SINTEF A1532 - Open

# REPORT

# Evaluations of the methodology and tool used during the 8th SECURIS field trail

Atle Refsdal, Bjørnar Solhaug

**SINTEF ICT**
Cooperative and Trusted Systems

June 2007

SINTEF

**SINTEF ICT**

Address: NO-7465 Trondheim, NORWAY
Location: Forskningsveien 1
Telephone: +47 22 06 73 00
Fax: +47 22 06 73 50

Enterprise No.: NO 948 007 029 MVA

# SINTEF REPORT

TITLE

**Evaluations of the methodology and tool used during the 8th SECURIS field trial**

AUTHOR(S)

Atle Refsdal, Bjørnar Solhaug

CLIENT(S)

FLO/IKT

| REPORT NO. | CLASSIFICATION | CLIENTS REF. | |
|---|---|---|---|
| SINTEF A1532 | Open | Sten-Vidar Eikrem | |
| CLASS. THIS PAGE | ISBN | PROJECT NO. | NO. OF PAGES/APPENDICES |
| A | 978-82-14-04058-6 | 40332800 | 23 |
| ELECTRONIC FILE CODE | | PROJECT MANAGER (NAME, SIGN.) Ketil Stølen | CHECKED BY (NAME, SIGN.) Ketil Stølen |
| FILE CODE | DATE | APPROVED BY (NAME, POSITION, SIGN.) | |
| | 2007-06-04 | Bjørn Skjellaug, Research director | |

ABSTRACT

This report presents the evaluation of the risk analysis in the 8th SECURIS field trial carried out the autumn 2006 and early 2007. FLO/IKT was the client and the target of the analysis was work with/handling of information with security level up to BEGRENSET outside controlled areas.

The CORAS methodology and the CORAS tool were evaluated in addition to the CORAS modelling language.



| KEYWORDS | ENGLISH | NORWEGIAN |
|---|---|---|
| GROUP 1 | ICT | IKT |
| GROUP 2 | Information systems | Informasjonssystemer |
| SELECTED BY AUTHOR | Model-based risk analysis | Modellbasert risikoanalyse |
| | Evaluation | Evaluering |
| | | |

# TABLE OF CONTENTS

## Executive Summary

This report presents the evaluation of the risk analysis in the 8th SECURIS field trial which was carried out during the autumn of 2006 and early 2007. FLO/IKT was the client and the target of the analysis was work with/handling of information with security level up to BEGRENSET outside controlled areas.

| Issue | Experiences and recommendations |
|---|---|
| Context identification, three (four) meetings | • The client should be made aware that he is expected to give an extensive presentation of the target of analysis at the introductory meeting at least two weeks in advance of the meeting.<br>• The analysis team should be able to extract as much information as possible from the client during the first meeting. This could for example be achieved by having a list of prepared standard questions for the relevant kind of target.<br>• Some questions concerning the distinction between direct vs indirect assets need to be addressed by the CORAS method.<br>• Having prepared suggestions for assets, consequence scales and frequency scales worked well.<br>• The ranking of assets created much discussion, but had little impact on the analysis.<br>• Using the modelling languages recommended in the modelling guidelines for target descriptions in step 2 worked very well. |
| Risk analysis, two workshops | • The first workshop was found to be by far the most difficult meeting to conduct. Disproportionately much time was spent in connection to this workshop by the analysis team.<br>• The analysis team deviated from the modelling guidelines by preparing suggestions for threat diagrams in advance of the first workshop. This was not a success; it took up a lot of time before the workshop and lead to an ineffective brainstorming session.<br>• Finding the right level of abstraction for the threat diagrams and granularity for the unwanted incidents was difficult. These choices have a major impact on the risk estimation and evaluation.<br>• The lack of AND/OR-gates and a decomposition mechanism in the CORAS language was seen as a problem. |
| CORAS tool | • The main part of the tool was not used. This decision was taken on the basis of another team's experiences with the same version of the tool.<br>• The diagram editor worked well except for the lack of reliable cut/paste functionality. This functionality was seen as so important that the team decided to switch to Visio when restructuring the diagrams after the first workshop. |
| Modelling guidelines | • The guidelines were found to be useful, but not sufficient for inexperienced analysts. Some questions that are not addressed are: How should time be prioritized? What is the right level of abstraction/granularity in threat diagrams? How should the target description be exploited during the risk identification workshop?<br>• The modeling guidelines require that likelihoods are assigned to all threat scenarios. But they do not tell how to calculate likelihoods for unwanted incidents based on the preceding threat scenarios. |

## Terminology List

| | |
|---|---|
| **Analysis leader:** | The person who leads the structured brainstorming sessions and maintains contact with the client between sessions. |
| **Analysis secretary:** | The person who documents the results during the sessions with the client. |
| **Analysis team:** | The group of persons who conduct the risk analysis, including the analysis leader and analysis secretary. |
| **Participants:** | The representatives of the client who took part in the risk analysis meetings |
| **Risk analysis:** | The process of context identification, risk identification, risk analysis, risk evaluation and risk treatment. Sometimes referred to as risk assessment. |
| **Structured brainstorming:** | Brainstorming session involving experts on the target of analysis to identify and analyse risks. |
| **Target of analysis:** | The organization or system that is the object of the analysis. External entities that may affect the security of the target are also analysed. |

# 1 Introduction

This section gives an introduction to the analysis background and context and the scope of the evaluation. It furthermore provides an overview of the structure of the report.

