Cynamics

Network Fundamental

Whitepaper

Authors: Dr. Aviv Yehezkel, Eyal Elyashiv

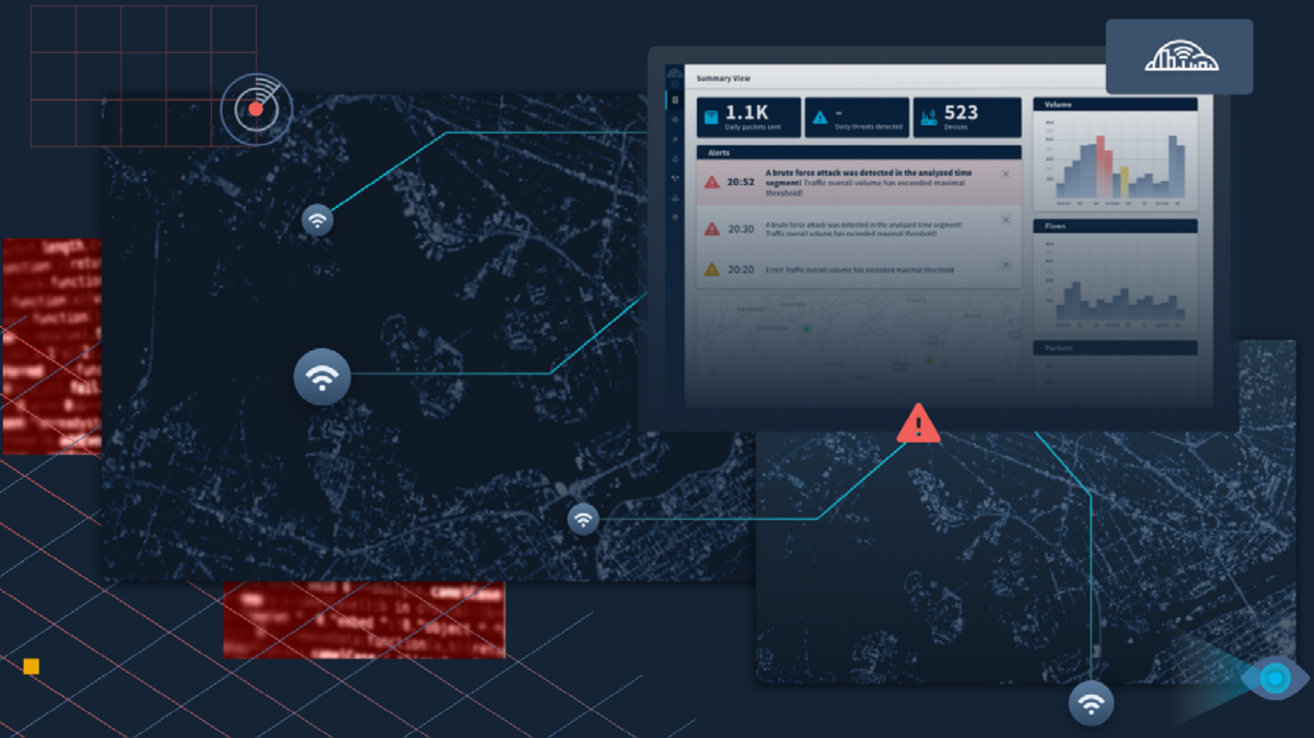
Summary View
1.1K
Daily packets sent
Daily threats detected
523
Devices
Alerts
20:52
A brute force attack was detected in the analyzed time segment! Traffic overall volume has exceeded maximal threshold!
20:30
A brute force attack was detected in the analyzed time segment! Traffic overall volume has exceeded maximal threshold!
20:20
Error! Traffic overall volume has exceeded maximal threshold
Volume
Flows
Packets

cynamics.ai

In our previous white paper, we introduced a Cynamics solution that provides complete network visibility for the most complex and complicated networks, regardless of whether it consists of legacy routers and switches, virtual devices, the cloud, or all of them together. The Cynamics solution is built upon the methodology of sampling a radically small fraction of network traffic. This capability is a standard 'built-in' in every network device and doesn't require any hardware or software modifications in the monitored network. The crux of the approach is to apply innovative machine learning techniques to infer the required monitoring information from a very small sample.

