Components and Architecture of Mandiant Automated Defense

Overview

Mandiant Automated Defense, a critical component of extended detection and response (XDR) and the Mandiant Advantage platform, features decision automation software pre-built with the reasoning and decision-making skills needed to tackle the complexity and high volume of data facing security teams today. Automated Defense automates the analysis and triage of security data at machine speed with depth and consistency. Its proprietary intelligent decision engine provides built-in reasoning and judgment to make better decisions faster.

Automated Defense evaluates the event data stream in real-time from an organization's existing security detection sensors and learns about its security infrastructure and network context. The solution can analyze all ingested events and alerts, regardless of volume; to build evidence and context around malicious activity. Automated Defense processes every event, not just alerts labeled "important" or "critical," and performs extensive checks against an internal repository of context to appropriately escalate incidents.

Automated Defense uses probability-based reasoning and provides 24x7 continuous monitoring removing the need to filter, tune-down or ignore security alerts resulting in a significantly reduced number of false positives. Automated Defense eliminates human bias or fatigue of monitoring security alerts and maximizes the effectiveness of security teams by enabling analysts to go threat hunting and other security related activities.

Designed to easily integrate into any security infrastructure, Automated Defense brings additional value to existing investments by providing the capacity to thoroughly analyze all security events that are detected—without any learning mode or security rules to maintain.

Using the latest advancements in artificial intelligence, machine learning, modern streaming architectures and unique Integrated Reasoning,™ Automated Defense acts autonomously—without a heavy system management burden, security engineering or long learning cycles.

Functional Application

Automated Defense was built to support the exponential growth in security data and the variety of sensors and infrastructures deployed in today's security environments.

- Automated Defense is a cloud based application that includes:

- Data Processors that enrich, normalize and clean data across multiple types of security sensors, threat intelligence and company context.

- Integrated Reasoning that analyzes all network, endpoint, web filtering, cloud events and alerts to equate malicious attacks and determine which incidents need investigation.

Automated Defense extracts the data it needs from the existing security event collection infrastructure to conduct deeper analysis with Integrated Reasoning.

