OpenTC Newsletter
June 2007

From the Open Trusted Computing (OpenTC) research project, sponsored by the European Union.

In this issue:

Editorial: Why this newsletter?

By: Arnd and Dirk Weber, ITAS, Forschungszentrum Karlsruhe, Germany

Dear reader,

Welcome to the first newsletter of the Open Trusted Computing (OpenTC) project. It was created to inform the interested public about downloadable prototypes, project achievements and other up-to-date information, and it is meant as a mean to start discussions about the underlying issues. We aim to publish this newsletter irregularly during the course of the project.

The Open Trusted Computing (OpenTC) project aims at building a secure hypervisor, using Trusted Computing (TC) technologies. In other words, OpenTC aims at creating a hypervisor for the virtualisation of hardware used by operating systems that are running in different compartments and in parallel, on top of the hypervisor. These compartments, including the hypervisor, are secured using the Trusted Platform Module (TPM). In such an environment, operating systems such as Linux and Windows will be able to run in strongly isolated compartments. Inside the compartments, applications may use TC or not - it will be for the user to choose. The source code of the hypervisor will be published as Open Source Software in order to be transparent, and show to the public what the hypervisor does and how it achieves it.

In this first issue of the OpenTC newsletter, the first article provides a short overview of the project objectives, written by the project leader Dirk Kuhlmann. A larger high level overview of the project is available on the project website: http://www.opentc.net/index.php?option=com_content&task=view&id=14&Itemid=29

The second article, written by Arnd Weber, is a report about one of the first Trusted Computing conferences, that took place in Berlin in October 2006. It provides an

overview of some TC applications and highlights some of the key issues debated by the audience.

At the end of the newsletter, we provide two lists of publications that are now available on the OpenTC website. One is a list of consortium “deliverables”. These are documents that have been delivered by the OpenTC consortium to the European Commission. A list of scientific publications written by project members is also given. Some of these papers can also be found on the project website, while others are available in journals and in conference proceedings.

About the authors: Arnd and Dirk Weber work with the Institute for Technology Assessment and Systems Analysis (ITAS) at Forschungszentrum Karlsruhe, Germany. They work on requirements and dissemination activities of the OpenTC project.

Contact: {arnd, dirk}.weber at itas.fzk.de

OpenTC – an open approach to trusted virtualisation

By: Dirk Kuhlmann, Hewlett Packard Laboratories, Bristol, UK

Editor´s note: This article was originally published in the INDICARE Monitor, a newsletter of the EU-project INDICARE (INformed DIalogue about Consumer Acceptability of DRM Solutions in Europe), on January 2, 2006, licensed under Creative Commons (http://www.indicare.org/tiki-read_article.php?articleId=183).

Introduction

The advent of "Trusted Computing" (TC) technology as specified by the Trusted Computing Group (cf. references) has not met much enthusiasm by the Free/Open Source Software (FOSS) and LINUX communities so far. Despite this fact, FOSS based systems have become the preferred vehicle for much of the academic and industrial research on Trusted Computing. In parallel, a lively public discussion between proponents and critics of TC has dealt with the question of whether the technology and concepts put forward by the TCG are compatible, complementary or potentially detrimental to the prospects of open software development models and products.

Common misconceptions of TC technology are that it implies or favours closed and proprietary systems, reduces options of using arbitrary software, or remotely controls users' computers. It has long been argued, though, that these and similar undesirable effects are by no means unavoidable, not least because the underlying technology is passive and neutral with regard to specific policies. The actual features displayed by TC equipped platforms will almost exclusively be determined by the design of the operating systems and software running on top of it. With appropriate design, implementation and validation of trusted software components, and by using contractual models of negotiating policies, negative effects can be circumvented while improving the system's trust and security properties. This is the intellectual starting

point of the EU-supported, collaborative OpenTC research and development project that started in November 2005.

