

Developing an Effective Incident Cost Analysis Mechanism

David A. Dittrich 2002-06-12

Developing an Effective Incident Cost Analysis Mechanism

by David A. Dittrich

last updated June 12, 2002

When it comes to calculating damages from computer security incidents, some in the media will tell you that it is impossible to come up with a value. At the same time, others will tell you that the Melissa Virus caused \$80 million in damages to US businesses. Who is right? Can these damages be calculated, and if so, how?

A project by representatives of the Big Ten Universities (plus a few others) in the late 90's undertook to systematically examine the real costs of security incidents. The results of this project were an incident cost model and examples of costs for typical security incidents at these institutions. This model has been used successfully in computer intrusion cases involving federal law enforcement, and by the HoneyNet Project for comparison of entries in the Forensic Challenge. It proves that fair and accurate damage estimates can be produced, and with very little work, provided that those doing the work are disciplined and diligent in keeping track of time, at the time of incident response. Unfortunately, this is where the system often breaks down. As we shall see, the need for diligence in collecting time data for every security incident response calls for policies and procedures to be set at the institutional level, and enforced as a regular part of incident handling, in order to have meaningful figures on institutional losses due to security incidents.

Can Computer Crime Damages be Calculated?

Pretty much anyone who reads electronic mail on a Microsoft Windows computer has heard of or had first-hand experience with the Melissa Virus (or SirCam, or Klez, or Code Red, or Nimda...) and knows that these things affect millions of users at a time. At the very least, they require time to run anti-virus software, but they sometimes require re-installation of the operating system and applications (not to mention lost user files that were deleted or damaged by the virus).

Ask any of these people how much "damage" they suffered, and you will often hear, "Well... none, I guess." After all, how much does it cost to re-install an operating system? You didn't have to buy a new copy, you just had to take the ti... that's right, take the time! Ever hear the

phrase "Time is money"? Well, in the real world, it is. Put simply, if you lost time, you lost money.

Add up all the time lost to handling incidents across a large institution, let alone around the globe, and you're talking big money. Nevertheless, an appallingly high percentage of University Regents, corporate CEOs, and top government administrators, have no clue how much money is being lost. No clue translates to no appropriate allocation of resources, which means an ever increasing problem as more and more intrusions take place. This is a classic computer security rant. (Indeed one of the major challenges of the security community is to develop an effective return on investment calculation that will allow them to justify security expenditures before catastrophic incidents occur.)

The uncertainty surrounding the cost of incidents was evident in the case of the Melissa Virus. Mark Rasch recently wrote [a column on the sentencing of David Smith](#), the author of the "Melissa" virus, to 20 months of incarceration in a federal penitentiary for his conviction for causing more than \$80 million in "loss" to affected businesses and computer users. Mr. Rasch pointed out that none of the companies who helped produce this estimated loss actually modified their SEC filings to include the loss. A [follow-up post by Bugman](#) restates the question, "If no victim companies can show actual damage costs, was there a loss?" A [follow-up to Bugman's post](#) by Anonymous suggests that victims can show a loss. Who is right?

Let's look at another example. In an article in the [Budapest Business Journal](#) published April 15, 2002, Mr. Robert Smyth writes: It is generally agreed that it is almost impossible to quantify the extent of damage of internet-related crime. "I don't think anyone has numbers on how much is lost. Most cybercrime goes unreported because of PR considerations," said [Dániel] Nemes. "It could be a very high figure, taking into account how online banks and brokerages can be seriously disrupted by denial of service."

I strongly disagree: not knowing something is not the same as it not existing. Perhaps there is general agreement that it is almost impossible to quantify the extent of damage; but I would argue that this is simply because people don't know how to do it – or they cannot be bothered doing it - not that it is really impossible to do. It would be ludicrous to suggest there is general agreement that companies can't tell how much labor costs go into software development projects or construction projects, and yet the two tasks - calculating productivity costs and calculating production losses - are essentially the same. Sure, many sites may not want to say how much they lost (for fear of negative publicity, to keep share-holders in the dark about

losses, etc.), but that doesn't imply that it is impossible to quantify the damage from computer crime!

