

Incorporating Threat Hunting Into Your Security Program




Contents

Introduction	3
The reality	3
A stitch in time	3
Where to begin	3
Threat Hunting vs. _____	4
Incident response	4
Penetration testing	4
Risk management	4
Compromise assessment	4
The 5 'W's	5
Why?	5
Identify who	5
When to hunt	6
The what and the where	6
The Pyramid of Pain	7
How to Hunt	8
Analyze the logs	8
Testing a theory	9
Going after the source	10
The Aftermath	11
Conclusion	11
Tools for Threat Hunting	12
About the Cisco Cybersecurity Series	13

Introduction

It's 1 o'clock and all is well. You're back from lunch, and as the company's senior SOC threat researcher, you've just reviewed your SIEM dashboards for security alerts. Nothing out of the ordinary has caught your attention. A recent automation project has drastically cut the time it takes to do this security sweep, freeing up valuable time that would have previously been spent on manual tasks. So how do you spend this newfound time?



Threat hunting is an activity you deliberately plan and regularly carry out to help strengthen your security posture.

Maybe it's time to consider threat hunting. Threat hunting involves going beyond what we already know or have been alerted to. Security software only alerts us to the risks and behaviors that we know are malicious. Threat hunting is about venturing into the unknown.

Threat hunting is an active security exercise, with the intent of finding and rooting out attackers that have penetrated your environment without raising the alarm. This is in contrast to traditional investigations and responses that stem from alerts that appear after potentially malicious activity has been detected.

The reality

Of course, this scenario could sound somewhat idealized. I mean, who really finds themselves with a free afternoon? There's always something else that needs doing, right?

The reality is that, most of the time, threat hunting isn't an activity you do on a whim. Nor is it something you do in an ongoing investigation as the next step in a procedure. Rather, it's an activity you deliberately plan and regularly carry out to help strengthen your security posture. Essentially, it's another tool in your security arsenal.

None of this sounds easy when your schedule is packed and your to-do list is as long as your arm. However, there are some key benefits to setting aside time on the calendar to perform threat-hunting activities.

A stitch in time

For starters, the identification and eradication of unknown and undetected threats is always a good thing. Even when a particular threat isn't discovered, threat hunting exercises often identify weaknesses in your environment that you can shore up and set new policies. Ultimately, the fruit borne from regular threat hunting is that it can significantly shrink the attack surface for future malicious actors.

There are also substantial opportunities to build upon what's learned during a threat hunting campaign. These exercises can identify areas where alerting for malicious behavior could be put in place, as well as where to develop automation to repeat a particular threat hunting scope. From there, you can carry out additional threat-hunting exercises, building up and extending your protections and capabilities.

Where to begin

The goal of this paper is to provide an overview of the threat hunting discipline. We'll explore the ins and outs of threat hunting, highlight why it's a worthwhile endeavor, who should be involved, what and where you should look, and when you should do it.

There are also a number of security disciplines with tasks that overlap with threat hunting. We'll compare and contrast disciplines, showing that while threat hunting is similar to other tasks, it deserves a place in your security arsenal.

Finally, we'll discuss how you can build out effective threat hunting campaigns within your organization. One of the toughest things to determine is where to start. To assist, we begin with the simple steps you can take to begin to build up your threat hunting posture, strengthening your organization's security in the process.

Threat Hunting vs. _____

As far as security disciplines go, threat hunting is a comparatively young specialty. Given this, there are overlaps with other security-related practices. In fact, many folks currently involved in threat hunting have experience with these other roles within their careers. The following are some quick comparisons to other disciplines.

Incident response

This role is perhaps the most similar to threat hunting. Both disciplines deal directly with threats in your environment. The primary difference is that incident response is reactive—you know something is on the network, or at least has tried to access the network, due to security alerts, network or endpoint behavior, or other evidence. In contrast, in threat hunting, there isn't necessarily any evidence of a threat. Instead, you're actively looking for something instead of trying to contain and remediate what you know is there.

Penetration testing

Threat hunting and penetration testing also share some similarities. At their heart, both attempt to seek out weaknesses in a network. However, penetration tests generally look for configuration problems or known vulnerabilities in order to gain access to a network or sensitive information. The goal of threat hunting isn't necessarily to gain access to anything, but rather identify hidden threats present in an environment, eradicate them, and set up policies to prevent them in the future.

