

The Safety-Critical Systems Club Newsletter

Safety Systems

Vol 33 No. 3 – Oct 2025

HOW TO BE AGILE ... SAFELY

The key role
of Product
Owner for
safe Agile

AGILITY TRAINING

The Lean
Learning
Machine

MIXIN
THINGS UP!
Our ontological
odyssey continues!



thescsc.org



SCSC Publication Number: SCSC-206

While the authors and the publishers have used reasonable endeavours to ensure that the information and guidance given in this work is correct, all parties must rely on their own skill and judgement when making use of this work and obtain professional or specialist advice before taking, or refraining from, any action on the basis of the content of this work. Neither the authors nor the publishers make any representations or warranties of any kind, express or implied, about the completeness, accuracy, reliability, suitability or availability with respect to such information and guidance for any purpose, and they will not be liable for any loss or damage including without limitation, indirect or consequential loss or damage, or any loss or damage whatsoever (including as a result of negligence) arising out of, or in connection with, the use of this work. The views and opinions expressed in this publication are those of the authors and do not necessarily reflect those of their employers, the SCSC or other organisations.

© SCSC CIC 2025 except where explicitly stated for specific articles, reports or images.

Front cover: AI generated image
 Editorial: The man engine of Dolcoath mine, John Charles Burrow, Public Domain
 Editorial: The Grenfell Miner’s hat, © Paul Hampton, licensed as Public Domain
 In Brief images may be subject to copyright – contact the SCSC Newsletter Editor for details.

Contents

WELCOME

Editorial

Opening words from the SCSC Newsletter Editor.

3

In Brief

Recent system safety news items from around the world.

4

FEATURES

Product Owners – The key to Safe Agile

Bill Blackburn discusses how key the Product Owner is for safe Agile

5

The Lean (but not mean) Learning Machine

Paul Lister discusses applying Agile approaches to make training more effective.

13

Mixin Things Up!

The SCSC Ontology Working Group provides the next episode in a series explaining ontological concepts.

17

SCSC Crossword

The famous annual SCSC Crossword returns once more!

46

REPORTS

Seminar: How Safety Culture has to Change with AI

Highlights from this SCSC seminar discussing how Safety Culture is impacted by AI.

23

Seminar: Safe Agile Developments

Highlights from this SCSC seminar discussing the application of Agile approaches on safety projects.

29

Tech Report

Paul Hampton and Dewi Daniels report on their experiences being pulled by the iconic Flying Scotsman Steam Engine

43

60 Secs with ... Nancy Leveson

Nancy answers some quick-fire questions on system safety and the future.

37

GROUPS

Working Groups

Details of the SCSC Working Groups.

49

SCSC Steering Group

Who's who in the Steering Group.

57

EVENTS

Calendar

59

Events Diary

60

SSS'26 Programme

34th Safety-Critical Systems Symposium



10 - 12 February 2026
The Milner York, UK

www.scsc.uk/ssc

Bookings for the Safety-Critical Systems Symposium (SSS'26) are now open.

scsc.uk/ssc

Editorial

This month marks the anniversary of the Cornish Levant mining disaster that took place over 100 years ago in 1919. The accident occurred when the “Man Engine” carrying about 100 miners broke, leading to 31 fatalities and many injuries. The Man Engine was installed over 60 years earlier to improve the speed at which miners could be conveyed to and from the mine. The mine was so extensive that it could take 90 minutes by ladder to reach a working location. Until the accident, it had also improved the safety of miners – previously, fatalities from falls from ladders were significant.



The Man Engine was ingenious in its simplicity – a series of inter-connecting wooden rods with platforms every 12 feet. The powerful beam engine housed at the top simply moved the whole interconnected assembly up and back down 12 feet. Men would step off and onto platforms as the array travelled in their desired direction. This reduced transit time to 30 minutes and meant the miners were safer and of course less fatigued for their work.

After 60 years of mostly unblemished operation, what precipitated the accident? From an engineering perspective, it was simply that the main iron strap holding the top rod of the linkage broke causing the whole assembly to collapse. Digging a bit deeper though, we'll see there were several other factors that will be surprisingly familiar. Firstly, the rods on the engine had been extended beyond the initial design intent and clearly had a number of single points of failure. Because of the First World War, and underfunding of the mine operations, there had been no maintenance of the equipment for five years. Many miners and the beam engine driver had reported concerns about the Man Engine's safety with it making unusual noises during operation, presumably all going unheeded by management. After the accident, the rescue response was poor, mostly undertaken by the miners themselves. Some waited over two days to be rescued and operations were hindered by the lack of carbide lamps.

Despite obvious failings in safety management (eg. lack of funding, maintenance, incident reporting, emergency planning) the subsequent inquiry concluded that no-one was to blame and it was a production flaw in the metal strap that led to the metal fatigue. With no one to blame, the victims or their families received scant compensation. Surely, in our modern world of mature safety engineering disciplines, standards and regulations these poor practices should be firmly in the past?

It was however poignant that one of the miners who died that day was called Grenfell. I'll leave it to the reader to see how many parallels they can draw between the events of 1919 and the Grenfell Tower disaster almost a century later. Are we learning from the past? Perhaps not as much as we should.



I have pleasure in presenting a couple of articles on Agile practices as a follow up to the very interesting and engaging Safe AI Seminar in May this year (see page NN). We also have an article from the Ontology Working Group, the first of a series, expanding on how we can formally specify concepts. Our 60 second interview is with Nancy Leveson, Professor at MIT.

Paul Hampton
SCSC Newsletter Editor
paul.hampton@scsc.uk



In Brief



Anger in Australia after telecom outage linked to deaths

Australian officials have promised telecommunications giant Optus will face "significant consequences" over a systems outage linked to multiple deaths. The incident last week left hundreds of people across more than half of the country unable to call emergency services for 13 hours.

Optus - one of the country's two major providers - says at least three people died as a result, and its chief executive has apologised to their families and the public for the "completely unacceptable" failure.

bbc.co.uk



Flying cars crash into each other at Chinese air show

Two flying cars crashed into each other at a rehearsal for an air show in China. The Xpeng AeroHT vehicles collided in mid-air, with one catching fire during landing, the company said in a statement to Reuters. The company said people at the scene were safe, but CNN reported, one person was injured in the crash, citing an anonymous company employee. [bbc.co.uk](https://www.bbc.com/news/technology-65103401)

Dating safety app Tea suspends messaging after hack



Messaging has been turned off on the women's dating safety app Tea, following a hack which has exposed thousands of images, posts and comments.

In a new statement the company said: "As part of our ongoing investigation into the cybersecurity incident involving the Tea App, we have recently learned that some direct messages (DMs) were accessed as part of the initial incident." [bbc.co.uk](https://www.bbc.com/news/technology-66111111)

The AI Was Fed Sloppy Code. It Turned Into Something Evil

Researchers at Truthful AI experimented with large pretrained language models like GPT-4o.

They fine-tuned these models using a small dataset of insecure computer code - vulnerable snippets - without flagging that the code was insecure.

Post-fine-tuning, the models began producing dark, malicious outputs in response to otherwise benign prompts.

quantamagazine.org

Error leaves 55,000 diabetes patients needing new tests

Errors by machines used to diagnose diabetes mean at least 55,000 people in England will need further blood tests, a BBC investigation has discovered.

Some patients have been wrongly diagnosed with type 2 diabetes and even prescribed medication they don't need – and there could be more people affected, say NHS England. [bbc.co.uk](https://www.bbc.com/health/longer-living/diabetes-misdiagnosis)

004558	MSA, Beta Subunit, Quant	84733	06/02/18
004416	MSA, Beta Subunit, Quant	84733	06/11/18
001925	HDL Cholesterol	83718	06/04/18
001453	Hemoglobin A1c	83536	07/05/17
144286	HIV Antibody w/ reflex to high	86709	06/23/16
144473	HIV Screening and Diagnosis	86709	06/11/17
000510	HIV Surface Antigens	87346	00/14/18
144050	HCV Ab RPH to Quant, P, PCR	86825	06/23/16
083935	HIV Ab p24 Ag w/ RPH	87360	06/23/16
140836	Hepatitis Virus Urine	83012	06/23/16

Product Owners – The key to Safe Agile



Bill Blackburn develops the theme of Agile in the context of high integrity and safety systems, exploring the role of the Product Owner. This role is pivotal in Agile and is relied upon to orchestrate and oversee the build aspects of the lifecycle, ensuring that requirements are reflected in the product and that relevant stakeholders are represented.

The Safety-Critical Systems Club (SCSC) ran an event – [Safe Agile Developments](#) – on 1st May 2025, held at the Wellcome Collection in London. See page 29 in this edition of the newsletter for the event report. This one-day seminar looked at techniques and approaches for managing Agile developments for safety-related systems.

This seminar was for those introducing or using Agile techniques with high integrity and safety-related developments, in e.g. Healthcare, Policing, Transport, Government or Defence areas. It explored the balance between agility and the demands of systems operation in situations where safety and business criticality requirements applied. The discussions touched on the challenges of blending the safety engineering disciplines with an Agile approach. Speakers and participants explored how the Agile manifesto can be interpreted and adopted/adapted.

To give its full title, the Manifesto for Agile Software Development” [1] has its origins in 2001 when a group of experts sought to establish a practice-driven philosophy for more effective software delivery and came up with their four key values. This is often quoted by agile evangelists, however the manifesto has much more to say elaborated through a set of 12 principles.

We are uncovering better ways of developing software by doing it and helping others do it.
Through this work we have come to value:

Individuals and interactions over processes and tools
Working software over comprehensive documentation
Customer collaboration over contract negotiation
Responding to change over following a plan

That is, while there is value in the items on the right, we value the items on the left more.

A thread running through the Agile Manifesto is about having the right skilled and empowered people involved throughout and collaborating in self-organising teams. Done poorly, with Agile evangelists proclaiming the benefits of agile vs “old school” waterfall can result in (FR)Agile delivery and chaos. Not a good recipe in the world of high integrity and safety.

So how can appropriate working practices be established that provide a level of confidence for the safety community?

Evolution of Agile

My paper called Scaled Agile for High Integrity Services [2], at the SCSC seminar talked about the evolution of Agile from the groundwork done by Tom Gilb in the 1980s through to Scaled Agile Framework that has become a well-respected set of tools and techniques covering the range of needs across portfolios, programmes and projects.


Gilb's work centred on the concepts of evolutionary delivery [3], and he also complemented this with his detailed book on Software Inspection [4]. It might appear odd to mix the flexibility of agile and the rigour of inspections however Gilb's approach was more about programme level evolutionary delivery rather than software development where the Agile manifesto tends to focus. In a large space systems project I found that the discipline of formally run inspections was extremely helpful to the development team who were able to use the cadence to drive their build iterations.

During the Safe Agile Development seminar, we discussed the various roles that are common across the Agile methods and in particular the Product Owner. With the focus on incremental delivery of a “product” towards a releasable state, the Product Owner ensures that requirements are reflected in the product backlog and oversees the sprint refinement sessions and delivery of product increments.



The Product Backlog contains user stories that translate functional and non-functional requirements into descriptive explanations and acceptance criteria. These are prioritised and elaborated with stakeholders through analysis and design sessions. Once features can be demonstrated, 'show & tell' sessions allow early feedback and refinement. Requirements are typically documented as functional or non-functional statements or outcomes to be achieved. They should avoid prescribing solutions or unnecessarily constraining design and build, though wider architectural principles may provide context.

User Stories are commonly captured in web-based collaborative tools such as Azure DevOps or Jira, which are industry standards used by business and technical teams to support the cycles of build, integration and deployment, as illustrated in the diagram below taken from Microsoft DevOps documentation [5].



Nothing fundamentally unsafe so far?

Platforms and Products

Building a product entirely from scratch is very different from building a product on an existing platform. Platforms need architectural and design skills that are wider than the skills needed for product increments once a platform exists, which can take years to mature. Most sectors including those where safety is paramount such as aerospace and automotive have taken a platform-based approach for many years. Modular and component-based products are commonplace to promote reusability and the advantages that can bring around cost management.

Recent examples exist in the electric vehicle domain where manufacturers are investing heavily in platforms. For example, in May 2025 Volkswagen [6] announced, "The Scalable Systems Platform (SSP) will represent the next generation of an all-electric, fully digitalized, highly scalable mechatronics platform based on a uniform system architecture." Platforms can underpin a whole range of products and not just those of one manufacturer but shared with collaboration partners.

This has some downsides from a safety viewpoint, if for example a component is used in a different way or if there is an inherent issue not apparent in one product, but manifests itself over time across the product range.

