Certification and Evaluation: A Security Economics Perspective

Ross Anderson
Computer Laboratory
Cambridge University
Ross.Anderson@cl.cam.ac.uk

Shailendra Fuloria
Computer Laboratory
Cambridge University
Shailendra.Fuloria@cl.cam.ac.uk

Abstract

There has been some discussion in the industrial control system security community of evaluation and certification. There are already at least two independent third party evaluators, and some have advocated Common Criteria certification of products used in critical systems. The broader IT security community has considerable experience of evaluation and certification, which we seek to summarise and share in this paper. Certification is not a silver bullet, and can very easily end up as spin rather than substance: as 'security theatre' designed to reassure customers or regulators rather than a genuine risk-reduction mechanism. It can also be very expensive, and once entrenched it can impose deadweight costs on industry that are difficult to eliminate even when certification processes are widely seen as failing. We discuss a number of further issues such as perverse incentives, usability and liability and argue that the industry should proceed with great caution.

1. Introduction

Evaluation and certification schemes have a long history. From about 1800, the growing use of high-pressure steam boilers led to an increasing number of boiler explosions, which in turn led insurers and others to devise certification procedures; the risks also drove innovation (such as Babcock and Wilcox's "non-explosive" water-tube boiler). The commercialisation of electricity led William Merrill to found Underwriters' Laboratories in 1894; it started evaluating products such as fire doors and fire extinguishers on behalf of the insurance industry (Merrill was later treasurer, and president, of the National Fire Protection Association). UL now has over 1,000 safety standards and evaluates over 20,000 products. Other countries have similar arrangements, such as the Loss Prevention Council in the UK, while the European Union has 'CE' marks that signify compliance with health and safety legislation.

Economists have developed an extensive theory, of 'asymmetric information', to explain the underlying phenomena. The seminal paper, on the 'Market for Lemons', won George Akerlof the 2002 Nobel Prize. In this paper, Akerlof imagines a town with 100 used cars for sale, of which 50 are reliable (and worth $2000) while the other 50 are 'lemons' that are forever breaking down (and worth $1000). What will be the equilibrium price of used cars in this town? One might surmise $1500; but at that price, no-one with a good car will be prepared to sell it unless they have to. The market price will rapidly approach $1000 and almost all the cars on the market will be lemons. This market failure occurs because of asymmetric information – the sellers know which car is a lemon, but the buyers don't. Evaluation and certification schemes mostly exist in order to deal with failures cause by asymmetric information. For example, car vendors try to mitigate the used-car market failure by having 'approved used car' schemes under which they inspect used cars and offer guarantees for them.

The insurance market (with which certification is often linked) is beset with two types of asymmetric information problem: adverse selection and moral hazard. For example, Volvo cars have a reputation for safety, yet Volvo drivers have more accidents than average[1]. This may occur because people who know they're bad drivers buy Volvos (so they're less likely to get killed), or because ordinary people who buy Volvos feel safer and thus drive faster, in order to bring their risk exposure back up to the level at which they feel comfortable. The former is called adverse selection, and the latter moral hazard; we'll meet them again in what follows.

2. Evaluation and Certification in Information Security

Evaluation and certification have proved more difficult in the world of information security, where they have a reasonably long history. By the early 1970s, government users in particular had realised that the security offered by commercial computer systems was poor, and was not getting any better. As soon as one vulnerability was fixed, another one would be found. This led the U.S. government to commission a report by James Anderson[2], which recommended that the security of operating systems should be reduced to that of a protection component that could be made small and simple enough to be subject to anal-

ysis and tests, the completeness of which could be assured. This led to the promulgation of a security standard, Trusted Computer Security Evaluation Criteria, also known as the "Orange book"[3].

The Orange book provided criteria for classifying system security into a series of levels – C1, C2, B1, B2, B3 and A1 – depending on how carefully engineered were the mechanisms for assuring the confidentiality of classified information. Orange Book certification became a requirement for computers processing classified information at more than one level (such as a stores system with both CONFIDENTIAL and SECRET data) and this created a market for evaluated systems. However evaluation meant having a system closely examined by engineers at the National Computer Security Center, a division of the NSA. A vendor needed a government sponsor to get a candidate system into this process, and once it was in, it often took 2–3 years. So the entry costs were high, and certified products always lagged well behind the commercial state of the art.

