

NAGIOS

**Building Enterprise-Grade
Monitoring Infrastructures
for Systems and Networks**

David Josephsen

Foreword by Ethan Galstad, creator of Nagios

FREE SAMPLE CHAPTER

SHARE WITH OTHERS



NAGIOS

This page intentionally left blank

NAGIOS

Building Enterprise-Grade Monitoring Infrastructure for Systems and Networks

Second Edition

David Josephsen



PRENTICE
HALL

Upper Saddle River, NJ • Boston • Indianapolis • San Francisco
New York • Toronto • Montreal • London • Munich • Paris
Madrid • Cape Town • Sydney • Tokyo • Singapore • Mexico City

Many of the designations used by manufacturers and sellers to distinguish their products are claimed as trademarks. Where those designations appear in this book, and the publisher was aware of a trademark claim, the designations have been printed with initial capital letters or in all capitals.

The author and publisher have taken care in the preparation of this book, but make no expressed or implied warranty of any kind and assume no responsibility for errors or omissions. No liability is assumed for incidental or consequential damages in connection with or arising out of the use of the information or programs contained herein.

The publisher offers excellent discounts on this book when ordered in quantity for bulk purchases or special sales, which may include electronic versions and/or custom covers and content particular to your business, training goals, marketing focus, and branding interests. For more information, please contact:

U.S. Corporate and Government Sales
(800) 382-3419
corpsales@pearsontechgroup.com

For sales outside the United States, please contact:

International Sales
international@pearsoned.com

Visit us on the Web: informit.com/aw

Library of Congress Cataloging-in-Publication Data is on file.

Copyright © 2013 Pearson Education, Inc.

All rights reserved. Printed in the United States of America. This publication is protected by copyright, and permission must be obtained from the publisher prior to any prohibited reproduction, storage in a retrieval system, or transmission in any form or by any means, electronic, mechanical, photocopying, recording, or likewise. To obtain permission to use material from this work, please submit a written request to Pearson Education, Inc., Permissions Department, One Lake Street, Upper Saddle River, New Jersey 07458, or you may fax your request to (201) 236-3290.

ISBN-13: 978-0-13-313573-2

ISBN-10: 0-13-313573-X

Text printed in the United States on recycled paper at R.R. Donnelley in Crawfordsville, Indiana.

First Printing: April 2013

*For Cynthia, for enduring and encouraging my incessant curiosity.
And for Tito, the cat with the biggest heart.*

This page intentionally left blank

CONTENTS

Foreword by the Nagios Creator, Ethan Galstad **xiii**

Introduction **I**

Do It Right the First Time	1
Why Nagios?	2
What's in This Book?	4
Who Should Read This Book?	7
End Notes	7

CHAPTER 1 **Best Practices** **9**

A Procedural Approach to Systems Monitoring	9
Processing and Overhead	12
Remote Versus Local Processing	12
Bandwidth Considerations	13
Network Location and Dependencies	14
Security	16
Silence Is Golden	19
Watching Ports Versus Watching Applications	20
Who's Watching the Watchers?	21
End Notes	22

CHAPTER 2 **Theory of Operations** **23**

The Host and Service Paradigm	24
Starting from Scratch	24
Hosts and Services	26
Interdependence	26
The Downside of Hosts and Services	27
Plug-ins	28
Exit Codes	28
Remote Execution	31

Scheduling	34
Check Interval and States	34
Distributing the Load	36
Reapers and Parallel Execution	38
Notification	39
Global Gotchas	39
Notification Options	40
Templates	41
Time Periods	41
Scheduled Downtime, Acknowledgments, and Escalations	42
I/O Interfaces Summarized	43
The Web Interface	43
Monitoring	45
Reporting	46
The External Command File	48
Performance Data	48
The Event Broker	49
End Notes	50

CHAPTER 3 Installing Nagios 51

OS Support and the FHS	51
Installation Steps and Prerequisites	53
Installing Nagios	54
Configure	54
Make	55
Make Install	56
Installing the Plug-ins	57
Installing NRPE	59
End Notes	60

CHAPTER 4 Configuring Nagios 61

Objects and Definitions	62
nagios.cfg	64
The CGI Config	67
Templates	68
Timeperiods	70
Commands	71
Contacts	73
Contactgroup	74

Hosts	75
Services	77
Hostgroups	79
Servicegroups	79
Escalations	80
Dependencies	81
Extended Information	83
Apache Configuration	83
GO!	85
End Notes	85

CHAPTER 5 Bootstrapping the Nagios Config Files 87

Scripting Templates	87
Autodiscovery	91
Check_MK	91
Nagios XI	92
Autodiscovery Is Dead: Long Live Autodiscovery	92
NagiosQL	92

CHAPTER 6 Watching: Monitoring Through the Nagios Plug-ins 95

Local Queries	95
Pings	96
Port Queries	98
Querying Multiple Ports	100
(More) Complex Service Checks	102
E2E Monitoring with WebInject and Cucumber-Nagios	104
Watching Windows	111
The Windows Scripting Environment	111
COM and OLE	113
WMI	113
To WSH or Not to WSH	118
To VB or Not to VB	119
The Future of Windows Scripting	119
Getting Down to Business	121
NRPE	122
Check_NT	123
NSCP	124

Data Visualization Strategies: A Tale of Three Networks	185
Suitcorp: Nagios, NagiosGraph, and Ddraw	185
singularitygov: Nagios and Ganglia	192
Massive Ginormic: Nagios, Logsurfer, Graphite, and Life After RRDTool	200
DIY Dashboards	209
Know What You're Doing	210
RRDTool Fetch Mode	212
The GD Graphics Library	214
NagVis	215
GraphViz	217
Sparklines	218
Force Directed Graphs with jsvis	220
End Notes	221

CHAPTER 9 Nagios XI 223

What Is It?	223
How Does It Work?	224
What's in It for Me?	226
One Slick Interface	226
Integrated Time Series Data	227
Modularized Components	228
Enhanced Reporting and Advanced Visualization	228
Integrated Plug-ins and Configuration Wizards	230
Operational Improvements	234
How Do I Get My Hands on It?	235

CHAPTER 10 The Nagios Event Broker Interface 237

Function References and Callbacks in C	237
The NEB Architecture	239
Implementing a Filesystem Interface Using NEB	242
DNX, a Real-World Example	255
Wrap Up	258
End Notes	259

Index 261

This page intentionally left blank

FOREWORD

People often say that Nagios is “flexible,” by which I think they mean that it is easily extended, but that misses the point. The power inherent in Nagios’ design derives not from its extensibility, but rather from its insistence on being extended. This is an admittedly small but important distinction. Many pieces of software can be extended to do new things, but very few pieces of software do nothing until you’ve extended them, and it is exactly because of this—this inherent demand that you customize it to suit your needs—that Nagios has always been a synthesis of contributions from engineers and administrators working to solve their own individual problems. No two installations are alike, and that is by design.

In the years since I first created Nagios, it has grown in breadth and scope beyond anything I’d imagined. With over 1 million users worldwide, Nagios Core has found a home everywhere from huge Fortune-500 conglomerates to state of the art scientific research labs. The Nagios user community is one of the healthiest and most actively contributing open source communities out there, with nearly 4,000 published plug-ins, add-ons, and extensions—many of which are sufficiently complex to warrant books of their own. The community is so large, diverse, and active, that Nagios now has its own annual conference where contributors, users, and educators come together to share ideas, learn tips and tricks, and find out about upcoming developments in the project.

There is also a thriving community of corporations at work on extending and supporting Nagios. In 2007 I joined them, founding Nagios Enterprises. Our flagship product, Nagios XI, is both an evolutionary step forward, and (as it should be) a fully-reverse compatible extension to Nagios Core. XI embraces the extend-by-design lineage of Core, preserving the power and flexibility of Core, while expanding its accessibility and usability.

But even given the wonderful success Nagios has enjoyed, I’m the first to admit that flexibility comes with a price. It can be difficult for newcomers and experienced admins alike to build and deploy a successful monitoring solution, and many of the challenges have nothing whatsoever to do with computers. Luckily, David is one of the few technical writers that are able to cover a complex subject like this in an easy-to-understand format. Whether you’re a newcomer to the world of network, system, and IT monitoring, or you’re an experienced Nagios admin, David’s work is sure to be helpful to you.

—Ethan Galstad, Nagios Founder and President

This page intentionally left blank

ACKNOWLEDGMENTS

My lovely wife, Cynthia, is patient and encouraging and pretty, and I love her.

Ethan Galstad, whose interest prompted the second edition, and without whom there would be no Nagios.

The tech reviewers on this project were outstanding—thanks, guys.

Last, my editors at Prentice Hall have been great. They aren't at all like the editors in Spiderman or Fletch. Debra Williams Cauley and Kim Boedigheimer are a hardworking, on the ball, and clued-in pair of professionals. They've been patient and helpful, and I appreciate their time and attention.

Thanks.

This page intentionally left blank

ABOUT THE AUTHOR

David Josephsen is the Director of Systems Engineering at DBG, Inc., where he maintains a collection of geographically dispersed server farms. He has more than a decade of hands-on experience with UNIX systems, routers, firewalls, and load balancers in support of complex, high-volume networks. In addition to this book, he authored several chapters in the O'Reilly book *Monitoring with Ganglia*, and currently writes “iVoyer,” the systems monitoring column for *;login* magazine. Josephsen is just one of many thousands of avid Nagios users.

