

Report

DeSPoT: A Method for the Development and Specification of Policies for Trust Negotiation

Author(s)

Tormod Vaksvik Håvaldsrud
Birger Møller-Pedersen
Bjørnar Solhaug
Ketil Stølen

SINTEF IKT
SINTEF ICTAddress:
Postboks 124 Blindern
NO-0314 Oslo
NORWAYTelephone: +47 73593000
Telefax: +47 22067350postmottek.ikt@sintef.no
www.sintef.no
Enterprise /VAT No:
NO 948 007 029 MVA

Report

DeSPoT: A Method for the Development and Specification of Policies for Trust Negotiation

KEYWORDS:trust negotiation, trust
management, risk
management, access
control, trust policy**VERSION**
4**DATE**
2012-01-25**AUTHOR(S)**
Tormod Vaksvik Håvaldsrud
Birger Møller-Pedersen
Bjørnar Solhaug
Ketil Stølen**CLIENT(S)**
Research Council of Norway**CLIENT'S REF.**
180052/S10**PROJECT NO.**
90B245**NUMBER OF PAGES/APPENDICES:**
18/0**ABSTRACT**

Information systems are ever more connected to the Internet, which gives wide opportunities for interacting with other actors, systems and resources and for exploiting the open and vast market. This pushes the limits for security mechanisms which in general are too rigorous to fully adapt to such a dynamic and heterogeneous environment. Trust mechanisms can supplement the security mechanisms in this situation to reduce the risk by means of trusted evidences. We propose DeSPoT, a method for the development and specification of policies for trust negotiation. The method supports the capturing of requirements for the trust policy as a specification of acceptable risk, and the specification of trust policies that fulfill the requirements. DeSPoT is created to be easy to use for business level experts, yet demonstrated in an industrial study to be useful for those who develop and maintain the system conducting trust negotiation within acceptable risk. Adherence to a DeSPoT policy should ensure that the target fulfills the organizational level requirements to the trust behavior, and that the target is not exposed to unacceptable risk. The paper gives an example-driven presentation of the method.

PREPARED BY
Tormod Vaksvik Håvaldsrud

SIGNATURE

CHECKED BY
Mass Soldal Lund

SIGNATURE

APPROVED BY
Bjørn Skjellaug, Research director

SIGNATURE

REPORT NO. SINTEF A20174
ISBN 978-82-14-04988-6**CLASSIFICATION**
Unrestricted**CLASSIFICATION THIS PAGE**
Unrestricted

Document history

VERSION	DATE	VERSION DESCRIPTION
1	2011-08-25	Created
2	2011-09-16	First full draft
3	2012-01-03	Minor changes
4	2012-01-25	Final version

CONTENTS

I	Introduction	7
II	Introduction to Trust Negotiation	8
III	Overview of the Method	9
	III-A Three Abstraction Layers	9
	III-B Conceptual Model	10
	III-C Language and Process	10
IV	Step 1: Characterizing the Target	11
V	Step 2: Capturing Requirements for the Trust Policy	11
	V-A Capturing Requirements for the Trust Formation Policy	11
	V-B Capturing Requirements for the Asset Exposure Policy	11
VI	Step 3: Modeling the Trust Policy	12
	VI-A Modeling the Trust Formation Policy	12
	VI-B Modeling the Asset Exposure Policy	13
VII	Step 4: Analyzing the Current Trust Policy with Respect to its Requirements	14
VIII	Step 5: Updating the Trust Policy to Reflect its Requirements	14
IX	Conclusion	14
	IX-A Contribution	14
	IX-B Discussion	15
	IX-C Related Work	15
	References	16

DeSPoT: A Method for the Development and Specification of Policies for Trust Negotiation

Tormod Håvaldsrud^{a,b}, Birger Møller-Pedersen^b, Bjørnar Solhaug^a, Ketil Stølen^{a,b}

^aSINTEF ICT, ^bDepartment of Informatics, University of Oslo

tormod.havaldsrud@sintef.no, birger@ifi.uio.no, bjornar.solhaug@sintef.no, ketil.stolen@sintef.no

Abstract

Information systems are ever more connected to the Internet, which gives wide opportunities for interacting with other actors, systems and resources and for exploiting the open and vast market. This pushes the limits for security mechanisms which in general are too rigorous to fully adapt to such a dynamic and heterogeneous environment. Trust mechanisms can supplement the security mechanisms in this situation to reduce the risk by means of trusted evidences. We propose DeSPoT, a method for the development and specification of policies for trust negotiation. The method supports the capturing of requirements for the trust policy as a specification of acceptable risk, and the specification of trust policies that fulfill the requirements. DeSPoT is created to be easy to use for business level experts, yet demonstrated in an industrial study to be useful for those who develop and maintain the system conducting trust negotiation within acceptable risk. Adherence to a DeSPoT policy should ensure that the target fulfills the organizational level requirements to the trust behavior, and that the target is not exposed to unacceptable risk. The paper gives an example-driven presentation of the method.

Index Terms- Trust negotiation, trust management, risk management, access control, trust policy

I. INTRODUCTION

Internet systems exploit the potential of the open market and the possibility of interacting with a vast number of other systems for the purpose of realizing opportunities. Trust mechanisms are introduced to mitigate the fact that security mechanisms usually are too rigorous to fully adapt to this dynamic and heterogeneous environment. When a system is exposed in an environment it is subject to risk, which we define as the combination of the likelihood of an incident and its consequence for an asset [10]. Security mechanisms are introduced to reduce the risk, but they are not in general able to eliminate it entirely. Moreover, increased security tends to be at the cost of interoperability. Security mechanisms should be used to achieve the necessary security level, i.e. keep risk below a certain critical level, and leave it to trust mechanisms to treat the residual risk by means of trust and uncertain information which is perceived to be true.