The Toyota accelerator pedal used across numerous vehicles springs to mind and resulted in worldwide vehicle recalls in the period 2010-12, one of the largest in automotive history. Along with reputational damage to one of the world's largest corporations [7]. In this case, the safety issue was unintended acceleration.




Product Manager / Product Owner

Product Managers handle the overall product lifecycle often from inception to end-of-life and that generally follows an established path. This role is fairly commonplace. For example, in 1998, as part of its initiative ambitiously titled Smart Acquisition, the UK Ministry of Defence introduced an approach termed Integrated Project Teams (IPT). This was part of a new way for defence equipment acquisition. These weren't really dealing with projects and were in fact Integrated Product Teams, where the IPT lead played the client-side role of a Product Manager with supplier-side equivalents. IPTs eventually operated across the whole lifecycle and in fact were pivotal in through life support for equipment. A helpful analysis is provided in the NAO review Implementation of Integrated Project Teams [8].

The Agile Product Owner can be regarded as a Product Manager, or potentially one of a handful of roles working with a product team, depending on the nature of the system or service.

It is best to think of the Agile Product Owner as the responsible owner for the lifecycle stage within which the build is undertaken, through to acceptance and transition/release into operational use. Following which, the Product Owner handles the ongoing evolution, cycling through product increments (or not if the budget has run out). Generally the Product Owner would sit on the business or client side with counterparts in the in-house or outsourced delivery team.

During build, the Product Owner is the person who orchestrates the process by managing the product backlog and getting the right stakeholders involved at the appropriate points. The Product Owner works collaboratively with the Scrum Master (in Scrum terminology, the scrum team are a self-organising delivery team) who oversees and coordinates the resources within the scrum team or other specialists as required.



In the Scaled Agile Framework (SAFe) [9], the definition is:

The Product Owner (PO) is the Agile team member primarily responsible for maximizing the value delivered by the team by ensuring that the team backlog is aligned with customer and stakeholder needs.

Reads like a nice simple definition but lurking within are wide implications and demands on skills. Furthermore, recruiting the right Product Owner can be a challenge, particularly so if looking for one already with a safety engineering background. Ensuring that a Product Owner coming into an organisation dealing with high integrity or safety systems is familiar with the disciplines and requirements. Realistically the Product Owner needs to be skilled at awareness level across a range of domains and safety aspects would need to be treated in a similar manner. Keeping in mind that the Product Owner would be relying upon those skilled and qualified people with safety responsibilities and manage them as stakeholders to be involved at the right time.

Requirements Traceability and Safety User Stories

From a safety engineering viewpoint, the relevant safety requirements (standards, policies, etc) need to be identified and agreed from the outset. The necessary skilled resources need to be lined up and available to provide direction and potentially act in verification & validation roles. This is within the Product Owner responsibilities.

Agile methods rely on “User Stories” and these are written, reviewed and approved then used to direct build activities whether those are manufacture or other form of product development. User Stories are generally thought of as functional specifications written as acceptance criteria usually in the structure:

“Realistically the Product Owner needs to be skilled at awareness level across a range of domains and safety aspects would need to be treated in a similar manner”

“As a [type of user], I want [goal] so that [reason/benefit]” Taking the earlier accelerator pedal scenarios: “As a **driver** I want to be able to **control engine power** so that **the vehicle speed can be varied**”.

User Stories can equally be written as non-functional requirements and instruct how aspects such as safety requirements are to be implemented using an equivalent structure:

“As an automotive safety engineer, I want the engine control system to comply with ISO 26262, so that risks X, Y, Z in the Product Omega FMECA are mitigated”

More elaborate acceptance criteria would be required, but that's done through refinement sessions as analysis and design progresses. User stories are progressed through the backlog and the Product Owner would need to ensure that those representing essential requirements are prioritised alongside the functional aspects.

Functional user stories tend to be implemented through sprints where functionality is configured/built, and the first stages of testing are performed.

Product Assurance

Agile methods tend to skip over the broader aspects of Product Assurance and the disciplines of verification and validation. Testing within sprints is performed and the integration and build activities ensure that those aspects are assured. This is where a Product Assurance Manager needs to step in, working alongside Product Owner or part of the product team.

My experience where the role is scoped as a “Test Manager”, their perspective is subsequently somewhat narrow, and the resulting test-based assurance is not sufficient for high integrity products. Product assurance on the other hand is a set of carefully planned and executed activities running in parallel with the build/configuration, integration and release cycle that runs through the core of Agile methods.

“Agile methods tend to skip over the broader aspects of Product Assurance and the disciplines of verification and validation”

Product Assurance standards are particularly well-defined in the aerospace sector for example, ISO 22893 [10], and these can be readily applied alongside Agile build as found in research [11] that examined how agile attributes and principles can help achieve high European Cooperation for Space Standardization (ECSS) compliance levels.

A Few Takeaways:

This article has been based on personal experience across a range of sectors, while keeping an eye on the ups and downs of Agile project delivery.

- Be wary of evangelists, especially those who pin up sheets proclaiming waterfall bad / Agile good
- Agile requires highly skilled and empowered people who know how to collaborate where the Product Owner takes a leadership role and needs at least awareness-level of safety knowledge
- Platforms upon which to be Agile are vital. Business, information and technology architecture is needed to create a platform upon which product increments can happen
- With the right pre-requisites in place, safety engineering can be applied. The Agile team need to appreciate the demands of safety, and the safety engineer needs to feel comfortable with Agile approaches, adapting techniques where appropriate

References

- [1] "Manifesto for Agile Software Development" Jim Highsmith et al, Agile Alliance <https://Agilemanifesto.org>, February 2001
- [2] Scaled Agile for High Integrity, Bill Blackburn SCSC Safe Agile Seminar, May 2025, <https://scsc.uk/r3107>
- [3] Evolutionary Delivery vs the "Waterfall Model", Tom Gilb ACM SIGSOFT Software Engineering Notes Vol 10 No 3 Jul 1985 <https://dl.acm.org/doi/pdf/10.1145/1012483.1012490>
- [4] Software Inspection, Tom Gilb and Dorothy Graham, Addison Wesley, 1993 ISBN 0-20163181-4
- [5] Microsoft DevOps <https://learn.microsoft.com/en-us/devops/overview/>
- [6] Volkswagen Future Plan – Scalable Systems Platform <https://www.volkswagen-group.com/en/articles/volkswagen-future-plan-brand-gives-first-glimpse-of-entry-level-electric-model-at-works-meeting-18997>
- [7] FBI press statement on Toyota Motor Corporation <https://archives.fbi.gov/archives/newyork/press-releases/2014/criminal-charge-against-toyota-motor-corporation-and-deferred-prosecution-agreement-with-1.2-billion-financial-penalty>
- [8] Implementation of Integrated Project Teams, National Audit Office 2022 <https://www.nao.org.uk/reports/ministry-of-defence-implementation-of-integrated-project-teams/>
- [9] Scaled Agile Framework (SAFe) Product Owner, <https://framework.scaledAgile.com/product-owner> Scaled Agile Inc.
- [10] ISO 22893:2022 Space systems — Software product assurance (SPA)
- [11] An Agile Compliance Framework for the European Cooperation for Space Standardization Al-Mhdawi et al, 2023 IEEE Aerospace Conference

Image Attribution:

leading image: AI generated by ChatGPT

Agile manifesto from <https://agilemanifesto.org>

DevOps Image Dirk Wouters on Pixabay

Pedal adapted from: ID 105521154 | Car Pedals © Asawin Kanakasai | Dreamstime.com

Bill Blackburn, Process Renewal Group Limited and SCSC Member

Bill has a background in mechanical engineering and worked across aerospace, manufacturing, defence and other high integrity sectors for over 40 years. He is presently engaged in advisory roles with clients in UK policing and capability development in Europe. He has extensive experience in business and information architecture, with a focus on process and data management.



For Everyone Working in System Safety

Safety-Critical Systems Club (SCSC)

New Developments in Rail Safety Systems

Thursday 4 December, 2025

Eurostars Book Hotel, Munich, Germany

This seminar will feature technical speakers presenting new technologies for improving safety on the railways.

The seminar will be useful for rail practitioners working in safety engineering including Safety Engineers, Safety Consultants and Safety Managers.

Speakers include

Jelena Frtunikj, NVIDIA

Hardening Machine Learning (ML) Inference for Safety-Critical Applications: From Static Libraries to ML Graph Compilers

Lennart Kilian, Siemens Mobility GmbH

Safe.trAIIn: AI based obstacle detection for trains - risk analysis, functional scalability, safety on system level and lessons learned

Bernhard Nölte, TÜV SÜD Rail GmbH

Challenges in the safety demonstration and assessment of a fully digitized rear-view mirror system

Carsten Thomas and Philipp Jaß, HTW Berlin


N-Version Architectures for Railway Track Segmentation with Deep Neural Network

Volker Wegener, DATADEV

Railway geodesy - 5 mm at 80 km/h: A modern method for efficient and precise surveying and digitalisation of railway infrastructure without any safety risk

www.thescsc.org

The Lean (but not mean) Learning Machine



How often have we undertaken training as a corporate tick box exercise? Perhaps sometimes months before it is needed or actually ever used at all? Have the curricula seemed bloated with patchy relevance to the job at hand? Drawing from the Agile Manifesto for product development, Paul Lister applies the ideas to establish the principles for improving the effectiveness of training by "Training Lean".

Introduction

An important part of any Safety protocol is training. This is often undertaken at an onboarding stage for an employee but is this the best thing to do? Over the years I have been involved in training in both the role of teacher and student and have found that it can often be treated as a tick box exercise.

So how can Agility help to make it more effective?

A significant part of the formulation of the Agile Manifesto [1] and its twelve principles comes from Lean Thinking [2] and we are going to use some of the ideas from this methodology to examine training, in particular utilising the seven Lean Wastes.

Wasting away

Let's look at these wastes and see what they can teach us about training.

Over Production

I've listed this one first because it is probably the biggest 'training crime' perpetuated by organisations. Overproduction does what it says on the tin, making something before it is needed. In training terms this often means that people are taught to do something that they either don't need or will not need for a considerable amount of time. I experienced this in my early career as a software developer on a large piece of financial software. I was learning about every single part of the system which took months. When I joined a team I ended up working on the same part of the system for three years. Lean is a 'pull' system i.e. you do the work that is needed now and then pull the next piece of work that is needed when you finish. Training should be the same. Understand what people need to be trained in now and then train them again later if needs be.

Overproduction often happens due to 'checklist' thinking, training departments being able to say that employees have all learned everything. But learning only comes with reinforcement. If the knowledge isn't used then it will eventually be forgotten.

"Lean is a 'pull' system i.e. you do the work that is needed now and then pull the next piece of work that is needed when you finish. Training should be the same."

Waiting

Sometimes training is scheduled for months ahead of time for individuals, especially when they are new. This is normally due to the training being undertaken by a different department or the dates of an external course being the first available. But pragmatism tends to ignore this and some sort of 'on the job' training will take place so the employee can be effective in terms of the organisation's needs. Often by the time they participate in a course they already know the bulk of the knowledge and skills that are required, making the course redundant or at the very least economically wasteful.

Over Processing

In normal production this would be adding features or technically finessing things that do not really satisfy the customer. In training terms this means adding knowledge to core training that is not needed. Lean is very much about prioritising what is needed first. A student learning CPR for instance does need to understand the mechanics of how to revive a person. They do not need a history of how the method was developed or a detailed explanation of the anatomy of the respiratory system.

Inventory

This is like overproduction but tends to reflect more on the trainers rather than the trainees. Often courses will be built which are based more on the expertise of the trainers and what they know rather than the need of the student. These courses build up as inventory but are seldom used. Similarly, the existing course will not often be audited and hence contain incorrect or redundant information that will not come to light until said information is utilised. Training should be designed for need not assumption of need. The lean waste '**Defects**' can fall into this category.


Motion

Within Lean this normally means the inefficient movement of people or machinery. It is to a lesser extent a problem today as more training is conducted online. But if a course is live and scheduled some distance from the student's office/home then this can lead to fatigue before the individual has even started learning.

The final lean waste or 'Muda' is '**Transportation**', unnecessary movement of product or materials but this does not easily fit into the training analogy.

In recent years **Under-utilised talent** has been added to the list of Lean Wastes. This is basically not allowing employees to interact with the process adding their own ideas to the workflow in order to improve it. Training that does not allow discussion, feedback or questions can effectively lose a layer of information and education which was not known at the start of the session.

Train Lean !



So how can Lean help provide better training? The lean thinking model has five basic steps:

Identify Value

Instead of creating generic courses, audit your potential students. Decide what training would be useful first, then second etc. building a list akin to a product backlog in Scrum.

Also understand the context in which these courses should be run. If I am running a Scrum Foundational course I will teach it in a very different way for Project Managers, that I would developers, that I would people with some experience in the framework.

