

Fluke Networks Special Edition

Limited Edition

# Application-Aware Network Performance Management

FOR  
**DUMMIES**<sup>®</sup>  
A Wiley Brand

With the compliments of

**FLUKE**  
networks.  
.....



# About Fluke Networks

## Fluke Networks is trusted

Fluke Networks offers best-of-breed solutions that span network deployment, network performance management and diagnostics, as well as security and performance monitoring. Fluke Networks solutions are trusted by:

- Network professionals, to help them deploy, manage and troubleshoot business-critical networks
- Agencies, enterprises and carriers around the world
- By 98 of the Fortune 100 companies

## Fluke Networks is growing

Backed by the financial strength of Danaher Corporation (NYSE: DHR), a \$18B corporate parent company, Fluke Networks has remained profitable and grown significantly, both organically and through acquisition. This has enabled it to expand in, and enter, new markets to deliver best-of-breed solutions that span IT networking, datacom cabling and telecoms.

## Fluke Networks is innovative

Twenty years ago, Fluke Corporation instituted a new program for innovation, built around carefully selected groups called Phoenix Teams. Each group was given 100 days and \$100,000 to do something remarkably simple: watch how customers work, analyze the problems they face and turn those insights into new businesses. After spending days with harried networking technicians, the original team had an epiphany: there was no easy way to determine whether network problems were related to hardware, software or cabling. So they created a new way.

Today we help simplify troubleshooting of complex networks with products that provide deep visibility and uncover previously unseen problems.

In 2014, Fluke Networks was recognised by Gartner Inc. as a leader in its Network Performance Management and Diagnostics Magic Quadrant, and by EMA Associates as a Value Leader in their 2014 "ANPM Radar" Report.

## Fluke Networks is global

Fluke Networks employ over 700 people worldwide.

- Service customers in more than 120 countries
- 45 percent of revenue from outside the US
- Worldwide Headquarters in Everett WA; Engineering Centers, sales offices and Technical Assistance Centers worldwide.
- Fluke Networks is a division of Fluke Electronics Corporation



***Application-Aware  
Network Performance  
Management***

FOR  
**DUMMIES**<sup>®</sup>  
A Wiley Brand

**By Fluke Networks**

FOR  
**DUMMIES**<sup>®</sup>  
A Wiley Brand

## Application-Aware Network Performance Management For Dummies®

Published by  
**John Wiley & Sons, Ltd**  
The Atrium  
Southern Gate  
Chichester  
West Sussex  
PO19 8SQ  
England

For details on how to create a custom *For Dummies* book for your business or organisation, contact [CorporateDevelopment@wiley.com](mailto:CorporateDevelopment@wiley.com). For information about licensing the *For Dummies* brand for products or services, contact [BrandedRights&Licenses@Wiley.com](mailto:BrandedRights&Licenses@Wiley.com).

Visit our Home Page on [www.customdummies.com](http://www.customdummies.com)

Copyright © 2015 by John Wiley & Sons Ltd, Chichester, West Sussex, England

All Rights Reserved. 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 under the terms of the Copyright, Designs and Patents Act 1988 or under the terms of a licence issued by the Copyright Licensing Agency Ltd, 90 Tottenham Court Road, London, W1T 4LP, UK, without the permission in writing of the Publisher. Requests to the Publisher for permission should be addressed to the Permissions Department, John Wiley & Sons, Ltd, The Atrium, Southern Gate, Chichester, West Sussex, PO19 8SQ, England, or emailed to [permreq@wiley.com](mailto:permreq@wiley.com), or faxed to (44) 1243 770620.

**Trademarks:** Wiley, the Wiley Publishing logo, For Dummies, the Dummies Man logo, A Reference for the Rest of Us!, The Dummies Way, Dummies Daily, The Fun and Easy Way, Dummies.com and related trade dress are trademarks or registered trademarks of John Wiley & Sons, Inc. and/or its affiliates in the United States and other countries, and may not be used without written permission. All other trademarks are the property of their respective owners. Wiley Publishing, Inc., is not associated with any product or vendor mentioned in this book.

**LIMIT OF LIABILITY/DISCLAIMER OF WARRANTY: THE PUBLISHER, THE AUTHOR, AND ANYONE ELSE INVOLVED IN PREPARING THIS WORK 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.**

Wiley also publishes its books in a variety of electronic formats. Some content that appears in print may not be available in electronic books.

ISBN: 978-1-119-03390-5

10 9 8 7 6 5 4 3 2 1

# Introduction

---

**W**here would your business be without your network? Where would you be without the applications that run on your network? The answer is . . . in a world of hurt. Organizations are increasingly dependent on the performance of their networks and applications alike — if things aren't running in tiptop shape, the business simply can't run smoothly.

It used to work just fine to monitor network and application performance separately, but these days the two are too interdependent for that to be a workable approach. It takes an integrated approach to troubleshoot effectively and minimize downtime. That's where AANPM comes in. The acronym is short for *application-aware network performance management*, and it describes an integrated approach that fully recognizes the fact that the network exists to support the applications that run on it.

*Application-Aware Network Performance Management For Dummies* is an introduction to this relatively new frontier. This book explains that the end user's experience is what really matters, and that end user needs the application to work well. From his or her perspective, if the application is slow, or down completely, it makes little difference where the problem is. What matters is solving the problem, and doing so quickly. Read on to find out more about solutions that offer end-to-end visibility, from the data center to the front line where the end user is — as well as the most efficient way to locate and eliminate any troubles.

## *Foolish Assumptions*

Yes, we know what they say about making assumptions. But we're going to do so, anyway. Seeing as how you have picked up this book, it's reasonable to presume a few things about you, the reader:

- ✔ You work in IT supporting a network for an organization or business, or you're a decision-maker who might be asked to sign off on acquiring a new and better way to keep your IT infrastructure running smoothly.

- ✔ You know there must be a better way to locate and resolve IT problems that are impacting applications.
- ✔ You don't have all the time in the world to find that better way, but would be glad to discover the basics in a book like this one.

## *How This Book is Organized*

Explore the concept of AANPM and the ever-intensifying need for an integrated approach to monitoring networks and applications. We'll delve into the reasons why the old ways no longer serve companies well, find out what makes AANPM work, discuss how to get the concept running in your organization, and share some case studies.

