bcs
The Chartered Institute for IT
IT Leaders Forum
Service Resilience
Phase II of the working group looking at the growing threat of software failure to the economy and society – eyeing the elephant in the room

Boosting confidence in the resilience of services dependent on software

September 2023

This Policy Think Piece has been developed by the Service Resilience Working Group (SR WG) of BCS (formerly, the British Computer Society) IT Leaders Forum. It is written by the co-chairs based on input from Working Group members and other experts: see Appendix 1.

The purpose of the SR WG is to improve the resilience of service delivery that is dependent on software systems - focusing on the UK. In this Policy Think Piece we identify initiatives in place and then go on to make recommendations for government and regulators, and for building appropriate capabilities among C-suite members, IT professionals and supporting disciplines.

The BCS is the UK's professional body for computing. It is governed by Royal Charter to advance education and practice in computing and IT for the benefit of the public. Its mission is "Making IT good for society."

Contents

Executive Summary

In making a living and in living our lives, we increasingly rely on services of digital systems. The software components of these systems are mostly sourced by the service providers from open source consortia or purchased from specialist suppliers, many based outside the UK. The digital systems built from these components are complex and tightly coupled. These characteristics mean that the combination of human error, cyber-attack and Natural Accident Theory will result in unpredictable service outages.[1]

These service outages are already a significant cost to the UK economy; to individual users, to citizens and to the public at large. All incur financial damage or loss and/or are otherwise seriously affected by service outages. Digital systems are not within users' choice, care or control, but they are becoming systems upon which users must rely.

The costs and consequences of software failure and service outages are increasing. This upward trend is due to technological as well as human factors. In addition to growing technological complexity, the widening use of AI and the escalating frequency of cyber-attacks carry increased risk of failures. The expanding reach of services to new users and environments means that more people are affected by service outages – the loss of "user hours" through service outages is increasingly recognised as a key measure.

Standards, guidelines, certifications, and regulation are helpful in improving service delivery. However, without increased transparency in recording and sharing instances of service outages, there is little data to support effective preventative and remedial actions.

The financial services sector leads in regulation to improve resilience of services. We suggest that the leverage effect of infrastructure on the UK economy suggests regulation or other approaches be considered across infrastructure sectors.

There are gaps in skills to manage this new environment. We recognise three affected communities: C-suite; IT professionals; and other professionals such as consultants, risk, quality and insurance managers, and business continuity managers.

The recommendations cover:

- Potential levers that government could choose to exercise;
- Role of regulators;
- Infrastructure sectors;
- Reducing skill gaps for the C-Suite and other leaders;
- Education and certification for IT and resilience professionals.

[1] Perrow, C. (1984), Normal Accidents: Living with High-Risk Technologies, Basic Books, New York is the origin of the term, which is still being debated, e.g., Turner, B. and Pidgeon, N. (1997), Man-Made Disasters, Butterworth Heinemann, Oxford.

In Appendix 1, we thank the members of the SR WG and the many people who contributed to this report. We also acknowledge gratefully the extensive support from the National Preparedness Commission.

Appendix 2 provides more detail on organisational resilience, service resilience and metrics. Appendix 3 has links to relevant standards and guidelines. Appendix 4 expands on our plans and proposals on learning and capability building. Appendix 5 records WG submissions to the Department of Science Innovation and Technology.

Introduction

This report proceeds from an examination of the background and context that inform the problem of software-based service resilience (Section 1) to findings about opportunities for action (Section 2), to conclusions on key areas (Section 3), and concludes with recommendations (Section 4).

1. Background and Context

1.1 Defining our focus

The BCS[2] Service Resilience Working Group (SR WG) held a joint Round Table with the National Preparedness Commission (NPC) in November 2022 to examine the risks to the UK economy from software failure.[3] It was concluded that:

- Software is different. It is intangible and obeys different rules from physical systems, creating problems in, for example, quality control and regulation.
- The software element of digital systems failure is a cost to the economy and society which increases as software becomes a utility, is in wider usage, and is more vulnerable to failure.
- More people and organisations need to be aware of the actual and potential impact of software failures.

Following this Round Table, the SR WG launched as multi-faceted investigation into what could be done to raise awareness of the risks of software failure and to explore what could be done to prevent failures or mitigate their consequences.

Our Terms of Reference directed us to identify means for reducing the impact of software failures on the UK economy. So we realised that we should focus on resilience – viewing systems from the outside - in order to define priorities for improvement. This has led us to actions during this year and further proposals to prevent and mitigate impacts of software failure, that is to improve resilience.

We quickly realised that many people – including IT professionals – had not understood the paradigm-shifting changes in software supply and usage over the past few years. We had covered these in our earlier Policy Think Piece[4] but found that these characteristics needed to be at the fore of our considerations. So we identify these in the next section.

[2] Was previously known as the British Computer Society

[3] National Preparedness Commission

[4] https://www.bcs.org/media/9679/itlf-software-risk-resilience.pdf

1.2 The organisational context of resilience

Most organisations deliver services which rely on software. This means that most organisations need to anticipate and plan for software failures, for how to handle disruptions to the running of their services, and for how to deal reasonably and fairly with the consequences. These are the basic elements of resilience.

The “typical” organisation has characteristics that constrain the actions needed to improve their resilience when software fails. Some of these are:

- Boards focus on functionality as the key commercial driver, not such things as resilience and security, both of which are wider than cyber threats.
- C-suite attention prioritises cyber threats and the impact of AI, not resilience.
- C-Suite may not receive regular updates on service outages.
- The education of IT professionals focusses on system development (rather than on operation).
- Budget constraints, and a view that resilience is a cost without immediate benefit, limit initiatives to address it.
- Systems delivering services are required 24/7, but strategy often neglects the need to ensure all the (3rd) parties needed to provide services can also provide 24/7 resilience.
- Tightly coupled, complex systems where no-one has the complete picture or collection of needed skills are liable to unpredictable failures.
- Combining old and new technology means services are dependent on multiple components or subsystems, many are developed externally, others, often internally developed, are inherited from the past, resulting in operational instability.
- New software releases or patches (updates) will usually not be synchronised between suppliers.
- Skill shortages at the level of boards and the non-tech C-suite as well as IT staff mean a lack of general knowledge, engineering or software training, or an insufficient range of competences needed for today’s roles.
- Limited ability to make changes to legacy software.
- IT frameworks are broad and deep, each with their own specialities and idiosyncrasies – network, cloud, data storage, endpoint security, perimeter security, email, social media, and with each having their own vulnerabilities.
- The customer experience is often neglected so that availability (loss of user hours), integrity of data (loss, corruption, access), threats to life or health, or financial damage to customers, are secondary considerations.
- Insurance premiums, exclusions, exceptions and other costs (e.g. of auditing) make it unattractive for small and medium sized organisations to insure (or ensure?) their business continuity.

1.3 Our focus in 2023

Having recognised the central theme of resilience, in early 2023 we realised that much attention was already being given to societal resilience[5]. The value that we could add as a BCS Working Group was firstly to focus on the resilience of software-enabled systems, secondly to understand learning and capabilities gaps in managing and improving the resilience of (IT enabled) services delivered by organisations, and thirdly to look for ways of bridging them. This involves:

- Understanding the policy landscape;
- Assessing work on organisational resilience;
- Identifying the role of intermediaries in developing learning and capability; and
- Initiating research to explore software risks in infrastructure sectors.

We set up four Task Groups within the SR WG with the aim of identifying actions to improve the resilience of service delivery of software dependent systems. The next section presents the findings and insights from these Task Groups.

[5] Scully, J, Shaw, D., and Powell, D., Operationalising Societal Resilience as a Local Resilience Capability, Report Prepared for the National Preparedness Commission, April 2023 available at https://nationalpreparednesscommission.uk/wp-content/uploads/2023/04/NPC__Operationalising-societal-resilience-as-a-Local-Resilience-Capability_APRIL-2023-1.pdf

2. *Findings and insights*

2.1 Policy and Communications

This Task Group aimed to identify organisations and government bodies with an interest in software resilience to understand their perspectives and establish links for ongoing communication and action. It was led by Gill Ringland, FBCS.

Among our policy and communications initiatives were:

- Discussions with the Digital Policy Alliance.
- Submission of views to DSIT (Department of Science Innovation and Technology) on software risk.[6] A summary of the submission and follow-up correspondence is in Appendix 5. The main points were:
  - Infrastructure services are at risk from software failures, with potentially major impact on the UK economy;
  - Guidelines published by the Prudential Regulation Committee for financial services could be adopted by regulators in other sectors;
  - Government could support sharing of information on service outages to accelerate improvement and work with the insurance industry on this.
- Conference planning for 25th October 2023 at BCS Copthall Avenue offices
- Publication for Z/Yen's Long Finance blog series, The Pamphleteers.[7]

2.2 Task Group – Operational Resilience

This Task Group was led by Paul Wiliams, National Preparedness Commissioner. The question for this Task Group was whether the approach advocated for financial services[8] could be beneficial in sectors with a different regulatory focus. That approach is:

- Identify important business services;
- Set impact tolerances for these services which define how much disruption can be absorbed before intolerable harm is inflicted on the users of the services;
- Undertake regular testing against severe but plausible (which goes beyond probabilistic assessment) operationally disruptive scenarios to identify vulnerabilities;
- Take mitigating action so that services can remain within tolerance.