This page intentionally left blank

ABOUT THE TECHNICAL REVIEWERS

Mark Bainter

Mark Bainter leads a team of sysadmins providing outsourced monitoring and management of high volume mail systems for Message Systems' clients, leveraging over 15 years experience as a sysadmin specializing in systems integration, monitoring, and automation. He is an autodidactic polymath and impenitent sesquipedalian. Mark currently resides in Texas with his lovely wife and four children and in his free time he enjoys reading, wood-working, and losing at Settlers to his wife.

Mike Guthrie

Mike Guthrie is the lead developer at Nagios enterprises and has developed new features and add-ons for Nagios Core, Nagios XI, and Nagios Fusion. Mike does the bulk of his programming in PHP and particularly enjoys front-end web development and data visualizations. When he's not at work, he enjoys spending time with his family, being outside, and working on his house.

Mathias Kettner

Mathias Kettner is known as the author of Check_MK, MK Livestatus, and other Nagios add-ons. He runs a fast growing company in Munich, Germany, which is dedicated to system monitoring based on Nagios, and offers professional support and software development.

This page intentionally left blank

Introduction

This is a book about untrustworthy machines—machines, in fact, that are every bit as untrustworthy as they are critical to our well being. But I don't need to bore you with a laundry list of how prevalent computer systems have become or with horror stories about what can happen when they fail. If you picked up this book, I'm sure you're well aware of the problems: layer upon layer of interdependent libraries hiding bugs in their abstraction, script kiddies, viruses, DDOS attacks, hardware failure, end-user error, backhoes, hurricanes, and on and on. It doesn't matter whether the root cause is malicious or accidental; your systems will fail, and when they do, only two things will save you from the downtime: redundancy and monitoring systems.

Do It Right the First Time

In concept, monitoring systems are simple: an extra system or collection of systems whose job is to watch the other systems for problems. For example, the monitoring system could periodically connect to a Web server to make sure it responds and, if not, send notifications to the administrators. Although it sounds straightforward, monitoring systems have grown into expensive, complex pieces of software. Many now have agents larger than 500MB, include proprietary scripting languages, and sport price tags above \$60,000.

When implemented correctly, a monitoring system can be your best friend. It can notify administrators of glitches before they become crises, help architects tease out patterns corresponding to chronic interoperability issues, and give engineers detailed capacity planning information. A good monitoring system will help the security guys correlate interesting events, show the network operations center personnel where the bandwidth bottlenecks are, and provide management with much needed high-level visibility into the critical systems that they bet their business on. A good monitoring system can help you uphold your service level

agreement (SLA) and even take steps to solve problems without waking anyone up at all. Good monitoring systems save money, bring stability to complex environments, and make everyone happy.

When done poorly, however, the same system can wreak havoc. Bad monitoring systems cry wolf at all hours of the night so often that nobody pays attention anymore; they install backdoors into your otherwise secure infrastructure, leech time and resources away from other projects, and congest network links with megabyte upon megabyte of health checks. Bad monitoring systems can really suck.

Unfortunately, getting it right the first time isn't as easy as you might think, and in my experience, a bad monitoring system doesn't usually survive long enough to be fixed. Bad monitoring systems are too much of a burden on everyone involved, including the systems being monitored. In this context, it's easy to see why large corporations and governments employ full-time monitoring specialists and purchase software with six-figure price tags. They know how important it is to get it right the first time.

Small- to medium-sized businesses and universities can have environments as complex as or even more complex than large companies, but they obviously don't have the luxury of high-priced tools and specialized expertise. Getting a well-built monitoring infrastructure in these environments, with their geographically dispersed campuses and satellite offices, can be a challenge. But having spent a good part of the past 13 years building and maintaining monitoring systems, I'm here to tell you that not only is it possible to get it done right the first time, but you can do it for free, with a bit of elbow grease, some open source tools, and a pinch of imagination.

Why Nagios?

Nagios is, in my opinion, the best system and network monitoring tool available, open source or otherwise. Its modularity and straightforward approach to monitoring make it easy to work with and highly scalable. Further, Nagios's open source license makes it freely available and easy to extend to meet your specific needs. Instead of trying to do everything for you, Nagios excels at interoperability with other open source tools, which makes it very flexible. If you're looking for a monolithic piece of software with check boxes that solve all your problems, this probably isn't the book for you. But before you stop reading, give me another paragraph or two to convince you that the check boxes aren't really what you're looking for.

Most commercial offerings get it wrong because their approach to the problem assumes that everyone wants the same solution. To a certain extent, this is true. Everyone has a large glob of computers and network equipment and wants to be notified if some subset of it fails.

So if you want to sell monitoring software, the obvious way to go about it is to create a piece of software that knows how to monitor every conceivable piece of computer software and networking gear in existence. The more gadgets your system can monitor, the more people you can sell it to. To someone who wants to sell monitoring software, it's easy to believe that monitoring systems are turnkey solutions and whoever's software can monitor the largest number of gadgets wins.

The large commercial packages I've worked with all seem to follow this logic. Not unlike the Borg, they are methodically locating new computer gizmos and adding the requisite monitoring code to their solution—or worse, acquiring other companies who already know how to monitor lots of computer gadgetry and bolting those companies' code onto their own. They quickly become obsessed with features, creating enormous spreadsheets of supported gizmos. Their software engineers exist so that the presales engineers can come to your office and say to your managers, through seemingly layers of white gleaming teeth, “Yes, our software can monitor that.”

The problem is that monitoring systems are not turnkey solutions. They require a large amount of customization before they start solving problems and herein lies the difference between people selling monitoring software and those designing and implementing monitoring systems. When you're trying to build a monitoring system, a piece of software that can monitor every gadget in the world by clicking a check box is not as useful to you as one that makes it easy to monitor what you need, in exactly the manner that you want. By focusing on *what* to monitor, the proprietary solutions neglect the *how*, which limits the context in which they may be used.

Take ping, for example. Every monitoring system I've ever dealt with uses ICMP Echo requests, otherwise known as pings, to check host availability in one way or another. But if you want to control *how* a proprietary monitoring system uses ping, architectural limitations become quickly apparent. Let's say I want to specify the number of ICMP packets to send, or I want to be able to send notifications based on the round-trip time of the packet in microseconds instead of simple pass/fail. More complex environments may necessitate that I use IPv6 pings, or that I portknock¹ before I ping. The problem with the monolithic, feature-full approach is that these changes represent changes to the core application logic and are, therefore, nontrivial to implement.

In the commercial monitoring applications I've worked with, if these ping examples could be performed at all, they would require reimplementing the ping logic in the monitoring system's proprietary scripting language. In other words, you would have to toss out the built-in ping functionality altogether. Perhaps being able to control the specifics of ping checks is of questionable value to you, but if you don't have any control over something as basic as ping, what are the odds that you'll have finite enough control over the most important checks

in your environment? They've made the assumption that they know *how* you want to ping things and from then on it was game over; they never thought about it again. And why would they? The ping feature is already in the spreadsheet, after all.

When it comes to gizmos, Nagios's focus is on modularity. Single-purpose monitoring applets called plug-ins provide support for specific devices and services. Rather than participating in the feature arms race, hardware support is community driven. As community members have a need to monitor new devices or services, new plug-ins are written and usually more quickly than the commercial applications can add the same support. In practice, Nagios will always support everything you need it to and without ever needing to upgrade Nagios itself. Nagios also provides the best of both worlds when it comes to support, with several commercial options, as well as a thriving and helpful community that provides free support through various forums and mailing lists.

Choosing Nagios as your monitoring platform means that your monitoring effort will be limited by your own imagination, technical prowess, and political savvy. Nagios can go anywhere you want it to and the trip there is usually pretty simple. Although Nagios can do everything the commercial applications can, and more, without the bulky, insecure agent install, it usually doesn't compare favorably to commercial monitoring systems because when spreadsheets are parsed, Nagios doesn't have as many checks. In fact, if they're counting correctly, Nagios has no checks at all, because technically it doesn't know *how* to monitor anything; it prefers that you tell it how. The question of "how" is difficult to encompass with a check box.

What's in This Book?

Although Nagios is the biggest piece of the puzzle, it's only one of the myriad of tools that make up a world-class open source monitoring system. With several books, superb online documentation, and lively and informative mailing lists, it's also the best-documented piece of the puzzle. So my intention in writing this book is to pick up where the documentation leaves off. This is not a book about Nagios as much as it is a book about the construction of monitoring systems using Nagios, and there is much more to building monitoring systems than configuring a monitoring tool.

I'll cover the usual configuration boilerplate, but configuring and installing Nagios is not my primary focus. Instead, to help you build great monitoring systems, I need to introduce you to the protocols and tools that enhance Nagios's functionality and simplify its configuration. I need to give you an in-depth understanding of the inner workings of Nagios itself, so you can extend it to do whatever you might need. I need to spend some time in this book exploring possibilities because Nagios is limited only by what you feel it can do.

Finally, I need to write about things only loosely related to Nagios, like best practices, SNMP, visualizing time-series data, and various Microsoft scripting technologies, such as WMI and WSH.

Most important, I need to document Nagios itself in a different way. By introducing it in terms of a task-efficient scheduling and notification engine, I can keep things simple while talking about the internals up front. Rather than relegating important information to the seldom-read advanced section, I'll empower you early by covering topics like plug-in customization and scheduling as core concepts.

