endace

WHITEPAPER

Psst! Your NOC is Flying Blind

How long did it take to get a permanent fix for your last network problem? Too long? Even worse, did you fall short of actually getting one? If so, it's not surprising. But because of network visibility problems, chances are good that your organization is leaking real money in network and business downtime.

You can't manage and protect what you can't see. While many hundreds of network monitoring tools exist, organizations today know surprisingly little about the traffic on their mission-critical high-speed networks. It takes them so long to fix network problems that the business suffers while the seconds, minutes, and hours of downtime accumulate.

This situation has never been ideal, but it is now becoming costlier than ever as we grow more dependent on networks to transact business. Network engineers are expected to quickly resolve network incidents and achieve as close to 100% network availability as possible. After all, your organization stands to lose thousands or even millions for every minute your network is down. And reliance on the corporate network is only increasing with the rollout of new applications such as latency-sensitive unified communications (UC) and virtual desktop infrastructure (VDI), which sends each user keystroke across the network to the data center for manipulation and processing.

Causes of poor visibility

With all the monitoring products on the market, why are your network troubleshooters flying blind? There are two main reasons:

1. While most enterprises are fully invested in incident detection and protection tools (1 and 2 in the figure on the right), they haven't yet embraced tools that facilitate incident resolution and root cause analysis. These are the tools that empower you to quickly find the real problem, resolve it, and eradicate it for good
2. Most monitors see just a fraction of network traffic. As a result, engineers receive incomplete information and are prone to following a misleading trail as they investigate incidents. Being sent off course slows down both your mean and maximum repair times in an economy when, more than ever, time is money.

These situations exist because traditional packet monitors are unable to accurately record and timestamp every single packet. So they leave network engineers to draw troubleshooting conclusions based on less than a complete picture of what occurred. This approach increases repair times and decreases network and business uptime because the information being acted upon is of the 'garbage in, garbage out' variety.

The consequences

Having only a partial picture makes it particularly difficult to troubleshoot intermittent problems that don't match up with the averages. Intermittent problems are usually the most time-consuming to resolve and cost enterprises a small fortune each year in downtime.

Specifically, the absence of incident response and root cause analysis tools in enterprises causes the following problems:

- Delayed responses
- Wasted resources as engineers chase traces
- Hit-or-miss corrective action
- Service degradation and resulting user dissatisfaction
- Escalation of network management costs
- Lack of an ability to contain problems and to minimize the reach of their impact.

Despite this state of affairs, today's network engineers need

Endace incident management model

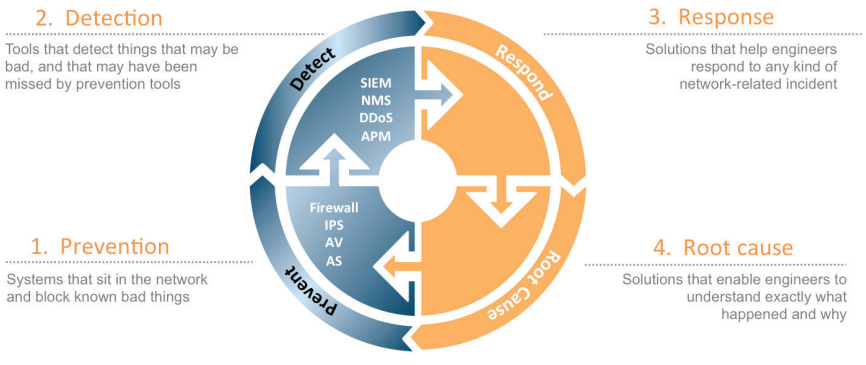

to respond to incidents and troubleshoot their root causes in minutes, not weeks, to keep the business running and profitable.

Why the urgency?

Network visibility and incident resolution struggles have been around for a long time, so why the scrutiny now? As mentioned, the relationship of the network to business revenue is only growing stronger as we rely on the network for more and more operations, transactions and applications. There are several other factors, too:

- **Network transmission rates now outpace monitoring equipment.** Many organizations currently are migrating to 10Gbps network cores – with 40Mbps and 100Mbps cores not far behind – and today's packet monitoring systems simply can't keep up. As a result, most drop packets, making it impossible to get an accurate picture of traffic behavior. Incident response times slow, and networks can't deliver their maximum ROI
- **Database systems don't scale.** Similarly, most monitoring systems store collected data in a database that can't scale in step with network growth. Again, this leads to the delivery of misinformation and lengthy troubleshooting times, network downtime and related financial losses
- **New applications are converging on the network.** UC apps enable collaboration among employees, partners, customers and others are placing more traffic and reliance on the network than ever. One key troubleshooting aspect of UC is the need to resolve intermittent voice-over-IP (VoIP) problems. This requires viewing the packet flow's exact behavior at the time of the issue, because real-time sessions like VoIP otherwise open up, close, and are gone forever. Employees and customers don't tolerate poor-quality voice calls, and business can suffer immediately from unresolved issues. Also, as noted, VDI is another app invading enterprises, making every user keystroke, in effect, a 'network incident' that requires tracking
- **Swivel-chair management is no longer adequate.** Most existing packet monitoring alternatives have been purpose-built for particular applications. The approach has required enterprises to spend a lot of money to get all the tools and analysis engines they need and also for the associated rack space and power required for each system. These stovepipe systems require swivel-chair viewing of separate displays, rather than integrated views of the network. These systems no longer suffice for today's very large, mission-critical networks. They leave room for too much guesswork and manual incident correlation, which wastes valuable time

What approach is needed?

Large enterprises that rely on their networks for business require a new approach to network monitoring – one that extends their visibility tools arsenal to include capabilities that accelerate incident response efforts and the ability to determine and fix root causes. To do that, the tools need to work pervasively, across your entire global high-speed network, and not miss a trick with overlooked or dropped packets.

The ability to visualize, search, and retrieve packets of interest from anywhere across the network empowers you to do the following:

- Drive down mean and max time to resolution on all kinds of network performance incidents
- Establish the true root cause of network incidents, which drives future incident prevention
- Develop meaningful network performance and security containment plans and strategies
- Monitor key network performance metrics in real time so issues can be addressed before they become service-affecting.

One implementation option is to deploy a searchable network visibility overlay, or 'fabric,' across the entire network. Just as you wouldn't want holes in your routing or switching fabric, it's detrimental to have holes in the fabric that provides the lens into those devices and the packets they send and receive. With the visibility fabric approach, you gain insight into the full story that the network packets have to tell. In this way, network engineers can plug the financial leakages caused by network interruptions and downtime.

For more information on Endace products visit: **endace.com**
For enquiries email: **enquiries@endace.com**