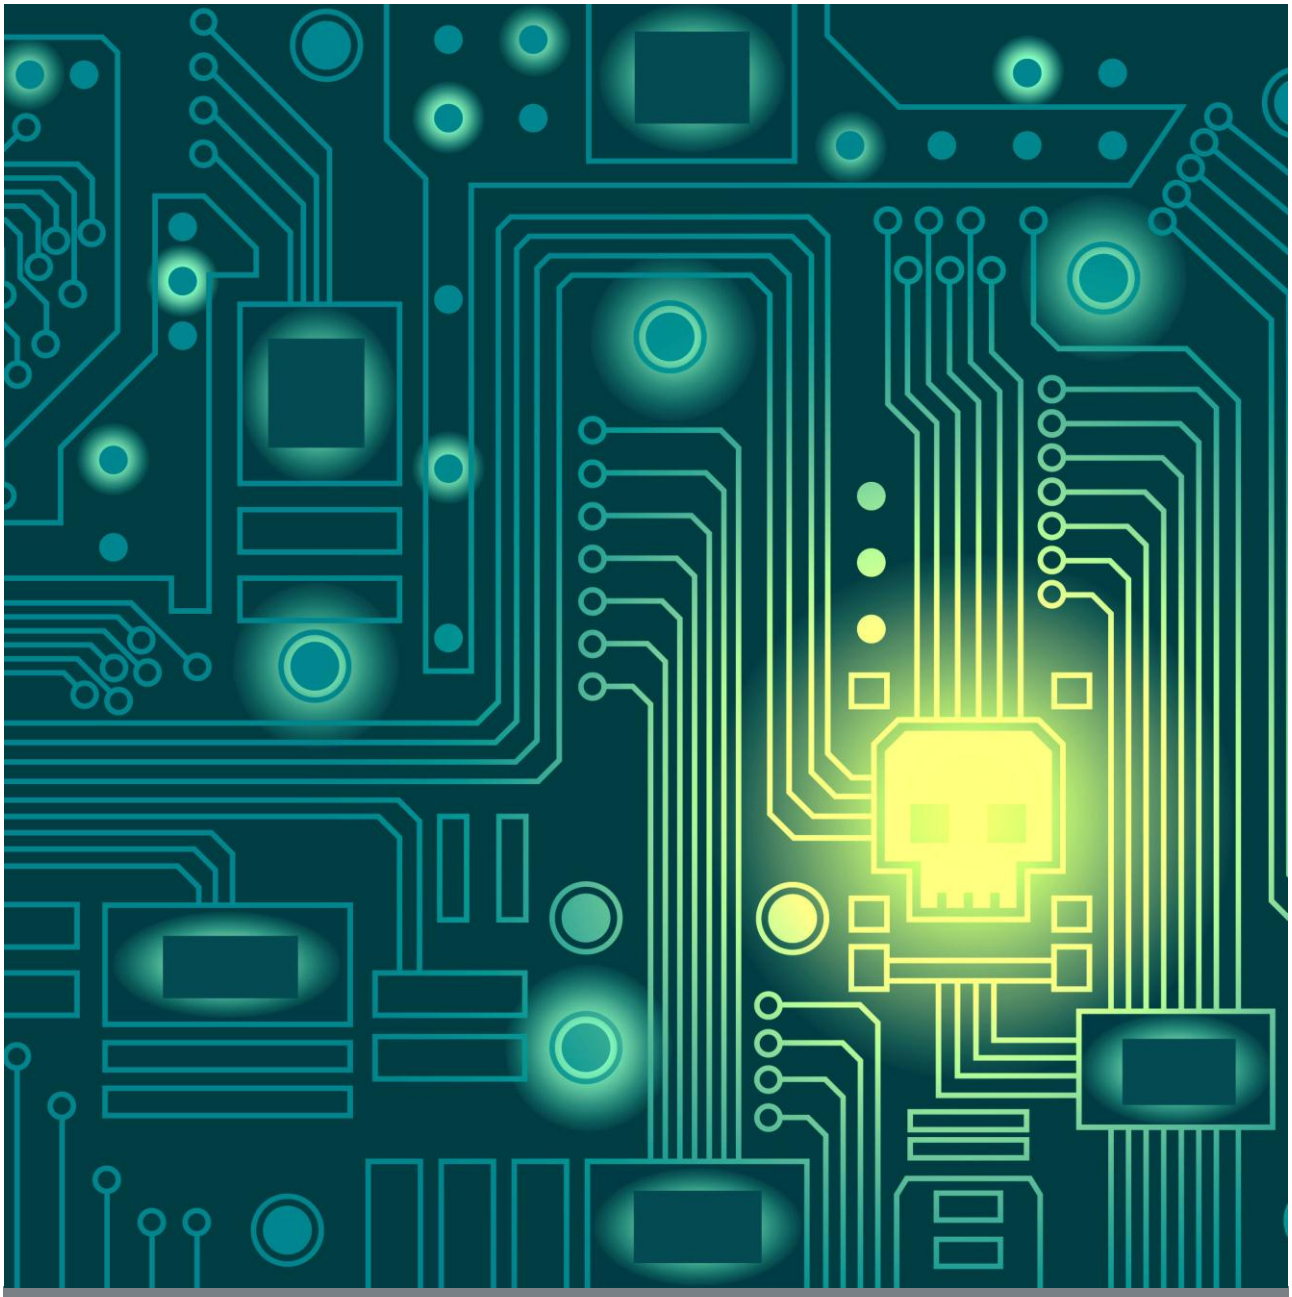

Cambridge Risk Framework

Profile of a Macro-Catastrophe Threat Type

Technological Catastrophe

Cyber Catastrophe

Centre for Risk Studies

UNIVERSITY OF CAMBRIDGE
Judge Business School

Simon Ruffle, Andrew Coburn,
Daniel Ralph, Gary Bowman

Working Paper 201307.02
July 2013
Availability for download at
www.risk.jbs.cam.ac.uk

