



Monitoring Application Response Time

APP NOTE

WHAT IS CAUSING POOR APPLICATION PERFORMANCE? IS IT THE NETWORK? THE SERVER? THE APPLICATION? A ROGUE APPLICATION OR USER HOGGING ALL THE BANDWIDTH? HOW CAN YOU QUICKLY TELL THE DIFFERENCE? CAN YOU PROVE THAT IT'S NOT THE NETWORK TO OTHER MANAGEMENT TEAMS? HOW DO YOU DEFEND AGAINST "THE NETWORK IS SLOW" CLAIMS?

Defending the Network

Regardless of the real cause of poor application performance, the first thing a network manager hears is "the network is slow." That's certainly the default assumption of end users. But it's also what the network staff hears from other management silos—the server and application teams—when their tools can't identify the source of an application response time problem. And, all too often, the network team has no way to "defend the network."

The problem is that none of the management systems that enterprises currently rely on can localize the cause of application performance problems. That requires visibility of all the components of application response, which none of these systems can furnish. Data center management systems, which monitor various aspects of server and back-end performance, can't see anything but the server components of end-to-end application response. Private network management software relies on polling or traps from individual devices, so unless the problem lies in a particular network device (i.e. a "hard" fault or failure), they reveal nothing and can't be used to defend the network. And application response time management systems can see end-to-end response, but can't drill down to pinpoint where the problem is: which user, server, network device, or application. Furthermore, they are generally limited in the number of applications they can monitor and the number of locations they can monitor them from.



Customer Problem:

- > Unable to quickly distinguish server, network, and application problems
- > Too much time wasted defending against "the network is slow" claims
- > Fingerpointing and lack of cooperation with other management silos

Network Physics Solution:

- > Flow-based non-invasive monitoring of all applications, all the time, in real time
- > Visibility into every component of application response for every application and every user
- > Correlation between performance and utilization


Customer Benefits:

- > Quick identification of source of application response problems: server, network, application
- > Much faster problem isolation, less guesswork
- > Greatly improved IT staff productivity, better cooperation between management silos



Go with the Flow

The Network Physics appliance solves this problem without agents, SNMP, or synthetic transactions. Simply and non-intrusively installed via spanning port or tap, its flow-based technology reveals the various components of application response for every application and every user, in real time.

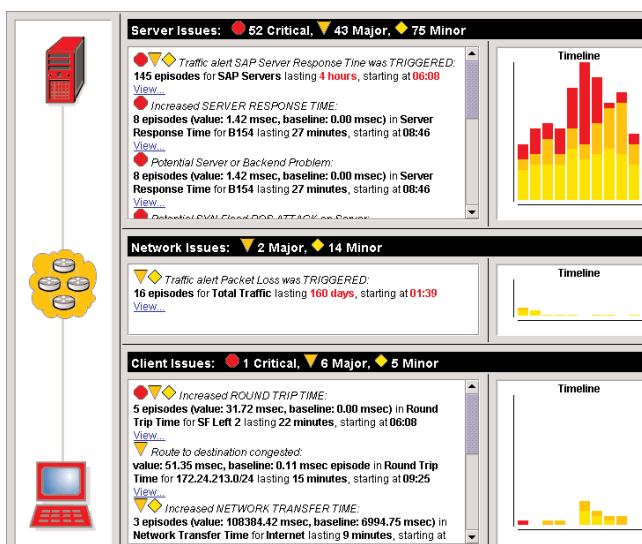


The Network Physics Response Time Composition Chart breaks down end-to-end application response into its component parts: TCP setup time, server response time, network time, time lost to packet loss, and latency, enabling network managers to very quickly distinguish between server, network, and application problems.

For example, the figure above displays the progress of the January 2003 Slammer virus attack on a customer's SAP server farm. At the outset, SAP application response is dominated by Connection Setup Time and Server Response Time, indicating the impact of the infection on server functions. Then, after the SAP server farm was cleared of the infestation, the end-to-end response time is dominated by Network Transfer Time and Outbound Retransmission Delay, indicating the continuing impact on the network of other infected server farms that are using up bandwidth replicating the virus.

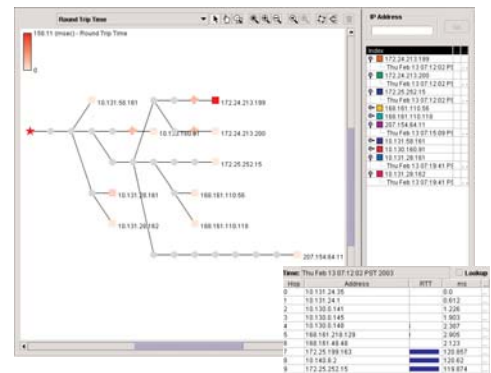
The Network Physics Response Time Composition Chart gathers all of this information—for any grouping of traffic, from individual IP addresses to business groups aggregating traffic with a common business significance, such as a server farm—into one easily interpreted display, enabling the rapid localization of application response time problems.

Furthermore, the Network Physics Problem Management Dashboard automatically correlates these metrics to deliver automatic alerts identifying the severity and likely source of application response time problems, without programming.



For instance, if the Connection Setup Time (the time to complete the three-way TCP handshake that establishes a connection) rises but the Connection Rate (which correlates to the number of application users) does not, the Dashboard will alert the user to a possible problem with the front-end server. Likewise, a rise in Server Response Time (the time it takes the server to respond to the client's request for data after the connection is established) rises without a concomitant rise in connections, this indicates a possible server back-end problem. In either case, the network team can turn the problem over to the appropriate team without wasting time trying to diagnose something that is not their responsibility.

And, if the problem is network-based, drilldown links from the Dashboard enable managers to quickly localize the problem, even in networks they do not control, using a sophisticated topological display, by correlating performance and utilization metrics to reveal rogue application or users, and other advanced techniques.



By correlating flow and traceroute data, the Network Physics topological display provides a graphical indication of route quality with drilldown to the latency contributed by specific nodes.

The Problem Management Dashboard applies a number of advanced statistical and correlative techniques to automatically flag problems, identify their likely origin, and allow further, focused drill-down analysis and resolution.