Cyber-risk and cyber-controls:

Modelling the cyber gap

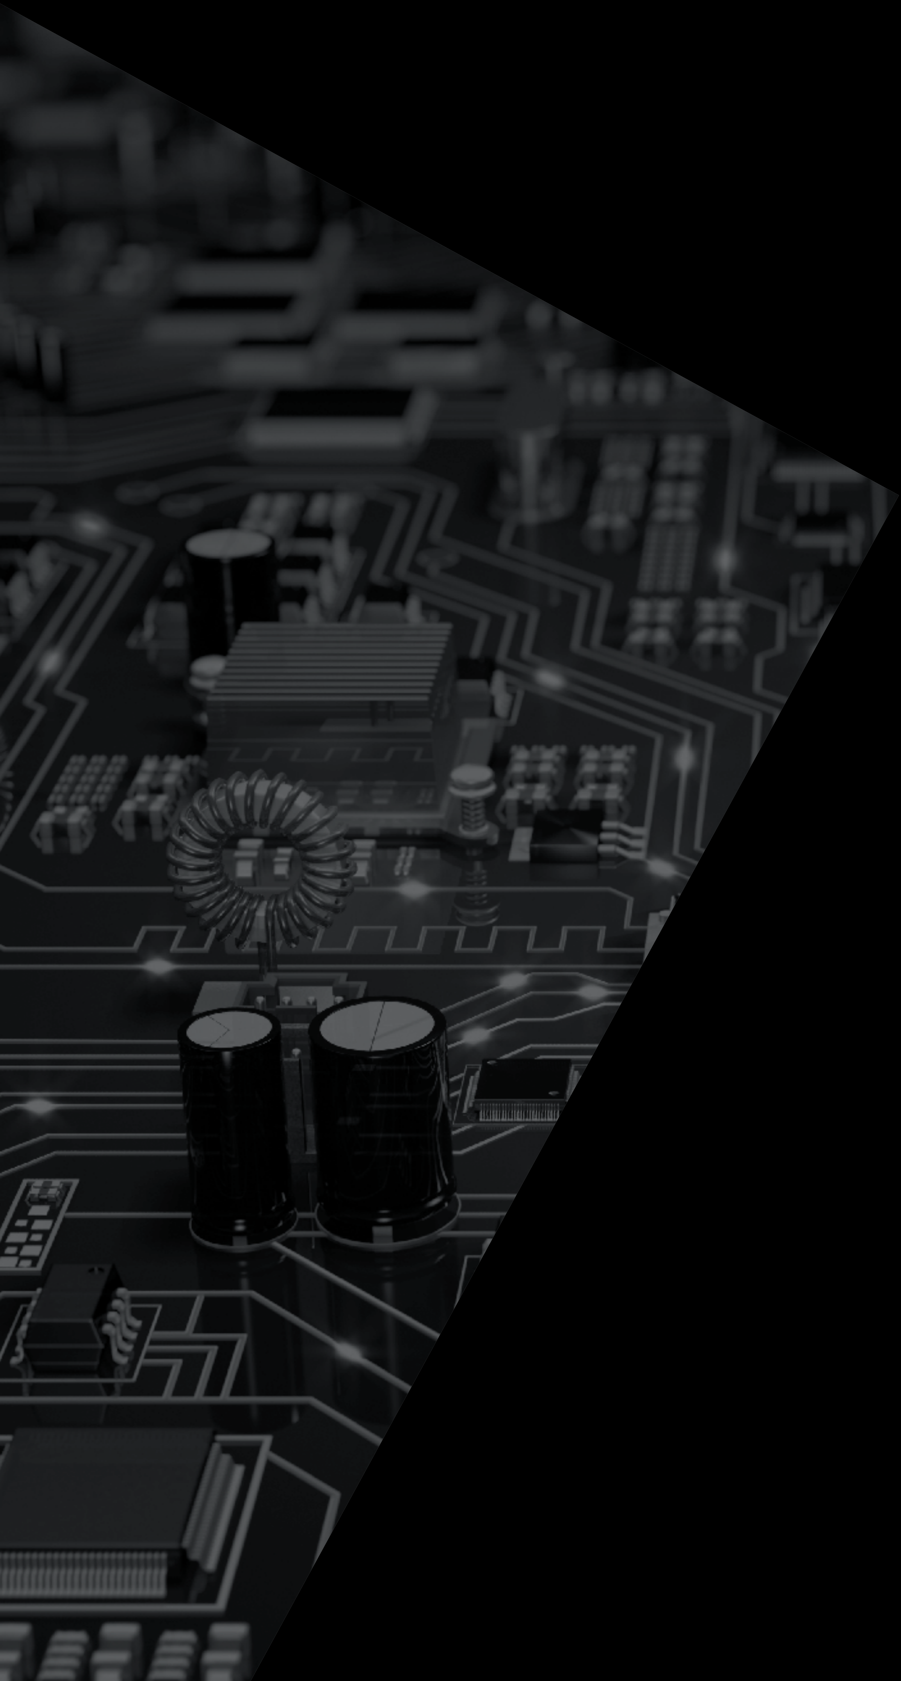

Insurance alone is not enough

Cyber-risk has become one of the most significant topics in boardrooms around the world. The threat is indeed, very real. Consequently, in a very short time cyber insurance and cyber-risk management have become important new insurance industry offerings, especially in Lloyd's.

