Empirical Software Engineering (2026) 31:100
https://doi.org/10.1007/s10664-026-10847-x

Crafting effective boundary artefacts in software engineering: A guideline-based approach

Raquel Ouriques[1] · Fabian Fagerholm[2] · Daniel Mendez[3,4] · Tony Gorschek[3,4] · Baldvin Gislason Bern[5] · Victoria Vucic[5]

Received: 23 May 2025 / Accepted: 9 March 2026

Abstract

Context Boundary artefacts are shared artefacts that support collaboration by allowing different groups to interpret the same information in different ways. Software development activities benefit from them, as a single artefact can support stakeholders across different organisational boundaries. When these artefacts contain inconsistencies, such as incorrect information, practitioners' trust in them may decrease, leading to inefficiencies in task execution.

Objective This study developed and evaluated a guideline to support the creation of boundary artefacts in software engineering contexts.

Method We conducted a longitudinal, multi-phase study embedded in an industrial setting. The guideline was developed based on a literature review and prior findings from a previous case study and was then submitted for practitioner evaluation. A post-implementation analysis of the guideline was carried out after a period without researcher intervention.

Results Our guideline consists of 10 principles grouped into three categories: (1) Scope: stakeholders, boundaries, and terminology; (2) Structure: artefact format, transference, granularity, and additions; and (3) Management: evaluation, ownership, governance, and integration. Practitioner evaluations suggested that these principles support the creation of reliable, predictable, and functional boundary artefacts. However, practitioners also noted challenges during use, including the time-consuming nature of the activity and difficulties in understanding the concept of boundary artefact.

Conclusions Overall, the guideline was well received. After the non-intervention period, it was adopted as a standard by the partner company for artefacts such as security testing, standards documentation, and requirements specifications. Adoption challenges persisted, including cultural barriers and comprehension issues. Further applications across different artefacts could clarify how the principles influence their reliability, functionality, and predictability.

Keywords Boundary artefacts · Trust · Software engineering · Guidelines

Communicated by: Kelly Blincoe.

Extended author information available on the last page of the article

1 Introduction

"*As a customer, I want to receive email notifications when my order ships, so I can track its delivery.*"

This short sentence ultimately travels through many different parts of a software project. For the product owner, it reflects what matters to the customer. For developers, it signals the features to build. For architects, it points to systems that must interact. And for testers, it suggests what needs to be verified before release. Each group focuses on something different, yet all rely on the same piece of information. Because it helps these different teams make sense of the same artefact in their own way, a user story acts as a boundary artefact.

Boundary artefacts are objects that cross the boundaries of social contexts, carrying multiple meanings that address the information needs of various stakeholders (Star and Griesemer 1989). In this paper, we use the term "artefact" to refer specifically to boundary artefacts. In software engineering, such artefacts are widely used for purposes including requirements specifications, architectural descriptions, and project planning documents (Zaitsev et al. 2014; Barrett and Oborn 2010). They provide teams with concise information in different formats and support the retention and transfer of relevant knowledge, enabling coordination across diverse groups. For example, software requirements specifications are used by and therefore support roles such as architects, project managers, and customer-side stakeholders by conveying information that serves different purposes for each of them.

These artefacts can support collaboration by providing timely and accurate content throughout a software development project. When these artefacts contain inconsistencies, however, practitioners' confidence in them may decline (Ouriques et al. 2023a; Blomqvist 1997). Reduced trust can affect how practitioners rely on and use artefacts in their daily work. This loss of confidence can lead to several negative consequences for software projects, including delays, uncertainty about information availability, time spent searching for missing details, and the incorrect execution of tasks (Ouriques et al. 2023a).

This lack of confidence can also lead practitioners to avoid using artefacts and instead re-create their solutions, resulting in frustration. Such inefficiency in managing a company's knowledge through artefacts can lead to increased costs and wasted time (Ouriques et al. 2023b). As trust is a crucial factor guiding the utilisation of artefacts in software projects, it is important to understand which principles should be followed when creating these artefacts and to what extent these principles contribute to fostering reliability (Ouriques et al. 2023a).

Although there is a general understanding of the widespread production and use of artefacts in software engineering, to the best of our knowledge, the community still lacks necessary guidelines to support the effective creation and management of boundary artefacts. We contribute to addressing this gap through the following contributions:

- We develop a guideline focusing on the creation of boundary artefacts and submit it to evaluation with a company partner.
- We gather participants' perceptions on how the proposed guideline contribute to the creation of reliable, functional and predictable artefacts as well as its overall usefulness.
- We conduct a post-factum analysis to gain insights into participants' experiences and perceptions regarding the effectiveness and applicability of the guideline.

This manuscript is organised as follows: Section 2 presents a brief background to our study and related work. Section 3 describes our research methodology. Section 4 presents the developed guideline, the results of the practitioners' evaluation, and the post-factum analysis. Section 5 presents the perceived benefits and challenges with the application of the guideline. Section 6 provides a discussion of the results and implications that our current findings have for future research. Section 7 discusses the limitations and the threats to the validity of our study. Lastly, Section 8 presents our concluding remarks. We also provide Supplementary Material for transparency and potential re-application of the guideline in other contexts.

2 Background

The term "boundary artefact" (also referred to as "boundary object") was originally coined in a case study by Star and Griesemer (1989), conducted at the Museum of Vertebrate Zoology. The authors describe it as an object that carries information, flexible enough to adapt to local needs and to support shared understanding among different users, while maintaining its identity when used across multiple boundaries. Star (2010) further argues that a boundary object is not just any object that has multiple meanings. It must actively support cooperation between different groups through shared but adaptable structures, which is why the concept is most effective at the organisational level. In this study, we adopt the term boundary artefact, or, for brevity, artefact, due to its proximity to the software engineering discipline and its alignment with the general conceptualisation of artefacts in the field as stated by Méndez Fernández et al. (2019). The term "boundary object" can be understood to be broader in the sense that it covers objects of vastly different kinds, including artefacts typically encountered in software engineering. Thus, for the scope of the present study, all boundary artefacts are boundary objects, but of kinds specific to software engineering.

The ability of these artefacts to cross and adapt to multiple contexts highlights the dynamic nature of boundaries within organizations. Boundaries are not always synonymous with limits in organisations. To Star and Griesemer (1989), they are often permeable. That is, they do not block information within an organisation but rather create a transition where individuals on one side of the boundary interpret information from their perspective, while those on the other side view it from a different perspective, even if they are looking at the same information or working towards the same goals. These characteristics make boundary artefacts especially useful in complex domains such as software development, where they can serve multiple purposes across different teams and contexts.

In this study, we use the term boundary to refer to a socially or organisationally defined separation between groups, roles, or domains that differ in expertise, responsibilities, or perspectives. Boundaries are not necessarily rigid structural divisions but represent differences in interpretation, practices, or information needs that require coordination. Boundary artefacts therefore function as mechanisms that enable collaboration across these differentiated contexts.

Let us consider three examples from the software engineering literature. In a grounded theory study, Sedano et al. (2019) examined the role of the product backlog (an ordered list of work items, commonly referred to as user stories) and how it emerges. They found that the product backlog serves as a boundary artefact that bridges the gap between the boundary

where user stories are generated (for example, through requirements elicitation and mock-ups) and the boundary where these user stories are developed (including activities such as coding, testing, and integration).

Another example, involving more complexity in both the number of stakeholders and the boundaries crossed, comes from a longitudinal case study by Barrett and Oborn (2010). The authors found that the requirements specification document served an intricate network of stakeholders across three major boundaries: bridging geographically dispersed teams (in India and Jamaica), communicating user needs to development teams, and supporting project management and control. Lastly, a more recent example comes from DevOps practice. Matthies et al. (2023) conducted a multiple-case study with practitioners from nine companies, showing which artefacts act as boundary objects in DevOps settings and mapping the stakeholders, concerns, and artefact attributes that shape their effectiveness in cross-team collaboration. Software development benefits from the use of such artefacts, as a single artefact can provide information to stakeholders across different boundaries, particularly in geographically distributed environments.

Nevertheless, these artefacts can become ineffective. Wohlrab et al. (2019) found, through a case study, that boundary artefacts can fail when they become degraded, overloaded, or inconsistently maintained, making it difficult for stakeholders to find relevant information or rely on the artefact as a shared reference. In a case research approach, Zaitsev et al. (2014) similarly found that boundary artefacts may lose effectiveness when they are too rigid to accommodate change or when they erode over time due to frequent modifications, leading to misalignment and confusion among teams. In a case study examining virtual communities, Akoumianakis et al. (2011) found that boundary artefacts can also fail for social and organisational reasons, such as breakdowns in trust, unmet partner commitments, or unstable collaboration practices, all of which prevent the artefact from functioning as a meaningful shared object across groups.

These vulnerabilities are further compounded by the fact that boundary artefacts are embedded in power dynamics that can shape how they are interpreted and used. Star and Griesemer (1989) point out that these artefacts are inherently political and subjective, emerging from negotiations among actors with differing interests, values, and authority. To the authors, their creation and use are not neutral but involve managing conflicts and balancing power relations, allowing groups to cooperate without fully agreeing. Examples from the literature further illustrate how politics and trust affect boundary artefacts. Kimble et al. (2010) found that boundary artefacts can be used strategically by mediators to advance their interests, such as steering the direction of the collaboration; however, their effectiveness depends on the trust stakeholders place in those mediators. Huvila (2011) explains, through an interview study, that boundary artefacts embody subjective knowledge claims that stakeholders must trust to accept and use them. Without such trust, their authority and usefulness quickly diminish. Further extending this perspective, Levina and Vaast (2005) argue that boundary artefacts become effective only when they are enacted in practice as "boundary objects-in-use." Their work shows that successful boundary spanning depends not solely on the artefact itself but also on the collaborative practices through which different groups mobilise it to create a shared field of work. This highlights that the effectiveness of a boundary artefact is shaped by both organisational dynamics and the ways stakeholders engage with it.

