ExtraHop

SECURITY ADVISORY

Remote Desktop Protocol in a Distributed Workforce

EXECUTIVE SUMMARY

The COVID-19 pandemic has reshaped many aspects of business operations, including the need to secure and enable remote work on a large scale. For many organizations, the most immediate solution was to configure RDP– and VPN–based work-from-home environments that would allow workers to access applications and data via personal devices and home networks. But the rush to operationalize these resources for hundreds, if not thousands, of employees resulted in misconfigurations that cyber attackers were fast to exploit. This report explores common malicious RDP activities and strategies for detection and mitigation.

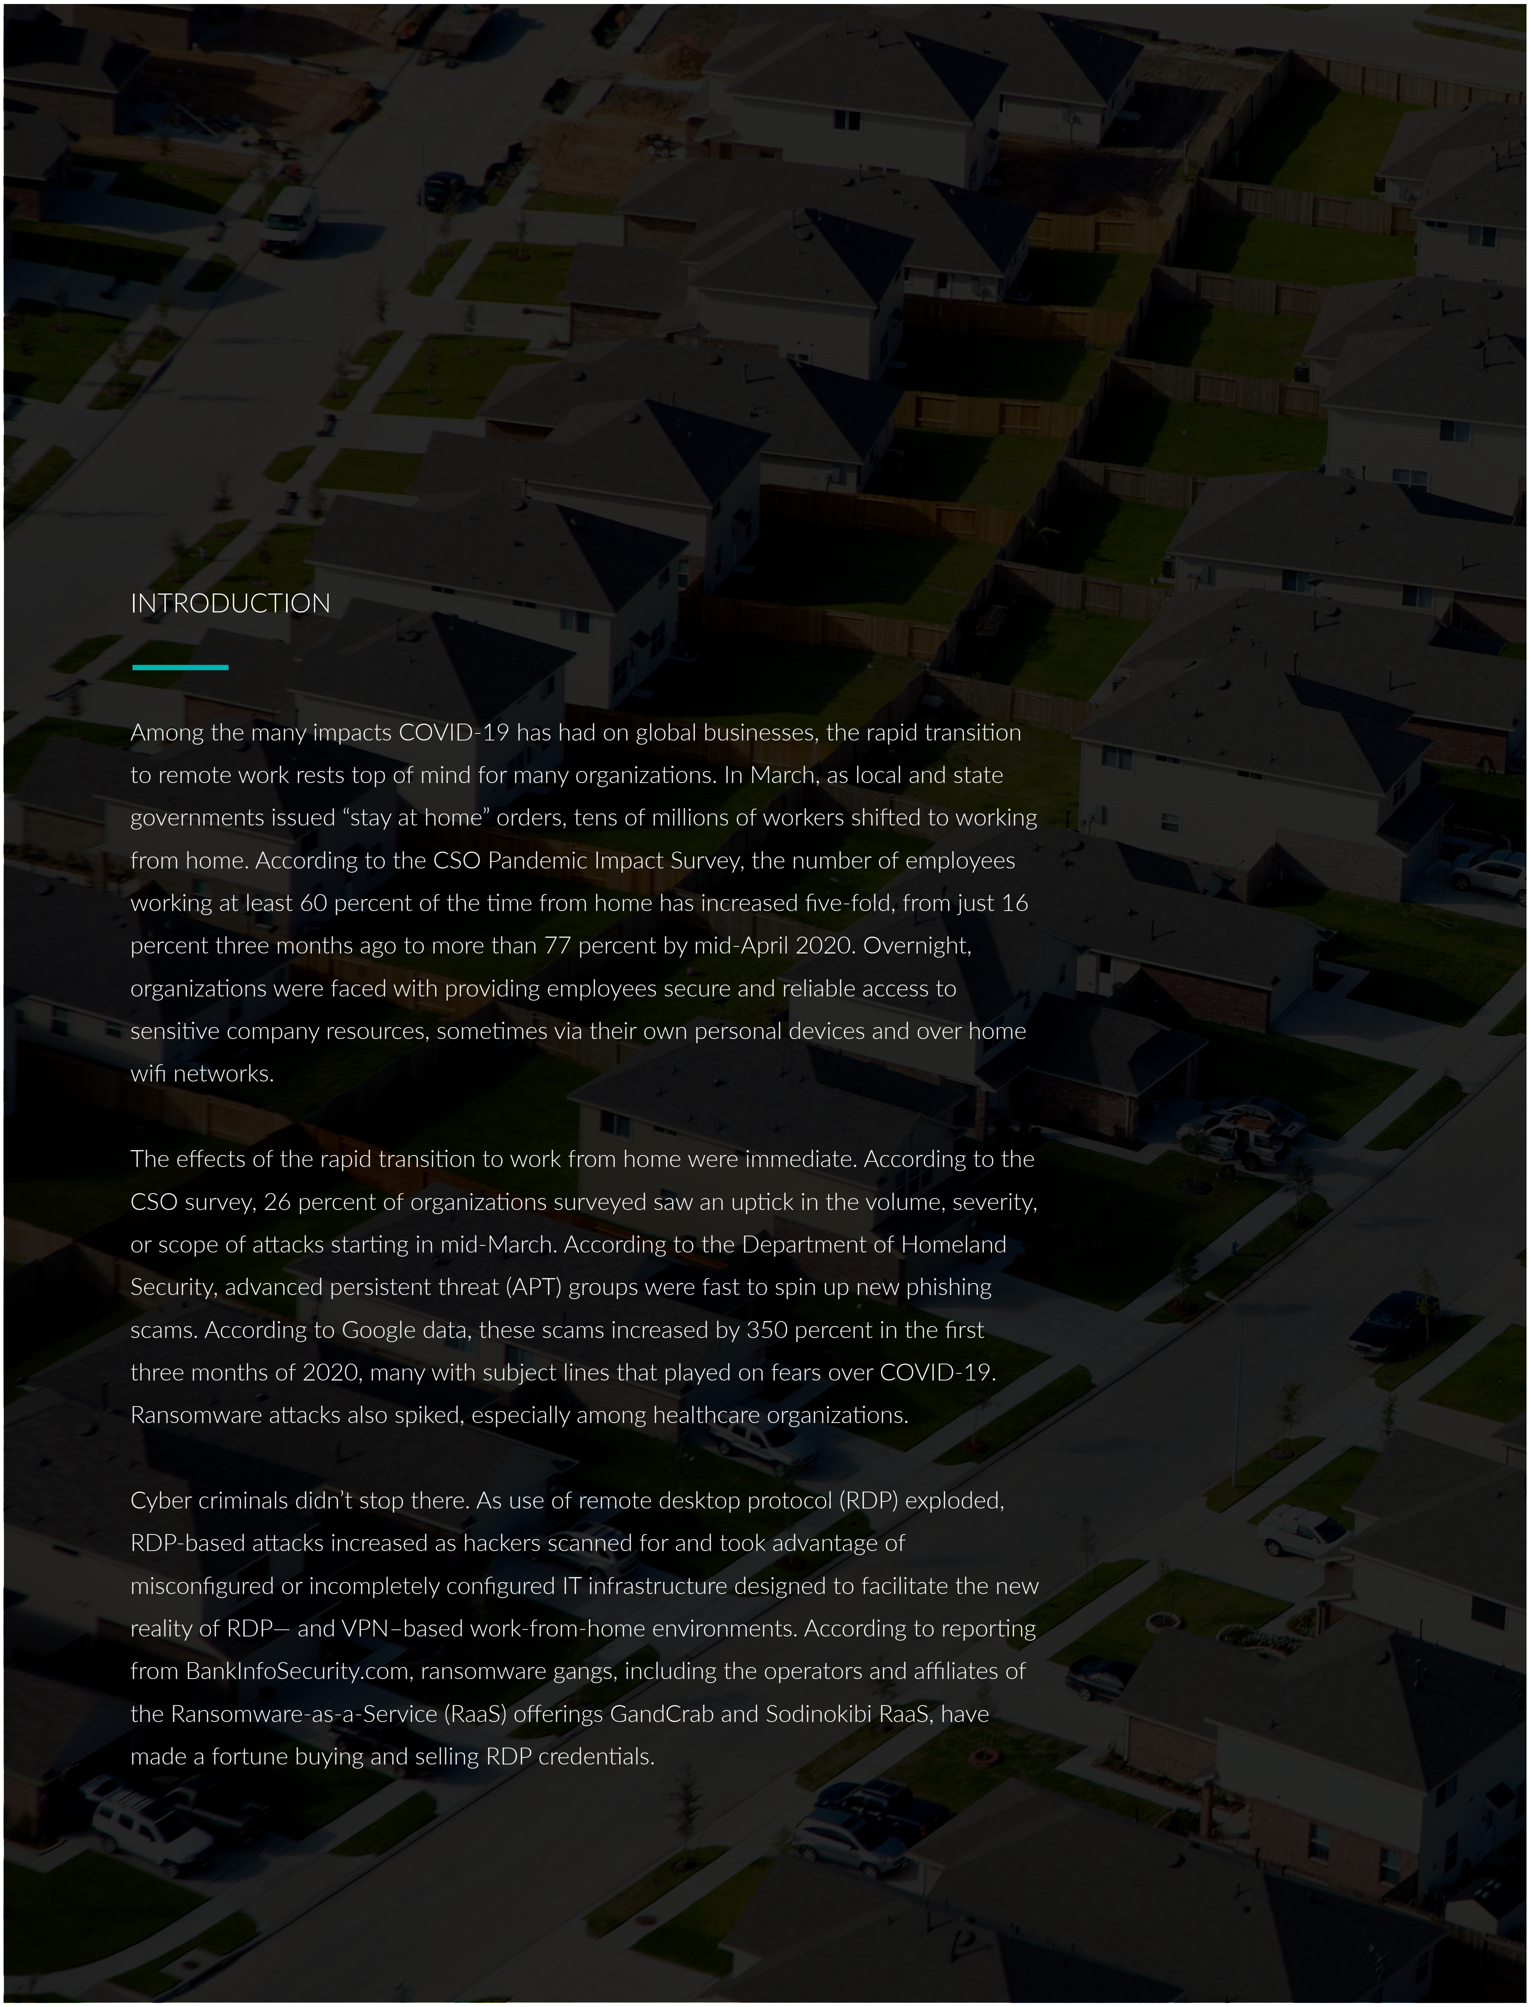

INTRODUCTION

Among the many impacts COVID-19 has had on global businesses, the rapid transition to remote work rests top of mind for many organizations. In March, as local and state governments issued “stay at home” orders, tens of millions of workers shifted to working from home. According to the CSO Pandemic Impact Survey, the number of employees working at least 60 percent of the time from home has increased five-fold, from just 16 percent three months ago to more than 77 percent by mid-April 2020. Overnight, organizations were faced with providing employees secure and reliable access to sensitive company resources, sometimes via their own personal devices and over home wifi networks.

The effects of the rapid transition to work from home were immediate. According to the CSO survey, 26 percent of organizations surveyed saw an uptick in the volume, severity, or scope of attacks starting in mid-March. According to the Department of Homeland Security, advanced persistent threat (APT) groups were fast to spin up new phishing scams. According to Google data, these scams increased by 350 percent in the first three months of 2020, many with subject lines that played on fears over COVID-19. Ransomware attacks also spiked, especially among healthcare organizations.

