VECTRA®

Data Center System Design Guide

About This Guide

This document guides you through design, planning, preparation, and deployment details for operating the Vectra X-series in your data center (DC). The following topics are covered:

- Design goals and considerations
- Preparation
- Network insertion
- Setup

Using the information in this guide, you will be able to accomplish the following planning and deployment tasks:

- Design the Vectra solution deployment to fit your DC environment.
- Configure the X-series appliance, the "Brain" of the Vectra solution, to support the Sensors and virtual Sensors (vSensors) you will place.
- Place Sensors to ensure optimal visibility to threat activity, as well as important machine identity artifacts, allowing Vectra to offer richer insights into suspicious machines in the environment.
- Connect Sensors on the network and "pair" them to the X-series Brain. Physical Sensors will capture network traffic forwarded to them (via a SPAN), while vSensors will capture network traffic passing through the VMware virtual switch on the same physical host. Both Sensor types will distill session information into metadata, and send the session metadata to the brain for detection processing and correlation.
- Configure the virtual switch infrastructure to attach vSensors' management interfaces and capture ports and span capability to enable copying of packets to the vSensor.

NOTE: The Vectra solution supports physical Sensors and virtual Sensors (vSensors). In text that refers specifically to vSensors, the term "vSensors" is used. In text that applies to both types of Sensors, the term "Sensor" is used. In text that applies specifically to physical Sensors, this is stated in the text.

Additional Information

See the resources listed in Table 1 for additional information or to give feedback.

Table 1: Additional information for Vectra X-series

Information	Access
Detection Models	Navigate to https://<your_Vectra_Brain's_Hostname>/resources/ Then click "download" for the *Understanding Vectra Detections* PDF.
Support	Navigate to www.vectra.ai/support
Latest X-series Release Notes	Navigate to https://<your_Vectra_Brain's_Hostname>/resources/ Then click "download" for the *Release Notes <version>* PDF.
Feedback	Send email to support@vectra.ai.

VECTRA

TABLE OF CONTENTS

Vectra System Overview

The Vectra Data Center Solution (the Vectra System) automatically detects network threats—in real time—by passively monitoring network traffic. The Vectra System consists of the following main components:

- **Brain:** The Brain is the central controller, which is paired with distributed Sensors.
- **Sensors:** Sensors capture network traffic flowing through the physical and virtual network, then distill the captured traffic into metadata about each traffic session. A session consists of the standard 5-tuple: source and destination IP addresses, Layer 4 protocol (TCP or UDP), and source and destination protocol ports.

Each Sensor sends its distilled metadata to the Brain. The Brain then analyzes and correlates the metadata to detect suspicious activity or potential threats. The Sensors are a set of eyes distributed throughout the network that provide the Brain with a relatively complete view of the hosts and compute machines in the environment. Figure 1 shows a high-level, logical view of how the Vectra System components relate to each other, and how they receive packet copies from the monitored network.

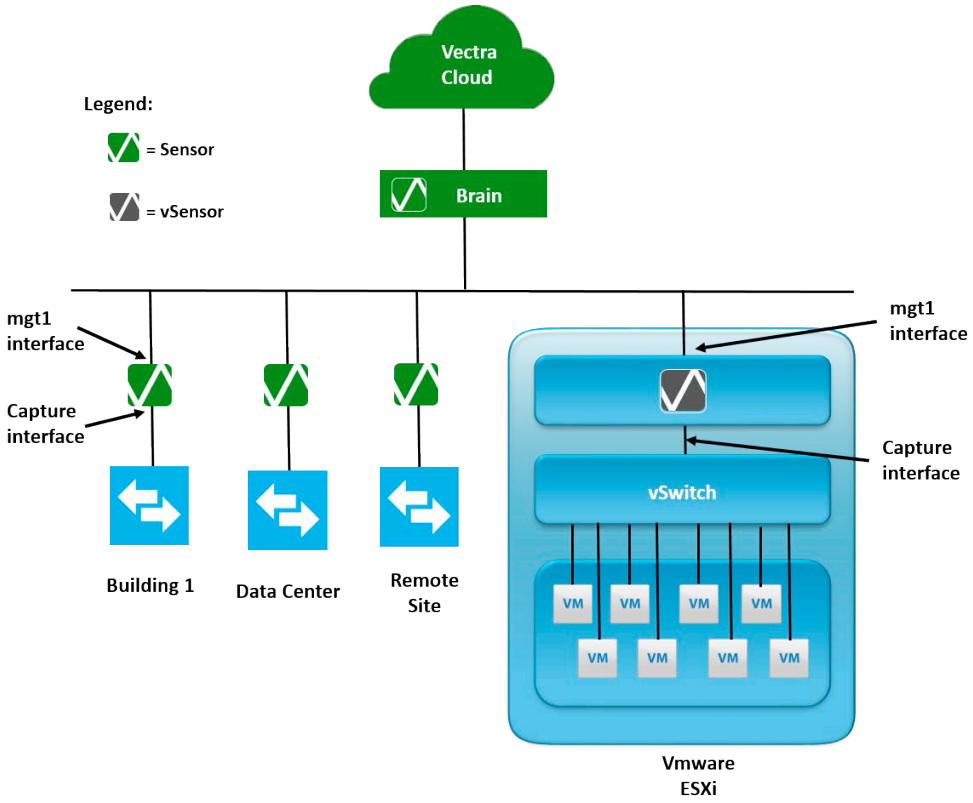

Figure 1: Vectra System logical view

Each Sensor ingests network traffic when (one of) its capture interface(s) connects to a switch port configured as a Switch Port Analyzer (SPAN), to a Test Access Port (TAP), or to a similar construct in a virtual environment such as a VMware port group configured in Promiscuous Mode.

Each Sensor also has a second interface, its management interface, over which it talks to the Brain.

In order to properly analyze network connections for suspicious behavior, each Sensor must see bidirectional traffic flows of the hosts/compute nodes for which monitoring is desired.

X-series platforms can be configured as Brain, Sensor or both. (See Table 2.)

When Sensor functionality is enabled, physical Ethernet ports that are dedicated to packet capture functions are plugged into SPANs or TAPs.

The following types of Sensor appliances are available: physical Sensors and virtual Sensors (vSensors). Physical Sensor functions can be performed by all X-series platforms when configured in Sensor-only mode, and by the S-Series dedicated Sensor appliances.

Table 2 describes each X-series deployment mode.

Table 2: X-series deployment options

X-series Deployment Mode	Description
Brain-only	• Functions as the Vectra system controller • Manages all Vectra System elements, including distributed Sensors • Analyzes metadata for suspicious activity and threats • Describes observed threat detections • Correlates detections to hosts • Provides management access (through Vectra UI and CLI)
Sensor-only	• Ethernet capture ports enabled • Brain function disabled • Must be paired with a Brain (separate device) • Captures all packets received on capture ports • Distills captured flows into session metadata • Sends metadata to Brain
Mixed	• Brain function enabled • Ethernet capture ports enabled for local packet capture and distillation • Additional physical Sensors or vSensors also may be paired with the Brain function

Figure 2 depicts a Vectra Data Center Solution that uses the following X-series components:

- One X80 platform in Brain-only mode
- One X24 in Sensor-only mode as a physical Sensor
- Two S2 physical Sensors
- Nine vSensors.

While this example shows model X80, any X-series platform model can serve as a Brain. Likewise, in addition to the S2 platform, any X-series platform model also can serve as a physical Sensor.

Figure 2 also shows the Vectra cloud services to which a Brain will need to connect, as well as a VMware vCenter server to which the Brain will make Application Programming Interface (API) calls.

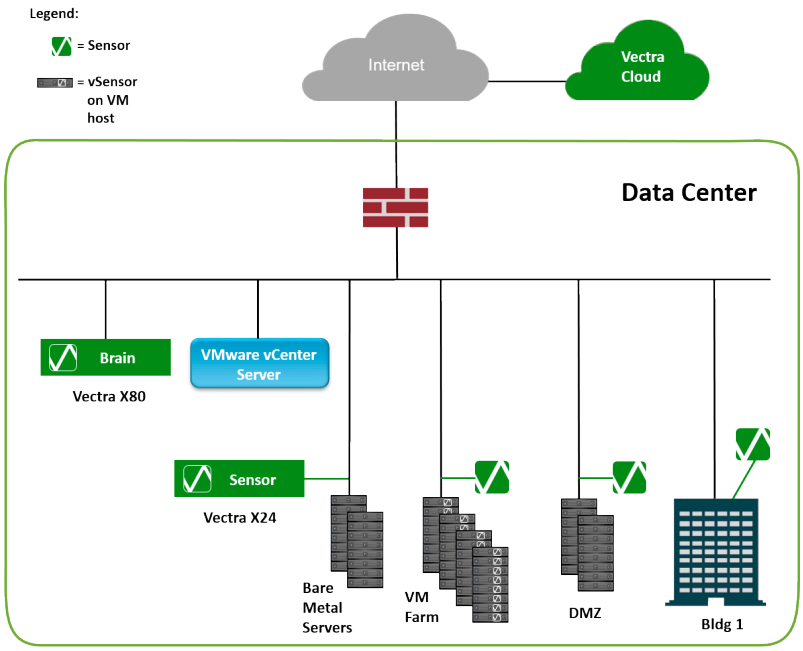

Figure 2: Vectra Data Center Solution

Every Vectra System component, Brain and Sensor, has a dedicated management port called “mgt1”. Table 3 describes how the Brain and Sensors use the mgt1 port.

Table 3: Uses of mgmt1 port

X-series Component	Mgt1 Use
Brain	• Vectra component management (liveness detection, updates, pairing, and so on over SSH) • Vectra Cloud Service connectivity • VMware vCenter API connectivity • Management through Vectra UI (web) and CLI (over SSH) • Network services (DNS, NTP, Syslog, PING, and so on) • Metadata and PCAP collection from Sensors
Sensor	• Management by Brain (liveness detection, updates, pairing, and so on over SSH) • Metadata and PCAP delivery to Brain • CLI for bootstrap configuration, pairing and provisioning (over SSH) • DNS and ping for connectivity verification

Understanding these message types helps during planning for Vectra System insertion into the data center network, as described in the Select Network Insertion section.

