ExtraHop

520 Pike Street, Suite 1600
Seattle, WA 98101

ExtraHop Reference Architecture Guide

Abstract

ExtraHop transforms IT by decoding and visualizing wire data transiting the network. This guide introduces basic architectural principles, critical design considerations, and decision points for use in deploying ExtraHop across a variety of physical and virtual platforms. The contents assume the reader is familiar with datacenter networking and application monitoring concepts, and enable the reader to design comprehensive solutions using ExtraHop products.

Executive summary

The ExtraHop platform is powerful and easy to use, but careful planning is necessary to create high-quality integrations in complex environments. By creating an architecture that is appropriate for the needs of the environment, the ExtraHop platform will capture the correct traffic and assess the appropriate metrics; the platform will support a logical, data-oriented IT model, and the user will gain insight from wire data that is trusted and actionable.

The ExtraHop platform consists of several components and relies upon a few unique concepts for the understanding of proper architecture. These components and concepts are described in the *Platform* section.

The ExtraHop platform integrates with a wide variety of network and application architectures, and experience has shown that different environments require differing architectural approaches. The *Architecture* section presents an architectural classification system - including the decision points used to classify environmental factors - and discusses important guidelines used to make decisions when designing ExtraHop integrations.

This document presents the following architectural models:

- Datacenter deployment, core layer
- Datacenter deployment, distribution layer
- Datacenter deployment, both sides of a Layer 3 boundary
- Packet aggregation
- Packet aggregation, multiple ExtraHop Discover appliances
- Transit link

- Multiple datacenter
- Datacenter and remote site
- Virtual deployment
- Amazon Web Services deployment: single vs. multiple availability zone architectures

For information about how to use the ExtraHop platform, please reference *the ExtraHop Training page at https://www.extrahop.com/support/training/.* To customize the platform to capture and visualize your network and applications, please reference the *Solutions Architecture Deployment Guide*.

Platform

The ExtraHop platform centers around passive, out-of-band network appliances that consume Ethernet packets from a network tap, port mirror, or packet aggregation device. The platform performs full-stream reassembly and content analysis of this traffic to extract and structure wire data. ExtraHop analyzes application transactions continuously and in real time, at speeds up to a sustained 40Gbps, and records thousands of metrics that describe network, device, and application performance. ExtraHop's flagship model (EDA9100) supports analysis for up to 1.3 million HTTP transactions per second, with bulk decryption at 40 Gbps and 64,000 SSL handshakes per second for 2048-bit keys. An open and extensible platform, ExtraHop enables IT teams to define and implement new metrics within minutes, and shares data with other IT operational analytics and data processing technology.

The traditional approach to obtaining visibility across all the tiers of an IT environment is to pull as many discrete metrics as possible from each tier, via logs, synthetic transactions, or point monitoring tools, and try to make sense of the collected data with analysis and reporting servers. ExtraHop takes a different approach, using wire data as the source for cross-tier insight. The network is the common element, tying all components of the application delivery chain together, even as those components become more numerous and distributed. Each component communicates with others using transport and application protocols. These protocols definitively describe events in the IT environment. Using this wire data as the foundation for IT decision-making provides a path to faster, better solutions.

To take advantage of these ideas, an architect must understand the structure of the network, the network activity generated by the application that sits on the network, and the capability of the network components to deliver a high-quality feed of network traffic to the correct ExtraHop appliance. Such a feed contains all the packets that encapsulate all the activity of all the applications under observation.

ExtraHop Appliances

ExtraHop Discover Appliance

The ExtraHop Discover appliance (EDA) is a high-speed Ethernet capture device with L2-L7 protocol fluency. It interprets the traffic it sees and creates a model of over 3500 metrics that describe each of the devices under observation. Users interact with the EDA via a Web interface that can be used to manage the appliance and to create rich visualizations of these metrics. The EDA is available in a variety of configurations, both physical and virtual.

ExtraHop Command Appliance

The ExtraHop Command appliance (ECA) is a virtual appliance that provides centralized Web management and visualization and supports multiple ExtraHop Discover appliances. For distributed environments, the ECA delivers a consolidated view of wire data from multiple ExtraHop appliances, enabling organizations visibility into communications of hundreds of thousands of devices across datacenters and branch offices.

ExtraHop appliances can be associated with multiple ECA instances. This ability allows role-specific ECAs to be implemented to serve specific roles and needs, including security team monitoring and appliance-firmware maintenance. ECA also provides an interface for scheduled reports to be created and sent via email.

ExtraHop Explore Appliance

The ExtraHop Explore appliance (EXA) provides an integrated, bottom-up record search facility for data extracted by an ExtraHop Discover appliance. One or more ExtraHop Discover appliances perform stream processing on Ethernet packet data to extract information before writing a record to the EXA. Users can access EXA queries through the EDA's UI, which allow them to perform queries across the records written to the EXA.

Data Acquisition

How to get packets into the ExtraHop

Several options are commonly used to feed Ethernet packets into an ExtraHop Discover appliance. Different network gear and different virtual environments offer different options and different limitations, and some networks may not support some of these features at all.

Port mirroring (SPAN)

Enterprise switches typically provide an option to designate one or more ports as a mirror port, which receives a copy of other traffic traversing the switch. Different platforms provide different criteria for selection of traffic to be mirrored, but selection by VLAN, by port, and by direction are typically possible. Cisco gave the acronym SPAN to this type of port, and the name SPAN port has become ubiquitous even though each manufacturer has a specific term.

Such ports are convenient because some of them may already be available in a network without having to add any additional equipment. They are inconvenient because they are often all spoken for by other monitoring or surveillance technology.

