FORGEROCK®

Installation Guide

/ Autonomous Identity 7.1

Latest update: 2020.10.2

,
, ,

Abstract

Guide to installing, and uninstalling ForgeRock® Autonomous Identity software. This software provides the entitlement analytics for your Identity and Access Management (IAM) systems.

Table of Contents

Overview

This guide shows you how to install and deploy Autonomous Identity for intelligent entitlements management in production environments. For hardware and software requirements, see the Release Notes.

ForgeRock® Autonomous Identity is an entitlements analytics system that lets you fully manage your company's access to your data.

An entitlement refers to the rights or privileges assigned to a user or thing for access to specific resources. A company can have millions of entitlements without a clear picture of what they are, what they do, and who they are assigned to. Autonomous Identity solves this problem by using advanced artificial intelligence (AI) and automation technology to determine the full entitlements landscape for your company. The system also detects potential risks arising from incorrect or over-provisioned entitlements that lead to policy violations. Autonomous Identity eliminates the manual re-certification of entitlements and provides a centralized, transparent, and contextual view of all access points within your company.

Quick Start

Architecture in Brief Learn about the Autonomous Identity architecture.	Deployment Architectures Learn about the different deployment architectures.	Install a Single-Node Target Install a single-node Autonomous Identity installation.
Install a Single-Node Air-Gapped Install a single-node air-gapped Autonomous Identity installation.	Install a Multi-Node Install a multi-node Autonomous Identity installation.	Install a Multi-Node Air-Gapped Install a multi-node Autonomous Identity air-gapped installation.
Upgrade Autonomous Identity Upgrade to the latest version of Autonomous Identity.	Appendix: Ports Learn about the Autonomous Identity ports.	Appendix: vars.yml Learn about the main deployment configuration file.

For a description of the Autonomous Identity UI console, see the Installation Guide.

Chapter 1
Features

Autonomous Identity provides the following features:

- **Broad Support for Major Identity Governance and Administration (IGA) Providers**. Autonomous Identity supports a wide variety of Identity as a Service (IDaaS) and Identity Management (IDM) data including but not limited to comma-separated values (CSV), Lightweight Directory Access Protocol (LDAP), human resources (HR), database, and IGA solutions.

- **Highly-Scalable Architecture**. Autonomous Identity deploys using a microservices architecture, either on-prem, cloud, or hybrid-cloud environments. Autonomous Identity's architecture scales linearly as the load increases.

- **Powerful UI dashboard**. Autonomous Identity displays your company's entitlements graphically on its UI console. You can immediately investigate those entitlement outliers as possible security risks. The UI also lets you quickly identify those entitlements that are good candidates for automated low-risk approvals or re-certifications. Users can also view a trend-line indicating how well they are managing their entitlements. The UI also provides an application-centric view and a single-page rules view for a different look at your entitlements.

- **Automated Workflows**. Autonomous Identity reduces the burden on managers who must manually approve new entitlements, for example, assigning access for new hires, by auto-approving high confidence, low-risk access requests and automate the re-certification of entitlements. Predictive recommendations lends itself well to automation, which saves time and cost.

- **Powerful Analytics Engine**. Autonomous Identity's analytics engine is capable of processing millions of access points within a short period of time. Autonomous Identity lets you configure the machine learning process and prune less productive rules. Customers can run analyses, predictions, and recommendations frequently to improve the machine learning process.

- **Powerful Explainable AI Algorithms**. The Analytics Engine provides transparent and explainable results that lets business users get insight into why the end-user has the access they have, or what access they should have.

- **Broad Database Support**. Autonomous Identity supports both Apache Cassandra and MongoDB databases. Both are highly distributed databases with wide usage throughout the industry.

- **Improved Search Support**. Autonomous Identity now incorporates Open Distro for Elasticsearch, a distributed, open-source search engine based on Lucene, to improve database search results and performance.

Chapter 2
Architecture in Brief

The Autonomous Identity architecture has a simple three-layer conceptual model:

- **Application Layer**. Autonomous Identity implements a flexible Docker Swarm microservices architecture, where multiple applications run in containers. The microservices component provides flexible configuration and end-user interaction to the deployment. The microservices components are the following:

  - **Autonomous Identity UI**. Autonomous Identity supports a dynamic UI that displays the entitlements, confidence scores, and recommendations.

  - **Autonomous Identity API**. Autonomous Identity provides an API that can access endpoints using REST. This allows easy scripting and programming for your system.

  - **Self-Service Tool**. The self-service tool lets users reset their Autonomous Identity passwords.

  - **Backend Repository**. The backend repository stores Autonomous Identity user information. To interface with the backend repository, you can use the **phpldapadmin** tool to enter and manage users.

  - **Configuration Service**. Autonomous Identity supports a configuration service that allows you to set parameters for your system and processes.

  - **Command-Line Interface**. Autonomous Identity supports a command-line interface for easy scripting and automation.

  - **Nginx**. Nginx is a popular HTTP server and reverse proxy for routing HTTPS traffic.

  - **Hashicorp Consul**. Consul is a third-party system for service discovery and configuration.

- **Data Layer**. Autonomous Identity supports Apache Cassandra NoSQL and MongoDB databases to serve predictions, confidence scores, and prediction data to the end user. Apache Cassandra is a distributed and linearly scalable database with no single point of failure. MongoDB is a schema-free, distributed database that uses JSON-like documents as data objects.

  Autonomous Identity also implements Open Distro for Elasticsearch and Kibana to improve search performance for its entitlement data.

- **Analytics and Administration Layer**. Autonomous Identity uses a multi-source Apache Spark analytics engine to generate the predictions and confidence scores. Apache Spark is a distributed, cluster-computing framework for AI machine learning for large datasets. Autonomous Identity also uses a deployer wrapper script to launch Ansible playbooks for easy and quick deployment of the components.

Figure 1: A Simple Conceptual Image of the Autonomous Identity Architecture

Chapter 3
Deployment Architectures

To simplify your deployments, ForgeRock provides a deployer script to install Autonomous Identity on a target node. The deployer pulls in images from the ForgeRock Google Cloud Repository (gcr.io) and uses it to deploy the the microservices, analytics machine, and database for Autonomous Identity on a target machine. The target machine only requires the base operating system, CentOS 7 or later.

There are five basic deployments, all of them similar, but in slightly different configurations:

- **Lightweight Single-Node Target Deployment**. Deploy Autonomous Identity on a single target machine without the analytics pipeline. This alleviates any direct data load on your target machine, which can be memory intensive. You can deploy this configuration on a 4-core 16 GB virtual machine on a cloud service, such as Google Cloud Platform (GCP), Amazon Web Services (AWS), or others. This configuration is only for evaluation purposes and is outlined in the *Getting Started Guide*.

  Figure 3: A lightweight single-node target deployment.

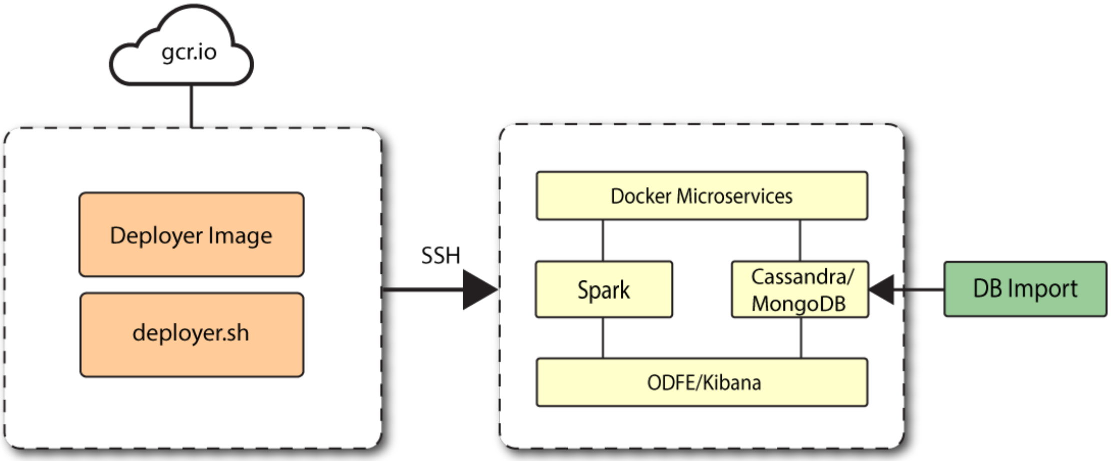

- **Single-Node Target Deployment**. Deploy Autonomous Identity on a single Internet-connected target machine. The deployer script lets you deploy the system from a local laptop or machine or from the target machine itself. The target machine can be on on-prem or on a cloud service, such as Google Cloud Platform (GCP), Amazon Web Services (AWS), Microsoft Azure or others. For installation instructions, see "*Install a Single Node Target*".

  Figure 4: A single-node target deployment.

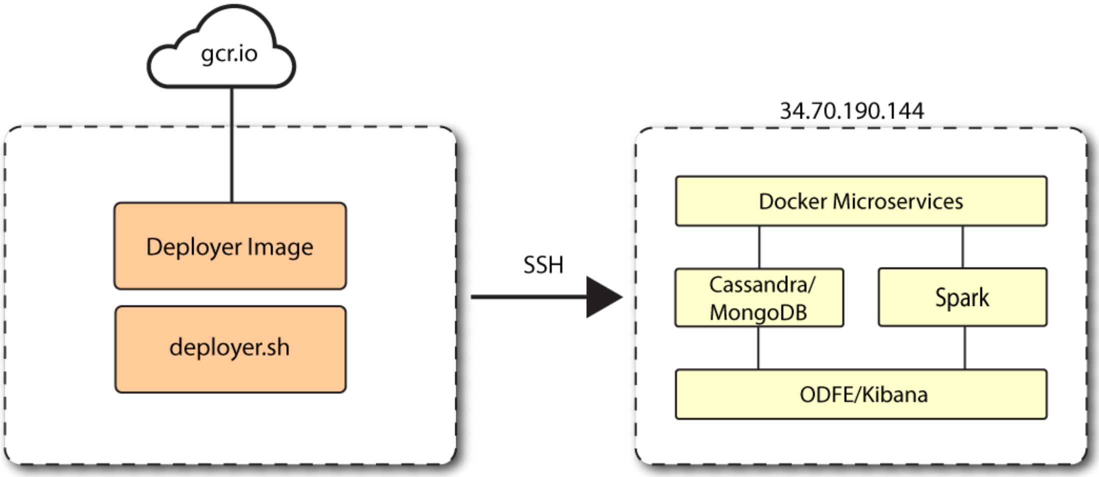

- **Single-Node Air-Gapped Target Deployment**. Deploy Autonomous Identity on a single-node target machine that resides in an air-gapped environment. In an air-gapped environment, the target machine is placed in an enhanced security environment where external Internet access is not available. You transfer the deployer and image to the target machine using media, such as a USB stick or portable drive. Then, run the deployment within the air-gapped environment. For installation instruction, see "*Install a Single Node Air-Gap Target*".

  Figure 5: An air-gapped environment.

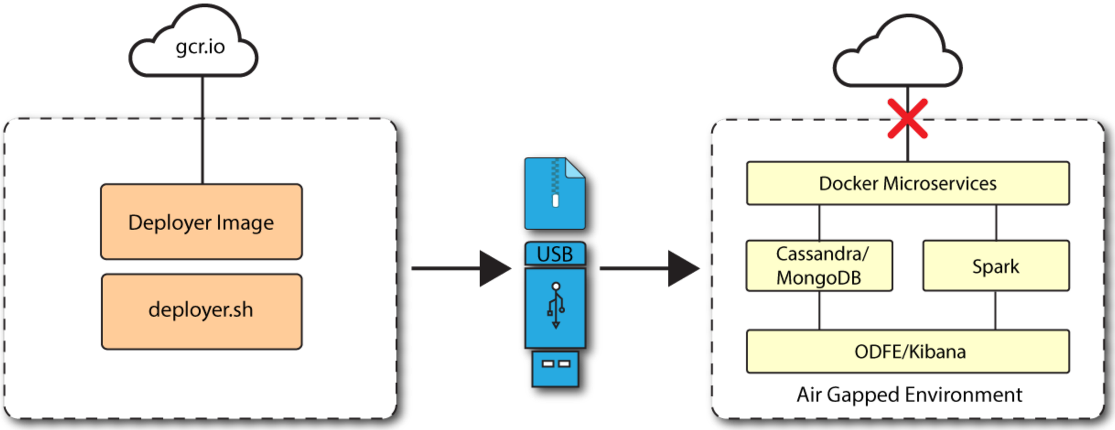

- **Multi-Node Deployment**. Deploy Autonomous Identity on multi-node deployment to distribute the process load on the servers. For installation instruction, see "*Install a Multi-Node Deployment*"

Figure 6: A multi-node target environment.

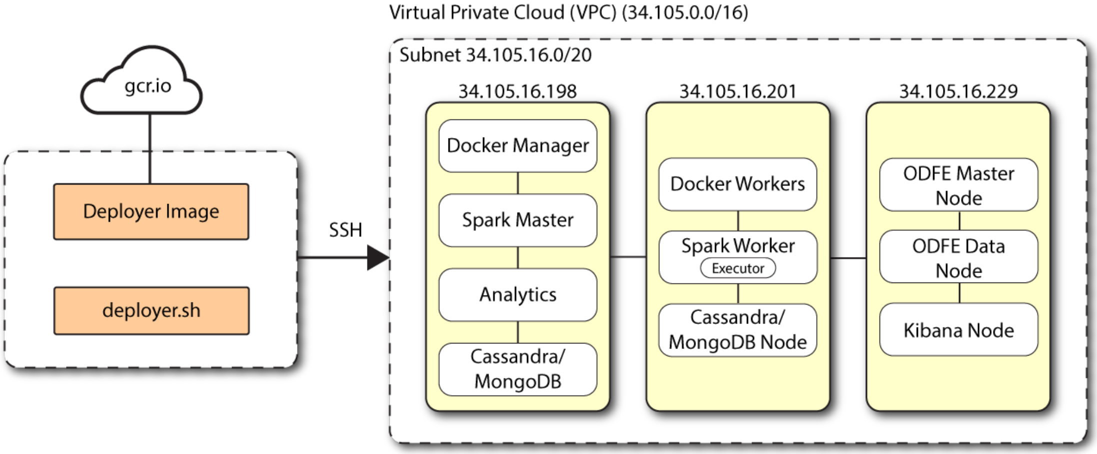

- **Multi-Node Air-Gapped Deployment**. Deploy Autonomous Identity a multi-node configuration in an air-gapped network. The multinode network has no external Internet connection. For installation instruction, see "*Install a Multi-Node Air-Gapped Deployment*".

Figure 7: A multi-node air-gapped target environment.

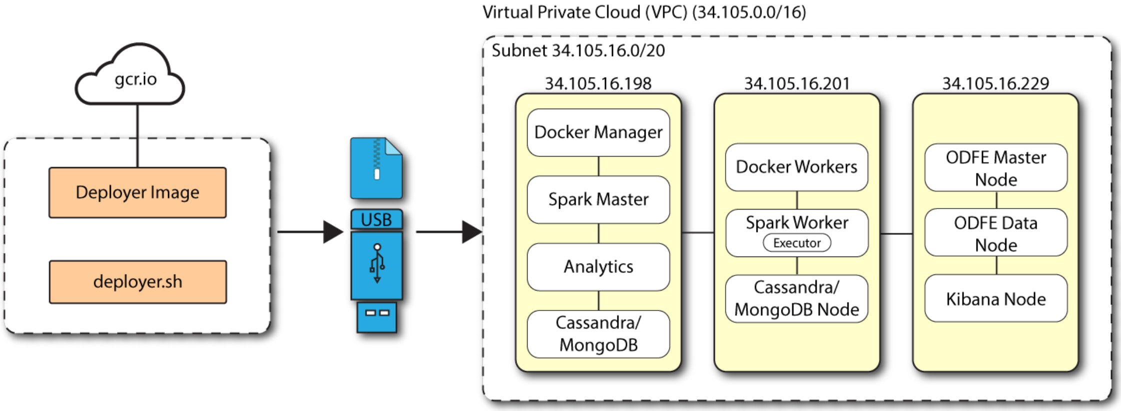

Chapter 4
Install a Single Node Target

This chapter presents instructions on deploying Autonomous Identity in a single-target machine that has Internet connectivity. ForgeRock provides a deployer script that pulls a Docker container image from ForgeRock's Google Cloud Registry (gcr.io) repository. The image contains the microservices, analytics, and backend databases needed for the system.

This installation assumes that you set up the deployer script on a separate machine from the target. This lets you launch a build from a laptop or local server.

Figure 8: A single-node target deployment.

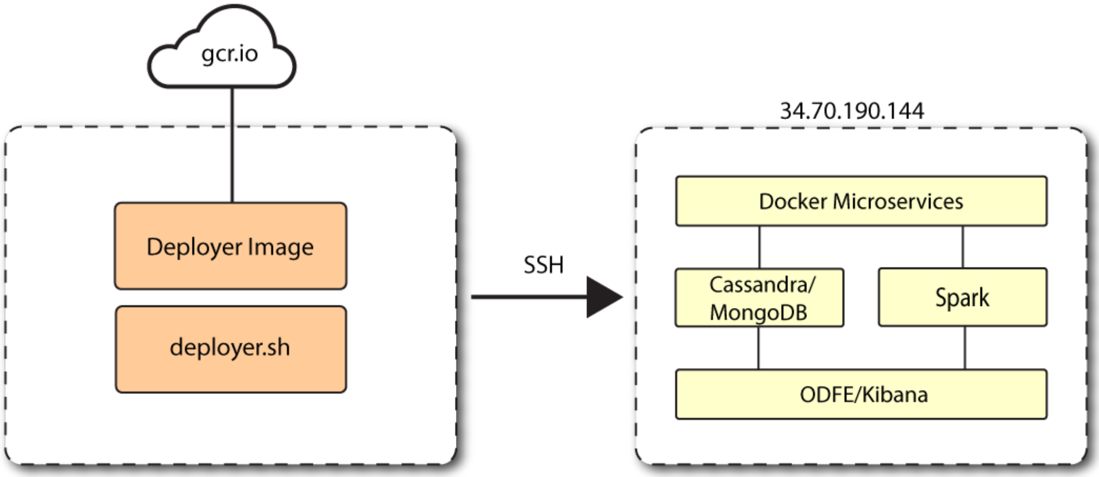

Let's deploy Autonomous Identity on a single-node target on CentOS 7. The following are prerequisites:

- **Operating System**. The target machine requires CentOS 7. The deployer machine can use any operating system as long as Docker is installed. For this guide, we use CentOS 7 as its base operating system.

