

Justifying the Expense of IDS, Part Two: Calculating ROI for IDS

Kevin Timm 2002-08-27

This article is the second of a two-part series exploring ways to justify the financial investment in IDS protection. In [part one](#) of this series we discussed general IDS types and expanded on the impact that the logical location of a company's critical networked assets could have on the risk equations. To this end we introduced the Cascading Threat Multiplier (CTM) to expand on the Single Loss Expectancy (SLE) equation. We also reviewed implementation and management costs based on various support profiles and reviewed the commonly accepted risk equations. Finally, we left off with the basic formula for calculating ROI for security, otherwise commonly known as Return on Security Investment (ROSI).

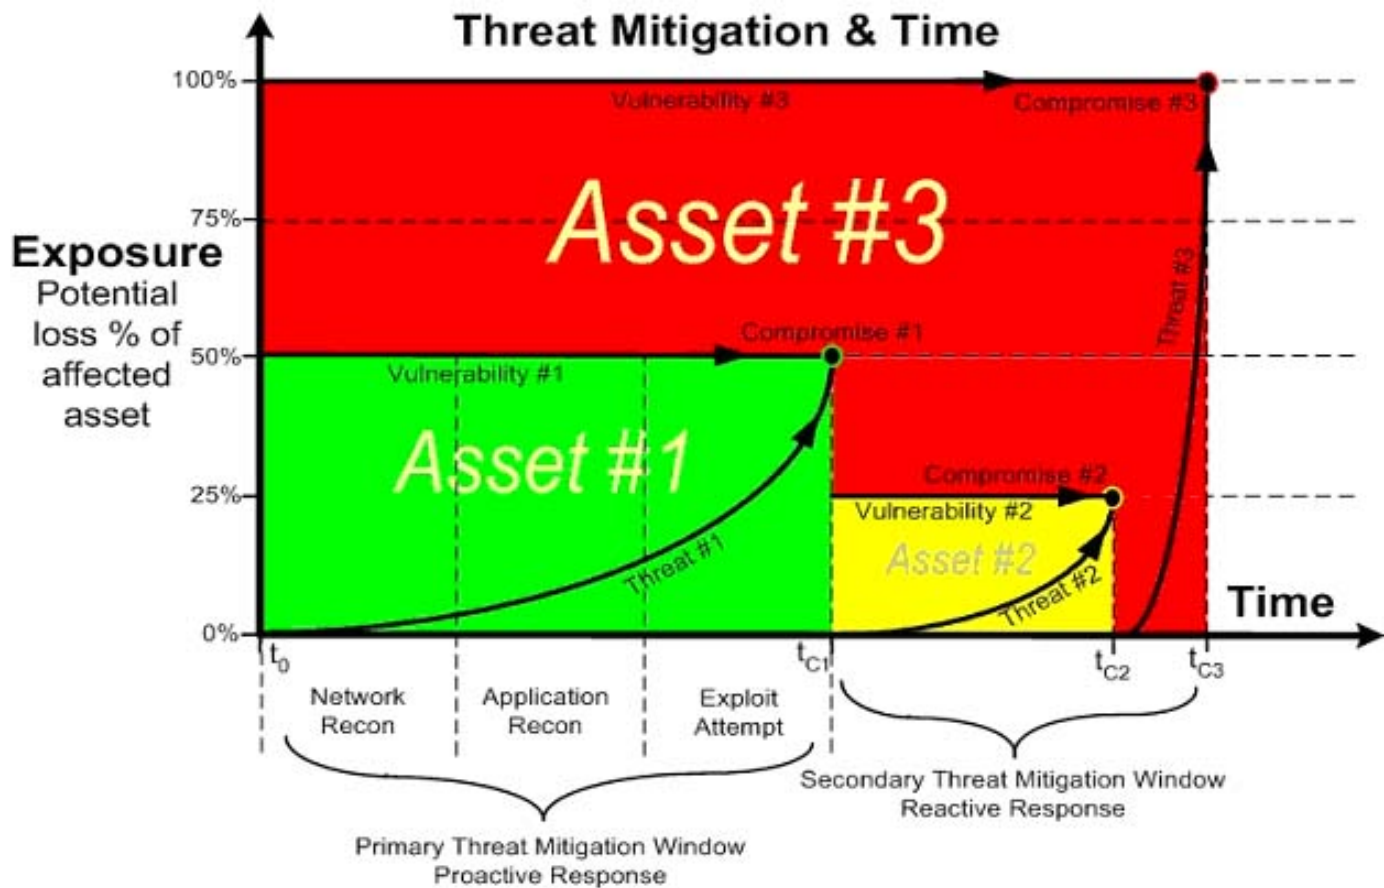
In this article we will discuss proactive and reactive management methodology and how this methodology affects our analysis of risk. This will set us up nicely for the number crunch for ROI on IDS. We'll culminate this exercise with a brief overview of our hypothetical company named Wally's Building Supplies (WBS) and, finally, we'll put all the numbers together to demonstrate our technique for calculating ROI for the WBS IDS deployment of one network-based IDS and two host-based IDSs.

