

Making Everything Easier!™

SolarWinds Special Edition

Network Monitoring

FOR
DUMMIES®
A Wiley Brand

Learn:

- Why you need monitoring
- Monitoring frameworks and technologies
- Best practices
- Monitoring tools

Brought to you by

solarwinds 

**Leon Adato
Kong Yang
Brad Hale**



About SolarWinds

SolarWinds provides powerful and affordable IT operations management software to more than 150,000 customers worldwide — from Fortune 500 enterprises to small businesses. SolarWinds' products are downloadable, easy to use and maintain, and provide the power, scale, and flexibility needed to manage today's IT environments. SolarWinds' growing online community, thwack, offers users problem-solving and technology-sharing for all of SolarWinds' products. This active user-community input is combined with decades of IT management experience to deliver a wide range of solutions and tools to address the real-world needs of IT professionals.

Network Monitoring

FOR
DUMMIES[®]
A Wiley Brand

SolarWinds Special Edition

**by Leon Adato, Kong Yang,
and Brad Hale**

FOR
DUMMIES[®]
A Wiley Brand

Network Monitoring For Dummies®, SolarWinds Special Edition

Published by
John Wiley & Sons, Inc.
111 River St.
Hoboken, NJ 07030-5774
www.wiley.com

Copyright © 2016 by John Wiley & Sons, Inc.

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. SolarWinds and the SolarWinds logo are registered trademarks of SolarWinds. 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: THE PUBLISHER AND THE AUTHOR 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 WARRANTIES OF FITNESS FOR A PARTICULAR PURPOSE. NO WARRANTY MAY BE CREATED OR EXTENDED BY SALES OR PROMOTIONAL MATERIALS. THE ADVICE AND STRATEGIES CONTAINED HEREIN MAY NOT BE SUITABLE FOR EVERY SITUATION. THIS WORK IS SOLD WITH THE UNDERSTANDING THAT THE PUBLISHER IS NOT ENGAGED IN RENDERING LEGAL, ACCOUNTING, OR OTHER PROFESSIONAL SERVICES. IF PROFESSIONAL ASSISTANCE IS REQUIRED, THE SERVICES OF A COMPETENT PROFESSIONAL PERSON SHOULD BE SOUGHT. NEITHER THE PUBLISHER NOR THE AUTHOR SHALL BE LIABLE FOR DAMAGES ARISING HEREFROM. THE FACT THAT AN ORGANIZATION OR WEBSITE IS REFERRED TO IN THIS WORK AS A CITATION AND/OR A POTENTIAL SOURCE OF FURTHER INFORMATION DOES NOT MEAN THAT THE AUTHOR OR THE PUBLISHER ENDORSES THE INFORMATION THE ORGANIZATION OR WEBSITE MAY PROVIDE OR RECOMMENDATIONS IT MAY MAKE. FURTHER, READERS SHOULD BE AWARE THAT INTERNET WEBSITES LISTED IN THIS WORK MAY HAVE CHANGED OR DISAPPEARED BETWEEN WHEN THIS WORK WAS WRITTEN AND WHEN IT IS READ.

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 [Branded Rights&Licenses@Wiley.com](mailto:BrandedRights&Licenses@Wiley.com).

ISBN: 978-1-119-27427-8 (pbk); ISBN: 978-1-119-27428-5 (ebk)

Manufactured in the United States of America

10 9 8 7 6 5 4 3 2 1

Publisher's Acknowledgments

Some of the people who helped bring this book to market include the following:

Project Editor: Carrie A. Johnson

Production Editor: Selvakumaran
Rajendiran

Editorial Manager: Rev Mengle

Business Development Representative:
Kimberley Schumacker

Introduction

Picture this scenario: You get to your desk at 9:00 a.m. sharp, having had a great morning workout, followed by a shower, a fantastic cup of coffee, and a frustration-free drive to the office. You're fresh and focused and ready to make a serious dent in that growing to-do list, which includes curious items like users complaining that "the Internet" gets really slow every so often, and the CFO thinks we're overpaying for WAN bandwidth. How much are we using?

