



Deployment Guide

SUSE Cloud Application Platform 1.1



Deployment Guide

SUSE Cloud Application Platform 1.1

by Carla Schroder

Publication Date: June 21, 2018

SUSE LLC

10 Canal Park Drive

Suite 200


Cambridge MA 02141

USA

<https://www.suse.com/documentation> 

Copyright © 2006– 2018 SUSE LLC and contributors. All rights reserved.

Permission is granted to copy, distribute and/or modify this document under the terms of the GNU Free Documentation License, Version 1.2 or (at your option) version 1.3; with the Invariant Section being this copyright notice and license. A copy of the license version 1.2 is included in the section entitled “GNU Free Documentation License”.

For SUSE trademarks, see <http://www.suse.com/company/legal/> . All other third-party trademarks are the property of their respective owners. Trademark symbols (®, ™ etc.) denote trademarks of SUSE and its affiliates. Asterisks (*) denote third-party trademarks.

All information found in this book has been compiled with utmost attention to detail. However, this does not guarantee complete accuracy. Neither SUSE LLC, its affiliates, the authors nor the translators shall be held liable for possible errors or the consequences thereof.

Contents

About This Guide vi

1 About SUSE Cloud Application Platform 1

- 1.1 SUSE Cloud Application Platform Architecture 2
- 1.2 New in Version 1.1 4

2 Production Installation with Optional High Availability 6

- 2.1 Prerequisites 6
- 2.2 Choose Storage Class 11
- 2.3 Test Storage Class 11
- 2.4 Configuring the SUSE Cloud Foundry Production Deployment 12
- 2.5 Deploy with Helm 13
- 2.6 Install the Kubernetes charts repository 14
- 2.7 Create Namespaces 14
- 2.8 Copy SUSE Enterprise Storage Secret 15
- 2.9 Deploy UAA 15
- 2.10 Deploy SUSE Cloud Foundry 15
- 2.11 Deploying and Managing Applications with the Cloud Foundry Client 16
- 2.12 Installing the Stratos Web Console 19
- 2.13 Upgrading SUSE Cloud Foundry, UAA, and Stratos 22
- 2.14 Example High Availability Configuration 23
 - Upgrading a non-High Availability Deployment to High Availability 25

3	Setting up and Using a Service Broker Sidecar	27
3.1	Prerequisites	27
3.2	Configuring the MySQL Deployment	27
3.3	Deploying the MySQL Chart	29
3.4	Create and Bind a MySQL Service	30
3.5	Deploying the PostgreSQL Chart	30
3.6	Removing Service Broker Sidecar Deployments	31
4	Backup and Restore	33
4.1	Installing the cf-plugin-backup	33
4.2	Using cf-plugin-backup	34
5	Preparing Microsoft Azure for SUSE Cloud Application Platform	36
5.1	Prerequisites	36
5.2	Create Resource Group and AKS Instance	37
5.3	Enable Swap Accounting	39
5.4	Create a Basic Load Balancer and Public IP Address	40
5.5	Configure Load Balancing and Network Security Rules	41
5.6	Example SUSE Cloud Application Platform Configuration File	43
6	Installing SUSE Cloud Application Platform on OpenStack	45
6.1	Prerequisites	45
6.2	Create a New OpenStack Project	45
6.3	Deploy SUSE Cloud Application Platform	47
6.4	Bootstrap SUSE Cloud Application Platform	47
6.5	Growing the Root Filesystem	48

7 Running SUSE Cloud Application Platform on non-SUSE CaaS Platform Kubernetes Systems 49

8 Minimal Installation for Testing 50

8.1 Prerequisites 51

8.2 Create hostpath Storage Class 57

8.3 Test Storage Class 59

8.4 Configuring the Minimal Test Deployment 59

8.5 Deploy with Helm 60

Install the Kubernetes charts repository 60 • Create Namespaces 61 • Copy SUSE Enterprise Storage Secret 61 • Install UAA 62 • Install SUSE Cloud Foundry 62

8.6 Install the Stratos Console 63

8.7 Updating SUSE Cloud Foundry, UAA, and Stratos 63

9 Troubleshooting 64

9.1 Using Supportconfig 64

9.2 Deployment is Taking Too Long 64

9.3 Deleting and Rebuilding a Deployment 65

9.4 Querying with Kubectl 65

A GNU Licenses 67

A.1 GNU Free Documentation License 67

About This Guide

SUSE Cloud Application Platform is a software platform for cloud-native application development, based on Cloud Foundry, with additional supporting services and components. The core of the platform is SUSE Cloud Foundry, a Cloud Foundry distribution for Kubernetes which runs on SUSE Linux Enterprise containers.

The Cloud Foundry code base provides the basic functionality. SUSE Cloud Foundry differentiates itself from other Cloud Foundry distributions by running in Linux containers managed by Kubernetes, rather than virtual machines managed with BOSH, for greater fault tolerance and lower memory use.

SUSE Cloud Foundry is designed to run on any Kubernetes cluster. This guide describes how to deploy it on [SUSE Container as a Service \(CaaS\) Platform 2.0](https://www.suse.com/documentation/suse-caasp-2/book_caasp_deployment/data/book_caasp_deployment.html) (https://www.suse.com/documentation/suse-caasp-2/book_caasp_deployment/data/book_caasp_deployment.html) ↗.

1 Required Background

To keep the scope of these guidelines manageable, certain technical assumptions have been made:

- You have some computer experience and are familiar with common technical terms.
- You are familiar with the documentation for your system and the network on which it runs.
- You have a basic understanding of Linux systems.

2 Available Documentation

We provide HTML and PDF versions of our books in different languages. Documentation for our products is available at <http://www.suse.com/documentation/> ↗, where you can also find the latest updates and browse or download the documentation in various formats.

The following documentation is available for this product:

Deployment Guide

The SUSE Cloud Application Platform deployment guide gives you details about installation and configuration of SUSE Cloud Application Platform along with a description of architecture and minimum system requirements.

3 Feedback

Several feedback channels are available:

Bugs and Enhancement Requests

For services and support options available for your product, refer to <http://www.suse.com/support/>.

To report bugs for a product component, go to <https://scc.suse.com/support/requests>, log in, and click *Create New*.

User Comments

We want to hear your comments about and suggestions for this manual and the other documentation included with this product. Use the User Comments feature at the bottom of each page in the online documentation or go to <http://www.suse.com/documentation/feedback.html> and enter your comments there.





Mail

For feedback on the documentation of this product, you can also send a mail to doc-team@suse.com. Make sure to include the document title, the product version and the publication date of the documentation. To report errors or suggest enhancements, provide a concise description of the problem and refer to the respective section number and page (or URL).

4 Documentation Conventions

The following notices and typographical conventions are used in this documentation:

- /etc/passwd: directory names and file names
- PLACEHOLDER: replace PLACEHOLDER with the actual value
- PATH: the environment variable PATH
- ls, --help: commands, options, and parameters
- user: users or groups
- package name: name of a package
- Alt, Alt-F1: a key to press or a key combination; keys are shown in uppercase as on a keyboard

- *File, File > Save As*: menu items, buttons
-  This paragraph is only relevant for the AMD64/Intel 64 architecture. The arrows mark the beginning and the end of the text block. 
-  This paragraph is only relevant for the architectures z Systems and POWER. The arrows mark the beginning and the end of the text block. 
- *Dancing Penguins* (Chapter *Penguins*, ↑Another Manual): This is a reference to a chapter in another manual.
- Commands that must be run with root privileges. Often you can also prefix these commands with the sudo command to run them as non-privileged user.

```
root # command
tux > sudo command
```

- Commands that can be run by non-privileged users.

```
tux > command
```

- Notices



Warning: Warning Notice

Vital information you must be aware of before proceeding. Warns you about security issues, potential loss of data, damage to hardware, or physical hazards.



Important: Important Notice

Important information you should be aware of before proceeding.



Note: Note Notice

Additional information, for example about differences in software versions.



Tip: Tip Notice

Helpful information, like a guideline or a piece of practical advice.

5 About the Making of This Documentation

This documentation is written in SUSEDoc, a subset of DocBook 5 (<http://www.docbook.org> [↗]). The XML source files were validated by **jing** (see <https://code.google.com/p/jing-trang/> [↗]), processed by **xsltproc**, and converted into XSL-FO using a customized version of Norman Walsh's stylesheets. The final PDF is formatted through FOP from Apache Software Foundation (<https://xmlgraphics.apache.org/fop>) [↗]. The open source tools and the environment used to build this documentation are provided by the DocBook Authoring and Publishing Suite (DAPS). The project's home page can be found at <https://github.com/openSUSE/daps> [↗].

The XML source code of this documentation can be found at <https://github.com/SUSE/doc-cap> [↗].

1 About SUSE Cloud Application Platform

SUSE Cloud Application Platform is a software platform for cloud-native application deployment based on SUSE Cloud Foundry and SUSE CaaS Platform 2.0. It serves different but complementary purposes for operators and application developers.

For operators, the platform is:

- Easy to install, manage, and maintain
- Secure by design
- Fault tolerant and self-healing
- Offers high availability for critical components
- Uses industry-standard components
- Avoids single vendor lock-in

For developers, the platform:

- Allocates computing resources on demand via API or Web interface
- Offers users a choice of language and Web framework
- Gives access to databases and other data services
- Emits and aggregates application log streams
- Tracks resource usage for users and groups
- Makes the software development workflow more efficient

The principle interface and API for deploying applications to SUSE Cloud Application Platform is SUSE Cloud Foundry. Most Cloud Foundry distributions run on virtual machines managed by BOSH. SUSE Cloud Foundry runs in SUSE Linux Enterprise containers managed by Kubernetes. Containerizing the components of the platform itself has these advantages:

- Improves fault tolerance. Kubernetes monitors the health of all containers, and automatically restarts faulty containers faster than virtual machines can be restarted or replaced.
- Reduces physical memory overhead. SUSE Cloud Foundry components deployed in containers consume substantially less memory, as host-level operations are shared between containers by Kubernetes.

SUSE Cloud Foundry packages upstream Cloud Foundry BOSH releases to produce containers and configurations which are deployed to Kubernetes clusters using Helm.

1.1 SUSE Cloud Application Platform Architecture

This guide details the steps for deploying SUSE Cloud Foundry on SUSE CaaS Platform 2. CaaS Platform is a specialized application development and hosting platform built on the SUSE MicroOS container host operating system, container orchestration with Kubernetes, and Salt for automating installation and configuration.

Important: Review the SUSE CaaS Platform Deployment Guide

Setting up SUSE Cloud Foundry correctly depends on setting up SUSE CaaS Platform correctly. Review the [SUSE CaaS Platform Deployment Guide \(https://www.suse.com/documentation/suse-caasp-2/\)](https://www.suse.com/documentation/suse-caasp-2/) to understand how it operates, and configuration and administration options. You should understand basic Linux, Docker, and Kubernetes administration and use.

A supported deployment includes SUSE Cloud Foundry installed on CaaS Platform. You also need a storage backend, such as SUSE Enterprise Storage, a DNS/DHCP server, and an Internet connection to download additional packages during installation and ~10GB of Docker images on each Kubernetes worker after installation.

A production deployment requires considerable resources. SUSE Cloud Application Platform includes an entitlement of SUSE CaaS Platform 2 and SUSE Enterprise Storage 5. SUSE Enterprise Storage alone has substantial requirements; see the [Tech Specs \(https://www.suse.com/products/suse-enterprise-storage/\)](https://www.suse.com/products/suse-enterprise-storage/) for details. SUSE CaaS Platform requires a minimum of four hosts: one admin and three Kubernetes nodes. SUSE Cloud Foundry is then deployed on the Kubernetes nodes. Four CaaS Platform nodes are not sufficient for a production deployment. *Figure 1.1, "Minimal Example Production Deployment"* describes a minimal production deployment with SUSE Cloud Foundry deployed on a Kubernetes cluster containing three Kubernetes masters and three workers, plus an ingress controller, administration workstation, DNS/DHCP server, and a SUSE Enterprise Storage cluster.

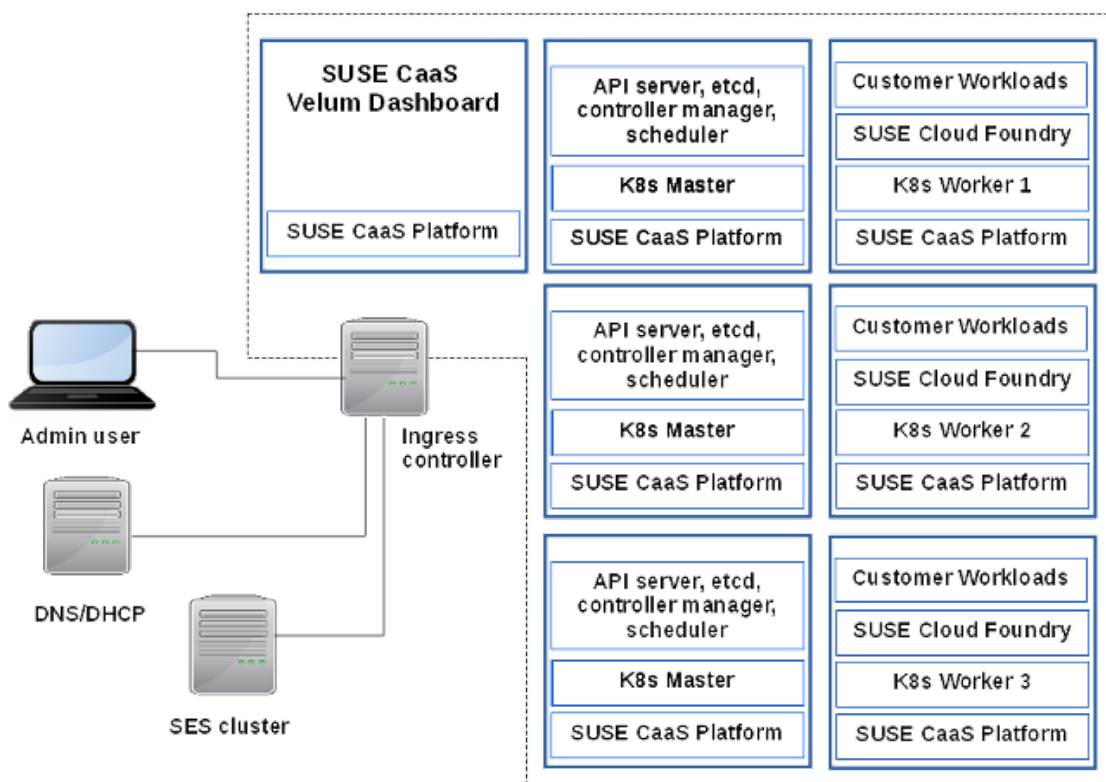


FIGURE 1.1: MINIMAL EXAMPLE PRODUCTION DEPLOYMENT

The minimum 4-node deployment is sufficient for a compact test deployment, which you can run virtualized on a single workstation or laptop. *Chapter 2, Production Installation with Optional High Availability* details a basic production deployment, and *Chapter 8, Minimal Installation for Testing* describes a minimal test deployment.

Note that after you have deployed your cluster and start building and running applications, your applications may depend on buildpacks that are not bundled in the container images that ship with SUSE Cloud Foundry. These will be downloaded at runtime, when you are pushing applications to the platform. Some of these buildpacks may include components with proprietary licenses. (See *Customizing and Developing Buildpacks* (<https://docs.cloudfoundry.org/buildpacks/developing-buildpacks.html>)⁷ to learn more about buildpacks, and creating and managing your own.)

1.2 New in Version 1.1

These are some of the changes in the 1.1 release (April 2017). See the [Release Notes \(https://www.suse.com/releasenotes/\)](https://www.suse.com/releasenotes/) for a complete list and known issues. See the Release Notes and [Section 2.13, "Upgrading SUSE Cloud Foundry, UAA, and Stratos"](#) for upgrade instructions.

