CATAPULT Digital | A new perspective on Cyber Risk

Applying the benefit harm index, 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

With an increasingly complex and interconnected world, the perception of risk needs to be reconsidered as it applies to cyber threats. Traditional risk models rely heavily on probabilistic approaches, which demand stable distributions and almost complete knowledge of 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. Digital Catapult[1], as part of the Hermeneut project, has proposed a new approach – the benefit harm index (BHI). The index which integrates ideas from economics and complexity science into a new approach to understanding dynamic and emergent threats. The Hermeneut project is a part of the European Community's Horizon 2020 programme.

In this white paper Digital Catapult shows how this new perspective on cyber risk can be applied to the cyber ecosystems that form many of today's critical national infrastructures (CNI). This report illustrates how these complex systems of systems exhibit emergent behaviour and require a new approach to cyber risk assessment.

This report uses the example of the UK connected autonomous vehicle (CAV) ecosystem to bring this to life. This complex system is highly interconnected, tightly coupled and provides a detailed example of how disruptions in one area can cascade easily and in unexpected ways that form systemic cyber risks to the UK economy.

This report focuses on the evolution of this example CNI ecosystem over time and uses the BHI approach to illustrate how the rate of growth of socio-economic benefits over time can be overtaken by the rate of growth of harm associated with such systemic cyber risks.

Finally, the report introduces example approaches to mitigating emergent risks in these CNI ecosystems.

CONTENTS

INTRODUCING THE BHI – A NEW PERSPECTIVE TO CYBER RISK

The BHI modelling methodology is designed to provide new insights into the potential risks associated with the cyber ecosystems that underpin the complex and dynamic markets 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 is evolving into something new that is not fully predetermined. In the complex biological world a single virus can mutate, evolve and spread through bio ecosystems and could result in a global pandemic infecting significant numbers of the human population. The ability for such microscopic changes in the bio ecosystem to propagate rapidly and create macroscopic effects (which can be positive or negative) highlights the uncertainty associated with such complex systems.

Since cyber ecosystems are complex dynamic environments, they evolve rapidly and feature high levels of uncertainty. These ecosystems can generate similar emergent behaviours that can often not be predicted by studying the way in which the constituent parts interact. The emergent behaviour is seen to be manifesting in many forms, including the murmurations of birds in the biosphere and in the emergence of new sociopolitical collective behaviours through the use of social media in cyber space.

Attempting to apply traditional risk assessment methodologies to cyber ecosystems will typically involve pretence of knowledge of all the risks. Traditional risk assessment methodologies assume a complete knowledge of all possible states of the system being assessed and that a mathematical likelihood can be applied to each event. Such an approach to risk does not address the complex dynamics and the associated uncertainties of cyber ecosystems.

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

Modelling the dynamic complexity provides a perspective for exploring the rate of growth of the 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 they 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.

The event driven scenario approach enables the exploration into the implications of cyber chain reactions to help identify hidden risks (and benefits) using tools such as the implication wheel. This helps with mitigating the fact that in complex dynamic systems, all the risks are unknown, some of which are emergent and may be very significant.

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

USING BHI TO MITIGATE TO EMERGENT THREATS

Cyber ecosystems are complex and such systems exhibit emergent behaviour. There are different levels of complexity and as the complexity increases different types of emergent behaviour come into play, these being:

- Simple dynamic behaviour (e.g. clock, measuring time)
- Weak emergent behaviour (e.g. flocking of birds/drones)
- Strong emergent behaviour (e.g. bubbles in the financial markets)
- Spooky emergent behaviour (e.g. conscious thought in humans/AI)

The first two are associated with deterministic systems. These types of emergent behaviour can be easily reproduced using simulations of the system. The third and fourth are associated with stochastic (random interactions defined by probability distributions) systems. Stochastic systems can exhibit strong emergent behaviour that cannot be fully reproduced by simulations. Spooky emergent behaviour cannot be reproduced even by detailed simulations of the systems.

If you are governing/operating a cyber ecosystem the extent to which you (as its defender) can control it is intrinsically linked to the level of complexity of that ecosystem. The stability of the system is related to its level of complexity, changes at the micro level can result in a dramatic change at the macro level. An attack on the system can trigger a significant cyber chain reaction which will appear as emergent behaviour.

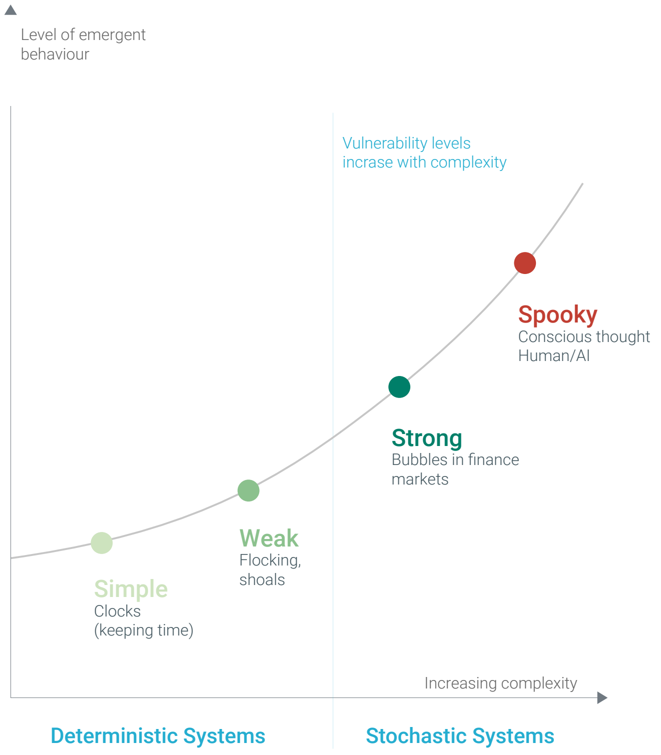

Figure 1 – Complexity and emergent behaviour

In the case of strong and spooky emergence (stochastic systems), the system is fundamentally uncontrollable! In simple terms, the higher the complexity of the ecosystem the more vulnerable it is to emergent threats.

The BHI methodology proposes a taxonomy for the vulnerability level (VL) of a system. This details states of a system in terms of a given scope and phase space (which represents all possible states of the system) with a given resolution, and uses this as a measure of its intrinsic lack of controllability, from the perspective of the defenders, who legitimately operate the system.

As shown in Table 1 below, threats and vulnerabilities to components in the system vary fundamentally by class. Each VL requires radically different types of mitigation.

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

Vulnerability level (VL)	Threat class	Attacker's control
● 5	Emergent system	The nature of the system can show emergent behaviour and cannot be controlled, since the state space of the system changes as emergent properties manifest.
● 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 systems state space.
● 3	Uncontrolled system	The scope of the system or the nature is such that the system is not under control, although it would be possible to control the system in principle.
● 2	Uncontrolled inputs	An attacker uses a legitimate control input in the systems scope, but outside of its expected or normal range.
● 1	Unauthorised activities	An attacker uses a legitimate control inputs within the control system of the system in scope.

Table 1 – Vulnerability levels and their associated class of threats.

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

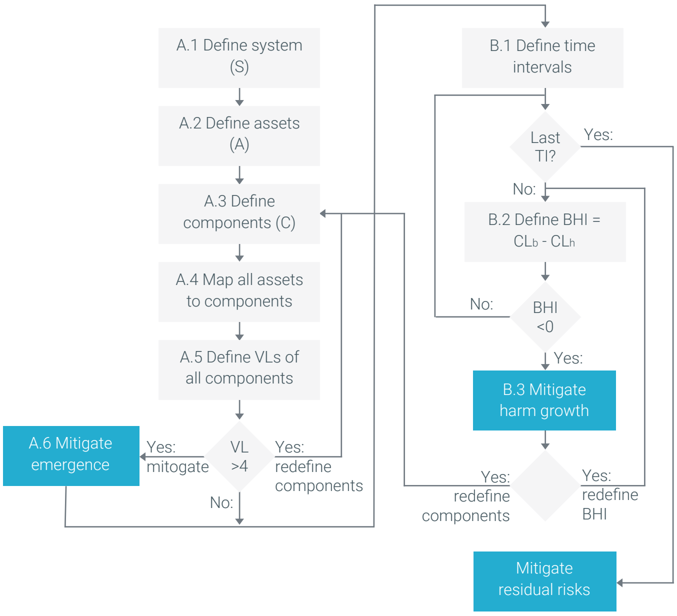

Figure 2 - BHI process for mitigating emergent threats

The first three steps for addressing emergent threats (A.1-A.3) as shown in Figure 2 define:

A.1: The ecosystem being considered.

A.2: The set of assets, whose sensitivity 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 behaviours or activities.

A.3: The set of components, into which the system is decomposed. 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 fourth step defines the association between each asset and the component(s) that directly influence the security of the asset.

A.5: This defines the VL for each component. In the BHI approach there are distinct levels of vulnerability that increase with the complexity of the ecosystem.

This activity is performed considering the nature of a component and its vulnerabilities, as well as the threats from the environment and other components.

If any component has VL value under four, which corresponds to emergent threat, the process takes one of two paths:

- Redefine the components, for example to localise an associated asset in a component that has a lower VL value. This results in reiterating over steps A.3-A.5
- Mitigate emergence (A.6) by designing a set of security controls that seek specifically to mitigate risks from emergence. These controls will need to detect, and potentially isolate and neutralise the impact of an attack

Using BHI, it is expected to distinguish those characteristics that can be localised, from those that cannot. One cannot expect companies to rationally mitigate the latter, 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 minimize impact on the individual organisations adopting such recommendations.

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

USING BHI TO MITIGATE TO THE GROWTH OF HARM

Modelling the dynamic complexity provides a perspective for exploring the rate of growth of the socio-economic benefits generated by an evolving cyber ecosystems over time. It also provides a perspective for exploring the rate of growth of threats to that ecosystem and the associated socio-economic harm they 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 (Figure 3).

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 growth rates. First, each ecosystem will evolve through a number of distinct phase transitions.

For example, the introduction of a new product or class of products - penetrates a market. Initially, there is near exponential growth, which is 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, and the bass diffusion distribution starts to manifest its asymptotic growth complexity level of zero – a constant.

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

TI0
From product introduction 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).

Second, 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:

- The time of events that mark the start and end of relevant changes, such as investment rounds, introduction of new products etc
- The time at which the distribution of growth is likely to be discontinuous, for example as a result of some material event, such as change in a 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 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 that seek to mitigate the growth of harm. In the case that a plausible mitigation is found, the process recomputes the BHI value and iterate to the next time interval.

In some cases, for example, where an effective mitigation cannot be found, it may be considered appropriate to redefine the components.

In this case, the process returns to the right-hand side of the diagram at step (A.3).

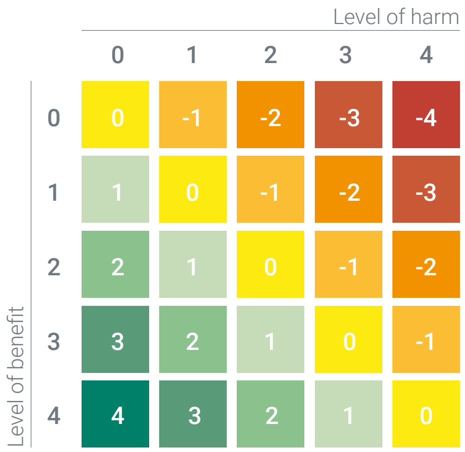

Level of benefit \ Level of harm	0	1	2	3	4
0	0	-1	-2	-3	-4
1	1	0	-1	-2	-3
2	2	1	0	-1	-2
3	3	2	1	0	-1
4	4	3	2	1	0

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

For BHI over zero systemic (ecosystem) level mitigations are required.

Once all members of CI have been processed, the mitigation of risks from growth are 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, each one of which 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. This wider environment will generate political, economic, social, technological, environmental and legal (PESTL) influences on the operation and growth of that ecosystem.

In the approach 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 for example the operational systems within the ecosystem and the supply chain systems that support the manufacture and production of the components that eventually populate that operational systems 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 that contains the governance systems and regulatory frameworks that are used to set and police the policies rules and standards associated with governing the cyber ecosystem
- The value added services domain that includes the systems and processes associated with services that add value to the operational services, for example insurance services

The cyber system domains will all have vulnerabilities. Threats to the ecosystem will exploit these vulnerabilities through attack vectors originating from threats sources (for example hostile states) attacking via threat actors (external and internal), as illustrated schematically in Figure 4 above. The BHI approach exploits methodologies such as the implications wheel to investigate the VL of components in such complex systems and the potential for cyber chain reactions being generated through multiple iterations. Targeted scenario analysis is used in this context to help identify such events through systematically exploring the implications of interaction/contagion through multiple first, second and nth order interaction flows.

