
Cynamics
How to See
the Unseen
Network
Dr. Aviv Yehzkel, Eyal Elyashiv
www.cynamics.ai

Smart networks power some of the most sensitive organizations on the planet. These highly complex and interconnected networks enable governments, healthcare facilities, public safety, and utility providers to manage and monitor essential services. They are the backbone behind critical infrastructure sectors, helping industries vital to national survival to function uninterrupted.

But there's a troubling lack of transparency across these increasingly complex smart networks. Without total visibility, CSOs, CIOs, network operators, and security analysts cannot get an accurate picture of what's taking place, leaving room for threats to infiltrate unseen.
Since the beginning of the smart network revolution, solutions have attempted to address this shortfall with network monitoring and detection solutions. Problematically, few are purpose-built with the intricacies of extensive smart networks in mind and all are prohibitively costly.

In this white paper, we'll explore a new approach to obtaining complete visibility for your smart networks using Artificial Intelligence (AI).
For the first time, you can get global visibility of the most complex and interconnected networks, using a fraction of the network traffic, at a fraction of the cost of traditional network monitoring tools.

(Not) Managing the Complexity

The increase in scale and complexity of smart networks poses a seemingly impossible management challenge. A network failure, a traffic bottleneck, or an architectural mis-configuration can start in any one of hundreds—or even thousands—of network devices, causing rapid deterioration of network performance or significantly compromising network security.

The resulting lack of transparency leads to unanticipated attacks, unmitigated threats, and other potentially harmful security anomalies.

Average Cost per Cyber Attack Incident - $13M

Utilities $18.4M

US Federal $13.7M

Health $11.9M

Public Sector $7.9M

So what's the solution to cutting through the complexity of smart networks?
Some organizations assume that by adding specialized monitoring to each network device—including edge routers, firewalls, AWS/Azure/GCP clouds, internal switches and private networks—coupled with some network monitoring and detection solutions, they should be okay.

However, detection of sophisticated cyberattacks requires a global view and analysis of patterns between a number of devices. The reality is that you need to cover your entire network with monitoring and detection solutions. But this is highly expensive, requires a great deal of network modification and device configurations, and can affect their performance.

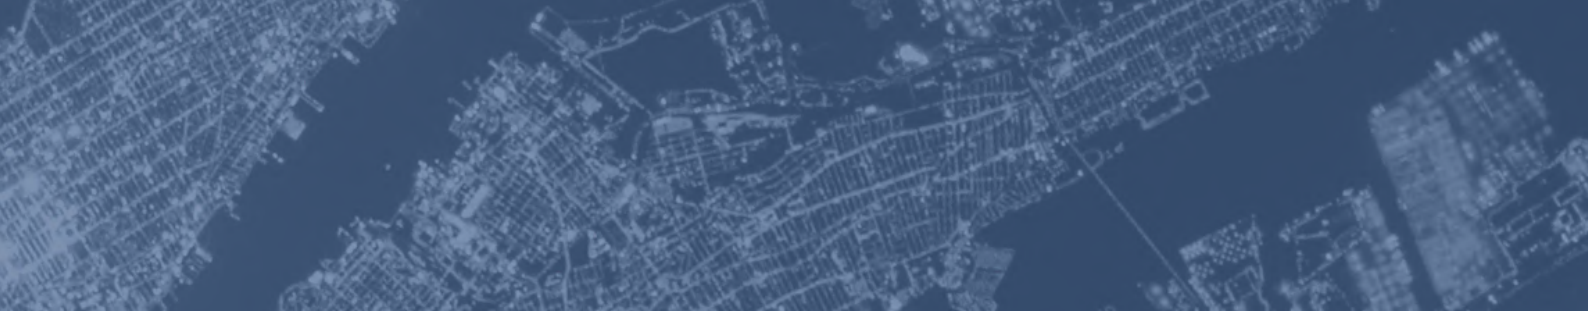

So the unfortunate reality is CIOs, CSOs, IT managers, network operators and cyber analysts are left to choose which assets are most critical—and which ones are not. They then put all resources into supporting and protecting those deemed most critical, while simply hoping the other ones don't get hit or suffer performance issues.
This renders most of the network unprotected, and leaves security teams just praying they never get that "network disaster" midnight phone call. But as the stats clearly demonstrate, monitoring only some of the devices is not enough to get a comprehensive and nuanced view of what's happening on your smart network.

To address this challenge, experts from Israel's leading universities and most prestigious cybersecurity units, along with industry leaders, teamed up to put an end to this problem and allow security teams of these complex networks to get a decent night of sleep.

With Cynamics, you can achieve 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.

No More Compromises - Here's How We Do It

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.

