



# XCA EDGE

## FAQ





# Summary

<b>OPERATION</b> .....	2
<b>PROBING</b> .....	2
<b>DDOS</b> .....	4
<b>BLACK OUT DETECTION:</b> .....	4
<b>DISTRIBUTED XCA EDGE</b> .....	5
<b>API</b> .....	5
<b>NATURAL BGP</b> .....	5
<b>PERFROMANCE / COST OPERATING MODES</b> .....	5
<b>IPV6 SUPPORT</b> .....	6



## OPERATION

### **1. Does XCA Edge act as a Route Reflector Client collecting data from the customer network through the system collector? Can XCA Edge be configured to collect data via Port Span, Net Flow and SFlow?**

The BGP session between customer routers and the XCA Edge collector is configured as a route reflector client. This way XCA Edge collects all BGP routes. This has no relation to flow data, though, which is handled by entirely different protocols (netflow/sflow/jflow). XCA Edge needs both: a BGP session (route reflector client) to each of the customer routers and a live export of sampled traffic through netflow or sflow. XCA Edge does not support direct mirroring nor port spanning.

### **2. After collecting the data, will XCA Edge send the data to the controller where the top prefixes carrying most of the traffic are selected for probing?**

The collected data is sent from the collector to the controller as an aggregated/compressed set.

### **3. Can the customer adjust the maximum number of selected prefixes in the Controller configuration?**

Yes. There is a hard limit that Expereo will control through the license, and then there is a second limit that is controlled by the (user-side) XCA configuration.

## PROBING

### **1. Will the Controller probe the prefixes through the customer's ISPs using ICMP, UDP and TCP\_SYN probes?**

XCA Edge does TCP stealth half-handshakes as this is the only generic and reliable way to measure latency towards unknown systems.

- ICMP we don't do, because ICMP doesn't have ports, and we rely on ports for our PBR configuration.
- UDP we don't do, because it would require us to implement the exact UDP protocol required for communication, but then the RTT would not reflect the network state anyway, but mostly the application's response time on the remote server, which is no good to us.
- We do, however, provide MANUAL UDP-based SIP and DNS probing subsystems within XCA Edge. XCA Edge won't ever use them for its own needs, but for a customer that has some exotic probing needs they can configure SIP or DNS probes themselves through the GUI, and attach these probes to some of the destinations.

We also provide the possibility of configuring manual ICMP probes, but since these are port-less by nature, such probes are limited to testing only the current path towards their destination, and won't be able to provide visibility over alternative paths.

### **2. Is the controller's behaviour configurable? Can we configure how often the probes should be sent by setting the probing period?**



Yes. In the menu probing/engine/settings you will find a section called 'automatic probes', with two parameters: "autoprobe frequency" and "autoprobe tolerance". Both these settings have help messages explaining their behaviour.

### **3. Can the customer adjust the probing algorithms according to their needs?**

The only probing system that we are able to leverage in a fully automated way is the TCP half-handshake method (SYN/SYN-ACK/RST). Other probings are possible (DNS, SIP), but need to be configured manually on a per-destination basis by the customer.

### **4. Will an IP alias be configured on the platform for each ISP?**

Not necessarily. We can use a different IP address for probing through each transit, but since XCA Edge v2.0 we can just as well use a single IP for all, since we rely on source ports now for PBR. This development enabled us to support probing IX points.

### **5. Is the PBR configured on the Edge Routers for probing originating from the IP's to be routed over different transit providers?**

The PBR is configured on the edge routers for routing XCA Edge probes over different providers based on the source port of the probes. So when XCA Edge sends a probe with port 'x', it knows exactly what transit provider (or IX peer) this probe will go through, since it knows the PBR that have been deployed on routers.

### **6. After probing each prefix, are the performance metrics returned to the controller, which will calculate the best available path/provider?**

Yes, all computations are performed by the controller, while all probing is done by the collector. At 'init time', the controller sends orders to the collector "please probe destination x.y.z every t seconds", and then the collector does it continually and reports back results to the controller in real time.

### **7. As soon as the controller identifies the best provider for a specific prefix, will the platform's BGP DAEMON inject it into the Edge Router as a regular BGP announcement with an updated next-hop?**

Yes, this is how it works. The actual order may be delayed, depending on the history of this new path - typically we avoid putting traffic on paths that were known to be unstable during the last few minutes, and we prefer using slightly less optimal paths during a grace period to make sure the path has stabilized.

### **8. Can a customer specify a list of prefixes or ASNs that will be probed by XCA Edge much more often than other destinations, in case the customer has critical business destinations or customers that they would like XCA Edge to pay more attention to?**

If there are important destinations, then the customer can create manual probes in XCA Edge (with custom timings) and attach these probes to the 'important' destinations.



## DDOS

### 1. How is the DDoS mitigation done?

XCA Edge provides DDoS detection and automation, not mitigation. XCA Edge can be used as a triggering mechanism within a wider anti-DDoS mitigation policy.

We have at least two different scenarios working in our current customer base:

a) Customer X has a DDoS mitigation service with one of their transit providers. When XCA Edge detects a DDoS, it then reconfigures the customer's routers so they apply an AS-PREPEND to all transits, but not to the one transit that provides DDoS protection. This way, during a DDoS all inbound traffic is forced to flow in through the transit with DDoS protection.

b) Customer Y subscribed to a remote (cloud-based) scrubbing center. When XCA Edge detects a DDoS, it sends a UDP signal to the remote scrubbing center, and this center starts announcing the customer's prefixes, taking in all the customer's traffic and cleaning it. Then, the clean traffic is fed to the customer through a GRE tunnel.

