Progress® | ipswitch®

A PROGRESS PROFESSIONAL GUIDE

IT Pro's Guide to Faster Troubleshooting

Gaining more control in an increasingly complex networking world.

Introduction

This ebook addresses challenges and best practices for troubleshooting applications, servers and networks.

Here's what you will learn:

- The challenges of pinpointing performance problems
- Preventing alert storms
- Managing Service Level Agreements
- Application root cause analysis
- Wireless root cause analysis
- Why unified monitoring matters

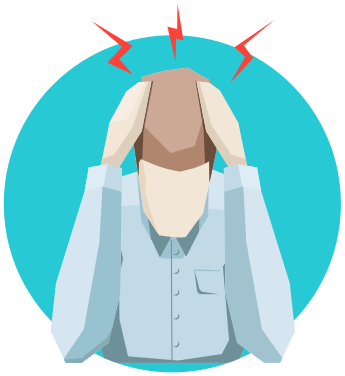

The Challenges of Pinpointing Performance Problems

The biggest headache for IT is dealing with intermittent performance problems. These are those problems that make themselves apparent and disappear before you can identify the source, only to happen again and again, but randomly. In most cases, these intermittent performance problems look like they are rooted in a certain area of your network where in fact they are stemming from a completely different one. Let's take Tom the hardworking sysadmin, for example:

It's Tuesday morning and Tom starts getting a bunch of tickets from staff complaining about emails getting trapped in their outboxes.

Examining the error logs on the Exchange server shows too many requests are being received at once. Complaints from users are stacking up. Emotions are running higher and higher. Tom figures his 8-year-old server is to blame and decides to buy a brand new one.

That weekend Tom installs the new server. He spends 8 hours running the server in a test environment to make sure all is working well. It looks like all systems are a "go". The new server goes live. So far, so good.

Monday morning arrives and everything is running smoothly for a few hours when all of a sudden Tom starts receiving an influx of tickets complaining about the same issue that led him to buy the server the week before.

Even worse, the log files on the new Exchange server show the same errors as those on the old server. Not only has Tom spent $5000 on a server he didn't need but he's also lost several hours of work trying to fix a problem that didn't need fixing.

That afternoon the source of the problem is finally discovered, about a week after it first made itself known. The culprit? A buggy plug-in that a consultant installed on Outlook three weeks before. How much downtime was incurred? Too much.

Scroll through any IT forum like Spiceworks and you'll always see stories like this. We don't like to admit it as IT pros, but if you've been in the field long enough, something similar to Tom's experience has happened to you at least once.

If Tom had a way to monitor his network end-to-end he would have easily been able to make the connection between the buggy Outlook plugin and the overload of server requests. Time and money would not have been wasted.

Tom's nightmare is a perfect example of intermittent performance problem on the network affecting business and user productivity. We polled more than 400 IT pros to determine the most common sources of intermittent performance problems.

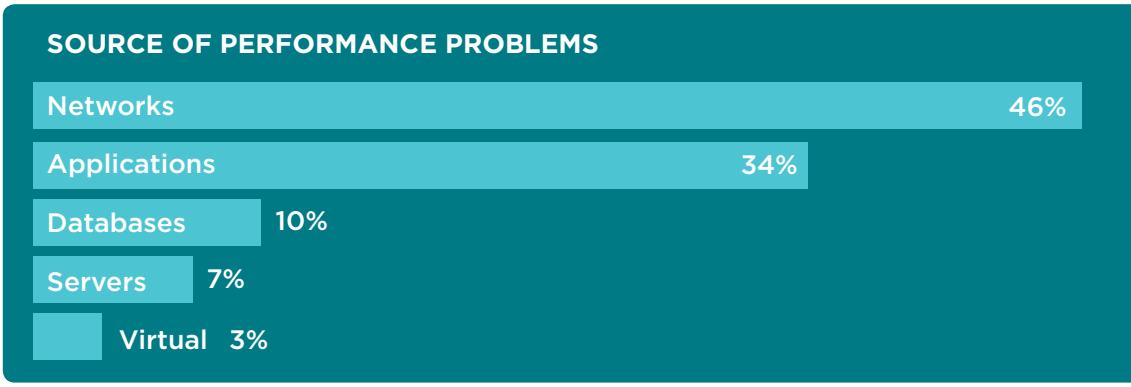

Networks, which includes switches, routers, and firewalls accounted for nearly half of intermittent performance issues. Applications weren't far behind, netting slightly more than one-third of all responses.