- **Memory Requirements**. Make sure you have enough free disk space on the deployer machine before running the `deployer.sh` commands. We recommend at least a 40GB/partition with 14GB used and 27GB free after running the commands.

- **Default Shell**. The default shell for the autoid user must be bash.
- **Deployment Requirements**. Autonomous Identity provides a Docker image that creates a deployer.sh script. The script downloads additional images necessary for the installation. To download the deployment images, you must first obtain a registry key to log into the ForgeRock Google Cloud Registry (gcr.io). The registry key is only available to ForgeRock Autonomous Identity customers. For specific instructions on obtaining the registry key, see How To Configure Service Credentials (Push Auth, Docker) in Backstage.
- **Database Requirements**. Decide which database you are using: Apache Cassandra or MongoDB.
- **IPv4 Forwarding**. Many high-security environments run their CentOS-based systems with IPv4 forwarding disabled. However, Docker Swarm does not work with a disabled IPv4 forwarding setting. In such environments, make sure to enable IPv4 forwarding in the file etc/sysctl.conf:

  ```
  net.ipv4.ip_forward=1
  ```

Set Up the Target Machine

Autonomous Identity is configured on a target machine. Make sure you have sufficient storage for your particular deployment. For more information on sizing considerations, see *Deployment Planning Guide*.

1. The install assumes that you have CentOS 7 as your operating system. Check your CentOS 7 version.

   ```
   $ sudo cat /etc/centos-release
   ```

2. Set the user for the target machine to a username of your choice. For example, autoid.

   ```
   $ sudo adduser autoid
   ```

3. Set the password for the user you created in the previous step.

   ```
   $ sudo passwd autoid
   ```

4. Configure the user for passwordless sudo.

   ```
   $ echo "autoid  ALL=(ALL)  NOPASSWD:ALL" | sudo tee /etc/sudoers.d/autoid
   ```

5. Add administrator privileges to the user.

   ```
   $ sudo usermod -aG wheel autoid
   ```

6. Change to the user account.

   ```
   $ su - autoid
   ```

7. Install yum-utils package on the deployer machine. yum-utils is a utilities manager for the Yum RPM package repository. The repository compresses software packages for Linux distributions.

```
$ sudo yum install -y yum-utils
```

Set Up the Deployer Machine

Set up another machine as a deployer node. You can use any OS-based machine for the deployer as long as it has Docker installed. For this example, we use CentOS 7.

1. The install assumes that you have CentOS 7 as your operating system. Check your CentOS 7 version.

   ```
   $ sudo cat /etc/centos-release
   ```

2. Set the user for the target machine to a username of your choice. For example, `autoid`.

   ```
   # sudo adduser autoid
   ```

3. Set the password for the user you created in the previous step.

   ```
   $ sudo passwd autoid
   ```

4. Configure the user for passwordless sudo.

   ```
   $ echo "autoid  ALL=(ALL)  NOPASSWD:ALL" | sudo tee /etc/sudoers.d/autoid
   ```

5. Add administrator privileges to the user.

   ```
   $ sudo usermod -aG wheel autoid
   ```

6. Change to the user account.

   ```
   $ su - autoid
   ```

7. Install yum-utils package on the deployer machine. yum-utils is a utilities manager for the Yum RPM package repository. The repository compresses software packages for Linux distributions.

   ```
   $ sudo yum install -y yum-utils
   ```

8. Create the installation directory. Note that you can use any install directory for your system as long as your run the **deployer.sh** script from there. Also, the disk volume where you have the install directory must have at least 8GB free space for the installation.

   ```
   $ mkdir ~/autoid-config
   ```

Install Docker on the Deployer Machine

Install Docker on the deployer machine. We run commands from this machine to install Autonomous Identity on the target machine. In this example, we use CentOS 7.

1. On the target machine, set up the Docker-CE repository.

```
$ sudo yum-config-manager \
    --add-repo https://download.docker.com/linux/centos/docker-ce.repo
```

2. Install the latest version of the Docker CE, the command-line interface, and containerd.io, a containerized website.

   ```
   $ sudo yum install -y docker-ce docker-ce-cli containerd.io
   ```

3. Enable Docker to start at boot.

   ```
   $ sudo systemctl enable docker
   ```

4. Start Docker.

   ```
   $ sudo systemctl start docker
   ```

5. Check that Docker is running.

   ```
   $ systemctl status docker
   ```

6. Add the user to the Docker group.

   ```
   $ sudo usermod -aG docker ${USER}
   ```

7. Reset the privileges on the Docker socket.

   ```
   $ sudo chmod 666 /var/run/docker.sock
   ```

Set Up SSH on the Deployer

1. On the deployer machine, change to the SSH directory.

   ```
   $ cd ~/.ssh
   ```

2. Run **ssh-keygen** to generate an RSA keypair, and then click Enter. You can use the default filename. Enter a password for protecting your private key.

   ```
   $ ssh-keygen -t rsa -C "autoid"
   ```

   The public and private rsa key pair is stored in `home-directory/.ssh/id_rsa` and `home-directory/.ssh/id_rsa.pub`.

3. Copy the SSH key to the `autoid-config` directory.

   ```
    $ cp id_rsa ~/autoid-config
   ```

4. Change the privileges and owner to the file.

   ```
   $ chmod 400 ~/autoid-config/id_rsa
   ```

5. Copy your public SSH key, `id_rsa.pub`, to the target machine's `~/.ssh/authorized_keys` file.

> **Note**
>
> If your target system does not have an `/authorized_keys` directory, create it using **mkdir -p ~/.ssh/authorized_keys**.

```
$ ssh-copy-id -i id_rsa.pub autoid@<Target IP Address>
```

6. On the deployer machine, test your SSH connection to the target machine. This is a critical step. Make sure the connection works before proceeding with the installation.

   ```
   $ ssh -i ~/.ssh/id_rsa autoid@lt;Target IP Address>
   Last login: Wed Sep 23 14:06:06 2020
   ```

7. While SSH'ing into the target node, set the privileges on your `~/.ssh` and `~/.ssh/authorized_keys`.

   ```
   $ chmod 700 ~/.ssh && chmod 600 ~/.ssh/authorized_keys
   ```

8. If you successfully accessed the remote server and changed the privileges on the `~/.ssh`, enter **exit** to end your SSH session.

Install Autonomous Identity

1. On the deployer machine, change to the installation directory.

   ```
   $ cd ~/autoid-config
   ```

2. Log in to the ForgeRock Google Cloud Registry (gcr.io) using the registry key. The registry key is only available to ForgeRock Autonomous Identity customers. For specific instructions on obtaining the registry key, see How To Configure Service Credentials (Push Auth, Docker) in Backstage.

   ```
   $ docker login -u _json_key -p "$(cat autoid_registry_key.json)" https://gcr.io/forgerock-autoid
   ```

   You should see:

   ```
   Login Succeeded
   ```

3. Run the **create-template** command to generate the `deployer.sh` script wrapper and configuration files. Note that the command sets the configuration directory on the target node to `/config`. The **--user** parameter eliminates the need to use **sudo** while editing the hosts file and other configuration files.

```
$ docker run --user=`id -u` -v ~/autoid-config:/config -it gcr.io/forgerock-autoid/deployer:2020.10.2
 create-template
     ...
d6c7c6f3303e: Pull complete
Digest: sha256:15225be65417f8bfb111adea37c83eb5e0d87140ed498bfb624a358f43fb48bf
Status: Downloaded newer image for gcr.io/forgerock-autoid/autoid/dev-compact/
deployer@sha256:15225be65417f8bfb111a
dea37c83eb5e0d87140ed498bfb624a358f43fb48bf
Config template is copied to host machine directory mapped to /config
```

4. Make the script executable.

```
$ chmod +x deployer.sh
```

5. To see the list of commands, enter `deployer.sh`.

```
$ ./deployer.sh
Usage: deployer <command>

Commands:
  create-template
  download-images
  import-deployer
  encrypt-vault
  decrypt-vault
  run
  create-tar
  install-docker
  install-dbutils
  upgrade
```

Configure Autonomous Identity

The **create-template** command from the previous section creates a number of configuration files, required for the deployment.

1. On the deployer machine, open a text editor and edit the `ansible.cfg` to set up the remote user and SSH private key file location on the target node. Make sure that the `remote_user` exists on the target node and that the deployer machine can ssh to the target node as the user specified in the `id_rsa` file. In most cases, you can use the default values.

```
[defaults]
host_key_checking = False
remote_user = autoid
private_key_file = id_rsa
```

2. On the deployer machine, open a text editor and edit the `~/autoid-config/vars.yml` file to configure specific settings for your deployment:

    a. **Domain and Target Environment**. Set the domain name and target environment specific to your deployment by editing the `/autoid-config/vars.xml` file. By default, the domain name is set to forgerock.com and the target environment is set to autoid. The default Autonomous Identity URL will be: https://autoid-ui.forgerock.com. For this example, we use the default values.

```
domain_name: forgerock.com
target_environment: autoid
```

If you change the domain name and target environment, you need to also change the certificates to reflect the new changes. For more information, see "*Customize the Domain and Namespace*" in the *Admin Guide*.

b. **Analytics Data Directory and Analytics Configuration Direction**. Although rarely necessary for a single node deployment, you can change the analytics and analytics configuration mount directories by editing the properties in the ~/autoid-config/vars.yml file.

```
analytics_data_dir: /data
analytics_conf_dif: /data/conf
```

c. **Dark Theme Mode**. Optional. By default, the Autonomous Identity UI displays its pages with a light background. You can set a dark theme mode by setting the enable_dark_theme property to true.

d. **Database Type**. By default, Apache Cassandra is set as the default database for Autonomous Identity. For MongoDB, set the db_driver_type: to mongo.

```
db_driver_type: mongo
```

e. **Private IP Address Mapping**. If your external and internal IP addresses are different, for example, when deploying the target host in a cloud, define a mapping between the external IP address and the private IP address in the ~/autoid-config/vars.yml file.

If your external and internal IP addresses are the same, you can skip this step.

On the deployer node, add the private_ip_address_mapping property in the ~/autoid-config/vars.yml file. You can look up the private IP on the cloud console, or run **sudo ifconfig** on the target host. Make sure the values are within double quotes. The key should not be in double quotes and should have two spaces preceding the IP address.

```
private_ip_address_mapping:
  external_ip:  "internal_ip"
```

For example:

```
private_ip_address_mapping:
  34.70.190.144:  "10.128.0.71"
```

f. **Authentication Option**. Autonomous Identity provides a single sign-on (SSO) feature that you can configure with an OIDC identity provider.

g. **JWT Expiry and Secret File**. Optional. By default, the session JWT is set at 30 minutes. To change this value, set the jwt_expiry property to a different value.

```
jwt_expiry: "30 minutes"
```

h. **Elasticsearch Heap Size**. Optional. The default heap size for Elasticsearch is 1GB, which may be small for production. For production deployments, uncomment the option and specify `2G` or `3G`.

```
#elastic_heap_size: 1g   # sets the heap size (1g|2g|3g) for the Elastic Servers
```

i. **OpenLDAP**. Optional. Autonomous Identity installs an OpenLDAP Docker image on the target server to hold user data. Administrators can add or remove users or change their group privileges using the **phpldapadmin** command. You can customize your OpenLDAP domain, base DN, and URL to match your company's environment. For more information, see "*Configuring LDAP*" in the *Admin Guide*.

3. Open a text editor and enter the target host's public IP addresses in the `~/autoid-config/hosts` file. Make sure the target machine's external IP address is accessible from the deployer machine.

+ *Click to See a Host File for Cassandra Deployments*

If you configured Cassandra as your database, the `~/autoid-config/hosts` file is as follows for single-node target deployments:

```
[docker-managers]
34.70.190.144

[docker-workers]
34.70.190.144

[docker:children]
docker-managers
docker-workers

[cassandra-seeds]
34.70.190.144

[cassandra-workers]
34.70.190.144

[spark-master]
34.70.190.144

[spark-workers]
34.70.190.144

[analytics]
34.70.190.144

[mongo_master]

[mongo_replicas]

[mongo:children]
mongo_replicas
mongo_master

# ELastic Nodes
```

```
[odfe-master-node]
34.70.190.144

[odfe-data-nodes]

[kibana-node]
34.70.190.144
```

+ *Click to See a Host File for MongoDB Deployments*

If you configured MongoDB as your database, the `~/autoid-config/hosts` file is as follows for single-node target deployments:

```
[docker-managers]
34.70.190.144

[docker-workers]
34.70.190.144

[docker:children]
docker-managers
docker-workers

[cassandra-seeds]

[cassandra-workers]

[spark-master]
34.70.190.144

[spark-workers]
34.70.190.144

[analytics]
34.70.190.144

[mongo_master]
34.70.190.144  mongodb_master=True

[mongo_replicas]
34.70.190.144

[mongo:children]
mongo_replicas
mongo_master

# ELastic Nodes
[odfe-master-node]
34.70.190.144

[odfe-data-nodes]

[kibana-node]
```

```
34.70.190.144
```

4. Open a text editor and set the Autonomous Identity passwords for the configuration service, LDAP backend, and Cassandra database. The vault passwords file is located at `~/autoid-config/vault.yml`.

   > **Note**
   >
   > Do not include special characters & or $ in `vault.yml` passwords as it will result in a failed deployer process.

   ```
   configuration_service_vault:
     basic_auth_password: Welcome123

   openldap_vault:
     openldap_password: Welcome123

   cassandra_vault:
     cassandra_password: Welcome123
     cassandra_admin_password: Welcome123

   mongo_vault:
     mongo_admin_password: Welcome123
     mongo_root_password: Welcome123

   elastic_vault:
     elastic_admin_password: Welcome123
     elasticsearch_password: Welcome123
   ```

5. Encrypt the vault file that stores the Autonomous Identity passwords, located at `~/autoid-config/vault.yml`. The encrypted passwords will be saved to `/config/.autoid_vault_password`. The `/config/` mount is internal to the deployer container.

   ```
   $ ./deployer.sh encrypt-vault
   ```

6. Download the images. This step downloads software dependencies needed for the deployment and places them in the `autoid-packages` directory.

   ```
   $ ./deployer.sh download-images
   ```

   Make sure you have no failed processes before proceeding to the next step

   ```
   PLAY RECAP *****************************************************************************
   localhost  : ok=24  changed=17  unreachable=0  failed=0  skipped=0  rescued=0  ignored=0
   ```

7. Run the deployment. The command installs the packages, and starts the microservices and the analytics service. Make sure you have no failed processes before proceeding to the next step.

   ```
   $ ./deployer.sh run
   ```

   Make sure you have no failed processes before proceeding to the next step.

```
PLAY RECAP ***********************************************************************************
34.70.190.144  : ok=450  changed=236  unreachable=0  failed=0  skipped=30  rescued=0  ignored=2
localhost      : ok=11   changed=5   unreachable=0  failed=0  skipped=6   rescued=0  ignored=0
```

Resolve Hostname

After installing Autonomous Identity, set up the hostname resolution for your deployment.

Resolve the hostname:

1. Configure your DNS servers to access Autonomous Identity dashboard and self-service applications on the target node. The following domain names must resolve to the IP address of the target node: `<target-environment>-ui.<domain-name>` and `<target-environment>-selfservice.<domain-name>`.

2. If DNS cannot resolve target node hostname, edit it locally on the machine that you want to access Autonomous Identity using a browser. Open a text editor and add an entry in the `/etc/hosts` file for the self-service and UI services for each managed target node.

   ```
   target-ip-address  <target-environment>-ui.<domain-name> <target-environment>-selfservice.<domain-name>
   ```

   For example:

   ```
   34.70.190.144  autoid-ui.forgerock.com autoid-selfservice.forgerock.com
   ```

3. If you set up a custom domain name and target environment, add the entries in `/etc/hosts`. For example:

   ```
   34.70.190.144  myid-ui.abc.com  myid-selfservice.abc.com
   ```

   For more information on customizing your domain name, see "*Customize the Domain and Namespace*" in the *Admin Guide*.

Access the Dashboard

Access the Autonomous Identity console UI:

1. Open a browser, and point it to `https://autoid-ui.forgerock.com/` (or your customized URL: `https://myid-ui.abc.com`).

2. Log in as a test user: **bob.rodgers@forgerock.com**. Enter the password: **Welcome123**.

Check Apache Cassandra

Check Cassandra:

1. On the target node, check the status of Apache Cassandra.

   ```
   $ /opt/autoid/apache-cassandra-3.11.2/bin/nodetool status
   ```

2. An example output is as follows:

   ```
   Datacenter: datacenter1
   =======================
   Status=Up/Down
   |/ State=Normal/Leaving/Joining/Moving
   --  Address       Load       Tokens       Owns (effective)  Host ID                               Rack
   UN  34.70.190.144 1.33 MiB   256          100.0%            a10a91a4-96e83dd-85a2-4f90d19224d9  rack1
   ```

Check MongoDB

Check the status of MongoDB:

- On the target node, check the status of MongoDB.

  ```
  $ mongo --tls --host <Host IP Address> --tlsCAFile /opt/autoid/mongo/certs/rootCA.pem   --
  tlsAllowInvalidCertificates  --tlsCertificateKeyFile /opt/autoid/mongo/certs/mongodb.pem
  ```

Check Apache Spark

Check Spark:

- SSH to the target node and open Spark dashboard using the bundled text-mode web browser

  ```
  $ elinks http://localhost:8080
  ```

  You should see Spark Master status as ALIVE and worker(s) with State ALIVE.

