# Input to the Commission on Enhancing National Cybersecurity

**Submission date:** September 6, 2016

**Joint submission made by:**

**Benjamin Gittins**
Chief Technical Officer
b.gittins@synaptic-labs.com
+356 9944 9390

**Ronald Kelson**
Chief Executive Officer
r.kelson@synaptic-labs.com
+356 9944 9390

**Synaptic Laboratories Ltd.**
www.synaptic-labs.com
13 Nadur Heights,
Nadur NDR-1390,
MALTA, Europe

*Designers of safe and secure computing and communication architectures. Developers of general-purpose soft IP for FPGA devices, to increase security and performance, and to reduce circuit area.*

**Topic of this submission:**

**Today's Internet Ecosystem Has Toxic Conceptual Design Flaws That Undermine Its Security —**
**Change The Game By Designing and Incrementally Deploying An Internet Ecosystem That Is Universally Trustworthy and Dependable**

**RFI topic areas this submission relates to:**

- Cybersecurity Research and Development
- Cybersecurity Insurance
- Critical Infrastructure Cybersecurity
- Identity and Access Management
- Internet of Things
- International Markets

**Input submission contents:**

(1) A 1 page executive summary for this comment, in the format requested by the RFI, which "identifies the topic addressed, the challenges, and the proposed solution, recommendation, and/or finding." Citations in the Executive Summary map back to the references listed at the end of the 16 page security article attached to this submission. We have inserted headings that match these points in the executive summary.

(2) A 16 page security publication. B. Gittins and R. Kelson. "Verifying Secure Systems is also Not Reasonable (Today)". An Invited Presentation to the Eighth IBM Haifa Verification Conference. Nov. 2012. Full text subsequently published online on the IBM website. (www.research.ibm.com/haifa/conferences/hvc2012/papers/Security_Gittins.pdf)

(3) Brian Snow. We need assurance! In ACSAC '05: Proceedings of the 21st Annual Computer Security Applications Conference, pages 3–10, Washington, DC, USA, Dec. 2005. IEEE Computer Society. Full text published online on the ACASC website (https://www.acsac.org/2005/papers/Snow.pdf).

The "only way" to address the looming cybersecurity crisis is "to build more trustworthy secure components and systems," .. "What's needed is a new approach of building computing systems that are designed from scratch to be fundamentally trustworthy" — Ron Ross (NIST Senior Fellow) addressing the president's commission on long-term cybersecurity, August 23, 2016.

# Today's Internet Ecosystem Has Toxic Conceptual Design Flaws That Undermine Its Security — Change The Game By Designing and Incrementally Deploying An Internet Ecosystem That Is Universally Trustworthy and Dependable

## 1 Page Executive Summary:

***RFI Topics:*** *Cybersecurity Research and Development, Cybersecurity Insurance, Critical Infrastructure Cybersecurity, Identity and Access Management, Internet of Things, International Markets.*

***Problem:*** The US [4] and UK [69] Governments assert their respective nations are at strategic risk of failure due to security problems plaguing the ICT ecosystem. World-leading cyber security experts claim this is because of serious conceptual design flaws throughout our ICT foundations [64], [40], [13]. In particular: 1) Flawed security assumptions lead to incorrect requirements definition, conceptually flawed designs and the inevitability of security failure [45]. 2) Most secure systems use (defacto) standards based ciphers, protocols, and COTS hardware, that collectively suffer from these problems and have no formal security model [38]. 3) Conceptual flaws exist in the way safety and security systems are formally modelled. Temporal properties [43] and human trust factors are often ignored [64]. Brian Snow (former US NSA technical director IAD) warns: "Today's Trust Bubble [ed. ICT] products are rife with a huge pile of crippling unaddressed conceptual and implementation debt. ... we are ripe for a Trust Bubble melt-down with the same scale of consequences that the Credit Markets suffered." [64] M. Hathaway laments: "We have not designed systems for failure for over 40 years." ... "We are not designing and investing into an infrastructure ... that could succeed through a major disaster." … "I would argue that we need to be thinking about designing for a more secure and resilient architecture." [40] B. Snow observes: "the Security Industry has yet to fully internalize how much CHANGE is required in the DESIGN environment given that MALICE rather than benign failure is the major driver for their products … The creators of the Internet knew that MALICE was a serious issue. ... However, the creators of the Internet pushed security aside due to the perceived difficulties, or cost, and that is the start of our problems today. To put it bluntly, the Internet was not built to address the known risks [16]. By design, the Internet naïvely relies on the honesty of every network user, and places far too little emphasis on healthy mutual suspicion! The cost and risks were not eliminated – rather they were both shifted away from the designers and the manufacturers, and transferred to the Global user base." [64] The U.S. National Cyberspace Policy Review states: "An advisory group for [DARPA] describes defense of current Internet Protocol-based networks as a losing proposition." [4]

