



D5.3 Assessment Method

Mass Soldal Lund, Bjørnar Solhaug, Ketil Stølen (SIN), Frank Innerhofer-Oberperfler (UIB), Massimo Felici, Valentino Meduri, Alessandra Tedeschi (DBL)

Document information

Document Number	D5.3
Document Title	Assessment method
Version	1.8
Status	Final
Work Package	WP 5
Deliverable Type	Report
Contractual Date of Delivery	31 January 2011
Actual Date of Delivery	31 January 2011
Responsible Unit	SIN
Contributors	SIN, UIB, DBL
Keyword List	Risk assessment, risk modeling, changing systems, changing risks
Dissemination level	PU

Document change record

Version	Date	Status	Author (Unit)	Description
0.1	26.04.10	Draft	M. S. Lund, B. Solhaug (SIN)	Abstract and table of contents
0.2	12.08.10	Draft	B. Solhaug (SIN)	Methodological needs, risk assessment process, and risk assessment techniques
0.3	01.10.10	Draft	B. Solhaug and K. Stølen (SIN)	Restructuring of contents
0.4	13.10.10	Draft	B. Solhaug (SIN)	First draft of all sections on new artifacts in main part of deliverable
0.5	09.11.10	Draft	B. Solhaug (SIN), F. Innerhofer-Oberperfler (UIB), M. Felici, V. Meduri, A. Tedeschi (DBL)	First full draft of ATM case study report added
0.6	17.11.10	Draft	B. Solhaug and K. Stølen (SIN)	Second draft of all sections on new artifacts in main part of deliverable
0.7	19.11.10	Draft	B. Solhaug (SIN)	New draft of introduction
0.8	23.11.10	Draft	B. Solhaug (SIN), F. Innerhofer-Oberperfler	Inserted section on integration with testing and appendix on HOMES case study
0.9	26.11.10	Draft	B. Solhaug (SIN)	First full draft of appendices on maintenance perspective, continuous evolution perspective and instantiation in CORAS

0.10	03.12.10	Draft	F. Innerhofer-Oberperfler (UIB), B. Solhaug	First full draft of section on integration with testing and of appendix on HOMES case study
0.11	07.12.10	Draft	B. Solhaug (SIN)	Second draft of introduction, first full draft of evaluation and conclusion
1.0	08.12.10	Draft	B. Solhaug, K. Stølen (SIN)	First full draft of deliverable
1.1	17.12.10	Draft	B. Solhaug (SIN)	Updated executive summary
1.2	04.01.11	Draft	B. Solhaug (SIN), F. Paci (UNITN)	Various amendments based on internal review by UNITN
1.3	07.01.11	Draft	F. Innerhofer-Oberperfler (UIB)	Changes to chapter 8
1.4	10.01.11	Draft	F. Innerhofer-Oberperfler (UIB)	Chapter 8: Updated definitions related to testing
1.5	10.01.11	Draft	B. Solhaug (SIN)	Polished and made ready for final quality check
1.6	11.01.11	Draft	M. Angeli (UNITN)	Quality check completed; minor remarks
1.7	13.01.11	Pre-final	B. Solhaug (SIN)	Finalization
1.8	21.01.11	Final	B. Solhaug (SIN)	Finalization

Executive Summary

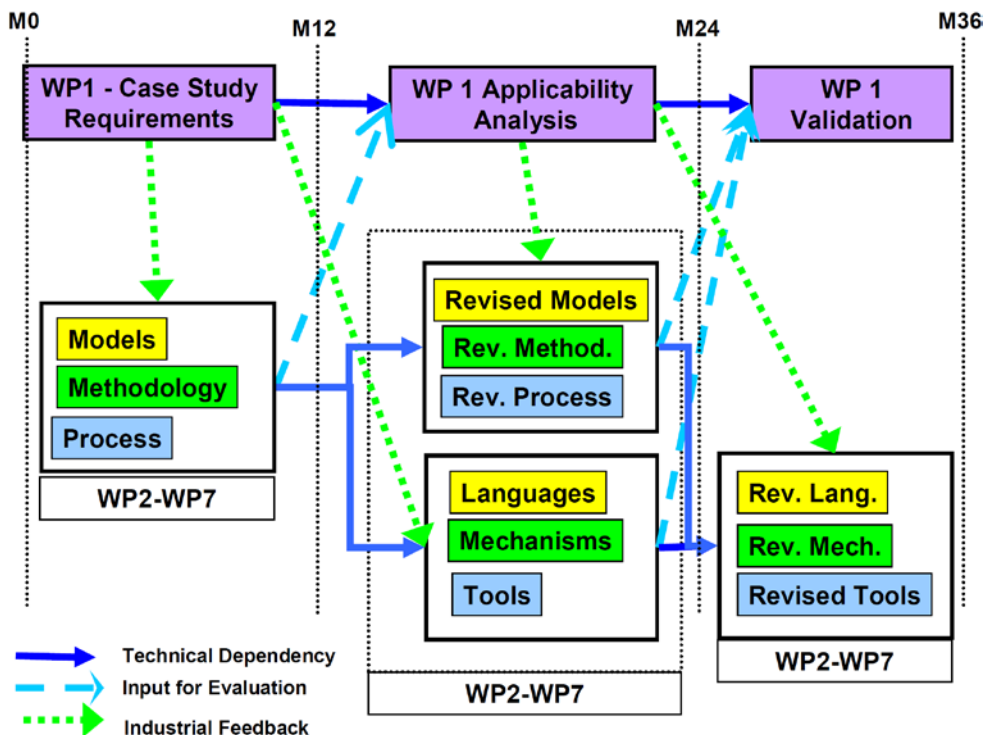
This deliverable presents a risk assessment method that meets the methodological needs of assessing changing systems. The guiding principle of the method is that by the occurrence of risk relevant changes, only the parts of the risk picture that may be affected by the changes should be assessed anew. Moreover, in order to properly understand the risks of changing systems as changing risks, the method facilitates the understanding and documentation of the changes to the identified risks. The main artifacts that are presented are the following:

- A risk assessment method for long-life evolving systems
- A language for the modeling and documentation of changing risks
- Techniques for tracing changes from target system to risk models

The method is formally founded by the formalization of the risk modeling language. The syntax of the language for the modeling of changing risks is formally defined, and is underpinned by a formal semantics. The precise reasoning about and analysis of risks are moreover supported by analysis rules that applies to the risk models. The applicability of the approach is demonstrated by the ATM and HOMES case studies. The former case study is the main WP5 case study and has been subject to a full risk assessment that is reported in this deliverable.

Position of the deliverable in the project timeline

This deliverable reports on the results of WP5 task T5.3, Assessment methods. According to the SecureChange description of work, the timeframe of T5.3 is M12-M24 with the milestone at M24 of delivering D5.3.



The main artifacts of WP5 are the risk assessment methodology and process, the risk modeling language and the risk assessment tools. Considering the SecureChange project timeline depicted above, the risk model artifact mainly belongs to the M0-M12 timeframe. However, although the risk modeling language was reported in D5.2 at M12, this artifact is further elaborated in this deliverable. Similarly, the risk assessment method reported in this deliverable was outlined already during M0-M12. This deliverable hence gives the full presentation of the method and process for the risk assessment of changing systems, as well as a revised risk modeling artifact that includes the formal foundation.

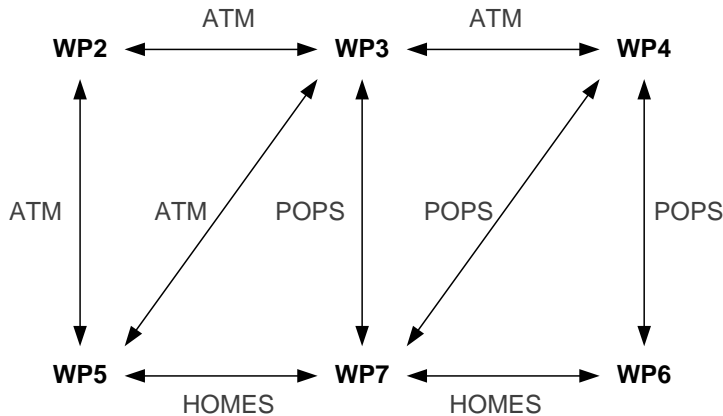
Validation

The WP5 artifacts of the risk assessment method, the risk modeling language and the prototype tool support should be understood as three related parts that integrate into an overall approach to the risk assessment of changing and evolving systems. The language serves as a technique for risk identification, risk assessment and risk documentation, and the modeling and documentation are in turn supported by the tools. Each of these artifacts is subject to the validation activities in SecureChange. The documentation framework prototype of D5.4 is due at M24 and will be subject to validation in year three. The revalidation prototype of D5.5 is due at M36, but we aim at validating preliminary versions during year three, i.e. during M24-M36.

The validation activities for the risk assessment method and the risk modeling language have been initiated by conducting risk assessment case studies. A full ATM risk assessment was conducted during year two, as documented in this deliverable. The change requirement addressed is the *organization level change*, and the security properties are *information protection* and *information provision*. WP5 uses also the HOMES case study, but to a lesser extent. It nevertheless serves as part of the validation. The change requirement that is addressed is *bundle lifecycle operations*, and the security properties are *policy enforcement* and *security expandability*. The HOMES risk assessment is documented in this deliverable. As shown by the project timeline depicted above, the WP5 case studies during M12-M24 have mainly demonstrated applicability.

Integration

The strategic position of WP5 in terms of case studies and integration with technical artifacts of the other work packages is shown in the figure below. The ATM case study serves as the example for demonstrating the integration with artifacts of WP2 and artifacts of WP3. The HOMES case study is used for exemplifying the integration with artifacts of WP7.



WP5-WP2 The integration link between WP2 and WP5 is reported in D2.2, and proposes a connection between the Integrated SecureChange Process developed in WP2 with the risk assessment method. In particular, the integration is made by instantiating the artifacts of the risk assessment method and the risk models in the Integrated SecureChange Process. The integration is demonstrated in the ATM case study, addressing the *organizational level change* and the security properties of *information protection* and *information provision*.

WP5-WP3 The integration link between WP3 and WP5 is reported in D3.2. The integration is both at conceptual level and at process level. At the conceptual level, an integration of concepts is presented and it is explained how requirement model artifacts should be mapped to risk model artifacts and vice versa. The process level integration leverages on the conceptual level integration for the integration of the requirements elicitation and risk assessment methodologies. The integration is demonstrated in the ATM case study, addressing the *organization level change* and the security properties of *information protection* and *information provision*.

WP5-WP7 The integration link between WP5 and the security testing approach of WP7 is reported in this deliverable. The integration is in terms of mapping artifacts from the risk model domain to the test model domain, and vice versa. Based on these options for mapping of model artifacts, the risk assessment activities and the testing activities are integrated so as to allow the two domains to leverage on each other. The integration is demonstrated in the HOMES case study, addressing the change requirement of *bundle lifecycle operations* and the security properties of *policy enforcement* and *security expandability*.

Index

DOCUMENT INFORMATION	1
DOCUMENT CHANGE RECORD	2
EXECUTIVE SUMMARY	4
1 INTRODUCTION	11
1.1 Artifacts of the Deliverable	11
1.2 Relation to Other WP5 Deliverables	12
1.3 Position of the Deliverable in SecureChange	13
1.4 Structure of the Deliverable	14
2 METHODOLOGICAL NEEDS AND CRITERIA	15
2.1 Methodological Needs	15
2.2 Evaluation Criteria	18
2.2.1 Scientific Criteria	18
2.2.2 Industrial Criteria	20
3 RISK ASSESSMENT PROCESS	21
3.1 Assessment Steps	21
3.1.1 Establish the Context	23
3.1.2 Identify Risks	26
3.1.3 Estimate Risks	29
3.1.4 Evaluate Risks	30
3.1.5 Treat Risks	30
3.2 Assessment Techniques	30
4 TARGET DESCRIPTION	32
5 RISK MODELING – FORMAL FOUNDATION	34
5.1 Risk Graphs	34
5.1.1 The Syntax of Risk Graphs	34
5.1.2 The Semantics of Risk Graphs	36
5.2 Risk Graphs for Changing Risk	37
5.2.1 The Syntax of Risk Graphs with Change	37



5.2.2	The Semantics of Risk Graphs with Change	42
5.3	Reasoning about Likelihoods in Risk Graphs	42
5.3.1	Rules for Likelihood Calculation	42
5.3.2	Guidelines for Consistency Checking Likelihoods	44
5.3.3	Reasoning about Likelihoods in Risk Graphs with Change	45
6	RELATING RISK MODEL TO TARGET DESCRIPTION	46
6.1	Indexing of Target Model	46
6.2	Specification of Relations between Target Model and Risk Model	49
6.3	Visualization of Relations to Target Model in Risk Models	52
6.3.1	Trace Model in Risk Graphs	52
6.3.2	Trace Model in Risk Graphs with Change	54
7	RISK ASSESSMENT METHOD	57
7.1	Overview	57
7.2	Conducting the Risk Assessment	59
7.2.1	Establish the Context	60
7.2.2	Identify Risks	71
7.2.3	Estimate Risks	80
7.2.4	Evaluate Risks	85
7.2.5	Treat Risks	87
8	INTEGRATION OF RISK ASSESSMENT AND TESTING	89
8.1	Mapping between CORAS and TTS	89
8.1.1	Risk Concepts	89
8.1.2	Testing Concepts	90
8.1.3	Integration	91
8.2	Information Flow from Risk Model to Test Model	93
8.2.1	Security Functionality Tests	93
8.2.2	Regression Tests	95
8.2.3	Misuse Case Test	96
8.2.4	Prioritization of Tests	97
8.3	Information Flow from Test to Risk Model	97
8.3.1	Confirmation of Risk Reduction by Treatments	98
8.3.2	Confirmation of Closure of Vulnerabilities	99
9	EVALUATION	101
9.1	Scientific Criteria	101
9.1.1	Evaluation of Risk Assessment Methodology	101
9.1.2	Evaluation of Risk Modeling Language	102



9.2	Industrial Criteria	103
9.2.1	Evaluation of Risk Assessment Methodology	104
9.2.2	Evaluation of Risk Modeling Language	105
10	CONCLUSION	106
	APPENDICES	107
11	A – MAINTENANCE PERSPECTIVE	108
12	B – CONTINUOUS EVOLUTION PERSPECTIVE	111
13	C – INSTANTIATION OF METHOD IN CORAS	116
13.1	CORAS Threat Diagrams as Specialized Risk Graphs	116
13.2	Generalizing the CORAS Language to Changing Risks	119
13.2.1	Standard CORAS Diagrams	119
13.2.2	CORAS Diagrams with Change	122
13.3	Assessment of Changing Risks using CORAS	130
13.3.1	Context Establishment using CORAS Asset Diagrams	130
13.3.2	Risk Identification using CORAS Threat Diagrams	131
13.3.3	Risk Estimation using CORAS Threat Diagrams	131
13.3.4	Risk Evaluation using CORAS Risk Diagrams	132
13.3.5	Risk Treatment using CORAS Treatment Diagrams	132
14	D – REPORT ON ATM CASE STUDY WITH CORAS	134
14.1	Context Establishment	134
14.1.1	Analysis Background and Motivation	134
14.1.2	Target Description	135
14.1.3	Asset Identification	167
14.1.4	High-level Risk Analysis	168
14.1.5	Establishing the Risk Evaluation Criteria	173
14.2	Risk Identification	175
14.2.1	Risk Identification before Changes	176
14.2.2	Risk Identification after Changes	177
14.3	Risk Estimation	180
14.4	Risk Evaluation	187
14.4.1	Risk Evaluation before Changes	192
14.4.2	Risk Evaluation after Changes	192
14.5	Risk Treatment	193
15	E – REPORT ON HOMES CASE STUDY WITH CORAS	197



15.1	Context Establishment	197
15.1.1	Business Needs	198
15.1.2	Actors	199
15.2	Change Requirements and Security Properties	199
15.2.1	Change Requirements	199
15.2.2	Security Properties	200
15.3	Timeline for the HOMES case study	200
15.4	T1: Risk Identification before Change	201
15.4.1	System Model	201
15.4.2	Risk Model	202
15.5	T2: Risk Identification under the Maintenance Perspective	203
15.5.1	Rationale for a New Risk Analysis under the Maintenance Perspective	203
15.5.2	Risk Model – Maintenance Perspective	203
15.5.3	Treatment Identification	204
15.6	T3: Risk Identification under the Before-After Perspective	205
15.6.1	System Model after the Change	205
15.6.2	Risk Model – Before-After Perspective	206
15.6.3	Identification of Risk to Change	208
16	F – GLOSSARY	209
	REFERENCES	211

1 Introduction

In this section we first give an overview of the artifacts that are presented in this deliverable. Thereafter we explain the relation between the artifacts presented in this deliverable and the other deliverables of the risk assessment work package, and we explain the position of the artifacts in the general setting of the artifacts that are delivered project wide in SecureChange. Finally we provide an overview of the contents.

1.1 Artifacts of the Deliverable

This deliverable presents a method for risk assessment of changing and evolving systems, with particular focus on risk with respect to security, privacy and dependability. For long-lived and evolvable systems the environment of which is heterogeneous and evolving, also the risks and the security threats are changing and evolving. The method that is presented is a systematic approach to identify, assess and document evolving risks so as to ensure that the risk picture and the risk assessment results are kept valid under change. Risk management supported by such a method should provide support for maintaining an acceptable risk level while the system evolves.

For systems that change and evolve, there is always the possibility of using established state-of-the-art risk assessment methods and conduct new iterations of risk assessments whenever changes have occurred. This is, however, not optimal, as it would require a full risk assessment to be conducted from scratch every time some change has occurred. Instead, the risk assessment of such systems should be supported by methods for how to identify the parts of the system that need to be reassessed after the changes, and how to indentify the previous risk assessment results that are still valid.

The risk assessment method that is introduced is based on established methods that are applicable in the traditional setting where change is not taken into account. The deliverable focuses on the additional guidelines, principles and procedures that are needed for addressing the particular challenges of assessing changing risks, and only refer to traditional methods when these are straightforwardly applicable. Additional methodological support is needed, not only to identify and understand the risks of changing systems, but also to take into account that stakeholders and other interested parties may change, stakeholders may become more or less risk aversive, security requirements may change, assets and asset values may change, and so forth. An adequate risk assessment method needs to take all such aspects into account in the identification, estimation and evaluation of the risks that are also changing and evolving.

A risk assessment method provides guidelines and principles for the identification and evaluation of risks. In the practical setting of conducting risk assessments, however, the users, i.e. the risk analysts, need several risk assessment techniques to support and facilitate the various activities of the risk assessment. Risk assessment techniques

may, for example, be techniques to support the risk identification, techniques for the quantitative or qualitative estimation and reasoning about likelihoods, techniques for consequence estimation, and so forth. Risk assessment techniques furthermore include language support for adequate ways of risk modeling and documentation, where the risk models usually serve as a basis for the aforementioned techniques.

This deliverable presents several novel risk assessment techniques that are developed to support activities of the risk assessment method that is introduced. The main artifacts are a language for the modeling and documentation of changing risks, support for establishing and documenting the traceability of change between the target system and the risk models, and support for identifying and reasoning about the propagation of changes through risk models.

The risk modeling language supports the specification of risks that change, and the specification of the relations to the target system. The latter facilitates the identification of the parts of the risk models that may be affected by system changes, and therefore need to be reassessed. The language furthermore provides support for the reasoning about likelihoods and for the consistency analysis of likelihood estimates. The syntax is formally defined and is underpinned by a formal semantics.

The main part of the deliverable presents a general approach to risk assessment of changing systems in the sense that some of the techniques that we generalize to the setting of changing risks can be understood as a common abstraction of several state-of-the-art techniques. This means that the artifacts that are introduced may be instantiated by these latter techniques, provided that also these are generalized to the setting of changing risks by the same principles. The advantage of this is that the techniques and underlying formalism presented in the main part can be transferred to several approaches. In the appendix we demonstrate this by instantiating the methods and techniques in CORAS.

For the purpose of illustrating the various artifacts presented in this deliverable, we use as a running example a full risk assessment from the Air Traffic Management (ATM) domain. The ATM risk assessment was conducted as a case study in the SecureChange project, and the full report of the results is presented as a separate appendix. The running example in the main part of the deliverable uses only extracts from the full case study documentation in the appendix.

1.2 Relation to Other WP5 Deliverables

Deliverable D5.1 provides an evaluation of the state-of-the-art within risk assessment, and identifies a number of requirements that should be fulfilled by the artifacts that are delivered in WP5. The main artifact of this deliverable is the assessment method for changing risks, and aims at contributing to filling the gap that was identified in D5.1.

The main artifacts of deliverable D5.2 are languages for the modeling of changing risks. These artifacts are closely related to the risk assessment methods, as the latter make extensive use of risk modeling. The risk modeling languages furthermore serve as a basis for other risk assessment techniques, such as risk identification and likelihood estimation. Deliverable D5.2 focuses on the abstract syntax of the risk modeling languages. In this deliverable, concrete syntaxes are introduced and provided a formal semantics. This deliverable furthermore explains how to utilize the

risk modeling languages during risk assessment, and also presents risk assessment techniques that use models of changing risks.

Deliverable D5.4 and D5.5 are both prototypes. These are tools and frameworks that are developed to support various risk assessment activities, such as risk identification and risk documentation. Tool support will furthermore be provided in order to automate some of the tasks in order to increase efficiency. The need for tool support is discussed for the various artifacts throughout the main part of this deliverable.

1.3 Position of the Deliverable in SecureChange

The strategic position of WP5 in terms of case studies and integration with technical artifacts of the other work packages is shown in Figure 1. The main case study of WP5 is the ATM case study for which a full risk assessment has been conducted and documented. The ATM case study moreover serves as the example for demonstrating the integration with artifacts of WP2 and artifacts of WP3. The HOMES case study is addressed to a lesser extent; the case study is used for further demonstrating the applicability of the risk assessment artifacts of WP5, and for exemplifying the integration with artifacts of WP7.




Figure 1 Case studies in WP integration

The integration link between WP5 and WP7 is reported in this deliverable, and explains how risk assessment and testing of changing systems can leverage on the technical solutions of each other. The integration link between WP3 and WP5 is reported in deliverable D3.2, and explains how risk assessment results can serve as input to the requirement engineering process and vice versa in a process that integrates the two respective methodologies. The integration link between WP2 and WP5 is reported in D2.2, and explains the integration of the risk assessment process into the overall Integrated SecureChange Process.

1.4 Structure of the Deliverable

The structure of the deliverable is as follows. In Section 2 we discuss the general methodological needs of a method for the risk assessment of changing and evolving systems, and we present evaluation criteria for the artifacts. In Section 3 we introduce the process for risk assessment of changing systems and identify the artifacts of risk assessment techniques that are needed for supporting this process. In the subsequent sections we introduce such assessment techniques. In particular, Section 4 presents the techniques we assume for adequately describing the target of analysis. In Section 5 we introduce the techniques for the modeling of changing risks, as well as the formal foundation of risk modeling with change. In Section 6 we introduce techniques to support the traceability of changes from the target system to risk models. These artifacts facilitate the identification of the parts of the risk picture that are affected by change and therefore need to be reassessed. In Section 7 we present the method for the risk assessment of changing systems. The method is based on the assessment process presented in Section 3 and makes use of the various assessment techniques that we introduce. The method is illustrated by means of a running ATM example. In Section 8 we propose and explain approaches to integrate risk assessment with testing, demonstrated and exemplified by the HOMES case study. In Section 9 we evaluate the artifacts of this deliverable with respect to the evaluation criteria presented in Section 2. Finally, we conclude in Section 10.

Following this main part of the deliverable are a number of appendices. The specific methodological needs of the risk assessment of changing and evolving systems depend somewhat on the kinds of changes. In the appendices of Section 11 and Section 12 we discuss two particular kinds of system changes and present assessment methods that are adequate for these. In the appendix of Section 13 we revisit the main risk assessment methods and techniques as presented in the main part of the deliverable and present their instantiation in CORAS, thus generalizing CORAS to a method for risk assessment of changing systems. In the appendix of Section 14 we give the full report of the ATM risk assessment case study, and in the appendix of Section 15 we report on the HOMES risk assessment. The appendix of Section 16 is a glossary with definitions of concepts related to risk assessment.

2 Methodological Needs and Criteria

In this section we generally discuss the methodological needs for adequately conducting risk assessments of changing and evolving systems. Thereafter we present the evaluation criteria for the proposed assessment method and related artifacts.

2.1 Methodological Needs

The international risk management standard ISO 31000 [22] defines risk management as coordinated activities to direct and control an organization's risk. Risk may be expressed in terms of the consequences of an event (unwanted incident) and the likelihood for the event to occur [22][23]. The risk management process as defined in ISO 31000 is illustrated in Figure 2. The five activities in the middle constitute the core activities of a risk analysis, and are described as follows:

- *Establish the context* is to define the external and internal parameters to be accounted for when managing risk, and to set the scope and risk criteria for the risk management policy.
- *Risk identification* is to find, recognize and describe risks.
- *Risk estimation* is to comprehend the nature of risk and to determine the risk level.¹
- *Risk evaluation* is to compare the risk estimation results with the risk criteria to determine whether the risk and its magnitude are acceptable or tolerable.
- *Risk treatment* is the process of modifying the risk.

The remaining two activities are continuous activities of the overall risk management process, and are described as follows:

- *Communicate and consult* are the continual and iterative processes an organization conducts to provide, share or obtain information, and to engage in dialogue with stakeholders about risk management.
- *Monitoring* involves the continuous checking, supervising and critically observing the risk status in order to identify changes from the performance level required or expected, whereas *review* focuses on the activity undertaken to determine the suitability, adequacy and effectiveness of the subject matter necessary to achieve established objectives.

A risk management process aligned with the ISO 31000 principles and guidelines will typically generate a risk picture that focuses on a particular system configuration at a particular point in time, and the results and conclusions will therefore be valid only under the current configuration and assumptions. However, the target system and its environment may evolve and vary over time, as may also the stakeholders and the risk

¹ The ISO 31000 standard refers to this activity as *risk analysis*.

criteria. At a very general level the ISO 31000 standard addresses changes by the *Monitor and review* activity, since the activity aims at detecting “changes in the external and internal context, including changes to the risk criteria and the risk itself, which can require revision of risk treatments and priorities” [22]. However, the ISO 31000 standard does not provide guidelines for how to manage such changes in a systematic way. In particular, a risk management process explicitly targeting changing and evolving systems should provide guidelines for how to address change within the core activities from context establishment, through risk assessment, to risk treatment.

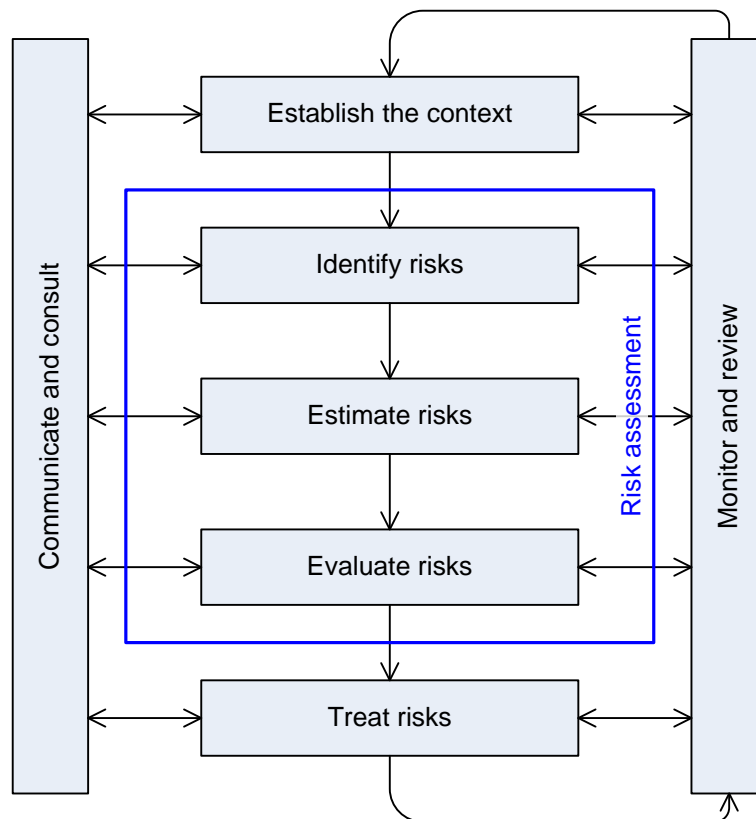


Figure 2 Risk management process

When targeting long-lived and evolving systems for the purpose of risk management and risk assessment the challenge is to ensure that the assessment and analysis results are kept valid under change. A straightforward way to ensure this is to conduct a full risk assessment from scratch whenever a potentially risk relevant change has occurred. Needless to say, such a strategy is not to prefer as it is time and resource consuming, and as it often implies conducting exactly the same assessments again to the extent that the risk picture is persistent. Instead, a customized assessment method for changing and evolving systems should provide guidelines and techniques for how to systematically trace the relevant changes from the system to the risk picture, and thereby for how to update only the part of the risk picture that is affected by the changes. Such a method should furthermore be supported by adequate risk and threat modeling languages with the expressiveness to represent changing and evolving risks,

and there should also be support for the assessment of changing and evolving risks, as well as the detection of dependencies of the risks on the parts of the target or the target environment that may be subject to change.

A crucial part of the risk management and risk assessment of evolving systems is the ability to efficiently and systematically trace system changes to changes in the risk picture. By identifying and documenting the relations between a representation of the target of analysis and the risk models, we can efficiently identify which parts of the system changes that affect the risk picture and consequently update the risk assessment results for the relevant parts only.

How to adequately deal with changes in a risk assessment depends, however, on the nature of the changes in each specific case. The adequate method for assessing risks of changing and evolving systems thereby also depends on the nature of the changes. We have identified three different perspectives on change, each with its specific methodological needs. These are the maintenance perspective, the before-after perspective, and the continuous evolution perspective.

By the *maintenance perspective* we refer to smaller changes that evolve more or less unnoticed over time and that eventually may accumulate to substantial changes that make previous risk assessment results outdated. An outdated risk analysis may give a false and invalid picture of the risks associated with the current system, and therefore requires a new risk assessment to be conducted. Conducting a risk assessment from scratch is time and resource consuming, so an adequate assessment method should allow the maintenance of the previous risk assessment results by addressing only the parts that are affected by the system changes.

By the *before-after perspective* we refer to substantial system changes that are planned or anticipated, and that may themselves motivate a risk assessment. From this perspective we need a clear understanding of the target of analysis and the risk picture as-is, and we need an understanding of the target of analysis and the risk picture to-be. Since the process of change, i.e. the transition from the current to the future system, may itself involve risk, we furthermore need to conduct a risk assessment of the change process in the before-after perspective.

By the *continuous evolution perspective* we refer to systems that gradually change and evolve, and where the changes are planned or where they can be anticipated. Such changes may, for example, be the plan to gradually increase the number of components working in parallel, the gradual inclusion of more and more sites into a system, the anticipation of wear and tear of hardware, the prognosis of increase of system users, the prognosis of increase of cyber attacks, and so forth. What is common to such cases is that the target of analysis can be described as a function of time. The objective is then to understand and describe also the risks as a function of time. A risk assessment method adequate for the continuous evolution perspective would give a risk picture not for one or a few, but for any future point in time.

In this deliverable, we mainly address the before-after perspective. This is because the before-after perspective is the perspective that is mainly addressed in the SecureChange project, and is also the main perspective of all of the SecureChange case studies. The central WP5 case study is ATM, and the full ATM risk assessment that was conducted as part of the WP5 activities is from the before-after perspective, as documented in the appendix.

The process for risk assessment of changing systems that is presented in Section 3 is independent of the various perspectives. The assessment methods and techniques that are presented subsequently, however, specifically address the before after perspective.

Considering the methodological challenges of risk management and risk assessment of changing systems, also the maintenance perspective and the continuous evolution perspective are important and interesting. Adequate methodologies for risk assessment of changing systems should therefore provide guidelines and techniques also for these. While being outside our chosen focal point in this deliverable, we address also these perspectives separately, only more brief, in the appendix.

2.2 Evaluation Criteria

We refer to deliverable D5.1 [31] for the criteria that were used for the evaluation of state-of-the-art risk assessment methods and techniques in the setting of changing and evolving systems. As concluded there, there is little or no explicit support in state-of-the-art risk analysis methods for handling changing and evolving risks. Briefly summarized, the CORAS method [24] provides guidelines for identifying parts of risk analysis documentation affected by changes and for maintaining risk analysis documentation. ProSecO [20][21] provides guidelines for relating risk analysis documentation to target descriptions, for identifying parts of the risk analysis documentation affected by changes, and for identifying parts of the target in need for additional risk analysis in the face of change. Such guidelines may facilitate the overall management of risks of systems or organizations that may change and evolve. However, both approaches are restricted to component-based systems and system descriptions, and to discrete changes.

We distinguish between scientific criteria and industrial criteria. The scientific criteria are case study independent, whereas the industrial criteria are related to the application of the WP5 artifacts in ATM and HOMES. The industrial criteria presented in the following are described in more detail in SecureChange deliverable D1.2.

The main artifact of this deliverable is the risk assessment method. However, as the assessment method is tightly interwoven with the assessment techniques, in particular the risk modeling language, we provide criteria also for the latter. The criteria moreover address aspects in relation to tool support.

2.2.1 Scientific Criteria

The scientific apply to the two WP5 artifacts of the risk assessment method and the risk modeling language.

2.2.1.1 Risk Assessment Methodology

The criteria for the risk assessment methodology are divided into the categories of criteria for a well-defined methodology, criteria for a potentially computer-aided methodology, and criteria for linkage of artifacts.

Well-defined methodology

- The risk assessment methodology should be defined in terms of procedural steps.
- It should be precisely defined for each step of the risk assessment method which artifacts that are input and which artifacts that are output of the step.
- For each change requirement, the risk assessment methodology should provide explicit techniques and guidelines for how to trace changes from system to risk models.

Computer-aided methodology

- Each of the risk analysis techniques supporting the various steps can lend itself to tool support.
- Each of the model artifacts have a formally defined syntax and can lend itself to tool support.

Explicit linkage of artifacts

The artifacts can be artifacts that are used as input to the risk assessment (such as system models or requirements models) or they can be artifacts that are produced as output of the risk assessment (such as models of changing risks).

- The traceability between target system and risk models should be explicitly represented as syntactic links of a model artifact in itself.
- The syntactic links should be based on the semantics of the artifacts that are linked, providing means for reasoning about the kinds of changes that are traced from target system to risk models.
- The notion of dependency between risk elements should be formally defined, providing means for tracing the propagation or risk changes through risk models.

2.2.1.2 Risk Modeling Language

The criteria for the risk modeling language are divided into the categories of criteria for well-formedness and consistency, criteria for tool support, criteria for formalization, and criteria for local usability.

Well-formedness rules and consistency rules of constructs

- The risk modeling language should have a formally defined syntax that precisely captures the set of syntactically correct specifications in the language.
- The risk modeling language should have a formally defined semantics that precisely captures the set of consistent specifications in the language.

Computer-aided support for syntactically correct and consistent specifications

- The syntax of the modeling language should lend itself to tool support for the detection of syntactical errors in the specifications.
- The semantics of the modeling language should lend itself to tool support for the detection of inconsistencies in the specifications.



Formal characterization of specifications

- The semantics of the risk modeling language should enable a precise and formal characterization of the specified behavior that is acceptable or unacceptable.

Local usability of specifications

- The specifications should be self-contained, i.e. the user should be able to determine the syntactical correctness, the consistency and the semantics of the specifications without the need to consult or understand other artifacts than the specification itself.

2.2.2 Industrial Criteria

The industrial criteria are evaluation criteria for the WP5 artifact of the risk assessment methodology and the risk modeling language in the case studies. The main case study in WP5 is the ATM, for which a full risk assessment has been conducted and documented. The HOMES case study is also addressed, however to a lesser extent.

The criteria concern the effective usage of the artifacts and express requirements to their applicability in the industrial case studies, as well as requirements to human effort.

2.2.2.1 Risk Assessment Methodology

Applicability

The first criterion is that the risk assessment methodology and its assessment techniques can be applied on the case studies for the assessment, modeling and documentation of changing risks.

Human effort

The second criterion is that the risk assessment methodology and its techniques can produce the desired results with less effort than by using alternative, traditional methods.

2.2.2.2 Risk Modeling Language

Applicability

The first evaluation criterion is that the risk modeling language can be applied on the case studies for modeling and assessing changing risks. The use of the risk modeling language should result in consistent and syntactically correct specifications that are well understood.

Human effort

The second evaluation criterion is that the modeling of changing risks in the case studies can be conducted with less effort than by using traditional risk modeling languages or techniques.

3 Risk Assessment Process

In this section we describe a general process for risk assessment of changing and evolving systems. The process is general in the sense that it provides guidelines that are applicable to all the three perspectives on change, i.e. the maintenance perspective, the before-after perspective and the continuous evolution perspective.

Practitioners usually rely on several risk assessment techniques in order to carry out the activities of a risk assessment process. These techniques typically facilitate activities such as likelihood analysis, consequence estimation, consistency checking, and treatment evaluation. In turn, these techniques are usually based on customized languages or other specification means for representing the subject matter. In Section 3.1 we focus on the risk assessment process and what such a process needs to provide in terms of principles and guidelines for targeting changing systems in an adequate way. In Section 3.2 we discuss some of the assessment techniques and modeling support that should be provided for supporting the risk assessment method.

Subsequently, in Section 4 through Section 6 we introduce techniques for risk assessment of changing systems. In Section 7 we present a risk assessment method based on the assessment process that makes use of the assessment techniques. The risk assessment process, method and techniques are in Section 7 exemplified by the ATM case study. These sections in the main part of the deliverable mainly address the before-after perspective. The remaining two perspectives are addressed in the appendix.

3.1 Assessment Steps

We take as a starting point the risk assessment process defined by ISO 31000 [22] and depicted in Figure 2. Our main concern is the analysis process depicted in the middle from the first activity *Establish the context*, via the activities of *Risk assessment*, through the last activity *Treat risk*.

The first activity of establishing the context is where the premises for the subsequent risk assessment are made. Establishing the context should result in a target description, which is the documentation of all the information that serves as the input to and basis for the risk assessment. This means that any information about the target of analysis that is relevant for the risk assessment and its outcome needs to be included in the target description. It also means that any risk relevant change in the target of analysis must be reflected by changes to the target description; in order to incorporate system change into the risk assessment process, the method must therefore come with guidelines for how to include the specification of change in the target description.

When we are considering the actual risk assessment that succeeds the context establishment, there is then a need for guidelines for how to take change into account during the activities of risk identification, risk estimation and risk evaluation. Given a target description that incorporates system change, one could of course perceive the target of analysis as two different systems and use standard risk assessment methods



for conducting two separate assessments from scratch. This is, however, not an optimal approach for several reasons.

- In most cases a substantial part of the risk assessment results of the target of analysis before the changes are still valid after the changes. An adequate risk assessment method should therefore provide guidance for how to identify only the parts of the target that need to be reassessed after the changes, and how to identify the previous risk assessment results that are still valid.
- Conducting the risk assessment from scratch every time a potentially risk relevant change has occurred will generally be more time and resource consuming than risk assessments of only the parts that are affected by the change.
- Conducting new risk assessments from scratch will yield updated risk pictures as new snapshots at given points in time. This may not be sufficient for capturing and assessing the dynamics of risks that may continuously evolve over time.
- Separate risk pictures as snapshots at different points in time will not give explicit support for identifying which risks change and due to which causes. There should be support for relating risk models or other risk documentation from different points in time such that changes in risks are explicitly represented.
- Representing the target of analysis as two separate systems will not show the change process itself. For substantial changes that are planned, there is a need to make a description of the change process, both in order to better understand the result of the changes and to identify and assess the risks of implementing the change process itself.

We now introduce a risk assessment process that incorporates change in all activities, with guidelines for how to track changes throughout the process. We focus in this section on the principles and guidelines that are common for the three perspectives on change.

Figure 3 gives an overview of a risk analysis process based on the ISO 31000 standard where the assessment of changing risks has been incorporated. The boxes at the right hand side depict the activities that need to be added in order to take change into account throughout the risk assessment process. These activities should be understood as an integrated part of the respective activities of the integrated process; they have been extracted in the figure only for the purpose of highlighting.

As shown by the diagram, the risk assessment process is an iterative process. In practice, a traditional risk assessment process is usually conducted sequentially, possibly with some backtracking. When support for the assessment of changing risks is incorporated into the process it may be useful from a practical point of view with more iterations in order to reassess parts that are affected by change. When we in the following describe each of the five activities in turn we focus on the iterations and activities that particularly address change.

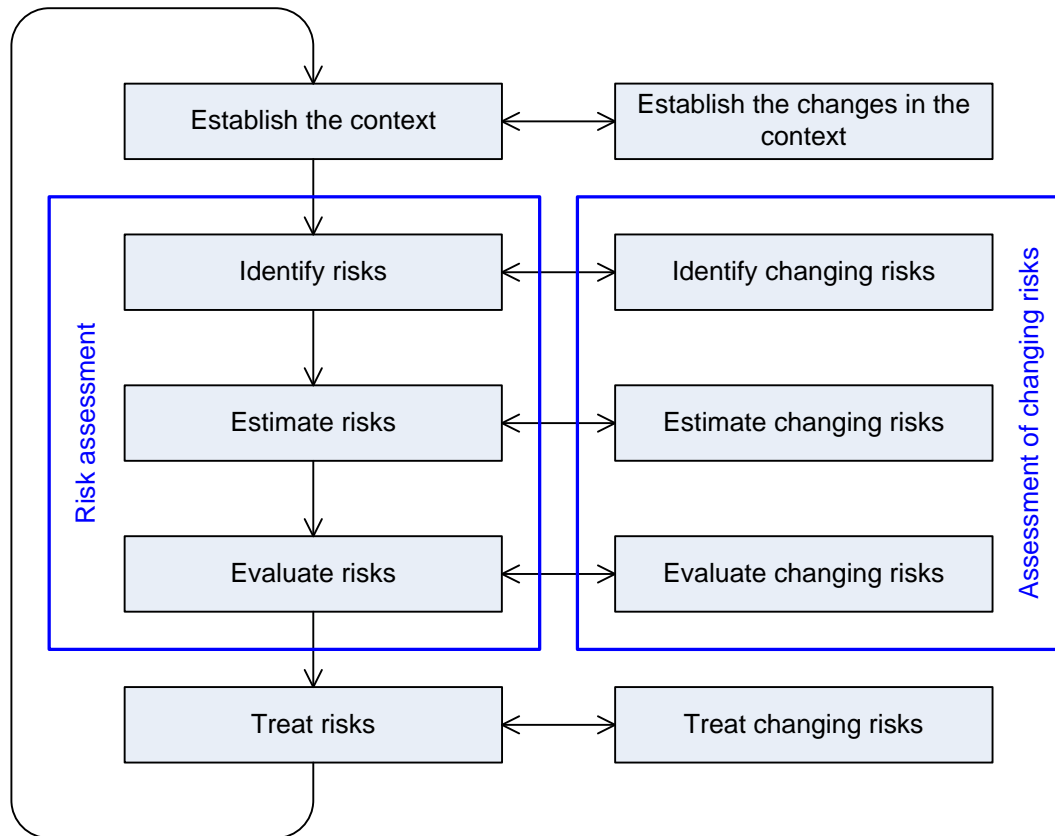


Figure 3 Risk analysis process for changing systems

3.1.1 Establish the Context

Establishing the context of the analysis includes articulating the goals and objectives of the analysis and deciding the focus and scope of the analysis. In particular, when we are establishing the context we need to determine precisely what the target of analysis is and what the assets that need to be protected are. The risk assessment is conducted with respect to the identified assets, and it is only by precisely understanding what the assets are that we can conduct a risk assessment that meets the overall goals and objectives of the assessment.

In a risk assessment, the notions of party, asset and risk are closely related. A party is an organization, company, person group or other body on whose behalf the risk assessment is conducted. An asset is something to which a party assigns value and hence for which the party requires protection. A risk is the likelihood of an unwanted incident and its consequence for a specific asset. This means that if there is no party, it also makes no sense in speaking about assets. And without assets there can moreover be no risks.

Most commonly it is the customer of the analysis that is the party of the risk assessment, although in some cases we also need to take other parties into account. If the customer is a service provider, for example, it may be that the customer wishes to include the end-users as one of the parties. When the parties of the analysis have

been identified, we proceed by establishing and documenting the target description and establishing and documenting the risk evaluation criteria.

3.1.1.1 Establishing the Target Description

The target description is the documentation of all the information that serves as the input to and the basis for a risk assessment. This includes the documentation of the target of analysis, the focus and scope of the analysis, the environment of the target, the assumptions of the analysis, the parties and assets of the analysis, and the context of the analysis. The class diagram of Figure 4 gives an overview of the elements of a target description.

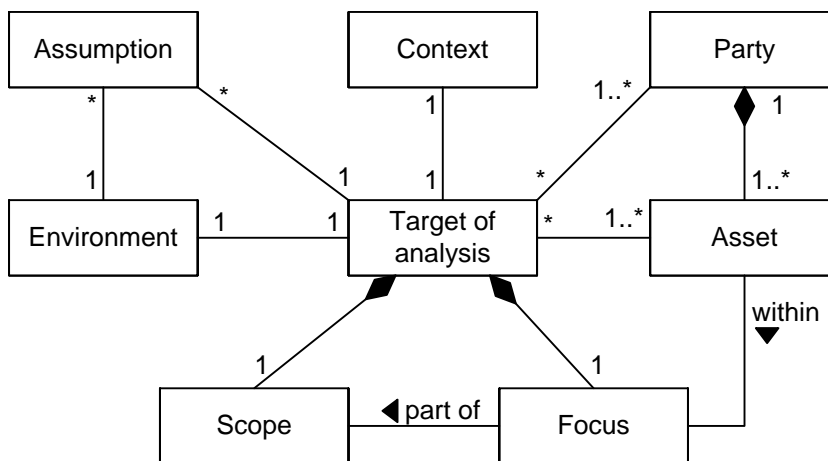


Figure 4 Target description

The target of the analysis is the system, organization, enterprise, or the like that is the subject of a risk analysis. The focus of the analysis is the main issue or central area of attention. The focus is within the scope of the analysis, which the extent or range of the analysis. The scope defines the border of the analysis, i.e. what is held inside and what is held outside of the analysis. The environment of the target is the surrounding things of relevance that may affect or interact with the target; in the most general case the environment is the rest of the world. The assumptions are something we take as granted or accept as true, although it may not be so; the results of a risk assessment are valid only under the assumptions. The context of the analysis is the premises for and the background of the analysis. This includes the purposes of the analysis and to whom the analysis is addressed.

For detailed methodological guidelines on how to establish the target description in a traditional setting we refer to existing methods such as the ISO 31000 standard or the approaches described in deliverable D5.1. In the following we focus on the methodological guidelines for how to address change during this activity.

For a given target of analysis we assume that we can use a traditional risk assessment method to conduct the context establishment while not taking into account changes. Given the resulting target description, we then need to take into account a set of change transactions. A change transaction brings the target to a new and different state and may imply changes to the risk picture. We distinguish between two kinds of

change transactions, namely change requests and change logs. A change request refers to changes that are planned or foreseen and that can be controlled. A change log, on the other hand, refers to observed changes that have already occurred and that need to be reacted to.

The additional task of establishing the changes in the context as depicted in Figure 3 includes making a description of the target of analysis when the change transactions have been implemented. This extended target description should include both a description of the change transaction and the result of the transaction, although it should be possible to deduce the latter from the current target description and the description of the change transaction. Precisely how and which parts of the target description and change transactions that should be documented depends on the relevant perspective on change. For the continuous evolution perspective, for example, the specification of how the target evolves over time should be incorporated in the target description.

For most change transactions, the changes concern the target of analysis. Such changes can be new or different work processes, the introduction of new services or applications, changes in users or roles, etc. These may imply changes in vulnerabilities, threats, threat scenarios, and so forth, and therefore require new risk assessments of parts of the target. There may, however, also be changes in parties, changes in assets or asset priorities, changes in the environment or in the assumptions, changes in the focus or scope, and so on. A set of change transactions therefore triggers a new iteration of the context establishment in order to identify and document all the relevant issues. Because all elements of the target description are relevant for and affects the subsequent risk assessment, each of them needs to be addressed when considering the change transactions; otherwise the results of assessing the risks when changes are taken into account may be false.

In order to conduct the activity of making the target description of a changing and evolving target of analysis, there is a need for the following artifacts:

- Language for documenting the target description when change transactions have not been taken into account.
- Language for documenting the change transactions.
- Language for documenting the target description of the changed target that results from the change transactions.

3.1.1.2 Establishing the Risk Evaluation Criteria

The risk evaluation criteria are a specification of the risk levels that the parties of the risk assessment are willing to accept. The criteria will later be used to evaluate the significance of risk, and should reflect the values, objectives and resources of the parties in question.

When we are deciding the risk evaluation criteria we need to take into account not only the views of the parties, but also the nature of the assets, the types of consequences and how they should be measured and described. We furthermore need to take into account how likelihoods should be defined, and the timeframe of the likelihoods. Specifically, we need for each asset to define a consequence scale where each consequence value describes a level of impact of an unwanted incident on an asset in



terms of harm or reduced asset value. We furthermore need to define a likelihood scale of a suitable time frame, the values of which will be used to describe the frequency or probability of unwanted incidents and threat scenarios to occur.

Recall that a risk is the likelihood of an unwanted incident and its consequence for a specific asset. The risk level is the level or value of a risk as derived from its likelihood and consequence. The risk level of each combination of a likelihood and a consequence is calculated by a risk function. Since it is only the party of a given asset that can determine the severity of a risk, it is the party that must determine an adequate risk function.

Essentially, the risk evaluation criteria are a specification of the level at which risks become unacceptable. When we have established and documented the consequence scales, the likelihood scale and the risk function, we establish and document the risk evaluation criteria as a mapping from risk levels to one of the categories of acceptable and unacceptable. For intermediate risk levels, we may also operate with categories such as “accept, but monitor risk”.

Some change transactions are of a kind that does not affect the assets or other values, objectives or resources of the parties. In that case, there is also no need to reconsider the risk evaluation criteria. For other change transactions, the value or priorities of assets may change, new assets may arise, the parties may become more or less risk averse, and so forth. In that case we need a new iteration on establishing and documenting the risk evaluation criteria.

3.1.2 Identify Risks

Risk identification means to identify unwanted incidents, threat scenarios that may lead to unwanted incidents, threats that initiate threat scenarios and the vulnerabilities that make it possible for scenarios and incidents to arise.

An unwanted incident is an event that harms or reduces the value of an asset. A threat is a potential cause of an unwanted incident, and may be both human and non-human. A human threat may furthermore be both deliberate and accidental, where a deliberate human is an adversary of malicious intent and an accidental human threat is someone that may cause unwanted incidents, for example, by accident or sloppiness. A threat scenario is a chain or series of events that is initiated by a threat and that may lead to an unwanted incident. A vulnerability is a weakness, flaw or deficiency that opens for, or may be exploited by, a threat to cause harm to or reduce the value of an asset.

The risk identification should involve people with appropriate expert knowledge about the target of analysis. The activity of extracting the relevant information relies on techniques and tools for identifying risk relevant information, for structuring the information in a sensible way, and for adequately documenting the information. While the documentation of the risks that are identified should serve as a means for reporting the finding to the relevant stakeholders, it should at the same time facilitate the subsequent estimation and evaluation. In this section we focus on the methodological guidelines for risk identification of changing and evolving systems. In Section 3.2 we will discuss more closely the required documentation techniques.

For now we assume that the risk assessment process is supported by artifacts for target modeling and artifacts for risk modeling. The former artifact supports making the



target description and the latter artifact supports making the risk documentation. From the first activity of establishing the context we have established a description of the target where the relevant change transactions are not taken into account, and we have established a target description where the change transactions are taken into account. If relevant, we have also made a description of the change process or change transaction itself.

As mentioned above, it is the target description that serves as the input to and basis for the subsequent risk assessment. The objective of the risk identification is to identify and document the changing risks given the description of the changing target. The guidelines for how to conduct the risk identification based on the description of the changing target depends somewhat on the relevant perspective on change. However, the main principle remains: To the extent that we have identified and documented the risks for the target of analysis without taking into account changes, we only address the parts of the target that are affected by the change when identifying the changing risks.

This means that when considering the target description without the changes, the risk identification and the risk documentation are conducted according to traditional risk assessment methods. When this is completed we need to update the resulting risk documentation according to the change transactions. This is conducted by making a walkthrough of the current target description and risk documentation and identifying the risks that are persistent under change. This part of the risk documentation can then immediately be included in the documentation of the risks when the change transactions are taken into account, with no further investigation. The risks that may be affected by change need to be considered again: Previous scenarios, events, etc. may change, new may arise, and others may disappear.

The methodological problem of identifying and documenting the changed risks is illustrated in Figure 5. The rounded rectangle at the upper left corner illustrates the target description before the change transaction has been taken into account, and the rounded rectangle at the lower left illustrates the documentation of the identified risks given this target description. When moving to the right hand side we see the target description where the change transaction has been considered. We see that the target element T4 has changed to T4', whereas the remaining elements are not affected. The problem is then how to update the risk documentation without conducting the risk identification from scratch using the full changed target description as input and basis.

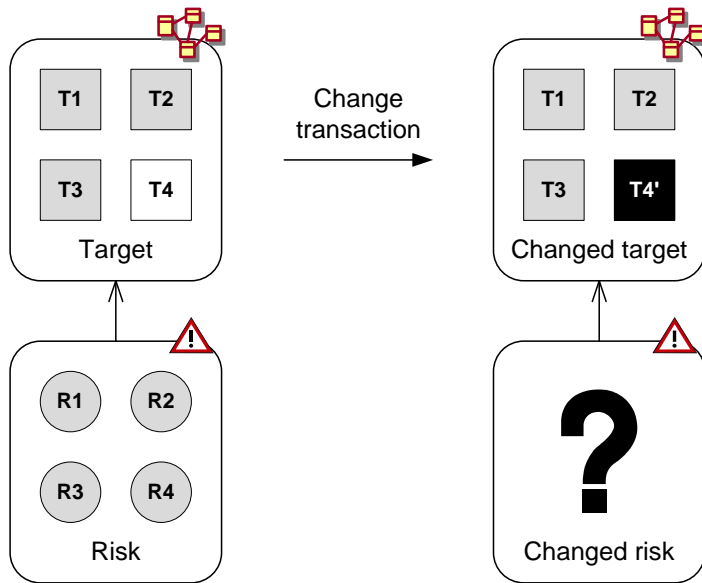


Figure 5 Identification of changed risks

In order to enable and facilitate the risk identification of the changed target, we need techniques for establishing and documenting traceability between the target description and the risk documentation. This process is illustrated in Figure 6. We see, for example that risk R1 traces to target element T1. Because R1 traces to T1 only also after the change transaction, R1 is persistent under change and can therefore be immediately included in the documentation of the changed risks. We furthermore see that risk R4 traces to target element T4, and that T4 is affected by the change transaction. The documentation of R4 must then be reconsidered in order to identify possible changes to this part of the risk picture. The changes in the risk picture regarding R4 are depicted by the fragment R4' of the documentation of the changed risks.

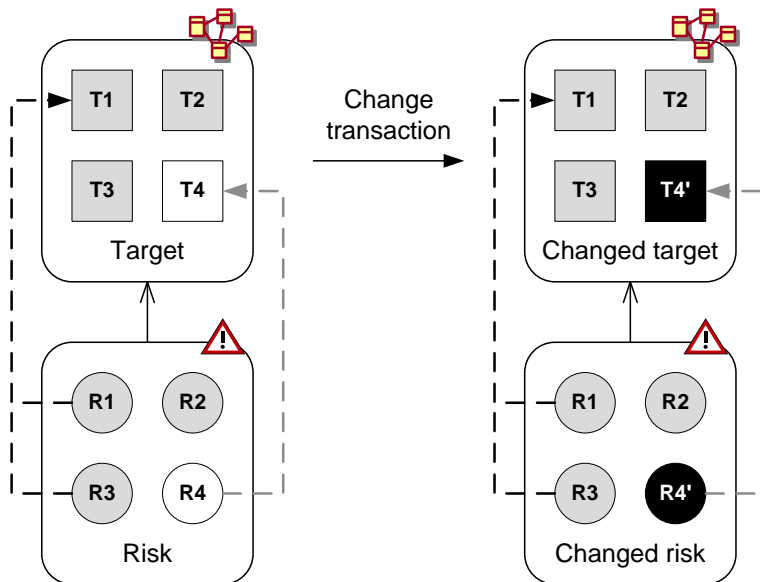


Figure 6 Identification of changed risks based on traceability

The methodological guidelines for risk identification of a changing target of analysis are summarized as follows:

1. Identify and document risks by using as input the target description before change transactions have been taken into account.
2. Establish and document the traceability between the target description before change and the risk documentation resulting from the previous step.
3. Based on the traceability and the description of the changed target, identify the parts of the risk documentation that are persistent under change.
4. Conduct the risk identification of the changed target only with respect to the parts of the target and the risks that are affected by the change transaction.

In order to conduct the activity of risk identification of a changing target of analysis, there is a need for the following artifacts:

- Techniques for risk identification and language for risk documentation (risk modeling) of risks that change.
- Techniques for establishing traceability between target description and language for documenting the traceability (trace model).

3.1.3 Estimate Risks

The objective of the risk estimation is to establish an understanding of the identified risks, and to provide the basis for the subsequent risk evaluation and risk treatment. By considering the causes and sources of risk, including the threats, threat scenarios and vulnerabilities, the risk estimation amounts to estimating and documenting the likelihoods and consequences of the identified unwanted incidents. It is the likelihoods of unwanted incidents and their consequences for assets that constitute risks, and by making estimates of the likelihoods and consequences we can understand which risks are the most important and which risks are less relevant.

It is, however, not enough to consider the unwanted incidents alone in order to reach an adequate understanding of the risks. We also need to understand the main causes for risks to arise. For this purpose we estimate and document the likelihood for the identified threat scenarios to occur and the likelihood for the identified threats to initiate threat scenarios and unwanted incidents. We may furthermore make estimations of the conditional likelihoods for threat scenarios or unwanted incidents to lead to other threat scenarios and unwanted incidents. The result of such an analysis will serve as a basis for determining the most important sources of risks, and thereby also the most efficient and appropriate options and strategies for risk treatment and mitigation.

Given the documentation of the identified risks from the previous step, including the documentation of the changing risks, the risk estimation of a changing and evolving target is quite similar to traditional risk analyses: The estimation is conducted by a walkthrough of the risk documentation addressing each of the relevant elements in turn. To the extent that risks are persistent under the change transactions, the estimation is not repeated.

The estimates need to be continuously documented, which means that there must be adequate support for including the estimates in the risk documentation. In order to



conduct the activity of risk analysis and the documentation of the results, there is a need for the following artifact:

- Techniques for making estimates of likelihoods and consequences of changing risks, and language for documenting the results.

3.1.4 Evaluate Risks

The objective of the risk evaluation is to determine which of the identified risks that need treatment, and to make a basis for prioritizing the treatment options. Basically, the risk evaluation amounts to estimating the risk levels based on the likelihood and consequence estimates, and to compare the results with the risk evaluation criteria. The need for treatment can be considered on the basis of this comparison.

The risk evaluation of a changing and evolving target of analysis is conducted in the same way as risk evaluation of traditional risk assessments. Given the risk documentation of the changing risks with the risk estimates, the risk evaluation is conducted by calculating the risk level of each pair of an unwanted incident and asset that is harmed by the incident. The calculation is straightforwardly done by using the risk function defined during the context establishment. For changing systems, the criteria may of course be different before and after some given change transactions.

3.1.5 Treat Risks

A risk treatment is an appropriate measure to reduce risk level. The risk treatment succeeds the risk assessment activities, and the objective is to identify and select a set of treatment options for the risks that are not acceptable according to the risk evaluation criteria. The implementation of the selected treatments should bring the risk level down to an acceptable level. Before the identified treatments are selected and implemented, we need to conduct a cost-benefit analysis. If a treatment option is more costly than its benefit in terms of reducing risk level, the treatment should obviously not be implemented.

The adequate strategy for identification and implementation of treatments depends on the perspective on change. For changes that have already occurred, there is obviously no use of identifying treatments for the risks of the target before the change transactions. For changes that are planned or predicted, however, it may be that we are only concerned about the future risks and to ensure that the planned or foreseen change transactions results in a system with an acceptable risk level. For risks that continuously evolve and for which we make risk prognoses, we may need to identify treatments for which we make a plan for how and when to consecutively implement in the future in order to maintain an acceptable risk level.

3.2 Assessment Techniques

A risk assessment technique provides support for conducting one or more of the various activities of the overall risk assessment process. Brainstorming techniques such as Hazard and Operability (HazOp) studies [19], for example, are applicable and widely used during risk identification. Examples of well known techniques for risk



estimation are, for example, Failure Mode Effect Analysis (FMEA) [4] and event tree analysis (ETA) [17] for consequence estimation, fault tree analysis (FTA) [18] for probability estimation and consequence-likelihood matrix for risk level estimation. We refer to deliverable D5.1 [31] for an overview of state of the art risk assessment techniques for traditional risk assessments.

In this section we discuss the techniques and other artifacts that in particular are needed for risk assessments of changing and evolving systems. Existing, traditional techniques may, of course, be used for several activities also in the setting of changing analysis targets, but we will focus only on the additional techniques that are needed. When presenting the risk assessment process in the previous subsection, we identified some of the artifacts that are needed on order to conduct the various activities. In the following we discuss these more closely and also identify further artifacts and techniques that are needed.

The target description serves as a basis for the risk identification. When the target of analysis is a changing and evolving system, we also need to understand and represent it as such. Generally speaking and without considering the specific perspective on change, we therefore need the following artifact:

- A language for specifying a changing target of analysis.

Such a target description should facilitate the understanding of what is the target of analysis before and after any change transaction, and also of what is the change transaction or change process itself. The language should have a well defined syntax that describes the rules for making correct specifications, and it should have a well defined semantics that describes the precise meaning of specifications.

When conducting the risk identification on the basis of such a target description, the risks that are identified may also be changing. There is therefore a need for understanding and representing also the risks as such. We therefore need the following artifact:

- A language for specifying changing risks.

The risk models should facilitate the understanding of the risks, and the understanding of how risks change. The models should furthermore facilitate the various tasks of assessing the changing risks, and the language should have a well defined syntax and semantics.

Because many risks may be persistent under change, there is a need for techniques for how to identify these so as to not repeat assessment tasks from scratch when the results are the same. Such a technique relies on an identification and specification of the relations between the target description and the risk models. A notation or language for specifying these relations is then an artifact in itself that is needed in the process of risk assessment of changing and evolving systems:

- A language for specifying the relations between the target description and the risk models.

The specification of these relations should facilitate the tracing of changes in the target description to changes in the risk models: For a given change transaction, which risks are affected and therefore also may change? Conversely, the specification of these relations should facilitate the tracing of elements of the risk models to the target description: For a given risk, is the risk affected by the change transaction or not?



4 Target Description

In order to properly understand the target of analysis, we need to make a description of the target that can be well understood and that precisely describes the system, organization, enterprise or the like that is the subject of the risk analysis. The appropriate or suitable way of modeling the target may vary depending on the kind of target, the required level of details, the involved stakeholders, and so forth. In any case, the target of analysis should be modeled in a precise and unambiguous way in order to avoid misunderstandings and to ensure the correctness of the target description. The target description serves as the basis for the subsequent risk identification and risk assessment. Misunderstandings about the target or errors in the target description may therefore lead to erroneous risk assessment results.

While the chosen language for modeling and describing the target of analysis may vary, it is recommended to use a formal or semi-formal notation with a well-defined syntax that is well understood, such as the UML or similar. In this deliverable we do not assume any specific language for modeling the target. However, we need to assume that the chosen modeling language is suitable for representing the aspects of the world that we need to understand.

When considering risks in general and security risks in particular, what we need to understand and represent are the system actors that are within the scope of the analysis, the relevant system behavior involving these actors, as well as the events that may occur. Our notion of actor is very general and includes all the relevant entities that are involved in the system behavior. Such entities may be users and roles, devices and other components, applications and networks, and so forth. An actor may even be a system of other actors that together form a sub-system within the target of analysis. The system behavior is the interactions between actors, where an interaction can be described by the sequences or traces of events that occur in the interaction. A particular behavior can typically be conducted in different ways, and there are therefore generally several event traces describing a behavior. We refer to such a set of traces as a scenario. Describing behavior by trace sets allows underspecification of system behavior.

As our main concern is the representation of actors and the interactions between them, we assume a trace semantics for explaining the target models. A trace is a finite or infinite sequence of events, and we let \mathcal{H} denote the set of all traces.

The UML class diagram of Figure 7 can serve as a common meta-model for the modeling languages used for specifying target models or system models. The system model consists of a non-empty set of scenarios, a non-empty set of events and a non-empty set of actors. A scenario is a non-empty set of traces, which in turn is a non-empty set of ordered events. Each event is associated with exactly one actor.

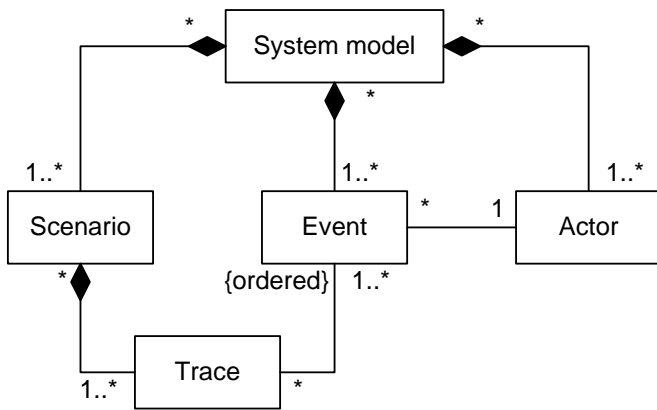


Figure 7 Meta-model for system models

As an example of a specific modeling language with trace semantics, we consider UML sequence diagrams. The sequence diagram of Figure 8 is a small fragment of the interactions involved in arrival management of ATM. It shows the interaction between the air traffic controllers of the Tactical Controller (TCC) and the Planner Controller (PLC), as well as their respective Controller Working Positions (CWPs). The diagram shows the parallel (**par**) composition of two interactions. The upper interaction is the sequence of the two events of transmitting the radar data by CWP_TCC, which we denote $!r$, and the reception of the same message, which we denote $?r$. The sequence of these two events is denoted by $\langle !r, ?r \rangle$. This is the only trace representing the upper interaction, so the trace set is the singleton set $\{\langle !r, ?r \rangle\}$. Similarly, the semantics of the lower interaction is the singleton set $\{\langle !t, ?t, !i, ?i \rangle\}$. The semantics of the sequence diagram is the parallel composition of the two trace sets, which yields all the interleavings where the ordering of the events of the operands is maintained. The two traces $\langle !t, ?t, !r, !i, ?i, ?r \rangle$ and $\langle !r, !t, ?r, ?t, !i, ?i \rangle$ are hence examples of traces in the resulting trace set.

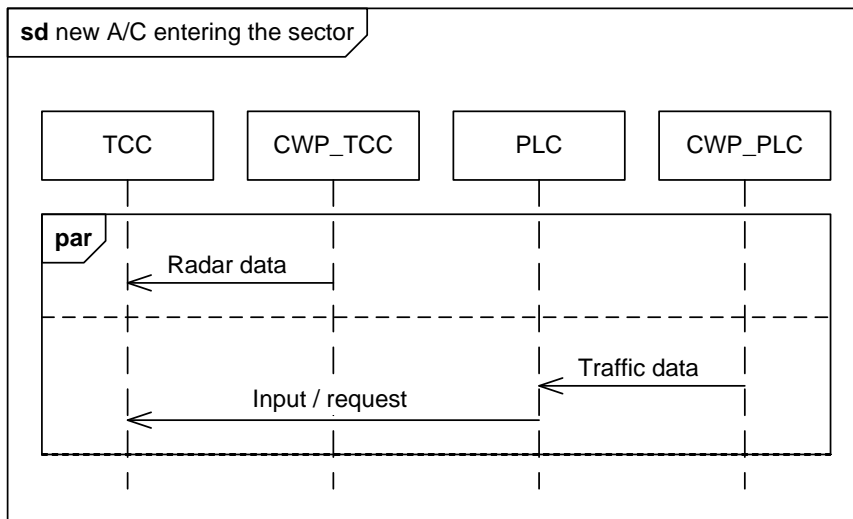


Figure 8 UML sequence diagram

5 Risk Modeling – Formal Foundation

Risk analysis involves the process of understanding the nature of risks and determining the level of risk [22]. Risk modeling refers to techniques that are used to aid the process of identifying, documenting and estimating likelihoods and consequences of unwanted incidents. A risk model is a structured way of representing an unwanted incident and its causes and consequences by means of graphs, trees or block diagrams [29]. In this section we introduce risk graphs as a means for risk modeling. A risk graph can be understood as a common abstraction of several well known and more specific approaches to risk modeling. By defining a formal semantics for risk graphs, we thereby also provide a risk model semantics that can be used to explain and reason about several approaches to risk modeling.

The overall aim of this section is to introduce a risk modeling language for the specification of changing risks. Once we have introduced the syntax and semantics of risk graphs, we generalize these to enable the modeling and reasoning about changing risks.

5.1 Risk Graphs

The introduction of risk graphs in this section is based on [7]. We introduce and exemplify the syntax before we present the semantics.

5.1.1 The Syntax of Risk Graphs

A risk graph consists of a finite, non-empty set of vertices (threat scenarios) and a finite set of directed relations (leads-to relations) between them.

Each vertex in a risk graph is assigned a set of likelihood values representing the estimated likelihood for the scenario to occur. The assignment of several likelihood values, typically a likelihood interval, represents underspecification of the likelihood estimate.

A relation from threat scenario t_1 to threat scenario t_2 means that t_1 may lead to t_2 . The relation from one threat scenario to another can also be assigned a set of likelihoods. These are conditional likelihoods that specify the likelihood for the former scenario to lead to the latter scenario when the former occurs. One threat scenario may lead to several other threat scenarios, so when operating with probabilities for likelihood estimates, the probabilities on the relations leading from a threat scenario may add up to more than 1. A risk graph may furthermore not be complete in the sense that a given threat scenario may lead to more scenarios than what is accounted for in the risk graph. The probabilities of the relations leading from a threat scenario may therefore add up to less than 1.

The meta-model for risks graphs is given in Figure 9. A vertex has an identifier, which is the description of the scenario, and a likelihood. A relation also has a likelihood, which is the conditional likelihood. The source of a relation is the vertex that leads to the vertex that is the target of the relation.




Figure 9 Meta-model for risk graphs

An example of a risk graph adapted from [7] is given in Figure 10. The risk graph describes two ways in which confidential information on a laptop may be exposed, either through theft or through the execution of malware. Data can be exposed through theft if the laptop is stolen, and the thief either has observed the login credentials or the laptop was not locked when stolen.

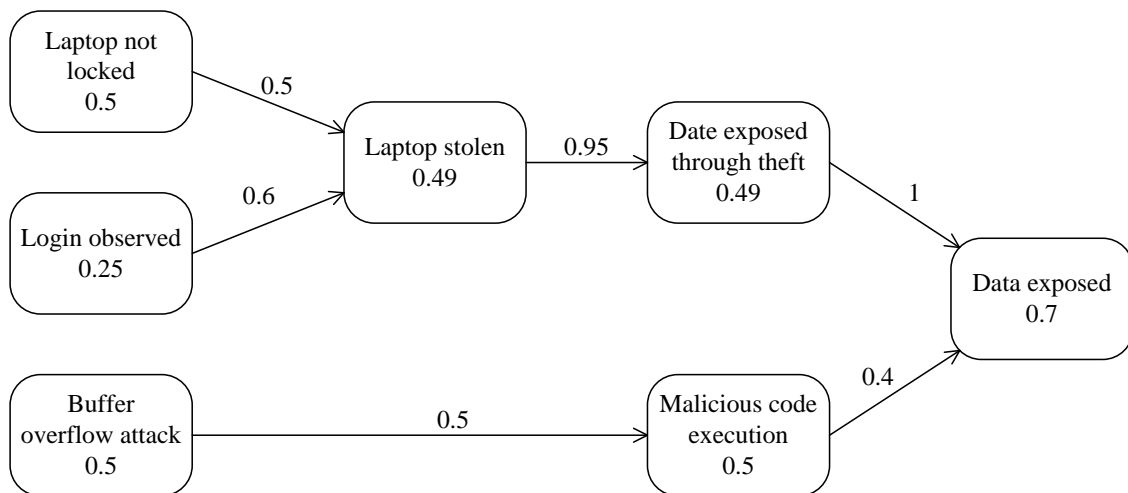


Figure 10 Example risk graph

There exist several modeling techniques that can be used for such structuring of scenarios and incidents, and for the reasoning about likelihoods of incidents. Robinson et al. [29] distinguish between three kinds of modeling techniques, namely trees, blocks and integrated presentation diagrams. The most common kinds of techniques are trees and integrated presentation diagrams. Some examples of state-of-the-art risk modeling techniques are Fault Tree Analysis (FTA) [18], Event Tree Analysis (ETA) [17], attack trees [32], cause-consequence diagrams [25][29], Bayesian networks [9] and CORAS threat diagrams [24]. The reader is referred to SecureChange deliverable D5.1 [31] for a presentation of these and other state-of-the-art risk modeling techniques.

