Chapter XVIII
Assessing Enterprise Risk Level:
The CORAS Approach

Fredrik Vraalsen
SINTEF, Norway

Tobias Mahler
SINTEF, Norway

Mass Soldal Lund
SINTEF, Norway

Ida Hogganvik
SINTEF, Norway

Folker den Braber
SINTEF, Norway

Ketil Stølen
SINTEF, Norway

ABSTRACT

This chapter gives an introduction to the CORAS approach for model-based security risk analysis. It presents a guided walkthrough of the CORAS risk analysis process based on examples from risk analysis of security, trust and legal issues in a collaborative engineering virtual organisation. CORAS makes use of structured brainstorming to identify risks and treatments. To get a good picture of the risks, it is important to involve people with different insight into the target being analysed, such as end users, developers, and managers. One challenge in this setting is to bridge the communication gap between the participants, who typically have widely different backgrounds and expertise. The use of graphical models supports communication and understanding between these participants. The CORAS graphical language for threat modelling has been developed especially with this goal in mind.

INTRODUCTION

Businesses face an increasing number of security risks in the online world, not limited to those of a technical nature. At the enterprise level, technical aspects of security are tightly interwoven with other aspects such as trust and legal issues. This is particularly true for the new breed of networked, virtual organisations. A virtual organisation (VO) can be understood as a temporary or permanent coalition of geographically dispersed individuals, groups, organisational units or entire organisations that pool resources, capabilities and information to achieve common objectives (Dimitrakos, Goldby, & Kearney, 2004).

Virtual organisations' dependency on information and communication technology for performing their daily work leads to a number of risks related to security, trust and legal issues. One area where VOs face risks is the protection of intellectual property (IP) and confidential information, which is the focus of the case study presented in this chapter. Confidentiality is the property that information is not made available or disclosed to unauthorised individuals, entities or processes (ISO/IEC 13335, 2004). The individual stakeholders in a VO desire to protect their IP and maintain confidentiality of information, but at the same time they need to share some of this information with the other partners in the VO in order to fulfil common objectives as well as specific obligations laid down in a co-operation contract. The risks are exacerbated by the international nature of many VOs, as well as their dynamic nature where participants can join and leave the VO at any point during its lifetime.

There is no general international legal framework for the establishment and operation of virtual organisations, and legal issues in relation to VOs are still a topic for research. A strategic roadmap for advanced virtual organisations points out that the analysis of legal risks arising from operating VOs, and the development of legal strategies to overcome them, is an important research task in order to support collaborative networked organisations (Camarinha-Matos et al., 2004). Such legal strategies for VOs should focus both on the contracts that need to be put into place and on the technology that may be utilised in order to facilitate and support the collaboration. When drafting the VO collaboration contract, parties need to identify and address risks that may arise from the collaboration. This risk analysis should preferably follow a clear methodology. Some approaches have considered project risk management in a more general setting (Baccarini & Archer, 1999; Raz & Michael, 1999), focusing mainly on the risk of project failure. However, collaborators also need to assess risks to their own assets.

To reduce the risks involved with establishing, joining and operating a VO, an approach for analysing and managing enterprise security risks is needed which takes into account both technical and nontechnical aspects. The collaboration of different experts, like computer scientists and lawyers, is necessary when analysing what may go wrong in a co-operation (Heymann, 2005; Müller-Hengstenberg, 2005). The CORAS model-based risk analysis approach facilitates the integration of these different perspectives, and focuses also on incorporating the context of the system into the analysis, that is, the organisations, processes and people which interact with the system.

CORAS is a framework for model-based security risk analysis. This framework consists of a method, a language, and a computerised tool. The method integrates aspects from different risk analysis techniques with state-of-the-art system modelling methods based on UML 2.0 (OMG, 2005b), the de facto standard modelling language for information systems. The CORAS graphical language for threat modelling is an extension of the UML 2.0 specification language. It is defined as a UML profile (Lund, Hogganvik, Seehusen, & Stølen, 2003), and has recently become part of an OMG standard (OMG, 2005a).

The goal of this chapter is to make the reader familiar with the CORAS method for model based

risk analysis as well as the graphical language used for threat modelling, and explain how they may be employed in the analysis of VOs. The next section provides more background on the CORAS approach, followed by a walkthrough of the CORAS risk analysis process. Finally, we present some concluding remarks.

The example case is based on a risk analysis which was performed in the TrustCoM IST project (*http://www.eu-trustcom.com/*) using CORAS (Mahler, 2005). The analysis focuses on a collaborative engineering project in the aerospace industry, where a group of companies establish a VO to collaborate on the upgrade of an airplane design. The focus of the analysis is on intellectual property rights (IPR) and confidentiality issues in relation to the sharing of trade secrets between the partners of the virtual organisation. Hence, the analysis is also an attempt to contribute to the investigation of methods for legal risk management, which are "in their infancy" (Burnett, 2005).

THE CORAS APPROACH

The CORAS risk management method is based on the AS/NZS 4360 standard for risk management (AS/NZS 4360, 2004). *Risk management* is the sum of the culture, processes and structures that are directed towards effective management of potential opportunities and adverse effects. The *risk management process* consists of systematic application of management policies, procedures and practice to the tasks of establishing the context and identifying, analysing, evaluating, treating, monitoring and communicating risks. Risk management thus covers the entire life cycle of the system or organisation, and may include several risk analyses with different focus areas and abstraction levels as the system or organisation and its surroundings evolve over time.

Risk analysis requires a clear understanding of the system or organisation to be analysed. This can only be achieved through the involvement of stakeholders and other interested parties with different backgrounds and knowledge about the system or organisation (e.g., decision makers, security experts, legal experts, system owners, developers, and users). These participants are involved in the identification and evaluation of risks and treatments through structured brainstorming sessions.

The effectiveness of such sessions depends on the extent to which the participants are able to communicate with and understand each other. The CORAS graphical language for threat modelling (den Braber, Lund, Stølen, & Vraalsen, 2005) has been designed to mitigate this problem within the security domain. The CORAS language covers notions like asset, threat, risk and treatment, and supports communication among participants with different backgrounds through the definition of easy-to-understand symbols associated with the modelling elements of the language. A recent study has shown that the graphical symbols allow the participants to understand and read the diagrams more quickly (Hogganvik & Stølen, 2005). Recent work has also focused on application of the CORAS language and method to the analysis of security, trust and legal issues (Brændeland & Stølen, 2004; Mahler & Vraalsen, 2005; Vraalsen, Lund, et al. 2005), as well as continuous improvements of the language based on experiences from use and empirical investigations.

Figure 1 shows the main elements of a risk analysis and gives examples of the graphical symbols used by the CORAS language. The *target* is the system or organisation, or parts thereof, which is the focus of the analysis. *Assets* are the parts or features of the target which have value to the client commissioning the analysis, such as physical objects, know-how, services, software and hardware, and so on. A *vulnerability* is a weakness of the system or organisation. A *threat* may exploit a vulnerability and cause an *unwanted incident*, an event which reduces the value of one or more of the assets. A *risk* is an unwanted incident along with its estimated likelihood and

Figure 1. Elements of a risk analysis

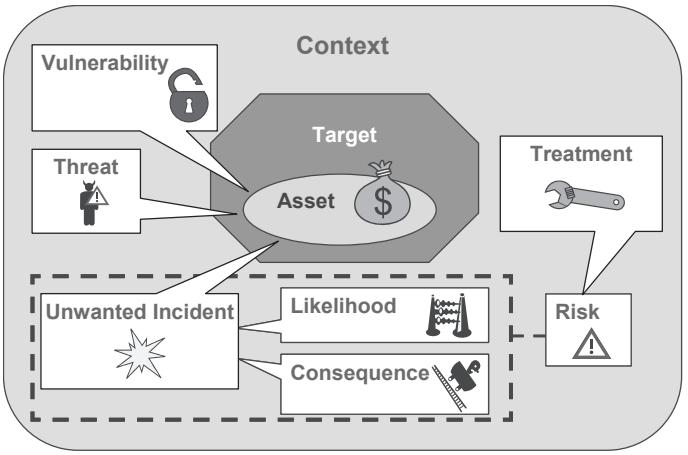

Figure 2. CORAS meetings

Meeting 1: Introduction
- Clients present the system or organisation they wish to analyse
- Identify the focus and scope for the analysis
- Set up analysis plan

Meeting 2: High-level analysis
- Risk analysts present their understanding of the target of analysis
- Identify assets
- Establish initial threats and vulnerabilities

Meeting 3: Approval
- Target of analysis documentation
- Assign values to assets
- Identify risk evaluation criteria

Meeting 4: Risk identification
- Identify risks through structured brainstorming

Meeting 5: Risk estimation and evaluation
- Estimate likelihood and consequence of risks
- Evaluate risks with respect to risk evaluation criteria

Meeting 6: Risk treatment
- Identify and evaluate treatments

Meeting 7: Finalisation meeting (if necessary)
- Present results and get any missing input

consequence values. *Treatments* represent various options for reducing risk.