## 1.1 Background and Context

The major objective of the SECURIS project is to establish a tool-supported methodology for model-driven specification, analysis and development of secure IT systems. SECURIS builds on the results of the EU projects CORAS and COMBINE and the COBRA project funded by the Norwegian Research Council.

The experiences derived from the field trials serve as a basis for an ongoing evaluation of the CORAS methodology for risk assessment [4] as well as contributing to the improvement of how risk assessment should be carried out in practice. The trials furthermore contribute to the evaluation and improvement of the CORAS UML profile [1] for security assessment and the computerized CORAS tool designed to support the risk analyses.

This report is an evaluation of the 8th SECURIS field trial, which was conducted in cooperation with FLO/IKT. The target of the analysis was work with/handling of information with security level up to BEGRENSET outside controlled areas.

The FLO/IKT field trial started in October 2006 and lasted until January 2007. The activities included three context identification meetings and two workshops. The analysis team consisted of two persons, and a total of 9 persons from FLO/IKT were attending meetings at different stages, but none of these were users. The project leader from SINTEF also participated in some meetings.

## 1.2 The Scope of the Evaluation

The CORAS risk assessment methodology consists of seven steps:

1. *Introductory meeting*: The representatives of the client present their overall goals of the analysis and the target they wish to have analysed.
2. *High level analysis*: The analysts present their understanding of the target. This step also involves a rough, high-level security analysis.
3. *Approval*: A more refined description of the target along with any assumptions and preconditions are approved by the client.
4. *Risk identification*: Unwanted incidents, threat scenarios, threats and vulnerabilities are identified.
5. *Risk estimation*: Consequences and frequency values are estimated for each unwanted incident.
6. *Risk evaluation*: The first overall risk picture is presented. Risk values are computed and placed in the risk matrix. Likelihood and consequence estimates are confirmed or adjusted, along with the acceptable area in the risk matrix.
7. *Treatment identification*: Treatments are identified and their cost/benefit is addressed.

Each of the seven steps will be evaluated separately. The evaluation is based on the experiences from the analysts and (except from the third and fifth meeting) the project leader. Positive and negative experiences are documented for each of the phases and suggestions are made with respect to the conduction of future trials.

Cost/benefit analyses are an important aspect of security risk assessments. A crucial initial question for a potential client is what costs the analysis in itself will involve. This report provides an overview of the resources that were needed to conduct the FLO/IKT field trial. A significant

number here is the number of hours that was spent by the analysis team to carry out the analysis. This number of working hours will be evaluated against the competence and experience of the individuals of the analysis team. Based on this evaluation we will give an estimate of the time and resources that will be required in the general case.

The computerized risk analysis tool[2] is designed to support the risk analysis by providing means to structure and store information from the assessment. This tool consists of two parts. The main tool is used for storing results and producing the analysis report. The diagram editor is a graphical editor is used for creating CORAS diagrams, and can be used independently from the main tool. The editor part of the tool will be evaluated with respect to the feasibility of using it during/between the meetings with the client. At the outset the intention was to evaluate also the main tool, but during the analysis it was decided not to use this tool. It had been used and evaluated in an earlier analysis. Several improvements had been recommended, but not yet implemented. In particular, with the existing version it had proved difficult to produce an acceptable analysis report; this had taken much time and required assistance from a tool expert (the person responsible for maintenance and development). Therefore the analysis team believed that use of the main tool would require too much time, and that the gain from another evaluation would be small since suggested improvements had not been implemented. Besides, the tool expert would not be available for assistance. It was therefore decided to use Word for writing the analysis report. This decision was taken at a point where it was clear that the analysis would exceed its estimated time frame.

During the security risk analysis of the FLO/IKT field trial the CORAS language was used to model the security risks. We will in this report evaluate the language with respect to its strength and weaknesses considering understandability and usefulness as a means to support communication between the participants of the risk analysis meetings.

### 1.3 Main Hypotheses

The following hypotheses have been tested in this trial:

1. The modelling guidelines in the summary of each of the seven steps of the CORAS method for security analysis as presented in [5] gives valuable instructions on how to perform the relevant step of the analysis.
2. The use of CORAS diagrams facilitate the risk analysis process and documentation
3. The CORAS diagram editor is suitable for producing and modifying CORAS diagrams before, during and after workshop meetings.

The testing of the first hypothesis was done by trying to follow the guidelines throughout the analysis process, while the second hypothesis was tested by using CORAS diagrams to document results during and after the workshops (brain storming sessions). The last hypothesis was tested by applying the CORAS diagram editor during and after the sessions with the client. The evaluation of the hypotheses is based on the experience of the analysis team and feedback from the project leader.

### 1.4 Report Structure

In Section 2 a brief description of the data collection method is given. Section 3 gives an evaluation of the CORAS methodology for risk assessment based on the experience from this field trial. Section 4 reports on the practical issues relevant for conducting a risk analysis. In Section 5 we document the use of resources (working hours) that were required to carry out the trial. Section 6 evaluates the three hypotheses presented above. Finally, in Section 7, we conclude the report.