Next, we asked the same IT pros how much time it took to find and fix intermittent performance problems.

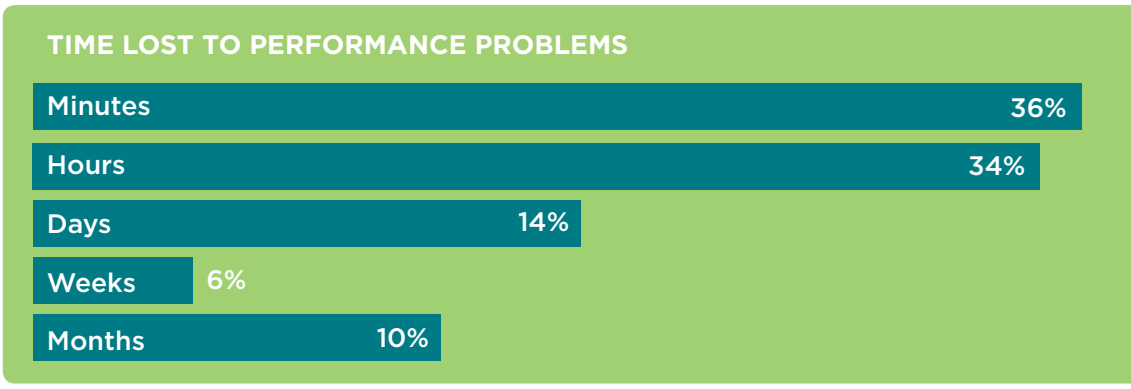

More than one-third of the IT pros surveyed were able to fix intermittent performance issues within minutes. However, almost the same number of respondents spent hours finding the source of the problem, others taking days and even months to resolve. The harder the problem is to find, the more downtime can accumulate over the course of a year.

So, how do you go about combatting intermittent performance problems before they become serious issue? The answer is arming your IT team with a flexible IT monitoring tool that lets them be proactive instead of reactive. The tool will ideally provide an early warning system when an issue starts to arise that could lead to unhappy users and downtime.

Preventing Alert Storms

In larger networks, sysadmins will daisy-chain multiple switches otherwise known as “cascading”. A failed switch at the head of a chain will generate hundreds of unnecessary alerts throughout the chain. We call this an alert storm.

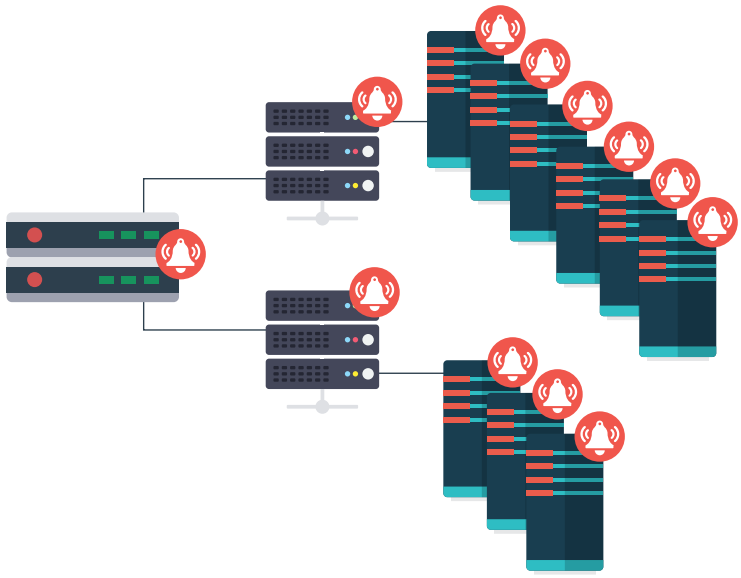

Alert storms can overwhelm an IT team and waste several hours of valuable time. Some storms involve the generation of several thousand alerts within a single hour.

An IT monitoring tool should identify network dependencies to automatically suppress redundant alerts. In other words, it would know what devices are connected to the failed switch and only issue an alert for the failed switch, suppressing all others.

