WS 11/12

Seminar „Vertrauenswürdige Systeme“

Organisatorisches

Ansprech-partner	**Dr. Edzard Höfig** Raum 018, Tel. 838 75277 edzard.hoefig@fu-berlin.de
Termine	Voraussichtlich 12 Termine im Zeitraum 20.10.11 bis 9.2.12. Termine am 29.12.11 und 2.2.12 fallen aus, es gibt eine Pause (etwa drei Termine) vor Beginn der Vorträge.

Geforderte Leistungen	• **Vortrag** über 30 Minuten + 15 Minuten Diskussionsleitung
	  • Einreichen der Folien eine Woche vor Vortrag
	• Schriftliche **Ausarbeitung** mit 10-15 Seiten in LaTeX
	  • LaTeX Vorlage wird gestellt
	  • Abgabe der Ausarbeitung spätestens am 29.2.12 per EMail an Edzard Höfig

Kursform	Seminar mit 2 SWS, bzw. 4 ECTS Punkten
Sprache	Vortragssprache und Sprache für die schriftliche Ausarbeitung wahlweise Englisch oder Deutsch.
Fehlzeiten	Maximal zwei Fehltermine sind möglich.

Zur Auswahl der Themen

Jeder Teilnehmer und jede Teilnehmerin stellt einen Artikel aus der folgenden Liste vor. Die Artikel sind dabei in acht Kategorien geteilt: „Trusted Computing“, Modellierung von Vertrauensaspekten, Vertrauen in digitale Information, Kontrolle von Datenverwendung, Datenschutz und Privatsphäre, gesellschaftliche und wirtschaftliche Aspekte, Engineering, sowie anwendungsbezogene Artikel. Mit Ausnahme des ersten Artikels (Vorstellung des ISO Standards zum „Trusted Platform Module“) der für eine Vorstellung durch zwei Personen im Rahmen einer 90 Minuten Veranstaltung vorgesehen ist, sollte jeder Teilnehmer genau einen Artikel vorstellen. Da Überschneidungen bei der Auswahl nicht zu verhindern sind, wählen Sie bitte jeweils drei Artikel aus die Sie potentiell vorstellen möchten. In der nächsten Veranstaltung am 27.10.11 werden wir dann gemeinsam die Artikel durchgehen und versuchen eine faire Zuordnung zu Ihren Vorlieben zu finden.

Es stehen die im folgenden aufgeführten Artikel zur Verfügung. Jeder Artikel ist mit einer Kennung gekennzeichnet, die dem Dateinamen des Artikels entspricht (z.B. Mod01) und es werden der Titel und eine Zusammenfassung aufgeführt. Die vollständigen Artikel können sie in einer passwortgeschützten Zip-Datei von der Veranstaltungs-Homepage runterladen. Bitte beachten Sie, dass die Artikel urheberrechtsgeschützt sind und nicht weitergegeben werden dürfen. Das Passwort wird während der Veranstaltungstermine bekannt gegeben.

Artikel zu „Trusted Computing“

TPM01: ISO/IEC Standard 118899 Trusted Platform Module (2 Students)

Consists of 4 documents: Overview (A) - Design Principles (B) - Structures (C) - Commands (D)

The Trusted Platform Module (TPM) is both the name of a published specification detailing a secure cryptoprocessor that can store cryptographic keys that protect information, as well as the general name of implementations of that specification. The TPM specification is the work of the Trusted Computing Group and also available as the international standard ISO/IEC 11889. The Trusted Platform Module offers facilities for the secure generation of cryptographic keys, and limitation of their use, in addition to a hardware pseudo-random number generator. It also includes capabilities such as remote attestation and sealed storage. Software can use a Trusted Platform Module to authenticate hardware devices. Since each TPM chip has a unique and secret RSA key burned in as it is produced, it is capable of performing platform authentication. For example, it can be used to verify that a system seeking access is the expected system.

TPM02: Cryptography and Competition Policy – Issues with ‘Trusted Computing’

The most significant strategic development in information technology over the past year has been ‘trusted computing’. This is popularly associated with Microsoft’s ‘Palladium’ project, recently renamed ‘NGSCB’. In this paper, I give an outline of the technical aspects of ‘trusted computing’ and sketch some of the public policy consequences.

