Defending Network Infrastructure Against Attack

Kevin Dooley

ABOUT AUVIK NETWORKS

Auvik's cloud-based network management software keeps IT networks around the world running optimally. By automating and simplifying network management, Auvik helps rocket an IT team's efficiency and capacity, while protecting the business from network risk.

Contents

THE GOOD NEWS—AND THE BAD NEWS—ABOUT NETWORK SECURITY

The news, it seems, is always thick with reports of major attacks on corporate networks. In the cases of the Panama Papers, the OPM leak, and the Hacking Team leak, the results were catastrophic leaks of extremely confidential information.

In each case, the organizations that were hacked spent a great deal of time and energy on PR after the fact downplaying the significance of the leaks. But in each case, the hacks were made possible because of basic flaws in the network infrastructure and a failure to take security seriously.

Perhaps the worst attack in recent years was the case of the Bangladeshi central bank. Reports suggest they did almost nothing to secure their infrastructure prior to the attack, and it wound up costing them many tens of millions of dollars.

In truth, a determined and well-resourced attacker can always find a way in. If your information is of value or interest to a foreign government, you should probably assume they've already taken it. But in the Panama Papers and Hacking Team leaks, the attackers were probably independent, non-government hackers.

In all of the cases, though, there were critical errors made in securing network infrastructure.

Right now, the most worrisome network threats fall into four main categories:

- Malware
- Phishing
- Denial of service (DoS) attacks
- Advanced persistent threats (APTs)

The good news is that it's neither extremely hard nor overly expensive to mount a reasonably effective defense.

The bad news is that it's impossible to create a *perfect* defense.

In particular, you can't keep determined, skilled, and well-resourced attackers like government agencies away from your data. They'll always be able to find and exploit vulnerabilities in your defensive security infrastructure.

So let's focus on the more manageable task of keeping out the routine criminals.

PART 1 - Network Threats

1. MALWARE AND RANSOMWARE

Malware is the modern term for what we used to call a computer virus. The term has changed because the threat has changed—today's malware is much more dangerous. It's usually deployed starting with a small "dropper" program that then contacts a command-and-control host (variously abbreviated C&C, CC, or C2) to get further instructions and download additional malware.

One of the most dangerous and growing types of malware is crypto-ransomware, which immediately sets about encrypting all your files, including all files on any network shares. It then offers to give you the key to unlock your files in exchange for paying a ransom.

DEFENSES AGAINST MALWARE AND RANSOMWARE

Most malware isn't targeted. By that I mean it's produced with the general intent of catching somebody—anybody—not you in particular.

Malware writers often exploit software vulnerabilities in common applications like your web browser, Flash, Word, or Excel. In many cases, they also exploit operating system vulnerabilities.

First line of defense

There are two critical elements in a first line of defense against malware.

First, keep up with software patches. If you apply all patches as soon as they're released, you'll be ahead of just about all malware threats.

Second, use a good endpoint protection system. Endpoint protection is what we used to call antivirus software. Traditional antivirus software relied heavily on file signatures.

Whenever a new file appeared on your computer, the antivirus software would scan it. It would calculate an overall file checksum that it could compare against a database of known malware, and it would scan through the file to see if it contained a sequence of bytes associated with any known malware.

Antivirus packages are still valuable tools, but the problem is that malware writers have started adopting clever tricks like encrypting the code with random and frequently changing keys. That means the malware will never appear with the same checksum twice, and the internal byte sequences will be obscured.

To combat random keys, modern endpoint protection software generally includes some sort of sandboxing feature. The suspected malware is unpacked in a safe, virtualized environment and allowed to install itself and run while the scanning software carefully monitors it for signs of any malicious actions.

But the fight against malware is an ever-escalating battle. Malware writers have started building special sandbox avoidance techniques that detect when they're running in a sandbox instead of a real system.

So most good endpoint protection systems also monitor the real endpoint workstations for signs of malicious software actions and try to block the malware before it's too late.

Second line of defense

The second line of defense in protecting infrastructure from malware is to assume the first piece of malware will actually be a dropper, and that it will reach back across the internet to a command-and-control server for further instructions and further software packages.

Modern malware is often highly modular, so we can often catch the infections by implementing good scanning on the network edge. This is different from a traditional Intrusion Detection System (IDS), which generally monitors inbound connections. Here, we're monitoring outbound connections.