Know Your Network

When an alert is issued, the first thing you want to see are the network maps, often considered the most valuable diagnostic tool. Being able to visualize your network can save hours, and even days troubleshooting problems. However, when your server closet becomes a repository for tangled wiring, problems take long to resolve.

An IT monitoring tool should be able to discover layer 2 and layer 3 network information to automatically generate maps, a great first response tool that enables you to visualize the network. Network maps provide an orderly representation of the server closet or data center, and dynamically display up-to-date device status.

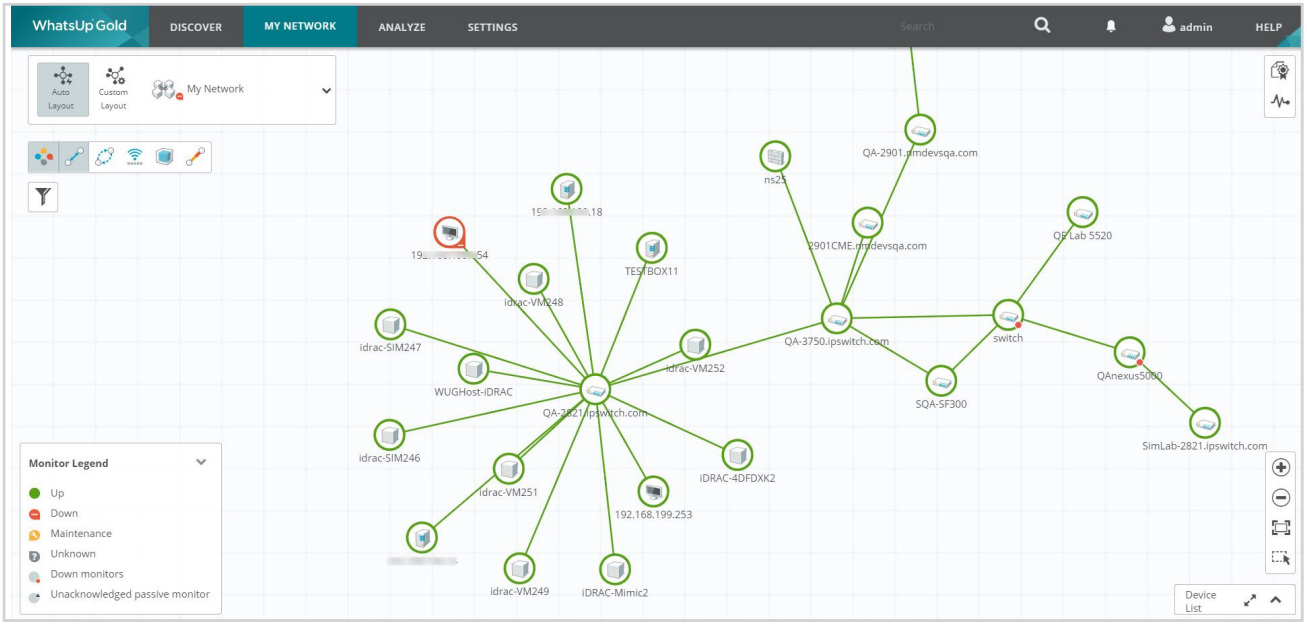

Managing Service Level Agreements

The importance of network performance management has been emphasized by the assignment of SLAs. Based on a series of recent Ipswitch surveys on social media, about half of all IT teams are being held contractually liable for downtime in the form of an SLA.

SLAs are often factored into an IT team's compensation plans which means that not meeting them can negatively affect morale, and ultimately, IT staff turnover. An SLA isn't always a bad thing, though. They are a great way to demonstrate the improvements you and your IT team have made to employee productivity and network performance.

Even a few minutes of downtime can have serious repercussions, considering that more than two nines (>99%) of uptime have become the SLA threshold for many IT teams. A few minutes here and there can add up to hours over the course of the year.

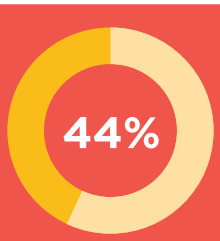

NEARLY HALF OF IT TEAMS DON'T HAVE SLAs!

It's rather striking that so many IT teams don't have SLAs in place, considering how downtime can hurt their business operations.

