DARKTRACE

Cloud Detection and Response:

The Ultimate Guide to Automating Incident Response

Contents

Introduction

Automating the collection of incident evidence following detection helps ensure cybersecurity events are appropriately handled before they are at risk of escalating. The lack of automation coupled with alert fatigue often means things are missed and what may seem like a low-severity detection may actually be connected to something far more malicious. Leveraging Cloud Detection and Response (CDR) to remove many of the complexities and manual steps in kicking off a more thorough investigation means security teams can dive deep more often and better protect their environments.

When an alert is generated from a solution such as CDR, Endpoint Detection and Response (EDR), or eXtended Detection and Response (XDR), evidence can be automatically collected and processed across systems of interest and then run against a deeper set of threat intelligence to validate severity level.

Often, a first pass of collection is done by grabbing a subset of triage artifacts and analyzing those events before a more in-depth full disk acquisition may be needed. Both the triage and full disk collections should be automated such that the data is available even if the ephemeral system is terminated. This also ensures the data is accessible as soon as an analyst needs to perform a deeper investigation.

Analyzing the collected data completes the response fully, allowing organizations to apply lessons learned and improve overall security posture.

Automating triage collection

Triage artifact collection is an important technique used during IR. It involves collecting a subset of artifacts like system files, running processes, network connections, registry hives, volatile memory, and event logs before performing more time-consuming activities like full disk analysis and digital forensics. This speeds up Mean Time to Respond (MTTR) by triaging the most important artifacts on a set of systems first.

It's no surprise that increasing the speed at which incident responders perform their triage and analysis can drastically reduce the risk to an organization during an incident.

Traditional methods of performing IR rely on a collection of tools, scripts, and manual analysis, often performed sequentially: grab a full disk and memory capture from system A, process and analyze evidence from system A, then rinse and repeat as you discover other systems of interest.

Though this works, it is very time consuming, and security and IR teams need a better way to collect key artifacts and quickly determine whether to dive deeper into a system or move on to others.

With advancements in the IR space, organizations can start to automate these processes today. Managed Security Service Providers (MSSPs) and many enterprise organizations have already adopted Security Orchestration, Automation, and Response (SOAR) platforms to help with this. But regardless of whether you choose to perform evidence collection and processing in a semi-manually or fully-automated way, the guidelines and best practices recommended here are applicable in all instances.

Automating full disk collection

Similar to triage acquisitions, full disk collections have historically been a manual process involving bootable USB sticks used to image volumes or physically shipping a device to a secure location for forensic analysis.While these methods are effective, they are even more time consuming than trying to perform a triage collection.

With the adoption of cloud, capturing images of full volumes, running in the cloud, becomes a lot easier via snapshotting and cloud provider APIs. But this still requires an in-depth understanding of each cloud providers' APIs and the skillset to write the scripts to call these APIs. And even after you capture the disk images, you still have the challenge of getting that image into an environment where you can process and analyze the data before the ephemeral server is nuked. You have your cloud evidence in your cloud storage, why bring it back on-premises if you don't need to?

Fast forward to today and you can leverage the cloud-native APIs to automate this process. Or you can use a cloud-native investigation platform which:

- Abstracts away the complexity of the cloud
- Does not impact production workloads because no agent is needed
- Fully automates acquisition, processing, and analysis of cloud volumes and more.

Best practices

Many of the guidelines for evidence collection, published by SANS [1] [2] and NIST [3] [4], over the past 15 years are still applicable today.

Though the principles still hold true, the technology and tools have evolved significantly. A summary of evidence collection best practices is listed below. We include examples later to help show how you can start implementing these practices in your organization:

Best practice #1

Identify data sources and be prescriptive in what you collect

Reduce acquisition resources and processing time even further by collecting the right set of artifacts. The goal being to perform an initial triage, helping to identify additional data sources and ruling out others. In line with SANS and NIST guidance, best practice for live data triage collection should be based on:

- Artifacts' likely value to the investigation
- Volatility of the data
- Amount of effort required to acquire that data

Triage artifact collection should include artifacts such as network connection state, logged on users, currently executing processes, event logs, $MFT, registry hives, and volatile memory. After initial analysis of triage evidence, if a full disk image is deemed necessary, this can be acquired, processed, and analyzed automatically.