Proactive vs. Reactive Management

Independent of implementation and management costs, the method in which the devices are managed can have a serious affect on ROI. To this point, the key question to answer is: is the system going to be proactive or reactive as security events are detected? The following table depicts the normal event flow in each method. A proactive implementation response is automated by the system while a reactive implementation response is done once personnel have been enlisted.

Method	System actions	Personnel actions	Follow up information
Reactive	Log -> Alert ->	Respond -> Analyze -> Eradicate	Forensics and Evidence
Proactive	Respond -> Log -> Alert	Analyze -> Eradicate if necessary	Forensics and Evidence

By examining the Annual Loss Expectancy ($ALE = ARO * SLE$, where $SLE = Exposure\ Factor * Asset\ Value * Cascading\ Threat\ Multiplier$) we can determine which variables are affected by each of these two management methods. In a reactive design, where personnel must be engaged to respond to each event, the exposure factors (primary [EF] and secondary [EFS]) will be affected. In a proactive design there will be similar benefits to the exposure factors (re: a reduction) and, in addition, the Annual Rate of Occurrence (ARO) will be influenced in a beneficial way as well. To demonstrate the impact of threat vs. time we will use the concept of primary and secondary mitigation windows. In the following graph the primary mitigation window affects ARO while the secondary mitigation window affects Exposure Factor and Cascading Threat Multiplier. An effective way of impacting ARO is through automated response.



Auto-response can take many forms. On host-based IDS this is sometimes called shielding, where a specific process is terminated. Network-based IDS generally employs TCP resets or shunning. TCP resets effectively kill one specific session based on suspicious activity, but it still allows other activity from that same IP. Shunning, on the other hand, changes firewall rules or router access lists and effectively denies all traffic from that host for a specific period of time. In essence, shielding will protect a single host from one process, resets will protect a host from a specific session, and shunning will protect the entire network from a specific host for a pre-determined amount of time.

The accuracy of automated response can vary tremendously. This is dependent on the skill level of the engineers managing the devices. If the engineers are moderately skilled then auto-response will not be very effective, which may adversely affect the ROI of the IDS deployment. This adverse effect may manifest itself in the form of a loss of productivity from network-related problems due to improperly implemented auto-response, as well as the additional fallout related to a false sense of security throughout the company.

With skilled engineers managing the devices, auto-response can be very accurate and effective. Because few statistics exist that illustrate the accuracy of automated response we will use statistics generated from our analysis of one month's worth of data on networks that NetSolve, Incorporated manages (the authors, Kevin Timm and David Kinn, both work for NetSolve, Incorporated located in Austin, TX). If we include Code Red and Nimda activity, in 99.96% of the attacks, where automated response was used to mitigate the threat, the activity was malicious. Excluding large-scale worms, the attacks were malicious in 95.8% of auto-response uses. Of the 4.2 % of the traffic that was not malicious, not all of it was desirable. Some of this traffic was peer-to-peer programs, on-line gaming, chat and other undesirable traffic that triggered alarms. The percentage of traffic that

was denied that was business related was very small. It should be noted that many of these devices provide numerous different techniques for ensuring that very little, if any, legitimate traffic is denied through the use of automated response.