- SUSE Cloud Application Platform 1.1 supports Azure Container Service (AKS).
- cf backup CLI plugin for saving, restoring, or migrating CF data and applications.
- PostgreSQL and MySQL service broker sidecars, configured and deployed via Helm.
- Cloud Foundry component and buildpack updates.
- Stratos UI 1.1 is required for SUSE Cloud Application Platform 1.1, and older versions will not work. Use the same `scf-config-values.yaml` file for both.
- The Helm command line client must be version 2.6.0 or higher.
- There are some changes in the `scf-config-values.yaml` configuration file. The variable `kube.external_ip` has been changed to `kube.external_ips`. Upgrades from older versions will fail unless the latter variable exists in the `scf-config-values.yaml` file. Both variables can exist at the same time, and be set to the same value for those in mixed version environments.

For example, when upgrading enter both variables in `scf-config-values.yaml`:

```
external_ip=1.1.1.1
external_ips=[1.1.1.1]
```

Going forward, `kube.external_ips` is an array, like this example:

```
external_ips=["1.1.1.1", "2.2.2.2"]
```

- All the secrets have been renamed from `env.FOO` to `secrets.FOO`, so all the appropriate entries in `scf-config-values.yaml` must be modified to align with that change. For example, move `CLUSTER_ADMIN_PASSWORD:` and `UAA_ADMIN_CLIENT_SECRET:` from the `env:` section to `secrets:`. Some of the `helm` commands for passing secrets also change, e.g. from `--set "env.FOO"` to `--set "secrets.FOO"`.

You must specify your secrets on each upgrade (e.g. the `CLUSTER_ADMIN_PASSWORD`) as they won't be carried forward automatically.

To rotate secrets, increment the `kube.secret_generation_counter`. (Please note: immutable generated secrets will not be reset.)

- Some roles (like `diego-api`, `diego-brain` and `routing-api`) are configured as active/passive, so passive pods can appear as Not Ready. Other roles (`tcp-router` and `blobstore`) cannot be scaled.

2 Production Installation with Optional High Availability

A basic SUSE Cloud Application Platform production deployment requires at least eight hosts plus a storage backend: one SUSE CaaS Platform admin server, three Kubernetes masters, three Kubernetes workers, a DNS/DHCP server, and a storage backend such as SUSE Enterprise Storage. This is a bare minimum, and actual requirements are likely to be much larger, depending on your workloads. You also need an external workstation for administering your cluster. You may optionally make your SUSE Cloud Foundry instance highly-available.



Note: Remote Administration

You will run most of the commands in this chapter from a remote workstation, rather than directly on any of the SUSE Cloud Foundry nodes. These are indicated by the unprivileged user Tux, while root prompts are on a cluster node. There are few tasks that need to be performed directly on any of the cluster hosts.

The optional High Availability example in this chapter provides HA only for the SUSE Cloud Foundry cluster, and not for CaaS Platform or SUSE Enterprise Storage. See [Section 2.14, “Example High Availability Configuration”](#).

2.1 Prerequisites

Calculating hardware requirements is best done with an analysis of your expected workloads, traffic patterns, storage needs, and application requirements. The following examples are bare minimums to deploy a running cluster, and any production deployment will require more.

Minimum Hardware Requirements

- 8GB of memory per CaaS Platform dashboard and Kubernetes master nodes.

- 16GB of memory per Kubernetes worker.

- 40GB disk space per CaaS Platform dashboard and Kubernetes master nodes.

- 60GB disk space per Kubernetes worker.

Network Requirements

Your Kubernetes cluster needs its own domain and network. Each node should resolve to its hostname, and to its fully-qualified domain name. Typically, a Kubernetes cluster sits behind a load balancer, which also provides external access to the cluster. Another option is to use DNS round-robin to the Kubernetes workers to provide external access. It is also a common practice to create a wildcard DNS entry pointing to the domain, e.g. *.example.com, so that applications can be deployed without creating DNS entries for each application. This guide does not describe how to set up a load balancer or name services, as these depend on customer requirements and existing network architectures.

Install SUSE CaaS Platform 2

After installing [CaaS Platform](https://www.suse.com/documentation/suse-caasp-2/) (<https://www.suse.com/documentation/suse-caasp-2/>) and logging into the Velum Web interface, check the box to install Tiller (Helm's server component).

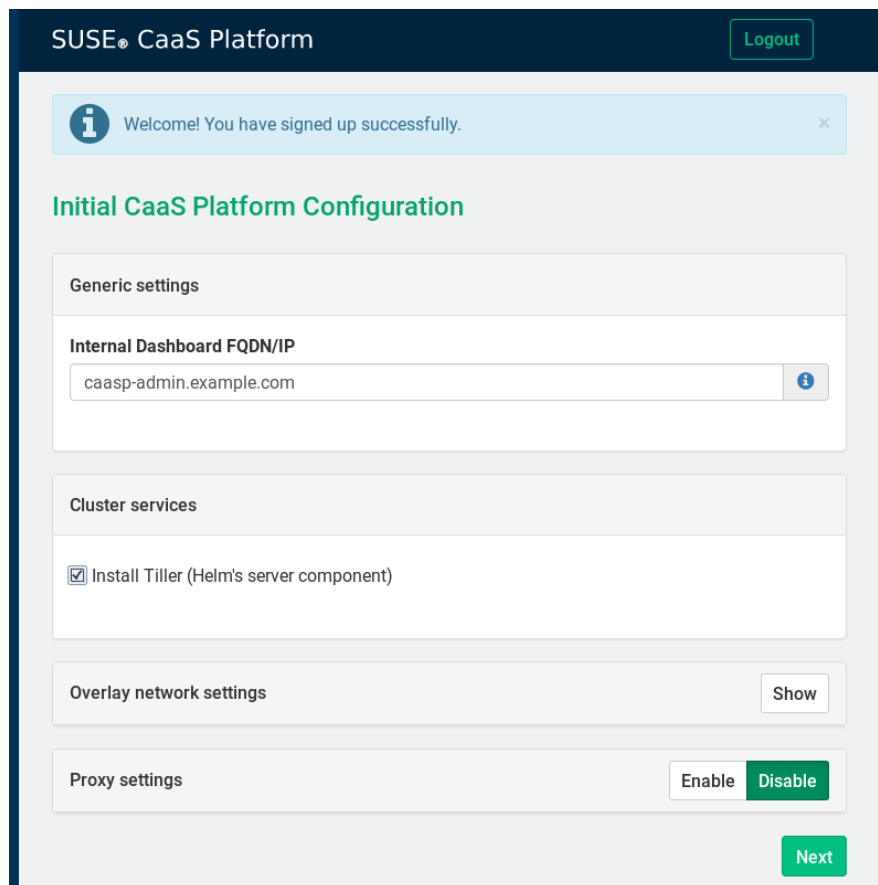


FIGURE 2.1: INSTALL TILLER

Take note of the *Overlay network settings*. These define the networks that are exclusive to the internal Kubernetes cluster communications. They are not externally accessible. You may assign different networks to avoid address collisions.

There is also a form for proxy settings; if you're not using a proxy then leave it empty. The easiest way to create the Kubernetes nodes, after you create the admin node, is to use AutoYaST; see [Installation with AutoYaST](https://www.suse.com/documentation/suse-caasp-2/book_caasp_deployment/data/sec_caasp_installquick.html#sec_caasp_installquick_node_ay) (https://www.suse.com/documentation/suse-caasp-2/book_caasp_deployment/data/sec_caasp_installquick.html#sec_caasp_installquick_node_ay). Set up CaaS Platform with one admin node and at least three Kubernetes masters and three Kubernetes workers. You also need an Internet connection, as the installer downloads additional packages, and the Kubernetes workers will each download ~10GB of Docker images.

Select nodes and roles

6 nodes found ✓ Select remaining nodes

After choosing the master and clicking "Next" all the other selected nodes will be set to the worker role.

ID	Hostname	Role
9c737b9a57c64db2ac4c34df0ba85a7a	VT1-node5	Master Worker Unused
2dd527766708445db00252c2e269b88c	VT1-node2	Master Worker Unused
898a8a1da27741258b4d0f1c7d257ce4	VT1-node1	Master Worker Unused
fa38abf50adc4d67a70d3f2499236028	VT1-node4	Master Worker Unused
961c280e45db4f86be31ea723a1e8aef	VT1-node3	Master Worker Unused
32bb549557154a44beacde4f508752bb	VT1-master	Master Worker Unused

Back Next

FIGURE 2.2: ASSIGNING ROLES TO NODES

When you have completed [Bootstrapping the Cluster](https://www.suse.com/documentation/suse-caasp-2/book_caasp_deployment/data/sec_caasp_installquick.html#sec_caasp_installquick_bootstrap) (https://www.suse.com/documentation/suse-caasp-2/book_caasp_deployment/data/sec_caasp_installquick.html#sec_caasp_installquick_bootstrap) click the `kubectl config` button to download your new cluster's `kubeconfig` file. This takes you to a login screen; use the login you created to access Velum. Save the file as `~/.kube/config` on your workstation. This file enables the remote administration of your cluster.

Cluster Status

Summary

Total nodes 6

Master nodes 3

New nodes ⓘ 0 (new)

Updates

Manual

of nodes w/ outdated software 0

Nodes

📄 kubectl config

FIGURE 2.3: DOWNLOAD KUBECONFIG

Install kubectl

Follow the instructions at [Install and Set Up kubectl \(https://kubernetes.io/docs/tasks/tools/install-kubectl/\)](https://kubernetes.io/docs/tasks/tools/install-kubectl/) to install **kubectl** on your workstation. After installation, run this command to verify that it is installed, and that it is communicating correctly with your cluster:

```
tux > kubectl version --short
Client Version: v1.9.1
Server Version: v1.7.7
```

As the client is on your workstation, and the server is on your cluster, reporting the server version verifies that **kubectl** is using `~/.kube/config` and is communicating with your cluster.

The following **kubectl** examples query the cluster configuration and node status:

```
tux > kubectl config view
apiVersion: v1
clusters:
- cluster:
  certificate-authority-data: REDACTED
  server: https://192.168.10.101:6443
  name: local
contexts:
[...]
```

```
tux > kubectl get nodes
NAME                                STATUS              ROLES    AGE   VERSION
b70748d.example.com                 Ready              <none>   4h   v1.7.7
cb77881.example.com                 Ready,SchedulingDisabled <none>   4h   v1.7.7
d028551.example.com                 Ready              <none>   4h   v1.7.7
[...]
```

Install Helm

Deploying SUSE Cloud Foundry is different than the usual method of installing software. Rather than installing packages in the usual way with YaST or Zypper, you will install the Helm client on your workstation to install the required Kubernetes applications to set up SUSE Cloud Foundry, and to administer your cluster remotely.

Helm client version 2.6 or higher is required.



Warning: Initialize Only the Helm Client

When you initialize Helm on your workstation be sure to initialize only the client, as the server, Tiller, was installed during the CaaS Platform installation. You do not want two Tiller instances.

If the Linux distribution on your workstation doesn't provide the correct Helm version, or you are using some other platform, see the [Helm Quickstart Guide \(https://docs.helm.sh/using_helm/#quickstart\)](https://docs.helm.sh/using_helm/#quickstart) for installation instructions and basic usage examples. Download the Helm binary into any directory that is in your PATH on your workstation, such as your `~/bin` directory. Then initialize the client only:

```
tux > helm init --client-only
Creating /home/tux/.helm
Creating /home/tux/.helm/repository
Creating /home/tux/.helm/repository/cache
Creating /home/tux/.helm/repository/local
```

```
Creating /home/tux/.helm/plugins
Creating /home/tux/.helm/starters
Creating /home/tux/.helm/cache/archive
Creating /home/tux/.helm/repository/repositories.yaml
Adding stable repo with URL: https://kubernetes-charts.storage.googleapis.com
Adding local repo with URL: http://127.0.0.1:8879/charts
$HELM_HOME has been configured at /home/tux/.helm.
Not installing Tiller due to 'client-only' flag having been set
Happy Helming!
```

2.2 Choose Storage Class

The Kubernetes cluster requires a persistent storage class for the databases to store persistent data. Your available storage classes depend on which storage cluster you are using (SUSE Enterprise Storage users, see [SUSE CaaS Platform Integration with SES \(https://www.suse.com/documentation/suse-caasp-2/book_caasp_deployment/data/integration.html\)](https://www.suse.com/documentation/suse-caasp-2/book_caasp_deployment/data/integration.html)). After connecting your storage backend use `kubectl` to see your available storage classes:

```
tux > kubectl get storageclasses
```

See [Section 2.4, “Configuring the SUSE Cloud Foundry Production Deployment”](#) to learn where to configure your storage class for SUSE Cloud Foundry. See the Kubernetes document [Persistent Volumes \(https://kubernetes.io/docs/concepts/storage/persistent-volumes/\)](https://kubernetes.io/docs/concepts/storage/persistent-volumes/) for detailed information on storage classes.

2.3 Test Storage Class

You may test that your storage class is properly configured before deploying SUSE Cloud Foundry by creating a persistent volume claim on your storage class, then verifying that the status of the claim is `bound`, and a volume has been created.

First copy the following configuration file, which in this example is named `test-storage-class.yaml`, substituting the name of your `storageClassName`:

```
---
kind: PersistentVolumeClaim
apiVersion: v1
metadata:
  name: test-sc-persistent
spec:
  accessModes:
```

```
- ReadWriteOnce
resources:
  requests:
    storage: 1Gi
    storageClassName: persistent
```

Create your persistent volume claim:

```
tux > kubectl create -f test-storage-class.yaml
persistentvolumeclaim "test-sc-persistent" created
```

Check that the claim has been created, and that the status is bound:

```
tux > kubectl get pv,pvc
NAME                                     CAPACITY  ACCESS MODES  RECLAIM POLICY
STATUS  CLAIM                                     STORAGECLASS  REASON  AGE
pv/pvc-c464ed6a-3852-11e8-bd10-90b8d0c59f1c  1Gi      RW0           Delete
Bound   default/test-sc-persistent  persistent                    2m

NAME          STATUS  VOLUME                                     CAPACITY
ACCESS MODES  STORAGECLASS  AGE
pvc/test-sc-persistent  Bound   pvc-c464ed6a-3852-11e8-bd10-90b8d0c59f1c  1Gi
RW0           persistent    2m
```


This verifies that your storage class is correctly configured. Delete your volume claims when you're finished:

```
tux > kubectl delete pv/pvc-c464ed6a-3852-11e8-bd10-90b8d0c59f1c
persistentvolume "pvc-c464ed6a-3852-11e8-bd10-90b8d0c59f1c" deleted
tux > kubectl delete pvc/test-sc-persistent
persistentvolumeclaim "test-sc-persistent" deleted
```

If something goes wrong and your volume claims get stuck in pending status, you can force deletion with the **--grace-period=0** option:

```
tux > kubectl delete pvc/test-sc-persistent --grace-period=0
```

2.4 Configuring the SUSE Cloud Foundry Production Deployment

Create a configuration file on your workstation for Helm to use. In this example it is called `scf-config-values.yaml`. (See the [Release Notes \(https://www.suse.com/releasenotes/\)](https://www.suse.com/releasenotes/)  for information on configuration changes.)

```
env:
```

```

# Enter the domain you created for your CAP cluster
DOMAIN: example.com

# UAA host and port
UAA_HOST: uaa.example.com
UAA_PORT: 2793

kube:
# The IP address assigned to the kube node pointed to by the domain.
external_ips: ["192.168.10.101"]

# Run kubectl get storageclasses
# to view your available storage classes
storage_class:
  persistent: "persistent"
  shared: "shared"

# The registry the images will be fetched from.
# The values below should work for
# a default installation from the SUSE registry.
registry:
  hostname: "registry.suse.com"
  username: ""
  password: ""
  organization: "cap"

# Required for CaaS 2
auth: rbac

secrets:
# Create a password for your CAP cluster
CLUSTER_ADMIN_PASSWORD: password

# Create a password for your UAA client secret
UAA_ADMIN_CLIENT_SECRET: password