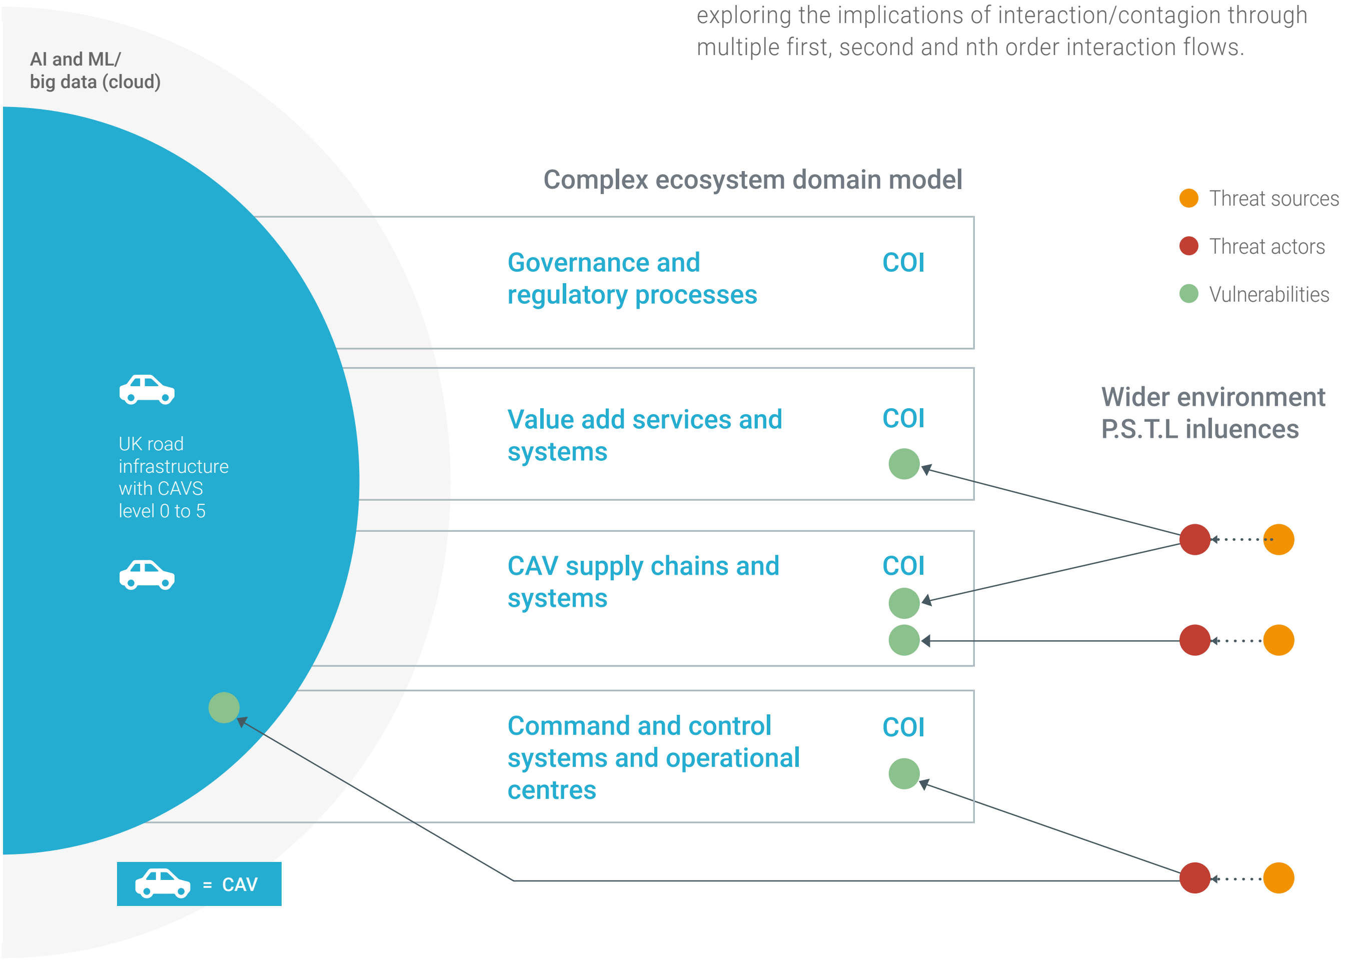
AI and ML/
big data (cloud)
Complex ecosystem domain model
Threat sources
Threat actors
Vulnerabilities
Governance and regulatory processes
COI
Value add services and systems
COI
CAV supply chains and systems
COI
Command and control systems and operational centres
COI
Wider environment P.S.T.L inluences
UK road infrastructure with CAVS level 0 to 5
= CAV

Figure 4 - Cyber ecosystem high level domain model

The BHI dynamic approach to risks also enables the construction of multiple phase states for each cyber ecosystem model to reflect the 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.

APPLYING THE BHI TO THE UK CAV ECOSYSTEM

The UK's transport system is one of the UK's 13 critical national infrastructure (CNI) components. In this paper are focused on the UK road network and the connected and autonomous vehicle (CAV) transportation services that it supports.

The UK road transport system is evolving and is moving strategically towards being an intelligent transportation system that supports a hybrid population of vehicles with varying degrees of autonomy, ranging from full human control to fully autonomous artificial intelligence (AI) control.

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

The CAVs operate within the context of the UK intelligent road infrastructure, which itself is an integral part of the UK CAV ecosystem.

In simple terms the intelligent road infrastructure can be thought of as being comprised of a number of different sub domain types such as smart motorways, country roads and metropolitan (urban cores).

These subdomains will evolve over time to support, accelerate and exploit the higher levels of CAV autonomy. Such sub domains can be designed to try and reduce the hybrid mix of low level legacy CAVs with next generation CAVs as higher levels of autonomy enter the UK CAV ecosystem. Adaptive computer operated speed and traffic control systems are examples of what is meant by an intelligent road infrastructure, such control systems will of course be more prevalent in sub domains such as smart motorways and urban cores.

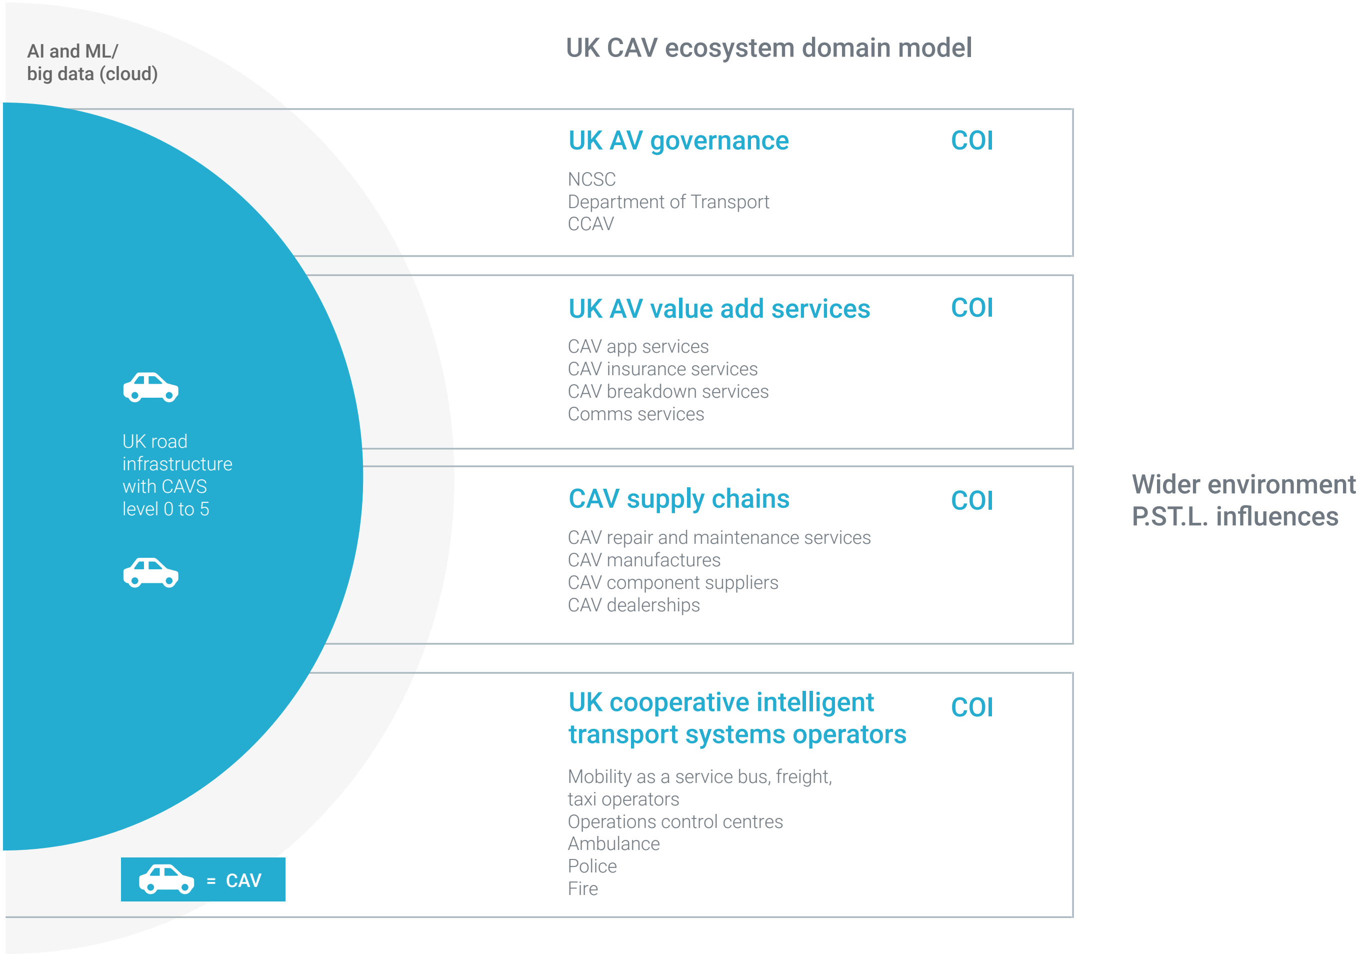
AI and ML/
big data (cloud)
UK CAV ecosystem domain model
UK road infrastructure with CAVS level 0 to 5
🚗 = CAV
Wider environment P.ST.L. influences

UK AV governance — COI
NCSC
Department of Transport
CCAV

UK AV value add services — COI
CAV app services
CAV insurance services
CAV breakdown services
Comms services

CAV supply chains — COI
CAV repair and maintenance services
CAV manufactures
CAV component suppliers
CAV dealerships

UK cooperative intelligent transport systems operators — COI
Mobility as a service bus, freight, taxi operators
Operations control centres
Ambulance
Police
Fire

Figure 5 – UK CAV ecosystem domain model

In Figure 5, in line with approach to modelling each cyber ecosystem, the UK CAV ecosystem features the following domains:

- The CAV supply chain community of interest. This is the domain that includes the entire end to end global supply chain of each CAV manufacturer. This includes all the CAV manufactures that produce CAVs that can operate on the UK road transport system. This in the absence of any regulatory constraints includes most major CAV manufacturers
- The UK cooperative intelligent transport systems operators/ operations centres community of interest. This is the domain that includes all the organisations that actually operate the smart road infrastructures and those that operate the road transportation services such as traffic control, emergency services and mobility as a service including operating fleets of autonomous taxis, buses and lorries. It includes the many associated operations centres, for example, the National Traffic Control Centre (NTCC) and the London Streets Traffic Control Centre (LSTCC)
- The UK CAV value added services community of interest. This domain includes all those organisations that provide value added services such as motor insurance, breakdown recovery and remote vehicles monitoring and diagnostics telematics services
- The UK CAV governance community of interest. This domain includes all those UK Government organisations that are responsible for the regulatory, safety, security and legal policies and governance of the UK road transport infrastructure. The Department of Transport plays a key role here together with its 22 agencies in particular the Centre for Connected & Autonomous Vehicles (CCAV)[3]. The NCSC is also makes a key contribution to this domain as does the CPNI

In defining the 'UK CAV ecosystem', it makes sense to adopt the UK Parliamentary Office of Science and Technology definitions of the five levels of CAV autonomy. These are defined in Table 2 below:

Level	Name	Description
0	No automation	Human driver completely controls the vehicle.
1	Driver assistance	Individual activities which assist steering or acceleration/ deceleration are partially automated.
2	Driver assistance	Several simultaneous activities which assist steering or acceleration/ deceleration are partially automated.
3	Conditional automation	In certain driving scenarios, all dynamic, non-strategic, driving activities (for example vehicle control but not route choice) are automated but human is expected to intervene when requested.
4	High automation	In certain driving scenarios, all dynamic driving activities are automated and vehicle can cope with human not intervening if and when requested.
5	Full automation	Always and everywhere, all dynamic driving activities are automated with no need for human intervention.

Table 2 – UK levels of CAV autonomy

Over time the hybrid mixture of CAVs at different levels of autonomy will change from the current state of almost, all being at levels (zero or one) to the majority being at (level three, four and five). As can be seen from the definitions of level three, four and five autonomy in Table 1 above, such CAVs can operate in different modes for example they can operate autonomously or under human control, dependant on circumstances.

A SIMPLE HIGH LEVEL ARCHITECTURE OF THE UK CAV ECOSYSTEM

In order to explore this report's new perspective on cyber risk in this context, a simple high level architecture model of the UK CAV ecosystem is produced, which enables the exploration of cyber attack scenarios at two different time periods during the UK operational CAV market evolution. The BHI is very different in these two cases.

Figure 6 illustrates a simple high level architecture model focused on the operational CAVs in the context of the UK intelligent road infrastructure. The CAV population in this simple model is a hybrid-mix of different autonomy levels (zero, to five), the percentage of each CAV level being a function of time with the higher level autonomy population increasing.

Each CAV at level 1 and above will feature varying degrees of vehicle to vehicle (V2V) and more generally vehicle to everything (V2X) modes of communication. V2V will either use dedicated short range communication (DSRC) or C-V2X depending on which of these two competing and incompatible protocols are selected. This paper assumes the adoption of 5G C-V2X for the UK CAV ecosystem, with the first significant 5G networks predicted to start going live in the UK by 2020.