We at Novae Group believe that insurance alone cannot entirely manage cyber-risk. Much more needs to be done to understand the risk environment and halt the potential damage to organisations that this new threat can inflict. Our approach must be holistic. The dangers posed by cyber-attackers must be answered through new collaborations to educate stakeholders, evaluate risk, test compliance, monitor the ever-changing threat environment, prepare and implement best in class crisis management plans in response to cyber-threats.

To that end, Novae Group has partnered with the University of Oxford's Department of Computer Science and Saïd Business School to conduct ongoing research. The initial product of this work probes the effectiveness of our current defences against cyber-attacks, and the standards set by international bodies which businesses use to measure the sufficiency of their cyber security efforts. This document summarises the findings of Oxford's research. You will find the complete report on Novae Group's website. If you would like to discuss the findings, please get in touch with me.

Dan Trueman
Chief Innovation Officer and Head of Cyber,
Novae Group
DTrueman@novae.com
21 February 2017

This paper summarises 'The relative effectiveness of widely used risk controls and the real value of compliance', published by the Department of Computer Science, University of Oxford, and the Saïd Business School, Oxford, sponsored by Novae Group plc.

SAID BUSINESS SCHOOL

UNIVERSITY OF OXFORD

DEPARTMENT OF COMPUTER SCIENCE

Harm arising from cyber-attacks is on the rise. Organisations must be able to show that they are aiming to reduce cyber-risk, and typically do so by working towards being compliant with one or more of many established sets of standards, such as ISO/IEC 27000. However, such standards are often not backed up by objective, empirical research, and therefore cannot be shown to have quantifiable benefits. This shortfall weakens the value of compliance to risk control standards, because a compliant organisation may not be protected from cyber-harm.

Businesses, particularly SMEs are not well prepared for data/software damage, which lead to large exposure following data incidents.

$1.7 trillion

Data loss and downtime cost enterprises $1.7 trillion around the globe in 2014[1]

4x

In the UK, data loss grew by 4 times between 2012 and 2014

78%

78% of UK organisations are still not fully confident in their ability to recover after a disruption

40%

SMEs are more vulnerable to data/software loss

- 40% of SMEs don't back up their data at all and 60% of business data is held on PC and does not get regularly backed up
- 40-50% of those backups are not fully recoverable[2]

51%

Business trends, such as mobile, big data and hybrid cloud create new challenges for data protection

- 51% of organisations lack a disaster recovery plan for emerging workloads
- Just 6% have plans for big data, hybrid cloud and mobile[3]

Sources: (1) TechWeek Europe; (2) workspace; (3) SecurityWeek

This research examines the relative effectiveness of cyber-risk controls, and the real value of compliance to cyber-risk control standards. Are cyber controls effective? Do they reduce cyber-risk, so less harm will be done to an organisation that suffers an attack? If so, by how much? This study has considered these questions through four lenses:

1. What controls do organisations typically need to protect themselves, given their assets and attack surface (the sum of the points where an attacker can try to enter or extract data from a software environment)? How should risk controls be deployed to achieve protection?

2. If the controls are the most appropriate ones, how well do they protect a given organisation, if optimally deployed?

3. What are the current threats, based on hackers' evolving abilities, and what harm do they actually cause, despite controls which meet industry best-practice?

4. What is the broader perspective of organisational cyber-harm (the set of detrimental impacts resulting from cyber-attacks), both inside and outside the organisation?

Through these lenses, this research proposes a model that reveals the effectiveness of various risk controls and responses by considering and comparing their relative benefits. It is essential that the entire exercise is conducted in a full organisational context for each application of the model. Analysis with the model allows improved management decision-making and better-designed insurance products. By relating risk controls to assets and value-at-risk (VaR), the model has a double benefit. It sheds light on the potential harm that could be inflicted by successful cyber-attacks, and it can be used to assess the preventative benefits of compliance to security standards and frameworks.