```

2.5 Deploy with Helm

Run the following Helm commands to complete the deployment. There are six steps, and they must be run in this order:

- Download the SUSE Kubernetes charts repository
- Create the UAA and SCF namespaces

- Copy the storage secret of your storage cluster to the UAA and SCF namespaces
- Deploy UAA
- Copy the UAA secret and certificate to the SCF namespace
- Deploy SUSE Cloud Foundry

2.6 Install the Kubernetes charts repository

Download the SUSE Kubernetes charts repository with Helm:

```
tux > helm repo add suse https://kubernetes-charts.suse.com/
```

You may replace the example `suse` name with any name. Verify with `helm`:

```
tux > helm repo list
NAME          URL
stable        https://kubernetes-charts.storage.googleapis.com
local         http://127.0.0.1:8879/charts
suse          https://kubernetes-charts.suse.com/
```

List your chart names, as you will need these for some operations:

```
tux > helm search suse
NAME          VERSION  DESCRIPTION
suse/cf       2.8.0    A Helm chart for SUSE Cloud Foundry
suse/console  1.1.0    A Helm chart for deploying Stratos UI Console
suse/uaa      2.8.0    A Helm chart for SUSE UAA
```

2.7 Create Namespaces

Create the UAA (User Account and Authentication) and SCF (SUSE Cloud Foundry) namespaces:

```
tux > kubectl create namespace uaa
tux > kubectl create namespace scf
```

2.8 Copy SUSE Enterprise Storage Secret

If you are using SUSE Enterprise Storage you must copy the Ceph admin secret to the UAA and SCF namespaces:

```
tux > kubectl get secret ceph-secret-admin -o json --namespace default | \
sed 's/"namespace": "default"/"namespace": "uaa"/' | kubectl create -f -

tux > kubectl get secret ceph-secret-admin -o json --namespace default | \
sed 's/"namespace": "default"/"namespace": "scf"/' | kubectl create -f -
```

2.9 Deploy UAA

Use Helm to deploy the UAA (User Account and Authentication) server. You may create your own release `--name`:

```
tux > helm install suse/uaa \
--name susecf-uaa \
--namespace uaa \
--values scf-config-values.yaml
```

Wait until you have a successful UAA deployment before going to the next steps, which you can monitor with the `watch` command:

```
tux > watch -c 'kubectl get pods --all-namespaces'
```

When the status shows RUNNING for all of the UAA nodes, proceed to deploying SUSE Cloud Foundry. Pressing `Ctrl-C` stops the `watch` command.

2.10 Deploy SUSE Cloud Foundry

First pass your UAA secret and certificate to SCF, then use Helm to install SUSE Cloud Foundry:

```
tux > SECRET=$(kubectl get pods --namespace uaa \
-o jsonpath='{.items[*].spec.containers[?(.name=="uaa")].env[?
(.name=="INTERNAL_CA_CERT")].valueFrom.secretKeyRef.name}')

tux > CA_CERT="$(kubectl get secret $SECRET --namespace uaa \
-o jsonpath='{.data['internal-ca-cert']}'" | base64 --decode -)"
```



```
tux > helm install suse/cf \  
--name susecf-scf \  
--namespace scf \  
--values scf-config-values.yaml \  
--set "secrets.UAA_CA_CERT=${CA_CERT}"
```

Now sit back and wait for the pods come online:

```
tux > watch -c 'kubectl get pods --all-namespaces'
```

When all services are running use the Cloud Foundry command-line interface to log in to SUSE Cloud Foundry to deploy and manage your applications. (See [Section 2.11, “Deploying and Managing Applications with the Cloud Foundry Client”](#))

2.11 Deploying and Managing Applications with the Cloud Foundry Client

The Cloud Foundry command line interface (cf-cli) is for deploying and managing your applications. You may use it for all the orgs and spaces that you are a member of. Install the client on a workstation for remote administration of your SUSE Cloud Foundry instances.

The complete guide is at [Using the Cloud Foundry Command Line Interface \(https://docs.cloud-foundationry.org/cf-cli/\)](https://docs.cloud-foundationry.org/cf-cli/), and source code with a demo video is on GitHub at [Cloud Foundry CLI \(https://github.com/cloudfoundry/cli/blob/master/README.md\)](https://github.com/cloudfoundry/cli/blob/master/README.md).

The following examples demonstrate some of the commonly-used commands. The first task is to log into your new SUSE Cloud Foundry instance. When your installation completes it prints a welcome screen with the information you need to access it.

```
NOTES:  
Welcome to your new deployment of SCF.  
  
The endpoint for use by the `cf` client is  
https://api.example.com  
  
To target this endpoint run  
cf api --skip-ssl-validation https://api.example.com  
  
Your administrative credentials are:  
Username: admin  
Password: password
```

Please remember, it may take some time for everything to come online.

You can use

```
kubectl get pods --namespace scf
```

to spot-check if everything is up and running, or

```
watch -c 'kubectl get pods --namespace scf'
```

to monitor continuously.

You can display this message anytime with this command:

```
tux > helm status $(helm list | awk '/cf-([0-9]).([0-9]).*/{print$1}') | \
sed -n -e '/NOTES/, $p'
```

You need to provide the API endpoint of your SUSE Cloud Foundry instance to log in. The API endpoint is the DOMAIN value you provided in scf-config-values.yaml, plus the api. prefix, as it shows in the above welcome screen. Set your endpoint, and use **--skip-ssl-validation** when you have self-signed SSL certificates. It asks for an email address, but you must enter admin instead (you cannot change this to a different username, though you may create additional users), and the password is the one you created in scf-config-values.yaml:

```
tux > cf login --skip-ssl-validation -a https://api.example.com
API endpoint: https://api.example.com

Email> admin

Password>
Authenticating...
OK

Targeted org system

API endpoint: https://api.example.com (API version: 2.101.0)
User: admin
Org: system
Space: No space targeted, use 'cf target -s SPACE'
```

cf help displays a list of commands and options. **cf help [command]** provides information on specific commands.

You may pass in your credentials and set the API endpoint in a single command:

```
tux > cf login -u admin -p password --skip-ssl-validation -a https://api.example.com
```

Log out with `cf logout`.

View your current API endpoint, user, org, and space:

```
tux > cf target
```

Switch to a different org or space:

```
tux > cf target -o org
tux > cf target -s space
```

List all apps in the current space:

```
tux > cf apps
```

Query the health and status of a particular app:

```
tux > cf app appname
```

View app logs. The first example tails the log of a running app. The `--recent` option dumps recent logs instead of tailing, which is useful for stopped and crashed apps:

```
tux > cf logs appname
tux > cf logs --recent appname
```

Restart all instances of an app:

```
tux > cf restart appname
```

Restart a single instance of an app, identified by its index number, and restart it with the same index number:

```
tux > cf restart-app-instance appname index
```

After you have set up a service broker (see [Chapter 3, Setting up and Using a Service Broker Sidecar](#)), create new services:

```
tux > cf create-service service-name default mydb
```

Then you may bind a service instance to an app:

```
tux > cf bind-service appname service-instance
```

The most-used command is `cf push`, for pushing new apps and changes to existing apps.

```
tux > cf push new-app -b buildpack
```

2.12 Installing the Stratos Web Console

Stratos UI is a modern web-based management application for Cloud Foundry. It provides a graphical management console for both developers and system administrators. Install Stratos with Helm after all of the UAA and SCF pods are running. Start by preparing the environment:

```
tux > kubectl create namespace stratos
```

If you are using SUSE Enterprise Storage as your storage backend, copy the secret into the Stratos namespace.

```
tux > kubectl get secret ceph-secret-admin -o json --namespace default | \
sed 's/"namespace": "default"/"namespace": "stratos"/' | \
kubectl create -f -
```

You should already have the Stratos charts when you downloaded the SUSE charts repository:

```
tux > helm search suse
NAME          VERSION DESCRIPTION
suse/cf       2.8.0   A Helm chart for SUSE Cloud Foundry
suse/console  1.1.0   A Helm chart for deploying Stratos UI Console
suse/uaa      2.8.0   A Helm chart for SUSE UAA
```

Install Stratos, and if you have not set a default storage class you must specify it:

```
tux > helm install suse/console \
  --name susecf-console \
  --namespace stratos \
  --values scf-config-values.yaml \
  --set storageClass=persistent
```

Monitor progress:

```
$ watch -c 'kubectl get pods --namespace stratos'
Every 2.0s: kubectl get pods --namespace stratos

NAME                                READY   STATUS    RESTARTS   AGE
```

console-0	3/3	Running	0	30m
console-mariadb-3697248891-5drf5	1/1	Running	0	30m

When all statuses show Ready, press `Ctrl-C` to exit and to view your release information:

```

NAME:      susecf-console
LAST DEPLOYED: Thu Apr 12 10:28:34 2018
NAMESPACE: stratos
STATUS: DEPLOYED

RESOURCES:
==> v1/Secret
NAME                                TYPE      DATA  AGE
susecf-console-mariadb-secret      Opaque    2       2s
susecf-console-secret              Opaque    2       2s

==> v1/PersistentVolumeClaim
NAME                                STATUS    VOLUME
CAPACITY  ACCESSMODES  STORAGECLASS  AGE
console-mariadb      Bound        pvc-ef3a120d-3e76-11e8-946a-90b8d00d625f
1Gi          RWO          persistent    2s
susecf-console-upgrade-volume      Bound        pvc-ef409e41-3e76-11e8-946a-90b8d00d625f
20Mi        RWO          persistent    2s
susecf-console-encryption-key-volume Bound        pvc-ef49b860-3e76-11e8-946a-90b8d00d625f
20Mi        RWO          persistent    2s

==> v1/Service
NAME                                CLUSTER-IP      EXTERNAL-IP      PORT(S)          AGE
susecf-console-mariadb      172.24.181.255  <none>           3306/TCP         2s
susecf-console-ui-ext      172.24.84.50   10.10.100.82    8443:32511/TCP  1s

==> v1beta1/Deployment
NAME                                DESIRED  CURRENT  UP-TO-DATE  AVAILABLE  AGE
console-mariadb      1         1         1             0           1s

==> v1beta1/StatefulSet
NAME      DESIRED  CURRENT  AGE
console  1         1         1s

```

In this example, pointing your web browser to <https://10.10.100.82:8443> opens the console. Wade through the nag screens about the self-signed certificates and log in as `admin` with the password you created in `scf-config-values.yaml`. If you see an upgrade message, wait a few minutes and try again.

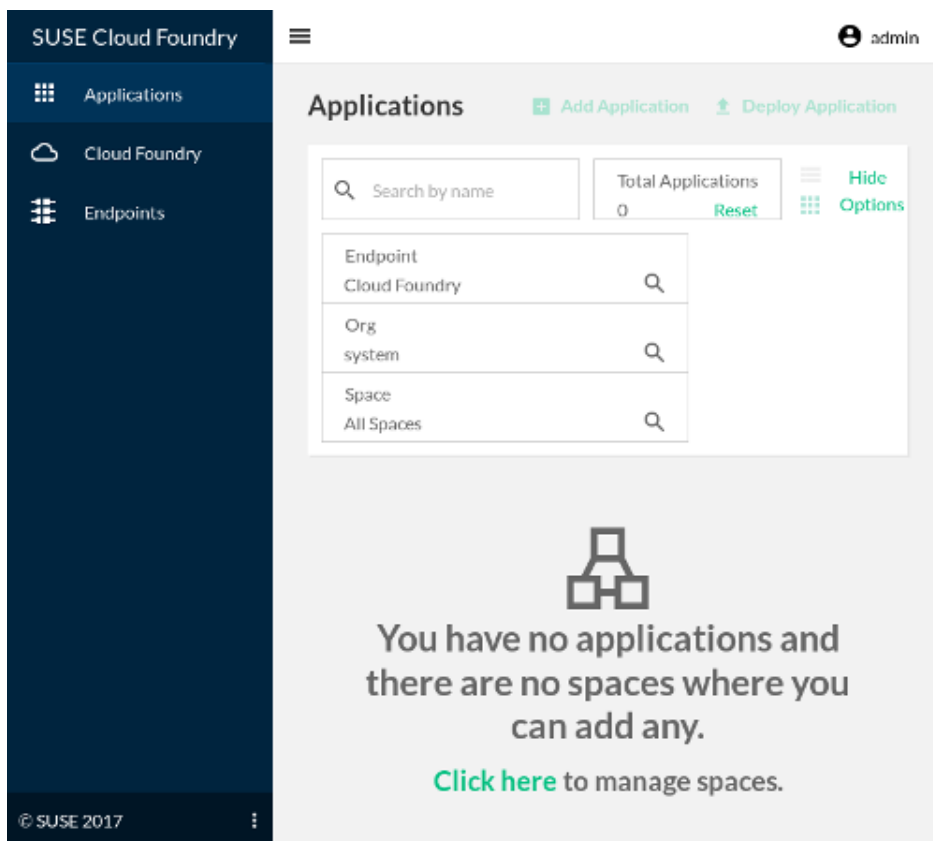


FIGURE 2.4: STRATOS UI CLOUD FOUNDRY CONSOLE

Another way to get the release name is with the `helm ls` command, then query the release name to get its IP address and port number:

```
tux > helm ls
NAME                REVISION  UPDATED                               STATUS  CHART              NAMESPACE
susecf-console      1         Thu Apr 12 10:28:34 2018    DEPLOYED  console-1.1.0     stratos
susecf-scf          1         Wed Apr 11 14:55:23 2018    DEPLOYED  cf-2.8.0          scf
susecf-uaa          1         Wed Apr 11 14:48:01 2018    DEPLOYED  uaa-2.8.0         uaa

tux > helm status susecf-console
LAST DEPLOYED: Thu Apr 12 10:28:34 2018
NAMESPACE: stratos
STATUS: DEPLOYED
[...]
==> v1/Service
NAME                CLUSTER-IP      EXTERNAL-IP    PORT(S)          AGE
susecf-console-mariadb 172.24.181.255 <none>         3306/TCP         19m
susecf-console-ui-ext 172.24.84.50    10.10.100.82  8443:32511/TCP  19m
```

2.13 Upgrading SUSE Cloud Foundry, UAA, and Stratos

Maintenance updates are delivered as container images from the SUSE registry and applied with Helm. (See the [Release Notes \(https://www.suse.com/releasenotes/\)](https://www.suse.com/releasenotes/) for additional upgrade information.) Check for available updates:

```
tux > helm repo update
Hang tight while we grab the latest from your chart repositories...
...Skip local chart repository
...Successfully got an update from the "stable" chart repository
...Successfully got an update from the "suse" chart repository
Update Complete. * Happy Helming!*
```

For the SUSE Cloud Application Platform 1.1 release, update your `scf-config-values.yaml` file with the changes for secrets handling and external IP addresses. (See [Section 2.4, "Configuring the SUSE Cloud Foundry Production Deployment"](#) for an example.)