***The recommendation:*** We respectfully propose that the Commission's detailed recommendations to strengthen cybersecurity should include the following points (that we argue in greater detail in the Peer-Reviewed Technical Publication attached to this submission):

1. Perform a high-level survey to identify and catalogue the conceptual design flaws that undermine the security of all common Internet protocols, including: HTTP, email, X.509, SSL/TLS, S/MIME, DNS. (Note: this is a well researched subject, so not an impossibly scoped or extremely expensive ask) This includes performing a high-level security aware Failure Mode and Effects Analysis that considers the impact of those flaws wrt. the stakeholders in multi-jurisdiction, multi-stakeholder Internet-scale environments.
2. Building on the results of the above analysis, Survey and Define the core elements in a "trustworthy and dependable communication and computation vision that seeks to protect the legitimate interests of all stakeholders in multi-jurisdiction, multi-stakeholder Internet-scale environments". Propose a grand-design strategy for viably achieving that vision that includes: (a) Aiming for end-to-end trustworthiness and dependability within systems of systems; (b) Addressing the human trust issues to protect the stakeholders; (c) Decentralising power across stakeholders in a fault tolerant way that builds trust between distrusting and potentially competing participants; (d) Eliminating and/or controlling all hazards at their source, rather than merely reduce its severity; (e) Designing clean-slate building blocks/foundations that can be rapidly deployed to protect the existing infrastructure with a minimum of disruption while enable future capabilities / systems to be built on them; (f) Employing high assurance development methods and target high certification levels where these are appropriate; (g) Ensuring that traditionally low-assurance systems can be viably enhanced with or built using high assurance components/protocols/foundations, thereby raising the base-line level of security for those low-criticality systems; (h) In high criticality systems, systems with a very large number of stakeholders, and for COTS IP that is employed by many organisations, using formal methods, high-assurance development methodologies and independent audits to help prevent against insider attacks in the design; (i) Build (a base-line level of) emission security in from the onset [54]; and (j) Defining a minimum base-line of security that all products and systems must employ, as many are aware of risks but deliberately ignore them (this being also important to achieve in low risk products to ensure they are not a weak link access point to subvert higher risk systems they are informing or connected to).
3. Perform an in-depth survey to identify, catalogue and evaluate the viability of all candidate next-generation Internet protocols that are credibly trustworthy and dependable in multi-jurisdiction, multi-stakeholder Internet-scale environments that can be incrementally deployed to protect existing security systems while permitting the transition to higher levels of security assurance and improved capabilities.

Please read Synaptic Labs' proposals for a universally trustworthy and dependable Internet ecosystem in section 5 of "Verifying Secure Systems is also Not Reasonable (Today)" attached to this submission.

Sincerely, Benjamin Gittins and Ron Kelson.

# Verifying Secure Systems is also Not Reasonable (Today)

Benjamin Gittins (CTO), Ronald Keslon (CEO)

[1] Synaptic Laboratories Limited
**b.gittins@synaptic-labs.com, r.kelson@synaptic-labs.com, http://synaptic-labs.com**

[2] ICT Gozo Malta Project
**http://ictgozomalta.eu/vision-and-projects**

**Abstract.** Many problems undermine the formal verification of security in (non-trivial) ICT systems: 1) Flawed security assumptions lead to incorrect requirements definition, conceptually flawed designs and the inevitability of security failure [45]. 2) Most secure systems use (defacto) standards based ciphers, protocols, and COTS hardware, that collectively suffer from these problems and have no formal model [38]. 3) Conceptual flaws exist in the way safety and security systems are formally modelled. Temporal properties [43] and human trust factors are often ignored [64]. This paper collates expert assessments of the current global ICT security status and presents the ICT Gozo Malta Project Technology Roadmap (see figure 1), developed by Synaptic Laboratories Limited. This Roadmap offers a grand collaborative clean-slate ICT vision designed to address many known security problems, to viably bolster existing ICT systems, and to enable more verifiably secure, trustworthy and dependable, systems of systems in practice.