Trust can play a significant role in driving stakeholders' behaviour towards a boundary artefact. Lansing and Sunyaev (2016), through a comprehensive literature review study on information technology, investigated which trust beliefs best describe trust in information technology artefacts. They identified three key trust beliefs: reliability, the perception that the artefact provides accurate content; predictability, confidence that the artefact content is always provisioned as requested; and functionality, the artefact performs as needed for the task environment. Because their study offers a robust conceptual model for understanding trust in inanimate artefacts, we adopt these three beliefs to examine practitioners' perceptions of how our guideline supports creating artefacts that are reliable, predictable, and function as planned.

2.1 Practices for Creating and Managing Boundary Artefacts

This emphasis on what makes boundary artefacts effective, including the role of stakeholder trust, highlights the importance of understanding how these artefacts can be created and managed effectively. We found studies in software engineering that have investigated or developed solutions or practices targeting boundary artefacts in software engineering, either for creating them or enhancing their effectiveness. The number is limited but offers valuable insights into their role in supporting collaboration across stakeholder groups.

Jain et al. (2014) conducted a multiple-case study to investigate how boundary artefacts supported requirements engineering in both conventional and agile contexts. They identified three principles that shaped artefact effectiveness: shared terminology, which enabled stakeholders with different expertise (customers and developers) to coordinate around requirements and design representations; controlled changes, where evolving artefacts were balanced with periodically updated, stable ones to maintain platform integrity; and balanced formality, in which teams combined lightweight, informal artefacts with selectively formal documentation to avoiding complexity in creating and maintaining.

Wohlrab et al. (2018) examined, through a design-science study involving six automotive companies, how boundary artefacts and other systems-engineering artefacts were managed in large-scale agile automotive development. Their analysis highlighted several practices that shaped effective artefact use, including finding an appropriate level of detail so that information remained useful to stakeholders and handling scattered information distributed across tools and departments. They showed that teams struggled particularly with architecture descriptions, high-level requirements, and variability information when these artefacts were either too detailed, too vague, or fragmented across locations, all of which undermined their function as boundary objects. To address these issues, the authors suggested practical guidelines such as creating lightweight high-level artefacts during early stages of the development, establishing groups of representatives to coordinate changes across boundaries, and ensuring artefacts were continuously reviewed and refined so that degradation and inconsistency could be mitigated.

Van Loggem and Van Der Veer (2014) examined, in a conceptual study, how documentation could function as a boundary artefact to align heterogeneous stakeholders in large software development teams. They argued that effective documentation practices required clear ownership at each stage of development, with responsibility explicitly transferred as the project moved from requirements analysis to architectural design, implementation, and finally user support. Their approach further required that the documentation remain con-

tinuously extensible, capable of incorporating new content such as evolving requirements and features, system logic, user interaction details, and even low-level coding solutions, ensuring that it reflected the current state of the system throughout development. Through this documentation-centered workflow, the authors provided practical guidance for shaping documentation to sustain communication, support shared understanding, and serve as a robust boundary artefact across multiple communities of practice.

Ouriques et al. (2023a) conducted an exploratory case study examining how trust in a boundary artefact shapes stakeholders' behaviour within a large company that provides network solutions in video surveillance, access control, intercom, and audio systems. Their findings revealed that ineffective content management and misalignment between contributors and users resulted in inconsistencies that eroded trust and prompted stakeholders to develop unofficial workarounds. Based on these observations, the authors proposed several practices for improving the design and management of boundary artefacts: establishing a formal strategy for managing the entire content so that responsibilities and updates are clearly governed; performing periodic evaluations and feedback cycles to detect outdated or unreliable information; mapping stakeholders and their needs to ensure that the artefact's content supports all roles that rely on it; defining an approach for handling terminology so that meanings remain consistent across social groups; and finally, enabling the artefact to accommodate experimental and evolutionary content, ensuring that emerging features and temporary information can be incorporated without fragmenting knowledge.

These studies provide valuable insights into how boundary artefacts function in specific settings. The practices they propose are inherently situational, shaped by the organisational contexts, artefact types, and stakeholder constellations under investigation. What remains missing in the literature is a set of generic, context-independent, design-oriented principles that can guide the creation of boundary artefacts across software engineering environments. To address this gap, this study offers a complementary perspective by synthesising findings from diverse empirical cases into a consolidated guideline and by empirically examining its use in practice. To this aim, we investigate the following research questions:

- RQ1: What principles should guide the creation of boundary artefacts in software engineering?
- RQ2: How do practitioners perceive the developed guideline for creating reliable, predictable, and functional boundary artefacts?
- RQ3: What are the outcomes of using our proposed guideline in the creation of boundary artefacts?

3 Research Methodology

This study aimed at develop and evaluate a guideline for creating boundary artefacts in software engineering. Our motivation stemmed from discussions with our industry partner and from gaps identified in prior research, particularly the work of Ouriques et al. (2023a), which suggests that insufficiently planned boundary artefacts often struggle to meet stakeholders' needs consistently. This indicated a need for a deeper understanding of how such artefacts can be more deliberately designed and governed in practice.

This study follows a longitudinal, multi-phase study embedded in a single industrial case. The organisational setting of Axis Communications provided the context in which the guideline was developed, evaluated, and later observed in use over time. Rather than analysing the organisation itself as the primary object of study, the case serves as the setting through which we examine boundary artefact creation and governance practices in a real-world software engineering context.

The longitudinal dimension allowed us to observe how engagement with the guideline unfolded over time within the organisational context. A longitudinal perspective is particularly well-suited for examining how new practices are incorporated into ongoing work, how they stabilise or change, and how practitioners' perceptions evolve (Hall 2007; Fucci et al. 2018). In this study, we followed how practitioners engaged with and utilised the guideline across an extended period, including more than a year of independent use without researcher intervention, enabling us to capture how boundary artefact practices developed over time.

The study comprised multiple phases conducted within the same organisational setting and employed multiple sources of empirical evidence. We employed three empirical data collection methods to develop and evaluate the guideline: interactive workshops, a questionnaire, and a post-factum group interview. The workshops were chosen because they enable practitioners to collaboratively interpret the guideline principles and articulate their reasoning, providing insight into how these principles are understood in practice. The questionnaire was included to generate a structured reflection tool to capture individual perceptions following the workshops. The post-factum interviews allowed us to explore practitioners' experiences after an extended period of independent use of the guideline, providing insight into its practical impact and long-term applicability. The practitioners selected for the workshops and questionnaire all had extensive experience in the company and held roles with the authority to make decisions regarding the creation, maintenance, and use of boundary artefacts. Their involvement therefore provided informed and contextually grounded perspectives that are appropriate for our case. The questionnaire responses were analysed qualitatively to identify tendencies and contrasts in perception, complementing the data obtained through workshops and interviews and supporting methodological triangulation.

To address our research questions, we utilised a mixed-methods approach conducted in four steps (see Fig. 1). First, we collected principles from a literature review our prior case study in the same company (Ouriques et al. 2023a). Second, we applied and evaluated the guideline through two workshops and a questionnaire. Third, we refined the guideline based on feedback collected during the workshops. Finally, we allowed practitioners to utilise the revised guideline independently for one year and eight months, after which we conducted interviews to gain insights into their experiences with the guideline. The details of each step are presented in the following subsections.

3.1 Step 1 - Development of the Guideline

In the first step (see Step 1 in Fig. 1), we developed the initial version of the guideline. We conducted four searches in the Scopus database using distinct search strings (see Table 1). Our aim was to identify studies that provided actionable guidance such as guidelines, frameworks, checklists, or documented practices related to the creation or management of boundary artefacts (or boundary objects more broadly). Because relevant discussions

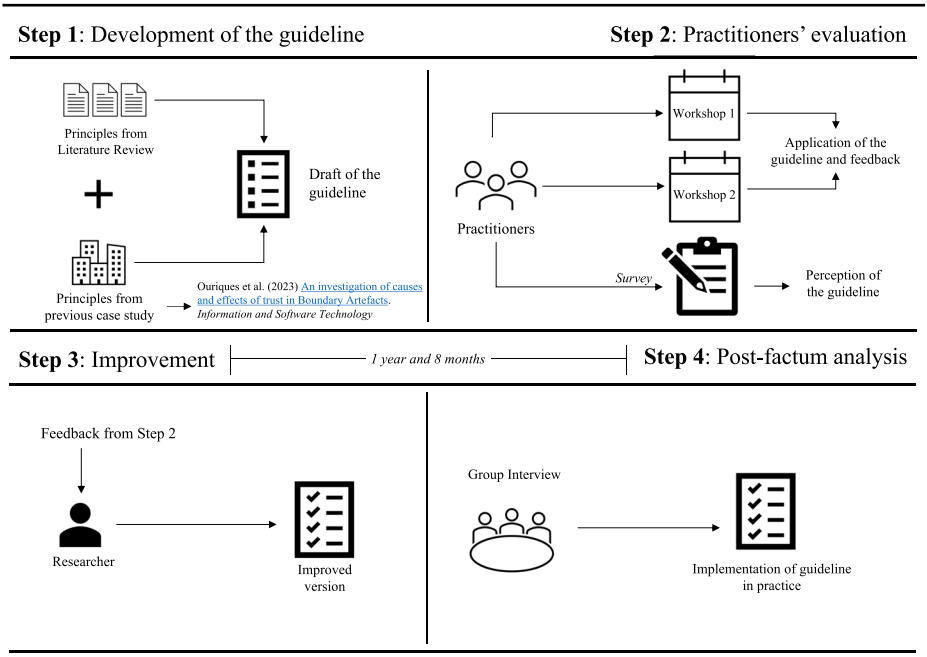

Fig. 1 Research approach followed for developing and evaluating the guideline

Table 1 Search strings used in the exploratory literature search informing guideline development

Search strings	Date	Hits	Number of papers selected
1 - ((“boundary artefact” OR “boundary object”) AND (feedback OR “feedback loop”))	10/2022	27	1
2 - (“designing boundary artefact” OR “designing boundary object”)	10/2022	2	0
3 - (“creating boundary artefact” OR “creating boundary object”)	10/2022	4	2
4 - ((“boundary artefact” OR “boundary object”) AND (“software development” OR “software engineering”))	10/2022	65	9

appear across several disciplines, we intentionally broadened our search beyond software engineering.