Security mechanisms are designed to eliminate uncertainty. However, for this reason they become too rigorous in some situations as they cannot base decisions on uncertain information. In this respect trust mechanisms are more flexible than security mechanisms since uncertainty is an explicit and important property of trust. Our notion of trust is based on the definition proposed by Gambetta [7] and defined as the subjective probability by which an actor (trustor) expects another entity (trustee) to perform a given action on which the welfare of the trustor depends. Trust is hence a probability estimate that ranges from 0 (complete distrust) to 1 (complete trust). Trust is subjective, which means that it is a belief held by the trustor. Welfare in a trust relation refers to an associated asset of the trustor. If the trustee performs as expected, it results in a positive outcome for the trustor. There is, however, always the possibility of deception, and in that case there will be a negative impact. For a trust level of probability p , the likelihood of deception is $1 - p$. This likelihood combined with the consequence is what constitutes subjectively perceived risk [14].

We often need to make assumptions regarding uncertain information, and this forces us to take uncertainty into account when making decisions. Even though information is uncertain it provides important indications of the actual situation. The challenge is to make the uncertainty sufficiently visible so that we are aware of its extent and not just of its existence. In trust mechanisms uncertainty is the focal point of judgment, whereas security mechanisms hide uncertainty by the assumption that it is sufficiently small to be ignored in the clearly defined situation.

To achieve trust, systems may conduct trust negotiation [20] utilizing trust mechanisms. The systematic use of trust mechanisms may be formulated in a trust policy. The system developers need specialized methods to support the development and maintenance of the trust policy in the same way as they need specialized methods for the development and maintenance of security policies. A natural way to describe trust behavior is, as for security behavior, by means of rules. Many security systems do not explicitly define the rules, but rather embed them in the implementation of the system as actions triggered by events. For this reason it is natural to aim for a rule-based policy specification language.

The contribution of this paper is a *method for the Development and Specification of Policies for Trust negotiation* (DeSPoT). Correctly enforcing such a policy should ensure that the trust behavior realizes opportunities while keeping risks at an acceptable level. Focusing on the requirements to risk and the criteria for trust formation at an organizational level, the method aims to support decision makers in understanding the potential implications of trust mechanisms without going into the low level

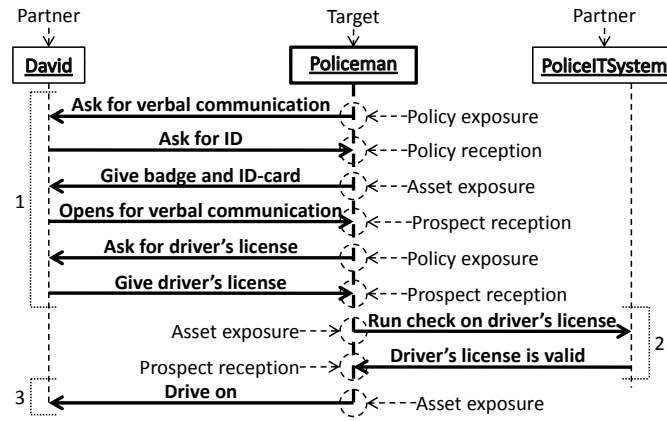


Fig. 1. Example of trust negotiation

details of trust negotiation protocols. This is achieved by systematically linking the high level organizational requirements and criteria to the developed trust policy, which in turn is linked to the low level trust behavior.

The method proposed in this paper should at least have the following characteristics:

- **Risk related:** The method should relate trust decisions to risk.
- **Adherence:** The method should support the ability to ensure that the target's trust behavior adheres to a given trust policy, and that the trust policy meets the requirements specified by the trust policy requirements.
- **Generality:** The method should have the ability to formulate complex authorization mechanisms.
- **Heterogeneity:** The method should allow for flexible definition of evidences.
- **Conceptual foundation:** The method should offer a clear conceptual foundation to ensure that the meaning of the policy is clear.
- **Development guidance:** The method should guide the developers in making a trust policy reflecting the intended trust behavior of the target.

As part of the validation, we introduce in this paper an example from power production and distribution in which all parts of the method are demonstrated. We moreover refer to our report on a recent case study where the method was applied to an ICT system in an industrial setting [9].

The rest of the paper is organized as follows. We give an introduction to trust negotiation in Section II. In Section III we provide an overview of our method. In Section IV through Section VIII, we present the five steps of our method in an example-driven manner. We conclude in Section IX by characterizing our contribution, discussing the suitability of our approach and presenting related work.

II. INTRODUCTION TO TRUST NEGOTIATION

To introduce trust negotiation and our terminology we use an everyday example. Assume a policeman wants to check the validity of David's driver's license. The scenario is illustrated by the sequence diagram in Fig. 1. The diagram has three lifelines, David, the Policeman and PoliceITSystem. The situation is observed from the policeman's perspective, and for that reason we tag the policeman as the **target** of our analysis and David and the PoliceITSystem as **partners**. All events at the policeman's lifeline are tagged to show what kind of event it is with respect to trust negotiation from the policeman's perspective. In general, the chosen target of analysis is the system or organization for which the method aims to develop and analyze a policy to govern the trust negotiation. The partners are the trustees in potential interactions with the target system. Both partners and target are agents performing actions, but with respect to the analysis we differentiate between the target, as the focal point of the analysis, and the partners communicating with the target.