To determine how effective the device is in recognizing attacks we will use the most recent [NSS study](#). In this test the worst NIDS detected 67 of 109 attacks or 61.5%, while the best detected 94 of 109 attacks for an 86.2 % detection rate. Even the worst case, the 61.5% detection rate was out of the box and [NSS reported](#) that it would not be difficult to improve this with some custom signatures and tuning?.

What does all this mean? It means that the worst IDS tested can still detect at least 61.5% of attacks.

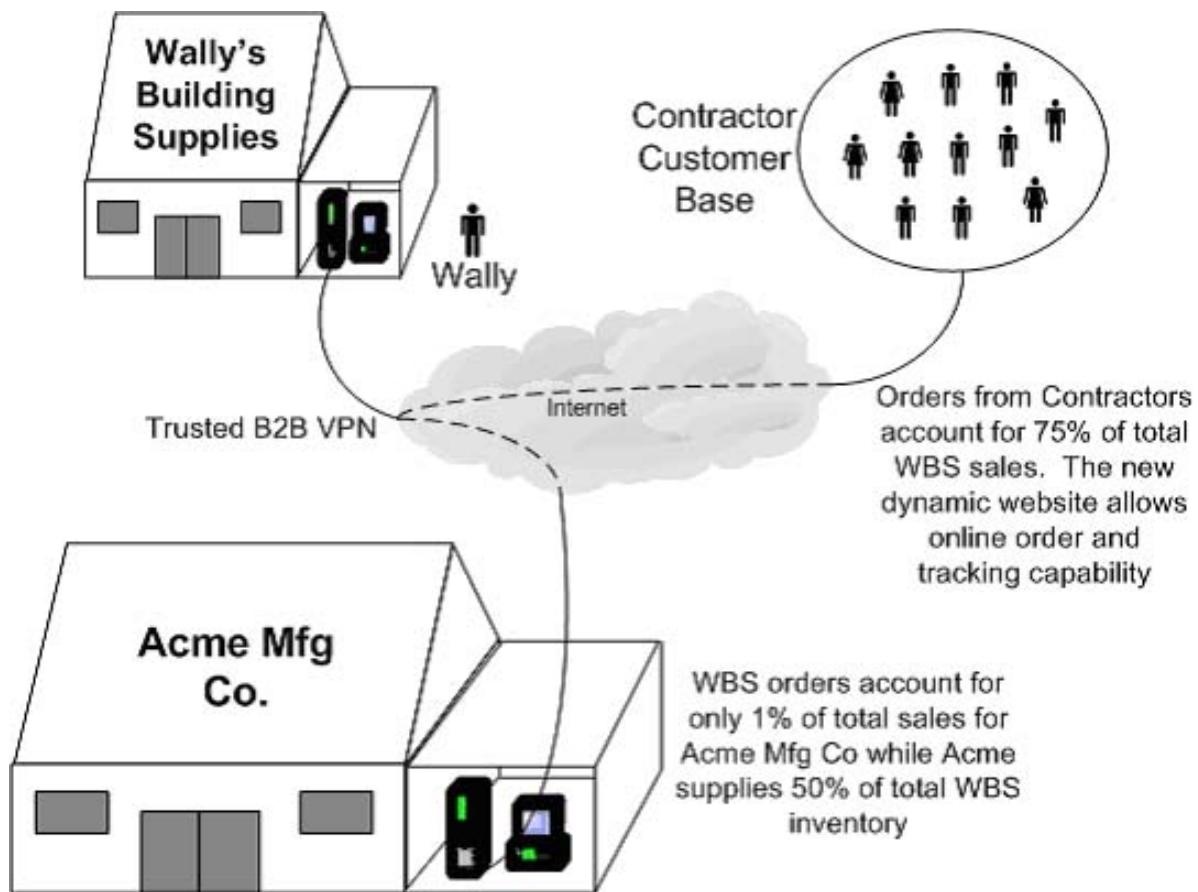
Realistically that number should be closer to 70% when a skilled engineer or technician manages the device. The auto-response feature, when properly used, can be a very effective method of reducing the Annual Rate of Occurrence (ARO). This provides us with some general numbers we can plug into our equations for calculating a ROI for Wally's Building Supplies.