Map the Value Stream

A value stream is a timeline where you understand which pieces of a workflow directly contribute to the outcome you want, and which do not (normally classed as one of the wastes above). Here, depending on the type of students, you would modularise a piece of training and then order it depending on the subject. To go back to our first aid example, you would teach CPR¹ first and sprains last. With other subjects it might be smarter to give an overview first and then drill down into the detail. (I created a 'COMET' system (Concept, Methodology, Tools) exactly for this type of system.) The important thinking here is that if something occurred that meant the student had to leave the course they would have learned 'the important bit'.

Create Flow and Establish Pull

These are normally two separate concepts, but they would be combined in a training context. If the courses are designed in the manner described in the previous two sections, then the students would be able to select both the courses or parts of courses to achieve a knowledge flow in order to gain the information and skills they actually need.

Seek Perfection

Feedback mechanism would be built into the sessions so that they could be improved both immediately and retrospectively. This would mean that on the next round of design such improvements could be added.

Branching out

The above suggestion would lead to a treelike structure where students could move down various branches depending on who they were, what they needed to know and the order they needed to know it, dependent on the urgency.

Safety often depends on understanding and managing risk. Increasing the efficiency of learning utilising Lean ideas can help to mitigate such risks.

References

- [1] Kent Beck et al, "Manifesto for Agile Software Development", agilemanifesto.org
- [2] James Womack, Daniel Jones, "Lean thinking" Simon & Shuster, 2003, Version, ISBN 0743231643
- [3] Paul Lister, "Agile keys for locked down cages", 2022, ISBN 979-8446432813

Dr Paul Lister, Agile Coach with the Scrum Alliance


Certified Scrum Trainer and former physicist with a PhD in Theoretical Physics from Cambridge, Paul Lister brings over 20 years of fintech experience to helping teams minimise risk and embrace Agility. He is passionate about making work better, sharing knowledge through training, coaching, writing, and community building. He is the author of the book 'Agile Keys for locked down cages' [3] and can be found at agilekeys.co.uk.

Article and Graphics © Paul Lister

"A value stream is a timeline where you understand which pieces of a workflow directly contribute to the outcome you want, and which do not."

¹ CPR (Cardiopulmonary resuscitation) an emergency procedure used during cardiac or respiratory arrest that involves chest compressions, often combined with artificial ventilation.

Mixin Things Up!



In the February 2025 edition of the SCSC Newsletter, the Ontology Working Group (OWG) provided an introduction to ontology, explaining how relevant and important the topic was and the OWG's objectives in developing new ontological models.

The OWG aims to make the topic even more accessible through a series of connected articles to be published in serialised form in the coming editions of the newsletter. Each episode will help expand on ontological aspects and by the end, will hopefully provide the reader with enough knowledge to read and contribute to models in the future.

What Is an Ontology? – a Recap

As we mentioned in the Feb/25 SCSC Newsletter article (scsc.uk/scsc-205), an ontology is not some esoteric formalisation, which is nice to have, but is essential to people's day-to-day work on real systems.

When we build a system, we are modelling our world and this involves developing concepts, terms and their relationships – this is called an ontology, or more formally: “**an explicit specification of a conceptualization**” [1].

If not explicitly specified, establishing a shared understanding of interrelated concepts will be much harder and miscommunication may be more common. For the system to operate correctly, any model used to support a part of the lifecycle of a system needs to be correct and consistent. Hopefully the concepts in the model are defined in writing, but at worst, meaning can just be held in the minds of the designers.

Capturing Models

The OWG uses a specialised framework called the Unified Foundation Ontology (UFO) [2]. This is a theoretical framework developed to provide a basis for conceptual modelling and implemented in OntoUML [3] (this is a customisation of the Unified Modelling Language (UML) [4] with associated software support via application plugins).

In the introductory article we modelled two different perspectives of the concept of a Flight of an aircraft; one from the perspective of a small private airfield, and one that considers flights by a commercial transport airline. The two models were quite different but, importantly, neither was incorrect with both reflecting the particular concerns and perspectives relevant to the associated stakeholders.

To recap on the concepts introduced in the previous article:

- We use boxes (more formally **Classes** using the modelling terminology, that are highlighted in this section in **bold**) to represent the concepts that are important to us in our view of the world. For readability, classes mentioned in the text of this article are **bold purple**, for example: **Aircraft**
- Ontological aspects are expressed as **Stereotypes** (the part of the class definition with Guillemets like this: «role»). **Stereotypes** mentioned in the text in this article are **this bold blue colour** for example, **Role**. (Note that in the diagrams, classes have different colours depending on their **Stereotype**, this is determined by the tooling and not to be confused with text colouring)
- Lines between the boxes show their relationships with other things
- Arrows with a ▽ closed white triangle are **Generalizations**, or in other words, one 'thing' subtypes the more general 'thing' that is being pointed to (for example, a **Cessna** is a subtype of **Aircraft** and **Aircraft** is a **Generalization** of **Cessna**)
- An open arrow ➤ relationship shows an **Association** between things. For example, in the phrase "Pilot 'flies' Aircraft", the concept 'flies' would be modelled as an **Association**. Note that **Associations** also have **Stereotypes**.
- **Association ends** are the numerical expressions at the end of each association, such as **0..*** and **1..*** (read as *zero or more* and *one or more* respectively). These show **Multiplicity** – the range of the number of each type of thing that can be related. This is used to help express constraints, such as a **Pilot** can be trained to fly one or more **Aircraft**, but a given **Aircraft** can usually be flown by at most two **Pilots** at any one time.

We also introduced some core UFO concepts, expressed as stereotypes:

- **Kinds** reflect the objects we find around us like **Humans**, **Aircraft** and **Charts**
- **Roles** and **Phases** reflect things that **Kinds** might be during their lifetime. For example, a **Human** may be a **Student Pilot** for some part of their life (a **Role**), but will cease to be a **Student Pilot** when they leave training school. Similarly, an object's failed state is a **Phase** it might go through, which is transitory and only for part of their life (usually the latter part) such as a **Broken Propeller**.

The concept of a **Mixin** was also briefly introduced and that is where our story will soon continue. A **Mixin** is something we use to describe shared properties of objects that do not all belong to the same **Kind** and so can span multiple **Kinds** that do not share the same *identity principles*. Before expanding on **Mixin**, let's explore *identity principles* in more detail.

Identity Principles

If I replace a tyre on my car, I doubt few would dispute that the car is still identifiable as the same car. This might also be safely assumed if I replace the clutch, the drive shaft or even the engine. If, however, I scrap the car so it is just a pile of metal, screws and plastic in a box, we'd likely say it is no longer a car; so where is the line between these extremes? What if *all* the parts in the car were replaced over time? Is it still the same car?

In 1990 there was a dispute around the authenticity of a vintage racing car called "Old Number One". The original, a Bentley Speed Six vehicle, won the Le Mans 24 hours endurance race in 1929 and 1930 and the purchaser, having agreed to buy the car for £10m, researched the vehicle and realised it had been significantly modified since then and disputed its authenticity. In fact, the chassis, front axle, gearbox, differential and, quite possibly, the steering column and engine crankcase had already been replaced between 1929 and 1930 and over the years other modifications and replacements parts where necessary. The case went to court and the outcome was that the car was deemed the "direct descendant" of Old Number One and that *no other car can claim its identity*.



I am Spartacus!

The reason the box of scrap is not my car anymore but Old Number One remains the same vehicle is to do with our common-sense identity principles. These are functions used to distinguish individuals. Unfortunately, there is seldom a single unique function that faithfully distinguishes things and so we need a series of conditions that, in combination, give us the ability to separate one individual from another (*individuate* more formally).



The famous quote "I am Spartacus!" arose from the eponymous 1964 Kirk Douglas film where a number of people stood up claiming they were the individual named Spartacus. This caused a headache for the Romans, but this is fairly straightforward with the right information. Having the birth date, place of birth and birth name is usually sufficient to identify any given person². Only in combination do these conditions

reliably identify and differentiate individuals.

² Unless of course parents give newly born twins the same name, which apparently has happened. Twins are problematic as they pretty much share the same DNA too. It's left to the reader to puzzle over what might be a completely foolproof method of individuating human beings!


Identity principles are important to us as they help guide us in building type hierarchies by imposing constraints on meaning through abstraction and other techniques in conceptual models. This provides constraints and structure that are a “good thing” as they help us build meaningful and self-consistent models that we can use to reason about our world.

Some stereotypes provide an identity principle and some do not. For example, **Kind**s provide identity principles but a **Mixin** does not. Other classes such as **Role** and **Phase** don’t provide an identity principle themselves but must inherit from a class that does, such as **Kind**. The identity principle allows us to clearly individuate different instances of a class and effectively allows us to do things like count the number of individuals. For example, we can count the number of **Aircraft** (a **Kind**) on the planet, count the number of **Pilots** (a **Role**) employed by an airline and perhaps worry about the number of **Broken Propellers** (a **Phase**) we have on our **Aircraft**.

Mixins

After covering the principles of Identity, we can now return to **Mixins**. These were introduced briefly in the previous introductory article where the **Mixin: Insurable Item** was shown to span different **Kinds**.

We can now show this explicitly on a revised ontology diagram. We can now see that the **Pilot**, the **Passenger** and the **Air Traffic Controller** in their respective roles are insurable. For example, the **Pilot** and **Passenger** might be insured against death and injury and the **Air Traffic Controller** might have professional indemnity insurance. The **Flight** itself might also be insured – compensation might be claimable if the flight doesn’t take place.



There are a lot of lines at the top of the diagram but they are simply saying that:

Pilot, **Passenger** and **Air Traffic Controller** are roles of **Human**
Pilot, **Passenger**, **Air Traffic Controller** and **Flight** are specialisations of the generalization: **Insurable Item**

Importantly though for **Mixins**, things with the **Role** and **Relator** stereotypes are quite different but are still types of **Insurable Item**.

Note that although **Role** and **Relator** follow different identity principles, the **Mixin** cannot therefore provide an overarching identity principle.

By their nature, **Mixins** are a little too flexible so in some cases it helps to be a bit more specific about what is intended. There are therefore stereotypes called **PhaseMixin** and **RoleMixin**. As the names suggest, these apply only to subtypes of **Phases** and **Roles** respectively. For example, An **Individual Customer** and a **Corporate Customer** (both **Roles**) are types of **Customer** (a **RoleMixin**)

Mixins are used in ontologies to reuse patterns of properties and relationships across different concepts, promoting reuse without forcing strict hierarchical inheritance. By defining shared characteristics in a **Mixin**, a concept can "mix in" these properties, leading to more modular, maintainable, and expressive ontologies that can more accurately represent complex domain knowledge.

What We've Learned

In this episode we refreshed our memories on the modelling concepts introduced in the previous article covering:

- **Classes, Stereotypes, Generalizations and Relationships**
- The UFO concepts of **Kind**, **Role** and **Phase**

We then discussed the *Identity Principle* and covered the UFO **Mixin** concept in more detail through the example of an **Insurable Item**

In the next episode we will explore relationship stereotypes in more detail and cover the few remaining class stereotypes we've not mentioned so far. We'll do this by expanding our example to include aspects of flight planning.



References

- [1] T. R. Gruber, A Translation Approach to Portable Ontology Specifications, in Knowledge Acquisition, 5(2):199-220, Academic Press, 1993, <http://www-ksl.stanford.edu/kst/what-is-an-ontology.html>
- [2] G. Guizzardi, A. Botti Benevides, C. M. Fonseca, D. Porello, J. P. A. Almeida, T. P. Sales, Ufo: Unified foundational ontology, Applied ontology 17 (2022) 1–44, <https://doi.org/10.3233/AO-210256>
- [3] Guizzardi, Giancarlo & Fonseca, Claudenir & Almeida, João & Prince Sales, Tiago & Benevides, Alexander & Porello, Daniele. (2021). Types and taxonomic structures in conceptual modeling: A novel ontological theory and engineering support. Data & Knowledge Engineering. 134. 101891. 10.1016/j.datak.2021.101891. https://www.researchgate.net/publication/351557008_Types_and_taxonomic_structures_in_conceptual_modeling_A_novel_ontological_theory_and_engineering_support
- [4] Unified Modelling Language, <https://www.omg.org/spec/UML>
- [5] American Airlines Flight 965, https://en.wikipedia.org/wiki/American_Airlines_Flight_965 , accessed August 2025.
- [6] European Geostationary Navigation Overlay Service, https://en.wikipedia.org/wiki/European_Geostationary_Navigation_Overlay_Service, accessed Sept 2025
- [7] Halliwell, J., Chambers, S., Cropper, P. and Foulsham, M. (2024) Independent Review of NATS (En Route) Plc's Flight Planning System Failure on 28 August 2023 – Final Report. Independent Review CAP2993 (Ver. 1). UK Civil Aviation Authority. <https://www.caa.co.uk/our-work/publications/documents/content/cap2993>

Image attributions:

Old Number One: | El Caganer / Craig Howell, CC BY 2.0 <https://creativecommons.org/licenses/by/2.0>, via Wikimedia Commons

Thracian soldiers: AI generated

OntoUML diagram: OWG provided as public domain

American Airlines Flight 965 crash scene: FAA, Public domain, via Wikimedia Commons

The Ontology Working Group (OWG)

The OWG was set up in 2020 to help address issues with the lack of formalisation in terms and conceptual relationships used in standards and guidance documents. In particular, the similar but at times different, interpretations of terms and concepts used between safety and security domains is an ongoing cause of confusion and misunderstandings and presents barriers to more formal integration of these disciplines.