Risk management

The idea with risk management is to determine weaknesses within the network or on systems, determine their severity, prioritize, and then take appropriate steps to correct them. This may involve identifying threat sources, and threat hunting may help to inform a risk assessment. However, such assessments generally cover far more ground than threat hunting, looking at all potential risks, both known and unknown.

Compromise assessment

Also similar to threat hunting, compromise assessment is about finding out if your network has been breached by unknown, bad actors. However, it is a much broader exercise than threat hunting. During compromise assessments various tools are installed across a network, looking across the board for anything out of the ordinary. In contrast, threat hunting begins with a very particular idea or scenario and maintains focus on that scope.

The 5 "W's"

Figuring out where to start can be challenging when establishing threat hunting exercises within your organization. Utilizing the five "W's," often used in journalism, can be a good way to begin planning out the process.

Why?

The up-front investment in proactive threat detection can strengthen an organization's security posture significantly. The fact is, organized, skilled, and well-funded attackers exist. If you become a target of one of these such groups, they can work diligently looking for a weakness to get in. Sadly, you can't possibly uncover everything with even the best security tools. This is where threat hunting comes in—its primary mandate is to find just these types of attackers.

An added bonus to threat hunting is that carrying out such exercises breeds familiarity with tools and techniques that are so important when an outbreak or breach occurs. Your threat hunting team will likely overlap with your incident response team and threat hunting sharpens their skills and response times when faced with a real incident. It can be looked at as practice for when things go wrong.

Identify who

Building that threat hunting team may seem as daunting as assembling a team of superheroes to work towards defeating a common enemy. Part of assembling that team is pulling folks together with different skill sets and backgrounds.

If you're a large organization, then the first step may be as simple as setting aside a block of time during the month for a group, or

tiger team, to plan, perform, and report on a threat hunting campaign. However, if you're a small organization with only a couple (perhaps only one!) dedicated IT person, this may not be so easy. Given this, you may want to bring in a third-party, external expertise to help. This carries advantages and disadvantages. On the plus side, you'll likely get access to people that fulfill the skills requirements of threat hunting. However, an external threat hunting team will not be as familiar with the ins-and-outs of your specific network as internal personnel will be.

Regardless, there are a mix of core skills needed in a team in order to carry out a threat hunting campaign:

- **Familiarity with endpoint and network security**


This almost goes without saying. You'll need seasoned members of your SOC or IT team that have an extensive breadth and depth of knowledge around security issues and best practices.

- **Understanding of data analytics**

Oftentimes threat hunting requires teasing patterns out of raw data. Having an understanding of statistical analysis will help to identify patterns in the data. Data visualization is equally important in order to spot and share anomalies that are found.

- **An innate curiosity**

Threat hunting is not a cut-and-dry exercise. It can sometimes be likened to an artistic pursuit. It requires a certain amount of creative thinking, connecting seemingly unrelated items or asking, "I wonder what would happen if..."



Your threat hunting team will likely overlap with your incident response team and threat hunting sharpens their skills and response times when faced with a real incident.

One bonus to threat hunting, from the perspective of a security professional, is it's fun. Threat hunting gives folks in your SOC or IT department a break from the day-to-day reactive nature of their roles and a chance to go on the offense. Such active, fulfilling tasks for employees can often lead to higher retention rates for SOC employees, retaining them in a field where highly qualified people can be hard to come by and often move around.

When to hunt

Ultimately, the most successful hunts are those that are planned. You need to set a scope for the hunt, identify clear goals, and set aside a block of time to perform the hunt. When you're done, you need to assess steps to improve your security posture, establishing security playbooks to address the results moving forward.

At other times you may also wish to undertake a threat hunting exercise when you suspect risky behavior may have occurred.

- **Is a particular user downloading far more data on a given day than normal?**
- **Does a user attempt to log into a system that he or she doesn't have access to?**
- **Did an administrator appear to clear his or her bash logs?**

Many of these behaviors could indicate the actions of a malicious actor having compromised a device and is a fairly straightforward place to begin a threat hunt.

Finally, there are times where a threat hunt may crop up unexpectedly. Has a cybersecurity news story that caught your CIO's attention ever lead to an email or phone call inquiring if

the company is vulnerable? This is a perfectly valid question and having a process in place to field inquiries like this can save a significant amount of time and resources.