The fact is, it is rather simple to estimate damage due to security incidents if you know a few simple facts about the personnel who are responding to, or are affected by, the incident. Such information can be ascertained by answering the following questions:

- Who worked on responding to or investigating the incident?
- How many hours did each of them spend?
- How many people were prevented from working because of the incident?
- How much productive time did each of them lose?
- How much do you pay each of those people to work for you?
- How much overhead do you pay (insurance, sick leave, etc.) for your employees?

Once you know these facts (and they are all pretty easy to determine), it takes simple mathematics to come up with a pretty accurate damage estimate.

Problems in Making an Estimate

There are numerous problems with the cost approximation scheme outline above. Lost productivity by users of the affected system(s) is harder to calculate. For this, you may have to do interviews with those affected to find out if they were able to perform other tasks during the incident, or if they were entirely unable to work for periods of time. It is usually a serious over-estimate of damages to simply multiply the number of user accounts by the amount of down-time for a compromised system. Inflated damages are not fair - and the judicial system is geared toward fairness, or at least in theory strives to attain it - and can have a back-lash effect if public opinion, or a jury, does not believe it, or if the figures are found out to lack validity.

Something else that must be added in (for companies who use computers as part of their business) is lost revenue, loss of reputation, and insurance deductibles. These are the kinds of costs that businesses are much better equipped to calculate. Even so, I am going ignore them for the purposes of this article and stick with wage related costs for simplicity.

The biggest problem in generating such an estimate is getting people to keep track of time. Incident handlers and system administrators are often advised to keep notebooks and to take

careful notes of what they did, when they did it, etc., but the reality is that most do not do it. If a case ever makes it to court, these notes are often the best means of ensuring that testimony is accurate as to events and facts surrounding evidence collection and processing. They also can be used to back up estimates of time spent (or to calculate it outright), and for future reference in case someone else must work on the same system.

Underlying this final problem is that most organizations do not require personnel who perform incident handling or own compromised systems keep track of the time they spend dealing with these problems. I have tried on many occasions, with multi-system intrusions - even the "Melissa" virus! - to try to get people to give me time estimates and wage information, and the response rate is very low. Most systems administrators consider this request to be an unnecessary bother. Unfortunately, with an organizational security policy stipulating they do so, I have no leverage to require them to provide such information.

Large portions of staff time is spent administering anti-virus software, and dealing with cleanup of things like Code Red and Nimda, but almost nobody is assigned the task with keeping track of the costs for all of this damage (and spending time cleaning up a virus instead of providing a real service *is* a loss of productivity). Internal Audit organizations are often not tasked with auditing incident costs. Top-level managers and business administrators often don't require that computer security losses (along with depreciation, revenue, and capitol expense losses) be tracked and accounted for. I think this represents a serious failure in organizational leadership that must be corrected, or computer crime will simply increase until it becomes a crisis (at which point the government may be forced to step in, and that may be even worse.)

But how to do this?

A Security Incident Cost Model

In the mid 90's, the Big Ten Universities got together and formed the Incident Cost Analysis Modeling Project (I-CAMP). In 1999, they followed up with a second round, involving a few more institutions and improving on their prior results. The [I-CAMP-II report](#) came out in early 2000.

The I-CAMP projects used a cost model that is very straightforward and easy to follow. It takes into account time, wages, overhead, and incident costs involved in security incidents. I wrote a [FAQ](#) on how to use this model, and included example spreadsheets to make life easy for those

reporting.

The model was first used successfully at the University of Washington (UW) for a multi-system intrusion/sniffer incident in 1998, which stemmed from the "largest security incident in New Zealand history", which was also prosecuted by federal law enforcement in the US. (For more information, see the related articles "[New Zealand Hacker Convicted in Landmark Case](#)" and "[Internet Unfair Trade Practices \(New Zealand\)](#)".) In this incident, the loss for 18 Linux computers on the UW campus was calculated (using the first I-CAMP model) to be \$27794.54 +/- \$4169.18 (an average of \$1544 per host), and the systems at the UW were just part of hundreds of systems around the world that were compromised around the same time period.

Part of the reason the 1998 UW incident was so costly was that the majority of the 18 systems had sniffers, "root kits" [10], and IRC bots installed on them. The sniffers were successful in obtaining many passwords (ironically the intruders' sniffers sniffed *their own logins*, leading to their identification and conviction.) Without a thorough analysis of files left on these systems, many would have remained in the intruders' hands and the incident could have spread to dozens more computers on the UW network, as well as to federal facilities where UW researchers had their passwords sniffed. Such a scenario can result in "down-stream liability" exposure, as your systems are used as stepping stones to attack other sites.

