

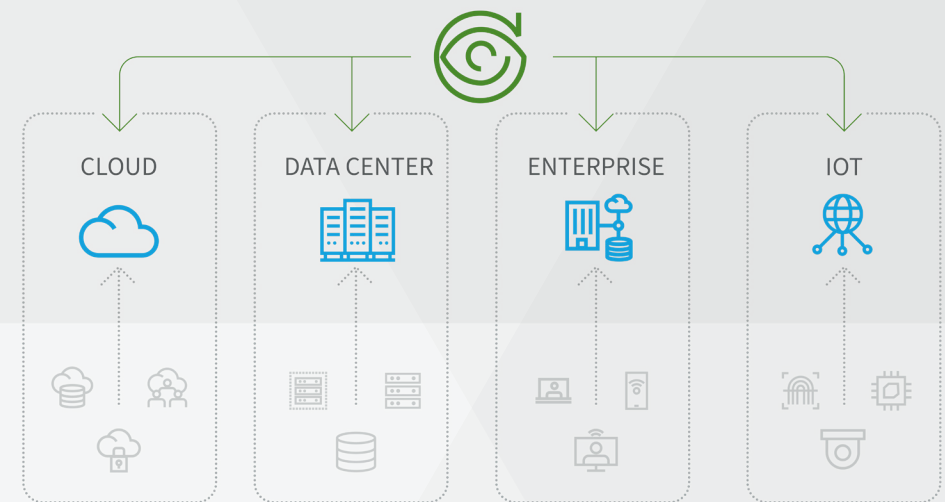
Appliance and Sensor specifications

Vectra AI can be deployed in Cloud, On-Prem, and Virtual environments to satisfy customer security requirements. For the Cloud, we support vSensors in AWS, Azure, and GCP. For On-Prem, we have both hardware sensors (S-Series, X-Series), and virtual sensors (VMware ESXi, Hyper-V, and KVM) to provide flexible deployment and coverage.

Included in your Vectra AI subscription, software updates with new threat detection algorithms are delivered to your system on a regular basis to ensure continuous protection from the latest advanced threats.

To extend the reach of Vectra AI, X-Series and S-series Sensors are easily deployed at remote sites or with access switches on internal network segments. They passively monitor network traffic, extract critical metadata and forward it to Vectra for analysis and threat detection. Sensors can be deployed out-of-band on a SPAN port, network TAP, or packet broker.

With Vectra, attackers have nowhere to hide



Virtual Sensors

Virtual Sensors, known as vSensors, run in VMware ESXi, Microsoft Hyper-V, and KVM, making it easy to extend Vectra threat detection coverage across the physical network and into virtualized data centers. VMware vSensors support both VSS and VDS VMware switch types to provide visibility into all traffic and detect threats that pass between workloads in the virtual environment. Vectra AI also integrates with VMware vCenter so it has authoritative, always up-to-date views of the virtual environment.

Physical Appliance Performance

Appliances	MGT Interfaces	Capture Ports	Paired Sensors	Sensor Mode Performance	Mixed Mode Performance	Brain Mode Performance
X3	2 x 10/100/1000 Copper Interfaces	2 x 1 GbE Copper 2 x 10 Gb SFP+	150	9 Gbps	8 Gbps	15 Gbps
X29	2 x 10/100/1000 Copper Interfaces or 1 x 10Gbps SFP+ Interface	2 x 1 GbE 2 x 10 Gb SFP+	150	15 Gbps	8 Gbps	20 Gbps
S1	2 x 1 GbE Copper	4 x 1 GbE Copper	N/A	1 Gbps	N/A	N/A
S11	2 x 10/100/1000 Copper Interfaces	2x1 GbE	N/A	2Gbps	N/A	N/A
S101	2 x SFP+ Interfaces (included in purchase, model selected by customer)	2 x 10G SFP+ and 2 configurable to: SFP28 (10/25 G) or QSFP (40 G) or QSFP28 (100 G)	N/A	50 Gbps	N/A	N/A
B101	2 x SFP+ Interfaces (included in purchase, model selected by customer)	N/A	500	N/A	N/A	75 Gbps

