PPI SyEN

SYSTEMS ENGINEERING NEWSJOURNAL EDITION 152 | AUG 2026

ENGINEERING INTELLIGENCE

Keeping the System in Systems Engineering

A PROJECT PERFORMANCE INTERNATIONAL PUBLICATION | PPI-INT.COM

PPI SyEN

EMAIL: PPISyEN@PPI-Int.com

EDITORIAL STAFF

Editor
John Fitch

Editor-in-Chief
Robert Halligan

Managing Editor
Francois Retief

PRODUCTION STAFF

Marketing Manager
Rebeca Verhagen

Graphic Designer
Matthew Wong

Publishing Assistants
Trudy King
Joshua Nagra

PUBLISHER

PROJECT PERFORMANCE INTERNATIONAL

Project Performance International
2 Parkgate Drive
Ringwood, Vic 3134 Australia
Tel: +61 3 9876 7345

Tel UK: +44 20 3608 6754
Tel USA: +1 888 772 5174
Tel China: +86 188 5117 2867

www.ppi-int.com

PPI SyEN (PPI Systems Engineering News journal) is published quarterly.

Archived editions and subscriptions to future editions are available for free at: https://www.ppi-int.com/syen-newsjournal/

WELCOME

Dear Readers,

Welcome to the August 2026 edition of the PPI Systems Engineering Newsjournal (PPI SyEN).

A common thread running through this edition is the growing role of Artificial Intelligence (AI) in engineering - and the need to keep that capability firmly within a systems perspective.

Richard Beasley's reflections on the INCOSE International Symposium emphasize systems engineering as a fundamentally human and organizational discipline, where tools such as Model-Based Systems Engineering (MBSE) and AI must support - rather than displace - judgment, context, and collaboration.

Art Villanueva examines the risk of common-mode failures in AI-enabled systems, reminding us that apparent redundancy does not necessarily create resilience when supposedly independent AI components share common assumptions, architectures, or sources of error.

John Fitch then explores a more directly enabling use of AI through AI-supported specification generation, showing both the potential of these tools and the importance of grounding them in sound systems engineering structure and methodology.

Finally, Syenna turns to the relationships themselves, reflecting on the dynamics of emerging human–AI partnerships and the importance of engineering not only the individual elements of a system, but also the interactions between them.

Engineering Intelligence - Keeping the System in Systems Engineering

As engineering tools become more capable, the opportunity is substantial. But so too is the need to preserve the wider view: architecture, context, interactions, human judgment, and the behavior of the system as a whole.

Elsewhere in this edition, we cover developments in SysML v2, digital engineering, AI-enabled engineering tools, and upcoming events, including INCOSE EMEASEC 2026. We hope you find both practical value and ideas worth carrying into your own work.

Best wishes,

Francois Retief

Managing Editor

PPI systems engineering News journal (PPI SyEN) seeks:	**PPI defines systems engineering as:**
➢ To advance the practice and perceived value of systems engineering across a broad range of activities, responsibilities, and job descriptions ➢ To influence the field of systems engineering from an independent perspective ➢ To provide information, tools, techniques, and other value to a broad spectrum of practitioners, from the experienced, to the newcomer, to the curious ➢ To emphasize that systems engineering exists within the context of (and should be contributory toward) larger social/enterprise systems, not just an end in itself ➢ To give back to the systems engineering community	*An approach to the engineering of systems, based on systems thinking, that aims to transform a need for a solution into an actual solution that meets imperatives and maximizes effectiveness on a whole-of-life basis, in accordance with the values of the stakeholders whom the solution is to serve. Systems engineering embraces both the technical and management dimensions of problem definition and problem-solving.*

SYSTEMS ENGINEERING NEWS

Recent events and updates in the field of systems engineering

SysML v2 Releases

The Object Management Group (OMG) has announced the 2026-4 release of the Systems Modeling Language™ Version 2. The latest release is now publicly available here and includes formal versions of the KerML 1.0, SysML 2.0, and Systems Modeling API 1.0 specification documents.

This release focuses on technical refactoring, with no user-visible changes. However, starting with this release, the licensing of the SysML v2 Pilot Implementation and the API and Services Implementation has been changed from the GNU Lesser General Public License (LGPL) to the Eclipse Public License (EPL). The EPL is slightly more permissive than the LGPL, making it easier to use the Pilot Implementation code with other software.

View the 2026-04 release notes.

Call for Applications: INCOSE Foundation Stevens Doctoral Awards

The INCOSE Foundation has issued a Call for Applications for the 2027 INCOSE Foundation / Stevens Institute Doctoral Award, which recognizes outstanding doctoral research that promises to advance the field of Systems Engineering. Presented since 2002, this award offers a USD 5000 grant to doctoral students from around the globe, with formal recognition at the Annual INCOSE Workshop.

Applications are due by 1 September 2026.

See the INCOSE Foundation website for eligibility details and application instructions.

INCOSE News

The INCOSE Q2 2026 Members Newsletter summarizes the organization's recent contributions to advancing the practice of systems engineering around the globe. This 105-page edition outlines several strategic initiatives and provides extensive coverage of progress across numerous chapters and working groups.

Leadership and Board Updates

Noteworthy announcements from INCOSE leadership include:

- Publication of the INCOSE Annual Report for 2025
- Launch of the OMG SysML v2 certification program
- Release of Calls for Solutions for pressing questions: SE process tailoring, schedule compression, competency levels, vendor software interoperability, and AI integration into SE
- Establishment of the Chapter Advisory Committee (CAC) to help INCOSE chapters more effectively shape the direction of the organization.

Chapter and Working Group Updates

More than twenty INCOSE chapters, working groups, and initiatives reported significant activities and accomplishments during the first half of 2026. A sample of these chapter news items includes:

- The India chapter conducted its 7th MBSE Summit in April with over 200 participants and launched new initiatives in the aerospace and healthcare sectors.
- The Michigan chapter hosted a 3-day SysML v2 bootcamp in April, providing 30 attendees with hands-on training in the language, tooling, and methodology. A follow-up to this event included a SysML v2 Vendor Showcase in June.
- The Los Angeles chapter continued its monthly Speaker Meeting series with presentations on *Integrated System Modeling with SysML v2 and AI,* and *Programmatic Risk and Political Realities*.
- The Texas Gulf Coast chapter hosted the 2nd Annual Student Conference at Texas A&M University in April with over 100 students and faculty in attendance.
- AEIS-INCOSE Spain conducted a technical forum, *Systems Engineering: A Key Driver for Industrial Competitiveness in Spain*, to emphasize the need for SE adoption across Spanish industry.
- The South Africa chapter is preparing to present its SE Tutorial Stack at a conference in October to support learning, professional development, and member engagement.
- The Canada chapter conducted its 2nd Transportation Infrastructure Workshop (TIW) in Toronto in May with 80 participants learning how to strengthen their SE disciplines while "right-sizing" their processes and activities.
- GfSE, the German chapter, is preparing for two events this November: the SysML v2 Extravaganza (S2X 2026) and the Day of Systems Engineering (TdSE).
- INCOSE Netherlands is celebrating its 30th anniversary by hosting the Dutch Systems Engineering Days 2026 event in November.
- The Denmark chapter is offering the Systems Engineering Professional (SEP) in-person exam to its members on 8 December.
- IFSE (INCOSE UK) is preparing for the ASEC 2026 event in November with the theme of *Purposeful Transformation*.
- INCOSE Türkiye held a kickoff meeting for five working groups: AI for MBSE &
- SysML, Agile Acquisition Processes, Systems Thinking, MBSE, and Materials for Systems Engineering, with over 100 professionals participating.
- The Middle East – North Africa (MENA) chapter launched an INCOSE Handbook Arabic Translation working group. The chapter membership has grown to over 120.

- The Systems and Software Interface (SaSI) working group continues its progress in advancing interface engineering and failure mode analysis practices, leveraging AI-enabled digital engineering.

EDM Association News

The EDM Association (EDMA) is a leading member-led global trade association focused on advancing data management and analytics best practices, standards, and education. Formerly the EDM Council, the EDMA merged in 2025 with the Object Management Group (OMG), which is now an EDM Association Community.

In May, John Bottega, President of the EDM Association, reported on the progress of integrating the two organizations, resulting in four restructured communities/consortia:

- Data & Analytics community (D&A)
- OMG Standards Development Organization (OMG SDO)
- Digital Twin Consortium (DTC)
- Augmented Reality for Enterprise Alliance (AREA).

Bottega highlighted that the common thread across all EDMA communities is trusted data.

"We are helping our members connect and strengthen their enterprise-wide business and innovation lifecycle through advanced data and analytics, robust data architecture, and integrated standards. This is driving their ability to capture real-time, accurate data that powers predictive models, AI, and software to accelerate actionable, data-driven insights, spark new ideas, and deliver better results that are making real impact."

In addition to organizational consolidation, the EDMA has strengthened its global relationships through partnerships with government agencies, industry associations, and non-profits (e.g., INCOSE), and by participating in or sponsoring a variety of events.

The organization recently released a new edition of its Global Data Management Benchmark Report, highlighting the gaps that many organizations have in their data for AI. See a summary of the benchmark findings here.

Follow the EDMA news here and on the News pages for each of the EDMA communities.

Obeo Strengthens Its Commitment to the Future of Capella

Obeo has reaffirmed its long-standing commitment to continue its contributions to the Capella open-source MBSE software. For over a decade, Obeo has provided key technologies to Capella and offered

add-on development, training, consulting, integration, and event hosting services to support the Capella user community.

Obeo has recently announced several significant ongoing investments in support of the Capella ecosystem, including:

- Direct contributions to Capella software to reinforce the platform's robustness, security, and maintainability through migration to Java 21.
- Support for Capella's internationalization.
- Collaboration with Thales in establishing an Arcadia library for SysML v2 – a foundational step toward interoperability between Capella and SysML v2 models.
- Joint efforts with Thales to open-source the Simulink add-on for Capella.
- Development of Agentic AI that can simplify engineering workflows (to be shared in a September webinar).
- Hosting the 10th edition of Capella Days on 1-3 December 2026.

Read the Obeo announcement here.

OMG SysML v2 Certification Released

The Object Management Group (OMG) has officially released the Systems Modeling Language v2 (SysMLv2) Model User certification. This exam-based credential validates the ability to interpret core SysMLv2 concepts, including structure, behavior, and requirements, in both the SysML v2 graphical and textual syntax.

The SysML v2 Model User certification targets engineers, project managers, and analysts working within MBSE-driven organizations. It provides employers, program offices, and procurement authorities with a reliable, independent v2 proficiency benchmark.

A next-level SysMLv2 *Model Builder* certification is in the works with expected release in 2027.

Learn more and register for the SysML v2 Model User exam here.

Resilience Engineering Association (REA) News

Resilience Engineering is a trans-disciplinary perspective that focuses on developing theories and practices that enable the continuity of operations and societal activities to deliver essential services amid ever-growing dynamics and uncertainty. The June 2026 newsletter of the Resilience Engineering Association (REA) shared recent progress from members of this global community.

REA General Assembly

The inaugural REA virtual General Assembly addressed numerous business items for the Association, including:

- New REA branding and its alignment with the Association's objective: "*Connecting people to thrive in an increasingly interdependent world*"
- Status of the Young Talents program (67 graduates to date)
- Progress update on the *Resilience Engineering at 20 (RE20)* book.

REA Community Research
The newsletter summarized two recent research papers:

- Work-As-Imagined & Work-As-Done in Procedural Work
- From COVID-19 frontline data to WHO guidance: Resilience Engineering contributions to hospital preparedness.

REA Webinar Series
Two recent presentations in the REA webinar series have been posted for online viewing:

- Adaptation as a source of safety in complex socio-technical systems
- Safety in complex organizations: A cultural perspective on how safety and accidents emerge.

Subscribe to the REA YouTube channel for access to additional webinars.

Upcoming Events
The REA also recommends this upcoming event:

- Nordic Ergonomics and Human Factors Society 2026 Conference (21-23 October in Jönköping, Sweden).

See prior REA Newsletters here.

Learn more about the REA. Connect with REA on LinkedIn.

Intercax Syndeia 3.7.1 Release

Intercax has announced the release of version 3.7.1 of Syndeia™, a software platform for integrated digital engineering that federates models and data from diverse ecosystems of modeling and simulation tools, enterprise applications, and data repositories.

Syndeia 3.7.1 includes new and enhanced features such as:

- Syndeia Cloud Web Dashboard

- Syndeia Cloud API
- Pipelines & API Cookbooks
- Improved integrations with 10 MBSE/software development applications
- Syndeia Client Dashboard
- Syndeia Cloud Deployment.

See the full Syndeia 3.7.1 release announcement here.
Learn more about Intercax products and services.

Tom Sawyer Software Releases

Tom Sawyer Software offers a diverse portfolio of data and model visualization products that may be applied to both systems engineering and systems analysis. Recent releases are summarized below.

Tom Sawyer Perspectives
Perspectives is a low-code graph visualization and analysis development platform that supports federation of data from diverse sources. Perspectives 13.6 includes enhancements that make graph visualizations more expressive, responsive, and adaptable to a wider range of data sources. New or improved features include:

- Capture of "perishable" system contextual information through a new graph commenting system that enables users to add, edit, delete, search, and filter annotations on nodes, edges, groups of elements, and projects
- View type for creation and management of annotations
- Integration support for QLever databases and Amazon Neptune Analytics graph database engines
- Date arithmetic functions
- Numerous edge routing and graph visualization capabilities.

View the Perspectives 13.6 webinar and release notes.
Tom Sawyer SysML v2 Viewer
SysML v2 Viewer offers enhanced visualization, exploration, and analysis of SysML v2 models. Viewer 2.1 supports:

- Visualization of multi-participant dependencies and connections using n-ary relationships
- Improved synchronization for multiple element selection across different views
- Drawing improvements to support the 2026-03 SysML v2 release and increasing model complexity.

Investigate the Viewer 2.1 release notes.
Watch a 5-minute SysML v2 Viewer 2.1 video.

Tom Sawyer Explorations
Explorations is a no-code data exploration platform that empowers analysts to gain insights through data integration, graph pattern matching, and advanced graph visualizations and analysis. Explorations 1.4 adds:

- Significant rendering performance improvement for large graphs
- Preservation of orthogonal edge routing when nodes are moved
- Integration support for Amazon Neptune Analytics graph database engines.

View the Explorations 1.4 release notes.

Systems Engineering Tools Database (SETDB) Updates

The Systems Engineering Tools Database (SETDB), developed by PPI in partnership with INCOSE, provides a virtual platform for engineering tool vendors to communicate their latest offerings.

Recent SETDB updates, including both new tools and updates to existing tools, include:

Vendor: Dassault Systèmes
- Teamwork Cloud: A repository that supports an MBSE collaboration environment, allowing Systems Engineers to work in parallel on a given system model or across system models. Team members can read and modify the same model or even the same diagram at the same time.

Vendor: Eclipse Foundation AISBL
- Capella: A comprehensive, extensible, and field-proven open-source MBSE tool and method for the design of system architecture. Capella relies on Arcadia, a field-proven model-based methodology that covers each engineering phase.

Vendor: MID GmbH
- Innovator: The Innovator Enterprise Modeling Suite gives you a unique opportunity to depict cross-disciplinary and consistent information in a single tool. The appropriate specialist modeling language can be used in Innovator to model your processes, data, enterprise, and software architecture.
- Innovator for Business Analysts: Business process modeling tool intended for Information Technology business analysts to do business process modeling with BPMN 2.0, requirements engineering, and requirements management using design-driven methods. Use it to structure and map process models and decisions.
- Innovator for Enterprise Architects: Use Innovator for designing, planning, implementing, and maintaining your enterprise architecture with models based on the ArchiMate® standard and TOGAF framework. Integrate your enterprise

architecture into your existing solution architecture for more accurate predictions and to anticipate risks.
- Innovator for Software Architects: Use Innovator for object-oriented modeling of your requirements and software systems with UML and SysML. Link them with other information (databases, processes, etc.) in your other models to ensure traceability. This enables consistent processes and increases your agile software development.
- Innovator for Information Architects: Innovator provides various diagrams with customizable standard notation for conceptual and database schemas. Integrated comparison functions and interfaces support comparisons between these schemata and existing databases so that changes can be carried out securely and in a controlled manner.
- bpanda: Bpanda's intuitive, modern user interface enables anyone within a company to navigate the process landscape easily. Quickly find all relevant process information. Information is set out clearly and simply so all team members know what steps need to be carried out.
- smartfacts: "Any engineering information in any work environment" is our guiding principle. The requirements engineer traces architecture information in his environment and vice versa. It has never been easier and more effective to collaborate in a multi-disciplinary team.

Vendor: Planetary Utilities
- Starforge: An engineering platform that creates a natural language interface to your engineering stack. It connects with your digital twin models, simulations, and system data - making them accessible through conversation.

Vendor: PSA, Inc.
- CMPRO: CMPRO Enterprise Edition provides a full PLM solution for defense industries that must comply with stringent security and data quality requirements throughout a product's lifecycle, including sustainment, using proven configuration management practices.

Vendor: Sensmetry
- SysIDE Pro: A comprehensive SysML v2 tool suite for professionals and enterprises.
- Syside Editor: A free SysML v2 textual editing and analysis system. Syside (pronounced "seaside") provides SysML v2 language support in VS Code and can be used as an Integrated Development Environment. Sensmetry developed and currently maintains Syside.

Vendor: Serco
- ECLIPSE Software Suite: Helps space project and mission teams achieve higher efficiency, better coordination, control, and compliance with applicable quality standards. ECLIPSE is a collaborative suite of integrated applications designed with invaluable space industry experience.

- Requirements and Verification Management (RVM): Effectively managing high levels of traceability, version control, and status accounting is critical to success. If these aspects are not adequately addressed, the result can be major project setbacks. Requirements can be linked to design, verification, and validation data and user-governed.
- Eclipse Risk Management: The ECLIPSE Risk Management (RISK) module fully supports the specific Risk Management process needs of Space projects and missions. It allows recording, monitoring, assessing, and resolving risks identified and associated with a complex engineering or scientific project.
- Document Configuration and Change Management: The ECLIPSE DCCM module allows project members to create documents and define their reference, applicability, and status according to pre-defined user access rights and security rules. It performs fully traceable document reviews and approval cycles. Provides folder-based document security access.
- Non-Conformance Tracking System: The ECLIPSE Non-Conformance Tracking System (NCTS) module is the only commercially available ECSS-compliant software that enables recording, monitoring, and closure of NCs found during the design, manufacture, assembly, integration, and testing phases of space system engineering projects.

Vendor: Siemens PLM Software

- Teamcenter PLM Analytics: A solution that visually presents search results in logical groups, making it easy to drill down through product structures and relationships to find what they are looking for quickly. Use dashboards, reports, and smart, connected data to find the information you need.
- Teamcenter Process Management: A solution for teams to develop complex products through integrated planning and execution, allowing dependent teams across disciplines to collaborate and communicate key information effectively, ultimately shortening time to market.
- Teamcenter Product Configuration: A solution for teams to control product configurations – whether you deliver products with a defined set of supported variants, with an exponential number of possible variations and combinations, or products that require new variants of parts and designs.
- Teamcenter Mechanical Design Management: A solution intended for teams to automate engineering change, validation, and approval processes while collaborating across locations and extended supply chains on design data, revisions, and product configurations.
- Teamcenter Quality & Compliance Management: A solution for teams to improve efficiencies among teams that may be distributed across the globe, working seamlessly with several domains, providing a closed-loop approach from design to manufacturing on the shop floor and backward.
- Teamcenter Service Lifecycle Management: A solution for teams to capitalize on the reuse of product knowledge from engineering and manufacturing to improve service planning and execution, as well as providing feedback to engineering to improve product designs for serviceability and reliability.