The what and the where

Ultimately data is key to any threat hunt. Before you can do anything threat-hunting related, you'll need to ensure you have adequate logging enabled to carry out the hunt. The fact is, if you can't see what's happening on your systems, then you can't respond in kind.

Choosing which systems to pull from will often depend on the scope of the hunt—one hunt it could be endpoints in the Finance department, another could focus on web servers. In some cases, you may even want to install tools within the environment to monitor particular types of traffic. The logs pulled by these temporary systems will then be utilized in the hunt.

Of course, enabling logging can quickly fill up storage assets and gathering logs can easily eat into your team's time. This may require setting aside physical resources to store logs and setting up basic automation to send them there. In the short term you may have to be selective about how extensively you configure the systems to log. Utilizing tools such as security information and event management (SIEM) software can go a long way towards making the analysis of logs faster and easier.

In the first few threat hunting exercises, the result may include a list of questions that couldn't be answered, based on the logs available. In time it will become clearer which systems need to have logging enabled, and at what level, in order to get the results that are desired.



Before you can do anything threat-hunting related, you'll need to ensure you have adequate logging enabled to carry out the hunt.

The Pyramid of Pain

Security researcher David Bianco came up with an approach titled the [Pyramid of Pain](#), which outlines how to cause adversaries the most difficulty when attacking your network. Each of the six layers represent different approaches you can take, starting with the simple and working your way up to the most difficult.

For instance, at the base of the pyramid are hashes. Files bearing known malicious hashes are simple to detect, and also simple for the attacker to replace. The same goes for IP addresses, although this takes a little more work, both to find, and for an attacker to replace, hence a smaller piece of the pyramid. Domains are a little bit harder, network artifacts harder still, etc.

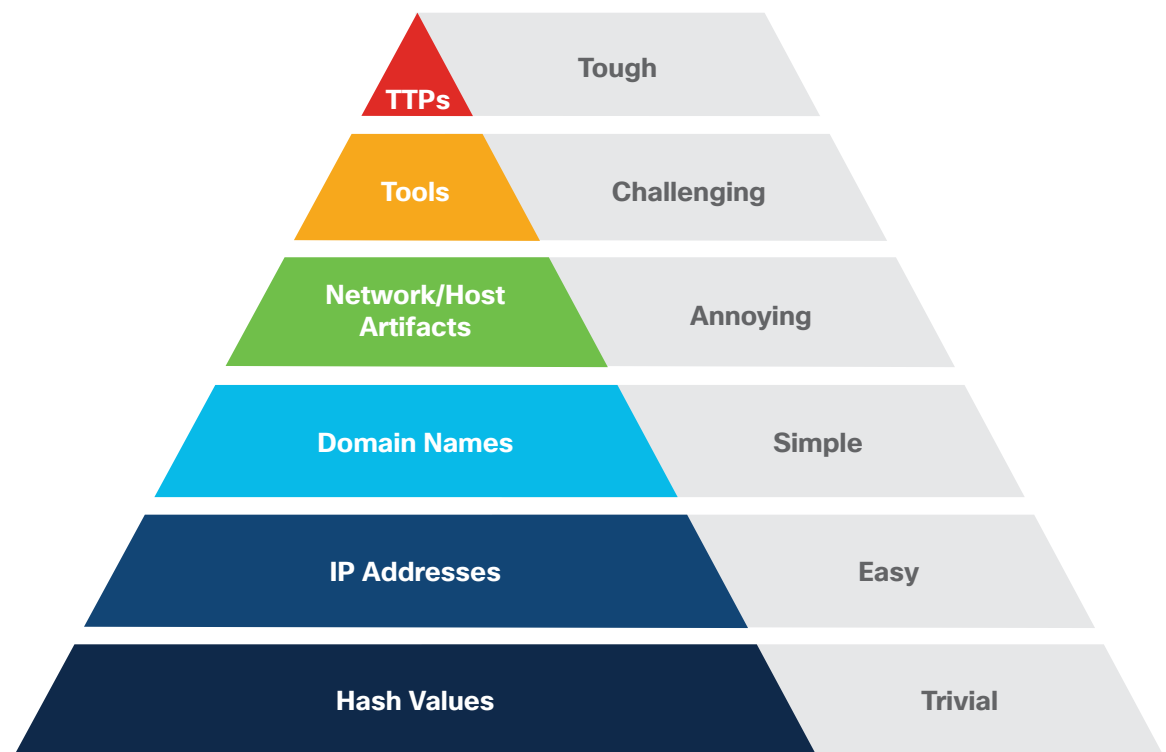
The goal of your threat hunting exercise should be to uncover an attacker's tactics,

techniques, and procedures (TTPs). These are the most valuable because they are hard for the attacker to replace. It's often the most difficult and/or time consuming to identify, mainly because it requires comparing data points from different data sets and making connections where the relationship isn't apparent at first.

