

Assurance Cases as Knowledge Graphs

SOFTWARE CERTIFICATION CONSORTIUM MEETING

Annapolis, MD, May 9th, 2024

Mauricio Castillo-Effen, Ph.D.



Overview

- A. Background
- B. From Toulmin-inspired to “Digital” Assurance Cases
- C. CertGATE
- D. Knowledge Graphs



DISCLAIMER


The views and opinions presented in this presentation are solely the author's, and they do not represent the official policy or position of the Lockheed Martin Corporation or the Lockheed Martin Advanced Technology Laboratories.

A. Background

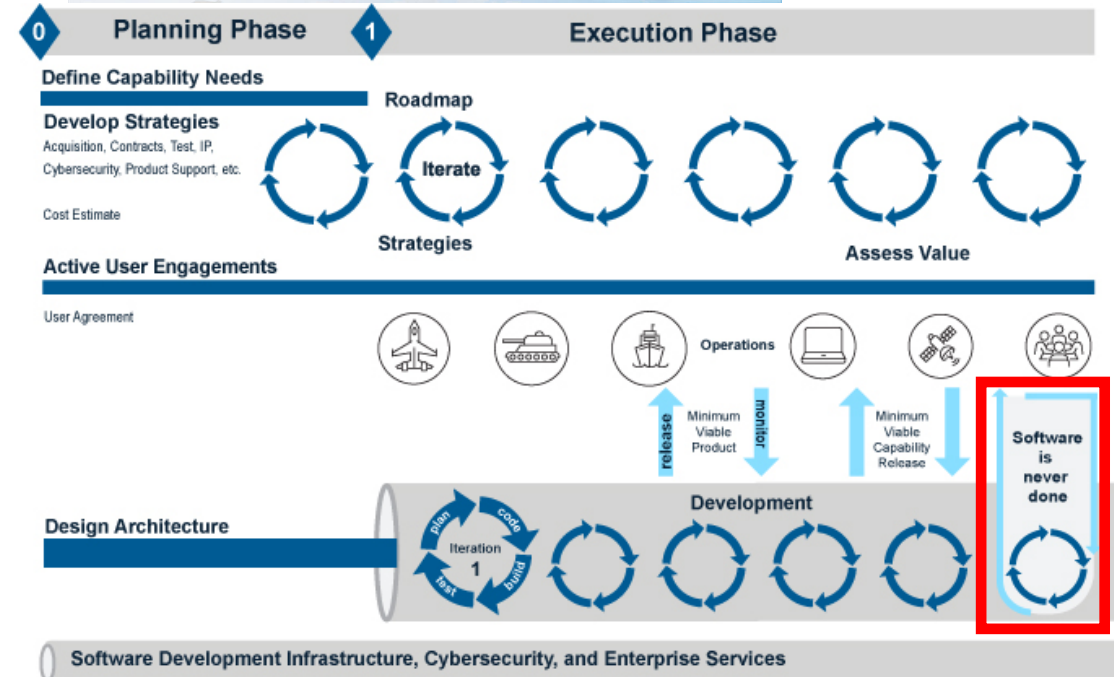


Acquisition Demands

- Capabilities are software/AI-enabled
- Capabilities will be assessed and assembled **just-in-time** and **at the edge** to address specific mission needs
- Capabilities are added/composed and weaknesses addressed continually (“**software is never done**”) based on dynamic needs while **maximizing reuse!**
- Capabilities may be obtained from **multiple vendors and providers**
- Decisionmakers need to **assess risk-benefit tradeoffs** based in part on the information obtained from assurance activities
- Transition from traditional waterfall V-model systems engineering to Dev*Ops transforms the view of **assurance from compliance- to value-driven**
- This assurance “**ConOps**” may apply to **dynamic aspects of other high consequence systems** that may contain less dynamic components



LM
Project Hydra



New acquisition regimes are characterized by **complexity** and **agility**

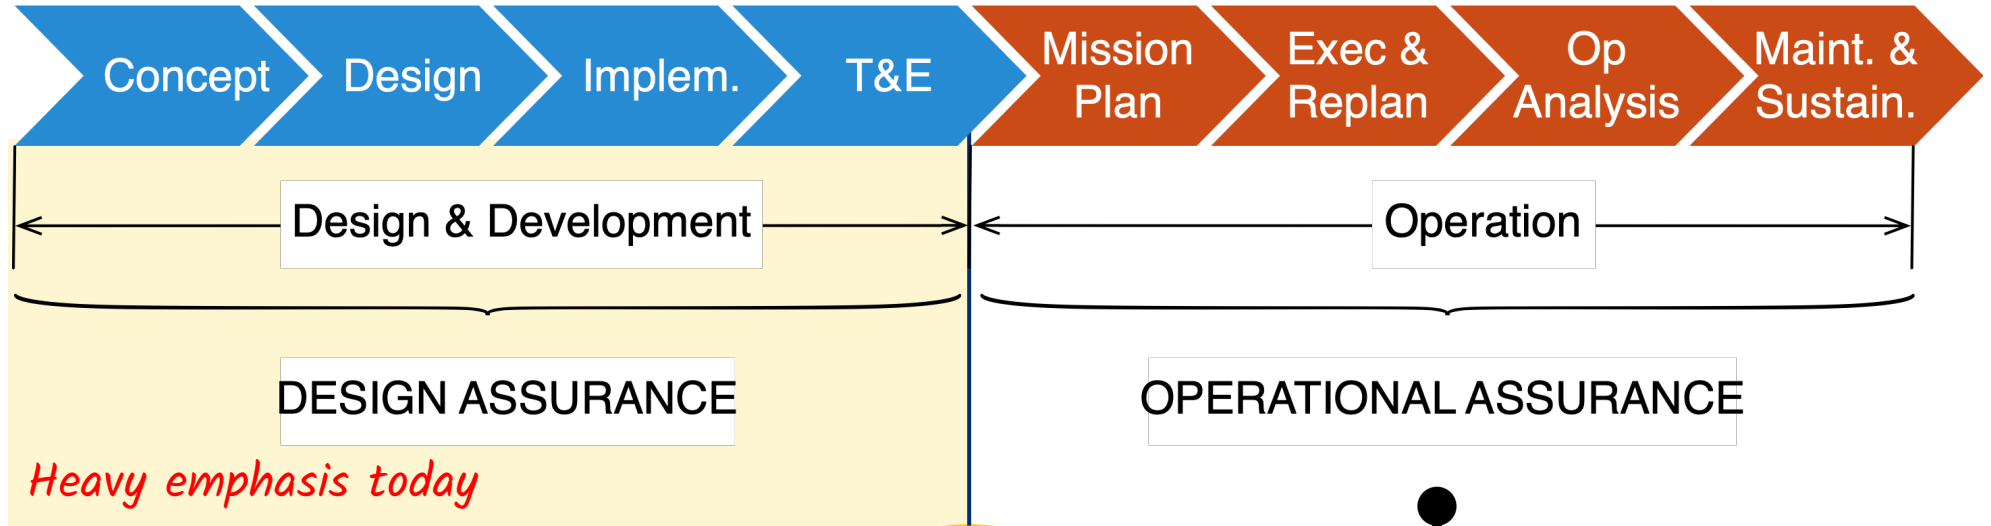
From DoD instruction 5000.87: OPERATION OF THE SOFTWARE ACQUISITION PATHWAY