Although the chapters more or less stand on their own, and I've tried to make the book as reference-friendly as possible, I think it reads better as a progression from start to finish. I encourage you to read from cover to cover, skipping over anything you are already familiar with. The text is not large, but I think you'll find it dense with information and even the most seasoned monitoring veterans should find more than a few useful nuggets of wisdom.

The chapters tend to build on each other and casually introduce Nagios-specific details in the context of more general monitoring concepts. Because many important decisions need to be made before any software is installed, I begin with "Best Practices" in Chapter 1. This should get you thinking in terms of what needs to take place for your monitoring initiative to be successful, such as how to go about implementing, who to involve, and what pitfalls to avoid.

Chapter 2, "Theory of Operations," builds on Chapter 1's general design guidance by providing a theoretical overview of Nagios from the ground up. Rather than inundating you with configuration minutiae, Chapter 2 will give you a detailed understanding of how Nagios works without being overly specific about configuration directives. This knowledge will go a long way toward making configuration more transparent later.

Before we can configure Nagios to monitor our environment, we need to install it. Chapter 3, "Installing Nagios," should help you install Nagios, either from source or via a package manager.

Chapter 4, "Configuring Nagios," is the dreaded configuration chapter. Configuring Nagios for the first time is not something most people consider to be fun, but I hope I've kept it as painless as possible by taking a bottom-up approach, documenting only the most used and required directives, providing up front examples, and specifying exactly what objects refer to what other objects and how.

Most people who try Nagios become attached to it² and are loathe to use anything else. But if there is a universal complaint, it is certainly configuration. Chapter 5, “Bootstrapping the Nagios Config Files,” takes a bit of a digression to document some of the tools available to make configuration easier to stomach. These include automated discovery tools, as well as graphical user interfaces.

In Chapter 6, “Watching: Monitoring Through the Nagios Plug-ins,” we are finally ready to get into the nitty-gritty of watching systems, including specific examples with Nagios plug-in configuration syntax solving real-world problems. I begin with a section on watching Microsoft Windows boxes, followed by a section on UNIX, and ending with the “other stuff” section, which encompasses networking gear and environmental sensors.

Chapter 7, “Scaling Nagios,” is new to the second edition. Scaling Nagios for large networks has been one of the most interesting problems Nagios sysadmins have had to deal with over the past five or six years. The explosion of machine virtualization and cost-effective cloud services have created a lot of interest in large parallel processing architectures that are composed of lots of little nodes. In this chapter, I cover several tools and strategies that will enable you to distribute the monitoring load and build a stable large-scale monitoring infrastructure for tens of thousands of nodes and beyond.

Chapter 8, “Visualization,” covers one of my favorite topics: data visualization. Good data visualization solves problems that couldn’t be solved otherwise, and I’m excited about the options that exist now, as well as what’s on the horizon. With fantastic visualization tools like RRDTool, Ganglia, and Graphite, graphing time series data from Nagios is getting easier every day, but this chapter doesn’t stop at mere line graphs.

Also new in the second edition is Chapter 9, “Nagios XI,” which is dedicated to the new commercial version of Nagios. Built from many of the tools covered in this book by the guys who originally wrote Nagios, XI is truly a masterpiece of integration and usability. They’ve made monitoring with Nagios so simple my mom could do it (well, my mom writes optimizing cross-compilers for embedded FLIR systems, but you get my point).

And finally, now that you know the rules, it’s time to teach you how to break them. At the time of this writing, Chapter 10, “The Nagios Event Broker Interface,” is the only print documentation I’m aware of that covers the new Nagios Event Broker interface. The event broker is the most powerful Nagios interface available. Mastering it rewards you with nothing less than the ability to rewrite Chapter 2 for yourself by fundamentally changing any aspect of how Nagios operates or extending it to meet any need you might have. I describe how the event broker works and walk you through building an NEB module.

Who Should Read This Book?

If you are a systems administrator with a closet full of UNIX and Windows systems and assorted network gadgetry, and you need a world-class monitoring system on the cheap, this book is for you. Contrary to what you might expect, building monitoring systems is not a trivial undertaking. Constructing the system that potentially interacts with every TCP-based device in your environment requires a bit of knowledge on your part. But don't let that give you pause; systems monitoring has taught me more than anything else I've done in my career and, in my experience, no matter what your level of knowledge, working with monitoring systems has a tendency to constantly challenge your assumptions, deepen your understanding, and keep you right on the edge of what you know.

To get the most out of this book, you should have a pretty good handle on the text-based Internet protocols that you use regularly, such as SMTP and HTTP. Although it interacts with Windows servers very well, the Nagios daemon is meant to run on Linux, which makes the text pretty Linux heavy, so a passing familiarity with Linux or POSIX-ish systems is helpful. Although not strictly required, you should also have some programming skills. The book has a fair number of code listings, but I've tried to keep them as straightforward and as easy-to-follow as possible. With the exception of Chapter 8, which is exclusively C, the code listings are written in either UNIX shell or Perl.

Perhaps the only strict requirement is that you approach the subject matter with a healthy dose of open curiosity. If something seems unclear, don't be discouraged; check out the online documentation, ask on the lists, or even shoot me an email; I'd be glad to help if I can.

Have fun!

—Dave

End Notes

¹ www.portknocking.org

² Dare I say, love it?

This page intentionally left blank

Nagios XI

In 2009, Nagios Enterprises, the corporation formed by Nagios creator Ethan Galstad, launched Nagios XI, a commercial version of Nagios. XI truly is an amazing accomplishment. You need to know next to nothing to use it, and yet the first eight chapters of this book are prerequisite to your understanding it. But now that you have a good handle on how Nagios and the various add-ons surrounding it work, we can finally examine XI and see if it might be a good fit for you.

What Is It?

After the release of 3.0, Nagios was, it seemed, in danger of becoming a victim of its own success. Sysadmins who knew and loved it were happy to see it continue in the way it always had, but its popularity had risen to the point that a different and more populous group of potential end users had taken notice, and with them, Nagios wasn't comparing favorably with newer, prettier, and less flexible commercial competitors.

This new breed of user was quite vocal and had a few very specific gripes. First they found Nagios's configuration syntax unwieldy, to say nothing of the intolerable notion of (gasp) editing text files by hand. Second, they found the Nagios web interface, with its C-based CGI and lack of integrated time-series data, unforgivably old-fashioned. Finally they had no idea what to make of the fact that there was no database back-end. Jiminy Christmas—wrist watches and garbage disposals run MySQL these days! How was one to take seriously a monitoring system that didn't?

For this considerable subset of users, Nagios's price tag didn't make up for its abhorrent lack of bling, and answers to the effect that all these things could be rectified with add-ons fell on deaf Bluetooth earpieces. Add-on options were birds in the bush, and they would rather pay for a bird in the hand than go beating around the bush themselves for free.

XI might best be called the perfect compromise between maintaining the power and flexibility of Nagios and providing a turnkey monitoring system that more than satiates the desires of the PHP proletariat. But that description sells it short; XI is much more than just a shiny interface; it represents a huge amount of custom development and integration work. Further, there is real functionality in XI that simply can't be found in Nagios Core. But neither can it be called a new monitoring system in its own right, because in very important ways, it remains Nagios and retains all the flexibility and power that I've described in the previous chapters. Everything I've written up to now about the underlying architecture, plug-ins, scalability, and even advanced visualization, is applicable to Nagios XI.

It'll be easier to just show you.

How Does It Work?

Figure 9.1 is a rough sketch of the Nagios XI architecture. As you can see, all host and service monitoring, as well as notification, escalation, and so on, relies on an unmodified Nagios Core daemon, so any preexisting plug-ins or customization you might have can be made to work under XI. The NDOUtils plug-in (described in Chapter 7, "Scaling Nagios") has been enabled and configured to replicate state information from Nagios Core into a MySQL database. Here is the primary information hand-off between Core and XI; Nagios XI reads this database to glean information about the current state of hosts and services, as well as the Core daemon itself. This adds an information layer to Core that can be consumed by third-party UIs as well as your own custom integration scripts.

The NagiosQL add-on (described in Chapter 5, "Bootstrapping the Nagios Config Files") provides the hooks necessary to modify the Nagios Core configuration from the XI interface. Every parameter that can be configured in the flat files may be set via the web interface using the customized NagiosQL forms in the "Advanced Configuration" section of the XI interface. Although these forms are well integrated into XI, and retain an XI look and feel, there is a bit of a line in the sand between NagiosQL-driven core configuration, which is referred to as "advanced" in the XI interface, and the configuration parameters that are specific to XI itself.




Figure 9.1 The Nagios XI architecture, simplified

XI goes beyond presenting a simple web wrapper to the Nagios Core configuration files, providing in addition a litany of semiautomated wizards and autodiscovery tools to ease the burden of initial and ongoing host and service configuration. I talk more about these later, but suffice to say that it is the intention of the XI creators to isolate the majority of XI users from the intricacies of the Nagios core configuration to the extent that they never need to know what a check command is, much less a template. This makes it possible for monitoring configuration, traditionally an operations task, to be delegated to first-level support types, or in some environments, even to normal users. More clueful administrators who need to customize this or that can still do so, without editing the config files by hand, using the NagiosQL-driven advanced configuration tool.