The CORAS risk analysis process is typically organised as a set of meetings, as summarised in Figure 2. Between the meetings, the risk analysts need time to process the collected information, gather additional necessary documentation, and prepare for the next step of the analysis.

The meeting schedule should be tailored to the needs of each individual risk analysis. Not all activities need to be conducted as face-to-face meetings, but may be performed through for ex-

Figure 3. The CORAS tool

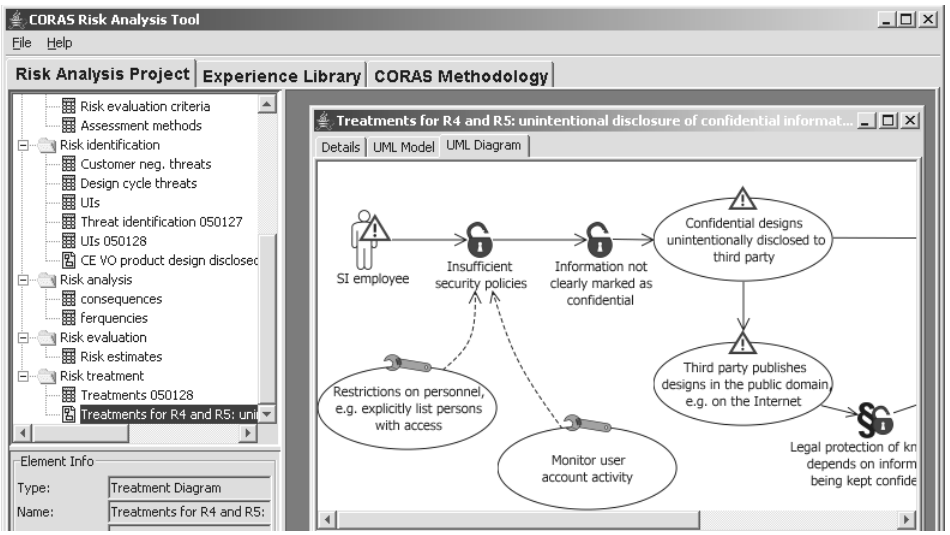

ample phone or video conferences or via e-mail discussions. Some meetings may be combined to save time or costs. For instance, in several cases we have combined meetings four through six into a 2- or 3-day workshop in order to reduce travel costs for the involved participants. This requires careful planning and preparation, however, as well as scheduling time during the workshop to give the analysts a chance to process the output, for example at the end of each day. On other occasions, additional meetings are needed, for example if new information comes up during the analysis which necessitates extra risk identification work to properly identify all the relevant risks.

Risk analysis is an elaborate and prolific process which involves many different types of documentation from different sources, such as UML models, tables with analysis data, and natural language descriptions of the target of analysis. All this information needs to be organised and accessible. In addition, it is important to maintain consistency between all the elements to prevent errors, and we also wish to be able to reuse elements from previous analyses where appropriate to avoid starting from scratch every time. Computerised support for documentation, maintenance and reuse of analysis results is thus of high importance.

The CORAS Tool (Vraalsen, den Braber, Lund, & Stølen, 2005) is a Java-based risk analysis tool which is publicly available as open source (*http://coras.sourceforge.net/*). The client-server architecture of the tool enables multiple risk analysts to collaborate on the risk analysis projects. The risk analyst uses the CORAS client application to create new analysis projects, document and edit risk analysis results in tables and diagrams, generate analysis reports, and manage and reuse experiences from previous analyses. Information can be imported from various modelling and risk analysis tools used by the analyst through standardised data exchange formats, such as XMI for UML. The tool also contains a built in diagram editor for the CORAS graphical language. Help is provided to the user in the form of integrated online versions of the CORAS method and user guides. A screenshot of the CORAS client application is shown in Figure 3.

In the following sections, we will present the risk analysis meetings and activities in more detail. This will be illustrated with examples taken from the TrustCoM risk analysis.

INTRODUCTORY MEETING

The introductory meeting aims at achieving an initial understanding of what the client wishes to have analysed and what kind of risks the client is most concerned about. Some of the questions that should be answered include:

- For whom is the analysis carried out?
- For what purpose do we perform this analysis?
- What do we want to protect?
- What is the scope?

An in-depth analysis can be a time consuming and costly process, and the client typically has limited resources available for risk management. By clearly characterising the target and focus of the study, including identifying what falls outside the scope of the analysis, the available resources can be utilised in the most effective and efficient manner.

During a risk analysis, we make several assumptions and choices with regard to the system or organisation under analysis as well as its surroundings. Documenting these choices and assumptions is necessary in order to determine in which contexts the analysis results are valid. As the system or organisation and its surroundings change over time, these assumptions may no longer hold true. In this case, the analysis may need to be updated to determine whether the risk level of any of the previously identified risks has changed and to identify any new risks which may have arisen. Mechanisms thus need to be put in place to monitor the risks and assumptions and determine when a new risk analysis is necessary.

The introductory meeting should include the risk analysts and the client of the analysis, typically represented by a person with decision making powers with respect to the system or organisation being analysed. The meeting may also involve other stakeholders or parties who have an interest in or knowledge about the system or organisation.

The risk analysts should give a brief presentation of CORAS to familiarise the client with the risk analysis process and some of the methods and techniques which may be used, such as structured brainstorming and the graphical language.

Client Presents System or Organisation

The client presents the system or organisation they wish to have analysed and what kind of incidents they are most worried about. This presentation will typically include a mix of text (prose, tables, etc.), informal diagrams, such as rich pictures (Checkland & Scholes, 1990), and models describing the system or organisation to be analysed. Depending on what the client wishes to analyse, this presentation would normally cover a number of different areas, such as business goals and processes, users and roles, contracts and policies, hardware and software specifications, network layout, and so on.

SI is a company specialising in the integration of different aircraft subsystems. SI wants to win a contract with an airliner for the upgrade of their business jets with a new feature – an in-flight entertainment system. In order to be able to fulfil this objective, it joins a group of aerospace companies to form a virtual organisation in order to pool their resources and know-how and have a better chance of winning the contract. However, before joining the VO, SI wants to perform a risk analysis in order to determine the potential risks involved in this venture, and hires a consultant company to carry out the analysis.

The three main actors in this business scenario are:

- The airliner that operates a fleet of business jets.

- The proposed collaborative engineering VO (CE VO) which has the technical expertise to specify, design and integrate systems into complex products, and which may also manufacture the solution for the customer. Three partners would be involved in this VO; an avionics manufacturer, an in-flight system entertainment provider, and the aforementioned system integrator (SI) – the client of the risk analysis.
- An analysis consultancy which support design activities within engineering companies by performing general analysis work across engineering and scientific sectors.

Figure 4 shows a diagram presented by SI, depicting the actors and their relationships. The various subsystem designs and integrated designs produced and shared during the design process are stored in the product design database (PDD).

Characterise Focus and Scope of the Analysis

The client and the risk analysts should characterise the focus and scope of the analysis. Characterising the focus and scope is important to ensure both a common understanding of the problem at hand and to ensure an efficient use of the available resources by focusing on the aspects of the system or organisation that are of real importance to the client. This includes defining the borders between what is to be part of the analysis (target) and what is to be left out. Part of defining the scope is selecting which security properties are to be considered in the analysis, such as confidentiality, integrity, and availability, as well as other aspects of interest. The risk analysts should interact with the client to clarify any questions or uncertainties with regards to the target of analysis to avoid misunderstandings later on.

