



ExtremeAnalytics[®] Deployment Guide

05/2023
23.04.11
PN: 9037825-00
Subject to Change Without Notice



Copyright © 2023 Extreme Networks, Inc. All Rights Reserved.

Legal Notices

Extreme Networks, Inc., on behalf of or through its wholly-owned subsidiary, Enterasys Networks, Inc., reserves the right to make changes in specifications and other information contained in this document and its website without prior notice. The reader should in all cases consult representatives of Extreme Networks to determine whether any such changes have been made.

The hardware, firmware, software or any specifications described or referred to in this document are subject to change without notice.

Trademarks

Extreme Networks and the Extreme Networks logo are trademarks or registered trademarks of Extreme Networks, Inc. in the United States and/or other countries.

All other names (including any product names) mentioned in this document are the property of their respective owners and may be trademarks or registered trademarks of their respective companies/owners.

For additional information on Extreme Networks trademarks, please see:
www.extremenetworks.com/company/legal/trademarks/

Contact

If you require assistance, contact Extreme Networks using one of the following methods.

- [Global Technical Assistance Center \(GTAC\) for Immediate Support](#)
 - **Phone:** 1-800-998-2408 (toll-free in U.S. and Canada) or 1-603-952-5000. For the Extreme Networks support phone number in your country, visit:
www.extremenetworks.com/support/contact
 - **Email:** support@extremenetworks.com. To expedite your message, enter the product name or model number in the subject line.
- [GTAC Knowledge](#) — Get on-demand and tested resolutions from the GTAC Knowledgebase, or create a help case if you need more guidance.
- [The Hub](#) — A forum for Extreme customers to connect with one another, get questions answered, share ideas and feedback, and get problems solved. This community is monitored by Extreme Networks employees, but is not intended to replace specific guidance from GTAC.
- [Support Portal](#) — Manage cases, downloads, service contracts, product licensing, and training and certifications.

Table of Contents

ExtremeAnalytics® Deployment Guide	1
Table of Contents	3
About This Guide	4
Introduction	5
Deployment Overview	5
Deployment with NetFlow and First N Mirror (Legacy devices only)	6
Deployment with XIQ Controller sending enhanced IPFIX	7
Deployment with Application Telemetry based on ACL mirror and sFlow	8
Deployment with IPFIX and K-Mirror	9
Deployment with Traffic Sensor and Raw Data Analysis	10
ExtremeAnalytics Deployment Summary	10
Deployment Requirements	11
Designing Your ExtremeAnalytics Deployment	12
Home Engine for Multiple Endpoint Locations and Traffic Domains	12
Traffic Sensor Deployment and Multiple Monitored Points	13
Deployment with K-mirror and IPFIX	13
Deployment with Fabric Connect	13
Deployment in the Core Network	14
Deployment in the DMZ	14
Common Flow Collection Scenarios	14
Unidirectional Flows	14
Duplicate Flows	15
Asymmetric Routing	16
Network Load Balancing	16
Jumbo Frames	16
WAN links	17
ExtremeAnalytics Engine Deployment	17

About This Guide

This document describes how to design your ExtremeAnalytics solution deployment.

This document is intended for experienced network administrators who are responsible for implementing and maintaining communications networks.

Introduction

This section provides an overview and lists the requirements for your ExtremeAnalytics deployment.

Deployment Overview

There are three components for the ExtremeAnalytics deployment:

ExtremeCloud IQ - Site Engine (Site Engine)

In the context of ExtremeAnalytics, the Site Engine provides a GUI for monitoring and managing the Application Analytics Engine or Application Analytics Traffic Sensor, or both. Site Engine provides long-term storage of the aggregated flow data.

Application Analytics Engine (Analytics Engine) or Application Analytics Traffic Sensor (Traffic Sensor)

Both engines process information from the network infrastructure. On the engine, the application stream is assembled and fingerprints are applied to identify the applications. Engines provide a real-time flow cache to the Site Engine.

The complete record of IP and TCP/UDP information, application name and category, network and application response times, and application metadata is sent to the Site Engine for graphing and storage.

Network Infrastructure Device (Switch, or Wireless Controller)

Network infrastructure Devices provide data to Application Analytics Engine. To scale properly, the engine does not analyze the complete network traffic. The network infrastructure can pre-process the traffic by either calculating the NetFlow/IPFIX or selecting the interesting traffic by ACLs.

