FORESIGHT

# Cyber Trust and Crime Prevention project

# Technology Forward Look: User Guide

OFFICE OF SCIENCE AND TECHNOLOGY

# Fore*sight* Cyber Trust and Crime Prevention Project

# Technology Forward Look: User Guide

While the Office of Science and Technology commissioned the work, the findings are independent of Government and do not constitute Government policy.

This document was produced for the OST Foresight project on Cyber Trust and Crime Prevention by Bill Sharpe, The Appliance Studio Ltd, bill@appliancestudio.com, and Stefek Zaba, Hewlett Packard Laboratories, stefek_zaba@hp.com, and edited by Martin Ince, martin@martinince.com, in January, 2004. A full version of the report is available on the Foresight website www.foresight.gov.uk



# Scope and Purpose

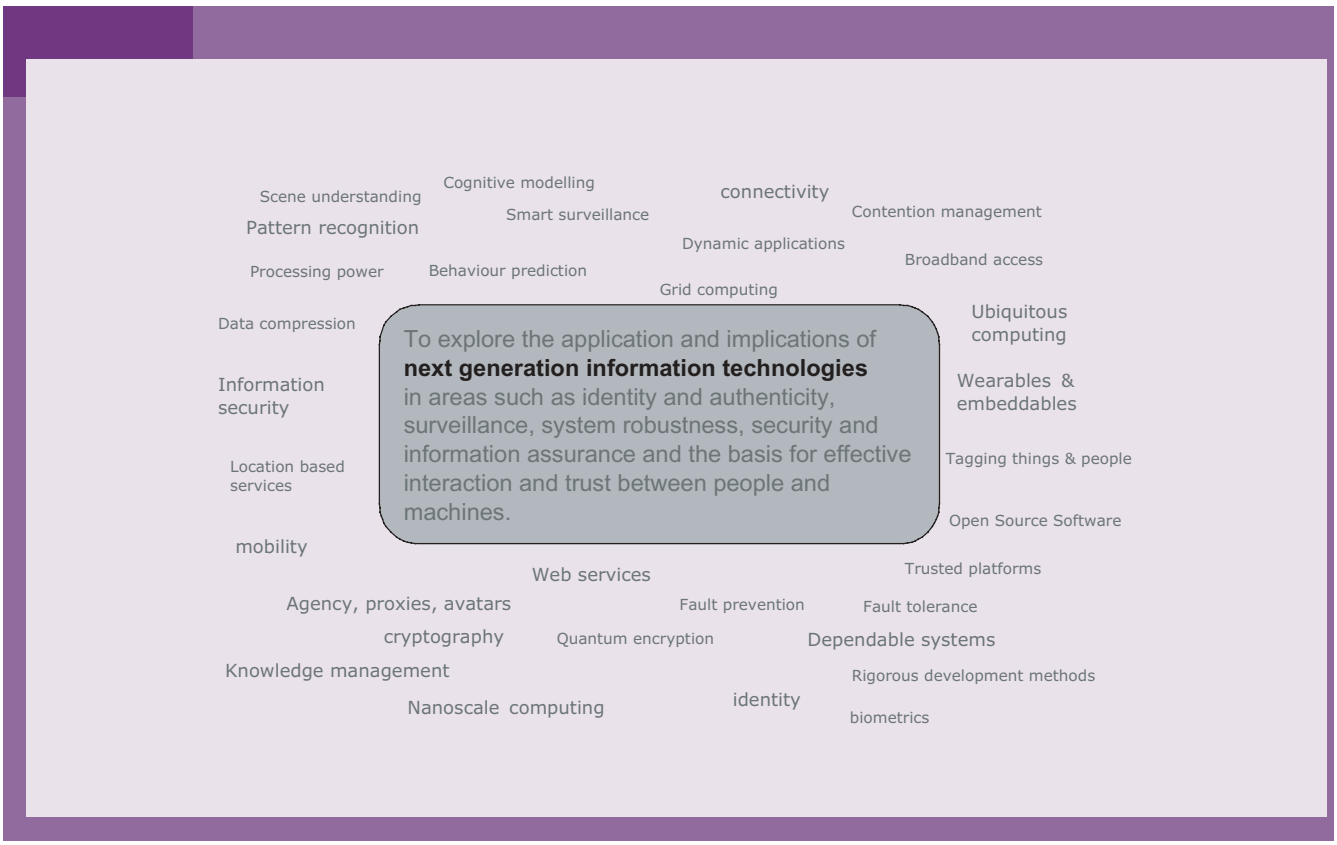


This report has been produced as part of the UK OST Foresight project on Cyber Trust and Crime Prevention (CTCP). Its purpose is to complement the state-of-the-art reviews of science that the project has commissioned with an account of how the technologies

relevant to the CTCP project might evolve. We hope this document will be of wide interest, but its main aim is to feed into the scenario building and modelling work being undertaken for the project by RAND Europe. We hope to reduce the unmanageable range of technology futures to a more tractable number of representative topics – to turn confusion into well-structured uncertainty for the purposes of the Foresight project.

## Structure of This Report

### *Part I*

Describes our approach and the six organising topics selected for detailed discussion.

### *Part 2*

Provides an overview of the emerging world of pervasive computing which provides a setting for the analysis of individual topics, and looks at the challenge of achieving a trustworthy cyberworld.

### *Part 3*

Analyses the technologies under a limited number of topic headings.

Part 1

# Analysis Approach and Selected Topics

Our analysis needs to take place within a view of the future shape of the computing industry, which will shape the way in which information technology reaches into everyday products and services. This approach will allow us to answer questions such as:

- What sort of cyber trust issues will be of dominant concern?
- What new sorts of vulnerability will appear?
- How can cyber trust be enhanced?

**Figure 1.1**

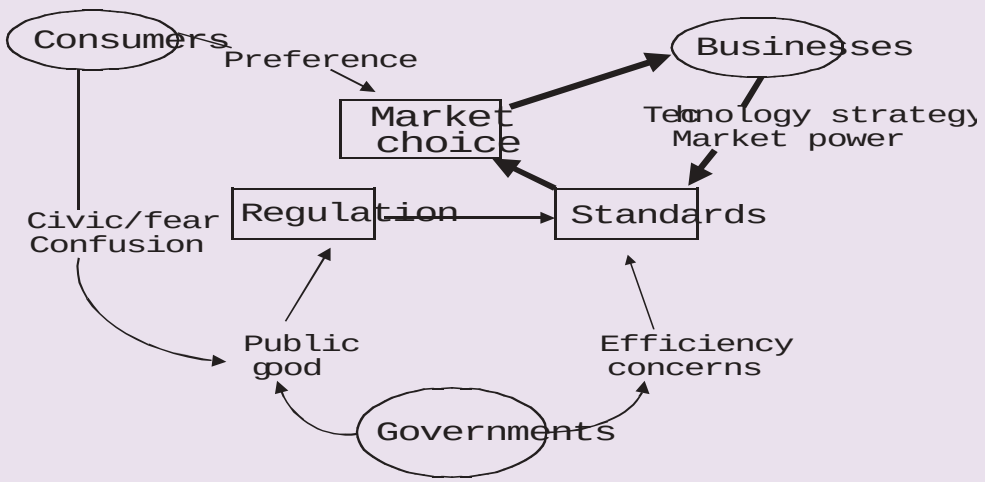


Originated by Ian Page at HP Labs Bristol.

The analysis is based on the model shown above, which illustrates the way that systems of value creation appear and acquire a self-reinforcing structure. The components of the model are:

- the players: governments, public, consumers, businesses
- the fields of play: regulation, standards, market choice
- motivations and actions.