  + *Click to See an Example of the Spark Dashboard*

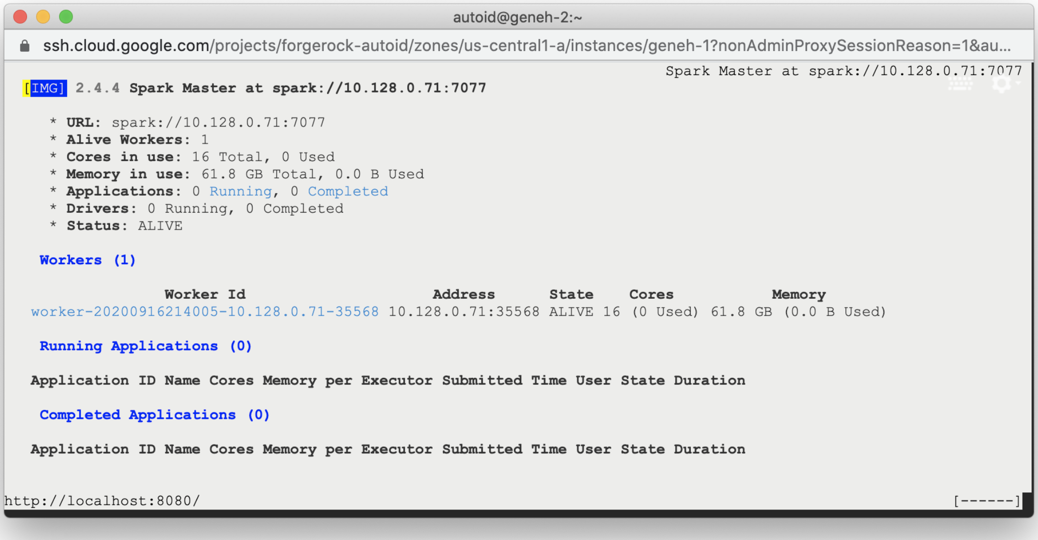
autoid@geneh-2:~
ssh.cloud.google.com/projects/forgerock-autoid/zones/us-central1-a/instances/geneh-1?nonAdminProxySessionReason=1&au...

```
                                                        Spark Master at spark://10.128.0.71:7077
[IMG] 2.4.4 Spark Master at spark://10.128.0.71:7077

   * URL: spark://10.128.0.71:7077
   * Alive Workers: 1
   * Cores in use: 16 Total, 0 Used
   * Memory in use: 61.8 GB Total, 0.0 B Used
   * Applications: 0 Running, 0 Completed
   * Drivers: 0 Running, 0 Completed
   * Status: ALIVE

 Workers (1)

               Worker Id                     Address      State    Cores          Memory
worker-20200916214005-10.128.0.71-35568 10.128.0.71:35568 ALIVE 16 (0 Used) 61.8 GB (0.0 B Used)

 Running Applications (0)

Application ID Name Cores Memory per Executor Submitted Time User State Duration

 Completed Applications (0)

Application ID Name Cores Memory per Executor Submitted Time User State Duration


http://localhost:8080/                                                                [------]
```

Access Self-Service

The self-service feature lets Autonomous Identity users change their own passwords.

Access self-service:

- Open a browser and point it to: `https://autoid-selfservice.forgerock.com/`.

  + *Click to See an Example of the Self-Service Dashboard*

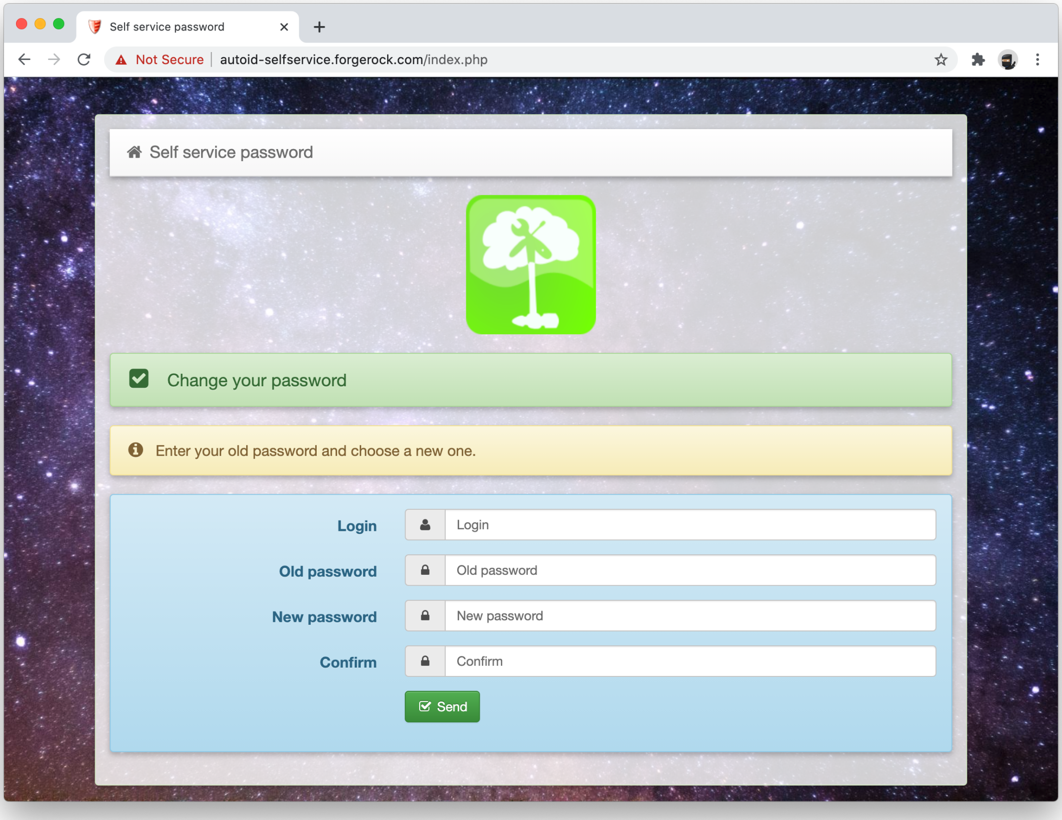

Self service password

☑ Change your password

Enter your old password and choose a new one.

Login	Login
Old password	Old password
New password	New password
Confirm	Confirm

Start the Analytics

If the previous steps all check out successfully, you can start an analytics pipeline run, where association rules, confidence scores, predications, and recommendations are generated. Autonomous Identity provides a small demo data set that lets you run the analytics pipeline on. Note for production runs, prepare your company's dataset as outlined in "*Data Preparation*" in the *Admin Guide*.

Start the analytics service:

- Run the analytics pipeline commands. This may take a bit longer than the install, depending on the size of your dataset. For specific information, see "*Run the Analytics Pipeline*" in the *Admin Guide*.

Chapter 5

Install a Single Node Air-Gap Target

This chapter presents instructions on deploying Autonomous Identity in a single-node target machine that has no Internet connectivity. This type of configuration, called an *air-gap* or *offline* deployment, provides enhanced security by isolating itself from outside Internet or network access.

The air-gap installation is similar to that of the single-node target deployment with Internet connectivity, except that the image and deployer script must be saved on a portable media, such as USB drive or drive, and copied to the air-gapped target machine.

Figure 9: A single-node air-gapped target deployment.

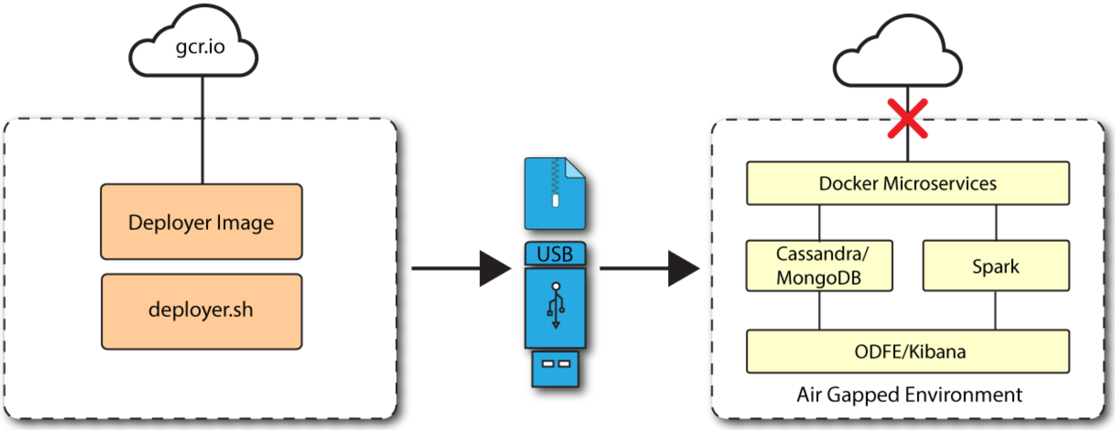

Let's deploy Autonomous Identity on a single-node air-gapped target on CentOS 7. The following are prerequisites:

- **Operating System**. The target machine requires CentOS 7. The deployer machine can use any operating system as long as Docker is installed. For this guide, we use CentOS 7 as its base operating system.
- **Memory Requirements**. Make sure you have enough free disk space on the deployer machine before running the `deployer.sh` commands. We recommend at least a 40GB/partition with 14GB used and 27GB free after running the commands.
- **Default Shell**. The default shell for the `autoid` user must be bash.

- **Deployment Requirements**. Autonomous Identity provides a Docker image that creates a `deployer.sh` script. The script downloads additional images necessary for the installation. To download the deployment images, you must first obtain a registry key to log into the ForgeRock Google Cloud Registry (gcr.io). The registry key is only available to ForgeRock Autonomous Identity customers. For specific instructions on obtaining the registry key, see How To Configure Service Credentials (Push Auth, Docker) in Backstage.
- **Database Requirements**. Decide which database you are using: Apache Cassandra or MongoDB.
- **IPv4 Forwarding**. Many high-security environments run their CentOS-based systems with IPv4 forwarding disabled. However, Docker Swarm does not work with a disabled IPv4 forwarding setting. In such environments, make sure to enable IPv4 forwarding in the file `etc/sysctl.conf`:

  ```
  net.ipv4.ip_forward=1
  ```

Set Up the Deployer Machine

Set up the deployer on an Internet-connect machine.

1. The install assumes that you have CentOS 7 as your operating system. Check your CentOS 7 version.

   ```
   $ sudo cat /etc/centos-release
   ```

2. Set the user for the target machine to a username of your choice. For example, `autoid`.

   ```
   $ sudo adduser autoid
   ```

3. Set the password for the user you created in the previous step.

   ```
   $ sudo passwd autoid
   ```

4. Configure the user for passwordless sudo.

   ```
   $ echo "autoid  ALL=(ALL)  NOPASSWD:ALL" | sudo tee /etc/sudoers.d/autoid
   ```

5. Add administrator privileges to the user.

   ```
   $ sudo usermod -aG wheel autoid
   ```

6. Change to the user account.

   ```
   $ su - autoid
   ```

7. Install yum-utils package on the deployer machine. yum-utils is a utilities manager for the Yum RPM package repository. The repository compresses software packages for Linux distributions.

   ```
   $ sudo yum install -y yum-utils
   ```

8. Create the installation directory. Note that you can use any install directory for your system as long as your run the **deployer.sh** script from there. Also, the disk volume where you have the install directory must have at least 8GB free space for the installation.

```
$ mkdir ~/autoid-config
```

Install Docker on the Deployer Machine

1. On the target machine, set up the Docker-CE repository.

   ```
   $ sudo yum-config-manager \
        --add-repo https://download.docker.com/linux/centos/docker-ce.repo
   ```

2. Install the latest version of the Docker CE, the command-line interface, and containerd.io, a containerized website.

   ```
   $ sudo yum install -y docker-ce docker-ce-cli containerd.io
   ```

3. Enable Docker to start at boot.

   ```
   $ sudo systemctl enable docker
   ```

4. Start Docker.

   ```
   $ sudo systemctl start docker
   ```

5. Check that Docker is running.

   ```
   $ systemctl status docker
   ```

6. Add the user to the Docker group.

   ```
   $ sudo usermod -aG docker ${USER}
   ```

7. Reset the privileges on the Docker socket.

   ```
   $ sudo chmod 666 /var/run/docker.sock
   ```

Set Up SSH on the Deployer

While SSH is not necessary to connect the deployer to the target node as the machines are isolated from one another. You still need SSH on the deployer so that it can communicate with itself.

1. On the deployer machine, run **ssh-keygen** to generate an RSA keypair, and then click Enter. You can use the default filename. Enter a password for protecting your private key.

   ```
   $ ssh-keygen -t rsa -C "autoid"
   ```

   The public and private rsa key pair is stored in `home-directory/.ssh/id_rsa` and `home-directory/.ssh/id_rsa.pub`.

2. Copy the SSH key to the `~/autoid-config` directory.

   ```
    $ cp ~/.ssh/id_rsa ~/autoid-config
   ```

3. Change the privileges to the file.

   ```
   $ chmod 400 ~/autoid-config/id_rsa
   ```

Prepare the Tar File

Run the following steps on an Internet-connected host machine:

1. On the deployer machine, change to the installation directory.

   ```
   $ cd ~/autoid-config/
   ```

2. Log in to the ForgeRock Google Cloud Registry (gcr.io) using the registry key. The registry key is only available to ForgeRock Autonomous Identity customers. For specific instructions on obtaining the registry key, see How To Configure Service Credentials (Push Auth, Docker) in Backstage.

   ```
   $ docker login -u _json_key -p "$(cat autoid_registry_key.json)" https://gcr.io/forgerock-autoid
   ```

   You should see:

   ```
   Login Succeeded
   ```

3. Run the **create-template** command to generate the `deployer.sh` script wrapper. Note that the command sets the configuration directory on the target node to `/config`. Note that the **--user** parameter eliminates the need to use **sudo** while editing the hosts file and other configuration files.

   ```
   $ docker run --user=`id -u` -v ~/autoid-config:/config -it gcr.io/forgerock-autoid/deployer:2020.10.2
    create-template
   ```

4. Make the script executable.

   ```
   $ chmod +x deployer.sh
   ```

5. Download the Docker images. This step downloads software dependencies needed for the deployment and places them in the `autoid-packages` directory.

   ```
   $ ./deployer.sh download-images
   ```

6. Create a tar file containing all of the Autonomous Identity binaries.

   ```
   $ tar czf autoid-packages.tgz deployer.sh autoid-packages/*
   ```

7. Copy the `autoid-packages.tgz`, `deployer.sh`, and SSH key (`id_rsa`) to a USB drive or portable hard drive.

Install from the Air-Gap Target

Before you begin, make sure you have CentOS 7 installed on your air-gapped target machine.

1. Create the `~/autoid-config` directory if you haven't already.

   ```
   $ mkdir ~/autoid-config
   ```

2. Copy the `autoid-package.tgz` tar file from the portable storage device.

3. Unpack the tar file.

   ```
   $ tar xf autoid-packages.tgz -C ~/autoid-config
   ```

4. On the air-gap host node, copy the SSH key to the `~/autoid-config` directory.

5. Change the privileges to the file.

   ```
   $ chmod 400 ~/autoid-config/id_rsa
   ```

6. Change to the configuration directory.

   ```
   $ cd ~/autoid-config
   ```

7. Install Docker.

   ```
   $ sudo ./deployer.sh install-docker
   ```

8. Log out and back in.

9. Change to the configuration directory.

   ```
   $ cd ~/autoid-config
   ```

10. Import the deployer image.

    ```
    $ ./deployer.sh import-deployer
    ```

    You should see:

    ```
    ...
    db631c8b06ee: Loading layer [=================================================>]   2.56kB/2.56kB
    2d62082e3327: Loading layer [=================================================>]  753.2kB/753.2kB
    Loaded image: gcr.io/forgerock-autoid/deployer2020.10.1
    ```

11. Create the configuration template using the **create-template** command. This command creates the configuration files: `ansible.cfg`, `vars.yml`, `vault.yml` and `hosts`.

    ```
    $ ./deployer.sh create-template
    ```

    You should see:

    ```
    Config template is copied to host machine directory mapped to /config
    ```

Configure Autonomous Identity Air-Gapped

1. Open a text editor and edit the `ansible.cfg` to set up the target machine user and SSH private key file location on the target node. Make sure that the remote_user exists on the target node and that the deployer machine can ssh to the target node as the user specified in the `id_rsa` file.

   ```
   [defaults]
   host_key_checking = False
   remote_user = autoid
   private_key_file = id_rsa
   ```

2. Open a text editor and edit the `~/autoid-config/vars.yml` file to configure specific settings for your deployment:

   a. **Domain and Target Environment**. Set the domain name and target environment specific to your deployment by editing the `/autoid-config/vars.xml` file. By default, the domain name is set to `forgerock.com` and the target environment is set to `autoid`. The default Autonomous Identity URL will be: `https://autoid-ui.forgerock.com`. For this example, we use the default values.

      ```
      domain_name: forgerock.com
      target_environment: autoid
      ```

      If you change the domain name and target environment, you need to also change the certificates to reflect the new changes. For more information, see "*Customize the Domain and Namespace*" in the *Admin Guide*.

   b. **Analytics Data Directory and Analytics Configuration Direction**. Although rarely necessary for a single node deployment, you can change the analytics and analytics configuration mount directories by editing the properties in the `~/autoid-config/vars.yml` file.

