

Proposal for a Sentinels project
IPID: Integrated Policy-based Intrusion Detection
Version 1.0, June 29, 2004 (public version; budget details removed)

Title

Integrated Policy-based Intrusion Detection (IPID)

Address

Main applicant

Prof.dr. Roel Wieringa
University of Twente, EEMCS-IS
P.O. Box 217, NL-7500 AE ENSCHEDE, The Netherlands
Email roelw@cs.utwente.nl

Other applicants

Dr. Pascal van Eck, University of Twente, EEMCS-IS
Dr. Sandro Etalle, University of Twente, EEMCS-DIES
Prof.dr. Pieter Hartel, University of Twente, EEMCS-DIES
Prof.ir. Eddie Michiels, University of Twente, EEMCS-DIES
Dr. Andreas Wombacher, EEMCS-IS (from October 1st 2004)
Email {vaneck,etalle,pieter,michiels}@cs.utwente.nl

Embedding

The project will be part of the STW SENTINELS programme. It will be embedded in the CTIT (Centre for Telematics and Information Technology), which is one of the spearhead research institutes of the University of Twente). Within the CTIT, it will be part of the ISTRICE (Integrated Security and Privacy in a Networked World) programme.

Suitability

The scope of the proposal is fundamental research in intrusion detection (ID). The ambition is to extend secure systems engineering with conceptual modelling and monitoring techniques for ID that fully integrate digital and physical security. This complies with the SENTINELS aim of developing secure systems engineering.

Focus area

The focus is on area 2, specifically Security for vital ICT infrastructures in government and industry.

Resources

Of the total effort, 50% will be fundamental research and 50% method and tools.

Other funding applications

We have not submitted other grant applications to support this project.

Keywords

Computer security, I(CT) security, intrusion detection, patch management, policy based security, security event analysis, security event filtering, security incident management, security incident response, security management, security policy specification, vulnerability management.

1 Project summary

1.1 Research

Currently available intrusion detection tools monitor events at a relatively low level of abstraction. Due to the large number of events that occur at that level, and due to the low abstraction level, these tools are either ineffective (by generating a large number of false negatives) or inefficient (by generating a large number of false positives). The objective of IPID is to increase both effectiveness and efficiency of these tools by relating low-level events to a smaller number of events at a high level that are meaningful to the business. We will do this by:

1. Developing an enterprise security policy based method for the specification and deployment of ICT security mechanisms for intrusion detection and prevention;
2. Developing tools and mechanisms for high-level intrusion detection in realistic environments.

The research and engineering challenges are as follows:

1. Translating strategic ICT security policies, which are defined in business terms, into detailed tactical and operational ICT security policies, which are defined in technical terms. This translation enables the dynamic and static event analysis tools to determine whether or not security policy violations are occurring or have occurred.
2. Mapping large sets of low-level events to meaningful high-level events that might violate the high-level security policies.
3. Developing methods and tools for static and (near) real-time dynamic event analysis, which are effective (generating few false negatives) and efficient (generating few false alerts).

The project will result in the development of a set of methods, tools and techniques that make security management more business-oriented (based on enterprise-level security policies), more efficient (generating fewer incident reports that need to be taken care of) and more effective (allowing fewer incidents to go unnoticed).

1.2 Utilisation

The applicability of the developed methodologies and tools will be determined by testing them in two demanding environments: (1) A complex telecom environment, which consists of multiple ICT domains that are managed by different partners. This environment and the required telecom (security) know-how will be provided by TNO Telecom. (2) A large banking environment, where the security requirements are severe and where a comprehensive security policy needs to be enforced; in addition, the number of events to be analysed is huge due to high-volume processing and strict ICT security settings. This environment and the required (security) know-how will be provided by Rabobank. In both environments the effectiveness and efficiency of intrusion detection methods and tools is of paramount importance.

TNO Telecom and Rabobank will be involved in all stages of the research project, from the initial requirements analysis to the final assessments. Both TNO Telecom and Rabobank intend to integrate the results of this project into their (customer's) ICT infrastructures, so as to increase the level of security in an effective and efficient manner.

1.3 Summary in Dutch

1.3.1 Onderzoek

De op dit moment voor intrusion detection beschikbare tools monitoren de gebeurtenissen op een laag niveau van abstractie. Omdat op dit niveau grote hoeveelheden gebeurtenissen optreden en omdat het niveau van abstractie laag is, zijn dergelijke tools niet effectief (een grote hoeveelheid incidenten wordt niet gemeld) en/of niet efficiënt (een grote hoeveelheid valse alarmen wordt afgegeven). Doelstelling van IPID is om de effectiviteit en de efficiëntie van deze tools te verbeteren door de gebeurtenissen op laag niveau van abstractie te relateren aan een kleinere hoeveelheid gebeurtenissen op een hoog niveau van abstractie, welke relevant zijn voor de onderneming. We zullen dit als volgt doen:

1. Het ontwikkelen van methoden voor de specificatie en de implementatie van ICT beveiligingsmechanismen, welke zijn gebaseerd op het beveiligingsbeleid van de onderneming;
2. Het ontwikkelen van tools and mechanismen voor intrusion detection op een hoog niveau van abstractie en in realistische omgevingen.

Hierbij worden de volgende onderzoeks- en engineering uitdagingen onderkend.

1. De vertaling van het strategische ICT beveiligingsbeleid (gedefinieerd in termen van de ondernemingsdoelstellingen) naar tactisch en operationeel ICT beveiligingsbeleid (gedefinieerd in technische termen). Deze vertaalslag is nodig om de tools voor statische en dynamische analyse van gebeurtenissen vast te laten stellen of overtredingen van het beveiligingsbeleid al dan niet hebben plaatsgevonden of onderweg zijn.
2. Het afbeelden van grote hoeveelheden gebeurtenissen van een laag niveau van abstractie op significante gebeurtenissen op een hoog niveau van abstractie, voor zover deze met (mogelijke) overtredingen van het beveiligingsbeleid te maken hebben.
3. Het ontwikkelen van methoden en tools voor statische en (near) real-time dynamische analyse van gebeurtenissen, welke effectief zijn (slechts een beperkt aantal incidenten wordt niet gemeld) en efficiënt zijn (er wordt slechts beperkt aantal valse alarmen afgegeven).

