

Brought to you by



Anomaly Detection in Cybersecurity

for
dummies[®]
A Wiley Brand

ManageEngine Special Edition



Score risks and
prioritize threats

Be alerted instantly to
abnormal activity

Prevent internal and
external attacks

Ram Vaidyanathan

About ManageEngine

As the IT management division of Zoho Corporation, ManageEngine prioritizes flexible solutions that work for all businesses, regardless of size or budget. ManageEngine has 90+ products and free tools that comprehensively cover your IT needs, at prices you can afford. More than 180,000 companies worldwide rely on ManageEngine products to ensure the optimal performance of their critical IT infrastructure, including networks, servers, applications, desktops, Active Directory, and more. With over 4000 employees working round the clock to make product requests a reality, ManageEngine focuses on simplifying IT for everyone.



Anomaly Detection in Cybersecurity

ManageEngine Special Edition

by Ram Vaidyanathan

**for
dummies[®]**
A Wiley Brand

Anomaly Detection in Cybersecurity For Dummies[®], ManageEngine Special Edition

Published by

John Wiley & Sons, Inc.

111 River St.

Hoboken, NJ 07030-5774

www.wiley.com

Copyright © 2022 by John Wiley & Sons, Inc., Hoboken, New Jersey

No part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, scanning or otherwise, except as permitted under Sections 107 or 108 of the 1976 United States Copyright Act, without the prior written permission of the Publisher. Requests to the Publisher for permission should be addressed to the Permissions Department, John Wiley & Sons, Inc., 111 River Street, Hoboken, NJ 07030, (201) 748-6011, fax (201) 748-6008, or online at <http://www.wiley.com/go/permissions>.

Trademarks: Wiley, For Dummies, the Dummies Man logo, The Dummies Way, Dummies.com, Making Everything Easier, and related trade dress are trademarks or registered trademarks of John Wiley & Sons, Inc. and/or its affiliates in the United States and other countries, and may not be used without written permission. All other trademarks are the property of their respective owners. John Wiley & Sons, Inc., is not associated with any product or vendor mentioned in this book.

LIMIT OF LIABILITY/DISCLAIMER OF WARRANTY: WHILE THE PUBLISHER AND AUTHORS HAVE USED THEIR BEST EFFORTS IN PREPARING THIS WORK, THEY MAKE NO REPRESENTATIONS OR WARRANTIES WITH RESPECT TO THE ACCURACY OR COMPLETENESS OF THE CONTENTS OF THIS WORK AND SPECIFICALLY DISCLAIM ALL WARRANTIES, INCLUDING WITHOUT LIMITATION ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. NO WARRANTY MAY BE CREATED OR EXTENDED BY SALES REPRESENTATIVES, WRITTEN SALES MATERIALS OR PROMOTIONAL STATEMENTS FOR THIS WORK. THE FACT THAT AN ORGANIZATION, WEBSITE, OR PRODUCT IS REFERRED TO IN THIS WORK AS A CITATION AND/OR POTENTIAL SOURCE OF FURTHER INFORMATION DOES NOT MEAN THAT THE PUBLISHER AND AUTHORS ENDORSE THE INFORMATION OR SERVICES THE ORGANIZATION, WEBSITE, OR PRODUCT MAY PROVIDE OR RECOMMENDATIONS IT MAY MAKE. THIS WORK IS SOLD WITH THE UNDERSTANDING THAT THE PUBLISHER IS NOT ENGAGED IN RENDERING PROFESSIONAL SERVICES. THE ADVICE AND STRATEGIES CONTAINED HEREIN MAY NOT BE SUITABLE FOR YOUR SITUATION. YOU SHOULD CONSULT WITH A SPECIALIST WHERE APPROPRIATE. FURTHER, READERS SHOULD BE AWARE THAT WEBSITES LISTED IN THIS WORK MAY HAVE CHANGED OR DISAPPEARED BETWEEN WHEN THIS WORK WAS WRITTEN AND WHEN IT IS READ. NEITHER THE PUBLISHER NOR AUTHORS SHALL BE LIABLE FOR ANY LOSS OF PROFIT OR ANY OTHER COMMERCIAL DAMAGES, INCLUDING BUT NOT LIMITED TO SPECIAL, INCIDENTAL, CONSEQUENTIAL, OR OTHER DAMAGES.

For general information on our other products and services, or how to create a custom *For Dummies* book for your business or organization, please contact our Business Development Department in the U.S. at 877-409-4177, contact info@dummies.biz, or visit www.wiley.com/go/custompub. For information about licensing the *For Dummies* brand for products or services, contact BrandedRights&licenses@Wiley.com.