There's also a chance that if a monitor port is planned poorly or supported incompletely by the network device, that the collection of a mirrored stream of packets might cause packet loss on the other ports. Packet loss on other ports can disrupt production traffic, and can be avoided by considering the specific limitations of the environment and equipment at hand.

Taps

When an optical or electrical link exists between two network devices, a tap can passively duplicate all of the traffic on the link and send it to a collection device. Taps have the benefit of being completely separate from the switching infrastructure, so they won't interfere with production traffic even if they're overloaded or misconfigured.

Downsides include the fact that they represent another layer of infrastructure to deploy and maintain, and the nature of taps to typically be more numerous than switches, which often requires the use of packet aggregation.

Packet aggregation

When a stream of packets is collected from numerous sources and is to be delivered to one or more tools that have limited port density, packet aggregation appliances can be used. Huge packet streams can be divided into sub-streams and directed according to various metadata at one or more tool ports.

Many packet aggregation devices provide a sophisticated layer of filtering to incoming traffic. This allows for a more selective and organized approach for delivering data to a monitoring device such as ExtraHop.

RSPAN, ERSPAN

Certain switches and virtual switches have a capability similar to SPAN, where a set of traffic can be collected for monitoring. Instead of being sent to a particular interface, however, the traffic is assigned to a specific monitoring VLAN and mirrored back into a Layer 2 domain (remote SPAN or RSPAN), or encapsulated in IP and sent across the wire to a target, which captures and processes the feed (encapsulated remote SPAN or ERSPAN).

This technique can add flexibility to a packet capture approach because packets can be collected across significant distances in locations where packet capture hardware does not exist. It can also put considerable strain on a network because it causes substantial duplication on the wire if not used judiciously.

RPCAPD

RPCAPD (Remote Packet Capture Daemon) is a software package that runs a lightweight packet forwarder. The software includes a server daemon that runs on a host, uses operating system facilities to gather packets in promiscuous mode, encapsulates those packets in UDP, and sends them to an ExtraHop Discover

appliance acting as an RPCAP client for stream processing. RPCAPD is available from ExtraHop for Windows and Linux environments.

In cloud deployments, where physical access to infrastructure is not possible, RPCAPD becomes the primary method for data acquisition.

Capture ports and management ports

ExtraHop Discover appliances copper and optical Ethernet ports, which have different capacities and restrictions, can be assigned to different functional roles depending on appliance model and the requirements of the integration. The highest-capacity optical ports are used as capture ports, with Ethernet packets delivered to these ports from switches, taps, or packet aggregation systems. The lower capacity copper ports can be used as capture ports and/or management ports, as well as destinations for RSPAN, ERSPAN, or RPCAP streams.

Figure 1. below represents an ExtraHop Discover appliance (EDA6100). We can see that one of two 10Gbps optical ports (highlighted in green) is in use for data collection. We can also see that one of the 1Gbps copper ports (highlighted in blue) is in use as a management connection.

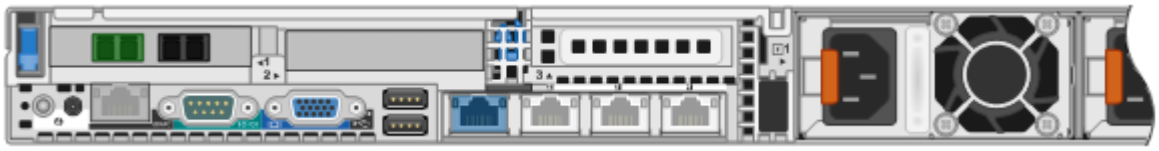

Stream processing

Several pre-processing facilities for improving feed quality are provided in the ExtraHop Discover appliance. It's often possible to solve these problems upstream of the ExtraHop, either by choosing more accurate rules to include traffic in the feed, or being intentional about the locations on the network from which feeds are sourced.

De-duplication

The EDA can tell the difference between packets retransmitted because of TCP rules and duplicate packets introduced into the packet stream because feeds are taken from multiple points in the network, all of which are transited by the packets in question. When the EDA detects multiple identical packets within a short period, it can optionally discard the duplicates, and process only the original packet. This automatic de-duplication provides more accurate accounting of TCP retransmissions while reducing unnecessary resource utilization on the EDA. De-duplication can be performed at Layer 2 (the Ethernet header and the entire IP packet must match) or at Layer 3 (only the IP packet must match).

Protocol de-encapsulation

MPLS, TRILL, and Cisco FabricPath protocols are de-encapsulated automatically. VXLAN and NVGRE overlays can also be de-encapsulated if needed. Note that removal of these packet overlays does place additional load on the ExtraHop appliance.

Feed Quality

Feed quality allows us to discuss the characteristics that make a feed useful. A *high-quality feed* consists of all of the packets (and only the packets) that contain all communication between all devices that make up the application targeted for monitoring. A *low-quality feed* might be missing all of the packets from certain devices because of improper configuration. Or, some of the packets might be missing or out-of-order because of packet loss or disruption in the feed delivery system. There might be too many discovered devices in the feed, causing resource exhaustion on the EDA. Similarly, too many packets or too much bandwidth can also exhaust EDA capture resources.

The process of deploying the ExtraHop platform can be conceptualized as determining the right feed, engineering a feed of sufficient quality, and delivering the right part of the feed to the right ExtraHop appliance.

Device Count, Packet Rate, and Throughput

Part of determining what a high-quality feed looks like is to understand how you intend to limit the feed. Users should consider the following factors when creating a deployment architecture:

Device Count - Different ExtraHop platform sizes and configurations determine the number of discoverable devices. An optimal data feed is designed with the correct device limits maintained.