      ```
      analytics_data_dir: /data
      analytics_conf_dif: /data/conf
      ```

   c. **Dark Theme Mode**. Optional. By default, the Autonomous Identity UI displays its pages with a light background. You can set a dark theme mode by setting the `enable_dark_theme` property to `true`.

   d. **Database Type**. By default, Apache Cassandra is set as the default database for Autonomous Identity. For MongoDB, set the `db_driver_type:` to `mongo`.

      ```
      db_driver_type: mongo
      ```

   e. **Private IP Address Mapping**. An air-gap deployment has no external IP addresses, but you may still need to define a mapping in the `~/autoid-config/vars.yml` file, if your internal IP address differs from an external IP, say in a virtual air-gapped configuration.

      If your external and internal IP addresses are the same, you can skip this step.

      Add the `private_ip_address_mapping` property in the `~/autoid-config/vars.yml` file. You can look up the private IP on the cloud console, or run **sudo ifconfig** on the target host. Make sure the values are within double quotes. The key should not be in double quotes and should have two spaces preceding the IP address.

```
private_ip_address_mapping:
  external_ip:  "internal_ip"
```

For example:

```
private_ip_address_mapping:
  34.70.190.144:  "10.128.0.71"
```

f. **Authentication Option**. Autonomous Identity provides a single sign-on (SSO) feature that you can configure with an OIDC identity provider.

g. **JWT Expiry and Secret File**. Optional. By default, the session JWT is set at 30 minutes. To change this value, set the `jwt_expiry` property to a different value.

```
jwt_expiry: "30 minutes"
```

h. **Elasticsearch Heap Size**. Optional. The default heap size for Elasticsearch is 1GB, which may be small for production. For production deployments, uncomment the option and specify `2G` or `3G`.

```
#elastic_heap_size: 1g   # sets the heap size (1g|2g|3g) for the Elastic Servers
```

i. **OpenLDAP**. Optional. Autonomous Identity installs an OpenLDAP Docker image on the target server to hold user data. Administrators can add or remove users or change their group privileges using the **phpldapadmin** command. You can customize your OpenLDAP domain, base DN, and URL to match your company's environment. For more information, see "*Configuring LDAP*" in the *Admin Guide*.

3. Open a text editor and enter the target host's private IP addresses in the `~/autoid-config/hosts` file. The following is an example of the `~/autoid-config/hosts` file:

+ *Click to See a Host File for Cassandra Deployments*

If you configured Cassandra as your database, the `~/autoid-config/hosts` file is as follows for single-node air-gapped target deployment:

```
[docker-managers]
10.128.0.34

[docker-workers]
10.128.0.34

[docker:children]
docker-managers
docker-workers

[cassandra-seeds]
10.128.0.34

[cassandra-workers]
```

```
10.128.0.34

[spark-master]
10.128.0.34

[spark-workers]
10.128.0.34

[analytics]
10.128.0.34

[mongo_master]
#ip# mongodb_master=True

[mongo_replicas]

[mongo:children]
mongo_replicas
mongo_master

# ELastic Nodes
[odfe-master-node]
10.128.0.34

[odfe-data-nodes]

[kibana-node]
10.128.0.34
```

+ *Click to See a Host File for MongoDB Deployments*

If you configured MongoDB as your database, the `~/autoid-config/hosts` file is as follows for single-node air-gapped target deployment:

```
[docker-managers]
10.128.0.34

[docker-workers]
10.128.0.34

[docker:children]
docker-managers
docker-workers

[cassandra-seeds]

[cassandra-workers]

[spark-master]
10.128.0.34

[spark-workers]
10.128.0.34

[analytics]
10.128.0.34
```

```
[mongo_master]
10.128.0.34  mongodb_master=True

[mongo_replicas]
10.128.0.34

[mongo:children]
mongo_replicas
mongo_master

# ELastic Nodes
[odfe-master-node]
10.128.0.34

[odfe-data-nodes]

[kibana-node]
10.128.0.34
```

4. Set the Autonomous Identity passwords, located at `~/autoid-config/vault.yml`.

> **Note**
>
> Do not include special characters & or $ in `vault.yml` passwords as it will result in a failed deployer process.

```
configuration_service_vault:
  basic_auth_password: Welcome123

openldap_vault:
  openldap_password: Welcome123

cassandra_vault:
  cassandra_password: Welcome123
  cassandra_admin_password: Welcome123

mongo_vault:
  mongo_admin_password: Welcome123
  mongo_root_password: Welcome123

elastic_vault:
  elastic_admin_password: Welcome123
  elasticsearch_password: Welcome123
```

5. Encrypt the vault file that stores the Autonomous Identity passwords, located at `~/autoid-config/vault.yml`. The encrypted passwords will be saved to `/config/.autoid_vault_password`. The `/config/` mount is internal to the deployer container.

```
$ ./deployer.sh encrypt-vault
```

6. Run the deployment.

```
$ ./deployer.sh run
```

Resolve Hostname

After installing Autonomous Identity, set up the hostname resolution for your deployment.

Resolve the hostname:

1. Configure your DNS servers to access Autonomous Identity dashboard and self-service applications on the target node. The following domain names must resolve to the IP address of the target node: `<target-environment>-ui.<domain-name>` and `<target-environment>-selfservice.<domain-name>`.

2. If DNS cannot resolve target node hostname, edit it locally on the machine that you want to access Autonomous Identity using a browser. Open a text editor and add an entry in the `/etc/hosts` file for the self-service and UI services for each managed target node.

   ```
   target-ip-address  <target-environment>-ui.<domain-name> <target-environment>-selfservice.<domain-name>
   ```

   For example:

   ```
   34.70.190.144  autoid-ui.forgerock.com autoid-selfservice.forgerock.com
   ```

3. If you set up a custom domain name and target environment, add the entries in `/etc/hosts`. For example:

   ```
   34.70.190.144  myid-ui.abc.com  myid-selfservice.abc.com
   ```

   For more information on customizing your domain name, see "*Customize the Domain and Namespace*" in the *Admin Guide*.

Access the Dashboard

Access the Autonomous Identity console UI:

1. Open a browser, and point it to `https://autoid-ui.forgerock.com/` (or your customized URL: `https://myid-ui.abc.com`).

2. Log in as a test user: **bob.rodgers@forgerock.com**. Enter the password: **Welcome123**.

Check Apache Cassandra

Check Cassandra:

1. On the target node, check the status of Apache Cassandra.

   ```
   $ /opt/autoid/apache-cassandra-3.11.2/bin/nodetool status
   ```

2. An example output is as follows:

```
Datacenter: datacenter1
=======================
Status=Up/Down
|/ State=Normal/Leaving/Joining/Moving
--  Address       Load       Tokens       Owns (effective)  Host ID                               Rack
UN  34.70.190.144 1.33 MiB   256          100.0%            a10a91a4-96e83dd-85a2-4f90d19224d9  rack1
```

Check MongoDB

Check the status of MongoDB:

- On the target node, check the status of MongoDB.

```
$ mongo --tls --host <Host IP Address> --tlsCAFile /opt/autoid/mongo/certs/rootCA.pem   --
tlsAllowInvalidCertificates  --tlsCertificateKeyFile /opt/autoid/mongo/certs/mongodb.pem
```

Check Apache Spark

Check Spark:

- SSH to the target node and open Spark dashboard using the bundled text-mode web browser

```
$ elinks http://localhost:8080
```

You should see Spark Master status as ALIVE and worker(s) with State ALIVE.

+ *Click to See an Example of the Spark Dashboard*

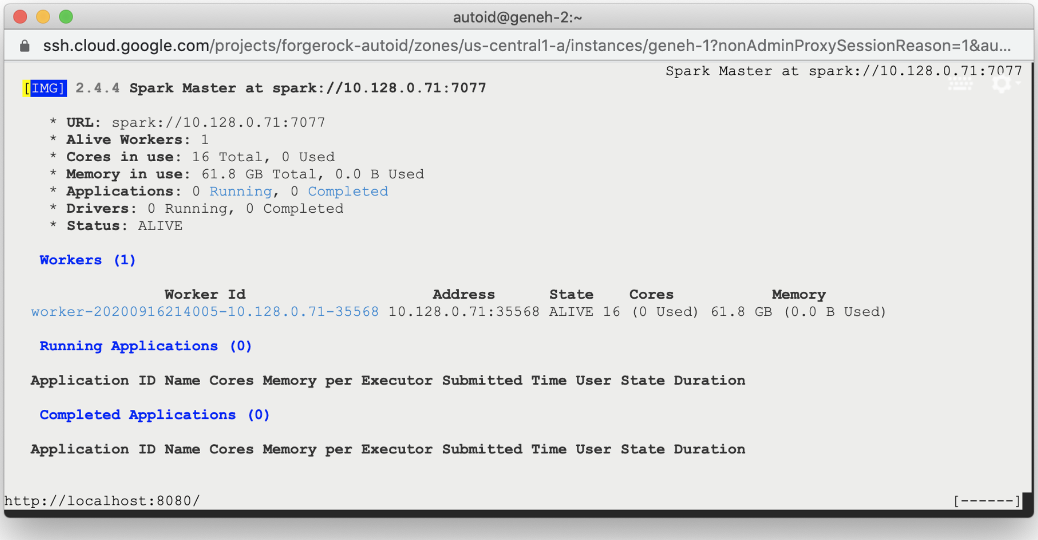

autoid@geneh-2:~
ssh.cloud.google.com/projects/forgerock-autoid/zones/us-central1-a/instances/geneh-1?nonAdminProxySessionReason=1&au...

```
                                                                   Spark Master at spark://10.128.0.71:7077
[IMG] 2.4.4 Spark Master at spark://10.128.0.71:7077

   * URL: spark://10.128.0.71:7077
   * Alive Workers: 1
   * Cores in use: 16 Total, 0 Used
   * Memory in use: 61.8 GB Total, 0.0 B Used
   * Applications: 0 Running, 0 Completed
   * Drivers: 0 Running, 0 Completed
   * Status: ALIVE

 Workers (1)

               Worker Id                      Address      State    Cores          Memory
worker-20200916214005-10.128.0.71-35568 10.128.0.71:35568 ALIVE 16 (0 Used) 61.8 GB (0.0 B Used)

 Running Applications (0)

Application ID Name Cores Memory per Executor Submitted Time User State Duration

 Completed Applications (0)

Application ID Name Cores Memory per Executor Submitted Time User State Duration

http://localhost:8080/                                                                              [------]
```

Access Self-Service

The self-service feature lets Autonomous Identity users change their own passwords.

Access self-service:

- Open a browser and point it to: `https://autoid-selfservice.forgerock.com/`.

  + *Click to See an Example of the Self-Service Dashboard*

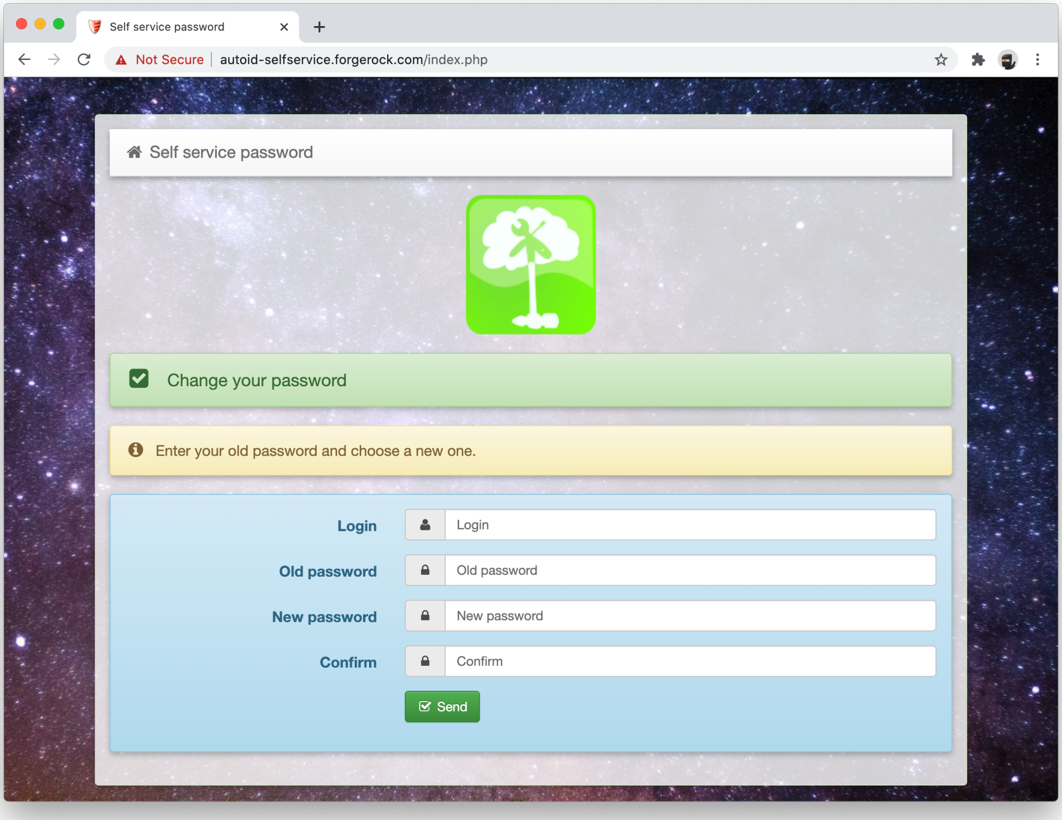

Change your password

Enter your old password and choose a new one.

Login: Login

Old password: Old password

New password: New password

Confirm: Confirm

Start the Analytics

If the previous steps all check out successfully, you can start an analytics pipeline run, where association rules, confidence scores, predications, and recommendations are generated. Autonomous Identity provides a small demo data set that lets you run the analytics pipeline on. Note for production runs, prepare your company's dataset as outlined in "*Data Preparation*" in the *Admin Guide*.

Start the analytics service:

- Run the analytics pipeline commands. This may take a bit longer than the install, depending on the size of your dataset. For specific information, see "*Run the Analytics Pipeline*" in the *Admin Guide*.

Chapter 6
Install a Multi-Node Deployment

This chapter presents instructions on deploying Autonomous Identity in a multi-node target deployment that has Internet connectivity. ForgeRock provides a deployer script that pulls a Docker container image from ForgeRock's Google Cloud Registry (gcr.io) repository. The image contains the microservices, analytics, and backend databases needed for the system.

This installation assumes that you set up the deployer on a separate machine from the target.

The deployment depends on how the network is configured. You could have a Docker cluster with multiple Spark nodes and Cassandra or MongoDB nodes. The key is to determine the IP addresses of each node, which the deployer uses to set up the overlay network for your multinode system.

Figure 10: A multi-node deployment.

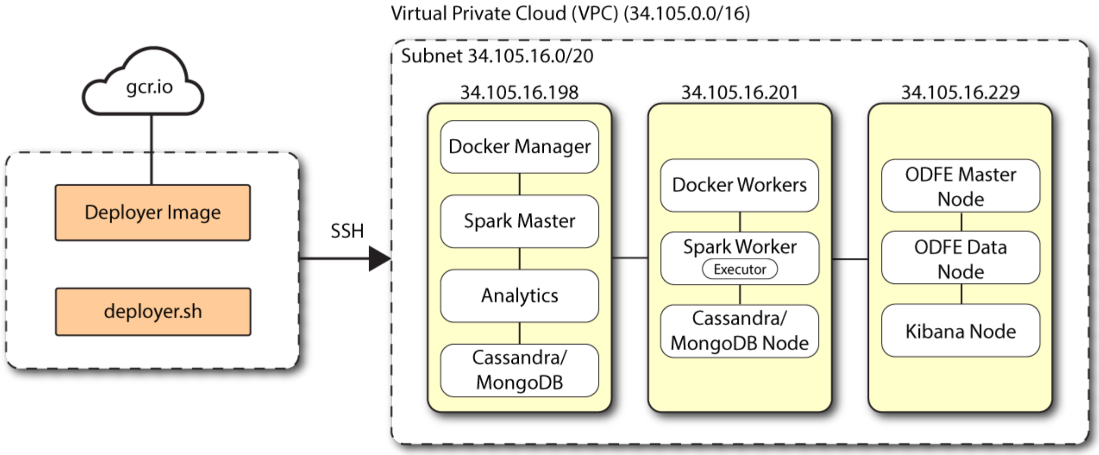

Let's deploy Autonomous Identity on a multi-node target on CentOS 7. The following are prerequisites:

- **Operating System**. The target machine requires CentOS 7. The deployer machine can use any operating system as long as Docker is installed. For this guide, we use CentOS 7 as its base operating system.
- **Memory Requirements**. Make sure you have enough free disk space on the deployer machine before running the `deployer.sh` commands. We recommend at least a 40GB/partition with 14GB used and 27GB free after running the commands.

- **Default Shell**. The default shell for the `autoid` user must be bash.
- **Subnet Requirements**. We recommend deploying your multinode instances within the same subnet. Ports must be open for the installation to succeed. Each instance should be able to communicate to the other instances.

  > **Important**
  >
  > If any hosts used for the Docker cluster (docker-managers, docker-workers) have an IP address in the range of 10.0.x.x/16, they will conflict with the Swarm network. As a result, the services in the cluster will not connect to the Cassandra database or Elasticsearch backend.
  >
  > The Docker cluster hosts must be in a subnet that provides IP addresses 10.10.1.x or higher.