[6] https://www.gov.uk/government/publications/call-for-views-on-software-resilience-and-security-for-businesses-and-organisations/call-for-views-on-software-resilience-and-security-for-businesses-and-organisations

[7] https://www.longfinance.net/news/pamphleteers/what-is-cutting-uk-productivity/ , https://www.longfinance.net/news/pamphleteers/software-the-elephant-in-the-room/

[8] https://nationalpreparednesscommission.uk/2021/09/operational-resilience-in-financial-services/

Interviews were held with regulators and senior IT professionals with the following findings:

Consequences of software failures that produce service discontinuities range from inconvenience to major financial loss or loss of life.[9] At present, there is some progress in formal reporting of outages, which is a stimulant to avoiding and mitigating the consequences of software failures. This progress is, however, limited across sectors and focusses on high threshold events.[10] This leaves many sectors out of the positive learning that can come from systematic and metric-based attention to software failure risks and their consequences for services.

Incentives do exist to reduce risk exposure and mitigate the costs of recovery. However, these incentives are relatively weak and again vary across sectors and sizes of organisation. Better reporting would help to improve resilience by identifying causal and contributory factors and their relative incidence in operational experience. It would also provide a stronger foundation for extending insurance by providing a stronger actuarial base for assessing risk. It could set in motion the incentives for improvement that accompany a desire to reduce insurance premiums, e.g. encouraging the adoption of auditable preventative and loss reduction practices to lower assessed risk. Aside from improved reporting, better standards for the actions needed to reduce risk and mitigate the consequences of software failures could improve resilience. Failing to adopt such actions means that the extent and scale of consequences of service failure will only grow in significance in the coming years as the operating environment becomes more complex.

Regulation of operators of essential services and broader government actions that have been stimulated by cyber-security threats need a broader agenda regarding service resilience. Threats to service resilience are becoming more intimately tied to software failures because all services increasingly rely upon software. The operational control of systems is becoming less localised as these systems are integrated into larger networks through the internet, including data communications, peer-to-peer interconnections and the internet of things.

This evolving architecture presents a broader and more complex 'attack surface' for cyber-attack risk, but it is also creating greater internal complexity between coupled systems which raises the likelihood of natural accidents. This suggests a broadening of the agenda of regulators and government attention from merely responding to the incidence and consequences of service outages, to putting in place proactive, ideally preventative, monitoring and alerting protocols that are sensitive to the early warning signs of systems failures and outages that occur in everyday operations.

[9] An example of extraordinary financial loss is the case of Knight Capital, see https://www.henricodolfing.com/2019/06/project-failure-case-study-knight-capital.html .

[10] 'High threshold' events, such as those defined by the NIS framework are events that affect large numbers of people and incur large financial losses.

Experience suggests that training/rehearsing for managing outages is worthwhile. Although what happens is seldom what was rehearsed for, a combination of familiarity with response protocols and challenges, and having established communications between key stakeholders is really helpful in mitigating the consequences of outages when they do occur.

More communication about operational experience between companies and with regulators and other parts of government would assist in early detection and mitigation of failure, plus early proactive assessment of failure risks so that consequences are averted, rather than suffered. This communication process also suggests greater communication within government concerning operational resilience experience across operating domains and government departments.

Insight into operational resilience challenges of two sectors important to Scotland, the oil and gas sector and the space sector was gained through discussions with Scotland House.

Both sectors have old, tried and tested technology (akin to legacy systems in the banking sector) that are reliable and manage inherently high risks because it has been developed and implemented by relatively few people, all experts who know each other.

Now there is new technology in both sectors, created by many new organisations that are highly innovative, but lack knowledge and experience. The drivers for developing new products are functionality and speed to market, aiming to make the new technology suitable for many sectors, whilst specialising in none. That is a high risk strategy for national critical infrastructure and inherently risky in sectors like energy and space.

The really high risk is when the old and the new tech are bolted together: since they were built for different operational environments they have different underlying risk assumptions. Comprehensive testing is needed to avoid bizarre outcomes and unpredictable failures.

2.3 Intermediaries

This Task Group was led by Sue Milton, MBCS. The Group researched definitions of service resilience, founding broad areas of consensus – see Appendix 2 for findings on the resilience definition. In addition, this group sought to identify key intermediary actors and to elicit their views on what actions could be taken to improve resilience.

Members of the SR WG have extensive experience as IT expert witnesses in disputes and litigation (in UK High Courts, before Arbitration Tribunals, and internationally) over ‘failed’ or ‘faulty’ software and systems. They have been involved in forensically investigating the defects and deficiencies in their specification, design, development, deployment, operation, maintenance and security, across a very wide range of application domains, and industry, business and services sectors. This has involved analysing and assessing the consequences of such faults and failures which can result in financial claims for restitution, losses and damages, running to the tens and hundreds of millions of pounds. These expert witness experiences have delivered many clear metric-centred lessons and much guidance – 'early warning red flags' – as to how to recognise warning signs, and how to avoid software and systems failures and operational outages, and how to mitigate the consequences.[11]

However, computer technology litigation has concerned only a small fraction of the deployment of computer software and inter-working. Further, only a relatively small proportion of cases reach a public court trial. So it may come as no surprise that there is little general appreciation of the potential severity of the catastrophes, liabilities and consequences that can occur. These issues suggest the need for an increase in attention to the forensics of failure, exercised earlier and at smaller scale than the forensics applied after a major failure has occurred and the consequences become the subject of litigation. The need for such practices has been identified in several contexts, including the chemical and nuclear industries[12].

Consistent with this, we found that, in most sectors, IT professionals; risk professionals, consultants, educators; and C-suite generally were not aware of the demands of service resilience, the importance of systematic forensic analysis or the potential risks of large-scale outages. In one important economic sector, the FCA/PRA's Senior Management Regime and the consultation on operational resilience and third parties[13] has made Boards/C-suites in Financial Services aware of their responsibilities for ensuring operational resilience.

We found little evidence of focus on service resilience in user organisations outside the FS sector or in providers of software. Technology providers focus on functionality and 'speed to market', using functionality as the benchmark around which resilience, safety and security are built. Organisations, whether using their own or 3rd party tech, are still focusing on technical operations, not on the outcomes of interactions between technical operations, internal policies and processes, 3rd-party dependencies and supply chain inter-dependencies, and their effect on service delivery. The assumption is that if the functionality works, then we have resilience. As NATS found out on the 28th August 2023, how

[11] See https://www.cutter.com/article/forensic-systems-analysis-methodology-assessment-and-avoidance-it-disasters-and-disputes and https://barristermagazine.com/computer-evidence-presume-nothing-trust-no-software-or-data-engage-an-expert-costly-just-look-at-the-cost-if-you-dont/

[12] https://riskcenter.wharton.upenn.edu/wp-content/uploads/2014/07/03-01-JP.pdf

[13] https://www.bankofengland.co.uk/prudential-regulation/publication/2022/july/operational-resilience-critical-third-parties-uk-financial-sector

the software reacts at a given moment can be at odds with designed functionality, leading to bizarre outcomes. Software has embedded faults, akin to cyber vulnerabilities, that can be accidentally or deliberately activated.

We found broad consensus that skills need to be more agile. Yes, we need deep specialisms, but we also need people to acquire breadth of skills. Some skills, often deemed specialist, e.g., coding, need to be part of a help desk/support role as means to interrogate data to address queries that are not routine.

Not surprisingly, we also found a lack of understanding on just how good, how mature an organisation is when it came to resilience. Resilience is a broad term and organisations need to know the interdependencies, for example, between their website failure and their reputation. They may be able to repair the website but that might be too late to save their reputation, leading to a company closing.

After contact with a number of intermediary organisations, including Business Continuity Institute (BCI), Chartered Quality Institute, Decoded.com, Deloitte, Worshipful Company of Information Technologists (WCIT), ISACA, we found that the BCI and ISACA are thinking along similar lines and offered to collaborate.

We consider that the BCS and BCI together can provide both materials and access to appropriate resources to deliver learning and capability building to the C-suite, IT professionals and others. With ISACA, we could encourage boards to assess their maturity using a maturity model (see Capability Maturity Model Integration (CMMI) Solutions | ISACA). Further, Scotland House would like to be involved in running workshops.

2.4 Research

This Task was led by Professor Liz Varga, a National Preparedness Commissioner. Research proposals have been made to BCI and UKRI (UK Research and Innovation), but there are no decisions yet.