Our objective was not to perform a systematic literature review in the sense defined by Kitchenham and Charters (2007). Instead, the goal was to collect a set of sources that could inform an initial version of our guideline. Accordingly, we applied the following inclusion criteria during screening:

– The study discussed boundary artefacts/objects in a way that included practical implications for how they are created, structured, maintained, or governed.
– The study provided explicit principles, recommendations, or design considerations that could be translated into guideline elements.

We also applied the following exclusion criteria:

- Studies that discussed boundary objects only in theoretical or conceptual terms without offering actionable guidance.
- Studies where boundary objects were mentioned only tangentially.

We screened titles, abstracts, and introductions against these criteria. Studies that met the inclusion criteria were read in full, and we extracted the principles, guidelines, or actionable insights they proposed. This process resulted in a set of 17 principles synthesized from the existing literature. To complement this, we also incorporated five principles derived from our previous case study (Ouriques et al. 2023a), which identified concrete, practice-oriented challenges in creating and maintaining boundary artefacts in an industrial setting. After extraction, similar principles from all sources were analysed and grouped into three overarching categories Scope, Structure, and Management, as detailed in Section 4.2. A complete list of selected studies and the corresponding extracted principles is provided in the Supplementary Material.

In this study, a literature review was chosen to ground our guideline in prior empirical work. Research on boundary artefacts is dispersed across several fields, including software engineering, organisational studies, and knowledge management, and we did not find any consolidated set of design principles for creating boundary artefacts. Conducting a literature review allowed us to identify and synthesise existing practices and recommendations relevant to boundary artefact creation.

3.2 Step 2 - Practitioners Evaluation

In step two, we evaluated our guideline with practitioners from our industry partner and collected feedback for further improvement (see Step 2 in Fig. 1). We conducted two interactive workshops and a questionnaire to better understand how the guideline could be used in practice and how practitioners perceived it in detail.

We conducted our workshops at Axis Communications, a company specialising in network solutions for video surveillance, access control, intercom, and audio systems. With over 4,000 employees across more than 50 countries, Axis is headquartered in Lund, Sweden, where the evaluation activities were focused. Each workshop lasted approximately 90 minutes. A total of six practitioners participated in the workshops (see Table 2), applying the guidelines to two boundary artefacts at different maturity stages, one per workshop. In

Table 2 Practitioners and their roles

Participant	Role	Years in the company	Work. 1	Work. 2	Questionnaire	Post-factum
P1	Expert engineer	Over 15 years	✓	✓	✓	✓
P2	Test automation engineer	Over 4 years	✓	✓	✓	
P3	Expert test engineer	Over 5 years	✓	✓	✓	✓
P4	Product information specialist	Over 12 years	✓			✓
P5	Expert engineer	Over 7 years	✓		✓	
P6	Test engineer	Over 15 years	✓		✓	

the first workshop, six practitioners attended, three of whom also participated in the second workshop. The first author led the workshops. We recorded the audio from both sessions to facilitate detailed analysis. The practitioners were selected by the company because of their senior experience and their authority to make decisions related to boundary artefact creation and maintenance.

3.2.1 Workshop 1

During this workshop, practitioners evaluated Boundary Artefact 1 (A1) by mapping the guideline's principles to their current practices. This allowed them to assess the alignment of each principle, its current status, and its importance for future development. Six practitioners participated in this first workshop. A1 (see example in Fig. 2) has been extensively utilized within the company for many years. They use it daily, creating or applying its content as input for their tasks. This boundary artefact presents the core list of features of all the products in a software product line of the company (covering over 200 products). This content is stored in an XML file that includes the developed features for the device software used in the company's products, whether they are new, modified, or deprecated.

3.2.2 Workshop 2

In the second workshop, the practitioners analysed the relevance of the principles for creating Boundary Artefact 2 (A2) and briefly discussed how they might implement them. A total of three practitioners participated in this workshop. They had also attended the first workshop and were selected by the company due to their close involvement with A2, which was initially intended to complement A1. The A2 (see example in Fig. 3) is a type of product feature model that outlines the logical relationships between product-related facts expected to be true. For example, a rule might state, 'If a product in the software product line has Feature A, then the product must also have Feature B.' These rules are defined in YAML format and can be used to automatically determine whether products in the product line adhere to the specified rules.

```
<?xml version="1.0" encoding="utf-8"?>
<?xml-stylesheet type="text/xsl" href="features.xsl"?><xs:schema xmlns:xs="http://www.w3.org/2001/XMLSchema" xmlns:axis="http://www.

<xs:include schemaLocation="types.xsd"/>

  <xs:complexType name="AbsoluteBrightness">
    <xs:annotation>
      <xs:documentation>
        <AddedDate>2015-10-01</AddedDate> <!-- Before PIA imported data from AFI -->
        <Groups>
          <Group type="FunctionTeam">PT PTZ</Group>
          <Group type="FeatureType">Firmware</Group>
        </Groups>
        <Producer>FeatureDetection</Producer>
        <Description>Control image brightness using absolute coordinates in combination with pan/tilt/zoom moves</Description>
        <Mappings>
          <Mapping name="Supported" type="ParameterValue"><![CDATA[root.PTZ.Support.*.AbsoluteBrightness == true]]></Mapping>
        </Mappings>
      </xs:documentation>
    </xs:annotation>
    <xs:complexContent>
      <xs:extension base="axis:Feature">
        <xs:anyAttribute processContents="lax"/>
      </xs:extension>
    </xs:complexContent>
  </xs:complexType>
...
```

Fig. 2 Example snippet of boundary artefact 1: XML-based feature catalogue describing features across products in the software product line

```
rule_onvif_profile_s_videostream:
  description:
     "If Videostream is supported in the product, then ONVIF Profile S should be supported"
  statement:
    precondition:
   - Videostream/Supported:true
 condition:
   - ONVIF/Parameters/Profiles/Profile:S
  exclude_products:
    - Developer Board
 - Network Video Recorder

rule_https:
  description:
    "HTTPS should be supported"
  statement:
    precondition:
   - None
 condition:
   - HTTPS/Supported:true
  exclude_products:
    - None

...
```

Fig. 3 Example snippet of boundary artefact 2: YAML-based rule set expressing logical relationships between product features in the software product line

3.2.3 Questionnaire

The questionnaire was used as a reflective complement to the workshop discussions, allowing us to capture individual practitioner perceptions rather than to produce statistically generalisable results. We gathered their perceptions of the guideline based on the following rubrics (details of the workshop dynamics can be found in the Supplementary Material):

- Contribution to reliability, functionality, and predictability - This rubric assesses the extent to which practitioners perceive that the guideline supports the creation of artefacts that are reliable, predictable, and functional. As explained in the Background section, these three dimensions correspond to the trust beliefs identified by Lansing and Sunyaev (2016). We used these beliefs as the basis for formulating the related questionnaire questions (see Table 3)
- Ease of use, usefulness, and intention to use - This criterion examines the guideline based on practitioners' perception of its ease of use (the extent to which it is free of effort), its usefulness (the degree to which it effectively achieves their intended objectives), and their intention to use the proposed guideline (Moody 2003).

At the end of the second workshop, the first author distributed a questionnaire (see Table 3) to all participants from both workshops. In the questionnaire, participants assessed how they perceived each principle's contribution to the reliability, predictability, and functionality of the boundary artefacts. They also evaluated the guideline's usefulness, ease of use, and intention to use. A total of five practitioners completed the questionnaire. Participants were asked to indicate the extent to which they agreed with the questionnaire statements using a 5-point Likert scale, except for the 'Intention to Use' rubric, which was an open-ended question. We analysed the responses provided by the practitioners qualitatively, grouping each

Table 3 Questions asked in the questionnaire

Rubrics	Questions
Predictability	To what extent do these principles contribute to the content always being available?
Reliability	To what extent do these principles contribute to the correctness of content in an artefact?
Functionality	To what extent do these principles contribute to the proper functionality of the artefact?
Ease of use	Please indicate the degree you believe the guideline are easy to follow.
Usefulness	Please indicate the degree you believe the guideline are useful.
Intention to use	How likely are you intended to use the guideline? Under which conditions would you use it?

principle into three categories: strong positive contribution, moderate/mixed contribution, and low or no contribution (see our analysis in Section 4.3).

3.3 Step 3 - Guideline Improvement

We collected feedback on the guideline during both workshops and made adjustments to its design. The feedback focused on the relevance of each principle, the overall structure, and the clarity of the content. After both workshops, we refined our guideline, which we detail in Section 4.2.

3.4 Step 4 - Post-Factum Analysis of the Guideline Implementation

In the final step, we analysed how the guideline was implemented in practice. This process began with practitioners independently utilising the guideline over a period of one year and eight months. Following this implementation phase, we conducted a group interview with two on-site participants and a virtual interview with one to gain insights into practitioners' experiences and perceptions regarding the effectiveness and applicability of the guidelines in their work. This evaluation aimed to assess the practical impact of the guideline and identify areas for further development (see Table 4).

We interviewed two practitioners for 1 hour and 30 minutes. One was involved in the first two workshops and is now applying the guideline in a different context (cybersecurity). The other practitioner recently started using the guideline. During this group interview, they explained the challenges they faced when applying the guideline independently, as well as how it has been adopted as a formal process in the creation of new features. Additionally, we conducted an online interview with a third practitioner who was unable to be on-site during our visit. This practitioner had participated in both workshops. We presented the results of this post-factum analysis in Section 4.4.

3.5 Ethical Concerns

Our ethical approach focused on three main aspects: data collection, data analysis, and the handling of the company's confidential information. We designed and executed our study in accordance with the ethical research guidelines provided in the Swedish Research Council (2017).