Artikel zur Modellierung von Vertrauensaspekten

Mod01: Deriving Trust from Experience

In everyday life, trust is largely built from experience. Reputation-based trust models have been developed to formalize this concept. The application to networks like the Internet where a very large number of predominantly unknown principal identities engage in interactions is appealing considering that the evaluation of trusted experience may result in a more successful choice of trusted parties to interact with.

In this paper we pick the SECURE framework, as developed within the equally named EU project on Global Computing, which builds upon event structures to model possible outcomes of interactions. We extend it by three concepts: (i) a flexible way to determine a degree of trust from given past behavior, (ii) a basic notion of context, exemplarily in the form of roles the interacting parties may occupy, and (iii) we explicitly equip observed events with a time component to refine the granularity of observations.

We extend definitions of concepts used in SECURE in order to incorporate our notion of context information, we provide the syntax and semantics of an LTL-like logic, in its basics similar to the one proposed by Krukow, Nielsen and Sassone, that allows for layered reasoning about context information. We then show how this new language relates to the one used in SECURE and we determine under which conditions our concept of deriving trust from experience may be used within SECURE’s computational model to obtain a global state of trust.

Mod02: CertainLogic: A Logic for Modeling Trust and Uncertainty

The evaluation of the trustworthiness of complex systems is a challenge in current IT research. We contribute to this field by providing a novel model for the evaluation of propositional logic terms under uncertainty that is compliant with the standard probabilistic approach and subjective logic. Furthermore, we present a use case to demonstrate how this approach can be applied to the evaluation of the trustworthiness of a system based on the knowledge about its components and subsystems.

Mod03: Trust Metrics in Recommender Systems

Recommender Systems based on Collaborative Filtering suggest to users items they might like, such as movies, songs, scientific papers, or jokes. Based on the ratings provided by users about items, they first find users similar to the users receiving the recommendations and then suggest to her items appreciated in past by those like-minded users. However, given the ratable items are many and the ratings provided by each users only a tiny fraction, the step of finding similar users often fails. We propose to replace this step with the use of a trust metric, an algorithm able to propagate trust over the trust network in order to find users that can be trusted by the active user. Items appreciated by these trustworthy users can then be recommended to the active user. An empirical evaluation on a large dataset crawled from Epinions.com shows that Recommender Systems that make use of trust information are the most effective in term of accuracy while preserving a good coverage. This is especially evident on users who provided few ratings, so that trust is able to alleviate the cold start problem and other weaknesses that beset Collaborative Filtering Recommender Systems.

Mod04: A Formal Approach Towards Measuring Trust in Distributed Systems

Emerging digital environments and infrastructures, such as distributed security services and distributed computing services, have generated new options of communication, information sharing, and resource utilization in past years. How- ever, when distributed services are used, the question arises of to what extent we can trust service providers to not violate security requirements, whether in isolation or jointly. Answering this question is crucial for designing trustworthy distributed systems and selecting trustworthy service providers. This paper presents a novel trust measurement method for distributed systems, and makes use of propositional logic and probability theory. The results of the qualitative part include the specification of a formal trust language and the representation of its terms by means of propositional logic formulas. Based on these formulas, the quantitative part re- turns trust metrics for the determination of trustworthiness with which given distributed systems are assumed to fulfill a particular security requirement.

Artikel zu Vertrauen in digitale Informationen

Inf01: An Information Flow Verifier for Small Embedded Systems

Insecurity arising from illegal information flow represents a real threat in small computing environments allowing code sharing, dynamic class loading and overloading. We introduce a verifier able to certify at loading time Java applications already typed with signatures describing possible information flows. The verifier is implemented as a class loader and can be used on any Java Virtual Machine. The experimental results provided here support our approach and show that the verifier can be successfully embedded. As far as we know, this is the first information flow analysis adapted to open embedded systems.

Inf02: Trust in Digital Information

Trust in information is developing into a vitally important topic as the Internet becomes increasingly ubiquitous within society. Although many discussions of trust in this environment focus on issues like security, technical reliability, or e-commerce, few address the problem of trust in the information obtained from the Internet. The authors assert that there is a strong need for theoretical and empirical research on trust within the field of information science. As an initial step, the present study develops a model of trust in digital information by integrating the research on trust from the behavioral and social sciences with the research on information quality and human– computer interaction. The model positions trust as a key mediating variable between information quality and information usage, with important consequences for both the producers and consumers of digital information. The authors close by outlining important directions for future re- search on trust in information science and technology.