Determine Network Traffic To Monitor

Table 4 lists the types of network traffic you will want to monitor.

NOTE: When the term “host” is used in this document, it refers to a node on the network that originates and terminates connections, such as end points and servers. Under this definition, “hosts” may be either physical or virtual machines.

When the term “physical host” is used, it refers to server hardware on which virtual machines reside.

System to detect devices involved in an in-progress attack in a data center environment.

Table 4: Traffic types monitored by X-series

Traffic Type	Types of Detection
DC host to Internet	• Command-and-control communication • Botnet monetization (click fraud, bitcoin mining, spam, and so on) • Data exfiltration
DC host to DC host This includes the following types of connections: • Physical node to virtual node • Virtual node to virtual node • Physical node to physical node • Application font-end nodes • Application processing nodes • Data stores	• Management by Brain (liveness detection, updates, pairing, and so on over SSH) • Metadata and PCAP delivery to Brain • CLI for bootstrap configuration, pairing and provisioning (over SSH) • DNS and ping for connectivity verification
DC host to/from user machine Example: application front-end systems	• Reconnaissance • Lateral movement • Data acquisition
DC host to/from DMZ host Example: between application front-end nodes and application processing nodes	• Reconnaissance • Brute-force login attempts • Lateral movement • Data acquisition
DC host and users to authentication servers	• Brute-force login attempts • Lateral movement Also used for host identification.
Users accessing the DC from remote access VPN	• Command-and-control communications • Botnet monetization (click fraud, bitcoin mining, spam, and so on) • Data exfiltration Also used for host identification.
Compliance-governed machines Example: Payment Card Industry (PCI) data security, Health Insurance Portability and Accountability (HIPAA)	• Full monitoring required for compliance

Consider where in your network each of these interactions occurs. The next section will explore the physical placement of the Brain device and Sensors, in order to capture an appropriate amount and appropriate cross-section of traffic.

Select Network Insertion Points for Vectra Components

This section describes the types of insertion points where Vectra physical and virtual components (Brain and Sensors) should be installed for maximum protection.

These are the interfaces that need to be inserted into the network topology:

- Brain's management interface (mgt1)
- Brain's capture interfaces (if they are to be used)
- Each Sensor's or vSensor's mgmt1 interface
- Each Sensor or vSensor's capture interface(s)

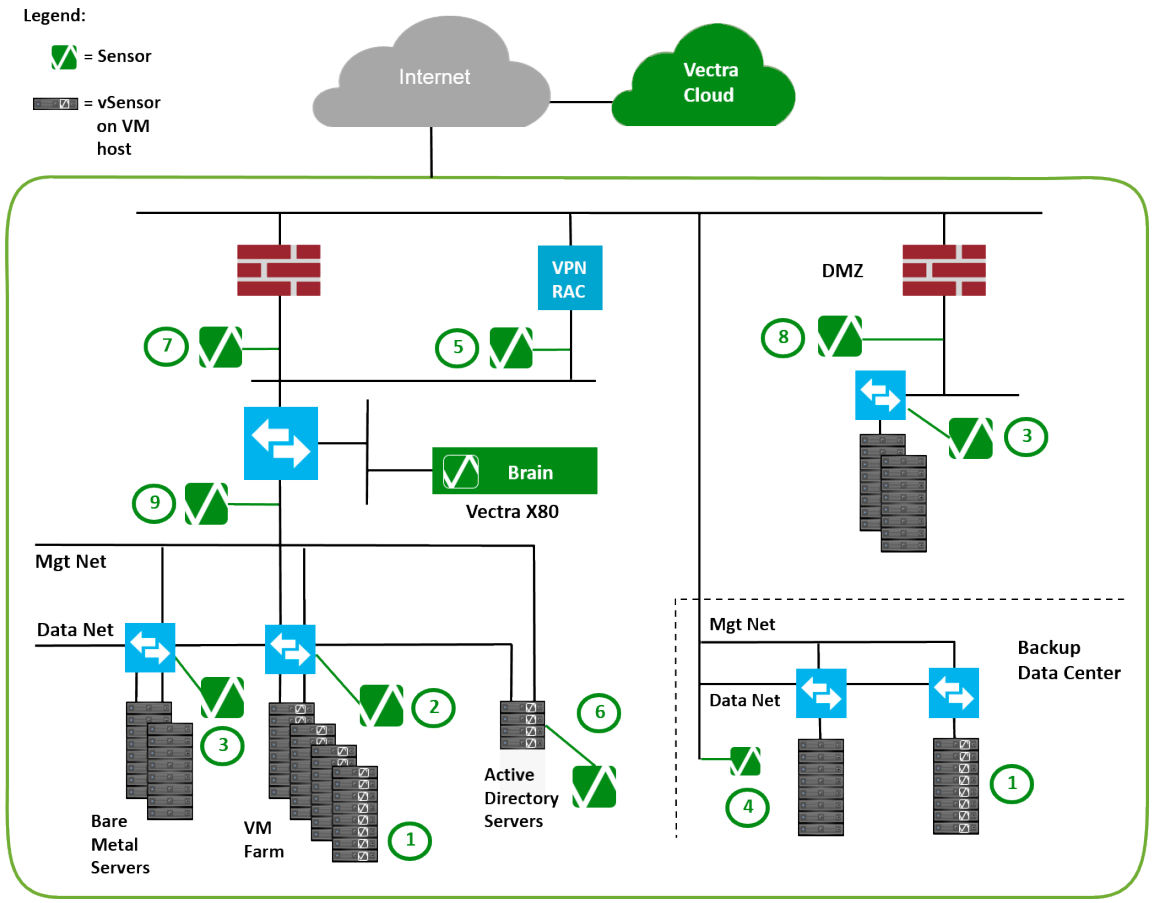

Figure 3: Sensor/vSensor insertion points for data center topology

Figure 3 shows an example of a DC topology. Each potential insertion point for Vectra components is labeled (1-9). Each of these insertion points is summarized in Table 5 and described in more depth in the following subsections.

Details about each type of insertion point are provided in Table 5. For more in-depth descriptions, see the sections following the table.

Table 5: vSensor/Sensor insertion points

Insertion Point	Network Location	Sensor Type	Resulting Visibility
1	Virtual switch of hypervisor	vSensor	• DC hosts to DC hosts • User to DC hosts • 100% visibility to all VMs in all ports configured to copy to the monitor port group where the vSensor is attached • Intra-hypervisor connections
2	Switch connecting physical hosts of VMs	Physical Sensor	• DC hosts to DC hosts • User to DC hosts • All VM traffic that exits the physical host **Note:** Zero of the intra-hypervisor VM traffic that stays on physical host is captured.
3	Switch connecting bare metal servers	Physical Sensor	• DC hosts to DC hosts • User to DC hosts • Server traffic that traverses the physical switch **Note:** None of the traffic traversing switches below the switch where the SPAN/TAP is connected is visible.
4	Other DCs (such as DR sites or regional DCs)	Physical Sensor or vSensor	• DC hosts to DC hosts **Note:** Sensors can be deployed in a very distributed way, as long as they can reach the Brain over IP. • Server and VM traffic
5	Inside of Remote Access VPN terminator	Physical Sensor or vSensor	• User to DC hosts • 100% visibility of connections from remotely connected employees into the corporate network
6	Ahead of AAA servers	Physical Sensor or vSensor	• DC host to DC host • User to DC host Visibility into 100% of authentication connections dramatically aids host identification and correlation.
7	Inside DC perimeter firewalls	Physical Sensor or vSensor	• DC host to Internet Ensures visibility into any incoming or outgoing (North/South) transmissions from DC resources
8	At DMZ borders	Physical Sensor or vSensor	• DC host to DMZ host Provides 100% visibility to connections between DMZ(s) and DC subnets
9	Management subnets	Physical Sensor or vSensor	• Management workstation to infrastructure equipment

* In most cases, a physical Sensor is used.

The following subsections provide more information about each of the “insertion points” listed above.

Management Network

Each Vectra component—the Brain and every Sensor and vSensor—has a management port, labeled “mgt1”. The mgt1 interfaces are layer 3, IP interfaces. Each mgt1 interface will require a unique unicast IP address.

The IP addresses of the mgmt1 ports are not required to be on the same Layer 2 Ethernet segment, in the same virtual LAN (VLAN), or in the same IP subnet. They are designed to reach one another over routed networks.

Determine the appropriate network segments / subnets for the management interfaces of the Brain and all Sensors, physical and virtual, wherever you intend to place the components. (This is different from insertion point 9 in Table 5 above, covered in Monitoring Management Subnets below.)

Monitoring Virtual Machines

North-South traffic into or out of the virtual switch on a virtual machine (VM) can be captured on the virtual switch itself, using a vSensor installed on the VM's hypervisor, or using a physical Sensor placed on a physical segment within the traffic path into and out of the VM.

If East-West traffic between VMs on the same physical host travels only through that host's virtual switch and does not traverse a link outside of the physical host, using a vSensor installed on the host hypervisor is the only way to capture the traffic between those two VMs.

If the East-West traffic is forwarded through a router that is outside the VM, a Sensor inserted in the traffic path between the router and the VM can capture the traffic. However, the only way to ensure that 100% of the traffic that passes through a virtual switch is captured is to install a vSensor on the VM's hypervisor.

If a vSensor is used, it must be installed on the hypervisor using an Open Virtualization Format (OVA) file downloaded from the Brain that will manage the vSensor. (Typically, the network will have a single Brain. The OVA files for all vSensors deployed within that network must come from this Brain.)