Risk graphs can be understood as a common abstraction of these modeling techniques. A risk graph combines the features of fault trees and event trees, but does not require that causes of one scenario are connected by logical and-gates or or-gates. A risk graph may furthermore have more than one root vertex. Risk graphs can also be annotated with likelihoods on both vertices and relations, whereas in fault trees only

vertices are assigned likelihoods. The estimation of a likelihood of a vertex in a risk graph can therefore be supported by both the parent vertices and the relations from them. Another important difference between risk graphs and fault trees is that the former allow intervals of likelihoods to be assigned to vertices and relations, and thereby the underspecification of risks. Underspecification may be important in practical settings where it is difficult to come up with exact likelihoods.

5.1.2 The Semantics of Risk Graphs

Risk graphs are used for the purpose of documenting and reasoning about risks, particularly the documentation and analysis of threat scenarios and unwanted incidents and their likelihoods. The approach in [7] assumes that scenarios and their probabilities are represented by a probability space [10] on traces. As for the system models we let \mathcal{H} denote the set of all traces. Notice that this means that we can use system modeling techniques for specifying threat scenarios if desired. We let $\mathcal{H}_{\mathbb{N}}$ denote the set of all finite traces. A probability space is a triple $(\mathcal{H}, \mathcal{F}, \mu)$. \mathcal{H} is the sample space, i.e. the set of possible outcomes, which in our case is the set of all traces. \mathcal{F} is the set of measurable subsets of the sample space, and μ is a measure that assigns a probability to each element in \mathcal{F} . The semantics of a risk graph is statements about the probabilities of the trace sets that represent vertices or the composition of vertices. In other words, the semantics is a set of statements about the measure μ .

For composition of vertices, $v_1 \sqcap_1 v_2$ denotes the occurrence of both v_1 and v_2 where the former occurs before the latter. We let $v_1 \sqcup v_2$ denote the occurrence of at least one of v_1 and v_2 . A vertex is atomic if it is not of the form $v_1 \sqcap_1 v_2$ or $v_1 \sqcup v_2$. We use lower case v_i as the naming convention for arbitrary vertices, and upper case V_i as the naming convention for the set of finite traces representing the vertex v_i .

In order to formally define the semantics of risk graphs we need the auxiliary function $tr(_)$ that yields a set of finite traces from an atomic or combined vertex. Intuitively, $tr(v)$ are all possible traces that lead up to and through the vertex v , without continuing further. The function is defined as follows:

$$tr(v) \stackrel{\text{def}}{=} \mathcal{H}_{\mathbb{N}} \succ V \text{ when } v \text{ is an atomic vertex}$$

$$tr(v_1 \sqcap_1 v_2) \stackrel{\text{def}}{=} tr(v_1) \succ tr(v_2)$$

$$tr(v_1 \sqcup v_2) \stackrel{\text{def}}{=} tr(v_1) \cup tr(v_2)$$

where \succ the operator for sequential composition of trace sets, for example weak sequencing in UML sequence diagrams. Notice that the definition of the composition $v_1 \sqcap_1 v_2$ does not require v_1 to occur immediately before v_2 . The definition implies that $tr(v_1 \sqcap_1 v_2)$ includes traces from v_1 to v_2 via finite detours.

A probability interval P assigned to v , denoted $v(P)$, means that the likelihood of going through v is a value $p \in P$, independent of what happens before or after v . The semantics of a vertex is defined as follows:

$$[[v(P)]] \stackrel{\text{def}}{=} \mu_c(tr(v)) \in P$$

The expression $\mu_c(S)$ denotes the probability of any continuation of the trace set $S \subseteq \mathcal{H}$, and is defined as follows:

$$\mu_c(S) \stackrel{\text{def}}{=} \mu(S \succ \mathcal{H})$$

A probability interval P assigned to a (leads-to) relation $v_1 \rightarrow v_2$ means that the likelihood of v_2 occurring after an occurrence of v_1 is a value in P . This likelihood is referred to as the conditional likelihood. The semantics of the relation is defined as follows:

$$[[v_1 \xrightarrow{P} v_2]] \stackrel{\text{def}}{=} \mu_c(\text{tr}(v_1 \sqcap v_2)) \in \mu_c(\text{tr}(v_1)) \cdot P$$

Multiplication of two intervals $[p_i, p_j]$ and $[p_k, p_l]$ is defined by

$$[p_i, p_j] \cdot [p_k, p_l] \stackrel{\text{def}}{=} [p_i \cdot p_k, p_j \cdot p_l].$$

When multiplying an exact value p with an interval, the value p is replaced by the interval $[p, p]$.

We use D as naming convention for arbitrary risk graphs. Hence, D denotes a set of vertices v and relations $v_1 \rightarrow v_2$. We refer collectively to vertices and relations as elements, and use e as the naming convention for the latter.

The semantics $[[D]]$ of a risk graph is the conjunction of the expressions defined by the elements in D , formally defined as follows:

$$[[D]] \stackrel{\text{def}}{=} \bigwedge_{e \in D} [[e]]$$

A risk graph is said to be correct (with respect to the world or a specification of the relevant part of the world) if each of the conjuncts of $[[D]]$ is true. We say that D is inconsistent if it is possible to deduce \perp (False) from $[[D]]$.

5.2 Risk Graphs for Changing Risk

In this section we generalize the syntax and semantics of risks graphs as presented in [7] and summarized in Section 5.1 to risk graphs with change. We introduce and exemplify the syntax before we present the semantics.

5.2.1 The Syntax of Risk Graphs with Change

In order to support the modeling of changing risks we need to generalize risk graphs to allow the simultaneous modeling of risks both before and after the implementation of some given system changes or change requirements. For this purpose we extend the risk graph notation of vertices and relations to three kinds of vertices and three kinds of relations, namely *before*, *after* and *before-after*. When an element (vertex or relation) is of kind *before* it represents risk information before the changes, when it is of kind *after* it represents risk information after the changes, and when it is of kind *before-after* it represents risk information that holds both before and after the changes.



Table 1 gives an overview of the language constructs and the naming conventions we use for referring to them. The symbols written in **bold face** and the arrows denote the specific language constructs, whereas v denotes an arbitrary vertex of any kind.

Variable	Diagram construct
v	Vertex before and after
vb	Vertex before
va	Vertex after
v	Vertex
$v_1 \rightarrow v_2$	Relation before and after
$v_1 \rightarrow_b v_2$	Relation before
$v_1 \rightarrow_a v_2$	Relation after

Table 1 Naming conventions

As before, vertices can be assigned likelihoods, and relations can be assigned conditional likelihoods. Table 2 gives an overview of the various ways of specifying likelihoods. The before-after elements can be assigned a pair of likelihoods, the former specifying the likelihood before and the latter specifying the likelihood after. Notice that any of the likelihoods can be undefined, in which case they are completely underspecified.

Likelihood specification	Interpretation
$v(P_1 P_2)$	v occurs with likelihood P_1 before, and v occurs with likelihood P_2 after
$vb(P)$	vb occurs with likelihood P before
$va(P)$	va occurs with likelihood P after
$v_1 \xrightarrow{P_1 P_2} v_2$	v_1 leads to v_2 with conditional likelihood P_1 before, and v_1 leads to v_2 with conditional likelihood P_2 after
$v_1 \xrightarrow{P}_b v_2$	v_1 leads to v_2 with conditional likelihood P before
$v_1 \xrightarrow{P}_a v_2$	v_1 leads to v_2 with conditional likelihood P after

Table 2 Denoting likelihoods

A meta-model for the risk graphs with change is given in Figure 11. Vertices and relations are given a mode attribute, where the mode is one of *before*, *after* and *before-after*.



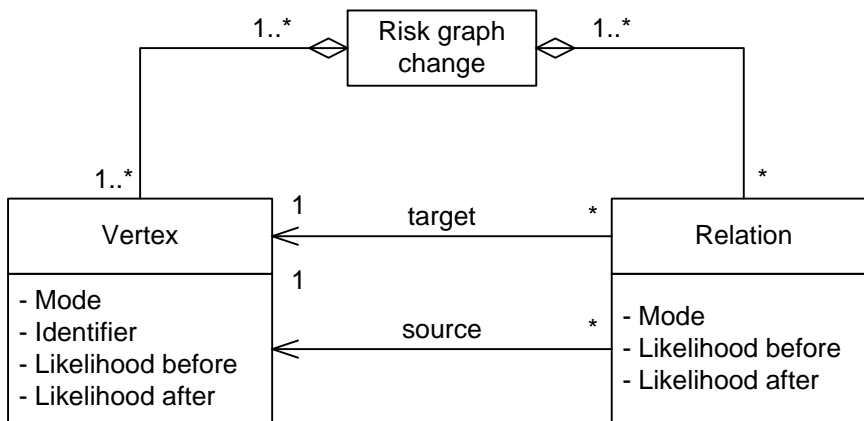


Figure 11 Meta-model for risk graphs with change

Explicitly distinguishing between the three kinds or modes of the risk graph elements is useful when modeling changing risks, because it allows the explicit modeling and documentation of risks that disappear after change, risk that emerge after change, and risks that are persistent under change. Furthermore, by operating with pairs of likelihoods, we can explicitly model the changing risk levels of risks that are present both before and after the changes.

We can, nevertheless, understand risk graphs with change as the combination of two risk graphs, one representing risks before changes and one representing risks after changes. More formally, this means that the *before-after* language constructs are syntactic sugar for specifying one element in the before risk graph and one element in the after risk graph.

Such a combination of two risks graphs into one representation imposes some restrictions on the vertices and relations that are not captured by the meta-model. For example, if the mode of a vertex is *before*, it cannot have a likelihood after. And a before vertex cannot be related to an after vertex, as the two do not occur at the same time. A before-after vertex, on the other hand, can be related to all three kinds of vertices. The additional restrictions on the risk graphs with change are the following:

- If the mode of a vertex is *before*, the attribute “Likelihood after” does not apply.
- If the mode of a vertex is *after*, the attribute “Likelihood before” does not apply.
- If the mode of a relation is *before*, the attribute “Likelihood after” does not apply.
- If the mode of a relation is *after*, the attribute “Likelihood before” does not apply.
- If the mode of the target of a relation is *before*, the mode of the relation is *before*.
- If the mode of the target of a relation is *after*, the mode of the relation is *after*.
- If the mode of the source of a relation is *before*, the mode of the relation is *before*.

- If the mode of the source of a relation is *after*, the mode of the relation is *after*.

Notice that by these restrictions we can determine the mode of the relation from the modes of the source and target of the relation. When it is clear from the context we therefore represent all three kinds of relations by $v_1 \rightarrow v_2$.

An example of a risk graph with change is given in Figure 12. Vertices of kind *before* are represented by dashed, shaded rounded rectangles, whereas vertices of kind *after* are represented by solid, white rounded rectangles. Vertices of kind *before-after*, on the other hand, are represented by the two-layer rounded rectangles to convey the combination of the two other kinds of vertices.

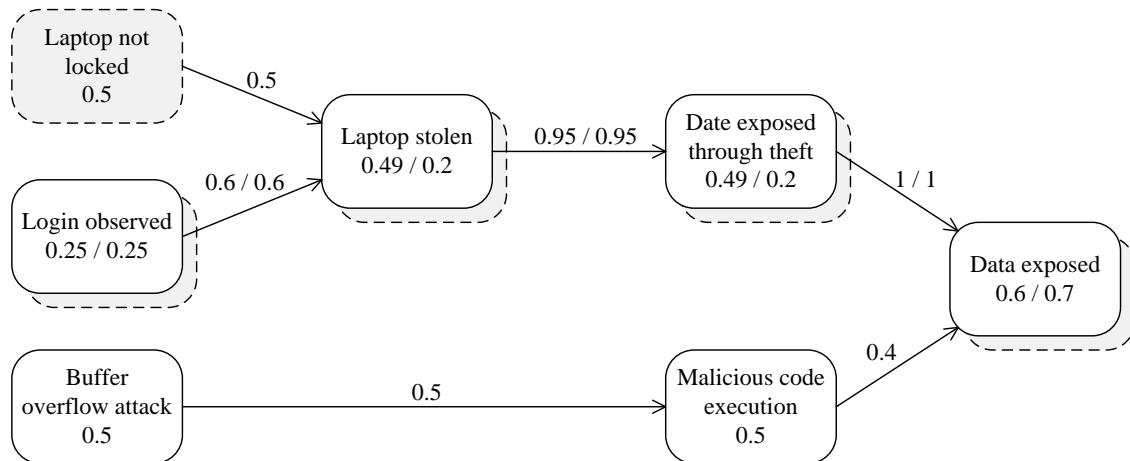


Figure 12 Example risk graph with change

As mentioned above, since the vertices of kind *before-after* represent threat scenarios that occur both before and after the changes, they are assigned a pair of likelihoods. The former denotes the likelihood before, and the latter denotes the likelihood after.

Observe that there is no distinction between the three kinds of relations in the graphical representation. The kind of the relation, however, can be determined by the source and/or target of the relation in question. It is furthermore only the relations of kind *before-after* that are assigned a pair of likelihoods.

Because the before-after elements can be understood as syntactic sugar for separate representations of before elements and after elements, a risk graph with change can be translated to a pair of regular risk graphs that together are equivalent to the former.

The risk graph of Figure 13 shows the lower layer of the risk graph with change from Figure 12. This shows the documentation of the risks before the changes. The risk graph of Figure 14 shows the upper layer, i.e. the documentation of the risks after the changes.

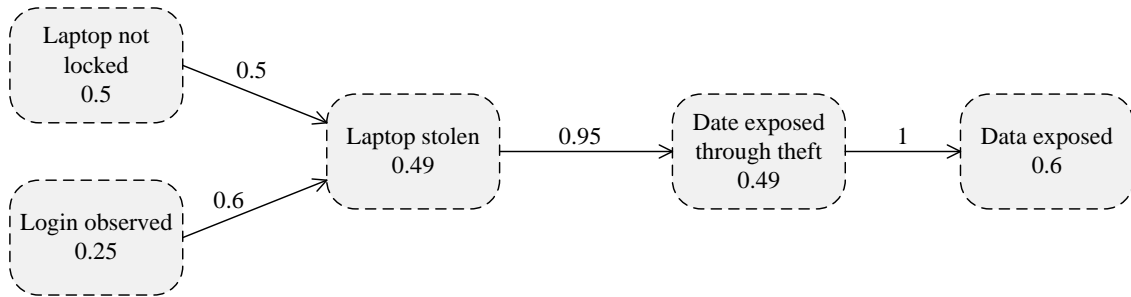


Figure 13 Risk graph before change

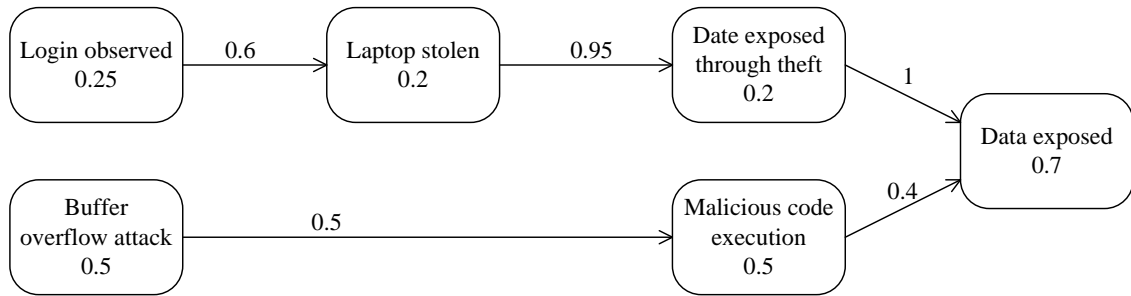


Figure 14 Risk graph after change

A regular risk graph consists of a finite non-empty set of vertices and a finite set of relations between them. That is to say, a risk graph is a set D of elements e . When generalizing risk graphs to risks graphs with change, they are instead represented by a pair (D_b, D_a) of sets of elements, the former consisting of the vertices and relations of kind *before* and the latter consisting of vertices and relations of kind *after*.

Since we are operating with vertices and relations of kind *before-after* as language element of their own, we also allow the representation of risk graphs with change as a single set D of vertices and relations, where each element is of one of the kinds *before*, *after* and *before-after*. This single set of elements is then syntactic sugar for the equivalent representation of a pair of sets of elements. For such a combined representation D we use the functions $before(_)$ and $after(_)$ to filter the combined risk graph with respect to the elements of kind *before* and *after*, respectively. The following define the function $before(_)$ for singleton sets of elements.

$$before(\{v(P_1 P_2)\}) \stackrel{\text{def}}{=} \{vb(P_1)\}$$

$$before(\{vb(P)\}) \stackrel{\text{def}}{=} \{vb(P)\}$$

$$before(\{va(P)\}) \stackrel{\text{def}}{=} \emptyset$$

$$before(\{v_1 \xrightarrow{P_1 P_2} v_2\}) \stackrel{\text{def}}{=} \{v_1 \xrightarrow{P_1} v_2\}$$

$$before(\{v_1 \xrightarrow{P} v_2\}) \stackrel{\text{def}}{=} \{v_1 \xrightarrow{P} v_2\}$$

$$\text{before}(\{v_1 \xrightarrow{P} v_2\}) \stackrel{\text{def}}{=} \emptyset$$

The filtering of a risk graph with change D with respect to the *before* elements is then defined as follows:

$$\text{before}(D) \stackrel{\text{def}}{=} \bigcup_{e \in D} \text{before}(e)$$

The definition of the function $\text{after}(_)$ is symmetric. For a risk graph with change D of elements of the three different kinds, the representation as a pair of elements of kind *before* and elements of kind *after* is then given by $(\text{before}(D), \text{after}(D))$.

5.2.2 The Semantics of Risk Graphs with Change

Given the syntax of risk graphs with change as defined above, we can define the semantics as a straightforward generalization of the semantics of regular risk graphs defined in [7] as summarized in Section 5.1.2.

The semantics $[[D_b, D_a]]$ of a risk graph with change is defined as follows:

$$[[D_b, D_a]] \stackrel{\text{def}}{=} [[D_b]] \wedge [[D_a]]$$

For a combined representation D of a risk graph with change, the semantics is defined as follows:

$$[[D]] \stackrel{\text{def}}{=} [[\text{before}(D), \text{after}(D)]]$$

5.3 Reasoning about Likelihoods in Risk Graphs

In this section we introduce rules for calculating probabilities of vertices in risk graphs, and we provide guidelines for consistency checking probabilities that are assigned to risk graphs.

5.3.1 Rules for Likelihood Calculation

The rules we introduce are of the following form:

$$\frac{R_1 \quad R_2 \quad \dots \quad R_i}{C}$$

We refer to $R_1 \dots R_i$ as the premises and to C as the conclusion. The interpretation is that if the premises are valid, so is the conclusion.

The first rule is referred to as the *relation rule*, and captures the conditional likelihood semantics of a risk graph relation. For a vertex v_1 that leads to v_2 , the vertex $v_1 \sqcap_1 v_2$ denotes the occurrences of v_2 that happen after an occurrence of v_1 .



Rule 1 (Relation). If there is a direct relation from v_1 to v_2 , we have:

$$\frac{v_1(P_1) \quad v_1 \xrightarrow{P_2} v_2}{(v_1 \sqcap v_2)(P_1 \cdot P_2)}$$

The second rule is referred to as the *mutual exclusive vertices rule*, and yields the probability of either v_1 or v_2 occurring when the two vertices are mutually exclusive:

Rule 2 (Mutually exclusive vertices). If the vertices v_1 and v_2 are mutually exclusive, we have:

$$\frac{v_1(P_1) \quad v_2(P_2)}{(v_1 \sqcup v_2)(P_1 + P_2)}$$

Addition of two intervals $[p_i, p_j]$ and $[p_k, p_l]$ is defined by

$$[p_i, p_j] + [p_k, p_l] \stackrel{\text{def}}{=} [p_i + p_k, p_j + p_l].$$

When adding an exact value p with an interval, the value p is replaced by the interval $[p, p]$.

The third rule is referred to as the *statistically independent vertices rule*, and yields the probability of either v_1 or v_2 occurring when the two vertices are statistically independent.

Rule 3 (Statistically independent vertices). If vertices v_1 and v_2 are statistically independent, we have:

$$\frac{v_1(P_1) \quad v_2(P_2)}{(v_1 \sqcup v_2)(P_1 + P_2 - P_1 \cdot P_2)}$$

Subtraction of two intervals $[p_i, p_j]$ and $[p_k, p_l]$ is defined by

$$[p_i, p_j] - [p_k, p_l] \stackrel{\text{def}}{=} [p_i - p_k, p_j - p_l].$$

When one of the operands is an exact value p , it is replaced by the interval $[p, p]$. Notice that subtraction of intervals by this definition is used only in the context of the rules presented in this section. The definition ensures that every probability in $P_1 + P_2 - P_1 \cdot P_2$ can be obtained by selecting one probability from P_1 and one probability from P_2 .

As a small example of probability calculation, consider the risk graph in Figure 15. We let ll abbreviate *Laptop not locked*, lo abbreviate *Login observed* and ls abbreviate *Laptop stolen*. The risk graph then consists of the three vertices $ll(0.5)$, $lo(0.25)$ and $ls(p)$, where p is the probability we need to calculate, and the two relations $ll \xrightarrow{0.5} ls$ and $lo \xrightarrow{0.6} ls$.

By Rule 1 we calculate $(ll \sqcap ls)(0.25)$ and $(lo \sqcap ls)(0.15)$. Assuming that ll and lo , as well as $ll \sqcap ls$ and $lo \sqcap ls$, are statistically independent, we use Rule 3 to calculate $((ll \sqcap ls) \sqcup (lo \sqcap ls))(0.3625)$ by $0.25 + 0.15 - 0.25 \cdot 0.15$.






Figure 15 Probability calculation in risk graphs

Assuming that the likelihood estimates in Figure 15 are correct, there is still one issue to consider before we can conclude about the likelihood of the vertex *Laptop stolen*. The issue is whether or not the risk graph is complete. If the risk graph is complete, the graph shows all the possible ways in which the laptop can be stolen. In that case we have that $ls = (ll \sqcap_l ls) \sqcup (lo \sqcap_l ls)$ and that 0.3625 is the correct likelihood of this vertex. If the risk graph is incomplete, there may be further scenarios that can lead to the theft of the laptop. In that case we only know that 0.3625 is the lower bound of the probability p .

5.3.2 Guidelines for Consistency Checking Likelihoods

Consistency checking of risk models is important, as it is a useful means for detecting errors or misunderstandings of the risk estimates that are documented during a risk assessment. The basis for the consistency checking is the likelihood values that are already assigned to the vertices and relations of a risk graph.

The guidelines for consistency checking depend on whether the risk graph in question is complete, and whether the likelihoods are given as exact probabilities or as probability intervals. The guidelines are given in Table 3.

Exact values in complete diagrams

Assigned value: $v(p)$

Calculated value: $v(p')$

Consistency check: $p = p'$

Exact values in incomplete diagrams

Assigned value: $v(p)$

Calculated value: $v(p')$

Consistency check: $p \geq p'$

Intervals in complete diagrams

Assigned interval: $v([p_i, p_j])$

Calculated interval: $v([p_i', p_j'])$

Consistency check: $[p_i', p_j'] \subseteq [p_i, p_j]$ or, equivalently $p_i \leq p_i'$ and $p_j \geq p_j'$

Intervals in incomplete diagrams

Assigned interval: $v([p_i, p_j])$

Calculated interval: $v([p_i', p_j'])$

Consistency check: $p_j \geq p_j'$

Table 3 Guidelines for consistency checking probabilities

As an example of consistency checking, consider the risk graph in Figure 10, assuming first that the graph is complete. By the example from Section 5.3.1, we know that the probability of the vertex *Laptop stolen* is 0.3625 given the vertices and relations that lead to this vertex. The assigned probability 0.49 is therefore inconsistent with the preceding probability estimates. This indicates that the estimates must be reconsidered. Consistency can be restored by changing the probability of *Laptop stolen* to 0.3625 or by changing the probability of several of the vertices and relations in question.

If the evidence for the probability estimates in questions is very strong, it may on the other hand indicate that the assumption of the risk graph being complete is erroneous. Discarding this assumption gives the consistency requirement that the assigned probability 0.49 must be greater than or equal to the calculated probability 0.3625, which indeed it is.

5.3.3 Reasoning about Likelihoods in Risk Graphs with Change

Given the generalization of the syntax and semantics of risk graphs to risk graphs with change presented in Section 5.2, the rules and guidelines for reasoning about likelihoods in risk graphs with change can be straightforwardly applied. The only constraint is that the rules and guidelines must be applied separately for the *before* and *before-after* elements on the one hand, and the *after* and *before-after* elements on the other hand.

For example, when reasoning about the *before-after* vertex *Data exposed* in Figure 12, we address on the one hand the *before* likelihood 0.6, and consider the *before* layer of the vertices that may lead to it. On the other hand we address the *after* likelihood 0.7, and consider the *after* layer of the preceding vertices. Given the syntax of the risk graph with change, the two layers can be easily kept separate during the likelihood reasoning and estimation. If desired, however, the separate layers can be extracted and presented separately before the likelihood reasoning, as depicted in Figure 13 and Figure 14.

6 Relating Risk Model to Target Description

The purpose of identifying and documenting the relations between the target description and the risk models is to facilitate the tracing between system elements and risk model elements. Such traceability supports the tracing of system changes to changes in the risk picture. On the one hand, this will in turn support techniques for identifying the parts of the risk picture that are affected by changes to a specific part of the system and therefore need to be reassessed, as well as identifying the parts of the risk picture that are not affected and therefore are valid also after the changes. On the other hand, this will conversely support tracing in the other direction and determining whether or not a given part of a risk picture is affected by the system changes.

Two of the key artifacts that are used during a risk assessment process are the languages for target modeling and risk modeling. The target models are a core part of the overall target description, and documents the events, scenarios and actors that are the subject for the risk assessment. In this section we assume these two artifacts as described in Section 4 and Section 5, respectively. The meta-model for the target models (or system models) is given in Figure 7, and the meta-model for risk models (or risk graphs) is given in Figure 9. Given these artifacts, we need a third artifact, namely a trace model, for specifying the relationships between the former two.

In order to specify the relation between a target model element and a risk model element, each of the elements must have a unique identifier. For some modeling languages, each element already has an identifier, and if these are unique they can be used for the trace modeling. In the general case, however, we may need to index target model elements and/or risk model elements before the trace model can be specified. In this section we assume that the risk models have unique identifiers.

In the following we first introduce an approach to and a format for indexing the target model in an adequate way. Thereafter we introduce the trace model artifact for specifying and documenting the relations between target model and risk models. Finally, we extend the risk modeling syntax with a separate construct for specifying the relations from risk model elements to target model elements. The purpose of the latter is to visualize in the risk models the tracing from risk models to target models, so as to support and facilitate the identification of the change affected risks during the risk assessment process.

6.1 Indexing of Target Model

There are three basic categories of elements in the target models, as defined by the meta-model in Figure 7, namely actors, events and scenarios. A straightforward way of indexing these elements is simply to give each of them a unique index or identifier. However, in order to enable a better basis for reasoning about the kinds of system changes that affects the risks, the specification of the relations between the target model and the risk models should include information about the categories of the target elements. We therefore attach to each target model index the category of the target element in question.



The indexing of the target elements is conducted on the target models at the level of the syntax. With this approach, tool support for the automation of the indexing can be developed, where the indexing is conducted by parsing the specification.

As an example of the indexing of the target model, consider the fragment of the ATM target specification given in Figure 16. This UML sequence diagram is a part of the full documentation of the ATM risk assessment that is given in the appendix of Section 14. Because the purpose of this section is to exemplify the modeling artifacts we do not explain in detail the ATM target of analysis or the risk assessment results. For now we only mention that the change requirement that is addressed is the organization level change of ATM with the introduction of the Arrival Manager (AMAN) tool, and the security properties that are addressed are information protection and information provision.

The sequence diagram shows a part of the ATM arrival management process before the changes, i.e. before the introduction of the AMAN and the Automatic Dependent Surveillance Broadcasting (ADS-B). In particular, the diagram shows the arrival management task T1 – Controlling the A/C in the sector. This task is conducted in parallel by the two Air Traffic Controllers (ATCOs) of Tactical Controller (TCC) and Planner Controller (PLC). The TCC and the PLC operates their own Controller Working Position (CWP). As shown by the diagram, the TCC and the PLC conducts more or less the same activities during this task. The difference is that the PLC works on a wider scope in time and space than the TCC, so as to observe and plan ahead and to support the TCC.

There are four actors in this specification, namely TCC, CWP_TCC, PLC and CWP_PLC. There are moreover seven scenarios, one of them specified by the sequence diagram itself, namely T1 – Controlling the A/C in the sector. The remaining six are the various scenarios specified by the **ref** construct. Finally, there are four events, namely the two transmissions of messages by the PLC (send events) and the two receptions of the messages by the TCC (receive events).

The indexing of target models is specified and documented in the table format exemplified in Table 4. Each index is specified in a separate row of four columns. The first column is the unique id. The second column is the name of the target model element as specified in the model. The third column is the category of the element in question, where the category is one of *scenario*, *actor* and *event*. Notice that for the events, the name is in this example prefixed with an exclamation mark (!) or a question mark (?) to denote the kind of the event, i.e. whether it is a send or receive event, respectively. The fourth column is an entry that optionally can be manually filled by the user with an informal description of the element in question. If tool support is provided for making the index, the three first columns should be filled in automatically, where the ids are tool generated.

The indexing example shown in Table 4 is the indexing of the part of the target model depicted in Figure 16. Notice that the level of details in the indexing is determined by the level of details in the target model. For example, in this sequence diagram the TCC task of monitoring the flights in the sector is represented by an interaction use (the **ref** construct) that hides the events that occur in the interaction. If the specific events were specified directly into this sequence diagram instead of the interaction use, the indexing would also have these events instead of the corresponding scenario. This ensures that the level of details that is chosen in the target specification, and therefore

is held as suitable for the objectives of the risk assessment, is maintained in the indexing.

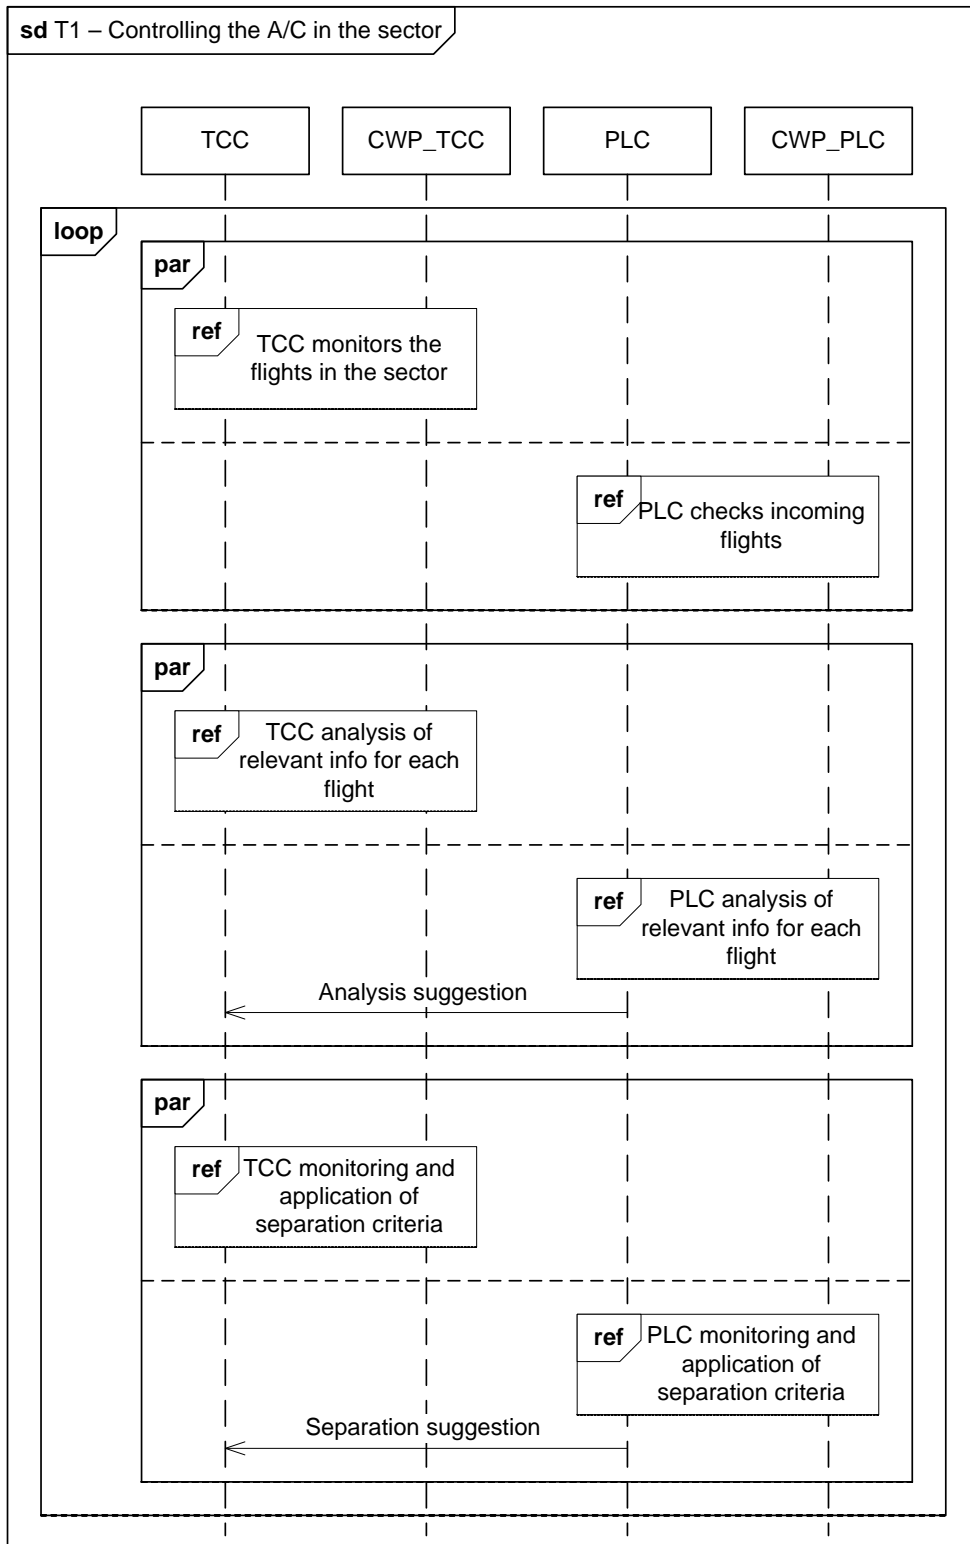


Figure 16 Example part of ATM target model

ID	Name	Category	Description
t1	TCC	Actor	Air Traffic Controller
t2	CWP_TCC	Actor	Controller working position of TCC
t3	PLC	Actor	Air Traffic Controller
t4	CWP_PLC	Actor	Controller working position of PLC
t5	T1 – Controlling the A/C in...	Scenario	Task 1 of arrival management
t6	TCC monitors the flights in...	Scenario	
t7	PLC checks incoming flights	Scenario	
t8	TCC analysis of relevant info...	Scenario	
t9	PLC analysis of relevant info...	Scenario	
t10	TCC monitoring and...	Scenario	
t11	PLC monitoring and...	Scenario	
t12	! Analysis suggestion	Event	From PLC to TCC
t13	? Analysis suggestion	Event	
t14	! Separation suggestion	Event	From PLC to TCC
t15	? Separation suggestion	Event	

Table 4 Example indexing of target model

The target models are in some cases very extensive and detailed, and in many cases it may not be necessary to do a complete indexing. As the purpose of the indexing is to provide a basis for tracing changes between the target model and the risk models, the index should be sufficiently rich and detailed to enable the detection of all changes that may affect the risk picture. The sufficient level of details must in each case be determined by the risk analysts.

6.2 Specification of Relations between Target Model and Risk Model

The identification of the relationships between the target system and the risk models is a manual analysis task that is conducted during the risk assessment process. In this section we introduce a separate artifact to support the specification and documentation of the relations, namely a trace modeling artifact.

When specifying a trace model we assume that we already have an indexed target model and a risk model of elements with unique identifiers. (If the risk model elements do not have unique identifiers, also these must be indexed.)

The trace model is of a table format that allows the tracing from target model elements to risk model elements, and vice versa. In order to explain and motivate the format, we introduce it by referring to an example.



The sequence diagram of Figure 16 shows a part of the ATM target description that served as a basis for a risk assessment of the arrival management in ATM. The risk graph of Figure 17 shows a fragment of the risk identification that is fully documented in the appendix of Section 14.

For each of the three risk graph vertices, we need to identify and document the relevant parts of the target model. The risk graph vertices *Malfunctioning of radar* and *Loss of radar signal in MRT* (Multi Radar Tracking) are both related, for example, to radar and surveillance in the target system. Considering the part of the target model shown in Figure 16, we need to address the vertex *Monitoring of A/C in the sector fails*.




Figure 17 Example risk model in ATM before changes

The relations between the exemplified part of the target model and the exemplified risk model are given in Table 5. The first row refers to the target model elements by their indexes as specified in Table 4, and the second row refer to the risk model elements by their unique identifier. The table represents an excerpt that only shows the relations including the risk graph vertex in question.

Target index	Risk element identifier
...	...
t1	Monitoring of A/C in the sector fails
t2	Monitoring of A/C in the sector fails
t3	Monitoring of A/C in the sector fails
t4	Monitoring of A/C in the sector fails
t5	Monitoring of A/C in the sector fails
t6	Monitoring of A/C in the sector fails
t7	Monitoring of A/C in the sector fails
t10	Monitoring of A/C in the sector fails
t11	Monitoring of A/C in the sector fails
...	...

Table 5 Examples of relations between target model and risk model

From a pragmatic point of view, there are two obvious shortcomings of the table format of the trace models as given in Table 5. For end-users and other stakeholders to make efficient use of the trace model during risk assessments, the trace model should convey information about the relations in an intuitive way. The use of possibly tool generated indexes for the target model elements is not intuitively informative. Furthermore, in some cases several target model elements are logically understood as

a whole. Without some means of grouping several relations into one compound relation, such structures of the target model will be obscured.

To mitigate this we introduce a third row to the table representing the trace model for tagging the specified relations. The grouping of several relations is then conducted by inserting the same tag on several rows. The name of the tag will be chosen by the end-user, and should be a unique name that conveys intuitive information about the relations (i.e. the pairs of target index and risk element identifiers) that are grouped. Table 6 shows the table format for the trace model.

The table consists of the three columns of target index, risk element identifier and tag. In this example the target elements of the TCC, the PLC and their CWP's (indexed t1 through t4) are combined by the tag Sector team. This is a convenient and adequate grouping of elements, because in the ATM setting these four actors actually form what is referred to as precisely a sector team. (See Figure 71 and Figure 86 of the ATM target specification in the appendix of Section 14.) Instead of referring to the less intelligible indexes t1 through t4 when relating the risk graph vertex *Monitoring of A/C in the sector fails*, the end-user will relate the vertex to the target elements by referring to the tag Sector team.

Target index	Risk element identifier	Tag
...
t1	Monitoring of A/C in the sector fails	Sector team
t2	Monitoring of A/C in the sector fails	Sector team
t3	Monitoring of A/C in the sector fails	Sector team
t4	Monitoring of A/C in the sector fails	Sector tem
t5	Monitoring of A/C in the sector fails	Task T1
t6	Monitoring of A/C in the sector fails	A/C monitoring
t7	Monitoring of A/C in the sector fails	A/C monitoring
t10	Monitoring of A/C in the sector fails	A/C monitoring
t11	Monitoring of A/C in the sector fails	A/C monitoring
...

Table 6 Example trace model

More formally, the trace model is a set of tuples (t_{id}, r_{id}, t) of a target model index, a risk model index and a tag. The table format shown in Table 6 is one possible data structure for representing the trace model. With this choice of data structure, there should for reasons of usability be provided tool support with sort & filter, as well as find & select functionality. A database could also be an adequate way of organizing the trace model.

6.3 Visualization of Relations to Target Model in Risk Models

During a risk assessment of a changing system, there is a need continuously to keep track of the traceability between the target description and the risk models. In order to facilitate the traceability during the activities of risk identification, risk estimation and risk evaluation, the risk assessment should be supported by means and techniques for representing the trace model in an intuitive way that is easily comprehensible. Whereas the representation of the set of tuples (t_{id}, r_{id}, t) of a trace model in a table format is adequate for documentation purposes, it may not be suitable for example in a workshop setting with structured brainstorming.

To mitigate this, we extend the risk graph notation with a language construct for explicitly specifying the relations to the target description. The construct is used for annotating risk graphs with links to the target description, where each link refers to a subset of the relations that are documented in the trace model. We first define the extension for traditional risk graphs, and thereafter generalize it to risk graphs with change.

6.3.1 Trace Model in Risk Graphs

We use the risk graph example depicted in Figure 17 and the trace model depicted in Table 6 to exemplify the visualization of the relations to target models in risk graphs. The risk graph in Figure 18 shows the visualization as annotations on risk graph vertices. The annotations are in the form of rectangles with a description of the part of the target model that is related to the vertex in question.

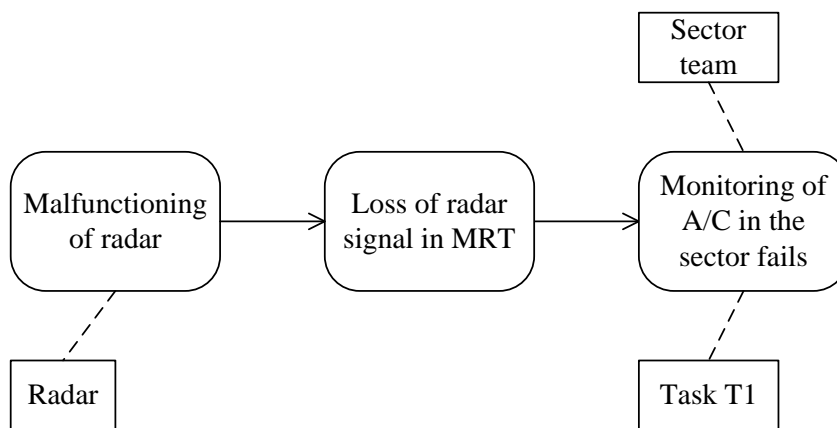


Figure 18 Risk graph with tracing to target model

More precisely, the description in each annotation is the name of a tag from one or more entries in the trace model. In Figure 18, for example, the annotation *Sector team* on the vertex *Monitoring of A/C in the sector fails* refers in Table 6 to the relations that are documented in the rows that are tagged *Sector team*. The annotation *Task T1* is also a tag from the same example, whereas the annotation *Radar* is assumed to be a further tag specified elsewhere in the trace model.

Notice, importantly, that since the initial vertex *Malfunctioning of radar* in the risk graph example is related to the radar, so are the other vertices in the graph that can be reached from the initial vertex via the risk graph relations. This means that not only this initial scenario may be affected by target system changes that involve the radar, but also each of the subsequent relations. We may optionally annotate also one or more of the subsequent vertices with the same tag *Radar*, but due to the dependencies in the risk graph such annotations are redundant. In large risk models, it may still be useful to repeat annotations in order to make the relations to the target model more visible.

The UML class diagram of Figure 19 shows the meta-model for risk graphs extended with the construct *Target element* and the relation *Trace relation*. The target element is a reference to a part of the target model and has an identifier. The identifier should be a tag from a trace model that has already been specified. Otherwise, a trace model can later be specified based on these risk graph annotations. The trace relation is the relation from risk graph vertices to the target element annotations, and is shown as dashed lines in the concrete risk graphs.

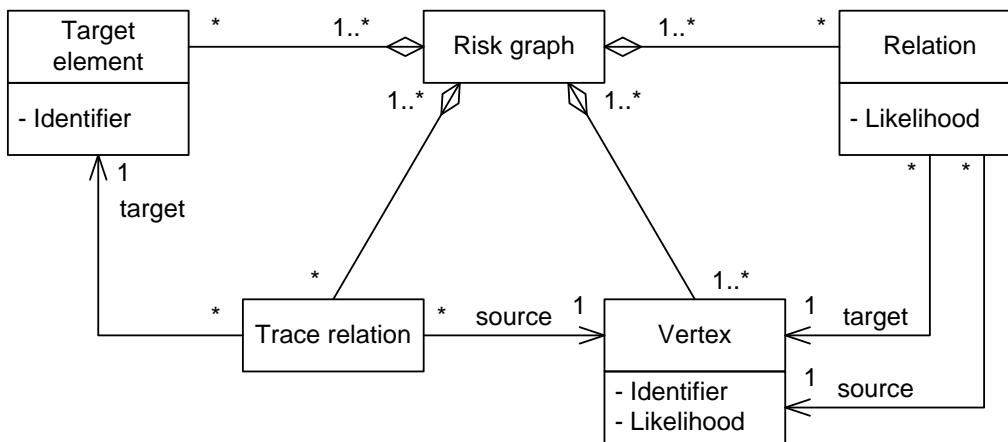


Figure 19 Meta-model for risk graphs extended with trace annotation

We understand the target element construct as a mere annotation on risk graphs, and these annotations are therefore not part of the formal semantics of risk graphs. The semantics as defined in Section 5.1 is therefore maintained. If a trace model TM has been specified, each target element with tag t as an annotation on a risk graph element with identifier r_{id} is then a reference to the set of elements $\{(t_{id}, r_{id}, t) \mid (t_{id}, r_{id}, t) \in TM\}$.

As mentioned above, the relations from target model elements to risk graph vertices propagate through risk graphs via dependencies. In order to precisely capture this, we introduce a formal notion of dependencies in risk graphs. For this purpose, we first introduce the notions of well-formed risk graphs and interfaces between risk graphs. The notions of dependencies, well-formedness and interfaces in risk graphs are previously presented in [7].

For a risk graph D of vertices and relations to be well formed, we require that if a relation is contained in D , so are the source and target vertices of the relation:

$$v_1 \rightarrow v_2 \in D \Rightarrow v_1 \in D \wedge v_2 \in D$$

When we speak of dependencies in a risk graph D , we speak of dependencies between sub-graphs D_1 and D_2 of D . The notion of interface is defined for such sub-graphs that do not necessarily fulfill the well-formedness requirement. Given two sub-graphs D_1 and D_2 , we let $i(D_1, D_2)$ denote D_1 's interface towards D_2 . The interface is obtained from D_1 by keeping only the vertices and relations that D_2 depends on directly, formally defined as follows:

$$i(D_1, D_2) \stackrel{\text{def}}{=} \{ v_1 \in D_1 \mid \exists v_2 \in D_2: v_1 \rightarrow v_2 \in D_1 \cup D_2 \} \cup \{ v_1 \rightarrow v_2 \in D_1 \mid v_2 \in D_2 \}$$

Given the notion of interface between sub-graphs, we can define the notion of dependency. For this purpose we introduce the relation $D_1 \ddagger D_2$ which means that D_2 does not depend on any vertex or relation in D_1 . In turn, this means that D_1 and D_2 have no common vertices or relations, and that D_1 has no interface towards D_2 . In the definition, we assume a risk graph D with sub-graphs D_1 and D_2 such that $D_1 \cup D_2 = D$:

$$D_1 \ddagger D_2 \stackrel{\text{def}}{=} D_1 \cap D_2 = \emptyset \wedge i(D_1, D_2) = \emptyset$$

Observe that \ddagger is not commutative, i.e. $D_1 \ddagger D_2$ does not imply $D_2 \ddagger D_1$.

Identifying dependencies in risk assessments of changing systems is important, since we need to identify the possibilities for changes to propagate through risk models; it is only by understanding how changes propagate that we can completely determine which parts of the risk picture that are affected by changes and therefore need to be reassessed. This applies, however, not only to the risk models, but also of the target system models. The latter issue is not addressed here, since modeling and reasoning about target systems is outside the scope of this deliverable.

6.3.2 Trace Model in Risk Graphs with Change

The visualization of the trace model in risk graphs with change is exemplified in Figure 20. The risk graph shows that malfunctioning of radar may lead to failure of A/C monitoring both before and after the system changes. However, failure of A/C monitoring due to malfunctioning of ADS-B is relevant only after the changes, as introduction of the ADS-B is part of the changes.

In risk assessments of changing systems we assume a target model of the system before the changes and a target model of the system after the changes. In the same way, we assume that a trace model is specified separately for the target model and the risk model before changes on the one hand, and the target model and the risk model after the changes on the other hand. With support for risk modeling with change, the risk model before and the risk model after are combined into one representation, which could of course also be the case for the target models although we do not address the latter issue here.

In the risk graph with change depicted in Figure 20, there are three target element annotations that are relevant both before and after, which is captured by the two-layered appearance of the target element construct. These are *Radar*, *Sector team* and *Task T1*. These identifiers should therefore be specified as tags both in the trace model before and the trace model after. The annotation *ADS-B*, on the other hand, refers only

to the target of analysis after the changes, and the identifier should be a tag in the corresponding trace model.

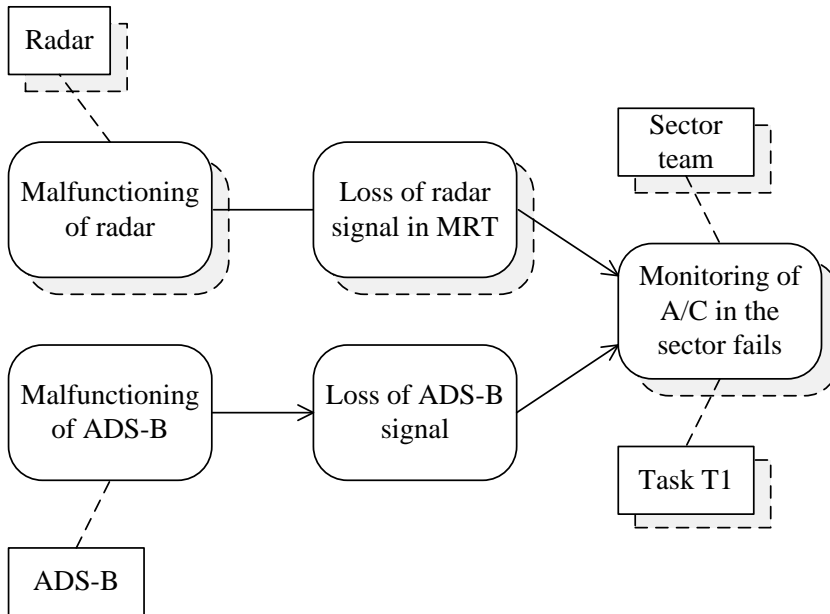


Figure 20 Risk graph with change with tracing to target model

The class diagram of Figure 21 shows the meta-model for risk graphs with change extended with the construct for annotating vertices with references to target model elements. The target element and the trace relation from risk model vertices to target elements have a mode that is one of *before*, *after* and *before-after*. Additionally, the syntactical restrictions listed in Section 5.2 apply.

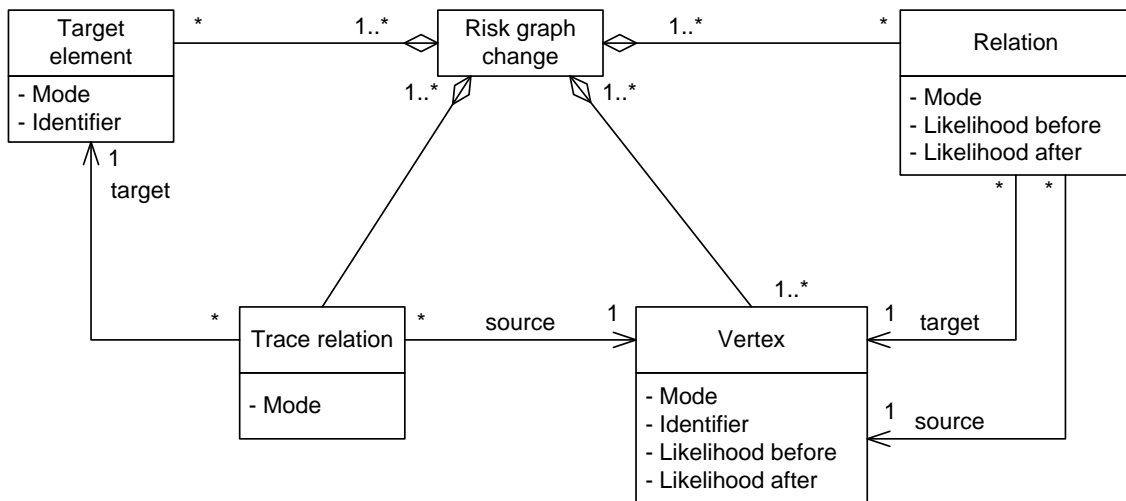


Figure 21 Meta-model for risk graphs with change extended with trace annotation

Assuming a target model before changes and a target model after changes, as well as a trace model before changes and a trace model after changes, the more formal interpretation of the extended risk graphs with change is a quite straightforward generalization of the presentation for traditional risk graphs in Section 6.3.1. We

understand the target element constructs as mere annotations, and therefore maintain the semantics as defined in Section 5.2. By including the target elements and trace relations in the function $before(_)$ and $after(_)$ as defined in Section 5.2 we can extract from a risk graph D with change the separate risk graphs $before(D)$ and $after(D)$ with their respective annotations. If a trace model TM_b has been specified for the target model before changes, each target element in $before(D)$ with tag t annotating a risk graph element with identifier r_{id} is then a reference to the set of elements $\{(t_{id}, r_{id}, t) \mid (t_{id}, r_{id}, t) \in TM_b\}$, and symmetrically for a trace model TM_a for the target model after changes. The notion of dependency carries over similarly, and is applied separately for the parts $before(D)$ and $after(D)$.

7 Risk Assessment Method

In this section we present a risk assessment method for the risk assessment of changing systems under the before-after perspective. The method is based on the general and perspective independent risk assessment process introduced in Section 3, and makes use of the assessment techniques introduced above.

As a running example we use the Air Traffic Management (ATM) risk assessment case study. The results of the case study are fully documented in the appendix of Section 14. In this section we use only extracts for exemplifying selected issues. Notice that whereas the appendix presents the results of instantiating the methods and techniques presented in this deliverable in CORAS, we use in this section the more general techniques that we have introduced in the preceding sections.

7.1 Overview

A typical scenario in the before-after perspective is risk assessors or risk analysts that are asked to predict the effect of implementing certain changes on the current risk picture. The changes that are addressed are planned and/or anticipated, and could be radical with significant impact on the risk picture. Such changes can, for example, involve rolling out a new system, or making major organizational changes such as implementing a merger agreement between two companies. We must therefore understand the current risk picture, the risks that may arise from the change process or change transaction itself, and the future risk picture after the change transaction.

Figure 22 shows the principles by which a risk assessment in the before-after perspective is conducted. Assuming that we have a description of the current target and a description of the change transaction that brings the target from the current state to the future state, we can make a coherent risk picture for the current and future target of analysis, as well as for the change process itself.

From a methodological viewpoint, the main challenges involve obtaining and documenting a risk picture that describes the current and future risk picture and the impact of the change transaction on the risk picture. This requires techniques for modeling the current target and the future target, techniques for modeling the change transaction, and – importantly – techniques for identifying, estimating, evaluating and documenting current and future risks without doing double work.



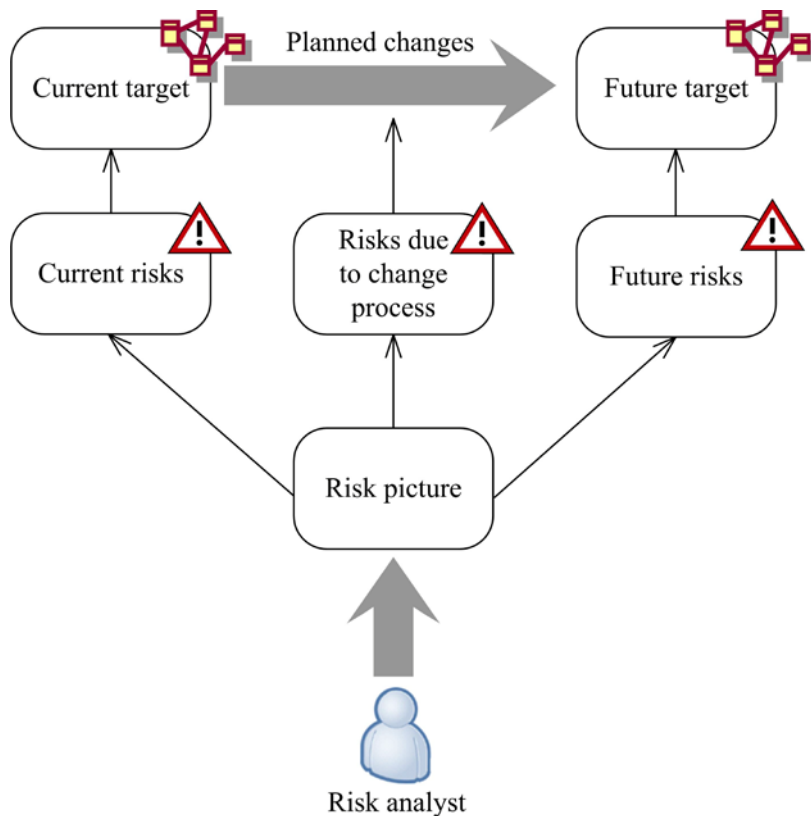


Figure 22 Risk picture in the before-after perspective

Having assessed and documented the current risk picture, the guidelines and techniques for how to derive the future risk picture without conducting the full assessment from scratch is at the core of the method for risk assessment under the before-after perspective. This core work process is illustrated by the UML activity diagram in Figure 23.

We assume here that we have conducted a risk assessment of the current target system and documented the results in a risk model (RM). Based on the documentation of the change transaction, we then proceed along three paths.

The two first paths involve deriving the before-after risk picture. On the one hand we need to determine for each of the current RM elements whether it is affected by the system model (SM) changes. If it is affected, it must be reassessed and replaced by a RM element that is valid for the SM after the changes. If it is unaffected, it is still valid and can be kept in the RM. On the other hand we need to assess from scratch any parts or features that are introduced during the change transaction. When all RM models are checked and made valid, the result is the before-after risk model (RM before-after).

The third path is a separate risk assessment of the change transaction itself, which yields a risk model of its own. Together, the before-after risk model and the risk model for the change transaction are the documentation of the before-after risk assessment, which is the output of the before-after risk assessment.

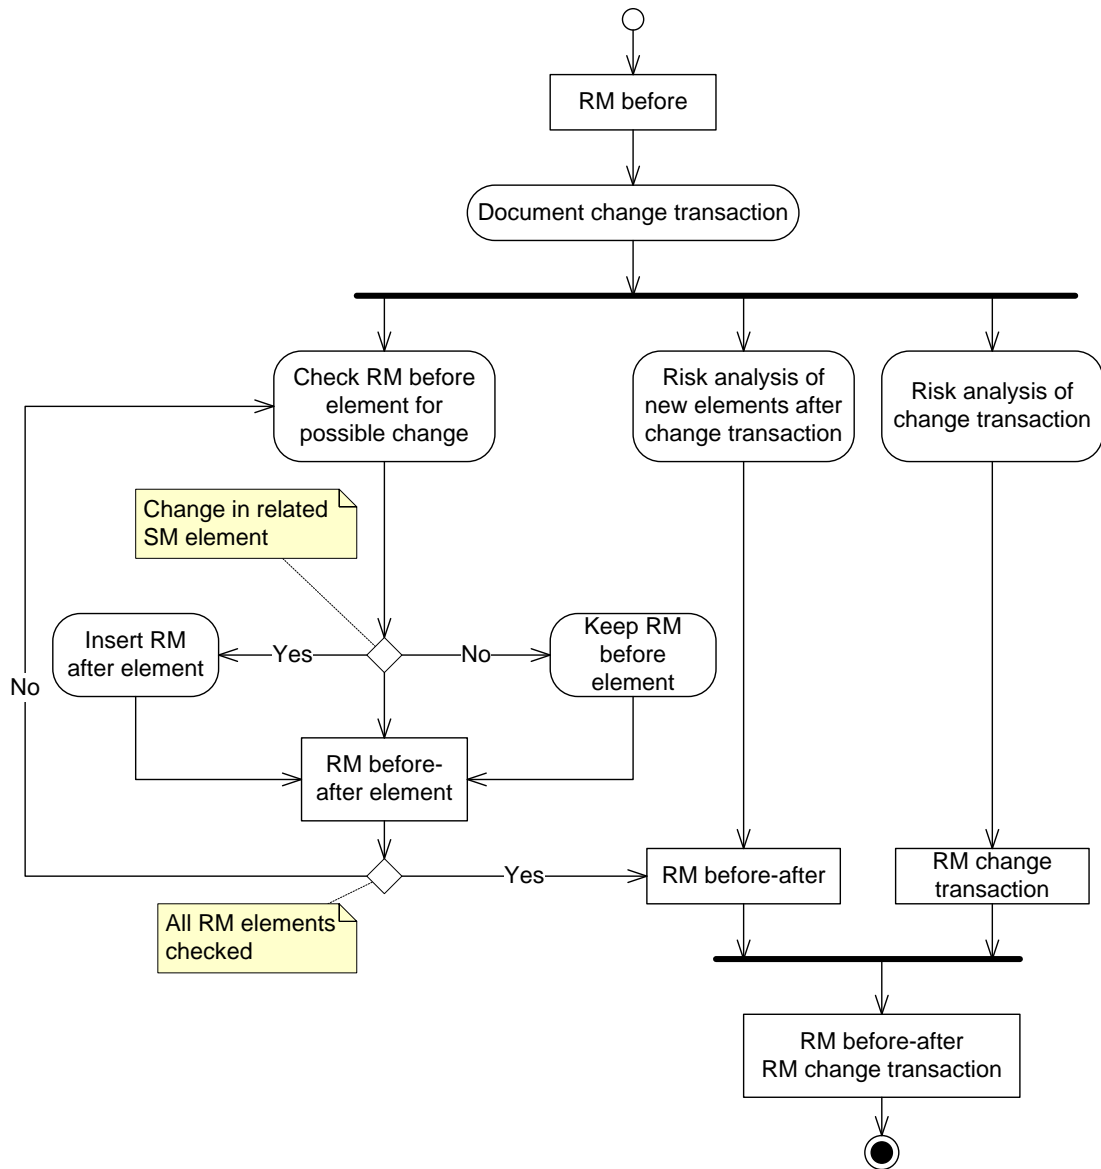


Figure 23 Core of risk assessment process in the before-after perspective

7.2 Conducting the Risk Assessment

In the following we present and exemplify in more details the various activities of the risk assessment process under the before-after perspective. The activities are conducted according to the general overview of the risk management process of changing systems depicted in Figure 3 in Section 3.

7.2.1 Establish the Context

The context establishment includes articulating the overall goals and objectives of the risk assessment, and deciding its focus and scope. This includes making a description of the target of analysis, identifying the assets and deciding the risk evaluation criteria.

In the before-after perspective, we can use traditional risk assessment methods for establishing the context before the changes. Having completed the target description for the target as-is, we proceed by specifying the change requirements and making a description of the change transaction. Based on the change transaction, we make a description of the target to-be. The result is a before-after target description that serves as a basis for the subsequent before-after risk assessment.

The following gives an overview of the context establishment in the before-after perspective:

1. Articulate the goals and objectives of the overall before-after risk assessment.
2. Make a target description and description of change transaction.
 - Description of target before changes.
 - Description of change transaction.
 - Description of target after changes.
3. Conduct asset identification for target before and after changes.
4. Conduct high-level risk analysis of target before and after changes.
5. Decide and specify the risk evaluation criteria before and after changes.

If a separate risk assessment is conducted for the change process, task 3 through 5 must be conducted also for this.

7.2.1.1 Goals and Objectives

Articulating the overall goals and objectives of the assessment includes a specification of the background and motivation for the assessment, providing an initial characterization of the focus and scope of the assessment, and initially describing the planned or anticipated changes to the target system.

ATM example: Goals and objectives. The ATM domain involves an aggregation of services provided by ground-based Air Traffic Controllers (ATCOs). One of the main critical responsibilities of ATCOs is to maintain horizontal and vertical separation among aircrafts and between aircrafts and possible obstacles. They must ensure an orderly and expeditious air traffic flow by issuing instructions and information to aircrafts, and by providing flight context information to pilots, such as routes to waypoints and weather conditions.

An important characteristic of the ATM domain of today is that there are limited interactions with the external world, and therefore also limited security problems in relation to information flow to and from the environment. A further characteristic is that humans are at the center of the decision and work processes, with limited role of automated decision support systems and tools. Current safety problems are therefore mainly due to human errors, air-ground communication problems and degradation of



technical and human services, all possibly combined with adverse atmospheric conditions that could raise safety problems.

The planned and ongoing introduction of new information systems and decision support systems, as well as the reorganization of ATM services, raise new security issues and security concerns with immediate impact on safety issues. The overall objective of this risk assessment is to understand, document and assess security risks of ATM with particular focus on the arrival management process. More specifically, the chosen target of analysis is an Area Control Center (ACC) and the ATCOs. The ACC is a ground-based control center with responsibility of managing the traffic of a given airspace. The actual traffic management is conducted from the operation room (OPS room), which is the operational environment of the ATCOs. The ATCOs have different roles, some of which have their own Controller Working Position (CWP). The CWP makes a range of tools available to the ATCOs for surveillance, communication and planning. The focus of the analysis is the arrival management process with the involved activities, tasks, roles, components and interactions.

Included in the target of analysis are the organizational level changes that are implied by the introduction of the arrival manager (AMAN). The timeframe of introducing the AMAN tool is from today and until 2020. The aim of this analysis is on the one hand to understand, assess and document the current risk picture before the introduction of the AMAN. On the other hand, the aim is to try and foresee risks that may emerge as a consequence of introducing the AMAN and to identify means for risk treatment in order to ensure an acceptable risk level both before and after the implementation of the changes.

The client of the risk assessment is the ATM service provider, which is also the party of the assessment. A risk assessment party is an organization, company, person, group or other body on whose behalf the assessment is conducted. The risks that are identified and evaluated are therefore risks from the perspective of the ATM service provider.

7.2.1.2 Description of Target and Change Transaction

An important objective of the context establishment is to ensure that the risk assessors correctly understand the target of analysis, and that the client, the risk assessors and other involved stakeholders reach a common understanding of the goals, focus and scope of the risk assessment. In order to reach a precise understanding and documentation of the target of analysis, the target should be modeled in a suitable formal or semi-formal modeling language that is well understood, such as the UML. Any misunderstandings or errors must be identified and corrected, because otherwise the results of the risk assessment may be invalid. In the before-after perspective we need to model the target as-is, the change transaction, and the target to be.

The development of modeling artifacts for the specification of changing and evolving systems is outside the scope of WP5 and this deliverable, because WP5 concerns methods and techniques for the assessment and modeling of changing and evolving risks. When conducting the case studies we have used standard UML for making the target description before and after the changes.



ATM example: Target modeling before changes. We document the target of analysis by different kinds of UML diagrams. The target of analysis before the changes is the ATCOs of an ACC, focusing on the arrival management process before the introduction of the AMAN.

We refer to Section 14.1.2.1 for the full documentation of the target of analysis before the changes. To summarize, the target models are divided into three parts, each part documented by means of a specific kind of UML diagram:

- A conceptual overview of the ATM target by means of UML class diagrams.
 - These diagrams give a conceptual overview of the various roles, components and networks involved in the ACC, as well as the most important relationships between them.
- A specification of the internal structure of components by means of UML structured classifiers.
 - These diagrams document in a hierarchical way the various roles and components of the ACC, the communication links between components, and the internal structure of components.
- A specification of the relevant ATM activities by means of UML interactions.
 - These diagrams document the activities involved in arrival management by showing the interactions between roles and components; the diagrams describe the components, events and scenarios at different levels of detail.

The target models describe the Operational Room (OPS Room) as the operational environment of the ACC. The OPS Room is divided into dedicated operative zones or ACC Islands, where each island consists of a number of Controller Working Positions (CWPs). Each CWP is operated by exactly one Air Traffic Controller (ATCO).

The ATCOs have one of four different roles, namely Supervisor (SUP), Planning Controller (PLC), Tactical Controller (TCC) and Coordinator (COO). The PCL and the TCC forms a Sector team and are together responsible for operating and managing the traffic of their sector. The TCC is in charge of all air-ground communication. He monitors the aircrafts in the sector and provides pilots with instructions/clearances on aspects such as speed, altitude and routing to maintain a safe separation with other aircraft flying in the sector, and with other possible obstacles that are present. He also gives pilots weather and air traffic information. When the aircraft approaches the sector boundary, he passes it off to the TCC of the adjacent sector (not always belonging to the same ACC). The PLC assists the TCC, coordinating entry and exit flight level and entry and exit flight point with adjacent sectors in order to ensure a smooth air traffic flow. He also monitors the traffic within the sectors and in most of cases updates the air traffic control system with the instructions given by the TCC.

Groups of neighboring sectors are coordinated by a SUP, who is in charge of managing the sector configurations under his responsibility according overall traffic forecast. The SUP can split and merge sectors depending on the traffic. The SUP is moreover responsible for the formation of the sector teams. The COO is involved only in islands where there is a Terminal Area (TMA). The COO does not work on a CWP, but moves between sector teams to survey the arrival management process and coordinating the tasks between sectors.



The UML structured classifiers of Figure 24 and Figure 25 give two small extracts of the full documentation of the ATM target of analysis. The first diagram shows the OPS Room as consisting of a number of ACC islands that are connected to the ACC network. The OPS Room furthermore has a number of SUPs that communicate with the ACC islands, and that also are connected to the ACC Network via the controller working position CWP_SUP.

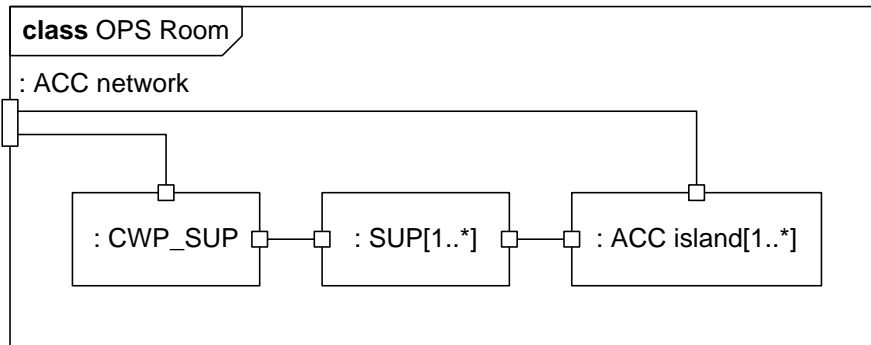


Figure 24 Internal structure of ACC Operational Room before changes

The second diagram shows the internal structure of the ACC islands. There are a number of sector teams that are connected to the ACC network and that also communicate with the SUPs. The COO is also a part of the ACC island, and communicates both with the sector team and with the SUP.