Inf03: Trustworthy Information: Concepts and Mechanisms

We used to treating information received (from recognized sources) as trustworthy, which is unfortunately not true because of attacks. The situation can get worse with the emerging shift of information sharing paradigm from “need to know” to “need to share.” In order to help information consumers make the “best” decision possible, it is imperative to formulate concepts, models, frameworks, architectures, and mechanisms to facilitate information trustworthiness management in distributed and decentralized environment. In this paper we initiate a study in this direction by proposing an abstraction called information networks as well as two supporting mechanisms called provenance digital signatures and optimal security hardening of information network.

Artikel zur Kontrolle von Datenverwendung

Use01: A Trustworthy Usage Control Enforcement Framework

Usage control policies specify restrictions on the handling of data after access has been granted. We present the design and implementation of a framework for enforcing usage control requirements and demonstrate its genericity by instantiating it to two different levels of abstraction, those of the operating system and an enterprise service bus. This framework consists of a policy language, an automatic conversion of policies into enforcement mechanisms, and technology implemented on the grounds of trusted computing technology that makes it possible to detect tampering with the infrastructure. We show how this framework can, among other things, be used to enforce separation-of-duty policies. We provide a performance analysis.

Use02: Representation-Independent Data Usage Control

Usage control is concerned with what happens to data after access has been granted. In the literature, usage control models have been defined on the grounds of events that, somehow, are related to data. In order to better cater to the dimension of data, we extend a usage control model by the explicit distinction between data and representation of data. A data flow model is used to track the flow of data in-between different representations. The usage control model is then extended so that usage control policies can address not just one single representation (e.g., delete file1.txt after thirty days) but rather all representations of the data (e.g., if file1.txt is a copy of file2.txt, also delete file2.txt). We present three proof-of-concept implementations of the model, at the operating system level, at the browser level, and at the X11 level, and also provide an ad-hoc implementation for multi-layer enforcement.

Use03: State-based Usage Control Enforcement with Data Flow Tracking using System Call Interposition

Usage control generalizes access control to what happens to data in the future. We contribute to the enforcement of usage control requirements at the level of system calls by also taking into account data flow: Restrictions on the dissemination of data, for instance, as stipulated by data protection regulations, of course relate not to just one file containing the data, but likely to all copies of that file as well. In order to enforce the dissemination restrictions on all copies of the sensitive data item, we introduce a data flow model that tracks how the content of a file flows through the system (files, network sockets, main memory). By using this model, the existence of potential copies of the data is reflected in the state of the data flow model. This allows us to enforce the dissemination restrictions by relating to the state rather than all sequences of events that possibly yield copies. Generalizing this idea, we describe how usage control policies can be expressed in a related state-based manner. Finally, we present an implementation of the data flow model and state-based policy enforcement as well as first encouraging performance measurements.

Artikel zu Datenschutz und Privatsphäre

Pri01: Some Like It Private - Sharing Confidential Information Based on Oblivious Authorization

Privacy-Preserving Policy-Based Information Transfer (PPIT) lets entities that lack mutual trust share sensitive information. The authors discuss the security of two efficient PPIT constructs, then propose an innovative construct that allows entities to efficiently verify the equality of their information.

Pri02: Anonymous Authentication with TLS and DAA

Anonymous credential systems provide privacy-preserving authentication solutions for accessing services and resources. In these systems, copying and sharing credentials can be a serious issue. As this cannot be prevented in software alone, these problems form a major obstacle for the use of fully anonymous authentication systems in practice. In this paper, we propose a solution for anonymous authentication that is based on a hardware security module to prevent sharing of credentials. Our protocols are based on the standard protocols Transport Layer Security (TLS) and Direct Anonymous Attestation (DAA). We present a detailed description and a reference implementation of our approach based on a Trusted Platform Module (TPM) as hardware security module. Moreover, we discuss drawbacks and alternatives, and provide a pure software implementation to compare with our TPM-based approach.

Pri03: Escrowed Data and the Digital Envelope

As computers continue to permeate all aspects of our lives, there is a growing tension between the requirements of societal security and individual privacy. Societal security encompasses all ways in which we try to make the world more secure, including transport security, financial security, infrastructure security, etc. A prime mechanism for achieving this security involves collecting quantities of data about individuals, for example via ISP logs, mobile phone logs, ticketing systems, and banking systems. We propose escrowed data as an approach that may be capable of providing an appropriate balance between the requirements of individual privacy and societal security.

