Nutanix Brain Deployment Guide

Version: October 23, 2025

Table of Contents

Introduction

This guide is intended to help customers or partners deploy a virtual Brain appliance in Nutanix environments. A Nutanix Brain appliance can be used in Vectra AI Platform deployments that use either the Respond UX or the Quadrant UX. The Respond UX is served from Vectra's cloud and the Quadrant UX is served locally from the Brain appliance. For more detail on Respond UX vs Quadrant UX please see Vectra Analyst User Experiences (Respond vs Quadrant).

This guide will cover basic background information, connectivity requirements (firewall rules that may be needed in your environment), licensing, deployment, and next steps. One of the below guides should be the starting point for your overall Vectra deployment:

- Vectra Respond UX Deployment Guide
- Vectra Quadrant UX Deployment Guide

Either of the above guides cover basic firewall rules needed for the deployment and initial platform settings. Virtual Sensor (Nutanix, Hyper-V, KVM, AWS, Amazon, and GCP) configuration and pairing and covered in their respective guides. Physical appliance pairing is covered in the Vectra Physical Appliance Pairing Guide. Please see the Vectra Product Documentation Index on the Vectra support site for additional documentation including deployment guides for CDR for M365 / IDR for Azure AD, CDR for AWS, and CDR for Azure.

About Nutanix Brain Images

The .zip file that contains the .ova image used to deploy a Brain in Nutanix along with a deployment script is made available on the Vectra Customer Portal which is part of Vectra Support. Vectra periodically updates the base image used for Nutanix Brain deployment. It is a best practice to always download the latest .zip from the Vectra Customer Portal prior to deployment of a new Nutanix Brain. Brains that are connected to Vectra are updated automatically according to the settings on that Brain. Offline updates are also possible for Quadrant UX deployments only. Please see the Vectra Quadrant UX Deployment Guide for more details regarding offline updates. Please see the Offline Updates article on the Vectra Support site for instructions on how to apply offline updates.

Nutanix Brain Requirements and Throughput

For use in Respond UX or Quadrant UX deployments:

Resource Type	Requirement		
Performance[1]	Coming Soon	Coming Soon	10 Gbps
CPU	8 Cores	16 Cores	32 Cores
Memory	64 GB RAM	128 GB RAM	256 GB RAM
Drive (OS, Data) Requires 260 MB/s	128 GB, 512 GB	128 GB, 512 GB	128 GB, 512 GB
Max Paired Sensors	Coming Soon	Coming Soon	100
Max Simultaneous Tracked Hosts[2]	Coming Soon	Coming Soon	150,000
Nutanix Versions Supported	AOS 6.8.1 and higher with Prism Central (and v3 API) available		

For use ONLY in Respond UX for Network deployments (using network Sensors with the Respond UX):

Resource Type	**Requirement** (Respond UX for Network Deployments Only)	
Performance[1]	Coming Soon	Coming Soon
CPU	4 Cores	6 Cores
Memory	48 GB RAM	48 GB RAM
Drive (OS, Data) Requires 260 MB/s	128 GB, 512 GB	128 GB, 512 GB
Max Paired Sensors	Coming Soon	Coming Soon
Max Simultaneous Tracked Hosts[2]	Coming Soon	Coming Soon
Nutanix Versions Supported	AOS 6.8.1 and higher with Prism Central (and v3 API) available	

- Please note: At the initial availability of the Nutanix Brain image, only the 32-core version is supported by Vectra. Vectra plans to make the other sizes (greyed out in the chart) available in the future.

[1] Performance represents the aggregate bandwidth observed on the capture interfaces of any Sensors that are paired to the Brain. Guidance is for average traffic mixes. Traffic mixes that skew toward larger flows (like file transfers) will perform better than traffic mixes that skew towards smaller flows (like DNS) as they produce more metadata.

[2] Refers to how many hosts the Brain can track simultaneously (open host sessions). Brains retain and display data for larger numbers of hosts, this only refers to how many hosts the system can process metadata for simultaneously.

- The **virtual CPU MUST** support the pdpe1gb cpu flag (1GB Large Pages) – More information, and a minimum SSE instruction level of 4.2, and must support the POPCNT (population count) instruction. This requires the hypervisor host to be running one of the following processors or later:
    - Intel Nehalem (2008) processors and newer
    - AMD Bulldozer (2011) processors and newer
- Vectra Nutanix based Brains do not support Mixed Mode deployment. They can only be used in Brain mode.
- Vectra Nutanix based Brains support running in FIPS mode. Note that the underlying hardware must also be FIPS compliant (it must support the RDRAND CPU instruction). To configure FIPS more, once deployed login to the CLI of the Brain and use the following commands to enable/disable FIPS mode.

```
set security-mode fips
set security-mode default
```

- Vectra Nutanix based Brains do not support Direct PCI or SR-IOV passthrough.
- Vectra Nutanix based Brains do support paravirtualized NICs. Vectra uses a VirtIO NIC for the Nutanix Brain.
- Vectra recommends that Brains are configured to use storage local to the hypervisor and are not stored on a SAN. Vectra Brains require extremely high throughput from their disk storage and this throughput cannot normally be sustained by SAN systems without impact to other SAN users.
- Live Migration is not explicity supported by Vectra.
    - If you do use Live Migration and encounter any issues, support from Vectra will be best effort only.
    - It is a best practice to set VM-Host affinity to pin the Brain to a node with adequate resources where satisfactory Brain performance test results are achieved.

Connectivity Requirements (Firewall Rules)

The Vectra AI Platform uses several TCP/UDP ports for different communication purposes. This document will detail basic requirements for initial setup and pairing. Many features and integrations are optional and not in scope for this guide. Additional connectivity requirement guidance is in the Vectra Respond UX Deployment Guide or Vectra Quadrant UX Deployment Guide. For full detail on all possible firewall rules that might be required in your environment please see Firewall requirements for Vectra appliances on the Vectra support site.

Vectra Cloud Connectivity

- For this document, the portions of the Vectra AI Platform that reside in Vectra's cloud are referred to as the Vectra cloud.
    - This does not refer to any specific service offering.
- Please check each category below to see if it is applicable to your deployment and if rules are required in your environment to enable the required connectivity.
    - For rule categories that have multiple region options, it is only necessary to put rules in place to allow connectivity to the region that your Vectra tenant is deployed in. This region should be visible in the URL used to access the Respond UX.
        - i.e. [tenant_id].ew1.prod.vectra-svc.ai is used for EU deployments (ew1).
- RUX for Network refers to a RUX deployment that has enabled network data sources (sensors).
    - This means you have a Brain somewhere in your premises (data center or public cloud) that is paired with network sensors (virtual or physical) to capture network traffic and distill a metadata stream for processing by the Brain appliance.
    - Please refer to the Vectra Respond UX Deployment Guide for more details.
- Please refer to the table below to see applicability of the various categories.
- The "For Brain or User's Host" column should be interpreted as follows:
    - Brain – Rules required for the Brain to the Vectra Cloud.
    - User's Web Browser – Rules required for the User's Web Browser to the Vectra cloud.

Rule Category	Required For	For Brain or User's Host
RUX for Network GUI Synchronization	RUX for Network Deployments	Brain
Auth Gateways	RUX for Network Deployments Quadrant UX deployments of CDR for M365, IDR for Azure AD, and CDR for AWS.	Brain
RUX Metadata Forwarding	RUX for Network Deployments	Brain
RUX Research Metadata Forwarding	RUX for Network Deployments	Brain
RUX Analyst/Admin Access	All RUX Deployments	User's Web Browser
RUX Static Asset CDN	All RUX Deployments	User's Web Browser
RUX Customer File Upload	All RUX Deployments	User's Web Browser

RUX for Network GUI Synchronization

- **Required for:** All RUX for Network deployments.
- This is used to synchronize configurations between the Brain appliance and your Vectra tenant.
- This communications channel is initiated from the Brain to the endpoint in your Vectra tenant's region.
- The protocol and ports in use for each entry is the same: **Websocket and HTTPS over TCP/443**

Fully Qualified Domain Name (FQDN)	IP(s)	Region	Initiated From
main-cbi-tunnel-uw2.app.prod.vectra-svc.ai	Dynamic	US	Brain
main-cbi-tunnel-ew1.app.prod.vectra-svc.ai	Dynamic	EU	Brain
main-cbi-tunnel-ec2.app.prod.vectra-svc.ai	Dynamic	Switzerland	Brain
main-cbi-tunnel-cc1.app.prod.vectra-svc.ai	Dynamic	Canada	Brain
main-cbi-tunnel-as2.app.prod.vectra-svc.ai	Dynamic	Australia	Brain

Auth Gateways

- **Required for:**
  - All Respond UX for Network Deployments.
  - Quadrant UX deployments of CDR for M365, IDR for Azure AD, and CDR for AWS.
    - Your Brain must be able to securely access the Vectra cloud over TCP/443 HTTPS connections to enable detection events from these products to be reported to your UI.
- In Respond UX for Network deployments, the Brain forwards network detections, entities, host sessions, and any selective PCAPs (Vectra Packet Capture) to your Vectra tenant via this connection.
- This communications channel is initiated from your Brain to the endpoint in your Vectra tenant's region.