Combining FOSS and TC technology

OpenTC aims to demonstrate that a combination of TC technology and FOSS has several inherent advantages that are hard to meet by any proprietary approach. Enhanced security at the technical level tends to come at the expense of constraining user options, and the discursive nature of FOSS-development could help to find the right balance here. Trusted software components have to be protected from analysis during runtime, so it is highly desirable that their design is documented and that the source code is available to allow for inspection and validation. Finally, any attempts to introduce TC technology are likely to fail without the buy-in of its intended users, and openness could prove to be the most important factor for user acceptance.

OpenTC sets out to support cooperative security models that can be based on platform properties without having to assume the identifiability, personal accountability and reputation of platform owners or users. For reasons of privacy and efficiency, these models could be preferable to those assuming adversarial behaviour from the outset. A policy model based on platform properties, however, requires reliable audit facilities and trustworthy reporting of platform states to both local users and remote peers. The security architecture put forward by the TCG supplies these functions, including a stepwise verification of platform components with an integral, hardware-assisted auditing facility at its root. In OpenTC, this will be used as a basic building block.

Trusted virtualisation and protected execution environments

The goal of the OpenTC architecture is to provide execution environments for whole instances of guest operating systems that communicate to the outside world through reference monitors guarding their information flow properties. The monitors kick into action as soon as an OS instance is started. Typically, the policy enforced by a reference monitor should be immutable during the lifetime of the instance: it can neither be relaxed through actions initiated by the hosted OS nor overridden by system management facilities. In the simplest case, this architecture will allow two independent OS instances to be run, with different grades of security lock-down on an end user system. Such a model, with an unconstrained "green" environment for web browsing, software download and installation, and a tightly guarded "red" side for tax record, banking communications etc., has recently been discussed by Carl Landwehr (2005). More complex configurations are possible and are frequently needed in server scenarios.

OpenTC is borrowing from research on trusted operating systems that goes back as far as 30 years. The underlying principles – isolation and information flow control – have been implemented by several security hardened versions of Linux, and it has been demonstrated that such systems can be integrated with Trusted Computing technology (see e.g. Maruyama et al. 2003). However, the size and complexity of these implementations is a serious challenge for any attempt to seriously evaluate their actual security properties. The limited size of developer communities, and difficulties of understanding and complexity of managing configurations and policies,

continue to be road blocks for the deployment of trusted platforms and systems on a wider scale.

Compared to full-blown operating systems, the tasks of virtualisation layers tend to be simpler. This should allow OpenTC to reduce the size of the Trusted Computing Base. The architecture separates management and driver environments from the core system and hosted OS instances. They can either be hosted under stripped-down Linux instances, or they can run as generic tasks of the virtualisation engines. The policy enforced by the monitors is separated from decision and enforcement mechanisms. It is human readable and can therefore be subjected to prior negotiation and explicit agreement.

OpenTC chose (para-)virtualisation as the underlying architecture for a trusted system architecture, which allows standard OS distributions and applications to be run side by side with others that are locked down for specific purposes. This preempts a major concern raised with regard to Trusted Computing, namely, that TC excludes components not vetted by third parties. The OpenTC architecture allows the imposing of constraints on components marked as security critical, while unconstrained components can run in parallel.

OpenTC builds on two virtualisation engines: XEN and L4. Both are available under FOSS licenses and are boosted by active developer and user communities. Currently, it is necessary to compile special versions of Linux that cooperate with the underlying virtualisation layer. However, the development teams will improve their architectures to also support unmodified, out-of-the-box distributions. This will be simplified by hardware support for virtualisation as offered by AMD's and INTEL's new CPU generations. Prototypic results have shown that this hardware support could also allow hosting unmodified operating systems other than Linux (see e.g. Shankland 2005).

From trusted to trustworthy computing

TCG hardware provides basic mechanisms to record and report the startup and runtime state of a platform in an extremely compressed, non-forgeable manner. It allows the creation of a digitally signed list of values that correspond to elements of the platform's Trusted Computing Base. In theory, end users could personally validate each of these components, but this is not a practical option. End users may have to rely on other parties to evaluate and attest that a particular set of values corresponds to a system configuration with a desired behaviour. In this case, their reason to trust will ultimately stem from social trust users put in statements from specific brands, certified public bodies, or peers groups.