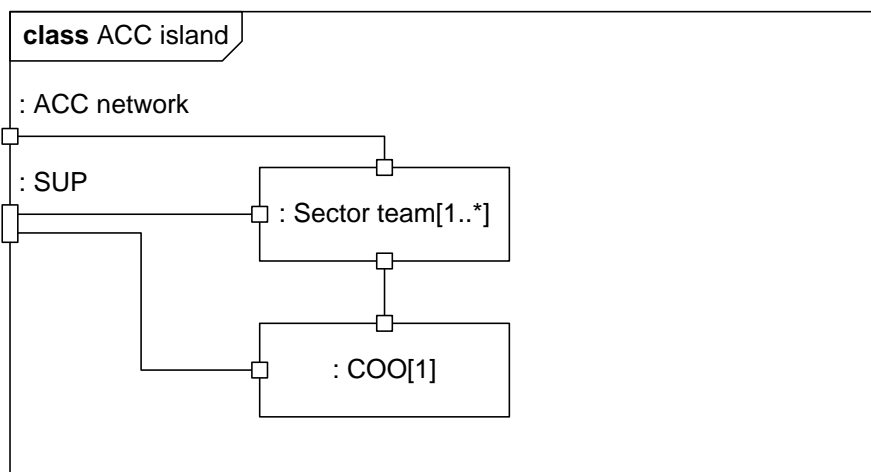


Figure 25 Internal structure of ACC island before changes

In order to properly understand the target of analysis with the focus on the arrival management process, it is important to properly understand the various activities of the arrival management process. We refer to Section 14 for UML interaction overview diagrams and the UML sequence diagrams that document the arrival management tasks. The various arrival management tasks are summarized as follows:

- **Task 1:** Controlling the aircraft (A/C) in the sector
- **Task 2:** A/C data analysis for starting the sequence creation
- **Task 3:** Sequence finalization

- **Task 4:** Clearances to the A/C for building the planned sequence
- **Task 5:** Progressive transfer of the whole sequence to the adjacent sector

Having completed the documentation of the target of analysis before the changes, we turn to the specification of the change transaction. The change transaction is the changes that bring the target system from its current state to its future state after the implementation of the planned or anticipated changes.

Making a precise and correct specification of the change transaction is important for two reasons. First, it is only by precisely knowing the planned or anticipated change process that we can precisely and properly understand the future target, and thereby also predict the risks that may arise in the future. Second, risks may arise due to the risk process itself, and in order to identify and evaluate the risks of the change transaction, we need to include the description of the change transaction in the target description.

The extent to which the change transaction can be described, as well as the level of details of the description, depends on the extent to which the changes are planned or known. Less knowledge and uncertain anticipations yield more underspecification of the change transaction, and consequently more underspecification of the description of the target of analysis after the changes. More underspecification of the target description in turn means that the conclusions from the risk assessment are weaker. The predictions about the future risks therefore depend on the knowledge about the change transaction.

As for the description of the target system, the change transaction should be precisely specified in a formal or semi-formal language that is well-understood.

ATM example: Change requirements. The change requirements that are addressed in the ATM risk assessment are selected from the change requirements of the SecureChange deliverable D1.1.1 [30]. The changes are in the operational processes of managing air traffic in Terminal Areas (TMAs). In particular, the introduction of the Arrival Manager (AMAN) affects the ATM system as a whole both at a process level and at an organizational level.

This ATM risk assessment addresses the organizational level change. The introduction of the AMAN affects the controller working positions (CWPs), as well as the area control center (ACC) as a whole. The main foreseen changes from an operational and organizational point of view are the automation of tasks (i.e. the usage of the AMAN for the computation of the arrival sequence) that currently are carried out by air traffic controllers (ATCOs).

The organizational level changes moreover require the redefinition of the ATCO role of the coordinator (COO), who will be responsible for monitoring and modifying the sequences generated by the AMAN, and for providing information and updates to the sectors. In order to highlight this redefinition of the ATCO role, the COO is renamed to Sequence Manager (SQM).

Having completed the description of the target of analysis before the changes and the description of the change transaction, we proceed to making a description of the target of analysis after the changes. The target of analysis after the changes should be described and documented in the same way as the target of analysis before the changes, as precise as possible in a language that is well understood.

ATM example: Target modeling after changes. As for the target description before changes, we use UML class diagrams to give a conceptual overview, we use UML structured classifiers to document the internal structure of components, and we use UML interactions to document the relevant activities.

The main changes are the introduction of the AMAN, the redefinition of the ATCO role of COO to that of SQM, and the changes to the activities of the arrival management tasks. Additionally, the risk assessment takes into account the introduction of the automatic dependent surveillance-broadcast (ADS-B). The latter is actually independent of the AMAN introduction, but is taken into account in the assessment because ADS-B will be introduced during the same time frame and may affect security issues. ADS-B is a cooperative GPS-based surveillance technique for air traffic control where the aircrafts constantly broadcast their position to the ground and to other aircrafts.

The full description of the target of analysis after the changes is given in the appendix of Section 14. The structured classifiers of Figure 26 and Figure 27 give a small extract. The former depicts the internal structure of the OPS room after the changes, and shows the introduction of the AMAN as a separate component. The AMAN is connected to the CWPs via the ACC network. The latter depicts the internal structure of the ACC island after the changes, and shows the replacement of the COO role by the SQM role. The diagram furthermore shows that after the changes, also this ATCO role operates a CWP. Before the changes it was only the TCC, the PLC and the SUP roles that operated their respective CWPs.

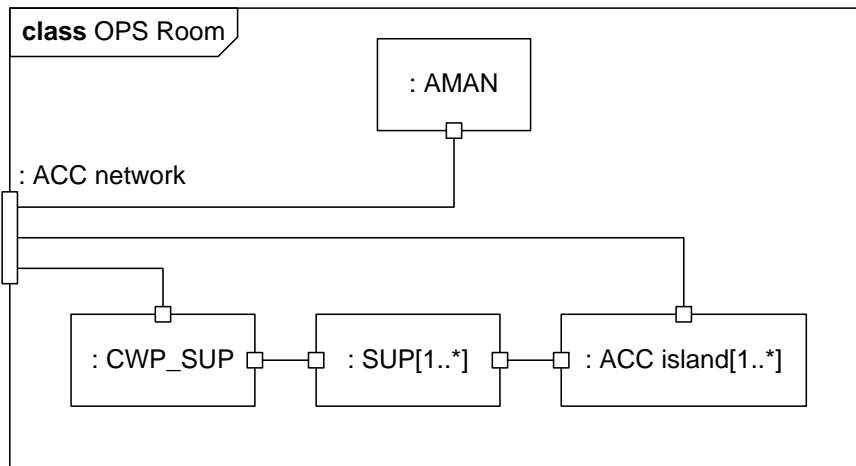


Figure 26 Internal structure of ACC Operational Room after changes

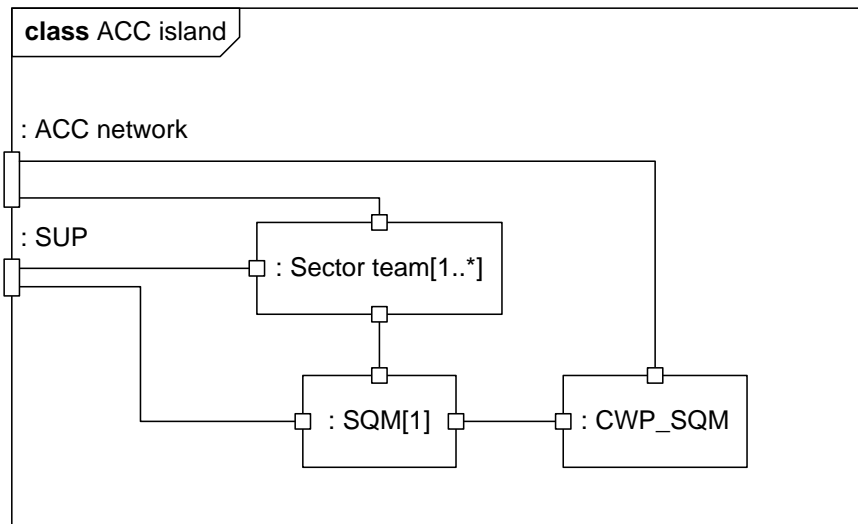


Figure 27 Internal structure of ACC island after changes

The arrival management tasks after the changes are described by means of UML interactions. To summarize, the arrival management with the AMAN consists of the following activities:

- **Task 1:** Monitoring of aircraft (A/C) in the sector
- **Task 2:** Acquisition of the AMAN provided sequence
- **Task 3:** AMAN sequence monitoring and verification
- **Task 4:** Clearances to the A/C for building the planned sequence
- **Task 5:** Progressive transfer of the whole sequence to the adjacent sector

It is mainly Task 2 and Task 3 that are affected by the changes. Task 4 and Task 5 are not affected, and Task 1 is affected only by the ADS-B providing additional support for A/C monitoring.

7.2.1.3 Asset Identification

Making a precise and well-understood description of the target of analysis is an important prerequisite for the actual risk assessment. However, in order to do the risk identification in a focused and directed way, we need to precisely identify the assets. An asset is something to which a party assigns value, and hence for which the party requires protection. The risks that we seek to identify are therefore risks with respect to the identified assets. The asset identification is an important technique for defining the focus of the risk assessment. The purpose is to identify the parts, aspects or properties of the target with respect to which the risk assessment will be conducted.

Since we cannot speak of assets without speaking of parties, we need to explicitly identify and document the parties of the assessment during the asset identification. Furthermore, when addressing changing and evolving systems, we need to determine whether parties, assets, asset values or asset priorities also may change. We therefore conduct separate asset identification both before and after the change transaction. If a

risk assessment is conducted for the change process, assets must be identified also for this part.

ATM example: Asset identification. The party of the ATM risk assessment is the ATM service provider who owns the Area Control Center in question. The assessment addresses security issues, focusing on the following two security properties selected from SecureChange deliverable D1.1.1 [30]:

- **Information protection:** Unauthorized actors (or systems) are not allowed to access confidential queue management information.
- **Information provision:** The provisioning of information regarding queue management sensitive data by specific actors (or systems) must be guaranteed 24 hours a day, 7 days a week, taking into account the kind of data shared, their confidentiality level and the different actors involved.

The risk assessment is conducted with respect to these security properties by operating with the two corresponding assets of confidentiality and availability, respectively. The precise interpretation of these assets throughout the risk analysis is confidentiality of queue management information and availability of queue management information. Because the focus of the analysis is arrival management, the queue management information is restricted to arrival management information.

In the ATM risk assessment, both the party and the assets are the same both before and after the changes. The identified assets are documented in Table 7. In some cases the identified assets have different value, and should have different priority in the risk assessment. In this case the two assets are considered equally important and have a high priority.

Party	Asset	Priority
ATM service provider	Confidentiality	High
ATM service provider	Availability	High

Table 7 Documentation of parties and assets before and after changes

7.2.1.4 High-level Analysis

The purpose of the high-level risk analysis is to complement the target models and the asset diagrams in increasing our understanding of the focus and scope of the risk analysis. This is a rough, initial risk identification that aims to identify the main worries and main incidents so that we can better decide what to include and not, and also to get a better grip of the very motivation for the risk analysis in the first place.

We use table formats for documenting the results of the high-level risk analysis. The high-level risks analysis before and after the changes are conducted and documented separately. A high-level analysis for the risks of the change transaction may also be conducted and documented if required or necessary.

The table format is of four columns, where each row specifies circumstances that may give rise to one or more risks. The first column specifies threats, documenting the initial cause of scenarios or incidents. The second column specifies scenarios and/or incidents, and describes what is harmed. The third column specifies what makes it



possible for the scenarios and incidents to occur by documenting vulnerabilities. The fourth column specifies the parts or elements of the target system that are related to the identified threats, scenarios, incidents and vulnerabilities. The documentation of the relations to the target system facilitates the identification of the parts of the risk picture that may be affected by system changes.

ATM example: High-level risk analysis before changes. Table 8 shows an extract of the documentation of the high-level risk analysis before the changes.

Initial cause	Scenario/incident	Vulnerability	Target element
Component failure	Provisioning of information to ATCO fails due to loss of CWP	Insufficient CWP maintenance	CWP
Software error	The consolidation of data from several radar sources fails, leading to duplication of labels		Surveillance
Component failure; radar disturbance	Malfunctioning of radar antenna leads to loss or degradation of radar signal	Insufficient radar maintenance	Radar
Software bugs	False or redundant alerts from safety tool	Insufficient software testing	OPS room

Table 8 High-level analysis before changes

Table 9 shows an extract of the documentation of the high-level risk analysis after the changes.

Initial cause	Scenario/incident	Vulnerability	Target element
System failure	Loss of the AMAN leads to loss of provisioning of information to ATCO		AMAN
Attacker	Attacker broadcasts false ADS-B signals which leads to the provisioning of false arrival management data	Use of ADS-B; dependence on broadcasting	ADS-B
Attacker	Confidentiality breach by attacker eavesdropping on ADS-B	Use of ADS-B; dependence on broadcasting	ADS-B
Software fail	Provisioning of unstable or incorrect sequences by the AMAN leading to ATCO reverting to manual sequencing	Immature software	AMAN
SQM	SQM fails to build stable sequence or make optimal coordination	High workload on SQM after AMAN introduction	SQM; AMAN

Table 9 High-level analysis after changes

The first step of the high-level analysis after the changes is to conduct a walkthrough of the high-level analysis table before changes to identify risks that are persistent under the changes. This task is facilitated by the fourth column that refers to the relevant parts and elements of the target. In this analysis all entries in Table 8 applies also after the changes. This table therefore also serves to document the high-level analysis after



the changes. Importantly, these risks may change in severity, i.e. their risk levels may increase or decrease. In the full risk assessment, these risks are therefore evaluated both before and after the changes.

The results of the high-level analysis show that there are two main kinds of worries. On the one hand the high-level analysis focuses on component, system and communication failures that can lead to loss of availability. On the other hand, the analysis focuses on the human factor. The human factors are particularly interesting for this assessment, since the change requirements concern the introduction of decision support systems that should mitigate related risks. A part of the analysis therefore aims at investigating to what extent such risks change with the introduction of the AMAN.

7.2.1.5 Risk Evaluation Criteria

The risk evaluation criteria define the level of risk that the parties of the risk assessment are willing to accept. Basically, the criteria are a mapping from risk levels to the decision of either accepting the risk or evaluating the risk further for possible treatment.

In order to speak of risk levels, we need first to define the risk function. The risk function is a mapping from pairs of consequence and likelihood to risk levels. Recall that a risk is the likelihood of an unwanted incident and its consequence for a specific asset, and that the risk level is the level or value of a risk as derived from its likelihood and consequence. Before we can define the risk function we hence need to define the consequence scales and the likelihood scales. Since the kinds of consequences may be different for different assets, we define one consequence scale for each kind of asset. A separate risk function and separate risk evaluation criteria must in turn be specified for each asset and consequence scale.

When addressing a changing target of analysis, it may be that the risk evaluation criteria also change. This can be because parties change, assets change, asset values change, or because the parties become more or less risk averse due to the changes. We therefore need to establish the risk evaluation criteria for the target of analysis both before and after the changes. If a separate risk assessment is conducted for the change transaction, risk evaluation criteria must be established also for that part.

ATM example: Risk evaluation criteria. In the ATM risk assessment, the same risk evaluation criteria apply both before and after the changes. The consequence scales, the likelihood scale, the risk function and the risk evaluation criteria are therefore defined and documented in combination for before and after changes. The consequence and likelihood scales are partly based on requirements and advisory material provided by EUROCONTROL [12][13].

The consequence scales for the confidentiality and availability assets are documented in Table 10 and Table 11, respectively. The scales use qualitative values that range from insignificant to catastrophic, and the meaning of each value is given by a description that serves as a reference point for the degree of severity.



Consequence	Description
Catastrophic	Loss of data that can be utilized in terror
Major	Data loss of legal implications
Moderate	Distortion of air company competition
Minor	Loss of aircraft information data (apart from A/C position data)
Insignificant	Loss of publically available data

Table 10 Consequence scale for confidentiality before and after changes

Consequence	Description
Catastrophic	Catastrophic accident
Major	Abrupt maneuver required
Moderate	Recovery from large reduction in separation
Minor	Increasing workload of ATCOs or pilots
Insignificant	No hazardous effect on operations

Table 11 Consequence scale for availability before and after changes

The likelihood scale of five quantitative values is documented in Table 12.

Likelihood	Description
Certain	A very high number of similar occurrences already on record; has occurred a very high number of times at the same location/time
Likely	A significant number of similar occurrences already on record; has occurred a significant number of times at the same location
Possible	Several similar occurrences on record; has occurred more than once at the same location
Unlikely	Only very few similar incidents on record when considering a large traffic volume or no records on a small traffic volume
Rare	Has never occurred yet throughout the total lifetime of the system

Table 12 Likelihood scale before and after changes

In the ATM analysis the risk functions turned out to be equal for the two assets, and are therefore documented by one risk matrix. The risk matrix shows for each combination of a likelihood and consequence the resulting risk level. The risk function is documented in Table 13 and use three risk levels, namely low (green), medium (yellow) and high (red).

		Consequence				
		Insignificant	Minor	Moderate	Major	Catastrophic
Likelihood	Rare					
	Unlikely					
	Possible					
	Likely					
	Certain					

Table 13 Risk function before and after change

The risk evaluation criteria for both availability and confidentiality of queue management information are as follows in the ATM risk assessment:

- **High risk:** Unacceptable and must be treated.
- **Medium risk:** Must be evaluated for possible treatment.
- **Low risk:** Must be monitored.

Establishing the risk evaluation criteria completes the context establishment.

7.2.2 Identify Risks

The risk identification is the first of the three activities in the actual risk assessment, and is followed by the risk estimation and the risk evaluation. Whereas the preceding context establishment to a large extent can be conducted according to traditional risk management guidelines, the risk assessment of changing systems must be supported by customized techniques and artifacts. In particular, it is during the risk assessment that we make use of the artifacts that we have introduced and presented in Section 5 and Section 6:

- Techniques for the identification and modeling of changing risks with support for reasoning about likelihoods and dependencies before and after the changes.
- Techniques for the indexing and categorization of the various parts of the target system.
- Techniques for the identification of the relationships between the target system and the risk picture, and for the specification of the corresponding trace model.
- Techniques for the explicit visualization of the trace model in the models of the changing risks.

The following gives an overview of the methodological guidelines for conducting the risk identification in the before-after perspective:

1. Identify and document risks based on the target description before changes.
2. Establish and document the trace model for the target model and the risk models before the changes.



3. Based on the trace model and the description of the target of analysis after the changes, identify the parts of the risk picture that are persistent under change.
4. Conduct the risk identification of the changed target only with respect to the parts of the target and the risks that are affected by the changes.

Once the final step is concluded, a trace model for the target model and the risk models after the changes should be established and documented. This will not only serve to better explain where and how the risks arise, but also provides a better basis for conducting new risk assessments in the future in case further changes will be planned or anticipated.

In case a risk assessment of the change transaction itself is required, a separate risk identification must be conducted based on the description of the change transaction from the context establishment. This part of the assessment can be conducted by means of traditional risk identification guidelines and techniques.

The risk identification before the changes is conducted and documented according to traditional guidelines and techniques. Once this task is concluded, however, we need to establish and document the trace model as preparation for the identification of the changing and new risks.

ATM example: Risk identification before changes. The identification of risks involves the identification of threats, unwanted incidents, threat scenarios and vulnerabilities, as well as the relationships between them. The risk identification in the ATM risk assessment was conducted by structured brainstorming involving personnel with expert background from the ATM domain.

The identified risks are fully documented in Section 14 (appendix). In this section we only give some samples for explaining, exemplifying and illustrating the methodology and other artifacts. Notice that the risk assessment as documented in the appendix is by instantiating the methodology and techniques in the CORAS approach, whereas in this part we use the more general approach of risk graphs.

The documentation of the risks before the changes is done by using traditional risk graphs, since changes are still not taken into account. The risk graph in Figure 28 shows some possible causes for the incident of failure of information provisioning to occur. The vertices are structured such that the unwanted incident is to the right and the scenarios that may lead to it are sequentially ordered from left to right.

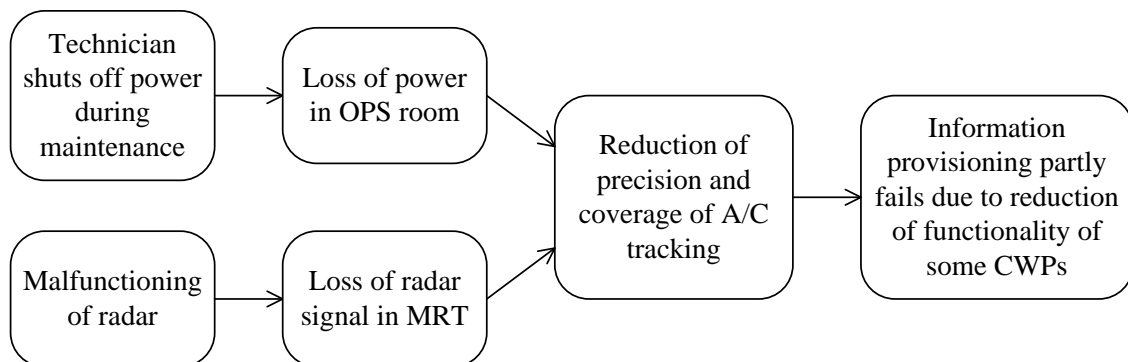


Figure 28 Risk identification before changes – Loss of functionality

Once the risk models for the target of analysis before the changes have been specified, the trace model must be established and documented. This can either be conducted successively for each risk model while they are made or as a separate task for all the risk models after the completion of the identification of the risks before the changes.

Before the trace model can be established, we need a way of referring to the various parts of the target model. For this purpose we conduct the indexing of the target model. The target model indexing can be conducted already during the context establishment once the risk models are completed. However, it is only after the completion of the risk identification that we can identify the relevant parts of the target models. In order to save the effort of indexing the complete target model, the task can with advantage be postponed to after the risk identification. For purposes of efficiency and user-friendliness, the indexing should be automated by tool support, although it can be conducted manually. The format of the target model indexing is defined in Section 6.

ATM example: Indexing the target model before changes. The ATM target of analysis includes a number of actors, events and scenarios. For the purpose of exemplifying the indexing we only show a fraction of the full index in Table 14.

ID	Name	Category	Description
t1	ATCO	Actor	Air traffic controller
t2	TCC	Actor	ATCO role of tactical controller
t3	PLC	Actor	ATCO role of planner controller
t4	COO	Actor	ATCO role of coordinator
t5	SUP	Actor	ATCO role of supervisor
t6	CWP	Actor	ATCO controller working position
t7	CWP_TCC	Actor	CWP of TCC
t8	CWP_PLC	Actor	CWP of PLC
t9	T1	Scenario	Task T1 of controlling the A/C in the sector
t10	T4	Scenario	Task T4 of clearances of sequence to A/C
t11	Radar	Actor	Radar antennas for surveillance
...

Table 14 Fraction of ATM target model index before changes

Recall from Section 4 that we assume a target system model as a specification of actors, events and scenarios. The categorization of the various target model elements is accordingly. This means that the notion of actor includes all components and entities that interact with other components and entities.

Given the target model index and the finalized risk models before the changes, we can establish and document the trace model. Making the trace model is a manual task and amounts to identifying and documenting the relations between the target system and the risk models. The format of the trace model is defined in Section 6. It documents

relations between the target model and the risk models, and consists of pairs of target model elements and risk model elements. The format furthermore allows the use of tags for grouping of sets of relations. The names for the tags are selected by the risk assessment participants for their own convenience.

ATM example: Trace model before changes. Table 15 shows a fraction of the trace model before the changes. The trace model uses the target indexes for referring to the target model elements and the risk element identifiers (name of risk graph vertices) for referring to the risk model elements. The risk element identifiers can be used provided they are unique. Otherwise, unique indexes must be specified also for the target model elements.

The fraction of the trace model shows some of the relations between the target model and two of the vertices of the risk graph in Figure 28, namely *Malfunctioning of radar* and *Information provisioning partly fails due to reduction of functionality of some CWP*s. The former is related to the radar, whereas the latter is related to the TCC, the PLC and their respective CWPs, as well as to the arrival management task T1.

Notice that four of the relations can be referred to collectively by means of the tag Sector team. The names of the tags can furthermore be chosen for the convenience of the user so that the relations can be referred to by intuitively understandable names.

Target index	Risk element identifier	Tag
...
t11	Malfunctioning of radar	Radar
t2	Information provisioning partly fails due to deduction...	Sector team
t3	Information provisioning partly fails due to deduction...	Sector team
t7	Information provisioning partly fails due to deduction...	Sector team
t8	Information provisioning partly fails due to deduction...	Sector team
t9	Information provisioning partly fails due to deduction...	Task T1
...

Table 15 Fraction of trace model before changes

In order to facilitate the identification of the parts of the risk picture that may be affected by changes in the target system, we make use not only of the trace model, but also of the support we have provided for explicitly visualizing the trace model in the risk models. By using the extended risk graph syntax defined in Section 6, we specify the trace model by annotating the risk graphs with the target element construct. The name or identifier of a target element construct is a tag from the trace model.

ATM example: Visualization of trace model in risk graph before changes. Figure 29 shows the risk graph from Figure 28 annotated with relations between risk graph elements and target model elements. The relations correspond to the relations specified in the trace model, a fraction of which is shown in Table 15.



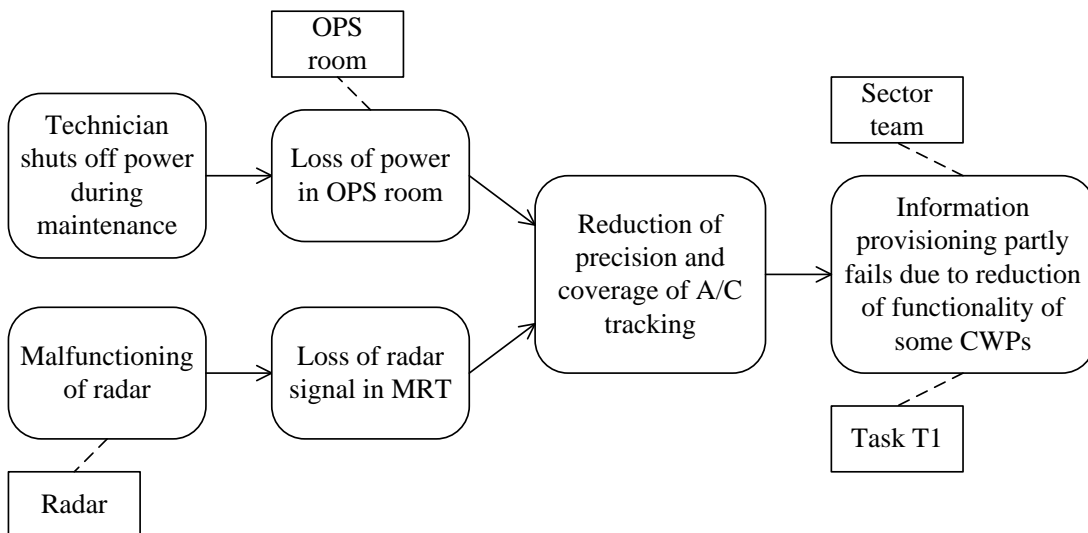


Figure 29 Risk identification with traceability before changes - Loss of functionality

The risk graph depicted in Figure 30 shows a further example the identified risks, and shows the two incidents of degradation of A/C position data and delays in sequence provisioning. Relevant relations to the target model are furthermore specified by the annotations.

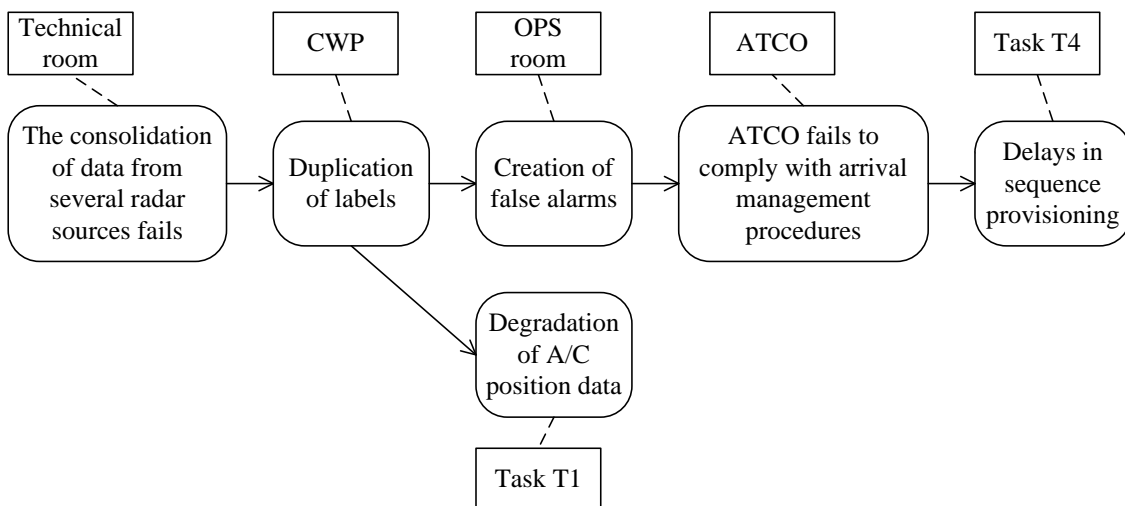


Figure 30 Risk identification with traceability before changes - Label duplication

The indexing of the target model and the specification of the trace model can be conducted manually. This can, however, be a tedious and time consuming task. For efficient use of the techniques and artifacts that we introduce for the purpose of supporting and facilitating the risk assessment of changing systems, tool support should be provided.

First, the user should have the possibility of automatic indexing of the target model, leaving open only the column for the description of each target model element that optionally can be filled in by the user.

Second, the user should have tool support for making the trace model. On the one hand it should be possible to fill in the two first columns of the trace model table by selecting from a list of target model indexes and risk model identifiers. When annotating the risk models, the tags of the already specified entries from the trace model table should then be provided as options that are automatically inserted if selected. On the other hand it should be possible to create new tags directly in the risk models and select target model indexes from a list. The annotation of a vertex v with a tag t and the selected target indexes ti_1, \dots, ti_n should then result in the new entries $(ti_1, v, t) \dots (ti_n, v, t)$ in the trace model table.

With the finalization of the identification and documentation of the risks and the relations between the target model and the risk models before the changes, we proceed to the identification of risks after the changes. When addressing the target system after the changes, we aim at addressing only the parts of the risk picture that are affected by the changes, so as to not conducting a full risk identification from scratch. We furthermore aim not only at documenting the risk picture after the changes, but also at explicitly documenting how the risks change as a consequence of the system changes.

The identification of the risks that are persistent under change is conducted by a walkthrough of risk models and the trace model. The risks that may be affected by system changes need to be reconsidered for the identification incidents, scenarios, etc. that may emerge, change or disappear. Furthermore, any new parts or features of the target of analysis must be considered from scratch for the identification of new risks that may emerge. During this activity, any dependencies in the target models or risk models must be carefully considered in order to take into account the possible propagation of changes.

Consider, for example, the risk graph depicted in Figure 29. If the radar is not affected by the changes, the threat scenarios that are related only to the radar are also not affected. In that case, the vertices *Malfunctioning of radar* and *Loss of radar signal in MRT* are persistent under change. If the OPS room is affected by the system changes, the scenario *Loss of power in OPS room* must be reconsidered, and also the subsequent scenarios due to the risk graph dependencies as defined in Section 6.

The identified risks are documented by using risk graphs with change as introduced and defined in Section 5. These diagrams support the simultaneous documentation of risks before and after changes, showing risks that emerge, risk that disappear and risks that remain. For the risk elements that represent risks after the changes, we furthermore need to make a separate trace model for documenting the relationships to the target system after the changes.

ATM example: Risk identification after changes. The risk graphs in Figure 31 and Figure 32 shows some of the results of the risk identification after changes.

The risk graph with change depicted in Figure 32 shows a sample of the result of the risk identification after changes. The diagram builds on the risk graph of Figure 30. Whereas the latter addresses issues related to radar data, the former also takes the ADS-B into account. In the risk graph with change, the risk elements that are relevant both before and after the changes are represented by the two-layered before-after vertices. The risk elements that emerge after the changes are represented by the



white, solid rounded rectangles, whereas the elements that disappear are represented by the grey, dashed rounded rectangles.

Also the annotated target model elements for documenting the trace model are of one of the kinds before, after or before-after. A target element of kind before-after, such as OPS room, means that the element is part of the target system both before and after the changes. The target element ADS-B, on the other hand, is an example of an element of the target system only after the changes.

The names or identifiers of the target element annotations after the changes are tags from the trace model that must be established and documented for the target model after changes. As for the target model before changes, we need to do the indexing of the target model after the changes before we make the trace model.

ATM example: Indexing and trace model for the target model after changes. The indexing of the target model after the changes is exemplified in Table 16. The elements are provided unique indexes (IDs), also for the elements that are the same before and after the changes. In this example, we see that the SQM has replaced the previous COO, and that AMAN and ADS-B are new elements.

ID	Name	Category	Description
u1	ATCO	Actor	Air traffic controller
u2	TCC	Actor	ATCO role of tactical controller
u3	PLC	Actor	ATCO role of planner controller
u4	SQM	Actor	ATCO role of sequence manager
u5	SUP	Actor	ATCO role of supervisor
u6	CWP	Actor	ATCO controller working position
u7	CWP_TCC	Actor	CWP of TCC
u8	CWP_PLC	Actor	CWP of PLC
u9	T1	Scenario	Task T1 of controlling the A/C in the sector
u10	T4	Scenario	Task T4 of clearances of sequence to A/C
u11	AMAN	Actor	Arrival Manager
u12	Radar	Actor	Radar antennas for surveillance
u13	ADS-B	Actor	Automatic dependent surveillance broadcast
...

Table 16 Fraction of ATM target model index after changes

The relations between the target model and risk models after the changes are exemplified by the fraction shown in Table 17. For convenience, the names of the tags from the trace model before changes should be reused for the elements that present both before and after. These relations to the target model can then be represented by the before-after target elements in the risk graphs with changes, such as the annotation named Radar in Figure 31.



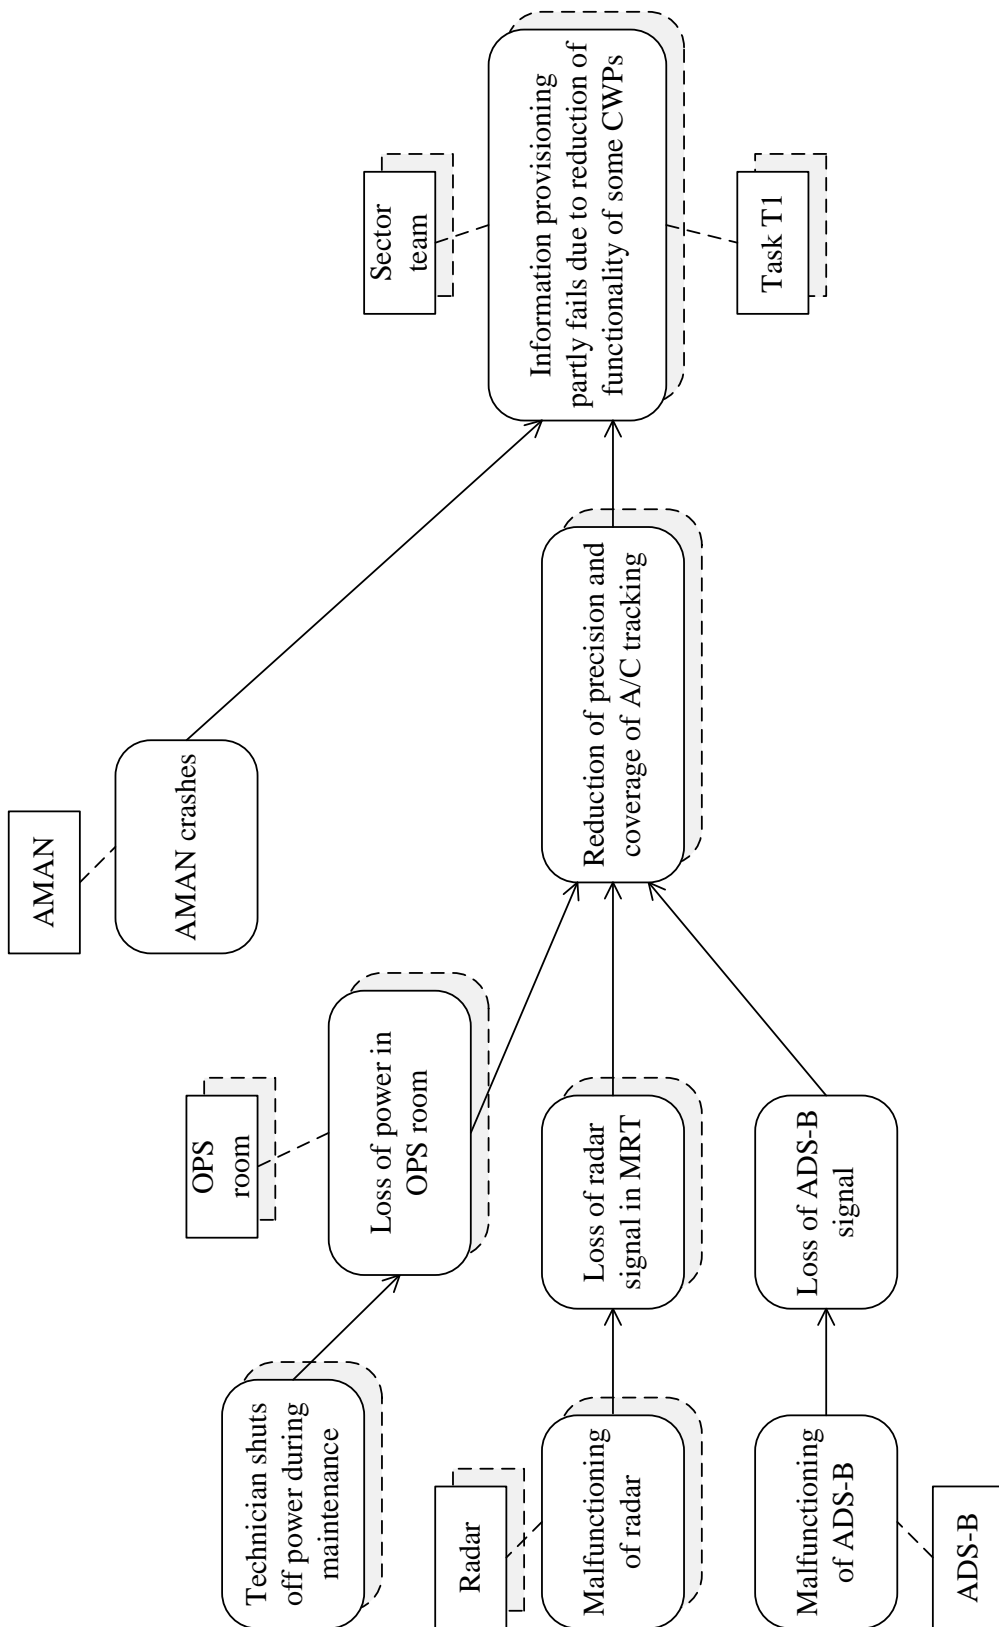


Figure 31 Risk identification with traceability after changes - Loss of functionality

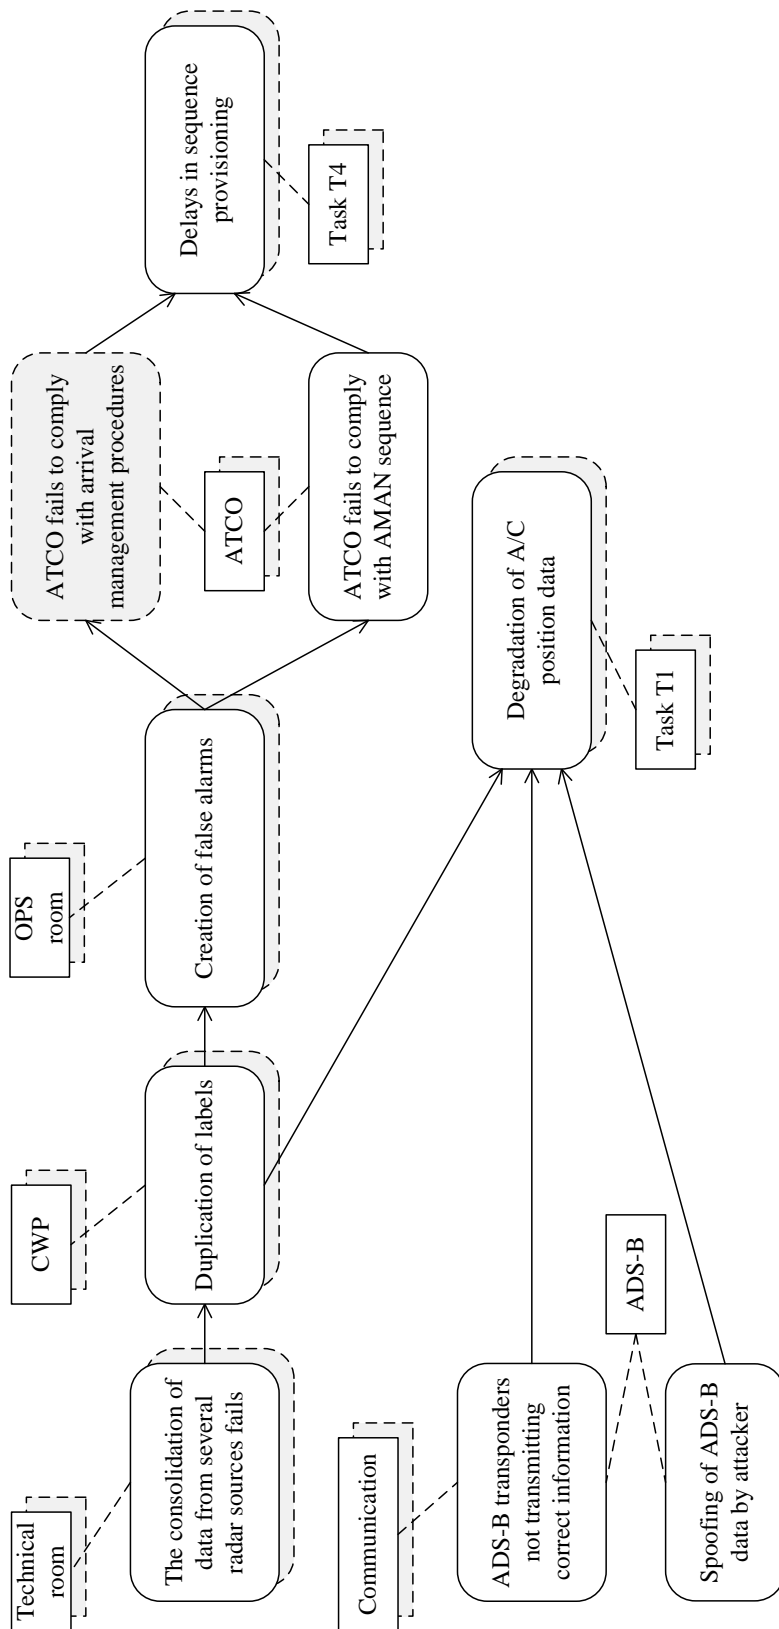


Figure 32 Risk identification with traceability after changes – Label duplication and ADS-B

The tag named ADS-B is used only in the trace model after the changes, and we see that the annotation of that name in Figure 32 is correspondingly of kind after.

Target index	Risk element identifier	Tag
...
u12	Malfunctioning of radar	Radar
u13	Malfunctioning of ADS-B	ADS-B
u2	Information provisioning partly fails due to reduction...	Sector team
u3	Information provisioning partly fails due to reduction...	Sector team
u7	Information provisioning partly fails due to reduction...	Sector team
u8	Information provisioning partly fails due to reduction...	Sector team
u9	Information provisioning partly fails due to reduction...	Task T1
...

Table 17 Fraction of trace model after changes

Having completed the risk identification we have identified and documented risks before changes, risk after changes, as well as risks that occur both before and after changes. Referring to the latter as risks that are persistent is, however, somewhat imprecise. Although certain scenarios may occur both before and after the given change transaction, it may still be that these scenarios evolve by changes in the likelihood of their occurrence. The identification of such changes to the risk picture is a topic for the next activity, namely risk estimation.

7.2.3 Estimate Risks

Risk estimation basically amounts to estimating the likelihood and consequence of unwanted incidents. Additionally, risk estimation should include the estimation of the likelihood for the occurrence of scenarios that may lead to unwanted incidents. This will increase the understanding of the most important sources of risks, and it will also provide a better basis for estimating the likelihood of the unwanted incidents. Likelihood estimation may furthermore include the estimation of the conditional likelihood of scenarios to lead to other scenarios.

With the documentation of the changing risks in risk graphs, the risk estimation can be conducted more or less as in traditional risk assessments. The estimates can be based on historical data, statistics, expert judgments and so forth. The techniques for likelihood calculation and consistency checking of likelihood estimates introduced in Section 5.3 can also be utilized.

ATM example: Likelihood estimation. The likelihood estimation of the identified risks in the ATM risk assessment was conducted as a structured brainstorming involving personnel with expert background from the ATM domain. The estimates were made by a walkthrough of the risk graphs with change. The estimates for the vertices before the



changes were made first. Next, the experts judged whether the likelihoods would be different after the change transaction, and if so, new estimates were made. Finally, the likelihoods for the vertices after the changes were estimated. All the results were documented on-the-fly by annotating the risk graphs.

Some of the results of the likelihood estimation are shown in the risk graphs depicted in Figure 33 through Figure 36. The likelihood values are from the likelihood scale defined in Table 12.

Vertices of kind *before* are assigned one likelihood, as exemplified by the likelihood *rare* of the vertex *ATCO fails to comply with arrival management procedures* in Figure 34. Vertices of kind *after* are also assigned one likelihood, as exemplified by the likelihood *rare* of the vertex *ATCO fails to comply with AMAN sequence*. Vertices of kind *before-after* are assigned a pair of likelihoods, as exemplified by the likelihoods *possible* and *unlikely* of the vertex *Delays in sequence provisioning*. The former is the likelihood of the scenario before the changes, and the latter is the likelihood of the scenario after the changes.

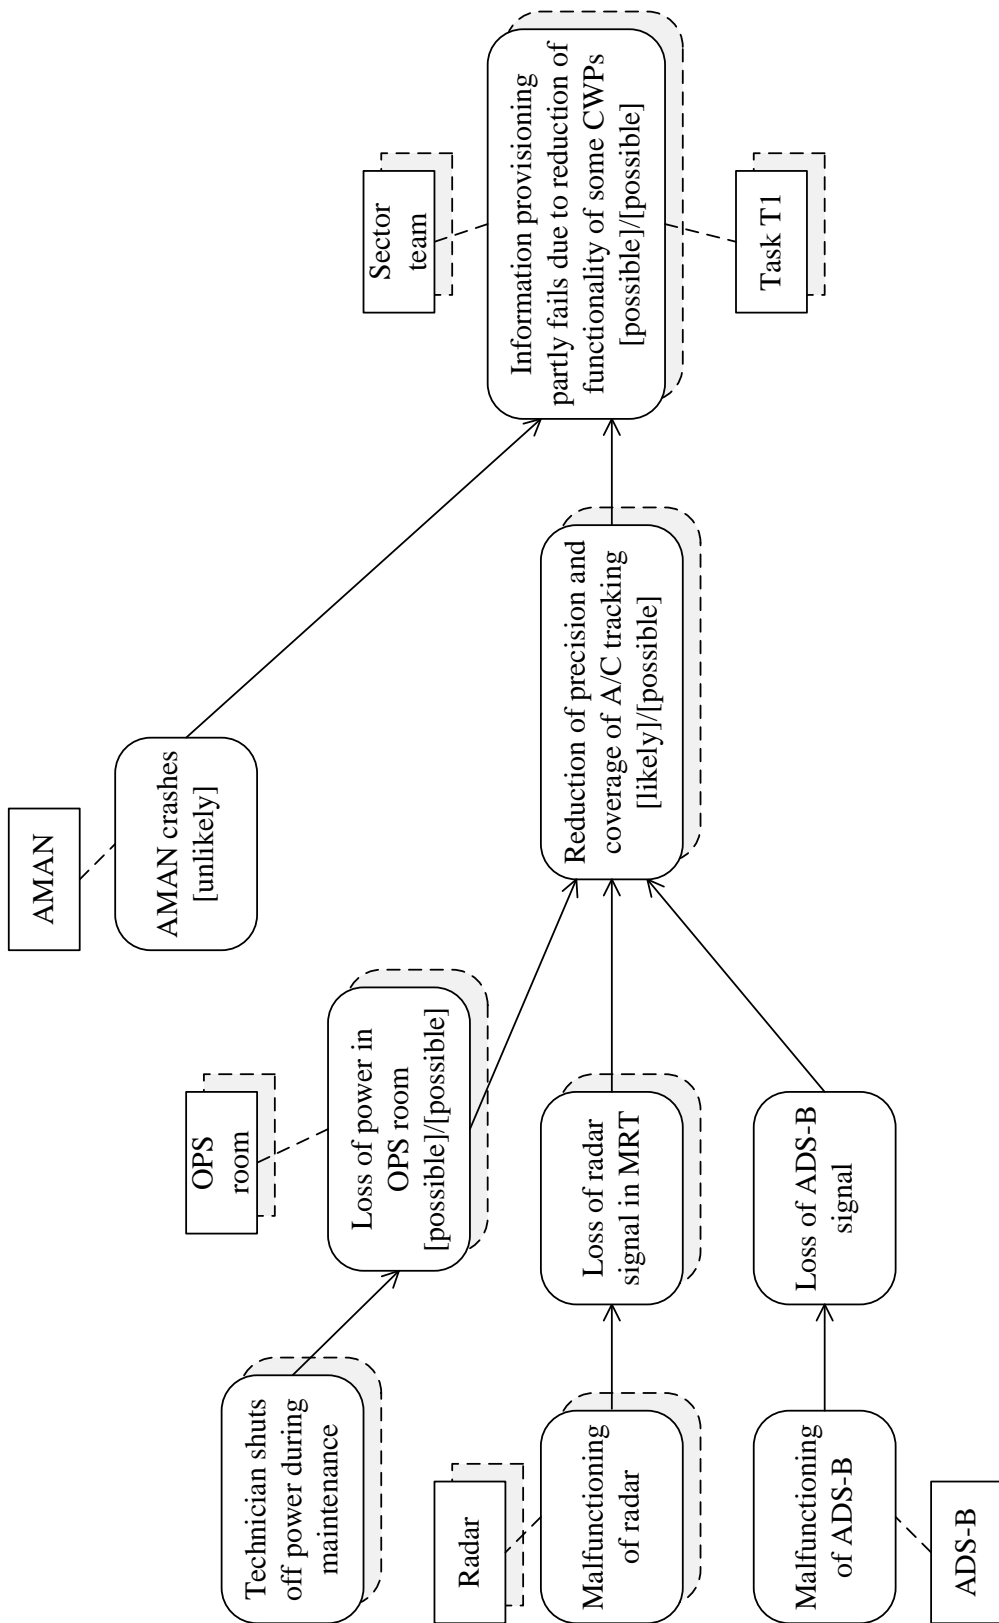


Figure 33 Likelihood estimation before and after changes – Loss of functionality

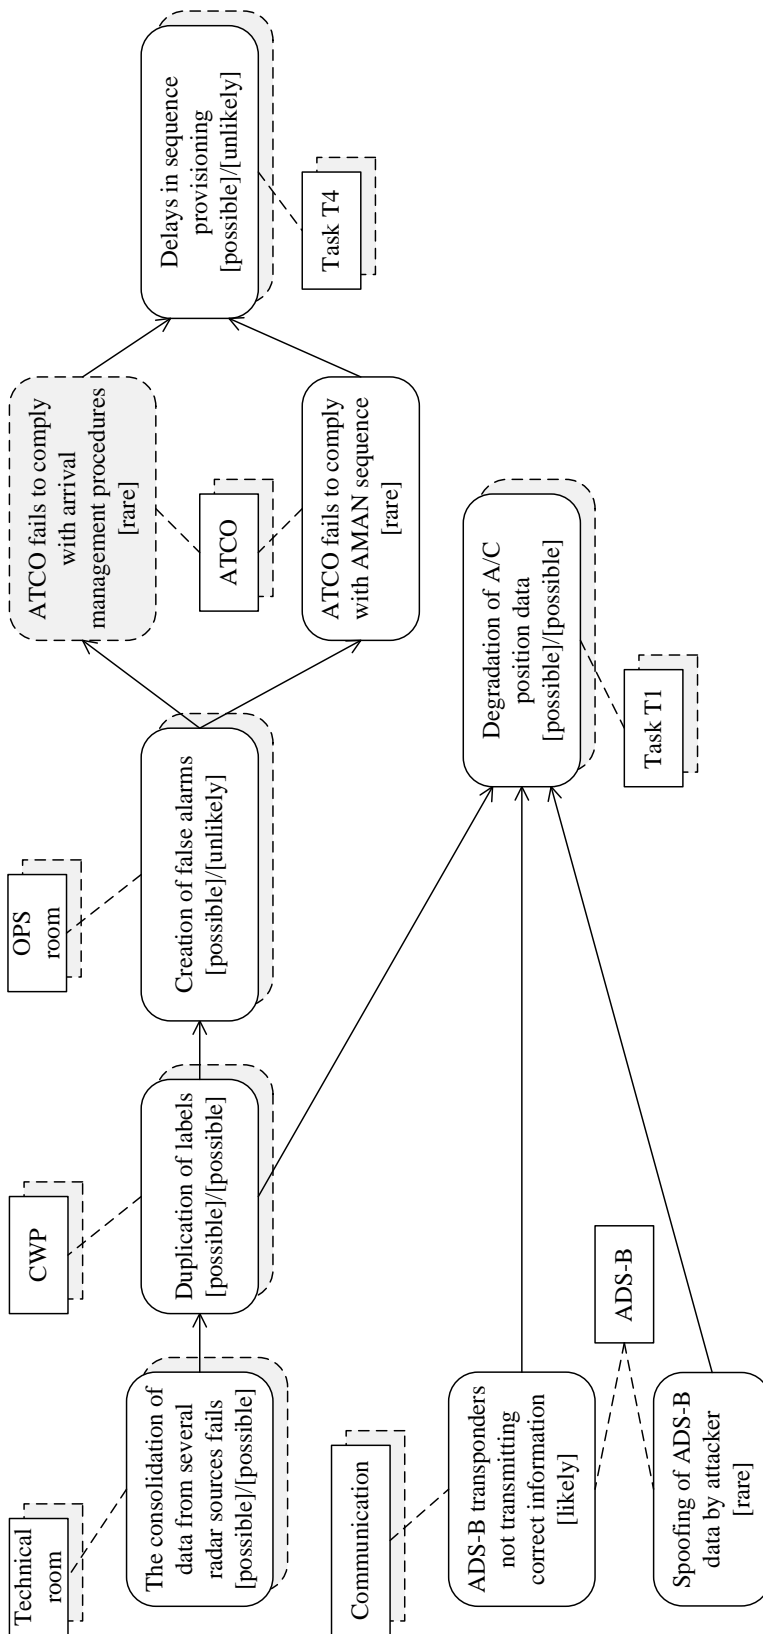


Figure 34 Likelihood estimation before and after changes – Label duplication and ADS-B

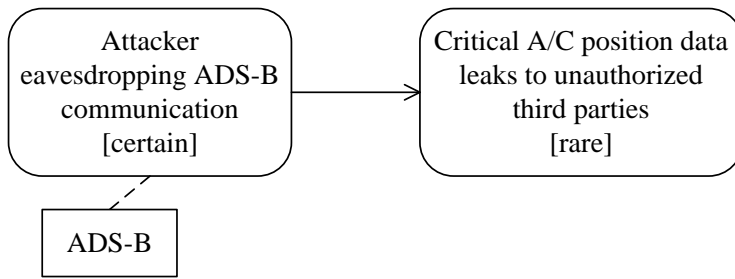


Figure 35 Likelihood estimation before and after changes - Leakage of ADS-B data

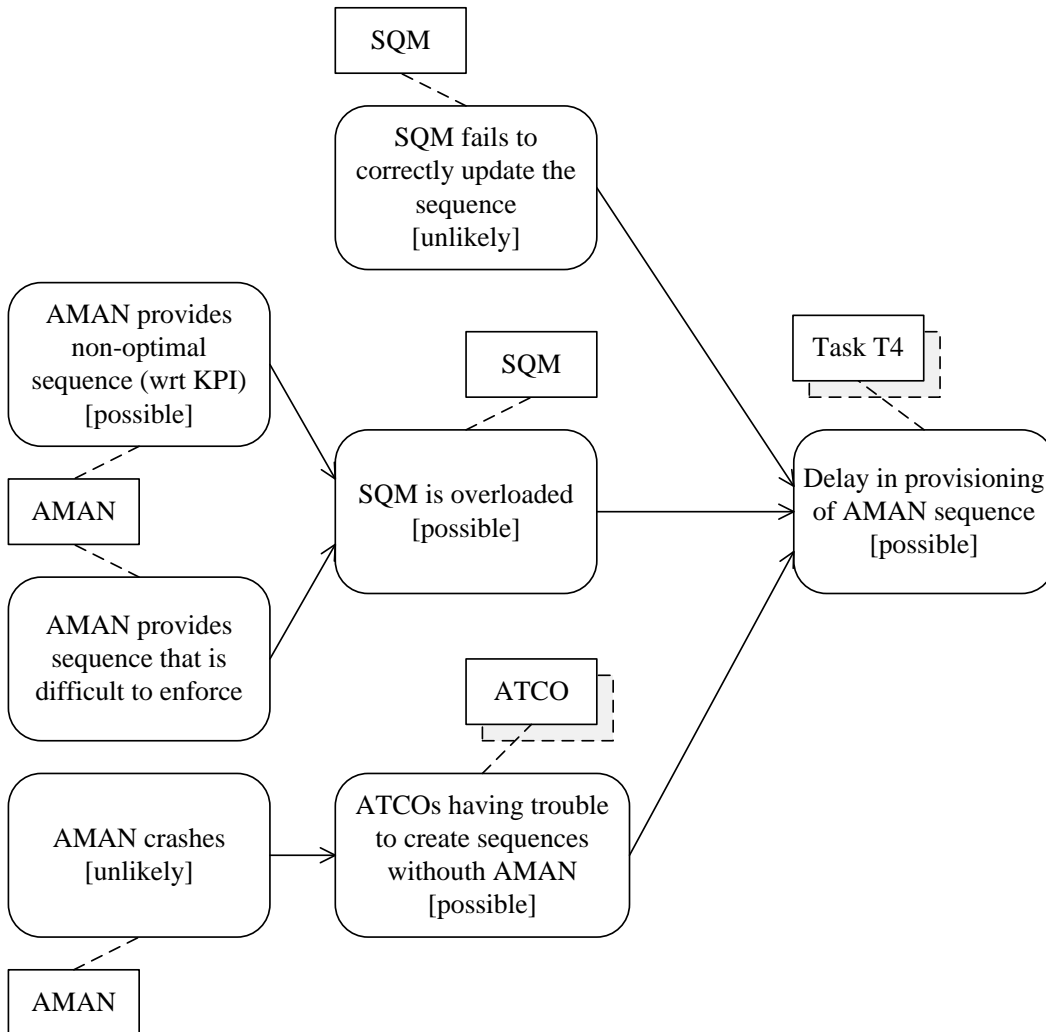


Figure 36 Likelihood estimation before and after changes - Human factors

The consequence estimation is also conducted by a walkthrough of the risk graphs, where the unwanted incidents are assigned consequences. A consequence is the impact of an asset in terms of harm or reduced asset value. For each unwanted incident, it is only the party associated with the asset in question that can determine the consequence. Whereas the likelihoods of scenarios or incidents to occur are independent of the parties and assets, the consequences depend on the parties; if a

risk assessment is conducted from the perspective of several parties, it may be that one and the same unwanted incident has different consequences for the different parties.

Risk graphs do not provide explicit support for documenting the consequence estimates. On the one hand we need for each unwanted incident to identify and document the assets that are harmed, and on the other hand to estimate and document the consequence for each of these assets. In the following we choose to estimate and document the unwanted incidents, the assets they harm and the consequence for each asset separately by using a table format. The reader is referred to Section 14 for examples of the explicit documentation of unwanted incidents, assets and consequences in risk models by using the CORAS approach.

ATM example: Consequence estimation. In the risk graphs exemplified in Figure 33 through Figure 36, the identified unwanted incidents are placed to the right. For each of these, the assets that are harmed and the consequence for each asset before and after changes are documented in a table format. Table 18 shows the consequence estimations for the unwanted incidents exemplified in this section.

Risk ID	Unwanted incident	Asset	Consequence before	Consequence after
R1	Information provisioning partly fails due to reduction of functionality of some CWP's	Availability	Minor	Minor
R1	Delays in sequence provisioning	Availability	Minor	Minor
R3	Degradation of A/C position data	Availability	Minor	Minor
R4	Critical A/C position data leaks to unauthorized third parties	Confidentiality	N/A	Major
R5	Delay in provisioning of AMAN sequence	Availability	N/A	Major

Table 18 Consequence estimation before and after changes

Because each pair of an unwanted incident and an asset represents a risk, each row in the consequence estimation table also represents a risk. Each risk is in the table given a risk ID that can later be used for referring to the risks.

7.2.4 Evaluate Risks

During the risk evaluation we first calculate the risk levels by using the risk function defined during the context establishment and the likelihood and consequence estimates of the previous risk assessment step. We then compare the risk levels with the risk evaluation criteria. The risk evaluation is conducted separately for the risks before the changes and the risk after the changes. Beyond that, the risk evaluation of changing systems is as for traditional risk assessments.



ATM example: Calculation of risk levels. The calculation of the risk levels of the risks documented in this section is shown in Table 19. The three risks that occur both before and after the changes remain at level low, where as the two risks that emerge after the changes are of level medium and high, respectively.

Risk ID	Risk level before			Risk level after		
	Likelihood	Consequence	Risk level	Likelihood	Consequence	Risk level
R1	Possible	Minor	Low	Possible	Minor	Low
R2	Possible	Minor	Low	Unlikely	Minor	Low
R3	Possible	Minor	Low	Possible	Minor	Low
R4	N/A	N/A	N/A	Rare	Major	Medium
R5	N/A	N/A	N/A	Possible	Major	High

Table 19 Risk levels before and after changes

Having calculated the risk levels, we do the risk evaluation by simply plotting the identified and estimated risks into the risk matrix defined during the context establishment. The risk evaluation is documented separately for the risks before changes and the risks after changes.

ATM example: Risk evaluation. The three risks that occur before the changes are plotted into the risk matrix of Table 20, and the five risks that occur after the changes are plotted into the risk matrix of Table 21. From the matrices we see that risk R2 is slightly lower after the changes, but remains of level low.

		Consequence				
		Insignificant	Minor	Moderate	Major	Catastrophic
Likelihood	Rare					
	Unlikely					
	Possible		R1 R2 R3			
	Likely					
	Certain					

Table 20 Risk evaluation before changes

		Consequence				
		Insignificant	Minor	Moderate	Major	Catastrophic
Likelihood	Rare				R4	
	Unlikely		R2			
	Possible		R1 R3		R5	
	Likely					
	Certain					

Table 21 Risk evaluation after changes



According to the risk evaluation criteria, it is only the two risks that emerge after the changes that must be treated. Risk R5 is of level high and therefore unacceptable, whereas risk R4 is of level medium and must be evaluated for possible treatment. The three remaining risks are low and must be monitored.

7.2.5 Treat Risks

For changes that are planned and/or anticipated, the risk assessment and the risk treatment should ensure that the risk level is maintained at an acceptable level through the risk transaction, and that the risk level of the resulting system is acceptable. Whether or not the current risks should be subject to treatment depends on the time frame of the change transaction, as well as the priorities of the parties and other stakeholder. If the planned or anticipated changes are very immediate, it may not make sense to invest in treatments for risks that disappear after the changes.

Treatments for the unacceptable risks are identified by conducting a walkthrough of the risk models documenting the unacceptable risks. The treatment identification is conducted separately for the risks before and after the changes, and for the risks of the change transaction itself in case a risk assessment has been conducted for the risk process. In some cases, treatments that are identified for the risks before the changes are applicable also after the changes, and should therefore be considered for both of these risk pictures.

ATM example: Treatment identification. The ATM domain of today is characterized by limited interactions with the external world. There are therefore also limited security problems in relation to information flow to and from the environment. The ATM risk assessment to a large extent confirmed this, and the treatment identification therefore foremost targeted the ATM system after the changes.

Table 22 shows a few samples of the treatment options that were identified. The first column documents the various treatments. The second column documents the threat scenarios (vertices in risk the graphs) to which the treatments can be applied. The fourth column documents the risks that are mitigated by implementing the treatments.

Notice that the effect of each treatment follows the dependencies of risk graphs. For example, by improving the testing of the AMAN software, the likelihood of the scenario *AMAN crashes* depicted in the risk graph of Figure 36 may decrease. Because this scenario may lead to the scenario *ATCOs having trouble to create sequences without AMAN*, also the likelihood of the latter may decrease. The risk level of R5 represented by the incident *Delay in provisioning of AMAN sequence* may then finally decrease because of the dependency on the preceding scenarios.

Treatment	Scenario	Risk
Implement backup or improve maintenance of the ADS-B transponder	ADS-B transponders not transmitting correct information	R3
Implement encryption of ADS-B signals	Spoofing of ADS-B data by attacker	R3
	Attacker eavesdropping ADS-B communication	R4
Ensure necessary and sufficient criteria for AMAN sequence calculation	AMAN provides non-optimal sequence (wrt KPI)	R5
	AMAN provides sequence that is difficult to enforce	R5
Ensure thorough training of ATCOs on using the AMAN	SQM fails to correctly update the sequence	R5
Improve software testing	AMAN crashes	R5

Table 22 Risk treatment after changes

8 Integration of Risk Assessment and Testing

In this section the integration of technical solutions from WP5 on risk assessment and WP7 on testing is outlined. The particular solutions that will be used are the CORAS risk assessment methodology of WP5 and the TellingTestStories TTS approach of WP7. The reader is referred to Section 13 of the appendix for the presentation of the instantiation in CORAS of the assessment method and techniques introduced in the preceding sections.

Telling TestStories (TTS) [14] is a tool-based methodology developed at the University of Innsbruck (UIB) for model-driven system testing of service-centric systems. TTS has separated but interrelated requirements, system and test models which are validated by consistency, completeness and coverage checks. The framework guarantees full traceability between the requirements, system and test models, and the executable services of the system which is crucial for efficient test evaluation. The test model integrates the tabular definition of tests and is very abstract to be defined by domain experts and automatically transformed to executable test cases in Java. TTS supports the definition and execution of tests for functional and non-functional requirements especially security requirements [15]. The evolution of the requirements, the system and the tests are handled by state machines attached to model elements from which regression test suites are generated.

The conceptual integration of both methodologies is outlined below using a running example of the HOMES case study. In the first section an overview of the mapping between the two solutions is outlined. Then the links are explained and exemplified one by one using the HOMES example. For a short introduction to the relevant aspect of the HOMES case study the reader is referred to Section 15 of the appendix. In this section we only focus on the integration examples drawn from that case study.

8.1 Mapping between CORAS and TTS

There are several points which allow the integration of the CORAS risk assessment methodology and the TTS testing approach. Some of these are a direct conceptual integration at the level of artifacts, while others represent an indirect conceptual integration also on the level of artifacts, but requiring interpretation steps by the stakeholders.

8.1.1 Risk Concepts

Figure 37 depicts the meta-model containing the concepts of the CORAS risk language.

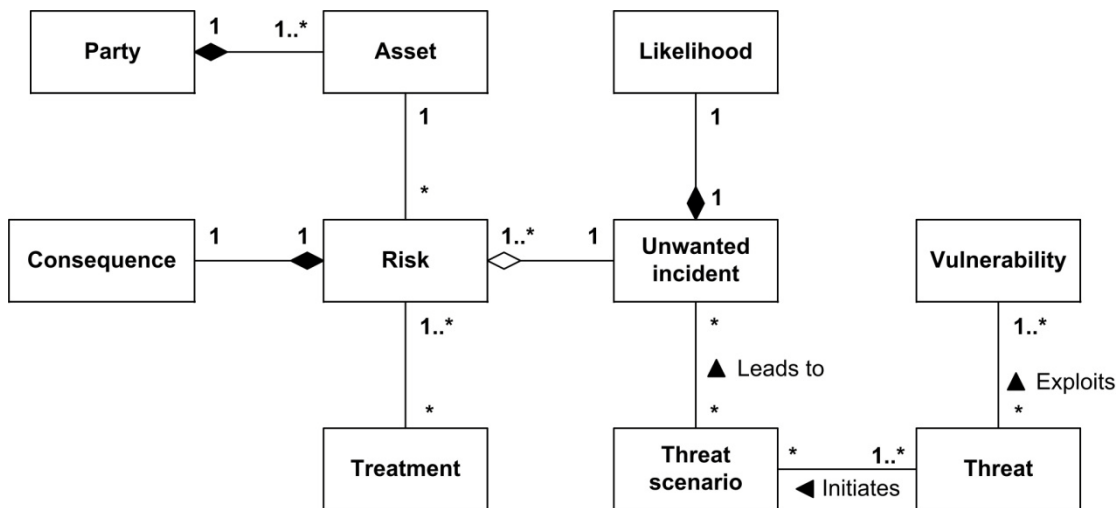


Figure 37 Basic risk concepts from the CORAS language

The following list shortly explains the concepts of the CORAS risk model:

- **Asset:** Something to which a party assigns value and hence for which the party requires protection.
- **Consequence:** The impact of an unwanted incident on an asset in terms of harm or reduced asset value.
- **Likelihood:** The frequency or probability of something to occur.
- **Party:** Stakeholder; an organization, company, person, group or other body on whose behalf a risk analysis is conducted.
- **Risk:** The likelihood of an unwanted incident and its consequence for a specific asset.
- **Threat:** A potential cause of an unwanted incident.
- **Threat scenario:** A chain or series of events that is initiated by a threat and that may lead to an unwanted incident.
- **Treatment:** An appropriate measure to reduce risk level.
- **Unwanted incident:** An event that harms or reduces the value of an asset.
- **Vulnerability:** A weakness, flaw or deficiency that opens for, or may be exploited by, a threat to cause harm to or reduce the value of an asset.

8.1.2 Testing Concepts

Figure 38 depicts the testing concepts defined in the Telling TestStories (TTS) meta-model. For a more detailed explanation the reader is referred to the deliverable D7.3, Section 7.

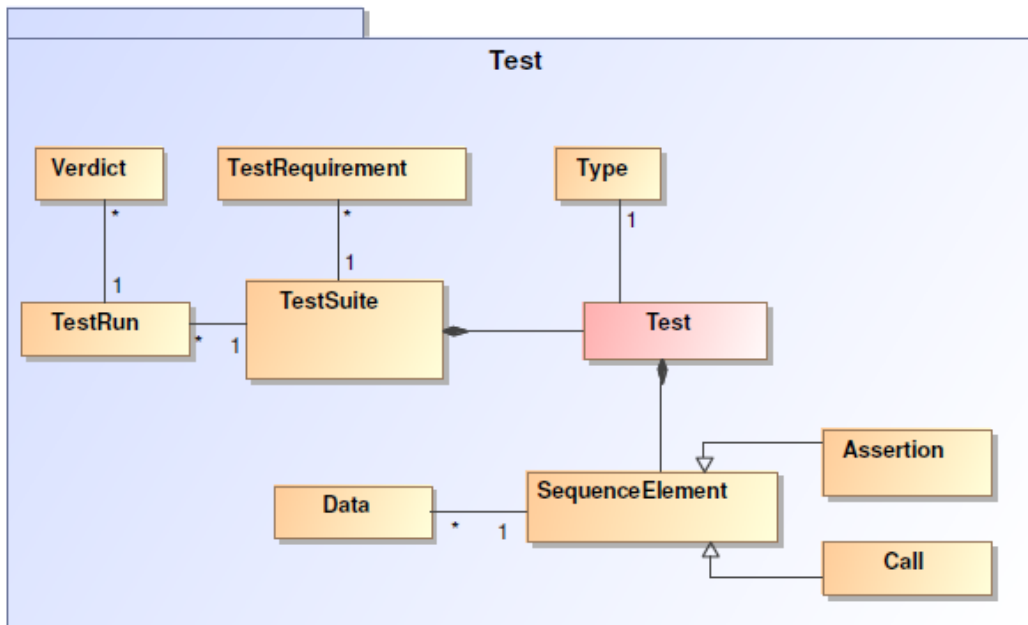


Figure 38 Basic testing concepts of the Telling TestStories test meta model

The following list shortly explains the concepts of the Telling TestStories meta model:

- **TestSuite** is a sequence of **Test** elements.
- **TestRequirement** is assigned to a TestSuite and defines test selection or test exit criteria.
- **TestRun** is the executed TestSuite.
- **Verdict** is the result of an Assertion and can have the values pass, inconclusive, fail, or error.
- **Test** has a **Type**, which can be evolution, stagnation, regression, or obsolete, and consists of SequenceElement artifacts.
- **SequenceElement** is an abstract element which can either be an Assertion or Call and has assigned Data.
- **Assertion** defines an evaluation criterion to compute a verdict.
- **Call** invokes a service operation.
- **Data** defines test data for a specific service Call or an Assertion.