A CAV can for example use V2V 5G to communicate breaking information electronically to other CAVs in its vicinity. V2V is important for supporting CAV collision avoidance. Other scenarios that V2X 5G can be used for include enabling remote control driving systems, where a real time 360 degree view of a CAVs live road traffic environment is transmitted to a remote control room via on board cameras and sensors. This enables the CAV to be driven under remote control from a remote location/country[4].

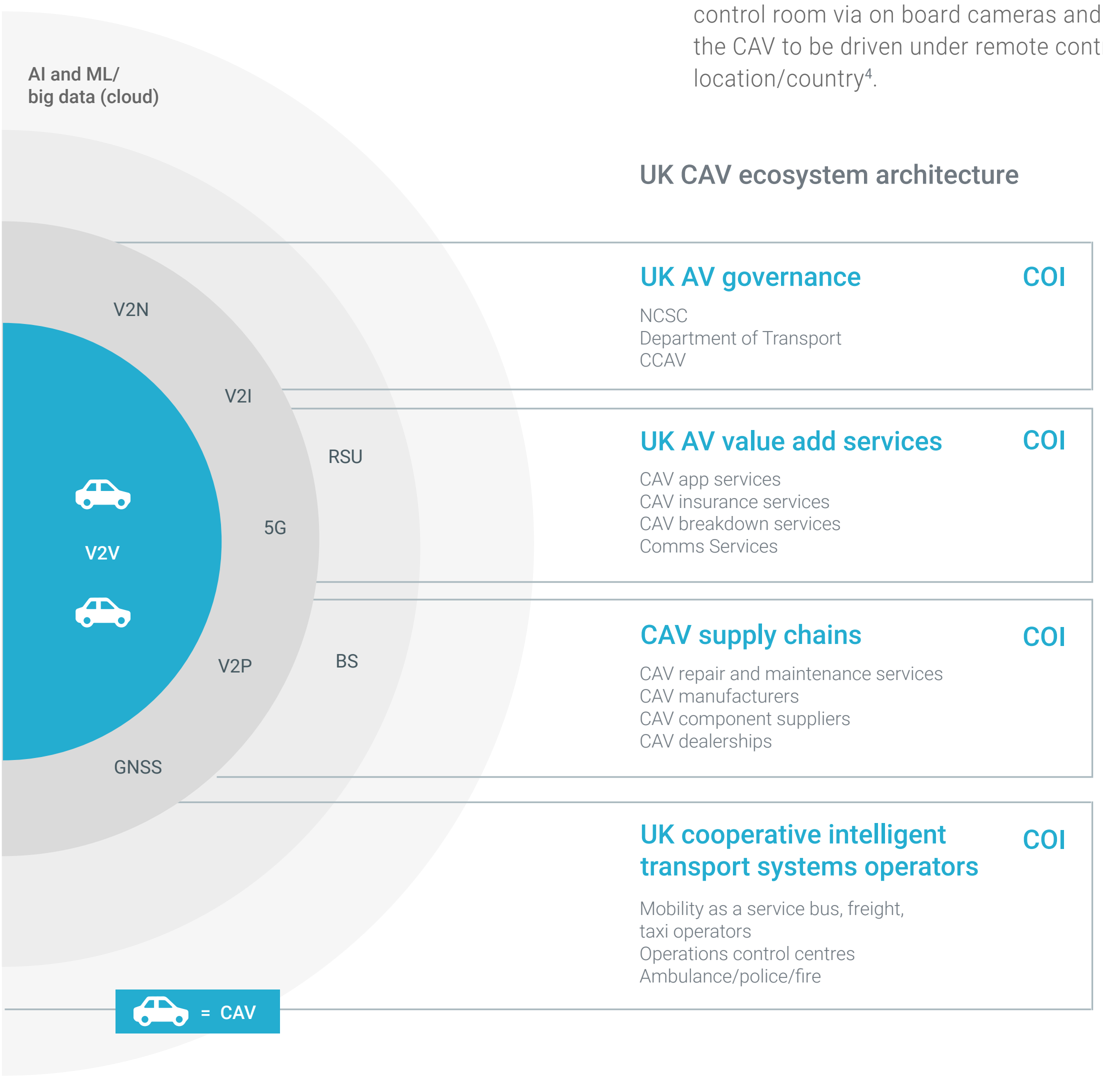

UK CAV ecosystem architecture

UK AV governance — COI

NCSC
Department of Transport
CCAV

UK AV value add services — COI

CAV app services
CAV insurance services
CAV breakdown services
Comms Services

CAV supply chains — COI

CAV repair and maintenance services
CAV manufacturers
CAV component suppliers
CAV dealerships

UK cooperative intelligent transport systems operators — COI

Mobility as a service bus, freight, taxi operators
Operations control centres
Ambulance/police/fire

Figure 6 – UK CAV ecosystem high level architecture

The intelligent road infrastructure itself will include road side units (RSU) and 5G base stations. These RSU's will in effect form a low latency edge computing environment (or FOG) that can intermediate between the CAV's and the cloud systems that will host the aggregated big data lakes associated with the overall UK CAV ecosystem. There will also be other cellular modes such as 4G especially outside of urban cores.

The RSU's will be found in greater numbers in urban cores and smart motorways than in rural areas. They provide important traffic information to CAV's and can relay CAV information to ITS operations centres via cloud services and help to both monitor and control traffic flow in these core environments.

The ITS cloud will host big data lakes of information containing data on individual CAV's, their past and current journeys, owners and passengers as well as vast amounts of related data including aggregated traffic flow and emissions data. AI and machine learning using deep neural networks will be used in this context for example to predict traffic congestion in these ITS core areas.

The CAV's have onboard units (OBU's) that support transmission and reception of V2X communications such as 5G and the receipt of global navigation satellite system (GNSS) geo-location data. The high level architecture view provides a simple logical model of the CAVs themselves as shown below in Figure 7.

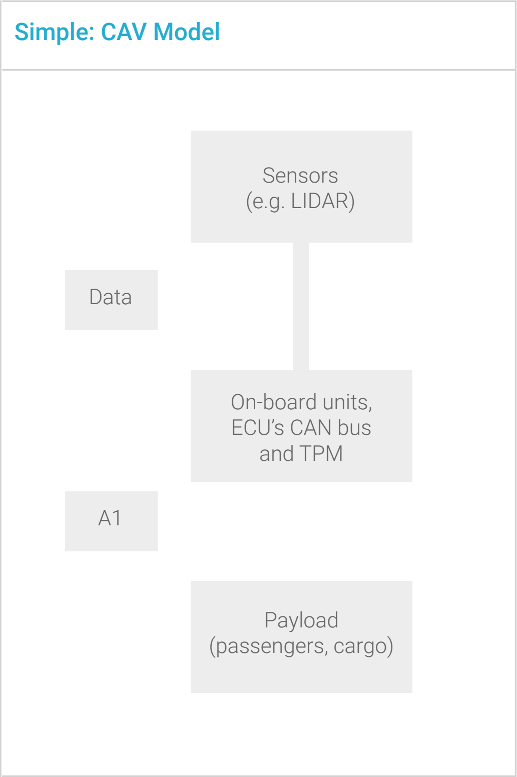
Simple: CAV Model
Sensors (e.g. LIDAR)
Data
On-board units, ECU's CAN bus and TPM
A1
Payload (passengers, cargo)

Figure 7 – Simple logical model of a CAV

This simple CAV model highlights the basic class of components that are interesting at this level. The sensor network includes a range of potential types of sensor including cameras, radar and light detection and ranging (LIDAR) sensors. Each CAV has a large number of electronic control units (ECUs) distributed over an on-board network bus that perform different functions including some that perform safety critical driving control functions as part of the control area network. The data in this model includes vehicle data and potentially personal data related to passengers.

These control functions will be under human control at lower CAV autonomy levels and under the control of AI for example machine learning typically using deep neural networks at the higher levels. The AI control function can be viewed simplistically as an observe, orientate, decide and act (OODA) loop. The human is in loop at lower levels of autonomy, on the loop (for example can interject if needed to take back control) at level three and four and off the loop at level five (for example full AI control). The trusted platform module (TPM) provides cryptographic security functions such as encryption, decryption, signing, and verification which can be used to help authenticate and secure CAV over the air software/data updates.

The final component depicted in Figure 7 is the payload for example the passengers and or cargo (for example goods, fuel) being transported by the CAV at given point in time. A level five CAV may travel on the road network for periods of time with no humans on board. A level five CAV could in theory also carry cargo that is valuable and or dangerous such as petrol tanker without any humans on board.

EVOLUTIONARY STATES OF THE UK CAV ECOSYSTEM

This white paper applies the BHI to two different time periods during the UK operational CAV market evolution. This considers predicted future states of the UK CAV cyber ecosystem (using the simple high level architecture model) for the period 2025 – 2030 as the first period and in its predicted future state in the period 2030 to 2035 as the second period. These will be based on predicted future states on market forecasts for the UK CAV population[5] and 5G C-V2X and associated technologies. Included at each stage will be the UK socio-economic benefits predicted as being generated by the UK CAV ecosystem in these time frames.

In Figure 8 below the illustrative growth rate prediction for the population of level four and five CAVs on the UK intelligent road infrastructure between 2020 and 2035. The report assumes a bass diffusion for this illustrative prediction. It has been assumed a high uptake scenario for level four and five. It is recognised that there will be a range of possible growth rate curves however, this one will simply be used as the context for the BHI illustration.

The bass diffusion model[6] is a contribution to the understanding how new products are introduced and become more widely adopted by customers. Initially the concept of diffusion of innovations was developed by Everett Rogers in 1962[7]. Frank Bass published his paper on new product growth a few years later, proposing a more mathematical formulation of the dynamics of new product adoption. The model presents a rationale of how current adopters and potential adopters of a new product interact. The basic premise of the model is that adopters can be classified as innovators or as imitators and the speed and timing of adoption depends on their degree of innovativeness and the degree of imitation among adopters.

The report now looks at the two evolutionary states (2025-2030) and (2030-2035).

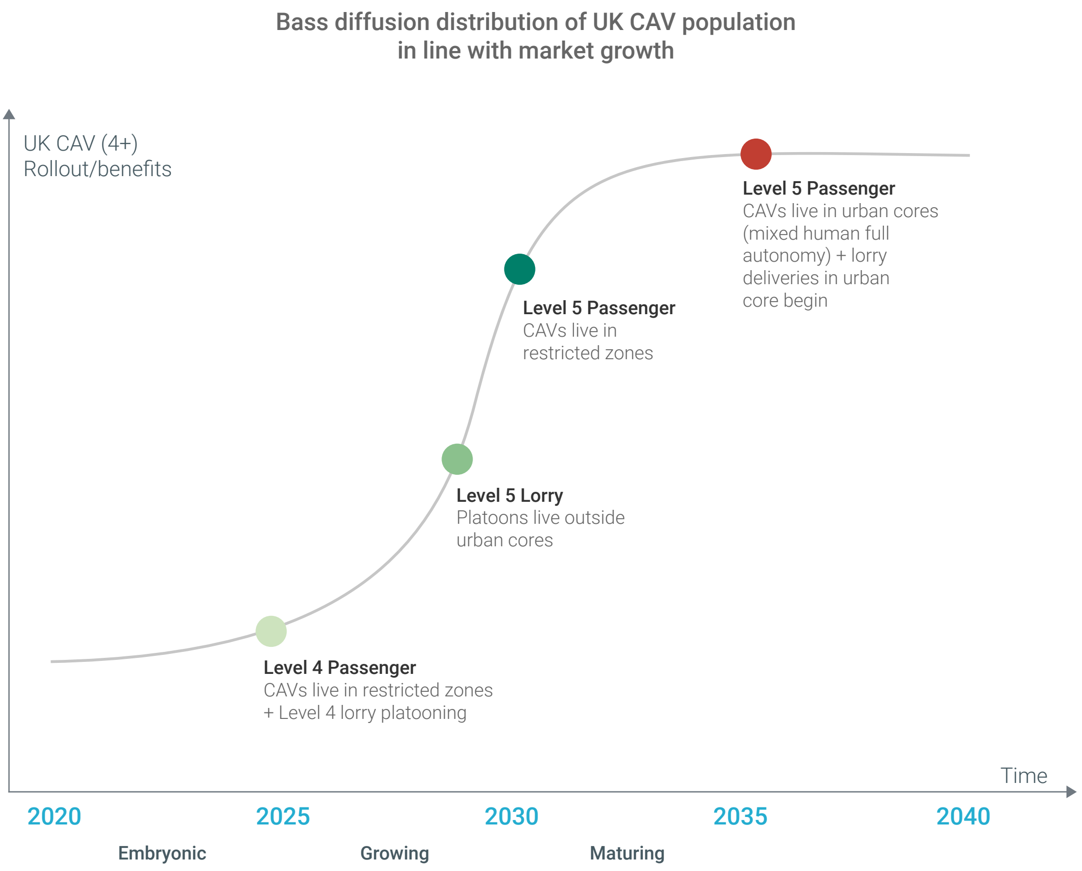

Figure 8 – Illustrative prediction of the UK CAV population rollout

EVOLUTIONARY STATE 2025-2030 OF THE UK CAV ECOSYSTEM

The UK CAV ecosystem in this period is modelled on the UK CAV ecosystem high level architecture that was introduced earlier in section 2.1 of this paper. There will now be a focus on a particular time interval for this ecosystem that being the period between 2025 and 2030.

As was shown in Figure 8 the rate of growth of the population of level four and five CAVs is growing exponentially and corresponds to the period of that product introduction on the UK road infrastructure.

