

Neptune: installation guide

Overview	3
Introduction	3
Hardware requirements	3
Neptune services	3
Storage	3
Installation	4
Prerequisites	4
Installation artifacts	5
Configuration	5
General configuration	5
Kubernetes cluster configuration	5
Storage configuration	6
Storage in local deployment	6
Storage in cluster deployment	6
S3-compatible object storage	7
Database configuration	7
Required database schemas	7
Elasticsearch configuration	8
Kafka configuration	8
Additional configuration	8
Exposing Neptune	9
Installation process	10

Load balancer configuration	11
Upgrading	13
Upgrading to 2.2.x	13
Migration to S3-compatible object storage	13
Upgrading to Elasticsearch 7.10+	14
Standard upgrade procedure	14
Effect on end-user	15
Backup	15
Client packages	15

Overview

Introduction

Neptune is composed of a set of microservices (Neptune Services) distributed as a Helm chart deployable on Kubernetes. If you don't have your own Kubernetes cluster deployed, our installer is able to set up a single-node cluster.

Aside from Kubernetes, Neptune requires the following components to function. All of them are delivered as part of the installation process so you don't need to worry about them.

- MySQL, version **5.7** or **8.x**
- Elasticsearch, version **7.x**, not older than **7.10**
- Kafka, version **2.8**

Additionally, Neptune can utilize S3-compatible object storage to store the bulk of your data. If you don't have such storage available, Neptune is able to store your data simply on a disk.

You can use the components delivered with Neptune, but in some cases, it makes sense to use your own. E.g. in cloud deployments, you can (and usually should) use a managed database.

Hardware requirements

Neptune services

Depending on how exactly users interact with Neptune, the hardware requirements can vary. A minimal recommended installation of Neptune requires 8 vCPUs and 32GB of RAM.

Storage

The requirement for storage is quite difficult to estimate since it depends heavily on how much data users store in Neptune.

For simplest deployments, in which all of Neptune and its data live on a single VM, the recommended minimum is 1TB SSD.

If possible, using S3-compatible object storage is strongly encouraged. It lowers the requirement for disk size to a recommended minimum of 100GB.

In case you provide your own MySQL, Elasticsearch, Kafka, and S3-compatible object storage, Neptune doesn't require any additional storage space.

You can find more information on storage requirements in the *Prerequisites* section below.

Installation

Prerequisites

Neptune installer has the following prerequisites (some requirements may be optional, depending on the type of installation):

- A host with Ubuntu 18.04 or 20.04 installed (depending on the type of Neptune installation, it will serve as either a temporary host to run the installer on or as the actual deployment host).
- Access to the Internet (required only for installation). For air-gap installation, please contact our support.
- Ansible (<https://www.ansible.com/>) in version **2.8.0** or higher. This is how to install it on Ubuntu 18.04

```
sudo apt-add-repository --yes ppa:ansible/ansible
sudo apt-get update
sudo apt-get install --yes 'ansible=2.9.*' --install-recommends
```

- yq (<https://github.com/mikefarah/yq/>), installable with `pip install yq==2.10.1`
- jq (<https://stedolan.github.io/jq/>), installable with `apt-get install jq`
- Persistent storage for Neptune services and users' data.
 - When you configure the Neptune installer to deploy a single-node Kubernetes cluster, it needs an unformatted SSD attached to the VM where the installation takes place.
 - When you deploy Neptune to an existing Kubernetes cluster, it assumes that the cluster has the capability to provision storage. In particular, managed Kubernetes

services (like Amazon's EKS, Google's GKE, and Azure's AKS) usually have this capability out of the box.

- Regardless of whether you deploy to an existing cluster or set up a new one using the Neptune installer, you can store the bulk of your data on S3-compatible object storage.
- When deploying to an existing Kubernetes cluster, the Neptune installer requires a configured `kubectl` executable available in `PATH`.
- When deploying to an existing Kubernetes cluster, the Neptune installer requires Kubernetes 1.19 or newer.

Installation artifacts

Neptune installer consists of two files.

The `neptune_installation_{version}.tgz` file needs to be unpacked. It contains a single directory `neptune_installation`.

The `configuration.yaml` file contains the installation configuration. In particular, it can contain Neptune's docker registry credentials and your license for using Neptune. **It should therefore be treated confidentially.**

Configuration

Before starting the installation you should modify the `configuration.yaml` file.

Aside from what was mentioned above, the `configuration.yaml` file may contain the following options (some are optional, depending on the installation type):

General configuration

This section describes some basic options you need to fill before installing Neptune.

