The importance of resolving all alerts

CRITICALSTART
They're good. We're better.

CRITICAL**START MDR service is built on our experience that to eliminate risk you need to resolve all alerts. BUT WHY?**

Today's threat actors continue to incorporate new and sophisticated techniques, tactics, and procedures (TTPs) into their attacks. They apply hands-on techniques and purpose-built malware to scrape credentials, move laterally, establish persistence and exfiltrate data. All while evading detection from today's security tools.

Consider this, Cybercriminals are organized, multi-billion-dollar enterprises. Their labs are equipped with the same security tools that you are using. They reverse engineer security tools and study the TTPs used by security teams, to identify and exploit weaknesses in defense. A common tactic is to design attacks that hide in the noise of the lower priority alerts generated by security tools. Why? Because they know that security teams cannot scale to investigate and remediate the volume of lower priority alerts.

Ragnarlockbit - A real world attack

CRITICAL**START's Incident Response (IR) team was recently engaged by a mid-size bank that was victimized by the RagnarLockbit ransomware. The client paid the ransom, only to be re-infected twice and making a second payment.**

This was a hands-on attack utilizing multiple techniques. Like most attacks, it most likely started off with phishing to gain access to the network. The human element continues to be the weak link in security. No matter how much training you provide, it only takes one mistake to compromise your business.

The first evidence we observed was the use of PowerShell scripts to download malicious code from sslip.io, a legitimate content hosting website that was not blacklisted at the time. PowerShell is a common tool used by system admins for legitimate IT operations. Threat actors take advantage of its unlimited access to Windows resources including, Windows APIs, full access to Window Management Instrumentation (WMI), and the .NET framework. They often disguise malicious PowerShell scripts with encryption and obfuscation techniques to evade detection. The EDR solution reported this activity with a Low priority alert.