A much discussed dilemma arises if trusted components become mandatory prerequisites for consuming certain services. Even if such components are suspicious to the end user, they might still be required by a provider. This problem is particularly pronounced if named components come as binaries only and do not allow for analysis. The recent history of DRM technology has shown that trojans can easily be inserted under the guise of legitimate policy enforcement modules. Clearly, a mechanism that enforces DRM on a specific piece of content acquired by a customer must not assume an implicit permission to sift through the customer's hard disk and report back on other content.

This highlights an important requirement for components that deserve the label "trusted": at least in principle, it should be possible to investigate their actual trustworthiness. A clearly stated description of function and expected behaviour should be an integral part of their distribution, and it should be possible to establish that they do not display behaviour other than that stated in their description – at compile time, runtime, or both. A socially acceptable approach to Trusted Computing will require transparency and open processes. In this respect, a FOSS based approach looks promising, as it might turn openness into a crucial competitive advantage.

The TCG specification is silent on procedures or credentials required before a software component can be called "trusted". OpenTC works on the assumption that defined methodologies, tools, and processes to describe goals and expected behaviour of software components are needed. In this way, it will become possible to check whether their implementation reflects (and is constrained to) their description. Independent replication of tests may be required to arrive at a commonly accepted view of a component's trustworthiness, which in turn requires accessibility of code, design, test plans and environments for the components under scrutiny.

Trust, risk, and freedom

Most of us have little choice but to trust IT systems where more and more things can go wrong, while our actual insight into what is actually happening on our machines gets smaller by the day. Users are facing a situation of having to bear full legal responsibility for actions initiated on or by their machines while lacking the knowledge, tools and support to keep these systems in a state fit for purpose. Due to the growing complexity of our technology, we will increasingly have to rely on technical mechanisms that help us to estimate the risk prior to entering IT based transactions. Enhanced protection, security and isolation features based on TCG technology will become standard elements of proprietary operating systems and software in due time.

This evolution is largely independent of whether FOSS communities endorse or reject this technology. OpenTC assumes that mutual attestation of the platforms' "fitness for purpose" will become necessary for proprietary systems as well as FOSS based ones. The absence of comparable protection mechanisms for non-proprietary operating or software systems will immediately create problems for important segments of professional Linux users. In fact, many commercial, public or governmental entities have chosen non-proprietary software for reasons of transparency and security. These organizations tend to be subjected to stringent compliance regulations requiring state-of-the-art protection mechanisms. If FOSS based solutions do not support these mechanisms, the organizations could eventually be forced to replace their non-proprietary components with proprietary ones: a highly undesirable state of affairs that OpenTC might help to avoid.

From this perspective, the current discussion about the next version of the GNU public license raises serious concerns. Some of the suggested changes could impact on the possibility of combining Trusted Computing technology and Free Software licensed under GPLv3 - this refers to the GPLv3 Draft, status 2006-02-07 16:50 (cf. references). Section 3 of this draft concerns Digital Restrictions Management, a term

that has been used by Richard Stallman in discussions about Trusted Computing. For example, the current draft excludes “modes of distribution that deny users that run covered works the full exercise of the legal rights granted by this License”. It is an open question whether this might apply to elements of a security architecture such as OpenTC. A Trusted Computing architecture does not constrain the freedom of copying, modifying and sharing works distributed under the GPL. However, it can constrain the option of running modified code as a trusted component, since previously evaluated security properties might have been affected by the modifications. Unless a re-evaluation is performed, the properties of modified versions cannot be derived from the attestation of the original code; security assurances about the original code become invalid.