The system integrator is particularly concerned about loss of intellectual property and confidential information and the possibility of industrial espionage in connection with exchange of information with other partners, both internal and external to the VO. Retaining confidentiality of the design information communicated with the partners and stored in the Product Design Database is therefore of utmost importance. To prevent other companies from competing with the CE VO proposal, it is also important to protect the confidentiality of the requirements which have been gathered from the airline through the discussions and initial design meetings.

Figure 4. Actors in CE scenario

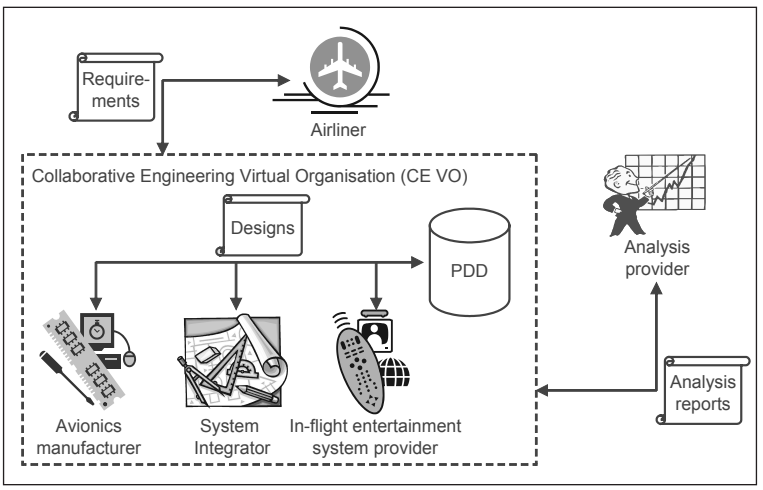

To limit the size of the analysis, other aspects such as data integrity issues, for example, malicious modification or deletion of information because of industrial sabotage or for example virus attacks, are left outside the scope of this particular analysis.

Plan the Analysis

Finally, the rest of the risk analysis should be planned in more detail, including identifying participants and meeting times and venues, based for example on the suggested meeting schedule presented in Figure 2. To achieve continuity in the risk analysis process it is important that the core group of participants commit to the risk analysis and are able to participate during the whole process so that the risk analysts do not have to interact with new and different people at every meeting. Additional persons may be involved in the different meetings based on the competence which is required.

The risk analysis team typically consists of one or two risk analysts who perform the actual risk analysis. One risk analyst should be responsible for leading the risk analysis sessions, and an additional person may act as a secretary during the sessions, recording the results and assisting the risk analysis leader when necessary.

The analysis team should include a representative of the client with decision making power with regards to the target of analysis. In addition, it should include other stakeholders, domain experts, and other interested parties with knowledge about the target of analysis, such as system managers, developers, users, lawyers, security experts, and so on. The goal is to involve people with different backgrounds and different insight into the problem at hand in order to elicit as much relevant information about potential risks as possible. If the risk

Table 1. Risk analysis roles

Role	Name	Organisation	Background/Expertise
Risk analysis leader	Thomas	CORAS Ltd.	Risk analysis, security
Risk analysis secretary	Frank	CORAS Ltd.	Risk analysis, security
Target owner	David	AirFrame Inc.	Aerospace industry
Domain expert	Peter	EngiCorp	Engineering & design
Domain expert	Irene	U. of Oslo	Intellectual property law
Domain expert	Claire	U. of London	Socio-economy and trust

Table 2. Risk analysis plan

Date	Tasks	Participants
29th November	Target identification Asset identification	Analysis leader & secretary, legal expert
11th January	High-level analysis	Analysis leader & secretary, legal expert
27th January	Approval Risk identification	Whole risk analysis team
28th January	Risk estimation and evaluation Risk treatment	Whole risk analysis team
2nd February	Cleanup of results Risk analysis report	Analysis leader & secretary, legal expert

analysis team becomes large, it may be beneficial to split it into smaller groups during parts of the process (e.g., the brainstorming sessions described below). The point is to give everyone a chance to participate and feel useful, as well as to be able to control the group when necessary.

The risk analysis team consisted of two risk analysts with backgrounds in security. The risk analysis team also included two representatives from the client company, the project leader for the aircraft upgrade project and an engineer with good knowledge of the engineering design processes. In addition, it involved an IP lawyer and an expert on socio-economy and trust. The participants of the risk analysis are documented in the risk analysis roles table as shown in Table 1.

The risk analysis for SI was performed over the course of 2 ½ months. Because the participants were spread across several countries, the main part of the analysis was performed during a two day workshop involving the whole analysis team. Other activities were performed in smaller groups and through phone conferences and e-mail discussions. The plan for the analysis is summarised in Table 2.

HIGH-LEVEL ANALYSIS

One of the goals of the second meeting is to ensure a common understanding of the focus and scope of the analysis, as well as to identify the client's main assets in the system or organisation. Assets are central to the CORAS risk analysis method and help guide the entire risk analysis process. The assets are used to assist in identifying risks and estimating their consequences in terms of loss of (monetary) value of the different assets. A high level analysis of threats, vulnerabilities and unwanted incidents is performed to help identify what the client is most worried about happening, and thus to ensure a correct characterisation of the focus and scope of the analysis.

Risk Analysts Present Target of Analysis

Based on the background documentation from the client and the presentations and discussions from the introductory meeting, the risk analysts start by presenting their understanding of the target of analysis, inviting comments and cor-

Figure 5. Target of analysis

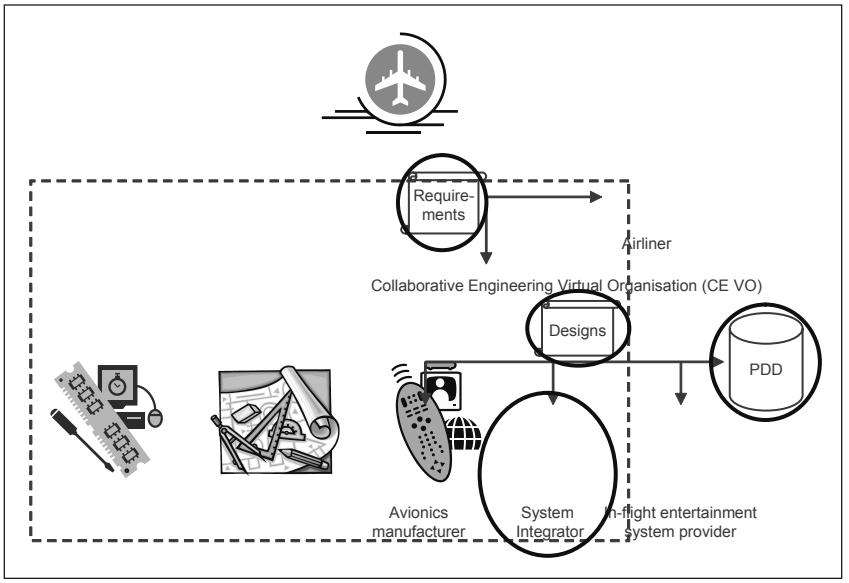

Figure 6High-level CE VO design process

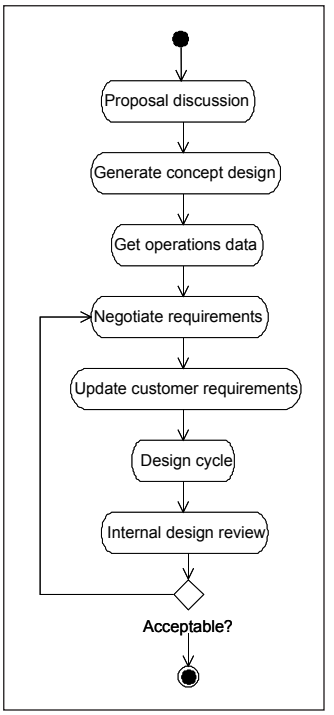

rections from the client. This is done to ensure a common understanding of what is to be analysed and what is to be considered outside the scope of the analysis. The target is characterised using for example UML diagrams or other types of models to specify the target and its relations with the surroundings.