Fully Qualified Domain Name (FQDN)	IP(s)	Protocol / Ports	Region	Initiated From
authgateway.uw2.public.app.prod.vectra-svc.ai	54.245.33.175 52.42.70.176 100.21.109.72 52.26.91.157	HTTPS - TCP/443	US	Brain
authgateway.ew1.public.app.prod.vectra-svc.ai	54.171.40.108 54.246.213.148 54.75.47.147	HTTPS - TCP/443	EU	Brain
authgateway.ec2.public.app.prod.vectra-svc.ai	16.62.18.237 16.62.142.98 51.96.54.201	HTTPS - TCP/443	Switzerland	Brain
authgateway.cc1.public.app.prod.vectra-svc.ai	3.96.112.208 52.60.211.221 15.222.69.161	HTTPS - TCP/443	Canada	Brain
authgateway.as2.public.app.prod.vectra-svc.ai	13.54.11.66 13.55.79.24 13.55.106.102	HTTPS - TCP/443	Australia	Brain

RUX Metadata Forwarding

- **Required for:** All Respond UX for Network Deployments.
- Network metadata is forwarded to AWS S3 buckets and processed to make it available for features such as Instant Investigation and Advanced Investigation in the Respond UX.
- This communications channel is initiated from your Brain to the endpoint in your Vectra tenant's region.
- The protocol and ports in use for each entry is the same: **HTTPS over TCP/443**

Fully Qualified Domain Name (FQDN)	IP(s)	Region	Initiated From
cbo-upload-network-metadata-forwarder-uswt2-371371611652.s3-accesspoint.us-west-2.amazonaws.com	Dynamic	US	Brain
cbo-upload-network-metadata-forwarder-euwt1-371371611652.s3-accesspoint.eu-west-1.amazonaws.com	Dynamic	EU	Brain
cbo-upload-network-metadata-forwarder-eucl2-371371611652.s3-accesspoint.eu-central-2.amazonaws.com	Dynamic	Switzerland	Brain
cbo-upload-network-metadata-forwarder-cacl1-371371611652.s3-accesspoint.ca-central-1.amazonaws.com	Dynamic	Canada	Brain
cbo-upload-network-metadata-forwarder-apse2-371371611652.s3-accesspoint.ap-southeast-2.amazonaws.com	Dynamic	Australia	Brain

RUX Research Metadata Forwarding

- **Optional but highly recommended for:** All Respond UX for Network Deployments.
- Research metadata from precursor algorithms are used to improve model quality and reduce detection noise.
- This communications channel is initiated from your Brain to the endpoint in your Vectra tenant's region.
- The protocol and ports in use for each entry is the same: **HTTPS over TCP/443**

Fully Qualified Domain Name (FQDN)	IP(s)	Region	Initiated From
cbo-upload-network-precursors-uswt2-371371611652.s3-accesspoint.us-west-2.amazonaws.com	Dynamic	US	Brain
cbo-upload-network-precursors-euwt1-371371611652.s3-accesspoint.eu-west-1.amazonaws.com	Dynamic	EU	Brain
cbo-upload-network-precursors-eucl2-371371611652.s3-accesspoint.eu-central-2.amazonaws.com	Dynamic	Switzerland	Brain
cbo-upload-network-precursors-cacl1-371371611652.s3-accesspoint.ca-central-1.amazonaws.com	Dynamic	Canada	Brain
cbo-upload-network-precursors-apse2-371371611652.s3-accesspoint.ap-southeast-2.amazonaws.com	Dynamic	Australia	Brain

RUX Analyst/Admin Access

- **Required for:** All Respond UX deployments.
- Any analyst or admin that wishes to access the Respond UX will need to ensure that their browser can reach their Vectra tenant to login and access the UI.
- This communications channel is initiated from the user's host.
- The protocol and ports in use for each entry is the same: **HTTPS over TCP/443**

Fully Qualified Domain Name (FQDN)	IP(s)	Region	Initiated From
[tenant_id].uw2.portal.vectra.ai	Dynamic	US	User's Web Browser
[tenant_id].ew1.portal.vectra.ai	Dynamic	EU	User's Web Browser
[tenant_id].ec2.portal.vectra.ai	Dynamic	Switzerland	User's Web Browser
[tenant_id].cc1.portal.vectra.ai	Dynamic	Canada	User's Web Browser
[tenant_id].as2.portal.vectra.ai	Dynamic	Australia	User's Web Browser

RUX Static Asset CDN

- **Required for:** All Respond UX deployments.
- The Respond UX has certain static assets (HTML, CSS, JS) that are required to serve the web application hosted by a CDN (Content Delivery Network).
- This communications channel is initiated from the user's host.

Fully Qualified Domain Name (FQDN)	Protocol / Ports	IP(s)	Region	Initiated From
dd6462tdmvp79.cloudfront.net dpew7prsvwbf0.cloudfront.net	HTTPS - TCP/443	Dynamic	All	User's Web Browser

RUX Customer File Upload

- **Required for:** All Respond UX deployments.
- This communications channel is used for:
  - Vectra Match deployments and will allow upload of rulesets.
  - PCAP download from the Vectra Cloud for Selective PCAP (Vectra Packet Capture)
  - Additional capabilities are planned for future releases.
    - It is recommended to put rules in place even if you don't use Match or Selective PCAP.
- This communications channel is initiated from the user's host.

Fully Qualified Domain Name (FQDN)	Protocol / Ports	IP(s)	Region	Initiated From
prd-main-customerfiles-580786928539-uswt2.s3.amazonaws.com	HTTPS - TCP/443	Dynamic	US	User's Web Browser
prd-main-customerfiles-580786928539-euwt1.s3.amazonaws.com	HTTPS - TCP/443	Dynamic	EU	User's Web Browser
prd-main-customerfiles-580786928539-eucl2.s3.amazonaws.com	HTTPS - TCP/443	Dynamic	Switzerland	User's Web Browser
prd-main-customerfiles-580786928539-cacl1.s3.amazonaws.com	HTTPS - TCP/443	Dynamic	Canada	User's Web Browser
prd-main-customerfiles-580786928539-apse2.s3.amazonaws.com	HTTPS - TCP/443	Dynamic	Australia	User's Web Browser

Connectivity Requirements – General

Source	Destination	Protocol/Port	Description	QUX-RUX-Both
Admin hosts	Brain / Sensors	TCP/22 (SSH)	CLI access for Brain and Sensors.	Both
Admin hosts	Brain	TCP/443 (HTTPS)	Web UI of Brain appliances (Quadrant UX). Redirect / Status of Brain (Respond UX).	Both
Brain	update2.vectranetworks.com (54.200.156.238)	TCP/443 (HTTPS)	Automatic updates. Pairing keys for physical sensors.	Both
Brain	api.vectranetworks.com (54.200.5.9)	TCP/443 (HTTPS)	Health monitoring, algorithm support, reverse lookups for external IPs, Vectra Threat Intelligence, additional detection content.	Both
Brain	rp.vectranetworks.com (54.200.156.238)	TCP/443 (HTTPS)	Used only for Brains deployed in IaaS clouds. Used for authentication and verification (integrity check of the file system).	Both
Brain	rs.vectranetworks.com (74.201.86.229)	TCP/443 and UDP/9970	Remote Support. OpenVPN type if using firewall with App ID rules.	Both
Brain	DNS servers (as configured)	TCP/53, UDP/53	Both TCP and UDP are required for normal operation.	Both
Brain	NTP servers (as configured) Default is ntp.ubuntu.com	UDP/123	Time synchronization.	Both
Brain	SMTP servers (as configured)	TCP (as configured)	Email alerting (optional but suggested).	Quadrant UX
Sensors, Stream	Brain	TCP/22 (SSH), TCP/443 (HTTPS)	Pairing, metadata transfer, and ongoing communication.	Both
Brain	Sensors, Stream	TCP/22 (SSH)	Remote management and troubleshooting.	Both
Brain	Recall collector	TCP/443 (HTTPS)	Destinations provisioned after enabled.	Quadrant UX
Brain	metadata.vectra.ai (100.20.236.31, 44.229.57.246, 44.228.37.60, 44.228.101.87)	TCP/443 (HTTPS)	Optional anonymized metadata sharing to contribute to future algorithm development.	Quadrant UX

- Customers should note that the following IP ranges will conflict with remote support capability:
  - 192.168.72.0/21 and 192.168.80.0/21
  - For remote support outside of screen sharing sessions, care should be taken to number the management network interface (MGT) used on any appliance (Brains and Sensors) outside of the above ranges. If your management network interface (MGT) is numbered in either of these ranges, remote support access will not function.

Licensing and Deployment Overview

Vectra appliance code has been encrypted to protect Vectra's intellectual property, and a license is required to enable successful decryption of the file system during boot. The licensing for Vectra NDR (formerly Detect for Network) running on Nutanix Brains also governs the ability of the system to create Detections. After deployment, if your Vectra NDR license expires, Detection algorithms will cease operation until a valid license is applied. Metadata related functionality such as Instant Investigation, Investigate, Recall (QUX only), and Stream operation are unaffected if your Vectra NDR license expires and other licenses are valid.