- **Deployment Requirements**. Autonomous Identity provides a Docker image that creates a `deployer.sh` script that downloads and installs the images necessary. To download the deployment images, you must first obtain a registry key to log into the ForgeRock Google Cloud Registry (gcr.io). The registry key is only available to ForgeRock Autonomous Identity customers. For specific instructions on obtaining the registry key, see How To Configure Service Credentials (Push Auth, Docker) in Backstage.
- **Filesystem Requirements**. Autonomous Identity requires a shared filesystem accessible from the Spark master, Spark worker, analytics hosts, and application layer. The shared filesystem should be mounted at the same mount directory on all of those hosts. If the mount directory for the shared filesystem is different from the default, `/data`, update the `/autoid-config/vars.yml` file to point to the correct directories:

  ```
  analytics_data_dir: /data
  analytics_conf_dif: /data/conf
  ```

- **Architecture Requirements**. Make sure that the analytics server is on the same node as the Spark master.
- **Database Requirements**. Decide which database you are using: Apache Cassandra or MongoDB. The configuration procedure is slightly different for each database.
- **Deployment Best-Practice**. For best performance, dedicate a separate node to Elasticsearch, data nodes, and Kibana.
- **IPv4 Forwarding**. Many high-security environments run their CentOS-based systems with IPv4 forwarding disabled. However, Docker Swarm does not work with a disabled IPv4 forwarding setting. In such environments, make sure to enable IPv4 forwarding in the file `etc/sysctl.conf`:

  ```
  net.ipv4.ip_forward=1
  ```

Set Up the Target Nodes

Make sure you have sufficient storage for your particular deployment. For more information on sizing considerations, see *Deployment Planning Guide*.

For each target node, run the following commands.

1. The install assumes that you have CentOS 7 as your operating system. Check your CentOS 7 version.

   ```
   $ sudo cat /etc/centos-release
   ```

2. Set the user for the target machine to a username of your choice. For example, `autoid`.

   ```
   $ sudo adduser autoid
   ```

3. Set the password for the user you created in the previous step.

   ```
   $ sudo passwd autoid
   ```

4. Configure the user for passwordless sudo.

   ```
   $ echo "autoid  ALL=(ALL)  NOPASSWD:ALL" | sudo tee /etc/sudoers.d/autoid
   ```

5. Add administrator privileges to the user.

   ```
   $ sudo usermod -aG wheel autoid
   ```

6. Change to the user account.

   ```
   $ su - autoid
   ```

7. Install yum-utils package on the deployer machine. yum-utils is a utilities manager for the Yum RPM package repository. The repository compresses software packages for Linux distributions.

   ```
   $ sudo yum install -y yum-utils
   ```

Set Up the Deployer Machine

Set up another machine as a deployer node. You can use any OS-based machine for the deployer as long as it has Docker installed. For this example, we use CentOS 7.

1. The install assumes that you have CentOS 7 as your operating system. Check your CentOS 7 version.

   ```
   $ sudo cat /etc/centos-release
   ```

2. Set the user for the target machine to a username of your choice. For example, `autoid`.

   ```
   $ sudo adduser autoid
   ```

3. Set the password for the user you created in the previous step.

   ```
   $ sudo passwd autoid
   ```

4. Configure the user for passwordless sudo.

   ```
   $ echo "autoid  ALL=(ALL)  NOPASSWD:ALL" | sudo tee /etc/sudoers.d/autoid
   ```

5. Add administrator privileges to the user.

   ```
   $ sudo usermod -aG wheel autoid
   ```

6. Change to the user account.

   ```
   $ su - autoid
   ```

7. Install yum-utils package on the deployer machine. yum-utils is a utilities manager for the Yum RPM package repository. The repository compresses software packages for Linux distributions.

   ```
   $ sudo yum install -y yum-utils
   ```

8. Create the installation directory. Note that you can use any install directory for your system as long as your run the **deployer.sh** script from there. Also, the disk volume where you have the install directory must have at least 8GB free space for the installation.

   ```
   $ mkdir ~/autoid-config
   ```

Install Docker on the Deployer Machine

Install Docker on the deployer machine. We run commands from this machine to install Autonomous Identity on the target machine. In this example, we use CentOS 7.

1. On the target machine, set up the Docker-CE repository.

   ```
   $ sudo yum-config-manager \
       --add-repo https://download.docker.com/linux/centos/docker-ce.repo
   ```

2. Install the latest version of the Docker CE, the command-line interface, and containerd.io, a containerized website.

   ```
   $ sudo yum install -y docker-ce docker-ce-cli containerd.io
   ```

3. Enable Docker to start at boot.

   ```
   $ sudo systemctl enable docker
   ```

4. Start Docker.

   ```
   $ sudo systemctl start docker
   ```

5. Check that Docker is running.

   ```
   $ systemctl status docker
   ```

6. Add the user to the Docker group.

   ```
   $ sudo usermod -aG docker ${USER}
   ```

7. Reset the privileges on the Docker socket.

   ```
   $ sudo chmod 666 /var/run/docker.sock
   ```

Set Up SSH on the Deployer

1. On the deployer machine, change to the `~/.ssh` directory.

   ```
   $ cd ~/.ssh
   ```

2. Run **ssh-keygen** to generate an RSA keypair, and then click Enter. You can use the default filename. Enter a password for protecting your private key.

   ```
   $ ssh-keygen -t rsa -C "autoid"
   ```

   The public and private rsa key pair is stored in `home-directory/.ssh/id_rsa` and `home-directory/.ssh/id_rsa.pub`.

3. Copy the SSH key to the `autoid-config` directory.

   ```
    $ cp id_rsa ~/autoid-config
   ```

4. Change the privileges to the file.

   ```
   $ chmod 400 ~/autoid-config/id_rsa
   ```

5. Copy your public SSH key, `id_rsa.pub`, to each of your target machine's `~/.ssh/authorized_keys` file.

   > **Note**
   >
   > If your target system does not have an `/authorized_keys` directory, create it using **mkdir -p ~/.ssh/authorized_keys**.

   For this example, copy the SSH key to each node:

   ```
   $ ssh-copy-id -i id_rsa.pub autoid@lt;Node 1 IP Address>
   ```

   ```
   $ ssh-copy-id -i id_rsa.pub autoid@lt;Node 2 IP Address>
   ```

   ```
   $ ssh-copy-id -i id_rsa.pub autoid@lt;Node 3 IP Address>
   ```

6. On the deployer machine, test your SSH connection to each target machine. This is a critical step. Make sure the connection works before proceeding with the installation.

   If you can successfully SSH to each machine, set the privileges on your `~/.ssh` and `~/.ssh/authorized_keys`.

a. SSH to first node:

```
$ ssh autoid@lt;Node 1 IP Address>
Last login: Sat Oct 3 03:02:40 2020
```

Set the privileges.

```
$ chmod 700 ~/.ssh && chmod 600 ~/.ssh/authorized_keys
```

Enter Exit to end your SSH session.

b. SSH to the second node:

```
$ ssh autoid@lt;Node 2 IP Address>
Last login: Sat Oct  3 03:06:40 2020
```

Set the privileges.

```
$ chmod 700 ~/.ssh && chmod 600 ~/.ssh/authorized_keys
```

Enter Exit to end your SSH session.

c. SSH to the third node:

```
$ ssh autoid@lt;Node 3 IP Address>
Last login: Sat Oct  3 03:10:40 2020
```

Set the privileges.

```
$ chmod 700 ~/.ssh && chmod 600 ~/.ssh/authorized_keys
```

Enter Exit to end your SSH session.

Install Autonomous Identity

Before you begin, make sure you have CentOS 7 installed on your target machine.

1. On the deployer machine, change to the installation directory.

   ```
   $ cd ~/autoid-config/
   ```

2. Log in to the ForgeRock Google Cloud Registry (gcr.io) using the registry key. The registry key is only available to ForgeRock Autonomous Identity customers. For specific instructions on obtaining the registry key, see How To Configure Service Credentials (Push Auth, Docker) in Backstage.

   ```
   $ docker login -u _json_key -p "$(cat autoid_registry_key.json)" https://gcr.io/forgerock-autoid
   ```

   You should see:

   ```
   Login Succeeded
   ```

3. Run the **create-template** command to generate the `deployer.sh` script wrapper. Note that the command sets the configuration directory on the target node to `/config`. Note that the **--user** parameter eliminates the need to use **sudo** while editing the hosts file and other configuration files.

   ```
   $ docker run --user=`id -u` -v ~/autoid-config:/config -it gcr.io/forgerock-autoid/deployer:2020.10.2
    create-template
   ```

4. Make the script executable.

   ```
   $ chmod +x deployer.sh
   ```

5. To see the list of commands, enter `deployer.sh`.

   ```
   $ ./deployer.sh
   Usage: deployer <command>

   Commands:
     create-template
     download-images
     import-deployer
     encrypt-vault
     decrypt-vault
     run
     create-tar
     install-docker
     install-dbutils
     upgrade
   ```

Configure Autonomous Identity

The **create-template** command from the previous section creates a number of configuration files, required for the deployment.

1. The **create-template** commands creates a number of configuration files, including `ansible.cfg`. Open a text editor and edit the `ansible.cfg` to set up the remote user and SSH private key file location on the target node. Make sure that the remote_user exists on the target node and that the deployer machine can ssh to the target node as the user specified in the `id_rsa` file.

   ```
   [defaults]
   host_key_checking = False
   remote_user = autoid
   private_key_file = id_rsa
   ```

2. On the deployer machine, open a text editor and edit the `~/autoid-config/vars.yml` file to configure specific settings for your deployment:

   a. **Domain and Target Environment**. Set the domain name and target environment specific to your deployment by editing the `/autoid-config/vars.xml` file. By default, the domain name is set to `forgerock.com` and the target environment is set to `autoid`. The default Autonomous Identity URL will be: `https://autoid-ui.forgerock.com`. For this example, we use the default values.

```
domain_name: forgerock.com
target_environment: autoid
```

If you change the domain name and target environment, you need to also change the certificates to reflect the new changes. For more information, see "*Customize the Domain and Namespace*" in the *Admin Guide*.

b. **Analytics Data Directory and Analytics Configuration Direction**. For a multi-node Spark deployment, Autonomous Identity requires a shared filesystem accessible from Spark Master, Spark Worker(s), and Analytics hosts. The shared filesystem should be mounted at same mount directory on all of the above hosts. If the mount directory for shared filesystem is different than `/data`, update the following properties in the `vars.yaml` file to point to the correct location:

```
analytics_data_dir: /data
analytics_conf_dif: /data/conf
```

c. **Dark Theme Mode**. Optional. By default, the Autonomous Identity UI displays its pages with a light background. You can set a dark theme mode by setting the `enable_dark_theme` property to `true`.

d. **Database Type**. By default, Apache Cassandra is set as the default database for Autonomous Identity. For MongoDB, set the `db_driver_type:` to `mongo`.

```
db_driver_type: mongo
```

e. **Private IP Address Mapping**. Define a mapping between the external IP and private IP addresses. This occurs when your target host is in a cloud, so that your external and internal IP addresses are different.

For each target node, add the `private_ip_address_mapping` property in the `~/autoid-config/vars.yml` file. You can look up the private IP on the cloud console, or run **sudo ifconfig** on the target host. Make sure the values are within double quotes. The key should not be in double quotes and should have two spaces preceding the IP address.

```
private_ip_address_mapping:
  external_ip:  "internal_ip"
```

For example:

```
private_ip_address_mapping:
  34.105.16.198:  "10.128.0.51"
  34.105.16.201:  "10.128.0.54"
  34.105.16.229:  "10.128.0.71"
```

f. **Authentication Option**. Autonomous Identity provides a single sign-on (SSO) feature that you can configure with an OIDC identity provider.

g. **JWT Expiry and Secret File**. Optional. By default, the session JWT is set at 30 minutes. To change this value, set the `jwt_expiry` property to a different value.

```
jwt_expiry: "30 minutes"
```

h. **MongoDB Configuration**. For MongoDB clusters, enable replication by uncommenting the `mongodb_replication_replset` property.

```
# uncomment below for mongo with replication enabled. Not needed for single node deployments
mongodb_replication_replset: mongors
```

Also, enable a custom key for inter-machine authentication in the clustered nodes.

```
# custom key
# password for inter-process authentication
# please regenerate this file on production environment with
# command 'openssl rand -base64 741'
mongodb_keyfile_content: |
  8pYcxvCqoe89kcp33KuTtKVf5MoHGEFjTnudrq5BosvWRoIxLowmdjrmUpVfAivh
  CHjqM6w0zVBytAxH1lW+7teMYe6eDn2S/O/1YlRRiW57bWU3zjliW3VdguJar5i9
  Z+1a8lI+0S9pWynbv9+Ao0aXFjSJYVxAm/w7DJbVRGcPhsPmExiSBDw8szfQ8PAU
  2hwRl7nqPZZMMR+uQThg/zV9rOzHJmkqZtsO4UJSilG9euLCYrzW2hdoPuCrEDhu
  Vsi5+nwAgYR9dP2oWkmGN1dwRe0ixSIM2UzFgpaXZaMOG6VztmFrlVXh8oFDRGM0
  cGrFHcnGF7oUGfWnI2Cekngk64dHA2qD7WxXPbQ/svn9EfTY5aPw5lXzKA87Ds8p
  KHVFUYvmA6wVsxb/riGLwc+XZlb6M9gqHn1XSpsnYRjF6UzfRcRR2WyCxLZELaqu
  iKxLKB5FYqMBH7Sqg3qBCtE53vZ7T1nefq5RFzmykviYP63Uhu/A2EQatrMnaFPl
  TTG5CaPjob45CBSyMrheYRWKqxdWN93BTgiTW7p0U6RB0/OCUbsVX6IG3I9N8Uqt
  l8Kc+7aOmtUqFkwo8w30prIOjStMrokxNsuK9KTUiPu2cj7gwYQ574vV3hQvQPAr
  hhb9ohKr0zoPQt31iTj0FDkJzPepeuzqeq8F51HB56RZKpXdRTfY8G6OaOT68cV5
  vP1O6T/okFKrl41FQ3CyYN5eRHyRTK99zTytrjoP2EbtIZ18z+bg/angRHYNzbgk
  lc3jpiGzs1ZWHD0nxOmHCMhU4usEcFbV6FlOxzlwrsEhHkeiununlCsNHatiDgzp
  ZWLnP/mXKV992/Jhu0Z577DHlh+3JIYx0PceB9yzACJ8MNARHF7QpBkhtuGMGZpF
  T+c73exupZFxItXs1Bnhe3djgE3MKKyYvxNUIbcTJoe7nhVMrwO/7lBSpVLvC4p3
  wR700U0LDaGGQpslGtiE56SemgoP
```

On production deployments, you can regenerate this file by running the following command:

```
$ openssl rand -base64 741
```

i. **Elasticsearch Heap Size**. Optional. The default heap size for Elasticsearch is 1GB, which may be small for production. For production deployments, uncomment the option and specify `2G` or `3G`.

```
#elastic_heap_size: 1g   # sets the heap size (1g|2g|3g) for the Elastic Servers
```

j. **OpenLDAP**. Optional. Autonomous Identity installs an OpenLDAP Docker image on the target server to hold user data. Administrators can add or remove users or change their group privileges using the **phpldapadmin** command. You can customize your OpenLDAP domain, base DN, and URL to match your company's environment. For more information, see "*Configuring LDAP*" in the *Admin Guide*.

3. Open a text editor and enter the public IP addresses of the target machines in the `~/autoid-config/hosts` file. Make sure the target host IP addresses are accessible from the deployer machine. The following is an example of the `~/autoid-config/hosts` file:

   + *Click to See a Host File for a Multi-Node Cassandra Deployment*

If you configured Cassandra as your database, the `~/autoid-config/hosts` file is as follows for multi-node target deployments:

```
[docker-managers]
34.105.16.198

[docker-workers]
34.105.16.201

[docker:children]
docker-managers
docker-workers

[cassandra-seeds]
34.105.16.198

[cassandra-workers]
34.105.16.201

[spark-master]
34.105.16.198

[spark-workers]
34.105.16.201

[analytics]
34.105.16.198

[mongo_master]

[mongo_replicas]

[mongo:children]
mongo_replicas
mongo_master

# ELastic Nodes
[odfe-master-node]
34.105.16.229

[odfe-data-nodes]
34.105.16.229

[kibana-node]
34.105.16.229
```

+ *Click to See a Host File for a Multi-Node MongoDB Deployment*

If you configured MongoDB as your database, the `~/autoid-config/hosts` file is as follows for multi-node target deployments:

```
[docker-managers]
34.105.16.198
```

```
[docker-workers]
34.105.16.201

[docker:children]
docker-managers
docker-workers

[cassandra-seeds]

[cassandra-workers]

[spark-master]
34.105.16.198

[spark-workers]
34.105.16.201

[analytics]
34.105.16.198

[mongo_master]
34.105.16.198  mongodb_master=True

[mongo_replicas]
34.105.16.201

[mongo:children]
mongo_replicas
mongo_master

# ELastic Nodes
[odfe-master-node]
34.105.16.229

[odfe-data-nodes]
34.105.16.229

[kibana-node]
34.105.16.229
```

4. Open a text editor and set the Autonomous Identity passwords for the configuration service, LDAP backend, and Cassandra database. The vault passwords file is located at `~/autoid-config/vault.yml`.

> **Note**
>
> Do not include special characters & or $ in `vault.yml` passwords as it will result in a failed deployer process.

```
configuration_service_vault:
  basic_auth_password: Welcome123

openldap_vault:
  openldap_password: Welcome123

cassandra_vault:
  cassandra_password: Welcome123
  cassandra_admin_password: Welcome123

mongo_vault:
  mongo_admin_password: Welcome123
  mongo_root_password: Welcome123

elastic_vault:
  elastic_admin_password: Welcome123
  elasticsearch_password: Welcome123
```

5. Encrypt the vault file that stores the Autonomous Identity passwords, located at `~/autoid-config/vault.yml`. The encrypted passwords will be saved to `/config/.autoid_vault_password`. The `/config/` mount is internal to the deployer container.

   ```
   $ ./deployer.sh encrypt-vault
   ```

6. Download the images. This step downloads software dependencies needed for the deployment and places them in the `autoid-packages` directory.

   ```
   $ ./deployer.sh download-images
   ```

7. Run the deployment.

   ```
   $ ./deployer.sh run
   ```

Resolve Hostname

After installing Autonomous Identity, set up the hostname resolution for your deployment.