Get your release and chart names (your releases may have different names than the examples), and then apply the updates:

```
tux > helm list
NAME                REVISION  UPDATED                               STATUS  CHART              NAMESPACE
susecf-console     1         Thu Apr 12 10:28:34 2018    DEPLOYED  console-1.0.2     stratos
susecf-scf         1         Wed Apr 11 14:55:23 2018    DEPLOYED  cf-2.7.0          scf
susecf-uaa         1         Wed Apr 11 14:48:01 2018    DEPLOYED  uaa-2.7.0         uaa

tux > helm repo list
NAME      URL
stable   https://kubernetes-charts.storage.googleapis.com
local    http://127.0.0.1:8879/charts
suse     https://kubernetes-charts.suse.com/

tux > helm search suse
NAME      VERSION  DESCRIPTION
suse/cf   2.8.0    A Helm chart for SUSE Cloud Foundry
suse/console  1.1.0    A Helm chart for deploying Stratos UI Console
suse/uaa  2.8.0    A Helm chart for SUSE UAA
```

Run the following commands to perform the upgrade. Wait for each command to complete before running the next command. Note the new commands for extracting and using secrets and certificates.

```
tux > helm upgrade --recreate-pods susecf-uaa suse/uaa \
--values scf-config-values.yaml
```

```
tux > SECRET=$(kubectl get pods --namespace uaa \
-o jsonpath='{.items[*].spec.containers[?(.name=="uaa")].env[?
(.name=="INTERNAL_CA_CERT")].valueFrom.secretKeyRef.name}')

tux > CA_CERT="$(kubectl get secret $SECRET --namespace uaa \
-o jsonpath="{.data['internal-ca-cert']}" | base64 --decode -)"

tux > helm upgrade --recreate-pods susecf-scf suse/cf \
--values scf-config-values.yaml --set "secrets.UAA_CA_CERT=${CA_CERT}"

tux > helm upgrade --recreate-pods susecf-console suse/console \
--values scf-config-values.yaml
```

2.14 Example High Availability Configuration

This example High Availability configuration needs two separate configuration files, one for UAA and one for SCF. The first example is for UAA, [uaa-sizing.yaml](#).

```
sizing:
  api:
    count: 2
  cf_usb:
    count: 2
  consul:
    count: 3
  diego_access:
    count: 2
  diego_api:
    count: 3
  diego_brain:
    count: 2
  diego_cell:
    count: 3
  doppler:
    count: 2
  etcd:
    count: 3
  loggregator:
    count: 2
  mysql:
    count: 1
  nats:
    count: 2
  router:
```



```
count: 2
routing_api:
count: 2
```

The second example is for SCF, `scf-sizing.yaml`.

```
sizing:
  api:
    count: 2
  cf_usb:
    count: 2
  consul:
    count: 3
  diego_access:
    count: 2
  diego_api:
    count: 3
  diego_brain:
    count: 2
  diego_cell:
    count: 3
  doppler:
    count: 2
  etcd:
    count: 3
  loggregator:
    count: 2
  mysql:
    count: 3
  nats:
    count: 2
  router:
    count: 2
  routing_api:
    count: 2
```

Follow the steps in [Section 2.4, “Configuring the SUSE Cloud Foundry Production Deployment”](#) until you get to [Section 2.9, “Deploy UAA”](#). Then deploy UAA with this command:

```
tux > helm install suse/uaa \
--name susecf-uaa \
--namespace uaa \
--values scf-config-values.yaml \
--values uaa-sizing.yaml
```

When the status shows RUNNING for all of the UAA nodes, deploy SCF with these commands:

```
tux > SECRET=$(kubectl get pods --namespace uaa \
```

```

-o jsonpath='{.items[*].spec.containers[?(.name=="uaa")].env[?
(.name=="INTERNAL_CA_CERT)].valueFrom.secretKeyRef.name}'')

tux > CA_CERT="$(kubectl get secret $SECRET --namespace uaa \
-o jsonpath="{.data['internal-ca-cert']}" | base64 --decode -)"

tux > helm install suse/cf \
--name susecf-scf
--namespace scf \
--values scf-config-values.yaml \
--values scf-sizing.yaml
--set "secrets.UAA_CA_CERT=${CA_CERT}"

```

The HA pods with the following roles will enter in both passive and ready states; there should always be at least one pod in each role that is ready.

- diego-brain
- diego-database
- routing-api

You can confirm this by looking at the logs inside the container. Look for `.consul-lock.acquiring-lock`.

Some roles cannot be scaled. `mysql-proxy` needs a proper active/passive configuration. `tcp-router` has no mechanism for exposing ports correctly. `blobstore` needs shared volume support and an active/passive configuration.

Some roles follow an active/passive scaling model, meaning all pods except the active one will be shown as NOT READY by Kubernetes. This is appropriate and expected behavior.

2.14.1 Upgrading a non-High Availability Deployment to High Availability

You may make a non-High Availability deployment highly available by upgrading with Helm:

```

tux > helm upgrade suse/uaa \
--name susecf-uaa \
--namespace uaa \
--values scf-config-values.yaml \
--values uaa-sizing.yaml

tux > SECRET=$(kubectl get pods --namespace uaa \

```

```
-o jsonpath='{.items[*].spec.containers[?(.name=="uaa")].env[?
(.name=="INTERNAL_CA_CERT")].valueFrom.secretKeyRef.name}'

tux > CA_CERT="$(kubectl get secret $SECRET --namespace uaa \
-o jsonpath="{.data['internal-ca-cert']}" | base64 --decode -)"

tux > helm upgrade suse/cf \
--name susecf-scf
--namespace scf \
--values scf-config-values.yaml \
--values scf-sizing.yaml
--set "secrets.UAA_CA_CERT=${CA_CERT}"
```

This may take a long time, and your cluster will be unavailable until the upgrade is complete.

3 Setting up and Using a Service Broker Sidecar

The Open Service Broker API provides your SUSE Cloud Foundry applications with access to external dependencies and platform-level capabilities, such as databases, filesystems, external repositories, and messaging systems. These resources are called services. Services are created, used, and deleted as needed, and provisioned on demand.

3.1 Prerequisites

The following examples demonstrate how to deploy service brokers for MySQL and PostgreSQL with Helm, using charts from the SUSE repository. You must have the following prerequisites:

- A working SUSE Cloud Application Platform deployment with Helm and the Cloud Foundry command line interface (cf-cli).
- An Application Security Group (ASG) for applications to reach external databases. (See [Understanding Application Security Groups \(https://docs.cloudfoundry.org/concepts/asg.html\)](https://docs.cloudfoundry.org/concepts/asg.html).)
- An external MySQL or PostgreSQL installation with account credentials that allow creating and deleting databases and users.

For testing purposes you may create an insecure security group:

```
tux > echo > "internal-services.json" '[{"destination": "0.0.0.0/0", "protocol":  
"all" }]'  
tux > cf create-security-group internal-services-test internal-services.json  
tux > cf bind-running-security-group internal-services-test  
tux > cf bind-staging-security-group internal-services-test
```

You may apply an ASG later, after testing. All running applications must be restarted to use the new security group.

3.2 Configuring the MySQL Deployment

Start by extracting the `uaa` namespace secrets name, and the `uaa` and `scf` namespaces internal certificates with these commands. These will output the complete certificates. Substitute your secrets name if it is different than the example:

```
tux > kubectl get pods --namespace uaa \
```

```

-o jsonpath='{.items[*].spec.containers[?(.name=="uaa")].env[?
(.name=="INTERNAL_CA_CERT)].valueFrom.secretKeyRef.name}'
secrets-2.8.0-1

tux > kubectl get secret -n scf secrets-2.8.0-1 -o jsonpath='{.data.internal-ca-cert}' |
base64 -d
-----BEGIN CERTIFICATE-----
MIIE8jCCAtqgAwIBAgIUT/Yu/Sv4UHl5zHZYZKCy5RKJqmYwDQYJKoZIhvcNAQEN
[...]
xC8x/+zT0QkvcRJBio5gg670+25KJQ==
-----END CERTIFICATE-----

tux > kubectl get secret -n uaa secrets-2.8.0-1 -o jsonpath='{.data.internal-ca-cert}' |
base64 -d
-----BEGIN CERTIFICATE-----
MIIE8jCCAtqgAwIBAgIU5I02lj0a0InLb/zMrjNgW5d8EygwDQYJKoZIhvcNAQEN
[...]
to2GI8rPmb9W9fd2WwUXGEHTc+PqTg==
-----END CERTIFICATE-----

```

You will copy these certificates into your configuration file as shown below.

Create a `values.yaml` file. The following example is called `usb-config-values.yaml`. Modify the values to suit your SUSE Cloud Application Platform installation.

```

env:
  # Database access credentials
  SERVICE_MYSQL_HOST: mysql.example.com
  SERVICE_MYSQL_PORT: 3306
  SERVICE_MYSQL_USER: mysql-admin-user
  SERVICE_MYSQL_PASS: mysql-admin-password

  # CAP access credentials, from your original deployment configuration
  # (see Section 2.4, "Configuring the SUSE Cloud Foundry Production Deployment")
  CF_ADMIN_USER: admin
  CF_ADMIN_PASSWORD: password
  CF_DOMAIN: example.com

  # Copy the certificates you extracted above, as shown in these
  # abbreviated examples, prefaced with the pipe character

  # SCF cert
  CF_CA_CERT: |
    -----BEGIN CERTIFICATE-----
    MIIE8jCCAtqgAwIBAgIUT/Yu/Sv4UHl5zHZYZKCy5RKJqmYwDQYJKoZIhvcNAQEN
    [...]
    xC8x/+zT0QkvcRJBio5gg670+25KJQ==
    -----END CERTIFICATE-----

```

```

# UAA cert
UAA_CA_CERT: |
-----BEGIN CERTIFICATE-----
MIIE8jCCAAtqgAwIBAgIUUSI02Lj0a0InLb/zMrjNgW5d8EygwDQYJKoZIhvcNAQEN
[...]
to2GI8rPmb9W9fd2WwUXGEHTc+PqTg==
-----END CERTIFICATE-----

kube:
organization: cap
registry:
  hostname: "registry.suse.com"
  username: ""
  password: ""

```

3.3 Deploying the MySQL Chart

The 1.1 release of SUSE Cloud Application Platform includes charts for MySQL and PostgreSQL:

```

tux > helm search suse
NAME                                VERSION DESCRIPTION
suse/cf                              2.8.0   A Helm chart for SUSE Cloud Foundry
suse/cf-usb-sidecar-mysql            1.0.1   A Helm chart for SUSE Universal Service
Broker ...
suse/cf-usb-sidecar-postgres        1.0.1   A Helm chart for SUSE Universal Service
Broker ...
suse/console                         1.1.0   A Helm chart for deploying Stratos UI Console
suse/uaa                             2.8.0   A Helm chart for SUSE UAA

```

Create a namespace for your MySQL sidecar:

```
tux > kubectl create namespace mysql-sidecar
```

Install the MySQL Helm chart:

```

tux > helm install suse/cf-usb-sidecar-mysql \
  --devel \
  --name mysql-service \
  --namespace mysql-sidecar \
  --set "env.SERVICE_LOCATION=http://cf-usb-sidecar-mysql.mysql-sidecar:8081" \
  --values usb-config-values.yaml \

```

```
--wait
```

Wait for the new pods to become ready:

```
tux > watch kubectl get pods --namespace=mysql-sidecar
```

Confirm that the new service has been added to your SUSE Cloud Applications Platform installation:

```
tux > cf marketplace
```

3.4 Create and Bind a MySQL Service

To create a new service instance, use the Cloud Foundry command line client:

```
tux > cf create-service mysql default service_instance_name
```

You may replace *service_instance_name* with any name you prefer.

Bind the service instance to an application:

```
tux > cf bind-service my_application service_instance_name
```

3.5 Deploying the PostgreSQL Chart

The PostgreSQL configuration is slightly different from the MySQL configuration. The database-specific keys are named differently, and it requires the `SERVICE_POSTGRESQL_SSLMODE` key.

```
env:
  # Database access credentials
  SERVICE_POSTGRESQL_HOST: postgres.example.com
  SERVICE_POSTGRESQL_PORT: 5432
  SERVICE_POSTGRESQL_USER: pgsql-admin-user
  SERVICE_POSTGRESQL_PASS: pgsql-admin-password
  # The SSL connection mode when connecting to the database. For a list of
  # valid values, please see https://godoc.org/github.com/lib/pq
  SERVICE_POSTGRESQL_SSLMODE: disable

  # CAP access credentials, from your original deployment configuration
  # (see Section 2.4, "Configuring the SUSE Cloud Foundry Production Deployment")
  CF_ADMIN_USER: admin
  CF_ADMIN_PASSWORD: password
```

```

CF_DOMAIN: example.com

# Copy the certificates you extracted above, as shown in these
# abbreviated examples, prefaced with the pipe character

# SCF cert
CF_CA_CERT: |
-----BEGIN CERTIFICATE-----
MIIE8jCCAtqgAwIBAgIUT/Yu/Sv4UHL5zHZYZKCy5RKJqmYwDQYJKoZIhvcNAQEN
[...]
xC8x/+zT0QkvcRJBio5gg670+25KJQ==
-----END CERTIFICATE-----

# UAA cert
UAA_CA_CERT: |
-----BEGIN CERTIFICATE-----
MIIE8jCCAtqgAwIBAgIU02lj0a0InLb/zMrjNgW5d8EygwDQYJKoZIhvcNAQEN
[...]
to2GI8rPMb9W9fd2WwUXGEHTc+PqTg==
-----END CERTIFICATE-----
SERVICE_TYPE: postgres

kube:
organization: cap
registry:
  hostname: "registry.suse.com"
  username: ""
  password: ""

```

Create a namespace and install the chart:

```

tux > kubectl create namespace postgres-sidecar

tux > helm install suse/cf-usb-sidecar-postgres \
  --devel \
  --name postgres-service \
  --namespace postgres-sidecar \
  --set "env.SERVICE_LOCATION=http://cf-usb-sidecar-postgres.postgres-sidecar:8081" \
  --values usb-config-values.yaml \
  --wait

```

Then follow the same steps as for the MySQL chart.

3.6 Removing Service Broker Sidecar Deployments

To correctly remove sidecar deployments, perform the following steps in order.

- Unbind any applications using instances of the service, and then delete those instances:

```
tux > cf unbind-service my_app my_service_instance
tux > cf delete-service my_service_instance
```

- Install the CF-USB CLI plugin for the Cloud Foundry CLI:

```
tux > cf install-plugin \
https://github.com/SUSE/cf-usb-plugin/releases/download/1.0.0/cf-usb-
plugin-1.0.0.0.g47b49cd-linux-amd64
```

- Configure the Cloud Foundry USB CLI plugin, using the domain you created for your SUSE Cloud Foundry deployment:

```
tux > cf usb-target https://usb.example.com
```

- Remove the services:

```
tux > cf usb delete-driver-endpoint "http://cf-usb-sidecar-mysql.mysql-sidecar:8081"
```

- Find your release name, then delete the release:

```
tux > helm list
NAME                REVISION  UPDATED              STATUS   CHART
NAMESPACE
susecf-scf          1         Mon May 21 10:59:57 2018  DEPLOYED  cf-2.8.0
scf
susecf-uaa          1         Mon May 21 10:32:13 2018  DEPLOYED  uaa-2.8.0
uaa
mysql-service 1         Mon May 21 11:40:11 2018  DEPLOYED  cf-usb-sidecar-
mysql-1.0.1 mysql-sidecar

tux > helm delete --purge mysql-service
```

4 Backup and Restore

`cf-plugin-backup` backs up and restores your cloud controller database (CDDDB), using the Cloud Foundry command line interface (`cf-cli`). (See [Section 2.11, “Deploying and Managing Applications with the Cloud Foundry Client”](#).) Use it after a fresh, clean SUSE Cloud Application Platform deployment has been completed. Use the restore function to return your deployment to its original clean state, or to replicate your deployment.

`cf-plugin-backup` creates a JSON file that contains your SUSE Cloud Application Platform data in the current directory, `cf-backup.json`, and your application data in a directory called `app-bits/`.

4.1 Installing the `cf-plugin-backup`

Download the plugin from `cf-plugin-backup/releases` (<https://github.com/SUSE/cf-plugin-backup/releases>).

