A new perspective on cyber risk, applied to the evolving UK energy grid ecosystem

Applying the benefit harm index (BHI), a new approach to modelling risk assessment of cyber ecosystems and their socio-economic impacts to the UK's evolving connected and autonomous vehicle ecosystem.

CHARLES FOX — SECURITY LEAD, DIGITAL CATAPULT
BRIAN MACAULAY — LEAD ECONOMIST, DIGITAL CATAPULT

Executive summary

Within an increasingly complex and interconnected world, the way in which cyber-threats are perceived and responded to needs to be reconsidered. Traditional risk models rely heavily on probabilistic approaches, which demand stable distribution and almost complete knowledge of all possible states.

New advances in digital technologies, combining huge data, rapidly evolving automated algorithms and the prospect of a generational shift in network speed and capacity, pose serious challenges to traditional risk modelling. Through the Hermeneut project (part of the European Community's Horizon 2020 programme) Digital Catapult has proposed a new approach to understanding dynamic and emergent threats: the benefit harm index (BHI), which integrates ideas from both economics and complexity science.

This report shows how this exciting new perspective on cyber risk modelling can be applied to the cyber ecosystems that form many of today's critical national infrastructures (CNI) - complex systems of systems that exhibit emergent behaviour and require a new approach to cyber risk assessment. This study looks at the systemic socio-economic impacts that can result from cyber attacks associated with emergent threats to CNI cyber ecosystems, and uses the UK energy smart grid ecosystem as a case study for the new BHI approach.

The UK energy smart grid ecosystem is part of the UK economy and is one of the UK's 13 CNI components. To a large extent, the energy grid is essential to the operation of the UK's entire socio-economic system, and therefore, any prolonged nationwide power outage would have a systemic impact on the UK economy.

A high-level ecosystem for 2020-30 has been modelled to focus on the energy grid's critical operational systems domain, and on the associated domains of UK governance, the supply chain and wider non-critical core services. This model provides the context for applying the BHI approach to an illustrative multi-vector cyber attack that would have a systemic impact on the UK energy sector.

This report also describes the approaches that can be used to mitigate the growth of harm within these complex systems of systems, and highlights the use of Implication Wheel™ methodology to uncover emergent systemic threats to the UK energy grid cyber ecosystem.

In this report

INTRODUCING THE BHI – A NEW PERSPECTIVE TO CYBER RISK

BHI modelling methodology is designed to provide new insights into the potential risks associated with the cyber ecosystems which underpin complex and dynamic markets that are driven by the exploitation of emerging technologies. These rapidly evolving markets typically contribute significantly to national and international economies, and often form an integral part of CNI.

Unlike a controlled (deterministic) system with a known set of risks and a well-defined future state, a complex system features many unknown risks and will evolve in ways that cannot be fully predetermined. For example, within the biological ecosystem, microscopic changes can propagate rapidly and create a huge-scale impact, such as when a single virus mutates, evolves and spreads to cause a pandemic, demonstrates Cyber ecosystems are also complex dynamic environments that evolve rapidly and feature high levels of uncertainty. They can generate emergent behaviours which cannot always be predicted by studying the way in which constituent parts interact. Emergent behaviours manifest themselves in many forms (as seen in the murmurations of birds in the biosphere, and new socio-political collective behaviours through social media use online).

Traditional risk assessment methodologies - which assume a complete knowledge of all possible states of the system being assessed and that a mathematical likelihood can be applied to each event - cannot address the complex dynamics, emergency behaviours and associated uncertainties of cyber ecosystems.

The Hermeneut BHI introduces a new approach to risk assessment, by modelling the growth of benefits and risks in the context of complex cyber ecosystems. It also features event-driven scenario analysis methods, recognising the evolution of such systems over time.

Modelling dynamic complexity provides a perspective for exploring the rate of growth of socio-economic benefits generated by an evolving cyber ecosystem over time. It also provides a perspective for exploring the rate of growth of threats to that ecosystem, and the associated socio-economic harm that those threats could generate over time. The difference between the level of benefit and the level of harm at any given time period is a key output of the BHI model.

An event-driven scenario approach enables exploration of the implications of cyber chain reactions, helping to identify hidden risks (and benefits) using tools such as the Implication Wheel™[1]. This helps mitigate the fact that the risks for complex dynamic systems cannot fully be predicted as some will be emergent, and could be significant.

The BHI methodology applies many of the principles used in the latest economics research[2], recognising that the economy is a complex system within other systems. When the BHI methodology is applied to a cyber ecosystem, the balance between benefit and harm, and how that balance changes over time, can be explored. BHI is used to identify and mitigate emergent threats, and then to explore ecosystem-level mitigation strategies for those scenarios where the socio-economic harm outweighs the benefits. Any residual risks can then be managed using traditional risk assessment methodologies.

USING BHI TO MITIGATE TO EMERGENT THREATS

Cyber ecosystems are complex, and therefore exhibit emergent behaviour. As the level of complexity increases, different types of emergent behaviour will appear:

- Simple dynamic behaviour (such as a clock measuring time)
- Weak emergent behaviour (such as the flocking of birds or shoaling of fish)
- Strong emergent behaviour (such as bubbles within financial markets)
- Spooky emergent behaviour (such as conscious thought in humans or AI)

The first two emergent behaviour types are associated with deterministic systems, and can be easily reproduced using system simulations. The third and fourth are associated with stochastic (random interactions defined by probability distribution) systems. Stochastic systems can exhibit strong emergent behaviour that cannot be fully reproduced in simulations; spooky emergent behaviour cannot be reproduced by even the most detailed simulation.

The extent to which a cyber ecosystem can be controlled - and defended - is intrinsically linked to its level of complexity. The stability of the system is also related to its level of complexity, and changes at micro level can result in dramatic change at macro level. Therefore, an attack on a cyber ecosystem can trigger a significant chain reaction that will appear as emergent behaviour.

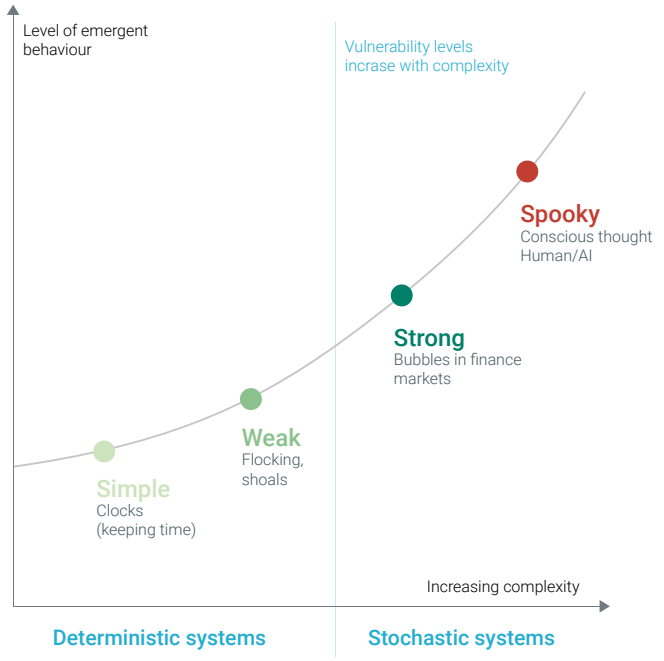

Figure 1 – Complexity and emergent behaviour

In the case of strong and spooky emergence, the stochastic systems) the system is fundamentally uncontrollable.

BHI methodology proposes a taxonomy for the vulnerability level (VL) of a system. This defines the states of a system in terms of a given scope and phase space (representing all possible states of the system) with a given resolution. BHI uses this as a measure of a system's intrinsic lack of controllability, from the perspective of those who are defending it (those who legitimately operate the system).

Table 1 shows how threats and vulnerabilities to components in a system will vary by class. Each VL requires a different type of mitigation.

The VL of a component may be changed by reconfiguring other components in the system. Some levels of vulnerability must be mitigated across the ecosystem.

Vulnerability level (VL)	Threat class	Attacker's control
5	Emergent system	The system can show emergent behaviour and cannot be controlled, since its phase space changes as emergent behaviours manifest themselves.
4	Stochastic system	The system cannot be controlled, but vulnerabilities can be reliably modelled using closed-form probability distributions over a fixed (and finite) set of state variables in the system's phase space.
3	Uncontrolled system	The system is not under control, but could be controlled in principle.
2	Uncontrolled inputs	An attacker uses a legitimate control input within the system's scope, but outside its expected or normal range.
1	Unauthorised activities	An attacker uses legitimate and in-scope control inputs within the control system.

Table 1 – Vulnerability levels and their associated class of threat

One of the key components of the BHI approach to dynamic risk involves mitigating emergent threats within complex ecosystems. Figure 2 illustrates the process for doing this.

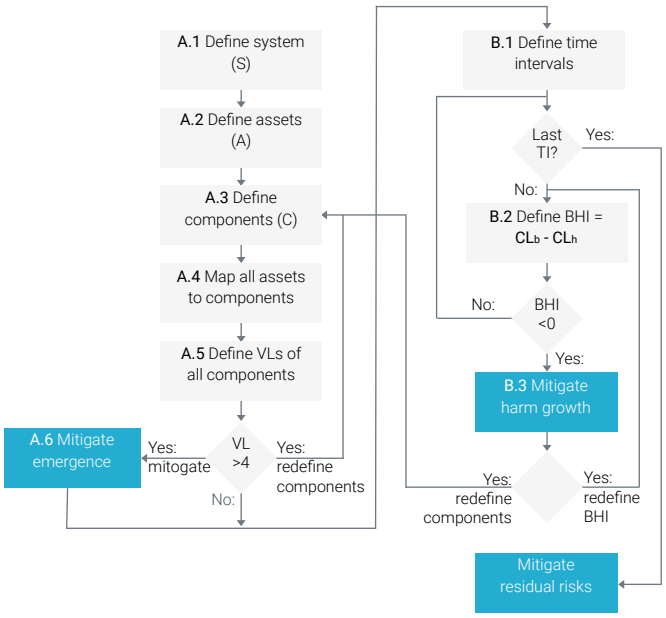

Figure 2 - BHI process for mitigating emergent threats

As shown in Figure 2, the first steps for addressing emergent threats (A.1 to A.5) are to define:

A.1: The ecosystem being considered.

A.2: The set of assets (the sensitivity of which is such that their loss or compromise would cause significant harm, and which - as a whole or in part - may be of interest to a threat agent for malicious, fraudulent and criminal activities).

A.3: The set of components which comprise the system - a component must contain hardware and may contain software and data (it is assumed that components can communicate with each other using sufficiently secure protocols).

A.4: The association between each asset and any component that directly influences its security.

A.5: The VL for each component.

These definitions should take into account the nature of each component and its vulnerabilities, as well as the threats from the environment and other components. If any component has VL greater than four (corresponding with emergent threat), the process takes one of two paths:

- Redefinition of the components, for example, to localise an associated asset in a component that has a lower VL value - this requires in reiteration over steps A.3 to A.5
- Mitigate emergence (A.6) by designing a set of security controls that seek to mitigate associated risks - these controls need to detect, and potentially isolate and neutralise the impact of an attack

Using BHI, characteristics that can be localised need to be distinguished from those which cannot. Organisations cannot be expected to mitigate non-local characteristics, so other classes of intervention must be applied to safeguard the ecosystem. For the latter class, mitigations must be a set of governance, standard, and other interventions across the ecosystems, and key criteria for adoption must seek to minimise impact on the individual organisations adopting such recommendations.

Once this iterative process is complete, the process of considering emergent threat is also complete, and analysis passes to using BHI to mitigate threats from growth.

USING BHI TO MITIGATE TO THE GROWTH OF HARM

Modelling dynamic complexity provides a perspective for exploring the rate of growth of the socio-economic benefits that an evolving cyber ecosystem generates over time. It also provides a perspective for exploring the rate of growth of threats to that ecosystem, and any associated socio-economic harm could be generated as a result. The difference between the level of benefit and the level of harm at any given time is a key output of the BHI model.

Benefit and harm can grow at different rates within a cyber ecosystem. There are two key features of complex ecosystems that help to refine understanding of these growth rates.

1. Each ecosystem will evolve through a number of distinct phase transitions as it evolves

For example, the introduction of a new product or class of products that penetrates a market. Initially there is near exponential growth, often modelled as compound growth in business plans, with a constant or slowly varying compound annual growth (CAGR) parameter. As penetration of the market occurs and saturation approaches, the Bass diffusion distribution eventually manifests its asymptotic growth complexity at constant of zero.

It is therefore appropriate to consider the BHI at three distinct time intervals:

TI0: From product introduction to when the complexity level is four (exponential)
TI1: From when the complexity level transitions from four to zero
TI3: From market saturation onwards, when the complexity level is zero (constant)

2. Each ecosystem will typically have multiple domains, each of which can feature different levels of complexity and associated growth rates

The right-hand side of Figure 2 shows the process for using BHI to mitigate threats from growth.

The first step (B.1) defines the set of time intervals, that are relevant to the various developments of both the benefit and harm over time.

In particular, these time intervals will consider for example:

- Times of events marking the start and end of relevant changes, such as investment rounds or the introduction of new products

- Times at which the distribution of growth is likely to be discontinuous, for example as a result of some material event such as a change in product or the channel it uses to access the market

The second step (B.2) iterates over the intervals to compute the benefit to harm index (BHI) for each sub-interval, by determining the complexity index (CI) for each growth distribution. If the BHI is negative, indicating that the CI for growth of harm exceeds that of benefit, the process proceeds to mitigate harm growth (B.3), which specifies security controls. If a plausible mitigation is found, the process re-computes the BHI value and iterates to the next time interval.

In some cases, for example where an effective mitigation cannot be found, it may be appropriate to redefine the components. In this case, the process returns to the right-hand side of the diagram at step (A.3).