This is by no means specific to the Trusted Computing approach; it also applies to commercial Linux server distributions with protection profiles evaluated according to the Common Criteria. The source code for the distribution is available, but changing any of the evaluated components results in loss of the certificate. Whether or not software is safe, secure, or trustworthy is independent of the question of how it is licensed and distributed. The option to choose between proprietary and FOSS solutions is an important one and should be kept open. This is one of the reasons why several important industrial FOSS providers and contributors are participating in OpenTC. The project aims at a practical demonstration that Trusted Computing technology and FOSS can complement each other. This is possible in the context of the current GPLv2. Whether it will be so under a new GPLv3 remains to be seen.

References:

- GPLv3 Draft, status 2006-02-07 16:50: http://gplv3.fsf.org/draft
- Landwehr, Carl (2005): Green Computing. IEEE Security&Privacy, Vol 3, No 6, Nov/Dec 2005, p. 3
- Maruyama et al. (2003): Linux with TCPA Integrity Measurement. IBM Research Report RT0575, January 2003; http://www.research.ibm.com/trl/people/munetoh/RT0507.pdf
- Shankland, Stephen (2005): XEN passes Windows Milestone. CNET News.com, August 23, 2005; http://news.com.com/Xen+passes+Windows+milestone/2100-7344_3-5842265.html
- Trusted Computing Group: http://www.trustedcomputinggroup.org

Disclaimer
The content of this paper is published under the sole responsibility of the author. It does not necessarily reflect the position of HP Laboratories or other OpenTC members.

About the author:

Dirk Kuhlmann is a senior research engineer for Hewlett Packard Laboratories in Bristol, UK, where he works as a member of the Trusted Systems Laboratory. He acts as the overall technical lead for the OpenTC project.

Contact: dirk.kuhlmann at hp.com

Report from the conference "The World of Trusted Computing – Hightech in Europe", Berlin, Germany, October 19-20, 2006

By Arnd Weber, ITAS, Forschungszentrum Karlsruhe, Germany

The German Federal Ministry of Economics and Technology, together with Ruhr-Universitaet Bochum and ISITS (International School of IT Security), organised a conference on the needs of European public and private organisations regarding Trusted Computing (TC) and the future of this global initiative. The conference took place in October 2006 and was sponsored by Hewlett Packard, Sirrix Security Technologies, Computer-Zeitung (a German computer magazine) and Linux Magazin. It was organised in the Ministry's spacious facilities of the former Kaiser Wilhelm-Akademie for military physicians, and was chaired by Ahmad-Reza Sadeghi from Ruhr-Universitaet Bochum (Germany). 85 participants from Europe, Japan, Israel and USA discussed over two days a variety of issues related to TC and its deployment. It is outside the scope of this report to provide a detailed review of the content of all of the 18 talks that were given by speakers from Germany, Belgium, Japan, UK, and US. Instead, three key aspects of the discussions will be highlighted in this article.

The first significant aspect developed during the conference is the high level of security provided by the combination of new processor architectures together with TC technology. David Grawrock from Intel, the originator and editor of the TCG's Best Practices Paper, spoke about the security of the new Intel architecture, formerly referred to as Intel LaGrande Technology and now called Trusted Execution Technology, which contributes to the provision of a trusted platform environment. David Grawrock explained that the new processor architecture (similar to OpenTC partner AMD's AMD-V architecture, formerly called "Pacifica") is aimed at a strict separation of the system kernel and applications by supporting the new virtualisation technology. By making use of TC technology, the new processors should be capable of providing virtualisation similar to mainframe computer task separation. Additionally, the new Intel processors will contain a special functionality block (together with a separate dedicated memory cache module) for measuring code and data structures, one of the main functions in the TCG standards that is currently implemented in the Trusted Platform Module (TPM). Using this architecture, security attacks such as Trojan horses mentioned by Udo Helmbrecht, President of the German Federal Office for Information Security, can hopefully be prevented.

A second aspect discussed during the conference is the topic of the applications of Trusted Computing. The applications mentioned by the speakers can be grouped as follows:

- Protection against theft of hard disk data. If user data are encrypted using TC, theft of the hard disk is no longer a threat. In addition, as mentioned by Peter Biddle from Microsoft, decommissioning of hard disks is eased significantly, as one simply needs to reset the TPM.