Core ITS domain

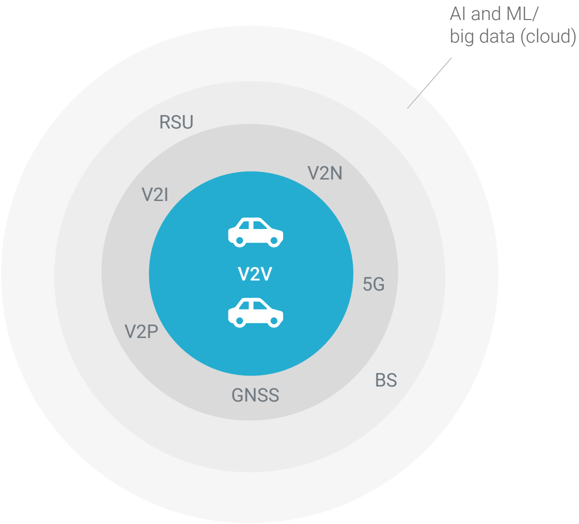

During this period the level four and five CAVs that have been undergoing trials in trial environments will start to go live on the UK road infrastructure.

In this model 0.4% of new car sales are at level four and level five CAVs in 2025 rising to 8% by 2030. The intelligent road infrastructure during this state of the model features 5G V2X inside urban cores and sections of some smart motorways.This is in addition to other cellular services such as 4G and GNSS.

In bass diffusion model (Figure 8) the growth rate of the on road level four - five CAV population will be exponential during this period. The model assumes platoons of level four and eventually level five lorries outside of urban cores feature in this period. The model in this state also features fully autonomous (level five) passenger cars including autonomous taxi and bus services operated in restricted zones within urban cores.

The cloud based big data lakes will provide the intelligent transport service (ITS) operations centres with growing amounts of CAV vehicle data and journey data as well as direct and indirect forms of CAV driver/passenger personal data, during this period. There will be different CAV priorities, for example emergency vehicle CAVs will have higher priority than private CAVs of the same autonomy level.

ITS operations control domain

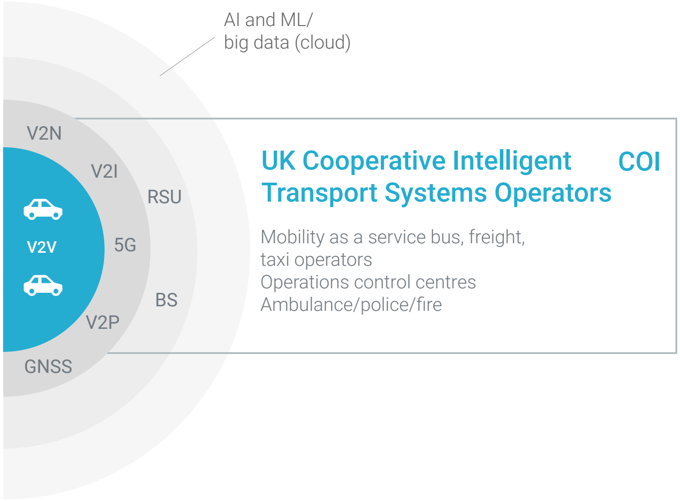
AI and ML/
big data (cloud)
V2N
V2I
RSU
V2V
5G
BS
V2P
GNSS
UK Cooperative Intelligent Transport Systems Operators
COI
Mobility as a service bus, freight, taxi operators
Operations control centres
Ambulance/police/fire

Also during this model state, there is significant cooperation between the various UK ITS operations centres including emergency services, NTCC, city transport controllers and to a lesser extent private CAV fleet-operators.

The model assumes multiple commercial fleet operators with their own private data lakes and limited sharing of anonymised data.

During this period the UK CAV ITO centres domain will start to include private operations centres that operate new business ownership models such as mobility as a service (MaaS). These will feature live CAV five MaaS for restricted routes within zones in urban cores.

The emergency services and national and city transport operations centres will also feature procedures and capabilities to deal with CAV level four and five incidents in place during this period. This is in addition to smart traffic congestion, emissions and parking management capabilities using the ITS data cloud services.

UK CAV Governance Domain

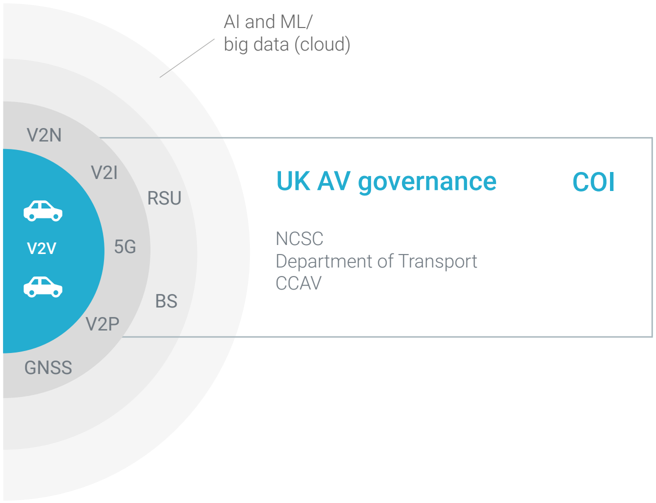

The UK governance domain of the CAV ecosystem is assumed in this evolutionary model state to have put in place a set of regulatory guidance and controls to support the live operations of level four and five CAVs on the UK intelligent road infrastructure.

The governance domain of the UK CAV cyber ecosystem is responsible for regulatory policy and guidance and that includes driving the policy on cyber security for the overall UK CAV ecosystem.

The Department for Transport (DfT), in conjunction with the Centre for the Protection of National Infrastructure (CPNI), have already created the following key principles for use throughout the automotive sector, the CAV and the ITS domain of the ecosystem and their supply chains.

Principle 1
Organisational security is owned, governed and promoted at board level.

Principle 2
Security risks are assessed and managed appropriately and proportionately, including those specific to the supply chain.

Principle 3
Organisations need product aftercare and incident response to ensure systems are secure over their lifetime.

Principle 4
All organisations, including sub-contractors, suppliers and potential third parties, work together to enhance the security of the system.

Principle 5
Systems are designed using a defence-in-depth approach.

Principle 6
The security of all software is managed throughout its lifetime.

Principle 7
The storage and transmission of data is secure and can be controlled.

Principle 8
The system is designed to be resilient to attacks and respond appropriately when its defences or sensors fail.

However by the 2025 - 2030 period there needs to be a level of compliance determined for example through an assurance framework such as 5StarS to provide an assurance rating for the cybersecurity of CAVs. It is assumed in the model that as a result of the three year law review currently under way that changes to UK law are in place in this period that include:

- The allocation of civil and criminal responsibility by law where there is shared control between humans and computers
- The role of automated vehicles in public transport, car sharing and on-demand passenger services; any need for new criminal offences
- The impact on other road users and how they can be protected from risk, and determining who the responsible person is in a self-driving vehicle

UK CAV value added services domain

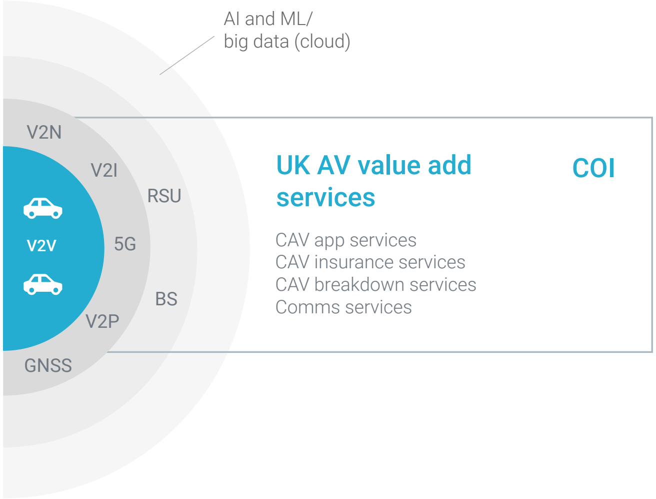

This paper assumes that at this 2025 - 2030 model state of the UK CAV ecosystem that the government has ensured there is an appropriate insurance legal framework in place so that appropriate motor insurance cover is available for CAVs.

The UK CAV value added services domain during this evolutionary state will feature many telematics services intelligent solutions extending existing telematics offerings and provide services that help operators to optimise fleets and reduce their environmental footprint and for insurance providers to optimise their motor insurance services.

UK CAV supply chain domain

The global supply chain domain of the UK CAV ecosystem in this evolutionary state is assumed in the model to be relatively weak in terms of supply chain security. The principle that supply chain security risks are assessed and managed appropriately and proportionality will be difficult to enforce cost effectively in the this time frame given the global scale, complexity and fragmentary structure of the CAV supply chains.

The supply chain risks include those where a nation state actor installs backdoors in the CAVs through microchips on motherboards on components that will be used as part of major brands of CAV used on UK roads.

The supply chain risks to the CAV ecosystem include supply chains associated with the 5G technology including that used for the intelligent road infrastructure. Global risks from potential adversaries are of concern; for example Huawei is a major supplier of broadband equipment and mobile networks in the UK, meaning its products are used in critical national infrastructure which could be targeted.

"Finding such extra chip is something very few companies can do since the level of details of the design and implimentation of the motherboard aren't generally distributed."

APPLYING THE BHI USING AN ILLUSTRATIVE CYBER ATTACK SCENARIO

Having defined the system of interest and identified the time interval in focus, there can now be a demonstration of how the BHI approach can be applied to the UK CAV ecosystem using an illustrative cyber attack scenario.

The system in this case is the UK CAV ecosystem defined in section 2 and the time interval is the period 2025 to 2030 which corresponds to an evolutionary state of that system which corresponds to the period of CAV four and CAV five product introduction onto the UK road infrastructure. This product introduction will feature an exponential growth during this period of the overall bass diffusion distribution curve, as shown in section three Figure 8.

In simple terms, applying the BHI here is to model the growth characteristics of benefits in the context of the UK CAV ecosystem businesses (for example revenues and profits as well as intangible assets) minus the possible harm arising from classes of cyber risk as a function of time. Each cyber risk is characterised by the product of the impact of their consequences and their likelihood which is associated with a given threat.

Benefits are defined in terms of the positive business/socio-economic impacts multiplied by their likelihood. Harm is defined in terms of the negative business/socio-economic impacts multiplied by their likelihood. A simple discrete formulation of how to calculate the associated growth is
shown below:

$$B_{tn+1} = B_t + [(b_t P_{bt}) - (h_t * P_{ht})]$$

For benefits b(t) and harm (h(t) with probabilities Pb(t) and Ph(t) respectively.

During the period 2025 to 2030 the benefits in the model are associated with the introduction of the level four and five CAVs onto the UK road infrastructure. As stated, the product introduction will feature exponential growth during this period of the overall bass diffusion distribution curve, and this in turn will drive the growth in benefits. In the BHI model this exponential order of growth is associated with a complexity level of four so the nature of the overall UK economic business impacts will be stochastic. In other words as shown earlier in Figure 1, strong emergent behaviour is expected. For example unpredicted value added spin off services generating new revenue streams are likely to occur.

According to the Society for Motor Manufacturers and Traders (SMMT)[8] the overall economic benefits of CAVs to the UK are expected to be in the region of £51bn per year by 2030, of which £16bn accrue to adjacent industries such as telecoms, technology, digital services and freight. It is also expected that up to 320,000 new jobs will be created, 25,000 of which are in automotive manufacturing, in the same period.

Furthermore, given that 94% of traffic accidents occur due to human error, significant social benefits are expected to be realised in increased safety that comes with automation, which could see 2,500 lives saved and 25,000 serious accidents prevented in the UK between 2014 and 2030.

Projected market value from CAV sales in the UK in 2025 is £35bn, rising to £46bn in 2030 according the market forecast for connected and autonomous vehicles by transport Systems Catapult[9].

The likelihood of these socio-economic and business benefits being realised is to some extent influenced by the governance domain of the UK CAV ecosystem since it can help stimulate the market for example through its approach to providing a CAV supportable updated transport legal framework and a light touch regulatory regime and through investments in CAV test and intelligent road infrastructure pilot environments. For this illustrative model it will simply assume 50% likelihood for the forecast benefits.

THE ILLUSTRATIVE CYBER ATTACK SCENARIO

The report now defines a viable illustrative cyber attack scenario on the UK CAV ecosystem and look at two time intervals, the first being the year 2025 and the second being the year 2030.

The threat source selected is a nation state for example China, which could potentially be an adversary at some point as the geopolitical landscape evolves. The scenario selected is an insider supply chain attack that exploits vulnerabilities in the complex global supply chain of CAV manufactures and the 5G intelligent road infrastructure providers. The objective here is to have the ability to remotely access and or control CAVs and the associated UK CAV ecosystem data.

The supply chain risk here is that of a nation state threat source in this case the nation states advance persistent threat (APT) group installing backdoors in the CAVs through microchips on or backdoors designed into motherboards on components that will be used as part of major brands of CAV used on UK roads. The compromised components in this scenario are manufactured by Chinese owned or controlled companies (the threat actors). These are generic components used by multiple European and UK CAV manufactures and in this illustrative scenario also include 5G systems used in the UK intelligent road infrastructure.

A high level view of the initial phase of the attack vector associated with this illustrative attack scenario is depicted in Figure 9 above. In line with conventional cyber risk assessments, the likelihood of attack can be determined by assessing the capability of the threat source/threat actors and their motivation/priority and associate that with the vulnerability being targeted by the attack vector.