Reaching agreement with your managers on how many nines your SLA will factor will determine what kind of leeway you will get when downtime does occur. Like any occupational contract, when it comes to SLAs the devil is always in the details.

There are 3 major challenges that SLAs present to IT:

- Determining what factors count the most when determining the number of 9's
- Negotiating reasonable numbers with line of business owners and executives
- Visibility across the entire IT service stack in order to respond to problems quickly

Once an SLA is established and goes into effect, how are you going to be able to meet it? One place to start is to identify the things that waste your time the most that keep you from supporting and educating staff.

SLA LEVEL	ACCEPTABLE UNPLANNED DOWNTIME
< Two nines (<99%)	
Three nines (99.9%)	9 hours / year
Four nines (99.99%)	52 minutes / year
Five nines (99.999%)	Five minutes / year

Application Root Cause Analysis

Remember Tom and his shiny new server? If he had an IT monitoring tool he would have identified the problem in Outlook quickly. No new servers required.

Taking a look at a historical graph of internet information services (IIS) performance against an SLA with three nines (99.9% uptime), we can see that the IIS front end went into a warning state multiple times over three days. This is something that warrants examination before it goes down and your SLAs take a hit.

If you can drill down even deeper to analyze that historical status of a specific monitor, the data can be correlated to helpdesk events to further isolate intermittent performance problems.

Controlling Application States

Applications are dependent on enabling technologies like web servers and databases. They are also dependent on other applications. Some IT monitoring tools enable IT managers to define and monitor these dependencies when assessing the state of an application.

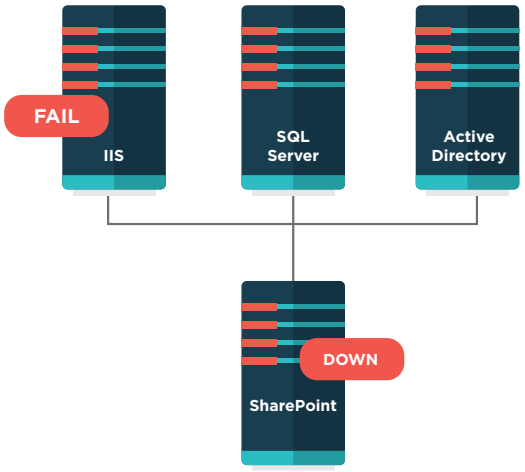

If IIS fails, it can no longer serve up SharePoint webpages. From a user's perspective SharePoint is down.

An IT monitoring system can support several application states– the up state, the down state, the warning state and the maintenance state. This allows IT to define an application state by assigning threshold values to monitored performance metrics.

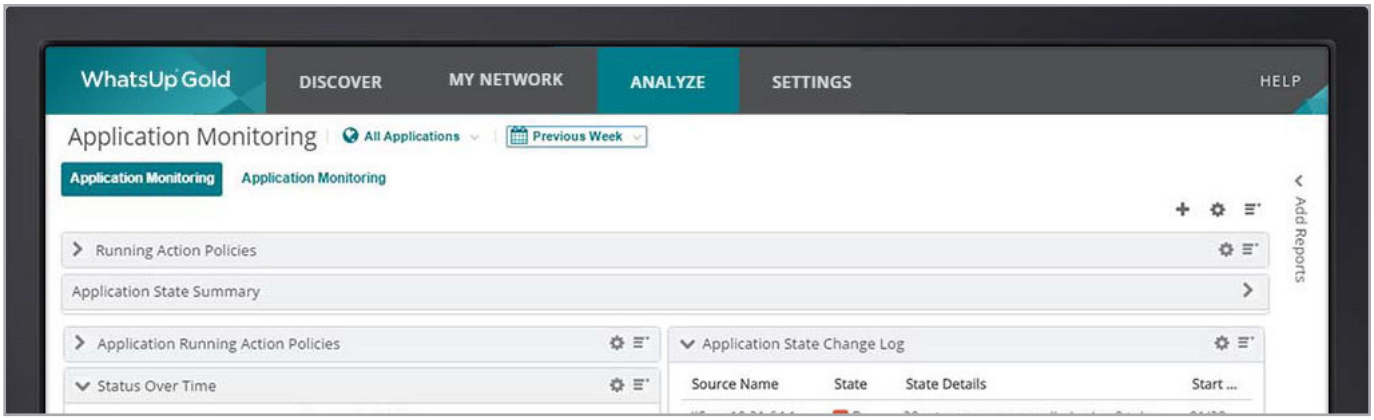

Another example, when CPU utilization for a process exceeds 75% on a server, this should put the application in the warning state. When CPU utilization exceeds 90%, IT should be alerted that the application is in the down state. This provides IT managers with early warning and ample time to respond to performance problems before impacting users and the business.