The OWG has its origins in the Data Safety Initiative Working Group (DSIWG), which, between 2016 and 2019, explored the development of an ontology model for the concepts described in the Data Safety Guidance (DSG). However, it was realised that the work couldn't be completed without a wider ontological model for risk, and so the OWG was established and is currently working on developing a clear, systematic terminology, and unambiguous model of risk and value concepts and their relationships.

SCSC Seminar: How Safety Culture has to Change with AI



The “How Safety Culture Has to Change with AI” seminar was held at the Hilton Euston hotel in London, on 19th June 2025 and hosted by Mike Parsons. The speakers presented their work and insights into the ways in which AI impacts existing safety culture, as well as methods to integrate AI and human expertise, address concerns around accountability and recognise the importance of human factors.

The six speakers presented interesting and lively sessions on development and frameworks for safe AI systems, covering a range of topics such as:

- Michael Wright (Wright Human Factors Ltd): **AI: Emergent Safety Culture Issues** – how AI impacts human factors, how accidents can be attributable to both human and AI failures
- Ben Fulford (BMT): **Merging Human Expertise with AI for Better Safety Management** – how can we make the output of generative AI more accurate and effective, integrating risks and mitigation strategies into our processes
- Paul Leach (RSSB) – **Human Factors Principles for Artificial Intelligence Systems in Rail** – a psychologist’s perspective on rail safety, how human factors can inform AI integration
- Adam Johns (Marsh McLennan) – **Redefining Accountability: Just Culture in the Age of AI** – maintaining psychological safety of workers, the cultural and organisational challenges of integrating AI into socio-technical systems
- Kathy Syfret (RAF) – **AI and Aircraft Maintenance** – AI’s role in aircraft maintenance, and the importance of understanding organisational and human factors
- Paul Traub (Paul Traub Associates) – **Automation Policy and Human Performance in Autonomous Weapon Application** – levels of autonomy, risks associated with AI and human control

Michael Wright – AI: Emergent Safety Culture Issues

Michael warned us all that there were currently more questions around AI integration and culture than there were answers, which proved to be a recurring theme throughout the day! Taking us through some preliminary definitions including automation vs AI, Michael provided a number of potential AI applications including both collaborative and cooperative systems. He noted that while humans perform badly at “monitoring roles”, this could be an area where AI can excel, and illustrated this with some examples of accidents resulting from human disengagement with monitoring.

Michael then moved on to discuss how accountability within safety culture can be affected by AI, providing us with some more enlightening illustrations of how over-reliance on AI can lead to serious consequences. Not only do we experience loss of ownership of decision making, but we risk de-skilling, loss of trust and the potential for some truly unsettling incorrect AI learning and adaptation. Michael concluded by showing us an AI generated essay on AI and safety culture from Microsoft Copilot, which was – he and the audience estimated – probably about 70% right. Certainly “pretty good”, but as Michael warned us, the difficulty is in validating which parts are in that 70%!



Ben Fulford – Merging Human Expertise with AI for Better Safety Management



Next up was Ben Fulford, discussing the risks on the opposite side of the question: the risks of not using AI. Ben showed us that AI adds most value where both the user’s job experience and the job complexity are either both high – for example in safety-critical systems assessment – or both low. He illustrated this with a compelling example describing how better outcomes from a System-Theoretic Process Analysis (STPA) were achieved by using a combination of experts and generative AI, as opposed to either working alone. However, he did caveat this with the observation that ChatGPT is not trying to tell its users the truth, it is trying to sound plausible. Definitely something to keep in mind there!

Ben also presented BMT’s work on a prompt creation framework to get the best out of generative AI, requiring that each query is provided with context, the relevant role the quest is applicable for, the question itself, the required format, the required tone and a template to show the generative AI what a good answer would look like. Ben presented the BMT Co-pilot tool which uses both the prompt creation framework and a knowledge graph to generate accurate and effective answers. He finished with a timely quote which sums up much of the work presented today: “it’s not about AI replacing expertise, it’s about merging expertise with AI”.

Paul Leach – Human Factors Principles for Artificial Intelligence Systems in Rail

Paul Leach then took over to provide us with a psychologist's perspective on rail safety, serving as a very good illustration for the multi-disciplinary nature of safety engineering! Paul emphasised the need to identify human-centred design principles relevant to AI and consider the socio-technical aspects of human-AI interaction, noting that we should not jump headfirst into using AI technology just because it is "new and shiny": a relevant reminder for us all.



Paul showed us some recent research on identifying a set of human factors principles that could be used to inform the design, implementation and operation of AI systems based on literature review, expert interview and repeated internal review and refinement. These principles included recognising that AI should support rather than sideline human strengths, working towards creating effective human-AI teams where the AI is a team-mate rather than a tool, and ground-up recognition of ethical challenges and potential mitigations. Along with this, Paul emphasised the importance of recognising "work as done" vs "work as imagined, or prescribed". In other words, when we discuss AI integration into safety culture, we need to consider failures of the system rather than only concentrating on those situations where the system performs ideally!

Adam Johns -Redefining Accountability: Just Culture in the Age of AI



After lunch Adam Johns picked back up on the theme of "questions, not answers" by asking us all to define what we thought the primary purpose of a just culture was. Adam followed up on this by reminding us all of the importance of definitions, including – especially salutary given the day's topic – the necessity of recognising that safety culture is about how to make trade-offs between safety and other competing goals. He emphasised that culture is about humans and that accountability should not be regarded as a

synonym for blame. This is particularly important when considering how to foster an effective safety culture even without the complexity of AI integration.

As with several of the other presenters, Adam had coaxed AI into generating a sample image: in this case a beautiful image of a tree as an ecological metaphor for aviation safety (although the less said about the words that the AI had provided to go along with the image, the better: a humorous and very realistic example of the importance of monitoring AI output!). The importance of language is a crucial point here: when we think of AI using environmental metaphors this opens our minds to new ways of describing both its risks and its benefits. Adam finished with the important reminder that we must fully understand a system before intervening in it, particularly where this intervention can make things both worse and better at the same time, as AI is particularly wont to do!

Kathy Syfret (RAF) – AI and Aircraft Maintenance



Kathy Syfret was next, bringing us right back to practical elements with a very enlightening look at how AI might be used for aircraft maintenance. Kathy emphasised the importance of understanding both organisational and human factors, particularly how culture is made up of a kaleidoscope of ideas, customs and social behaviour. Referencing the Haddon-Cave report³, which will no doubt be familiar to everyone reading, Kathy showed us how culture can be changed in order to improve safety, identifying that such a shift will be important to ensure successful integration of AI.

Kathy then moved on to provide us with some examples of how AI might be used in practice for aircraft maintenance, looking at predictive maintenance and reliability-centred maintenance as particular areas where AI could be integrated. Supply chain management was another, which she noted tended to be overlooked in favour of communications: while it's certainly true that "if you have no comms then you have no bombs", the same could be said if you indeed *have* no bombs due to a supply chain failure! Kathy also referenced the importance of adaptive safety culture when integrating AI, and the need to address both complexity of organisations and complacency in our existing processes.

Paul Traub – Automation Policy and Human Performance in Autonomous Weapons Application

Paul Traub demonstrated technology in action by dialling in to give the final talk of the session, returning to the question which had cropped up again and again during the day: is AI the panacea for human performance limitations? The answer, he warned us, is "not necessarily": there are risks from automation too. For example, AI usage has been linked to loss of situational awareness, where the operators of a system do not understand the extent of automation in that system, nor how best to manipulate this. As automation occurs, Paul demonstrated, the operator becomes more "out of the loop", thus further increasing the chances for error.




Paul next discussed how a robust automation policy can reduce risks of automation, and that this should start at the ground up with assessing how automation is to be integrated into design. He illustrated how the implementation of automation should include a task and error analysis, and emphasised the importance of identifying and understanding a system's critical tasks and potential points of human failure. He finished with another important reminder that problems in AI must not just be simply reported, but that the report must be acted on: this is not always something which is sufficiently highlighted in safety culture, but is going to become increasingly important as more systems integrate AI.

³ "The Nimrod Review: an independent review into the broader issues surrounding the loss of the RAF Nimrod MR2 aircraft XV230 in Afghanistan in 2006 report", See <https://www.gov.uk/government/publications/the-nimrod-review>

Mike Parsons – Wrap Up

Mike Parsons brought the threads of the day together in a final general discussion, noting some important areas we hadn't yet had time to discuss. These were wide-ranging, including regulation, drones, autonomous vehicles, malicious action and what a near-miss might look like in AI. Other participants also raised concerns such as the maturity of AI: is it really mature enough to be used in safety-critical systems? Audience members also contributed, highlighting the "human" aspects of emotional resistance to AI and the importance of avoiding anthropomorphism when we use AI, a timely reminder to everyone who has ever become irritated by ChatGPT's hallucinations!



The audience were in general agreement that AI integration into safety-critical systems is as much a human question as it is a technological question. Safety culture is, inevitably, a by-product of human culture, and safe AI integration will require an adaptation of the ways we think, interact and relate to systems. All in all, this was a very interesting and lively day, with some fantastic insights into AI frameworks and the challenges that must be overcome.

Report by Catherine Menon.

Catherine Menon is a principal lecturer at the University of Hertfordshire. Her research focuses on the intersection of trust, safety and ethics in public-facing autonomous systems. She has a PhD in category theory and has been involved in the development of regulatory standards and guidance for safety, security and ethics assurance within autonomous systems.

Upcoming Seminar: Soft Skills for Safety Practitioners

Seminar: Soft Skills for Safety Practitioners



This workshop will be held on Thursday 23rd April 2026 at the BCS, 25 Copthall Ave, London EC2R 7BP.

The day will cover the soft skills that a safety practitioner needs to support safety engineering activities, and will have a variety of speakers followed by a workshop.

This event will be useful for all those who have to deal with safety on a day-to-day basis including safety engineers, consultants, auditors and managers.

Topics include:

- How to ask the right questions
- How to be assertive when required
- Use of prompting techniques
- Management of meetings where safety is discussed
- Management of difficult behaviours regarding safety issues
- Knowing when to stop
- How to support an active safety culture
- How to convince budget-holders for the need to attend events


This event will be held at the BCS, 25 Copthall Ave, London EC2R 7BP

Event Information

Event Date	23/04/2026 9:30 am
Event End Date	23/04/2026 5:00 pm
Cut Off Date	23/04/2026
Individual Price	£265.00
Location	BCS London

scsc.uk/e5055

Seminar: Safe Agile Developments



The “Safe Agile Developments” seminar was held at the Wellcome Collection in London, on 1st May 2025 and hosted by Mike Parsons. This seminar was for those introducing or using Agile techniques with safety-related developments. It was aimed at developers, change managers, safety engineers and project managers.

Mike Parsons opened the seminar saying that there is a perception that Agile development and safety do not mix well and he hoped the seminar would help determine if this indeed was the case or whether techniques and approaches have now been developed to allow Agile principles to be used successfully on safety-related endeavours.

There were seven speakers in total:

- **Dewi Daniels**, Software Safety Ltd - 24 years on from the Agile Manifesto: What have we learned?
- **Bill Blackburn**, Consultant - Scaled Agile for High Integrity Services
- **Oumaima Fasih**, Functional Safety Engineer & Scrum Master - Bringing Agility to Functional Safety: A Practical Approach Aligned with ISO 26262
- **Fred Heath**, Faria Education Group - Can agility and safety coexist?
- **Paul Lister**, Equiniti - How Agile helps you hit the complexity bullseye
- **Núria Mata**, Fraunhofer Institute for Cognitive Systems IKS - Building blocks for the development of safety-critical automated driving functions in a holistic DevOps approach
- **Thor Myklebust**, Sintef - The AI Act and The Agile Safety Plan

Dewi Daniels

Dewi Daniels kicked off the seminar covering the history of software development. Dewi first mentioned Dave Thomas who he had worked with in the past, and was a cofounder of the *Agile Manifesto*. Dewi's early engagement with Agile was actually in "lean" development and its principles that came from the Japanese car industry as described in books like "The Machine That Changed The World".