The I-CAMP model was used publicly in 2001 for the [Honeynet Project Forensic Challenge](#). Twelve entries conformed with the requirements of time tracking and cost estimation called for in the Challenge. The average time spent per investigation was 48.0 hours, with the average time spent per investigator being 33.9 hours. Entrants were asked to use a standard salary figure of US\$70,000 and to follow the model as described in my cost analysis FAQ (which works out to be US\$33.65/hour). The [results](#) were quite interesting, and serve to provide an example of the model in use. Here is the "costs.txt" file for one entrant, Addam Schroll of Purdue (one of the I-CAMP Universities, by the way):

Incident Cost Estimate – A Case Study

Indirect Costs

Software

All software used for system analysis consisted of freely available tools

- Coroner's Toolkit free
- Mandrake Linux free
- Redhat Package Manager free
- Tripwire 1.2 for Linux free

Worker Costs Table

Title	Hours	Cost/ Hr.	Total	-15%	15%
Incident Investigator	37	\$33.65	\$1,245.05	\$1,058.29	\$1,431.81
System Administrator(*)	3	\$33.65	\$100.95	\$85.81	\$116.09
Benefits @ 28%			\$348.61	\$296.32	\$400.91
Subtotal (Salary and Benefits)			\$1,694.61	\$1,440.42	\$1,948.81
Indirect Costs			\$0.00	\$0.00	\$0.00
Total Labor Cost			\$1,694.61	\$1,440.42	\$1,948.81
Median Cost +/- 15			\$1,694.61	+/- \$254.20	

(*) Expected time for system reinstallation

As you can see, Addam's analysis came out to just under \$1700, with a margin for error of just over \$250, for 40 hours of work. (That is not too far off from the UW incident costs in 1998.) This is typical of an analysis that requires searching for sniffer logs, lists of compromised hosts, etc., which is often necessary to ensure that an intruder who has gained access to an indeterminate number of your systems is completely removed from the host network. Anything less, and the victim organization may potentially suffer even greater losses (including possible destruction of data to cover the intruder's tracks, or in retaliation for spoiling their fun.)

Comparing all entries for the Forensic Challenge, the average cost per investigation turned out to be US\$2,067.46 (+/- US\$310.12), with the minimum being US\$430.72 (+/- US\$64.61) and the maximum being US\$4,479.50 (+/- US\$671.92). If the same forensic analysis was done by a typical independent consulting firm charging US\$300.00 per billable hour (benefit costs included in this rate), and using the five most extensive investigations for estimating the average (75.4 hours), the "damage" would escalate significantly to US\$22,620.00 (+/- US\$3,393.00).

Conclusion

So it should be clear now that with a little preparation, some policies and procedures in place, and a little discipline during incident handling, fairly accurate damage estimates are not only possible to produce, but rather straightforward.

It is worth noting that the I-CampII model provides a basic cost estimate mechanism. While this is not necessarily a comprehensive measure, it is certainly sufficient to gauge the direct costs of lost labour and productivity, which would certainly provide a tangible starting point for estimating the overall costs of an incident. Furthermore, I don't know of any other models that are publically available to assist a victim in estimating these costs.

Damage estimates are one of the first things law enforcement requires in order to prioritize their response to computer crime. And accurate amage estimates are necessary to ensure that appropriate penalties are levied for these crimes. Without them, the whole legal process is negatively affected and justice is not properly served. By making this a regular part of incident handling, a better result is obtained all around.

Leaving law enforcement aside (because it's a convenient excuse for not estimating losses), if incident cost estimation and tracking were required within all institutions, public and private, the sorely needed data on total security incident costs and trends would be there, and I strongly believe the benefits of investment in prevention would clearly be shown. As computer crime losses become more and more acute, this problem may correct itself because insurance companies or the government demand it, but there is no need to let it get that bad. The current attitude of top managers - see no incidents, hear no incidents, feel no cost from incidents - simply has to change.

Like most things in computer security, the problem - not being able to quantify security losses - is more about lack of preparation, lack of political will, and lack of requirement, than anything else. It only takes a little education and a little discipline to be able to produce results.

[Privacy Statement](#)

Copyright 2006, SecurityFocus