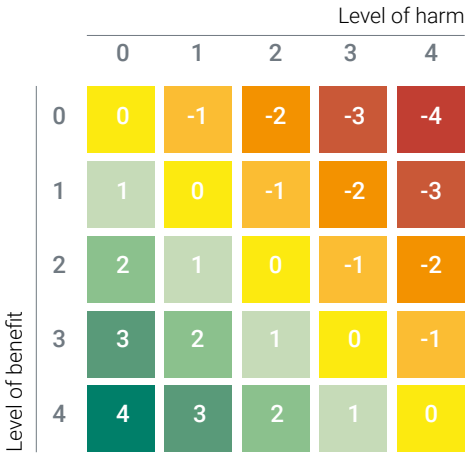

Figure 3 – The BHI for distinct time intervals (TI)

For any BHI greater than 0 systemic (ecosystem) level mitigations are required.

Once all members of CI have been processed, mitigation of risks from growth is complete and the process can continue by using traditional risk management techniques to address any residual risks.

APPLYING BHI TO CYBER ECOSYSTEMS

To apply the BHI methodology to a target cyber ecosystem, the following high-level ecosystem domain model is used.

A cyber ecosystem is a complex system of systems, where each system can be modelled in terms of a set of interacting components. Each ecosystem will have a scope/system boundary and will typically be embedded in a wider environment. Political, economic, social, technological, environmental and legal (PESTL) influences from this wider environment affect the ecosystem's operation and growth.

Each cyber ecosystem is structured into a number of domains that support different dynamic communities of interest (COI). As shown in Figure 4, these domains reflect the distinction between operational systems within the ecosystem and the supply chain systems that support the manufacture and production of the components that will eventually populate that operational system's domain.

The other domains shown include the command and control systems domain, and the underlying system components, processes and interactions that comprise them. The governance and regulatory processes domain contains the governance systems and regulatory frameworks used to set and police the policies, rules and standards associated with governing the cyber ecosystem. The final domain is the value-added services domain, which includes the systems and processes associated with services that add value to the operational services, for example, insurance services.

All cyber system domains will have vulnerabilities. Threats to the ecosystem will exploit these vulnerabilities through attack vectors originating from threat sources (for example, hostile states), and attacking via threat actors (external and internal), as shown schematically in Figure 4. Through multiple iterations, the BHI approach exploits methodologies such as the Implementation Wheel™ to investigate the vulnerability levels of components and cyber chain reactions being generated in complex systems. Targeted scenario analysis is used to help identify such events by systematically exploring the implications of interaction/contagion through multiple first, second, and nth order interaction flows.

The BHI dynamic approach to risks also enables the construction of multiple phase states of each cyber ecosystem model to reflect its different evolutionary states. This is then used to help create the BHI growth model across those different time intervals, resulting in an output of the form shown earlier in Figure 3.

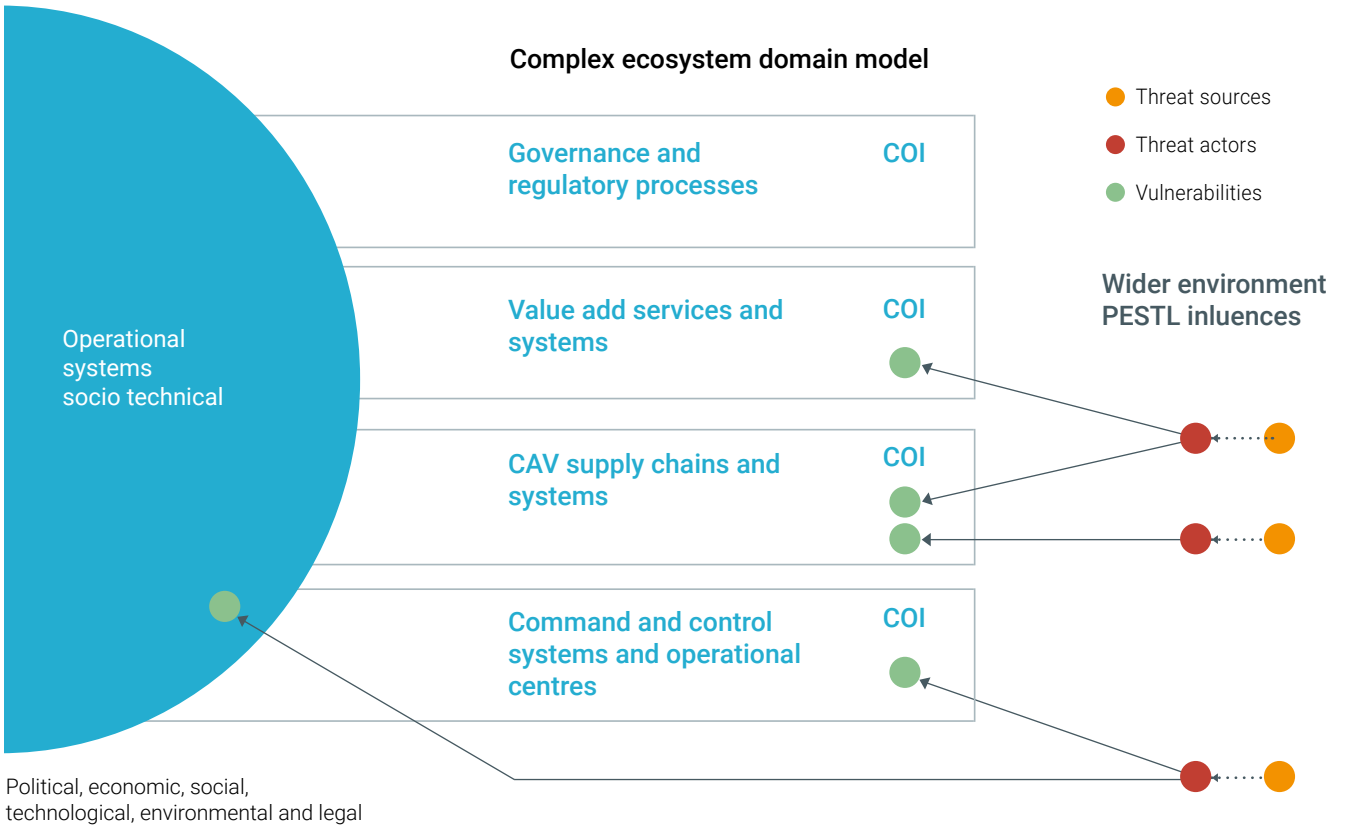

Figure 4 - Cyber ecosystem high level domain model

HIGH-LEVEL MODEL OF THE UK ENERGY GRID ECOSYSTEM

The UK's energy grid is one of the UK's thirteen CNI components. This paper focuses on the critical operational systems that underpin the resilience of the UK energy grid.

Using the ecosystem domain model, the UK energy grid ecosystem can be represented at a conceptual level, as shown in Figure 5.

Each of the domains shown in Figure 5 represents a distinct dynamic socio-technical community of interest (COI) within the UK energy grid ecosystem.

The central core in Figure 5 represents the evolving energy smart grid infrastructure, from power generation,through transmission and distribution, to consumption. This includes both gas and electricity distribution networks.

The evolution of the UK energy grid from analogue to smart is being driven by the increasing adoption of diverse energy sources, including wind and solar, together with a fundamental shift from being a monopoly of grid operators and utilities generating power to a system where prosumers play a key role. At the technological heart of this digital transformation from operational technology (OT) to IT is the internet of things (IoT). The UK energy ecosystem is embedded in the wider global energy ecosystem and is subject to global political, economic, social, technical and legal (PESTL) influences.

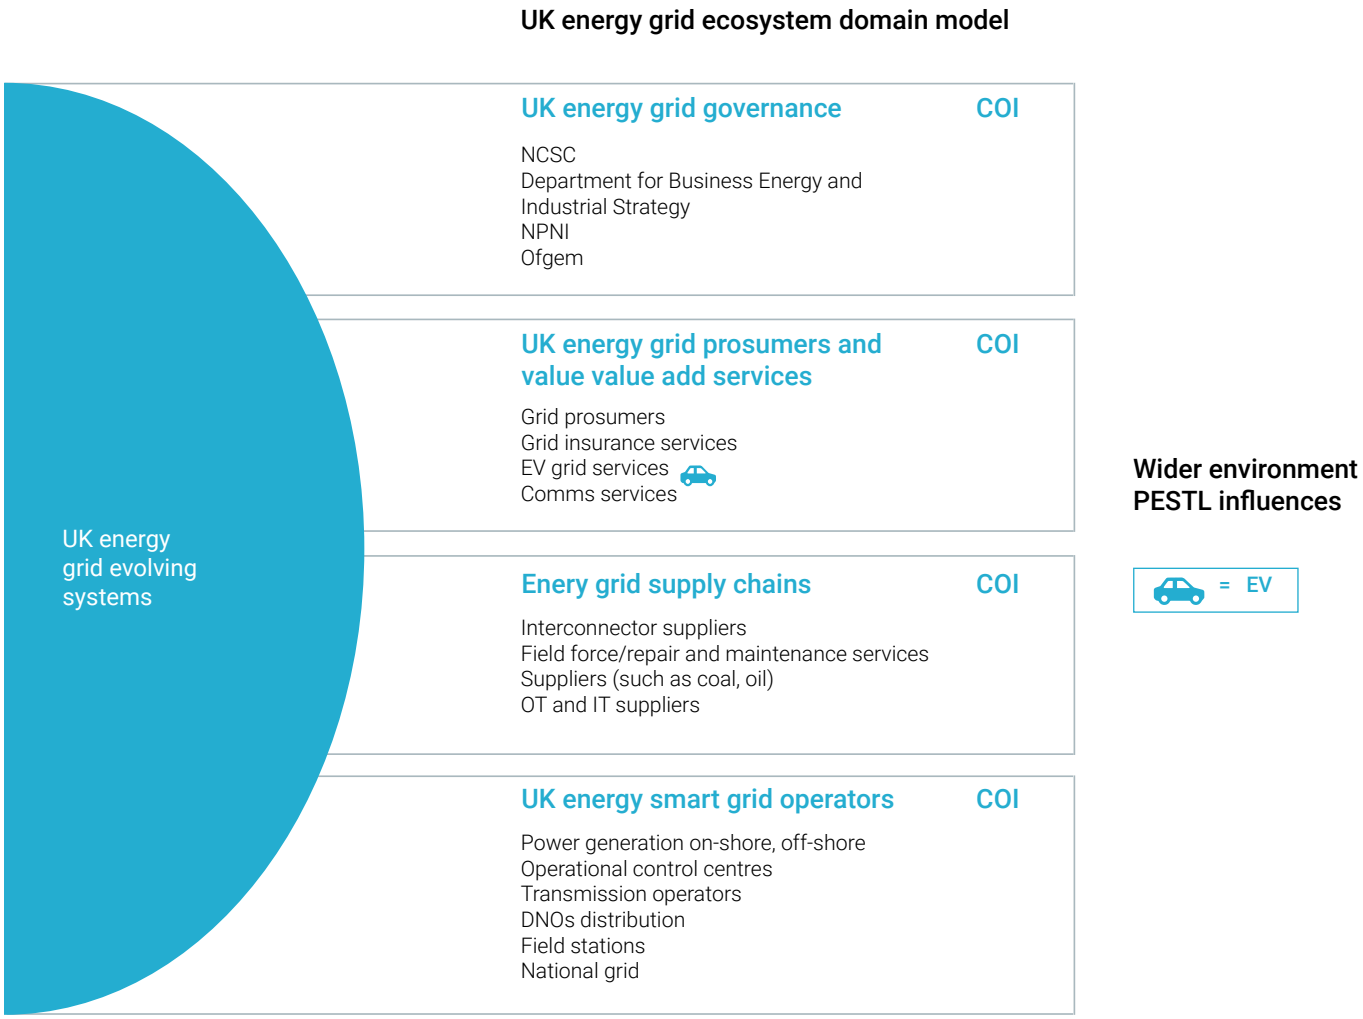

Figure 5 – UK CAV ecosystem domain model

THE UK ENERGY GRID OPERATIONAL SYSTEM COI

The COI depicted by the domain on the left hand side of Figure 5 comprises the operational systems of each of the critical members of the UK energy grid. These critical energy grid members include:

National transmission network and system operators:
- Transmission Network Control Centre, for example National Grid Control Centre (NGCC)
- Gas National Control Centre (GNCC)
- National Grid (group) Systems Operator

Distribution network operators (DNOs):
There are 7 DNOs in the UK, one of which is in Northern Ireland.

Power generation plant operators:
- Nuclear, coal (until 2025), oil, gas, wind, hydro and diesel farms (backup)

Together, these entities supply gas and electricity to industrial, commercial and domestic consumers who, in the case of electrical energy, can be prosumers. Figure 6 provides a schematic view of the end-to-end electric power transmission network.

These critical infrastructures keep the UK running, and any cyber-attack that successfully disrupts them for a significant period of time would potentially have a systemic impact on the UK economy.

