EXTRAHOP

WHITE PAPER

How ExtraHop RevealX™ Network Performance Monitoring Works

Behind the Curtain on ExtraHop's Approach to Transforming Network Data into Real-Time Insights

Take a peek behind the curtain to better understand the features that enable ExtraHop RevealX Network Performance Monitoring (NPM) to deliver comprehensive insights that help ensure consistent uptime and reliable performance across hybrid environments. In this white paper, you'll find information on: ExtraHop stream processing for full content analysis and line-rate decryption; machine learning and anomaly detectors; data indexing and storage; data visualization and exploration, including querying and live activity maps; and the more than 90 protocols that RevealX decodes.

With ExtraHop RevealX NPM, you gain complete network visibility and streamlined workflows that enable faster troubleshooting and improved mean time to resolution.

Table of Contents

Stream Processing

Full-content analysis and decryption at the speed of the data center

The backbone of our technology is the real-time stream processor that transforms unstructured packets into structured wire data at up to 100 Gbps. Architected for parallel processing, the stream processor splits processing tasks across multiple computing cores so you get deeper insight at a fraction of the cost per Gbps of analysis compared to other real-time analytics platforms.

Once the real-time stream processor receives a copy of network traffic from a tap or port mirror, here's what goes on beneath the hood:

1. Line-Rate Decryption

The stream processor decrypts SSL/TLS-encrypted traffic, including TLS 1.3 and TLS cipher suites that support perfect forward secrecy. Additionally, ExtraHop supports decrypting protocols encrypted via Kerberos or NTLM, as seen in Active Directory environments, as well as LDAP, RPC, SMB, and WSMAN. This bulk decryption occurs at line rate.

2. High-Performance TCP State Machines

Starting at the most fundamental level, the stream processor recreates the TCP state machines for every sender and receiver communicating on the network. A prerequisite for deeper application-protocol and universal payload analysis, this allows the platform to understand all TCP mechanisms and their impact. Because TCP is where the network and application meet, this approach helps you clearly identify whether problems are a network or an application issue right from the start.

3. Wire-Protocol Decoding and Full-Stream Reassembly

The real-time stream processor decodes IP-based protocols (skip to Protocols We Decode) in order to understand, define, and act on that protocol's unique application boundaries. This allows the processor to construct complete flows, sessions, and transactions for total application fluency, which in turn allows for higher-order content analysis through full-stream reassembly into wire data (derived from the wire protocol itself).

While in a perfect world this would all run pretty smoothly from start to finish, in reality traffic patterns like microbursts might result in packet loss from tap or SPAN or cloud-native traffic mirroring; in those cases, the processor will automatically resynchronize and recover.

4. Full-Content Analysis

After reassembling packets into full streams, the stream processor analyzes the payload and content from layers 2-7, auto-analyzing and classifying any device or client communicating on the network. The processor also continuously maps the relationships between all clients, applications, and infrastructure communicating on the network with over 5,000 metrics measured and recorded out-of-the-box.

Full-content analysis supports more than 90 of protocols, providing key performance indicators such as database methods used and their process time, file access by user, storage access time and errors, DNS response time and errors, web URI processing time and status codes, SSL certificates with expiration, and load-balancer and firewall latency. The stream processor also gathers sophisticated network metrics such as receive-window throttles, retransmission timeouts, and Nagle delays.

We get that not everyone is interested in knowing every detail about every layer of their environment, however, so don't worry—while the full analytics capabilities are always available to you, it's also easy to tailor your experience so you only see the precise metrics and insights you need.