The vulnerability being targeted in this illustrative attack scenario is the global supply chain domain of the UK CAV ecosystem. The UK DfT/CPNI principle that supply chain security risks are assessed and managed appropriately and proportionality will be difficult to enforce cost effectively in the time frame given the global scale, complexity and fragmentary structure of the CAV supply chains.

The ability to implant a malicious microchip on a motherboard that controls some data lines on the vehicles control area network (CAN) occur at the manufacturing phase of a component in the supply chain. The number of chips going into cars is steadily increasing and is likely to increase dramatically with the advent of 5G.

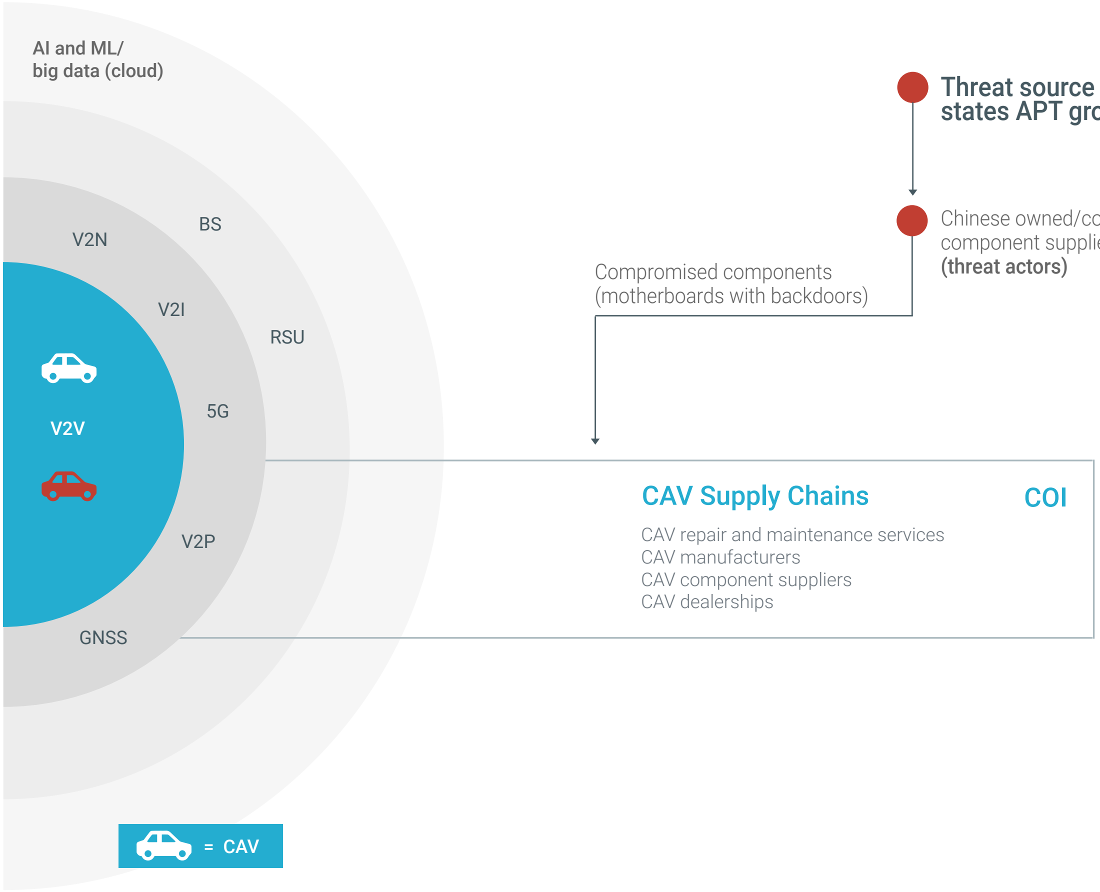
AI and ML/
big data (cloud)
V2N
BS
V2I
RSU
5G
V2V
V2P
GNSS
= CAV
Threat source the nation states APT group
Chinese owned/controlled CAV component suppliers
(threat actors)
Compromised components (motherboards with backdoors)
CAV Supply Chains
COI
CAV repair and maintenance services
CAV manufacturers
CAV component suppliers
CAV dealerships

Figure 9 – Illustrative cyber attack scenario on UK CAV-ecosystem supply chain

The second phase of this attack scenario occurs when the motivation is in place, for example because of a hostile geopolitical situation developing between the nation state threat source and the UK in this illustrative scenario. This is depicted schematically below in Figure 9a where the threat actor SSF (or their proxies) takes remote control of the infected CAVs. This could be to drive them into crowds or into each other or simply to activate a pre-programmed manoeuvre such as swerve right or emergency stop so that it occurs simultaneously in multiple CAVs at a set time.

The ability for state controlled auto parts manufacturers to design in backdoors to hardware components such as micro-processors and motherboards is well within their technical capabilities and very hard to detect.

EXPLORING THE VULNERABILITY AND CONTROL ASPECTS OF THE SYSTEM

The complexity of the CAV supply chain influences the associated complexity of the individual CAV platforms. The vulnerability of the CAV's increases in line with this complexity which is driven in part by their rapidly evolving interconnectivity with other CAVs and with the evolving intelligent road infrastructure for example through 5G and with telematics services and through 5G and 4G communication channels.

Classic risk mitigations imply that the system can be controlled in the presence of threat actors, so that their threat is reduced or effectively removed. However as exemplified by the case of the UK CAV ecosystem the systems that support ecosystems of businesses, and the information and computer technologies that support a given organisation's business, are increasingly complicated.

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. In contrast, control of complex systems - including distributed networks of actors and components - and control of systems of systems are poorly understood and mostly poorly characterized. 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.

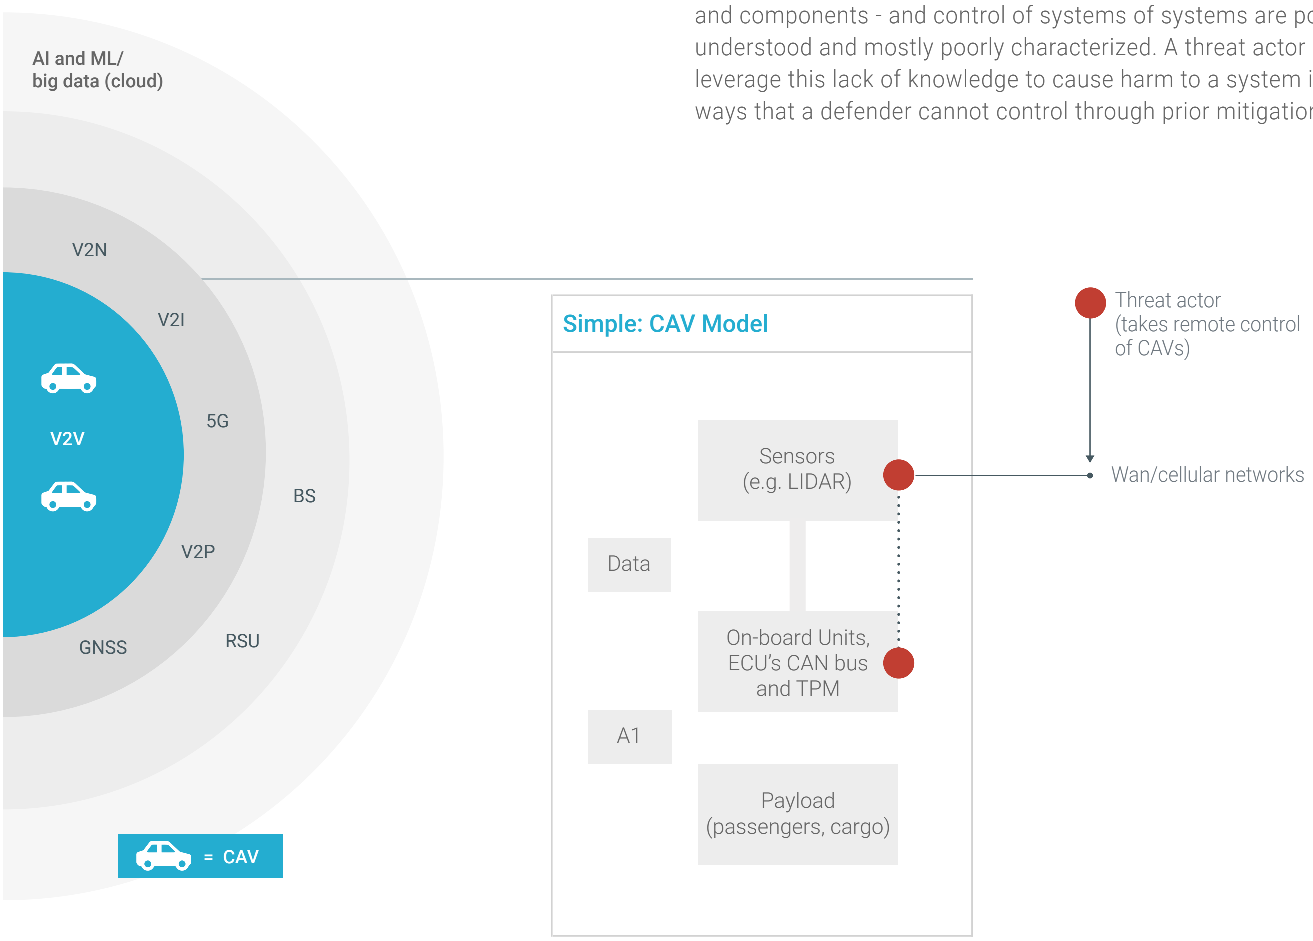

Figure 9a – Second phase of the attack scenario, where threat actor takes remote control

In the BHI model the vulnerability level (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 2025 the VL of the UK CAV ecosystem is VL (four) in line with its level of complexity. The VL four is illustrated from a control perspective in Table 3.

Vulnerability level	4
Threat class	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 a state space.
Attackers control	Radical ignorance – black swan events may occur as preparation for these events are frequently hindered by a pretence of knowledge of all the risks. Scenario modelling using Shackle Potential Surprise.
Economic rationale	Localised significant intra CAV Domain operational disruption. Minor UK wide disruption of the overall UK AV ecosystem operations. No significant loss of life. Minor impact on most intangible assets.

Table 3 – VL 4

In simple terms this VL is representative of the fact that the CAV manufactures and buyers actual knowledge of whether the microchips or motherboard of a CAV component has a backdoor designed in by the supplier is effectively hidden. In this case by the prohibitive cost of the analysis and associated reverse engineering at this level of detail that would be required to detect such latent threats across the diverse range of CAV components.

This lack of knowledge differentiates the risk decision from that of a gambler who knows the odds against winning but is prepared to take the risks as they feel lucky. In the context of cyber threats, it is the attacker who possess the knowledge with which they can harm the defender. In other words there is a move from a knowledge mode of rational ignorance to one of radical ignorance.

EXPLORING THE LIKELIHOOD OF THE ATTACK SCENARIO

In classic risk assessments the likelihood of a cyber attack on a particular target of interest is modelled 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 their level of motivation/priority for attacking that target of interest.

Threat Level = F(Capability of Threat Source/Actor(t), Motivation/Priority(t))

The capability of the threat source and associated threat actors in this scenario are those associated with a nation state, in this case the capabilities of the SSF. The capability of such nation state actors for launching sophisticated cyber attacks is very high.

The attack vector in the scenario exploits significant vulnerabilities in the supply chain that are relatively easy to compromise given direct or indirect control of the supplier companies. So the actual capability relative to the difficulty in exploiting those vulnerabilities is indeed very high.

There is motivation present in this scenario also. In 2017 the DoD reported that, “the SSF may seek to use its cyber warfare capabilities to collect data for intelligence and cyber attack purposes; to constrain an adversary’s actions by targeting network based logistics, communications, and commercial activities; or to serve as a force multiplier when coupled with kinetic attacks during times of crisis or conflict”.

In this illustrative scenario the priority to put the backdoors (effectively latent zero day attacks) into this target of interest will be high since over time the UK CAV ecosystem will become a key part of the UK transport critical national infrastructure. This is distinct to the motivation to exploit this latent zero day to actually launch an attack.

Having the latent zero day in place gives a hostile nation state the potential to launch a cyber attack to generate a significant socio economic impact on the UK if and when it deems that’s necessary. Alternatively it could exploit the vulnerability to exfiltrate potentially vast quantities of data which could include personal data and valuable CAV related intellectual property.

In other words the likelihood of the CAV components being compromised with such latent zero day backdoors is ‘very likely’.

The likelihood of an actual attack being executed that exploits (and thus exposes) the zero day vulnerability will depend on the state of the geopolitical relationship between the UK and the nation state threat source. In simple terms zero day exploits are powerful cyber weapons that you don’t want to exploit until you really need them since once you use them they soon become known and mitigated.

EXPLORING THE POTENTIAL IMPACT/HARM OF THE ATTACK SCENARIO

This illustrative scenario will assume the consider deleting or changing launch an attack exploiting the latent zero day vulnerabilities that they have designed into the compromised CAV components. The objective is to cause economic damage to consider the UK as part of a cyber warfare campaign that is escalating in the year 2025.