Artikel zu Gesellschaftlichen und Wirtschaftlichen Aspekten

EcoSoc01: The Economics of Click Fraud

Click fraud is a substantial threat in the cyberworld. Here, the author examines the contexts, mechanisms, and processes associated with the click-fraud industry from an economics viewpoint. The nature of electronic channels, characterized by asymmetric hypermediation, provides a fertile ground for such fraud.

EcoSoc02: Trust Economy: Aspects Of Reputation And Trust Building For SMEs In E-Business

The lack of direct communication is a problem in E-Business. It often leads to financial disadvantages for small and medium sized companies. This paper investigates the interrelationships between reputation, trust, risk, and costs. It presents a framework combining these four parameters. In addition, it develops a trust building process, a trust cycle, and presents some instruments to engender trust. Companies may consider these solutions to be helpful in order to overcome the lack of direct communication, to build trust between business partners, to understand the interrelationships between the four mentioned parameters, and to avoid financial disadvantages.

EcoSoc03: Privacy Requirements Engineering for Trustworthy e-Government Services

Several research studies have applied information systems acceptance theories in order to examine issues related to the acceptance of e- services by users. Their application in the e-government systems has revealed that trust is a prerequisite for their usage. Moreover, it has been proved that privacy concerns are a main antecedent of trust in e-government systems intention of use. Therefore, information systems that are not privacy aware are not trusted and thus not accepted by users. Currently there are many different attacks that can be realized by malicious users for compromising the confidentiality of private data and thus putting at stake the trustworthiness of the systems. The conventional way for preventing such attacks is mainly the employment of Privacy Enhancing Technologies (PETs). However, PETs are employed as ad hoc technical solutions that are independent from the organizational context in which the system will operate. We argue that we need privacy requirements engineering methods for capturing the context dependent privacy requirements and for selecting the appropriate technical, organizational and procedural countermeasures which will help building privacy aware systems that can offer electronic services which users can trust.

EcoSoc04: Axiomatic and Behavioural Trust

Academic discourse on trust is fractured along disciplinary lines. Security theorists routinely use a definition of trust which, apparently, has little in common with any of the definitions of trust that appear in the sociological and psychological literature. In this essay, we extend a recently-proposed framework for the technical analysis of se- cure systems, so that its notion of trust is roughly congruent with the sociological theories of Parsons, Luhmann, Barber, Lewis and Weigert. This congruent extension suggests some ways in which a computerised system might, appropriately, inspire trust in its non-technical users.

Artikel mit Schwerpunkt „Engineering“

Eng01: Modeling Trust Negotiation for Web services

As Web services become more widely adopted, developers must cope with the complexity of evolving trust negotiation policies spanning numerous autonomous services. The Trust-Serv framework uses a state-machine-based modeling approach that supports life-cycle policy management and automated enforcement.

Eng02: Scenario-Driven Role Engineering

Scenario-driven role engineering is a systematic approach for defining customized role-based access control models. Based on his role-engineering experiences, the author discusses the relations between different role-engineering artifacts, the need for process tailoring, and the use of preexisting documents in role-engineering activities.

Eng03: Harmony: Integrated Resource and Reputation Management for Large-Scale Distributed Systems

Advancements in technology over the past decade are leading to a promising future for large-scale distributed systems, where globally-scattered distributed resources are collectively pooled and used in a cooperative manner to achieve unprecedented petascale supercomputing capabilities. The issues of resource management (resMgt) and reputation management (repMgt) need to be addressed in order to ensure the successful deployment of large-scale distributed systems. However, these two issues have typically been addressed separately, despite the significant interdependencies between them: resMgt needs repMgt to provide a cooperative environment for resource sharing, and in turn facilitates repMgt to evaluate multi-faceted node reputations for providing different resources. Current repMgt methods pro- vide a single reputation value for each node in

providing all types of resources. However, a node willing to provide one resource may not be willing to provide another resource. In addition, current repMgt methods often guide node selection policy to select the highest-reputed nodes, which may overload these nodes. Also, few works exploited node reputation in resource selection in order to fully and fairly utilize resources in the system and to meet users' diverse QoS demands. We propose a system called Harmony that integrates resMgt and repMgt in a harmonious manner. Harmony incorporates two key innovations: integrated multi-faceted resource/ reputation management and multi-QoS- oriented resource selection. The trace data we collected from an online trading platform confirms the importance of multi-faceted reputation and potential problems with highest-reputed node selection. Trace-driven experiments performed on PlanetLab show that Harmony outperforms existing resMgt and repMgt in terms of the success rate, service delay, and efficiency.