We're looking for several key indications of compromise including:

- Connecting to known malware domains or C&C systems
- Downloading suspicious files
- Traffic patterns that appear to indicate interactive VPN-like activities, which might indicate a remote access Trojan (RAT)

But be aware! In the case of common ransomware infections, additional command-and-control action often isn't necessary. The initial piece of malware just starts encrypting every file it can read and keeps going until you stop it. If it's been allowed to unpack itself and start running, it's probably already too late.

The best defense against encrypting ransomware is good old-fashioned backups. Find and shut down the infected machine, then start restoring files from the backup.

If you keep the backups offline or otherwise inaccessible to normal users, and if you take backups at least daily, then your exposure is limited to whatever has changed in the last 24 hours. It isn't great, but it's usually not a disaster. And it's certainly preferable to paying ransom, since there's little reason to believe the criminals will actually give you the decryption keys after you've paid.

2. PHISHING

I actually don't consider phishing to be an IT security problem. It's a social engineering attack in which an attacker contacts a person inside the organization, often through email, and tricks that person into doing something.

In some cases, the person is tricked into transferring money to the attacker's account. In other cases, the target is tricked into installing or running software that helps the attacker with the next stage of their attack.

Because phishing isn't really a technological attack, technological solutions are generally ineffective. Closing the door to an email-based phishing attack doesn't necessarily close the door against a similar attack conducted over the telephone or through the mail. This is old-style fraud—it has always existed and it will always exist.

DEFENSES AGAINST PHISHING

There are two main defenses against phishing.

The first is education. You can reduce the chances of suffering from a phishing attack if everyone is vigilant and aware of what phishing looks like. But this really only reduces the chances. An extremely clever attacker will always be able to come up with a convincing ruse.

What would happen if they emailed a realistic looking invoice to somebody in the Accounts Payable department, designed to look exactly like it came from one of your suppliers? It's almost certain the message would be opened. The same would be true of a realistic resume sent to the HR department.

So the other defense against phishing is procedural. Make it a well established and rigidly adhered to process that money transfers are never done without verbal confirmation from a small number of specifically named individuals. The CEO will never send an email to the head of accounting requesting a money transfer to a mysterious supplier in a foreign country. And even if they do, standard procedure is to call the CEO's cell phone and verify the instructions.

If the phishing attack is a method of deploying malware, then you can at least fall back on the malware defenses mentioned in the previous section.

3. DENIAL OF SERVICE

Denial of service (DoS) attacks can be launched fairly easily by unskilled attackers, and they can be carried out without needing access to your internal infrastructure.

The simplest DoS attacks try to overwhelm your internet link by sending huge amounts of traffic so that legitimate business traffic gets shut out. More sophisticated DoS attacks involve less traffic, instead using up other network resources.

DoS attacks are sometimes done to disrupt business and sometimes to extract a ransom to make the attacks stop.

DEFENSES AGAINST DENIAL OF SERVICE

The trouble with most DoS attacks is that once they hit, they've already used up your internet resources. It doesn't help to throw away malicious packets if there are so many of them that the link is full.

You really can't protect against most DoS attacks. The best approach is to use a protection service provider like CloudFlare to interrupt the attacks somewhere upstream from your infrastructure.

Protection services typically work by directing your traffic through their infrastructure before it gets to you. They'll either automatically detect attacks or allow you to specify you're under attack. Then they simply redirect the malicious packets into the trash can, and only forward the legitimate traffic to you.

Another good and popular way of mitigating DoS attacks is to put public-facing infrastructure on a cloud service provider with ample resources. That way, if you're attacked, there's no effect on your real infrastructure, just on the web hosting provider, which will generally have robust DoS mitigation systems.

More sophisticated DoS attacks seek to disrupt web infrastructure without necessarily using lots of traffic as the attack mechanism. Instead, these mechanisms use software vulnerabilities in the web hosting systems to take the systems off-line, or they use up all resources on those systems to prevent them from accepting new connections.

Since the more sophisticated attacks are generally based on software vulnerabilities, it's hard to create permanent defenses against them. But regularly patching internet-facing infrastructure goes a long way to minimizing risk.

4. ADVANCED PERSISTENT THREATS

Advanced persistent threats (APTs) are the attacks everybody fears. In an APT, the attackers manage to build a backdoor into your infrastructure, then carefully extract your most valuable data. The effectiveness of this type of attack depends on the skill of the attackers and your ability to detect and stop them.