ISBN 978-1-119-83861-6 (pbk); ISBN 978-1-119-83862-3 (ebk)

Publisher's Acknowledgments

Project Editor: Daniel Mersey

Development Editor: Faithe Wempen

Acquisitions Editor: Ashley Coffey

Editorial Manager: Rev Mengle

Business Development

Representative: Matt Cox

Production Editor:

Vivek Lakshmikanth

Table of Contents

Introduction	1
About This Book	1
Foolish Assumptions	1
Icons Used in This Book	2
Where to go From Here	2
 CHAPTER 1: Understanding Why Anomaly Detection Is Critical	 3
Comparing Anomaly Types	4
Time Anomalies	4
Count Anomalies	5
Pattern Anomalies	5
Baselining to Establish Normal Activity	6
Robust Principal Component Analysis	6
Markov Chains	7
User Identity Mapping	9
 CHAPTER 2: Scoring Risks to Achieve Better Security	 11
Arriving at a Risk Appetite	12
Identifying the Different Types of Risk	13
Risk of Insider Threats	13
Risk Based on Logons	14
Risk of Data Exfiltration	14
Risk of Account Compromise	16
Calculating the Overall Risk Score	17
Scoring Risks Accurately with Peer Group Analysis	17
Benefits of Peer Group Analysis	18
Building Peer Groups	19
Scoring Risks Even More Accurately with Seasonality	20
 CHAPTER 3: Five Ways to Make Anomaly Detection Work for You	 23
Customizing Anomaly Models	23
Using Peer Group Analysis	24
Accurately Scoring Risks	25
Analyzing with a Watchlist	25
Being Alerted in Real Time	26

Introduction

Any deviation from what's considered normal is an *anomaly*. Not all anomalies are bad, though! For example, it's an anomaly when a retail company that has never sold much on Sundays breaks its all-time daily sales record on a Sunday, or when the pitcher has the highest runs batted in on their baseball team.

It would also be anomalous if an organization's employee accessed a sensitive file server for the first time ever and deleted major chunks of data. That's the kind of anomaly that can signal a major security breach—and that's the kind of anomaly that this book is all about.

About This Book

One of the best ways to defend against both internal and external attacks is to integrate *anomaly detection*, a.k.a. user and entity behavior analytics (UEBA) capabilities, into your security analytics solution.

In this book, we make a strong case that anomaly detection is essential for effective cybersecurity defense. We compare the different types of security anomalies—time, count, and pattern—and explain what each one looks like in real-life situations. You discover how to create baselines using techniques such as Robust Principal Component Analysis (RPCA) and Markov chains, and we explain how to determine your risk appetite and calculate your risk score. There are many risk types, and we break them down for you, and then show how to use two different methods to score risks more accurately: peer group analysis and seasonality.

Foolish Assumptions

This book isn't for everyone—but it may be for you! As we wrote this book, we made a few assumptions about you, our reader. Nothing too personal, we promise! But we did assume that you know a bit about the different kinds of activities that users and hosts can do on a network, and that you're interested in cybersecurity.

Icons Used in This Book



REMEMBER

We've included some handy pointers to make reading easier, so take note of these as you're following along.

Discover key definitions and essential take-aways.



TIP

Seek out the target for tips that can save you time.



WARNING

Watch out for these pitfalls on your journey.

Where to go From Here

This book is written as a reference guide, so you can read it cover to cover or jump straight to the topics that most interest you. You can't go wrong with either choice—both give you a better understanding of anomaly detection and its relevance to your business's success.

Want to know more? Check out <https://www.manageengine.com/log-management/ueba/resources.html>, which offers some interesting use cases and real-life examples of anomaly detection for effective cybersecurity.

In addition, we found the following resources helpful as we wrote this book, and you might find them helpful as well:

<https://www.zdnet.com/article/ghost-in-the-wires-the-kevin-mitnick-interview/>

<https://www.manageengine.com/log-management/ueba/user-and-entity-behavior-analytics-software/data-exfiltration.html>

<https://www.investopedia.com/terms/s/seasonality.asp>

- » Comparing the different types of anomalies
- » Establishing a baseline
- » Mapping user identities

Chapter 1

Understanding Why Anomaly Detection Is Critical

It's 1 a.m. Richard furiously types commands into his laptop, leverages an open Remote Desktop Protocol (RDP) port, and accesses the Kimble Hospital's intranet. He navigates into the chief cosmetic surgeon's workstation and gains further network privileges, and then he connects to several servers that house gigabytes of medical records of patients.

He then goes on to the dark web and types into a forum "For Sale: Records from Kimble Hospital." With a final flourish, he releases malware in Kimble's network that demands \$10 million in ransom in exchange for releasing the sensitive data. Richard then ponders: "Should I also send ransom demands to the patients?"

Let's suppose that you are the person in charge of Kimble Hospital's IT system. You notice several anomalies as Richard performs this ransomware attack. The chief cosmetic surgeon's user account shows a logon at 1 a.m.—a time anomaly. You notice that multiple servers were accessed, one after the other—a count anomaly. Finally, you notice the username, accessed hosts, and time of access have never been seen together before—a pattern anomaly.



WARNING

Anomalies are warning signs that can help you shut this attack down before it happens—or at least before it gains a significant foothold.

So, Kimble Hospital could have saved its reputation, saved its money, and secured patient data by improving its anomaly detection.



REMEMBER

Fast, decisive action can potentially save you millions of dollars. But that's only possible if you happen to notice what's going on in time. And that, in a nutshell, is *anomaly detection*: noticing when things aren't normal and finding out why.