Cyber criminals didn’t stop there. As use of remote desktop protocol (RDP) exploded, RDP-based attacks increased as hackers scanned for and took advantage of misconfigured or incompletely configured IT infrastructure designed to facilitate the new reality of RDP— and VPN–based work-from-home environments. According to reporting from BankInfoSecurity.com, ransomware gangs, including the operators and affiliates of the Ransomware-as-a-Service (RaaS) offerings GandCrab and Sodinokibi RaaS, have made a fortune buying and selling RDP credentials.

Why Now?

Last summer, the BlueKeep exploit had IT organizations scrambling to find and disable RDP servers exposed to the Internet. But the need to maintain business operations during the COVID-19 pandemic has required some teams to open RDP so that employees at home can access their computers in the office. During an April panel discussion put on by ISC², Glenn Leifheit, Senior Security Program Manager of Customer Security and Trust at Microsoft, noted that Microsoft has seen a record spike in the number of virtual desktops — which use the RDP protocol.

COMMON MALICIOUS RDP ACTIVITIES

Remote Desktop is a common target for attackers, and it's easy to see why. It's prevalent in enterprise environments, it provides remote access to a Windows device, and it leaves credentials exposed in memory. Devices with active Remote Desktop sessions can also be enumerated easily with attack tools.

When a cyber criminal gains access to a windows device, it can be used for a variety of malicious activities capable of compromising not just the windows device, but devices and data across the network. Below is a list of some of the most common types of malicious RDP activity.

Exfiltration

If an attacker gains access to a poorly-secured RDP device, they can easily transfer data. Unusual data transfers can be associated with suspicious activity such as sharing malicious files between compromised devices or data staging. Data staging is the process of collecting and preparing data for exfiltration. Depending on the sensitivity of the transferred files, the impact can be devastating if important, proprietary, or customer data is leaked.

Command & Control

Attackers may use RDP to gain access to a network using Command & Control techniques. These include devices outside of the network attempting or making connections with in-network Windows devices using compromised credentials. Attackers may also attempt to use RDP to control a Windows device via an external endpoint.

