VECTRA

NDR (Detect) and Network Identity Architecture Overview

Definitions / Deployment

Definitions

- **Vectra UI / Deployment Types**
  - Please see Vectra Analyst User Experiences (Respond vs Quadrant) for additional details.
  - **Respond UX (RUX)** - The Vectra UI is served from Vectra's cloud. Shown in diagram on next page.
    - The UI serves as the central point of web-based management for your RUX deployment(s).
    - When used with network Sensors, it communicates with Brain appliance(s) deployed in customer premises.
  - **Quadrant UX (QUX)** - The Vectra UI is served from Brain appliance(s). Not shown in diagram on next page.
    - The UI serves as the central point of web-based management for your QUX deployment(s).
- **Vectra Cloud** - The portions of the Vectra AI Platform that reside in Vectra's cloud.
- **Customer Premises** - Private/Shared Data Centers, Public Cloud, Campus/Office environments where Brain or Sensor appliances will be installed to capture network traffic (including network identity).
- **Appliance** - Physical or virtual (including public IaaS cloud) Brain, Sensor, or Stream appliance.
- **Vectra AI Platform** - Vectra's cloud platform, delivered as SaaS with the RUX UI.
  - Deployments with network Sensors will include at least one Brain appliance and one or more Sensor(s). For smaller deployments, this can be a single mixed-mode appliance.
- **Brain Appliance** - Can be a physical appliance or virtual appliance.
  - Pairs with Sensors (network data sources) and processes / deduplicates and optionally forwards the metadata received from Sensors (when licensed for Stream or Recall (Recall is for QUX deployments only).
  - Serves as communications broker between Vectra's cloud and local integration points for RUX deployments.
- **Sensor Appliance** - vSensor is a virtual Sensor (for hypervisors or IaaS cloud), Sensor is a physical Sensor.
  - Must be paired to a Brain.
  - Captures and deduplicates raw network traffic.
  - Forwards metadata to the Brain for processing.
  - Houses rolling capture buffer to enable PCAP retrieval when requested from the Brain.
  - Optionaly runs Vectra Match and Suspect Protocol Anomaly Detections.
- **Mixed-mode Appliance** - Can perform both Brain and Sensor functions when used in smaller deployments.
- **Network Identity** - Vectra covers IDR (Identity Detection and Response) use cases through analysis of identity as observed in the network through protocols such as Kerberos, DCE/RPC, LDAP, NTLM, etc.

Brain and Sensor Deployment

- Start Here:
  - RUX Deployments - Vectra Respond UX Deployment Guide
  - QUX Deployments - Vectra Quadrant UX Deployment Guide
  - Additional documentation is available on the Vectra Support Portal
  - Formal product docs are tracked in the Vectra Product Documentation Index (KB article in the above portal)
- The general process is:
  - RUX - Receive your welcome email for a RUX deployment, login and configure user accounts / SSO, deploy your Brain appliance, perform configuration / integration, deploy Sensors, direct traffic to the Sensors.
  - QUX - This deployment is largely the same except there is no RUX UI, and your UI is served from the Brain.
- Vectra offers a full complement of services including technical support, professional services / training, managed services, and offensive security services. Please see the Vectra AI Services datasheet and speak with your Vectra account team for details and pricing.
- Post deployment - Enable backups, integrations, enable other features/products, configure notifications/reporting, built groups and triage rules to suppress unwanted detections for authorized behaviors. AI-Triage will also learn the environment and automatically suppress detections as part of AI-driven Prioritization.

General Deployment Description

- The diagram to the right represents a high-level Respond UX deployment of Vectra NDR (including network identity).
- A Brain appliance is installed in the customer premises and the RUX UI is served from the Vectra cloud.
- Sensors are paired to the Brain and capture network traffic in the data center, campus, and IaaS public clouds.
  - The Sensors can also run Vectra Match and Suspect Protocol Anomaly Detections.

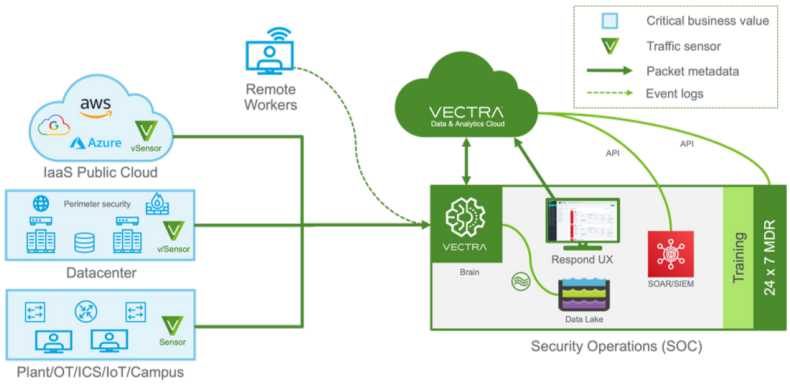

- The Brain appliance processes the metadata locally to create detections.
  - The Brain sends metadata and detections to the Vectra cloud for presentation in the Respond UX.
  - For QUX deployments, the UI is local, and metadata can be sent to Vectra Recall.
  - Vectra Stream sends network metadata to the customer's data lake for both RUX and QUX deployments.
- Logs are analyzed by the Brain to link remote workers with traffic seen by app connectors in the data center.
- SOAR/SIEM integration uses API connections to communicate with the Vectra cloud.