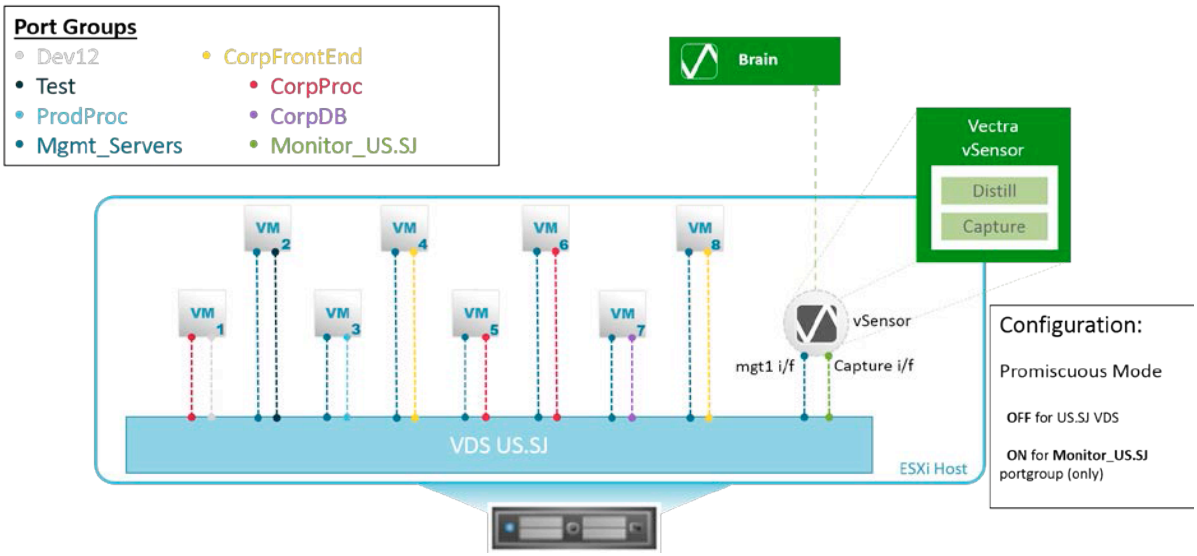

Figure 4: vSensor port group connections

The vSensor has a management port (mgmt1), and one or more capture ports. Place mgt1 in the port group used by the virtual switch for management traffic.

Place the capture port in a port group that has Promiscuous Mode enabled. Promiscuous Mode enables the capture port to copy all packets traversing the virtual switch to the port group. Any ports attached to the port group also will receive copies of the captured traffic.

For environments that make use of multiple virtual switches, up to three virtual switches can be managed using a single vSensor

An example is depicted in Figure 4. The vSensor's mgt1 interface's vNIC (above in red) will attach to the management port group, "Mgmt_Servers", while the capture interface of the vSensor's vNIC (above in green) will connect to the port group dedicated for monitoring, "Monitor_US.SJ.07".

Additionally, to capture all traffic, assign the capture port group to all VLANs. If using VMware Virtual Standard Switch (VSS), define the port group's VLAN setting as "All (4095)". If using Virtual Distributed Switch (VDS), as depicted in Figure 4 above, configure the port group's VLAN setting as "0-4094". If the port group is assigned to only a subset of VLANs, traffic for only those VLANs will be captured.

In the Figure 4 example above, assuming port group “VDS US.SJ” has VLAN set to “0-4094”, any packets traversing any of the ports in the port groups depicted in Figure 4 will be captured by the vSensor, including packets traversing the Dev12, Test, ProdProc, CorpFrontEnd, CorpProc, CorpDM, and Mgmt_Servers port groups. 100 percent of the guest VMs’ (VM1 through VM8) network connections will be inspected by the Vectra System. If 100 percent visibility is required, vSensors are the design choice.

Cases Where a Physical Sensor Can Be Used Instead

Some environments may have little or no intra-hypervisor network traffic. This may be the case if there is not much synchronization between workloads in the same subnet, or if disaster recovery or deployment rules always spread synchronizing workloads in the same subnet across different physical hosts.

It also is the case in some environments that two VMs on the same hypervisor speak to each other, but their respective ports are on different VLANs or subnets configured such that traffic between the two VM’s ports must first leave the physical host’s NIC in order to be routed to the destination subnet, only to come right back to the destination VM on the same physical host, but on a different VLAN or subnet.

If there are little or no intra-hypervisor flows, or if the network topology normally routes packets outside the physical host before returning to a different VM on the same hypervisor, then vSensors may not be necessary. In this case, physical Sensors may adequately monitor the VM’s traffic.

In such a case, Vectra recommends inserting the physical Sensors’ capture ports on the internal side of the DC network, behind the DC edge firewall and proxy, and as close to the compute machines’ physical access switches as possible.

It is advantageous to place the physical Sensor as close as possible to the physical host to maximize visibility of all relevant traffic.

Monitoring Bare Metal Physical Servers

For physical servers (non-VM hosts), the following considerations affect their placement:

- Proximity to physical servers
- Bandwidth capacity

Proximity to Physical Servers

Place Sensor capture interfaces as close to the monitored servers as possible. The optimal placement is in front of a Top-of-Rack (ToR) switch or leaf node.

The next best option is to place the Sensor in front of a spine switch. In this case, make sure that the switch fabric can send any traffic from any switch comprising the spine to the SPAN port connected to the Sensor. The entire spine must be visible in a leaf-spine architecture.

Also make sure that the SPAN configuration delivers bidirectional traffic, both the sent and received packets for any (optimally, all) network connections.

Bandwidth Capacity

Bandwidth capacity involves the following:

- Traffic Volume: The switch where the Sensor is placed must be able to copy the volume of traffic desired at the SPAN port without slowing down the forwarding function fundamental to the switch’s operation. Make sure to not oversubscribe a SPAN port when spanning VLANs.
- Traffic Spikes: The Sensor must be rated to handle the volume of traffic that may be emitted from a second- or third-tier switch. This includes capacity to handle traffic spikes that may be emitted from the switch.

  Traffic volumes greater than the supported level will not cause harm to the Sensor. They also will not disrupt the surrounding network.

  However, the excess packets will be dropped, causing a black-out period for some sessions. The Brain will not have complete traffic flow information to analyze for threats.

TAP Aggregator

Use of a TAP aggregator (such as those from Arista, Gigamon, or Cisco's Nexus product) makes it possible to cost-effectively and without loss monitor all network traffic and divert it to whichever physical port is desired.

In this case, the traffic can be aggregated from anywhere in the multi-switch fabric, and pumped through a specific port. Proximity of the Sensor is no longer a consideration. Instead, the TAP aggregator must have the capacity to handle the traffic.

Monitoring VPN Hosts

A Remote Access Concentrator (RAC) connects remote hosts to the DC. Typically, these are hosts that connect through a Virtual Private Network (VPN).

The RAC assigns IP addresses from a pool to the VPN NICs of the remote devices. The IP addresses belong to the internal network. Usually, the address pool is in a dedicated subnet for the RAC, and the IPs are re-used regularly as the remote devices pin up and tear down VPNs to the RAC.

Placing a Sensor between the inside interface of the RAC and the rest of the DC allows the Vectra System to see all the remote host traffic into the network, which is important because remote hosts are a common attack vector.

This Sensor placement also the Vectra System to view the authentication activity from end-point devices and use the information to clearly identify the devices, even if a device is using an IP address that was used recently by several other devices. This visibility is highly recommended, as the ability to monitor authentication activity will benefit a variety of Vectra detection models.

Place the Sensor on the inside of the DC, where the Sensor can capture the VPN traffic after it is decapsulated by the RAC.

Monitoring Authentication Servers

Place Sensor capture ports so that 100% of authentication requests to and from authentication servers (such as Active Directory) are monitored.

The visibility provided by the authentication traffic significantly aids the Vectra System's mapping of sessions to specific machines/devices, and their owners. This may be accomplished with either vSensors or physical Sensors.

Monitoring the Internal Side of the Data Center Perimeter Firewall

To steal internal data, an outside attacker must have some command-and-control communication with compromised machines inside the data center, and must be able to ship or exfiltrate the collected data out of the data center network.

Some data centers have only one Internet ingress/egress point that internal systems' communications must cross, while others have many. Place Sensors capture ports just inside any such Internet ingress/egress points, such that 100% of North-South traffic traveling through the DC firewall perimeters is inspected by the Vectra System.

This Sensor placement ensures the ability to observe any unauthorized or suspicious connections involving outside attackers, regardless of where other Sensors are placed internally.

Usually, a Physical Sensor is placed at the internal side of a perimeter firewall but a vSensor can be used, depending on your environment.

In the case of a perimeter employing NAT, place the capture function on the internal segment. Because most DC use private IP addresses (as described in RFC 1918), perimeter firewalls typically use network address translation (NAT) to rewrite the Layer 3 (IP) and Layer 4 (TCP/UDP) information so that communication can exit to the Internet. In order to identify compromised systems inside your network, insert capture functionality where the Sensors can see the internal, private addresses of the machines that make outbound connections. This enables you to pinpoint the exact compromised machine when a detection is made.

A Sensor placed outside the firewall would see all detections from the same IP address(es), the firewall's external and routable IP interface(s), and would have no idea which internal machine actually was involved in a connection.

If the data center uses public IP addresses and only routes across its perimeter, and does not use NAT, the Sensor can be placed either inside or outside the firewall. However, a Sensor placed inside also detects any connection attempts that are denied by the firewall, in addition to those that succeed.

Fluctuations in unacceptable connection attempts can indicate compromised machines or misconfigurations. A Sensor placed outside the firewall sees only those internally initiated connections that succeed past the firewall. Moreover, a Sensor placed outside the firewalll would see all ingress, Internet-side initiated connections, including those that are dropped by the firewall. These ingress connections are likely to be very large in number due to the constant onslaught of attackers against high-value DC assets, and will likely create far more noise in the Vectra System than your operations are prepared to wade through.

For this reason, even in routed, non-NAT environments, Vectra recommends placing the Sensor on the network segment inside the firewall.

Monitoring Inside DMZ Borders

Often, a DC design includes a DMZ security zone. The DMZ provides an additional layer of access control between machines that are allowed to communicate with the Internet, and those that are not.

Examples of service machines placed in DMZs include the following:

- Web service front-end servers
- Collectors of external web services
- SMTP gateways
- DNS servers doing zone transfers
- NTP collector

Tight and highly monitored access controls are placed on connections to/from machines in the DMZ to machines inside the DC.

For the same reasons as noted for Sensor placement at the internal side of perimeter firewalls, Vectra recommends placing Sensors on the DC side of DMZ borders, where the internal DC connects to any DMZ security zones.

Monitoring Management Subnets

Management subnets are used exclusively for managing infrastructure equipment in the computing environment. They are differentiated from, for example, application delivery or enduser subnets. A management subnet connects management and monitoring workstations to infrastructure devices such as servers, storage devices, switches, routers, and firewalls for the purpose of configuration, maintenance, and monitoring. An example is the network on which reside physical servers' Intelligent Management Platform Interfaces (IPMI).