Based on SI's concerns, the focus of the risk analysis is defined as confidentiality of designs and customer requirements in relation to interaction between the partners of the CE VO and other external partners. The product design database (PDD) is central to the exchange of designs between the different CE VO partners and is regarded as a main focus point for the analysis. The target of the analysis is highlighted in the "rich picture" provided by the client of the VO and its partners, as shown in Figure 5.

The documentation provided by the client also contains descriptions of the main business processes related to the aircraft design process. A few of these are selected for analysis, based on where exchange of confidential information between the different participants occurs. The processes are modelled using UML activity diagrams, such as the high-level design process shown in Figure 6.

The legal expert and risk analysts also perform an analysis of the potential contractual obligations and rights of the VO and VO partners. It can be assumed that a number of different contracts will govern the internal and external relations and activities of the CE VO. These will most probably include at least three types of contracts:

- Consortium agreements, which establish a consortium of organisations with a common goal.
- Services or goods related contracts, e.g. outsourcing contracts, which govern the

Figure 7. Contracts in CE VO scenario

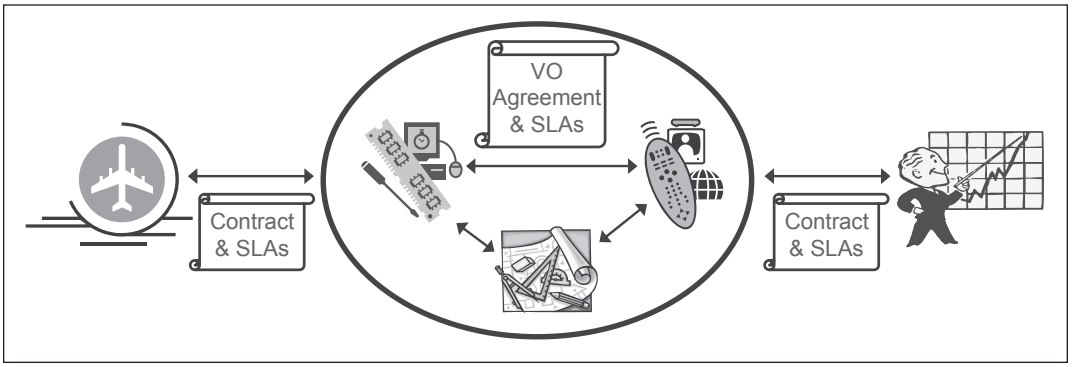

provision of services or the purchase of goods without establishing a consortium.

- Service level agreements (SLAs),that is, (mostly electronic) contracts that deal with the specific rules that partners in an operational business process are bound to.

An overview of these contracts and agreements are shown in Figure 7.

Identify Assets

Assets are the parts or features of the target of analysis that have value to the client and that the client wants to protect, such as physical objects, key personnel, services, software and hardware, or more intangible things such as know-how, trust, market share, and public image. By directing the analysis towards the assets of highest value to the client, one ensures that the available resources are spent on identifying the risks of highest impact on these assets. If the system or organisation does not contain any assets of value to the client commissioning the analysis, there is nothing that can be harmed and lose value for the client, and hence no point in a risk analysis.

The risk analysts typically perform an initial identification of assets based on the information provided by the client in presentations and target documentation. During the meeting, the list of assets is discussed and updated together with the client. To limit the size of the analysis, the number of assets should not grow too large; typically the four to six most important assets suffice.

As mentioned in the target characterisation, the integrated aircraft designs and customer requirements were identified as the most important IP from the viewpoint of the system integrator. In addition, based on the discussion, it becomes clear that the system integrator is also concerned about its public image and how trust may be affected, both the trust of the other VO partners and the trust of customers of the system integrator. The identified assets are shown in the asset diagram in Figure 8.

High-Level Risk Analysis

Sometimes it may be difficult to determine exactly what should and should not be included in the risk analysis. For instance, identifying the most important assets may be hard without also looking at the relevant risks at the same time. Furthermore, the client is often tempted to include as much as possible. However, the result of this may be an inability to analyse anything at all in sufficient detail due to lack of time and resources for the analysis.

A preliminary high-level analysis of the target may be performed to identify the most important assets, threats, vulnerabilities and unwanted incidents to ensure that the focus of the analysis

Figure 8. Asset diagram

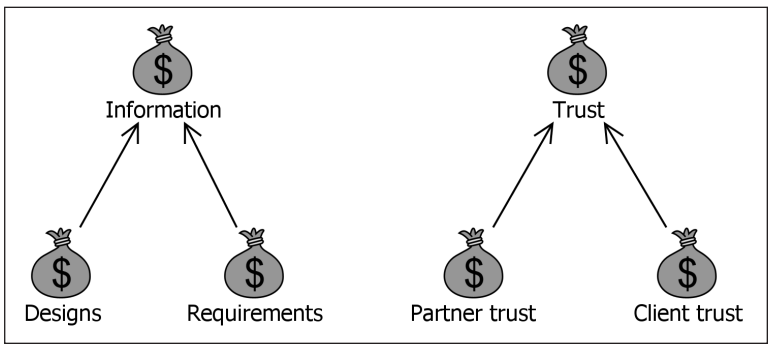

will be on the risks that the client is most worried about. The results of this analysis may help refine the focus and scope of the analysis and also serve as a starting point for the risk identification activity, where the results may be further refined and expanded upon.

This high-level analysis can thus be seen as a first iteration of the risk analysis. The same techniques for risk identification as described in the following sections may be utilised, but more informally. For example, one could use structured brainstorming as described in the section on risk identification below to identify focus areas but leave out the more detailed analysis of likelihood and consequence of the identified risks. The results of the high-level analysis are documented in a table such as the one below.

APPROVAL MEETING

The goal of the approval meeting is to ensure that the background documentation for the rest of the analysis, including the target, focus and scope as characterised by the risk analysts, is correct and complete as seen by the client of the analysis. The documentation of the target of analysis, assets and risk evaluation criteria must be approved by the client.

The client does not have unlimited resources to implement risk reducing measures. We therefore need a mechanism to prioritise the risks and select risks for further attention and treatment. To facilitate this, we must identify the level of risk that the client is willing to tolerate, in terms of loss of asset value over a given time interval. In order to assess the potential loss, we also need to determine the value of the assets.

This meeting should also include people who will be involved in the following risk meetings, such as domain experts, users, and so on, in order to give them an introduction to the analysis.

In preparation for the approval meeting, the risk analysts need to clean up the documentation of the target of analysis and assets. CORAS diagrams should be created for the results of the

Table 3. High-level analysis table

 Threat (deliberate)  Threat (accidental)  Threat (non-human)	 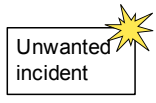  Asset	 Vulnerability
Who/what causes it?	How? What is the incident? What does it harm?	What makes this possible?
…	…	…

Table 4. CE VO analysis asset table

Asset ID	Description	Asset category	Asset value
Designs	SI's share in the designs of the passenger aircraft	Information	Very high
Requirements	The requirements of the VO's customer	Information	High
Partner trust	The VO partner's trust in SI	Other	High
Client trust	The client/customer's trust in SI	Other	Very high

high-level analysis. The resulting documentation should be sent to the client for perusal prior to the meeting.

Documentation of Target of Analysis

The documentation of the target of analysis, i.e. the system or organisation being analysed and the focus and scope of the analysis, forms the basis for the rest of the analysis activities. It is therefore essential that it correctly describes the target of analysis and captures the aspects that the client is most concerned about. A walkthrough is conducted of the documentation, and any errors or omissions are pointed out and recorded. Changes may be performed on the fly or by the risk analysts later on.

Asset Values

After identification, the assets should be ranked according to value or importance to the client, in order to facilitate selection of the most important assets and also prioritising the risks later on. Not all assets can be measured in monetary value, such as human life and health. In these cases, other criteria for risk evaluation may be needed. The assets should be documented in an asset table, as shown in Table 4.