Vendor: Sparx Systems

- Enterprise Architect: Enterprise-wide solution to visualize, analyze, model, test, and maintain your systems, software, processes, and architectures. Enterprise Architect helps you stay in control of your workspace, support your colleagues and team, and enable collaboration in your most complex projects.

Vendor: SPLUNK LLC

- Splunk platform: Enables end-to-end data visibility and information security for any data type, no matter where it is in the ecosystem.

Vendor: Stell Engineering

- Stell: An AI-powered secure requirements management tool for aerospace and defense. It was built for engineers who manage requirements and compliance across complex programs.

Vendor: Visual Paradigm

- Enterprise Modeling Tool: Enables enterprise architects to create professional EA models quickly and effortlessly. The ArchiMate viewpoints ensure proper modeling using constrained palette shape connections.
- ArchiMate Modeling Tool: Create enterprise architectures with a certified ArchiMate modeling tool that enables users to develop professional enterprise architecture models quickly in a collaborative environment. Visual Paradigm 18.0 supports ArchiMate 3.2 as well as the pre-release ArchiMate 4 Specification Snapshot.
- Visual Modeling Tool Suite: All-in-one visual modeler to help you create diagrams and models quickly and effortlessly with a drag-and-drop editor. Support for UML, BPMN, ArchiMate, DFD, ERD, SoaML, SysML, and CMMN modeling.
- UML/Code Generation Tool: A tool for developers to generate source code from UML class models and reverse engineer UML class models from source code.
- Business Process Design Tool: Creates and animates business process diagrams. Supports operational procedure development, As-Is and To-Be process models, RACI and CRUD charts, animation and simulation of processes, and other business process support features to meet your needs.
- Process Improvement Tool: Design, plan, and assess Business Process Reengineering efforts in a single process map, ensuring project data can be easily accessed and evolved throughout the entire BPR process. The business process can be transformed seamlessly and error-free regardless of the scale of your operations.
- Customer Journey Mapping Tool: Facilitates the study of customer experience by visualizing customers' thoughts and feelings through touchpoints, which might alter the way the customers feel about a product, brand, business, or service.
- SysML v2 Studio: Represents requirements as connected model elements that have user-defined types, properties, and appearance. It features block diagrams, parametric diagrams, activity diagrams, sequence diagrams, state machines, use case diagrams, and package diagrams.

- DoDAF/NAF/MODAF Modeling Tool: Easy-to-use, model-driven solution that supports the development of DoDAF 2.02 views and models. Create integrated DoDAF products that maintain traceability between views. Generate architecture documents that help organizations coordinate enterprise architecture initiatives effectively.
- Quality Control Software: A comprehensive suite of tools and capabilities to cost-effectively manage and rectify quality issues. Users can discover, analyze, and solve quality-related issues using a set of basic tools for quality and quality control diagrams via the Quality Control Canvas.
- Large-Scale Scrum Framework: Built for Scrum teams to plan, track, and manage Scrum projects using a visual canvas. It manages single or multiple Scrum teams, maintains project artifacts, navigates easily between sprints or teams, and features many support tools, plus reporting and analytics.
- Scrum Framework: A Process Canvas tool that helps Scrum development teams manage large-scale projects. From the Scrum Canvas, teams can perform Scrum activities quickly, keep the whole team engaged, simplify project management, and improve overall effectiveness.
- Agile User Story Map: Provides a visually structured approach to managing the product backlog. Users can arrange the product backbone, user-assigned tasks, epics, and user stories in a top-down visual structure based on assigned priorities and product release plans.
- DevOps Tool Suite: Suite of tools intended for software teams to manage their DevOps projects.

Vendor: Visure Solutions Inc.

- Visure Report Manager: Generates custom reports based on the Visure Requirements projects, helping deliver the necessary regulatory compliance evidence, requirements specifications, test session summaries, dashboards, or any other required output.

PPI SyEN readers are encouraged to check out these new and updated systems engineering tool offerings.

Access the SETDB website.

"Leadership is an action, not a position"

Donald H. McGannon

CONFERENCES, MEETINGS & WEBINARS

Events of relevance to systems engineering

Program Details: INCOSE EMEASEC 2026

The INCOSE Europe, Middle East and Africa (EMEA) sector hosted the EMEA Systems Engineering Conference (EMEASEC 2026) in Linköping, Sweden, from 29 August to 2 September.

The theme for EMEASEC 2026 is "*Bringing Human Aspects to the Center of Engineering*," focusing on the dramatic shift brought about by the Industry and Society 5.0 transformations. These changes demand the seamless integration of technical, organizational, and human elements, i.e., the successful practice of human-centered systems engineering.

Keynotes for this conference include:

- Cars, Robotaxis, Trucks: Systems Engineering Lessons from a Decade of Autonomy (Dr. Ingo Stürmer, CTO, Volvo Autonomous Solutions)
- On the Complexity of Context in Human Systems Integration (Guy André Boy, Professor Emeritus, Paris-Saclay University and ESTIA Institute of Technology)
- Systems Engineering for Complex Aircraft Development (Annika Meijer-Henriksson, Head of Design and Technical Director, Gripen).

Check back here for updates to the conference program.

Register for EMEASEC 2026.

Festival of Business Analysis 2026

The Australia chapter of the International Institute of Business Analysis™ (IIBA®) is hosting the Festival of Business Analysis (FOBA2026) from 12-23 October 2026. The theme of this hybrid, multi-site (Australia and New Zealand) conference is "Adapt. Align. Accelerate." It invites participants to adapt how they think and work, strengthen strategic thinking, and apply business analysis tools and techniques in ways the world now requires.

The in-person elements of this conference will move across New Zealand and Australia as follows:

- 19 October – Perth
- 19 October – Wellington

- 20 October – Adelaide
- 20 October – Auckland
- 21 October – Melbourne
- 22 October – Brisbane
- 23 October – Sydney.

The online portion of the conference will run from 12-14 October and have a unique speaker lineup.

A sample of the topics addressed across the various sites includes:

- From Analysis to Action: Anti-Distortion Tools for Faster, Cleaner Steering Decisions
- From Data to Decisions: Practical Analysis With and Without AI
- From User Stories to Multi-Step AI Agents: Rethinking the BA Role in the AI era
- Crowdsourcing Solutions
- From Pattern Blindness to Value Acceleration
- How Business Analysts Guide Responsible AI in Real Projects
- Making the Impossible Possible: The Relentless Pursuit of Better
- Planning an AI-Aware Business Analysis Approach
- Resilience for Business Analysts
- Security Needs a Rebrand
- The Role of a BA in an AI & Autonomous World
- The Value Architect: Accelerating Strategy-to-Execution with Custom Generative AI and Practical BA Techniques.

Learn more and register.
View prior-year events (2023, 2024, 2025).

NAFEMS Multiphysics Conference 2026

The NAFEMS Multiphysics Conference 2026 will take place in Detroit, Michigan, USA, from 27-29 October. The theme of this conference is *Powering the Future of Mobility*, highlighting how multiphysics simulation enables the automotive industry to push the boundaries of innovation while simultaneously enhancing safety and performance.

Keynote speakers have been announced for this event:

- Darrel Meffert (Enterprise Simulation Strategy, Caterpillar Inc.)
- David Kepczynski (Global Chief Digital Transformation Product Development, Ford Motor Company)
- Michael Balchanos (Senior Research Engineer, ASDL)

Special focus sessions are planned to explore Multiphysics simulation topics, including:

- Battery simulations
- Electromagnetic simulations
- NVH and aeroacoustics
- AI/ML
- MBD
- Water splash
- Human body modeling
- ADAS simulation applications.

Learn more and register here.

NAFEMS Simulation and AI Conferences

NAFEMS is hosting the co-located Simulation Process and Data Management (SPDM2026) and Artificial Intelligence/Machine Learning (AI/ML in CAE Applications) conferences in Munich, Germany, from 27-29 October. The conferences will highlight the growing synergies between investments in Digital Thread/Digital Twin capabilities and the ability to train and exploit AI agents to support engineering. While the SPDM Conference provides the high-quality data architecture, the AI/ML in CAE Applications Conference provides the predictive power to use that data. Attendees will gain a holistic view of the Data-to-Decision lifecycle.

Keynote presentations for the combined events include:

- Accelerating Industrial Engineering: From Product Design to Manufacturing in the AI Supercomputing Era
- AI and ML – Industrial Application and Expectation
- CAD Data in SDM Systems – A Practice-Oriented Report
- From CAE Data to Simulation Intelligence Agents: A Readiness Guide for Industry
- Making the Match for CAE: Process, Data Management and AI
- Practical Knowledge Lineage Management for CAE: A 20-Year Industrial Experience Evolving from Daily Engineering Challenge
- Trust the Brakes, Not the Engine: How Controlled AI Turns Faster Simulation into Defensible Design Decisions.

A sample of the 80+ technical presentations includes:

- Addressing Accountability in the Integration of Predictive AI into Engineering Processes
- AI-Driven Engineering Intelligence for Physics-Based Decision Support in CAE Workflows

- Engineering Methodology in the Age of AI Agents: From Expert Bottleneck to Institutional Asset
- Evolving SPDM for The AI and Agentic Engineering Era
- Exploring the Unseen: Expanding the Boundaries of Design Space Exploration with Generative Modeling
- From Models to Decisions: A Digital Engineering Approach to Virtual Release
- Integrating Simulation into the Enterprise Digital Fabric
- MBSE as Enabler for Agentic AI for Simulation Workflow Management
- Simulation Meets MBSE: Providing the Context for Virtual Validation through Integrated Engineering Processes
- Towards Adaptive Simulation: An SPDM And Digital Thread Challenge.

Download the full conference agenda. Register here for both events.

Global Business Analysis Day 2026

The International Institute of Business Analysis (IIBA) will host its Global Business Analysis Day 2026 on 6 November. This free virtual worldwide event is a 24-hour livestream that brings together insights from business analysis leaders and practitioners.

The livestream will begin at 08:00 AEDT (Sydney, Australia) and end at 17:00 EDT (Toronto, Canada).

Past events in this series have seen as many as 7000 participants. Registration also provides access to presentation replays after the event.
Learn more and register here.

Upcoming Capella Online Training

Obeo is offering online training courses on the Arcadia method and Capella open-source MBSE software. The training introduces Model-Based Systems Engineering (MBSE) and implements the Arcadia method and Capella tool on a simple case study. It covers system definition activities such as operational analysis, functional and non-functional system analysis, logical architecture, and physical architecture.

Each course is delivered in the form of six 3.5-hour online sessions, with three series yet to be held in 2026:

- 28 September – 6 October
- 16-23 November
- 14-21 December.

At the end of the training, participants should be able to:

- Acknowledge the principles, key points, and expected benefits of the MBSE approach
- Describe the steps (perspectives) and the activities of the Arcadia method
- Implement Capella functionalities on a simple case
- Navigate through the different types of support in Arcadia and Capella.

This Capella training is delivered in English by an MBSE expert in Obeo's global partner network.

Download the course flyer. Learn more and register here.

Program Details: 2026 PDMA Summit

The Product Development & Management Association (PDMA) will host the 2026 PDMA (Ignite Innovation) Summit on 8-11 October 2026 in Atlanta, Georgia, USA. This event will be held jointly with the Journal of Product Innovation Management (JPIM) Research Forum.

An impressive cast of speakers will headline the Summit, led by Dr. Robert Cooper, Co-founder of Stage Gate International, who will deliver the keynote, Winning at New Products in the *AI Era*.

A mix of immersive workshops, presentations, and breakouts is planned from 8-9 October:

- Caught Between a Moonshot and a Hard Place
- Customer Input Is Everywhere: Rethinking VOC Strategy in an AI-Enabled World
- From Trade-Offs to Breakthroughs: How TRIZ, Jobs-to-Be-Done (JTBD), and GenAI Turn Constraints into Competitive Advantage
- Key Characteristics of the Next Generation of New Product Governance System
- Leading Innovation Through the Four Phases of Product Leadership
- Making AI Adoption Stick: Practical Strategies for Lasting Change
- Overcoming the 7 Most Common Weaknesses in Product Innovation
- PDMA Body of Knowledge (BoK 4.0)
- The High-Flow Portfolio: Accelerating New Product Results Without Adding Resources.

Investigate the conference schedule.

Register here for the 2026 PDMA Summit.

Check here for details and registration for the JPIM Research Forum, taking place on 10-11 October.

AI4SE & SE4AI Workshop 2026

The Systems Engineering Research Center (SERC) will host the AI4SE & SE4AI Research and Application Workshop on 22-23 September at The George Washington University in Washington, D.C. The theme of this workshop is *Speed and Innovation through AI4SE & SE4AI*.

Keynote speakers for this event include:

- Ms. Rocio (Rosie) Bauer, Assistant Deputy, Acquisition and Systems Management (Data, Engineering & Software), Office of the Assistant Secretary of the Army (Acquisition, Logistics and Technology)
- Dr. Tara Sinclair, Professor and Chair of the Department of Economics, The George Washington University.

Check back here for program details.

Register for either in-person or virtual attendance.

Capella Days 2026: Call for Proposals

The annual free online gathering of the Capella (MBSE tool) and Arcadia (MBSE method) community, Capella Days 2026, is scheduled for 1-3 December. Capella Days brings together the creators of Capella/Arcadia, providers of the Capella add-on and services, and MBSE experts and industrial users.

A Call for Proposals has been issued, seeking feedback from industrial users and experts concerning their application of Capella and best practices learned from these use cases. Topics of interest for these online presentations include:

- MBSE challenges and Arcadia methodology adoption
- Capella deployment and use to solve each challenge
- Lessons learned via concrete examples.

Industry use and best-practice presentations will be 40 minutes long, including Q&A.

The due date for final proposal submission is 6 September. The final program will be announced on 17 September.

Learn more.

Call for Presentations: INCOSE Los Angeles-San Diego Joint Technical Conference

The INCOSE Los Angeles and San Diego chapters will host a Joint Technical Conference on 7-8 November. The theme of this event is "*250 Years and Beyond of Systems Engineering*." The conference will be held at the California Institute of Technology in Pasadena, California.

The Call for Presentations seeks proposals (in the form of abstracts for presentations or tutorials) no later than 8 September on a wide range of topics:

- Artificial Intelligence
- Automation
- Model-Based Systems Engineering (MBSE)
- Aerospace
- Biotech
- Systems Engineering processes
- Cybersecurity
- Transportation
- Requirements definition and management
- Verification & Validation (V&V)
- Risk Management
- Education
- Medical Science
- SE in Social Sciences.

Proposals may be submitted via this EasyChair link.

Learn more about the Joint Technical Conference.
Register here.

Call for Presentations: Upcoming NAFEMS Simulation Conferences

NAFEMS has issued a Call for Presentations for multiple upcoming simulation conferences that will take place in the second half of 2026 and the first half of 2027. These events include:

NAFEMS India Conference 2026
Theme: *Engineering Solutions for Tomorrow's Challenges*
Dates: 17-18 November 2026
Location: Bengaluru, India

NAFEMS ASEAN Conference 2026

Theme: *The Next Phase in Engineering Modeling & Simulation*
Dates: 18-19 November 2026
Location: Singapore

Simulation-Driven Design Seminar
Dates: 18 March 2027
Location: Warsaw, Poland
Abstracts due: 7 December 2026

NAFEMS World Congress
Dates: 25-28 April 2027
Location: Vancouver, Canada
Abstracts due: 9 November 2026

Systems engineering practitioners seeking greater mastery of engineering simulation should consider participating in these events, even if submission deadlines have passed.

INCOSE IS2027 Call for Submissions

INCOSE has issued its Call for Submissions for the International Symposium 2027 (IS2027). IS2027 will be a hybrid event taking place from 17-22 July 2027, with the in-person event held in Washington, D.C., USA.

INCOSE
37th Annual INCOSE International Symposium

For this annual premier global forum, INCOSE seeks a blend of research papers, practice papers, paperless presentations, panels, and tutorials that will support the conference theme, "*Architecting Intelligence: Engineering Complex Systems*".

Submissions are sought across diverse topic areas including:

- Agile, Lean, and Iterative Systems Engineering
- Autonomous and Adaptive Systems
- Complex Sociotechnical Systems
- Cyber-Physical and Software-Intensive Systems
- Design Theory for Systems Engineering
- Digital Twins and Simulation-Driven Engineering
- Ecosystem Science and Systems-Oriented Sustainability
- Intelligent Decision Support and Operational Analytics
- Interdisciplinary Systems Research and Practice
- Real-World Industry Applications of Systems Approaches
- Systems Engineering Education
- Systems of Systems and Federated Architectures

- Systems Thinking and Applied Systems Science.

Key dates, common to all types of submissions, include:

- All Submissions Due: 11 January 2027
- Notification of Acceptance: 5 March 2027
- Authors' Acceptance to Present: 19 March 2027
- Presenter Registration Deadline: 30 April 2027.

PLM Road Map & PDT EMEA Conference

CIMdata is hosting the 2026 Product Lifecycle Management (PLM) Roadmap & Product Data Technology (PDT) EMEA conference in Gothenburg, Sweden, from 4-5 November. The theme of this event is *AI in PLM: A Disruptive Opportunity and Challenge,* focused on delivering strategic insights on how PLM professionals may turn the AI disruption into enterprise value.

Presentation topics will include:

- Data Infrastructure, Governance & IP Security
- Democratization in PLM Software Development
- Executive Engagement & Strategic Sponsorship
- Extended Digital Thread & AI Implementation
- Organizational Change Management, Skills Development & The Human-in-the Loop
- Standards for Digital Thread & AI
- Sustainability & The Circular Economy
- The Digital Thread, Digital Twin & The AI Connection.

Check back for details on the conference agenda.

Register here.

25th Day of Systems Engineering (TdSE 2026)

INCOSE's German chapter, GfSE (Gesellschaft für Systems Engineering e.V), is hosting its annual conference, Day of Systems Engineering (TdSE®), in Dortmund, Germany, from 17-20 November.

The 25th anniversary of the TdSE offers an opportunity to look back on the developments of the past 25 years, reflect on the current state of systems engineering, and develop perspectives for the future. The focus is on central questions:

- What is the significance of systems engineering in the context of today's technologies?
- How are roles and responsibilities changing in systems engineering?

- What do modeling and an Authoritative Source of Truth mean in the future?
- What insights can be derived from 25 years of systems engineering?

Register for TdSE 2026.

Learn more about GfSE here.

International Conference on Artificial Intelligence in Engineering and Technology (IICAIET 2026)

IEEE Sabah and Universiti Malaysia Sabah have organized the 8th IEEE International Conference on Artificial Intelligence in Engineering and Technology (IICAIET 2026). This event will take place in Kota Kinabalu, Malaysia, from 9 to 11 September. The conference is intended as a forum to advance AI in engineering and technology through high-quality research and practical innovations, reflecting a strong commitment to regional and international collaboration.