Table 4 Question guide for the post-factum analysis

Rubrics	Questions
Evolution	How has the implementation evolved since the development of the schema template?
Implementation	Which parts of the guideline were difficult or challenging to follow? What were the main obstacles (e.g., technical, operational, cultural)?
	Did you feel that there was adequate preparation for implementing the guideline? Have you missed any preparation?
Usefulness and Impact	Which parts of the guideline did you find most useful or relevant to your work?
	Were there any categories you felt were not applicable or that you needed to modify to fit your context?
	Has using these guideline changed how you approach certain tasks or processes? In what ways?
Unexpected Outcomes	Did any unexpected results or consequences arise from following the guideline (either positive or negative)?
General Feedback	Overall, how satisfied are you with the outcomes after implementing the guideline?
	Are there any specific elements that could be improved or added to the guideline?

During the data collection phase, we carefully crafted the workshop questions to avoid eliciting intense emotions or causing emotional harm (Allmark et al. 2009). The complete list of workshop questions and the workshop agenda are available in the Supplementary Material. We also communicated the study's goals and informed participants of their right to withdraw at any time. Additionally, we explained how the data would be used and how anonymisation would be ensured. These procedures are described in the informed consent form and in the formal agreement with the company (see Supplementary Material for the full informed consent text and the questionnaire items).

In the data analysis phase, we anonymized quotes from practitioners and any confidential information related to the boundary artefacts utilized in the workshops. When reporting the results, we paid close attention to avoid stigmatizing or harming specific populations. Furthermore, a company employee reviewed our manuscript to ensure proper handling of confidential information prior to submission.

4 Results

In this section, we present the results of our study. They are organized into three subsections. First, we outline the organisation of the guideline into categories and principles (RQ1). Second, we provide the practitioners' perception of the guideline (RQ2). Finally, we summarize the outcomes of applying the guideline after the workshops and a post-factum analysis (RQ3).

4.1 The Organisation of the Guideline and Feedback from the Workshops

The version of the guideline presented in this article is the result of iterative adjustments made after both workshops. The original guideline derived from the literature underwent several refinements based on practitioners' feedback, which led to reorganising principles, clarifying descriptions, and improving terminology. Table 5 presents the final structure of the guideline. The subsequent paragraphs detail the revisions introduced in response to practitioner feedback and how these shaped the final version of the guideline.

The categories for the guideline emerged from our analysis as we collected principles from the literature. During this process, we identified similarities among the principles, which naturally led us to group them into categories, which we detail in Table 5. For example, principles 2a and 2b 'Decide on the format of the artefact' and 'Establish the means of spreading the artefact' are closely related to decisions about content arrangement, layout, and the most effective channels for distribution. As a result, principles related to content manipulation were grouped under the 'structure' category.

The first version included four categories (shared understanding, scope, structure, and management). The category 'shared understanding,' contained a single principle: 'Define an approach to address different interpretations of terminology.' Based on feedback from practitioners during the first workshop, we moved this principle to the 'Structure' category due to its close connection to that theme. When designing the boundary artefact, understanding stakeholders' needs can provide valuable insights into how terminology may vary across contexts. Furthermore, based on the same feedback, we also moved principle 3c 'Agree upon a balanced formality,' from the 'Structure' category to the 'Management' category, as it aligned more closely with the scope of the latter.

Feedback from the second workshop focused primarily on improving the clarity and usability of the guidelines. Practitioners who had participated in the first workshop were able to provide detailed input on how to make the principles more accessible to those unfamiliar with boundary artefacts. Revisions included simplifying overly academic terminology (particularly in the description of ownership), improving consistency in writing style, and using active verbs to begin each principle. These changes aimed to ensure that the guideline communicated its intent clearly to practitioners in different roles.

Beyond these refinements, it is important to highlight that the primary contribution of this work lies in the systematic consolidation of principles into a coherent, design-oriented guideline for creating and managing boundary artefacts. While the workshops were conducted within a single company, this organisational context serves primarily to demonstrate how the guideline can be applied in practice. The industry setting thus functions as an example of use, illustrating how practitioners interpret, adapt, and implement the principles, thereby showing the practical value and applicability of the guideline. The guideline itself, however, is grounded in the synthesis of findings from diverse empirical studies and represents a contribution that extends beyond the specific organisational case. In the next section, we present the refined version of the guideline, which incorporates all feedback from the workshops while maintaining its intended generalisability.

Table 5 The BA-Guide: Guideline for creating boundary artefacts in software engineering. Each principle is labelled using a category–letter notation (e.g., 2a, 2b). The first number refers to the category (1 = Scope, 2 = Structure, 3 = Management)

Category	Principles	Description
1. Scope	1a. Identify the stakeholders and the boundaries	Identify the stakeholders who will use, depend on, or contribute to the artefact, along with their information needs (Chow and Leiringer 2014; Zurba et al. 2019; Ouriques et al. 2023a). Boundaries between groups, such as teams, departments, organisational units, or domains are often permeable spaces where terminology and interpretations may differ. Mapping stakeholders first, and then identifying the boundaries among them, helps to reveal these differences. As part of this initial scoping, define the high-level purpose of the artefact and assign it a clear name or identifier to distinguish it from other artefacts with related goals.
	1b. Define an approach to deal with different interpretations of terminology	Determine how terminology will be interpreted and communicated across boundaries (Ouriques et al. 2023a; Zurba et al. 2019; Jain et al. 2014). This may involve including a glossary, documenting variations in meaning across groups, or providing supplementary material to help stakeholders understand key terms. Clear approaches for managing terminology reduce misinterpretation and support shared understanding.
2. Structure	2a. Decide on the format of the artefact	Define how the content will be represented and structured within the artefact (Zurba et al. 2019). The chosen format, such as text-based documentation, diagrams, XML, JSON, or mixed formats affects how stakeholders access, interpret, and use the artefact. The selected format should support users' informational needs and enable efficient retrieval and comprehension of content.
	2b. Establish the means of spreading the artefact	Determine how the artefact will be disseminated across boundaries and accessed by stakeholders (Zurba et al. 2019). Appropriate dissemination mechanisms may include documentation portals (e.g., Confluence), Git repositories, web interfaces, internal wikis, or shared whiteboards. The dissemination strategy should ensure that the artefact is easy to locate, retrieve, and use.
	2c. Decide an appropriate level of content detail	Determine how much detail the artefact should contain. The content should include enough information for stakeholders to find what they need without becoming overly detailed or difficult to navigate. An excessive level of detail can decrease usability, while insufficient detail may lead stakeholders to rely on alternative representations. Such workarounds can result in duplicated or scattered information. Wohlrab et al. (2018) report similar challenges, where inappropriate detail contributed to fragmentation and inconsistency in systems engineering artefacts.
	2d. Accommodate additional and experimental content	Ensure the artefact can incorporate new or evolving content as software development progresses (Van Loggem and Van Der Veer 2014; Ouriques et al. 2023a). Development activities often involve provisional or experimental information. Artefacts should therefore allow for temporary additions and enable this content to be refined, formalised, or removed without compromising coherence or stability.
3. Management	3a. Set up an evaluation and feedback process	Establish a recurring process to assess whether the artefact continues to meet stakeholders' needs and to gather feedback systematically (Chow and Leiringer 2014; Zurba et al. 2019; Jain et al. 2014; Ouriques et al. 2023a). Evaluation may include automated checks, review meetings, surveys, or analysis of usage patterns. Regular evaluation helps identify outdated, missing, or inconsistent content and ensures that improvements are incorporated in a timely manner.

Table 5 (continued)

Category	Principles	Description
	3b. Establish ownership of the artefact	Assign responsibility for maintaining the artefact and coordinating changes (Zurba et al. 2019; Van Loggem and Van Der Veer 2014). Clear ownership ensures that updates are managed consistently, content remains accurate, and gaps or inconsistencies are addressed promptly. Ownership may reside with a single role or be distributed across contributors depending on the artefact's scope.
	3c. Agree upon a balanced formality	Determine the level of governance required to manage changes to the artefact (Zurba et al. 2019). This includes defining approval processes, review steps, and criteria for incorporating new content. The degree of formality should safeguard quality and accountability while avoiding excessive overhead that could hinder timely updates.
	3d. Handle scattered information	Prevent important information from becoming duplicated or dispersed across multiple tools or documents (Wohlrab et al. 2018). When stakeholders create alternative sources due to missing or unclear content, information becomes fragmented. Monitoring for scattered information and consolidating content within the artefact helps maintain coherence and reduces the risk of parallel artefacts emerging.

4.2 The BA-Guide: Guideline for Creating Boundary Artefacts in Software Engineering

The BA-Guide (see Table 5) provides ten principles supporting the creation of boundary artefacts in software engineering. We describe the principles and supplement them with examples of implementation. We organise the guideline into scope, structure, and management categories.

4.2.1 1 - Scope

The scope category displays the principles referring to the definition of the target audience in each boundary, their needs regarding the artefact, and how different terminology is held across boundaries.

4.2.2 2 - Structure

The principles in this category relate to how the artefact's content is structured to meet stakeholders' needs regarding format, level of detail, and formality.

4.2.3 3 - Management

The category refers to principles that can guide the establishment of practices to manage the artefact throughout time.

4.3 Practitioners' Perception of the Guideline

After conducting the workshops, we distributed a questionnaire to collect practitioners' perceptions of how the guideline's principles contributed to creating functional, reliable, and

predictable boundary artefacts (RQ2). Given the small number of participants, our analysis focused on qualitative patterns rather than statistical interpretation.

To analyse the responses, we grouped each principle into three categories: strong positive contribution, moderate/mixed contribution, and low or no contribution, based on the Likert responses. Principles were classified as "strong" when the majority of responses (three or more) fell in the two highest agreement categories (to a moderate extent or to a large extent). Principles with responses distributed across positive, neutral, and negative categories were classified as "moderate/mixed". Principles for which most responses fell in the lowest categories (not at all or to a small extent) were classified as "low or no contribution". This approach allowed us to identify tendencies without implying statistical significance, which would not be appropriate given the small sample size.

