

Enhancing Security Operations

Sophos NDR

Introduction:

In today's ever-changing threat landscape, organizations must adopt a proactive approach to identify and respond to potential cyber-attacks. Network Detection and Response (NDR) technology plays a critical role in this strategy.

NDR technology leverages deep learning analytics, traditional rule-based matching, and risk-based flow statistics to analyze raw network traffic and identify suspicious and potentially malicious activities on the network. This enables security teams to take proactive measures to prevent cyber-attacks and minimize their impact.

However, the high rate of false positives is a common challenge associated with NDR technology. Sophos NDR addresses this limitation by utilizing patented clustering and scoring technology that combines evidence from multiple threat detection engines.

While NDR technology has been available since the 1990s, the complexity and accuracy vary between vendors. It is crucial for organizations to consider a robust NDR solution like Sophos NDR, which provides advanced levels of threat detection and supports convictions while minimizing false positives. In this white paper, we will delve into the features and benefits of Sophos NDR and explain why it should be a fundamental component of any organization's security operations.

- Introduction: 1
- Evolution of Network Security Monitoring: A Timeline of NDR Technology 3
- Sophos NDR: Advanced Network Monitoring for Modern Threats 4
- NDR Sensor conceptual architecture 5
- Network Packet Processing (NPP) 5
 - NPP: Packet Header Data 6
 - NPP: Application Layer Data 7

Sophos NDR Detection Engines 8

IDS – Intrusion Detection System Engine 9

 Miscellaneous Activity 9

 Policy Violation 9

 Bad Unknown 9

 Malware Download 9

 Trojan Activity 9

 TLS Blacklist 9

SRA – Session Risk Analytics Engine 10

EPA - Encrypted Payload Analytics Engine and Machine Learning [ML] 11

DGA – Domain Generation Algorithm Engine 12

DDE – Data Detection Engine 13

CSS – Clustering and Severity Scoring 14

APPENDIXES 15

 APPENDIX A: SRA Flow Risks 15

 APPENDIX B: NPP Protocols 17

Evolution of Network Security Monitoring: A Timeline of NDR Technology

Sophos NDR is a critical component of modern security operations, but the history of NDR dates to the 1990s when network-based intrusion detection systems (NIDS) first emerged. Early NIDS systems focused on identifying and blocking network-based attacks, but they lacked the ability to correlate multiple events or detect advanced threats that spanned multiple systems.

In the early 2000s, NDR technology evolved to address these limitations. Instead of simply identifying individual network-based attacks, NDR solutions began analyzing network traffic and correlating events across multiple systems to identify advanced threats. Sophos NDR is a leading NDR solution that utilizes deep learning analytics, traditional rule-based matching, and risk-based flow statistics to identify suspicious and potentially malicious activities on the network.

Over time, NDR technology has become more sophisticated, providing near real-time visibility into network activity, and integrating seamlessly with other security solutions. The following timeline table outlines the key milestones in the evolution of NDR technology:

Year	Milestone
1980s	Network security products begin to emerge, with firewall technology becoming widely adopted
1990s	Network-based intrusion detection systems (NIDS) first emerge, marking the beginning of network security monitoring
2000s	Network Detection and Response (NDR) technology evolves to analyze network traffic and correlate events across multiple systems
2010s	Advanced machine learning algorithms are integrated into NDR solutions to identify complex threats and reduce false positives
2016	The Mirai botnet, which leverages IoT devices, launches one of the largest distributed denial of service (DDoS) attacks in history, underscoring the need for advanced network security
2019	Gartner introduces the term Network Detection and Response (NDR) to replace the previous term Network Traffic Analysis (NTA)
2020s	NDR solutions offer real-time visibility and flexible deployment options, enabling organizations to deploy them in any environment

NDR solutions like Sophos NDR enable organizations to effectively detect and respond to advanced threats.

Sophos NDR: Advanced Network Monitoring for Modern Threats

Sophos NDR is an advanced network monitoring solution designed to address the complex and evolving threat landscape.

Unlike traditional NDR solutions, Sophos NDR combines multiple proprietary detection engines with deep learning analytics, resulting in real-time and actionable intelligence on a wide range of network threats.

Sophos NDR's proprietary detection engines classify network traffic based on over 330 protocols, 50 flow risks, and thousands of indicators of compromise (IOCs). These engines also incorporate predictions from multiple deep learning models, providing an unprecedented level of threat detection accuracy while minimizing false positives.

Key advantages of Sophos NDR:

Traditional NDR	Sophos NDR	Improvement
Limited protocol coverage	Over 330 network protocols	Sophos NDR classifies traffic using over 330 protocols, allowing for a more comprehensive view of network traffic, which is crucial in identifying new and emerging threats. See Appendix B for a full list of Protocols.
Basic IOCs	Thousands of IOCs	Sophos NDR utilizes thousands of IOCs to detect indicators of compromise, resulting in a higher level of threat detection accuracy.
Minimal flow risk identification	50 flow risks	Sophos NDR incorporates 50 flow risks in its proprietary detection engines, allowing for the detection of more complex threats that may go undetected by other NDR solutions. See Appendix A for a full list of Flow Risks.
Rule-based matching	Deep learning analytics	Sophos NDR utilizes deep learning analytics to provide an unprecedented level of threat detection accuracy while minimizing false positives.
High false positive rates	Patented clustering and scoring technology	Sophos NDR uses patented clustering and scoring technology to reduce false positives, providing actionable intelligence on a wide variety of network threats.