- Unsampled NetFlow or IPFIX provides an accurate statistical representation of all flows mirrored for application identification
- The first 15 packets of each flow, through a forensic policy mirror (First N Mirror or K-mirror)


Mirrored traffic can be delivered to the Application Analytics Engine by:

- A local traffic mirror.
- Remote mirroring through GRE or ERSPAN L2 tunneling.

NOTE: Remote mirroring allows for a single Application Analytics Engine to receive traffic feeds from multiple switches in the network without being directly connected to all Network Infrastructure Devices.

The following sections show a simplified architecture and information flow in ExtremeAnalytics deployments.

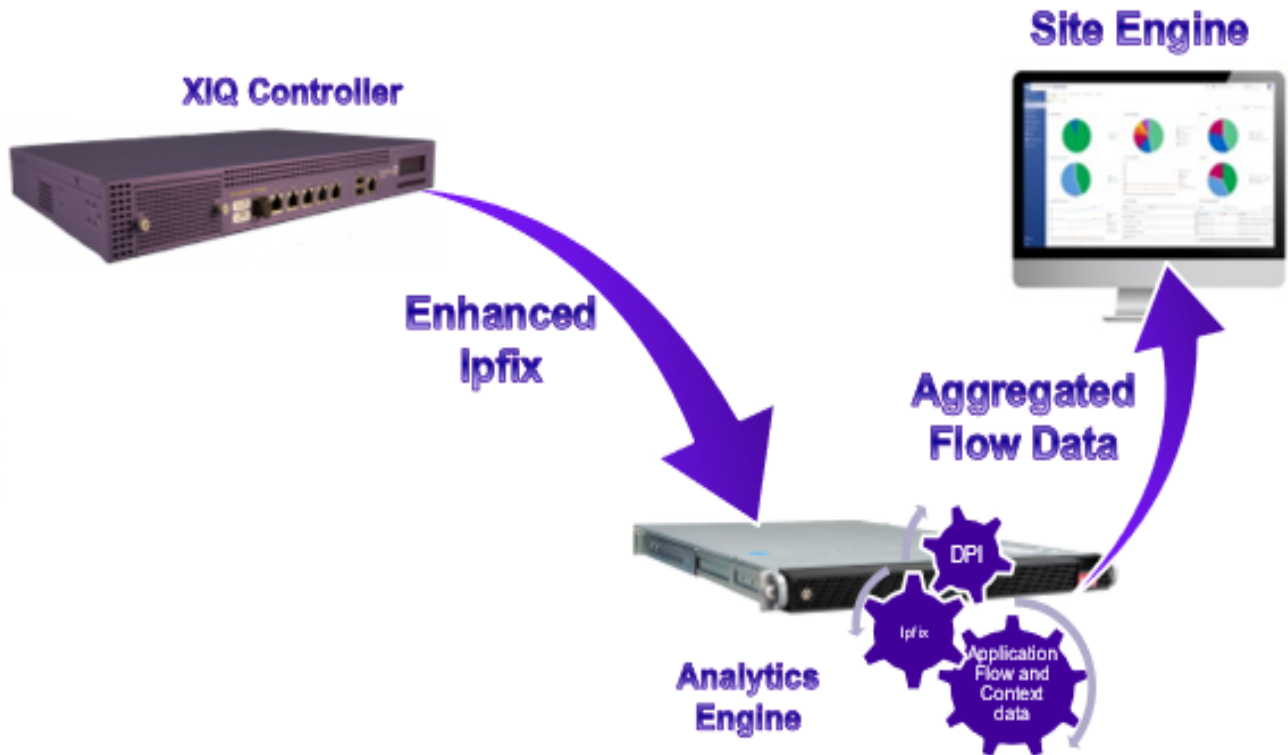
Deployment with NetFlow and First N Mirror (Legacy devices only)



Deployment with NetFlow and First N Mirror

The Network Infrastructure Device sends unsampled NetFlow to the Application Analytics Engine with the first 15 frames from each flow. The Application Analytics Engine calculates the L7 information based on packets from the First N Mirror, and combines that information with the unsampled NetFlow records.


Deployment with XIQ Controller sending enhanced IPFIX



Deployment with XIQ Controller sending enhanced IPFIX

Wireless Controller aggregates statistics and metadata from access points and sends enhanced IPFIX records to the Application Analytics Engine.