Resolve the hostname:

1. Configure your DNS servers to access Autonomous Identity dashboard and self-service applications on the target node. The following domain names must resolve to the IP address of the target node: `<target-environment>-ui.<domain-name>` and `<target-environment>-selfservice.<domain-name>`.

2. If DNS cannot resolve target node hostname, edit it locally on the machine that you want to access Autonomous Identity using a browser. Open a text editor and add an entry in the `/etc/hosts` file for the self-service and UI services for each managed target node.

   ```
   target-ip-address  <target-environment>-ui.<domain-name> <target-environment>-selfservice.<domain-
   name>
   ```

   For example:

   ```
   34.70.190.144  autoid-ui.forgerock.com autoid-selfservice.forgerock.com
   ```

3. If you set up a custom domain name and target environment, add the entries in `/etc/hosts`. For example:

   ```
   34.70.190.144  myid-ui.abc.com  myid-selfservice.abc.com
   ```

   For more information on customizing your domain name, see "*Customize the Domain and Namespace*" in the *Admin Guide*.

Access the Dashboard

Access the Autonomous Identity console UI:

1. Open a browser, and point it to `https://autoid-ui.forgerock.com/` (or your customized URL: `https://myid-ui.abc.com`).

2. Log in as a test user: **bob.rodgers@forgerock.com**. Enter the password: **Welcome123**.

Check Apache Cassandra

Check Cassandra:

1. On the target node, check the status of Apache Cassandra.

   ```
   $ /opt/autoid/apache-cassandra-3.11.2/bin/nodetool status
   ```

2. An example output is as follows:

   ```
   Datacenter: datacenter1
   =======================
   Status=Up/Down
   |/ State=Normal/Leaving/Joining/Moving
   --  Address       Load       Tokens       Owns (effective)  Host ID                               Rack
   UN  34.70.190.144 1.33 MiB   256          100.0%            a10a91a4-96e83dd-85a2-4f90d19224d9   rack1
   ```

Check MongoDB

Check the status of MongoDB:

- On the target node, check the status of MongoDB.

  ```
  $ mongo --tls --host <Host IP Address> --tlsCAFile /opt/autoid/mongo/certs/rootCA.pem   --
  tlsAllowInvalidCertificates  --tlsCertificateKeyFile /opt/autoid/mongo/certs/mongodb.pem
  ```

Check Apache Spark

Check Spark:

- SSH to the target node and open Spark dashboard using the bundled text-mode web browser

  ```
  $ elinks http://localhost:8080
  ```

  You should see Spark Master status as ALIVE and worker(s) with State ALIVE.

  + *Click to See an Example of the Spark Dashboard*

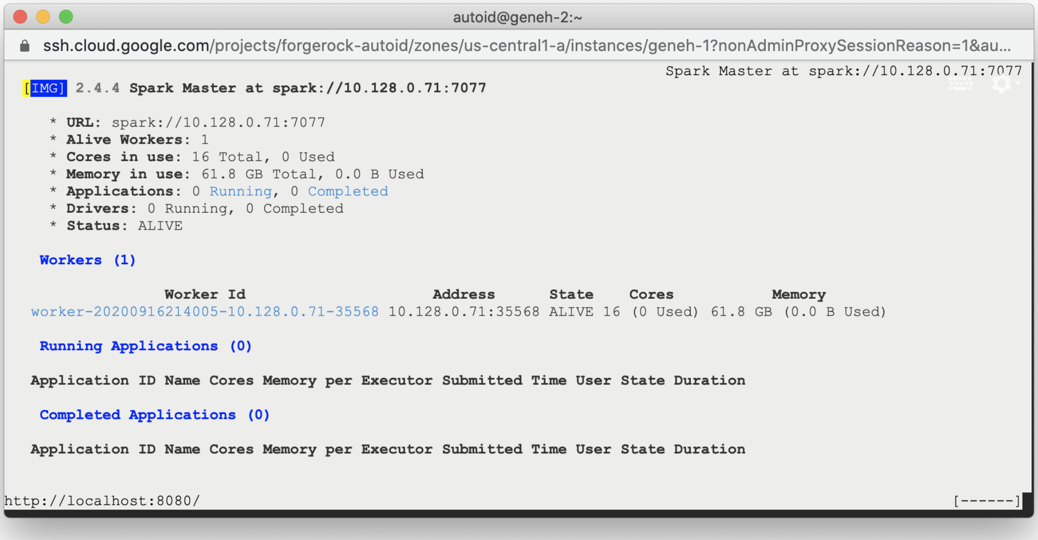
autoid@geneh-2:~
ssh.cloud.google.com/projects/forgerock-autoid/zones/us-central1-a/instances/geneh-1?nonAdminProxySessionReason=1&au...

```
                                                    Spark Master at spark://10.128.0.71:7077
[IMG] 2.4.4 Spark Master at spark://10.128.0.71:7077

   * URL: spark://10.128.0.71:7077
   * Alive Workers: 1
   * Cores in use: 16 Total, 0 Used
   * Memory in use: 61.8 GB Total, 0.0 B Used
   * Applications: 0 Running, 0 Completed
   * Drivers: 0 Running, 0 Completed
   * Status: ALIVE

 Workers (1)

              Worker Id                        Address       State    Cores            Memory
worker-20200916214005-10.128.0.71-35568 10.128.0.71:35568 ALIVE 16 (0 Used) 61.8 GB (0.0 B Used)

 Running Applications (0)

Application ID Name Cores Memory per Executor Submitted Time User State Duration

 Completed Applications (0)

Application ID Name Cores Memory per Executor Submitted Time User State Duration


http://localhost:8080/                                                                  [------]
```

Access Self-Service

The self-service feature lets Autonomous Identity users change their own passwords.

Access self-service:

- Open a browser and point it to: `https://autoid-selfservice.forgerock.com/`.

  + *Click to See an Example of the Self-Service Dashboard*

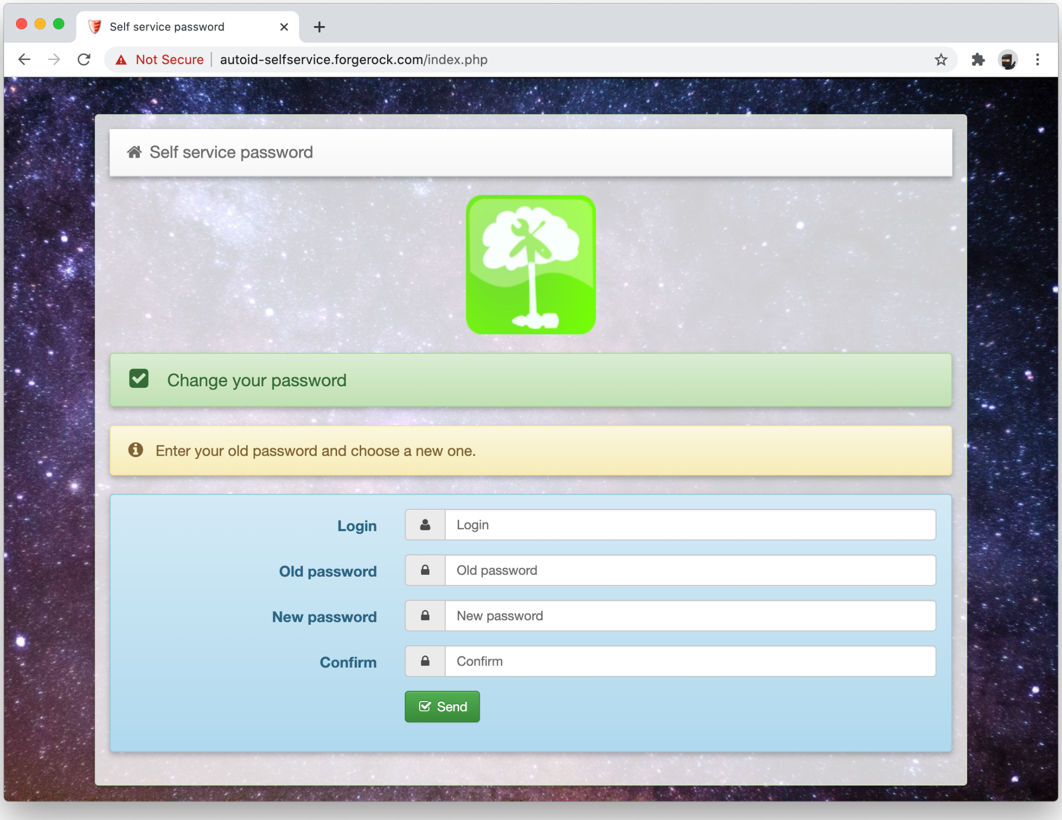

Change your password

Enter your old password and choose a new one.

Login	Login
Old password	Old password
New password	New password
Confirm	Confirm

Start the Analytics

If the previous steps all check out successfully, you can start an analytics pipeline run, where association rules, confidence scores, predications, and recommendations are generated. Autonomous Identity provides a small demo data set that lets you run the analytics pipeline on. Note for production runs, prepare your company's dataset as outlined in "*Data Preparation*" in the *Admin Guide*.

Start the analytics service:

- Run the analytics pipeline commands. This may take a bit longer than the install, depending on the size of your dataset. For specific information, see "*Run the Analytics Pipeline*" in the *Admin Guide*.

Chapter 7
Install a Multi-Node Air-Gapped Deployment

This chapter presents instructions on deploying Autonomous Identity in a multi-node air-gapped or offline target machine that has no external Internet connectivity. ForgeRock provides a deployer script that pulls a Docker container image from ForgeRock's Google Cloud Registry (gcr.io) repository. The image contains the microservices, analytics, and backend databases needed for the system.

The air-gap installation is similar to that of the multi-node deployment with Internet connectivity, except that the image and deployer script must be stored on a portable media, such as USB drive or drive, and copied to the air-gapped target environment.

The deployment depends on how the network is configured. You could have a Docker cluster with multiple Spark nodes and Cassandra or MongoDB nodes. The key is to determine the IP addresses of each node.

Figure 11: A multi-node air-gap deployment.

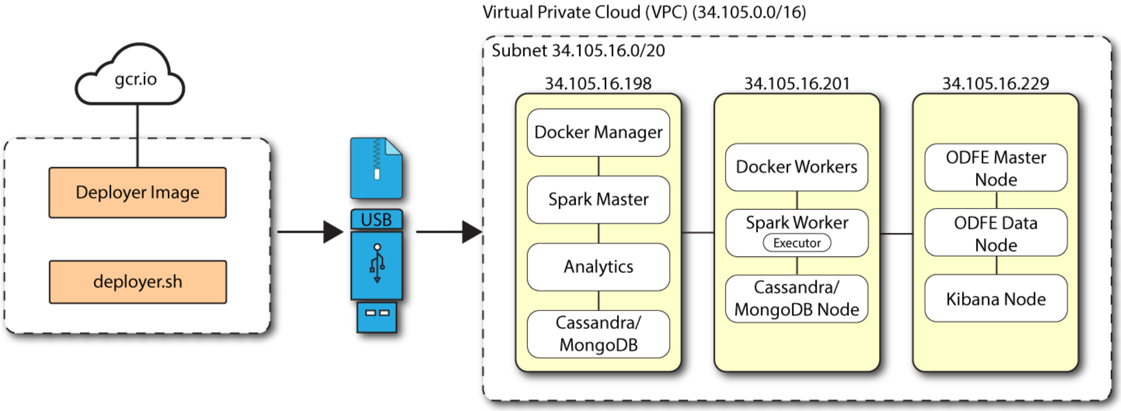

Let's deploy Autonomous Identity on a single-node target on CentOS 7. The following are prerequisites:

- **Operating System**. The target machine requires CentOS 7. The deployer machine can use any operating system as long as Docker is installed. For this guide, we use CentOS 7 as its base operating system.

- **Memory Requirements**. Make sure you have enough free disk space on the deployer machine before running the `deployer.sh` commands. We recommend at least a 40GB/partition with 14GB used and 27GB free after running the commands.
- **Default Shell**. The default shell for the `autoid` user must be bash.
- **Subnet Requirements**. We recommend deploying your multinode instances within the same subnet. Ports must be open for the installation to succeed. Each instance should be able to communicate to the other instances.

  > **Important**
  >
  > If any hosts used for the Docker cluster (docker-managers, docker-workers) have an IP address in the range of 10.0.x.x/16, they will conflict with the Swarm network. As a result, the services in the cluster will not connect to the Cassandra database or Elasticsearch backend.
  >
  > The Docker cluster hosts must be in a subnet that provides IP addresses 10.10.1.x or higher.

- **Deployment Requirements**. Autonomous Identity provides a Docker image that creates a `deployer.sh` script. The script downloads additional images necessary for the installation. To download the deployment images, you must first obtain a registry key to log into the ForgeRock Google Cloud Registry (gcr.io). The registry key is only available to ForgeRock Autonomous Identity customers. For specific instructions on obtaining the registry key, see How To Configure Service Credentials (Push Auth, Docker) in Backstage.
- **Filesystem Requirements**. Autonomous Identity requires a shared filesystem accessible from the Spark master, Spark worker, analytics hosts, and application layer. The shared filesystem should be mounted at the same mount directory on all of those hosts. If the mount directory for the shared filesystem is different from the default, `/data`, update the `/autoid-config/vars.yml` file to point to the correct directories:

  ```
  analytics_data_dir: /data
  analytics_conf_dif: /data/conf
  ```

- **Architecture Requirements**. Make sure that the analytics server is on the same node as the Spark master.
- **Database Requirements**. Decide which database you are using: Apache Cassandra or MongoDB.
- **IPv4 Forwarding**. Many high-security environments run their CentOS-based systems with IPv4 forwarding disabled. However, Docker Swarm does not work with a disabled IPv4 forwarding setting. In such environments, make sure to enable IPv4 forwarding in the file `etc/sysctl.conf`:

  ```
  net.ipv4.ip_forward=1
  ```

Set Up the Deployer Machine

Set up the deployer on an Internet-connect machine.

1. The install assumes that you have CentOS 7 as your operating system. Check your CentOS 7 version.

   ```
   $ sudo cat /etc/centos-release
   ```

2. Set the user for the target machine to a username of your choice. For example, `autoid`.

   ```
   $ sudo adduser autoid
   ```

3. Set the password for the user you created in the previous step.

   ```
   $ sudo passwd autoid
   ```

4. Configure the user for passwordless sudo.

   ```
   $ echo "autoid  ALL=(ALL)  NOPASSWD:ALL" | sudo tee /etc/sudoers.d/autoid
   ```

5. Add administrator privileges to the user.

   ```
   $ sudo usermod -aG wheel autoid
   ```

6. Change to the user account.

   ```
   $ su - autoid
   ```

7. Install yum-utils package on the deployer machine. yum-utils is a utilities manager for the Yum RPM package repository. The repository compresses software packages for Linux distributions.

   ```
   $ sudo yum install -y yum-utils
   ```

8. Create the installation directory. Note that you can use any install directory for your system as long as your run the **deployer.sh** script from there. Also, the disk volume where you have the install directory must have at least 8GB free space for the installation.

   ```
   $ mkdir ~/autoid-config
   ```

Install Docker on the Deployer Machine

1. On the target machine, set up the Docker-CE repository.

   ```
   $ sudo yum-config-manager \
       --add-repo https://download.docker.com/linux/centos/docker-ce.repo
   ```

2. Install the latest version of the Docker CE, the command-line interface, and containerd.io, a containerized website.

   ```
   $ sudo yum install -y docker-ce docker-ce-cli containerd.io
   ```

Set Up SSH on the Deployer

While SSH is not necessary to connect the deployer to the target node as the machines are isolated from one another. You still need SSH on the deployer so that it can communicate with itself.

1. On the deployer machine, run **ssh-keygen** to generate an RSA keypair, and then click Enter. You can use the default filename. Enter a password for protecting your private key.

   ```
   $ ssh-keygen -t rsa -C "autoid"
   ```

   The public and private rsa key pair is stored in `home-directory/.ssh/id_rsa` and `home-directory/.ssh/id_rsa.pub`.

2. Copy the SSH key to the `autoid-config` directory.

   ```
    $ cp ~/.ssh/id_rsa ~/autoid-config
   ```

3. Change the privileges to the file.

   ```
   $ chmod 400 ~/autoid-config/id_rsa
   ```

Prepare the Tar File

Run the following steps on an Internet-connect host machine:

1. On the deployer machine, change to the installation directory.

   ```
   $ cd ~/autoid-config/
   ```

2. Log in to the ForgeRock Google Cloud Registry (gcr.io) using the registry key. The registry key is only available to ForgeRock Autonomous Identity customers. For specific instructions on obtaining the registry key, see How To Configure Service Credentials (Push Auth, Docker) in Backstage.

   ```
   $ docker login -u _json_key -p "$(cat autoid_registry_key.json)" https://gcr.io/forgerock-autoid
   ```

   You should see:

   ```
   Login Succeeded
   ```

3. Run the **create-template** command to generate the `deployer.sh` script wrapper. Note that the command sets the configuration directory on the target node to `/config`. Note that the **--user** parameter eliminates the need to use **sudo** while editing the hosts file and other configuration files.

   ```
   $ docker run --user=`id -u` -v ~/autoid-config:/config -it gcr.io/forgerock-autoid/deployer:2020.10.2
    create-template
   ```

4. Make the script executable.

   ```
   $ chmod +x deployer.sh
   ```

5. Download the Docker images. This step downloads software dependencies needed for the deployment and places them in the `autoid-packages` directory.

   ```
   $ sudo ./deployer.sh download-images
   ```

6. Create a tar file containing all of the Autonomous Identity binaries.

   ```
   $ tar czf autoid-packages.tgz deployer.sh autoid-packages/*
   ```

7. Copy the `autoid-packages.tgz` to a USB drive or portable hard drive.

Install from the Air-Gap Target

Before you begin, make sure you have CentOS 7 installed on your air-gapped target machine.