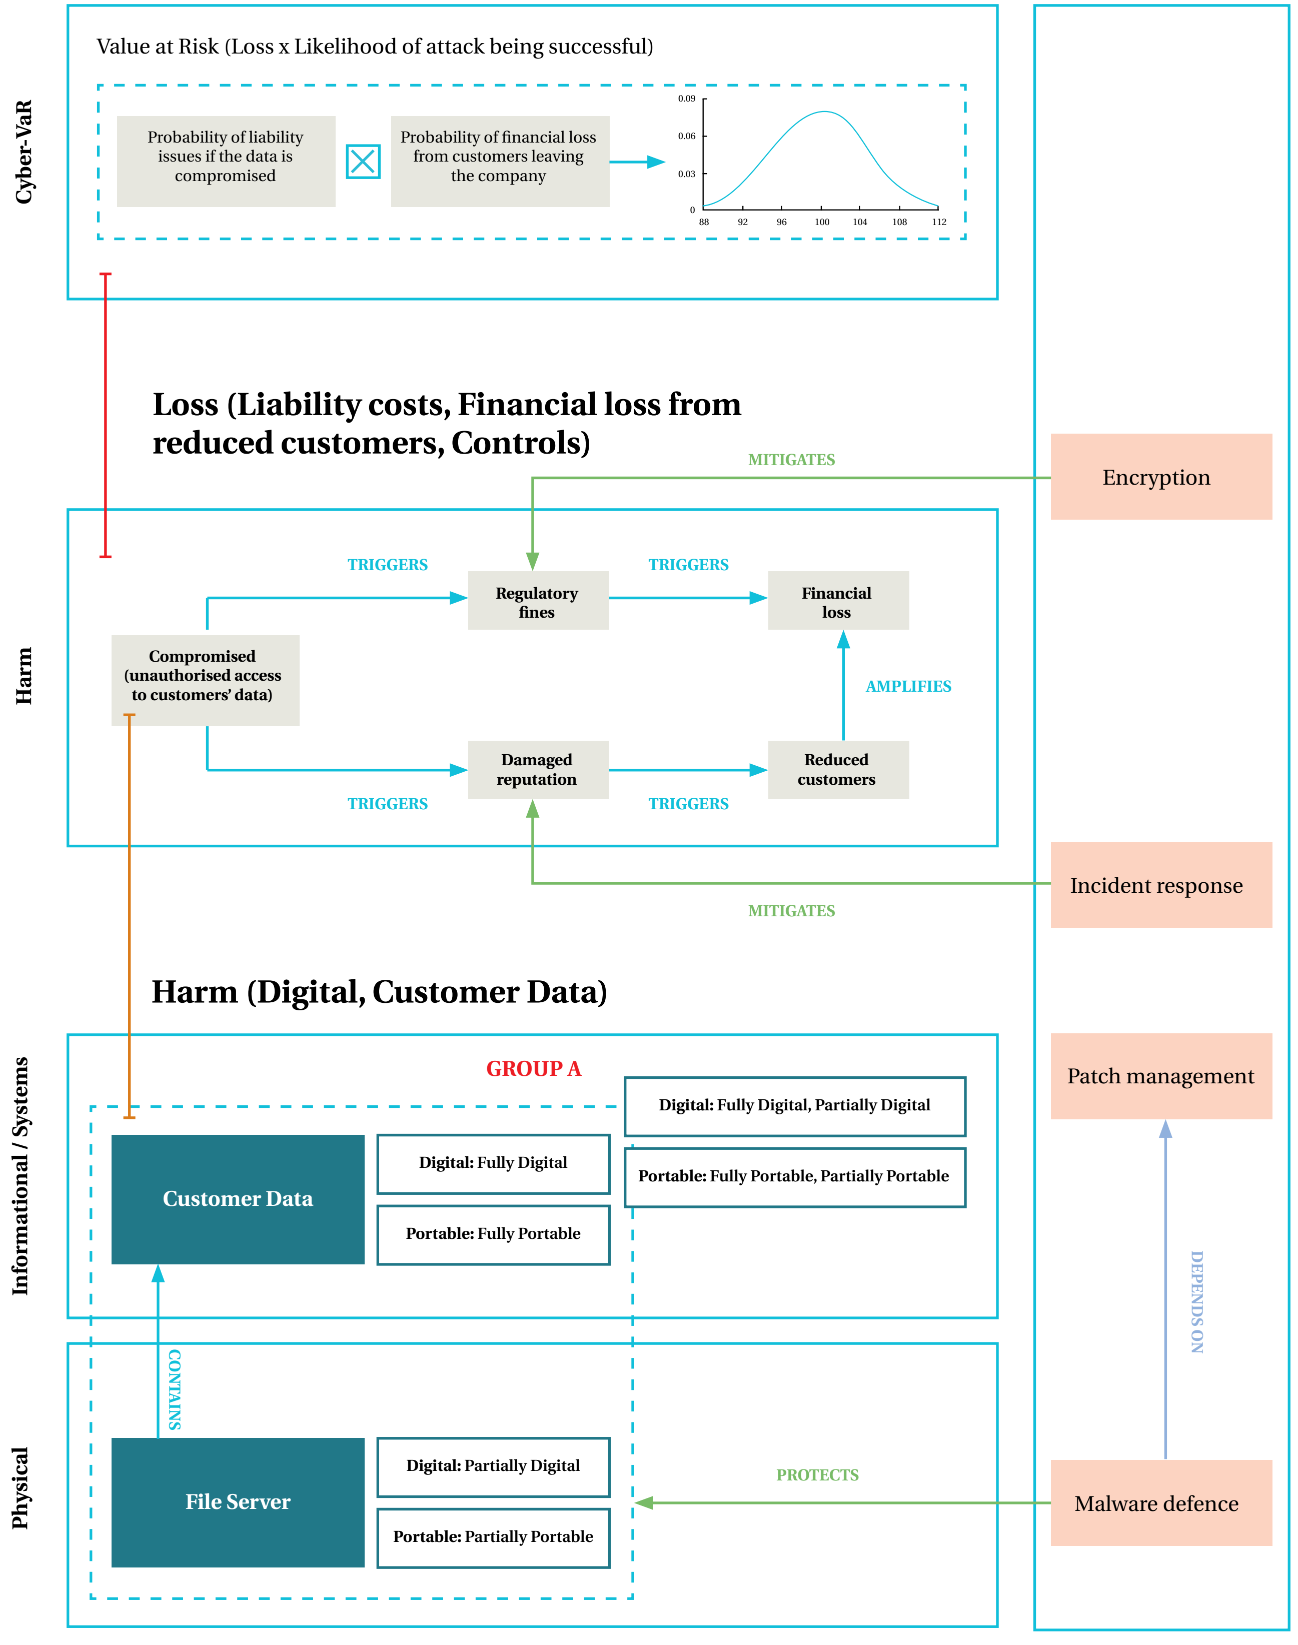
Cyber-VaR
Value at Risk (Loss x Likelihood of attack being successful)
Probability of liability issues if the data is compromised
☒
Probability of financial loss from customers leaving the company
0.09
0.06
0.03
0
88
92
96
100
104
108
112
Harm
Loss (Liability costs, Financial loss from reduced customers, Controls)
MITIGATES
Encryption
TRIGGERS
Regulatory fines
TRIGGERS
Financial loss
Compromised (unauthorised access to customers' data)
AMPLIFIES
Damaged reputation
Reduced customers
TRIGGERS
TRIGGERS
Incident response
MITIGATES
Informational / Systems
Harm (Digital, Customer Data)
GROUP A
Digital: Fully Digital, Partially Digital
Portable: Fully Portable, Partially Portable
Patch management
Customer Data
Digital: Fully Digital
Portable: Fully Portable
DEPENDS ON
CONTAINS
Physical
File Server
Digital: Partially Digital
Portable: Partially Portable
PROTECTS
Malware defence
Controls