Cambridge Centre for Risk Studies
University of Cambridge Judge Business School
Trumpington Street
Cambridge, CB2 1AG
United Kingdom
enquiries.risk@jbs.cam.ac.uk
www.risk.jbs.cam.ac.uk

Acknowledgement

This work was carried out with partial support from

NANYANG TECHNOLOGICAL UNIVERSITY

Institute of Catastrophe Risk Management

Cambridge Centre for Risk Studies Working Paper Series

Cambridge Risk Framework

Profile of a Macro-Catastrophe Threat Type

Cyber Catastrophe

Simon Ruffle, Andrew Coburn, Daniel Ralph, Gary Bowman

July 2013

Abstract

The cyber threat landscape is evolving with new aspects of the threat emerging all the time. This report examines the risk from a business perspective. We have built a framework for classifying cyber threat and compiled a database of significant past attacks. We examine approaches to measuring vulnerabilities to attacks and whether the risk has industry sector and geographic variations.

Keywords:
Catastrophe, Scenario, Cyber, Risk, Threat, Vulnerability, Theft, Interruption, Damage, Magnitude.

Authors:

1. Director of Technology Research, Centre for Risk Studies, Judge Business School, University of Cambridge
2. Senior Vice President, Risk Management Solutions; and Director of External Advisory Board of the Centre for Risk Studies, University of Cambridge
3. Professor of Operations Research, Judge Business School, and Academic Director of the Centre for Risk Studies, University of Cambridge
4. Research Associate, Centre for Risk Studies, Judge Business School, University of Cambridge; Cambridge-Singapore Fellow; partial support from the Institute of Catastrophe Risk Management, Nanyang Technical University, Singapore, gratefully acknowledged

* Correspondence to: Simon Ruffle, University of Cambridge Judge Business School – Centre for Risk Studies, Cambridge, UK, CB2 1AG. Email address: s.ruffle@jbs.cam.ac.uk

1 Overview of the Threat

1.1 Definition

A Cyber Catastrophe is an information technology based attack, either malicious or accidental, that afflicts multiple companies or social sectors, causing interruption to business activities, theft of high value information and damage to systems, which results in a substantial interruption to normal commercial productivity for more than one week, costs $10 billion or more to repair and restore to pre-event levels of functionality, or causes at least one country to lose at least 1% of its Gross Domestic Production.

1.2 Summary of the Threat

Cyber threats cover a wide range of malicious activity that can occur through cyberspace. Such threats include web site defacement, espionage, theft of intellectual property, denial of service attacks, and destructive malware[1].

There are numerous recent examples demonstrating the breadth and complexity of the cyber threat landscape: individual computers are attacked in people's home with viruses that attempt to extort money; 'hacktivists' post videos that threaten governments; China is accused of sustained cyber espionage directed at western companies; malicious software damages equipment in a nuclear facility; organised crime employs hackers to enable them to steal $45 million in cash from ATMs in just twelve hours; stock markets react to hoax information posted on news feeds by state sponsored hackers; and a young hacker claims to have found a way to interrupt navigation systems in aircraft whilst in flight using a smart phone.

1.3 Stakeholders

Cyber risk is a fast growing emerging threat. Different constituencies have an interest in – and often a self interest in – this topic.

- IT / Security: Concerned mainly with day to day defence against cyber attack and particularly interested in the technology of the threat. Security companies abound and are looking to sell cyber defence products and services to IT departments.
- Military science: Concerned with understanding the battle going on in cyber space. Interested in attack and defence postures and the resources and covertness of attackers.
- Criminology: Concerned with understanding what crimes have been committed, methods of prosecution and sentencing policy.
- Regulation / Standards: Looking to improve cyber security through regulations and standards.
- Policy: Governments, industry bodies and institutions such as the EU looking to improve resilience to the cyber threat through policy decisions.

[1] Caitlin Hayden, spokeswoman for the White House National Security Council quoted in The Verge 14th February 2013

2 The Threat to Companies

2.1 Overall Losses

A 2011 report by the Cabinet Office and Detica Limited [7] estimated a total cost to the UK economy of £27bn annually, and published a distribution amongst different crimes. Attempts to estimate an annual US total cost have resulted in figures ranging from $250bn to $1tn. However these estimates are controversial. The Detica report was greeted with widespread scepticism and its estimates of substantial losses due to IP theft and espionage have been criticised as lacking in evidence [16].

2.2 Company Losses

Reliable data on individual company losses from cyber attacks is difficult to obtain. Companies are often concerned about reputation damage if they go public with losses due to a cyber attack. Even within companies, IT departments may want to shield senior management from details of breaches in security. This may change with a trend for regulators to start demanding disclosure of cyber breaches as is happening in the US with the Securities Exchange Commission [5] and forthcoming in the EU according to ENISA [10].

The average direct cost varies widely but according to the annual Ponemon Cost of Cyber Crime study [12], which surveyed 199 companies in five countries, it averaged $9m per company per year in the US in 2012 (see figures 1 and 2). To these costs should be added the indirect costs – lost business opportunities, staff morale and company reputation that although difficult to estimate, can be greater than the direct costs suffered.