- `administrator_username` - the desired username of your administrator account. It can contain letters, numbers, and hyphens. We recommend `administrator`.
- `administrator_password` - the password for the administrator account. Select a strong password here.
- `organization_name` - the name of your organization in Neptune. We suggest using something simple, related to your company. It can contain letters, numbers and hyphens.

- `deployment_type` - set to `local` in case of single-node deployment or `cluster` in case of deployment on existing Kubernetes cluster.

Kubernetes cluster configuration

Use options in this section to tell the Neptune installer how to interact with your Kubernetes cluster. Use these options only for `cluster` deployment.

- `kubeconfig_path` - path to an existing `kubectl` configuration file pointing to the cluster you want to deploy Neptune to (only for `cluster` deployment).
- `namespace` - Kubernetes namespace to which Neptune will be deployed. It's created if it doesn't exist. Defaults to `neptune`.
- `node_tolerations` - an array of node taint tolerations that Neptune installation can use. This is copied directly to Kubernetes manifests.
- `node_selector` - a key-value map that Neptune installation should use to select nodes. This is copied directly to Kubernetes manifests.

Storage configuration

Using options in this section you can configure how to provision storage for Neptune in your deployment.

The first thing to understand is what types of data Neptune stores. It can be most generally split into two categories.

1. The storage required by MySQL, Elasticsearch, and Kafka.
This data is always stored on a POSIX-compliant file system.
2. The storage for the bulk of your data: numeric series, logs, images, etc.
This data can be stored either on a POSIX-compliant file system or in S3-compatible object storage. In particular, when deploying Neptune to an existing cluster, you can choose whether to provide an S3-compatible storage service via `s3_*` options, or a POSIX file system via `storage_pvc_name`. If both are present, S3 takes precedence.

Regardless of the `deployment_type`, you can always provide MySQL, Elasticsearch, Kafka, and S3-compatible object storage in which case Neptune doesn't require additional storage. If you prefer to use the above services provided as part of the Neptune installer, they need a way to provision storage for themselves. Depending on the `deployment_type`, there are several options that take care of that.

Storage in `local` deployment

- `storage_device` - this should point to an SSD device. You can retrieve the value from the output of the `lsblk` command. For a clean VM, it's often `/dev/sdb`.

Take a look at the *S3-compatible object storage* section below for more options.

Storage in `cluster` deployment

If you decide to use MySQL, Kafka or Elasticsearch delivered as part of Neptune installation or the default file storage, you need to specify the following options (they allow Neptune to provision disk space for the mentioned services):

- `ssd_storage_class` - the name of the storage class to be used by MySQL and Elasticsearch service.
- `hdd_storage_class` - the name of the storage class to be used for Kafka.

You can also choose to provide your own file storage by setting the following option:

- `storage_pvc_name` - the name of the Persistent Volume Claim present in the namespace to which Neptune will be installed. This claim needs to be of type `ReadWriteMany` and the underlying storage has to be POSIX-compliant. This option is disregarded if S3-compatible object storage is provided (unless you're migrating from Neptune version 2.1 to 2.2, in which case it will serve as the source for the migration; no new data will be written there).

Take a look at the *S3-compatible object storage* section below for more options.

S3-compatible object storage

Options in this section work regardless of whether you're deploying Neptune in a `local` or `cluster` mode.

Neptune stores most of your data either on a POSIX file system or in object storage. In most situations, we strongly recommend using object storage

- `s3_bucket_name` - the name of the bucket Neptune uses for storing most of your data.
- `s3_service_endpoint` - the S3 service endpoint to connect to.
 - On AWS, one of the S3 service endpoints described here: <https://docs.aws.amazon.com/general/latest/gr/rande.html>
 - If you're using an S3-compatible service, you should set this to the service's endpoint, e.g. on GCP `https://storage.googleapis.com/`

- `s3_region` - the AWS region corresponding to the `s3_service_endpoint` option. For S3-compatible services, the value depends on the service.
- `s3_access_key_id` - an S3 access key.
- `s3_secret_access_key` - an S3 secret key.

Database configuration

Neptune installer can set up a MySQL database as part of the installation process, however, if you prefer to provide one yourself, this section describes how to do that. This may be desirable, especially in cloud environments, where the storage used by the database is automatically scaled.

- `db_host` - External MySQL database host for Neptune to use.
- `db_port` - External MySQL database port. Optional, defaults to `3306`.
- `db_username` - The username of a user with access to all schemes used by Neptune. Required if `db_host` is set.
- `db_password` - The password of the user specified in `db_username`. Required if `db_host` is set.