Sample Company: Wally's Building Supplies

Wally's Building Supplies (WBS) has six supply outlets, with the business office located within the primary outlet. WBS has several business-to-business (B2B) VPN connections to its suppliers. Their small staging department procures most of WBS items for all six outlets over these B2B connections by running an over-the-counter order procurement software application agreed on by each of the suppliers. Of the dozen or so suppliers, ACME is WBS most important one, accounting for 50% of all WBS procurement needs. ACME and WBS have built their trust relationship over the course of many years doing business together. ACME has experienced phenomenal growth over the past decade and supplies scores of building suppliers around the country. WBS orders account for a mere 1% of total ACME sales.

For several years WBS has maintained a simple informational Web page showing store locations and directions, general goods and services available and monthly specials. The primary target market for WBS consists of residential and commercial building contractors. Contractors comprise 75% of total WBS sales, with the remaining 25% generated from do-it-yourself consumers.

Recently WBS had contracted out the development of a dynamic database-driven Web site that allows contractors to order supplies on-line, check the status of their orders, and confirm deliveries to the construction site. The dynamic Web site has already had a positive effect on the operational ROI of WBS by improving efficiencies related to its' antiquated order fulfillment and delivery confirmation process. Inventory turnover has increased as a result of these efficiency gains, which in turn has improved WBS bottom line. That's the good news. WBS maintains Internet connectivity through a T1. Most of WBS servers, routers and infrastructure have been set up by outside IT contractors. See the diagram below to visualize the WBS world a little bit better. So what's the bad news?



The Compromise

WBS primary supplier, ACME, recently informed WBS that a malicious attacker gained access to ACME's data and network through the VPN tunnel with WBS. It is unknown to ACME if this was an outside attacker or an ill-willed employee from WBS. Because of this, ACME has disconnected the B2B VPN with WBS and temporarily discontinued service with WBS until the issue is resolved. They have agreed to fulfill all outstanding orders in the interim. Since WBS has very little technical expertise, they called in ABC Security Consulting Services (ABC) for a thorough analysis of the alleged compromise. Before we talk about the ABC analysis results, let's talk a bit about why this particular incident is so bad for WBS and also point out some of the negative fallout that will most likely occur as a result of this compromise.

For one thing, it's bad that ACME was the one that informed WBS that the attack seemed to have originated from their trusted B2B VPN connection. ACME is probably thinking: "Does WBS have a clue about a basic security policy"? Rest assured that ACME would be very diligent in getting to the bottom of what happened. Next, the trust relationship between WBS and ACME has always been strong; but after this compromise, who knows what the future holds. Remember that WBS relies heavily on ACME for doing business while ACME could drop WBS and hardly notice the hit on its revenue stream. Let's say it took years for WBS to establish the favorable credit terms it currently enjoys with ACME. How much is this business relationship worth to WBS? What's the value of highly favorable credit terms with an established and reliable supplier like ACME?

In addition to this, there are the contractors, who constitute 75% of WBS's revenue stream to consider. This compromise, and the ensuing fallout with ACME, could delay many contractor orders, which could result in missing

contract deadlines and substantial penalties. Many of WBS larger contractors may be required to procure supplies from some of WBS competitors to meet their building schedules and avoid these costly overruns. Lastly, there is potential litigation to consider. If it turns out that one of WBS employees was the perpetrator of the ACME compromise, the whole ship could sink. We could go on and on about all the negative effects of this type of trusted peer compromise but we'll stop here. It should be obvious to observers that Mr. Wally is in quite a challenging situation! Life was so simple before the Internet.

