Ericsson Review

The communications technology journal since 1924

2014 • 9

Trusted computing for infrastructure

October 24, 2014

ERICSSON

Trusted computing for infrastructure

The Networked Society is built on a complex and intricate infrastructure that brings distributed services, data processing and communication together, combining them into an innovative and more meaningful set of services for people, business and society. But combining services in such an advanced way creates new requirements in terms of trust. Trusted computing technologies will play a crucial role in meeting the security expectations of users, regulators and infrastructure owners.

MIKAEL ERIKSSON, MAKAN POURZANDI AND BEN SMEETS

Today's industries are in transformation and ICT is changing the game. New applications built from a combination of services, communication and virtualization are being rolled out daily, indicating that the Networked Society is becoming reality.

Communication is transitioning from a person-to-person model to a system where people, objects and things use fixed and mobile connections to communicate on an anything-to-anything, anywhere and anytime basis. But even though people and businesses are beginning to use and benefit from a wide range of innovative applications, the potentially massive benefits that can be gained by combining modern computing, web services and mobile communication have yet to be realized.

As we progress deeper into the Networked Society, people, systems and businesses will become ever more dependent on an increasingly wider range of internet and connected services. And so the fabric of the Networked Society needs to be built on solutions that are inherently secure, socially acceptable and reliable from a technical point of view.

Modern internet services rely on web and cloud technology, and as such they are no longer independent packages with in-built security, but are constructed through the combination and reuse of other services distributed across the web. This creates new issues in terms of security. One of the most fundamental of these issues is securing processing in the communication infrastructure so that it can be trusted. Solving this issue is a prerequisite for building trust relationships into a network fabric for data communication and cloud computation. The red arrows in **Figure 1** illustrate possible trust relationships in such a network fabric that connects servers, data centers, controllers, sensors, management services, and user devices.

Trusted computing concepts

Users and owners of processing nodes use trusted computing to assess the certainty of one or several of the following aspects:

- what the processing nodes do;
- how nodes protect themselves against threats; and
- who is controlling the nodes.

This includes determining where data is stored and processed – which can be significant when legal or business requirements related to data handling need to be met.

This article presents an overview of the technical solutions and approaches for implementing trusted computing in a telecommunications infrastructure. Some of the solutions follow the concepts outlined in the Trusted Computing Group (TCG) specifications. Together the solutions described here enable what is often referred to as a Trusted Execution Environment (TEE), and with the addition of platform identities they provide a means for secure access control and management of platforms.

BOX A Terms and abbreviations

BIOS	basic input/output system
CBA	Component Based Architecture
DoS	denial-of-service
DRM	Digital Rights Management
DRTM	dynamic RTM
HE	Homomorphic Encryption
MME	Mobility Management Entity
OS	operating system
PKI	public key infrastructure
ROM	read-only memory
RoT	Root of Trust
RTM	RoT for measurement
RTR	RoT for reporting
RTS	RoT for storage
SDN	software-defined networking
SGSN	Serving GPRS Support Node
SGSN-MME	Network node combining SGSN and MME functions
SGX	Software Guard Extensions
SICS	Swedish Institute of Computer Science
SLA	Service Level Agreement
SRTM	static RTM
SSLA	Security Service Level Agreement
TCB	trusted computing base
TCG	Trusted Computing Group
TEE	Trusted Execution Environment
TLS	Transport Layer Security
TPM	Trusted Platform Module
TXT	Trusted eXecution Technology
UEFI	Unified Extensible Firmware Interface
VM	virtual machine
VMM	virtual machine manager (hypervisor)
vTPM	virtual TPM

In this article, the term platform is used to refer to the technical system for computational processing, communication and storage entities; which can be physical or virtual. The term infrastructure is used to refer to a wider concept, normally consisting of a collection of platforms and networks that is designed to fulfill a certain purpose.

Ensuring that the implementation of a technical system can be trusted calls for assurance methodologies. How to apply a security assurance methodology to every stage of product development, so that the implementation of a security-assurance product is in accordance with agreed guidelines has been discussed in a previous Ericsson Review article[1].

A model for trust

The infrastructure, which is illustrated in Figure 1, consists of servers, routers, devices and their computational, communication and storage aspects. This complex set of relationships can be redesigned using a cloud-based model – as shown in **Figure 2**. While the cloud model also consists of devices, access nodes, routing units, storage, servers and their respective management processes, the principles of trusted computing have been applied, and so the building blocks of each entity include trusted computing sub-functions.