APT attacks often start out as malware or phishing attacks. The attacker has to somehow get a foothold inside the infrastructure. Then, typically, the attacker instructs the initial dropper software to download a remote access Trojan, which is like a VPN that allows them interactive access.

I say typically because there have been a few cases where APT attacks have happened entirely without external feedback, but it's a ridiculously difficult way to attack a network, and nobody would do it this way if they had the option of interactive access.

DEFENSES AGAINST APTs

The defenses against APT attacks include everything we said about malware attacks. It's also useful to monitor your internet links for the typical signatures of RAT-like traffic patterns. However, if your attackers are skilled, an APT attack can go undetected for considerable lengths of time as they move laterally within your network looking for valuable data.

For this reason, in addition to prevention and detection, it's useful to have good forensic abilities when it comes to APT attacks. In particular, it can be very useful to maintain thorough logs of every access to every file and system on your infrastructure so you can reconstruct what user IDs accessed what resources from what systems. This is most easily done using the Active Directory or LDAP server logs.

Another useful forensic capability is some sort of packet capture or traffic flow monitoring tool. NetFlow-based systems can keep track of every conversation that takes place, including source and destination addresses, protocols, and amount of data transferred. This is often supplemented by detailed packet capture data, which can show you exactly what was transferred.

PART 2 - Network Defenses

DEFENDING A NETWORK AGAINST INBOUND ATTACKS

By inbound attacks, I mean traditional hacking attempts against a network's front-door elements like web servers, web applications, email systems, and remote access (VPN) systems.

FIREWALLS

The first and most obvious security element against inbound attack is a firewall. Pretty much any commercial firewall is suitable, but I generally recommend using a more sophisticated next-generation or unified threat management (UTM) firewall.

A firewall is a network device more than it is a security device. It facilitates connection between your internal network and the public internet. Since everybody these days uses private IP addressing, a firewall is the right place to do address translation between internal and external address spaces.

A firewall also has the ability to filter incoming or outgoing connections based on simple Layer 3 and Layer 4 elements in the packet header. These include internal and external IP addresses, as well as TCP and UDP port numbers.

A firewall also has to be stateful, which means it keeps track of every individual session passing through it. You shouldn't be able to get packets past a firewall simply by constructing them to look like part of a pre-existing session.

The problem with basic firewalls is that they assume everything that looks like a duck must be a duck. If it's a TCP session on port 443, then it must be HTTPS. But this assumption completely neglects the possibility that somebody might shift an illegal application to a legal port. Ports are just numbers—they're easy to change.

INTRUSION DETECTION AND PREVENTION SYSTEMS

To get around a firewall's essential flaw, it's useful to couple a firewall with an intrusion detection system or intrusion prevention system (IDS/IPS). An IDS/IPS is a device that watches every packet and session to look for signs of malicious activity. Generally, more serious problems cause the IDS/IPS to use prevention mode, where it drops the session. Less serious problems and things that are suspicious but not necessarily bad are merely detected, resulting in an alert.

Most IDS/IPS devices use a combination of factors to detect malicious behavior. They monitor for what might be called application orthodoxy, which means they reject sessions that don't appear to be following the established rules for the protocol they're using. And they use signature-based detection to look for known patterns of malicious behavior. To be really effective, you need to make sure those signatures are kept up to date, which generally means some sort of subscription model.

The reason I like next-generation or UTM firewalls is because they include an IDS/IPS in the same box, and display IDS alerts on the same management interface. Such a two-in-one device is cost-effective and simplifies management. Some UTM firewalls also look for things like virus signatures during file transfers.

REVERSE PROXIES AND WEB APPLICATION FIREWALLS

The next important inbound protection you can deploy at the network edge is a reverse proxy, a device that masquerades as a web server or similar internet-accessible server. The real server sits somewhere inside the network, and the reverse proxy passes the data to and from that real server. (Obviously, you only need this type of protection if you have some sort of server that's accessible from the internet.)

In many cases, a reverse proxy is used to translate between an insecure protocol like HTTP, which might be all that a legacy application server supports, and a more secure one like HTTPS, which is more appropriate for anything on the public internet.

