



[www.isg.rhul.ac.uk](http://www.isg.rhul.ac.uk)



[www.mobilevce.com](http://www.mobilevce.com)

# Open Workshop on Trusted Computing

Information Security Group  
Royal Holloway  
University of London

*in association with:*

Mobile VCE  
the Virtual Centre of Excellence in  
Mobile & Personal Communications

30th March 2004



[www.isg.rhul.ac.uk](http://www.isg.rhul.ac.uk)



[www.mobilevce.com](http://www.mobilevce.com)

## Workshop programme

9:30 Registration (refreshments provided)

10:00 *Welcome* Professor Fred Piper (Director, ISG) and  
Dr Walter Tuttlebee (Chief Executive, Mobile VCE)

### Part 1: Trusted / trustworthy computing – background and technology

10:10 *An introduction to trusted computing* Chris Mitchell

10:30 *The TCPA/TCG trusted platform* Eimear Gallery

11:20 Coffee

11:40 *NGSCB* Eimear Gallery

12:30 Buffet lunch

### Part 2: Applications of trusted computing

13:30 *Single Sign-on using trusted platforms* Andreas Pashalidis

14:10 *Secure Content Management using trusted computing* Allan Tomlinson

14:50 Tea

15:10 *Protecting User Privacy using trusted computing* Anand Gajparia

15:50 *Distributing PKI Functionality using trusted computing* Alex Dent

16:30 *Closing Remarks*



[www.isg.rhul.ac.uk](http://www.isg.rhul.ac.uk)



[www.mobilevce.com](http://www.mobilevce.com)

## *Mobile VCE*

Walter Tuttlebee

Chief Executive

Mobile VCE

<http://www.mobilevce.com>

## **Virtual Centre of Excellence in Mobile and Personal Communications - Mobile VCE -**

### **MAXIMISING BENEFITS FOR OUR INDUSTRY MEMBERS**

Mobile VCE's member companies value and utilise different aspects of their membership very differently, reflecting their size, the nature of their business, their product and R&D strategies, recruitment needs, etc. By pro-actively managing their involvement to reflect their own specific objectives they achieve strong and different benefits.

Mobile VCE encourages all our industrial members to consciously consider their goals and priorities in this respect. This enables them make appropriate choices regarding participation in events and to identify appropriate procedures - eg establishment of appropriate in-company 'gatekeepers' and communication routes – to maximise the benefits of membership to their organisation.

#### **MEMBERSHIP BENEFITS**

Examples of emphases and benefits currently enjoyed by specific companies include:

- access to R&D to set the strategic direction of their own in-house R&D
- (or conversely) access to R&D in place of in-house R&D
- high financial leverage - typically 40:1 gearing - secured by matching the company's subscription with those of ~20 others, plus government grants
- access to software tools spinning out from the research programme
- access to IPR and to new technology ideas – 'seed & feed'
- opportunity to recruit from the top research talent pool, familiar with industry
- access to specialist research consultants
- the ability to commission company-specific research – 'elective R&D'
- training of their own research staff by working with Mobile VCE's researchers
- building relationships with other industry players
- identifying future trends, threats and opportunities – 'window on the future'
- advice on sources of specific technical expertise

#### **YOUR GOALS**

We encourage new (and existing) members to talk with us about how we can work together to maximise the specific benefits of membership to your own company. Through its innovative structure and mechanisms, Mobile VCE has not only created a unique and effective academic-industry research partnership but is also strengthening industry relationships and creating new opportunities for all our members.





**future.wireless@mobilevce.com**

*Walter Tuttlebee, Executive Director, Mobile VCE*

*The shape of future mobile and personal communications may not be the simple generational shift that characterised past transitions. In October 2003 the ITU-R Working Party 8F, responsible for 'IMT2000 and Systems Beyond' hosted a Workshop on Services & Market Aspects, jointly organised with Mobile VCE, to consider such issues [1]. In this article Walter Tuttlebee describes the role of Mobile VCE and its research programmes. An important aspect of this is the way that the differing regional perspectives on the future held by its member companies shape Mobile VCE's own industry-led research programme, which aims to enable new services and revenue streams for the industry.*

## What is Mobile VCE?

Mobile VCE – the Virtual Centre of Excellence in Mobile & Personal Communications – is a not-for-profit company, established as a joint initiative of the mobile communications industry and leading research universities in 1996. Initially a national initiative, Mobile VCE today has a global industry membership base and its research is having international impact.

The sustained involvement of member companies from Asia, America and Europe means that Mobile VCE is unique in undertaking an integrated industry-driven research programme, which accommodates diverse perspectives of 'B3G/4G'. This permits the identification of new synergies between these approaches which would otherwise be overlooked.

The innovative nature of the organisation provides companies with a unique environment for the development of shared visions and with a highly cost-leveraged approach to industry-led academic research, supporting the future development of the wireless telecoms industry. This brief introduction to the company – its origins, its objectives, its membership and its modus operandi – is complemented by an overview of the evolution of its research priorities at a

time of significant change, as the industry begins to implement 3G services and systems.

Mobile VCE offers to its industry members a range of benefits including *inter alia* high financial gearing of research funding, a favourable IPR arrangement, ready access to highly talented academic research teams and opportunities to exploit research spin-off for shorter-term objectives.

## Origins & Objectives

Historically the mobile telecoms industry, like others, viewed research as something to be undertaken by company research departments or subcontracted to universities. In Europe collaborative research with universities utilised mechanisms such as the Framework programmes (RACE, ACTS, IST, etc) [2]. The apparent dichotomy between the long-term academic research culture and the short-term development-focused mentality of industry is well recognised. It was from this background that representatives of the mobile phone industry began to explore alternative models as part of the UK Government's Foresight exercise in late 1995. At that time companies were placing direct research contracts with Universities, but such work was often uncoordinated and disconnected from mainstream industry thinking. From this came the concept that something more effective and substantial could be done together – but what?

By late 1996 ideas had clarified and a nucleus of industry players had coalesced around the idea which subsequently became the Virtual Centre of Excellence in Mobile & Personal Communications – Mobile VCE. The emerging concept was for a 'virtual company', funded by industry, pioneering an industry-steered programme of research, undertaken by integrated, but geographically-dispersed, research teams drawn from the country's top universities specialising in mobile telecoms research. Mobile VCE was thus proposed as part of the Foresight Challenge competition. Clear industry support and the evident benefits of the proposed approach were key factors in the subsequent Government decision to award Foresight funding, which led to the formal establishment of the company in November 1996.

Mobile VCE's strategic objective is to facilitate industry growth, through pioneering long term research directed by shared industry vision, and, within this framework, to create commercially valuable IPR relevant to mobile / wireless communications for its industrial members and, most importantly, with their active involvement. The motivation, creativity and collaborative approach of the research teams, matched by the commitment of the major industry players, has created a momentum and environment that encourages open participation from the mobile telecommunications companies – perhaps one of the key factors in its success in anchoring the research into industry evolution.

|                            |                                   |              |
|----------------------------|-----------------------------------|--------------|
| BAE Systems                | BBC                               | Crown Castle |
| Degree2 Innovations        | Fujitsu                           | Hutchison 3G |
| Inmarsat                   | Independent Television Commission | Lucent       |
| NEC                        | Nokia                             | Nortel       |
| Orange                     | Panasonic                         | Philips      |
| Radiocommunications Agency | Samsung                           | Siemens      |
| SK Telecom                 | Sony                              | Thales       |
| Toshiba                    |                                   | Vodafone     |

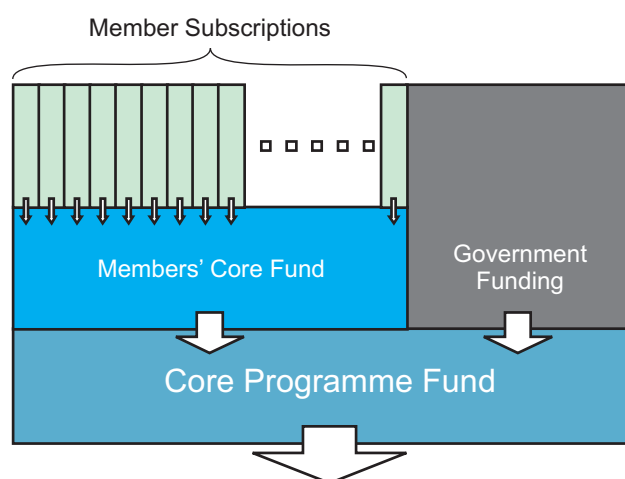
Table 1: Industrial Members of Mobile VCE, FY2003

- ♦ Industrially Relevant, Long Term, Research – defined and steered by the industrial member companies
- ♦ Highly Cost Effective – due to the significant Financial Gearing of Members' Subscriptions (See Figure 1)
- ♦ Intellectual Property – royalty-free access to the IPR portfolio is available to full industrial members
- ♦ Short Term Technology Spin Offs – structures in place for commercial exploitation
- ♦ Elective Research Opportunities – funded by a single company, or by a small group of companies on a shared-cost, shared-benefit, basis, to address specific R&D needs outside of the Core programme
- ♦ Development of Company Staff – through working alongside highly competent academic researchers
- ♦ Recruitment – access to identify and recruit leading academic research talent
- ♦ Networking – to identify and shape, with industry colleagues, the trends which will influence the short and long term

Table 2: Benefits of Industrial Membership

|                           |                              |                              |
|---------------------------|------------------------------|------------------------------|
| University of<br>Bristol  | University of<br>Edinburgh   | Kings College<br>London      |
| Royal Holloway<br>College | University of<br>Strathclyde | University of<br>Southampton |
|                           | University of Surrey         |                              |

Table 3 Academic Members of Mobile VCE



**Figure 1** High Financial Gearing: The Core Research Programme is funded by the aggregation of members' subscriptions, further augmented by Government support

## Membership


Industrial Members join Mobile VCE on payment of an annual membership fee and acceptance of the Memorandum and Articles of Association of the company, and upon signing a Deed of Adherence to the company's IPR Agreement. In 2003 Mobile VCE had 23 industrial members, including many top international players in the mobile telecommunications industry – see Table 1 – each paying an annual subscription of £36k / £72k, for Class 1 / Class 2 membership respectively. Members' annual subscriptions sustain 'Core' research programmes, with membership viewed as a long-term commitment, rather than simply as a one-year decision; the costs, taken over a typical Core Programme period of 3-4 years, are however remarkably low compared to the benefits.

Mobile VCE delivers value to its Industrial Members in many ways, the main ones being summarised in Table 2. As noted earlier, the philosophy of Mobile VCE ties its research into industrial reality – the contribution of the in-

dustry staff, whilst relatively low in terms of manpower, is crucial in shaping the direction and value of the research. Different companies choose to participate with different levels of manpower, with those who play a more active role deriving proportionally greater benefits.

To date Mobile VCE has generated approaching 40 patent applications, of which 4 have reached the grant stage (3 European and 1 USA); 18 others have progressed to the national/regional stage. Industrial members have royalty-free exploitation rights to this IPR. A comprehensive library of technical reports is delivered each year on CD-ROM, covering all aspects of the Core Programme. This is complemented by a series of Members' Technical Seminars, providing opportunities for a much broader base of staff from Member companies to be familiarised with the results, tools and resources available to them through Mobile VCE. Raw data and processed results from propagation research campaigns (SISO, SIMO and MIMO) are also available to members, as are software tools relating to multimedia traffic planning, channel models and other topics.

These membership benefits are seen as of varying importance by different members, reflecting their size, position and role in the industry. For some companies membership provides a highly cost-effective means to monitor new research directions, threats and opportunities across a wide field, without the need to fund large in-house research activities. For others, quite the converse is true – the programme of Mobile VCE is used to inform a larger internal programme, to help shape its strategic direction. Another group sees royalty-free access to the growing pool of IPR as important, whilst for others, the interaction with the academic researchers is seen as a way of keeping their staff sharp and aware of new research trends or as a means of recruiting key research staff.



**Figure 2** Evolution of Mobile VCE's Core Research Agenda

Academic Members, who staff the research teams, must meet a tough set of objective criteria, a key one being a *proven and clearly demonstrable track record* – not just potential – of *excellence in relevant research*. Ongoing performance, assessed through formal review by the Industrial Members, determines the level and areas of participation in future research programmes, thereby ensuring Mobile VCE maintains the highest quality, as befits a true Centre of Excellence. Whilst such procedures mean that membership of Mobile VCE is not a ‘cosy club’ for academics – far from it – the universities clearly benefit from their participation, not least through the explicit recognition of quality that accompanies membership and which thus helps them to attract international, high quality, staff. Perhaps one of the most strategically important benefits for the universities, however, is the input regarding research directions that they receive from a wide breadth of Industry representatives through Mobile VCE’s Steering Groups.