Dewi introduced the Correctness by Construction idea he developed when at Praxis and ran through the principles. Dewi then went on to discuss his time at Silver Software and then Verocel and Callen-Lenz Associates, his agile work and applying these to safety critical projects. Dewi said that in 2015 Dave Thomas stated that "Agile is Dead" as people had developed methodologies for agile which was what the manifesto was trying to avoid (treating agile as a noun).

Dewi then went on to discuss:

- "The Good" – testing is now fashionable, it is more responsive to change, delivery is more predictable and there are excellent tools for version control
- "The Bad & Ugly" – misunderstandings rather than failures with Agile in itself and short termism – preoccupation with current "Sprints"

Dewi then went on to discuss applying Agile principles to safety-critical software development and said it is fundamentally different from commercial software, particularly in testing, where 40-60% of a safety project will be spent on verification. Formal methods and supporting tools are used more often in safety projects.

Dewi said that after 24 years of the manifesto, the age of early adoption is over and he concluded that, in his experience, Agile principles can be applied to safety critical systems.

Bill Blackburn

Bill was our second speaker of the morning covering high-integrity services and in particular how Agile can scale, not specifically on safety programmes but where high integrity is paramount. Bill provided a recap of Agile and earlier methods such as DSDM and its principles which are similar to Agile principles.

Bill then introduced a diagram which is a framework for scaled agile, which showed the collection of components such as roles, and the Agile Release Train and the continuous delivery pipeline.

Bill then covered three case studies in Life Sciences, Financial Markets and Public Sector.



The Life Sciences case study was for a pharmaceutical company operating in a highly regulated industry that recognises the importance of safety, which is called Pharmacovigilance. For this project, the focus was on product assurance and compliance. Bill worked on evolving the business to be more Agile such as improving management of supply chains and ensuring there is an Agile platform to work upon.

The second project was for an Electronic Swiss Exchange in Switzerland. For this project, Bill took a safety-related approach to business criticality where Agile processes were applied to trading systems.

The third case study was a policing system being replaced with a new customer management system, where he took an architecture-led approach. One misunderstanding he encountered was treating Agile as a silver bullet. Bill mentioned the government's Technology Code of Practice and he said although it doesn't mention Agile, it is aligned with those principles.

Bill concluded with the key mobilisations aspects that realise the overall framework and introduced SAFe (a knowledge base of proven, integrated principles, practices, and competencies for achieving business agility using Lean, Agile, and DevOps).

- Risk-based approaches help deliver benefits of agile with assurance
- There are good industry practices
- Collaborative culture and right skills are imperative
- Regulators need to be on board and need a platform

Paul Lister

After a coffee break, Paul Lister from Equiniti discussed how agile can help manage complexity. He started by defining complexity with models of interactive "agents" in complex networks. Paul introduced the idea of "Safety First" where risks of unwanted agents are mitigated and detected and removed. But no system is truly closed, and work "as imagined" is not often "work as done" and doesn't cover emergent behaviours.

The new processes such as "Safety II" accept this and protocols are emergent based on work done and accidents don't have a single cause.

Paul introduced a four-level scheme for the levels of understanding and control for an endeavour and mapped this to the 4 important aspects of Why?, What?, When? and How? as seen on a 'bullseye' diagram and he gave real-life examples of different scenarios.

Paul showed how the principles of the Agile Manifesto can help in each of the four area and summarised the principled as Outcome of Output and Trust over Control. He then covered "Lean" principles and in particular cutting out waste and focussing on generating value.

Paul then introduced "Scrum" with the focus on building small things and having small teams with the appropriate skills and doing the work in small amounts of time.

The next stage is then examining what has been done that requires transparency, inspection and adaptation and a number of "Scrum events" (planning, daily scrum, sprint review and retrospective) effectively using complexity to fight complexity. Paul concluded with key take-aways:

- no matter the task it happens in the open
- understand complexity with visualisation
- utilities Agility to mitigate risk and inspect and adapt



Oumaima Fasih

After lunch we heard from Oumaima Fash from Stellantis presented remotely and discussed a practical approach to bringing agility to functional safety aligned with ISO 26262.

Oumaima introduced her talk by asking 'Why Agile?' and said that agility is now highly prevalent in everything we now do from online ordering to over the air updates in cars and so there is a desire and benefits for agile approaches. But this level of agility brings challenges for safety and she showed that there has been a marked increase in vehicle recalls.

Oumaima described the history of ISO 26262 – a safety standard for the automotive industry and showed a process overview and 'V' model for ISO 26262. She then focussed on the V model in more detail and concluded that it is ostensibly a linear waterfall methodology. She then reprised the Agile Manifesto and highlighted the similarities and conflicts with ISO 26262, such as agile having minimal upfront planning where ISO 26262 requires detailed planning before development begins.

Oumaima then showed how safety roles and terminology apply to the Scrum methodology. For example, there is a Safety backlog (outstanding work to be done) along with Safety Validation acting as a kind of Sprint Retrospective.


Oumaima concluded with six tips:

- Start small, scale gradually
- Engage cross-functional teams early
- Establish safety backlog early
- Regular reviews and feedback loop
- Address compliance in the Agile Safety Model
- Use digital tools to help generate the documentation required by ISO 26262.

Fred Heath

Fred Heath from Faria Education Group was next up and asked the broad question of whether Agile and safety can co-exist. Fred said he was not a safety engineer as such but works in a safety-related educational environment where children's data needs to be protected.

Fred started by explaining the issues that led to the crash of the Mars Climate Orbiter where there were discrepancies on the units of measurements being used between two teams. Fred stressed that nothing in the manifesto conflicts with safety objectives and he thought Scrum was the best framework for safety-critical systems. Fred covered the benefits of Agile and mentioned SafeScrum as a successful adaption of RTCA/DO-178C. There are concerns however for using Agile in safety-related projects and Fred dispelled some of the myths and misconceptions.



Fred provided more details of the Scrum framework and the concept of splitting activities into a “sprint”, the “backlog” of agreed work tasks and the general process of Plan, Implement and Inspect & Adapt and the important aspect of the “Definition of Done”.

For safety-critical projects Fred said that sprint planning requires more refinement and will take longer. He then discussed and gave examples of Behaviour-Driven Development (BDD) which is a superset of test-driven development applied at the system level. As the definition is structured, a test stub can be automatically generated. BDD provides automated validation, documentation and traceability.

Fred then returned to the earlier concerns on using Agile in safety-critical projects and for each showed how each can be addressed. Fred concluded with references to his books on safety.

Thor Myklebust

Thor Myklebust from Sintef then discussed the AI Act and the Agile Safety Plan, which is encapsulated in a published book based on interviews with 35 risk experts and 18 companies.

He described the wide range of safety projects he is involved in such as Hyperloop, transportation and offshore. Thor said that there have been a lot of changes since he started work on IEC 61508 with more companies becoming involved over time.



He described the wide range of safety projects he is involved in such as Hyperloop, transportation and offshore. Thor said that there have been a lot of changes since he started work on IEC 61508 with more companies becoming involved over time.

Thor discussed the EU Artificial Intelligence Act and presented the 4-tiered risk levels in the Act. He said there are different interpretations of the Act and showed the process for developing harmonised standards, which might take many years.

AI safety plans and cases involving AI are still in their infancy and the act is likely to change again but as with previous speakers, he stressed the importance and benefits of Agile.

Thor referred to “SafeScrum” an adaptation of Scrum for safety-critical projects but said there were challenges with collaboration between the different disciplines and in particular, there are safety planning challenges. Thor also covered challenges with software, data, documentation, compliance, procurement (reluctance to buy certified products) and AI with, for example, being bad at outliers, and explainability for different stakeholders. Thor said AI systems are more like a grey-box and you should use the information you have although some treat it as a black box.

Thor covered the pitfalls with producing a safety argument and said good design might reduce the burden on argumentation. SafeScrum can be adapted for AI lifecycles with new AI roles. Thor said that prototypes can be developed quickly in Python to show the client and then developed in RUST. There has been more use of tools but this can be a disadvantage as it may mean they need to be qualified.

Mike Parsons – Discussion

Mike Parsons then hosted a general discussion starting with asking whether existing standards can be adapted for Agile?

Stephen felt this would be a challenge; for certain things we need the right level of documentation. There may be cases where you need to understand the system and its functionality before you can do anything Agile.

Mike then asked if there is a danger that engineers focus on particular sprints and lose the overall picture?

Fred thought there is a misconception that the focus is just on iterations, and there is still planning for the bigger picture.

Andy said there are challenges with attitudes between Safety and Development teams; with release pressures, safety can be seen as just compliance and a blocker. Bill thought the focus is on the product owner and perhaps training and guidance might be good.



Núria said safety cultures are critical and essential for AI and Agile/Devops helps promote the awareness through collaboration.

Tim said that he's seen high degrees of reliability on equipment and it is more often the people causing the problem and perhaps AI could be used to improve human errors.

Bill said big companies are doing more on responsible AI and firms like Amazon are doing more in this area on documentation.


Núria said it is about trustworthiness of AI. Although not 100% reliable, the risk may be acceptable.

Stephen also thought it might be more about trust and how it has been developed and confidence rather than having lots of documentation at the end.


Paul asked if we really need a paradigm change from a regulatory perspective with more active engagement from regulators on the ground during the developments – a form of “continuous regulation”. Stephen agreed and said there was a precedent for this in the railway sector he was working in.

Report by Paul Hampton, SCSC Newsletter Editor

Recent Safety Publications




Safety-Critical Systems
eJournal vol.4 no.1
Spring Issue 2025
scsc.uk/scsc-204




The Safety-Critical Systems
Club Newsletter, Volume 33,
Number 2
scsc.uk/scsc-205




Through-Life Co-Assurance
of System Safety
and Cyber Security
scsc.uk/scsc-173



Proceedings of the Thirty third
Safety-Critical Systems Sym-
posium
scsc.uk/scsc-199




Guidance on the Through-
Life Assurance of Complex
Systems
scsc.uk/scsc-200



Guidance on the management of
data safety risks (V3.7)
scsc.uk/scsc-1271



The Safety-Critical Systems
Club Newsletter, Volume 32,
Numbers 1-3
scsc.uk/scsc-197



A Safety Case Report Format
Andrew Eaton & Stephen
Barker
[www.ama-
zon.co.uk/dp/B0DPL3C84P](https://www.ama-zon.co.uk/dp/B0DPL3C84P)



Safety-Critical Systems eJournal
vol.3 no.2
Summer Issue 2024
scsc.uk/scsc-196

60 Seconds with ...

Nancy Leveson



Nancy Leveson is Professor of Aeronautics and Astronautics at MIT. She is an elected member of the National Academy of Engineering (NAE). Prof. Leveson conducts research on the topics of system safety, software safety, software and system engineering, and human-computer interaction.

She has received many awards for her work including the IEEE Medal for Environmental and Safety Technologies, the ACM Allen Newell Award for outstanding computer science research and an AIAA Information Systems Award for "developing the field of software safety and for promoting responsible software and system engineering practices where life and property are at stake."

She has published over 400 research papers and is author of three books, "Safeware: System Safety and Computers" [1] published in 1995 by Addison-Wesley, "Engineering a Safer World" [2] published in 2012 by MIT Press, and "An Introduction to System Safety Engineering"[3] published in 2023 by MIT Press. She consults extensively in many industries on the ways to prevent accidents.

What first attracted you to working in System Safety?

To answer this question, I need to start before that time. I first earned a B.A. in Mathematics and then an M.S. in management (operations research). Then I joined the working world as a system engineer for IBM. I didn't like working in industry, so I quit and hitchhiked one and a half times around the world for a couple of years until I could decide what I wanted to do. After returning, I knew that I could not avoid the normal working world without a doctorate so I started a Ph.D. program in Cognitive Psychology. In the end, I got a Ph.D. in Computer Science, specializing in formal methods for software, but I remained interested in human cognition and human factors. At that point I had never heard of safety engineering.

Immediately after finishing my degree, my dissertation was plagiarized. I was pretty bummed out. I started my first academic position, but I had no idea what I wanted to do beyond knowing that I did not want to continue to work on the same topic (database integrity). Within a week of starting my academic position, I got a phone call from someone at Hughes Aircraft Company about a torpedo they were building for the U.S. Navy.

The torpedo had 15 microprocessors on it, which was very unusual at that time. Previously, computers had primarily been used for computation only, while control was done by humans or by simple electronic PID (Proportional-Integral-Derivative) controllers. The engineers were very nervous about the potential for this "heat-seeking" torpedo to turn around 180 degrees and hit the ship that launched it.