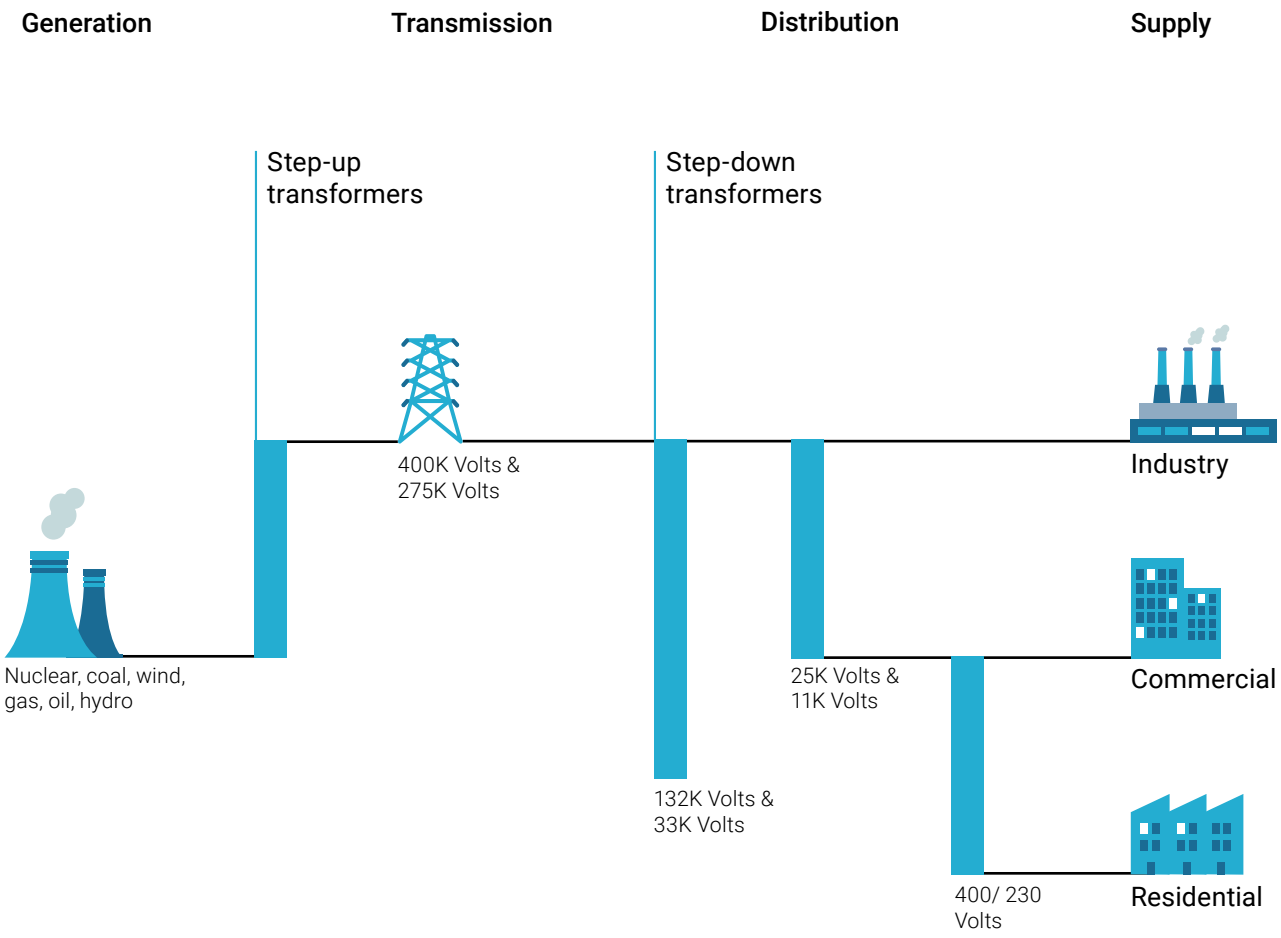

Figure 6 – Schematic view of the electricity transmission network

THE ENERGY GRID TRANSMISSION NETWORK AND SYSTEM OPERATORS

National Grid is the operator of the gas and electricity national transmission systems within the UK and is responsible for managing the response to any gas and electricity supply emergencies. There is a National Grid-operated control centre for the UK electricity transmission network (NGCC) in Wokingham and one for the UK gas network (GNCC) in Warwick.

This paper focuses on the National Grid electricity transmission network. This is owned and maintained by regional transmission companies, while the system as a whole is operated by a single system operator (SO). This role is performed by National Grid electricity transmission plc (NGET), and is responsible for ensuring the stable and secure operation of the entire transmission system.

The National Grid systems operator (SO) is tasked with progressing whole-system smart solutions, and for working closely with the industry to develop new services and approaches to a smarter energy system.

THE DISTRIBUTION NETWORK OPERATORS

The distribution network operators (DNOs) are companies licensed to distribute electricity in the UK. These companies own and operate the system of cables and towers that bring electricity from the national transmission network to the end users, including homes and businesses.

There are seven DNO's operating across various regions of the UK, as shown in Figure 7. Each DNO interfaces with the National Grid transmission network, which provides the UK national backbone of the electricity transmission network.

Seven DNOs in the UK:

- SSE
- SP Energy Networks
- Northern Ireland Electricity
- Electricity North West
- Northern Power Group
- Western Power Distribution
- UK Power Networks

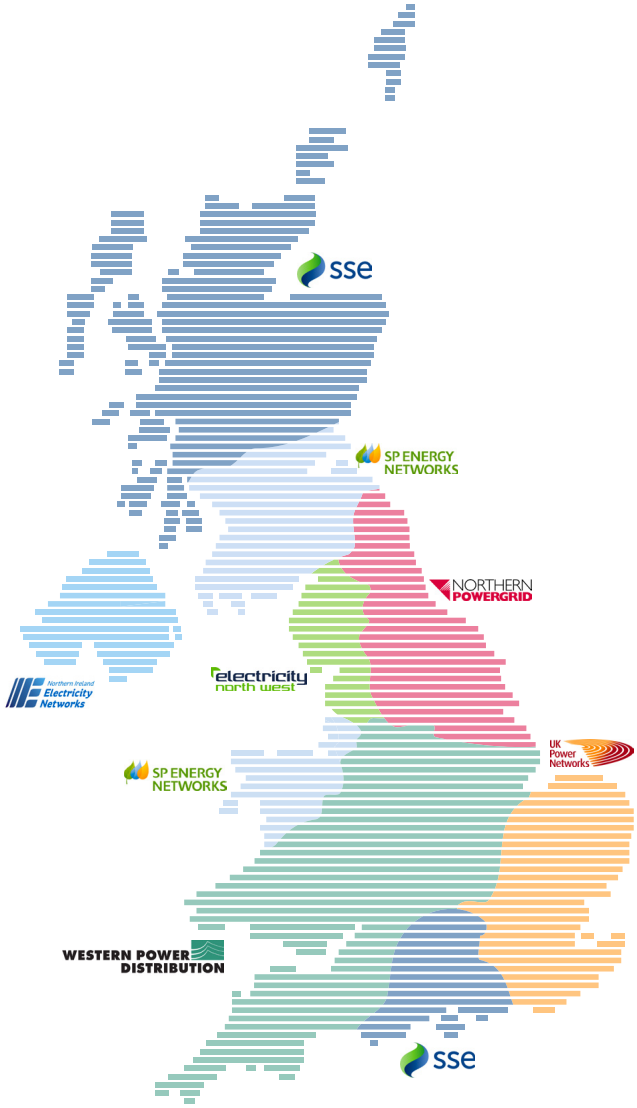

Figure 7 – The seven UK distribution network operators

POWER GENERATION PLANT OPERATORS

There are many diverse forms of power stations distributed across the UK, ranging from nuclear power stations to old-world coal-fired power stations. and including a growing number of green energy sources, such as wind, solar and hydro.

The National Grid also has a number of contingency power generation sources, including international interconnectors and diesel farms.

The operator companies owning these power generation plants are also many and diverse.

THE UK ENERGY GRID GOVERNANCE COI

The UK energy grid governance community of interest depicted in Figure 5 is comprised of the key UK government agencies responsible for, or involved in, the regulation and governance of the UK energy services.

The Department for Business, Energy & Industrial Strategy (BEIS) works with industry, regulators, sector bodies and other stakeholders to improve and maintain the resilience of the energy infrastructure, networks and assets; to reduce vulnerabilities; and ensure an effective response to actual or potentially disruptive incidents.

BEIS, as the UK competent authority and lead government department (LGD) for gas and electricity emergencies, is responsible for the development, review, updating and testing of the arrangements outlined in this document. Updated documents will be approved by the appropriate BEIS Minister.

The Office of Gas and Electricity Markets (Ofgem), is the National Regulatory Authority in Great Britain, and is responsible for regulating the gas and electricity markets in England, Wales and Scotland. Ofgem is also responsible for ensuring market arrangements are established and maintained, to minimise the possibility of gas or electricity supply disruptions.

In an emergency impacting the operation of the regulated or licenced gas and electricity markets, Ofgem would form part of the emergency response team. Its role includes providing guidance on market operation, industry codes and regulatory arrangements.

The National Cyber Security Centre (NCSC) would be part of any cross-government response to a category 1 or 2 cyber attack on the UK energy CNI. This would be in the context of the National Security Council (threats, hazards, resilience and contingencies) which operates from the Cabinet Office Briefing Rooms (COBR) and manages the central government response to emergencies, including those in the energy sector.

The Centre for the Protection of National Infrastructure (CPNI) is also a member of the UK energy grid governance COI, as the energy grid is a UK CNI. The CPNI is the government authority for protective security advice to the UK national infrastructure. Its role is to protect national security by helping to reduce the vulnerability of national infrastructure to terrorism and other threats, and is are accountable to the Director General of MI5. Although the National Grid is an operator, it would also participate as a member of the emergency response team in the event of a significant cyber attack on the energy sector.

THE UK ENERGY GRID PROSUMERS AND VALUE ADDED SERVICES COI

The wider energy services COI of the UK energy grid ecosystem, as depicted in Figure 5, is comprised of the diverse entities that provide energy services outside the core critical transmission grid services. These include heavy industry, light industry, commercial and domestic consumers and prosumers. Over time, the role of prosumers will grow as the UK energy grid evolves to a more distributed and adaptive form, where prosumers provide a significant proportion of the overall energy generating capacity. For example, electric vehicle (EV) users are prosumers of growing importance in the evolving UK energy smart grid, and form part of the growing IoT.

Other members the COI include insurance providers for the energy sector, and other value-adding participants, such as communications providers.

THE UK ENERGY GRID SUPPLY CHAIN COI

The UK energy grid supply chain COI, shown in Figure 5, is comprised of the global supply chains associated with the evolving UK energy grid. There different types of supply chains with this COI, range from environmental services through fuel supplies and shipping to those involving the long-term development cycles of new smart grid technologies.

Members of this supply chain COI provide the hardware and software components of UK energy grid service provider operational platforms, and include international interconnector services, cloud services and data centres.

Other key members of the supply chain COI include the field force organisations and their personnel, who provide on-site repair and maintenance of the UK operational energy grid infrastructure.

Together these supply chains comprise a massive socio-technical threat landscape that could be exploited by threat actors wishing to target the UK energy grid.

COMPLEXITY AND EVOLUTION OF THE UK ENERGY ECOSYSTEM

The UK energy transmission grid is dependent on the topology of the network formed by electrical buses and their interconnections, and its operating conditions, such as supply and demand distribution.

The drive towards different energy generation types and a prosumer model is partly being accelerated by a digital transformation of the way in which the energy grid operates. The UK energy grid ecosystem is evolving into a complex smart model, where real time data from the IoT will enable it to intelligently control overall electricity supply and demand, based on network analysis, simulations and both behavioural and predictive analysis. This transformation of the ecosystem is being driven by a number of emerging technologies, including:

- Artificial intelligence and machine learning
- Big data analytics
- Cloud services
- The IoT anvd associated 5G
- New communications protocols in addition to OT modbus and profibus, such as IEC 60870

This paper will explore how these emerging technologies bring with them not only significant potential benefits, but also a whole new threat surface with associated cyber risks.

As the complexity within the UK energy ecosystem will increase over time, its cyber resilience is a fundamental requirement for protection of the UK economy.

Figure 8 provides a high-level view of the conceptual architecture of the UK energy grid ecosystem. It depicts the key emerging technologies, together with its power and communications core.