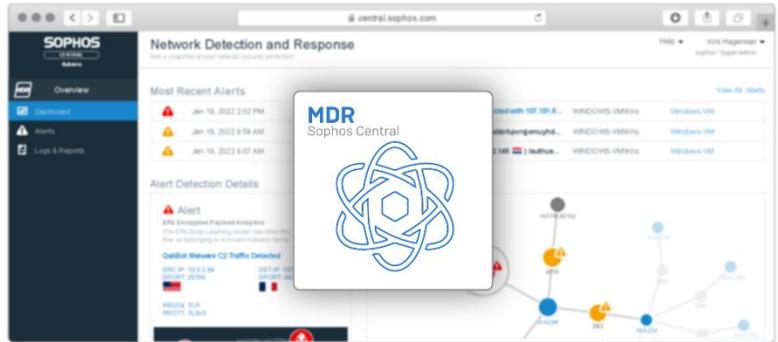
These improvements are particularly relevant to NDR because they enable Sophos NDR to accurately identify and respond to network threats without generating an excessive number of false positives. Sophos NDR's focus on speed, accuracy, and its ability to handle encrypted traffic without having to perform decryption of the traffic makes it an essential component of any comprehensive security strategy.

Sophos NDR offers organizations an advanced network monitoring solution designed to effectively detect and respond to the ever-evolving threat landscape. By combining multiple proprietary detection engines with deep learning analytics, Sophos NDR provides actionable intelligence that is both accurate and relevant to today's modern threats.

NDR Sensor conceptual architecture

The Sophos NDR solution deploys as a passive traffic monitor listening on a SPAN/Mirror port and does not add any latency to the network traffic or create a point of failure in the network if it becomes overloaded or is offline.

As the data flows into the Sensor meta data is collected and the network flow details are sent to a series of detection engines before being clustered and scored. Results of the clustered network flows are sent to the Sophos data lake and presented in Central in the detections dashboard.



Network Packet Processing (NPP)

Effective network flow metadata collection is critical to the success of NDR solutions. This process involves aggregating the network packets into a single communication or flow and collecting metadata from each network packet using Deep Packet Inspection (DPI). The collected metadata is then enhanced with geolocation information and other metrics, such as unpopular destinations, periodicity, and packet dynamics. The final step is detecting risk indicators such as bad TLS information, unidirectional traffic, large DNS packets, and more.

To better understand the packet header and application layer data that is collected during this phase, the following tables outline examples of what can be determined from each category, and why they are important for threat hunting.