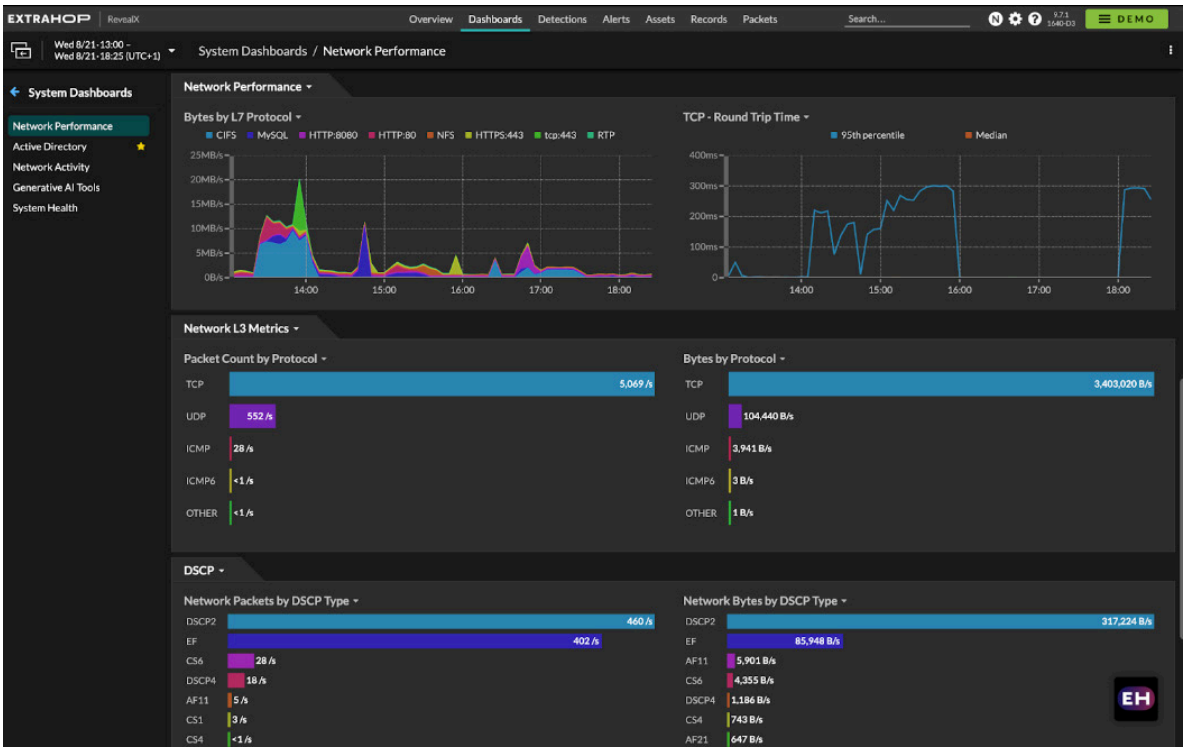

5. Fully Programmable Insights

Once the stream processor has done its thing and begun supplying wire data metrics, it's time to take control of which insights you see and at what depth.

ExtraHop uses an event-driven programmable interface called Application Inspection Triggers to connect you to the stream processor and all stream transactions. Triggers allow you to programmatically extract wire data events and correlated metrics that are specific to your business, infrastructure, network, clients, and applications.

With Application Inspection Triggers, you can be as surgical or as verbose as you want and extract nearly anything from a header to the full application payload. For example, with HTTP payloads, this data can include revenue, order IDs, unique user IDs found in cookies or URIs, and even titles for web pages or error descriptions embedded by a developer in a 500 status code. And it doesn't matter what's traversing HTTP—it could be SOAP/XML, REST, JSON, JavaScript, or HTML.

The same principle and functionality hold true for all of our natively decoded protocols. You can also use triggers to extract, measure, and visualize data from defined fields, or to decode proprietary protocols based on TCP and UDP.

Machine Learning & Anomaly Detectors

Advanced behavioral analytics guided on wire data metrics

Our cloud-based machine learning service tracks detections in eight categories across your environment:

- Authentication, authorization, and access control Network file system
- Network
- Infrastructure
- Database
- Email server
- Web server
- Remote access servers and methods
- Internet Communications and Telephony

Within each of these categories, our machine learning and anomalous behavior detectors evaluate dozens of protocols and hundreds of ExtraHop metrics, all with custom logic, in order to find and correlate active problems.