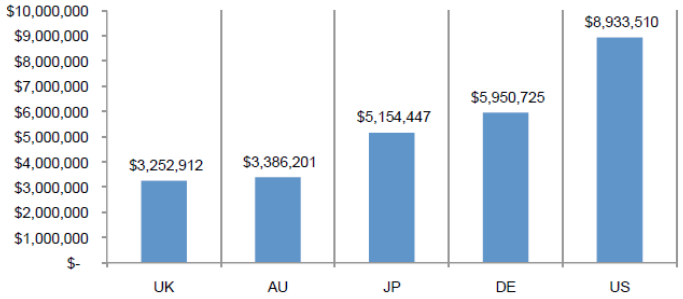

Figure 1: 2012 total average cost of cyber crime in five countries (USD, n=199 companies)

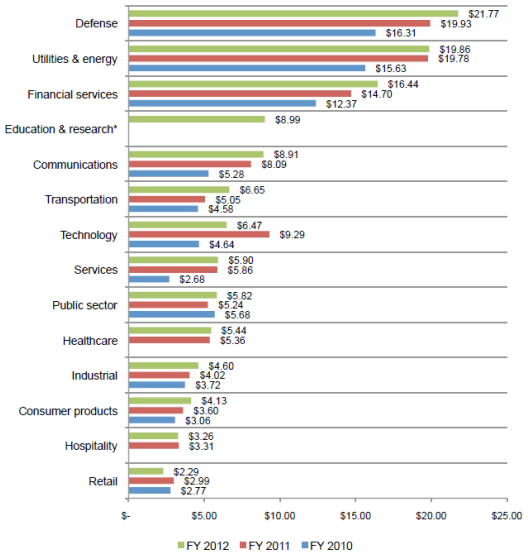

Figure 2: Average annualized cost US by industry sectors 2010 – 2012 (millions of USD)

2.3 Target of the Cyber Threat

It is useful to distinguish between the general classes of target of a cyber attack:

- ICT systems: the traditional target of cyber attacks causing essentially 'virtual' losses.
- Physical systems: Increasing numbers of physical services are being connected to the internet, from power stations through to smart meters in people's homes. Cyber attacks can cause actual physical damage.

2.4 Primary Characteristics of the Cyber Threat

Cyber threats can be categorised very generally as acts of warfare, terrorism or crime, or as non-malicious (e.g. the case of a system malfunction). This report focuses on cyber crime.

There is a shortage of robust loss data from past cyber events. There is little agreement on the overall costs – although some cost estimates have been published they are highly controversial and estimates for the same event from different commentators vary by orders of magnitude. For this reason we do not attempt to characterize attacks by estimated cost but by their consequences.

Broadly speaking there are three categories of harms from cyber threats:

- **Theft**

  Actions that extract data items that are of *value to the perpetrator* and breach the *confidentiality and duty of care* of the data holder

  - Espionage of industrial secrets, intellectual property, corporate know-how
  - Theft of money; transfer of funds; appropriation of assets, investments, stocks and bonds
  - Obtaining customer records; Databases of personal information; trading records; confidential business transaction data
  - Obtaining identity information; passwords; credit card details; consumer data

- **Disruption**

  Actions that *interrupt business functionality* for a period of time, or *degrade productivity* of commercial operations, transactions, or communications

  - Denial of service for internet-based businesses
  - Blocking or degrading communications, emails, transaction orders
  - Downtime of public facing websites, internal servers, cloud resources and individual workstations

- **Damage**

  Actions that *corrupt data*, or *damage software, systems, or physical equipment*, and require resources to repair or restore, and *incur costs, liabilities and reputational damage*

  - Hacks that corrupt or delete data or software
  - Attacks that disable servers, hard drives, individual computers
  - Subverting control systems to trigger damage to physical equipment or systems

2.5 Cyber Crime

It is difficult to finding a rigorous definition of cyber crime as it particularly impacts business. In this this list, which is a combination of [7], [8] and [16], ***I*** designates events that have resulted in insurance claims, and ***C*** designated events that have the potential to be catastrophic:

- Malware and hacking attacks causing damage to computers and/or networks and/or data. Includes sabotage. ***IC***
- Computer-based theft using stolen personal account information or by direct hacking of computer systems. Might usefully be split between financial sector and non- financial sector.
- Scareware such as fake antivirus
- Extortion and scams such as stranded traveller. ***I***
- Physical Infrastructure attack causing physical damage and/or personal injury. ***C***
- Denial of Service Attacks. ***IC***
- Breaches of Personally Identifiable Information (PII), Private Health Information (PHI) and Credit Card and Other Financial Data: Fines and class action suits for privacy violation are increasing. ***I***
- Theft of IP and trademarks. ***IC***
- Espionage. ***C***
- Fiscal fraud
- System or network malfunction accidental or due to negligence. ***C***

2.6 Sectorisation of the Cyber Threat

Cyber threats can either be indiscriminate or they can be targeted at particular industrial sectors or organizations. Cyber threats can therefore also be characterised by the industrial sector or sectors targeted.

2.7 Threat groups

The motivation of perpetrators of cyber attacks can be political, military, financial, revenge, or just curiosity or notoriety. As cyber attacks become more sophisticated the resources behind a particular actor become an important measure of magnitude, as are the degree of covertness and the duration of the attack before discovery. Perpetrators can be divided into groups as follows:

- Intelligence services / electronic armies: Many states now operate cyber intelligence specialists (GCHQ in the UK; NSA in the US) which are actively involved in cyber offence and defence.
- Terrorists
- Industrial Spies
- Organised crime: Organised crime has moved into cyber space with 'backer – hacker – casher' style operations.
- Insiders: Still a worrying element in any organization is the disgruntled employee who has access to passwords and sensitive systems.
- Hacktivists: Groups with an activist or anarchist agenda now have many channels for expounding their views and launching attacks.

- Individual Hackers: Of less concern as cyber attacks become more sophisticated and require increasing resources, but the individual hacker still has potential.

Of course it may prove impossible to identify perpetrators.

2.8 The Costs to Companies of a Cyber Attack

Direct Costs include (from [12]):

- External consequences (information loss, business disruption, revenue loss, and equipment damage)
- Internal consequences (detection, recovery, ex-post response, containment, investigation and incident management)

Coverage from existing cyber liability insurers includes (from [8]):

- Crisis Services (forensics, notification, credit monitoring, legal counsel)
- Legal Damages (defence and settlement)
- Fines (PCI and regulatory)

2.9 System Malfunctions

A cyber catastrophe may not be the result of a malicious attack but may result from a system malfunction, such as a routine system upgrade that introduces errors into the system. This may be accidental or the result of negligence.