Using these qualitative categories, we identified patterns in practitioners' perceptions across the three trust-related rubrics. In Table 6, we report the practitioners' perception for each rubric (see Appendix E for the distribution of the responses).

Functionality is primarily supported by Scope and Management. Practitioners assessed the extent to which the guideline principles support the proper functionality of their artefact. The responses indicate that functional performance relies primarily on the principles associated with Scope, which establish who the artefact is for and how its content should be interpreted, and Management, which ensures the artefact is maintained and governed over time.

Table 6 Perceived contribution of guideline principles to functionality, reliability, and predictability

Principle	Functionality	Reliability	Predictability
1a. Identify the stakeholders and the boundaries	Strong	Low	Low
1b. Define an approach to deal with different interpretations of terminology	Strong	Strong	Low
2a. Decide on the format of the artefact	Strong	Strong	Low
2b. Establish the means of spreading the artefact	Low	Low	Strong
2c. Decide an appropriate level of content detail	Low	Strong	Low
2d. Accommodate additional and experimental content	Low	Strong	Low
3a. Set up an evaluation and feedback process	Strong	Strong	Strong
3b. Establish ownership of the artefact	Strong	Strong	Strong
3c. Agree upon a balanced formality	Strong	Strong	Low
3d. Handle scattered information	Low	Low	Low

The Scope principles 1a (Identify stakeholders and boundaries) and 1b (Terminology interpretation) were considered strong in creating artefacts that stakeholders can understand and use effectively. They emphasise the importance of identifying who the artefact serves and how concepts are interpreted across boundaries. P6 mentioned "*It is unclear to the developers that they are not the primary consumers of the artefact. I think it could be made clear to the stakeholders who the stakeholders are. We are the people who know the stakeholders, but we haven't made it clear to everybody else. We think it is really important.*"

When stakeholders and boundaries are not explicitly identified, artefacts tend to reflect implicit assumptions about who the primary users are. This can result in missing information for peripheral groups or misaligned expectations across roles. In such cases, stakeholders may struggle to interpret the artefact correctly or may use it in unintended ways, which weakens its function as a shared coordination mechanism. Explicitly mapping stakeholders and terminology therefore reduces ambiguity and strengthens the artefact's ability to function across boundaries.

The Management principles 3a (Evaluation), 3b (Ownership), and 3c (Balanced formality) were also strongly tied to operational effectiveness. Their shared contribution lies in ensuring that the artefact is not only produced correctly but also governed in a way that supports everyday use. Without governance mechanisms, artefacts may drift from their intended purpose as new content is added or responsibilities shift. This can lead to gradual misalignment between stakeholder needs and artefact content. Governance principles therefore act as stabilising mechanisms that maintain coherence over time, ensuring that the artefact remains usable in daily coordination rather than becoming outdated or fragmented. Evaluation ensures an artefact remains aligned with stakeholders' needs; ownership ensures accountability; and balanced formality ensures that change procedures do not obstruct the artefact's functional role. Regarding the level of content detail, P5 mentioned: "*if you add more information and suddenly you might have more stakeholders, and if you have more stakeholders, there will be more information. In some cases, less information is needed. In other, more information is needed*"

Together, these principles explain practitioners' overall perception, which suggests that the functionality of the artefact analysed, besides its structure, is supported by how well it captures shared meaning and how effectively it is maintained.

Reliability is primarily supported by content Structure and Management. Practitioners indicated the extent to which each principle contributes to the correctness of the artefact's content. Their responses showed that content correctness is supported primarily by the Structure principles, which indicate how information is organised and detailed, and by the Management principles, which govern verification and maintenance over time.

Within the Structure category, principles 2a (Decide on the format of the artefact), 2c (Decide an appropriate level of content detail), and 2d (Accommodate additional and experimental content) were perceived as essential for ensuring the accuracy and integrity of the artefact's content. An appropriate format can improve clarity, adequate detail reduces ambiguity, and the ability to accommodate evolving or experimental content can help maintain correctness as stakeholders needs change.

When format or level of detail is poorly defined, stakeholders may misinterpret information or rely on external clarifications. Overly rigid artefacts may also fail to incorporate evolving requirements, causing discrepancies between practice and documentation. These breakdowns directly undermine perceptions of correctness. By structuring content appropri-

ately and allowing controlled evolution, the artefact can maintain internal consistency and better support trust in its reliability.

The Management principles also played a central role. 3a (Set up an evaluation and feedback process) was seen as crucial for identifying and correcting inaccuracies. P6 stated "*we do not have any of those. It is good that we introduce regular checks*" 3b (Establish ownership of the artefact) ensures that responsibility for content quality is clearly defined, and 3c (Agree upon a balanced formality) helps maintain appropriate control over changes.

Without a recurring evaluation process, outdated or inconsistent content may remain unnoticed, leading stakeholders to question the accuracy of the artefact. In cross-boundary settings, such inaccuracies can propagate quickly, as different groups rely on the artefact as a shared reference point. When errors are not systematically detected and corrected, trust in the artefact's reliability erodes, and stakeholders may develop alternative sources of information, undermining its coordinating function.

In contrast, the Scope principles were perceived as less relevant for reliability: although identifying stakeholders and boundaries provides important contextual understanding, practitioners did not see these aspects as strongly influencing content correctness. Likewise, 2b (Establish the means of spreading the artefact) was not viewed as contributing to reliability. This suggests that dissemination mechanisms primarily affect accessibility rather than the accuracy of the artefact's content.

Predictability is primarily supported by dissemination and maintenance-related principles. Practitioners indicated the extent to which each principle contributes to the artefact's content always being available. Their responses revealed the most selective pattern, with only a small subset of principles meaningfully supporting the continuous availability of the content of the analysed artefact.

The strongest contributors were 2b (Establish the means of spreading the artefact), 3a (Set up an evaluation and feedback process), and 3b (Establish ownership of the artefact). These principles all relate to how the artefact is distributed and maintained over time. Predictability depends less on what the artefact contains and more on whether stakeholders can consistently access updated content. If dissemination channels are unclear, or ownership and evaluation mechanisms are absent, updates may not propagate across boundaries. This creates uncertainty about whether the artefact reflects the current state of work. By ensuring clear distribution paths and maintenance responsibilities, these principles support confidence that the artefact will remain available and up to date.

Principle 2b ensures that the artefact can reach stakeholders across boundaries. During our first workshop, P1 referred to their practice regarding this principle "*It is not good, we do not get updates. It could be better. The format is good but it requires some technical competence*"; principle 3a supports predictability by enabling regular updates and evaluations; and principle 3b provides clear responsibility for sustaining the artefact in practice. Together, they align directly with the idea of predictability as the ongoing provisioning and maintenance of content.

The remaining Scope and Structure principles were considered less relevant for predictability. While they may enhance clarity or correctness, practitioners noted that these qualities do not, on their own, guarantee that the content remains consistently available.

Overall, two principles emerged as consistently strong contributors: establishing ownership and setting up an evaluation. In contrast, principle 3c (Handle scattered information) was not seen as contributing to any of the rubrics analyzed during the workshops. During

discussions in the first workshop, practitioners noted that if principle 1a (Identify the stakeholders and the boundaries) is properly addressed, scattered information is unlikely to occur. In the second workshop, they did not consider this principle when creating the new artefact because the new artefact had not yet been implemented in practice within the organisation, and therefore scattered information was not yet an issue.

Practitioners were also asked about the guideline's ease of use, usefulness, and their intention to use it. Responses regarding **ease of use** showed no clear consensus. Ratings were dispersed across "to a small extent (2)," "neutral(1)," and "to a moderate extent(2)," indicating mixed perceptions (see Supplementary Material for results).

This dispersion reflects a challenge already observed during the workshops, where practitioners found the concept of boundary artefacts itself difficult to grasp. The abstraction of the term and the complexity of the written language used in the guideline appeared to hinder ease of adoption. Thus, ease of use appears to be influenced less by the structure of the guideline and more by familiarity with the underlying concept. This is illustrated by P3 "*Given the academic language, people might find it hard to get it. It would take some time to understand what a boundary artefact is. It needs to be easily written. It's good to have a concrete example; it is much easier for me. Too abstract is not very good*".

Practitioners were also asked to assess the degree to which they believed the guideline was useful. In contrast to the mixed views on ease of use, perceptions of **usefulness** were generally strongly positive. Most participants rated the guideline as useful "to a large extent (3)," indicating that they see value in having a structured approach to creating boundary artefacts.

Participants highlighted that the guideline brings consistency, clarity, and structure to artefact creation, which can otherwise vary widely across teams or domains. Several practitioners emphasised that using a shared guideline or template would make it easier for stakeholders to use artefacts and would promote alignment across organisational boundaries. This reinforces the notion that usefulness is closely tied to the guideline's ability to support coordination, reduce variability, and provide shared expectations. On this, P1 commented "*what I would like to be able to do with the guidelines is, for example, if you have a boundary artefact, these are the things you should look at. Everybody has to look to scope, structure…but everybody does not have to necessarily go to every principles below or even interpret exact the same description, and then they would look different when applying.*"

Regarding the **intention to us**e the guideline, practitioners expressed a positive intention to apply it to both artefacts used in the workshops. As P1 noted, "*one benefit of having some guideline like this is that if many boundary artefacts follow the same guideline or template, that will make it easier for people to consume the artefact. That makes me more convinced that we should use this.*" This reflects the perceived value of consistency across artefacts and its influence on practitioners' intention to adopt the guideline.

These findings offer important initial insights into how boundary artefacts can be created and supported through a guideline. Because the study was designed to explore practitioners' perceptions rather than produce statistically significant results, our focus was on identifying directional tendencies in how the guideline contributes to reliability, predictability, and functionality. We acknowledge that the small sample size of five respondents limits the generalisability of the findings or further statistical analysis. Consequently, while the results may offer valuable preliminary observations, they should be interpreted with caution and not considered as representative of the broader practitioner community.