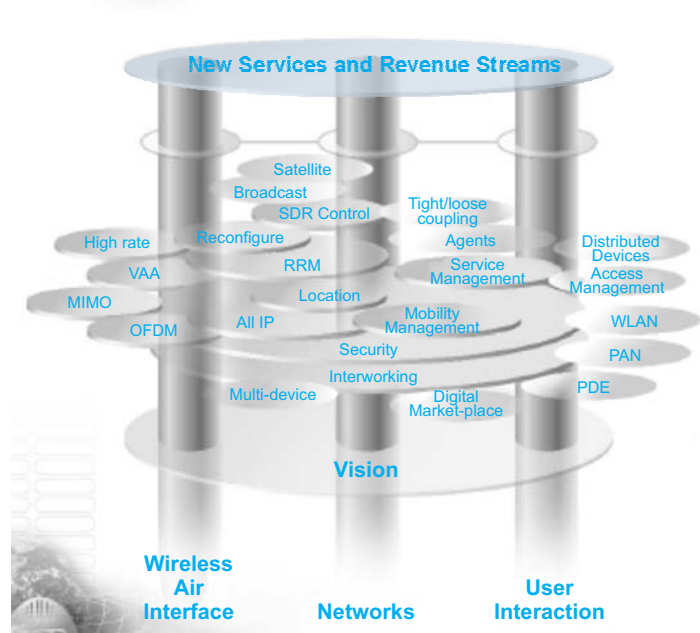
## Modus Operandi

Industrial Members have full access to the Mobile VCE Core Research Programme, technical reports and research teams, as well as the opportunity to participate in Industrial Steering Groups, established for each of the Core Programme work areas. Steering Groups typically comprise some 8-15 Industrial Members who meet quarterly to monitor, review and direct the research activities, identifying patent proposals and approving publications. Each Steering Group is chaired by a senior Industrial Member.

The appointment of senior industry technologists, from the operational as well as the research parts of their companies, to chair and participate in these groups allows for a very strong input to the academics. It secures a higher degree of industrial relevance than is achieved in most other models of collaborative research. And yet, the structure and ownership of Mobile VCE ensures that such industrial relevance is achieved without compromising either the long-term nature of the research or its academic integrity.

Whilst such involvement demands a high degree of commitment, experience has shown that those Industrial Members who participate actively in this way reap substantial benefits. Steering Group membership provides an avenue to direct the research, to secure early access to its results, to network with other members of the mobile industry on a regular basis and, from time to time, the opportunity to recruit top quality research staff.

The research within each work area is undertaken by collaborative teams, comprising personnel usually from two or more universities. Co-ordination of these pan-university teams is managed by an Academic Co-ordinator, a senior staff member – Lecturer or Professor – from one of the member Universities who acts in support of the Industrial Steering Group Chairman. Innovative ‘carrot and stick’ mechanisms have proven to be very effective in encouraging the universities to deploy high quality creative staff within the Mobile VCE teams, to collaborate closely across traditional boundaries and to identify valuable IPR. Joint pan-university research has achieved more than simply the sum of the parts, as unanticipated opportunities to leverage research activities have been identified. Complex software simulation tools have also been created in a composite manner across multiple research institutions. When Mobile VCE was first established such collaborative achievements were just hopes, but have become reality.



**Figure 3** Key Elements and Pillars of Mobile VCE's Core 3 Research Programme

## Core Research

Two types of research activities are undertaken by Mobile VCE, integrated research programmes available to all Members – ‘Core Research’ – and research for just one or a few Industrial Members – ‘Elective Research’. Both types of programme (described more fully below) are defined jointly by industrial members (usually specifying objectives) and by lead academics (usually proposing technical approaches). They are monitored and steered by Steering Groups and are undertaken usually by combined teams from more than one university, selected to match expertise to requirement. Formal research contracts are placed by Mobile VCE with its Academic Members. The research is monitored by the Industrial Steering Groups and overseen by the Mobile VCE office, which provides a central hub for management, communication and information dissemination.

Core Research Programmes represent substantial integrated programmes of research, funded by pooling members’ subscriptions, augmented by competitively-won Government grants (see Figure 1). The evolution of the scope of the Core Research Programmes is shown in Figure 2, which illustrates the way in which the research agenda has developed, to reflect the evolution of mobile services and consequent industry requirements.

The Core 1 Programme (1997-2000) was a 50 manyear programme, with four distinct work areas – Networks, Services, Terminals and Radio Environment [3]. Whilst the focus of this was primarily 3G systems, some of its outputs were used, for example, for 2.5G deployment, and other firmly beyond 3G. The vision for the Core 2 Programme (1999-2003) was described in [4], and was more substantial in scope, with an effort of over 100 man-years, structured into three work areas – Networks & Services, Wireless Access, Software Based Systems. The enlargement of the staffing for Core 2 enabled Mobile VCE to widen its academic base, embracing an increased emphasis on mobile computing, middleware, agent technology and security, complementing, and building upon, foundational strengths in wireless access, software radio, networks and services.



As we look forward at future evolution of the cellular industry, beyond the relatively short term evolutions being developed by 3GPP, three distinct geographical emphases have emerged over the past two years – these have been presented at various international conferences; indeed these perspectives have also been hotly debated within the ITU-R Working Party 8F over this same period [1]. From Asia a requirement has been identified for high bit rate wireless access (100Mb/s wide area, 1Gb/s short range), reflecting the rapid rollout of wired broadband and a sense that in the relatively near future users will require such access to be provided wirelessly. In Japan, NTT DoCoMo have implemented a testbed using their VSF-OFCDM wireless access technology [5], which continues to be trialled at the present time. In North America the dramatic growth of wireless LAN, combined with the wireless industry downturn, has resulted in a focus on shorter term issues such as WLAN-based technology concepts [6] and IEEE 802.20 wireless WAN [7]. Meanwhile, in Europe the focus, largely led by the early 6th Framework agenda, has tended to be upon reconfigurable networks [8] and interworking between mobile and other networks, such as digital broadcasting networks [9].

Reflecting the geographical composition of its industrial membership, Mobile VCE's Core 3 Programme (2002-2005), contains key elements of all the above themes, see Figure 3, which illustrates interrelationships between the main three work areas:

- ♦ 'Wireless Enablers' encompasses work on novel air interface technologies and implementations, for both high-speed PAN and wide area coverage. Part of this approach includes a novel MIMO-based propagation measurement campaign and exploration of flexible multi-standard hardware implementation.
- ♦ 'Interworking of Networks' includes research on secure internetworking of mobile networks with digital broadcast, WLAN and personal networks. This is based upon an 'arms-length', rather than convergent, philosophy, reflecting the industrial reality of the distrust that exists between the broadcast and telecommunications industry – both believe in convergence, but each industry would prefer it to be under their own control [10].
- ♦ The 'Personal Distributed Environment' distinctive is its focus upon the challenges of the user's individual information environment, to provision new services in a world of multiple short-range interlinked wireless devices and ubiquitous service access. Such personalised services will become increasingly essential to ARPU growth in the saturated voice markets of the developed world, but to succeed will require easy, flexible and efficient management of user services and devices.

## Elective Programmes

In addition to its Core research, Mobile VCE also undertakes 'Electives' – research programmes of interest to a subset of the industrial membership and in which members may choose ('elect') or decline to participate. Such programmes are funded by either a single company or jointly by a group of companies with common interest. For such programmes, access to results and any consequent IPR is restricted to those who elect to fund the work. Such programmes may build on prior Core Research, providing



Figure 4 Mobile VCE session at recent ITU-R WG8F meeting in Edinburgh

members with an easy route for cost-effective pre-development of emerging technology, bridging the gap between research and product development and ensuring that the companies have direct access to the relevant researchers, easing the transition from research to product. Recent research in this category has included an adaptive antenna handset demonstrator, which has been shown to eliminate fades. New elective programmes just commencing include research on ultrawideband and a software defined radio experimental programme.

## The International Dimension

Mobile VCE began life as a national initiative but today has global impact, reflecting the international composition of its industrial membership and the focus of its research agenda, with member companies making full use of today's IT tools to circulate its outputs to their research groups across national boundaries.

Formal collaboration exists with Yokosuka Research Park, YRP, in Japan, who hosted a Mobile VCE event in 2002 [11]; a second regional event in Asia is under consideration, as are researcher secondments to Japan.

A similar MoU exists with the SDR Forum, which has included Mobile VCE research in its responses to the FCC's Notices of Inquiry into SDR. The Forum supported a recent 'by-invitation' Round Table event on the impact of SDR [12], attended by spectrum regulators from Europe, Asia and North America.

Strategic research overviews from Mobile VCE's academic members are regularly contributed into the Wireless World Research Forum, whilst indeed to ITU-R WP8F is also increasing. Whilst overview publications are fed into such fora, as well as into international research conferences and refereed journals, access to the detailed technical research reports, software tools etc, remains limited to industrial member companies.

## Moving Forward...

Mobile VCE – the Virtual Centre of Excellence in Mobile & Personal Communications – is now 8 years old. Initially trading on promise and the goodwill of the industry, the company has delivered tangible benefits to its Members and has evolved from merely a national initiative to become an important player on the world stage. Effective structures and mechanisms have resulted in a high and sustained degree of research achievement, whilst the value and effectiveness of its Industrial Steering Groups has encouraged

companies to resource these with high quality staff. Together these result in a very effective industry input to the work and enable member companies to derive direct benefits themselves. International recognition of the Core Programme research outputs has been accompanied by a fresh thrust into important new areas reflecting key priorities of today's industry.

For further information on industrial membership please see the Mobile VCE website [www.mobilevce.com](http://www.mobilevce.com) or e-mail [future.wireless@mobilevce.com](mailto:future.wireless@mobilevce.com)

## References

- [1] ITU-R Working Party 8F 'IMT2000 and Systems Beyond', Edinburgh, Scotland, 8<sup>th</sup> October 2003, including the Services & Market Aspects Workshop, jointly organised by Mobile VCE with ITU-R WP8F
- [2] 'European research in mobile and personal communications perspectives for the future', J. M. Pereira, *International Seminar on Broadband Communications*, 2002
- [3] Several articles summarising the research achievements of the Core 1 Programme appeared in *IEEE Electronics & Communications Engineering Journal*, December 2000
- [4] 'Visions of 4G', K Baughan, BG Evans & WHW Tuttlebee, *IEEE Electronics & Communications Engineering Journal*, December 2000
- [5] 'Variable Spreading Factor Orthogonal Frequency and Code Division Multiplexing (VSF-OFCDM)', H. Atarashi, M. Sawahashi, *Third International Workshop on Multi-Carrier Spread Spectrum (MC-SS 2001) & Related Topics*, pp. 113-122, Oberpfafenhofen, Sept. 2001
- [6] '4G: An Idea whose Time has Come?', P Henry, *4G Forum*, London, UK, April 2003 [http://www.research.att.com/areas/wireless/Mobile\\_Interdomain\\_Roaming/Cellular\\_WLANs/4GForum.pdf](http://www.research.att.com/areas/wireless/Mobile_Interdomain_Roaming/Cellular_WLANs/4GForum.pdf)
- [7] *IEEE 802.20 Mobile Broadband Wireless Access (MBWA)*, <http://grouper.ieee.org/groups/802/20/>
- [8] *IST Mobile, Satellite & Personal Communications*, Programmes and Action Lines <http://www.cordis.lu/ist/ka4/mobile/calls/areas.htm>
- [9] *Broadcast – Mobile Convergence/Synergy Conference*, Metz, France, 22-23<sup>rd</sup> October 2003, <http://tdfint.iside.net/article/archive/1591/>
- [10] 'Interworking – not Convergence', WHW Tuttlebee *et al*, *European Broadcasting Union Technical Review*, January 2003
- [11] *Future Mobile Evolution Symposium*, Yokosuka Research Park, Japan, 15<sup>th</sup> April 2002
- [12] *International Regulatory Round Table on Software Defined Radio*, London, England, 15<sup>th</sup> September 2003, <http://www.radio.gov.uk/topics/research/topics/converge-new-emerging/sdr/sdr-index.htm>

*Walter Tuttlebee joined Mobile VCE as Executive Director in December 1999, with a background in R&D management and business development in personal communications. With his previous company Dr Tuttlebee was a key player in the initiation of major research programmes, including the European FRAMES programme that significantly shaped 3G standards. His activities in short range wireless communications and software radio – editing several books and creating a global web community ([www.dectweb.com](http://www.dectweb.com)) – involve him in chairing and speaking at industry conferences on DECT, Bluetooth, SDR and related technologies.*

*Dr Tuttlebee has BSc and PhD degrees from Southampton University and an MBA from Cranfield. He is a Senior Member of the IEEE.*

Virtual Centre of Excellence in Mobile & Personal Communications – Mobile VCE,  
Grove House, Basingstoke, RG24 8AG, UK  
Email: [walter.tuttlebee@mobilevce.com](mailto:walter.tuttlebee@mobilevce.com)  
Web: [www.mobilevce.com](http://www.mobilevce.com)



[www.isg.rhul.ac.uk](http://www.isg.rhul.ac.uk)



[www.mobilevce.com](http://www.mobilevce.com)

## *An introduction to trusted computing*

Chris Mitchell

Information Security Group

Royal Holloway

University of London

<http://www.isg.rhul.ac.uk/~cjm>

# An introduction to trusted computing

Chris Mitchell

Information Security Group

Royal Holloway

<http://www.isg.rhul.ac.uk/~cjm>

## Contents

- What is trusted computing?
- The need for trusted computing
- A multiplicity of specifications ...
- Agenda of workshop

## Computer security

- Computer security has a long history, and many secure computer systems have been produced and sold.
- Almost all of them depend on the assumption that the computer hardware will be physically secure, and managed by trusted personnel.
- Physical access to the machine will typically allow software integrity to be compromised

## Multi-user systems

- Many systems (e.g. Unix, Windows 2K/XP) designed to allow users to protect their data and resources against other users of same machine.
- All based on access control systems.
- Again typically dependent on physical security of machine.