The levels of harm is modelled in the example of the impact levels on UK CAV ecosystem as shown below in Table 4:

Level	Impact
● Very high	Overall capability of the UK CAV ecosystem brought to a halt. significant socio-economic scale disruption and/or significant large scale (100+) loss of life. high impact on all intangible assets.
● High	Total disruption of one or more UK CAV domains (for example an urban core) and/or (5 to 100) loss of life. Some socio-economic disruption. Significant impact on most intangible assets.
● Medium	Localised significant intra CAV domain operational disruption. Minor UK wide disruption of the overall UK AV ecosystem operations. No significant loss of life. Minor impact on most intangible assets.
● Low	Localised Intra CAV domain short term operational disruption.

Table 4 – UK CAV ecosystem level impact levels

As highlighted above in Table 4 in assessing the impact of a successful cyber attack on the UK CAV ecosystem the potential harm to both tangible and intangible assets needs to be included and also given the kinetic nature of the target of interest the potential loss of human lives. An example of what is meant here by an intangible asset is brand equity which can be lost as a result of the reputational damage caused by falling victim to a successful cyber attack.

This illustration assumes the SSF (via a proxy) take remote control of some level four and five CAVs. In this model of the UK CAV ecosystem in the year 2025 there are level four lorries operating (for example platooning) on smart motorways but not operating at level four in urban cores. There will be level five taxi CAVs and buses but only operating in restricted zones in urban cores.

By taking remote control of some of these CAVs the threat actors can cause the CAVs to simply stop operating or to drive into roadside objects, other vehicles or pedestrians. In this scenario it is assumed the threat actors get the vehicles to simply swerve into oncoming traffic resulting in say, 20 crash events across the UK at the same time and resulting in say, the loss of less than five lives and the temporary closure of smart motorway lanes.

In the model this would be classed as having a medium impact. The impact level will however be different at different points in time, as will the motivation of the attacker. For example it would potentially be very high if carried out aggressively in 2030 when there are is significant amount of level five as well as a large amount of level four traffic live on the UK road infrastructure.

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

THE BENEFIT HARM INDEX PERSPECTIVE ON OUR SCENARIO

The overall socio-economic benefits of the UK CAV ecosystem grow over time in line with a bass diffusion distribution as shown earlier in Figure 8. As highlighted in the illustrative cyber attack scenario the harm that can be inflicted on the ecosystem by a specific threat can also grow with time. The associated threat level will also vary with time.

Defined earlier is a discrete formulation of benefit as:

$$B_{tn+1} = B_t + [(b_t P_{bt}) - (h_t {*} P_{ht})]$$

Where $\mathbf{B_t P_{bt}(tn)}$ is proportional to the threat level

$$\textbf{Threat Level = F(capability of threat source/actor(t),}$$
$$\textbf{motivation/priority(t))}$$

The BHI relates to differences in the complexity levels of benefit (CLb), and harm (CLh), over some time interval, TIi assuming M distinct threats (j) where j ranges from 1 to M.

$$CL_b(T1_i) - CL_h(T1_i)$$

Where:

$$CL_b(TI_i) = \text{MAX}\{\text{Level}(\text{Distribution}(b(TI_i)),$$
$$\text{Level}(\text{Distribution}(\ P_b(TI_i))\}$$

And:

$$CL_h(TI_i) = \text{MAX}\{\text{Level}(\text{Distribution}(h(TI_i))\}, \text{MAX}$$

$$\sum j\ \{\text{Level}(\text{Distribution}(j(TI_i))),(\text{Distribution}(\text{priority}\ (j(TI_i)))\}$$

In simple terms for this illustrative cyber attack scenario on the UK CAV ecosystem there are an overall set of socio-economic benefits that are growing in line with a bass diffusion distribution curve, as described earlier in Figure 8. During the strong growth period 2025 to 2030 the benefit growth rate is exponential which equates to a benefit complexity level four.

During the maturing growth period 2030 to 2035 the benefit growth rate decreases rapidly from exponential to asymptotic which equates to a benefit complexity level four decreasing to zero during this period.

Given the UK government's strategy of making the UK a leading player in this field, the significant investments by CAV manufactures in what is an active globally competitive market it can be assumed that the associated probability of following that distribution is high and flat so for simplicity here it is assumed it is close to one. The accuracy of the market forecasts assumed are also high for this illustrative example.

For the period 2025 to 2030 is:

$$\mathbf{CL_b(2025\text{-}2030) = 4}$$

And for the period 2030 to 2035 is:

$$\mathbf{CL_b(2030\text{-}2031) = 3,\ CL_b(2031\text{-}2032) = 2,\ CL_b(2032\text{-}2034) = 1,\ CL_b(2034\text{-}2035) = 0}$$

Again for this illustrative cyber attack scenario on the UK CAV ecosystem an attack scenario has been selected for the impact (harm) of which also grows with a bass diffusion distribution curve. However this bass diffusion curve for harm lags in time behind the benefits bass diffusion curve since to a significant extent the level of systemic harm that can be inflicted is dependent on the level of maturity of the benefits that can be impacted.

The capability associated with the illustrative threat scenario will grow exponentially (for example complexity level four) in line with the growth in complexity of the CAV supply chain, CAV platform and intelligent road infrastructure system of systems. This brings with it an explosive growth in the size of the overall threat surface for the selected attack vector. It is expected that the nation state threat actors that feature in the illustrative cyber attack scenario to exploit this to embed their backdoors such as by design in the motherboard or through malicious microchips and or firmware.

However the motivation/priority of the nation state source in the illustrative scenario to launch an attack that exploits these implanted zero day vulnerabilities will remain low until the potential impact (harm) reaches a significant level, of socio-economic gain to itself and or damage to the UK. The report will focus on the case where the nation state intends to do harm. In this scenario the motivation/priority will remain low (for example close to zero) until there is a geopolitical tension between the nation state threat source and the UK sufficient for it to launch such an attack.

To bring the scenario to life, a hypothetical geopolitical tension arising in 2030 from an earlier escalating trade war resulting now in significant forms of aggression between the nation state threat source and the UK can be considered. This changes the motivation/priority from low to high and takes the threat level to a very high value compounding the resulting cyber risk of the exponential rate of growth of the harm/impact.

Consequently for the illustrative cyber attack scenario the threat level will be low during the period 2025 to 2030 and the level of harm that can be inflicted is also less than the benefits being generated during this period.

For the period 2025 to 2030 is:
CL_h(2025-2030) < 3 which reflects the late embryonic phase of bass diffusion distribution of the growth of harm.

For the period 2030 to 2035 is:
CL_h(2030-2035) = 4 which reflects the strong growth phase of bass diffusion distribution of the growth of harm.

For this illustrative cyber attack scenario the BHI during these growth periods of the UK CAV ecosystem can be represented schematically as shown below in Table 3 where
BHI = $CL_b(TI_i)$ - $CL_h(TI_i)$.

There is particular interest in the case when BHI <= 0, when the growth order (CL) of the harm exceeds the growth order of benefit. In this case, unless there is mitigation, it can reasonably be expected that however the benefit grows it will be overtaken by harm.

	Period/time interval				
	2025 - 2030	2030 - 2031	2031 - 2032	2032 - 2034	2034 - 2035
BHI value	1	-1	-2	-3	-4

Table 3 – benefit harm index of the illustrative cyber attack scenario on the UK CAV ecosystem

Even though there is a focus on just one illustrative 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. Section 5.2 of this paper illustrates this in more detail and also show how the implications wheel methodology can be used to try and detect emergent threats in this context.

The mitigation actions for this type of risk need to be put in place much earlier during the embryonic growth phase of the CAV ecosystem. Putting in security controls retrospectively after the event would be costly and time consuming in this scenario given the need to identify and replace malicious CAV components which would potentially require their recall to CAV maintenance and repair operators and subsequent loss of live CAV traffic and associated benefits.

In applying the BHI formally one would of course look at the BHI systematically across a significant number of risks rather than just the one illustrative risk highlighted here.

MITIGATING GROWTH OF HARM AT THE UK CAV ECOSYSTEM LEVEL

There are a number of ways to try and mitigate the emergent threats associated with the growth of harm in the complex UK CAV ecosystem. Typically these involve designing a set of security controls that seek specifically to mitigate risks from emergence. These controls will necessarily need to detect and potentially isolate and neutralise the impact of an attack.

This paper highlights two mitigation examples, first illustrating the sharing of cyber threat Information across the cyber ecosystem and second illustrating an approach to predicting black swan events that are characteristic of the radical ignorance inherent in these VL four complex systems of systems.

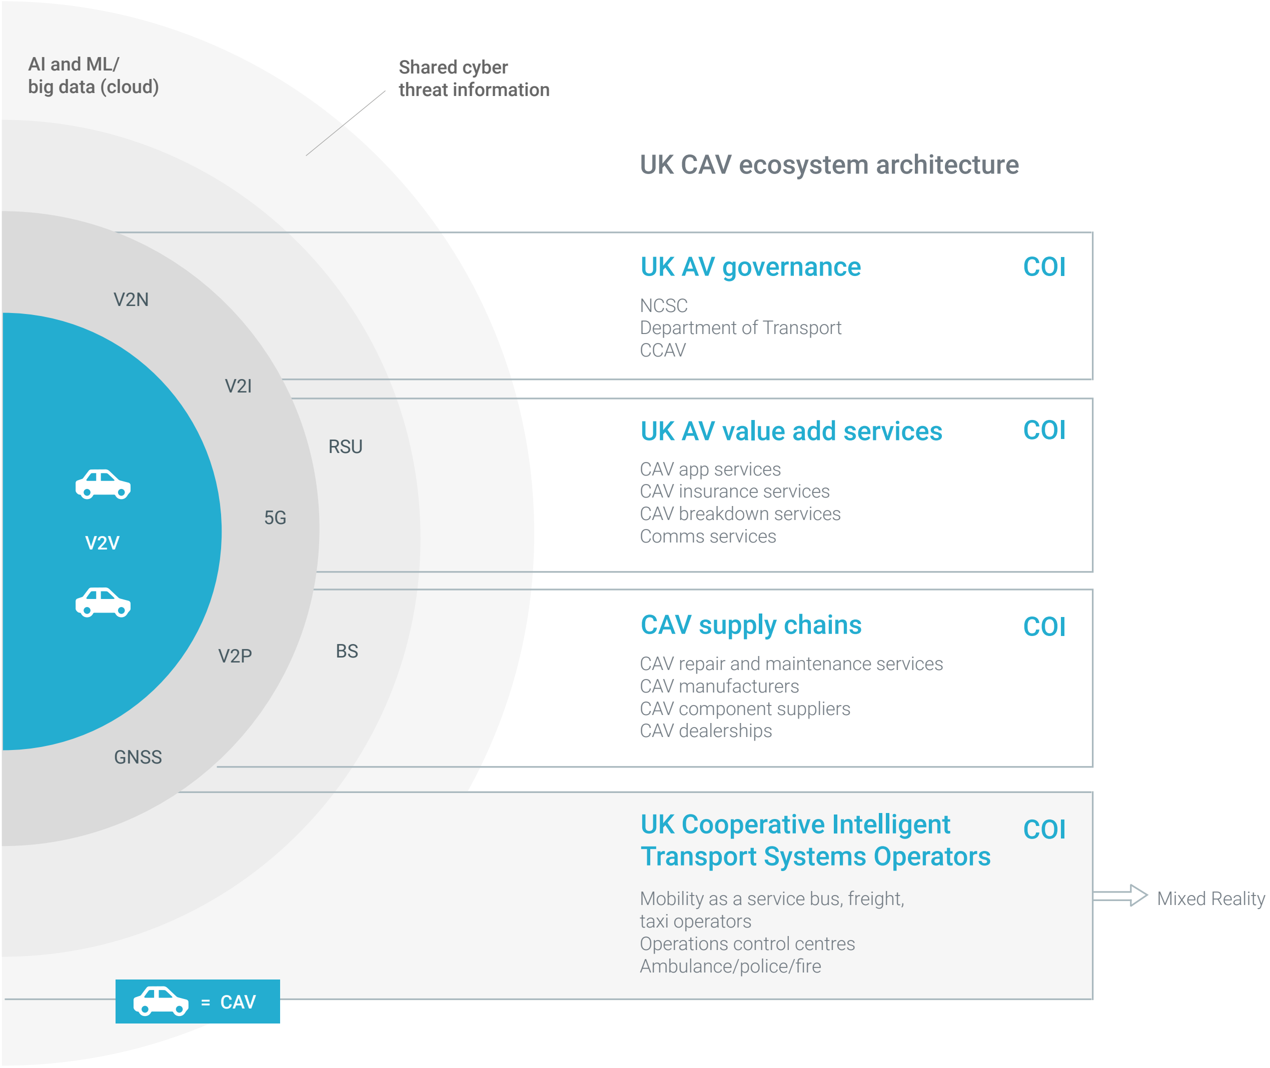

UK CAV ecosystem architecture

UK AV governance COI

NCSC
Department of Transport
CCAV

UK AV value add services COI

CAV app services
CAV insurance services
CAV breakdown services
Comms services

CAV supply chains COI

CAV repair and maintenance services
CAV manufacturers
CAV component suppliers
CAV dealerships

UK Cooperative Intelligent Transport Systems Operators COI

Mobility as a service bus, freight,
taxi operators
Operations control centres
Ambulance/police/fire

Figure 10 – The UK CAV ecosystem - Mitigation through sharing CTI

MITIGATING EMERGENT RISK BY SHARING CYBER THREAT INFORMATION

In the complex VL four 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 of those are likely to be successful.

However by sharing cyber threat information across the ecosystem, effective cyber security can be provided which requires cooperation and collaboration among all the entities involved. Increasing the information available for analysis allows better prediction, prevention and mitigation of cyber attacks.