Notes regarding performance chart above:

- **Paired Sensors** – Refers to how many Sensors (physical, virtual, or cloud) an appliance can pair with.
- **N/A – Refers** to unsupported modes. For example, an S11 cannot be used as a Brain, therefore the only relevant metric for an S11 is the performance in Sensor mode.
- **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. Please contact Vectra AI to discuss further.
- 10Gbps management Interface supported in 6.19 or later. Consumes one of the 10Gbps traffic interfaces to operate as a management interface
- Console access to Vectra appliances detailed in: [Support](#)

Vectra Match Performance

Appliances	Mode	Match Throughput (Detect and Match)
S1	Sensor	400 Mbps
S2	Sensor	600 Mbps
S11	Sensor	1.2 Gbps
S101	Sensor	33 Gbps
X3	Sensor	3 Gbps
X3	Mixed	1 Gbps
X29	Sensor	9 Gbps
X29	Mixed	4.6 Gbps
X80	Sensor	11 Gbps
2 core vSensors (VMware, Hyper-V, KVM, Nutanix)	Sensor	250 Mbps
4 core vSensors (VMware, Hyper-V, KVM, Nutanix)	Sensor	500 Mbps
8 core vSensors (VMware, Hyper-V, KVM, Nutanix)	Sensor	1 Gbps
16 core vSensors (VMware, Hyper-V, KVM, Nutanix)	Sensor	2.5 Gbps
32 core vSensor (VMware)	Sensor	10 Gbps
2 core vSensors (AWS, Azure, GCP)	Sensor	500 Mbps
4 core vSensors (AWS, Azure, GCP)	Sensor	1 Gbps
8 core vSensors (AWS)	Sensor	2 Gbps
16 core vSensors (AWS)	Sensor	4 Gbps
16 core vSensor (GCP)	Sensor	2.5 Gbps
32 core vSensor (GCP)	Sensor	5 Gbps