## Computer security – external view

- If a (secure) computer digitally signs a message, then trust in messages depends on:
  - trust in computer software, and
  - trust in physical security of hardware (and in correct application of security procedures by administrators).
- Makes sense in conventional ‘computer centre’.

## PC security

- Perhaps an inherent contradiction!
- PCs are not stored in a physically secure environment.
- Even though modern versions of Windows (and Linux) have multi-user security features, users and programs often run as administrator.
- There are many ways that the operating system integrity can be damaged.

## Contents

- What is trusted computing?
- **The need for trusted computing**
- A multiplicity of specifications ...
- Agenda of workshop

## Trusting a PC

- Today, neither the user of a PC nor a communicating party can trust very much at all about a PC.
- This is despite major efforts to improve security of Windows.
- Anyone with access to the PC hardware can modify Windows (e.g. by removing hard disk and changing files).

## Trusting a PC – more bad news ...

- Even if the user looks after the physical security of their PC, there are many other threats to system integrity.
- Modern operating systems and applications are highly complex and it is almost impossible to remove all vulnerabilities.
- Users can easily accidentally run malicious software which can damage system integrity.

## Need for trust I

- User may want to trust the integrity of their PC.
- For example, the PC may be used for:
  - managing a bank account,
  - performing e-commerce transactions,
  - managing personal information,
  - ...

all of which require *user* trust in the PC.

## Need for trust II

- Third party may want to trust integrity of PC.
- This could be for a variety of reasons, e.g.:
  - 3rd party is a bank: PC being used for e-commerce,
  - 3rd party is a content provider: PC performing DRM,
  - PC performing other security functions (e.g. authentication, key management) on behalf of 3rd party,all of which require *third party* trust in the PC.

## Role of Trusted Computing

- Enables trust in integrity of PC based on combination of software and hardware.
- Third parties can measure PC integrity.
- Trusted Computing does not just apply to conventional PCs: equally relevant to PDAs, mobile phones, broadcast receivers, ...

## Contents

- What is trusted computing?
- The need for trusted computing
- A multiplicity of specifications ...
- Agenda of workshop

## TCG

- Trusted Computing in the sense of this workshop dates back to late 1990s.
- Consortium of major manufacturers started TCPA (Trusted Computing Platform Alliance).
- This has morphed into TCG, the Trusted Computing Group.

## NGSCB

- *Next Generation Secure Computing Base* (NGSCB) is Microsoft's take on Trusted Computing.
- Version of Windows that uses trusted hardware (e.g. hardware conformant to TCG specifications) to build a trusted kernel.
- Allows trusted applications to run under control of a trusted operating system, in parallel to 'regular' Windows applications.

## LaGrande

- Set of enhancements to Intel chip sets incorporating everything needed to build a Trusted Computing Platform.
- Also provides a potential platform for NGSCB-enabled PCs.

## Contents

- What is trusted computing?
- The need for trusted computing
- A multiplicity of specifications ...
- **Agenda of workshop**

## Workshop agenda

- Two main parts to workshop.
- This morning:
  - Introduction to Trusted Computing technology, including TCG, NGSCB, etc.
- This afternoon:
  - Four examples of possible applications of Trusted Computing technology.

## Follow up

- We hope that this will be a worthwhile day for all involved.
- If you have questions for any of the presenters, please either contact them directly or email me at [c.mitchell@rhul.ac.uk](mailto:c.mitchell@rhul.ac.uk)
- We hope to run similar workshops in the future – please let us know of topics you would like covered in future events.





[www.isg.rhul.ac.uk](http://www.isg.rhul.ac.uk)



[www.mobilevce.com](http://www.mobilevce.com)

## *The TCPA/TCG trusted platform*

Eimear Gallery  
Information Security Group  
Royal Holloway  
University of London  
`e.m.gallery@rhul.ac.uk`

# The TCPA/TCG Trusted Platform

Eimear Gallery  
Mobile VCE Research Group  
Information Security Group  
Royal Holloway, University of London  
[e.m.gallery@rhul.ac.uk](mailto:e.m.gallery@rhul.ac.uk)

## Contents

- History
- Trusted platforms
- TCG specifications, version 1.1b
- TCG specifications, version 1.2

## Contents

- History
- Trusted platforms
- TCG specifications, version 1.1b
- TCG specifications, version 1.2

3

## TCPA: The early years

- TCPA: An industry working group
- Focus: Enhancing trust and security in computing platforms
- Original alliance of promoter companies (HP, IBM, Intel and Microsoft): Founded 1999
- Initial draft standard unveiled: Late 1999
- Invitation then extended to other companies to join the alliance
- Specification eventually became an open industry standard
- By 2002: The TCPA had over 150 member companies

4

## Specification output

- February 2001: Generic Platform Specification, version 1
- August 2001: Generic Platform Specification, version 1.1
- September 2001: PC Specification, version 1
- 2002 onwards: Planned specifications for:
  - Servers
  - Mobile phones
  - Internet appliances.

5

## TCG: The early years

- TCG: announced April 8, 2003
- TCPA recognised TCG as successor organisation for the development of trusted computing specifications
- Adopted the specifications of the TCPA, including:
  - Main specification version 1.1 which evolved into a version 1.1b specification; and
  - PC implementation specification version 1, which evolved into a version 1.1 specification.
- Aim:
  - To extend the specifications for multiple platform types
  - To complete software interface specifications to facilitate application development and interoperability
  - To ensure backward compatibility.

6

## Current position

- May 2003: Operational technical working groups for:
  - Future TPM, trusted platform module
  - PC specific implementation specifications
  - New TSS, TCG software stack specifications as well as for
  - The development of common criteria protection profiles.
- Followed closely by formation of working groups for:
  - Server, PDA, mobile phone platform specific implementation specifications.

7

## Specification output

- TCG TPM main specification (general platform specification) version 1.2:
  - Design principles
  - Structures of the TPM
  - TPM commands.
    - (superseded TCG main specification version 1.1)
- TCG software stack specification version 1.1
- TCG software stack specification header file
- TCG PC specific implementation specification version 1.1.

8

## Contents

- History
- **Trusted platforms**
- TCG specifications, version 1.1b
- TCG specifications, version 1.2

9

## Trusted Platforms

- A trusted platform:
  - A computing platform that has a trusted component
  - Usually in the form of built-in hardware which is
  - Used to create a foundation of trust for software processes.

10

## Trusted Platform functionality (1)

- Trusted platform technologies aim to provide:
  - Confidentiality and integrity of application code and data
  - Confidentiality and integrity of application code and data during storage
  - Integrity of the operating system and underlying hardware such that the properties above can be satisfied

11

## Trusted Platform functionality (2)

- Platform authentication to external entities
- Trusted path to user ensuring confidentiality of user input
- Secure channels to devices and between applications to ensure confidentiality, integrity, and authenticity of inter-application communication
- Ensure reliability by restricting size of trusted critical components:
  - Stafford estimates 1 security related bug per 1000 LoC.

12

## Contents

- History
- Trusted platforms
- TCG specifications, version 1.1b
- TCG specifications, version 1.2

13

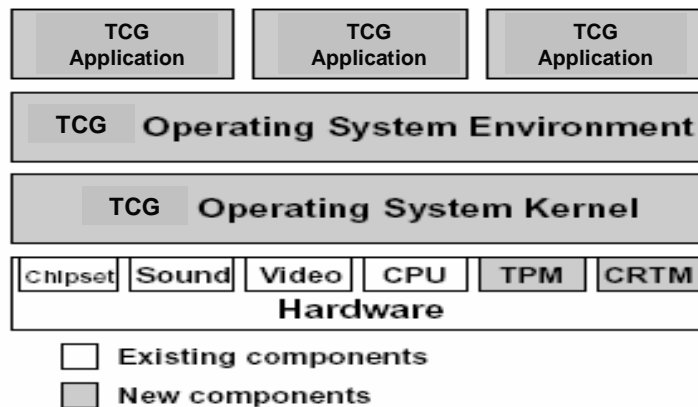
## A TCG trusted platform

- Measurement and storage of platform configuration
- Platform attestation: reporting of platform configuration
- Data protection:
  - Confidentiality
  - Access
    - Access control data
    - Platform configuration.

14



## TCG architecture



15

## TCG components

Two roots of trust:

- Core root of trust for measurement, CRTM
- Trusted platform module TPM:
  - Root of trust for storage
  - Root of trust for reporting.

Trusted Platform Support Service, TSS:

- Trusted platform measurement store
- TCPA validation data
- Measurement agents
- Trusted platform agent.

16

## Core root of trust for measurement, CRTM

- First piece of code to execute on the platform:
  - Bios; or
  - Bios boot block in an IA-32 platform.
- It is trusted to accurately measure at least 1 integrity metric that indicates the software environment of the platform:
  - Records measured integrity metrics to a TPM
  - Records details of the measuring process to a 'trust platform measurement store'.

17

## Trusted platform module, TPM (1)

- Comprised of:
  - Root of trust for storage: accepts integrity measures and records them
  - Root of trust for reporting: supplies an accurate digest of all sequences of integrity metrics presented to it.
- Uniquely bound to a single platform by either physical or cryptographic means
- All TPM functions and storage are isolated from all other components on the platform.

18

## Trusted platform module, TPM (2)

- The TPM also incorporates the following functionality:
  - Asymmetric key generation
  - Asymmetric co-processor: Signing and encryption
  - Computing engine
  - Keyed hash functions: MACs
  - Hashing functionality
  - Power detection
  - Random number generation
  - Non-volatile memory
  - Memory
  - PCRs, platform configuration registers.

19

## TCG entities

- Trusted platform module entity
  - Validation entity
  - Conformance entity
  - Platform entity
  - Privacy-CA
- } Other roots of trust
- 
- TPM owner
  - TPM user
  - Challenger
  - Protected object owner
  - Intermediary

20

## TCG mechanisms

- Storing and reporting on integrity
- Authenticated and secure boot processes
- Platform attestation
- Protected storage

21

## Storing and reporting on integrity (1)

Integrity metric registers:

- Platform configuration registers, PCRs
  - Used to store platform software integrity metrics
  - Usually a TP has several PCRs (sixteen in the case of the version 1.1 specifications) and uses those to record different aspects of the state of the trusted platform
  - Each storage register has a length equal to the sha-1 digest: 20 bytes.

22

## Storing and reporting on integrity (2)

- Each PCR holds a summary value of all the measurements presented to it:
  - Less expensive than holding all individual measurements in the TPM; and
  - This means that an unlimited number of results can be stored.
- A PCR value is defined as:
  - SHA-1 (existing PCR value | latest measurement result )
- A PCR must be a TPM shielded location, protected from interference and prying
  - The fewer sequences/PCRs there are, the more difficult it is to determine the meaning of the sequence.
  - The more there are, the more costly it is to store sequences in the TPM.

23

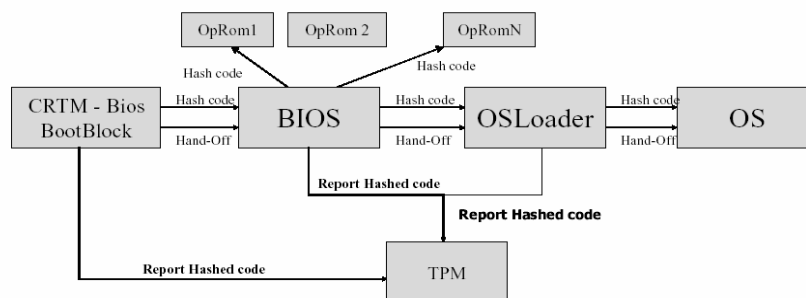
## Storing and reporting on integrity (3)

- Data integrity register, DIR
  - DIRs are storage registers that hold digest values
  - Hold values which represent the expected outcome of the measurement process/ expected PCR values
  - Version 1.1 of the TCGA specifications require that the TPM contains at least one DIR in a TCG-shielded location
  - A value is written to the DIR by the authorised TPM owner
  - Each storage register has a length equal to the sha-1 digest, 20 bytes.

24

## Authenticated boot process (1)

The Authenticated boot process



25

## Secure boot process (2)

Secure boot:

- The expected PCR values are written by the TPM owner to DIRs
- The CRTM or a measurement agent computes PCR values and then compares them to DIR values
- If the two values match, control is passed to the measured software and the boot process continues, otherwise an exception is called and the boot process halted.

26

## Secure boot process (3)

Secure boot:

- Alternatively if the TPM has access to non-volatile memory, all expected PCR values can be held in unprotected non-volatile memory and their summary or accumulative digest held in a lone DIR.
- When a PCR value has been calculated, the RTM or measurement agent checks that:
  - The accumulative digest of the expected table of PCR values matches that held in the DIR; and
  - The calculated PCR value then matches its expected value in the table.

27

## Platform attestation (1)

The Endorsement Key Pair:

- Required that:
  - Each TPM has a private endorsement key embedded in it
  - The public half of endorsement key pair is certified by the TPME/manufacture (endorsement credential)
- Used to recognise that TPM is indeed genuine.
- Never used for signing
- It is not a platform identity
- Only ever used for decryption in 2 scenarios:
  - Taking ownership of TPM
  - To derive platform attestation identities/platform identities.

28

## Platform attestation (2)

### Platform credentials:

- Endorsement credential:
  - Certifies that the public encryption key from the pair in an endorsement credential belongs to a genuine TPM
  - Constructed by: Trusted platform management entity.
- Conformance credential:
  - A document that vouches that a particular design of TPM and platform meets the TCG's specifications
  - Constructed by: Conformance laboratories.
- Platform credential:
  - A document that proves that a TPM has been correctly incorporated into a design which conforms to the specifications
  - Proves the trusted platform is genuine
  - Constructed by: Platform entity.