The trick is, as you go up the pyramid, you force the adversaries to spend more resources in attacking your network, making it more difficult and increasing the chances that they will be caught doing so. The ultimate goal of the Pyramid of Pain is that, by following its principles, your network becomes so challenging to hack that the attackers move on to other, simpler targets.



The goal of your threat hunting exercise should be to uncover an attacker's TTPs - the most valuable IoCs because they are hard for the attacker to replace.



Source: David J. Bianco, personal [blog](#)

How to Hunt

As to the how, there are a number of ways to approach a threat hunting exercise. The resources and skills you have available will play into how detailed a threat hunting campaign is carried out.

In the following section we start out with simple, basic ways to get started in threat hunting, then work our way up in complexity. The idea here is that, after every threat hunting exercise, you can build upon what you've learned. Establishing playbooks, automation, and policy changes where needed gives you a foundation to move on into more advanced techniques.


Analyzing the logs

Sometimes, the simplest threat hunting activities stem from research or reports on newly discovered threats. It's common practice these days to include indicators of compromise (IoCs) alongside research for others to use. These data points are generally comprised of IP addresses, URLs, domains, file hashes, or other IoCs that comprise a threat.

One of the simplest ways to kick off a threat hunting exercise is to check the logs from your systems against IoCs. Command line tools or simple scripts can be enough to get you started. Using a SIEM is another method to quickly compare IoCs to logs. There are also more advanced security products that can help facilitate threat hunting by allowing you to copy and paste IoCs into a dashboard to see if they have been seen in your environment.

Once you become comfortable with these activities, it's time to dive deeper into the logs and start discovering new IoCs that may exist. This is where data analytic skills come into play. Applying statistical models to logs,

such as [clustering](#) or [frequency distribution](#), can help shed light on anomalies. Ultimately you're hoping to reach the top of the Pyramid of Pain and identify the TTPs of an attacker.



Threat Hunting in Action

Jeff Bollinger manages security investigations for CSIRT here at Cisco. The following is a first-hand account of a threat hunting exercise his team carried out.

"Looking through historical Cisco AMP endpoint data for indicators of compromise, we saw a suspicious binary dropper that had been deleted by the user.

We recovered the binary by restoring the single file from the user's (corporate) backup archive and were able to reverse it and extract additional indicators (C2 hostnames) that we then applied to all our network telemetry.

This yielded additional affected hosts that didn't trigger on the original dropper's hash."



If you're the seasoned vet on the team, don't think you've seen it all before. Instead, try to prove that it isn't a threat. If you can't do so offhand, then dig deeper.

Testing a theory

Some may argue that simply checking the logs against known IoCs isn't true threat hunting. The reasoning goes that you're simply matching 1:1. In these cases, to qualify as threat hunting, you have to dig deeper than that.

This is where creativity plays a part. You have to come up with a theory about where a threat may reside, the vectors it may have used to get there, or the techniques it exploited. The following are a few ideas of the sort of investigations you might carry out.

• Read security news

The latest threat landscape news can be full of material for a threat hunt. For instance, if there was a critical vulnerability in a Windows process recently disclosed, investigate if any odd activity has occurred surrounding that process. Of course, pay particular attention to material that applies to your industry. For example, if you work in aviation, a credit card stealer wouldn't be a high priority. Conversely, if you work in banking, a threat found attacking an ICS wouldn't apply.

• Look into reports of strange behavior

Investigate unusual reports of activity from staff. Are sleeping systems suddenly waking up during the night? Investigate what is waking them. Has an office reported that they found internal data on an external source? Look for indications of data exfiltration.