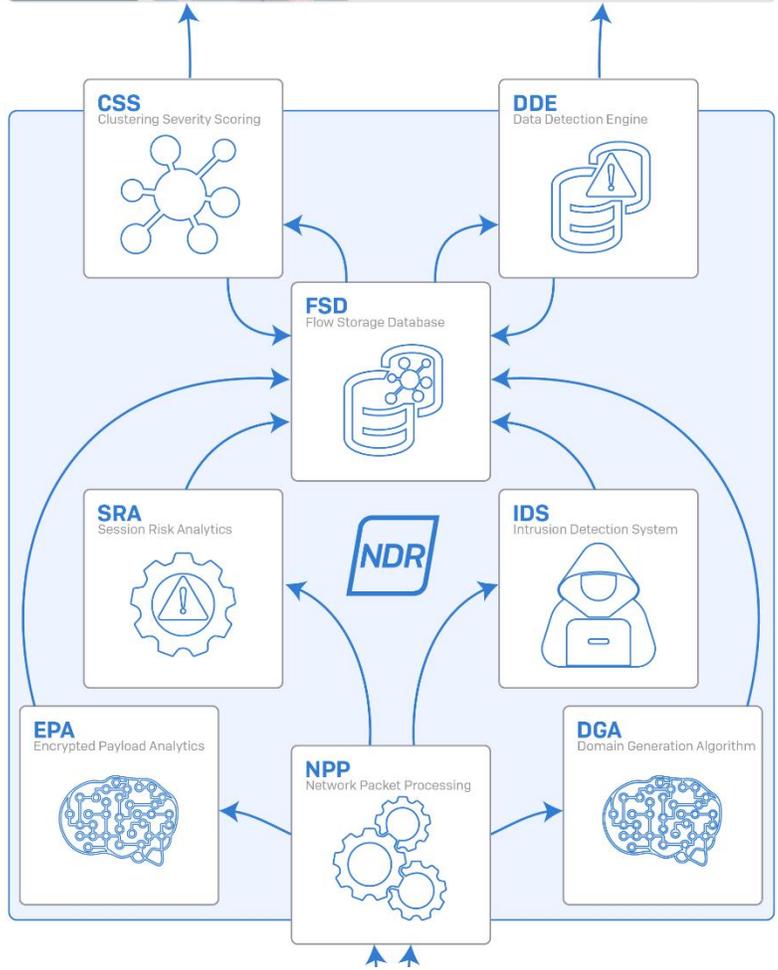


Figure 1: Sophos NDR Architecture Diagram



NPP: Packet Header Data

Packet header data provides information about the network communication such as source and destination addresses, transport protocol, duration, and size. It helps NDR solutions to identify the source of the communication and the potential threat that it might pose. Some of the information that can be determined from packet header data includes:

Packet Header Data	Description	Importance to NDR Threat Hunting
Source IP	The IP address of the sender	Identifies the source of the communication, which can be used to track suspicious activity or locate infected hosts
Source MAC address	The Media Access Control (MAC) address of the sender	The MAC address can be used to identify the physical device associated with network traffic and used with other sensor information to correlate alerts from multiple sensors to a specific device.
Source port	The port used by the sender for the communication	Helps identify the specific service or application associated with the communication, which can be used to detect suspicious or unauthorized activity
Destination IP	The IP address of the receiver	Helps identify the target of the communication, which can be used to identify external threat sources
Destination MAC address	The MAC address of the receiver	Helps identify the physical device associated with the communication, which can be used to track suspicious activity or locate infected hosts
Destination port	The port used by the receiver for the communication	Helps identify the specific service or application associated with the communication, which can be used to detect suspicious or unauthorized activity
TCP flags	Indicates the status of a TCP connection, such as SYN, ACK, FIN, RST, etc.	Can be used to detect suspicious network activity or attacks, such as port scanning or denial-of-service attacks
Duration of communication	The length of time the communication lasted	Helps identify suspicious activity, such as connections that last longer than expected, unusually short-lived connections and periodic communication related to beaconing
Bytes received	The amount of data received during the communication	Can be used to detect suspicious activity or attacks, such as data exfiltration or malware downloads
Layer 3 (network) and Layer 4 (transport) protocols	The protocols used for the communication, such as IP, OSPF, ICMP, TCP, UDP	Helps identify the type of traffic and associated services, which can be used to detect suspicious or unauthorized activity
Network VLAN (Virtual Local Area Networks) ID	The VLAN tag associated with the communication	Helps identify the specific network segment associated with the communication

NPP: Application Layer Data

Application layer data provides insight into the content of the network communication, allowing NDR solutions to identify potential threats that might be hiding within. It provides information about the applications and services used in network communication and helps identify clear-text usernames and passwords. Some examples of the information that can be determined from application layer data include:

Application Layer Data	Explanation	Importance for NDR Threat Hunting
Application layer protocol	The protocol being used at the application layer, such as HTTP, TLS, or SMB (Server Message Block)	Knowing the application layer protocol being used can help identify potentially malicious traffic or abnormal behavior for that protocol
Source and destination hostnames	The hostnames associated with the source and destination IP addresses, resolved through DNS or other means	This can help identify potentially malicious traffic or abnormal behavior associated with hosts or domains
HTTP Content Type	The type of content being transferred over HTTP, such as text, image, or video	This can help identify potentially malicious traffic or abnormal behavior associated with types of content
Response Code	The HTTP status code returned by the server in response to an HTTP request	This can help identify potentially malicious traffic or abnormal behavior associated with particular response codes, such as 404 Not Found or 500 Internal Server Error
URL	The full URL being requested or accessed	This can help identify potentially malicious traffic or abnormal behavior associated with URLs or domains
User Agent	The software agent used by the client to make the request, such as a web browser or mobile app	This can help identify potentially malicious traffic or abnormal behavior associated with user agents, such as those associated with known malicious software
Clear-text Usernames and Passwords	Any usernames or passwords transmitted in clear text, such as in an unencrypted HTTP request	This can help identify potential security issues or unauthorized access attempts
TLS Certificate Information	Information about the TLS certificate used in a secure connection, including JA3 hashes	This can help identify potentially malicious or spoofed certificates, as well as provide insight into the nature of encrypted traffic
SSH client and server HASH	A fingerprinting method for identifying SSH clients and servers	This can help identify potentially malicious SSH traffic or detect unauthorized SSH access attempts
CAPWAP Encapsulation	The Control and Provisioning of Wireless Access Points (CAPWAP) protocol, used for managing wireless access points	This can help identify potentially malicious or unauthorized wireless access attempts or unusual wireless network activity

In conclusion, the collection of network flow metadata is an essential component of NDR solutions, providing insight into the network communication and enabling the detection of potential threats. The packet header and application layer data collected provides valuable information that helps to identify the source, type, and potential risk of the communication. By leveraging this information, NDR solutions can effectively detect and respond to network threats.