Configuration created by NagiosQL is automatically written to text configuration files in `etc/nagios` and is read by the Core daemon from these flat files in the usual fashion. Although it's technically possible to hand edit these configuration files, you will gain nothing

because NagiosQL will eventually overwrite any changes you make. If you have your own configuration-generating automation (like `Check_MK`), or preexisting configuration that you do not want to import into NagiosQL, or even if you're a curmudgeon who just prefers to manually edit the configuration, you can still maintain static config files in `etc/nagios/static`, and your files will still be parsed by the Core daemon while being left alone by NagiosQL. That runs both ways; statically configured hosts and services can't be modified via the UI unless you manually import them into NagiosQL (at which point they cease to be static).

Finally, Nagios XI maintains its own Postgresql database to store various configuration parameters such as user-settings, custom dashboards, authentication info, and the like. Given the shiny new PHP interface, the simplified configuration options, and the open database back-end, Nagios XI should satisfy the complaints I'm used to hearing from corporate administrators who are in the market for a "grown up" commercial monitoring product; however, there's a lot more functionality than what I've encompassed in the architecture diagram.

What's in It for Me?

Now that we've taken a quick look at what XI is and how it works, let's take a look at how XI compares to Nagios Core and the various commercial monitoring systems with which it was designed to compete.

One Slick Interface

Given the general quality of the alternative PHP interfaces we find in the Nagios Exchange repository, the XI interface is shockingly excellent. It is certainly not yet another effort to bring the CGI interface "up to date" by replacing it with a PHP version of itself. The XI user interface is a complete rethinking of the UI, which truly takes advantage of the strengths of a web programming platform like PHP at every opportunity. Elements within dashboards can be unlocked, moved around, or even deleted to suit the preferences of the user. AJAX is employed, both to update individual information elements and to provide feedback, so that when I send a command via the UI to reschedule a service check or acknowledge an alert, a box momentarily appears to let me know my command has been accepted. One of my least favorite things about the Core UI is the way it dumps me to an acknowledgment page after I've issued a command, forcing me to manually navigate back to somewhere useful.

The traditional Nagios tables like "service detail" and "hostgroup grid" still exist, but are implemented as repurposable widgets that I can use to build custom dashboards. New tables have been added, a few of which are very dense and handy, like the "minemap" visualization pictured in Figure 9.2.

CPU and provides a snappy, feature-rich data visualization, allowing you to scale the graph by dragging to select a range and providing pop-up numerical values when you mouse over any data areas. The stacked time series graphs include time-shifted historical data, so you can easily compare today's data to that of yesterday, and so on.

Modularized Components

The UI as a whole is highly modular, incorporating add-on components to implement extra features. This enables the XI developers to quickly react to the needs of the user community by adding features to the UI as needed or even adding custom developing features for larger end users with special needs. A notable example is the Operations screen depicted in Figure 9.3, which is intended to be displayed on a dedicated screen in a Network Operations Center. In addition to this and other single-page summaries, custom views can be configured to rotate between pages with more detailed information on timed intervals. I bring up these little summary views because seeing them so prominently displayed in the XI interface hits home for both the extent to which the Nagios developers are listening to the needs of the community and their eagerness to satisfy those needs now that incremental progress in the UI is possible.

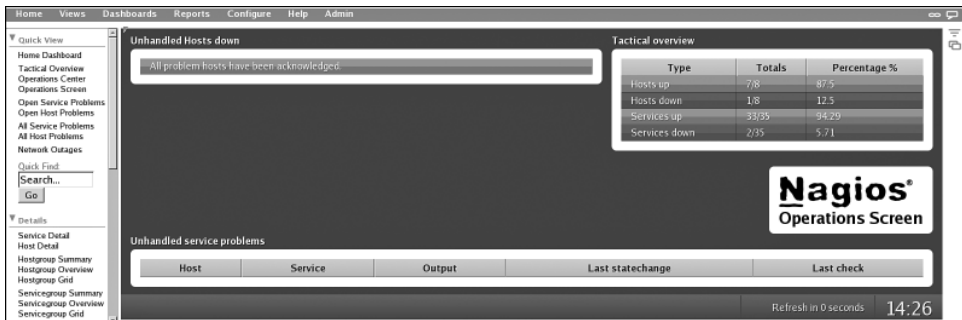


Figure 9.3 Nagios XI Operations screen

Finally! Acknowledgments and Scheduled Downtime for Multiple Hosts

Another component that implements a feature for which the core community has been begging for years is the Mass Acknowledgment Component. This allows an admin to schedule downtime and acknowledge problems for groups of hosts and services. I know more than one sysadmin who would purchase XI for this feature alone.

Enhanced Reporting and Advanced Visualization

The XI developers are not solely focused on the community, however, as a quick glance at the Reporting tab in XI shows; they are proactively exploring some interesting data visualization

techniques from the neoformix data-visualization field. Components that implement heat maps, force directed graphs, and stream graphs, as depicted in Figure 9.4, have been added to the classic reporting options. Several shiny new implementations of the core reports are also provided, each of which I find generally cleaner than their legacy counterparts and more likely to impress the wearers of neckties and high heels in our lives. The new reports may be exported in CSV and PDF formats with the click of a button. The button, which links to a predictable URL, makes it possible for the shorts and t-shirt wearers among us to automatically grab the reports with tools like curl and wget.

From: 2012-08-17 16:46:56 To: 2012-08-18 16:46:56

The alert stream provides a visual representation of host and service alerts over time.

Clicking on a host name will cause the graph to drill down to show service alerts for that particular host.

Nagios Neoformix

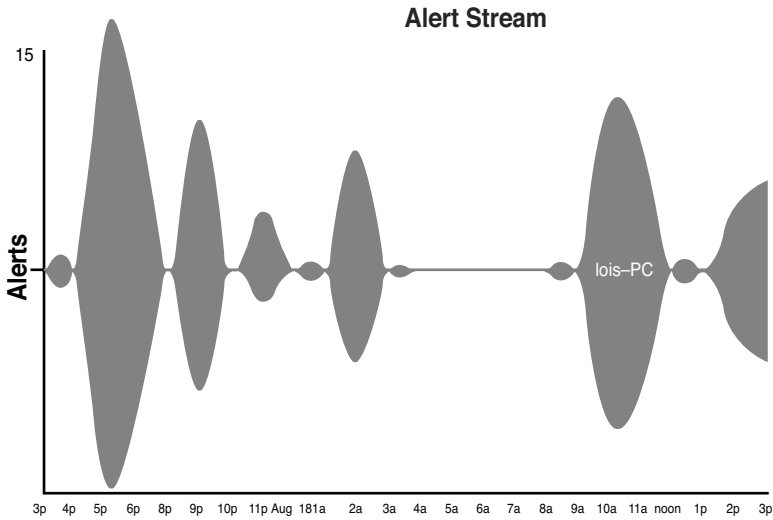


Figure 9.4 Nagios XI Stream Graph component

Nagvis

Nagvis, (described in Chapter 8) is installed and available in the Maps section of the Home view. Setting up your own NagVis diagrams couldn't be easier. First, copy your map or diagram graphic to /usr/local/nagvis/share/userfiles/images/maps, launch the Nagvis tool in the XI UI, select Manage Maps from the options menu, and create a new map, pointing the Background to the map you uploaded. Finally, open your map using the Open menu, and add status icons to it by selecting Add Icon from the Map menu.

Business Processes

Nagios XI contains wrapper logic for grouping individual services into higher-level entities called business processes. The intent here is to implement what the Gardiner Group calls Business Application Monitoring, or BAM. BAM attempts to provide real-time status for critical business entities like a sales catalog web site or corporate email. Nagios XI implements BAM by breaking a high-level concept like “corporate email,” into its requisite pieces, such as Mail Transfer Agents, Mail Exchangers, Groupware systems, and Databases, and then quantifying the relative importance of each of the services that make up those pieces as well as describing dependency relationships between them.

XI Business Process groups contain services that are said to be “essential” or “non-essential.” A database service in our example might be considered essential, whereas the SMTP port on a single mail exchanger might be “non-essential” (because they are usually redundant, and even if they go down, the mail will queue somewhere else). When any essential service or the combination of all non-essential services goes critical, the XI business process logic registers this as a “problem.”

Each business process group contains critical and warning thresholds that depend on the number of problems that are occurring in the group. In our example, we might imagine two business process groups, one for SMTP speakers (MXs and MTAs) and one for SQL-speakers (groupware systems and DBs). If the latter group registers a single problem because a database is down, that might throw the whole group into a warning state.

Business process groups can contain other nested business process groups, and so on. Our top-level entity, corporate email, is therefore just a business process group that contains the two groups previously described. It is configured like the other two groups so that a single “problem” in any of the nested groups causes it to go into a warning state. Finally, notification commands can be assigned on each business process group in the same way they are assigned to individual host and service events. Additionally, visualization widgets exist for the top-level groups. These can be added to any dashboard or view, and they allow the user to drill down into the groups to see what services or subgroups constitute them.

Integrated Plug-ins and Configuration Wizards

The core installation of Nagios XI includes all the plug-ins in the standard plug-ins package, as well as NRPE, NSCA, and NRDP. In addition to all the plug-ins being preinstalled, the XI developers have provided a plethora of semiautomated configuration wizards, which, given the bare-minimum information about a host, take care of the initial setup as well as adding and modifying services on already configured hosts.

If you consult the official XI documentation at

<http://library.nagios.com/library/products/nagiosxi/documentation>,

you'll quickly discover that the wizards are the preferred method for host and service configuration. With names like Exchange Server, website, and Windows Workstation, they make setting up new hosts and services easy enough that these tasks can be delegated to first-level support techs, or even end users. The autodiscovery wizard is capable of bootstrapping an environment given only a CIDR netblock to start with, and it does a good job of initial setup. To add NRPE-based host checks or other services after the fact, run the appropriate wizard on the preexisting host.