Risk Evaluation Criteria

The goal of this activity is to determine what level of risk the client is willing to accept, in terms of what losses can be tolerated over a given period of time. Risk level is expressed in terms of likelihood, that is, what are the chances of this risk occurring, and consequence, what is the loss with regards to the asset which is affected by the risk. The likelihood and consequence values can be expressed in terms of quantitative values, such as statistical probability or amount of money lost. However, often we do not have the necessary data needed to calculate accurate values. Instead, we may use qualitative values for likelihood and consequence (e.g., low, medium, high), together with examples illustrating what these values mean. The values used for likelihood and consequence can be documented in a value definition table, such as the one shown in Table 5.

The risk evaluation criteria specify what level of risk the client is willing to accept, and should be expressed in terms of the likelihood and consequence values defined above. Based on the

Table 5. Value definition table from CE VO analysis

Value type	Values	
Likelihood	Rare:	Less than once per ten years.
	Unlikely:	Less than once a year.
	Possible:	About once a year.
	Likely:	2-5 times a year.
	Certain:	More than 5 times a year.
Consequence	Insignificant:	No impact on business. Minor delays.
	Minor:	Loss of profits. Lost project phases.
	Moderate:	Loss of project/client.
	Major:	Loss of business sector. Close department.
	Catastrophic:	Out of business

consequence and likelihood, a risk may either be accepted, or selected for further evaluation and treatment. Typically, this is done by setting up a matrix which shows the mapping of consequence and likelihood values to either "accept" or "evaluate," as shown in Table 6. Note that not all risks that end up in the "evaluate" region will necessarily be treated, depending on the availability and cost of effective treatments. Likewise, risks which end up in the "accept" region may still be treated if simple and inexpensive treatments are available.

RISK IDENTIFICATION

This meeting seeks to identify the risks to be managed, i.e. where, when, why and how incidents could prevent the achievement of objectives or reduce the value of an asset. The activity makes use of selected techniques and elements of conventional risk analysis methods which have been adjusted to fit the model-based approach of CORAS. The risk identification session is organised as a structured brainstorming, inspired by HazOp – Hazard and Operability Analysis (Redmill, Chudleigh, & Catmur, 1999).

The goal is to involve people with different backgrounds and different insight into the problem at hand in order to elicit as much relevant information about potential risks as possible. In addition to the risk analysts and the client, the meetings should include people with an interest in and knowledge of the system or organisation under analysis, such as security experts, lawyers, users, system managers, and so on.

Based on the identified assets, models describing the target, and the threats and weaknesses identified by the high-level analysis, the risk analysts should prepare the session by first selecting suitable models as a basis for the analysis, such as use cases, network diagrams, and so on, that match the desired level of abstraction. These should be illustrated using e.g. UML class, sequence or activity diagrams. The risk analysis leader should also prepare for vulnerability identification by selecting suitable checklists. The background documentation, in the form of models, checklists, and so on, should be sent out to the whole risk analysis team prior to the meeting.

Structured Brainstorming

The risk identification activity is organised as a structured brainstorming. The risk analysis team tries to identify scenarios describing how threats exploit vulnerabilities, leading to unwanted incidents which may reduce the value of one or more assets. The risk analysis leader uses the assets of highest value in conjunction with the diagrams of the target to guide the identification process, e.g. by asking relevant questions to the risk analysis team. The use of graphical diagrams also facilitates understanding and communication between the participants. The identification of threats and vulnerabilities may be supported with the use of

Table 6. Risk matrix from CE VO analysis

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

Accept risk

Evaluate risk

Figure 9. Model-based structured brainstorming

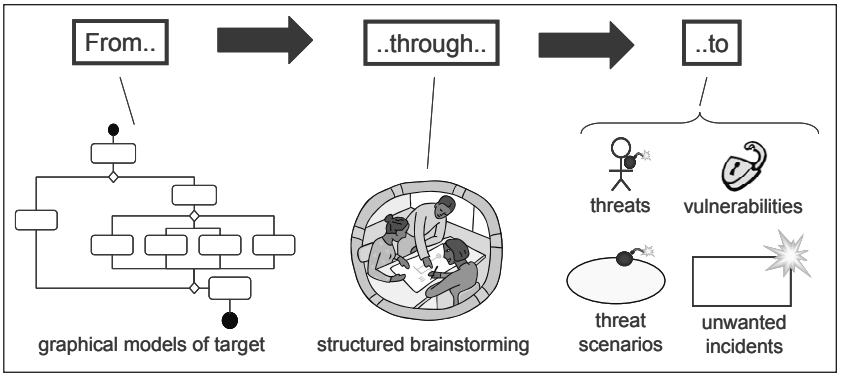

pre-defined questionnaires and checklists. This process is illustrated in Figure 9.

Vulnerabilities can be thought of as control mechanisms that ideally should be in place, but for some reason are missing or not sufficiently robust. Using this metaphor, vulnerabilities can be regarded as unsatisfactory controls, or exceptional circumstances that have not been planned for or that nullify the effect of existing, satisfactory, controls. Vulnerabilities can also be system characteristics that are impossible to treat; an internet connection that is crucial to the system, for example. Identifying new vulnerabilities is often a matter of finding the "blind spot". It is usually necessary to consider all aspects of the target (e.g., the organisational, judicial, physical, and computational characteristics) and compare these findings with the relevant policies.

During the meeting, one person from the risk analysis team should have the responsibility to record and document the results of the structured brainstorming. Following the risk identification meeting, the risk analysts structure the results and document the findings in diagrams using the CORAS graphical language for threat modelling. These diagrams are used later on as a basis for estimating the risk level as well as for identification of treatments. In the CORAS language, a threat (e.g., a disloyal employee or a computer virus) is related to a threat scenario, which is a sequence of events or activities leading to an unwanted incident. A vulnerability may be attached to this relation. An unwanted incident is an event resulting in a reduction in the value of the target asset. Furthermore, an unwanted incident may initiate or lead to other unwanted incidents, forming chains of events.

The risk analysts should also assess the need for further threat or vulnerability identification. For each unwanted incident the risk analysts should decide whether it is described at an appropriate level of abstraction, or whether additional analysis is required. The reason for the latter could be the need for more detailed incidents to make the assignment of frequencies feasible, or that the unwanted incident seems to require a higher priority than originally assigned. Additional information may be elicited from the client or other participants of the risk identification session, or the risk analysis leader may determine that an additional risk identification meeting is needed, but this time focusing on a smaller part of the target of analysis.

As a basis for the analysis, a number of models of the business processes in the organisation were selected. Figure 6 shows a high-level view of the iterative design process used by the CE VO. This process includes a lot of collaboration between the different VO partners, as well as interaction with the airliner at a number of points, such as in the concept and requirements phases.

The identified risks relate to different IPR issues, including the protection of confidential information (i.e., know-how and trade secrets), the ownership of IP, and liability for IPR infringements by other VO partners. The internal collaboration in the CE VO and its cooperation with the analysis company and the airliner, respectively, implies that confidential information is shared or otherwise disclosed to VO partners or to external parties. This involves the risk that such confidential information is disclosed to third parties or used by VO members for purposes that are not related to the VO.

Figures 10, 11 and 12 show use of the CORAS graphical language for describing some ways in which confidential information can be disclosed along with potential consequences this disclosure may have. For example, an employee may have access to confidential information which he/she could disclose to a third party, either willingly or by mistake. This disclosure could lead to the information being used for competitive purposes, or it could reach the public domain and thereby lose its legal protection and value as a trade secret.