Het project voorziet in de oplevering van een verzameling methoden, tools en technieken welke er toe zullen bijdragen dat het beveiligingsbeheer meer ondernemingsgericht is (d.w.z. gebaseerd op het ondernemingsbeveiligingsbeleid), meer efficiënt is (d.w.z. minder incidentmeldingen af te handelen) en meer effectief is (d.w.z. minder incidenten blijven onopgemerkt).

1.3.2 Utilisatie

De toepasbaarheid van de ontwikkelde methoden en tools zal worden getoetst in twee veeleisende omgevingen: (1) Een complexe telecom omgeving, bestaande uit meerdere ICT domeinen die door verschillende partners worden beheerd. Deze omgeving en de benodigde telecom (beveiligings) know-how wordt geleverd door TNO Telecom. (2) Een grote bank, waar zware eisen worden gesteld aan de beveiliging en waar een uitgebreid beveiligingsbeleid van toepassing is; verder moet in deze omgeving een zeer grote hoeveelheid gebeurtenissen worden afgehandeld, als gevolg van massale gegevensverwerking en zeer stricte beveiligingsinstellingen. Deze omgeving en de benodigde (beveiligings) know-how wordt geleverd door Rabobank. In beide omgevingen staat het belang van de effectiviteit en de efficiëntie van methoden en tools voor intrusion detection voorop.

TNO Telecom and Rabobank zullen in alle fasen van het onderzoek worden betrokken, vanaf de initiële inventarisatie van beveiligingseisen tot en met de uiteindelijke evaluaties. Zowel TNO Telecom als Rabobank hebben het voornemen om de resultaten van dit onderzoek toe te passen in (de) (hun) ICT infrastructuur (van hun klanten), om het beveiligingsniveau ervan te verhogen op een wijze die zowel effectief als efficiënt is.

2 Research groups

2.1 Present groups

The research will be carried out in a cooperation of two groups within the Faculty of Electrical Engineering, Mathematics and Computer Science (EEMCS) of the University of Twente (UT): the group Information Systems (IS), lead by Prof.dr. Roel Wieringa, and the group Distributed and Embedded Systems (DIES), lead by Prof.dr. Pieter Hartel. The team responsible for this research consists of Dr. Pascal van Eck (IS), Dr. Sandro Etalle (DIES) and Prof.ir. Eddie Michiels (DIES).

3 Scientific project description

3.1 Aims and Objectives

The main research question addressed in this proposal is whether the large number of reported security vulnerabilities, security patches provided by software vendors, and security events detected by ICT

components (applications, operating systems, middleware, firewalls, intrusion detection systems, etc.) relating to an ICT infrastructure can be automatically correlated and analysed for violations of the ICT security policy specifications, which are defined from the business point of view. The answer to this question should allow development of methods, tools and techniques that make security management more efficient (fewer incident reports that need to be taken care of) and more effective (fewer incidents that go unnoticed).

In our view, ICT security can be seen as a measure of how robust an ICT infrastructure is with respect to a particular ICT security policy. So security events are intimately related to the ICT security policy: Any event that is not forbidden by the ICT security policy is considered (explicitly or implicitly) not to pose a threat to the security of the ICT infrastructure.

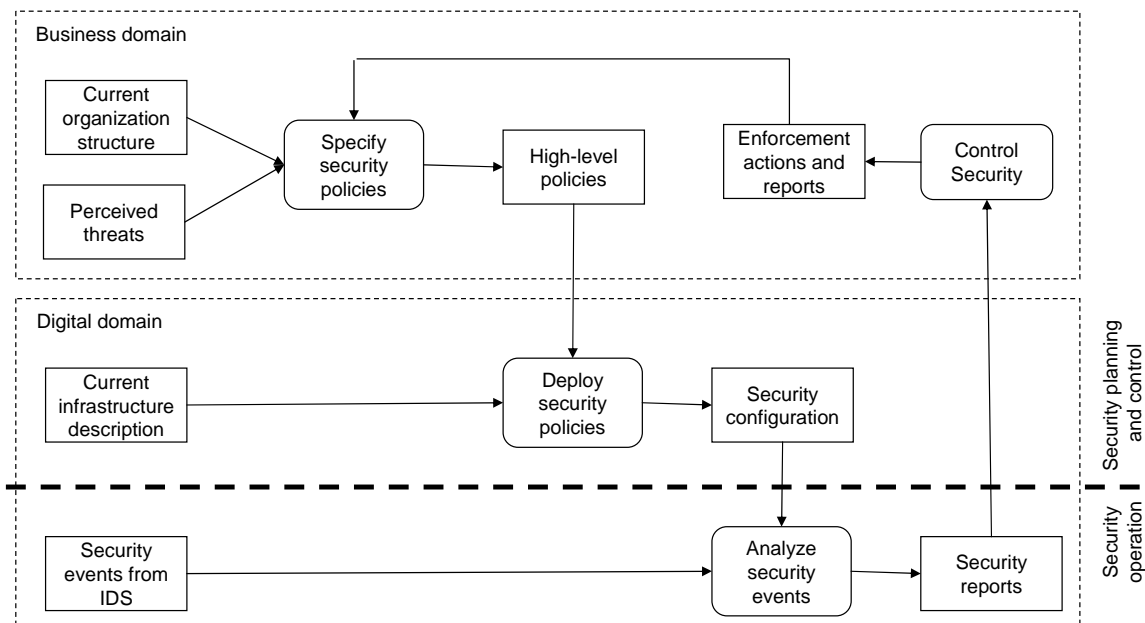


Figure 1: Security management.

Figure 1 focuses on the role of security policies in security management. We identify four processes (drawn as rounded boxes). Two processes take place during the design of the system, (*Specify* and *Deploy*) and two take place during run time (*Analyze* and *Control*). Two processes (*Specify* and *Control*) belong to the business domain and two processes (*Deploy* and *Analyze*) belong to the digital domain. The results of these processes are reports and other artefacts (drawn as rectangular boxes) that are input to further processes.

The first process *Specify security policies* takes place at design time in the *business domain*, where policies are defined at the strategic, the tactical and the operational level. These definitions are based on the current organizational structure (i.e., mission, departments, business processes), perceived threats, and information obtained via a feedback loop from security control activities. In the business domain, security policies are stated in terms of business entities and business events. This distinguishes the business domain from the *digital domain*, in which policies are defined in terms of software entities and software events. (We do not explicitly distinguish between hardware and software).