Architecture Overview

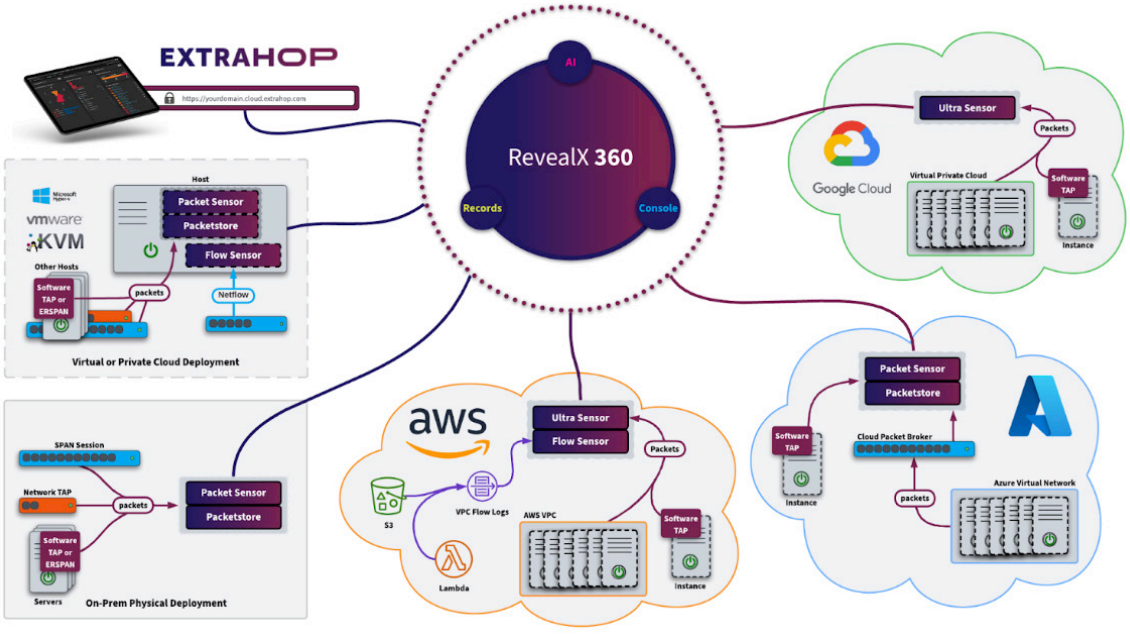

ExtraHop has received SOC 2, Type 2 and SOC 3 compliance certification for our machine learning technology. ExtraHop uses the following combination of on-premises tech and cloud services to support the full ML process:

1. A physical or virtual sensor, either self-managed or SaaS-based, analyzes network traffic to extract and store 5,000+ metrics including IP addresses, URIs, database queries, CIFS filenames, VoIP phone numbers, and other potentially sensitive data; you can configure this device to collect custom metrics as you choose.

2. When the ML service is enabled, a subset of these metrics are de-identified and sent to a customer-dedicated cloud-computing instance in Amazon Web Services, which is operated by ExtraHop.

3. ExtraHop then builds predictive models for how we expect devices and applications to behave, and detects significant deviations from these predictions as anomalies.

4. Anomaly events are sent back to your on-premises device, although you can also opt-in to receive email alerts (which don't include any sensitive data). Once events are back inside your environment, you can re-identify and decrypt them with your private key for alerting and investigation.

Performance Detections

On the performance side, we detect issues such as quality problems in VoIP linked to increased latency and errors, system boot-up and login delays associated with DHCP server errors, and more.