Eng04: Engineering Attestable Services

Web services require complex middleware in order to communicate using XML standards. However, this software increases vulnerability to runtime attack and makes remote attestation difficult. We propose to solve this problem by dividing services onto two platforms, an untrusted front-end, implementing the middleware, and a trustworthy back-end with a minimal trusted computing base.

Eng05: Dynamic Trust Management

Trust management forms the basis for communicating policy among system elements and demands credential checking for access to all virtual private service resources—along with careful evaluation of credentials against specified policies—before a party can be trusted.

Eng06: Verifying Trustworthiness of Virtual Appliances in Collaborative Environments

Often in collaborative research environments that are facilitated through virtual infrastructures, there are requirements for sharing virtual appliances and verifying their trustworthiness. Many researchers assume that virtual appliances — shared between known virtual organisations — are naturally safe to use. However, even if we assume that neither of the sharing parties are malicious, these virtual appliances could still be mis-configured (in terms of both security and experiment requirements) or have out-of-date software installed.

Based on formal methods, we propose a flexible method for specifying such security and software requirements, and verifying the virtual appliance events (captured through logs) against these requirements. The event logs are transformed into a process model that is checked against a pre-defined whitelist — a repository of formal specifications. Verification results indicate whether or not there is any breach of the requirements and if there is a breach, the exact steps leading to it are made explicit.

Eng07: Managing application whitelists in trusted distributed systems

Many distributed batch systems, such as computational grids, require a level of integrity protection to guarantee the proper execution of a job or workflow. One way of achieving this, implicit in many trusted computing proposals, is to use application whitelisting to prevent unknown and untrusted applications from being executed on remote services. However, this approach has significant shortcomings across multiple administrative domains, as conflicts between locally managed whitelists will result in many useful services appearing untrustworthy to users. This has the potential to limit availability and prevent trusted distributed systems from ever being successfully deployed.

We propose a set of requirements for a system which will manage these conflicts, and provide a mechanism for updating application whitelists that will increase service availability and trustworthiness. We also suggest and specify a set of components, including a centralised configuration manager, which will meet these requirements.

Anwendungsbezogene Artikel

App01: Using Trust for Secure Collaboration in Uncertain Environments

The SECURE project investigates the design of security mechanisms for pervasive computing based on trust. It addresses how entities in unfamiliar pervasive computing environments can overcome initial suspicion to provide secure collaboration.

App02: Lessons Learned from Building a High-Assurance Crypto Gateway

The construction of a complex secure system composed from both secure and insecure components presents a variety of challenges to the designer. The example system described here highlights some lessons learned from first-hand experience in attempting such a task.

App03: Querying Trust in RDF Data with tSPARQL

Today a large amount of RDF data is published on the Web. However, the openness of the Web and the ease to combine RDF data from different sources creates new challenges. The Web of data is missing a uniform way to assess and to query the trustworthiness of information. In this paper we present tSPARQL, a trust-aware extension to SPARQL. Two additional keywords enable users to describe trust requirements and to query the trustworthiness of RDF data. Hence, tSPARQL allows adding trust to RDF-based applications in an easy manner. As the foundation we propose a trust model that associates RDF statements with trust values and we extend the SPARQL semantics to access these trust values in tSPARQL. Furthermore, we discuss opportunities to optimize the execution of tSPARQL queries.

App04: The Challenge of Validation for Autonomic and Self-Managing Systems

The research community has achieved great success in building autonomic systems (AS) in line with the vision of autonomic computing (AC) released by IBM in 2001. The success is gaining ground in addressing the perceived concerns of complexity and total cost of ownership of information technology (IT) systems. But we are now faced with a challenge springing from this very success. This challenge is trustworthiness and there are limited research results published in this direction. This, if not addressed will definitely undermine the success of AC. How do we validate a system to show that it is capable of achieving a desired result under expected range of contexts and environmental conditions and beyond? This paper identifies the challenges and significance of AS validation and proposes a roadmap towards achieving trustworthiness in autonomic systems.

App05: Analysis of a Botnet Takeover

This article describes an effort to take control of a particularly sophisticated and insidious botnet and study its operations for a period of 10 days. It summarizes what the authors learned and reports on what has happened to that botnet since.