

Impact of a mass cyber security event

Written by Dave Hogan

Overview

There has been intense discussion about the likelihood and impact of a mass cyber security event that compromises several corporations and affects the functioning of critical organizations such as financial institutions, power systems, telecommunication systems etc.

The exponential growth of computer and network usage for critical functions has caused the problem to be of immense concern because one attack can cause several critical systems harm, either because of shutdown or because of loss of data and time. The periodic release of malicious code that exploits vulnerabilities in applications and operating systems and spreads rapidly around the world has intensified this fear. Estimates of losses after each of these events vary widely and are usually overstated tremendously. For e.g. the cost of the recent **Mydoom** worm is expected to cross 4 billion*(Mcafee) this year. Most events that affect a large number of systems on a global level, have till date been restricted to malicious code based attacks, with a few exceptions – the most important of which was the distributed denial of service attack in Feb 2000. While there have been individual cases of intrusion attacks on financial institutions and some utility companies, none have resulted in a large scale shut down till date. There have been attempts made to affect the operation of the Internet by attacks directed at the so called root servers that maintain the Internet backbone and through which a large percentage of Internet traffic flows. None of these attacks have succeeded beyond slowing down the response times for a limited period.

Most scenarios that attempt to illustrate the possibility of a mass event resulting in massive failure, rely on assigning complex and unlikely attributes to the threat and source of threat, and assume extremely flawed design and paralysis at the sites being attacked. Assigning extraordinary capabilities to the attacker and minimizing all defenses is not an accurate simulation of a likely scenario.

The nature of the security marketplace is such that most of the firms that release numbers related to the likelihood and impact of mass events tend to overstate for market reasons. This phenomenon has been called spreading FUD (Fear Uncertainty and Doubt). FUD works against rational risk analysis and results in inaccurate data being generally accepted.

Companies have traditionally approached the concern by relying on technology alone to protect them from such an event. While this approach is inadequate and has to be updated to include a risk management perspective with estimation, mitigation and transfer being considered, the result has been that the effort placed in mitigation and response has been disproportionately high. This results in the risk transferred being less likely than in the case of other exposures and the impact of the event being cushioned by response mechanisms.

This paper will outline the various types of threats that are present, the sources of these threats and undertake an analysis of the likelihood and the severity of the impact of a mass event.

Sources of Threat

There are several sources of threat that are identified as a mass event. It is important to keep in mind that all sources of cyber attack are not necessarily sources of a mass event. Particular kinds of attackers can be specifically highlighted as being most likely to be the source of attack. Sources such as insider attacks on a corporation are restricted by motive, opportunity and capability from causing a mass event.

Terrorism

Terrorist groups have the greatest motive to precipitate a mass cyber attack and the consequent chaos that is caused by the attack. There are a number of hacker groups that are affiliated with known terrorist organizations and while capabilities for full scale information warfare are unlikely to be present, coordinated attacks aimed at causing disruptions in infrastructure and economic damage can be envisaged. ¹ The combination of a physical attack and a cyber attack is especially feared. Till date, however, there has been no mass terrorist cyber attack.

Virus/Malicious code creators

Malicious code has been the main form of threat that has affected a significant percentage of corporations worldwide. Malicious code creators are in the majority found to be young men, usually teenagers, who aim to do one of two things. They are looking to validate their skill in being able to create code that is “smart” and can outmaneuver existing defenses or they are engaged in competitions with other writers and want to prove themselves. In spite of the high frequency of these threats, the actual damage caused has been difficult to estimate, with varying numbers being proffered.

Hacktivists

Hacktivists are groups that are focused on publicizing a political or special interest point of view. An example of a group that may become hacktivists is Greenpeace. The aim of these groups is to spread their message by getting to as many machines as possible. Until now, most hacktivism has been restricted to defacements and targeted attacks at specific corporations. Mass event hacktivism has not yet occurred.

Criminal Groups/Individuals

While several cases of criminal groups attacking financial institutions and other targets with the aim of extortion or theft have occurred ³, these tend to be targeted attacks and not mass events. The use of a mass event as a cover to perpetuate a crime has been hypothesized, but has not occurred yet.