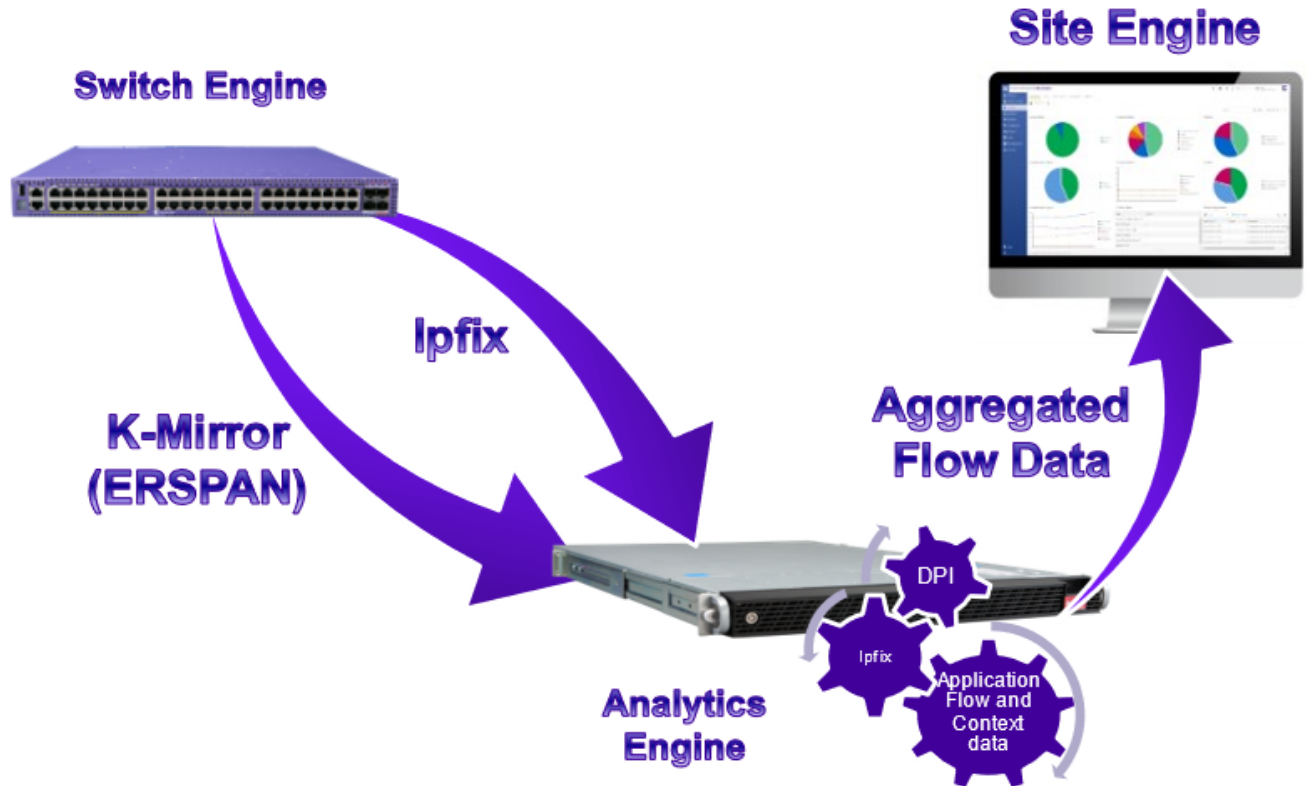
Deployment with Application Telemetry based on ACL mirror and sFlow



Deployment with Application Telemetry based on ACL mirror and sFlow

ExtremeCloud IQ - Site Engine configures the ACL mirror in the network infrastructure. The ACLs carry the defined interesting traffic for transfer to the Application Analytics Engine. The Application Analytics Engine calculates the L7 information based on packets from the ACL mirror. The traffic volume information is derived from the sampled sFlow information. Each sFlow record is multiplied by a sampling rate to estimate the volume information. For example, two sFlow packets for the same flow with a sampling rate of 1:1024 translates to 2048 packets for the flow.

