Understanding On-premises Device, Hostname, and IP Mapping in XDR-A

Contents

Introduction

This document describes how to understand XDR-Analytics behavior in relation to device hostname, and IP mapping.

Background

XDRA attempts to track a logical devices behavior over time, known as a Device.

It uses various techniques to correlate network traffic to these logical devices over time.

However, particularly in an on-premises environment, there are limits to how well the system can associate traffic to a Device.

XDRA primarily gathers telemetry for on-premises environments through netflow via the ONA, CTB, or Cisco Meraki integration (the “new” Meraki integration). Secondarily, it can get hostname resolution through:

- Active hostname resolution via reverse DNS lookups and optionally SMB queries via the ONA
- ISE integration via the ONA
- The “old” Meraki integration
- NVM integration, with additional caveats

Netflow has IP addresses without hostname information.

Without hostname information, it assumes each internal IP address (see definition below) seen is a Device, since it has no further information to make a more intelligent Device association.

In a case where hostname collection is configured, XDRA uses hostnames, when seen, to tie it to an internal representation of a Device.

This allows XDRA to group multiple IP addresses over time to one Device.

NVM telemetry can be optionally configured as part of XDR.

This telemetry source provides a netflow-like data feed, but also provides endpoint information with unique identifiers.

The way XDRA leverages this information has the net effect of Device tracking behaving similarly to the case where hostname collection is enabled on the ONA.

All of these setups have limitations based on the limitations of the available telemetry.

Please note XDRA assumes the nature of IP address and hostname mappings is a many-to-one relationship

(many IPs can map to one hostname).

One logical device can have multiple IPs simultaneously (for example two physical interfaces or IPv4 and IPv6).

Due to the nature of the monitoring XDRA can never assume to have all relationships of the actual network at any given moment in time.

Overlapping subnets

In the case that a single XDRA tenant is monitoring multiple on-premises subnets simultaneously, the system cannot distinguish between the same IP seen in each of them.

As such it over-correlate IPs to Devices. Hostname availability does not improve this situation.

One way around this is to have more than one XDRA portal (one per subnet). Another is to use the "New" Cisco Meraki Integration due to the namespace isolation this integration brings.

Environment with no available hostname information

As a side effect of the limited telemetry information the system can come to an incorrect understanding of a Devices history.

One scenario is when IPs are dynamically assigned, XDRA does not have a way to know that the underlying logical device has changed for example a laptop on WIFI leaves, and the IP is assigned to a new laptop.

In the absence of hostname or other identifying information, the system associates the activities of multiple logical devices to one Device. This can lead to confusing device profile information.

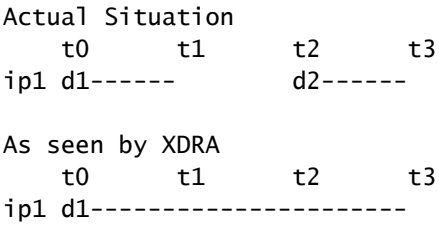

Conversely, in cases where one logical device has more than one IP address (for example two physical interfaces or IPv4 and IPv6), there is no information with which we can reliably tie these to the same Device, so the system does not.

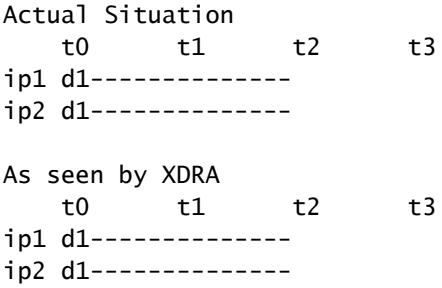

Environment with hostname information

Where XDRA can see hostname information the system has the ability to associate more than one IP address with one Device. However, given the nature of the data there are still limits to what the system can reliably determine. This can lead to over-correlation of IPs to Devices in the system.

If a Device that has an IP to hostname association in XDRA, and then the logical device changes IP address, the telemetry eventually reflects the new IP to hostname mapping.

However, because of the potential for this to be a many-to-one relationship XDRA can NOT safely assume the previously seen IP is no longer associated with the hostname (and thus the Device).

It could, for example, be a separate physical interface to the same logical device. So XDRA keeps both the previously seen IPs along with the most recently seen IP, until telemetry is seen that positively maps the IP address to a different hostname.