Sophos NDR Detection Engines

Sophos NDR incorporates five distinct detection engines to provide comprehensive threat detection capabilities. These detection engines work together to identify and correlate various indicators of compromise, which are then scored and presented as actionable threat intelligence in Sophos Central for customers and analysts.

For increased performance, the machine learning detection engines (EPA – Encrypted packet Analytics, and DGA – Domain Generation Algorithm) are not run on all network flows and are instead triggered based on findings from the other detection engines. Enabling the detection engines to collaborate in the classification is critical to maintaining performance and reducing false positives.

The results of the detection engines are then fed into a clustering and severity scoring algorithm (CSS) to generate an overall threat score for presentation to the administrator as a detection in the Sophos Central Detections dashboard. The detection record contains the results of each engine's contribution.

IDS – Intrusion Detection System Engine

This proprietary IDS engine is a streamlined, more efficient engine with the capability to identify Indicators of Compromise (IOC)s in unencrypted traffic. Many security vendors continue to use overly robust content matching systems even with the loss of visibility, due to encryption.

Sophos NDR uses carefully selected threat intelligence to create IDS rules classified into six groups, based on the type of IOC. The following are these rule classifications and their descriptions:

Miscellaneous Activity

This rule classification has a Low Severity and is used for detecting network traffic not associated with the other classifications. Examples include traffic to public DNS servers, traffic to content delivery networks, or traffic to trusted cloud services. Identifying miscellaneous activity helps establish a baseline of normal network traffic and highlights any deviations from that baseline.

Policy Violation

This rule classification has a Low Severity and is used for detecting traffic that potentially violates a corporate policy. Examples include traffic to unauthorized websites or services, or traffic from unauthorized devices. Detecting policy violations helps organizations enforce their security policies and prevent unauthorized access or data exfiltration.

Bad Unknown

This rule classification has a Medium Severity and is used for identifying network communication with a potentially bad destination. This can include communication with a known malicious IP address or domain, or communication with a sinkhole domain used to redirect traffic to malicious infrastructure. Detecting bad unknown traffic can help identify compromised endpoints and prevent data exfiltration or further compromise.

Malware Download

This rule classification has a High Severity and is used to identify network communications with a known malware distribution source. This can include communication with a known command-and-control (C2) server used to download or distribute malware, or communication with a known malware distribution site. Detecting malware downloads helps organizations identify and isolate infected endpoints to prevent further spread of malware.

Trojan Activity

This rule classification has a High Severity and is used to identify network communications with a known malware C2 server. This can include communication with a C2 server used for remote control of a compromised endpoint, or communication with a C2 server used to exfiltrate data. Detecting Trojan activity helps organizations identify and isolate compromised endpoints and prevent data exfiltration or further compromise.

TLS Blacklist

This rule classification has a Critical Severity and is used to identify network communications with a known malicious actor based on TLS certification match. This can include communication with a known malicious domain using a compromised TLS certificate or communication with a known malicious domain that is not using a valid TLS certificate. Detecting TLS blacklisted traffic helps organizations prevent communication with known malicious infrastructure and protect against cyber-attacks.

SRA – Session Risk Analytics Engine

The SRA engine detects when network traffic deviates from documented protocol standards, which could indicate suspicious or risky network activity. This is important in a threat hunt because it helps identify non-standard behavior that may indicate an attack. When the SRA engine observes such activity, it adds information about the behavior to the flow metadata. These flow risks are not considered indicators of compromise on their own, but when combined with detections from other engines, they can help in identifying malicious activity.

The following is a list of general flow risks, which can be found across multiple protocols, and what they indicate:

Type	Flow Risk	Description
General	Possible Exploit	Indicates a possible exploit was detected, such as Log4J/Log4Shell. Important for detecting exploit activity and preventing/mitigating attacks.
General	Known Protocol on Non-Standard Port	Indicates a protocol is being used on a non-standard port, such as HTTP on TCP/8000 instead of the standard TCP/80. Important for detecting attackers who use non-standard ports to evade detection.
General	Risky ASN	Indicates network traffic was exchanged with a server belonging to an ASN (Autonomous System Number) that is considered risky. Important for identifying malicious hosts or networks.
General	Unidirectional Traffic	Indicates a session is only one direction, which could indicate C2 activity to a server that is no longer operating at the address. Important for identifying compromised hosts or C2 servers.
General	Desktop or File Sharing Session	Indicates the flow carries desktop or file sharing data, such as TeamViewer or AnyDesk. Important for detecting attackers who use these tools to remotely control a compromised host.
General	Unsafe Protocol	Indicates the protocol used is insecure and should not be used, such as Telnet instead of SSH. Important for detecting attackers who can intercept and read traffic sent over insecure protocols.
General	Clear Text Credentials	Indicates credentials were transmitted in clear text over a known protocol, such as FTP, HTTP, IMAP, POP3, or SMTP. Important for detecting attackers who can intercept and read clear text credentials.
General	Malformed Packet	Indicates a packet has an unexpected format, which could indicate a protocol error or a takeover of a valid protocol to carry another type of data. Important for detecting attacks that use packet manipulation or protocol misuse.
General	TCP Issues	Indicates issues were found in the TCP settings of the network session. Important for detecting attackers who use TCP issues in attacks to disrupt or evade detection.
General	Periodic Flow	Indicates the network session is repeating at a scheduled interval, which could indicate C2 activity from a Trojan or Botnet. Important for detecting attackers who use periodic communication to maintain control over compromised hosts.

