SANS ANALYST PROGRAM

Sponsored by BreakingPoint

Network Security: Theory Versus Practice

A SANS Survey – May 2011

Written by James Tarala

Advisor: David Rice

Introduction

Network security professionals rightly worry about the threats facing the network and data center infrastructures that they must defend. Cyber attacks have never been easier to launch, and successful attacks are more costly than ever. Yet the answers to a recent SANS Institute survey of IT security professionals reflect a chasm between the perception of these risks and the measures being taken to harden network security—even within highly security-conscious sectors of government and private enterprise.

Results of the SANS survey make it clear that network security personnel are not consistent about validating the resiliency—performance, security and stability—of the devices and systems that go into their network and data center infrastructures. In far too many cases, their organizations either lack processes to validate resiliency or attempt to ensure resiliency inadequately, without subjecting devices to the conditions of attack, heavy user load, and complex application traffic that they will face once in the real world. While network security practitioners understand the risks to their operations, they continue to neglect a valuable measure of the health of their networks—resiliency. For the purposes of this report, network and data center *resiliency* is defined as *the performance, security and stability of devices and systems under attack, high-stress load, or both*.

Consider these survey findings that illustrate how organizations believe resiliency is an important component to security and stability of their networks:

- Of respondents, 53 percent have suffered denial-of-service (DoS) attacks.
- Approximately 68 percent say that device resiliency is an important consideration when making network and data center equipment purchases.
- A mere eight percent say they trust the claims of network and security equipment vendors.

Yet IT policies and practices do not reflect these realities, as indicated by findings like these:

- An incredible 82 percent of respondents have no formal resiliency validation program in place.
- Nearly 62 percent of respondents do not yet validate the resiliency when new infrastructure devices are added to the network.

A paradox emerges. On one hand, there is common agreement about the importance of network and security device resiliency; the inability to trust vendors' claims; and the prevalence, risks, and costs of threats such as malware and DoS attacks. On the other hand, however, many major government and corporate IT organizations do not consistently perform formal resiliency reviews, nor do they perform routine validation of devices under the high-stress conditions, such as attack or heavy traffic load, found in production environments.

This is a real problem. Lax practices create unnecessary risks and force organizations to overpay for both IT products and their maintenance. Networks and data centers must deliver high performance and secure services to foster competitive advantage. They are also under the strain of increasingly aggressive and persistent threats. Cyber attacks, especially when combined with heavy network traffic, threaten many network and security devices with performance degradation or outright failure. Such events can lead to harmful business impacts, including lost revenue, leakage of sensitive data, legal penalties, fines for regulatory noncompliance and damage to the brand.

This research report unveils the gap between theory and practice when it comes to building resiliency into the network and data center infrastructures of government agencies, service providers and major enterprises. Ultimately, the results of this survey make it clear that network security practitioners must go beyond *understanding* the relevant risks and *take action* to develop standard processes and frameworks for maintaining the resiliency of their infrastructures.

About the Survey

Nearly 200 IT professionals responded to the survey, which was conducted by the SANS Institute. Most respondents (90 percent) were hands-on network and systems operations personnel whose roles put them in positions to understand the relationships among security, performance and stability on their networks. These roles included system engineers, security engineers and similar staff positions. There was some crossover between these positions and management roles, which were also strongly represented in this survey base (see Figure 1).