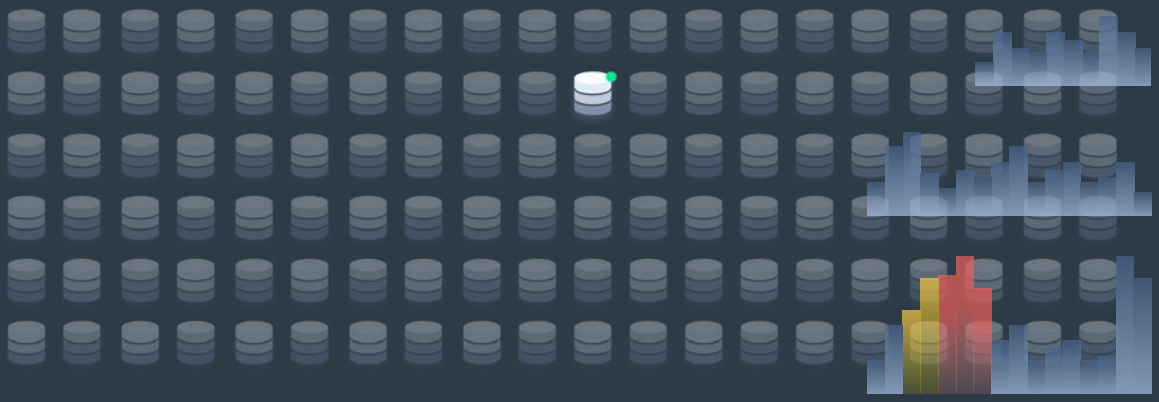

If this sounds impractical, remember; this is exactly how blood tests work, inferring a complete picture of blood properties from an almost negligible sample. So too, analyzing just a small percentage (less than 1%) of network traffic minimizes the impact of monitoring activities on network performance and reduces costly processing time compared to other solutions.

Network Visibility - How to See the Unseen

To get a better understanding, let's explore some actual networking examples. To start, know this; An important network field is called **"flow"** or **"connection"**. It's defined as a collection of packets with the same "network ID": source and destination IP address, source and destination port, and protocol. A good example of this is a user's login to a given website. Each user connects from their specific home IP address to the website's specific IP address over HTTP port (80) and TCP protocol (6). Thus, each user uses a different and unique connection—namely, a different flow.

For simplicity, let's assume that every action the user performs while on the website (e.g., login, navigating to a certain page, filling a form, logging out, etc.) is executed by sending a single **packet** from the user, to the website over its connection.
An inherent difficulty in any sample-based solution is that the sample is likely to miss low probability events. Think of packets coming to a website from a "passive" user, i.e., a user who performs very little activity while on the website. For example, the user only logs in and then immediately closes the browser.

Now assume that we sample at 1/100 rate, so that 1 out of every 100 packets is sampled. This user contributed only 1 packet (when he logged in). Thus, the probability that this packet will be sampled is 1%, which means that it is very likely to be missed, and therefore our sample won't include this user.

The immediate question is then, if multiple users' packets are missed, how can we know how many users visited our website in actuality? Or in other words, how can we know how many flows entered the website?
This problem is known as **"Estimating the number of different flows"**. Monitoring the value over different time intervals and different network devices is a basic principle in network monitoring and can be related to several network attacks. For example, in a **Distributed Denial of Service (DDoS)** attack, a simple, yet effective, technique is to use many bots to send high-volume traffic to the network, thus saturating its resources until they crash it.

Going back to our website example, this attack can be executed by using numerous bots to send requests to the website and preventing legitimate users from accessing it. Translating this to networking language, it means the number of users will rise significantly over a short period of time—or in other words, the number of flows will rise. Thus, by monitoring the number of flows, we can detect such attacks by looking for suspicious changes in the number of flows.

How do we monitor this number? We have developed a novel AI technique that models the function of a number of different flows with a given number of packets in the sample (frequency). This function gives an estimate of the number of flows with a specific frequency by learning the most likely behavior of the total network traffic, based on the small samples we observe. Back to our example, it learns how "unseen users" most likely behave.

How does this relate to smart network attacks?

Every attack has a pattern that precedes it; whether longer or shorter, there is some type of unusual behavior that occurs before the attack. Consider our website's DDoS example—the pattern, in this case, is the sudden increase in the number of flows over a short period of time.

This is very similar to a DDoS attack we recently detected at one of our municipality customers, where we monitor all Internet gateways and firewalls. The following graph from their Cynamics dashboard demonstrates it very clearly: a sudden, unusual spike in traffic volume. Due to our sample-based approach, we can significantly reduce detection and response time by over a minute, when compared to existing solutions that need to collect, process, and analyze each and every packet. Thanks to our real-time detection, the municipality was able to mitigate the attack swiftly and suffered no harm.