When the policeman approaches David's car he notices that David is intimidated by his pose and checks that the doors of the car are locked. The policeman knocks on the closed window and makes it clear that he wants verbal contact. David indicates that he wants the policeman to identify himself. These two social messages are revealing parts of their respective trust policies; the policeman reveals some of his intentions and David requests some evidence from the policeman. As a reaction to David's request the policeman shows his ID card and his badge so that David can see for himself that the policeman really is a police officer and that it is his badge. The badge and the ID card are **assets** to the policeman, which he exposes to authenticate himself to get the authority he is entitled to. An asset is something of value for the target and therefore needs protection against risk.

David rolls down the window and hands over his driver's license on request from the policeman. David's driver's license is a **prospect** for the policeman; it can tell him whether David is allowed to drive or not. Generally, a prospect is something of value for the target that can be provided by a partner. It is obvious that the driver's license is of value to David, but bear in mind that the target of the analysis is the policeman and everything is seen through his eyes.

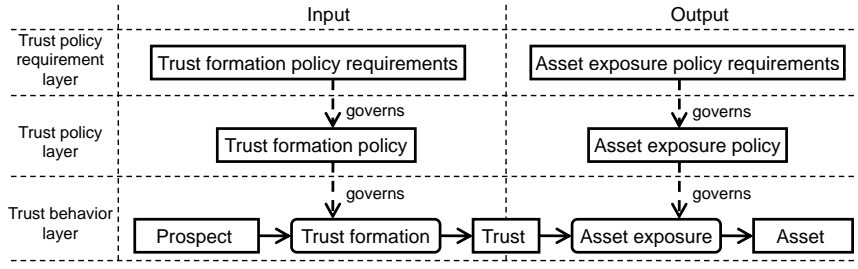


Fig. 2. Trust policy concepts in three abstraction layers

To verify the validity of the driver’s license, the policeman employs the PoliceITSystem and runs a validity check using the license number. For the policeman in this situation, the license number is an asset and the potential response from the PoliceITSystem is a prospect that may validate David’s driver’s license. After the interaction with the PoliceITsystem the policeman knows that David is allowed to drive a car.

The generalization of the policeman’s activity in this scenario is that he starts of with an initial trust level and then receives prospects, and based on these he forms new **trust**. We refer to this as the input side because it is a reaction to input. Based on the trust he has formed, he then exposes his assets to achieve his goals. This is referred to as the output side, since it has to do with outgoing information. The policeman receives David’s driver’s license and after verifying its validity he allows David to drive on. Usually a target communicates with and trusts several partners at the same time. This may be seen as a complex trust and risk situation, but we consider it as several distinct trust negotiations; if the target is communicating with five partners at a given time, this may be projected into five distinct trust negotiations. The state of such a trust situation represents the level of trust the target has at the time in one of the partners. In our little example in Fig. 1 we have numbered three trust negotiation fragments represented by the dotted square brackets. Every trust negotiation fragment starts with an initial trust level. The trust negotiation fragment has access to the pool of all the prospects and assets of the target.

The method presented in this paper aims to support the development and specification of policies that govern such interactions. Adherence to such a policy should ensure that that the target fulfills the organizational level requirements to the trust behavior, and that the target is not exposed to unacceptable risk.

III. OVERVIEW OF THE METHOD

In the following we describe the different parts of our method. In Section III-A we break down the trust negotiation into three abstraction layers with respect to the trust policy: the policy requirements, the policy itself, and the behavior governed by the policy. Further on, in Section III-B, we describe a conceptual model focusing on the target of the analysis, which is a system conducting trust negotiation. After this we describe our table-based modeling approach and the the process on which the method is based in Section III-C.

A. Three Abstraction Layers

The parts of the trust negotiation between the target and one partner as addressed by the method is broken down into the three abstraction layers as illustrated in Fig. 2.

The **trust policy requirement layer** captures and characterizes the trust policy requirements for the target. The requirements specify the highest acceptable risk to which the target can be exposed, and consequently the boundaries for the target’s trust-based behavior. The trust policy requirements consist of two parts. One part belongs to the input side and is referred to as the trust formation policy requirements. They define the boundaries for how the target can form trust based on prospects. The other part belongs to the output side and is referred to as the assets exposure policy requirements. They define the boundaries for how the target is allowed to expose assets based on trust.

The **trust policy layer** defines the target’s trust policy by specifying the rules for the target’s trust-based behavior. The trust policy is also divided in two parts. The input side is the trust formation policy, which is a policy for how trust should be formed based on incoming prospects. The trust formation policy is governed by and must meet the trust formation policy requirements. The output part of the trust policy is referred to as the asset exposure policy, which is a policy for how assets can be exposed based on the trust level. The asset exposure policy is governed by and must meet the asset exposure policy requirements.

The **trust behavior layer** is the actual trust-based behavior of the target, and should adhere to the trust policy. The validation of the trust-based behavior may be realized in at least two ways. One way is that every suggested action is consulted with the trust policy layer. The other way is to realize the trust mechanism as a controller that controls every input and output and enforces the policy. The latter is real policy enforcement, while the former is suitable for situations in which enforcement mechanisms are not implemented.