- Trusted Network Connect (TNC), allowing corporations to verify the security of machines which are accessing their network. Michael Hartmann of SAP highlighted the significant progress that TNC brings.

- Marit Hansen from the Independent Centre for Privacy Protection Schleswig-Holstein (Germany) talked about the possibilities and opportunities for using Trusted Computing to protect the privacy of users, for example in eHealth.

- Digital Rights Management was another example of a Trusted Computing application. Supply chains can be designed to protect the rights of the different companies involved, for example in the automotive industry. Furthermore, secure communication between patent holders and patent lawyers can also be implemented using DRM.

- Protecting applications from each other was mentioned by several speakers from different companies, for example for consolidating servers, in grid computing, in car navigation systems, as well as in eGovernment, eCommerce and eHealth applications. Applications used on behalf of one's employer on a private home PC can be separated from applications used for private reasons. Speakers also mentioned secure printing, multi-player games, and mobile phone applications.

A third group of issues debated at the conference comprises open questions raised by the speakers and the participants, such as:

- Ahmad-Reza Sadeghi, the leader of the EMSCB project (European Multilaterally Secure Computing Base), pointed out that Trusted Computing is an emerging technology and that there are still many technical and non-technical challenges for Trusted Computing to face, and that the Trusted Computing Group (TCG) should work more closely with academia and support research and education in this area.

- An important question is whether Trusted Computing is open and will remain so in the future. This set of issues starts with the question of whether a TPM is compliant with the TCG standard (a topic developed in the presentation given by Georg Rankl from Infineon). Ruhr-Universitaet Bochum (Germany) tested most of the TPMs available on the market and discovered that some TPMs are not fully compliant with the TCG specifications. Another issue is whether computer architectures will remain open in the future. European governments want this openness of specification and design, as Ulrich Sandl from the Ministry of Economics and Technology discussed. The EMSCB and OpenTC projects are important projects that will help to ensure that Trusted Computing systems are open. Jacques Bus from the European Commission mentioned research on TC supported by his department, i.e. the projects OpenTC, Re-TRUST and Robin.

- Stefan Bechtold from the Max Planck Institute pointed out that Trusted Computing systems may lead to possible legal and economic problems, such as the lack of competition when a single, all-encompassing infrastructure is set up.

- Seth Schoen from the Electronic Frontier Foundation (EFF) explained some concerns regarding certain aspects of Trusted Computing Technology such as

attestation. In particular he brought up the issue of who decides what is trustworthy.

- A hot topic on the agenda is the implementation of a secure hypervisor, a trustworthy implementation of the virtualisation concept, as confirmed by Peter Biddle from Microsoft. This is a major objective of the OpenTC project.

- Dirk Kuhlmann, the OpenTC project's technical leader, working at the HP laboratories, brought up the question of how to host Windows on such a hypervisor, an issue where he would like to obtain support from European governments.

- For David Grawrock from Intel, one of the important research issues is to figure out exactly what evidence a relying party needs in order to be able to trust a remote platform.

- The author brought up the question of assessing the level of assurance that can be guaranteed to the user if the new processors really provide their claimed level of security, in particular with regard to isolation properties.

Is the debate over Trusted Computing over? Well, if people with differing views such as Peter Biddle from Microsoft and Seth Schoen from the Electronic Frontier Foundation can openly discuss the underlying issues, one can believe that progress is taking place. The number of useful applications as well as the mood of the discussion suggests that the debates have become more objective and less polemical. However, the open issues and the possibility of charging more for software and content by using TC technology means that the topic of Trusted Computing will remain on the agenda, from technical but also political and economic perspectives. Jacques Bus from the European Commission demanded the continuation of the dialogue with all stakeholders.

