Procedure to Kerberize HDP 3.1

Below is the procedure to setup Kerberos on HDP 3.1.0.0

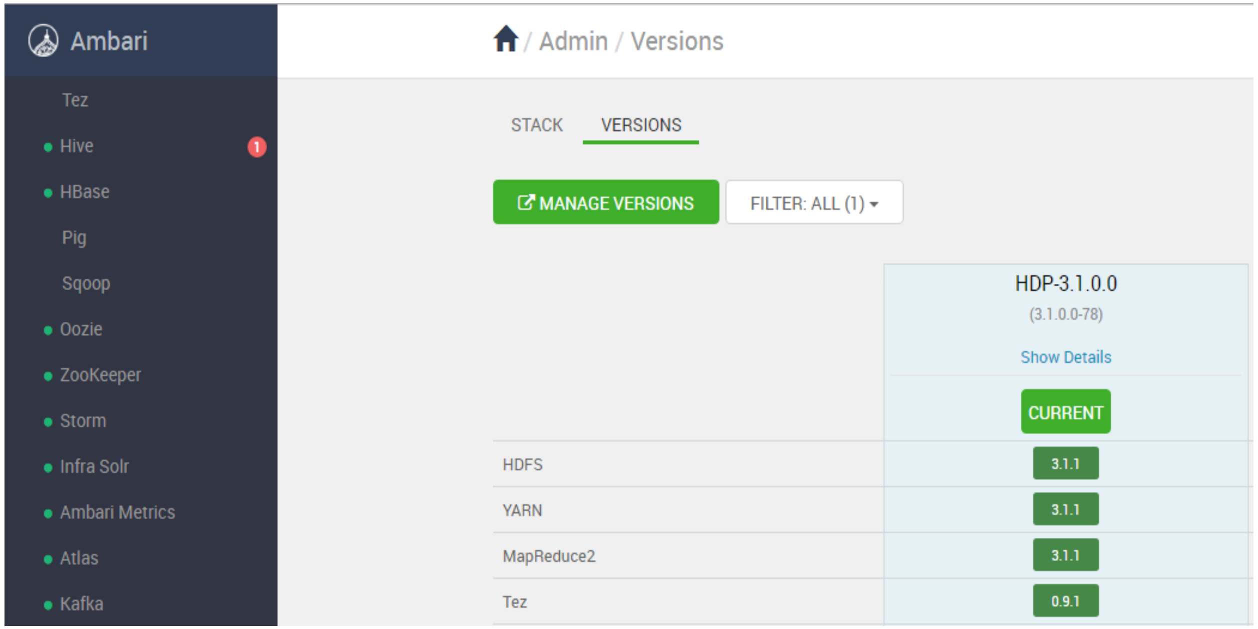

	HDP-3.1.0.0 (3.1.0.0-78)
HDFS	3.1.1
YARN	3.1.1
MapReduce2	3.1.1
Tez	0.9.1

Fresh install all green

To ensure successful install all the HDP components should be in good health ….**GREEN**