Source: University of Oxford Department of Computer Science

Assets and attack surfaces

An asset is ‘anything of value to the organisation’. It can be **physical** (such as electronic equipment, buildings, cash, chairs), **information or systems** (data, financial information, software), **people** (employees, suppliers, customers), **routines** (like a sales programme or an R&D programme), or **enterprise** (including reputation, culture, patents, profitability). Importantly, the research shows that assets can be located on ‘dimension’ scales. These include:

- Connected to Isolated
- Digital to Analogue/Non-Electronic
- Human to Non-Human
- Intelligent to Unintelligent
- Portable to Fixed
- Novel to Established
- Persistent to Transient
- Physical to Non-Physical
- System to Component
- Tacit to Explicit

For use in the model, assets must be rated along these and other dimensions, for example on a scale of one to three. The dimensions are important to consider because they can provide some insight into how likely an asset is to be successfully attacked. While most organisations have a good understanding of the value of most of their assets, the VaR and specifically the cyber-VaR are very poorly understood. This makes proper assessment of the worth of investments in risk mitigation, and thus risk transfer mechanisms such as insurance, very difficult. Understanding of the risk transfer solution is weak as a result, creating the urgent need for a process that provides clear organisational understanding of cyber-VaR and the effectiveness of risk controls.

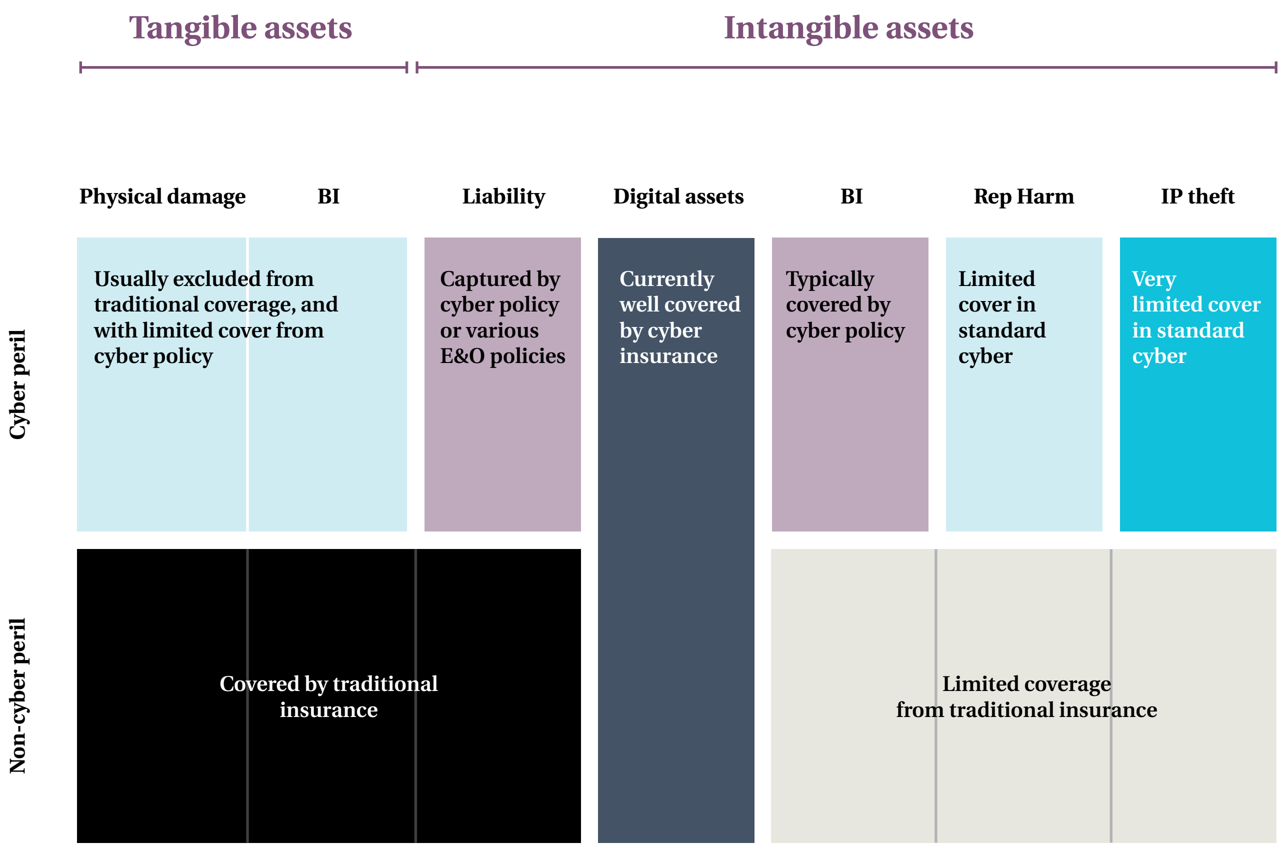
Tangible assets

Intangible assets

Physical damage | BI | Liability | Digital assets | BI | Rep Harm | IP theft

Cyber peril

Usually excluded from traditional coverage, and with limited cover from cyber policy