Figure 1.1 shows in a very simple way, that the motivations of the different players in society resolve themselves into a particular way in which things are done, which then becomes stable for a lengthy period of time – the reinforcing loop shown at the top right of the diagram. These stable systems are significant because the computing industry operates in this way. Because it co-ordinates a huge number of technologies within agreed interfaces, it has to have de facto or de jure standards. These evolve from generation to generation as the technology shifts and the players act to change the current order of things, but at any one time there will be some dominant organising themes. At the present time, some examples are:

- stable: desktop corporate and home computing organised around the 'Wintel' model – Microsoft Windows and Intel
- stable: Internet Protocol (IP) established as the global networking standard
- unstable: there is no set standard for secure music distribution and payment, but services are emerging and gaining market traction – some structure will emerge
- unstable: mobile communications are making the transition to 3G in which data services are delivered alongside voice – the industry structure is now in an unstable zone as players vie for positions in data services.

We have chosen to focus on technology which is in wide use and has direct relevance to CTCP. We limited our work to six areas to concentrate attention on key issues for the coming decade. They are:

- web services
- software liability, operating system and open source issues
- trusted platforms, content, digital rights management (DRM)
- privacy and surveillance: ID cards, loyalty cards, tagging anonymity, anything authenticated
- critical infrastructure
- e-cash and why it has not happened.

Part 2

# Context – Pervasive Computing

Even after the dotcom bust, everyone agrees that e-commerce has had profound effects and that computing and communications have restructured both business and many everyday activities.

In the 1980s and 1990s the personal computer was at the heart of that revolution as the tool through which computing reached into everyday uses. The PC provided a de facto standard around which massive investments could be made in every branch of business.

The emergence of the World Wide Web was a defining transition in the computing industry. It shifted the emphasis from individually crafted islands of computing to a global utility that can be accessed, using common standards, by any networked computer anywhere. This is accelerating the shift to a new era when the notion of personal computing embodied by the PC loses its central role. The new era, variously called 'pervasive computing', 'ubiquitous computing', or 'ambient intelligence' is built around developments in three areas:

- the core computing and communications components become cheap enough for 'anything' to connect to the web; today and in the near future that means smart phones, TVs, cameras, hi-fi, car navigation. In the next two decades this will be any artefact for which a reason to connect can be found – if it makes sense for my coffee cup to be online, it will be

- the web with which we are all familiar was built primarily so that people access information on individual websites, and communicate with each other via e-mail. The technical community is now busy working on standards and components for a new generation of the web under such headings as 'web services', 'semantic grids' and 'autonomic infrastructure'. This new layer of technology will allow the exploding population of smart devices, information services and applications to interact directly with each other in dynamic and flexible ways so that the end user will be drawing on many different layers of services and sources of applications at once.



- this shift towards massive numbers of things interacting together in ways that nobody can completely design or control in advance is causing the computing research community to explore a new approach to system design known as agent-based computing. Agents are autonomous software entities with the ability to act and react with one another in a distributed environment. This approach will find its way into many applications, and bring with it a more complex online world with unexpected emergent features.

Taken together, these developments will create a profoundly different information landscape from the one with which we are familiar, and will entwine online digital technology ever more closely into everyday life:

- we will be using the language of autonomy and intelligence to describe wide classes of everyday things which operate under partial human supervision. There will be implications for product liability, security, and service definition. Pinning down responsibility will be particularly difficult as performance will be the outcome of interactions between many agents, combining embedded and online capabilities from a range of providers

- individual things will display intelligent behaviour achieved through the interaction between the thing itself and the capabilities of the web services operating as a utility 'behind the wall' – just as the electricity grid and the power stations that drive it provide electrical power to individual appliances. A simple rule of thumb is that if today a thing relies on power to operate, in the future it will also rely on the intelligence of the pervasive web. This will make everyday life more dependent on pervasive but invisible infrastructure. Questions of the ownership and governance of these services will loom large

- there will be new and massive economic activity in the trading of 'smarts' – pieces of software that improve the performance of things. The web has already fuelled the swapping of music, of virtual characters for games, and of software for intelligent toys. Such precedents suggest that these possibilities will be created well in advance of the legal structures for them. At the same time digital rights technologies that allow control of content within web services will be rolled out. Expect vigorous discussions of who has rights over what information for what purpose. The status quo will not survive unchanged

- there will be a constant struggle to defend this world of ambient intelligence against attacks from viruses, spam, fraud, masquerade, cyber terrorism etc. The risk of new vulnerabilities may prove to be one of the biggest brakes on the deployment and adoption of new capabilities.

The next two sections expand the description of this emerging environment.

## Pervasive Web Services

*"A web service is a software system designed to support interoperable machine-to-machine interaction over a network. It has an interface described in a machine-processable format … Other systems interact with the web service in a manner prescribed by its description…"*
(www.w3.org)

Information technology has allowed business to restructure into virtual networks of organisations with a high level of outsourcing. However, the technology that drove this transformation largely pre-dated the web. It was designed primarily around applications that belonged to, and ran on, the computing infrastructure of each organisation separately. The emerging generation of web services will accelerate this trend by the dynamic linking of service and software components. The first steps down this path have already been taken. For example, an e-tailing portal will use multiple online service providers for such things as information feeds, credit card verification and shipping.



Figure 2.1 compares web services to a familiar, present-day system, a lift in a high-rise building. At the bottom level the lift relies on pervasive electricity delivery. Then in the middle layer there are many generic components and services supplied by a wide range of providers. These all come together in a lift system that is (we hope) trustworthy. By analogy, on the right, the two lower boxes show the new utility view. The Internet is pervasive, and dumb – it is the basic connectivity structure used as a common carrier. Web services are the components and services that can be brought together in systems that connect to the Internet to realise a particular system. For example, city-wide traffic management would integrate all the possibilities of adaptive routing, charging, surveillance and emergency response.

**Figure 2.1** **Trustworthy Web Services Structure**

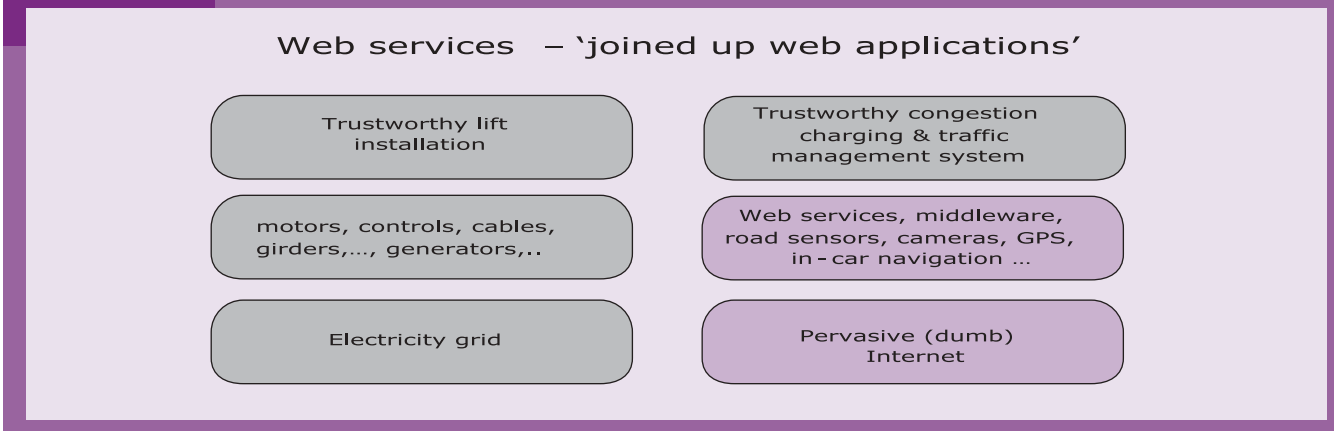


These systems need common standards to work together and these are still under development. A visit to the website of any major IT company will provide white papers and visions of how standards are going to be brought about. There is a lot of technical debate, mixed up with commercial jostling for position. Not much distinction is maintained between:

- technical descriptions of possible standards (such as the 'semantic web') by which services will interoperate
- business and organisational descriptions of the software needed to build web services, or the operational services themselves
- users' experiential views and their visions of the future environment.