Autonomous Systems for Firefighting


Exemplary Assurance Challenges

- Specifying and verifying correct and effective crewed-uncrewed heterogeneous team behavior in the presence of unique and rapidly evolving wildfire conditions
- Ensuring combined capabilities and resources (sensing, endurance, agility, payload, reliability, accuracy, etc.) applied in the firefighting mission are adequate and sufficient
- Maximize reuse
- Minimize unintended consequences
- Ensure the system is governable by a commander
- Ensure effective crewed-uncrewed collaboration
- Cyber resiliency






System's Lifecycle Today



- Designers
- Architects
- Developers
- Testers




ATO Officer / Regulator



Operator/User Supervisor

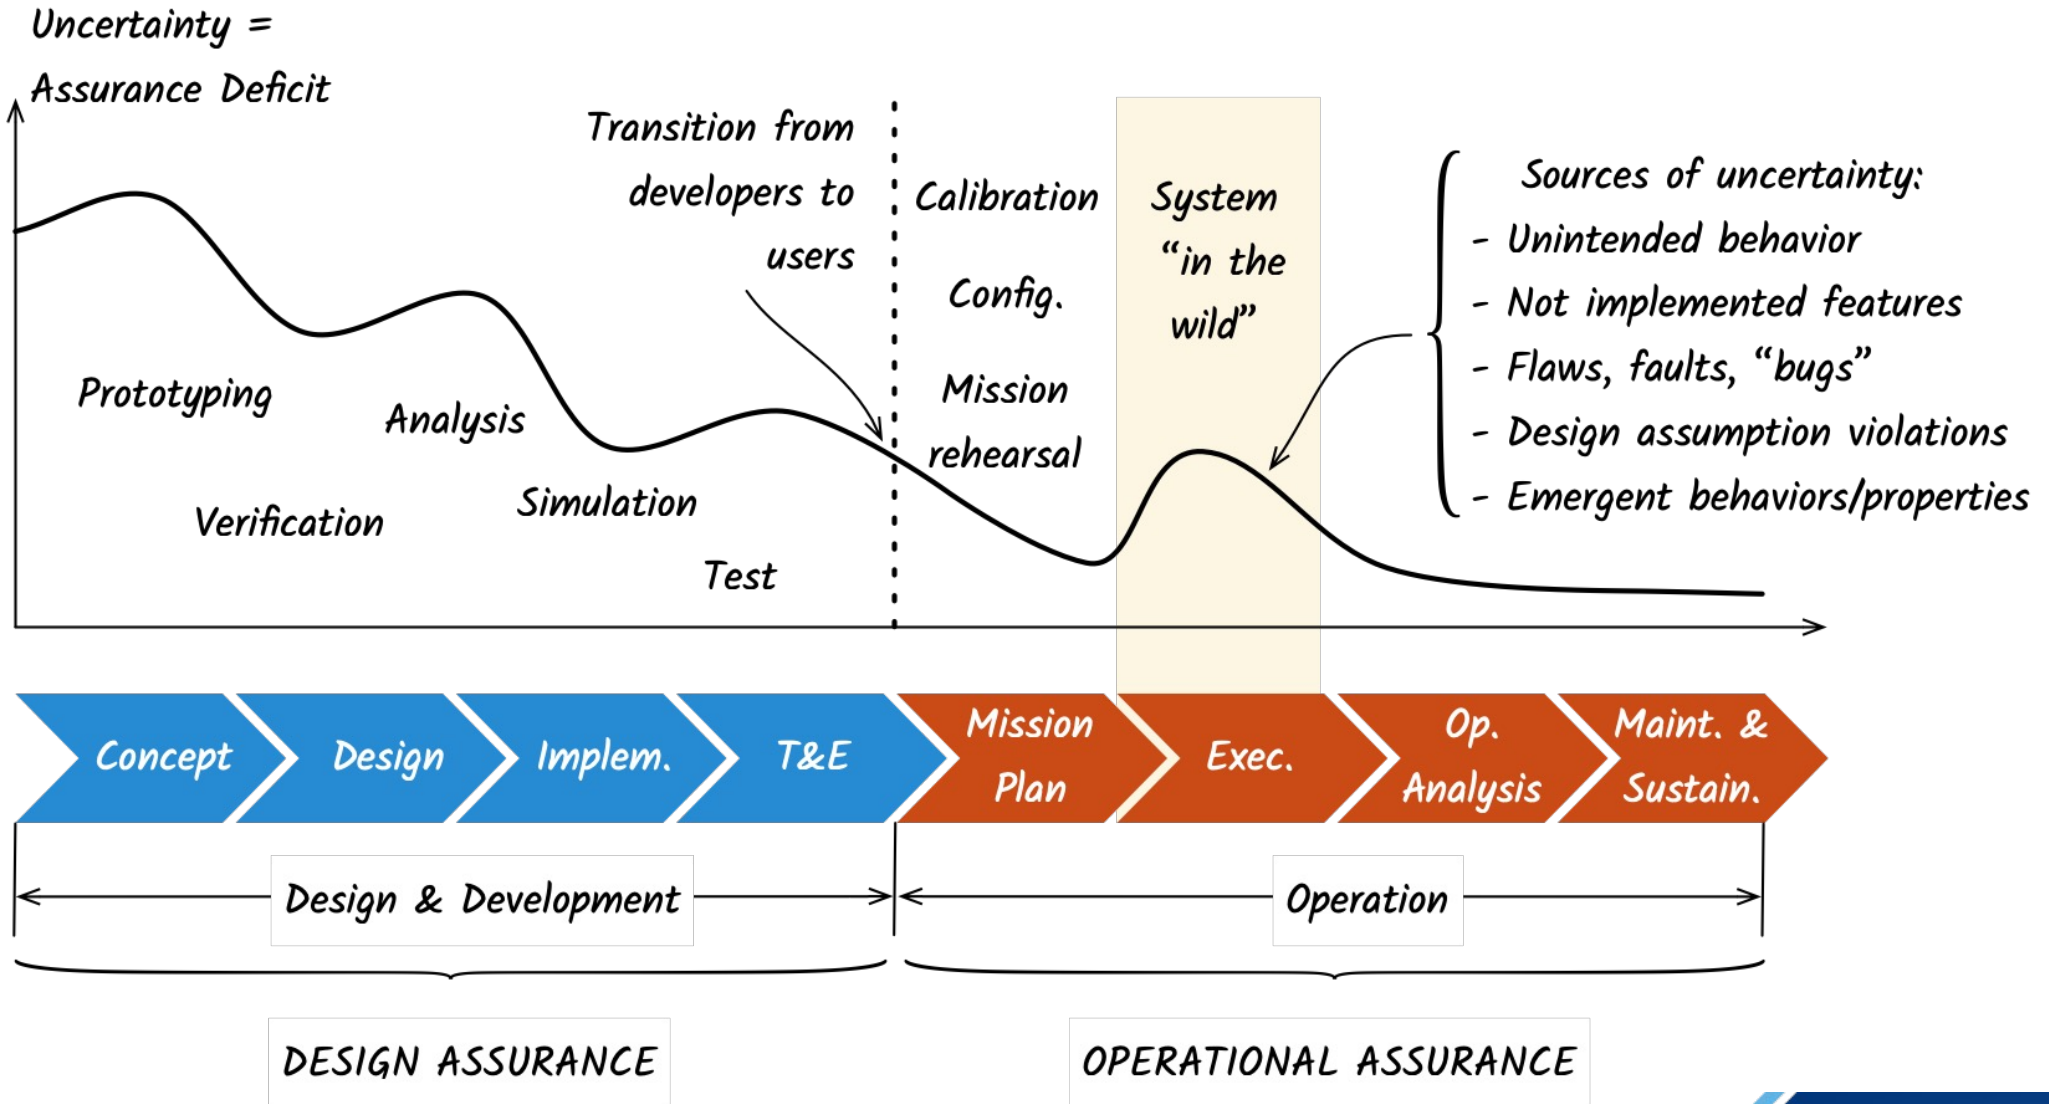


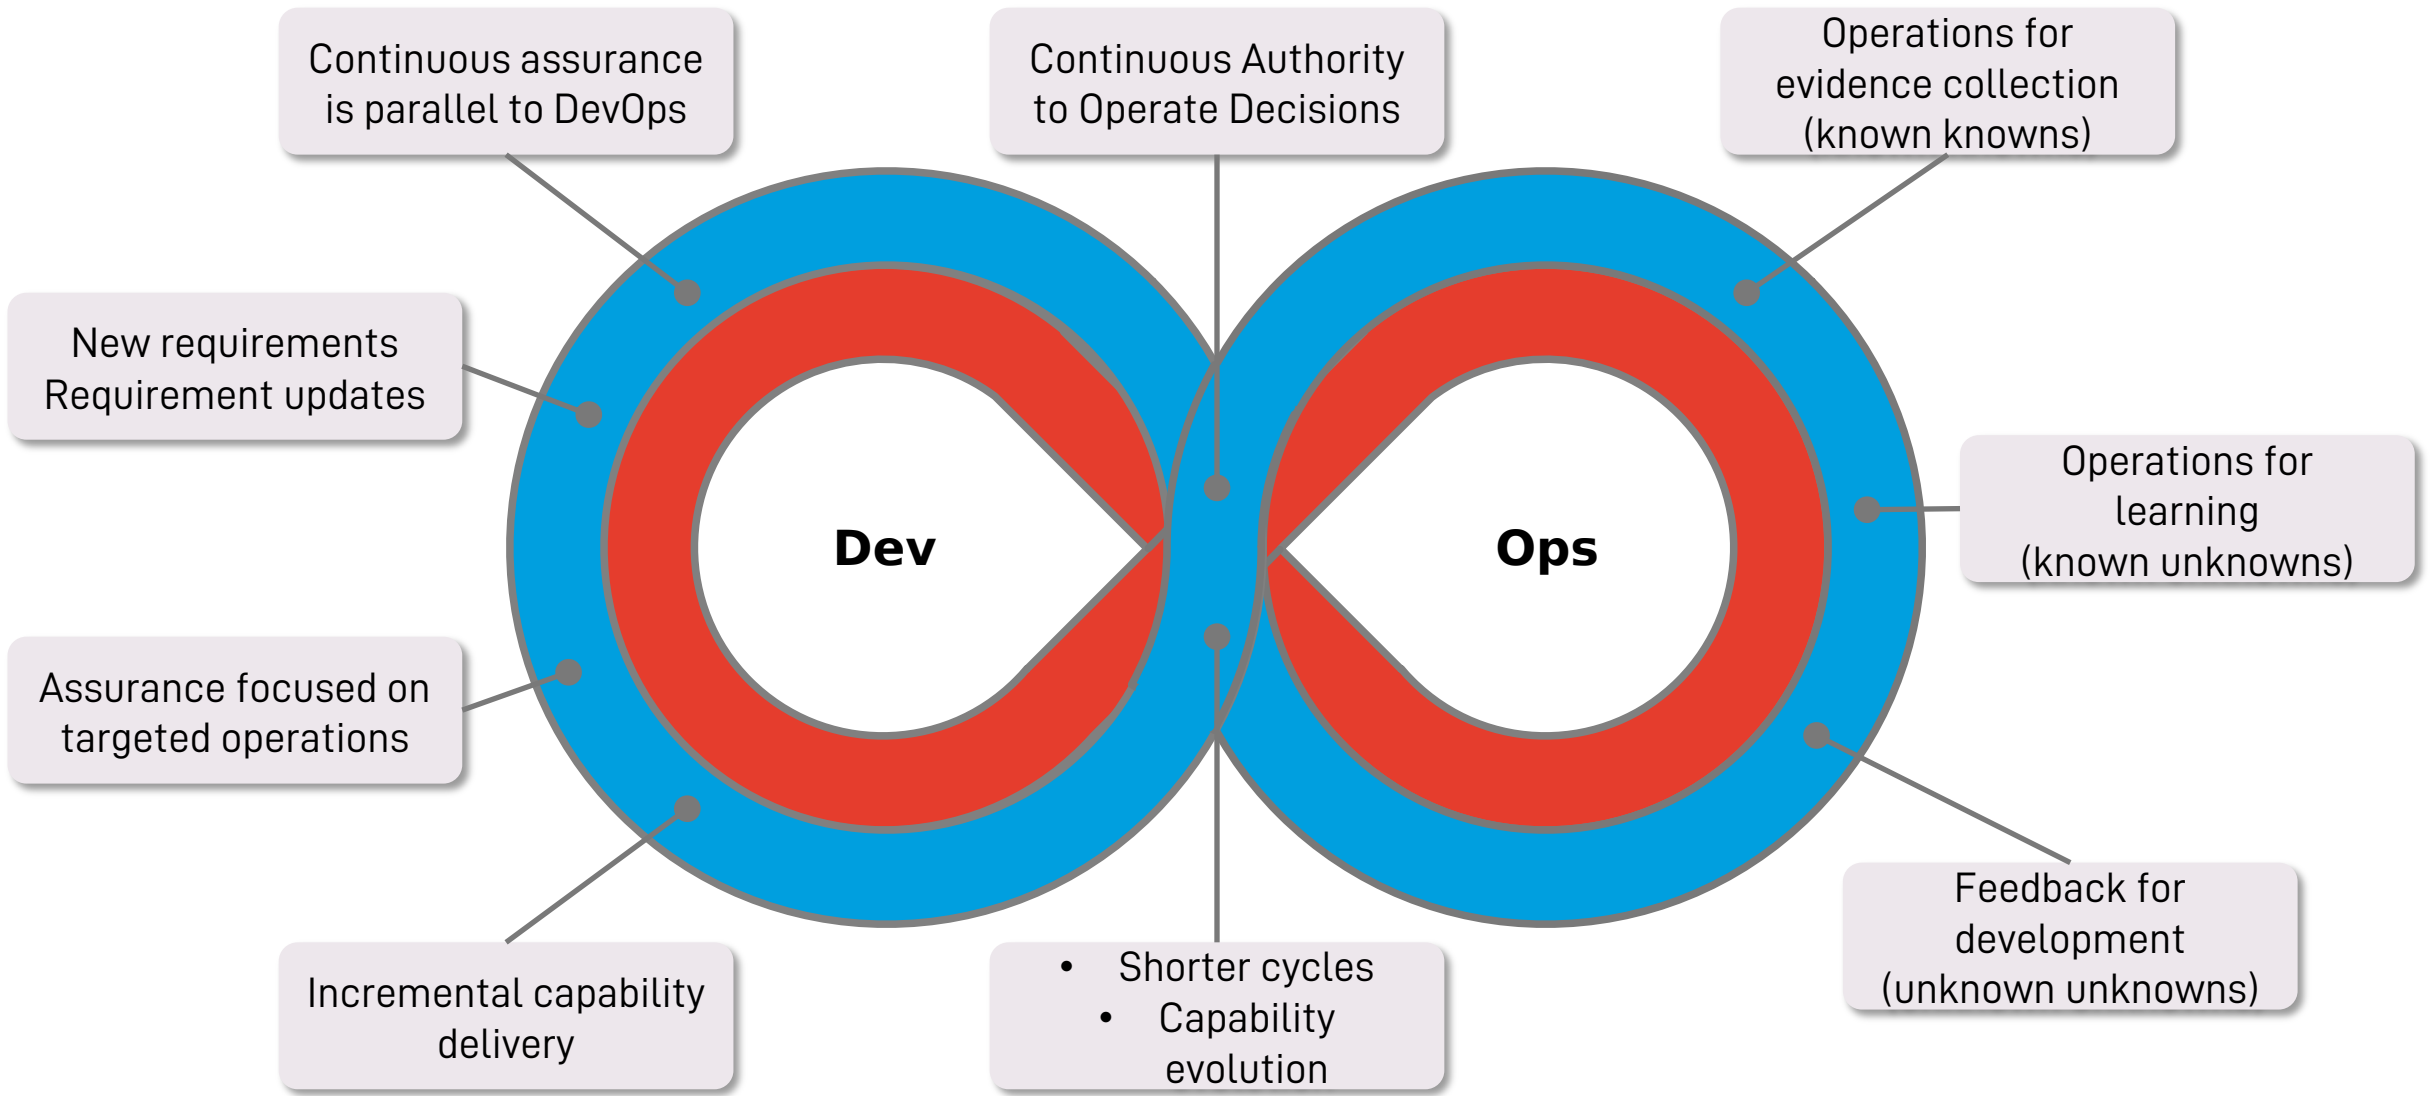
Customer / Program Executive Officer



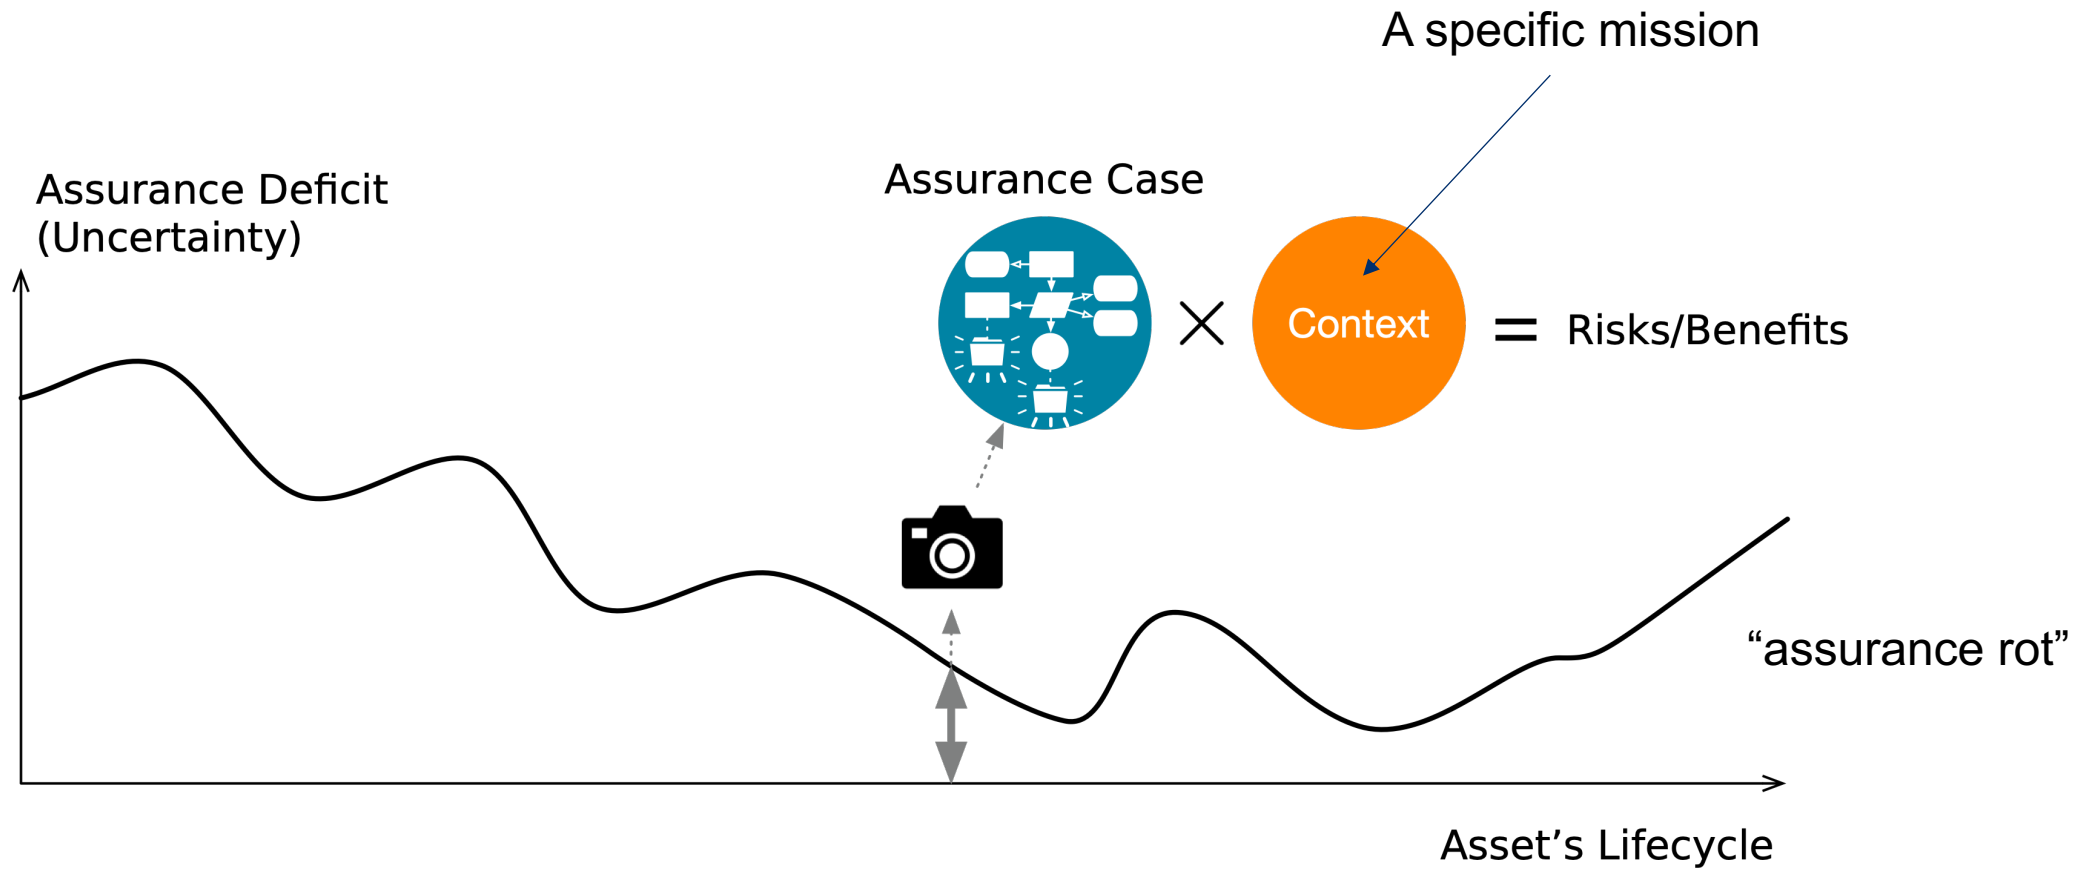
Society

Assurance as Uncertainty Reduction





Vision: Continuous Assurance Evaluation



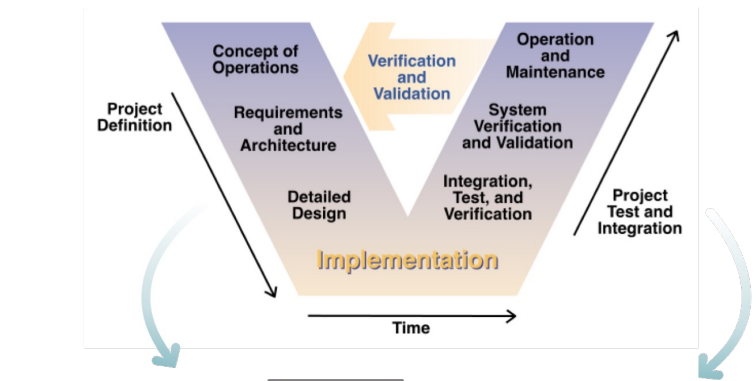
This is what we want but can't do with assurance cases today

From Toulmin-inspired to Digital ACs

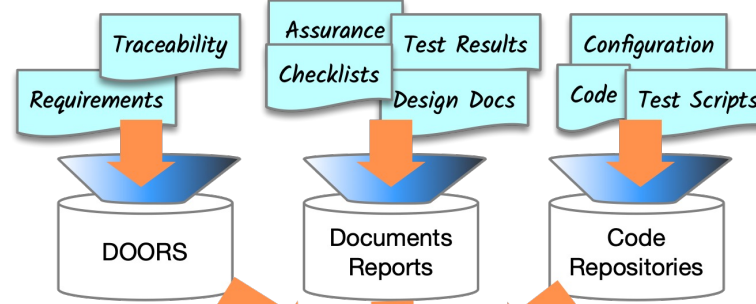


Methodology

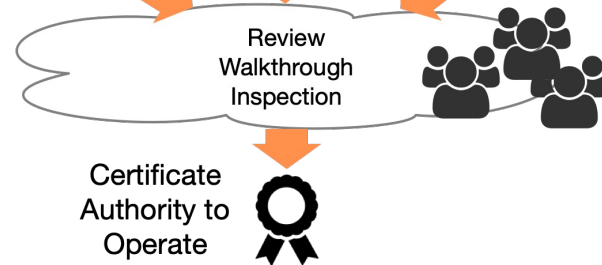
TODAY



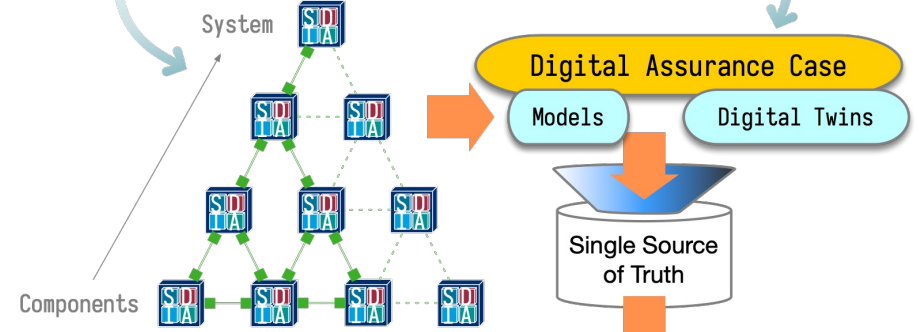
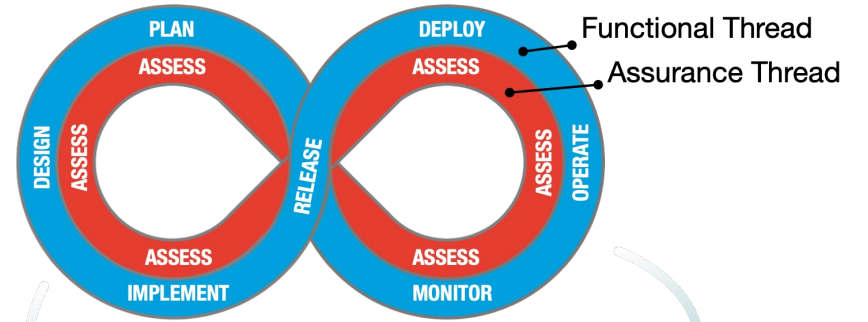
Artifacts



Assessment



VISION