- ✔ **Chapter 1: It's Complicated** — Every day there are new applications, new network features, and new ways to access the network and applications. Keeping things running requires the ability to look through all the layers, and that's no small accomplishment.
- ✔ **Chapter 2: AANPM is the Answer** — You need visibility from end to end of the network and application infrastructure. This chapter explores how AANPM delivers that visibility.
- ✔ **Chapter 3: Building Blocks of AANPM** — An AANPM solution pulls a wide range of performance data into one place, correlates and analyzes it, and allows you to pinpoint where in the path the roadblock exists.
- ✔ **Chapter 4: The AANPM Solution** — Troubleshooting with AANPM is a step-by-step process, and it's made all the more powerful when you connect a centralized system with local views of user and application activity.
- ✔ **Chapter 5: Implementing Your AANPM Solution** — As you choose the right AANPM solution for you, be sure it can really connect all the dots and focus on the end user.
- ✔ **Chapter 6: Ten More Thoughts About AANPM** — You're about to embark on a better way to make your infrastructure function for the people whose work depends on it. Here are some parting thoughts.

## Icons Used in This Book

You'll see icons from time to time in the margins of *Application-Aware Network Performance Management For Dummies*. They're not just there to be pretty: Their purpose is to let you know you're about to encounter some important information you won't want to miss. Here are the ones you'll find in this book:



Next to this icon you'll discover an interesting fact or suggestion that will help you better understand the concepts of AANPM and how it will help your IT team improve the user experience.



We'd like you to read and recall every word between the covers of this book, of course — but if you don't have time to study every detail, please at least pay close attention to these paragraphs.



No one has to tell you how serious the business of glitches and downtime can be; here is a thought intended to help you avoid trouble.



There are all kinds of technical details that are beyond the scope of this book, but it is hard to resist slipping in a few; if you like such details, here they are!

## Where to Go From Here

Enough of the formalities, it's time to turn the page! Once you've read the next page, then move on to the next. That's how books are supposed to be read, right?

Well, independent thinkers like you don't always follow such conventions, and that's perfectly fine. With this book, it's quite permissible to read a chapter here, another one there, skip from one topic to another, whatever your needs and interests dictate. Want to read the last chapter first? If that's your kind of thing, more power to you! This book is organized to make sense however you proceed.



# Chapter 1

---

## It's Complicated

.....

### *In This Chapter*

- ▶ Stirring trouble with apps
  - ▶ Making networks more complex
  - ▶ Tracking the problem
  - ▶ Looking across the layers
  - ▶ Monitoring performance
  - ▶ Tallying the costs
- .....

You could say there's a love-hate relationship between networks and the applications that run on them. They need each other to fulfill their purposes in life, but they tend to cause each other lots of trouble.

The IT teams are kind of like marriage counselors, bridging the sometimes-rocky relationship and helping sort through the problems. But as with marital difficulties, you've got to figure out the problem before you can craft a solution. In this chapter, we explore why figuring out the problem is more challenging than ever.

### *A Tangled Web*

What business today would thrive — or even survive — without the countless applications that run on its network and server infrastructure? There are new applications every day, it seems, created or upgraded to keep up with the changing needs of the business and the realities of the competitive marketplace.



Better features are added with regularity, more users are plugged in, new ways to access the applications are brought online. The BYOD — bring your own device — concept is just one of the ideas having an increasing impact on the landscape.

It's also taking an increasing toll on networks and servers. And not just BYOD, which complicates matters by introducing new sets of applications with each new device, every application demanding access to a different set of network assets. Voice over IP is another example, cutting costs by asking voice and data to peacefully coexist on the same network. These are just two of the countless examples of the increasing challenges and strains that today's organizational networks face.

Add in the expectations of increasingly demanding and sometimes even impatient users. Rapid response times are an expectation — not just a luxury but an absolutely vital factor in business success.



More than ever, the network is the lifeblood of the organization, supporting business-critical applications, providing an essential communication pathway linking coworkers and customers and partners, delivering the data needed to conduct transactions and make decisions. If there's a slowdown or shutdown of the network, or any kind of problem involving the applications running on it, the impact on the business can be huge and costly.

## Virtually Complex

It's not just the applications that are becoming increasingly complicated. The networks themselves are becoming more complex every day, too, and that just adds to the challenges.



Virtualization reaches across the network, from the data center to the desktop and into the cloud. Indeed, a hybrid cloud environment brings more players into the picture, because your users are relying not only on the infrastructure you oversee, but also that of the cloud provider.

All of those WiFi devices mean all the more potential headaches and performance problems. Bandwidth is strained by the uses you expect plus a plethora of unauthorized applications, more voice and more video usage.

Configuration issues can confound, application delivery infrastructure can fail, and problems can arise in countless other ways. Network issues can occur anywhere from the physical layer to the wireless. And it just keeps getting more complex all the time, it seems.

Truth be told, the network is transforming from a statically configured thing to more of a living, breathing fabric that interconnects its various components. Monitoring is transforming into management, and the tools must evolve to meet the changing environment.

## *Whose Problem Is It?*

To an end user, the application is not working properly. It seems simple enough — “Just get the technology to work, the sooner the better,” the user requests (or demands). That end user likely has absolutely no idea just how complicated that fix might be, no idea that proper functioning is a harmony of network and application performance.




In fact, one of the biggest complications may be simply figuring out whose problem it is to solve. If the application isn't functioning the way it should be, is the problem in the application? Could well be, given the complexity of applications and how often new ones show up.

Or, is the problem in the network? Is the application or the data it needs not successfully making the journey from somewhere in the cloud to the device that's somewhere out there on the network? It's easy enough to imagine how that could happen, with all of the twists and turns that are part and parcel of today's networks.

What if the user is having trouble but the various performance indicators don't even show it? What if all the lights are green — all except for that frustrating, spinning “hold on for just a minute” icon on the user's device?

When you're alerted to a problem, you embark on a multistep journey that leads eventually to the solution. The journey as outlined in Figure 1-1 seems simple enough, but the increasing complexity of the network and application infrastructure has made the process anything but straightforward.

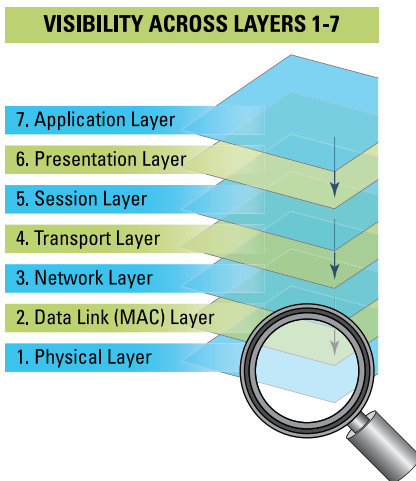


**Figure 1-1:** The workflow for problem-solving.

## Looking Through the Layers

You have no doubt heard someone talk about the need to peel back the layers of the onion. Metaphorically speaking, we're talking about the process of peering through the many layers of complexity that obscure a full understanding of a topic or a person. From a more literal perspective in the kitchen, peeling back layers of an onion may bring tears to the chef's eyes.