## 2 Data collection method

After each meeting the members of the analysis team took notes on their experience of the meeting, addressing issues such as:

- Was the meeting successful?
- What worked well?
- What did not work well, and why?
- Did the discussions stay on track throughout the meeting?
- Did the participants understand diagrams that were presented?
- Were the concepts clearly understood by the participants?

The evaluations presented in this report are based on these notes and on discussions between the members of the analysis team. In addition to discussing the meeting itself, the analysis team also discussed their experiences with the work before and after each meeting.

# 3 Evaluation of the CORAS methodology for risk assessment.

We will in the following discuss our experiences with each of the seven steps of the CORAS methodology. Recommendations for future trials and real cases will be provided on the basis of the experiences we made. We also give an evaluation of the modelling guidelines given in [5]. At the end of the section we report on some general experiences and recommendations that are relevant to most of the steps of the process.

## 3.1 Introductory meeting

The main point of this meeting is to get the representatives of the client to present their overall goals of the analysis and the target they wish to have analysed. The analysts gather information based on the client's presentations and discussions.

### 3.1.1 Experiences

The initial meeting started with a brief presentation of the security analysis method given by the analysis leader. After that, the analysis team expected the client to present the target of the analysis. However, at this point the client could give very little information about the target, for several reasons. Firstly, the client was reluctant to give information about their system for security reasons, and envisaged a very general or hypothetical analysis. Secondly, not all members of the analysis team had gotten their security clearance, so confidential information could not have been given away in any case. The project leader made it clear that a general/hypothetical analysis could be performed, but that the result would probably not be very relevant, and that the process would then need little involvement from the client. After some discussion it was agreed that the client should choose a specific and suitable target, and that another introductory meeting was needed.

The next meeting (which served as the introductory meeting and will be called meeting #1) was more successful. The client used informal figures to give an overview of the target, and gave good guidelines and delimitations for what should be the focus of the analysis were given. The meeting lasted for about 130 minutes, which was 50 minutes less than had been booked. At this point the analysts had no more questions for the client.

### 3.1.2 Recommendations

It is important that the client is made aware that he is expected to give a presentation of the target of analysis at the introductory meeting, and that this should take up most of the meeting. He is also expected to produce any relevant documentation. We recommend that the client is explicitly informed about this at least two weeks in advance of the meeting, so that preparations can be made.

The introductory meeting (meeting #1) lasted about 50 minutes shorter than had been booked. At this point the analysis team had no more questions for the client. This situation should have been avoided; the analysis team should be able to extract as much information as possible from the client while they have the opportunity. One way to achieve this could be to have a list of prepared standard questions for the relevant kind of target. At least for inexperienced analysts such a list could be useful.

### 3.1.3 Evaluation of the modelling guidelines

Assuming the client has been made aware in advance what is expected from this meeting, the guidelines for this step are adequate.

### 3.2 High level analysis

In this step the analysts present their understanding of what they learned at the first meeting and from documentation that has been made available to them. The step also involves a rough, high-level security analysis. During this analysis the first threats, vulnerabilities, threat scenarios and unwanted incidents are identified. This step was conducted during meeting #2 with the client.

#### 3.2.1 Experiences

At the high level analysis meeting (meeting #2) the analysis team presented their understanding of the target. UML collaboration diagrams, activity diagrams and class diagrams were used to present the target, in addition to some text presented as simple "bullet points". The participants seemed to understand the target description well, which was confirmed by the fact that they gave corrections to details in the specifications. All target specifications were accepted as relevant and correct after some minor modifications.

The analysis team gave some suggestions for assets and the scales used for measuring the assets. The suggestions were based on information given in meeting #1. Having suggestions to discuss worked well, and establishing assets posed no particular problems.

Following the modelling guidelines in the summary of step 2 in [5] the assets were categorized as either direct or indirect. However, there are some points related to this distinction that the two members of the analysis team are not certain how to handle:

1. Should the value/importance of a direct asset (or the consequence of incidents affecting the direct asset) be increased if the direct asset also affects an indirect asset? If so, what do we do in cases where an additional unwanted incident need to occur before the indirect asset is harmed? And if not, should the indirect assets be ignored altogether, or are they also subject to further analysis?
2. In the guidelines it says that indirect assets should be placed outside the region that logically or physically represents the target of analysis. Does this mean that indirect assets are considered to be out of scope of the analysis and need no further analysis?
3. Is it not too early to decide at this point if an asset is indirect or not? Since we have not yet identified the unwanted incidents, we do not know whether any unwanted incident will harm a given asset directly. During the context identification one asset was considered to be indirect. But during the later risk identification some unwanted incidents were identified that could possibly harm this asset directly.

Arrows were used to indicate how assets affect other assets. Work was also started on ranking assets according to their importance, which created much discussion. This may have been partly because not all participants seemed to have a clear understanding of the concepts "asset", "consequence" and "frequency" – one of them arrived late and missed the explanation of these concepts. There was a tendency to rank assets according to (expected) risk or frequency of related incidents rather than importance. One participant found it hard to rank the assets, as this was "like comparing apples and pears". The ranking of the assets turned out to have little impact on the rest of the analysis; its only consequence was that diagrams related to the least important asset were the last to be walked through in the two workshops.