Continuous Authority to Operate (c-ATO) with ribbon icon

Digital Assurance Cases are Aligned with Digital Engineering Methodology, Artifacts, and Agility

From Toulmin to Digital Assurance Cases

Digital Assurance Cases	Toulmin-inspired Assurance Cases
Concern Separation: Four separate and potentially distributed processes : i) Evidence Generation; ii) Assurance Case Construction; iii) Evaluation; iv) Presentation	ii)-iv) are intermingled and disjoint from i)
Inductive arguments: Suggest or hypothesize claim satisfaction. Assurance case evaluation leads to informed benefit/risk evaluation	Strive for deductive validity
Efficient Computational Constructs: focus efficient algorithmic storage and manipulation (querying, transformation) of evidence, arguments, and other metadata	Graphical/textual interactive editing (GSN, ACE, FAN, etc.)
Semantics: Use of controlled vocabularies (knowledge graphs), algorithmic semantic reasoning	Mostly, natural language for human readability

Digital Assurance Cases (DACs) are computational constructs designed to store, manage, and manipulate assurance information. **They are meant to be used by machines.**

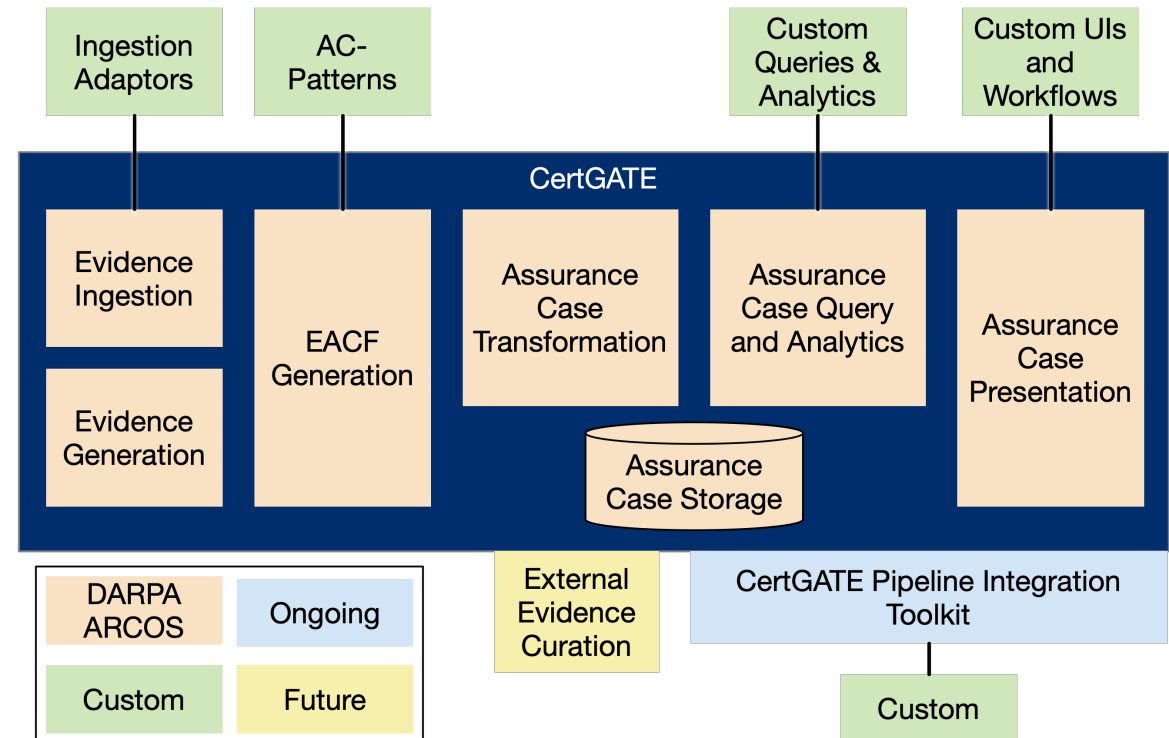
“Assurance information”: not just data items/work products but also metadata (arguments, traceability linkages, linkages to configurations/baselines, etc.)