8.1.3 Integration

In this section the integration points are grouped according to the direction of the information flow. Table 23 presents possible links in the direction from CORAS models to TTS models. In the other direction, i.e. from TTS models to CORAS models, the links presented in Table 24 are possible. The various integration links depicted in Table 23 and Table 24 are applied to the HOMES case study in the following Sections 8.2.1 to 8.2.4 and Sections 8.3.1 to 8.3.2, respectively.

	Artefact from the risk model	Related artefacts in the test model	Section
a)	CORAS treatments	Security functionality tests	8.2.1
b)	CORAS risk diagram to change	Regression tests	8.2.2
c)	CORAS threat scenarios	Misuse case test	8.2.3
d)	CORAS risk values	Prioritisation of tests	8.2.4

Table 23 Inputs from the risk model to the test model

- a) This conceptual integration directly links the CORAS concept of treatment which is any measure of reducing the identified risk level with the TTS concept of test which tests the security functionality of the specific treatment.
- b) A CORAS diagram depicting the risk to change in the before-after perspective can be used to derive regression tests. This is an indirect integration requiring interpretation on behalf of the test engineers.
- c) This conceptual integration directly links the CORAS concept of threat scenario with the TTS concept of test. Misuse case tests allow checking whether a specific vulnerability is still exploitable.
- d) CORAS risk values are modeled as a combination of likelihood and consequences. These risk values can be used to prioritize defined tests. This is an indirect integration requiring prioritization on behalf of the test engineers.

	Artefact from the test model	Related artefacts in the risk model	Section
e)	Result of security functionality test	Confirmation of successful deployment of treatment in the risk model Reduction of related risk values	8.3.1
f)	Result of misuse case tests	Elimination of threat scenarios and vulnerabilities which are not anymore exploitable	8.3.2

Table 24 Feedback from the test model to the risk model

- e) This direct conceptual integration links the TTS concept of verdict with the CORAS concept of treatment. A successful test of the security functionality of a treatment can be used to confirm deployment of a treatment in the CORAS risk diagram and in addition confirm the intended reduction of risk values.
- f) This direct conceptual integration links the TTS concept of verdict with the CORAS concept of threat scenarios. If a misuse case test or penetration test which check whether a specific vulnerability or threat scenario are still

exploitable fails, then these threat scenarios or vulnerabilities can be deleted from the CORAS risk diagram or set to an inactive state.

8.2 Information Flow from Risk Model to Test Model

The various integration links described in the previous section are discussed below outlined as an example on the basis of the HOMES case study.

8.2.1 Security Functionality Tests

A point of integration between CORAS risk models and TTS test models is the direct derivation of tests from a treatment. In a CORAS risk model, treatments may describe technical solutions which are then tested accordingly.

A treatment is any measure that helps to reduce identified risks. Several of these measures can be tested directly using security functionality tests which test the proper operation of these measures.

An example for such a treatment could be the implementation of a specific protocol to protect from eavesdropping. The test engineers would then define security functionality tests that check whether the proposed treatment works as expected.

Another example could be the deployment of a network filter to only accept connections from certain hosts. A security functionality test could be used to check whether the filter is correctly denying connection attempts from untrusted hosts.

HOMES case study example

In the example application of the risk assessment methodology on the HOMES case study described in Section 15, a specific treatment has been proposed, namely the deployment of a non-repudiation service. See the CORAS treatment diagram of Figure 39.

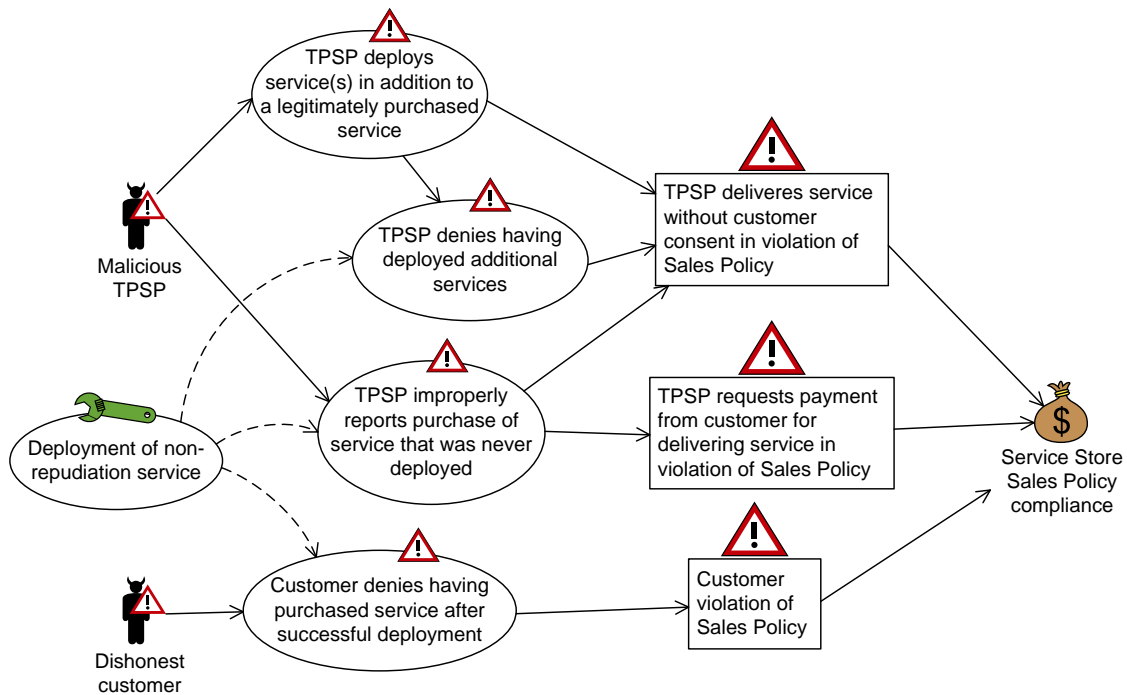


Figure 39 Treatment identified during the HOMES risk analysis

This specific treatment can be used to derive security functionality tests. The test engineers build new tests on the basis of the updated risk models received by the risk analysis team. To guarantee the proper functionality of the treatment, which defines a new security service, test engineers design and execute a test suite on the adapted system. In particular this test suite contains *functional security tests*, which test the functionality of the treatment proposed by the risk engineers (cf. Figure 39). Table 25 lists the functional security test related to the treatment. This functional security test checks the deployed treatment to its compliance with a non-repudiation protocol.

Test ID	Test	Test result
3	NEW: Test NonRepudiation	

Table 25 Functional security test related to treatment NonRepudiation service

Figure 40 graphically depicts the test NonRepudiation which is used to test the functionality of the newly deployed Non-Repudiation Service. The test runs through the defined non-repudiation protocol and checks at the end if the respective transaction is contained in the logs. The test result in the exemplified table is empty because the tests are not run yet and the results not available yet (cf. Section 8.3.1 for the test results).

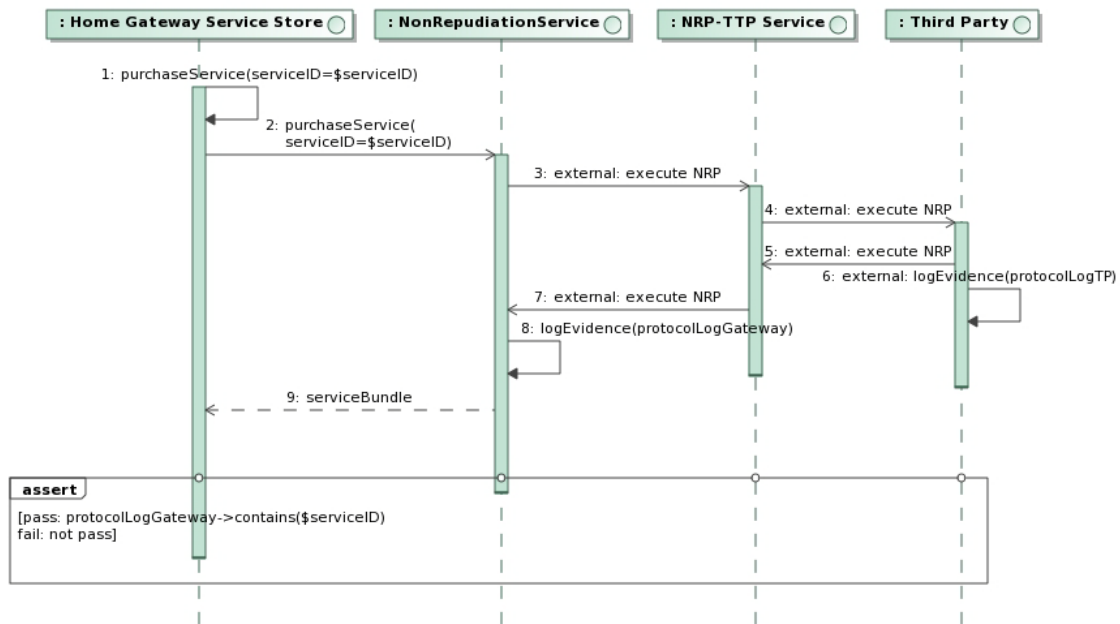


Figure 40 Test NonRepudiation related to the treatment Non-Repudiation Service

8.2.2 Regression Tests

If a risk assessment methodology is applied under the before-after perspective, also the risk to the change itself is analyzed and modeled. The assessment of the risks to change (see Figure 23 of Section 7) provides a basis for the derivation of regression test.

Regression testing is the selective retesting of a system or component to verify that modifications have not caused unintended effects and that the system or component still complies with its specified requirements [33].

The integration link based on the identified risk to change is not a direct conceptual integration, but it requires interpretation and reasoning by the test engineers.

HOMES case study example

Figure 41 depicts risks to the change transaction as identified as part of the risk assessment of introducing the new treatment discussed in the previous section. The treatment of deploying a non-repudiation service could lead to two unwanted incidents related to the integrity of existing security services.

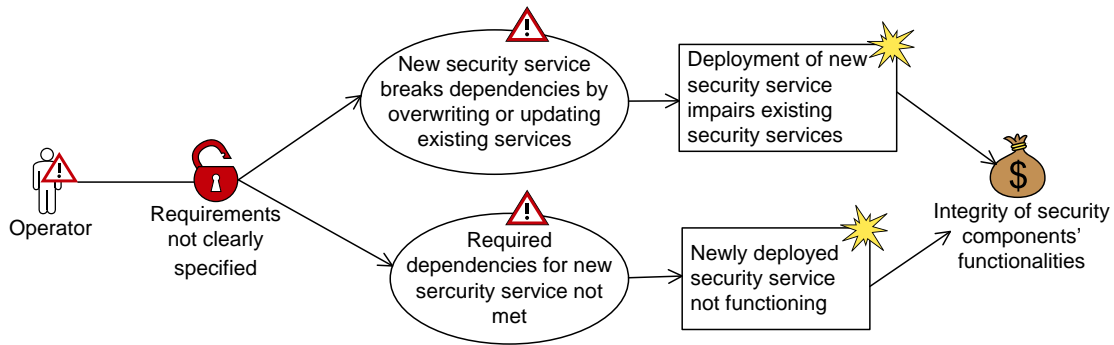


Figure 41 CORAS threat diagram of risk to change

Related regression tests could be defined for testing the functionality of an already existing security service. In our example the existing security service is a Confidentiality service which allows encryption and decryption of information.

The test cases related to the risk of change outlined in Figure 41 are listed below in Table 26. The test result is empty because the tests are not run yet and the results not available yet (cf. Section 8.3.1 for the test results).

Test ID	Test	Test result
1	Confidentiality Service Encryption	
2	Confidentiality Service Decryption	

Table 26 Regression tests of pre-existing security services

8.2.3 Misuse Case Test

Similar to the derivation of security functionality tests described in the previous section, the identified threats and unwanted incidents documented in a CORAS risk model can be used to derive misuse case tests or penetration tests. Penetration tests are a method to evaluate the security of a system by simulating an actual attack on the target of analysis. An example of such a link could be a newly identified threat scenario related to a specific vulnerability.

Misuse case tests are considered to check whether vulnerabilities can still be exploited and threat scenarios can still occur. The test engineers derive specific tests that check whether an identified vulnerability is still exploitable or an identified scenario may still occur in the system after the deployment of a treatment.

HOMES case study example

In Figure 39, several threats related to misbehavior of customers or third party service providers are depicted. A misuse case test could check whether these malicious behaviors, i.e. the threat scenarios, are still possible in the system after the deployment of a non-repudiation service. Table 27 contains an unspecified example of such misuse case tests. The test result is empty because the tests are not run yet and the results not available yet (cf. Section 8.3.1 for the test results).

Test ID	Test	Test result
4	NEW: Misuse Tests	

Table 27 Misuse case tests

A concrete example of such a misuse case test could be to conduct a purchase on the HOMES gateway. If the non-repudiation protocol can be bypassed, the threat scenarios may still occur. Otherwise according to the functional security tests the customer should not be able to deny the purchase as there is now a logged evidence.

8.2.4 Prioritization of Tests

Another integration mechanism between CORAS risk models and TTS test models is the use of risk values. Risk values are determined from the likelihoods and consequences of unwanted incidents, and are affected by vulnerabilities, threats, threat scenarios and treatments, which in turn are related to various parts or elements of the target system. The documentation of risk levels can be utilized in test prioritizing; the mapping from risk levels to prioritization of tests paves the way for a risk-based testing approach in which highly critical vulnerabilities, countermeasures or system components are tested first.

HOMES case study example

The application of the risk assessment methodologies on the HOMES case study focused only on risk identification. Due to the limited time and scope of the second case study in WP5, the risk estimation and evaluation as such were left out.

Therefore the prioritization of tests based on risk values is not shown on concrete examples but just conceptually explained. In the case of a more complex system with many components, risk values indirectly related to these components could be used to prioritize test.

As a concrete example, suppose we have a set of security services which are deployed on the HOMES gateway and have to be tested. If only a limited time span and resources are available for testing, the tests related to the critical security services should be run first. This could be the case for a test of the non-repudiation security service which has a high risk value associated and thus gets precedence over the test of the low risk confidentiality service.

8.3 Information Flow from Test to Risk Model

The integration from the test model to the risk model is based on the provisioning of feedback from the test model to the risk model. As outlined in the previous section, there are several starting points in the risk model from which to derive tests. Vice versa, there are several feedback links possible which are discussed in the following subsections.

8.3.1 Confirmation of Risk Reduction by Treatments

Successful test results of the security functionality of treatments can be used to confirm the associated risk reduction in the risk model. That way, the risk reduction related to a specific treatment is then also backed by concrete tests, which can be connoted in the risk model.

HOMES case study example

The functional security test related to the treatment of deploying a non-repudiation service validates whether the purchase of a service is logged properly such that the purchase cannot be denied later.

The Home Gateway Service Store invokes the purchaseService operation with a specific service identifier on the Non-RepudiationService. After the confirmation, the Non-RepudiationService also logs the non-repudiation and triggers serviceBundle on the Home Gateway Service Store. The test passes if the purchasing has been logged properly and fails otherwise.

If the test passes as outlined in Table 28 (Test 3), this information can subsequently be fed back to the risk model. In the risk model the related treatment is thus confirmed by a functional security test. The risk reduction related to this treatment can thereby also be confirmed.

Test ID	Test	Test result
1	Confidentiality Service Encryption	pass
2	Confidentiality Service Decryption	pass
3	NEW: Test NonRepudiation	pass
4	NEW: Misuse Tests	fail

Table 28 Results of the test runs

Table 28 highlights the results of all the tests used in the examples of the previous sections. In Figure 42 the treatment “Deployment of non-repudiation service” is now actually confirmed by the Test 3. Since the treatment reduces the probabilities of exploiting the threat scenarios to zero, the related threats can be considered as completely treated (cf. Figure 42), i.e. as eliminated.

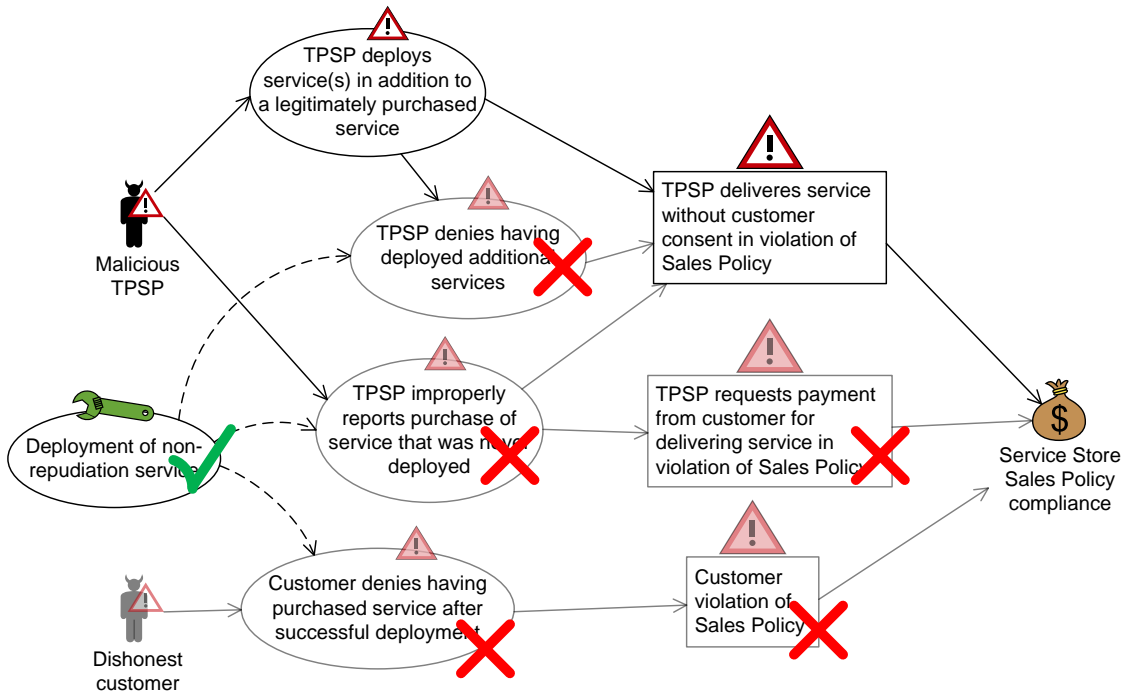


Figure 42 Test results fed back to the risk model

8.3.2 Confirmation of Closure of Vulnerabilities

In this regard also vulnerability and penetration tests are relevant for vulnerabilities that should not be exploitable anymore after the implementation of a treatment. In this case the threat scenarios and vulnerabilities themselves are tested.

HOMES case study example

The incorporation of the test results and the corresponding updates of the risk model result in the new risk model outlined in Figure 43 where the previously identified threats and unwanted incidents are now confirmed to be not exploitable anymore and are thus completely removed. This removal is explicitly seen by comparing the treatment diagram of Figure 42 with the threat diagram of Figure 43 that shows risks after the change, i.e. after the implementation of the identified treatment.

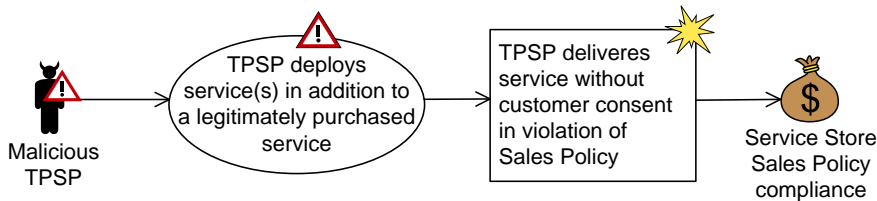


Figure 43 CORAS threat diagram after deletion of treated threat scenarios

Notice that by the language support for the modeling of changing risks that is provided in this deliverable, we may explicitly describe the change of risk picture in one and the same threat diagram. The reader is referred to Figure 134 in Section 15 of the

appendix for the modeling of the risk changes that are exemplified by the shift from Figure 42 to Figure 43.



9 Evaluation

In this section we evaluate the artifacts presented in this deliverable by discussing the extent to which the evaluation criteria presented in Section 2 are fulfilled. The evaluation is structured according to the presentation of the evaluation criteria.

9.1 Scientific Criteria

The scientific evaluation criteria apply to the two WP5 artifacts of the risk assessment methodology and the risk modeling language.

9.1.1 Evaluation of Risk Assessment Methodology

The criteria for the risk assessment methodology are divided into the categories of well-defined methodology, computer-aided methodology and explicit linkage of artifacts.

Well-defined methodology

The risk assessment method should be defined in terms of procedural steps. Such procedural steps are defined by the ISO 31000 standard on risk management [22], and the risk assessment methodology presented in this deliverable instantiates the process defined by this standard. The standard does not provide much level of details for how to conduct the various steps and activities in practices.

The approach of the WP5 work is to extend and generalize traditional methodologies by providing the additional artifacts that are needed for handling change. We therefore refer to other literature for detailed procedural steps for handling the traditional risk assessment problems that still are a substantial part of risk assessment of changing systems. The book on the CORAS method [24], for example, comes with detailed, practical guidelines that are broken down to concrete sub-task. The instantiation in CORAS of the risk assessment method for changing systems presented in Section 13 of the appendix thus leverage on previous work on CORAS. The modeling language and techniques are tightly interwoven with the methodology, which describes the required input and output for each step. The instantiation in CORAS is, however, not completely self-contained, as there in some case may arise specific problems or assessment needs that require other well-known assessment techniques.

When addressing changing systems, there is a need for techniques and guidelines for how to trace changes from target system to risk models, and this requirement is specified by a separate criterion. This is supported by the artifacts of this deliverable by the artifacts and techniques for trace modeling, and the guidelines for how to use the trace model in identifying the parts of the risk picture that are affected by system change.

Computer-aided methodology

The risk assessment method presented in this deliverable, and particularly the instantiation in CORAS, makes extensive use of risk modeling as an assessment



technique. Each of the five activities of the overall process is supported by designated kinds of risk diagrams, and the risk modeling language is defined such that tool support can be provided. This is actually implemented by the prototype of D5.4.

The formal syntax and the underlying semantics of risk graphs and CORAS threat diagrams, the formalization of the notion of dependencies and the rules for likelihood calculation and consistency checking moreover means that the risk assessment method can lend itself to further tool support in many directions.

Explicit linkage of artifacts

The explicit linkage of artifacts that is needed as part of risk assessment of changing system is the linkage between the target of analysis and the risk models. This criterion is fulfilled by the artifact of the trace model that relates system elements and target model elements.

The trace model moreover captures information about the semantics of the system elements by expressing the kind of system model element. Only by consulting the trace model, we can thereby deduce whether a specific change involves system actors, system events or system scenarios.

Furthermore, the precisely and formally defined notion of dependencies in risk models serves as a basis for identifying and reasoning about the propagation of changes through risk models. The risk assessment method proposed in this deliverable comes with guidelines for how to take dependencies into account when changes are traced from the target system to the risk model. In order to increase efficiency and decrease the human effort, tool support should be provided for the full automation of dependency detection.

9.1.2 Evaluation of Risk Modeling Language

The criteria for the risk modeling language are divided into the categories of well-formedness and consistency rules, computer aided support, formal characterization and local usability.

Well-formedness rules and consistency rules of constructs

The syntax of risk graphs is defined by a meta-model, and the notion of well-formed risk graphs is formally defined. As the meta-model that defines the generalization of risk graphs to risk graphs with change does not completely capture all the syntactic constraints, it comes with a set of additional restrictions. Together, the meta-model and the additional restrictions define precisely the set of syntactically correct specifications.

Some consistency checking of specification is supported by the formal foundation of risk graphs by the rules for likelihood reasoning and consistency of likelihood estimates. The underlying semantics of risk graphs is in terms of a probability space on traces. In the risk graphs, the scenarios are specified at a high level with textual descriptions. For the purpose of a more low level investigation of the scenarios, the traces must be explicitly spelled out, for example by means of sequence diagrams.

Computer-aided support for syntactically correct and consistent specifications

Given the formal syntax and the semantics of the risk graphs, the risk modeling can be supported by tools that prevent syntactically incorrect specifications. The D5.4

prototype serves as a proof of concept for the instantiation of the risk modeling with change in the CORAS language.

Although tool support is currently not provided for consistency checking, the syntax and semantics allows some automation, in particular of consistency checking of likelihood estimates.

Formal characterization of specifications

The criterion of the formal characterization of specifications concerns the precise and formal characterization of the behavior that is acceptable or unacceptable. In risk models, the unacceptable behaviors are the scenarios that represent unacceptable risks. Understanding this behavior is of course decisive if the system risks are to be properly understood and if appropriate risk treatments are to be identified.

The textual annotations on the threat scenarios of risk graphs are informal descriptions of the behavior, and in most cases this high-level characterization suffices for risk and treatment identification. If a more formal and detailed characterization is required, the trace sets that describes the behavior may be spelled out, for example by means of sequence diagrams.

Local usability of specifications

Local usability of specifications means that the risk models are self-contained, i.e. that the user can determine the syntactical correctness, the consistency and the semantics of a specification without consulting other artifacts than the specification itself.

Considering the integration of risk assessment with other domains such as testing (cf. Section 8) or requirement engineering (cf. D3.2) local usability is ensured. For example, there is no need for a risk analyst to include requirement models in the documentation of the risk assessment results.

Considering the setting of a risk assessment alone, it is a prerequisite that the involved stakeholders have a common and correct understanding of the target system. This may lead to details being omitted in the risk models when these details are common knowledge or when they are implied. For those that are not in the know, they may need to consult, for example, the target description (which is prepared as part of the context establishment) in order to properly understand the risk models. This is, however, not only a problem of the risk modeling language, but also a problem of risk modeling, as there is always the possibility of making models that are not self-contained, no matter the expressiveness of the language. Because risk graphs (and CORAS risk models) can be specified at any level of abstraction and details, the degree to which the specifications are self-contained is therefore to some extent a matter of choice.

9.2 Industrial Criteria

The industrial criteria are evaluation criteria for the WP5 artifacts in the case studies. The main WP5 case study is ATM, for which a full risk assessment has been conducted. The HOMES case study is addressed to a lesser extent, and therefore provides only some basis for evaluation. In the following we briefly report on the application of the WP5 artifacts in the case studies as this is presented in much more details within WP1.



9.2.1 Evaluation of Risk Assessment Methodology

The first criterion for the risk assessment method is that it is applicable in the case studies. This means that the method can be applied to the risk assessment of changing systems such as the ATM and HOMES case studies, and that the method indeed supports and facilitates the identification, estimation, evaluation and documentation of changing risks.

Notice that the case studies were conducted by using the instantiation in CORAS, as documented in Section 14 and Section 15 in the appendix for ATM and HOMES, respectively.

The ATM risk assessment was conducted over two workshops, each of two full days. The main purpose of risk assessment workshops is to gather personnel with first-hand and expert knowledge about the target of analysis and extract the necessary risk relevant data and information. The extracted information serves as the basis for the risk assessment tasks that are conducted by the risk analysts between the workshops.

The ATM risk assessment demonstrated the applicability of the risk assessment method and techniques. The assessment was lead by two of the researchers involved in developing the WP5 artifacts, and both of these furthermore have background in risk analysis. Properly evaluating the applicability of the risk assessment method would require a case study involving a WP5 external risk analyst that conducted the risk assessment in complete independence.

The applicability of the risk assessment methodology moreover requires that the involved participants and other stakeholders understand the tasks and that they understand the artifacts that serve as input and output to the various tasks. The ATM case study involved the participation of ATM domain experts that are external to the SecureChange project. These were involved in all of the risk assessment activities, which included making the models of changing risks during risk identification and risk estimation. The experience was that the various tasks and their objectives were well-understood, as were the model artifacts that were continuously produced.

The second criterion is that the risk assessment method and its techniques can produce the desired result with less human effort than by using alternative, traditional method. Properly evaluating the human effort criterion would require the same risk assessment to be conducted several times by using different approaches. For now, it is demonstrated that conducting the risk assessment is doable. Due to constraints on time and resources, the ATM risk assessment was very much compressed, and therefore not immediately comparable to real life, industrial risk assessments. Nevertheless, the fact that the parts of the risk assessment results that were not affected by the system changes were identified and therefore not reassessed indicates less effort than if traditional methods with a full risk assessment from scratch were conducted for the target of analysis after the changes.

The HOMES case study did not involve proper analysis workshops, and focused more on extracting realistic examples and scenarios. It nevertheless contributed to demonstrate the applicability of the method, and to demonstrate that the produced artifacts of the models of changing risks are understood by relevant stakeholders.

9.2.2 Evaluation of Risk Modeling Language

The first criterion for the risk modeling language is that it is applicable in the case studies. The risk modeling should result in syntactically correct and consistent specifications that are well understood by the relevant stakeholders.

All the risk modeling activities in the ATM risk assessment were done by the two WP5 researchers that acted as risk analysts. Ensuring consistency and syntactically correct diagrams were therefore not really an issue. Successfully conducting a risk assessment requires, however, that the risk models that are produced are correctly and well-understood by all the involved stakeholders. During the workshops, the risk models were made on-the-fly based on the instructions and information provided by the participants. The participants of the ATM risk assessment included personnel external to SecureChange, and the experience was that they were able to both communicate their opinions and ensure that the risk assessment results were correctly documented in the risk models.

The second criterion is that the modeling of changing risks in the case studies can be conducted with less effort than by using traditional risk modeling languages or techniques. The case studies made use of the generalization of the CORAS language to the setting of changing risks, and it is obvious that using the standard CORAS language would not require less human effort; if anything, the standard CORAS language would require more effort. Comparing CORAS risk modeling with change and standard CORAS risk modeling, the latter requires keeping track of pairs of risk models before and after change, whereas the former explicitly models changing risks in one and the same diagram.

The modeling of changing risks was much less extensive, but nevertheless demonstrated the applicability of the risk modeling language in the HOMES risk assessment.

10 Conclusion

For systems that are changing and evolving, also the risk are changing and evolving and should be understood as such. Traditional risk assessment methods and techniques typically focus on the target of analysis at a particular point in time, and therefore yields risk assessment results that are valid for the current configuration of the target system. Should a potentially risk relevant system change occur, the validity of the risk assessment results can no longer be guaranteed. Considering state-of-the-art approaches to risk assessment, the occurrence of such changes would require a full risk assessment to be conducted from scratch in order to ensure the validity of the results.

In this deliverable we have presented a risk assessment method that meets the methodological needs of assessing changing systems. The guiding principle of the method is that by the occurrence of risk relevant changes, only the parts of the risk picture that may be affected by the changes should be assessed anew. Moreover, in order to properly understand the risks of changing systems as changing risks the method should facilitate the understanding and documentation of the changes to the identified risks.

The deliverable furthermore presents a number of novel risk assessment techniques that support various activities of the risk assessment process for changing systems. The most important of these artifacts is the language for the modeling and documentation of changing risks. This risk modeling language serves as a basis for the further risk assessment techniques of identification, estimation, evaluation and treatment of changing risks.

A further important artifact is the support for establishing and specifying a trace model between the target system and the risk model. The trace model specifies the relations between system elements and risk element and serves as a technique for tracking changes from the target system to the risk model.

The risk assessment process and the risk modeling language as presented in the main part of the deliverable are generic in the sense that they can be instantiated by several specific approaches. In the appendix we present the instantiation of the approach in the CORAS method and language.

The strategic position of the risk assessment method in the general setting of the SecureChange process is demonstrated and exemplified by the integration of risk assessment into the overall Integrated SecureChange process of WP2, by the integration with the requirement engineering method of WP3, and by the integration with testing of WP7. The integration is exemplified in the ATM and HOMES case studies.

These case studies furthermore serve as the case studies of the WP5 work package. The applicability of the risk assessment method and risk modeling is demonstrated by these case studies which are documented in the appendix.

Appendices



11 A – Maintenance Perspective

In this appendix we briefly present a risk assessment method for the risk assessment of changing systems under the maintenance perspective. The method is a specialization of the general and perspective independent method introduced in Section 3, and we focus here on the issues that are specific for the maintenance perspective.

The following scenario exemplifies risk assessment from the maintenance perspective: Some risk assessors conducted and documented an assessment three years ago, and are now requested by the same client to reassess and update the risk picture to reflect changes to the target system or environment, thereby restoring the validity of the assessment.

The changes we address in the maintenance perspective are typically smaller changes that accumulate more or less unnoticed over time, but eventually may have significant impact on the risk picture. Such changes can be bug fixes and security patches, increase in network traffic, increase in the number of attacks, and so forth. In this case, the risk picture remains more or less the same, but risk levels may still have changed such that previously acceptable risks could now become unacceptable, or vice versa. The objective is to maintain the documentation of the previous risk assessment by conducting an update.

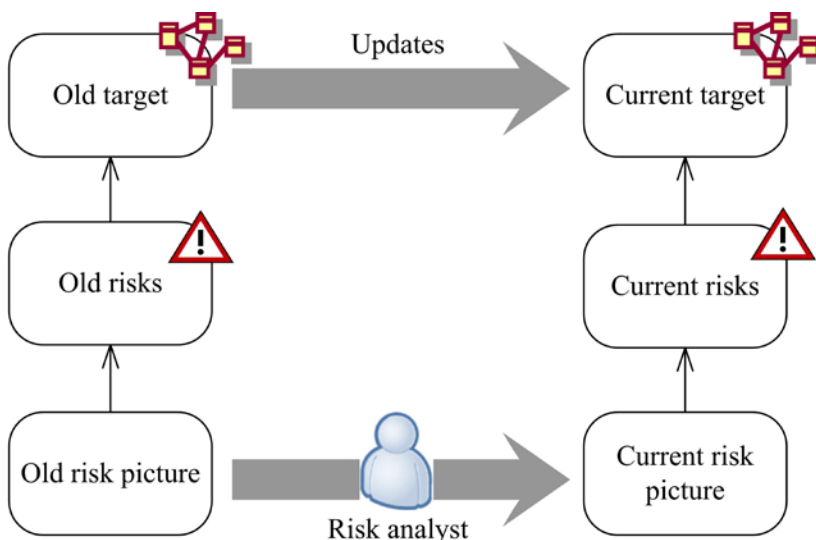


Figure 44 Risk picture in the maintenance perspective

Figure 44 illustrates the principles by which a risk assessment is conducted from the maintenance perspective. Based on the description of the old target of analysis, we make the description of the current target of analysis by doing updates based on the changes that have occurred. We then conduct a new risk assessment only of the parts of the old risk picture that are affected by the changes, thereby identifying the current risks and documenting the current risk picture.

The challenge of risk assessments from the maintenance perspective is to maintain the old documentation of the risk picture without having to conduct the full risk assessment from scratch. The key technique to facilitate this is the trace model introduced in Section 6. The trace model specifies the relations between the target system and the risk models, and facilitates the traceability of changes from the system to the risks.

Identifying such relations and making the trace model is therefore a separate task during the risk assessment from the maintenance perspective. Assuming that a trace model is established and documented for the previous risk assessment, the core of the risk assessment process is illustrated by the UML activity diagram of Figure 45.

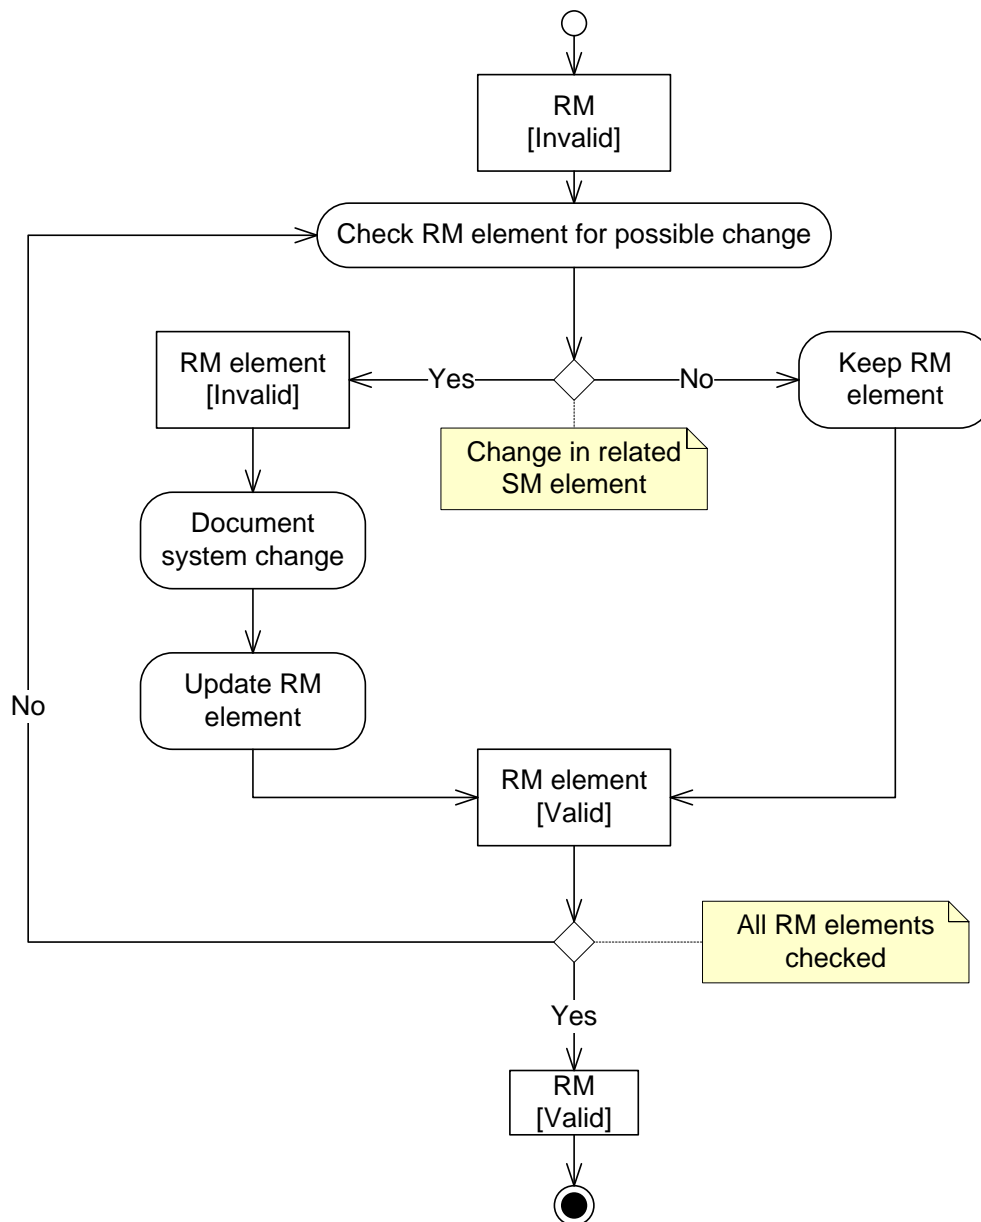


Figure 45 Core of risk assessment process in the maintenance perspective

We assume here that the previous risk assessment is documented by the risk model RM with respect to the system model SM. Due to changes that have occurred since the previous assessment the old risk model may currently be invalid. Based on the trace model, each of the risk model elements must be checked for possible change by checking whether it is affected by system changes. If there is a change in the related system model elements, the system model must be updated accordingly and a new risk assessment must be conducted to update the relevant parts of the risk models, so as to restore the validity of these risk model elements. If there is no change in the related system model elements, the risk model elements are still valid and can be kept in the documentation. When all risk model elements have been checked, the validity of the risk picture has been restored.

12 B – Continuous Evolution Perspective

In this appendix we briefly present a method for the risk assessment of changing systems under the continuous evolution perspective, focusing on the issues that are specific for this perspective when specializing the general method introduced in Section 3.

A scenario for exemplifying risk assessments from the continuous evolution perspective is risk assessors that are requested to predict future evolutions of risks, based on predictions on how the target system will evolve in the future. It mandates that the risk assessors make a dynamic risk picture that reflects the dynamics of the target of analysis.

The kind of changes we address from the continuous evolution perspective is predictable and gradual evolutions that can be described as functions of time. Such predictions can, for example, be based on well-founded forecasts or prognoses, or on planned developments. Examples include a slow increase in the number of components working in parallel, or gradually including more sites in a system. Examples of well-founded forecasts and prognoses are the expected steady increase of end-users, attacks or annual turnover.

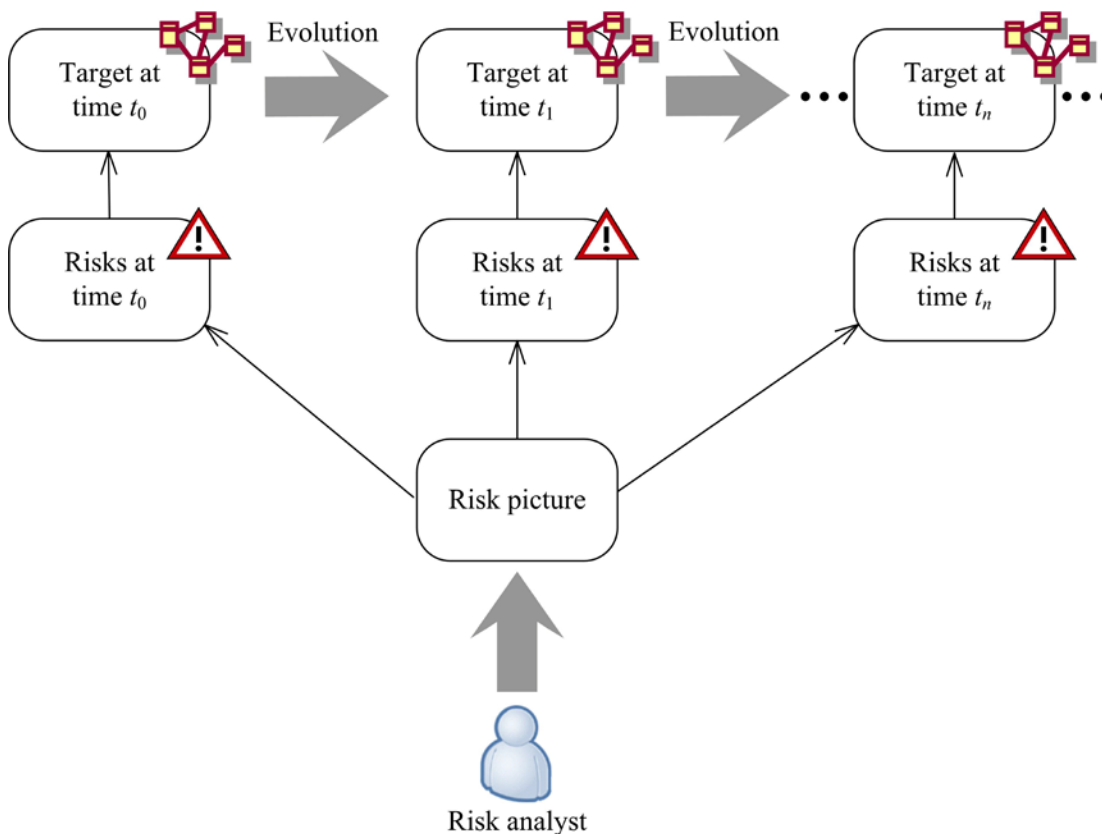


Figure 46 Risk picture in the continuous evolution perspective

Figure 46 illustrates the principles by which a risk assessment is conducted from the continuous evolution perspective. Assuming that we have a description of the target of analysis as a function of time, such that we can extract the (expected) target at any future point in time, we use this as input to the risk assessment. Knowing how the target and its environment evolve, we seek to establish a risk picture as a function of time that shows the according evolvement of risks.

Recall from Section 3 that the risk assessment activities that succeed the context establishment are based on the target description. In the continuous evolution perspective the aim is to understand and represent the risks as evolving risks. It is therefore a prerequisite that the future evolutions of the target of analysis are sufficiently well known to be described as such in the target description.

In order to explain the risk assessment process, let us assume that we have made a target description with a system model *SM* that is a snapshot of the current system with no prediction of the future evolvments. We furthermore assume that we conduct and document a risk assessment based on this target description to produce the risk model *RM*, and that we establish the trace model between the *SM* and the *RM*. In order to enable an assessment of the future evolvments of the risks, the target description must first be generalized so as to characterize the future evolution of the target system and its environment. In other words, we need to transform the system model *SM* to a system model *SM(t)* as a function of time. Based on the trace model we then need to identify the elements of the risk model *RM* that are affected by the system evolutions and make risk predictions that are documented by the generalized risk model *RM(t)* as a function of time.

The core of the risk assessment process from the continuous evolution perspective is illustrated by the UML activity diagram of Figure 47. Starting with the risk model *RM* that shows the current risk picture as a snapshot, we use the trace model to determine for each *RM* element whether the related *SM* elements evolves over time. If so, the *RM* element must be generalized to a representation of an *RM* element that evolves over time. If the *RM* element is related to *SM* elements that do not evolve, the *RM* element can be kept as it is in the risk picture. When all risk model elements have been checked and possibly generalized, the resulting risk model *RM(t)* gives the predictions of how the risks will evolve. The validity of these predictions depends, of course, on the validity of the predictions of the system evolvments.

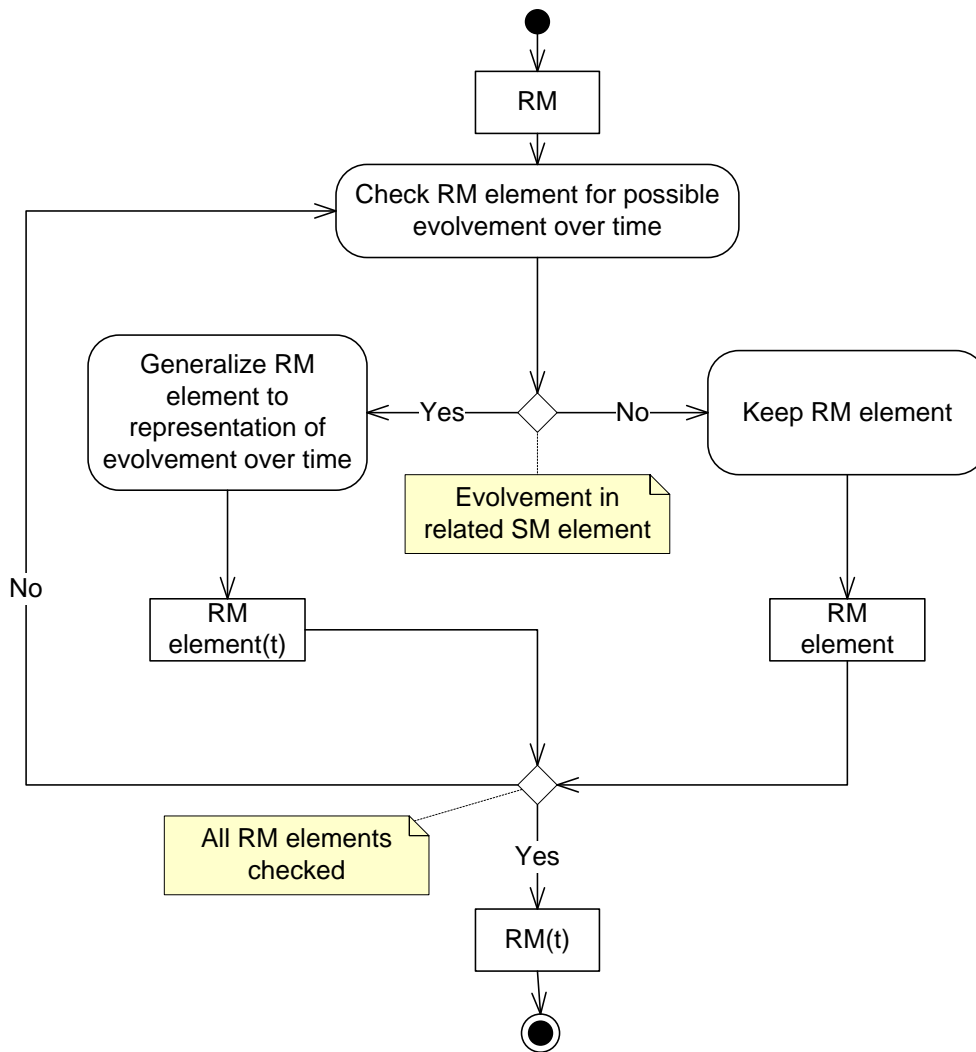


Figure 47 Core of risk assessment process in the continuous evolution perspective

As a small example of a risk assessment under the continuous evolution perspective, consider the assessment of the online store system of a company that does Internet retailing. A fragment of the target description is shown in Figure 48, where we see that the online store includes a web application with an interface towards the customers.

The CORAS threat diagram of Figure 49 shows some of the results of the risk assessment. It shows the unwanted incident of the online store to go down due to software flaw, caused by system developers. The risk level is represented by the likelihood *rare* and its consequence *moderate* for the asset *Online store*. The diagram furthermore documents the asset value *high*.

The CORAS threat diagram depicts parts of the trace model by the annotation to the threat *Developer* and the threat scenario *Developer causes flaw in SW*. These annotations refer to elements of the target model in Figure 48. Notice that due to the dependencies as formally defined in Section 6 all the threat model elements that the threat *Developer* may lead to, possibly via other threat model elements, are related to

the system element of developer. The same is the case for the threat model elements that succeed the threat scenario *Developer causes flaw in SW* with regards to the system element of online store.

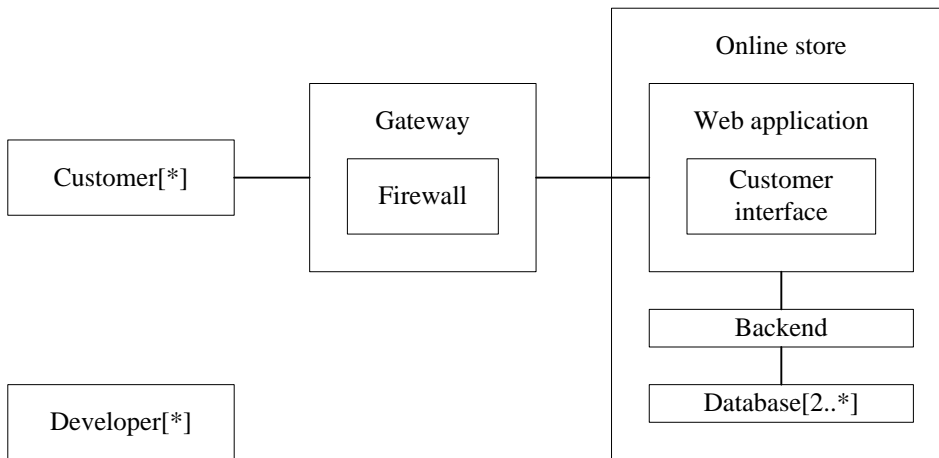


Figure 48 Part of target description of Internet retailing system

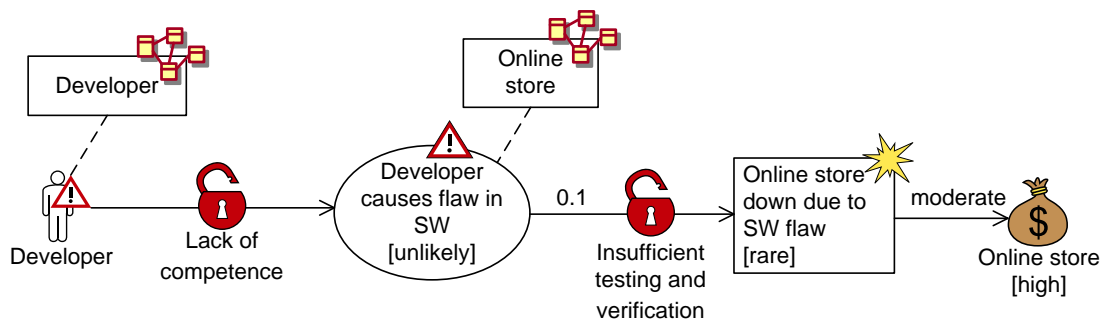


Figure 49 Risks with respect to online store of Internet retailing system

In order to make predictions about how the current risk picture exemplified by the threat diagram will evolve in the future, the target description must first be generalized to the description of an evolving target system as a function of time.

The system owner is a company that has an in-house software development department that is responsible for developing and maintaining the online store and the web application. The company plans to increase the competence of this department hiring new personnel and by offering courses at a regular basis. Based on this plan, they make predictions about how the level of competence will evolve in the future. The company furthermore makes some predictions about future software testing time and number of expected bugs. Finally, they have some expectations and prognoses about the future sales volume, which will affect the future asset values and potential for losses in case something goes wrong.

Predictions and prognoses such as these are included in the target description as indicators, which are values the level of which can be given as functions of time and possibly other indicators. This is shown for the online store example in Figure 50. Based on this description of the target as a function of time, as well as the trace model and the risk model dependencies, the risk model need to be generalized so as to characterize the identified risks as evolving risks.

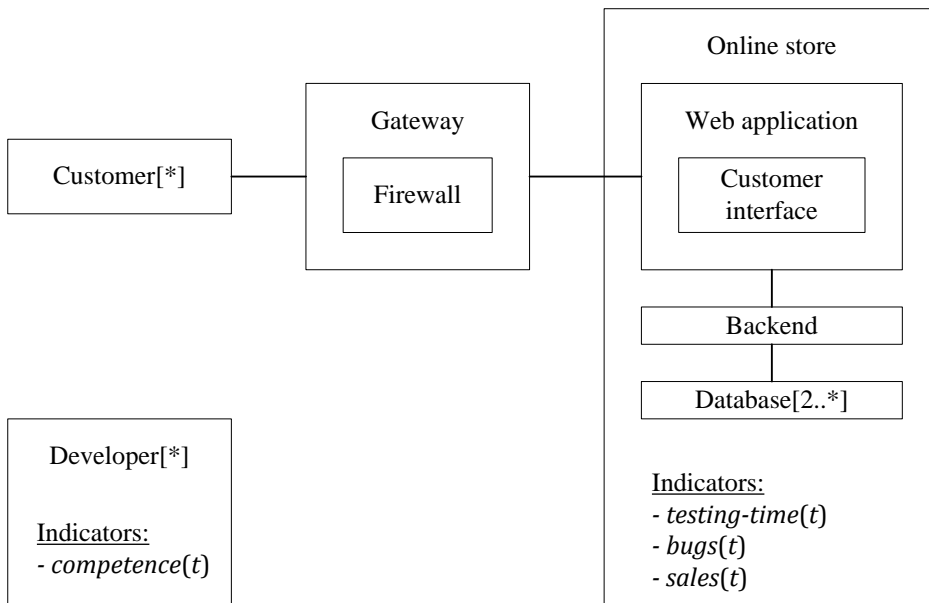


Figure 50 Part of target description of evolving Internet retailing system

The generalized risk model is exemplified by the threat diagram of Figure 51. Instead of the previous likelihoods and consequence estimates and asset value are now functions to deduce these for any point in time.

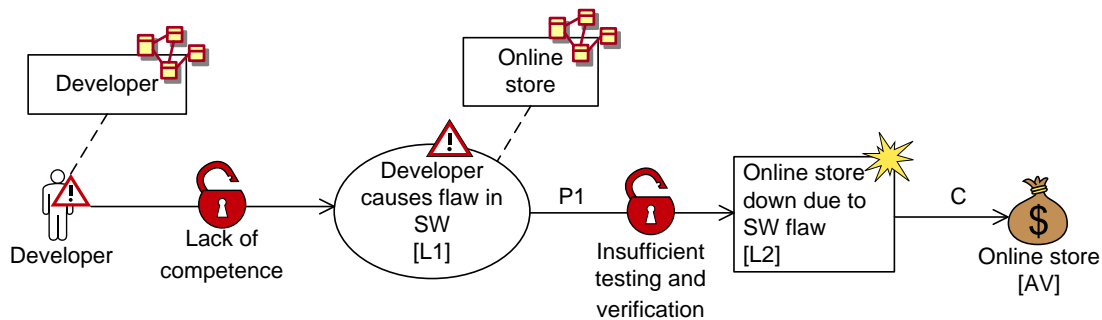


Figure 51 Evolving risks with respect to online store of evolving Internet retailing system

The evolution of the risks depends on the evolution of the target of analysis, which is captured by the indicators. Given the trace model, the threat diagram dependencies and the indicators, the functions to derive the values in the threat diagram may, for example, be as follows for some sensible functions f .

$$L1 = f_1(\text{competence}(t))$$

$$P1 = f_2(\text{testing-time}(t))$$

$$L2 = f_3(L1, P2, \text{bugs}(t))$$

$$C = 0.3 \cdot AV$$

$$AV = f_4(\text{sales}(t))$$

13 C – Instantiation of Method in CORAS

In this appendix we present an instantiation of the method for the risk assessment of changing systems in CORAS, particularly the method for risk assessment under the before-after perspective introduced in Section 7.

Basically, the risk assessment process remains the same, independent of the specific instantiation. That is to say, instantiating the method in CORAS means that the risk assessment is a process of the five activities of context establishment, risk identification, risk estimation, risk evaluation and risk treatment.

A characteristic feature of CORAS is the tight integration of the activities of the risk assessment process and the CORAS risk modeling language. The risk modeling language consists of five kinds of diagrams that serve as a basis for many of the risk assessment techniques that are used during the risk assessment process. A large part of instantiating the method for risk assessment of changing systems in CORAS therefore amounts to instantiating the risk modeling techniques in the CORAS language.

In Section 13.1 we explain the relation between CORAS threat diagrams and risk graphs. CORAS threat diagrams are one of the five kinds of CORAS diagrams, and are the most important diagrams in risk assessment using CORAS. In particular, we explain and show how threat diagrams instantiate risk graphs, and thereby also the rules for likelihood reasoning presented in Section 5 are instantiated. In Section 13.2 we extend the CORAS language to the setting of changing risks, similar to the extension of risk graphs in Section 5. In Section 13.3 we explain how to use the extended CORAS language in the process of assessing changing risks.

13.1 CORAS Threat Diagrams as Specialized Risk Graphs

CORAS threat diagrams are intensively used during risk assessment to facilitate risk identification and risk estimation. The diagrams are furthermore used as a part of the documentation and reporting of the assessment results. The diagrams describe how threats may exploit vulnerabilities to initiate threat scenarios, how threat scenarios may lead to unwanted incidents or other threat scenarios, and which assets that are harmed by the unwanted incidents. The language constructs are threats (deliberate, accidental and non-human), vulnerabilities, threat scenarios, unwanted incidents and assets as depicted in Figure 52. Threat scenarios and incidents can be annotated with likelihoods.

There are furthermore three kinds of relations in threat diagrams, namely initiates relations, leads-to relations and impacts relations. An initiates relation has a threat as source and a threat scenario or unwanted incidents as target. It can be annotated with a likelihood that describes the likelihood for the threat to initiate the related scenario or incident. A leads-to relation has a threat scenario or unwanted incident as both source and target. It can be annotated with a conditional likelihood. An impacts relation has an

unwanted incident as source and an asset as target, and can be annotated with a consequence value that describes the harm of the incident on the asset when the incident occurs.

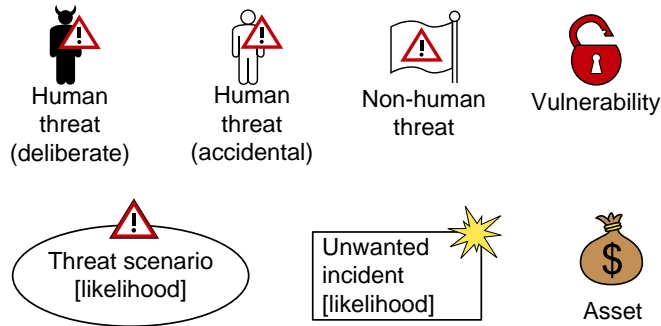


Figure 52 Constructs of CORAS threat diagrams

Figure 53 depicts an example of a threat diagram. In fact, this threat diagram shows the scenarios that are modeled by means of a risk graph in Figure 10 of Section 5 only with qualitative likelihood values on the scenarios and incident instead of exact probabilities. While the same set of scenarios and relations between them are depicted in the two diagrams, there are some significant differences. The threat diagram explicitly shows the initial threats, it distinguishes the incident *Data exposed* from the other scenarios as an unwanted incident, and it explicitly shows the asset that is harmed.

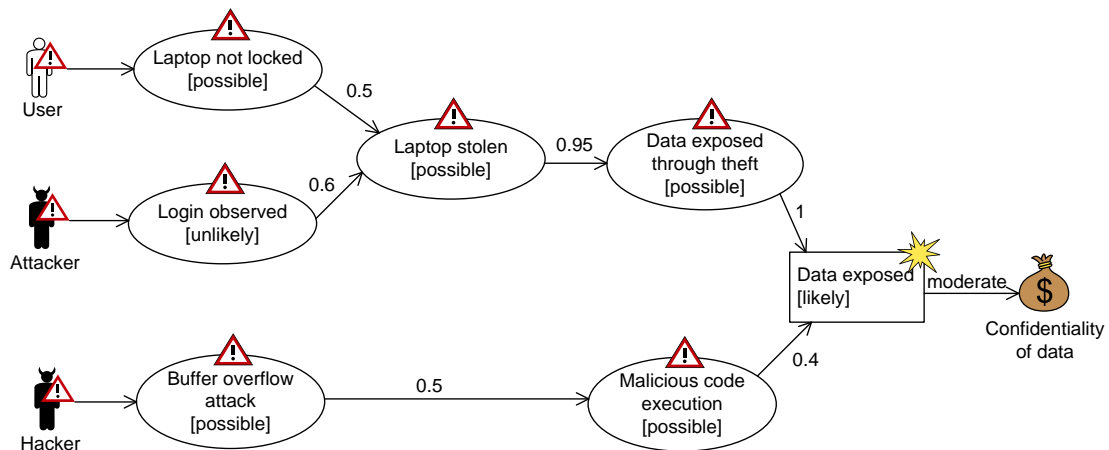


Figure 53 Example CORAS threat diagram

The differences between threat diagrams and risk graphs are summarized as follows:

- Initiate relations and leads-to relations in threat diagrams can be annotated with vulnerabilities, while the relations in risk graphs cannot.
- Threat diagrams distinguish between four kinds of vertices, namely threats, threat scenarios, unwanted incidents and assets, while risk graphs have only scenarios.

- Threat diagrams distinguish between three kinds of relations, namely initiates relations, leads-to relations and impacts relations, while risk graphs have only leads-to relations.

Given the differences between threat diagrams and risk graphs, the techniques for reasoning about likelihoods and dependencies still carry over to the CORAS instantiation. The vulnerabilities are mere annotations on relations, and can be ignored in the formal representation of the diagrams. Moreover, the various vertices and relations of threat diagrams can be interpreted as special instances of the risk graph vertex and relation:

- An unwanted incident of a threat diagram is interpreted as a scenario of a risk graph.
- A set of threats t_1, \dots, t_n with initiates relations to the same threat scenario s is interpreted as follows: The threat scenario s is decomposed into n parts, where each resulting sub-scenario s_j , $j \in \{1, \dots, n\}$, corresponds to the part of s that is initiated by threat t_j . As a threat is not an event, but rather an actor, it cannot be assigned a likelihood. Instead, the initiate relation from the threat may be so. A threat t_j with initiates relation of likelihood l_j to sub-scenario s_j is then combined into the risk graph scenario *Threat t_j initiates s_j with likelihood l_j* .
- An impacts relation from unwanted incident u to asset a with consequence c in a threat diagram is interpreted as follows: The impacts relation is interpreted as a risk graph relation with likelihood 1; the asset a is interpreted as the risk graph scenario *Incident u harms asset a with consequence c* .

With this interpretation, we refer to Section 5 and Section 6 for the techniques for reasoning about likelihoods and dependencies, respectively, in CORAS threat diagrams. Notice only that Rule 1 (Relation) applies to the CORAS leads-to relations only and that Rule 2 (Mutually exclusive vertices) and Rule 3 (Independent vertices) apply to the CORAS threat scenarios and unwanted incidents.

With the above interpretation of CORAS threat diagrams as risk graphs we can use Rule 1 to reason about the likelihoods of threats initiating threat scenarios or unwanted incidents as annotated on the initiates relations. However, in order to allow all likelihood reasoning to be conducted directly in CORAS diagrams, we introduce a separate rule for the initiates relation. We let t denote a threat, v denote a vertex (threat scenario or incident), and $t \rightarrow v$ denote the initiates relation from threat t to vertex v .

Rule 4 (Initiates). If there is an initiates relation from threat t to vertex v , we have:

$$\frac{t \xrightarrow{P} v}{(t \sqcap v)(P)}$$

By $t \sqcap v$ we denote the occurrences of vertex v that are initiated by the threat t .

13.2 Generalizing the CORAS Language to Changing Risks

In this section we define the language extensions for generalizing the CORAS language to the setting of changing risks in the before-after perspective. The definition of the syntax extends the definition of the standard CORAS syntax [24]. The extension is defined in exactly the same way as for the risk graphs in Section 5.2, and we therefore refer to that section for further explanations. In addition to generalizing the CORAS language to handle the modeling of changing risks, we extend it with the construct for relating risk model elements and system elements, i.e. with the construct for visualizing the trace model.

13.2.1 Standard CORAS Diagrams

The CORAS language consists of five kinds of diagrams, and we define the extension for all of them. The diagrams are asset diagrams, threat diagrams, risk diagrams, treatment diagrams, and treatment overview diagrams. Before we define the extensions, we briefly introduce each of them.

An asset diagram is used for defining and documenting the assets of a risk assessment. An asset is always associated with a party (stakeholder), which is the organization, company, person, group or other body on whose behalf the risk assessment is conducted. Because there may be several parties in one risk assessment, the party in question is explicitly shown in the asset diagram. Asset diagrams can furthermore specify dependencies between assets, and they can specify asset values.

Figure 54 shows an example of an asset diagram. It is extracted from the ATM risk assessment case study that is fully documented in Section 14. The party of the risk assessment is the ATM service provider, and the assets are availability and confidentiality of arrival management information.




Figure 54 Example asset diagram

A threat diagram is used during risk identification and risk estimation. The constructs and relations are introduced and explained above, and we therefore only give a further example. The threat diagram of Figure 55 shows some of the results of the ATM risk identification and risk estimation. The threats, vulnerability and threat scenarios explain some of the issues that can lead to the occurrence of two unwanted incidents. Each of them represents a risk, and the risk estimates are given by their likelihoods and their consequences for the assets.

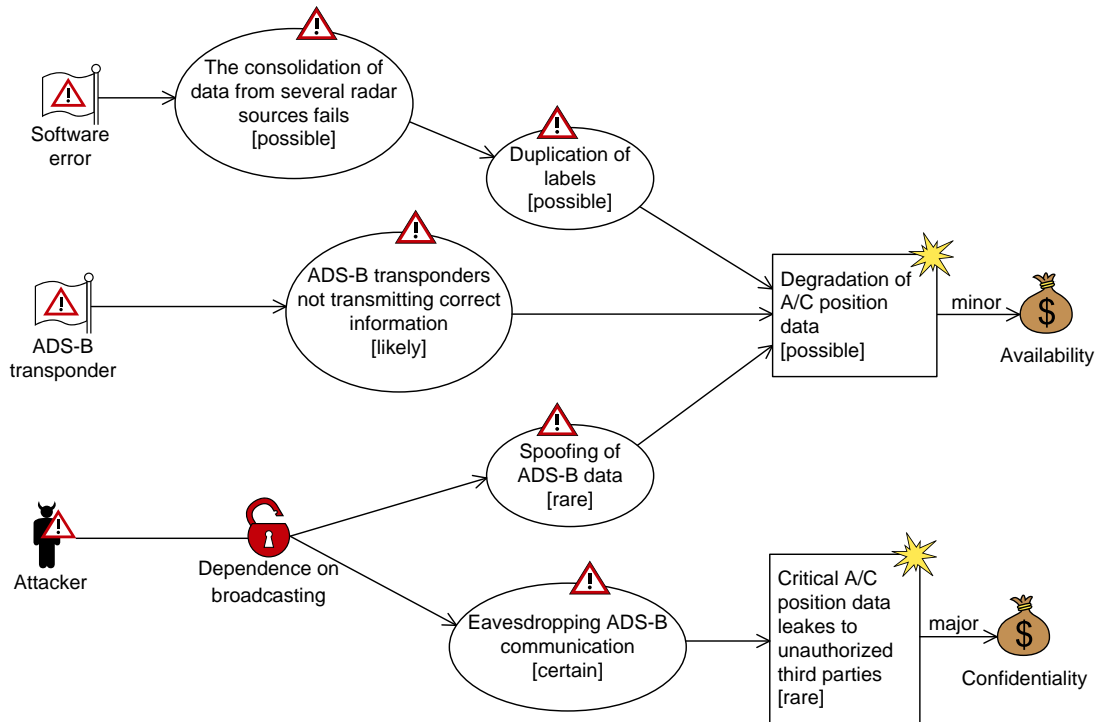


Figure 55 Example threat diagram

A risk is the likelihood of an unwanted incident and its consequence for a specific asset. Because one unwanted incident may harm several assets when it occurs, each pair of an unwanted incident and an asset constitutes a risk. A risk diagram is used to explicitly show all the risks of a threat diagram, where each pair of an unwanted incident and an asset from the threat diagram is replaced by a separate risk symbol. The risk can be annotated with the risk level as derived from the likelihood and consequence of the unwanted incident for the asset in question. The purpose of the risk diagrams is to give an overview of the risks, and therefore shows only the threats that initiate them and the assets they harm.

The risk diagram in Figure 56 is derived from the threat diagram of Figure 55. Since the two unwanted incidents harm only one asset each, they also constitute only one risk each.

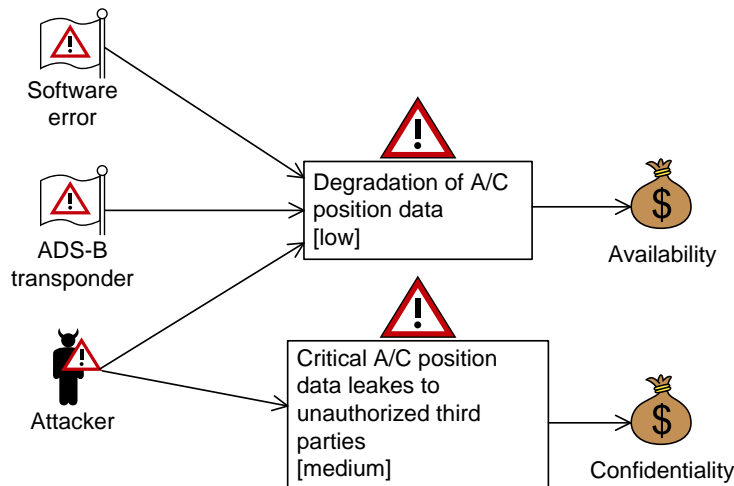


Figure 56 Example risk diagram

Treatment diagrams are a kind of extended threat diagrams that are used to identify and document treatments to unacceptable risks. Basically, a treatment diagram is a threat diagram annotated with treatments, with the difference that the unwanted incidents of the threat diagram are replaced by the risks that they constitute.

The treatment diagram of Figure 57 shows some treatment options for the risks that were identified and documented in the threat diagram of Figure 55. In this case we have shown only the parts of the threat diagram that are relevant for the treatments. During risk assessments, we commonly also remove all the parts of the threat diagrams that do not contribute to the unacceptable risks. Showing only the threats, vulnerabilities and threat scenarios that may lead to the unacceptable risks facilitates the identification of treatments for the parts that really matter.