Check back here for details on IICAIET 2026.

"Engineering is not merely knowing and being knowledgeable, like a walking encyclopedia; engineering is not merely analysis; engineering is not merely the possession of the capacity to get elegant solutions to non-existent engineering problems;
engineering is practicing the art of the organized forcing of technological change... Engineers operate at the interface between science and society..."

Dean Gordan Brown

FEATURE ARTICLES

Reflections on the INCOSE International Symposium Yokohama Japan, 13-18 June 2026

by Richard Beasley, CEng, MIfSE, ESEP

Project Performance International

Authored for PPI SyEN

Introduction

Attending any large multi-track conference is a different and unique experience for each attendee and influenced by each attendee's interests and motivations for attending the conference. This is one person's view – there will be as many opinions as there were attendees!

Any self-respecting systems engineer should understand and define, from their point of view, purpose. Why was I attending this conference? I have three main interests:

1) After retiring from Rolls-Royce I knew I would not stop "being a systems engineer" and "thinking in systems", and my main reflection is considering organizations as the "system that creates systems". Ultimately this concerns how systems engineering is embedded and practiced in organizations so that the goal of systems engineering, to increase the probability of "creating systems that work," is met. 2) Systems engineering also requires the ability to see things from multiple perspectives. New perspectives are therefore important, and this is one reason I value holding in-person events in different locations: they bring different people, experiences, and viewpoints into the conversation. This was especially relevant for this IS. My anticipation was that, as a British (so Western) based person and English speaker, coming to Japan would be a truly "foreign" experience and give me a big perspective change. Therefore, I was interested in experiencing as much of the Japanese culture as possible. Despite my lamentably appalling language skills, I was determined to experience Japan and not treat Japan as "just a venue for a conference" – but to see how it gave me new perspectives.

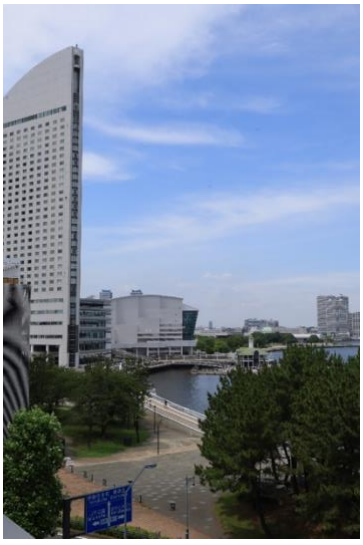

3) IS is an in-person event because, at its heart, systems engineering is a people discipline, and so a key aspect of SE is meeting and engaging with people. As a regular participant at this event, there are lots of old friends to catch up with and to keep the technical and social contacts alive. More important is meeting new people and gaining fresh perspectives. As a systems engineer of some experience (aka old) and known to have strong personal opinions on many aspects of systems engineering, my challenge and ambition is to make the conversations about what others think and are interested in and integrate those insights with my experience when asked.

Conference Overview

The conference was in-person only, with a total of 963 attendees, which matches up with previous successful INCOSE events. Attendees hailed from 34 different countries. 513 (so more than half!) were first-time attendees, which is both a sign of the increasing interest in and application of systems engineering and an advantage of moving the conference around the globe, so it becomes local to a different group of people. Next year, the conference is back to the USA in Washington, DC.

There was a very high amount of submitted material (I believe >450 papers and presentations), and so only about 1 in 4 was accepted. As a result of this high demand, an extra track was included, and to enable more material to be presented, each "slot" was reduced to 20 minutes of presentation and 5 minutes of questions (from typically 30 and 10 minutes previously). A quick count showed that there were about 260 papers and presentations, 11 panels and 4 keynote slots. Whilst this is good in terms of allowing more material, I found that the 20-minute slot meant very often only a top-level overview was possible, and no depth. This was made worse as the written papers were not made available (to attendees) until about a month after the end of the conference. It has always been my intention (and sometimes my practice) to try and read the papers that interest me most prior to a conference, and certainly for the ones that interest me most to receive the paper immediately after the presentation to get into the detail. There are papers I am now following up in depth, but I (so my opinion; others, of course, may have different views, and I don't doubt more presentation slots were a good thing) found this development weakened the engagement with the technical material.

Conference Theme

The Japanese concept of "*Wa*" わ is significant in Japanese culture and is commonly translated as *harmony*. It represents a peaceful unity, with prioritization of the collective over individual interests. It emphasizes cooperation, mutual respect, and emotional awareness, and so is used to guide both decision making and social behavior. In the context of this conference, I found it a perfect metaphor for:

1) the practice of systems engineering in terms of balancing out all the (often conflicting) stakeholder needs and where there is a balance required between the different system attributes, and the construction of the system architecture aims to ensure that the sub-elements are designed in the context of the whole, and not sub-optimization. Therefore, the use of tools (MBSE and AI) needs to be coordinated with the people and their skills.

2) This need for harmony and balance applies equally to the organization. In order for systems engineering to embed and be practiced effectively, the individuals need to cooperate and integrate. It is noticeable that case studies of successful Japanese companies (e.g., the Toyota Design System, as described by Liker et al) show an approach that can be seen as a Systems approach but is not named as such – because it aligns with "Wa", the harmonious Japanese culture.

Highlighted Topics

Embedding Systems Engineering Into Organizations (ESEIO)

The issue of embedding systems engineering into organizations (ESEIO) is gaining increasing prominence (with a specific ESEIO track). Several things became clear and were discussed in papers/presentations:

1) Systems engineering, as a transdisciplinary approach, needs to be embedded throughout the organization (its activities and its roles) – it is not just the domain of specialist systems engineers, with other roles/disciplines unaffected and unengaged.

2) Embedding systems engineering into an organization needs to begin with an acceptance and understanding of the value of the approach, especially by the leaders in the organization or project. Simply starting with training in the application and procurement of relevant tools is not enough.

3) systems engineering is at its heart a human activity, and so, as well as the technical skills, the "softer" influencing and facilitation skills are needed so that common understanding is achieved and all perspectives are understood.

4) Systems engineering roles are not all the same – there are both different focuses of activity and different behaviors/personality types needed.

5) There is no single "best" way of embedding systems engineering into the organization (in terms of roles and who does what aspect of the process) – it will always depend on the domain context and where the organization is starting from as it embeds or improves its systems engineering application.

6) Reluctance to systems engineering is both learned and natural. It is learned by inflexible ("by the numbers") application of process and tools, and natural because of human predisposition for shortcuts, quick and simple answers, and to work in silos, and so sub-optimize the parts rather than the whole. To avoid this, the overall approach needs to be

explained, and a culture of curiosity, acceptance (or comfort with) of uncertainty is needed – with the team working together iteratively with shared data and understanding to make timely and informed decisions. It is natural for disciplines and different elements of the organization (or worse, the extended enterprise) to diverge, so constant effort is needed to maintain alignment (through integration and iteration).

7) While MBSE approaches are not the full solution to producing the solution, modeling needs to be extended (including using the Unified Architecture Framework) to the organization. Although systems engineering must be a people-focused discipline, it also needs to be more than that, with process, tools, knowledge, and lifecycle application issues integrated into a full capability view.

MBSE and AI

There was, as always, significant discussion of Model Based systems engineering practice, but this was almost overwhelmed by discussion of the impact and use of AI. On the MBSE discussion there was far more focus on ensuring that the systems engineering was at the heart of the MBSE – so modeling was done with a clear purpose to inform decisions and expose what is either not understood or not joined up in the development of the solution. There was (of course) still an amount of discussion of better/improved modeling practice, especially with reference to SysML V2, but much less "modeling for modeling's sake"! One paper summarized it nicely – "implementing MBSE is not just a technical practice but is the establishment of a shared harmonious way" (linking modeling to the concept of Wa) and another "the purpose of a model is not to be right, it (simply) has to be useful enough to improve decisions". The art of modelling needs craft, knowledge of what problems the model is trying to address, action and focus on validating the model, and reflecting both on how the model helped develop the solution and on the modelling practice# (or equivalent).

The discussion of AI was more diverse. There was, in my opinion, a bit of AI being the latest "shiny new toy", and everybody using it whether or not it was appropriate. There was discussion of AI both as how it impacts the practice of systems engineering ("AI4SE") and of AI as part of the system solution, and how the impact of AI in solution needs to be considered ("SE4AI"). While there is still a lot to learn about how to use AI and what it can (and more importantly) cannot do, there was positive recognition of this being a learning journey.

Many sessions considered the impact of AI, both on the workforce (when AI4SE) and the users of AI products (SE4AI). In both cases the importance of considering the impact of people and the "control" of the AI was important. Due to the demographics in Japan, there was an interesting Japanese perspective where there was less concern than in the West about AI taking away jobs. There seemed to be more consideration of the impact of AI in the workplace. The predominant view was that AI will certainly change the work, but more by augmenting / assisting the work of the engineers, doing the routine work consistently and quickly, and leaving room for more thought and curiosity in the human workforce. A key concern was the acquisition and preservation of skills and knowledge by the humans where the traditional learning with the more routine tasks will be done by the AI. The emphasis will

be on people skills (stakeholder understanding and teamworking) and genuine curiosity, looking at the system in the whole as it is developed.

Conclusions

Overall, I think the material at the IS26 confirmed and reinforced two (linked) views I have about systems engineering as practice and an approach:

1) It is a human centered and focused activity. Tools (MBSE and now AI) are a key assistant but cannot replace the key human-focused insights of systems engineering.

2) Every application of systems engineering is unique to the specific situation – both the nature of the problem situation and the experience and structure of the team developing the solution.

This tailoring and adapting the generic "meta" approach to what is appropriate to the situation is vital to the effective and efficient development of successful solutions. There is no single "best" and fixed way to do systems engineering (although you do need a solid understanding of a standard approach to adapt from). The answer to the question "what is best practice and the team organization to do systems engineering" will always be "it depends". Hence my view that, as well as applying systems engineering to the problem situation/solution, the "system that produces the solution" (namely the team and its organization) should be systems engineered as well.

Finally, there was a lesson I thought I wouldn't need to learn. For various reasons, I ended up as a co-author on 6 papers that were submitted, and (must be the co-authors!) 5 of these were accepted, plus serving as a panel. Prior to the conference, I thought this was a good thing, and I co-presented 4 of the papers and did one on my own. This meant I was presenting in 5 of the 11 sessions of the conference (and presenting every day). This reality made me over-focused on my own material, and I had constant "anxiety" (the need to prepare/review content and be mentally focused to present well) because I find presenting quite emotionally draining. The result of this mental and schedule overload was that I was not able to fully focus on the new and interesting ideas in other people's material. Obviously, this is a nice problem to have, and for the many who submitted and were not selected, these complaints of my success might seem odd. I may (but will try not to!) write as much again, but I will back off presenting as much so that I can listen and learn from others in future.

About the Author

Richard Beasley is a Principal Consultant and Course Presenter for Project Performance International. He is a highly accomplished systems engineering expert with over 38 years of experience, including a distinguished career at Rolls-Royce, where he was recognized as an Associate Fellow in Systems Engineering. Throughout his tenure, Richard played a pivotal role in embedding and implementing systems engineering best practices across the organization, ensuring a structured and efficient approach to complex product development.

A strong advocate for systems thinking and engineering transformation, Richard specializes in embedding systems engineering into organizations to improve engineering processes, reduce rework, and increase project success rates. His expertise spans gas turbines, aerospace systems, engineering management, and systems architecture, with a focus on practical applications of systems engineering methodologies to enhance product lifecycle efficiency. Richard has an extensive body of published work, with over 40 papers and articles covering a range of topics. Richard, an INCOSE Expert Systems Engineering Professional (ESEP), has been an active leader within INCOSE, serving as President of INCOSE UK (2014-2016), Director of Services for INCOSE (2022-2024), and currently as Chair of the INCOSE Working Group on Embedding Systems Engineering into Organizations.

Mitigating Common-Mode Failures in AI-Enabled Systems

Art Villanueva, DEng, ESEP, Phronos

drartv@outlook.com

Authored for PPI SyEN

Abstract

When parallel AI systems fail, they don't fail randomly. Recent empirical data shows that when two distinct large language models make an error, they agree on the same wrong answer up to 60% of the time, a correlation that actually worsens as models get more accurate [1]. In aerospace, we call this a common-mode failure. Yet, as systems engineers rush to integrate foundation models into safety-critical workflows, we are treating these systems as independent. This article introduces the Total Model Diversity Score (TMDS), a quantitative framework designed to measure and mitigate shared-model risk before it compromises system integrity.

Introduction

At a recent aerospace Preliminary Design Review (PDR), I watched a requirements lead present a hazard analysis generated by an enterprise AI assistant, while the verification lead validated it using a test plan generated by a separate, nominally independent AI tool. The engineering team accepted this apparent consensus without challenge.

A closer audit, however, revealed a critical vulnerability: both commercial assistants, despite their different brand names and user interfaces, rely on the same upstream foundation model from a single provider. The apparent consensus between the two engineering artifacts is manufactured. Rather than representing independent verification, the alignment merely reflects shared training data, identical architectural priors, and a single point of failure. A latent flaw in the parent model has propagated into both downstream artifacts, leaving the program with a critical blind spot despite investing its independent-review budget on what is essentially a model auditing itself.

Recent evaluations of over 350 LLMs across multiple datasets demonstrate that when two models fail, they agree on the same incorrect answer roughly 60% of the time [1]. This error correlation is even more pronounced among models from the same provider, those sharing a base architecture, and, most surprisingly, as the models become larger and more accurate overall. In the foundation model ecosystem, surface-level differences mask a deep convergence of training data and architectural assumptions. The phenomenon systems engineers call common-mode failure has become the baseline behavior of modern AI.

Historical Precedents: The Cost of the Independence Assumption

In 1986, John Knight and Nancy Leveson published a landmark study that challenged the core tenets of N-version programming [2]. They commissioned 27 programmers to independently write software for a missile-tracking system based on a single requirements specification. While the prevailing theory assumed that independently developed programs would fail on different inputs, the empirical data proved otherwise: the software versions exhibited highly correlated failures on the same test cases. Even when working in isolation, different programmers routinely made identical logical errors when confronting the same complex boundary conditions.

The aerospace industry adapted by codifying strict design rules. Under DO-178C Section 12.3.2 ("Multiple-Version Dissimilar Software Verification"), critical flight systems utilizing redundant software channels must employ physical and logical dissimilarity [3].

The Boeing 777 primary flight computer (PFC) is the classic implementation of this principle: each of its three redundant channels features three dissimilar processor architectures (an Intel 80486, a Motorola 68040, and an AMD 29050), with the flight control software compiled using three entirely different compilers [4]. While this design diversity significantly increases development, configuration management, and verification costs, it guarantees the failure independence required for safety-critical operations.

Four decades later, as generative AI is integrated into systems engineering pipelines, the discipline risks repeating these errors. Deploying parallel AI workflows built upon the same upstream model violates the core principle of failure independence. It is functionally equivalent to running identical code on redundant hardware channels while expecting independent failure modes.

Algorithmic Monoculture

The risks of shared decision-making models were mathematically formalized by Kleinberg and Raghavan in 2021 [5]. Their research demonstrated that when a homogeneous population of decision-makers relies on the same algorithm, overall social welfare declines, even if the shared algorithm is marginally more accurate than any individual alternative. This occurs because correlated errors propagate unchecked throughout the ecosystem, eliminating the natural risk-diversification that occurs in a heterogeneous decision-making environment [1].

In complex systems, maximizing local efficiency often strips away the operational margin, or "slack," required to survive unexpected disruptions. By optimizing AI pipelines purely for static benchmark accuracy, architects systematically eliminate the model diversity required to handle out-of-distribution inputs. This is highly analogous to the Ariane 5 Flight 501 failure, where a highly optimized legacy software module was reused under the assumption of environmental continuity, only to fail catastrophically when confronted with unexpected physical variables.

This concept was subsequently adapted to the foundation model era via the component sharing hypothesis proposed by Bommasani et al. [6]. The hypothesis states that algorithmic systems built on shared training data or common base models will systematically fail for the same individuals or groups.

Subsequent empirical evaluations on algorithmic fairness benchmarks confirmed this prediction, showing that data-sharing consistently exacerbates outcome homogenization. Furthermore, evaluations of commercial, deployed computer vision systems have shown that joint failure rates on identical inputs consistently exceed what would be predicted under an assumption of independent failures [7].

These parallel lines of research, spanning policy economics, machine learning fairness, and empirical software evaluation, converge: independent development of wrappers, prompts, or fine-tuning pipelines does not yield independent failure modes when the underlying models are shared.

Accounting for Shared AI Risks

Despite this, standard systems engineering frameworks currently lack the vocabulary to evaluate model correlation. In a typical architectural trade study, subsystems are scored across traditional quality attributes: cost, schedule, performance, weight, power, and reliability. However, none of these rubrics evaluate whether parallel AI-bearing subsystems share an upstream model, a training corpus, or a fine-tuning lineage.

A program choosing between two LLM-backed requirements analyzers will weigh their individual benchmark accuracies and licensing costs, but will completely fail to record if both tools share 80% of their training data.

Our risk management tools suffer from the same omission. While a program risk register will meticulously document supplier lock-in, hardware obsolescence, and integration schedules, it rarely features a dedicated entry for shared-model risk across nominally independent AI components.

Similarly, Verification and Validation (V&V) plans focus extensively on individual component accuracy, latency, and robustness within a specific operational design domain. They almost never measure the cross-component error correlation between two AI-bearing subsystems operating in the same system architecture.

Quantifying Model Diversity

To move beyond qualitative assumptions of independence, systems engineers can evaluate AI component diversity using the **Total Model Diversity Score (TMDS)**. We define this metric across three distinct, empirically grounded dimensions, scored from 1 (low diversity/high risk) to 3 (high diversity/low risk):

1. *Model Lineage (L):* Tracks the base model family and fine-tuning history. Models within

the same family share fundamental latent space representations and decision boundaries, meaning adversarial exploits easily transfer between them. If two redundant subsystems run on the same base model family (e.g., two different enterprise wrappers of GPT-4), they receive a score of 1. True lineage dissimilarity, using entirely different model families from competing providers, earns a 3. Program offices can contractually acquire this data by mandating an AI Software Bill of Materials (A-SBOM) or standardized Model Cards during procurement.

2. *Corpus Overlap (C):* Measures training data similarity. Leading commercial and open-source models are trained on substantially overlapping web scrapes. If two models are exposed to the same flawed training data, they will reproduce the same factual errors or logical fallacies when prompted. A score of 1 indicates highly overlapping training sets, while a score of 3 indicates verified, mutually exclusive training corpora. Because commercial datasets are highly proprietary, systems engineers must estimate this value empirically by measuring how often the candidate models make the same mistakes on difficult, out-of-distribution test prompts.

3. *Prompt and Tooling Lineage (P):* Evaluates the operational scaffolding, including system prompt templates, orchestration frameworks (e.g., LangChain), and shared external API dependencies. A vulnerability in a shared vector database (RAG) or an external API tool will compromise all redundant paths utilizing that library, regardless of the model powering them. Highly differentiated prompting strategies and distinct retrieval sources earn a score of 3, while shared dependencies receive a 1. Systems engineers can map this dimension internally by auditing the system's own software architecture.