## 1 Executive summary

The US [4] and UK [69] Governments assert their respective **nations are at strategic risk of failure** due to security problems plaguing the ICT ecosystem. World-leading cyber security experts claim this is because of serious conceptual design flaws throughout our ICT foundations [64], [40], [13]. According to Brian Snow (former US NSA IAD): *"We must change our toxic environment."* [64] Over the past ≈ 12 years Synaptic Labs has been systematically addressing the conceptual, functional and security flaws in today's ICT ecosystem, including: global-scale networking, global-scale cryptographic key and identity management, and secure computing. Many of our conceptual cross-domain designs have been independently, positively, peer reviewed by world-leading companies and experts; some have also been openly published. The Roadmap begins by converging high-assurance safety and security requirements in universal computing designs to create dependable platforms that seek to protect the legitimate interests of all stakeholders, globally. We can *change the game* by realising this Roadmap, using high assurance formal methods from the onset, to enable applications built on them to be formally verifiable down to the processor core level. We invite the formal methods community to join this collaboration with other world-leading ICT organizations and domain experts, to realise verifiably secure  trustworthy and dependable  systems of systems in practice.

## 2 Structure of this paper

This paper collates assessments by world-class domain experts' on our global cyber safety and security status §3 and their views on the condition of our ICT pillars §3.3. §4 outlines the ICT Gozo Malta / Synaptic Labs' Technology Development Roadmap (fig. 1), based on ≈ 12 years cross-domain research and design, to realise a universally trustworthy and dependable ICT ecosystem that can be formally verified. §5 briefly surveys design strategies for success. §6 outlines the application of our Roadmap design strategies in each of the ICT pillars. In §7 we invite you to join the revolution and help realize a globally inclusive, universally trustworthy and dependable, ICT ecosystem.

## 3 Cyber safety and security assessment

### 3.1 Experts claim our cyber foundations are fundamentally flawed

Our videos and publications provide a wide summary on this subject [37], [38]. Examples include: in 2011 Brian Snow (35 years, U.S. National Security Agency NSA, incl. 12 years as Technical Director of Information Assurance Directorate IAD), asserted: *"There are problems today in cyber security practice that impact the community as a whole, and we need to solve those problems soon. They are pervasive, ongoing, and getting worse, not better."* ... *"the community at large is applying the wrong or inadequate engineering practices, and taking a lot of short cuts. ... your cyber systems continue to function and serve you NOT due to the EXPERTISE of your security staff, but solely due to the SUFFERANCE of your opponents."* [64] The Director of U.S. National Intelligence testified (2010) that the public and private information infrastructure was 'threatened'. Melissa Hathaway (leader of the U.S. National Cyberspace Policy Review [4]) added: "***And I would say that it is compromised.***" [40]. *"I think it is unconscionable that our leaders are not talking about what is really happening. Some of it is because of the fear that we are going to lose trust in the core infrastructure and/or that we are going to lose public confidence."* [40] Debora Plunkett, Director of the U.S. NSA IAD, stated: *"we are not at all overstating the threat."* [13]

### 3.2 Why the flawed cyber foundations are a Trust Bubble

**3.2.1 Conceptual design flaws throughout the ICT ecosystem:** B. Snow warns: *"Today's Trust Bubble [ed. ICT] products are rife with a huge pile of crippling un-addressed* ***conceptual*** *and implementation debt. ... we are ripe for a Trust Bubble melt-down with the same scale of consequences that the Credit Markets suffered."* [64] M. Hathaway laments: *"We have not designed systems for failure for over 40 years."* ... *"We are not designing and investing into an infrastructure ... that could succeed through a major disaster."* [40]