A treatment overview diagram is similar to a risk diagram, and can be understood as a collapsed version of a treatment diagram. The purpose is to give an overview of the treatments and the risks that they mitigate.

An example of a treatment overview diagram is given in Figure 58. Notice that the relations from the treatments point directly on the risks that they mitigate. For example, because the treatment *Implement encryption of ADS-B signals* is a treatment for the threat scenario *Spoofing of ADS-B data*, it is indirectly a treatment of the risk *Degradation of A/C position data* since the threat diagram may lead to the risk.

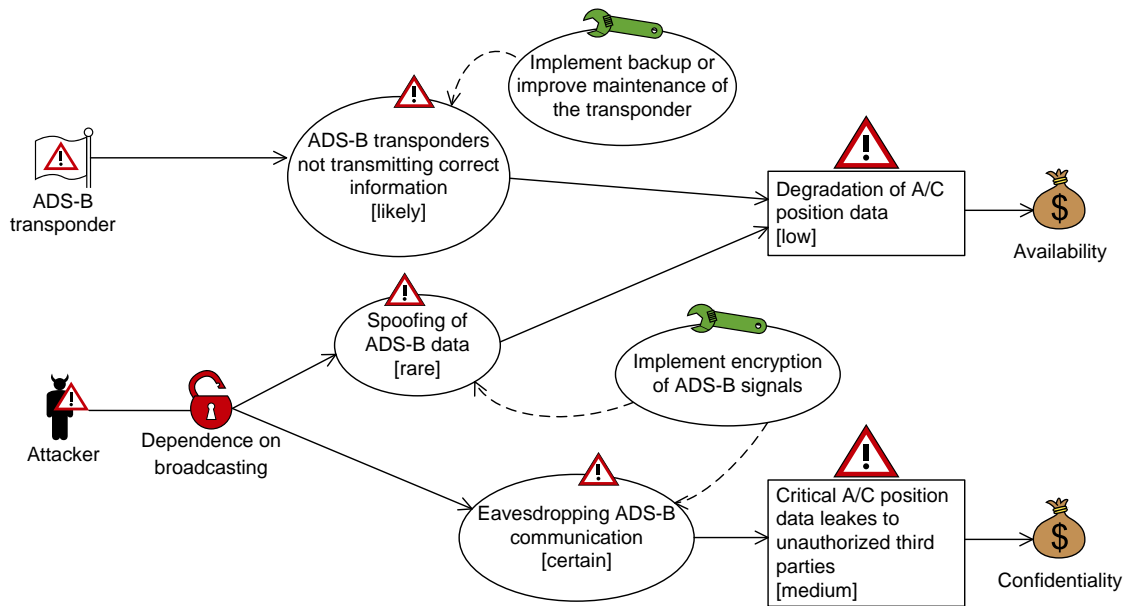


Figure 57 Example treatment diagram

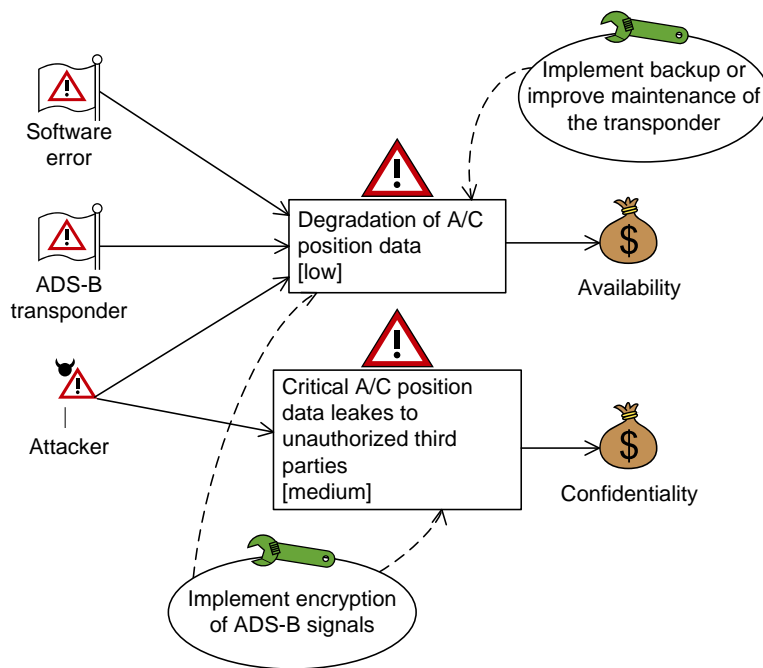


Figure 58 Example treatment overview diagram

13.2.2 CORAS Diagrams with Change

In this section we define the extension of the syntax of the CORAS language to provide the expressiveness for the modeling of changing risks. It is defined as an extension of the CORAS language syntax defined in [24].

The principle by which the extension is defined is exactly as the principle for the extension of risk graphs defined in Section 5. The meta-model for standard risk graphs is given in Figure 9, and the extension to risk graphs with change is defined by the meta-model in Figure 11. The extension introduces the mode attribute for risk graph vertices and relations, where the mode is one of *before*, *after* and *before-after*. In Section 6 the risk graph syntax is further extended to include the construct for annotating risk graph vertices with references to related system elements, the meta-model of which is given in Figure 21.

The meta-model for the standard CORAS language [24] is a specialization of the meta-model for traditional risk graphs in Figure 9. The extension of the CORAS language defined in the following results in a meta-model that similarly can be understood as a specialization of the risk graphs meta-model in Figure 21. We define the meta-model for the CORAS language elements with change and the CORAS language relations with changes separately. Thereafter we combine them in the definition of the meta-model for the CORAS language with change.

The meta-model for the CORAS language with change is implemented by the prototype tool of WP5 deliverable D5.4.

13.2.2.1 Meta-model for CORAS elements with Change

The class diagram of Figure 59 defines the meta-model for the elements of the CORAS before-after language. The class *Element* has an identifier and a mode. The mode attribute of an element has one of the values of the set {before, after, before-after}. If the mode attribute of an element is “before”, the element represents a fragment of the risk picture before the change requirements in question have been implemented. If the mode attribute is “after”, the element represents a fragment of the risk picture after the change requirements have been implemented, and if the mode attribute is “before-after” the element represents a fragment of the risk picture that is relevant both before and after the changes.

The concrete classes *Threat scenario* and *Unwanted incident* are *Elements with likelihood*. By the definition of this abstract class, these concrete classes have a pair of likelihood values, namely Likelihood before and likelihood after. It is only when the elements with likelihood are in mode before-after that they can be assigned both likelihood before and likelihood after. If the mode is before, the likelihood after is undefined. Similarly, if the mode is after, the likelihood before is undefined.

Risk and *Asset* are classes that likewise have pairs of attributes, but it is only in mode before-after that both attributes may have defined values.

In addition to the extension of the standard CORAS language to represent the three possible modes of the diagram elements, the abstract class *Target element* with the specialization of the concrete class *Target segment* is added. The target segment is also in one of the three modes and serves as a reference to a segment of the target description.

Together, we have the following additional restrictions that are not captured by the meta-model:

- If the mode of *Element with likelihood* is “before”, the attribute “Likelihood after” is undefined.

- If the mode of *Element with likelihood* is “after”, the attribute “Likelihood before” is undefined.
- If the mode of Risk is “before”, the attribute “Risk level after” is undefined.
- If the mode of Risk is “after”, the attribute “Risk level before” is undefined.
- If the mode of Asset is “before”, the attribute “Asset value after” is undefined.
- If the mode of Asset is “after”, the attribute “Asset value before” is undefined.

13.2.2.2 Meta-model for CORAS relations with Change

The class diagram of Figure 60 defines the meta-model for the relations of the CORAS before-after language. As for the language elements, also the relations are in one of the modes “before”, “after” and “before-after”. The abstract class *Relation with likelihood* has one attribute for likelihood before and one attribute for likelihood after. If the mode is “before”, the likelihood after is undefined, and vice versa. The same is the case for *Relation with consequence* and the attributes for consequence before and consequence after. Together, this gives the following restrictions that are not captured by the meta-model:

- If the mode of *Relation with likelihood* is “before”, the attribute “Likelihood after” is undefined.
- If the mode of *Relation with likelihood* is “after”, the attribute “Likelihood before” is undefined.
- If the mode of *Relation with consequence* is “before”, the attribute “Consequence after” is undefined.
- If the mode of *Relation with consequence* is “after”, the attribute “Consequence before” is undefined.

We have the following restrictions on the relations depending on the modes of the source or target:

- If the mode of the target is “before”, the mode of the relation is “before”.
- If the mode of the target is “after”, the mode of the relation is “after”.
- If the mode of the source is “before”, the mode of the relation is “before”.
- If the mode of the source is “after”, the mode of the relation is “after”.

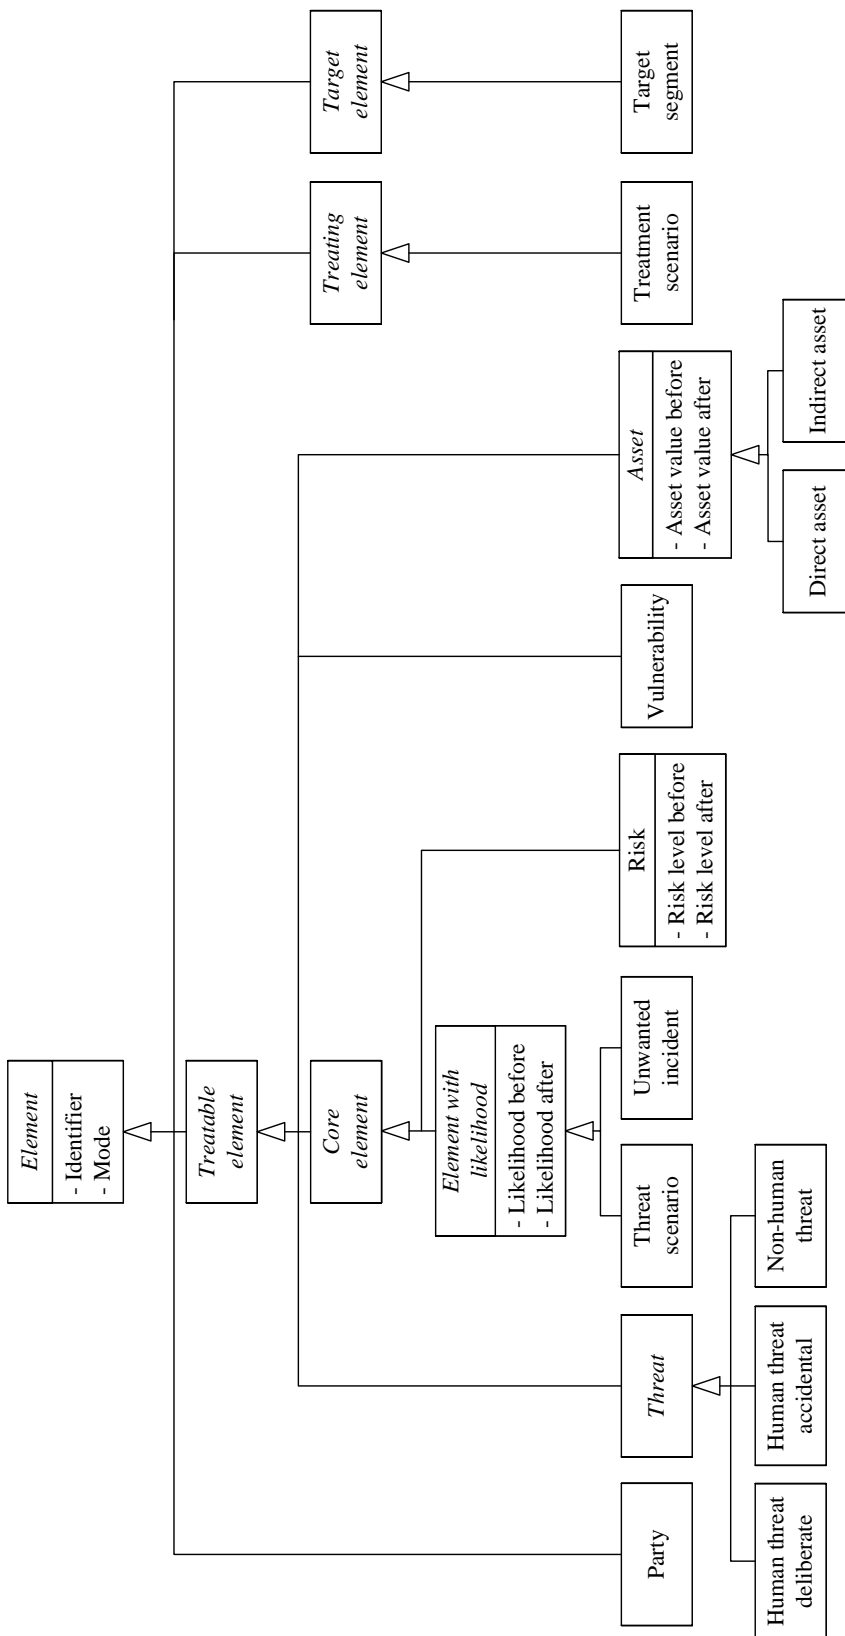


Figure 59 Meta-model for CORAS elements with change

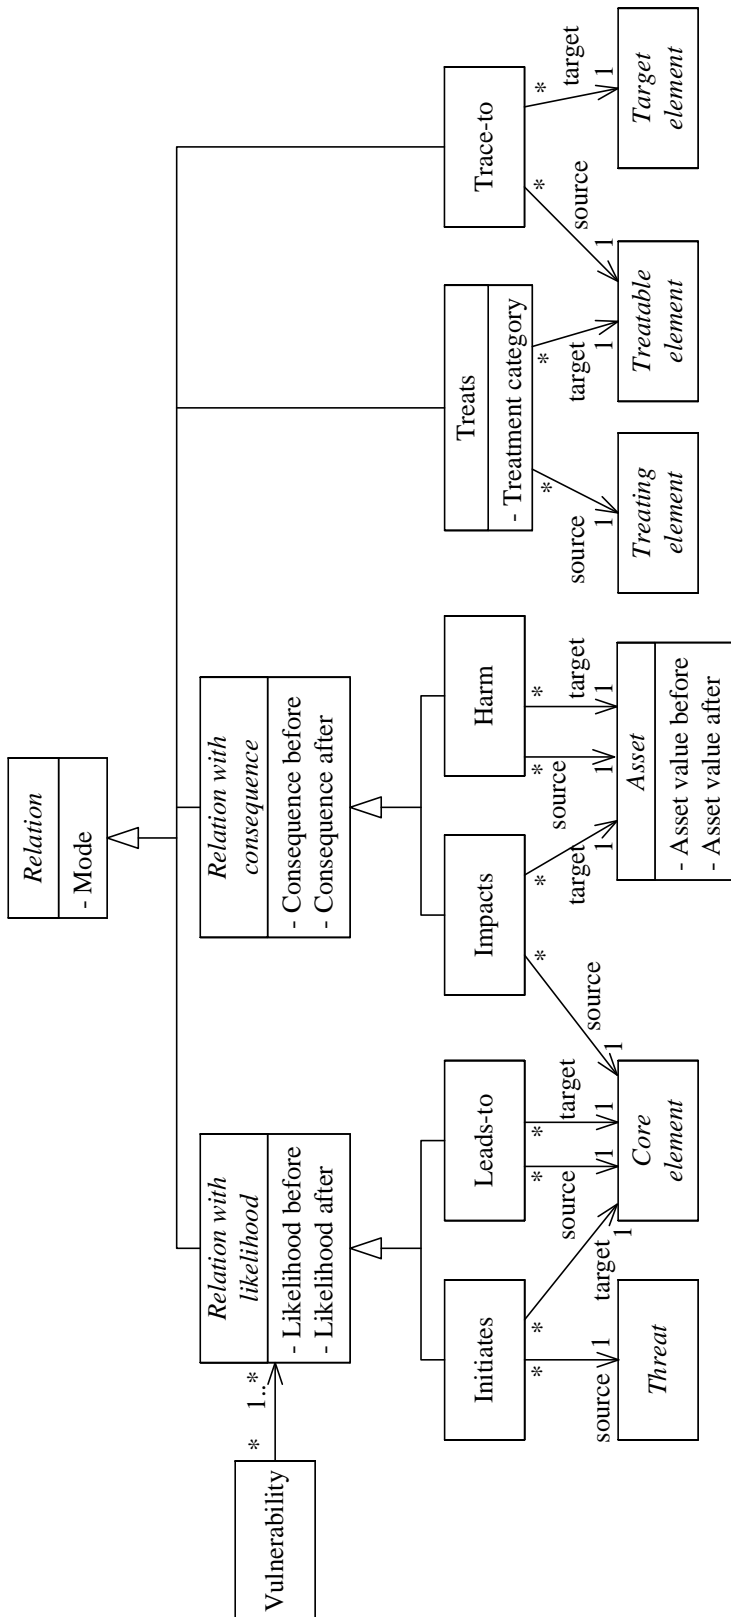


Figure 60 Meta-model for CORAS relations with change

13.2.2.3 Meta-model for CORAS diagrams with Change

The meta-model for CORAS diagrams with change is given in Figure 61. A diagram is composed of one or more elements and zero or more relations. The diagrams are defined by the abstract class *CORAS diagram change*, which is specialized into the five concrete diagrams of the language.

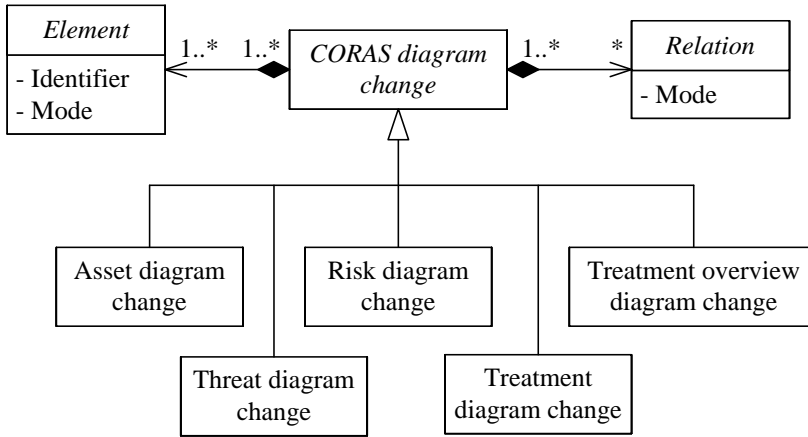





Figure 61 Meta-model for CORAS diagrams with change

For the concrete diagrams there are some restrictions that apply that are not captured by the meta-models:

- Asset diagrams can contain only parties, assets, and harm relations, and must contain exactly one party and at least one asset.
- Threat diagrams can contain any elements or relations except for parties, risks, treatment scenarios, and treats relations.
- Risk diagrams can contain only threats, risks, assets, initiates relations, leads-to relations and impacts relations.
- Treatment diagrams can contain any elements or relations except parties.
- Treatment overview diagrams can contain only threats, risks, assets, treatments, initiates relations, leads-to relations, impacts relations, and treats relations.

The graphical symbols for the various CORAS language elements as they appear in the diagrams are shown in Table 29. We use grey color for the elements that depict parts of the risk picture before changes, we use the standard, colored CORAS symbols for the elements that depict parts of the risk picture after the changes, and we use the two-layered symbols for the elements that depict parts of the risk picture both before and after the changes.

Element name	Before	After	Before-after
Party			

















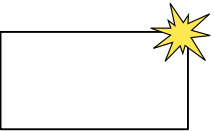







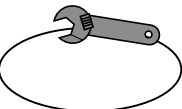
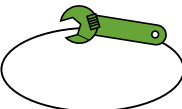

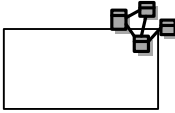
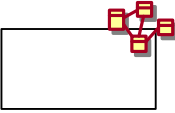
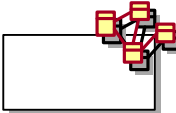
Asset			
Human threat (deliberate)			
Human threat (accidental)			
Non-human threat			
Threat scenario			
Unwanted incident			
Vulnerability			
Risk			
Treatment			
Target segment			

Table 29 Overview of CORAS language elements with change

13.2.2.4 Example of CORAS Diagram with Change

The ATM risk assessment that is reported in Section 14 was conducted using the CORAS instantiation of the risk assessment method for changing systems. We refer to that section for numerous examples of the use of CORAS diagrams with change. In this section we reuse the risk graph example from Section 5.

Figure 12 shows an example of using risk graphs with change for modeling the risk of data exposure before and after certain changes. The same before-after scenarios are modeled by using CORAS threat diagrams with change in Figure 62. The unwanted incident *Data exposed* occurs with the same likelihood both before and after changes. However, the threat scenario of unlocked laptop occurs only before the changes, whereas buffer overflow attack and malicious code execution occurs only after the changes. The threat scenario *Login observed* and its subsequent scenarios occur both before and after the changes. Notice the changes of the likelihoods of the threat scenarios *Laptop stolen* and *Data exposed through theft* from before to after.

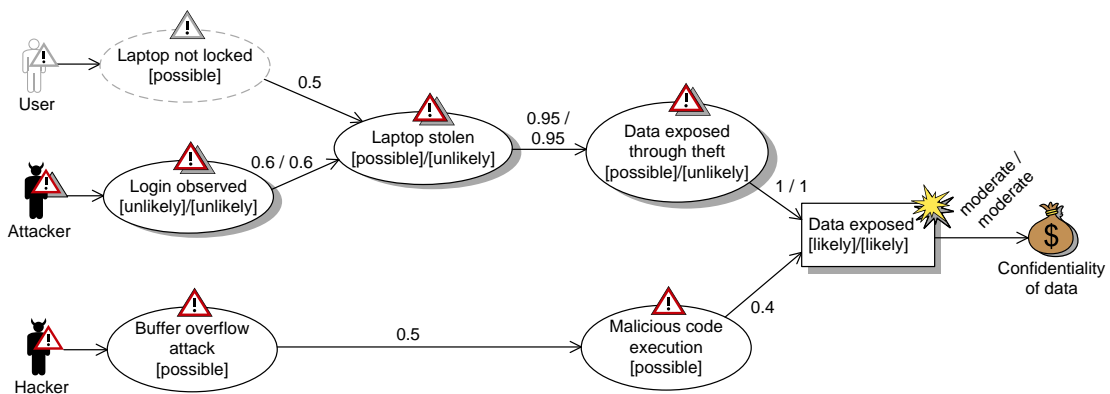


Figure 62 Example CORAS threat diagram with change

We explained in Section 5 that the risk graphs with change can be understood as syntactic sugar for two separate risk graphs, one risk graph showing risks before changes and one risk graph showing risks after changes. This was exemplified for the risk graph with change in Figure 12 by the separate traditional risk graphs of Figure 13 and Figure 14, respectively. This understanding of risk modeling with change as the combination of diagrams applies also to the CORAS language with change.

The threat diagram of Figure 63 is (apart from the appearance of the symbols) a standard CORAS threat diagrams that shows the before part of the threat diagram with change of Figure 62. The threat diagram of Figure 64 shows the after part of the threat diagram with change.



Figure 63 CORAS threat diagram before change

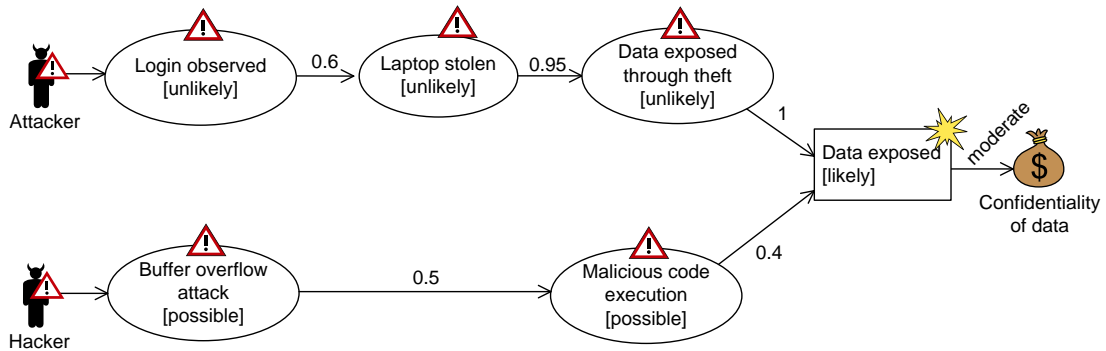


Figure 64 CORAS threat diagram after change

13.3 Assessment of Changing Risks using CORAS

The method for the risk assessment of changing and evolving systems is presented and exemplified in Section 7 (for the before-after perspective). Instantiating the method in CORAS does not introduce any changes to the overall process and its activities and tasks. However, by instantiating the risk modeling technique of risk graphs in CORAS, we are provided a number of risk assessment techniques that facilitates each of the five activities of the risk assessment process.

The CORAS approach offers a risk assessment method, a risk modeling language and a risk assessment tool, and these three parts are tightly integrated in the overall risk assessment process. In the following we explain the use of the CORAS language with change in the risk assessment process, and compare this with the techniques in Section 7. We moreover refer to Section 14 which uses the CORAS instantiation on the ATM risk assessment.

13.3.1 Context Establishment using CORAS Asset Diagrams

The context establishment of a risk assessment of change instantiated in CORAS is conducted according to the method of Section 7, with one exception: With CORAS, we use asset diagrams as a technique for asset identification.

Asset identification is a technique for determining and pinpointing the focus of the risk assessment. CORAS is an asset-driven risk assessment process, which means that all

the activities that follow the context establishment are directed towards the identified assets. This ensures that the risk assessment focuses on the issues that really matter.

In Section 7 the assets, asset value/priority and related parties are documented using a table format. Using CORAS we document the assets for each party in turn. In addition to the documentation of party, assets and asset values, asset diagrams support the identification and documentation of relations between assets. These relations describe which assets that may be harmed via harm to other assets, and are useful for understanding the wider impact of risks, as well as the relative importance of assets. An example of an asset diagram is given in Figure 101 in Section 14.

13.3.2 Risk Identification using CORAS Threat Diagrams

Risk identification involves the identification of risk with respect to the identified assets, and the identification of potential sources of these risks, as well as the documentation of the results, both before and after changes. Understanding the risks and their sources includes an understanding of the unwanted incidents that may occur, the assets they harm, the threat scenarios that may lead to unwanted incidents, the threats that initiate the threat scenarios, and the vulnerabilities that opens for these things to occur.

The approach to risk identification presented in Section 7 uses risk graphs with change as a technique for identifying and documenting risks. The approach with CORAS is similar, but with the richer expressiveness of the CORAS language, we can explicitly model the various aforementioned aspects that constitute the risk picture. In fact, the CORAS language is closely related to the underlying concepts of risk assessment. Risk graphs, on the other hand, only describe scenarios that may occur and other scenarios that these may lead to.

The CORAS approach uses structured brainstorming as a technique for risk identification, involving personnel of various backgrounds and with expert knowledge about the target of analysis. The findings are documented on-the-fly using CORAS threat diagrams. An important purpose of the graphical icons and the relations between them is to make intuitive and easily understandable risk models that serve as a basis for the discussions.

Some of the results of the risk identification are documented in the risk graphs with change of Figure 31 and Figure 32. The threat diagrams of Figure 107 and Figure 112 documents some of the same findings with the CORAS approach.

13.3.3 Risk Estimation using CORAS Threat Diagrams

Risk estimation involves the estimation of likelihoods and consequences of unwanted incidents, as well as the estimation of the likelihoods for the occurrence of the scenarios that may lead to the unwanted incidents. Conditional likelihoods may furthermore be estimated for the leads-to relations between scenarios and unwanted incidents.



Likelihood estimation and documentation with CORAS is conducted as explained in Section 7 with risk graphs by annotating vertices. The difference is that CORAS distinguishes between threat scenarios and unwanted incidents, and that CORAS allows the specification of likelihoods for threats to initiate scenarios as annotations on initiates relations.

Risk graphs do not provide explicit support for consequence estimation and documentation, however. In Section 7 the consequence estimates are documented by using tables. In CORAS, consequence estimation and documentation are supported in the language by annotating the impacts relations from unwanted incidents to assets with consequences. Threat diagrams are therefore offers techniques for both risk identification and risk estimation, and also serves as a means for documenting the results.

Several examples of likelihood and consequence estimates documented in CORAS threat diagrams with change are provided in Section 14.3.

13.3.4 Risk Evaluation using CORAS Risk Diagrams

Risk evaluation involves calculating the risk levels based on the likelihood and consequence estimates, and comparing the results against the risk evaluation criteria. The approach with CORAS is the same as the approach presented in Section 7. However, whereas the risk levels before and after the changes are calculated and documented in tables in Section 7, the CORAS approach uses risk diagrams.

Risk diagrams document all the identified risks, and can be annotated with risk levels. If desired, the risks can also be annotated with the results of the risk evaluation, explicitly showing which risks are acceptable and which risk that need to be evaluated for possible treatment. In addition to serving as a technique for risk evaluation, the CORAS risk diagrams provide an overview of the identified risks and how they change by depicting only the risks together with the threats that initiate them and the assets they harm.

Several examples of risk evaluations documented in CORAS risk diagrams with change are provided in Section 14.4.

13.3.5 Risk Treatment using CORAS Treatment Diagrams

Risk treatment is the identification of treatment options for the unacceptable risks. This is conducted as a structured brainstorming with a walkthrough of the risk models that documents the risk that are unacceptable according to the risk evaluation criteria.

In Section 7 the identified treatments are documented in a table format that for each treatment lists the threat scenarios that are mitigated by the treatment. In the CORAS approach we use treatment diagrams as a technique for treatment identification and documentation. In these diagrams the identified treatments are inserted as annotations to threat diagrams. Using treatment diagrams is advantageous as it allows us to distinguish between threats, threat scenarios, vulnerabilities and risks in the search for adequate treatment options. At the same time the treatment diagrams serve as a

means for documenting the results. Examples of treatment diagrams with treatments for risks before and after changes are given in Section 14.5.

Treatments that apply to threats, vulnerabilities or threat scenarios are treatments to risks indirectly, as they provide treatment to the risks that are caused by such elements. In treatment diagrams, this indirect effect of treatment to risks is documented by relating treatments to the elements to which they apply. For the purpose of providing an overview of the identified treatments, we can use CORAS treatment overview diagrams. These are risk diagrams annotated with treatments that are related directly to the risk that they mitigate. An example of a treatment overview diagram is given in Figure 58.

14 D – Report on ATM Case Study with CORAS

This appendix gives the full report on the risk assessment conducted as part of the Air Traffic Management (ATM) case study of SecureChange. The risk assessment was conducted according to the risk analysis process described in Section 7, and instantiated on the CORAS approach. The report is structured according to the five activities of the risk assessment process as depicted in Figure 3 and is documenting the outcome of these activities.

14.1 Context Establishment

This section documents the context establishment of the ATM risk assessment, and includes the target description before and after the changes, a description of the changes themselves, a high-level risk analysis, and the documentation of the risk evaluation criteria.

14.1.1 Analysis Background and Motivation

The ATM domain involves an aggregation of services provided by ground-based Air Traffic Controllers (ATCOs). One of the main critical responsibilities of ATCOs is to maintain horizontal and vertical separation among aircrafts and between aircrafts and possible obstacles. They must ensure an orderly and expeditious air traffic flow by issuing instructions and information to aircrafts, and by providing flight context information to pilots, such as routes to waypoints and weather conditions.

An important characteristic of the ATM domain of today is that there are limited interactions with the external world, and therefore also limited security problems in relation to information flow to and from the environment. A further characteristic is that humans are at the center of the decision and work processes, with limited role of automated decision support systems and tools. Current safety problems are therefore mainly due to human errors, air-ground communication problems and degradation of technical and human services, all possibly combined with adverse atmospheric conditions that could raise safety problems.

However, the planned and ongoing introduction of new information systems and decision support systems, as well as the reorganization of ATM services, raise new security issues and security concerns with immediate impact on safety issues. The widespread deployment of innovative information system technologies at every stage of the air transport value chain, from ticket purchase to flight management, raises major security concerns with regards to the vulnerabilities of these new information technologies. Traditionally, security aspects have not been fully and thoroughly taken into account in the development and deployment of components of the ATM, but in a few years this will become a central problem to be solved.

The overall objective of the risk analysis reported in this appendix is to understand, document and assess security risks of ATM with particular focus on the arrival management process. Arrival management is a process that involves several actors and roles in the planning and organization of the air traffic flow. The ATCOs with their tasks and responsibilities are at the core of the overall arrival management process. In order to identify relevant security issues and to understand the security risks, it is therefore necessary to properly understand the ATCO roles and work processes, the interactions that involve the ATCOs in the arrival management tasks, and the information that is passed between ATCOs, between ATCOs and the aircrafts and between ATCOs of different Air Traffic System (ATS) units.

More specifically, the chosen target of analysis is an Area Control Center (ACC) and the ATCOs. The ACC is a ground-based control center with responsibility of managing the traffic of a given airspace. The actual traffic management is conducted from the operation room (OPS room), which is the operational environment of the ATCOs. The ATCOs have different roles, some of which have their own Controller Working Position (CWP). The CWP makes a range of tools available to the ATCOs for surveillance, communication and planning. The focus of the analysis is the arrival management process with the involved activities, tasks, roles, components and interactions.

Included in the target of analysis are the organizational level changes that are implied by the introduction of the arrival manager (AMAN). The AMAN is a queue management tool that aids the arrival management tasks of ATCOs. In particular, the AMAN is a sequencing tool helping to manage and better organize the air traffic flow in the approach phase. The AMAN calculates sequences on the basis of predicted times of arrival at a sequencing point, which is a navigation point usually five to ten minutes before landing. The aim of the AMAN is to achieve a more precisely defined flight profile and traffic flow management, in principle from off-block to arrival at the destination airport, in order to minimize the airport delay leading to better efficiency in terms of flight management, fuel consumption, time and runway capacity utilization.

The timeframe of introducing the AMAN tool is from today and until 2020. The aim of this analysis is on the one hand to understand, assess and document the current risk picture before the introduction of the AMAN. On the other hand, the aim is to try and foresee risks that may emerge as a consequence of introducing the AMAN and to identify means for risk treatment in order to ensure an acceptable risk level both before and after the implementation of the changes.

14.1.2 Target Description

The target of the analysis is a specific Area Control Center and the activities of the Air Traffic Controllers in the arrival management process. The party of the analysis, i.e. the stakeholder with respect to which the analysis is conducted, is the ATM service provider. In the following we first document the target of analysis and the assets before the organization level changes of the AMAN introduction are taken into account. Thereafter we describe the change requirements before we document the target description where the AMAN introduction is reflected.



14.1.2.1 Target Description before AMAN Introduction

The documentation of the target of analysis is divided into different parts. We first use UML class diagrams to provide a conceptual overview of the target of analysis. Second, we use UML structured classifiers to document the internal structure of the roles and components, as well as the communication links between them. Third and finally, we use UML interactions to describe the relevant activities.

Conceptual Overview

The UML class diagram of Figure 65 gives a conceptual overview of the various roles, components and networks involved in the ACC.

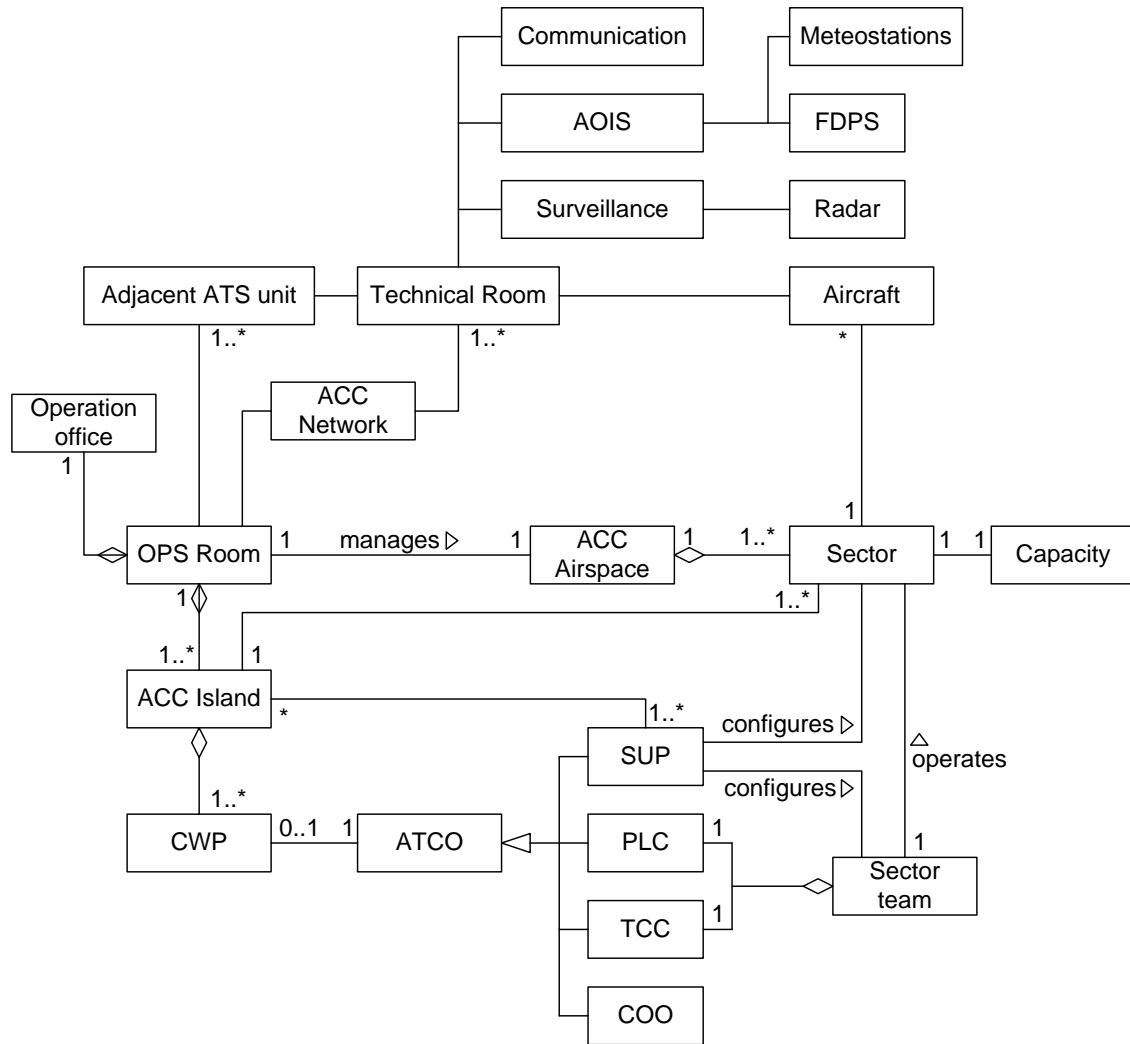


Figure 65 Conceptual overview of ACC before changes

The Operational Room (OPS Room) is the operational environment of the ACC. The OPS Room is divided into dedicated operative zones or ACC Islands, where each island consists of a number of Controller Working Positions (CWPs). Each CWP is operated by exactly one Air Traffic Controller (ATCO).

As seen from the class diagram, the OPS Room is responsible for managing the ACC Airspace, which is the segment of the airspace that is allocated to the ACC. The ACC Airspace is in turn divided into a number of Sectors. Each ACC Island is at any given point in time responsible for managing the traffic of a number of Sectors. The number of aircrafts in a Sector constantly varies, but it can never exceed the Capacity, which is the maximum number of aircrafts that can be managed by the sector.

The ATCOs have one of four different roles, namely Supervisor (SUP), Planning Controller (PLC), Tactical Controller (TCC) and Coordinator (COO). The PCL and the TCC forms a Sector team and are together responsible for operating and managing the traffic of their sector. The TCC is in charge of all air-ground communication. He monitors the aircrafts in the sector and provides pilots with instructions/clearances on aspects such as speed, altitude and routing to maintain a safe separation with other aircraft flying in the sector, and with other possible obstacles that are present. He also gives pilots weather and air traffic information. When the aircraft approaches the sector boundary, he passes it off to the TCC of the adjacent sector (not always belonging to the same ACC). The PLC assists the TCC, coordinating entry and exit flight level and entry and exit flight point with adjacent sectors in order to ensure a smooth air traffic flow. He also monitors the traffic within the sectors and in most of cases updates the air traffic control system with the instructions given by the TCC.

Groups of neighboring sectors are coordinated by a SUP, who is in charge of managing the sector configurations under his responsibility according overall traffic forecast. The SUP can split and merge sectors depending on the traffic. The SUP is moreover responsible for the formation of the sector teams. The COO is involved only in islands where there is a Terminal Area (TMA). The COO does not work on a CWP, but moves between sector teams to survey the arrival management process and coordinating the tasks between sectors.

The UML class diagram of Figure 66 gives a conceptual overview of the ATCO roles. The roles of SUP, PLC and TCC have dedicated CWPs with functionalities and interfaces adapted to the needs and tasks of each ATCO.

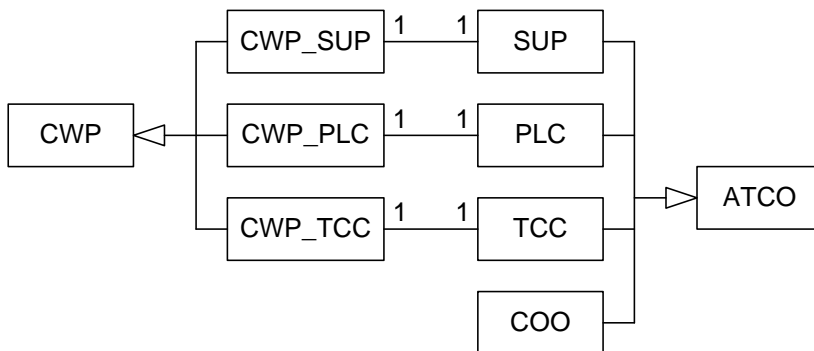


Figure 66 Conceptual overview of ATCO roles before changes

From the class diagram of Figure 65 we furthermore see that the OPS Room is linked to the Technical Room, and thereby also to the aircrafts, via the ACC network. The technical room provides phone lines and radio frequency antennas for communication, and radar antennas for surveillance. The technical room furthermore provides an Aeronautical Operational Information System (AOIS) that includes a Flight Data Processing System (FDPS) and weather information from Meteostations.

The UML class diagram of Figure 67 shows that the ACC Network is divided into two partitions that correspond to the communication and surveillance of the technical room.

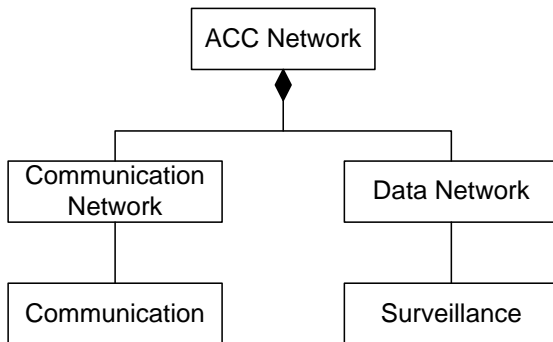


Figure 67 Conceptual overview of ACC Network before changes

Components and Communication

We use UML structured classifiers to document the internal structure of components and the communication lines between components. The diagrams are hierarchically organized to show the structures at various levels of detail.

The UML structured classifier of Figure 68 shows the structure and communication links of the ATM components.

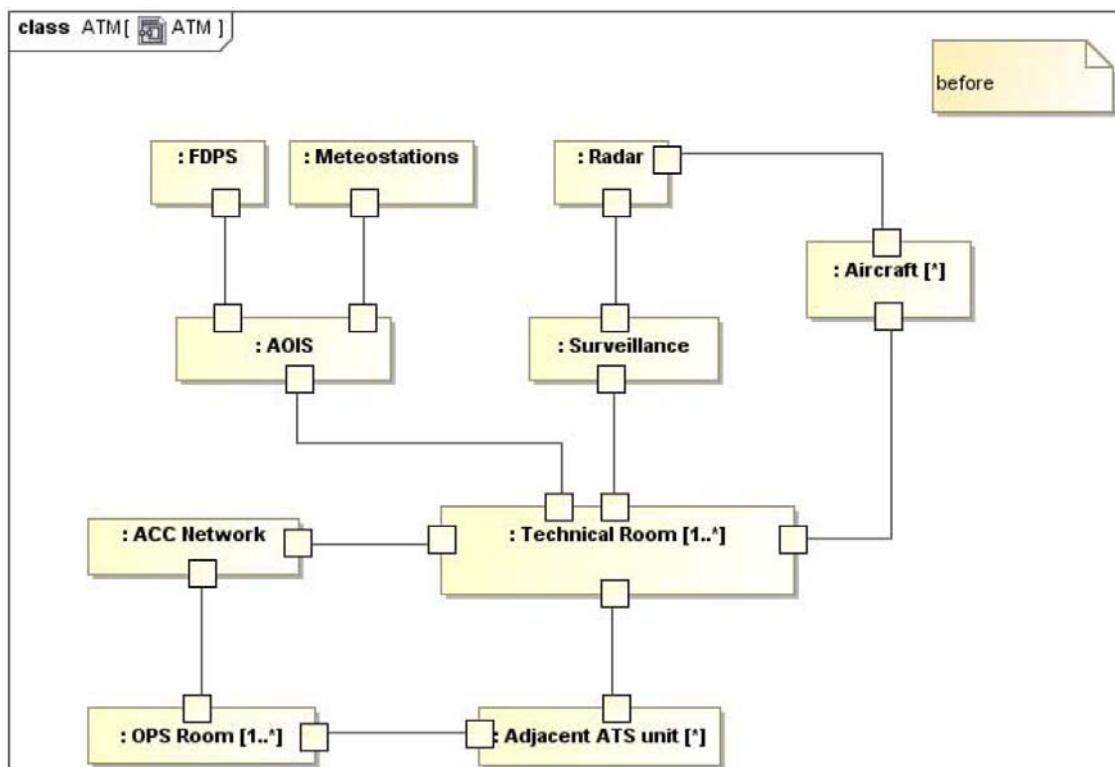


Figure 68 Structure of ATM components before changes

The structured classifier of Figure 69 shows the OPS Room as consisting of a number of ACC islands that are connected to the ACC network. The OPS Room furthermore

has a number of SUPs that communicate with the ACC islands, and that also are connected to the ACC Network via the controller working position CWP_SUP.

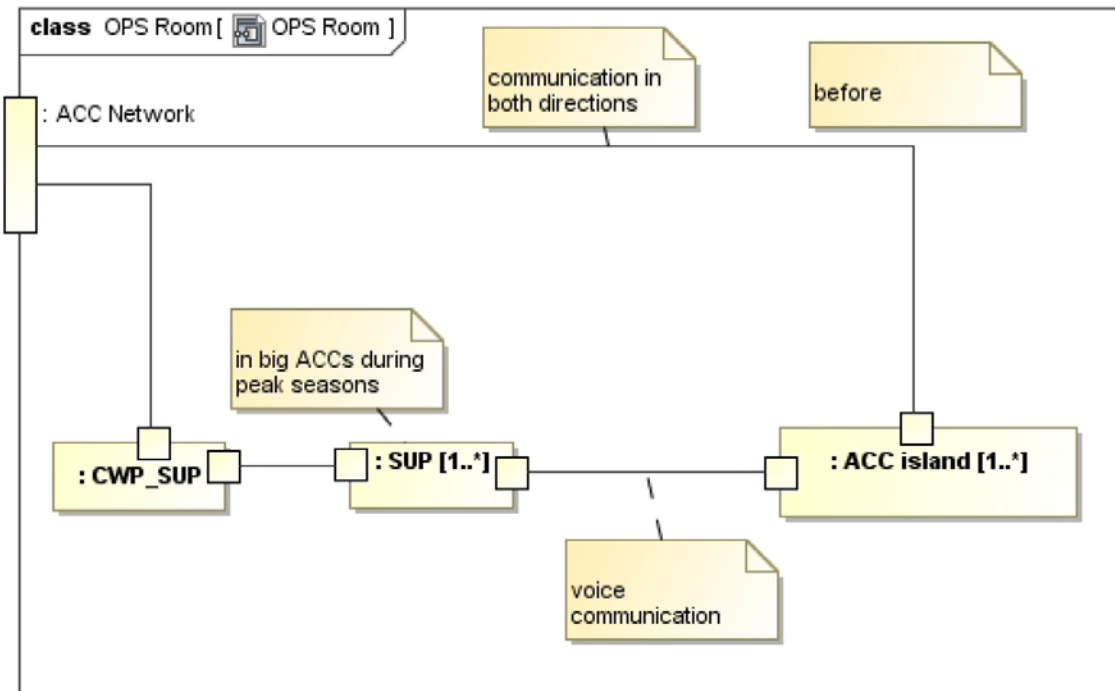


Figure 69 Internal structure of OPS Room before changes

The structured classifier of Figure 70 shows the internal structure of the ACC islands. There are a number of sector teams that are connected to the ACC network and that also communicate with the SUPs. The COO is also a part of the ACC island, and communicates both with the sector team and with the SUP.

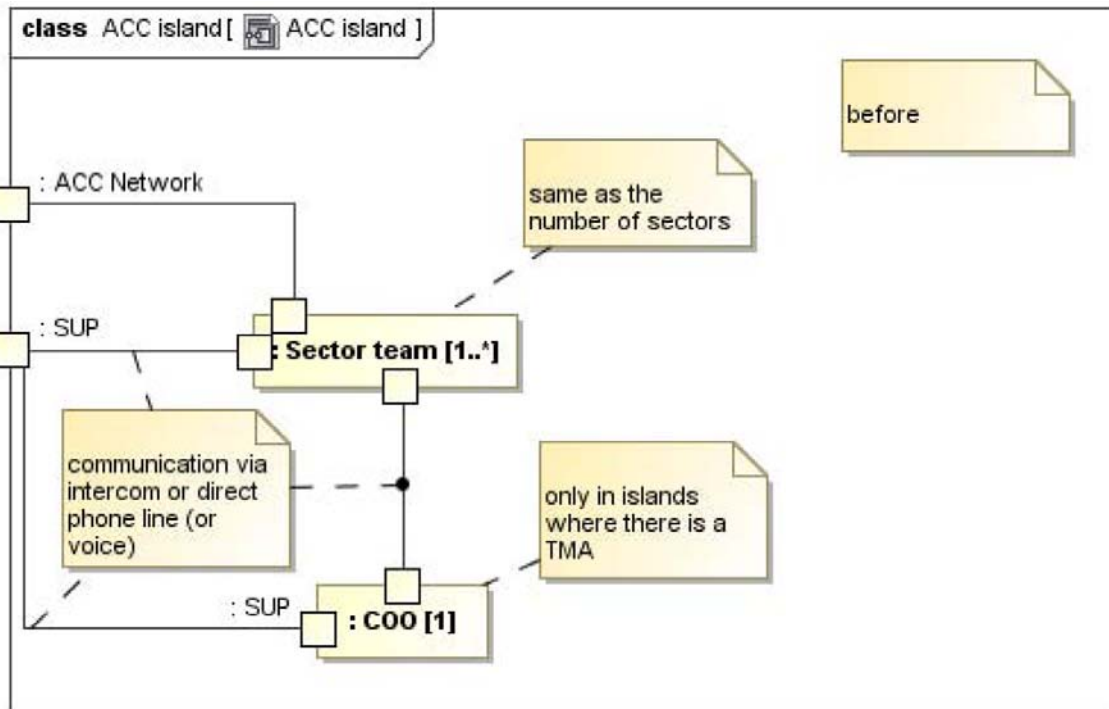


Figure 70 Internal structure of ACC island before changes

The structured classifier of Figure 71 finally shows the structure of the sector team. The sector team consists of a PLC and a TCC that can communicate directly between them by voice. These ATCOs are furthermore operating their own CWP which connects them to the ACC network. They communicate with the COO and the SUPs by voice communication, but the communication and information flows between the PLC, the TCC and the SUPs are also passed between them via the ACC network and the respective CWPs.

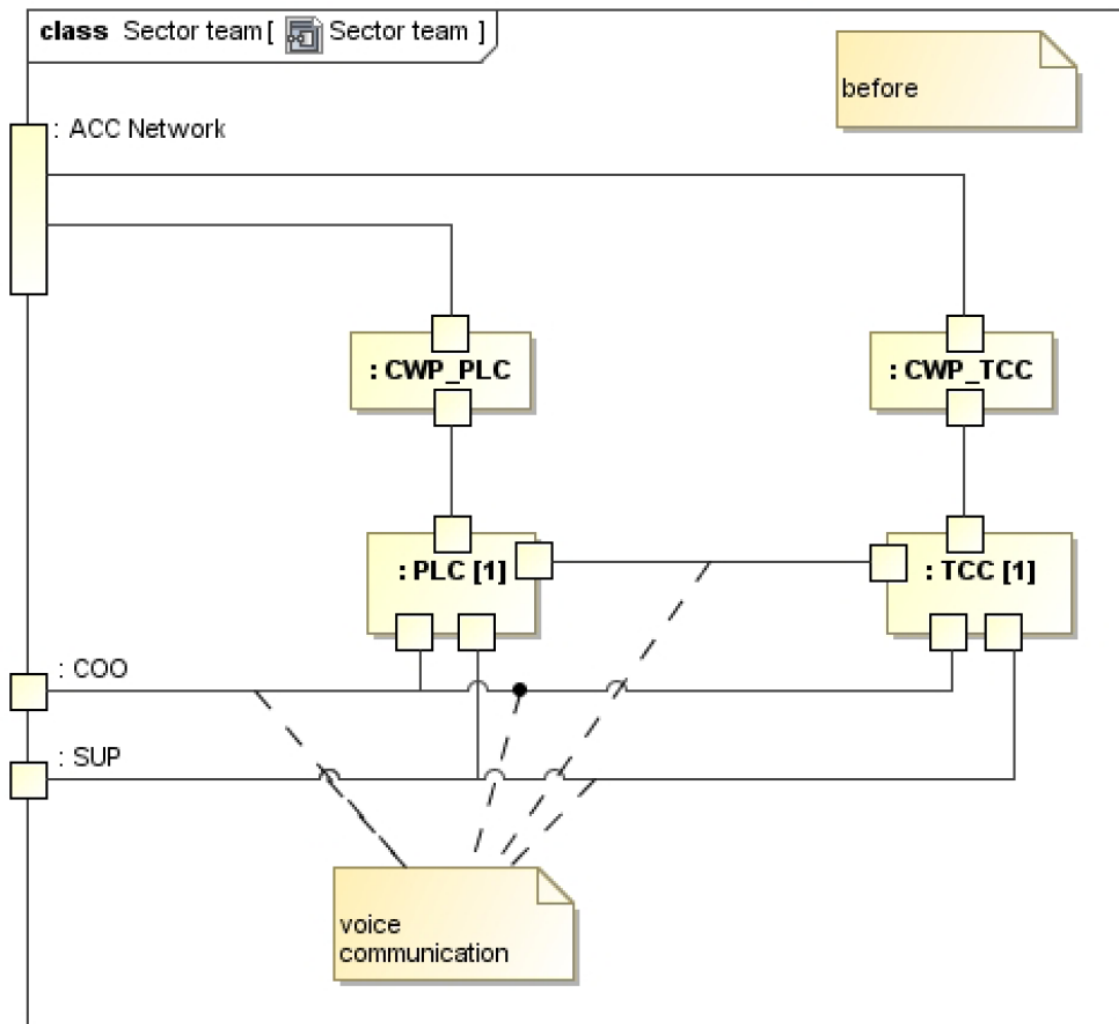


Figure 71 Internal structure of sector team before changes

Interactions

In order to properly understand the target of analysis with the focus on the arrival management process, it is important to properly understand the various activities of the arrival management process. The UML interaction overview diagram of Figure 72 gives a high-level overview of the various tasks without explicitly showing the involved roles. As the naming of the tasks indicates, they can be structured into five main tasks, summarized as follows:

- **Task 1:** Controlling the aircraft (A/C) in the sector
- **Task 2:** A/C data analysis for starting the sequence creation
- **Task 3:** Sequence finalization
- **Task 4:** Clearances to the A/C for building the planned sequence
- **Task 5:** Progressive transfer of the whole sequence to the adjacent sector

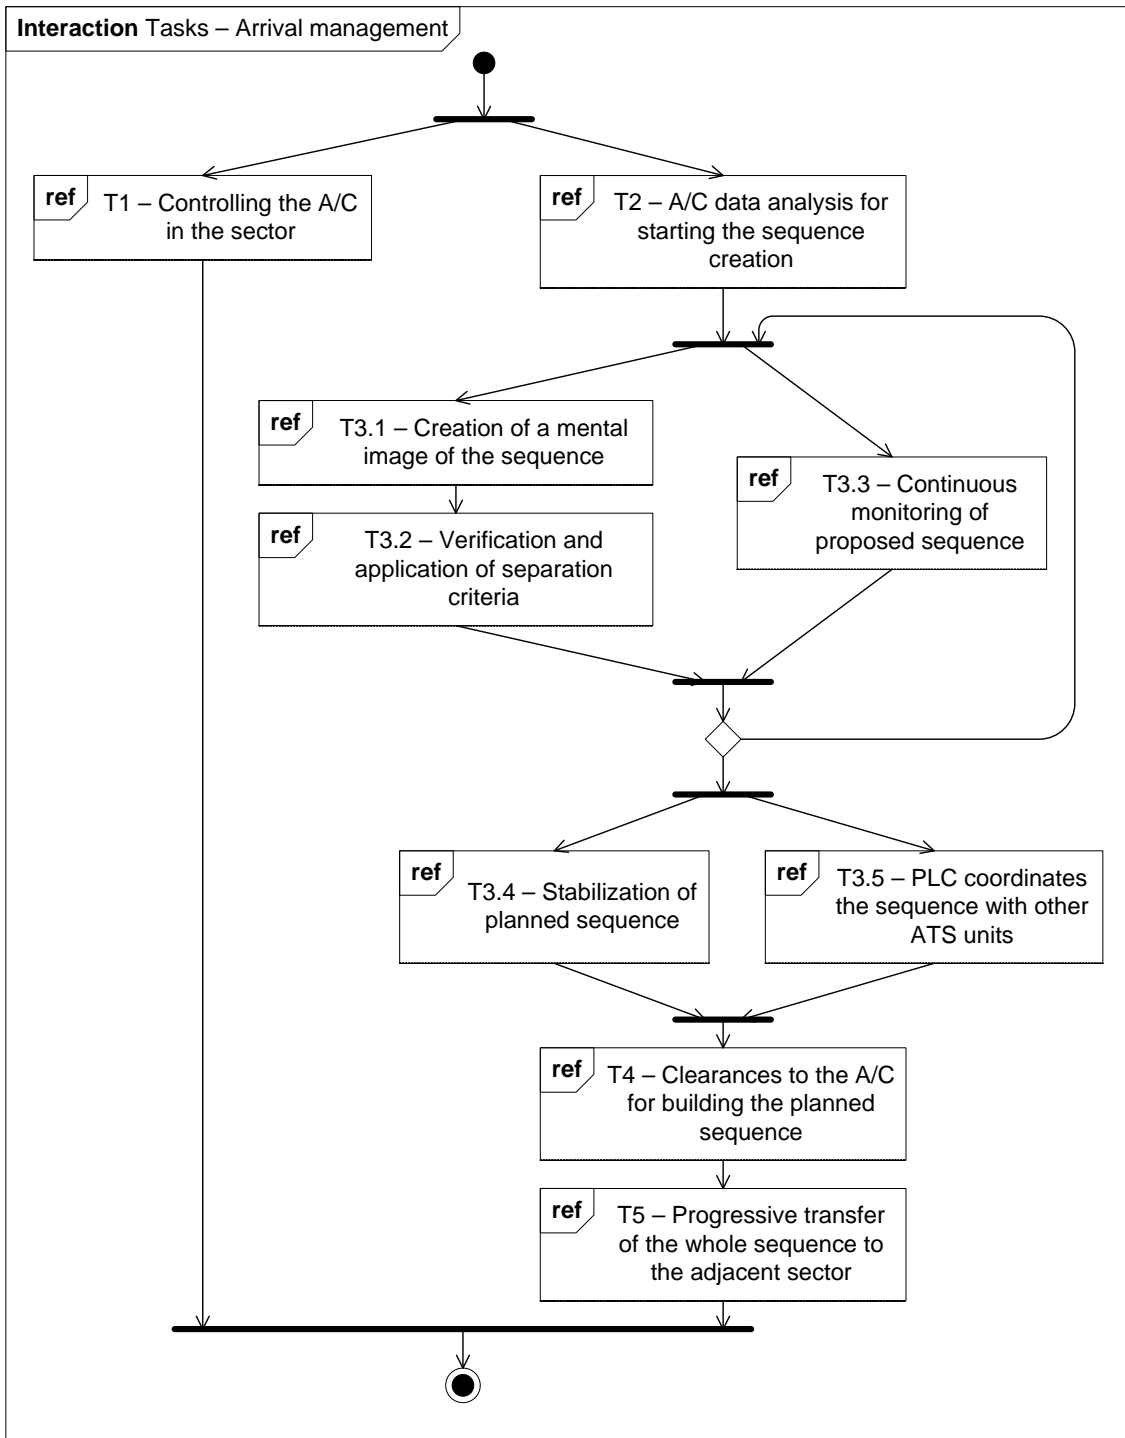


Figure 72 Overview of arrival management tasks before changes

The arrival management is conducted according to the ACC internal procedures, but the exact way of conducting the tasks may vary over time and may vary between different sector teams. External factors such as traffic intensity and meteorological conditions may affect how the tasks are conducted, and the various sector teams and

ATCOs may have their own preferences and habits for how to conduct the tasks. The description of the arrival management tasks documented here shows typical examples of how the tasks are usually conducted. For the purpose of the risk analysis, the important thing is to document the information flows and the involved roles and components in the various tasks.

We use UML sequence diagrams to document the roles and components that are involved in the various tasks. For some interactions it suffices to know and document the involved entities, and for those cases the messages that are passed between them are hidden. For other interactions we specify the information that is passed between roles and components by specifying the messages that are sent and received.

The sequence diagram of Figure 73 shows the roles and components involved in task T1, Controlling the A/C in the sector. This is an ongoing task that is conducted in several iterations and involves several sub-tasks that are conducted by the TCC and the PLC in parallel. The tasks of these two sector team members are basically the same, but the difference is that the PLC works on a wider scale in terms of time and space. The role of the PLC is to plan ahead and to assist the TCC who operates on a more narrow scale. The ATCOs operate their CWP's for conducting these tasks. In the diagram we have not specified the interactions between the ATCOs and their respective CWP's, since that level of details is not necessary for our purposes.

Whereas task T1 is a continuous task for controlling and monitoring the aircrafts in the sector, the remaining tasks concern building flight sequences. The sequences define the sequencing of the aircraft in the traffic flow and must always comply with the separation criteria for maintaining sufficient horizontal and vertical separation between aircrafts.

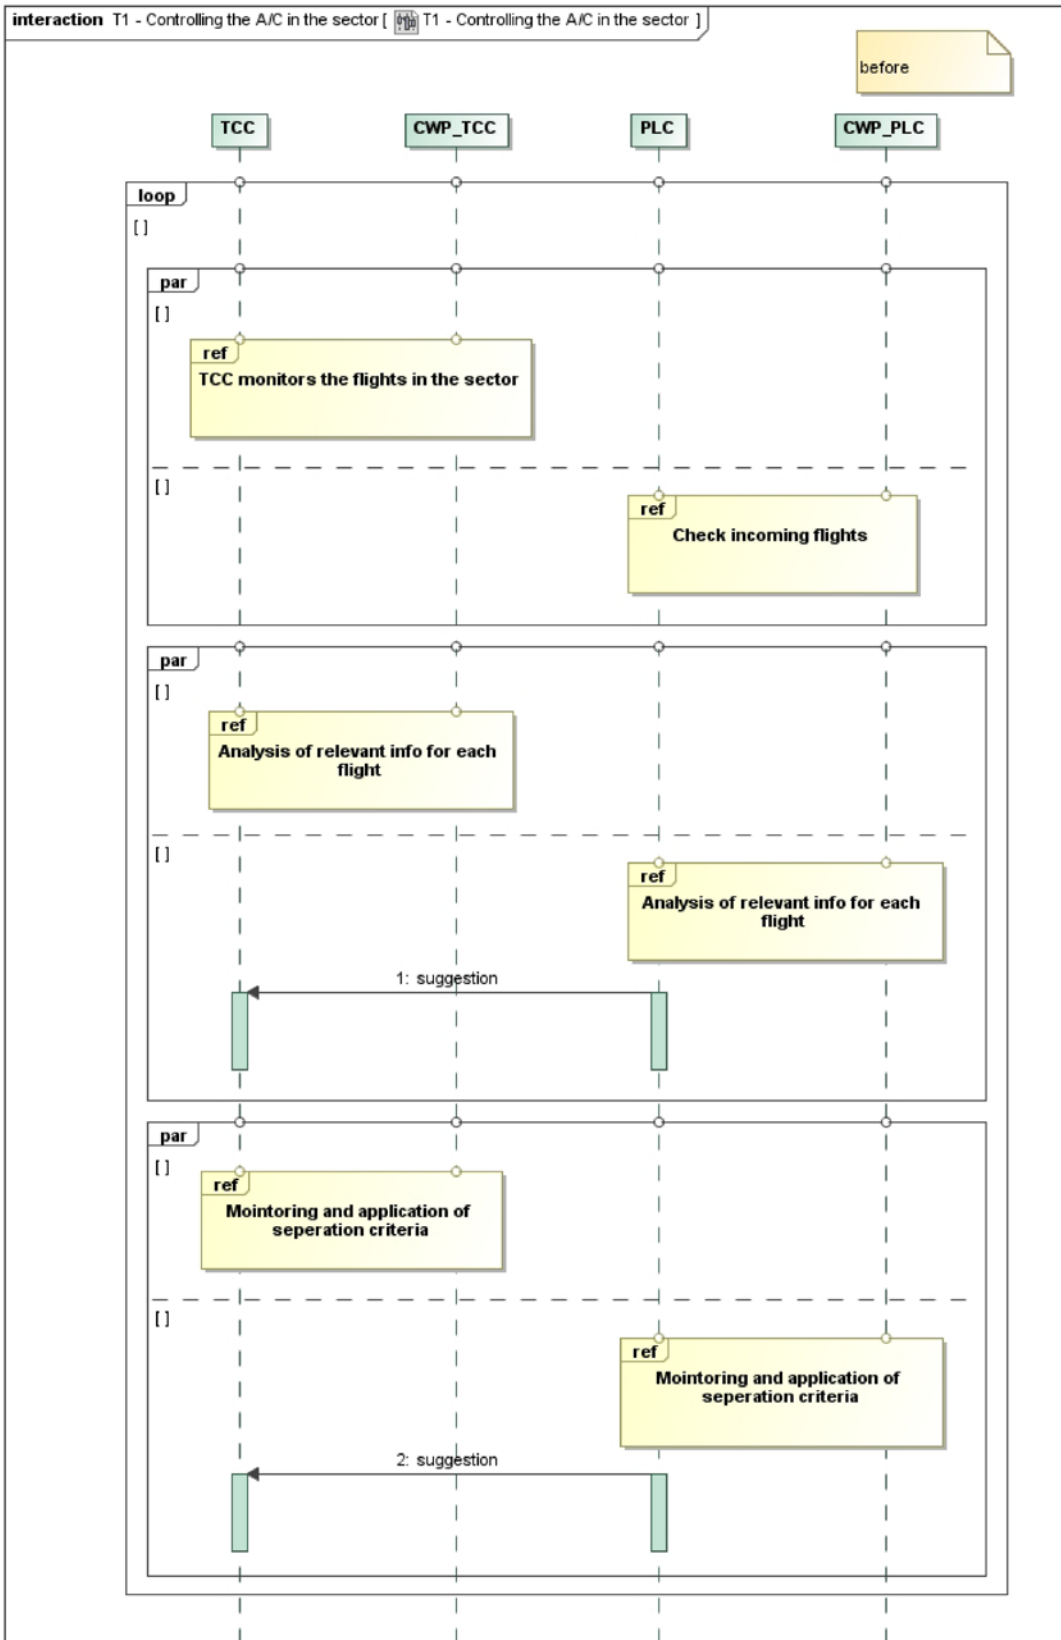


Figure 73 Task T1 before changes

The sequence diagram of Figure 74 shows that task T2, A/C data analysis for starting the sequence creation, is conducted by the TCC, mainly by operating the CWP. The TCC also relies on the flight data processing systems (FDPS) for providing flight information. The first sub-task is an A/C classification, the second a comparison of the A/C position with the point of top of climb (TOC), and the third task is a comparison of A/C speed and flight level (FL).

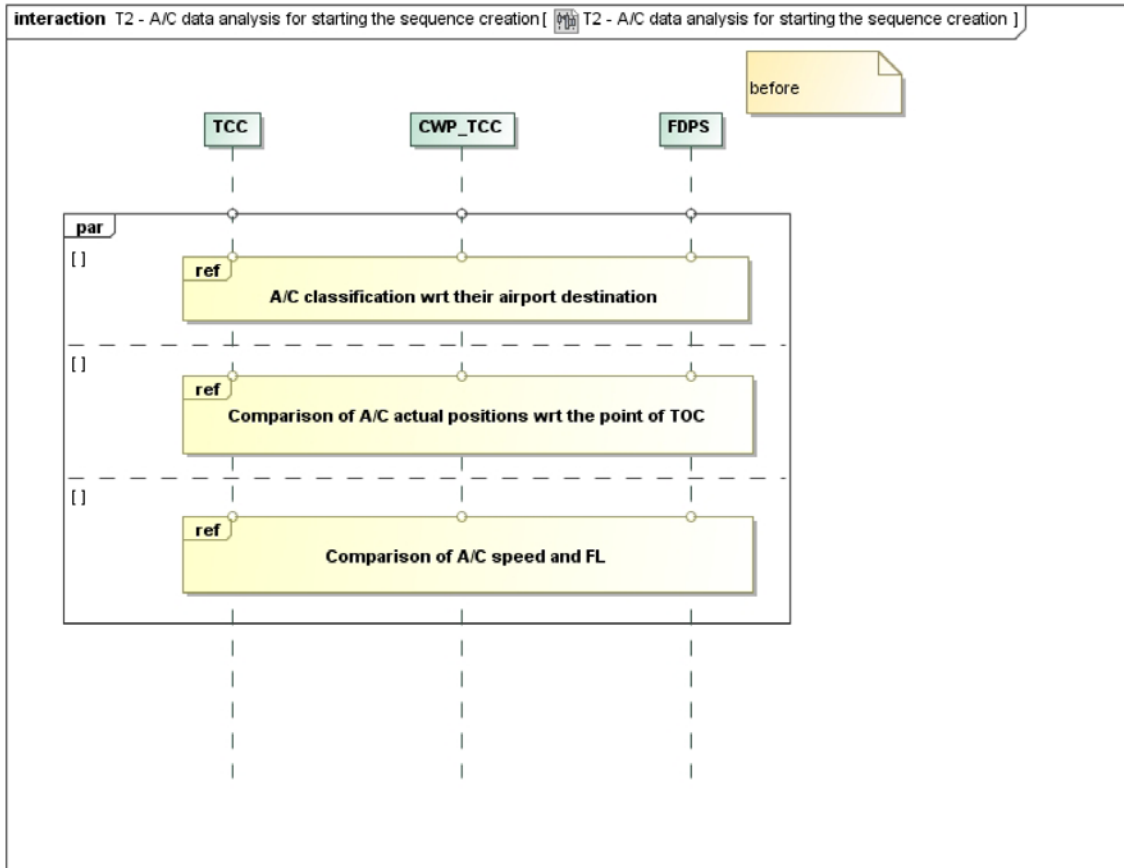


Figure 74 Task T2 before changes

Task T3, Sequence finalization, is where the sequences are planned and built while being monitored, before they are stabilized and continuously passed to the A/Cs and to the adjacent sectors. As depicted by the interaction overview of Figure 72, the third task consists of five separate sub-tasks.