Management functions govern the behavior of the platforms through a number of Security Service Level Agreements (SSLAs). For example, an SSLA might impose policies for booting, data protection or data processing. Through a trustworthy component known as Root of Trust (RoT), each entity locally enforces and checks for SSLA compliance. An RoT may be referred to as a trusted computing base (TCB) or trust anchor. It can be implemented as a hardware component, or exposed through a trusted virtual entity.

The RoT is one of the fundamental concepts of trusted computing for providing protection in the cloud model illustrated in Figure 2. Together with a set of functions, an RoT is trusted by the controlling software to behave in a predetermined way. The level of trust may extend to external entities, like management functions, which interact remotely with the RoT and contribute to establishing a trustworthy system.

FIGURE 1 Examples of trust relationships in the Networked Society

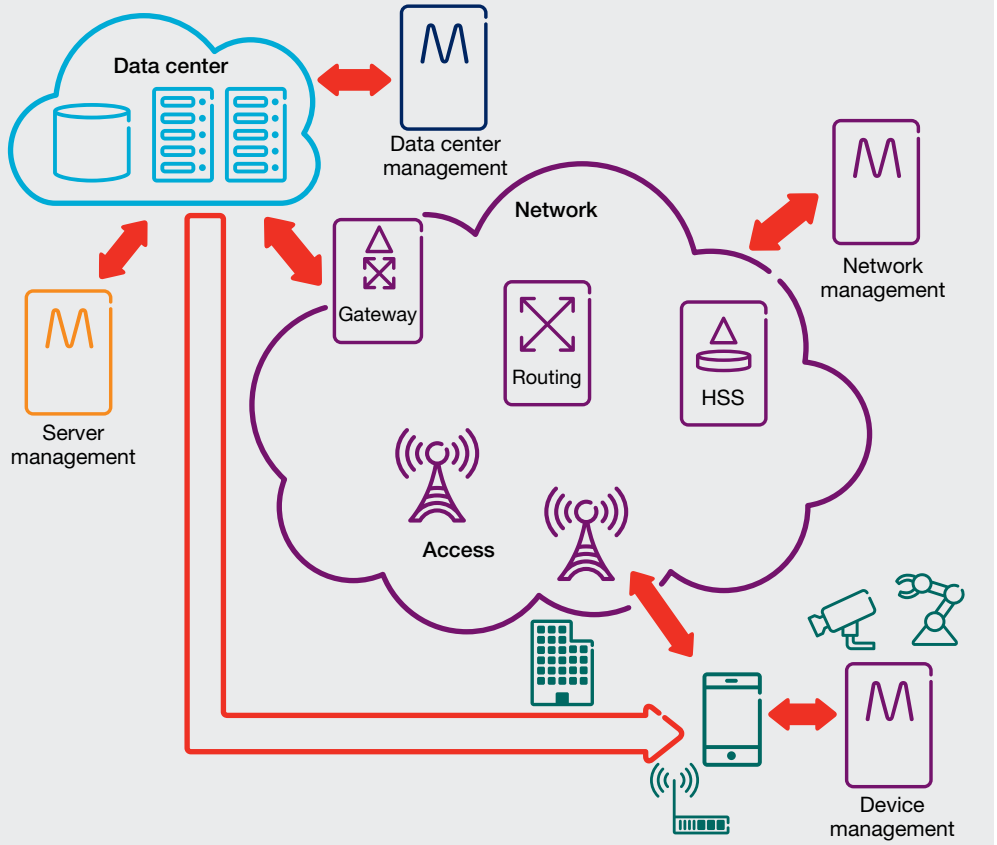

How the terms trust and trustworthiness are interpreted can be quite complex. They may depend on the results of an evaluation (such as Common Criteria methodology for Information Technology Security Evaluation[1]), or of a proof, and may even depend on the reputation of the organization or enterprise delivering the RoT. An RoT can provide several functions, such as:

- verification of data authenticity and integrity;
- provision and protection of secure storage for secret keys;
- secure reporting of specific machine states; and
- secure activation.

In turn, these functions allow features such as boot integrity, transparent drive encryption, identities, DRM protection, and secure launch and migration of virtual machines (VMs) to be built.

The implementation of an RoT must be able to guarantee a certain level of assurance against modification. A good example of this is the ROM firmware that loads and verifies a program during a boot process. The TCG approach to trusted computing relies on the interaction of three RoTs to guarantee protection from modification – each one with a specific task (**see Box C**):

- storage – the RoT for storage (RTS);
- measurement – the RoT for measurement (RTM); and
- reporting – the RoT for reporting (RTR).

