Optimizing enterprise security via application and network monitoring

By deploying capable application and network analysis tools throughout your infrastructure, you weave in the ability to handle attacks regardless of which defenses fail or where the attacks originate. Visual UpTime® Select™ provides the detailed visibility and analysis to detect attacks, limit the damage and restore operations which limit the risk and exposure of security breaches.

Table of contents

FLUKE networks®

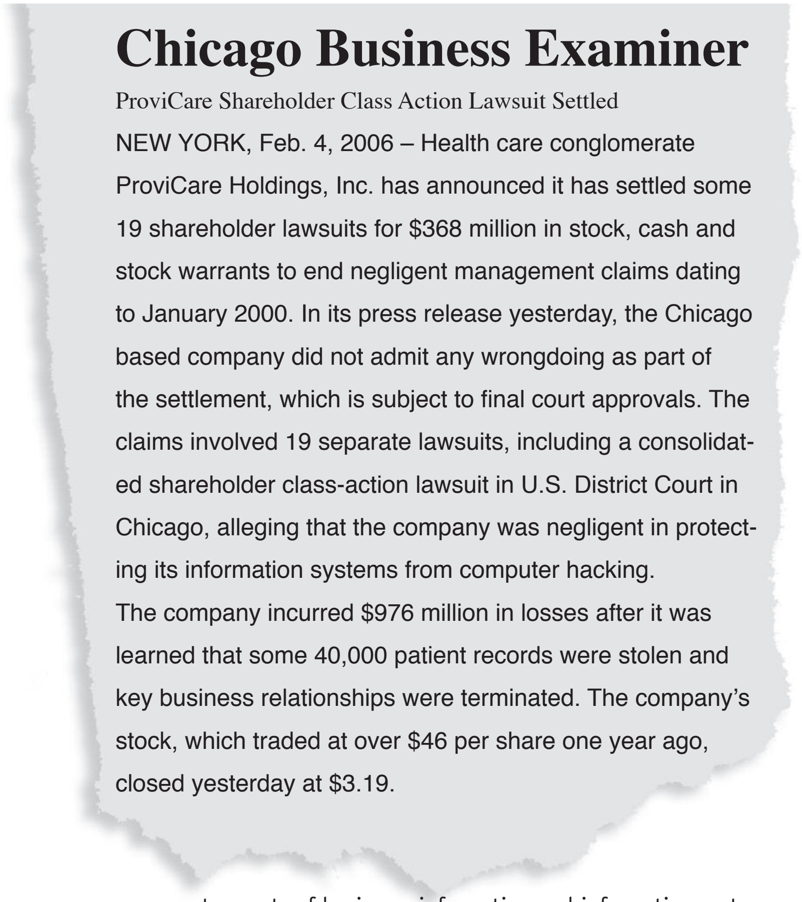

Chicago Business Examiner

ProviCare Shareholder Class Action Lawsuit Settled

NEW YORK, Feb. 4, 2006 – Health care conglomerate ProviCare Holdings, Inc. has announced it has settled some 19 shareholder lawsuits for $368 million in stock, cash and stock warrants to end negligent management claims dating to January 2000. In its press release yesterday, the Chicago based company did not admit any wrongdoing as part of the settlement, which is subject to final court approvals. The claims involved 19 separate lawsuits, including a consolidated shareholder class-action lawsuit in U.S. District Court in Chicago, alleging that the company was negligent in protecting its information systems from computer hacking.

The company incurred $976 million in losses after it was learned that some 40,000 patient records were stolen and key business relationships were terminated. The company's stock, which traded at over $46 per share one year ago, closed yesterday at $3.19.

Sound hard to believe? The story is completely plausible and it's only a matter of time before a similar story becomes front-page news. Standards set by recent legislation such as the Sarbanes-Oxley Act of 2002 are backed with the threat of fines and criminal prosecution. These same standards boost civil lawsuits whenever substandard practices expose corporate earnings to undue risks – or actual security breaches lead to real business losses. These types of regulations will affect virtually every business in the near future.

A regulation for every occasion

Even before the terrorist attacks and corporate accounting scandals, legislators and other governing bodies were more closely scrutinizing and monitoring information-age businesses. Since then, the pace has quickened substantially, perhaps with good reason.