For simplicity in this paper we use the term pervasive web services as a generic description of the systems that are likely to be built over the next decade. Figure 2.2 shows some of the terms that are used by the actors in this area and gives some sense of their context.

**Figure 2.2**

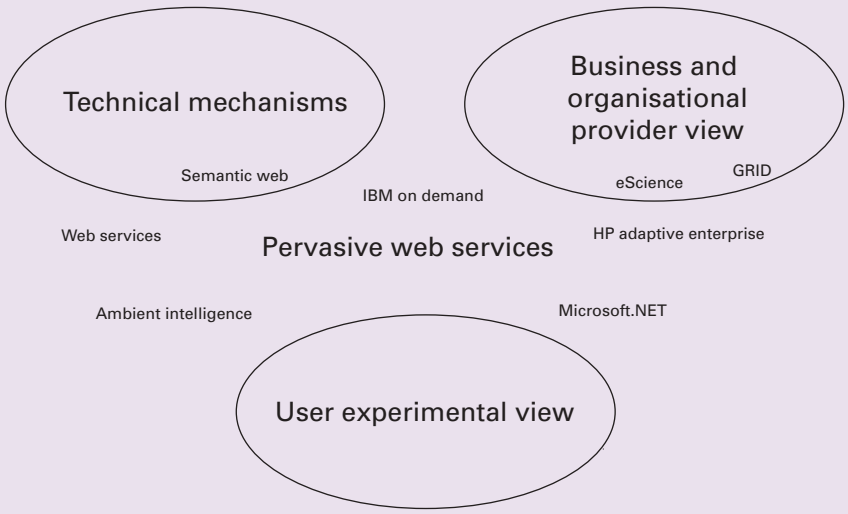


Different regions of the network are likely to be built within boundaries so that they can provide particular, guaranteed levels of functionality. Different sectors of business and society will solve different problems with different priorities, in a different order. So this world will be built by a gradual convergence of many different solutions. A good academic example of what can be achieved within a strongly managed network across many organisations is the e-science Grid network currently under construction in the UK. Here the priorities are such things as access to on-demand computing power for very demanding problems, and very deep sharing of data sets between research groups working on different aspects of a problem. A commercial example would be a typical on-line shopping portal. Here the goal is to provide breadth of access to many suppliers, with consistent mechanisms for secure consumer transactions and logistics. Both the academic and business worlds are building these new networks and solving the practical problems of joining up the many services needed.



Since pervasive web services provide the fundamental setting for the next generation of the online world, their dependability affects all the individual topic analyses in this report. A description of the challenges is the new Dependability Roadmap for the Information Society in Europe produced by the EU AMSD project. Its authors point to these issues (lightly paraphrased for this report):

- scale and complexity. As hardware capabilities improve and costs reduce, there is continuing pressure to attempt to build systems of even greater scope and functional sophistication, especially for software components

- the boundary-less and interconnected nature of the systems, which rarely have a clear-cut frontier. There are often systems within systems, within larger systems. In addition, all nodes are always reachable from everywhere, which results in unpredictable emergent behaviours. They are connected through a common underpinning infrastructure that becomes a critical factor

- heterogeneity and blurring of the human/device boundary – wristheld gadgets, wearable devices, implantable devices

- incremental development and deployment – systems are never finished, evolution is incessant, upgrades, changes in functionality and new features are being added continuously

- self-configuration and adaptation – systems are expected to be able to respond to the changing circumstances of the ambient devices where they are embedded

- multiple innovative types of networking architectures and strategies for sharing resources – Grid-like, peer-to-peer, 'on-the fly' services, etc

- on the down side, there will be more types of fault – in particular a growing danger of malicious faults, both due to individual or organised external attackers, and due to deceitful insiders.

The lesson is that these problems will not be solved in one way across all sectors – they will be tackled in ways specific to each sector, depending on the degree of control that can be exerted over

all the component services and the risk assessments associated with each vulnerability. War, health, science and shopping have very different degrees of control available to them and very different priorities, and will therefore use very different solutions built from the generic technologies.

## From PDA to PDE – The Personal Digital Environment

The move to the post-PC world of the pervasive web is bringing an explosion in the number of connected devices that an individual can own and use in their personal space. A thriving research and commercial community is exploring this new world of wearable and ultra-portable computing. The defining market transition is the move that is currently underway towards 2.5 and 3G phones that will result in individuals being always instantly connected to the web.

The cost of computing technologies means that, at the moment, there is a forced bundling of all sorts of capabilities into smart phones and WDAs (wireless digital assistants). This is a passing phase, to be replaced by what we might call the personal digital environment in which fully connected technology is at the scale of credit cards, jewellery, glasses, coins, wristwatches, etc, and can be embedded into whatever form is most convenient to the user. Individuals will take many different paths through the space this opens up, from eager adoption to total rejection of the 'wired world':

- people will be able to participate in online worlds continuously if they choose. The absorption of people today in their personal entertainment and communication devices suggests that this will significantly change the way people go about their daily lives

- many people will adopt a range of online persona or avatars to provide them with interfaces and buffers to the otherwise overwhelming intrusion of the online world. We may see surprising crossovers of media. with people using 'artificial' communication in face-to-face situations. Just as the physical world has a range of safe and dubious environments, the online world will proliferate venues that will need to evolve the equivalents of Neighbourhood Watch to ensure communal safety. People at work already choose to send emails to colleagues within speaking distance



- assistive technologies can be fully on-line. Continuous tele-monitoring of physical and cognitive sensors and feedback mechanisms will be available and will form part of many people's lives. This could significantly enhance the provision of care in the community but will pose new challenges for dependability and privacy

- we are all just getting used to the huge impact of the e-mail 'trail' that is left behind from daily interactions, with the loss of privacy and control that it involves. This will intensify by orders of magnitude, with many consequent problems – 'You weren't paying attention when you crossed the road and caused that accident – your brain trace proves it'.

The evidence of technology adoption so far suggests that there will be many people willing to move rapidly into this world, well ahead of understanding the issues it raises. Others will be very reluctant, and some will challenge the explicit or implicit agendas of the commercial or government interests that are driving the changes.

# Trustworthy systems in the Cyberworld

*'Cyber trust is a confident expectation in the reliability and value of the Internet and related ICTs, such as the equipment, people and techniques essential to the use of online services'*

Dutton and Shepherd, *'Confidence and risk on the Internet'*, (In Press)*

## Digital Threat Space – The Attacker's View

Any computer user is now all too familiar with the problems of spam and viruses. These are just the most obvious instances of the range of problems, from minor mischief to major attack, that must be considered. They show how new systems both bring new targets, and also put new tools in the hands of attackers. We believe that these are the main types of threat to be considered:

*The source material for this paper is published in full on the Foresight website, *www.foresight.gov.uk* and will be published in book format in due course.

*Faults*

- generational problems that only emerge after innovation has occurred – e.g. Y2K
- major outages from cascading failure
- major negative emergent behaviours such as programmed trading by agents
- simple but pervasive bugs in the infrastructure.

*Mischief*

- viruses, worms, DoS (Denial of Service) and Hacktivism (which interacts with fears about rights of free speech). There is the inherent problem that the speed of transmission outpaces speed of response
- script kiddies – computing allows attacks to be automated, just like anything else
- but such attacks need to be visible to succeed in their own terms and confer peer bragging rights. Detection and response may be easier than when the motivation is fundamentally criminal.