In this chapter, you will learn about the three key types of anomalies and how to detect them, which is a key step in cybersecurity defense. You will also learn how user identity mapping can build a security context around all detected anomalies.

Comparing Anomaly Types

The three types of anomalies mentioned in the preceding scenario—time, count, and pattern—are the most common symptoms of a cybersecurity attack. Let's have a closer look at them.

Time Anomalies

A *time anomaly* occurs when an activity takes place at an unexpected time.



TIP

To detect these anomalies, you should establish a baseline time for all activities of every user and host on the network based on historical behavior. You can then check whether the observed activity is happening at an unusual time compared to the baseline. If the deviation occurs outside the predefined threshold, this would be an anomaly.

Here are a couple of real-life time anomalies:

» **Example of a time anomaly for a user:** An employee who generally logs on between 9 a.m. and 10 a.m. suddenly logs on at 5 p.m.

- » **Example of a time anomaly for a host:** A file is modified on a particular host between 8 and 8:15 a.m., much earlier than the expected window of 4 to 4:15 p.m.

Count Anomalies

A *count anomaly* happens when an unusually high number of activities are performed in a short period of time, either by a user or on a host.



TIP

To detect count anomalies, you should first establish a baseline count for all activities in your system performed by users or on a host. If the observed count is higher than the baseline by more than a certain range, that's a count anomaly.

Check out these count anomaly scenarios:

- » **Example of a count anomaly for a user:** A user with a baseline of three data manipulation language (DML) queries executes more than 20 DML queries on a Structured Query Language (SQL) server to get their hands on sensitive data.
- » **Example of a count anomaly for a host:** A particular router has over 50 configuration changes while the expected number is only 13.

Pattern Anomalies

A *pattern anomaly* takes place when there is an unexpected sequence of events.



REMEMBER

Each of these events may not seem anomalous when considered in isolation. However, when you consider them together in one sequence, and that deviates from the expected, it becomes a pattern anomaly.

Establish a *pattern baseline*. Then you compare all observed activities with this baseline to determine whether an anomaly exists.

Unlike with time and count anomalies, you don't need to consider anomalous patterns for users and hosts separately; a pattern under examination may already contain both host and user activities.

Some examples of pattern anomalies are when:

- » **A user account performs a software installation on a host at an unusual time:** The pattern analyzed is Host name > Username > Time. The combination of host name and username is not unusual, but the time of installation is, and triggers an anomaly.
- » **An unexpected user reads a file containing sensitive information on a host through a USB port at a particular time:** The pattern being analyzed is Host name > Username > Event ID > Time. The Host name > Username event triggers an anomaly, although everything else is normal.
- » **A firewall rule is changed at an unexpected time:** The pattern analyzed is Host name > Rule ID > Time. The RuleID > Time event is anomalous.

Baselining to Establish Normal Activity

You can't identify anomalies without established baselines of expected activity. An anomaly detection system that uses machine learning algorithms can create those baselines for you effectively. The system undergoes a training period during which it learns the baseline behavior of every user and host.

There are numerous techniques or statistical models used to decipher anomalous behaviors in a network. This section looks at two of them: robust principal component analysis and Markov chains.

Robust Principal Component Analysis

You can use *robust principal component analysis (RPCA)* to detect time and count baselines.

The RPCA algorithm is a variation of the technique of *principal component analysis (PCA)*. When you train the machine learning model using PCA, you have it look at all the historical data and find the line of best fit. RPCA improves the machine learning model by accounting for any outliers.

The RPCA model views the observed events as a matrix summation of expected and outlier events. This enables you to identify

the anomalies within a set of values. (the anomalies would be the outlier events). Figure 1-1 shows how this algorithm works.




FIGURE 1-1: The RPCA algorithm predicts expected values by using a “line of best fit.”

Markov Chains

A *Markov chain* is a sequence of events where the probability of the next event in a sequence is solely dependent on the state of the current event. It compares each action a host or user performs to a list of possible actions. If it finds an event to have a low probability, the algorithm identifies that event as an anomaly.

The anomaly detection system learns over time what a low probability event is. Markov chains can be used for distinguishing between normal and abnormal patterns of behavior.



TIP

To implement this algorithm, break up the pattern of interest into sets of two consecutive activities and check whether the second activity in the set has a good probability of occurring after the first. Your machine learning-powered anomaly detection system defines “good probability” by studying past behavior.

For example, let’s say you want to look for anomalies in this pattern: A software installation is performed by a user on a host at an unusual time. The general pattern that you want to analyze here is:

Username > Host name > Time

To implement the Markov chain algorithm, you would break up this pattern into two parts:

- » Part 1: Username > Host name
- » Part 2: Host name > Time

Then the machine learning model checks whether the probability of the host being accessed by the user (Part 1) and the probability of the host being accessed during the time window (Part 2) are both acceptable. If either is not, that implies that the corresponding activity is unexpected, and the pattern as a whole is anomalous. Figure 1-2 shows how this process works.