Information technology has fundamentally changed most aspects of business; information and information system processes are a key asset and a fundamental competitive differentiator. Loss or compromise of that information – or the ability to use it – is a major risk to shareholder value, and is often of a very private and sensitive nature to customers and must be safeguarded to protect client relationships.

IT-impacting regulations affecting virtually every business are now in place and more are on the horizon. One of the most publicized regulations is the Sarbanes-Oxley Act of 2002, which applies to all companies with publicly traded stock and focuses mostly on financial governance of the corporation. Since modern businesses rely heavily on IT for accounting support, Sarbanes-Oxley also contains the mandatory "Computer and Security Controls" provisions.

Even without fines, lawsuit settlements and prison terms, the costs of inadequate computer security are significant. According to a recent article in *Optimize*, the indirect costs of computer security incidents (e.g. lost sales, damaged customer relations, legal damages) far outweigh the direct expenses of security equipment and personnel. The study shows that cyber crimes involving a breach of confidentiality (e.g. disclosure of medical or financial account information) caused the victim company's market valuations to decrease by more than 5% on average.[1] The market perception is that any breach of trust will convert directly into loss of future revenues as customers take their business elsewhere.

Another study of 162 companies by research firm Aberdeen Group found the average firm was losing $2 million in revenue each year because of Internet attacks. The businesses typically faced one significant disruption each year with its systems requiring an average of 22 hours of downtime to recover.[2] Even for companies that have avoided disruptions on this scale, it appears they may just have been lucky so far. According to Symantec's Internet Security Threat Report, "more than 40% of Fortune 100 companies controlled IP addresses from which worm-related attacks propagated."[3]

The message is clear – all businesses will be successfully attacked from time to time, whether they know it or not. So how can today's businesses protect themselves?

Heading off trouble

The key to protecting your company is the actions your organization takes before, during and after an attack. Should the unthinkable happen and you must defend yourself, you will need to demonstrate a detailed understanding of the risks and a clear pattern of actions to address them. Your success depends on your organization's approach to everyday planning. Keep in mind three keys as you formulate your plans:

Not unlike total cost of ownership and scalability, security is a dimension you must consider in every IT-related decision.

Security is a dimension. Not unlike total cost of ownership and scalability, security is a dimension you must consider in every IT-related decision. Continuously weaving security throughout the fabric of your IT infrastructure builds a strong foundation and limits the need for expensive custom add-ons.

Attacks can come from anywhere. Most of the hype about security solutions today still focuses on the network perimeter. Most attacks walk right through the front door – they never encounter the network perimeter at all.
For example:

- An employee brings an infected floppy disk from his home network to update the latest football pool
- A contract employee goes online with her infected laptop
- A vendor's infected PC inadvertently attaches to a rogue wireless access point on your network while she attends a meeting in your conference room
- A disgruntled former employee leaves behind a time bomb

Count on your defenses being breached. Build the best perimeter defenses you can afford. Install the most sophisticated firewalls, authentication mechanisms, intrusion detection systems, e-mail filters and VPNs for secure, remote communications. Run the latest anti-virus software as a last line of defense. Configure them all as conservatively as your organization can bear and your patience will allow. Assume all your carefully constructed walls will be outflanked by something as simple as a new way to disguise a trojan-bearing worm attack in an e-mail, or by a well-intentioned employee running something he thought was harmless. What then?

Prevent attacks from occurring

Your plans for responding to security incidents are just as important as your plans to prevent them from happening. From a practical standpoint, all modern security planning flows from the classical "protect, detect, react and restore" paradigm. Incident response focuses on the final three areas of the paradigm and yields three distinct action phases:

- Detecting attacks
- Limiting the damage
- Restoring operations

Detecting attacks

There is no single technique, mechanism, or tool for detecting every type of attack. Attacks come in all shapes and sizes. Several new ones have probably been invented since you got your morning cup of coffee. Some attacks are automated – such as viruses, worms and distributed denial of service attacks (DDoS) – while other attacks are manual, such as when a hacker sets up shop on one of your high-performance servers. The automatic attacks are often indiscriminate and attack when their algorithmic profile is met. Most manual attacks are targeted based on who you are (e.g. a huge, monopolistic corporation) or what you have (i.e. high bandwidth connectivity and processors and lots of storage). Almost all attacks have one thing in common: they do things your users don't ordinarily do.