The Analysis

After extensive network and operational analysis by ABC, several key operational deficiencies were uncovered that revealed the inadequacies of the WBS network. Furthermore, ABC found several key security vulnerabilities and design flaws. These vulnerabilities included:

- The original static Web site was on a NT 4.0 server that was several service packs behind. This server ran multiple vulnerable programs. This server was also configured as a part of the NT domain so it had network access to many of the internal devices as if it was on the internal network. There was also no network access control between this server and the internal network.
- The newer dynamic Web site was on a modern Unix server that was relatively current. This was also the only Unix host on the network. This host was connected to an older internal database that contains inventory, availability, and pricing information. This was the contractor Web site and is a large part of the WBS business plan going forward. Even though this host was relatively current, it still had vulnerabilities within Apache and SSH and had FTP and RPC daemons running.
- WBS has a router that had the WEB management interface open on the internal interface (which was still accessible from the internet). This gives away complete router configuration, including VPN information and internal network architecture.

ABC recommended that there is a definite need for WBS to implement IDS technology to monitor the content of each connection. The recommendation is that a host-based IDS be run on the Web servers and a network-based IDS run at the border. See Table 1 below for a summary of ABC's analysis and Table 2 for general statistics on how often these types of attacks occur (note: numbers in Table 2 are based on networks that NetSolve manages).

Table 1

Scenario	Descriptions of Compromised Asset (AV) and Underlying Exposed Assets (UEA)	Considerations in assigning estimates for Asset Valuations (AV & UEA), Exposure Factors (EF & EFS) and Annual Rate of Occurrence (ARO)	AV	EF	UEA	EFS	ARO
----------	--	--	----	----	-----	-----	-----

One	WBS NT 4.0 Web server (AV); WBS NT Domain (UEA)	Cost of lost productivity and revenue from downtime? Cost of compromised underlying data and assets? Cost of rebuilding web server? Potential cost of compromise of NT domain resources?	\$2,000	75%	\$20,000	75%	3
Two	WBS UNIX-based Web server (AV); old internal WBS database containing inventory data and pricing for customers and suppliers (UEA)	Cost of lost productivity and revenue from downtime? Cost of loss of trust or confidence of WBS online customers? Cost of compromised data and assets? Cost of rebuilding web server? Immediate cost of fulfilling current orders with whatever it takes to satisfy customers?	\$2,000	50%	\$50,000	50%	2
Three	WBS router; Primary supplier ACME's network	Cost of supply interruption from primary supplier? Cost of loss of trust or confidence of primary supplier? Potential cost of compromised data at ACME? Cost for WBS to replace ACME as a supplier? Difference in credit terms for new supplier (compared to highly favorable terms that WBS currently enjoys with ACME)? Difference in pricing between normal (new) supplier and ACME's pricing? Potential cost of litigation if ACME determines WBS employee is at fault? Cost of WBS	\$3,000	75%	\$200,000	50%	1

		compromise? Potential cost of liability for attacks directed at other non-partner networks?					
--	--	---	--	--	--	--	--

Table 2: Average Attack Occurrence per Network

Attack	Per Network Attempts (April)	Per Year Attack Attempts	Scenario
General Cmd.Exe	9492	113,904	1
Root.exe backdoor	1869	22,428	1
Ida overflow	105	1260	1
SSH Attacks	.2	2.0	2
DNS Bind Attacks	.7	8.0	2
FTP Attacks	.3	4.0	2
Apache Chunked	.6*	7.0	2
IOS HTTP Unauth	3	36	3

*Data taken from July since this vulnerability was not common knowledge in April

Let's now recap all the variables and risk equations we have covered so far before we pull the lever and churn out the ROI for WBS IDS deployment. In Table 3 below we itemize each variable and equation that we will utilize in our calculations. Review the table noting how we have tied the traditional ROSI equation back to our ALE containing the CTM factor. Thus, $ROSI = R - ALE$, where the commonly accepted $ALE = (R - E) + T$ is now replaced with $ALE = ARO * SLE$, where $SLE = AV * EF * CTM$. Confused yet? See Table 3 below for a visual representation.