Captured by cyber policy or various E&O policies

Currently well covered by cyber insurance

Typically covered by cyber policy

Limited cover in standard cyber

Very limited cover in standard cyber

Non-cyber peril

Covered by traditional insurance

Limited coverage from traditional insurance

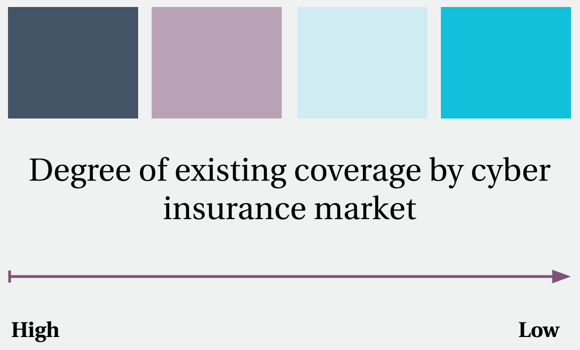
Degree of existing coverage by cyber insurance market

- Current cyber insurance focuses on losses related to digital assets caused by cyber perils, while traditional insurance mainly captures risks with non-cyber perils in tangible assets. **Some large loss areas such as BI, IP theft are still left under-insured.**
- With the proportion of tangible assets reducing and cyber perils becoming more prevalent, the existing **traditional insurance segment is likely to need to change.**
- Cyber insurers are **likely to innovate and expand coverage into a wide scope of intangible assets following several waves of development.**

Source: Novae Group

Controls and protection

A control is a security mechanism put in place to reduce an asset's attack surface and protect it from harm. Its purpose is either to reduce the likelihood of a successful attack by completely avoiding a risk, or diminish it by reducing or removing the attack surface (including by making it more difficult to attack), or both. Controls typically comprise multiple sub-controls. Such measures protect the VaR contained in assets, and therefore prevent harm.

Risk mitigation through the implementation of controls is essential. Controls include impact mitigation, which is the ultimate goal of controls. Adoption of the numerous security controls available is supported by control standards, each of which provides a unique approach, and guidance on usage. These controls are derived from their creators' understanding of the threat environment, and current technologies used by organisations.

Each control seeks to protect an asset or a set of interconnected assets of a particular class (physical, information, etc). Often multiple controls protect single assets; often they are layered and interdependent. Some have inherent design vulnerabilities, others are poorly implemented. Little empirical data on their effectiveness has been compiled, so a gap exists in our knowledge of exactly which controls are most effective at protecting against attacks. This means decisions regarding the implementation of controls are not based on facts about the true performance of controls. This problem is exacerbated by the interdependencies of controls.

Interdependence of controls

Cyber-risk controls must work together to ensure that they do not conflict with each other, and allow for layers of security and defence-in-depth. The research found that sub-controls – the components of individual controls – sometimes provide essential small steps towards the bigger aim of the implementation of a control, and therefore imply an implementation sequence. Less often, sub-controls are independent of each other.

Controls themselves sometimes rely on other controls. Some are high-priority controls, which many others rely upon. Therefore, the extent to which an organisation is effectively protected from risks is dependent on the system of controls. Inadequacies in basic controls (such as maintaining an inventory of authorised devices) could impact an organisation's ability to implement more complex controls properly. An ineffective control may leave a higher residual risk, which dependent controls also carry. Propagation of risk in this way can be dramatic.