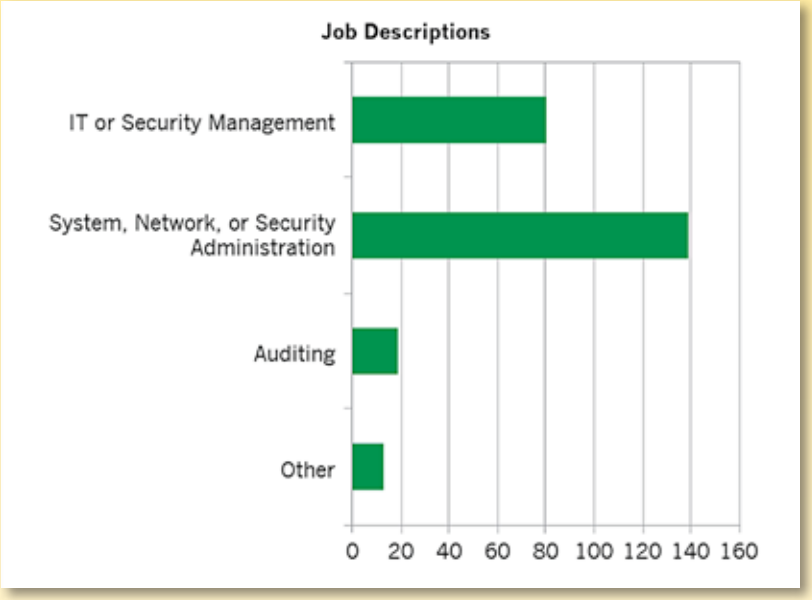

Figure 1: Breakdown of Administrators and Management

Large organizations were strongly represented in this survey, with 50 percent, of respondents representing enterprises of 2,000 end users or more, as illustrated in Figure 2.

Figure 2: Size of Respondents' Organizations

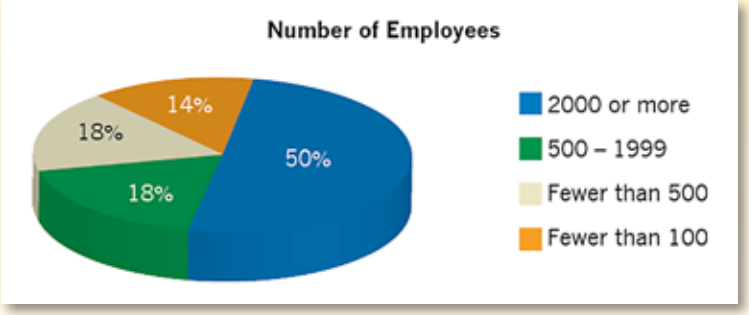

The financial services industry and

government agencies were most strongly represented, with another large segment of respondents coming from the telecommunications sector. The industry verticals of respondents are summarized in Figure 2.

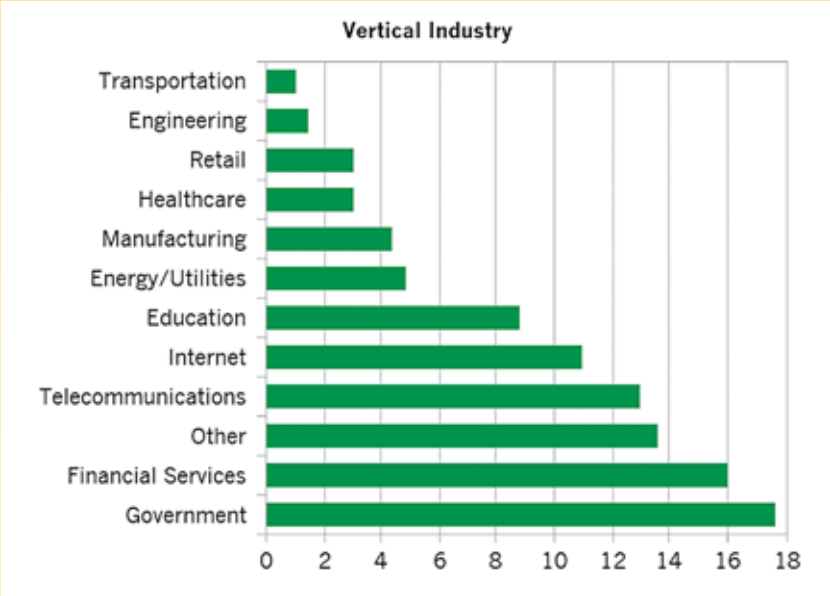

Figure 3: Vertical Industries Represented

Large enterprises, particularly within the government and financial sectors, generally have more rigorous demands for network uptime and security. Given that these groups represented the survey's largest response base, it would indicate that respondents to this survey are better informed about resiliency than the smaller organizations in less sensitive vertical markets. Even with this advanced survey base, however, responses show that enterprises have a long way to go to incorporate proper resiliency testing methodologies across their networks. This will be discussed further in the following sections.