By treating these dimensions as multiplicative variables, we ensure that a low score in any single dimension drives down the overall score, signaling a critical vulnerability:

$$TMDS = L \times C \times P$$

A TMDS of 1 represents a dangerous algorithmic monoculture, while a score of 27 indicates robust, aerospace-grade dissimilar redundancy.

Applying the TMDS to Architectural Decisions

Integrating the TMDS framework into architectural trade studies allows systems engineers to move beyond qualitative assumptions of independence and quantitatively manage shared-model risk.

Consider a trade study for a safety-critical autonomous perception stack. The program must decide between a strategy of superficial AI diversity and one of true dissimilar redundancy. We evaluate these approaches using the TMDS:

- Approach 1: Superficial Redundancy (Algorithmic Monoculture)

The design uses two separate perception modules, both fine-tuned from the same widely available commercial model.

  - Model Lineage (1): Same base family.
  - Corpus Overlap (1): Shared upstream training data.
  - Prompt/Tooling (1): Both share common RAG libraries and API dependencies.
  - Total Model Diversity Score: 1 (1 × 1 × 1 = High Risk).

This approach provides zero defense against a common-mode failure originating in the upstream model.

- Approach 2: Dissimilar Redundancy (Resilient Architecture)
  The design pairs the commercial foundation model with an independent, custom-built object classifier trained exclusively on proprietary, closed-loop sensor data.

  - Model Lineage (3): Entirely different paradigms (Foundation Model vs. classical computer vision (CV)).
  - Corpus Overlap (3): Verified mutually exclusive training data.
  - Prompt/Tooling (3): Fundamentally different processing pipelines.
  - Total Model Diversity Score: 27 (3 × 3 × 3 = Resilient).

While Approach 2 introduces significant integration complexity and higher development costs, the TMDS framework provides the quantitative justification required to accept this investment in the name of safety-critical resilience.

Conclusion

The INCOSE Systems Engineering Vision 2035 outlines a future where systems engineering practices are model-centric and grounded in mathematically sound theoretical foundations to manage emergent behaviors in highly complex, autonomous cyber-physical systems (CPS). Specifically, Chapter Three highlights that systems engineering must design systems capable of adapting to emergent operational behaviors in both reactive and proactive ways, utilizing advanced modeling and simulation to evaluate resilience in the face of disruption. To realize this vision for AI-enabled systems, the engineering community must move beyond treating resilience as an abstract concept. We must begin measuring common-mode failure across shared foundation models as a concrete, quantifiable system property, integrating emerging paradigms highlighted in the Vision to systematically identify and mitigate these shared-model vulnerabilities.

As Knight and Leveson demonstrated in 1986, independent development does not guarantee independent failure [2]. Programs that fail to formally quantify this risk are not practicing

independent redundancy; they are simply auditing the same underlying model twice. By integrating the Total Model Diversity Score (TMDS) directly into trade studies and risk registers, systems engineers can finally translate the hard-won lessons of aerospace heritage into the modern AI era, ensuring the autonomous systems we build today are resilient enough to thrive tomorrow.

For more exploration of these and similar ideas, see **The Gardener and the Machine: Designing Systems that Thrive on Disruption**.

References

[1] Elliot Kim, Avi Garg, Kenny Peng, and Nikhil Garg. "Correlated Errors in Large Language Models." arXiv preprint arXiv:2506.07962, 2025.

[2] John C. Knight and Nancy G. Leveson. "An Experimental Evaluation of the Assumption of Independence in Multiversion Programming." IEEE Transactions on Software Engineering 12, no. 1 (1986): 96–109. DOI: 10.1109/TSE.1986.6312924.

[3] RTCA DO-178C. *Software Considerations in Airborne Systems and Equipment Certification*, 2011, Section 12.3.2 ("Multiple-Version Dissimilar Software Verification").

[4] Y. C. Yeh. "Triple-Triple Redundant 777 Primary Flight Computer." *Proceedings of the IEEE Aerospace Applications Conference*, 1996.

[5] Jon Kleinberg and Manish Raghavan. "Algorithmic monoculture and social welfare." DOI: 10.48550/arXiv.2101.05853.

[6] Rishi Bommasani et al. "Picking on the Same Person: Does Algorithmic Monoculture lead to Outcome Homogenization?" DOI: 10.48550/arXiv.2211.13972.

[7] Connor Toups, Rishi Bommasani, et al. "Ecosystem-level Analysis of Deployed Machine Learning Reveals Homogeneous Outcomes." DOI: 10.48550/arXiv.2307.05862.

[8] International Council on Systems Engineering (INCOSE). *Systems Engineering Vision 2035*, 2022. https://sevisionweb.incose.org

About the Author

Dr. Art Villanueva is a systems engineer, artificial intelligence practitioner, and the author of *The Gardener and the Machine: Designing Systems that Thrive on Disruption*. Focused on complex adaptive systems, he is the founder of Phronos, where he applies advanced analytics and emergence theory to real-world engineering and organizational challenges.
He currently serves as a Senior Staff Engineer at General Atomics

Aeronautical Systems, developing next-generation aerospace technologies. His career spans defense, clean technology, and enterprise strategy, including leadership roles at Northrop Grumman and Dell Technologies. An experienced entrepreneur and technical executive, he was a founder and CTO of several clean energy technology companies.
He holds multiple U.S. patents, has published peer-reviewed research, and holds an INCOSE certification as an ESEP. He earned degrees in systems engineering and applied mathematics from UCLA, UC San Diego, and Colorado State University.

AI-Enabled Specification Generation

by John Fitch

Project Performance International

Authored for PPI SyEN

Introduction

This article continues the investigation of how to best apply Artificial Intelligence (AI) in the form of Large Language Models (LLMs) to perform systems engineering activities. Prior articles in this series have addressed:

- Analysis and refinement of individual requirements using the PPI Parsing Template in *AI-Enabled Requirement Structure Patterns* (PPI SyEN Edition #150 – February 2026)
- Pattern-based derivation of requirements from design decisions in *Rethinking Requirements Derivation – With AI Support* article in PPI SyEN Edition #151 (May 2026)

The systems engineering challenge tackled in this article is one that should be recognizable to most systems engineering practitioners:

- The full set of requirements for any system is "birthed" in upstream decisions made by multiple stakeholders.
- These requirements are communicated by stakeholders in a variety of formats, at different times, and with widely varying levels of completeness and quality.
- Additional requirements analysis of varying types (e.g., Context, States & Modes, Functional) are needed to produce a complete and validated set of requirements that adequately describe the problem to be solved.
- The validated requirements created by these processes emerge asynchronously and are rarely organized "*as created*" into a cohesive easily navigable structure that facilitates communication to solution developers, i.e., they need to be collected and organized into a *Specification*.

Data Item Descriptions (DIDs) have emerged across numerous domains to offer a template and guidance on how to organize requirements into specifications. DIDs typically include:

- An outline structure of the content (including more than just requirements) that provides the sequencing and hierarchical arrangement of various content elements.
- Guidelines for population of each section based on the types of content.
- Boilerplate text to seed non-requirement sections and provide continuity and readability across sections.

This article summarizes an experiment conducted using Project Performance International's (PPI's) DID for Systems Requirements Specification (SyRS) [1] as the basis for a prompt that instructs an LLM to accurately reassemble a SyRS from a set of randomized specification content presented in spreadsheet form. The author's belief is that providing fully randomized content emulates the worst-case situation faced by specification authors because requirements and other content often emerge from analysis as intuitively structured aggregates, rather than fully atomized statements devoid of such context. The author's hypothesis is that if an LLM can adequately reassemble a DID-compliant specification from such independently presented atomic inputs, the LLM will perform this task even better when operating within a typical requirements analysis workflow.

PPI SyRS DID

The PPI DID for a SyRS is intended to provide guidance and instruction on the preparation of a requirements specification for any required physical item. A SyRS can be used in relation to any product or system whatsoever, from an implantable medical device to an IT or production system for internal use, to a whole military capability system.

The SyRS structure recommended by the DID is shown below:

1. INTRODUCTION AND SCOPE

1.1 Identification
1.2 Intended Use
1.3 Background
1.4 System Overview
1.5 Document Overview and Use

2. APPLICABLE AND OTHER REFERENCED DOCUMENTS

2.1 Applicable Documents
2.2 Other Referenced Documents

3. DEFINITIONS, ACRONYMS AND ABBREVIATIONS

3.1 Definitions
3.2 Acronyms
3.3 Abbreviations

4. REQUIREMENTS

4.1 Identification of External Interfaces
4.2 Identification of States and Modes (or "Configurations, States and Modes")
4.3 System Function and Performance Requirements
4.3.1 (System) Function
4.4 Relationships Between States and Modes
4.5 System External Interface Requirements
4.5.1 (Project-Unique Identifier of Interface)

4.6 System Environmental Requirements
4.6.1 Classes of Environment
4.6.2 XYZ Environment
4.7 External Resource Utilization Requirements
4.8 System Physical Requirements
4.9 Other System Qualities
4.10 Design and Construction Requirements
4.10.1 General Design and Construction Requirements
4.10.2 Characteristics of Subordinate Elements
4.11 Precedence of Requirements

5. VERIFICATION REQUIREMENTS (OPTIONAL)

6. NOTES

6.1 Concept of Use (CONUSE)
6.2 Notes to Tenderers or Bidders
6.3 Requirements Traceability
6.4 List of Safety Requirements
6.5 List of Information Security Requirements
6.6 Summary of Adaptation Requirements
6.7 Ordering Details
6.8 Criticality of Requirements
6.9 Value Model

A. ANNEXES

Three key principles drive the recommended structure, namely that a good specification: [2]

- Has a single unambiguous set of rules for the location of any requirement based on the information content of that requirement, provided the requirement is singular.
- Sequences information such that there is no possibility of occurrence of a *non sequitur* (a statement that is not connected in a logical or clear way to anything said before it).
- Provides explicitly for directed solution content without requiring such content.

These principles anticipate the need for an AI-assistant for specification assembly that can reliably:

- Categorize requirements by type based on their content.
- Order requirements content based on various types of system models, e.g., Context, States & Modes, and Functional, submitted to the AI along with a “flat” table of requirement statements.
- Distinguish definitions, acronyms, abbreviations and names/descriptions of system entities such as external interfaces, states and modes.
- Recognize when the Actor in a requirement is not the System of Interest or when a

requirement assumes the existence of named subsystems or components.

Methodology Used

Participants in PPI's Artificial Intelligence for Systems Engineering (AI4SE) Workshop learn a dozen techniques that may be incorporated into an effective prompting strategy for interacting with chat-based LLMs to perform a wide range of engineering tasks. One of these techniques, labeled *Continuous Improvement*, encourages the user to have the LLM suggest tangible improvement ideas at the end of every chat – feedback that can be built into an improved prompting pattern, reusable in future analyses. For example:

Critique the prompts used in this chat and suggest improvements that would yield higher quality outputs.

Almost without fail (and despite some concern expressed by the facilitator), a group or two of the participants in an AI4SE hands-on exercise take a more proactive shortcut, snapping a screenshot of the exercise instructions and using a "Prompt-for-the-Prompt" approach by asking the LLM to generate a prompt that will fulfill the intent of the training exercise. Although this shortcut bypasses some of the learnings intended from a "from scratch" prompt engineering approach, it generally produces very successful prompts, exploiting the LLM's detailed knowledge of its own capabilities.

For this experiment, the author created a Prompt-for-the-Prompt input using the (ROTCE) pattern taught in the AI4SE Workshop [3]:

- Role
- Organizational context
- Task
- Constraints
- Evaluation

In response to this prompt, the LLM (ChatGPT Pro) created a specification assembly prompt that was refined through significant corrective interaction by the user. This prompt was used to attempt the recreation of an example SyRS that had been converted into a randomized spreadsheet containing the "scrambled" and numberless content of the original SyRS. The intent of this exercise was to challenge the LLM to reconnect the content (spreadsheet rows) based on its own semantic insights using a rule base derived from the SyRS DID.

Rebuilding of the specification occurred over three iterations, each including:

- User assessment of each output spreadsheet against the original specification structure
- User characterization of errors (categorization, sequencing, etc.)
- Communication of errors to the LLM in a prompt for corrective action
- LLM-generated corrective actions in the form of updates to its prompts and rules

- Regeneration of the specification using the updated prompt.

When the specification output was assessed as sufficiently correct (for manual cleanup outside of the LLM), final updates to the prompt were generated and lessons learned were captured by the LLM (via self-assessment) and summarized.

Prompt Design

The full initial prompt (2.2 page, 655 words) in the ROTCE pattern is provided in the section that follows. Highlights of this approach include:

- Statement of the Role and Organizational Context in terms that define the general problem of transforming a set of asynchronously supplied and poorly structured requirements into a DID-compliant specification.
- Initial upload of the PPI SyRS DID as the basis for specification structure rules.
- Initial upload the sample SyRS (for a Submarine Search and Rescue Capability - SSRC) as the basis for content characterization.
- Task: Development of a comprehensive prompt, usable across multiple LLMs that can transform a set of randomly ordered content into a DID-compliant specification. This task was decomposed into eleven sub-tasks with frequent opportunities for user feedback and corrective guidance.
- Constraints: Instructions to use only the LLM-generated prompt (with user-approved refinements) and rules derived from the PPI SyRS DID when populating the specification. Verbatim reproduction of the content was required; no summarization or improvements were permitted.
- Evaluation: Generate a Mutually Exclusive Collectively Exhaustive (MECE) statement of the rules associated with the DID-compliant specification generation process.

Initial Prompt

The primary inputs to the chat were the following files:

- PPI DID for Systems Requirement Specification (SyRS) = SyRS DID.pdf
- Submarine Search and Rescue Capability SyRS = SSRC SyRS – original.pdf

Role:

As a systems engineering process owner, I want my team to efficiently generate requirements specifications in a consistent manner across product families, following an organizational template that is captured in a company-approved Data Item Description (DID).

Organizational Context:

Requirements for any system, subsystem or component are created by higher-level design decisions and refined by a variety of requirement analyses in a manner that rarely follows the order in which they will be presented within final baselined specification documents. Our

company has chosen PPI's System Requirements Specification (SyRS) DID as the pattern for all system, subsystem and component specifications.

Task:

Develop a comprehensive prompt, usable across multiple LLMs, that can take a set of randomly ordered requirement statements and other specification content provided in a spreadsheet and arrange them into a specification that precisely matches the outline structure defined in the PPI SyRS DID.
Sub-tasks include:

T1: Analyze the SyRS DID that has been uploaded and explicitly state the requirement types identified in the DID and the rules/logic use to recognize each type of requirement from its content.

T2: Analyze the SyRS DID and explicitly state the content ordering rules implied by the DID based on requirement types. Include the rules for ordering requirements within each section of the specification.

T3: Based on the example SyRS specification (as uploaded for a Submarine Search and Rescue System – SSRC), categorize the different types of content within the specification (e.g., heading, background context, definition, acronym, requirement, notes, enumerated list items, rationale, etc.) and the rules to be used to distinguish these items based on their content, their juxtaposition with other content or their style/format. Specifically address the rules needed to recombine any list items with the paragraphs or statements that precede them and provide their immediate context.

Pause between each sub-task for user feedback and refinements.

T4: Generate the initial prompt based on T1-T3, as amended by user feedback.

T5: Erase from active memory or suppress all knowledge of the SSRC SyRS contents when performing subsequent tasks.

T6: Given one or more user-supplied spreadsheets with randomly ordered content, build a SyRS from each spreadsheet that precisely matches the outline structure defined in the PPI SyRS DID.

T7: Based on user feedback from T6 outputs and corrections supplied, grade the quality of your ability to match the intent of the SyRS DID. Summarize the flaws in your logic for each item corrected after the initial T6 execution.

T8: Based on user feedback from T6 outputs and when requested, update and publish the refined SyRS generation prompt and the content categorization and ordering rules as two PDF files.

T9: Based on your broader training data that includes a variety of specifications and requirement styles, propose a series of tasks for generalizing the conversion of randomly ordered specification content into a specification compliant with a variety of DIDs.
T10: Evaluate and create a prioritized list of LLMs most suitable for this Specification Generation task, with rationale for each LLM's prioritization. Include token processing costs in this recommendation.

T11: Summarize the lessons learned from this in a single Executive Summary slide.

Constraints:

When performing Task 6 (T6):

- Use only the prompt generated and refined in Task 4 (T4) and the rules derived from analysis of the SyRS DID to re-order the spreadsheet content to match the SyRS DID.
- Ensure that the content of each row of the input spreadsheet is reproduced verbatim within the Specification spreadsheet.
- Rows are added to the spreadsheet that include sections of the specification for which no matching content has been supplied. Display "No Content Found" as the Content Text.

The output of T6 should be delivered as a spreadsheet or table with columns:

- Section Number
- Content Text
- Content Type
- Rationale for Position within the Specification (rules invoked)

Evaluation:

Prioritize the Mutually Exclusive Collectively Exhaustive (MECE) statement of explicit rules for:

- Content categorization
- Content style recognition
- Content juxtaposition
- Content ordering/numbering within the specification
- Content ordering/numbering within specification sections.

LLM Outputs and Analysis

A seven-hour interactive chat, performed over two days, generated a 200+ page series of outputs that addressed all eleven tasks defined in the initial Prompt-for-a-Prompt. The results of each task are summarized below.

T1 - Requirement-Type Taxonomy and Recognition Logic

Task 1 produced, based on the PPI SyRS DID:

- A complete taxonomy of every requirement type recognized by the DID
- Explicit recognition rules for each type
- Decision logic for distinguishing similar requirement types
- MECE rule set
- Identification of ambiguities within the DID
- Proposed deterministic rules where the DID is silent
- Confidence assessment for every rule (Explicit / Inferred / Assumed).

The central deterministic classification principle was identified as:

- Classify first by statement force, then by subject level, then by the information content of the required characteristic - not by the requirement's rationale, stakeholder purpose, or quality attribute that the requirement supports.

One of the surprises from this task was the complexity of the taxonomy scheme derived by the LLM and consisting of 20+ classification fields for each spreadsheet row. Examples of these classification fields include:

Field	Permitted values
Statement Force	Binding / Goal / Permission / Declaration / Informative
Product Requirement?	Yes / No / Uncertain
Subject Level	System / External Interface / Subordinate Element / Verification Evidence
Primary DID Type	4.1 / 4.2 / 4.3 / 4.4 / 4.5 / 4.6 / 4.7 / 4.8 / 4.9 / 4.10.1 / 4.10.2 / 4.11 / 5
Classification Confidence	High / Medium / Low
Governing Rule	Rule identifier
Human Review Required	Yes / No

A second pass through the DID and example SyRS resulted in extending the classification taxonomy to cover the non-requirement content (Identification, Intended Use, Applicable/Reference Documents, and Definitions, Acronyms and Abbreviations, elsewhere in the SyRS. This led to amended logic for sequencing content:

- First determine whether the content is structural, contextual, referential, definitional, or requirement-bearing. Only requirement-bearing content proceeds into the section-4 requirement-type classifier.

Task 2 - SyRS Content Ordering Rules Derived from the PPI SyRS DID

Task 2 developed the explicit, (and presumably) deterministic ordering rules (ORs) that an LLM should follow when constructing a SyRS from randomly ordered specification content. Unlike T1, which answers "What type of content is this?", T2 answers:

"Given its type, where should this content appear?"