Best practice #2

Collect and process data efficiently

Document and standardize the collection and processing of evidence data. Where possible, collect and process evidence from systems of interest in parallel. The faster you can analyze key events, the faster you can work to resolve an incident. This, in turn, lowers the overall risk to your organization when an incident occurs.

Best practice #3

Standardize the preservation of data

The lifecycle management of data is typically based on its value and volatility. Organizations should define and document where data will be stored, who will have access to it, and for how long it should be retained. When possible, ensure that the lifecycle management of evidence takes into account hot and cold storage requirements and full chain of custody, including proper tagging and labeling of evidence. This is helpful for supporting audits and legal regulations.

Best practice #4

Analyze data in a holistic manner

Having a holistic view, across all pieces of evidence during an investigation, increases the speed at which security teams can move to containment, eradication, and recovery. Consider how you will collect and aggregate data at scale allowing not only the collection of evidence, but also the ability to view and drill down into that data in either a timeline (or other user friendly) view, across all systems.

Best practice #5

Refine and sharpen your toolset

Stay up to date on the latest trends and technological advancements in the industry. For example, the rapid adoption of cloud has forced many organizations to either shoehorn their current IR processes for use in cloud investigations or they simply acknowledge and accept the risk that they have a lack of visibility and response capabilities in their cloud environments. But the cloud can be an asset to your security team. Taking it a step further, consider using cloud resources to collect, process, and store evidence in a secure, flexible, and efficient way.

Practical example

Imagine you have an environment with 50,000 hosts and you are informed of a new zero-day vulnerability that needs to be investigated. By following an effective IR plan, you should be able to narrow the scope of your investigation significantly by performing the following, as an example:

- Leverage your threat detection and response solution, SIEM, and other detection tools to hunt and narrow the scope of the systems you need to investigate from the tens of thousands to a few hundred. At this stage you typically have a pretty well-defined set of Indicators of Compromise (IoCs) that you are looking for across your environment. Defining those IoCs would trigger alerts for devices that match the indicators.
- Next, prioritize the smaller set of systems and use your triage collection capabilities to further reduce the scope of the investigation. Collect, process, and analyze the data from the initial systems of interest so you can narrow the focus further, to key systems of interest. This collection of triage artifacts can be automated based on alerts from a threat detection and response tool or SIEM.
- A more in-depth full disk analysis is used to confirm earlier findings, and if possible, provide further context of attacker activity. At this point, we have a high level of confidence that ~5 systems were involved in the incident. This smaller number of systems would be key attacker systems where they have spent most of their time or are channels for C2.