3 Protection against cyber risk

3.1 Improving Cyber Security

Cyber security is of wide concern and many agencies and stakeholders are involved. Policy is directed from governments and EU level; standards agencies have produced standards (such as ISO 27001:2005); there is academic research; and numerous private security companies offer their services. Companies are urged to improve their day to day security (see for example figure 3) and to make cyber security a corporate governance concern.

1. Network security
2. Malware protection
3. Secure configuration
4. Manage user privileges
5. Information risk management regime
6. Monitoring
7. Removable media controls
8. Incident management
9. User awareness and education
10. Home & mobile working

Figure 3: '10 Steps to Cyber Security – Executive Companion', CESG, BIS, CPNI, Cabinet Office [14]

3.2 Policing and Prosecution

In a recent paper by Anderson at al, *Measuring the Cost of Cyber Crime* [16] they conclude: "…we should perhaps spend less in anticipation of computer crime (on antivirus, firewalls etc.) but we should certainly spend an awful lot more on catching and punishing the perpetrators." and "…cybercrime is now the typical volume property crime in the UK and the case for more vigorous policing is stronger than ever".

3.3 Insurance for business interruption

'Business Interruption' (BI) insurance is traditionally linked with insurance policies for building damage – the insured property has to be damaged (by fire, flood etc.) for compensation to be paid. However companies are increasingly getting insurance coverage for loss of earnings due to failures in their communications, utilities or essential services even if their own property isn't damaged (known as 'Contingent BI'). Network issues are probably dominating their thinking in this context. In addition, insurers are beginning to offer more specific insurance products to cover network failures and virus infection. Some insurers see this as a major growth area.

3.4 Insurance for liability

Chubb, Wells Fargo and Chartis are examples of companies offering cyber liability cover [8].

4 History

We have constructed a database of significant cyber attacks since 2000. Loss figures are taken from various sources and must be regarded with great suspicion.

Table 1 Catalogue of major cyber events from 2000 to 2013

ILOVEYOU	2000	$15Bn
MafiaBoy	2000	$1.2Bn
Code Red	2001	$2Bn
SQL Slammer	2003	$750m - $1Bn
MyDoom	2004	$38Bn
Sasser	2004	$500m
Titan Rain	2004	
TJX	2005	$250m
APT1	2006	
Conficker	2007	$9Bn
Zeus	2007	$70m
Estonian Cyber attack	2007	
Heartland	2008	$140m
RBS WorldPay	2008	$9m
Stuxnet	2010	
Aurora	2010	
Epsilon	2011	£225m to $4Bn
Sony Playstation	2011	$1 – 2Bn
Citigroup	2011	$2.7m
RSA	2011	$66m
Operation Ababil	2012	
Shamoon	2012	
Flame / Skywiper	2012	
The Unlimited Operation	2012	$45m

CloudFlare	2013
ObamaTwitter Scare	2013

Statistical analysis of the frequency and severity of virus impacts over the past 13 years provides an estimate of the probability of experiencing different scales of loss in the next twelve months. The loss probability distribution is also adjusted to take into account the increasing reliance of commercial activity on information technology, the increased investment in network security and the growing prevalence and sophistication of modern cyber attacks. The loss level that would be achieved with 1% probability is estimated at over $90 Bn.

4.1 Historical event examples

Conficker Worm (2008 to present)

This worm is most closely related to the attacks that affected the early internet, such as ILOVEYOU and MyDoom. It originally targeted a vulnerability in Microsoft Windows which was fixed by Microsoft a long time ago and no longer exists in the latest versions of Windows. Once the worm has infected its host it no longer needs the vulnerability. The worm has numerous methods of spreading itself and has the ability to be upgraded by its originators (perpetrators unknown, but thought to reside in Ukraine). There have been five versions of Conficker to date (A through E) each becoming increasingly malicious. The worm is currently 'in remission' having had its links cut to its command and control servers. 1.8 million PCs are still infected, years after the initial infection. McAfee estimates the total global loss from Conficker at $9bn.

Stuxnet (2009)

Stuxnet was a game changer – although losses were not large, it made headlines because malicious code was seen deliberately targeting physical critical infrastructure. Stuxnet targeted industrial systems under control of the Siemens PCS7 SCADA (Supervisory Control and Data Acquisition) system. The specific target appears to be the Natanz Nuclear Facility in Iran where it spun 1000 nuclear centrifuges past their operating limits and destroyed them. It also caused damage to other industrial systems under control of the Siemens system; the oil industry seems to have been particularly affected. The perpetrators are generally considered to be the US and Israel.

The Unlimited Operation (2012 – 2013)

This was an organized crime of the type known as a 'backer – hacker – casher' attack. Hackers were paid to compromise two banks in the Middle East — the National Bank of Ras Al-Khaimah PSC in the United Arab Emirates and the Bank of Muscat in Oman — where prepaid debit card accounts were breached and the withdrawal limits normally placed on debit card accounts were removed. Then teams of operatives on the ground (the cashers) were provided with corresponding compromised debit cards which were used to extract cash from ATMs in various places around the world. A map of their extractions from ATMs in Manhattan is shown at figure 4.

There were two attacks – the first in December 2012 lasted 3 hours and netted $5m. A second attack in April 2013 netted $40m in 9 hours. Some cashers were later arrested.

Figure 4: ATM cashing by the Unlimited Operation cell in Manhattan.

APT1 (2006 onwards)

APT1 ('Advanced Personal Threat' 1) is a large scale economic espionage attack by China on western nations that seems to have taken place over many years. Companies in industry sectors that match the strategic industries identified in the current Chinese Five Year Plan are particularly targeted. The preferred mode of attack is 'spear phishing' where individuals are targeted in organizations. The key document that has identified this attack is by the Mandiant Corporation [17] – it identifies the perpetrators as Unit 61398 of the Chinese People's Liberation Army (PLA) and goes as far as naming individual actors and the building in Shanghai from which they operate. There is concern that the quantity of western intellectual property and strategic company information that has fallen in to Chinese hands has cost the West dearly.