For example, if Server1 was created with the autodiscovery wizard, and you now want to add NRPE checks to get CPU, memory, and disk information from the host, you must first install NRPE on Server1. If Server1 doesn't already have NRPE on it, and is one of several common server types, such as a Windows 200X server, Red Hat, or Ubuntu, the XI developers have an agent package designed to work with XI specifically at:

<http://assets.nagios.com/downloads/nagiosxi/wizards>

After the agent is installed on Server1, run the NRPE Wizard on the server from the configuration tab of the XI user interface, as shown in Figure 9.5, entering the IP or FQDN of the server, and choosing the type from the drop-down list. The wizard will then display a preconfigured subset of available check commands relevant to your server type, and provide text-entry fields for you to specify custom settings or additional commands if you'd like.

As I said earlier, static configuration files may still be maintained in `etc/nagios/static`. So it's entirely possible to run your own scripts, or autogeneration tools like those included with `check_mk`, provided you configure them to write their configuration to the static directory. I can't deny that the automated configuration features in XI have, perhaps ironically, complicated things a bit for those of us who have reason to maintain the configuration manually. In the Nagios Core universe, there is a single way to configure Nagios (text files). However, there are three ways to configure Nagios Core in the XI universe (text files, NagiosQL, and XI Wizards), and although the three coexist well enough, it can become burdensome to ensure a uniformity of parameters if the administrators mix and match their configuration methodologies in XI. I'll give you an example.

Larry, his brother Darryl, and his other brother Darryl all work at Bloody Stump Lumber Mill, where they recently purchased a Nagios XI server to monitor their growing sales web-application server farm. Larry was a UNIX admin in college, so he prefers to edit the config files. Darryl likes to have fine-grained control over the config, but isn't very good

in vim, so he uses the XI advanced configuration section, and other Darryl would rather be watching football, so he just runs the wizard for everything. Each of the brothers has a server running sshd that he wants to configure in XI.

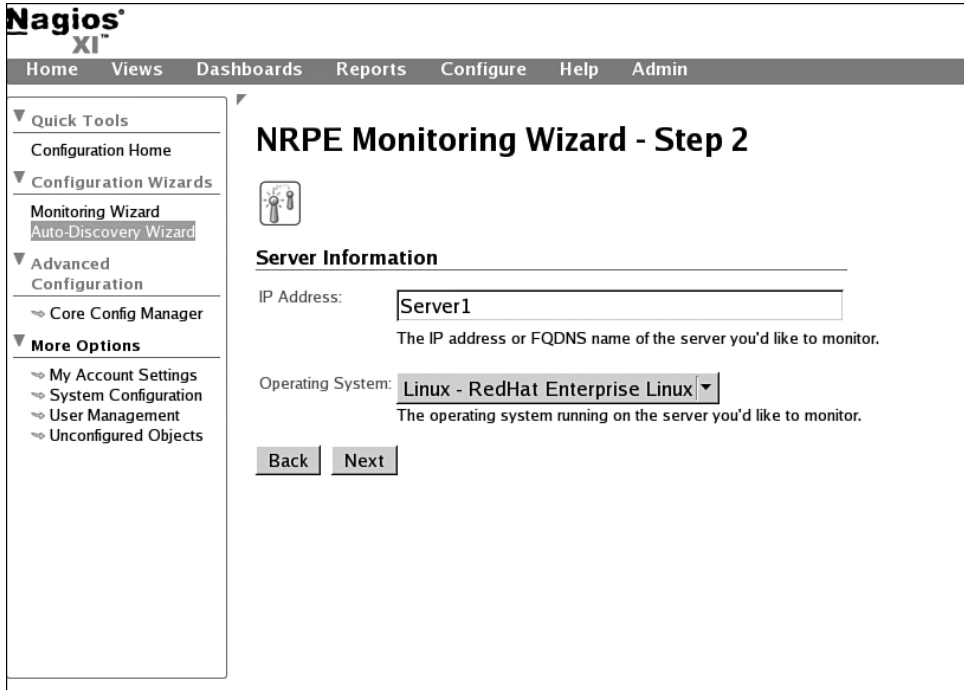


Figure 9.5 The Nagios XI NRPE Wizard

When other Darryl runs the Autodiscovery Wizard on his server's IP, XI scans the host and automatically configures a host check and a `check_tcp` service check for the SSH port. It then pushes the config to NagiosQL, which commits it to the DB, writes out the configuration, and restarts the daemon.

Darryl meanwhile, sets up his host using the NagiosQL forms directly, but instead of choosing `check_tcp`, he chooses the `check_ssh` service, which does pretty much the same thing, but returns slightly different output. He also names the service "ssh" instead of "SSH" like the wizard does.

Larry, meanwhile, has really done his homework. He already has a servicegroup for ssh servers in the static config files he created, so rather than doing all the typing and clicking that his brothers do, he simply adds his server to the `ssh_servers` servicegroup, and the rest

takes care of itself. The problem is, his servicegroup inherits a different set of templates than NagiosQL, so although his service check uses the same name and check command as the wizard, his polling interval is different, and he has a different notification target for service warnings.

In this way, the brothers end up with three different definitions for the same service, which might not be a problem immediately, but will cause all manner of headaches if and when they want to integrate Nagios with another tool, or generally try to do any sort of automation using their monitoring server.

I admit these sorts of disconnects are possible with text configuration files, but my point is the text configuration encourages administrators to use templates to normalize the configuration, like Larry did in the previous example. The automated tools by comparison encourage isolating the configuration at the host level, because it's easier for the automated tools to parse them that way. Thus, in Larry's configuration, we find a single `services.cfg` wherein every service is defined and assigned a hostgroup, whereas in NagiosQL's configuration, we find a services directory with a single file for each host. The former makes it pretty easy to verify that all the service checks for every host are implemented in the same way. The later makes it much more difficult.

Further, in my experience, the disdain that people like Larry naturally feel for people like other Darryl generally discourages them from paying close attention to what people like other Darryl are doing. In fact, merely inviting other Darryl to configure the monitoring server with wizards might trigger a tendency in Larry to go off on his own and "do it the right way" using well-written static config files, which only exacerbates the problem by more widely diverging the configuration paths.

Whether this will be a problem in your shop will depend on how many hands are stirring the pot and the extent to which the more clueful users are aware of the potential problem. The idea of delegating the configs is certainly tempting, and I'm not saying you shouldn't. If you do, my advice would be to use either the wizards or static config for service and host *creation*, and avoid using NagiosQL directly if you can avoid it (you could still safely use it for host and service *modification*). That way, you can carefully set up the static config to ensure that it references the wizard templates, or simply copy definitions from the NagiosQL files, and everything should remain pretty much uniform.

Automated Configuration for Passive Checks

Another very cool bit of functionality that is related to automated configuration in Nagios XI is the Unconfigured Objects feature. In the event that XI receives a passive check result for a host or service that it doesn't know about, it automatically generates an inert configuration for that host or service and places it in the Unconfigured Objects section of the Configure tab.

Administrators may then approve the inert objects, and they will become part of the running configuration. Good stuff.

Operational Improvements

In addition to the myriad functional improvements in Nagios XI, several maintenance-related features exist that make it easier to manage the Nagios server itself.

Backups

Out of the box, XI takes a snapshot of the running configuration each time it changes. These configuration snapshots can be downloaded from the UI in an automated fashion using tools like curl or wget. It can be used to restore the configuration in the event the monitoring system kicks the bucket, or it can roll it back to a prior version if someone made an inappropriate change. A real system backup, including historical state and metric data, involves a lot more than just the configuration files, however. Remember, XI maintains three databases and has untold amounts of performance data stored in RRDs, not to mention the Nagios Core state file and logs. For detailed instructions on properly backing up your XI install, see:

http://assets.nagios.com/downloads/nagiosxi/docs/Backing_Up_And_Restoring_XI.pdf

User Management

Account management is more important in XI, especially when individual users are encouraged to change configuration parameters and create new hosts and services. Individual users in XI also have the ability to configure the interface with custom views and dashboards as they see fit. For these reasons, XI must track users in its own database rather than leaving it up to Apache to sort out like the Nagios Core UI does. Account management is well done in XI and generally behaves in a manner that enterprise users expect. Access control exists to prevent individual accounts from making modifications, and components exist to enable XI to use LDAP servers. Nagios has published official documentation on multitenant setups, where, for example, access to a Nagios server hosted by a service provider is shared by multiple customers. This documentation resides at:

http://assets.nagios.com/downloads/nagiosxi/docs/XI_Multi-Tenancy.pdf

Daemon Status

As depicted in Figure 9.6, the XI interface provides an array of detailed information about the Core daemon process. This includes metric values for the server hardware as well

as performance metrics internal to the daemon itself. A real-time graph of the event queue displays reaper and service check events scheduled 5 minutes into the future. This really is fantastic capacity planning info of a quality I've never seen in any monitoring system.