The impact to a business can be low if the RDP connection is not authenticated, or if the attacker connects to a device with limited privileges. However, these activities should be examined before they facilitate critical and costly attacks.

Exploit

Brute force attacks on RDP are low cost and relatively easy to perform. Even though this type of brute force attack is noisy, it can be highly effective due to the commonality of weak and repurposed passwords. An attacker might run tools such as Ncrack or Hydra to perform a brute force attack on RDP accounts to find weak passwords or valid login credentials—or perform a scan

Devices with active Remote Desktop sessions can be enumerated easily with attack tools.

to discover and exploit known vulnerabilities. Once an attacker has accessed passwords or valid login credentials, they can easily open multiple RDP sessions from a single device to control many devices on the network.

One of the most famous types of RDP exploits is BlueKeep, a well-known Remote Desktop vulnerability that has been implemented in a number of exploits. It allows an unauthenticated attacker to remotely run arbitrary code on an RDP server to grant themselves administrator access to a network-accessible Windows system without user credentials. The attacker can then tamper with data or automate the process by installing malware that could propagate to other Windows devices across the network. Remote Desktop is commonly enabled on devices, which increases the likelihood that this threat will significantly affect or even halt business services.

Reconnaissance

Devices with active Remote Desktop sessions can be enumerated easily with attack tools. This type of reconnaissance does not negatively affect network performance, but an attacker could locate Windows devices to target.

EXAMPLES OF RDP ATTACKS

RDP-based threat activity can take many forms, and can require a spectrum of detections, from signature-based to behavior modeling, to determine whether the activity is malicious, whether it has resulted in a breach, and how many systems are affected. Starting in March 2020, a number of ExtraHop customers started to see an uptick in these types of detections.