Wireless Root Cause Analysis

When it comes to wireless networks, displaying wireless LAN controllers (WLC), access points and clients is crucial. These maps should get updated with every polling cycle to show new clients as they log onto the wireless network.

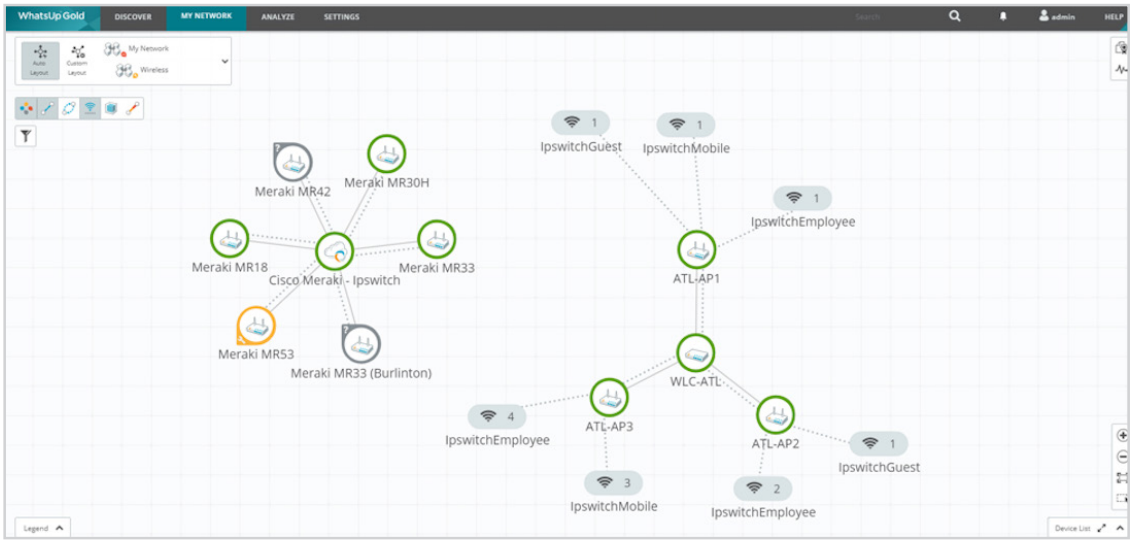

When a wireless network end user complains about performance, a wireless network map lets you follow the connection between the client, access point and WLC. You can also see all the other clients connected to the same access point, possibly indicating an oversubscription problem.

The first question you should ask when a network issue arises is: "*Do I have an access point capacity problem?*"

Historical data should be collected and presented in a way that exposes patterns in client count and bandwidth usage. This allows you to correlate graphs to the time when a performance problem was reported. By analyzing patterns in the number of clients connected to an access point, and the corresponding bandwidth usage, you can determine if the access point can handle wireless volumes at peak usage.