Packet Rate - Maximum packet rates for an ExtraHop appliance depend on the model in use. Also, it is best practice to leave room in the maximum rate for future growth across the platform. Leaving 25 percent of the maximum is a good starting point. Heavy customization of the appliance (Triggers, Trends, Alerts, etc.) warrants leaving closer to 50 percent of the maximum rate.

Throughput - Maximum throughput is also variable depending on the model in use. Best practice is again to leave headroom for each interface in use. For example, sending 9 Gbps down a 10 Gbps link may lead to packet loss due to microbursting across the link. A better design might be to split this feed and send 4.5 Gbps down two 10 Gbps links.

ARPs and device discovery

This document discusses the ExtraHop platform's automatic device discovery at length below (link). ExtraHop’s ability to use ARP (Address Resolution Protocol) frames to distinguish local and remote traffic is central to device discovery . If ARPs are not part of the data feed sent to ExtraHop, device discovery will be limited. Therefore, as part of a deployment architecture and data acquisition plan, ARP traffic should be considered.

If ARP traffic is not possible, ExtraHop has other options available to make device discovery possible:

Remote Networks - This setting allows specific devices or subnets to be declared as devices to be discovered. While this option is easiest in configuration, the appropriate list will need to be created and maintained

ARPless Discovery - This setting allows device discovery without ARPs. It relies on observing devices with consistent traffic to be considered for discovery. This option tends to be more inclusive than Remote Networks and can grow the device list too large if not monitored and controlled correctly.

Features

Device Discovery

Upon receiving packets, the ExtraHop Discover appliance will immediately begin to understand conversations which are taking place between pairs of devices on the network which it is monitoring. When one or both of the devices are “discovered”, it means that the ExtraHop platform creates a container to hold metrics that describe the conversations in which that device participates. Each discovered device takes up some platform resources that are used to store these metrics. Some devices do participate in conversations without being discovered. For example, when monitoring an online retail website, discovering each of the end users' PCs and storing all the metrics associated with their transactions would consume significant resources. Instead, ExtraHop discovers the Web server and aggregates the metrics associated with the transactions in the Web server's device object.

Certain packet activity will provide clues regarding the nature and position of network devices. If the MAC to IP address translation provided by the Address Resolution Protocol (ARP) is visible to the ExtraHop, its presence can be used to infer that the participants are not separated by a Layer 3 / IP boundary; that is, they are on the same Ethernet segment. The first principle of device discovery on the ExtraHop platform is that every local device should be discovered.

Devices that are not local fall into two categories: those valuable enough to spend EDA resources to create a separate container for metrics (e.g., devices in an IP block assigned to a retail store that uses resources on a local device), and those not valuable enough to spend those resources (e.g., individual client devices from the public Internet which are too numerous to track individually). Non-local devices important enough to warrant separate collection of metrics are designated in the UI as Custom Devices or Remote Networks.

All conversations are run through a suite of protocol parsers to determine the L7 / application-level protocol. This protocol analysis is configurable for environments that use custom ports and protocols. Dynamic functional groups (e.g., all HTTP servers, all DNS clients) are automatically discovered in this way.

Packet capture

At the core of the ExtraHop platform is high-speed packet capture and analysis. Features exist to extract both global packet captures from the entire ExtraHop feed (useful for verifying the quality of the upstream feed) and precision packet captures from a single transaction. These packet captures can be downloaded from the platform and imported into external analysis tools using common file formats.

Open Data Stream

The ExtraHop philosophy is that wire data should not be locked inside the ExtraHop platform, but free to be visualized, contextualized, and aggregated in any platform. The Open Data Stream allows ExtraHop to export wire data to many of the industry's most popular data stores to support correlation with complementary data sources. Open Data Stream enables IT teams to send policy-based, event-driven metrics to any IT management console, custom Big Data analysis store, SIEM product, or third-party management tool such as Keynote, Kafka, Syslog, SevOne, or Splunk.

ExtraHop API

The EDA, ECA and EXA can all be controlled and queried by a powerful, REST-based API. All of the data access techniques that the EDA uses to build its Web UI are available in the API to pull data out of the ExtraHop platform. The API can also be used to apply metadata from other information sources such as infrastructure catalogs or scanners.

The API must run in an environment where it can access the appropriate appliance via HTTPS on port 443.

Cloud Atlas

ExtraHop Discover appliances can connect via an encrypted HTTP connection to a cloud service which allows ExtraHop services and support personnel to make queries to attached appliances for data visualization reports. Packet data is not exported to Cloud Atlas, and the connection between the ExtraHop Discover appliance and Cloud Atlas is encrypted and authenticated. This facility is used to allow ExtraHop personnel to access a deployed ExtraHop appliance to make customizations or to provide support. Cloud Atlas connectivity requires that the ExtraHop Discover appliance be able to perform DNS resolution on the domain .a.extrahop.com and to connect to hostnames resolved in that domain on TCP port 443 over SSL.

Architecture

The ExtraHop platform architecture is based on capture and analysis of Ethernet packet data. Designing a system to deliver the right packets to the right ExtraHop Discover appliance is essential to deploying ExtraHop platform.

Virtually all networks have some ability to copy a certain group of traffic to a port where a tool can listen to monitor the network traffic or its contents. Cisco calls this a SPAN port; other vendors have other terminology, but we use the term SPAN generally in this document. While it is possible to gather packets with software running on a host (e.g., with the RPCAPD software forwarder) it is generally advisable to use SPAN ports, taps, or packet aggregation infrastructure to feed Ethernet packets to an ExtraHop Discover appliance.