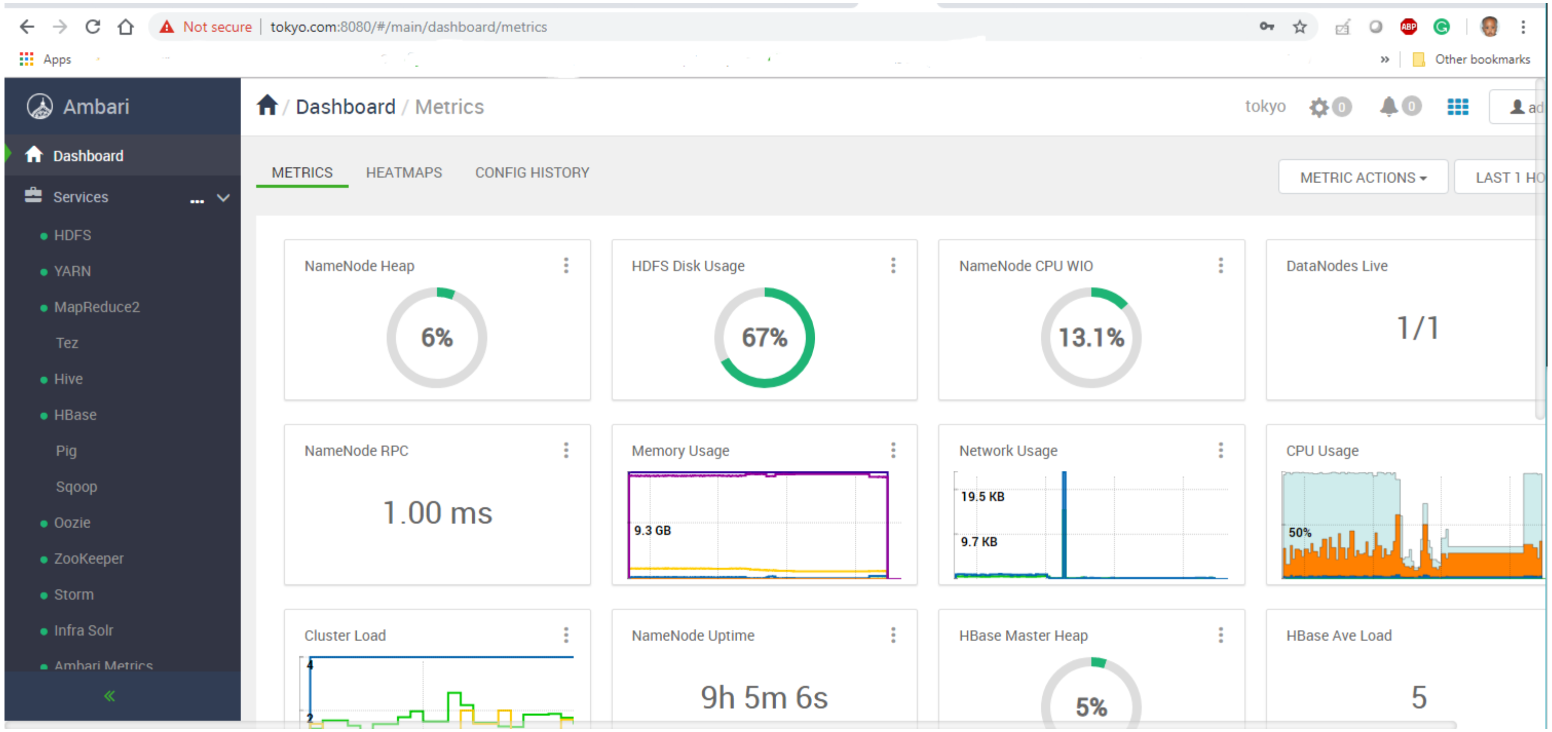

Un-Kerberized

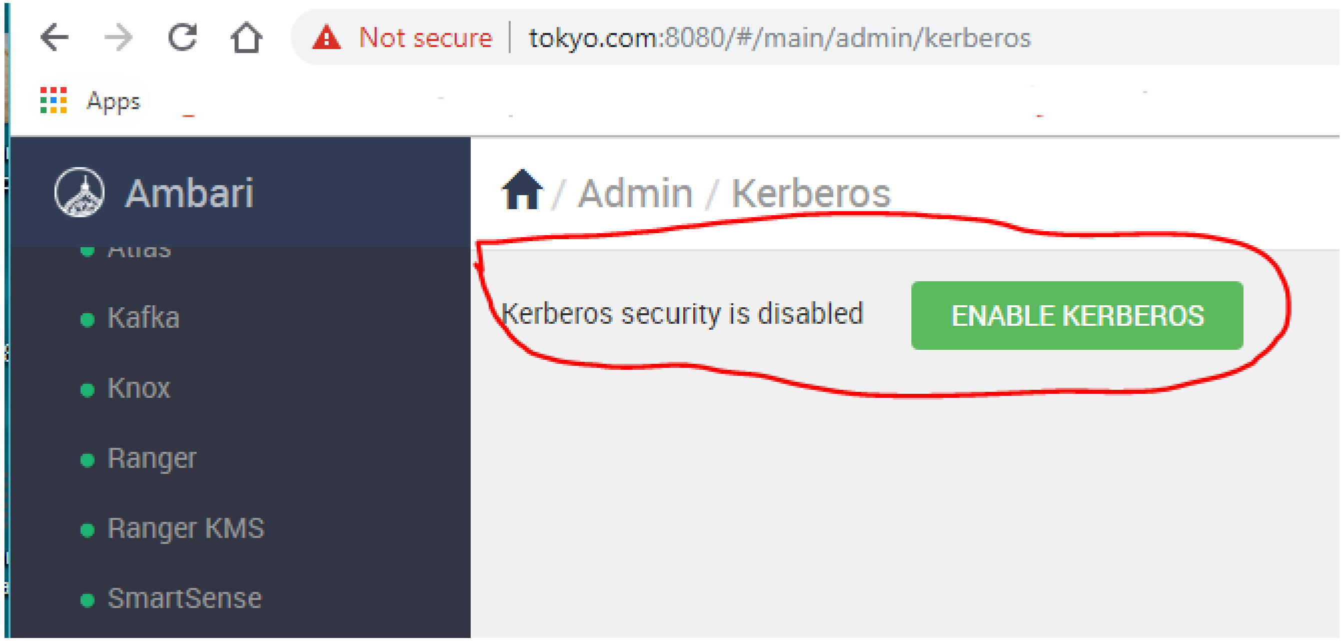

Assumption

All the commands should be run as root or someone with sudo privileges

The Cluster has been pre-installed

- Cluster Name	**Tokyo**
- Cluster Mode	**Single Node**
- KDC host	**Tokyo.com** (Kerberos uses FQDN)
- REALM	**Tokyo.com**
- Kerberos master password	**welcome1**
- OS	**Centos 7**

Kerberos Installation

The below install steps should be performed on the KDC server ONLY, on all the client nodes i.e. In a multi-node cluster the KDC can run on one of the nodes or outside the cluster but MUST be accessible to the cluster nodes.

Install the KDC Server RHEL/CentOS/Oracle Linux

```
# yum install -y krb5-server krb5-libs krb5-workstation
```

Desired output Installed: krb5-server.x86_64 0:1.15.1-34.el7 Complete!

Install the kerberos clients

Install the Kerberos client package on all the cluster client

```
# yum install krb5-libs krb5-workstation
```

Edit the krb5.conf

Configuration snippets may be placed in this directory as well includedir /etc/krb5.conf.d/

Note the REALM upper and lower case, this modified krb5.conf should be copied to all the hosts in the cluster and should replace the files delivered by client install step in **/etc/krb5.conf**