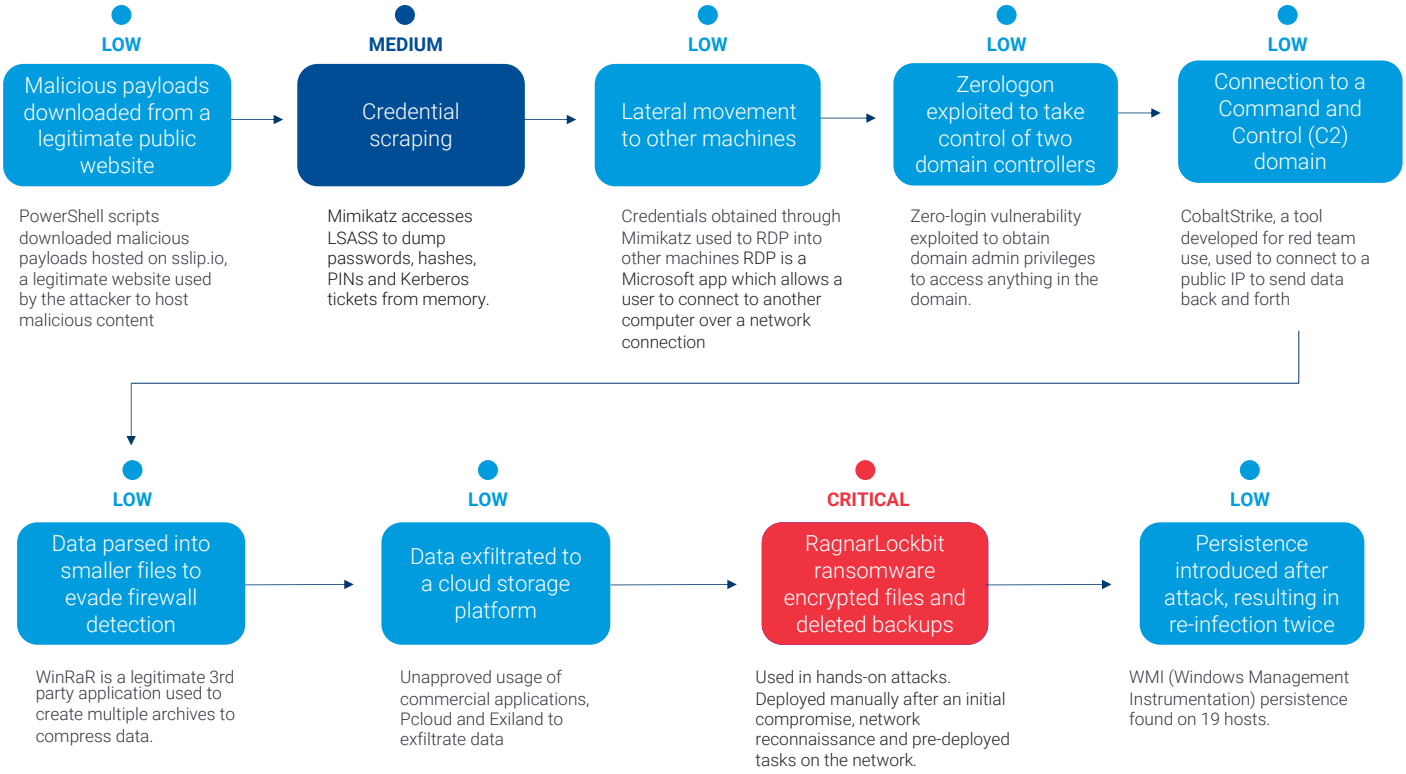

Ragnarlockbit - A real world attack

Once in, Mimikatz was used to scrape credentials and escalate privileges. Mimikatz is an open-source application that allows users to view and save authentication credentials. It was initially developed as a tool for Pen Testers and Red Teams to detect vulnerabilities in a network. Mimikatz accesses LSASS, the local Windows security authority that manage passwords and authentication for all processes on the computer. It performs credential gathering techniques such as Pass-the-Hash, Pass-the-Ticket, Kerberos Golden and Silver Ticket and more. EDR looks for abnormal uses of LSASS. Attackers apply sophisticated methods to hide this. In this case, the use of Mimikatz generated Medium priority alert.

Credentials collected by Mimikatz were used to move laterally to other machines in the network using RDP (Remote Desktop Protocol). Lacking more detailed context for legitimate versus malicious RDP usage, the EDR determined this was a Low priority alert.

The threat actors exploited a Microsoft Zerologin vulnerability (CVE-2020-1472). It allowed them to access two domain controllers, giving them full control of the domains. It gave them the ability to steal credentials from individual Windows accounts across the domains. Once attackers gain control of domain controllers, they've won.

Attackers are quick to take advantage of Microsoft and other application vulnerabilities. They have a limited window to launch new attacks before patches can be released and installed. Even then, they will look for networks that have not been upgraded. The EDR solution was not configured to look for this vulnerability and issued another Low priority alert.

At this point, the CRITICAL**START** IR team identified a C2 connection. The threat actors used a Red Team tool, CobaltStrike, to connect to a public IP address to send data back and forth. We see further evidence of their intentions with WinRaR, a third-party application used to parse data into smaller files to evade firewall detection. The attackers followed this by applying Pcloud and Exiland, two commercial applications, to exfiltrate device information, credentials, and other data for use in future attacks. All three of these actions created Low priority alerts.

It is not until RagnarLockbit executes that we see a Critical alert. By this time, it is too late. After launching the ransomware, the threat actors used WMI to establish persistence on 19 hosts. Persistence mechanisms can include registry changes and installing malware in a start-up folder. Their goal was to leave a kill chain in place to launch future attacks. And in fact, they re-infected the client two more times and extracted a second ransomware payment before engaging CRITICAL**START**.

What does Resolve All Alerts mean?

Many traditional MDRs claim they resolve all alerts. But what does that mean?

For CRITICAL**START**, it begins with configuring security tools to collect all possible alerts, across all priorities – Critical, High, Medium, and Low. This will generate a tremendous volume of alerts requiring investigation and resolution. Over a 6-month period, the CRITICAL**START** SOC recorded an average of 14,400 alerts per client per day. When we look at the alert traffic of our busiest clients, we see over 359,000 alerts per client per day. We have recorded an astounding **890,300 alerts for a single client in a single day.**

Key Takeaways

- ✓ Threat actors reverse engineer security tools and study security tactics, techniques, and procedures to identify and exploit defensive weaknesses.
- ✓ They combine fileless and living-off-the-land techniques with purpose-built malware to gain access, escalate privileges, move laterally, establish persistence and exfiltrate data.
- ✓ **They are relentless.** If you don't detect and remediate the entire kill chain, they will continue to re-infect and threaten your business.