Table 3

Variables	Equations
Asset Value (AV)	$AV = \text{hardware} + \text{comm. software} + \text{proprietary software} + \text{data}$
Exposure Factor (EF)	EF is the % estimation of the exposure of the initial compromised asset
Underlying Exposed Assets (UEA)	UEA is the estimation of the \$ value of the assets behind the compromised initial asset
Secondary Exposure Factor (EFs)	EFs is the % estimation of the exposure of the UEAs
Cascading Threat Multiplier (CTM)	$CTM = 1 + ((UEA \times EFs) / AV)$
Single Loss Expectancy (SLE)	$SLE = EF \times AV \times CTM$

Annual Rate of Occurrence (ARO)	ARO is estimated number, based on available industry statistics or experience
Annual Loss Expectancy without IDS (ALE1)	$ALE1 = SLE \times ARO$
Annual Loss Expectancy with IDS using auto-response (ALE2)	ALE2 = conservative 50% reduction of ARO when IDS is managed skillfully with auto-response
Annual Loss Expectancy with IDS using auto-response & incident response (ALE3)	ALE3 = conservative 25% reduction of EF & EFS when IDS is managed skillfully with auto-response and incident response
Annual Cost (T) of IDS Technology and Mgmt	T
Annual Recovery Cost ('R) from Intrusions without IDS	$R = ALE1$
Annual Dollar Savings (E) gained by stopping intrusions with IDS	$E = ALE1 - (ALE2 \text{ or } ALE3)$
Traditional Return on Security Investment (ROSI) equation	$ROSI = R - ALE$, where $ALE = (R - E) + T$
WBS ROI of IDS with auto-response (ROI1)	$ROI1 = ALE1 - ((ALE1 - (ALE1 - ALE2)) + T)$
WBS ROI of IDS with auto-response & incident response (ROI2)	$ROI2 = ALE1 - ((ALE1 - (ALE1 - ALE3)) + T)$

As stressed throughout this article, for these numbers to make better sense, the IDS technology must be managed by highly skilled engineers or technicians that have a sound understanding of the technology, including the inherent strengths and weaknesses. Also, it is not unreasonable to assume that WBS IDS deployment of one NIDS and two HIDSs would be supported by a single in-house engineer or technician. What is not so reasonable is to assume that one person can support this highly dynamic technology on a continual 24/7/365 basis with active auto-response and real-time incident response for every security event. Multi-shift internal support (although not really an option for WBS considering this small deployment) as well as Managed Security Service Provider (MSSP) support are the preferred ways of providing definitive 24/7/365 support and real-time incident response.

Given that this deployment encompasses both NIDS and HIDS, a 50% reduction in ARO facilitated by the utilization of auto-response should be considered a conservative estimate. The 25% reduction in both exposure factors (EF & EFS) should also be considered a conservative estimate when coupling auto-response with prompt incident response. That being said, refer to Tables 4 to analyze the reductions in these variables (see highlighted cells under ARO, EF and EFS) and our calculation for ROI of WBS IDS deployment based on single in-house support and MSSP support related to the three compromises described above (in Table 1). The support costs used to calculate these ROIs were taken from the table entitled "Implementation & Management of one Network IDS and two Host IDS" from [part one](#) of this article. Those costs are \$83,217/year for single support coverage and \$44,217/year for MSSP support coverage.