Then install it with `cf`, using the name of the plugin binary that you downloaded:

```
tux > cf install-plugin cf-plugin-backup-1.0.7.0.g0217eef.linux-amd64
Attention: Plugins are binaries written by potentially untrusted authors.
Install and use plugins at your own risk.
Do you want to install the plugin
backup-plugin/cf-plugin-backup-1.0.7.0.g0217eef.linux-amd64? [yN]: y
Installing plugin backup...
OK
Plugin backup 1.0.7 successfully installed.
```

Verify installation by listing installed plugins:

```
tux > cf plugins
Listing installed plugins...

plugin  version  command name  command help
backup  1.0.7    backup-info   Show information about the current snapshot
backup  1.0.7    backup-restore Restore the CloudFoundry state from a
        backup created with the snapshot command
backup  1.0.7    backup-snapshot Create a new CloudFoundry backup snapshot
        to a local file

Use 'cf repo-plugins' to list plugins in registered repos available to install.
```

4.2 Using cf-plugin-backup

View the online help for any command, like this example:

```
tux > cf backup-info --help
NAME:
  backup-info - Show information about the current snapshot

USAGE:
  cf backup-info
```

Create a backup of your SUSE Cloud Application Platform data and applications. The command outputs progress messages until it is completed:

```
tux > cf backup-snapshot
2018/06/18 12:48:27 Retrieving resource /v2/quota_definitions
2018/06/18 12:48:30 org quota definitions done
2018/06/18 12:48:30 Retrieving resource /v2/space_quota_definitions
2018/06/18 12:48:32 space quota definitions done
2018/06/18 12:48:32 Retrieving resource /v2/organizations
[...]
```

Your CAP data is saved in the current directory in cf-backup.json, and application data in the app-bits/ directory.

View the current backup:

```
tux > cf backup-info
- Org system
```

Restore from backup:

```
tux > cf backup-restore
```

There are two additional restore options: --include-security-groups and --include-quota-definitions.

The following tables lists the scope of the cf-plugin-backup backup. Organization and space users are backed up at the SUSE Cloud Application Platform level. The user account in UAA/LDAP, and the service instances and their application bindings are not backed up.

Scope	Restore
Orgs	Yes

Scope	Restore
Org auditors	Yes
Org billing-manager	Yes
Quota definitions	Optional
Spaces	Yes
Space developers	Yes
Space auditors	Yes
Space managers	Yes
Apps	Yes
App binaries	Yes
Routes	Yes
Route mappings	Yes
Domains	Yes
Private domains	Yes
Stacks	not available
Feature flags	Yes
Security groups	Optional
Custom buildpacks	No

5 Preparing Microsoft Azure for SUSE Cloud Application Platform

SUSE Cloud Application Platform version 1.1 and up supports deployment on Microsoft Azure Kubernetes Service (AKS), Microsoft's managed Kubernetes service. This chapter describes the steps for preparing Azure for a SUSE Cloud Application Platform deployment, with a basic Azure load balancer. (See [Azure Kubernetes Service \(AKS\) \(https://azure.microsoft.com/en-us/services/container-service/\)](https://azure.microsoft.com/en-us/services/container-service/) for more information.)

In Kubernetes terminology a node used to be a minion, which was the name for a worker node. Now the correct term is simply node (see <https://kubernetes.io/docs/concepts/architecture/nodes/>). This can be confusing, as computing nodes have traditionally been defined as any device in a network that has an IP address. In Azure they are called agent nodes. In this chapter we call them agent nodes or Kubernetes nodes.

5.1 Prerequisites

Install **az**, the Azure command-line client, on your remote administration machine. See [Install Azure CLI 2.0 \(https://docs.microsoft.com/en-us/cli/azure/install-azure-cli?view=azure-cli-latest\)](https://docs.microsoft.com/en-us/cli/azure/install-azure-cli?view=azure-cli-latest) for instructions.

See the [Azure CLI 2.0 Reference \(https://docs.microsoft.com/en-us/cli/azure/reference-index?view=azure-cli-latest\)](https://docs.microsoft.com/en-us/cli/azure/reference-index?view=azure-cli-latest) for a complete **az** command reference.

You also need the **kubect1**, **curl**, **sed**, and **jq** commands, and the name of the SSH key that is attached to your Azure account.

Log in to your Azure Account:

```
tux > az login
```

Your Azure user needs the User Access Administrator role. Check your assigned roles with the **az** command:

```
tux > az role assignment list --assignee login-name  
[...]  
"roleDefinitionName": "User Access Administrator",
```

If you do not have this role, then you must request it from your Azure administrator.

You need your Azure subscription ID. Extract it with **az**:

```
tux > az account show --query "{ subscription_id: id }"
{
  "subscription_id": "a900cdi2-5983-0376-s7je-d4jdmsif84ca"
}
```

Replace subscription-id in the next command with your subscription-id. Then export it as an environment variable and set it as the current subscription:

```
tux > export SUBSCRIPTION_ID=a900cdi2-5983-0376-s7je-d4jdmsif84ca"

tux > az account set --subscription $SUBSCRIPTION_ID
```

Verify that the Microsoft.Network, Microsoft.Storage, Microsoft.Compute, and Microsoft.ContainerService providers are enabled:

```
tux > az provider list | egrep -w 'Microsoft.Network|Microsoft.Storage|Microsoft.Compute|
Microsoft.ContainerService'
```

If any of these are missing, enable them with the **az provider register -n provider** command.

5.2 Create Resource Group and AKS Instance

Now you can create a new Azure resource group and AKS instance. Set the required variables as environment variables, which helps to speed up the setup, and to reduce errors.



Note: Use different names

It is better to use unique resource group and cluster names, and not copy the examples, especially when your Azure subscription supports multiple users.

1. Create and set the resource group name:

```
tux > export RGNAME="cap-aks"
```

2. Create and set the AKS managed cluster name. Azure's default is to use the resource group name, then prepend it with MC and append the location, e.g. MC_cap-aks_cap-aks_eastus. This example command gives it the same name as the resource group; you may give it a different name.

```
tux > export AKSNAME=$RGNAME
```

3. Set the Azure location. See [Quickstart: Deploy an Azure Kubernetes Service \(AKS\) cluster \(https://docs.microsoft.com/en-us/azure/aks/kubernetes-walkthrough\)](https://docs.microsoft.com/en-us/azure/aks/kubernetes-walkthrough) for supported locations. Current supported Azure locations are eastus, westeurope, centralus, canadacentral, and canadaeast.

```
tux > export REGION="eastus"
```

4. Set the Kubernetes agent node count. (CAP requires a minimum of 3.)

```
tux > export NODECOUNT="3"
```

5. Set the virtual machine size (see [Sizes for Cloud Services \(https://docs.microsoft.com/en-us/azure/cloud-services/cloud-services-sizes-specs\)](https://docs.microsoft.com/en-us/azure/cloud-services/cloud-services-sizes-specs)):

```
tux > export NODEVMsize="Standard_D2_v2"
```

6. Set the public SSH key name associated with your Azure account:

```
tux > export SSHKEYVALUE="~/.ssh/id_rsa.pub"
```

7. Create and set a new admin username:

```
tux > export ADMINUSERNAME="scf-admin"
```

Now that your environment variables are in place, create a new resource group:

```
tux > az group create --name $RGNAME --location $REGION
```

Create a new AKS managed cluster:

```
tux > az aks create --resource-group $RGNAME --name $AKSNAME \  
--node-count $NODECOUNT --admin-username $ADMINUSERNAME \  
--ssh-key-value $SSHKEYVALUE --node-vm-size $NODEVMsize
```

This takes a few minutes. When it is completed, fetch your **kubectl** credentials. The default behavior for **az aks get-credentials** is to merge the new credentials with the existing default configuration, and to set the new credentials as the current Kubernetes context. You should first backup your current configuration, or move it to a different location, then fetch the new credentials:

```
tux > az aks get-credentials --resource-group $RGNAME --name $AKSNAME
```

```
Merged "cap-aks" as current context in /home/tux/.kube/config
```

Verify that you can connect to your cluster:

```
tux > kubectl get nodes
NAME                                STATUS    ROLES    AGE     VERSION
aks-nodepool1-47788232-0           Ready    agent    5m      v1.9.6
aks-nodepool1-47788232-1           Ready    agent    6m      v1.9.6
aks-nodepool1-47788232-2           Ready    agent    6m      v1.9.6

tux > kubectl get pods --all-namespaces
NAMESPACE   NAME                                READY   STATUS    RESTARTS   AGE
kube-system  azureproxy-79c5db744-fwqcx         1/1     Running   2           6m
kube-system  heapster-55f855b47-c4mf9           2/2     Running   0           5m
kube-system  kube-dns-v20-7c556f89c5-spgbf      3/3     Running   0           6m
kube-system  kube-dns-v20-7c556f89c5-z2g7b      3/3     Running   0           6m
kube-system  kube-proxy-g9zpk                    1/1     Running   0           6m
kube-system  kube-proxy-kph4v                    1/1     Running   0           6m
kube-system  kube-proxy-xfngh                     1/1     Running   0           6m
kube-system  kube-svc-redirect-2knsj             1/1     Running   0           6m
kube-system  kube-svc-redirect-5nz2p             1/1     Running   0           6m
kube-system  kube-svc-redirect-hlh22             1/1     Running   0           6m
kube-system  kubernetes-dashboard-546686-mr9hz   1/1     Running   1           6m
kube-system  tunnelfront-595565bc78-j8msn       1/1     Running   0           6m
```

When all nodes are in a ready state and all pods are running, proceed to the next steps.

5.3 Enable Swap Accounting

Identify and set the cluster resource group, then enable kernel swap accounting. Swap accounting is required by CAP, but it is not the default in AKS nodes. The following commands use the `az` command to modify the GRUB configuration on each node, and then reboot the virtual machines.

1.

```
tux > export MCRGNAME=$(az group list -o table | grep MC_"$RGNAME"_ | awk '{print $1}')
```
2.

```
tux > vmnodes=$(az vm list -g $MCRGNAME | jq -r '[] | select (.tags.poolName | contains("node")) | .name')
```
3.

```
tux > for i in $vmnodes
do
  az vm run-command invoke -g $MCRGNAME -n $i --command-id RunShellScript \
```



```
--scripts "sudo sed -i 's|linux.*./boot/vmlinuz-.*|& swapaccount=1|' /boot/grub/
grub.cfg"
done
```

```
4. tux > for i in $vmnodes
do
    az vm restart -g $MCRGNAME -n $i
done
```

When this runs correctly, you will see multiple `"status": "Succeeded"` messages for all of your virtual machines.

5.4 Create a Basic Load Balancer and Public IP Address

Azure offers two load balancers, Basic and Standard. Currently Basic is free, while you have to pay for Standard. (See [Load Balancer \(https://azure.microsoft.com/en-us/services/load-balancer/\)](https://azure.microsoft.com/en-us/services/load-balancer/).) The following steps create a Basic load balancer (Basic is the default.) Look for `"provisioningState": "Succeeded"` messages in the command output to verify that the commands succeeded.

1. Create a static public IPv4 address:

```
tux > az network public-ip create \
--resource-group $MCRGNAME \
--name $AKSNAME-public-ip \
--allocation-method Static
```

2. Create the load balancer:

```
tux > az network lb create \
--resource-group $MCRGNAME \
--name $AKSNAME-lb \
--public-ip-address $AKSNAME-public-ip \
--frontend-ip-name $AKSNAME-lb-front \
--backend-pool-name $AKSNAME-lb-back
```

3. Set the virtual machine network interfaces, then add them to the load balancer:

```
tux > NICNAMES=$(az network nic list --resource-group $MCRGNAME | jq -r '[][.name]')
tux > for i in $NICNAMES
```

```
do
  az network nic ip-config address-pool add \
    --resource-group $MCRGNAME \
    --nic-name $i \
    --ip-config-name ipconfig1 \
    --lb-name $AKSNAME-lb \
    --address-pool $AKSNAME-lb-back
done
```

5.5 Configure Load Balancing and Network Security Rules

1. Set the required ports to allow access to SUSE Cloud Application Platform. Port 8443 is optional for the Stratos Web Console.

```
tux > export CAPPORTS="80 443 4443 2222 2793 8443"
```

2. Create network and load balancer rules:

```
tux > for i in $CAPPORTS
do
  az network lb probe create \
    --resource-group $MCRGNAME \
    --lb-name $AKSNAME-lb \
    --name probe-$i \
    --protocol tcp \
    --port $i

  az network lb rule create \
    --resource-group $MCRGNAME \
    --lb-name $AKSNAME-lb \
    --name rule-$i \
    --protocol Tcp \
    --frontend-ip-name $AKSNAME-lb-front \
    --backend-pool-name $AKSNAME-lb-back \
    --frontend-port $i \
    --backend-port $i \
    --probe probe-$i
done
```

3. Verify port setup:

```
tux > az network lb rule list -g $MCRGNAME --lb-name $AKSNAME-lb|grep -i port
```

```

"backendPort": 8443,
"frontendPort": 8443,
"backendPort": 80,
"frontendPort": 80,
"backendPort": 443,
"frontendPort": 443,
"backendPort": 4443,
"frontendPort": 4443,
"backendPort": 2222,
"frontendPort": 2222,
"backendPort": 2793,
"frontendPort": 2793,

```

4. Set the network security group name and priority level. The priority levels range from 100-4096, with 100 the highest priority. Each rule must have a unique priority level:

```

tux > nsg=$(az network nsg list --resource-group=$MCRGNAME | jq -r '[][.name]')
tux > pri=200

```

5. Create the network security rule:

```

tux > for i in $CAPPORTS
do
  az network nsg rule create \
  --resource-group $MCRGNAME \
  --priority $pri \
  --nsg-name $nsg \
  --name $AKSNAME-$i \
  --direction Inbound \
  --destination-port-ranges $i \
  --access Allow
  pri=$(expr $pri + 1)
done

```

6. Print the public and private IP addresses for later use:

```

tux > echo -e "\n Resource Group:\t$RGNAME\n \
Public IP:\t\t$(az network public-ip show --resource-group $MCRGNAME --name
$AKSNAME-public-ip --query ipAddress)\n \
Private IPs:\t\t\"$(az network nic list --resource-group $MCRGNAME | jq -r '
[][.ipConfigurations[].privateIpAddress' | paste -s -d " " | sed -e 's/ /", "/
g')\""\n"

```

```

Resource Group:      cap-aks
Public IP:           "40.101.3.25"

```

```
Private IPs: "10.240.0.4", "10.240.0.6", "10.240.0.5"
```

5.6 Example SUSE Cloud Application Platform Configuration File

The following example `scf-config-values.yaml` contains parameters particular to running SUSE Cloud Application Platform on Azure Kubernetes Service. You need the IP addresses from the last command in the previous section. This is a simplified example that does not use Azure's DNS services. For quick testing and proof of concept, you can use the free wildcard DNS services, xip.io (<http://xip.io/>) or nip.io (<http://nip.io/>). See [Azure DNS Documentation](https://docs.microsoft.com/en-us/azure/dns/) (<https://docs.microsoft.com/en-us/azure/dns/>) to learn more about Azure's name services.



Warning: Do not use xip.io or nip.io on production systems

Never use xip.io or nip.io on production systems! You must provide proper DNS and DHCP services on production clusters.

```
secrets:
  # Password for user 'admin' in the cluster
  CLUSTER_ADMIN_PASSWORD: password

  # Password for SCF to authenticate with UAA
  UAA_ADMIN_CLIENT_SECRET: password

env:
  # Use the public IP address
  DOMAIN: 40.101.3.25.xip.io

  # uaa prefix is required
  UAA_HOST: uaa.40.101.3.25.xip.io
  UAA_PORT: 2793

  #Azure deployment requires overlay
  GARDEN_ROOTFS_DRIVER: "overlay-xfv"

kube:
  # List the private IP addresses
  external_ips: ["10.240.0.5", "10.240.0.6", "10.240.0.4"]
  storage_class:
```

```
# Azure supports only "default" or "managed-premium"
persistent: "default"
shared: "shared"

registry:
  hostname: "registry.suse.com"
  username: ""
  password: ""
  organization: "cap"

auth: none
```

Now Azure is ready, and you can deploy SUSE Cloud Application Platform on it. Note that you will not install SUSE CaaS Platform, which provides a Kubernetes cluster, because AKS provides a managed Kubernetes cluster. Start with the "Helm Init" sections of the [Chapter 2, Production Installation with Optional High Availability](#) or [Chapter 8, Minimal Installation for Testing](#) guides.