4.4 Application of Guideline and Post-Factum Analysis

In some parts of Axis, practitioners use schemas (i.e., common definitions of structured data in a computer-readable format) to facilitate the automation of information flow, which is also the case for the boundary artefacts referenced in this study. After our last workshop, the practitioners developed a preliminary version of a template based on our guideline, designed to support the creation of schema documentation (see Fig. 4).

A1 (a schema) was originally poorly documented when it was created. Several issues stemming from this inadequate documentation were revealed during the workshop. Practitioners emphasised the need for clearer documentation, particularly regarding stakeholders, terminology, and use cases, which were missing and had led to reduced trust in the artefact's content. They expected that improving these aspects would help new schemas, such as A2, avoid some of the trust-related problems experienced with A1. The template was therefore developed to help schema creators identify a minimum set of questions that should be answered and documented, ensuring consistency across schemas.

To support the creation of new schemas and prevent issues similar to those observed with A1, practitioners used our guideline to structure the documentation template (see Fig. 4). Five principles guided the design of the template's components: ownership (3b), stakeholder identification (1a), dissemination (2b), format decisions (2a), and terminology management (1b). By mapping these principles to corresponding template sections, practitioners established a minimum set of information that schema creators must provide. This alignment aims to improve transparency, trust, and traceability across schemas. It also demonstrates

Schema Documentation Template		Related Principles in the Guideline
Owner Question (s) to be answered: - *Who are the owners?* - *Is it the same owner of the content and code?*	→	3b. Establish ownership of the artefact
Stakeholders Question (s) to be answered: - *Who are the stakeholders?*	→	1a.Identify the stakeholders and the boundaries
Location Questions to be answered: - *Where is the schema file located (e.g. git?)* - *Is there more information to be found elsewhere (e.g. confluence space)?*	→	2b. Establish the means of spreading the artefact
Format Question (s) to be answered: - *What is the format?* - *Why are the elements/instances in the schema formatted as they are?*	→	2a.Decide on the format of the artefact
Terminology Questions to be answered: - *What is the long name and short name of the schema?* - *Are there terms that need to be explained?* - *What guidelines regarding the terminology will the schema follow?* - *How/where will the terminology used in the schema be explained?*	→	1b.Define an approach to deal with different interpretations of terminology
Use cases Question (s) to be answered: - *What are the use cases (i.e., why is the schema needed?)*	→	1a.Identify the stakeholders and the boundaries

Fig. 4 Schema documentation template derived from the guideline

Table 7 Overview of boundary artefacts and their corresponding statuses over time

Period	Boundary artefact	Status
March 2023	A1	Assessment of current practices mapped against the guidelines
	A2	Creation and implementation of the template
1 year and 8 months: The period during which the guideline was used independently.		
November 2024	A1	The template adopted as the official process for creating new features.
	A2	Template in operation
	A3 - JSON file containing instructions for standards and security testing.	Implementation phase of the template for security testing and standard compliance.
	A4 - Requirements specification	Currently working on a proof of concept.

how the guideline can be operationalised in practice, translating abstract principles into concrete documentation elements that support more reliable and predictable boundary artefacts.

The remaining principles (2c, 2d, 3a, 3c, and 3d) were not explicitly incorporated into the schema template at this stage. This does not imply that they were considered unimportant; rather, their relevance depends on the specific characteristics and maturity of the artefact under development. For example, principles related to evaluation processes, balanced formality, or handling scattered information become more salient once an artefact is in use. As the guideline is not intended to be prescriptive or uniformly applied, organisations are not expected to implement all principles for every artefact. Instead, the principles are meant to be interpreted and adapted according to contextual needs.

During the discussions for the creation of the template for A2, the practitioners identified immediate stakeholders, including people within the Quality Assurance, R&D, and Product Management departments. They also foresaw potential interest in A2 outside these departments, such as in Sales or Marketing. After the workshops, the research team left the company for a period of one year and eight months. At that point, the key outcomes were the creation of the template and the identification of problems in the original design of A1. The status in March 2023 is shown in Table 7.

The first author returned to the company in November 2024 and conducted interviews with three practitioners to observe what had happened since the last intervention and to assess their experiences with the guideline. Since then, the template for A1 has been adopted as the official process for creating new features. Additionally, the use of the guideline has expanded to include security testing and standards, such as A3 (see Table 7 for the current status). The company has adopted the guideline to create numerous schema descriptions that instruct on various standards. Each standard includes different information (such as owners, use cases, etc.), but all maintain a consistent backbone derived from the guideline. In relation to their data-driven automatic tests for security, the current discussion centers around the adoption of elements from the schema description, including content and technical ownership.

Lastly, there is ongoing discussion about using the template for requirements specification (A4) to standardize information across documents, though this is still in the early stages.

In summary, the application of our guideline has significantly expanded. It has become the official process for creating new features, extended to security testing and standards, and is now being used to create consistent schema descriptions across various domains. There are also discussions about incorporating elements from these schemas into their data-driven security tests.

5 Benefits and Challenges

Overall, our approach was well received by our company partner, particularly in contexts where boundary artefacts were long-lived and used across multiple roles. While the guideline has already been put into operation, some practitioners expressed the need for more time to fully utilise the developed templates before drawing final conclusions. As P1 noted, "*we need a longer period of use to really see the full effect.*" Even so, several benefits were identified during this period of use without further researcher intervention. For example, P3 described challenges emerging from the lack of ownership of the first artefact: "*since there were no clear owners, the things you had researched started to drift away, and maybe some features in boundary artefact 1 aren't really features because there was no clear guideline.*" This indicates that without explicit responsibility structures, artefacts may gradually lose coherence over time.

Practitioners also recognised that planning boundary artefacts is not an isolated activity. The design of a boundary artefact is a collaborative process involving people from different social contexts. As we brought together individuals who sometimes belonged to the same organisational boundaries, their diverse perspectives helped address the challenges of determining what information to represent and how to represent it. This collaborative negotiation reflects one of the core functions of boundary objects described by Star and Griesemer (1989), namely their ability to facilitate cooperation across different social worlds despite differing viewpoints.

During the workshop for A1, practitioners realised that several assumptions had been made while creating the artefact. When they compared their existing practices with the guideline, they identified missing decisions related to terminology handling, stakeholder identification, and management responsibilities. They concluded that mapping the guideline onto a new or existing boundary artefacts helps reveal preconceived assumptions. As P1 explained, "*you have to actually think about what you need to do, not just say, 'Oh, I'll code something or add features.' You have to ask yourself, 'Why am I doing this? Who will it benefit, and why is it valuable?'*"

Another benefit of the guideline was its support in visualising initial boundaries. In early phases of a boundary artefact creation, boundaries are not always immediately clear, and even when they are, managing them can be complex. Starting with a small set of boundaries and expanding them gradually as the artefact evolves helped shape how the guideline was implemented.

Practitioners also found that reviewing the information sources stakeholders used for their tasks improved identification of stakeholder needs. In the A1 workshop, some questioned whether to retain a particular piece of information because it lacked ownership and was outdated. Through discussion, they discovered that stakeholders had no alternative

access to this information. As a result, they decided to retain the content, assign an owner, and ensure it would be updated regularly.

A particularly important benefit, aligned with the fundamental purpose of boundary artefacts, was that the template made the boundary artefact content more accessible to stakeholders who do not work with code. This helped bridge a communication gap between development teams and non-technical actors. As P4 described, "*it makes it easier for people who want to access it, as everything is available in one place… When we discuss this with others, it's much easier to start with the description and talk around it rather than showing people code. Not everyone is comfortable discussing code.*"

Despite these benefits, practitioners also identified key challenges. One of the most significant was the difficulty of understanding the concept of "boundary" itself, even within a well-established artefact. P1 reflected, "*I underestimated the need to explain clearly to people who were unprepared. The problem is that it is deceptively simple, so people think they get it. But they don't fully get it. It takes time.*" Understanding the boundary concept is essential for developing shared understanding, which Star and Griesemer (1989) identify as a critical characteristic of effective boundary objects.

Finally, practitioners noted that creating a boundary artefact utilising the guideline is a time-consuming process. Developing the structure and reconciling diverse perspectives requires much more than a single meeting. Additional stakeholders often need to be involved before consensus is reached. Early discussions typically focused on brainstorming how to apply the guideline to the artefact, while later sessions gradually shaped concrete practices and decisions. This suggests that successful adoption depends on organisational willingness to allocate coordination time.

The reported benefits were observed in artefacts that were cross-functional and persistent over time. In settings where artefacts are short-lived or confined to a single team, the overhead of applying the guideline may outweigh its benefits.

6 Contribution to the Literature and Future Research

We followed a multi-step, longitudinal study embedded in an industrial context to develop and evaluate our guideline. Our work expands on previous research on boundary artefacts by providing a structured guideline organised into three main categories (Scope, Structure, and Management) and ten principles (see Table 5), which support practitioners in creating boundary artefacts in software engineering contexts. As the boundary artefact literature in software engineering contexts is scarce, this guideline is relevant because it offers an original perspective on design-oriented principles that practitioners should observe upon the conception of such artefacts.

In developing the guideline, we also incorporated practices identified in studies of product family development and automotive software engineering, where boundary artefacts play a central role in coordinating complex development work. For example, Jain et al. (2014) discuss practices for managing variability and shared understanding in product family contexts, while Wohlrab et al. (2018) present principles for handling information fragmentation and change in the automotive domain. Although these studies focus on specific industries, the practices they describe are extendable and were relevant to the context of our

case company. Their inclusion enriches the guideline and supports its potential applicability beyond a single organisational setting.

The current version of our guideline evolved based on practitioners' feedback after conducting two workshops. Besides rearranging principles into categories (see Section 3.3), the latest version differs from the first to meet practitioners' needs for clarity in describing each principle, especially regarding the concept of a boundary. During the workshops, practitioners agreed with the importance of the principles, although not all of them were implemented in this specific context.

