VECTRA

DATA SHEET

Vectra Recall - Security and Privacy Statement

Introduction

Vectra Recall is a cloud service, provided by Vectra, which stores network metadata in support of security investigation and threat hunting. This document gives an overview of how Vectra develops, tests and secures Vectra Recall, including the following:

- An overview of Vectra Recall
- The types of data we store and process
- How we protect your data
- How we securely develop and operate the service

Due to the ever-changing threat landscape and business environment we reserve the right to modify our security and privacy practices as needed. The current version of this document is available at https://support.vectra.ai/s/article/KB-VS-1120. Further questions may be directed to privacy@vectra.ai.

An overview of Vectra Recall

Vectra Recall is part of the Vectra platform, which is shown in Figure 1:

- Sensors monitor the customer's network and extract metadata which describes the observed activity.
- Vectra Detect is an appliance, hosted in the customer's datacenter or virtually in their cloud, which provides AI-powered automated threat detection.
- Vectra Recall is a cloud service, provided and managed by Vectra, which stores historical network metadata to support security investigation.
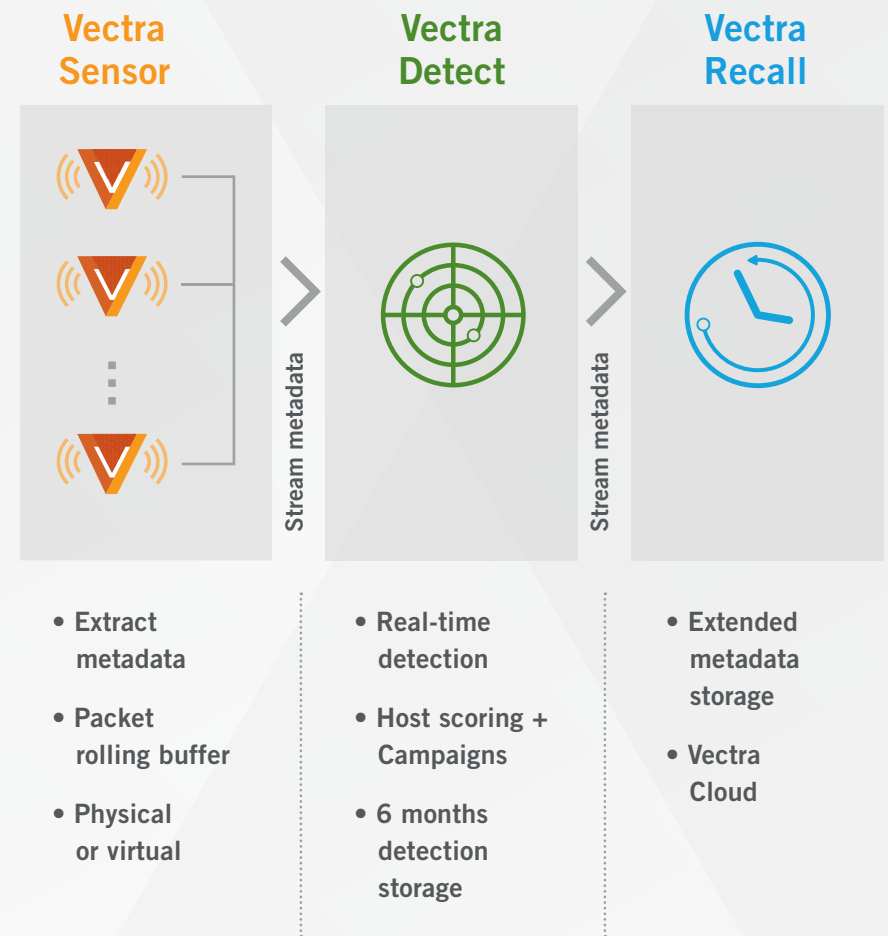

Vectra Sensor

Stream metadata

Vectra Detect

Stream metadata

Vectra Recall

- Extract metadata
- Packet rolling buffer
- Physical or virtual

- Real-time detection
- Host scoring + Campaigns
- 6 months detection storage

- Extended metadata storage
- Vectra Cloud

Figure 1: Components of the Vectra platform

The types of data we store and process

Vectra Recall stores the metadata that is extracted by the network sensors and makes it available to analysts to help them efficiently hunt for threats. Metadata is available for a customer-selected period, starting at 2 weeks. There is no user function to delete the data, it is automatically deleted after the retention period or when the customer cancels their subscription to the service.