Threats to Networks and Data Centers

Government agencies and financial services companies—the two sectors most strongly represented among survey respondents—are often particularly sensitive to network security risks because of the nature of their work and the sensitive information handled by their network and data center infrastructures. Almost all large enterprises, regardless of sector, are also prime targets for cyber threats. All of these organizations understand that network and data center performance degradation, not to mention system outage, can lead to costly impacts.

While insider threats were also a major concern, when respondents were asked to choose the top three risks to their organizations, system outage due to device failure also emerged as a top area of risk. Figure 4 provides a view of the top risks to network and data center security.

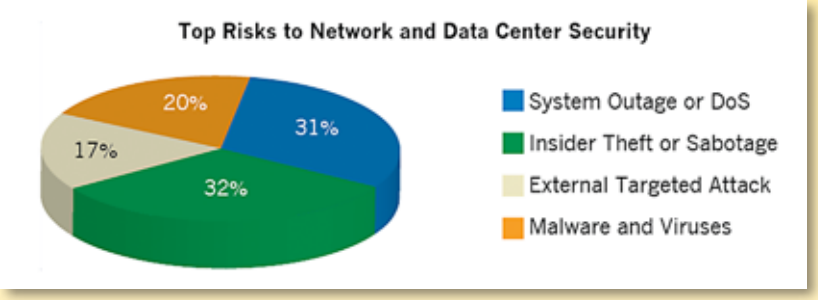

Figure 4: Top Security Concerns

When respondents were asked to set aside insider threats and consider only the impact of external attacks, device or system outage emerged as an even greater concern. When looking strictly at external threats, 53 percent of respondents considered system outage due to device failure, system stress or denial of service (DoS) attack as their biggest worry, as illustrated in Figure 5.

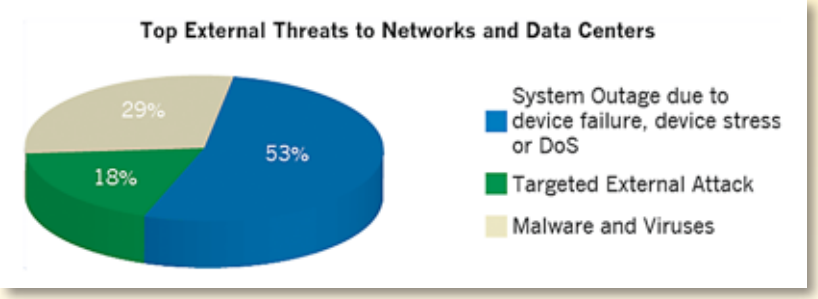

Figure 5: Top External Threats

These concerns compare directly with the reality these respondents are experiencing. A follow-up question asking about the threats that respondents have actually faced revealed that their worries are well founded. More than half of all respondents had directly experienced a DoS attack against their enterprise network or data center. In fact, concern about the impact of external events on resiliency correlates to the frequency of DoS attacks experienced by enterprises. (It is worth noting that this concern over DoS attacks preceded the December 2010 attacks against major enterprises such as MasterCard, Visa, Amazon and PayPal.)

A DoS can arise from both benign and malicious spikes in traffic. But no matter the origin of a particular DoS, every organization must understand precisely the resiliency of its network and data center infrastructures when faced with unexpectedly high load. Regardless of the reason, the inability of a business network to respond to users carries a negative impact. This lesson was brought home sharply by the DoS attacks launched by "hacktivist" supporters of Wikileaks in December 2010. Even without ultimately threatening the transaction-processing infrastructures of the commercial entities they targeted, the hacktivists threatened the normal commercial operations of companies as large as PayPal/eBay and MasterCard

Where Is the Commitment to Harden Network and Data Center Resiliency?