Logging on to your PC, you notice that no emails have come in overnight. "That's odd," you're thinking. Seeing you arrive, your buddy now walks over and says, "Looks like something's wrong with email." You log on to the email server and find out that it's . . . well, you don't actually log on to the email server. The remote desktop won't make a connection. You try ping-
ing the box, and there's no response. You wonder to yourself if the problem is in the network or somewhere else in the system. With a sinking feeling, you make the long journey to the computer room. All hope of working on your to-do list is now gone as you stab a finger at the server's power switch. A few moments later, you're logged on at the console. A pop-up alert on the screen tells you that one of the drives is completely full.

Much . . . (much!) later in the day, a picture forms of what happened. Sometime during the night (2:30 a.m. to be exact) the data drive filled up, causing mail services to stop. Shortly after that, errors on the system drive reached a critical point, and the entire system crashed. Meanwhile, in the heat of fighting this fire, you didn't dig deeper to note that the data drive has been hovering at 95 percent capacity for over a week. And the drive that contains the operating system has been throwing read/write errors every 15 minutes for the last 17 days.

About this time, your manager, who's been keeping a respectful distance while you worked, lets you know that the CEO is back from his contract discussions overseas. During the flight home, the CEO needed to send some follow-up documentation

to the customer. When the corporate email wasn't responding, he resorted to creating a professional-sounding Gmail account and sent the files from there. The three of you are scheduled to sit down and debrief the situation in 30 minutes. You start to pull some notes together for what you predict will be an uncomfortable conversation. Well, it *was* going to be a great day.

About This Book

The situation you read above may be a typical one for you in your Information Technology (IT) monitoring scope. If you can relate, then this book is for you! *Network Monitoring For Dummies, SolarWinds Special Edition*, provides an introduction to IT monitoring for someone who is familiar with IT in general but not with monitoring as a discipline. As such, (almost) no former knowledge or experience is required before delving into the chapters of this book. If you already have experience with monitoring, this may not be the book for you. But then again, couldn't we all use a refresher? It couldn't hurt.

We have attempted to make this book tool-agnostic. The purpose of this book is to give you a basic understanding of why you need monitoring, what the monitoring tools are, and some best practices of networking monitoring.

Icons Used in This Book

This book uses the following icons to call your attention to information you may find helpful in particular ways.



The information marked by this icon gives you certain details that are important to remember. This way, you can easily spot noteworthy information when you refer to the book later.



This icon points out extra-helpful information, including ways to save time, money, and headaches.



Paragraphs marked with the Warning icon call attention to common pitfalls that you may encounter.

Chapter 1

Monitoring as a Discipline

.....

In This Chapter

- ▶ Discovering that monitoring isn't side work
 - ▶ Seeing the benefits of network monitoring
 - ▶ Understanding how an effective monitoring solution is built
-

“**M**onitoring as a discipline” means devoting your focus as an IT professional to ensuring your network, servers, applications, and so on are all stable, healthy, and running at peak efficiency. It means not just being able to tell that a system has crashed, but more importantly to tell when a system *will* crash, and intervening so the crash is avoided.

This chapter gives you insight into monitoring as a discipline, the benefits of monitoring, and the difference between monitoring and managing.

It's Not Work You Do "On the Side"

About a decade ago, there were no InfoSec professionals, no “white hat hackers,” no pen testers. Network security, such as it was, was typically handled by a network or server admin who was drawn to security issues, who had an interest, and who felt passionately about keeping his or her environment safe. Ten years later, no company would think of excluding information security from the list of must-have in-house expertise.

We believe that the same is happening for monitoring professionals. Currently, many IT shops run without any

significant monitoring solution. Others go about it in a piecemeal fashion, allowing teams or even individuals to deploy solutions with no thought to interoperability, scalability, or standards.

But in the not-so-distant future, we imagine a world where the idea of having a monitoring team is as natural as the teams of network, server, virtualization, storage — and yes, security — administrators we have today.

To get to that future, people who are drawn to monitoring, who have an interest and a passion for it, need the information to get up to speed on common terms, concepts, and techniques, and then they need the tools to turn that knowledge into results. This book is dedicated to imparting knowledge and experience gleaned from years of focus on building up our monitoring expertise, and from thousands of engagements with customers who had the same goal as you do.