How these RoTs are implemented is highly dependent on the Trusted Platform Module (TPM) and the cryptographic keys that are used to secure device hardware.

FIGURE 2 **A trusted computing cloud model**

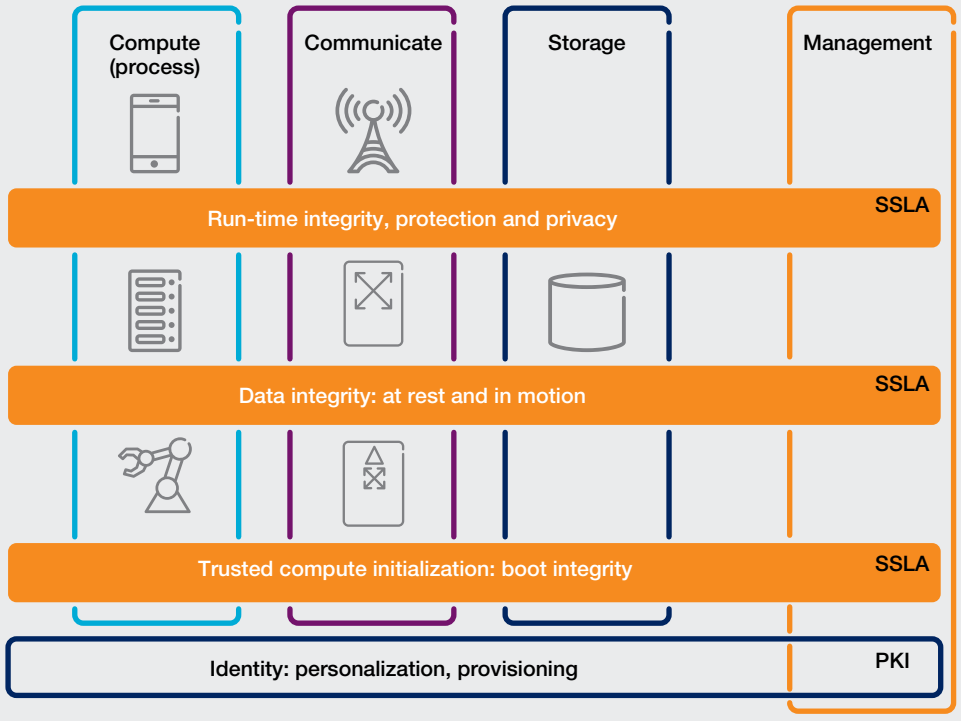

Measurement

The RoT for measurement – RTM – is defined in the platform specification and provides the means to measure the platform state. It comes in two flavors: static and dynamic – SRTM and DRTM, respectively. Intel's TXT, for example, is a DRTM; it supports platform authenticity attestation and assures that a platform starts in a trusted environment. The RTM is a crucial component for ensuring that a platform is in a trusted state. In contrast to the reporting and storage RoTs, the RTM resides outside the TPM – see Box C. A DRTM can be used to bring a platform into a trusted state while it is up and running. Whereas the static flavor starts out from a trusted point, based on a fixed or immutable piece of trusted code as part of the platform boot process.

Chipset vendors and platform manufacturers decide what flavor the RTM should be implemented in – static or dynamic. The implementation of Intel's TXT, for example, includes many adaptations in the chipset, and even uses Intel propriety code.

A TPM is often implemented as a separate hardware component that acts as a slave device. However, it can be virtualized, and in this case is often referred to as a vTPM (see[2], for example). To implement an RoT, there are other solutions than strictly following the TCG approach, such as those built using the ARM TrustZone concept. TrustZone can itself be used to implement an RoT as an embedded TPM with the functions mentioned in Box C.

Business aspects

In the Networked Society, cloud computing and cloud-based storage will be widely deployed. These technologies rely on a trustworthy network fabric; however, in a recent survey of the Open Data Center Alliance, 66 percent of the members stated that they are concerned about data security[3]. The upshot of this has been a delay in the adoption of cloud computing. Consequently, the use of trusted computing in existing and emerging cloud solutions is highly desirable, as it will help to dispel the fears associated with data security, lead to increased service use and new business models, and create opportunities for technological leadership.

Other business aspects influencing trusted computing solutions include requirements for scalability and elasticity of cloud computing, and the extent to which processing will be self-governed.