Respondents perceive the top external risk to their organizations to be system outage due to device failure. Because the costs of such failures are clear, what are companies doing about this risk?

Despite the fact that most organizations represented in this survey are strongly aware of many of the risks inherent in their network and data center infrastructures, 82 percent of respondents said that they have no formal review process to determine the resiliency of devices under load and attack. An amazing 26 percent of respondents lack even an informal review process, as illustrated in Figure 6. This result highlights the previously mentioned chasm between theory and practice when it comes to determining the resiliency of devices and systems that make up the IT infrastructures of today's enterprises.

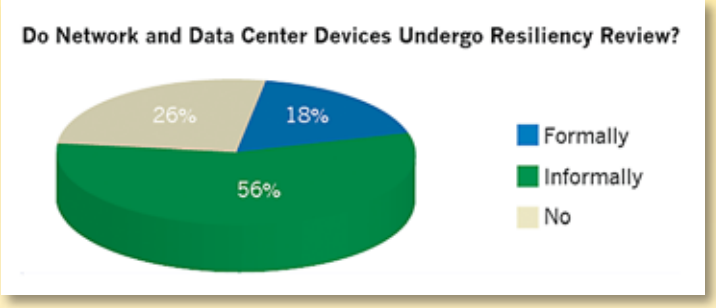

Figure 6. Formal Review Processes Lacking

Digging deeper into organizational behaviors only verifies this trend. For example, although 68 percent of respondents said that resiliency in the face of attack and high-stress application load is a consideration prior to purchasing critical network systems, only 38 percent said that their organizations are reviewing each new critical device under attack and load prior to installation. In other words, most organizations are not putting their theories into practice by proactively determining how their critical network devices will operate (or fail) under heavy stress. Figure 7 illustrates how common performance testing is.

Figure 7: Performance Testing Important

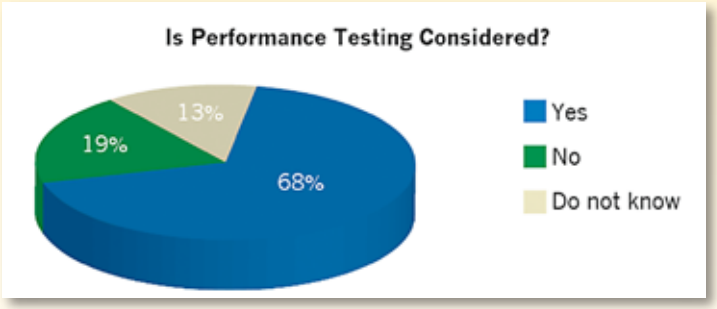

This gap between beliefs and practices is even more striking, considering respondents' overall lack of trust in vendor and analyst claims about network devices' ability to withstand external attacks, DoS, insider action and other threats. In the survey, 84 percent of respondents agreed that they needed to verify vendor and analyst claims, yet only 18 percent said that they were already performing this verification with any formal, repeatable process. Figure 8 illustrates the mistrust of performance claims.

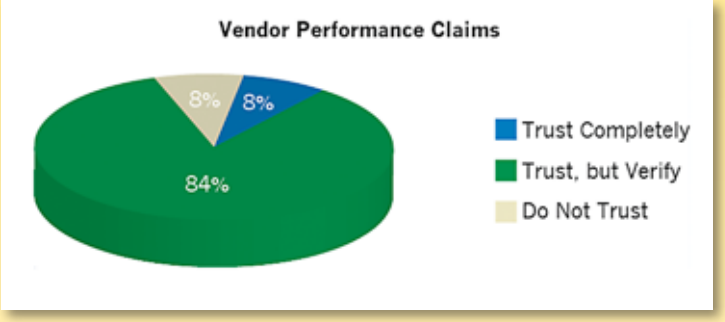

Figure 8: Widespread Mistrust of Performance Claims

So, if IT buyers do not trust vendors, where do they look for guidance? Clearly they do not look to analysts. Answers to a similar question about industry analysts indicate that only 6 percent of respondents trust analyst claims. Unable to trust either vendors or analysts, organizations are left to verify the performance, security and stability of devices for themselves.

