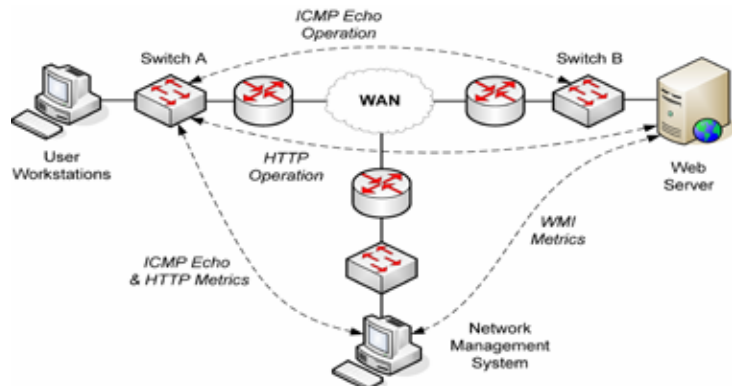


## IP SLA, WMI and Service Health

When a user calls to complain that a service is slow, how easily can you diagnose the problem? Net-centric services like email or a web-based application can involve multiple services and devices, forcing network managers down a checklist of possibilities: Is it the web service itself? The DNS service? The web server? The WAN link?

Most often, services like a web application have two “layers” of operation: the performance of the application elements (servers, app, desktops) and the connections between them (routers, switches, links, etc.). Different metrics in each zone complicates troubleshooting, so the more you can do to simplify and normalize (or at least integrate) those metrics, the faster you’ll eliminate the problem. This creates a challenge if your only troubleshooting tool is ping, but for environments with Cisco equipment, using IP SLA™ instrumentation can make life a lot easier. And if your apps are running in a Microsoft environment, that’s even better because you can leverage WMI instrumentation as well.



**Figure 1:** In a typical configuration for a net-centric application like a web service, the web server hosts the HTTP and DNS services and MS SQL Server, and the switches and routers enable WAN connectivity for servers and workstations.

## Using IP SLA to Access the User Experience

IP SLA (Internet Protocol Service Level Agreements) is embedded in the Cisco IOS (Internet Operating System) for most Cisco routers and switches. IP SLA operations can measure delay (round trip time), jitter, packet loss, connectivity, voice quality scores, and many other key metrics for monitoring and troubleshooting network elements.

Additionally, threshold levels can be set for most metrics. When a metric crosses a threshold level, IP SLA sends an SNMP trap to the specified IP addresses.

You can configure an IP SLA HTTP operation to monitor the overall user experience for the “connectivity layer” of a web application (or any other net-centered application such as email, VoIP or videoconferencing). This operation uses a synthetic web transaction to measure the total round trip time (RTT) to perform a DNS query, establish a TCP connection to the HTTP service, and retrieve the web site’s home page. By configuring the HTTP operation on the LAN switch closest to users, the total RTT (or latency) is an accurate measure of the users’ experience (as opposed to measuring RTT from a central network management server).

Next, configure an IP SLA ICMP Echo operation to monitor RTT between the switch on the user LAN and the switch to which the web server is connected. This way, if the HTTP operation indicates the web transaction is slow or unresponsive, you can check the WAN RTT between the switches to see whether the problem is related to the WAN link or something on the web server.

## Watching the Apps and Servers: Adding WMI

WMI (Windows Management Instrumentation) is an instrumentation tool similar to IP SLA that Microsoft has created for its products. WMI provides thousands of performance metrics for applications such as MS Exchange and MS SQL Server, as well as for server hardware and operating system components. Combining WMI with IP SLA provides performance information about both layers and gives an end-to-end view of your web app or other net-centered service.

Microsoft has a built-in performance administration tool for monitoring WMI data for applications and servers. Using the tool you can view each server’s CPU utilization, physical memory and free disk space. Each of these sub systems are critical to the server’s performance regardless of the application running. Lack of memory, CPU cycles and low disk space are common causes of slowdowns on a server. You’ll have to go into each server to view the individual performance counters or you can use dopplerVUE to simplify the process by collecting any of the thousands of available WMI counters from across multiple servers.

## Getting the End-To-End View

Because dopplerVUE integrates fault and performance data from a variety of sources, including SNMP, syslog, WMI and IP SLA, you can integrate metrics from both layers of a web service into a single end-to-end dashboard view.

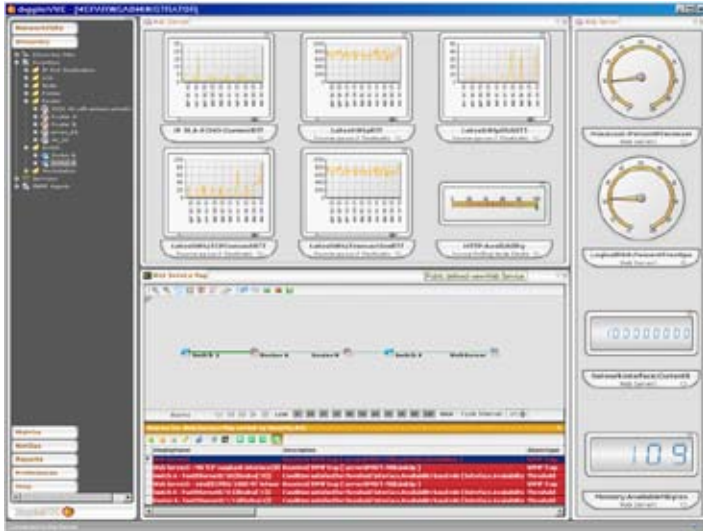


Figure 2: dopplerVUE lets you easily create custom views for critical services, here combining WMI and IP SLA metrics for a web application, including Total RTT for the DNS look up, TCP connection and web page transaction.

Using dopplerVUE's drag-and-drop interface, you can quickly create an integrated view of both layers of the service without having to shift between tools or viewers. Figure 2 is an example that shows all of the key WMI and IP SLA metrics involved using line charts and gauges, alarms for each element of the service, and a map that includes both color-coded alarm status for each element as well as the Doppler effect that color-codes each element's performance health.

Since dopplerVUE automatically creates IP SLA polling jobs for all discovered IP SLA devices, you just need to set those jobs to collect the HTTP and ICMP Echo metrics. It's also easy to use the built-in wizards to create a custom WMI polling job for the web server.

Additionally, you can also have dopplerVUE automatically notify you whenever a critical performance measure for the service exceeds acceptable limits. For example, you could use the built-in wizard to create a rule that sends an email whenever the HTTP RTT exceeds 800 ms or if the web server performance is suffering in any way.

---

## About dopplerVUE

dopplerVUE is a powerful yet easy to use network management solution for managing up to 5000 network elements. To see how dopplerVUE can solve your network management needs, visit to [www.dopplerVUE.com](http://www.dopplerVUE.com). For more information call 888-dvue-now (888-388-3669).

© 2008 SYS Technologies. All Rights Reserved. All trade names referenced are the service mark, trademark or registered trademark of the respective manufacturer.