When analyzing the network traffic, there are multiple artifacts that can refer to personal information (e.g. IP-address, account name, URIs, etc.). Vectra processes this information to find malicious activity in the network. The following types of metadata are processed:

Metadata type	Information stored by Vectra Recall	Examples	May be considered to contain personal information
Beacons	Client and server hostname and IP	www.vectra.ai 204.12.57.132	Yes
DCE/RPC	Endpoint name, client and server hostname and IP, username	RpcSqlOpenPrinter johndoe	Yes
DHCP	DHCP and DNS server IPs, hostname, MAC address	00:21:44:23:45:67	Yes
DNS	Domain Name System (DNS) queries and responses, e.g. hostname and IP	www.vectra.ai 204.12.57.132	Yes
HTTP	URL, cookie, referrer, user agent	http://www.vectra.ai?user=johndoe	Yes
Kerberos	Client and server hostname and IP, realm, name	johndoe@vectra.ai, krbtgt, admin.vectra.ai	Yes
LDAP	Client and server hostname and IP, query, results (incl usernames)	DC=vectra,DC=ai, CN=John* givenName=John Doe	Yes
Network connections (iSession)	Client and server hostname and IP, vlan	www.vectra.ai 204.12.57.132	Yes
NT Lan Manager (NTLM)	Domain, client and server hostname, username	Admin, MicrosoftAccount Johndoe	Yes
Remote Desktop Protocol	Client version, RDP Cookie, desktop width and height, client and server IP, keyboard layout	RDP 7.1 "United States-Dvorak for right hand"	Yes
SMB Files & Mappings	File operation, path, size	open, \\server\docs\report.doc	Yes
SSL handshakes	Client and server hostname and IP supported cypher suites	TLS_RSA_EXPORT1024_WITH_RC2_CBC_56_MD5	No
X509 Certificates	Information from the certificate public keys including subject, SAN, issuer & expiry	www.vectra.ai servername.vectra.ai	Yes

Table 1: Data stored and processed by Vectra Recall

The following types of personal information are contained within the data:

- **IP Addresses:** The metadata contains the IP address of the source and destination of the flow. This is an IP address from the internal address space of the customer.
- **Username:** Metadata streams such as HTTP, Kerberos, NTLM, RDP, LDAP can contain usernames
- **Device names:** The metadata can contain the device names which may in turn contain the username in them depending on how they are constructed, e.g. "Bob-MBP".
- **HTTP Cookies:** HTTP metadata will contain the cookie value if the HTTP is not encrypted.
- **Keyboard layout:** Remote Desktop communication that is not encrypted may contain the keyboard layout of the originating host.

How we protect your data

Protecting your data is of paramount importance to Vectra. As Vectra Recall is a SaaS (Software as a Service) service, security is the shared responsibility of Amazon Web Services (AWS), Vectra and the customer. AWS secure the cloud platform, Vectra secures the data and services within the cloud and the customer secures the credentials they use to access the service.

SOC2 cloud infrastructure

Vectra Recall is hosted in AWS datacenters. These datacenters are built with layered security controls to protect the physical equipment (such as servers, disks and network equipment) from tampering and to monitor and protect against unauthorized access. They also contain redundant services to protect against natural disasters and mitigate against outage of critical services such as the power-grid. These controls are monitored as part of AWS's SOC2 audit. SOC2 is a set of controls developed by the American Institute of CPAs (AICPA) to govern security, availability and privacy of information systems. For more details, see AWS's compliance program.