The recognizable signs of an attack are often referred to in the computer security trade as "indicators." These signs include overt evidence of an attack such as screen and web site graffiti as well as more covert evidence like modified files (file indicators) and server log file entries (system indicators). All of these indicators are "after the fact." They are signs that damage has already been done. The final type of indicator is the network indicator. Network indicators are signs of attacks that you see on the network itself and can provide you with the early warning that someone or something is trying to do harm while there is still time to react.

Making the most of network indicators

The key to identifying an attack before bad things happen and users are affected is to recognize abnormal (suspicious) behavior in networks. The first generation of Intrusion Detection Systems (IDS) is trying, but is still in its infancy and it will be years before it can spot abnormalities as well as the human brain can. IDS now needs to be backed up with human intellect and tools that help it scale. How do the humans do it?

> Application traffic and network performance are the easiest areas in which to spot abnormal behavior associated with an attack.

Application traffic and network performance are the easiest areas in which to spot abnormal behavior associated with an attack. For example, if your organization uses Microsoft Exchange Server as its mail platform, alarm bells should sound when a significant amount of SMTP traffic suddenly hits your internal network. This might be an indication that one of the many mass-mailing worms is starting an infection. To know what is abnormal, you must first know what is normal. The general process is as follows:

1. Construct a baseline. The baseline is your record of normal network conditions. It is a combination of application flow, usage and status information from across the network to create a complete picture that you can use later when things have changed. The things you want in your baseline are the things most likely to be changed by an attack of some sort. For each site on your network, and for every hour of the week, you need to know the following:

- Which addresses act as "authorized" servers?
- What are the "authorized" protocols?
- What are the top protocols?
- What are the top sources and destinations of traffic?
    - How much traffic volume do they generate?
    - What protocols do they use?
- How much traffic is exchanged with other sites?
    - What protocols do they use?

Because you decide to create a baseline today doesn't make today's conditions "normal." Malefactors could already be at work in your organization, contributing to application traffic loads and performance degradation. Don't accept your baseline until you believe it is an accurate record of normal conditions.

2. Analyze traffic using the baseline. Now that you have established a baseline, you need to routinely analyze network activity against it.

Start with the application protocols. Are there new protocols that you don't recognize? Do any of the baseline protocols have significantly greater traffic than before? This could be an indication that the protocol is being used to tunnel peer-to-peer traffic or worse. Are there applications active at odd hours of the day for your business, such as telnet at 3 a.m? This is made easier by analysis tools that analyze and baseline protocol usage over time.

Now look at servers. First, are there new servers that weren't there before? Were you expecting them? Are the servers seeing the same relative amount of inbound and outbound traffic? Have any client stations now become servers of one sort or another? Any of these could be worm activity, new peer-to-peer file sharing or hacking. Analysis tools can help by automatically identifying the servers in your network.

Know the flow

An application flow is the sequence of packets related to a single application session in one direction between a pair of endpoints. This means that for most applications, there will be at least two application flows for every session, one from client to server, and the other from server back to client. Each application flow is usually characterized by several distinguishing factors that include source and destination IP address, source and destination port, TCP handshake status flags, and the application protocol in use. These items make it possible to reconstruct the key attributes of an application session for both security and performance analysis.

To create application flow data, agents are deployed that passively inspect the user traffic as it goes by on the network. Agents can be part of network analysis devices or built into devices such as routers and CSUs. More important than the application flow data itself is the analysis applied to it once it is collected. Given the overwhelming amount of data produced by applications on even the smallest network, it is essential to have an intelligent tool that reconstructs flows and presents analysis in ways that help you do your job most effectively. Features such as top N lists of protocols, hosts, and flows as well as automatic server identification and alerting on new protocols dramatically enhance your ability to catch abnormal behavior.

Next, look at the application flows. Are certain end-user's computers or stations unknowingly initiating many small flows to a large number of addresses? Are lots of connection attempts failing and being repeated? Perhaps a station is initiating flows to invalid addresses. These stations are likely conducting network scans or probes, looking for vulnerabilities. Are connections being made at unusual times? Are certain clients downloading more traffic than makes sense to you? A tool that provides application flow volumes and organizes the flow information by protocol, source and destination makes the analysis quite easy.