RISK ESTIMATION AND EVALUATION

As mentioned in the approval meeting section, the client does not have unlimited resources to implement risk reducing measures. We therefore need to prioritise the risks and select a subset of them for further attention and treatment. Risk estimation is the systematic use of available information to determine how often specified events may occur and the magnitude of their consequences. A *risk* is an unwanted incident along with its estimated likelihood and consequence values. These values are the basis for risk evaluation. The goal of the risk evaluation is to prioritise the risks and identify which ones are in need of treatment by

Figure 10. Hacker steals designs and sells them to competitor

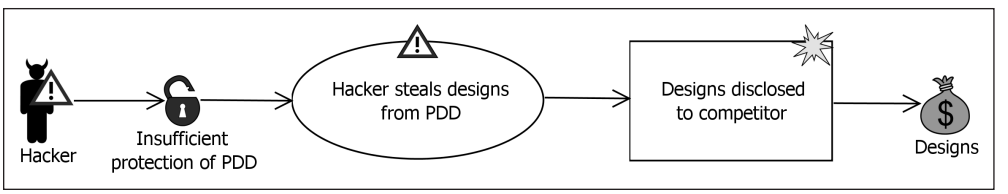

Figure 11. Unfaithful employee discloses customer requirements

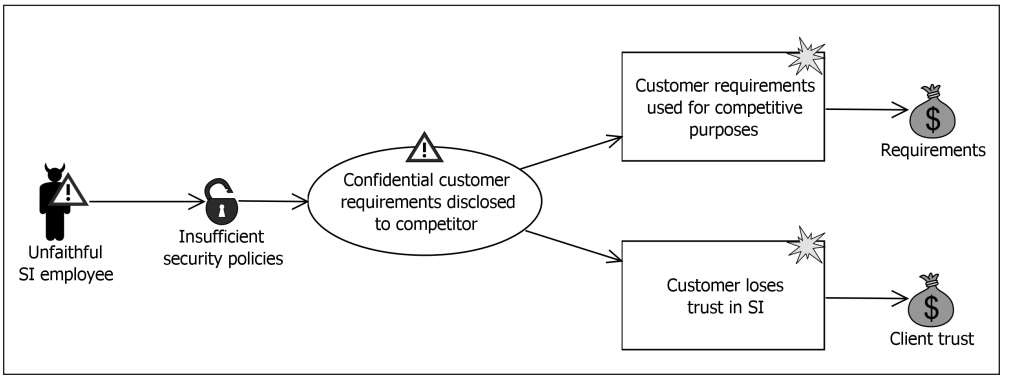

Figure 12. Loss of legal protection for know-how

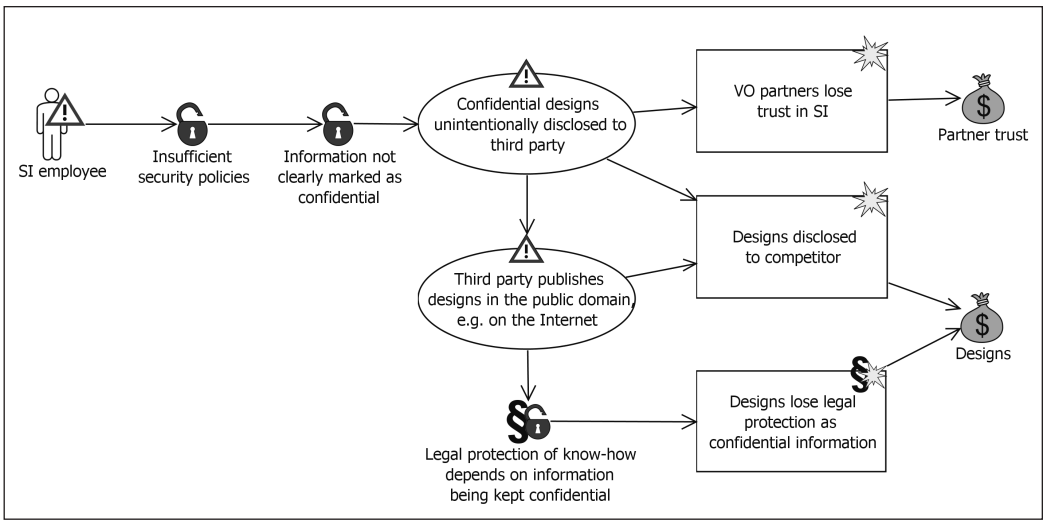

comparing against the preestablished risk evaluation criteria.

Estimate Risk Level

The goal of this activity is to estimate the level of risk for the identified unwanted incidents. This consists of evaluating the likelihood and consequence of the incident. The consequence is a measure of loss of asset value when the incident occurs, while the likelihood is a measure of how often an unwanted incident occurs. The diagrams output by the risk identification activity are used as a basis for the likelihood and consequence evaluation. These document the identified threat scenarios, and may also contain consequence values which have been provided by the risk analysis team during the risk identification.

The methods chosen for consequence and likelihood evaluation depend on the results from the risk identification, the historical and statistical information available, and the analysis group's ability to assign consequence and likelihood values. In many cases, estimates are elicited from the client, domain experts or other people with knowledge of the target of analysis. If statistical

Table 7. Consequence and likelihood table

Risk	Asset	Unwanted incident	Consequence	Likelihood
R1	Designs	Designs disclosed to competitor	Moderate	Unlikely
R2	Requirements	Customer requirements used for competitive purposes	Moderate	Unlikely
R3	Client trust	Customer loses trust in SI	Major	Unlikely
R4	Partner trust	VO partners lose trust in SI	Major	Possible
R5	Designs	Designs lose legal protection as confidential information	Moderate	Possible

or historical data is available, more sophisticated methods may be used, for instance Fault Tree Analysis (IEC 61025, 1990) for calculating the frequency of an incident.

The risk analysis leader presents the CORAS diagrams. For each diagram, consequence and likelihood values are estimated for the different threat scenarios and unwanted incidents, based on expert judgements made by the system owner in collaboration with the risk analysis team. Some of the risks identified in the CE VO analysis are listed in Table 7 along with their consequence and likelihood values.

An example of how calculation of the likelihood of risk R1 could be performed using fault tree analysis is shown in Figure 13. For each event, a probability is given for it occurring during a time period of one year. The resulting probability of 0.28 fits the likelihood category 'unlikely' ("less than once a year").

Fault trees may also be used as a mechanism to decompose and structure scenarios and events without necessarily needing to perform the probability calculations.

Evaluate Risks

The risk evaluation compares the estimated risk level against the pre-established criteria which were identified in the approval meeting. This enables a prioritisation of risks, which is the basis for the subsequent decision about which risks should be targeted for treatment. Note that we may not be in a position to treat all risks, depending on the resources available for establishing risk reducing measures.

Prior to the evaluation, risks may be grouped or categorised. This categorisation can be done according to different concerns, for instance grouping risks which affect the same assets or which stem from the same vulnerability. This may reduce the work necessary for treatment identification and evaluation as the different risks in a category can often be treated using the same approach. An example based on the CE VO risk analysis is shown in Figure 14.

We then apply the risk evaluation criteria specified earlier during the approval meeting.

Figure 13. Fault tree

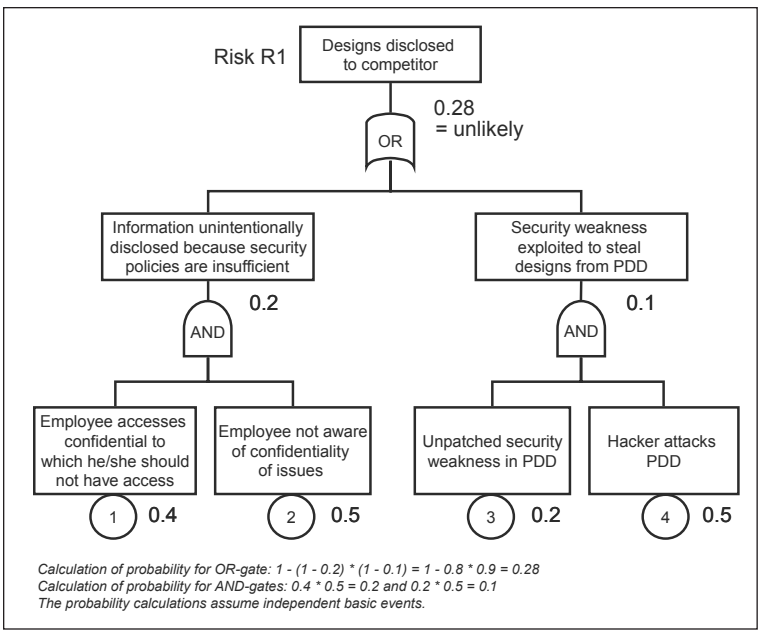

Figure 14. Risk category

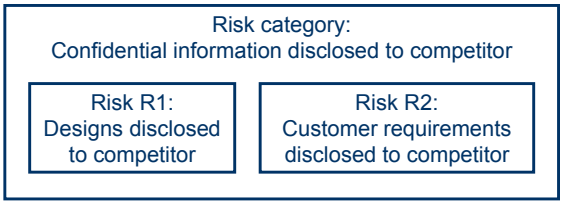

After estimating the likelihood and consequence of the risks, they are plotted into the preestablished risk matrix, as shown in Figure 15. As can be seen, risks R3, R4 and R5 need further evaluation, whereas risks R1 and R2 are accepted and may only need to be monitored to see if their risk level changes in the future. In the evaluation of R3-R5 it was decided that they are all in need of treatment.