#### 3.2.2 Recommendations

A better explanation on how to deal with indirect assets is needed. The questions raised in Section 3.2.1 need to be answered.

### 3.2.3 Evaluation of the modelling guidelines

The modelling guidelines for asset diagram raised some questions and caused some confusion with the analysis team, as indicated in Section 3.2.1. Also, we do not understand why point 4 (indicate with arrows which assets may affect other assets) comes after point 2 (Place the direct assets within the region). It seems more natural to decide how assets affect other assets before separating indirect assets from direct assets. The importance of step 5 (ranking of the assets) is not clear, since this information is not used in the guidelines for the later steps.

The modelling guidelines for target descriptions were very helpful; by following these closely and using suggested languages (UML class diagrams and activity diagrams) we were able to produce target descriptions that were well understood by the participants and useful for later analysis.

## 3.3 Approval

This step involves a more refined description of the target to be analysed, and also all assumptions and preconditions being made. The step is terminated after all the relevant documentations have been approved by the client. This step was conducted during meeting #3 with the client, and terminated between meeting #3 and meeting #4.

### 3.3.1 Experiences

The updated target description was presented and the final assets along with their ranking and categorization as direct/indirect were established without any problems. Scales for consequence and frequency were also established. Again the analysis team had prepared suggestions that were modified by the participants. The steps of the frequency scale were qualitative, such as "daily" or "weekly". The analysis team gave exact intervals for each step after the meeting. A risk evaluation criteria matrix was agreed upon, again based on a suggestion by the analysis team that were modified by the participants. One participant proposed to have a medium (orange) category of risks between green (acceptable) and red (must be evaluated for treatment). In the end this was not introduced. A risk is either acceptable or it must be evaluated for treatment; no gain was seen in introducing a third category.

The meeting was completed on time. However, not all decision makers were present, so acceptance was given offline between meeting #3 and meeting #4.

Later in the analysis it turned out that one asset was not very relevant to the analysis as an asset – it was very different from other assets, and was not closely connected to the target description. Instead the item in question played a significant role as vulnerability. The analyst team was uncertain how this should be handled. The asset was kept, but little time was spent on finding unwanted incidents related to the asset.

### 3.3.2 Recommendations

Having prepared suggestions for assets, consequence scales and likelihood scales may contribute to quickly get the discussion going among the participants. Apart from this, the analysis team do not have any particular recommendations regarding this step.

### 3.3.3 Evaluation of the modelling guidelines

The guidelines for this step are adequate. Possibly, they could advice that suggestions for assets, consequence scales and likelihood scales be prepared in advance of the meeting.

### 3.4 Risk identification

The goal of this step is to identify as many as possible potential unwanted incidents, as well as threats, vulnerabilities and threat scenarios. This step was conducted during meeting #4 with the client, i.e. during the first structured brainstorming session (workshop 1).

#### 3.4.1 Experiences

The analysis team were concerned that time would be too short for this meeting. Therefore they prepared suggestions for threat diagrams in advance based on information received from the client in advance. This is not what is instructed by the CORAS guideline, and was not a success for two main reasons.

Firstly, it meant that the brainstorming session was not very effective. The participants spent much time discussing different elements of the suggested diagrams. The discussion was allowed to take off in all directions, often without leading to any significant corrections of the prepared diagrams. One reason for this is probably that the diagrams in many cases were fairly complex. It would probably have been better if the participants had been encouraged to come up with unwanted incidents and threat scenarios during the meeting. Another reason was that the participants started to discuss treatments, which they were allowed to do without being stopped by the analysis leader.

Secondly, the analysis team spent much time preparing the diagrams before the meeting. This time was not well spent, as the result would most likely have been better if the diagrams had been produced during the meeting with more involvement from the participants.

Overall, the analysis team had great difficulties in choosing a suitable level of abstraction for the diagrams, both when preparing the diagrams before meeting #4 and when modifying them afterwards. During the workshops it was desirable to have simple diagrams, but in many cases there was too much relevant information to be conveyed in a simple diagram. The members of the analysis team were uncertain how to handle such cases. One solution that was discussed was to split the unwanted incidents into several more specific unwanted incidents, based on the scenarios leading up the incident. However, the team was worried what effect this might have on the risk evaluation later. Clearly, if an incident is split into several more specific incidents in this way, then the frequency of each specific incident may become so low that the risk is deemed acceptable, even if the more general incident (which would have a higher frequency) would have resulted in an unacceptable risk. Another drawback with this solution was that the number of diagrams would become very high – so high that it would be almost impossible to be able to assign frequencies and consequences to all the diagrams during the next meeting.

Another solution regarding the level of abstraction that was discussed was to decompose the scenarios leading up to the incidents. This solution was not chosen since no rules for correctly decomposing a CORAS diagram are established. In the end, the team decided to go for a fairly high level of abstraction in most cases. The result included ca 65 unwanted incidents distributed on ca 30 diagrams, which are still relatively high numbers.