A tool that provides application flow volumes and organizes the flow information by protocol, source, and destination makes the analysis quite easy.

Finally, look at traffic volumes over your key network links. Are volumes of traffic higher during what are normally off-peak times? Are certain sites generating more traffic than you would expect? Any of these can be signs of undesirable activity that could be related to an attack.

3. Investigate differences. Beware of rationalizing the differences between the baseline and current conditions. Any one of them could be your early warning of a larger problem. Take steps to investigate them. It can be as simple as waiting a short time to see if traffic levels return to normal. Call a division contact to learn if they deployed a new server or application. Without alerting the IT department, some IT shops shoot first and ask questions later – they block new applications and servers that show up unannounced knowing that if a legitimate purpose was being served the out-of-process owners will get in touch soon enough.

4. Update the baseline. Your applications, systems and networks are changing every day. This complexity created the holes that let the attackers in and they are counting on it to cover their tracks. A periodic update of your baseline is necessary for tuning your detection process; it is also a good time to identify unintended, unauthorized changes to your infrastructure. Tools that continually monitor network activity and store a historical record of it automate much of this process for you, but you still need to give it a critical eye and validate it yourself.

Besides using baseline analysis to detect network indicators of an attack, you can also employ various types of trip-wires. A trip-wire is a virtual or actual network resource that sends you an alarm whenever particular conditions are met. If you know what activity is normal in your environment, you can devise simple trip-wires to detect the presence of some types of intruders by the presence of abnormal activity. Of course, trip-wire conditions are carefully chosen so that honest people and applications have no reason to trigger them. Consider them the network equivalent to the holdup alarm that is triggered when the bottom dollar bill is removed from a bank teller's cash drawer.

Instrumenting trip-wires is easy with the automatic alarm generation capabilities of some network equipment and probes. Once configured, the agents will alert you whenever the trip-wire conditions are met on the network segments being monitored. Choose trip-wire guides based on your knowledge of your application environment, baseline and the ways that attacks operate. For instance, a common tactic used by many worms is to send mail using an internal SMTP mail-forwarding relay to reach the outside. If you don't already use the common "mail.yourdomain.com" address these worms often target, you can add a bogus entry for it to your internal DNS and trap on any SMTP (tcp/25) traffic to that destination. Alternately, you can set trip-wires on any probing activity to some of the more popular back doors and other vulnerable ports and IP addresses that attackers tend to use.

Limiting damage

Once you know you are under attack, limit the damage by denying further access to the resources the attack is trying to use, preventing the attack from reaching fresh targets. The challenge is to understand the attack you are dealing with well enough to employ the right type and level of response for the particular kind of attack and its stage of development. For example, a live hacker at work on your server cluster requires an entirely different response than a worm spreading by instant messenger. Understanding the attack for this purpose requires that you answer three basic questions:

1. How did it arrive on your network?
2. What does it target?
3. How does it spread and how far has it gone?

The search for answers to these questions should start with a historical analysis of network traffic. How was the attack first detected and where? Is there evidence that the attack was underway earlier and perhaps at other places, but was somehow overlooked? Chances are the point of origin is near the location where evidence of the attack was first detected. If you can identify the origin, by all means, take it offline if possible. Failing that, it may be practical to quarantine the network branch containing the origin to achieve some relief if the attack has not spread very far. This will also limit your exposure to other attacks or re-infection if the point of origin has weaknesses that are being exploited.

To quickly answer those questions and especially to determine how the attack is being carried out and how far it has progressed, you will need to rely on network analysis tools. If the malefactor involved a known agent, your up-to-date virus scanning software would have almost certainly detected it by this point. If not, you are likely dealing with a zero-day mutant, or hacker attack. Application flow analysis makes quick work of summarizing the traffic along links and between endpoints to expose abnormal patterns. Link utilization and top talker analysis let you track down the most active victims. Once you have the infection in your sights, capturing and analyzing infected packets may be the key to your salvation.

> If your visibility is limited to only a portion of the network, your response can't really begin limiting the damage until the infection reaches that point.