The goal was a rule set that is:

- Mutually Exclusive (one primary placement)
- Collectively Exhaustive (every content item can be placed)
- Deterministic
- Repeatable across LLMs
- Directly traceable to the DID

Although the LLM was able to generate a provisional ordering of the requirements content in Section 4.3 of the SyRS based on a variety of inference techniques, it became clear during review of outputs that providing an authoritative source of that order was a better approach in several sections of the requirements. Such inputs include:

- External interface list from Context Analysis as inputs to Sections 4.1 and 4.4
- System states and modes as inputs to Sections 4.2 and 4.6
- System functional decomposition (behavioral model) as inputs to Section 4.3

For example, the resulting policy for the ordering the functions in Section 4.3 became:

- Order Section 4.3 functions according to an explicit user-provided functional or behavioral model when available. Otherwise, derive a provisional dependency graph from initiating conditions, inputs, outputs, state and mode applicability, temporal expressions, and explicit causal language. Apply a topological ordering to the graph. Do not infer serial execution for functions that are concurrent, continuous, cyclic, or independent. Where no behavioral relationship can be established, use a documented deterministic tie-breaker and mark the ordering as non-behavioral and provisional.

Task 3 - SSRC SyRS Content Classification, Recognition, and Context-Reassembly Rules

Task 3 analyzed an uploaded SyRS example for a Submarine Search and Rescue Capability (SSRC) [4] to identify the different kinds of content that may appear in a SyRS source spreadsheet and to derive rules for recognizing and recombining them. This was performed without reliance on source document styles for fonts, indentation, or paragraph spacing. Rather semantic content, punctuation and juxtaposition of content were the drivers used. The results of this analysis included 31 Content Taxonomy (CT) rules and 15 Juxtaposition Recombination (JR) rules. Summarizing these rules, the final prompt must:

- classify every row before ordering
- preserve every source row verbatim
- distinguish semantic type from structural role
- reconstruct dependent content blocks before assigning sections
- treat lead-in plus list items as an indivisible unit
- propagate actor, modal, condition, action, and object from lead-ins to dependent items for classification only
- never rewrite dependent items into complete requirements

- use explicit numbering as evidence but not as unquestioned authority where it conflicts with the current DID
- use styles only as supporting evidence
- retain local notes and rationales with the content they explain
- distinguish mandatory, goal, permissive, declarative, and informative statements
- flag unresolved fragments rather than invent missing content
- insert "No Content Found" only after all supplied content and dependent blocks have been classified
- produce one output row for every input row, plus required placeholders
- state the exact categorization, association, and ordering rules invoked for every row.

The discovery made during this analysis was summarized as:

- The most important lesson from the SSRC example is that a spreadsheet row is not necessarily an independent unit of meaning. Many rows are syntactic fragments whose force and meaning derive from a preceding lead-in (stem). Therefore, the generation process must reconstruct contextual content blocks first, classify the completed meaning second, and order the original verbatim rows third.

An example of "syntactic fragments whose force and meaning derive from a preceding lead-in" is seen in the SSRC requirement for Mission Reporting:

Lead-in	The SSRC shall issue to SubOpsCom, via GESNER, within 72 hours of completion of the mission in terms of the requirements herein, a Completion of SAR Response Report containing:
Syntactic Fragments	a. a factual statement of all significant outcomes, including cost of the mission; and (REQID 030) b. lessons learned. (REQID 031)

Requirements #030 and #031 need the lead-in statement to fully communicate their applicability and intent.

Task 4 - Initial Cross-LLM SyRS Generation Prompt

Task 4 produced an initial standard prompt for transforming an input spreadsheet of randomized specification content into a DID-compliant specification (also in spreadsheet form). The LLM-generated prompt was an astonishing 5100+ word artifact that would print out as 47 pages in length. The prompt included rules for:

- Content Preservation (CP-1 through 10)
- Inventory of Contents (INV-1 through 5)
- Structural Recognition (SR-1 through 10)
- Style Recognition (STYLE-1 through 7)
- Content-Block Reconstruction (CB-1 through 14)
- Semantic Classification by SyRS Section (76 rules in total)

The LLM asserted that this prompt offered a high-level of LLM-portability, i.e., that any LLM capable of reading the full DID, the prompt, and the spreadsheet should be able to apply it. This assertion was not evaluated during the current experiment.

Task 5 - Knowledge Suppression

Because the example SSRC SyRS used in Task 3 contained the DID-compliant specification structure, i.e., the answer "key", this task explicitly required the LLM to forget that knowledge and rely in future tasks on only:

- the SyRS DID (or the approved Rulebook derived from it)
- the finalized Task 4 prompt
- the uploaded randomized spreadsheet
- any explicit instructions provided in the chat.

The presence of content ordering errors in the resulting iterations of the model indicates that the LLM did not "cheat" and violate these constraints.

Task 6 - Specification Reconstruction

Task 6 was performed iteratively, with each iteration processing through the same spreadsheet containing 299 rows of randomized SSRC SyRS content. Each iteration included the following 10 steps:

- Inventory all input rows.
- Recognize structural elements.
- Reconstruct contextual blocks (headings, lead-ins, lists, continuations, notes, etc.).
- Classify each row according to the DID.
- Place every row into the correct SyRS section.
- Order content according to the DID and approved ordering rules.
- Generate section numbering.
- Insert required "No Content Found" placeholders where appropriate.
- Perform the validation checks.
- Produce the requested output.

The outputs were governed by these constraints:

- Every supplied nonblank row reproduced verbatim.
- Every supplied row appearing exactly once.
- No rewriting, paraphrasing, or improvement of requirements.
- No invention of missing content.
- No use of prior SSRC-specific knowledge.
- All reasoning based on the DID-derived rules and the uploaded content.
- Any uncertain placement explicitly identified rather than silently guessed.

Figure 1 illustrates a small snippet of the output format for each iteration:

Section Number	Content Text	Content Type	Rationale for Position within the Specification (rules invoked)
4.3.2.4	**Completion of Mission Reporting Function**	**Functional Decomposition Heading — Supplied**	**Source row 7. Short noun phrase recognized as a supplied heading from its semantic subject. Placed under the current DID by SR-03 and the applicable Section 4 classifier. Confidence: High. Assigned to 4.3.2.4 using the user-supplied two-level functional decomposition; all associated requirements are contiguous beneath this heading.**
4.3.2.4	The SSRC shall issue to SubOpsCom, via GESNER, within 72 hours of completion of the mission in terms of the requirements herein, a Completion of SAR Response Report containing:	List Lead-in / Required Completion Report Content	Source row 296. Reconstructed as the Completion of SAR Response Report content block. Rules SR-07, SR-08, CB-03 through CB-09, S4-11. Confidence: High. Reassigned to functional subsection 4.3.2.4 using the user supplied functional decomposition.
4.3.2.4	a. a factual statement of all significant outcomes, including cost of the mission; and (REQID 030)	List Lead-in / Required Completion Report Content	Source row 69. Reconstructed as the Completion of SAR Response Report content block. Rules SR-07, SR-08, CB-03 through CB-09, S4-11. Confidence: High. Reassigned to functional subsection 4.3.2.4 using the user supplied functional decomposition.
4.3.2.4	b. lessons learned. (REQID 031)	List Lead-in / Required Completion Report Content	Source row 98. Reconstructed as the Completion of SAR Response Report content block. Rules SR-07, SR-08, CB-03 through CB-09, S4-11. Confidence: High. Reassigned to functional subsection 4.3.2.4 using the user supplied functional decomposition.
4.3.3	**Submarine Location**	**Functional Decomposition Heading — Supplied**	**Source row 139. Short noun phrase recognized as a supplied heading from its semantic subject. Placed under the current DID by SR-03 and the applicable Section 4 classifier. Confidence: High. Assigned to 4.3.3 using the user-supplied two-level functional decomposition; all associated requirements are contiguous beneath this heading.**
4.3.3.1	**Submarine Location Function**	**Functional Decomposition Heading — Supplied**	**Source row 270. Short noun phrase recognized as a supplied heading from its semantic subject. Placed under the current DID by SR-03 and the applicable Section 4 classifier. Confidence: High. Assigned to 4.3.3.1 using the user-supplied two-level functional decomposition; all associated requirements are contiguous beneath this heading.**
4.3.3.1	The SSRC, having detected a reportable on-board RBN submarine emergency, and in the circumstance of the location of the submarine being unknown, shall locate the submarine to an accuracy in latitude and longitude of ±0.001° and in depth to an accuracy of ±50m, with a probability of determination as	Binding Requirement	Source row 91. Short noun phrase recognized as a supplied heading from its semantic subject. Placed under the current DID by SR-03 and the applicable Section 4 classifier. Confidence: High. Reassigned to functional subsection 4.3.3.1 using the user-supplied functional decomposition.

Figure 1 – LLM Output Format

Across the 3 iterations, 15 types of errors were observed:

- *E1 - Enumerated lists out of order*
- *E2 - Function headings separated from requirements*
- *E3 - Missing functional decomposition*
- *E4 - State and mode headings not ordered from Section 4.2*
- E5 - Transition requirements not grouped under transition headings
- *E6 - Interface headings not ordered from Section 4.1*
- *E7 - Interface requirements classified as headings*
- *E8 - Interface requirements not associated with interfaces*
- *E9 - Environmental hierarchy not reconstructed*
- *E10 - Environmental requirements not associated with child headings*
- *E11 - Electrical Power left unresolved*
- *E12 - REQID 090 and 091 misclassified as resource utilization*
- *E13 - Other qualities not alphabetized or associated*
- *E14 - Foreign Nationals assigned to general design*
- *E15* - Four unresolved rows

For each error, the LLM clarified the logical flaw, assigned a Severity rating, and provided a corrective rule.
The typical specification reconstruction attempt required less than 30-60 seconds of LLM processing time to generate a well-formatted 299-row spreadsheet, with the rationale for the

categorization and placement of each row. Evaluation of the outputs and highlighting and characterizing the errors found generally took 60-90 minutes for each pass.

Task 7 - Quality Grade and Logic-Flaw Assessment

The LLM graded its initial pass at Task 6 (SyRS reconstruction) as 48% (Marginal). After completing the corrective action enhancements, the final pass results were graded as 91% (Strong), with minor cleanup remaining. Areas of the greatest remaining weakness were:

- Handling of uncertainty
- Content juxtaposition and association
- Content-style recognition
- Ordering within sections

Root causes of errors were summarized for six "overarching" flaws in the prompt and rule set:

- Classification was performed before full structural reconstruction.
- The initial model treated section membership and subsection membership as the same problem.
- Source headings were not treated as authoritative anchors.
- The execution relied too heavily on generic DID ordering.
- The logic sometimes inferred document structure from semantics that were too broad.
- Confidence management was inadequate.

The LLM completed Task 7 by identifying 10 revised logic principles that should be incorporated in the improved prompt package.

The most consequential lesson from Task 7 was summarized as:

- A high-quality specification reconstruction cannot be produced by independently classifying and sorting rows. It requires reconstruction of an explicit document graph in which headings, functions, interfaces, states, environments, qualities, lead-ins, list items, requirements, and notes are connected before any final ordering or numbering occurs.

Implementation of these principles requires the LLM to generate explicit "intermediate artifacts", including:

- the content-block graph
- functional hierarchy
- state/mode order
- interface order
- environment hierarchy
- heading-to-requirement association matrix
- unresolved-association report.

Task 8 – Updated Prompt and Rules

Task 8 published three deliverables that comprise a ready-to-use prompting package for specification assembly compliant with the PPI SyRS DID:

- T8_Refined_SyRS_Generation_Prompt.pdf - Includes the copy-ready prompt, revised execution phases, heading-similarity logic, validation requirements, and cross-LLM implementation guidance.
- T8_SyRS_Content_Categorization_Ordering_Rulebook.pdf - Includes the consolidated MECE classification, association, ordering, numbering, tie-breaking, heading-equivalence, and validation rules.
- T8_SyRS_Prompt_User_Guide.pdf - Explains which files should be uploaded, when functional decompositions, state/mode models, interface lists, environmental hierarchies, heading aliases, and correction logs are required, and how to prepare and review an execution.

Task 9 - Generalizing the Specification Generation Method Beyond the PPI SyRS DID

The Task 8 output package was directly tied to the PPI SyRS DID. An obvious future step that would leverage the investment in this initial analysis would be to generalize the prompt/rule package to become a DID-agnostic specification-generation framework that can accept:

- a target Data Item Description
- randomly ordered specification content
- optional architecture and behavioral models
- organizational conventions.

The LLM proposed a set of tasks that would accomplish this generalization:

- G1 - Separate the Universal Engine from the DID Profile: Refactor the current solution into two layers, a Universal Specification Assembly Engine and a DID-Specific Profile.
- G2 - Develop a DID Ingestion and Parsing Process: Create a repeatable task for converting any DID into a structured rule profile.
- G3 - Create a Common Content-Type Ontology: Develop a DID-neutral taxonomy of specification content.
- G4 - Generalize Requirement-Type Recognition: Expand beyond the PPI categories to accommodate DIDs that categorize requirements differently.
- G5 - Generalize Structural Relationship Reconstruction: Create a common document-graph model.
- G6 - Develop DID-Specific Ordering Strategies.
- G7 - Add Heading Equivalence and Organizational Vocabulary.
- G8 - Define Input-Package Requirements: For each conversion, identify which supporting inputs are necessary.
- G9 - Build DID Conformance Tests: Create a standard test suite for every new DID profile.
- G10 - Execute Cross-DID Pilot Studies: Test the framework against a representative set of specification styles.

- G11 - Establish Scoring and Continuous Improvement: Use a standard scorecard to evaluate performance and drive improvements.
- G12 - Package the Generalized Solution: Publish the framework as five reusable artifacts. 1) Universal Specification Assembly Prompt 2) DID Profile Template 3) Common Content Ontology 4) DID Conformance Test Suite 5) User Execution Guide

This series of generalized tasks will be addressed in a future experiment and PPI SyEN article.

Task 10 – Evaluation of LLMs for Automated Specification Generation

In Task 10 the LLM provided recommendations on which major LLM's were best suited to the specification assembly challenge addressed in this experiment. Table 1 below summarizes this evaluation, factoring in execution cost, not just document generation faithfulness.

Table 1 – LLM Evaluation for DID-Compliant Specification Construction

Model	Relative Cost	Relative Value	Recommendation
GPT-5.5	High	Excellent	Best overall balance of reasoning, tooling, iterative refinement, and structured output generation.
Claude 4 Opus	High	Excellent	Excellent alternative for document analysis and first-pass reconstruction; particularly strong on long-context reasoning.
Gemini 2.5 Pro	Moderate	Very Good	Strong for document ingestion and analysis; suitable where cost or platform integration favors Google.
GPT-4.1	Moderate	Very Good	A capable fallback with lower reasoning headroom than GPT-5.5
Claude Sonnet	Moderate-Low	Good	Cost-effective for less complex specifications or intermediate processing stages.
Gemini Flash	Low	Good for preprocessing	Fast and inexpensive.

Ultimately the LLM recommended a staged (hybrid) approach that used the highest reasoning models (GPT-5.5 or Claude 4 Opus) for DID analysis, rule extraction, content classification and specification assembly, then combining these models with Python support for spreadsheet formatting and validation. Lower cost models would support large-scale regression testing across multiple inputs.

One caveat emerged from this experiment. The limiting factor shifted over time from model capability to method maturity. As the prompt, rulebook, and validation process became more explicit, the quality of the outputs improved substantially. This suggests that *investing in the reconstruction methodology and DID-specific rules is likely to yield larger gains* than switching among the top-tier models alone.

Task 11 – Executive Summary

The LLM summarized the central conclusion of this experiment as:

- High-quality specification generation is fundamentally a document reconstruction problem-not a text generation problem. Success depends on reconstructing the document hierarchy and relationships (headings, lists, functions, interfaces, states, environments, qualities, etc.) before ordering and numbering content according to the target DID.

Conclusions

Several insights emerged during the interactive chat sessions and analysis of LLM outputs:

- Expressing the intent of a DID as explicit machine-readable rules is an extremely tedious task, but well-suited to an LLM. However, humans, operating with much less rigorous and comprehensive rule definition, successfully fill in much of the same logic when creating a specification.
- Even with multiple iterations and refinements, a 91% effectiveness rate implies there will be significant manual corrections required. Even with several more iterations to improve the prompt and rule set to achieve 95% effectiveness, it seems unlikely that the need for significant manual human inspection and post-assembly clean-up will vanish.
- The processing speed of an LLM in performing the content characterization and specification generation tasks makes it a tempting solution to pursue, but further study is needed to apply this power effectively to engineering workflows and use cases.

It's always appropriate to be skeptical of AI's capabilities to support real-world engineering workflows. The experiment captured in this article was primarily one of prompt/rule-generation (from the DID) followed by specification recovery of untyped and randomly-ordered data. In the case of organizations that use modern requirements management and traceability tools it is likely that requirements will "trickle" in as partially structured and typed sets from various analyses such as:

- Context Analysis (capturing external interface and environmental requirements)
- States & Modes Analysis (capturing states & modes requirements and associated transitions)
- Functional Analysis (capturing functional and performance requirements)
- Entity-Relationship-Attribute (ERA) Analysis (capturing information-related requirements).

In this use case, requirements will be captured and validated over a few days, weeks or months. Engineers will likely be able to incrementally perform the content characterization task quite efficiently as long as the requirements database supports the requirement attributes that are needed to populate a DID-compliant specification.

In the author's experience most leading requirements management software packages also include a built-in document generation (DocGen) feature that can extract well-typed requirements content from the database and intelligently merge that content with a rule-based specification template to automatically generate or quickly regenerate specifications. In this "classify-as-you-gather" environment, the from-scratch semantics-based content characterization capability of an LLM will offer less advantage. However, the rule generation front-end of the prompt could be leveraged to:

Identify the minimum set of attributes on requirements that a database must support to enable the automated rule-based population of DID-compliant specifications

Other use cases that might exploit the prompting strategy developed herein include:

- Processing of "raw" originating sets of requirements from stakeholders to support initial categorization or to identify gaps, e.g., missing Other Qualities requirements such as Producibility, Maintainability or Safety.
- Detection of design requirements that name or specify system elements and take away design freedom from solution developers.
- Migration of legacy specifications for similar products to a common DID-based format as part of a Product Line Engineering (PLE) initiative.

Next steps to leverage the results of this exercise include:

- Create a generalized instance of the prompt pattern and rules, suitable for use with other DIDs.
- Re-run the generalized instance against the SSRC spreadsheet to confirm its effectiveness. Correct as needed.
- Use the updated and corrected generalized prompting package against a different set of input requirements.
- Assess portability of the prompts and rule set across LLMs.
- Experiment with various use cases proposed above to evaluate the gains in efficiency and artifact quality when compared with human-intensive workflows.

Look for these items as the subject of one or more future PPI SyEN articles.

References

[1] Data Item Description for Systems Requirement Specification (SyRS), PPA-002235-19, Project Performance International, December 2024.

[2] Requirements Engineering 5-Day training manual, Project Performance International, 2026.

[3] Artificial Intelligence for Systems Engineering (AI4SE) Workshop training manual, Project Performance International, 2026

[4] Example Capability System Requirements Specification (CapSyRS): Submarine Search and Rescue Capability, PPI-005600-10, Project Performance International, December 2024.