Scripted deployment of a Brain appliance in Nutanix environments requires Prism Central and v3 APIs. Prism Element is not supported. Vectra provides a .zip archive on the Vectra Support Portal that contains the Nutanix Brain .ova image along with a python script that performs the deployment using Prism Central API calls. For customers who wish to deploy manually instead of using the python script, alternate deployment instructions are also provided.

Licensing Enforcement

Vectra NDR (Detect for Network) supports licensing functionality regardless of the type of deployment (physical appliance, cloud IaaS, Nutanix, etc). All versions will be able to see license status and enable requests for and application of licenses. Enforcement of NDR licensing is only enabled on Nutanix and VMware Brains. Any other Brain type does NOT currently have licensing for NDR enforced. Vectra does plan to add licensing enforcement for other Brain types in the future. It is recommended that all customers work with their account teams to ensure their licensing is up to date. Please refer to the following table for additional detail:

Product	Deployment Type	License Enforcement
NDR (Detect for Network)	Nutanix or VMware Brain	Algorithms stop producing Detections when expired.
NDR (Detect for Network)	Physical or Cloud Brains	Not currently enforced. Planned for future (timing TBD).

Other Vectra products such as Recall, Stream, CDR for Azure, CDR for M365, and IDR for Azure AD are also licensed but enforcement of the license is a matter of contract compliance between sales teams and customers.

Deployment Overview

The main steps for the deployment are summarized below. For additional detail see Brain Deployment in Nutanix.

- Download Nutanix Brain appliance image from https://support.vectra.ai/vectra/additional-resources and upload the Brain image to your Prism Central OVA store.
    - You must be logged into your Vectra support account, to see the download option.
    - Unzip the file to access the OVA and Python deployment script.
    - Upload the OVA image to Prism Central.
- Deploy the OVA in Nutanix and power on the appliance using one of the below methods.
    - For either method below, please continue reading for additional details and requirements.
        - Deploying the OVA Using the Python Script (Recommended)
        - Deploying the OVA Manually (Alternate Method)
- Browse to the IP assigned to the Brain's management interface to see the initial boot status messages. The status message screen will update but a manual refresh is required to display any new information.

  - When you are able to see the “System Setup and Provisioning” screen, enter proxy information if required for your environment and then enter the “License configuration” screen.
- Copy the licensing request code from the Vectra UI.
- Retrieve a license by pasting the code in the licensing portal.
- Copy the license once generated and paste it into the Vectra UI “Licensing Information” box.
- If your Brain deployment will not be online (connected to Vectra’s provisioner/updater system, this is only supported for Quadrant UX deployments), check the “Offline” box to enable an offline deployment. This means all licensing functions will be done manually offline.
  - If the “Offline box is not checked prior to hitting the “Save” button, they deployment will fail and will need to be started over with a fresh deployment of the OVF template.
  - Offline updates mode is automatically enabled when selecting the offline deployment mode.
    - Nutanix Brains deployed in offline mode can **never** be updated online.
- Click “Save” on the licensing configuration screen.
- After the license is validated, the file system will be decrypted, a performance test will be run, the Brain will reboot, and the Brain will reach out to the Vectra provisioning server to complete provisioning (online Brains). Finally, a success message will be presented with a button to redirect to the main UI login screen. Offline Brains but do not need to communicate with the provisioning server and can validate the license locally.
  - For Respond UX deployments, please see the rest of the deployment process in the Vectra Respond UX Deployment Guide. You should not login and configure anything in the Quadrant UX (which would be available at this point) if you will be performing a Respond UX deployment.
- Initial login credentials for the UI are **`admin / changethispassword`**. Initial login credentials for the SSH access to the CLI are **`vectra / changethispassword`**.
  - You will be asked to change the password after the initial login.
- Complete your configuration using instructions available in the Vectra Respond UX Deployment Guide or Vectra Quadrant UX Deployment Guide.

Brain Deployment in Nutanix

Requirements

- IP address, subnet mask, default gateway, and hostname for the Management interface of the Brain (DHCP is also supported).
  - If DHCP is used, a reservation should be created to keep the IP consistent for Sensors that may pair via IP instead of Hostname.
- DNS server addresses.
- Nutanix Prism Central with v3 API accessible to the user who will perform the deployment. The specific permissions required are:
  - Cluster – View Cluster
  - OVA – View OVA
  - Subnet – View Subnet
  - AHV VM – Create Virtual Machine
- Current login to a fully approved Vectra Support Portal account.
  - Accounts that are self-registered and not fully approved on the Vectra Support Portal will not have the license request option enabled.
- An open Proof of Value (Proof of Concept or Trial) that you are working with Vectra or a Vectra partner or a valid entitlement to Vectra NDR through purchase.
  - The licensing system cannot provide licenses for customers who are not currently entitled to a license through a trial or purchase.

Downloading the Latest Nutanix Brain Files

The current Nutanix Brain files (.zip containing .ova and python deployment script) can be downloaded from the Vectra Customer Support Portal after logging in. The URL for the download page is: https://support.vectra.ai/vectra/additional-resources and then click on the “Download” tab. Click on the “Nutanix Brain Cverify the download completed successfully. Always download a current copy when you go to deploy a new Nutanix Brain for your organization. This will save time during deployment as fewer updates will need to be downloaded after deployment. Make this file available via a URL or the local filesystem where you will run the vSphere client from.

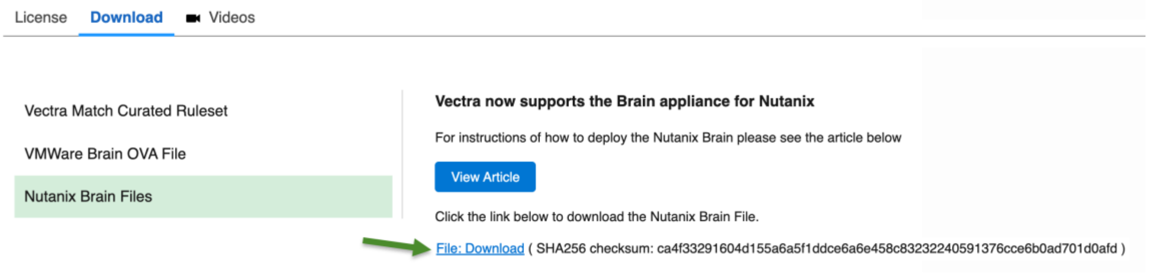

Vectra now supports the Brain appliance for Nutanix

For instructions of how to deploy the Nutanix Brain please see the article below

Click the link below to download the Nutanix Brain File.

File: Download (SHA256 checksum: ca4f33291604d155a6a5f1ddce6a6e458c83232240591376cce6b0ad701d0afd)

Uploading the OVA Image to Prism Central

- After downloading the .zip file from the Vectra support portal, unzip it using the tool of your choice.
  - Save the python script to the location of your choice. The location must have reachability to your Prism Central v3 APIs.
- Upload the .ova file to your OVA store in Prism Central.
  - Navigate to *Infrastructure > Compute > OVAs*.
  - Click “Upload OVA” and fill out the dialog.
    - Choose your AHV cluster.
    - Pick a name for the OVA. This will be used later selecting the image with the deployment script.
    - The Checksum file is not required as the OVA will automatically be verified by Nutanix based on information included in the OVA file.
    - Choose your file and click “Upload”.
  - It will take a few minutes to upload the OVA and then Prism Central will perform validation of the OVA.
  - You can check the status of these tasks in the All Tasks as seen below:

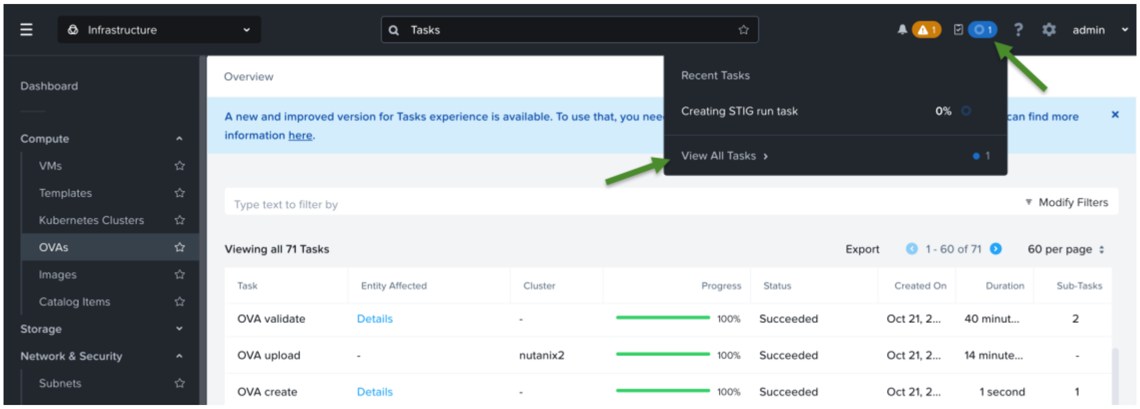

Creating STIG run task 0%

A new and improved version for Tasks experience is available. To use that, you nee… can find more information here.