How it works

Functional architecture
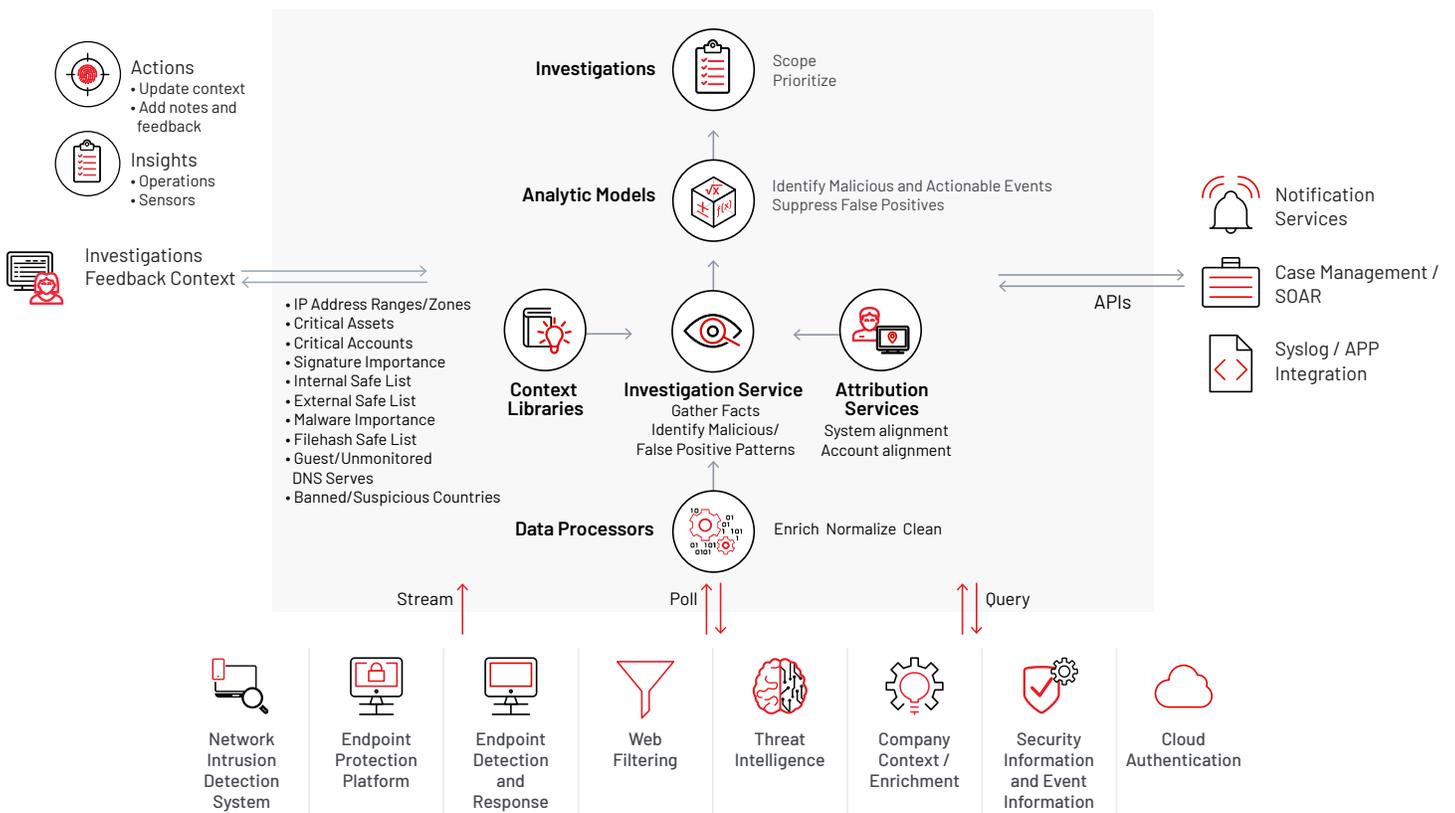
Event Handling

Automated Defense evaluates the event data stream from existing security sensors and learns about company context, the IT network and cloud environment. Automated Defense analyzes all events and alerts, regardless of volume, building evidence and context around malicious activity.

Automated Defense comes pre-structured with expert judgement, but learns and adapts while maintaining tribal knowledge. Automated Defense runs 24x7 performing without fatigue, loss of attention or staff attrition. This mix of expert judgment and self-adaption enables Automated Defense to immediately produce high-fidelity results and improve quickly as it works with a security response team.

Events and context

During onboarding, the Automated Defense administrator is asked to provide important context about the IT environment through the management dashboard including:

- The company's publicly owned IP space

- Critical assets and accounts

- Security and network infrastructure (vulnerability scanners and load balancers)

- Network IDS/IPS signatures (high or low importance)

- Dynamic Host Configuration Protocol (DHCP) leases

- Asset vulnerability information

Not all context is required for Automated Defense to be operational; however, each additional contextual element incrementally increases the certainty about whether the detected activity is malicious and actionable or benign.

During the initial setup, the administrator configures event sources that Automated Defense will use.
To help organizations use their existing event-processing infrastructures, Automated Defense supports event sources including enterprise analytics solutions, data storage, data lakes, cloud providers, SIEM forwarders and connectors from

popular SIEM and SOAR vendors. Automated Defense will also accept events streamed directly from the management consoles of network IPS/IDS, endpoint protection platforms, web proxy and filtering solutions and endpoint detection and response systems. Unlisted event sources are also possible since Automated Defense listens for events on TCP-6060, UDP- 514, TLS-6514, and HTTP-6080.
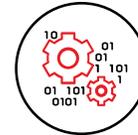
Data Processors   Enrich Normalize Clean

Company context/enrichment allows Automated Defense to adapt to each customer's unique environment.

Event processing

As Automated Defense receives events, it checks IP addresses and hostnames against sensitive contextual references such as the critical asset list, the critical account list, file name checks, geolocation information and vulnerability data.

The Automated Defense dashboard exposes the number of events reviewed, false positives that were suppressed, internal assets escalated and new investigations escalated. The dashboard includes the sensor event volume monitored, false positives suppressed, the automated actions and escalation, and lists the new and re-scoped incident investigations by priority.

Integrated Reasoning

Integrated Reasoning has three features that codify the foundational knowledge, complex decision-making process and ongoing learning of a highly skilled security analyst.

Integrated Reasoning

Investigation Service
- System Identification
- Gathers Facts
- Classifies Assets and Accounts

Attribution Services
- Feedback TTP Importance
- Attack Patterns System and User Importance
- Tribal Knowledge