Alert Volume

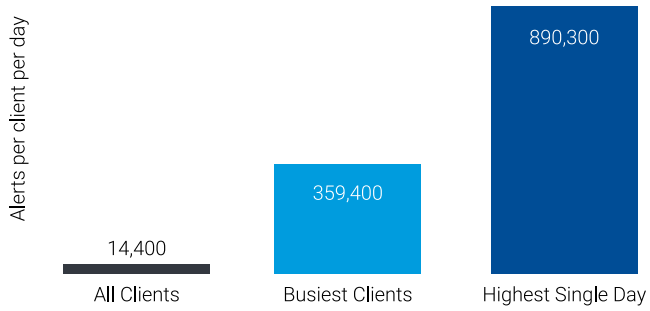

*Recorded by CRITICAL**START** SOC over 6-months*

Assuming an average of ten minutes to investigate and resolve each of these alerts, you are looking at over 2,400 hours.

Resolution Time (Hrs)

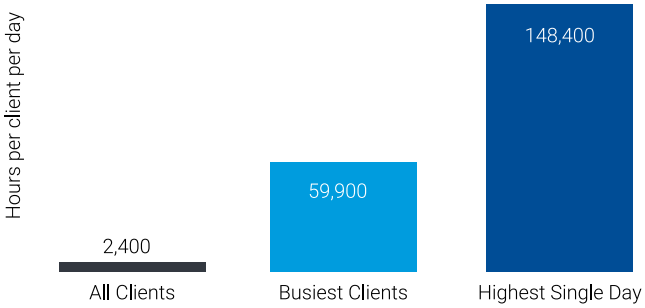

Based on 10 minutes investigation/resolution time per alert

This would require 300 SOC analysts - for a single client.

What happens when another client suffers through an alert storm?

How will this impact risk exposure to your organization?

And how does an MDR provider support this alert volume across all their clients?

The answer is they don't. Traditional MDRs control alert volumes through alert suppression.

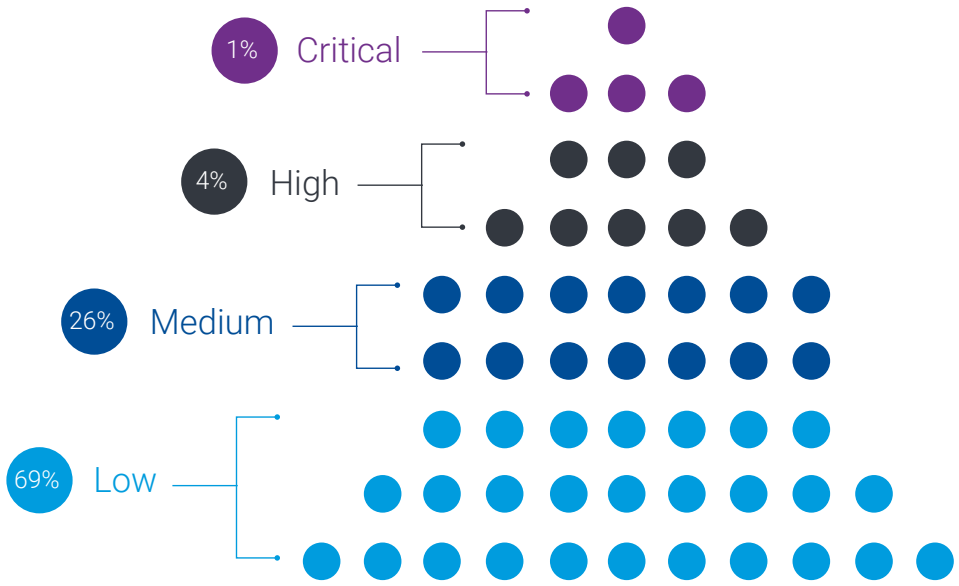

Based on 6-month record of alerts escalated by priority

They disable lower priority alerts, change thresholds and prioritize Critical and High priority alerts while ignoring the Mediums and Lows. Based on a 6-month record of unknown alerts escalated by the CRITICAL**START** SOC, this means 95% of all alerts are escaping investigation. As we stated above, the threat actors know this. The RagnarLockbit proves they are taking advantage of this approach.