Viewing all 71 Tasks

Task	Entity Affected	Cluster	Progress	Status	Created On	Duration	Sub-Tasks
OVA validate	Details	-	100%	Succeeded	Oct 21, 2...	40 minut...	2
OVA upload	-	nutanix2	100%	Succeeded	Oct 21, 2...	14 minute...	-
OVA create	Details	-	100%	Succeeded	Oct 21, 2...	1 second	1

  - When the validation has completed, you can move on to deployment.
- **!! Choose one of the following two deployment methods for the OVA.**

Deploying the OVA Using the Python Script (Recommended)

- The python deployment script can be run from any location that has reachability to your Prism Central APIs.
  - Your python interpreter must be v3.7 or higher.
  - Please see below for a sample interactive run using DHCP (yellow highlights indicate user input):

```
❯ python Vectra\ Nutanix\ Brain.py --interactive --insecure
Enter Prism Central url (include port if applicable): https://your_prism_central:9440/
Enter Prism Central username: admin
Enter Prism Central password: ****Redacted****
INFO:vectrantnx:logging into prism central
To which cluster should the VM be deployed?
Choose one of the following objects:
UUID                                    Name
000639fc-7386-2da4-7938-bc2411a17fee    nutanix2
4a061048-29a6-4b6b-8f6d-f5bca68241e7    prism_central
Enter a UUID or name: nutanix2
Which ova should be used to deploy the VM?
Choose one of the following objects:
UUID                                    Name
42b71f41-60cc-43a2-ba72-b8c1b9975ad8    nutanix-vhe-image-10-8-2025
c593d1cb-1813-42a9-8c62-731b5c3f3e4b    Nutanix Brain - GA 9.6.0-16-10
Enter a UUID or name: Nutanix Brain - GA 9.6.0-16-10
What should be the name of this VHE? Example_Nutanix_Brain
Which subnet should be used for the management interface?
Choose one of the following objects:
UUID                                    Name
f5d8c6bb-8ad9-42c5-b977-8bafbba10b81    default-subnet
5a4a32d9-76f6-4b40-8eda-eb2fa46899bc    test
Enter a UUID or name: default-subnet
What will be the size of this VHE? Profiles marked with RespondUX may only be used with RespondUX.
Valid choices are:
32CORE
Enter choice: 32CORE
Enable dhcp? Static networking options will be asked if False.
Valid choices are:
True
False
Enter choice: True
Should the VHE boot directly into RespondUX mode? This enables pairing with the Vectra AI Platform utilizing the Respond
UX and disables the local UI.
Valid choices are:
True
False
Enter choice: False
INFO:vectrantnx:gathering info for deployment...
INFO:vectrantnx:deploying...
deployment started, check Prism Central for status
```

- Help is available if you use the `--help` argument when running the script:

```
❯ python Vectra\ Nutanix\ Brain.py --help
usage: Vectra Nutanix Brain.py [-h] [--interactive] [--insecure] [--nutanix-host NUTANIX_HOST] [--nutanix-user
NUTANIX_USER] [--nutanix-pass NUTANIX_PASS] [--cluster CLUSTER] [--ova OVA] [--name NAME] [--mgt-network MGT_NETWORK] [--
profile {32CORE}] [--dhcp {True,False}] [--hostname HOSTNAME] [--ip-address IP_ADDRESS] [--ip-netmask IP_NETMASK] [--ip-
gateway IP_GATEWAY] [--dns DNS] [--respondux {True,False}]

Vectra Nutanix deployment tool

options:
  -h, --help            show this help message and exit
  --interactive         Ask interactively for missing options
  --insecure            Disable tls certificate verification
  --nutanix-host NUTANIX_HOST
                        URL of your Prism Central instance, including port if applicable
  --nutanix-user NUTANIX_USER
                        Prism Central username
  --nutanix-pass NUTANIX_PASS
                        Prism Central password
  --cluster CLUSTER     UUID or name of cluster to which to deploy the VM
  --ova OVA             Name or uuid of ova to deploy
  --name NAME           Name of the virtual machine to be created
  --mgt-network MGT_NETWORK
                        Name or uuid of subnet to which to connect the management NIC
  --profile {32CORE}    Deployment profile (size) of VHE
  --dhcp {True,False}   Set to True to enable DHCP. Static networking options can be omitted if DHCP is enabled.
  --hostname HOSTNAME   The hostname or FDQN for the virtual machine.
  --ip-address IP_ADDRESS
                        Static IP address to assign to the management interface.
  --ip-netmask IP_NETMASK
                        The netmask for the management interface.
  --ip-gateway IP_GATEWAY
                        Default gateway for the management interface.
  --dns DNS             The DNS servers for the VM (comma separated). You may specify up to 2 DNS servers.
  --respondux {True,False}
                        Set to True to boot the VHE directly into RespondUX mode. This enables pairing with the Vectra AI
Platform utilizing the Respond UX.
```

- Additional guidance for using the script:
    - If you would like to pass all arguments to the script instead of running interactively, it is supported.
    - Interactive deployment is recommended because the script will display options for the cluster, OVA image, and subnet names or UUIDs and you will not need to look them up elsewhere.
    - The URL can include a trailing / or it can be left off, both formats will work but you should include the full URL including port, and not just the hostname.
    - After the last question is answered or if all arguments were passed to script for a non-interactive deployment, the execution is nearly instantaneous, and you can then find the VM in Prism Central at *Infrastructure* > *Compute* > *VMs*.
    - When choosing static IP addressing, additional questions will be asked during the script execution to capture the static IP details.
        - Brains that start with DHCP can be set to static addressing after deployment. See Setting a static IP and DNS after initial DHCP deployment for details.
- Per the Nutanix Brain Requirements and Throughput section, Vectra plans to make additional configurations available in the future.
    - Initially, only the 32 core Brain is supported.
    - Once available, the 4 and 6 core Respond UX specific configuration can only be used in Respond UX deployments. They are NOT supported for Quadrant UX deployments.
    - **RespondUX** – Choose this option if you are doing a Respond UX for Network deployment.
        - When this option is selected, the Brain will boot directly into a state that is ready to be linked to the Vectra Cloud for use with the Respond UX. There will be no local Quadrant UX GUI served as there normally would be for a standard Nutanix Brain deployment before it is linked with Vectra for use with the Respond UX. Vectra personnel will still need to link your Brain to your Respond UX tenant to complete your deployment.
        - This option should be selected for any Respond UX for Network deployment. i.e. You still need to pick this option even if you previously chose the “6CORE_RespondUX” configuration.
- Once deployed, you can move on to Initial boot up and licensing.

Deploying the OVA Manually (Alternate Method)

- Select the check box next to the OVA you uploaded and choose *Actions* > *Deploy as VM*. Follow the instructions below for each screen.
    - The OVA should be visible in *Infrastructure* > *Compute* > *OVAs* in Prism Central.
- A dialog will open with Configuration, Resources, Management, and Review panes.

- 1 - Configuration
  - Set Name and Description as desired.
  - Choose the cluster you wish to deploy on.
  - Under VM Properties, set the CPUs and Memory to one of the following configurations. In all cases, set Cores Per CPU to 1.
    - 4 CPUs / 48GiB RAM (Respond UX ONLY) – Coming soon, do NOT use until Vectra approved by Vectra!
    - 6 CPUs / 48GiB RAM (Respond UX ONLY) – Coming soon, do NOT use until approved by Vectra!
    - 8 CPUs / 64GiB RAM – Coming soon, do NOT use until approved by Vectra!
    - 16 CPUs / 128GiB RAM – Coming soon, do NOT use until approved by Vectra!
    - 32 CPUs / 256GiB RAM
  - Under Advanced Settings, do NOT enable Memory Overcommit.
- 2 - Resources
  - Do NOT change Disks, Boot Configuration, or Shield VM Security Settings!
  - Select the edit (pencil) icon next to the NIC and configure it.
    - Choose the desired subnet.
    - For Network Connection State, choose “Connected”.
    - For NIC Configuration, choose Access and then “Save” you NIC configuration.
- 3 - Management
  - “Default-Storage” policy can be enabled if this is your standard.
    - If Performance testing indicates too low of a disk throughput, you may need to alter or remove the policy later if the policy is the reason for the lower performance.
  - Categories are NOT required.
  - Timezone must be set to UTC.
    - Timezone can later be set in your Vectra UI if you desire to have a different timezone.
  - Check "Use this VM as an Agent VM" if you want the Brain to boot before all other VMs.
    - This is **not required**, but is supported if desired.
  - Under Guest Customization, choose Script Type: Cloud-init (Linux) and Configuration Method: Custom Script
  - Add your configuration to the template below and paste it into the "Startup Script" box.
    - Download the script from the support KB as it won’t format correctly if copied from the .pdf.
    - Look for the “Startup_Script_For_Manual_Deployment.txt”.

```
#vectra

# Set this to True to enable DHCP. Static networking options can be left blank if DHCP is enabled.
dhcp=False
# The hostname or FDQN for the virtual machine.
hostname=
# Static IP address to assign to the management interface.
ip_address=
# The netmask for the management interface.
ip_netmask=
# Default gateway for the management interface.
ip_gateway=
# The DNS servers for the VM (comma separated). You may specify up to 2 DNS servers.
dns=

# Set this to True to boot the VHE directly into RespondUX mode. This enables pairing with the Vectra AI
Platform utilizing the Respond UX.
network_to_cloud=False
```

- 4 – Review
  - Review your configuration, make changes if required, and when ready, click “Create VM”.
  - After creation, you can then find the VM in Prism Central at *Infrastructure* > *Compute* > *VMs.*

Initial Boot Up and Licensing

After deploying the Brain VM into your Nutanix cluster using one of the two methods above, it will need to be licensed during the initial boot process.