Because only a small percentage of respondents indicated that they have formal review processes, the question arises: How are organizations examining systems and validating resiliency? In addition to the 18 percent that do have formal processes, other respondents, at least, are analyzing some devices some of the time, including:

- **When new security or network devices are added:** Only 23 percent of respondents perform network-wide assessment when new devices are added.
- **During production changes and upgrades:** Only 15 percent said that they conduct resiliency reviews consistently prior to most production changes or upgrades.
- **When performing capacity planning and risk management:** Approximately 50 percent said that networkwide assessment of performance and security are part of their capacity planning and risk management strategies, though not for every new network system or change.
- **When changes are made to the network:** Almost 37 percent perform resiliency reviews for network changes only on an ad-hoc basis.
- **During major changes:** Only 21 percent conduct resiliency reviews as part of their change cycle.

Current Network and Data Center Resiliency Practices

With risk, cost and mistrust so high, what are organizations doing to ensure the resiliency of the devices in their network and data center infrastructures? Informal measurement is being performed to some degree, but not often. It is also unclear what organizations are measuring and what standards they are measuring against.

Current Practices

Of those respondents that review at least some of their devices for resiliency, there are variances in their lab, load and traffic practices:

Lab Environments. Despite their lack of formal review processes, 41 percent of respondents say their organizations have test lab environments with the ability to perform the resiliency exercises described in this paper. The remaining 59 percent of organizations lack the lab resources to validate vendor claims for themselves.

Simultaneous Attack and Load. When asked if they run security attacks and heavy traffic loads at the same time to assess devices, 24 percent say they do so formally, but 25 percent do no such reviews at all.

Types of Traffic. Those organizations that validate their devices under high-stress load generate different traffic types: 12 percent simply use HTTP traffic (which is probably inadequate in nearly every case), 26 percent use blended application traffic, and 31 percent use custom traffic sets, as illustrated in Figure 9.

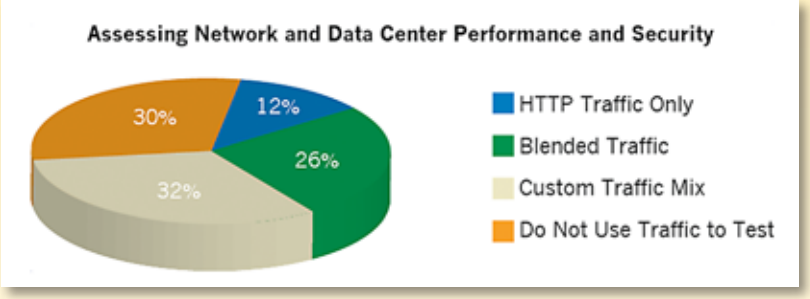

Figure 9: What Traffic They are Assessing

Although 58 percent use blended or custom traffic to simulate network load, it is unclear how enterprises are generating this custom or blended traffic. However, these results do indicate that organizations attempting to validate devices are at least in a position to create their own network traffic types, loads and behaviors for stressing network devices.

The only way to carry out a proper resiliency review is to use real traffic protocols and profiles. Only such a mix of genuine application load can emulate the conditions that network and security must withstand once deployed. In particular, simulations using only HTTP traffic are bound to be of little value to an organization. True resiliency validation must represent an accurate mix of application traffic, including protocols commonly used on corporate networks such as AIM, MSN Messenger, Gmail, FIX/FIXT, SMB/CIFS, database applications and even social networking applications.