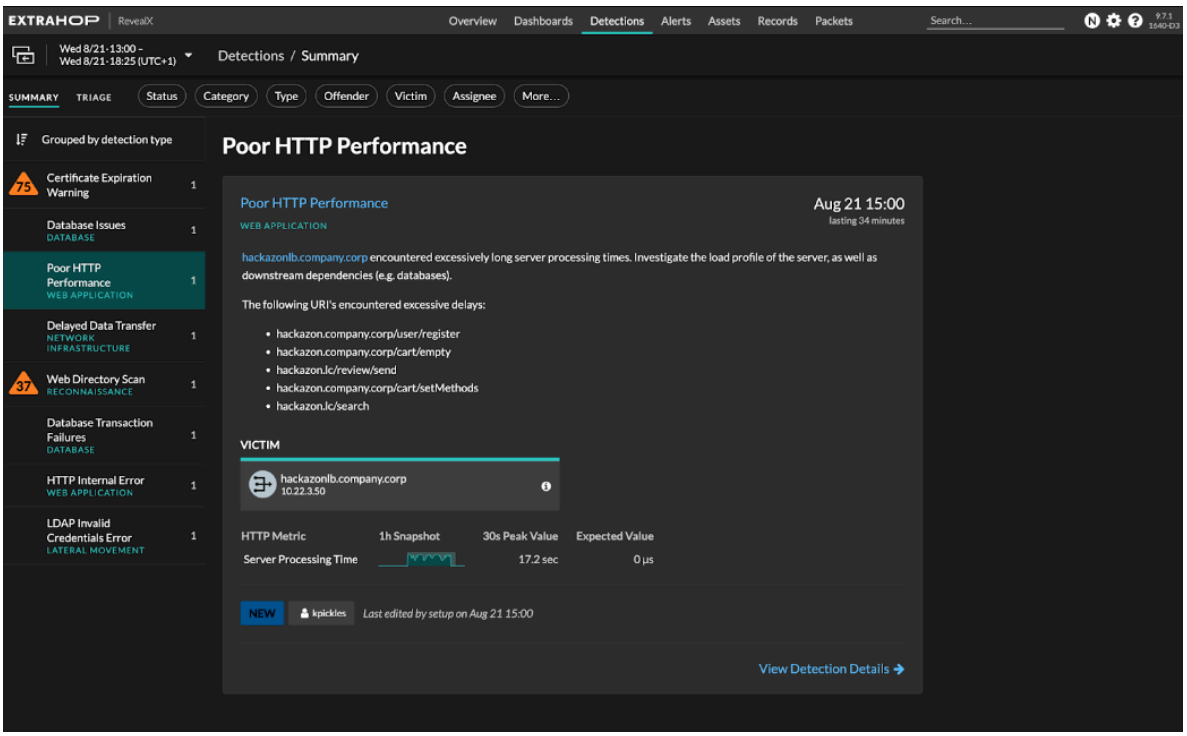

Data Indexing & Storage

Complementary formats to index and store your data

ExtraHop uses three complementary formats to index and store your wire data:

1. Correlated, cross-tier metrics in the streaming datastore

Optimized for time-sequenced telemetry, the streaming datastore enables customizable dashboards that can be populated with more than 5,000 possible metrics in real time. This way you can easily see all communications across your entire environment, or narrow your focus to specific datasets.

As metrics are indexed in the datastore, newly discovered devices are automatically classified based on heuristic analysis of machine information and behavior, and ExtraHop begins building activity baselines for all systems, applications, and networks.

You can use your existing NAS infrastructure to extend the streaming datastore for long-term lookback, which is helpful for capacity planning, proving compliance efforts or continuous improvement, and analyzing business activity over time. By default, your datastore will store fast (30-second), medium (5-minute), and slow (1-hour) metrics locally. You can, however, store 5-minute, 1-hour, and 24-hour metrics externally.

The datastore also allows you to create alerts based on current or past behaviors and events such as unusual payload size or expiring SSL certificates.

2. Transaction, message, and flow records

ExtraHop allows you to conduct multidimensional analysis of your wire data even if you don't know any query languages. Think of this like the search capabilities you'd find in a log analytics platform, except you're searching and analyzing wire data—a much richer, more consistent, and more reliable source of information than machine logs can provide.

There are two basic types of records in the ExtraHop UI: flow and transaction. Flow records show network-layer communications between two devices over an (L3) IP protocol, while L7 records show details from individual messages or transactions over any of the three types of supported L7 protocols (transactional, message-based, and session-based). ExtraHop allows you to search and filter for L7 traffic only, or to query both flow and L7 records.