In the digital domain, we identify the second, design-time process, *Deploy security policies*, and a third, run-time process, *Analyze security events*. In the design-time process, people from the IT department configure security systems based on the current infrastructure present and the policies defined in the business domain. This may only consist of setting configuration parameters of, e.g., firewalls, but in many cases this process amounts to translating security policies from the business domain into policies in terms of software entities and software events.

The run-time part of the digital domain is concerned with intrusion detection. Current intrusion detection systems report large numbers of events at a low abstraction level. These systems are either ineffective (by generating a large number of false negatives) or inefficient (by generating a large number of false positives). Therefore, apart from conventional intrusion detection, an additional process is needed that analyzes security events and generates reports on a higher level of abstraction, in terms of the business domain. These reports

are then used by a fourth, runtime process, *Control security* (in the business domain), in which security officers perform enforcement actions based on these reports and provide input to the definition of updated security policies.

The objective of the IPID project is to increase the efficiency and effectiveness of security control as depicted in Figure 1 by supporting the four processes identified above in the following way:

- Objective 1: to provide a security policy based *method* for the *specification* and *deployment* of ICT security policies for intrusion detection and prevention;
- Objective 2: to develop *tools and techniques* for high-level *analysis* of security events and *control* of the security in realistic environments.

To reach these objectives, a number of research questions have to be answered. We structure these research questions according to the processes identified in Figure 1.

Business domain: processes *Specify security policies* and *Control security*

In the business domain, ICT security policies are specified in an informal manner, using natural language. Consequently, business ICT security policies at different abstraction levels are not integrated; consistency and completeness of a lower level specification with respect to a higher-level specification is usually not checked, neither automatically (the specifications are informal) nor manually (the specifications are large). We intend to develop guidelines for elaborating strategic business security policies into operationalised policies by collecting best practices already used by our business partners, as well as by making an inventory of guidelines based on existing literature in security as well as in safety. For example techniques such as hazard analysis may yield useful insights in desired security policies [Kle99]. Furthermore, our research can build upon results in requirements engineering, which is a main topic of the EEMCS-IS group.

With respect to *Security control*, it would be unrealistic to assume that the processing of events can be completely automated. Some human intervention will be required to make decisions as to whether or not specific actions (e.g. blocking certain communication paths) should be taken to protect the ICT infrastructure from ICT security policy violations. To be able to make “well-informed” decisions, decision support must be available in the form of guidelines and current, accurate, high-level information on observed security violations.

Most ICT infrastructures are evolving rapidly, as a result of organisational mergers/acquisitions, deployment of new application systems and new ICT components, etc. In addition, the environment in which these systems operate is changing e.g., by the discovery and exploitation of new vulnerabilities. As a result, the tools that are being used for filtering security vulnerabilities, security patches and security incidents need to be adaptable. Current ICT management best practices, for example ITIL (IT Infrastructure Library) and Cobit (Control Objectives for Information and related Technology) specify that very strict change management procedures should be followed when applying changes to an ICT infrastructure.

The following research questions need to be answered to address the issues raised in this section and provide a security policy based method for specification of ICT security policies:

Question 1: How to find and specify security policies that are integrated (i.e., consistent and complete from a strategic down to an operational level)? The security-related parts of process standards such as ITIL and Cobit will be relevant starting points here.

Question 2: How to map security policies stated in terms of business events to policies stated in terms of software events and back again? The challenge is both informal (how to write the definitions effectively) as well as formal (how to recognize a business event in a trace of low-level events).

Question 3: It will be a challenge to find practical means to manage the hierarchy of specifications. Can goal-oriented requirements engineering techniques [Lamsweerde95] [Dardenne93] [Anton98] be used to link business-level security goals to low-level operational security policies? Can traceability tools known from requirements engineering, such as Telelogic DOORS, be used to support security management?

Question 4: How to specify business procedures in such a way that rapid response to reported security vulnerabilities, security patches and security incidents is possible, e.g. by taking controlled shortcuts (or expedite/automate) strict management procedures? In practice, there is a trade-off between costs and benefits

of acting or not acting on a violation and so we will investigate introducing a measure of severity of a security breach.

Digital domain: process *Deploy security policies*

Deployment of security policies amounts to translating security policies from the business domain to configuration of the security systems that are part of the ICT infrastructure. The configuration of ICT components (including configuration changes as a result of reported security vulnerabilities, security patches and security incidents) will have to be derived from these business-level ICT security policy specifications (this is commonly referred to as *policy based* configuration/patch management etc.). This means that traceability should be established between policies at the business level and at the software level. In a modern and extensive infrastructure, this means that formal rules are defined in terms of software entities and events that implement the business-level policies. The challenge here is to manage this process so that the relation between these rules and business-level policies is traceable and that it is provable that these rules indeed implement the policies specified in the business domain. Tracing business-level security policies against software-level ones is informal even though tool-supported. At the lowest level of software-monitorable events, though, we will formalize security policies and use model checkers for checking their correctness against the next higher-level policy specification. Feasibility of this approach has been shown by us [Har04]. This work is based on a combination of techniques developed by Alan Mycroft and his PhD students from Cambridge, UK, for dynamic security policies in the logical domain (security policy abstraction [Mad03]) and from the work of Cheng et al [Che03] on security patterns.

Based on our earlier work, we identify the following research questions:

Question 5: How do these low level operationalised software management instructions relate to the high-level strategic business policies?

Question 6: In our earlier work [Eshuis02a], a tool for the specification of processes using UML activity diagrams is used as a front-end for a model checker. In what way can this be extended to the domain of security policies?

Digital domain: process *Analyze security events*

Intrusion detection systems monitor low-level software events that are not directly related to meaningful high-level business events. Since there are large numbers of low-level events, intrusion detection software generates large numbers of false negatives (intrusions not recognized as such) and false positives (reports about non-intrusions). To automate the processing of reported security vulnerabilities, security patches and security events, the low-level software events should be related to meaningful events in the business domain. One of the main objectives of the proposed project is to develop a tool that is able to provide this relation automatically.