RISK TREATMENT

This phase aims at treating the non-acceptable risks by developing and implementing specific cost-effective strategies and action plans for reducing the risk level.

Identify Treatments

For each risk which is not accepted, potential treatment options are explored in a similar manner to the structured brainstorming used for risk identification. This session typically involves the same participants as the risk identification. A walkthrough is performed of the CORAS diagrams created from the risk identification sessions, and the participants are asked to come up with suggestions for different ways to reduce the risk.

There are four main approaches to risk treatment:

- Reduce the likelihood of the incident occurring
- Reduce the consequence if the incident should occur
- Transfer the risk to another party (e.g., through insurance or outsourcing)
- Avoid the activity leading to the risk

The outcome of the treatment identification is documented using the CORAS graphical language by adding treatments to the existing diagrams. The type of treatment is specified according to the main approaches previously listed (e.g., <<ReduceLikelihood>>).

For each of the risks which were not accepted during risk evaluation, potential treatments are explored by the risk analysts and the other participants. A selection of treatments to the risks described above is shown in the CORAS diagrams in Figure 16 and Figure 17. The figures show some

Figure 15. Risks plotted into risk matrix

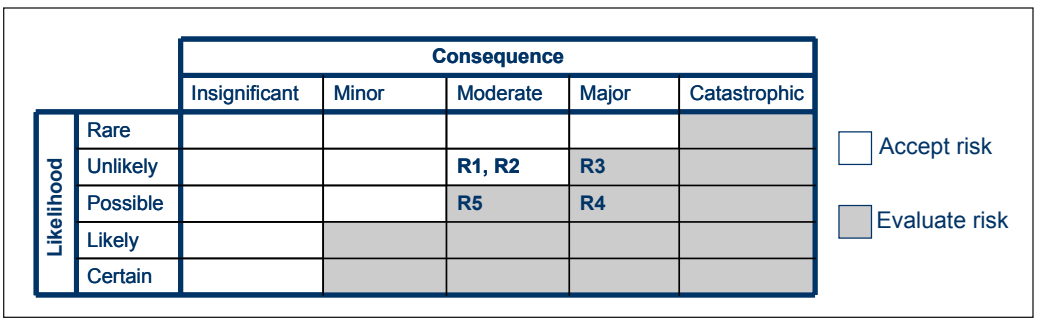

threat scenarios from Figure 11 and Figure 12 and some options for treating them.

The aim in the CE VO analysis was to develop an integrated set of treatments, where legal and other measures are seen together. In this context, the focus was on proactive legal mechanisms, which try to solve legal issues before they arise. Various access rights policies can be imposed via contractual clauses in the agreement between the CE VO partners as well as with the analysis provider (e.g., requiring that access is limited to only those people involved in the project), as well as requiring that access to the confidential information is monitored to allow for auditing. This is shown as two treatments in the figures below, which reduce the likelihood that some of the vulnerabilities from Figure 11 and Figure 12 will be exploited.

Furthermore, if the technology is available, a VO-internal enterprise digital rights management (DRM) system could also reduce the risk of confidential information being disclosed, particularly if some of the contractual obligations could be enforced through technological means. Information security mechanisms like limitations to storage time and the deletion of data after use were also identified as possible treatments.

Figure 16. Treatments for risk R3

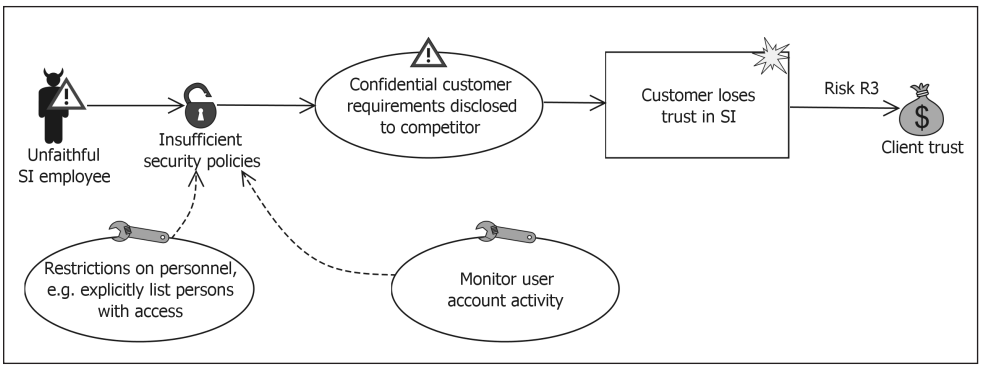

Figure 17. Treatments for risks R4 and R5

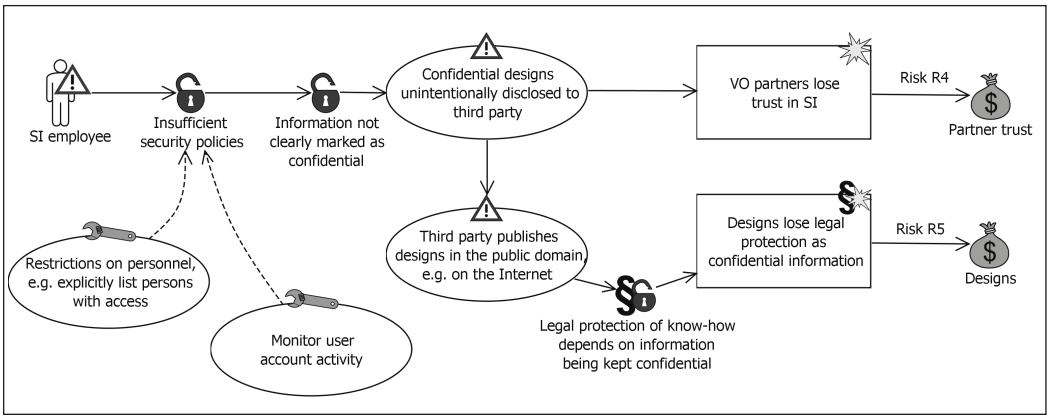

Evaluate Treatments

To determine the best expenditure of the resources available for risk reducing measures, the identified treatments are evaluated with respect to their usefulness. The degree to which the treatment reduces the level of risk is estimated, and a cost/benefit analysis is performed. Table 8 shows some examples of treatment evaluations from the CE VO analysis. Based on these results, the treatments can then be prioritised and implemented based on the available resources.

FINALISATION MEETING

For the risk analysis to have value, the findings of the risk analysis also need to be communicated to the relevant stakeholders to raise awareness and to ensure that relevant measures are put in place to prevent harmful events from occurring. In addition, the results may provide important input to future analyses, serving as a starting point and avoiding the need to start analysing from scratch every time.

The content of this meeting, and whether it is held at all, depends on how the client wants the findings of the risk analysis to be presented. To cut down on costs, the client may forego a written report in favour of a slide presentation of the main findings. Other clients want a written report, or a combination of both.