The problem with reverse proxies is that they don't inspect the contents of the traffic they're relaying. That's fine for protocol-based attacks, but not for application-based attacks. For example, two of the most dangerous and common types of attacks against web servers are SQL injection and cross-site scripting. In both cases, the HTTP traffic looks perfectly fine, but is specially constructed to trigger a bug in the implementation of the web server that gives the attacker access they shouldn't have.

For this reason, I'm not a huge fan of reverse proxies unless they're also web application firewalls (WAF). A WAF is also a device that sits between the internet and a web server, but it explicitly sanitizes every request coming from the remote user. At a minimum, it will remove quotation marks and other special characters that are typical of SQL injection attacks.

Some WAF devices can also be configured to know exactly what types of data are allowed in specific fields on a web page that accepts input. Some WAFs will even monitor the connection between a web server and a database to ensure that an otherwise innocent-looking request didn't somehow result in a huge table dump.

EMAIL SCANNERS

Another device frequently exposed to the public internet is an email server. I've seen attackers directly trying to do malicious things to an email server, but more typically, they're trying to deliver malware or spam. So I like to deploy an email scanner in front of an email server. (Note that you can completely avoid the need for this type of defense if you use an outsourced email service.)

Email scanners are concerned with two problems: They want to eliminate or at least reduce spam, and they want to reduce malware attacks by scanning for viruses. An emerging use of these devices is a reduction in phishing attacks as well. However, a well-constructed phishing attack looks so much like a legitimate email message that it's really only feasible to eliminate the bad ones.

Imagine that an incoming email message has a virus in an attachment. If the email scanner catches it, it quarantines the message and the end user doesn't see it. If the email scanner doesn't catch it, then you have to hope whatever anti-malware system you have on that user's workstation will catch it. For this reason, it makes sense to use completely different malware scanning systems on email scanners and workstations. I'll discuss malware scanning systems in more detail later in this post

DEFENDING A NETWORK AGAINST OUTBOUND ATTACKS

Outbound attacks include malicious traffic that originates inside a network and goes out to the internet. It probably sounds like a pretty low priority. Why should you spend any special effort looking for evidence that your staff or clients are hacking somebody else?

While that's one of the things we're looking for, it's far from the most important. There are two other things we really care about here:

- Attempts by malicious software in the infrastructure to reach out to its controllers
- The leaking of sensitive data to external parties

DNS

The simplest and least expensive thing you can do to protect against outbound attacks is to use a really robust domain name system (DNS). Malware has three ways to call home for instructions. It can use:

- Hard-coded IP addresses—but they're very easy to block once discovered
- Domain names the malware rotates through on a programmed schedule
- Compromised legitimate sites or advertising services

DNS-based protection can be a useful first line of defense against the second type of attack, but not the others.

Interestingly, though, there do appear to be some malware systems that use DNS queries for communicating with other command-and-control traffic. That is, the results of a DNS query might be interpreted by the malware to mean something other than an IP address. This is another way a good DNS filtering system can be helpful in combatting malware.

However, the problem with relying on DNS for malware is that the malware is still sitting there inside your environment. It hasn't been eliminated, just temporarily silenced.

IDS/IPS

Another useful line of defense against outbound attacks is an IDS/IPS, perhaps the same one deployed for inbound protection. But this time we're interested in different types of attacks with a completely different set of signatures. Now we're looking for malware inside a network as it tries to connect to command-and-control servers on the internet. We're also looking for indications of unauthorized VPNs or remote access Trojans (RATs).

LOGS

If you use a central authentication system like LDAP or Active Directory, (and I believe you should) then you get another excellent defensive tool essentially for free. Simply monitor LDAP or Active Directory logs. Look for repeated failed logins, which might be an indication of somebody attempting a brute force attack. Look for people logging in when they shouldn't be or logging into systems they shouldn't be on, particularly users with special privileges, like system administrators.

Probably the most effective type of advanced persistent threat (APT) attack is one where the attackers steal legitimate user credentials. Then they don't need to devise any exceptionally clever hack—they simply log in and poke around until they find something interesting.

WEB PROXY SERVERS

The next thing to look at for protection against outbound attack is a web proxy server, a system that intercepts all outbound web requests and tries to serve them locally. If somebody just loaded a particular web page, the proxy can take the page from its cache and send it back to the user immediately.

Proxy servers provide better performance, as well as an opportunity to centrally scan all web content, including encrypted SSL content, and reject things that look malicious.