Your transaction, message, and flow records are all stored in a resilient cluster built on scalable Elasticsearch technology so you can easily add nodes as your data grows.

In addition to ExtraHop hardware appliances or third-party storage infrastructure, users can choose to securely store records for 30, 90, or 180 days in a cloud-hosted record store via SaaS-based RevealX 360 Standard Investigation. Users can purchase records capacity for Standard Investigation in 25 GB, 100 GB, and 1 TB increments, and they have the ability to "stack" storage increments for additional capacity.

3. Packets for the full payload

You can either begin with individual metrics, users, devices, or packets associated with a particular transaction, or easily drill down to that information from a high-level view. ExtraHop supplies packets that offer the full payload, which you can download and analyze further as needed.

ExtraHop offers a highly scalable packetstore for extended forensic lookback, or you can send packets to the storage infrastructure of your choice.

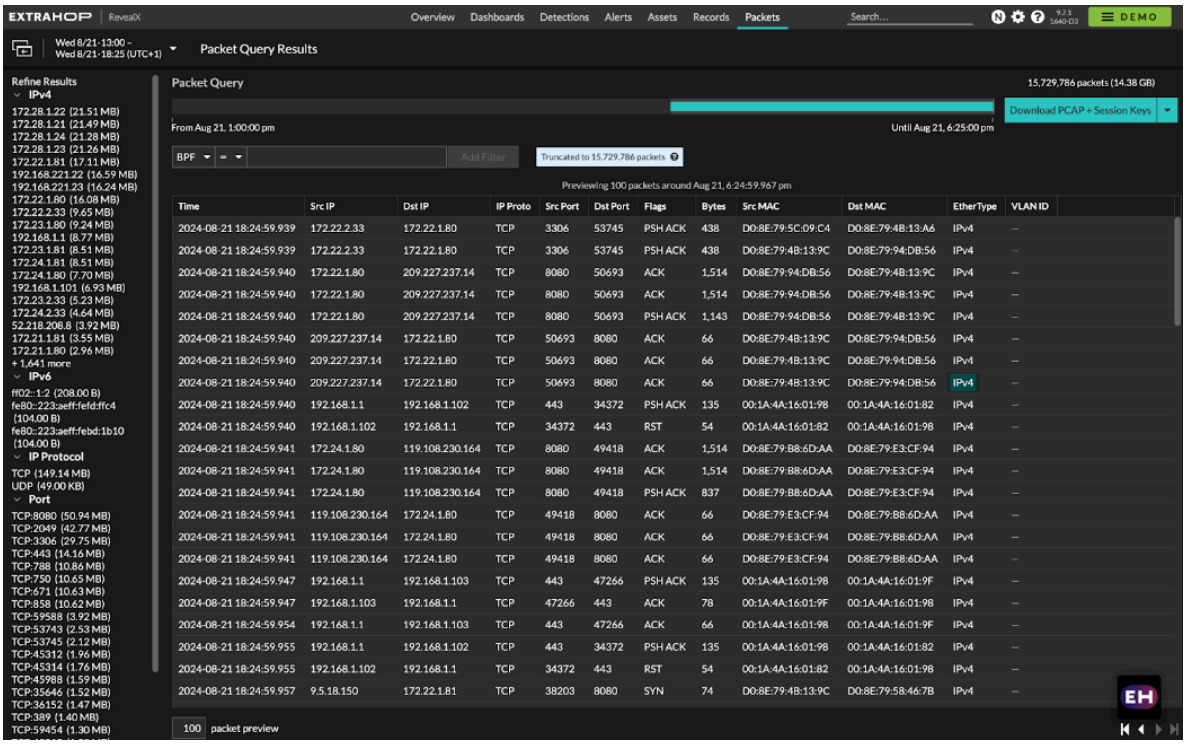

Data Visualization and Exploration

Intuitive querying, live activity maps, and open data stream

One of the most challenging aspects of real-time analytics at enterprise scale is, well, the scale itself. At ExtraHop, we do our best to make this easy for you as a user to parse the immense wealth of information that is wire data and derive meaningful insights no matter which perspective you're coming from.