## BLACK OUT DETECTION

### 1. Does this feature allow XCA Edge to automatically disconnect a BGP session with a provider suffering from excessive levels of packet loss, after considering configurable preferences set by the user?

Yes, if Blackout Detection is enabled.

### 2. When a Black Out issue is detected, will XCA Edge send alerts to the customer?

Yes, if configured to do so (and if Blackout Detection is enabled).

### 3. Can it be configured to automatically drop BGP sessions with the providers that suffer from excessive Packet loss?

Yes. The trigger to defining what a "blackout" is, is configurable in term of % of packet loss. Note that we recommend performing AS-PREPENDs instead of BGP session shutdown.

### 4. Does XCA Edge continue to monitor these providers?

Yes. But if the BGP session was shut down and the transit providers does some kind of ingress filtering, then the probing traffic might be blocked out. This is one of the arguments in favour of performing AS prepends rather than shutting sessions.

### 5. Does XCA Edge restore BGP sessions once the loss averages return to normal?

Yes. It's all configurable through custom scripts attached to XCA Edge events (this is all part of a wider framework we call "Smart Actions").



## DISTRIBUTED XCA EDGE

**1. Does this mean that one single XCA Edge can optimize traffic across multiple physical locations if they are relatively close to each other and the connection between them is stable and predictable?**

Yes. XCA Edge doesn't care about geographical distribution. Hence whether two edge routers are sitting close to each other or are separated by 1000 kms of black fiber doesn't make any difference.

**2. Can XCA Edge gather traffic information across multiple points of presence and probe available destinations and re-route traffic accordingly?**

Yes, although in such situations we usually implement a single XCA Edge controller, and as many collectors as there are POPs. Then, POPs act as "virtual transits" to each other. We call this "inter-dc links".

## API

**1. Is XCA Edge equipped with an API that allows the customer to integrate our analytics to any third-party applications?**

All data is available through SQL queries (if the SQL db is enabled to be listening for external queries). We also provide a web API to perform configuration actions and poll basic states of the platform.

## NATURAL BGP

**1. Does Natural BGP back up the main XCA Edge platform in case of operation failure?**

Yes, the entire BGP platform will simply revert to its legacy working mode, without optimization, until XCA Edge is restored.

**2. If it happens, will XCA Edge simply withdraw its announcements and the routing decisions be handed back to BGP?**

If there is failure, the BGP sessions between XCA Edge and routers will go down, meaning that the routers will immediately discard all orders (optimisations) performed by XCA Edge.

## PERFORMANCE / COST OPERATING MODES

**1. Can XCA Edge operate in two different modes (Performance and Cost)?**

Yes, it's all a matter of configuring triggers to set the balance of optimisation the customer requires.



## **2. Does the Performance Improvement Mode choose the best path/provider in terms of Packet Loss and Latency?**

The RDE (the Routing Decision Engine) can be configured to favour performance above costs, yes.

## **3. Does the Cost Improvement Mode choose the cheapest provider to reach a destination while making sure the performance is kept within acceptable levels?**

Yes.

## **IPV6 SUPPORT**

### **1. Does XCA Edge support IPv6 for Juniper and Cisco?**

Xca Edge supports IPv6 in the same manner as IPv4 for all supported devices.

1. Do we offer better SLA on XCA vs standard DIA
  - Service Level Guarantees are included on service availability and time to repair.
2. What is the value add of XCA for SD-WAN
  - XCA and SD-WAN are a perfect marriage. XCA optimizes the underlay routing destination based. This optimized underlay route is used by the SD-WAN Overlay tunnels between the SD-WAN devices.
3. How does XCA move traffic to good path (ie how do we route)
  - Out of each of the Expereo Cloud Acceleration Hubs (CXH) all customer traffic destination are learned and the respective prefixes probed across all available IP transits and IP Peers. This data is analyzed and the best available paths calculated and injected into the BGP routers. The XCA choice is destination performance based.
4. How fast can XCA change path to the secondary better path
  - Path analysis, calculation and BGP router injection is done under 2 minutes
5. Can there be a conflict between SDWAN and XCA as they are 2 independent dynamic systems ie counter acting.
  - No, XCA is an underlay service and SD-WAN an overlay. SD-WAN creates tunnels between SD-WAN (virtual) devices across an underlay. Optimizing the underlay enables a better performing SD-WAN



7. What technology do we use for XCA
  - The XCA platform is powered by the XCA Edge software. This is a propriety software package with Expereo developed algorithms capable of analyzing the probe statistics and calculating the best performing path to each prefix on the public internet (600,000 known routes today and increasing)
8. Can a carrier or customer owned internet circuits be tunneled to Expereo hub to use XCA
  - Yes, the carrier or customer can use GRE or IPSec tunnels to reach the XCA Hubs. Alternatively a carrier or customer can consider a direct interconnect with the XCA platform at the main DC or Hub site as the optimized gateway to the public cloud applications or SD-WAN underlay.
9. Is XCA using a private IP or MPLS backbone globally?
  - No, unlike other global cloud optimization provider, the Expereo XCA platform is not limited to a private IP or Ethernet back-bone. XCA is an SDN (Software defined network) that runs on top of the public internet. This allows the SDN to use any available path of the internet itself, making it the largest and most flexible cloud optimization platform existing today