By collecting small samples from the customer's network devices, Cynamics' innovative approach provides complete visibility and threat prediction down to the end-point level, which is needed now more than ever as hackers and bad actors don't take any corona-cation. Just over the last few weeks, our predictions have saved millions of dollars of recovery for our customers from various public-safety domains.

Cynamics' core breakthrough is the ability to keep learning and evolving automatically and autonomously from each new customer by a combination of 2 deep learning frameworks:

The local expert algorithm: a specialized learner for each customer, learning its behavior over time in different layers, such as different networks, devices, protocols, trends over time, inter-devices behavior, etc.

The global expert algorithm: normalizing all networks' behavior to create a global basic fundamentals of the network, which we call the Network Blueprint.

What does it mean?

Each network has its own characteristics and properties, which may vary significantly from each other. Consider, for example, traffic volume, i.e., the total number of bytes entering the network each second. In small networks (e.g. a town), it can be hundreds of megabytes (Mb), while in large networks (e.g., the state of New-York), it can be hundreds of gigabytes (Gb) or more.

Now, let's consider a DDoS attack on the town, consisting of 10 Gbps traffic while its network resources are built for 1Gb maximum. It will easily crush the town's network, however, it will hardly affect the state's network.

So how can a single global detector handle such different traffic patterns?

Let's continue with the above town/state example. In this simple scenario, a naive solution would be to normalize the traffic volumes according to their maximal values: if the maximal volumes over the 'training period' were 100 Mb for the town's network and 100 Gb for the state's network, then on the 'inference period' a value of 1 Gb will be 0.01 in the state's network (accounting for 1%), but 10 in the town's (accounting for 1000%), so the 1 Gb will be flagged in the town's, but also won't cause a false alarm in the state. This simple normalization allows us to transfer the traffic volumes in each network to one single baseline, even though they are significantly different from each other.

However, in reality, things are of course much more complicated - there are an endless number of network features, which may be common in one network but rare in another one (even if both are from the same industry), and the patterns and behaviors are changing constantly.

Take a hospital's network, for example - IoT protocols will account for the majority of internal medical devices, while the network gateways (routers and firewalls) will hardly see these protocols and will use mainly the 'classical' IP protocols that will be found in municipality networks, for example.

To solve this challenge, Cynamics developed a novel deep-learning approach that is done by normalizing the various different customers' networks to create one single basic fundamental of the networks.

By doing this, the local experts of the different networks and devices with (possibly significantly) varying characteristics, properties and behaviors are 'translated' to a unified language that is then forwarded to global detection models that can detect and classify threats in a generic way which is agnostic to the specific deployment.

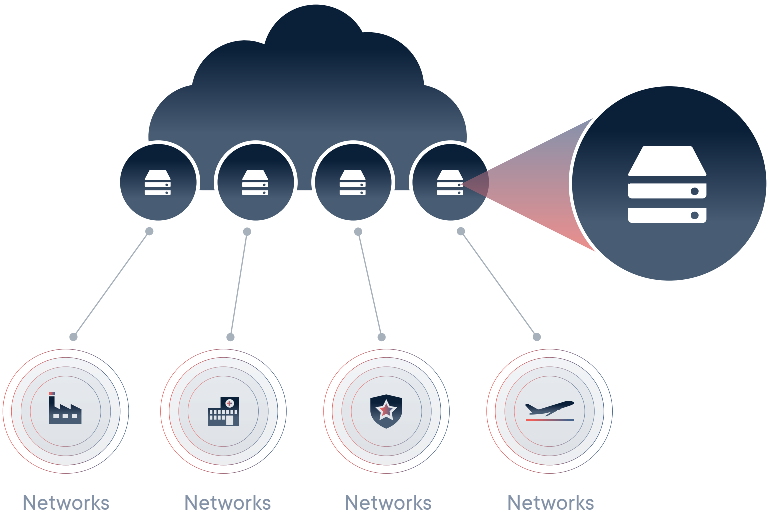

The Network Blueprint's breakthrough is two-fold:

First, it allows exponentially faster learning from threats and attacks detected, such that after a new threat was detected in a public-safety network, it will immediately update the Network Blueprint to detect 'similar' threats in other networks, such as a municipality.

Second, and no less important, it allows the Network Blueprint to recognize numerous normal behaviors, even if some of them may seem as anomalies, if looking from the local perspective of the specific customer. For example, a software update, a daily back-up, etc. - all these procedures are anomalous (different from the network's normal behavior), yet they are not malicious and shouldn't be flagged.