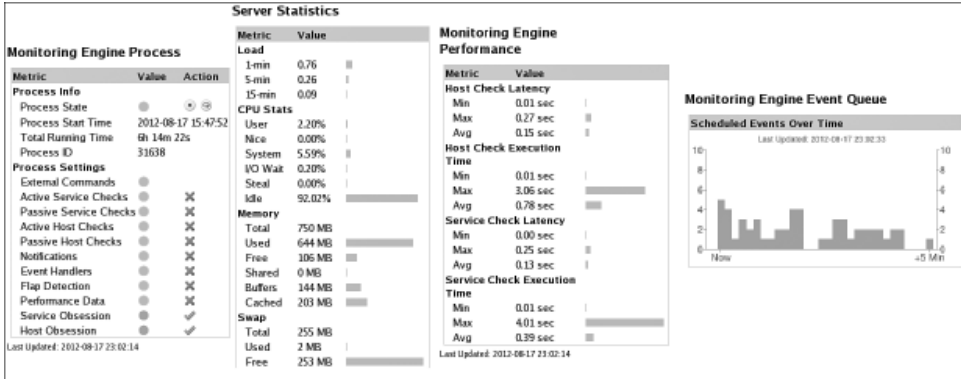


Figure 9.6 Detailed daemon statistics

How Do I Get My Hands on It?

Fully functional demo versions of XI (60-day expiration) are available from nagios.org. You may download self-contained installers, or VMware disk images with XI preinstalled. The latter can be run by any system that supports the free `vmplayer` utility, while the former requires a relatively recent Red Hat or CentOS install.

The reason for the RHEL dependency is the dizzying array of packages that must exist for XI to run. The XI developers have chosen to rely heavily on YUM to satisfy the requisite dependencies, so although it's possible to run XI on other distros, you won't be able to use the official installation script to get it up and running on anything other than a Red Hat or CentOS system.

This page intentionally left blank

Symbols

- .1.3.6.1 prefix, 135
- * (asterisk), 67
- { } (curly braces), 64
- \$ (dollar signs), 96
- . (dot), 116

A

- abnormal utilization, 169
- acknowledgement, notification, 43
- action_url, 198-199
- active_checks_enabled, 150
- address directive, 75
- Afterglow, 218
- alarms, false alarms, 19
- Alert summary, 47
- Apache, configuration, 83-85
- applications versus ports, watching, 20-22
- architecture
 - Event Broker, 239-241
 - Nagios XI, 225
- AREA, RRDTool, 181
- argument passing, command definitions (check_load), 128
- asterisk (*), 67
- authorized_for_all_host_commands, 68
- authorized_for_all_hosts, 68

- authorized_for_all_service_commands, 68
- authorized_for_all_services, 68
- authorized_for_configuration_information, 68
- authorized_for_system_information, 68
- authorized_for_system_commands, 68
- autodiscovery, 91-92
 - Check_MK, 91
 - Nagios XI, 92
- automated configuration for passive checks, 233
- AVERAGE, RRDTool, 180
- averageSeries, 209
- awk, 196

B

- backups, Nagios XI, 234
- bandwidth, monitoring systems, 13-14
- bar charts, data visualization, 210
- baselines, 19
- benefits of Nagios XI
 - advanced reporting and advanced visualization, 228-230
 - integrated plug-ins and configuration wizards, 230-233
 - integrated time series data, 227-228
 - interface, 226-227
 - modularized components, 228
 - operational improvements, 234

bootstrapping Nagios config files, 87
 business processes, Nagios XI, 230

C

- C
 - callbacks, 237-239
 - function references, 237-239
 - c, SNMP, 141
 - callbacks, 157
 - C, 237-239
 - NEB callback types, 240
 - carbon, 202
 - CDEF, 185
 - data summarization, 184
 - RRDTool, 181
 - syntax, 182
 - cgi.cfg, 63, 67-68
 - directives, 68
 - check_cluster, 116
 - check_command directive, 76-78
 - check_disk, 189
 - command definition, 131
 - check_dllhost
 - command definitions, 122
 - service definitions, 123
 - check_dllHost, 114
 - check_host_regix.sh, 197
 - check_http, 99-100, 104
 - check_http service definition, 98-100
 - check_load
 - command definitions with argument passing, 128
 - service definitions, 129
 - Check_MK, 13, 91, 131-134
 - Check_NT, 123-124
 - check_nt_cpuload, 124
 - check_ping, 96, 189
 - check_ping service definition, 97
 - check_snmp, command definitions, 141
 - check_ssl service definition, 104
 - check_swap, command definitions, 130
 - check_tcp, 98-99
 - check_tcp command definition, 99
 - check_tcp wrappers, 101
 - child instance, ping service
 - definitions, 152
 - child/parent relationships, 27
 - CIM (Common Information Model), 114
 - Cisco routers, enabling SNMP, 138
 - COM, 113
 - command, 62
 - command definition, check_ping, 96
 - command definitions, 68
 - check_disk, 131
 - check_dllhost, 122
 - check_load with argument passing, 128
 - check_nt_cpuload, 124
 - check_snmp, 141
 - check_swap, 130
 - check_tcp, 99
 - WebInject, 108
 - command_line, 72
 - command_name, 72
 - command objects, 72
 - commands, 71-73

- compile-time options, 55
- complex service checks, local queries, 102-104
- configuration wizards, Nagios XI, 230-233
- configuring Nagios, 54-55
- config.xml, WebInject, 106
- consolidation functions, RRDTool, 177
- contact, 62
- contactgroup, 62
- contact_groups, 78
- contactgroups, 74-75
- contact objects, 73
- contacts, 73-74
- COUNTER, 171-172
- CPU, UNIX, 126-129
- Cscript, 113
- Cucumber-Nagios, 108-111
 - installing, 110
- cut, 196

D

- daemon, OS support, 51
- daemon status, Nagios XI, 234
- daemon, 51
- dashboards, 209
 - creating, 210-212
 - distributed. *See* distributed dashboards
 - force direct graphs with jsvis, 220-221
 - GD graphics library, 214-215
 - GraphViz, 217-218
 - NagVis, 215-216
 - RRDTool fetch mode, 212-214
 - sparklines, 218-220

- data, sending
 - with Event Broker, 249
 - from Nagios to Ganglia, 194-197
- data polling glitches, heartbeat, 173
- data sources (DS), 171
- data summarization, CDEF, 184
- data visualization, 167, 185
 - dashboards
 - creating*, 210-212
 - force directed graphs with jsvis*, 220-221
 - GD graphics library*, 214-215
 - GraphViz*, 217-218
 - NagVis*, 215-216
 - RRDTool fetch mode*, 212-214
 - sparklines*, 218-220
 - Massive Ginormic, 200-209
 - Nagios XI, 228-230
 - singularity.gov, 192-193
 - displaying graphs from Ganglia in Nagios UI*, 198-200
 - Ganglia*, 193-194
 - monitoring Ganglia metrics using Nagios*, 197-198
 - sending data from Nagios to Ganglia*, 194-197
 - Suitcorp, 185-187
 - drraw*, 190-192
 - NG (NagiosGraph)*, 187-190
- DEF, RRDTool, 180
- default_user_name, 68
- definitions, 64
- definition skeleton, services template, 89

- dependencies, 81-83
 - monitoring systems, 14-16
- dependent_host_name, 82
- DERIVE, 172
- directives
 - address, 75
 - cgi.cfg, 68
 - check_command, 76-78
 - event_handler, 76
 - first_notification, 81
 - hostgroup, 79
 - host_name, 77
 - _interval type, 78
 - last_notification, 81
 - parents, 76
 - service_description, 77
 - servicegroup_members, 80
- disk, UNIX, 130-131
- displaying graphs from Ganglia in Nagios UI, 198-200
- distributed architecture with passive checks, 151
- distributed dashboards, 159-165
 - Fusion, 160-161
 - Livestatus. *See* Livestatus
 - MNTOS (Multi Nagios Tactical Overview System), 160-161
- distributed passive checks with secondary Nagios daemons, 150-153
- distributing loads, scheduling, 36-38
- DNX (Distributed Nagios Executor), 154-155, 255-258
- dnxPluginInit(), 257
- dollar signs (\$), 96
- dot (.), 116
- downtime
 - scheduled downtime, 42
 - scheduling (Nagios XI), 228
- drdraw, 190-192
- DS (data sources), 171
- ds struct, 251-252
- E**
 - e, 99
 - E2E (End to End), 20
 - E2E monitoring, 104
 - Cucumber-Nagios, 108-111
 - WebInject, 105-108
 - ehProcessData, 256
 - ehSvcCheck function, 258
 - EMU, 143
 - End to End (E2E), 20
 - environmental sensors, 142-143
 - escalations, 80-81
 - notification, 42
 - Event Broker, 237
 - architecture, 239-241
 - DNX, 255-258
 - implementing file system interfaces, 242-255
 - I/O interfaces, 49-50
 - event broker modules, 153
 - DNX (Distributed Nagios Executor), 154-155
 - Mod Gearman, 154-157
 - Op5 Merlin, 154, 157-159
 - event_handler directive, 76

event scheduling, 35
execution_failure_criteria, 82
Exit Codes, 28-32
extended information, 83
external command file, I/O interfaces, 48

F

false alarms, 19
fetch mode, RRDTool, 212-214
FHS (File System Hierarchy Standard), 52
file locations, 52
File System Hierarchy Standard (FHS), 52
filesystem interfaces, implementing (Event Broker), 242-255
filter headers, LQL, 165
first_notification directive, 81
force directed graphs with jsvis, 220-221
fPointer, 239
function pointers, 237-239
function references, C, 237-239
functions, Massive Ginormic, 208
Fusion, 160-161

