

How to Optimize Network Performance While Avoiding Unnecessary Investments

A Network Visibility Guide

What is this guide about?

As the title suggests, the fundamental objective of this guide is to help you optimize the performance of your network so that you're delivering a consistently high quality of service to all networked users within your company or organization. Equally important, this guide is intended to help you get the most out of what you have. Before you consider costly investments to upgrade workgroup server performance or increase bandwidth across the network backbone, you should first ensure that your current network resources are delivering maximum return on investment. But that's easier said than done, right? After all, the corporate enterprise network is a complex web of technologies, systems, applications, and protocols provided by an endless number of vendors. Just getting it to work together as a system is a miracle in itself. How can you ever hope to optimize its performance? The solution to this complex problem is surprisingly simple: visibility. To be able to see into your network, to discover why problems occur and how to solve them — that's the starting point for optimizing network performance. And network visibility is a function of a well-crafted network fault and performance management strategy.

What is fault and performance management?

Fault and performance management are two of the five categories of network management applications defined by the ISO (International Standards Organization). (See figure 1.)

ISO Network Management

Configuration
Security
Accounting
Performance
Fault

Figure 1. The ISO has defined five categories of network management applications. This guide focuses exclusively on fault and performance management.

NH Northeast
Consulting
Resources, Inc



TOTAL NETWORK VISIBILITY

What a Fault and Performance Management System Should Do For You

- Interpret all major LAN and WAN protocols used throughout the enterprise, including encapsulated LAN protocols over leased lines, Frame Relay and X.25.
- Locate problems, identify unique internetwork characteristics and use expert analysis to automatically recommend quick problem resolution. You should be able to convert this information to a text file to appear on a pager or electronic mail file.
- Display complete seven-layer protocol interpretations in clear, concise language.
- Provide internetwork bandwidth statistics that indicate utilization based on LAN protocols as well as end users.
- Embrace an intelligent, scalable client/server architecture to minimize network traffic across a large enterprise. It must also be flexible to support both centralized and multiple monitoring stations throughout the enterprise.
- Automatically learn node names, addresses and connections between existing nodes and protocols.
- Automatically learn routing paths and measure WAN bandwidth utilization.
- Provide comprehensive reports that include historical trend analysis of peak network performance and offer in-depth network management reports that can be used to plan expansion and establish budgets.
- Integrate seamlessly with your existing enterprise-wide management system and consoles such as Hewlett-Packard OpenView, SunNet Manager and IBM NetView.
- Extract data from the various sites and combine them logically for a complete end-to-end view of user traffic patterns and all the devices it traverses.

Fault management has to do with keeping the network up and running. To make sure that happens, network monitoring and analysis applications are deployed to detect, diagnose, and isolate potential problems on the network. Then, through a series of processing functions, these applications help managers determine appropriate corrective action. Performance management tools offer a variety of capabilities, such as baselining your network, analyzing network utilization and trends, and identifying performance bottlenecks. When used in context of a proactive management strategy, these tools will help you make intelligent, well-informed investment decisions in the future when optimizing performance across the enterprise. (See accompanying sidebar for more information on what to look for in fault and performance management solutions.)

Why is this issue of "visibility" into the Enterprise Information Infrastructure so important?

To appreciate the need for visibility, one must also appreciate how the enterprise network has evolved over the last 10 years. Distributed client/server computing has replaced centralized host-based computing as the dominant networking model for the enterprise. While this new model has increased user productivity, it has also created a very complex corporate information system. Today, the network consists of far more than just the physical connectivity and internetworking devices enabling the flow of data. It consists of multiple technologies, multiple applications, multiple systems, and multiple protocols from multiple vendors. The result is a complex new entity called the Enterprise Information Infrastructure (see figure 2). And it comes complete




Figure 2. A comprehensive fault and performance management system must provide visibility into the major components of the Enterprise Information Infrastructure.

with multiple points of potential failure between any two devices communicating within the infrastructure. So unless you have total visibility into this infrastructure, it will be very difficult — if not impossible — to deal with fault and performance management issues.