Required database schemas

Before running the Ansible installer you should create the following database schemas and grant the user defined in `db_username` access to those schemas:

- `neptune_instance`
- `neptune_notifications`
- `neptune_discussions`
- `neptune_leaderboard`
- `neptune_keycloak`
- `neptune_artifacts`

Make sure that the schemas have `utf8` as default character set

Elasticsearch configuration

Neptune installer can set up an Elasticsearch instance as part of the installation process, however, if you prefer to provide one yourself, this section describes how to do that.

- `elasticsearch_address` - An address of the external Elasticsearch server host in format `https://<address>[:<http api port>]`. The `<http api port>` is optional and defaults to 443 for https connections.
- `elasticsearch_cluster_name` - The name of your Elasticsearch cluster. Defaults to `elasticsearch`, which is the default cluster name in Elasticsearch.

Kafka configuration

Neptune installer can set up a Kafka instance as part of the installation process, however, if you prefer to provide one yourself, this section describes how to do that.

- `kafka_address` - An address of the external Kafka cluster in a comma-separated list. Required format: `<address 1>:<port>, <address 2>:<port>, ... <address N>:<port>`,

Additional configuration

In some cases you want to connect Neptune to your own identity management system (like LDAP). If your identity management system uses a certificate issued by your own Certificate Authority, you need to provide a set of trusted certificates to Neptune (so that it's able to verify the security of the connection).

- `trusted_certificates` - A list of absolute paths to files containing certificates that are to be trusted
- `keycloak_java_opts` - A rarely used option that allows overriding some default behaviors of Java. It's a string containing Java options that are to be added to Keycloak (the component responsible for connecting to an external identity management system). Example value: `"-Djdk.tls.client.protocols=TLSv1.0,TLSv1.1,TLSv1.2"`

Exposing Neptune

This section deals with exposing Neptune from your cluster/VM to the outside.

The first thing you need to decide is whether to use an ingress controller embedded within the Neptune installer, or provide your own.

For more information on ingress controllers, see <https://kubernetes.io/docs/concepts/services-networking/ingress-controllers>

- `ingress_controller` - this option determines whether the installer should deploy an ingress controller, or use an ingress controller already present in the cluster. If you're installing Neptune within a VM or on a fresh cluster, chances are you should go with the default.
 - `embedded` (the default) - the installer will install an NGINX ingress controller
 - `provided` - the ingress controller will not be installed. It's assumed that you have an ingress controller working in the cluster.
 - `provided-aws` - a special value applicable to deployments on Amazon's EKS clusters. The ingress controller will not be installed and the Neptune services will be deployed in a way compatible with AWS Application Load Balancer. Please note, that AWS-specific annotations still need to be configured via the `ingress_annotations` option (see below). Additionally, keep in mind, that prior to the installation your EKS cluster needs to be configured to provision ALBs for ingress resources.

For more information on exposing ingress resources as ALBs, see <https://docs.aws.amazon.com/eks/latest/userguide/alb-ingress.html>

- `service_exposition_type` - valid only if `ingress_controller` is set to `embedded` (which it is, by default). Determines the way Neptune (actually Neptune's embedded ingress controller) is exposed. Can take one of the following values:
 - `LoadBalancer` - usually the preferred option when running in a cloud environment. It provisions a load balancer that points to Neptune.

For more information on LoadBalancer services, see <https://kubernetes.io/docs/concepts/services-networking/service/#loadbalancer>

- `NodePort` (the default) - commonly used when deploying Neptune on your own hardware (especially when using the single VM setup). In this case, Neptune will be exposed on port 30080 on every node of your Kubernetes cluster.

For more information on NodePort services, see <https://kubernetes.io/docs/concepts/services-networking/service/#nodeport>

See the "Load balancer configuration" section below on how to set up a load balancer that directs traffic to Neptune exposed as a `NodePort` service.

The options below require some understanding of Kubernetes' concepts. You can skip them if you're installing Neptune within a VM and plan to add a separate load balancer to strip TLS.