In the cloud

Trusted computing in a cloud environment is a special case. Web services and programmable routing technology (SDN based) using infrastructures like the one illustrated in Figure 1, will be deployed on platforms that exploit virtualization. To ensure overall security in the cloud, both the launch and the operation of virtualized resources need to be secure.

With respect to Figure 2, three core features are essential for building trusted computing in a cloud environment:

- **boot integrity – so that the hardware platform can guarantee a trustworthy RoT for the overall cloud environment;**
- **secure management of VMs – to secure the launch and migration of VMs in the cloud environment; and**
- **secure assessment of VMs – to attest the security and trustworthiness of VMs throughout their life cycles.**

Boot integrity

To boot a platform in a trustworthy way, a bootstrap process that originates from an immutable entity – an RoT – must be used. If the RoT provides proof of the progress of the bootstrapping process to the user in some transparent way, it acts as a measurement RoT.

There are two main approaches to the bootstrapping process: a verified boot or a measured boot.

A verified boot actively attests each component before it is loaded and executed. Using this approach, a platform will either boot or fail, depending on the outcome of the verification of each component's cryptographic signature.

Measured boot, on the other hand, is passive. Here, each component is measured and progress reports are saved into safe storage. The controlling process can then parse the recorded measurements securely and determine whether to trust the platform or not. Of the two approaches, only measured

boot complies with TCG; measurements combined with attestation are referred to as a trusted boot.

Both approaches can be used independently, or combined in a hybrid version to extend the integrity of the boot to client applications – which is illustrated in **Figure 3**. At Ericsson, ongoing work in Component Based Architecture (CBA) aims to establish a common approach to boot solutions and signed software; coordinating use in products.

Secure launch

Security-sensitive users need assurance that their applications are running on a trustworthy platform. Such a platform provides a TEE and techniques for users to attest and verify information about the execution platform.

In some cases, clients may want to receive an attestation directly from the platform. To do this, users need to be provided with a guaranteed level of trust in hardware or the virtualization layer during the initial VM launch, as well as throughout the entire VM life cycle – migration, cloning, suspension and resumption.

To launch a VM in a secure way, the security and trustworthiness of the hardware platform and virtual layer first need to be attested. For certain sensitive applications, like financial transactions or handling legal intercept, the VM or the owner of the VM need to be advised on the trustworthiness of the hardware platform each time the hardware platform is changed – for example following the migration, suspension or resumption of a VM.

In a cloud environment, some additional security constraints may apply to a VM launch. For example, due to the risk of a side channel or a DoS attack, some customers may require their virtual resources to be separated (not co-located) from any other customer's resources.

There are basically two of ways of attesting a secure VM launch to clients:

- the cloud provider can deploy the trusted cloud and prove its trustworthiness to the client; or
- trustworthiness measurements can be conveyed to the client – either by the cloud provider or by an independent trusted third party.

In the first approach, customers must trust the cloud provider. The difficulty with the second approach is the ability of a customer or trusted third party to collect the trustworthiness evidence related to the cloud providers – given the dynamic nature of the cloud and the diverse set of hardware, operating systems (OSs) and VM managers (VMMs) used. This task becomes even more complex because trustworthiness needs to be reestablished and checked every time a change occurs in the underlying layers: hardware, OS, and VMM. It seems inevitable that for the second approach to work, cloud providers would have to expose some, or all of their internal hardware and software configuration, including, say, hardware platform specifics, OS, VMM, and even configuration information and IP addresses. This may conflict with a cloud provider's policy to keep its internal architecture private.

The solution presented in Huebner on Intel TXT[4] is of the first type – based on trust. Here, attestation is achieved inside the cloud environment, and the results are then provided to users.

The BIOS, OS, and hypervisor of the hardware platform are measured, and the results are sent to an attestation server. The server in turn verifies their trustworthiness by comparing them against a known database of measurements. Following successful verification, the secure VM launch can then be carried out.

When attestation is achieved through the trust model, users cannot remotely attest the hardware platform and consequently have to trust the cloud provider and its attestation through SLAs.

To attest a secure VM launch using the second approach – based on measurement – Ericsson security researchers have created a framework[5,6] in OpenStack to verify the trustworthiness of VM host system software through remote attestation with a trusted third party.