On the KDC edit

This modified krb5.conf should be copied to all the hosts in the cluster it's usually a good practice to backup original files

```
cp /etc/krb5.conf  /etc/krb5.conf.bak
```

vi /etc/krb5.conf

```
[logging]
 default = FILE:/var/log/krb5libs.log
 kdc = FILE:/var/log/krb5kdc.log
 admin_server = FILE:/var/log/kadmind.log

[libdefaults]
 dns_lookup_realm = false
 ticket_lifetime = 24h
 renew_lifetime = 7d
 forwardable = true
 rdns = false
 pkinit_anchors = /etc/pki/tls/certs/ca-bundle.crt
 default_realm = TOKYO.COM
 default_ccache_name = KEYRING:persistent:%{uid}

[realms]
TOKYO.COM = {
kdc = tokyo.com
admin_server = tokyo.com
}

[domain_realm]
.tokyo.com = TOKYO.COM
tokyo.com = TOKYO.COM
```

Edit the kdc.conf

Adjust /var/kerberos/krb5kdc/kdc.conf on the KDC

```
[kdcdefaults]
 kdc_ports = 88
 kdc_tcp_ports = 88
```

```
[realms]
 TOKYO.COM = {
  #master_key_type = aes256-cts
  acl_file = /var/kerberos/krb5kdc/kadm5.acl
  dict_file = /usr/share/dict/words
  admin_keytab = /var/kerberos/krb5kdc/kadm5.keytab
  supported_enctypes = aes256-cts:normal aes128-cts:normal des3-hmac-sha1:normal arcfour-hmac:normal camellia256-cts:normal camellia128-cts:normal des-hmac-sha1:normal des-cbc-md5:normal des-cbc-crc:normal
 }
```