29

## Platform attestation (3)

### Platform Identity/attestation key pairs:

- Used by TPM to attest to platform properties to external entities
- Used to verify that TPM is indeed genuine without identifying a specific TPM
- The functionality of a privacy certification authority, P-CA is used to achieve this:
  - TPM chooses an arbitrary key pair and an 'identity'
  - The chosen P-CA verifies all TPM credentials
  - The P-CA certifies that the particular identity chosen by the TPM belongs to a genuine TPM.
    - The P-CA however becomes a point of weakness.

30



## Protected storage (1)

- Each TP contains a key hierarchy:
  - At the root is the storage root key, SRK, stored securely in the TPM.
- Data or keys can be encrypted such that it can only be decrypted by the TPM
- Asymmetric encryption is used.

31

## Protected storage (2)

- Binding (data)
  - This capability allows for external data to be encrypted under a public TPM parent key such that it can only be decrypted by the TPM
- Wrapping (keys)
  - TSS Wrap Key: This capability allows an externally generated key to be encrypted under a parent key

### Wrapping variants:

- TSS Wrap key to PCR: Similar to above but the externally generated key is wrapped to PCR values
- TPM Create wrap key: Creates a TPM key, which may or may not be locked to PCRs

32

## Protected storage

- Sealing (data)
  - This is an important aspect of protected storage
  - The seal operation can bind a secret to an individual TPM
  - External data is concatenated with the value of integrity metric sequence at the time the seal operation is performed and encrypted under the public key of a parent key pair
  - It provides the capability to store a secret such that it can only be revealed by the TPM when the platform is in an acceptable software state
  - The caller of the seal operation may choose not to wrap the secret to any PCR values.

33

## Demonstrating privilege

- Physical presence
- Cryptographic authorisation:
  - 20 bytes, for example a hashed password or 20 bytes from a smartcard submitted to a hash algorithm may be used
  - Separate authorisation data must exist for the TPM owner, TPM users as well as each protected object:
    - No concept of 'super-user'.

34

## Criticisms

- P-CA:
  - Point of weakness: Are capable of:
    - User/TPM activity tracking; or
    - Making unwanted disclosures of platform information.

35

## Contents

- History
- Trusted platforms
- TCG specifications, version 1.1b
- TCG specifications, version 1.2

36

## Scope and purpose

Examination of:

- New features appearing in version 1.2 of the main specification
- Version 1.1b TPM features which have been improved.

37

## Motivations

- In order to provide support for trusted software processes
- To address feedback received from users, developers and others concerning version 1.1 of the specification.

38

## New features

### Direct Anonymous Attestation, DAA:

- Removes the necessity to disclose the public value of the endorsement key to a P-CA
- Allows a platform to attest to platform characteristics while removing the necessity of a TTP altogether
- A zero knowledge proof based method
- DAA can be used to convince a remote entity (a verifier) that a particular TPM is indeed genuine without disclosing the public endorsement key or any unique identifier.

39

## New features

### Locality:

- Allows TPM owner to assign privileges to external processes based on their locality
- Allows the characteristics (integrity metrics) of the external software processes to be recorded in locality specific PCRs
- When a trusted process sends commands to the TPM:
  - A non-spoofable modifier is sent with it which indicates the locality of the process and thereby its trust value:
  - This can be used as a qualifier for more granular access to any TPM resources.

40

## New features

### Delegation (1):

- Allows an owner to have fine-grained control over the use of specific owner authorised TPM commands
- Previously, in version 1.1b specification, an owner that wished to authorize a software module to perform an owner-authorized TPM function would have been required to provide the software with the TPM's owner password.

41

## New features

### Delegation (2):

- With the new delegation function, the TPM owner may delegate to a software object or other entity the ability to use any individual owner-authorized TPM command or subset of TPM commands, without granting them the ability or permission to use any other TPM commands.

42

## New features

### Non volatile storage:

- The use of fixed-size non-volatile, NV, storage registers (such as the DIR) are replaced with a more general NV storage facility
- NV storage facility also has a modest amount of access-protected storage which may include:
  - Direct authorisation of read and write;
  - Authorisation of read and write based on locality; and
  - Authorisation of read and write based on platform integrity measurements.

43

## New features

### Transport protection:

- Implemented to improve the security of the communication channel between the TPM and trusted processes
- Transport session provides integrity and confidentiality protection to commands sent to the TPM:
  - Integrity is provided by the MACing; and
  - Confidentiality by the encryption of the command sent using the one time pad, generated inside the TPM.
- The logging of commands sent to the TPM within a transport session is also facilitated.

44

## New features

### Monotonic counters (1):

- Help to prevent replay attacks against the TPM
- Applications add to the encrypted data an additional field containing the current value of a monotonic counter
- This dedicated monotonic counter is incremented when the data is modified, invalidating any stored encrypted version of the data
- When the encrypted data is decrypted, the stored value of the monotonic counter is compared to the current value of the counter

45

## New features

### Monotonic counters (2):

- If values don't match, then the decrypted data is rejected
- The monotonic counter must itself be protected
- Version 1.2 TPM: Provides a limited number of protected hardware monotonic counters for use by system software
- Trusted software can provide dedicated 'virtual monotonic counters' for applications based on a single hardware counter.

46



## New features

### Tick counter:

- With some external support a tick counter can enable secure timing
- By correlating the v1.2 TPM tick counter value with an external time-stamping source it becomes possible to use the TPM to do secure time stamping using an external time source, without the expense of adding a real-time clock to the TPM.

47

## New features

### Context save and restore (1):

- Addresses the need to share the TPM between multiple users operating at different locality levels
- As a user of the TPM loads keys or other objects into the TPM, resources are consumed
- At some point a new user of the TPM may not be able to find sufficient free resources to perform their intended operation.

48

## New features

### Context-save and restore (2):

- To allow easy sharing of the TPM:
  - When a TPM user encounters a resource problem the software can context save the appropriate objects off the TPM;
  - Release and use the resources; and then
  - Restore the TPM state using Context Restore before relinquishing the TPM to another user.

49

## New features

### Clear endorsement key:

- A function that erases the Endorsement Key (EK):
  - providing the EK was created as an erasable EK.

### I/O functionality:

- Implements a low-bandwidth access-controlled physical communication channel that originates at the TPM
- This channel can be used to provide an access-protected channel between the TPM and other hardware components in the platform.

50

Thank you  
Questions?



[www.isg.rhul.ac.uk](http://www.isg.rhul.ac.uk)



[www.mobilevce.com](http://www.mobilevce.com)

## ***NGSCB***

Eimear Gallery  
Information Security Group  
Royal Holloway  
University of London  
`e.m.gallery@rhul.ac.uk`

# NGSCB

Eimear Gallery  
Mobile VCE Research Group  
Information Security Group  
Royal Holloway, University of London  
e.m.gallery@rhul.ac.uk

## Contents

- History
- NGSCB
- LaGrande Technology (LT)

## Contents

- History
- NGSCB
- LaGrande Technology (LT)

3

## Palladium: The early years

- The Palladium architecture has been under development since 1997:
  - Peter Biddle, Microsoft representative at CPTWG, copyright protection, and DVD-CCA:
    - Sceptical of technical efficiency of software based DRM
  - Patents have been filed which cover portions of it:
    - DRM OS;
    - Loading and identifying a DRM OS; (both approved December 2001)
      - (No link has been officially identified between the DRM OS and Palladium)

4

## NGSCB: The early years

- January 2003: The name Palladium was dropped for 2 fundamental reasons:
  - Officially: The palladium name had already been trademarked by another company; and also
  - Unofficially: Because of the controversy and what Microsoft considered to be misdirected criticism surrounding it.
- The work has continued under the name NGSCB, next generation secure computing base.

5

## Specification output

- Currently no technical information had been made available on NGSCB except for:
  - A sequence of white papers upon which this talk is based

6

## Current position

- It was initially announced that Palladium (now NGSCB) was due to ship with the next major version of Windows, the Longhorn operating system in 2004/5;
- As of yet still waiting.....

7

## Relationship with TCG

- Microsoft have declared that NGSCB is not an implementation of the existing specifications by the TCG
- It must be noted however that Microsoft is one of the founding members of TCG and actively participates in the organisation.

8



## Version 1.1 TCG specifications

- Version 1.1 of the TPM is:
  - Not supported by Microsoft Windows; nor will it
  - Enable functionality in an NGSCB machine;
- Version 1.1 TPMs may however be supported on a Microsoft windows machine with addition of:
  - A CSP, cryptographic service provider; provided by the hardware manufacturer, which has been written to support:
  - The Windows CAPI, cryptographic applications programming interface.

9

## Version 1.2 TCG specifications

- The TPM as defined in version 1.2 of the specifications is however:
  - Expected to work as the central secure component in NGSCB, described below.

10

## Contents

- History
- **NGSCB**
- LaGrande Technology (LT)

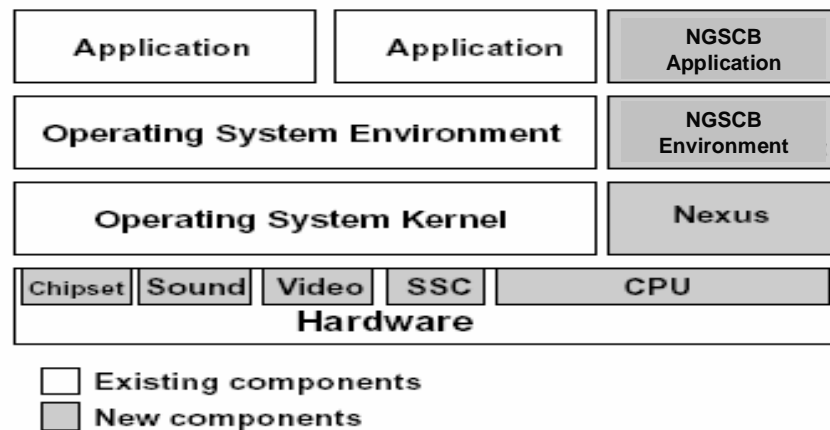
11

## The NGSCB Trusted Platform

- Palladium/NGSCB is defined as:
  - A system which combines software and hardware controls to create a trusted computing platform

12

## NGSCB architecture



13

## TCG components

- In order to develop an NGSCB, Microsoft stipulates that modifications must be made to the following components of the traditional computing platform:
  - The CPU;
  - The memory controller or chipset;
  - The keyboard;
  - The video graphics card; and
  - The graphics adaptor.
- A trusted platform module, labelled the SSC, must also be added.

14

## TCG components

Three fundamental components:

- The security support component (SSC) implemented in hardware
- The nexus or security kernel and
- The NGSCB trusted applications, also called nexus computing agents, NCAs, which may constitute:
  - An application;
  - Part of an application; or
  - A service;which are run securely within the protected operating environment.

15

## Security Support Component SSC (1)

- The SSC is a tamper-resistant cryptographic chip
- It contains a variety of cryptographic functionality, including a minimum of:
  - SHA-1 operations
  - RSA public key operations for encryption/decryption as well as signing
  - AES encryption/decryption
  - A small amount of memory and
  - A monotone counter.

16

## Security Support Component SSC (2)

- It contains a PCR, a "process control register", used to store the image of the nexus security kernel or the nexus hash/identity
- Version 1.2 of the TCG's TPM is expected to serve as the SSC in the NGSCB
- As was the case with the TCG's TPM, the SSC is permanently attached to the motherboard.

17

## Security Support Component SSC (3)

- Provides the following primitives to a nexus:
  - TPM\_Seal                                 } Sealing
  - TPM\_Unseal                                 }
  - TPM\_PKSeal                                 } Attestation
  - TPM\_GetMonotonicCounter
  - TPM\_IncrementMonotonic Counter
  - TPM\_GetEntropy
  - TPM\_GetCertificate                 } Returns the SSC certificate if  
the nexus is authorised to get it

18

## Security Support Component SSC (4)

- **Minimum keys a SSC stores: Never leaves the SSC**
  - One AES symmetric key:
    - Used to secure data such that it is only retrievable by the SSC in question
  - One RSA private key from a pair:
    - Unique to SSC
    - Could be used to identify the motherboard and computer containing the SSC
    - Only used to certify that new public key pairs were indeed generated by a specific nexus on an NGSCB enabled machine
    - The RSA private key may also be used when migrating secrets from one NGSCB system to another.

19

## Security Support Component SSC (5)

- Digital certificate for RSA public key:
  - Issued by SSC manufacturer
  - It attests to the fact that the certificate signer (the SSC manufacturer) burned the corresponding private key in to the SSC
  - This public key certificate is the only information which can identify the platform therefore:
    - This certificate is only released to nexuses named by the user as authorised recipients and
    - This certificate is expected to be used in a manner consistent with the privacy policy set.

20

## The Nexus (1)

- Small security kernel (not a complete OS kernel)
- It manages the security hardware and the protected operating environment in which the trusted applications run
- Nexus is authenticated during system start-up.

21

## The Nexus (2)

### Functionality (once it has been authenticated):

- The nexus can identify and authenticate NCAs and can store secrets which are then only accessible to an identifiable NCA running on a specific hardware platform
- Also provides basic services to NCAs such as:
  - Memory management
  - Inter-process communication (IPC) mechanism and
  - Thread management
- The nexus controls access to trusted mechanisms offered by the SSC to NCAs via a set of APIs and services.