Topics that demand further research include: identifying and characterising types of software failure types; developing strategies to deal with these types of failure (and which will be contingent on contexts and exposures), strategies for discovery of how software failures interact with other modes of failure, and means of improving testing practice.

Continuous revision is now a leading feature of the operational software environment, particularly as software-based services are extended and further developed. This ever-changing environment opens new vulnerabilities for software accidents and cyber-attack. Many of these vulnerabilities are related to data exchanges between systems that are intended to be interoperable but, in practice, are fragile. Others are related to the real time operation, the increasing array of interconnections and interdependencies, and the growing diversity of users of services. Research on better means of characterising failure types and causes – from both theory and practice are needed.

Because failures are often contingent on context and exposure, research is needed to identify strategies that are effective in identifying vulnerabilities and better assuring 'silent running' systems.

Services based upon software are often meant to control physical equipment through interfaces, each of which can generate or fall victim to software errors or cyber-attack. Empirical and theoretical research is needed to better characterise the risks and to devise means for mitigating failures and outages of inter-dependent sub-systems.

Testing is an essential action for improving operational resilience. The complexities of testing increases as the architecture of service delivery evolves to include many more interdependent sub-systems and software modules. Traditionally, testing was integral to software development and such testing is still relevant. However, testing of operational systems in which interdependencies such as those we identified create combinatorial thickets of interaction between software modules using multiple data communication channels. This raises new challenges and suggests new methods. These new methods include digital twinning of systems, architectures with redundancy and fallback options, and the design of new test suites that can unveil potential failures at the system level. Because of the rapid changes in system architectures, research into new testing methods and the efficacy of testing strategies is a priority for the research community.

3. *Key Conclusions*

3.1 Practices for Improving Service Resilience

Effective methods for improving service resilience exist, but these methods are often poorly understood within organisations and are unevenly deployed across sectors. Our work on raising awareness indicates that awareness of the certainty of failure and the potential scale of consequences is often confined to organisations' IT professionals. To the extent that resilience-related issues – e.g. confidence in the reliability of service or assurance of business continuity – are acknowledged, responsibilities for addressing these issues are most often relegated to these IT professionals. Unfortunately, investing in resilience is often seen as a cost lacking offsetting benefit – either reducing profits or constraining other activities.

The most direct and pragmatic set of methods for improving resilience we encountered came from the practices established for the regulation of financial services and involve four distinct processes:

- Identify important business services;
- Set impact tolerances for these services which define how much disruption can be absorbed before intolerable harm, is inflicted on the users of the services;
- Undertake regular testing against severe but plausible (which goes beyond probabilistic assessment) operationally disruptive scenarios to identify vulnerabilities;
- Take mitigating action so that services can remain within tolerance.

In interviewing IT professionals operating outside the financial services sector, we found similar understandings and a broad alignment of practice with these four processes.

What differed was:

- The degree to which the first two of these practices – identifying key services and establishing impact tolerances - involved collaboration with higher level decision-makers in C-suites or equivalent leadership positions.
- The availability of adequate financial and human resources to undertake regular testing and extend testing to the rapidly evolving technological or organisational complexity.
- The degree of commitment to mitigating action which would involve cost but, importantly, a better appreciation of the scale of possible consequences of service failures.

These differences help to define a strategy for improvement with the following elements:

- Continuing efforts to improve awareness of the risks accompanying the proliferation of software-based services that can be translated into specific agendas for action and methods of accountability that can be shared broadly (e.g. with C-suite and other leaders) in the organisation.
- Identifying incentives to take initiatives that will improve resilience - the most appropriate will depend upon the sector or part of society in which organisations are located.
- The development and deployment of training and certification in resilience planning and improvement for IT professionals in order to augment existing human resources.
- Extending awareness and capability building to ancillary professionals who can assist or provide further motivation for improving resilience planning and practice.

These elements are discussed in more detail in the following sub-sections.

3.2 Awareness and the Role of Leadership Engagement

The "typical" organisation has some characteristics that may not be immediately apparent. These inform the design and delivery of interventions to increase learning and capability and were identified earlier in section 1.2.

The demands of the current operational environments mean that the C-Suite and other organisational leaders need to be able to ask the right questions of the relevant professionals about the resilience of software dependent services, and to calibrate the answers.

The IT literature has many articles with titles such as "The Top 10 Things Executives Should Know About Software". They suggest that the CEO and other executives should understand software, e.g. what it is reasonable to expect software to do, how it is made, how software projects are managed, and how a Web-based service is run. Other approaches suggest that all CEOs should learn to code.

We asked whether all CEOs should be able to do financial book-keeping. The answer was no, but they should know what questions to ask the Financial Director. How can the C-suite ask the right questions on resilience?

We suggest, as an example, the development of a 3 hour C-suite workshop which asks:

- What are the most important (software based) services for us?
- Can we measure resilience using the framework of availability, integrity, risk and material damage?

- Who is responsible for what and when could we seed results?

This is framed as "Five questions to improve Board confidence", and a draft agenda is outlined in Appendix 4.

Education and training for CEOs could also include:

- War gaming aiming to disrupt digital twins (copies) of software systems in use;
- Training in systems thinking;
- Role playing on the effect of software failure: training/rehearsing for managing service outages

The BCI uses a simulation game[14] that could be a good introduction.

3.3 The Incentive Roles of Regulation, Standards, Certification and Insurance

This section is about levers to improve service resilience via regulation, standards, guidelines, certification (of people, processes and software) and insurance.

Regulation

Regulation can be an incentive to address resilience issues and there are important examples of effective regulatory practice.[15] Regulation, however, is not a panacea because it involves costs that may be excessive or, more importantly, may constrain innovation by dictating methods as well as outcomes.

In heavily regulated sectors such as financial services, the concept of organisational and service resilience is mandated via the regulator. The regulation of Regulated Data Service Providers (RDSPs, see Appendix 2) includes the ability to impose financial penalties on data service providers – and defines consequence levels for service outages that could be used as the basis of reporting.

In many infrastructure industries, the current focus is on the shortfalls in investment needed for the physical elements of service continuity and achieving carbon neutrality.[16] With regard to resilience, the priority assigned to physical elements may divert attention from the software-based systems that must also

[14] https://www.thebci.org/training-qualifications/bci-simulation-game.html

[15] As mentioned in the next paragraph.

[16] See National Infrastructure Commission, Infrastructure Progress Review 2023 at https://nic.org.uk/ipr-2023-final/

operate continuously to deliver key services.[17] Deeper consideration of the balance between priorities appears to be indicated.

In considering extending regulation, there are proposals to employ the Registered Data Service Provider (RDSP) regulatory approach[18] of capturing information on outages over a threshold magnitude, to Operators of Essential Services' (OES)[19] service outages. However, regulators warn that for tightly coupled complex systems, data points covering small or "near miss" outages or events (not currently disclosed to regulators) are important to understand the robustness of the underlying system[20].

In the absence of pressure from the regulator, it is up to organisations' Boards to set priorities. Awareness, such as leadership engagement as suggested above may prompt initiatives that make regulation unnecessary.

Standards

The external sourcing of software and other constraints of the operational environment for most UK organisations (see Sections 1 and 2 above) mean that standards certified for individual software components can only marginally contribute to service resilience. The procurement functions of most organisations are unable to assess the characteristics of components or the record of maintenance and upgrades, etc.

For these reasons, we focus on standards for external behaviour of a system – these are listed in Appendix 3. The DORA regulations will affect UK organisations although they are part of EU legislation - we include a description in Appendix 3.

Certification

Certification, alone, is a relatively weak incentive. Nevertheless, it can be an important complement to resilience improvement initiatives, either by providing a standard against which to gauge achievement or a signal enhancing the reputation or credibility of a software component or an organisation. Certification can be applied to products, processes or people.

Software product certification

[17] Physical infrastructure is nonetheless important. Examples include the NHS, https://www.independent.co.uk/news/health/cost-nhs-hospital-buildings-b2202118.html and wastewater treatment, https://www.theguardian.com/environment/2023/jul/04/thames-water-fined-33m-for-pumping-sewage-into-rivers. Physical infrastructure can also be crippled by software error, for example train services https://www.bbc.co.uk/news/uk-england-manchester-64051621.

[18] https://ico.org.uk/for-organisations/the-guide-to-nis/incident-reporting/

[19] https://www.ncsc.gov.uk/collection/caf

[20] Private communication with regulator

Several organisations support software product certification schemes[21], usually aligned with published ISO/IEC or BSI standards.

The US government issued an executive order requiring companies selling to the federal government to take precautionary measures to identify and remediate vulnerabilities in software and to provide government customers with a software bill of materials (SBOM) enumerating various software components, including open-source components, contained in their products[22]. After discussion with senior BCS Members, in our submission to the UK Government's call for views on the impact of software risk and resilience[23], we did not support a similar approach in the UK due to the rapid outdating of such SBOMs as systems evolve and are extended.