For a complete list of Flow Risks, see Appendix A.

EPA - Encrypted Payload Analytics Engine and Machine Learning [ML]

Machine learning is increasingly being used in network detection and response (NDR) solutions to detect suspicious traffic on enterprise networks. NDR tools continuously analyze raw traffic and/or flow records, such as NetFlow, to build models that reflect normal network behavior, according to Gartner. Deep learning takes this approach further by allowing for the detection of patterns across multiple attributes, enabling detections without IOC-based threat intelligence.

Sophos has developed a specific solution called Encrypted Payload Analytics (EPA) to address the challenge of detecting threats in encrypted traffic using older technologies. Network flows consist of packets with header and payload data, and when inspecting an encrypted communication, only the payload data is encrypted, making it impossible to know the contents without decryption. EPA is a multi-class deep learning prediction model trained to detect patterns in network flows based on Sequence of Packet Length and Interarrival Time (SPLIT). These SPLIT attributes are simple to compute and are used to train a Convolutional Neural Network (CNN) for classification. Sophos NDR uses a patented process for normalizing, transforming, and presenting this data to the CNN for classification.

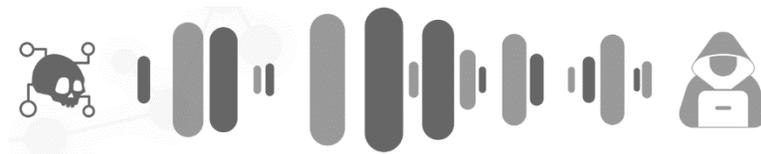


Figure 2: Sequence of Packet Length and Interarrival Time

By using live malware samples, the EPA model can identify malicious activity in real-time, including zero-day or unknown malware variants and C2 servers based on the patterns in the network flows between them. The EPA engine also enriches the flow metadata with the detected malware family and a confidence score to reduce the number of false positives. Overall, EPA enables organizations to detect and respond to encrypted threats that would have previously gone undetected. This approach is particularly useful when the endpoint devices cannot run a traditional endpoint protection product and when network communications should not be decrypted because of requirements to protect PII (Personal Identifying Information)

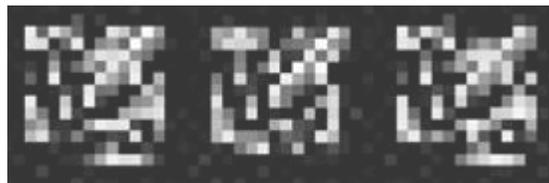


Figure 3: Cobalt Strike variant after processing as image for EPA CNN.

The Encrypted Payload Analytics (EPA) engine enhances the flow metadata by identifying the specific Malware Family (such as Bumblebee, Cobalt Strike, Emotet, Dridex, QakBot) and provides a Confidence Score ranging from 0-100. To reduce the number of false positives, the model also includes an "unknown" classification.

DGA – Domain Generation Algorithm Engine

Domain Generation Algorithms (DGA) are used by malicious actors to generate domain names that can be used for Command-and-Control (C2) purposes without being blacklisted. Using these algorithms, malware can generate a list of potential domain names that the C2 server might be hosted on. After numerous attempts, the algorithm will find a domain that exists and establish a connection.

husbbrkpvraqjomuyhdpd.com

Figure 4: Example DGA domain

Historically, DGA has been used in several high-profile attacks. For example, in the Conficker worm outbreak of 2008, DGA was used to generate a list of over 50,000 domain names each day that could be used as C2 servers. This made it extremely difficult for security researchers to shut down the worm's C2 network. Another example is the use of DGA in the Gameover Zeus malware, which was used to generate up to 1,000 domain names per day for C2 purposes. The Gameover Zeus botnet was responsible for stealing over \$100 million from victims around the world.

The Sophos NDR's DGA detection engine is crucial for identifying malicious activity in real-time. Sophos NDR's DGA detection engine is powered by a Deep Learning Long Short-Term Memory (LSTM) neural network that evaluates every domain name queried and accessed. It is important to note that not all DGA activity is malicious; many legitimate services use DGA regularly. Therefore, Sophos NDR does not generate an alert every time a DGA is detected. Instead, a confidence score (0-100) is added to the flow metadata and is used by the Clustering and Severity Scoring (CSS) engine to determine if the activity involving DGA detections is indeed malicious.

DDE – Data Detection Engine

The Data Detection Engine (DDE) is a component of Sophos NDR that runs on each sensor. It is a lightweight correlation engine that utilizes the onboard database storage of network flows and clusters of flows. The DDE performs scheduled data mining activities on this information to identify complex network threats such as enumeration activities. This information is then sent to Sophos Central and used to generate network information reports.