CertGATE

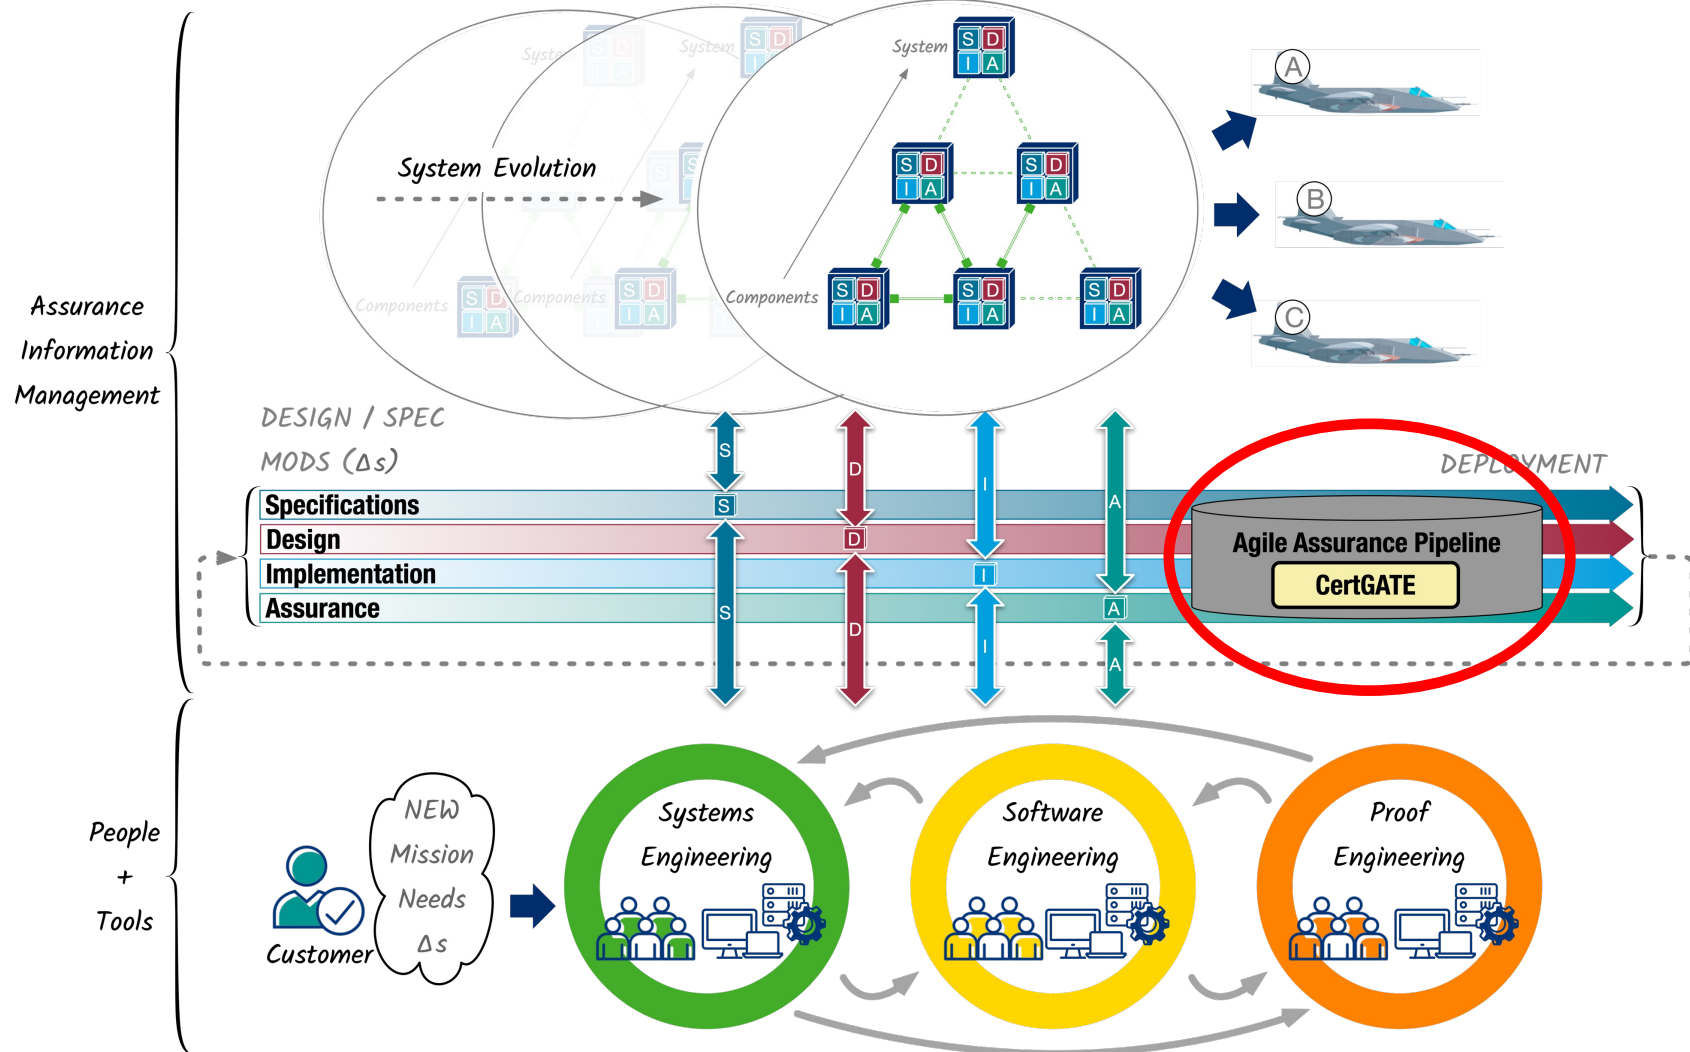


Separation of Concerns


- Separating concerns can help scale the application of assurance case technology for rapid certification
- It allows for development of best-in-class tools vs “jack of all trades” solutions
- Evidence Generation → build composable AC modules
- Assurance Case Construction
- Assurance Case Evaluation
- Assurance Case Presentation
- This can be achieved most effectively if all four concerns apply standard representations, e.g., SACM