Looking at the Benefits of Network Monitoring

If you've worked in IT for more than 15 minutes, you know that systems crash unexpectedly, users make bizarre claims about how the Internet is slow, and managers request statistics, which leaves you scratching your head wondering how to collect them in a way that's meaningful and doesn't consign you to the headache of hitting Refresh and spending half the day writing down numbers on a piece of scratch paper just to get a baseline for a report.

The answer to all these challenges (and many, many more) lies in effectively monitoring your environment, collecting statistics, and/or checking for error conditions so you can act or report effectively when needed. This goes well beyond a passive "make sure everything is green" approach to one that includes resource optimization, performance optimization, and proactive prevention and remediation.



Industry studies peg the cost of downtime in the hundreds of thousands of dollars per hour so the benefits of monitoring are indisputable:

- ✓ Improved operational efficiency and reduced costs
- ✓ Improved time-to-resolution and reduced downtime
- ✓ More efficient use of resources

Building an Effective Monitoring Solution

Attaining the benefits of monitoring (see the preceding section) is easier said than done. Saying “let’s monitor our IT environment” presumes that you know what you should be looking for, how to find it, and how to get it without impacting the system you’re monitoring. You’re also expected to know where to store the values, what thresholds indicate a problem situation, and how to let people know about a problem in a timely fashion.

Yes, having the right tool for the job is more than half the battle. But, it’s not the whole battle, and it’s not even where the skirmish started.

To build an effective monitoring solution, the true starting point is learning the underlying concepts. You have to know what monitoring is before you can set up what monitoring does.



Network monitoring is the phrase used to describe the practice of continuously monitoring the network and providing notifications to an administrator (probably you if you’re reading this book) when an element of the network fails. Monitoring is usually performed by software or hardware tools and doesn’t have an effect on the operation or condition of the network. Monitoring can be performed passively or actively:

monitor

[mon-i-ter]

verb

to observe, record, or detect (an operation or condition) with instruments that have no effect upon the operation or condition

6

Network Monitoring For Dummies, SolarWinds Special Edition ____

This is in contrast to management in which the administrator governs or controls the environment:

manage

[man-ij]

verb

to handle, direct govern or control in action or use

Chapter 2

Monitoring 101

In This Chapter

- ▶ Looking into monitoring basics
- ▶ Knowing monitoring technologies
- ▶ Diving deeper into monitoring your network

Every monitoring system, regardless of the vendor or packaging, utilizes basic monitoring principles and technologies. This chapter lays out those core techniques and then gives you a deeper look into monitoring your network.

Defining the Monitoring Basics



A few fundamental aspects of a monitoring system exist across the board, no matter what software you use, or the protocol, or the technique. These basic technologies used for monitoring include the following:

- ✔ **Element:** An *element* is a single aspect of the device you're monitoring, which returns one or more pieces of information.
- ✔ **Acquisition:** How you get information is another key concept. This process is called *acquisition*. Does your monitoring routine wait for the device to send you a status update (push), or does it proactively go out and poll the device (pull)?
- ✔ **Frequency:** Closely tied to acquisition (see the preceding section) is how often information comes back — aptly named *frequency*. Does the device send a “heartbeat” every few minutes? Does it send only data when there’s a problem?

- ✔ **Data retention:** Monitoring, by its very nature, is data-intensive. Whether the acquisition method is push or pull, those statistics typically have to go somewhere and they pile up pretty quickly. At its simplest level, *data retention* is a Yes or No option. Either the statistic is 1) collected, evaluated, acted on, and then forgotten, or 2) data is kept in a data store.
- ✔ **Threshold:** One of the core principals of monitoring is that you collect a statistic and see if it has crossed a line of some kind. It can be a simple line (is the server on or off?), or it can be more complex. Regardless, that line, which is crossed, is called a *threshold*.
- ✔ **Reset:** *Reset* is the logical opposite of threshold. It marks the point where a device is considered “back to normal.”
- ✔ **Response:** What happens when a threshold is breached? Response defines that aspect. A *response* could be to send an email, play a sound file, or run a predefined script.
- ✔ **Requester:** With all the talk about monitoring, little has been said (yet) about where the monitoring is occurring — meaning, from what point in the environment are the monitoring statistics being requested. In its simplest terms, you have two choices: either a piece of software running on the monitored device itself (for example, an agent), or some location outside of the monitored device (agentless).