22

## The Nexus (3)

### Minimum key set a SSC stores for each nexus: Generated when nexus is initially loaded:

- AES keys, generated randomly when nexus is started for the first time:
  - These keys never leave the TPM
  - Are linked to the particular nexus identity
  - Can be used by nexuses in seal and unseal operations
- Private RSA keys:
  - Generated by the nexus for its use
- Public key certificates:
  - The SSC private key is used to certify that nexus public key pairs were indeed generated by a specific nexus on an NGSCB enabled machine

23

## Nexus computing agents (1)

- An NCA may constitute:
  - An application
  - Part of an application or
  - A service
- The NCAs run in curtained memory, where they are isolated from:
  - Other NCAs and
  - All software with which they do not have a trusted relationship
- An NCA can make requests to the nexus for security related services and for NGSCB services such as memory management

24



## Nexus computing agents (2)

- Each NCA has a unique code identity, which equates to a combination of:
  - The cryptographic identity of the nexus upon which the NCA is running and
  - The cryptographic identity of the NCA
- IPC, inter process communication, management is provided by the nexus to NCAs such that:
  - Communication can occur among NCAs and
  - Between NCAs and untrusted programs running either locally or remotely
- NCAs can communicate directly with the nexus.

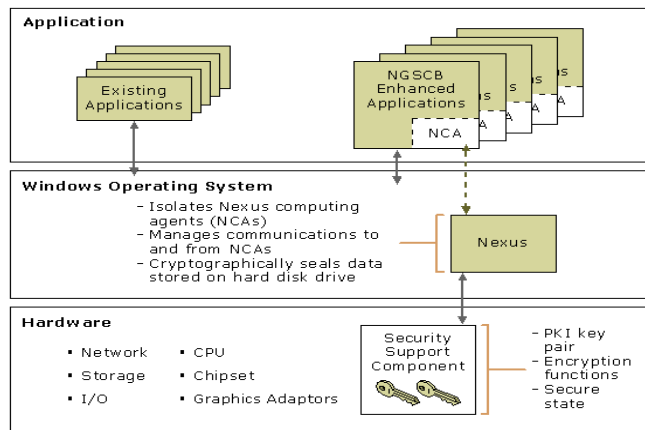
25

## Nexus computing agents (3)

- NCAs have access to the following security mechanisms:
  - Attestation
  - Sealed storage
  - Curtained execution and
  - Secure I/O

26

## Inter-working of components



27

## NGSCB security mechanisms

- Storing and reporting on integrity
- Authenticated boot processes
- Attestation
- Sealed storage
- Curtained execution and
- Secure I/O

28

## Storing and reporting on integrity (1)

- **Process Control Register, PCRs**
  - Used to store nexus integrity metrics
  - A TP is expected to have at least one PCR
  - On nexus start-up, the nexus image is:
    - Sent to the SSC; which
    - Makes a cryptographic hash of the image; and
    - Stores the hash in a register that cannot be modified.

29

## Authenticated booting of nexus (1)

Listed as an enhanced feature of NGSCB, not offered by TCG specifications.

- Not explicitly defined anywhere:
- Speculation:
  - Use of the concept of authenticated code.....described in the high level description of Intel's LaGrande technology.

30

## Platform attestation (1)

- Allows an NCA and nexus to prove its identity to other NCAs
- With this mechanism an NCA signs and attests to a piece of data, thereby confirming that the data was constructed by a cryptographically identifiable software stack.

31

## Platform attestation (2)

- The NCA requests that attestation on a constructed blob, be it arbitrary data or the public key from a newly generated key set:
  1. The nexus attests to the NCA identity and the created blob,  $y$ , by:
    - Concatenating the NCA identity and
    - The blob,  $y$ , and
    - Signing it with the nexus private key.

32

## Platform attestation (3)

2. In conjunction with this, an attestation statement which:
  - Concatenates the nexus identity; and
  - The nexus RSA public key corresponding to the one used to attest above
  - This is then signed by the SSC.
3. Finally an attestation statement/certificate which:
  - Authenticates the SSC identity key pair; which is
  - Signed by the hardware manufacturer, is sent.
4. Through the three separate 'attestation' statements an authentication chain is constructed.

33

## Platform attestation (4)

### Identity Service Providers

(mentioned by England et al, IEEE Computer (Volume 36, no. 7, July 2003, pp.55-62). but not subsequently)

- Act as trusted intermediaries between service providers and their customers
- Allow trusted platforms to acquire an arbitrary number of pseudonyms and key pairs:
  - Which are then used for attestation statements.
- Microsoft have also expressed their plans for the implementation of zero-knowledge protocols.

34

## Sealed storage (1)

- Information is securely stored and made accessible only to the NCA which stored it and other applications and services the NCA and user deem trustworthy
- Sealed data can only be read when the SCC is present
- This sealed data cannot be read if another OS is started or if hard disk is moved to another computer.

35

## Sealed storage (2)

- Each nexus generates a random key set when first loaded
- The nexus uses:
  - TPM-Seal and TPM-Unseal primitives to securely retrieve or store this key set.
- In turn:
  - NCAs use the nexus facilities and key-sets to seal and unseal their own private secrets/data.

36

## Curtained execution

| Current computer memory (RAM)  | Curtained memory   |
|--|--|
| Divided into 2 sections:<br>1. operating system in ring 0<br>2. user space in ring 3,<br>to which 2 addressing mode bits control access. | Isolates a specific portion of RAM within the address space:<br><ul style="list-style-type: none"> <li>An NGSCB addressing bit is set to address this memory portion</li> <li>This bit is added to the NGSCB CPU.</li> </ul>                             |
|  | Blocks direct memory access (DMA) devices from reading and writing;<br>Use of a DMA exclusion vector, which states for each page of memory whether or not it is a curtaigned area of memory and therefore cannot accept the transfer of DMA information. |
| Provide virtual memory protection.   | Provides physical memory protection.   |
| Relatively easy for an attacker to add malicious code to both OS and user space.   | Ensures trusted applications running behind the curtain are not modified or observed by any other program or the OS.   |
|  | In curtaigned memory, standard ring and virtual memory protections are used to isolate the nexus from NCAs and NCAs from each other.   |

37

## Secure I/O

- Refers to a secure path from, for example:
  - The mouse and keyboard to a secure I/O NCA; or
  - From an NCA to the screen.
- It is envisaged that many secure input devices, e.g. USB devices, smart card readers or biometric devices, and secure output devices will become available.

38

## Contents

- History
- NGSCB
- LaGrande Technology (LT)

39

## LaGrande Technology (LT)

- Defined as:
  - ‘A set of enhanced hardware components designed to help protect sensitive information from software-based attacks.
  - LT features include capabilities in the microprocessor, chipset, I/O subsystems, and other platform components.’ (Intel)

40



## Security mechanisms supported

- Protected execution
- Sealed storage
- Protected input
- Protected graphics
- Attestation and
- Protected launch

41

## Features

- CPU extensions
- The chipset extensions
- The keyboard and mouse extensions
- The video graphics card extensions
- TPM version 1.2

42

Thank you  
Questions?



[www.isg.rhul.ac.uk](http://www.isg.rhul.ac.uk)



[www.mobilevce.com](http://www.mobilevce.com)

## *Single sign-on using trusted platforms*

Andreas Pashalidis  
Information Security Group  
Royal Holloway  
University of London  
`a.pashalidis@rhul.ac.uk`

# Single Sign-On using Trusted Platforms

Andreas Pashalidis

Information Security Group

Royal Holloway,

University of London

a.pashalidis @ rhul.ac.uk

## Outline of Talk

- What is SSO?
- The TPM services we need.
- How to do SSO with TPMs.
- Conclusions.

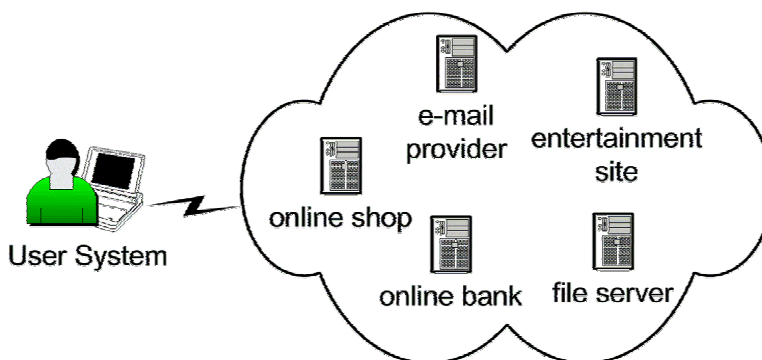
## Outline of Talk

- What is SSO?
- The TPM services we need.
- How to do SSO with TPMs.
- Conclusions.

3

## Why do we need SSO?

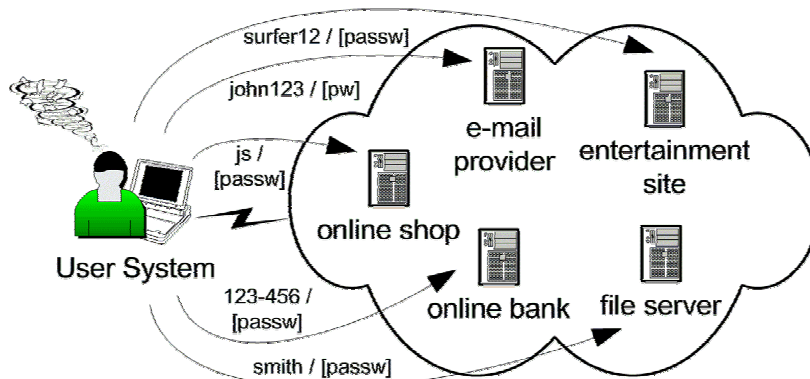
Current Situation:  
Network users interact with multiple service providers.



4

## Why do we need SSO?

Problems: Usability, security, privacy...



5


## What is SSO?

A mechanism that allows users to authenticate themselves to **multiple** service providers, using only **one** identity.

6

## SSO – How?


Introduce Trusted Third Party,  
Authentication Service Provider (ASP).



7

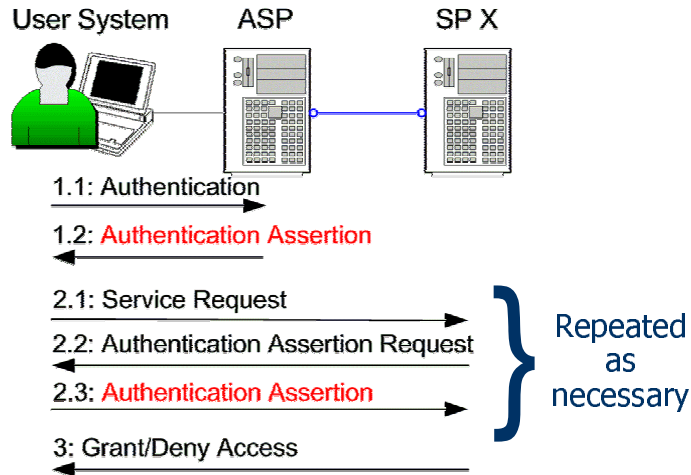
## SSO – How?

Establish trust relationships, common security  
infrastructure (e.g. PKI), sign agreements...



8

## SSO – How?



9

## Authentication Assertion

Cryptographically protected token that describes the user's authentication act.

- User's **Identifier**
- Timestamp
- Description of authentication method  
(e.g. username/password, smartcard, biometrics, etc)
- Description of authentication context  
(e.g. initial user registration, protection measures, policies...)

10



## SSO – some examples

- Microsoft Passport
  - ASP = [www.passport.com](http://www.passport.com)
  - Assertion = (symmetrically) encrypted cookie
  - Identifier = global
- Liberty Alliance
  - ASP = “Identity Provider”
  - Assertion = signed XML document
  - Identifier = per-SP (“unlinkable”)
- Kerberos
  - ASP = Kerberos server
  - Assertion = ticket (+ proof of knowledge of session key)
  - Identifier = global

11

## SSO - advantages

- User registration, identification (if necessary) and authentication is outsourced to the ASP.
- ASP can deploy multiple user authentication mechanisms for SPs to choose from, SPs can flexibly change their preferred method.
- Changes of the user authentication mechanism only affect the ASP.
- User mobility is preserved.

12

## SSO - problems

- ASP has to be trusted universally;
  - (security) it can impersonate users at SPs at will.
  - (privacy) it knows all identifiers, who logs in where and when.
- ASP is a single point of failure – it must be online and available at all times.

13

## SSO - problems

- ASP has to be trusted universally;
  - (security) it can impersonate users at SPs at will.
  - (privacy) it knows all identifiers, who logs in where and when.
- ASP is a single point of failure – it must be online and available at all times.
- We can address these problems using Trusted Computing

14

## Outline of Talk

- What is SSO?
- The TPM services we need.
- How to do SSO with TPMs.
- Conclusions.

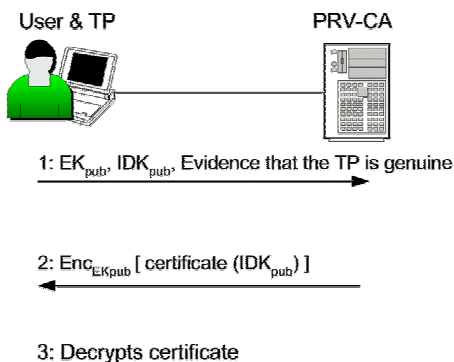
15

## The TPM services we need

- TPM Identities
  - Instead of using its **unique** EK, the TPM generates any number of RSA “Identity Keys” (IDKs). These are used for signing.
- Integrity Metrics
  - The TPM forms a hash chain at boot time which measures all the software that runs on the platform.
- Integrity Challenge
  - A third party can check the software state of a TP.