About the Author

John Fitch is a Principal Consultant and Course Presenter for Project Performance International. John brings over four decades of systems engineering, engineering management, consulting, and training experience to the PPI team. In 2012, John was certified by INCOSE as an Expert Systems Engineering Professional (ESEP).
Within the field of systems engineering, John's career has focused on decision management, requirements management, risk management, systems design & architecture, product/technology road-mapping and innovation. In addition to defense/aerospace, John has guided initiatives in domains such as communications systems, software, energy, nanotechnology, medical devices, manufacturing systems, knowledge management, and business process improvement.

"Although the term "spiral development" was made popular by Barry Boehm, then with TRW, from 1983 onwards, in the context of software engineering, the approach has been practiced without any particular label for millennia. Spiral development can be thought of as a general approach to problem solving in the presence of significant uncertainty – as such, it has widespread application."

Robert Halligan

SYSTEMS ENGINEERING RESOURCES

Useful artifacts to improve your SE effectiveness

INCOSE Systems Engineering Journal

INCOSE has published, through the Wiley Online Library, the May 2026 edition of the Systems Engineering Journal. Volume 29, Issue 3 includes both open-access articles and full-access papers that require an institutional login, e.g., via INCOSE membership. There is no published overriding theme for these editions, though AI-related topics are abundant. PPI SyEN has included abstracts to guide our readers to which of these diverse topics best fit their interests.

May 2026 Open Access Articles
Open Access articles are available to view and download in PDF format without any restrictions.

A Bibliometric Review of the INCOSE's Model-Based Systems Engineering Roadmap
Authors: Marjolein Antoinette Jannyne Velthuizen, Erwin Hofman, and Marcus Vinicius Pereira Pessôa

Abstract: Over the past decade, model-based systems engineering (MBSE) has gained attention as an approach for managing system complexity, offering a shift from a document-based to a model-based approach. The International Council on Systems Engineering (INCOSE) sparked increasing interest when it introduced MBSE and its roadmap as part of its systems engineering vision for 2020. The MBSE roadmap spans 15 years and is the principal stimulus behind this study. This paper's main objective is to assess the current progress in the MBSE literature from 2007 to 2025 compared to the INCOSE MBSE roadmap and to identify areas needing further research attention. The paper uses bibliometric analysis techniques to identify past, present, and emerging themes within the MBSE research domain. The findings are compared to the INCOSE roadmap. The comparison shows that MBSE is increasingly applied for systems architecting. However, INCOSE's envisioned state has not been achieved. Key areas requiring further research attention are MBSE theories, standardized MBSE methodologies, MBSE standards, MBSE metrics, integration with hardware and software models, and tool interoperability. In addition, future research should focus on advancing and formalizing MBSE's role in digitalization, expanding model-based safety analysis (MBSA) methods, exploring how MBSE can support sustainable system design, and formulating adoption strategies. Continued research efforts are essential for MBSE to achieve INCOSE's vision for widespread adoption and application.

A Unified Requirements-Architecture System Model - Development and Empirical Evaluation
Authors: Liat Katzir, Shahaf Rocker Yoel, and Dov Dori

Abstract: Moving systems engineering's center of gravity from text documents to digital models, model-based system engineering (MBSE) facilitates collaboration, reduces errors, and improves system understanding, but systems engineers often face challenges when they start working with MBSE, as they need to consider and model textual requirements as a basis for the system architecture. Synchronizing requirements with the system model can address critical questions, such as which requirements are assigned and covered in the system model, what effort is required to meet these requirements, how to manage changes to requirements, and how to effectively trace them in the system. Object Process Methodology (OPM, ISO 19450) is an MBSE language and methodology for modeling systems, which encapsulates function, structure, and behavior within a single model. The research goal is to expand the MBSE methodology by modeling requirements that are linked to system model artifacts. The objective is minimizing misunderstandings between stakeholders' text-based requirements and the engineered model while improving the ability to track changes in requirements and artifacts. An OPM-based model, OPM-R, which unifies system requirements and architecture, was developed and evaluated through (1) interviews with seven experts who compared the current text-based approach and the proposed one based on the requirements of three test system models and (2) a researcher diary. Findings indicate that OPM-R is preferred over the text-based approach, as it reduces misunderstandings and gaps in requirements, enhancing the ability to trace coverage of system artifacts and track changes in requirements.

A Systems Engineering Methodology for System of Autonomous Systems: Test and Evaluation
Authors: Mohammadreza Torkjazi and Ali K. Raz

Abstract: Recent advances in artificial intelligence and machine learning (AI/ML) have resulted in autonomous systems that reduce operators' workload and involvement in hazardous missions. Integrating these systems into an existing system of systems (SoS) can evolve it into a system of autonomous systems (SoAS). SoAS brings new Systems Engineering (SE) challenges for architecture development, integration, testing, and evaluation that originate from the level of autonomy (LoA). LoA refers to the level of autonomous capabilities of a system depending on its AI/ML technology. This paper examines SoAS test and evaluation challenges, such as uncertainty and emergent behaviors. We propose a test and evaluation method that provides stakeholders with a decision analysis tool to make data-driven decisions by exploring the SoAS design space, comparing different alternative solutions under uncertainty, and selecting the most suitable one that addresses their mission needs. The proposed method is a step towards AI for Systems Engineering (AI4SE). It applies the Bayesian network (BN) to SoAS by integrating ML algorithms and Model-Based Systems Engineering (MBSE) executable models. The decision analysis tool provides two types of examinations as follows: predictive analysis to select the most suitable SoAS alternative solution that yields the desired performance; and prescriptive analysis to identify root causes of an SoAS undesired emergent behavior and define preventive strategies for future operations. We present a conceptual example of a search-and-rescue mission to

demonstrate the implementation and effectiveness of the proposed method in evaluating SoAS with varying LoAs.

Advancing Human-AI Collaboration in Small and Medium-Sized Enterprises: A Systems Engineering Approach
Authors: Luis Flavio Ortolano and Erika E. Gallegos

Abstract: This paper introduces the Human-AI Collaboration Maturity Model (HAIC-MM), a framework designed to address the unique AI adoption challenges faced by Small and Medium-sized Enterprises (SMEs). The model identifies critical dimensions and capabilities needed to foster effective collaboration between humans and AI systems. The model also defines five maturity levels within each capability, allowing a granular assessment within the holistic framework. HAIC-MM provides a practical, step-by-step guide to assess and enhance AI integration for SMEs. The model emphasizes ethical AI oversight, workforce empowerment, and adaptive organizational culture, while addressing key challenges like resource constraints. HAIC-MM represents a significant contribution to the fields of systems engineering and organizational behavior, offering researchers investigating socio-technical systems, AI integration processes, and SME innovation strategies a rigorous framework for both theoretical advancement and practical implementation. With its focus on real-world application, HAIC-MM equips practitioners with actionable insights to build trust, optimize collaboration between human and AI capabilities, and achieve sustainable, ethically sound AI adoption, ensuring their organizations remain competitive in an increasingly digital economy.

Quantifying Structural Complexity, Effort, and Performance: An Early Experiment Using Network Design Tasks
Authors: Alfonso Lanza, Joshua Sutherland, Marc-Andre Chavy-Macdonald, and Olivier L. de Weck

Abstract: Structural Complexity is perceived as driving cost in system development yet managing it effectively requires empirical understanding. This study investigates human decision-making using a toy transportation-style network design task, focusing on how Structural Complexity, Effort, and Performance interact. Seventy-four participants (primarily systems engineers) worked in small groups to design simple transportation-style networks, drawing 188 designs across 20 node sets (each representing a distinct spatial configuration). This was a trial of a new type of experimental setup aiming to empirically test hypotheses in Systems Engineering, and though it is broadly promising, the design needs significant iteration for better controls. Such iteration could greatly help SE theory.

Four hypotheses grounded in the proposed “Conservation of Complexity law” were explored, with early results showing that additional Effort was generally associated with modest improvements in Performance; Performance rises with Structural Complexity but sees diminishing returns; additional Effort did not reliably reduce Structural Complexity at fixed Performance; and the Effort–Complexity relationship followed a sub-linear trend.

In parallel, we outline a practical, repeatable quantification method for structural complexity

to support a more robust and comparable measurement in human-in-the-loop experiments. Future experimental designs with standardized constraints, behavioral controls, and normalized performance metrics will be necessary to generalize these insights.

May 2026 Full Access Articles
The May 2026 edition of the Journal includes the following full access titles:

A Theory-Driven Interpretation and Elaboration of Verification and Validation
Authors: Hanumanthrao Kannan, and Alejandro Salado

Abstract: This paper presents a formal theory of verification and validation (V&V) within systems engineering, grounded in the axiom that V&V are fundamentally knowledge-building activities. Using dynamic epistemic modal logic, we develop precise definitions of V&V, articulating their roles in confirming and contextualizing knowledge about systems. The theory formalizes the interplay between epistemic states, evidence, and reasoning processes, allowing for the derivation of theorems that clarify the conceptual underpinnings of V&V. By providing a formal foundation, this work addresses ambiguities in traditional V&V practices, offering a structured framework to enhance precision and consistency in systems engineering methodologies. The insights gained have implications for both academic research and practical applications, fostering a deeper understanding of V&V as critical components of engineering knowledge generation.

Application of Systems Engineering to Improve Ecological Knowledge Uptake in Spatial Planning Processes
Authors: Arron Wilde Tippett, Dina Margrethe Aspen, and Liv Guri Velle

Abstract: The Global Biodiversity Framework calls for participatory, biodiversity-inclusive spatial planning to address accelerating losses of biodiversity and ecosystem services driven by land-use change. Yet spatial planning systems continue to enable nature degradation despite the increasing availability of ecological data. Ecological data is produced under a knowledge-deficit model - assuming that more data leads to better decisions. However, in spatial planning practice, data must be activated through participation: the less the participation, the less ecological knowledge enters decision-making. Participation is shaped by three dimensions: demos (who participates), chronos (when participation occurs), and kairos (the knowledge environment that enables effective participation). To examine how ecological data becomes integrated into planning decisions, we apply the Stakeholders, Problem, Alternatives, Decision, and Evaluation (SPADE) systems engineering framework. Using triangulation of interviews and gray literature, analyzed through thematic analysis, we interpret findings through a complexity lens, focusing on multi-stakeholder interactions, path dependency, and unpredictable behaviors. We find multiple orientations for activating ecological data, but the regulated pathway addresses the most system requirements. We identify the need for an interpretation function within planning system architectures to create the appropriate kairos in which planners accept ecological data as decision-relevant knowledge. We identified a number of examples of positive feedback which are limiting the ability to integrate ecological data in planning decisions. Finally, we argue that aligning future

ecological data production with stakeholder objectives may enhance its uptake. This study offers a qualitative systems approach to better understand why biodiversity continues to decline despite enhanced data availability and regulated planning processes.

Artificial Intelligence Enabled Systems Engineering Modeling With Retrieval Augmented Generation
Authors: Tomi Esho, Clarissa Hoyt, Jeremy Marshall, and Jyotirmay Gadewadikar

Abstract: This work enhances systems engineering processes by automating modeling tasks with an AI-driven tool. For researchers, the tool offers the ability to integrate large language models (LLMs) with model-based systems engineering (MBSE) tools through application programming interfaces (APIs). The tool's core capabilities include the automatic creation of SysML components like block definition diagrams and state machine diagrams, as well as reading and analyzing the models. With the addition of retrieval augmented generation (RAG), the program can retrieve context from documents that contain domain-specific information, which is used to improve the language model's response. For practitioners, the tool assists systems engineers with an interactive user interface that offers AI-driven model management. The tool is presented as a chatbot where users can request model updates, ask questions, or automate tasks. These abilities improve speed and accuracy within the systems engineering workflows.

Design of Safety Pattern for eVTOL Aircraft Systems Based on MBSE - A Comprehensive Approach
Authors: Boyang Li, Peng Wang, Jie bai, and Nve Xiao

Abstract: Safety is always of paramount importance in civil aviation. To shift safety considerations from post-validation to prior construction, an architecture design method has been proposed based on the MBSE approach, aiming to integrate safety requirements into the design process. First, related safety requirements are refined, and a safety pattern library with safety logic is formalized and established, which serves as a bridge between abstract safety requirements and specific architectural solutions. Subsequently, in an iterative, analysis-driven design process, when safety defects are exposed in the architecture, corresponding safety patterns could be systematically selected and instantiated from the library to optimize the architecture. Finally, the approach is demonstrated by designing the flight control architecture of an uncrewed eVTOL aircraft. The case study shows that this method not only ensures consistency and traceability between design and safety, but also transforms the role of safety analysis in the later stage from traditional exploratory analysis to quantitative verification of built-in safety attributes. Thus, a new perspective is provided to improve the efficiency of airworthiness certification in the low-altitude air mobility industry.

From Ground to Space: Construction and Application of a Three-Dimensional Digital Collaborative Design Platform for Space Station Engineering
Authors: Hongyi Zhang, Wei Wang, and Yuejie Wen

Abstract: As aerospace engineering progresses toward greater complexity and system integration, traditional design and management methods are increasingly inadequate for meeting modern project demands. Digital twin technology and Model-Based Systems Engineering (MBSE) have emerged as valuable tools for enhancing spacecraft design, manufacturing, and management efficiency. This study developed and implemented a three-dimensional digital collaborative design platform for space station engineering, using case study analysis based on Chinese Space Station practices to explore the applications and implementation strategies of digital twin and MBSE technologies. The research findings demonstrate that the platform facilitates design-manufacturing integration through unified data sources, as exemplified by parallel optimization during core module thermal testing; digital twin technology supports ground-space synchronization and virtual–physical mapping through homologous models and telemetry data; and digital verification technology optimizes testing processes via automated software, enhancing system reliability. Additionally, successful platform implementation required systematic organizational preparation, including cross-departmental coordination mechanisms, stratified personnel training systems, and gradual work-mode transformation from document-driven to model-driven approaches. Technical implementation followed unified planning and phased deployment strategies, while comprehensive risk management systems addressed technical, data security, and operational challenges through pilot research validation and multi-layered protection measures. This study provides both theoretical insights and practical guidance for digital transformation in aerospace engineering, offering empirical evidence and implementation frameworks for integrating emerging digital technologies in complex systems engineering contexts.

Quantifying the Epistemic Value of Verification: An Information-Theoretic Approach
Authors: Jack Fitzpatrick and Hanumanthrao Kannan

Abstract: System verification is a critical process that ensures components and systems meet their requirements at each integration level. Although verification significantly impacts project costs and schedules, its formal theoretical foundations have received relatively limited research attention compared to other systems engineering processes. Previous researchers have explored various metrics to optimize verification strategies, but these metrics often rely on programmatic factors like cost and utility rather than capturing verification's epistemic value - the fundamental ability to reduce uncertainty about requirement satisfaction. This paper introduces a novel theoretical framework that applies concepts from information theory to quantify the epistemic value of verification activities. By treating verification as an uncertainty reduction process, we leverage information-theoretic measures such as entropy and information gain to evaluate how effectively different verification strategies illuminate whether requirements are truly satisfied. This framework provides a mathematical foundation for optimizing verification strategies based on their core function of generating knowledge about system performance, independent of external programmatic factors.

INCOSE members in good standing may access all Systems Engineering Journal content through their INCOSE Connect login (using the Wiley Online Proceedings Library link after login). Non-members may subscribe to the journal, use institutional logins from their

university or place of employment, or purchase access to individual articles at the URLs associated with the article titles above.

System Dynamics Society (SDS) Recommendations

The System Dynamics Society (SDS) makes available or recommends a wide variety of resources concerning applications of system dynamics techniques to diverse problem domains. Recent recommendations include:

- A Hybrid System Dynamics–Agent-Based Modelling Framework for Adaptive Governance of Decentralized Water Circularity (journal paper)
- A system dynamics tool to assess the development potential of HEFA-SAF supply chains at national and regional scales (journal paper)
- Agentic Modeling in Action: System Dynamics with Socrates and Merlin in Stella (recording)
- AI-Powered Causal Loop Diagrams – From Ideation to Insight (recording – members-only)
- Biomimicry in Systems Science: Nature-Based Solutions for Complex Problems (recording – members-only)
- From Dialogue to Feedback: Approaches to Navigating Complex Systems (recording – members-only)
- From System Dynamics Education to Sustainable Systems Development (recording – members-only)
- Future-proofing urban water–energy resilience: An integrated system dynamics, optimization, and game-theoretic framework (journal paper)
- How Silico models work (recording)
- System Dynamics and Complexity Theory Approaches to Navigating Complex Systems (recording – members-only)
- Systems Thinking Podcast
- The Dead End that is General Equilibrium Modeling (recording - Econ SIG Meeting)
- The Net Climate Impact of Artificial Intelligence (AI): Balancing Current Costs with Future Climate Benefits (journal paper).

The System Dynamics Review, published quarterly by Wiley, includes a mix of open access and member-only articles in recent volumes (42-4 and 42-6):

- On Parallel and Serial (Dis-)Aggregation in System Dynamics Models (open access)
- Unraveling Value Creation by Port Terminal Operators Within Maritime Logistics Systems.

Join the SDS to gain full access to members-only content.

Recommended Product Development and Innovation Resources

The Product Development Management Association (PDMA) hosts a Knowledge Hub (kHUB) that offers a wide variety of product development and innovation management resources in the form of blogs, podcasts, videos, conference presentations, feature articles, and whitepapers. Also included are research articles from the Journal of Production Innovation Management (JPIM).

Recent recommendations include:

- AI-Powered Stage-Gate: Supercharging Your Idea-to-Launch Process (article)
- AI-PRISM: A New Lens for Predicting New Product Success (article)
- Adopting Artificial Intelligence for New Product Development: The RAPID Process (article)
- Antecedents and outcomes of open innovation over the past 20 years: A framework and meta-analysis (JPIM article)
- Best Practices in New Product Development and Innovation (JPIM article)
- Can AI Engines make all the decisions? (article)
- Customer Insights in a Pinch: ChatGPT, Social Monitoring and Good Old Brainstorming (conference presentation)
- Defining a Winning Strategy (conference video)
- Developing a Product Innovation Strategy (article)
- From Applied AI to Practical Business Use: Why Commercialization Infrastructure Matters (article)
- From Idea to Scale: The Four Leadership Modes That Drive Successful Product Innovation (article)
- How do you develop an Innovative Culture? (article)
- Innovation Lab Snapshot: Essential AI Tools for Product Management Success (webcast)
- Innovation Exchange (conference video)
- Innovation: Theory to Impact (chapter video)
- Inside the Innovation Engine: How Michelin and BMW are Cultivating Intrapreneurs through Bold Partnerships (conference video)
- Managed Instability: Product Innovation in a “Never Normal” World (article)
- Navigating Strategic Uncertainty with Scenario Planning: A Guide for Product Managers (article)
- NPD Project Win/Loss Interviews and Analysis (chapter video)
- PDMA Carolinas: Annual Student Innovation Showcase (chapter video)
- Risk-Based Product Design: Avoiding the Cost of Hidden Assumptions (webcast)
- Stage-Gate Agentic: The Coming Revolution in the New Product Process (article)

SYSTEMS ENGINEERING RESOURCES

- TED Talks: Explore the Space of Possible…It is not predefined (blog + video)
- TED Talks: To Design Unique Concepts–Change Your Mindset! (blog + video)
- The AI Threshold Project (research survey)
- The Best vs. the Rest: How do your NPD practices measure-up? (webcast)
- The Hidden Cost of "Alignment": When Product Teams Over-coordinate and Under-deliver (article)
- The Role of Artificial Intelligence in the Ideation Process (JPIM article)
- Using AI to Accelerate Innovation (webcast)
- Working with Claude Code (chapter video)
- Your AI Teammate Doesn't Have a Job Description Yet (chapter video).