Additionally, the data collected by the DDE can be correlated with data from Sophos XDR (Extended Detection & Response) endpoint sensors to identify unmanaged assets on the network. This correlation happens in the Sophos Data Lake and provides a comprehensive view of the network, allowing administrators to identify potential security risks and take appropriate action. It is important to note that the DDE performs data mining activities on a set schedule and not in real-time.

CSS – Clustering and Severity Scoring

The Clustering and Severity Scoring (CSS) feature is an essential part of Sophos NDR's threat detection capabilities. During network sessions between clients and servers, the system observes a wide range of threat indicators. These indicators, when analyzed alone, may not accurately represent a problem or malicious activity. Therefore, Sophos NDR utilizes a patented process to cluster these indicators over time, providing a higher level of confidence in identifying threats.

The clustering process groups network flows based on basic network information, such as source and destination IP/Port and protocol information. By clustering multiple flows that have occurred over time the system is able to generate a more comprehensive view of suspect activity and to aggregate related network flows into a single detection event where the correlation between the flows aids in understanding of the suspect activity.

Once these clusters are created, they are scored based on the information gathered from each of the detection engines. The CSS algorithm evaluates all activities within a cluster to provide additional context, improving accuracy and reducing false positives.

The scoring system within CSS is based on various factors, including severity levels and threat indicators identified by the various detection engines. By combining this information, Sophos NDR assigns a score to each cluster, reflecting the potential risk posed by the network activity. This scoring system provides network administrators with valuable information on potential threats, allowing them to prioritize responses based on the severity of the risk.

APPENDIXES

APPENDIX A: SRA Flow Risks

Protocol	Flow Risk	Description
GENERAL	Possible Exploit	Possible exploit was detected (e.g., Log4J/Log4Shell)
GENERAL	Known Protocol on Non-Standard Port	Protocol is being used on a non-standard port (e.g., HTTP over TCP/8000, standard is TCP/80)
GENERAL	Risky ASN	Network session was exchanged with a server belonging to a risky ASN (Autonomous System Number)
GENERAL	Unidirectional Traffic	Session is only one direction. This could indicate C2 activity to a server that is no longer operating at the address.
GENERAL	Desktop or File Sharing Session	Flow carries desktop or file sharing data (e.g., TeamViewer, AnyDesk)
GENERAL	Unsafe Protocol	Protocol used is insecure and should not be used (e.g., Telnet vs SSH)
GENERAL	Clear Text Credentials	Credentials were transmitted in clear text over a known protocol (e.g., FTP, HTTP, IMAP, POP3, SMTP)
GENERAL	Malformed Packet	Network packet has an unexpected format. This could indicate a protocol error or a takeover of a valid protocol to carry another type of data
GENERAL	TCP Issues	Issues were found in the TCP settings of the network session
GENERAL	Anonymous Subscriber	Source IP address has been anonymized and cannot be used to identify the subscriber (e.g., flow generated by an iCloud-private-relay exit node)
GENERAL	Periodic Flow	Network session is repeating at a scheduled interval. This could indicate C2 activity from a Trojan or Botnet
TLS, HTTP, DNS	Suspicious DGA Domain	Domain name may be a DGA, which is used to generate domain names often used by malware
TLS, HTTP, DNS	Risky Domain	Network traffic occurred with a domain that is considered risky
TLS, HTTP, DNS	Invalid Characters	Decoded protocol contains characters not allowed in that protocol field (e.g., a DNS hostname must only contain a subset of all printable characters)
TLS, HTTP, DNS	Punycode IDN	Domain name was observed in IDN format. Punycode IDN domains could indicate a homograph phishing attack
HTTP, DNS	Error Code Detected	Error was detected in the protocol
DNS	Suspicious Traffic	Unexpected or obsolete DNS record type was observed
DNS	Large Packet	DNS over UDP packet exceeded the size limit of 512 bytes. This could indicate DNS Tunneling or Exfiltration
DNS	Fragmented	UDP DNS was fragmented. This could indicate DNS Tunneling or Exfiltration
SSH	Obsolete Client Version or Cipher	SSH client used an obsolete protocol version or insecure ciphers
SSH	Obsolete Server Version or Cipher	SSH server used an obsolete protocol version or insecure ciphers
SMB	Insecure Version	Insecure version of SMB was observed (e.g., SMBv1)
ICMP	Suspicious Entropy	Suspicious entropy was observed in ICMP packets. This could indicate data exfiltration over ICMP
TLS	Self-Signed Certificate	Self-signed certificate was used
TLS	Malicious SHA1 Certificate	Observed TLS certificate was found on a malicious certificate
TLS	Certificate Mismatch	TLS certificate does not match the hostname being accessed.
TLS	Missing SNI	The SNI of the accessed server is missing.
TLS	Suspicious ESNI Usage	Encrypted SNI was observed. This could indicate a domain fronting attack.
TLS	Not Carrying HTTPS	TLS flow was not being used to transport HTTPS.
TLS	Malicious JA3 Fingerprint	JA3 fingerprint was found on a malicious JA3 blacklist
TLS	Suspicious Extension	The domain name in the SNI extension was not printable.
TLS	Uncommon ALPN	Uncommon ALPN extension was observed in the TLS flow (Example: HTTP/1.1).
TLS	Certificate Expired	TLS certificate used in the flow is expired
TLS	Certificate About to Expire	TLS certificate used in the flow is about to expire.