While we may sometimes use terminology such as "the Web server is part of this feed," the components of a feed are better understood as network links or portions of links like VLANs. It's important to understand this because the traffic from a device may flow over a number of links. If not all of them are properly brought to the ExtraHop, the feed may be incomplete.

Unlike other technology from some vendors, ExtraHop Discover appliances don't sample. We listen to every packet, and then build a simulation of every device's network stack-state to infer activity in the environment. Every packet is important, and having a plan for exactly which packets needs to be delivered to the ExtraHop is key.

Example Deployment Scenarios

The sections below describe several deployment scenarios for the ExtraHop platform. This guide presents each example along with reasons why it is applicable to certain environments and not desirable in other cases. These examples provide guidance on integration with the ExtraHop Platform. In most cases, these idealized examples will be hybridization. For example, in a medium-sized enterprise, it would be reasonable for a single high-powered EDA to capture a feed from a transit link to capture cloud traffic, from the core to capture most application traffic, and from behind the Layer 3 boundary of a load balancer, effectively combining three examples into one architecture.

Datacenter deployment, core layer

A basic ExtraHop architecture takes a feed from a SPAN port on a core switch. The switch defines the traffic that goes to the SPAN port, by source or destination port or by VLAN, and the responsibility of feed design includes understanding how to specify the set of traffic that includes all desired application traffic and nothing else, from the point of view of switch configuration. Putting all of the traffic from the core on the ExtraHop is rarely desirable; picking applications, understanding their traffic patterns, and designing suitable SPAN rules to deliver just that traffic to the ExtraHop is usually more useful.

For use when:

- This is the simplest, preferred architecture, useful when all of the traffic being considered is handled in the core switch.

Reasons not to use:

- A network might not contain any equipment capable of providing a data feed. Some switches that can be used as the core of a small network can’t, and there may be no facility for physically tapping any of the links in such a network. Or maybe all the SPAN ports on the core are in use; some only have two. In this case, try to deploy network taps and consider the Packet aggregation example or, in extreme cases, consider deploying RPCAPD to servers where applications live to get a direct packet feed from each server.
- Part of a network may be behind a routable Layer 3 boundary, and it might be necessary to monitor both sides. In this case, see the datacenter deployment, both sides of a Layer 3 boundary example.
- A network’s data feed may come from many more ports than can be plugged into the back of your ExtraHop. In this case, consider the Packet aggregation example.

Datacenter deployment, distribution layer

Sometimes, much of the traffic for some parts of an application does not traverse the core switch and consists primarily of "east-west" traffic which only travels as far as the distribution switch before being directed to the peer, which is also connected to the distribution switch. ExtraHop can provide better visibility into applications like this when instead of SPANing the core switch, we SPAN multiple distribution switches.

For use when:

- Not all of the traffic you want to monitor travels through the core switch.

Reasons not to use:

- There are too many distribution switches to plug a SPAN port from each directly into an ExtraHop. In this case, consider the Packet aggregation example.

Datacenter deployment, both sides of a Layer 3 boundary

The diagram in Figure 4 illustrates a common deployment scenario similar to that shown above, where traffic is captured from SPAN ports from switches on both sides of the load balancers, which are characterized as a Layer 3 boundary. This allows troubleshooting of data traffic issues both from the access network (i.e. above the load balancers) and also all datacenter internal traffic between the various application and databases servers.

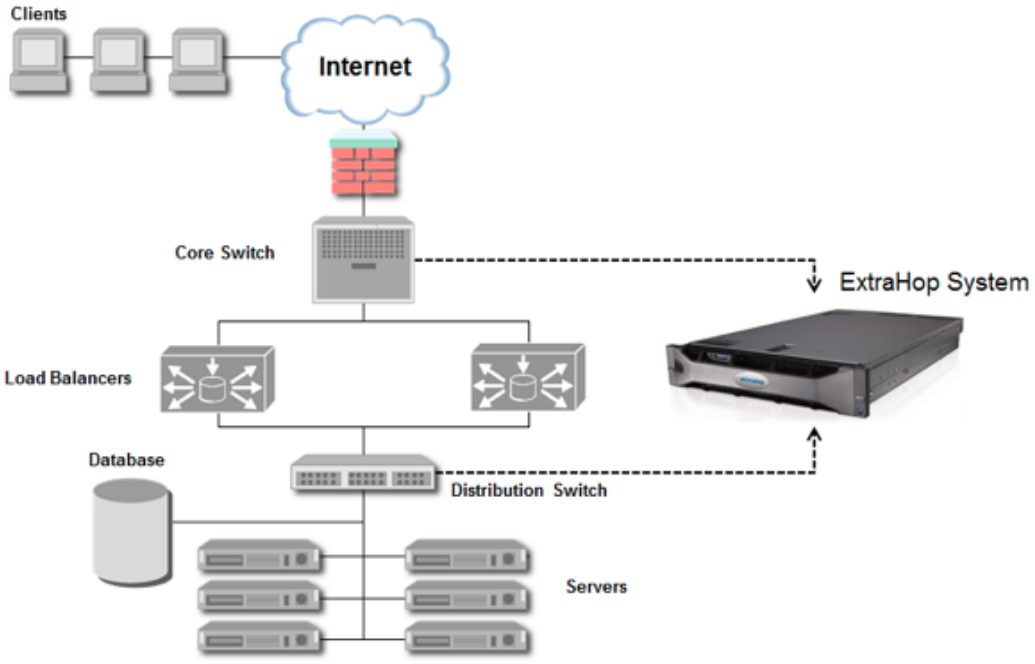

Figure 4 Example Datacenter Deployment

In such a configuration, the ExtraHop listens to feeds coming from two separate Ethernet broadcast domains, and can be configured to either combine the metrics from all the conversations in both Layer 2 domains together or to keep metrics separately for the separate views separately, at the cost of increased resource

requirements. Keeping the metrics separate enables isolation of the front-side and back-side environments, which can aid greatly in troubleshooting.