- Power on your Brain VM.
  - The VM should now be in Prism Central at *Infrastructure* > *Compute* > *VMs.*
  - Right click on it, chose Power Operations > Power On.
- Once the UI is available (this will be a few minutes after power on), use your browser to connect to your Brain VM over HTTPS (using the IP assigned statically, via DHCP, or via hostname if your Brain is in your DNS).
  - If watching your console and you hit ESC when you see the Ubuntu boot screen you will see system messages.
  - You may need to switch between console views in Nutanix to see boot messages or access the login prompt after the initial licensing is completed. Typically, this can be done with ALT + arrow keys (see this KB for more details).
  - Please note that login at the CLI is not available until after the full licensing/provisioning process has been completed.
  - Once you see this particular message below, the UI should be available for licensing:

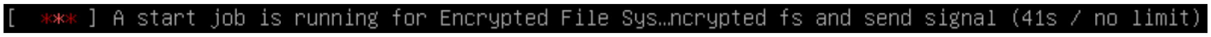

```
[  *** ] A start job is running for Encrypted File Sys…ncrypted fs and send signal (41s / no limit)
```

  - Navigate to https://<your_brain_IP_or_hostname> and you'll see a screen like this:

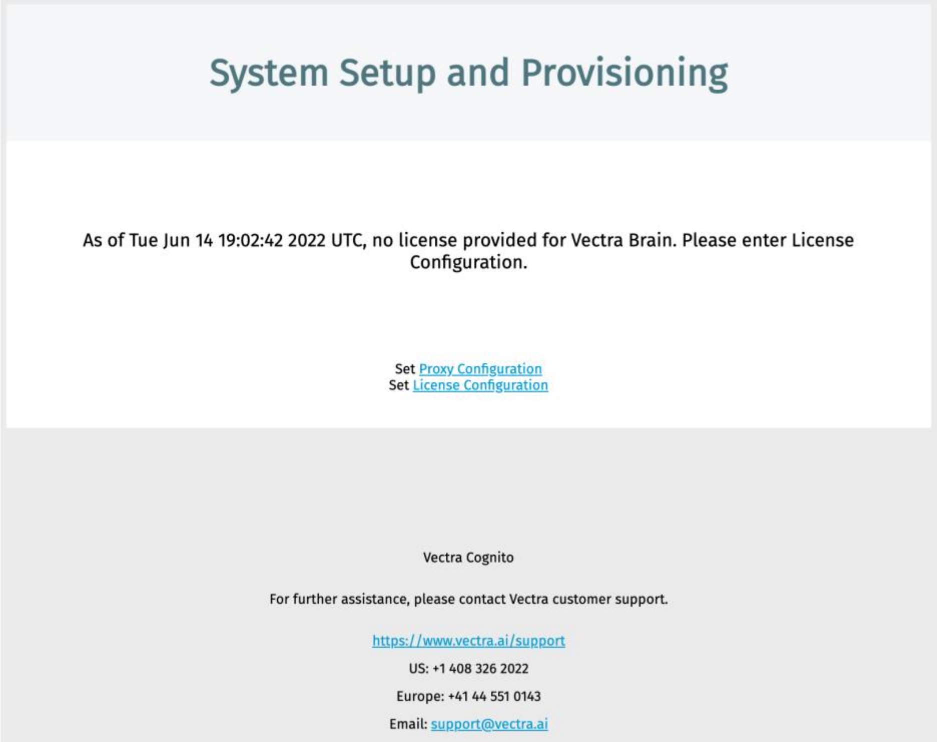

System Setup and Provisioning

As of Tue Jun 14 19:02:42 2022 UTC, no license provided for Vectra Brain. Please enter License Configuration.

Set Proxy Configuration
Set License Configuration

Vectra Cognito

For further assistance, please contact Vectra customer support.

https://www.vectra.ai/support
US: +1 408 326 2022
Europe: +41 44 551 0143
Email: support@vectra.ai

- At this point the initial boot process is paused until a valid license is entered. This process will not time out.
- If a proxy is required to communicate with Vectra for provisioning from your network, please enter the "Proxy Configuration" screen and enter your proxy information:
  - If this is not done now, it can be done after a license is saved but the provisioning process will time out, and time will be wasted until a required proxy is configured.
  - **Please note:** This proxy configuration screen is only used to communicate with Vectra's provisioning server and must utilize an HTTPS proxy. HTTP only proxies are not supported for this use. Other proxy configuration in the main Vectra UI (***Data Sources > Network > Brain Setup > Proxy & Status***) after deployment accepts HTTP proxies and is used by non-provisioning related items.

- ○ **Please note:** If you are doing a Respond UX deployment and require a proxy for non-provisioning related services and integrations (this includes linking to Vectra's cloud for use with the Respond UX), you should configure that proxy at the CLI of your Brain AFTER you progress through this initial configuration and get to the "Success!" message at the end of this section. Please see the Respond UX Deployment Guide in the *Deployment* > *Proxy Support* section for more detail.

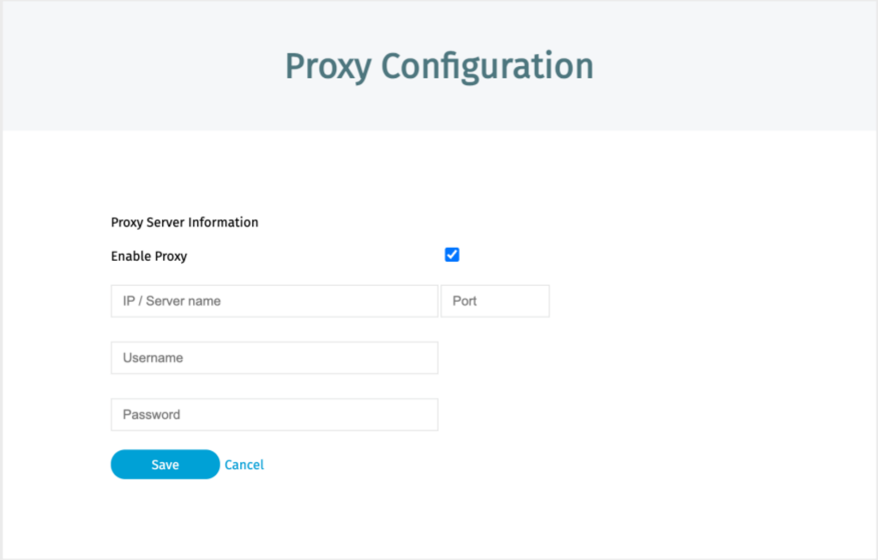

Proxy Configuration

Proxy Server Information

Enable Proxy ☑

- Click into the "License Configuration" screen after saving a proxy configuration (if required):

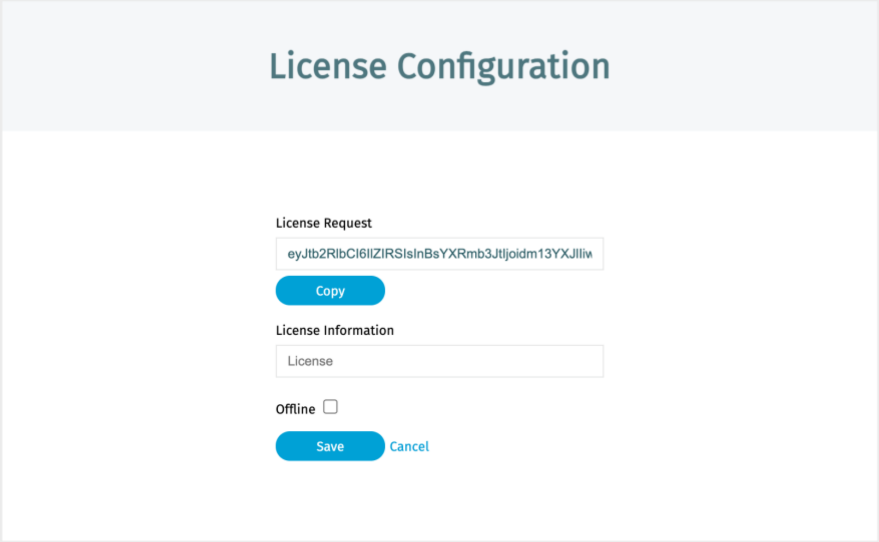

License Configuration

License Request

eyJtb2RlbCI6IlZIRSIsInBsYXRmb3JtIjoidm13YXJlIiw

Offline □

- Copy the "License Request" information by clicking the "Copy" button.
- In another browser tab or window, navigate to https://support.vectra.ai/vectra/additional-resources and then the "License" tab.
  - ○ If you are not already authenticated, you will be redirected to authenticate to your Vectra support account.
  - ○ If you do not have a Vectra support account, you can self-register at the login screen, but licensing will not be available until your account is validated as being a Vectra customer or prospect involved in a trial.
- Paste the license information you copied into the "Enter Authorization Code" section of the page and click "Generate License Key". You should get a "Success" message at the top and a key in the "License Key" box. Copy the license key using the "Copy License Key" button and go back to your Brain in your other tab or window.