Capturing data from the management subnets allows the Vectra system to watch administrative connections for suspicious activity and report anomalies in management patterns.

Additional Sensor Placement Considerations

The following sections provide some additional considerations for Sensor placement in a DC network.

Capturing Bidirectional Flows in Asymmetric Forwarding Planes

The Vectra System needs to see both the ingress and the egress legs of all flows in order to construct a complete flow session (standard 5-tuple), and to extract the flow metadata used to detect suspicious activity and threats.

When selecting Sensor insertion points, make sure that each Sensors' capture port(s) are placed where it can see bidirectional traffic (both the ingress and egress legs of a connection flow).

Many DC center perimeters and switching spines are designed with full, active-active redundancy, including the switch, firewall, and routing elements. Such designs may have asymmetric forwarding paths. If this is the case in your network, make sure a single Sensor receives multiple SPAN feeds, one from each switch at this network point, such that the single

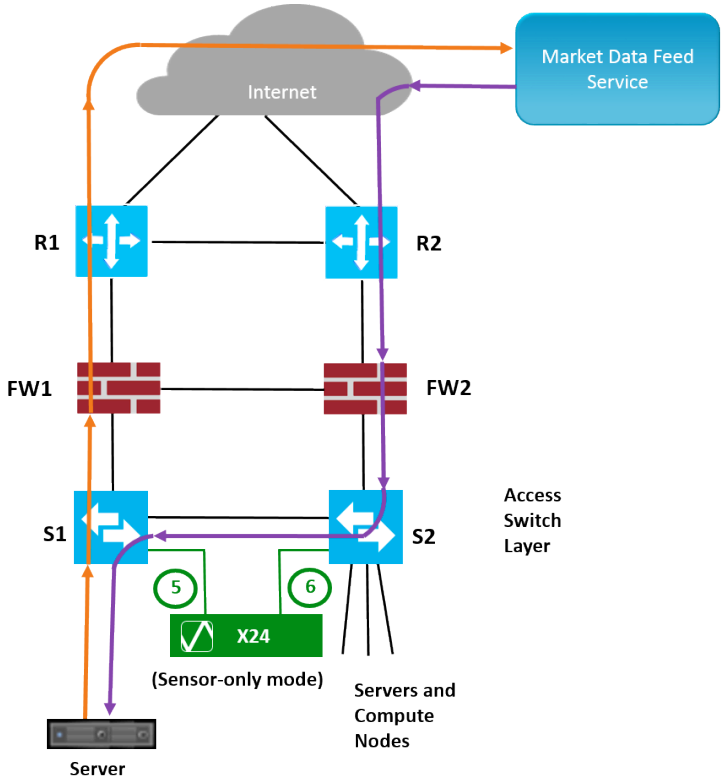

Figure 5: Capturing asymmetric traffic

Sensor receives both legs of any flow. Figure 5 shows an example of a fully redundant, activeactive perimeter with asymmetric forwarding.

The Sensor in this example, an X-series platform model X24 in Sensor-only mode, is placed for full, bidirectional flow capture.

This Sensor receives SPAN traffic from both switch 1 (S1) and switch 2 (S2) on its two 10 GE optical ports, eth5 and eth6. The SPAN rules on the switches are set to copy and forward packets from the same VLANs.

When the Server's flow exits the network through S1, firewall 1 (FW1), and router 1 (R1), the X24 Sensor sees the beginning of the flow on its port eth5.

When the response from the vendor's market data feed service arrives on router 2 (R2), firewall 2 (FW2), and then S2, the X24 Sensor sees the response on its port eth6, and can construct the full flow.

More detail on this topic is found in the Vectra System Design Guide, in the section "Best Practices for Configuring SPAN".

Session Duplication Handled Automatically by Vectra System

Session duplication occurs when multiple Vectra Sensors observe the same single session.

For example, session duplication occurs when a VM is talking to an Internet-dwelling system, and Sensors are capturing packets at the virtual switch, at a spine switch, and also at the DC firewall egress. In this example, the Vectra System will observe the same flow at three different locations.

You do not need to worry about session duplication when planning for Vectra insertion. The Vectra System is designed to identify and handle these cases automatically.

Sensor Performance

As a general rule, for best visibility, place Sensors as close to the end compute machines or end-user machines as possible. This consists of placing the Sensors on the access layer switches to which the machines you would like to protect are connected. This type of distributed deployment typically leads to a larger number of smaller Sensors, and better overall coverage.

Within the Vectra System, the following types of Sensors are available. The best Sensor type to deploy in a given DC network depends upon the environment and on scale requirements.

Throughput is the primary resource to keep in mind when determining the type of Sensor to deploy.

Some of the Sensor types available are dedicated Sensors while the remaining two are X-series platforms that can be run in a Brain-only, Sensor-only, or Mixed mode.

- vSensors integrate with VMware vSphere, are deployed on each hypervisor, and have a full view of the visualized infrastructure.
- Physical Sensor deployment requires more in-depth planning. This is due to the more distributed nature of the physical infrastructure (core, distribution, and access) when compared to a virtual infrastructure.

If you have not already, please review the above section Select Network Insertion Points for Vectra Components to determine which of the insertion points described you want to mimic in your environment. Once this determination is made, analyze the traffic and determine what size and model of physical Sensor needs to be deployed.

Table 6 describes each Sensor device type.

Table 6: Sensor device types

Sensor Device	Description
vSensor	Dedicated capture device for VMware environments. Installed as a virtual appliance, a guest VM on a hypervisor. Deployed on a per-host basis to see all VM traffic, including intrahypervisor connections. Maximum sustained throughput of 2 GB with 8 vCore vSensor. Can be used as replacement for a Physical Sensor. Interfaces: one management port, two or three captures interfaces depending on vSensor size (see Table 7)
S2	Dedicated physical Sensor comparable to a vSensor. Maximum sustained throughput of 1 GB. Interfaces: one management port and two capture ports
X24 in Sensor Mode	Although the X24 is typically deployed as a Brain only or in Mixed mode, it can also be deployed as a stand-alone Sensor. Deployed in Sensor Mode, the X24 supports a total of 8 GB sustained throughput. Interfaces: • One management port • Six capture ports: • Four 10/100/1000 Mbps copper ports • Two optical 10 GB SFP+ ports
X80 in Sensor Mode	The X80 typically is deployed in Brain-only or Mixed mode, but also can be deployed as a stand-alone Sensor. Deployed in Sensor Mode, the X80 supports a total of 20 GB sustained throughput. Interfaces: • One management port • Four optical 10 GB SFP+ capture ports

vSensor as Physical Sensor Replacement

This vSensor deployment model allows for the same sort of insertion as a physical Sensor, but in a virtual form factor. The goal is to capture frames from physical switches, using a virtual form factor to accomplish the capturing.

Use this deployment model to gain visibility into physical machines' traffic, or virtual infrastructure where you do not want to place a vSensor on each hypervisor. This use case is in contrast to the vSensor deployment model where the vSensor sits on each hypervisor's virtual switch in order to gain visibility to 100% of the local VM's traffic.

The benefit of this deployment model is that it removes the need for physical device racking and stacking in order to expand detection scope.

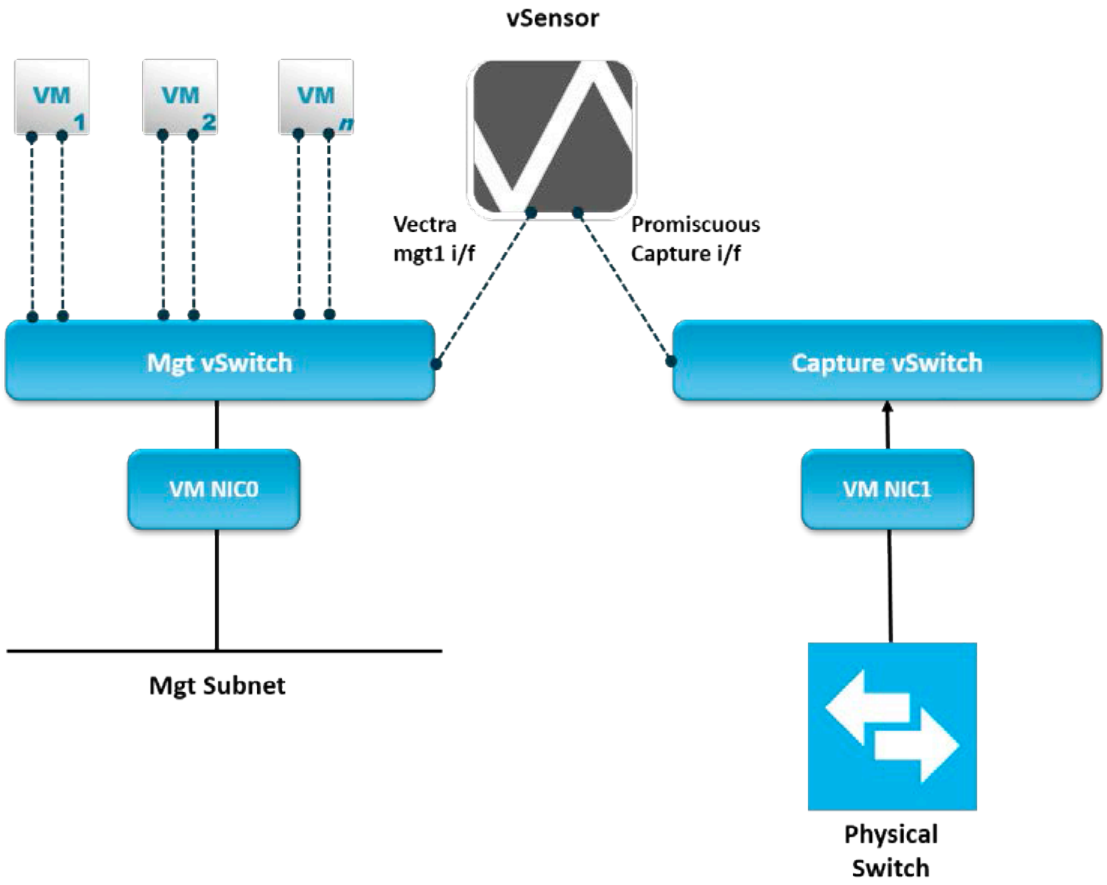

Figure 7: vSensor interface connections

The limits of this deployment model include:

- Partial visibility: Only traffic traversing the physical switch will be seen. Any intra-hypervisor traffic on this or other hypervisors will not be captured. Also, any traffic flows where the path between the source and destination passes switches, topologically excluding the spanning switch, will not be captured.
- Total bandwidth consumable from the physical SPAN / TAP is the performance limit of the vSensor. Currently, the limit is 2 Gbps with 8 vCore vSensor.