FIGURE 3 Hybrid boot process using an RoT for measurement

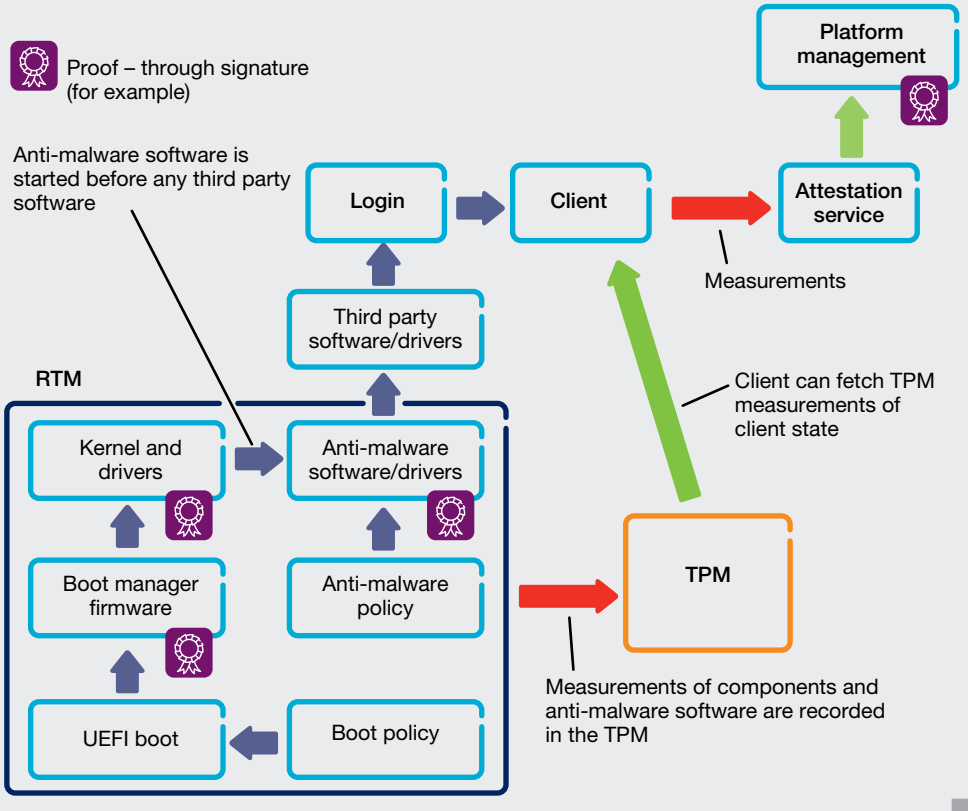

FIGURE 4 **Trusted computing attestation process**

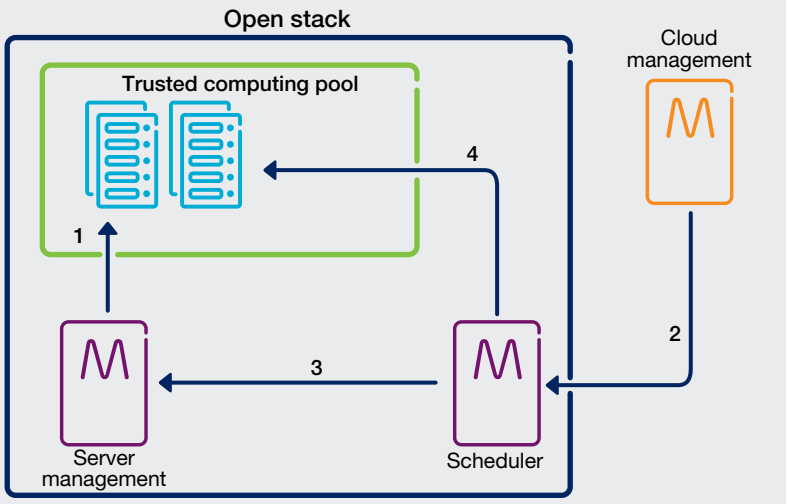

Secure migration

In a cloud environment, VM migration is often necessary to optimize the use of resources and ensure optimal power consumption. This is a highly dynamic process that depends on many factors, including application needs, host loads, traffic congestion, and hardware and software failures. A secure VM migration ensures the security of the VM both at rest and during the migration – guaranteeing the same level of trust before and after. Similarly, cloud federation use cases require interoperability guarantees among the different cloud service providers. To achieve this, mechanisms need to be in place to ensure the same level of trust when a VM is migrated from one cloud provider to another.