Brute Force in Healthcare

One notable example took place on March 20 during a proof of concept with a company in the healthcare device industry. Within five minutes of configuring VPC Traffic Mirroring and tagging the customer's cloud instances, Reveal(x) alerted on an RDP Brute Force attack in progress against one of their important systems. The brute force attack showed inbound attempts from every continent but Antarctica. Reveal(x) also found several successful connections indicating that credentials had been compromised.

The company in question was able to identify root cause immediately: a misconfiguration. The wrong security group had been applied to this instance, leaving inbound access on TCP and UDP 3389 (the ports for RDP) open to the internet. With the misconfiguration identified, the security team was able to apply the correct security group to rapidly contain the damage and stop the attack.

Within five minutes of configuring VPC Traffic Mirroring and tagging the customer's cloud instances, Reveal(x) Cloud alerted on an RDP Brute Force attack in progress

Suspicious Insider Traffic

Another notable example comes from an insurance provider based in North America and illuminates how RDP can be used improperly by employees. Reveal(x) detected an application server attempting to install new code on a production database (SQL) server during non-approved hours. No change request was filed for the update, so the security team decided to take a deeper look.

Looking at the application server device overview in Reveal(x), the security team identified peers coming in from VPN space. In total, Reveal(x) identified half a dozen different sessions, all from VPN space, all from a single laptop assigned to a developer. The developer used VPN to get a command shell on the app server, and then attempted to use that command line to access special resources on the database server — something that's not typically done in production. By looking at the packets, the security team was also able to investigate exactly what the VPN user did from the app server to the database server.

While this activity ultimately wasn't malicious, it was outside the accepted business process. It underscores the ways in which RDP can be misused even by employees in ways that could expose an organization and its data.

Below are additional examples of RDP activity detected by Reveal(x) in March across different customer environments.

- A law firm based in Europe started to experience inbound, established RDP connections that weren't expected. The behavioral models Reveal(x) had established for the environment hadn't seen this traffic pattern before, and ExtraHop also got a threat intelligence hit on the source of the inbound connections on a known bad source. The customer was alerted due to unexpected behaviors as well as a threat intelligence hit, and was able to shut down the RDP connections.
- A North American financial services organization was alerted to unusual remote RDP traffic, combined with a spike in RDP session attempts. The traffic was identified as interactive and likely related to command and control activities, as the server in question had reached out to the internet. This same environment saw the spike in RDP session attempts from a public source. Both of these behavioral models with Reveal(x) fired detections.
- An educational institution uncovered unusual remote and interactive traffic from an external endpoint, reaching back into the environment. This traffic had not been seen before and wasn't expected in the customer's environment.
- A large retail company was experiencing RDP brute force attempts, followed by subsequent suspicious RDP sessions which indicated a successful login. Using Reveal(x), the security team dug deeper into the traffic and identified other malicious activity including scanning, LDAP

authentication attempts, WMI calls to server systems in the environment, and CIFS exfiltration. Collectively, this activity served as an early indication of an active intrusion. Lateral movement, privilege escalation and other malicious behaviors were correlated back to the server with the suspicious RDP sessions.

DETECTING & MITIGATING MALICIOUS RDP ACTIVITY WITH REVEAL(X)

ExtraHop Reveal(x) can detect a broad spectrum of malicious activity associated with RDP, as well as devices configured to enable RDP access. Below is a list of detections and associated descriptions.

The list of detections in this doc is not comprehensive, and the titles and descriptions might change.

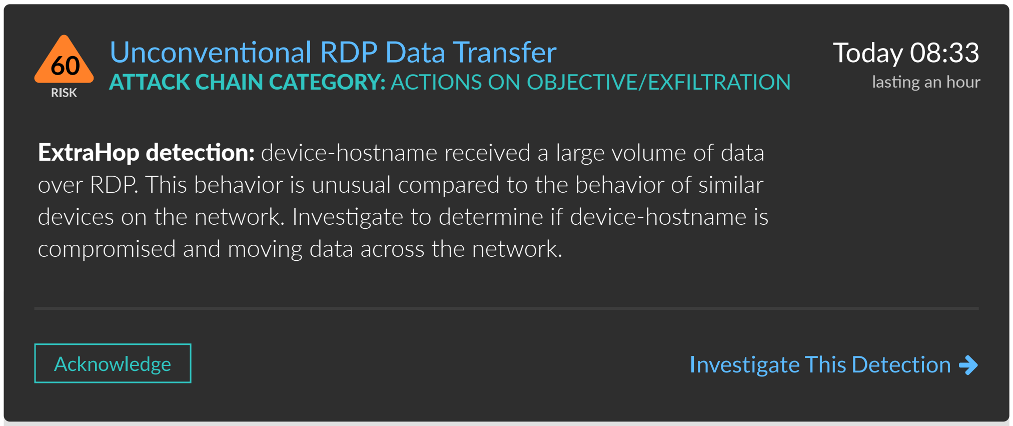

Unconventional RDP Data Transfer

ATTACK CHAIN CATEGORY: ACTIONS ON OBJECTIVE/EXFILTRATION

Today 08:33

lasting an hour

ExtraHop detection: device-hostname received a large volume of data over RDP. This behavior is unusual compared to the behavior of similar devices on the network. Investigate to determine if device-hostname is compromised and moving data across the network.