*Crime*

- parasite/host ecology, tries to stay hidden (e.g. Trojan horse virus), needs the system to operate just well enough to exploit its vulnerabilities
- always has the initiative in exploiting vulnerabilities
- insiders have many ways to exploit systems. Outsourcing and the nature of web services mean that it is hard to have confidence in all the systems you are relying upon
- automation makes it worthwhile to use a large number of small transactions, and possible to prepare and launch huge simultaneous attacks



- data mining can find key target information – criminals can do sophisticated market research
- action at a distance changes the game.

*Terrorism*

- visible destruction is the goal
- no need for sophistication in the means
- critical infrastructure or symbolic services are likely targets
- exploit confusion – whole areas of societal infrastructure can be attacked to provide cover for the attack.

## Security – The Cryptographic Toolbox

*'Anyone who creates his or her own cryptographic primitive is either a genius or a fool. Given the genius/fool ratio for our species, the odds aren't very good'*

Bruce Schneier, *Secrets and Lies*

There is a central core of powerful cryptographic technology that is well understood and available, and is the result of many years research and development by the technical community. It offers building blocks from which secure mechanisms can be built for a very wide range of applications. The main ones are:

- symmetric encryption
- message authentication codes
- public-key encryption
- digital signatures
- one-way hash functions
- random number generation
- zero-knowledge protocols.

These technologies are described in the Appendix on 'The Cryptographer's Toolbox' in the full version of this report (see *www.foresight.gov.uk*).

The accepted best practice in designing such systems is to treat their method of construction as being entirely public. Particularly in the case of widely deployed systems it is only reasonable to assume that they will be examined in detail by a capable attacker. Thus, rather than keeping the whole system secret, only the keys – that vary with each use of the system – are kept secret, while the procedures for using them are typically widely published and subjected to hostile academic and technological evaluation. It is extremely hard to produce security technologies that will withstand sustained attack and only prolonged testing in the public domain by the security community can build confidence in a particular method.

Unfortunately such best practice is not nearly as widespread as it should be. It is still common to see vendors claiming breakthroughs for their own proprietary security technologies, and for system and service vendors to opt for closed, proprietary approaches rather than open ones.

When this established toolbox is used to meet the emerging demand for trustworthy systems, we find that:

- we are entering an era of abundance – the basic trend of Moore's law (the doubling of computer power available at a constant price every 18 months or so) that is driving pervasive computing favours code makers over code breakers. It will be feasible to find cryptographic schemes for pervasive embedding for any desired solution

- such systems have, until recently, been limited to highly controlled environments – particularly military or very demanding commercial requirements. And current system architecture ideas are largely inherited from those applications. The dominant norm in these high-risk environments has been the prevention of attack

- investment in cyber trust will create new, application-level, dependable components, and fit them into a new culture that emphasises detection and response as well as attack, and a culture of appropriate risk assessment at the level of the application.



## Trustworthy Applications – Solutions in Practice

*'Whoever thinks his problem can be solved using cryptography, doesn't understand his problem and doesn't understand cryptography'*

Attributed by Roger Needham and Butler Lampson to each other
(Quoted in R. J. Anderson, Security Engineering, Wiley Computer Publishing, 2001).

The cryptogaphic toolbox provides a very general and very powerful set of techniques that can be used for any of the things that we would like to do in the cyberworld. But building trustworthy systems relies on the much wider set of technologies by which computer systems are built and operated, and on the means to ensure they do what they are meant to do – including running cryptographic mechanisms appropriately.

We use the term 'solution in practice' to describe the combination of:

> a trustworthy system – the artefact that people engage with, built to enable certain levels of trustworthiness in principle to be achieved (the building with locks on the doors, etc.)
>
> **plus** the people who commission, build, operate, maintain, use, and attack them within organisations and wider socio-cultural systems
>
> **plus** a 'liability/responsibility structure' that describes who is responsible for what risk at every stage, from commissioning, building and maintenance to use.

For example, a system that allows voting over the Internet involves several levels of protection, to ensure that only registered voters can vote, that no one can vote more than once, and that no one can learn anyone else's vote. It is well understood that even if the hard problems of trustworthy systems are solved at the technical level, the other issues provide the greatest challenge to trustworthy solutions in practice. The remainder of this section summarises the challenges.

The first issue is that systematic handling of issues of trustworthiness is not at all widely established in the systems development community. In particular, there is no stable, shared

notion of a 'liability and responsibility structure' that can be used in systems development. This is an emerging area, and a focus for development and education[1]. The lack of this understanding can have major negative consequences in the roll-out of new technologies for trustworthy systems.

Secondly, the scale and complexity of the emerging generation of 'systems of systems' – the most complex things built by the human race – challenge our understanding of the social process by which such things can be built and maintained. Computer systems are immensely complex things. Making them work without fault is beyond the state of engineering art in all but the most well-constrained cases. In this area Moore's law is not working to our advantage. Systems are becoming more complex, and more interconnected, faster than we are learning how to build them in dependable ways. This issue is highlighted in the summary of the AMSD Dependability report.

The faults to which actual networked systems are prone offer vulnerability for attackers to exploit. In addition, even systems that are highly dependable when used as intended are vulnerable to malicious attack which exploits their unintended properties. In the world of systems, attackers and defenders get similar advantages from the progress of technology. Trustworthiness is the outcome of continual improvement in which attention is paid to detecting and responding to faults and attacks.

There are three main, complementary, approaches to making software that is trustworthy. The first is the use of good software engineering practice, which unfortunately is far less widespread than it should be. Much existing software would be far more trustworthy if good discipline were pervasive. Second are strong engineering approaches based on rigorous (formal) development methods. These have been used mostly in safety-critical situations to achieve much higher levels of reliability than the norm and they are unlikely to become widespread. The third approach is the one embodied in the open source software movement, in which code is developed by an open community and is under the ongoing scrutiny of 'a thousand eyes'.

[1] See for example a leading text in the field: Security Engineering, Ross J Anderson, Wiley, 2001.



The emergence of the open source approach into the mainstream of commercial systems and into policy thinking on cybersecurity has the potential for major, beneficial, impact on the trustworthiness of the cyberworld, by enabling core protocols and components embodying them to be the subject of continuous testing by the whole technical community. However this cannot be expected to change the overall landscape of deployed systems. Open source approaches do not of themselves create trustworthy systems: there is plenty of poor open source software. Only a very limited amount of the critical software that is needed for securing the cyberworld is ever likely to be produced in this way because of the sheer scale of investment that is already embodied in legacy systems and the production of new commercial software[2].

In summary, although the core cryptographic technology available is very powerful, the rapid growth in the scale and complexity of the cyber environment means that even the widespread use of the available technology inherently involves huge, so far unsolved, engineering problems in achieving robust, reliable, trustworthy systems. Understanding basic Newtonian mechanics does not mean that you have the engineering materials and disciplines to build a skyscraper.

## Settling the Cyberworld

*'Cyber Trust will have to be won over and over again'*

The emergence of the Internet into mainstream daily life is a one-time shift in society's infrastructure. As the previous section has highlighted, this shift is happening before we have discovered how to build trustworthy systems, and before what we do know is fully deployed. In looking at the technical challenges involved in this transition, the following consequences provide the context for cyber trust and crime prevention for the individual topic analyses that follow:

- a new frontier or an emerging economy needs well-understood mechanisms and institutions, alongside everyday norms and behaviours for the new situation. Many lessons can be drawn from the first generation of the Internet, the dotcom crash, and the subsequent emergence of successful e-commerce. There will be a lot that is old about the new cyberworld

[2] See Raymond, E.S. The Cathedral and The Bazaar, 1999, O'Reilly for full discussion of Open Source.