In the proposed project, we will focus on using model checking as the starting point for such a tool as follows: A low-level security event in the digital domain often consists of multiple inter-related smaller events, which may be separated in time and space. It is therefore necessary to keep a trace of related events (*event trace*). These event traces will be checked in (near) real time whether or not they might progress in such a way that the ICT security policy is violated. We envision two approaches to use model checking for this:

- Encode low-level security policies as an automaton; check whether an event trace is a path through this automaton. This approach is an extension of our previous work [Har04], but see also [Rog01]. This approach is akin to, but not the same as the problem of dynamic integrity constraint monitoring studied in the 1980s [Lip87,Lip90]. We will take those techniques as our point of departure for developing a tool for monitoring policy violations.
- Identify a fragment of an event trace with a minimal automaton that generates this trace; check whether this automaton is a model of the low level security policies encoded as logic formulae. Counter examples found by the model checker can be interpreted as event traces which show behaviour that violates the ICT security policy. These counter examples are related to event traces that are observed. Given an initial state, the model checker is capable of finding all the counter examples, so it should be possible to find the counter example that corresponds to a particular event trace.

Given the ICT security policies at the software and at the business levels and traceability relations between them, and given an event trace, an extended intrusion detection system should be able to return the set of policies (at the business level) that is violated by this event trace. Sometimes it will be impossible to continue an event sequence without violating at least one ICT security policy element. In that case, the least harmful event sequence can be selected. The engineering challenge is to deal with very large numbers of low-level events and still yield responses in near real time. There is hope of achieving this if the automata involved are sufficiently simple.

Question 7: Which of the two approaches to use model checking outlined above is the most suitable? Is it possible to do the proposed analysis in near real time? If not, how to perform off-line analysis of a given event trace to determine whether a combination of events has occurred that constitutes a violation of the ICT security policy?

Question 8: Following traceability relations backwards to present security violations in terms of the business domain is probably not enough. How can information from configuration databases, process definitions in workflow management systems and/or database schemas be used to give more meaning to event traces that violate a security policy?

3.2 Personnel and equipment

We request funding for two PhD students, one of whom will be employed in the IS group (promotor: Wieringa, daily supervisor: Van Eck and Wombacher), and one of whom will be employed in the DIES group (promotors: Hartel and Michiels, daily supervisor: Etalle).

The IS PhD student will mainly focus on the business domain, but also on the tools and techniques needed for security event abstraction. He/she will be working to answer research questions 1-4.

The DIES PhD student will mainly focus on the digital domain and cooperate with the IS PhD student on the tools and techniques needed for security event abstraction. The DIES PhD student will be trying to answer research questions 5-8.

3.3 Programme and timetable

Year 1:

- The DIES PhD student will study the state of the art in ICT security policy specification.
- The IS PhD student will study the state of the art in dynamic event correlation and analysis and static event analysis (decision support).
- The PhD students will perform an initial requirement analysis, partly based on input provided by TNO Telecom and Rabobank.

Deliverables:

- State of the art reports, first papers
- Requirements analysis report.

Year 2:

- The DIES PhD student will develop a methodology and prototype tools for ICT security policy specification.
- The IS PhD student will develop a methodology and prototype tools for dynamic event correlation and analysis and static event analysis.
- Refinement of requirements analysis.

Deliverables:

- Methodologies;
- Prototype toolset;
- Requirements specification, papers

Year 3:

- The DIES PhD student will refine the methodology and supervise the development of tools (by Master students) for ICT security policy specification.
- The usability of the ICT security policy specification methodology and tools will be evaluated by means of a case study or pilot implementation at Rabobank/TNO Telecom.

- The IS PhD student will refine the methodology and supervise the development of tools (by Master students) for dynamic event correlation and analysis and static event analysis.
- The usability of the methodology and tools for dynamic event correlation and analysis and static event analysis will be evaluated by means of a case study or pilot implementation at Rabobank/TNO Telecom.

Deliverables:

- toolset;
- evaluation reports, papers

Year 4:

- Both PhD students will finish ongoing work and write a PhD thesis.

Deliverables:

- PhD theses, papers

3.4 Track record

The research will be carried out by Information Systems (IS) and Distributed and Embedded Systems (DIES) research groups.

IS will contribute its expertise in requirements engineering, software specification, policy specification and business-ICT alignment, The IS group performs research in requirements engineering and software specification techniques [Wie96] [Wie03]. Members of the group have done earlier research in policy specification [Mey93] [Mey98] [Wie89]. Current research includes the investigation of the practice of business-ICT alignment [Wie03a] [Eck04].

DIES has a strong expertise and experience in secure systems engineering as reported in nearly 100 publications (See <http://dies.cs.utwente.nl/~pieter/refs/DIES-Security.html>). We participate in 9 externally funded security projects, including a project on security incident response: project SIRE, sponsored by Rabobank, and two projects on security policies: project Inspired, sponsored by EU-FP6 and project ACCOUNT funded by NWO (For all DIES projects see <http://dies.cs.utwente.nl/research>).

In previous work, the IPID team presents SPIN models of four case studies where considerations from the physical domain is shown to spoil security policies that are appropriate in the purely digital domain . In each case the model captures both the system of interest and its security policy. The model is then formally checked against a property that represents a design principle from the problem domain. The model checking activity shows many examples of policies that are too weak to cope with integrated domains [Har04].

3.5 Related research

3.5.1 Security policy specification and model checking

Security policies are important both in the physical and the logical world; the problem is that they have hardly been studied in an integrated fashion. By a *policy* we mean ‘a rule that defines a choice in the behaviour of a system’ [Dam01]. A *security policy* rules out behaviour ‘that has been deemed unacceptable’ [Sch00]. A *spatial security policy* constrains this further by ruling out behaviour tied to particular locations [Sco03]. Our view on policies is broader than that of other authors because we consider not only the logical world, but also the physical world.

Ahmed and Tripahti's [Ahm03] describe a translation of security policies from XML based collaboration descriptions into Promela, and Cheng et al translate of security patterns into Promela [Che03]. These could be a useful starting point for our model checking activity (Promela is the modelling language of the Model Checker Spin [Hol04]). We will also build upon the application of organisational control principles to security policy specification by Schaad and Moffett [Sch01x].

Intrusion detection

Worldwide, there are several research groups that perform research in the area of intrusion detection and prevention. Unlike IPID, these research projects focus on technical issues and do not take the business

perspective into account. Some representative examples of existing intrusion detection/prevention research are given below.