During the preparation of diagrams before meeting #4 the analysts saw a need to introduce AND- and OR-gates in the CORAS diagrams, since in some cases two different threat scenarios needed to occur (at the same time or one before the other) in order to initiate an unwanted incident, while in other cases it was sufficient if only one of the threat scenarios took place. Such cases could in theory have been modeled by making a separate branch for each possible ordering of the threat scenarios, and interpreted two branches leading to a single unwanted incident as an OR. But this would have given very large diagrams. Another possibility would be to merge all the relevant

threat scenarios into i single threat scenario. But this would give a higher level of abstraction, which may not be suitable for assigning probabilities.

### 3.4.2 Recommendations

Follow the modelling guideline as given in [5]! Do not prepare complete threat diagrams in advance of the meeting; this may hinder the brain storming and direct the participants' focus to issues of little relevance. The question on how to model cases where a number of things need to happen (in any order or in parallel) for an unwanted incident to take place need to be answered by the CORAS method. It would also be desirable to have established decomposition mechanisms for CORAS diagrams.

### 3.4.3 Evaluation of the modelling guidelines

Since the analysts chose not to follow the guidelines for this step it is not possible to give a proper evaluation of the guidelines. Therefore the considerations in this section are necessarily speculative. Still, with hindsight we believe that by following the guidelines a better result could have been achieved – probably the result would be simpler diagrams, giving less work before and after the meeting. If so, the guidelines are certainly helpful.

However, the guidelines do not explicitly address the question of what level of abstraction should be used and what should be the granularity of the unwanted incidents. These questions may well come up even if the guidelines are being followed – for example if a participant during the meeting starts to describe a complicated web of scenarios leading up to an unwanted incident in every detail. In such a case we might end up spending very much time on creating a large and unwieldy diagram. The guidelines would give no help, and the analysis leader would have to decide how to handle the case based on experience.

Related to the question of abstraction level/granularity is the question of how to ensure that the risk identification can be completed during the meeting. The guidelines give no guidance on how to prioritize the available time. As long as the guidelines for earlier steps advice assets to be ranked, it would be natural for the guidelines for the risk identification step to advice that incidents affecting the most important assets are prioritized.

Also, the guidelines say nothing about how to make use of the target description during the meeting. Since the target description plays such an important role in the process, one would expect that that it is actively used in meeting #4. For example, a natural thing to do would be to organize the brain storming around each part/diagram of the target specification, by showing one part/diagram at a time and asking the participants what are the relevant incidents and scenarios here.

The members of the analysis team did not understand how the use of regions as instructed in the modelling guidelines for threat diagrams contributes to the risk identification, and why assets belong outside the region. Furthermore, the meaning of "...add more regions if necessary" in point 1 was not entirely clear to the analysts.

All in all the analysis team believes that the modelling guidelines for the risk identification step are helpful, but that they do not give sufficient guidance to ensure that an inexperienced analysis team is able to conduct a successful risk identification step.

### 3.5 Risk estimation

This step involves estimation of likelihood and consequence values to the already identified unwanted incidents, and possibly to the threat scenarios leading up to the unwanted incidents. This step was conducted in the second workshop, i.e. in meeting #5.

#### 3.5.1 Experiences

Values for frequencies and consequences were decided from the estimates of the participants, since in most cases there were no available statistics or historical data to draw values from, and also no experiences from actual users (who were not represented). Frequencies were assigned directly to each unwanted incident without first assigning values to the relevant threat scenarios, which is a small deviation from the summary of step 5 given in [5]. When the analysis team went through the estimates after the meeting, it was discovered that in some cases incidents that where very similar were given different consequence values for the same asset. This was probably because the participants did not remember the earlier estimate, and had no data on which to base the estimate. In some cases the participants chose high estimates in order to ensure that a risk would be further considered for treatment.

Assigning values to the first diagrams took much longer time that the rest of the diagrams. This may have been partly because the first diagrams were the ones deemed most important, and partly because the participants became more confident about the task after going through a few diagrams. Besides, the analysis leader made it very clear that time would run out unless the pace was increased.

When estimating frequencies and consequences for unwanted incidents the following question came up: Should we consider the worst case or an average case? Most unwanted incidents were really too general to be given a single consequence value. One possibility was to split every incident in two – one for the worst case and one for the average case. But this would mean doubling the number of incidents, which would require much more work when estimating consequences and frequencies. The question is of course closely connected to the question about the granularity of unwanted incidents discussed in Section 3.4.1.

#### 3.5.2 Recommendations

Set aside more time for the initial diagram than the rest, so that the participants get a chance to warm up to the task. Group the diagrams so that similar incidents are handled after each other. If estimates cannot be based on data, then the estimates should be checked for possible inconsistencies; if similar incidents are given different consequence values for a given asset, then these estimates should be verified with the client after the second workshop.

#### 3.5.3 Evaluation of the modelling guidelines

According to the guidelines every threat scenario must be given a likelihood estimate. This should not be necessary in cases where the participants are able to estimate likelihoods directly on the unwanted incidents. Furthermore, no guidance is given on *how* to add likelihoods for unwanted incidents based on the likelihoods of the threat scenarios leading up the incidents. Therefore the usefulness of giving estimates for threat scenarios is not clear. Are they only intended to be "nice to know" when giving estimates for the unwanted incidents, or can the values for the unwanted incidents be calculated according to some rule?

In this field trial this was not a problem, since frequencies were estimated directly on unwanted incidents. Strictly speaking this is not in accordance with the guidelines. The guidelines should allow direct assignment of likelihood values to unwanted incidents in cases where such estimates can be just as good as estimates for threat scenarios.