Investigations
- Makes Decisions
- Scopes Incidents
- Determines Priorities
- Notifies Incident Response
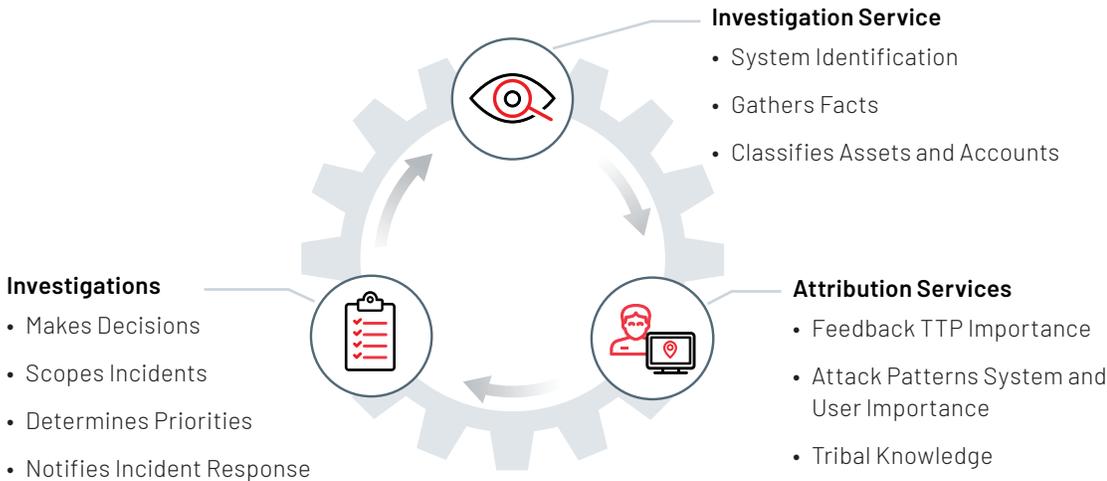
What is Integrated Reasoning?

At the core of Automated Defense is its Integrated Reasoning capability, developed by security experts and data scientists to analyze all ingested network, endpoint, web filtering, cloud events and alerts to equate malicious attacks and determine which incidents need investigation. Integrated Reasoning uses the most critical variables a security analyst considers relevant and decides if an event is malicious and actionable.

Integrated Reasoning is a patent-pending, multi-layered technology developed at the unique intersection of applied mathematics, security expertise, knowledge engineering and proprietary algorithms. With machine-level scalability, Integrated Reasoning uses all four of these elements to monitor, analyze, and determine which events are malicious across the organization's entire infrastructure. Through continuous learning and adaptation to an organization's environment, Integrated Reasoning becomes more efficient at prioritizing events and making actionable decisions. It is purpose-built to emulate the decision-making process of an experienced security analyst, considering 40-70 relevant factors per decision model. Integrated Reasoning is foundational to Automated Defense decision models, delivering efficient and effective security.

For more information on Integrated Reasoning, download the Integrated Reasoning solution brief.

Investigation Service

The Automated Defense Investigation Service uses events from DHCP and endpoint protection to constantly keep an up to date record of an IP address' associated hostname. An accurate mapping between hostname and an IP address ensures that subsequent investigations appropriately map context and behaviors to the correct system.

The Investigation Service evaluates if the systems, accounts, external IPs and domains, signatures and hashes (or other event attributes) in the security event result in an affirmative reference check that is maintained within the Knowledge Base.

Additionally, system attributes (such as open ports, operating systems) are used to classify the type and function of the internal system involved in the event.

System types inferred through the asset classification service include identifying if the internal system is a workstation or a server, or if the server is a domain controller, DNS server, web server, database server, or file server. For example, account attributes, such as administrative access, are identified through integration with Microsoft Active Directory and Azure Active Directory.

This feature gathers the information required for Automated Defense to answer a series of analytical questions for every event.

Attribution Services

Events from Automated Defense are further annotated with checks made against the Knowledge Base, a repository of both local "tribal knowledge" about a customer's unique environment as well as global threat intelligence.

Within the Network Intrusion Analytic Model, Automated Defense maintains a history of communications between sources and destinations (both internal and external to the company) in order to identify patterns and anomalies which indicate either suspicious or benign behavior.

Within the Malware Event Analytic Model, Automated Defense keeps a record of attributes shared across systems within your organization and uses this knowledge base to look for patterns indicating malware may be spreading or isolated within the environment.

Within the Web Filter Analytic Model, Automated Defense maintains a history of web requests per system to identify if traffic patterns are suspicious and possibly an indication of command and control.

Within the Endpoint Detection and Response Analytic Model, Automated Defense classifies each process and evaluates the process's relationship with the parent and child processes for suspicious behaviors.

Within the Cloud Monitoring Analytic Model, Automated Defense profiles authentication behavior and incorporates threat intelligence to identify suspicious logins that may indicate account compromise.

Additionally, Automated Defense keeps track of repeat offending systems and accounts, corroborating suspicion garnered from the Network Intrusion Model, Malware Event Analysis Model, Endpoint Detection and Respond Model, Web Filter Model, and the Cloud Monitoring Model.