Seminal research

The Information Security Center of Lucent Laboratories (former AT&T Laboratories) in New Jersey played a paramount role in initial research efforts in the area of intrusion detection [Amo98] [Amo99].

“Traditional” research approaches

There are many research groups that perform research in the area of intrusion detection/prevention. The most well known is undoubtedly Stanford Research Institute (SRI), which is involved in several intrusion detection/prevention research projects. Early research at SRI led to the development of a prototype Intrusion Detection Expert System (IDES). The IDES prototype evolved into a production quality intrusion detection system called Next-Generation Intrusion Detection Expert System (NIDES). Current SRI research focuses on the [Event Monitoring Enabling Responses to Anomalous Live Disturbances](#) (EMERALD) project, which extends the NIDES concept to distributed computing environments. This effort includes profile-based analysis, signature-based analysis and localized results fusion with automated response capability [Neu99]. Another well-known research project is NetSTAT (Network-based State Transition Analysis Tool), which is performed at the University of California at Santa Barbara [Eck02].

Data mining approach

MADAM ID (Mining Audit Data for Automated Models for Intrusion Detection) and JAM are research projects from Columbia University. The objective is to use data mining techniques for adaptive intrusion detection systems [Fan01]. This approach is relevant for static intrusion event analysis.

Adaptive/agent based approach

Modern approaches take into account the complexity and scalability requirements of intrusion detection/prevention for large distributed ICT infrastructures. AAIRS (Adaptive Agent-based Intrusion Response System) and AHA! IDS (Adaptive Hierarchical Agent-based Intrusion Detection System) are projects undertaken by Texas A&M University and the United States Military Academy. Their objective is to develop adaptive agent based intrusion detection systems [Nas01]. IDIOT (Intrusion Detection In Our Time) and AAFID (Autonomous Agents for Intrusion Detection) are projects of the Center for Education and Research in Information Assurance and Security of Purdue University, which use artificial intelligence genetic approaches and agent architectures for detecting coordinated and sophisticated attacks on large and distributed ICT infrastructures [Ker02].

Standardisation

The Common Intrusion Detection Framework (CIDF) is a [Defense Advanced Research Projects Agency](#) (DARPA) effort to develop protocols and application programming interfaces so that intrusion detection research projects can share information and resources and so that intrusion detection components can be reused in other systems [Fei00]. Some of the ideas involved in CIDF have encouraged the creation of an [Internet Engineering Task Force](#) (IETF) working group, named the [Intrusion Detection Working Group](#) (IDWG) [IDWG].

With a few exceptions, intrusion detection/prevention research efforts, including the ones listed above, are limited to intrusion detection/prevention systems per se. Furthermore, the focus of their research is entirely on (specific) technical issues while the business perspective is not at all taken into account. In contrast, the IPID proposal represents a novel way forward, as we intend to follow an integrated approach and to take the business view fully into account, by relating security relevant events to a formally specified ICT security policy. We believe that IPID is unique in doing so, but this does not imply that we have to start from scratch as there are many bits and pieces of existing research that we could encompass in our project (these will be identified during the first year, see § 3.3)

4 Utilisation

4.1 Practical challenges and solutions

Most organisations recognise the importance of ICT security but experience difficulty in implementing security properly, which requires processing of huge amounts of vulnerability disclosures, security patches/updates and security events (especially IDS events). For each of these, one must decide whether or

not it is necessary to take action. In order to do this effectively and efficiently, an appropriate methodology and automated tools are needed. Currently available (IDS) tools are either ineffective (by generating a large amount of false negatives) or inefficient (by generating a large amount of false positives).

The solution is to relate vulnerability disclosures, security patches/updates and security events to a specified ICT security policy by automated means. The ICT security policy needs to be defined at different layers of abstraction; at the highest layer, the security policy is expressed in business terms.

We recognise that it will not be probable to fully automate the processing of events and therefore, a decision support system needs to be made available for human decision making.

4.2 Identification of the users

Rabobank

Rabobank is one of the largest Dutch banks. Rabobank's business is completely dependent on their ICT infrastructure and ICT security is of paramount importance. Rabobank will contribute to the specification of requirements and to the validation of the methodology and tools that result from this project. Rabobank intends to use the project results to improve the security of its ICT infrastructure.

TNO Telecom

TNO Telecom performs research for telecom providers and vendors. In the telecom world, an infrastructure is often composed of multiple parts, which are managed by different parties. TNO Telecom is concerned how the security of such an integrated infrastructure has to be managed by the different partners and how it must be certified and controlled by an independent third party. TNO Telecom will bring in their experience and their security laboratory facilities and tools for high speed, real time network emergency detection and response, with a view to extend these facilities and tools with the results of the project.

4.3 Implementation

The research project has the active support of its industrial users. These companies will have staff participating in the work and/or will make their facilities and products available to the research team. This involvement will improve the adoption of the results in commercial products and will act as an additional motivation for the researchers. We will maintain close contacts with TNO Telecom and Rabobank throughout the course of the project. Initially, TNO Telecom and Rabobank will provide input for the requirements analysis. The requirements analysis will be refined and the resulting requirements specification will be reviewed by TNO Telecom and Rabobank.

Cases studies and/or pilot projects will be performed at TNO Telecom and Rabobank, to determine the suitability of the methodologies and tools that are being developed.

4.4 Past Performance

The UT (Hartel) and KUN (Jacobs) have initiated a major national funding program, SENTINELS <http://www.sentinels.nl>, which aims to foster security research in the Netherlands. This shows that we set the national research agenda in security.

The UT (Etalle) is one of the founders of SAFE-NL (the platform for Security: Applications, Formal aspects and Environments in the NetherLands).

Wieringa is currently chair of the Steering Committee of the IEEE International Requirements Engineering Conference. He is also member of the board of the NAF (<http://www.naf.nl>), a forum of Dutch business users, suppliers, and research institutes to promote IT architecture in the Netherlands. He was a co-founder of the biannual workshop on Deontic Logic in Computer Science [Meyer93].

CTIT, which manages all ICT and related research of the UT, has funded a strategic research orientation on Conformable Privacy and Security (ISTRICE), see http://www.ctit.utwente.nl/research_program/sro/security/. ISTRICE is lead by Sandro Etalle.