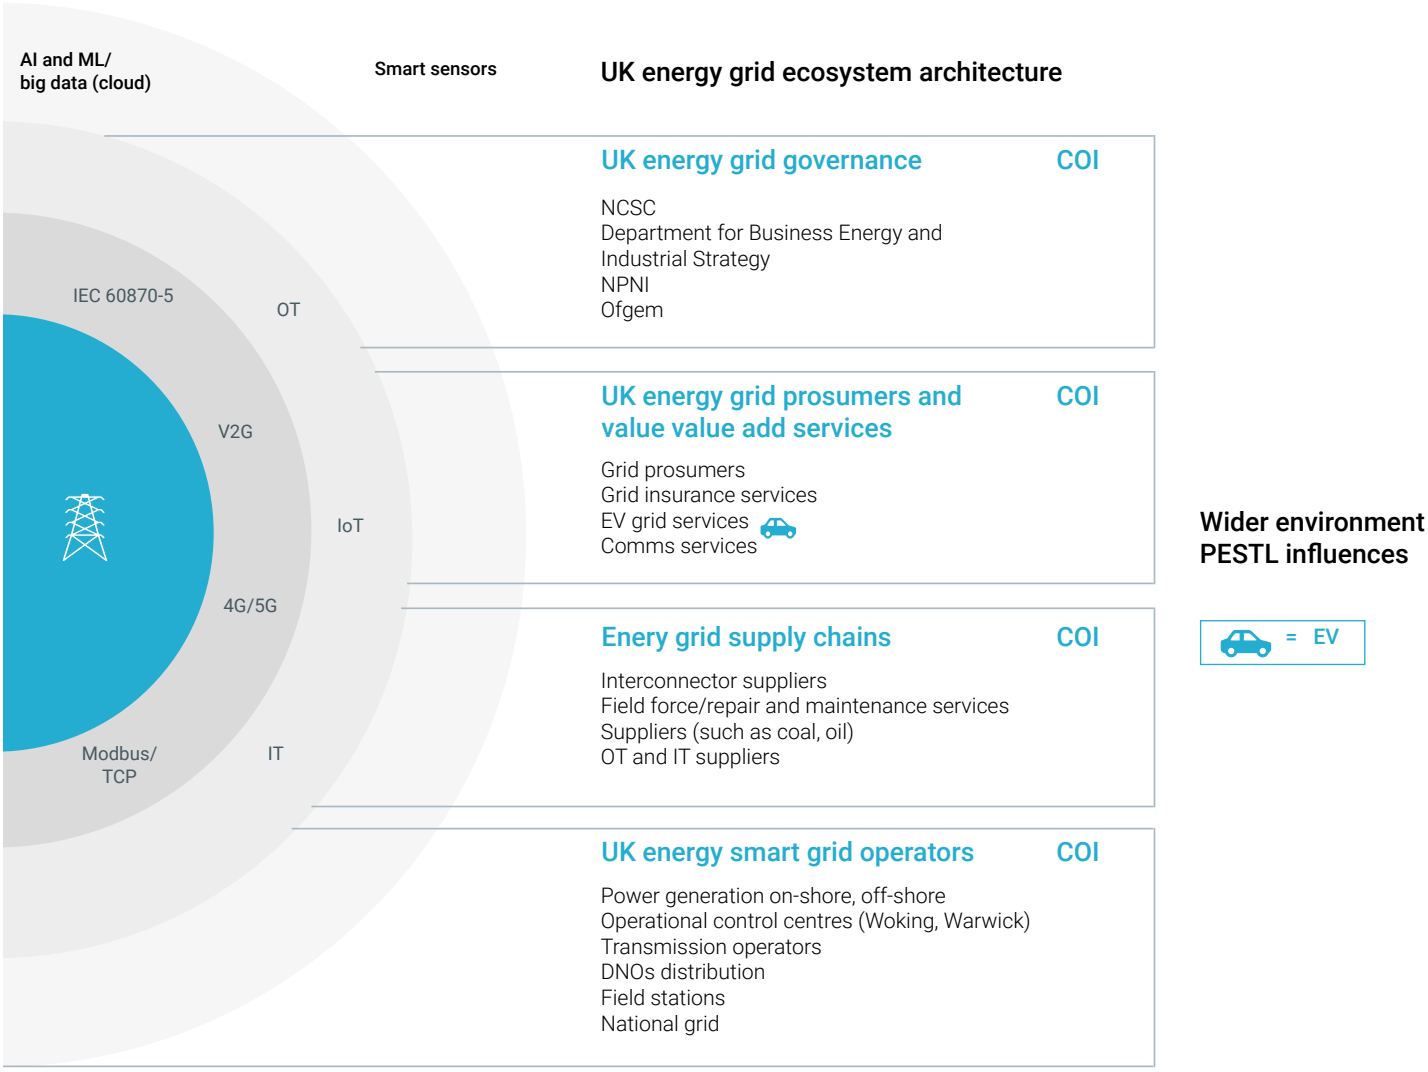

Figure 8 –UK energy grid ecosystem conceptual architecture view

The core UK energy grid infrastructure

As shown in Figure 8, the heart of the energy grid core is the electricity transmission network, which is managed by the National Grid Control Centre (NGCC) operations team using their integrated energy management system (iEMS). To provide high availability, this system is distributed over two secure data centre locations, with access/backup available from a third.

These locations are:

- NGCC Data Centre 1: Wokingham Bearwood Rd, Sindlesham, Wokingham RG41 5BN
- NGCC Data Centre 2: Warwick Technology Park, Gallows Hill, Warwick CV34 6DA
- NGCC Data Centre 3: Reading

These sites each act as backup disaster recovery sites for the other and are a critical component of the UK energy grid.

Supporting the iEMS is National Grid's core network infrastructure and operational system, which has as its foundation an optical fibre network (Optel), originally provided by Cable & Wireless, and now by Vodafone.

Disruptive technology services

The conceptual architecture of the communication core of the UK energy grid ecosystem also features a layer of technology capabilities highlighted by the yellow annulus in Figure 8. These transformative technology capabilities include:

- Cloud services
- Artificial intelligence (AI) and machine learning (ML)
- Big data analytics services
- IoT and 5G

Cloud services

Energy companies are rapidly adopting cloud services, which provide a powerful set of tools to manage data needs. This shift introduces new opportunities for combining public and proprietary data into big data, which can be used to generate innovative new analytical insights. Over time, increasing numbers of core services are likely to migrate to cloud hosting.

Big data analytics services

The use of big data analytics in the UK energy ecosystem is enabling transformation in a number of areas, including behavioural analysis of energy usage. However, the emerging technologies delivering such transformation are the AI and ML algorithms that feed on big data.

Artificial intelligence (AI) and machine learning (ML)

The use cases for AI and ML are constantly changing, but there are a number of ways that deep learning neural networks can be applied to support the intelligent operations of the energy smart grid including:

- Distributed generation management
- Outage management

AI and ML can be developed to exploit big data to make informed real-time operational decisions, for example, quantifying energy variations, and optimising energy demands and customer behaviours across the ecosystem - right down to the IoT 5G level of granularity. However, although AI and ML can provide the energy ecosystem with many advantages, they also present new challenges, such as the governance of and reporting on AI- and ML-driven transactions that can out-perform human understanding, as experienced with deep neural networks.

While energy companies are using AI and ML to improve their energy distribution processes and to better detect suspicious activity, malicious actors are also using AI to create new cyber threats, for example, by injecting biased data into the training sets of ML algorithms that can then be exploited by an attacker.

IoT and 5G

The UK smart energy grid ecosystem is largely comprised of smart interconnected devices. This includes devices such as line sensors in the transmission network (for example, such as phasor management units (PMU) for voltage and current), intelligent end devices and remote terminal units. In the domain of connected substations, the smart grid includes devices such as transformers, switches and protective relays. In the distribution network, it includes devices such as intelligent inverters, switches and power quality meters. Prosumers will also connect electric vehicles to the grid.

Smart grid control systems have special latency and performance requirements for their underlying communication networks. They may use communications standards such as IEC 61850-5, an international standard defining communication protocols for intelligent electronic devices at electrical substations. However, old-world OT-based protocols like Modbus are still also used in these supervisory control and data acquisition (SCADA) operations. Modbus is typically used between the supervisor computer and the remote terminal unit (RTU) for data transmission and control functions.

As the UK energy ecosystem evolves by extending its connected IoT, even more data will be generated and used. This is where 5G comes in, to ensure that both the radio and core network performance requirements can be met in terms of (end-to-end) latency, reliability and availability for different services.
The wider IoT includes micro-grids such as those operated by smart cities or by home/business owners.

APPLYING THE BHI USING AN ILLUSTRATIVE CYBER ATTACK SCENARIO

The BHI approach models the growth of benefits and harm in the context of complex cyber ecosystems. The high-level model of the UK energy grid ecosystem can be used as a context for showing how this approach can be applied.

The first step is to model the growth of benefits over a time period. The time period selected here is 2019 to 2030, which corresponds to the period during which the smart digital transformation of energy services will take place.

A simple model of the overall benefits growth associated with this digital transformation is shown in Figure 9, and assumes that the benefits grow as a Bass diffusion in line with the projected rate at which emergent technologies are adopted into the UK energy grid ecosystem.

The second step is to model the growth of risks (as a combination of likelihood and adverse impact) over the same period. To do this, the scenario of a multi-vector cyber attack on the UK energy grid ecosystem is used to illustrate and explore the associated risks as a function of time during the evolution of the ecosystem.

Firstly, the evolution of the risk likelihood (threat level) is explored by assuming the threat source is a nation state (with associated capabilities); by factoring in the growth of the threat surface (vulnerabilities such as opportunities) over time; and by modelling variations in PESTL influences, such as motivation. Figure 9 illustrates this risk likelihood in red, for example, the threat level assuming constant motivation but an exponential growth in vulnerabilities during the transformation period.

The potential impact is then modelled by exploring the potential of the cyber attack to generate a cyber chain reaction that poses a systemic risk to the UK. A systemic risk is generally seen as the potential for a major financial crisis adversely affecting the real economy. The vulnerability level and stochastic nature of the UK energy grid ecosystem during the transition period exposes it to so-called 'black swan' events that can result in systemic impacts.

In 2017, the UK energy services sector generated £31.7 billion of value for the UK economy[3]. The benefits associated with the smart grid evolution have been estimated as saving the UK £19 billion by 2050, exports of £5 billion by 2030, and the creation of over 8,000 jobs.

Bass diffusion distribution of UK energy grid distruptive technologies benefit growth and the association of the liklihood of systemic cyber attacks

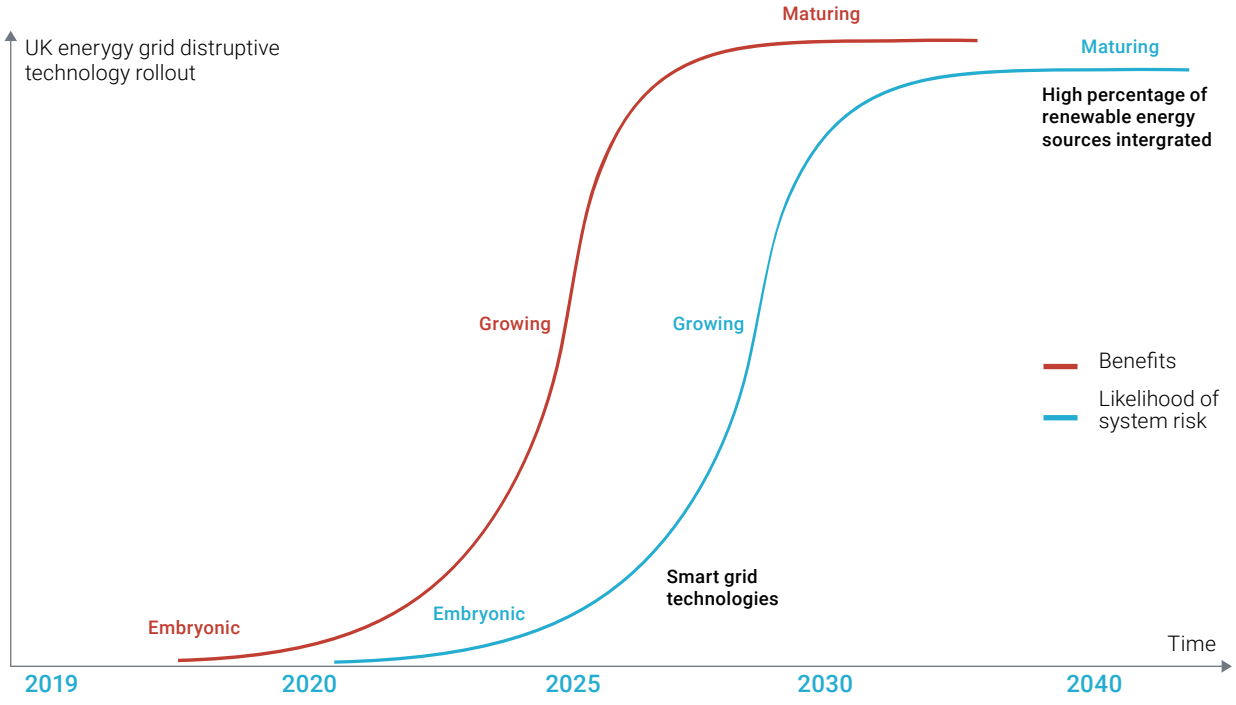

Figure 9 – Model of the growth in the UK energy sector benefits generated by distruptive technologies, versus the associated growth of liklihood of systemic risk

Secondary industries, such as electric vehicles, will be enabled by the smart grid; the benefits of electric vehicles have been estimated as between £25 billion and £60 billion by 2030[4].

However, the UK energy grid ecosystem also underpins the operation of the UK economy. Without electricity, more or less everything stops - including other CNIs such as the banking system, transport, communications, hospitals and water supplies. If you are a hostile nation state and you want to take down the UK economy, then taking down the UK energy grid ecosystem for a prolonged period is a natural attack vector. Any successful cyber attack generating a systemic impact on the UK energy grid ecosystem would therefore result in significant impact on the UK economy. In 2018 the UK gross domestic product (GDP) was £2.14 trillion, placing the United Kingdom fifth in the GDP ranking of 196 countries[5].

In a simplistic model, it can be assumed that the UK generates around £0.0058 trillion each day, which is £5.8 billion per year. However, taking the UK power down for just one day would cause at least two days of disruption, including significant potential losses of data. Taking the power down for three days would have a much more significant socio-economic impact, as standby generators would start to fail, resulting in deaths in hospitals, and a general lack of TV and internet communications would cause social unrest.

The impact of such a prolonged outage would be difficult to quantify, but is likely to be more than the pro-rata £18 billion of lost GDP, and would need to take into account impacts on intangible assets such as the UK's national brand equity, reputation and trust, resulting in potential downturns in UK investment.

The third step in the BHI process, therefore, evaluates the difference between the growth in benefits and growth in harm during the transformation period.

Category level	Category definition	Who responds?	What do they do?
Category 1 National cyber emergency	A cyber attack which causes sustained disruption of UK essential services or affects UK national security, leading to severe economic or social consequences or to loss of life.	Immediate, rapid and co-ordinated cross-government response. Strategic leadership from ministers/Cabinet Office (COBR), tactical cross-government co-ordination by NCSC, working closely with law enforcement.	Co-ordinated on-site presence for evidence gathering, forensic acquisition and support. Co-location of NCSC, law enforcement, lead government departments and others where possible for enhanced response.
Category 2 Highly significant incident	A cyber attack which has a serious impact on central government, UK essential services, a large proportion of the UK population, or the UK economy.	Response typically led by NCSC (escalated to COBR if necessary), working closely with law enforcement (typically NCA) as required. Cross-government response coordinated by NCSC.	NCSC will often provide on-site response, investigation and analysis, aligned with law enforcement and criminal investigation.