1. Create the `~/autoid-config` directory if you haven't already.

   ```
   $ mkdir ~/autoid-config
   ```

2. Unpack the tar file.

   ```
   $ tar xf autoid-packages.tgz -C ~/autoid-config
   ```

3. On the air-gap host node, copy the SSH key to the `~/autoid-config` directory.

   ```
    $ cp ~/.ssh/id_rsa ~/autoid-config
   ```

4. Change the privileges to the file.

   ```
   $ chmod 400 ~/autoid-config/id_rsa
   ```

5. Change to the configuration directory.

   ```
   $ cd ~/autoid-config
   ```

6. Install Docker.

   ```
   $ sudo ./deployer.sh install-docker
   ```

7. Log out and back in.

8. Change to the configuration directory.

   ```
   $ cd ~/autoid-config
   ```

9. Import the deployer image.

   ```
   $ ./deployer.sh import-deployer
   ```

10. Create the configuration template using he **create-template** command. This command creates a configuration file, `ansible.cfg`.

    ```
    $ ./deployer.sh create-template
    ```

11. Make the script executable.

    ```
    $ chmod +x deployer.sh
    ```

12. To see the list of commands, enter `deployer.sh`.

    ```
    $ ./deployer.sh
         Usage: deployer <command>

    Commands:
      create-template
      download-images
      import-deployer
      encrypt-vault
      decrypt-vault
      run
      create-tar
      install-docker
      install-dbutils
      upgrade
    ```

Configure Autonomous Identity

The **create-template** command from the previous section creates a number of configuration files, required for the deployment.

1. Open a text editor and edit the `ansible.cfg` to set up the remote user and SSH private key file location on the target node. Make sure that the remote_user exists on the target node and that the deployer machine can ssh to the target node as the user specified in the `id_rsa` file.

   ```
   [defaults]
   host_key_checking = False
   remote_user = autoid
   private_key_file = id_rsa
   ```

2. On the deployer machine, open a text editor and edit the `~/autoid-config/vars.yml` file to configure specific settings for your deployment:

   a. **Domain and Target Environment**. Set the domain name and target environment specific to your deployment by editing the `/autoid-config/vars.xml` file. By default, the domain name is set to `forgerock.com` and the target environment is set to `autoid`. The default Autonomous Identity URL will be: `https://autoid-ui.forgerock.com`. For this example, we use the default values.

      ```
      domain_name: forgerock.com
      target_environment: autoid
      ```

      If you change the domain name and target environment, you need to also change the certificates to reflect the new changes. For more information, see "*Customize the Domain and Namespace*" in the *Admin Guide*.

   b. **Analytics Data Directory and Analytics Configuration Direction**. For a multi-node Spark deployment, Autonomous Identity requires a shared filesystem accessible from Spark Master, Spark Worker(s), and Analytics hosts. The shared filesystem should be mounted at same mount directory on all of the above hosts. If the mount directory for shared filesystem is different than `/data`, update the following properties in the `vars.yaml` file to point to the correct location:

```
analytics_data_dir: /data
analytics_conf_dif: /data/conf
```

c. **Dark Theme Mode**. Optional. By default, the Autonomous Identity UI displays its pages with a light background. You can set a dark theme mode by setting the `enable_dark_theme` property to `true`.

d. **Database Type**. By default, Apache Cassandra is set as the default database for Autonomous Identity. For MongoDB, set the `db_driver_type:` to `mongo`.

```
db_driver_type: mongo
```

e. **Private IP Address Mapping**. An air-gap deployment has no external IP addresses, but you may still need to define a mapping if your internal IP address differs from an external IP, say in a virtual air-gapped configuration.

If the IP addresses are the same, you can skip this step.

On the target machine, add the `private_ip_address_mapping` property in the `/inventory/vars.yml` file. Make sure the values are within double quotes. The key should not be in double quotes and should have two spaces preceding the IP address.

```
private_ip_address_mapping:
  external_ip:  "internal_ip"
```

For example:

```
private_ip_address_mapping:
  34.105.16.198:  "10.128.0.51"
  34.105.16.201:  "10.128.0.54"
  34.105.16.229:  "10.128.0.71"
```

f. **Authentication Option**. Autonomous Identity provides a single sign-on (SSO) feature that you can configure with an OIDC identity provider.

g. **JWT Expiry and Secret File**. Optional. By default, the session JWT is set at 30 minutes. To change this value, set the `jwt_expiry` property to a different value.

```
jwt_expiry: "30 minutes"
```

h. **MongoDB Configuration**. For MongoDB clusters, enable replication by uncommenting the `mongodb_replication_replset` property.

```
# uncomment below for mongo with replication enabled. Not needed for single node deployments
mongodb_replication_replset: mongors
```

Also, enable a custom key for inter-machine authentication in the clustered nodes.

```
# custom key
# password for inter-process authentication
# please regenerate this file on production environment with
# command 'openssl rand -base64 741'
mongodb_keyfile_content: |
  8pYcxvCqoe89kcp33KuTtKVf5MoHGEFjTnudrq5BosvWRoIxLowmdjrmUpVfAivh
  CHjqM6w0zVBytAxH1lW+7teMYe6eDn2S/O/1YlRRiW57bWU3zjliW3VdguJar5i9
  Z+1a8lI+0S9pWynbv9+Ao0aXFjSJYVxAm/w7DJbVRGcPhsPmExiSBDw8szfQ8PAU
  2hwRl7nqPZZMMR+uQThg/zV9rOzHJmkqZtsO4UJSilG9euLCYrzW2hdoPuCrEDhu
  Vsi5+nwAgYR9dP2oWkmGN1dwRe0ixSIM2UzFgpaXZaMOG6VztmFrlVXh8oFDRGM0
  cGrFHcnGF7oUGfWnI2Cekngk64dHA2qD7WxXPbQ/svn9EfTY5aPw5lXzKA87Ds8p
  KHVFUYvmA6wVsxb/riGLwc+XZlb6M9gqHn1XSpsnYRjF6UzfRcRR2WyCxLZELaqu
  iKxLKB5FYqMBH7Sqg3qBCtE53vZ7T1nefq5RFzmykviYP63Uhu/A2EQatrMnaFPl
  TTG5CaPjob45CBSyMrheYRWKqxdWN93BTgiTW7p0U6RB0/OCUbsVX6IG3I9N8Uqt
  l8Kc+7aOmtUqFkwo8w30prIOjStMrokxNsuK9KTUiPu2cj7gwYQ574vV3hQvQPAr
  hhb9ohKr0zoPQt31iTj0FDkJzPepeuzqeq8F51HB56RZKpXdRTfY8G6OaOT68cV5
  vP1O6T/okFKrl41FQ3CyYN5eRHyRTK99zTytrjoP2EbtIZ18z+bg/angRHYNzbgk
  lc3jpiGzs1ZWHD0nxOmHCMhU4usEcFbV6FlOxzlwrsEhHkeiununlCsNHatiDgzp
  ZWLnP/mXKV992/Jhu0Z577DHlh+3JIYx0PceB9yzACJ8MNARHF7QpBkhtuGMGZpF
  T+c73exupZFxItXs1Bnhe3djgE3MKKyYvxNUIbcTJoe7nhVMrwO/7lBSpVLvC4p3
  wR700U0LDaGGQpslGtiE56SemgoP
```

On production deployments, you can regenerate this file by running the following command:

```
$ openssl rand -base64 741
```

i. **Elasticsearch Heap Size**. Optional. The default heap size for Elasticsearch is 1GB, which may be small for production. For production deployments, uncomment the option and specify `2G` or `3G`.

```
#elastic_heap_size: 1g   # sets the heap size (1g|2g|3g) for the Elastic Servers
```

j. **OpenLDAP**. Optional. Autonomous Identity installs an OpenLDAP Docker image on the target server to hold user data. Administrators can add or remove users or change their group privileges using the **phpldapadmin** command. You can customize your OpenLDAP domain, base DN, and URL to match your company's environment. For more information, see "*Configuring LDAP*" in the *Admin Guide*.

3. Open a text editor and enter the public IP addresses of the target machines in the `~/autoid-config/hosts` file. Make sure the target host IP addresses are accessible from the deployer machine. The following is an example of the `~/autoid-config/hosts` file:

   + *Click to See a Host File for Cassandra Deployments*

   If you configured Cassandra as your database, the `~/autoid-config/hosts` file is as follows for single-node target deployments:

   ```
   [docker-managers]
   10.10.15.240

   [docker-workers]
   10.10.15.201
   ```

```
[docker:children]
docker-managers
docker-workers

[cassandra-seeds]
10.10.15.212

[cassandra-workers]
10.10.15.94

[spark-master]
10.10.15.212

[spark-workers]
10.10.15.94

[analytics]
10.10.15.212

[mongo_master]

[mongo_replicas]

[mongo:children]
mongo_replicas
mongo_master

# ELastic Nodes
[odfe-master-node]
10.10.15.201

[odfe-data-nodes]
10.10.15.201

[kibana-node]
10.10.15.201
```

+ *Click to See a Host File for MongoDB Deployments*

If you configured MongoDB as your database, the `~/autoid-config/hosts` file is as follows for single-node target deployments:

```
[docker-managers]
10.10.15.240

[docker-workers]
10.10.15.201

[docker:children]
docker-managers
docker-workers

[cassandra-seeds]

[cassandra-workers]
```

```
[spark-master]
10.10.15.212

[spark-workers]
10.10.15.94

[analytics]
10.10.15.212

[mongo_master]
10.10.15.240  mongodb_master=True

[mongo_replicas]
10.10.15.212

[mongo:children]
mongo_replicas
mongo_master

# ELastic Nodes
[odfe-master-node]
10.10.15.201

[odfe-data-nodes]
10.10.15.201

[kibana-node]
10.10.15.201
```

4. Set the Autonomous Identity passwords, located at `~/autoid-config/vault.yml`.

> **Note**
>
> Do not include special characters & or $ in `vault.yml` passwords as it will result in a failed deployer process.

```
configuration_service_vault:
  basic_auth_password: Welcome123

openldap_vault:
  openldap_password: Welcome123

cassandra_vault:
  cassandra_password: Welcome123
  cassandra_admin_password: Welcome123

mongo_vault:
  mongo_admin_password: Welcome123
  mongo_root_password: Welcome123

elastic_vault:
  elastic_admin_password: Welcome123
  elasticsearch_password: Welcome123
```

5. Encrypt the vault file that stores the Autonomous Identity passwords, located at `~/autoid-config/vault.yml`. The encrypted passwords will be saved to `/config/.autoid_vault_password`. The `/config/` mount is internal to the deployer container.

   ```
   $ ./deployer.sh encrypt-vault
   ```

6. Run the deployment.

   ```
   $ ./deployer.sh run
   ```

Access the Dashboard

Access the Autonomous Identity console UI:

1. Open a browser, and point it to `https://autoid-ui.forgerock.com/` (or your customized URL: `https://myid-ui.abc.com`).

2. Log in as a test user: **`bob.rodgers@forgerock.com`**. Enter the password: **`Welcome123`**.

Check Apache Cassandra

Check Cassandra:

1. On the target node, check the status of Apache Cassandra.

   ```
   $ /opt/autoid/apache-cassandra-3.11.2/bin/nodetool status
   ```

2. An example output is as follows:

   ```
   Datacenter: datacenter1
   =======================
   Status=Up/Down
   |/ State=Normal/Leaving/Joining/Moving
   --  Address       Load       Tokens       Owns (effective)  Host ID                               Rack
   UN  34.70.190.144 1.33 MiB   256          100.0%            a10a91a4-96e83dd-85a2-4f90d19224d9  rack1
   ```

Check MongoDB

Check the status of MongoDB:

- On the target node, check the status of MongoDB.

  ```
  $ mongo --tls --host <Host IP Address> --tlsCAFile /opt/autoid/mongo/certs/rootCA.pem   --
  tlsAllowInvalidCertificates  --tlsCertificateKeyFile /opt/autoid/mongo/certs/mongodb.pem
  ```

Check Apache Spark

Check Spark:

- SSH to the target node and open Spark dashboard using the bundled text-mode web browser

  ```
  $ elinks http://localhost:8080
  ```

  You should see Spark Master status as ALIVE and worker(s) with State ALIVE.

  + *Click to See an Example of the Spark Dashboard*

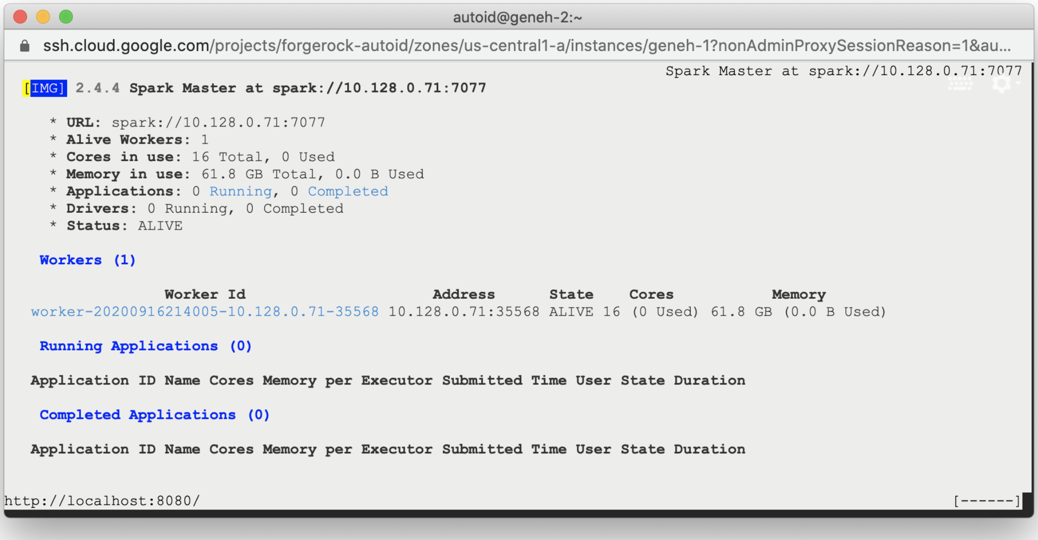

autoid@geneh-2:~
ssh.cloud.google.com/projects/forgerock-autoid/zones/us-central1-a/instances/geneh-1?nonAdminProxySessionReason=1&au...

```
                                                                Spark Master at spark://10.128.0.71:7077
[IMG] 2.4.4 Spark Master at spark://10.128.0.71:7077

   * URL: spark://10.128.0.71:7077
   * Alive Workers: 1
   * Cores in use: 16 Total, 0 Used
   * Memory in use: 61.8 GB Total, 0.0 B Used
   * Applications: 0 Running, 0 Completed
   * Drivers: 0 Running, 0 Completed
   * Status: ALIVE

 Workers (1)

               Worker Id                      Address       State    Cores            Memory
worker-20200916214005-10.128.0.71-35568 10.128.0.71:35568 ALIVE 16 (0 Used) 61.8 GB (0.0 B Used)

 Running Applications (0)

Application ID Name Cores Memory per Executor Submitted Time User State Duration

 Completed Applications (0)

Application ID Name Cores Memory per Executor Submitted Time User State Duration

http://localhost:8080/                                                                          [------]
```

Access Self-Service

The self-service feature lets Autonomous Identity users change their own passwords.

Access self-service:

- Open a browser and point it to: `https://autoid-selfservice.forgerock.com/`.

  + *Click to See an Example of the Self-Service Dashboard*

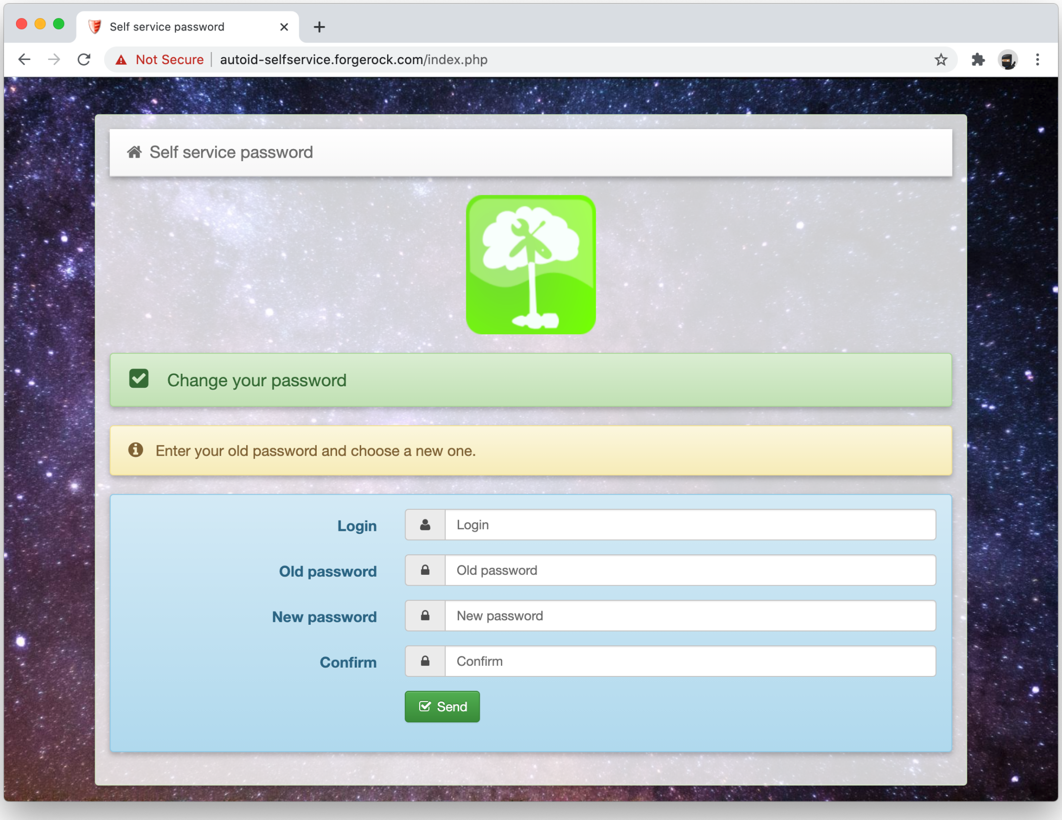

Self service password

Not Secure | autoid-selfservice.forgerock.com/index.php

Self service password

Change your password

Enter your old password and choose a new one.