It's a pretty good way of thinking about the challenge of keeping your networks and applications in tip-top shape. After all, as illustrated in Figure 1-2, it's customary to think of the network infrastructure in terms of layers — seven of them, running from the data center to the branch office where the application is being used.



**Figure 1-2:** Peeling back the layers.

What makes this work tricky is the fact that a problem can arise on any of these layers, yet wreak havoc somewhere else. Your ability to quickly resolve the problem depends upon

being able to get to the root of it, on whichever layer it happens to reside.

## Separate Monitoring

Once upon a time, it made sense to have separate groups keeping watch over the disparate pieces of the puzzle. One team could oversee application performance, using its own monitoring systems. Another could focus on network performance, with another set of systems for monitoring. The rationale was hard to dispute, as the functions and disciplines were (and are) quite different from one another.



On the network side, analyzing performance seems straightforward enough. The tools for the task boil down to two primary categories: *network management systems*, or NMS, and *packet capture and analysis tools*.



NMS has typically been focused on the infrastructure itself. It's all about planning capacity, monitoring devices, managing configurations, analyzing interface traffic, conducting fault management, that kind of thing. What it does not involve is analyzing application response time, tracking TCP errors, or following other things that can cause trouble for applications.

Even so, it can get complex, so much so that an NMS may manage down only as far as Layer 3 devices. Switches at Layer 2? Nope. On top of that, polling data can miss the impact of usage spikes if it's aggregated over too many minutes, as can be the case. And as a centrally located system, the NMS may not be able to accurately measure those things that can impact end user response time on a different place in the network.

That's where packet capture and analysis tools come into play. You get a lot more detail this way, which you need, but you may wind up with data overload — millions of packets in need of analysis, showing up in a variety of user interfaces. It takes a lot of skill and experience on the part of the engineers, not to mention a lot of time.



NMS and packet analysis tools can certainly prove that the problem is not within the network, but that doesn't solve the problem, does it? Network engineers don't typically have a

way to see and analyze live traffic where transactions are happening, so their ability to troubleshoot an application problem is limited. They may be able to get at past data to analyze, but that offers no answers as to what's happening now, and it certainly isn't much use when it comes to identifying and addressing problems before they get out of hand. The horse is out of the barn by that point.

On the application side, application performance management systems have been the go-to for keeping an eye on how things are working. They can analyze transactions, track usage, evaluate the end-user experience, conduct user-defined transaction profiling, and in general keep an eye on the health and performance of the various application infrastructure assets. Even so, if these systems find that an application is running slowly, they can't always tell if the problem is within the application itself, or in the network.



As you can see, it's no longer sufficient to monitor networks and applications separately. Networks and applications are just too interdependent for a silo approach. It's neither efficient nor effective, and it can make it quite tricky to track down a problem, especially when the teams on both sides of the equation have performance indicators showing green, while the frustrated end user is complaining.

The network and application performance management tools of the past are simply not designed to track the interplay between networks and applications. And you absolutely must understand and manage the interactions between these environments if you're really going to optimize the experience for the end user.

This is the kind of thing that keeps network professionals awake at night. Fluke Networks polled some 3,000 of them and found that 82 percent ranked network and application performance problems as a critical issue or concern. Just over half said an NMS simply can't get to the root of the problem most or all of the time. And about as many respondents said that troubleshooting these kinds of problems forced them to get up from their desks some or most of the time to go in search of the trouble.

## Costly Old Ways

Satisfaction on the part of the end user is of paramount importance, but so is cost. And as it turns out, the old ways just don't cut it anymore from a cost perspective, either.



Consider that according to Cisco research, as much as 49 percent of the total cost of ownership of a network is tied up in labor. The more time your experts spend chasing their tails trying to track down and resolve issues using inefficient tools and processes, the more it's going to cost from a labor perspective. Your IT teams need better ways to work together, with access to correlated data, so they can find root causes of issues more quickly and get the solutions in place fast.



Wasted IT labor is just the beginning of the cost of downtime. The Ponemon Institute conducted research a few years ago on just how much downtime can cost from a user perspective. Here are some of the sobering results:

- ✓ Every minute that a data center is down can cost \$5,600. Yes, every single minute.
- ✓ The average reported downtime lasts 90 minutes.
- ✓ Doing the math, that means the average incident costs just over half a million dollars.
- ✓ A partial data center outage averages 59 minutes and costs about \$258,000.
- ✓ A total data center outage takes an average of 134 minutes for recovery, running \$680,000.
- ✓ The percentage of engineers who admit they don't have the tools they need for quick and accurate problem resolution is 55 percent.
- ✓ The percentage confessing that troubleshooting gets bogged down by finger-pointing and non-collaborative results: 70 percent.

Another perspective on downtime comes from a 2012 study by Aberdeen Group. The study calculated the true cost of downtime for a company with a thousand or more employees at more than a million dollars per hour. What's more, the cost grows every year, according to Aberdeen Group, by more than

15 percent. That's because IT shops are being asked to do more work with fewer people, making them more reliant upon complicated, automated technology.

According to that study, enterprise companies experience significant IT outages three times a year on average. Add that all up, and the costs become downright scary. Surely there must be a better solution, one that can take a downtime or slowdown lasting hours or even days, and turn it into a few seconds or minutes of pain.

As it happens, there is. Of course there is, or there wouldn't be a reason to produce this book, or invite you to read it. There can, indeed, be happy Hollywood endings, even in the complicated world of IT.



## Chapter 2

# AANPM Is the Answer

---

### *In This Chapter*

- ▶ Seeing the application
  - ▶ Appreciating AANPM
  - ▶ Tearing down the silos
- 

It's clearly not enough to ensure that your network infrastructure is operating at its best. After all, it's the applications that run on those networks that are doing the tasks that users require, and if they're not up to the job, then the job is not getting done.

This chapter introduces up-and-coming concepts for ensuring that both networks and applications are behaving the way they're supposed to behave. Read on and you'll learn why tearing down the silo walls is the answer for keeping end users happy.

## *Application-Aware*

Chapter 1 talked about the layers of the onion, the seven layers of the network and application infrastructure. To get a handle on the network infrastructure and the applications, you need visibility from end to end, across the seven layers. It's important to be able to peel back those layers and have a good view of each of them in the process. And unlike peeling that onion in the kitchen, you'll need to be able to do so in a way that doesn't make you cry.



The solution needs to bridge the gap between an NMS and the complex packet capture and analysis tools that deliver more information. It needs to be able to sniff its way through the

onion to find those hidden problems that sap productivity and give the IT department a black eye.