Detection Mitigation Options:

- Disable remote desktop services unless required.
- Implement microsegmentation by adding secure zones based on the zero-trust security model: partition network traffic with endpoint firewalls, virtual or software-defined networks, or physical networks.
- Review access controls to ensure that only necessary users can connect to remote access services.
- Review authentication methods and enforce policies for secure credential creation and multi-factor authentication.

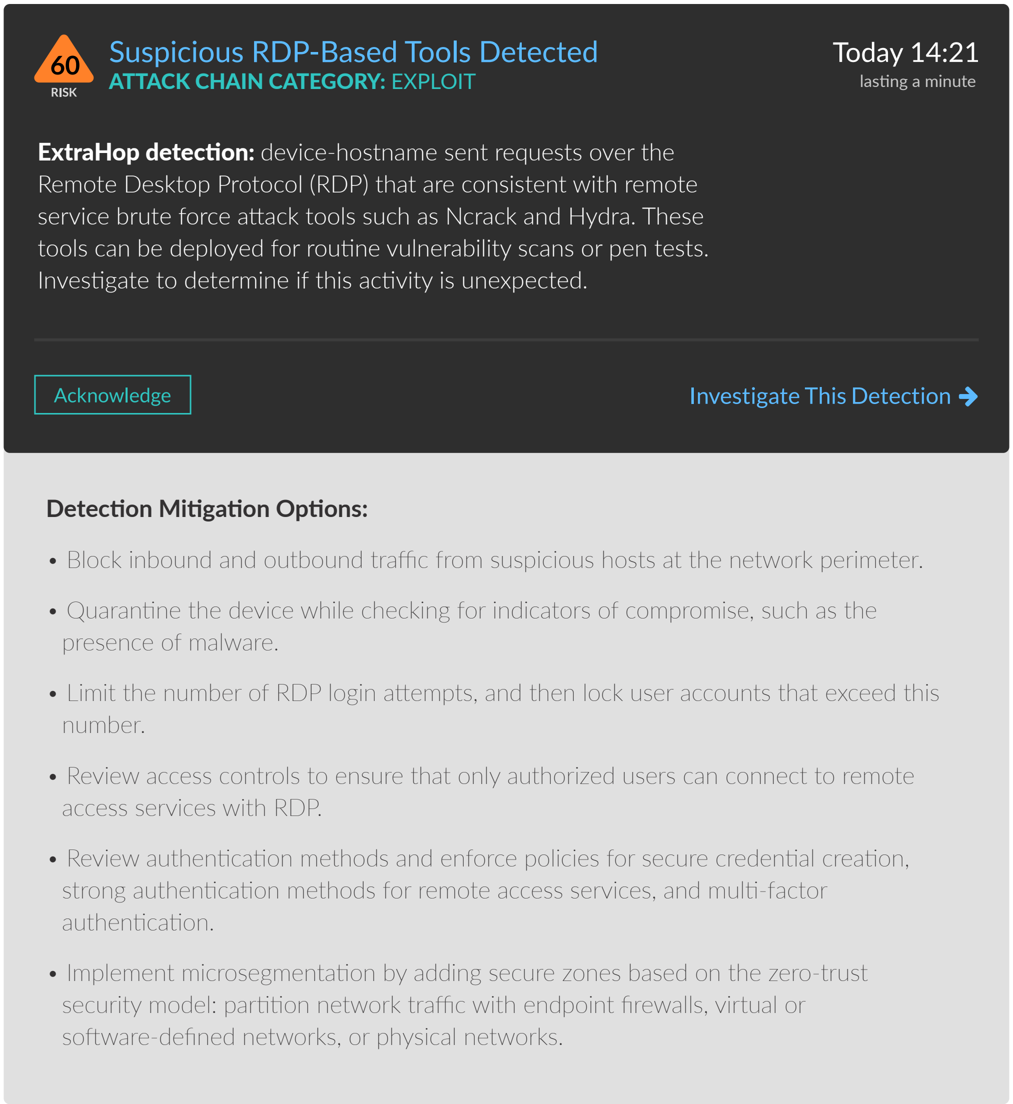

Suspicious RDP-Based Tools Detected

ATTACK CHAIN CATEGORY: EXPLOIT

Today 14:21

lasting a minute

ExtraHop detection: device-hostname sent requests over the Remote Desktop Protocol (RDP) that are consistent with remote service brute force attack tools such as Ncrack and Hydra. These tools can be deployed for routine vulnerability scans or pen tests. Investigate to determine if this activity is unexpected.