Adding frequency values to the unwanted incidents and consequence values to the unwanted incident-asset relations during the meeting worked well.

### 3.6 Risk evaluation

In this step the client is given the first overall risk picture, and gets an opportunity to adjust the likelihood and consequence estimates and the risk evaluation criteria. Since only 5 meetings with the client were scheduled, this step was conducted towards the end of the second workshop, i.e. in meeting #5.

#### 3.6.1 Experiences

The risks were inserted in their proper place in the risk evaluation matrix. This was not done until after workshop 2 (meeting #5) in order to save time during the meeting. The client was then given the opportunity to adjust the risk evaluation matrix and the frequency and consequence estimates related to the risks. This had to be done "offline" after the final meeting. Since incidents that where very similar had been given different consequence values for the same asset (as noted in Section 3.5.1), the analysts grouped similar incidents together, so that the client could easily check if the consequence values for similar incidents should be harmonized or the initial estimates were correct. A "combined incident" was also introduced representing a set of incidents that were almost identical but triggered from different scenarios. This was done because the analysis team believed that these incidents had to be considered together when judging whether the risk level was acceptable.

#### 3.6.2 Recommendations

For this step it would be useful if the CORAS tool was able to automatically fill in risks in the risk evaluation table based on the assigned likelihood and consequence values. This would save time for the analysts, and would be particularly useful in cases where steps 5 and 6 have to be performed during a single meeting.

#### 3.6.3 Evaluation of the modelling guidelines

As described above, this step had to be completed after the final meeting. This posed no problem, except that the treatment discussion took place before it was established what risks were unacceptable. The analysis team believes the guidelines should give sufficient guidance to perform this step. A minor point is that it is not clear why threats are included but not vulnerabilities in the risk diagram. At least with respect to treatments they are just as important as threats?

### 3.7 Risk treatment

The risk treatment step involves the identification of treatments for the risks that cannot be accepted without further analysis. Since only 5 meetings with the client were scheduled, this step was conducted in the second workshop, i.e. in meeting #5.

#### 3.7.1 Experiences

No separate risk treatment meeting was set up for this risk analysis, so the only time the participants were available for treatment discussion was in meeting #5. About 60-70 minutes of meeting #5 was used for treatment discussions. At this point the risks had not been inserted in the risk evaluation matrix. Therefore we did not have an overview over what risks needed to be considered for treatment according to the criteria. This was not seen a big problem, since there was a small number of threats and vulnerabilities that were involved in most of the unwanted incidents and were obvious candidates for treatment. However, when filling in risks in the risk evaluation matrix after the meeting it turned out that some of the treatments that were identified

during the discussion only applied to risks within the acceptable range. Still, these treatments were related to several risks, so all in all they might still be worthwhile to implement.

The risk treatment discussion was performed by first asking the participants about what treatments they believed would be suitable, and then going through a list of suggested treatments that the analysts had prepared in advance. Some of these suggestions were rejected immediately as too costly or unpractical, while other were considered to be likely candidates for implementation. The threat diagrams were not used to any large extent in the treatment discussion, since there was little time, and the risks had not yet been evaluated against the criteria.

There was no time to estimate the effectiveness of the identified treatments, so the client would have to follow this up on his own after the analysis.

### 3.7.2 Recommendations

The analysis team do not have any particular recommendations regarding this sub-process.

### 3.7.3 Evaluation of the modelling guidelines

As indicated above, there was no time to complete the risk treatment step. No proper cost/benefit estimates were made. The analysts believe the guidelines should give sufficient guidance to perform this step if time is set aside for this.

## 3.8 General experiences and recommendations

At the very first meeting the concepts and methodology was presented and explained for the participants, and the meeting plan was presented. The client commented that it was useful to go through the concepts since they are often used in other ways in other contexts. In order to ensure that all participants have a clear understanding of the concepts, use the terms in the same way and know how far the process has progressed, we recommend that the first 15-20 minutes of each meeting are used to repeat the central concepts, the steps of the methodology and the meeting plan.

This analysis clearly suffered from having an analysis team consisting only of two persons, who both had very little practical experience with the methodology. An office was reserved for the analysis team at the client's workplace. For confidentiality reasons all work and discussions on the analysis had to take place in this office, and the analysts could not discuss the work with any other SINTEF employees except for the project leader. This meant that the analysis team was not able to benefit much from the experience of other analysts, apart from getting general advice before each meeting. The risk identification workshop suffered most from the analysis team's lack of experience, and proved to be by far the most difficult meeting to conduct. All other meetings were considered by the analysts to be fairly successful. We recommend that each analysis team includes at least one member with significant practical experience with the methodology. This is particularly important if the team cannot discuss the analysis with other analysts. If the analysis team is inexperienced, then extra effort should be made to ensure that they are well prepared for the risk identification workshop – preferably by having a rehearsal with experienced analysts.

The guidelines say nothing about how the available time should be distributed to the seven steps. It would be very helpful to have a rough guidance on how much of the available time should be spent on the various activities. This would allow for much better planning, and could serve as a warning if the team spends too much time on a certain task. We recommend that a plan is made for how the assigned number of hours should be distributed over the whole process in advance of every analysis.