- `ingress_host` - defines a domain at which the ingress controller will listen for requests to Neptune. translates directly to `.spec.rules[].host` in Neptune's ingress. Note that if this option is set, Neptune will be available only through this domain.
- `ingress_annotations` - a dictionary, that translates directly to `.metadata.annotations` in Neptune's ingress. You can use this option to configure behavior specific to your provided ingress controller, like forcing an SSL redirect.
- `ingress_labels` - a dictionary, that translates directly to `.metadata.labels` in Neptune's ingress. Please, note, that the Neptune installer might add a few labels. In particular, label names `app`, `chart`, `release`, and `heritage` are reserved, and providing them will cause installation failure.
- `ingress_tls_secret` - a name of a secret in the target namespace containing a valid TLS certificate and key for ingress to use. If defined, values of `ingress_tls_cert` and `ingress_tls_key` are ignored.
- `ingress_tls_cert`, `ingress_tls_key` - a base64-encoded TLS certificate and private key. Using those two values the installer will build a secret and provide it to `.spec.tls.secretName` in Neptune's ingress.

Installation process

The actual installation can be performed with a single command executed from within the `neptune_installation` directory.

```
ansible-playbook neptune.yml --extra-vars @/path/to/configuration.yaml
```

Remember about the `@` character in the above command! It's a part of ansible's syntax.

If you're installing Neptune with `deployment_type` set to `local`, you need to run the above command as root.

The installation will take from 5 minutes up to an hour, depending on the cluster's Internet connection speed - installation downloads Neptune's docker images.

Read the next section to learn more about exposing and accessing Neptune.

Load balancer configuration

If you're deploying Neptune on your own hardware with `ingress_controller` set to `embedded` (the default) and `service_exposition_type` set to `NodePort`, Neptune will be exposed on port

30080 on each node of the cluster. An external load balancer has to be configured to terminate SSL/TLS connection and forward traffic to all nodes to port 30080.

Obviously, in the case of deployment on a single VM, there's only a single node.

Any SSL-stripping load balancer should work, provided it adds `x-forwarded-for`, `x-forwarded-port`, `x-forwarded-proto`, and `x-forwarded-host` headers properly and allows for long-lasting WebSocket connections.

Note, that `neptune-client` can send large HTTP request bodies, so putting (for nginx) `client_max_body_size 30m`; is required to make sure client works properly

We recommend using NGINX - you can use the sample configuration below as an inspiration:

```
server {
    listen 80;
    return 302 https://$host$request_uri;
}

upstream kubernetes-nodes {
    server 127.0.0.1:30080;
}

server {
    listen 80;
    return 302 https://$host$request_uri;
}

map $http_upgrade $connection_upgrade {
    default upgrade;
    '' close;
}

map $http_x_forwarded_proto $new_x_forwarded_proto {
    default $http_x_forwarded_proto;
    "" $scheme;
}

map $http_x_forwarded_host $new_http_x_forwarded_host {
    default $http_x_forwarded_host;
    "" $host;
}

map $http_x_forwarded_port $new_http_x_forwarded_port {
    default $http_x_forwarded_port;
    "" $server_port;
}

server {
    listen 443;
    server_name _;
```

```
ssl_certificate /etc/nginx/certs/cert.crt;
ssl_certificate_key /etc/nginx/certs/cert.key;
ssl on;
ssl_session_cache builtin:1000 shared:SSL:10m;
ssl_protocols TLSv1 TLSv1.1 TLSv1.2;
ssl_ciphers HIGH:!aNULL:!eNULL:!EXPORT:!CAMELLIA:!DES:!MD5:!PSK:!RC4;
ssl_prefer_server_ciphers on;

access_log /var/log/nginx/access.log;

location / {

    client_max_body_size 10G; # prevents 413 header too large for proxy

    # enables websockets. Needs all 3 to work. missing variables causes
header to not be set
    # so non-ws connections dont have Upgrade and Connection headers set
    proxy_set_header Upgrade $http_upgrade;
    proxy_set_header Connection $connection_upgrade;
    proxy_set_header Host $host;

    proxy_set_header X-Real-IP $remote_addr;

    # needed for keycloak and backend to read client side address of Neptune
    proxy_set_header X-Forwarded-For $proxy_add_x_forwarded_for;
    proxy_set_header X-Forwarded-Proto $new_x_forwarded_proto;
    proxy_set_header X-Forwarded-Host $new_http_x_forwarded_host;
    proxy_set_header X-Forwarded-Port $new_http_x_forwarded_port;

    proxy_pass http://kubernetes-nodes$request_uri;

    # websocket/upload connection kill prevention
    proxy_read_timeout 86400;
    proxy_send_timeout 86400;

    proxy_http_version 1.1;
}
}
```

Upgrading

If instead of installing Neptune, you're upgrading a previous installation to the new version, this section describes how to do it. Make sure to read the sections below on upgrading between particular versions.