The "Information Technology Security Evaluation Criteria" (ITSEC) was the European response to the Orange book. In this model, only the highest levels of certification are performed directly by government labs; lower levels are performed by commercial labs that are paid by the vendor but regulated by the government. After the end of the Cold War, the European model prevailed in the form of the Common Criteria for Information Technology Security Evaluation[4]. The Common Criteria not only use commercial labs, called Commercial Licensed Evaluation Facilities (CLEFs) for the lower levels of evaluation, but introduced a further innovation. While the Orange book had focussed solely on protecting classified information from compromise, the Common Criteria permit systems to be evaluated against a "protection profile" that specifies what sort of threats are to be assumed, and what sort of protection must be provided against them.

The authorities' objective in setting up this structure was to broaden the system from defence computing to a much greater range of applications, in the hope that this would increase the number of evaluated products and bring down costs. An interesting book chapter on 'Why the Security Market has Not Worked Well', written in 1990, reflects official thinking at the time[5]. Common Criteria evaluations have indeed started to be used in a number of new areas, such as smartcards. Here, the protection profile may not be concerned so much with data confidentiality as with assuring the integrity of transactions by making the device difficult to tamper with or copy.

3. How Evaluation and Certification Fail

By now, we have learned quite a lot about what goes wrong with evaluations. In this section we will attempt a rough taxonomy.

3.1. Inadequate testing criteria

Physical security relies on locks, and it has recently transpired that most of the high-security locks on the market are easy to open covertly using techniques such as 'bumping' – in which a cut-down key is inserted into the keyway, torsion is applied, and the key then 'bumped' with a soft hammer to bounce the lock pins up to the shear line, allowing the cylinder to rotate[6]. This has caused consternation in the industry. High-security locks in the USA are evaluated to UL 437 or BHMA 156.30; yet these standards specify that a lock should resist picking for a given period of time. They do not specify resistance to bumping (or to other advanced techniques such as vibration or mechanical bypass). In effect, the standards were written by the vendors; they tested against threats the vendors could deal with fairly easily, rather than against real-world threats. UL has now set up a task force to rewrite the standards for locks, safes, vaults and ATMs.

Testing what the vendor wants tested rather than what the customer (or other relying party) needs tested is a pervasive problem with the Common Criteria. There are two major problems – whether we are certifying assurance or process, and what security policy gets used. By assurance, IT security folks mean 'an estimate of the likelihood that a system will fail in a particular way' – examples being a safety failure, or a failure that will break the security policy. Evaluation is the process of collecting evidence that a given system meets a given assurance target. However, all too often, security certifications focus on process: did the developers use some particular methodology and tick all the right boxes? In the old days of the Orange Book, the higher levels of evaluation involved rigorous penetration testing, providing a reasonable level of assurance. However, the Common Criteria provide the flexibility for the writer of the protection profile to emphasise process assurance instead.

3.2. Inappropriate protection profiles

But that is not all. Security is not a scalar; it has meaning only in the context of a threat model and protection goals, a succinct statement of which is known as the security policy. This is critical, as it's far too easy to protect the wrong thing. Control system engineers wouldn't think much of a VPN encryption system that protected the confidentiality of signalling but left devices open to service-denial attacks, for example; yet there are firms trying to sell such products. You have to understand what you're trying to protect, and against what attacks.

Locks provide one illustration of what goes wrong: the security policy is about right but the assurance – the testing – is wrong. Another example comes from Iceland, which set up a national medical database system in the late 1990s in the teeth of opposition, on privacy grounds, from most doctors and many members of the public. The security targets specified standard protection mechanisms, such as passwords and audit, while avoiding detailed consideration of the inference controls that would be needed

to prevent abuse of the database – inference controls are a hard problem[7]. In this case, it was the security policy that was wrong.