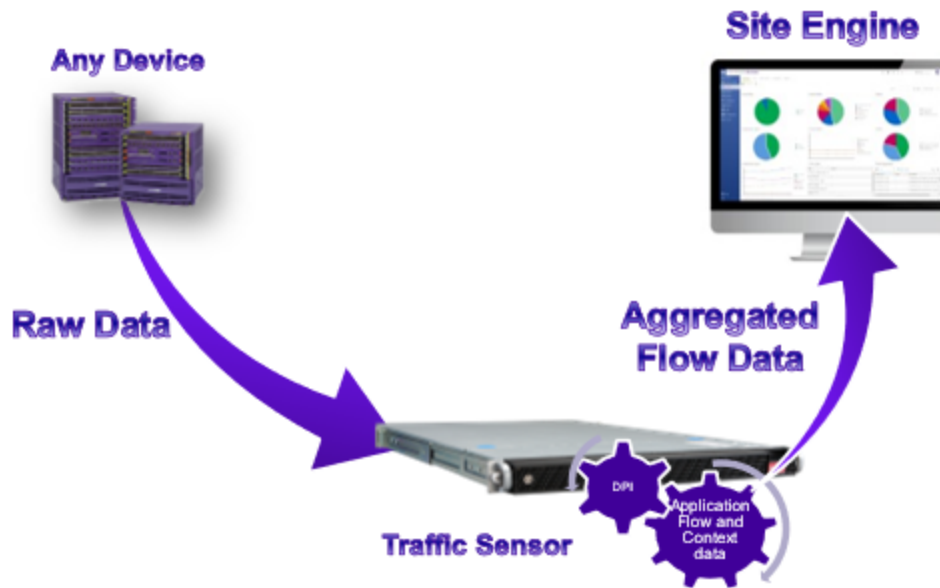
Deployment with IPFIX and K-Mirror



Deployment with IPFIX and K-mirror

The Network Infrastructure Device sends unsampled IPFIX to Application Analytics Engine together with the first 15 frames from each flow. The Application Analytics Engine calculates the L7 information based on packets from the K-mirror and combines that information with the unsampled IPFIX records.

Deployment with Traffic Sensor and Raw Data Analysis



Deployment with Traffic Sensor and Raw Data Analysis

The Network Infrastructure Device mirrors the raw data to the Application Analytics Traffic Sensor. The Application Analytics Traffic Sensor processes and calculates all information based on the raw data.

ExtremeAnalytics Deployment Summary

The following table summarizes the main differences between the ExtremeAnalytics deployment options:

	NetFlow and First N Mirror	XIQ Controller	ACL Mirror and sFlow	IPFIX and K-Mirror	Traffic Sensor and Raw Data
Network Response Time calculation	Yes	Yes	Yes	Yes	Yes
Application Response Time calculation	Yes	Yes	Yes	Yes	Yes

Available on cost-optimized edge	No	No	Yes	No	No
L7 analysis for most common applications	Yes	Yes	Yes	Yes	Yes
Available in hardware of commercial silicon chips	No	Extreme controller	Yes	Yes	Barebone server
All flows transferred through device reported	Yes	Yes	Limited visibility to UDP traffic	Yes	Yes but not inline
Advanced fingerprints identify the application	Yes	No	No	Yes	No
Accurate volume calculation	Yes	Yes	Calculation based on sampling	Yes	Yes
Resource demanding	No	No	No	No	Yes

Deployment Requirements

Deploying an ExtremeAnalytics solution requires the following:

- Network Infrastructure Devices with compatible firmware. You can find the tested and supported platforms in the Extended Firmware Support document:
For additional information, see [Extended Firmware Support](#).
- Application Analytics Engine or Application Analytics Traffic Sensor — The Application Analytics Engine is available to download as a virtual appliance for Microsoft Hyper-V and VMware ESXi. The hardware specifications required for a barebones server deployment is available for both the Application Analytics Engine and Application Analytics Traffic Sensor.

NOTE: Virtual appliances are bandwidth-bound by the underlying host interface bandwidth. If the host interface can operate at 10Gbps, you can reassign the virtual interface to a new hardware interface without making any changes in the Application Analytics Engine.

- ExtremeCloud IQ - Site Engine management — Beyond the configuration and monitoring of the ExtremeAnalytics solution, Site Engine provides the long term storage and reporting, presenting the correlated data with contextual information. ExtremeCloud IQ - Site Engine can also provide that correlated and contextual data to other IT systems through the Connect module.
- Licenses:
 - XIQ-PIL-S-C subscription license is mandatory for each component: ExtremeCloud IQ - Site Engine, Application Analytics Engine, and Application Analytics Traffic Sensor.
 - XIQ-PIL-S-C or XIQ-NAV-S-C subscription is required for each Network Infrastructure Device. For additional information, see [ExtremeCloud IQ - Site Engine Licensing](#).