They told me that this was a “software safety problem.” I told them I had never heard of software safety, but it sounded like a software reliability problem. They were adamant about it not being reliability, but rather involved safety.

This was the first time I had heard about safety engineering. As you can imagine, I wasn’t able to help them much with the torpedo, but I learned an enormous amount about safety engineering from these exceptionally qualified engineers.

When the torpedo was finished, they took it to a torpedo testing area and test fired it. Every time it came out of the torpedo tube, it turned itself off, went to the bottom of the ocean, and just laid there. I said, “well, it’s safe.” They replied that the Navy did not want to pay for this safe torpedo. This was the first time I started to truly understand that safety and reliability were different system properties and, in fact, they could be conflicting—increasing one did not necessarily increase the other and in fact could decrease it. Forty-five years later I am still trying to explain this distinction to computer scientists and to many engineers.

To make the torpedo useful, the engineers took off safety devices, one by one, until it was acceptably reliable and safe. This approach didn’t seem great to me, especially as the amount of software and complexity was surely going to increase in the future. I decided I would take six months to figure out the solution and then decide what area of research I really wanted to work in. It took a little longer than I expected (20 years) to find a successful solution and to convince others that it worked.

Software engineers have always, especially at that time, focused on assurance after the design is complete, so I started with trying to enhance traditional test and assurance activities and to extend the standard hazard analysis techniques that existed at that time (FTA, ETA, HAZOP, FMECA, Bow Tie, etc.).

It quickly became clear that after-the-fact assurance was not practical nor cost-effective as software increased in size. Today, embedded control software can have tens of millions or even hundreds of millions of lines of code. Finding that there is a problem after the system and software has been created is too late—the options available after the design is complete are limited, very expensive, and not very effective.

Instead, safety needs to be designed into systems and software from the very beginning. Focusing on assuring it at the end is a dead-end approach and leads to an unnecessarily and often unacceptable level of risk. In addition, the pressure is great at the end to find any argument (whether bogus or not) to demonstrate that the system is safe “enough.” The focus is on making a convincing argument that what you have is acceptable rather than on increasing safety in the design and the system.

Besides starting too late in the life cycle, the other problem I encountered in my early work is that traditional hazard analysis techniques were developed to analyze the relatively simple systems existing at the time, which did not include software or cognitively complex human activities. The traditional techniques are based on the assumption that accidents are primarily caused by hardware component failure or by humans not following prescribed procedures. They also assume that “software failure” is just like hardware failure.

But the world of engineering has changed and new causes of losses are becoming increasingly important, many due to flawed software requirements and system design as well as human inability to cope with increasing complexity.

The old hazard analysis techniques, which assume that accidents are caused by component failure alone, are not very useful on these new systems. The focus on component failure leads to techniques that can identify component failures leading to hazards, but nothing else: They cannot be expanded to identify the very different causes of accidents today. That fact doesn't seem to prevent many people from using those techniques, it is just preventing them from providing much useful information to the system designers or assurers. This fact seems to be ignored by those filling out endless forms and tables of information about random component failures and computing their probability while ignoring design flaws.

I believe that the only way to make real progress is to change to solution paradigms that solve the problems of today, not those of yesterday. To accomplish this, I decided that we needed to start from System Theory, and I created a new model of accident causality (STAMP) that focuses on control rather than failures and reliability. On that foundation, very different types of hazard analysis techniques can be created, including STPA (System-Theoretic Process Analysis) and CAST (Causal Analysis based on System Theory). These new analysis techniques seem to be working well and are being used widely in most industries, especially on the most complex and potentially dangerous systems being built today.

What aspect of your career are you most proud of?

I'm clearly proud of my contributions to system safety engineering, but another important aspect has been my work throughout my career mentoring and assisting women to be successful in a profession primarily focused on and catering to men.

What has been one of the key challenges you've experienced in your career?

Being female has far and away been the biggest challenge. The second has been to get people to stop just doing what they have always done—although it no longer works—and to look at new ways of approaching problems. We need more scepticism about what is considered “truth.” I tell my students “You are here not to worship what is known, but to question it” (Jacob Bronowski, *The Ascent of Man*).

If you could change one aspect of system safety, what would it be?

We need to stop relying on probabilities as the goal in safety engineering or even as a way to reach the goal. They may provide a way to help management make decisions, but the probabilities on which those decisions are based are always wrong so the resulting decisions are not necessarily very good. I've read hundreds of accident reports, and almost every one of them contained a Probabilistic Risk Assessment (PRA) that had showed the accident could not happen—usually exactly the way it did happen. The complacency generated by the numbers was often a major contributor to the loss. There have been only two scientific evaluations of PRA that I have ever found. Both found it to be extremely inaccurate.

Beyond not being accurate, the results do not provide any information to engineers today to design systems to be safer. The risk calculation can be improved by adding a lot of redundancy, but that doesn't have any impact on safety when the problem is a design flaw, the most prevalent contribution today of software and humans to accidents. We need to design systems to be safe, not just convince ourselves that they are safe by producing computations that cannot be validated but make us feel better about the potential risk.

What are the most infuriating misconceptions about the discipline that you still encounter?

Most misconceptions arise from the fact that engineering is very different than it used to be and that old solutions no longer work. Here are a few:

Misconception 1: Risk should be defined as likelihood combined with severity. In fact, likelihood is impossible to determine for most hazards and losses and even severity is more difficult to determine than is usually assumed. The problems in defining severity arise from heuristic biases [4]. There are alternative, improved ways of defining risk, but they require a paradigm change.

Misconception 2: Most accidents are the result of component failures.

Misconception 3: Software “fails” and thus reliability concepts can be applied to software.

Misconception 4: Safety can be assured after the design phase on complex systems today.

What advice would you give to younger people seeking to build a career in systems safety?

- Don't work on “toy” problems. Everything works on such examples, particularly if you ignore all the non-straightforward (“ugly”) parts of the problem. Work on real problems at real scale. I make my students apply their new solutions to real systems in their theses and dissertations. The proprietary parts of the design do not need to be made public, but validating new approaches and ideas on real systems is necessary before we encourage people to use those approaches. The important problems in very large systems are different than those in smaller ones.
- Start from the problem, not from solutions. One of my very early papers—that I regret ever writing—applied Petri nets to hazard analysis. Why did I choose Petri nets? I knew them and liked using them. Since that time, I don't know of any real system on which Petri nets have been or could be used for hazard analysis. Amazingly, however, this paper is one of my most highly cited, and I still come across government handbooks, etc., that recommend the use of Petri nets for hazard analysis. They all cite that silly paper I wrote 40 years ago. People are writing lots of conference papers on using Petri nets for safety analysis, citing my paper as a reference. Do not start by applying previous solutions, that you know or like, to a different problem. Instead, start with understanding the problem you are trying to solve. Use that information to find solutions tailored to that problem.
- Learn the history of engineering and of safety engineering—it goes back at least 150 years—and determine what still works and what never worked. Stop reinventing the same things (often with new names) that did not work in the past. We don't make progress that way. At the same time, realize that the world of engineering is changing at lightning pace: We should not robotically keep doing what was done in the past just because it used to work and we know how to do it.

We should not robotically keep doing what was done in the past just because it used to work and we know how to do it.

How do you see the rise of Artificial Intelligence affecting the future of safety engineering?

There are some important engineering problems that will be solved usefully by machine learning and LLMs, but safety is not necessarily one of them

We are living in one of the periodic AI hype bubbles. I've lived through several of them. In the 90's, "expert systems" were the focus of almost all engineering research. They were touted as the solution to all important engineering problems. How many engineering researchers today are focused on using expert systems?

There are some important engineering problems that will be solved usefully by machine learning and Large Language Models (LLMs), but safety is not necessarily one of them. Their use is more likely to increase risk than to decrease it and to create new types of risk. I don't see any way to ensure safety in software when we cannot predict or control the outputs and where the software

periodically "hallucinates" (i.e., lies). It may be possible to design the parts of the system around the AI components in some cases to protect against potentially dangerous AI behavior. But I don't see any obvious general solution.

People are also talking about using AI to perform safety engineering tasks (e.g., doing hazard analysis), which seems unethical to me except for the most trivial tasks.

References

- [1] Nancy Leveson, *Safeware: System Safety and Computers*, Addison Wesley, 1995
- [2] Nancy Leveson, *Engineering a Safer World*, MIT Press, 2012.
- [3] Nancy Leveson, *An Introduction to System Safety Engineering*, MIT Press, 2023.
- [4] Nancy Leveson, *Improving the Standard Risk Matrix, Part 1*, unpublished white paper, downloadable from <http://sunnyday.mit.edu/Risk-Matrix.pdf>

Upcoming Event: High Integrity Software Conference



The mission of the High Integrity Software Conference (HISC) is to bring software engineering practitioners together to share their challenges, best practices and experiences. HISC will return to the International Convention Centre (ICC) Wales on 13th November 2025.

Registration is now open. Get your tickets and learn more:

<https://www.his-conference.co.uk/>

We are in an age of digital dependency, where more and more aspects of our everyday lives will rely on the correct function of software-intensive electronic systems.

These systems may be tangible, at-hand, embedded systems, but more frequently are remote and highly distributed, relying on multiple layers of communication and infrastructure for correct operation. We expect these systems to deliver ever-more complex functionality on demand, and need to be able to trust this functionality to be delivered with appropriate levels of safety and security.

The need for trustworthy software – software that does exactly what it is intended to do – is well-established in certain sectors such as defence, aerospace, and transportation, along with a mature approach to its design and assurance. But the increased use of software in more and more domains – embedded, connected, consumer and infrastructure systems – means that similar approaches are needed across a broader range of industries.

In addition, there are disruptive technological influences which mean that traditional approaches to verification and assurance in the established industries will no longer remain applicable or sufficient.

The aim of the High Integrity Software conference is to share the experience and know-how that will address these challenges within our industry and help to build the foundations of a trustworthy software ecosystem for the future.

Flying Visits!

An impromptu SCSC Tech Trip



Paul Hampton and Dewi Daniels report on their experiences being pulled by the iconic “Flying Scotsman” locomotive, now over 100 years old, but still going strong!

West Somerset Railway (WSR) operate a heritage line between Bishops Lydeard and Minehead. This is a 20-mile stretch in beautiful Somerset countryside and includes a coastal stretch with views across the Bristol Channel to Wales. When I heard WSR were offering passenger trips on the line with the Flying Scotsman, I thought it would be a rare opportunity to see the almost 100 tonne locomotive in action.

Static museum exhibitions of vehicles are of course interesting and worthwhile, but to see, hear (and smell!) such a large vehicle in operation brings another dimension of emotional engagement. I was not alone in this ambition and it also appealed to Dewi Daniels who travelled on a separate day so we have a sort of impromptu SCSC Tech Trip to report on!

The Flying Scotsman was designed by Sir Nigel Gresley and built in Doncaster in 1923, the first locomotive of the newly formed London and North Eastern Railway (LNER). In the next year, after an appearance at the British Empire Exhibition in London, it was given the name “Flying Scotsman” after the daily London to Edinburgh rail service. In 1928 it was given a new type of tender (the carriage hauled immediately behind the loco carrying the coal) with corridors so new crew could reach the engine without stopping the train. May that year then saw the first non-stop London to Edinburgh service with a journey time of eight hours.

In 1934 it reached a record-breaking speed of 100 mph on a special test run and 15 years later, following the nationalisation of Britain's railways, it was given its now famous 60103 designation.

After travelling more than 2 million miles in almost 40 years of service, the loco was at risk of being scrapped but was saved through its purchase by railway enthusiast Alan Pegler. After restoration, Pegler continued to run services and in 1969 the loco embarked on its remarkable international career, running services in the US, Canada and even Australia. Its future has never been certain with its many owners over the years facing financial difficulties and various campaigns and appeals have been required to pay for expensive overhauls. It is a testament to its enduring appeal that despite these setbacks, it has survived and is still running today.

The Journey



Our days both started at Bishops Lydeard. Having picked up our commemorative tickets, we waited to board the outbound locomotives. We had both opted to go out in a standard loco and have the return journey in the Flying Scotsman. Both days were incredibly hot with 28°C for me on 30th June and a sweltering 32°C for Dewi on the 12th July! My trip out was a treat in its own right, being pulled by the "Erlestoke Manor".



Dewi's trip out was on a diesel; we think this was to reduce the risk of trackside fires due to the extraordinary temperature that day.

The SWR heritage line has a speed restriction of 25 mph and so the journeys out and back took around 1 hour, quite sedate compared to



what the vehicles were actually capable of. This however gave plenty of time to take in the experience, enjoy the scenery and stand by the door for the breeze to keep cool! The carriages themselves were heritage of course and clearly designed more for heating rather than cooling!