Process certification

Several standards exist that improve business processes for resilience. For example, ISO 22301:2019 'Security and resilience — Business continuity management systems — Requirements' sets a framework for reviewing and certifying the preparedness of an organisation to recover from disruptions.[24] This and similar frameworks address the generalities of disruption without a specific focus on the source or consequences of disruption. In this respect, the guidance provided by the NIS framework and its extension in OES sectors is a more direct and detailed approach to addressing software resilience issues.

Process certification also involves an explicit assessment of risk in which identified vulnerabilities are translated into the metrics associated with risk analysis – likelihood and scale of consequence.

- ISO 9001:2015 includes risk analysis as an important step in identifying potential problems and deciding how to deal with them[25].
- ISO 31010:2019 provides guidance on selecting and applying techniques for assessing risk in various situations[26].
- ISO 27001 requires performing a risk assessment as part of implementing an Information Security Management System[27].

[21] See the Appendix 3 on standards.

[22] https://www.lawfareblog.com/open-source-security-how-digital-infrastructure-built-house-cards

[23] https://www.gov.uk/government/publications/call-for-views-on-software-resilience-and-security-for-businesses-and-organisations/call-for-views-on-software-resilience-and-security-for-businesses-and-organisations

[24] https://www.iso.org/standard/75106.html

[25] https://advisera.com/9001academy/blog/2015/09/01/methodology-for-iso-9001-risk-analysis/

[26] https://www.iso.org/standard/72140.html

[27] https://www.iso.org/standard/27001

- ISO 31000:2018 provides principles and guidelines for managing risks that could negatively impact organizations[28].The recommendations in ISO 31000 can be customized to any organization and its context.

People certification
Software is anomalous among the engineering professions in that, in the UK and elsewhere, software engineers do not have to be licenced. We do not know if the licencing of software engineers by Canadian provinces has led to fewer software failures in Canadian software systems or has had any other effects: this may be a useful subject for research.[29] The rapidly developing sector of AI systems software engineering raises novel concerns about service resilience, software failures and their consequences. These include those related to identifiable systems responsibility and accountability, safety, 'bias' and ethics. There is escalating global discussion about the potential need for licencing of AI software engineers (who themselves might come to be replaced substantively by AI-generated software coding) in the context of proposed legislative developments targeted at regulation of AI as a whole.[30]

Similarly, the crucial role of software in delivering services in the UK economy suggests that the profession might review the scope and nature of approaches to cost and the safety/impact of outages, learning from other engineering disciplines.

The BCS and Business Continuity Institute are in discussion about how best to combine resources to provide certification on (digital) service resilience knowledge and skills.

Insurance

In many areas of risk management, the insurance industry creates effective incentives and assists organisations in risk management and reduction (e.g. electrical safety, fire prevention, and health and safety practice.) At present, insurance policies are available to compensate for service outages, primarily through cyber-attack, but also for other forms of business discontinuity.

In practice, insurance claims are the consequence of extended downtime and/or significant financial loss; with a proportion of such claims related to consequences of software and systems failure arising because of alleged 'non-fitness for purpose' defects and deficiencies. In general, insurance aims to cover the direct losses of the insured.

From a policy perspective, social well-being and productivity can be severely affected by the aggregate effect on service users. In general, service users have

[28] https://www.iso.org/iso-31000-risk-management.html/

[29] https://www.jobbank.gc.ca/marketreport/requirements/5485/ca

[30] https://www.techtarget.com/searchitoperations/feature/The-promises-and-risks-of-AI-in-software-development

little recourse for remediating these losses because they are often individually below thresholds worthy of legal representation and few mechanisms exist for aggregating the claims of those affected. Whether and how insurance coverage might be extended to cover loss claims of service users or customers is worthy of further consideration.

Larger organisations in the UK often have cover against cyber-attack, with premiums dependent on implementation of a variety of security practices including multi-factor identification, e-mail access, evidence of robust backup procedures and security of endpoints. The 'cyber incidents' that fall within cyber insurance coverage are defined broadly – beginning with consequences of unauthorised access and detailing various forms of cyber-attack and then extending to system failures more generally.

Cyber insurance coverage has been complicated by exclusions regarding a specific type of attack, state terrorism. A system of "event declaration" in which an event is categorised by size of impact independent of type of attack is being discussed among insurers: Level 1 might be the largest – with government providing financial backup for extraordinary losses, while lower levels would be addressed entirely through insurance premiums. This approach is logically similar to parametric triggers that pay out, if an event happens, an amount depending on the size of the event, with few or no requirements for demonstrating loss.

To summarise, the least complicated incentives for improving software-based service resilience originate from concern with the consequences of service outage and the desire for greater confidence in business continuity. Regulation can be a powerful incentive, but entails costs of demonstrating compliance and may over-specify the means for compliance. By comparison to regulation, standards and certification processes provide relatively weak incentives, but can complement and enhance service resilience improvement initiatives. Insurance can, in principle, provide stronger incentives, but is currently limited to larger organisations, impeded by exclusions, and overly focussed on cyber-attacks, an important, but not the only, challenge to resilience.

3.4 Learning and capabilities for IT professionals

Currently much of the training of IT professionals is on software development methods and tools. There is increasing need for technical skills to deliver services based on complex tightly coupled systems.

Additionally, IT management and professionals should become much more focused on service and business outcomes, particularly software and systems failures and service outages, and their consequential effects on users and other 'innocent' third parties.

Software is anomalous among engineering professions in that, in the UK and many other countries, software engineers do not have to be licenced.

We conclude that developing a Service Resilience qualification would be a major contribution to improving service resilience practices. It would address key current limitations in the human resources available to address resilience issues.

The BCS provides CPD training based on the SFIAplus[31] framework. The Service Management section provides a basis for adding service resilience as a topic.

An example of a course of study that could be delivered through online modules would focus on the following topics (see Appendix 4 for the detailed outline)

1. Operational resilience
2. Software environment for service resilience
3. Existing standards and codes of practice including the role of resilience in risk analysis
4. Description of case study or use output from C-Suite Workshop
5. Responsibilities and accountabilities (internal and external)
6. The elements of failure forensics – how to learn from one's own and other's experience of software failure.
7. Measuring resilience achieved using service outages and associated 'user hours lost' as indicator and metric – ambition (target)
8. Measuring resilience – data gathering, comparisons: forensics after a failure
9. Sources of vulnerabilities
10. Architectures for resilience
11. Testing systems in a 24/7 environment
12. Training/rehearsing for managing service outages: When it fails; what are the likely financial and other consequences to those affected by the failure, and how should the matter of recovery, restoration, remediation and potential compensation to them best be addressed?
13. Drafting and costing a plan
14. Presenting the plan to C-suite

3.5 Government Practice

The Cabinet Office has published a technology road map to 2025[32] with six cross government missions:

- Mission One - Transformed public services that achieve the right outcomes
- Mission Two - One Login for government
- Mission Three - Better data to power decision making
- Mission Four - Secure, efficient and sustainable technology

[31] https://www.bcs.org/it-careers/sfiaplus-it-skills-framework/

[32] Transforming for a digital future: 2022 to 2025 roadmap for digital and data - GOV.UK (www.gov.uk)

- Mission Five - Digital skills at scale
- Mission Six - A system that unlocks digital transformation

Mission One is:
"By 2025, at least 50 of the government's top 75 identified services will move to a 'great' standard, against a consistent measure of service performance. The Central Digital and Data Office will work with partners across government to transform the critical services which are frequently used by citizens, businesses, and civil servants. By 2025, these prioritised services will have great user experience and efficient processes that reduce their cost to run."

We cannot find any definition of service performance or evidence of Government taking a service resilience approach to software failures – The focus seems to be on supply chain issues of software components.

The National Audit Office (NAO) has a remit to comment on the implementation of government policies but does not set policy.

Regulators are directed by government on priorities. We noted that in some cases customer complaints were directed to intermediaries like service Ombudsmen and that the volume and nature of these complaints was outside the regulator remit. It is worth reviewing whether all the possible data sources for identifying and diagnosing threats to resilience are effectively managed which means that their review should be integrated and reported to the regulator. The possibilities for such disconnections are a direct consequence of not taking a service resilience approach to software failure.

3.6 Accounting and Auditing

The BIG 4 and other accounting firms will respond to recent announcements on the role of resilience reporting in Annual Reports (see Appendix 2). There seems at the time of writing to be no proposal on a standard set of metrics for resilience. This could lead to another ESG – type reporting exercise with organisations complaining at the effort involved and with questionable interpretations[33].

Government might consider specifying the NIS framework for reporting on service delivery in Annual Reports. (Appendix 2).

3.7 Other Relevant Professionals in Organisations

Understanding the nature of complex tightly coupled systems and their failures needs to be part of skills development for finance, legal and audit staff.