Monitoring Technologies

Regardless of what monitoring vendors will have you believe, a finite and limited number of technologies can be used to monitor. Where the sophistication comes in is with the frequency, aggregation, the relevance of displays, the ease of implementation, and other aspects of packaging.

Ping

Ping sends out a packet to the target device, which (if it's up and running) sends an “I'm here” type response. The result of a ping tells you whether the device is responding at all (up) and how fast it responded.

SNMP

Simple Network Management Protocol (SNMP) has a few pieces that combine to provide a powerful monitoring solution. SNMP is comprised of a list of elements that return data on a particular device. It could be CPU or the average bits per second transmitted in the last five minutes. SNMP provides data based on either a Trap trigger (when one of the internal data points crosses a threshold) or an SNMP poll request.

ICMP

The Internet Control Message Protocol (ICMP) is used by network devices like routers and switches to send error messages indicating that a host isn't reachable along with some diagnostics.

Syslog

Syslog messages are similar to SNMP traps. A syslog service or agent takes events that occur on the device and sends them to a remote listening system (Syslog destination server).

Log file

An application or process writes messages to a plain text file on the device. The monitoring piece of that comes in the form of something that reads the file and looks for trigger phrases or words.

Event log

Event log monitoring is specific to Windows. By default, most messages about system, security, and (standard Windows) applications events are written here. Event log monitors watch the Windows event log for some combination of EventID, category, and so on, and perform an action when a match is found.

Performance monitor counters

Performance monitor (or PerfMon) counters are another Windows-specific monitoring option that can reveal a great deal of information, both about errors on a system and ongoing performance statistics.

WMI

Windows Management Instrumentation (WMI) is a scripting language built into the Windows operating system that focuses on collecting and reporting information about the target system.

Script

Running a script to collect information can be as simple or complicated as the author chooses to make it. In addition, the script might be run locally by an agent on the same device and report the result to an external system. Or, it might run remotely with elevated privileges.

IP SLA

Internet Protocol Service Level Agreements (IP SLAs) are a pretty comprehensive set of capabilities built into Cisco equipment (and others nowadays, as they jump on the bandwagon). These capabilities are all focused on ensuring the WAN, and more specifically VoIP, environment is healthy by using the devices that are part of the network infrastructure instead of requiring you to set up separate devices to run tests.

Flow

Standard monitoring can tell you that the WAN interface on your router is passing 1.4 Mbps of traffic. But who is using that traffic? What kind of data is being passed? Is it all HTTP, FTP, or something else? *Flow* (most commonly referred to as *NetFlow*) monitoring answers those questions. It sets up the information in terms of conversations and monitors who, what, and how network traffic is being used.

Chapter 3

Going beyond Monitoring Basics

In This Chapter

- ▶ Looking into device availability, fault, and performance
 - ▶ Working with traffic and bandwidth
 - ▶ Understanding WAN and IP address monitoring
-

Monitoring your network allows you to be alerted to possible pot holes before your users hit them at top speed. In this chapter, we provide insight into monitoring your network.

Device Availability, Fault, and Performance

In most modern network monitoring systems, devices are monitored for the following:

- ✓ Availability (is the device reachable?)
- ✓ Faults (detection, isolation, correction, and logging of network events)
- ✓ Performance (efficiency of the network, including throughput, utilization, error rates, and response time)



Monitoring here relies primarily on SNMP and ICMP with more advanced monitoring taking advantage of packet inspection. Some of the key metrics that you should look at include response time and packet loss, CPU load and memory utilization, and hardware health details.

Traffic and Bandwidth

Understanding how network bandwidth is being used is critical in ensuring the availability and performance of business services. Bandwidth and traffic usage are most often monitored using the Flow (most commonly referred to as *NetFlow*) technology that is built into most routers by looking at “conversations” between devices.