There are many other examples. Secure signature creation devices were provided for in the European Electronic Signature Directive, and the relevant protection profiles were aimed at smartcards, following vigorous lobbying by the vendor community. This is unfortunate, as smartcards don't possess a trustworthy user interface; any smartcard that acts as a peripheral to a PC can trivially by attacked by infecting the PC with malware and getting the smartcard to sign inappropriate transactions. Again, this was a case of getting the security policy wrong.

3.3. Target scope too narrow / ambiguous

The smartcard case leads us naturally to a broader problem; that many software vendors have obtained Common Criteria evaluations for products that are valid only in unrealistically stripped-down configurations[8]. For example, Windows NT was evaluated for diskless workstations that were essentially useful only as thin clients. Such evaluations ignore both real-world configurations, and the need to update software regularly in response to the discovery of vulnerabilities. Oracle opted for the Common Criteria for several of its products and in the wake of 9/11 hyped security as a marketing strategy, calling their products 'unbreakable'[9]. Several vulnerabilities were soon found in these products, leading to negative publicity but senior Oracle staff maintain that overall the security campaign was a huge success in terms of sales. And Windows Server 2003 SP2 was certified as EAL 4+ in Feb 2008; several security vulnerabilities have been identified after that as well[10].

A further, and striking, example comes from evaluations of cryptographic hardware. These are carried out not just under the Common Criteria, but also under an older but similar U.S. scheme, the federal information processing standard FIPS-140. This mechanism is narrowly focussed on certifying the design, development and implementation of cryptographic equipment, which can be certified from level 1 (the lowest) up to level 4 (the highest). At the highest level, the certification claim is that the equipment is tamper-proof: that is, there is no known way for an opponent to extract the keys protected by the equipment, whether by drilling, power analysis, or any other known technique. The IBM 4758 was the first system ever to get a FIPS-140 level 4 rating. However, once we looked at it closely, we found critical security vulnerabilities that we could exploit to extract key material[11]. We did not attack the device's hardware (which was certified) but the software that ran on it (which wasn't). Most of IBM's customers had been unaware that the certification related only to the hardware, and that the overall system they purchased carried no guarantee at all. What's more, IBM didn't exactly go out of its way to educate them.

3.4. Race to the bottom

When a number of certification authorities compete, customers will choose the easiest ride, and this can lead to downward pressure on standards. For example, there is much controversy in the UK about the school leaving exam, the "A-level". A-level exams are sold by a number of competing companies, and naturally schools enter their pupils in the exams they think are easiest. By 2007, even the chief executive of one of the firms was admitting that declining standards and public confidence were a problem[12]. The exam vendors are regulated by the government, so universities don't want to discriminate in favour of one brand of A-level or against another.

The framers of the Common Criteria were aware of this hazard, and thought it could be dealt with by having national authorities license and regulate the CLEFs. However, these authorities are curiously reluctant to revoke a local CLEF's license. In the late 1990s, the UK and German governments were tussling with the French government over standards for digital tachographs, objecting that French proposals (which would favour French suppliers) were insecure. The French finessed this by writing a protection profile that suited their industry and having it evaluated by an English CLEF. The UK government was then faced with a dilemma: it could either let the French prevail in Brussels, at some cost to road safety, or it could challenge the evaluation and thus undermine confidence in the CLEF system. It chose the former[13].

There have been many other cases; the promoters of the Iceland database selected a CLEF with little relevant expertise to evaluate its security target. In fact we know of no instance of a CLEF losing its license, regardless of how many vulnerabilities are later found in products it evaluated. If the operators of the Common Criteria were serious, they would shoot a CLEF from time to time pour encourager les autres. We'll come back to this topic again later.

3.5. Adverse selection

We mentioned that adverse selection is a pervasive problem in insurance. Sick people are more likely to insure their lives, while firms in rough areas are more likely to insure all their assets against fire and theft. If the insured has more information about risk than the insurer does, then premiums will rise and low-risk individuals may self-insure, leaving the insurer with a risk pool of declining quality. Insurance companies mitigate this problem using certification mechanisms: an applicant for a large life insurance policy, for example, will be asked to undergo a medical exam.