Deployment details for this use of a vSensor can be found in the *Sensor Installation Guide*.

Prepare the Network and Virtual Environment for vSensor Insertion

In a virtualized compute environment, using a vSensor is the recommended method of deployment. Using a vSensor has the advantage of integration with vSphere, which allows Vectra to ingest host, guest, and network information (performance, VLANs, and so on) to analyze monitoring coverage and performance.

This section describes the system and environment requirements for inserting vSensors into your DC network.

DC Deployment Preparation Worksheet

Vectra has created a deployment preparation worksheet to help you gather all the necessary information and make changes in the environment (for example, firewall access rules) required in order to complete initial configuration for the deployment. The worksheet is reproduced in printable format (Appendix: Data Center Deployment Worksheet Hard Copy), and also is available in Microsoft Excel format (find link in the Appendix), should an editable format be desired to record information as it becomes available from other teams in your organization.

Create a DNS Entry for Brain IP

It is highly recommended that a DNS entry be made for the IP address of the Brain's management interface (mgt1). The Brain's hostname should have both forward and reverse entries in the DNS system. This is especially true in a DC deployment, where the number of attached Sensors will be large.

With a DNS entry, each Sensor may be configured to reach the Brain using the Brain's hostname, for example: vectra01.company.com. This allows the Brain's IP address to be changed at will, and, as long as the DNS record is updated with the new IP address, the Sensors will continue to connect to the Brain without issue. Otherwise, if the Sensors' configurations use the Brain's IP address instead of the hostname, and the Brain's IP changes, each Sensor would then need to be accessed via SSH to manually change its Brain configuration setting.

NOTE: For the network identities of devices, Vectra recommends using hostnames instead of IP addresses. If the IP address of the Brain ever needs to be changed, you will not need to manually log in and change all of the vSensors' Brain settings.

vSensor Resource Requirements

Table 7: vSensor resource requirements

Resource Type	Requirement		
Performance	400 Mbps	1 Gbps	2 Gbps
Capture Interfaces	2	2	3
CPU	2 vCores	4 vCores	8 vCores
Memory	8 GB vRAM	8 GB vRAM	16 GB vRAM
Drive	100 GB	150 GB	150 GB
VMware vSphere	5.0 or later		
Virtual switch type	VMware Virtual Standard Switch (VSS) or Distributed Virtual Switch (VDS), also called dvSwitch		

Table 7 lists the resources required on each VM host machine where a vSensor will be installed.

The above resources will be allocated for the VM when it is created from the OVA template in vSphere.

Firewall/Access Control

Access control mechanisms may be in effect in your network environment that, left unchanged, would prevent the Vectra System from operating correctly once the components are deployed.

NOTE: Use hostnames instead of IP addresses if your access control systems/firewalls support it. This allows the IP address of the interface to be renumbered, the DNS entry to be updated accordingly, and the Vectra System's access control to work automatically through such a change, without reconfiguration of the access control system.

Table 8: Access control connections used by X-series (sourced from Brain mgt1 hostname)

Destination	Purpose	Detail
api.vectranetworks.com TCP 443 (TLS/SSL)	Functions and lookups performed in Vectra's cloud for the Brain	Recommended Necessary for Virus Total lookups of contextual information on suspect IPs and domains, as well as other cloudbased services that deliver beneficial features.
update2.vectranetworks.com TCP 443 (TLS/SSL)	Software updates, including detection model and Vectra packages, tech support statistics	Required Necessary for security patches, model enhancements, pro-active support monitoring, as well as new features. Enable from **Settings > System**.
metadata.vectranetworks.com TCP 443 (TLS/SSL)	Transmission of anonymized metadata to Vectra data cluster	Optional Data used to create new detection models, and improve existing ones. Sent when "metadata sharing" is enabled in **Settings > System**.
vpn.vectranetworks.com TCP 443 (TLS/SSL)	Remote access for Vectra Support	Optional
or UDP 9970	Initiated from Brain to Vectra Cloud Service	Allows authenticated and authorized Vectra staff access to the Brain, and from the Brain to the vSensors. **Note:** UDP will operate faster than TCP.*
DNS server(s) UDP 53	Domain name resolution	Required
Network Time Protocol (NTP) TCP 123	System time updates	Recommended
Syslog TCP 514	Log collection	Optional The Syslog connection options are configurable within the Vectra UI in order to meet your environment's Syslog use. The corresponding firewall rule detail will depend on those options.
Simple Mail Transfer Protocol (SMTP) TCP 25	Alert notification	Optional Email notification options are configurable within the Vectra UI. The corresponding firewall rule detail will depend on those options.
vSensors' mgt1 hostnames TCP 22 (SSH)	Troubleshooting	For connection from the Brain CLI to the CLI on a vSensor. May be used by customer support, or during the product Beta cycle.
DNS server(s) ECMP Echo Request (ping)	Test network forwarding, route path, aliveness, and so on	Recommended

* Such access may be needed for technical support, professional services, or Beta programs.

Table 9: Access control connections used by Sensor (sourced from Sensor's mgt1 hostname)

Destination	Purpose	Detail
Brain mgt1 IP address TCP 22 (SSH)	Management Metadata passing	Required
DNS server(s) UDP 53	Resolve Brain hostname Resolve any hostnames pinged during troubleshooting	Required
DNS server(s) ECMP Echo Request (ping)	Test network forwarding, route path, aliveness, and so on	Recommended

Table 8 lists the connections used by the Brain. Table 9 lists the connections used by Sensors.

Deciding Which VLANs To Monitor

As part of the insertion process for a vSensor, you will need to add its Capture port to a port group that is a member of all the VLANs you want to monitor.

For optimal coverage, all VLANs carrying application or management traffic should be captured.

Exceptions

VLANs dedicated to network-based I/O traffic, such as iSCSI or vMotion, do not need to be monitored. If possible, such VLANs should NOT be captured, for the following reasons:

- The Vectra System does not detect threats at an I/O level, so there is no use to consume such packets.
- The volume of these connections is orders of magnitude larger than that of application and management traffic. Excluding these connections therefore will improve the chances that the capture volume arriving at the vSensor's capture port will be within the supported volume.

NOTE: In the case of a VMware VSS switch, VLANs cannot be isolated from the promiscuous setting – the configuration is either one VLAN or all VLANs (4095). For VSS, the I/O-related VLANs will need to be taken into account when assessing the bandwidth requirements.

Monitoring IPMI Interfaces

One additional note in a virtualized environment is the presence of Intelligent Platform Management Interface (IPMI) interfaces. Because IPMI interfaces are out-of-band management interfaces, they will not be monitored by vSensors. For this reason, it will be necessary to deploy a physical Sensor to monitor any traffic on IPMI networks.

Determining Network Utilization on a Hypervisor

A vSensor can support up to 2 Gbps of overall throughput. If more than this amount of traffic regularly passes through the hypervisor, the vSensor will not be able to capture all the traffic. In this case, a physical Sensor may be needed instead.

To determine the network utilization on a hypervisor, you can use either of the following:

- vSphere Web Client
- Vectra UI (if vSphere connection is configured in Vectra). See Table 10 in Brain Initial Configuration Settings for Data Center Deployment.

vSphere Web Client

To see the network statistics of a given host in the vSphere Web Client (shown here for 6.0):

1. Login to vSphere.
2. Navigate to **Hosts and Clusters**.
3. On the Hosts and Clusters page, select the host to be analyzed in the left pane. In the right pane, select the Monitor tab.
4. Once the main pane updates, select Performance.
5. Select Overview and scroll down until the network statistics are visible. (For a more detailed view of the traffic, select **Advance**.)

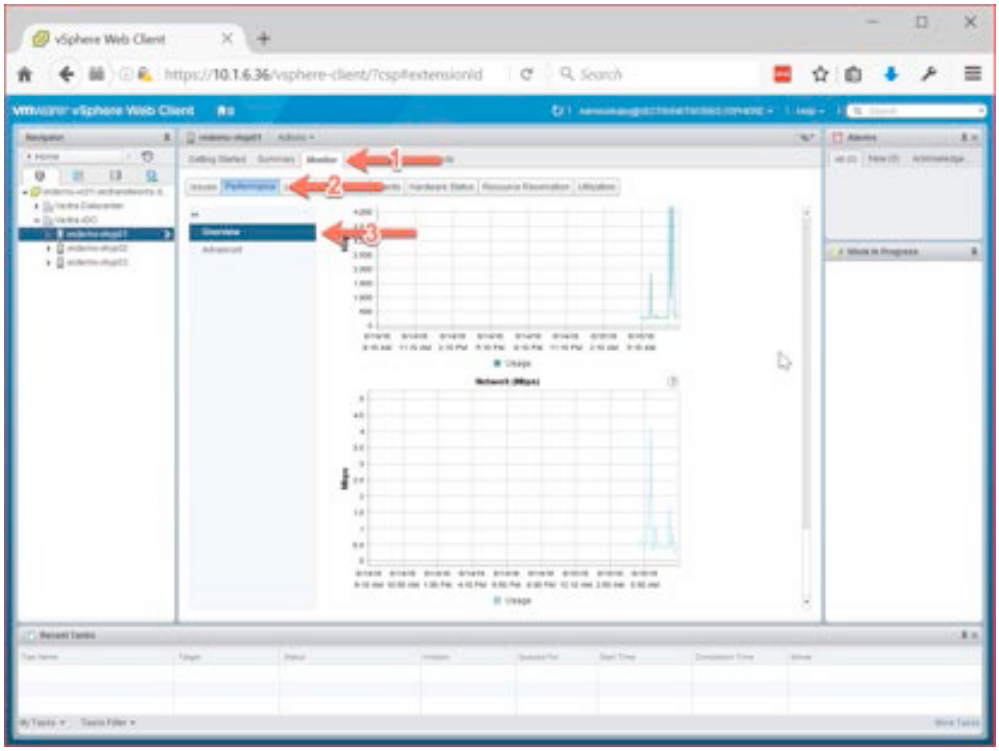

Vectra UI

If the Vectra appliance is configured and is synchronized with vSphere, the same information can be gathered from the Vectra UI. (See Table 10 in Brain Initial Configuration Settings for Data Center Deployment.)