Cynamics' ability to recognize these 'anomaly but not threat' behaviors removes noise and leaves our customers to respond only to significant events. In particular, it makes Cynamics' false-positive rate (out of 100 detections, how many of them are false?) negligible, much below the 0.5%-1% 'golden rule' of current state-of-the-art machine learning threat detectors.

This is done using a new type of 'transfer learning' invented by Cynamics. Transfer learning is one of the most powerful tools in the deep-learning toolbox, allowing knowledge gained (or learned) while solving one problem to be applicable to a different problem.

A nice, well-known example is the dogs and cats classification.
First, a deep-learning neural network was trained over millions of images in learning how to classify hundreds of different types, such as a table, chair, person, car, etc., becoming an expert classifier. Then came the transfer learning part. The expert network was used to learn new data, this time consisting only of dogs and cats, and was able to become fully trained after a fraction of the dataset's size it needed initially. The deep-learning intuition is that after being trained over so many different types, the initial network became an expert in its ability to uncover specific underlying low-level features.

Thus, when re-trained over the dogs/cats dataset, it needed a very small number of examples in order to fine-tune its classification mechanism to classify dogs and cats, too. This is the 'usual' transfer learning done today in endless applications - computer vision, speech recognition, natural language processing, etc., training a neural network to become an expert on one problem and then fine-tuning it into a new, close, problem using much less training data.

Cynamics developed a new kind of transfer learning, normalizing different customer's network behaviors to one global model that keeps evolving and improving.

Now, let’s now see some examples from the last couple of weeks...

Malicious traffic in a municipal network

Suspicious, **stealth-looking**, communication was detected over a specific port between a user’s workstation and an IP associated with a **foreign cloud messenger application which is known to have several vulnerabilities** that is taking place almost every day throughout the day. Because the traffic was of negligible volume, accounting for no more than a few hundreds of packets per day, the existing security postures used by the client were unable to notice it, but it was immediately detected by Cynamics.

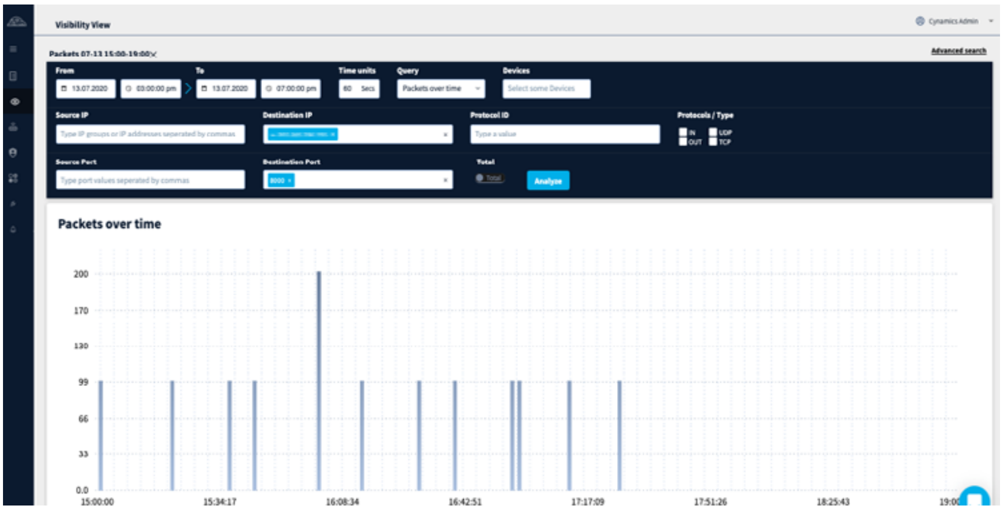

Visibility View

Packets 07-13 15:00-19:00

From: 13.07.2020 03:00:00 pm
To: 13.07.2020 07:00:00 pm
Time units: 60 Secs
Query: Packets over time
Devices: Select some Devices
Source IP: Type IP groups or IP addresses seperated by commas
Destination IP: [illegible]
Protocol ID: Type a value
Protocols / Type: □ IN □ OUT □ UDP □ TCP
Source Port: Type port values seperated by commas
Destination Port: 8000
Total: Total

Unauthorized network sync in public safety

Cynamics’ threat prediction autonomously spotted a daily update that was previously unknown to the customer. Several times a week, at 3 am ET, a highly unusual HTTP traffic was entering the client’s network from a local IT company. Thanks to Cynamics, autonomous root-cause analysis spotted the specific update details: traffic type, origin, etc. The client was then able to stop this unauthorized sync process.

