

THE ATTACKER'S WINDOW

Why detection alone isn't enough,
and how modern SOC's close the gap



Introduction

Most security incidents start with identity misuse. Attackers rely on compromised credentials, quiet privilege manipulation, persistent access in cloud services, and small configuration gaps that are easy to overlook. These patterns do not reveal themselves through isolated events or logs. They only make sense when detection and investigation work together to establish context and respond to contain threats quickly.

SOC teams often struggle with identifying meaningful threats among large volumes of activity, piecing together evidence from multiple systems, and executing containment actions quickly and consistently. Effective SOC operations depend on the ability to detect suspicious activity in context, investigate it efficiently, and respond without unnecessary delay. Context-aware detection helps highlight events that align with known techniques or behavioral changes. AI-assisted investigation brings together related evidence, identifies relationships between entities, and provides a clear view of how an incident unfolded. Automated response helps execute the necessary containment actions once a threat has been validated. Together, these capabilities help address common operational challenges and reduce the time an attacker has to move through the environment.

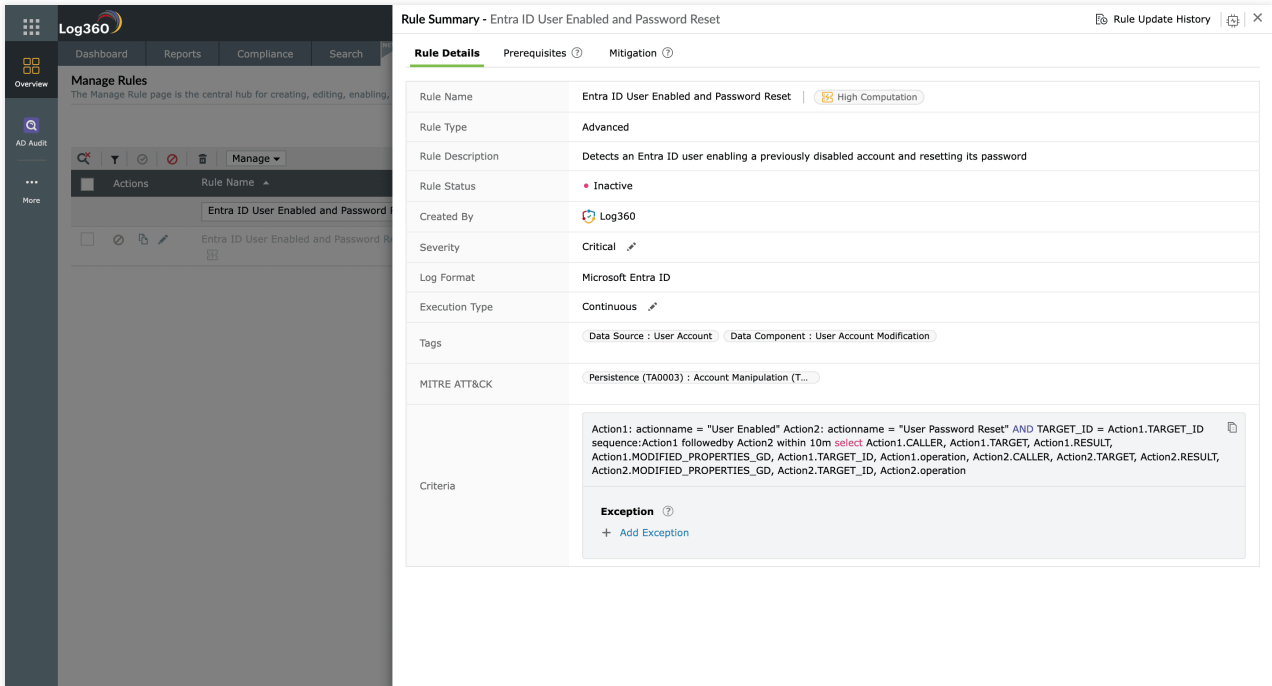
Detection that reads the full attack story

Threat detection used to revolve around finding high severity events. That model does not work when attackers' activities are spread across identity systems, cloud workloads, and network controls. Understanding the relationships between events is critical for identifying attacks that unfold across multiple systems and stages. Correlation, threat intelligence, and behavioral baselines help analysts separate operational noise from attacks that develop gradually. The following use cases show where this shift provides measurable value.