What you need is a solution that addresses a number of shortcomings of the old way of doing business. For example:

- ✔ Unmanaged equipment may have seemed like a cost-cutting bargain to purchase, but turns out to be expensive to troubleshoot because the health of this part of the infrastructure can't be monitored.
- ✔ Undocumented networks are troublesome yet not uncommon, given the fact that frequent changes make any documentation out-of-date in no time. Effective troubleshooting requires the ability to discover the real-time path through the network.
- ✔ Newer Ethernet and WiFi technologies aren't always monitored because it seems like the boost in capacity will overcome any problems (and that, of course, doesn't always turn out to be the case).
- ✔ Mountains of data can obscure the source of the problem, which probably lies in just a few packets. You need the ability to quickly sort through the captured packets and find the bad eggs.
- ✔ Intermittent problems from the past that become visible only hours later can cause real headaches, unless you can go back to capture and analyze large quantities of granular data over an extended period of time.
- ✔ Wireless devices including BYOD can be tricky to identify and monitor, as can interference from Bluetooth devices.
- ✔ Problems outside the network can cause headaches. You need to be able to identify them and come up with supporting evidence to hand to whoever is responsible for achieving a quick resolution.



In short, what you're looking for is a holistic solution to network and application performance — a solution that captures all of the data in the network and delivers intelligent analysis of that data. This dream system will allow your engineers to narrow down that root cause, or perhaps even find that the problem is not in your network at all.

It will aggregate and correlate and mediate a wealth of information, including flow, SNMP data, information gathered from other devices, and it'll do so with granularity of a millisecond. It'll put all of the high-value information into a handy interface that meets the needs of the IT engineer looking at it.

Is this asking too much? What solution can help your IT department overcome the countless challenges posed by increasing complexity?



The solution goes by the name of *application aware network performance management*, or AANPM for short. Some people refer to it as NPMD, which stands for *network performance management and diagnostics*, which is a perfectly fine term, too, but for now it seems appropriate to stick with AANPM, because the “AA” part of it underscores the whole point of being aware of application performance.

AANPM, thus, is an application-centric way of looking at what's happening across the network. It allows engineers to really stay connected to the end-user experience — not from an application coding perspective, but from the perspective of how applications are deployed and how they're performing.

## Case Study: i2i Infrastructure Insights

Network performance is critical to Dallas-based i2i Infrastructure Insights, a network consulting company with a wide range of clients that rely on the firm to keep infrastructure functioning smoothly. Gaining visibility into large enterprise networks and the applications running on them was a top priority.

AANPM, in the form of Visual TruView from Fluke Networks, allows i2i to keep watch over its customers' networks and applications, and quickly

isolate problems. Application-level detail helps the company address performance discrepancies and deal with ongoing issues using accurate root cause data. No need for finger-pointing, just optimized productivity.

What's more, as the company expands its geographic reach, it's critical to streamline remote monitoring, which the AANPM solution allows. The results are good for both i2i and its clients.

Think of AANPM as a bridge of sorts. It links the network-infrastructure-focused monitoring on one side with a focus on applications and servers from the other side. It leverages data from both methodologies to allow cross-platform visibility. It keeps an eye on all of the layers of the onion, so that all branches of IT can be sure that overall performance is at its peak — and it reduces that onion-caused shedding of tears.



AANPM is an answer to some of the challenges created by virtualization, BYOD and cloud-based services, because it can identify problems wherever they might occur along the network path. Application performance data sheds light on when users are experiencing poor response times, as well as which application component might be playing a role in the delays. This is truly actionable information that can help the application team understand which component needs attention.

## *The Benefits of AANPM*

Consider the way things often go when there's an issue for the end user. The finger is pointed first at the network, and the logical place to begin the search is at Layer 1. But pinpointing the actual cause is time-consuming and labor-intensive, even if you troubleshoot every aspect simultaneously.

And then comes the blame game. There's a big debate about where the problem is, and the environment grows toxic as the various players try to figure out where the problem is and who's responsible for fixing it. The clock keeps ticking as the solution remains elusive.



AANPM, on the other hand, provides end-to-end visibility of the entire IT infrastructure. Right there in one place are all of the various data points that one would expect to find in traditional network management systems and application performance management systems. It's a single-dashboard view covering both critical applications and the underlying network infrastructure.

That means much faster problem-solving. Become aware of problems with a key application and right away your engineers can use that cross-platform visibility to track down the root cause.



AANPM is a window looking through the entire network infrastructure (physical and virtual alike), across the LAN and WAN connections, into the wireless environment, and also into all the tiers of the server and application environment. No need for debating or arguing or blaming, because the roadmap for finding and solving the problem is right there.

Because finding the root cause is so much faster, achieving a solution is much speedier, too. That, of course, means a better experience for the end user, no small feat given the ever-increasing complexity of the infrastructure, with so many tiers and applications residing in so many places, and users accessing applications from so many different directions.



Indeed, AANPM keeps a real-time focus on the experience of the end user, by monitoring the user's transactions and how they're being delivered by the infrastructure. You're able to spot such issues as slow performance and identify where things have bogged down — starting your troubleshooting at the place where the trouble is being felt.

AANPM makes visible all of the infrastructure elements and tracks the flow of data and messages between the various elements. Not only does that help get to the source of any problem more efficiently, it also provides an element of proactivity when it comes to performance. That allows engineers to spot and fix small problems before they grow into monsters.

That, in turn, means greater productivity. If engineers aren't battling monsters, they're speeding up the MTTR, or *mean time to resolution*. AANPM offers the ability to identify the problem domain in the data itself. Information about the end user experience is already broken down into how the server, network, and application are each contributing to issues facing the user.

By removing the guesswork from identifying and isolating the problem domain, you're cutting down the downtime, which means everyone is more productive, not just the engineers. And if they're not slaying monsters, they've got more time for strategic work, tackling such worthwhile aims as optimizing performance.

Then there's money, a topic of importance to just about everyone. Poor network and application performance can be quite expensive, in many different ways. Downtime is very costly, for one thing, as outlined in Chapter 1. So is inefficient troubleshooting, as well as buying multiple systems when one would suffice.



And the loss of productivity is costly, too. Consider this productivity quandary. Say that two routers are performing poorly. One is at a branch office, the other supporting a critical business application. AANPM helps you to see both of those problems in a big-picture kind of way, helping you to fix the most impactful problem first. That gets more users back in the productive game more quickly, and that, once again, saves money.

One more thing that AANPM brings is a boatload of data. "Right," you're thinking, "I'm already up to my ears in data." But this is really useful data that'll make you look like a star. We're talking data that can be used in many different ways to help optimize your network and applications, and make the case for investments that are needed in various places.



AANPM helps you correlate and analyze performance data, pinpointing areas of slow, poor performance. It becomes clearer where server upgrades are needed, and when those upgrades are made, whether they've resolved the issue or perhaps made something else worse. This flood of data helps with capacity planning and clarifies the need for added bandwidth. You can track the good (or occasionally not so good) impacts of WAN optimization, data center consolidation, or virtualization.