Practitioners considered the **scope category** logical and easy to understand. However, during the workshops it became evident that, in this particular case company, existing practices did not explicitly address stakeholders' information needs or the interpretation of terminology across boundaries. Practitioners also appeared unfamiliar with how organisational boundaries could shape the meaning and use of boundary artefacts. As discussed in prior research, when boundary artefacts are not intentionally planned or managed, they tend to evolve organically as they are used by different groups (Henderson 1991; Star and Griesemer 1989; Zaitsev et al. 2016; Ouriques et al. 2023a). Under such circumstances, both in our case study and as described in the literature, stakeholders are not clearly identified, leading to neglected information needs. The Scope category therefore helped practitioners recognise these issues and emphasised the importance of establishing an appropriate level of plasticity that supports both local flexibility and robustness.

The **structure category** brings new insights into principles particularly relevant to software development contexts, such as accommodating additional and experimental content (Van Loggem and Van Der Veer 2014; Ouriques et al. 2023a). Practitioners already recognised some of the principles in this category, such as selecting an appropriate format and establishing the means of spreading the artefact, and noted how these considerations enhance the transfer of information across different social worlds. However, they were less familiar with the full range of stakeholders who rely on the artefact, often focusing primarily on immediate or local users. A further observation related to this category concerns the principle "Decide an appropriate level of content detail." Practitioners highlighted that explicitly addressing the level of detail can help prevent unnecessary duplication of content, as well as the creation of parallel artefacts by individuals or teams attempting to compensate for missing or insufficient information.

The **management category** had a surprising and positive reaction from practitioners in both the questionnaire and workshop, especially regarding the principle of "establish ownership of the artefact". Ownership has been explored in software engineering literature but primarily focuses on code (Maruping et al. 2009; Eldh and Murphy 2015; Zabardast et al. 2022; Nordberg 2003). We emphasise that this principle, aligned with evaluation and feedback, shapes a structured strategy for managing boundary artefacts and keeping them consistent (Star and Griesemer 1989). Besides, their implementation creates a favourable environment for stakeholders to feel confident about relying on boundary artefacts, as they trust the content is correct and updated.

Artefacts, in general, are recognised as necessary, yet practitioners often perceive them as burdensome to create and maintain (Stettina and Heijstek 2011). When an artefact is not actively maintained, trust in its content may decrease, and its usefulness in practice can be significantly reduced. In our study, practitioners indicated that the proposed guideline

helped support the creation of boundary artefacts that they perceived as more reliable, functional, and predictable, although the strength of this contribution varied across principles.

Across the three rubrics of reliability, functionality, and predictability, practitioners saw Management principles as central to both functionality and reliability, while Scope principles contributed primarily to functionality. Reliability, in contrast, was supported mainly by Structure and Management principles, which practitioners associated with organising, verifying, and sustaining accurate content. Predictability showed a more selective pattern, with only a few principles, particularly those related to dissemination and maintenance, such as establishing ownership and defining how the artefact is spread, being perceived as directly influencing the continuous availability of content. Although these tendencies cannot be generalised due to the small sample size, they offer useful insight into how practitioners interpret the roles of the guideline's principles. Such insights may help organisations decide which aspects of the guideline to prioritise when addressing specific concerns, such as improving the accuracy, functionality, or availability of an existing artefact.

It is important to emphasise that the guideline is not intended to prescribe a linear process nor require all principles to be implemented uniformly. Instead, each principle should be interpreted and adapted according to the specific context and purpose of the boundary artefact. The relative importance of principles and the practices selected to implement them will necessarily vary across cases.

Although the guideline contributes to the software engineering literature by consolidating and operationalising principles that were previously dispersed, it also raises several open questions. For example, would the same set of principles be equally relevant for different types of boundary artefacts or in other organisational settings? Would additional principles be required, or might some become less important? We consider that the current guideline will likely evolve as it is applied in new contexts. Further applications may reveal the need to refine, expand, or adjust the principles.

Building on our initial questions, the perception of how the guideline's principles contribute to boundary artefacts being more reliable, functional, and predictable may also vary. But how do they differ? Would principles like 3d, 'Handle scattered information,' yield different results, for example? We suspect that they would. Therefore, future studies could focus on identifying how the contribution of principles to trust changes when the type of boundary artefact changes. The findings could provide valuable insights into which principles are crucial for maintaining consistency across different boundary artefacts.

Relatedly, practitioners' perceptions of how individual principles contribute to reliability, functionality, and predictability may also change depending on the type of artefact being considered. For instance, a principle such as 3d "Handle scattered information" may play a more prominent role in artefacts that naturally accumulate content from multiple sources. Future work should therefore examine how the perceived contribution of different principles varies across artefact types, which may provide clearer insight into which principles are universally valuable and which are context dependent.

Taken together, the findings suggest that the guideline principles operate not merely as descriptive recommendations but as mechanisms addressing recurrent coordination breakdowns. When principles related to scope are absent, artefacts risk failing to reflect stakeholder needs; when structural principles are neglected, ambiguity and inconsistency arise; and when management principles are missing, artefacts deteriorate over time. In each case, the consequence is reduced trust in the artefact as a cross-boundary mechanism. The prin-

ciples therefore contribute to trust by mitigating specific failure modes in boundary artefact creation and governance.

7 Limitations and Threats to Validity

As with any study embedded in a single industrial case, this research has limitations. Evaluating the guideline through a single case restricts the generalisability of our findings. While the guideline is grounded in literature from multiple domains, practitioners' perceptions of how each principle contributes to reliability, functionality, and predictability may vary depending on the type and use of a boundary artefact.

Our goal was to propose a flexible guideline rather than a universally prescriptive approach, recognising that different boundary artefacts require distinct strategies. We also acknowledge that the guideline is time-bound, as it reflects the state of the literature at the time of the study. Future work may refine or expand the set of principles as they are applied in other contexts.

We discuss threats to validity following Wohlin's classification of construct, internal, external, and conclusion validity (Wohlin et al. 2012).

Construct validity concerns whether the theoretical constructs were interpreted correctly by participants. A potential threat arose from practitioners' varying levels of familiarity with concepts such as boundary artefacts, reliability, and predictability. To minimize this threat, during workshops, the researchers clarified these concepts to reduce misinterpretation. However, some practitioners still found the notion of boundary artefacts difficult to apply in practice, as described in Section 5.

A further construct threat relates to the design of the questionnaire and interviews. With only five questionnaire respondents, there is a risk that individual interpretations of the questions may have influenced the results. This risk was mitigated by the fact that the questionnaire was administered after two workshops in which each trust belief (reliability, functionality, predictability) and each guideline principle was extensively discussed with the practitioners. This prior shared grounding reduced the likelihood that participants interpreted the questionnaire items inconsistently. Even so, subtle variations in how respondents understood these concepts may still have introduced bias. To minimise this threat, the questionnaire items were reviewed by all authors to ensure clarity and close alignment with the constructs defined in the Background section.

Interviews also carry threats related to researcher influence and social desirability bias. Participants may have responded in ways they believed were expected or favourable to the research team, particularly given the collaborative relationship established during the workshops. To mitigate this, interviews were semi-structured, included open reflective questions, and were conducted using neutral phrasing to avoid leading responses.

Internal validity relates to factors other than the intended ones that may have influenced the empirical results. One potential threat concerns researcher involvement in the data collection process. The first author facilitated the workshops and later conducted the follow-up interviews. This dual role may have introduced researcher influence or confirmation bias, and participants may have responded in ways they perceived as aligned with the study's goals. To mitigate this risk, workshops followed a structured format, interviews were semi-

structured and open-ended, all sessions were audio recorded, and interpretations were reviewed collaboratively by multiple authors.

Another threat relates to participant selection. Practitioners were selected by the company based on seniority and decision-making authority regarding boundary artefacts. On the one hand this ensured informed perspectives, on the other it may have resulted in viewpoints that do not fully represent all stakeholder groups. A further internal validity threat concerns the process of collecting principles from the literature. To minimise selective bias, we deliberately included studies from multiple domains and did not limit ourselves to software engineering. The intention was to ensure that potentially relevant practices for creating and managing boundary artefacts were not excluded.

However, as the literature review was not conducted as a full systematic review, there remains a possibility that relevant studies were overlooked or that the synthesis of principles reflects subjective interpretation. To mitigate this risk, the extracted principles and their categorisation were reviewed and discussed among multiple authors.

External validity concerns generalisability. While we incorporated knowledge from diverse fields, the results of applying the guideline, particularly practitioners' perceptions, cannot be generalised due to the single case design. Nonetheless, generalisability was not the primary aim; rather, our intention was to propose a guideline that can be incrementally adapted and extended.

Conclusion validity addresses whether the conclusions drawn from the empirical data are adequately supported by the evidence. A potential threat arises from the relatively small number of participants involved in the workshops, questionnaire (five respondents), and follow-up interviews (three practitioners), which limits the robustness of patterns identified.

In addition, the qualitative categorisation of questionnaire responses and interpretation of workshop discussions may involve analytical judgement. To mitigate this risk, we triangulated data across multiple empirical sources (workshops, questionnaire responses, and post-factum interviews). Questionnaire responses were analysed systematically and grouped into qualitative categories. Workshop and interview sessions were audio recorded to ensure accurate interpretation. Furthermore, multiple authors were involved in reviewing and discussing the analysis, which helped reduce individual bias in interpreting the findings.

8 Conclusion

Boundary artefacts are essential in software development as they help bridge multiple boundaries within an organization, delivering relevant information to various stakeholders. However, they are doomed to fail if they contain significant inconsistencies, causing stakeholders to avoid using them. These inconsistencies can also negatively affect software projects, such as providing misleading information that leads to incorrect task execution or increasing lead times as people waste time searching for reliable data.

We developed a guideline for creating boundary artefacts. We then tested and evaluated this guideline through a longitudinal study, gathering practitioners' perceptions on how each principle contributes to the creation of reliable, functional, and predictable artefacts. The principles identified in our literature review were grouped into three categories. The first category, Scope, focuses on defining each boundary's target audience, needs, and terminology. The second category, Structure, addresses how the artefact's content should be

organized to meet stakeholder needs. The third category, Management, refers to principles for establishing practices to manage the artefact over time.