[33] Delmas, M. A., & Burbano, V. C. (2011). The Drivers of Greenwashing. *California Management Review*, *54*(1), 64–87. https://doi.org/10.1525/cmr.2011.54.1.64 .

In the audit professions – quality, risk, health and safety, and finance – auditors need education and training to measure, and to recommend, ways of mitigating software risk and managing software resilience. They also should put in place protocols for dealing with the consequential effects and potential damages to users and other 'innocent' third-parties. The Y2K crash that never happened was, it is thought, at least partly because auditors were reluctant to sign off "going concern" statements unless the organisation was able to identify their Y2K plan.[34]

There are now many skilled and experienced specialist IT lawyers and firms, but the legal profession at large - for instance as in-house legal team - has few members comfortable with digitalisation. Similarly, procurement staff are less competent to advise on software purchases than in other areas. This is a handicap in contractual negotiations and purchasing of resilient software components or system engineering services. It is also not clear in many organisations what the process is for assessing and advising on the potential (for example, tortious[35]) liabilities of software systems failures, in particular, the consequential impacts on third parties.

Professional associations could have an important role in working with private and public sector organisations including universities. They should encourage relevant skills education including risk analysis, causes of software failure and measurement of the resilience of digitalised systems through service outages, the mitigation of, and remediation for, failures and adaptation of the organisation to their dependence on complex tightly coupled systems, with unpredictable failure modes.

BCS could make the online certification material available to non- BCS members, possibly though a licence with other professional associations.

3.8 Business Schools, other education and training providers

These could be channels for education and training on resilient service delivery, working with the BCS, BIG4 or other professional associations.

However, providers of C-Suite training advise that demand may only follow some sort of catastrophe.

The BCI has developed a simulation game of the consequences of software failure that could be an entry point to discussions with providers.

[34] In addition, of course, major efforts by IT professionals to identify and correct vulnerabilities were another key reason that the predicted potential problems did not emerge.
[35] Due to wrongful acts

4. *Recommendations*

4.1 Potential levers that government could exercise.

Government could promote and support greater transparency and information sharing on failures of digital systems. The sharing should include both breaks caused by cyber-attacks, but also by software accidents. Government departments could take a lead on publishing failure data on their own services, using the NIS framework, viz availability; integrity, authenticity or confidentiality; risk; and material damage to users.

There is an emerging cross-government focus on improving the resilience of the UK economy. We suggest in addition to the above: Working with insurers on catastrophe insurance for both cyber-attacks and software accidents.

Outcome:
Improvement in confidence that government and infrastructure enterprises provide (digitally supported) services at or above specified levels of reliability.

4.2 Role of regulators

The regulator of Regulated Data Service Providers has the remit to require reporting of service outages, as well as data breaches, within the Network and Information Systems (NIS) framework. The regulator has implemented fines for data breaches, and publication of their occurrence, but not for service outages.

We recommend that the regulator of Regulated Data Service Providers publish information on service outages and consider a structure for fines.

We recommend that data on service outages should be published by regulators.

Outcome:
An auditable resilience plan in regulated enterprises.

4.3 Infrastructure sectors

Across infrastructure sectors, new technology is being combined with legacy systems. The drivers for developing new products are functionality and speed to market, aiming to make the new technology suitable for many sectors, whilst specialising in none. These products have different risk profiles to the existing legacy systems: data exchanges are fragile. Other challenges are real time operation, the increasing array of interconnections and interdependencies, and the growing diversity of users of services.

We recommend research on better means of characterising failure types and causes – in theory and practice.

Given the importance of infrastructure to the economy, we recommend that government develop guidelines on operational resilience for infrastructure sectors, based on those published by the Prudential Regulation Committee for financial services.

We recommend that regulators of Operators of Essential Services should consider using the NIS framework, vis availability; integrity, authenticity or confidentiality; risk; material damage to users; to set resilience levels for enterprises.

Outcome:
Visibility and auditability of the role of software in delivering infrastructure.

4.4 Workshop style learning and capability development for C-suite and other leaders

Measures taken by government and regulators as above would start to increase awareness of the likelihood and costs of software failure and the need for resilience – the ability to avoid service outages and to recover from them. In addition, we recommend that Boards and C-suite should engage in conversations on improving their confidence in their service resilience. In smaller organisations, means for facilitating these types of conversations should be developed. An important role for intermediary organisations such as the Business Continuity Institute would be to stimulate and facilitate these discussions within and across organisations.

The workshop outlined in the report provides an agenda for such a conversation in larger organisations that should then be adapted for small and medium sized organisations.

Government should promote such capability development to their own staff and via other channels.

Outcome:
Increased understanding of the value of resilience among senior managers and the potential causes of service outages. The ultimate aim is to integrate knowledge about the value of resilience and means to improve it within organisations and to share experience across organisations sharing this aim.

4.5 Education and certification for IT and resilience professionals

BCS and BCI are planning to develop capability development materials as outlined in the report. This builds on the technical skills of BCS and the business continuity skills of the BCI. The expectation is that certification courses will evolve.

We recommend that BCS and BCI should promote this certification to government, professional associations, and other enterprises.

Outcome:

A growing community of professionals with an informed interest and competence in resilience in the current systems environment.

Appendix 1: For the record

A1.1 Terms of Reference

The Terms of Reference of the Working Group were:

Preamble to the Terms of Reference
There is clear if anecdotal evidence that our economy is increasingly dependent on software and that software failures are occurring in operational systems, leading to loss of service with a range of consequences from inconvenience to major financial loss. We have not found any systematic effort in the UK to collect case studies of failures leading to economic impact and/or their cost to the economy and/or trends which may increase or decrease frequency or impact.

Purpose
The BCS IT Leaders Forum has set up a Working Group (WG) to:

- In the short term, create a network of people and organisations with an understanding of software risks and their potential impact. It will focus initially on the six infrastructure sectors (energy, transport, water and wastewater (drainage and sewerage), waste, flood risk management and digital communications).
- In the longer term the aim is to work with relevant bodies to provide a framework for action to reduce the impact of software failures on the UK economy.

Responsibilities
To create a network and gather data to provide in 2022

- an event for BCS IT Leaders and outsiders, and think-pieces for relevant channels
- a BCS/ITLF White Paper to communicate about software risks to those without an IT background."

The Policy Think Piece itlf-software-risk-resilience.pdf (bcs.org) is the "BCS/ITLF White Paper".

The output from the RoundTable with the National Preparedness Commission held in November 2022 was published as NPC BCS Software-Risk -the-Elephant-in-the-Room Dec-2022-Upload.pdf (nationalpreparednesscommission.uk)

We gave a seminar in the Long Finance series, All Events - FSClub (zyen.com), and published a blog Software – The Elephant In The Room - Long Finance

In Phase 2 we moved from awareness to starting to answer the question, who can do what to reduce the impact of software failures on the UK economy.

In the context of

- much excellent work on operational and national resilience commissioned by NPC,
- a focus by government on resilience and interest in software risk,
- the initiation of standards work in ISO, BSI and the EC
- new initiatives by the Association of British Insurers,

we realised that our value add contribution could be to developing the skills needed to operate in this new system environment.

A1.2 Working Group

The Members of the Working Group are:
Katie Barnes
Colin Butcher
Stephen Castell
Andy Ellis
Tom Gilb
Jon Hall
Lucy Hunt
Adeel Javaid
Neville de Mendonca
David Miller
Sue Milton
Jeff Parker
Gill Ringland (co-chair)
Adam Leon Smith
Ed Steinmueller (co-chair)
Gordon Thompson
Liz Varga
Paul Williams
Yusuf Woozer

A1.3 Contributors

The Working Group of volunteers has been supported throughout by many staff of the BCS. Rashik Parmar and James Woodward have provided encouragement. Support for the submission to DSIT was provided by Arnoldis Nyamande, Martin Cooper and Dale Titcombe. Pat Barlow and her team have enthusiastically picked up the challenge of certification of service resilience professionals. The BCS IT Leaders Forum which commissioned the Working Group has provided a framework throughout, through Chair David Miller and Executive Committee sponsor Jon Hall.

The BCS Information Security Specialist Group - Steve Sands and members of the group, the Quality Special Interest Group -- Margaret Ross and the

Information Risk Management and Assurance Group – Durgesh Gaitonde as well as Professor John A. McDermid -- have provided critiques and helped us to define the problem; and comments through the BCS Community Portal have given extra insights. Adam Leon Smith has provided input and encouragement as both Chairman of the BCS's F-TAG and a member of the Working Group.