Finally, AANPM is good for IT departments that need a little love from the executive side of the business. An AANPM system helps execs understand the business cost of critical applications, as well as the costs when they go offline.

The information generated by the system helps demonstrate how critical applications and the IT infrastructure are interrelated and interdependent — and ultimately, the more information you can feed to the number crunchers, the more likely it is they'll have a healthy appreciation for the work that you do.

## Case Study: Southeastern Freight Lines

Based in South Carolina, Southeastern Freight Lines employs nearly 7,000 people in more than 80 locations. Its network supports 5,000 applications developed in-house along with some third-party applications, focusing on everything from shipments to billing to maintenance. Most of these applications are accessed from remote locations, through VPN or wireless clients.

That's a lot of users, and good customer service leaves little room for trouble. If network or application issues start to impact the user experience, the technical service department hears about it pretty quickly

and loudly. With that in mind, the company employs Virtual TruView™ from Fluke Networks to create a single interface covering a wide range of network and application performance analytics. The solution follows such data sets as stream-to-disk packet storage, application response time, transactional decode, IPFIX (NetFlow), and SNMP, to name just a few.

IT specialists are able to promptly identify the root causes of network delays and unresponsiveness. For a company constantly focused on process and quality improvement, it's a welcome tool.

## No More Silos



It's hard to overstate the need for network management tools that are application-aware. It's the roadmap out of the silo mentality (and let's be honest, when's the last time anyone other than a grain farmer spoke of silos in a positive way?).

In the case of IT, you just can't get the big picture needed for troubleshooting network and application performance issues if your teams are completely siloed in terms of network, server, database, and application. Giving each support team its own monitoring tools, with no integration or infrastructure visibility, is a recipe for finger-pointing, with a bunch of players all claiming to see green lights and no one seeing a real solution.

AANPM, on the other hand, spreads knowledge and visibility across the organization. The result is less friction, more productivity, and a lot less costly downtime.



## Cast Study: Cimarex

With three main offices, multiple data centers, and some two dozen additional locations linked with MPLS and VPN connections, troubleshooting could be quite a challenge for the IT team at Cimarex, a Denver-based oil and energy company. Beyond standard business productivity applications, the company uses VoIP and high-bandwidth video over IP solutions.

The challenge is to provide the support needed at each of the data centers without having to engage vast levels of engineering support. Before adopting AANPM, the three-person

engineering team had been finding that its method of troubleshooting was not as effective as it needed to be, and was hampering other initiatives. Oftentimes, it was necessary to diagnose and repair problems remotely.

Cimarex has employed multiple Fluke Networks solutions to harness AANPM power. The company began by using OptiView XG to monitor its biggest data center and eventually expanded to use Visual TruView in order to identify, test, and resolve issues across its network and application infrastructure.

## Chapter 3

# Building Blocks of AANPM

---

### *In This Chapter*

- ▶ Collecting the data
  - ▶ Drawing the map
  - ▶ Recreating the problem
  - ▶ Monitoring the cloud
  - ▶ Getting a good vantage point
- 

**T**he whole idea of AANPM is to, in essence, simplify the complicated. Okay, perhaps that statement is a bit of an oversimplification in and of itself. There are a lot of complex pieces that go into AANPM, and that's the focus of this chapter. AANPM assembles a wide range of performance data, pulls it into a powerful dashboard, allows you to figure out what happened in the past, and gives you remarkable visibility.

## *A Wealth of Data*

With an AANPM system, your engineers have at their fingertips all of the information they need to analyze what's happening all across the enterprise network as well as the applications that it supports. The system pulls together critical performance data from the network and the applications alike.



Performance data assembled through AANPM include stream-to-disk packet storage, analytics on application response, IPFIX (NetFlow), and SNMP. The system keeps a watch of the health of every link, device, and interface across the network, and provides traffic analysis that includes virtual traffic flows.

A solid AANPM system provides these volumes of data on a real-time basis, which is more than you can expect from some traditional systems of yesteryear. Yet it also stores historic data to enable powerful, retrospective data analysis.

## Mapping the System

You can custom-configure your AANPM system to display on one dashboard all of the network and application performance metrics needed by users in all IT functions. It's cross-platform visibility, and it can be tailored to the needs and levels of responsibility of individual users (they can also be granted visibility to adjacent areas of responsibility).

With their individualized dashboards and glimpses into the responsibility areas of others, members of cross-functional teams are better equipped to work together on performance problems. They've got the tools they need to work backwards from the effect to the cause — quickly and collaboratively, resolving problems promptly.



A performance map supplements individual dashboards, allowing engineers to see the entire enterprise network. The better visibility they have into what's happening on the network, who's using it, where they're connected, and what path leads from here to there, the faster the engineers can isolate the root cause of any problem.

As an example, engineers can see links that are showing errors or high utilization, and with a click of the mouse they can take a deep dive into the transaction that's introducing latency to all upstream transactions. A good system offers automated discovery of all network elements and connectivities as well as applications. It creates performance base-lines from actual packets, and that makes configuration more simple.

Through AANPM, it's possible to analyze every device in the path — across the WAN, through virtual infrastructure, and on into the application environment. Engineers can identify the problem domain and the root cause of the problem through simple, logical workflows and a few mouse clicks.

## Case Study: Global Call Center

How many electronic devices and other appliances have you bought in your lifetime? Odds are you've been offered, and quite possibly have purchased, an extended-service contract on some of those items. One company that provides that kind of coverage operates a call center that handles more than 20 million calls a year, from more than 150 million customers. Just imagine the kind of infrastructure required to run that kind of operation.

There are some 5,000 users at any given time, across multiple locations, accessing a customer relationship

management database, using VoIP to answer calls, and connecting to an interactive voice response platform. No time is a good time for downtime.

This company has implemented AANPM solutions that have made a significant difference. The network team, for example, can now prove whether or not the network is the problem, and can offer up critical information revealing if an issue lies within the application or a server, or the network itself. Issues that once took days to resolve can now be dealt with in hours or even minutes.

Whether a performance event is real-time or historic, engineers can isolate it right down to the individual network element, transaction, or packet. If more detailed visibility is needed, graphical packet filtering does the trick.

## *Recreating the Problem*

How many times have you taken your car in for service, but that nagging noise or hesitation just isn't happening when the mechanic is watching? Doesn't mean it's not a real problem, but it's tough to solve if the service professional can't make it happen.

Similarly, solving intermittent or historic problems network or application problems requires the ability to reconstruct events. AANPM comes in handy in this regard, because it stores all data flows, transactions, and packets.



With this kind of data at hand, it's a whole lot easier for engineers to reconstruct web content, database queries, and N-Tier applications transactions. Your team can use flow forensics to go back in time and identify traffic on key links, and even play back video streams and VoIP calls.