We start you off with a simple, intuitive user interface that includes automatically populated role-based dashboards for teams across your organization. These dashboards function on a drag-and-drop model so you can customize them further with unique widgets; if you want to create your own widget, all you have to do is select your desired data source and metrics, pick a visualization type, and save it to your dashboard of choice. You can quickly and easily export charts and background data points to PDF, Excel, or CSV.

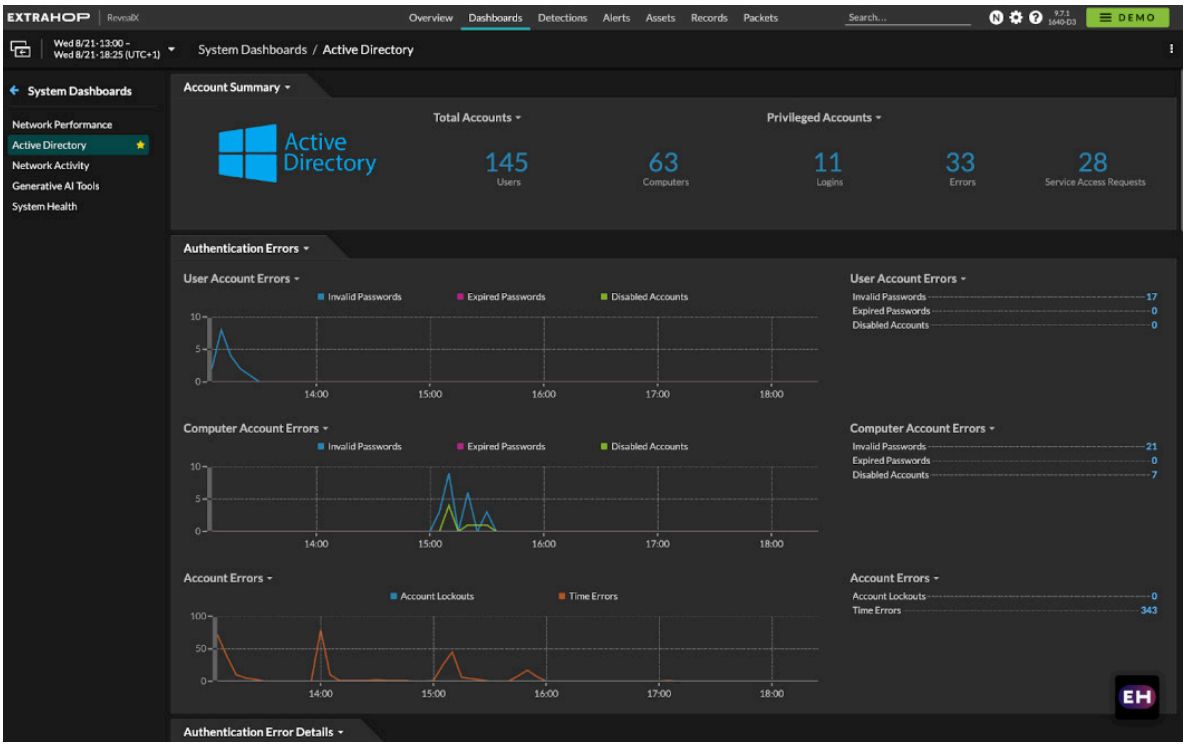

No Scripting? No Problem.

Our visual query language gives you the power to refine or change your search queries by clicking UI elements that control everything from grouping, to filtering, to time-range selection. Whether you stick with the hundreds of built-in record attributes or branch out and define your own, this functionality means any user can quickly answer performance and security questions without needing to learn a query language.

For example, if you're experiencing poor audio quality over VoIP, you can search VoIP traffic for expedited forwarding tags and quickly determine which packets might be lagging behind less time-sensitive traffic due to a misconfigured tag.

Live Activity Maps

Along with traditional methods of data visualization like charts and graphs, ExtraHop uses live activity maps to present a dynamic and intuitive view of your environment. Instead of manually creating and updating network diagrams as your IT environment changes, you can use live activity maps to visualize protocol-based connections between devices and applications in real time.