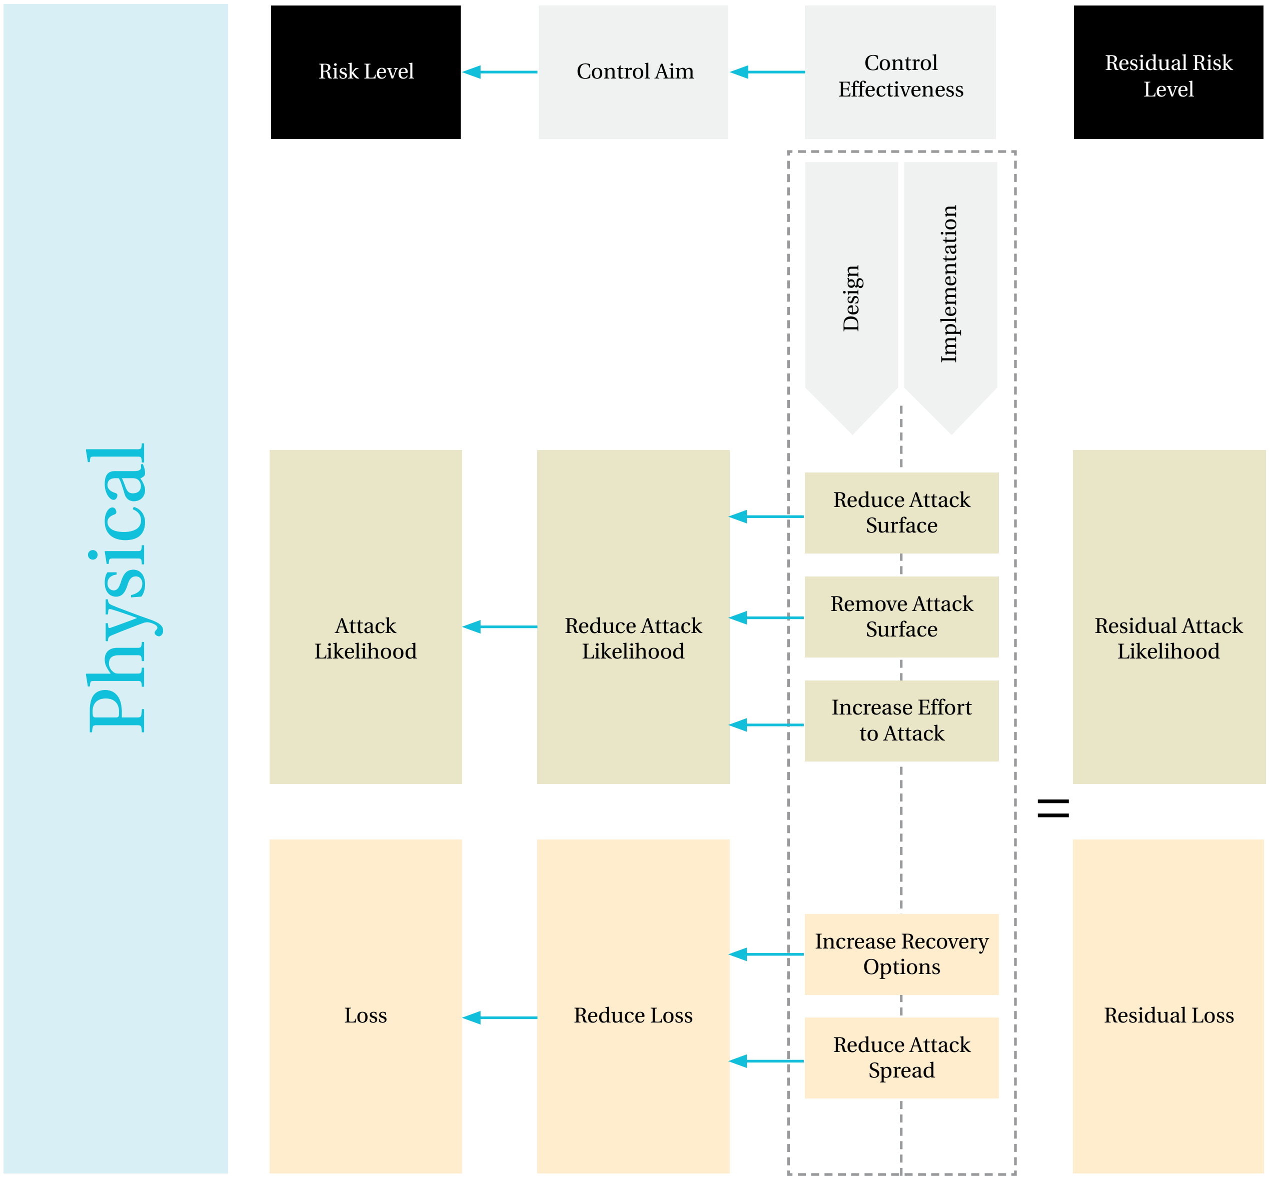

Source: University of Oxford Department of Computer Science

Attackers

The threat actor – the hacker – selects and controls the attack. They have intent, for example to steal, sabotage, or simply gain access and persist. The effectiveness of a risk control is variable in relation to the attacker faced.

Some risk controls are intended to prevent an attack by removing the attack surface, although it is not always easy or even possible to know that all attack surface has been removed. Many organisations rely on testing to provide confidence that it has, but the question of knowing whether testing has been sufficient remains unsolved. Risk controls, like firewalls, seek to remove the likelihood that a threat can reach an exploitable attack surface. However, it must catch all possible attacks that are aimed at an exploitable attack surface, and it is a great challenge to anticipate all types of attack. A practical solution might be to develop a measure of likelihood of an attack's success. However, limited data exists to underpin the assignment of probability. It can only be estimated to present a range of possible outcomes based on a scenario.

Alternately, risk controls may seek to detect and limit attacks. Such measures are essential to organisations that face a large or frequent threat, and a substantial risk. The effectiveness of controls can be measured by their ability to detect threats quickly enough to allow time for a response which can limit the harm. Such controls vary in effectiveness given the nature of the threat faced. It may be possible to measure and then predict more accurately the degree of harm exposure for a given threat capability, independent of the attacker's intent.

Unfortunately attackers pose a creative threat capable of reason and innovation, which can seek to adapt to take account of risk controls. Creative attacks employing distraction have often occurred, for example when multiple Denial of Service attacks across a system mask a more aggressive hack. The creativity of attackers means the predictability of risk controls may render them much less effective.

Modelling cyber-risk control effectiveness

Our model hypothesises about the relationships between risk controls on the one hand, and assets, cyber-VaR, and cyber-harm on the other. It allows analysis of areas where value and harm are unaddressed by current controls. It is based on three analytical requirements:

1. Identify and predict where value and harm are unaddressed by controls and responses. A key aim is to identify where controls do not mitigate harm and protect VaR.