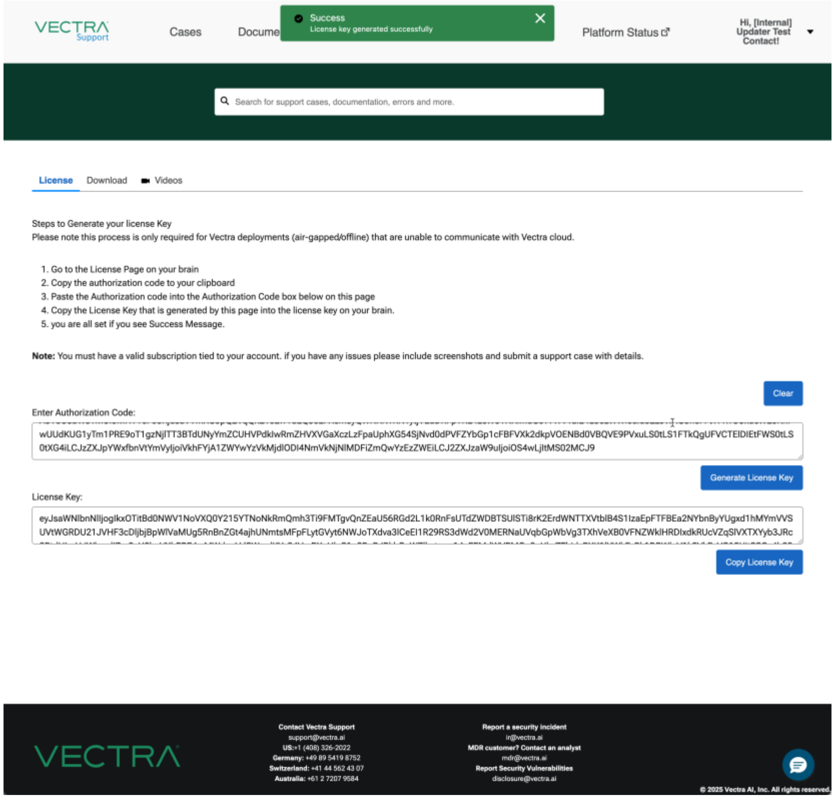

VECTRA Support
Cases
Docume
Platform Status
Hi, [Internal] Updater Test Contact!

Success
License key generated successfully

Search for support cases, documentation, errors and more.
License
Download
Videos

Steps to Generate your license Key
Please note this process is only required for Vectra deployments (air-gapped/offline) that are unable to communicate with Vectra cloud.

1. Go to the License Page on your brain
2. Copy the authorization code to your clipboard
3. Paste the Authorization code into the Authorization Code box below on this page
4. Copy the License Key that is generated by this page into the license key on your brain.
5. you are all set if you see Success Message.

Note: You must have a valid subscription tied to your account. if you have any issues please include screenshots and submit a support case with details.

Enter Authorization Code:

wUUdKUG1yTm1PRE9oT1gzNjlTT3BTdUNyYmZCUHVPdklwRmZHVXVGaXczLzFpaUphXG54SjNvd0dPVFZYbGp1cFBFVXk2dkpVOENBd0VBQVE9PVxuLS0tLS1FTkQgUFVCTElDIEtFWS0tLS
0tXG4iLCJzZXJpYWxfbnVtYmVyIjoiVkhFYjA1ZWYwYzVkMjdlODI4NmVkNjNlMDFiZmQwYzEzZWEiLCJ2ZXJzaW9uIjoiOS4wLjltMS02MCJ9

License Key:

eyJsaWNlbnNlIjogIkxOTitBd0NWV1NoVXQ0Y215YTNoNkRmQmh3Ti9FMTgvQnZEaU56RGd2L1k0RnFsUTdZWDBTSUlSTi8rK2ErdWNTTXVtblB4S1IzaEpFTFBEa2NYbnByYUgxd1hMYmVVS
UVtWGRDU21JVHF3cDljbjBpWlVaMUg5RnBnZGt4ajhUNmtsMFpFLytGVyt6NWJoTXdva3lCeEI1R29RS3dWd2V0MERNaUVqbGpWbVg3TXhVeXB0VFNZWklHRDlxdkRUcVZqSlVXTXYyb3JRc

Contact Vectra Support
support@vectra.ai
US:+1 (408) 326-2022
Germany: +49 89 5419 8752
Switzerland: +41 44 562 43 07
Australia: +61 2 7207 9584

Report a security incident
ir@vectra.ai
MDR customer? Contact an analyst
mdr@vectra.ai
Report Security Vulnerabilities
disclosure@vectra.ai

© 2025 Vectra AI, Inc. All rights reserved.

- Paste the license key into the “License Information” box.
  - **DO NOT CLICK SAVE YET!!!** You must first determine if your Brain will be “Online” or “Offline”.

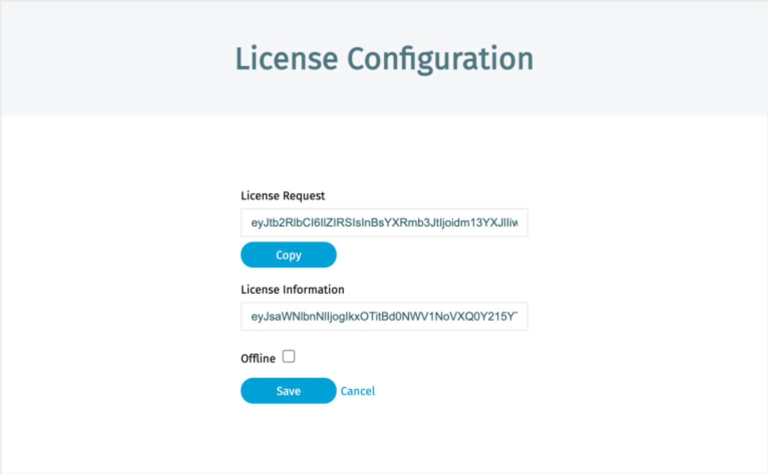

License Configuration

License Request

eyJtb2RlbCI6IlZIRSIsInBsYXRmb3JtIjoidm13YXJlIiw

License Information

eyJsaWNlbnNlIjogIkxOTitBd0NWV1NoVXQ0Y215Y

Offline □

- **If you Brain will be “Offline” (only supported for Quadrant UX deployments) you MUST click the “Offline” checkbox now before clicking “Save” or your deployment will fail**, and you will need to redeploy a new Brain VM and start the process over.
  - Offline Brains do not communicate to Vectra’s provisioning service.
  - Offline Brains are typically used in “air gapped” environments where customers do not have internet access and the ability to communicate with Vectra directly.
  - For additional information regarding offline Brains see the following Vectra Support articles:
    - Offline Updates (v8.9+)
    - Vectra Respond UX Deployment Guide or Vectra Quadrant UX Deployment Guide
  - After determining if your Brain will be online (normal for the vast majority of customers) or offline, click “Save”. You may need to refresh the resulting page that may still say there is no license in place. After the page refresh, you should see screens as shown on the next page.

- After the license is validated, the file system will be decrypted, a performance test will be run, the Brain will reboot, and the Brain will reach out to the Vectra provisioning server and complete provisioning. Finally, a success message will be presented with a button to redirect to the main UI login screen. Offline Brains follow a similar process but do not need to communicate with the provisioning server and validate the license locally.
  - Status messages will be updated but you must manually refresh the browser to see new messages.
  - Below are some examples:

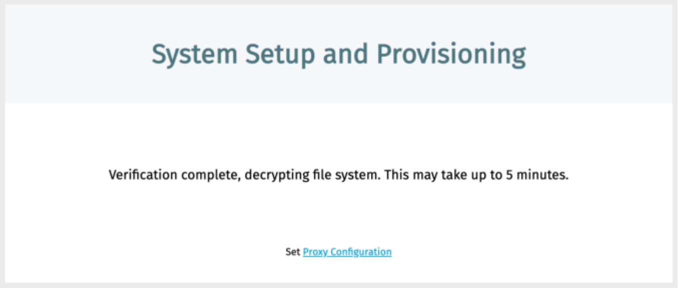

System Setup and Provisioning

Verification complete, decrypting file system. This may take up to 5 minutes.