To view bandwidth utilization:

1. Log in to the Vectra UI on the Brain.
2. Select **Manage** in the menu on the left.
3. From the Manage page, select **Physical Hosts** along the top.
   This will populate the primary pane with all of the hosts that are being managed by vSphere.
4. Click the arrow next to a given host to expand and provide details for the host, including network throughput.

Preparing vSphere Port Groups

For a vSensor to be able to capture traffic, Vectra recommends the use of a port group set to Promiscuous Mode. The Capture port group is the port group into which the vSensor's capture port will be placed. All frames passing over the virtual switch will be duplicated and pushed across the Capture port group.

A port group is an object set in the context of a virtual switch. The following types of VMware virtual switches exist:

- VSS: This is the basic Virtual Switch that comes with a VMware installation
- VDS: More robust networking capabilities come along with the ability to define one Virtual Switch with port groups and networking objects that will then stretch across multiple physical hosts. VDS is available with VMware's Enterprise Plus feature pack.

NOTE: A vSensor is a virtual appliance. Though its VM hardware resources (including processor, memory, and drive) and can be reconfigured (increased or decreased) through vSphere like other VMs, such changes **will not be recognized** by the vSensor's OS and potentially will have an adverse effect on the vSensor. **DO NOT CHANGE THE vSENSOR's HARDWARE RESOURCES** from those set by the OVA at installation.

The following instructions apply to vSphere 6.0.0. While the exact windows and options may differ slightly in your version of vSphere, the concepts will be very similar.

Adding the Capture Port Group

This section describes how to add the Capture port group to the virtual switch. Use the procedure for the type virtual switch (VSS or VDS) that is used by the hypervisor.

Adding a Capture Port Group to VSS

In VSS there can be only one VLAN assigned to a given port group. In this case, there is a special VLAN number that will allow the capture port group to receive traffic from all VLANs on the virtual switch: All (4095). Set the Capture port group's VLAN to "All (4095)".

1. From within the vSphere Web Client, select the host that needs to be configured from the navigator pane on the left.
2. Once the details for this host have populated in the primary pane on the right, select the **Manage** tab.
3. Within the **Manage** view, select **Networking > Virtual Switches** and select the virtual switch that you want to monitor from within the Virtual Switch table.
4. Click on the Add networking icon.

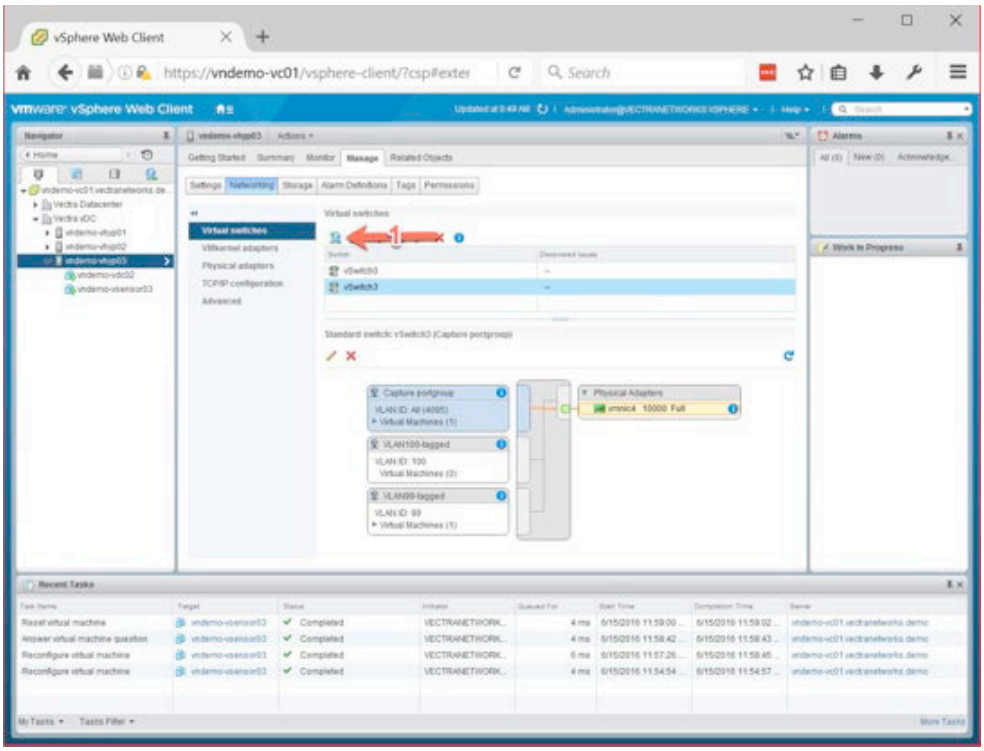

5. A popup appears, to walk you through the process of configuring a new port group.
   Define the port group type as a Virtual Machine Port Group for a Standard Switch, and click **Next**.

6. Next will be the assignment of the port group to the appropriate virtual switch. Assuming the appropriate virtual switch was selected in the step above, the correct virtual switch should be defined.

   If you accidentally select an incorrect virtual switch, click **Browse** and select the appropriate switch.

   Click **Next**.

7. In **Connection Settings**, enter a Network label of your choice to serve as the name for this port group (example: “Capture” or “Monitor”), and set the VLAN ID to “All (4095)”.

   Click **Next**.

8. Finally, review the configuration and click **Finish** to create the Capture port group.

Adding a Capture Port Group to VDS

In the case of a VDS, it is recommended to specify VLANs 0-4094 (all traffic) if possible.

Exception: As mentioned in Exceptions, if any VLANs are dedicated to I/O traffic or vMotion traffic, exclude those VLANs from the port group. In this case, in the port group configuration, a list of VLANs or VLAN ranges (example: “200-500”) can be provided that will allow for selective interaction between the port group and the VLANs.

1. From within the vSphere Web Client, select the host that needs to be configured from the navigator pane on the left.
2. Once the details for this host have been populated into the primary pane on the right, select the **Manage** tab.
3. In the **Networking** tab, select the distributed virtual switch that the port group will be assigned to. This should be the VDS that you want to monitor.
4. Select **Actions > Distributed Port Group > New Distributed Port Group**. This will launch a window to walk you through the configuration of the port group.
5. Assign a name to the Capture port group and click **OK**.
6. Configure the port group. In the VLAN section, select a type of VLAN **Trunking** and provide a list of VLANs to be monitored. VLANs can be entered as single VLAN IDs or as ranges. Separate the VLAN IDs or ranges with commas (example: 10,12,15-25,27-4094.)

   For all possible VLANs (both now and for anything that may be added in the future) use 0-4094.)

7. Review the configuration and click **Finish**.

Enabling Promiscuous Mode on the Capture Port Group

Enabling promiscuous mode on the Capture port group causes any frame passing over the virtual switch to be replicated to the vSensor, like on a SPAN port. Using a single port group that can listen to all VLANs eliminates the need to have an interface for each port group on the vSensor. Using a single port group for all VLANs also eliminates the need to make monitoring object configuration changes every time VM ports are added to the environment or moved.

Enabling Promiscuous Mode in VSS

1. Login to the vSphere client, and select the host that needs to be configured from the navigator pane on the left.
2. Once the details for this host have populated in the primary pane on the right, select the **Manage** tab.
3. Within the **Manage** view, select **Networking > Virtual Switches** and the switch you would like to configure. This will then update the primary pane with a view of the virtual switch.
4. Once the virtual switch is displayed, select the Capture port group that needs to be configured.
5. Click on the pencil to edit the properties of the port group.
6. Navigate to Actions > Edit Settings.
7. Select **Security** in the left pane and change **Promiscuous** from **Reject** to **Accept**.
8. Click **OK**.

Enabling Promiscuous Mode In VDS

1. Login to the vSphere client and navigate to the Networking view from the left pane.
2. Once in the Networking view, select the capture port group that was created on the distributed switch.
3. Navigate to Actions > Edit Settings.
4. Select Security in the left pane and change Promiscuous from Reject to Accept.
5. Click OK.

Prepare vSensor Configuration Settings

To prepare for installing a vSensor, collect or prepare the following settings. The preparation worksheet (see Appendix: Data Center Deployment Worksheet Hard Copy) contains spaces to record these settings.

Recommended Initial Deployment: Three vSensors

Based on deployment experience with other customers, Vectra recommends deploying three vSensors at first, across a variety of physical hosts, and doing so through the vSphere client. This approach allows the multiple teams within your organization to get used to, understand, and observe the vSensors in action before agreeing to a broader deployment in a production environment.

Though Vectra provides CLI on the X-series for rapid creation and provisioning of vSensors, Vectra recommends that you create and deploy the first three vSensors through the vSphere client interface, one at a time, so that you understand the constructs and settings at play.

After deploying and verifying the first 3 vSensors manually, you may want to use the vSensor deployment CLI to quickly create and provision the additional vSensor virtual appliances required for your monitoring design. Details about the vSensor deployment CLI are located in the Sensor Installation Guide.

Sheet 2 on the preparation worksheet (see Appendix: Data Center Deployment Worksheet Hard Copy) contains spaces to fill in the above list of configuration settings for each additional vSensor to be deployed, so that many vSensors' details can easily be added, copied, and moved. Gathering and noting this information before you begin deployment will significantly ease and hasten the deployment.

vSensor General Settings

Table 10 lists the general settings that will be used over multiple vSensor installations.