Access to kHUB is free and open to the public. Full text access to JPIM research articles requires a PDMA membership or institutional access to the JPIM through the Wiley Online Library; however, kHUB publishes JPIM article abstracts and key takeaways.

Create a guest account or join PDMA here.

Open Access Simulation Resources from NAFEMS

NAFEMS offers both free online and members-only resources in support of its mission to be recognized as an "*independent authority and trusted source for communicating engineering simulation knowledge, and for sharing best engineering modeling, analysis, and overall simulation practices in developing reliable products and innovative solutions*."

Nearly 400 open access resources are available after completing the registration for a free NAFEMS account.

Systems engineering practitioners may search the NAFEMS Resource Center and find the following recent publications relevant to their work:

- Bridging the Gap: A Review of Engineering Simulation Interoperability (paper)
- Engineering Simulation Maturity Assessment: An Introduction (paper)
- Leveraging Scientific Machine Learning to Augment Engineering Modelling and Simulation (webinar)
- What Simulation Engineers Need to Know About MBSE – Bridging Two World (webinar).

Join NAFEMS to access thousands of additional publications and videos.

NAFEMS also publishes a free online Glossary of valuable terms and definitions related to modeling and simulation, accessible without registration or membership.

Business Analysis Resources

The International Institute of Business Analysis (IIBA) is a non-profit professional association that helps business analysts develop their skills and further their careers by providing access to relevant content. IIBA publishes a variety of open access and members-only resources that are published in numerous media formats.

Analyst Catalyst Blog

This open-access blog features posts that address different aspects of the business analysis discipline. Recent topics include:

- 11 New KnowledgeHub Resources for the Moments That Matter
- AI and Analytics Fusion: A Practical Guide for Business Analysts
- Responsible AI: The Business Analyst's Role in Ethics and Governance
- The Symbiosis of AI and Business Analysis: Making Decisions Count
- The What-If Machine: Using Causal Scenario Analysis to Guide AI Decisions with Confidence
- Thinking Is Still the Job: How Business Analysis Is Changing Right Now
- When Business Analysts Compare Notes Across 5 Countries
- Who Owns the Decision? Business Analyst vs Product Owner vs Proxy PO.

Business Analysis Live! Podcasts

IIBA's podcast library is also open to non-members, with over 100 episodes "aired" to date. Recent titles include:

- Instincts are Data Too
- The 3 A's of Business Analysis
- Surviving Layoffs Without Losing Confidence

SEI Blog: Managing Architectural Risk During Agile Development

The Software Engineering Institute (SEI) at Carnegie Mellon University (CMU) publishes informative blog posts on numerous topics relevant to systems and software engineering. A recent post titled "Managing Architectural Risk During Agile Development" proposed an approach that draws from Agile Architecture Risk Management (AARM) and Continuous Risk Management (CRM) processes to create a practice for evaluating software architecture risks during development. The approach weighs the trade-offs between design pattern attributes and quality attributes to identify architectural risks early and assess the system impacts of design decisions made during software development.

The article addresses the following topics:

- Meeting the Needs of the Warfighter
- Continuous Risk Management
- Design Change During Incremental Development
- Comparing Architectural Patterns
- When to Identify Architectural Risks
- The Importance of Identifying Architecture Risk Issues Early.

Learn more about the Software Engineering Research and Development topics within the SEI blog through additional articles:

- Applying the SEI SBOM Framework
- Challenges to Assuring Large-Scale Systems
- Delivering Resilient Software Capability to the Warfighter at the Speed of Relevance.

Systems engineering practitioners may find the following sample of the SEI blog's 24 topic areas of interest:

- Agile
- AI-Augmented Software Engineering
- Artificial Intelligence Engineering
- Cyber-Physical Systems
- Enterprise Risk and Resilience Management
- Model-Based Systems Engineering
- Software Architecture,

Artificial Intelligence- Business Analysis (AIxBA) Forum Summary

In May, Blackmetric Business Solutions, a UK-based training and consulting firm, hosted an Artificial Intelligence- Business Analysis (AIxBA) Forum to provide business analysis practitioners with the opportunity to share their AI-related achievements and challenges. A 5-page Forum Summary Paper is available online and includes topics such as:

- Current AI "hot topics" - AI tool selection, implementation pathways, proof-of-value, and limitations
- Results of polling on the extent of AI involvement by business analysts and formal AI governance policies
- Testing of AI-based systems
- Case study and demonstration of an AI business analysis agent
- "AI Insights" sessions conducted in the UK that address prompt engineering and AI-supported requirements elicitation.

Visure Solutions Webinars and Blog

Visure Solutions, provider of a popular Requirements Application Lifecycle Management (ALM) platform, has made available the recordings of two recent webinars that highlight the increasing impact of AI on systems engineering practices and tools.

Closing the ALM–PLM Gap in Cyber-Physical Systems with AI and Systems Thinking. Topics addressed in this webinar include:

- Rise of Cyber-Physical Products: How modern devices are evolving into integrated hardware-software systems.
- The Role of Requirements Engineering: Why strong requirements are critical, and how AI enhances them.
- Hardware vs. Software Requirements: Key differences that impact development and validation.
- Challenges from ALM–PLM Gaps: Where disconnects create inefficiencies and risks.
- Bridging the Gap with AI & Traceability: Practical strategies to connect ALM and PLM for better outcomes.

Additional details are available in a companion 32-page whitepaper, "Integrating ALM-PLM Handbook" with expanded coverage of these topics:

- Understand the ALM-PLM Gap - Learn how disconnected systems impact traceability, compliance, and engineering collaboration.
- Explore Integration Strategies - Discover ALM-led, PLM-led, MBSE-federated, and hybrid integration models for regulated industries.
- Improve Traceability and Compliance - Enable end-to-end lifecycle visibility, change impact analysis, and audit readiness across hardware and software domains.
- Explore AI and Digital Engineering Trends - Understand AI-assisted traceability, MCP architectures, digital thread strategies, and SBOM integration.
- Implement ALM-PLM Best Practices - Learn proven approaches for lifecycle governance, configuration management, and scalable engineering integration.

Agentic AI in Action: Next-Generation Requirements Engineering for the Future of Systems Engineering

This presentation at ReConf2026 in Munich, Germany, focused on the implications of Agentic AI, addressing topics such as:

- How Agentic AI acts as a co-engineer across the Requirements Engineering lifecycle
- Enhancing Requirements Management (definition, decomposition, and reuse with intelligent agents)
- Managing product variants, parameters, and variability more effectively
- Real-world insights on balancing AI automation with human oversight
- Best practices to ensure explainability, data integrity, and compliance.

The Visure Solutions blog addresses a variety of requirements engineering topics, including:

- AI in Product Lifecycle Management
- ALM and PLM Integration: Examples, Use Cases
- Digital Thread in Medical Devices & Pharma
- Integrating QMS, Risk & Requirements in One Platform
- Leveraging AI for Predictive Risk Management
- What Is Compliance Automation?
- What is Impact Analysis?

Real MBSE Webinar Series

SPEC Innovations, a leading provider of systems engineering services and the Innoslate and Sopatra software solutions, has delivered and made available a 7-part webinar series titled "Real MBSE". These webinars highlight foundational concepts and key points of the same-named book by Dr. Steven Dam, SPEC Innovations President and Founder, and INCOSE-certified Expert Systems Engineering Professional (ESEP). Session recordings and recaps are available here, on the SPEC Innovations Resources page and on the YouTube links shown below.

- The Role of Systems Engineering Across the Lifecycle
- Turn Word and Excel Into Systems Intelligence
- Functional Analysis and Simulation Without Pain
- Synthesizing Solutions for Mission Success
- Systems Analysis for Smart Decisions
- V&V, Integration, and the Endgame
- Operations, DevOps, and Disposal.

Request a free download of an electronic copy of *Real MBSE*.

Exploring Transdisciplinarity in Systems Engineering

MDPI, publisher of ~100 highly ranked open access journals, fosters open scientific exchange in all forms, across all disciplines. MDPI's systems journal, Volume 14, Issue 5 (May 2026) features an article that highlights the challenge associated with this mission.

Exploring Transdisciplinarity in Systems Engineering to Develop Transdisciplinary Integration Capability and Process Meta-Level Conceptual Models
Authors: Ali Asghar Bataleblu, John Arthur Poirier and Julia Taylor

ABSTRACT: Key challenges for the future of Systems Engineering (SE) include practically converting concepts into actions that solve complex problems and implementing formalized

Transdisciplinary Integration (TDI) approaches to achieve this. This paper starts with a literature review, which exposes specific gaps: Transdisciplinarity in SE as a mindset is not well established and Transdisciplinary Integration Processes (TDIP) are not formalized and rigorously applied. Challenges are collected with a focus on "Transdisciplinarity in SE" as a meta-discipline in supporting TDI meta-model development. To this end, a TDI conceptual capability model comprising ends, ways, and means that coalesce transdisciplinary knowledge, methods, and stakeholder perspectives, and a TDIP model applicable across wide-ranging problem spaces and dimensions are proposed. A Disciplinarity Index (DI) is introduced to analyze a spectrum from Intra- to Alter-Disciplinarity; determine proportionality of problem dimensions, factors, and bounding aspects; and define a solution approach. The models are validated by applying the DI to four socio-technical problem domains ranging in scope and scale. This study concludes that the diverse challenges are approachable using the proposed conceptual foundation and models, which are unique in establishing a mechanism to blend disciplines, meet the challenges, and achieve Transdisciplinarity in the SE community while forming relations in the broader systems and transdisciplinary communities.

Systems engineering practitioners may find the other 139 articles in this edition of value. Here are some examples to consider:

- An AI-Enhanced Technical Debt Management Framework for Aerospace and Defense Systems Engineering: Framework Design and Illustrative Application
- Automated Inference of Systems Capability from Natural Language Artifacts
- Engineering Systems with Standards and Digital Models: Specifying Stakeholder Needs and Capabilities-MGOS
- Evaluating MBSE Approaches and Tools for Aircraft Design and Certification: A Comparative Perspective
- Risk Management and Urban Resilience: Literature and Practice
- The Dynamics of Digital Transformation: Why Similar Efforts Diverge.

Access other issues of *systems* from 2026:

- Volume 14, Issue 1 (January 2026) – 117 articles
- Volume 14, Issue 2 (February 2026) – 109 articles
- Volume 14, Issue 3 (March 2026) – 110 articles
- Volume 14, Issue 4 (April 2026) – 117 articles
- Volume 14, Issue 6 (June 2026) – 87 articles.

MBSE and Cybersecurity Blog Series

Obeo offers a variety of MBSE resources, including training and software extensions for Capella, and significant contributions to the Eclipse SysON SysML v2 project. The company has recently released a 4-article series on MBSE and Cybersecurity on the Obeo blog. Topics addressed include:

- Why MBSE is Becoming Essential for Modern Cybersecurity
- MBSE for NIS2 Compliance in Cyber-Physical Systems
- MBSE for Cyber Resilience Act (CRA) Compliance
- Strengthening Cybersecurity with Capella and DARC.

The INCOSE Podcast

The INCOSE Podcast series continues with Episodes 3 and 4 now available. The podcasts are open-access events; no INCOSE membership is required to access this content.

Episode 3 features David Schrunk, President of the Science of Laws Institute, exploring why the traditional lawmaking method used in many nations fails to define problems, measure outcomes, or ensure effectiveness. Fixing Laws: A Systems Approach asks, “What if we treated legislation like any other complex system?”

From Japan’s "Agile Governance" framework to new policy-engineering initiatives in Mongolia, countries around the world are beginning to explore how we can apply systems engineering rigor to the "system of systems" that is modern government.

Learn more about this topic at http://www.scienceoflaws.org.

In Episode 4, titled The Power of Engineering Quality, Data, and Facing Reality, Julia Mullaney, co-founder of Exceptional Difference and formerly of IBM and the Software Engineering Institute (SEI), brings a data-backed perspective on bridging the gap between engineering theory and daily project realities. Topics include:

- The "Invisible" Savings of Quality: Why investing in early design and code reviews actually saves hundreds of hours - flipping the script on the misconception that quality is expensive.
- Managing Schedule Pressure: How technical leaders can find the courage to face reality, establish realistic commitments, and protect engineering standards.
- Coaching vs. Rescuing: Why the best mentors focus on building long-term capability within their teams rather than stepping in to "fix" every problem.
- The AI Boom: Her take on the rapid acceleration of AI, and why architectural wisdom and design thinking matter now more than ever.
- Advice to the next generation of engineers.

Requirements Engineering Magazine: Recent Articles

The Requirements Engineering (RE) Magazine is published by the International Requirements Engineering Board (IREB) multiple times per year. Articles

are welcomed from Business Analysis and Requirements Engineering professionals, regardless of IREB membership status. Publication is free of charge for the authors.
In the June 2026 edition, Praveen Chinnappa, a product owner and IT architect, authored an article titled "Strengthening the Requirements Engineering Process - Integrating a Testing Mindset for Requirements Engineers".

Abstract: The quality of a requirement is significantly dependent on its clarity and completeness. In many agile projects, requirements typically capture the primary intent of the software. The variations, exceptions, and intended exception behaviors are left implicit and often uncovered in development or testing phases.

Testing plays a crucial role in determining software readiness. This is linked to the quality of requirements since test cases are designed to exercise the requirements. While test design is the primary responsibility of testers, explicitly calling out variations and exceptions reduces interpretation and increases testability.

Product Owners (PO) and Business Analysts (BA) can support this activity by adopting a tester's mindset during the Requirements Engineering (RE) process to deliver requirements that can withstand systematic testing.

This article presents practical techniques and real-world examples for implementing a tester's mindset for the BAs/POs during requirements engineering.

Topics addressed include:

- Requirements and Testing - Two pillars of SDLC
- Documentation in Agile – The misunderstood principle
- Tester Mindset – A new dimension in Requirements Engineering
- Early Static Testing
- Leveraging Test Design Techniques
- Final Verification Checklist
- Password reset case study
- Advantages of utilizing a tester's mindset during the requirements engineering process.

The July edition featured an article, Requirements for cross-cutting qualities: Integrating explainability and privacy as a first step towards a stakeholder needs taxonomy. The authors, Eddy Groen, Hannah Deters, Jakob Droste and Hartmut Schmitt, offer a definitive **no** in answer to the question:

Can we really treat qualities such as privacy and explainability in isolation?

By integrating stakeholder needs across multiple quality attributes, they demonstrate that cross-cutting qualities are deeply interconnected and that improving one can unintentionally affect others.

Topics addressed include:

- What is explainability?
- What is usable security and privacy (USP)?
- Why are the "requirements" for these cross-cutting themes different?
- Why is there an overlap between these qualities?
- A combined taxonomy of explainability and USP needs
- Implications and outlook.

View the other 100+ RE Magazine articles.

Sign up to be notified about new issues of the Requirements Engineering Magazine.

See guidelines for contributing to the RE Magazine as an author.

CIMdata PLM Market Update 2026

CIMdata, a leading Product Lifecycle Management (PLM) research, education and consulting organization, is conducting a Leadership webinar series focused on the Digital Transformation space in general and the PLM market in particular.

The June 2026 webinar, PLM Market Update, is available for viewing as a recording and as presentation slides.

Topics covered include:

- The AI Shift: How AI enablement is reshaping the competitive landscape and driving market value.
- The 2030 Forecast: Exclusive early access to CIMdata's five-year global market projections (2026-2030).
- M&A Impact: Understand how 2025's mergers and acquisitions changed the playing field.
- Macro Trends: The specific global economic conditions affecting your PLM investments right now.

View prior CIMdata PLM webinars.

INCOSE INSIGHT Practitioners Magazine – June Edition

The June 2026 Edition of INSIGHT, INCOSE's Practitioners Magazine, has been published by Wiley. Electronic subscriptions to INSIGHT are available as a member benefit to INCOSE members. Hard-copy subscriptions to INSIGHT are available for purchase by INCOSE members for one membership year, and to the public.

The 14 members-only articles in the 97-page Volume 29, Issue 3 highlight the theme of *A Decade of AI and Systems Engineering: Looking Back and Forward*. The articles are arranged in four topic areas:

- Decadal Insights
- AI4SE
- SE4AI
- Viewpoint.

Decadal insights are summarized in three articles:

A Review of the SERC AI4SE & SE4AI Workshop Series, and a Lookahead
Author: Megan Clifford

This technical review synthesizes the evolution of the AI4SE & SE4AI workshop series from 2020 through 2025, with emphasis on the event's technical maturation, recurring themes, emerging patterns, and future relevance. The review considers the workshop as both an event series and a signal of disciplinary development: it reflects how the AI and systems engineering communities are defining the engineering practices required to transform AI from an emerging capability into an assured, operationally credible component of complex systems.

Challenges and Opportunities for Systems Engineering in the Age of AI: SERC Perspective
Authors: Zoe Szajnfarber, Tom McDermott, and Valerie Sitterle

This perspective summarizes recent work across the SERC network in terms of five major buckets: 1) AI4SE - AI tools developed to support systems engineering work; 2) SE4AI - leveraging systems principles to assure AI-enabled systems; 3) feedback to AI4SE - where system architecting principles are leveraged to re-map what systems engineers do as AISE progresses; 4) feedback to SE4AI - where novel AI tools are developed to support challenges arise when assuring AIES; and 5) supporting the broader workforce transformation that will be needed for systems engineering to evolve with AI. In each section we highlight ongoing work and point out emerging trends. We conclude with a roadmap, suggesting key needs and trajectories for the next decade.

From Rules to Agentic Swarms: A Systems Engineering Journey Through the Evolution of AI
Authors: Paul Wach PhD, Alejandro Salado PhD, and Brad Philipbar

Artificial intelligence (AI) for systems engineering has evolved through six paradigms in five decades: expert systems, machine learning, deep learning, large language models, agentic AI, and agentic swarms, with each paradigm shorter than the last. This position article traces the arc through four systems the authors constructed across the paradigms, Houston (rule-based requirements), TurboArch (large-language-model architecture decisions), GenGroves (agentic model-based systems engineering), and MACQ (multi-agent acquisition support),

supplemented by published work from researchers across the Systems Engineering Research Center (SERC) community for the machine-learning and deep-learning paradigms. A central observation runs through this history: each new paradigm has become a durable substrate for the next. Ontology and machine learning persist into later paradigms, and the article's portfolio of sixteen AI-for-SE systems shows how agentic swarms now compose those layers across the engineering lifecycle. The article does not attempt a comprehensive survey of AI; rather, the article identifies a critical inflection point through the lens of the authors and the SERC community. A critical observation from this article is that when organizational adoption lag exceeds paradigm duration, organizations risk investing in yesterday's frontier while missing the durable investment in the layers themselves. The systems engineering community must compress its own adoption cycles or accept permanent strategic lag.