Set Proxy Configuration

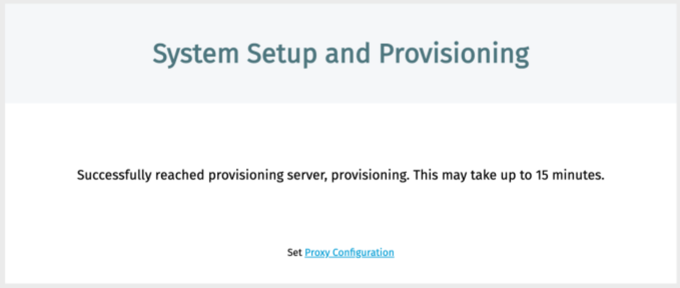

System Setup and Provisioning

Successfully reached provisioning server, provisioning. This may take up to 15 minutes.

Set Proxy Configuration

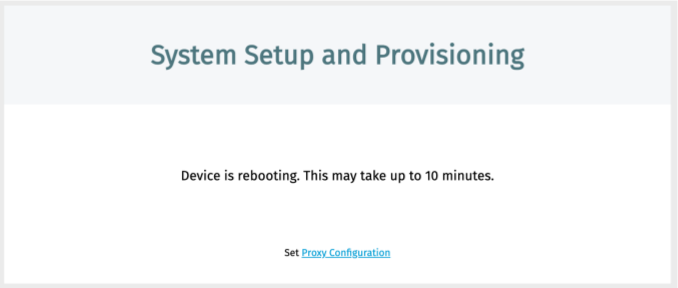

System Setup and Provisioning

Device is rebooting. This may take up to 10 minutes.

Set Proxy Configuration

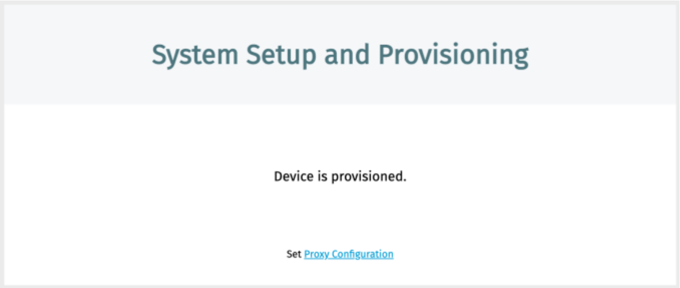

System Setup and Provisioning

Device is provisioned.

Set Proxy Configuration

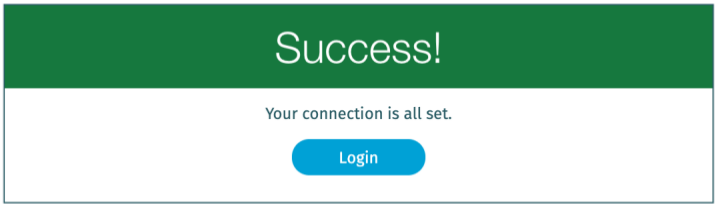

Success!

Your connection is all set.

- Once you see the success message and blue “Login” button you are ready to login to the main Vectra UI if performing a Quadrant UX based deployment.
  - Click the “Login” button and enter the default credentials.
  - Initial login credentials for the UI are: **`admin / changethispassword`**.
  - You will be asked to change the default password upon initial login.
- At this point, SSH to the CLI is available: Initial login credentials are: **`vectra / changethispassword`**
  - You will be asked to change the default password upon initial login.
- For Respond UX deployments, please follow the deployment process as detailed in the Vectra Respond UX Deployment Guide. You should NOT login to the Quadrant UX that is served from the Brain at this time!

Post Deployment Guidance

Setting a Static IP and DNS After Initial DHCP Deployment

If you used DHCP for initial deployment but would like to configure a static IP for production use, you will need to login to the CLI of the Brain to set a static interface assignment. DNS for Brain VMs can be configured at the CLI or in the UI.

Logging in can be done via your hypervisor console function or using SSH to the management port if it was preconfigured with DHCP.

- Connect to your Brain CLI using your hypervisor console or "**`ssh vectra@<IP or Hostname>`**" if you use DHCP and already know the address or hostname.
- Once logged in to the Brain you can view command syntax for the "set interface" command:

```
set interface -h
Usage: set interface [OPTIONS] [mgt1] [dhcp|static] [IP] [SUBNET_MASK]
[GATEWAY_ADDRESS]

Sets mgt1 to either dhcp or static ip configuration

Options:
-h, --help Show this message and exit.
```

- Setting the IP address statically:
  - IPv6 is supported for the MGT1 interface. For full details, including information regarding dual stack support, please IPv6 Management Support for Vectra Appliances on the Vectra support portal. Below we will show how to enable IPv6 support (it's off by default) and the syntax to use when setting an IPv4 or IPv6 address.
  - To enable/disable IPv6 support

```
# show ipv6 enabled
IPv6 is disabled

# set ipv6 enabled
Response: ok

# show ipv6 enabled
IPv6 is enabled

# set ipv6 disabled
Response: ok
```

  - Setting IPv4 and IPv6 syntax examples:

    Execute the following command to set the MGT1 interface to the desired static IP address:

```
IPv4 Syntax:
set interface mgt1 static x.x.x.x y.y.y.y z.z.z.z

Where:
x.x.x.x is the desired interface IP address
y.y.y.y is the desired interface network mask
z.z.z.z is the desired gateway

IPv6 Syntax:
set interface mgt1 static [IPv6 IP] [Subnet Mask] [Gateway]
```

```
Example:
set interface mgt1 static 2001:0db8:0:f101::25 64 2001:0db8:0:f101::1
```

- To change back to DHCP (default):

```
set interface mgt1 dhcp
```

- Configure DNS for the appliance:

Command syntax to set DNS (up to 3 nameservers are supported):

```
set dns [nameserver1 <ip>] [nameserver2 <ip>] [nameserver3 <ip>]
```

Example:

```
set dns 10.50.10.101 10.50.10.102
```

Verifying DNS Configuration:

```
show dns
```

To set DNS in the UI, navigate to *Data Sources* > *Network* > *Brain Setup* > *DNS Entries* and edit the settings.

Setting static IP and DNS at the CLI example:

```
vscli > set interface mgt1 static 172.16.12.11 255.255.255.0 172.16.12.1
Interfaces updated successfully
vscli > set dns 10.50.10.101
DNS Set: success
vscli > show interface
mgt1:
    Running:
        Gateway: 172.16.12.1,
        Ip: 172.16.12.11,
        Link Speed: 10Gbps,
        Link State: up,
        Mac: 00:0c:29:89:ad:a6,
        Mode: static,
        Netmask: 255.255.255.0
vscli > show dns
Id|Server      |Description
1  10.50.10.101 Configured DNS nameserver
```

Performance Testing

As discussed earlier in this guide, a performance test is run during the initial boot process. This is to test the performance of the Brain against baselines that Vectra has established for the different configuration options.

Cached results from the initial performance test run can be retrieved from the command line while logged in as the **"vectra"** user. Additional performance tests can be run by using the **--force** switch on the performance test command. Please note the following:

- **Running the performance test is an intensive operation which stops most services on the Brain**.
- Additional performance tests should only be run when your security team knows the Brain will be unavailable.
  - Paired Sensors will buffer metadata that can't be sent to the Brain so there should ultimately be no detection gap, although this could introduce a delay in detection publishing while the test is run.
- Baselines are set by Vectra for each of the various configurations of Brain.
  - Warning is for 10% below expectations. Critical is for 20% or more below expectations.
    - 260 MB/s is the minimum required throughput for all disks (OS and Data) and is represented by a score of 10.0 on the performance test in the "disk" category.
  - Critical is considered a failure and performance is not expected to be satisfactory. Vectra engineering considers systems which fail the performance test to be invalid configurations and customers should use more performant base hardware to ensure supportability, reliability, and performant operation.

Example:

```
vscli > performance-test --help
Usage: performance-test [OPTIONS]

Run a system performance test

Options:
--force Run all tests regardless of cached results.
-h, --help Show this message and exit.

vscli > performance-test
This may take up to five minutes. Most system services will be down for the duration of the test.
Test |Score |Result |Time
cpu 10.00 / 10.0 pass 30.04
cpu_steal 10.00 / 10.0 pass 0.06
disk 10.00 / 10.0 pass 47.94
memory 10.00 / 10.0 pass 0.00
memory_balloon 10.00 / 10.0 pass 0.05
overall 10.00 / 10.0 pass 78.09
```

Integrity Checks

Vectra performs file system integrity checks to make sure that core libraries have not been altered. If the system detects issues during boot, a system setup and provisioning dialog will appear that is similar to the licensing screen.

- Click "Set File System Configuration".
- Copy the "Error Code" and send it to Vectra support for decryption.
- Vectra has tooling to determine what has been changed, and if warranted can provide a whitelist code to the customer to allow the system to continue booting.
- Whitelist codes work one time. If the system again fails a file system integrity check, a new whitelist code will be required. Please work with Vectra support to ensure compliance.