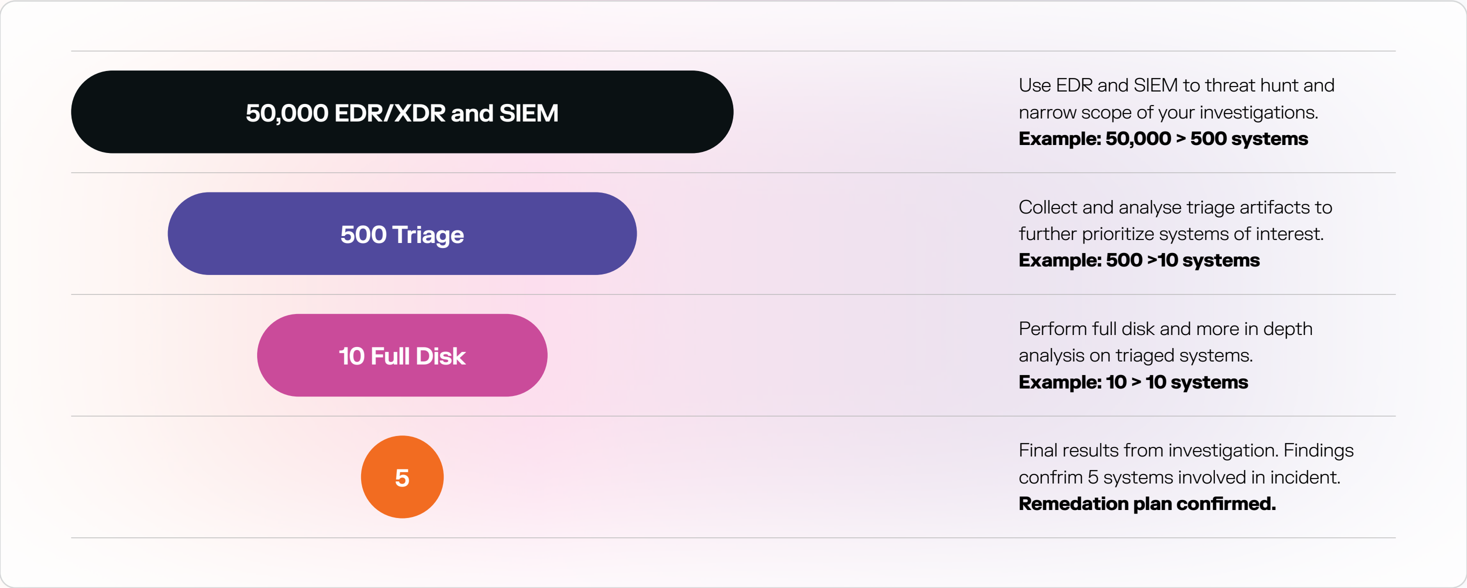

50,000 EDR/XDR and SIEM

Use EDR and SIEM to threat hunt and narrow scope of your investigations.
Example: 50,000 > 500 systems

500 Triage

Collect and analyse triage artifacts to further prioritize systems of interest.
Example: 500 >10 systems

10 Full Disk

Perform full disk and more in depth analysis on triaged systems.
Example: 10 > 10 systems

5

Final results from investigation. Findings confrim 5 systems involved in incident.
Remedation plan confirmed.

Putting best practices to use

So, what can you do to put these guidelines into practice in your environment?

To start, standardize on a base set of artifacts to collect during an initial triage. Follow industry guidance, including network connection state, logged-on users, running processes, event logs, $MFT, volatile memory, and registry hives. Documenting and standardizing these data sources enable more streamlined and automated collection.

Whichever tool you use, ensure it offers flexibility—customizable acquisition options and broad OS support are key.

Next, evaluate how data collection processes can be automated. If you're leveraging an EDR, XDR, or CDR solution, most platforms now support remote access via a live response shell or equivalent interface. These remote command capabilities are typically accessible via API, e nabling automated collection of triage artifacts during an incident—across endpoints and cloud workloads alike. Some solutions also support API-based acquisition of cloud storage volumes or containers.

To maximize these capabilities, many teams rely on a SOAR platform to orchestrate automated actions. These platforms ingest signals from across the environment—network, endpoints, cloud logs, and CDR platforms—and use predefined playbooks to drive response.

A SOAR platform can correlate the event and initiate automated triage. By automatically capturing this data at the time of alert, security teams are better equipped to respond quickly and completely.

Critically, CDR plays a central role in this process for cloud-native workloads. For cloud-first or hybrid environments, CDR solutions monitor runtime activity and use behavioral analysis to detect anomalies across cloud infrastructure, identities, and services. If EDR is not present, alerts from CDR or native cloud services (like AWS GuardDuty) can serve as the trigger for triage or full-disk acquisition—closing the response gap in cloud environments.

If deeper analysis is warranted after initial triage, a SOAR platform can escalate to full disk capture of affected cloud or on-prem systems, enabling complete forensic investigation.

To help demonstrate how to proactively detect, investigate, and respond to incidents across hybrid environments, we'll walk through an example using a SOAR tool, a cloud-native forensic investigation platform, and CDR or EDR running across your workloads. While EDR is useful, it is not required—CDR solutions can provide equivalent detection and API-driven response capabilities in cloud environments without agents, enabling comprehensive incident handling even where traditional endpoint tooling falls short.

Step #1

An alert is triggered by your threat detection tool, informing you of malicious activity. This may be multiple events or even multiple systems (either cloud or on-premises) which your tools flag as malicious. The important point is that we have a high-fidelity alert which means there is a high likelihood of needing a triage collection performed. The detection alerts are sent to the SOAR, in near real time to be actioned upon by predefined playbooks.

Step #2

The SOAR platform correlates the various alerts, and a playbook is automatically triggered.

An API call is made from the SOAR platform to a cloud-native investigation platform to request a triage acquisition command.

JSON is returned, which includes a command to run on the endpoint. The command also includes a pre-signed URL which allows the triage package to be automatically uploaded to an AWS S3 bucket or Azure storage container for automatic processing after collection.