CONCLUDING REMARKS

In this chapter we have presented the CORAS method for model based security risk analysis and the CORAS graphical language. The risk analysis process has been illustrated with results from the analysis of a collaborative engineering VO scenario, where a number of risks and treatments were identified. The focus of this scenario was an integrated analysis of security, trust, and legal issues. The risk analyses conducted in the TrustCoM project indicate how legal risks, such as the loss of protection of confidential information, can be treated by an integrated solution, including contractual elements, trust management and security management (Mahler, 2005; Vraalsen, 2006). Interestingly, many of the relevant contractual treatments are also included in a general manner in the ALIVE contract template for VOs (ALIVE, 2002a). The risk analyses provide indications about how these rules can be adapted to the specific target under analysis. Since the

Table 8. CE VO treatment evaluation

Risk	Unwanted incident	Asset	Treatment	Risk reduction	Cost
R3	Customer loses trust in SI	Client trust	Monitor user account activity	Major → Moderate	Low
R3	Customer loses trust in SI	Client trust	Access restrictions	Major → Moderate	High
R4	VO partners lose trust in SI	Partner trust	Monitor user account activity	Major → Moderate	Low
R4	VO partners lose trust in SI	Partner trust	Access restrictions	Major → Moderate	Medium
R5	Designs lose legal protection as confidential information	Designs	Monitor user account activity	No	N/A
R5	Designs lose legal protection as confidential information	Designs	Access restrictions	No	N/A

graphical representation implies a simplification, a lawyer would have to integrate analysis results into the contractual document in an appropriate way, taking into account the terminology and the system of the contractual template.

The analysis results presented in this chapter were generated during a number of brainstorming sessions involving participants with varied backgrounds, including law, computer science, engineering, economics, and formal methods and languages. Based on our experiences, the graphical models can indeed facilitate the communication and understanding with respect to security and legal issues in a multidisciplinary context, and this is also supported by other studies (Hogganvik & Stølen, 2005).

As a result of experiences and feedback from the risk analyses conducted in TrustCoM and other research projects, a number of improvements have been made both to the CORAS method and the graphical language (Vraalsen, 2006). Some of these improvements have been aimed at better support for legal risk analysis (Vraalsen, Lund, et al., 2005). Facilities have been added to enable modelling of legal risks and treatments, and reusable elements have been created in the form of e.g. checklists for legal risks. A number of general improvements have also been made. For instance, users were confused by the different types of lines and arrows in the diagrams, and these have now been cleaned up.

Current work focuses on updating the CORAS Tool with support for the new method and graphical language features. The built-in diagram editor has been extended with the new language facilities to support modelling of legal risks and treatments. Work is also being done on improving the reporting facilities in the tool and on updating the integrated online method handbook and tutorials.

ACKNOWLEDGMENT

The developments presented in this chapter were partly funded by the European Commission through the IST programme under Framework 6 grant 001945 to the TrustCoM Integrated Project and partly financed under the Research Council of Norway through the projects SECURIS (152839/220) and ENFORCE (164382/V30).

We would like to acknowledge the work done by David Goldby from BAE Systems, who has defined the collaborative engineering scenario for TrustCoM. We would also like to thank David Goldby, Xavier Parent and Claudia Keser for participating in the risk analysis sessions.

REFERENCES

AS/NZS 4360. (2004). *Risk management*. Standards Australia/Standards New Zealand.

ALIVE IST Project. (2002a). *Intellectual and industrial property rights (Tech. Rep. D13).* Retrieved June 5th, 2006, from http://www.vive-ig.net/projects/alive/docs.html

ALIVE IST Project. (2002b). *VE model contracts (Tech. Rep. D17a)*. Retrieved June 5th, 2006, from http://www.vive-ig.net/projects/alive/models.html

Baccarini, D., & Archer, R. (2001). The risk ranking of projects: A methodology. *International Journal of Project Management, 19*, 139-145.

den Braber, F., Lund, M. S., Stølen, K., & Vraalsen, F. (2005). Integrating security in the development process with UML. In M. Khosrow-Pour (Ed.), *Encyclopedia of information science and technology* (pp. 1560-1566). USA: Information Resources Management Association.

Brændeland, G., & Stølen, K. (2004). Using risk analysis to assess user trust – A net-bank scenario. In C. Jensen, S. Poslad, & T. Dimitrakos (Eds.), *Proceedings of the Second International Conference on Trust Management (iTrust 2004)* (pp. 146-160). Oxford, England: Springer LNCS 2995.

Burnett, R. (2005). Legal risk management for the IT industry. *Computer Law & Security Report, 21*, 621-67.

Camarinha-Matos, L., Afsarmanesh, H., Löh, H., Sturm, F., & Ollus, M. (2004). A strategic roadmap for advanced virtual organizations. In L. Camarinha-Matos & H. Afsarmanesh (Eds.), *Collaborative networked oganizations: A research agenda for emerging business models*. New York: Springer.

Checkland, P., & Scholes, J. (1990). *Soft systems methodology in action.* New York,: Wiley.

Dimitrakos, T., Goldby, D., & Kearney, P. (2004). Towards a trust and contract management framework for dynamic virtual organizations. In *E-Adoption and the knowledge economy: eChallenges 2004*. Vienna, Austria: IOS Press.

Heymann, T. (2005). Outsourcing in Deutschland – eine Bestandsaufnahme zur Vertragsgestaltung. Die Grundtypen des Outourcing und ihre Konsequenzen für die Vertragsgestaltung. *Computer und Recht, 10,* 706-710.

Hogganvik, I., & Stølen, K. (2005). On the comprehension of security risk scenarios. In *Proceedings of the 13th International Workshop on Program Comprehension (IWPC 2005)*, 115-124.

IEC 61025. (1990). *Fault tree analysis (FTA).* International Electrotechnical Commission.

ISO/IEC 13335. (2004). *Information technology – Security techniques – Management of information and communications technology security – Part 1: Concepts and models for information and communications technology security management.* International Organization for Standardization/ International Electrotechnical Commission.

Lund, M. S., Hogganvik, I., Seehusen, F., & Stølen, K. (2003). *UML profile for security assessment* (Tech. Rep. STF40 A03066). Oslo, Norway: SINTEF Telecom and Informatics.

Mahler, T. (Ed.). (2005). *Report on legal issues* (Tech. Rep. D15). Oslo, Norway: TrustCoM EU IST Project 01945.

Mahler, T., & Vraalsen, F. (2005). Legal risk analysis with respect to IPR in a collaborative engineering virtual organization. In *Proceedings of the 6th IFIP Working Conference on Virtual Enterprises (PRO-VE 2005).* Valencia, Spain.

Müller-Hengstenberg, C. D. (2005). Der vertrag als mittel des risikomanagements. Ein plädoyer für die dynamische projektbegleitung im vertrag. *Computer und Recht, 5*, 385-392.

OMG. (2005a). *UML profile for modeling quality of service and fault tolerance characteristics and mechanisms, available specification* (OMG Document: ptc/2005-05-02) Author.

OMG. (2005b). *Unified modeling language: Superstructure, version 2.0* (OMG Document: formal/2005-07-04). Author.

Raz, T., & Michael, E. (1999). Use and benefits of tools for project risk management. *International Journal of Project Management, 19*, 9-17.

Redmill, F., Chudleigh, M., & Catmur, J. (1999). *HazOp and software HazOp.* Wiley.

Vraalsen, F. (Ed.). (2006). *Methods and Tools for legal risk management* (Tech. Rep. D17). Oslo, Norway: TrustCoM EU IST Project 01945.

Vraalsen, F., den Braber, F., Lund, M. S., & Stølen, K. (2005). The CORAS tool for security

risk analysis. In P. Herrmann, V. Issarny, & S. Shiu, (Eds.), *Proceedings of the 3rd International Conference on Trust Management (iTrust 2005)* (pp. 402-405). Paris: Springer LNCS 3477.

Vraalsen, F., Lund, M. S., Mahler, T., Parent, X., & Stølen, K. (2005). Specifying legal risk scenarios using the CORAS threat modelling language – experiences and the way forward. In P. Herrmann, V. Issarny, & S. Shiu (Eds.), *Proceedings of the 3rd International Conference on Trust Management (iTrust 2005)* (pp. 45-60). Paris: Springer LNCS 3477.