The train worked its way up through the Somerset countryside passing places such as Watchet and Dunster before arriving at Minehead.



After a well-earned ice cream there was a short wait at the station watching the ground crew spray water on the coal to be used to refuel the Flying Scotsman prior to its return trip. The station master said that as the Flying Scotsman would be running in reverse on the return journey, the coal dust would otherwise be blown into the driver's eyes. He also explained that the imported coal, which is the only coal that is available nowadays, is dustier and doesn't burn as well as the Welsh anthracite coal that was used in the heyday of steam locomotives.

The Main Event!

We were then treated to the main event of the day as the Flying Scotsman pulled in slowly from its previous leg coming up from Bishops Lydeard. This was a hugely exciting part of the whole experience with amazing sounds of venting steam and the majestic movements of the pistons and wheels as the driver brought 100 tonnes of steam and metal into the buffers.

I was amazed however that the final moments of manoeuvring the immense engine had the driver with his head stuck out of the window and twisting a small knob so hot he needed a rag to protect his hand!

In this modern era where driver assist sensors are commonplace in vehicles, it's a stark reminder as to how much skill and touch would be required to operate these locomotives. This was starkly illustrated in 2023 when the loco collided with a stationary train during shunting operation at Aviemore station.

As mentioned earlier, for the return trip, the engine was simply moved to the rear of the set of carriages and then operated in reverse all the way back to Bishops Lydeard. The visibility for the driver must have been severely limited by the tender.



Overall, a hugely enjoyable experience albeit on a rather uncomfortably hot day.

It really is a delight to see this 100-year marvel of engineering still operating in its full glory and hopefully its national treasure status will let it continue to do so to enthral many future generations to come.

Many thanks to the WSR for hosting the event and thanks to Dewi for sharing his experiences and pictures of his day.

Report by Paul Hampton and Dewi Daniels


Images © Paul Hampton and Dewi Daniels

Safety Systems Crossword

Welcome to the 4th Safety Systems Crossword prize competition where you can test your knowledge of the world of System Safety!

Send you completed grids to crossword@scsc.uk or drop into the crossword entry box at SSS'26 by 1pm on 6th Feb 2026.

The first correct answer drawn from a hat at SSS'26 will win a £100 gift voucher!



Across

- 6 Branch of philosophy that explores moral principles (6)
- 8 A devastating flow of fast-moving hot gas and volcanic matter (11)
- 9 Daredevil famous for an epic parachute jump (11)
- 11 Common name for the aircraft that crashed 32 seconds after takeoff in June 2025 (10)
- 12 First name of a female automotive pioneer (6)
- 14 Project that developed non-speech communication between conversational AI Agents (10)
- 16 Surname of a female champion of electrical safety (7)
- 18 A massive AI infrastructure project (8)
- 20 Concorde-like supersonic airliner currently under development (8)
- 21 City that hosted the first SCSC seminar to be held outside of the UK (6)
- 22 Type of gate used in quantum computing (8)
- 24 Name of the 2023 UK Safety Act to protect children from harmful internet services (6)
- 25 Form of testing to check if changes negatively impact existing functionality (10)
- 26 Surname of a leading campaigner for justice following IT failures (5)
- 27 A family of AI models that includes Haiku, Sonnet and Opus (6)

Down

- 1 Type of buttress usually providing lateral support to large structures like cathedrals (6)
- 2 A launch and landing site for a new breed of aircraft (9)
- 3 A large-scale complex, primarily for the production of electric vehicle batteries (11)
- 4 Type of fibre used in the manufacture of a submersible that imploded in 2023 (6)
- 5 Name of the fictional rogue computer that said "I'm sorry Dave, I'm afraid I can't do that" (3)
- 7 Type of system misuse that might be reasonably anticipated (11)
- 8 Capable of being done or effected with the available mean (11)
- 10 City holding the first ever World Humanoid Robot Games (7)
- 11 Location of European waypoint "DVL" whose duplication led to an air navigation system outage (9)
- 13 Crane-like device often used to lower a ship's lifeboat (5)
- 15 Type of malware that encrypts the victim's personal data (10)
- 17 Failed blood testing company whose female CEO was imprisoned for fraud (8)
- 19 A new remotely piloted aircraft system being developed in Europe (9)
- 20 Name of a general-purpose robotic humanoid with ambitions to leave Earth (7)
- 23 An unlucky accident (6)

Name: _____

Contact Nos/email: _____

Connect

The Newsletter and eJournal

Do you have a topic you'd like to share with the systems safety community? Perhaps an interesting area of research or project work you've been involved in, some new developments you'd like to share, or perhaps you would simply like to express your views and opinions of current issues and events. There are now two publishing vehicles for content – shorter, more informal content, can be published in the Newsletter with longer, more technical peer-reviewed material more suitable for the eJournal. If you are interested in submitting content, then get in touch with Paul Hampton for Newsletter articles: paul.hampton@scsc.uk or John Spriggs for eJournal papers: john.spriggs@scsc.uk

The SCSC Website

Visit the Club's website thescsc.org for more details of the Safety-Critical Systems Club including past newsletters, details of how to get involved in working groups and joining information for the various forthcoming events.



Facebook



Follow the Safety-Critical Systems Club on its very own Facebook page.

www.facebook.com/SafetyClubUK

X/Twitter

Follow the Safety-Critical Systems Club's X/Twitter feed for brief updates on the club and events: @SafetyClubUK



LinkedIn



You can find the club on LinkedIn. Search for the Safety-Critical Systems Club or use the following link:

www.linkedin.com/groups/3752227

Advertising

Do you have a product, service or event you would like to advertise in the Newsletter? The SCSC Newsletter can reach out to 1,000's of individuals involved in Systems Safety and so is the perfect medium for engaging with the community. For prices and further details, please get in touch with the Newsletter Editor.

SCSC Working Groups



[Security Informed Safety](#)



[Assurance Cases](#)



[Service Assurance](#)



[AI / Autonomous Systems](#)



[Data Safety Initiative](#)



[Safer Complex Systems](#)



[Multicore](#)



[Safe System Architecting](#)



[Safety Culture](#)



[Ontology](#)



[Safety Futures Initiative](#)



[Systems Approach to Safety of the Environment](#)



[Safety Management & Safety Management Systems](#)

The Safety-Critical Systems Club is committed to supporting the activities of working groups for areas of special interest to club members. The purpose of these groups is to share industry best practice, establish suitable work and research programmes, develop industry guidance documents and influence the development of standards.

Security Informed Safety



The Security Informed Safety Working Group (SISWG) aims to capture cross-domain best practice to help engineers find the 'wood through the trees' with all the different security standards, their implication and integration with safety design principles to aid the design and protection of secure safety-critical systems and systems with a safety implication.

The working group aims to produce clear and current guidance on methods to design and protect safety-related and safety-critical systems in a way that reflects prevailing and emerging best practice.

The guidance will allow safety, security and other stakeholders to navigate the different security standards, understand their applicability and their integration with safety principles, and ultimately aid the design and protection of secure safety-related and safety-critical systems.


Earlier this year, the working group published its first piece of guidance: [Co-Assurance of Safety and Security](#) and presented a [poster](#) at SCSC Symposium in Feb'25.

Lead Stephen Bull stephen.bull@scsc.uk

Assurance Cases

The Assurance Cases Working Group (ACWG) has been established to provide guidance on all aspects of assurance cases including construction, review and maintenance. The ACWG will:

- Be broader than safety, and will address interaction and conflict between related topics
- Address aspects such as proportionality, rationale behind the guidance, focus on risk, confidence and conformance
- Consider the role of the counter-argument and evidence and the treatment of potential bias in arguments



One of the working group's activities is the maintenance of the Goal Structuring Notation (GSN) Community standard.

See scsc.uk/gsn for further details.

In May 2021, the group published v3.0 of the standard: scsc.uk/scsc-141C

In Aug 2021, the group published v1.0 of the Assurance Case Guidance: scsc.uk/scsc-159

Lead Jane Fenn jane.fenn@baesystems.com with support from **Phil Williams** phil.williams@scsc.uk

Service Assurance



Risks presented by safety-related services are rarely explicitly recognised or addressed in current safety management practices, guidelines and standards.

It is likely that service (as distinct from system) failures have led to safety incidents and accidents, but this has not always been recognised. The Service Assurance Working Group (SAWG) has been set up to produce clear and practical guidance on how services should be managed in a safety-related context, to reflect emerging best practice.

The group published guidance v4.0 in Jan 2024: scsc.uk/scsc-156D


Lead Kevin King kevin.king@baesystems.com

Safety of AI / Autonomous Systems

It is clear that AI and autonomous systems can introduce many new paths to accidents, and that autonomous system technologies may not be practical to analyse adequately using accepted current practice. Whilst there are differences in detail, and standards, between domains many of the underlying challenges appear similar and it is likely that common approaches to core problems will prove possible.



The SCSC Safe AI Working Group (SAIWG) aims to capture cross-domain best practice and guidance on key topics within the design, evaluation, assurance, and approval of safety systems that use or are developed using AI, bringing together emerging standards and key results from the incredible amount of research being conducted into AI safety.



The working group was kicked-off at SSS'24 and is led by Alan Simpson. The SAIWG will conduct regular meetings, workshops, and publications to share knowledge and experience on various topics related to AI and safety systems, such as coordination of safety with other disciplines, evaluation of risk, and mapping of terminology and language.


The group builds on the earlier work of the Safety of Autonomous Systems Working Group (SASWG) that culminated in production of the 4th version of its guidance: Safety Assurance Objectives for Autonomous Systems, in Feb 2024 scsc.uk/scsc-153C

Lead Alan Simpson alan.simpson@ebeni.com

Data Safety Initiative

Data in safety-related systems is not sufficiently addressed in current safety management practices and standards.

It is acknowledged that data has been a contributing factor in several incidents and accidents to date and there is foreseeable harm that can arise from Machine Learning and Large Language Models' use of data by Artificial Intelligence (AI) systems that are subject to issues of biasing, interpretation and, arguably, falsification. There are clear business and societal benefits, in terms of reduced harm, reduced commercial liabilities and improved business efficiencies, in investigating and addressing outstanding challenges related to safety of data.



The Data Safety Initiative Working Group (DSIWG) aims to have clear guidance on how data (as distinct from the software and hardware) should be managed in a safety-related context, which will reflect emerging best practice.

An update to the guidance (v3.7) was published in Jan 2025: scsc.uk/scsc-127J

Lead Mike Parsons mike.parsons@scsc.uk

Safer Complex Systems



The Safer Complex Systems Working Group (SCSWG) builds on the IET/RAE work already done in this area. It is recognised that the RAE work is ongoing and collaboration is encouraged. The group's mission is to produce practical guidance on developing, managing and assuring complex systems throughout their lifecycle (so as to achieve and justify their safety).

The following provides a statement of the problem:

- It is acknowledged that complex systems are becoming more prevalent with more opportunities to cause harm
- There are also new complex systems arising from using combinations of existing systems and services which are then used for safety purposes
- Complex systems are not sufficiently addressed in current safety management practices and standards
- In particular, complex interactions and emergent behaviours are not currently assessed and managed sufficiently
- There could be benefit in developing new analysis, tools and techniques to manage complex system risks
- There are clear business and societal benefits, in terms of reduced harm, reduced liabilities and improved business efficiencies, in improved management of complex systems risk


The group produced the first version of Guidance on the Through-Life Assurance of Complex Systems (scsc.uk/scsc-200) in January 2025.

Contact the group lead if you would like to attend future meetings or find out more.

Acting Lead Mike Parsons mike.parsons@scsc.uk

Multi- and Manycore Safety

It is becoming harder and harder to source single-core devices and there is a growing need for increased processing capability with a smaller physical footprint in all applications. Devices with multiple cores can perform many processes at once, meaning it is difficult to establish (with sufficient evidence) whether or not these processes can be relied upon for safety-related purposes.



Parallel processes need to access the same shared resources, including memory, cache and external interfaces, so they may contend for the same resources. Resource contention is a source of interference which can prevent or disrupt completion of the processes, meaning it is difficult to know with a defined uncertainty the maximum time each process will take to complete (Worst Case Execution Time, WCET) or whether the data stored in shared memory has been altered by other processes.

The Multi- and Manycore Safety Working Group (MCWG) has been established to explore the future ways of assuring the safety of multi- and manycore implementations.