Physical Appliance Specifications

	X3 Appliance	X29/M29 Appliance	S1 Sensor	S11 Sensor	S101 Sensor	B101
Input Voltage	Dual modular power supplies; auto-sensing 100-240 VAC, 50-60 Hz	Dual modular power supplies; auto-sensing 100-240 VAC, 50-60 Hz	Single external power supply, auto-sensing 100-240VAC, 50-60 Hz	Single power supply, auto-sensing 100-240VAC, 50-60 Hz	Dual modular power supplies; auto-sensing 100-240 VAC, 50-60 Hz	Dual modular power supplies; auto sensing 100-240 VAC, 50-60 Hz
Power	375 watts	550 watts	45W (Max), 35W (Normal)	450W (Max), 200 Watts (Normal)	1600 watts	1600 watts
Current	3.4 amps at 120 vac 1.7 amps at 240 vac	7.4 A at 120 VAC, 3.7 A at 240 VAC	2.0 A to 120 VAC, 1.0 A at 240 VAC	3.5 A to 6.5A (Max), 2 A (Normal) at 110V	10 A	8.7 amps
Dimensions	42.8 mm (1.685 in) H x 482 mm (18.976 in) W x 662.19 mm (26.070 in) D	45 mm (1.75 in) H x 432 mm (17 in) W x 660 mm (26 in) D	52 mm (2.04 in) H x 208 mm (8.18 in) W x 200 mm (7.87 in) D	42.8 mm (1.69 in) H x 434 mm (17.1 in) W x 535mm (22.6 in) D	42.8 mm (1.69 in) H x 482.0 mm (18.98 in) W x 808.5 mm (31.8 in) D	42.8 mm (1.69 in) H x 482 mm (18.98 in) W x 808.5mm (31.8 in) D
Weight	16.6 kg (36.6 lb)	12 kg (27 lbs)	1.4 kg without PSU	12.2 kg, 26.9 lb	21.9kg (48.3 lbs)	21.9 kg (48.3 lbs)
Environment	Operating temperature: <ul style="list-style-type: none"> • 0° to 35° C (50° to 95° F) Non-operating temperature: <ul style="list-style-type: none"> • -40° to 65° C (-40° to 149° F) Airflow: <ul style="list-style-type: none"> • Front to back 	Operating temperature: <ul style="list-style-type: none"> • 0° to 35° C (32° to 95° F) Non-operating temperature: <ul style="list-style-type: none"> • 0° to 50° C (32° to 122° F) Airflow: <ul style="list-style-type: none"> • Front to back 	Operating temperature: <ul style="list-style-type: none"> • 0°C-40°C (32°F-104°F) Non-operating temperature: <ul style="list-style-type: none"> • -40°C to 70°C (-40°F to 158°F) Airflow: <ul style="list-style-type: none"> • In bottom, out sides and back 	Operating temperature: <ul style="list-style-type: none"> • 10° to 35° C (50° to 95° F) Non-operating temperature: <ul style="list-style-type: none"> • -40° to 65° C (-40° to 149° F) Airflow: <ul style="list-style-type: none"> • Front to back 	Operating temperature: <ul style="list-style-type: none"> • 10° to 35° C (50° to 95° F) Non-operating temperature: <ul style="list-style-type: none"> • -40° to 65° C (-40° to 149° F) Airflow: <ul style="list-style-type: none"> • Front to back 	Operating temperature: <ul style="list-style-type: none"> • 10° to 35° C (50° to 95° F) Non-operating temperature: <ul style="list-style-type: none"> • -40° to 65° C (-40° to 149° F) Airflow: <ul style="list-style-type: none"> • Front to back

Virtual Brains – Customer Premise Hypervisor Deployment

Hypervisor	VM Type	Cores	Memory	Storage (OS, Data) in GB	Paired Sensors	Performance
VMware	vSphere 6.5 or later	8	64 GB	128,512	15	2 Gbps
VMware	vSphere 6.5 or later	16	128 GB	128,512	25	4 Gbps
VMware	vSphere 6.5 or later	32	256 GB	128,512	100	10 Gbps

Virtual Sensors - Customer Premise Hypervisor Deployment

Hypervisor	VM Type	Cores	Memory	Storage	Performance
VMware	vSphere 6.5 or later	2	8 GB	100 GB	500 Mbps
VMware	vSphere 6.5 or later	4	8 GB	150 GB	1 Gbps
VMware	vSphere 6.5 or later	8	16 GB	150 GB	2 Gbps
VMware	vSphere 6.5 or later	16	64 GB	600 GB*	5 Gbps
VMware	vSphere 6.5 or later	32	114 GB	830 GB	20 Gbps
Hyper-V	Windows Server 2016 w/ HW v8 or higher	2	8 GB	100 GB	500 Mbps
Hyper-V	Windows Server 2016 w/ HW v8 or higher	4	8 GB	150 GB	1 Gbps
Hyper-V	Windows Server 2016 w/ HW v8 or higher	8	16 GB	150 GB	2 Gbps
Hyper-V	Windows Server 2016 w/ HW v8 or higher	16	64 GB	500 GB	5 Gbp
KVM	Standard PC (Q35 + ICH9, 2009)	2	8 GB	100 GB	500 Mbps
KVM	Standard PC (Q35 + ICH9, 2009)	4	8 GB	150 GB	1 Gbps
KVM	Standard PC (Q35 + ICH9, 2009)	8	16 GB	150 GB	2 Gbps
KVM	Standard PC (Q35 + ICH9, 2009)	16	64 GB	500 GB	5 Gbps
Nutanix	AOS Version: 5.20.3.5 or later AHV Version 2021105.2267 or later	2	8 GB	100 GB	500 Mbps
Nutanix	AOS Version: 5.20.3.5 or later AHV Version 2021105.2267 or later	4	8 GB	150 GB	1 Gbps
Nutanix	AOS Version: 5.20.3.5 or later AHV Version 2021105.2267 or later	8	16 GB	150 GB	2 Gbps
Nutanix	AOS Version: 5.20.3.5 or later AHV Version 2021105.2267 or later	16	64 GB	500 GB	5 Gbps