Detection Mitigation Options:

- Block inbound and outbound traffic from suspicious hosts at the network perimeter.
- Quarantine the device while checking for indicators of compromise, such as the presence of malware.
- Limit the number of RDP login attempts, and then lock user accounts that exceed this number.
- Review access controls to ensure that only authorized users can connect to remote access services with RDP.
- Review authentication methods and enforce policies for secure credential creation, strong authentication methods for remote access services, and multi-factor authentication.
- Implement microsegmentation by adding secure zones based on the zero-trust security model: partition network traffic with endpoint firewalls, virtual or software-defined networks, or physical networks.

RDP ports opened to the internet jumped from around 3 million in Jan. to around 4.5 million in March 2020.

—ZD Net

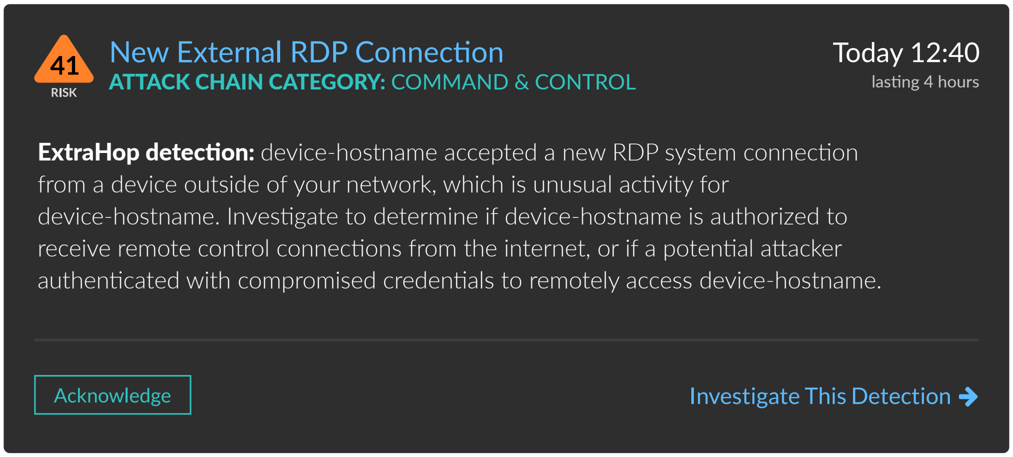

New External RDP Connection

ATTACK CHAIN CATEGORY: COMMAND & CONTROL

Today 12:40

lasting 4 hours

ExtraHop detection: device-hostname accepted a new RDP system connection from a device outside of your network, which is unusual activity for device-hostname. Investigate to determine if device-hostname is authorized to receive remote control connections from the internet, or if a potential attacker authenticated with compromised credentials to remotely access device-hostname.

New External RDP Connection (continued)

Detection Mitigation Options:

- Disable Remote Desktop Services unless required.
- Remove the default local Administrators group from the list of approved RDP groups and add specific users to the list.
- Enforce multi-factor authentication for remote logins.
- Only allow incoming external RDP connections from trusted devices.
- Limit the number of RDP login attempts, and then lock user accounts that exceed this number.
- Enforce security zones by implementing network segmentation and firewall policies to limit how devices can communicate.

> In the US, RDP brute force attacks peaked at 1.4 million attempts per day
>
> **—SC Media**

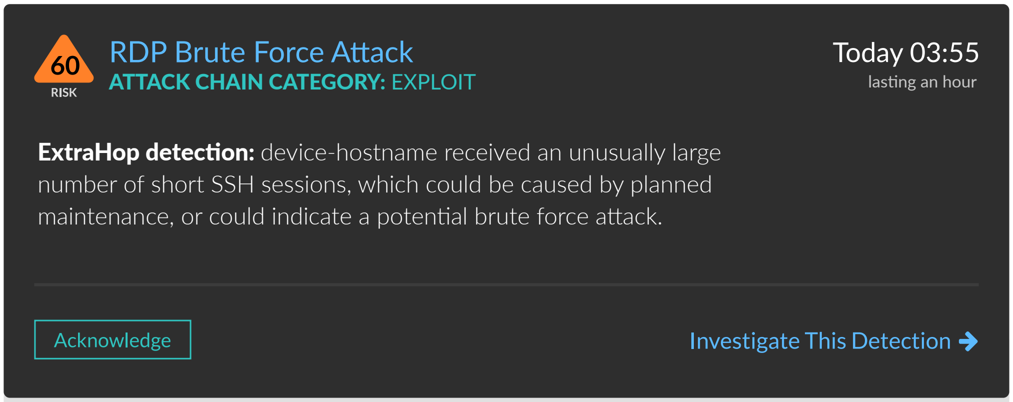

RDP Brute Force Attack

ATTACK CHAIN CATEGORY: EXPLOIT

Today 03:55

lasting an hour

ExtraHop detection: device-hostname received an unusually large number of short SSH sessions, which could be caused by planned maintenance, or could indicate a potential brute force attack.

Detection Mitigation Options:

- Limit the number of Remote Desktop Protocol (RDP) login attempts and then block users that exceed this number.
- Implement strong authentication methods for remote access services.
- Implement network segmentation and firewall policies to limit how devices can communicate and enforce security zones.
- Review access controls to ensure that only necessary users can connect to remote access services with RDP.
- Review authentication methods and enforce policies for secure credential creation and multi-factor authentication.