CURRENT WORK

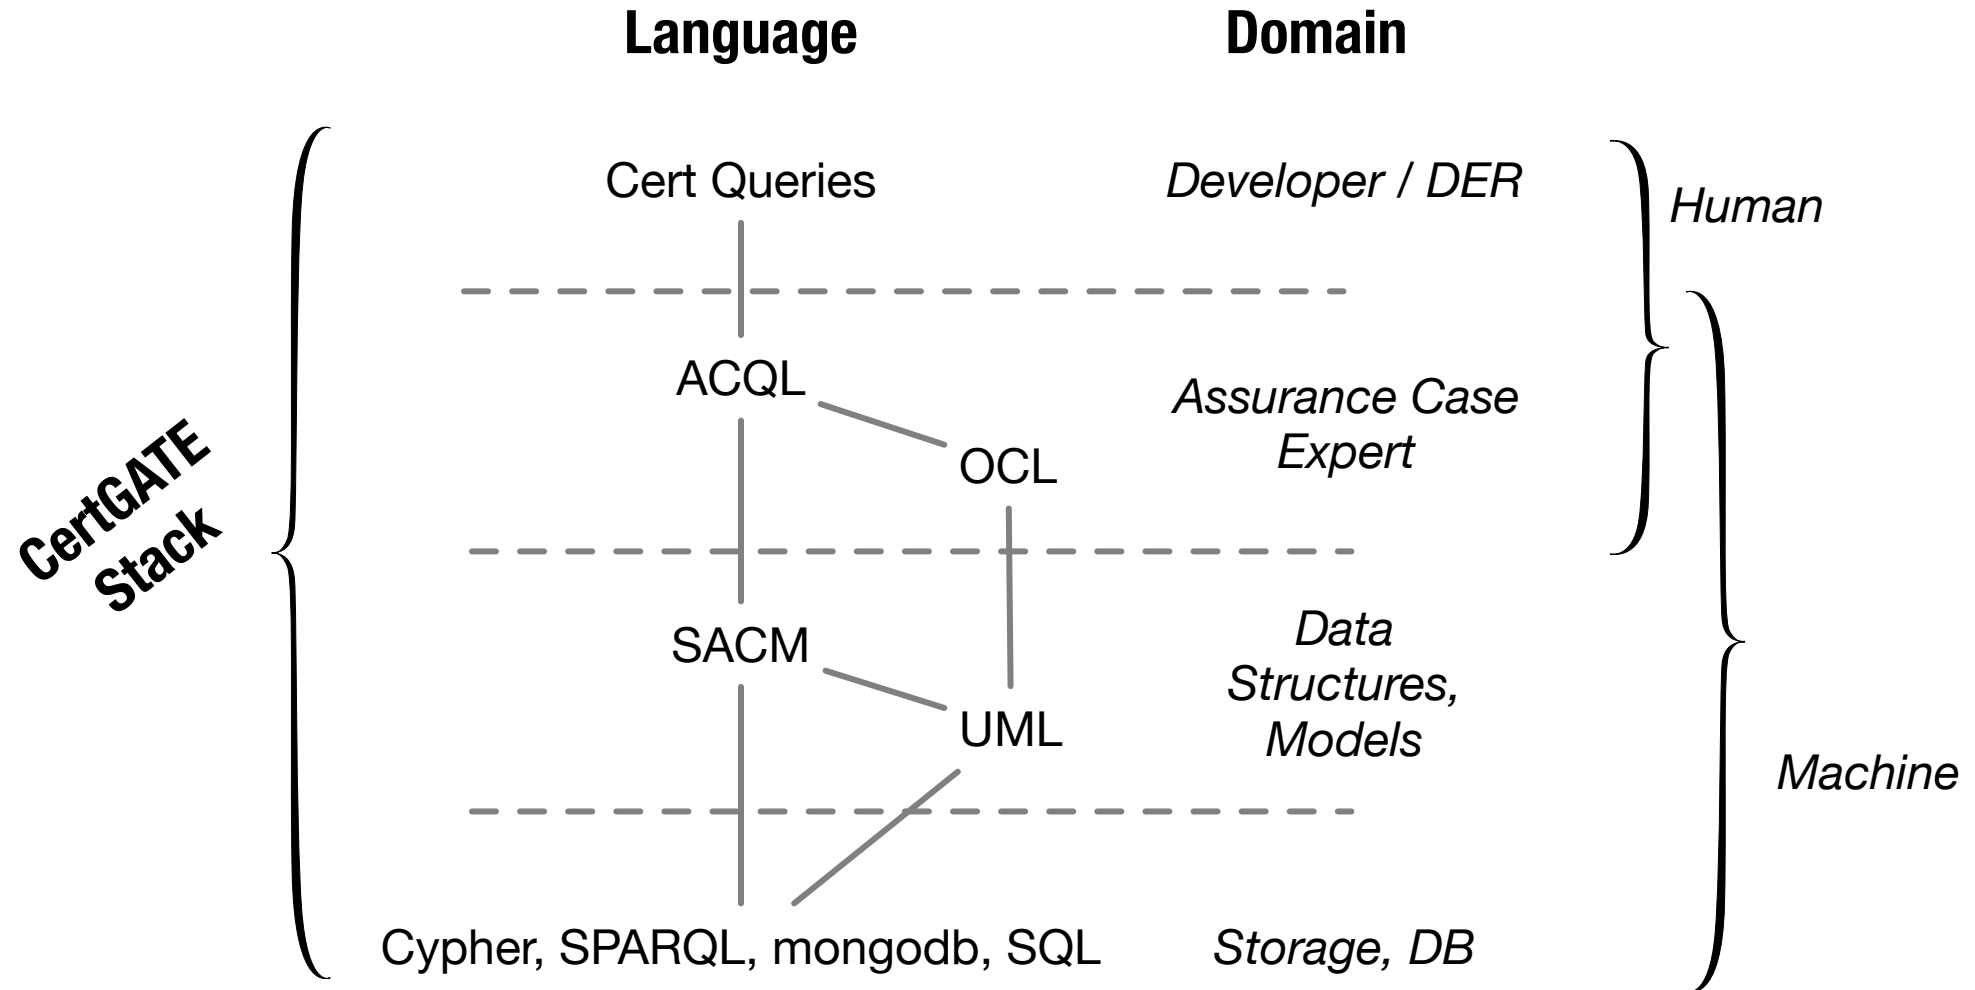


Four types of components



Type	Examples
S: Specifications	Requirements Contracts Properties Formal Specifications
D: Design	Models Architecture Diagrams
I: Implementation	Code Binaries Configuration files
A: Assurance	Proofs Tests Analysis results Inspection results

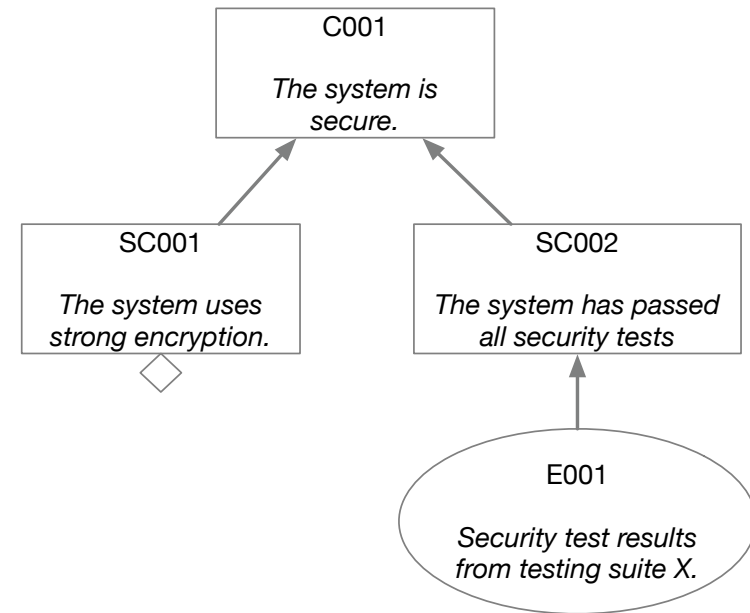
Assurance Case Query Language



Simple Example

SACM pseudo-code

```
{
  "Claim": {
    "id": "C001",
    "description": "The system is secure.",
    "supportedBy": [
      {
        "SubClaim": {
          "id": "SC001",
          "description": "The system uses strong encryption.",
          "status": "undeveloped",
          "evidence": []
        }
      },
      {
        "SubClaim": {
          "id": "SC002",
          "description": "The system has passed all security tests.",
          "status": "developed",
          "evidence": [
            {
              "id": "E001",
              "description": "Security test results from testing suite X."
            }
          ]
        }
      }
    ]
  }
}
```