Entire traffic:

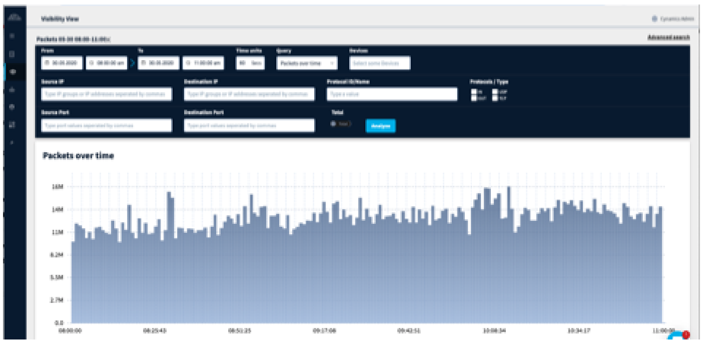

HTTP-only:

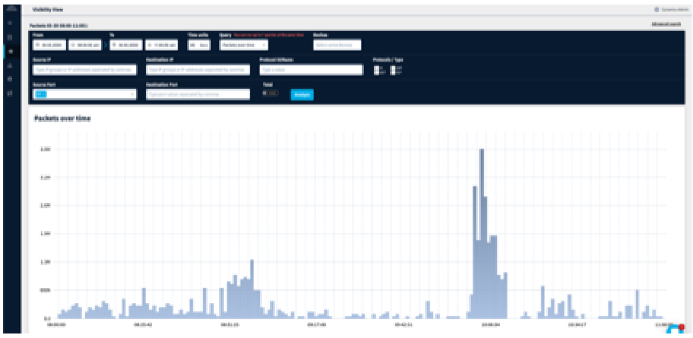

While the traffic patterns seem perfectly normal when looking at the entire traffic (upper image), drilling down to HTTPS traffic specifically showcases Cynamics AI in action: **automatically analyzing multi-vector network patterns, thus enabling the ability to predict and spot anomalies and threats that otherwise go unnoticed on the numerous different network patterns.**

Ransomware attack on a hospital chain

The suspicious traffic pattern was detected coming towards a specific end-user's workstation. A short follow up investigation with the customer yielded **"This host was compromised"**, as preparation for malicious activity such as **ransomware attack.**

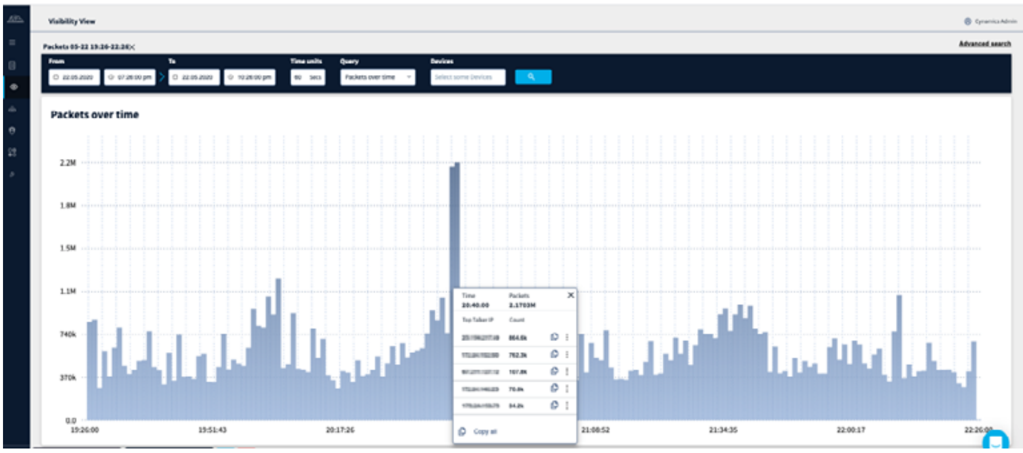

Packets 05-22 19:26-22:26

From: 22.05.2020 07:26:00 pm

To: 22.05.2020 10:26:00 pm

Time units: 60 secs

Query: Packets over time

Devices: Select some Devices

Packets over time

Time	Packets
20:40:00	2.1703M
Top Talker IP	Count
[illegible]	864.8k
[illegible]	762.3k
[illegible]	107.8k
[illegible]	70.4k
[illegible]	64.2k

Copy all

Use case - Data Leakage

This example shows Cynamics' ability to learn network behaviors and predict threats down to the endpoint level. Suspicious HTTPS traffic coming from an endpoint IP was detected **towards Google and Microsoft - seems legit right?**

The user said he was transferring large files from his workstation at that time to Google GSuite (explaining Google's traffic, but not Microsoft's).