16

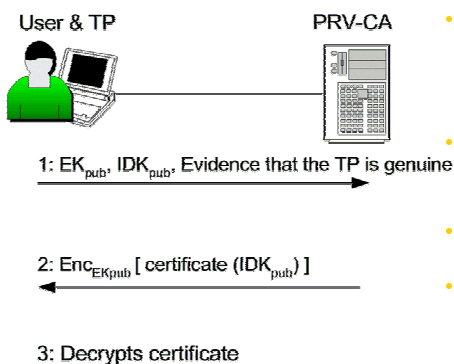
## TP – TPM Identities (v.1.1)



(simplified)

17

## TP – TPM Identities (v.1.1)



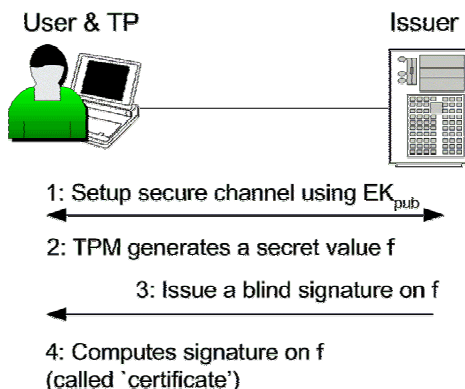
(simplified)

- The PRV-CA certifies that the IDK belongs to a genuine TPM.
- The certificate can be verified by anyone.
- It is unique (serial number).
- But it does not uniquely identify the user's TP or the user.
- It can be used as a pseudonym.

18

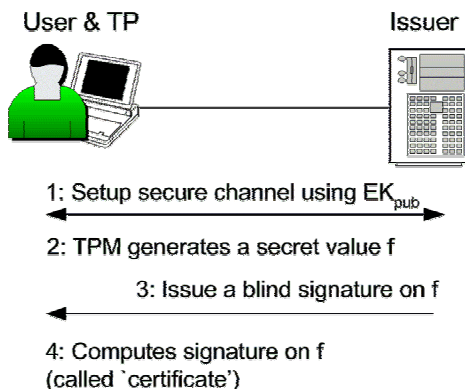
## TP – TPM Identities (v.1.2) & DAA

TPM obtains “attestation” only once, not for each IDK.



19

## TP – TPM Identities (v.1.2) & DAA



- The TPM proves that it is genuine by a **zero-knowledge proof of knowledge** of the Issuer's signature on (secret)  $f$ .
- In this way, the platform is not uniquely identified.
- The TPM can also sign messages using the value  $f$  – e.g. it can directly sign the public part of IDKs.

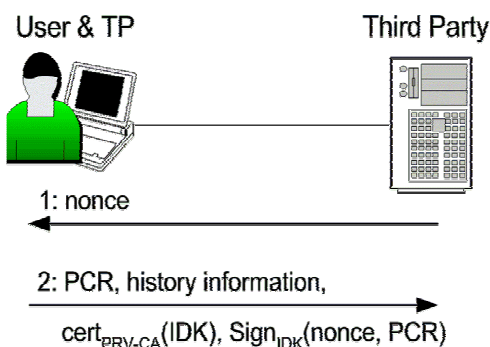
20

## TP – Integrity Metrics

- The TPM keeps SHA-1 values of its current software state in **Platform Configuration Registers (PCRs)**.
- This hash is the result of a hash chain that includes SHA-1 hashes of all critical software that has been executed on the TP (e.g. BIOS, OS, Applications).
- A structure called the **history information** contains the list of this software.
- It also contains links to **Validation Certificates**, certificates issued signed by the software manufacturer that bind the software to its hash.

21

## TP – Integrity Challenge (v.1.1)

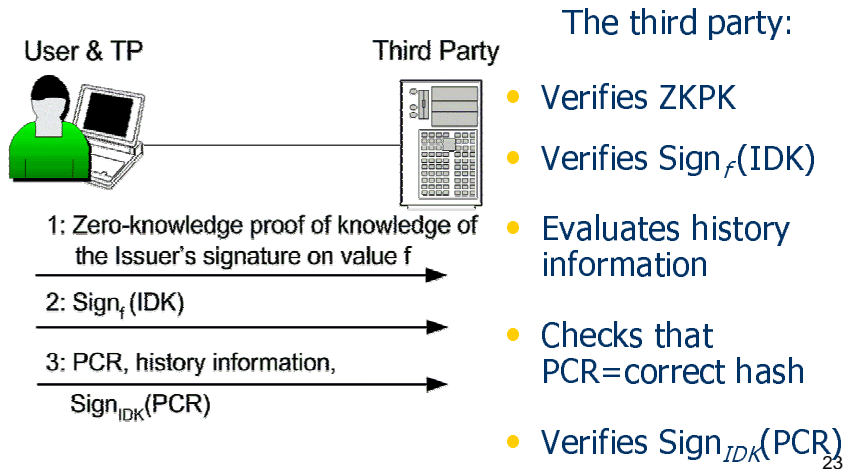


The third party:

- Verifies cert<sub>PRV-CA</sub>
- Verifies Sign<sub>IDK</sub>
- Evaluates history information
- Checks that PCR=correct hash

22

## TP – Integrity Challenge (v.1.2)




## Outline of Talk

- What is SSO?
- The TPM services we need.
- How to do SSO with TPMs.
- Conclusions.

## SSO using TPs – 1/3

A **local** User Authentication System (AS) authenticates the user.



The AS has a Validation Certificate from its manufacturer. It is measured in the PCR (i.e. its hash is included in the hash chain).

25

## SSO using TPs – 2/3

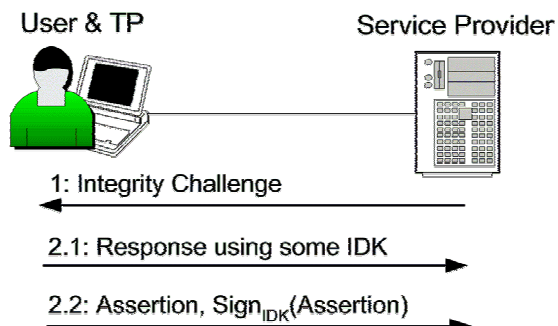
- Generate a **designated set of IDKs** for each user.
- Certify the IDKs at one or more PRV-CAs (only TCG v.1.1).
- User is known to SPs under **pseudonyms**.
  - TCGv.1.1 The **certificates**  $\text{certificate}_{\text{PRV-CA}}(\text{IDK})$  will be used as **pseudonyms** at SPs.
  - TCGv.1.2 The DAA scheme includes specification of pseudonyms (function of a user-selected *basename* and *f*)
- Using the same pseudonym at more than one SP enables linking of accounts (-> roles).

26



## SSO using TPs – 3/3

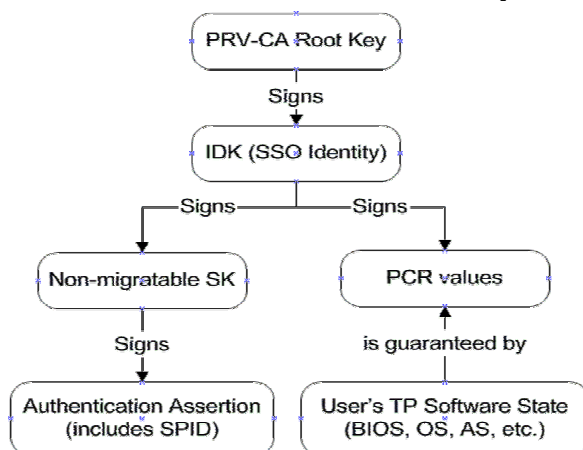
### Integrity Challenge/Response + Assertion



The assertion does not contain any identifier;  
The user's **pseudonym** is implicit in step 2.1

27

## Data Structure Relations (v.1.1)



28

## Security for Service Providers

- Service Provider decides which software configurations are acceptable.
- SSO enabled for all SPs that have chosen to trust the user's particular TPM / TP / Software AS configuration.
- If it becomes known that a particular software configuration is vulnerable/compromised, simply reject it until patched (no explicit revocation needed).

29

## Privacy (for TCG v.1.1)

- The response does not contain personally identifying information.
- Different IDKs are unlinkable, except for the PRV-CAs who certified them.
- So, if SPs collude with PRV-CAs, they can correlate users' pseudonyms (compromise privacy).

30

## Privacy (for TCG v.1.2)

- The response does not contain personally identifying information.
- There are no PRV-CAs – the user selects which of his pseudonyms are linkable as part of the TCG v.1.2 DAA scheme.
- However, care must be taken not to use the same IDK for different pseudonyms – otherwise unlinkability is compromised.


31

## Privacy (for both versions)

- The AS has access to sensitive information.  
  
(All pseudonyms, SP-associations, login times, etc...)
- AS should not disclose information to unintended parties. How can we guarantee this?

32

## Review




User TP  
with AS

Before

33

## Review




User TP  
with AS

We can get rid of the external TTP.

34

## Review



Trust is distributed among the user's local software (OS, AS, etc), TPM/TP manufacturers, testing labs and (in v.1.1) the PRV-CAs.

35

## Outline of Talk

- What is SSO?
- The TPM services we need.
- How to do SSO with TPMs.
- Conclusions.

36

## Conclusions

SSO without a continuously online TTP is possible with TCG-conformant platforms, **but:**

- The system is inherently complex - SPs must verify the user's software configuration (i.e. BIOS, OS, AS...).
- In open environments these configurations may vary considerably.
- User mobility is essentially lost.

37

## Extensions

- Possible use of NGSCB in order to reduce the complexity of verifying software configurations.
- Permission-based attribute sharing:
  - The AS could permanently store frequently used information (e.g. name, address, phone #).
  - Release only according to privacy policy (e.g. P3P).
- More on this topic later!

38

---

# Thank You!

# Questions?

e-mail: [a.pashalidis@rhul.ac.uk](mailto:a.pashalidis@rhul.ac.uk)

web: [www.xrtc.com](http://www.xrtc.com)

39



[www.isg.rhul.ac.uk](http://www.isg.rhul.ac.uk)



[www.mobilevce.com](http://www.mobilevce.com)

## *Secure content management using trusted computing*

Allan Tomlinson

Information Security Group

Royal Holloway

University of London

`allan.tomlinson@rhul.ac.uk`




On the left side of the slide, there is a vertical graphic. It features a mobile device, possibly a PDA or early smartphone, with a screen displaying a news image of two people. The device is set against a background that includes a globe of the Earth and binary code (0s and 1s). The word "news" is visible in a small font near the bottom of the device, and "audio" is partially visible below it.



*March 2004*


**Secure content  
management using trusted  
computing**

*Allan Tomlinson,  
Eimear Gallery,  
Mobile VCE Research Group,  
Information Security Group,  
Royal Holloway, University of London,*



## Contents

- Background
  - Motivation for this research
- Protection mechanisms for broadcast content
  - Limitations of current protection mechanisms
- Application of Trusted Platform technology
  - Provision of mechanisms for future content protection




www.mobilevce.com

© 2004 Mobile VCE

3


Open Workshop on Trusted Computing, 30<sup>th</sup> March



## Background

### Mobile VCE work areas

- Personal Distributed Environment
  - Multiple devices
  - Multiple services
  - Multiple networks
  - Heterogeneous and dynamic network
  - User centric
- Inter-working of Networks
  - Delivery of services to mobile users
    - Point-to-point, Broadcast, Multicast
  - Inter-working
    - Use functionalities in one network to assist other network in delivering services



www.mobilevce.com


© 2004 Mobile VCE

4

Open Workshop on Trusted Computing, 30<sup>th</sup> March

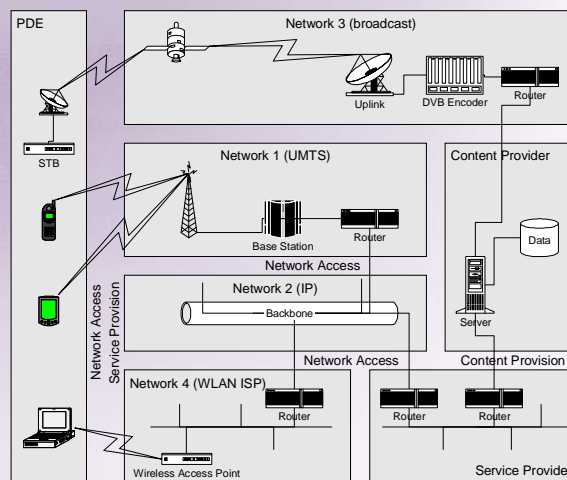
## Background

### Personal Distributed Environment



## Background

### Inter-working of Networks



## Motivation

### Technological advances

- Potential exists to deliver complex content to *mobile consumers*

### Trust

- Required between collaborating network operators

### Piracy

- Content providers are increasingly aware of, and concerned about, copyright management

### Business development

- Businesses that can trust each other and manage and protect content will be in a better position to exploit the above potential

## Contents

- Background
  - Motivation for this research
    - New business models require trusted hosts
- Protection mechanisms for broadcast content
  - Limitations of current protection mechanisms
- Application of trusted platform technology
  - Provision of mechanisms for future content protection