It definitely eases the challenge of solving problems from the past, something that hasn't always been easy to do if several hours pass before problems come to the attention of the engineer. AANPM can shine the light on those nagging, intermittent problems that seem to hide in the network and clog up productivity.

## Case Study: FloraHolland

You've seen the stunningly beautiful photos of acre upon acre of flowers growing in The Netherlands. It's far more than pretty — the floral industry is a big slice of the Dutch economy. FloraHolland sells more than 20,000 different flowers and plants at auction, some 12.5 billion items a year, worth more than \$4 billion. The infrastructure supporting this trade is critical, and there's zero tolerance for downtime or network performance delays.

The organization encountered a response-time delay between the database and the auction clock, slowing the clock display information by a couple of seconds between each auction, which is far too long in an environment where some 50 million items are auctioned to onsite and international bidders every day.

So where was the problem — in the database, the application, or the network? Using Visual TruView™ by Fluke Networks as an AANPM solution, FloraHolland's engineers quickly determined that the database was taking an extra second or two to respond to queries. They were able to optimize the database and solve the problem quickly, without going through the time-consuming trouble of reviewing the network and application.

How big a problem is a second or two here and there? Bigger than you might imagine. Added up over a year, the cost savings afforded by the AANPM-driven solution totaled more than \$200,000, and that doesn't include the value of potentially improved brand image.

## Eye on the Cloud



You've got to meet your service-level agreements if you want your IT team to look good. AANPM can help, by providing instant, real-time reports on bandwidth usage by top applications, hosts, and conversations. It's a window into where and how the network is busy, with important insight into which is critical business usage and which is less critical, or even recreational.

It's information that can help identify which links need more bandwidth and which can afford less — and such information over time can reveal growth areas where expansion should be considered in the future. In addition, you're able to manage VoIP quality of service, assessing traffic levels before and after deployment. Performance and playback analysis helps you promptly act upon the causes of poor call quality.

You're also monitoring the link to the cloud through quality of service and class of service reporting. That helps you keep tabs on cloud service delivery, ensuring that providers are meeting their own SLAs.

## Excellent Visibility

As we mentioned at the outset, you need visibility throughout all seven layers of our metaphorical onion, and that's what you get with AANPM. It's visibility across the LAN, WAN, and data center environments, including all tiers of the server and application environment, whether they're physical or virtual.



To use the AANPM solution from Fluke Networks as an example, your glance across Layers 1 through 7 supports rates from 1 to 10 Gbps. Added to that is the ability to monitor and troubleshoot the wireless infrastructure, and to support remote locations through a portable form factor if you need more in-depth visibility. In short, you can see from the data center to the user device, and everything in between.

## Case Study: Credit Union

A Southern U.S. credit union with more than 20 branches and about 100,000 members has a small IT team that was having trouble quickly determining whether issues were caused by network or application glitches. Sometimes it would take days to track down a problem, as the team relied on guesswork and trial and error.

An AANPM solution turned out to be the answer. These days the IT team can pinpoint problems in minutes, or even seconds. Just as important, the team has been able to proactively assess and update various parts of the network infrastructure to help prevent problems.



## Chapter 4

# The AANPM Solution

---

### *In This Chapter*

- ▶ Getting the word about the problem
  - ▶ Launching an investigation
  - ▶ Fixing the problem
  - ▶ Optimizing the infrastructure
  - ▶ Seeing across a big network
  - ▶ Sharing local and centralized information
- 

**S**o how does it all work? How does AANPM make life easier for engineers? It starts with knowing there's a problem before the phone rings with an end user on the line. In this chapter, we'll quickly walk through a typical scenario as it's addressed by a network application and performance solution.

We'll also spotlight some of the shortcomings of centralized monitoring and troubleshooting in a widely distributed environment, and dream of true "end-to-end" visibility.

### *It Begins With an Alert*

Problems happen, but from the perspective of an engineer, it sure is nice to be the first to know about the problem. Wouldn't you rather have a system that alerts you when there's a problem, rather than field a call from a disgruntled user? Or if you're a manager or chief information officer, you probably have at some point experienced that sinking feeling when you walk into a meeting and someone asks about a performance issue that is news to you. It's far better to know of the issue already

so that you can tell the group your team is already on the case. Providing helpful alerts is one of the first goals of AANPM.



Sure, setting up alerts is possible even without AANPM, but that often requires manually configuring network management tools for each network, telling them to ping or discover all of the devices in each broadcast domain. By contrast, an always-on network and application performance solution automates the discovery and reduces the setup time.

The system collects performance data, stores it in a database and feeds the display on the customizable performance dashboard. Set your own baselines, such as the service-level agreement, and the dashboard can then sound the alarm when performance varies from the standard. You may be able to integrate an AANPM system with existing network management systems.

## *Begin the Investigation*

Once the alarm has sounded, it's time to look into the scope of the problem. Imagine how helpful it would be to have quick access to all of the pertinent data, such as SNMP, flows, packets, and end-user-response time. An AANPM system has already socked all of this info away for future analysis.



The system also allows you to discover the path from the client to the service or application, which saves time. Once the path between devices is found, it can be monitored for problems across internal and external networks and devices. It's ideal if the system can provide interfaces with both 1Gbps and 10Gbps connectivity, and also be able to capture data at line speed on the wire. Depending on the system, you may be able to trace a path that identifies Layer 2 and Layer 3 devices along the way, with granularity needed to spotlight the source of the problem.

Is the issue with a client or a group of clients? If so, the engineer can learn whether it's a wired or wireless network issue by carrying out a performance or application response test. AANPM makes this simple. It also can spotlight the source of a malware outbreak.

## *Isolate, Find and Fix the Problem*

In this scenario, the problem has now been isolated down to a single network segment, switch, router, server, or application. The path and all of the devices and ports in the path have been identified — now it's time to analyze the path, which will require traffic statistics for each link in order to figure out whether the problem is a faulty device, link media, traffic overload, noise or interference.

Ultimately, your engineer can confirm the cause of the issue, determine how to fix it, implement the solution, and validate that it has worked. All of this, of course, can happen much more effectively with the help of an AANPM solution.

## *Optimize the Network*

Troubleshooting is one thing. Preventing trouble is even better, and one great benefit of more efficient troubleshooting is that you'll be left with more time for network optimization.



Your network application and performance solution is a great tool for that work, too. It provides the visibility you'll need to document and audit the health of the network, pinpointing poor performance, identifying where paths of applications or servers are running slowly, and addressing those that are the most troublesome or critical.