Migrating a VM can sometimes result in a change of underlying hardware that the VM is not aware of. This is significant, as the RoT function can depend on both hardware and VMM (when it comes to virtual TPM deployment for VMs). Migrations are often performed programmatically by cloud orchestration or management in a manner that is transparent to the VM. So, cloud orchestration and management need to be involved to choose the right physical hosts and VMMs with adequate levels of trust expressed in SSLAs to run VMs.

For regulation or auditing purposes, preserving proof of trustworthiness of the platform needs to be provided for security-sensitive applications. This use case can be extended to a remote attestation of HW-VMM-VM to the tenant's auditor. There are two aspects related to preserving trustworthiness:

- ensure that the hardware and VMM after the migration can be trusted to preserve the same level of trust (trusted computing base) for VM before and after the migration; and
- provide the same RoT functionality to a VM before and after migration: for example, protection and storage of secret keys in a virtual TPM.

So far, secure VM migration has received less attention than the secure launch from both academia and industry. Despite this lack of interest, secure VM migration is an essential part of the overall secure life cycle of VMs, if satisfactory levels of security for applications in the cloud are to be achieved.

Secure assessment

From a management point of view, the platform needs to provide trustworthiness information and provide assurance that it responds correctly to management commands. Remote assessment of the platform state is of particular importance to ensure that the launch or migration of a virtual machine is carried out securely.

BOX B

Trusted computing attestation process

1) the Open Attestation Server determines a trusted computing pool;
2) cloud management requests new workloads from the scheduler;
3) the scheduler requests the list of trusted computing nodes in the trusted computing pool; and
4) the workload is initiated on a computing node inside the trusted computing pool.

Obtaining assurance for every single functional aspect of the platform and the services it hosts can be difficult. Obtaining assurance for just a limited set of functions can reduce the complexity of this task and be an acceptable trade-off. Ideally, those aspects that have security relevance should be expressed in an agreement between the provider and the user – typically detailed in an SSLA, which might demand the support of remote assessment procedures. For this, a platform should have a set of mechanisms, like RTM coupled to RTR, that allow a remote entity to securely assess certain properties recorded by monitoring capabilities of the platform's local trustworthy subsystem. Yet proper assurance methodologies have to be applied to ensure that these mechanisms deliver what is needed without any blind spots, which would result in a false sense of security.

Implementation aspects

Standards

Although extensive academic work has been carried out in the field of trusted computing, only a few implementation standards exist for interoperable trusted computing solutions. The TCG has specified a framework and components for implementing trusted computing, which are used by chipset vendors such as Intel and AMD. However, the TCG specifications can result in varying implementations by the different vendors, which is good, as different vendors can optimize their solutions for different capabilities such as for performance or for storage. While this flexibility is advantageous, it also creates interoperability issues[7]. Flexibility has been further increased in TPM 2.0 through implementation and choice of cryptographic primitives. Currently, the TCG specifications remain the most comprehensive standards for implementing RoTs.

Another important set of specifications has been issued by the GlobalPlatform organization. Its TEE specifications include architecture for secure computation and a set of APIs. Although these specifications provide trusted computing for mobile devices, they can also be used for infrastructure nodes such as base stations. How the secured environment is actually

implemented is left to the discretion of the hardware vendors and can be system-on-chip or a dedicated separate component; ARM TrustZone is an example implementation of this technology. As of mid-2014, the GlobalPlatform specifications do not address how a system reaches a trustworthy state and how trust properties can be asserted. With this in mind, the GlobalPlatform and TCG specifications complement each other.

Hardware aspects

As illustrated by the Intel TXT implementation of the TCG DRTM concept, several components in general purpose chipsets must be modified to achieve the needed protection. Similarly, the protection provided by TrustZone affects the ARM core as well as its subsystems. This level of invasiveness results in hardware vendors sticking to their chosen approach to trusted computing, and changes to functionality tend to be implemented in a stepwise fashion. Intel and AMD have been using TCG functionality, and ARM has pursued its TrustZone concept and announced cooperation with AMD.

Unfortunately, the TCG specifications do not really cover the aspects of isolation of execution. To fill this gap, Intel introduced the SGX concept, which is a set of new CPU instructions that applications can use to set aside private regions of code and data.

Isolation during execution is an important principle, and future hardware will have more functionality to improve isolation and control of the execution environments. The SGX concept also supports attestation and integrity protection, as well as cryptographic binding operations per private region.