Yet adverse selection also affects the business of certification, as was established by Ben Edelman . In 2006 he looked at websites certified using 'TRUSTe', an industry scheme for endorsing websites as non-malicious; he found that certified websites were more than twice as likely to actually be malicious (in that they would try to infect users with spyware, or otherwise violate their privacy)

than random websites. In other words, if it's cheaper for dirty website operators to buy a privacy certification than to actually clean up their act, then certification is what they'll buy – and the certification will become worthless. Economists who have studied this phenomenon model a certificate as a signal of the quality of the underlying product or service; and for a certification scheme to work, the signal must be cheaper for a high-quality vendor to purchase than it is for a low-quality vendor.

3.6. Moral hazard

Moral hazard also affects the insurance industry; why should someone who has insured the full value of his car bother to lock it? This problem is mitigated by the design of insurance contracts: the insured will typically have an 'excess' or 'deductible' amount that the insurer will deduct from any claim, and contract prices also depend on previous claims history. In this way, the interests of the insured and the insurer are better aligned.

What forms of misbehaviour can be expected from principals whose systems have been certified secure? A good example comes from banking, where in some countries the banks have considerable success in passing the risks and costs of card fraud on to merchants and customers; the EMV smartcard payment system ("chip and PIN cards") issued by banks in Europe, and which is now being rolled out in Canada, is an example of this. Banks rely on security claims about this system to dump liability: "your PIN was used so you must have been negligent or complicit"[15]. This in turn undermines security; it's a fundamental principle of security economics that if Alice guards a system but Bob pays the costs of failure, Alice won't work hard enough. Sure enough, UK card fraud has increased by half in the three years since this misconceived liability-engineering system was introduced.

3.7. Framework abuse

A curious abuse of the Common Criteria has emerged in the financial sector. The PIN entry devices (PEDs) used by merchants to capture transactions from the EMV smartcards are claimed to be "Common Criteria Evaluated" against the threat that merchant staff or others may tamper with them in order to harvest card and PIN data for use in fraud. According to the relevant protection profile, it should cost at least $25,000 to tamper successfully with a single PED. Yet penetration tests on the two most common PEDs, showed that it was trivial to hack them[16].

GCHQ, the UK government agency responsible for regulating CLEFs under the Common Criteria, disclaimed all responsibility: the PEDs had not been Common Criteria Certified, but merely evaluated by a CLEF according to a protection profile produced by the banking industry. For a proper certification, the vendor would have had to file the evaluation report with GCHQ, which would have published it. APACS, the bankers' trade association that sponsored the evaluations, had kept the results confidential, along with the name of the CLEF responsible, so the CLEF was shielded from discipline (and ridicule). So the banking industry's claim that PEDs are "Common Criteria Evaluated" was of no real value – other than perhaps to persuade the ignorant that the card payment system is more secure than it actually is. In short, it's not security – but security theatre. It is also telling that GCHQ (and the other agencies that regulate the Common Criteria in other countries) are insufficiently vigorous at protecting their brand. If they were private-sector firms, they'd have stopped the bankers infringing their trademark years ago.

3.8. Other framework problems

We already remarked that the Common Criteria don't deal well with products that are regularly upgraded – such as operating systems with their monthly patch cycle. Another bugbear is usability. Security usability is a significant problem and the focus of growing research efforts; more and more fraud is based in deceiving users rather than achieving purely technical penetration of systems. Yet the Common Criteria were not designed to assess usability; they focus on technical aspects of systems and handle usability issues badly.

3.9. Other security-economics failures

Certification is often used where a third party is expected to rely on the protection provided by the evaluated product. But where the relying party isn't the principal who buys or maintains a security product, there is an ever-present risk that the express or implied liability transfers will undermine either the system's security or even its business justification.