When monitoring traffic and bandwidth, pay attention to

- ✓ Interface utilization
- ✓ Applications, users, and protocols generating traffic (who and what are generating traffic)
- ✓ Endpoints (where traffic is coming from and going to)
- ✓ Conversations (who is talking to whom)

WAN

You may not own the WAN between your sites and remote locations and can't directly monitor the fault, availability, and performance of the devices within the WAN. If that's the case, you can use a technology such as IP SLA to generate synthetic traffic or operations to measure the performance between two locations or devices, determining the performance of the WAN.



IP SLA is especially beneficial when monitoring applications that are particularly sensitive to delay, jitter, or packet loss such as VoIP or video streaming.

IP Address Monitoring

A network can have thousands of IP addresses in use at any given time. A duplicate IP assignment, exhausted subnet or DHCP scope, or misconfigured DHCP or DNS service will cause a network fault.



Look for a solution that monitors these IP resources and that can proactively alert you of problems to help you plan for orderly expansion.

Discovering the Different Monitoring Tools

After all is said and done, you still need to buy or build a tool or set of tools that help you monitor all the elements of the IT stack. This can be done with discrete specialized tools that monitor a specific element (for example, network monitoring, storage monitoring, virtualization monitoring, and so on) or with a fully integrated suite of products that provides a common platform across the entire stack. Each approach has its advantages and disadvantages.

Regardless of which approach you choose, all software vendors are selling solutions that work from the same basic playbook. What should you look for as a differentiating factor? What is it, exactly, that makes brand X so much better than brand Y? The answer has as much to do with you and your organization as it does with how monitoring gets done.

Will your monitoring team be one person who is also your server team and network team and helpdesk team and database team? If so, you probably need a tool that sacrifices comprehensive options for simplicity and manageability. Does your organization need absolute flexibility so that the monitoring solution is the one-stop-shop for all your needs? You will pay more, and require more staff, but at the end of the day (or month, or more likely year) you will have a software suite that fits you like a glove.

With all of that said, the nontechnical items you should consider include the following:

- ✓ **Cost to purchase and install:** This includes hardware requirements and the specific needs for your environment. Do you need a separate system to monitor devices in your firewall and/or remote sites? How many monitoring systems do you need for all the devices in your company? And so on.
- ✓ **Ongoing maintenance cost:** These include license costs in year two and beyond.

14 Network Monitoring For Dummies, SolarWinds Special Edition ____

- ✔ **Support requirements:** How many people are needed to maintain the system? This is one of those questions that you should *never* trust the vendor to answer. Talk to some other companies that are using the software.
- ✔ **How much customization is needed?** Again, talking to other companies is extremely useful here.

To learn more about SolarWinds network monitoring solutions, visit www.solarwinds.com.

Chapter 4

Getting to Know the SolarWinds Framework: DART

In This Chapter

- ▶ Finding out what's going on
- ▶ Knowing when something breaks
- ▶ Fixing the problem
- ▶ Pinpointing the root cause of the problem

This chapter offers practical advice that helps you do your job every day. To that end, we now introduce the SolarWinds framework, DART:

- ✔ Discovery
- ✔ Alerting
- ✔ Remediation
- ✔ Troubleshooting

Discovery

Discovery is finding out what's going on. This simple principle should guide you in understanding the health and risks of your network assets. Discovery begins by establishing a point-in-time baseline for the health and risks of your network. Once you understand what's in your network and each component's health and risks, you should look at addressing changes that occur in the network.

Discovery serves three key functions:

- ✔ Identifies all your assets and resources and shows their connected context
- ✔ Provides a point-in-time baseline for the performance and risk of your network
- ✔ Populates the data used to calculate how efficient your network implementation is

Alerting

Alerting is finding a simple way to know when something breaks. The essence of this skill is to ensure that you're not constantly in front of a monitor because, frankly, no one has time for that. The noise should be filtered from the signal such that only the most important information is presented to you. The information that's highlighted should allow you to take corrective actions on a much narrower problem set. As you gain more experience, you'll be more adept at creating more meaningful alerts to bypass even more noise.