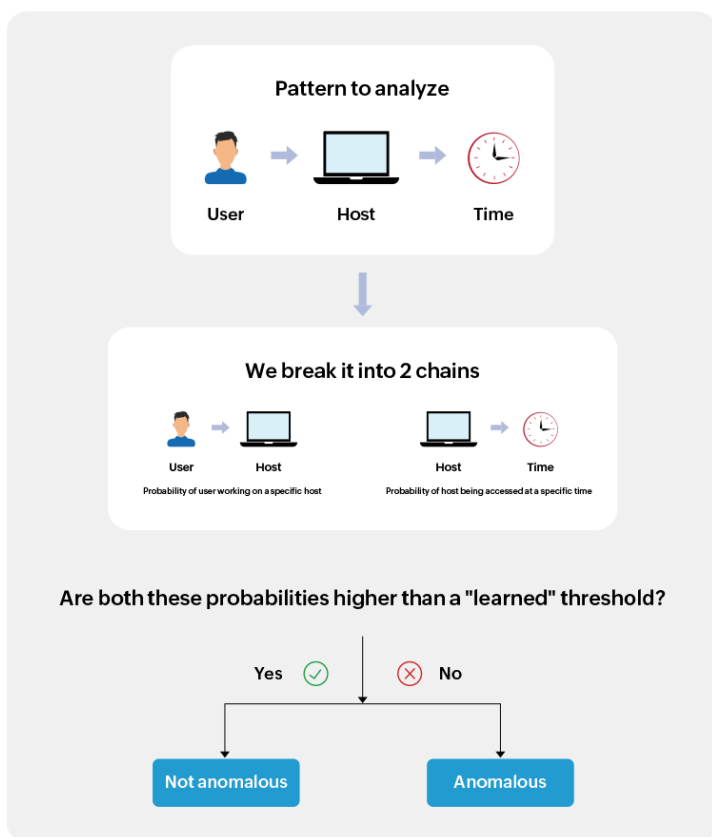


FIGURE 1-2: Using Markov chains to detect anomalies in a pattern.

User Identity Mapping

To track time, count, or pattern anomalies more effectively, you need visibility into what users are doing in various parts of the network.



REMEMBER

You should be able to map all users' activities into a single thread. This is called *user identity mapping (UIM)*.

Consider the following sequence of activities:

1. An employee logs on to the organization's domain from a remote location through successful Active Directory authentication.
2. Within five minutes of the logon, the employee changes a couple of firewall policy rules.
3. Suddenly, the employee executes an abnormal process.
4. They log onto a database server that holds business-critical information.
5. They modify files in the database server.

Unless you linked this sequence of actions to the particular employee, you wouldn't know that something malicious was taking place. Your anomaly detection system needs to link the user ID, or other user identifiers, to build a complete picture of the attack. Figure 1-3 shows how this works.

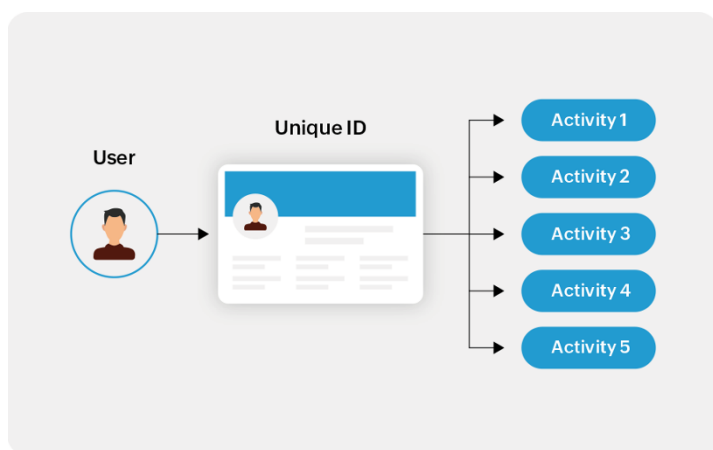


FIGURE 1-3: How user identity mapping works.

- » Determining your risk appetite
- » Calculating your risk score
- » Understanding the different types of risks
- » Using peer group analysis to score risks
- » Using seasonality to score risks

Chapter 2

Scoring Risks to Achieve Better Security

Kevin Mitnick, arguably the world's most famous hacker, once said, "You can never protect yourself 100%. What you do is protect yourself as much as possible and mitigate risk to an acceptable degree. You can never remove all risk."



WARNING

Ever-present risk makes IT risk assessments critical. What level of risk is your company willing to tolerate? By estimating the risk posed by different users and hosts, you can better manage your company's security.



TIP

The level of tolerable risk can vary from company to company, and you should decide how much is okay based on factors such as the nature of business, the industry, the market, the pricing and revenue model, size of the customer base, and more. In addition to this, you should also consider the risk posed by specific users and hosts.

This chapter explains how to calculate risk scores for users and hosts. You find out about some of the risks and how to score them using different methods.