Table 2 - NCSC definition of Category 1 and 2 cyber attacks/incidents

THE ILLUSTRATIVE CYBER ATTACK SCENARIO

The illustrative and hypothetical cyber attack scenario assumes that the threat source is a nation state entity, and for this sake of this exercise that fictional role has been given to Russia and its Federal Security Service (FSB) acting through an insider threat actor and one or more proxy APT groups, including APT 29 (a Russian hacker group, Cosy Bear).

The potential capability of this nation threat source is high, although the actual motivation to carry out a major cyber attack on the UK energy grid ecosystem is assumed by default to be low. This fictional scenario is of a hypothetical escalation of any Anglo-Russian geo-political tension. This hypothetical and escalation is imagined to be the result of Russia and Iran being targeted by America and the UK with a major new round of sanctions in the year 2020, as a result of an escalation in the weaponisation of gas and oil supplies. The cyber attack scenario models here are based on a sustained multivector cyber attack targeting the UK energy grid ecosystem with the objective of causing a systemic impact on the UK. In NCSC terms, this would equate to a category 1 or 2 cyber attack, as defined in Table 2.

Our hypothetical multi-vector cyber attack is intended to cause a systemic impact on the UK energy grid ecosystem with the aim of generating severe damage to the UK economy.

As part of the threat intelligence approach, it is noted that Russia views cyber attacks as a sub-component of information warfare, which covers a broad range (including computer network operations, electronic warfare, psychological operations and disinformation operations). The fictional cyber attack scenario has been designed around this broader multi-vector approach.

The cyber attack in this scenario features the following three attack vectors:

- An NGCC insider attack exploiting the integrated energy management system (iEMS), a SCADA system for managing the National Grid's transmission assets
- A carefully timed kinetic attack on the optical fibre network connections to the NGCC Data Centres in Warwick and Wokingham
- A synchronised external APT 29-led attack on the wider electricity grid, exploiting electrical vehicles, solar farm inverters and SCADA Modbus vulnerabilities on targeted backup generators and primary substations

Nation state threat sources will invest in building and imbedding their information warfare capabilities over many years, including placing insider threat actors as sleepers into the CNI operational companies and associated supply chains of target nation states.

ATTACK VECTOR 1: AN NGCC INSIDER ATTACK EXPLOITING THE IEMS

In this purely hypothetical scenario, it is assumed that the Russian FSB has embedded an insider threat actor within the NGCC in Woking, as part of the iEMS system monitoring team. With physical access to the iEMS system and authentic identity credentials, the threat actor is able to install sophisticated malware supplied by APT 29 onto the iEMS/SCADA supervisory computers.

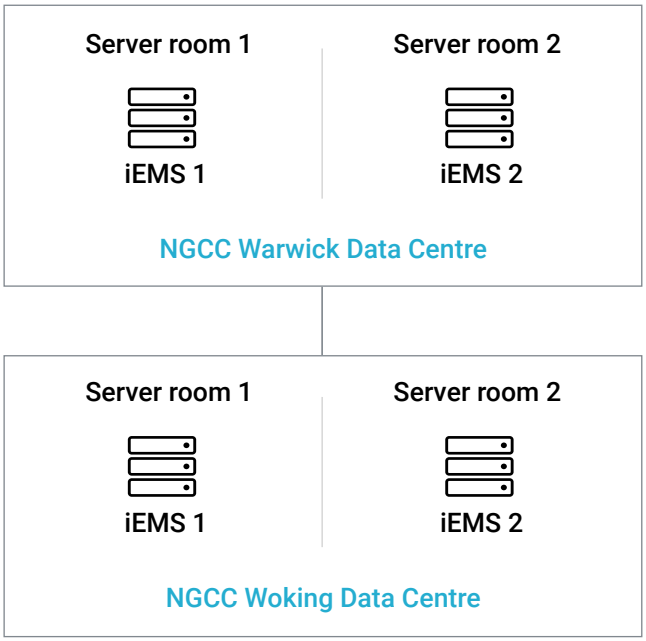

Figure 10 - Primary and secondary iEMS in the two key NGCC data centres

As shown in Figure 10, there are two iEMS production systems: one in the NGCC Woking data centre and the other in the NGCC Warwick data centre. Each production iEMS has its own backup in a separate server room.

The insider attack is designed to ensure that each of the iEMS production systems and their associated backup systems is infected with the APT 29 malware. In this scenario, the malware has been designed to evade intrusion detection systems (as operated by the NGCC security team) until it executes. A schematic of the APT 29 toolkit used by the inside threat actor is shown in Figure 11.

Exploiting their physical access and credentials, the inside attacker uses response injection attack tools on the human machine interface (HMI) system to generate a false reassuring (situation normal) user view in, for example, CIMPLICITY - the system which provides the wider operations team with client server visualisation and control.

The actor then infects the master terminal unit (MTU) with command injection malware that can escalate privileges to execute all switchgear commands. It can also spread backdoor access to the remote terminal unit's (RTU) slave programmable logic controllers (PLC)through over the air (OTA) firmware updates.

Intrusion detection is avoided by using appropriate message sizes and timing, and ensuring that command sequences are all legitimate. The inside threat actor has the necessary knowledge of the UK power grid layout, and intends to cause maximum disruption to the network's performance by attacking as few nodes (electrical buses) as possible. The insider also exfiltrates design data, such as PLC physical addresses and IP information to APT 29.

The iEMS attack strategy exploits telecommand switching to target switchgear in the most critical or most vulnerable nodes, where removal significantly disrupts the functioning of the UK network, while still displaying a false 'situation normal' status on the HMI.

Each transmission line in a power grid is associated with a maximum safe flow-carrying capability. If these limits are exceeded, the situation is detected by protection relays, circuit breakers are tripped and the corresponding element is taken out of service. The possibility and negative impact of cascading failures in power grids increases when the operating point of a power grid is close to the flow-carrying capabilities of its links.

Case studies in both real-world and test power grids[6] show that they are highly vulnerable to targeted attacks. Sequentially removing the nodes with the highest centrality is an effective strategy to fragmenting power grids, and decreasing their operational performance. In almost all power grids, the removal of approximately 15% of the nodes according to flow centrality will result in almost complete destruction of the network.

The suite of malware used in the first hypothetical attack vector includes the ability to overwrite firmware in critical systems, using tools such as Killdisk to destroy data in essential files in all four iEMS server rooms - including overwriting master boot records. This suite of malware is invoked as the final phase of the first attack vector, prolonging the recovery time required to get the system back up and operating.

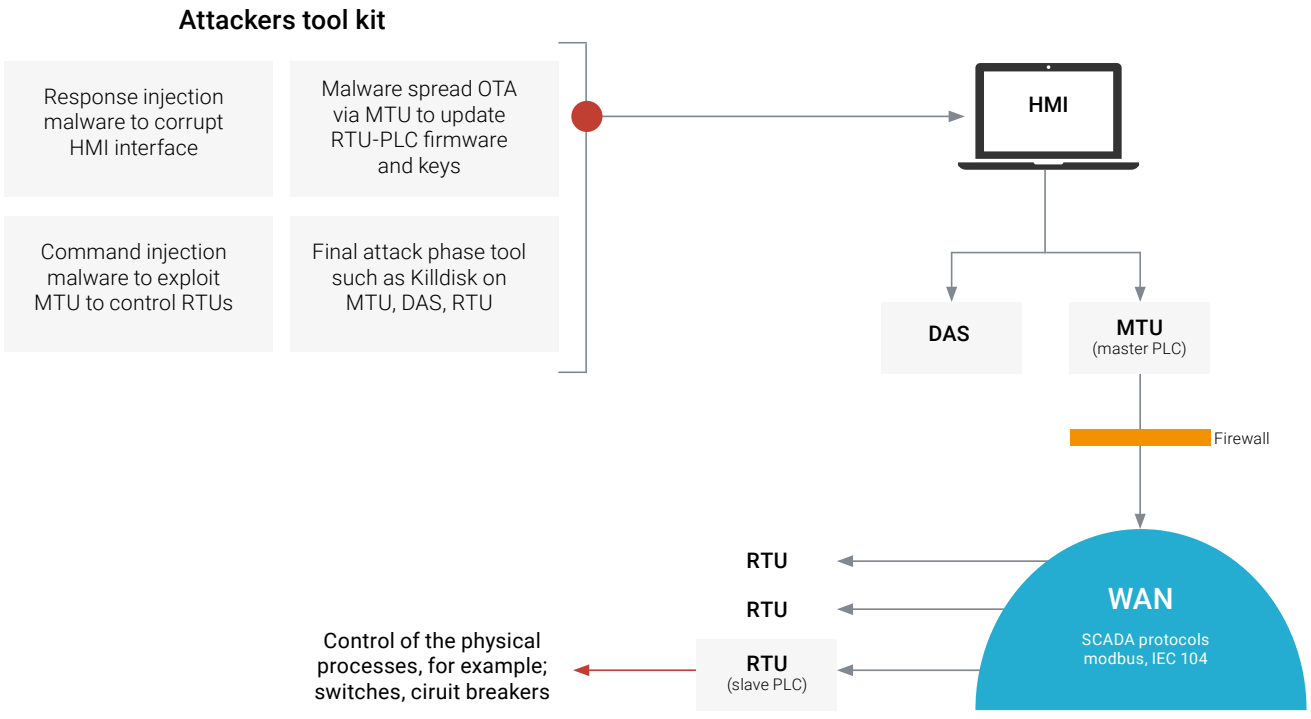

Figure 11 - Insider threat actor iEMS/SCADA attack vector/tool kit

ATTACK VECTOR 2: KINETIC ATTACK ON DC OPTICAL FIBRE NETWORK CONNECTIONS

The second attack vector is launched almost immediately after the first, and is aimed at prolonging the blackout period by exploiting the physical vulnerability of the fibre optic cables serving the NGCC Woking and Warwick data centres.

It has been assumed that these cables are located in ducts located at different ends of each of the NGCC DC sites. Once these ducts leave the site boundary, they merge with the main multi-tenant cable ducts running along nearby main roads. By cutting all the cables in these ducts at the nearest road-based maintenance point (for example, by using fake telecoms provider vans or roadworks as a cover for the operation) both NGCC data centres can be effectively put offline for a period of hours, just when they are needed to try and recover the energy grid.

ATTACK VECTOR 3: AN EXTERNAL APT 29 ATTACK ON THE WIDER GRID

The third hypothetical attack vector exploits both the vulnerabilities associated with the SCADA systems in generators and substations and those associated with prosumer charge station interfaces (such as the growing number of connected electrical vehicles (EVs) and solar farms).

As part of its advanced persistent threat activities, the threat actor APT 29 will have exploited the EV supply chain to create back doors into the motherboards of battery management systems (BMS) for a number of major EV brands, and in solar inverters. They will also have invested significant time and effort in the reconnaissance of SCADA vulnerabilities in the wider UK energy grid ecosystem.

In this third attack vector, the APT 29 threat actor will launch targeted attacks on internet connected SCADA devices in a way that amplifies the impact of the first attack vector. For example, they will target the diesel backup generators of key UK facilities at the same time as the NGCC insider attack takes down major parts of the UK electricity transmission network. The APT techniques here would typically involve address resolution protocol (ARP) poisoning to facilitate a 'man in the middle' (MIM) attack that continues to cause damage as the NGCC operations team try to regain control. This will prevent high profile NHS hospitals and other critical industry players, such as key data centres, from using backup generators, thus causing loss of both life and data.

In this fictionalised scenario, the APT 29 threat actor also launches a synchronised attack via C-2VX/LTE-V connectivity to set a timer on the BMS malware it has installed. This then causes all electric vehicles that are connected to charge points on the UK energy grid overnight to synchronise their charging, creating a peak load, while permanently disabling the capability of any electric vehicle to supply power back to the grid.

The threat actor would co-ordinate this attack vector with supply chain exploitations in other inverters, such as those enabling solar power farms to supply power into the UK energy grid. By remotely controlling the power flow, the actor can cause peaks and troughs of several gigawatts, resulting in disruptive power-balancing issues for the grid, in conjunction with the first attack vector.

This third attack vector is exploited just after the first two attack vectors are complete, in order to maximise and prolong the damage to the energy grid.

EXPLORING VULNERABILITY AND CONTROL ASPECTS OF THE SYSTEM

The complexity of the UK smart energy grid ecosystem is growing as it undergoes digital transformation driven by the emerging technologies described in Section 2 (such as AI and ML, big data analytics, IoT, 5G and cloud services).

The vulnerability of the UK energy grid ecosystem is increasing in line with its complexity. This is because of the expanding threat surface presented by the interconnectivity and collaboration that these technologies enable, and their links with legacy Modbus/SCADA technologies. The illustrative cyber attack scenario shows how these evolving vulnerabilities can be exploited in different ways, via multi-vector attack paths.

As shown in Figure 9, the digital transformation period associated with smart grid technologies is 2020 to 2025. During this period, the Bass diffusion of the introduction of these smart grid technologies is growing exponentially, as is the corresponding threat surface, which means that the associated vulnerabilities which can be exploited by the first two attack vectors are also growing exponentially. However, the vulnerabilities associated with the third attack vector (such as the physical optical fibre network connections to the two NGCC data centres) are relatively well-known and constant.

Classic risk mitigations work on the basis that a system can be controlled in the presence of threat actors, reducing or removing the threat. Control of small systems is a mature discipline: controllability of linear systems is well understood, and understanding for non-linear systems has been developing steadily. However, ecosystems that include prosumers, and the information and computer technologies that support smart energy distribution and demand management, are becoming increasingly complex. Control of such complex systems - including distributed networks of actors and components - and control of systems of systems are poorly understood, and mostly poorly characterised. A threat actor can leverage this lack of knowledge to cause harm to a system in ways that a defender cannot control through prior mitigation.