G

Galstad, Ethan, 223
Ganglia, 193-194
 displaying graphs in Nagios UI, 198-200
 monitoring using Nagios, 197-198
 sending data from Nagios to, 194-197
ganglia_service_name, 199
GAUGE, 171-172

GD graphics library, 214-215
GET columns, 163
GET hosts, 163
global enablers, nagios.cfg, 65
global notification settings, 39-40
global timeouts, nagios.cfg, 66
gmetric, 194
Gmond.conf, 194
Graphite, 202-204
graph mode, RRDTool, 180-182
graph.php, 198
graphs
 forced directed graphs, jsvis, 220-221
 from Ganglia displaying in Nagios UI, 198-200
GraphViz, 217-218

H

-H switch, 96
headers, LQL, 163
 OR headers, 166
 stats headers, 166
heartbeat, RRDTool, 172-173
host and service paradigm, 24
 downside of, 27-28
 hosts and services, 26
 interdependence, 26-27
 starting from scratch, 24-25
host definition skeletons, 88
hostdependency, 63
hostescalation, 63
hostextendedinfor, 63
hostgroup, 63, 79

- hostgroups, 79
 - host_name, 82
 - host_name directive, 77
 - hosts, 26, 62, 75-77
 - downside of, 27-28
 - host templates, 69
 - host template skeletons, 88
- I**
- i2c, 142
 - ICMP Echo requests, 3
 - implementing filesystem interfaces (Event Broker), 242-255
 - include statements, 245
 - installing
 - Cucumber-Nagios, 110
 - Nagios, 54
 - configuring*, 54-55
 - make install*, 56-57
 - make targets*, 55-56
 - steps for*, 53-54
 - NRPE, 59-60
 - plug-ins, 57-58
 - integrated plug-ins, Nagios XI, 230-233
 - integrated time series data, Nagios XI, 227-228
 - Intelligent Platform Management Interface (IPMI), 145-146
 - interdependence, host and service paradigm, 26-27
 - interesting events, 20
 - interfaces
 - I/O interfaces
 - Event Broker*, 49-50
 - external command file*, 48
 - monitoring*, 45-46
 - overview*, 43
 - performance data*, 48-49
 - reporting*, 46-47
 - web interfaces*, 43-44
 - Nagios XI, 226-227
 - internal RRDTool metric averaging, 214
 - _interval type directives, 78
 - intervals, scheduling, 34-36
 - I/O interfaces
 - Event Broker, 49-50
 - external command file, 48
 - monitoring, 45-46
 - overview, 43
 - performance data, 48-49
 - reporting, 46-47
 - web interfaces, 43-44
 - IPMI (Intelligent Platform Management Interface), 145-146
- J**
- jsvis, force directed graphs, 220-221
- K**
- Kettner, Mathias, 131, 161
 - Klein, Dan, 143
- L**
- last_notification directive, 81
 - LINE, RRDTool, 181
 - Linux, installing daemon, 59

listings

- Apache Sample VirtualHost Config, 84
- Broker5s smake_callbackm code for SERVICE_STATUS_DATA, 251
- CDEFs for Data Summarization, 184
- CDEF Syntax, 182
- A check_cluster Plug-in in Perl/WMI, 116
- Check_disk Command Definition, 131
- A check_disk Definition for NG, 189
- Check_dllHost, 114
- Check_dllhost Command Definition, 122
- Check_dllhost Service Definition, 123
- Check_http service Definition, 98
- Check_load Command Definition with Argument Passing, 128
- The check_load Service Definition, 129
- Check_nt_cpuload Command Definition, 124
- Check_nt_cpuload Service Definition, 124
- Check_ping Command Definition, 96
- Check_ping Service definition, 97
- The Check_snmp Command Definition, 141
- The check_ssl Service Definition, 104
- Check_swap Command Definition, 130
- A check_tcp Wrappper, 101
- Command Example, 71
- A Command to Perform an SMTP Handshake, 103
- The config.xml for WebInject, 106
- Contact Example, 73-74
- Creating a Multicounter RRD, 178
- Creating a Single-Counter RRD, 175
- A Cucumber Feature File, 109
- dnxPluginInit() Function, 257
- Enabling SNMP on Cisco Routers, 138
- The Event Broker Sending Data, 249
- Fully MIBd snmpwalk Output, 140
- The Generic check_tcp Definition, 98
- The Host Definition Skeleton, 88
- A Host Template and Consumer Definition, 69
- Hostdependency Example, 81
- Hostescalation Example, 80
- Host Example, 75
- Hostextendedinfo Example, 83
- Hostgroup Example, 79
- A Host Template Skeleton, 88
- Includes, 245
- The init Function, 246
- Installing Nagios for the Impatient Person, 54
- Internal RRDTool Metric Averaging, 214
- A List of Boxes, 120
- A List of Hosts, 89
- A Merlin Load-Balanced Peer Configuration, 159
- Modifying RRAs in Nagios Graph, 188
- My qls Script, an Interactive Shell for MK-Livestatus, 163

- The nebmodule struct, 247
- An NEB Module That Implements a Filesystem Interface, 242
- The nebstruct_service_status_data struct, 252
- NGNs check_ping Definition, 189
- A Notification Command Definition, 74
- OCSF Configuration in the nagios.cfg on the Child, 152
- Our Event Handler Function, 250
- Output from sconfigrec, 57
- Output from Plug-ins Oconfigrec, 58
- Output from the rrdtool Fetch Command, 212
- Output from the Sensors Program, 144
- A Performance Data Wrapper for All Plug-ins, 49
- A Ping Plug-in, 30
- Ping Service Definition for the Child (Poller) Instance, 152
- Ping Service Definition for the Parent Instance, 151
- Ping with Summary Output, 30
- The process-service-perfdata Command for Use with NG, 188
- Protocol_Specific check_tcp Command Definition, 99
- A Realistic Nagios Installation, 56
- A Remote Load Average Checker, 31
- A Remote Load Average Checker with Exit Codes, 32
- A Sample Host Definition, 64
- A Script That Calls load-checker and Parrots Its Output and Exit Code, 33
- ServiceDependency Example, 82
- Service Example, 77
- Servicegroup Example, 79
- Servicescalation Example, 80
- A Services Definition Skeleton, 90
- A Services Template for Use with a Definition Skeleton, 89
- The service_struct Def from nagios.h, 252
- A Shell Script to Create a hosts.cfg from the Skeletons and Host List, 89
- A Shell Script to Parse the Output from the fetch Command, 213
- Shiny New check_http Service Definition, 100
- A Solution for Ted, 104
- Some Required Tidbits, 245
- Specifying Object Config Files by Directory, 65
- Specifying Object Config Files Individually, 65
- Step Definition Example, 109
- The submit_service_check.sh Shell Script on the Child, 153
- The submit_service_check_to_parent Definition on the Child, 153
- The Test Case File for WebInject, 106
- Timeperiod Example, 70
- Unrecognizable SNMP Gobbledygook, 138
- Using a Function Pointer, 238
- Verbose Output from WebInject, 107
- A WebInject Command Definition, 108
- A WebInject Service Definition, 108

- Livestatus, 161-163, 166
 - filter headers, 165
 - tables, 162
- Livestatus Query Language (LQL), 162
- lm78 sensor chip, 142
- lm-sensors, 144-145
- local installs, file locations, 52
- local processing versus remote processing, 12-13
- local queries, 95
 - complex service checks, 102-104
 - pings, 96-98
 - port queries, 98-100
 - querying multiple ports, 100-102
- LQL (Livestatus Query Language), 162
 - headers, 163
 - OR headers, 166
 - stats headers, 166

M

- macros, 73
- make cgis, 56
- make contrib, 56
- make install, 56-57
- make install-command mode, 65
- make install-config, 64
- make modules, 56
- make nagios, 56
- make targets, 55-56
- Mass Acknowledgement Component, 228
- Massive Ginormic, 200-209

- max, RRDTool, 174
- max_check_attempts, 78
- max check attempts option, 35
- memory, UNIX, 129-130
- MIBs, SNMP, 140
- Microsoft Visual Basic, script Edition, 112
- Microsoft Windows, installing NRPE, 60
- min, RRDTool, 174
- Minemap, Nagios XI, 227
- minimizing overhead, 14
- MK-Livestatus, 163
- MK-Multisite, 161
- MMCs (Microsoft Management Consoles), 113
- MNTOS (Multi Nagios Tactical Overview System), 160-161
- Mod Gearman, 154
 - event broker modules, 156-157
- modularized components, Nagios XI, 228
- monitoring
 - Ganglia metrics, using Nagios, 197-198
 - I/O interfaces, 45-46
- monitoring systems, 1-2
 - bandwidth considerations, 13-14
 - dependencies, 14-16
 - network locations, 14-16
 - procedural approach to, 9-12
 - processing, remote versus local, 12-13
 - silence, 19-20
- MRTG, 170
- multicounter RRD, 178

Multi Nagios Tactical Overview System (MNTOS), 160-161

Multisite, 166

MX outages, 28

N

Nagios, 2-4

installing, 54

configuring, 54-55

make install, 56-57

make targets, 55-56

steps for, 53-54

nagios.cfg, 62-65

global enablers, 65

global timeouts, 66

nagios.cfg file, 62

Nagios Core, XI, 231

NagiosGraph, 187-190

NagiosQL, 92-94, 225

Nagios Remote Plugin Executor.