Table 4

Scenario	ROI Scope	AV	EF	UEA	EFS	CTM	SLE	ARO	ALE1	ALE2	ALE3	Single Support ROI		MSSP Support ROI	
									(no IDS)	(w/ AR)	(w/ AR & IR)	\$	%	\$	%
One	no IDS	\$2,000	75%	\$20,000	75%	8.5	\$12,750	3	\$38,250	N/A	N/A	N/A	N/A	N/A	N/A
	IDS w/ Auto-Response	\$2,000	75%	\$20,000	75%	8.5	\$12,750	1.5	\$38,250	\$19,125	N/A	-\$64,092	-77%	-\$25,092	-57%
	IDS w/ Auto-Response & Incident Response	\$2,000	56%	\$20,000	56%	6.6	\$7,453	1.5	\$38,250	N/A	\$11,180	-\$56,147	-67%	-\$17,147	-39%
Two	no IDS	\$20,000	50%	\$50,000	50%	2.3	\$22,500	2	\$45,000	N/A	N/A	N/A	N/A	N/A	N/A
	IDS w/ Auto-Response	\$20,000	50%	\$50,000	50%	2.3	\$22,500	1	\$45,000	\$22,500	N/A	-\$60,717	-73%	-\$21,717	-49%
	IDS w/ Auto-Response & Incident Response	\$20,000	38%	\$50,000	38%	1.9	\$14,531	1	\$45,000	N/A	\$14,531	-\$52,748	-63%	-\$13,748	-31%
Three	no IDS	\$3,000	75%	\$200,000	50%	34.3	\$77,250	1	\$77,250	N/A	N/A	N/A	N/A	N/A	N/A
	IDS w/ Auto-Response	\$3,000	75%	\$200,000	50%	34.3	\$77,250	0.5	\$77,250	\$38,625	N/A	-\$44,592	-54%	-\$5,592	-13%
	IDS w/ Auto-Response & Incident Response	\$3,000	56%	\$200,000	38%	26.0	\$43,875	0.5	\$77,250	N/A	\$21,938	-\$27,905	-34%	\$11,096	25%
WBS ROI with IDS Auto-Response (ROI1)									\$160,500	\$80,250	N/A	-\$2,967	-4%	\$36,033	81%
WBS ROI with IDS Auto-Response & Realtime Incident Response (ROI2)									\$160,500	N/A	\$47,648	\$29,635	36%	\$68,635	155%

Summary

Auto-response affects primary mitigation windows, which has a direct impact on partially reducing the Annual Rate of Occurrence (ARO). This is illustrated in the ROI table above, where we see a beneficial conservative reduction in ARO of 50% (highlighted in yellow in the "IDS w/Auto-Response" rows for each of the three scenarios). Incident response affects the secondary mitigation window, which impacts exposure factor (EF) and secondary exposure factor (EFS), which in turn impacts the Cascading Threat Multiplier (CTM). This is also illustrated in the ROI table above, where we see a beneficial conservative reduction in EF and EFS of 25% respectively (highlighted in yellow in the "IDS w/ Auto-Response & Incident Response" rows for each of the three scenarios).

These reductions have positive effects on the ROI of IDS. Once the aggregate annualized savings (ALE1 - ALE2 or ALE1 - ALE3) occurring from IDS deployment equals the support costs associated to the deployment a positive ROI should materialize. In the case of WBS, the two ROIs (ROI1 & ROI2) for each support profile are as follows:

- Single support with IDS using auto-response (ROI1) = -4%;
- Single support with IDS using auto-response and incident response (ROI2) = 36%;
- MSSP support with IDS using auto-response (ROI1) = 81%; and

- MSSP support with IDS using auto-response and incident response (ROI2) = 155%.

These ROIs are based on the aggregate annualized savings from deploying and effectively managing the IDS technology and the resulting impact the IDS technology could reasonably have on the combined effect of the three compromise scenarios described above (see Table 1 to review these again).

In the end, to argue the ROI case for IDS, you need to have a sound understanding of your company including, at a minimum, how it does business, how its connected, where the asset value really lies and what vulnerabilities and associated threats (equating to risk and exposure) need to be analyzed and addressed through sound security policy and risk mitigation techniques. Deploying IDS technology using a comprehensive management methodology involving skilled personnel is directly correlated to the attainable goal of a positive ROI on IDS.

Relevant Links

[Justifying the Expense of IDS, Part One: An Overview of ROIs for IDS](#)

David Kinn and Kevin Timm

[Privacy Statement](#)

Copyright 2006, SecurityFocus