Task T3.1 is the initial creation of a mental image of the sequence that is conducted by the TCC, and needs no further detailed specification. Task T3.2, the verification and application of the separation criteria is also conducted by the TCC, the further specification of which is not necessary as it is a mental task. Task T3.3, the continuous monitoring of proposed sequence, is a more compound task. The UML interaction overview diagram of Figure 75 shows the further decomposition of T3.3 into five subtasks that may be conducted in several iterations. The continuous monitoring of the sequence is the responsibility of the TCC, but the TCC is supported by the PLC and the COO, and the TCC furthermore uses flight data and surveillance data provided via the CWP.

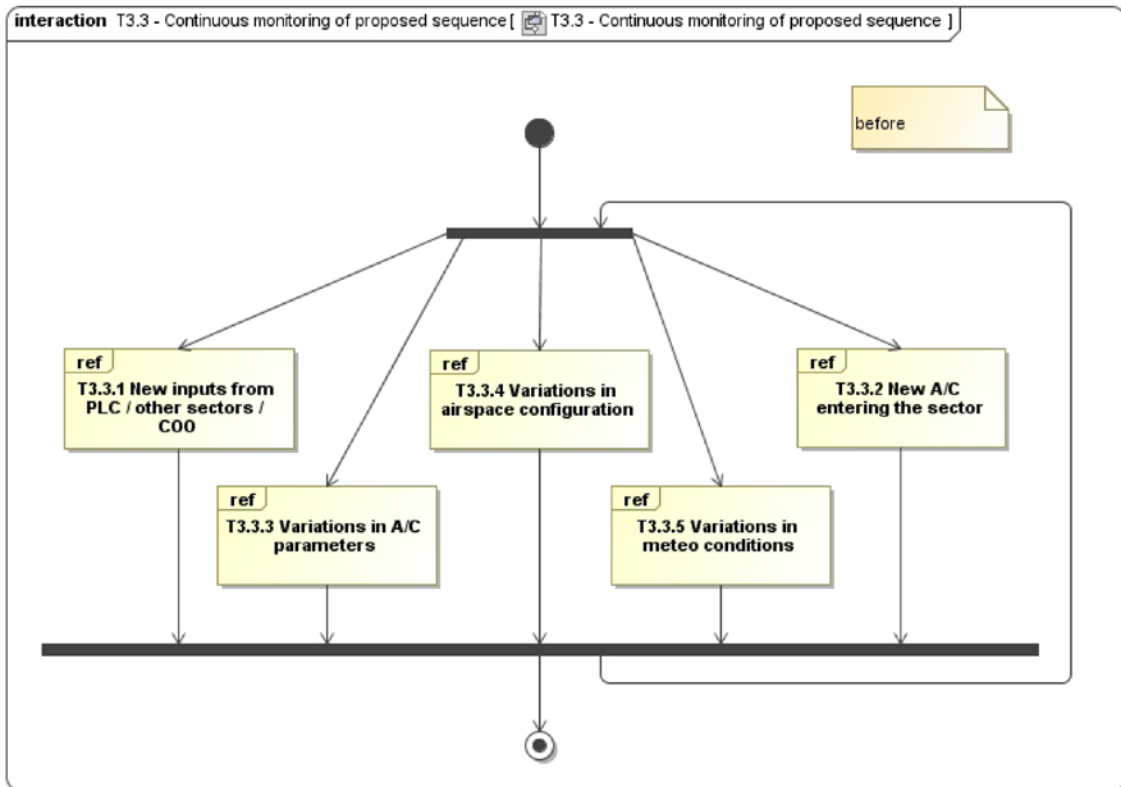


Figure 75 Overview of task T3.3 before changes

The sequence diagram of Figure 76 shows Task T3.3.1 and the input that is fed to the TCC for supporting the sequence monitoring. The PLC and the COO conducts sequence monitoring in parallel, and may provide input to the TCC or give specific requests for changing the sequence if they find it necessary.

The sequence diagram of Figure 77 shows Task T3.3.2 and the information that is provided to the TCC when a new aircraft is entering the sector.

The sequence diagram of Figure 78 shows task T3.3.3 and the flight data that is constantly provided to the TCC via the CWP. The flight data is monitored for identifying possible variations in the aircraft parameters that may be relevant for the sequencing.

The sequence diagram of Figure 79 shows task T3.3.4 and the possible reconfigurations of the airspace that must be taken into account by the TCC. The SUP is responsible for reconfiguring the airspace in case this is required, and subsequently informing the TCC.

The sequence diagram of Figure 80 shows task T3.3.5 and the variations in meteo conditions that are fed to the TCC via the CWP.

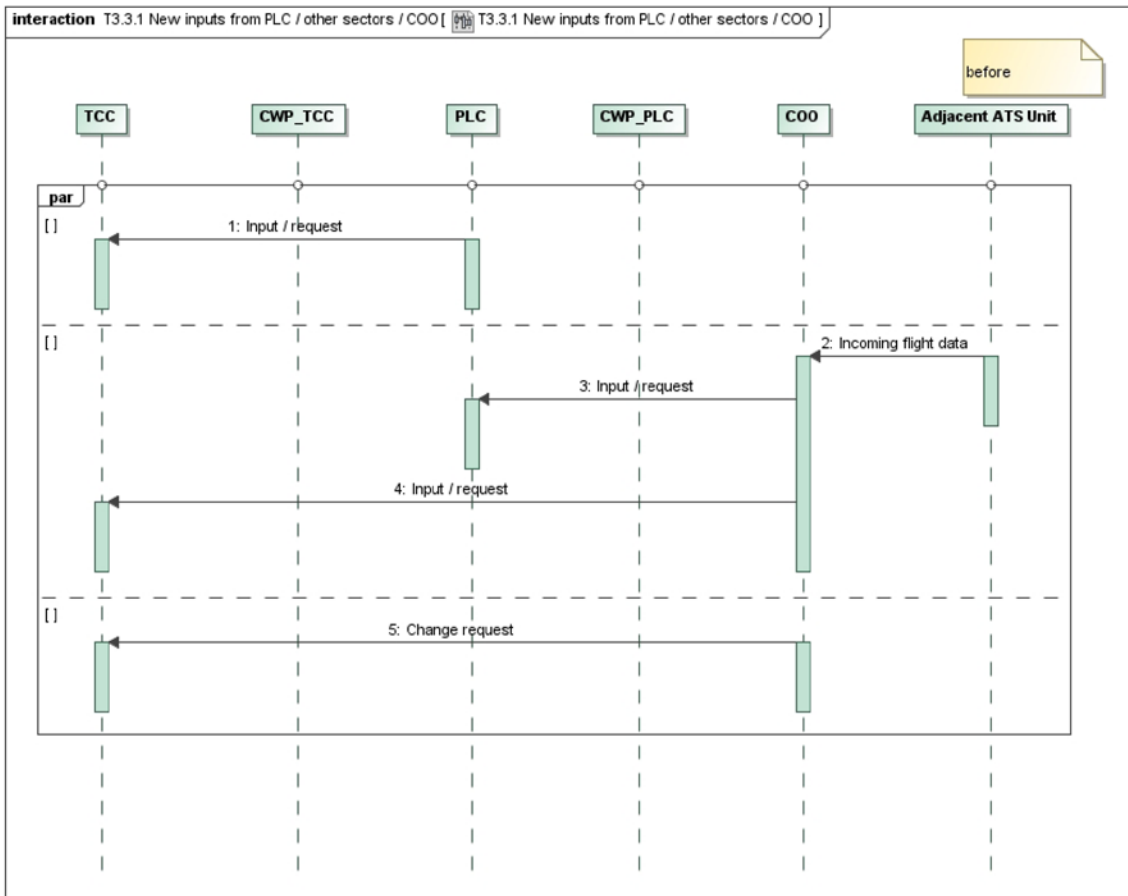


Figure 76 Task T3.3.1 before changes

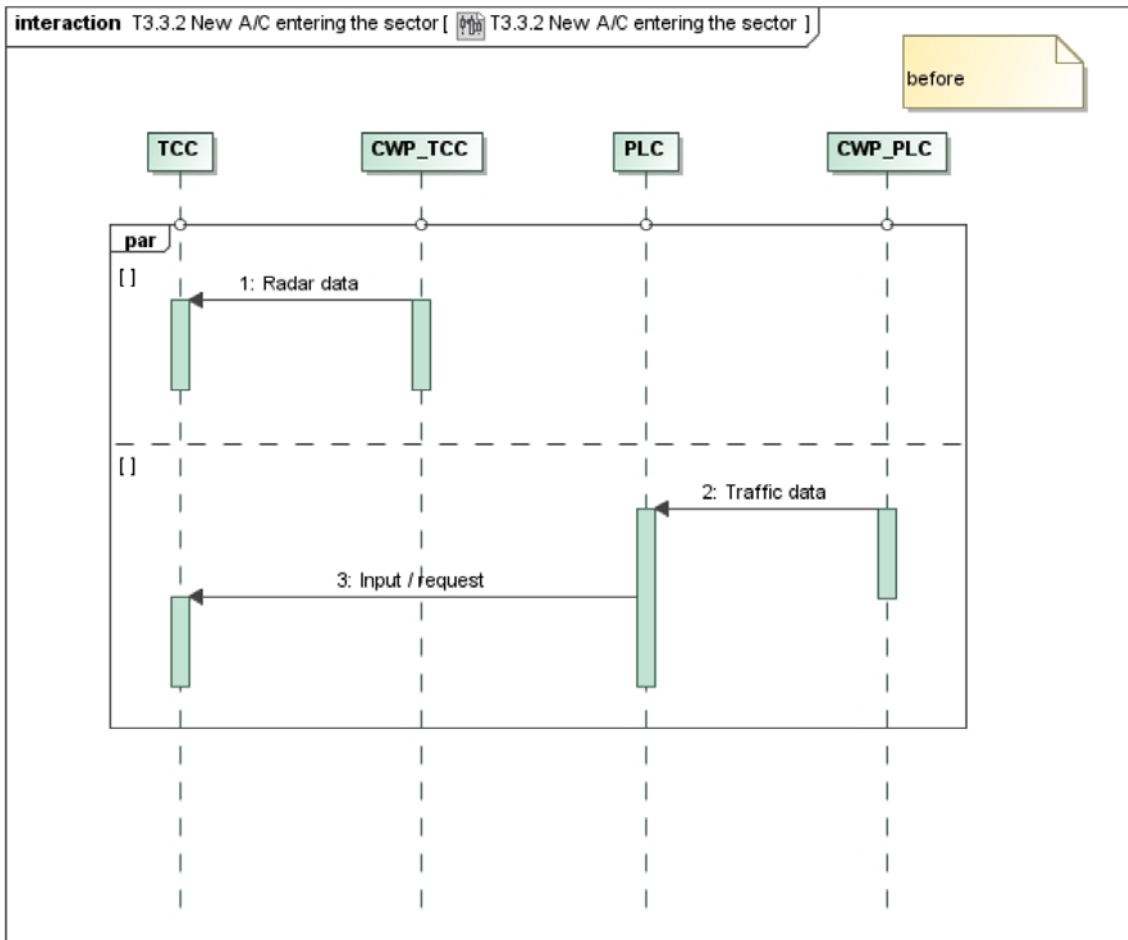


Figure 77 Task T3.3.2 before changes

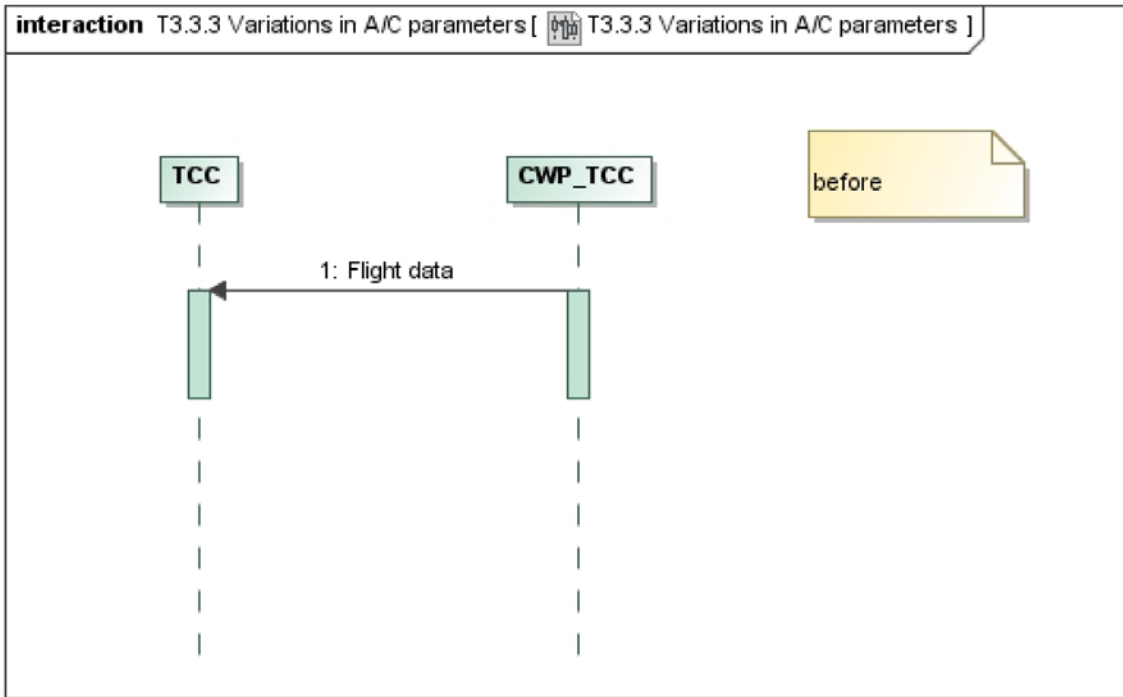


Figure 78 Task T3.3.3 before changes

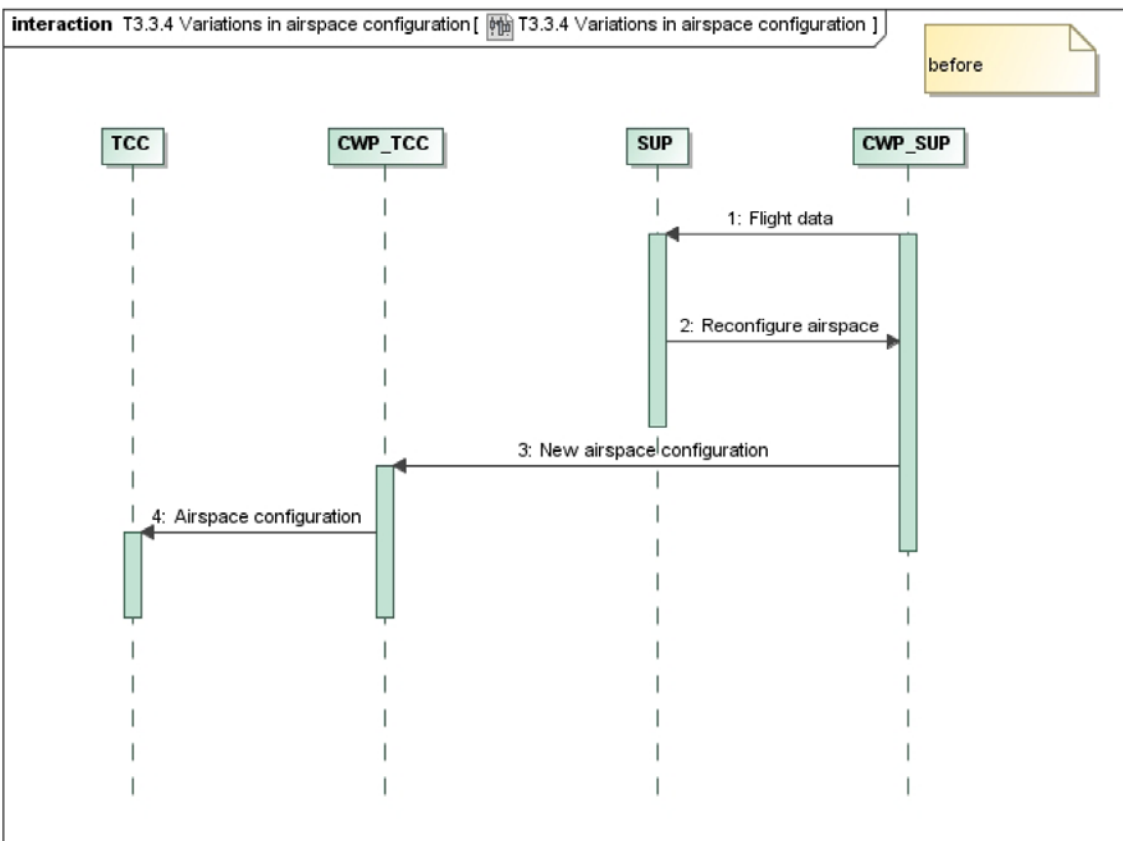


Figure 79 Task T3.3.4 before changes

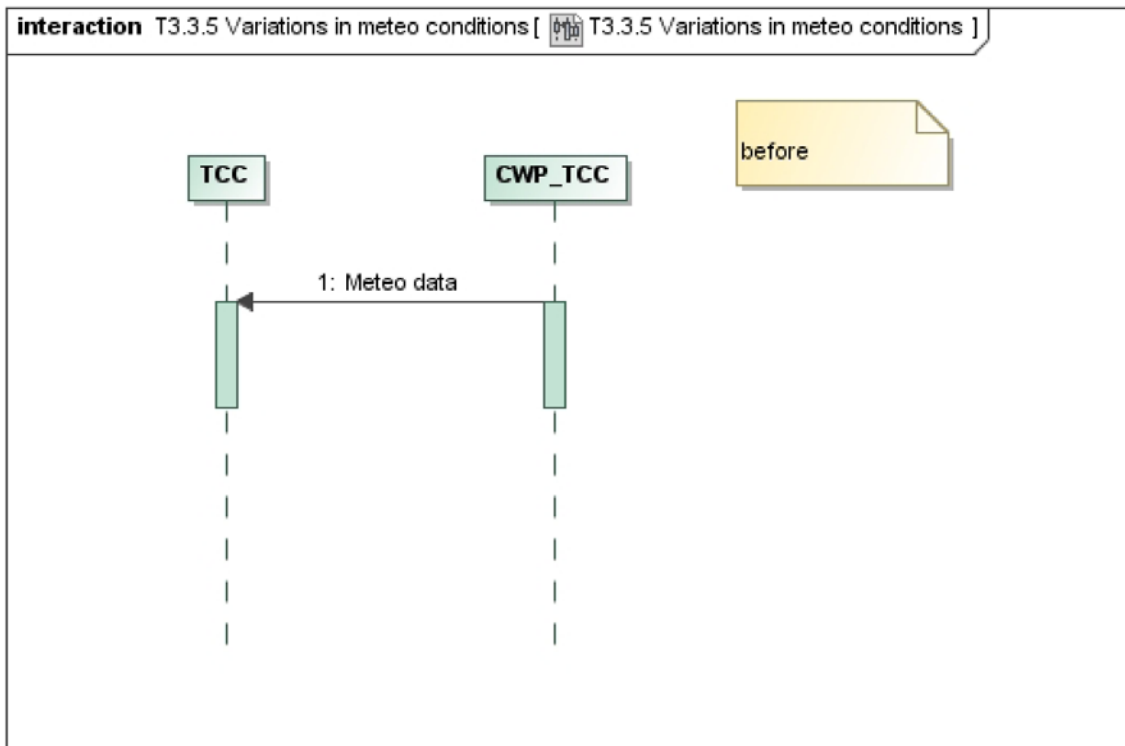


Figure 80 Task T3.3.5 before changes

The tasks T3.4, stabilization of planned sequence, and T3.5, PLC coordinates the sequence with other ATS units, are the final sub-tasks of the sequence finalization. These two tasks are the responsibility of the TCC and the PLC, respectively. The former is conducted by operating the CWP, and the latter by operating the CWP and communicating with the adjacent ATS units, the detailed specification of which is omitted in the documentation since it suffices with the high-level specification of Figure 71.

The detailed specifications of task T4, clearances to the A/C, and T5, progressive transfer for the sequence to the adjacent sector, are also omitted. Both tasks are the responsibility of the TCC and are conducted by communication with the pilots and with the other sectors.

14.1.2.2 Change Requirements

Having completed the documentation of the target of analysis before the changes, we turn to the change requirements. These are selected from the change requirements documented in SecureChange deliverable D.1.1.1 [30].

The changes we are addressing are changes in the operational processes of managing air traffic in Terminal Areas (TMAs). In particular, the introduction of the Arrival Manager (AMAN) affects the ATM system as a whole both at a process level and at an organizational level.

This risk analysis addresses the organizational level change. The introduction of the AMAN affects the controller working positions (CWPs), as well as the area control center (ACC) as a whole. The main foreseen changes from an operational and

organizational point of view are the automation of tasks (i.e. the usage of the AMAN for the computation of the arrival sequence) that currently are carried out by air traffic controllers (ATCOs).

The organizational level changes moreover require the redefinition of the ATCO role of the coordinator (COO), who will be responsible for monitoring and modifying the sequences generated by the AMAN, and for providing information and updates to the sectors. In order to highlight this redefinition of the ATCO role, the COO is renamed to Sequence Manger (SQM).

In the following we present the documentation of the target of analysis where the relevant changes have been taken into account.

14.1.2.3 Target Description after AMAN Introduction

The target of analysis after the changes is documented in the same way as before. We use UML class diagrams to give a conceptual overview, we use UML structured classifiers to document the internal structure of components, and we use UML interactions to document the relevant activities.

Conceptual Overview

The class diagram of Figure 81 gives a conceptual overview of the ACC after the introduction of the AMAN. The most important change is the introduction of the AMAN itself, which is connected to the ACC network and also linked to the CWPs. The diagram also shows that the ATCO role of the SQM has replaced the previous role of COO. This redefinition of the ATCO role also affects the CWPs; whereas the COO did not operate a CWP, the SQM do. The CWPs of the PLC and the TCC are also modified in order to accommodate to changes in their tasks and responsibilities.

The introduction of the automatic dependent surveillance-broadcast (ADS-B) is actually independent of the AMAN, but is taken into account in the analysis because ADS-B will be introduced during the same time frame and may affect security issues. ADS-B is a cooperative GPS-based surveillance technique for air traffic control where the aircrafts constantly broadcasts their position to the ground and to other aircrafts.

In order to highlight the changes, we use colors in the diagrams. The yellow color (light shading in black-and-white) indicates elements that are introduced, whereas the purple color (darker shading in black-and-white) indicates elements that are modified.

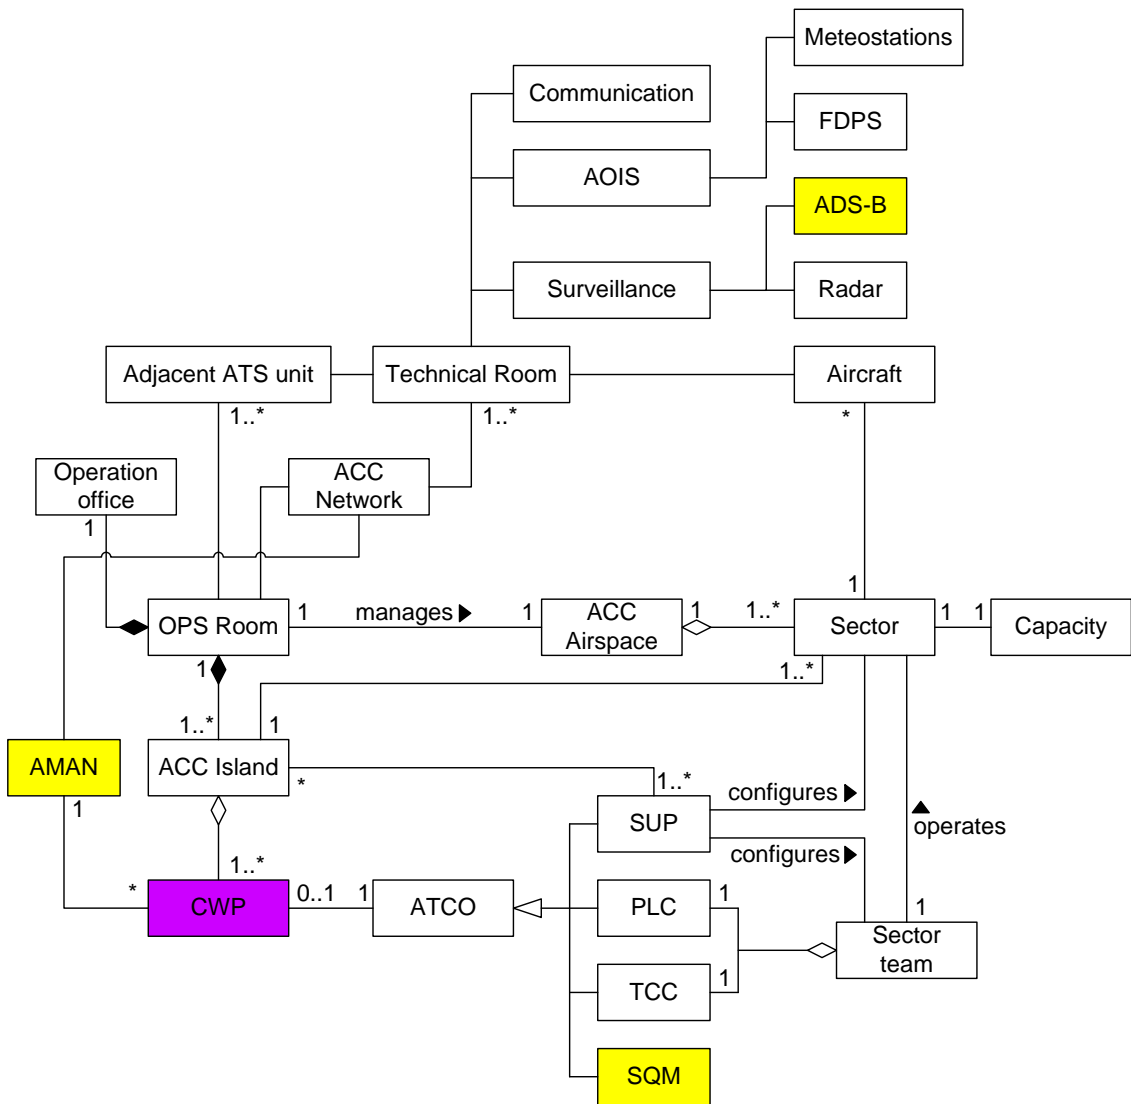


Figure 81 Conceptual overview of ACC after changes

The class diagram of Figure 82 gives an overview of the ATCO roles after the introduction of the AMAN. It shows that all the roles but the SUP uses the AMAN and that all the ATCOs operate a dedicated CWP.

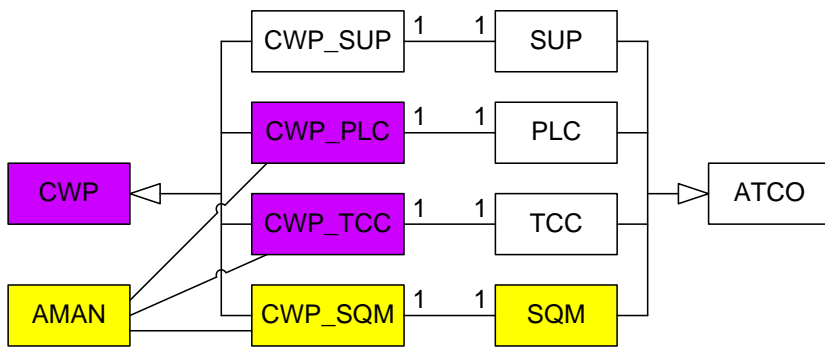


Figure 82 Conceptual overview of ATCO roles after changes

The overview of the ACC network as depicted in Figure 67 remains the same after the AMAN introduction.

Components and Communication

The UML structured classifier of Figure 83 shows the structure and communication links of the ATM components after the changes. At this level we only see the introduction of the ADS-B.

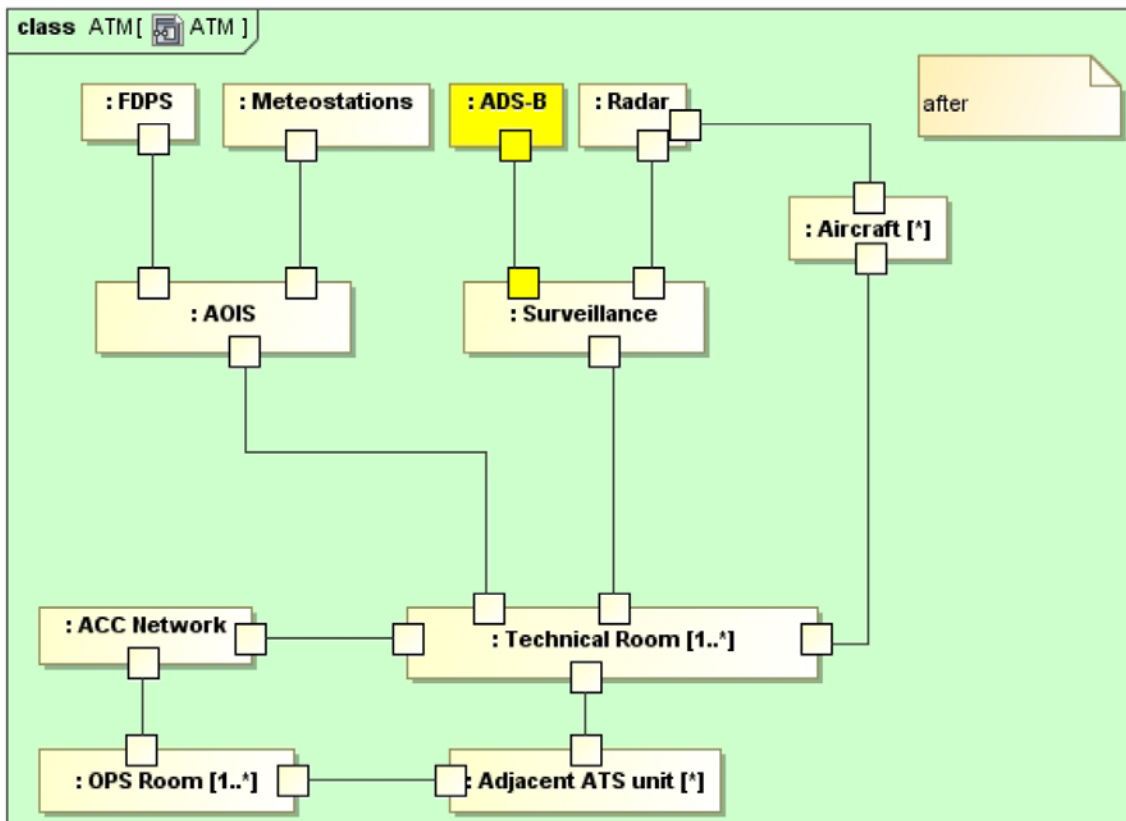


Figure 83 Structure of ATM components after changes

The diagram of Figure 84 shows the internal structure of the OPS room after the introduction of the AMAN. The AMAN is connected to the ACC network, and thereby also to the ACC islands and the CWPs.

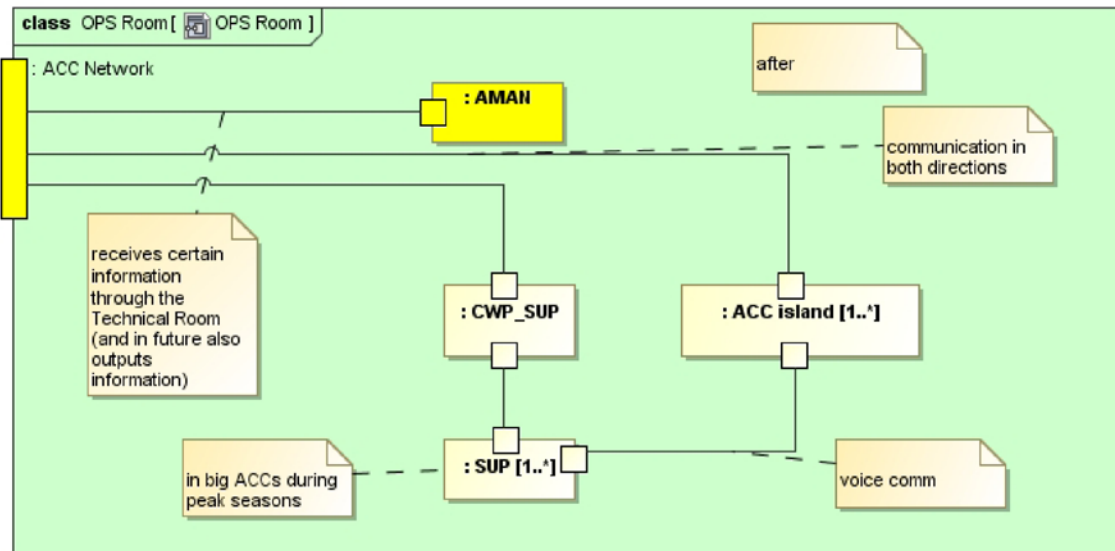


Figure 84 Internal structure of OPS Room after changes

The diagram of Figure 85 shows the internal structure of the ACC island after the changes. The SQM (previously COO) is now operating a CWP.

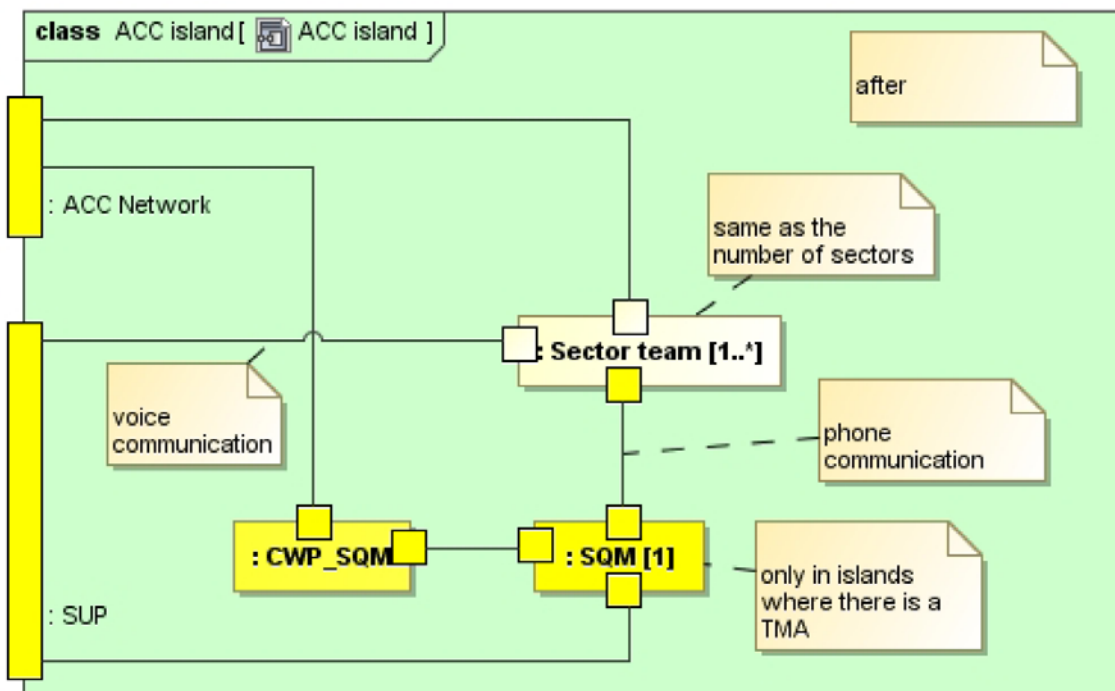


Figure 85 Internal structure of ACC island after changes

The diagram of Figure 86 shows the internal structure of the sector team after the changes. The internal structure is the same as before, only that now the PLC and the TCC communicate with the SQM by phone, whereas the communication with the COO previously was by voice.

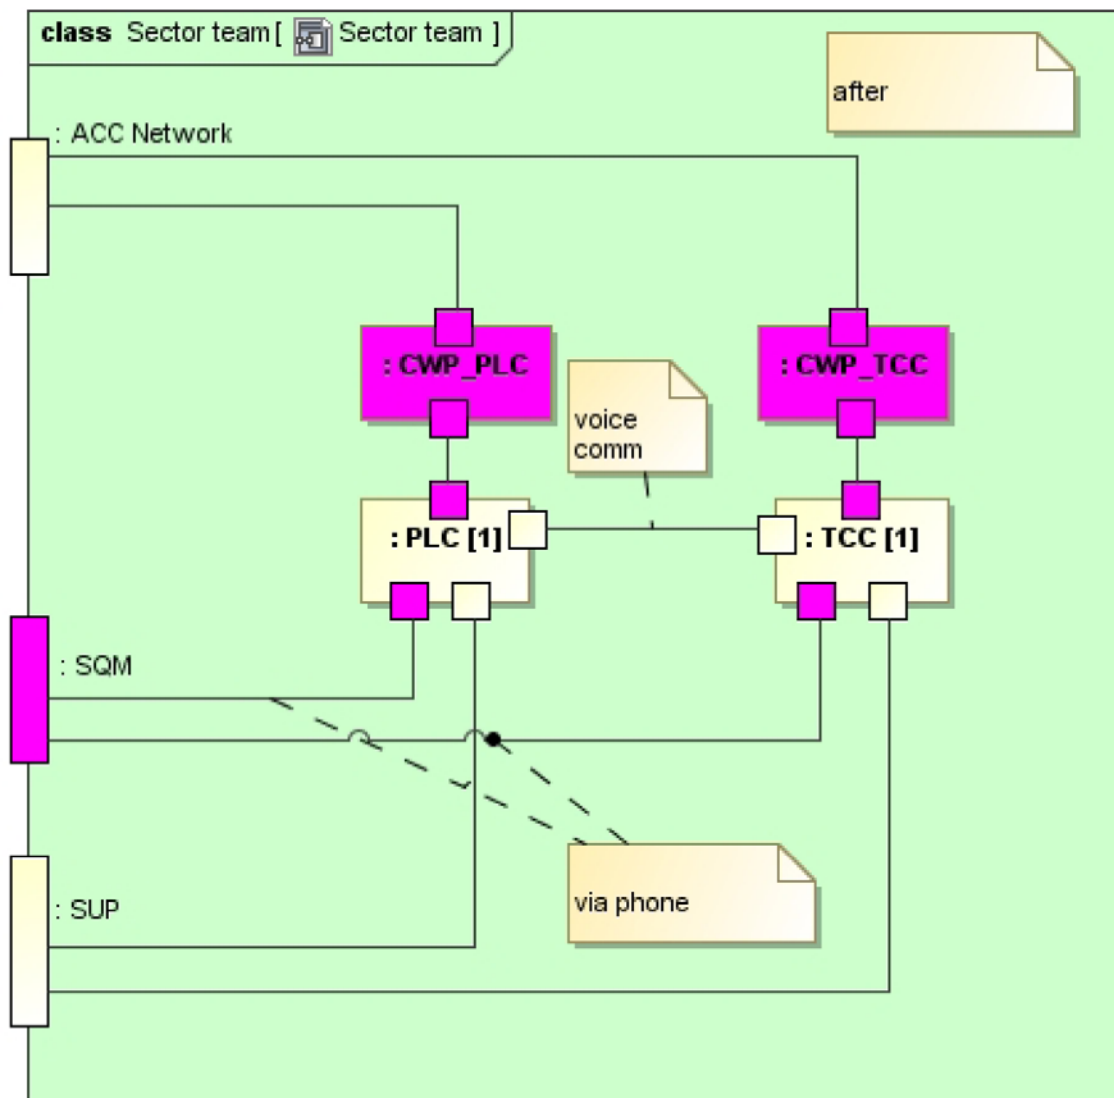


Figure 86 Internal structure of sector team after changes

Interactions

The UML interaction overview diagram of Figure 87 gives a high-level overview of the various tasks arrival management tasks after the changes. As the naming of the tasks indicates, they can be structured into five main tasks, summarized as follows:

- **Task 1:** Monitoring of aircraft (A/C) in the sector
- **Task 2:** Acquisition of the AMAN provided sequence
- **Task 3:** AMAN sequence monitoring and verification
- **Task 4:** Clearances to the A/C for building the planned sequence
- **Task 5:** Progressive transfer of the whole sequence to the adjacent sector

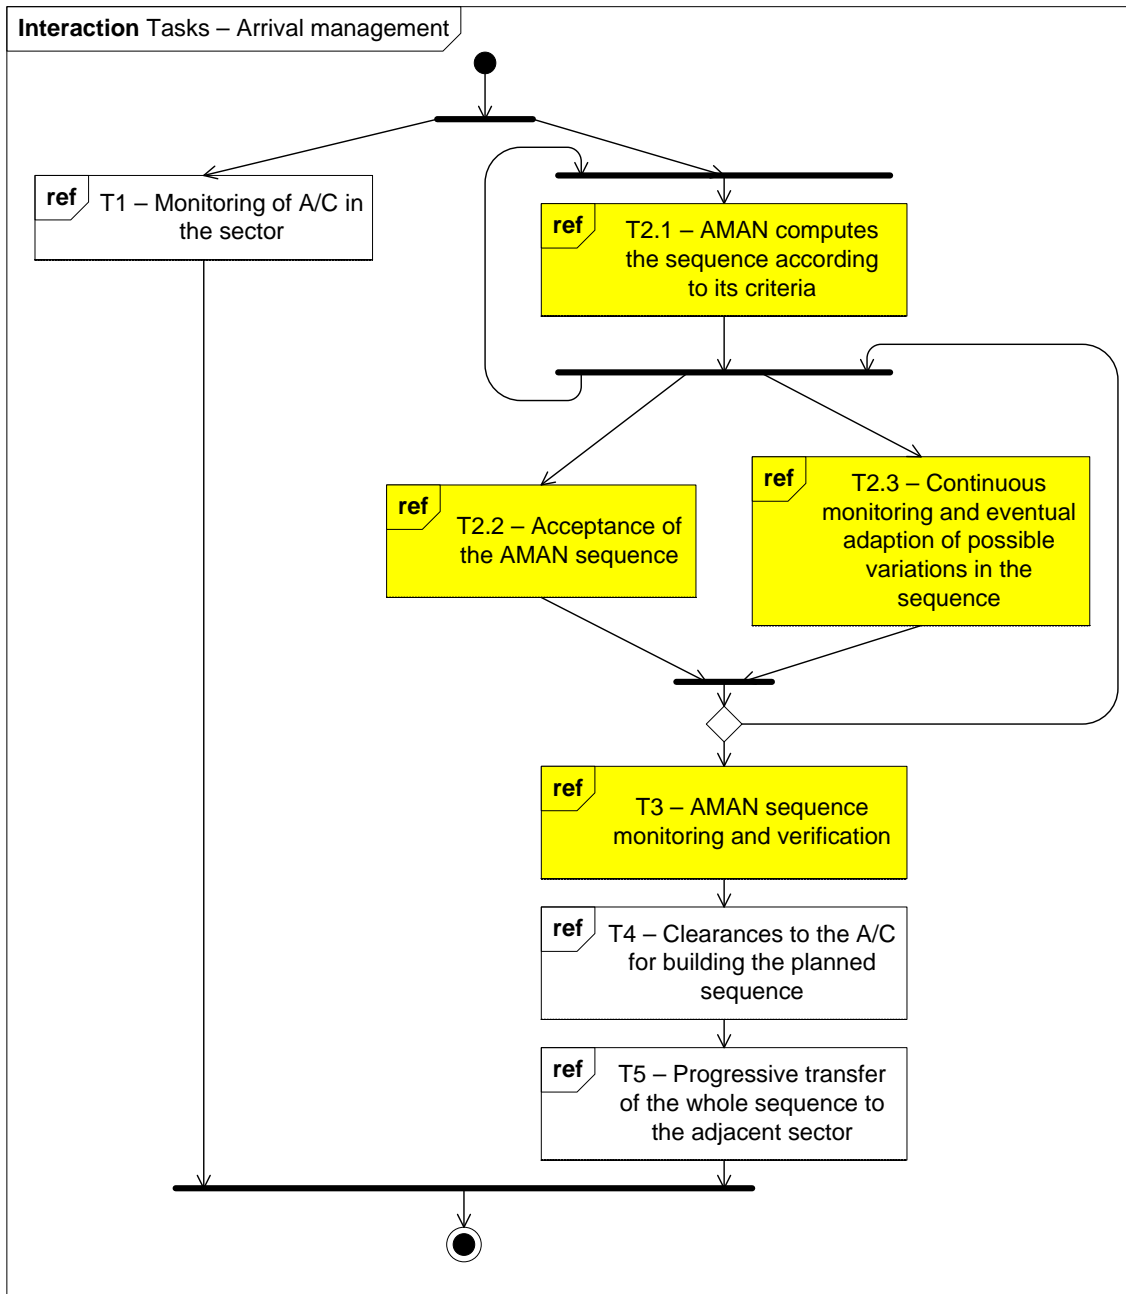


Figure 87 Overview of arrival management tasks after changes

Task T1, T4 and T5 remain the same before and after the changes, apart from some minor changes to T1 with insignificant impact on the analysis.

The sequence diagram of Figure 88 shows that activity T1 in principle remains the same after the changes.

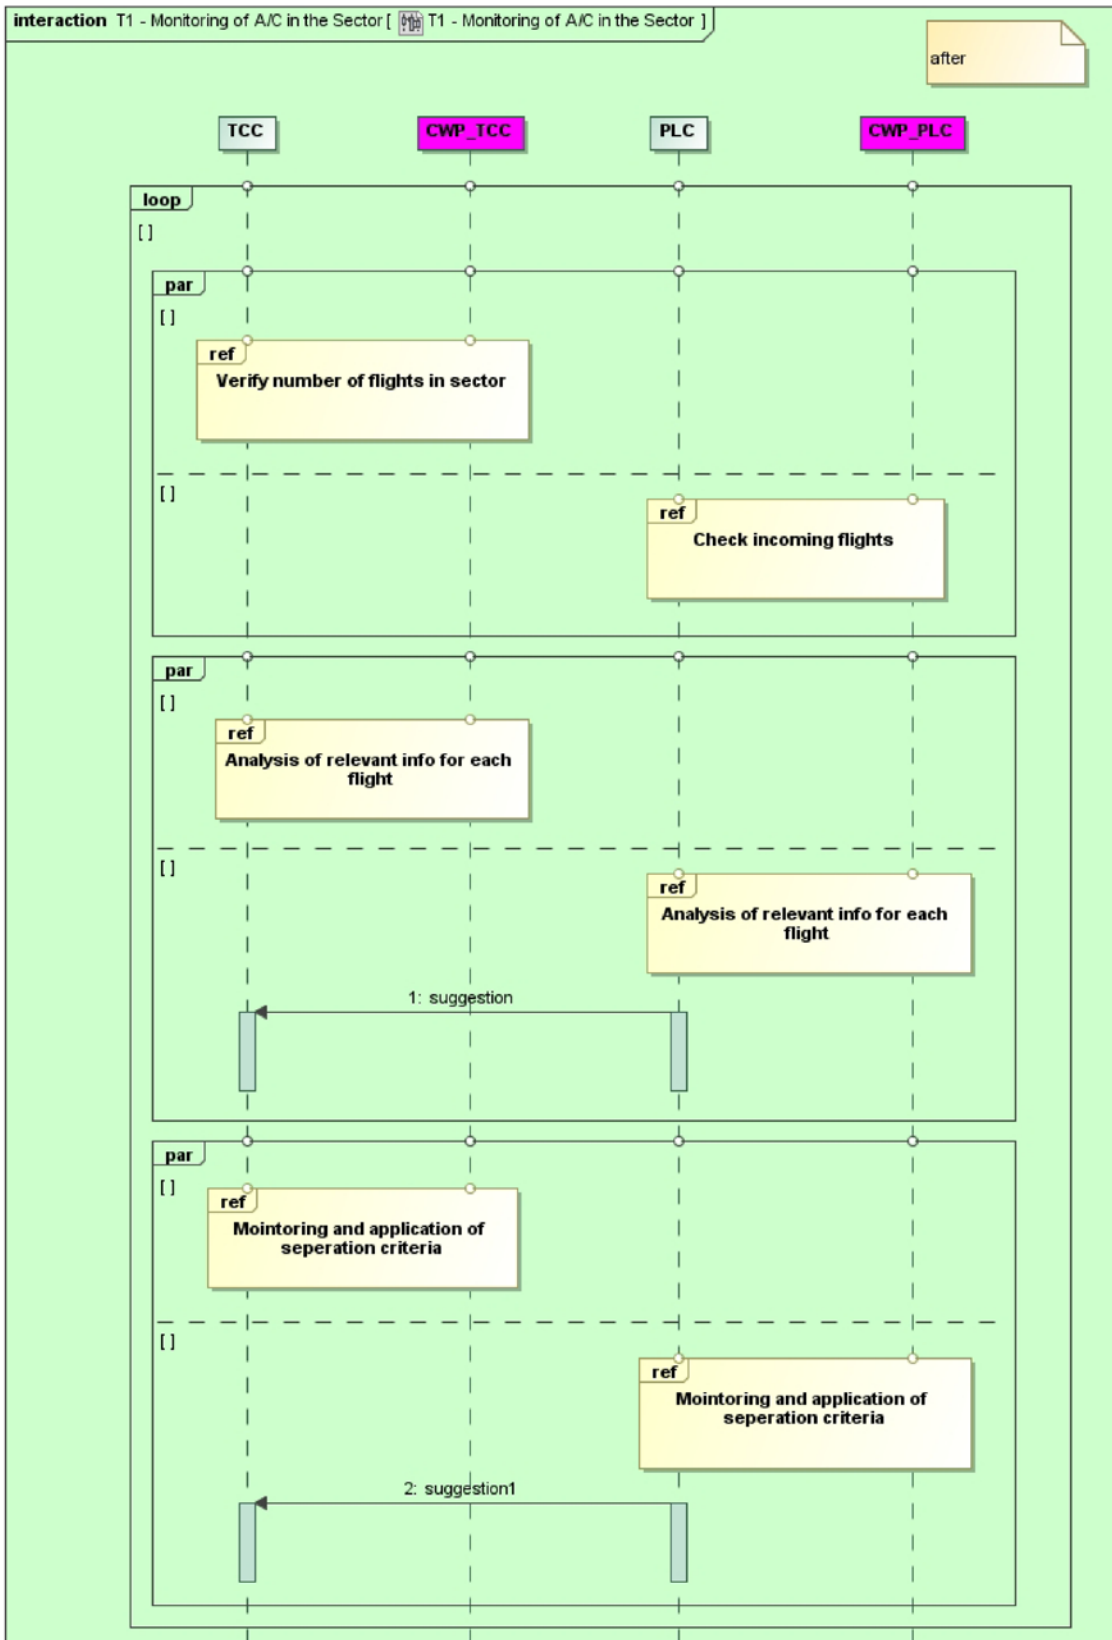


Figure 88 Task T1 after changes

The sequence diagram of Figure 89 shows task T2.1, AMAN computation of the sequence. The AMAN implements an algorithm that calculates sequences according to certain criteria. The criteria are partially based on the aircraft positions as provided by radar and ADS-B and other flight data. The AMAN computed sequences are provided to ATCOs via the CWPs.

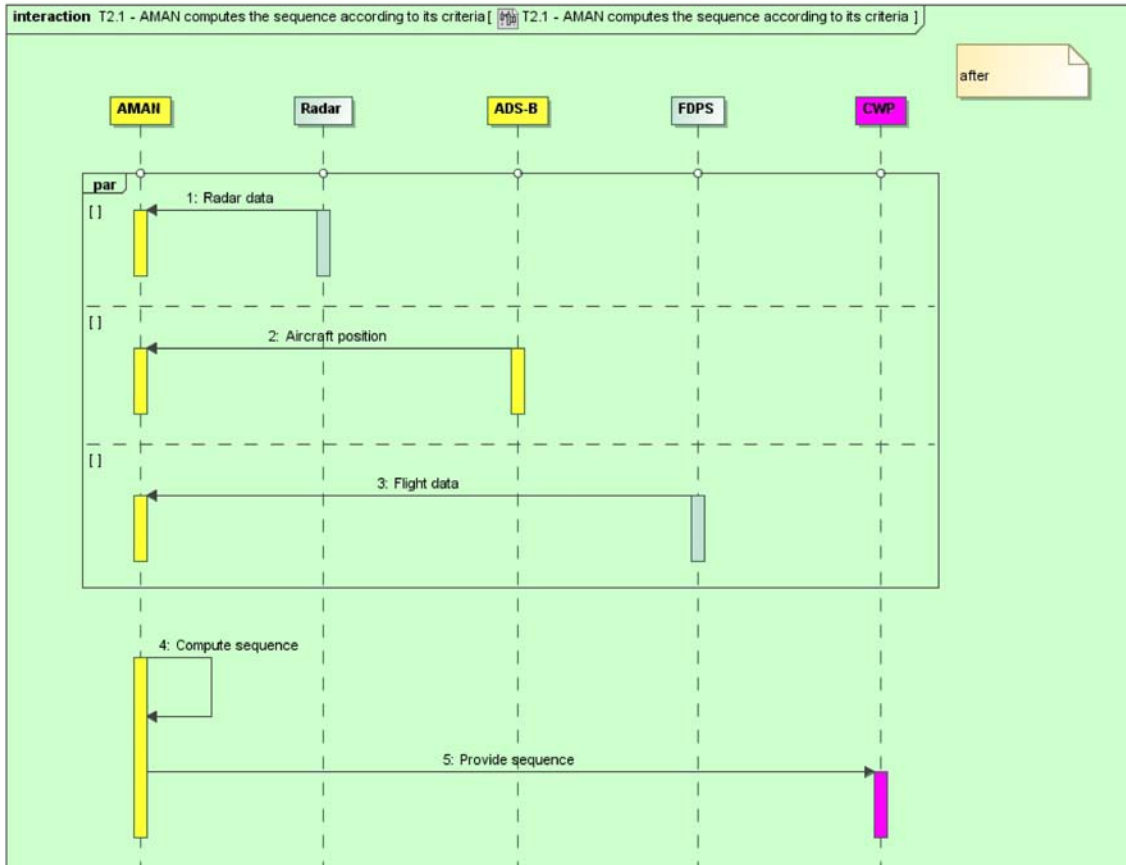


Figure 89 Task T2.1 after changes

The sequence diagram of Figure 90 shows task T2.2, acceptance of the AMAN sequence, as a task conducted by the TCC. Task T2.3, the continuous monitoring and eventual adaption of possible variations in the sequence, is a more compound task. The UML interaction diagram of Figure 91 shows the further decomposition of T2.3 into five subtasks that may be conducted in several iterations. The continuous monitoring of the AMAN proposed sequence is conducted by the TCC, but the TCC is supported by the other ATCOS, and particularly by the PLC in the team. The TCC furthermore uses flight and surveillance data, and also receives requests from pilots.

The formal acceptance and the continuous monitoring and updating of the AMAN sequence are up to the SQM. The TCC manages the traffic and in principle has just to follow AMAN instructions. Only in case he observes a very strange or contradictory sequence, he can call the SQM by phone and ask for clarifications. In case of emergency he shall ask for a sequence modification.

The increased responsibility of the SQM (previously COO) is furthermore shown in task T2.3, as well as in task T3 below. The SQM is now responsible for manually updating the AMAN sequence in case this is required as a consequence of what is monitored.

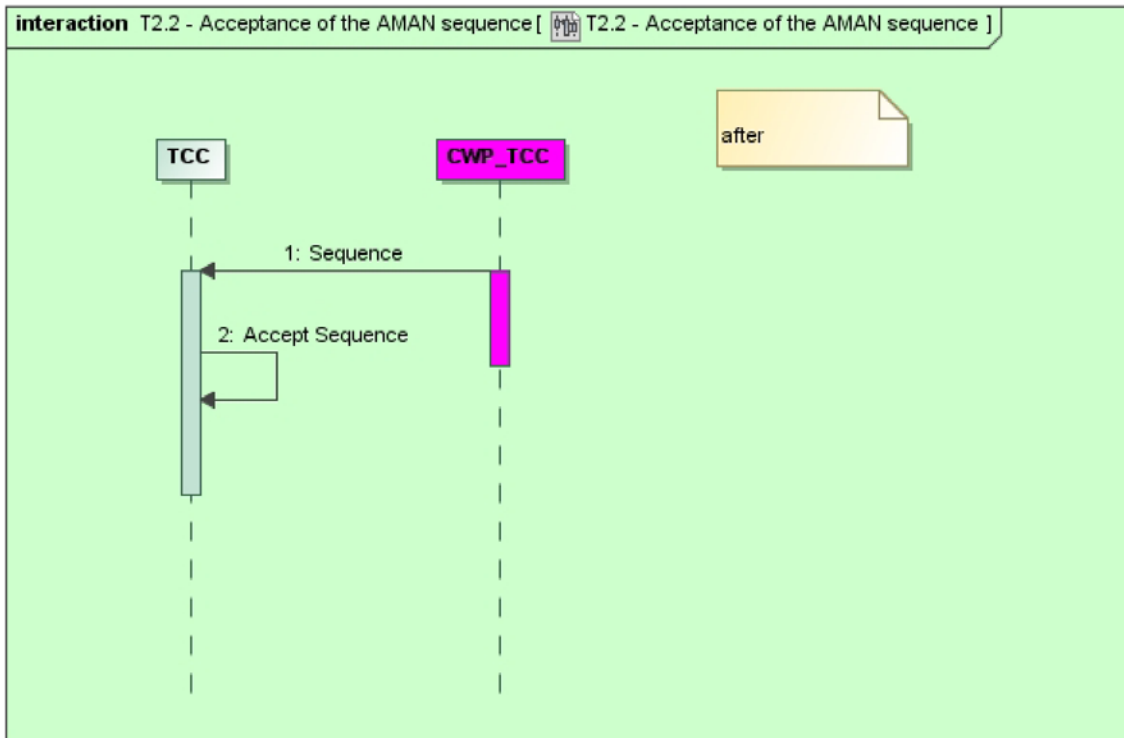


Figure 90 Task T2.2 after changes

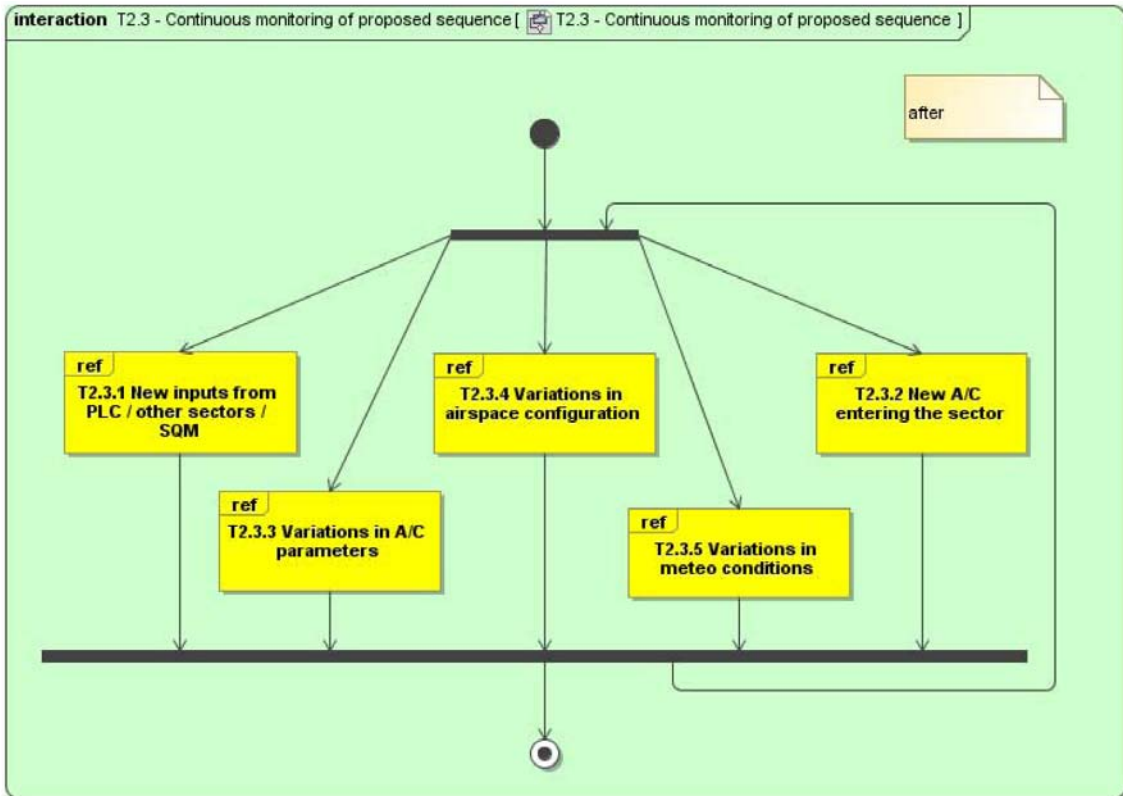


Figure 91 Overview of task T2.3 after changes

The sequence diagram of Figure 92 shows task T2.3.1 and the input that is provided to the TCC for supporting the sequence monitoring. Possible requests for changing the sequence are passed from the TCC to the SQM that updates the AMAN sequence.

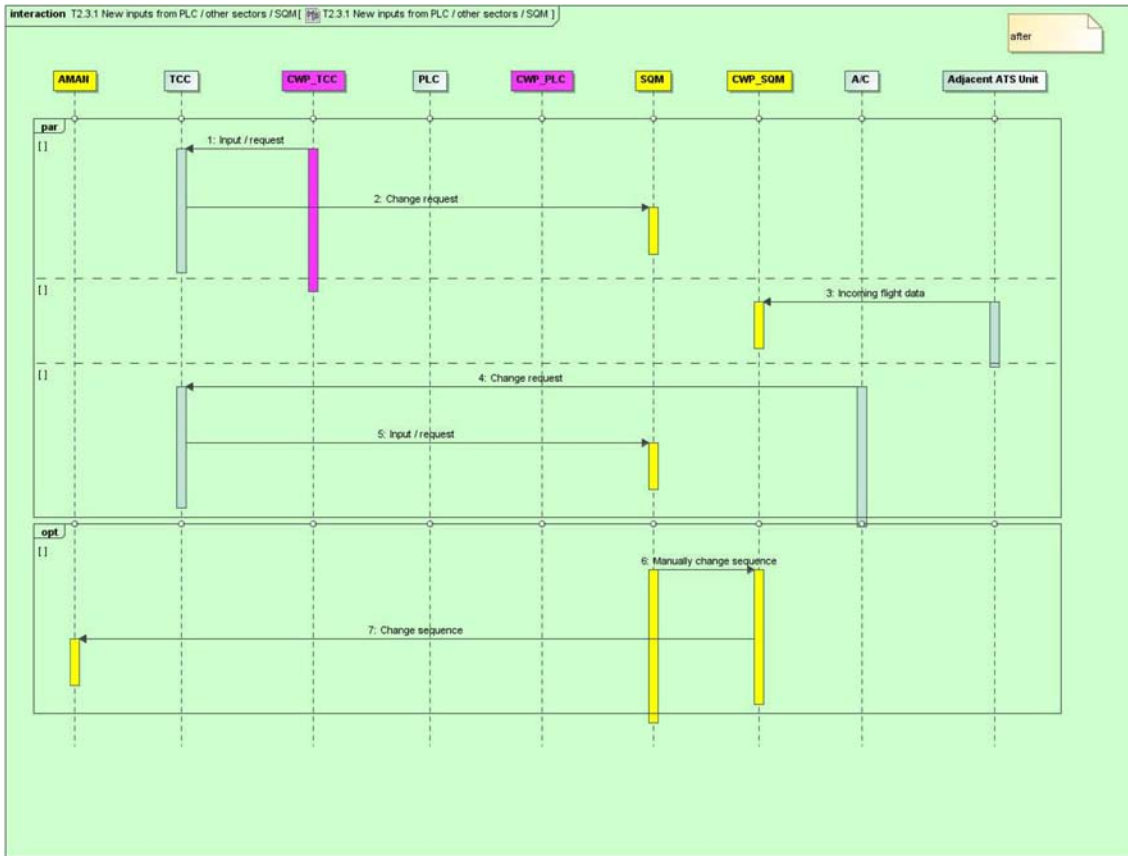


Figure 92 Task T2.3.1 after changes

The sequence diagram of Figure 93 shows task T2.3.2 and the information that is provided to the TCC when a new aircraft is entering the sector. In emergency cases, possible change requests are passed to the SQM that updates the AMAN sequence.

The sequence diagram of Figure 94 shows task T2.3.3 and the variations in aircraft parameters as fed directly to the AMAN for possible recalculation of the sequence.

The sequence diagram of Figure 95 shows task T2.3.4 and the possible reconfigurations that must be taken into account. The SUP is responsible for configuring the airspace, and possible changes are fed to the AMAN.

The sequence diagram of Figure 96 shows task T2.3.5 and the variations in the meteo conditions that are fed to the AMAN which in turn recalculates and updates the sequence if necessary.

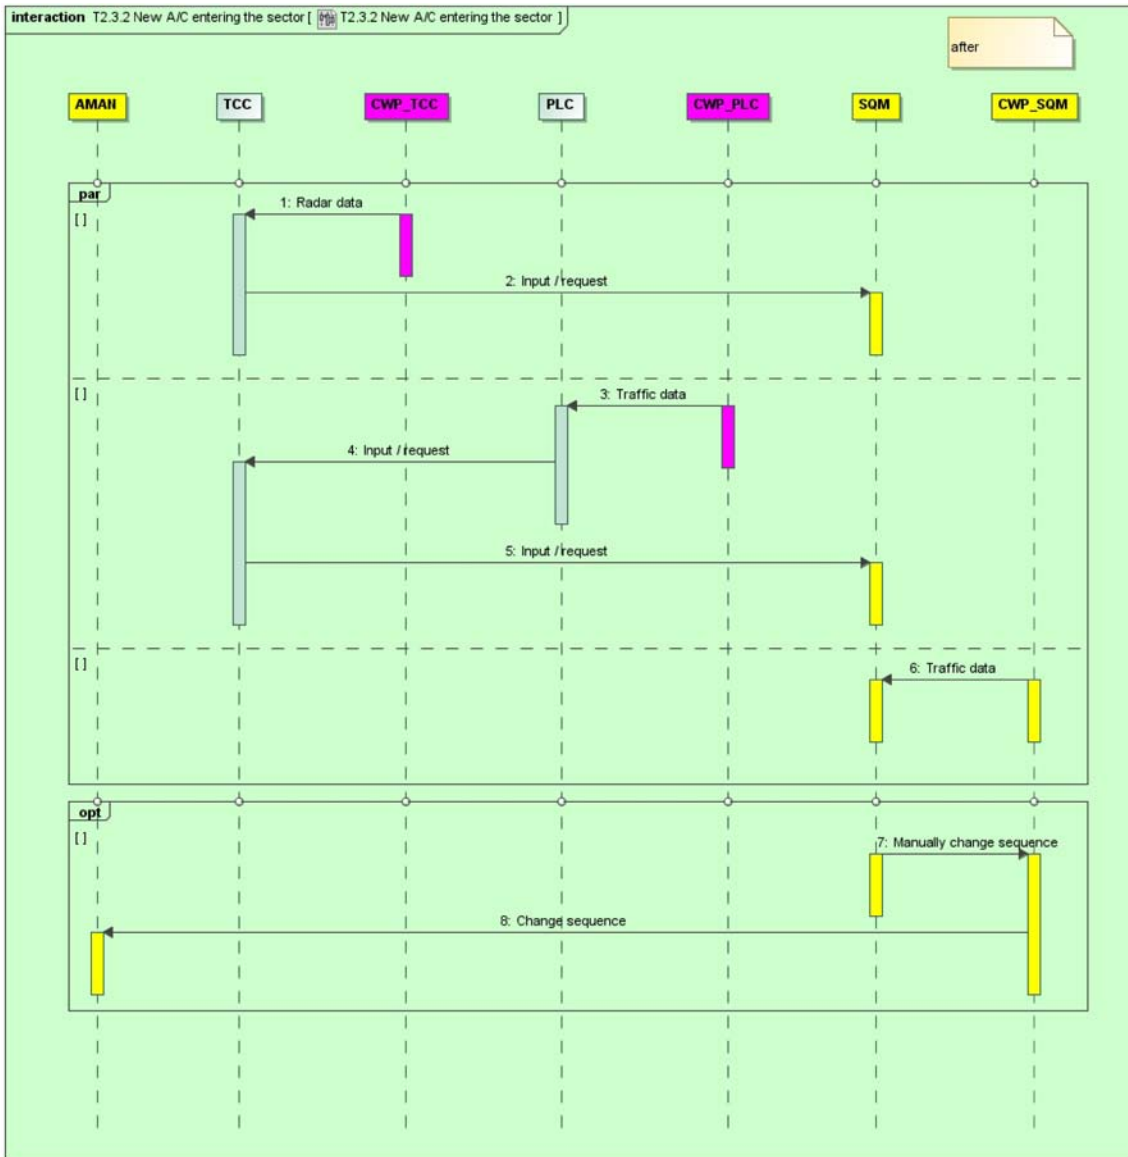


Figure 93 Task T2.3.2 after changes

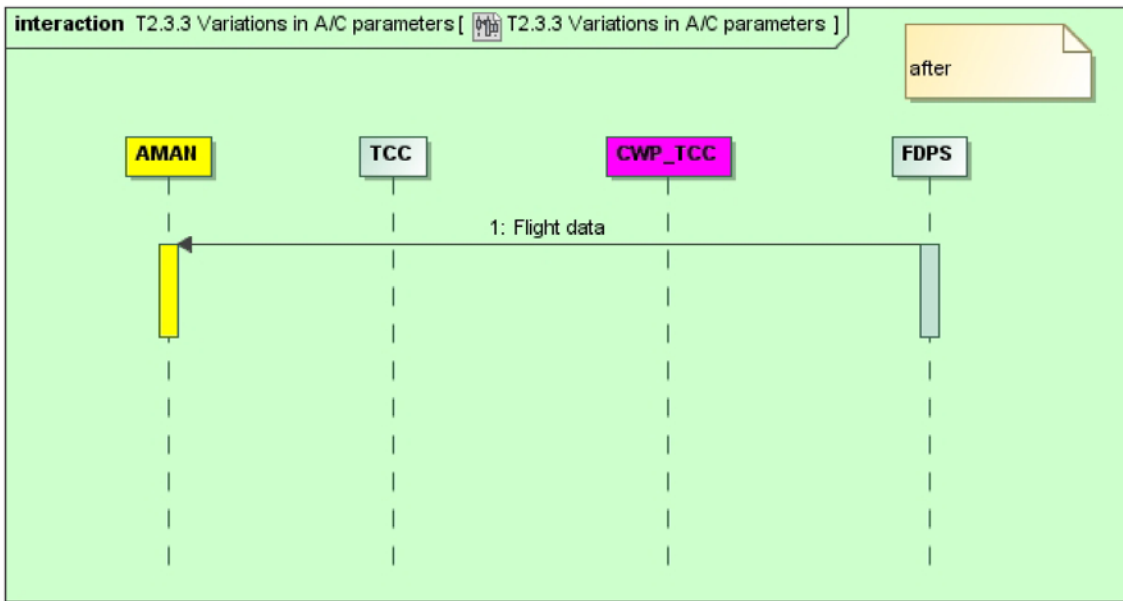


Figure 94 Task T2.3.3 after changes

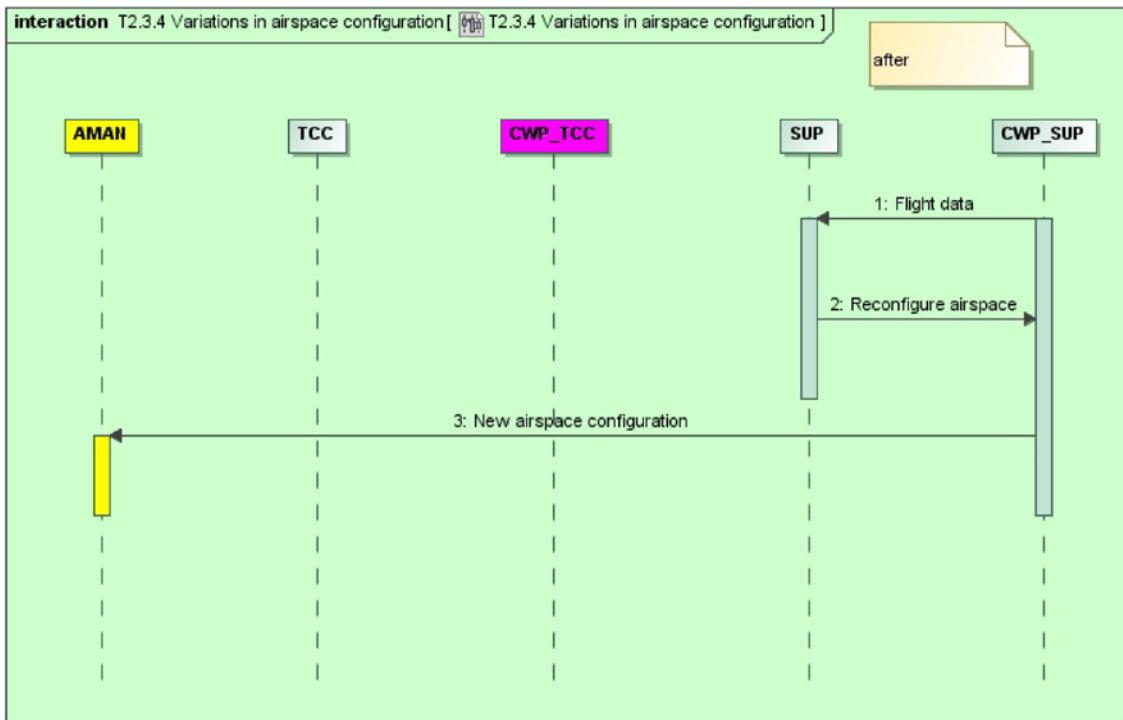


Figure 95 Task T2.3.4 after changes

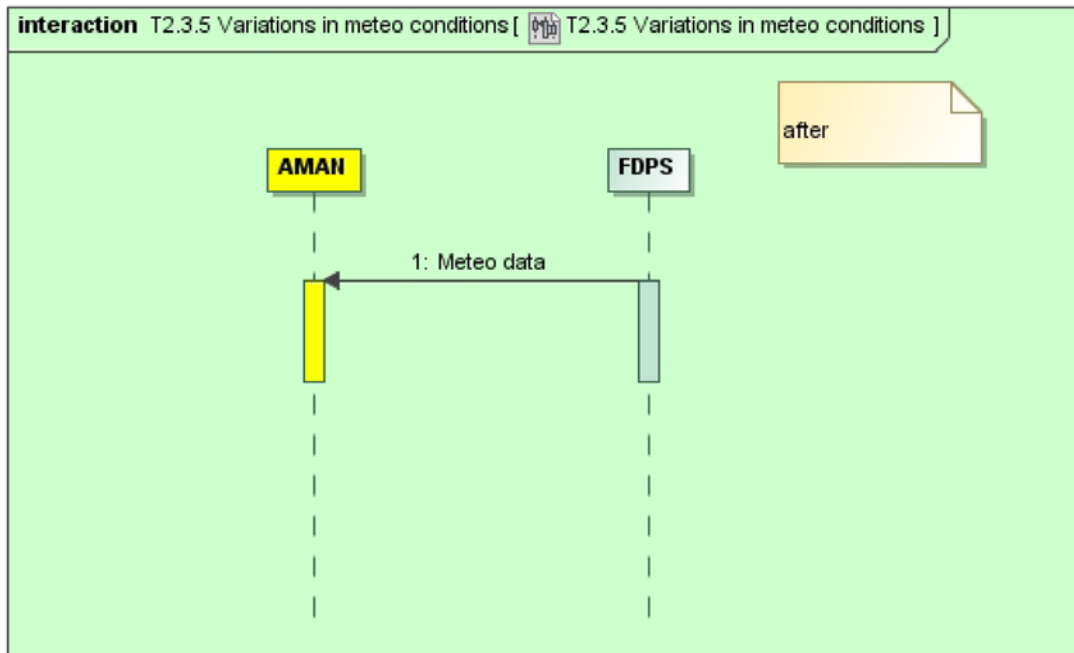


Figure 96 Task T2.3.5 after changes

Task T3, the AMAN sequence monitoring and verification is also a compound task that consists of three sequential sub-tasks as depicted in Figure 97.

