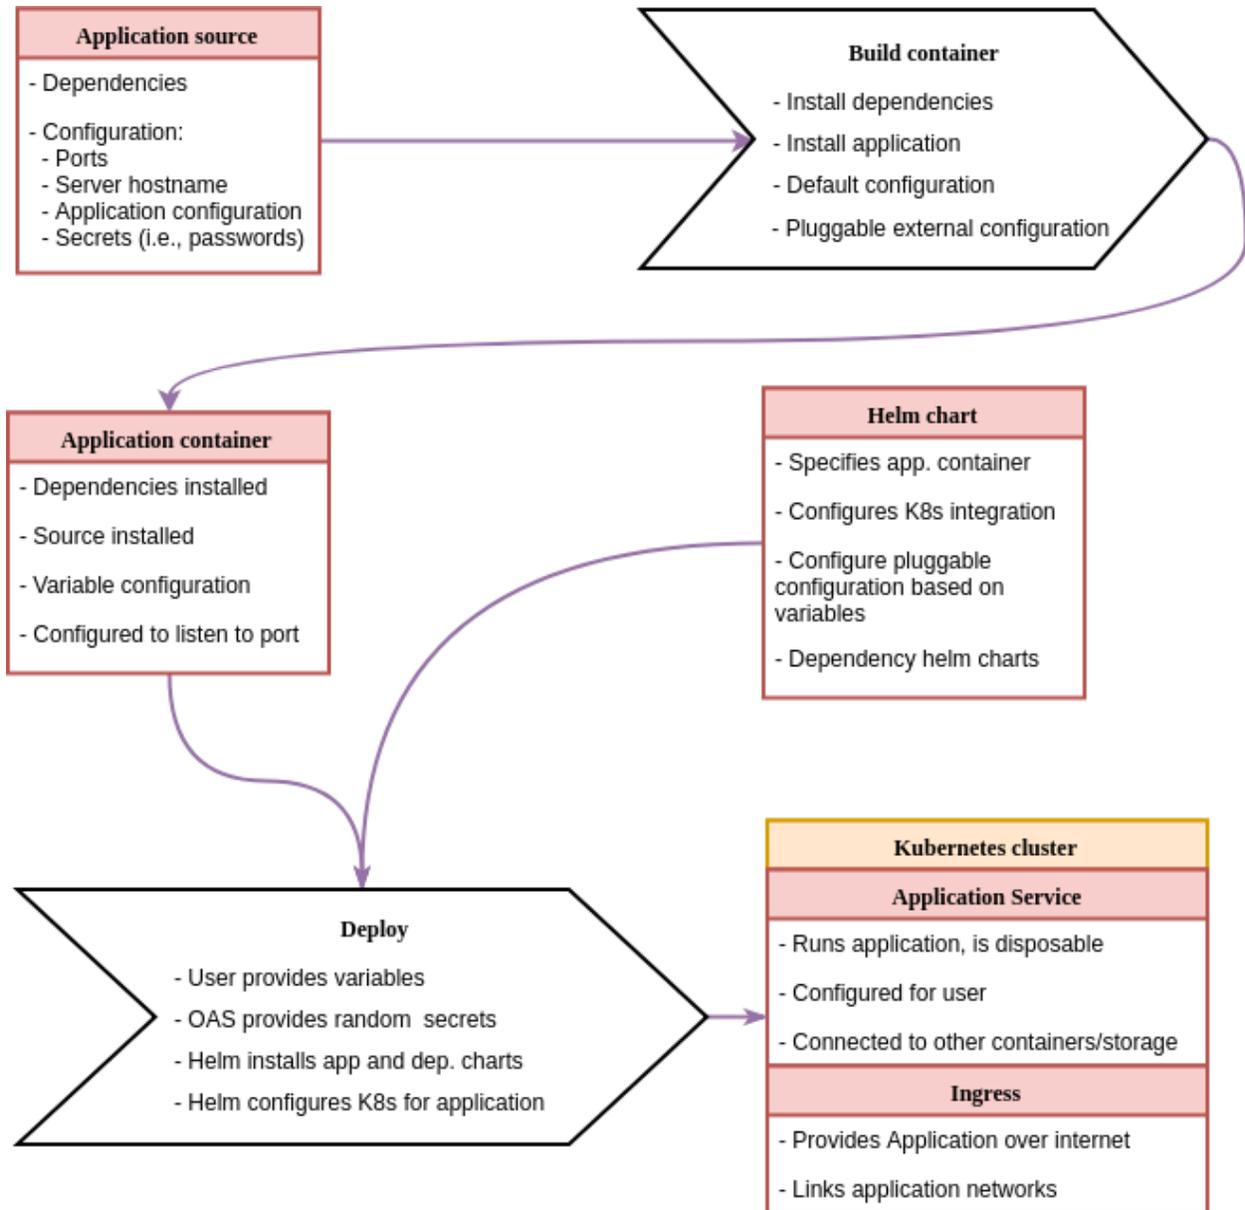
```
---
apiVersion: v1
kind: secret
metadata:
  name: stackspin-nginx-override
data:
  values.yaml: |
    controller:
      config:
        # https://kubernetes.github.io/ingress-nginx/user-guide/nginx-configuration/
↔ annotations/#whitelist-source-range
        # comma separated list of CIDRs, e.g. 10.0.0.0/24,172.10.0.1.
        whitelist-source-range: 1.2.3.4/24
```