Not only do you need to acquire tools that give you the necessary information, you also need to deploy them in the right places ahead of time. Remember the incident response planning paradigm "Attacks can come from anywhere?" If your visibility is limited to only a portion of the network, your response can't really begin limiting the damage until the infection reaches that point. Hopefully that isn't your headquarters site. A distributed analysis system that provides visibility to all of your remote sites is your best tool for both early detection and damage control. A cost effective way to approach distributed analysis functionality is by embedding it into the network infrastructure. By choosing network gear with analysis capabilities that meet your security and day-to-day operations needs, you get stronger, more integrated capabilities at a lower price.

Restoring operations

With the damage contained and the attack's spread halted, users will be breathing down your neck to restore operations as soon as possible. While doing the job quickly will win you supporters, nothing will convert them into detractors like a re-infection that occurs when you miss something. Be meticulous now.

Take advantage of intelligence from a variety of online resources to learn about your attack and how to restore your server. Visit DoSHelp (www.doshelp.com/trojanports.htm) as well as the vendor of your company's antivirus software, such as Symantec, McAfee or Trend Micro. The vendor sites all have online databases to help you identify, troubleshoot and recover from the attack. They are also excellent clearinghouses for information regarding the latest infections.

Besides repairing compromised and damaged systems, close the exposed holes checking the steps in the "Limiting the Damage" phase. It's also a good time to update your plan and monitoring procedures based on what you have learned. When you are ready to start bringing things back online, it is worthwhile to do it in phases using your network surveillance tools to watch for signs the infection or some part of it may still be present.

How Visual UpTime Select can help

Many of the capabilities you need for maintaining strong surveillance over your network to ward off attacks are the very same capabilities you depend on for operating mission-critical parts of your network on a daily basis. Visual UpTime Select fills both roles with its detailed seven-layer analysis, granular historical baselining and flexible alarm-generation capabilities.

A network-wide summary views that provide a complete accounting of all the applications everywhere on the network, including the unknowns, is critical.

Deployed as a part of your infrastructure, Visual UpTime Select uses a systems approach to managing networks based on rich and deep instrumentation of the LAN, WAN and router components that support your applications. Instrumentation can take the form of specialized Analysis Service Elements (ASEs/DSUs) or hardware-based agents, which can be upgraded with additional intelligence as needs change, as well software agent technology embedded in common network components, such as Cisco routers and other access devices.

Regardless of how they are deployed, Visual UpTime Select agents provide deep-packet inspection for all data flowing past them to collect detailed application performance and usage information. The agents upload this data to a Visual UpTime Select server, which stores, analyzes and presents a complete view of network usage.

The server's database enables network-wide summary views that provide a complete accounting of all the applications everywhere on the network, including the unknowns. This data stored historically forms the basis for your network usage baseline by detailing all the key factors:

- Servers and clients
- Protocols in use
- Top protocols
- Top sources and destinations including traffic volumes and protocols
- Site-to-site traffic volumes and protocols

Armed with a detailed understanding of your network's usage patterns, you can instruct Visual UpTime Select to warn you of abnormal traffic and usage patterns wherever they occur. To automate various aspects of your notification and incident response, you can export these warnings to third-party event management and trouble-ticketing systems such as Remedy.

Finally, when you have to take the gloves off to look inside packets for identifying zero-day and other unknown attacks, Visual UpTime Select provides traffic capture, protocol decode and packet export functions at every managed point. With Visual UpTime Select, you will be armed in the face of a network attack.