Table 10: vSensor configuration settings—general

vSensor General Settings	Description
Virtual switch	Name of the virtual switch to which the vSensor ports will connect. This will be the same for all hypervisors if using a VMware VDS. Otherwise, if using VMware VSS, it may be a different name from hypervisor to hypervisor, depending on your VMware environment. Check with your VMware administrator to determine which virtual switch is the correct one for each host.
Capture port group	Virtual switch port group that will be used to copy the packets from the virtual switch to the Vectra vSensor port connected to the Capture Port Group (the vSensor capture port). If you already have a port group configured on this virtual switch for monitoring, you can use that same port group. If not, you will need to create a new port group. Set the port group to Promiscuous Mode, and set the port group's VLANs to match the VLANs to be monitored. VMware VSS allows either a single VLAN ID or all VLANs (VLAN ID 4095). **NOTE:** The deployment planning spreadsheet already has "4095" as the default for VSS. Vectra recommends this setting, unless there is a reason why only one VLAN should be monitored. See Deploy vSensors for a more detailed discussion and instructions on VLAN configuration for the monitor port group. After determining which VLANs to monitor, enter them into the deployment planning spreadsheet.
vSensors' management interfaces port group	Name of the port group to use for the vSensors' management interface (mgt1). Make a note of the VLAN associated with the port group. If using VMware VDS, this port group will be the same for all hypervisors. Otherwise, if using VMware VSS, you will need to create the port group on each hypervisor's virtual switch. Nonetheless, Vectra recommends using the same port group object name across all relevant hypervisors. Ensure this port group's VLAN is configured on the corresponding ports of all upstream switches to which the target physical hosts' physical uplink ports are connected. See Enable Sensor Management VLAN on Upstream Switches for more information.

Individual vSensor VM Attributes

Table 11 lists the specific settings to prepare for each unique vSensor to be deployed.

Table 11: vSensor attributes

vSensor Collection Settings	Description
Physical Host	Hostname of ESXi host on which the VM for the vSensor virtual appliance will reside
VM name for vSensor	Name of vSensor VM in VMware
Hostname	DNS hostname for vSensor
vSensor name in Vectra System	Name of vSensor in Vectra System The Vectra UI allows you to enter an arbitrary string for the name of this vSensor. For consistency, Vectra recommends using either th VMware VM machine name or the fully qualified domain name (FQDN), for example, “test01.sjc.vectranetworks.com”, as the name in the Vectra system. By default the VM name from VMware will be used upon pairing to the Brain, but this label can be edited.
Virtual switch on hypervisor	Name of the virtual switch on the hypervisor where the vSensor is installed If using VMware VDS, this will be the same for all hypervisors. If using VMware VSS, the name may be different for each hypervisor.
Datastore	Name of the datastore on the physical host’s VM disk. This should be determined by your VMware administrator.
DHCP	Enabled (use DHCP) or disabled (do not use DHCP) If enabled, Dynamic Host Configuration Protocol (DHCP) will be used to automatically assign the management IP address, default gateway address, and DNS server addresses. In this case, these values do not need to be added to the deployment planning worksheet. Those rows can be skipped.
Management IP address	IP address of the management interface (not the data collection interface) on the vSensor
Management default gateway	IP address of the default gateway through which traffic from the vSensor’s management interface will travel to reach other Layer 3 IP subnets
DNS servers	IP address of each local DNS server to be used by the vSensor for resolution of DNS hostnames into their corresponding IP addresses
Pin vSensor	VMware affinity rule that “pins” the vSensor to its physical host If VMware Distributed Resource Scheduler (DRS) and vMotion is used to dynamically rebalance VM load across hypervisors, the vSensor should be pinned to the hypervisor on which you install it, to prevent the vSensor from automatically being migrated to another hypervisor. Have the VMware team create an affinity rule that pins the vSensor to its physical host, and ask them to not add it to vMotion. Only move the vSensor if you no longer want monitoring of the physical host. In this case, Vectra recommends destroying the vSensor and creating another one from scratch on the new hypervisor.

Enable Sensor Management VLAN on Upstream Switches

The vSensor’s port used for the management interface (mgt1) will be placed in a port group appropriate for infrastructure management traffic. That port group will have a specific VLAN associated with it, and will be configured in VMware’s vCenter.

Each physical host on which a vSensor resides will have many guest VMs, many port groups, and their many corresponding VLANs. These many VLANs likely will pass through the physical Ethernet ports (NICs, VM NICs) to an upstream switch(es) with a trunk configuration.

Be sure to enable the vSensor’s management interface (mgt1) port group’s VLAN on the upstream switch(es) port trunk connected to the physical host where the vSensor will reside. Without the VLAN enabled thus, the newly created vSensor will fail to connect and pair to the Brain.

NOTE: For each host where a vSensor is placed, make sure the ports on the upstream switches connected to the target host’s physical uplink ports are configured to pass the VLAN of the vSensor’s mgt1 subnet.

For example, assume three vSensors are to be placed on each of three physical hosts named “ussjcesx010”, “ussjcesx011” and “ussjcesx012”. For each of those physical hosts, the port associated with the vSensors’ mgt1 interface, Network adaptor 1 in vCenter, will be placed into a port group called “Infrastructure_Mgmt”.

The “Infrastructure_Mgmt” port group is configured for VLAN 10, and its uplink port is “vmnic2” on all three physical hosts. Assume the uplink ports are physically cabled to switch ports ussjcswitch022 eth 1/9, ussjcswitch024 eth 1/23, and ussjcswitch024 eth 1/24, respectively. The networking team will need to configure the trunks on all three of those switch ports to include VLAN 10.

Making this switch configuration **before** the vSensors are created helps deliver a smooth vSensor deployment and pairing process.

Brain Initial Configuration Settings for Data Center Deployment

Table 12 lists and summarizes the settings on the Brain relative to vSensor deployment and management. These settings are located on the **Settings > System** page in the Vectra UI. Each setting is discussed in more detail following the table.

Table 12: Brain initial settings for data center deployment

Setting	Description	Recommendation
Sensors > Auto-pairing	When vSensors first boot up they attempt to connect to the Brain that created the OVA from which they were spawned. Auto-pairing automatically accepts the pairing requests from all Sensors. Auto-pairing is particularly useful when spinning up multiple vSensors at the same time. **Note:** For a vSensor to be able to pair with the Brain, the vSensor must be installed from an OVA downloaded from that Brain.	Enable when you are deploying vSensors, particularly multiple vSensors using automated scripts, for deployment speed and convenience. Disable otherwise. Disabling prevents rogue, imposter vSensors from pairing to the Brain.
Sensors > Password	When Sensors first come online they have a default user name (**vectra**) and default password (**youshouldchangethis**) for their CLI access over SSH. It is strongly recommended that you change the password as soon as possible. This feature allows you to change the password for all Sensors at once. After the new password is specified here, the Brain will change the vectra user password for CLI access on all attached Sensors, and all subsequently paired Sensors.	This is a security versus convenience trade-off regarding CLI access on Sensors. Specify this password configuration if you prefer the convenience of a single password for all your Vectra Sensors. Leave this empty if you prefer the more secure option, to manually log in and set a unique password for the **vectra** user on each Sensor. **Note:** If you do neither option, the **vectra** user password will remain at its factory default setting.
VMware	Vectra System will query the VMware vCenter for device information, using the vCenter API. Enabling this feature provides a read-only view into the vSphere state that can otherwise be obtained only by logging into vSphere itself. Using the **VMware** option helps with vSensor deployment planning by indicating where vSensor coverage currently exists and where it does not exist (where monitoring currently is not occurring). This option also provides information about available resources on physical VMware hosts, and indicates configuration errors that might be affecting packet capture.	Enable. Strongly recommended. Helps Vectra operator to know more precisely what they need as they engage VMware operational teams. (See Prepare the Network and Virtual Environment for vSensor Insertion.)

Enabling VMware vSphere Integration

If using VMware vSphere, some configuration is required to enable the Vectra System to query the vCenter API. Enabling API access to vCenter provides a read-only view into the vSphere state, otherwise obtainable only by logging into vSphere itself.

Enabling the vCenter API query connectivity helps with vSensor deployment planning by identifying the physical hosts, clusters and data centers that currently have vSensor coverage, and those that do not have coverage.

Enabling the vCenter connection also shows available resources on physical VMware hosts, and exposes any configuration errors that might be affecting packet capture. This view, seen in the Vectra UI **Manage > Physical Hosts** page, helps the Vectra System operator identify the exact requirements that need to be conveyed to VMware operational teams.

Once this setting is enabled, the **Manage > Physical Hosts** page appears in the Vectra UI.

Through the vSphere connection, the Vectra Brain drives email notifications to the configured administrators about changes in the virtual environment that merit security consideration. For example:

- New physical server where a vSensor may be needed is added to the network
- vSensor has been moved or powered down
- VM is moved from a host that is monitored by a Sensor to a host that is not monitored by a Sensor

NOTE: Vectra strongly recommends enabling the VMware integration setting, as a best practice. More about this integration is covered in Deploy vSensors.

To enable VMware integration, enable the **VMware** option on the Brain, and prepare a vSphere account to use for read-only access by the Brain.

Prepare vSphere Account for Brain Access

To connect the Brain to vSphere, a vSphere user account and password must be configured into the Brain. The vSphere user account must have at least global, read-only rights.

To ensure that the vSphere user/group the Brain will use has global, read-only access, use the following steps in the vSphere UI:

1. From the vSphere Administration page select **Access > Global Permissions**.
2. Click the plus sign to display the global permissions dialog.
3. At the bottom of the left pane, click **Add**.
4. Ensure the domain is set to the proper domain, select the users or groups you intend to use in Vectra's configuration to connect to vCenter's API, and click **OK**.
5. In the Assign Role section, select **Read-Only** from the drop-down list.
6. Make sure the **Propagate to children** checkbox is selected, and click **OK**.

Deploy vSensors

This section describes setup tasks that are specific to deploying vSensors in a DC network.

NOTE: This section assumes that the vSensor VMs have been created and that their management and capture interfaces have been attached, based on the instructions in the *Sensor Installation Guide*. Only after the vSensors VMs are installed and attached, continue with the section below.

NOTE: This section does not cover insertion of Physical Sensors. Instead, see the *Sensor Installation Guide*.

Pairing the vSensors to the Brain

To be able to communicate with the Brain, a vSensor must be paired with the Brain. Pairing can be automated or performed manually.

To automate pairing, log onto the Brain and enable the **Settings > System > Sensors > Autopairing** option. Auto pairing allows the Brain and vSensor to pair automatically, without any user intervention.