Stealthy privilege escalation in hybrid environments

Privilege escalation usually builds through small identity changes that look normal on their own. A previously disabled user account may be reenabled unexpectedly, an admin group may lose or gain members, or a sensitive group may be modified in Microsoft 365. In a hybrid setup, these changes occur across AD and cloud services, so the full pattern becomes difficult to spot until the attacker already has access.

Log360 connects these signals by evaluating the events together with its correlation logic, UEBA risk indicators, and its cloud-delivered detection rules. Rules such as **Entra ID User Enabled and Password Reset, Multiple Admin Membership Removals, and Unauthorized Group Deletion Detected** help establish whether the changes match known privilege escalation behavior. Object level filters then restrict the analysis to selected users, groups, and OUs, which keeps routine activity out of view and highlights identity changes that matter.



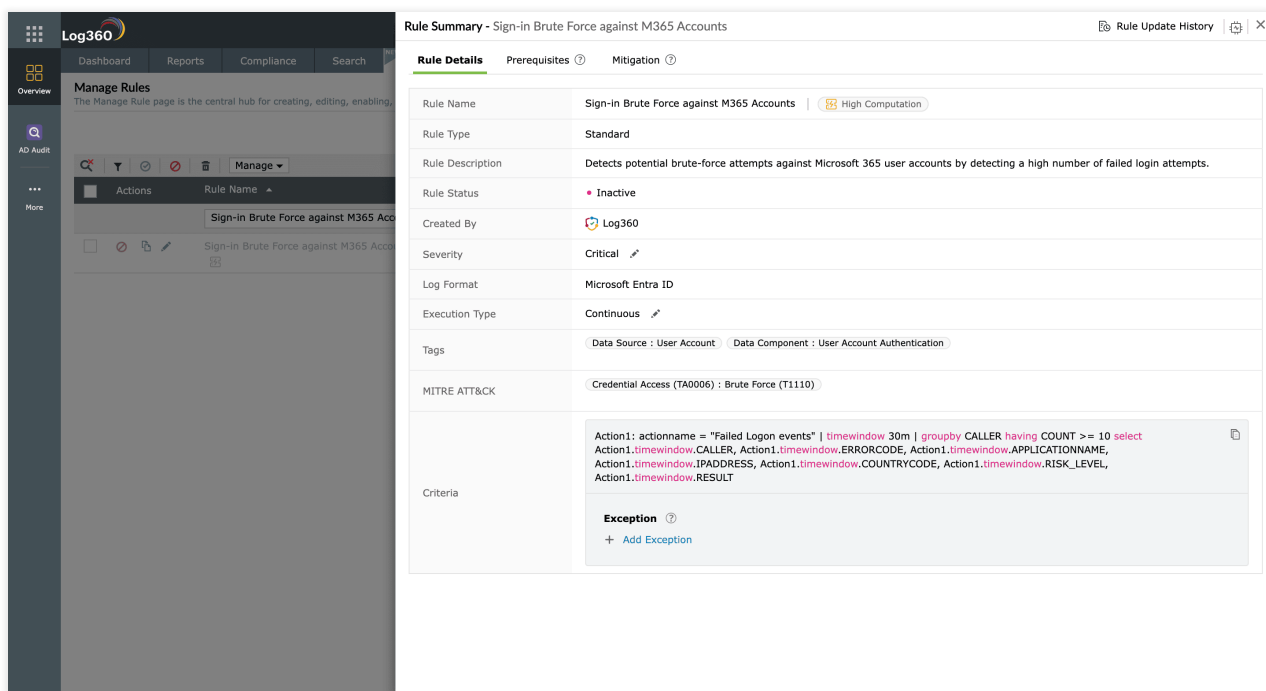
The screenshot displays the Log360 Manage Rules interface. The left sidebar shows the 'Manage Rules' section. The main panel shows the 'Rule Summary' for 'Entra ID User Enabled and Password Reset'. The rule is marked as 'Inactive' and 'High Computation'. The criteria section shows a sequence of actions: 'User Enabled' followed by 'User Password Reset' within 10 minutes, with specific field selections for caller, target, and result. The rule is configured with a 'Critical' severity, 'Microsoft Entra ID' log format, and 'Continuous' execution type. The data source is 'User Account' and the data component is 'User Account Modification'. The MITRE ATT&CK category is 'Persistence (TA0003) : Account Manipulation (T...)'. An 'Exception' section is also visible with an 'Add Exception' button.