Audited SOC2 service

Vectra Recall is a SOC2-audited service built on top of AWS. We have completed a SOC2 type 1 report, have implemented all required controls, and have successfully completed a SOC2 type 2 audit. For more details on our compliance, or to obtain a copy of our SOC2 type 2 audit report, please contact privacy@vectra.ai.

Guaranteed data sovereignty

Vectra Recall is available in multiple AWS datacenters across North America, Europe, the Middle-East and Asia. During provisioning you can select where your data is hosted. Vectra and AWS ensure that data will not leave the selected region even for availability or disaster recovery.

The specific AWS regions used are:

- AWS US West (Oregon)
- AWS Canada (Central)
- AWS US East (Ohio)
- AWS Europe (Ireland)
- AWS Europe (Paris)
- AWS Europe (Frankfurt)
- AWS Middle East (Bahrain)
- AWS Asia Pacific (Sydney)
- AWS Asia Pacific (Tokyo)
- AWS Middle East (UAE)
- AWS Asia Pacific (Singapore)

Per-customer virtual private cloud

Vectra Recall separates each customer's data using comprehensive controls. Each customer's service is logically isolated as a separate virtual private cloud (VPC) within the chosen datacenter. This ensures that each customer can only ever see their own data and that access, either administrative or as result of a compromise, to one customer's data doesn't enable access to the data belonging to other customers. All communication from Vectra Detect/X-series appliance to the Recall cloud is unidirectional using a Push/Pull model; there is no path from the Recall cloud back to the Vectra X-Series appliance in the customer's datacenter or cloud.

Strong encryption and authentication

To protect your data, we ensure that it is encrypted at rest and on the wire. AES-256 encryption is used (via AWS's Key Management Service) to protect data in storage and TLS encryption is used on each data connection. The connection from Detect to Recall is mutually authenticated and authorized with per-customer security. This ensures that each Recall instance will only accept connections from the corresponding Detect instance and Detect instances will not connect to anything other than the intended Recall instance.

How we securely develop and operate the service

Security is at the heart of Vectra's culture. Vectra Recall is developed using best-in-class software development and operations methodologies (including DevOps).

Using a practice known as Infrastructure-as-Code manual operations are replaced with automation by treating infrastructure components (compute instances, storage buckets etc.) as code. All changes are designed, implemented, reviewed and change-controlled with the same rigor as code. This includes two-factor authentication, an audit trail of who made the change, when the change was made, who reviewed the change etc.

Our engineering teams employ a very high level of rigor during development:

- All designs are reviewed by security and domain experts
- Paired programming is used for critical parts of the system
- Each change must pass automated test and be peer reviewed prior to being committed
- Changes are automatically deployed to ensure consistent system state.

State-of-the-art cloud management tools are used to deploy changes predictably and to ensure that the desired state is applied uniformly across the platform. All changes and login attempts are logged and retained for at least 12 months. Service components are actively monitored to maintain good performance and health.

Every Vectra employee is bound by confidentiality agreements which protect customer data. The use of raw customer data is strictly limited to those involved in customer support incidents. Each employee is bound by a code of conduct that describes acceptable behavior and use of information systems.

Retention

Retention policy on Vectra Recall is defined by the customer in increments of 2 weeks. Once the data is outside of the retention window, it is deleted.

Access to Vectra Recall

Access by Customers

Customers can access Vectra Recall through the user interface. Access to the user interface requires authentication with strong password over HTTPS.

Access by Vectra

Access to information in Vectra Recall is restricted to Vectra devops, security research and data science, and customer service teams. Access will be allowed for the purposes of troubleshooting, solving issues and improving the effectiveness of the offering. All access is recorded and audited.

Vectra's security research and data science teams may need to copy metadata from Vectra Recall for offline processing, for example to improve detection effectiveness. Any personal information within the data will be anonymized prior to it leaving the customer's Recall instance. All access privileges are managed by engineering leadership and audited for privilege access violations.