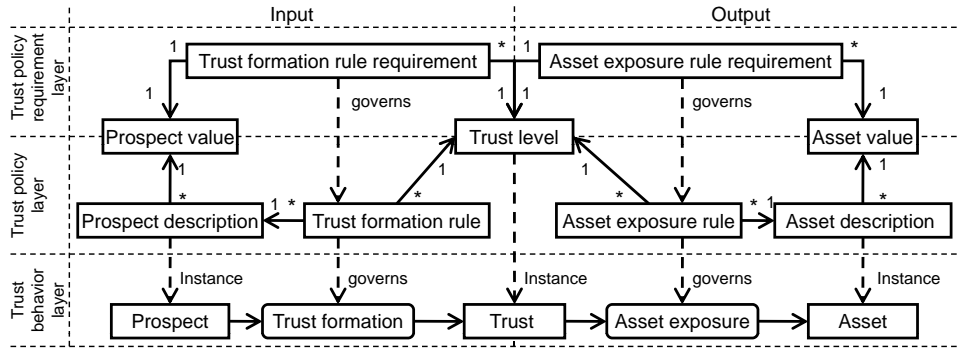


Fig. 3. Conceptual overview of the modeling language

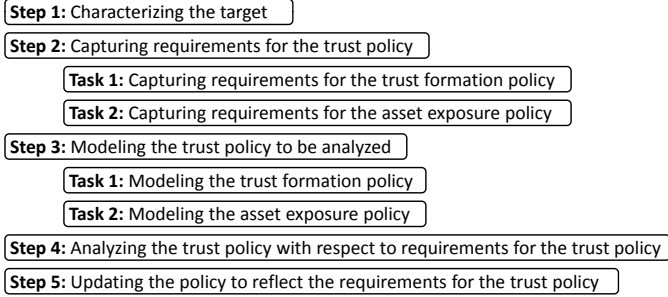


Fig. 4. The five step process

B. Conceptual Model

Our conceptual model is presented in Fig. 3. It concretizes the trust policy requirements layer and the trust policy layer in Fig. 2. Our method comes with modeling support for specifying and reasoning about each of these concepts. The trust behavior layer is as described before. We start explaining the concepts in the trust policy layer on the input side.

A **prospect description** is a general description of the prospects. It can be understood as classification for the purpose of referring collectively to a set of prospects that for the purpose of the modeling and analysis can be treated as the same. A **trust formation rule** is a mapping of a prospect description to an appropriate **trust level**. Such a rule captures how the target is allowed to form trust based on an incoming prospect matching the prospect description. A trust formation policy includes a set of trust formation rules. The value of the prospect is measured in terms of a **prospect value** from a well-defined scale. In the trust policy requirements layer at the same side we have the **trust formation rule requirement** which is a requirement specifying the maximum trust allowed to be formed based on a prospect with a specific value, thus restricting trust formation rules.

On the output side of Fig. 3, we find the **asset description** which is a general description of the assets for the purpose of referring collectively to several assets in the analysis. An **asset exposure rule** relates an asset description to a trust level to regulate when the target is allowed to expose the asset or not depending on the trust in the partner. An asset exposure policy includes a set of such rules. The value of the asset is measured in terms of an **asset value** from a well-defined scale. In the policy requirements layer we find **asset exposure rule requirement** which is a requirement characterizing the level of trust required to expose an asset of a specific value, thus restricting asset exposure rules.

C. Language and Process

We use a table-based modeling approach. The tables are exemplified in later sections. The trust policy requirement layer is described by means of two kinds of tables specifying the set of trust formation rule requirements and the set of asset exposure rule requirements, respectively. The trust policy layer is documented by means of seven kinds of tables; four for the input side and three for the output side. The four for the input side are the prospect value table, the trust formation rule table, the prospect property verification table and the trust policy reception table. The three for the output side are the asset value table, the asset exposure table and the trust policy exposure table.

The overview of the five steps of the DeSPoT process is given in Fig. 4. The process provides guidance in correctly capturing the target's trust policy requirements and the trust policy, as well as ensuring that the latter fulfills the former. The process supports both the analysis of an existing policy, as well as the development of a new. The next sections introduce the process and exemplify its use.

IV. STEP 1: CHARACTERIZING THE TARGET

This section is an example-driven presentation of how to characterize a target. This activity is conducted in close cooperation with the commissioning party, who is usually the target owner. Our example target is a power grid system where we focus on the system balancing of the production and the consumption of electrical power. This task is quite complex and for the purpose of this paper we simplify the system so that we are able to focus on the important features of DeSPoT. The balancing of the production and consumption is conducted in real time, but it also needs planning from day to day. It is this day to day planning on which we focus.

The selected target is the Power Production Organizer (PPO) which is a central software system collecting and spreading key information about power production and consumption. Several systems communicate with the PPO; some are providing necessary information, some are providing business crucial information and others are just providing supplementary information.

A power flow sensor measures the flow of power through a line. It informs the PPO how much power is transported through the power line during the last 24 hours. This is the main source for knowing how the power is distributed and consumed. A power station provides the PPO with information on produced power during the last 24 hours, as well as the expected production capacity. Sometimes the capacity is lower because of maintenance, while in other situations the power stations may have higher capacity than usual, for example due to heavy rain upstream a hydroplant. In addition to this, the power station needs to know how much power it is expected to feed into the grid the next 24 hours to adjust its production. When the PPO has collected all the information, the PPO calculates and finds the best suited production profile for the next 24 hours. It then assigns production quotas for the next 24 hours to all the power stations. This is important because a power station does not always produce at maximum.