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

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 constraints and inhibitors that companies face when sharing such data. These include:

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

For example sharing CTI data containing details of a data breach or personal identifiable information needs to be managed in a way that ensures regulatory compliance and may require anonymisation or homomorphic encryption to ensure confidentiality.

An ecosystem CTI sharing service that is viable is therefore one where:

- The member entities of that ecosystem are able to choose the type of confidentiality controls that are appropriate for safeguarding their CTI data on the CTI service, for example to go for either open access, or data anonymisation techniques, or even use homomorphic encryption based techniques
- Due to the availability of different data confidentiality and access options, the member entities can confidently share specific types of their CTI data via the CTI sharing service, with even non-trusted third parties
- The CTI sharing service should incorporate diverse techniques for supporting the protection of CTI data, so the member entity does not have to be aware of the inner workings of these techniques. Thus the member entities shall be able to choose and consume from the alternative techniques most suitable to them from their own perspective without worrying about their design and implementation
- The CTI sharing service can also incorporate diverse techniques for analysing the shared CTI without the member entities worrying about issues like information leakage, as this process should be transparent for the member entities

The report focuses here on the main European CTI sharing initiative for example C3ISP which an EU project that is part of the EU Horizon 2020 project. This platform addresses the concerns raised above. Digital Catapult is a member of this EU research project. The C3ISP concept is described on their website[11] as:

"Providing effective cyber security requires cooperation and collaboration among 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 the information in the most appropriate way depending on the scenarios. This is aligned with the main guidelines of the European Cyber Security Strategy."

The C3ISP mission is to define a collaborative and confidential information sharing, analysis and protection framework as a service for cyber security management. Of particular interest, in this white paper, is a C3ISP component that is focused on supporting small medium enterprises (SME's) to share CTI. This is important in domains like the CAV supply chain where SME's will typically not have as strong security capabilities as the larger enterprise member entities of the UK CAV 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 targetting SMEs

Security issue

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

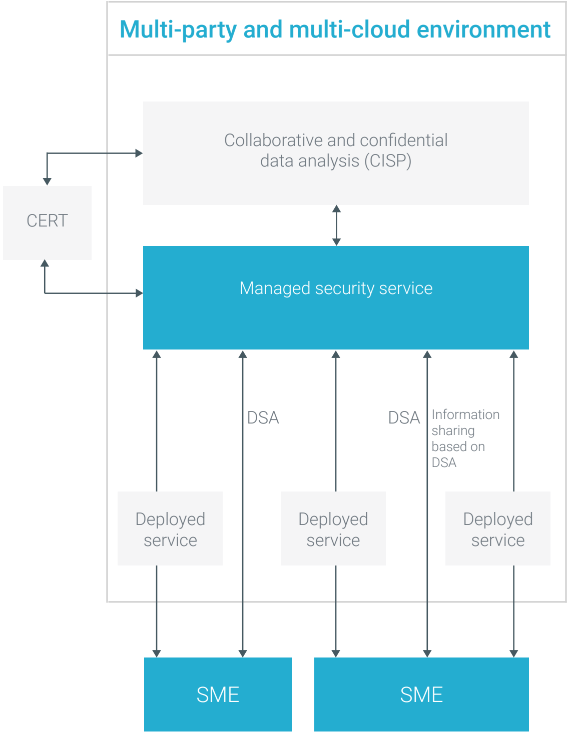
Multi-party and multi-cloud environment
Collaborative and confidential data analysis (CISP)
CERT
Managed security service
DSA
DSA
Information sharing based on DSA
Deployed service
Deployed service
Deployed service
SME
SME

Figure 11 – C3ISP Pilot for CTI sharing amongst SME's. (Ref the C3ISP EU project)

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

- Deploy the SME pilot providing a secure multi-party cloud environment for collaborative information sharing, performing collection and analysis of SME 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 SMEs. The capability of providing security intelligence obtained through the collaborative analysis will be evaluated. Also the capability of delivering this intelligence without disclosing private information and the compliance with DSA policies will constitute an important evaluation index

The sharing of CTI data in a way that enables the 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 interoperable automated exchanges of CTI.

The C3ISP Architecture also supports a shared platform where not just SME's but enterprise level participants, ISP's and CERTS can collaborate. This clearly can be used to support the model proposed in this white paper for an ecosystem level shared CTI capability, for each UK CNI ecosystem. (Including the test case the UK CAV ecosystem)

The C3ISP platform provides CTI analytics services (such as that provided by BT Saturn) that provides 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 the current threat in the operational context, for example of the UK CAV ecosystem.

AN APPROACH TO MITIGATING EMERGENT RISK/RADICAL IGNORANCE

The approach highlighted is the use of the implication wheel™ methodology to help uncover emergent threats. Figure 12 illustrates a context that will be used to introduce the implication wheel concept. It features one threat scenario as it could unfold in the UK CAV ecosystem.

Cyber ecosystems are complex 'systems of systems' like the UK CAV 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 of indirect and unanticipated consequences. The direct first order effects are often relatively easy to predict and mitigate. However the second and third order effects are much less obvious and may contain surprises some of which will be of significant concern, these are referred to as black swan events.

The implication wheel™ is a methodology which in simple terms is a participatory "smart group" method that uses a structured brainstorming process to uncover multiple levels of consequences which can lead to the discovery of black swan events. Each smart group is comprised of a diverse set of individuals that will bring different perspective to the task.

The team members of each smart group starts by considering an initial event, in this case as illustrated in Figure 12 the initial event could be the hostile state actor (SSF) installs backdoors/malware into the Chinese CAV component suppliers that they indirectly control (for example through the Chinese equivalent of the US Patriot Act). The threat actor is shown as the deep red circle event on the top left of Figure 12. They are then asked 'what might happen next?' This generates the direct first order consequences.

Figure 12 illustrates some potential first order consequences that propagate outwards from the initial event at the top left. These first order consequences include possibility of the infected ECU module being detected during assembly of the CAV four/five, through to multiple CAV four/five manufactures not detecting it resulting in multiple fleets of infected CAVs being released onto the UK intelligent transport infrastructure.

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

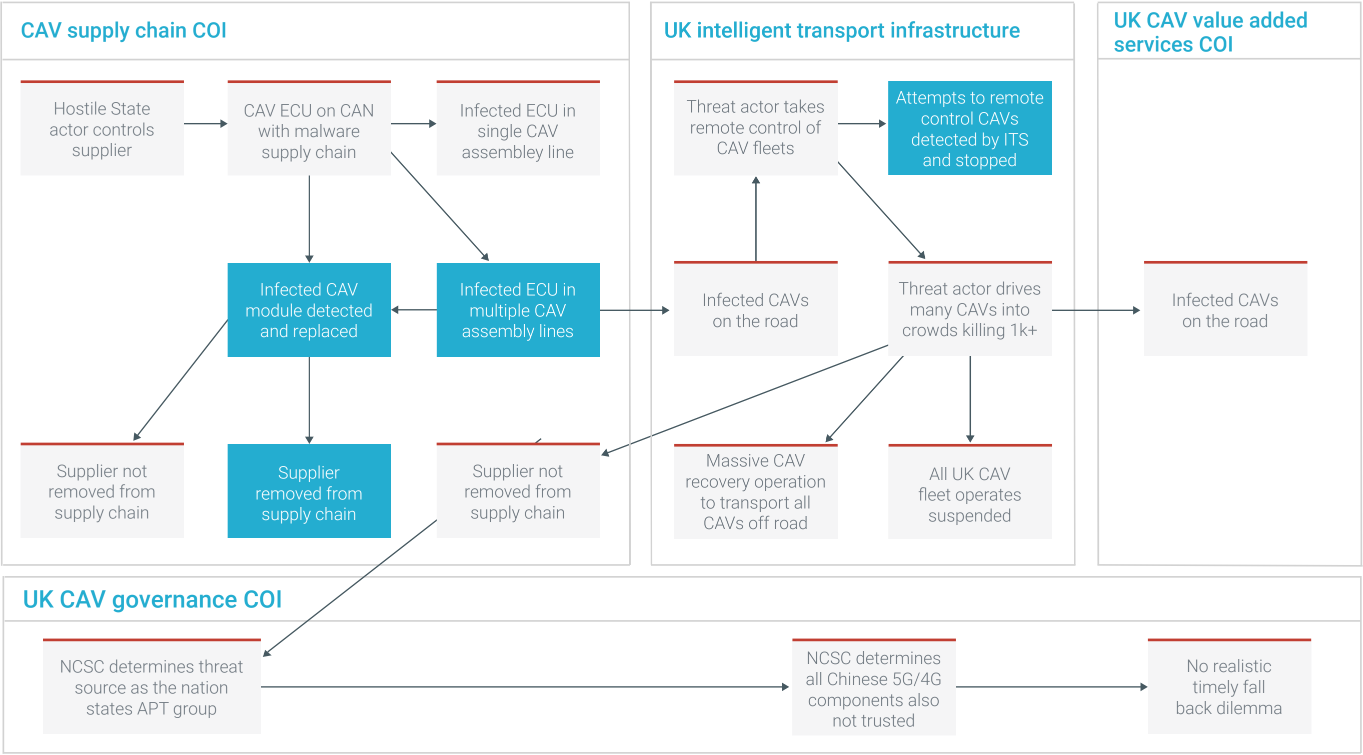
CAV supply chain COI
Hostile State actor controls supplier
CAV ECU on CAN with malware supply chain
Infected ECU in single CAV assembley line
Infected CAV module detected and replaced
Infected ECU in multiple CAV assembly lines
Supplier not removed from supply chain
Supplier removed from supply chain
Supplier not removed from supply chain
UK intelligent transport infrastructure
Threat actor takes remote control of CAV fleets
Attempts to remote control CAVs detected by ITS and stopped
Infected CAVs on the road
Threat actor drives many CAVs into crowds killing 1k+
Massive CAV recovery operation to transport all CAVs off road
All UK CAV fleet operates suspended
UK CAV value added services COI
Infected CAVs on the road
UK CAV governance COI
NCSC determines threat source as the nation states APT group
NCSC determines all Chinese 5G/4G components also not trusted
No realistic timely fall back dilemma

Figure 12 Illustrating the implication wheel approach applied to the UK CAV ecosystem

This process is then repeated at each first order consequence to create an associated set of second order consequences. In Figure 12 the second order consequences illustrated, builds on one of the first order consequences in this case that of multiple infect level four/five CAVs from multiple manufactures being live on the UK ITI. The second order consequence shown is that of the threat actor SSF taking remote control of those CAVs and attempting to drive them into crowds of people in different city locations across the UK. These consequences range from attempted attacks being detected in good time by the ITI operators, and the CAVs being stopped; to the case where the attack is successful and results in circa one thousand fatalities.

The worst case second order scenario (shown in black inside the UK intelligent transport infrastructure domain Figure 12). Although shown in black it is not a black swan event since it relatively easy to predict such a second order scenario. However things get more interesting when the event is used to move out to explore the third order consequences associated with it.

In figure 12 these third order consequences include the suspension of all UK level four/five CAV fleets, the massive recovery operations to remove CAVs (which can no longer be driven safely) off the live road infrastructure, and the third order consequences for the CAV manufactures in recalling all those CAVs and involving NCSC to try and determine the vulnerability the threat actors and any other vulnerabilities.

The black swan effect shown here is the potential macroeconomic level impact caused by the inability to recover the intelligent road infrastructure with level four/five CAV fleets, due to the time taken to verify forensically the extent of the zero day threats across all the CAV supply chains and then replacing them with trusted components in the CAVs.

When the implications wheel is used more formally in this context, a layered structure like the wheel is produced, shown below in Figure 13.

Here, one second order effect and its associated third order effects can be seen.

A part of the implication wheel methodology is to allow the smart group participants to propose levels of impacts/importance and likelihood for each consequence. For example the likelihood of the CAVs being deliberately used to kill by a nation state threat actor might be low, relative to them simply causing all the CAVs to stop running or misbehave. Although interestingly the third order effects in Figure 13 would still apply once the vulnerability of the operational CAVs to being controlled maliciously was demonstrated.