In the BHI model, the VL of a system – of a given scope and phase space with a given resolution – is a measure of its intrinsic lack of controllability, from the perspective of the defenders who legitimately operate the system. In the scenario here in 2020 to 2025, the vulnerability level of the UK energy grid ecosystem is at level 4, in line with its level of complexity. VL4 is shown from a control perspective in Table 3.

Vulnerability level 4 is representative of the fact that the UK energy grid system operators, and in particular its distribution network operators and prosumers, have no actual knowledge of, for example, the 'zero day' threats in the armoury of nation state- sponsored APT groups. The expanding threat surface associated with the emerging/collaborative technologies driving digital transformation makes it difficult, if not infeasible, to detect all such latent threats.

This lack of knowledge makes risk decisions far less certain even than gambling, as at least a gambler knows the odds against success. In the context of cyber threats, it is the attacker who holds the knowledge. In other words, the knowledge status supporting risk decisions has moved from rational ignorance to one of radical ignorance.

Vulnerability Level (VL)	Threat class	Attacker's control	Economic rationale
4	Stochastic system	The nature of the system is such that it cannot be controlled, but vulnerabilities can be reliably modelled using closed-form probability distributions over a fixed (and finite) set of state variables in the system's phase state space.	Radical ignorance: black swan events may occur, as preparation for such events is frequently hindered by an assumption of knowledge of all the risks. Scenario modelling using Shackle's potential surprise.

Table 3 - VL4

EXPLORING THE LIKELIHOOD OF AN ATTACK SCENARIO

Classic risk assessments model the likelihood of a cyber attack on a particular target of interest in terms of a threat level assessment at a given point in time. The threat level is typically modelled as a function of the capability of the threat sources/actors and the level of motivation and priority for attacking that target of interest.

Threat level = F(capability of threat source/actor(t), motivation/priority(t))

Where the capability and motivation are both functions of time.

The capability of the threat source and associated actors in our hypothetical example scenario are those associated with a nation state, in this case the capabilities of the FSB and their APT groups. The capability of such nation state actors for launching sophisticated cyber attacks is generally considered as being high.

The attack vector in the hypothetical scenario exploits significant vulnerabilities (attack opportunities) within the UK energy grid ecosystem, and these are growing exponentially in line with digital transformation and the explosive growth in the threat surface. Such vulnerabilities are relatively easy to exploit, so the likely capability of the threat actors relative to the difficulty in exploiting them is high.

Readiness of latent zero day threats to the UK's CNIs would give any hostile nation state the potential to launch a cyber attack with a significant socio-economic impact on the UK. Therefore, the likelihood of the UK energy grid ecosystem entities potentially being compromised via such latent zero day back doors is very likely.

The likelihood of an actual attack being executed that exploits (and thus exposes) any zero day vulnerability would depend on motivation and priority, which themselves would be driven by the state of the geopolitical relationship between the UK and the hostile nation state in question.

Although zero day threats are powerful cyber weapons, once they are used, they can soon become known and mitigated.

The motivation level in the year 2020 is high in our hypothetical attack scenario.

EXPLORING THE POTENTIAL IMPACT/HARM OF THE ATTACK SCENARIO

The illustrative scenario hypothesised in this report assumes that the Russian FSB and their APT groups launch an attack that exploits insider threat and latent zero day vulnerabilities to compromise UK energy grid ecosystem components. The objective is to cause economic damage to the UK as part of a geopolitical weaponisation of energy supplies campaign that begins to escalate in the year 2020.

The levels of harm involved are modelled using the examples of the impact on the UK energy grid ecosystem, as shown in Table 4.

Level	Impact
Very high	Overall capability of the UK energy grid ecosystem brought to a halt for 2 or more days. Significant UK-wide socio-economic disruption. High impact on all intangible assets. Systemic impact (for example NCSC Category 1 cyber attack, national cyber emergency).
High	Disruption of the UK energy grid resulting in one or more DNOs going down for one day or more. Significant impact on most intangible assets. Systemic impact (for example, NCSC Category 2 cyber attack).
Medium	Localised significant operational disruption of the intra-UK energy grid domain. Minor and prolonged (one day or more) UK-wide disruption of overall UK energy grid ecosystem operations. Minor impact on most intangible assets (for example, NCSC Category 3 or 4 cyber attack).
Low	Localised short-term operational disruption of the intra-UK energy grid ecosystem (for example, NCSC Category 5 cyber attack).

Table 4 - UK energy grid ecosystem impact levels

As shown in Table 4, when assessing the impact of a successful cyber attack on the UK energy grid ecosystem, the potential harm to both tangible and intangible assets must be included. For example, brand equity can be lost as a result of the reputational damage caused by succumbing to a successful cyber attack.

In our hypothetical scenario, the Russian FSB (and their APT groups) successfully launch their multi-vector cyber attack scenario in the year 2020. The impact of such a Category 1 cyber attack for example, a prolonged UK-wide blackout for a number of days, could extend to intangible assets in the following ways:

- An NGCC insider attack exploiting the integrated energy reputational damage to many of the players in the UK energy market, in particular the National Grid.
- Loss of trust in the UK for providing reliable energy infrastructure
- Potential migration of business and investment out of the UK, negatively affecting the UK economy

The wide-ranging impact on both tangible and intangible assets is an important aspect of such a cyber attack. Power loss would potentially cause maximum disruption to the operation of critical UK industries, including finance and transportation, as well as loss of life in sectors such as the NHS.

However, the impact level will be different at different points in time, as will the motivation of the attacker. For example, the impact in our hypothetical scenario would be more likely to be very high if carried out aggressively in 2025, when there will be significantly more dependence on prosumers (such as electrical vehicles and solar power farms). This would not only extend the impact of the attack, it could cause both planned and unplanned cyber chain reactions to propagate across the wider UK energy ecosystem.

An attack carried out in 2020, when the population of prosumers is much smaller, would cause a likely impact of high, rather than very high.

Having explored the illustrative cyber attack scenario in classic risk assessment terms, it can now be explored from the perspective of the BHI.

THE BENEFIT HARM INDEX PERSPECTIVE

The overall socio-economic benefits of the UK energy grid ecosystem grow over time in line with a Bass diffusion distribution, as shown earlier in Figure 9. As shown in the hypothetical cyber attack scenario, the harm which can be inflicted on the ecosystem by a specific threat can also grow with time, and the associated threat level will vary with time. Benefits are defined in terms of the positive business and socio-economic impacts multiplied by their likelihood. Harm is defined in terms of the negative business and socio-economic impacts multiplied by their likelihood. A simple discrete formulation of how to calculate the associated growth is shown below:

$$B_{t+n} = B_t + (b_t P_{bt} - h_t P_{ht})$$

For benefits B_t and harm h_t with probabilities P_{bt} and P_{ht} respectively.
Here P_{ht} is proportional to the threat level

Threat level = F(capability of threat source/actor(t), motivation/priority(t))

This shows threat level as a Function (F) of both the capability of the threat actor and their motivation to attack. Both these factors vary and are a function of time (t).

The BHI relates to differences in the complexity levels of benefit (CL_b), and harm (CL_h), over a time interval, TI_i, assuming M distinct threats (j) where j ranges from 1 to M.

$$BHI = CL_b(TI_i) - CL_h(TI_i)$$

Where:

$$CL_b(TI_i) = MAX \{ level(distribution(b(TI_i)), level(distribution(P_b(TI_i)) \}$$

$$CL_h(TI_i) = MAX \{ level(distribution(h(TI_i)) \}, \; MAX \; \Sigma distribution \; (priority \; (jTI_i)) \} \; j(level(distribution)(j(TI_i))),$$

In simple terms, for the hypothetical cyber attack scenario on the UK energy grid ecosystem there is an overall set of socio-economic benefits that are growing in line with a Bass diffusion distribution curve, as described in Figure 9.

During the growth period 2019 to 2025, the benefit growth rate is embryonic, which equates on average to a benefit complexity level 2. During the period 2025 to 2030, benefit growth is exponential, which equates to a complexity level 4.

During the period 2030 to 2035 the benefit growth rate decreases rapidly from exponential to asymptotic which equates to a benefit complexity level 4, decreasing to 0, during this period.

Given the UK government's strategy of ensuring that the UK remains a leading player in smart energy services, it can be assumed that the associated probability of following that distribution is high and flat, so for simplicity it is assumed that it is close to 1. The accuracy of the market forecasts is assumed to be high.

If initial assumptions are that the value of benefits and harm over each interval are the same then, in effect, there is a focus on the difference in the growth rates of benefit and risk likelihood, rather than on the actual quantitative benefit and harm multipliers.

For the period 2019 to 2025: **CL_b(2019-2025) = 2**

For the period 2025 to 2035: **CL_b(2025-2030) = 4 CL_b(2030-2032) = 2, CL_b(2032-2035) = 0.,**

The selected threat level (likelihood) for the example also grows with a Bass diffusion distribution curve. However, the Bass diffusion curve for potential harm grows in advance of the benefits Bass diffusion curve, since nation state threat actors will be targeting the UK's critical national infrastructure, by creating an arsenal of zero day threats for each CNI as part of their cyber warfare readiness capabilities. The potential for harm occurs prior to the generation of benefits, which take time to be realised.

The potential for harm associated with the hypothetical threat scenario will grow exponentially (for example, complexity level 4) in line with the growth in complexity of the UK energy grid ecosystem during the digital transformation period. This high level of complexity is associated with the explosive growth in the size of the overall threat surface for the selected attack vector. It can be expected that any nation state threat actors would exploit by further developing their arsenal of associated zero day threats.

As already mentioned, the motivation priority in the hypothetical threat scenario is low until the year 2020 when it becomes high, taking the threat level (likelihood) to a very high value.

So for the period 2020 to 2025: **CL_h(2020-2025) = 4** which reflects the strong growth phase of Bass diffusion distribution of the growth of potential harm. And for the period 2025 to 2030 the exponential growth in potential harm asymptotically decreases. The resulting BHI values simply show the difference in growth rate, as outlined in Figure 9: the potential harm is growing faster than the potential benefit in the earlier period 2020 to 2025, before evening out. Growth in potential benefits is faster in the later period, 2025 to 2030.

These calculations assume that the level of benefit and the level of harm were of equal magnitude for each time interval. However, the level of systemic harm that can be inflicted on the UK energy grid, and thus on the dependant UK economy, is largely relative to the incremental growth in benefit generated by the early stages of the digital transformation of the UK energy grid ecosystem.

Although there is no formal quantitative analysis in this paper, the value of the UK energy services sector contributed £31.7 billion to the UK economy in 2017[7].

Any single successful cyber attack generating a prolonged outage of the UK energy grid ecosystem would result in a systemic impact to the UK economy, which could easily result in a multi-billion pound loss to the UK, which has a GDP of £2.14 trillion[8]. In other words a two or three-day UK power blackout could easily wipe out all the benefits generated by the UK energy sector over a year.

This does not take into account impacts on intangible assets such as UK brand equity, reputation and trust for inward investment in areas such as electric vehicles. The CL_h values need to be multiplied to reflect this impact. The resulting BHI values are reflected in Table 5 where deeper red indicates increasingly negative BHI values.

	2020 - 2025	2025 - 2030	2030 - 2035
BHI value	< 0	< 0	< 0

Table 5 - of the hypothetical cyber attack scenario

In the case of BHI <= 0, the growth order (CL) of the harm exceeds the growth order of benefit. In such a case, unless there is mitigation, it is reasonable to expect that however the benefit grows, it will be overtaken by harm.

Even for just one hypothetical cyber threat scenario, the complexity of the ecosystem and the vulnerability levels of the components at these negative BHI time intervals make it hard to predict the full spectrum of associated cyber chain reactions. In Section 4 of this paper there is an illustration of this scenario in more detail, showing how the Implication Wheel™ can be used to try and detect emergent systemic threats in this context.

In applying the BHI formally, it can now be looked at systematically across a significant number of risks rather than just the one hypothetical example explored in this report.

AN APPROACH TO MITIGATING EMERGENT RISK/ RADICAL IGNORANCE

The approach highlighted here uses the Implication Wheel™ methodology to help uncover emergent threats. Figure 12 illustrates the context that will be used to introduce the Implication Wheel™ methodology. It features a hypothetical illustrative threat scenario as it could unfold in the UK energy grid ecosystem.

Cyber ecosystems are complex systems of systems, like the UK energy grid ecosystem explored in this paper. As described earlier, such ecosystems are constantly changing, often in surprising ways.

Cyber attacks on such systems can cause cascading cyber chain reactions with indirect and unanticipated consequences. The direct first order effects are often relatively easy to predict and mitigate, however, second and third order effects are much less obvious and may contain surprises, some of which will pose a systemic risk. These are referred to as 'black swan' events.

The Implication Wheel™ is participatory 'smart group' methodology that uses a structured brainstorming process to uncover multiple levels of consequences, and which can lead to the discovery of black swan events. Each smart group comprises a diverse set of individuals who will bring different perspectives to the task.