Lead Lee Jacques Lee.Jacques@leonardocompany.com

Safe System Architecting

The SCSC Safe System Architecting Working Group (SSAWG) is a joint initiative between the INCOSE UK Architecture Working Group and the Safety Critical Systems Club, with a vision of addressing systems safety through the promotion of choosing, using, and developing appropriate architecture. Two focal areas are defined as:



- Safety-driven architecting – what safety drivers are likely to exist when architecting is being considered
- Architecting system safety – what architectural factors may realise, enable, support, or preclude the achieving of safety considerations

The groups Mission Statement is:

- To produce practical guidance on addressing system safety during system architecting.
- To encourage and facilitate collaboration between INCOSE and the club.

The working group meets every 6 weeks online via Teams. Note that the INCOSE UK AWG meets every quarter.

If you would like to be involved in the group, please contact the co-chairs.

Co-chairs: Siyuan Ji s.ji@lboro.ac.uk and Jane Fenn jane.fenn@baesystems.com

SCSC Safety Culture

The Safety Culture Working Group (SCWG) has been established to provide guidance on creating and maintaining an effective safety culture. The group seeks to improve safety culture in safety-critical organisations focussed on product and functional safety, by sharing examples and latest approaches collated from real-life case studies.

Meetings provide an opportunity to discuss any particular aspects attendees are interested in taking forward, and to help set future directions for the group.

In Dec 2023, the group published a position paper for assessing and managing safety culture. scsc.uk/r2995



A seminar was held 19th June 2025 in London and focussed on how Safety Culture practices need to change to accommodate Artificial Intelligence. See page 23 for an event report and SCSC members can access the seminar presentations and videos here: scsc.uk/e1156

Lead Anne Seldon anne.seldon@wae.com

Ontology

The Ontology Working Group (OWG) develops ontologies that will form the basis of SCSC guidance, as well as having wider industrial and academic applications.

During system development, and especially with Model-Based Systems Engineering (MBSE), it is essential to build a model of the world that the system will inhabit, and this involves developing and structuring concepts, terms and their relationships. An ontology is an explicit specification of such a conceptualization and provides the foundation for MBSE and architectural descriptions in general. Ontologies are also increasingly important, if not essential, for explainable Artificial Intelligence (AI), enterprise knowledge systems and the standards that underpin them.



The OWG is currently working on the definition of an ontology of risk and value for application in guidance for risk-based decision making – notably safety and security. The framework for modelling is the Unified Foundation Ontology (UFO) and OntoUML, a domain-specific language (DSL) for modelling in UML tailored for ontological modelling based on UFO.

The OWG has also embarked on a series of SCSC Newsletter articles to gently explain ontological concepts and notations. The aim is to provide an accessible introduction to the domain so readers can be more informed and to help demystify a topic that may have a perception of being quite academic. The second article in the series can be found on page 17.

Lead Dave Banham ontology@scsc.uk

The Safety Futures Initiative

The Safety Futures working group meets once a month to bring together fresh perspectives, innovative ideas and insights.

The Safety Futures is diligently working on developing a comprehensive roadmap to guide new safety professionals in exploring career paths across various industries. To supplement this, they have an upcoming initiative to create a University 'League' Table to rank universities based on their safety engineering courses and the career pathways they support.



They are also looking to start a reverse mentoring programme where we can match more experienced people with new professionals for information exchange.

The group encourages anyone to get involved to help shape the future of safety engineering careers.

Lead Khadijah Khatun khadijah.khatun@scsc.uk

Systems Approach to Safety of the Environment



The Systems Approach to Safety of the Environment Working Group (SASEWG) is a new group intending to apply Systems Safety practices to systems that are embedded within the natural environment, while focussing on that environment.

The group aims to produce clear guidance on how engineered systems should be developed and managed throughout their entire lifecycle so as to preserve, protect and enhance the environment.

Please get in touch with the working group lead if you would like to join or find out more about this group.

Lead James Inge james.inge@scsc.uk

Safety Management and Safety Management Systems

This new working group aiming to cover Safety Management and also Safety Management Systems (SMSWG) is now in full swing with almost 40 members and held its second working group meeting on 20th August 2025.

Safety Management can be a difficult challenge with many different stakeholders with their own objectives pulling in different directions, perhaps one might draw parallels with the equally challenging task of herding cats! This analogy has given rise to the group's icon as shown.



Being an effective Safety Manager requires a balance of technical knowledge, regulatory understanding, soft skills, and leadership. Yet many are asked to lead safety without prior experience or clear support. This group will explore challenges such as overcoming organisational resistance, gaining senior buy-in, building trust across the workforce, and creating a positive safety culture where reporting is safe and encouraged. The group will also explore regulatory frameworks, going beyond compliance to fully understand the intent behind the rules.

The group is led by Si Hays so get in touch with him if you want to find out more about the group.

Lead Si Hayes si.hays@scsc.uk

SCSC Membership

The SCSC provides a range of services to the System Safety community including seminars, tutorials, leadership events, specialist topic working groups, the annual symposium and a comprehensive body of publications. Membership brings many valuable benefits such as free access to online events, the SCSC Newsletter and access to presentations and other resources from events.

Individual Membership

To become an individual member of the SCSC please register on the SCSC website using the  icon at the top right of any page and select "Register". Complete and save your account registration and then verify your email address. Once registered and logged in click the link "why not join the SCSC..." inviting you to become a member at the top right of the page or select "Pay membership" from the  icon.

Individual membership can be paid online using a credit/debit card through our secure payment partner Realex Global Payments or contact Alex King for other payment methods. For student or retired member rates please contact Alex King to get your account status changed.

Corporate Membership

Your company contact with the SCSC should arrange the membership and any renewals for your organisation. To join as a member covered by a corporate membership, register as per the instructions for an individual member and then contact Alex King to confirm your affiliation.

Renewing Membership

You should be notified by email when your membership is almost expired or shortly after it has expired. These notifications will contain a link to the online renewal page or you will be able to renew when logging onto the website through the 'click to renew' link.

Membership Fees

The following fees are applicable from January 2025 for new and renewing members:

- 1 year Individual Membership: £149
- 2 year Membership: 8% discount: £275
- 3 year Membership: 16% discount: £375
- 1 year SFI Membership: FREE for first year, £35 for years 2 & 3
- 1 year Membership, retired member rate: £35
- For Corporate Membership discounts contact Alex King

A one-month Publication Pass is also available for £15. This allows access to all SCSC website publications in a particular calendar month.

Contact Alex King using office@scsc.uk

The SCSC Steering Group

Current Members



Stephen Bull
stephen.bull@scsc.uk



Dewi Daniels
dewi.daniels@scsc.uk



Paul Hampton
paul.hampton@scsc.uk



James Inge
james.inge@scsc.uk



Khadijah Khatun
khadijah.khatun@scsc.uk



Mark Nicholson
mark.nicholson@scsc.uk



Wendy Owen
wendy.owen@scsc.uk



Davy Pissort
davy.pissoort@scsc.uk



Sean White
sean.white@scsc.uk



Carmen Carlan
carmen.carlan@scsc.uk



Jane Fenn
jane.fenn@scsc.uk



Giles Howard
giles.howard@scsc.uk



Brian Jepson
brian.jepson@scsc.uk



Alex King
alex.king@scsc.uk



Yvonne Oakshott
yvonne.oakshott@scsc.uk



Mike Parsons
mike.parsons@scsc.uk



John Spriggs
john.spriggs@scsc.uk

Club Positions

The current and previous (marked in italics) holders of club positions are as follows:

Managing Director

Mike Parsons 2019-

Tim Kelly 2016-2019

Tom Anderson 1991-2016

Steering Group Chair

Dewi Daniels 2024-

Roger Rivett 2019-2024

Graham Jolliffe 2014-2019

Brian Jepson 2007-2014

Bob Malcolm 1991-2007

Programme & Events

Coordinator

Mike Parsons 2014-

Chris Dale 2008-2014

Felix Redmill 1991-2008

Manager

Alex King 2019-

Honorary Solicitor

Dai Davis 2022-

Newsletter Editor

Paul Hampton 2019-

Katrina Attwood 2016-2019

Felix Redmill 1991-2016

University of York Coordinator

Mark Nicholson 2019-

Website Editor

Paul Hampton 2025-

Brian Jepson 2004-

eJournal Editor

John Spriggs 2021-

Administrator

Alex King 2016-

Joan Atkinson 1991-2016

Diversity, Equity and Inclusion (DE&I) Lead

TBC

Louise Harney 2024-2025

Wendy Owen 2023-2024

Safety Futures Initiative Leads

Khadijah Khatun 2024-

Zoe Garstang 2019-2024

Nikita Johnson 2019-2021, 2023-2024

Calendar

October '25

M	T	W	T	F	S	S
		1	2	3	4	5
6	7	8	9	10	11	12
13	14	15	16	17	18	19
20	21	22	23	24	25	26
27	28	29	30	31		

November '25

M	T	W	T	F	S	S
					1	2
3	4	5	6	7	8	9
10	11	12	13	14	15	16
17	18	19	20	21	22	23
24	25	26	27	28	29	30
30						

December '25

M	T	W	T	F	S	S
1	2	3	4	5	6	7
8	9	10	11	12	13	14
15	16	17	18	19	20	21
22	23	24	25	26	27	28
29	30	31				

January '26

M	T	W	T	F	S	S
		1	2	3	4	
5	6	7	8	9	10	11
12	13	14	15	16	17	18
19	20	21	22	23	24	25
26	27	28	29	30	31	

February '26

M	T	W	T	F	S	S
					1	
2	3	4	5	6	7	8
9	10	11	12	13	14	15
16	17	18	19	20	21	22
23	24	25	26	27	28	

March '26

M	T	W	T	F	S	S
						1
2	3	4	5	6	7	8
9	10	11	12	13	14	15
16	17	18	19	20	21	22
23	24	25	26	27	28	29

April '26

M	T	W	T	F	S	S
		1	2	3	4	5
6	7	8	9	10	11	12
13	14	15	16	17	18	19
20	21	22	23	24	25	26
27	28	29	30			

May '26

M	T	W	T	F	S	S
				1	2	3
4	5	6	7	8	9	10
11	12	13	14	15	16	17
18	19	20	21	22	23	24
25	26	27	28	29	30	31

June '26

M	T	W	T	F	S	S
1	2	3	4	5	6	7
8	9	10	11	12	13	14
15	16	17	18	19	20	21
22	23	24	25	26	27	28
29	30					

July '26

M	T	W	T	F	S	S
		1	2	3	4	5
6	7	8	9	10	11	12
13	14	15	16	17	18	19
20	21	22	23	24	25	26
27	28	29	30	31		

August '26

M	T	W	T	F	S	S
					1	2
3	4	5	6	7	8	9
10	11	12	13	14	15	16
17	18	19	20	21	22	23
24	25	26	27	28	29	30

September '26

M	T	W	T	F	S	S
	1	2	3	4	5	6
7	8	9	10	11	12	13
14	15	16	17	18	19	20
21	22	23	24	25	26	27
28	29	30				

Events Diary



21 Oct 2025 International Workshop on Assurance for Software-intensive Systems (ASSURE) Sao Paulo, Brazil sites.google.com/aiit.ac.jp/assure2025/	28 Oct 2025 SCSC Working Group Meeting Safer Complex Systems WG Meeting #22 online scsc.uk/e5061	4 Nov2025 SCSC Working Group Meeting OWG Working Group Meeting, #100 online scsc.uk/e5064	4 Dec 2025 SCSC Seminar New Developments in Rail Safety Systems Munich, Germany scsc.uk/e5004
13 Nov 2025 High Integrity Software Conference (HISC) International Convention Centre (ICC) Wales www.his-conference.co.uk	25-26 Nov 2025 High Integrity Software Conference (HISC) International Convention Centre (ICC) Wales www.his-conference.co.uk	10-12 Feb 2026 SCSC Symposium Safety Critical Systems Symposium SSS'26 York, UK scsc.uk/sss'26	23 Apr 2026 SCSC Seminar Soft Skills for Safety Practitioners London, UK scsc.uk/e5055



The logo consists of four red squares arranged horizontally, each containing a white letter: 'S', 'C', 'S', and 'C'.

SCSC

For Everyone Working in System Safety

Safety-Critical Systems Club (SCSC)

A top-down view of a car simulation showing a white car and a black car on a road, with various blue and white data overlays representing sensor data and system status.

The SCSC is the professional network for sharing knowledge about system safety

A sleek, white high-speed train with a colorful stripe, traveling on a track under a blue sky with clouds.

Bringing together engineers and specialists from a range of disciplines working on safety critical systems

A black and grey quadcopter drone with orange propellers, shown from a side-on perspective.

The SCSC provides learning events, with a 3-day conference every February. It also provides a website, supports industrial working groups, and publishes a newsletter and journal

A man and a woman wearing safety glasses and lab coats, working together on a piece of industrial machinery.

www.thescsc.org