Arriving at a Risk Appetite

To determine the acceptable risk posed by specific users and hosts, consider criteria such as:

- » **The role of the employee:** You may want to decrease the risk appetite if the employee is from the C-suite. These employees may have higher levels of privilege on their accounts, and the repercussions of a risky activity may be much higher.
- » **Employees serving a notice period:** Employees who are serving their notice period may pose a bigger risk. You may want to decrease exposure in these cases.
- » **Servers or systems storing sensitive information:** You may want to treat business-critical data—customer contact information, trademarks or copyrights, product design documents, and technology roadmaps—with more care. Servers or systems storing these data should not be exposed to much risk.
- » **Service accounts with higher privileges:** Service accounts usually have higher privileges so you may want to lower the risk. (See the sidebar, “What is a service account?” for more info.)

Once you know the acceptable levels of risk, find the scenarios where the risk score exceeds your threshold.

There are many methods to calculate a user or host risk score.

WHAT IS A SERVICE ACCOUNT?

Service accounts are non-human accounts used by critical applications or services to interact with their operating systems, as well as to execute batch files, schedule tasks, and work with applications hosted across databases, file systems, and devices. These accounts are controlled by non-human “users” such as systems, scripts, and applications. Service accounts typically require elevated privileges so that they can access business-critical applications, databases, web services, and APIs. That’s why it’s so risky for one to be compromised.



TIP

One way of calculating risk is to use a straightforward percentage score. Your machine learning algorithm assigns a risk score between 0 (no risk) and 100 (maximum risk) based on factors such as the significance of the action from a security standpoint, the extent of the deviation from the baseline, the frequency of the deviation, and the time elapsed since the deviation.

What logic should your algorithm use to assign risk scores? After all, every anomalous activity is not the same. For example, an anomaly involving a database server could be riskier than an anomaly involving a USB drive, even though the latter shows a larger deviation from the baseline. For instance, the database servers in an organization could hold vital information, and even a small deviation would be of concern. The next section considers this in more detail.

Identifying the Different Types of Risk

One way to assign the right risk score to each anomalous activity is by breaking up the overall risk into four sub-risks: insider threats, risk-based on logons, data exfiltration risk, and account compromise risk.

Let's have a look at each of these next.

Risk of Insider Threats

An *insider threat* is any malicious or unintended security threat to an organization's data or information systems posed by an individual in the organization or operating inside the organization.



WARNING

Many anomalous activities in a network could indicate an insider threat:

- » **Abnormal data deletions:** These happen when data is deleted unexpectedly from any part of the network, either in storage or during transit.
- » **Logon success anomalies:** These happen when unexpected logon successes happen on hosts, servers, databases, or cloud services.
- » **Application whitelisting anomalies:** These happen when abnormal EXE or DLL files are run.

The threat actor may use various techniques to execute the above anomalies. (Table 2-1 provides some examples)

TABLE 2-1 **Abnormal actions that could indicate an insider threat**

Abnormal data deletions	Logon success anomalies	Application whitelisting anomalies
Abnormal file deletes	Abnormal successful logons to AWS	Abnormal EXE and DLL files run by a user
Unusual directory deletion	Abnormal host logon event	Abnormal scripts run by a user
Abnormal database deletion by a user	Abnormal SQL server logon	Multiple EXE and DLL files disallowed to run by a user
Unusual database alteration on a host	Multiple hosts shut down by a user	Multiple scripts run by a user
Multiple files deleted	Unusual host startup	Abnormal software restricted

Risk Based on Logons

Anomalous logon activities occur throughout an attack, from the moment of initial access to privilege escalation, lateral movement, and other stages.



To carry out an attack, someone has to log on to the system, whether they’re a malicious insider or an external attacker, so it’s vital to monitor all logons. Your machine learning-powered anomaly detection system must flag users whose logon events are sufficiently different from their historic patterns.

Both logon successes and failures could constitute anomalies. Table 2-2 provides a partial list of logon anomalies.

Risk of Data Exfiltration

Attackers are usually after your organization’s data.

Data exfiltration is the unauthorized transfer of data from inside an organization to the outside by someone who may or may not be an employee. The transfer can originate from workstations, servers, databases, or network devices.

TABLE 2-2 **Actions that could indicate logon anomalies**

Logon failure anomalies	Logon success anomalies
Abnormal AWS authorization failures	Abnormal AWS successful logins
Abnormal AWS error events	Abnormal host logon event
Abnormal SQL Server logon for a user	Abnormal host logon type
Abnormal Windows logon failure for a user	Abnormal host startup
Numerous logon failures on the IIS server	Abnormal remote IP used for logon



WARNING

Watch the following activities for signs of data exfiltration:

- » Data hoarding.
- » Anomalies in data uploads.
- » Build-up of data exfiltration indicators such as logon success and failure anomalies, abnormal software installations, and application whitelisting anomalies.

Table 2-3 lists some techniques that attackers could use.