Accepting risk on your behalf

Practicing alert suppression come at a price – **Risk Acceptance.** The cost of a ransomware attack comes from business disruption which leads to lost revenues, lost productivity and potential SLA fines from partners and customers. Data breach costs from legal and restitution, brand damage and compliance can add up quickly. According to IBM's Cost of a Data Breach Report (2020) the average cost of a data breach is $3.86 million. In many cases, the cost can run up to the tens of millions.

How can you resolve all alerts?

Is it possible to resolve all alerts? **Yes.** CRITICAL**START**'s MDR service is based on this. We combine the industries most advanced analytics and automation platform with elite security expertise to resolve every alert, stop breaches and reduce risk acceptance.

CRITICAL**START** – We don't accept risk

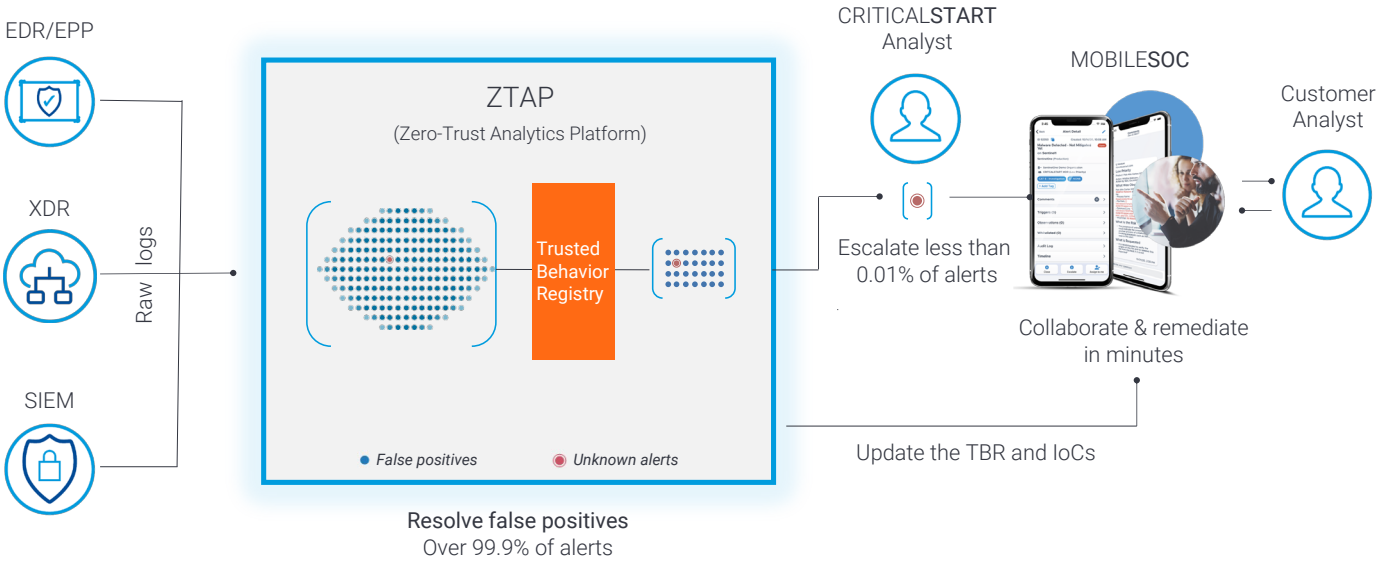

Traditional MDRs view lower priority alerts as "noise". CRITICAL**START** recognized that the "noise" is from false positives. The Zero-Trust Analytics Platform (ZTAP), featuring the Trusted Behavior Registry (TBR) is designed to auto-resolve false positives so our SOC and your security team can focus on the alerts that truly require attention.

The CRITICAL**START** implementation team works with your team to enable all possible alerts from your security tools. ZTAP ingests the raw logs. The TBR is the only purpose-built database of known good behaviors. Incoming alerts are analyzed by the TBR. If an alert matches the conditions defined in the TBR, it is auto resolved as a false positive. Over the last six months, ZTAP and the TBR auto resolved 99.94% of all ingested alerts.