Verify that all subclaims have evidence

```
context AssuranceCase::hasAllRequiredEvidence(): Boolean
derive: self.Claim.supportedBy->forall(subclaim | subclaim.evidence->notEmpty())
```

Result: False

Find IDs of undeveloped subclaims


```
context AssuranceCase::getUndevelopedSubclaimID(): Sequence(String)
derive: self.Claim.supportedBy->select(subclaim | subclaim.status = 'undeveloped').id
```

Result: SC001

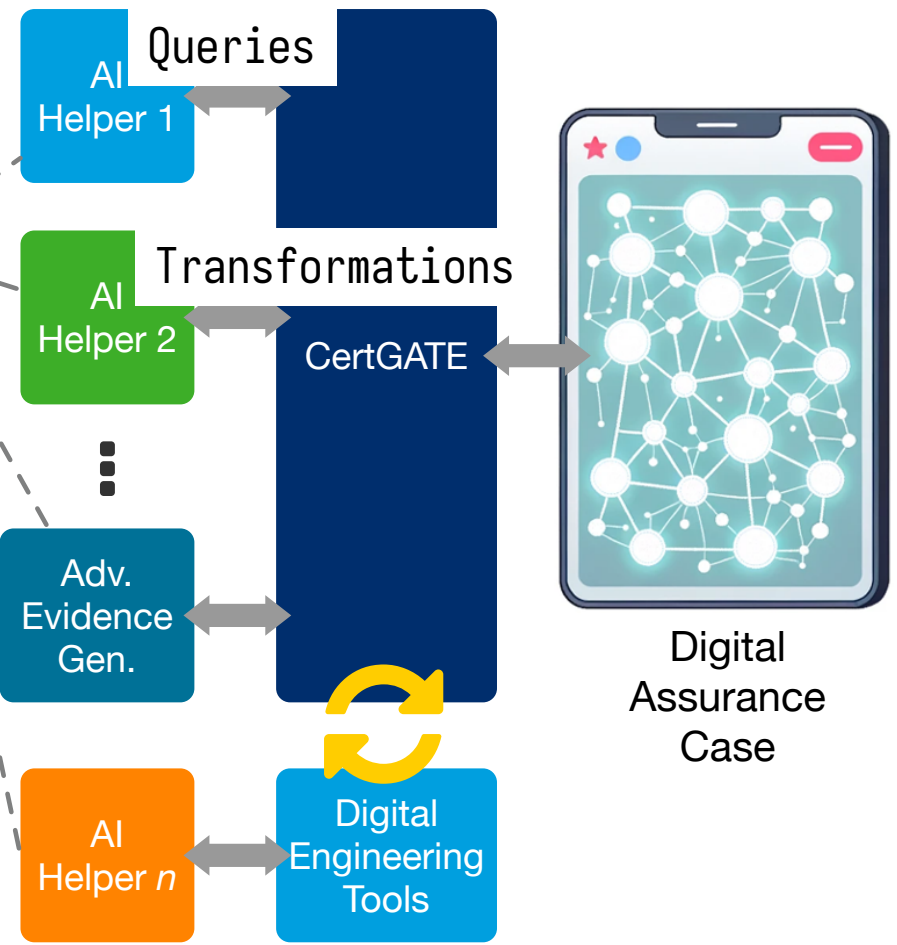
Obvious Next Step

Where are the most pressing assurance deficiencies for the firefighting mission given the current weather conditions?

What are the risks if I use a neural network-based flight control to deal with turbulence?



AI-enabled Interface



Summary So Far

Original intent of assurance cases

- **Capture our understanding** of the system’s ability to satisfy functional and non-functional properties and evidence supporting the confidence in these beliefs
- Make all assumptions, arguments, and justifications explicit **such that another human being can evaluate a “case”** supporting claims about these properties

Additional demands due to complexity and agility

- How do we capture information such that humans can arrive at conclusions with the help of machines more effectively?
- How do we make this process part of the design, development, and operation adding minimal friction?
 - Two aspects of integration: process (automation) and information (representation)
 - Knowledge is dynamic: it needs to be accumulated!

This is a knowledge representation problem!

Enter Knowledge Graphs



Brief History

- Early foundations (1950s–1970s): early foundations
- Semantic Web and Ontologies (1980s–1990s)
- Google’s Knowledge Graph (2012): “things not strings”
- Expansion in Use Cases, Advancements in Technology (2010s–Present)
- Examples today
 - Google Search, Imdb, Wikipedia, Thomson Reuters, Airbnb, Facebook, Elsevier, Etc.

The image shows a screenshot of a Wikipedia article for George Washington. On the right side, a red box highlights a Knowledge Card. The card includes a portrait of George Washington, his title as the 1st President of the United States, his term in office (April 30, 1789 – March 4, 1797), and a list of personal details such as his birth date (February 22, 1732), birthplace (Popes Creek, Virginia), and death date (December 14, 1799).

KGs are used widely and supporting technologies continue evolving

Knowledge Card

KG Definition

“A graph of data intended to **accumulate** and **convey** knowledge of the world or a portion of the world. They are designed to support sophisticated reasoning and to be **easily interpreted by humans and computers**. Knowledge graphs are typically constructed by integrating **data from many different sources** and are often used in applications where **complex reasoning and inferences** are required” (Aidan Hogan et al. “Knowledge Graphs Survey”, 2021)

The overlap between KGs and Assurance Cases is remarkable!

Advantages of KGs for Representing ACs

- Already co-located with other information that has graph (of graphs) structure, e.g.: requirements traceability, testing, terminologies, configurations control, version control, etc.
- Connectivity with other graphs enriches the semantics of assurance cases (“things not strings,” avoidance of “word salads” and “letter soups”)
- Ideally suited for ingestion of raw information. Most industry-relevant use cases are “brown-field.” There is already a big collection of artifacts, documents, code, etc. that need to be reused.
- Powerful automated reasoning, querying, and inferencing.
- We can “piggy-back” on continuously evolving technologies widely used in other sectors of industry → high performance, reasonable cost, demonstrated scalability, security, privacy, etc.

Future




KGs ❤️ LLMs

- Enhanced knowledge and factuality
 - Accurate and up-to-date information
- Improved contextual understanding
 - Entity recognition and disambiguation
 - Richer context
- Natural language interaction for ease of use
- Flexibility and personalization for reviews, inspection, and walkthroughs

Use of LLMs in conjunction with KGs will grow creating a richer substrate for AC reasoning!

Conclusions

We've started exploring the amenities of the ski resort. We'd like to meet other explorers.



LOCKHEED MARTIN 