**3.2.2 Flaws in the approach to ICT design:** We agree with M. Hathaway: *"I would argue that we need to be thinking about designing for a more secure and resilient architecture."* [40] B. Snow observes: *"The security professional faces an environment that adaptively and rapidly changes to nullify his efforts ... He must accept that standard design practices simply are not adequate in a malicious environment! ... the Security Industry has yet to fully internalize how much CHANGE is required in the* ***DESIGN*** *environment given that MALICE rather than benign failure is the major driver for their products."* [64]

**3.2.3 Global Risks Report:** Critical systems failure was identified by the World Economic Forum's Global Risks Report [10] as *"a key concern for world leaders from government, business and civil society"* and that this will *"most likely be caused by cyber attacks"*; currently ranked 4th out of 50 global risks.

**3.2.4 The bottom line is trustworthiness:** Jeannette Wing, U.S. National Science Foundation, states: *"We need to be able to trust our systems, digital and physical, because after all what protects our physical is often digital now."* [52]

### 3.3 Assessments of ICT pillars

**3.3.1 The state of ICT hardware:** B. Snow states: *"For a one-word synopsis of computer design philosophy, it was and is: SHARING. In the security realm, the one word synopsis is SEPARATION. .... So today, making a computer secure requires imposing a "separation paradigm" on top of an architecture built to share. That is tough! Even when partially successful, the residual problem is going to be covert channels."* [63] Real-time experts [49] state: *"In safety critical and mission critical systems ... it is important to assign applications with different requirements to different partitions with different criticality levels ... Partitions should be isolated functionally, temporally and securely ... Unfortunately, modern COTS architectures* ***are not*** *built to provide strong isolation guarantees."* From a safety perspective, in 2012 Airbus' Benoît Triquet [68] stated multicore processors represent *"a major challenge how to adequately deploy them for safety applications they were typically not designed specifically for. ... Temporal behaviour has been much less addressed ... Airbus ... have found very few multicore chips that can ever hope to be useable for avionics."* From a security perspective, B. Snow argues: *"And so it makes a lot of fun, they have good cryptography, they have little computers and chips, and they are radiating [compromising emissions] like swine."* [66] The ICT GM / Synaptic Labs Technology Roadmap seeks to specifically address these issues in a manner that makes safety and security viable in universal computing hardware.

**3.3.2 The state of ICT operating systems:** The paper titled "The Inevitability of Failure: The Flawed Assumption of Security in Modern Computing Environments" states [45]: *"Current security efforts suffer from the flawed assumption that adequate security can be provided in applications."* In 2005, B.

Snow goes wider and deeper: *"Given today's common hardware and software architectural paradigms, operating systems security ... is the current 'black hole' of security."* Today, B. Snow states: *"Consider the use of high-assurance ... operating systems ... as a way to reduce the attack surface of your critical systems, and to isolate one component from another. ... they can provide considerable gains in security and functionality for systems needing high-assurance or high-integrity or high-performance."* See: [32], [3]. The safety and security RTOS vendors are collaborating with us because they need much better hardware support.

**3.3.3 The state of ICT clouds:** CEBR [6] cautions that the full shift to cloud computing may not happen if **perceptions** in relation to security and resilience-related aspects of cloud computing solutions deteriorate [6]. According to an ENISA report, administrator roles in today's cloud architectures expose cloud customers to extremely high risk [23]. ENISA says these insider attacks have a Medium probability of occurrence and will have a Very High negative impact on stakeholders [23]. In a 2010 public cloud privacy breach, clients had to notify Google that insider attacks were defacing their accounts for months before Google took corrective action. [42] Dr. Howard Shrobe, Program Manager for the DARPA I2O Mission Orientated Resilient Clouds project argues *"Clouds are concentrated Vulnerability Amplifiers"* because they are monocultured, have huge concentration of hosts on high speed network without internal checks, have implicit trust among hosts, they have resource sharing and co-residence of unrelated computations, are an obvious target, are vulnerable to activity monitoring and other types of side-channel attack vulnerabilities [8].