4.4.1 EEMCS-DIES

The Distributed and Embedded Systems research group has developed and is developing a number of security components in various successful projects as follows:

1. DIES participates in the EU Inspired project, which develops the next generation smart card.
2. Development of applications based on hardware tokens [Cho02a]
3. Developing policy based methods and tools to reason about accountability in e-commerce protocols in the project Account (with VU, CWI), funded by NWO [Cor04b].
4. Development of CoProVe [Cor02], which is likely to be the fastest tool for the verification of security protocols.
5. Developing the security component in an ad-hoc sensor network in the context of the European project EYES (with Infineon, Nedap) [Law03].
6. Developing a novel transacted smart card memory manager with Sun Microsystems in Cupertino (USA) [Har00b] [Pol02].
7. Developing a simulation tool for side channel attacks [HVVVW:sec]
8. We have studied the added value of Computer Emergency Response Teams [Haf02] [Haf03a] [Haf03b].

Item 1 shows our active involvement in the international research arena. Item 2 shows that our security protocol verifier is highly competitive. Items 3-7 show that we have the relevant experience in variety of security areas intrusions are a major threat. Item 8 shows our involvement with CERTs.

4.4.2 EEMCS-IS

The IS group is performing methodological research in various projects.

1. In the GRAAL project, alignment of IT infrastructure and application architecture to business architecture is investigated by means of case study research in large organizations in the government and finance sectors [Wie03a,Eck04]. The project is funded by the Telematics Institute.
2. In the TCM project, a collection of user-friendly graphical editors is developed to support information system specification using all techniques from structured analysis and the UML (www.cs.utwente.nl/~tcm). TCM is part of the SuSe and Debian Linux distributions
3. In an NWO-funded project, we have developed a formal execution semantics for UML activity diagrams and interfaced the activity diagram editor TCM with nuSMV to model check activity diagram properties using nuSMV [Eshuis02a].
4. We have a long-standing research program in requirements engineering [Wie96]. Current research includes investigating patterns in evolutionary requirements engineering by means of case study research. This is funded by the BITE program of the UT.

5 Intellectual Property

5.1 Contracts

The UT has no contracts that limit the research in the area of this research proposal.

5.2 Patents

There are a significant number of patents on intrusion detection. During the first phase of the project we will analyse the existing patents carefully. Where possible we will apply for patents in collaboration with our industrial partners.

6 Budget

All costs specified below are in EUROS.

6.1 Personnel

We apply for the funding of two PhD students during 4 years (2 full-time equivalents per year). The personnel costs are listed in table 1.

Table 1

<i>function</i>	<i>cost</i>	<i>comment</i>
DIES PhD student		methodology/tools for ICT security policy specification
IS PhD student		methodology/tools for event analysis

6.2 Consumables and domestic travel

The consumables and domestic travel costs (mainly to visit our industrial partners in Groningen and Utrecht) are listed in table 2.

Table 2

<i>item</i>	<i>cost</i>	<i>comment</i>
domestic travel		Participation in project meetings (Utrecht, Delft etc).
software licenses		commercial software packages, e.g. for intrusion detection systems

6.3 Foreign travel

The costs of foreign travel to international conferences, workshops and summer schools are listed in table 3.

Table 3

<i>item</i>	<i>cost</i>	<i>comment</i>
foreign conferences participation/visits		12 international trips

6.4 Computer Equipment

The computer equipment costs are listed in table 4. The desktops and laptops are provided by the university. A powerful dedicated server is needed to support experiments with intrusion detection systems, and to support the model checking activity.

Table 4

<i>item</i>	<i>cost</i>	<i>comment</i>
2 powerful desktop computers	PM	one for each PhD student
2 laptops	PM	one for each PhD student
dedicated server		

6.5 Support from users

The costs contributed by users are listed in table 5.

Table 5

<i>user</i>	<i>cost</i>	<i>comment</i>
TNO Telecom		Approx 0,2 fte per year
Rabobank		Approx 0,15 fte per year

6.6 Budget summary

Table 6

<i>cost type</i>	<i>year 1</i>	<i>year 2</i>	<i>year 3</i>	<i>year 4</i>
domestic travel				
software license				
foreign travel				
desktop computers				
laptop computer				
dedicated server				
total costs				

Table 7

<i>cost type</i>	<i>year 1</i>	<i>year 2</i>	<i>year 3</i>	<i>year 4</i>
PhD EEMCS-DIES (1 fte/year)				
PhD EEMCS-IS (1 fte/year)				
total costs				

Table 8

	<i>year 1</i>	<i>year 2</i>	<i>year 3</i>	<i>year 4</i>
contribution TNO Telecom				
contribution Rabobank				
total contributions				

7 References

7.1 References from the group

[Cho02a] C. N. Chong, Z. Peng, and P. H. Hartel. Secure audit logging with Tamper-Resistant hardware. In D. Gritzalis, S. De Capitani di Vimercati, P. Samarati, and S. K. Katsikas, editors, 18th IFIP TC11 Int. Conf. on Information Security, Security and Privacy in the Age of Uncertainty (SEC), pages 73-84, Athens, Greece, May 2003. Kluwer Academic Publishers, Boston. <http://www.ub.utwente.nl/webdocs/ctit/1/00000099.pdf>.

[Cor02] R. Corin and S. Etalle. An improved constraint-based system for the verification of security protocols. In M. V. Hermenegildo and G. Puebla, editors, 9th Int. Static Analysis Symp. (SAS), volume LNCS 2477, pages 326-341, Madrid, Spain, Sep 2002. Springer-Verlag, Berlin. <http://www.ub.utwente.nl/webdocs/ctit/1/00000096.pdf>.

[Cor04b] R. Corin, S. Etalle, J. I. den Hartog, G. Lenzini, and I. Staicu. A logic for auditing accountability in decentralized systems. Technical report CTIT-TR 04-27, Centre for Telematics and Information Technology, Univ. of Twente, The Netherlands, Jul 2004.

[HVVVW:sec] J. I. den Hartog, J. Verschuren, E. P. de Vink, J. Vos, and W. Wiersma. PINPAS: A tool for power analysis of smartcards. In D. Gritzalis, S. De Capitani di Vimercati, P. Samarati, and S. Katsikas, editors, 18th IFIP TC11 Int. Conf. on Information Security, Security and Privacy in the Age of Uncertainty (SEC), pages 453-457, Athens, Greece, 2003. Kluwer Academic Publishers, Boston. <http://www.wkap.nl/prod/b/1-4020-7449-2>.