2. Elucidate and refine understanding of residual risk within systems after deployment of controls.

3. Identify data (on specific asset types at risk, and the effectiveness and inter-reliability of controls) urgently needed to quantify and refine understanding of the real risk from cyber-attacks, and the impact of adopting certain risk controls or responses.

The model has three levels: the asset level, the harm level, and the cyber-VaR level. Based on the data input, it reasons within and between the model levels. Finally, the effectiveness of controls is tested against the model's findings, based on an analysis of the three levels. Comprehensive details of the model can be found in the full report.

Relative effectiveness of risk controls and the value of compliance

Control interdependency is considered by some security experts, but the concept remains in its infancy. Further, control selection is often not driven by effectiveness, but by regulation, legislation, or threat trends. This highlights a potential disconnect in the controls selected by companies. The model is the first step in determining the effectiveness of controls, but more research is necessary.

The initial model in detail

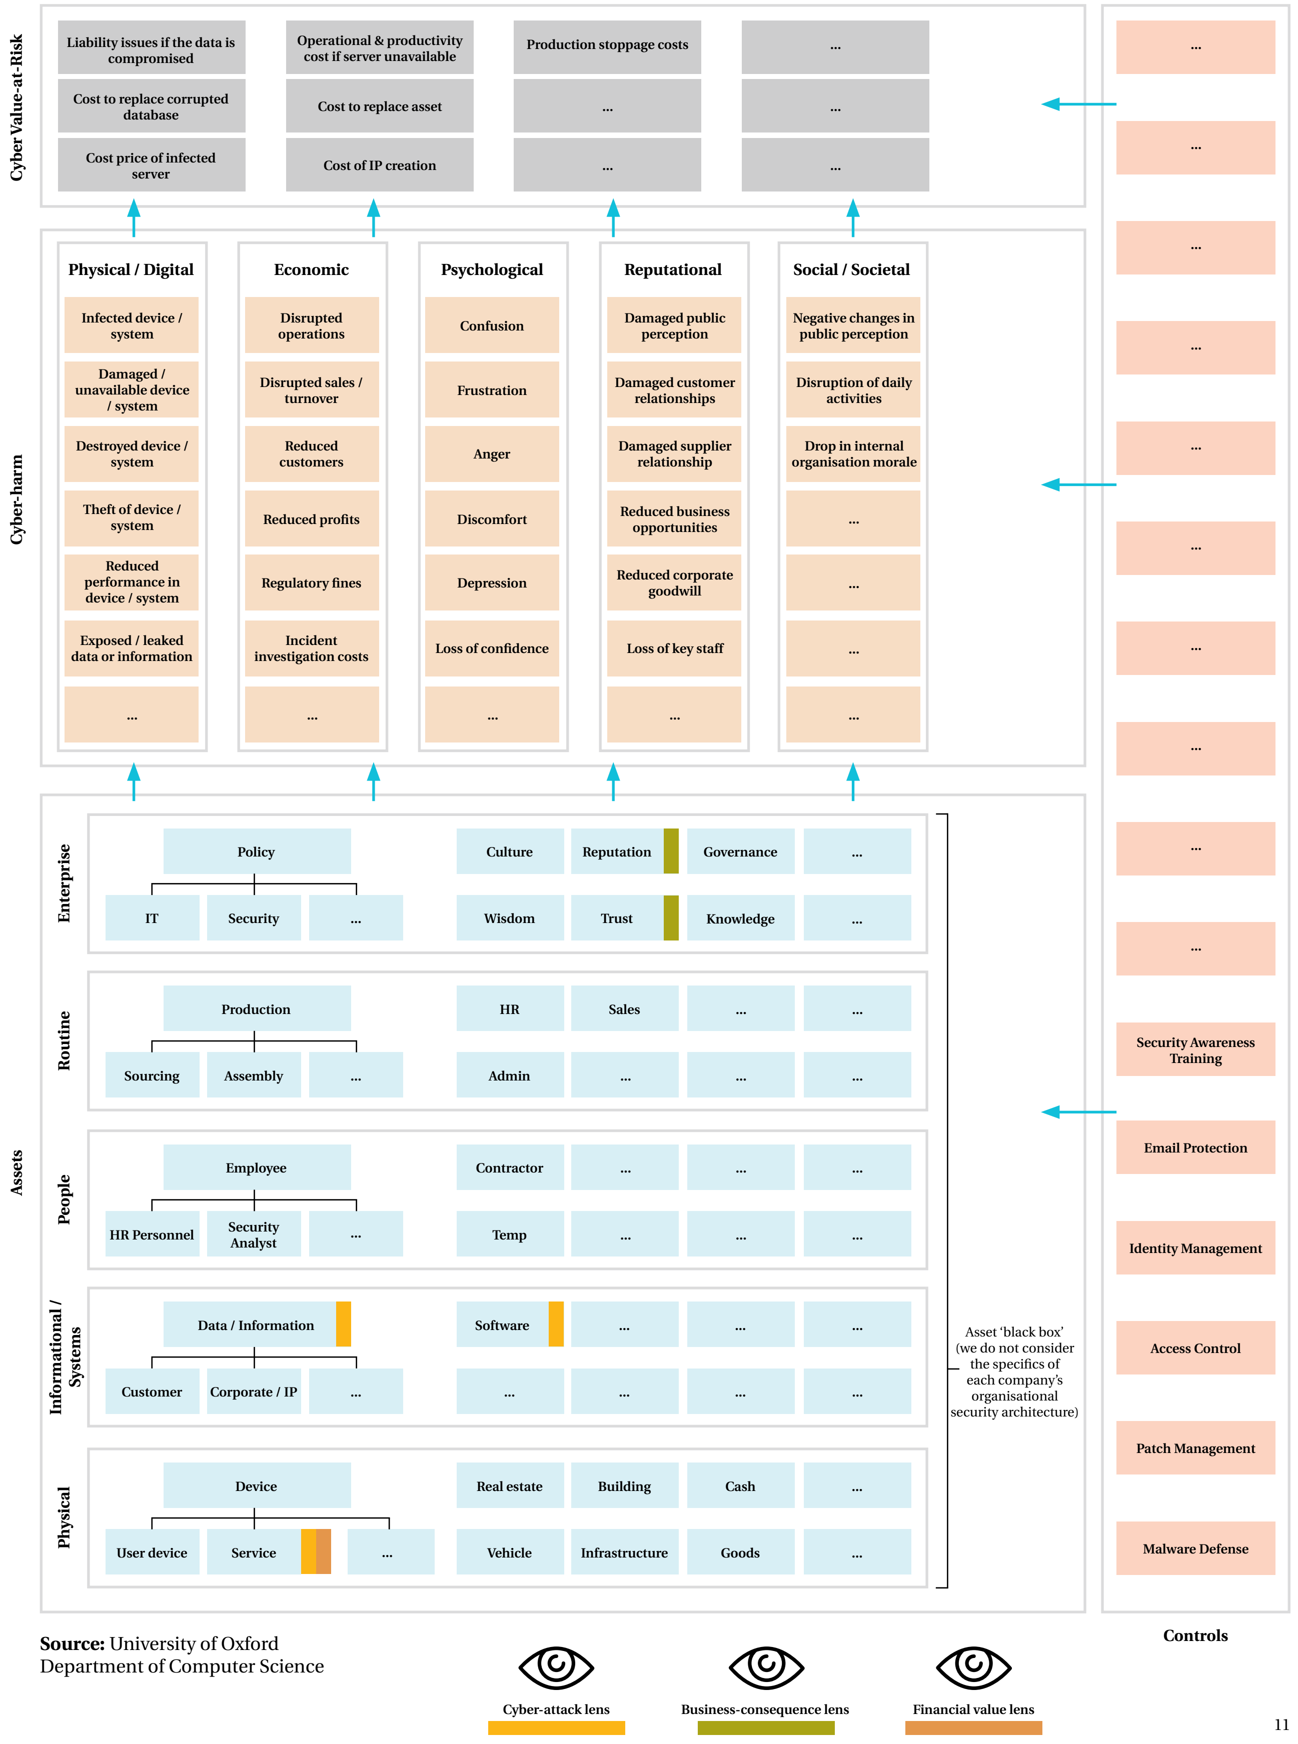

Cyber Value-at-Risk

Liability issues if the data is compromised	Operational & productivity cost if server unavailable	Production stoppage costs	...
Cost to replace corrupted database	Cost to replace asset	...	...
Cost price of infected server	Cost of IP creation	...	...

Cyber-harm

Physical / Digital	Economic	Psychological	Reputational	Social / Societal
Infected device / system	Disrupted operations	Confusion	Damaged public perception	Negative changes in public perception