Placement of Brain and Sensors

- It is generally recommended to deploy Brain and Sensor devices in locations not visible from the public internet.
  - Private connectivity or a VPN tunnel that is terminated outside of the Vectra appliances is preferred.
- Small remote branch locations do NOT all typically need a Sensor. Flows are still seen in central locations.
- Vectra NDR detections are focused on In-to-In and In-to-Out communications. It is preferred to NOT capture outside of your edge firewalls in the DMZ where a significant amount of traffic will be Out-to-In. Out-to-Out traffic should NOT be captured. Vectra Match or some metadata use cases may necessitate Out-to-In traffic capture.
  - Vectra allows you to define your internal IP space. Exceptions can be configured for internal labs where you may wish to simulate CnC traffic as an example. Networks can be fully ignored by VLAN or CIDR block.

Proxy Guidance

See Proxy Handling in Vectra for full details

- Traffic should be captured on the south side of any proxy or NAT device to fully recognize the real sending host.
- In some situations, you may need to capture traffic on the north side of proxies.
  - Flows from multiple hosts behind the proxy will be mapped to the north side proxy IP and would result in spurious detections. The Vectra system automatically suppresses detections that would fire inaccurately because of this. The *Manage > Proxies* in your Vectra UI page allows you to manage identified proxies.

Network Traffic Capture Guidance

General Traffic Types to Capture

Traffic Type	Purpose	Examples
North/South	C&C, Exfiltration, Botnet	Server to Internet, User to Internet
East/West	Recon, Lateral Movement	Server to Server, User to Server, User to User

Important to Capture

DCE/RPC
DHCP
DNS
HTTP
ICMP
Kerberos
LDAP
NTLM
Radius
RDP
SMB
SMTP
SSH
SSL/TLS
X509
Other session traffic

Should be Excluded from Capture

Core routing protocols
High-bandwidth backup data
High-performance computing (HPC) workloads high in bandwidth
HPC workloads that are well isolated
Multiprotocol Label Switching (MPLS)
Session Initiation Protocol (SIP)
Storage network file systems (SMB is ok)
Video Multicast

Supported Encapsulations

GENEVE
Generic Routing Encapsulation (GRE)
IEEE_802.1ad (known as QinQ)
IEEE_802.1Q (VLAN)
IPSec Authentication Header (IPSec AH)
Virtual Extensible LAN (VXLAN)

What Helps Improve Vectra's HostID

See Understanding Vectra Detect Host Naming for additional information.
DNS, Reverse DNS, Multicast DNS (mDNS)
Kerberos
DHCP
Netbios
EDR Integration, VWware integration, SIEM event forwarding, Windows Event Log Ingestion

Typical Traffic Capture Sources

Vectra NDR for Cloud (Gigamon)
SPAN/COPY/MIRROR Ports
Traditional network TAP devices
Packet brokers
Native Cloud mirroring options such as VPC Traffic Mirroring (AWS), GCP Packer Mirroring, VTAP (Azure)

Encrypted Traffic - Vectra AI models work with encrypted traffic (no decryption required)

- The vast majority of Vectra AI network traffic related detections do not require decryption. For more information, please see our whitepaper "The AI Behind Vectra Attack Signal Intelligence". The relevant section begins on page 12. Some customers may still have a need for decrypted data such as:
  - Some users of Vectra Recall, Vectra Stream, Instant Investigation, or Advanced Investigation may require the decrypted traffic metadata. These use cases can vary but are not very common and are supported.
  - If using Vectra Match, decryption can be useful and allow certain rules to fire that would otherwise not fire if the system processed only encrypted data. Again, this requirement will vary by customer implementation.

Guidelines when decrypting traffic to be sent for analysis to Vectra Sensors

- When you want both encrypted and decrypted traffic to be analyzed, this should be done in parallel pipelines.
  - A second Sensor (or as many as are needed for your specific throughput or architecture requirements) should be used that is paired to a different Brain appliance.
  - Vectra supports virtual Sensors, Brains, and Stream appliances and does not charge for their use outside of normal licensing metrics for your environment.
- If you do not use a parallel processing pipeline and send both decrypted and encrypted traffic to the same Sensors and paired Brain, Vectra's deduplication will see the same 5 tuple traffic in multiple data streams and will discard the duplicate data that arrives last. This would typically be the decrypted data as the decryption process adds overhead on the sending side.