STACKSPIN DESIGN

This article covers the basic design of Stackspin.

14.1 Application build pipeline

The following diagram explains the process to go from an application's source code to a deployment on Stackspin.



These are the steps in more detail:

- Build container (this process should be maintained by application developer by providing a Dockerfile with the application)
 1. Get application package (source code, installation package, etc.)
 1. If not part of the package: get default configuration for the application
 2. Build container with application package installed
 1. Install application dependencies
 2. Install application package
 3. Setup default configuration
 4. Setup pluggable configuration override, can be:
 - Reading environment variables

- Extra configuration file mounted into the container elsewhere
- Helm chart
 - Deployment configuration to specify:
 - * The container(s) that should be deployed.
 - * The port(s) that they expose.
 - * Volume mounts for configuration files and secrets.
 - * Live/readiness probes
 - * Persistent storage locations and methods
 - * A lot of other things
 - Service configuration to specify:
 - * Ports exposed to the user of the application
 - Ingress configuration to specify:
 - * How to proxy to the application (which hostname or URL)
 - * Some authentication plugins (http auth, for example)
 - Custom files:
 - * Add file templates for mountable application configuration files
 - * Files that specify integrations with other services
- Deploy
 1. Create `values.yaml` file with the variables for the Helm deployment to the Kubernetes cluster
 2. “Manually” add secrets to the Kubernetes cluster.
 3. Run `helm install` to install the customised application.

14.2 Configuration

As can be seen in the images, applications are expected to have two different types of configuration. Containers should provide a default configuration, that at least configures things like the port the application runs on, the locations for log files, etc.

What we call the *external configuration* is provided by the user. This includes overrides of the default application, as well as variables like the hostname that the application will run on and listen to and the title of the web interface.

Stackspin will use Helm charts to provide the external configuration for the “Deploy” step. Helm charts can contain configuration file templates with default values that can be overridden during the installation or upgrade of a helm chart.

14.3 Application containers

For inclusion in Stackspin, it is required that the application developers provide Docker containers for their applications. There are several reasons for this:

- If application developers do not provide a container, chances are they also do not think about how their application would update itself after a new container is deployed. This can lead to problems with things like database migrations.
- Maintaining the containerisation for an application can, in most cases, not be fully automated.

14.3.1 Container updates

When an application update is available, these updates need to be rolled out to Stackspin instances. This will be done according to the following steps:

1. Application container is built with new application source and tagged for testing.
2. Helm chart for application is updated to provide new container.
3. Helm chart is deployed to an Stackspin test cluster following the steps in the diagram above.
4. Application is tested with automated tests
5. If tests succeed, new container is tagged for release.
6. Stackspin automated update job fetches new Helm chart and upgrades current instance using Helm.

Most of these steps can be developed by configuring a CI system and configuring Kubernetes and Helm correctly. The automated update job that will run on Stackspin clusters will be developed by us.

14.4 Persistent data

Containerised applications are normally “stateless” (meaning no data is saved inside the containers). However, it is possible to mount persistent volumes to specific directories in the container, basically adding a persistent layer on top of the containerised application. To provide this in Stackspin’s simple setup, we use a [local storage provisioner](#) that automatically provides persistent data on the VPS running Stackspin to an application that requests it.

14.5 Automatic updates

Stackspin has an auto-update mechanism that performs unattended upgrades to applications. [Flux 2](#) is the system running in the cluster that is responsible for these updates.

Flux 2 tracks all the files in the `flux2` directory of the [Stackspin code repository](#). Once changes are pushed to the branch that Flux tracks, the changes are applied to the cluster.

We use Flux 2 in “read only” mode, which means that your Stackspin cluster does not push changes to our Git repository. You can read more about Flux 2 and its components in the [flux 2 documentation](#).

14.6 APIs

Stackspin includes several APIs that are available for some pieces of the software to talk to. Specifically the [Dashboard application](#) uses most of these APIs.