• Filter the normal to find the abnormal

Unusual activity is a good starting point, but not always easy to spot. Sometimes you have to dig through the tall grass to

find it. Look at a particular activity with a malicious goal in mind. For example:

- Look for long network connections, which could be a sign of data exfiltration. Filter out those that are expected and see if any of those that remain are suspicious.
- Look at CPU activity spikes and the processes that create them, which could indicate cryptomining or an infostealer logging activity. Filter out those that are well known and examine those that are not.
- What sort of files is the BITSAdmin tool downloading? It could be used to pull down malicious tools, as many threats use local tools to mask their actions. Clear out the regular downloads you're expecting and focus on the rest.
- Look at scheduled tasks. Attackers may add their own tasks to kick off certain malicious activity. Are there any that are not run by system administrators? Investigate any that seem suspicious.

Any cases where the behavior seems out of the ordinary are prime areas to dig deeper and find the root cause. However, it's important to approach anything found with an edge of caution. Just because something looks weird, doesn't necessarily mean it's a bad actor. Be sure to compare your findings against other data sources before reaching any conclusions. At the same time, if you're the seasoned vet on the team, don't think you've seen it all before. Instead, try to prove that it isn't a threat. If you can't do so offhand, then dig deeper.

Going after the source

You've managed to identify a threat within your network, pinpoint what allowed them to get in, and take measures to prevent it from happening again. However, the next time you run a threat hunting exercise, you find the attackers have gotten back in another way.

If you're constantly finding your organization the victim of attacks, it might be worth your while to investigate who is attacking, the infrastructure that they are using to attack, and attempt to get the group shut down.

However, this is not a suggestion to practice offensive hacking. As tempting as it may be, there are a number of problems with going that route.

For starters, if you attack a malicious infrastructure, there's a good chance the attackers will notice and hit back twice as hard. However, their motivation this time may not be to steal info, but rather revenge—disabling or destroying systems as they go.

Another reason not to hack back is that in most locations in the world doing so is illegal. Despite the fact that the systems in question are performing illegal activities, offensive hacking is still hacking.

The good news is that there is still plenty that can be done. The IoCs of an attack can reveal a lot about the attackers without even having to touch their networks.

The best approach to get malicious actors shut down is to gather up any IoCs you can uncover, from hashes all the way to TTPs, build a profile of the attacker, and then turn these details over to the appropriate law enforcement agencies. These authorities are the best method to pursue and shut down an attacker through legal means.



The best approach to get malicious actors shut down: gather up any IoCs you can uncover, build a profile of the attacker, and turn these details over to the appropriate law enforcement agencies.

Of course, for all but the largest and most targeted organizations, this isn't always something that can be easily done in-house. As a result, the lion's share of organizations can and should rely on external security research teams that have made investigating such attacks their mandate. Threat intelligence organizations, like [Talos Intelligence](#) or [Cisco's Incident Response Services](#), are here to help in such cases.



Leveraging Threat Hunting

Sean Mason, the Director of [Incident Response Services](#), reflects on how his teams have leveraged threat hunting.

"I really started to understand and value threat hunting in 2011 on the heels of the [RSA hack](#). I found myself in meeting after meeting discussing how we could detect this type of threat. It really made us think differently. We also realized what sorts of visibility gaps we had. Over the years the various teams I've been on have leveraged hunting in many different ways: either proactively following a hunch, responding to an incident, or being diligent after reading the latest security news. I can honestly say that after more than eight years of leveraging threat hunting in various capacities, it's a no-brainer that I consider it a critical component for every successful security program."

The Aftermath

As important as it is to identify and eradicate threats hidden in your network, figuring out how they got in and taking steps to prevent future attacks is perhaps the most important aspect of threat hunting. Plan to have a post-op meeting to discuss the hunt. In it show what's been found and discuss what needs to be done to fix it. Then implement network policy changes to lock it in.

Sometimes it's less about finding a threat, but rather uncovering weaknesses within the organization. A successful threat hunting campaign may uncover a misconfigured server or a policy violation that needs correcting. And as counter-intuitive as it may seem, sometimes the best threat hunting campaigns uncover nothing at all. The benefit here is you now tangibly know that the avenue investigated is not currently a risk to your organization.

Adding automation is another critical post-threat hunt step. After a threat hunt is complete, it's important to check periodically to see if the activity you've uncovered returns. Convert what has been found into a process that can be run again. Set up a trap with alerting when triggered. Over time this will fold into your security playbook.



Sometimes the best threat hunting campaigns uncover nothing at all. The benefit here is you now tangibly know that the avenue investigated is not currently a risk to your organization.