Information Warfare

A foreign government using a mass cyber event to cause damage to the country’s infrastructure and cause economic damage is one other source for a cyber event. However, the sophistication required to cause strategic harm and limiting global fallout is quite complex and similar results are achievable through use of non cyber strategies such as electro-magnetic weapons.

Threats

The actual threat vectors that can cause a cyber event fall into three categories

Malicious Code

Malicious code remains the most likely threat that can cause a mass event, primarily because of the rapid rate of transmission and the ease of creation. There have been several instances of malicious code that affected thousands of organizations. ⁴ Originally mass event causing malicious code took the form of viruses that were propagated through diskettes and over internal LANs. The advent of the Internet and the ubiquitous use of the Internet and email saw the appearance of mass mailing worms such as Melissa and Iloveyou. The worms rarely did any actual damage, but brought down email servers due to overloading. Recently more sophisticated worms that combine the exploitation of vulnerability in software and rapid propagation mechanisms are common as shown by Slammer and MyDoom. There has been parallel evolution in mitigating and preventing the threat from malicious code. Antivirus software that used to, at one point, be updated only once a month, is now updated almost daily and is automated and

centralized. Response times of antivirus companies and application vendors has also reduced significantly, with immediate workarounds being available, even if a patch or update is not immediately ready. In addition, regular patch management and vulnerability scans are fast becoming a standard practice in most companies.

Distributed Denial of Service

The Distributed Denial of Service threat relies on several so called zombies that are pieces of code which reside silently on many computers. These send out requested to targeted companies and overload the capacity of the web server or other device being addressed. Since the software that initiates the attacks is essentially automatic, these can occur at random. However, there has been only one instance in Feb 2000 when multiple companies were affected by this form of attack. The critical root servers that maintain in Internet backbone and are a critical part of the Internet have also been targeted, but unsuccessfully. Many companies face routine denial of service attacks and the hardware and software governing network infrastructure has evolved to greatly minimize the chances of a successful attack. There are configuration guides and specialized solutions also available to combat this threat.

Coordinated Intrusions

A possibility of coordinated intrusions resulting in a mass event can be theoretically envisaged. No such known event has occurred as intrusions usually are targeted attacks. However if multiple groups get together and target critical servers and companies in order to initiate a cascading sequence of failure, they could cause a mass event. Intrusions rely upon vulnerabilities in the applications and infrastructure of servers or other devices to allow unauthorized access. This would require several individuals collaborating and would take a significant amount of time and is likely to be overtaken by response actions before the event becomes critical.

Likelihood of a Mass Event

There are almost no models in existence to predict the likelihood of a mass event. While some extrapolation from history can be done to anticipate the occurrence of mass events such as estimating the probability of a new worm being released in the next year, based on identified flaws and ease of exploitation, these tend to be primarily speculative in nature. The evidence from the past shows that malicious code that affects a large percentage of companies appear very frequently.

In most cases the degree of damage that the code can do is directly correlated to the difficulty of creation and inversely proportional to the speed of propagation and the reliability of the code. In other words, if a given worm is designed to spread extremely rapidly in order to affect machines before they are patched or before an anti-virus fix is released, then the ability of the worm to carry and reliably execute destructive payloads is greatly reduced. Thus even though we have seen several instances of such attacks, none have effectively shut down communications for a significant period of time or have risen to the level of what we would call a cyber hurricane. The other constraint that is placed on the attackers is that of control. The global nature of a mass event implies that the ability to direct and control a threat vector, such as malicious code reduces as the geographic spread increases.

When considering denial of service attacks, the ability to trigger a mass event relies on the successful placing of code that initiates the attack on multiple computers. Currently almost all preventative software, be it antivirus, advanced firewalls, intrusion detection systems etc, have the ability to detect such software, making the first massive denial of service attack most likely being the most successful. This is borne out by the lack of repeat events on the scale of the initial attack in 2000.