Illustrating a cyber chain reaction leading to systemic risk in UK CAV ecosystem

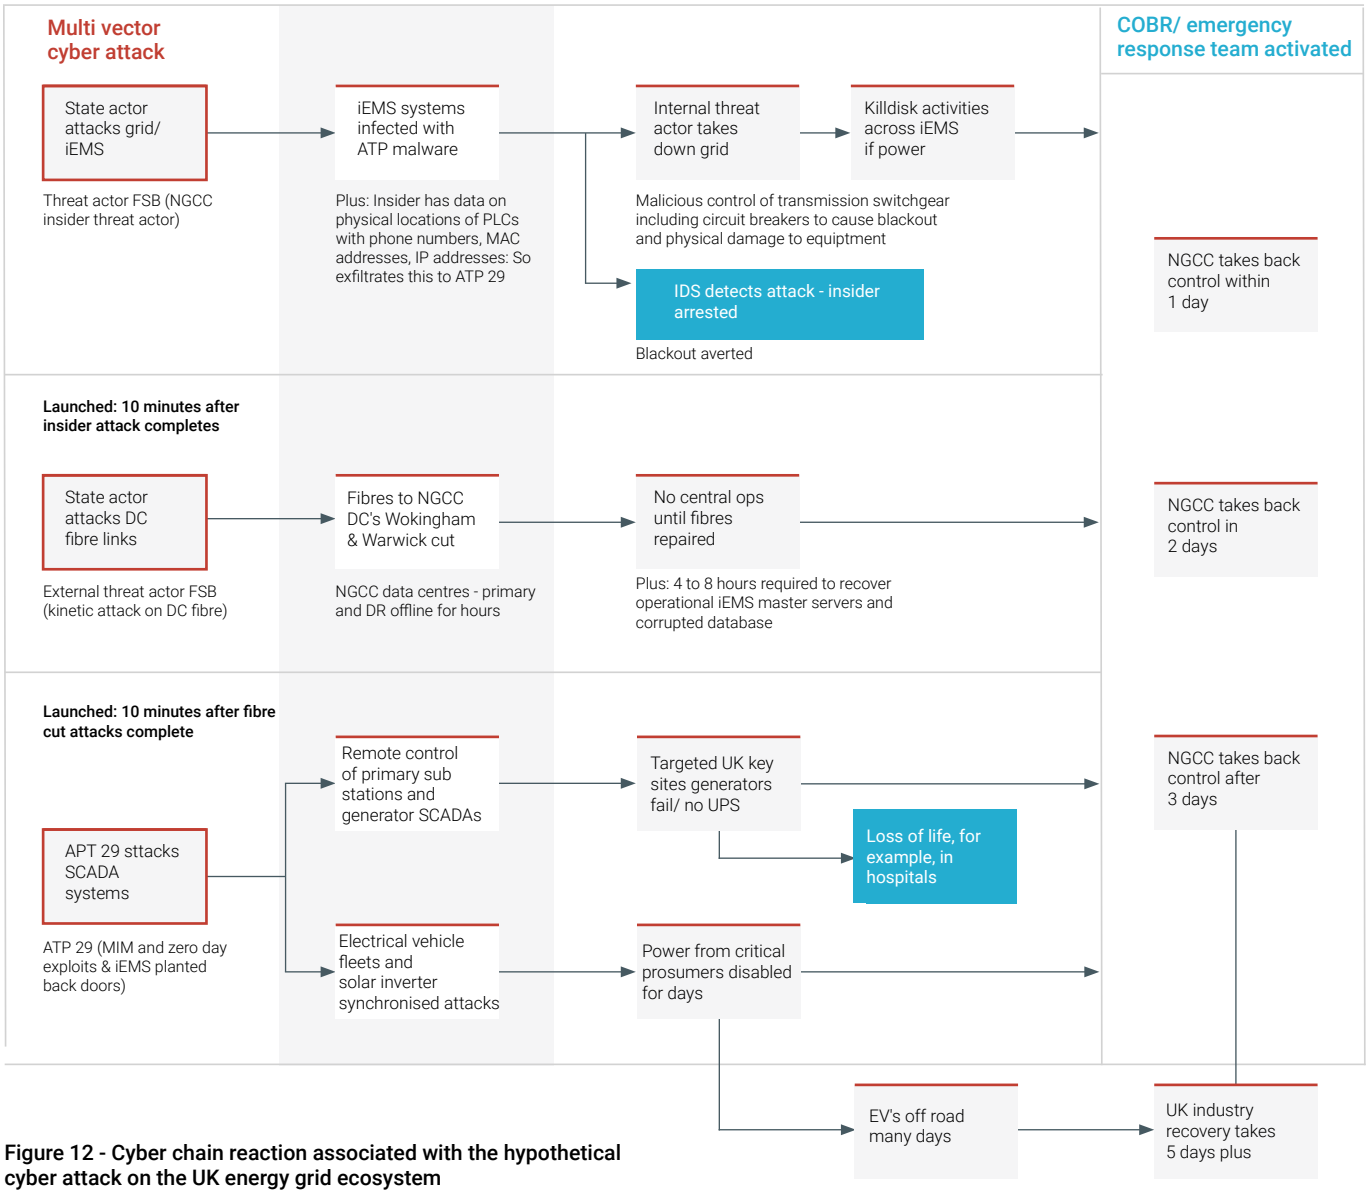

Figure 12 - Cyber chain reaction associated with the hypothetical cyber attack on the UK energy grid ecosystem

Each smart group starts by considering an initial event, such as the hostile state actor launching a hypothetical multi-vector cyber attack.

The example threat actors are shown in the red outlined squares on the left of Figure 12. The initial event resulting from the launch of the attack is represented by the set of white boxed activities in the grey highlighted column.

Each smart group is then asked "What might happen next?" This generates the direct first order consequences, as illustrated in the example in Figure 12.

These first order consequences include, for example, loss of connectivity to both NGCC data centres for two or three hours, followed by four to eight hours spent trying to recover the iEMS from the Killdisk system corruption. In this illustration, one of the significant first order consequences would be the taking down of the transmission grid by the insider. Another is the possibility that the insider's activities are detected by the NGCC monitoring team early enough to avert the blackout and arrest the insider threat agent. This "What might happen next?" process is then repeated by the smart groups for each first order consequence, creating an associated set of second order consequences. This process can be repeated to explore third order consequences, and so on.

Figure 12 shows a second order consequence equating to the NGCC being unable to restore power for three days, due to physical damage to remote switchgear equipment and ongoing APT 29 MIM SCADA attacks. This second order event could lead to a third order black swan event where many critical UK industries take over five days to fully recover normal operations, for example, due to irretrievable data loss during the prolonged blackout.

When the Implication Wheel™ is used more formally in this context, a layered structure is produced, like the wheel shown in Figure 13. This illustrates just one second order effect and its associated third order effects.

The Implication Wheel™ methodology permits smart group participants to propose levels of impacts and importance, and the likelihood of each consequence. For example, the likelihood of the NGCC operations team taking longer than one day to regain control of the transmission grid and end the blackout should be low. However, if the attack causes significant physical damage to remote switchgear equipment and the NGCC data centre fibre links are all cut, then the likelihood would increase.

Figure 13 - Formal Implications Wheel™ example showing layered structure

The smart group should include people with different perspectives and expertise, as the chain reaction involves technical, economic and social impacts. Impacts can range from macro level (the entire UK energy grid ecosystem) down to small localised consequences for a specific member entity.

As mentioned previously, when exploring the impacts of attacks on cyber ecosystems, the impact on both tangible and intangible assets needs to be included, as illustrated in Figure 14.

Cyber attacks can impact on intangible assets but are less likely to impact on physical assets, such as plant and machinery. However, this is not the case in the hypothetical example scenario, which has resulted in impact to physical assets and loss of life.

The intangible assets associated with the UK energy grid ecosystem include the brand equity of each of the participants, and the UK national grid operator in particular. Crucially, intangibles include trust in the UK itself as a suitable location in which to invest and operate businesses. The black swan event hypothesised in the example is therefore significant, since it demonstrates how such trust could be damaged, and how the illustrative multi-vector cyber attack could pose a systemic risk.

A whole spectrum of cyber attacks would need to be modelled in this way to help discover some of the many emergent systemic risks associated with the complex system of systems that forms the UK energy grid ecosystem.

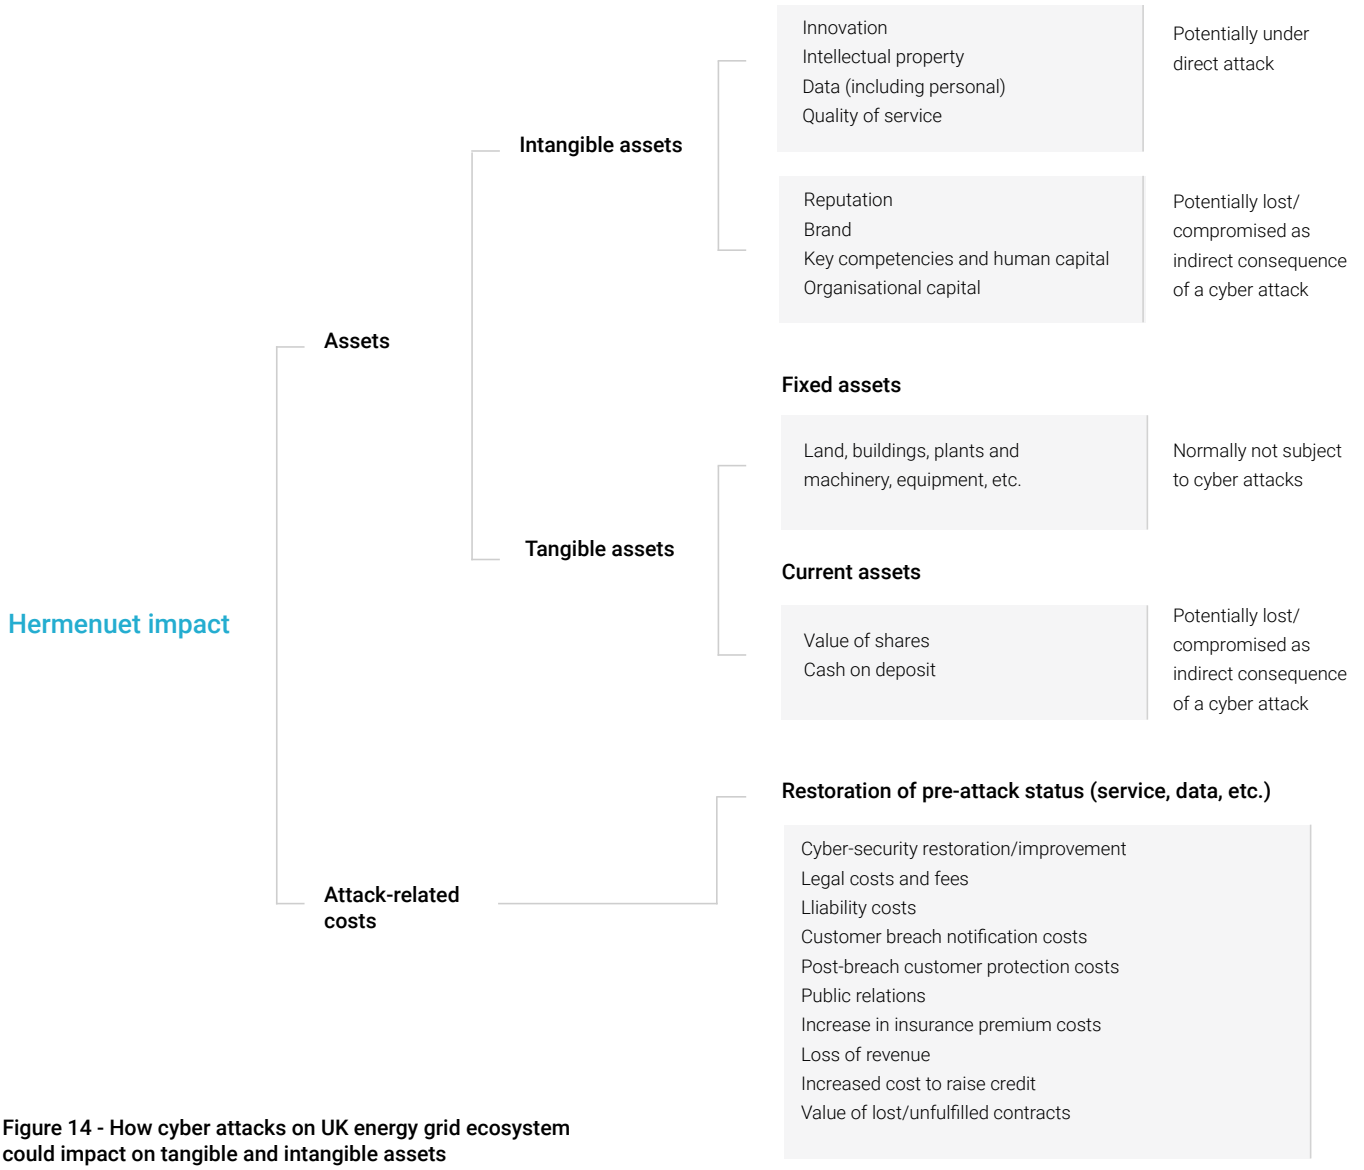

Figure 14 - How cyber attacks on UK energy grid ecosystem could impact on tangible and intangible assets

Conclusion

Find out more about the BHI and the Hermeneut project

This report shows how to apply the BHI to CNI cyber ecosystems, and uses a hypothetical cyber attack scenario to illustrate the process. The formal application of the BHI to CNI cyber ecosystems would uncover potentially significant emergent threats in advance of exploitation by hostile nation state actors and their proxies, as well as threat actors such as terrorists.

This report has shown how methodologies such as the Implications Wheel™[9] can be used to help discover systemic risks in this context.

Appendix A provides an example risk mitigation for complex CNI cyber ecosystems, based on a state-of-the-art cyber threat information platform produced by C3ISP[10] which, like Hermeneut, is an EU Horizon 2020 project.

Digital Catapult welcomes further discussion with CNI stakeholders on the potential benefits of such projects.

