



Deployment Planning Guide

/ Autonomous Identity 2020.6.4

Latest update: 2020.6.4

Copyright © 2020 ForgeRock AS.

Abstract

This guide is targeted to deployer, administrators, and system architects who must plan a production deployment of the Autonomous Identity system.



This work is licensed under the Creative Commons Attribution-NonCommercial-NoDerivs 3.0 Unported License.

To view a copy of this license, visit <https://creativecommons.org/licenses/by-nc-nd/3.0/> or send a letter to Creative Commons, 444 Castro Street, Suite 900, Mountain View, California, 94041, USA.

© Copyright 2010-2020 ForgeRock, Inc. All rights reserved. ForgeRock is a registered trademark of ForgeRock, Inc. Other marks appearing herein may be trademarks of their respective owners.

This product or document is protected by copyright and distributed under licenses restricting its use, copying, and distribution. No part of this product or document may be reproduced in any form by any means without prior written authorization of ForgeRock and its licensors, if any.

DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESSED OR IMPLIED CONDITIONS, REPRESENTATIONS, AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID.

DejaVu Fonts

Bitstream Vera Fonts Copyright

Copyright (c) 2003 by Bitstream, Inc. All Rights Reserved. Bitstream Vera is a trademark of Bitstream, Inc.

Permission is hereby granted, free of charge, to any person obtaining a copy of the fonts accompanying this license ("Fonts") and associated documentation files (the "Font Software"), to reproduce and distribute the Font Software, including without limitation the rights to use, copy, merge, publish, distribute, and/or sell copies of the Font Software, and to permit persons to whom the Font Software is furnished to do so, subject to the following conditions:

The above copyright and trademark notices and this permission notice shall be included in all copies of one or more of the Font Software typefaces.

The Font Software may be modified, altered, or added to, and in particular the designs of glyphs or characters in the Fonts may be modified and additional glyphs or characters may be added to the Fonts, only if the fonts are renamed to names not containing either the words "Bitstream" or the word "Vera".

This License becomes null and void to the extent applicable to Fonts or Font Software that has been modified and is distributed under the "Bitstream Vera" names.

The Font Software may be sold as part of a larger software package but no copy of one or more of the Font Software typefaces may be sold by itself.

THE FONT SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NON-INFRINGEMENT OF COPYRIGHT, PATENT, TRADEMARK, OR OTHER RIGHT. IN NO EVENT SHALL BITSTREAM OR THE GNOME FOUNDATION BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, INCLUDING ANY GENERAL, SPECIAL, INDIRECT, INCIDENTAL, OR CONSEQUENTIAL DAMAGES, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF THE USE OR INABILITY TO USE THE FONT SOFTWARE OR FROM OTHER DEALINGS IN THE FONT SOFTWARE.

Except as contained in this notice, the names of GNOME, the GNOME Foundation, and Bitstream Inc., shall not be used in advertising or otherwise to promote the sale, use or other dealings in this Font Software without prior written authorization from the GNOME Foundation or Bitstream Inc., respectively. For further information, contact: fonts@gnome.org.

Arev Fonts Copyright

Copyright (c) 2006 by Tavmjong Bah. All Rights Reserved.

Permission is hereby granted, free of charge, to any person obtaining a copy of the fonts accompanying this license ("Fonts") and associated documentation files (the "Font Software"), to reproduce and distribute the modifications to the Bitstream Vera Font Software, including without limitation the rights to use, copy, merge, publish, distribute, and/or sell copies of the Font Software, and to permit persons to whom the Font Software is furnished to do so, subject to the following conditions:

The above copyright and trademark notices and this permission notice shall be included in all copies of one or more of the Font Software typefaces.

The Font Software may be modified, altered, or added to, and in particular the designs of glyphs or characters in the Fonts may be modified and additional glyphs or characters may be added to the Fonts, only if the fonts are renamed to names not containing either the words "Tavmjong Bah" or the word "Arev".

This License becomes null and void to the extent applicable to Fonts or Font Software that has been modified and is distributed under the "Tavmjong Bah Arev" names.

The Font Software may be sold as part of a larger software package but no copy of one or more of the Font Software typefaces may be sold by itself.

THE FONT SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NON-INFRINGEMENT OF COPYRIGHT, PATENT, TRADEMARK, OR OTHER RIGHT. IN NO EVENT SHALL TAVMJONG BAH BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, INCLUDING ANY GENERAL, SPECIAL, INDIRECT, INCIDENTAL, OR CONSEQUENTIAL DAMAGES, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF THE USE OR INABILITY TO USE THE FONT SOFTWARE OR FROM OTHER DEALINGS IN THE FONT SOFTWARE.