Basic Communications/Pairing

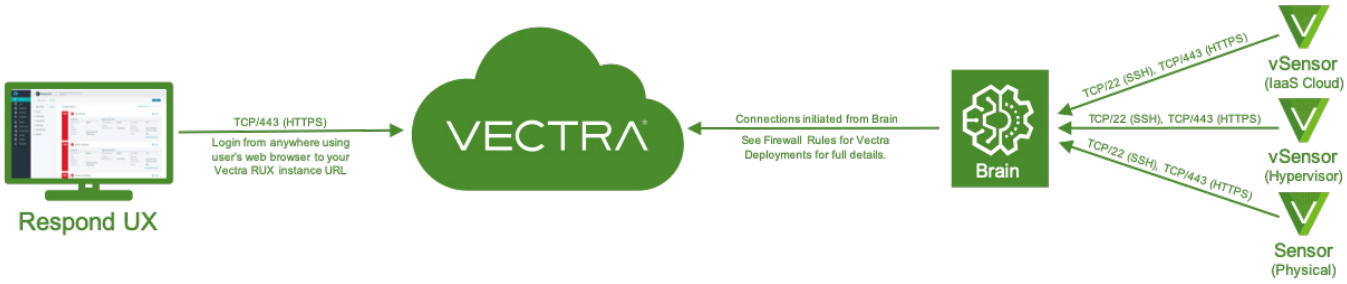

Note! - Please see Firewall Requirements for Vectra Deployments for full details. There are additional requirements depending on the specifics of your deployment, specific integrations, etc.

Note! - Air gapped deployments of Vectra NDR are possible (QUX only). NDR Detections are generated locally on the Brain. Please see Offline Updates (v8.9+) for details on how updates are easily managed in these situations.

- Users login using their web browser to the Respond UX at the URL given in their welcome letter.
  - Quadrant UX users login to their Vectra UI (served from the Brain) at the IP or hostname of their Brain.
- Sensors must be able to reach the Brain over TCP/22 (SSH), and TCP/443 (HTTPS).
- In general, Sensors only communicate with the Brain appliance and DNS. There is only one exception:
  - Physical Sensors can retrieve the IP of the Brain they need to pair with from the Vectra Cloud at initial boot. With DHCP enabled, a physical Sensor can automatically begin the pairing process with no login required to its CLI. This requires that TCP/443 be open to update2.vectranetworks.com from the physical Sensor.
- The Brain appliance is deployed in the customer premises and communicates with the Vectra Cloud over a variety of channels. These connections are initiated from the Brain.
  - Updates are served from the Vectra Cloud and installed automatically in an update window chosen by the customer. Sensor updates are served from the Brain appliance and are also installed automatically.
- Pairing Sensors or Stream appliances with a Brain is a simple process. Some details are below:
  - Physical Appliance Pairing Guide
  - vSensor pairing is covered in each vSensor deployment guide (see Product Documentation Index for guides)
  - A Sensor registration token, managed in your UI or CLI, allows IaaS cloud vSensors or any other Sensor to pair with a Brain even if it was not originally configured for use with a specific Brain.

Appliance Sizing

Note! - Vectra periodically releases new appliances to support new throughput requirements, hypervisors, or cloud providers. Always refer to Appliance and Sensor Specifications and your Vectra account team for the latest guidance.

Note! - Global View enables maximum scale for the largest global enterprises. The performance guidelines below are for individual appliances. Please see the Global View detail on the next page for additional detail.

Note! - Performance refers to the amount of network traffic observed by Sensors that a Sensor can produce metadata for, or the amount of traffic observed by Sensors that a Brain can process metadata for. The performance numbers are based upon average throughput a given Sensor/Brain can process. Actual performance may vary depending on traffic composition.

- Brain appliances are available in configs that support up to 75 Gbps of performance and 500 paired Sensors.
- Sensor appliances are available in configs that support up to 50 Gbps of performance.
- When running Match and/or Suspect Protocol Anomaly Detections in addition to Vectra NDR (Detect), Sensor appliances are available in configuration that support up to 33 Gbps of performance.
- Please work with your Vectra account team to determine the right mix of physical and virtual appliances required.