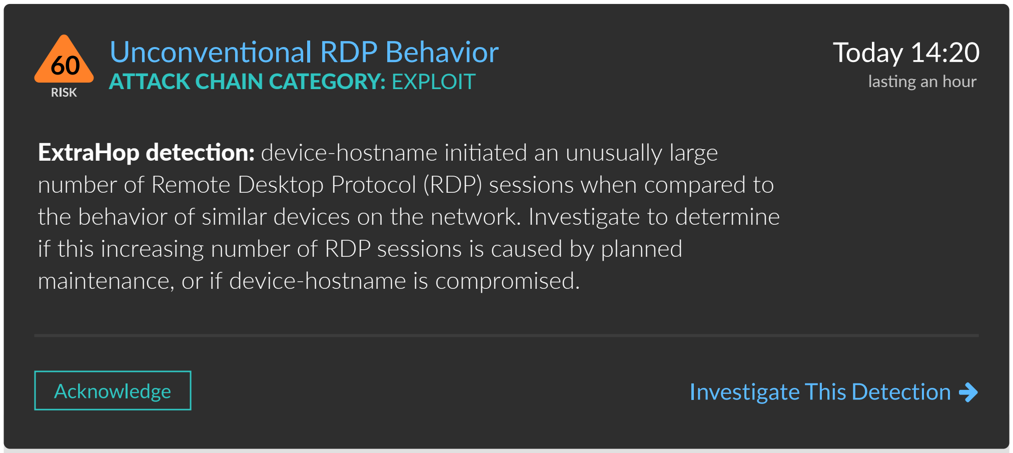

Unconventional RDP Behavior

ATTACK CHAIN CATEGORY: EXPLOIT

Today 14:20
lasting an hour

ExtraHop detection: device-hostname initiated an unusually large number of Remote Desktop Protocol (RDP) sessions when compared to the behavior of similar devices on the network. Investigate to determine if this increasing number of RDP sessions is caused by planned maintenance, or if device-hostname is compromised.

Detection Mitigation Options:

- Implement strong authentication methods for remote access services.
- Implement network segmentation and firewall policies to limit how devices can communicate and enforce security zones.
- Review access controls to ensure that only necessary users can connect to remote access services.
- Review authentication methods and enforce policies for secure credential creation and multi-factor authentication.

> The US and China have the most vulnerable ports open, around 1.3 million each.
>
> **—ZD Net**

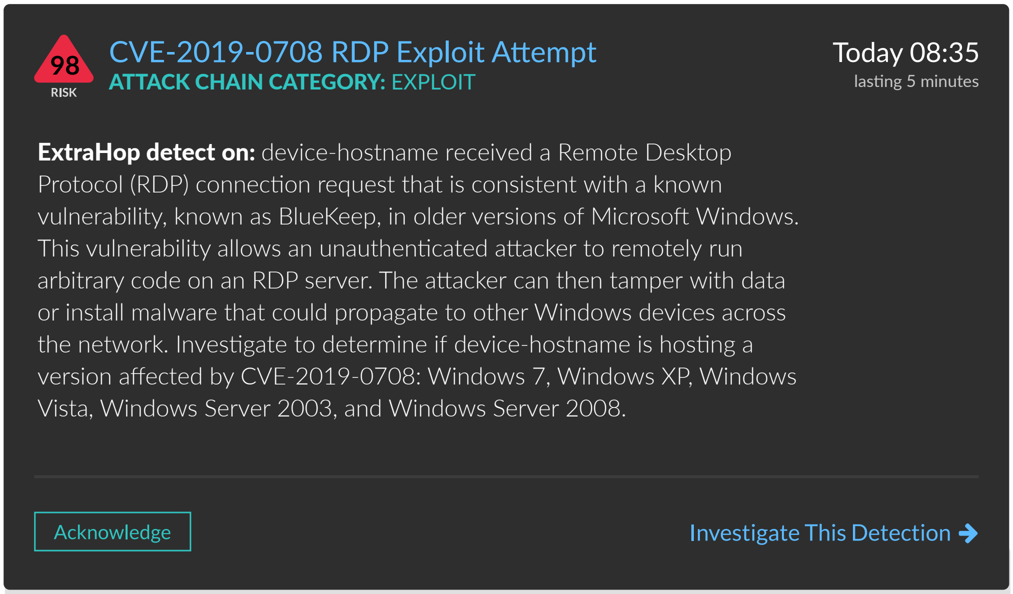

CVE-2019-0708 RDP Exploit Attempt

ATTACK CHAIN CATEGORY: EXPLOIT

Today 08:35
lasting 5 minutes

ExtraHop detect on: device-hostname received a Remote Desktop Protocol (RDP) connection request that is consistent with a known vulnerability, known as BlueKeep, in older versions of Microsoft Windows. This vulnerability allows an unauthenticated attacker to remotely run arbitrary code on an RDP server. The attacker can then tamper with data or install malware that could propagate to other Windows devices across the network. Investigate to determine if device-hostname is hosting a version affected by CVE-2019-0708: Windows 7, Windows XP, Windows Vista, Windows Server 2003, and Windows Server 2008.