When your UAA deployment has completed, test that it is operating correctly by running `curl` on the DNS name that you configured for your UAA_HOST:

```
tux > curl -k https://uaa.40.101.3.25.xip.io:2793/.well-known/openid-configuration
```

This should return a JSON object, as this abbreviated example shows:

```
{"issuer": "https://uaa.40.101.3.25.xip.io:2793/oauth/token",
"authorization_endpoint": "https://uaa.40.101.3.25.xip.io:2793/
/oauth/authorize", "token_endpoint": "https://uaa.40.101.3.25.
xip.io:2793/oauth/token",
```

6 Installing SUSE Cloud Application Platform on OpenStack

You can deploy a SUSE Cloud Application Platform on CaaS Platform stack on OpenStack. This chapter describes how to deploy a small testing and development instance with one Kubernetes master and two worker nodes, using [Terraform \(https://www.terraform.io/\)](https://www.terraform.io/) to automate the deployment. This does not create a production deployment, which should be deployed on bare metal for best performance.

6.1 Prerequisites

The following prerequisites should be met before attempting to deploy SUSE Cloud Application Platform on OpenStack. The memory and disk space requirements are minimums, and may need to be larger according to your workloads.

- 8GB of memory per CaaS Platform dashboard and Kubernetes master nodes
- 16GB of memory per Kubernetes worker
- 40GB disk space per CaaS Platform dashboard and Kubernetes master nodes
- 60GB disk space per Kubernetes worker
- A [SUSE Customer Center account \(https://scc.suse.com/login\)](https://scc.suse.com/login) for downloading CaaS Platform. Get `SUSE-CaaS-Platform-2.0-KVM-and-Xen.x86_64-1.0.0-GM.qcow2`, which has been tested on OpenStack.
- Download the `openrc.sh` file for your OpenStack account

6.2 Create a New OpenStack Project

You may use an existing OpenStack project, or run the following commands to create a new project with the necessary configuration for SUSE Cloud Application Platform.

```
tux > openstack project create --domain default --description "CaaS Platform Project" caasp
tux > openstack role add --project caasp --user admin admin
```

Create an OpenStack network plus a subnet for CaaS Platform (for example, *caasp-net*), and add a router to the external (e.g. floating) network:

```
tux > openstack network create caasp-net
tux > openstack subnet create caasp_subnet --network caasp-net \
--subnet-range 10.0.2.0/24
tux > openstack router create caasp-net-router
tux > openstack router set caasp-net-router --external-gateway floating
tux > openstack router add subnet caasp-net-router caasp_subnet
```

Upload your CaaS Platform image to your OpenStack account:

```
tux >
$ openstack image create \
--file SUSE-CaaS-Platform-2.0-KVM-and-Xen.x86_64-1.0.0-GM.qcow2
```

Create a security group with the rules needed for CaaS Platform:

```
tux > openstack security group create cap --description "Allow CAP traffic"
tux > openstack security group rule create cap --protocol any --dst-port any --ethertype
IPv4 --egress
tux > openstack security group rule create cap --protocol any --dst-port any --ethertype
IPv6 --egress
tux > openstack security group rule create cap --protocol tcp --dst-port 20000:20008 --
remote-ip 0.0.0.0/0
tux > openstack security group rule create cap --protocol tcp --dst-port 443:443 --
remote-ip 0.0.0.0/0
tux > openstack security group rule create cap --protocol tcp --dst-port 2793:2793 --
remote-ip 0.0.0.0/0
tux > openstack security group rule create cap --protocol tcp --dst-port 4443:4443 --
remote-ip 0.0.0.0/0
tux > openstack security group rule create cap --protocol tcp --dst-port 80:80 --remote-
ip 0.0.0.0/0
tux > openstack security group rule create cap --protocol tcp --dst-port 2222:2222 --
remote-ip 0.0.0.0/0
```

Clone the Terraform script from GitHub:

```
tux > git clone git@github.com:kubic-project/automation.git
tux > cd automation/caasp-openstack-terraform
```

Edit the `openstack.tfvars` file. Use the names of your OpenStack objects, for example:

```
image_name = "SUSE-CaaS-Platform-2.0"
internal_net = "caasp-net"
external_net = "floating"
```

```
admin_size = "m1.large"
master_size = "m1.large"
masters = 1
worker_size = "m1.xlarge"
workers = 2
```

Initialize Terraform:

```
tux > terraform init
```

6.3 Deploy SUSE Cloud Application Platform

Source your `openrc.sh` file, set the project, and deploy CaaS Platform:

```
tux > . openrc.sh
tux > export OS_PROJECT_NAME='caasp'
tux > ./caasp-openstack apply
```

Wait for a few minutes until all systems are up and running, then view your installation:

```
tux > openstack server list
```

Add your cap security group to all CaaS Platform workers:

```
tux > openstack server add security group caasp-worker0 cap
tux > openstack server add security group caasp-worker1 cap
```

If you need to log into your new nodes, log in as root using the SSH key in the `automation/caasp-openstack-terraform/ssh` directory.

6.4 Bootstrap SUSE Cloud Application Platform

The following examples use the `xip.io` wildcard DNS service. You may use your own DNS/DHCP services that you have set up in OpenStack in place of `xip.io`.

- Point your browser to the IP address of the CaaS Platform admin node, and create a new admin user login
- Replace the default IP address or domain name of the *Internal Dashboard FQDN/IP* on the *Initial CaaS Platform configuration* screen with the internal IP address of the CaaS Platform admin node

- Check the *Install Tiller* checkbox, then click the *Next* button
- Terraform automatically creates all of your worker nodes, according to the number you configured in `openstack.tfvars`, so click *Next* to skip *Bootstrap your CaaS Platform*
- On the *Select nodes and roles* screen click *Accept all nodes*, click to define your master and worker nodes, then click *Next*
- For the *External Kubernetes API FQDN*, use the public (floating) IP address of the CaaS Platform master and append the `.xip.io` domain suffix
- For the *External Dashboard FQDN* use the public (floating) IP address of the CaaS Platform admin node, and append the `.xip.io` domain suffix

6.5 Growing the Root Filesystem

If the root filesystem on your worker nodes is smaller than the OpenStack virtual disk, use these commands on the worker nodes to grow the filesystems to match:

```
tux > growpart /dev/vda 3
tux > btrfs filesystem resize max /.snapshots
```

7 Running SUSE Cloud Application Platform on non-SUSE CaaS Platform Kubernetes Systems

SUSE Cloud Application Platform is designed to run on any Kubernetes system that meets the following requirements:

- Kubernetes API version 1.8 +
- Kernel parameter `swapaccount=1`
- `docker info` must not show `aufs` as the storage driver
- `kube-dns` must be running
- Either `ntp` or `systemd-timesyncd` must be installed and active
- The Kubernetes cluster must have a storage class for SUSE Cloud Application Platform to use
- Docker must be configured to allow privileged containers
- Privileged container must be enabled in `kube-apiserver`. See [kube-apiserver \(https://kubernetes.io/docs/admin/kube-apiserver\)](https://kubernetes.io/docs/admin/kube-apiserver/).
- Privileged must be enabled in `kubelet`
- The `TasksMax` property of the `containerd` service definition must be set to infinity
- Helm's Tiller has to be installed and active, with Tiller on the Kubernetes cluster and Helm on your remote administration machine

8 Minimal Installation for Testing

A production deployment of SUSE Cloud Application Platform requires a significant number of physical or virtual hosts. For testing and learning, you can set up a minimal four-host deployment of SUSE Cloud Foundry on SUSE CaaS Platform on a single workstation in a hypervisor such as KVM or VirtualBox. This extremely minimal deployment uses Kubernetes' hostpath storage type instead of a storage server, such as SUSE Enterprise Storage. You must also provide DNS, DHCP, and a network space for your cluster. KVM and VirtualBox include name services and network management. *Figure 8.1, "Minimal Network Architecture"* illustrates the layout of a physical minimal test installation with an external administration workstation and DNS/DHCP server. Access to the cluster is provided by the UAA (User Account and Authentication) server on worker 1.

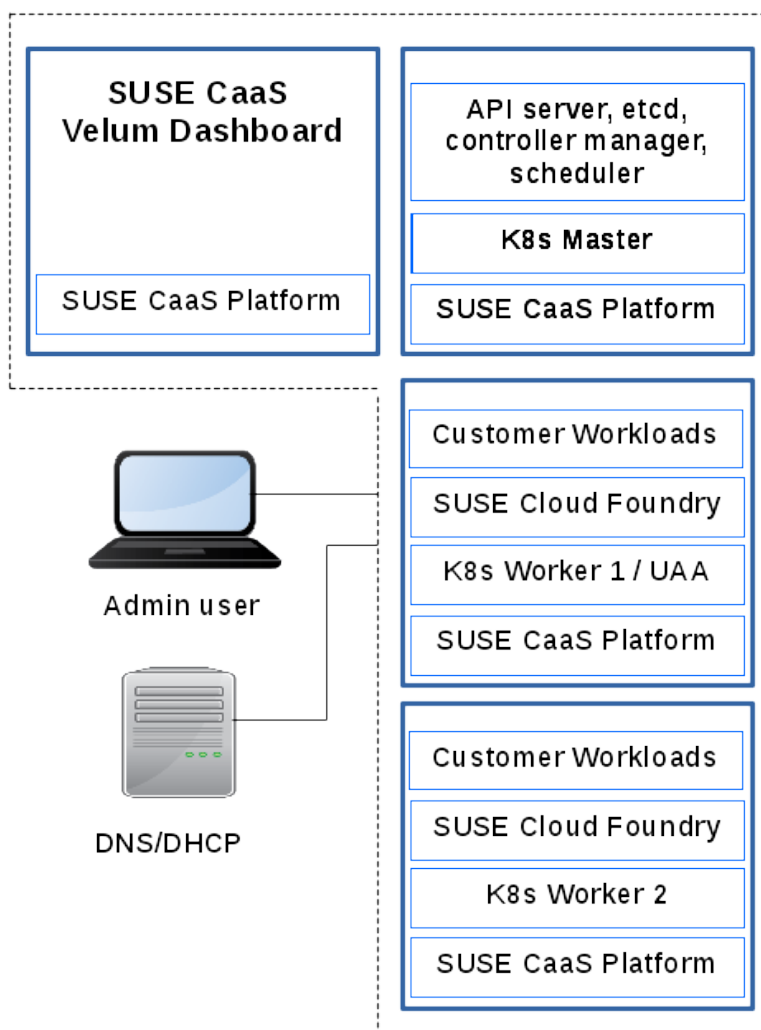


FIGURE 8.1: MINIMAL NETWORK ARCHITECTURE

This minimal four-node deployment will run on a minimum of 32GB host system memory, though more memory is better. 32GB is enough to test setting up and configuring SUSE CaaS Platform and SUSE Cloud Foundry, and to run a few lightweight workloads. You may also test connecting external servers with your cluster, such as a separate name server, a storage server (e.g. SUSE Enterprise Storage), SUSE Customer Center, or Subscription Management Tool. You must be familiar with installing and configuring CaaS Platform (see the [SUSE CaaS Platform 2 Deployment Guide \(https://www.suse.com/documentation/suse-caasp-2/book_caasp_deployment/index.html\)](https://www.suse.com/documentation/suse-caasp-2/book_caasp_deployment/index.html)).

After you have installed CaaS Platform you will install and administer SUSE Cloud Foundry remotely from your host workstation, using tools such as the [Helm package manager for Kubernetes \(https://docs.helm.sh/\)](https://docs.helm.sh/), and the Kubernetes command-line tool `kubectl` (<https://kubernetes.io/docs/tasks/tools/install-kubectl/kubectl/>).



Warning: Limitations of minimal test environment

This is a limited deployment that is useful for testing basic deployment and functionality, but it is NOT a production system, and cannot be upgraded to a production system. Its reduced complexity allows basic testing, it is portable (on laptops with enough memory), and is useful in environments that have resource constraints.

8.1 Prerequisites



Important: You must be familiar with SUSE CaaS Platform

Setting up SUSE CaaS Platform correctly, and knowledge of basic administration is essential to a successful SUSE Cloud Application Platform deployment. See the [SUSE CaaS Platform 2 Deployment Guide \(https://www.suse.com/documentation/suse-caasp-2/book_caasp_deployment/index.html\)](https://www.suse.com/documentation/suse-caasp-2/book_caasp_deployment/index.html)

CaaS Platform requires a minimum of four physical or virtual hosts: one admin, one Kubernetes master, and two Kubernetes workers. You also need an Internet connection, as the installer has an option to download updates during installation, and the Kubernetes workers will each download ~10GB of Docker images.

Hardware requirements

Any AMD64/Intel EM64T processor with at least 8 virtual or physical cores. This table describes the minimum requirements per node.

Node	CPU	RAM	Disk
CaaS Platform Dashboard	1	8GB	40GB
CaaS Platform Master	2	8GB	40GB
CaaS Platform Workers	2	16GB	60GB

Network and Name Services

You must provide DNS and DHCP services, either via your hypervisor, or with a separate name server. Your cluster needs its own domain. Every node needs a hostname and a fully-qualified domain name, and should all be on the same network. By default, the CaaS Platform installer requests a hostname from any available DHCP server. When you install the admin server you may adjust its network settings manually, and should give it a hostname, a static IP address, and specify which name server to use if there is more than one. CaaS Platform supports multiple methods for installing the Kubernetes workers. We recommend using AutoYaST, and then when you deploy the Kubernetes workers you will create their hostnames with a kernel boot option.

After your Kubernetes nodes are running select one Kubernetes worker to act as the external access point for your cluster and map your domain name to it. On production clusters it is a common practice to use wildcard DNS, rather than trying to manage DNS for hundreds or thousands of applications. Map your domain wildcard to the IP address of the Kubernetes worker you selected as the external access point to your cluster.

Install SUSE CaaS Platform 2

Install SUSE CaaS Platform 2 [CaaS Platform \(https://www.suse.com/documentation/suse-caasp-2/\)](https://www.suse.com/documentation/suse-caasp-2/). When you reach the step where you log into the Velum Web interface, check the box to install Tiller (Helm's server component).

SUSE® CaaS Platform Logout

Welcome! You have signed up successfully. ×

Initial CaaS Platform Configuration

Generic settings

Internal Dashboard FQDN/IP
caasp-admin.example.com i

Cluster services

Install Tiller (Helm's server component)

Overlay network settings Show

Proxy settings Enable Disable

Next

FIGURE 8.2: INSTALL TILLER

Take note of the *Overlay network settings*. These define the cluster and services networks that are exclusive to the internal cluster communications. They are not accessible outside of the cluster. You may change the default overlay network assignments to avoid address collisions with your existing network.

There is also a form for proxy settings; if you're not using a proxy then leave it empty.

The easiest way to create the Kubernetes nodes is to use AutoYaST see [Installation with AutoYaST \(https://www.suse.com/documentation/suse-caasp-2/book_caasp_deployment/data/sec_caasp_installquick.html#sec_caasp_installquick_node_ay\)](https://www.suse.com/documentation/suse-caasp-2/book_caasp_deployment/data/sec_caasp_installquick.html#sec_caasp_installquick_node_ay). Pass in these kernel boot options to each worker: `hostname`, `netsetup`, and the AutoYaST path, which you find in Velum on the "Bootstrap your CaaS Platform" page.