But that's not the only reason visibility is important. There is also the “mission critical” aspect of the Enterprise Information Infrastructure as a strategic business asset. The global nature of business today, coupled with the reliance upon enterprise network resources, means that large client/server systems must deliver optimum performance and maximum quality of service 24 hours a day, seven days a week. To ensure reliability, availability and peak performance around the clock, organizations are finding it necessary to implement continuous, automated fault and performance monitoring and analysis tools.

Give me an example of the need for fault and performance management.

OK, let's make it real. Let's look at a specific example where visibility is important. The XYZ company has a solid Enterprise Information Infrastructure in place at their corporate headquarters (see figure 3). They've got 2,000 PCs and Macintosh systems assigned to 40 different Ethernet and token ring LAN segments, each connected to Cabletron hubs and internetworked over an FDDI backbone via Cisco routers. They're all running under Novell NetWare, except for a few segments which are using IBM LAN Server or AppleTalk. They've got Sun SparcStation servers running UNIX in half the workgroups and Pentium servers supporting Windows NT in the other half.




Figure 3. With multiple potential points of failure within its enterprise Information Infrastructure, it was difficult for XYZ to isolate a performance problem without total network visibility.

Lotus Notes was deployed a few months ago and finally seems to be stabilized. But last week the folks in MIS rolled out a new Oracle database application across the entire enterprise and with it came a plethora of new problems. Is the Oracle application the cause of the problem? Or have the servers run out of gas? Do they need more bandwidth on each segment? Or is there too much latency caused by router hops? Maybe they made some errors in updating the config.sys files. Who knows? This all too common scenario involves more than a dozen protocols, along with products and technologies from at least 10 different vendors.

So what's the starting point? Where do you go to begin to identify the fault and restore the performance? Maybe the problem is the Oracle application. For corporate developers writing client/server applications, little emphasis has been placed on whether network applications will perform efficiently on the network. A new client/server application might offer acceptable response times in early testing and during its initial roll out in small workgroups. However, as a new application is distributed to a wider user audience and becomes crucial to daily business operations, its response time can erode. And the physical network or server hardware is usually blamed for failing to meet performance expectations. Unless you have visibility into the network — that is, the ability to monitor and analyze application traffic activity as well as physical and data-link statistics — these kinds of fault and performance issues may never be resolved.

Incidentally, after deploying a fault and performance management system the XYZ Company gained the visibility to discover the problem. Each Oracle record, in response to a query, was being sent in a separate Ethernet packet from the server. The transaction needed an acknowledgement from the client before sending the next record. In addition, the data in each record was a mere two to five bytes in size. Although a single Ethernet packet is capable of carrying up to 500 records, only one record at a time was being sent. A simple change to the application to support an increased record payload per packet resulted in a 98 percent improvement.

Isn't RMON sufficient to provide the level of network monitoring I need?

The RMON (Remote Monitoring) standard provides a key component of a comprehensive fault and performance management strategy. Whether it is implemented as a stand-alone hardware probe or embedded into network devices (such as hubs and switches), it's a relatively easy way to monitor network activity at layers 1 and 2 of the OSI (Open Systems Interconnection) model (see figure 4). RMON has its strengths, like the ability to provide segment statistics, long term trending, node traffic pattern analysis, top bandwidth users, alarm reporting, and packet capture. It's a useful tool to identify who caused a network problem and when it happened. And RMON is a very cost-effective solution for monitoring remote office segments. But RMON also has its limitations. For example, while RMON tells you who caused the problem and when, it really doesn't tell you what the problem is or why it happened. Perhaps more important, the current RMON standard addresses only the physical layers of the network — that is, layers 1-2, the basic connectivity of cables, hubs, and related devices. The RMON standard does not provide any meaningful information about the upper layers of the network (layers 3-7). All of these limitations make it difficult for RMON to address the root cause of a problem within the Enterprise Information Infrastructure.