Detection Mitigation Options:

- Install relevant patches for affected Windows versions. Due to the severity of this vulnerability, Microsoft has extended patches to products that are otherwise out of support.
- Disable Remote Desktop Services unless required.
- Implement Network Level Authentication (NLA) on systems running supported versions of Windows 7, Windows Server 2008, and Windows Server 2008 R2.
- Configure firewalls to block traffic on TCP port 3389.

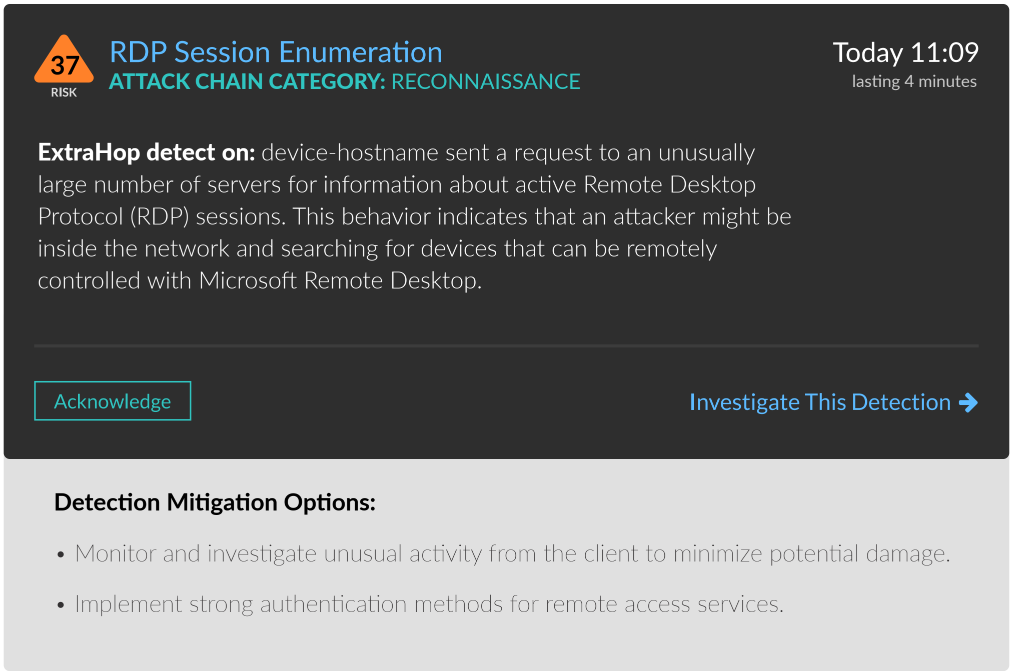

RDP Session Enumeration

ATTACK CHAIN CATEGORY: RECONNAISSANCE

Today 11:09

lasting 4 minutes

ExtraHop detect on: device-hostname sent a request to an unusually large number of servers for information about active Remote Desktop Protocol (RDP) sessions. This behavior indicates that an attacker might be inside the network and searching for devices that can be remotely controlled with Microsoft Remote Desktop.

Detection Mitigation Options:

- Monitor and investigate unusual activity from the client to minimize potential damage.
- Implement strong authentication methods for remote access services.

> RDP is the most common attack vector–used about 60% of the time.
>
> **—SC Media**

RDP BEST PRACTICES

When it comes to RDP exploitation, it's all about mitigation. The reality is that RDP will always be seen as a convenient attack vector, particularly now as more workers use VPNs and virtual desktops to access corporate resources. It's the steps you take to mitigate it that make the difference between an attempted intrusion and a major breach.

Practice good hygiene. Disable remote desktop services unless required and install relevant patches for any affected devices as fast as possible. Make sure you have a way to track in real time which devices are active and connected so that you don't inadvertently miss a new device. Review access controls to ensure that only necessary users can connect to remote access services.

Put up roadblocks. The more complex the credentials, the harder it is for cybercriminals to get into your network. Implement and regularly review authentication methods and enforce policies for secure credential creation and multi-factor authentication.

Breaches are inevitable. Plan accordingly. If we've learned anything by now, it's that no perimeter is going to keep out every threat. Assume the bad guys will get in and do everything possible to prevent lateral movement once they're inside. Implement microsegmentation by adding secure zones based on the zero trust security model and partition network traffic with endpoint firewalls, virtual or software-defined networks, or physical networks. Also closely monitor and investigate unusual activity from the client to minimize potential damage. While logs and agents can provide device-level insight, it's important to actively monitor network traffic for this activity. Logs can be erased and agents can be disabled, leaving a blind spot for attackers to exploit if you don't have network visibility.

For the time being, use of RDP will only grow as organizations around the world work to adjust to the new realities of IT. In time, increased use of cloud-based solutions may lessen the need for VPN and VDI, but RDP isn't going away anytime soon. With cybercriminals and APTs actively seeking to exploit the chaos, making sure you have both visibility and control has never been more critical.