We already remarked that banks' reliance on (low-quality) certification of card-payment system components was part of a larger scheme of liability transfer that led to an increase in fraud. As for electronic signature creation devices, both vendors and legislators hoped that a regime of certification would lead to their becoming universal in Europe. Yet today they are almost nowhere to be found. Part of the reason is that the Electronic Signature Directive gave electronic signatures created with them a presumption of validity. Translated out of legalese, that meant that if you bought such a device, then you agreed thereby to be liable, for an unlimited amount, to anyone in the world, for any transaction carrying a signature that appeared to have been made with your device, regardless of whether you (or your device) had actually made that signature. It is hardly surprising that both businesses and consumers have declined to purchase such devices. In short, certification – even when backed by legislation – doesn't provide a silver bullet for solving liability problems or even more general risk problems. You can move liability around, but this often has unexpected consequences.

4. Certification and Control Systems

The Sandia SCADA program came out with specifications in 2002 that would take the form of a Common

Criteria Protection Profile[17]. Later, the National Institute of Standards and Technology (NIST) started a forum called the Process Control Security Requirement Forum (PCSRF), which worked in collaboration with the industry to develop Protection Profiles for SCADA field devices. These profiles were to be connected "using the methodology defined in the Common Criteria"[18]. An English version that could be more easily understood was also produced[19].

The above discussion should highlight the risk of piggy-backing a proprietary certification scheme on the Common Criteria. Section 3.7 in particular gives a graphic example of how a "Common Criteria Evaluated" certification scheme, that falls short of the full Monty of "Common Criteria Certified", can end up deceiving the relying parties. In that case the relying parties were bank customers, who are neither well informed nor powerful; in the case of control systems, the relying parties are principally the asset owners, who are generally both. A scheme that provides the appearance of security rather than the reality and yet costs real money will be unlikely to appeal to them.

However, full Common Criteria certification would be expensive, adding both substantial evaluation costs and delays into the product development cycle. In addition, the Common Criteria fail to deal satisfactorily with systems that are patched frequently, as operating systems now are; observers of the operating-system patching cycle and vulnerability scene have come to the conclusion that the Common Criteria are no more than a bureaucratic exercise whose costs far outweigh the benefits. To quote Alan Paller, the director of research at the SANS institute, "If you are asking, if the effort is worth the money, the answer is a resounding no"[20]. As control systems start to be updated regularly, these words must have some resonance. In addition, the example of locks reminds us that attacks also evolve; and the Roadmap to Secure Control Systems in the Energy Sector launched by the US Department of Energy and Homeland Security identifies the growing strength of hacker tools as a prime challenge for energy security[21].

An equally significant consideration will be usability. Ease of safe use is one of the key design requirements for control systems, and as we noted above, usability is not well dealt with; and the complex interactions between safety and security are similarly outside the Criteria's scope. There is also the structure of control systems, while the Criteria were designed for essentially stand-alone products, such as a secure telephone, control systems are large, complex and evolving. At the vendor end, there are many small components (sensors, actuators, switches) for which the only practicable assurance strategy would focus on the vendor's internal quality-control processes. There may then be a few critical components, on which third-party assurance might focus, such as servers that control whole subsystems or firewalls that connect internal networks to the Internet. Then, on the plant side, there are further complex processes that are properly the asset owner's domain; many aspects (such as managing the evolution of business processes) are almost completely outside the framework that the Criteria provide.

So much for the bad news. Now the good news: much of what goes wrong with Common Criteria certification has to do with incentive problems. Firms want to persuade customers to accept liability by claiming that systems are secure, and look to third parties to certify this; in other applications, governments want their contractors to use secure systems, and let the vendors have their offerings certified by other third parties. In the world of control systems, the incentives are not quite so badly aligned: the asset owners who purchase the control systems have real liability if a saboteur manages to cause damage, injury of loss of production[22]. A certification scheme therefore only has to tackle one asymmetric information problem: the fact that vendors know more about their systems than the asset owners that buy them. It does not have to deal with all the other hidden-information and hidden-action problems that beset both the financial-systems and the government-systems worlds.