By allowing you to filter by time interval and broaden or narrow your scope as needed, activity maps make it easy to answer multi-part questions like, “How are devices interacting within a certain tier, and how have those devices been interacting across the network in the last hour?” Anomalous behavior detections also appear on live activity maps, so you can see the context of the detection before clicking down into the transaction or even into the precise packets straight from the map.

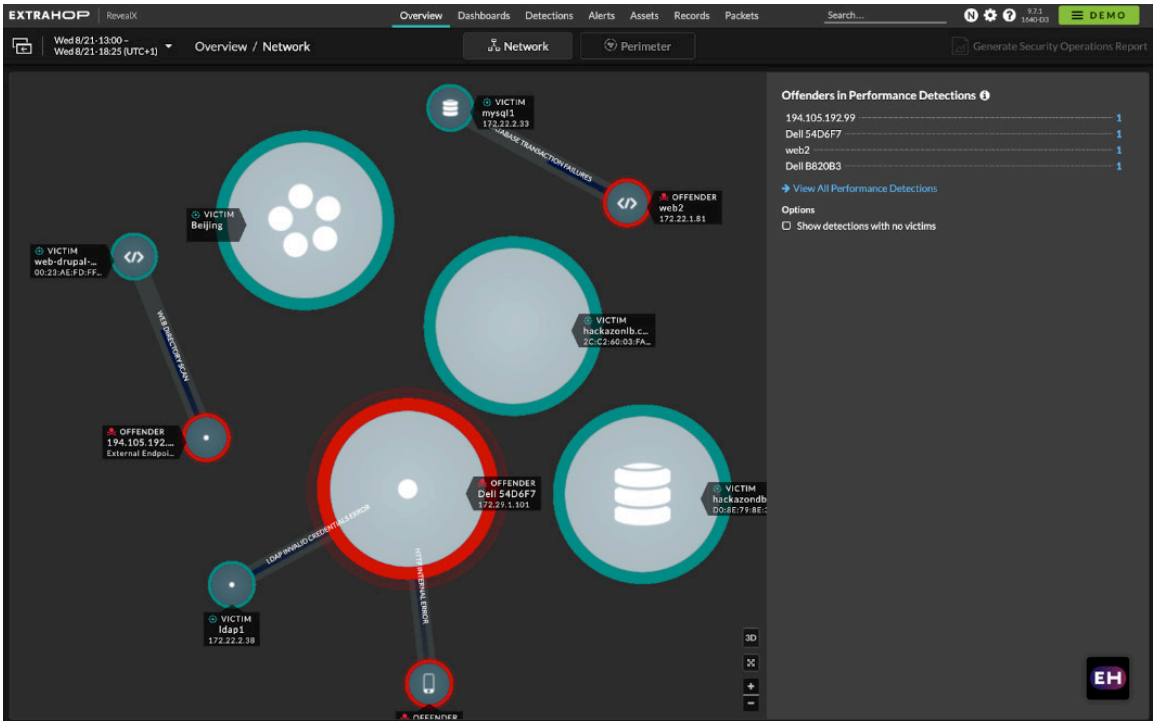

AI Search Assistant

The ExtraHop AI Search Assistant allows users to create expressive search queries for devices and records using natural human language. With AI Search Assistant, users of any experience level can quickly locate information using a search bar in the RevealX UI without the need to have knowledge of ExtraHop's query builder. This search feature currently uses the Claude 3 Haiku LLM model hosted in AWS Bedrock with no fine-tuning required. No customer data is shared with AWS Bedrock, except for the questions that the users explicitly send to the AI Search Assistant.