Managing an attack with Visual UpTime Select

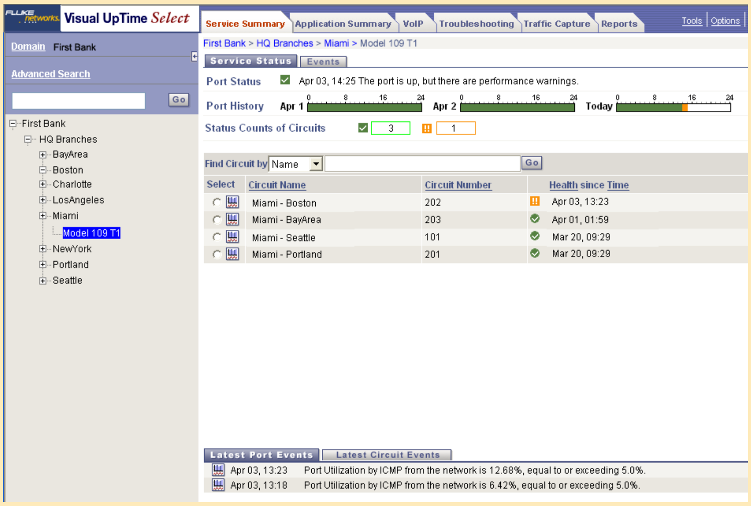

Domain First Bank

First Bank > HQ Branches > Miami > Model 109 T1

- First Bank
  - HQ Branches
    - BayArea
    - Boston
    - Charlotte
    - LosAngeles
    - Miami
      - Model 109 T1
    - NewYork
    - Portland
    - Seattle

Port Status ☑ Apr 03, 14:25 The port is up, but there are performance warnings.

Port History Apr 1 · Apr 2 · Today

Status Counts of Circuits ☑ 3 ‼ 1

Find Circuit by Name

Select	Circuit Name	Circuit Number		Health since Time
○	Miami - Boston	202	‼	Apr 03, 13:23
○	Miami - BayArea	203	✓	Apr 01, 01:59
○	Miami - Seattle	101	✓	Mar 20, 09:29
○	Miami - Portland	201	✓	Mar 20, 09:29

Apr 03, 13:23	Port Utilization by ICMP from the network is 12.68%, equal to or exceeding 5.0%.
Apr 03, 13:18	Port Utilization by ICMP from the network is 6.42%, equal to or exceeding 5.0%.

Figure 1: *Use customizable thresholds to set trip-wires across the network.*

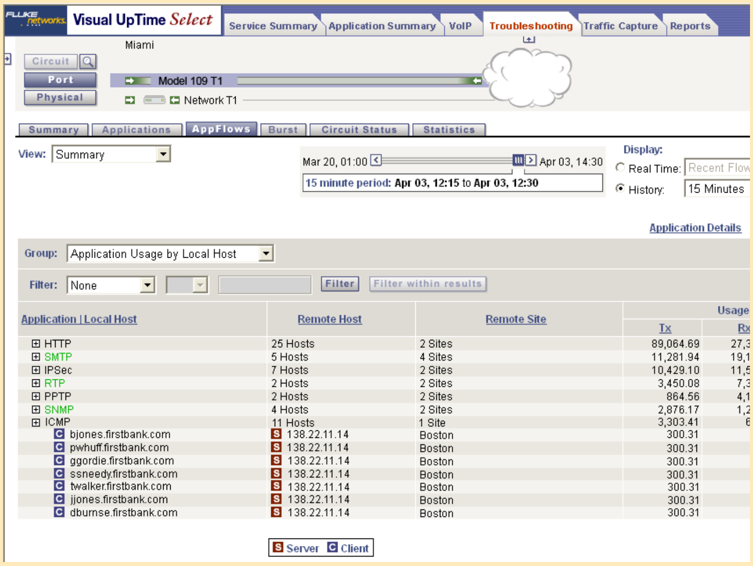

Miami

Model 109 T1

Network T1

View: Summary

Mar 20, 01:00 · Apr 03, 14:30

15 minute period: Apr 03, 12:15 to Apr 03, 12:30

Display:
○ Real Time: Recent Flow
● History: 15 Minutes

Group: Application Usage by Local Host

Filter: None

Application | Local Host	Remote Host	Remote Site	Usage	
			Tx	Rx
⊞ HTTP	25 Hosts	2 Sites	89,064.69	27,3
⊞ SMTP	5 Hosts	4 Sites	11,281.94	19,1
⊞ IPSec	7 Hosts	2 Sites	10,429.10	11,5
⊞ RTP	2 Hosts	2 Sites	3,450.08	7,3
⊞ PPTP	2 Hosts	2 Sites	864.56	4,1
⊞ SNMP	4 Hosts	2 Sites	2,876.17	1,2
⊞ ICMP	11 Hosts	1 Site	3,303.41	6
C bjones.firstbank.com	S 138.22.11.14	Boston	300.31	
C pwhuff.firstbank.com	S 138.22.11.14	Boston	300.31	
C ggordie.firstbank.com	S 138.22.11.14	Boston	300.31	
C ssneedy.firstbank.com	S 138.22.11.14	Boston	300.31	
C twalker.firstbank.com	S 138.22.11.14	Boston	300.31	
C jjones.firstbank.com	S 138.22.11.14	Boston	300.31	
C dburnse.firstbank.com	S 138.22.11.14	Boston	300.31	