Important: Be advised that upgrading to version 1.X is supported only from the previous version: 1.(X-1) (so, for example, upgrading to version 1.2 is supported only from version 1.1). Do not attempt to upgrade if you have an older version installed.

Upgrading to 2.2.x

This new version of Neptune comes with several important changes. Most notably, we're introducing S3 support for storing most of your data. If you'd like to use S3 or S3-compatible object storage (we strongly recommend that if you have this option), please read the migration overview below (if you decide to go ahead with it, contact us - we'd love to help!).

We're also changing the requirements on the version of Elasticsearch. Neptune 2.2.x works with ElasticSearch 6.7+ and 7.10+. The next version of Neptune will drop support of Elasticsearch 6.x completely and support 7.10+ exclusively.

Migration to S3-compatible object storage

Until version 2.1, most of the Neptune users' data was stored either on a persistent volume provided to Neptune via the `storage_pvc_name` option or on an internal POSIX file system. If you wish to continue using the storage you're using, you can skip this section. We do, however, recommend switching to S3-compatible object storage, especially in cloud environments, for improved scalability and performance.

If you do decide to go ahead with this migration, please contact our support team as we'd love to guide you through this process. Below is a brief overview of what will happen.

The data that needs migrating is (simplifying a bit) of two kinds: series data (numeric and text series) and files.

The migration happens in 2 phases:

1. Migrating series data to S3. This phase requires downtime. The length of this downtime depends on how much series data there is and on connection speed.
2. Migrating files. This works in the background during normal operations of Neptune 2.2. The duration, again, depends on the volume of data, but as it's a background process, it won't bother Neptune's users.

Upgrading to Elasticsearch 7.10+

This is only applicable if you're using your own Elasticsearch service (and not the one bundled within the Neptune installer).

Neptune 2.1 requires an Elasticsearch in version 6.7+, while Neptune 2.2 can work with either Elasticsearch 6.7+ or 7.10+. Before the upgrade to 7.10+ is performed, Neptune needs to run some checks on the data stored within the Elasticsearch index to make the data doesn't exceed the limits introduced in Elasticsearch 7.

The upgrade procedure should look like this:

1. Upgrade the Elasticsearch cluster to version 6.8.23
2. Upgrade Neptune to version 2.2 and wait for about one hour, during which time, Neptune will perform the checks on the Elasticsearch index (mentioned above) to ensure a smooth upgrade.
3. Upgrade the Elasticsearch cluster to version 7.10+

Standard upgrade procedure

A typical upgrade procedure consists of just a few steps (somewhat similar to a first-time installation):

1. Perform a backup (see section below). This step is optional, but strongly recommended.
2. Download the installer (`neptune_installation_{version}.tgz`) with the new version to a host where you will run the installer and unpack the archive.
3. From within the `neptune_installation` directory, run

```
ansible-playbook neptune.yml --extra-vars @/path/to/configuration.yaml
```

Remember about the `@` character in the above command! It's a part of ansible's syntax.

If you're upgrading Neptune with `deployment_type` set to `local`, you need to run the above command as root.

Remember that the configuration file used previously contains your license, credentials for downloading Neptune images, your Neptune administrator credentials and a few others - they shouldn't be changed without consulting with Neptune team.

4. That's it! After a while, you should be happily using the upgraded version of Neptune.

Effect on end-user

An upgrade should usually take from 5 minutes to an hour, depending on your Internet connection.

Some upgrades require an upgrade of the client libraries (like `neptune-client` and `neptune-notebooks`). Read the “Client packages” section to learn more.

Upgrades are often imperceptible to the end-user (the data scientist running experiments), however, there’s always a risk of Neptune being unavailable for some time. This risk is mitigated on the `neptune-client` library side - if the Neptune server is unavailable, runs switch to an offline mode, which stores all the metadata on the local disk. The data can be later manually uploaded using the `neptune sync` command.

Backup

Remember to back up your Neptune deployment before upgrading.

In particular, this means creating a snapshot of any persistent volume and/or database that you provide for Neptune.

Client packages

This release is compatible with the `neptune-client` library in versions between `0.16.3` (inclusive) and `0.17.0` (exclusive).

It can be installed with:

```
pip install "neptune-client>=0.16.3,<0.17.0"
```

The version of the `neptune-notebooks` package compatible with this release is `0.9.3`.

It can be installed with:

```
pip install neptune-notebooks==0.9.3
```