ManageEngine
Log360

Customizing event correlation rules for your organization: A best practice document

www.manageengine.com/log-management

Table of Contents

Introduction

Log360's event correlation engine alerts you to potential threats and attacks at their early stage, and helps you prevent or contain their impact. The solution detects these incidents by correlating logs across different network devices. Log360 comes with over 30 predefined correlation rules which cover several common attack patterns.

As with any technique, correlation works best when you adapt it to your network environment. By following the best practices outlined in this document, you can optimize the correlation engine to suit your security needs.

Correlation rule structure

Before listing the best practices, it is important to understand how a correlation rule is structured. A correlation rule is a pattern or a template that identifies a series of connected events across multiple logs to detect a security incident. The following figure illustrates the different components of a correlation rule that can be customized to meet your needs:

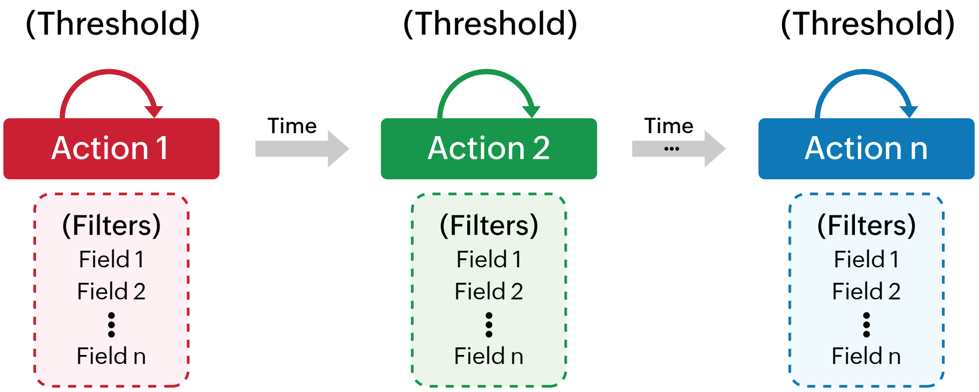

- Correlation rule: A correlation rule defines a sequence of actions (events that occur in network) that denote an attack pattern.

- Actions: Actions are individual events that make up a correlation rule. For example, logon failure, logon success, and file modification are all actions. You can combine these actions in this order to form a simple data breach rule: several logon failures, followed by logon success, followed by file modification. Log360 offers you the capability to add context to this rule with its field-level correlation, explained later in this section.

- Time between actions: The time interval between the occurrence of one action and the other. For instance, let's say a rule has the following sequence of actions: user logon, file permission modified, and user logoff. If these actions occur in quick succession, then it is a suspicious incident. So the rule is only meaningful if we specify times between the actions: user logon, followed within ten minutes by a file permission modification, followed within two minutes by a user logoff.

- Threshold for an action (optional): A threshold defines the minimum number of times a specific action in the rule has to occur. For instance, in case of brute force, the attacker might fail to log on at least five times before the successful logon. Instead of adding the failed logon action five times and providing five time intervals, you can add it just once and specify a threshold value of five times within ten minutes.

- Filters for an action (optional): Filters allow you to impose conditions on the fields within each action. This is known as field-level correlation. This capability empowers you to granularly customize rules to detect attacks specific to your environment. For instance, logon failures on critical database servers need more attention than those occurring on workstations. You can specify such conditions—logon failure on database server—with filters. Filters can be used to select specific users, devices, source/destination machine, and more for each action.

Best practice #1: Select rules based on your business context

Every organization has different security requirements. To understand which rules are most relevant to you, you can first look at the rule description on the Manage Rules page (accessible by going to the Correlation tab → Manage Rules). Look into the various rule descriptions to know which are valid to your business environment. You can even go into the rule definition page (access this by clicking on the update icon next to the required rule), and check how the rule is structured.

For example, if you have strict policies in place which prevent users from downloading third party software, you may not need to enable the rules related to software management. Correlation is a highly memory-intensive process. Select the rules you need carefully to ensure optimal performance.

Best practice #2: Customize rules granularly to monitor the most critical resources

When you create a custom correlation rule, ensure that it is applicable only for the specific users and/or devices you wish to monitor. This applies to predefined rules too. Predefined rules are configured to apply to a wide range of devices in your network.

For instance, the rules from the "Database threats" category check database logs across your network. However, you may need to monitor only one or two critical database servers which hold confidential data, or a specific set of blacklisted users. In such cases, when you choose a predefined correlation rule, it is best that you refine it to apply only to select users or devices, as required.

Best practice #3: Select suitable thresholds and time windows for each rule

It is important to select appropriate threshold limits and time windows for your correlation rule. The values in the predefined rules have been selected with care for each use case. However, every network environment differs in size, structure, and other factors. Hence, when you enable any predefined rule, the values may require some adjustment to suit your organization. Please define these based on your business context.

For example, the "Malicious URL requests" rule alerts you when your web servers receive more than five malicious requests from the same source, within two minutes. If you have a high number of web servers, you might expect a larger number of malicious requests. You can update the rule so that you are alerted only if ten such requests are received in two minutes.

Best practice #4: Fine tune correlation rules by adding/removing actions

The sequence of actions in a rule explains how an attack is carried out. The predefined rules are configured to represent the most popular ways in which each attack is carried out. For instance, the "Suspicious file access" rule describes an anomaly when a user fails to attempt a file several times, and is then suddenly able to access it. To suit your requirements, you can always add or remove actions to make this definition more or less specific.

In the given example, instead of just failed and successful file accesses, you can make the definition more specific by adding the "file permission modified" action in between the two. Alternatively, you can remove the successful file access and simply check for multiple failed attempts to access a file.

Best practice #5: Review performance of rules

Once you have enabled a set of correlation rules, it is important to review their performance from time to time. You may notice that a specific rule is giving you too many alerts. If this happens, review the reports and understand why it is happening. You can use the above best practices as a guide to identify why a rule may not be performing well - for instance, you may be getting too many alerts if a rule is applied to all Windows devices, as opposed to just a few critical servers.. Once you identify the reason, you can refine the rule definition and review its performance again after a while.

How to customize correlation rules

Several of the best practices listed above involve customizing correlation rules to suit your environment. To customize a correlation rule, go to:
Correlation → Manage rules → Select the Update icon next to the required rule.

This opens the rule builder page, which displays the existing structure of the rule. You can add, remove, or rearrange actions, change the threshold limits and time windows between actions, and use filters to apply the rule to specific users or devices. This video explains how you can use the rule builder.

Building custom correlation rules

If you would like to start building correlation rules from scratch, you can use our white paper to help you strategize and build custom use cases for your business.

If you have a use case in mind but are unsure how to create it as a correlation rule, please reach out to our support team at log360-support@manageengine.com and we will be glad to help you set up the new rule.