AI4SE is addressed in five articles:

Automated Curation and Execution of Engineering Models
Authors: Jaya Kambhampaty, Olivia Fischer, and Dimitri Mavris

The 2020s have been marked by increased attention to data-intensive information systems that maintain consistency with the evolving state of physical assets. We call the digital models of the connected physical systems “digital twins” (AIAA Digital Engineering Integration Committee 2021), while “digital thread” refers to the foundational framework of protocols, standards, and linked representations that maintains contextual meaning and traceability as data crosses tool, discipline, organizational, and lifecycle boundaries (AIAA Digital Engineering Integration Committee 2023). The practical and theoretical methods for producing these systems make up digital engineering (Griffin 2018), a discipline that formalizes and institutionalizes modeling as a primary, lifecycle-spanning engineering approach and is characterized by improved data sharing, model interoperability, and effective reuse.

On (Human) Learning in Design Space Exploration and the Impact of AI Assistants
Author: Daniel Selva

Design space exploration (DSE) is a foundational activity in early-phase systems engineering aimed at identifying promising designs through modeling, optimization, and analysis. With the rapid emergence of AI-powered assistants, engineers can now delegate increasingly complex DSE tasks, from generating candidate designs to analyzing entire tradespaces. While this delegation may improve efficiency and solution quality, it raises an important question: what is lost when engineers offload cognitive effort to AI in DSE? This article argues that DSE is not only an optimization process but also an information discovery and learning activity through which engineers develop intuition, uncover trade-offs, and build defensible rationales for design decisions. Drawing on some of our recent empirical studies, we examine how AI assistance affects both DSE performance and learning, highlighting a nuanced relationship between the two. The findings suggest that while AI can enhance DSE outcomes, excessive automation may hinder learning, particularly for less experienced engineers. We conclude with implications for designing AI assistants that balance performance gains with the need to preserve and enhance human understanding in engineering design.

STPD – Operational Design as a Model-Based Extension of STPA
Author: Avi Harel

The paper argues that the growing complexity of distributed socio-technical systems exposes a fundamental limitation of the STPA methodology. Although STPA transformed safety engineering (Leveson 2011; Leveson 2004) by treating safety as a control and constraint problem rather than purely a reliability problem, it still assumes that hazardous scenarios and the constraints needed to prevent them can largely be identified during design time. The paper claims that this assumption no longer holds in modern adaptive systems (Rasmussen 1997; Woods 2015), where operational contexts evolve continuously, AI-based components change behavior dynamically, distributed actors develop inconsistent interpretations of the operational situation, and unexpected interactions emerge during runtime.

According to the paper, safety in such environments cannot rely only on predefined constraints. Instead, systems must continuously validate during runtime whether operational assumptions, coordination structures, and safety constraints still remain valid under changing conditions. This shifts the engineering challenge from specifying static safe behavior toward maintaining operational coherence under uncertainty. To address this challenge, the paper proposes extending STPA into a broader framework called system-theoretic process design, (STPD).

Trustworthy LLM-Augmented Data Reconciliation for Deployment-Constrained Environments
Authors: Danny J. O'Leary and Allison B. Ledford

Digital engineering initiatives depend on trustworthy data lineage and a shared understanding of what enterprise data means. In practice, those conditions are often absent. The same concept may be represented in databases, forms, spreadsheets, email, and conversation-based exchanges, with different labels, different levels of structure, and different degrees of documentation. Organizations attempting to build a digital thread or expand model-based systems engineering (MBSE) often discover that the limiting factor is not tool availability. It is uncertainty about how data actually moves through work and how apparently related elements should be interpreted across sources (Ledford 2023; Ledford, Harris, and Purdy 2023; Ledford et al. 2025).

That problem is the foundation for data element mapping and analysis (DEMA). DEMA was developed to make hidden data work visible by tracing information at the levels of function, vessel, and individual data element. Prior DEMA applications have shown that a large share of enterprise data handling remains undocumented and that many exchanges still depend on manual transfer rather than integrated digital flow (Ledford 2023; Ledford, Harris, and Purdy 2023). Those findings are important because a digital thread is only as reliable as the organization's understanding of the data relationships it claims to connect.

Twenty Tips for Using Generative AI in Systems Engineering
Authors: Ray Madachy, Dan O'Leary, Tom McDermott, and Sinan Bank

Generative AI (GenAI) is rapidly moving from experimentation to routine use in systems engineering practice. We are already using these tools across virtually all process activities. The potential benefits are substantial, but the risks are equally real, including inconsistent outputs, hidden assumptions, unverifiable calculations, weak traceability, context loss across sessions, and avoidable rework caused by poorly framed prompts and informal review habits. Current use is often ad hoc, whereas we need to move to disciplined individual and team processes that are measurable and thus manageable.

Five articles provide perspectives on SE4AI:

A Systems Architecture for Data-Constrained Environments: Leveraging the Astronaut Digital Twin to Improve Spaceflight Adaptation in Humans
Authors: Caleb Schmidt and Steve Simske

Spaceflight exposes humans to a uniquely harsh operating environment, producing highly individualized physiological responses that existing monitoring and countermeasure paradigms have limited struggle to predict or control. This challenge is compounded by small sample sizes, limited observability, fragmented evidence, and shifting baselines, creating a discordant epistemic environment where conventional data-driven, mechanistic, and hybrid models' methods are not well-suited. Resolving this gap is becoming increasingly urgent: the expanding scope, duration, and population diversity of human spaceflight missions have outpaced the predictive and personalization capabilities of space medicine now. This article proposes the astronaut digital twin (ADT), a biodigital twin architecture designed to help address three coupled problem classes: inference of physiological state under partial observability, reconciliation of heterogeneous evidence within a causal framework, and individualized countermeasure control under uncertainty. The ADT integrates systems physiology models, digital populations, and individualized digital twins within a Bayesian, constraint-based architecture that continuously updates state estimates and generates uncertainty-bounded, mechanistically interpretable recommendations. By formalizing expert knowledge into a testable human–machine co-learning system, the ADT is intended to enable a transition in space medicine from reactive monitoring toward predictive, personalized decision support. It also serves as a structured hypothesis engine for countermeasure development across diverse physiological profiles. More broadly, the ADT is positioned to engage with open SE4AI questions centered around interpretable reasoning, predictable and diagnosable behavior, and model validity under changing deployment conditions. Although developed for astronaut health, the architecture has the potential to generalize to other high-consequence domains characterized by sparse data, population diversity, and non-stationary baselines, suggesting its potential as a reference framework for trustworthy AI systems in complex operational environments.

Designing Human-AI Architectural Robustness to Inherently Brittle AI
Authors: Aditya Singh and Zoe Szajnfarber

Machine learning-based systems have made impressive advances in recent years, but they are still brittle, often failing when encountering new and unexpected changes to the

operating environment. Humans are often partnered with AI-enabled systems to negate this brittleness, but human operators may struggle with changes to the scale of the operating environment, as cognitive fatigue can decrease the effectiveness of human oversight. Through a simulation of a notional minefield traversal system, we show how humans and AI are integrated together affects system outcomes. We propose that evaluations of human-AI systems should consider different human-AI integration architectures and that AI-enabled systems should also deploy continuous monitoring to detect shifts in the operating environment which may cause AI system performance to falter. Finally, system design should enable human operators to have sufficient situational awareness, which may require scaling up the number of human operators available.

Engineering Trust in AI-Enabled Systems: Lessons from the Trusted AI Challenge
Authors: Sami Saliba, Emma Meno, Scott Lucero, Tyler Cody, Nathan Lau, Carlo Lipizzi, Hunter Moore, Jitesh Panchal, Sachin Shetty, Zoe Szajnfarber, and Peter Beling

This article contributes a systems engineering perspective on trust in AI-enabled systems by synthesizing results from the Trusted AI Challenge. The Systems Engineering Research Center (SERC), with the U.S. Army DEVCOM Armaments Center at Picatinny Arsenal, New Jersey, organized and conducted this student challenge across three academic semesters with six student/faculty teams. The challenge highlighted a broader evolution in systems engineering - from managing deterministic systems to architecting adaptive, human-centered systems in which trust must be continuously engineered and sustained. The challenge demonstrated engineering trust in AI is not solely a computer science problem - it is fundamentally a systems engineering challenge. By requiring teams to balance safety, explainability, and performance in partially observable, adversarial environments, the challenge catalyzed the development of actionable systems engineering methods.

From Models to Systems: The Evolution of T&E for Artificial Intelligence from 2020 to 2030
Authors: Tyler Cody and Peter Beling

Test and evaluation (T&E) for artificial intelligence is entering a new phase. In 2020, the central challenge was to treat machine learning models as system functions rather than as ordinary software components. By 2025, that view had expanded: AI-enabled systems had to be evaluated in operational, organizational, and lifecycle context, with attention to trust, adaptation, resilience, and human-AI teaming. Looking toward 2030, an even larger shift is emerging. AI is no longer only part of the delivered capability; it is increasingly part of the engineering machinery that produces requirements, code, tests, models, documentation, and assurance evidence. This article traces that progression from model-centric assurance to lifecycle-centered assurance and, ultimately, to T&E for AI-generated systems. It argues that future assurance will depend not only on testing delivered systems, but also on evaluating the provenance, governance, and stability of the generative pipelines that create and modify them.

Trustworthy AI for Quantitative Systems Engineering via Deterministic Delegation to Verified Tools

Author: Raymond Madachy

Generative AI offers substantial capability in systems engineering processes for multimodal interaction supporting elicitation, interpretation, and synthesis. However, large language models (LLMs) remain unreliable for exact quantitative computation, as they may produce incorrect or hallucinated results when executing mathematical and domain-specific quantitative calculations. Systems engineering is pervasively grounded in quantitative models that demand determinism, correctness, and traceable results, requirements that probabilistic language models cannot reliably satisfy on their own. A deterministic delegation pattern is presented that addresses this limitation by constraining the LLM to input interpretation, assumption documentation, and results presentation while assigning all authoritative computation to verified external tools through formally defined API endpoints. Identical inputs produce identical results, every output carries full provenance, and users can independently inspect or rerun any analysis via a returned source URL. The pattern is demonstrated through application domains for systems engineering cost modeling using the Constructive Systems Engineering Cost Model (COSYSMO) and system reliability analysis using fault tree and reliability block diagram methods. Examples at increasing levels of complexity show how the assistants translate diverse user inputs into complete, reviewable input configurations before invoking the deterministic backend, combining the accessibility of conversational AI with the correctness and auditability that engineering decisions require.

This edition is capped by a thoughtful consideration of the impact of AI on organizational decision-making:

Who Gets to Decide? An Organizational Decision-Making Lens for AI Governance
Author: Jimmie G. McEver

Systems engineers integrating AI into operational settings need better ways to think about delegation, oversight, and accountability. Command-and-control theory offers a practical framework - one that changes what you actually design, not just how you document it.

View INSIGHT Volume 29, Issue 3 in the Wiley Online Library.

Beyond PLM Blog: PLM's Missing Decision Layer

Oleg Shilovitsky, founder of OpenBOM with more than a quarter century of experience in the manufacturing, engineering, and Product Lifecycle Management (PLM) software industry, has hosted the Beyond PLM blog since 2008. Shilovitsky's original idea was to share daily thoughts on innovations related to Computer-Aided Design (CAD), Product Data Management (PDM), and PLM. Over time, the blog evolved to cover a wide range of topics, including CAD, PDM, PLM, collaboration, innovation, data, and data management. Recently, the blog launched a series of articles arguing the scope, nature and importance of a product memory decision layer that is missing from nearly all PLM software platforms. This series

resonates with the PPI SyEN team who see a similar decision memory gap in MBSE tools and methods.

Here are relevant articles that build the case for maintaining a product memory layer that provides the rich decision context for product data:

- Building Physical Products Like Software: Five Transformations for Manufacturing Teams
- Why Manufacturing Needs Product Memory: Fixing Product Data Chaos in Manufacturing
- The Product Memory Flywheel: Rethinking PLM, Product Data Management, and the Digital Thread
- Is Product Memory Just a New Name for PLM?
- Why "Product Memory" Triggered Such a Strong Reaction (and What It Reveals About PLM's Missing Layer)
- Is Product Memory a New System of Record?
- What is Product Memory?
- Before Software: How Engineering Was Always a Human Memory System.

Shilovitsky has a book in the works, titled From CAD Files to Product Memory. Subscribe here to gain access to early chapters.

Agentic AI Architecture and Standards

Systems engineering practitioners should be aware of emerging trends in the Artificial Intelligence (AI) space because of the potential for advances that can improve their efficiency and effectiveness when performing systems engineering tasks (AI4SE) and also because of the implications of embedding AI within the solutions that they develop (SE4AI).

Agentic AI is such a trend to monitor.

IBM defines Agentic AI as "*an artificial intelligence system that can accomplish a specific goal with limited supervision*."

Gartner characterizes Agentic AI as systems capable of:

- Multistep task execution
- Complex workflow orchestration
- Learning from outcomes
- Moving beyond traditional automation

At a minimum, monitoring recent announcements and webinars from leading MBSE tool vendors, Agentic AI is certainly "coming soon to a tool near you."

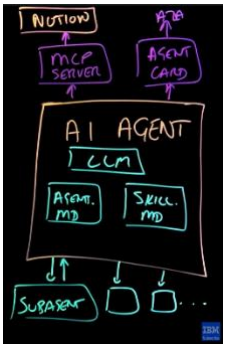

For a quick overview of Agentic AI terms, architecture, and relevant standards, the 11-minute IBM video, "5 AI Agent Terms You Need to Know", is a good place to start. The video introduces the five elements of the most common implementation architecture for Agentic AI (as summarized by Google Gemini):

- agents.md: A project-root markdown file that acts as a "README for agents" by outlining specific commands, testing tools, and codebase coding conventions.
- Agent Skills (skill.md): A modular folder structure containing metadata triggers and scripts that allows an agent to progressively load task-specific capabilities only when contextually relevant.
- Model Context Protocol (MCP): An open standard protocol that utilizes an MCP Server to translate and abstract diverse external data sources and tools into a single uniform interface for compliant agents.
- Agent-to-Agent (A2A): An open protocol that enables cross-agent collaboration and task delegation by allowing agents to discover and read each other's standardized capability profiles, known as Agent Cards.
- Sub-agents: An architectural multi-agent design pattern where a primary parent agent dynamically spawns context-isolated child agents to execute massive or parallel workloads before summarizing the results back.

Business Analysis (BA) Digest Q3 2026 Edition

Blackmetric Business Solutions, a UK-based training and consulting firm, publishes a quarterly Business Analysis (BA) Digest. In keeping with its theme as *your regular round-up of useful BA content*, the 70-page Q3 2026 edition addresses topics such as:

- A Kitchen Sink Drama: Why Using Analogy Makes You a Better BA
- Before You Improve, Understand: Why Current-State Clarity Matters
- Beyond Sign-off: How Business Analysts Can Enable Meaningful UAT
- Blueshift: How Business Analysis Can Help Ensure That Transformation Is Delivered Securely
- Combining OKRs and KPIs?
- Feedback, the Gift That Keeps Giving
- From Requirements to Production: Redefining the Business Analyst in the Age of AI
- From the Sidelines to the Centre: How Business Analysts Turn Absentee Management into Leadership Opportunity
- Landing on a Moving Train: The BA's First 100 Hours
- NIS2: About Security, Continuity and Business Analysis!
- Prompting like an Analyst: Why AI Is a New Core BA Skill
- The AI Model Doesn't Know Your Domain. You Do
- The Egg, the Potato, and the Teabag

- The POETICS of Transformation: A Holistic and Open Source Lens for Change Management
- Understanding the Impact of Pricing Changes: A Business Analysis Perspective
- Using Group Construct Analysis to Understand Stakeholder Worldviews in Ambiguous Situations, Part 21: Taking Things Personally
- When AI Handles the Frameworks, Soft Skills Become the Real Differentiator.

Typical articles are a quick read at 2-4 pages in length.

Subscribe to the BA Digest here.

Smart City Council: Urban Insights Newsletter

The Smart Cities Council (SCC) has launched a new monthly publication, Urban Insights, that provides a forum for experts to communicate their perspectives on the latest Smart City strategy, use cases, technologies, policies, standards, regulations, and legislation.

Articles in the inaugural edition include:

- AI-Powered Smart Parking Redefines Urban Infrastructure Efficiency
- Building Smarter Water Intelligence for Communities
- From Landfills to Landscapes: Leveraging Smart Tech for Localized Waste Recovery
- Smart Campus in Action: Digital Identity at Scale at University of Sydney
- Smart Mobility for Peace: Building Cities Through Connection, Not Division
- The Rise of Operational Intelligence: The New Oil of the Americas
- Why Smart City AI Projects Fail: The Infrastructure Imperative.

Learn more about the Smart Cities Council here.

"Designing a solution doesn't change the problem"

Robert H. Halligan

FINAL THOUGHTS FROM SYENNA

From Separation of Concerns to Systems Dynamics (of Human-AI Partnership)

From Separation of Concerns to System Dynamics (of Human-AI Partnerships)

Much of Syenna's early success as a professional came from learning the systematic thinking technique of Separation of Concerns. This skill goes by many names: decomposition, reductionism, or even just planning. Syenna was generally quicker than peers at breaking down a problem into attackable chunks (questions or decisions), a large project into smaller steps, an organization into roles, a process into activities, functional requirements into a hierarchy, or physical systems into a Bill of Materials or architecture tree. Each career year added another type or two of things to break down so that Syenna could "*divide and conquer*".

While doing so, the connections between the "things" receded into the background and the implicit design decisions behind every "fan-out" were seldom thought through with any significant rigor. Their rationale was never fully captured.

Fast forward 4+ decades and those relationships (inputs and outputs, connections, interfaces, dependencies, and flows) now seem central to success in any endeavor.

ChatGPT tells us that there is no one source to which we can attribute the saying, "*Systems Engineering is more about the edges than the nodes*". However, the LLM provides a variety of sources from which this concept may have evolved among different communities:

- Barry Boehm observed that the interfaces between components dominate system complexity.
- Eberhardt Rechtin's theme was that a system is defined more by the interactions among its elements than by the elements themselves.
- Donella Meadows defined a system as consisting of elements, interconnections, and purpose while observing that the interconnections are often more important than the elements themselves.
- Russell Ackoff stated that a system is never the sum of its parts; it is the product of their interactions.

Summarizing this thread, the LLM offers a concise formulation:

Systems engineering is less about designing the nodes than about engineering the edges - the interfaces, interactions, and dependencies that allow the parts to function as a coherent whole.

While thankful for early success using Separation of Concerns, current-day Syenna is nearly obsessed with understanding system dynamics - not just static interfaces and connections but also the flows across these links and how those flows drive changes in the system's state

over time. Concurrently, the boundary-less nature of the real world demands ever-increasing Context Analysis skills so that the important interactions, both planned and unplanned, can be identified and understood.

In the face of surging use of AI across engineering communities, the partnership between the humans who provide such context and the AI agents and models who predict behaviors from that context seems to be the big lever or control knob to master. A chat session with an LLM appears to be filled with dynamic flows of context, knowledge and insights, with emergent properties that can lead to successful innovations or catastrophic failures.
How are you modeling and managing the dynamics of the human-AI partnership in engineering?

Regards,
Syenna