For Use When:

- It is necessary to directly monitor a network on both sides of a routable Layer 3 boundary. In the ExtraHop UI, this will exhibit as missing devices, with all the missing traffic associated with one device representing the Layer 3 device.

Reasons not to use:

- This architecture is not necessary if there are no Layer 3 boundaries; in this case, see the previous example.
- A network's data feed may come from many more ports than can be plugged into the back of your ExtraHop. In this case, consider the next example.

Packet aggregation

If your network’s feed consists of too many physical connections to plug into the back of a single ExtraHop, use a packet aggregation device to create a single feed for the ExtraHop. This feed can be shared with other tools. Some packet aggregators allow the composition

For Use When:

- A network’s data feed comes from many more ports than can be plugged into the back of your ExtraHop, for example, when optical taps are in use, when many distribution switches are SPANed, or when there are many L3 boundaries with separate physical feeds.
- Other tools want to consume network feeds alongside the ExtraHop, and switches don’t have sufficient capacity to create a span.

Reasons not to use:

- Packet aggregation is another expense and another layer of complexity, so if a network is simple enough to capture a whole feed without aggregation, examples above will suffice.
- Sometimes, packet aggregators can create a composite feed that exceeds the total capacity of the fastest ExtraHop Discover appliances. In this case, consider the next example.

Packet aggregation, multiple ExtraHop Discover appliances

If your packet aggregation infrastructure produces feeds which exceed the capacity of an ExtraHop Discover appliance, send the feeds to a packet aggregator with sufficient capacity for your large feed. Then use advanced features on the packet aggregator to create multiple smaller feeds. Direct each of these feeds to its own tool port on the packet aggregator, and connect multiple EDAs, each to one tool port.

The critical factor in such a design is the manner in which the master feed is divided into sub-feeds. Customers have two options for managing and operating multiple ExtraHop Discover appliances. The EDAs can either operate alone, and users can log into each separately, for a completely separate experience.

Alternatively, the two EDAs can be joined to a community led by an ExtraHop Command appliance, and users can log into the ECA to address both EDAs. In the latter mode, the ECA can make all of the visualizations possible on each of the EDAs, and even allows users to create visualizations that show metrics from different EDAs in the same widget. However, the ECA does not seamlessly combine metrics from its member EDAs; it merely allows display of metrics alongside one another.

For example, in an environment where a very large feed is passed to a packet aggregation layer, where it is divided into two sub-feeds without regard to content or address (i.e., round-robin) before being passed to an EDA, each EDA maintains its own metrics for the content it sees. Therefore, there may end up being two separate HTTP Server Delay metrics: one average for each EDA. The ExtraHop Command appliance has access to both of these metrics, and can chart them against each other, but it has no way to combine them into an overall average for all the traffic. If the traffic on each of the ExtraHop Discover appliances cannot be meaningfully distinguished, the distinction between these two metrics will lead to confusion. On the other hand, if the master feed were deliberately divided so that one feed consisted of all traffic bound for one set of Web servers, while the other sub-feed consisted of traffic bound for a different set of Web servers serving a different application, the two different HTTP Server Delay metrics would be very useful.

Devices on your network must have affinity to a feed associated with a particular ExtraHop Discover appliance. If a device appears to one EDA sometimes and another EDA at other times, or on two EDAs simultaneously, resources will not be utilized optimally, and visualization will be impaired.

For use when:

- A network's data feed comes from a packet aggregation system, but the data presents on more physical connections that can be plugged into the back of your ExtraHop.
- Complex traffic modeling and/or security restraints require that subsets of data be processed separately, for example, when financial information must be maintained under tighter security constraints than general web traffic.

Reasons not to use:

- It can be difficult to divide a master feed in such a way that the resulting sub-feeds have clear conceptual boundaries and don’t overlap. Where possible, it’s best to use the largest ExtraHop Discover appliance available and minimize split feeds to minimize complexity. Consider the Packet aggregation model wherever possible.

Transit Link

An alternative deployment scenario is to capture all traffic carried over a transit link between two locations, such as between datacenters, or between two isolated sections of a single large datacenter. A SPAN port on one of the endpoint routers or a tap of the physical link would both allow for the extraction of such a feed to an ExtraHop Discover appliance. In this case, the ExtraHop appliance processes all traffic between these locations, independent of applications that are being used.