The BHI approach is described in full technical detail in the EU Hermeneut project document: D4.2 BHI (Benefit Harm Index) Report. This is available on the Hermeneut site at the following link: www.hermeneut.eu/resources/

Hermeneut's cybersecurity cost-benefit approach to risk assessment combines integrated assessment of vulnerabilities and their likelihoods with an innovative macro- and micro-economic model for intangible costs, delivering a quantitative estimation of the risks for individual organisations or a business sector, and investment guidelines for mitigation measures. Learn more about the wider Hermeneut project here: www.hermeneut.eu/about/

Glossary

ATM	Automated teller machine
BHI	Business harm index
FCA	Financial conduct authority
FMI	Financial markets infastructure
FSMA	Financial services and markets act
CNI	Critical national infrastructure
CREST	Certificateless registry for electronic share
GNSS	Global navigation satellite system
NCA	National crime agency
NGSC	National cyber security centre
PESTL	Political, economic, social, technical, legal

References

Endnotes

1. Countryeconomy.com **https://countryeconomy.com/gdp/uk**
2. C3ISP Collaborative and Confidential Information Sharing for Cyber Protection. **https://c3isp.eu/**
3. Identifying and Scoring Vulnerability in SCADA Environments https://www.researchgate.net/publication/318430054_Identifying_and_Scoring_Vulnerability_in_SCADA_Environments
4. Energy in the UK 2018 https://www.energy-uk.org.uk/files/docs/Research%20and%20reports/Energy_in_the_UK/EnergyintheUK2018finalweb.pdf
5. EY – Smart Grid: a race worth winning – a report on the economic benefits of the UK smart grid. **http://www.ourenergypolicy.org/wp-content/uploads/2016/03/EY-Smart-Grid-a-race-worth-winning.pdf**
6. Rogers, E.M. 1962 Diffusion of Innovations. New York: The Free Press
7. Identifying and Scoring Vulnerability in SCADA Environments **https://www.researchgate.net/publication/318430054_Identifying_and_Scoring_Vulnerability_in_SCADA_Environments**
8. EY – Smart Grid: a race worth winning – a report on the economic benefits of the UK smart grid. **http://www.ourenergypolicy.org/wp-content/uploads/2016/03/EY-Smart-Grid-a-race-worth-winning.pdf**
9. Countryeconomy.com **https://countryeconomy.com/gdp/uk**
10. An introduction to the Implications Wheel. **http://orgs.gustavus.edu/ric/pdfs/Introduction%20to%20the%20Implications%20Wheel.pdf**
11. C3ISP Collaborative and Confidential Information Sharing for Cyber Protection. https://c3isp.eu/
12. C3ISP Collaborative and Confidential Information Sharing for Cyber Protection. https://c3isp.eu/

Appendix

A: Mitigating emergent risk by sharing cyber threat infromation

In the complex VL4 systems of systems, the scale and dynamic nature of the threat landscape, coupled with the motivation of threat actors to focus on cyber ecosystems that provide critical national infrastructure, means that attacks will occur - and some are likely to be successful.

Increasing the information available for analysis allows better prediction, prevention and mitigation of cyber attacks, therefore, sharing cyber threat information across an ecosystem is likely to help make cyber security more effective. This requires cooperation and collaboration between all the entities involved.

Figure 10 highlights the concept of sharing cyber threat information (CTI) across the UK energy grid ecosystem. This could be facilitated using a cloud-based service, such as that proposed by the EU C3ISP research project[11].

The US National Institute of Standards and Technology (NIST) defines CTI as any information that can be used to identify, assess, monitor and respond to cyber threats. In order to be effective, any CTI sharing service needs to address the usual constraints and inhibitors that organisations face when sharing data, such as:

- Restrictions on the type of CTI to be shared
- The circumstances under which sharing CTI is acceptable
- Restrictions on parties with whom the CTI can be shared

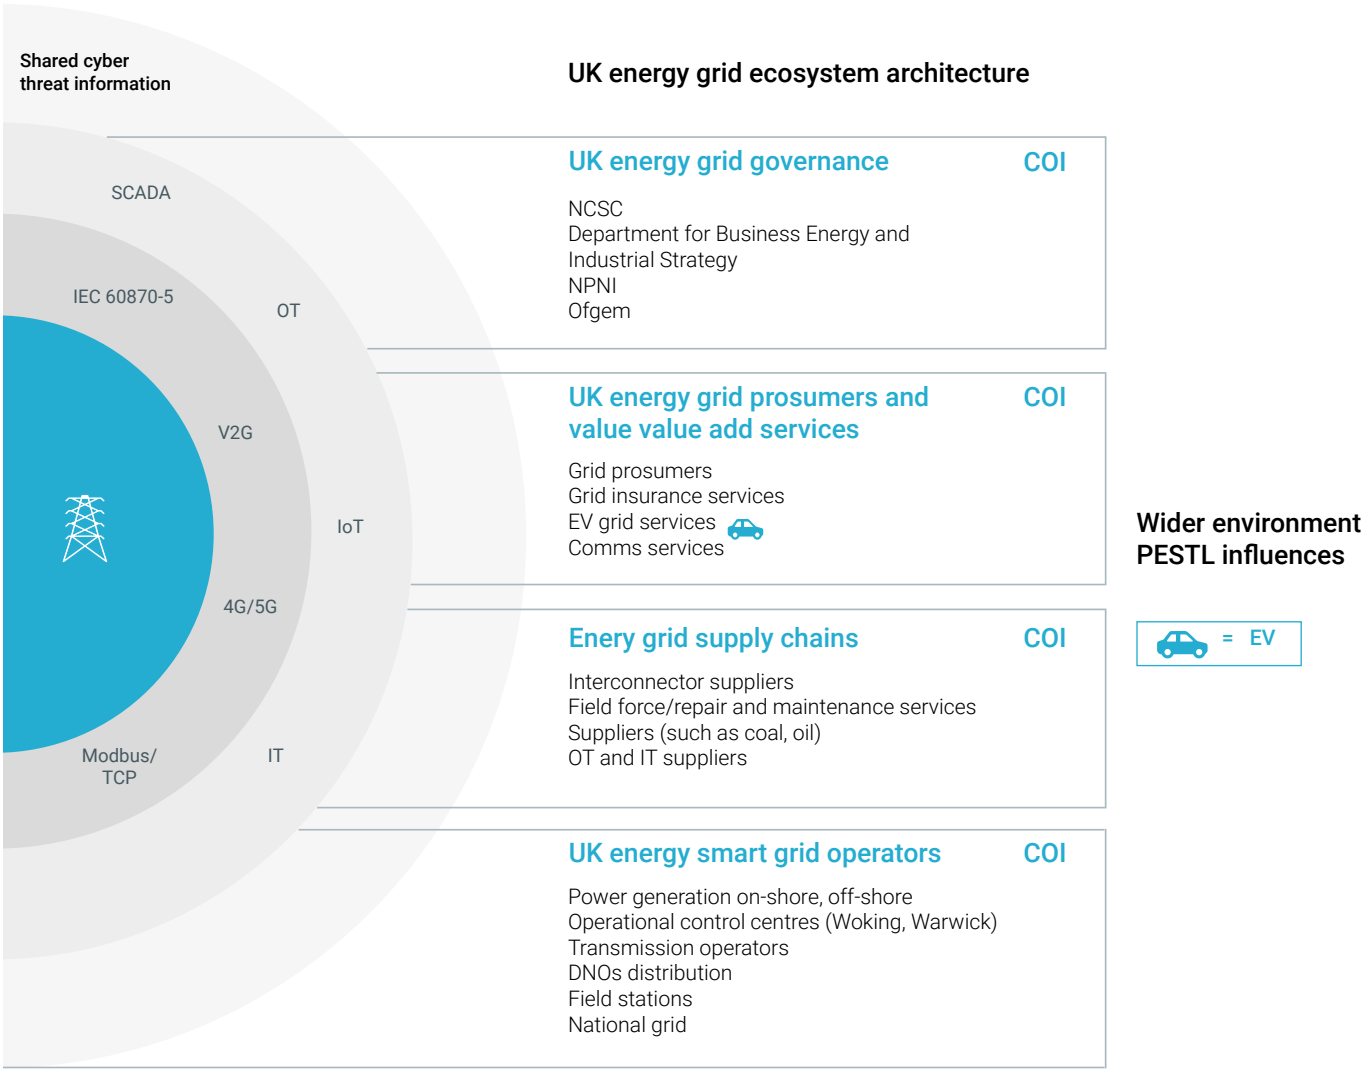

Figure 10 - The UK energy grid ecosystem – cyber threat mitigation through sharing CTI

For example, sharing CTI including details of a data breach or personal identifiable data needs to be managed to ensure regulatory compliance, and may require anonymisation or homomorphic encryption to ensure confidentiality. An ecosystem CTI sharing service that is viable is therefore one where:

- Member entities can choose the type of confidentiality controls appropriate for safeguarding their CTI data, for example, using open access, data anonymisation techniques, or even homomorphic encryption-based techniques
- Data confidentiality and access options enable member entities to confidently share specific types of CTI data, even with non-trusted third parties
- Members can choose and use the techniques most suitable to their needs without needing to be concerned about their design and implementation
- Diverse and transparent techniques for analysing shared CTI can be applied, without member entities needing to be concerned about issues such as information leakage

The main European CTI-sharing initiative, C3ISP, is part of the EU Horizon 2020 project and addresses the concerns raised within this paper. Digital Catapult is a member of this EU research project and describes C3ISP on their website[12] as:

'Providing effective cyber security requires co-operation and collaboration between all the entities involved. Increasing the information available for analysis allows better prediction, prevention and mitigation of cyber-attacks. However, concerns that sensitive and confidential information may be revealed currently deters organisations from sharing data. C3ISP addresses this concern by providing a set of flexible mechanisms, regulated by data sharing agreements, which allow owners to retain control of what is shared, and protect information in the most appropriate way, depending on circumstances. This is aligned with the main guidelines of the European Cyber Security Strategy.'

C3ISP's mission is to define a collaborative and confidential information sharing, analysis and protection framework as a service for cyber security management.

One of C3ISP's components is focused on supporting small and medium-sized businesses in sharing CTI. This is important in domains such as the supply chain, where smaller enterprises typically will not have as strong security capabilities as larger enterprise-level entities within the UK energy grid ecosystem.

Data analysis outcomes

- Early detection of attacks, based on pre-existing knowledge
- Distribution of best practices to avoid vulnerability exploitation
- Discovery of patterns for cyberattacks targeting SMEs

Security issue

- Risk of tampering SME reputation
- Risk of sharing privacy sensitive information
- Disclosure of private files
- Third party is not trusted

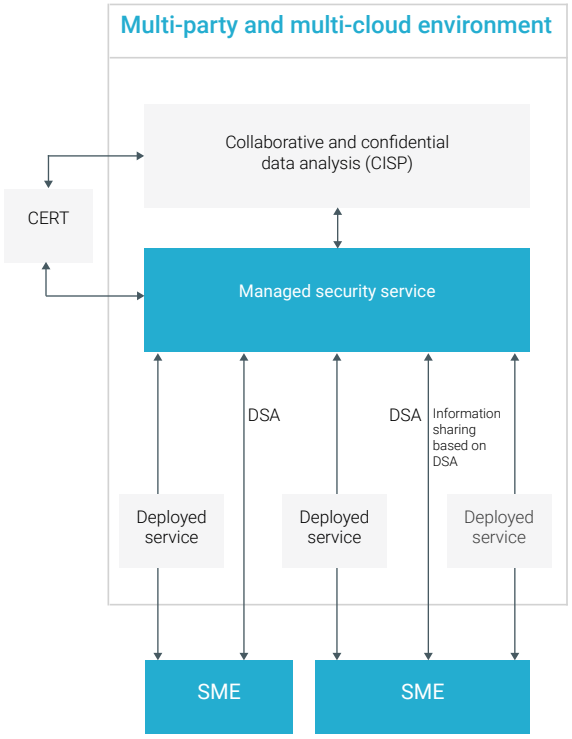

Figure 11 - C3ISP pilot for CTI sharing among small and medium-sized businesses (with reference to the C3ISP EU project)

The C3ISP project describes the objectives of their pilot as being to:

- Deploy the SME pilot, providing a secure multi-party cloud environment for collaborative information sharing, performing collection and analysis of small to medium-sized business data without disclosing privacy-sensitive information
- Use this prototype platform to evaluate and validate the C3ISP approach, architecture and technology in the context of a managed security analytics service provided to small to medium-sized businesses

- Evaluate the capability of providing security intelligence obtained through collaborative analysis

- Evaluable the capability of delivering this intelligence without disclosing private information while complying with compliance DSA policies

The sharing of CTI data in a way that enables participants to retain control is fully supported by C3ISP, which also exploits OASIS standards such as STIX (structured threat information expression) and TAXII (trusted automated exchange of indicator information) to support inter-operable automated exchanges of CTI.

The C3ISP architecture also supports a shared platform where small, medium and enterprise level participants, ISPs and CERTS can collaborate. This can be used to support the proposed model for an ecosystem-level shared CTI capability for each UK CNI ecosystem, including the energy grid.

The C3ISP platform provides CTI analytics services (such as that provided by BT Saturn) that show a real-time visualisation of the threat landscape, and active and historic cyber attack vectors across the ecosystem. These visualisation services can be provided in 3D immersive VR mode, so that ecosystem cyber security analysts can better comprehend any current threat in an operational context, such as within the UK energy grid ecosystem.

CATAPULT Digital

Digital Catapult is the UK's leading advanced digital technology innovation centre, driving early adoption of technologies to make UK businesses more competitive and productive to grow the country's economy.

We connect large established companies, startup and scaleup businesses and researchers to discover new ways to solve big challenges in the manufacturing and creative industries. Through this collaboration businesses are supported to develop the right technologies to solve problems, increase productivity and open up new markets faster.

Digital Catapult provides physical and digital facilities for experimentation and testing that would otherwise not be accessible for smaller companies.

As well as breaking down barriers to technology adoption for startups and scaleup, our work de-risks innovation for large enterprises and uncovers new commercial applications in immersive, future networks, and artificial intelligence technologies.

For more info please visit **digicatapult.org.uk**

Photography Markus Spiske