- trustworthy systems will develop fastest where commercial or governmental interest motivates action. The cost of insecurity is mostly an externality. This points to the importance of policy levers that are based on understanding and changing fundamentals in such areas as liability

- the environment for this shift into the widespread deployment of trustworthy systems is different from those in which these technologies were hatched (military and safety-critical), and need to be based on risk management in which detection and response play equal parts with prevention. Enabling compliance with norms (such as paying to copy content) will be as important as protection and denying access. There is no technology 'solution' to all the vulnerabilities and threats in the emerging cyberspace any more than there is in the more familiar arenas of everyday life – we still need detection and enforcement

- people appear not to care very much about security until it affects them directly. Culture and institutions respond to systemic problems and are pain-driven

- all the trends appear to be increasing the number of attackers, at least at the level of mischief and opportunistic crime – pain is going to increase. The complexity of new systems is opening up new space for entrepreneurial criminals

- the benefits to individuals of secure and trustworthy systems are an extra cost for system suppliers, while technical complexity can deter users from deploying appropriate security. Where there is market failure, public authorities will need to step in – the topic analyses in this paper suggest this should be at the overall application level (e.g. security and liability for financial transactions, data protection for transactions of tagged goods, safety of transport systems, etc.). However, there is no a priori way of discovering in which areas such intervention will be timely and will drive innovation, and which are the areas where such attempts will have a stifling effect on new developments



- society needs to find the new balance points between risks and benefits, and in some cases to establish a new 'settlement' e.g. in such areas as ownership and rewards. But the technology does not match today's social and political systems. The architecture of the Internet is inherently distributed across borders, and the defence is at the edges.

Finally, because of the rapidly developing range and scope of the new systems that will be developed, and the unsolved technical challenges ahead, we must recognise that cyber trust will have to be won over and over again.

Part 3

# Topic Analyses

## 1 – Web Services

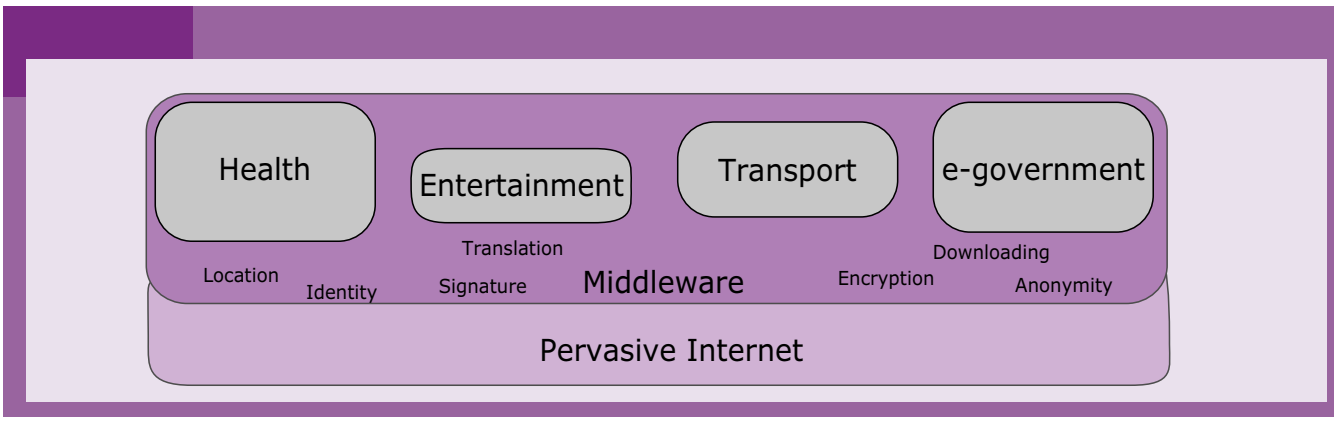


### *Pervasive web services are the major industry focus*

As we explained in the overview, we see the shift to a central role for pervasive web services as the dominant organisational trend in the IT industry. This trend means that issues of industry control and responsibility at this level will become at least as important as those of operating systems are today.

### *Web services built around applications – no general solution*

Different uses of web services by various sectors of society will place quite different demands on achieving an appropriate level of trustworthiness, with very different risk profiles. Consider, for example, the very different requirements of:

- content delivery for home entertainment
- tele-care medical monitoring services for primary health care
- city traffic management and road charging
- mobile online virtual gaming and gambling devices
- e-government
- inter-bank transfers



At the moment there appears to be no way that the technical community can deliver a single standard horizontal 'layer' of technology that can meet these varied demands in a common way. Instead, the industry will offer specific vertical services layered on top of the generic, open Internet services, and there will be strong forces driving inter-sector use of information as different services are aggregated for the end user. There is no generic solution to the problems of transfer of information between services that will arise.

**Example: My community-care tele-medical service ensures that I receive regular deliveries of key medical supplies wherever I am as I travel on worldwide business, and my condition is monitored to provide emergency response. This relies on standard mobile location services and international delivery services which have to share information with my medical provider. The medical service will have to carry responsibility for standards of data protection while aggregating the underlying mechanisms.**

The drive to vertical organisation will stay in constant tension with the drivers for maintaining open access from any client device to the full web – a critical uncertainty will be the degree to which each vertical market goes down a fully open, or controlled/ walled-garden path.

**Example: Compare the closed path of current games consoles, where individual vendors control the entire end-to-end environment, with the open path of home Internet services.**

*Trustworthiness*

End users will interact with this world through public or private branded services. They will have very little opportunity or ability to determine the nature of the back-end services with which they are engaging, and will be easily bewildered by rapidly changing reconfiguration of these services.

This account suggests that policy on the trustworthiness of web services will evolve primarily within the priorities of each vertical application and societal context. Wherever issues of public good are at stake in these services it is likely that the end user will rely on government regulation, just as we do now in many aspects of our everyday life.

*Delivered locally, implemented globally*

'Delivered locally, implemented globally' is the dominant organisational fact in analysing the evolution of web services. Suppliers will be able to choose the regime for back-end implementation that best suits their needs for risk management and commercial benefit. It is noteworthy that there is wide cultural divergence in many highly contentious areas such as freedom of speech, intellectual property rules, privacy and data ownership. This will place significant constraints on the ability of any one national government to govern completely the quality of delivered web services. Critical uncertainties therefore derive from the jurisdiction within which such issues are resolved.

**Example: The US/EU 'Safe Harbor' agreement**

**'In the absence of general legislation, the "Safe Harbor" agreement between the US and the EU, concluded in 2000 after an arduous and conflict-ridden transatlantic process of negotiation, is intended to fill the gap [in privacy legislation] for the private sector … American companies that gain entry to the 'Safe Harbor' are thereby deemed to provide adequate protection for the personal data of EU citizens that are transferred to the US, for they agree to be bound by a set of privacy principles that are overseen and enforced by the Federal Trade Commission. But, in the first place, their subscription to the "Safe Harbor" is voluntary.'**

**Charles D. Raab, *'The Future of Privacy Protection'* (In Press)***

*The source material for this paper is published in full on the Foresight website, *www.foresight.gov.uk* and will be published in book format in due course.



# 2 – Software Liability and Operating System trends

*The merging culture for systems liability*

High trustworthiness was not an issue in the broad technology development practice of the last two decades – the issues were connectivity, functionality, and time to market. Society has so far bought into this trade-off of liability versus flexibility and functionality. This is partly due to the fact that the applications were useful and valuable, but did not penetrate far into areas of high personal concern such as financial, health and legal records.