Except as contained in this notice, the name of Tavmjong Bah shall not be used in advertising or otherwise to promote the sale, use or other dealings in this Font Software without prior written authorization from Tavmjong Bah. For further information, contact: tavmjong@free.fr.

FontAwesome Copyright

Copyright (c) 2017 by Dave Gandy, <https://fontawesome.com/>.

This Font Software is licensed under the SIL Open Font License, Version 1.1. See <https://opensource.org/licenses/OFL-1.1>.

Table of Contents

Overview	iv
1. Architecture in Brief	1
2. Topology Planning	3
Suggested Number of Servers	3
Production Technical Specifications	3
3. Deployment Checklist	5
Glossary	8

Overview

Use this guide to plan your production Autonomous Identity deployment.

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.

Important

This guide is for deployers, technical consultants, and administrators who are familiar with Autonomous Identity and are responsible for architecting a production deployment.

Quick Start



Topology Planning

Learn about topology sizing considerations.



Deployment Checklist

Learn about the Autonomous Identity deployment checklist.



Glossary

Review the Autonomous Identity glossary terms

For installation instructions, see the [Installation Guide](#).

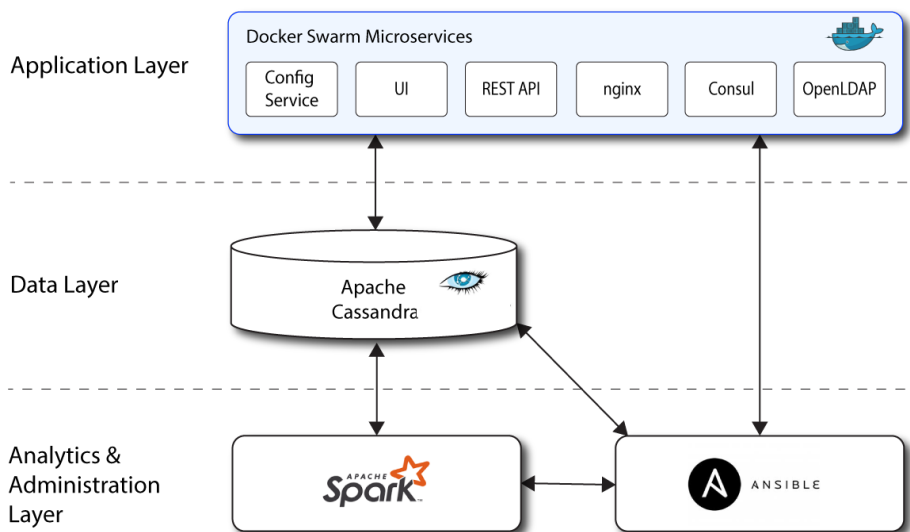
For component versions, see the [Release Notes](#).