Rule Summary - Entra ID User Enabled and Password Reset	
Rule Name	Entra ID User Enabled and Password Reset High Computation
Rule Type	Advanced
Rule Description	Detects an Entra ID user enabling a previously disabled account and resetting its password
Rule Status	Inactive
Created By	Log360
Severity	Critical
Log Format	Microsoft Entra ID
Execution Type	Continuous
Tags	Data Source : User Account Data Component : User Account Modification
MITRE ATT&CK	Persistence (TA0003) : Account Manipulation (T...)
Criteria	Action1: actionname = "User Enabled" Action2: actionname = "User Password Reset" AND TARGET_ID = Action1.TARGET_ID sequence: Action1 followedby Action2 within 10m select Action1.CALLER, Action1.TARGET, Action1.RESULT, Action1.MODIFIED_PROPERTIES_GD, Action1.TARGET_ID, Action1.operation, Action2.CALLER, Action2.TARGET, Action2.RESULT, Action2.MODIFIED_PROPERTIES_GD, Action2.TARGET_ID, Action2.operation Exception + Add Exception

Credential theft and cloud account takeover

Account takeover often begins with a mix of low severity actions. This includes repeated MFA denials, a rise in failed logins, unexpected activity from new locations, or mailbox configuration changes that hide the attacker's movements. These signals rarely appear together unless an account is being probed or controlled.

Log360 brings the identity, authentication, and cloud application events into one detection flow and evaluates them with its continuously updated rule content. Rules such as **Sign-in Brute Force against M365 Accounts** and **Multiple Denied MFA Requests** reflect common credential access techniques seen in real attacks. UEBA signals highlight unusual behavior for the affected account, and object filters keep non-production users out of the detection path so only meaningful account behavior contributes to the alert.



The screenshot displays the Log360 Manage Rules interface. On the left is a sidebar with navigation options like Overview, AD Audit, and More. The main area shows a 'Rule Summary' for 'Sign-in Brute Force against M365 Accounts'. The rule is marked as 'Inactive' and 'High Computation'. It is a 'Standard' type rule that detects potential brute-force attempts against Microsoft 365 user accounts by detecting a high number of failed login attempts. The rule was created by Log360 and has a 'Critical' severity. The log format is 'Microsoft Entra ID' and the execution type is 'Continuous'. The data source is 'User Account' and the data component is 'User Account Authentication'. The MITRE ATT&CK is 'Credential Access (TA0006) : Brute Force (T1110)'. The criteria section shows a KQL query: `Action1: actionname = "Failed Logon events" | timewindow 30m | groupby CALLER having COUNT >= 10 select Action1.timewindow.CALLER, Action1.timewindow.ERRORCODE, Action1.timewindow.APPLICATIONNAME, Action1.timewindow.IPADDRESS, Action1.timewindow.COUNTRYCODE, Action1.timewindow.RISK_LEVEL, Action1.timewindow.RESULT`. There is also an 'Exception' section with a '+ Add Exception' button.

