

While simple in concept, network baselining is often misunderstood.

Baselining involves recording network traffic and performance, saving it for future reference and/or reviewing it to see traffic patterns. Once baselines are saved, they can be used as a benchmark with which to compare other traffic patterns.

The technique provides the network administrator insight into expected behavior on the network and subsequently, the ability to notice changes. People often think of expected behavior as always being good traffic, meaning that expected behavior of a network reflects when everything is running perfectly. This is incorrect. Think of expected behavior as known vs. unknown traffic.

By understanding the behavior of a network and what has happened historically, one can begin to solve problems that arise. Baselining makes it easier to identify network attacks (internal or external) and even the people causing problems (staffers downloading movies at work, for example). Without a basis of comparison, how do you determine the difference between good and bad traffic?

But the reality is less than 5% of administrators make a practice of baselining, for reasons such as “we don’t have the time to do baselines” or “things change too much to do baselines” or “I’m not going to hire a person or multiple people to keep baselines organized.” In these tough economic times, such concerns need to be exposed for what they are: misconceptions.

Baselining is not a time sink — it’s actually quite the opposite. And consider the economics of foregoing baselining. Change on a network can cost from thousands to millions of dollars. For example, adding bandwidth to a network with multiple sites and WAN links to a thousand or more

stores might increase costs by \$500 to \$1,000 per site. It is imperative that organizations size their networks based on legitimate traffic before adding such significant recurring costs.

So, how should you approach baselining? While theoretically there could be thousands of baselines, the key to success is deciding what baselines are important to the organization.

There are many macro-level baselines to consider, such as how much bandwidth is going out to the Internet and how much bandwidth is in the core. And then there are many more granular views: How many people are talking to a particular network? What protocols are going across the network? How much bandwidth does a particular application use in general? What is latency on the network for a particular application? The list goes on and on.

Consider this basic list of baselines that everyone should start out with:

- All traffic on backbone links
- All traffic on WAN links
- All traffic to the Internet
- All traffic for particular business critical applications
- All traffic to/from critical systems
- All systems backup traffic
- All the above for each location if multiple locations exist

The security side of baselines also is important. Changes in the environment will reflect security posture because certain databases, applications, data and devices should only be accessed from particular locations, IP addresses, networks or people. By spending a few moments to create baselines for these entities, it takes only an instant to see if someone tries to access those protected resources. Similarly, change management and compliance issues should have baselines to back up the assumptions that are made about what is happening in that network environment.

Beyond determining what baselines are important, consider what should be baselined. Some people use baselining to understand and monitor network traffic in off hours and odd times in efforts to identify attack patterns. Others focus on how applications tax the network. Before implementing baselines, it’s crucial to determine the purpose and end goal.

If the end goal is security-related, consider the following list of baselines:

- All conversations that originate in a DMZ
- All traffic to/from network/security devices
- All VPN traffic after decryption
- DHCP traffic from unknown DHCP servers
- Mail traffic not sent to your mail servers
- DNS traffic not to your DNS servers
- Any internal addresses not matching your address space

No matter what baselines you choose to implement, the overall strategy remains the same: it is important to understand the baselines and the benchmarks of the particular application servers, clients on the network, and the overall health of the network from a utilization perspective. It takes just a little time and money to automate the baselining process, and results are invaluable to network management.

Every personality and aspect of your network that you know and understand substantially increases your chances to fix your next problem in seconds, rather than hours. And in these tough times, quicker resolution equates to less downtime, which equates to less dollars lost. That is where the true ROI of baselining can be seen.

*McCreery is president and CEO at WildPackets. He can be reached at [tim.mccreery@wildpackets.com](mailto:tim.mccreery@wildpackets.com).*