TABLE 2-3 **Anomalies indicative of data exfiltration**

Data hoarding	Anomalies in data uploads	Build-up of data exfiltration indicators
Abnormal file downloads	Abnormal creates on a removable disk in a domain	Logon failure anomalies such as the ones discussed in Table 2-2
Abnormal file reads	Multiple removable disk creates by a user	Logon success anomalies such as the ones discussed in Table 2-2
Multiple file reads by a user	Multiple removable disk creates on a host	Unusual software installations
Abnormal network share object permission modified	Unusual removable disk modifications by a user	Multiple services installed on a host by a user
Abnormal file type downloaded	Unusual removable disk modifications on a host	Application whitelisting anomalies such as those discussed in Table 2-1

THE RISK FORMULA

You can approximate the overall risk of a user or host using this formula:

$$OR = w_1 \times TDF_1 \times \text{Risk of insider threat} + w_2 \times TDF_2 \times \text{Logon risk} + w_3 \times TDF_3 \times \text{Risk of data exfiltration} + w_4 \times TDF_4 \times \text{Risk of account compromise}$$

Here are the variables used in that formula:

Variable	Description
OR	Overall risk of a user or host
w_1	Weightage assigned to risk of insider threat
w_2	Weightage assigned to logon risk
w_3	Weightage assigned to risk of data exfiltration
w_4	Weightage assigned to the risk of account compromise
TDF_1	Time decay factor assigned to insider threat
TDF_2	Time decay factor assigned to logon anomalies
TDF_3	Time decay factor assigned to risk of data exfiltration
TDF_4	Time decay factor assigned to risk of account compromise

Risk of Account Compromise

Account compromise takes place when an account is accessed by someone who is not authorized.



REMEMBER

Weak passwords are one of the main reasons for account compromise, but attackers can also use sophisticated tools to compromise accounts with stronger passwords.

Your anomaly detection system should be able to detect tell-tale signs of account compromise such as logon failure anomalies (Table 2-2) and malicious software installations.



WARNING

Here are some events that could indicate malicious software installations:

- » Windows Registry anomalies such as “Registry Permission Changed” or “Registry Value Modified.”
- » Abnormal software installed or uninstalled on a host.
- » Abnormal service installed or uninstalled on a host.
- » Unusual process creations.
- » Multiple software instances installed or uninstalled by a user.

Calculating the Overall Risk Score

Let’s say that your anomaly detection system has computed scores for each user and host that accounts for the four sub-risks: insider threat, logon, data exfiltration, and account compromise. The next step is to derive an overall score that aggregates the sub-risks and other risks that did not fall within the four sub-risks.



TIP

For each potential technique of attack, you need to specify a weight and a decay factor. These variables should be used in the risk calculation. For more on this, take a look at the sidebar, “The risk formula.”

The *weight* quantifies a technique’s importance, and the higher the weight, the greater its importance and higher the risk score.

The *decay factor* signifies how quickly the risk score decreases over time if no further anomalies are observed.

Scoring Risks Accurately with Peer Group Analysis

Your anomaly detection solution has assigned a risk score to each user and host on the network? Wonderful! Now, you can prioritize the risks based on these scores and address the most important challenges first.

But before doing so, there's a way to make risk scoring even more accurate: peer group analysis.

Peer group analysis is the act of identifying users or hosts with similar characteristics or behavioral patterns and classifying them as one group. You can build better security by comparing the observed behavior of a user or host to that of the relevant peer group. The user or host risk score can be impacted positively or negatively depending on the peer group (see Figure 2-1).



FIGURE 2-1: How peer group analysis works.

Benefits of Peer Group Analysis



REMEMBER

Peer grouping can help build more accurate risk scores in many situations. Here are some examples:

- » **First time access of a resource by a user:** A user accesses a critical database server for the very first time. Without peer group analysis, this activity would be deemed risky. But if the user belongs to the peer group of marketing analysts who typically access this database server, the activity will not be flagged.
- » **Logon time anomaly by a user:** A user logs on to the network at a time that deviates heavily from their baseline of expected behavior. If there is no peer group analysis, this could be considered risky. But if the user is a member of a peer group that does show logon activity at that time, the risk score will be lower.
- » **An IT administrator installs unusual software:** An IT administrator installs unusual software on a specific host. This could lead to a pattern anomaly and the IT administrator's risk score could rise. However, after peer group analysis, it's discovered that this user is a part of a peer group called "IT administrators" and this anomalous action does not really deviate from the average behavior of the group. Therefore, the risk score is not raised as much.

- » **There is an abnormal number of file reads on a host by a user:** A sensitive server holds numerous business-critical files on trademarks and product roadmaps. A user reads some of these files and deviates from their expected behavior. Without peer group analysis, the user's risk score would rise significantly. But after peer group analysis, you learn this behavior is typical of a group that contains 100 other members. Therefore, the risk score is not impacted that much.
- » **Thirty users belonging to different departments access a database over a weekend:** An engineering-related database is accessed by 30 users belonging to departments such as IT, pre-sales, sales, and product management. Without peer group analysis, this will seem like a risky activity and each user's risk score will rise. However, with peer group analysis, all of these users are classed into one group and the risk score is not increased as much.
- » **A user has an abnormally high number of successful logons into an Amazon Web Services (AWS) instance:** A user from the accounting team has an abnormally high number of successful logons to an AWS instance. This deviates heavily from the baseline and would trigger an anomaly and an associated risk score. Further, you learn that this user is part of only one peer group called Accounting and the members don't exhibit similar behavior. You increase the user's risk score by a greater amount than you would have otherwise.