Cloud - IaaS Deployment

Brain Deployment

Cloud	VM Type	Cores	Memory	Storage (OS, Data, Data, Data) in GB	Paired Sensors	Tracked Hosts	Performance
AWS	r5d.2xlarge	8	64 GB	256, 64, 128, 256	15	50,000	2 Gbps
AWS	r5d.4xlarge	16	128 GB	256, 64, 128, 256	25	50,000	5 Gbps
AWS	r5d.8xlarge	32	256 GB	256, 64, 128, 256	100	150,000	15 Gbps
AWS	r5d.16xlarge	64	512 GB	256 ² , 64, 512, 512	500	500,000	50 Gbps
Azure	Standard_E16s_v3	16	128 GB	256, 64, 128, 256	25	50,000	5 Gbps
Azure	Standard_E32s_v3	32	256 GB	256, 64, 128, 256	100	150,000	15 Gbps
GCP	n2-highmem-96	96	768 GB	4 TB (single partition)	100	500,000	50 Gbps

vSensor Deployment

Cloud	VM Type	Cores	Memory	Storage (OS, Data) in GB	Performance
AWS	r5(n).large *	2	16 GB	50, 128	1 Gbps
AWS	r5(n).large *	4	32 GB	50, 128	2 Gbps
AWS	r5(n).2xlarge *	8	64 GB	50, 512	4 Gbps
AWS	r5(n).4xlarge *	16	128 GB	50, 512	8 Gbps
AWS	c5n.18xlarge *	72	192 GB	50, 128 (No PCAP capability) *	Up to 10 Gbps *
Azure	Standard_DS11_v2	2	14 GB	50, 128	1 Gbps
Azure	Standard_DS3_v2	4	14 GB	50, 128	2 Gbps
GCP	e2-standard-2	2	8 GB	50, 128	1 Gbps
GCP	e2-standard-4	4	16 GB	50, 128	2 Gbps
GCP	e2-standard-16	16	64 GB	50, 128	5 Gbps
GCP	e2-standard-32	32	128 GB	50, 128	10 Gbps

- AWS vSensor configurations include both “n” and non “n” r5 instance types.
 - Networking performance is quoted as “up to 10Gbps” on the r5 instances by AWS and can be influenced by neighboring instances allocated to the same physical hardware in AWS.
 - Networking performance is quoted as “up to 25Gbps” on the r5n instances by AWS. These instances are still shared with neighbors but are optimized by AWS to have higher overall network throughput.
 - Customers can work with AWS to utilize dedicated instances and distribute instances to provide the required networking throughput to their vSensor instances on that dedicated hardware.
- The c5n.18x large vSensor instance type does not have a rolling capture buffer and can therefore not support PCAP generation for Detections that originate from traffic that is processed by those instances.
- Due to variability in customer cloud network configurations and how mirroring may be configured, it is not possible to guarantee performance on any instance with more than 2 cores (numbers are approximate and based on even distribution of packets across threads). Please contact Vectra to discuss further.
- If using port mirroring via VXLAN, please contact Vectra AI to discuss performance.

For more information please contact us at info@vectra.ai.

Email info@vectra.ai | vectra.ai

© 2024 Vectra AI, Inc. All rights reserved. Vectra, the Vectra AI logo, and Security that thinks are registered trademarks and the Vectra Threat Labs, the Threat Certainty Index, and Attack Signal Intelligence are trademarks of Vectra AI. Other brand, product and service names are trademarks, registered trademarks or service marks of their respective holders. Version: **111824**