You'll end up with information that will help prioritize improvement projects such as server upgrades, as well as make a solid business case for spending the money. The beauty is that you'll be making better-informed decisions on your IT spend. Instead of just diving in with an upgrade in a place where utilization is highest, you can first gain an understanding of how that bandwidth is being used, then decide what kind of upgrade makes sense.

## *Information from Everywhere*

“End to end” is a great term and a wonderful goal for monitoring and troubleshooting your network and application infrastructure. But what does it really mean? Complete visibility

and end-to-end application awareness are increasingly complicated in large, distributed environments.

A typical approach to AANPM has meant centralized instrumentation and data collection. Packet-based analysis will be deployed at areas of high traffic concentration, such as ingress and egress at the data center, the core backbone, and large remote site networks. It's the most definitive approach, albeit also the heaviest and most expensive. It's possible to deploy other technologies such as synthetic test and flow records more broadly, but the performance picture is less detailed.



The problem is, the larger and more distributed the environment, the more challenging it becomes to troubleshoot application performance and end-user experience. It's tough to achieve sustained monitoring and instrumentation of remote site networks. Such remote sites may simply be too numerous to deploy packet-based sustained monitoring appliances, and you'll typically gather flow records only from remote site edge routers.

And that leaves a gap in the big picture — there's no information on the conditions and activities within the remote site. Remote sites can be complex themselves, with a mix of wired and wireless access technologies connecting a whole bunch of endpoints. All of that activity at the remote site can have a direct impact on the experience of end users at that site, including how applications perform for them.

The answer for most organizations is to employ local troubleshooting tools that can assess the situation at the remote site on a user-by-user basis. Problem solved, right? Not quite.

## *Local is Not Enough*



This local troubleshooting is important and useful, but it's incomplete if it is unaware of what's happening at the central data center. Just as centralized tools may not have full visibility at the local level, local tools can only guesstimate what's happening at the central data center.

It's possible to integrate remote and local tools manually, or deploy heavy user/client agents. But these solutions are not

always practical. Try custom integration of local and centralized tools and you'll find it to be costly to achieve and just as expensive to maintain. Try correlating data sets manually and you'll find it's often too slow and cumbersome to be useful. And try putting direct measurement agents on every end-user device and you'll find that is also quite expensive and nearly impossible to administer.



You're better off if you can directly connect centralized AANPM monitoring and visibility with local views of user and application activity. Do that within a single solution and you'll have true end-to-end visibility.

It sounds like a dream, but it is possible. For example, Fluke Networks' Visual TruView offers a combination of packet-based inspection and analysis, flow record collection and analysis, and network device SNMP polling, and it comes built into a single appliance that can be deployed anywhere it's needed. That's often at the data center or at key sites of traffic aggregation. TruView now integrates fully with the company's OptiView XG network analysis tablet, which is designed for fast, local network and application analysis and troubleshooting.



The integration shares live status data and topology information to give greater visibility on both sides of the information divide. The centralized TruView even retains information gathered when an OptiView XG is connected, which adds to the understanding of local topologies. That is, indeed, the "end-to-end" dream.

## Chapter 5

# Implementing Your AANPM Solution

---

### *In This Chapter*

- ▶ Beginning with the end-user
  - ▶ Tying it all together
  - ▶ Encouraging collaboration
  - ▶ Getting the most for your money
- 

**A**ssuming you worked your way through at least some of the preceding chapters, it's reasonably safe to assume that you find AANPM to be an intriguing option. Otherwise, you would have put the book down by now.

This chapter explores some of the things to consider as you choose and implement an AANPM solution. You want a faster, more intuitive way to troubleshoot those pesky performance problems? Here are some things to keep in mind.

## *Focus on the End User*

Your solution has to monitor the situation from the application end-user's perspective. That's an absolute must.



If you're going to succeed at this, you'll need to be able to perform packet captures on the desktop of an end user, then time-correlate that information to all of the many application and network components. Your aim is to directly pinpoint the root cause of the end-user's issues.

## *Tie it All Together*

The whole point of AANPM is to bring together a bunch of performance data focusing on all the components of an application and the network infrastructure. What good would all that data do if you couldn't converge all that data into a single, unified management console? This convergence of data is the secret for improving the efficiency of troubleshooting, and the simplified workflows inherent in an AANPM tool are focused on quickly identifying and resolving performance issues.



With that in mind, make sure your AANPM tool really does integrate and correlate data between components. In some cases, a solution will offer multiple, disparate tools that perhaps can be configured through a single console, but don't actually correlate the data between the tools. Remember, this convergence is what allows for advanced analytics, and efficient and timely identification of problems. So don't settle for less.

## *Encourage Collaboration*

Having converged data at your fingertips can greatly increase communication and collaboration across different teams. Disparate IT support teams can work together resolving performance issues if they have a single performance-management interface with a holistic view of application performance. Collaboration is good.

What's more, your total cost of ownership goes down when you consolidate both the data and the diagnostics into a single tool.

## *You Get What You Pay For*



Organizations are always on the lookout for ways to save money, and that's certainly a good thing. But when shopping for AANPM, it's worth remembering that the cost to purchase is not the only important thing, and may be far from the most important.





What's critical is to choose the tool that really speeds up the mean-time-to-know and mean-time-to-resolve metrics for application performance issues. As noted in Chapter 1, time really is money when you're talking downtime.

The solution that restores application performance issues the fastest will end up saving the most money (not to mention, putting the IT department in the best possible light). So be sure you're getting the solution that will quickly get those end users back to the business of making money.

## Chapter 6

---

# Ten More Thoughts About AANPM

---

### *In This Chapter*

- ▶ Rolling with the changes
  - ▶ Moving past the war room mentality
  - ▶ Raising your expectations
  - ▶ Sorting out the acronyms
  - ▶ Pursuing perfection
- 

**T**here really is a better way to keep networks and applications running smoothly. There are answers that will reduce stress, promote better decision-making, bring team members together, and promote proactivity. Read on for more thoughts!

## *Change is Constant*

No one needs to tell you that business-critical applications and services are in a constant state of change. Every day they're expanding to support new users and features, or they're migrating to cut IT costs, or they're evolving to support BYOD. Every change has the potential of increasing the strain on the network and server infrastructure. Rather than scramble to stay ahead of costly performance problems and downtime, it's best to use a solution such as AANPM to anticipate where problems will arise.

## *Everybody Needs Visibility*

It makes sense that network professionals are concerned about the performance of their networks more than just about anything else. An Aberdeen Group survey found that 35 percent of respondents don't think they have adequate tools to resolve network events, and 34 percent said they don't have the visibility into network activities that they would like. Sounds like a job for AANPM!