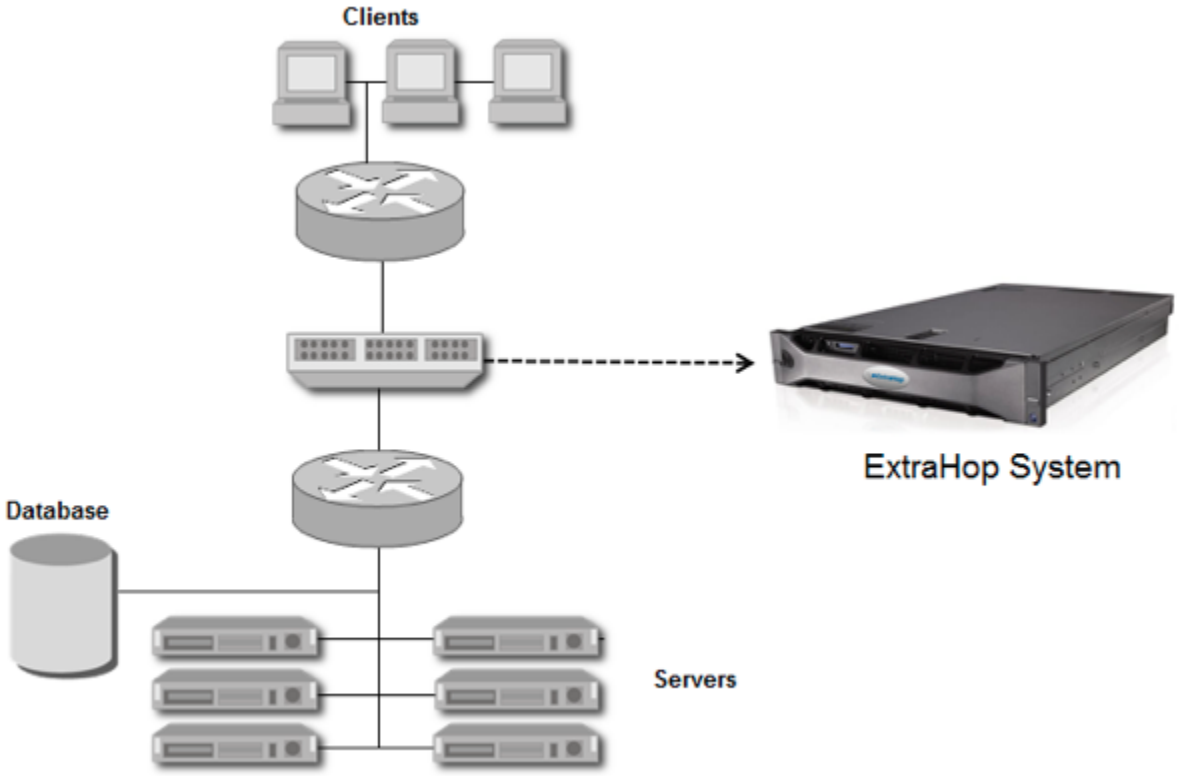

Figure 5 Example Transit Link Deployment

For use when:

- Taps are deployed on links between interesting networks
- There are natural choke points across which all data flows, but there are many switches that would need to be SPANed to collect the same traffic.

Reasons not to use:

- This architecture is limited in that it can not see ARP conversations for any devices other than the two routers on either side of the connection, which means that if device discovery and independent collection of metrics is required, Remote Networks or Custom Devices must be used.
- The ExtraHop platform's ability to differentiate between network delay and server delay is compromised in this architecture, since from the EDA's point of view, the network delay in the bottom of Figure 5 appears to be server delay happening in the bottom router. In this case, if observation of the network on both sides of the link is required, feeds should be taken from inside each of the Layer 2 domains on either side of the transit link (i.e., two implementations of the datacenter deployment, above.)

Multiple Datacenter Deployment

In many cases, large enterprises use multiple datacenters or have such large datacenters that multiple ExtraHop EDA's are needed. In these cases, one or more ECA appliances are used to access data from across multiple network areas via separate EDA appliances.

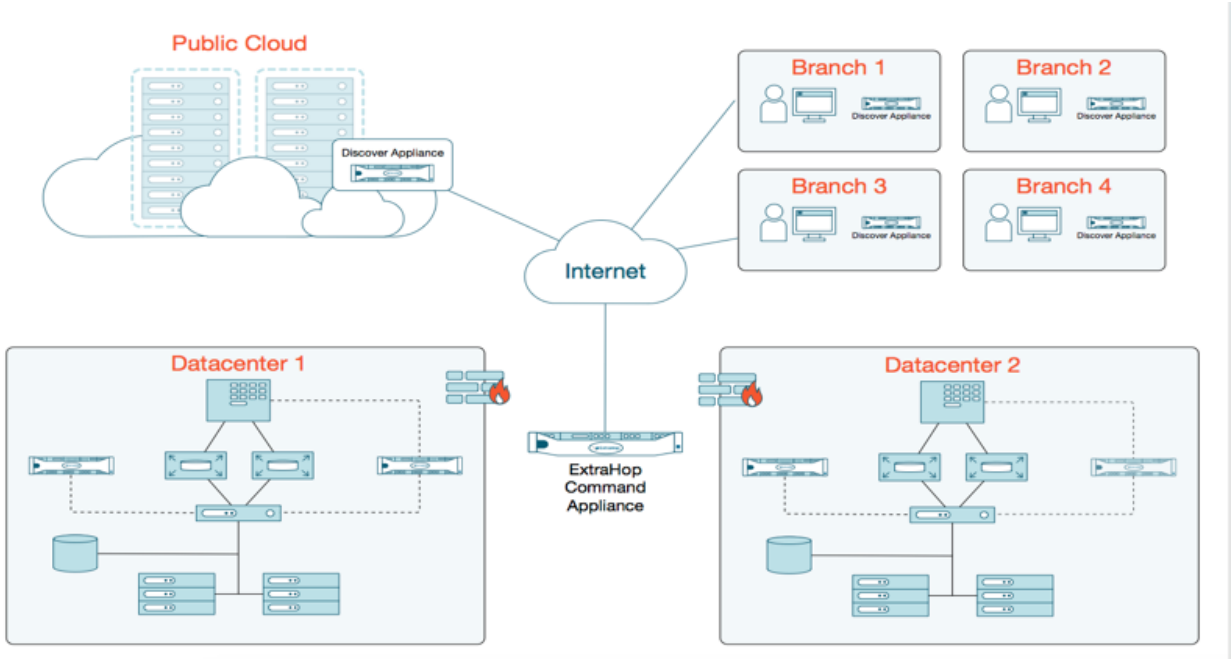

Figure 6 Multiple Datacenter Deployment

See the Packet aggregation, multiple ExtraHop Discover appliances example for a discussion of how the ExtraHop Command appliance aggregates metrics from multiple ExtraHop Discover appliances but does not combine those metrics.