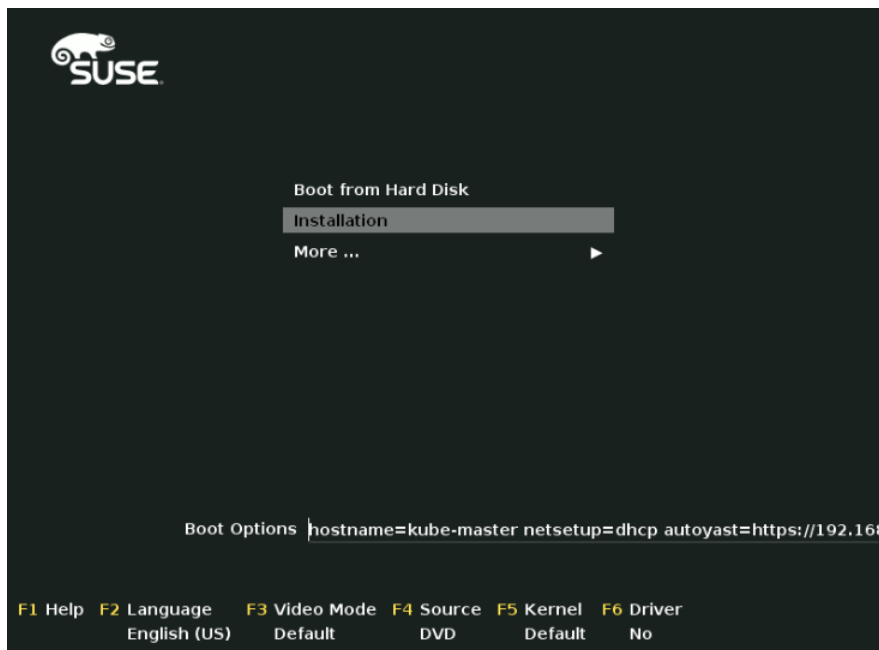


FIGURE 8.3: KERNEL BOOT OPTIONS

When you have completed [Bootstrapping the Cluster \(https://www.suse.com/documentation/suse-caasp-2/book_caasp_deployment/data/sec_caasp_installquick.html#sec_caasp_installquick_bootstrap\)](https://www.suse.com/documentation/suse-caasp-2/book_caasp_deployment/data/sec_caasp_installquick.html#sec_caasp_installquick_bootstrap) open a Web browser to the Velum Web interface. If you see a "site not available" or "We're sorry, but something went wrong" error wait a few minutes, then try again. Click the *kubectl config* button to download your new cluster's kubeconfig file. This takes you to a login screen; use the login you created to access Velum. Save the file as ~/.kube/config on your host workstation. This file enables the remote administration of your cluster.

Cluster Status

Summary

Total nodes 3**Master nodes** 1**New nodes** ⓘ 0 (new)**Updates** Manual**# of nodes w/ outdated software** 0

Nodes

 **kubectl config**

FIGURE 8.4: DOWNLOAD KUBECONFIG

Install kubectl



Note: Remote Cluster Administration

You will administer your cluster from your host workstation, rather than directly on any of your cluster nodes. The remote environment is indicated by the unprivileged user Tux, while root prompts are on a cluster host. There are few tasks that need to be performed directly on any of the cluster hosts.

Follow the instructions at [Install and Set Up kubectl \(https://kubernetes.io/docs/tasks/tools/install-kubectl/\)](https://kubernetes.io/docs/tasks/tools/install-kubectl/) to install **kubectl** on your host workstation. After installation, run this command to verify that it is installed, and that it is communicating correctly with your cluster:

```
tux > kubectl version --short
Client Version: v1.9.1
Server Version: v1.7.7
```

As the client is on your workstation, and the server is on your cluster, reporting the server version verifies that **kubectl** is using `~/.kube/config` and is communicating with your cluster.

The following examples query the cluster configuration and node status:

```
tux > kubectl config view
apiVersion: v1
clusters:
- cluster:
  certificate-authority-data: REDACTED
  server: https://192.168.10.101:6443
  name: local
contexts:
[...]
```



```
tux > kubectl get nodes
```

NAME	STATUS	ROLES	AGE	VERSION
4a10db2c.infra.caasp.local	Ready	<none>	4h	v1.7.7
87c9e8ff.infra.caasp.local	Ready,SchedulingDisabled	<none>	4h	v1.7.7
34ce7eb0.infra.caasp.local	Ready	<none>	4h	v1.7.7

Install Helm

Deploying SUSE Cloud Foundry is different than the usual method of installing software. Rather than installing packages in the usual way with YaST or Zypper, you will install the Helm client on your workstation to install the required Kubernetes applications to set up SUSE Cloud Foundry, and to administer your cluster remotely.

Helm client version 2.6 or higher is required.



Warning: Initialize Only the Helm Client

When you initialize Helm on your workstation be sure to initialize only the client, as the server, Tiller, was installed during the CaaS Platform installation. You do not want two Tiller instances.

If the Linux distribution on your workstation doesn't provide the correct Helm version, or you are using some other platform, see the [Helm Quickstart Guide \(https://docs.helm.sh/using_helm/#quickstart\)](https://docs.helm.sh/using_helm/#quickstart) for installation instructions and basic usage examples. Download the Helm binary into any directory that is in your PATH on your workstation, such as your `~/bin` directory. Then initialize the client only:

```
tux > helm init --client-only
Creating /home/tux/.helm
Creating /home/tux/.helm/repository
Creating /home/tux/.helm/repository/cache
Creating /home/tux/.helm/repository/local
Creating /home/tux/.helm/plugins
Creating /home/tux/.helm/starters
Creating /home/tux/.helm/cache/archive
Creating /home/tux/.helm/repository/repositories.yaml
Adding stable repo with URL: https://kubernetes-charts.storage.googleapis.com
Adding local repo with URL: http://127.0.0.1:8879/charts
$HELM_HOME has been configured at /home/tux/.helm.
Not installing Tiller due to 'client-only' flag having been set
Happy Helming!
```

8.2 Create hostpath Storage Class

The Kubernetes cluster requires a persistent storage class for the databases to store persistent data. You can provide this with your own storage (e.g. SUSE Enterprise Storage), or use the built-in hostpath storage type. hostpath is NOT suitable for a production deployment, but it is an easy option for a minimal test deployment.



Warning: Using the hostpath storage type on CaaS Platform

CaaS Platform is configured as a multi-node Kubernetes setup with a minimum of one master and two workers. Hostpath provisioning on CaaS Platform uses local storage on each of these nodes, therefore persistent data stored will only be available locally on the Kubernetes nodes. This impacts use cases where SUSE Cloud Foundry containers restart on a different Kubernetes worker, for example in high availability setups or update tests. If a container starts on a different worker than before it will miss its persistent data,

leading to various other side effects. In addition, hostpath-provisioner uses the local root filesystem of the Kubernetes node. If it runs out of disk space your Kubernetes node won't work anymore.

Open an SSH session to your Kubernetes master node and add the argument `--enable-hostpath-provisioner` to `/etc/kubernetes/controller-manager`:

```
root # vim /etc/kubernetes/controller-manager
      KUBE_CONTROLLER_MANAGER_ARGS="\
      --enable-hostpath-provisioner \
      "
```

Restart the Kubernetes controller-manager:

```
root # systemctl restart kube-controller-manager
```

Create a persistent storage class named hostpath:

```
root # echo '{"kind":"StorageClass","apiVersion":"storage.k8s.io/v1", "metadata":
{"name":"hostpath"},"provisioner":"kubernetes.io/host-path"}' | \
kubectl create -f -


storageclass "hostpath" created
```

Verify that your new storage class has been created:

```
root # kubectl get storageclass
NAME          TYPE
hostpath     kubernetes.io/host-path
```

Log into all of your Kubernetes nodes and create the `/tmp/hostpath_pv` directory, then set its permissions to read/write/execute:

```
root # mkdir /tmp/hostpath_pv
root # chmod -R 0777 /tmp/hostpath_pv
```

See the Kubernetes document [Storage Classes \(https://kubernetes.io/docs/concepts/storage/storage-classes/\)](https://kubernetes.io/docs/concepts/storage/storage-classes/)  for detailed information on storage classes.



Tip: Log in Directly to Kubernetes Nodes

By default, SUSE CaaS Platform allows logging into the Kubernetes nodes only from the admin node. You can set up direct logins to your Kubernetes nodes from your workstation by copying the SSH keys from your admin node to your Kubernetes nodes, and then you will have password-less SSH logins. This is not a best practice for a production deployment, but will make running a test deployment a little easier.

8.3 Test Storage Class

See [Section 2.3, "Test Storage Class"](#) to learn how to test that your storage class is correctly configured before you deploy SUSE Cloud Foundry.

8.4 Configuring the Minimal Test Deployment

Create a configuration file on your workstation for Helm to use. In this example it is called `scf-config-values.yaml`. (See the [Release Notes \(https://www.suse.com/releasesnotes/\)](https://www.suse.com/releasesnotes/) for information on configuration changes.)

```
env:
  # Enter the domain you created for your CAP cluster
  DOMAIN: example.com

  # UAA host and port
  UAA_HOST: uaa.example.com
  UAA_PORT: 2793

kube:
  # # The IP address assigned to the kube node pointed to by the domain.
  external_ips: ["192.168.10.101"]

  # Run kubectl get storageclasses
  # to view your available storage classes
  storage_class:
    persistent: "hostpath"
    shared: "shared"

  # The registry the images will be fetched from.
  # The values below should work for
  # a default installation from the SUSE registry.
```

```
registry:
  hostname: "registry.suse.com"
  username: ""
  password: ""
  organization: "cap"

# Required for CaaSP 2
auth: rbac

secrets:
# Create a password for your CAP cluster
CLUSTER_ADMIN_PASSWORD: password

# Create a password for your UAA client secret
UAA_ADMIN_CLIENT_SECRET: password
```

8.5 Deploy with Helm

Run the following Helm commands to complete the deployment. There are six steps, and they must be run in this order:

- Download the SUSE Kubernetes charts repository
- Create namespaces
- If you are using SUSE Enterprise Storage, copy the storage secret to the UAA and SCF namespaces
- Install UAA
- Copy UAA secret and certificate to SCF namespace
- Install SCF

8.5.1 Install the Kubernetes charts repository

Download the SUSE Kubernetes charts repository with Helm:

```
tux > helm repo add suse https://kubernetes-charts.suse.com/
```

You may replace the example `suse` name with any name. Verify with `helm`:

```
tux > helm repo list
```

NAME	URL
stable	https://kubernetes-charts.storage.googleapis.com
local	http://127.0.0.1:8879/charts
suse	https://kubernetes-charts.suse.com/

List your chart names, as you will need these for some operations:

```
tux > helm search suse
NAME          VERSION  DESCRIPTION
suse/cf       2.8.0    A Helm chart for SUSE Cloud Foundry
suse/console  1.1.0    A Helm chart for deploying Stratos UI Console
suse/uaa      2.8.0    A Helm chart for SUSE UAA
```

8.5.2 Create Namespaces

Use **kubectl** on your host workstation to create and verify the UAA (User Account and Authentication) and SCF (SUSE Cloud Foundry) namespaces:

```
tux > kubectl create namespace uaa
namespace "uaa" created

tux > kubectl create namespace scf
namespace "scf" created

tux > kubectl get namespaces
NAME          STATUS  AGE
default       Active  27m
kube-public   Active  27m
kube-system   Active  27m
scf           Active  1m
uaa           Active  1m
```

8.5.3 Copy SUSE Enterprise Storage Secret

If you are using the `hostpath` storage class (see [Section 8.2, "Create hostpath Storage Class"](#)) there is no secret so skip this step.

If you are using SUSE Enterprise Storage you must copy the Ceph admin secret to the UAA and SCF namespaces:

```
tux > kubectl get secret ceph-secret-admin -o json --namespace default | \
sed 's/"namespace": "default"/"namespace": "uaa"/' | kubectl create -f -
```

```
tux > kubectl get secret ceph-secret-admin -o json --namespace default | \
sed 's/"namespace": "default"/"namespace": "scf"/' | kubectl create -f -
```

8.5.4 Install UAA

Use Helm to install the UAA (User Account and Authentication) server:

```
tux > helm install suse/uaa \
--name susecf-uaa \
--namespace uaa \
--values scf-config-values.yaml
```

Wait until you have a successful UAA deployment before going to the next steps, which you can monitor with the `watch` command. This will take time, possibly an hour or two, according to your hardware resources:

```
tux > watch -c 'kubectl get pods --all-namespaces'
```

When the status shows RUNNING for all of the UAA nodes, then proceed to the next step.

8.5.5 Install SUSE Cloud Foundry

First pass your UAA secret and certificate to SCF, then use Helm to install SUSE Cloud Foundry:

```
tux > SECRET=$(kubectl get pods --namespace uaa \
-o jsonpath='{.items[*].spec.containers[?(.name=="uaa")].env[?
(.name=="INTERNAL_CA_CERT")].valueFrom.secretKeyRef.name}')

tux > CA_CERT="$(kubectl get secret $SECRET --namespace uaa \
-o jsonpath="{.data['internal-ca-cert']}" | base64 --decode -)"

tux > helm install suse/cf \
--name susecf-scf \
--namespace scf \
--values scf-config-values.yaml \
--set "secrets.UAA_CA_CERT=${CA_CERT}"
```

Now sit back and wait for the pods come online:

```
tux > watch -c 'kubectl get pods --all-namespaces'
```

When all services are running you can use the Cloud Foundry command-line interface to log in to SUSE Cloud Foundry. (See [Section 2.10, “Deploy SUSE Cloud Foundry”](#).)

8.6 Install the Stratos Console

Stratos UI is a modern, web-based management application for Cloud Foundry. It provides a graphical management console for both developers and system administrators. (See [Section 2.12, “Installing the Stratos Web Console”](#)).

8.7 Updating SUSE Cloud Foundry, UAA, and Stratos

Maintenance updates are delivered as container images from the SUSE registry and applied with Helm. See [Section 2.13, “Upgrading SUSE Cloud Foundry, UAA, and Stratos”](#).



Note: No Upgrades with Hostpath

Upgrades do not work with the hostpath storage type, as the required stateful data may be lost.

9 Troubleshooting

Cloud stacks are complex, and debugging deployment issues often requires digging through multiple layers to find the information you need. Remember that the SUSE Cloud Foundry releases must be deployed in the correct order, and that each release must deploy successfully, with no failed pods, before deploying the next release.

9.1 Using Supportconfig

If you ever need to request support, or just want to generate detailed system information and logs, use the **supportconfig** utility. Run it with no options to collect basic system information, and also cluster logs including Docker, etcd, flannel, and Velum. **supportconfig** may give you all the information you need.

supportconfig -h prints the options. Read the "Gathering System Information for Support" chapter in any SUSE Linux Enterprise Administration Guide to learn more.

9.2 Deployment is Taking Too Long

A deployment step seems to take too long, or you see that some pods are not in a ready state hours after all the others are ready, or a pod shows a lot of restarts. This example shows not-ready pods many hours after the others have become ready:

```
tux > kubectl get pods --namespace scf
NAME                                READY STATUS    RESTARTS  AGE
router-3137013061-wlhxb            0/1   Running    0         16h
routing-api-0                       0/1   Running    0         16h
```

The Running status means the pod is bound to a node and all of its containers have been created. However, it is not Ready, which means it is not ready to service requests. Use **kubectl** to print a detailed description of pod events and status:

```
tux > kubectl describe pod --namespace scf router-3137013061-wlhxb
```

This prints a lot of information, including IP addresses, routine events, warnings, and errors. You should find the reason for the failure in this output.

9.3 Deleting and Rebuilding a Deployment

There may be times when you want to delete and rebuild a deployment, for example when there are errors in your `scf-config-values.yaml` file, you wish to test configuration changes, or a deployment fails and you want to try it again. This has four steps: first delete the release or releases you want to re-deploy, delete its namespace, then re-create the namespace and re-deploy the release.

Use **helm** to see your releases:

```
tux > helm ls
NAME                REVISION  UPDATED              STATUS  CHART          NAMESPACE
susecf-console      1         Thu Apr 12 10:28:34 2018  DEPLOYED  console-1.1.0  stratos
susecf-scf          1         Wed Apr 11 14:55:23 2018  DEPLOYED  cf-2.8.0       scf
susecf-uaa          1         Wed Apr 11 14:48:01 2018  DEPLOYED  uaa-2.8.0      uaa
```

This example deletes the `susecf-console` release and namespace:

```
tux > helm delete susecf-console
release "susecf-console" deleted
tux > kubectl delete namespace stratos
namespace "stratos" deleted
```

Then you can start over.