Thus in considering the likelihood of a mass event a differentiation has to be made between a event that affects several organizations but is easily countered and an event that causes actual

failure in business and other critical processes for a sufficient amount of time as to have significant real impact. As mentioned before, there are several examples of the former case, with the likelihood of new events occurring being almost certain. There are as of now, no examples of the latter case.

The likelihood of a significant mass event with real impact is also affected by the response mechanisms that have grown, especially over the past three years. These responses now are magnitudes faster than before and there is a general awareness of where to go in case of an attack.

Increased likelihood of terrorist attacks will add to the chances of the targeting of the network infrastructure on a massive level, however in addition to the lack of evidence for the presence of sophistication levels and knowledge required, these groups face the same constraints outlined above, as other sources of threat.

The media, security experts and security vendors continue to hype the chances of a cyber event primarily from a vested interest perspective as large assumptions as to the capabilities of the attacker and the weakness of systems and networks are made. In some cases a doomsday scenario that conflates multiple failures of processes, people and technology, even if they are unrelated to the security breach, is created. Such scenarios additionally assume these failures to occur concurrently and/or opaquely so as to beat all defenses.

There is also an overestimation of the degree of internetworking that exists across companies, especially when it comes to mission critical processes and systems. Segregation is far more prevalent between processes and companies at least as far as critical exposures are concerned.

Based on the above, while there are no accurate quantitative means to measure likelihood, given the absence of an actual event in the past, the likelihood of a mass event currently remains low. This could change if there are significant increases in flawed application design and if network architectures are not created with appropriate stovepiping. Given the increased focus on security from several authorities, ranging from trade associations, governments, large institutions, insurance companies and auditors, the likelihood is probably trending to a decrease rather than an increase.

Impact of a Mass Event

If a mass event were to occur then, are there ways to measure the impact? This is another area where the methodology used to estimate costs vary significantly depending on the approach taken. One method of estimating cost is to be an aggregate the costs of one system being impacted – i.e. if a determination is made that on average a single email server being affected by a worm resulted in losses of x dollars, then n servers affected will cost $x*n$ dollars. This is by an large how global losses are estimated. This results in non real numbers that merely act as a comparative scale and not as a measure of the cost. Even on a comparative level these numbers are deceptive as they are not correlated to actual losses.

Thus numbers such as the \$4 billion estimated as costs incurred due to the MyDoom worm are derived independent of any actual costs and are effectively a theoretical construction.

Costs are also estimated based on surveys. Most surveys rely upon the estimates of the person filling out the survey. This results in a large degree of uncertainty and a difference in understanding of costs. For example, a CIO will probably estimate the costs to mainly include the number of hours spent and people required to fix the affected systems. A CFO may not accurately consider recovery costs, but might include costs of interruption and productivity throughout the company. Estimates such as these reduce to guesswork as none of them capture actual costs, but rely on speculation and assumption. The estimation of productivity costs pose a

particular problem as it is next to impossible to accurately gauge the impact on productivity. Salary based costing for productivity is flawed as alternate productive work accomplished in the period has to be considered and applied to reduce the cost. Also, salaries are in some cases not correlated to the productivity at all, for example, if a senior executive is involved in the recovery and earns \$1 million a year, a hourly breakdown of that number does not reflect losses due to the incident. These cost estimates are especially not useful when applied to mass events as the inaccuracies are inflated over a large population.

A bottom up analysis of the cost parameters show three primary categories of loss 1) Loss of data through theft or corruption 2) Interruption of business and recovery costs 3) Liability costs

Data loss is easily responded to by having appropriate back up and recovery processes in place. Such processes are nearly universal today, with at least backup to tape being an essential part of every business. In order to have a significant impact a mass event would have to also affect the backup systems, which are unlikely as these are designed with the assumption of primary systems being lost and are hence segregated. Business interruption and recovery costs are also addressed through the presence of business continuity plans, redundant systems and backup. In addition, several organizations have, or have been required to do so by regulation, invested in off site facilities that will act as alternate locations in case of the primary systems being affected. The presence of redundant systems is especially true in critical networks such as power networks or telecommunication networks. Liability costs have a lower impact in the case of a mass event as by definition a mass event assumes multiple points of failures and blame is not likely to attach to individual parties.