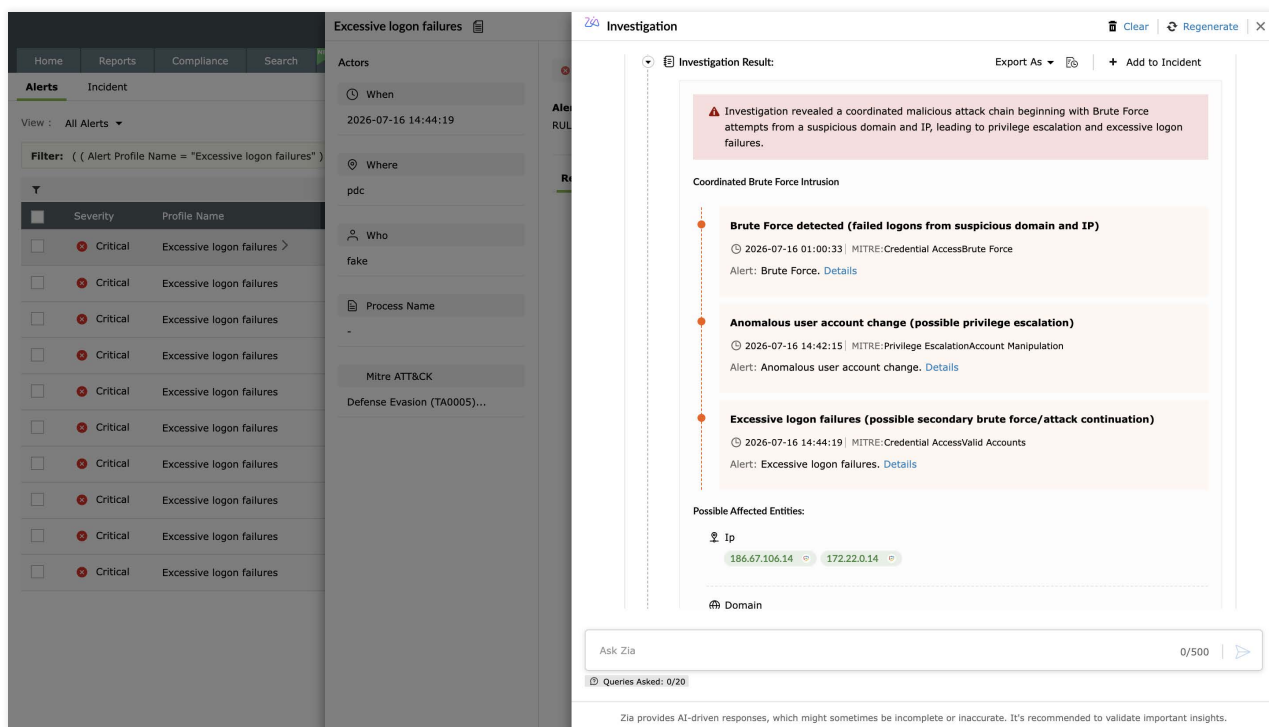
How Agentic AI reconstructs attacks automatically

Even when a threat has been detected, analysts still need to understand how the activity unfolded, what systems were affected, and whether additional actions occurred. Alerts rarely explain the full incident. They must gather logs, verify relationships between events, recreate timelines, and identify affected accounts. These steps slow down investigation and leave room for missed context. AI that organizes evidence, connects related events, and explains how an incident unfolded helps analysts validate conclusions without searching through data manually. The following examples illustrate where this accelerates real investigations.

AI-led investigation of a suspected credential compromise

A suspected credential compromise generates several related signals, like failed logins from different endpoints, password resets, mailbox changes, and repeated authentication against various services. Analysts have to stitch these events together to understand what happened, which is difficult when the events come from different systems.

Log360 brings these events into a unified investigation sequence. Detection signals such as **Mailbox Mail Forwarding Enabled** and **Unusual Mailbox Access** are added to the sequence as evidence. Log360's [AI-powered alert investigation agent](#) analyzes the alert, identifies the users, hosts, IP addresses, and other entities involved, correlates related alerts and authentication activity, and reconstructs the sequence of events leading to the compromise. It then generates a structured investigation report that highlights the attack chain, affected entities, supporting evidence, and recommended actions. This helps analysts understand how the compromise occurred and what actions followed without manually correlating data across multiple systems.



The screenshot displays the ManageEngine Log360 interface. On the left, a sidebar shows navigation tabs: Home, Reports, Compliance, and Search. Below these, there's a section for Alerts and Incidents, with a filter set to 'All Alerts'. A table lists several alerts, all with a severity of 'Critical' and the profile name 'Excessive logon failures'. The main panel shows an 'Investigation' report for the alert 'Excessive logon failures' (2026-07-16 14:44:19). The report title is 'Coordinated Brute Force Intrusion'. The investigation result states: 'Investigation revealed a coordinated malicious attack chain beginning with Brute Force attempts from a suspicious domain and IP, leading to privilege escalation and excessive logon failures.' The report details three key events:

- Brute Force detected (failed logons from suspicious domain and IP)**: 2026-07-16 01:00:33 | MITRE: Credential AccessBrute Force. Alert: Brute Force. Details.
- Anomalous user account change (possible privilege escalation)**: 2026-07-16 14:42:15 | MITRE: Privilege EscalationAccount Manipulation. Alert: Anomalous user account change. Details.
- Excessive logon failures (possible secondary brute force/attack continuation)**: 2026-07-16 14:44:19 | MITRE: Credential AccessValid Accounts. Alert: Excessive logon failures. Details.