The sequence diagram of Figure 98 shows task T3.1. Verification of the AMAN provided sequence is the responsibility of the SQM.

The sequence diagram of Figure 99 shows that task T3.2 of the sequence verification is conducted by the TCC.

The sequence diagram of Figure 100 shows task T3.3. Any change request by the TCC during the AMAN sequence monitoring and verification is passed to the PLC that in turn reports to the SQM. The SQM is responsible for manually updating the AMAN sequence.

The remaining tasks T4, clearances to the A/C for building the planned sequence, and T5, the progressive transfer of the whole sequence to the adjacent sector, are as before.

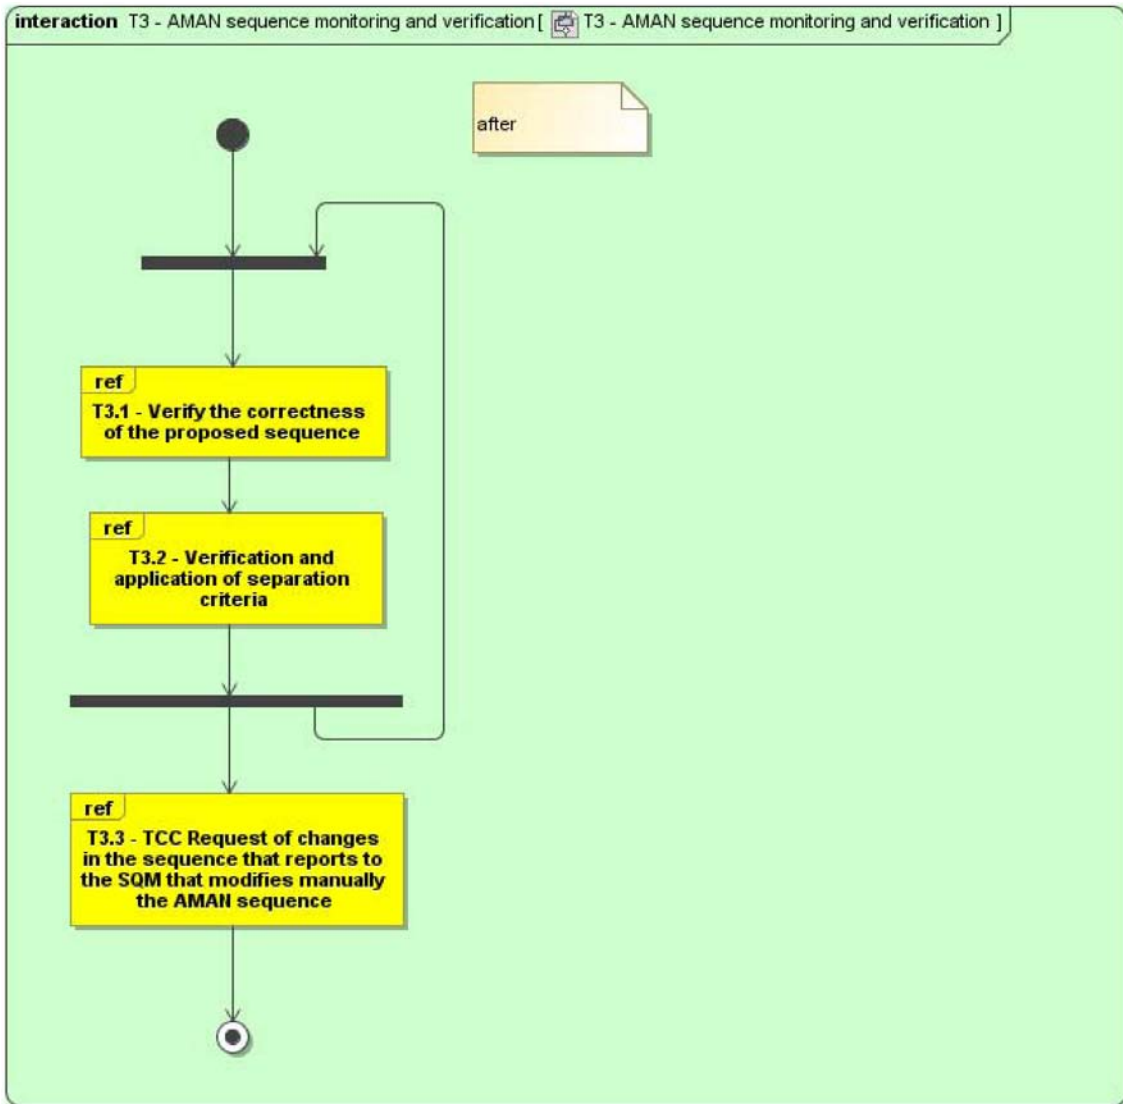


Figure 97 Overview of task T3 after changes

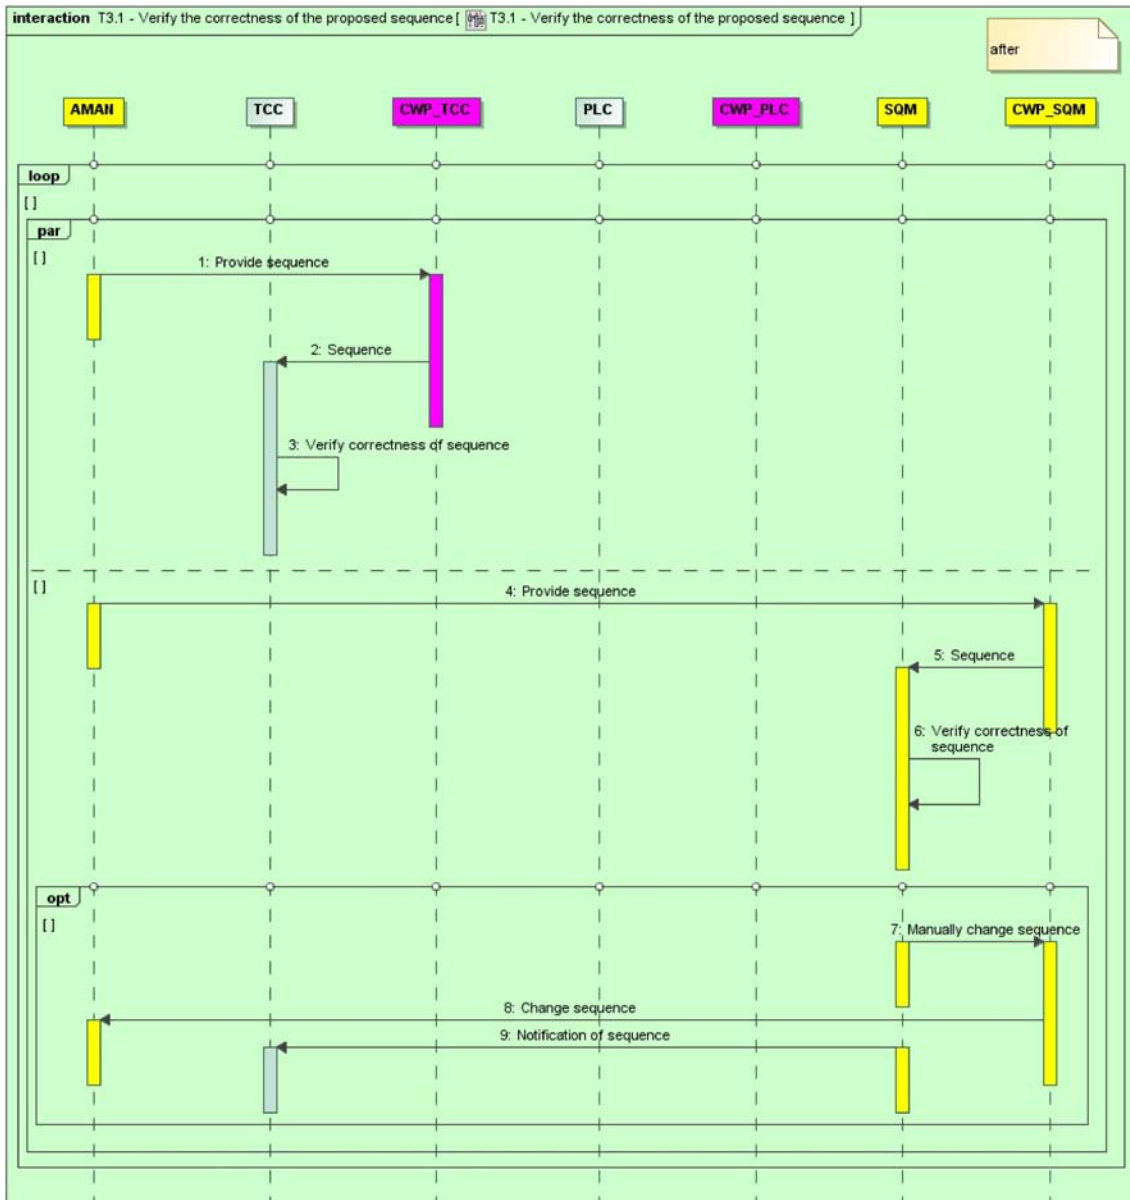


Figure 98 Task T3.1 after changes

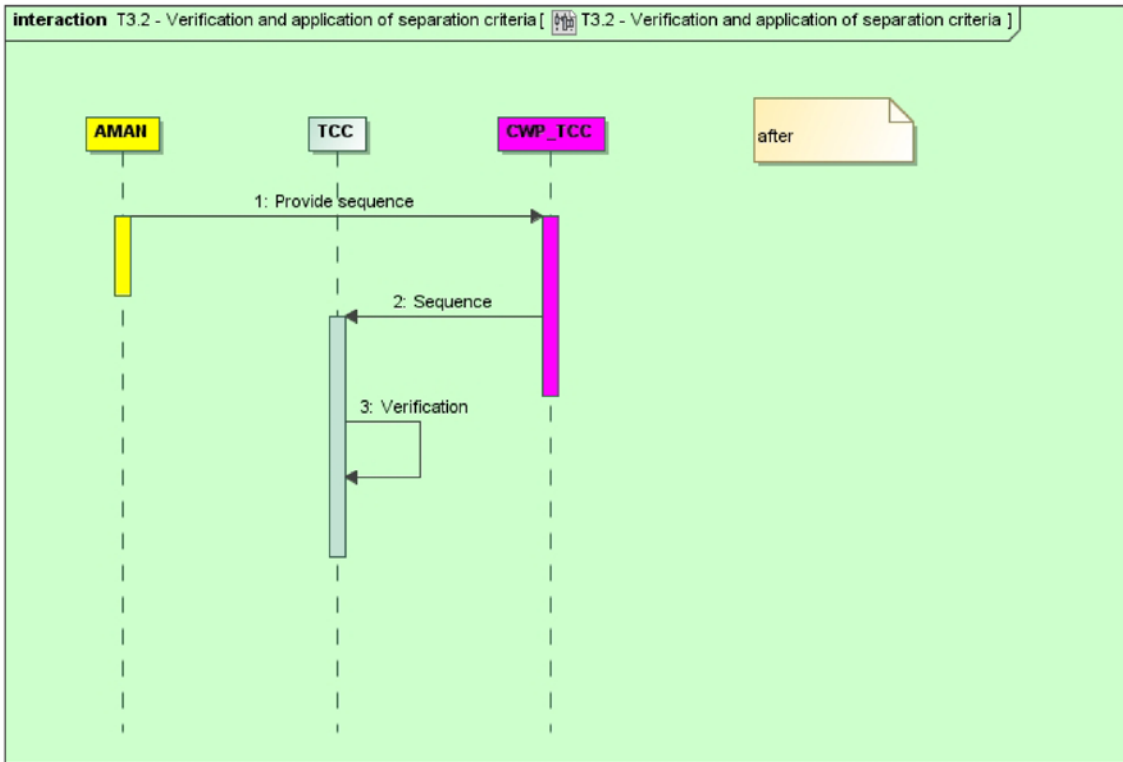


Figure 99 Task T3.2 after changes

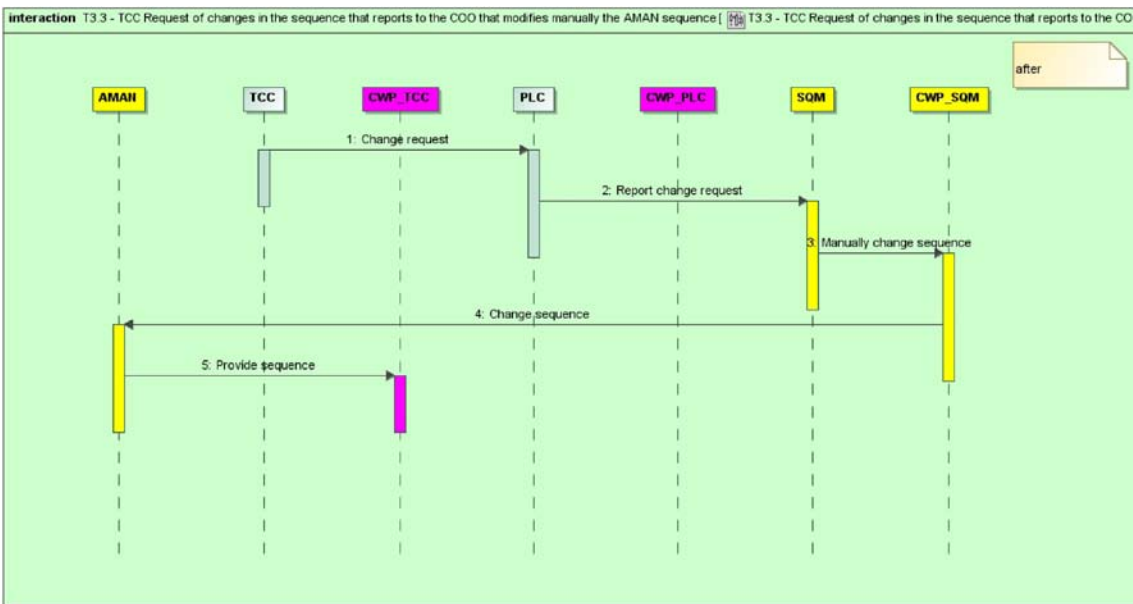


Figure 100 Task T3.3 after changes

14.1.3 Asset Identification

The purpose of the asset identification is to identify the parts, aspects or properties of the target with respect to which the risk analysis will be conducted. An asset is



something to which a party assigns value, and hence for which the party requires protection. A party is the organization, company, person, group or other body on whose behalf the risk analysis is conducted.

In this analysis, the party is the ATM service provider who owns the Area Control Center in question. The risk analysis addresses security issues, focusing on the following two security properties, selected from [30]:

- **Information protection:** Unauthorized actors (or systems) are not allowed to access confidential queue management information.
- **Information Provision:** The provisioning of information regarding queue management sensitive data by specific actors (or systems) must be guaranteed 24 hours a day, 7 days a week, taking into account the kind of data shared, their confidentiality level and the different actors involved.

The risk analysis is conducted with respect to these security properties by operating with the two corresponding assets of confidentiality and availability. The precise interpretation of these assets throughout the risk analysis is confidentiality of queue management information and availability of queue management information, respectively. Because the focus of the analysis is arrival management, the queue management information is restricted to arrival management information.

We use CORAS asset diagrams to document the assets and the relations between them. One asset is related to another if harm to the former may lead to harm to the latter. The two assets of availability and confidentiality are not related in such a way, as depicted by the asset diagram of Figure 101. The asset diagram also documents the party that requires protection of these assets.

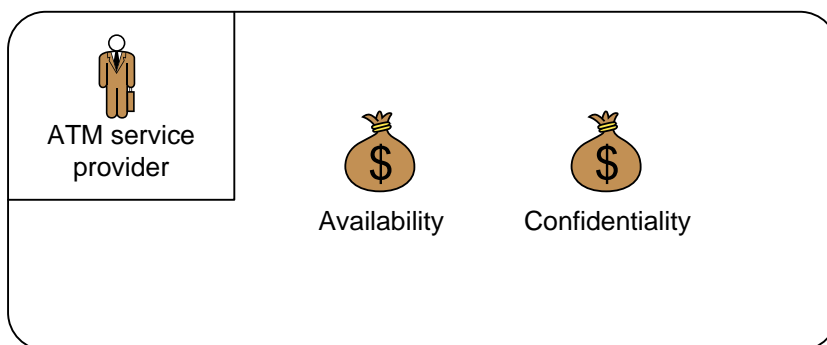


Figure 101 Asset diagram before and after

When we are considering changes, we need to take into account that not only the target of analysis may change, but also the parties and assets. Due to substantial changes, it may be that new assets emerge or previous assets disappear. It may also be that we need to take into account other parties after the changes. In this analysis, however, the party and the assets remain unchanged.

14.1.4 High-level Risk Analysis

The purpose of the high-level risk analysis is to complement the target models and the asset diagrams in increasing our understanding of the focus and scope of the risk analysis. This is a rough, initial risk analysis that aims to identify the main worries and

main incidents so that we can better decide what to include and not, and also to get a better grip of the very motivation for the risk analysis in the first place.

We use table formats for documenting the results of the high-level risk analysis. The high-level risks analysis is in the following documented separately for before and after the changes.

14.1.4.1 High-level Risk Analysis before AMAN Introduction

Table 30 documents the results of the high-level analysis before the changes are taken into account. The table format is of four columns. The first column documents the relevant threats for the corresponding row. A threat is the potential cause of an unwanted incident. The second column documents threat scenarios, unwanted incidents and the assets that may be harmed, i.e. it documents what can go wrong. A threat scenario is a chain or series of events that is initiated by a threat and that may lead to an unwanted incident. An unwanted incident is an event that harms or reduced the value of an asset. The third column documents vulnerabilities, i.e. weaknesses, flaws or deficiencies that opens for, or may be exploited by, a threat to cause harm or reduce the value of an asset. The fourth column documents the parts or elements of the target where the risks in questions may arise. This is to facilitate keeping track of which risks that may be affected by changes to the target.

The results of the high-level analysis show that there are two main kinds of worries. On the one hand the high-level analysis focuses on component, system and communication failures that can lead to loss of availability. On the other hand, the analysis focuses on the human factor, as documented by the last row. Although represented by only one row, there may be many security issues in relation to human factors, and these are therefore be more thoroughly addressed during the full risk identification.

The human factors are particularly interesting for this analysis, since the change requirements concern the introduction of decision support systems that should mitigate related risks. A part of the analysis therefore aims at investigating to what extent such risks change with the introduction of the AMAN.

A further interesting finding is that no specific incidents regarding confidentiality were documented during the high-level analysis. This may indicate that confidentiality issues are less critical in the traditional, closed ATM settings.

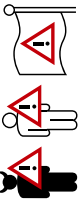



 Who/what causes it?	 What is the scenario or incident? What is harmed?	 What makes it possible?	 Target element
Component failure	Provisioning of information to ATCO fails due to loss of a single CWP	Insufficient CWP maintenance	CWP
Software error	The consolidation of data from several radar sources fails, leading to duplication of labels		Surveillance
Component failure; radar disturbance	Malfunctioning of radar antenna leads to loss of or degradation of radar signal	Insufficient Radar maintenance	Radar
System failure	ACC network interruption leads to loss of information provisioning to/from ATCOs	Insufficient network redundancy	ACC Network
Software bugs	False or redundant alerts from safety tool	Insufficient software testing	OPS Room
Communication Failure	Loss of Communication lines and/or loss of radio frequency	Lack of backup communication	Communication
Malicious employee at telecom service provider	External adversary corrupts ATM information flow over the dedicated communication system	ACC external communication relies on third party service provider	Communication
ATCO	ATCO fails to conduct ATM tasks due to accident, stress or sloppiness	High workload; stress	ATCO

Table 30 High-level analysis table before changes

14.1.4.2 High-level Risk Analysis after AMAN Introduction

The first step of the high-level analysis after the changes is to conduct a walkthrough of the high-level analysis table before changes to identify risks that are persistent under the changes. This task is facilitated by the fourth column that refers to the relevant parts and elements of the target. In this analysis all entries in Table 30 applies also after the changes. This table therefore also serves to document the high-level analysis after the changes. Importantly, these risks may change in severity, i.e. their risk levels may increase or decrease. In the full risk assessment, these risks are therefore evaluated both before and after the changes.

Table 31 documents the additional results of the high-level analysis where the changes are taken into account. The entries represent risks that may arise after the introduction of the AMAN and the ADS-B system.

The provisioning of the queue management information relies much on the AMAN, and risks in relation to loss or failure of the AMAN should therefore be considered. The use of the ADS-B may furthermore cause confidentiality issues. There are finally worries that the workload on the SQM, formerly the COO, could be critical for security under certain circumstances.









  	  		
Who/what causes it?	What is the scenario or incident? What is harmed?	What makes it possible?	Target element
System failure	Loss of the AMAN leads to loss of provisioning of information to ATCO		AMAN
Attacker	Attacker broadcasts false ADS-B signals which leads to the provisioning of false arrival management data	Use of ADS-B; dependence on broadcasting	ADS-B
Attacker	Confidentiality breach by attacker eavesdropping on ADS-B	Use of ADS-B; dependence on broadcasting	ADS-B
System failure	Loss of phone connection leads to SQM failing to update the TCC on manually updated sequences		Communication
Software fail	Provisioning of unstable or incorrect sequences by the AMAN leading to ATCO reverting to manual sequencing	Immature software	AMAN
SQM	SQM fails to build stable sequence or make optimal coordination	High workload on SQM after AMAN introduction	SQM; AMAN
AMAN	Sequence not optimal due to lack of flexibility in AMAN in taking all relevant factors into account	Simplistic AMAN algorithm	AMAN
TCC	TCC fails to provide arrival information to all relevant recipients simultaneously due to communication overload	Lack of routines for avoiding multitasking	TCC
SQM	Incorrect modification of sequence into the AMAN	Lack of routines for verification of modified sequence	SQM

Table 31 High-level analysis table after changes

14.1.5 Establishing the Risk Evaluation Criteria

The risk evaluation criteria define the level of risk that the party, i.e. the ATM service provider, is willing to accept for the given target of analysis. Basically, the criteria are a mapping from risk levels to the decision of either accepting the risk or evaluating the risk further for possible treatment.

In order to speak of risk levels, we need first to define the risk function. The risk function is a mapping from pairs of consequence and likelihood to risk levels. Before we can define the risk function we need to define the consequence scales and the likelihood scales. Since the kinds of consequences may be different for different assets, we define one consequence scale for each kind of asset.

When addressing a changing target of analysis, it may be that the risk evaluation criteria also change. In this analysis, however, the same criteria apply both before and after the changes. The criteria, as well as the consequence scales, the likelihood scale and the risk functions, are therefore documented only once.

14.1.5.1 Setting the Consequence Scales

In order to estimate and evaluate risks, we need to be able to describe and talk about the potential harm that may be caused by the risks. The consequence of a risk describes the level of damage the associated unwanted incident inflicts on an asset when the incident occurs. When we are setting the consequence scales we define the set of values that we use for describing the possible consequences.

The consequence scale for the confidentiality asset is documented in Table 32. The scale applies both before and after the changes.

Consequence	Description
Catastrophic	Loss of data that can be utilized in terror
Major	Data loss of legal implications
Moderate	Distortion of air company competition
Minor	Loss of aircraft information data (apart from A/C position data)
Insignificant	Loss of publically available data

Table 32 Consequence scale for confidentiality before and after change

We operate in this analysis with qualitative scales with values ranging from insignificant to catastrophic. Each value is defined by giving a description of its severity. The descriptions may not apply to all confidentiality incidents, but the purpose is only to provide an understanding of the severity of the various values.

The consequence scale for availability is documented in Table 33. This scale is based on the EUROCONTROL Safety Regulatory Requirement 4 (ESARR4) [13]. The consequence descriptions given in Table 33 summarize the descriptions given in ESARR 4 which are more elaborate and provide more examples.

Consequence	Description
Catastrophic	Catastrophic accident
Major	Abrupt maneuver required
Moderate	Recovery from large reduction in separation
Minor	Increasing workload of ATCOs or pilots
Insignificant	No hazardous effect on operations

Table 33 Consequence scale for availability before and after changes

14.1.5.2 Setting the Likelihood Scale

A likelihood is the frequency or probability for something to occur. The likelihood scale of five quantitative values is documented in Table 34. The definitions are based on EUROCONTROL advisory material [12].

Likelihood	Description
Certain	A very high number of similar occurrences already on record; has occurred a very high number of times at the same location/time
Likely	A significant number of similar occurrences already on record; has occurred a significant number of times at the same location
Possible	Several similar occurrences on record; has occurred more than once at the same location
Unlikely	Only very few similar incidents on record when considering a large traffic volume or no records on a small traffic volume
Rare	Has never occurred yet throughout the total lifetime of the system

Table 34 Likelihood scale before and after changes

14.1.5.3 Defining the Risk Function

The risk function yields for each combination of a likelihood and a consequence the resulting risk level. Since the risk function is a mapping from likelihoods and consequences to risk values, we must define a separate risk function for each of the consequence scales. In our case, however, the risk functions we defined turned out to be equal for the two assets and are therefore documented by one risk matrix. The risk matrix shows for each combination of a likelihood and consequence the resulting risk level.

The risk function is documented in Table 35. We use three risk levels, namely low (green), medium (yellow) and high (red).

		Consequence				
		Insignificant	Minor	Moderate	Major	Catastrophic
Likelihood	Rare					
	Unlikely					
	Possible					
	Likely					
	Certain					

Table 35 Risk function before and after change

14.1.5.4 Deciding the Risk Evaluation Criteria

The risk evaluation criteria for both availability and confidentiality of queue management information are as follows:

- **High risk:** Unacceptable and must be treated.
- **Medium risk:** Must be evaluated for possible treatment.
- **Low risk:** Must be monitored.

14.2 Risk Identification

The risk identification was conducted as a structured brainstorming involving personnel with first hand knowledge about the target of analysis. By conducting a walkthrough of the target description and using the results of the high-level risk analysis, the risk were identified by systematically identifying unwanted incidents, threats, threat scenarios and vulnerabilities. The results were documented on-the-fly by means of CORAS threat diagrams.

While the various parts of the threat diagrams were modeled and documented, the relations to the target of analysis were identified and documented at the same time. The relations were documented by annotating the threat diagrams with the dedicated target element icon. The documentation of these relations facilitate identifying the parts of the threat diagrams that are unaffected by the changes.

The documentation of the relations between the risk models and the target system by means of these annotations is the visualization of the underlying trace model. We refer to Section 6.2 for the presentation of the artifact of the trace model and to Section 6.3 for the presentation of the language extension of risk graphs to allow the graphical specification of the trace model. Section 13.2.2 defines the instantiation of the trace modeling in the CORAS language with change. Due to constraints on time and resources, the full underlying trace models were not worked out during the risk assessment workshops. In the CORAS models of this appendix we therefore present only the graphical representations of the trace models and refer to Section 6 for concrete examples.

The approach to the risk identification is to first identify and document risks for the target of analysis before the changes. This is conducted according to traditional risk

identification methods and techniques, with the additional activity of explicitly documenting the relations to the target description. Once this is completed, we proceed by identifying and documenting the risks after the changes. Based on the documented relations to the target description, i.e. the trace model, we identify the parts of the threat scenarios that are not affected by the changes and therefore do not have to be addressed again from scratch.

In the following we first document the results of the risk identification before the changes, and thereafter we document the results after the changes. The threat diagrams for the latter explicitly show how risks change from before to after the changes.

14.2.1 Risk Identification before Changes

In the following we give some examples of the results of the risk identification before the changes in order to show how it was conducted. The full documentation is given below under the risk estimation.

The CORAS threat diagram of Figure 102 documents two unwanted incidents concerning loss of functionality of the CWPs or of the full OPS room. Such functionality loss may reduce the availability of arrival management information.

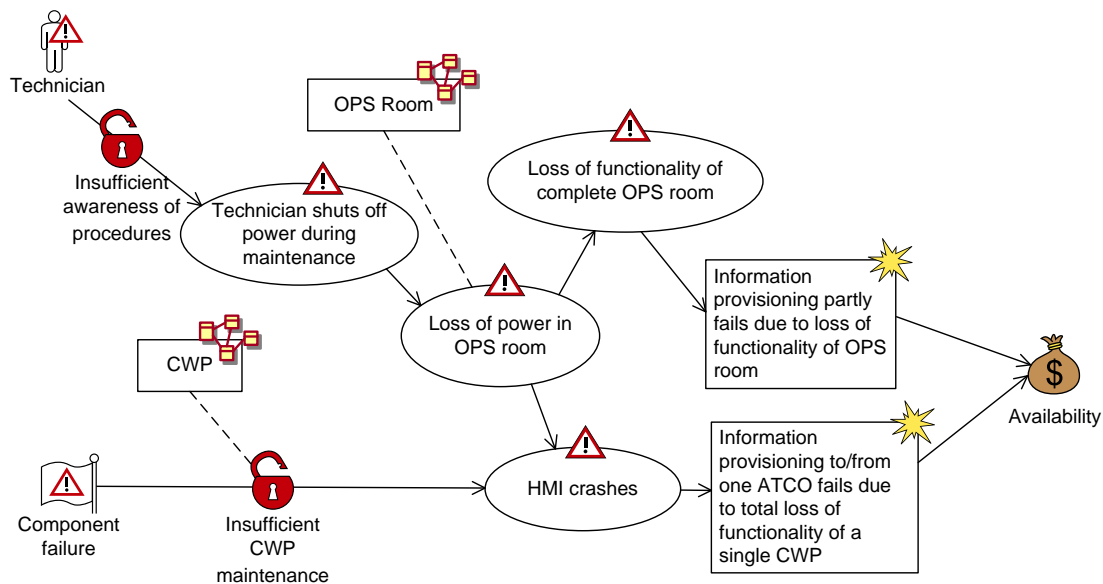


Figure 102 Loss of functionality in OPS Room before changes

The threat diagram of Figure 103 documents unwanted incidents that may arise due to duplication of labels on the CWP interface.

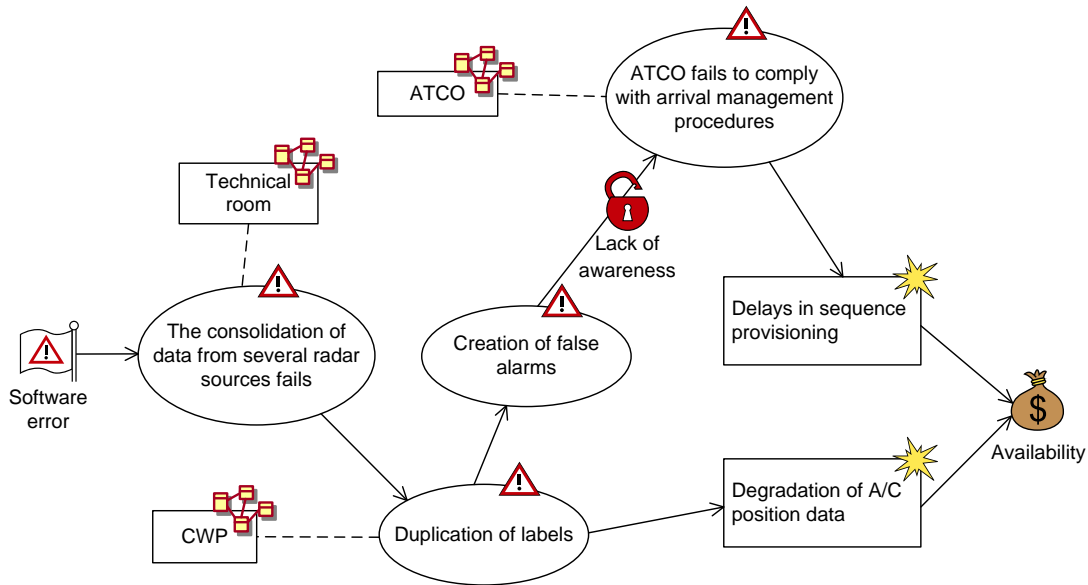


Figure 103 Label duplication

A label depicts an aircraft with its position data, and is derived from radar data. When several radar sources are used, the label is generated by automatically consolidating (merging) the data from the various sources. In some cases, software errors may yield duplicated labels that may lead ATCOs to believe there are two aircrafts. The duplication may also lead to false near miss alarms.

14.2.2 Risk Identification after Changes

The risk identification after the changes is conducted by first identifying and documenting the risks that are present both before and after the changes. The threat diagram of Figure 102, for example, concerns aspects of the CWPs and the OPS Room that are relevant both before and after the changes. This threat diagram therefore documents risks that are also relevant both before and after the changes.

The documentation of this threat diagram as persistent under changes is given in Figure 104. The “two-layered” icons convey that these elements are present both before (the below layer) and after (the front layer).

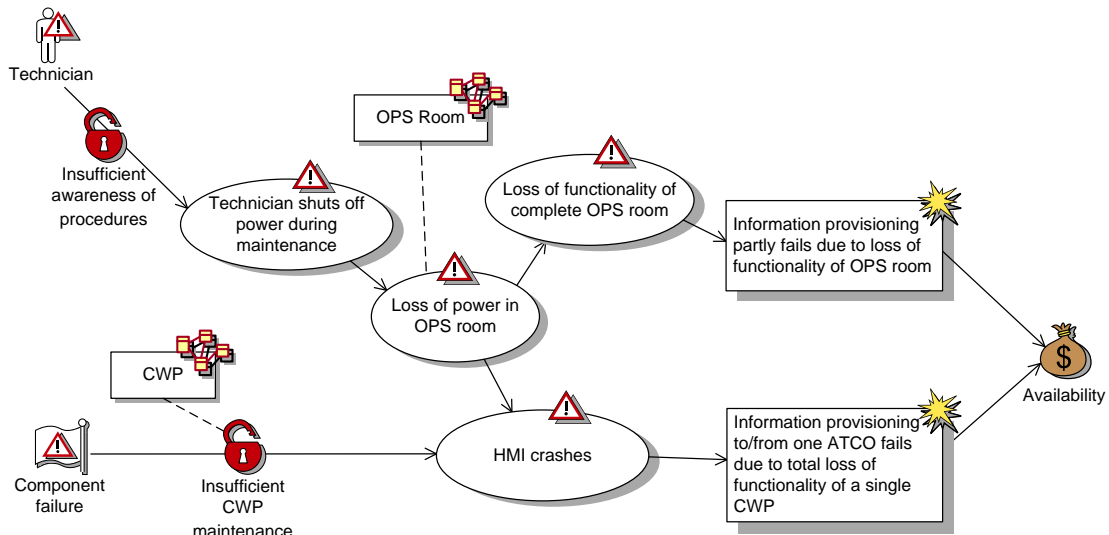


Figure 104 Loss of functionality in OPS Room before and after changes

This kind of threat diagrams gives immediate information about risks both before and after changes. Importantly, these risks may change in the sense that the likelihoods or consequences change. The latter changes are, however, identified and documented during risk estimation.

This kind of threat diagrams also documents elements of the risk picture that may emerge, and elements that may disappear after change. This is shown in the before-after threat diagram of Figure 105.

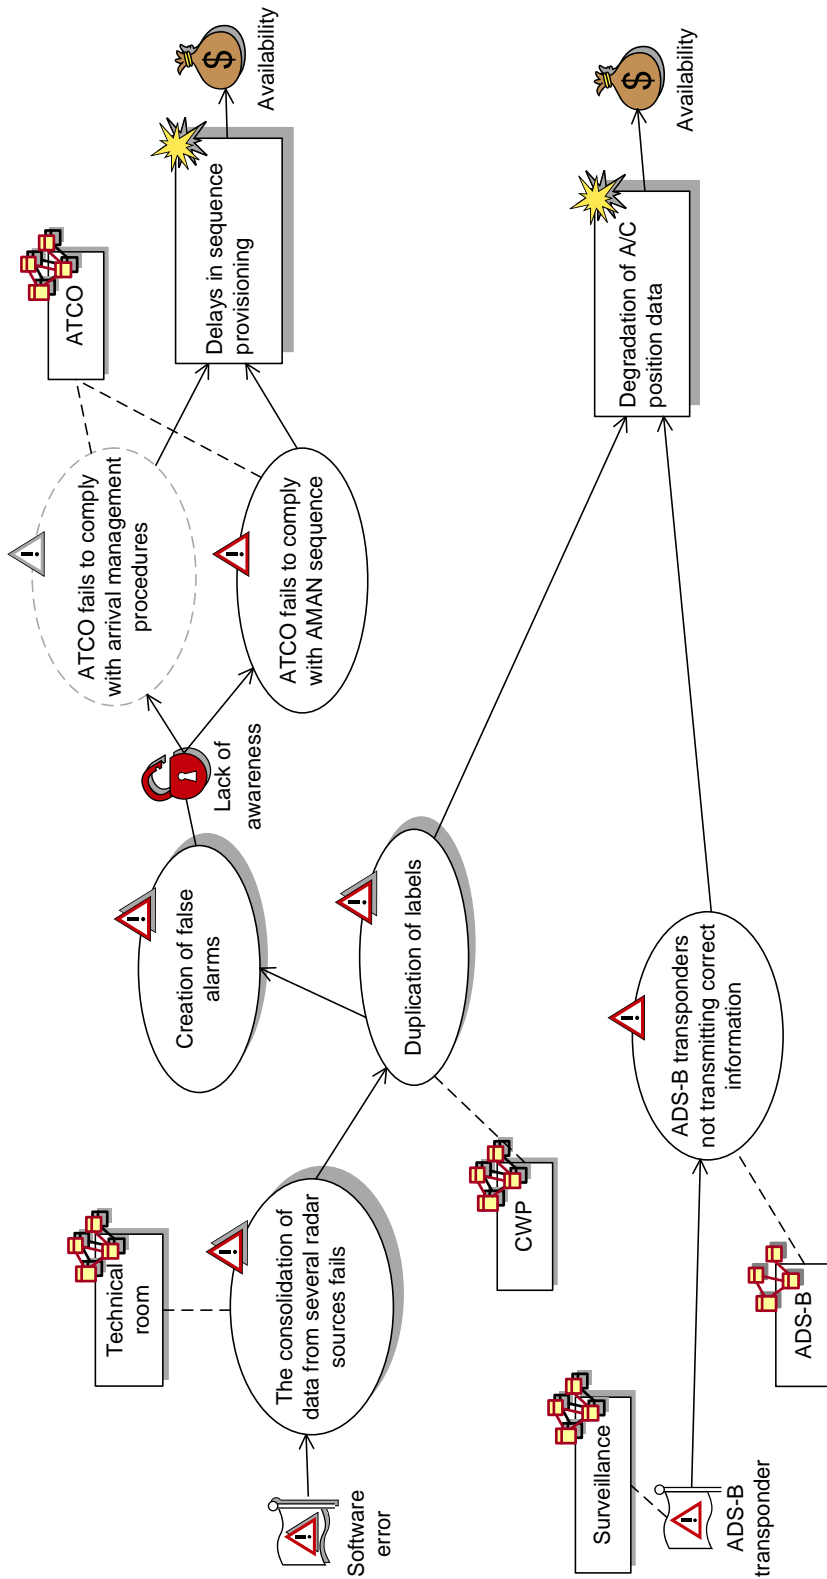


Figure 105 Label duplication and incorrect ADS-B data

The threat diagram of Figure 105 shows that the two incidents *Delays in sequence provisioning* and *Degradation of A/C position data* are relevant both before and after the changes. However, the former one may be caused by the threat scenario *ATCO fails to comply with arrival management procedures* only before the changes and by the threat scenario *ATCO fails to comply with AMAN sequence* only after the changes.

Due to the introduction of ADS-B as a means for surveillance, there are also further threats and threat scenarios that are relevant for the unwanted incident *Degradation of A/C position data* after the changes. This is documented by the threat *ADS-B transponder* and the threat scenario *ADS-B transponders not transmitting correct information*.

The before-after threat diagrams resulting from the risk identification after the changes document the risk both before and after, showing risk elements that are persistent, elements that disappear and elements that emerge. When proceeding with the risk estimation, risk evaluation and risk treatment, these diagrams may therefore be used for both before and after the changes.

14.3 Risk Estimation

The risk estimation basically amounts to estimating likelihoods and consequences for the unwanted incidents. Usually, we also estimate likelihoods for threat scenarios in order to get a better basis for estimating the likelihood of unwanted incidents and to understand the most important sources of risks. To the extent that risks before changes are completely unaffected by the changes, the risk estimates need not be conducted twice for these risks.

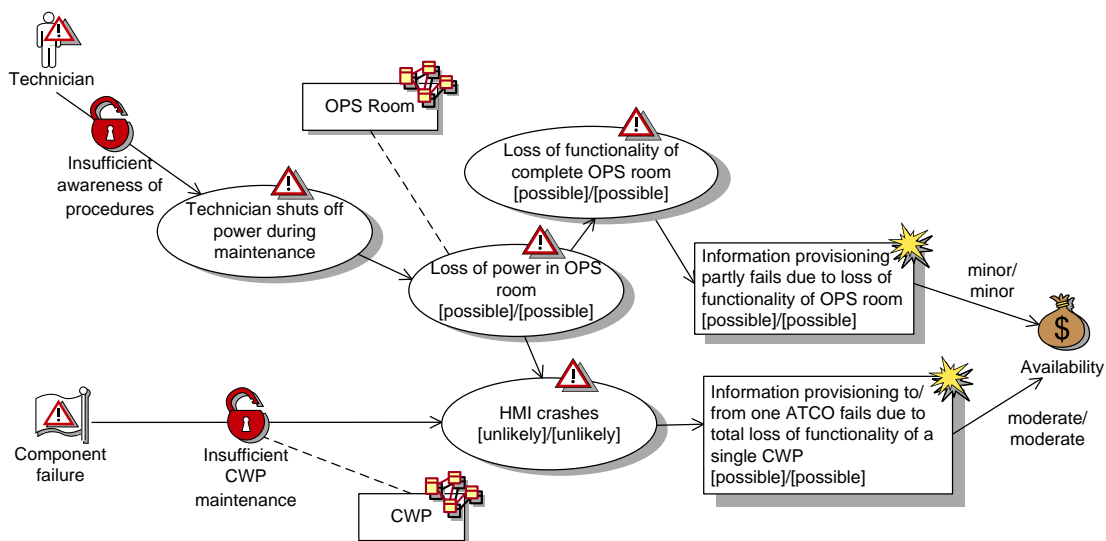


Figure 106 Risk estimation - Loss of functionality

Figure 106 shows that these unwanted incidents, as well as the threats, threat scenarios and vulnerabilities, that are documented in this threat diagram are persistent. The likelihoods and consequences are specified in pairs, where the former value is the estimate before the changes, and the latter value is the estimate after the changes.

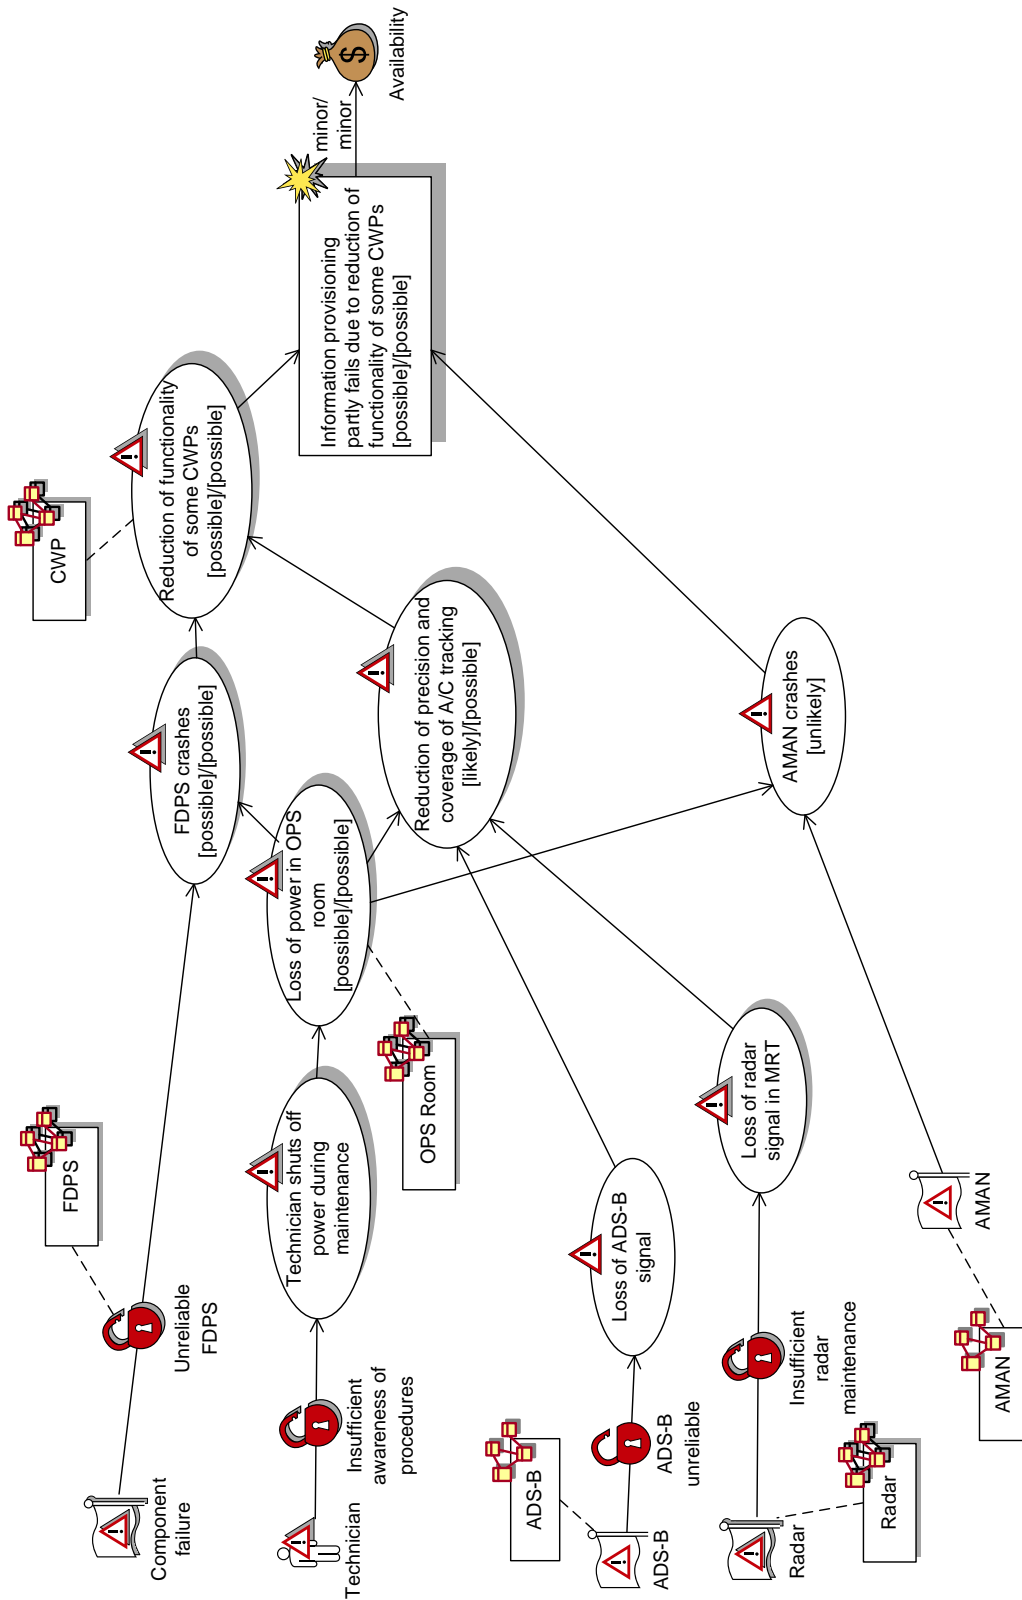


Figure 107 Risk estimation - Reduction of functionality

The threat diagram of Figure 107 gives the likelihood and consequence estimate for the unwanted incident *Information provisioning partly fails due to reduction of functionality of some CWP*s. Both estimates are the same before and after the changes. Notice, however that the introduction of the ADS-B nevertheless affects the likelihood of the threat scenario *Reduction of precision and coverage of A/C tracking*; the estimate changes from *likely* to *possible*. Notice also that the threats and threat scenarios that are related to the ADS-B and to the AMAN are relevant only after the changes. Because the threat scenario *AMAN crashes* may occur only after the changes, it has only one likelihood estimate.

The threat diagram of Figure 108 shows risks that are persistent under the changes. The estimates for before the changes are therefore immediately reused for after the changes. The same is the case for the threat diagrams of Figure 109 and Figure 110.

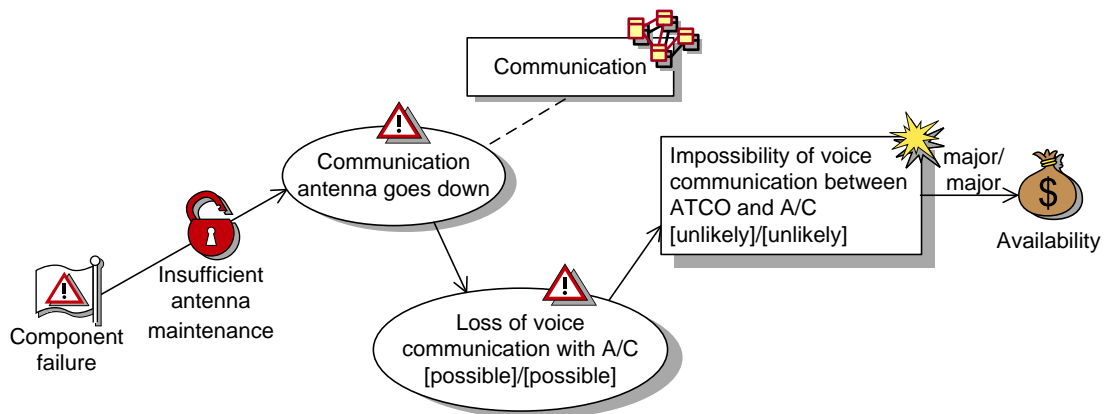


Figure 108 Risk estimation - Loss of voice communication

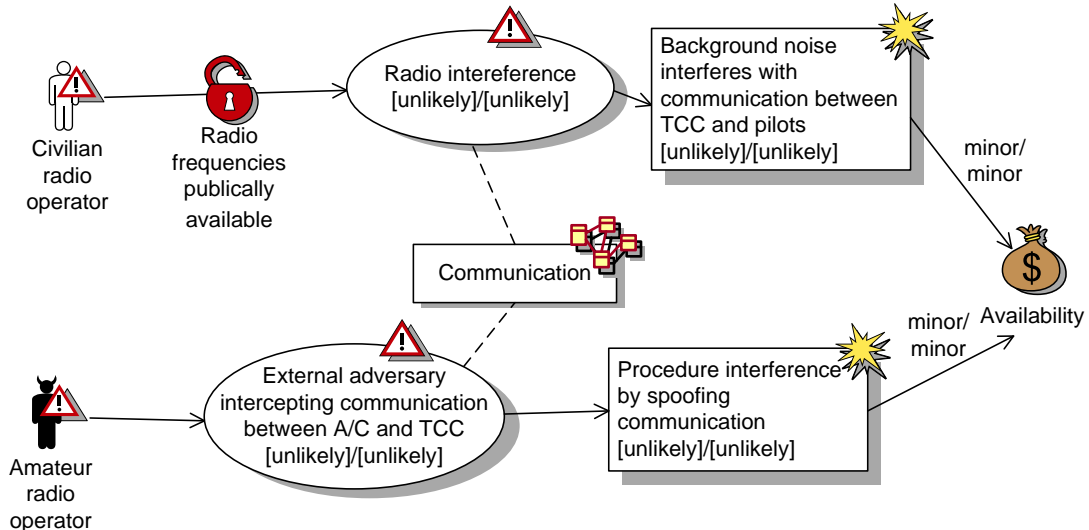


Figure 109 Risk estimation - Radio communication

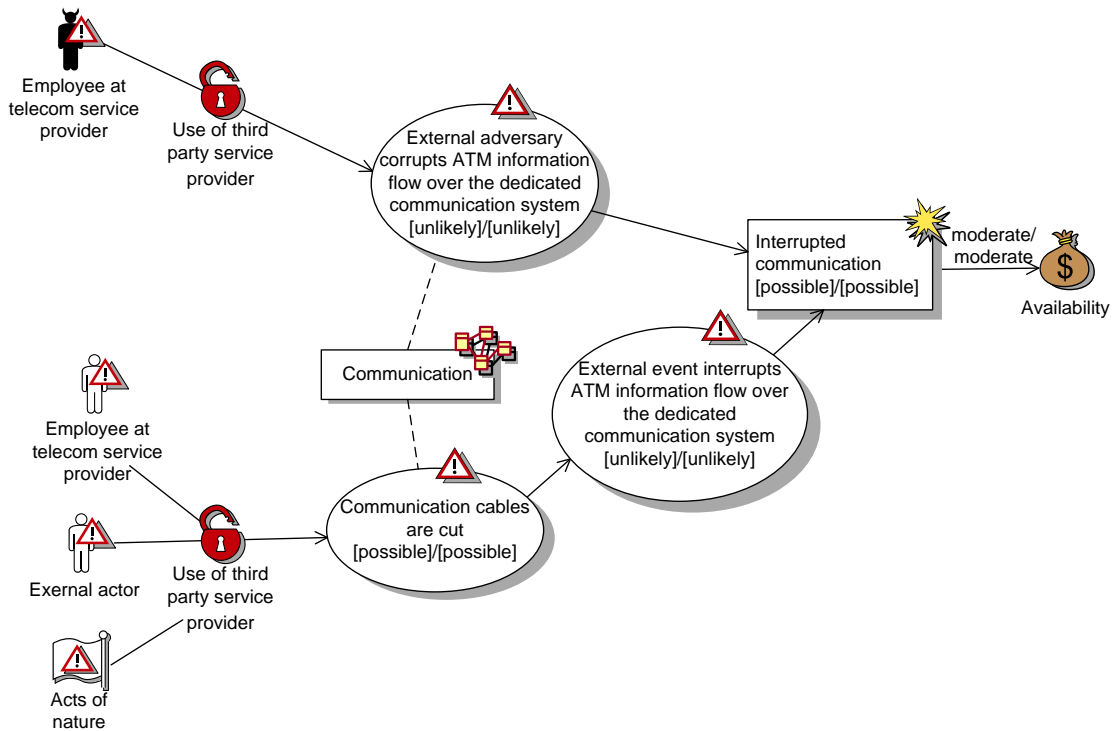


Figure 110 Risk estimation - Telecommunication

The threat diagram of Figure 111 documents issues in relation to human factors, in particular failures to provide arrival management information due to high workload. Noticeably, the threat scenario before described by *Sector team (feeder and approach) becomes overloaded* is split into two separate threat scenarios after the changes. This is because there turned out to be different likelihoods for different instances of this scenario after the changes. Notice furthermore that the likelihood of the unwanted incident *Miscoordination with adjacent ATS units/sectors* is estimated to change from *possible* to *unlikely* under the target changes.

In order to facilitate readability, we have used junction points for many to many relations in this threat diagram.

The threat diagram of Figure 112 documents issues in relation to radar and ADS-B. As shown by the diagram, the ADS-B is relevant only after the changes. The likelihood of one of the unwanted incident furthermore changes under the changes to the target of analysis. The diagram shows confidentiality issues that may arise due to the use of ADS-B.

The threat diagram of Figure 113 documents issues in relation to human factors that are relevant only after the changes. In particular, the diagram addresses the SQM role that is introduced with the introduction of the AMAN.

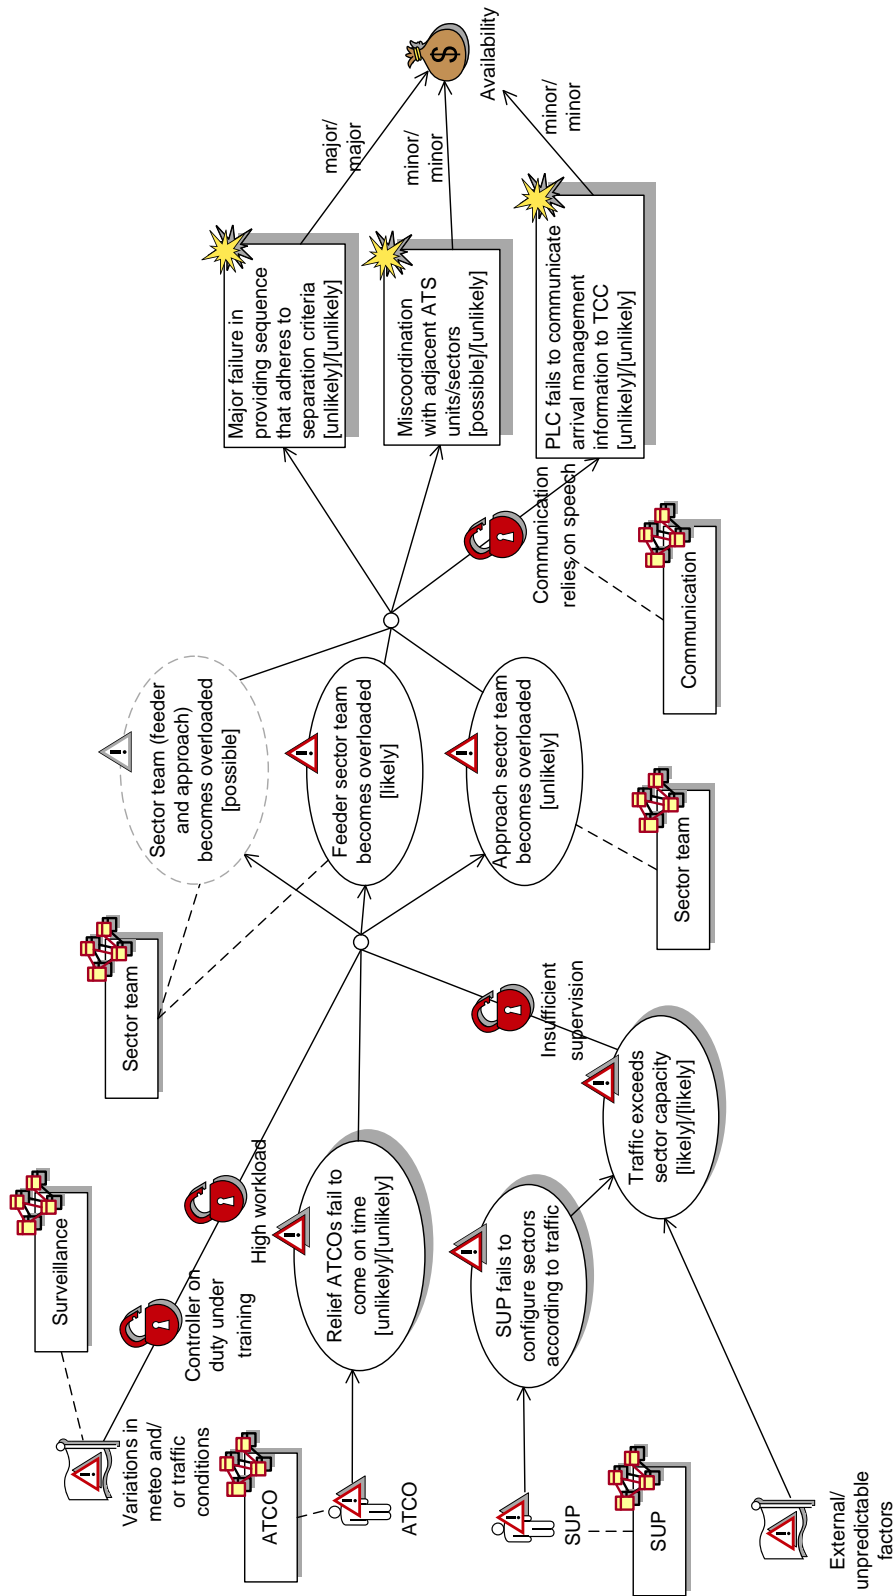


Figure 111 Risk estimation - Human factors

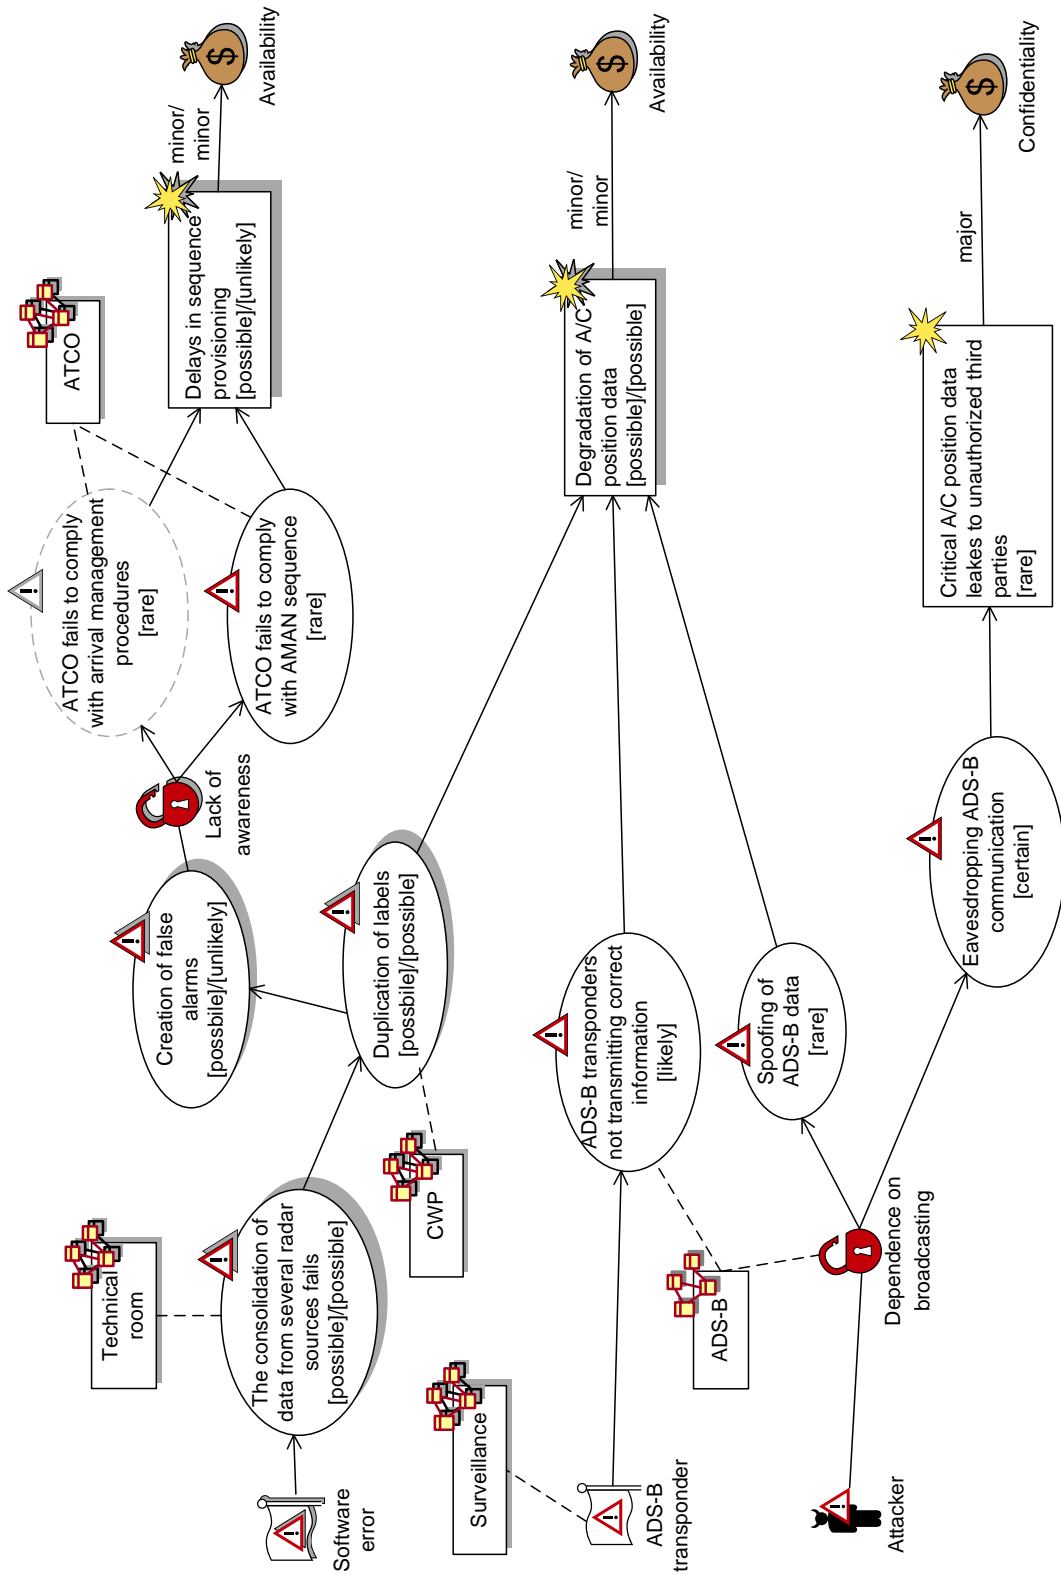


Figure 112 Risk estimation - Radar and ADS-B

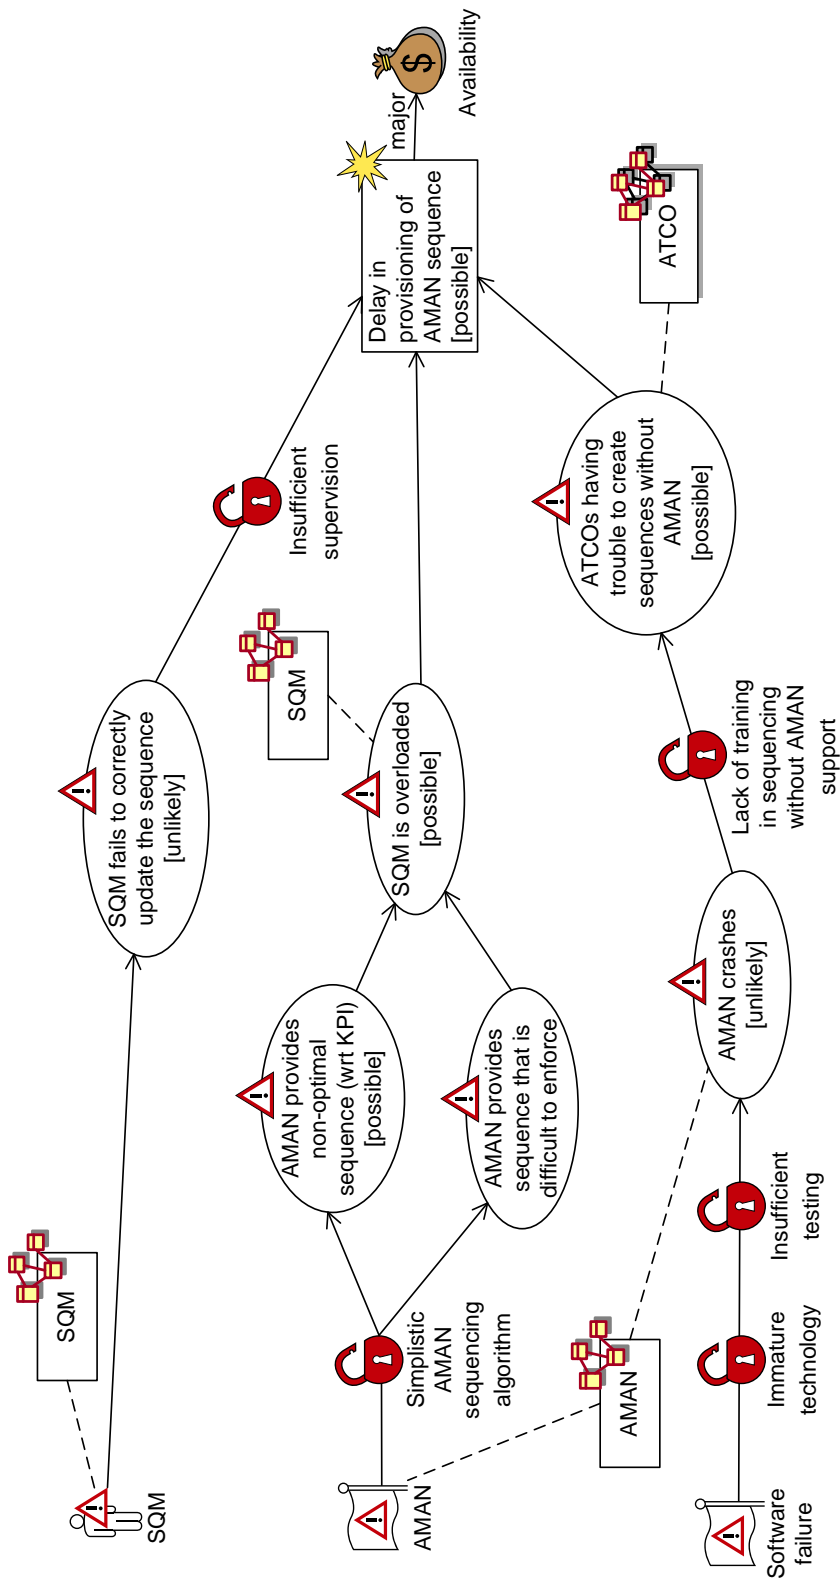


Figure 113 Risk estimation - Human factors after the changes

14.4 Risk Evaluation

During the risk evaluation we first calculate the risk levels by using the risk function defined in Section 14.1.5 and the likelihood and consequence estimates from the previous section. We then compare the risk levels with the risk evaluation criteria to determine which risks that must be treated or evaluated for treatment.