Designing Your ExtremeAnalytics Deployment

To design a reliable and accurate ExtremeAnalytics deployment, consider the following:

- [Home Engine and Endpoint Locations and Traffic Domains](#)
- [Application Analytics Traffic Sensor Deployment and Multiple Monitored Points](#)
- [Common Flow Collection Issues](#)
- [Application Analytics Engine Deployment](#)

Home Engine for Multiple Endpoint Locations and Traffic Domains

For scaling purposes, you can deploy multiple Application Analytics Engines or Application Analytics Traffic Sensors, or both. If there are multiple engines, the flow might be seen by more than one engine and you must apply a Home Engine to ensure that ExtremeAnalytics does not calculate the same traffic multiple times.

The Home Engine is defined on the site level. The flow is aggregated to the long-term storage from the Home Engine only.

Endpoint Locations are defined on the site level. Only flows with IP subnets defined in Endpoint Locations are stored for the long-term storage of aggregated flow data in ExtremeCloud IQ - Site Engine. If the IP address matches multiple records in the Endpoint Locations, then the best match is used.

To aggregate the flow to the long-term storage you must ensure the IP address is defined in the Endpoint Location and that the Home Engine is applied.

Use the following guidelines when creating your mapping sites to the Home Engine:

- Physical network layers: Such as Edge, Distribution, Core.
- IP network boundaries: Such as Internet, intranet, and DMZ. Network boundaries are easily identified and provide monitored points for the ExtremeAnalytics engine.
- Traffic domains: Such as sales and R&D. Isolating traffic from one functional domain or another can be difficult. To simplify your deployment, identify a single port for each functional domain to use for traffic monitoring.
- NAT boundaries: Use NAT boundaries as ExtremeAnalytics traffic domain boundaries.

NOTE: A NAT boundary inside a traffic domain can lead to unexpected results since the ExtremeAnalytics engine cannot identify the original flow and the NAT flow. If a traffic domain contains the original flows and the NAT flows, the flows are counted twice for the two different source addresses.

Traffic Sensor Deployment and Multiple Monitored Points

You must establish monitored points in the Network Infrastructure Device to provide a reliable and accurate traffic sample to an Application Analytics Traffic Sensor. In a reliable and accurate traffic sample, the statistics obtained are the same statistics that would be obtained from all the traffic in the domain. Ensure that your traffic samples meet the following requirements:

- Every flow in the traffic domain must be represented in the sample. To do this, you must carefully plan the points where the traffic is sampled or mirrored.
- Every flow in the traffic domain must appear only one time in the sample. Your deployment must avoid traffic samples that contain multiple copies of the same flow.

To minimize the collection of unidirectional or duplicate flows:

- Map all ingress (RX) pathways where network traffic can enter into a particular traffic domain.
 - The ingress pathways can traverse multiple switches and therefore will usually require a detailed and accurate network diagram to pinpoint them.
 - If any ingress ports to the traffic domain are left out of the deployed configuration, the result is inaccurately tagged unidirectional flows.
- Ensure that any multi-pathed traffic is delivered to the same Application Analytics Traffic Sensor.
- Do not include intra-traffic domain pathways built for redundancy between switches as these links can cause duplicate flows.
- Configure the ingress network pathways isolated during the planning process to capture only the inbound traffic to the port. For each switch involved in a traffic domain, ensure that the traffic mirror on the isolated ports is enabled in the RX direction only.

For scenarios that can lead to traffic capture issues, see [Common Flow Collection Scenarios](#).

Deployment with K-mirror and IPFIX

Extreme Networks universal switches running the Switch Engine network operating system can provide IPFIX statistics and traffic mirroring capabilities on all ports.

Deployment with Fabric Connect

Extreme Networks switches running Fabric Connect can provide ACL mirror and sFlow capabilities on all UNI (fabric edge) ports. NNI ports (fabric internal links) are excluded. The Application Analytics Engine must see both directions of each flow to deliver full details. Ensure the **Enable Fabric Mode** is checked in the Application Analytics Engine configuration.

If the Fabric Mode is enabled, the packet timestamps for both directions of a flow are stored in a global cache. If the Fabric Mode is disabled, the packet timestamps for both directions of a flow are stored in the cache on a per-switch basis. The calculation of network response time and application response time in a fabric deployment is most reliable with Fabric Mode enabled.

When running VOSS/Fabric Engine version 8.8 or later the telemetry source IP address can be either a Management VLAN, a Management CLIP - GRT, or a Management CLIP - VRF. When adding older versions of VOSS/Fabric Engine to the ExtremeAnalytics, a Management CLIP - GRT can only be used as a Telemetry Source IP address.