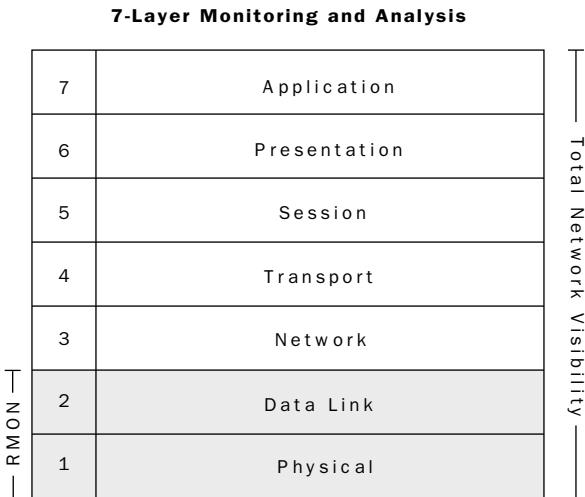


Figure 4. The OSI model organizes network communications into seven different layers. A fault and performance management system should provide visibility into all seven layers. RMON plays an important role yet limited in that system.

This limitation is important to understand because there is a danger in relying exclusively on RMON information to make strategic networking decisions. Too often, managers review these lower-layer management statistics and conclude that the way to solve an application performance problem is to offer more bandwidth or upgrade to more robust hardware. In our story of XYZ company, for example, RMON statistics would not have been able to tell you if the Oracle application really was the cause of the problem. Instead, those statistics might have pointed only to insufficient bandwidth, in which case you might need to upgrade key systems. While additional bandwidth and more powerful hardware is sometimes desirable, this narrowly-focused management approach can be costly, time-consuming, and often misleading. You must have visibility into all seven layers of the OSI model before considering any major investment decisions. Which is why you need to implement a comprehensive fault and performance management system that can help monitor, analyze, and measure client/server application traffic and performance. Using the right system, you can begin to identify the performance problems related to an application's internal design or higher-level protocol configurations.

Why so much concern about visibility into the upper layers of the OSI model?

In many organizations, new client/server applications are first run on a limited test network to measure the performance of the application (but not necessarily its relationship to the network). Once that phase is complete, the new application is rolled out on the Enterprise Information Infrastructure. This is where unexpected problems too often occur, resulting in unacceptable levels of application performance (and, often, degradation of overall network performance). Response times can suffer, and in some cases, the entire program will "time out." When this happens, the finger pointing begins because no one is really sure what's causing the problem. A typical knee-jerk reaction is to upgrade the application servers or selectively increase bandwidth. But that may not be necessary. The best solution is to first use the monitoring and analysis tools necessary to identify, diagnose, and isolate the problem causing

the sudden degradation in network performance. Often, the solution might be something as simple as tuning the autoexec.bat files, config.sys files, login scripts and packet payload size.

More and more organizations are applying these monitoring and analysis tools during both the applications testing and implementation phases to ensure a more thorough understanding of how any new client/server application will impact the performance of the network. Having the ability to perform monitoring and analysis of the upper layers is the only way to gain vital insights into the impact of applications on servers, the number of services being utilized to complete critical tasks, poorly configured protocols or application scripts, or misdirected data requests that make unnecessary use of costly WAN links. Well-defined corporate procedures and policies that guide the efforts of developers and network managers through all phases of client/server application design, development and implementation can enhance productivity while ensuring that applications are deployed on a timely basis.

The accompanying sidebar, *You Can't Fix What You Can't See*, further exemplifies the need for visibility into the upper layers of the OSI model.

What is the best strategy for implementing a fault and performance management system?

The best strategy for you is, of course, dependent on many variables that are unique to your organization, including the number of users on your network, the number of locations supported, and the role of your network in supporting your business. Nevertheless, there are certain guidelines worth following as you plan your fault and performance management strategy. Start by identifying which segments on your network are intended to support mission-critical applications. These applications typically touch specific workgroup segments as well as the network backbone. On these segments, it is important to deploy monitoring and analysis tools that can provide maximum visibility into all seven layers of network traffic. This will enable you to identify and solve problems before they can impact network performance.

You Can't Fix What You Can't See

One large bank solved a nagging problem that had the entire IT staff scratching their heads. The Microsoft Office suite of applications had been centrally installed and administered for a large group of users. Access to Excel and PowerPoint produced reasonable response times, but it took users several minutes to open a Word file. This was creating a serious impact on user productivity.