## *Make Peace, Not War*

Got a network problem? Call all of the various players together into a “war room” to battle the problem — makes sense, right? Not necessarily. War rooms can be toxic, with more finger-pointing than problem-solving, according to a white paper from the Center for Leadership Research & Development at Seattle Pacific University. A better approach is to give everyone on the team a better view of the big picture, not just their silo. A common monitoring system, such as AANPM, can make this happen.

## *Zero Tolerance*

Organizations can be amazingly tolerant of sub-optimal network and application performance. Database issues, slow response times, service disconnections — they're just part of doing business, some people think. It's anybody's guess why such things are tolerated, but too often they are, for days, weeks, months, or forever. Sometimes the solution seems too expensive, but it's crucial to note how expensive it is to *not* solve the problem, especially when you add in both hard and soft costs. The answers are out there, and it's worth the trouble and expense of seeking them.

## *New and Improved*

Once upon a time, software updates happened every quarter, perhaps every six months, maybe annually. Now, they happen nearly all the time. It's that way with smartphone apps, but

that's also becoming the norm with business applications on your network. This is one more reason to climb out of the silo. Many forward-thinking organizations are getting their network engineers together much more regularly with the software folks who are pushing out the constant updates. Better to be up-to-date on changes as they happen, rather than find yourself behind the times when trouble tickets come in.

## *Counting Dollars*

AANPM systems are great when it comes to minimizing costly downtime — that much is spelled out in the other pages of this book. In fact, they can also help you understand the financial impact when networks or applications do experience trouble. If these tools can help measure the revenue impact of downtime, it stands to reason that they can also gauge revenue implications when things are working perfectly. In the future, expect this type of data to play an ever-bigger role in business analytics, helping to understand how customers are using applications, and what revenues are being generated.

## *Alphabet Soup*

FWIW, if you're in IT, you get to know your letters and acronyms PDQ. Otherwise, you're not likely to LOL. By now you've got AANPM seared into your brain, but consider that there are other ways to abbreviate the same idea. Gartner Group has named a new "magic quadrant" in NPMD, short for network performance management and diagnostics. Enterprise Management Associates uses the acronym ANPM for pretty much the same thing (they're figuring that there's no need for two A's when one ought to suffice). Whatever you call it, it's worth the trouble to correlate your application performance management (APM) with your network performance management (NPM).

## *KISS for You*

While we're throwing around acronyms, let's not forget KISS, short for "keep it simple, stupid." That's the whole point of AANPM, when you get right down to it. It's offering a holistic view designed to let any engineer, regardless of experience or

training level, quickly locate the root cause of a problem and get started solving it. Given how difficult it is to find and retain a highly experienced network or application engineer, this simplicity is incredibly valuable. If a troubleshooting tool is too much trouble to understand and operate, what good is it?

## *Numbers Have Their Limits*

Are you trying to make bandwidth decisions based on utilization numbers alone, and finding that it's awfully time-consuming while lacking the type of analysis needed to make accurate decisions? You're not alone. You're better served by solutions such as TruView-Flow from Fluke Networks. Because it stores one-minute burst utilization data for a year, you can apply business hour filters and quickly figure out where you have the least room for growth on your network. A single click will identify who or what is consuming that bandwidth.

## *Nobody's Perfect*

Maybe not, but near-perfection is possible when it comes to network and application performance, if you have the right tools. If you can achieve end-to-end coverage of the network — from WiFi in the access layer to high-performance links in the data center — then you have right solution to resolve issues proactively, before they begin to hurt the bottom line. As a famous politician once said, “Yes, we can!”

# Take Dummies with you everywhere you go!

Whether you're excited about e-books, want more from the web, must have your mobile apps, or swept up in social media, Dummies makes everything easier.



Visit Us



Like Us



Follow Us



Watch Us



Join Us



Pin Us



Circle Us



Shop Us





## TIRED OF **GUESSING** WHERE THE **PROBLEMS** ARE? FIND AND **SOLVE** THEM **FAST** WITH FLUKE NETWORKS

Do you find your IT Department blaming each other for performance problems in critical systems? Is your management tool missing important events?

### Make end-to-end visibility a reality with Visual TruView

Fluke Networks has created a solution that is application and network aware, provides end-to-end visibility and both real-time and back-in-time data. Visual TruView combines critical packet, transaction, NetFlow/IPFIX and SNMP data and presents the results through a time-correlated dashboard view to help you quickly see how well your infrastructure is performing in the context of end user's experience.

### Fluke Network's AANPM solutions provide:

- > Holistic enterprise monitoring
- > Automated discovery and infrastructure diagramming
- > Problem domain isolation
- > Event reconstruction for root cause resolution



*Simple to deploy and easy to use – **Fluke Networks' Visual TruView** helps IT teams work collaboratively to solve problems fast!*

**Fluke Networks Application Aware Network Performance Management (AANPM) is:**

### ▶ FLEXIBLE

- > Wired/Wireless
- > Fixed/Portable
- > 1/10 Gbps
- > Layer 1-7

### ▶ INTELLIGENT

- > Self-learning Baselines
- > Time Correlated Views
- > Path Analysis and Visualization
- > Automated Configuration, Discovery and Diagramming

### ▶ COMPLETE

- > End-to-End Visibility
- > Packet, NetFlow, and SNMP
- > Correlated Dashboard View across Multiple Data Sources
- > Stream To Disk Packet Capture



## A guide for integrating your network and application management

It used to work reasonably well to have one team focused on the network and another with an eye on applications and the end user's experience. These days that's a recipe for finger-pointing and ineffective management. Application-aware network performance management can make life easier for users and troubleshooters alike.

- **Tackle the complexities** – keep up with the increasing demands on networks and applications
- **See the big picture** – use application-aware network performance management to gain visibility from the data center to the end user
- **Integrate and analyze** – pull performance data into one place to better identify and isolate the trouble spots
- **Get started** – be sure the solution you choose can really connect all the dots

Fluke Networks offers best-of-breed solutions that help simplify troubleshooting of complex networks with products that provide end-to-end visibility of the IT infrastructure to uncover previously unseen problems. In 2014, Fluke Networks was recognised by Gartner Inc. as a leader in its Network Performance Management and Diagnostics Magic Quadrant, and by EMA Associates as a Value Leader in their 2014 'ANPM Radar' Report.



### Open the book and find:

- How the complexity of today's applications and networks makes troubleshooting extra-troublesome
- How important it is to keep watch over the network from end to end, while also keeping an eye on the user's application experience
- How application-aware network performance management can give you that visibility
- How to get started with picking and implementing your AANPM solution

*Making Everything Easier!™*

Go to [Dummies.com](http://Dummies.com)<sup>®</sup>  
for videos, step-by-step examples,  
how-to articles or to shop!



# **WILEY END USER LICENSE AGREEMENT**

Go to [www.wiley.com/go/eula](http://www.wiley.com/go/eula) to access Wiley's ebook EULA.