Building Peer Groups

Anomaly detection solutions use these methods to create peer groups:

- » **Static:** Uses attributes and details available in databases like Active Directory. Your Active Directory database would contain substantial information about users, such as groups and organizational units they belong to, their reporting manager, geographical location, department, function, and more. Several static peer groups could be built using these attributes. For example, all employees who have the same reporting manager might constitute one peer group.

» **Dynamic:** Uses behavior-related information collected over time to build peer groups. If a user exhibits a behavior for the first time, the system checks whether a peer group that showcases such a behavior already exists. If it does, the user is classed into that peer group. To calculate the risk score, you would then analyze variables such as the number of members in the peer group and the frequency of occurrence of the action.

If a peer group that showcases the user's anomalous behavior does not exist, a new peer group is created and this user becomes the first member of the peer group. In this case, you'd end up with a much higher risk score.

Scoring Risks Even More Accurately with Seasonality

So, you've analyzed the activities that are happening across your network. You've created a list of anomalies. You've built risk scores for every host and user. And you've used peer grouping techniques to further hone the risk scores. There's one last step: you need to account for seasonality.

Seasonality is any activity in your network that happens at regular intervals, whether hourly, daily, weekly or monthly. If an activity that seems normal at first glance is in fact happening at an unusual time, your system should be able to recognize it as an anomaly.

For example, suppose you work for a bank that operates on the first and third Saturday of every month. On Saturday morning, your security analytics platform notices an employee logging into the network. Strangely, it's the *second* Saturday of the month. A lesser trained system would accept this; after all, the employee was online the previous Saturday, so why not today? But yours is well-trained to spot seasonal anomalies just like this. It knows the difference between the various Saturdays of a month. An alarm goes off and the risk score of the employee increases.

An effective anomaly detection solution should be able to show you the risk scores of all users and hosts at any point of time. Figure 2-2 shows how ManageEngine Log360 displays the risk score of users, broken down into the scores of the different sub-risks.




FIGURE 2-2: The overall and the sub-risk scores of a user in Log360.

- » Customizing anomaly models
- » Using peer group analysis
- » Accurately scoring risks
- » Watchlisting users and hosts
- » Getting real-time alerts

Chapter 3

Five Ways to Make Anomaly Detection Work for You

Anomaly detection is an extremely critical capability for defending against cyberattacks. ManageEngine Log360 is just one of the solutions that can help you with this. This chapter gives you five ways for making anomaly detection work for you more effectively.

Customizing Anomaly Models

All anomaly detection systems offer built-in anomaly models. These are nothing but built-in machine learning algorithms that learn the baseline of expected activity for every user and host in the network. To discover time anomalies, a built-in anomaly model might look for what is normal at 15-minute intervals, and it might aggregate the number of user and host activities every hour to discover count anomalies.



TIP

You may want to train your own anomaly model using different time intervals and aggregation periods. This capability is called *custom anomaly modeling*. It can enable you to cater to the specific security situation of your company in a better way.

Figure 3-1 shows the interface you can use for creating a custom anomaly model in Log360.

The screenshot shows the Log360 UEBA interface. The top navigation bar includes 'Dashboard', 'Anomaly Reports', 'Alerts', 'Settings', and 'Support'. The 'Settings' tab is active, and the left sidebar shows 'Settings' > 'Configuration' > 'Anomaly Modeling' > 'Risk Score Customization'. The main content area is titled 'Create New Model' and contains the following sections:

- Basic Details:**
 - Model Name:** A text input field with a 'Description' link to its right.
 - Select Source:** A dropdown menu with '- Select Source -', a '+', and a 'Filter' link.
 - Pivot Fields:** A dropdown menu with '- Pivot Fields -', a '+', and a help icon.
- Anomaly Parameters:**
 - Anomaly Report Views:** A box containing a message: 'No Pivot Fields Selected'.
 - Reporting:** A toggle switch that is currently turned on.
 - Report Name:** A text input field.
 - Custom Groups:** A dropdown menu with 'ela - snmp reports' selected.
 - Risk Scoring:** A toggle switch that is currently turned off.

At the bottom right, there are 'Save' and 'Cancel' buttons.

FIGURE 3-1: Building new anomaly models in ManageEngine Log360.

Using Peer Group Analysis

As explained in Chapter 2, *peer grouping* is the process by which you group users and hosts into distinct peer groups based on their past behavior. If your security analytics platform adopts peer group analysis, it will be able to determine whether a user or host behaves as expected based on the groups it is in. If it doesn't, the system triggers an anomaly alert.

Operating in addition to comparing a user's or host's behavior to its own baseline, peer group analysis helps reduce the number of false positives.

Accurately Scoring Risks

Your anomaly detection solution must be able to assign a risk to every user and host in the network. In ManageEngine Log360, this risk score can range anywhere from 1 to 100 and represents the degree of risk posed by an entity. The risk score depends on the anomalies that the user or host triggers.