I don't like to make a proxy my first line of defense. It can obscure and limit the effectiveness of some of the other tools. For example, if you see a DNS request or an outbound connection to a known malware domain, it's a multi-step process to track it back to a particular workstation.

And I haven't found that the scanning capabilities of most proxies are any better than those found on the best UTM firewalls anyway. The important thing a web proxy buys you is the ability to decrypt and inspect HTTPS (SSL) content.

FORENSIC PACKET CAPTURE

Finally, if I had all of the security tools we've already discussed and I really needed something more to help me deal with incident response following an attack, I'd look at forensic packet capture tools. These are carefully optimized and targeted network protocol analyzers that record interesting looking sessions. They then allow you to search through a massive historical database of these sessions.

Forensic tools are mostly useful during the incident response process when you're trying to figure out what systems might have been affected by an attack, what credentials might have been compromised, and what data stolen or altered. That's useful information, but it's obviously not the starting point.

ENDPOINT SECURITY

Malware sneaks onto a system in many different ways. The typical malware attack starts with some sort of dropper download. The dropper is a small piece of code whose main function is to get itself installed—somehow, somewhere—then to reach out to the C&C network for further instructions. Usually, the instructions will involve downloading additional malware modules, installing them, running them, and possibly reaching back again for more instructions.

The drop-and-reach-back methodology of most malware gives us many opportunities to catch attacks. First, we try to detect and block the suspicious domains. Then we try to detect and block the dropper download. Then we try to detect and quarantine the dropper software. Then we try to detect and block the C&C traffic. Then we try to detect and quarantine the other malware modules.

The first couple of items on that list are already covered by the UTM firewall and other tools we've already mentioned. So the most critical malware-specific tool in our defensive toolbox is good endpoint security.

An endpoint security tool is a mix of traditional antivirus and behavior-based malware detection. Don't bother with a straight antivirus system anymore. Traditional antivirus programs scan files and use signatures to spot anything that looks like a known virus. But modern malware isn't always file-based, and it's often able to change its appearance by repeatedly encrypting and re-encrypting itself. So on their own, signatures aren't very useful.

However, signature-based detection can be supplemented with behavior-based indicators of compromise (IoC). Malware does malicious things like encrypting files, making strange entries in Windows Registry, or modifying system files. A good endpoint security program will monitor for these types of actions, as well as traditional file signatures.

SANDBOXES

Another extremely useful modern anti-malware defense is a sandbox, which grabs a copy of anything being downloaded by any device. This could be email attachments, JavaScript code, a Windows executable file, a Flash animation—really anything.

The sandbox tries to unpack the files and run them in a special, isolated virtual machine. As it does so, it watches carefully for any signs of malicious behavior. The hope is to prevent malware from ever reaching a workstation, which is particularly important for those malware strains like crypto-viruses, which deploy as soon as they're downloaded and immediately start destroying files.

Sandbox detection can usually be done so fast that the end user isn't even aware their download was intercepted. However, sometimes malware can evade detection in a sandbox. Sandboxes keep getting better at detecting malware. Malware keeps getting better at detecting sandboxes. So a sandbox should never be considered a replacement for good endpoint security.

BRINGING NETWORK DEFENSES TOGETHER

Once you have several essential elements of network security up and running, you'll quickly find you don't have the resources to look at all of them. Firewall and IDS logs alone accumulate at rates of kilobits per second, even after filtering. Active Directory and DNS logs are often just as bad. There's no way a human can monitor these systems in real time. Instead, you need a way of storing, filtering, and correlating messages into something meaningful.

Initially, it might be enough to use the management tools that come with your equipment. A good UTM firewall generally has a management GUI. You can always use the GUI on your domain controller to look at Active Directory logs. And most endpoint security systems also have a central console. But at a certain point, there are just too many different consoles, and they aren't sharing information with one another.

This is where a security incident event management (SIEM) system becomes useful. Often the SIEM doubles as a searchable long-term storage system for log messages. Some organizations deploy these functions separately, but ideally, I like them to be together.

The SIEM is a single pane of glass that correlates information it receives from the various security systems you've deployed and presents you with all the relevant information about each event.

For example, if somebody is attacking a web server, you might receive relevant data about the attack from the firewall, the IDS/IPS, the WAF, and perhaps also the Active Directory server. The SIEM rolls all of that information together so you don't have to manually search through each of the different data sources to understand exactly what's going on.