Datacenter and remote site

A common architecture is one main datacenter which houses centralized applications used by many offices, stores, depots, or other remote locations. In some cases, it's possible to effectively monitor the network and applications in each remote location by placing an ExtraHop Discover appliance in the central datacenter, because all traffic from all remote locations comes back to the datacenter. However, this is not always the case.

This can be a useful application for ExtraHop's EDA1100 small form factor appliance if the remote sites' resource requirements are not excessive.

For use when

- There is traffic between tiers in the remote sites that cannot be seen without looking at the network in the remote site.
- Not all traffic from the remote sites returns to the datacenter where it can be monitored; some exits to the public internet directly from the remote site.

- Security requires monitoring of all network attack surfaces.

Reasons not to use

- If it's not necessary to see the network traffic in the remote site because all of that traffic comes to the datacenter, simply take the traffic from the datacenter and use Remote Networks or Custom Devices to discover the devices on the remote site's network.
- If the traffic from the remote network can be encapsulated and streamed back to ExtraHop in the datacenter, i.e., if there is very little traffic in the remote site.

Virtual deployment

VMWare ESXi technology provides a high-performance hypervisor environment that increasingly replaces physical computing environments for a variety of tasks. ExtraHop Discover appliances can themselves be deployed into VMWare hosts, provided they are configured so that they can not be automatically moved to a different host and provided they are assigned the required resources (including the dedication of at least 1 CPU).

Virtual EDAs source traffic differently than physical hosts because they lack physical capture ports. They can receive Ethernet feeds from the virtual switch (e.g., the VMMare virtual distributed switch provides a virtual bridged Ethernet interface); from a virtual connection to a hardware port to consume a feed directly; from a virtual connection to a hardware port to act as an RSPAN or ERSPAN target; or from RPCAP software taps deployed to virtual guests or elsewhere.

Amazon Web Services Deployment

The virtual ExtraHop Discover appliances monitor the AWS environment in real-time and can help you discover much richer intelligence than the limited machine data provided by CloudWatch. ExtraHop can show you hundreds of metrics, including full L2-L7 analysis, visibility into key performance indicators like application latency, EC2 processing time, RDS methods, and per-bucket S3 file access.

ExtraHop makes available several different AMIs deployable to customer AWS instances to support different numbers of devices. Multiple copies can be installed in a single region or across regions to provide analysis for more devices. ExtraHop Command appliance for AWS is an optional free AMI to deploy when you have multiple ExtraHop Discover appliances and want to monitor metrics gathered by different EDAs together in one UI.

The ExtraHop Packet Forwarder allows you to monitor instances using RPCAP. RPCAP forwards traffic from virtual environments without requiring a virtual tap. The usage of RPCAP removes the need to modify the virtual network-switching configuration within the VM environment. This is specifically required for operation within the AWS EC2 service.

The typical configuration for virtual ExtraHop Discover appliances is shown in Figure 7 Figure 7 below. The IP data stream from multiple AWS VMs is forwarded by the ExtraHop Packet Forwarder to the ExtraHop appliance using RPCAP.