5 Magnitude scale

We categorise cyber events by their type of harm – theft, disruption and damage – and we are developing a magnitude scale for each of these.

Magnitude Scale Value	Number of personal data records stolen (millions)	Amount of high value data breached (GBytes)	Value of IP stolen	Money stolen ($m)	Example events	How often this might occur worldwide Return period (years)
Magnitude **1** Theft Event	10m	50 Gb	Individual secrets of minor importance	$10m		Monthly (10 observed worldwide each year)
Magnitude **2** Theft Event	100m	500 Gb	Individual major secrets or many minor secrets	$100 m	Sony Playstation Heartland Unlimited Op RBS Worldpay	One a year (10 in past decade)
Magnitude **3** Theft Event	1 billion	5,000 Gb	Multiple major secrets from many companies	$1 billion	APT1	One a decade (largest event seen in past 10 yrs)
Magnitude **4** Theft Event	10 billion	50,000 Gb	Individual vital secrets or many major secrets from many companies	$10 billion		1 in 100 per year
Magnitude **5** Theft Event	100 billion	500,000 Gb	Many vital secrets from many companies	$100 Bn		1 in 1,000 per year

Figure 5: Theft Magnitude Scale

Magnitude Scale Value	Total workstation-hours downtime	Total server-hours downtime	Total website-hours downtime	Total business-hours downtime	Example events	How often this might occur worldwide Return period (years)
Magnitude 1 Disruption Event	1 day	1 day	1 day	1 day	SQL Slammer	Monthly (10 observed worldwide each year)
Magnitude 2 Disruption Event	1 week	1 week	1 week	1 week	ILOVEYOU MyDoom Conficker	One a year (10 in past 10 years)
Magnitude 3 Disruption Event	1 month	1 month	1 month	1 month		One a decade (largest event seen in past 10 yrs)
Magnitude 4 Disruption Event	6 months	6 months	6 months	6 months		1 in 100 per year
Magnitude 5 Disruption Event	3 years	3 years	3 years	3 years		1 in 1,000 per year

Figure 6: Disruption Magnitude Scale

Magnitude Scale Value	Number of PCs/ servers rendered inoperable	Data destroyed and unrecover-able (Gb)	Liability costs $m	Reputation loss costs + fines $m	Physical repair costs $m	Example events	How often this might occur worldwide Return period (years)
Magnitude 1 Damage Event	1,000 to 10,000	100		1	1-10 m	SQL Slammer	Monthly (10 observed worldwide each year)
Magnitude 2 Damage Event	10,000 to 1 million	1,000		10	10-100m	Sony Playstation	One a year (10 in past 10 years)
Magnitude 3 Damage Event	1 million to 10 million	1		100	100m to 1000m	Stuxnet	One a decade (largest event seen in past 10 yrs)
Magnitude 4 Damage Event	10 million to 100 million						1 in 100 per year
Magnitude 5 Damage Event	100 million to billions						1 in 1,000 per year

Figure 7: Damage Magnitude Scale

6 Vulnerability

There are various measures of the vulnerability of an organization to cyber threat, of which the Security Effectiveness Score (SES) [4] seem to be the most developed. The SES has been developed by PGP Corporation and Ponemon Institute and is used by Ponemon in its annual encryption trends survey to define the security posture of responding organizations. The SES is derived from the rating of 24 security features or practices. This method has been validated from more than 30 independent studies conducted since June 2005. The SES provides a range of +2 (most favourable) to -2 (least favourable). Hence, a result greater than zero is viewed as net favourable.

The Ponemon 2012 Cost of Cyber Crime: United States [12] ranks losses in companies by their SES. Their study shows that companies with a better SES, i.e. which are less vulnerable to cyber threat, tend towards lower losses (see figure 8).

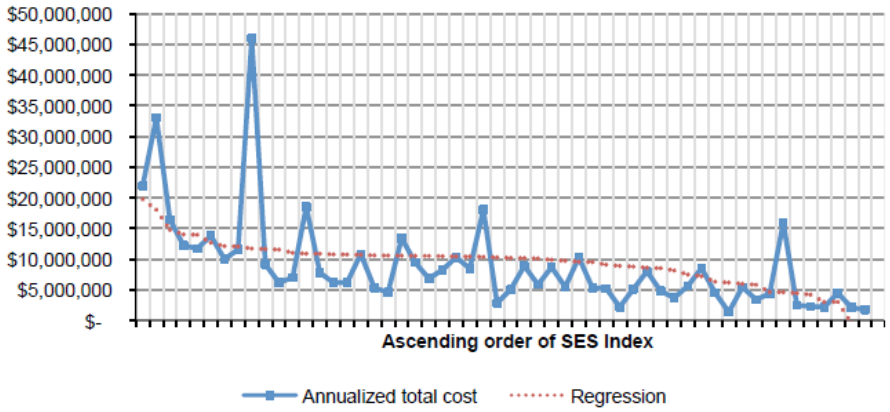

Figure 8: 2012 Annualised Cost in descending order by SES ranging from -1.19 to +1.69

7 Geography of cyber risk

Cyber attacks do not have 'footprints' in the same way as, say, earthquakes. However there are geographic variations in general computing infrastructure and general social and cultural attitudes to computer security.

[12] has attempted to distinguish different types of attack by geography (see figure 9) and does show a slight variation. For example companies in the US sample are more likely to experience attacks by malicious insiders than in Japan – is this showing something in the American vs Japanese corporate culture?