Start the Analytics

If the previous steps all check out successfully, you can start an analytics pipeline run, where association rules, confidence scores, predications, and recommendations are generated. Autonomous Identity provides a small demo data set that lets you run the analytics pipeline on. Note for production runs, prepare your company's dataset as outlined in "*Data Preparation*" in the *Admin Guide*.

Start the analytics service:

- Run the analytics pipeline commands. This may take a bit longer than the install, depending on the size of your dataset. For specific information, see "*Run the Analytics Pipeline*" in the *Admin Guide*.

Chapter 8

Upgrade Autonomous Identity

Autonomous Identity 2020.10.2 provides upgrade commands to update your core software to the latest version while migrating your data.

The upgrade assumes the following:

- **Database Systems are the Same**. If your current database is Apache Cassandra, you cannot upgrade to a MongoDB-based system. You will need to run a clean installation with the new version.
- **Host IPs should be the Same**. Host IP addresses must be the same for existing components. You must update the `~/autoid-config/hosts` file by adding the IP addresses for the Elasticsearch entries. See the instructions below.
- **Registry Key Required**. To download the deployment images for the upgrade, you still need a registry key to log into the ForgeRock Google Cloud Registry (gcr.io). The registry key is only available to ForgeRock Autonomous Identity customers. For specific instructions on obtaining the registry key, see How To Configure Service Credentials (Push Auth, Docker) in Backstage.
- **Additional CSV Files Requires**. The upgraded system requires the `app_attributes.csv` and `ent_attributes.csv` for attribute filtering on the Applications page. Copy these files to the `/data/input` directory. For more information on these files, see "*Data Preparation*" in the *Admin Guide*.
- **Upgrade Paths**. The upgrade paths are summarized as follows:
  - 2020.10.x -> 2020.10.2
  - 2020.6.4 -> 2020.10.2

Upgrade from 2020.10.x to 2020.10.2

1. On the deployer machine, back up the 2020.10.0 or 2020.10.1 `~/autoid-config` directory or move it to another location.

   ```
   $ mv ~/autoid-config ~/backup-2020.10.1
   ```

2. Create a new `~/autoid-config` directory.

   ```
   $ mkdir ~/autoid-config
   ```

3. Copy your `autoid_registry_key.json`, `ansible.cfg`, and `vault.yml` files from your backup directory to `~/autoid-config`. If your `vault.yml` file is encrypted, copy the `.autoid_vault_password` file to ~/autoid-config.

4. Remove your `known_files`.

   ```
   $ rm ~/.ssh/known_hosts
   ```

5. Copy your original SSH key into the new directory.

   ```
   $ cp ~/.ssh/id_rsa ~/autoid-config
   ```

6. Change the permission on the SSH key.

   ```
   $ chmod 400 ~/autoid-config/id_rsa
   ```

7. Check if you can successfully SSH to the target server.

   ```
   $ ssh autoid@<Target-IP-Address>
   ```

8. Enter **exit** to end your SSH session.

9. On the deployer node, change to the `~/autoid-config` directory.

   ```
   $ cd ~/autoid-config
   ```

10. Log in to the ForgeRock Google Cloud Registry (gcr.io) using the registry key. The registry key is only available to ForgeRock Autonomous Identity customers. For specific instructions on obtaining the registry key, see How To Configure Service Credentials (Push Auth, Docker) in Backstage.

    ```
    $ docker login -u _json_key -p "$(cat autoid_registry_key.json)" https://gcr.io/forgerock-autoid
    ```

    You should see:

    ```
    Login Succeeded
    ```

11. Run the **create-template** command to generate the `deployer.sh` script wrapper and configuration files. Note that the command sets the configuration directory on the target node to `/config`. The **--user** parameter eliminates the need to use **sudo** while editing the hosts file and other configuration files.

    ```
    $ docker run --user=`id -u` -v ~/autoid-config:/config -it gcr.io/forgerock-autoid/deployer:2020.10.2
     create-template
    ```

12. Make the script executable.

    ```
    $ chmod +x deployer.sh
    ```

13. Copy your `~/autoid-config/vars.yml`, `~/autoid-config/hosts`, and `~/autoid-config/vault.yml` files from your backup directory to the deployer machine.

> **Important**
>
> You must keep your configuration settings consistent from one system to another.

14. Download the images. This step downloads software dependencies needed for the deployment and places them in the `autoid-packages` directory. Make sure you are in the `~/autoid-config` directory.

    ```
    $ ./deployer.sh download-images
    ```

15. Run the upgrade.

    ```
    $ ./deployer.sh debug patch_log4j
    ```

Upgrade from 2020.6.4 to 2020.10.2

Upgrade to version from 2020.6.4 to 2020.10.2

1. On the deployer machine, back up the 2020.6.4 `~/autoid-config` directory or move it to another location.

   ```
   $ mv ~/autoid-config ~/backup-2020.6
   ```

2. Create a new `~/autoid-config` directory.

   ```
   $ mkdir ~/autoid-config
   ```

3. Copy your `autoid_registry_key.json`, `ansible.cfg`, `vault.yml`, and SSH keys. files from your backup directory to `~/autoid-config`. If your `vault.yml` file is encrypted, copy the `.autoid_vault_password` file to ~/autoid-config.

4. Remove your `known_files`.

   ```
   $ rm ~/.ssh/known_hosts
   ```

5. Copy your original SSH key into the new directory.

   ```
   $ cp ~/.ssh/id_rsa ~/autoid-config
   ```

6. Change the permission on the SSH key.

   ```
   $ chmod 400 ~/autoid-config/id_rsa
   ```

7. Check if you can successfully SSH to the target server.

   ```
   $ ssh autoid@<Target-IP-Address>
   Last login: Tue Jan 20 18:19:14 2022
   ```

8. Enter **exit** to end your SSH session.

9. On the deployer node, change to the `~/autoid-config` directory.

   ```
   $ cd ~/autoid-config
   ```

10. Log in to the ForgeRock Google Cloud Registry (gcr.io) using the registry key. The registry key is only available to ForgeRock Autonomous Identity customers. For specific instructions on obtaining the registry key, see How To Configure Service Credentials (Push Auth, Docker) in Backstage.

    ```
    $ docker login -u _json_key -p "$(cat autoid_registry_key.json)" https://gcr.io/forgerock-autoid
    ```

11. Run the **create-template** command to generate the `deployer.sh` script wrapper and configuration files. Note that the command sets the configuration directory on the target node to `/config`. The **--user** parameter eliminates the need to use **sudo** while editing the hosts file and other configuration files.

    ```
    $ docker run --user=`id -u` -v ~/autoid-config:/config -it gcr.io/forgerock-autoid/deployer:2020.10.2
     create-template
    ```

12. Make the script executable.

    ```
    $ chmod +x deployer.sh
    ```

13. Copy your `~/autoid-config/vars.yml`, `~/autoid-config/hosts`, and `~/autoid-config/vault.yml` files from your backup directory to the deployer machine.

    > **Important**
    >
    > You must keep your configuration settings consistent from one system to another.

14. SSH to the target server.

15. Stop the stack.

    ```
    $ docker stack rm configuration-service consul-server nginx openldap selfservice swagger-ui ui api
     consul-client
    ```

16. Delete the contents from the consul data. This step is required as the new consul server in the upgrade is not able to load the previous snapshots, because of a version incompatibility.

    ```
    $ mv /opt/autoid/mounts/consul-data/* <backup-directory>
    ```

17. Exit your SSH session.

18. Download the images. This step downloads software dependencies needed for the deployment and places them in the `autoid-packages` directory. Make sure you are in the `~/autoid-config` directory.

    ```
    $ ./deployer.sh download-images
    ```

19. Run the upgrade.

```
$ ./deployer.sh upgrade
```

20. Log out and then log in. SSH to the target server.

21. Take a backup of the */data/conf* or move it to another directory. The directory stores the 2020.6 configuration files.

```
$ mv /data/conf <backup-directory>
```

22. Create an analytics template. This step creates a template from the new analytics image.

```
$ analytics create-template
```

23. Edit the `/data/conf/analytics_init_config.yml` file if you made changes to this file in your previous deployment.

24. Apply the analytics template.

```
$ analytics apply-template
```

25. Run the analytics upgrade job.

```
$ analytics upgrade
```

26. Edit the *analytics_config* file to set up the ingestion process. The change is required to account for the applications and entitlements tables.

```

  etl: false
  reports: false
ingestion:
  drop_if_create: true
  tables: app_attributes,ent_attributes
  catalog_step: false
  staging: false
  connector:
     type: csv
  connector-oim:
     type: oim
     timeout: 15
     batchsize: 1000
     change reconciliation:
     ...

```

27. Run the analytics ingest job.

```
$ analytics ingest
```

28. Run the analytics publish job.

```
$ analytics publish
```

29. Run the create-assignment-index job.

```
$ analytics create-assignment-index
```

You have successfully upgrade from 2020.6.x to 2020.10.2!

Access the Dashboard

Access the Autonomous Identity console UI:

1. Open a browser, and point it to `https://autoid-ui.forgerock.com/` (or your customized URL: `https://myid-ui.abc.com`).
2. Log in as a test user: **`bob.rodgers@forgerock.com`**. Enter the password: **`Welcome123`**.

Appendix A. Appendix A: Autonomous Identity Ports

The Autonomous Identity deployment uses the following ports. The Docker deployer machine opens the ports in the firewall on the target node. If you are using cloud virtual machines, you need to open these ports on the virtual cloud network.

Autonomous Identity uses the following ports:

Autonomous Identity Ports

Port	Protocol	Machine	Description
2376	TCP	Docker	Secure Docker client communication. This port is required for the Docker machine, which orchestrates the Docker hosts.
2377	TCP	Docker	Communication between the nodes of a Docker swarm cluster. Only needed on manager nodes.
7946	TCP/UDP	Docker	Communication among nodes for container network discovery.
4789	TCP	Docker	Overlay network traffic.
7001	TCP	Cassandra	Internode communication.
9042	TCP	Cassandra	CQL native transport.
7077	TCP	Spark	Spark master internode communication port.
40040-40045	TCP	Analytics	Spark driver ports for Spark workers to callback.
443	TCP	Autonomous Identity	Port to access the dashboard and API.

Chapter 9

Appendix B: vars.yml

Autonomous Identity has a configuration file where you can set the analytics data and configuration directories, UI dark theme mode, private IP address mapping, LDAP/SSO options, and session duration during installation. The file is created when running the **create-template** command during the installation and is located in the `/autoid-config` directory.

The file is as follows:

```
domain_name: forgerock.com                                  # Default domain name
target_environment: autoid                                  # Default namespace
analytics_data_dir: /data                                   # Default data directory
analytics_conf_dir: /data/conf                              # Default config directory for analytics
enable_dark_theme: false                                    # Set true for dark UI theme mode

# Needed only if private and public IP address of
# target nodes are different. If cloud VMs the private
# is different than the IP address (public ip) used for
# SSH. Private IP addresses are used by various services
# to reach other services in the cluster
# Example:
# private_ip_address_mapping:
#   35.223.33.21: "10.128.0.5"
#   108.59.83.132: "10.128.0.37"
#   ...
private_ip_address_mapping:                                 # private and external IP mapping
#private_ip_address_mapping-ip-addesses#

api:
  authentication_option: "Ldap"                             # Values: "Ldap", "SSO", "LdapAndSSO"
  access_log_enabled: true                                  # Enable access logs
  jwt_expiry: "30 minutes"                                  # Default session duration
  jwt_secret_file: "{{ install_path }}/jwt/secret.txt"      # Location of JWT secret file

# mongo config starts
#uncomment below for mongo with replication enabled. Not needed for single node deployments
# mongodb_replication_replset: mongors
# custom key
# password for inter-process authentication
# # please regenerate this file on production environment with command 'openssl rand -base64 741'
#mongodb_keyfile_content: |
#   8pYcxvCqoe89kcp33KuTtKVf5MoHGEFjTnudrq5BosvWRoIxLowmdjrmUpVfAivh
#   CHjqM6w0zVBytAxH1lW+7teMYe6eDn2S/O/1YlRRiW57bWU3zjliW3VdguJar5i9
#   Z+1a8lI+0S9pWynbv9+Ao0aXFjSJYVxAm/w7DJbVRGcPhsPmExiSBDw8szfQ8PAU
#   2hwRl7nqPZZMMR+uQThg/zV9rOzHJmkqZtsO4UJSilG9euLCYrzW2hdoPuCrEDhu
#   Vsi5+nwAgYR9dP2oWkmGN1dwRe0ixSIM2UzFgpaXZaMOG6VztmFrlVXh8oFDRGM0
#   cGrFHcnGF7oUGfWnI2Cekngk64dHA2qD7WxXPbQ/svn9EfTY5aPw5lXzKA87Ds8p
#   KHVFUYvmA6wVsxb/riGLwc+XZlb6M9gqHn1XSpsnYRjF6UzfRcRR2WyCxLZELaqu
#   iKxLKB5FYqMBH7Sqg3qBCtE53vZ7T1nefq5RFzmykviYP63Uhu/A2EQatrMnaFPl
#  TTG5CaPjob45CBSyMrheYRWKqxdWN93BTgiTW7p0U6RB0/OCUbsVX6IG3I9N8Uqt
#  l8Kc+7aOmtUqFkwo8w30prIOjStMrokxNsuK9KTUiPu2cj7gwYQ574vV3hQvQPAr
#  hhb9ohKr0zoPQt31iTj0FDkJzPepeuzqeq8F51HB56RZKpXdRTfY8G6OaOT68cV5
#  vP1O6T/okFKrl41FQ3CyYN5eRHyRTK99zTytrjoP2EbtIZ18z+bg/angRHYNzbgk
#  lc3jpiGzs1ZWHD0nxOmHCMhU4usEcFbV6FlOxzlwrsEhHkeiununlCsNHatiDgzp
#  ZWLnP/mXKV992/Jhu0Z577DHlh+3JIYx0PceB9yzACJ8MNARHF7QpBkhtuGMGZpF
#  T+c73exupZFxItXs1Bnhe3djgE3MKKyYvxNUIbcTJoe7nhVMrwO/7lBSpVLvC4p3
#  wR700U0LDaGGQpslGtiE56SemgoP
# mongo config ends

# set the following API parameters when                     # SSO and LdapAndSSO properties
# authentication_option is SSO or LdapAndSSO
#  oidc_issuer:
#  oidc_auth_url:
#  oidc_token_url:
#  oidc_user_info_url:
#  oidc_callback_url:
#  oidc_jwks_url:
#  oidc_client_scope:
#  oidc_groups_attribute:
#  oidc_uid_attribute:
#  oidc_client_id:
#  oidc_client_secret:
#  admin_object_id:
#  entitlement_owner_object_id:
#  executive_object_id:
#  supervisor_object_id:
#  user_object_id:

# Elastic Optional Config
#elastic_heap_size: 1g   # sets the heap size (1g|2g|3g) for the Elastic Servers
```

Glossary

anomaly report
A report that identifies potential anomalous assignments.

as-is predictions
A process where confidence scores are assigned to the entitlements that users have.

auto-certify
An action that an entitlement owner can do to approve a justification. Auto-certify indicates that anyone who has the justification is automatically approved for the entitlement.

auto-request
An action that an entitlement owner can do to approve a justification. Auto-request indicates that anyone who matches these justification attributes but may not already have access should automatically get provisioned for this entitlement.

confidence score
A score from a scale from 0 to 100% that indicates the strength of correlation between an assigned entitlement and a user's data profile.

data audit
A pre-analytics process that audits the seven data files to ensure data validity with the client.

data ingestion
A pre-analytics process that pushes the seven .csv files into the Cassandra database. This allows the entire training process to be performed from the database.

data sparsity
A reference to data that has null values. Autonomous Identity requires dense, high quality data with very few null values in the user attributes to get accurate analysis scores.

data validation
A pre-analytics process that tests the data to ensure that the content is correct and complete prior to the training process.

driving factor

An association rule that is a key factor in a high entitlement confidence score. Any rule that exceeds a confidence threshold level (e.g., 75%) is considered a driving factor.

entitlement

An entitlement is a specialized type of `assignment`. A user or device with an entitlement gets access rights to specified resources.

insight report

A report that provides metrics on the rules and predictions generated in the analytics run.

recommendation

A process run after the as-is predictions that assigns confidence scores to all entitlements and recommends entitlements that users do not currently have. If the confidence score meets a threshold, set by the `conf_thresh` property in the configuration file, the entitlement will be recommended to the user in the UI console.

resource

An external system, database, directory server, or other source of identity data to be managed and audited by an identity management system.

REST

Representational State Transfer. A software architecture style for exposing resources, using the technologies and protocols of the World Wide Web. REST describes how distributed data objects, or resources, can be defined and addressed.

stemming

A process that occurs after training that removes similar association rules that exist in a parent-child relationship. If the child meets three criteria, then it will be removed by the system. The criteria are: 1) the child must match the parent; 2) the child (e.g., [San Jose, Finance]) is a superset of the parent rule. (e.g., [Finance]); 3) the child and parent's confidence scores are within a +/- range of each other. The range is set in the configuration file.

training

A multi-step process that generates the association rules with confidence scores for each entitlement. First, Autonomous Identity models the frequent itemsets that appear in the user attributes for each user. Next, Autonomous Identity merges the user attributes with the entitlements that were assigned to the user. It then applies association rules to model the sets of user attributes that result in an entitlement access and calculates confidence scores, based on their frequency of appearances in the dataset.