We use CORAS risk diagrams to document the results of calculating the risk levels. These diagrams show the risks together with the threats that initiate them and the assets they harm.

The risk diagram of Figure 114 shows the two risks of loss of CWP functionality both before and after. In order to simplify referring to the risks, they are given unique indices. We see that the risk levels of risk *R1* and *R2* are the same before and after the changes

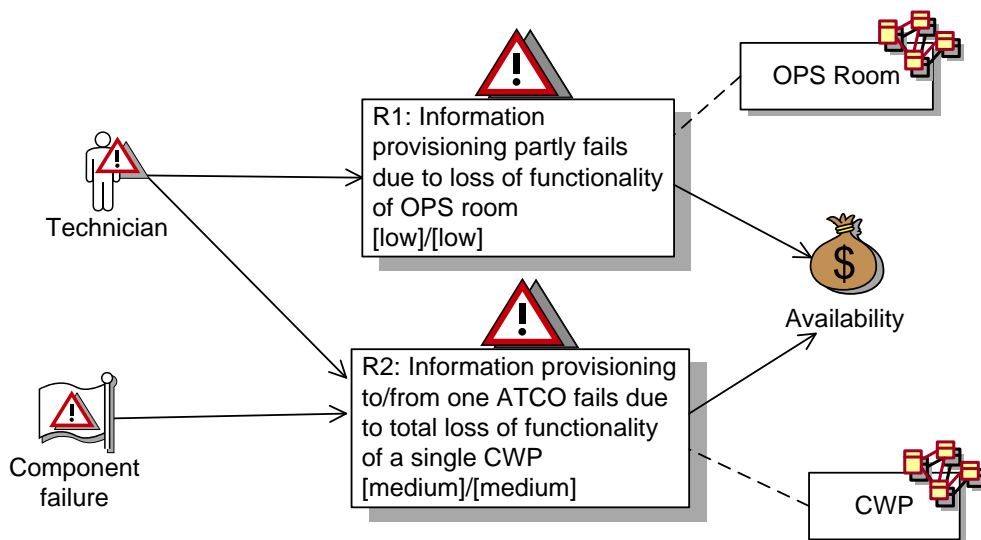


Figure 114 Risk levels - Loss of functionality

The risk diagram of Figure 115 shows that risk *R3* of reduction of functionality remains of level *low* under the changes. Risk *R4* of loss of voice communication shown in Figure 116 remains at level *medium*. Risk *R5* and *R6*, both related to radio communication, remain at level *low* as documented in Figure 117. Risk *R7*, related to telecommunication, remain at level *medium* as documented in Figure 118. The three risks *R8*, *R9* and *R10* as documented in Figure 119 are related to human factors. They remain at the levels *medium*, *low* and *low*, respectively. The three risks *R11*, *R12* and *R13* on radar and ADS-B issues are documented in Figure 120. The former two remain at level *low*. The latter is a risk only after the changes and is of level *medium*. The final risk, indexed *R14*, is shown in Figure 121. It is of level *high* and occurs only after the changes as it is related to the SQM role and to the AMAN.

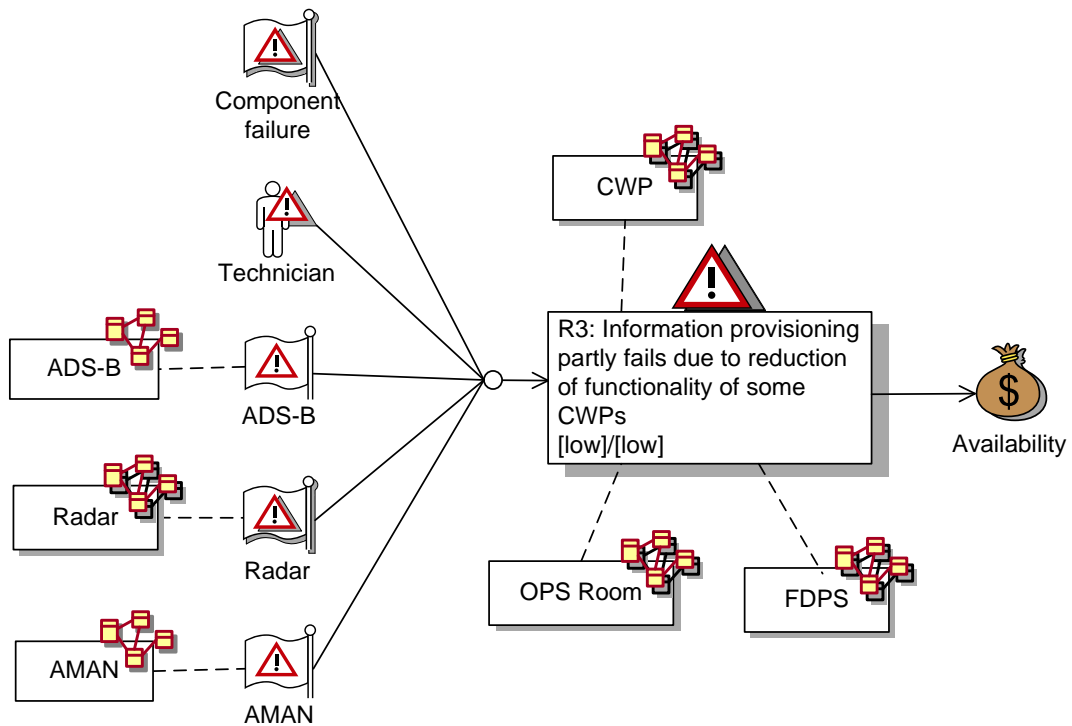


Figure 115 Risk levels - Reduction of functionality

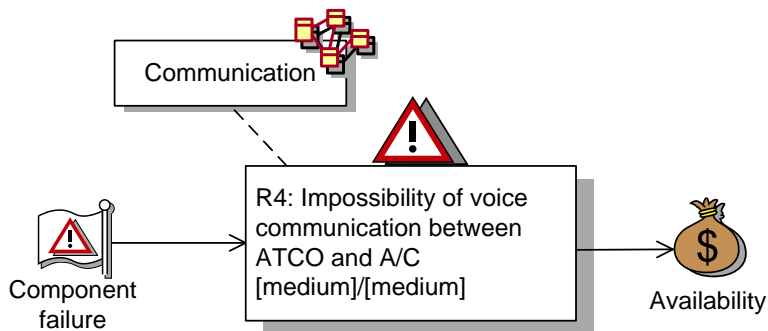


Figure 116 Risk levels - Loss of voice communication

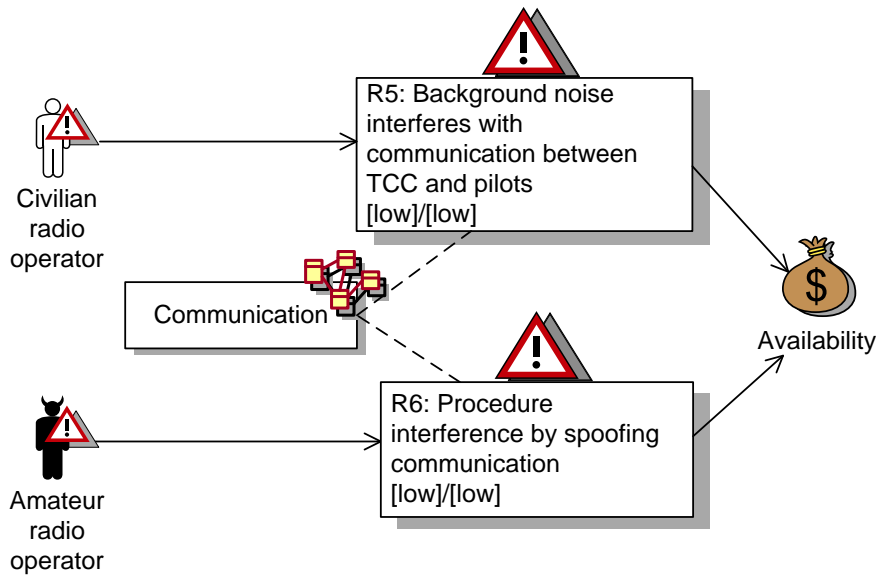


Figure 117 Risk levels - Radio communication

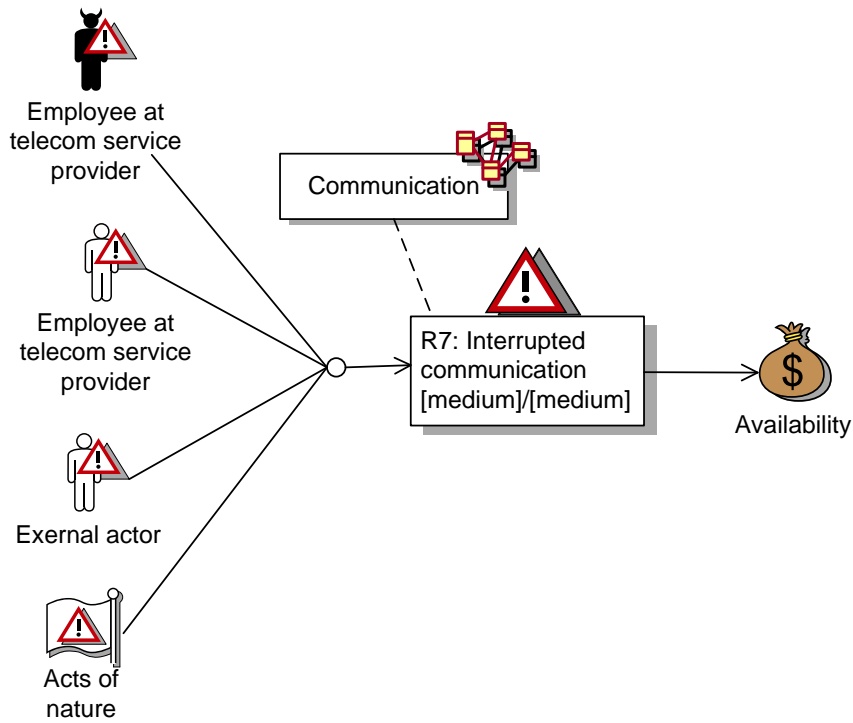


Figure 118 Risk levels - Telecommunication

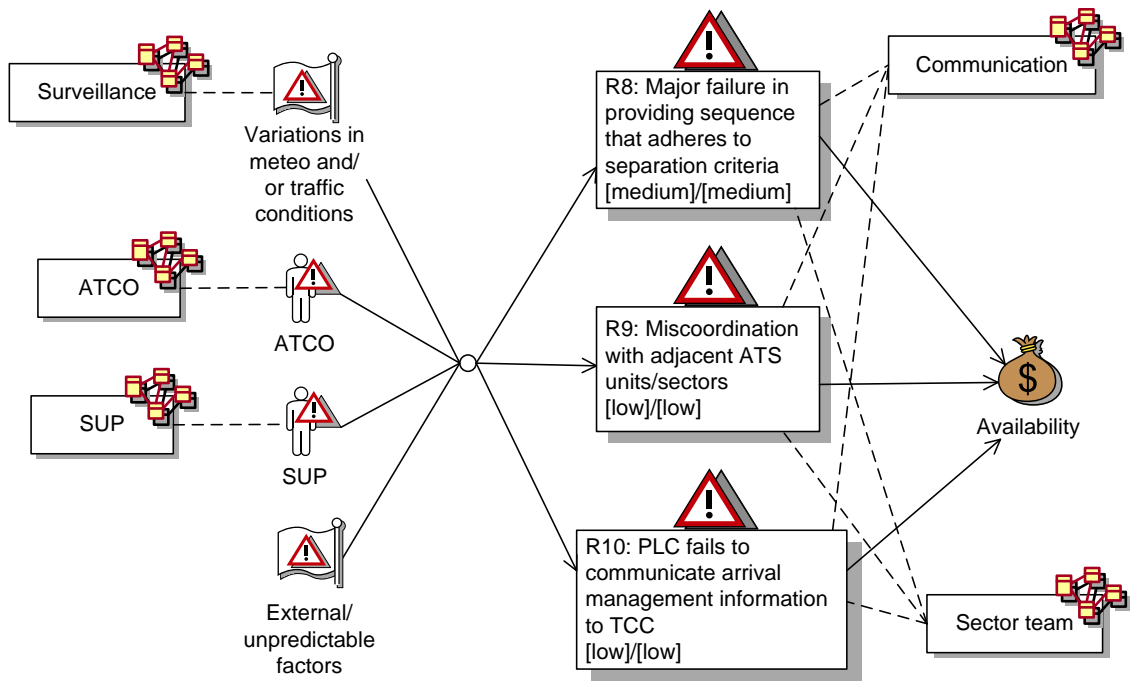


Figure 119 Risk levels – Human factors

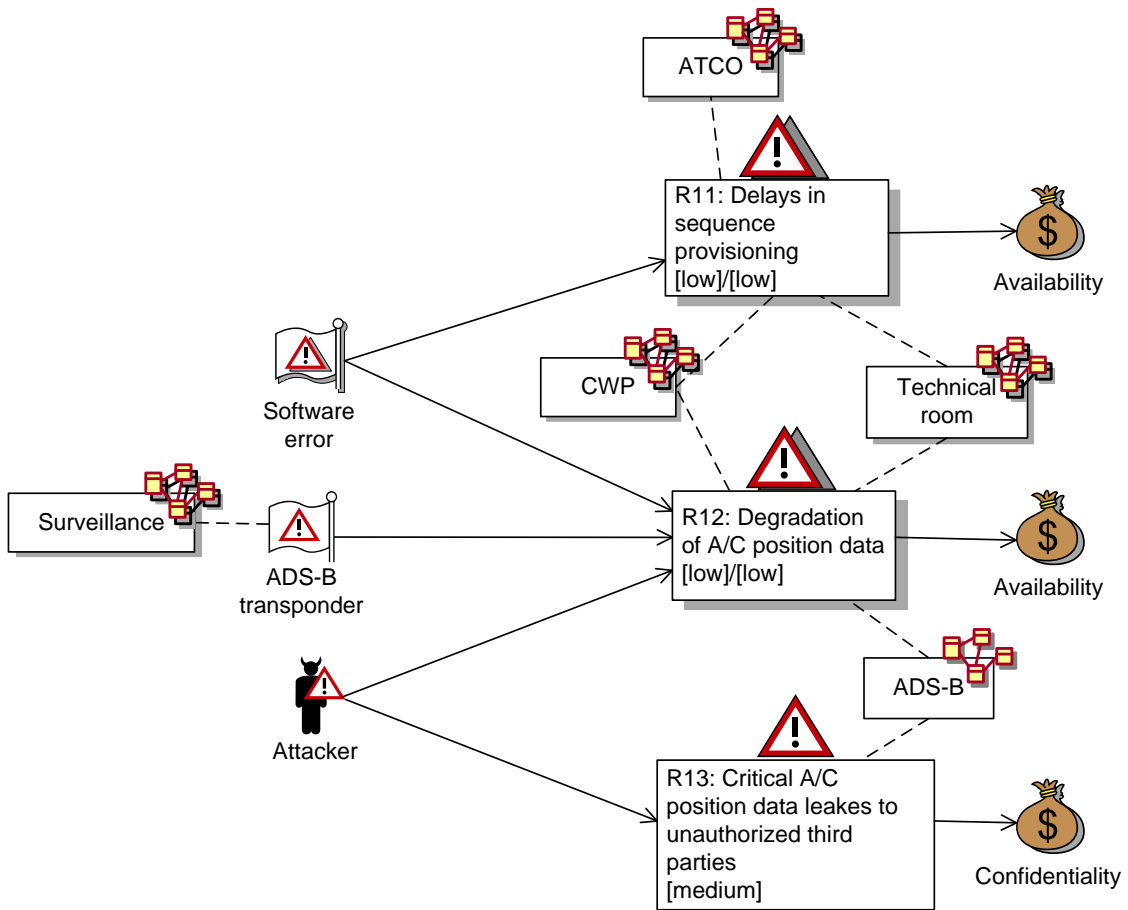


Figure 120 Risk levels – Radar and ADS-B

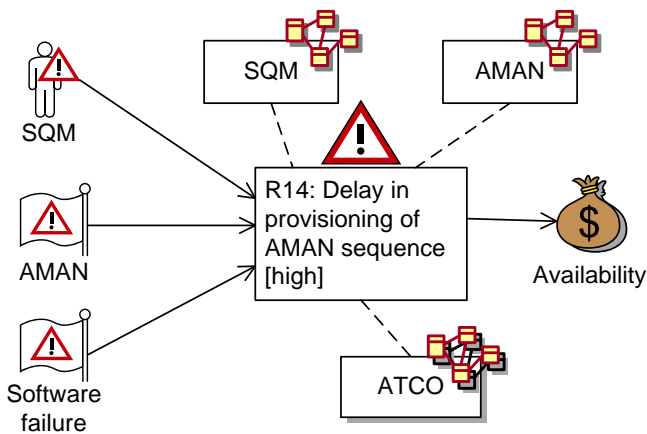


Figure 121 Risk levels – Human factors after the changes

We use the risk matrix to provide an overview of the results of the risk estimation and risk evaluation. We present the risk evaluation matrix for the risks before and after separately.

14.4.1 Risk Evaluation before Changes

The twelve identified risks before the changes are plotted into the risk matrix of Table 36. There are no risks that are estimated as *high*, but four of them are estimated as *medium*, and therefore need to be evaluated for possible treatment. These are risks *R2*, *R4*, *R7* and *R8*.

		Consequence				
		Insignificant	Minor	Moderate	Major	Catastrophic
Likelihood	Rare					
	Unlikely		R5 R6 R10		R4 R8	
	Possible		R1 R3 R9 R11 R12	R2 R7		
	Likely					
	Certain					

Table 36 Risk evaluation before changes

Whether or not to recommend the risks identified for the target of analysis before the changes may depend on whether these risks also occur after the changes and on the timeframe for implementing the change requirements. For short timeframes, it may not be reasonable to invest in risk mitigation for risks that nevertheless will become less severe or even obsolete.

14.4.2 Risk Evaluation after Changes

The fourteen identified risks after the changes are plotted into the risk matrix of Table 37. There is one risk that is estimated as *high* for which risk treatments must be identified, namely *R14*. The risks that are of level *medium* before the changes are still of the same level. In addition, risk *R13* is of level *medium*.

		Consequence				
		Insignificant	Minor	Moderate	Major	Catastrophic
Likelihood	Rare				R13	
	Unlikely		R5 R6 R9 R10 R11		R4 R8	
	Possible		R1 R3 R12	R2 R7	R14	
	Likely					
	Certain					

Table 37 Risk evaluation after changes

Risks *R13* and *R14* are risks that emerge after the changes. In the risk matrix this are written in bold face to highlight this. Risks *R9* and *R11* get reduced likelihoods after the changes, but remain with the same risk level of *low*. In the matrix, they are written in italics to highlight the slight change.

14.5 Risk Treatment

The risk treatment was conducted as a structured brainstorming following the risk estimation and risk evaluation. Some treatment options were identified, and are documented using CORAS treatment diagrams.

In order to come up with a recommended treatment plan, a more thorough treatment identification and treatment evaluation need to be conducted. This was outside the scope of the case study, and we therefore only document the treatment suggestions that came up during the brainstorming.

The identified treatments focus mainly on the risk picture after the changes, addressing risks that may arise as a consequence of implementing the process level changes in the arrival management.

The treatment diagram of Figure 122 shows a treatment option for mitigating the risk of interrupted telecommunication. The treatment diagram of Figure 123 addresses human factors, and identifies as a treatment option the increase of relief periods to decrease the likelihood of overloaded ATCOs. The treatment diagram of Figure 124 addresses risks related to the ADS-B. One of the treatment options, namely ADS-B encryption, can ensure both confidentiality and authentication. The treatment diagram of Figure 125 finally addresses human factors after the changes. After the introduction of the AMAN, it may be important that ATCOs maintain competence in ATM without tool support such as the AMAN, in case of tool crash. At the same time, the ATCOs need sufficient competence in using the AMAN to ensure that appropriate sequences are still built.

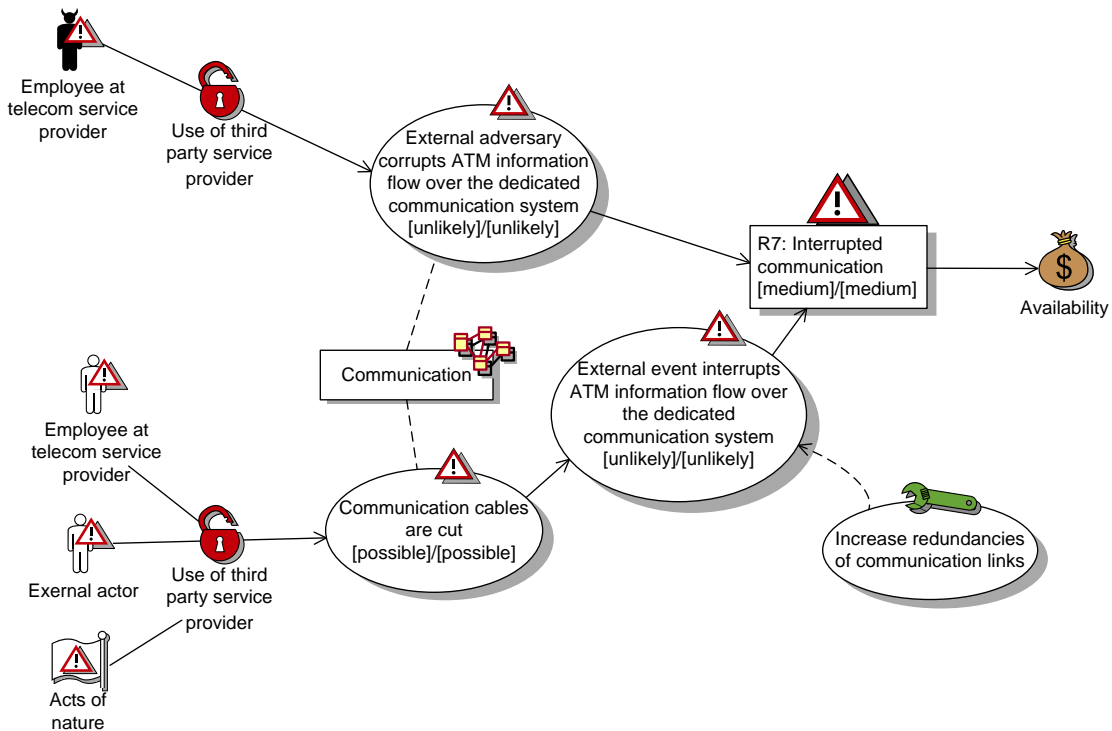


Figure 122 Risk treatment - Telecommunication

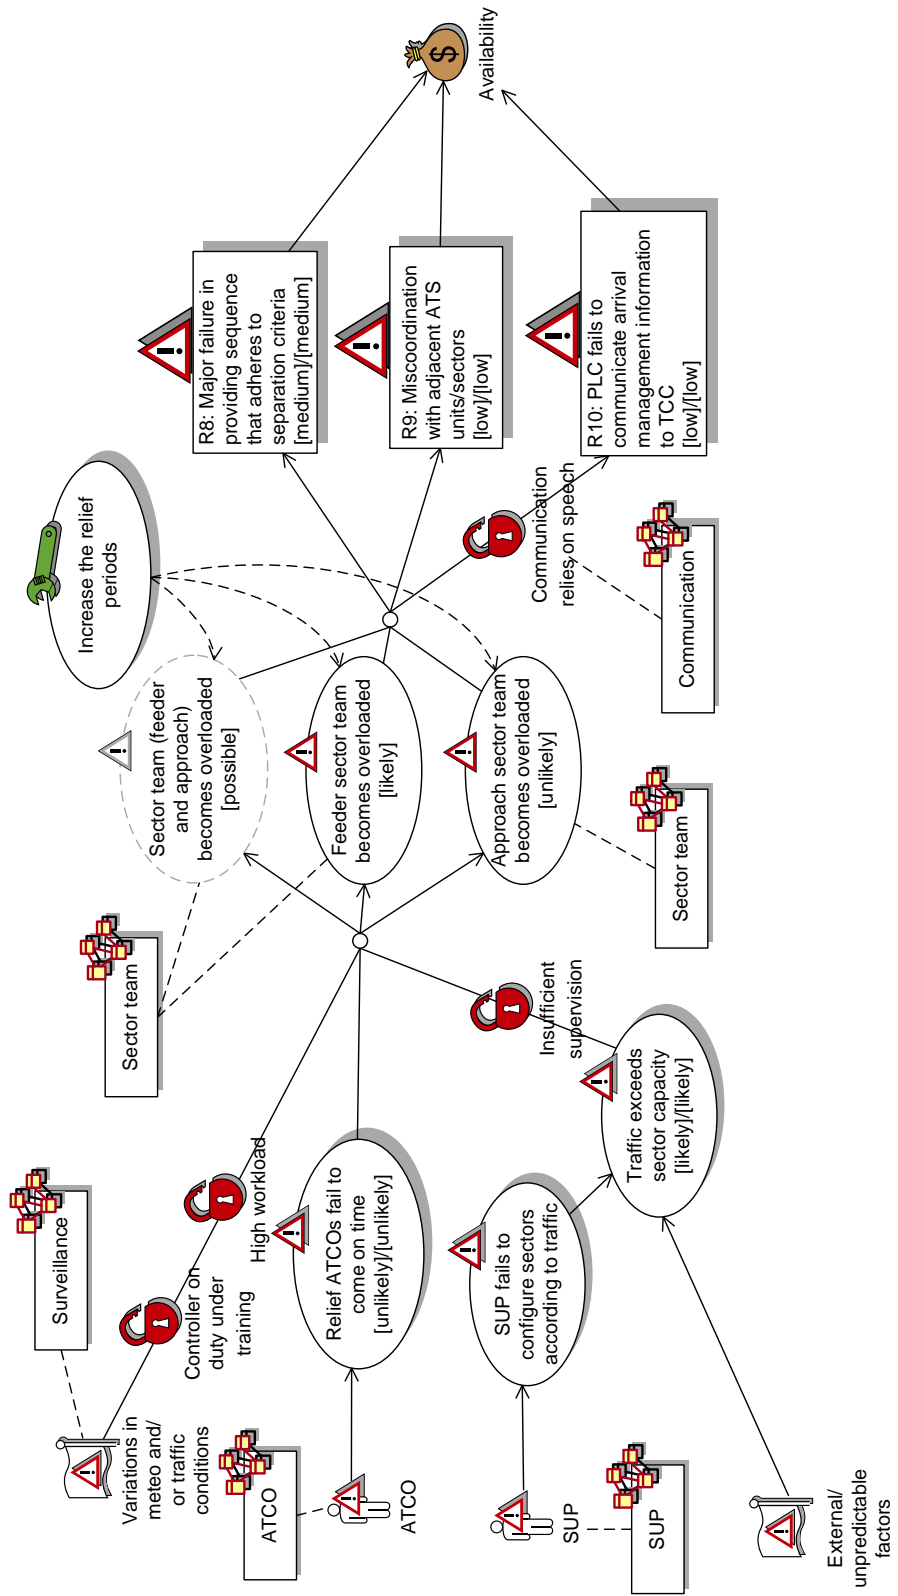


Figure 123 Risk treatment - Human factors

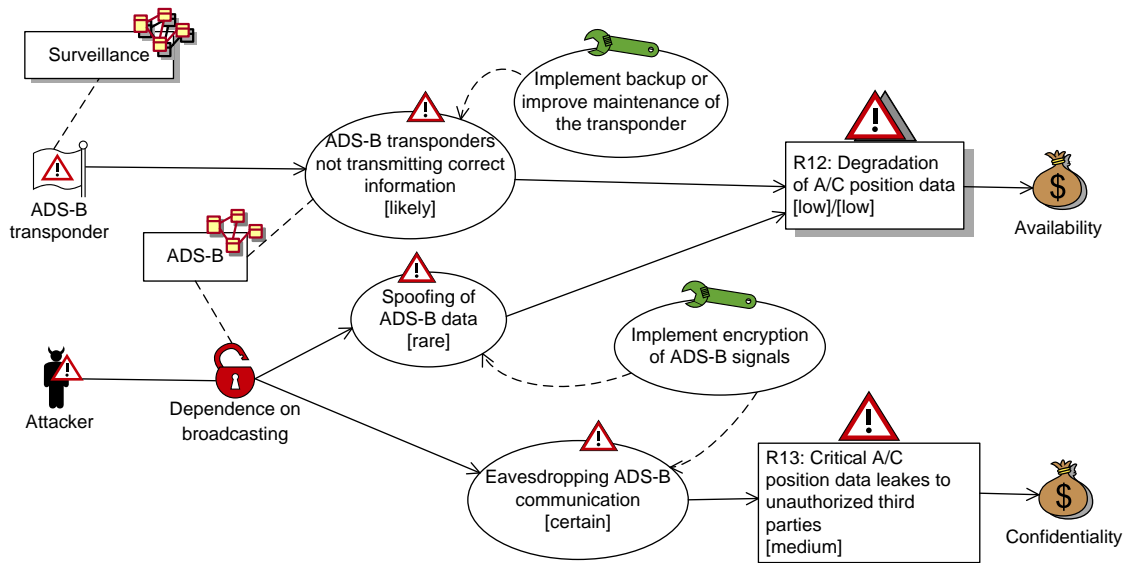


Figure 124 Risk treatment - ADS-B

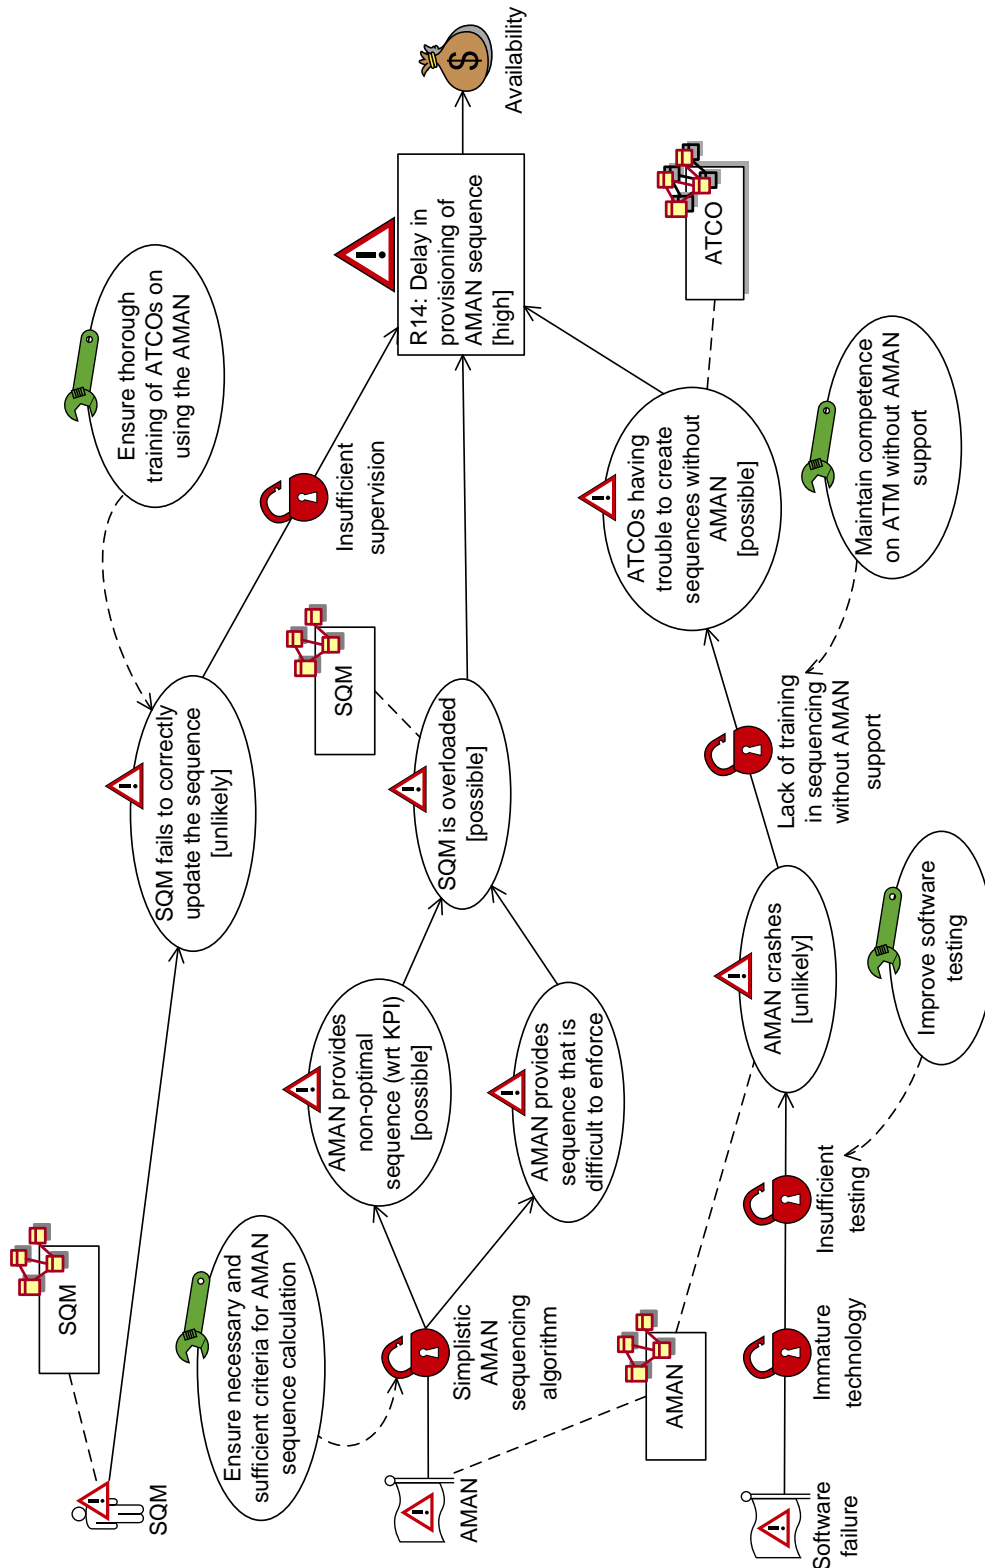


Figure 125 Risk treatment - Human factors after changes

15 E – Report on HOMES Case Study with CORAS

In this section the application of the CORAS risk assessment methodology on the HOMES case study will be described. HOMES is focused on digital home networks where some sensible changes take place from the point of view of the security.

After a general description and introduction to the HOMES case study, the change requirements and the security properties that will be addressed by applying WP5 methods and techniques will be presented.

The particular risk assessment solutions for changing systems that are applied to the HOMES case study are:

- Risk assessment under the maintenance perspective.
- Risk assessment under the before-after perspective.

In the application of the risk assessment methodologies on the HOMES case study no actual risk evaluation has been conducted, only risk identification. The risk identification has been checked with a domain expert. Conducting a proper risk evaluation requires meetings with domain experts and risk evaluation workshops. Due to pragmatic reasons and time constraints, this has been left out of the HOMES case study reported in this section.

The section is structured as follows: After a short description of the context of the HOMES case study, the change requirement that is addressed in the case study is presented. Thereafter, the security properties dealt with in the risk assessments are described, and finally the results of the risk assessment are presented and explained.

15.1 Context Establishment

The general environment is a home network wherein any connecting device shall be assessed by the Operator, following a Network Access Control (NAC) approach. Once the device is accepted we may consider the following interactions:

The Customer shall access an Operator's service store, (which is indeed a service installed in customer's Homes Gateway (HG)) and select any home service from the catalogue. Services are offered by Third party Service Providers. Once the customer selects a service, it is redirected to the proper Service Provider (SP) to proceed with the purchase. The SP shall deliver the service to the customer's home gateway once the customer accomplishes the payment. The delivered service shall be deployed as a Web Service client able to access the third party SP (Change Requirement: Bundle lifecycle operations).






Figure 126 General environment

15.1.1 Business Needs

Connection Level

Operator requires safe devices connected to the network. This is a prerequisite to any further interaction. A key part of the NAC implementation lies on the Home Gateway. Operator needs to keep updated the core security modules implementing those security functionalities (CR: Core Security Module Updates). Therefore it is critical to control the update procedures of those software modules. The update shall not affect any existing service running in the platform so Operator needs tools to check possible impact of the update of the security module before deploying it in the production environment to all the customers.

Service level

Operator requires a certain level of quality from the services offered by the SP. By default, once the Operator and the SP sign a commercial agreement, Operator trusts the SP and its services. This trust is translated into a basic level of control over the SP and its services, i.e. Operator does not impose strict constraints to the services. Nevertheless, this trust might **degrade** with the pass of time (Security property: Resilience to trust changes, Security property: Secure extensibility). Operator shall degrade the trust on a certain SP because of several reasons:

1. Reports on bad quality of the offered services: some SP may receive a noticeable amount of complaints from customers about malfunctions or low quality of the service, etc.
2. Critical bugs into the services or even malware.

3. Non delivery of services.

The trust degradation shall drive to the imposition of severe constraints to that SP in the form of strict security requirements that mitigates the mentioned threats:

1. **delivery of certified bundles only:** due to the new trust relationship between operator and third party service provider, the operator requests that only certified bundles of this operator may be deployed on the home network. (Security property: Policy enforcement.)
2. **deployment of a new security service:** due to the new trust relationship between operator, third party service provider, and customer, the operator requests that a non-repudiation protocol may be run between the parties to prevent denial of having subscribed, received or delivered a service. (Security property: Security expandability.)

If the SP continues presenting problems Operator could decide to ban that SP.

15.1.2 Actors

The following actors take part in the business case:

Customer / Service Requestor: the customer at his/her home installing and using services for his/her home.

Home Gateway, HG: device placed into the customer premises. Owned by the Operator and normally rented to the customer. This device acts as a service platform for the home.

Operator: broadband network provider and owner of the infrastructure connecting the home to it.

Third party Service Provider, SP: remote service provider, independent from the Operator but with a commercial agreement with it, offering home services to customers of the Operator.

15.2 Change Requirements and Security Properties

This section presents the change requirement and the security properties that are addressed from the HOMES case study.

15.2.1 Change Requirements

The change requirement on which the WP5 risk assessment methodologies are applied is "Bundle lifecycle operations".

Bundle lifecycle operations

A Home Gateway is a service platform for the home. Customers can install new home services, upgrade or delete existing ones. This type of change is similar to the previous one but here services do not usually implement security functionality. The bundles installed on the home gateway are used for higher level applications. The services may come from third parties and therefore some similar control over this software must



exist. Trust relationships among the customer, the service provider, and the third parties may evolve over time. However in some cases security bundles could be deployed (provided by the operator).

15.2.2 Security Properties

The security properties addressed by the technical solutions are:

Policy enforcement. The Policy Decision Point (PDP) is located in the security domain of the operator. The Policy Enforcement Point (PEP) is a core security module installed on the home gateway. The PEP always enforces policy decisions forwarded by the PDP so that only allowed actions can be carried out.

Security expandability. System security can be enhanced by taking advantage of the home gateway extension ability (mentioned in the Secure Extensibility property) through the deployment of new security services (e.g., deployment of a non-repudiation service bundle to ensure that neither service provider nor customer can later deny having sent/received a purchased service). The infrastructure shall be able to efficiently enforce such new requirements with a minimal impact on it.

15.3 Timeline for the HOMES case study

In this section we describe the timeline that provide the frame for the application of the risk assessment methods and techniques. The timeline is divided in three segments as depicted in Figure 127, each with its specific application of risk assessment methodologies to the aforementioned change requirement and security properties listed. The three segments correspond to the following points in time:

1. At this point in time a simple Home Gateway is already deployed and working successfully. The simple Home Gateway is already SeAAS capable and equipped with the simple security service “Confidentiality Service”. The Home Gateway is analyzed, tested and running without security problems. We have a system model and a risk model for that point in time.
2. At a certain point in time, the Operator notes increasing customer and third party service provider complaints. The Operator (depicted as User in the change story described in D2.2) orders the risk analysis team to update the existing risk analysis and find reasons and causes for the increasing number of complaints. The risk analysis team conducts a risk assessment from a maintenance perspective and provides an updated risk picture, including newly identified threat scenarios and proposed treatments.
3. The treatment is accepted as an actual change request to the system and therefore analyzed from a before-after perspective by the risk analysis team. The resulting risk models depict the risk before the application of the treatment and the risk after the application of the treatment. In addition the risk to the change transaction itself is analyzed.




Figure 127 Timeline of applying risk methodologies to the HOMES case study

In the following sections the single steps of the timeline is described in more detail, outlining the various artifacts that are produced.

15.4 T1: Risk Identification before Change

At time T1 a Home Gateway is deployed and running. A risk model of the system at this state is also available from a risk assessment that has been conducted for the target of analysis at this point in time. In the following the system model and the risk model are briefly described.

15.4.1 System Model

Figure 128 depicts the HOMES system model and its components at the beginning of the change story. The system is spread over three organizational boundaries, the Operator, the Third Party Service Provider and the HOMES Gateway itself, which is installed at the customer.

On the site of the Operator a Policy decision point – the PDP service – is running, which communicates with the Policy enforcement point – the SeAAS engine – on the HOMES gateway.

At the site of the Third Party Service Provider the component Feed Service is deployed, which provides content to the customer via the HOMES gateway.

The HOMES gateway itself is based on the Open Services Gateway initiative (OSGi) framework and contains various bundles. In particular there is the SeAAS Engine acting as a Policy enforcement point. To access the services of the Third Party Service Provider a Feed Server component is also deployed on the HOMES gateway. At this point in time there is only one security service deployed, the Confidentiality Service, which requires an additional component, namely the Cryptographic Resources used for managing the cryptographic keys.

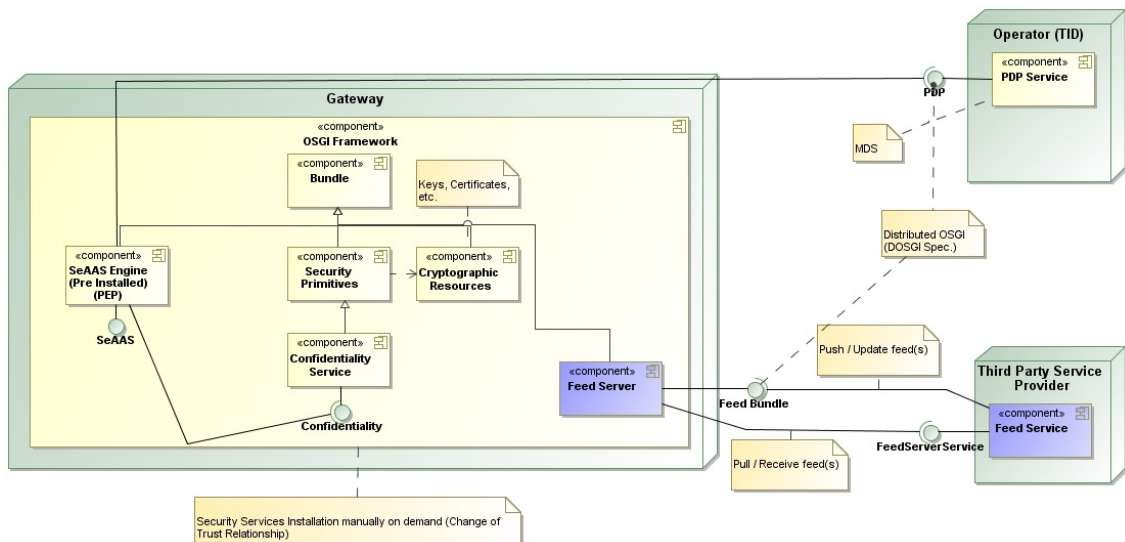


Figure 128 HOMES system model before the change

15.4.2 Risk Model

Figure 129 depicts a CORAS threat diagram describing parts of the risk picture in relation to the HOMES gateway at the beginning of the change story. At this point in time there was one potential unwanted incident documented. The unwanted incident represents a risk with respect to the asset *Integrity of security components' functionalities* which is related to the security property of Security expandability.




Figure 129 CORAS threat diagram for the HOMES system at T1

The unwanted incident which was identified is the incident of *Security services not functioning*. This unwanted incident can be caused by a threat scenario in which the PDP is not responding. This can have various reasons: On the one hand the PDP service component could fail in production use due to insufficient testing. On the other hand the PDP could stop responding in the case of a loss of connection to the PDP service. This again can be caused by a classical network failure since the whole setup depends on a working connection.

15.5 T2: Risk Identification under the Maintenance Perspective

At T2 the Operator as the primary stakeholder of the system orders a new risk analysis to analyze and understand the reasons for the increasing complaints. In the following, the business case for change is shortly discussed and the artifacts produced by the risk analysis team are presented.

15.5.1 Rationale for a New Risk Analysis under the Maintenance Perspective

The number of complaints at the operator as the primary stakeholder of the system has lately increased. The complaints were stemming from Third Party Service Providers and from Customers. Most of the complaints relate to accusations of violating the Service Store Sales Policy compliance. The Service Store Sales Policy compliance relates directly to the security property of Policy enforcement.

The Operator is providing a platform by which Third Party Service Providers can offer and deploy services on customer request. The increasing number of complaints is a direct threat to this business model and potentially undermines the reputation of the Operator on the market. Therefore the Operator issues an order to update the current risk analysis as partially documented in Figure 129.

Since there has not been any change to the HOMES gateway as a running system, there is no update of the system model necessary and the risk analysis team can directly identify new potential threats and risks based on the reporting of the new scenarios that have emerged.

The method for risk assessment under the maintenance perspective aims at updating and restoring the validity of a risk assessment that has been conducted earlier. Because the changes reported at T2 have already occurred, the maintenance perspective is the appropriate perspective in this scenario.

15.5.2 Risk Model – Maintenance Perspective

The risk analysis team conducts a risk assessment under the maintenance perspective to identify potential threat scenarios which are related to unwanted incidents resulting in breaches of compliance of the Service Store Sales Policy. The risk analysts furthermore need to determine whether the previously identified and documented risks are affected.

The threat diagram of Figure 130 documents a set of new potential threats that has been identified. The threats relate to the increasing number of complaints stemming from customers and third party service providers. The risk analysts furthermore determines that the risk documented by the threat diagram of Figure 129 is still valid

The new identified risks are represented by three unwanted incidents:

- Third Party Service Providers delivers service without customer consent in violation of Sales Policy.

- Third Party Service Provider requests payment from customer for undelivered service in violation of Sales Policy.
- Customer violation of Sales Policy.

These three unwanted incidents are possible because of different threat scenarios. First of all a malicious Third Party Service Provider might deploy unrequested services in addition to a legitimately purchased service. An example might be the deployment of an Ad-Service in addition to a Feed-Service. This is a violation of the Service Store Sales Policy.

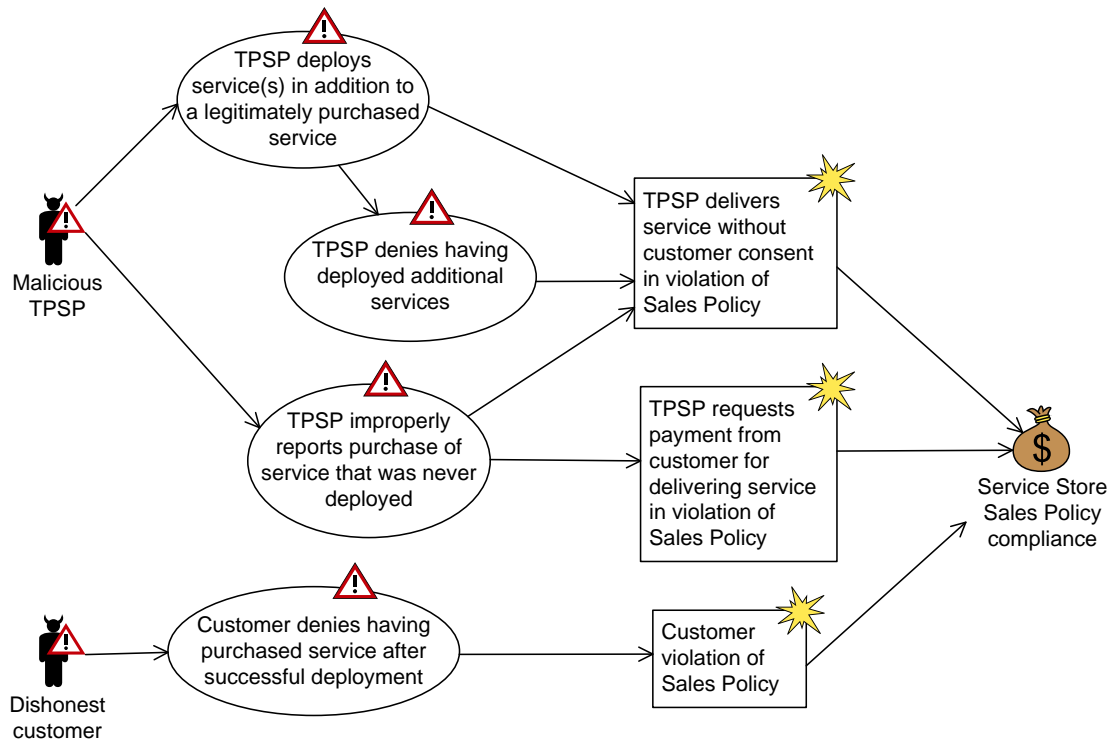


Figure 130 CORAS threat diagram for the HOMES system at T2

Directly related to this threat scenario is a scenario in which the Third Party Service Provider subsequently denies having deployed these additional service.

Another threat scenario related to a malicious Third Party Service Provider is the situation in which the purchase of a service that was never deployed is improperly reported and payment requested.

There are also scenarios of malicious customers which order a service which is successfully deployed, but then the customers deny having purchased the service at all.

15.5.3 Treatment Identification

The risk analysts have also identified a treatment which addresses most of the identified threat scenarios, namely the deployment of a non-repudiation security service on the SeAAS capable HOMES gateway.

Using this treatment, three threat scenarios are addressed, because neither a Customer nor the Third Party Service Provider can anymore deny having ordered, deployed and purchased a service. This particular treatment is documented in the treatment diagram of Figure 131, and may mitigate each of the three new risks that were identified.

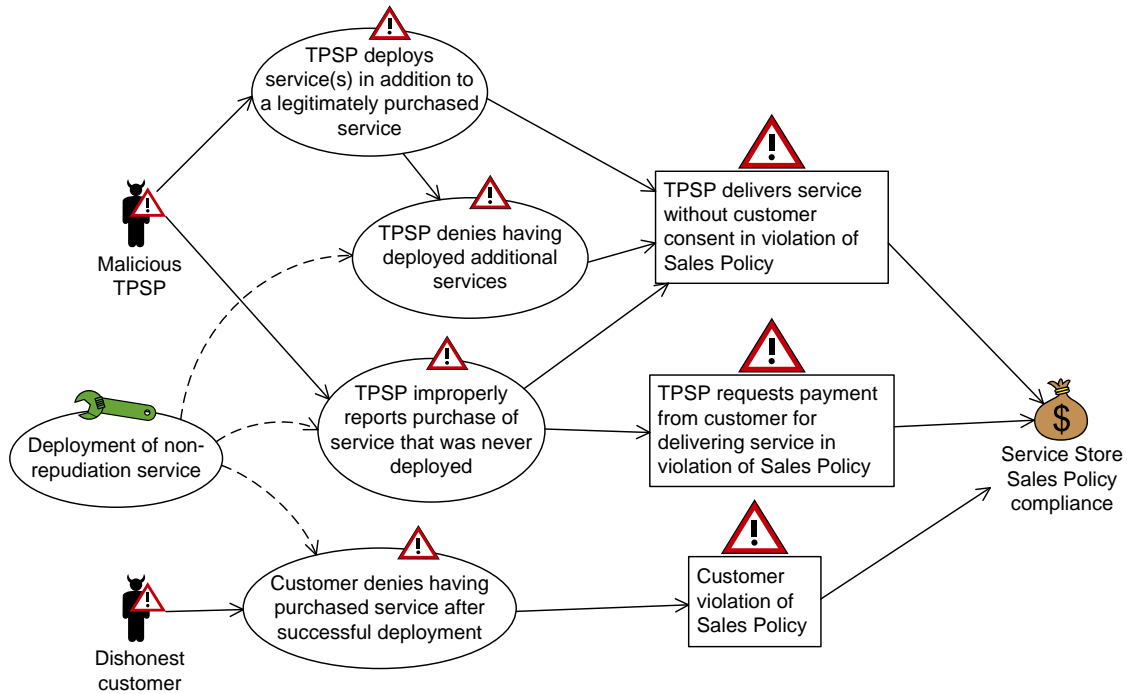


Figure 131 CORAS treatment diagram for the HOMES system at T2

15.6 T3: Risk Identification under the Before-After Perspective

The treatment proposed by the risk analysis team is accepted by the Operator who orders an analysis of the impact of the changes related to the application of the treatment from a before-after perspective. The risk assessment is to be conducted before this planned change so as to predict the possible changes to the risk picture.

As the first step of the before-after risk analysis an updated and changed version of the system model is produced. The system model contains the changes related to the deployment of the proposed treatment. As the next step the risk model after the changes is described. This includes new potential threats related to the treatment and the disappearance of threats which were present before but are not relevant after the application of the treatment.

15.6.1 System Model after the Change

Since the application of the treatment is considered a change request, the risk analysis team conducts a risk analysis under the before-after perspective. After the

implementation of the treatment, i.e. after the change transaction, the system will be changed as outlined in Figure 132.

The deployment of the non-repudiation security service on the HOMES gateway requires changes on the Operator site and on the Gateway itself. On the Operator side a NRP-TTP service is deployed. On the Gateway the Non-Repudiation security service is deployed which requires as a dependency two additional services, namely an Integrity Service and a Timestamp Service. The components which are added to the system are highlighted in yellow in the system model in Figure 132.

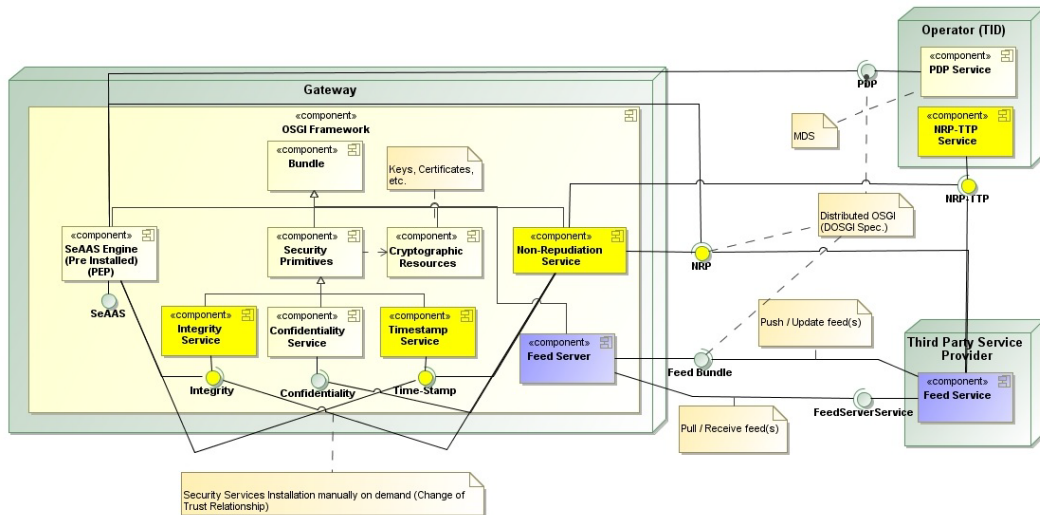


Figure 132 HOMES system model after the change

15.6.2 Risk Model – Before-After Perspective

Following the risk assessment method for the before-after perspective, the risk analysis team identifies new potential threats which are related to the deployment of the non-repudiation security service and protocol as a treatment. At the same time, the risk analysts need to determine which of the previously identified risks are persistent under the changes, and which of them that disappears.

The before-after threat diagram of Figure 133 builds on the threat diagram of Figure 129 from before the changes, and documents one new unwanted incident. With the decision of the Operator to base the overall Service store purchase protocol on a non-repudiation protocol there is a potential risk related to the Availability of the Service Store. If for any reason the non-repudiation service is not responding, then the purchase protocol cannot be executed anymore. This in result leads to a non-working Service Store on which no purchases can be made.

A particular threat which could lead to the threat scenario of a non-functioning non-repudiation security service is an attacker initiating a denial of service attack on the NRP-TTP service located at the Operator site.

The before-after risk modeling distinguishes between risks that are present only before changes, risk that are present only after changes, and risk that are present both before and after changes. In Figure 133 we explicitly see that the attacker and the related

scenarios and incident emerge after the changes, whereas the remaining elements are persistent under change.

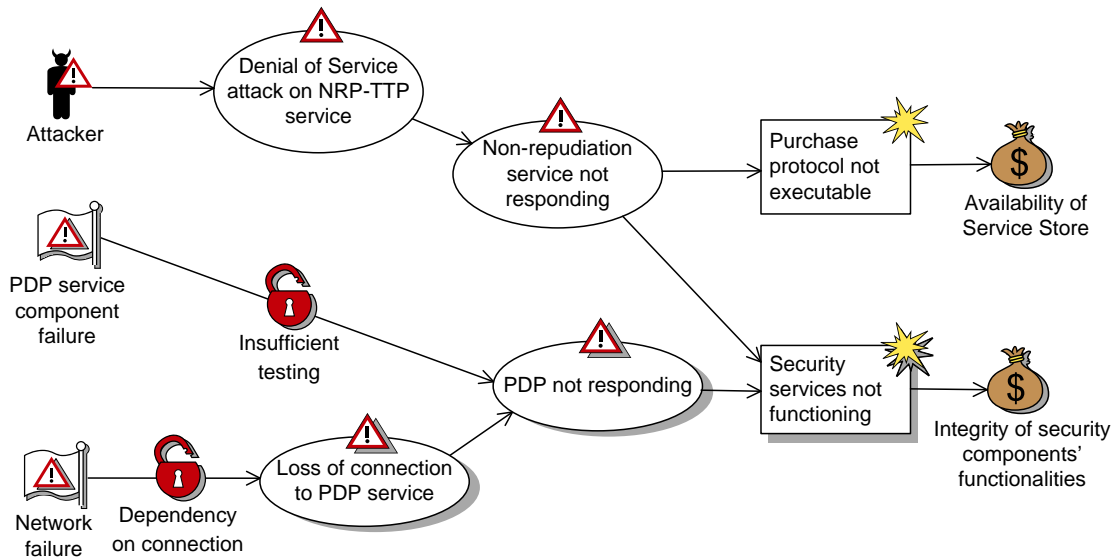


Figure 133 CORAS threat diagram for the HOMES system at T3

If the treatment is deployed and working as expected a series of scenarios will not occur anymore. In the before-after risk models, the elements that disappear are grayed out, so as to explicitly show how the risk picture changes. The before-after threat diagram of Figure 134 builds on the threat diagram of Figure 130 from before the changes, and documents the effect of implementing the treatment. Three of the threat scenarios and two of the unwanted incidents are expected to disappear. However, final confirmation will be awaited from the test engineers to be sure the treatment effectively addresses the threat scenarios.

Notice that fully understanding how risks change as the target system change, it is not enough to identify only the risks that disappear, the risks that maintain and the risks that arise. For risks that are present both before and after the changes, a risk estimation must be conducted in order to determine whether the risk levels change. For example, the unwanted incident *TPSP delivers service without customer consent in violation of Sales Policy* may be mitigated by the identified treatment. Due to other sources, the unwanted incident may not disappear, but one may expect a reduction of likelihood. The risk estimation is, however, outside the scope of the HOMES case study reported in this deliverable.

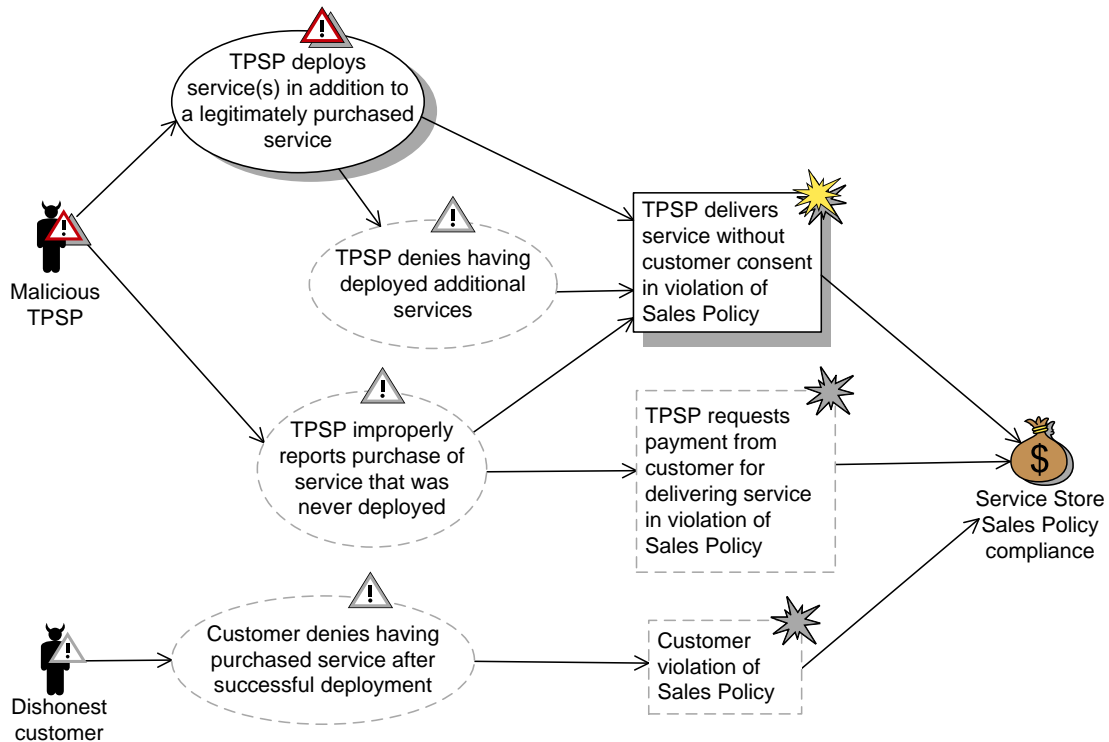


Figure 134 CORAS threat diagram for the HOMES system at T3

15.6.3 Identification of Risk to Change

The treatment is accepted as a change request to address the complaints. The risk analysis team conducts another risk analysis to identify the risks to change, i.e. the risks that may arise due to the implementation of the change request.

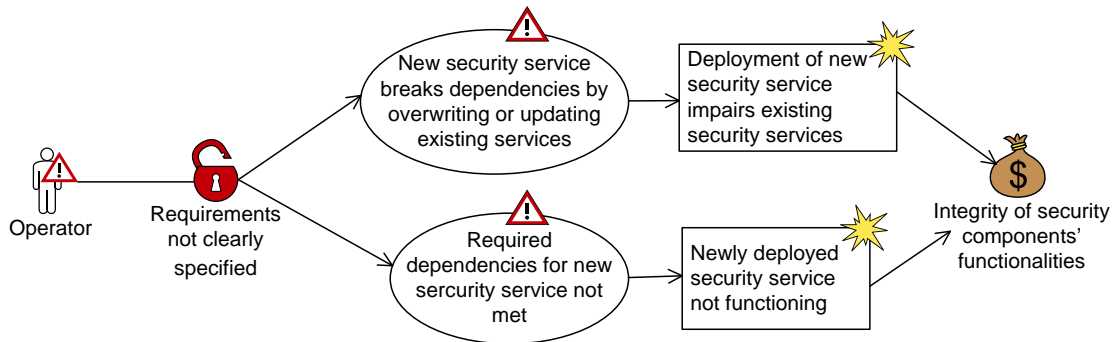


Figure 135 CORAS diagram of risk to change

The risks to change, which are related to the application of the treatment of deploying of a non-repudiation security service, are depicted in Figure 135. In particular, two potential security incidents have been identified which could impair already deployed security services or hinder the correct working of the non-repudiation security service.

16 F – Glossary

In this glossary we provide the definitions we apply for a number of central concepts in risk analysis.

Asset: Something to which a party assigns value and hence for which the party requires protection.

Assumptions: The assumptions of the analysis are what we take as granted or accept as true (although they may not be so); the assumptions may be about the target and about the environment; the results of the analysis are valid only under these assumptions.

Consequence: The impact of an unwanted incident on an asset in terms of harm or reduced asset value.

Context: The context of the analysis is the premises for and background of the analysis; this includes the purposes of the analysis and to whom the analysis is addressed.

Environment: The environment of the target is the surrounding things of relevance that may affect or interact with the target; in the most general case, the rest of the world.

Focus: The focus of the analysis is the main issue or central area of attention in the risk analysis; the focus is within the scope of the analysis.

Likelihood: The frequency or probability of something to occur.

Party: An organization, company, person, group or other body on whose behalf the risk analysis is conducted.

Risk: The likelihood of an unwanted incident and its consequence for a specific asset.

Risk level: The level or value of a risk as derived from its likelihood and consequence.

Scope: The scope of the analysis is the extent or range of the target of the analysis; the scope defines the border of the analysis, i.e. what is held inside of and what is held outside of the analysis, what is the target and what is the environment.

Target: The target of the analysis is the system, organization, enterprise, etc., or parts thereof, that is the subject of the risk analysis.

Target description: The target description is a description of the target including its focus, scope, context, environment, assumptions, parties and assets; only the parts or aspects of the environment that are relevant for the target and the analysis are included in the target description.

Threat: A potential cause of an unwanted incident.

Threat scenario: A chain or series of events that is initiated by a threat and that may lead to an unwanted incident.

Treatment category: A general approach to treating risks; the categories are avoid, reduce consequence, reduce likelihood, transfer and retain.



Treatment scenario: The implementation, operationalization or execution of appropriate measures to reduce risk level.

Unwanted incident: An event that harms or reduces the value of an asset.

Vulnerability: A weakness, flaw or deficiency that opens for, or may be exploited by, a threat to cause harm to or reduce the value of an asset.

References

- [1] CRAMM – the total information security toolkit. <http://www.cramm.com> Accessed 25 March 2010.
- [2] Alberts, C., J., Dorofee, A. J., “OCTAVE Criteria Version 2.0”, Tech. report CMU/SEI-2001-TR-016. ESC-TR-2001-016, 2001.
- [3] Asnar, Y., Moretti, R., Sebastianis, M., Zannone, N., “Risk as Dependability Metrics for the Evaluation of Business Solutions: A Model-driven Approach”, ARES, pp.1240-1247, 2008 Third International Conference on Availability, Reliability and Security, 2008.
- [4] Bouti, A., Kadi, A. D., “A state-of-the-art review of FMEA/FMECA”, International Journal of Reliability, Quality and Safety Engineering, vol. 1, pp. 515-543, 1994.
- [5] Breu, R., Innerhofer-Oberperfler, F., and Yautsiukhin, A., “Quantitative assessment of enterprise security system”. International Workshop on Privacy and Assurance. In Proceedings of ARES 2008, pp. 921-928, 2008.
- [6] Brændeland, G., Dahl, H. E. I., Stølen, K. “A modular approach to the modelling and analysis of risk scenarios with mutual dependencies”, technical report A8360, SINTEF ICT, 2008.
- [7] Brændeland, G., Refsdal R., Stølen, K., “Modular analysis and modelling of risk scenarios with dependencies”. Journal of Systems and Software, vol. 83, issue 10, pp. 1995-2013, 2010.
- [8] BSI (Federal Office for Information Security), “IT-Grundschutz Catalogues”, Version 2005, URL: <https://www.bsi.bund.de/english/gshb/download/index.htm>, 2005
- [9] Charniak, E., “Bayesian networks without tears: Making Bayesian networks more accessible to the probabilistically unsophisticated”, AI Magazine, vol. 12, issue 4, pp. 50-63, 1991.
- [10] Dudley, R. M., “Real Analysis and Probability”. Cambridge Studies in Advanced Mathematics, Cambridge, 2002.
- [11] EBIOS – Expression of Needs and Identification of Security Objectives, <http://www.ssi.gouv.fr/en/confidence/ebiospresentation.html>, retrieved 23/6-09
- [12] EUROCONTROL, ESARR advisory material/Guidance document (EAM 2/GUI 5) – Harmonisation of safety occurrence severity and risk assessment, 2005.
- [13] EUROCONTROL, EUROCONTROL safety regulatory requirements (ESARR) 4 – Risk assessment and mitigation, 2001.
- [14] Felderer, M. and Zech, P., Fiedler, F. and Breu, R.: A Tool-based methodology for System Testing of Service-oriented systems. In: Proceedings of the VALID 2010.
- [15] Felderer, M. and Agreiter, B. and Breu, R.: Security Testing by Telling TestStories. In: Proceedings of the Modellierung 2010.

- [16] Humprey, A., "SWOT - Strengths, Weaknesses, Opportunities, Threats", Stanford University, 1960-1970.
- [17] IEC 60300, Event Tree Analysis in Dependability Management – Part 3: Application Guide – Section 9: Risk Analysis of Technological Systems, 1995
- [18] IEC 61025, Fault Tree Analysis (FTA), 1990
- [19] IEC 61882, Hazard and operability studies (HAZOP studies) - Application guide, 2001.
- [20] Innerhofer-Oberperfler, F. and Breu, R., "Using an Enterprise Architecture for IT Risk Management". In Proc. of ISSA'06: Information Security South Africa Conference, 2006.
- [21] Innerhofer-Oberperfler, F. and Breu, R., "An empirically derived loss taxonomy based on publicly known security incidents", in Proc. ARES/CISIS 2009, Fukuoka, Japan, 2009.
- [22] International Organization for Standardization: ISO 31000 Risk management – Principles and guidelines, 2009.
- [23] International Organization for Standardization: ISO Guide 73 Risk management – Vocabulary, 2009.
- [24] Lund, M. S., Solhaug, B., Stølen, K., "Model-Driven Risk Analysis – The CORAS Approach", Springer, 2010.
- [25] Mannan S., Lees, F. P., "Lee's loss prevention in the process industries", vol. 1, 3rd ed., Butterworth-Heinemann, 2005.
- [26] MEHARI: Information risk analysis and management methodology, <https://www.clusif.asso.fr/en/production/mehari/>, retrieved 23/6-09
- [27] MODELPLEX deliverable D3.3.g: "DSML for security analysis", 2009.
- [28] Peltier, T. R., "How to Complete a Risk Assessment in 5 Days or Less", Auerbach Publications, 2008.
- [29] Robinson, R. M., Anderson, K., Browning, B., Francis, G., Kanga, M., Millen, T., Tillman, C., "Risk and Reliability – An Introductory Text", 5th ed. R2A, 2001
- [30] SecureChange Deliverable D1.1.1: Selected change requirements and security properties, 2010.
- [31] SecureChange Deliverable D5.1: Evaluation of existing methods and principles, 2009.
- [32] Schneier, B., "Attack trees: Modeling security threats". Dr. Dobb's Journal of Software Tools, vol. 12, issue 24, pp. 21-29, 1999.
- [33] IEEE, Standard Glossary of Software Engineering Terminology, 1990.