Global threat intelligence sourced and used by Mandiant includes known bad indicators such as external IP addresses and file hashes, IP geolocation information, IP anonymization services such as public VPNs or TOR nodes lists. Integration with STIX/TAXII enables customers to take advantage of additional third-party intelligence sources.

Knowledge Base, a repository of both local "tribal knowledge" about a customer's unique environment as well as global threat intelligence.

Mandiant Threat Intelligence

- Breach intelligence collected from Mandiant Consulting incident response engagement

- Adversarial intelligence from Mandiant threat researchers

- Machine intelligence from third-party security products

- Operational intelligence derived from Mandiant Managed Defense services

Investigations

The Investigations component of Integrated Reasoning uses decision automation to structure the reasoning and judgement of expert security analysts and makes decisions to:

- Escalate the case, with the recommendation that incident response should perform containment and remediation actions

- Ignore the case, as it is not a threat and needs no further action at this time

Each escalated case is prioritized based on the likelihood of the activity being malicious and actionable, current most progressed attack stage, number of internal systems involved, asset criticality, and the confidence level derived for the IC-Score.

If the escalated case is related to an ongoing and open incident (same system(s), attack techniques and so on), the case is added to the existing incident investigation and the investigation is scoped and prioritized using the new information.

IC-Score

Mandiant's defined confidence scoring for publicly known indicators, that combines expert knowledge with cutting-edge machine learning.

The IC-Score helps users reduce alert fatigue, prioritize resources, understand threat actor attribution and adequately align investigation resources.

12 OPEN INVESTIGATIONS - TOP 3 HIGHEST SEVERITY

PRIORITY 1   12, Int - Int Exploitation
PRIORITY 2   8, Ransomware Infection with Suspicious Repeated Connections
PRIORITY 3   14, Virus Infections, Suspicious Repeated Connections and Int -...

RECENTLY UPDATED OPEN INVESTIGATIONS

PRIORITY 3   14, Virus Infections, Suspicious Repeated Connections and Int -...
PRIORITY 4   13, RAT Infection with Int - Ext Network C2 Malware
PRIORITY 1   12, Int - Int Exploitation

FIGURE 1. Mandiant Automated Defense adds new events to existing investigations.

Integrated Reasoning

Automated Defense Analytic Models are currently available for data from network intrusion sensors, endpoint protection platforms, endpoint detection and response, web and URL filtering devices, and cloud authorization and alerting technologies. Analytic Models are streaming technologies with the ability to evaluate, scope and prioritize events in near real-time. Analytic Models do not filter or ignore any events; they use all collected data to create a picture of the incident timeline when and if they occur.

Each Analytic Model is delivered pre-built with Knowledge Base, Investigation Service and Model content.

The Network Intrusion Analytic Model analyzes network intrusion detection system (NIDS) data, providing automated decisions on incidents that are malicious and actionable and visibility across a broad range of attacks, including damaging inbound and lateral exploitation, command and control communications, internal reconnaissance and malware that is spreading across the network.

The Malware Event Analytic Model provides automated decisions on incidents based on whether malware is spreading, what the value of the system is and how dangerous the malware is; this enables rapid, efficient and effective incident response.

The Web Filter Analytic Model autonomously monitors and analyzes all outbound web requests reported by web filtering technologies for malicious attacks, including the discovery of targeted campaigns against the network, identification of client-side exploitation, command and control traffic, and data exfiltration.

The Endpoint Detection and Response Analytic Model autonomously evaluates the suspicious process alerts and, in some cases, performs an additional query against the endpoint agent to gather more contextual information before making a decision. This model is particularly suited to identify host intrusions that may go undetected by endpoint protection platforms, given that the model evaluates suspicious behaviors. Automated Defense goes beyond standalone EDR solutions by classifying the process behavior and incorporating other security events as corroborating evidence for the escalation.

The Cloud Monitoring Analytic Model autonomously monitors authentication logs, historical behavioral data and threat intelligence to identify account compromise. This model goes deeper by triaging alerts from cloud resources and applications to identify cloud services abuse and track the attack stage progression.

Summary

Mandiant Automated Defense uses the latest advancements in artificial intelligence, machine learning, modern streaming architectures and unique Integrated Reasoning to keep organizations protected and safe. It acts autonomously—without a heavy system management burden, security engineering or long learning cycles. Automated Defense is a critical component of an extended detection and response (XDR) solution and the Mandiant Advantage platform.

Learn more at **www.mandiant.com**

MANDIANT®
YOUR CYBERSECURITY ADVANTAGE