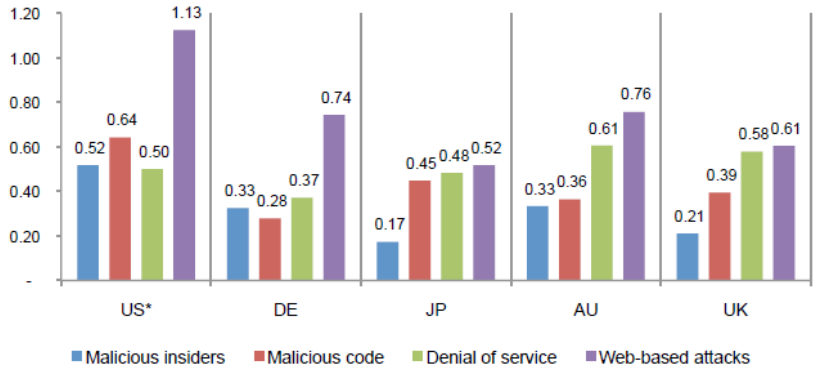

Figure 9: Adjusted frequency of four types of cyber attack by country

8 Scenario Specification

The Sybil Logic Bomb: A fictional example of a high-impact cyber attack

Background

The Sybil Corporation is a software vendor that produces the market leading relational database. Established in the late 1970's, Sybil's databases run on all common operating systems and are employed in most sectors of the global economy with a particularly strong uptake in the corporate world. Sybil databases tend to sit on servers at the heart of corporate IT systems storing data from all aspects of the business. Many third party vendors offer systems and services that are built upon the Sybil database.

Working in Sybil's software development division in Redwood City, CA, USA, is a thirty year old mathematician employee responsible for the computational and arithmetical software code. He is becoming increasing antagonistic to global capitalism and has recently become interested in and sympathetic to the activities of the Anonymous 'hacktivist' collective. He decides to covertly and maliciously modify some of the source code of the Sybil database to which he has access, knowing that the next routine upgrade (or 'patch'), which will be issued to all users of the product, will include his modified code.

Stage A: Preparation

The employee decides to modify the floating point algorithm of the Sybil database to produce errors in results that are in the range -10% to +10% away from the correct value[2]. The error is only to occur if any of the input variables match the last three numbers of the host computer's manufacturer's serial number.

This makes it hard to detect. By targeting the floating point algorithm, routine financial computations found in transaction processing will be largely unaffected – because errors in this type of calculation would be quickly spotted by the daily checks and balances of accountants and bookkeepers. Floating point tends to be used more in design, modelling, decision support and reporting activities, where small errors in the +-10% region are harder to detect. The additional filter of the input matching the machine serial number means that a specific problem cannot be replicated on a different machine but will be consistent on the host machine. This undermines attempts to replicate the issue, which will be the first thing support staff will try when attempting to diagnose the problem.

The employee then covers his tracks by altering the date on the source file, and the meta data in the code repository back to their original conditions before the modification was made, making it difficult for his managers to spot the change. The employee uses his knowledge of the Quality Assurance procedure in Sybil – specifically what tests are run on the floating point algorithm – to further optimise the compromise so it will not be detected by QA procedures.

The employee then decides that only companies in the western corporate world will be affected and adds a filter that reads the company name from the database license file and only applies the compromise if the company name includes 'plc', 'co' or 'inc'.

Finally the employee adds a 3 month time delay (or 'fuse') after the upgrade is applied before the compromise activates. This will reduce the likelihood that emerging issues are associated with the upgrade.

Stage B: Attack activation

The compromised software is released as part of a routine upgrade for all Sybil customers. There is widespread belief in the IT industry that software upgrades should be applied as soon as possible, and naturally customers will trust such an upgrade from a well respected company like Sybil. Good corporate cyber security practice which is on the lookout for worms, phishing and insider attacks will not spot a compromised upgrade; in fact it encourages rapid application of upgrades.

Customers vary widely as to when they actually apply the upgrade. Some do it immediately. Many run the upgrade through their own QA testing before they apply it but it is unlikely to include rigorous testing of the floating point algorithm. Some have been waiting for it and are keen to apply it because it contains a bug fix or new feature they need. Some wait to see if other companies experience problems with it, but as there appear to be none, due to the time delay, they install it.

[2] An example of a floating point number is '123.34567'. In computer code this is a different type of data to say an integer such as '123' or a fixed point monetary amount such as '123.45'.

Application of software updates varies by industry sector. On average the financial sector will install upgrades in 6 months, and the industrial sector 18 months.

Stage C: Latency Period

With the compromised upgrade released, and companies beginning to apply it, the scenario moves into the pre-detection or 'latency' period where the compromise is activated but not yet detected.

Once a company has installed the compromised upgrade, and the time delay has passed, all their systems based on the Sybil database start to give the error but it goes unnoticed for a while as it is largely affecting design, modelling, decision support and reporting activities in small ways. Some users never get the error – their data does not contain a value that matches the last three digits of the server's serial number.

In the end some users spot an error. Then begins the procedure of escalating the issue. Initially the user thinks it is an error in their own calculation. Once this is discounted they report it to their IT support desk. IT attempt the replicate the problem on a test machine – which having a different serial number does not produce the error. The issue is reported to Sybil's support team but they also are unable to replicate it. This leads to periods of frustration for users as their problems are not being taken seriously by support teams. Eventually many companies draw the conclusion that this is a hardware fault, so they replace the server. Unless the new machine's last three serial numbers happen to match the old machine, this appears to fix the issue.

Furthermore, these issues are only occurring in a segment of Sybil's market – the western corporate sector, with some sectors lagging behind in their uptake of the compromised upgrade – and Sybil's support team point to thousands of systems worldwide operating without problems. They disown the problem.

Some companies call in security consultants, but they draw a blank as the issue does not show any of the traces of a normal cyber attack – no unauthorised access, no detectable malware, and no known exploits of Sybil.

As it emerges in the IT world that servers running Sybil seem to be experiencing unexplained hardware issues, a certain brand of server is erroneously suspected, resulting in pre-emptive replacement of that brand and avoidance of purchasing that brand in the future by IT departments. This impacts the profitability, brand value and stock price of this server manufacturer.