Deployment in the Core Network

The best practice for ExtremeAnalytics is to analyze the traffic gathered at the edge of the core network, as close to the client as possible. If the edge of the network does not support ExtremeAnalytics, then you can deploy in the upper layers. Deploying ExtremeAnalytics at multiple layers in the core network can provide additional flow path reporting.

Deployment in the DMZ

The DMZ traffic domain is an independent network area, usually configured with VLANs or Network Infrastructure Devices that connect with the rest of the network through firewalls. These firewall interfaces can provide ExtremeAnalytics with the required traffic sample for all ingress and egress traffic. Consider that most of the traffic in the DMZ is with systems and clients outside of the DMZ. If you expect a large amount of intra-DMZ traffic, you must adapt the traffic sampling strategy to account for the intra-DMZ traffic and replicate the strategy used in an edge or core network deployment.

Common Flow Collection Scenarios

Ensure that your ExtremeAnalytics deployment accounts for the following potential scenarios:


- [Unidirectional Flows](#)
- [Duplicate Flows](#)
- [Asymmetric Routing](#)
- [Network Load Balancing](#)
- [Jumbo Frames](#)
- [WAN links](#)

Unidirectional Flows

By default, the Extreme Networks switches are configured by ExtremeCloud IQ - Site Engine to forward the data to Application Analytics Engine for all ingress ports. The manual configuration can cause unidirectional flows if not deployed properly.

In a fabric deployment, only UNI ports (at edge of the fabric) support ExtremeAnalytics. To avoid unidirectional flows issue, all fabric edge switch ports must report data to the same Application Analytics Engine.

The following figure shows three ports configured on a switch. When a user accesses the Internet, the traffic is captured by port 1 on the way out to the Internet and at port 3 for the return traffic. The same engine must see traffic from both ports.




Unidirectional Flows

Duplicate Flows

Duplicate flows (most common for Application Analytics Traffic Sensor deployment) occur when the same flow is seen at multiple monitored points included in an ExtremeAnalytics traffic domain.

In the following figure, User A's traffic, which is logged one time on each path to and from the Internet, is counted correctly. User B's traffic is seen twice on the way to the Internet and one time on the return path from the Internet. User C's traffic is logged twice in both directions. To avoid duplicate flows, limit the collection locations or create multiple traffic domains to keep the duplicate information separated.



Duplicate Flows

Asymmetric Routing

Asymmetric routing occurs when a packet takes a path from source A to destination B, but then the return packet from B takes a different path back to A. Because there are no perceived performance issues associated with asymmetric routing for normal applications traveling across the network, network administrators might be unaware that asymmetric routing is occurring. The issue becomes evident when ExtremeAnalytics is enabled for a specific path and only one half of a TCP conversation is seen, which causes ExtremeAnalytics to produce erroneous results.

Network Load Balancing

If your network includes a load balancing configuration, ensure that all necessary links are covered by the ExtremeAnalytics solution. If covering all links is impossible because of physical constraints or possible flow duplication, your ExtremeAnalytics deployment can require collection at specific network choke points.

Jumbo Frames

Enable Jumbo frames on the whole path from the switch to the Application Analytics Engine if any of the following is deployed:

- First N Mirror
- K-mirror
- ACL mirror
- sFlow.

If the maximum-sized frame is encapsulated to ERSPAN or GRE or sFLOW then the new frame size increases, and without Jumbo Frames enabled the frame is dropped.

WAN links

Due to the natural unbalance in throughput of LAN vs WAN, it is not recommended to transfer the First N Mirror / K-mirror / ACL mirror data through the WAN link. If the local traffic inside the branch needs to be analyzed by ExtremeAnalytics then consider the deployment of a local Application Analytics Engine or Application Analytics Traffic Sensor in the branch.

ExtremeAnalytics Engine Deployment

The ExtremeAnalytics engine supports multiple deployment modes to support different network environments and connectivity characteristics:

- Single Interface — A single interface is configured for both management and First N Mirrored traffic through GRE. You must configure a GRE tunnel for traffic monitoring.
- Single ERSPAN — A single interface is configured for both management and ACL mirror or K-mirror. Both ACL mirror and K-mirror use ERSPAN.
- Dual Tap Mirror-N — Separate interfaces are configured for management and First N Mirrored traffic goes directly to the interface. The monitoring interface uses tap mode for traffic monitoring.
- Dual Tunnel Mirror-N — Separate interfaces are configured for management and First N Mirrored traffic through GRE. The monitoring interface uses a separate IP address. You must configure GRE tunnels for traffic monitoring.
- Manual Mode - The interfaces are not configured by the script. You must manually assign each interface for the required role.