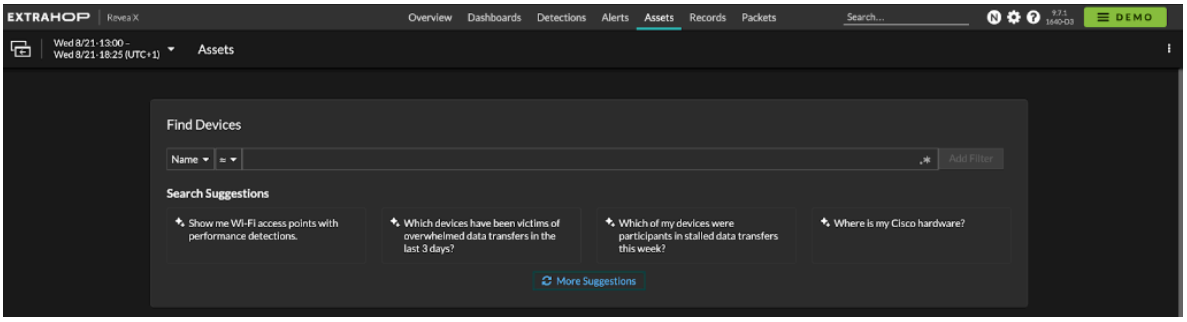

Open Data Stream

While we supply rich query and investigation workflows within the ExtraHop interface, we also make it easy for you to integrate wire data metrics with the other data stores, querying tools, and analytics platforms in your stack. Open Data Stream allows you to merge data from multiple sources into a single, rich set that can be queried and visualized using whatever tools your team prefers.

Visit our Technology Partners page to learn about specific integrations such as our partnerships with AWS, Splunk, F5, Palo Alto, Cisco, AppDynamics, ServiceNow, Elastic, MongoDB, and many others.

Protocols We Decode

Over 90 enterprise protocols decoded in real time

ExtraHop decodes the following enterprise protocols with real-time fluency at the application layer. Protocol modules offer varying levels of analysis, starting with L7 classification, and Application Inspection Triggers allow you to create a custom metric.

Protocols Supported

RevealX

- AAA: Diameter
- AAA: RADIUS
- ActiveMQ
- AJP
- ARP
- BitTorrent
- CFP
- CIFS
- Citrix ICA
- CoAP
- Cryptocurrency mining protocols
- Database: DB2
- Database: Informix
- Database: Microsoft SQL
- Database: MongoDB
- Database: MySQL
- Database: Oracle
- Database: Postgres
- Database: Redis
- Database: Riak
- Database: Sybase
- Database: Sybase IQ
- DHCP
- DICOM
- DNS
- DSCP
- FIX
- FTP
- GENEVE
- GRE
- HL7 (including FHIR and ICD-9/10)
- HTTP-AMF
- HTTP/S
- IBM MQ
- ICMP
- ICMP6
- IEEE 802.1X
- IIOP
- IKE
- IMAP
- IPFIX
- IPSEC
- IPX
- IRC
- ISAKMP
- iSCSI
- Java RMI
- Kerberos
- L2TP
- LACP
- LDAP
- LLDP
- LLMNR
- Memcache
- Microsoft NMF
- Modbus
- MPLS
- MS-RPC
- MSMQ
- Netbios
- NetFlow and SFlow
- NFS
- NTLM
- NTP
- NVGRE
- OpenVPN
- PCoIP
- POP3
- PPTP
- QUIC and GQUIC
- RDP
- RFB (VNC)
- RSH
- Skinny (SCCP)
- SMPP
- SMTP
- SNMP
- SOCKS
- Splunk
- SSH
- SSL
- STP
- Syslog
- TCP
- Telnet
- TFTP
- TRILL
- VNC
- VoIP: RTCP
- VoIP: RTCP XR
- VoIP: RTP
- VoIP: SIP
- VXLAN
- WebSocket
- Windows Update Delivery Optimization
- WireGuard
- WMI
- WSMAN

ABOUT EXTRAHOP

ExtraHop is the cybersecurity partner enterprises trust to reveal cyber risk and build business resilience. The ExtraHop RevealX platform for network detection and response and network performance management uniquely delivers the unparalleled visibility and decryption capabilities that organizations need to investigate smarter, stop threats faster, and move at the speed of risk. Learn more at extrahop.com.

EXTRAHOP®

info@extrahop.com
extrahop.com