In keeping with this overall culture, the software industry as a whole operates without levels of liability typical of mature engineering in other areas of society. In general, the software industry shows no signs of embracing liability in the way it would be commonly understood in other product categories.

As discussed above, pervasive web services, developed around the needs of particular sectors, are the new commercial battleground. Computer systems are now touching on areas of high citizen concern for trustworthiness such as finance, health and education. Web service vendors therefore want to offer service-level guarantees of trustworthiness spanning back-end systems and front-end devices (PCs, phones, TVs, etc.) cost-effectively to their clients. Liability is therefore coming into the IT industry at the level of services and systems rather than individual software components. Low trustworthiness of underlying components and services will be barriers to uptake, and will impose considerable costs on providers, and cause them to search for solutions to this problem.

*The changing role of the operating system (OS)?*

In order to achieve overall system trustworthiness, it is in the interests of service and system providers to build on a base of highly trustworthy core OS components. This is a central concern in the risk management of their service liabilities. This will force a response from the OS providers.

Up to now, both hardware and OS vendors have been able to use the evolution of a proprietary OS to embrace ever more functionality and as a powerful way of maintaining market position, but this approach is at odds with the modularity, stability, simplicity, and open scrutiny that are necessary for highly trustworthy systems. In particular, it is widely accepted by the technical community that security mechanisms should be open, so that they can be the subject of testing and analysis for vulnerabilities. This is not yet the industry norm, where many vendors put forward proprietary solutions to strengthen their commercial position, and is at odds with the closed approach to operating systems.

Two types of response are already in evidence, and might be reinforced by the policy choices of customers, vendors or regulatory authorities:

- convergence by the system vendor community on open secure protocols, whose definition is available to everyone for inspection and implementation. Enabling (rather than controlling, or mandating) policy initiatives to establish stable definitions of the open secure components might be seen by regulatory authorities as timely, and represents a key scenario uncertainty. The right enabling standards might have major catalytic effects on the uptake of pervasive web services, and could therefore drive powerful lock-in

- emergence of suppliers of highly reliable, minimal, modular, operating systems; this could easily happen in conjunction with open source operating systems consolidating a mainstream role in the industry. At the heart of such a shift would be the dynamics already operating in the open source software community (e.g. around Linux, Apache, OpenBSD etc), which creates an open community process for continuous testing, improvement and enhancement of modular, trustworthy core OS components.



**Example: The secure sockets layer (SSL) protocol**

**The development of the SSL protocol for securing the transmission of information between web browsers and servers provides one clear example of standards lock-in aided by open source dynamics. SSL was developed by Netscape, which at the time had a clear lead in the developing market for web browsers and servers. The spur to developing SSL was consumers' perceived worries about sending their credit card numbers over the Internet, and the initial released version of SSL provided adequate protection for this task, despite some security weaknesses in both the implementation and design. To make its design the de facto standard, Netscape publicly released the full specifications, and implementations that interoperated with Netscape's soon appeared in other commercial and free web software. Improvement and standardisation through the Internet standards body (IETF) came after the de facto standard had been established; and lock-in is now so complete that the SSL protocol is being widely used for such purposes as 'virtual private networking' – providing secure communications for an enterprise's remote workers – despite being technically less suited to the role than protocols developed for that specific task.**

# 3 – Surveillance and Privacy

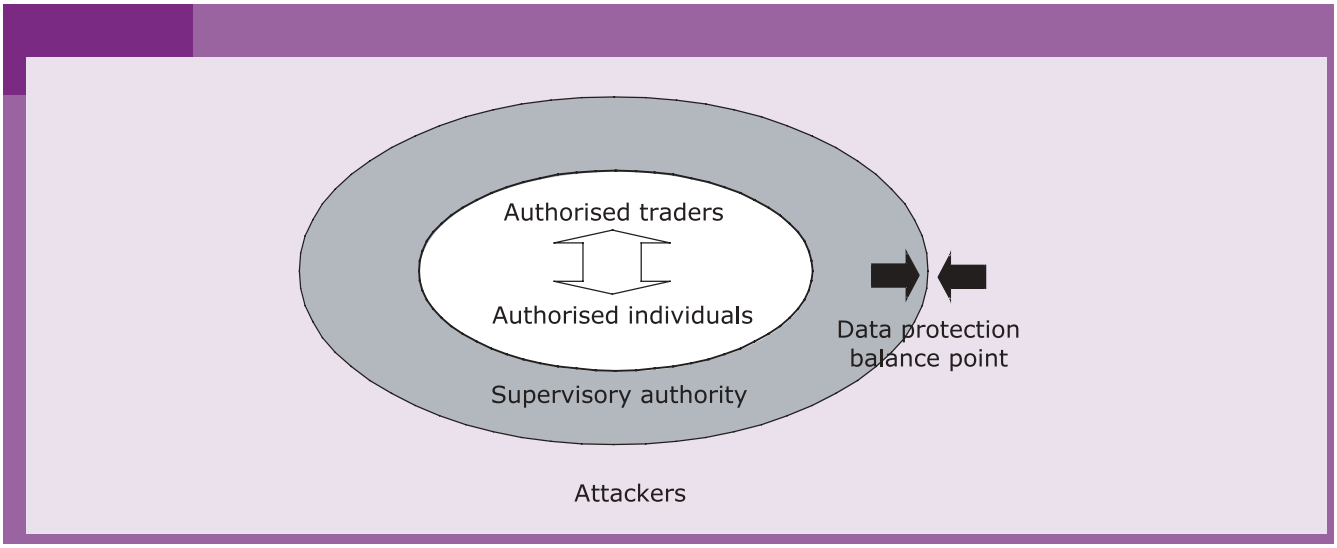


*More and more personal information in other people's hands*

The evidence of recent years is that most people will readily adopt more and more digital technology into their personal space for the value it has to them, and accept it in the public space for security. People are very ready to give up personal information to a wide range of services in return for various personalisation and price incentives.

Many people collude with a situation of low trustworthiness by interacting with potential attackers who operate legitimately in areas such as pornography that are legal but where people do not want the levels of security that would demand high levels of visibility.

The natural effect of authorised and legitimate trading is to vastly expand a large 'grey area' of information (e.g. location, traffic and transaction records), which supervisory authorities will demand be made available to them. Supervisory activity for the detection of and response to attack will further increase the deployment of privacy-destroying systems.

Precisely the same domain of information will inevitably become the target of attackers who can be presumed to equal the authorities in motivation, and skills (though not necessarily resources). Attack technology will use all the new digital threat capabilities. On the positive side, technology has a role in auditing misuse which could become much more extensive if government policy permits and the costs of doing so can be supported.



There is no natural boundary to this growing grey area of visibility that the authorities might wish to demand in the name of safety and enforcement – no 'red line'. There is no counterbalancing universal breakthrough technology 'fix' that can restore to individuals the sort of privacy they currently have. But privacy-enhancing technologies and regulation can change the dynamics of particular instances.

*Societal, technological, governmental responses*

Data protection legislation allows the flow of information, but places legal restriction on its use, rather than bringing in technologies that grant automatic anonymity or protection. This will probably be the path chosen. Technical approaches could be developed that change the dynamics in particular instances.

Loss of privacy to random third parties acting legitimately is preventable with modest effort – analogous to the everyday levels achieved for credit cards, and the restrictions embodied in data protection legislation which govern such things as the flow of information from one use to another. It is relatively easy to protect the content of transactions.

The trend towards cryptographic abundance means that widespread use of privacy-enhancing technologies becomes economical. This could enable some levels of anonymity within the digital domain by allowing such things as strong-use multiple anonymous cyber-identities detached from physical identities, anonymous e-cash, and zero-knowledge techniques. However, it is not in general feasible to protect against traffic analysis and this can be very revealing (as in the current use of cellular phone location information).