Example mitigation approach: Using the Implications Wheel

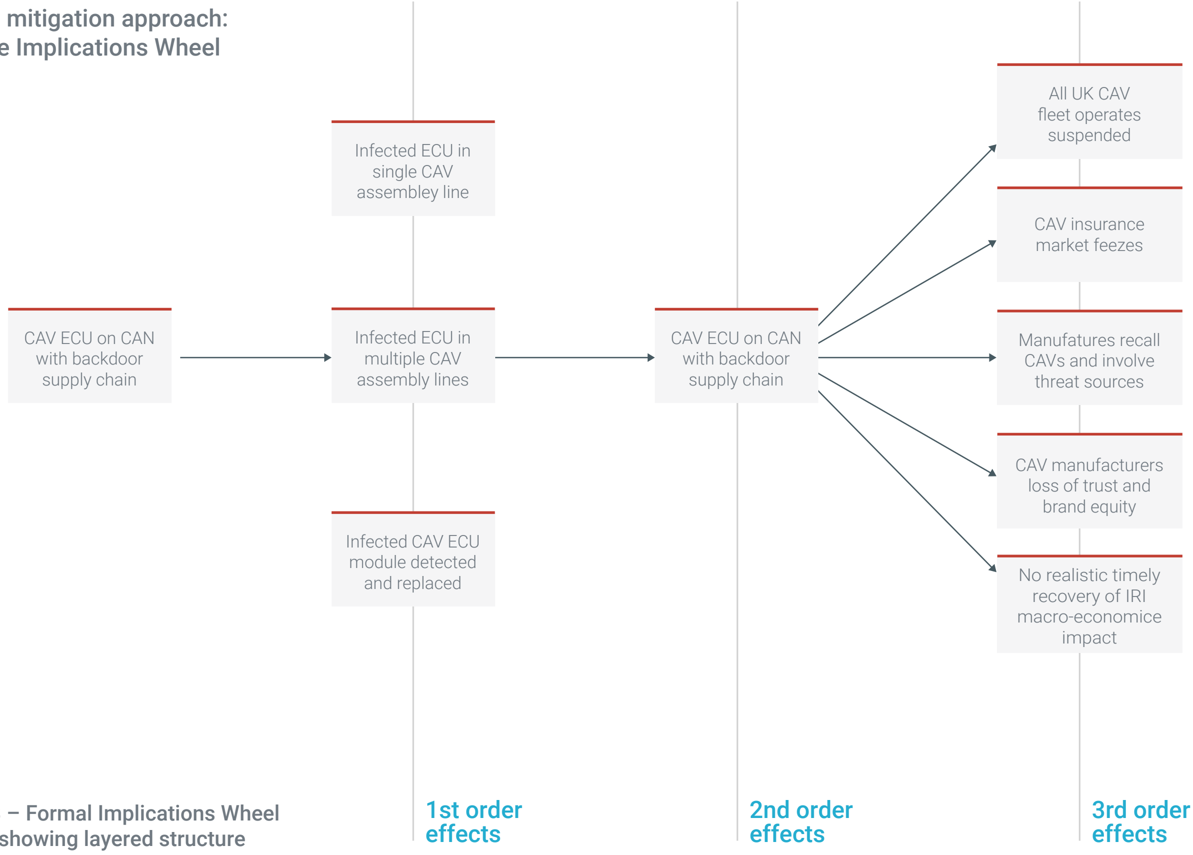

Infected ECU in single CAV assembley line
CAV ECU on CAN with backdoor supply chain
Infected ECU in multiple CAV assembly lines
Infected CAV ECU module detected and replaced
CAV ECU on CAN with backdoor supply chain
All UK CAV fleet operates suspended
CAV insurance market feezes
Manufatures recall CAVs and involve threat sources
CAV manufacturers loss of trust and brand equity
No realistic timely recovery of IRI macro-economice impact
1st order effects
2nd order effects
3rd order effects

Figure 13 – Formal Implications Wheel example showing layered structure

The impacts can be associated with a scale that ranges from macro level impacts on the entire UK CAV ecosystem down to small localised impacts on a specific member entity.

When exploring the impacts of attacks on cyber ecosystems, a number of facts needs to be included such as the impact on both intangible assets and on tangible assets. This scope is illustrated in Figure 14.

As shown in Figure 14 above cyber attacks can impact intangible assets but do not normally impact physical/tangible assets such as plant and machinery. However it is important to emphasise that attacks on cyber ecosystems like the UK CAV ecosystem that include kinetic components (in this case CAVs) can impact such physical systems. This means that there are risks to safety as well as to the usual data and IT systems confidentiality, integrity and availability (CIA). In other words cyber attacks on such systems could result in damage to physical infrastructure and multiple human fatalities.

To conclude, the illustrative black swan event is highlighted in Figure 13. The impact of this event on intangibles can be assessed in the context of the socio-economic value of the UK CAV ecosystem in the year 2030 when this cyber attack is being modelled.

In section four of this white paper it is noted that according to the Society for Motor Manufacturers and Traders (SMMT)[12] the overall economic benefits of CAVs to the UK are expected to be in the region of £51bn per year by 2030, of which £16bn accrue to adjacent industries such as telecoms, technology, digital services and freight. It is also expected that up to 320,000 new jobs will be created, 25,000 of which are in automotive manufacturing, in the same period.

As a result if the CAV's are out of action for say six months because of the time taken to unearth latent zero day threats that may still be in the CAV and 5G supply chain and restore the level four/five CAVs trust the socio economic impact could be of the order of £25bn plus potentially significant (1000+) job losses. The BHI process at this stage reverts to classic risk management, for example the provision of stronger supply chain security and recovery plans and strategies being put in place to mitigate such a black swan event.

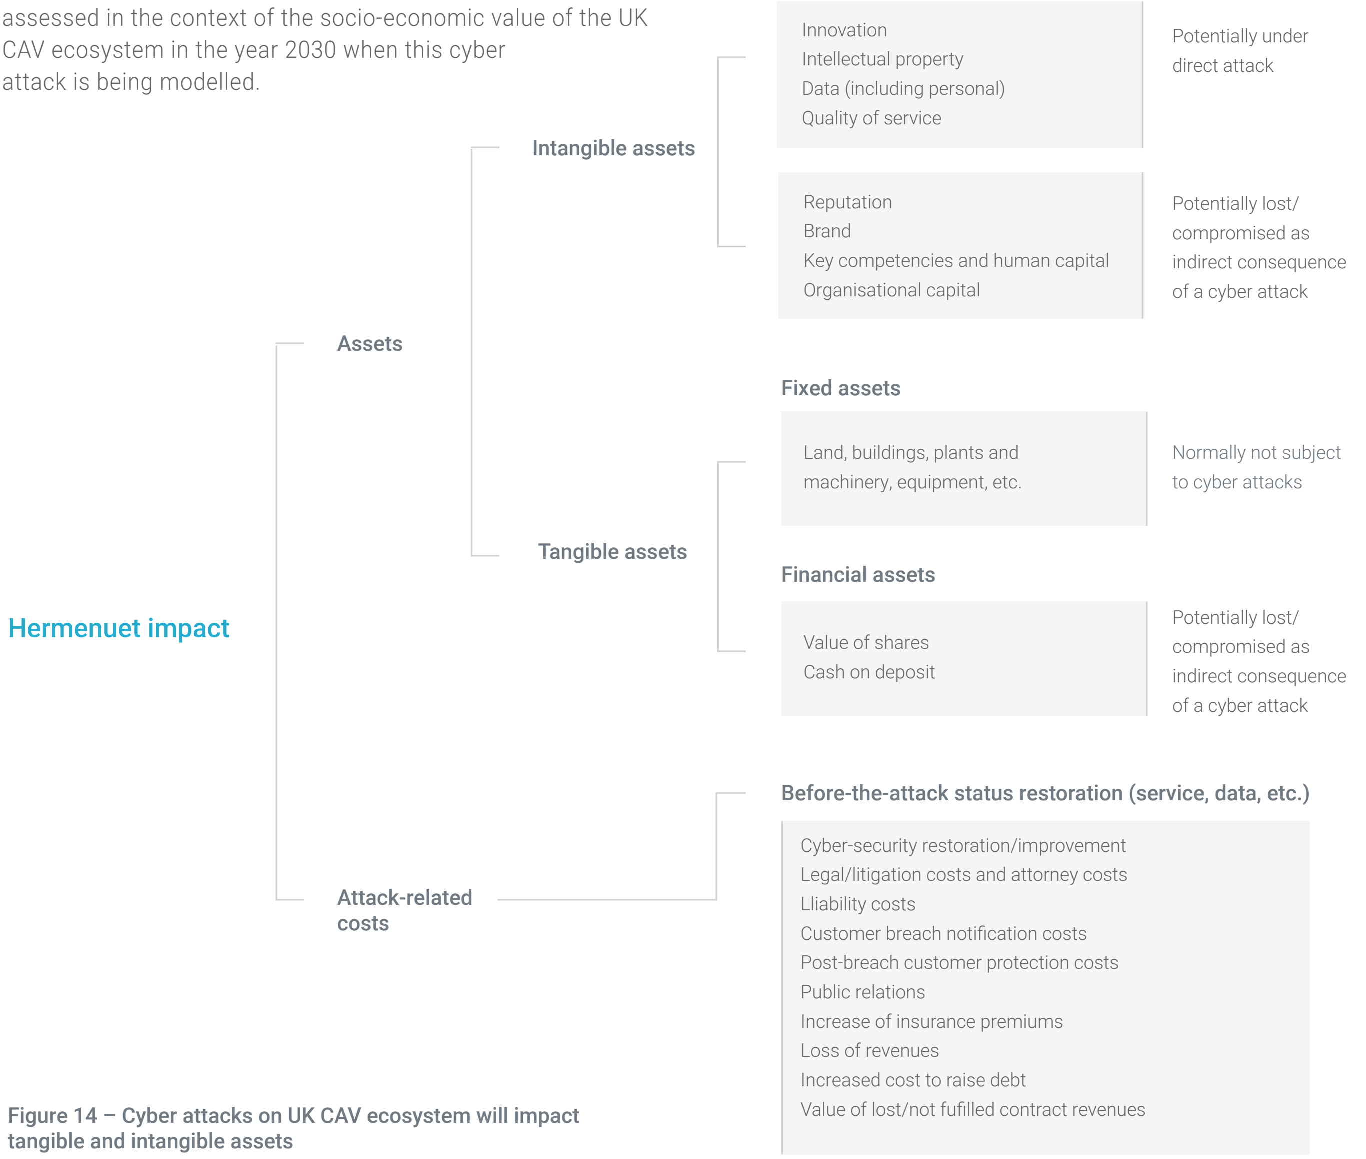
Hermenuet impact
Assets
Intangible assets
Innovation
Intellectual property
Data (including personal)
Quality of service
Potentially under direct attack
Reputation
Brand
Key competencies and human capital
Organisational capital
Potentially lost/ compromised as indirect consequence of a cyber attack
Tangible assets
Fixed assets
Land, buildings, plants and machinery, equipment, etc.
Normally not subject to cyber attacks
Financial assets
Value of shares
Cash on deposit
Potentially lost/ compromised as indirect consequence of a cyber attack
Attack-related costs
Before-the-attack status restoration (service, data, etc.)
Cyber-security restoration/improvement
Legal/litigation costs and attorney costs
Lliability costs
Customer breach notification costs
Post-breach customer protection costs
Public relations
Increase of insurance premiums
Loss of revenues
Increased cost to raise debt
Value of lost/not fufilled contract revenues

Figure 14 – Cyber attacks on UK CAV ecosystem will impact tangible and intangible assets

CONCLUSION

FIND OUT MORE ABOUT BHI/HERMENEUT PROJECT

This report has shown how to apply the BHI to CNI cyber ecosystems. In the UK CAV ecosystem case study, just one cyber attack scenario to illustrate the process. The formal application of the BHI to such CNI cyber ecosystems would uncover potentially significant emergent threats in advance of such threats being exploited by hostile nation state actors and their proxies, as well as threat actors such as terrorists.

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 EU Hermeneut project deliverable document, D4.2 BHI Index report. This is available on the Hermeneut site at the following link: https://www.hermeneut.eu/resources/

Hermeneut's cyber security 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: https://www.hermeneut.eu/about/

GLOSSARY

BHI	Business harm index
CAV	Connected autonomous vehicles
CAN	Control area network
CCAV	Centre for connected and autonomous vehicles
CNI	Critical national infrastructure
ECU	Electronic control unit
NCSC	National cyber security centre
P.E.S.T.L	Political economic social technical legal
RSU	Road side unit
V2V	Vehicle to vehicle
V2I	Vehicle to infrastructure
APT	Advance persistent threat

REFERENCES

The following research papers and other resources are referenced by this white paper:

Endnotes

1. Digital Catapult, 2019, Official Website. [online], Available at: https://digitalcatapult.org.uk/

2. Rogers, E.M. 1962 Diffusion of Innovations. New York: The FreePress

   Arthur, W. Brian, Durlauf Steven N and Lane David (Eds) 1997The Economy as an Evolving Complex System II. Proceedings Volume XXVII Santa Fe Institute Studies in the Science of Complexity, Reading, MA: Addison-Wesley.

3. Government UK CCAV Centre for Connected and Autonomous Vehicles, 2019, Official Website. [online] Available at: https://www.gov.uk/government/organisations/centre-for-connected-and-autonomous-vehicles/

4. Using 5G to remotely control a CAV https://www.telefonica.com/en/web/press-office/-/5g-can-make-remote-driving-a-reality-telefonica-and-ericsson-demostrate-at-mwc

5. Transport Systems Catapult, 2017, Market Forecast for Connected and Autonomous Vehicles. [pdf.] Available at: https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/642813/15780_TSC_Market_Forecast_for_CAV_Report_FINAL.pdf [Accessed on date]

6. Rogers, E.M. 1962 Diffusion of Innovations. New York: The FreePress

7. Bass, Frank M. (2004). “Comments on “A New Product Growth for Model Consumer Durables”: The Bass Model”. Management Science.50(12): 1833–1840.

8. The Society of Motor Manufacturers and Traders Limited, 2017, Connected and Autonomous Vehicles Position Paper. [pdf.] Available at: https://www.smmt.co.uk/wp-content/uploads/sites/2/SMMT-CAV-position-paper-final.pdf

9. Transport Systems Catapult, 2017, Market Forecast for Connected and Autonomous Vehicles. [pdf.] Available at: https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/642813/15780_TSC_Market_Forecast_for_CAV_Report_FINAL.pdf [Accessed on date]

10. C3ISP: collaborative and confidential information sharing and analysis for cyber protection https://c3isp.eu/

11. C3ISP: collaborative and confidential information sharing and analysis for cyber protection https://c3isp.eu/

12. The Society of Motor Manufacturers and Traders Limited, 2017, Connected and Autonomous Vehicles Position Paper. [pdf.] Available at: https://www.smmt.co.uk/wp-content/uploads/sites/2/SMMT-CAV-position-paper-final.pdf