Step #3a

SOAR calls the detection and response tool via API to execute the command on the host in question. This is typically done via an EDR, XDR, or CDR API call which will connect to the host and run the command or script returned in Step 2. Alternatively, if leveraging a cloud agent like SSM Agent for AWS or a Mobile Device Management (MDM) agent for managing software across your environment, these tools can be used as well.

Step #3b

Triage package is generated, uploaded to cloud storage, and automatically processed into the investigation platform. Using cloud workers which allow for rapid parallel processing, the ideal automated cloud IR tool will monitor a pre-defined AWS S3 bucket or Azure storage then import the triage packages automatically.

Then threat intel is run over the extracted artifacts to allow analysts to view, search, and collaborate with teammates to perform an investigation.

As an example, if collecting triage packages from ten separate hosts, all ten triage packages can be processed in parallel which means what used to take days to complete, can now be completed in minutes. Any alarms or suspicious indicators that are identified by the investigation platform are fed back into the SOAR platform to improve the playbooks.

After the triage packages are loaded successfully into the platform, analysts can be notified via email, Slack, etc. so they know when the evidence is available for investigation. Analysts should have a single pane of glass to search and filter on all timeline events across all evidence items, making it easier to correlate key events across large sets of data.

In addition to processing data efficiently and making that data available for analysis in a single timeline view, organizations should be able to easily control lifecycle management of collected data and hot and cold storage requirements in both AWS and Azure. AWS supports this through Object Locks and Lifecycle Rules associated with an S3 bucket. And Azure supports this through Life Cycle Management Policies for blob storage.

Steps #4 & #5

If further investigation is needed, a deeper set of artifacts can be collected from the host; for example, via a full disk acquisition. If this is the case, the SOAR platform would call upon the automated cloud IR tool to acquire the full disk images which should capture, process, and analyze the volumes automatically.

Once a thorough investigation is performed, remediation can commence.

Summary

Triage collection allows for the acquisition of key artifacts on a system to help identify the scope of an incident quickly and help focus time and effort in areas of an investigation that will yield the most efficient results. After triage collection, a deeper analysis may be needed, in which case you perform full disk analysis. And though the best practices for handling data collection have stayed relatively the same over the past decade, customers' environments and the tools and technologies have changed drastically.

Having a well-defined and practiced plan for performing your investigations is crucial. As part of that plan, triage acquisitions help significantly narrow down the systems of interest early in the process, reducing the number of systems on which you will need to perform deep-dive analysis.

Full disk acquisitions allow for deeper analysis of cloud resources without requiring any agents and without disrupting production workloads. Automating these steps adds consistency to the process and drastically reduces the mean time to resolution.

In summary, if you follow the best practices during collection, you can standardize on your artifact collection process, reduce the time to resolve incidents through automation, analyze data in a more holistic manner, and stay up to date with the latest technological advancements in the industry.

How Darktrace can help

Darktrace delivers a proactive approach to cyber resilience in a single cybersecurity platform, including cloud coverage. Darktrace / CLOUD is a real time CDR solution built with advanced AI to make cloud security accessible to all security teams and SOCs. By using multiple machine learning techniques, Darktrace brings unprecedented visibility, threat detection, investigation, and incident response to hybrid and multi-cloud environments.

Darktrace's cloud offerings have been bolstered with the acquisition of Cado Security Ltd., which enables security teams to gain immediate access to forensic-level data in multi-cloud, container, serverless, SaaS, and on-premises environments.

Darktrace is a valuable tool for automating IR processes.

Learn more about Darktrace / CLOUD for:

AWS → **Azure →**

Dive into how Darktrace / CLOUD works:

Read the solution brief →

See Darktrace in action with a personalized meeting:

Request a demo →

■ About Darktrace

Darktrace is a global leader in AI for cybersecurity that keeps organizations ahead of the changing threat landscape every day. Founded in 2013 in Cambridge, UK, Darktrace provides the essential cybersecurity platform to protect organizations from unknown threats using AI that learns from each business in real-time. Darktrace's platform and services are supported by 2,400+ employees who protect nearly 10,000 customers globally. To learn more, visit http://www.darktrace.com.

North America: +1 (415) 229 9100 | Europe: +44 (0) 1223 394 100 | Asia-Pacific: +65 6804 5010 | Latin America: +55 11 4949 7696

darktrace.com | info@darktrace.com