Bloor

ARTICLE REPRINT

Simplifying rules

Published: 8th April 2011

Rules are a pain. Whether it is data quality rules, correlation rules for SIEM (security information and event management), transformation rules or any other sort of rules, the big problem with rules is that they proliferate. And once they breed they not only continue to do so but each and every one of them needs maintaining. That means that you end up throwing more and more resource at them for maintenance and have less and less time for developing new capabilities.

The problem is that there are not a lot of alternatives to rules and these are often onerous in themselves. Even where these alternatives may be useful they often have to co-exist with rules-based approaches. So, there is a significant imperative towards making rules simpler to define and maintain. One company that has set out to tackle this issue, in the SIEM market, is LogRhythm.

The company has introduced, as an option, its Advanced Intelligence (AI) Engine which is designed to make the definition of correlation rules much simpler and easier, using a drag-and-drop based graphical approach rather than relying on scripting. These correlation rules are based on some 150,000 base rules that LogRhythm ships with its product, but it's the correlation rules that are the problem with SIEM because it is recognising what set of related events represents a particular type of attack that is the problem, not to mention doing so in as close to real-time as possible.

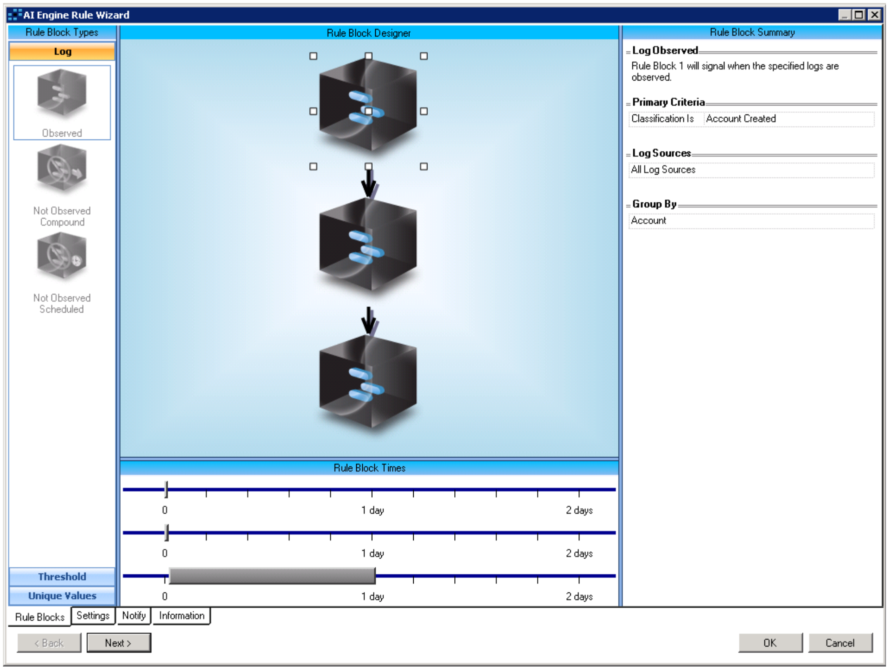

Log Observed

Rule Block 1 will signal when the specified logs are observed.

Primary Criteria

Classification Is: Account Created

Log Sources

All Log Sources

Group By

Account

Figure 1: LogRhythm screenshot showing the three blocks referred to in the text

The actual way that the product allows you to define these correlation rules is interesting because it does not employ what you would call your typical drag-and-drop interface, as Figure 1 illustrates.

The point here is that you have three "blocks": a log observed block, a threshold block and a unique values block. Each of these can be used in three different ways. For example, the log observed block can have three states: observed, not observed scheduled or not observed compound. Put simply, you can drag these blocks to specify primary criteria.

I have to say that this is a seriously cool way of building these correlation rules. It's easy to get the hang of and simple to execute. There are, of course, lots of options to, for example, define filters, set alerts and so forth and, needless to say, you can store rules for reuse. The company will also be setting up a community that will allow users to share correlation rules.

LogRhythm claims that difficulty with defining correlation rules is the reason why SIEM is not as widely deployed as it might be. I am not sure that I completely agree with that but it is certainly a reason, and it is good to see the company addressing this issue, especially in such an appealing fashion.

Philip Howard
Research Director - Data Management
Bloor Research