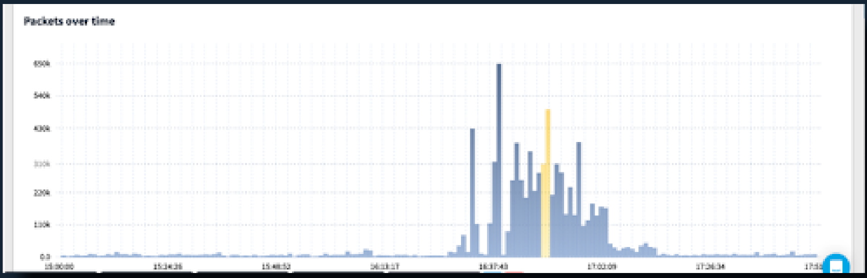
Packets over time

Cynamics further detected low-volume, **stealth-looking** traffic between the customer's **core network** and **the workstations** that were previously unknown to the customer, and were supposed to be blocked at the firewall, as the core network must be kept isolated from the internet.

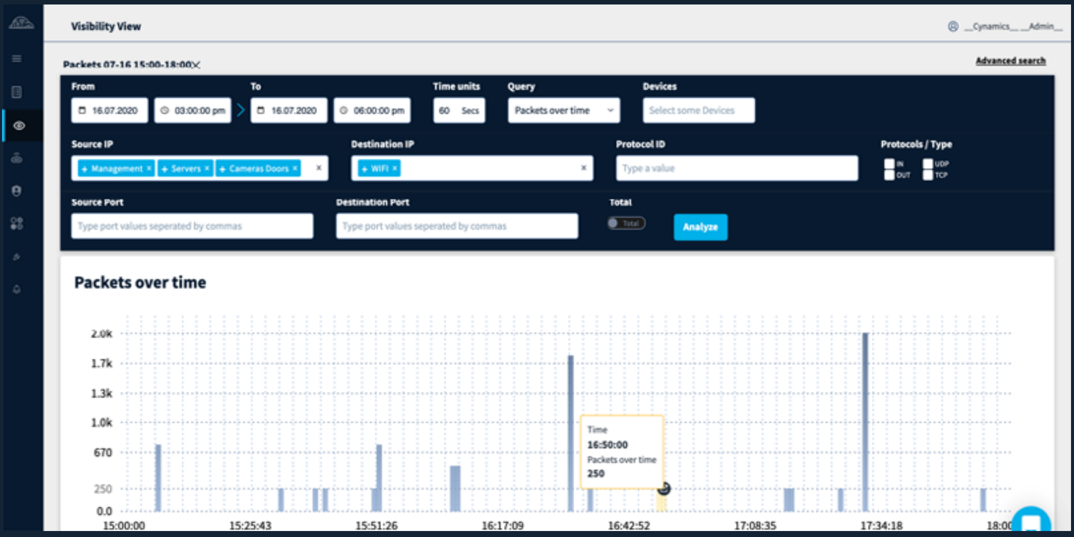

Visibility View

__Cynamics__ __Admin__

Packets 07-16 15:00-18:00

Advanced search

From: 16.07.2020 03:00:00 pm

To: 16.07.2020 06:00:00 pm

Time units: 60 Secs

Query: Packets over time

Devices: Select some Devices

Source IP: + Management × + Servers × + Cameras Doors ×

Destination IP: + WiFi ×

Protocol ID: Type a value

Protocols / Type: □ IN □ OUT □ UDP □ TCP

Source Port: Type port values seperated by commas

Destination Port: Type port values seperated by commas

Total: Total

Analyze

Packets over time

Time 16:50:00 Packets over time 250

In addition, the user's workstation's address was found to keep communicating with Microsoft's IP for a few hours afterward.