Effective security requires many tools working together. You'll often hear the expression "defense in depth", and this is what it means. One of the most important tasks is to identify your risk areas. Then allocate your budget and your time carefully to cover the highest risks first.

The other point that I really want to stress is that tools alone are not security. You can't lock the door and expect that will keep out all the burglars. Somebody has to be looking at the tools, investigating every single anomaly and eliminating the threats as they appear. Even the risks will change over time, so you need to continually re-evaluate your highest risk areas to make sure they're appropriately covered.

When networks run the world,

NETWORK MANAGEMENT IS EVERYTHING.

Gain true network visibility and control with Auvik.

USE FREE FOR 14 DAYS

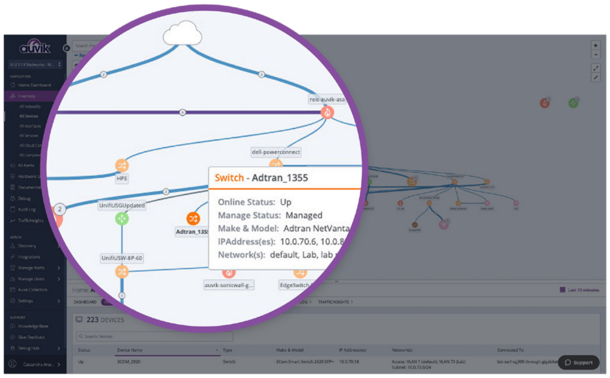

Switch - Adtran_1355

Online Status: Up
Manage Status: Managed
Make & Model: Adtran NetVanta
IPAddress(es): 10.0.70.6, 10.0.8
Network(s): default, Lab, lab

223 DEVICES

Search Devices

Status	Device Name	Type	Make & Model	IP Address(es)
Up	3COM_2920	Switch	3Com Smart Switch 2920 SFP+	10.0.70.18

Real-time network mapping & inventory

Quickly discover & audit new networks. Then, stay in the loop—you'll always know exactly what's where, even as users & devices move.

Automated config backup & restore on network devices

Mitigate network risk with no manual effort.

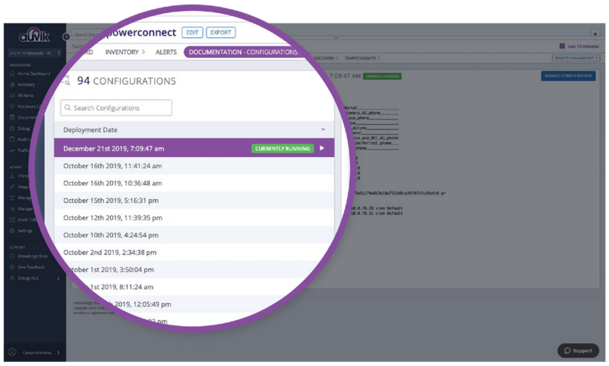

powerconnect

94 CONFIGURATIONS

Search Configurations

Deployment Date

December 21st 2019, 7:09:47 am CURRENTLY RUNNING
October 16th 2019, 11:41:24 am
October 16th 2019, 10:36:48 am
October 15th 2019, 5:16:31 pm
October 12th 2019, 11:39:35 pm
October 10th 2019, 4:24:54 pm
October 2nd 2019, 2:34:38 pm
tober 1st 2019, 3:50:04 pm
1st 2019, 8:11:24 am
2019, 12:05:49 pm

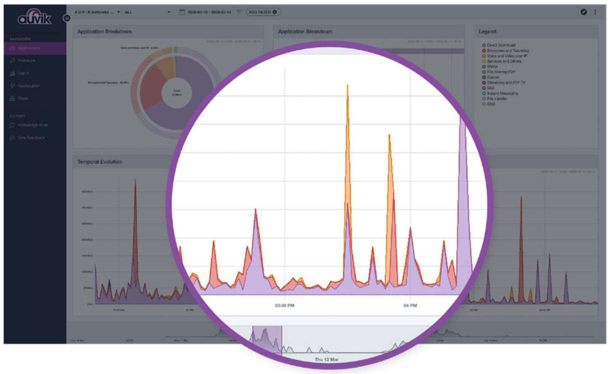

Deep insights into network traffic & flows

Quickly solve network bottlenecks & spot potential security vulnerabilities.

auvik

auvik.com/try-now