## Protection of Broadcast Content

### Broadcast content is currently protected by

- Conditional Access (CA) systems
- Scramble video
- Manage keys and viewing rights
  - Using proprietary security mechanisms
- DVB standards
  - Provide an interface to proprietary systems



www.mobilevce.com

© 2004 Mobile VCE

9

Open Workshop on Trusted Computing, 30<sup>th</sup> March

## Protection of Broadcast Content

### DVB Standards

- Common Scrambling Algorithm ETSI ETR 289
  - Used to scramble and descramble services (video)
  - Details available to all manufacturers
- Simulcrypt ETSI TS 103 197
  - Interface to proprietary systems at transmitter
  - Key encryption remains proprietary
  - Multiple CA systems in parallel at transmitter
  - Common key to scramble services
- Common Interface CENELEC 50221
  - Common Interface Modules – PC Cards
  - Changes proprietary CA at receiver



www.mobilevce.com

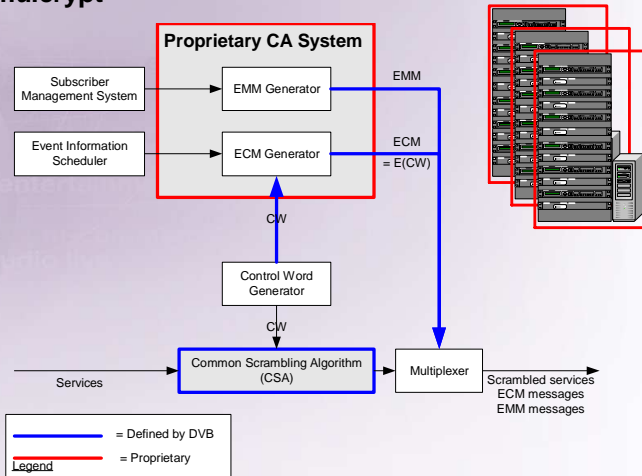
© 2004 Mobile VCE

10

Open Workshop on Trusted Computing, 30<sup>th</sup> March

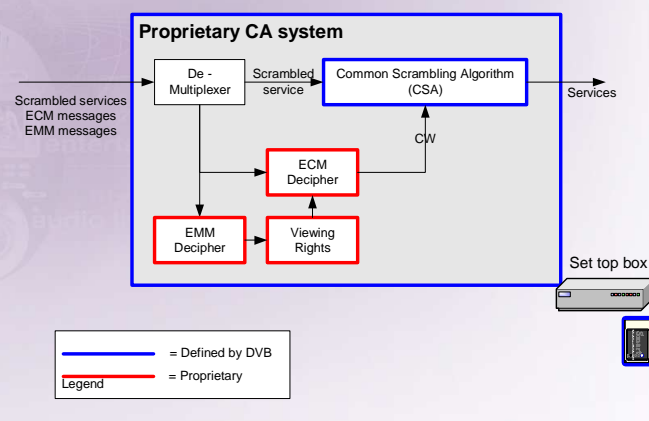
## Protection of Broadcast Content

### Simulcrypt



## Protection of Broadcast Content

### Common Interface



## Protection of Broadcast Content

### DVB Standards

- Provide a flexible interface to proprietary systems
- There are many proprietary systems

## Protection of Broadcast Content

### DVB Compliant Conditional Access (CA) systems

| ■ CA System   | Vendor       |
|---------------|--------------|
| ■ Viaccess    | Viaccess SA  |
| ■ NagraVision | Kudelski     |
| ■ Videoguard  | NDS          |
| ■ Mediguard   | Canal+       |
| ■ Mcrypt      | Irdeto       |
| ■ PiSys       | Irdeto       |
| ■ CryptoWorks | Philips      |
| ■ BetaCrypt   | BetaResearch |
| ■ Conax       | Telenor      |

## Limits of Current Protection Mechanisms

- New business model
  - Delivery of broadcast services to *mobile* receivers
    - Services available from many broadcasters
- Current protection mechanisms
  - Designed for relatively *static* receivers
    - Services available from a small number of broadcasters
- Common Interface
  - Consumers require multiple PC-Card modules
    - Cost, inconvenience, suitability for mobile devices
- Simulcrypt
  - Broadcasters install and maintain multiple CA systems
    - Cost, maintenance
- Current mechanisms not designed for mobile receivers

## Potential Solution

- Download proprietary applications to mobile devices on demand

### Problem

- These applications, and providers, are security sensitive
- Trust in the mobile host
  - Piracy: protection of proprietary algorithms, keys
- Host needs to demonstrate that it can be trusted
  - Application needs protection - not the host



## Contents

---

- Background
  - Motivation for this research
    - New business models require trusted hosts
- Protection mechanisms for broadcast content
  - Limitations of current protection mechanisms
    - Not designed for mobile receivers
- Application of trusted platform technology
  - Provision of mechanisms for future content protection



[www.mobilevce.com](http://www.mobilevce.com)

© 2004 Mobile VCE

17

Open Workshop on Trusted Computing, 30<sup>th</sup> March

## Requirements

---

### Demonstration of trustworthiness

- Integrity *challenge* mechanism
- Integrity *verification* mechanism

### Application protection

- Secure *delivery* mechanism
- Secure *execution* environment





[www.mobilevce.com](http://www.mobilevce.com)

© 2004 Mobile VCE

18

Open Workshop on Trusted Computing, 30<sup>th</sup> March

www.mobilevce.com

© 2004 Mobile VCE

## Application of TCG Trusted Platform technology

### Demonstration of trustworthiness



- Integrity metrics
  - Authenticated boot – CRTM
  - Configuration measurements – PCR
  - Attestation – TPM
    - current platform configuration

### Application protection

- Secure *delivery* mechanism
  - Key generation and exchange
- Secure *execution* environment
  - Sealed storage

19

Open Workshop on Trusted Computing, 30<sup>th</sup> March

www.mobilevce.com

© 2004 Mobile VCE

## Other security and trusted platform technology

### Demonstration of trustworthiness

- Integrity *verification* mechanism
  - Certificates and Certification Authorities

### Application protection

- Secure *delivery* mechanism
  - Encryption, Message Authentication Codes
- Secure *execution* environment
  - Physical separation of trusted and untrusted processes
    - Curtained memory – NGSCB, LaGrande
    - Compartmentalised OS - NGSCB Nexus

20

Open Workshop on Trusted Computing, 30<sup>th</sup> March

## General Approach to trusted download

### Demonstration of trustworthiness

- Authenticated boot
- Attestation of platform configuration
- Response to integrity challenge
- It is the challenger's responsibility to verify the response and determine whether to trust the platform or not
- Host must not change configuration

### Application protection

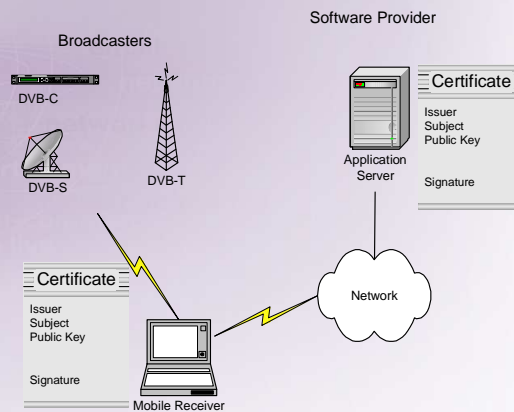
- Key exchange
- Keys in sealed storage to ensure consistent configuration
- Message Authentication Codes and Encryption
- Isolation of applications

## Protocol requirements

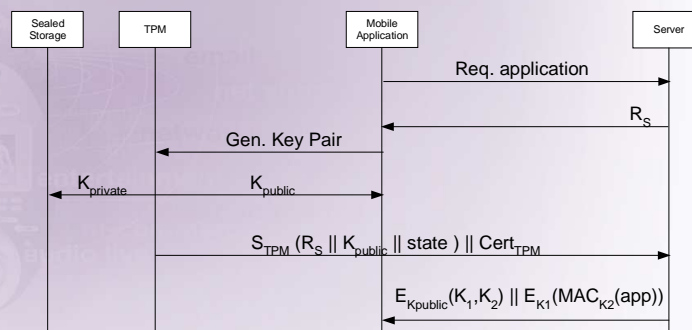
### The protocol must protect against

- Replay
  - A malicious host could replay attestation information from before the system was compromised
- Tampering
  - A malicious host could tamper with the integrity metrics before transmission to the challenger
- Masquerading
  - A malicious host could replace the original integrity metrics with data from another system
- Revealing the application
  - A malicious host could reveal the application and keys

## Model



## Protocol 1

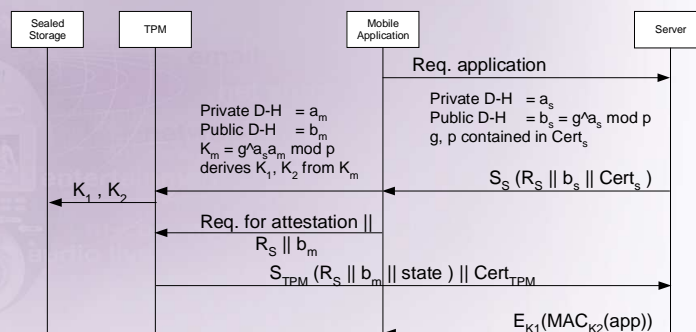


## Protocol provisions

### The protocol protects against

- Replay
  - The nonce,  $R_s$ , protects against replay
- Tampering
  - The TPM signature protects the integrity metrics
- Masquerading
  - The Certificate of the TPM protects against masquerading
- Revealing the application
  - $K_1, K_2$ , protect the application during transmission
  - Sealed storage and isolation protect during execution

## Protocol 2



## Summary

### Using Trusted Platform technology

- Host is able to demonstrate
  - It is running a secure execution environment
- Application provider
  - Has confidence that software and data will not be tampered
- User
  - Has access to a wider range of applications



[www.mobilevce.com](http://www.mobilevce.com)

© 2004 Mobile VCE

27

Open Workshop on Trusted Computing, 30<sup>th</sup> March

Thank you !



MOBILE  
VCE

For further information please contact:

Dr. Allan Tomlinson, Eimear Gallery

E-mail: {allan.tomlinson e.m.gallery}@rhul.ac.uk

Tel: +44 1784 414346

WWW: [www.mobilevce.com](http://www.mobilevce.com)



[www.isg.rhul.ac.uk](http://www.isg.rhul.ac.uk)



[www.mobilevce.com](http://www.mobilevce.com)

## *Protecting user privacy using trusted computing*

Anand Gajparia  
Information Security Group  
Royal Holloway  
University of London  
`a.gajparia@rhul.ac.uk`

# Protecting User Privacy using Trusted Computing

Anand Gajparia

Information Security Group

Royal Holloway

University of London

a.gajparia@rhul.ac.uk

Anand Gajparia is being supported by sponsorship from Toshiba Telecommunications Research Laboratory, UK **TOSHIBA**

## Contents

- Introduction
- Privacy
- Constraints
- Trusted Computing
- TCG enabling Privacy
- NGSCB enabling Privacy
- Conclusion



- Introduction
- Privacy
- Constraints
- Trusted Computing
- TCG enabling Privacy
- NGSCB enabling Privacy
- Conclusion

3

## Introduction

- Privacy of personal information is a major issue
- Particular concern is how users can control private data after it has entered the electronic world
- We show how Trusted Platforms (TPs) conforming to the Trusted Computing Groups (TCGs) specifications can be used to help enforce user control

4

- Introduction
- Privacy
- Constraints
- Trusted Computing
- TCG enabling Privacy
- NGSCB enabling Privacy
- Conclusion

5

## Privacy

- Types of personal information
  - Location Information (LI)
  - Medical records
  - Bank details
  - Phone number
  - Address
- Main focus of this talk is on LI, although similar mechanisms could be used to protect other personal information

6

## Privacy

- Users want personal information to remain private after it has been distributed
  - For example in application forms
- Users also want to retain some control of their personal information even after it has been distributed

7

## Privacy

- Range of mechanisms have been proposed which allow users to remain anonymous when using a service
  - This may be impractical for some applications
    - Vendors may require information for billing purposes
    - Contact details may be necessary to provide services at a later date
- We propose a mechanism which allows a user to retain a certain degree of control of personal information after it has been distributed

8

## Privacy model

- Private data
  - This is data containing personal information regarding a subject
- Privacy subject
  - the entity regarding which private data is being gathered, managed and used. This entity is most commonly a human user
- Service provider
  - This entity is willing to provide some service requiring personal information from a privacy subject

9

- Introduction
- Privacy
- Constraints
- Trusted Computing
- TCG enabling Privacy
- NGSCB enabling Privacy
- Conclusion

10

## Constraints

- Simple statements, bound to personal data, which may be used to help control its use, storage and distribution

11

## Constraint types

- Use
  - Purpose for which private data may be used
    - Example: Medical use only; Location based services only.
- Validity
  - Length of time private data may be stored
- Redistribution
  - Constraints for further distribution

12

## Constraint management

- Envisage that there will exist trusted software which will manage personal data in a manner trusted by the privacy subject
- Aim is to discuss how a user can establish whether or not this software is executing on a target platform, and further ensure that this software is executing when the private data is used in the future

13

- Introduction
- Privacy
- Constraints
- Trusted Computing
- TCG enabling Privacy
- NGSCB enabling Privacy
- Conclusion

14

## TCG and Trusted Computing

- Discussed in Department of Defense, Trusted Computer System Evaluation Criteria (Orange Book)
  - discusses a trusted computing base as a part of a computer protected by secure perimeter containing the parts of the system responsible for security protection of the system
- Trusted Computing Group specification could be viewed as an implementation of such a specification

15

## TCG and Trusted Computing

- We will show one way in which the use of mechanisms provided by the Trusted Computing Group specification can enhance the privacy of private data

16

- Introduction
- Privacy
- Constraints
- Trusted Computing
- TCG enabling Privacy
- NGSCB enabling Privacy
- Conclusion

17

## TCG enabling privacy

- Look at how described mechanisms may enable privacy
  - Consider a privacy subject who wishes to decide whether or not to divulge private data to a service provider with a Trusted Platform
- Using the mechanisms described by the TCG, a user can check to see if trusted management software was running on the platform when the platform booted and also specify that this personal information is only used when the platform is in a trusted state

18



## TCG enabling privacy

- Privacy subject first establishes if trusted software for management of their private data is on a service provider's Trusted Platform
- If trusted management software is found on the platform, the privacy subject then decides on limitations for future use of this private data

19

## TCG and Trusted Computing

- We use mechanisms found within the TCG TP which seal, measure, store and report integrity metrics in a trusted manner to ensure the privacy of data
- These mechanisms may be used to determine processes running on a target machine when it boots, and further ensure that data is only accessed in a trusted environment where only specified processes are running

20

## TPM Identity (version 1.1b)

- A Trusted Platform Module (TPM) identity is used to attest to aspects of the TP. This is provided by a Privacy CA
- Three certificates attesting to various aspects of the platform are sent to a Privacy CA
- If the certificates are valid for the TPM, the Privacy CA provides the TPM with a TPM identity

21

## TPM Identity (version 1.1b)

- TPM Endorsement Credential
  - attests that a Trusted Platform Module (TPM) conforms to the TCG specification
- Platform Credential
  - attests that the platform as a whole is a genuine TCG platform
- Conformance Credential
  - attests that the design and incorporation of the platform conforms to the TCG specification.