[Eshuis02a] R. Eshuis and R.J. Wieringa. Verification support for workflow design with UML activity graphs. In 24th International Conference on Software Engineering (ICSE 2002), pages 166-176, 2002.

[Haf02] W. Hafkamp. De 'mission impossible' van CERT's. *Informatiebeveiliging*, 2(6):page numbers?, Sep 2002.

[Haf03b] W. Hafkamp. Eindrapportage haalbaarheidsonderzoek 'interbancair computer security information sharing and analysis center' (ISAC). Technical report, Rabobank/Interpay, Utrecht, May 2003.

[Haf03a] W. Hafkamp. Thema cybercrime: gewaarschuwde organisatie telt voor twee. *Controlling*, Kluwer, Deventer, 18(6):19-24, 2003.

[Har00b] P. H. Hartel, M. J. Butler, E. K. de Jong, and M. Longley. Transacted memory for smart cards. In J. N. Oliveira and P. Zave, editors, 10th Formal Methods for Increasing Software Productivity (FME), volume LNCS 2021, pages 478-499, Berlin, Germany, Mar 2001. Springer-Verlag, Berlin. <http://www.dsse.ecs.soton.ac.uk/techreports/2000-9.html>.

[Har04] P. H. Hartel, P. van Eck, S. Etalle, and R. J. Wieringa. Modelling mobility aspects of security policies. Technical report TR-CTIT-04-06, Centre for Telematics and Information Technology, Univ. of Twente, The Netherlands, Jan 2004. <http://www.ub.utwente.nl/webdocs/ctit/1/000000ea.pdf>.

[Law03] Y. W. Law, R. Corin, S. Etalle, and P. H. Hartel. A formally verified decentralized key management architecture for wireless sensor networks. In M. Conti, S. Giordano, E. Gregori, and S. Olariu, editors, 4th IFIP TC6/WG6.8 Int. Conf on Personal Wireless Communications (PWC), volume LNCS 2775, pages 27-39, Venice, Italy, Sep 2003. Springer-Verlag, Berlin. <http://www.ub.utwente.nl/webdocs/ctit/1/000000b7.pdf>.

[Mey93] J.-J. Ch. Meyer and R. J. Wieringa (eds.). *Deontic Logic in Computer Science: Normative System Specification*. Wiley., 1993.

[Mey98] J.-J. Ch. Meyer, R. J. Wieringa, and F. P. M. Dignum. The role of deontic logic in the specification of information systems. In J. Chomicki and G. Saake, editors, *Logics for Databases and Information Systems*, pages 71-115, Dagstuhl, Germany, 1998. Kluwer.

[Meyer93] J.-J.Ch. Meyer and R.J. Wieringa, editors. *Deontic Logic in Computer Science: Normative System Specification*. Wiley, 1993.

[Pol02] E. Poll, P. H. Hartel, and E. K. de Jong. A Java reference model of transacted memory for smart cards. In 5th Int. IFIP wg 8.8 Conf. Smart card research and advanced

application (CARDIS), pages 75-86, San Jose, California, Nov 2002. Usenix Association, Berkeley, California. <http://www.ub.utwente.nl/webdocs/ctit/1/00000083.pdf>.

[Eck04] P. A. T. van Eck, H. Blanken, and R. J. Wieringa. Project GRAAL: Towards operational architecture guideline. *Int. J. of Cooperative Information Systems*, accepted for publication, to appear in 2004.

[Wie96] R. J. Wieringa. *Requirements Engineering: Frameworks for Understanding*. Wiley, 1996.

[Wie03] R. J. Wieringa. *Design Methods for Software Systems: Yourdon, Statemate and the UML*. Morgan Kaufmann, 2003. <http://www.mkp.com/dmrs>.

[Wie03a] R. J. Wieringa, H. M. Blanken, M. M. Fokkinga, and P. W. P. J. Grefen. Aligning application architecture to the business context. In J. Eder and M. Missikoff, editors, *Conf. on Advanced Information System Engineering (CAiSE)*, volume LNCS 2681, pages 209-225, Klagenfurt/Velden, Austria, 2003. Springer-Verlag, Berlin.

[Wie89] R. J. Wieringa, J.-J. Ch. Meyer, and H. Weigand. Specifying dynamic and deontic integrity constraints. *Data and Knowledge Engineering*, 4:157-189, 1989

7.2 Other references

[Ahm03] T. Ahmed and A. R. Tripathi. Static verification of security requirements in role based CSCW systems. In *8th ACM Symp. on Access Control Models and Technologies (SACMAT)*, pages 196-203, Como, Italy, Jun 2003. ACM Press, New York. <http://doi.acm.org/10.1145/775438>

[Amo98] E. Amoroso and R. Kwapniewski. Selection criteria for intrusion detection systems. In *14th Annual Computer Security Applications Conference (ACSAC)*, pages 280-292, Scottsdale, Arizona, Dec 1998. IEEE Computer Society.

[Amo99] E. G. Amoroso. In *Intrusion Detection: An Introduction to Internet Surveillance, Correlation, Traps, Trace back and response, Intrusion*. Net Books, New Jersey, 1999.

[Anton98] A.I. Antón and C. Potts. The use of goals to surface requirements for evolving systems. In *International Conference on Software Engineering (ICSE'98)*, pages 157-166. IEEE Computer Society, 1998.

[Che03] B. H. C. Cheng, S. Konrad, L. A. Campbell, and R. Wassermann. Using security patterns to model and analyze security requirements. In C. Hietmeyer and N. Mead, editors, *Int. Workshop on Requirements for High Assurance Systems (RHAS)*, pages 13-22, Monterey, California, Sep 2003. Software Engineering Institute, Carnegie mellon Univ. <http://www.sei.cmu.edu/community/rhas-workshop/rhas03-proceedings.pdf>.

[Dam01] N. Damianou, N. Dulay, E. Lupu, and M. Sloman. The ponder policy specification language. In M. Sloman, J. Lobo, and E. Lupu, editors, *Int. Workshop on*

Policies for Distributed Systems and Networks (POLICY), volume LNCS 1995, pages 18-38, Bristol, UK, Jan 2001. Springer-Verlag, Berlin.