Government could decide that the overwhelming volumes of personal information that are coming create the risk of a significant digital backlash, with people opting out of the cyber economy in many ways. This might be a spur to encouraging the use of anonymity-preserving systems. However, there is nothing apart from consumer backlash or regulation to drive this – and all the commercial and supervisory motivations, and current consumer attitudes, work the other way.

The notion of privacy as 'balance' between individuals and society is the conventional model, and may itself be part of the problem of recovering privacy. The paper by Raab (in press)* argues for recognising the social importance of privacy more explicitly, 'addressing the distributional question "who gets privacy" in much the same way as social policy addresses questions of deprivation and inequality.'

The governmental response to this trend towards reduced privacy is a critical uncertainty for scenario-building – it may be that, faced with radical loss of privacy and no broader framework of privacy as a societal good, there will be a widespread backlash that could hinder the deployment of whole classes of e-services which would otherwise be considered broadly beneficial. People may choose to become 'digital hermits'[3] to avoid what they see as an unwarranted loss of privacy and civil liberty.

There is the possibility of some areas of global convergence on international instruments for privacy protection, that provide framework for globally outsourced services (Raab, in press).

*The source material for this paper is published in full on the Foresight website, *www.foresight.gov.uk* and will be published in book format in due course.

[3] Thanks to Brian Collins for this evocative term.



**Privacy-Enhancing Technologies**

**Example: Managed Private Record[4]**

**Just as we have bank accounts that provide us with tightly managed and legally controlled mechanisms for money, we can imagine a similar system that holds records of all an individual's transactions and interactions with the on-line world. The controls on this record would maintain its integrity and access to it, which would provide the individual with certain rights over the ability to provide proofs under his or her control and defences versus malicious attacks. Access under warrant would satisfy law enforcement needs.**

**Example: eCards ('e' for entitlement)**

**Identity cards generate much debate, yet people make daily use of a wide range of membership, loyalty, credit cards, etc. which are used to grant them entitlements while at the same time enabling the provider to build up information about them.**

**Data protection enables people to opt in or out of some levels of use of such information, but only at a very coarse grain, and does not provide any anonymity in the use of services. Cryptographic abundance means that it is perfectly feasible that in future eCards (physical and virtual digital identities) might provide the user with choice over the level of information to be collected about them at the level of transactions. Vendors might be obliged to offer such choices in issuing cards, so that access to everyday services does not come with automatic loss of privacy.**

**Government might be able to lead the way with its desire to offer identity cards that act as guarantors of rights to government and other services. There are not many exemplars of using guarantees of anonymity to encourage participation for the public good, but voting in democratic society is one such: it is deemed to be an essential property of the system to ensure anonymity. This concept may need to be extended with the full capability of abundant cryptotechnology to encourage broad societal participation in the civic cyber society.**

[4] This builds on ideas from Robin Milner and forthcoming BCS (British Computer Society) Grand challenge (Brian Collins communication).

**Privacy Enhancing Technologies *(continued)***

**Example: sRFID**

**RFID tags have so many implications for the visibility of transactions that they will need to be the focus of explicit policy, which might:**

- **mandate their use for classes of product that are insured**
- **place legal restrictions on ownership of re-programming technology, but allow legitimate gift-giving, second-hand selling, charity-shop donation etc. – build on experience with phone SIM cards**
- **invent and introduce 'sRFID' – an imagined future security-oriented version that explicitly provides for domains of access that can be changed through the lifetime of the tag.**



# 4 – Trusted Platforms and Digital Rights Management (DRM)

Trusted computing and DRM are two very active and closely related areas of industry development of importance to cyber trust. DRM refers to approaches by which owners of digital content (information, software, movies, games, etc.) can control the copying and distribution of their content to protect copyright, and be paid for its use. This is technically very hard to do in any general way, especially when the content is used on general-purpose computing devices such as PCs. Trusted computing provides low-cost hardware integrated into common computers and other information appliances to improve the security achievable.

In essence a 'trusted platform'[5] is authenticated – you know that it is what you think it is, and that a binary image that is on it is what you think it is. It typically uses specialised hardware to overcome the security problems caused by the essential modifiability of software.

Early versions now available allow secrets such as cryptographic keys and sensitive documents to be held in protected storage, for a co-operating chain of software to securely record what programs are loaded into the computer, and for secure transmission of that record to another computer. Controversially, the 'protected storage' areas may include information secured against the user of the computer – for example, copyrighted material where the copyright owner proves access to the content only through 'approved' software.

Controlling what runs on which platform does not in itself tell you anything more about the reliability of the binary image, or the reliability of the combined system, or the interactions between the binaries and the system, or any of the other sources of insecurity or lack of dependability. However, in a situation where there is an overall system management environment which governs what should be running on the platform, trusted platforms can make a big contribution to security. They also support the enforcement of

[5] This is an unfortunate misnomer; it would be more appropriate to call the platforms 'trustworthy' – the intent is then that they become trusted by their users.

software licensing. For both these reasons a great deal of work has been put in by the computing industry to establish standards for trusted computing platforms. The issues in the uptake of this technology revolve around who it is that makes the decision about what should and should not be running on the platform – by definition the technology is only useful under conditions of strong system management.

The first adoption of trusted computing platforms will be in corporate enterprise computing environments for the internal operation of a company's IT systems. This will roll out with replacement cycles over the next five years. The same technology could play a major role in the consumer market, and here the dynamics are different.

Today's generic consumer PC does not operate within the sort of managed environment that benefits from the trusted platform approach, so early deployment is more likely on client devices sold as part of a highly managed service such as mobile telephony. Home PC's, which have peaked at around 50-60% penetration, might get drawn into the same logic later.

There appears to be a virtuous reinforcement between trusted platforms, trusted binaries, highly dependable components, and online service providers who want to guarantee levels of service and have to accept service liabilities. This suggests that there will be a continuing drive for service operators to provide managed client environments. A critical uncertainty is whether any of these service/client systems adopt the open source logic, or whether they attempt to stay within closed, proprietary models.

As discussed above, trusted platforms will generally be more secure if they are based on public technology, open to ongoing scrutiny and improvement – this is best security practice. However, the commercial needs of the consumer content and communication industries will continue to drive a variety of proprietary, 'good enough' platform solutions, such as those already seen in phones, TVs, DVD, etc.



The drive towards managed clients will be counterbalanced by demand for open access to the whole web. Consumers are unlikely to buy a music player that can only download from one record label, though they will buy a service that offers a range of labels of interest to them. The issues of the management control of systems and who decides what can be run will therefore loom very large in the consumer domain. On the vendor side, the investment in security will be in the context of the economic value of the delivered service, and risk management will be a matter of economics as well as technology. It might be cheaper to reimburse you if you lose your music collection than aim for very high reliability.

Trusted platforms and binaries are also an enabler of anonymous transactions from anonymous locations – see page 29, Surveillance and Privacy.

DRM issues are likely to become of concern in evermore areas as as web services intersect with more market sectors as they have done with entertainment, for example, and in the use of medical drug data.

# 5 – Critical Infrastructure

*An emergent future of crisis and response*

It appears to be inevitable that the global world of 'system of systems' will be implemented faster than we learn how to predict all its emergent properties or define appropriate standards of dependability. The experience of power and phone outages and of rail problems show how hard it is going to be to evolve major infrastructures into their new socio-technical environments. We must assume that we will evolve through 'best efforts' approaches but with a significant element of crisis and response.