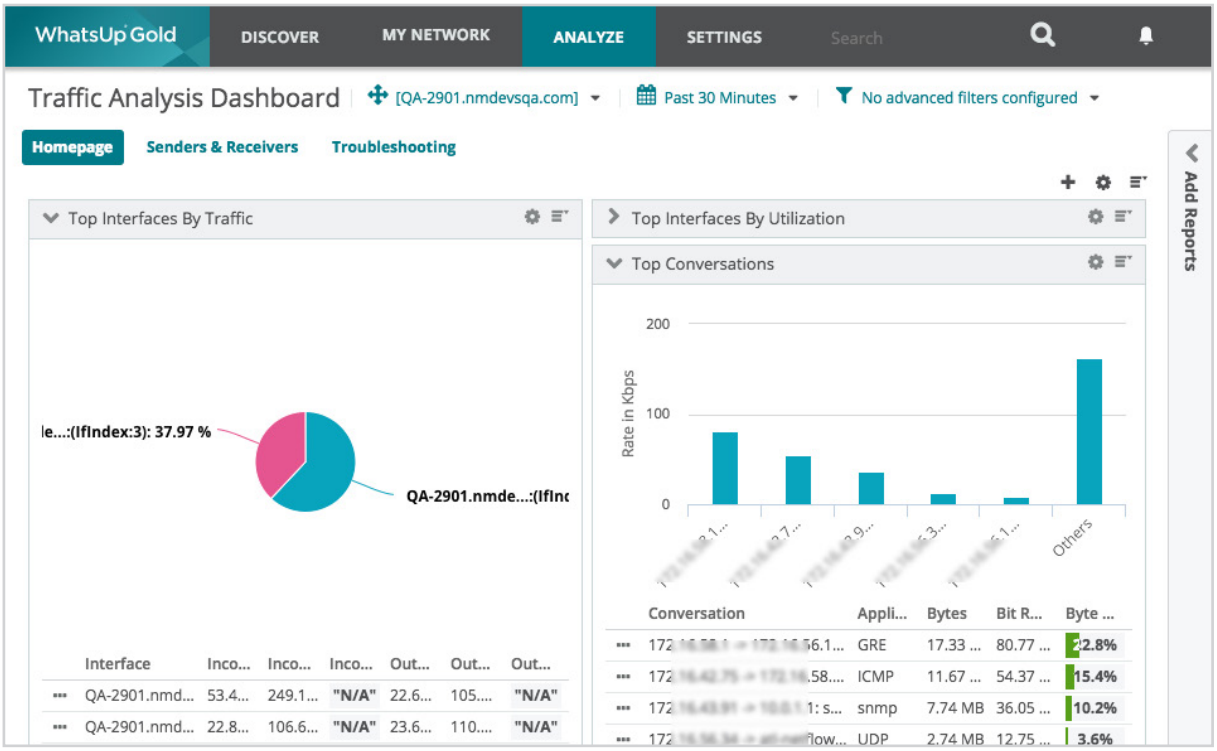

Now that you know that your wireless access point capacity is sufficient, you may want to dig deeper to see if you have a WLC capacity problem.

Historical graphs covering WLC, CPU, and memory utilization should also be viewed in multiple time measurements to expose patterns that can be correlated to the timing of reported performance problems. High utilization of either of these resources indicates that a WLC can't keep up with peak usage on the wireless network.

If you are satisfied with the wireless capacity your users have – but they aren't – you may want to ask yourself if you have a signal strength problem.

Taking a look at historical data in your IT monitoring tool on (SNR) and (RSSI) will expose patterns of excessive noise that is interfering with the wireless signal.

Since last we saw our friend, Tom, he has now implemented an IT monitoring tool to help him make more informed decisions. Now he's experiencing wireless performance problems. He and his team determined that the problem was experienced by end users attached to an access point located near a large conference room. Upon reviewing the historical graphs showing clients, bandwidth usage and resource utilization, his team concluded there was enough wireless capacity.

Tom then turned his attention to wireless signal strength and analyzed historical SNR graphs. This exposed several noticeable patterns:

- Monthly view: problem was two weeks old
- Weekly view: the SNR degraded almost daily between 11:30am and 1:30pm
- Hourly view: SNR degraded in 1 to 2 minute increments with no discernible pattern

Using the correlation between views, Tom identified root cause of the problem: noise generated from an old microwave oven in a nearby employee kitchen. Even legacy systems like kitchen appliances can affect network performance.

Covering the Complexity of Your IT Stack

If you're an IT admin, you now need to worry about your entire stack when delivering just a single service or application due to the complexity of the modern day business infrastructure. Not knowing how any given service can affect your network is like walking blindfolded off a cliff.

For instance, a single bug within an unsuspecting application could be causing memory leaks on the server side, resulting in everyone who is accessing that server running into resource issues. Your users may translate this as their computer being slow. Time can't be wasted on finding these issues, because that time will be needed for troubleshooting.

There are a multitude of locations where an error can cause downtime. If one part of your IT stack goes down a chain-reaction can bring down every other subsequent layer. The more time you spend searching for the root cause of the problem, the more you eat into your SLA.