As time passes the key disruptive consequences of this cyber attack become evident though still no one is making the connection to the upgrade or even to an issue within Sybil. These key consequences be characterised as *impacts on quality,* for example:

- Design systems (such as Aeronautical CAD systems) based on Sybil have started to introduce small random changes in manufactured parts which begin to fail or give degraded performance.
- Modelling and Decision Support Systems (such as a Commodity Trading or Oil Pricing Model) have started to give random erroneous results resulting in loss making trades and price setting that results in loss of profitability.
- Reporting systems (such as MIS and CRM systems) start feeding erroneous data back to managers and boards who make incorrect decisions. Company regulatory filing and annual reports appear with errors in them.
- Process Control Systems (such as can be found in manufacturing and industrial control systems) start producing erroneous threshold values resulting degradation in quality and, in the worst cases, equipment malfunctions.

- Logistics systems start causing shortages of parts to industry and products to consumers resulting in a fall in quality of service in these sectors.

Waves of unease are appearing through various sectors of the world economy. Many companies raise internal alarms about error-prone data and inaccurate internal routines. The rate of corruption is initially slow and difficult to detect but becomes more pronounced over time, but the extent of data corruption is not easily verifiable. Meanwhile as normal routine backup procedures have been running, erroneous data within Sybil, and other company systems to which erroneous data has propagated, is progressively corrupting backups.

News stories start circulating about regular inexplicable costly events taking place in the corporate world. Almost daily events like the following examples are being reported. *These are imaginary but are based on real events caused by erroneous data in corporate modelling, decision support and forecasting systems. See footnotes for references.*

- A book written by two prominent economists which has influenced major political and economic decisions of nation states has key assumptions later proved to be erroneous by a high school student. These assumptions are based on data taken from a compromised Sybil database. [4]
- For two years at a large bank they failed to notice an error in interest rate calculations caused by erroneous data from their compromised Sybil database. When finally revealed it knocks $6.5 (AUS) billion off its market value and costs it $755 (AUS) million.[5]
- Stock markets in the US and Europe are regularly plagued by sudden unexplained drops of around 10%. Analysis of trading reveals technical glitches in the reporting of prices on the exchanges and various alternative trading systems that might have contributed to the drying up of liquidity. These systems are taking their data from compromised Sybil databases. [6]
- An investment firm pays $250m to settle civil fraud charges that it used erroneous data in its quantitive investment model. Senior managers had concealed the error which unknown to them had originally come from a compromised Sybil database. [7]
- A struggling pharmaceutical company is forced to reiterate its yearly and midterm financial forecasts after admitting it contained "out of date planning information" which had resulted from data from a compromised Sybil database entered into a forecasting spread sheet. Its stock price falls as a result.[8]
- A major live event organiser has several occasions when their events are massively overbooked causing disruption as large numbers of people arrive at venues and cannot gain entry. This is being caused by the compromised Sybil database the runs their booking systems. Audiences at live events fall and the company's stock price crashes.[9]
- An "accounting error" forces the resignation of the CEO of a large outsourcing specialist after breaking banking agreements on debt. The error is unbeknown to them caused by incorrect valuation of a pension fund deficit caused by a compromised Sybil database.[10]
- A large publicly traded power generator and marketer of electricity and renewable energy takes a $25m charge after it lands more power transmission hedging contracts than it bargained for at

[4] http://www.cepr.net/index.php/blogs/beat-the-press/how-much-unemployment-was-caused-by-reinhart-and-rogoffs-arithmetic-mistake
[5] National Australia Bank (http://c3integrity.com/blog/posts/the-cost-of-bad-data)
[6] Flash Crash US Stock Market 6th May 2010
[7] AXA Rosenburg (http://www.eusprig.org/horror-stories.htm)
[8] AstraZenica (http://uk.reuters.com/article/2012/01/09/uk-astrazeneca-idUKTRE8080BX20120109)
[9] Locog / Ticketmaster "Spreadsheets behind Olympic data misentry"(http://www.eusprig.org/horror-stories.htm)
[10] Mouchel – (http://www.express.co.uk/posts/view/276053/Mouchel-profits-blow)

higher prices than it wanted to pay. The error came from ranking bids based on a compromised Sybil database.[11]

- In an automobile manufacturing facility, a 9-foot robotic swings around 180 degrees despite the controller for the arm being in standby mode. 3 workers are killed. The SCADA system controlling the facility was being fed operating parameters from a compromised Sybil database.[12]
- In an integrated circuits fabrication plant a system controlling the creation of integrated circuits in the fabrication plant hangs. The outcome is the destruction of $50m worth of wafers. The SCADA system controlling the facility was being fed operating parameters from a compromised Sybil database. [13]
- A gas utility is not able to send gas through its pipelines to its customers for 24 hours due to its Process Control System being fed incorrect operating parameters from a compromised Sybil database.[14]

Stage D: Detection

30 months after the release of the compromised upgrade Sybil finally recognises the problem as being theirs and quickly release an urgent security upgrade that removes the compromised code. Sybil apologises for the defect but announces it only affects a minority of its user base and points to its limited warranty clause in its software licence.

In their security bulletin Sybil describe the timeline of the compromise thus companies can identify the period over which they were infected based on the date of installing the upgrade. At this point the companies who installed the upgrade immediately will have been affected for 27 months (as the compromise had a 3 month time delay). Using industry averages of time taken to install upgrades, the financial sector will have been infected for 21 months and the industrial sector for 9 months. On average companies will have been affected for 15 months.

Stage E: Response by organisations

Awareness of the impact of the Sybil compromise dawns on the corporate sector. An investigative journalist writes an article 'The Sybil Logic Bomb' explaining how the previously assumed unrelated and unconnected events at can all be traced back to the Sybil compromise and how there is now corrupted data all over the corporate sector that has impacted decisions and quality and is now embedded deeply into backup systems.

He points out that the rectification of the defect by Sybil will have no effect on the data already corrupted and the problems will continue.

There is a collapse in trust. Events continue to occur – now, no one knows if are connected to corrupted data caused by the Sybil Logic Bomb or not.

Panic begins to spread around companies. No one knows which data is compromised and which is not. Because Sybil acts as a basis for so much business activity there is no guarantee that compromise is limited to the Sybil database – corruption could be any part of the business.