[Dardenne93] A. Dardenne, A. van Lamsweerde, and S. Fickas. Goal-directed requirements acquisition. *Science of Computer Programming*, 20:3-50, 1993.

[IDWG] M. Erlinger and S. Staniford-Chen (Eds.). *Intrusion Detection Exchange Format*. Internet Engineering Task Force, Nov 2003. <http://www.ietf.org/html.charters/idwg-charter.html>.

[Fan01] W. Fan, M. Miller, S. J. Stolfo, W. Lee, and Ph. K. Chan. Using artificial anomalies to detect unknown and known network intrusions. In N. Cercone and X. Wu T. Y. Lin, editors, *IEEE Int. Conf. on Data Mining*, pages 123-130, San Jose, California, Nov 2001. IEEE Computer Society.

[Fei00] R. Feiertag, S. Rho, L. Benzinger, S. Wu, T. Redmond, C. Zhang, K. Levitt, D. Peticolas, M. Heckman, S. Staniford, and J. McAlerney. Intrusion detection inter-component adaptive negotiation. *Computer Networks*, 34(4):605-621, Oct 2000. [http://dx.doi.org/10.1016/S1389-1286\(00\)00137-7](http://dx.doi.org/10.1016/S1389-1286(00)00137-7).

[Hol04] G. J. Holzmann. *The SPIN Model Checker: Primer and Reference manual*. Pearson Education Inc, Boston Massachusetts, 2004.

[Ker02] F. Kerschbaum, E. H. Spafford, and D. Zamboni. Using internal sensors and embedded detectors for intrusion detection. *J. of Computer Security*, 10(1/2):23-70, 2002.

[Kle99] T. Kletz. *Hazop and Hazan: Identifying and Assessing Process Industry Standards*. Taylor & Francis, 1999.

[Ko01] C. Ko. Logic induction of valid behavior specifications for intrusion detection. In *21th IEEE Symposium on Security and Privacy (S&P)*, pages 142-153, Berkeley, California, May 2000. IEEE Computer Society. <http://www.computer.org/proceedings/s&p/0665/06650142abs.htm>.

[Lamsweerde95] A. van Lamsweerde, R. Darimont, and P. Massonet. Goal-directed elaboration of requirements for a meeting scheduler: problems and lessons learnt. In *2nd IEEE International Symposium on Requirements Engineering (RE'95)*, pages 194-203, 1995.

[Lip90] U. W. Lipeck. Transformation of dynamic integrity constraints into transaction specifications'. *Theoretical Computer Science*, 76:115-142, 1990.

[Lip87] U. W. Lipeck and G. Saake. Monitoring dynamic integrity constraints based on temporal logic'. *Information Systems*, 12:255-269, 1987.

[Mad03] A. Madhavapeddy, A. Mycroft, D. Scott, and R. Sharp. The case for abstracting security policies. In H. R. Arabnia and Y. Mun, editors, *Int. Conf. on Security and Management (SAM)*, volume 1, pages 156-160, Las Vegas, Nevada, Jun 2003. CSREA Press.

- [Nas01] D. A. Nash and D. Ragsdale. Simulation of self-similarity in network utilization patterns as a precursor to automated testing of intrusion detection systems. *IEEE Transactions on Systems, Man, and Cybernetics, Part A*, 31(4):327-331, Jul 2001
- [Nec97] G. C. Necula. Proof-carrying code. In 24th Principles of programming languages (POPL), pages 106-119, Paris, France, Jan 1997. ACM, New York.
- [Neu99] P. G. Neumann and Ph. A. Porras. Experience with EMERALD to date. In 1ST Workshop on Intrusion Detection and Network Monitoring, pages 73-80, Santa Clara, California, Apr 1999. Usenix Association, Berkeley, California.
- [Rog01] M. Roger and J. Goubault-Larrecq. Log auditing through Model-Checking. In 14th IEEE Computer Security Foundations Workshop, pages 220-236, Cape Breton, Nova Scotia, Canada, 2001. IEEE Computer Society Press, Los Alamitos, California.
- [Sch01x] A. Schaad and J. D. Moffett. The incorporation of control principles into access control policies (extended abstract). In Workshop on Policies for Distributed Systems & Networks (Policy), Bristol, UK, 2001. Unknown Publisher. <http://www-users.cs.york.ac.uk/jdm/pubs/ICPACP.pdf>.
- [Sch00j] F. B. Schneider. Enforceable security policies. *ACM Transactions on Information and System Security*, 3(1):30-50, Feb 2000. <http://doi.acm.org/10.1145/353323.353382>.
- [Sco03] D. Scott, A. Beresford, and A. Mycroft. Spatial security policies for mobile agents in a sentient computing environment. In M. Pezzè, editor, 6th Fundamental Approaches to Software Engineering (FASE), volume LNCS 2621, pages 102-117, Warsaw, Poland, Apr 2003. Springer-Verlag, Berlin. <http://www.springerlink.com/link.asp?id=nyxyyrlkbe5c5acc>.
- [Sek01] R. Sekar, C. R. Ramakrishnan, I. V. Ramakrishnan, and S. A. Smolka. Model-Carrying code (MCC): A new paradigm for Mobile-Code security (extended abstract). In Workshop on new security paradigms, pages 23-30, Cloudcroft, New Mexico, 2001. ACM Press, New York.
- [Wag01] D. Wagner and D. Dean. Intrusion detection via static analysis. In 22nd Symp. on Security and Privacy (S&P), pages 156-169, Oakland, California, May 2001. IEEE Computer Society Press, Los Alamitos, California.

8 Appendices

8.1 Letters of support from the users

Signed letters of support have been provided by RaboFacet and TNO Telecom.

8.2 Other documents

None.

9 Abbreviations and acronyms

CTIT	Centre for Telematics and Information Technology
DIES	Distributed and Embedded Systems group
EEMCS	faculty of Electrical Engineering, Mathematics and Computer Science
fte	full-time equivalent
ISTRICE	Integrated Security and Privacy in a Networked World
ICT	Information and Communication Technology
ID	Intrusion Detection
IDS	Intrusion Detection System
IPID	Integrated Policy-based Intrusion Detection
IPS	Intrusion Prevention System
IS	Information Systems group
IT	Information Technology
ITIL	IT Infrastructure Library
SENTINELS	Security in ICT, Networks and Information Systems
STW	Stichting Technische Wetenschappen
UT	University of Twente