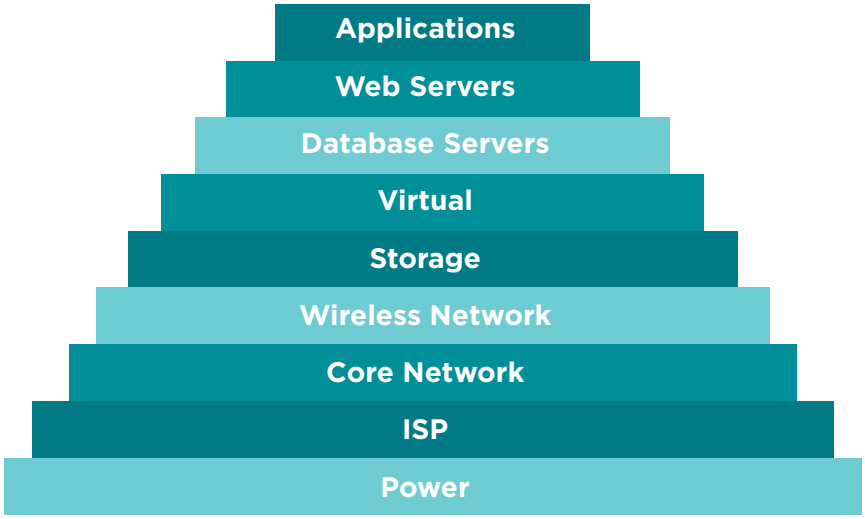

Why Unified Monitoring Matters

Solving performance problems on the business network and delivering on SLAs for critical business services is a race against time. To win the race against time when intermittent performance problems threaten your SLAs, IT needs to have visibility across all the silos in the IT infrastructure.

A unified IT monitoring solution that can cover the full surface area of your IT infrastructure should be able to perform the following:

- Device discovery
- Infrastructure mapping
- Monitor networks, applications and servers
- Proactive alerts
- Reporting mechanisms
- Simplistic Licensing Model

Unified IT monitoring starts with discovering all the devices on your network as well as the connectivity and dependencies between them. Network dependencies are important when it comes to preventing a flurry of alerts that don't do much more than cause a lot of noise.

IT monitoring tools poll devices for data at regular intervals, otherwise known as polling cycles. Should a switch fail, the network monitor can no longer poll it and, as a result, an alert is issued.

However, when a switch fails, the network monitor can't talk to all the devices connected to it and might think those devices are down as well. For example, on a 48-port switch, the network monitor could issue up to 48 alerts at once.

88%

> WANT IT MANAGEMENT SOFTWARE THAT OFFERS MORE MONITORING FLEXIBILITY, WITH FEWER LICENSING RESTRICTIONS.

Respondents overwhelmingly felt that today's IT environments were very complex — and that their growing complexity was making it increasingly difficult for them to do their jobs successfully. The research pointed to IT teams generally feeling that they are concerned about losing control of their company's IT environment in the face of the new technologies, devices and requirements.

Network errors and downtime can not only come from applications and hardware situated within your network, but also from outside sources within your facility. Without the right network monitoring solution that can adapt to the complexity of you network, any intermittent performance problem such as an unruly appliance in the office kitchen can be as indistinguishable as a buggy Outlook plug-in.

We just don't have time to chase down problems that don't exist, or miss the ones that do. So what can you do about it? Get proactive with a unified network monitoring tool that lets you have full control in an increasingly complex business network.

For a free trial of WhatsUp Gold,please visit:
www.ipswitch.com/forms/free-trials/whatsup-gold

About Progress

Progress (NASDAQ: PRGS) offers the leading platform for developing and deploying strategic business applications. We enable customers and partners to deliver modern, high-impact digital experiences with a fraction of the effort, time and cost. Progress offers powerful tools for easily building adaptive user experiences across any type of device or touchpoint, the flexibility of a cloud-native app dev platform to deliver modern apps, leading data connectivity technology, web content management, business rules, secure file transfer, network monitoring, plus award-winning machine learning that enables cognitive capabilities to be a part of any application. Over 1,700 independent software vendors, 100,000 enterprise customers, and two million developers rely on Progress to power their applications.

Learn about Progress at www.progress.com or +1-800-477-6473.