- [Ory Kratos](#) is the user management API that we use. It comes with several handy features, like sending password reset emails and (soon) MFA.
- [Ory Hydra](#) provides us with an OpenID Connect workflow.
- [Velero](#) is a backup mechanism that integrates well with Kubernetes: it includes “Custom Resource Definitions” for backups and backup locations. As such, we can access it through the Kubernetes API.
- [Flux 2](#) is used to configure and auto-update applications. App configuration happens mostly through Kubernetes [ConfigMaps](#) and [Secrets](#) and are applied through the Flux CRDs [Kustomization](#) and [HelmRelease](#).

INSTRUCTIONS

This document contains instructions for a few tasks that you might have to do as a Stackspin developer.

15.1 Building a new version of Zulip

As long as Zulip version 5 is not released yet, you need to build it from the source code's main branch every once in a while. In order to do so, you need the following:

1. The commit ref of the version of the Zulip repository you want to build, e.g. `5e90b32f40ad45d59e00a4e5c08cc6f180e04d30`
2. An abbreviated version of the commit ref. We take the first ten characters of the whole ref, `5e90b32f40` in our example.
3. A local copy of our fork of the `docker-zulip` repository (at least until our pull request is merged).

This is necessary because we have added the `SOCIAL_AUTH_OIDC_ENABLED_IDPS` variable to the `entrypoint.sh` docker entrypoint script. Without our changes, setting OIDC settings through environment variables is impossible.

Once you have these two things, you can run the following build command to build the image:

```
docker build \  
  --build-arg "ZULIP_GIT_REF=5e90b32f40ad45d59e00a4e5c08cc6f180e04d30" \  
  -t open.greenhost.net:4567/stackspin/stackspin/zulip:5e90b32f40 \  
  .
```

The build argument `ZULIP_GIT_REF` tells the Dockerfile to use that ref when checking out and building Zulip. The tag is necessary so you can push the new Zulip image to our registry after building with the following command:

```
docker push open.greenhost.net:4567/stackspin/stackspin/zulip:5e90b32f40
```

Once the push is completed, you can edit the tag in `flux2/apps/zulip/zulip-values-configmap.yaml` and create a PR to update Zulip in Stackspin.

Once our PR has been merged **and** Zulip version 5 has been released, this guide is obsolete and Stackspin can use docker images as provided by Zulip on Docker hub.

REFERENCE DOCUMENTATION

16.1 File structure

During the installation process, the following files and directories are created:

- `/var/lib/Stackspin/local-storage`: all application data (e.g., Nextcloud files) are stored here.
- `/var/lib/rancher/k3s`: this contains the files related to your “Kubernetes” cluster.
 - The `kubectl` configuration file is located at `/etc/rancher/k3s/k3s.yaml`

COMPARABLE PROJECTS

Other open source projects similar to Stackspin exist. Each of the platforms listed here, like Stackspin, provide a suite of open source cloud applications. Most of the platforms, like Stackspin, include their own user management dashboard and single sign-on (SSO) features.

17.1 What makes Stackspin unique

However there are changes in implementation that make Stackspin different from these alternatives. As far as we found, none of the projects listed below will automatically update your applications without you having to push a button. Also, all mentioned platforms require application modifications, for example to make use of the authentication system.

Stackspin avoids changing the included applications to prevent maintenance overhead and enable seamless native updates. It only requires OpenID Connect support, so the application can be integrated with single sign-on.

So far, Stackspin requires one intervention before an update can be pushed to users: It needs to be tested. We want to make sure that an application works well with the platform and other applications before we let you use it. We work towards a fully automated test suite to minimize failure while providing automatic updates.

17.2 Platforms

17.2.1 Fully Open

Sandstorm allows applications to be installed as so-called Grains. Each grain is a copy of the complete application made for a specific purpose. For example, a grain for document editor “Etherpad” contains not only the data written in the notepad, but also the Etherpad and database software. If you have two notepads, you also have two copies of the software. With many users, this approach runs into limits.

YunoHost is mostly based on Debian and its package management and user management system. As a result, Yunohost is relatively lightweight. However, if an application on your Yunohost contains a security hole, the data of other applications is likely compromised as well. This is less likely in Stackspin as applications run in sealed-off containers.

SelfPrivacy automatically sets up an open stack focused on personal use from your phone on an established cloud provider. From the documentation it seems to not add any integration (such as SSO) between the services and the only dashboard is the app itself.

Mail-in-a-Box is a narrowly focused project that tightly integrates Nextcloud and various applications to create a full-featured mailserver with a control panel. As such there is no SSO and no configuration options.

17.2.2 Open Core

Cloudrion offers a similar application suite. In contrast to Stackspin, Cloudrion *requires a paid account* if you want to use more than two applications or more than five users.

Univention Corporate Server is focused on business and does not make reproducible builds easily available or facilitate community contribution.