To truly appreciate the importance of alerting, you have to understand the pain that comes with incorrect alerting. The data from the network, systems, VMs, and applications being monitored can provide valuable insights into the ecosystem, but that data can easily overwhelm the admin. A constant stream of false alarms and data noise can result in paralysis by over-analyzing thresholds. Suffice it to say, when it comes to alerting, more isn't always better.



In addition to cutting through the metrics noise and data deluge, alerting serves two other critical functions:

- ✔ Records that a particular event has occurred, or a threshold has been reached or exceeded
- ✔ Triggers a notification to an admin for that given event

Alerting provides the first clues that an event is about to happen, is happening, or has happened. It guides the first steps on the path toward troubleshooting and remediating an event.

The alert life cycle spans three primary stages:

- ✔ **Alert creation** means deciding on key health and performance indicators and setting thresholds for those indicators.
- ✔ **Alert handling and routing** necessitates creating a meaningful notification in response to the alert trigger, and communicating that alert to the right person who can take the proper action to prevent or resolve the issue. These notifications can include emails, SMS messages, or automated calls to cellphones.
- ✔ **Alert feedback** involves being able to update alerts based on changes or trigger conditions to ensure the right balance of notification to false alarms.

Remediation

Remediation is fixing the problem. The core principle is to get the network in working order as fast as possible. For an IT admin, this is a race against time. Every minute an application or system is down equates to lost opportunity, and often, lost revenue.

As our Head Geek Thomas LaRock so eloquently stated, “As an IT administrator, you get paid for performance, but you keep your job with recovery.” Your job is on the line when stuff happens. So, when stuff happens, you must take a deep breath and repeat these three magic words: Stop. Drop. Roll. Yes, these are the same steps you take if you’re on fire. They work for IT fires as well:

- ✔ Stop
 - Assess the situation.
 - Focus on the steps that will lead to resolution.
- ✔ Drop
 - Drop all distractions like unnecessary and unconnected services and processes.
 - Remove all unnecessary pseudo-IT chefs from the kitchen. This means anyone not directly responsible for and connected to the stack that you’re trying to restore.

✓ Roll

- Roll out your recovery plan to get your network systems back in working order.
- Monitor key performance indicators to make sure everything is stable following the fix.

Troubleshooting

In *troubleshooting*, you're finding the root cause. Your focus should be on solving the right problem instead of chasing false positives or waterfalls. Being able to quickly uncover the root of a problem translates to being able to remediate it that much faster.

The basic troubleshooting flow to use in any situation includes the following steps:

- 1. Define the problem.**
- 2. Gather and analyze relevant information.**
- 3. Construct a hypothesis on the probable cause for the failure or incident.**
- 4. Devise a plan to resolve the problem based on the hypothesis.**
- 5. Implement the plan.**
- 6. Observe the results of the implementation.**
- 7. Repeat Steps 2–6.**
- 8. Document the solution.**

Follow these eight steps and you can troubleshoot anything.

Chapter 5

Ten Best Practices of Network Monitoring

In This Chapter

- ▶ Setting a baseline for behavior, thresholds, and alerts
 - ▶ Monitoring a delta
 - ▶ Defining parent-child relationships
 - ▶ Looking into event correlation
-



You can spend all the money in the world on fancy monitoring solutions, but if you don't follow some key best practices of network monitoring (we give you the top ten), your effort is bound to fall short of expectations — both yours and the people depending on you.

Baseline Behavior

To be able to identify potential problems even before users start complaining, you need to be aware of what's normal. Baseline behavior over a couple of weeks or even months will help you understand what normal behavior is.

Perform a Network Inventory

Keep an inventory of the network devices, ports, and interfaces being used for network connections, network hardware (links, network controllers, power supply, and so on), servers, virtual machines, and SAN devices.

Avoid Alert Flapping

Alerts help you monitor proactively. Most alerts are automatic email notifications when particular metrics thresholds are crossed. For each alert, you can set critical and warning threshold values. These threshold values are meant to be boundary values that when crossed indicate that the system is in an undesirable state.