Below are some example screenshots:

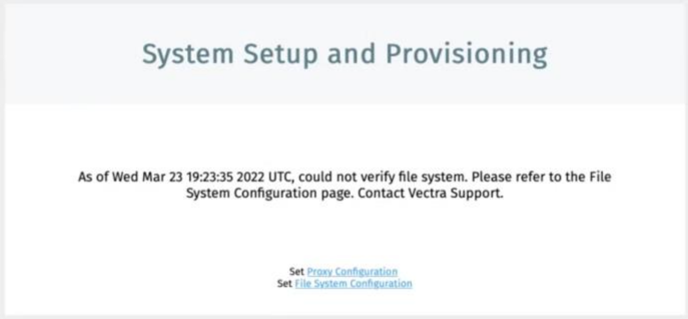

System Setup and Provisioning

As of Wed Mar 23 19:23:35 2022 UTC, could not verify file system. Please refer to the File System Configuration page. Contact Vectra Support.

Set Proxy Configuration
Set File System Configuration

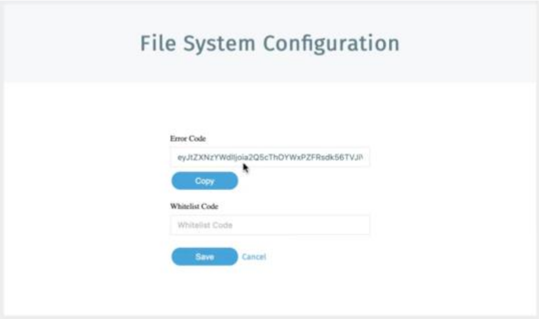

File System Configuration

Error Code: eyJtZXNzYWdlIjoia2Q5cThOYWxPZFRsdk56TVJI

Whitelist Code: Whitelist Code

Configuration Validation

During boot, the Brain determines which configuration it is running and sets some parameters differently depending on resource availability per configuration. This is an automatic process and requires no user input. The "**show system-health**" command can be run at the command line as the "**vectra**" user to see that your configuration is a supported option. Look for the "**[OK] VM Specifications**". The specific checks shown may not match your system. Vectra occasionally updates the specific checks used in the system-health command.

Example:

```
vscli > show system-health

======== Ran 8 check(s). 8 Passed, 0 Failed, 0 No Result ========

vscli > show system-health --verbose
[ OK ] Available Virtual Storage Space
[ OK ] Disk Writable
[ OK ] NIC Detection
[ OK ] Vectra User Password
[ OK ] Sensor Connectivity
[ OK ] Sensor Link Utilization
[ OK ] Sensor Tunnel
[ OK ] VM Specifications
```

License Checks and Renewal

Once a Brain is up and running, it will periodically check its license status. This will occur whether the Brain is online or offline (from the perspective of connection to Vectra). Once a Brain is 30 days from expiration, it will begin to send syslog messages with a count down until expiration. Once the license expires a new syslog message is sent (Quadrant UX). Respond UX deployments will write the message to the audit log which is available for query via API. Here are examples:

```
"License Checker: Detect License Expires in {days_until_expiration} days."
"License Checker: Detected Invalid/Expired License, disabling services"
```

The status of your license can be seen in the following locations in the Vectra UI:

- *Manage > Licensing*
- *Discover > System Health > Deployments*
    - If you license status does not show, click the "Update Now" button at the bottom of the section.

Examples:

Manage > Licensing:

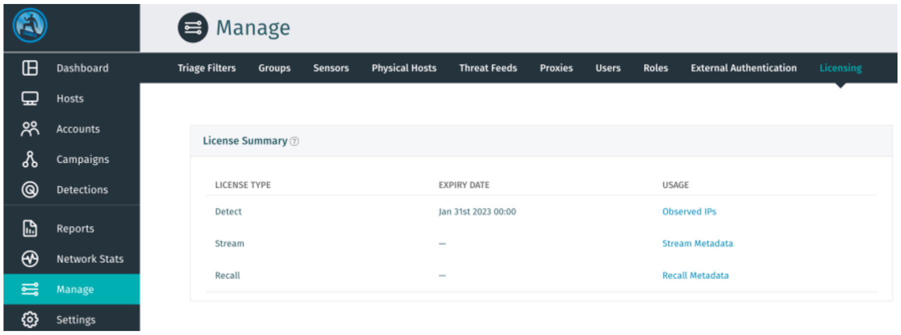

License Summary

LICENSE TYPE	EXPIRY DATE	USAGE
Detect	Jan 31st 2023 00:00	Observed IPs
Stream	–	Stream Metadata
Recall	–	Recall Metadata

Discover > System Health > Deployments:

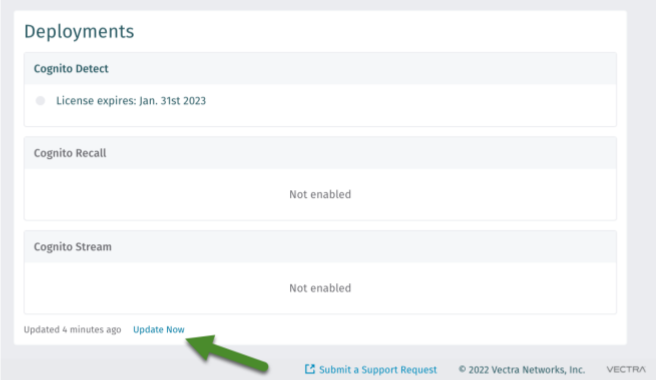
Deployments

Cognito Detect

License expires: Jan. 31st 2023

Cognito Recall

Not enabled

Cognito Stream

Not enabled

Updated 4 minutes ago Update Now

For Brains that are connected to Vectra, license renewal is an automated process that requires no user intervention. When you sales contract is renewed and the expiration date is updated, Vectra's provisioning service will provide a new license key to your Brain.

If your Brain is offline (not connected to Vectra or air-gapped), to renew your Vectra license simply browse to the *Manage > Licensing* screen, copy the authorization code, provide it to Vectra (support, sales team, etc), and Vectra will provide you a new license key for entry into the UI once your entitlement is verified:

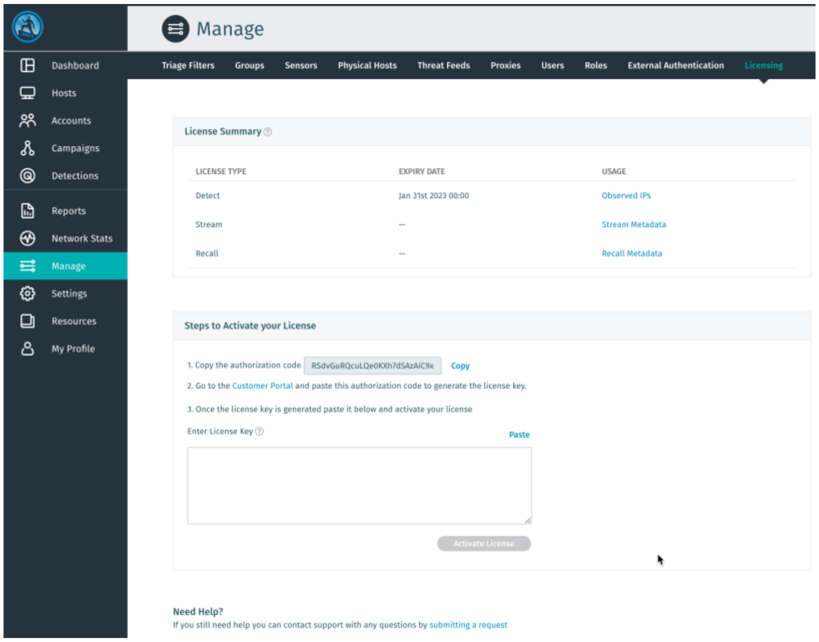

License Summary

LICENSE TYPE	EXPIRY DATE	USAGE
Detect	Jan 31st 2023 00:00	Observed IPs
Stream	–	Stream Metadata
Recall	–	Recall Metadata

Steps to Activate your License

1. Copy the authorization code RSdvGuRQcuLQe0KXh7dSAzAiC9x Copy
2. Go to the Customer Portal and paste this authorization code to generate the license key.
3. Once the license key is generated paste it below and activate your license

Enter License Key

Need Help?

If you still need help you can contact support with any questions by submitting a request

Resizing the Brain

In some environments, you may wish to start with a smaller Brain instance and then later move to a larger Brain instance to handle additional load (metadata coming from paired sensors or additional paired sensors).

- Please see: Resizing Virtual Sensors and Brains for details.

Next Steps

At this point your Nutanix Brain is fully deployed, and you can move on to other tasks associated with your overall deployment.

It is recommended to follow the Vectra Respond UX Deployment Guide or Vectra Quadrant UX Deployment Guide for additional information regarding initial settings for your deployment. You may wish to deploy and pair network Sensors or configure other Vectra offerings such as Recall, Stream, CDR for M365, IDR for Azure AD, CDR for AWS, CDR for Azure, etc. Additional documentation can be found in the Vectra Product Documentation Index on the Vectra Support site.

Worldwide Support Contact Information

- Support portal: https://support.vectra.ai/
- Email: support@vectra.ai (preferred contact method)
- Additional information: https://www.vectra.ai/support