Homomorphic Encryption

In some (cloud) processing cases, it might be possible to apply what is referred to as Homomorphic Encryption (HE) as an alternative to applying stringent secrecy demands on processing nodes. Current research in this subject and similar techniques appear to be promising – leading to reasonably fast cloud-based processing of secret (encrypted) data for certain operations without needing to make the data available in clear text to the processing node.

However, HE is a rather undeveloped technology; it only solves certain aspects of trusted computing, and involves a level of computational complexity that is, generally speaking, still too high. It may, however, become a complementary technique for trusted computing. If that happens, hardware support for HE operations will likely find its way onto server chipsets.

Examples of platform security

In cooperation with the Swedish Institute of Computer Science (SICS), Ericsson Research has modified OpenStack to use a TPM for secure VM launch and migration. A trusted third party was used for collecting and sending trustworthy information and control. Part of the solution has been used in a cloud-based test-bed setup for a regional health care provider in southern Sweden.

Ericsson security researchers have also implemented solutions for cloud-based protection of persistent storage[8]. Generally speaking, secure VM launch and migration are finding their way into OpenStack.

The coming release of the Ericsson SGSN-MME node is another example of how trusted computing has been implemented using TPM technology. Beyond the functionality discussed above, the TPM is used for secure storage of PKI credentials. These credentials are used for TLS connections and for encryption of sensitive data. Like other telco nodes, the SGSN-MME has high-availability requirements, which calls for the use of hardware redundancy and efficient maintenance procedures. As the TCG specifications do not address such use cases, special care must be taken when deploying TPMs in such a setting: production, personalization, rollout, and maintenance support have to be implemented before any of the trusted computing features can be enabled.

Conclusion

Ericsson recognizes that trusted computing is a technical approach that will enable secure infrastructures and services for the Networked Society. As the use of virtualization technologies and the cloud increases, maintaining trust is essential. In connection with the cloud, the use of a virtual trusted platform model as an RoT for different virtual machines has received some attention from both academia and industry. Despite this, further development is required to address issues related to establishment of trust models, trusted evidence collection, and real-time and dynamic attestation. Ericsson Research is active in this field and cooperates with Ericsson business units to incorporate such security solutions into products.

BOX C Three main TPM tasks

The TPM is responsible for protecting secret keys and sensitive functions. The bulk of the TPM's data is stored outside the TPM in so-called blobs. The RTS provides confidentiality and integrity protection for these blobs. The RTR is responsible for:

- reporting platform configurations;
- protecting reported values;
- providing a function for attesting to reported values; and
- establishing platform identities.

The interaction between the RTR and RTS relates to the responsibility for protecting measurement digests. The term measurement has a specific meaning in TCG and can be understood as verification in relation to RTM functions.

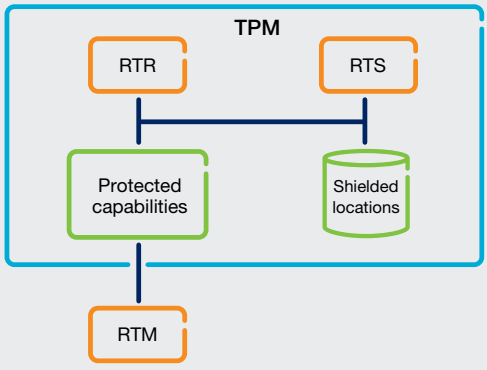

Mikael Eriksson

is a security architect at Business Unit Cloud & IP. He holds an M.Sc. in data and image communication from the Institute of Technology: Linköping University, Sweden. He joined Ericsson in 2009 to work with mobile broadband platforms after an 18-year career as a consultant, mostly in embedded systems. Since 2012, he has been with the Packet Core Unit, working on adaptation of security technology in mobile networks infrastructure. He is currently the study leader of a boot integrity integration project of Ericsson platforms.

Makan Pourzandi

works at Ericsson Security Research in Montreal, Canada. He has more than 15 years of experience in security for telecom systems, cloud and distributed security and software security. He holds a Ph.D. in parallel computing and distributed systems from the Université Claude Bernard, Lyon, France, and an M.Sc. in parallel processing from École Normale Supérieure (ENS) de Lyon, France.

Ben Smeets

is an expert in security systems and data compression at Ericsson Research in Lund, Sweden. He is also a professor at Lund University, from where he holds a Ph.D. in information theory. In 1998, he joined Ericsson Mobile Communication, where he worked on security solutions for mobile phone platforms. His work greatly influenced the security solutions developed for Ericsson Mobile Platforms. He also made major contributions to Bluetooth security and platform security related patents. In 2005, he received the Ericsson Inventors of the Year award and is currently working on trusted computing technologies and the use of virtualization.