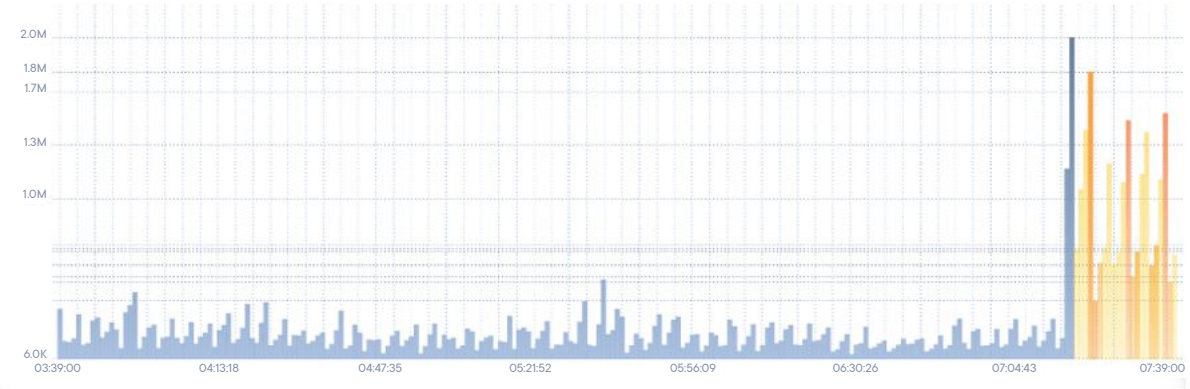

Frankly, this was a simple case. However, the patterns preceding attack are often not quite as easy to catch. Below, we have a case at another customer in the healthcare domain, where our solution was used to solve a significant network visibility bottleneck, caused by a mix of devices and high traffic volumes. Look at the following graph which shows the total traffic that entered a specific device.
Can you spot a suspicious pattern?

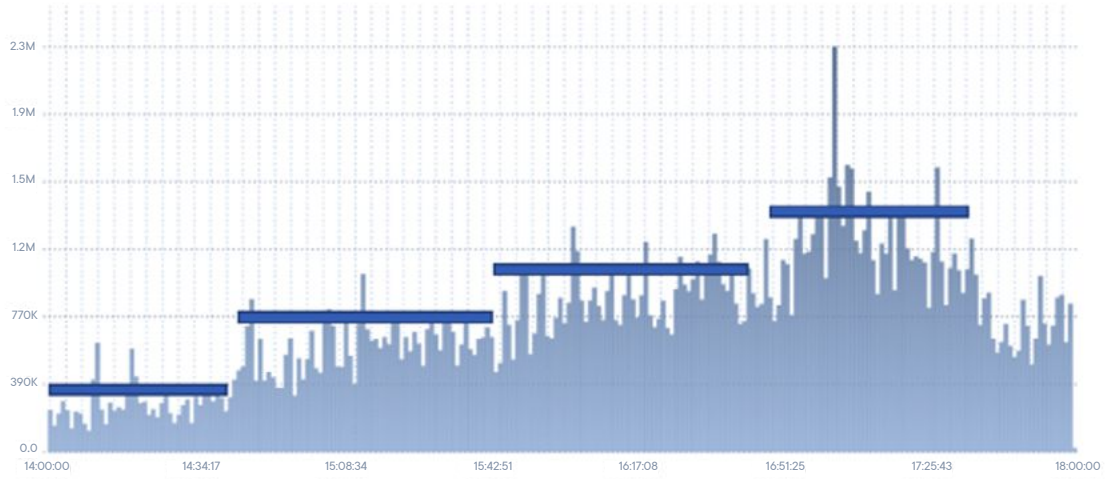

The four blue lines indicate an incremental growth of the traffic volume: from ~400K packets to ~1.35M over 3 hours (with a significant spike of 2.3M). Incremental growth like this is very common in DDoS attacks, where various attacking bots "wake up" on a delay from each other (in order to allow them to blend in with regular traffic) and begin to bombard the network, one after the other—until they all bombard it together, which eventually causes it to crash.

This is a perfect example of how our AI technology works behind the scenes, automatically learning the most important network fields and using them to summarize the network state at each timestamp. We call this the “traffic state vector”. Our detection models analyze these patterns over time in several layers, including each device by itself, the entire network level, groups of devices, etc., and compare their real-time behavior to find any suspicious changes and deviations from normal behavior.

Returning to our example, our automatic root-cause analysis capability, which looks for attack origins, found unusual traffic coming from a school district in Florida, even though the customer was located nowhere near there. We were able to detect the malicious activity and shut it down before the network could be harmed.

Cynamics - The Power to See the Unseen

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**.

Cynamics

www.cynamics.ai