22

## TPM Identity (version 1.2)

- Uses Direct Anonymous Attestation (DAA)
  - Removes need for Privacy CA

23

## TCG measuring, reporting and storing

- The Core Root of Trust for Measurement (CRTM)
  - Responsible for integrity measurement of the first component to execute on a platform
- Root of trust for measurement (RTM)
  - Responsible for the integrity measurement of following components

24

## Platform Configuration Registers (PCRs)

- When the platform starts up, the CRTM takes a measurement used to ensure the integrity of the first component to be executed on the TP
  - This is reported to the Platform Configuration Register (PCR)
- The measured component is then responsible for measuring the integrity of the next component to be executed and is called the RTM
  - This is also reported to a Platform Configuration Register (PCR)

25

## Challenging the Trusted Platform

- An entity which wishes to challenge the state of a Trusted Platform will receive the values found in the PCR together with some validation data
  - Validation data is data signed by an entity which vouches to an aspect of a platform and shows the values that should result when integrity measurements are made in the platform
- Using this validation data, a challenger recalculates PCR values and compares these to the ones sent by the challenger

26

## Sealing Data

- The TCG specification also provides a mechanism which may be used to dictate the state a platform must be in for data encrypted by the TPM to be decrypted
  - Sealing of data may also be without the additional requirement of dictating the state of the platform to decrypt data

27

## Sealing Data

- Sealing mechanism relies on three objects
  - the digestAtCreation
    - a hash of the list of PCR numbers and their corresponding PCR values when the sealed data item was created
  - the digestAtRelease
    - a hash of a list of PCR numbers and corresponding PCR values to which data may be released
  - a list of PCRs which must be considered when releasing data.

28

## Sealing Data

- When a request is made to unseal this data, the TPM decrypts the sealed item. The information held within this sealed item is only released if, when using the listed PCR values, the recalculation of digestAtRelease corresponds to the one found in the sealed item

29

## TCG enabling privacy

- Upon first interaction with the service provider's TP, the privacy subject will typically request proof of the state of the TP
  - The privacy subject will typically send the TP a challenge
- The TP then returns information in the form of a signed version of the challenge and PCR values and also validation data
- The signing key used here is from the TPM identity
  - The use of this key assures the privacy subject that the TP they are interacting with is a valid TP
  - The inclusion of the challenge prevents replay attacks

30

## TCG enabling privacy

- Using the validation data, the privacy subject may then recalculate the PCR values found within the scope of the signature, which will then assure them of the state of the service provider's platform
- Also, the user may also verify integrity of software found on the target TP
  - This may allow a user to ensure that target platforms use software which will manage their data in an appropriate manner

31

## TCG enabling privacy

- If the target platform is in a satisfactory state the user can then specify the state of the platform for future use of their private data
  - TPM Seals data to a platform state
  - The service provider's platform must be in the specified state for further use of the private data

32

- Introduction
- Privacy
- Constraints
- Trusted Computing
- TCG enabling Privacy
- NGSCB enabling Privacy
- Conclusion

33

## Next Generation Secure Computing Base (NGSCB)

- Trusted Computing Platform being developed by Microsoft
- Uses two kernels on the same platform
  - Standard operating system kernel
  - Security kernel
- Kernels and processes are partitioned by a machine monitor
  - Security essential processes run on the secure partition

34



## NGSCB

- NGSCB provides
  - Strong process isolation
    - Separates secure processes from un-trusted processes
  - Sealed storage
    - May be used to indicate platform states for unsealing data
  - Secure paths
    - Ensures existence of a secure path between devices and a platform
  - Attestation

35

## NGSCB

- The mechanisms provided in the NGSCB architecture may also be used to provide user privacy of personal information in a similar manner to that described using the TCG specification

36

- Introduction
- Privacy
- Constraints
- Trusted Computing
- TCG enabling Privacy
- NGSCB enabling Privacy
- Conclusion

37

## Conclusion

- We have shown how Trusted Computing Platforms may be used to ensure that a user's private data is managed according to the wishes of the user
- Additionally, we have extended this and shown how future use of this data can be dependent on the platform being in a secure state

38

---

**Thank you!**



[www.isg.rhul.ac.uk](http://www.isg.rhul.ac.uk)



[www.mobilevce.com](http://www.mobilevce.com)

## ***Distributing PKI functionality using trusted computing***

Alex Dent

Information Security Group

Royal Holloway

University of London

`a.dent@rhul.ac.uk`

# Distributing PKI functionality using trusted computing

Alex Dent

Information Security Group

Royal Holloway

a.dent@rhul.ac.uk

## *The short version...*

- Certification authorities (CAs) make nice central targets for attackers and, if also used as a directory service, may be a bottleneck.
- Distributing some of the work that a CA does (such as certificate generation or distribution) onto a trusted computing platform running on a user's machine may ease some of those problems.

## *The End*

*a.dent@rhul.ac.uk*

*<http://www.isg.rhul.ac.uk/~alex>*

3

## ***The longer version...***

- A certification authority (CA) is used:
  - to issue certificates about a user's public key
  - to distribute those certificates to other users
  - to supply revocation information
- Traditionally, there is one entity which provides the CA service (a central CA)...
- ... who signs documents that bind a user's identity to their public key.

4

## ***Distributed CAs***

- We envisage distributing part of the CA functionality to a trusted computing platform.
- Instead of issuing a certificate for a public key, a central CA could instead issue an applet to run on a user's trusted computing platform that would issue certificates for that user.
- The user could produce certificates for public keys without needing to contact the central CA

5

## ***The trusted computing platform***

- The trusted computing platform can be authenticated by any third party (either implicitly or explicitly).
- Data on the TCP cannot be accessed by applications on the host machine except via the proper interface, and vice versa.
- The TCP has a method of proving to any third party that it has executed a programme successfully (either implicitly or explicitly).

6

## ***Distributed CAs***

- The important questions are:
  - How are we going to initialise the CA-applet?
  - How is a user going to register a public key?
  - How is a user going to retrieve a public key?
  - How could we revoke a certificate?
  - How could we renew a certificate?
  - What if the CA-applet was compromised?

7

## ***Initialising the applet***

1. The user registers with a central CA in the normal manner, e.g. via a registration authority (RA).
2. When the user wants to be issued with a CA-applet, the user first authenticates himself to the central CA, and vice-versa.
3. Next, the user's trusted computing platform authenticates itself to the CA.

8



## ***Initialising the applet***

4. The CA generates a new signature key pair for the CA-applet and generates a certificate for the new public key using a master key pair.
5. The CA-applet is securely downloaded onto the trusted computing platform and made available to the user.

9

## ***Where are we now?***

- So we have securely downloaded a CA-applet to a user's machine in such a way that:
  - The applet knows that it is running on a trusted computing platform
  - The applet knows the identity of the user
  - The applet has its own key pair, for which the public key has been certified by the central CA

10

## ***Registering a public key***

1. The user authenticates himself to the trusted computing platform.
2. The user submits the public key to the trusted computing platform *along with proof that he is in possession of the private key (POP)*.
3. The applet returns a certificate, signed with the applet's key pair, that certifies that the public key belongs to the user.

11

## ***Retrieving a public key***

1. User obtains the certificate from a directory service (more on this later) and the certificate for the applet's verification public key.
2. The user checks that the certificate is valid, i.e. signed by the applet, using the applet's public key.
3. The user checks that the applet's public key is valid by checking the certificate that the central CA produced.

12

## ***Distributing certificates***

- A CA-applet could help certificate distribution.
- Instead of storing a certificate in a centralised directory service, it could be stored by the CA-applet.
- Users contact the appropriate applet to get certificate.
- Optional centralised backup of certificates.

13

## ***Where are we now?***

- We have a CA-applet installed on a user's trusted computing platform...
- ...that will issue certificates to that user alone and always in that user's name...
- ...and in such a way that the user's public key can be checked providing we have an appropriate certificate chain back to the central CA.

14

## ***Revoking a certificate***

- Two approaches to certificate revocation:
  - CRLs and variants
  - Online status checking
- Central CA could revoke CA-applet's key pair for user revocation, whilst, the CA-applet could handle individual certificate revocation.
- Optional backup of certificates at a central directory service could create synchronicity problems.

15

## ***The renewal problem***

- What if the central CA or a CA-applet's key pair expires?
- Simple solution: CA-applet simply deletes itself and a new applet is downloaded...
- ...this solution means new certificates have to be generated for all keys at the user's request...
- ...which could be computationally expensive.

16

## ***The renewal problem***

- More complex solution: handover between CA-applets includes re-signing of old certificates...
- ...and this may be slightly computationally cheaper as there would be no need for the POP checks.
- Final solution: CA-applet could generate a new key pair for itself...
- ...means one applet will be running indefinitely.

17

## ***Compromising a CA-applet***

- Running a CA-applet on a user's machine may make it easier to compromise the applet and retrieve the applet's key pair.
- The probability that this happens depends upon security of the underlying trusted computing platform.
- Best that can be done is to minimise the effects of a compromise.

18

## ***Compromising a CA-applet***

- Example counter-measures include:
  - Only certifying a CA-applet for a short period of time
  - Create a “certificate history” to be stored at the central CA, which records when certificates were created and revokes all certificates created by a compromised CA-applet
  - Place a “creation date” in a certificate and revoke all certificates created before a compromise
  - Use multiple CA-applets

19

## ***Where should we use CA-applets?***

- Distributed CAs not a “silver bullet” and not useful in all situations.
- But we should have learnt by now that the only silver bullets in cryptography should be the one used to shoot the PR execs that use the phrase.
- Distributed CAs not useful in very stable networks or where we do not have sufficient trust in the trusted computing platform.

20

## ***Short-lived certificate systems***

- A PKI using short-lived certificates only issues certificates for a short length of time and has no revocation method.
- Useful for dynamically changing populations.
- Distributed CAs could be issued to a user that are similarly short-lived – this reduces compromise problems.
- Distributed CAs could ease central CA's workload especially for renewal of certificates.

21

## ***Dynamically evolving networks***

- A central CA might not always be best placed in a dynamically evolving network.
- Distributed CAs allow a central authority to place CA functionality at the best position within a network at a particular point in time.
- Location of CA functionality can evolve too.
- Multiple CAs to support evolution and backward compatibility?

22

## ***Personal CAs***

- A personal CA is a particular example of a dynamically evolving network.
- A personal CA is a CA that supports the use of public keys within a PAN or PDE.
- Manufacturer could act as central CA and install distributed CAs on their devices.
- Manufacturer cross certification could allow a device to recognise distributed CAs on devices belonging to other manufacturers.

23

## ***Open problems***

- The biggest open problem is to do with synchronicity.
- Mentioned multiple CA-applets several times:
  - Help prevent problems with applet compromise
  - Useful in dynamically evolving networks
- Also mentioned the possibility of using the central CA to 'back-up' the services provided by the CA-applet (avoiding always-on issues).

24



## ***Open problems***

- If CA functionality is offered in several locations then the service must be the same in each locations, i.e. synchronous.
- Central storage facility would solve this problem but would also reintroduce old problems.
- Ideal solution: point-to-point update protocol that allows distributed CAs to be updated on the fly.

25

## ***Open problems***

- Other problems include customising system for use in specific situations, especially for use in a PDE or PAN.
- And can we extend the functionality by allowing a distributed CA to issue CA-applets? How will we cope with the revocation and renewal problems in such a system?

26

## *The End (Honestly)*

*alex@fermat.ma.rhul.ac.uk*

*<http://www.isg.rhul.ac.uk/~alex>*