To instead manually pair an individual vSensor to the Brain, Login to the vCenter and use the vSensor CLI to perform the pairing.

Enabling Auto Pairing

To enable auto pairing and allow the vSensors (or physical Sensors) to pair with the Brain in an automated fashion, complete the following steps.

1. Log in to the Vectra Brain and navigate to the **Settings > System** page.
2. In the **Sensors** pane, select the Edit icon.
3. The Sensors pane will expand and give the option to enable auto pairing. Enable auto pairing and click **Save**.

NOTE: It is recommended to disable the auto pairing setting once the deployed vSensors have been successfully paired to the Brain.

4. The Sensors, upon bootup, will try to connect with the Brain and pair. With this setting enabled, pairing occurs automatically.

Manually Pairing a vSensor to the Brain

To pair vSensors manually:

1. Launch the vSensor from vCenter and gain access to the CLI through the remote console or via SSH. The default username is **vectra**. The default password is **youshouldchangethis**.
2. Once logged in, set the Brain hostname (or IP address, if a hostname has not been set up) using the following command:

   $ set brain <hostname | IP>

NOTE: The IP address will already be configured, as this occurs when the OVA is created from the Brain. In fact, this vSensor cannot be paired with any other Brain than the Brain from which the OVA was downloaded, as that OVA includes keying material specific to the source Brain.

3. Login to the Brain and navigate to **Manage > Sensors**.
4. If auto pairing was enabled, using the previous instructions, the Sensor should show "paired". If the previous section was skipped, the vSensor should be present with a status of "Available". Click on this icon to begin pairing.
5. Within a few minutes, the vSensor status should change from "Pairing" to "Paired". Once this state change occurs, the vSensor is paired and ready to use.
6. To validate pairing and ensure that everything is properly configured, navigate to **Manage > Physical Hosts** and check the vSensor status in the right column. It should be present with a green check.

Pin the vSensor to its Hypervisor

There are many services, such as VMware's DRS and Turbonomic's VMTurbo, that provide dynamic resource management within a virtualized infrastructure. Their goal is to dynamically balance resources within an environment based on resource availability and requirements. For Vectra, this can lead to vSensors being moved from one hypervisor to another, leaving the first hypervisor unmonitored and thus unprotected.

To prevent this from occurring, it is important to create affinity rules that pin the vSensor VM to the hypervisor on which it was deployed, so that the vSensor does not move from one physical host to another automatically.

Changing vSensor CLI Password

Either of the following methods can be used to change the password on a vSensor:

Change the Password for all vSensors and Physical Sensors

Set the same single password across all existing vSensors and physical Sensors, and any new vSensors or Physical Sensors that are deployed.

1. From the Vectra UI of the Vectra Brain, navigate to **Settings > System**.
2. Click on the **Edit** button in the Sensor pane.
3. In the Sensor pane, type the new password in the **Sensor Password** field and click **Save**. This will propagate the new password to all of the vSensors and physical Sensors paired to this Brain.

Any future vSensor or physical Sensor that is created or deployed and is paired to this Brain will inherit the same password.

Change the Password on the Individual vSensor

Login to the vSensor and use the following command in that vSensor' CLI to change the password:

   $ set password

Appendix: Data Center Deployment Worksheet Hard Copy

The following worksheets provide a place to write down the Vectra deployment values for your DC network. Depending on your preference, either print these pages and write on the hardcopy, or download an Excel spreadsheet version of the worksheet. (Please contact Vectra Networks.)

Brain Settings

To plan for Brain insertion, fill out this worksheet (Table 13).

Table 13: Brain settings worksheet

Brain Setting	Description	Values in My Data Center
Brain General Settings General settings for each ESX/ESXi host, virtual switch, and port group where a vSensor will be installed.		
Hostname	DNS hostname of the Brain	
IP address and mask **Default gateway**	Unicast IP address and network mask to assign to the vSensor, and IP address of gateway router for reaching other subnets Enter these and other IP addresses in NetworkIP/CIDR format (example: *10.10.10.9/24*).	IP address: Network mask: Default gateway IP address:
Domain Name System (DNS) servers	IP addresses of one or more DNS servers that the Brain will use for resolving hostnames into IP addresses, in order to send traffic to the devices with those hostnames	DNS server 1 IP: DNS server 2 IP: DNS server 3 IP:
Network Time Protocol (NTP) servers	IP addresses of one or more servers to use as the source of the Brain's system date and time Enter these and other IP addresses in NetworkIP/CIDR format (example: *20.1.0.0/16*).	NTP server 1 IP: NTP server 2 IP: NTP server 3 IP: NTP server 4 IP:
Public IP subnets in DC network	If applicable, publically routable, global IP addresses (subnets or individual addresses) used in your DC network These will be classified by the Vectra System as "internal" for the sake of detection models that look for in-to-in or in-to-out connections. Enter in	
Remote access with Vectra Networks	Sharing of network metadata with Vectra Networks (On or Off) • **On**: Brain sends metadata from Sensors and vSensors to Vectra Networks. • **Off**: Brain does not send metadata to Vectra Networks. Set via Settings > Vectra Cloud > Share Metadata with Vectra	
	Remote access to Brain by Vectra Networks support staff (On or Off) • **On**: Brain allows remote access from Vectra Networks support staff • **Off**: Brain does not allow remote access Set via Settings > Vectra Cloud > Access for Remote Support	

vSphere API Access Settings

To plan for Brain access to the vSphere API, fill out this worksheet (Table 14).

Table 14: vSphere API access worksheet

vSphere Setting	Description	Values in My Data Center
VMware vSphere API Settings Settings required for the Brain to be able to learn device information from vSphere.		
vCenter server	Hostname or IP address of VMware vCenter	
Port number	TCP port to which the Brain should send API requests	
User ID	Username for the Brain to use when logging into vSphere	
Password	Password for the Brain to use when logging into vSphere	
Firewall	Hostname to IP address of device the Brain should use as its firewall Set this to the vCenter server's hostname or IP address	

vSensor Settings

To plan for vSensor insertion, fill out these worksheets (Table 15).

Table 15: vSensor Insertion Settings

vSensor Setting	Description	Values in My Data Center
vSensor General Settings General settings for each ESX/ESXi host, virtual switch, and port group where a vSensor will be installed.		
Resource requirements on host	Minimum system requirements for ESX/ESXi host	• CPU: 4 vCores • Memory: 8 GB vRAM • Drive: 150 GB • VMware vSphere 5.0 or later • Virtual switch type: VSS or VDS
vSwitch name	Name of the virtual switch on the hypervisor where the vSensor will be installed If using VMware VDS, this will be the same for all hypervisors. If using VMware VSS, the name may be different for each hypervisor.	
Mgt port group	Name and VLAN of port group to which the vSensor's management interface (Mgt1) will be attached	Mgt port group name: Mgt VLAN number:
Capture port group	Name and VLANs of port group to which the vSensor's capture interface will be attached	Capture port group name: Capture VLAN numbers: • VSS: 4095 (all VLANs) • VDS: 0-4094 (for all VLANs). If capturing subset of VLANs only, enter VLAN numbers or ranges:

Table 15: vSensor Insertion Settings

vSensor Setting	Description	Values in My Data Center
Physical Host #1 Settings Values for physical host where first vSensor will be installed		
Physical host name	The physical host's name within vCenter	
VM name for vSensor on host	VM name of the vSensor on this host	
DNS hostname for vSensor	DNS hostname for the vSensor	
vSensor name	vSensor name in Vectra System This can be the same as the vSensor's VM name or DNS hostname.	
Datastore (on physical host for VM disk usage)	Datastore that will be used for the vSensor VM on the physical host	
DHCP setting	Setting (enabled or disabled) of Dynamic Host Configuration Protocol (DHCP) DHCP automatically assigns IP settings, including IP address, default gateway IP address, and IP addresses of DNS servers	Enabled or Disabled (Circle one.)
vSensor mgt interface IP address and default gateway	Unicast IP address and mask to assign to the vSensor, and default gateway to be used to reach other subnets	Mgt interface IP: Default gateway IP: (If DHCP is enabled, leave blank.)
Domain Name System (DNS) server IP addresses	IP addresses of one or more DNS servers that the vSensor will use for resolving hostnames into IP addresses, in order to send traffic to the devices with those hostnames	DNS server 1 IP: DNS server 2 IP: DNS server 3 IP: (If DHCP is enabled, leave blank.)
Pin vSensor	Pinning the vSensor to the hypervisor prevents VMware from moving the vSensor to another hypervisor, such as when running vMotion.	Yes. Pin the vSensor to the hypervisor.

vSensor Setting	Description	Values in My Data Center
Physical Host #2 Settings Values for physical host where second vSensor will be installed		
Physical host name	**See descriptions for physical host #1 (above)**	
VM name for vSensor on host		
DNS hostname for vSensor		
vSensor name		
Datastore (on physical host for VM disk usage)		
DHCP setting		Enabled or Disabled (Circle one.)
vSensor mgt interface IP address and default gateway		Mgt interface IP: Default gateway IP: (If DHCP is enabled, leave blank.)
Domain Name System (DNS) server IP addresses		DNS server 1 IP: DNS server 2 IP: DNS server 3 IP: (If DHCP is enabled, leave blank.)
Pin vSensor		Yes. Pin the vSensor to the hypervisor.

Table 15: vSensor Insertion Settings

vSensor Setting	Description	Values in My Data Center
Physical Host #3 Settings Values for physical host where second vSensor will be installed		
Physical host name	**See descriptions for physical host #1 (above)**	
VM name for vSensor on host		
DNS hostname for vSensor		
vSensor name		
Datastore (on physical host for VM disk usage)		
DHCP setting		Enabled or Disabled (Circle one.)
vSensor mgt interface IP address and default gateway		Mgt interface IP: Default gateway IP: (If DHCP is enabled, leave blank.)
DNS server IP addresses		DNS server 1 IP: DNS server 2 IP: DNS server 3 IP: (If DHCP is enabled, leave blank.)
Pin vSensor		Yes. Pin the vSensor to the hypervisor.

VECTRA®
Security that thinks®

Email info@vectra.ai **Phone** +1 408-326-2020
vectra.ai