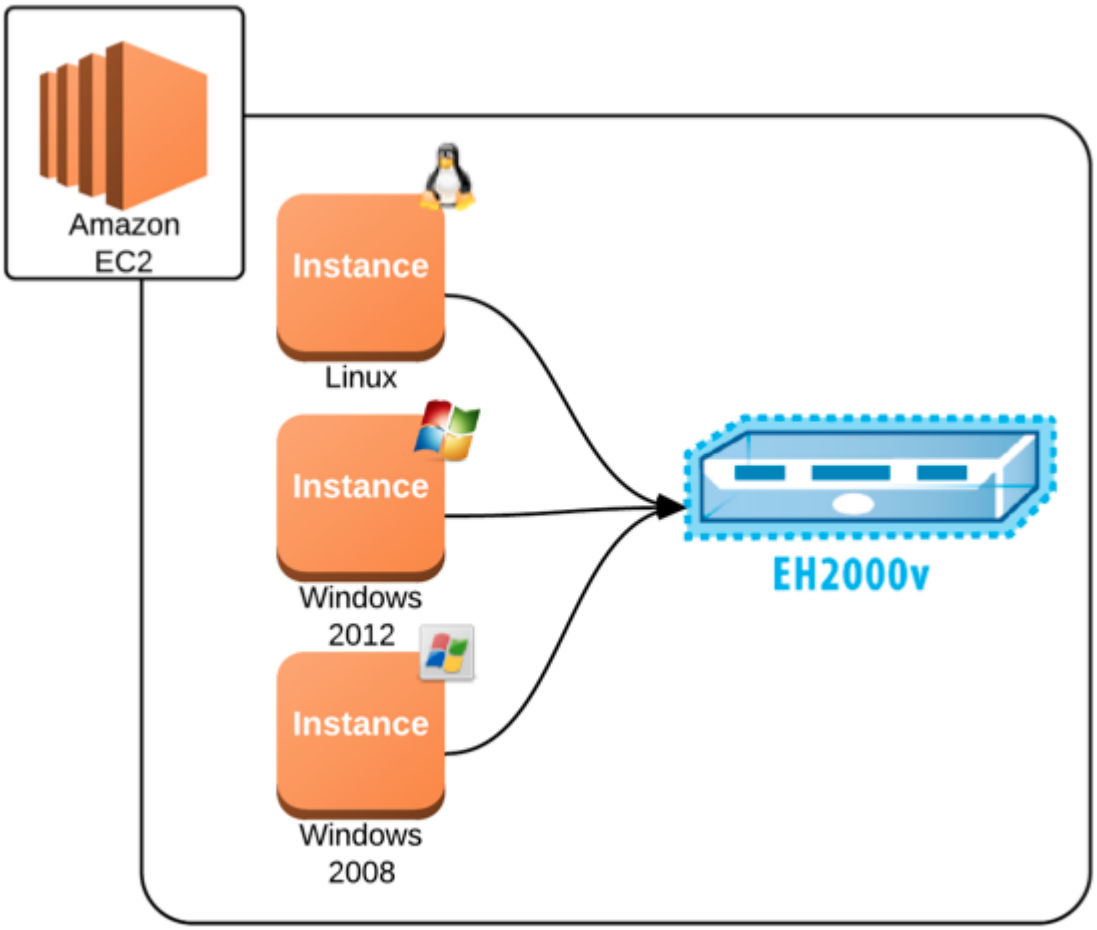

Figure 7 General ExtraHop Installation within AWS

Regions and availability zones are key concept to AWS deployments. AWS regions are geographically defined subsets of the overall AWS network. Each region is operationally isolated from all other regions. Often AWS deployments are selected to use specific regions due to the proximity of user or customer facilities.

Within each region are multiple availability zones (AZs). An AZ can be viewed as a virtual datacenter. Each AZ within a region is connected via low-latency, high-capacity data links. This allows VMs to be distributed across multiple AZ's but function as a single integrated system.

Options for using ExtraHop virtual appliances in support of customer AWS deployments vary based primarily on the use of single or multiple AZs. Figure 9 illustrates the scenario where all deployed AWS VMs being within a single AZ. In this case, a single ExtraHop virtual appliance, deployed in the same AZ, receives wire data from all customer VMs. This arrangement is beneficial since traffic inside of an AZ is free, and the

encapsulated packet data between traffic collection points and the ExtraHop virtual appliance can be considerable.

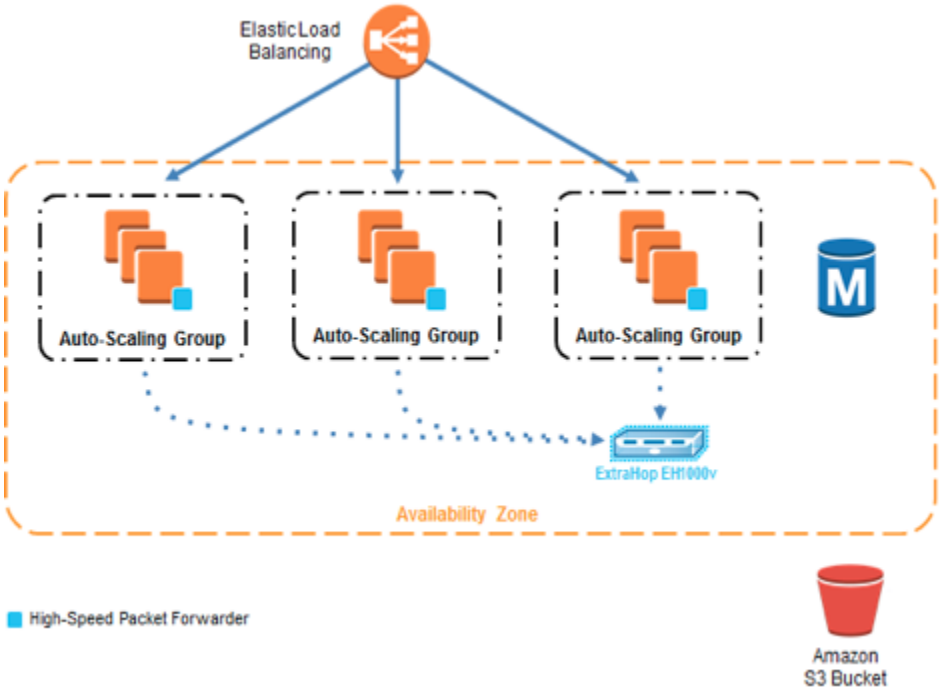

Figure 9 ExtraHop Deployment within a single AWS AZ

For increased availability, it is common that customer VMs are spread across multiple AZs within a region. This configuration takes advantage of the data connectivity between the AZs, as well as providing physical diversity for fault tolerance. In this case, it may be preferable to deploy multiple ExtraHop virtual machines so that one sits in each AZ and can catch all the traffic collected in its AZ.

Contact

ExtraHop Solutions Architecture has deployed the ExtraHop platform hundreds of times into every type of environment imaginable. If this guide gives you ideas that you'd like to talk about, we'd love to speak with you. Please contact your sales representative, or contact us directly using the information below.

ExtraHop Networks, Inc.
520 Pike Street, Suite 1700
Seattle, WA 98101 USA

www.extrahop.com
info@extrahop.com
T 877-333-9872
F 206-274-6393

Customer Support support@extrahop.com
877-333-9872 (US)
+44 (0)845 5199150 (EMEA)