Conclusion

There's no way to ever know if your network is completely free of threats. That doesn't mean the pursuit is futile. The benefit of threat hunting, besides uprooting threats that managed to get by your defenses, is that you can build up your security posture further.

Think of threat hunting as you would masonry. When building a house, start with that first ring of bricks, add mortar to hold them in place, then add another layer of bricks. Repeat the process layer by layer, building up the walls.

With threat hunting, that first layer of bricks could be turning on sufficient logging and storing it. The mortar is the automation that keeps those logs coming in regularly. The next layer of bricks is comparing logs against IoCs. Automate those processes to hold the bricks in place. Keep learning with layers of data analytics, testing theories, etc.

Pretty soon, you've built a strong and stable threat hunting process that will give you the peace of mind that your organization is as free of threats as the environment can be.



Tools for Threat Hunting

The following are some recommended tools that can be used for threat hunting. While the list is far from exhaustive, they will help when starting off.



Cisco Threat Response

Cisco Threat Response automates integrations across select Cisco Security products, applies threat intelligence from Cisco Talos and third-party sources against security events to automatically research indicators of compromise (IoCs) and confirm threats quickly. It also provides the capability to collect and store key investigation information, to manage and document your progress and findings, and remediate threats directly from the dashboard.



Cisco Threat Grid

Threat Grid combines advanced sandboxing with threat intelligence into one unified solution to protect organizations from malware. With a robust, context-rich malware knowledge base, you will understand what malware is doing, or attempting to do, how large a threat it poses, and how to defend against it.



Cisco Stealthwatch

Cisco Stealthwatch is a comprehensive visibility and network traffic and cloud security analytics solution. It can even detect malware in encrypted traffic without decryption. It provides advanced threat detection, accelerated threat response, and simplified network segmentation using multilayer machine learning and entity modeling. With advanced behavioral analytics, you can find out who is on your network or in your public cloud infrastructure and what they are doing.



Cisco Advanced Malware Protection (AMP) for Endpoints

Not only does AMP protect your endpoints, but it can assist in malware analysis and proactive threat hunting. AMP's robust search capabilities allow you to find various information, like file, hash, URL, IP address, registry keys, users, processes, applications, and much more. It can also show the lifecycle of a file in your environment, from the first time it was seen, what it did on the endpoint, and other intelligence.



Umbrella Investigate

Investigate provides the most complete view of the relationships and evolution of domains, IPs, autonomous systems (ASNs), and file hashes. Accessible via web console and API, Investigate's rich threat intelligence adds the security context needed to uncover and predict threats.

Security information and event management tools (SIEMs)

Having a SIEM is a key step in carrying out threat hunting activities, especially when starting out. A well configured SIEM can greatly reduce the amount of time spent gathering log files and performing basic analysis. Examples of well-known SIEMs include [Splunk](#), [IBM QRadar](#), and [Exabeam](#).

Endpoint monitoring tools

There are a variety of tools available to collect detailed logs from endpoints. Windows built-in Event Log is a good place to start, and more complex tools such as [Sysmon](#) and [Process Monitor](#) can extend your logging capabilities. (There are even [pre-built configurations](#) to help you get started.) On Apple Macs, check out [Console](#) to view logs.

Packet analyzers

These are tools that can be used to monitor your network traffic. Applications like [Wireshark](#) and [tcpdump](#), and APIs like [pcap](#) are popular tools for gathering information about the data being transferred across your network.

About the Cisco Cybersecurity Series

Throughout the past decade, Cisco has published a wealth of definitive security and threat intelligence information for security professionals interested in the state of global cybersecurity. These comprehensive reports have provided detailed accounts of threat landscapes and their organizational implications, as well as best practices to defend against the adverse impacts of data breaches.

In our new approach to our thought leadership, Cisco Security is publishing a series of research-based, data-driven publications under the banner: Cisco Cybersecurity Series. We've expanded the number of titles to include different reports for security professionals with different interests. Calling on the depth and breadth of expertise in threat researchers and innovators in the security industry, the previous collection of reports in the 2019 series include the Data Privacy Benchmark Study, the Threat Report, and the CISO Benchmark Study, with others to come throughout the year.

For more information, and to access all the reports and archived copies, visit www.cisco.com/go/securityreports.



Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA), Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

Published September 2019

THRT_05_0919_FINAL

© 2019 Cisco and/or its affiliates. All rights reserved.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: www.cisco.com/go/trademarks. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)