TLS	Certificate Validity Too Long	TLS certificate used in the flow has a lifespan longer than 13 months.
TLS	Obsolete Version	TLS version is older than 1.1.
TLS	Weak Cipher	Unsafe TLS cipher was used in the setup of the flow.
TLS	Fatal Alert	TLS protocol had a fatal alert in the flow.
HTTP	Numeric IP Host	Webserver was accessed using its IP address instead of the hostname.
HTTP	Suspicious URL	The access URL is suspicious. (Example: <code>http://127.0.0.1/msadc/..%25c../..%25c../winnt/system32/cmd.exe.</code>).
HTTP	Suspicious Header	The HTTP header contains suspicious entries that aren't expected in the HTTP header. (Example: UUID, TLS version, OS name).
HTTP	Suspicious User Agent	The User Agent string contained suspicious characters or formatting. (Example: <code><?php something ?></code>).
HTTP	Suspicious Content	The HTTP flow carried content in an unexpected format. (Example: The HTTP header indicates that the context is text/html, but the content is not readable due to being binary data).
HTTP	Binary Application Transfer	A binary application is being downloaded or uploaded. Detected files include Windows binaries, Linux executables, Unix scripts and Android apps.
HTTP	URL Possible XSS	A possible XSS (Cross Site Scripting) attack was observed.
HTTP	URL Possible SQL Injection	A possible SQL Injection attack was observed.
HTTP	URL Possible RCE Injection	A possible RCE (Remote Code Execution) attack was observed.
HTTP	Crawler Bot	A crawler/bot/robot was detected.
HTTP	Obsolete Server	A network session with an obsolete Apache or Nginx server was detected.

APPENDIX B: NPP Protocols

1KXUN	GIT	MICROSOFT_365	SPOTIFY
ACCUWEATHER	GITHUB	MICROSOFT_AZURE	SSDP
ACTIVISION	GITLAB	MINING	SSH
ADS_ANALYTICS_TRACK	GMAIL	MODBUS	STARCRRAFT
ADULT_CONTENT	GNUTELLA	MONGODB	STEAM
AFP	GOOGLE	MPEGDASH	STUN
AJP	GOOGLE_CLASSROOM	MPEGTS	SYNCTHING
ALIBABA	GOOGLE_CLOUD	MQTT	SYSLOG
ALICLOUD	GOOGLE_DOCS	MS_ONE_DRIVE	TAILSCALE
AMAZON	GOOGLE_DRIVE	MS_OUTLOOK	TARGUS_GETDATA
AMAZON_ALEXA	GOOGLE_MAPS	MSSQL_TDS	TEAMSPEAK
AMAZON_AWS	GOOGLE_PLUS	MSTEAMS	TEAMVIEWER
AMAZON_VIDEO	GOOGLE_SERVICES	MUNIN	TELEGRAM
AMONG_US	GOTO	MYSQL	TELNET
AMQP	GTP	NATPMP	TENCENT
ANYDESK	GTP_C	NATS	TENCENTVIDEO
APPLE	GTP_PRIME	NEST_LOG_SINK	TEREDO
APPLE_ICLOUD	GTP_U	NETBIOS	TFTP
APPLE_ITUNES	GUILDWARS	NETFLIX	THREEMA
APPLE_PUSH	H323	NETFLOW	TIDAL
APPLE_SIRI	HALFLIFE2	NFS	TIKTOK
APPLESTORE	HANGOUT_DUO	NINTENDO	TINC
APPLETVPLUS	HBO	NOE	TIVOCONNECT
ARMAGETRON	HOTSPOT_SHIELD	NTOP	TLS
AVAST	HPVIRTGRP	NTP	TOCA_BOCA
AVAST_SECUREDNS	HSRP	OCS	TOR
BADOO	HTTP	OCSF	TPLINK_SHP
BGP	HTTP_CONNECT	OOKLA	TRUPHONE
BITTORRENT	HTTP_PROXY	OPENDNS	TUENTI
BJNP	HULU	OPENVPN	TUMBLR
BLOOMBERG	I3D	ORACLE	TUNEIN
CACHEFLY	IAX	PANDORA	TUNNELBEAR
CAPWAP	ICECAST	PASTEBIN	TUYA_LP
CASSANDRA	ICLOUD_PRIVATE_RELAY	PINTEREST	TVUPLAYER
CHECKMK	IEC60870	PLAYSTATION	TWITCH
CISCOVPN	IFLIX	PLAYSTORE	TWITTER
CITRIX	IHEARTRADIO	PLURALSIGHT	UBNTAC2
CLOUDFLARE	IMO	POSTGRES	UBUNTUONE
CLOUDFLARE_WARP	INSTAGRAM	PPSTREAM	ULTRASURF
CNN	IP_EGP	PPTP	USENET
COAP	IP_GRE	PSIPHON	VEVO
COLLECTD	IP_ICMP	QQ	VHUA