Now we continue by defining the scales for prospect value, asset value and trust level from the perspective of the target. These scales may be both quantitative and qualitative, depending on the desired granularity of the analysis or what is otherwise suitable for the target in question.

The prospect value scale is used to measure the value of objects and information given to the target by partners. The value should reflect the direct value, the sensitivity and how easy it is to fake. For example, a credential that is very easy to fake should be given a relatively low value, whereas a message, which can not be forged, about the delivery of a crate with high quality goods can be given a high value. The prospect value scale chosen for the PPO is of five values from 1 to 5. Each of these values must be defined in terms of a precise interval or a qualitative description.

Next we define the scale used to measure the value of assets. As already indicated, an asset is something the target already possesses and that should be protected from harmful incidents. An asset may for instance be sensitive information. Breach of confidentiality, integrity or availability, for example, could be exploited to damage the reputation or revenue of the target. In this situation the potential total loss should correspond to the assigned asset value. The asset value scale chosen for the PPO is of four values from 1 to 4, each of which must be precisely defined.

The last scale needed to be defined is the trust level scale. We use this scale to measure the target's trust in a partner's capability and intention to protect the target's assets. If a partner is highly trusted the target trusts it to manage its assets well without compromising them. The trust level scale chosen for PPO is *No*, *Low*, *Medium* and *High*, each representing an interval of subjective probabilities between 0 and 1.

V. STEP 2: CAPTURING REQUIREMENTS FOR THE TRUST POLICY

This section is an example-driven presentation of how to capture and specify the requirements to the trust policy.

A. Capturing Requirements for the Trust Formation Policy

The **trust formation policy requirements** specify how the target is allowed to use incoming prospects as evidence to form trust. In general, the trust formation policy requirements restrict how the target is allowed to perceive the world with respect to trust. In particular, the trust formation requirements specify the trust level a prospect with a specific value is allowed to support. They also define the limits for how the target can perceive input given to it in the form of trust value. They are defined as a set of trust formation rule requirements each of which associates one trust level to each prospect value level.

The trust formation policy requirements for the PPO are defined in Table I and regulate how prospects are allowed to influence the trust level. The interpretation of the second row, for example, is that a prospect with prospect value 2 can at most support a trust level *Medium*.

B. Capturing Requirements for the Asset Exposure Policy

The **asset exposure policy requirements** specify the trust level the target must have in order to expose assets with a specific value. The trust level is an indirect measure of the likelihood of something going wrong, and the level of this risk can be deduced from the trust level and asset value. The asset exposure policy requirements hence specify the acceptable risk level in this setting.

TABLE I
THE TRUST FORMATION POLICY REQUIREMENTS

Prospect value	Maximum trust level
1	Low
2	Medium
3	Medium
4	High
5	High

TABLE II
ASSET EXPOSURE POLICY REQUIREMENTS

Trust level	Maximum asset value
No	1
Low	2
Medium	3
High	4

The asset exposure policy requirements for the PPO are defined in Table II where each row forms an asset exposure rule requirement. The interpretation of row three, for example, is that when the target has trust level *Medium* it is allowed to expose assets with value 3 or lower, while the last row allows the target to expose all assets (because 4 is the highest asset value) to partners in which its trust is *High*.

VI. STEP 3: MODELING THE TRUST POLICY

This section addresses the specification of the trust policy for the target. The trust policy consists of two parts, namely the trust formation policy and the asset exposure policy, each specified as a set of rules. We need to associate a value with every prospect and asset to determine what is on stake in every decision.

A. Modeling the Trust Formation Policy

The trust formation policy specifies how the target should form trust based on prospects and their properties. The first task is to specify the prospects and their values, as illustrated in Table III. The third row means that a prospect matching the prospect description *Consumer authentication* received by the target has the prospect value 3.

TABLE III
PROSPECT VALUES

Prospect description	Prospect value
Partner's access policy	1
Full postal address	2
Consumer authentication	3
Power certificate signature	4
Master certificate signature	5
Certificate validation	5

A **trust formation rule** defines what may form evidence and how this evidence should influence the target's trust level with respect to a partner. The evidence consists of a prospect, a prospect property and an evidence type. The prospect is received from a partner and can be anything that may give insight into this partner's properties, such as intention or capabilities to take care of the target's assets. A prospect may be the goal for the whole interaction, but it may also be a means to build trust. The **prospect property** may be as simple as the confirmed existence of the prospect itself, but also complex properties such as authenticity and validity of a chained signed electronic certificate. There are two different sorts of evidence, namely supporting evidence and exposing evidence. The **supporting** evidence may build trust, whereas **exposing** evidence on the other hand may reduce trust. Exposing prospect rules are overruling supporting prospect rules, such that trust is governed by the most exposing evidence and otherwise by the best supporting evidence.

The trust formation policy for PPO is presented in Table IV. In this case the rules use the properties *Existing*, *Valid*, *Invalid*, *Correct* and *Trusted*. The first row should be understood as follows: If a partner provides an invalid power certificate signature, it is perceived as an exposing evidence and results in a trust level no higher than *No*. The set of trust formation rules is part of the definition of the trust formation policy for the target.

Sometimes a prospect property must be verified by another prospect. This is typically the case in chains of certificates. In order to take into account and keep track of such relations, we document these in a designated table as exemplified in Table IV.