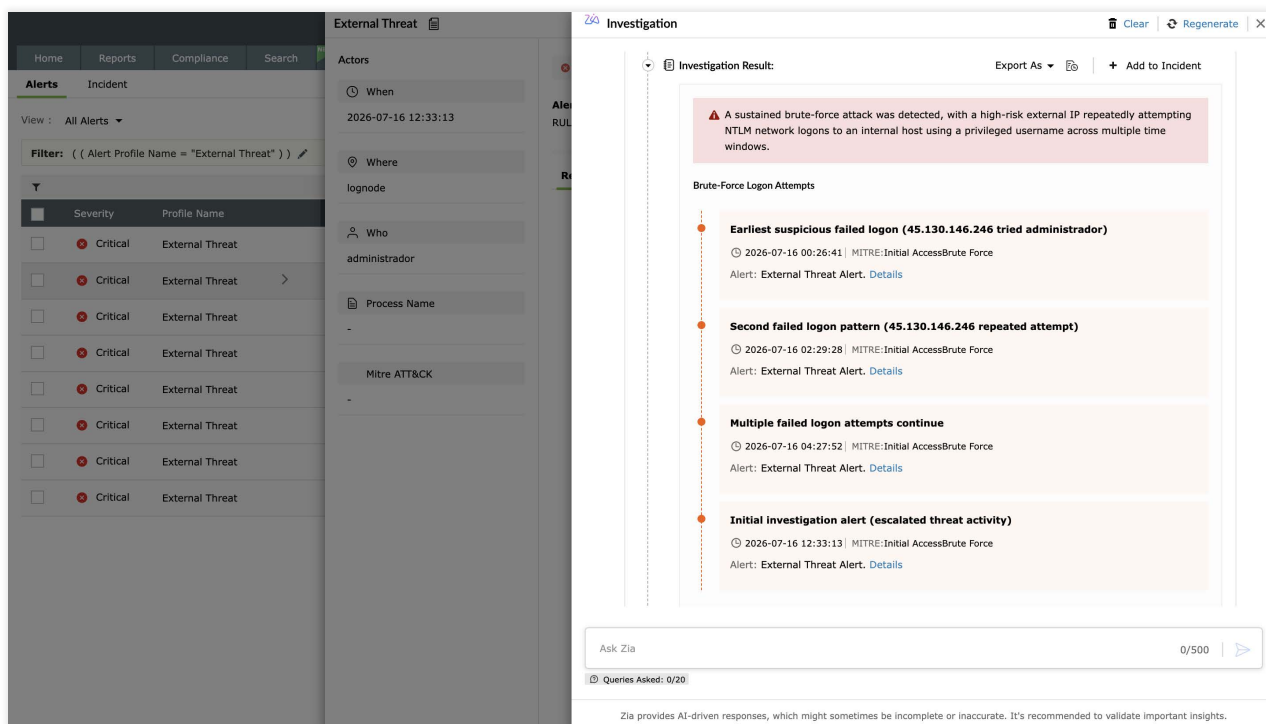
Below these events, a section titled 'Possible Affected Entities' lists IP addresses: 186.67.106.14 and 172.22.0.14. At the bottom, there's a search bar with the text 'Ask Zia' and a 'Queries Asked: 0/20' indicator.

Phishing-led cloud intrusion investigation

A successful phishing attempt can result in post-compromise activities such as the creation of mail-forwarding rules that redirect messages to attacker-controlled accounts, unusual token usage, privilege modifications, and suspicious file-sharing activity within cloud applications.

These appear across separate logs, which makes it hard to understand how the attacker moved through the account.