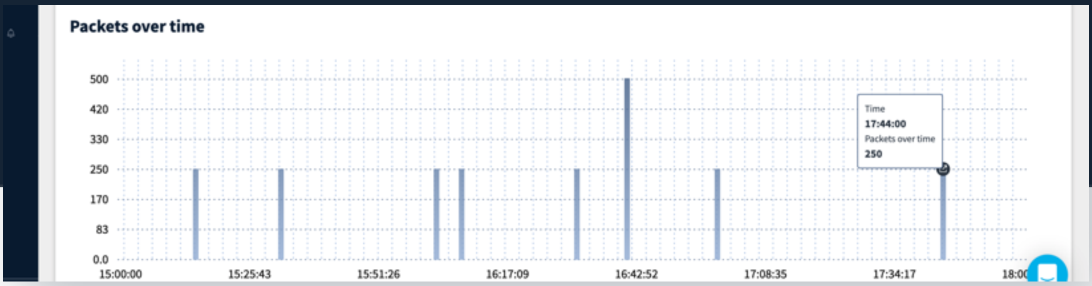

Packets over time

Time 17:44:00 Packets over time 250

Autonomously connecting the dots altogether, Cynamics provided the client with a clear comprehensive attack story detailing this data leakage:

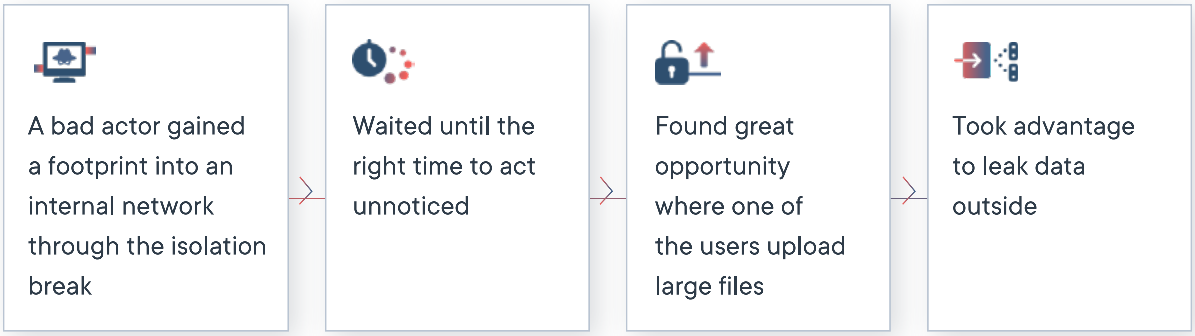

A bad actor gained a footprint into an internal network through the isolation break

Waited until the right time to act unnoticed

Found great opportunity where one of the users upload large files

Took advantage to leak data outside

Resolution and mitigation conducted by the customer:

- ✓ Firewall rules were updated
- ✓ Suspicious Microsoft IP was **blocked**
- ✓ Network isolation was forced and passwords were updated
- ✓ Cynamics **custom-alerts** were set up to monitor traffic from these groups and verify new rules

This is the strength of looking for what's hiding within the patterns. Previously unseen sequences reveal what's really taking place on networks in real time, without the need to monitor each and every device. At Cynamics, we're building the impossible: a cost-effective, scalable smart city network monitoring solution to help municipalities, critical infrastructure, healthcare, and other highly complex and sensitive organizations predict and prevent attacks and optimize network performance.

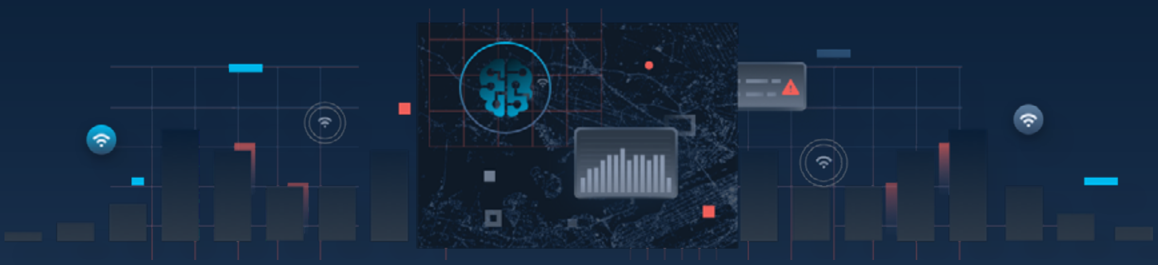

To learn more or to test-run the platform on your smart network, click here