The same focus on realism must also be present for emulating user load. Traditional software-based load testing can emulate only a thousand users or so. This is never enough in a world where even the most simplistic command-and-control botnets can easily generate a DoS attack from hundreds of thousands of computers. Companies must achieve realism with both the mix of applications and the level of user load that they emulate. In doing so they are preparing for worst-case scenarios.

Unfortunately, there are no common standards for the volume and types of simultaneous network events that devices should be able to withstand. Because networks and the particulars of their traffic patterns diverge widely, organizations need to be proactive about measuring and hardening systems using real traffic flows in the types applicable for their own environments. Better understanding these patterns would be a good topic to investigate in a follow-up survey.

Recommendation: Validate Resiliency under Real-World Conditions

Because vendor claims can't be completely trusted, and because each network requires validation against different types of traffic and levels of load, smart network security practices must ultimately include a formal resiliency review process for critical network systems. This process should be built around clear standards for the performance, stability and security of network and data center devices and systems under conditions that reflect a company's real network traffic and bandwidth. The process should be applied to all new devices and upgrades and should be implemented in periodic sweeps.

The review process should accurately represent daily operations and corresponding traffic types, as well as how that operational traffic will continue to flow during an adverse event or attack. Enterprise organizations will continue to look for more direct resiliency measurement as they deploy next-generation equipment. At the same time, these organizations must establish "resiliency lifecycle" practices that are both cost-effective and accurate.

Whatever specific steps they choose to take, enterprises must develop better approaches to online threats to their business—threats that are becoming more numerous, more complex and more dangerous by the day. Organizations should work systematically to protect what is most important to them, including the privacy of sensitive data and the integrity of their corporate brands, by hardening the resiliency of their network and data center infrastructures. This implies taking a proactive approach to the landscape of threats instead of merely reacting to new threats and technologies as they evolve.

Conclusion: Closing the Network and Data Center Resiliency Gap

Respondents to the SANS Institute survey on network security came primarily from organizations that place a high premium on network and data center security; they regard failures of devices, systems and services as among their top security concerns. Most of them have experienced DoS attacks, so these organizations have felt firsthand the cost that such outages cause.

Despite this painful awareness, organizational actions clearly contradict experiences and beliefs, because these same respondents lack well-defined standards, processes and resources for determining the resiliency of their critical network devices and systems. The source of this gap is unclear. It could be the result of pressures to pursue other priorities. It could arise from limited resources, or the from disconnect between network security and network operations teams within an organization. Or it could simply reflect a lack of execution.

Whatever the causes, this much is clear: Enterprises, particularly those with a vested interest in high security, should be setting and maintaining formal standards for repeatable network and data center resiliency reviews. This process should include:

- Resiliency reviews for every piece of network and data center equipment before purchase, before deployment, and after configuration changes
- Methodical resiliency validation using a combination of real traffic, heavy load, and security attacks
- A formal approach by which vendors are held accountable for the claims they make about their products

Whether those processes remain within individual organizations or eventually become industry standards, they will encourage—or force—organizations to be accountable for how they optimize and harden IT infrastructure resiliency. They will also force accountability on manufacturers of network and security equipment, spurring these vendors to raise their own standards for delivering resilient devices and systems.

About the Author

James Tarala is a principal consultant with Enclave Security and is based out of Venice, Florida. He is a regular speaker and senior instructor with the SANS Institute as well as a courseware author and editor for many SANS auditing and security courses. As a consultant, he has spent a large amount of time assisting organizations in their security management, operational practices, and regulatory compliance issues, and he often performs independent security audits and assists internal audit groups in developing their internal audit programs. Tarala completed his undergraduate studies at Philadelphia Biblical University and his graduate work at the University of Maryland. He holds numerous professional certifications.

SANS would like to thank its sponsor:

BreakingPoint™

Find it before they do.™