CORBA	IP_ICMPV6	QUIC	VIBER
CPHA	IP_IGMP	RADIUS	VIMEO
CRASHLYSTICS	IP_IP_IN_IP	RAKNET	VK
CROSSFIRE	IP_OSPF	RDP	VMWARE
CRYPNET	IP_PGM	REDDIT	VNC
CSGO	IP_PIM	REDIS	VUDU
CYBERSECURITY	IP_SCTP	RIOTGAMES	VXLAN
DAILYMOTION	IP_VRRP	RPC	WARCRAFT3
DATASAVR	IPP	RSH	WAZE
DAZN	IPSEC	RSYNC	WEBEX
DEEZER	IRC	RTCP	WEBSOCKET
DHCP	JABBER	RTMP	WECHAT
DHCPV6	KAKAOTALK	RTP	WHATSAPP
DIAMETER	KAKAOTALK_VOICE	RTSP	WHATSAPP_CALL
DIRECTV	KERBEROS	RX	WHATSAPP_FILES
DISCORD	KISMET	S7COMM	WHOIS_DAS
DISNEYPLUS	KONTIKI	SALESFORCE	WIKIPEDIA
DNP3	LASTFM	SAP	WINDOWS_UPDATE
DNS	LDAP	SD_RTN	WIREGUARD
DNSCRYPT	LIKEE	SFLOW	WORLD_OF_KUNG_FU
DOFUS	LINE	SHOWTIME	WORLDOFWARCRAFT
DOH_DOT	LINE_CALL	SIGNAL	WSD
DRDA	LINKEDIN	SIGNAL_VOIP	XBOX
DROPBOX	LISP	SINA	XDMCP
DTLS	LIVESTREAM	SIP	XIAOMI
EAQ	LLMNR	SIRIUSXMRADIO	YAHOO
EBAY	LOTUS_NOTES	SKINNY	YANDEX
EDGECAST	MAIL_IMAP	SKYPE_TEAMS	YANDEX_CLOUD
EDONKEY	MAIL_IMAPS	SKYPE_TEAMS_CALL	YANDEX_DIRECT
ELASTICSEARCH	MAIL_POP	SLACK	YANDEX_DISK
ETHERNET_IP	MAIL_POPS	SMBV1	YANDEX_MAIL
FACEBOOK	MAIL_SMTTP	SMBV23	YANDEX_MARKET
FACEBOOK_VOIP	MAIL_SMTPS	SMPP	YANDEX_METRIKA
FASTCGI	MAPLESTORY	SNAPCHAT	YANDEX_MUSIC
FIX	MDNS	SNAPCHAT_CALL	YOUTUBE
FORTICLIENT	MEGACO	SNMP	YOUTUBE_UPLOAD
FTP_CONTROL	MEMCACHED	SOAP	Z3950
FTP_DATA	MERAKI_CLOUD	SOCKS	ZABBIX
FTPS	MESSENGER	SOFTETHER	ZATTOO
FUZE	MGCP	SOMEIP	ZMQ
GENSHIN_IMPACT	MICROSOFT	SOUNDCLOUD	ZOOM

Statements contained in this document are based on publicly available information as of March 30, 2023. This document has been prepared by Sophos and not the other listed vendors. The features or characteristics of the products under comparison, which may directly impact the accuracy or validity of this comparison, are subject to change. The information contained in this comparison is intended to provide a broad understanding and knowledge of factual information about various products and may not be exhaustive. Anyone using the document should make their own purchasing decision based on their requirements and should also research original sources of information and not rely only on this comparison while selecting a product. Sophos makes no warranty as to the reliability, accuracy, usefulness, or completeness of this document. The information in this document is provided "as is" and without warranties of any kind, either expressed or implied. Sophos retains the right to modify or withdraw this document at any time.

United Kingdom and Worldwide
Sales

Tel: +44 (0)8447 671131

Email: sales@sophos.com

North American Sales

Toll Free: 1-866-866-2802

Email: na-sales@sophos.com

Australia and New Zealand
Sales

Tel: +61 2 9409 9100

Email: sales@sophos.com.au

Asia Sales

Tel: +65 62244168

Email: salesasia@sophos.com

The logo for Sophos, consisting of the word "SOPHOS" in a bold, blue, sans-serif font.

© Copyright 2023. Sophos Ltd. All rights reserved.

Registered in England and Wales No. 2096520, The Pentagon, Abingdon Science Park, Abingdon, OX14 3YP, UK

Sophos is the registered trademark of Sophos Ltd. All other product and company names mentioned are trademarks or registered trademarks of their respective owners.