The authors of this Policy Think Piece (the Working Group co-chairs) are immensely grateful to the many people who agreed to be interviewed or contributed vias email exchange during the course of this work. They are based in a range of organisations and enterprises with global reach. The interviews were conducted under Chatham House Rule and the individuals cannot be quoted individually; they were Alexander Woods, Anijuli Shere, Arthur Hill, Christine Ashton, Dalim Basu, David Ferguson, David Thorp, David Tynan, Dean Lonsdale, Emma Wright, Gemma Robson, Hank Marquis, Harold St John, James Burns, Joe Little, John McDermid, John Mitchell, Jonathan Pownall, Lisa Emery, Lorna Kirkby, Michael Burgess, Martin Hogg, Philip Wardle, Rachael Elliott, Reza Alawi, Rich Bishop, Richard Chilton, Resham Dillon, Richard Peters, Rob Wirszycz, Sam de Silva , Stuart Okin, Terry Downing, Tom Clementi, Tom Sykes.

Finally we gratefully acknowledge the extensive support of the National Preparedness Commission in and between the 2022 and 2023 RoundTables: Katie Barnes has developed our thinking at several critical points, and Paul Williams and Professor Liz Varga are National Preparedness Commissioners.

Appendix 2: Organisational Resilience, Service Resilience and Metrics

A2.1 Organizational Resilience

Reports commissioned and published in 2023 by the National Preparedness Commission on national and organisational resilience include:

- Resilience reimagined [36],
- Partnering with purpose [37],
- Unlocking value [38]
- Geopolitics, corporate governance and ESG [39]
- Operational resilience in financial services [40]
- Operational resilience applying the lessons of war [41]
- Resilience champion at Board level [42]
- Black, grey and white swans [43]
- Crises, resilience and complex systems[44]

The Prudential Regulatory Authority (PRA) has identified that four activities, which, taken together, can improve operational resilience of firms[45]:

- identify important business services;
- set impact tolerances for these services;
- undertake regular testing against severe but plausible operationally disruptive scenarios to identify vulnerabilities;
- take mitigating action.

[36] https://nationalpreparednesscommission.uk/category/reports/page/2/

[37] https://nationalpreparednesscommission.uk/2021/11/partnering-with-purpose/

[38] https://nationalpreparednesscommission.uk/2023/04/how-to-unlock-value-through-resilience-and-evolve-for-disruption/

[39] https://nationalpreparednesscommission.uk/2022/03/geopolitics-corporate-governance-and-esg/

[40] https://nationalpreparednesscommission.uk/2021/09/operational-resilience-in-financial-services/

[41] https://nationalpreparednesscommission.uk/2021/05/operational-resilience-applying-the-lessons-of-war/

[42] https://nationalpreparednesscommission.uk/2021/05/five-reasons-why-every-organisation-needs-a-resilience-champion-at-board-level/

[43] https://nationalpreparednesscommission.uk/2021/05/black-grey-and-white-swans/

[44] https://nationalpreparednesscommission.uk/2023/05/crises-resilience-and-complex-systems

[45] These are summarised by Paul Williams at https://nationalpreparednesscommission.uk/2021/09/operational-resilience-in-financial-services/

A2.2 UK Government Resilience Statement

Taken from [46]:

"The UK Government has released its long-awaited response to the consultation on strengthening the UK's audit, corporate reporting and corporate governance landscape.

The Government will introduce a new Resilience Statement to improve how organisations identify, manage and report on their resilience risks that are most material to their business. The new Resilience Statement will apply to Public Interest Entities (PIEs) with 750 or more employees and £750 million or more in annual turnover.

The Resilience Statement requirement means companies now need to engage in short and medium-term resilience risk assessment and management, as well as reverse stress testing and reporting for resilience.

What you need to know
There are three key areas that senior management with resilience responsibilities should focus on:

1) Assessing resilience: Companies will need to report on matters that they consider a material challenge to resilience over the short and medium term. Companies will be required to consider a number of specified issues likely to include financial resilience, cyber resilience and third-party resilience amongst others. They will also need to consider any material uncertainties that existed prior to the taking of mitigation actions which help users of the statement to understand the current position and prospects of the business.

In response to this, it may be necessary to update (or design) your resilience controls framework and establish Resilience Board reporting and KPIs to measure this throughout the period ahead of final reporting.

For each resilience issue identified, companies will be required to report on the following in the statement:

- the likelihood of the risk occurring and its impact on the company's operations or financial health if it were to materialise;
- the time period over which the risk is expected to remain, and potentially crystallise, if known;
- any mitigating action the company has put or plans to put in place to manage the risk;

[46] https://kpmg.com/uk/en/blogs/home/posts/2022/06/corporate-governance-reform-resilience-statement.html

- the length of the medium-term assessment period.

2) Performing at least one reverse stress test: Companies will be required to perform at least one reverse stress test – beginning with failure and working back the scenarios which could cause this to materialise. Whilst a regular practice in financial services, this will be a new exercise for many organisations outside of the financial services sector. Companies should ensure they understand their critical business services and processes in order to assess the greatest threats to their resilience. And based upon this, design the scenarios to perform the most relevant reverse stress tests.

3) Reporting and seeking independent assurance: The Resilience Statement will form part of the Strategic Report section of the annual report, and it is important to note that information provided by directors will be covered by the existing 'safe harbour' provision in Section 463 of the Companies Act 2006. The new Audit and Assurance policy (another reform announced by the Government) should set out whether, and if so, how a company intends to seek independent (external) assurance over the Resilience Statement."

The DORA regulations will affect UK organisations although it is part of EU legislation.

PwC describes DORA as follows [47]:

"The DORA regulation applies to more than 22,000 financial entities and ICT service providers operating within the EU, as well as the ICT infrastructure supporting them from outside the EU. The regulation introduces specific and prescriptive requirements for all financial market participants including (but not limited to) banks, investment firms, insurance undertakings and intermediaries, crypto asset providers, data reporting providers and cloud service providers.

DORA builds on previous industry-specific guidelines to define requirements around consistent ICT risk management; comprehensive resilience testing capabilities (including threat-led penetration testing); and third party risk management, ensuring a consistent provision of services across the entire value chain.

The five key topics at the centre of DORA are: ICT Risk Management; Reporting on ICT-related Incidents; Digital Operational Resilience Testing; Management of Third Party Risk; and Information and Intelligence Sharing.

The regulation is unique in introducing a Union-wide Oversight Framework on critical ICT third-party providers, as designated by the European Supervisory Authorities (ESAs).

[47] https://www.pwc.co.uk/industries/financial-services/insights/dora-and-its-impact-on-uk-financial-entities-and-ict-service-providers.html

DORA entered into force on 16th January 2023. With an implementation period of two years, financial entities will be expected to be compliant with the regulation by early 2025".

A2.3 Service resilience of systems dependent on software[48]

It is essential to accept that disruption and failure will occur. Resilience is the feature that allows complex systems to evolve and thrive in the face of external challenge and failures. This means the ability to:

- absorb and survive an initial shock: and
- adapt to continue to deliver services in the changed environment.

Four underlying pillars support response and adaptation:

- Redundancy: the availability of extra resources to maintain service delivery should one fail;
- Diversity: the availability, or ability to develop, multiple pathways to deliver a service should one fail;
- Modularity: an appropriate level of connectedness between components that seeks to minimise complexity in service delivery; and the risk that failure of one component will cause general failure; and
- Prudence: sensible risk analysis and planning (including business continuity) to avoid or decrease the impact of initial shock.

In many organisations, the ability to provide extra resources is constrained by cost; the ability to develop multiple pathways requires skills that may be in short supply; modularity or otherwise is defined by purchase decisions. The approach that we propose to improving service resilience of existing operational services is based on the fourth pillar, Prudence. By risk analysis and planning, resources can be allocated to improving the resilience of the critical services and diversity can be developed for critical services. The C-suite confidence building intervention focuses in surfacing the critical services under the three headings of financial damage to the organisation, reputation damage to the organisation, and consequential impact on customers.

A potential tool for measuring the resilience of services dependent on software focuses on setting ambition levels for the number and scope of service outages, see the section below on Metrics.

[48] This section based on correspondence with David Tynan, Director | Deloitte

A2.4 Metrics for Service Resilience

The NIS Regulations are the ‘Network and Information Systems Regulations 2018’ which came into force on 10 May 2018.

‘Network and information systems’ are any systems that process ‘digital data’ for operation, use, protection and maintenance purposes. Network and information systems play a vital role in the economy and wider society, and NIS aims to address the threats posed to them from a range of areas, most notably cyber-attacks. NIS requires these systems to have sufficient security to prevent any action that compromises either the data they store, or any related services they provide. Although NIS primarily concerns cybersecurity, it also covers physical and environmental factors.

NIS is regulated by sector-specific ‘competent authorities. NIS applies to two groups of organisations: ‘operators of essential services’ (OES) and ‘relevant digital service providers’ (RDSPs).

The Information Commissioners Office (ICO) is the ‘competent authority’ for RDSPs, with a range of powers to enforce NIS, including issuing fines of up to £17 million in the most serious cases.

“RDSPs are organisations that provide specific types of digital services: online search engines, online marketplaces, and cloud computing services. To be an RDSP, you must provide one or more of these services, have your head office in the UK (or have nominated a UK representative) and be a medium-sized enterprise”[49].