Servers showed an acceptable utilization rate with plenty of memory to spare. User workstations also indicated that there were no bottlenecks. Finally, a recommendation was made to upgrade one of the 30 workgroup servers to a Pentium-based architecture with 64 megabytes (MB) of RAM to see if performance would improve. Unfortunately, this did not solve the problem.

The real problem was discovered using a fault and performance management system that displayed text data and message types — not just low-level headers. The system revealed what occurred when a user double-clicked the Word icon — the open file command was going to the server only with `D:\Word`. There was no subdirectory or extension specified. The open file command should have specified `D:\MSOFFICE\WINWORD\WORD.COM`.

Since there were 19 subdirectories on each server and MSOFFICE is the 14th, the server was pouring over its entire path and using `Word.com`, `Word.bat`, and `Word.exe` searches three times for each subdirectory. More than 100 searches were being made and each was sent to the desktop operating system to let it know it had not yet found the executable. A simple change to include the subdirectory and extension in each user's login script was made. The Pentium-based server was redeployed and users were more productive.

To complement these tools, RMON agents and probes can be cost-effectively deployed on a variety of local and remote segments, providing continuous operation and information on lower-layer statistics and alarms. The combination of RMON and seven-layer data collectors can provide end-to-end visibility into the Enterprise Information Infrastructure (see figure 5).

Distributing RMON and seven-layer data collection intelligence to critical points on your network represents only part of the strategy. Equally important is the issue of what you do with the data after it has been collected. Information from these collectors are usually directed to a variety of management applications for action by the network manager or administrator. These applications make it easier to identify, diagnose, isolate, and resolve fault and performance management problems that arise during the course of a day. In addition, they can perform more detailed analysis on longer term network traffic and create reports that help you understand normal network traffic patterns, monitor the health of the network, and establish a network performance baseline.

Fault and Performance management applications may be launched from Integrated Network Management Systems (INMS) such as HP OpenView, SunNet Manager, and IBM NetView/6000. While the INMS does not perform any significant fault and performance management functions on its own, it does provide a useful framework for integrating dedicated fault and performance management applications. But even having all the right applications in place doesn't necessarily guarantee the ability to solve any problems within the Enterprise Information Infrastructure.




Figure 5. In formulating a fault and performance management strategy, consider deploying seven-layer monitoring and analysis tools in the backbone and mission critical segments, with RMON providing coverage on other segments.

The final, and perhaps most important, piece of the strategy is the investment that should be made in management expertise to ensure that the staff can take advantage of the visibility they now have into the network. This means that key members of the network management staff should be trained to deal with both proactive and reactive fault and performance management issues.

We can help.

Clearly, a strategic approach to Enterprise Information Infrastructure management requires an investment in time and resources. But evidence from companies that have made this investment indicate a relatively short payback period and a significant long term ROI (Return on Investment). (See accompanying sidebar, The Big Payback, for an example.) For more information on the ROI associated with fault and performance management, please contact NCRI or your authorized Network General sales representative.

What kind of process should I put in place to ensure network performance optimization?

This is probably the most important issue facing network managers today. Let's start by looking at phases of the network management lifecycle as they relate to the deployment of new systems, databases, applications and technologies onto the Enterprise Information Infrastructure. There are three main phases: planning, pre-staging, and implementation (see figure 6).




Figure 6. The three phases of the network management lifecycle represent a complementary ongoing process, critical to optimizing performance of the Enterprise Information Infrastructure.

The Big Payback

For a large national trucking company, rolling out a new client/server application before testing its performance over a WAN caused a complete breakdown in schedules and deliveries to customers. As a result, the trucking company had to reimburse customers for more than \$100,000 worth of perishable goods that were undelivered the first days of production operation.

After the trucking company successfully piloted its new client/server-based scheduling and delivery system in a regional office, the decision was made to implement it nationally on the same day. This would hopefully allow all scheduling and loading facilities to coordinate their delivery commitments across the country. However, on the deployment day no regional offices could access the corporate scheduling server. The application kept failing because the IT organization did not consider the impact of WAN lines on response times.

Developers immediately called for upgrading all WAN links from 64 Kilobytes per second (KB/s) to T-1 speeds. At the same time, someone in the corporate IT group had invested in fault and performance management tools to view application-layer client/server interactions. Using these tools, it was quickly determined that the length of time the server waited for an acknowledgement from the client was too short. By simply changing a single parameter — one that instructed the server wait longer to compensate for minor latencies in the multiple-hop network — there was a dramatic improvement in response time.