Edit kadm5.acl

Adjust /var/kerberos/krb5kdc/kadm5.acl on KDC:

```
    */admin@TOKYO.COM       *
```

Create the KDC database

Creating KDC database to hold our sensitive Kerberos data

Create the database and set a good password which you can remember. This command also stashes your password on the KDC so you don’t have to enter it each time you start the KDC:

Use the utility kdb5util to create the Kerberos database. The -s option, kdb5util will stash a copy of the master key in a stash file. A stash file allows a KDC to authenticate itself to the database utilities, such as **kadmin, kadmind, krb5kdc**, and **kdb5_util**.

```
# kdb5_util create -s
Loading random data
Initializing database '/var/kerberos/krb5kdc/principal' for realm 'TOKYO.COM',
master key name 'K/M@TOKYO.COM'
You will be prompted for the database Master Password.
It is important that you NOT FORGET this password.
```

```
Enter KDC database master key: {strong_password}
Re-enter KDC database master key to verify:{strong_password}
```

Don't forget or misplace this password as you will need it for the Ambari-Kerberos setup

Destroy the KDC database

In case you have made a mistake before generating the keytabs you can destroy and recreate the KDC database as below

```
# kdb5_util -r TOKYO.COM destroy
Deleting KDC database stored in '/var/kerberos/krb5kdc/principal', are you sure?
(type 'yes' to confirm)? yes
OK, deleting database '/var/kerberos/krb5kdc/principal'...
** Database '/var/kerberos/krb5kdc/principal' destroyed.
```

AutoStart KDC components

To ensure that Kerberos Service is always up incase of reboot of the server

```
# systemctl enable krb5kdc
```

Desired output

Created symlink from /etc/systemd/system/multi-user.target.wants/krb5kdc.service to /usr/lib/systemd/system/krb5kdc.service.

```
# systemctl enable kadmin
```

Desired output

Created symlink from /etc/systemd/system/multi-user.target.wants/kadmin.service to /usr/lib/systemd/system/kadmin.service.

Start the Kerberos components

These 2 components MUST be running to be able to Kerberize the cluster

```
# systemctl start krb5kdc
# systemctl start kadmin
```

Validate KCD component status

To ensure that the services started run the below commands you should see **SUCCESS**

```
# systemctl status kadmin
```

Desired output

```
● kadmin.service - Kerberos 5 Password-changing and Administration
   Loaded: loaded (/usr/lib/systemd/system/kadmin.service; enabled; vendor preset: disabled)
   Active: active (running) since Sat 2019-01-05 19:47:33 CET; 3h 34min ago
 Main PID: 31865 (kadmind)
   CGroup: /system.slice/kadmin.service
           └─31865 /usr/sbin/kadmind -P /var/run/kadmind.pid
Jan 05 19:47:33 tokyo.com systemd[1]: Starting Kerberos 5 Password-changing and Administration...
Jan 05 19:47:33 tokyo.com systemd[1]: Started Kerberos 5 Password-changing and Administration.
```

```
# systemctl status krb5kdc
```

Desired output

```
● krb5kdc.service - Kerberos 5 KDC
   Loaded: loaded (/usr/lib/systemd/system/krb5kdc.service; enabled; vendor preset: disabled)
   Active: active (running) since Sat 2019-01-05 19:47:24 CET; 3h 37min ago
 Main PID: 31694 (krb5kdc)
   CGroup: /system.slice/krb5kdc.service
           └─31694 /usr/sbin/krb5kdc -P /var/run/krb5kdc.pid

Jan 05 19:47:24 tokyo.com systemd[1]: Starting Kerberos 5 KDC...
Jan 05 19:47:24 tokyo.com systemd[1]: Started Kerberos 5 KDC.
```

Create administrative root user