*Global Internet vulnerability*

The Internet could be brought down by very knowledgeable people. The technical community 'holds the keys' to the global system and operates as a goodwill network, gradually engineering in more robust technology. This community is not responsive to any attempt to command or co-opt by commercial or government authorities. As in other domains (such as the emerging problems of bio-terrorism) the technical community will need to be more aware of responsible practice and of the need not to publicise potentially critical vulnerabilities.

As major infrastructures shift to running on top of the global digital infrastructure, they become equally exposed to the same class of digital attacks as the current Internet and web.

*No system is an island*

Critical services that used to be under a single-system management regime within well-controlled boundaries will be increasingly implemented as a 'system of systems', with many layers that stretch across sector and national boundaries (see Web Services, page 23).



*Dependability deficit*

Even the very best systems and software engineering practice cannot produce totally reliable software on the scale that will be deployed, and the expertise to produce the very best is in very short supply; so only a very small part of the 'system of systems' world will be anywhere near the standards we would like.

Major systems integrators (IBM, HP, EDS, etc.) are developing approaches to system integrity in the face of disasters and attack, driven by the needs of their commercial customers. Initiatives such as 'autonomic computing' reflect that intent, rather than being actual technology solutions.

*Strong technologies for dependable systems*

There are rigorous software engineering processes that can play a significant role in the delivery of dependable systems.[6]

- Although a lot is understood about strong engineering techniques for dependable systems, such as security, there has not been systematic uptake, because it has not driven the economics of the industry. Features, and being good enough, have won the day. The current growth in complexity means that good practice brings increasing rewards, as seen in Figure 3.1 below.

**Figure 3.1**

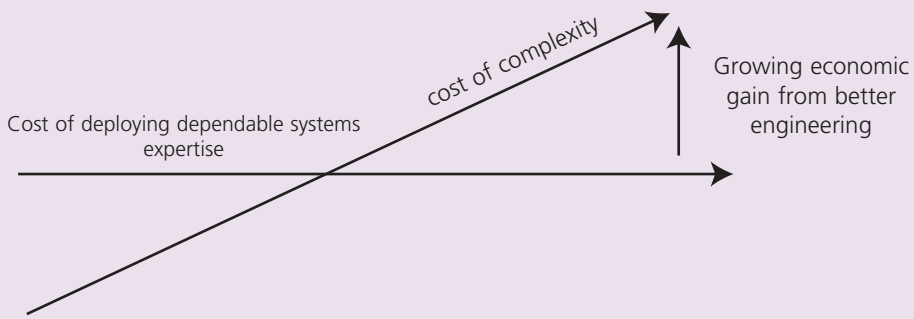


[6] See the Randell & Jones review paper and AMSD Dependability Roadmap for a full discussion.

- The default situation is that only safety-critical systems will drive the adoption of strong dependability engineering processes – otherwise good enough is good enough.

- It will be a slow generational shift, but there are policy opportunities to speed diffusion and a good case can be made for initiatives to build a distinctive UK position in this area which is already a strength.



# 6 – E-Cash

This topic has been chosen because e-cash has been the focus of considerable investment in new technology and trials, but has not emerged, while conventional payment mechanisms like credit cards are making a successful transition to the cyberworld. So it is a good case study in sociotechnical dynamics.

*e-cash versus e-payment*

For high-value items the conventional payment mechanisms are holding up well in the transition to the web, and are being extended (for example PayPal) to enable person-to-person payment as well as person-to-trader. We expect this trend to be consolidated through the period under consideration. For low-value items the dynamics of e-cash adoption in the offline world do not work:

- by definition, the value of items is small, so there is a big fixed cost of new infrastructure to be amortised over small amounts
- while penetration is low, the user has to have cash anyway, reducing the benefit
- trustworthiness of e-cash requires that tokens do not stay in circulation and out of touch with the central mechanisms for long (maybe only for one transaction), so the infrastructural difference from existing payment mechanisms is quite small, and the cost high. New mechanisms usually need a tenfold performance/benefit advantage to get going and this has no such 'killer app'.

Likewise, in the online world, attempts to launch new micro-commerce schemes have not taken off (Beenz, e-gold, Peppercoin). In addition to the low value/benefit and critical mass trader problem, the Internet service providers are not set up to offer transaction-level settlement structures to web traders. We should regard the emergence of micro-commerce e-payments on the web as a critical uncertainty.

**Example: Mondex**

**The Mondex system was an ambitious attempt to supplement and eventually replace physical cash with smartcard-based value transfers, which included unsupervised card-to-card transfers between individual cards. It was trialled on a number of occasions for general deployment: in Swindon, in Manhattan, and in two separate Canadian locations. The most detailed analysis of the unenthusiastic response from consumers is from the second Canadian trial (Stalder, 2001), which itself had the benefit of experiences from the three earlier trials.**

**Common to all the trials seems to have been a disinclination by consumers and merchants to adopt a new 'general' means of payment which is neither sufficiently general to replace cash in one fell swoop, nor offers any compelling advantage over existing means of payment. A second significant factor seems to have been consumer and advocacy-group worries over the privacy of this electronically mediated payment system: the Mondex developers seem initially to have been emphasising different aspects of the system's design to the regulators on the one hand (in that case, the potential auditability) and to the consumers on the other (in this case the 'privacy' of direct card-to-card transfers).**

**Open disclosure of the trial system's actual capabilities and design choices was further hampered by proprietary concerns on the part of the developers, and their well-meant (but perhaps naïve) wish to monitor transaction detail to maximise the smooth-running and lesson-learning of the trial. While the prospects of the Mondex system either replacing cash or becoming a dominant Internet payment system have at best receded into the middle-distant future, the spin-off technologies of rather more tamper-resistant smartcards and of operating systems for those cards which allow multiple applications to share the same card securely, are being used in a number of large-scale 'closed-user-group' deployments, particularly in the Asia-Pacific region. Notable among these is its use as an immigration/identity card in Hong Kong.**

**http://felix.openflows.org/html/ANT_Mondex.pdf, and published as Stalder, F.: *'Failures and successes: Notes on the development of electronic cash'*, The Information Society, Vol. 18, Nr. 3, 2002.**



*Mobile e-commerce*

By contrast, mobile infrastructure is built from the ground up to support transaction level micro-charging, and the market in such things as ring tones is well underway. Mobile phones already hugely outnumber PCs. Buying music (i-tunes) is taking off and mobile music players are a new technology gravity-well. The pervasive web and personal digital environment will rapidly expand a huge mobile market and there are no obvious barriers to an extension to full mobile micro-commerce over the next few years. This should probably be regarded as pre-determined, and analysed to discover what key issues and policy levers it involves.

*New cybermarkets*

There will be massive economic activity in the trading of all sorts of software and content between individuals. The web has already fuelled new areas of software swapping involving music, virtual characters for popular games, and software for intelligent toys. The precedents are that users will create these possibilities well ahead of the legal structures for them. It is easy to imagine that various sorts of e-cash could get established in one of these areas and 'break out' into general use – or simply come to represent so much value within its own domain that governments want to regulate and tax it.

**Example: Gaming Open Market**

**As reported by BBC Online News: 'Online games now have their own foreign exchange that lets players buy and sell different virtual currencies just like in the real world. The Gaming Open Market allows players who control characters in games such as Star Wars Galaxies, The Sims Online and Ultima, to buy and sell the currencies used in the different game worlds. Players can convert cash reserves in one game into a different currency in another world or sell their virtual money for US dollars. The market now has 29 characters in six different games that act as virtual bank managers in the separate worlds.'**

**http://news.bbc.co.uk/go/pr/fr/-/1/hi/technology/3368633.stm**

**http://www.gamingopenmarket.com/**

Printed in the UK on recycled paper with a minimum HMSO score of 75
First published June 2004. Department of Trade and Industry. www.dti.gov.uk