Utilizing the right monitoring and analysis tools, procedures and expertise saved the trucking company more than \$20,000 per month on WAN links alone. And the company had its deliv-

Planning – This is probably the most important, yet most often ignored, phase in the process. Why? Because planning takes time, and time is a network manager's most precious commodity. Nevertheless, a network without a plan is most likely a network in a constant state of chaos. At a minimum, you should spend enough time to understand the normal behavior of your network, including issues like network utilization, top users, and server performance. Collectively, this data — which is typically gathered over a two- or three-week period — will help you establish a “baseline” of your network's performance. From there, you will be able to establish performance thresholds, so that if a particular network variable exceeds the thresholds you've established you will be notified of the problem and be able to resolve it before it impacts the performance of the network.

That's very basic planning. If you've got more time, you will want to consider employing more advanced planning tools, such as capacity planning or network modeling. These tools allow you to take advantage of historical trend data to predict the future. But don't stop there. Take time to consider all relevant issues, including user behavior, new application traffic patterns, server placement, future applications, user training, management needs, and business drivers. This holistic approach to planning will enable you to better predict future requirements for the Enterprise Information Infrastructure.

Pre-Staging – The first step in pre-staging the deployment of any new system, application, database, or technology is to set up a controlled test network environment that closely emulates your Enterprise Information Infrastructure. Obviously, there will be limitations, but try to include as many production network protocols in your test network as is appropriate. As you introduce new variables into this otherwise controlled test environment, determine the interaction and impact of these variables on the network. Pay particular attention to application design, as indicated by such problems as too many read/writes, improper packet size, or poor Windows settings. An inefficient application design will bring a network to its knees very quickly. A seven-layer analysis tool will prove invaluable in helping you gain visibility into your test environment and the problems that may arise from any new variables. With this information, you will be able resolve problems and optimize performance so that the new system, application, database, or technology can be safely deployed within your Enterprise Information Infrastructure.

Implementation – Unfortunately, limitations of time, budget, and resources too often force network managers to focus almost exclusively on the implementation phase. This is where the network manager typically plays the role of “firefighter,” reacting quickly to isolated problems and solving them as quickly as possible before they bring down the entire network. While this reactive form of management will always be a reality, it can be minimized in two ways: first, by paying attention to the planning and testing phases outlined above and, second, by deploying a comprehensive fault and performance management strategy like the one discussed earlier. The key is to have complete visibility into the Enterprise Information Infrastructure and thus be able to “see” the source of potential performance problems before they become real performance problems.

It’s important to understand that the Enterprise Information Infrastructure is a constantly-changing entity. Which means that the three phases of the network management lifecycle discussed here must be continually reviewed. The process of fine tuning your network for optimum performance is an iterative, ongoing process.

What if there’s a more serious network performance problem?

If performance problems persist after you model application traffic and optimize the network configuration, then you might need to consider upgrading equipment or providing additional bandwidth. For example, many organizations start by modeling the impact of replacing shared-media hubs with switched hubs at strategic locations within the network. This gives a clear idea of how mission-critical application servers would perform with 10 Mb/s or more of dedicated bandwidth. It also enables you to determine whether a switched backbone hub providing dedicated connections to multiple shared workgroup hubs would improve performance. Some of the advanced planning tools and techniques described earlier would be very useful in helping you better understand the most effective solution to these kinds of issues. If redesigning a network and employing new technologies still do not solve performance problems, it may be necessary to rewrite applications to make better use of the existing information architecture. Depending on the design of the applications, network speeds, and user and application demographics, you might need to offload processing from a server or client. This gets into areas of enterprise network architecture that may exceed the

The Next Steps

So where do you go from here? Here are six steps to start creating a fault and performance management strategy and process that's right for you:

- **Build the Team** – Talk to all parties that should be involved in the planning — end-users, applications developers, system managers, and network managers. It's important to get buy-in and cooperation as you formulate strategy and build a solution.
- **Set Realistic Goals** – At first, it may not be possible to find the perfect solution. In that case, don't delay implementating an interim solution waiting to hit on the "silver bullet". Start with understanding the performance of your most important applications and segments.
- **Get the Facts** – Gather hard data from controlled experiments. Investigate and understand the critical tasks. Then make decisions based on facts rather than intuition. As a final measure, look beyond just the facts to really understand "why".
- **Review Your Options** – Once you've got the data in hand, work within the team to review and discuss all options before changing the network, protocols, or applications.
- **Develop a Plan** – Any modification to one component of the Enterprise Information Infrastructure will impact other components. Which means you need to develop a detailed short-term and long-term plan that considers these variables and has contingencies in case things go wrong.
- **Partner Wisely** – After building the team, setting your goals, collecting the facts, reviewing your options and developing a plan, you may also need some outside help. When looking to partner with a provider of products and/or services, understand exactly where you are and where you want to be. Partner with a company directly experienced in meeting objectives similar to yours.

capabilities you have on staff. In that case, it would be useful to consult with an organization that understands these complex design issues and can complement the expertise and resources you have on staff.

Bottom line: What are the real benefits associated with fault and performance management?

By identifying the precise cause and location of network performance problems — as well as potential solutions — a fault and performance management system will help you avoid making costly, unnecessary investments. For example, what appears to be a costly hardware or infrastructure problem might actually be a minor problem, requiring only a modification to an application script, protocol setting or routing table. The only way you'll know is by deploying a fault and performance management system that can provide visibility into all seven layers of the OSI reference model. Using this system proactively, you'll be able to accurately detect whether a problem is caused by an application, protocol software, network device, or WAN link. And you'll be able to take prompt corrective action to optimize performance anywhere on the enterprise.

By combining today's best fault and performance management technologies, the knowledge and skills of a seasoned network management staff, and the right set of proactive and reactive procedures, you gain control over the Enterprise Information Infrastructure. Optimizing the performance of your network, in turn, enables you to deliver a consistently high quality of service to all networked users within your company or organization. Consequently, you will spend less time dealing with fault and performance problems in a reactive mode, and have more time to address management issues from more of a strategic, proactive position.

Ideally, the time to invest in a fault and performance management system and strategy is *before* new client/server applications are deployed on the Enterprise Information Infrastructure. By developing procedures that encourage software developers and network managers to work closely together during the planning and staging phases of the management lifecycle, the result will be optimum network performance and high quality of service. And that will mean higher end-user productivity.

This *How to Optimize Network Performance While Avoiding Unnecessary Investments* guide was co-authored by Northeast Consulting Resources, Inc. (NCRI) and Network General Corporation to help you increase your network's performance and maximize the return on your network investment.

Northeast Consulting Resources, Inc. is a Boston-based consulting firm specializing in the application of information technology to the changing business environment. NCRI helps leading organizations with many facets of their IT strategy, including planning, advanced network architectures; establishing efficient and effective enterprise network management systems; application analysis and planning; and leveraging IT capabilities through outsourcing.

Network General is a worldwide leader in enterprise fault and performance management solutions. Network General is committed to providing businesses with Total Network Visibility™ — the ability to view the entire information enterprise from end to end and top to bottom within all seven layers of the OSI model. This is critical to optimizing network performance and ensuring a consistent quality of service to users. This visibility is delivered through a family of Sniffer Analyzer tools and systems that incorporate Experience Technology™, Network General's proprietary technology based on 10 years of internal R&D, as well as long-term working relationships with leading networking vendors and network management professionals. In addition, Network General offers a complete line of Foundation Manager/RMON solutions, reporting applications, and service options. Network General products are used by all of the Fortune 50 companies and more than 80% of the Fortune 500. Products are used in 40 countries and are available in several languages.

For more information, contact your authorized Network General sales representative or call 1-800-SNIFFER (1-800-764-3337).

GSA Schedule Number—GSOOK92AGS61Ø9 PSØ8

Network General and Sniffer are registered trademarks of Network General and/or its wholly owned subsidiaries. All other registered and unregistered trademarks in this document are the sole property of their respective owners. All specifications may be changed without notice.

©1995 Network General Corporation. All rights reserved. P/N 24152-01 8/95