Damaged / unavailable device / system	Disrupted sales / turnover	Frustration	Damaged customer relationships	Disruption of daily activities
Destroyed device / system	Reduced customers	Anger	Damaged supplier relationship	Drop in internal organisation morale
Theft of device / system	Reduced profits	Discomfort	Reduced business opportunities	...
Reduced performance in device / system	Regulatory fines	Depression	Reduced corporate goodwill	...
Exposed / leaked data or information	Incident investigation costs	Loss of confidence	Loss of key staff	...
...	...	...	...	...

Assets

Enterprise: Policy (IT, Security, ...); Culture, Reputation, Governance, ...; Wisdom, Trust, Knowledge, ...

Routine: Production (Sourcing, Assembly, ...); HR, Sales, ..., ...; Admin, ..., ..., ...

People: Employee (HR Personnel, Security Analyst, ...); Contractor, ..., ..., ...; Temp, ..., ..., ...

Informational / Systems: Data / Information (Customer, Corporate / IP, ...); Software, ..., ..., ...; ..., ..., ..., ...

Physical: Device (User device, Service, ...); Real estate, Building, Cash, ...; Vehicle, Infrastructure, Goods, ...

Asset 'black box' (we do not consider the specifics of each company's organisational security architecture)

Controls: ... ; ... ; ... ; ... ; ... ; ... ; ... ; ... ; ... ; ... ; Security Awareness Training; Email Protection; Identity Management; Access Control; Patch Management; Malware Defense

Cyber-attack lens — Business-consequence lens — Financial value lens

Source: University of Oxford Department of Computer Science

London Office
21 Lombard Street,
London EC3V 9AH
Tel: +44 (0) 20 7050 9000

Bermuda Office
Ideation House,
94 Pitts Bay Road,
Pembroke, HM 08
Bermuda
Tel: +1 441 278 86 00

enquiries@novae.com | www.novae.com