# 4 Practical Issues

For meetings #1-#3 and parts of #4-#5 PowerPoint was used to present results to the participants. Modifications on assets, frequency scales, consequence scales and risk evaluation criteria during the meetings were made directly to the PowerPoint file. PowerPoint files or printouts were also given as documentation to the participants between meetings, including the target description to be accepted after meeting #3. Using PowerPoint in this way gave no practical problems.

For meeting #4 the diagram editor of the CORAS tool was used to present and modify the suggested threat diagrams. This worked well during the meeting. Later we decided to switch to Visio, as stated in section 6.3.

For meeting #5 Visio was used to present threat diagrams. The diagrams were produced by the use of a CORAS stencil. During the meeting frequency values were inserted in square brackets on each unwanted incident, and consequence values were inserted on the arrows going from an incident to an asset. Before the meeting the Visio file had been prepared for presentation by ensuring that all font sizes were correctly set and that all diagrams were set to the correct view size and would be fully visible on the projected screen. Each tab contained a single diagram and the tabs were ordered in the way they would be presented. The use of Visio during the meeting worked well.

In order to get the most out of the participants while they are available it is important to ensure that meetings start on time and that breaks do not get too long. The analysis team should check what resources will be available in the meeting room in advance of each meeting and prepare accordingly. On the day of the meeting they should arrive early to ensure that projectors etc. are ready when the meeting is supposed to start. A plan for the meeting should be presented at the beginning of each meeting so that the participants know what will happen and when there will be breaks. The importance of returning on time after breaks should be stressed.

# 5 Required Resources

The table below summarizes the hours spent on the project by the analysis team and the project leader. Travel time is not included.

| Activity | Analysis team (total hours) | Project leader |
|---|---|---|
| Meeting #0 | 2 | 1 |
| Meeting #1 | 4 | 2 |
| Between meeting #1 and #2 | 48 | |
| Meeting #2 | 6 | 3 |
| Between meeting #2 and #3 | 27 | |
| Meeting #3 | 6 | |
| Between meeting #3 and #4 | 49 | |
| Meeting #4 | 10 | 5 |
| Between meeting #4 and #5 | 106 | |
| Meeting #5 | 10 | |
| After meeting #5 | 16 | |
| Writing the analysis report | 33 | |
| Writing the evaluation report | 39 | |
| **Total** | 356 | 11 |

The estimated time set aside for the analysis team was 250 hours. It is clear from the table that far too much time was spent before and after meeting #4. As has been stated in Section 3.4.1, the team struggled to produce satisfactory threat diagrams, both when preparing the diagrams before the meeting and when modifying them afterwards. To a large degree this has to be blamed on the lack of experience of the analysis team, including their decision to prepare threat diagrams in advance of the meeting and difficulties in deciding on the level of abstraction. Further time was wasted since the team also experienced some tool problems during this work, which will be described in section 6.

Also the preparations for meeting #2 took quite much time – this was for the most part spent on producing the target description. In this case the analysis team believes that the result was worth the effort, and that a satisfactory target description could not have been produced in considerably less time.

It is of course impossible to tell how much time an experienced team would have used on the analysis. It seems reasonable that the number of hours spent before and after meeting #4 could be at least halved, thereby reducing the total number of hours by (49+106)/2=76 hours. For the remaining activities the analysts are not aware of any particular tasks that took unduly long time, but believe that an experienced team would generally work a bit more effectively. Besides, the analysis team spent a total of 10 hours during the process on meetings with an experienced analyst outside the team to get general advice for the tasks ahead. Therefore the analysis could most likely have been completed within the estimated 250 hours by a more experienced team, without negatively affecting the result.

Representatives of the client also spent time on the analysis. Meetings were held for a total of 19 hours. From 4 to 6 participants from the client were present at each meeting, giving a total of ca 100 hours. In addition comes the time spent between meetings on going through the documentation produced by the analysis team, for which no numbers are available.

# 6 Evaluation of the hypotheses

We now evaluate the hypotheses presented in Section 1.3 based on the experiences recorded in Section 3, 4 and 5.

## 6.1 Hypothesis 1

The first hypothesis to be tested was the following:

- The modelling guidelines in the summary of each of the seven steps of the CORAS method for security analysis as presented in [5] gives valuable instructions on how to perform the relevant step of the analysis.

We consider this (admittedly rather weak) hypothesis to have been confirmed. Comments on the guidelines for each of the seven steps were given in Section 3. An overall comment is that the guidelines are useful, but far from sufficient to guide a team of inexperienced analysts through a successful analysis. Even if the guidelines are strictly followed, the analysts need to make some choices that strongly affect the result of the analysis. In particular this applies to the level of detail that is chosen when documenting threat scenarios and unwanted incidents. Here the guidelines give no help, so a good decision requires experience.

## 6.2 Hypothesis 2

The second hypothesis to be tested was the following:

- The use of CORAS diagrams facilitate the risk analysis process and documentation

We consider this hypothesis to have been confirmed, at least to the extent that this can be done without comparing the use of CORAS diagrams to alternative approaches. CORAS diagrams were used extensively in meetings #4 and #5. The participants seemed to have a good understanding of the diagrams, and there were very little communication problems between the participants and the analysts. The reason is probably that CORAS symbols are intuitive, and that the meaning and the central concepts (threat scenario, vulnerability etc.) were repeated briefly at the beginning of each meeting.