Figure 2: *Application flows highlight individual client to server relationships*

A management platform should be able to detect attacks, limit the damage, and restore operations. Visual UpTime Select provides the critical visibility needed to minimize the risk and exposure of security breaches.

Let's look at a scenario with a DDoS attack that has infected an enterprise. An attack of this nature is not identified until multiple locations bandwidth utilization exceeds port speed; now other applications and users are being impacted. With Visual UpTime Select, you can set customizable thresholds so you can detect potential attacks quicker and more efficiently.
An alarm has been generated highlighting ICMP traffic has increased beyond a baseline target **(See Figure 1)**.

A spike in ICMP traffic does not automatically mean you have an attack on your network; it can be a symptom of an attack or maybe just a spike in authorized usage of the application. Armed with this information, drill down into more detail and see application flows. Sorting by ICMP, Visual UpTime Select quickly identifies individual flows. You can see that many authorized users are hitting the same server with identical amounts of bandwidth **(See Figure 2)**. This is a trait of a virus attack.

By quickly identifying a potential attack and drilling down to application flows, enterprises can quickly limit the damage and rapidly restore normal operations.

Conclusion

Security in the IT infrastructure is no longer something enterprises can merely hope to achieve. The costs of poor security have never been higher. There is a growing trend toward government regulation to mandate minimum standards. Realizing a secure IT infrastructure requires planning and preparation long before an attack occurs.

The "protect, detect, react, and restore" paradigm is still an excellent way of creating your overall security plan. Think of security not as an add-on, but a quality you design into every aspect of the infrastructure. Don't focus security plans strictly on defending the perimeter, because a significant portion of today's threats bypass the perimeter. Strive for multiple layers of security, then expect your defensive measures to fail by designing and implementing an effective incident response to recover from any attack.

Strive for multiple layers of security and even then, plan for your defensive measures to fail by designing and implementing an effective incident response to recover from any attack.

Many of the keys to implementing your incident response plan are found in tools that provide surveillance and analysis on the networks throughout your organization. By maintaining awareness of how your networks are normally used through periodic base lining you learn what to look for. Performing routine application flow, source/destination and network utilization analysis, organizations can detect early attacks using network indicators. The setting of alarms on baseline deviations and anomalous conditions automates many of the more repetitive surveillance tasks. When an attack does eventually break through, you are armed with detailed information for a quick and effective response and restoration. By deploying capable network analysis tools throughout your infrastructure, you weave in the ability to handle attacks regardless of which defenses fail or where the attacks originate. Visual UpTime Select provides the detailed visibility and analysis to detect attacks, limit the damage and restore operations which limit the risk and exposure of security breaches.

About Fluke Networks

Fluke Networks is a leading provider of network and application performance management solutions. The company's technologies enable enterprises to reliably and securely manage the delivery of mission-critical applications across their infrastructure. Fluke Networks' products increase application and network availability, optimize the use of bandwidth, and reduce operating costs across traditional and IP-based infrastructures. For more information, visit **www.flukenetworks.com**.

[1] "The New Economics of Information Security", Lawrence Gordon and Robert Richardson, *Optimize Magazine*, April 2004, Issue 30

[2] *Corporate Losses From Internet-Based Attacks Average $2 Million*, TechWeb.com, July 6, 2004

[3] Symantec Internet Security Threat Report, Trends for January 1, 2004-June 30, 2004

NETWORKSUPERVISION

Fluke Networks
P.O. Box 777, Everett, WA USA 98206-0777

Fluke Networks operates in more than 50 countries worldwide. To find your local office contact details, go to **www.flukenetworks.com/contact**.

Printed in U.S.A. 11/2006 2791000 D-ENG-N Rev A