The framework and thresholds for capturing information on service outages is:

Parameter	Threshold
Availability	Your service was unavailable for more than 750,000 user-hours. The term “user hour” refers to the number of affected users in the UK for a duration of 60 minutes.
Integrity, authenticity, or confidentiality	The incident resulted in a loss of integrity, authenticity or confidentiality of: • the data your service stores or transmits, or • the related services you offer or make available via your systems. The loss affected more than 15,000 users in the UK.
Risk	The incident created a risk to public safety, public security, or of loss of life.

[49] https://www.gov.uk/government/collections/nis-directive-and-nis-regulations-2018

Material damage	The incident caused material damage to at least one user in the UK, and the damage to that user exceeded £850,000.

OES are organisations that operate services deemed critical to the economy and wider society. They include Communications, Energy, Health, Transport and Water. NIS is regulated by sector-specific 'competent authorities' for Operators of Essential Services.

The National Cyber Security Centre (NCSC) also has two functions: it is the UK's 'single point of contact' (SPOC), as well as the 'computer security incident response team' (CSIRT).

A2.5 Software resilience

Because of its focus on sustaining user functionality, 'software resilience' is often extended to include the quick recovery of systems after a disruptive incident. Scanning the numerous definitions of software resilience leads to:

A resilient software-intensive system can:

- experience a failure in one or more of its constituent components (hardware, software, network, etc.),
- and/or
  - encounter unexpected inputs or external conditions, and/ or
  - be under malicious attack from internal or external sources,
- and yet
  - continue to provide a useful level of functionality to the user, and
  - recover disrupted functions quickly after a disruptive incident.

Appendix 3: Standards and guidelines

There are two potential approaches to standards and guidelines for improving the service resilience of software dependent systems. The approaches are via

- External characteristics of the service delivery *system*
- Characteristics of software components

Standards are currently in place or being developed for each.

A3.1 Standards for external characteristics of the service delivery system

BS 65000:2022 - Organizational resilience. Code of practice. This standard provides guidance and recommendations on what constitutes organizational resilience, the defining attributes and the practical measures that should be considered or can be taken[50].

ISO 22316, Security and resilience – Organizational resilience – Principles and attributes, provides a framework to help organizations future-proof their business, detailing key principles, attributes and activities that have been agreed on by experts from all around the world[51].

ISO 22372 Security and resilience — Resilient infrastructure — Guidelines - This project will result in an International Standard that provides guidelines for developing, implementing, monitoring, and improving infrastructure resilience – end 2024[52].

ISO/IEC 25023 is a standard that provides measures related to the external behaviour of a (software) system such as downtime, incidents, and recovery speed. It does not provide measures related to system flaws that degrade resilience, or to architectural components that enhance resilience[53].

ISO/IEC 27001/2 is a great 'workhorse standard' for anything around IT security. It is an information security management standard which focusses first on understanding the risks, and providing a supporting control framework which addresses those risk (primarily concerned with Confidentiality, Integrity and Availability).

[50] https://www.bsigroup.com/en-GB/about-bsi/media-centre/press-releases/2014/November/Organizational-resilience-standard-published/

[51] ISO - Organizational resilience made simple with new ISO standard

[52] ISOTC292 (isotc292online.org)

[53] ISO/IEC 25023 - European Standards (en-standard.eu)

A3.2 Standards for software structure and components

'Resilience' has not been defined in any of the standard software product quality models. However, there are synonymous quality attributes in most of them.

The dominant model for software and system product quality is ISO/IEC 25010, soon to be revised as ISO/IEC 25010-2. ISO/IEC 25010 includes eight quality characteristics, each elaborated into sub characteristics.

The eight quality characteristics are

- Functional suitability
- Reliability: 'Reliability' is the ISO/IEC 25010 quality characteristic under which the concept of resilience best falls. The sub characteristics under Reliability most aligned with the typical descriptions of resilience are 'Availability', 'Fault-Tolerance', and 'Recoverability'.
- Performance
- Operability
- Security
- Compatibility
- Maintainability
- Portability

The CISQ paper[54] also identifies architectural attributes supporting resilience.

The EC Cyber Resilience Act (CRA)[55] addresses software (and hardware) components (products):

"Rules for the placing on the market of products with digital elements to ensure their cybersecurity;

- essential requirements for the design, development and production of products with digital elements, and obligations for economic operators in relation to these products;
- essential requirements for the vulnerability handling processes put in place by manufacturers to ensure the cybersecurity of products with digital elements during the whole life cycle, and obligations for economic operators in relation to these processes. Manufacturers will also have to report actively exploited vulnerabilities and incidents;
- rules on market surveillance and enforcement.

[54] https://www.it-cisq.org/cisq-files/pdf/How-Do-You-Measure-Software-Resilience-CISQ.pdf

[55] The European Commission has proposed a new Cyber Resilience Act that would introduce mandatory cybersecurity requirements for products with digital elements. Read more: (europa.eu)

- The new rules will rebalance responsibility towards manufacturers, who must ensure conformity with security requirements of products with digital elements that are made available on the EU market.

Note this excludes autos and medical devices as they are covered separately.

The Act will need to be supported by standards work in ISO/IEC Joint Technical Committee 1, Sub-committee 27 and a European Standardisation Organisation."

Appendix 4: Outline Learning and Capability Scoping

A4.1 Draft for C-suite workshop

Learning and Development Workshop on '(Digital) Service Resilience': to improve the organisation's confidence on resilience

Objectives

Board Members/C-Suite Leaders to gain

1. a common understanding of the resilience of the organisation's key digital (software, systems and data network based, driven or dependent) services.
2. increased confidence in how to measure and manage the risk, and potential labilities, to the organisation from service outages due to such interconnected software and systems failures, outages or degradations in performance

Who should participate

C-suite executives and possibly also senior Risk, Insurance, Legal and IT professionals.

Duration

3 hours

Who will lead the session

An experienced Board level facilitator, ideally with service delivery-oriented IT competence.

Target Outcomes

1. Board members gain a common understanding of the resilience of the organisation's key digital (software, systems and data network based) services.

2. Board members have increased confidence in measuring and managing the risk to the organisation from service outages due to software and systems failures, outages or degradations in performance.

3. Board members achieve awareness of the need for Board-level policies and management protocols, possibly involving appropriate enhanced liability

insurance cover, to meet the needs of reasonable remediation and/or reparation to those materially damaged because of such service outages, particularly third parties such as customers and society at large.

Agenda

- Brief – agenda and outcomes

- Brief –
  - why the risk to the organisation of service outages due to software and systems failure is increasing
  - an outline of how to tackle the issue

- Board discussion – which service dependent on software:
  - is the business most financially dependent on?
  - would, in the event of failure, outage or degradation in performance, create the most significant consequences to the business in terms of reputation, brand, market, third-party liability and/or other financial damage?
  - Has the biggest consequences for customers in case of failure?

- Brief – dimensions for measuring impact of outages
  - Availability – user hours lost
  - Ability to recover, repair and restart – quantifiable targets etc
  - Integrity of data
  - Risk to life or health of users
  - Financial damage incurred by users
  - Consequential losses/compensation/redress to third parties – insurance cover availability, premiums etc

- Board discussion – organisation's risk appetite – metrics against each of the three risks (financial, reputation, customer consequences)

- Board discussion – (Board-level?) responsibility within the organisation for measuring and improving resilience of each of the three services dependent on software

- Board discussion – agreement on next steps.

Background reading

To follow

A4.2 Service Resilience – draft for certification syllabus – online course

Each module is 90 minute duration and consists of a mixture of briefing information and student work.