During our evaluation, the practitioners perceived the principles as contributing to the creation of boundary artefacts, as well as to the relevance of the guideline and its usefulness. In the post-factum analysis, we discovered that the template derived from our guideline had been adopted as a formal process for developing content for A1. Additionally, it is in the process of being adopted for two other artefacts. However, these positive outcomes are accompanied by challenges. The main challenges we identified were: (1) cultural barriers related to the adoption of a new process, (2) difficulties in understanding the concept of boundary artefacts; and (3) the time-consuming nature of the activity, which often requires collaboration among people from different boundaries. This process demands time dedication from various individuals and support from management.

Our guideline connects fundamental traits of boundary artefacts, such as shared understanding, plasticity, and the ability to transfer information across boundaries. Practitioners can utilise the guideline to guide the creation or even evaluate current practices for existing artefacts. As the principles contribute differently to reliability, functionality, and predictability in our specific case, we believe that future application of the guideline with different types of artefacts, such as manual or automated, could provide valuable insights into how the perception about the principles vary depending on the type of boundary artefact.

Supplementary Information The online version contains supplementary material available at https://doi.org/10.1007/s10664-026-10847-x.

Acknowledgements We want to thank all practitioners at Axis Communications who participated and contributed to this study. We would also like to acknowledge that this work was supported by the KKS foundation through the S.E.R.T. Research Profile project at Blekinge Institute of Technology.

Author Contributions Conceptualization: Raquel Ouriques; Data curation: Raquel Ouriques, Baldvin G. Bern, Victoria Vucic; Funding acquisition: Tony Gorschek; Writing – original draft preparation: Raquel Ouriques; Writing – review and editing: Raquel Ouriques, Fabian Fagerholm, Daniel Mendez, Tony Gorschek, Baldvin G. Bern, Victoria Vucic; Methodology: Raquel Ouriques, Fabian Fagerholm, Daniel Mendez; Formal analysis and investigation: Raquel Ouriques.

Funding Open access funding provided by IT University. This work is supported by Stiftelsen för Kunskaps- och Kompetensutveckling.

Data Availability Raw data from the workshops are not publicly available in order to preserve individuals' privacy in accordance with the European General Data Protection Regulation, and the artefacts mentioned in this study are proprietary to Axis Communications.

Declarations

Ethical Approval Not applicable.

Informed Consent The collection, use, and handling of the data gathered from practitioners were conducted with informed consent, as outlined in the Supplemental Material.

Conflicts of Interest The authors declare that they have no conflict of interest.

Clinical Trial Number Not applicable.

References

Akoumianakis D, Vidakis N, Vellis G, Kotsalis D, Milolidakis G, Plemenos A, Akrivos A, Stefanakis D (2011) Transformable boundary artifacts for knowledge-based work in cross-organization virtual communities spaces. Intell Decision Technol 5(1):65–82

Allmark P, Boote J, Chambers E, Clarke A, McDonnell A, Thompson A, Tod AM (2009) Ethical issues in the use of in-depth interviews: Literature review and discussion. Res Ethics 5(2):48–54

Barrett M, Oborn E (2010) Boundary object use in cross-cultural software development teams. Human Relations 63(8):1199–1221

Blomqvist K (1997) The many faces of trust. Scand J Manag 13(3):271–286

Chow V, Leiringer R (2014) The translation of power: A study of boundary objects in public engagement processes. In 30th Annual ARCOM Conference Proceedings 2014. Association of Researchers in Construction Management (ARCOM)

Eldh S, Murphy B (2015) Code ownership perspectives. IEEE Softw 32(6):18–19

Fucci D, Romano S, Baldassarre MT, Caivano D, Scanniello G, Turhan B, Juristo N (2018) A longitudinal cohort study on the retainment of test-driven development. In: Proceedings of the 12th ACM/IEEE international symposium on empirical software engineering and measurement, pp 1–10

Hall T (2007) Longitudinal studies in evidence-based software engineering. In: Empirical software engineering issues. Critical Assessment and Future Directions: International Workshop, Dagstuhl Castle, Germany, June 26–30, 2006. Revised Papers. Springer, pp 41–41

Henderson K (1991) Flexible sketches and inflexible data bases: Visual communication, conscription devices, and boundary objects in design engineering. Sci Technol Human Values 16(4):448–473

Huvila I (2011) The politics of boundary objects: Hegemonic interventions and the making of a document. J Am Soc Inform Sci Technol 62(12):2528–2539

Jain R, Cao L, Mohan K, Ramesh B (2014) Situated boundary spanning: An empirical investigation of requirements engineering practices in product family development. ACM Trans Manage Inf Syst 5(3)

Kimble C, Grenier C, Goglio-Primard K (2010) Innovation and knowledge sharing across professional boundaries: Political interplay between boundary objects and brokers. Int J Inf Manage 30(5):437–444

Kitchenham BA, Charters S (2007) Guidelines for performing systematic literature reviews in software engineering. Technical Report EBSE 2007–001, Keele University and Durham University Joint Report

Lansing J, Sunyaev A (2016) Trust in cloud computing: Conceptual typology and trust-building antecedents. ACM SIGMIS Database Database Adv Inf Syst 47(2):58–96

Levina N, Vaast E (2005) The emergence of boundary spanning competence in practice: Implications for implementation and use of information systems. MIS Q 29(2):335–363

Maruping LM, Zhang X, Venkatesh V (2009) Role of collective ownership and coding standards in coordinating expertise in software project teams. Eur J Inf Syst 18(4):355–371

Matthies C, Heinrich R, Wohlrab R (2023) Investigating software engineering artifacts in devops through the lens of boundary objects. In: Proceedings of the 27th international conference on evaluation and assessment in software engineering, pp 12–21

Méndez Fernández D, Böhm W, Vogelsang A, Mund J, Broy M, Kuhrmann M, Weyer T (2019) Artefacts in software engineering: a fundamental positioning. Softw Syst Model 18(5):2777–2786

Moody DL (2003) The method evaluation model: a theoretical model for validating information systems design methods. European Conference on Information Systems 2003. Department of Information Systems, London School of Economics, pp 1–17

Nordberg ME (2003) Managing code ownership. IEEE Softw 20(2):26–33

Ouriques R, Fagerholm F, Mendez D, Bern BG (2023) An investigation of causes and effects of trust in boundary artefacts. Inf Softw Technol 158:107170

Ouriques R, Wnuk K, Gorschek T, Svensson RB (2023b) The role of knowledge-based resources in agile software development contexts. J Syst Softw 197:111572

Sedano T, Ralph P, Péraire C (2019) The product backlog. In: 2019 IEEE/ACM 41st international conference on software engineering (ICSE), pp 200–211
Star SL (2010) This is not a boundary object: Reflections on the origin of a concept. Sci Technol Human Values 35(5):601–617
Star SL, Griesemer JR (1989) Institutional ecology, translations' and boundary objects: Amateurs and professionals in Berkeley's museum of vertebrate zoology, 1907–39. Soc Stud Sci 19(3):387–420
Stettina CJ, Heijstek W (2011) Necessary and neglected? an empirical study of internal documentation in agile software development teams. In: Proceedings of the 29th ACM international conference on design of communication, SIGDOC '11, page 159–166, New York, NY, USA. Association for Computing Machinery
Swedish Research Council (2017) Good Research Practice. Swedish Research Council, Stockholm, Sweden
Van Loggem B, van der Veer GC (2014) A documentation-centred approach to software design, development and deployment. In Building Bridges: HCI, Visualization, and Non-formal Modeling: IFIP WG 13.7 Workshops on Human-Computer Interaction and Visualization: 7th HCIV@ ECCE 2011, Rostock, Germany, August 23, 2011, and 8th HCIV@ INTERACT 2011, Lisbon, Portugal, September 5, 2011, Revised Selected Papers. Springer, pp 188–200
Wohlin C, Runeson P, Höst M, Ohlsson MC, Regnell B, Wesslén A (2012) Experimentation in software engineering. Springer Science & Business Media
Wohlrab R, Pelliccione P, Knauss E, Larsson M (2018). Boundary objects in agile practices: Continuous management of systems engineering artifacts in the automotive domain. In: Proceedings of the 2018 international conference on software and system process, ICSSP '18. New York, NY, USA. Association for Computing Machinery, pp 31–40
Wohlrab R, Pelliccione P, Knauss E, Heldal R (2019) On interfaces to support agile architecting in automotive: An exploratory case study. In: 2019 IEEE international conference on software architecture (ICSA), pp 161–170
Zabardast E, Gonzalez-Huerta J, Tanveer B (2022) Ownership vs contribution: Investigating the alignment between ownership and contribution. In: 2022 IEEE 19th international conference on software architecture companion (ICSA-C), pp 30–34
Zaitsev A, Gal U, Tan B (2014) Boundary objects and change in agile projects. In: Proceedings of 25th Australasian conference on information systems, Auckland, New Zealand, 8th-10th Dec
Zaitsev A, Tan B, Gal U (2016) Collaboration amidst volatility: the evolving nature of boundary objects in agile software development. In: ECIS, page ResearchPaper172
Zurba M, Maclean K, Woodward E, Islam D (2019) Amplifying indigenous community participation in place-based research through boundary work. Prog Hum Geogr 43(6):1020–1043

Authors and Affiliations

Raquel Ouriques[1] · Fabian Fagerholm[2] · Daniel Mendez[3,4] · Tony Gorschek[3,4] · Baldvin Gislason Bern[5] · Victoria Vucic[5]

✉ Raquel Ouriques
raou@itu.dk

Fabian Fagerholm
fabian.fagerholm@aalto.fi

Daniel Mendez
daniel.mendez@bth.se

Tony Gorschek
tony.gorschek@bth.se

Baldvin Gislason Bern
baldvin.gislason.bern@axis.com

Victoria Vucic
victoria.vucic@axis.com

[1] IT University of Copenhagen, Copenhagen, Denmark

[2] Aalto University, Espoo, Finland

[3] Blekinge Institute of Technology, Karlskrona, Sweden

[4] Fortiss, Munich, Germany

[5] Axis Communications, Lund, Sweden