The first row in Table IV should be understood as follows. A *Certificate validation* prospect with the property *Trusted & CertValid* verifies the *Valid* property of a *Power certificate signature* prospect. Moreover, the *Trusted* property of the former can in turn be verified by the *Correct* property of *Master certificate signature* in the third row.

TABLE IV
TRUST FORMATION RULES

Prospect description	Property	Evidence Type	Trust level
Power certificate signature	Invalid	Exposing	No
Full postal address	Existing	Supporting	Low
Consumer authentication	Valid	Supporting	Medium
Power certificate signature	Valid & Correct	Supporting	High
Master certificate signature	Correct	Supporting	High
Certificate validation	Trusted	Supporting	High

TABLE V
PROSPECT PROPERTY VERIFICATION

Prospect description	Required property	Prospect description	Verified property
Certificate validation	Trusted & CertValid	Power certificate signature	Valid
Certificate validation	Trusted & CertInvalid	Power certificate signature	Invalid
Master certificate signature	Correct	Certificate validation	Trusted

The **policy reception rules** specify how to handle requests from the partner. These are referred to as such, because in trust negotiation the partners expose their policy when requesting assets from the target.

The policy reception rules for the PPO are documented in Table VI. The first row should be understood as follows. If the partner requests that the target system provides a *Power certificate signature* the target may raise the trust level up to *Low* based on this evidence. The prospect property constraint for this prospect is only that it exists, i.e. there are no additional requirements. This evidence is *Supporting* and may support a trust level up to *Low*.

B. Modeling the Asset Exposure Policy

The asset exposure policy specifies how the target may expose assets based on the trust level. The asset exposure policy includes a set of asset exposure rules that describe the actions or responses the target may take as restricted by the level of trust in the actual partner.

The assets identified for PPO are listed in Table VII together with their respective values. The asset *Get total power consumption*, for example, is assigned the asset value *1*. This is a service provided by the target system and it exposes an asset.

The part that remains in order to complete the documentation of the overall trust policy is to establish the asset exposure policy, which is defined as a set of asset exposure rules and a set of policy exposure rules. Each asset exposure rule has two parts, namely the asset being exposed and the trust level needed to perform the exposure.

The **asset exposure rule** for PPO is modeled in Table VIII. The first column specifies the minimum trust level for exposing the associated asset. Hence, the asset *Power certificate signature* can be exposed when the trust level is *Low* or higher. Or, given the trust level *Medium*, all the assets of the first five rows can be exposed.

In addition to expose assets, the PPO also needs to expose parts of its own trust policy to let the partners know its goals. These exposures may also be sensitive and for that reason we explicitly model how the target is allowed to expose policies through **policy exposure rules**. These are presented in Table IX. The first row should be understood as follows: The target

TABLE VI
POLICY RECEPTION RULES

Requested asset	Evidence Type	Trust level
Power certificate signature	Supporting	Low
Report power flow	Supporting	Low
Get assigned power quota	Supporting	Low
Report power consumption	Supporting	Low

TABLE VII
DEFINING ASSETS VALUES

Asset description	Asset value
Get total power consumption	1
Target's access policy	2
Report consumer consumption	3
Report power production	3
Get assigned power quota	3
Report power production capacity	4
Report power flow	4

TABLE VIII
ASSET EXPOSURE RULES

Needed Trust	Requested asset
Low	Power certificate signature
Low	Report consumer consumption
Low	Get total power consumption
Medium	Report power production
Medium	Get assigned power quota
High	Report power flow
High	Report power production capacity

TABLE IX
POLICY EXPOSURE RULES

Needed Trust	Asset description token
No	Power certificate signature
No	Consumer authentication
No	Postal address

must have at least trust level *No* in the particular partner to be allowed to request the partner for the asset *Power certificate signature*.

VII. STEP 4: ANALYZING THE CURRENT TRUST POLICY WITH RESPECT TO ITS REQUIREMENTS

At this point we have both the trust policy and its requirements. In this step we look for possible gaps between them. Every trust formation rule forms trust based on a prospect with a specific value. It can therefore be easily checked against the corresponding trust formation rule requirement which specifies the highest acceptable trust to be formed for a prospect of this value.

Consider, for example, the fourth trust formation rule in Table IV. The prospect description is *Power certificate signature*. According to Table III this prospect has value 4. Further on, evidence formed by the rule supports trust level *High*.

To sum up, the rule supports trust level *High* based on a prospect of value 4. To check whether this trust formation rule adheres to the trust formation requirements we look into Table I. This table states that evidences based on prospects with prospect value 4 and 5 can support trust level *High*. This means that the fourth prospect rule in Table IV meets the trust formation requirements.

The second asset exposure rule in Table VIII assigns access to the service *Report consumer consumption* and requires at least trust level *Low*. The service *Report consumer consumption* has the asset value 3 as shown in Table VII. This rule is then exposing an asset with asset value 3 based on a trust level *Low*. According to the PPO's asset exposure policy requirement shown in Table II, it is not allowed to give access to assets of value above 2 when the trust level is at *Low*. Hence, this is an example of an asset exposure rule that does not adhere to the asset exposure policy requirements. It turns out that this is the only breach of adherence in our example.

VIII. STEP 5: UPDATING THE TRUST POLICY TO REFLECT ITS REQUIREMENTS

The trust policy formulated above allows the consumers to report their power consumption just by giving their postal address, which is not hard to fake. If the asset exposure rule had required trust level *Medium* instead of *Low* the customer would be required to log on with their *Customer authentication* which is quite normal for this kind of service. In that case, adherence with respect to the asset exposure policy requirement would be ensured. This change is implemented by inserting *Medium* instead of *Low* in the second row of Table VIII.