See NRPE

Nagios UI, displaying Ganglia graphs, 198-200

Nagios XI, 92, 160, 223-224

architecture, 225

benefits of

enhanced reporting and advanced visualization, 228-230

integrated plug-ins and configuration wizards, 230-233

integrated time series data, 227-228

interface, 226-227

modularized components, 228

operational improvements, 234

business processes, 230

getting access to, 235

how it works, 224-226

Minemap, 227

NagVis, 215-216, 229

name-space collisions, preventing, 73

NANs, 201

Navigation bar, 44

NDOUtils, tuning, 150

NEB_API_VERSION, 245

NEBCALLBACK_PROCESS_DATA, 256

NEB call back types, 240

NEB module, 161

neb_register_callback, 248, 256

nebstruct_service_status_data struct, 252

.NET, 119

Network Interface Card (NIC), 18

network locations, monitoring systems, 14-16

network segments, security, 18

NG, 189

NIC (Network Interface Card), 18

normal_check_interval, 78

notification, 39

acknowledgement, 43

escalation, 42

global, 39-40

scheduled downtime, 42

templates, 41

time periods, 41-42

notification_commands, 74

notification_failure_criteria, 82

- notification_interval, 77-78, 81
- notification options, 40-41
- notification_options, 74, 78
- notification_period, 78
- notifications, service notifications, 78
- NRDP, tuning, 150
- NRPE (Nagios Remote Plugin Executor), 33, 59, 122-123
 - installing, 59-60
 - UNIX, 125
- NRPE-NT, 122-123
- NRPE Wizard, 231
- NSC++, 124
- NSCA, tuning, 150
- NSClient++, 60, 124-125
- NSCP, 124-125

O

- object configuration files, 62
- object definitions, 68
- object definition skeletons, 88
- objects, 62-63
 - summary of, 62
- object template skeletons, 88
- OCSF (Obsessive Compulsive Service Processor), 152
- Oetiker, Tobias, 170
- OIDs, SNMP, 136, 139-141
- OLE, 113
- OLE CPAN, 118
- On switch, SNMP, 139
- Op5 Merlin, 154, 157-159
- OR headers, LQL, 166
- OS support, 51
- overhead, minimizing, 14

P

- parallel executing, scheduling, 38-39
- parent instances, ping service definitions, 151
- parents directive, 76
- passive checks
 - automated configuration, 233
 - distributed passive checks, 150-153
- passive_checks_enabled, 150
- PDPs (primary data points), 174
- performance data, 168
- I/O interfaces, 48-49
- Perl, 111
- pie charts, data visualization, 210
- ping, 3
 - local queries, 96-98
 - service definitions
 - child instance*, 152
 - parent instances*, 151
 - summary output, 30
- PluginInit function, 257
- plug-ins
 - check_host_regex.sh, 197
 - check_http, 104
 - Check_MK, 131-134
 - check_ping, 189
 - check_swap, 130
 - Exit Codes, 28-31

- installing, 57-58
- integrated plug-ins, Nagios XI, 230-233
- redundant plug-ins, 13
- Remote Execution, 31-34
- PNP4Nagios, 187-190
- port queries, 98-100
- ports
 - versus applications, watching, 20-21
 - querying multiple, 100-102
- PowerShell, 120
- PowerShell cmdlets, 121
- preventing name-space collisions, 73
- primary data points (PDPs), 174
- problems, defining, 11
- procedural approach to monitoring systems, 9-12
- processing monitoring systems, remote versus local, 12-13
- process-service-perfdata, 188
- Python, 111

Q

- queries, local queries, 95
 - complex service checks, 102-104
 - pings, 96-98
 - port queries, 98-100
 - querying multiple ports, 100-102
- querying multiple ports, 100-102

R

- reapers, scheduling, 38-39
- redundant plug-ins, 13

- relationships, child/parent relationships, 27
- remote execution, NRPE, 59
- Remote Execution, 31-34
- remote processing versus local processing, 12-13
- reporting
 - I/O interfaces, 46-47
 - Nagios XI, 228-230
- resource_file, 72
- resources.cfg, 72
- retry_check_interval, 78
- Round Robin Archive (RRA), 174-175
- RPN, RRDTOol, 182-185
- RRA (Round Robin Archive), 174-175
- RRD data types, 171-172
- RRDs, 171
 - multicounter, 178
- RRDTOol, 149, 168-171
 - AREA, 181
 - AVERAGE, 180
 - CDEF, 181
 - consolidation functions, 177
 - create syntax, 175-179
 - DEF, 180
 - fetch mode, 212-214
 - graph mode, 180-182
 - heartbeat, 172-173
 - history of, 171
 - LINE, 181
 - max, 174
 - min, 174
 - RPN, 182-185

- RRA (Round Robin Archive),
 - 174-175
- RRD data types, 171-172
 - step, 172-173
- RRDTool wrappers, 187
- Ruby, 110
- S**
- s, 99
- scheduled downtime, notifications, 42
- scheduling, 34
 - check intervals and states, 34-36
 - distributing loads, 36-38
 - downtime, Nagios XI, 228
 - events, 35
 - parallel execution, 38-39
 - reapers, 38-39
- Scholz, Kyle, 221
- scripting templates, 87-90
- secondary Nagios daemons, distributed
 - passive checks, 150-153
- secondAxis, 209
- security, 16-19
 - network segments, 18
- sending data
 - Event Broker, 249
 - from Nagios to Ganglia, 194-197
- sensors
 - environmental sensors, 142-143
 - lm-sensors, 144-145
 - standalone sensors, 143
- service, 63
- SERVICE_CHECK_DATA, 258
- service checks, complex service checks, 102-104
- service definitions, 68
 - check_dllhost, 123
 - check_http, 98-100
 - check_load, 129
 - check_nt_cpuload, 124
 - check_ping, 97
 - check_ssl, 104
 - ping
 - child instance*, 152
 - parent instances*, 151
 - WebInject, 108
- servicedependency, 63
- service_description directive, 77
- serviceescalation, 63
- serviceextendedinfo, 63
- servicegroup, 63
- servicegroup_members directive, 80
- servicegroups, 79-80
- service notifications, 78
- service objects, 77-78
- SERVICEOUTPUT, 195
- service_perfdata_command, 194
- services, 26, 77-78
 - downside of, 27-28
- services definition skeleton, 90
- SERVICE_STATUS_DATA, 251
- services template, definition skeleton, 89
- shell wrappers, 100
- silence, monitoring systems, 19-20
- Simple Network Protocol. *See* SNMP

singularity.gov, 192-193
 displaying graphs from Ganglia in Nagios UI, 198-200
 Ganglia, 193-194
 monitoring Ganglia metrics using Nagios, 197-198
 sending data from Nagios to Ganglia, 194-197
skeleton config files, scripting templates, 87-90
SMTP handshakes, 103
SNMP (Simple Network Management Protocol), 17, 135-138
 -c, 141
 enabling on Cisco routers, 138
 MIBs, 140
 OIDs, 139-141
 -On switch, 139
 -v switch, 138
SNMP agents, 137
SNMP Version 1, 136
SNMP Version 2, 136
SNMP Version 3, 137
Solaris, vmstat, 130
sparklines, 218-220
SSH authentication, 33
standalone sensors, 143
states, scheduling, 34-36
stats headers, LQL, 166
step, RRDTool, 172-173
struct, 254
submit_service_check_to_parent command, 152

Suitcorp, data visualization, 185-187
 draw, 190-192
 NG (NagiosGraph), 187-190
summarize, 206
syntax
 CDEF, 182
 RRDTool, 175-179

T

tables, Livestatus, 162
templates, 69-70
 notifications, 41
 scripting, 87-90
test case files, WebInject, 106
Thermd, 143
threshold, 207
TIMED_EVENT_DATA, 258
timeperiod, 62
timeperiods, 70-71
 notification, 41-42
Tufte, Edward, 210, 218
tuning, 149
 NDOUtils, 150
 NRDP, 150
 NSCA, 150
tuning documentation, 149

U

UNIX, 125
 CPU, 126-129
 disk, 130-131
 memory, 129-130
 NRPE, 125

use_authentication, 68
use directives, 69
user management, Nagios XI, 234
user registrations, 208

V

-v switch, SNMP, 138
VBScript, 112
 choosing to use, 119
verbose output, WebInject, 107
vmstat (Solaris), 130

W-Z

watching ports versus watching applications, 20-21
watching the watchers, 21-22
web GUI, 43
WebInject, 105-108
 command definition, 108
 config.xml, 106
 service definitions, 108
 test case files, 106
 verbose output, 107
web interfaces, 43-44
Websensor EM01B, 143
Whisper, 201
Windows, 111
 Check_NT, 123-124
 COM, 113
 future of Windows scripting, 119-121
 NRPE-NT, 122-123
 NSClient++, 124-125

OLE, 113
VBScript, choosing to use, 119
Windows Scripting Environment, 111-113
WMI (Windows Management Instrumentation), 113-117
WSH, choosing to use, 118
Windows Management Instrumentation (WMI), 113-117
Windows Script Host (WSH), 111
Windows scripting, future of, 119-121
Windows Scripting Environment, 111-113
wizards
 configuration wizards, Nagios XI, 230-233
 NRPE Wizard, 231
WMI (Windows Management Instrumentation), 113-117
wrappers, RRDTool, 187
Wscript, 112
WSH (Windows Script Host), 111-112
 choosing to use, 118