When an alert repeatedly triggers (a device that keeps rebooting itself, a disk drive that hovers on the edge of full, and processes keep deleting/creating temporary files so that one moment it's over threshold, the next it's below), that condition is known as *flapping* or *sawtoothing*.



Here are a few techniques that can be used to avoid this, depending on the toolset:

- ✓ Suppress events within a window.
- ✓ Add a delay before triggering.
- ✓ Do not cut a new alert until the original has reset.
- ✓ Two-way communicate with a ticket or alert management system.

Never, Ever Set an Email Filter for Your Alerts



If you find yourself setting up a rule in email to filter or even delete alerts, you're admitting you've failed to set up the correct alert. You're bound to ignore critical alerts.

Monitoring a Delta

In some cases, such as when monitoring disk space, you may not be interested in a specific numeric threshold (alert when the disk is more than 90 percent utilized). Instead, in some cases, what you want to monitor is the delta, or the rate of change. Here, you might be interested in knowing that disk utilization has gone up by more than *xx* percent over *yy* minutes, which may indicate a spike in consumption.

Provide the Details

It is not enough to set and generate an alert if there are not enough details to begin troubleshooting. The more in-depth the alert, the faster the troubleshooting. In addition to the obvious ones (name and IP of the affected device, time of failure, statistic of the failed component at the time of failure, and so on), some other items that might come in handy include the following:

- ✔ The device that *detected* the fault
- ✔ The threshold which was breached (name and value)
- ✔ The name of the alert
- ✔ How long the problem has been in effect (duration)
- ✔ A link or other reference to the place where the current state of this element can be viewed

Escalation

One of the reasons why potential issues become an actual problem is because the alerts triggered based on a threshold are ignored or the right person isn't alerted. When setting up monitoring and reporting, the organization should have a policy on who has to be alerted when a malfunction occurs or a potential problem is detected. Based on the policy, the right person who administers the aspect that is having an issue can be alerted.

Parent-Child

Parent-child relationships (which typically have to be set up manually for each set of devices) are a way of telling the monitoring system what's connected and how. This way, when a parent device is down (the router is the parent of the switch, the switch is the parent of the server), any alerts related to the child devices are suppressed. When the router is down, the switch isn't (necessarily) down. It may be simply unreachable.

In more sophisticated tools, the software might actively check the upstream parent before marking a device as down, and continue all the way up the chain until it finds the highest level device that is down. This is known as *upstream verification* (or conversely, *downstream suppression*).

Event Correlation

Event correlation is a big topic. It's much bigger than one section of this document can accommodate, but it's important enough to merit a brief discussion.

Event correlation tools can perform the flap detection and suppression, as well as parent-child correlation. In addition, event correlation tools might perform the following:

- ✔ On event *X*, look for *Y*.
- ✔ On event *X*, wait *YY* minutes, and look for event *Z*.
- ✔ De-duplication.
- ✔ Alert after *XX* times.

If a problem occurs once, it's negligible, but repeated occurrences indicate the presence of a problem.

View Traffic from the Application Viewpoint

Your users don't care that a switch went down or a routing change was made; they only care that an application is performing poorly or failing. You need to view network traffic and performance and their impacts on the application. This can be accomplished using techniques such as packet inspection.

Monitoring is a discipline

See how you can become a disciple. IT's not glamorous, and few of you are getting rich off IT, but without IT, business would come to a screeching halt. Executive management visibility usually happens when things go wrong, so this book helps you steer clear of trouble and provides a primer on network monitoring to help you understand just why proactive monitoring is critical to the success of your business.

- **Monitor as a discipline** — see the benefits of monitoring and the difference between monitoring and managing
- **Learn frameworks and technologies** — understand the technologies and frameworks used to monitor networks
- **Find out which monitoring tools work for you** — gain a deeper insight into monitoring your network
- **Discover network monitoring best practices** — meet the expectations of everyone in your organization



Open the book and find:

- Why you should monitor
- Monitoring frameworks and technologies
- Beyond network monitoring basics
- Monitoring tools
- Information on the IT stack
- Best practices of network monitoring

Go to **Dummies.com**[®]
for videos, step-by-step examples,
how-to articles, or to shop!

WILEY END USER LICENSE AGREEMENT

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