Alert Volume

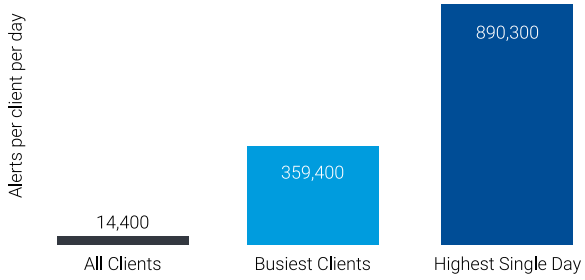

Alert Reduction

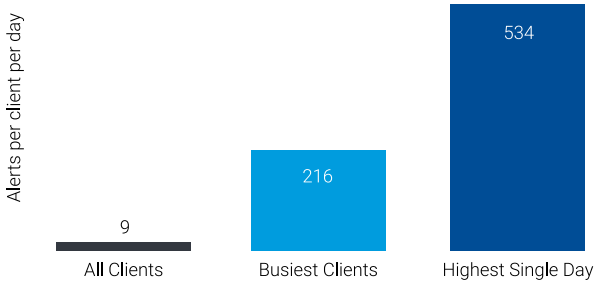

*Latest CRITICAL**START** results: Auto-resolving 99.94% of alerts*

This reduced the alerts requiring CRITICAL**START** SOC investigation from 14,400 alerts per client per day down to 9 alerts per client per day. This provides the means for CRITICAL**START** to resolve all alerts at scale. Other MDRs cannot.

Alerts not auto resolved are classified as 'Unknown". They are investigated by the CRITICAL**START** SOC. Over a six-month period, our SOC had to escalate only 0.006% of all alerts to our clients. This delivers unmatched scalability to our client security teams.

On average, we escalate one alert per client per day, requiring only ten minutes of investigation and resolution time for your team. This frees time to focus on other high priority and more proactive security projects.

Alert Escalation

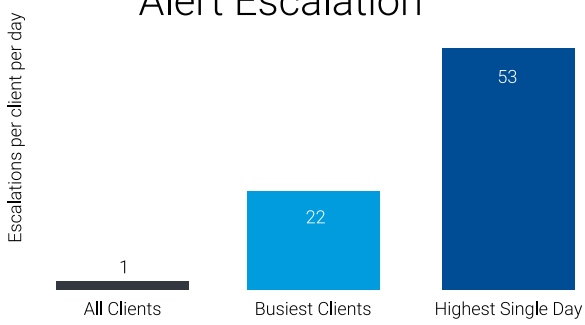

*Latest CRITICAL**START** results: Escalate 0.006% of alerts*

Client Resolution Time (Hrs.)

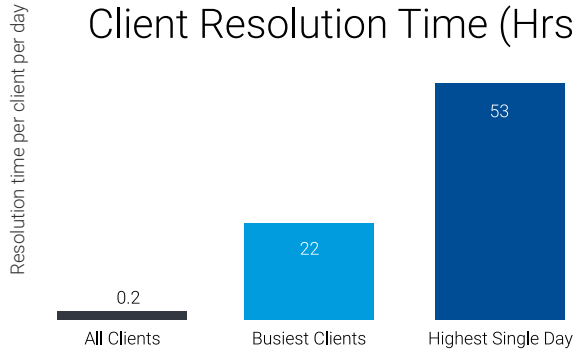

Assumes 10 minutes of investigation time per alert

Key Takeaways

- ✓ Resolving all alerts starts off with configuring security tools to collect all alerts. This creates an alert volume that overwhelm the capacity of most MDRs.
- ✓ Traditional MDRs solve their scalability problem by suppressing alerts, requiring their clients to accept risk, often without their knowledge.
- ✓ CRITICAL**START,** powered by ZTAP and the TBR is the only MDR service that scales to resolve every alert, stop breaches, and reduce risk acceptance.

We back our claims with the most aggressive SLAs in the industry. One hour Mean-Time-to-Detect and Median-Time-to-Resolve for all alerts and all priorities.

Contact CRITICAL**START** to run our Risk Acceptance Calculator to quantify the risk you're accepting today and how CRITICAL**START** MDR with ZTAP and our elite security analysts can reduce your risk acceptance to near zero – Cost effectively

Visit **www.criticalstart.com** for more information.