At this point XDR 'expires' the mapping and be listed as a previous IP address.

There is no way to tell the system to break an association 'early'.

Note on hostname matching

In order to try to better handle cases where a tenant has the same hostname configured in multiple domains, XDRA employs a ‘flexible’ matching and treats the entries show in this table as matching hostnames when looking to match an existing Device (that is in the case of a matching IP):

```
example
example.com
example.net
example.obsrvbl.com
example.invalid.obsrvbl.com
example.example.com
```

In other words, it considers just the hostname while ignoring the rest of the domain name.

Environment with NVM

This setup behaves very similar to Environment with hostname information section with hostname information, but there are a couple differences.

This data feed provides the added benefits of being able to provide some unique endpoint identifiers to the user, and these IDs potentially allow us to track a physical device that undergoes a change of hostname (which is not possible to track otherwise, we would create 2 different Devices).

While Devices are created based on the endpoint data feed (with unique endpoint IDs), there are no hostname or IPs associated with these Devices until an observation is made about that endpoint based on the flow data.

Environments with ISE

The benefits of ISE to Device tracking end up being identical to Environment with hostname information.

ISE data is used to associate hostname information it collects to IP addresses, but it does not create a new Device or track IPs that have not been seen in netflow.

Environments with Meraki

“Old” Meraki Integration (that is with XDRA)

This Meraki integration proactively gathers hostname information from Meraki devices, mapping those hostnames to IPs in the as usual for on-prem Devices (that is the “default namespace”).

This process creates Devices if they do not already exist.

It does not augment Device or IP information gathered from the other “new” Cisco Meraki integration due to namespace differences.

In effect, this causes this configuration to behave like an Environment with hostname information.

“New” Cisco Meraki Integration (that is with XDR)

This integration gets netflow from Meraki networking equipment, through the XDR data lake, into the standard XDRA netflow path.

As such it creates Devices like any other netflow; also like any other netflow, it does not contain hostname information.

In effect, this configuration behaves like Environment with no available hostname information, with one major exception.

This integration leverages the information sent to label the netflow from different Meraki equipment into different namespaces.

This avoids the usual Overlapping Subnets issues, but can introduce new difficulties if more than one integration is set up.

Most obviously, if both “Old” and “New” Meraki integrations are set up, they do not use the same namespaces and thus they create non-overlapping Devices, even in cases where the information represents the same physical device.

That is, you have 2 Devices, one in the default namespace with a hostname and no traffic, another with traffic in a specific Meraki namespace and no hostname.

Similar ‘splits’ can occur with other integrations if simultaneously enabled.

Definitions

1. **I**nternal IP address: XDRA considers IP addresses either internal or external, and this is configurable via the subnet settings. Subnets for on-prem default to the RFC internal subnets (RFC1918 and RFC4193) but subnets can be configured (added or removed).

2. Namespace: Additional information that is used to label netflow and Devices seen from different observation points, allowing Overlapping Subnets without overlapping IP issues.

ISE Hostname data flow

1. ONA collects ISE session data, uploads to S3 every 10 minutes
    1. this data contains user<->IP information, also sometimes includes hostname
2. IseSessionsMiner parses the uploaded data, and associates IPs with Devices where possible. It does NOT create a Device if one does not already exist. As it does this, it gathers available hostname<->IP mappings whenever we do have a Device already.
3. It then creates a file in s3 with those mappings in the same format as the ONA would upload one from it’s reverse DNS lookups
4. It then tells the system to load those hostnames just like it would load ONA hostnames.

FAQ

Why am I seeing IPs on an XDRA Device that are no longer associated with that logical device on my network?

Unfortunately, there is nothing we can do about this.

The system cannot know if the old association is invalid or the result of, for example an additional physical network interface.

I do not have any hostname information being sent to XDRA, why does my Device that is using both IPv4 and IPv6 addresses show as 2 distinct Devices?

Without hostname information we cannot know different IPs are associated with the same logical device on your network.

Why do I see multiple logical devices from different subnets appearing in the same XDRA device?

XDRA currently has no way to distinguish what subnet telemetry comes from, so the same IP is always grouped into one Device.