Chapter 1

Architecture in Brief

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

Figure 1: A Conceptual Image of the Autonomous Identity Architecture

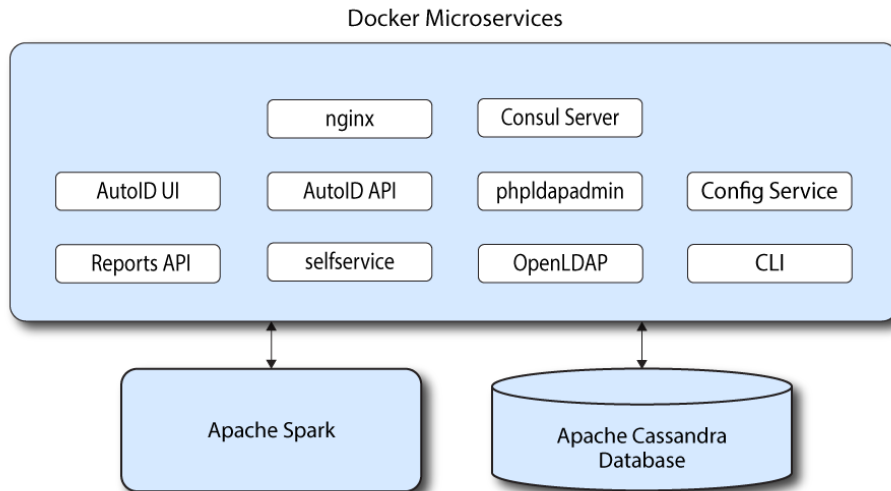


- **Application Layer.** Autonomous Identity implements a flexible Docker Swarm microservices architecture, where multiple applications run together in containers. The microservices component provides flexible configuration and end-user interaction to the deployment. The microservices components are the following:
 - **Configuration Service.** Autonomous Identity supports a configuration service that allows you to set parameters for your system and processes.
 - **Autonomous Identity UI.** Autonomous Identity supports a dynamic UI that displays the entitlements, confidence scores, and recommendations.
 - **Autonomous Identity REST API.** Autonomous Identity provides an API that can access endpoints using REST. This allows easy scripting and programming for your system.

- **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.
- **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.
- **Data Layer.** Autonomous Identity uses a Apache Cassandra NoSQL database 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.
- **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.

Autonomous Identity's flexible architecture can be deployed in any number of ways: single-node or multi-node configurations across on-prem, cloud, hybrid, or multi-cloud environments. For example, we typically configure a two-server deployment with the minimum hardware and software requirements for a pilot deployment.

Figure 2: An Image of a Two-Server Autonomous Identity Pilot Architecture



Chapter 2 Topology Planning

Based on existing production deployments, we have determined a suggested number of servers based on total entitlement assignments. These suggested number of servers are guidelines for your particular deployment requirements. Unique client requirements may require customization, which may differ from the listed number of servers.

For a description of possible production deployments, see "*Deployment Architectures*" in the *Installation Guide* in the *Installation Guide*.

Suggested Number of Servers

Data Set Ranges

	Small	Medium	Large	Custom
Total Assignments	<1M	1-5M	5-15M	15M+
Suggested # of Servers				
Application ^a	Discuss with Autonomous Identity Team (dependent on HA requirements)			
Database	2	2	3	Custom ^b
Analytics	1	2	3	Custom ^b
Admin ^c	1	1	1	1

^aDocker Swarm is designed to be highly available and may require 5 or more nodes. For a production deployment, the specific requirements can be discussed with the Autonomous Identity Team. Docker Swarm requirements are not correlated to data set size, but to front-end user requirements (for example, the number of users and frequency of use).

^bFor environments with more than 15 million assignments, server requirements will need to be specifically customized.

^cThe Ansible deployment server is seldomly used, except for initial deployment or potentially during future maintenance). If necessary, Ansible could be deployed on one of the front-end servers to eliminate the need for a standalone deployment server.

Production Technical Specifications

Autonomous Identity 2020.6.0 has the following technical specifications for production deployments:

Production Technical Specifications

	Application	Database	Analytics	Admin
--	-------------	----------	-----------	-------

Installed Components	Docker	Cassandra	Spark/Docker (Spark Master)	Docker
OS	CentOS	CentOS	CentOS	CentOS
# of Servers	See "Suggested Number of Servers"	See "Suggested Number of Servers"	See "Suggested Number of Servers"	1
RAM (GB)	32	32	64	32
CPUs	8	8	16	8
Non-OS Disk Space (GB) ^a	1000	1000	1000	1000
NFS Shared Mount	1 TB NFS mount shared across all Docker Swarm nodes (if more than 1 node is provisioned) at location separate from the non-OS disk space requirement. For example, <code>/shared</code> .	N/A	1 TB NFS mount shared across all Docker Swarm nodes (if more than 1 node is provisioned) at location separate from the non-OS disk space requirement. For example, <code>/shared</code> .	N/A
Networking	nginx: 443 Docker Manager: 2377 (TCP) Docker Swarm: 7946 & 4789 (UDP) + 7946 & 2049 (TCP)	Client Protocol Port: 9042 Cassandra Nodes: 7000	Spark Master: 7077 Spark Workers: Randomly assigned ports	N/A
Licensing	N/A using Docker CE free version	N/A	N/A	N/A
Software Version	Docker: 19.03.8	Cassandra: 3.11.2	Spark: 2.4.4 Docker: 19.03.8	Docker: 19.03.8
Component Reference	See below. ^b	See below. ^c	See below. ^d	See below. ^e

^aAt root directory "/"

^b<https://docs.docker.com/ee/ucp/admin/install/system-requirements/>

^c<http://cassandra.apache.org/doc/latest/operating/hardware.html>

<https://docs.datastax.com/en/dse-planning/doc/planning/planningHardware.html>

^d<https://spark.apache.org/docs/2.2.0/hardware-provisioning.html>

<https://spark.apache.org/docs/latest/security.html#configuring-ports-for-network-security>