1. Operational resilience
    a. Definitions of operational resilience
    b. Operational resilience as a subset of organisational risk
    c. Role of software in delivery of services
    d. Separating the how (to address) from the what
    e. Factors affecting operational resilience (capabilities and vulnerabilities intro)
2. Software environment for service resilience
    a. Complex tightly coupled system means emergent properties
    b. e.g. Natural accident theory
    c. 24/7 operation (implications for maintenance and testing)
    d. Learning from nuclear industry
    e. Learning from aircraft accident reports
    f. Sources of software failure - problems generated by third-party cloud or software providers, human or device errors or cyber attack, unpredicted traffic
    g. New source of software failure - AI
3. Existing standards and codes of practice
    a. Legislation
    b. Role of regulator
    c. Service (user) view vs system components (supply side) view
    d. The role of resilience in risk analysis
    e. Quality standards & guidelines – supply side view
    f. Quality standards and guidelines – operational performance view
4. Description of case study
    a. See below
    b. Discussion
5. Responsibilites and accountabilites (internal and external)
    a. Defining the most important services (financial, reputational, customer consequences)
    b. Stakeholders for each within the organisation
    c. Convergence/integration of practice at organisational level
6. Failure forensics
    a. Strategies for isolating failure causes
    b. Deconstructing cascade failures (point vs. multi-point sources)
    c. Identifying ephemeral conditions responsible for failure
7. Measuring resilence achieved using service outages as indicator – ambition (target)
    a. NIS framework
    b. Availability

    c. Integrity of data
    d. Threat to life or health
    e. Financial damage to customers

8. Measuring resilience – data gathering, comparisons
    a. Sources of data
    b. Baseline and recording events over time to measure improvement
    c. Cyber attacks and software failures
9. Sources of vulnerabilities
    a. Network, Cloud
    b. Data storage
    c. Endpoint security, Perimeter security (user access control)
    d. Email, social media (as sources of rogue code and as a means to expose vulnerabilities to malign actors)
    e. User error
    f. System/module upgrade
    g. Legacy and acquired sub-systems
    h. what else?
10. Architectures for resilience[56]
    a. Redundancy
    b. Circuit breaker
    c. Graceful degradation
    d. Canary services
    e. Health checks
    f. what else?
11. Testing
    a. Complex tightly coupled 24/7 system - approaches
    b. Use of AI
    c. Chaos engineering
    d. Test suites
    e. Wargame or challenge approaches
    f. What else?
12. When it fails
    a. Plan B – legacy software? manual methods?
    b. Digital twins or other parallel approaches
    c. Stopgap measures
    d. Testing Plan B
    e. Delivery of most important services under Plan B
    f. Damage in the aftermath of failures
13. Drafting and costing a plan
    a. Stakeholders at workshop to draft a plan
    b. Need to engage technical support for the 'how' of doing improvements
    c. Workshop agenda to draft a plan
    d. Outputs

[56] https://www.it-cisq.org/cisq-files/pdf/How-Do-You-Measure-Software-Resilience-CISQ.pdf

14. Presenting the plan to C-suite
    a. Most important services - (financial, reputation, customer consequences)
    b. Financial – targets – NIS framework, curent best guess, cost of improvement
    c. Reputation - targets – NIS framework, curent best guess, cost of improvement
    d. Customer consequences - targets – NIS framework, curent best guess, cost of improvement

Case study
Possible outline case study based on a supermarket: provides the operational context of the case study

Background:
Threshold reporting requirements benchmarks:

Parameter	Threshold
Availability	Your service was unavailable for more than 750,000 user-hours. The term “user hour” refers to the number of affected users in the UK for a duration of 60 minutes.
Integrity, authenticity, or confidentiality	The incident resulted in a loss of integrity, authenticity or confidentiality of: • the data your service stores or transmits, or • the related services you offer or make available via your systems. The loss affected more than 15,000 users in the UK.
Risk	The incident created a risk to public safety, public security, or of loss of life.
Material damage	The incident caused material damage to at least one user in the UK, and the damage to that user exceeded £850,000.

Context:
As an example case, a major UK food retailer Company S has 1,400 stores in the UK as of August 2023. UK supermarkets turnover is around £90 billion pa and Company S has about a 15% share making its turnover on grocerties around £13.5 billion. Company S serves around 4.5 million customers pa of whom 4 million pay by credit or debit card. The typical share of fresh food in total grocery sales for Company S is 40%.

Resilience value and ambition (target)
Financially most important: customer credit/debit card payments
    a. Availability – a 1 hour outage for Company S during business hours would affect 150,000 in-store users.

b. A loss of integrity, authenticity or confidentiality of payment data from 0.4% of Company S's customers would cross the data security threshold

c. Although little direct risk to public safety, public security or loss of life appears likely from issues with credit/debit card payments, it is possible to imagine remote possibilities (e.g. inability of an acutely ill diabetic to purchase food)

d. Material Damage - A one hour outage in credit/debit card payments during store opening hours that led to diversion of 30% of revenue to competitors would result in lost revenue of over £900,000[57]

Reputation risk highest: fresh food replenishment

a. Unavailability of fresh food causes a loss of customer time in procuring substitutes. For example, a half day outage that causes customers to see other sources of supply would quickly meet the availability threshold.

b. Fresh food replenishment is not directly linked to data security

c. A consequence of not replenishing fresh food is that a small percentage of the food remaining that is not removed from sale may go off leading to a threat to life and health.

d. Given that fresh food accounts for 40% of grocery sales, a several hour delay in replenishment would lead to material damage of over £1 million.[58]

Roles to assign to examine measures to prevent or recover from these outages:

Financially important – Financial Director
Reputation important – Legal Counsel and CEO
Customer impact – Sales & Marketing Director

[57] A crude estimate based on 4368 store hours per year and an even distribution of the £13.5 billion sales over store hours.

[58] Using the estimation method in the previous footnote.

Appendix 5: Summary of submissions to DSIT

A5.1 Letter to MPs

The following letter was sent to constituent MPs of several WG members as well as others identified as having a potential interest.

Dear xxxx

The BCS has recently answered the DCMS DSIT "Call for views on software resilience and security for businesses and organisations". In the answers we identified potential areas for government intervention. We thought that it might be useful to also write to you as we are clear that improving software resilience and security could have a positive leverage effect on UK productivity and growth.

As background, it is important to note that software services are now delivered through complex tightly coupled systems, with unpredictable failure modes. This requires new approaches. Being known for reliable digital services in this new complex environment would add to UK's competitiveness. Our report itlf-software-risk-resilience.pdf (bcs.org) describes the situation in more detail, and the National Preparedness Commission report NPC_BCS_Software-Risk_-the-Elephant-in-the-Room_Dec-2022-Upload.pdf (nationalpreparednesscommission.uk) concluded that "The software element of digital systems failure is a COST TO ECONOMY AND SOCIETY which will only increase as software has become a utility, is in wider usage, and more vulnerable to failure."

In the answer to the call, we have identified three complementary potential ways forward:

1. The BCS is currently undertaking a project targeted at reducing the software risk and improving the resilience of the UK's digital services. Our focus is on the resilience of operational digital systems in infrastructure sectors because:

- Failures in infrastructure services would have dramatic negative effects on the rest of the economy including impeding growth and reducing productivity
- The regulatory regimes of infrastructure sectors in the UK are oriented towards keeping costs to consumers down, rather than continuity of service or "keeping the lights on".

We are exploring whether guidelines for infrastructure sectors could be adapted from those published by the Prudential Regulation Committee for financial services.

2. We make recommendations on information sharing, so that organisations can make more informed decisions. Government could promote and support information sharing on failures of digitalised services. This would prompt Boards to take responsibility for resilience of the services supplied by their organisations. The sharing should include both breaks caused by cyber-attacks, and by software accidents. Government departments could take a lead on publishing failure data on their own services, using a framework based on that proposed for Regulated Data Service Providers by the Network and Information Systems Directive and Regulation, which addresses availability; integrity, authenticity or confidentiality; risk; material damage to users.

3. There is an emerging cross-government focus on improving the resilience of the UK economy. So, we suggest in addition to information sharing: Working with insurers on catastrophe insurance for cyber-attacks and software accidents. Insurers already play a key role in encouraging improved practice in safety and resilience.

We hope that in writing to you we can alert you to the issue of potential failures of digital services due to software, their impact on our economy, and measures that might be taken to improve resilience. We are happy to support in any way.

A5.2 Letter to Minister (Viscount Camrose) by invitation

We are following up on the letter from Viscount Camrose (attached) to inform you of work at the BCS (formerly the British Computer Society) related to cyber security. During the last eighteen months a Working Group of the BCS has been raising awareness and investigating actions that might improve the resilience of software-based services in the UK. We have followed the discussion of cyber security issues with interest as they are an important facet of the resilience of software-based services.

As IT professionals we are concerned that the cyber-attack or threat discussion threatens to overshadow a vitally important component of our information infrastructure – the inherent or intrinsic risk of software (and hence system) failure. All software contains errors and the construction of ever larger and more inter-dependent systems based upon software amplifies the potential for catastrophic system failures. These failures do not require a 'bad actor' (from overseas or at home); they are a feature of complex software systems. Cyber security and software-based service resilience discussions need to acknowledge and address the intrinsic risks of software. We are actively working this year to propose specific actions to improve the resilience of digital service providers including managed service providers who are often not UK companies.

There are three important overlaps between our work and the cyber-security efforts being promoted by NIS regulations.

1. The need for cyber-incident reporting. In order to make better corporate and government policy, cyber incidents that arise from 'accidents' as well as cyber-attacks by external parties should be recorded to gauge the nature and extent of problems that reduce resilience and productivity.
2. The complexity of software-based services is increasing, a key observation in the effort to expand the definition of organisations providing managed services. System complexity enlarges the scope for both cyber-attack and for accidents. The two sources of failure should be considered in parallel.
3. Discussion of software-based services are often framed as arising from threats external to organisation while the improvement of the resilience of these services requires attention to all of the sources of system failure and the differences between protecting against external threat and improving practices within organisations to prevent or deal with intrinsic software failure.

Yours etc