Acknowledgements

The authors gratefully acknowledge the colleagues who have contributed to this article: Lal Chandran, Patrik Ekdahl, András Méhes, Fredric Morenius, Ari Pietikäinen, Christoph Schuba, and Jukka Ylitalo

References

1. Ericsson Review, Setting the standard: methodology counters security threats, January 2014, available at: http://www.ericsson.com/news/140129-setting-the-standard-methodology-counters-security-threats_244099438_c
2. Stefan Berger, Ramón Cáceres, Kenneth A. Goldman, Ronald Perez, Reiner Sailer, Leendert van Doorn, vTPM: Virtualizing the Trusted Platform Module, RC23879 (W0602-126) February 14, 2006, Computer Science IBM Research Report, available at: https://www.usenix.org/legacy/event/sec06/tech/full_papers/berger/berger.pdf
3. Tech Times, Cloud computing is the future but not if security problems persist, June 2014, available at: http://www.techtimes.com/articles/8449/20140615/cloud-computing-is-the-future-but-not-if-security-problems-persist.htm
4. Christian Huebner, Trusted Cloud computing with Intel TXT: The challenge, April 16, 2014, available at: http://www.mirantis.com/blog/trusted-cloud-intel-txt-security-compliance/
5. Mudassar Aslam, Christian Gehrmann, Mats Bjorkman, Security and Trust Preserving VM Migrations in Public Clouds, 2012 IEEE 11th International Conference on Trust, Security and Privacy in Computing and Communications, June 25-27, 2012, available at: http://ieeexplore.ieee.org/xpl/articleDetails.jsp?reload=true&arnumber=6296062
6. Nicolae Paladi, Christian Gehrmann, Mudassar Aslam, Fredric Morenius, Trusted Launch of Virtual Machine Instances in Public IaaS Environments, 15th Annual International Conference on Information Security and Cryptology, 2013, available at: http://soda.swedish-ict.se/5467/3/protocol.pdf
7. TrouSerS, the open source TCG Software Stack, I've taken ownership of my TPM under another OS..., available at: http://trousers.sourceforge.net/faq.html#1.7
8. Nicolae Paladi, Christian Gehrmann, Fredric Morenius, Domain-Based Storage Protection (DBSP) in Public Infrastructure Clouds, 18th Nordic Conference, NordSec, October 18-21, 2013, available at: http://link.springer.com/chapter/10.1007%2F978-3-642-41488-6_19#page-1

Ericsson Review

To bring you the best of Ericsson's research world, our employees have been writing articles for Ericsson Review – our communications technology journal – since 1924. Today, Ericsson Review articles have a two-to-five year perspective and our objective is to provide you with up-to-date insights on how things are shaping up for the Networked Society.

Address:
Ericsson
SE-164 83 Stockholm, Sweden
Phone: +46 8 719 00 00

Publishing:
Ericsson Review articles and additional material are published on: www.ericsson.com/review. Use the RSS feed to stay informed of the latest updates.

Ericsson Technology Insights

All Ericsson Review articles are available on the Ericsson Technology Insights app available for Android and iOS devices. The link for your device is on the Ericsson Review website:www.ericsson.com/review. If you are viewing this digitally, you can:
download from Google Play or
download from the App Store

Publisher: Ulf Ewaldsson

Editorial board:
Hans Antvik, Ulrika Bergström, Joakim Cerwall, Stefan Dahlfort, Åsa Degermark, Deirdre P. Doyle, Dan Fahrman, Anita Frisell, Geoff Hollingworth, Jonas Högberg, Patrick Jestin, Cenk Kirbas, Sara Kullman, Börje Lundwall, Hans Mickelsson, Ulf Olsson, Patrik Regårdh, Patrik Roséen, Gunnar Thrysin, and Tonny Uhlin.

Editor:
Deirdre P. Doyle
deirdre.doyle@jgcommunication.se

Subeditors:
Paul Eade, Nathan Hegedus and Ian Nicholson

Art director and layout:
Carola Pilarz

Illustrations:
Claes-Göran Andersson

ISSN: 0014-0171

Volume: 91, 2014

Ericsson
SE-164 83 Stockholm, Sweden
Phone: + 46 10 719 0000

ISSN 0014-0171
284 23-3233 | Uen