^e<https://docs.ansible.com/ansible/latest/index.html>

Chapter 3

Deployment Checklist

Use the following checklist to ensure key considerations are covered for your 2020.6.4 deployment:

Deployment Checklist

Check <input checked="" type="checkbox"/>	Requirement	Details
Access		
<input type="checkbox"/>	Remote Access	The Autonomous Identity Team is a global team. To support the needs of client teams, remote access to all servers is required for deployment and support of product.
<input type="checkbox"/>	Root Access	<p>Root access is required to run required package installations (YUM), perform Docker installation, Docker Swarm-based installation applicable boxes, and potential troubleshooting.</p> <p>Please discuss with delivery team if this requirement is a concern. If so, submit a specified contact to run admin tasks.</p>
<input type="checkbox"/>	Service Account	The Autonomous Identity Team should have access to a single service account user (e.g., "autoid"), which will be used throughout as the primary owner of Autonomous Identity specific directories. Specific requirements regarding the service account are specified in this section.
<input type="checkbox"/>	File Transfer Process	The Autonomous Identity Team require access to a file transfer process, which lets specified packages be transferred from the vendor to the client infrastructure.
Service Account		
<input type="checkbox"/>	Autonomous Identity Team Access	Autonomous Identity team members must be able to switch to this user after logging in to the servers
<input type="checkbox"/>	SSH Ability	The service account must be able to passwordless SSH between all Autonomous Identity servers; preferred method is RSA SSH key authentication.
<input type="checkbox"/>	Default Shell	The default shell of the service account must be Bash.
<input type="checkbox"/>	Directory Ownership	<p>Ownership of the following directories must be given to the Service Account.</p> <ul style="list-style-type: none"> • /data (all servers) • /opt/autoid (all servers) • /shared (if applicable - Docker & Spark servers)

		<ul style="list-style-type: none"> • /tmp/zoran (R, W, E required at least)
<input type="checkbox"/>	Docker Commands	The service account must have permissions to run Docker commands. Note that Docker should NOT need to be installed as a prerequisite; this will be installed by deployment team.
Networking/Internet		
<input type="checkbox"/>	Access to the Internet	If available, the front-end servers downloads the required Docker images from the official Autonomous Identity image repository.
<input type="checkbox"/>	SSL Certificates	<p>If SSL is being implemented, SSL certificates are required for the UI, Cassandra nodes, and Spark nodes. These certificates can be generated using one of the following four options:</p> <ul style="list-style-type: none"> • Self-signed certificates for all 3 components • Valid certificate for the UI and self-signed certificates for Cassandra and Spark nodes (self-signed certs only used in server-server traffic) • Valid and separate certificates for the UI, Cassandra, and Spark • *.domainname.com certificate (wildcard)
<input type="checkbox"/>	Ports Open (Internal)	All internal ports specified in the Networking section of the Environment Specifications need to be opened for the specified servers.
<input type="checkbox"/>	Ports Open (external browser)	<p>The following ports must be accessible from a web browser within the client network:</p> <ul style="list-style-type: none"> • 443 (Front-end) • 8080 (Spark) • 8081 (Spark)
Required Packages		
<input type="checkbox"/>	Dependencies	<p>The following packages must be installed on specified servers as prerequisites:</p> <ul style="list-style-type: none"> • All Servers: Python 2.7 or Python 3 (3.5+) • Cassandra Servers: java-1.8.0-openjdk-devel.x86_64 • Analytics Servers: java-1.8.0-openjdk-devel.x86_64
Other		
<input type="checkbox"/>	Infrastructure Support POC	A point-of-contact (POC) with sufficient access to the infrastructure is required. The POC can support in case of infrastructure blockers arise (e.g., proxy, account access, or port issues).
<input type="checkbox"/>	SELinux	SELinux must be disabled on the Docker boxes. The package "container-selinux" must be present (this can be done as part of the root scripts described in the "Root Access" category).
<input type="checkbox"/>	Components Not Pre-installed	<p>The following software must NOT be pre-installed on the box:</p> <ul style="list-style-type: none"> • Docker

- Cassandra

- Spark

If any do come pre-installed, discuss the details with the Delivery Team ahead of time.

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.
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 <code>conf_thresh</code> 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.