References:

- Computer-Zeitung: Trusted Computing hat als rotes Tuch ausgedient (in German). October 31, 2006: http://www.computer-zeitung.de/themen/sicherheit/article.html?thes=&art=/articles/2006045/30859942_ha_CZ.html
- Intel Trusted ExecutionTechnology: http://www.intel.com/technology/security/
- ISITS (International School of IT Security): https://www.is-its.org/
- Ruhr-Universitaet Bochum, Applied Data Security Group: http://www.prosec.rub.de/
- Trusted Computing Group, Best Practices and Principles: https://www.trustedcomputinggroup.org/specs/bestpractices/

About the author:

Arnd Weber works with the Institute for Technology Assessment and Systems Analysis (ITAS) at Forschungszentrum Karlsruhe, Germany. He acts as editor of this newsletter and leads the work on requirements in the OpenTC project.

Contact: arnd.weber at itas.fzk.de

OpenTC publications available

The OpenTC project has produced several documents related to its various activities. These include public deliverables and scientific publications.

Public deliverables

You can access the public deliverables from the OpenTC project via this webpage: http://www.opentc.net/index.php?option=com_content&task=view&id=27&Itemid=41

The following deliverables are available:

- D01.3 - Executive Summary of the first project year of OpenTC
- D02.2 - Requirements and specifications report
- D04.1 – Basic Management Interface Specification V.01
- D07.1 - Security Requirements definition, Target Selection, Methodology Definition
- D10.1 - Intermediate report about all external cooperation and activities
- D10.4 – Training concepts and training plans

Scientific publications

Scientific publications of the project are listed on this webpage: http://www.opentc.net/index.php?option=com_content&task=view&id=27&Itemid=41

Among the scientific publications you can find:

- Alkassar, A.; Scheibel, M.; Sadeghi, A-R.; Stueble, C.; Winandy, M.: Security Architecture for Device Encryption and VPN, accepted for ISSE (Information Security Solution Europe) 2006.
- Birk, D.; Gajek, S.; Grobert, F.; Sadeghi, A-R.: Phishing phishers - oberserving and tracing organized cybercrime. In: IEEE Cyberfraud, 2007 (to appear).
- Chen, L.; Landfermann, R.; Loehr, L.; Rohe, M.; Sadeghi, A-R.; Stueble, C.: A Protocol for Property-Based Attestation, accepted for The First ACM Workshop on Scalable Trusted Computing (STC'06).(available at www.opentc.net)

- Gajek, S.; Sadeghi, A-R.: Client Authentication in Federations Using a Security Mode, accepted to be presented at Toward a More Secure Web - W3C Workshop on sability and Transparency of Web Authentication. (available at http://www.w3.org/2005/Security/usability-ws/program).

- Gajek, S.; Sadeghi, A-R.; Stueble, C.; Winandy, M.: Compartmented Security for Browsers - Or How to Thwart a Phisher with Trusted Computing, accepted for The Second International Conference on Availability, Reliability and Security (ARES 2007). (available at www.opentc.net)

- Gallery, E.; Balfe, S.: Mobile Agents and the Deus Ex Machina: Protecting Agents using Trusted Computing. In: Proceedings of the 2007 IEEE International Symposium on Ubisafe Computing (UbiSafe-07), Niagara Falls, Ontario, Canada, May 21-23, 2007. (available at www.opentc.net)

- Gallery, E.; Mitchell, C.: Trusted Computing Technologies and their Use in the Provision of High Assurance SDR Platforms. In: Proceedings of the 5th Software Defined Radio Technical Conference (SDR 2006), Orlando, Florida, USA, 13-17 November 2006. Invited paper. (available at www.opentc.net)

- Huber, U.; Sadeghi, A-R.; Wolf, M.: Security Architectures for Software Updates and Content Protection in Vehicles, accepted for Automotive Safety and Security 2006, Stuttgart, Germany.