4.1. What's needed

The above discussion should have shown how the Common Criteria are not well matched to the needs of the control systems world. At the technical level, a security certification scheme must be able to cope with dynamic systems, dynamic threats and real users working in real organisations. It must complement, rather than conflict with, existing safety certification mechanisms. But above all, its function is to provide assurance to asset owners that the systems and components they buy from the vendor community are fit for purpose.

It is perhaps unsurprising that some firms have set up independent schemes to evaluate the security of control system devices. MuDynamics provides MUSIC certification for network infrastructure protocols like ARP, DHCP, TCP/IP and the application level protocols like HTTP, DNP3, MODBUS, and SNMP. Another company, Wurldtech, runs the Achilles certification program[23] that provides a proprietary third-party control system evaluation platform for security benchmarking. It advertises to the customers to "insist on Achilles certifications—A NO cost effort to improve the security and reliability of your industrial operations"[24]. It appears that such marketing strategies are having some success. However, it's early days yet; the current schemes are limited in scope, certifying the basic protocols rather than the SCADA-specific implementation and system-level architecture.

How should the market for evaluation and certification develop?

One of the ways a certification scheme could do better is to add usability testing to its evaluation process. As already noted, the Common Criteria do not take into account the interaction of the system with the user; this can be particularly dangerous in industrial software, where ease of

safe use is a high priority.

Second, there should be a greater role for the end user in the evaluation process. This is difficult to implement in the traditional IT world since the customer is powerless and is hardly in a position to act collectively. However industrial asset owners are powerful customers with serious liability in case of a security breach and detailed systems knowledge. So the asset owner can be an important contributor.

Third, a certification scheme should take the whole product lifecycle into account. In our view, industrial systems vendors will inevitably move to a patching cycle, as IT platform vendors already have. The patching cycle is likely to be longer because of the nature of the business, but its management will be central to the level of assurance that asset owners enjoy. A certifier must therefore take into account the mechanisms that a vendor has for receiving vulnerability reports from asset owners and others, testing new releases of software, and supporting its customers in installing them.

The environment in which evaluation and certification take place is complex. The energy sector is already witnessing an interesting natural experiment in that the UK has opted for a light-touch regulation via CPNI while the US has adopted regulation in electricity through NERC and FERC. Now here is a second experiment: the development of private-sector certification in competition with the Common Criteria. Having two certification schemes does impose some costs. For example, the ABB AC800M-PM865 has attained certifications from both Achilles[25] as well as MUSIC[26]. If the number of certification companies were to rise significantly, vendors could face non-trivial costs getting their products through multiple evaluations.

Perhaps one of the certifiers will become the de facto standard. Wurldtech has set up a collaboration with Exida to deliver safe-secure certified systems[27]. It has also merged with Byres Research to form a safety-security certification company[28]. Asset owner Shell has signed a collaboration with Wurldtech according to which Shell would push for Wurldtech's certification in their vendor procurement requirements and factory acceptance tests[29].

However, what seems more important to us than the number of certification service providers is the incentive structure. The most common reason for the failure of certification programs is incorrect incentives, such as the lack of liability on any powerful party when things go wrong. However, given that asset owners carry liability, it is proper that they should drive the certification process, and quite reasonable that they should have more than one certifier to choose from.

5. Conclusion

The Common criteria have not worked particularly well, whether in their original role of certifying secure computer systems for government purchasers or in their new role of providing some assurance for products on which third parties have to rely. Imposing them on the world of industrial control systems would be unwise, as they are even less well matched to the industry's requirements.

Thankfully, a private-sector solution appears to be emerging, with two evaluation firms, and asset owners involved in certification. As asset owners currently bear most of the risk, and are also in a position to take appropriate precautions, a prudent regulator will seek to empower them rather than hobble them. Regulating the certifiers is very likely to be a bad idea, as we've seen with the Common Criteria (and elsewhere, as with the UK school exam market).

It may well be that some future applications will justify a more centralized approach. One we've heard mentioned is smart metering; misbehaviour of 100,000 domestic meters could cause power surges of the 300MW threshold of NERC, for example; certainly some thought should be given to security and resilience before such systems are deployed. However, as far as existing systems are concerned, the case for centralized regulation of security is clearly not proven.