9.4 Querying with Kubectl

You can safely query with **kubectl** to get information about resources inside your Kubernetes cluster. `kubectl cluster-info dump | tee clusterinfo.txt` outputs a large amount of information about the Kubernetes master and cluster services to a text file.

The following commands give more targeted information about your cluster.

- List all cluster resources:

```
tux > kubectl get all --all-namespaces
```

- List all of your running pods:

```
tux > kubectl get pods --all-namespaces
```

- See all pods, including those with Completed or Failed statuses:

```
tux > kubectl get pods --show-all --all-namespaces
```

- List pods in one namespace:

```
tux > kubectl get pods --namespace scf
```

- Get detailed information about one pod:

```
tux > kubectl describe --namespace scf po/diego-cell-0
```

- Read the log file of a pod:

```
tux > kubectl logs --namespace scf po/diego-cell-0
```

- List all Kubernetes nodes, then print detailed information about a single node:

```
tux > kubectl get nodes  
tux > kubectl describe node 6a2752b6fab54bb889029f60de6fa4d5.infra.caasp.local
```

- List all containers in all namespaces, formatted for readability:

```
tux > kubectl get pods --all-namespaces -o jsonpath="{..image}" |\ntr -s '[:space:]' '\\n' |\nsort |\nuniq -c
```

- These two commands check node capacities, to verify that there are enough resources for the pods:

```
tux > kubectl get nodes -o yaml | grep '\\sname\\|cpu\\|memory'  
tux > kubectl get nodes -o json | \  
jq '.items[] | {name: .metadata.name, cap: .status.capacity}'
```

A GNU Licenses

This appendix contains the GNU Free Documentation License version 1.2.

GNU Free Documentation License

Copyright (C) 2000, 2001, 2002 Free Software Foundation, Inc. 51 Franklin St, Fifth Floor, Boston, MA 02110-1301 USA. Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

0. PREAMBLE

The purpose of this License is to make a manual, textbook, or other functional and useful document "free" in the sense of freedom: to assure everyone the effective freedom to copy and redistribute it, with or without modifying it, either commercially or non-commercially. Secondly, this License preserves for the author and publisher a way to get credit for their work, while not being considered responsible for modifications made by others.

This License is a kind of "copyleft", which means that derivative works of the document must themselves be free in the same sense. It complements the GNU General Public License, which is a copyleft license designed for free software.

We have designed this License to use it for manuals for free software, because free software needs free documentation: a free program should come with manuals providing the same freedoms that the software does. But this License is not limited to software manuals; it can be used for any textual work, regardless of subject matter or whether it is published as a printed book. We recommend this License principally for works whose purpose is instruction or reference.

1. APPLICABILITY AND DEFINITIONS

This License applies to any manual or other work, in any medium, that contains a notice placed by the copyright holder saying it can be distributed under the terms of this License. Such a notice grants a world-wide, royalty-free license, unlimited in duration, to use that work under the conditions stated herein. The "Document", below, refers to any such manual or work. Any member of the public is a licensee, and is addressed as "you". You accept the license if you copy, modify or distribute the work in a way requiring permission under copyright law.

A "Modified Version" of the Document means any work containing the Document or a portion of it, either copied verbatim, or with modifications and/or translated into another language.

A "Secondary Section" is a named appendix or a front-matter section of the Document that deals exclusively with the relationship of the publishers or authors of the Document to the Document's overall subject (or to related matters) and contains nothing that could fall directly within that overall subject. (Thus, if the Document is in part a textbook of mathematics, a Secondary Section may not explain any mathematics.) The relationship could be a matter of historical connection with the subject or with related matters, or of legal, commercial, philosophical, ethical or political position regarding them.

The "Invariant Sections" are certain Secondary Sections whose titles are designated, as being those of Invariant Sections, in the notice that says that the Document is released under this License. If a section does not fit the above definition of Secondary then it is not allowed to be designated as Invariant. The Document may contain zero Invariant Sections. If the Document does not identify any Invariant Sections then there are none.

The "Cover Texts" are certain short passages of text that are listed, as Front-Cover Texts or Back-Cover Texts, in the notice that says that the Document is released under this License. A Front-Cover Text may be at most 5 words, and a Back-Cover Text may be at most 25 words.

A "Transparent" copy of the Document means a machine-readable copy, represented in a format whose specification is available to the general public, that is suitable for revising the document straightforwardly with generic text editors or (for images composed of pixels) generic paint programs or (for drawings) some widely available drawing editor, and that is suitable for input to text formatters or for automatic translation to a variety of formats suitable for input to text formatters. A copy made in an otherwise Transparent file format whose markup, or absence of markup, has been arranged to thwart or discourage subsequent modification by readers is not Transparent. An image format is not Transparent if used for any substantial amount of text. A copy that is not "Transparent" is called "Opaque".

Examples of suitable formats for Transparent copies include plain ASCII without markup, Texinfo input format, LaTeX input format, SGML or XML using a publicly available DTD, and standard-conforming simple HTML, PostScript or PDF designed for human modification. Examples of transparent image formats include PNG, XCF and JPG. Opaque formats include proprietary

formats that can be read and edited only by proprietary word processors, SGML or XML for which the DTD and/or processing tools are not generally available, and the machine-generated HTML, PostScript or PDF produced by some word processors for output purposes only.

The "Title Page" means, for a printed book, the title page itself, plus such following pages as are needed to hold, legibly, the material this License requires to appear in the title page. For works in formats which do not have any title page as such, "Title Page" means the text near the most prominent appearance of the work's title, preceding the beginning of the body of the text.

A section "Entitled XYZ" means a named subunit of the Document whose title either is precisely XYZ or contains XYZ in parentheses following text that translates XYZ in another language. (Here XYZ stands for a specific section name mentioned below, such as "Acknowledgements", "Dedications", "Endorsements", or "History".) To "Preserve the Title" of such a section when you modify the Document means that it remains a section "Entitled XYZ" according to this definition.

The Document may include Warranty Disclaimers next to the notice which states that this License applies to the Document. These Warranty Disclaimers are considered to be included by reference in this License, but only as regards disclaiming warranties: any other implication that these Warranty Disclaimers may have is void and has no effect on the meaning of this License.

2. VERBATIM COPYING

You may copy and distribute the Document in any medium, either commercially or non-commercially, provided that this License, the copyright notices, and the license notice saying this License applies to the Document are reproduced in all copies, and that you add no other conditions whatsoever to those of this License. You may not use technical measures to obstruct or control the reading or further copying of the copies you make or distribute. However, you may accept compensation in exchange for copies. If you distribute a large enough number of copies you must also follow the conditions in section 3.

You may also lend copies, under the same conditions stated above, and you may publicly display copies.

3. COPYING IN QUANTITY

If you publish printed copies (or copies in media that commonly have printed covers) of the Document, numbering more than 100, and the Document's license notice requires Cover Texts, you must enclose the copies in covers that carry, clearly and legibly, all these Cover Texts: Front-Cover Texts on the front cover, and Back-Cover Texts on the back cover. Both covers must also clearly and legibly identify you as the publisher of these copies. The front cover must present the full title with all words of the title equally prominent and visible. You may add other material on the covers in addition. Copying with changes limited to the covers, as long as they preserve the title of the Document and satisfy these conditions, can be treated as verbatim copying in other respects.

If the required texts for either cover are too voluminous to fit legibly, you should put the first ones listed (as many as fit reasonably) on the actual cover, and continue the rest onto adjacent pages.

If you publish or distribute Opaque copies of the Document numbering more than 100, you must either include a machine-readable Transparent copy along with each Opaque copy, or state in or with each Opaque copy a computer-network location from which the general network-using public has access to download using public-standard network protocols a complete Transparent copy of the Document, free of added material. If you use the latter option, you must take reasonably prudent steps, when you begin distribution of Opaque copies in quantity, to ensure that this Transparent copy will remain thus accessible at the stated location until at least one year after the last time you distribute an Opaque copy (directly or through your agents or retailers) of that edition to the public.

It is requested, but not required, that you contact the authors of the Document well before redistributing any large number of copies, to give them a chance to provide you with an updated version of the Document.

4. MODIFICATIONS

You may copy and distribute a Modified Version of the Document under the conditions of sections 2 and 3 above, provided that you release the Modified Version under precisely this License, with the Modified Version filling the role of the Document, thus licensing distribution and modification of the Modified Version to whoever possesses a copy of it. In addition, you must do these things in the Modified Version:

- A. Use in the Title Page (and on the covers, if any) a title distinct from that of the Document, and from those of previous versions (which should, if there were any, be listed in the History section of the Document). You may use the same title as a previous version if the original publisher of that version gives permission.
- B. List on the Title Page, as authors, one or more persons or entities responsible for authorship of the modifications in the Modified Version, together with at least five of the principal authors of the Document (all of its principal authors, if it has fewer than five), unless they release you from this requirement.
- C. State on the Title page the name of the publisher of the Modified Version, as the publisher.
- D. Preserve all the copyright notices of the Document.
- E. Add an appropriate copyright notice for your modifications adjacent to the other copyright notices.
- F. Include, immediately after the copyright notices, a license notice giving the public permission to use the Modified Version under the terms of this License, in the form shown in the Addendum below.
- G. Preserve in that license notice the full lists of Invariant Sections and required Cover Texts given in the Document's license notice.
- H. Include an unaltered copy of this License.
- I. Preserve the section Entitled "History", Preserve its Title, and add to it an item stating at least the title, year, new authors, and publisher of the Modified Version as given on the Title Page. If there is no section Entitled "History" in the Document, create one stating the title, year, authors, and publisher of the Document as given on its Title Page, then add an item describing the Modified Version as stated in the previous sentence.
- J. Preserve the network location, if any, given in the Document for public access to a Transparent copy of the Document, and likewise the network locations given in the Document for previous versions it was based on. These may be placed in the "History" section. You may omit a network location for a work that was published at least four years before the Document itself, or if the original publisher of the version it refers to gives permission.
- K. For any section Entitled "Acknowledgements" or "Dedications", Preserve the Title of the section, and preserve in the section all the substance and tone of each of the contributor acknowledgements and/or dedications given therein.
- L. Preserve all the Invariant Sections of the Document, unaltered in their text and in their titles. Section numbers or the equivalent are not considered part of the section titles.
- M. Delete any section Entitled "Endorsements". Such a section may not be included in the Modified Version.
- N. Do not retitle any existing section to be Entitled "Endorsements" or to conflict in title with any Invariant Section.
- O. Preserve any Warranty Disclaimers.

If the Modified Version includes new front-matter sections or appendices that qualify as Secondary Sections and contain no material copied from the Document, you may at your option designate some or all of these sections as invariant. To do this, add their titles to the list of Invariant Sections in the Modified Version's license notice. These titles must be distinct from any other section titles.

You may add a section Entitled "Endorsements", provided it contains nothing but endorsements of your Modified Version by various parties—for example, statements of peer review or that the text has been approved by an organization as the authoritative definition of a standard.

You may add a passage of up to five words as a Front-Cover Text, and a passage of up to 25 words as a Back-Cover Text, to the end of the list of Cover Texts in the Modified Version. Only one passage of Front-Cover Text and one of Back-Cover Text may be added by (or through arrangements made by) any one entity. If the Document already includes a cover text for the same cover, previously added by you or by arrangement made by the same entity you are acting on behalf of, you may not add another; but you may replace the old one, on explicit permission from the previous publisher that added the old one.

The author(s) and publisher(s) of the Document do not by this License give permission to use their names for publicity for or to assert or imply endorsement of any Modified Version.

5. COMBINING DOCUMENTS

You may combine the Document with other documents released under this License, under the terms defined in section 4 above for modified versions, provided that you include in the combination all of the Invariant Sections of all of the original documents, unmodified, and list them all as Invariant Sections of your combined work in its license notice, and that you preserve all their Warranty Disclaimers.

The combined work need only contain one copy of this License, and multiple identical Invariant Sections may be replaced with a single copy. If there are multiple Invariant Sections with the same name but different contents, make the title of each such section unique by adding at the end of it, in parentheses, the name of the original author or publisher of that section if known, or else a unique number. Make the same adjustment to the section titles in the list of Invariant Sections in the license notice of the combined work.

In the combination, you must combine any sections Entitled "History" in the various original documents, forming one section Entitled "History"; likewise combine any sections Entitled "Acknowledgements", and any sections Entitled "Dedications". You must delete all sections Entitled "Endorsements".

6. COLLECTIONS OF DOCUMENTS

You may make a collection consisting of the Document and other documents released under this License, and replace the individual copies of this License in the various documents with a single copy that is included in the collection, provided that you follow the rules of this License for verbatim copying of each of the documents in all other respects.

You may extract a single document from such a collection, and distribute it individually under this License, provided you insert a copy of this License into the extracted document, and follow this License in all other respects regarding verbatim copying of that document.

7. AGGREGATION WITH INDEPENDENT WORKS

A compilation of the Document or its derivatives with other separate and independent documents or works, in or on a volume of a storage or distribution medium, is called an "aggregate" if the copyright resulting from the compilation is not used to limit the legal rights of the compilation's users beyond what the individual works permit. When the Document is included in an aggregate, this License does not apply to the other works in the aggregate which are not themselves derivative works of the Document.

If the Cover Text requirement of section 3 is applicable to these copies of the Document, then if the Document is less than one half of the entire aggregate, the Document's Cover Texts may be placed on covers that bracket the Document within the aggregate, or the electronic equivalent of covers if the Document is in electronic form. Otherwise they must appear on printed covers that bracket the whole aggregate.

8. TRANSLATION

Translation is considered a kind of modification, so you may distribute translations of the Document under the terms of section 4. Replacing Invariant Sections with translations requires special permission from their copyright holders, but you may include translations of some or all Invariant Sections in addition to the original versions of these Invariant Sections. You may include a translation of this License, and all the license notices in the Document, and any Warranty Disclaimers, provided that you also include the original English version of this License and the original versions of those notices and disclaimers. In case of a disagreement between the translation and the original version of this License or a notice or disclaimer, the original version will prevail.

If a section in the Document is Entitled "Acknowledgements", "Dedications", or "History", the requirement (section 4) to Preserve its Title (section 1) will typically require changing the actual title.

9. TERMINATION

You may not copy, modify, sublicense, or distribute the Document except as expressly provided for under this License. Any other attempt to copy, modify, sublicense or distribute the Document is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.

10. FUTURE REVISIONS OF THIS LICENSE

The Free Software Foundation may publish new, revised versions of the GNU Free Documentation License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns. See <http://www.gnu.org/copyleft/>.

Each version of the License is given a distinguishing version number. If the Document specifies that a particular numbered version of this License "or any later version" applies to it, you have the option of following the terms and conditions either of that specified version or of any later version that has been published (not as a draft) by the Free Software Foundation. If the Document does not specify a version number of this License, you may choose any version ever published (not as a draft) by the Free Software Foundation.

ADDENDUM: How to use this License for your documents

```
Copyright (c) YEAR YOUR NAME.  
Permission is granted to copy, distribute  
and/or modify this document  
under the terms of the GNU Free  
Documentation License, Version 1.2  
or any later version published by the Free  
Software Foundation;  
with no Invariant Sections, no Front-Cover  
Texts, and no Back-Cover Texts.  
A copy of the license is included in the  
section entitled "GNU  
Free Documentation License".
```

If you have Invariant Sections, Front-Cover Texts and Back-Cover Texts, replace the "with...Texts." line with this:

```
with the Invariant Sections being LIST  
THEIR TITLES, with the  
Front-Cover Texts being LIST, and with the  
Back-Cover Texts being LIST.
```

If you have Invariant Sections without Cover Texts, or some other combination of the three, merge those two alternatives to suit the situation.

If your document contains nontrivial examples of program code, we recommend releasing these examples in parallel under your choice of free software license, such as the GNU General Public License, to permit their use in free software.