- Jansen, B.; Ramasamy, H. V.; Schunter, M.: Flexible Integrity Protection and Verification Architecture for Virtual Machine Monitors. Presented at the 2nd Workshop on Advances in Trusted Computing, Tokyo, Japan. (available at www.opentc.net)

- Loehr, H.; Ramasamy, H. V.; Sadeghi, A-R.; Schulz, S.; Schunter, M.; Stüble, C.: Enhancing Grid Security Using Trusted Virtualization. Proceedings of the 4th International Conference on Autonomic and Trusted Computing (ATC-2007), to appear. (available at www.opentc.net)

  A preliminary version of the above paper was presented (not published) at the:

  - 2nd Workshop on Advances in Trusted Computing (WATC-2006), Tokyo, Japan, December 2006.

  And at the:

  - 1st Benelux Workshop on Information and System Security, Antwerpen, Belgium, November, 2006.

- Manulis, M.; Sadeghi, A-R.: Property-based Taming Lying Mobile Nodes, accepted for International Workshop on Trusted and Autonomic Computing Systems (TACS 2006) at 20th IEEE International Conference on Advanced Information Networking and Applications (AINA 2006), 18.-20. April, Vienna, Austria.

- Ramasamy, H. V.; Schunter, M.: Architecting Dependable Systems Using Virtualization. Workshop on Architecting Dependable Systems: Supplemental

Volume of the 2007 International Conference on Dependable Systems and Networks (DSN-2007), to appear. (available at www.opentc.net)

- Sadeghi, A-R.; Scheibel, M.; Stueble, C.; Winandy, M.: Design and Implementation of a Secure Linux Device Encryption Architecture, accepted to be presented at Linux-Tag 2006.

- Sadeghi, A-R.; Selhorst, M.; Stueble, C.; Wachsmann, C.; Winandy, M.: TCG Inside? - A Note on TPM Specification Compliance, accepted for The First ACM Workshop on Scalable Trusted Computing (STC'06).

- Sadeghi, A-R.; Scheibel, M.; Stueble, C.; Wolf, M.: Play it once again, Sam - Enforcing Stateful Licenses on Open Platforms, accepted to be presented at The Second Workshop on Advances in Trusted Computing (WATC '06 Fall).

Edited by the Institute for Technology Assessment and Systems Analysis, Forschungszentrum Karlsruhe, Germany, on behalf of the OpenTC research project consortium, in co-operation with all partners.
Editor: Arnd Weber, Forschungszentrum Karlsruhe GmbH, ITAS, Hermann-von-Helmholtz-Platz 1, D-76344 Eggenstein-Leopoldshafen, Telephone: + 49 7247 82 3737.
Contact: editor at opentc.net

The OpenTC-project is a research project supported by the European Commission, project IST-027635. Its 23 partners are: Technikon Forschungs- und Planungsgesellschaft mbH (project coordination, AT); Hewlett-Packard Ltd (technical leader, UK); AMD Saxony LLC & Co. KG (DE); Budapest University of Technology and Economics (HU); Commissariat à l'Energie Atomique – LIST (FR); COMNEON GmbH (DE); Forschungszentrum Karlsruhe GmbH – ITAS (DE); Horst Goertz Institute for IT Security, Ruhr-Universitaet Bochum (DE); IBM Research GmbH (CH); Infineon Technologies AG (DE); INTEK Closed Joint Stock Company (RU); ISECOM (ES); Katholieke Universiteit Leuven (BE); Politecnico di Torino (IT); Portakal Teknoloji (TR); Royal Holloway, University of London (UK); SUSE Linux Products GmbH (DE); Technische Universitaet Dresden (DE); Technische Universitaet Graz (AT); Technische Universitaet Muenchen (DE); Technical University of Sofia (BR); TUBITAK – UEKAE (TR); and University of Cambridge (UK).

For more information about the project, see: http://www.opentc.net

Feedback to the consortium: http://www.opentc.net/feedback

Archive of newsletters: http://www.opentc.net/newsletter

Subscription: To subscribe or unsubscribe to the newsletter, write an email to <subscribe at opentc.net> or <unsubscribe at opentc.net>.