IX. CONCLUSION

In this section we conclude by summarizing the contribution of this paper, discussing the results and presenting related work.

A. Contribution

We have presented DeSPoT, a method for the development and specification of policies for trust negotiation. The trust policy is linked to risk assessment through the requirements for the trust policy. The trust policy must adhere to the trust policy requirements to delimit the trust behavior within acceptable risk. The method supports negative (exposing) as well as positive (supporting) evidences, sensitive assets (credentials), separation between the trust formation and the asset exposure, static adherence check of the policy with respect to the trust policy requirements, and prospects that verify the properties of other prospects which are the general mechanism behind delegation of trust (recommendation). The method is built around a five step process. Our rule-based approach enables the development of a trust policy the enforcement of which ensures trust

negotiations within the limits of acceptable risk. The method is independent from specific trust negotiation protocols, and does not assume such protocols to be predefined.

Our focus has been to create an easy-to-understand language for trust policies with few details, nevertheless containing the most important trust mechanisms. The language is made to be understandable for people knowing the target (e.g. a company) at the business level, while being useful and understandable for those that develop and maintain the business application. In this way we are able to assemble a trust policy that contains both the risk and asset knowledge from the business level and the technical knowledge about different security technologies in one trust policy. Such a combination may reveal inconsistencies in the perception of the trust domain internally in the company.

B. Discussion

In the following we discuss to what extent our method fulfills the characteristics outlined in Section I.

- **Adherence:** In our method every trust policy rule is directly related to the requirements for the trust policy. Every trust policy rule can be statically checked with respect to adherence to its requirements.
- **Risk related:** The requirements for the trust policy define the acceptable combinations of trust and exposed asset value. These two notions correspond to the notions of likelihood and consequence within risk analysis. In our method each unique combination of these two parameters is characterized as either acceptable or not. The similar relation exists between the combinations of trust and the prospect value.
- **Generality:** The method supports the ability to define any important property of a prospect and verify that property by means of either local checks or other prospects. The latter allows us to represent recommendations. Every prospect is given a value reflecting its strength and used in the context of a risk and trust analysis when the assets are exposed.
- **Heterogeneity:** The method supports description of prospects in a straightforward way, enabling vast number of prospects which is the reality we face on the Internet. Every kind of prospect that may be reflected digitally may be described in our approach, and its prospect value reflects both its potential value for the target and how certain we are that it originated from the alleged partner.
- **Conceptual foundation:** Our notions of trust, asset value and prospect value are well founded.
- **Development guidance:** The process describes a structured way of developing a trust policy based on a risk assessment. It makes sure that the trust policy adheres to the trust policy requirements which ensures that the trust policy does not lead to a trust behavior that take on unacceptable risk.

C. Related Work

Winslett [20] propose a distinction between trust negotiation and security systems as that the latter operates within a clear security domain with internally defined users and credentials, whereas the trust domain involves open systems that make use of foreign credentials held by unknown users to prove their trustworthiness. As such, trust negotiation is a useful means to establish trust. However, this approach does not extract trust as a separate element in the trust negotiation that is subject to analysis in itself. Moreover, the implicit trust level is not explicitly linked to the risk that is inevitably involved in trust decisions. On this aspects, our method is clearly distinguished from established approaches to trust negotiation.

We adopt the definition of trust from Gambetta [7], who discusses how trust affects behavior in social and cooperative situations. Trust has later been discussed in detail by McKnight and Chervany [13] who classify different types of trust and discuss their meaning. Various surveys [8], [15], [11] discuss trust in relation to ICT systems, and even though they address many important aspects of trust, we see the need for the explicit perception of trust as a subjective probability.

The approach to system authentication presented in [16] implements trust negotiation as described in [2] to build trust. Approaches like Trust-X [2], [3], [17], TrustBuilder [12], [19] and Protune [4], [5], [6] are examples of policy based access control systems for automated trust negotiation. While these approaches make use of mechanisms for building trust, the involved trust level is only implicit and not extracted as an explicit element of value in itself. Yao et al. [21] present a value and privacy scoring for credentials and find an optimal exposure with minimal privacy and sufficient value to achieve access. The automated trust negotiation proposed in [18] emphasizes that negative evidences cannot be supported because an agent only controls what it sends and not what it receives and therefore opens for Denial of Service (DoS) attacks. For this reason there are very few that support negative evidence. We believe, however, that negative trust evidence is important in trust management. Revocation of certificates, banning of accounts and blocking of credit cards are examples of activities based on new information resulting in lower trust. Hence, not all digital trust functions are non-decreasing. In our approach every use of evidence is based on the trust in the partner providing it. To conduct a DoS attack through negative evidence, one must exploit misplaced trust and be able to pose as a trusted communication partner. If this posing is possible, then the trust in the evidence is overrated. In a trust policy as well as in a security policy the risk may be underrated as this is a possibility in all risk analysis. The vital thing is to be aware of this fact, and try to avoid it.

Some may argue that negative evidence is not relevant because a partner will never disclose negative evidence about itself. This is not true; even though they do not see the evidence as negative themselves, it may be perceived as negative by the target.