- Some companies wipe hard drives and go back to the last clean backup, but consequently lose many months of work. Most companies do not even have this option.

[11] Transalta – (http://www.globeinvestor.com/servlet/ArticleNews/story/ROC/20030603/2003-06-03T232028Z_01_N03354432_RTRIDST_0_BUSINESS-ENERGY-TRANSALTA-COL)
[12] See reference [19]
[13] See reference [19]
[14] See reference [19]

- Most companies decide to carry on fixing issues as they arise, but if they suffer problems it is difficult to tell whether or not they are related to the Sybil Logic Bomb. This creates uncertainty and loss of confidence in management teams.
- A plane crashes at San Francisco airport – newspapers publish articles with headlines like 'Did Sybil Logic Bomb crash SFO jet?'
- Consumers and industry becomes mistrustful of supplies and products. Some companies decide to recall all products manufactured since the installation of the upgrade, and refund customers.
- Trust in the corporate sector is damaged and stock prices fall.
- The resulting malaise in the corporate world caused by fear and uncertainly has an impact on productivity and is seen in reduced GDP figures for many western nations. It takes many years for companies to recover.

Stage F: Rework

Each individual company carries out internal audits to establish what parts of their computer systems have been affected by the Logic Bomb. Many call in consultants to detect and analyze the problem. Data restitution is the priority. In extreme cases, some companies have to poll customers to rebuild data from scratch. Internal staff time is absorbed throughout the organization as IT departments scramble and senior managers attempt to minimize the impact on customers and business operations. Many companies re-install software and data systems, reconfigure firewalls and instigate new quality assurance measures at considerable expense. Legal counsel is brought in and consultants and staff spend time preparing a potential case for legal action against the perpetrator.

Stage G: Aftermath

Companies absorb most of the costs themselves. Although more than a third of major corporations have insurance policies that incorporate some protection against cyber crime losses, the number of medium and smaller companies that have insurance is less than 2%. Several individual companies affected face losses of over $100 million from lost revenues, shortfalls in assets, consultancy costs and extra expenses incurred for restitution and repair. The insurance recovery is less than 1% of the overall direct costs that result from the attack. Businesses hit by the virus take a long time to recover from the scale of the unexpected costs and the loss of revenues.

If the Sybil Corporation is seen to have handled the situation well they may suffer no more than reduced market share. If they didn't they will be the target of class actions.

New regulations are enacted aimed at improving quality control in software. Software companies are prohibited from hiding behind limited warranty clauses and this raises the cost of software by 20%.

In summary the consequences of the Sybil Logic Bomb can be likened to Asbestosis. Complicated problems will linger in global systems for years costing companies to sort out. Going forward, no one will be sure that the infection has been entirely eradicated, and it may never be declared to be formally fixed.

And...

The disgruntled employee, who left Sybil long ago, has since worked for two more database vendors…

9 Bibliography

Recommended reading

Cyber War: **The Next Threat to National Security and What to Do about It,** Richard A. Clarke.
http://www.amazon.co.uk/gp/product/0061962244/ref=oh_details_o01_s00_i02?ie=UTF8&psc=1#_

A Fierce Domain: **Conflict in Cyberspace 1986 to 2012**, Jason Healey.
http://www.amazon.co.uk/gp/product/0674015762/ref=oh_details_o03_s00_i00?ie=UTF8&psc=1#_

Cyber Crimes in a Globalized World: **Country Profiles, Rankings, Patterns & Risks,**
Eugen Mandrila
http://www.amazon.co.uk/gp/product/9730145806/ref=oh_details_o01_s01_i01?ie=UTF8&psc=1#_

Cyber War Will Not Take Place Thomas Rid
http://www.amazon.co.uk/gp/product/1849042802/ref=oh_details_o01_s00_i00?ie=UTF8&psc=1#_

10 Information Resources

[1] Categorizing Threat: Building and Using a Generic Threat Matrix; David P. Duggan, Sherry R. Thomas, Cynthia K. K. Veitch, and Laura Woodard

[2] Data Loss DB; DataLossDB is a research project aimed at documenting known and reported data loss incidents world-wide.

[3] Managing Digital Risk: Trends, issues and implications for business; Lloyd's 360 Risk Insight

[4] Security Effectiveness Framework Study; HP Information Security, Checkpoint and Ponemon Institute

[5] CF Disclosure Guidance: Topic No. 2: Cybersecurity; US Securities and Exchange Commission

[6] Quantitatively Assessing and Visualising Industrial Attack Surfaces; Eireann P. Leveret

[7] The Cost of Cybercrime; Detica / Cabinet Office

[8] Cyber Liability and Data Breach Insurance Claims; NetDiligence

[9] Cyber security guidance for business; BIS, Cabinet Office, FCO.

[10] Incentives and barriers of the cyber insurance market in Europe; ENISA, Neil Robinson, RAND Europe

[11] Cyber Threat Metrics; Mateski, Trevino, Veitch, Michalski, Harris, Maruoka, Frye; Sandia National Laboratories Report SAND2012-2427

[12] 2012 Cost of Cyber Crime Study: United States; Ponemon Institute Research Report

[13] Cybersecurity: Threats Impacting the Nation; GAO, Statement of Gregory C. Wilshusen, Director Information Security Issues

[14] 10 Steps to Cyber Security – Executive Companion; BIS, GCHQ

[15] Cyber risk management: a board level responsibility; BIS, CESG, Cabinet Office

[16] Measuring the Cost of Cybercrime; Anderson, Barton, Bohne, Clayton, van Eeten, Levi, Moore, Savage, WEIS 2012

[17] APT1: Exposing One of China's Cyber Espionage Units; Mandiant Corporation

[18] Reducing Systemic Cybersecurity Risk , Sommer and Brown, OECD

[19] Penetration Testing of Industrial Control Systems, Duggan, Berg, Dillinger, Stamp, 2005