However, making good diagrams requires experience or more detailed guidelines. It is of course possible to make a bad diagram even with a good language. The main challenges are to choose the right level of abstraction and a suitable structure for the diagrams.

The analysis team thinks the CORAS threat diagrams are not well suited to document cases where two or more threat scenarios need to occur (at the same time or in any order) for an unwanted incident to take place, since AND/OR-gates are not included in language.

## 6.3 Hypothesis 3

The second hypothesis to be tested was the following:

- The CORAS diagram editor is suitable for producing and modifying CORAS diagrams before, during and after workshop meetings.

This hypothesis was only partly confirmed. The diagram editor was easy to use for drawing a diagram. Using it during meeting #4 worked well, apart from the following minor inconvenience: According to the modeling guidelines for step 4 one should start by drawing a region, and then place symbols inside or outside the region. When clicking on any symbol inside the region, the region symbol is first marked. In order to mark the symbol you have to click once more. This is annoying since you usually want to mark the symbol inside the frame, and not the frame itself.

We therefore ended up minimizing the frame so that all symbols were lying outside it, and then resizing it after the rest of the diagram was completed.

We found no possibility to reorder the tabs of the different diagrams of a file in the diagram editor – the order was determined by the order in which they were created. This was very inconvenient during restructuring of the diagrams, since we were not able to group diagrams that logically belonged together after each other.

Even more serious was the lack of reliable cut/paste functionality in the editor; use of this functionality might corrupt the diagrams. This meant that restructuring the diagrams after the meeting would be very cumbersome, since many of the diagrams in this analysis were similar. We considered this functionality to be so important that we decided to switch to Visio, even if this meant drawing all diagrams from scratch. These problems obviously had an impact on the time spent.

# 7 Conclusions

The evaluation of the 8th SECURIS field trial divides into three main issues, namely the evaluation of the CORAS methodology, the evaluation of the modelling guidelines and the evaluation of the diagram editor in the CORAS tool. The two first points are most central in this evaluation. Since they are closely connected they should be considered together.

The experience from this trial is that the context identification part of the process went well. By following the method and the associated guidelines a target description was created and agreed upon with little problems. The models of the target were well understood by the participants, and gave a good basis for communication between the participants and the analysts. Assets, consequence scales and frequency scales were also established without problems. The analysts had prepared suggestions for these in advance of the relevant meetings. This worked very well; the suggestions were not immediately accepted, but triggered a fruitful discussion among the participants which soon lead to a conclusion, so that the proper adjustments could be made.

The risk identification workshop proved to be the most difficult part of the analysis, and suffered from the analysts' lack of experience. The team was concerned that time would run out during this meeting, and made the mistake of deviating from the modelling guidelines by preparing suggested threat diagrams to be presented for discussion in the workshop. This gave a lot of extra work and lead to a less effective brain storming session. The analysts spent much more time in connection to the risk identification workshop than any of the other meetings, and the allotted time for the analysis was exceeded. In hindsight the team believes that a better result would have been reached if the guidelines were followed. But there are still questions regarding the risk identification workshop that are not answered by the guidelines or the method, so that experience is required in order to conduct a successful workshop. Choosing a suitable level of abstraction and granularity for the threat diagrams and unwanted incidents was the biggest problem for the analysis team, and for this the modelling guidelines give little help.

The risk estimation workshop went fairly well, and guidelines were adequate for performing this step. Little statistic data were available, so the challenge here was to come up with realistic estimates. Luckily, the participants were quickly able to agree on both frequency and consequence values in most cases, even if these had to be based on "gut feeling" and qualified guesses. In cases where similar incidents had been given different consequence values, this was pointed out to the client after the meeting, so that estimates could be confirmed or adjusted.

No separate meetings had been set aside for the risk evaluation and risk treatment steps, so this was done partly in the last meeting, partly by the analysis team after the meeting, and partly by the client alone. A proper evaluation for this part of the analysis could therefore not be made. Still, the analysis team believes that these steps are unproblematic, and that the guidelines are sufficient.

The diagram editor of the CORAS tool was found to work well, apart from the lack of reliable cut/paste functionality. In order for the tool to be truly useful and more than just a drawing tool, it needs to be tightly integrated with the main tool. For example, it should be possible to import incidents, consequences, frequencies etc. directly from a diagram to the main tool.

## 8 References

[1] The CORAS Language for Risk Modeling. Part of the UML Profile for QoS and FT standard by the Object Management Group (OMG)

[2] The CORAS project homepage: http://coras.sourceforge.net/

[3] Seehusen, F. and Stølen, K.: *Graphical Specification of Dynamic Network Structure.*

[4] Vraalsen, F., den Braber, F., Hogganvik, I., Lund, M. S. and Stølen, K.: *The CORAS tool-supported methodology for UML-based security analysis.* SINTEF Technical report STF90 A04015, SINTEF ICT, 2004

[5] den Braber, F., Hogganvik, I., Lund, M. S., Stølen, K and Vraalsen, F.: *Model-based security analysis in seven steps – a guided tour to the CORAS method.* BT Technology Journal, 25(1), 2007.