Log360 brings these cloud events together and reconstructs the activity path. Rules such as **Mail Flow Rule for Forwarding Created** and **Anonymous Sharing Link Created** help identify the changes associated with persistence and data access after a phishing event. Log360's alert investigation agent analyzes the alert, correlates related cloud activities, identifies the users, applications, IP addresses, and other entities involved, and reconstructs the sequence of events associated with the intrusion. It presents the findings in a structured investigation report, helping analysts understand how the phishing attack progressed, what actions were performed, and which resources were affected.



The screenshot displays the ManageEngine Log360 interface. On the left, a sidebar shows navigation options like Home, Reports, Compliance, and Search. The main area is titled 'External Threat' and shows a list of alerts. The 'Investigation' tab is active, displaying a detailed investigation result for a brute-force attack. The result includes a summary of the attack, a timeline of logon attempts, and specific details about the earliest suspicious failed logon, a second failed logon pattern, multiple failed logon attempts, and an initial investigation alert. The interface also features a search bar and a 'Queries Asked' counter.

Consistent, cross-system response at machine speed

Once an incident is verified, containment often depends on a precise sequence of actions. These include account controls, network rules, session revocation, host isolation, and notifications. Manual execution increases delays and causes inconsistencies between analysts. Automation ensures these steps run in the right order, with the right parameters, every time. The following use cases show how this improves response accuracy and speed.

Containing a compromised account automatically


Once an attacker gains access to an account, they usually attempt lateral movement, privilege abuse, or repeated authentication attempts. Fast containment is necessary to prevent further access to systems or data.

After a credential compromise has been detected and investigated, Log360's SOAR playbooks can carry out containment steps through automated workflows. Based on the investigation findings, Log360 can automate actions such as disabling the affected account, revoking active sessions, resetting credentials, blocking the source IP, creating incident tickets, notifying stakeholders, and documenting each step in the incident record for follow up analysis. These actions can be coordinated across identity systems, security controls, and ticketing platforms through a single workflow, with analyst approval steps included where required.

Automated response to malicious traffic or connections

Malicious outbound communication, unexpected inbound traffic, or repeated probing attempts suggest reconnaissance or contact with an attacker controlled server. Rapid action is required to block further communication and secure the affected host.

After investigation confirms malicious activity, Log360's SOAR playbooks help automate containment and remediation actions across multiple systems. Based on the investigation findings, Log360 can automate actions such as blocking the identified IP or domain, updating enforcement settings where required, isolating the host when necessary, creating tickets, notifying responders, and logging every action in the incident record so the investigation can proceed with complete data. Analyst approval checkpoints can also be incorporated into the workflow when additional validation is required before executing critical actions.



What connected TDIR means for your team

A modern SOC depends on the connection between detection, investigation, and response. When these functions operate as a unified workflow, analysts can understand incidents faster, make informed decisions with greater confidence, and contain threats before attackers gain momentum.

The challenge for many security teams is not a lack of visibility, but the operational friction created by disconnected tools and manual handoffs. Detection identifies suspicious activity, investigation provides the context needed to assess impact, and response acts on validated findings. When these processes are connected, analysts spend less time gathering information and coordinating actions, and more time reducing risk.

Log360 brings these capabilities together in a single operational workflow. By combining AI-powered investigation with unified detection, response, compliance, and security analytics on a shared data foundation, Log360 helps SOC teams accelerate investigations, streamline response, and strengthen resilience against identity-focused attacks, cloud threats, and lateral movement.



Our Products

Log360 | ADAudit Plus | EventLog Analyzer | DataSecurity Plus
Exchange Reporter Plus | M365 Manager Plus

About Log360

Log360 is a unified SIEM solution with integrated DLP and CASB capabilities that detects, prioritizes, investigates and responds to security threats. Vigil IQ, the solution's TDIR module, combines threat intelligence, an analytical Incident Workbench, ML-based anomaly detection and rule-based attack detection techniques to detect sophisticated attacks, and it offers an incident management console for effectively remediating detected threats. Log360 provides holistic security visibility across on-premises, cloud and hybrid networks with its intuitive and advanced security analytics and monitoring capabilities.

For more information about Log360, visit manageengine.com/log-management/ and follow the [LinkedIn page](#) for regular updates.

\$ Get Quote

↓ Download