Several of the above mentioned approaches are introduced as contributions to the risk management domain while not explicitly extracting risk levels and relating them to the trust levels. In our approach we use combinations of consequence and trust to estimate risk, and categorize every combination as either acceptable or unacceptable as part of the policy requirements that are specified at the top level. Many aspects of our approach are mentioned in the approach of Alcade et al. [1], but it lacks explicit reflection on trust negotiation and the use of requirements to relate decisions to either acceptable or unacceptable risk.

Acknowledgments: The research on which this paper reports has partly been funded by the Research Council of Norway through the DIGIT (180052/S10) project, and partly by the European Commission through the NESSoS network of excellence.

REFERENCES

- [1] Baptiste Alcalde, Eric Dubois, Sjouke Mauw, Nicolas Mayer, and Saša Radomirović. Towards a decision model based on trust and security risk management. In *Proceedings of the Seventh Australasian Conference on Information Security - Volume 98, AISC '09*, pages 61–70. Australian Computer Society, Inc., 2009.
- [2] Elisa Bertino, Elena Ferrari, and Anna Squicciarini. Trust negotiations: Concepts, systems, and languages. *Computing in Science and Engineering*, 6:27–34, 2004.
- [3] Elisa Bertino, Elena Ferrari, and Anna Squicciarini. Trust-Xscr::: a peer-to-peer framework for trust establishment. *IEEE Transactions on Knowledge and Data Engineering*, 16(7):827–842, 2004.
- [4] Piero Bonatti, Juri Luca De Coi, Daniel Olmedilla, and Luigi Sauro. A rule-based trust negotiation system. *IEEE Transactions on Knowledge and Data Engineering*, 22:1507–1520, 2010.
- [5] Piero Bonatti and Daniel Olmedilla. Driving and monitoring provisional trust negotiation with metapolicies. In *Proceedings of the Sixth IEEE International Workshop on Policies for Distributed Systems and Networks (POLICY'05)*, pages 14–23. IEEE Computer Society, 2005.
- [6] Juri Luca De Coi, Daniel Olmedilla, Sergej Zerr, Piero Bonatti, and Luigi Sauro. A trust management package for policy-driven protection & personalization of web content. In *Proceedings of the 2008 IEEE Workshop on Policies for Distributed Systems and Networks (POLICY'08)*, pages 228–230, Washington, DC, USA, 2008. IEEE Computer Society.
- [7] Diego Gambetta. Can we trust trust? In *Trust: Making and Breaking Cooperative Relations*, chapter 13, pages 213–237. Department of Sociology, University of Oxford, 2000. Electronic edition.
- [8] Tyrone Grandison and Morris Sloman. A survey of trust in Internet applications. *IEEE Communication Surveys*, 2000.
- [9] Tormod Håvaldstud, Bjørnar Solhaug, and Ketil Stølen. Evaluation of a method for the analysis and development of policies for trust management. Technical Report A18834, SINTEF ICT, 2011.
- [10] International Organization for Standardization. *ISO 31000 Risk management – Principles and guidelines*, 2009.
- [11] Audun Jøsang, Claudia Keser, and Theo Dimitrakos. Can we manage trust? In *iTrust 2005*, volume 3477 of *LNCS*, pages 93–107. Springer, 2005.
- [12] Adam Lee, Marianne Winslett, and Kenneth Perano. TrustBuilder2: A reconfigurable framework for trust negotiation. In *Trust Management III*, volume 300 of *IFIP Advances in Information and Communication Technology*, pages 176–195. Springer, 2009.
- [13] D. Harrison McKnight and Norman L. Chervany. The meaning of trust. Technical Report MISRC Working Paper Series 96-04, University of Minnesota, Management Information Systems Research Center, 1996.
- [14] Atle Refsdal, Bjørnar Solhaug, and Ketil Stølen. A UML-based method for the development of policies to support trust management. In *Trust Management II - Proceedings of the 2nd Joint iTrust and PST Conference on Privacy, Trust Management and Security (IFIPTM'08)*, volume 263 of *IFIP*, pages 33–49. Springer, 2008.
- [15] Sini Ruohomaa and Lea Kutvonen. Trust management survey. In *iTrust 2005*, volume 3477 of *LNCS*, pages 77–92. Springer, 2005.
- [16] Jean-Marc Seigneur, Stephen Farrell, Christian Damsgaard Jensen, Elizabeth Gray, and Yong Chen. End-to-end trust starts with recognition. In *Security in Pervasive Computing*, volume 2802 of *LNCS*, pages 130–142, 2004.
- [17] Anna Squicciarini, Elisa Bertino, Elena Ferrari, Federica Paci, and Bhavano Thuraisingham. PP-trust-X: A system for privacy preserving trust negotiations. *ACM Trans. Inf. Syst. Secur.*, 10, 2007.
- [18] William H. Winsborough, Kent E. Seamons, and Vicki E. Jones. Automated trust negotiation. *DARPA Information Survivability Conference & Exposition*, 1:88–102, 2000.
- [19] M. Winslett, T. Yu, K.E. Seamons, A. Hess, J. Jacobson, R. Jarvis, B. Smith, and L. Yu. Negotiating trust in the Web. *Internet Computing, IEEE*, 6(6):30–37, 2002.
- [20] Marianne Winslett. An introduction to trust negotiation. In *iTrust 2003*, volume 2692 of *LNCS*, pages 275–283. Springer, 2003.
- [21] D. Yao, K.B. Frikken, M.J. Atallah, and R. Tamassia. Private information: To reveal or not to reveal. *ACM Transactions on Information and System Security (TISSEC)*, 12(1):6, 2008.



Technology for a better society
www.sintef.no