6. Acknowledgement

The second author's research is funded by ABB. The contents of this article do not necessarily express the views of ABB.

References

[1] "Risk", John Adams, 1995.

[2] James Anderson, Computer Security Technology Planning Study, ESD-TR-73-51, U.S. Air Force Electronic Systems Division (1973), http://csrc.nist/gov/publications/history

[3] Department of Defense, "Department of Defense Trusted Computer System Evaluation Criteria," DoD 5200.28-STD (Dec 1985).

[4] National Institute of Standards and Technology, "Common Criteria for Information Technology Security," Version 2.0/ISO IS 15408 (May 1998), http://www.commoncriteria.org.

[5] Why the Security Market has Not Worked Well, chapter 6 of Computers at Risk: Safe Computing in the Internet Age, National Academies Press, 1991

[6] Marc Weber Tobias, Tobias Bluzmanis, Locks, Safes and Security, from www.security.org; see also Ross Anderson, below [13]

[7] Ross Anderson, Comments on the Security Targets for the Icelandic Health Database, 2001, at www.ross-anderson.com

[8] Common Criteria Certified Product List, http://www.commoncriteriaportal.org/products_OS.html#OS

[9] Unbreakable: Oracle's commitment to security, http://www.oracle.com/technology/deploy/security/pdf/unbreak3.pdf

[10] Microsoft Security Bulletin MS08-067 – Critical, http://www.microsoft.com/technet/security/Bulletin/MS08-067.mspx

[11] Ross Anderson, Mike Bond, Jolyon Clulow and Sergei Skorobogatov, "Cryptographic processors – a survey", Cambridge University Computer Laboratory Technical Report no. 641 (July 2005), shortened version in Proc. IEEE}v 94 no 2 (Feb 2006) pp 357–369

[12] Alexandra Frean, "A-level reputation in severe decline ... now even an exam board chief doubts their value", The Times, Nov 10 2007

[13] Ross Anderson, Security Engineering – A Guide to Building Dependable Distributed Systems, chapter 26, Second edition, Wiley 2008

[14] Ben Edelman, Adverse Selection in Online 'Trust' Certifications, Fifth Workshop in the Economics of Information Security, 2006

[15] Anderson, op. cit – 13 above

[16] Ross Anderson, Saar Drimer and Steven Murdoch, "Thinking inside the box: system-level failures of tamper proofing", Computer Lab Technical Report UCAM-CL-TR-711

[17] Sandia SCADA Program High-Security SCADA LDRD final report. http://www.sandia.gov/scada/documents/020729.pdf

[18] Process Control Security Forum (PCSRF), SCADA Protection Profiles http://www.isd.mel.nist.gov/projects/processcontrol/members/documents.html

[19] White Papers and Articles, digital bond. http://www.digitalbond.com/index.php/resources/white-papers-and-articles/

[20] Common Criteria Under Attack, http://gcn.com/Articles/2007/08/10/Under-attack.aspx?Page=2

[21] Roadmap to Secure Control Systems in the Energy Sector, Department of Energy, Department of Homeland Security, January 2008, http://www.controlsystemsroadmap.net/

[22] Ross Anderson, Shailendra Fuloria, Security Economics and Critical National Infrastructure, 2009 Workshop on the Economics of Information Security

[23] The Achilles Certification Program, http://www.wurldtech.com/certification/index.php

[24] Are your Control Systems Achilles Certified?, http://www.wurldtech.com/library/pdf/Insist_On_Achilles_Certification_(Electric).pdf

[25] Achilles Certification Program, http://www.wurldtech.com/certification/devices/800M.php

[26] MUSIC-Certification, http://www.mudynamics.com/assets/docs/ABB_Certification.pdf

[27] Wurldtech and Exida announce Global partnership, http://www.wurldtech.com/news/archives/170309.php

[28] Byers Research and Exida merge, http://exida.com/news.asp?ID=55

[29] Wulrdtech and Shell announce Global cooperation for increased infrastructure protection, http://www.wurldtech.com/news/archives/310309.php