TIP

To make the risk score more accurate, your anomaly detection solution should consider:

- » Seasonality factors
- » Anomalous activity weights
- » Time decay factors
- » Peer group analysis

Analyzing with a Watchlist

A *watchlist* is a list of hosts and accounts that you want to keep an eye on, for whatever reason. It typically consists of your riskiest users and hosts, as well as accounts with greater than normal permissions. That list might include admins, users from the C-suite, users who are serving a notice period, and users who are under investigation for certain activities. A watchlist can let you know about entities of interest and their associated risk score.



REMEMBER

Your anomaly detection solution should be able to add users and hosts into a watchlist, and to constantly monitor the risk scores of watchlisted users and hosts on a dashboard.

ManageEngine Log360 provides you the ability to add users to a watchlist, as shown in Figure 3-2.

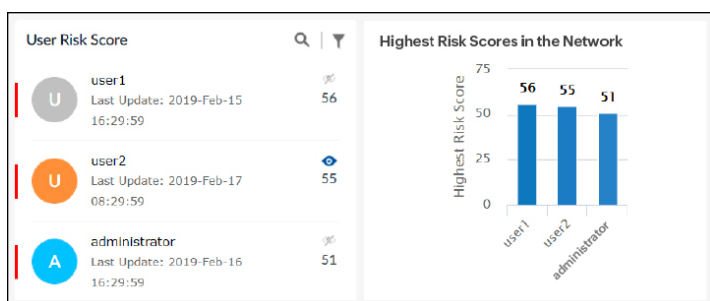


FIGURE 3-2: As the shaded eye symbol shows, user2 has been added to a watchlist in Log360.

Being Alerted in Real Time

With alerts, you can receive notifications about anomalies that happen on the network in real time. For example, you might receive a notification email as soon as an anomaly is identified.

Without real-time alerts, you'd need to log on to your anomaly detection solution every time you wanted to check to see whether there was a new risk your network is exposed to. Figure 3-3 shows the interface for configuring an alert in ManageEngine Log360.

Log360 UEBA

Dashboard | Anomaly Reports | **Alerts** | Settings | Support

Add Alert Profile

*Alert Name: [Description](#)

Severity:

Alert Based On: ☒ Report ☐ Entity ☐ Risk Card

*Select Report: [Add Filter](#)

*Select Entity: [Add](#)

*Alert Message: [Add](#)

Advanced Configuration [▶](#)

Alert Actions: ☐ Email Notification

[Save Changes](#) [Cancel](#)

FIGURE 3-3: Configuring alerts in ManageEngine Log360.



REMEMBER

In Log360, you can configure alerts in three ways:

- » **Alerts based on reports:** You select the anomaly report that contains the information for which you want to be alerted. Then you select users or hosts for which you want the alert.
- » **Alerts based on entities:** You select all the users and hosts for which you want alerts. With this option, the system alerts you when any of the selected users or hosts perform any anomalous activity.
- » **Alerts based on risk score:** You set a threshold risk score for any user in the network, and the anomaly detection system alerts you when anyone exceeds that score. You can set a threshold risk score for overall anomalies, insider threats, data exfiltration, compromised accounts, and logon anomalies.

Comprehensive UEBA Solution

ManageEngine[®]
Log360

ManageEngine Log360 is a UEBA tool which enforces tighter security measures by detecting behavior anomalies, and strengthening your defenses against insider threats and data breaches.



With over 1000 predefined report and alert profiles, Log360 uses machine learning to defend against:



Insider
Threats



Account
Compromise



Data
Exfiltration

Try Log360 now!

<https://zoho.to/ManageEngine-Log360>

Download now

Active Directory Security Solution

ManageEngine[®]
ADAudit Plus

ADAudit Plus is the real-time change auditing and analytics component of Log360 that helps secure your Active Directory, Azure AD, file servers, and other Windows servers.



ADAuditPlus leverages machine learning to notify security personnel in the event of:



Malicious
logins



Anomalous
change activity



Data
breaches

Start your 30-day free trial now!

<https://zoho.to/ADAuditPlus>

Download now

Improve cybersecurity with anomaly detection

Anomaly detection or user and entity behavior analytics help organizations detect cyber crimes before they can cause a catastrophe. A security analytics solution with anomaly detection capabilities uses machine-learning algorithms to learn what behavior is normal and not normal for every user and host in a network. In this way, it can alert you when an anomaly occurs, even though the activity seems benign. This book explains the intricacies of how anomaly detection works.

Inside...

- Discover different types of anomalies
- Identify different types of risk
- Understand scoring risks of users and hosts
- Monitor risk thresholds and generate alerts
- Stop internal and external attacks

ManageEngine 

Ram Vaidyanathan is a cyber risk expert and technical evangelist at ManageEngine. He is an expert in the various ways in which cyber attackers can strike organizations and exfiltrate data, and about how organizations can defend themselves. He also speaks at seminars and conferences about security best practices.

Go to **Dummies.com™**
for videos, step-by-step photos,
how-to articles, or to shop!

ISBN: 978-1-119-83861-6
Not For Resale

**for
dummies®**
A Wiley Brand



WILEY END USER LICENSE AGREEMENT

Go to www.wiley.com/go/eula to access Wiley's ebook EULA.