Global View

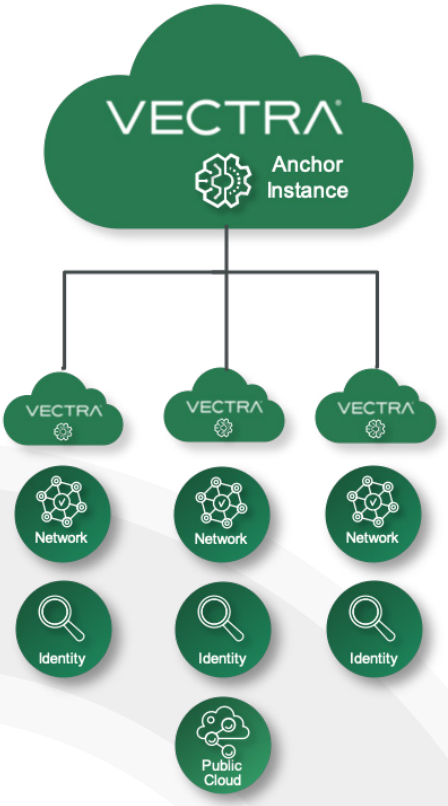

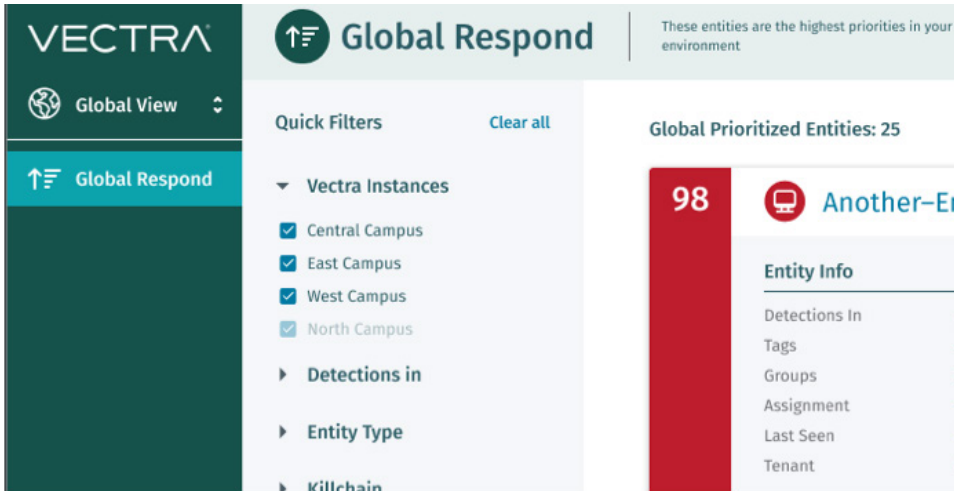

- Global View provides a Global Respond view of prioritized entities from child RUX instances in an anchor instance of RUX.
  - Available as a standard feature with any RUX deployment.
- Allows for overlapping IP space in child deployments.
- Enables maximum scale while still providing a single pane of glass.
- The anchor instance retrieves prioritized entities from child instances using an API client that is assigned the "Global Analyst" role.
  - Communications are all encrypted and contained within the Vectra AI Platform.
  - Data is not stored in the anchor instance. It is retrieved on demand from child instances.

Additional Guidance and Resources

- Detection data (including associated micro pcaps) are generally retained for 6 months. Metadata retention is dependent on your contracted retention options for Advanced Investigation (RUX) or Recall (QUX). When sending metadata via Stream, retention is based on the configuration of your downstream data lake.
- Vectra Packet Capture allows you to capture PCAPS for subsequent download and analysis with local tools.
- Guidance for optimizing your deployment when used with VPN clients is available here.
  - Zscaler specific information: ZPA Log Ingestion and Configuration, ZIA Integration and Optimization
- Response options
  - EDR Host Lockdown
  - Account Lockdown and Azure AD (Entra ID) Account Lockdown
  - SOAR integration - search for your SOAR vendor on the Vectra Support Portal
  - Vectra Automated Response - Click the link to see the GitHub repo with dozens of integrations
- Vectra is a very open platform with integration options for SIEM, SOAR, etc.
  - API guides are available for Respond UX deployments and Quadrant UX deployments.
  - Syslog/Kafka is NOT supported in RUX deployments but is available for QUX deployments.
  - If you require Syslog for RUX, this can be done through a customer deployed intermediary server with guidance from: SIEM Connector for the Vectra AI Platform (pulls data via API and sends to your SIEM).
- MITRE ATT&CK and D3FEND details are available here: Vectra's Coverage of MITRE ATT&CK and D3FEND

About Vectra AI

Vectra AI, Inc. is the leader in AI-driven extended detection and response (XDR). The Vectra AI Platform delivers integrated signal across public cloud, SaaS, identity, and data center networks in a single platform. Vectra AI's patented Attack Signal Intelligence empowers security teams to rapidly detect, prioritize, investigate and stop the most advanced hybrid cyber-attacks. With 35 patents in AI-driven detection and the most vendor references in MITRE D3FEND, organizations worldwide rely on the Vectra AI Platform and MXDR services to move at the speed and scale of hybrid attackers. For more information, visit www.vectra.ai.

VECTRA

For more information please contact us: Email: info@vectra.ai | vectra.ai