**3.3.4 The state of the Internet protocol/deployment:** B. Snow states: *"The creators of the Internet knew that MALICE was a serious issue." ... "However, the creators of the Internet pushed security aside due to the perceived difficulties, or cost,* **and that is the start of our problems today**. *To put it bluntly, the Internet was not built to address the known risks [16]. By design, the Internet naïvely relies on the honesty of every network user, and places far too little emphasis on healthy mutual suspicion! The cost and risks were not eliminated – rather they were both shifted away from the designers and the manufacturers, and transferred to the Global user base."* [64] To quote Vice Admiral J. Mike McConnell (USN Ret): *"The Internet has introduced a level of vulnerability that is unprecedented ... The nation is at strategic risk."*

The U.S. National Cyberspace Policy Review states: *"An advisory group for [DARPA] describes defense of current Internet Protocol-based networks as a* ***losing proposition****."* [4] Vint Cerf says: *"A new version of the Internet might be the best way to defend against cyber attacks."* [48]

**3.3.5 Conflicts of interest between cyber offence and defense:** Some governments seem to be to determined to exploit these strategic vulnerabilities rather than seek to deploy trustworthy ICT ecosystems. Prof. Ross Anderson argues that there is a fundamental conflict of interest inherent in the UK policy. In the USA, DARPA's global-scale cyber offensive initiative ***"Plan X"*** will *"support development of fundamental strategies and tactics needed to dominate the cyber battlespace."* [11] Effective cyber offense requires collective weakness.

**3.3.6 The state of the civilian identity management federation:** There are serious design and implementation flaws [39], [46], [47] that have plagued the civilian global-scale public key infrastructure (PKI) and **fundamentally undermine** its utility [36]. The following two citations provide an indication of the level of expert dissatisfaction: Richard R. Brooks' paper: "Liars and the lying liars that tell them" and Peter Gutmann's book "Engineering Security" [39] section titled: *"SSL certificates: Indistinguishable from Placebo."* According to Landon Noll, Cryptologist and Security Architect at CISCO: *"PKI ... In practice is it snake oil? It is somewhat indistinguishable in practice because of the problems.*" [36] Andrew McLaughlin, White House Deputy CTO of Internet Policy states: *"Fake secure websites ... are a danger the government is* **powerless** *to control.*" B. Snow states: *"Cyber trust, as implemented today, does not map to the way humans naturally reason about trust.*" ... *"The issuing of identity assertions is uncoordinated among many different certificate authorities, none of whom I have a personal relationship with. This means there are many system nodes that can make false assertions that would be accepted as truth within the global system.*" [64] Elaine Barker, project leader of the NIST global-scale Cryptographic Key Management (CKM) project [17] states on p. 31 and p. 52 of [18]: CKM designers *"must look at means other than public key-based key management schemes; they must look at quantum computing-resistant algorithms and schemes.*" Note: Today's public key algorithms catastrophically fail due to derivatives [21] of Shor's algorithm [19], no trusted alternative available.

### 3.4 Severe risk of global strategic failure

The U.S. National Cyberspace Policy Review states: *"[Security] Threats to cyberspace pose one of the most serious economic and national security challenges of the 21st Century for the United States and our allies.*" [4] The 2011 EU Commission funded FP7 RISEPTIS Report says: *"The trustworthiness of our increasingly digitised world is at stake.*" [58] The 2011 UK Cyber Security Strategy states: *"Any reduction in trust towards online communications can now cause serious economic and social harm to the UK.*" [69] Also see: §3.2, [37], [38].

## 4 The ICT GM / Synaptic Labs cyber design strategy

### 4.1 Statement of goal

World-leading experts [15], [52], [63], [40], and some Governments [4], [9], [69], [58], are calling for trustworthy and dependable global-scale ICT systems. The authors argue that such systems must be designed to protect the needs and legitimate interests of **all stakeholders** [2] with regard to services provided. They must be acceptable to mutually suspicious entities, irrespective of their relative power relationships, and not rely on (violent) sanctions to build acceptance.

## 4.2 The ultimate project for the formal methods

Today, literally billions of people rely on low-assurance technologies such as PKI X.509, the Internet and COTS computing hardware developed using low-assurance techniques. It is time to employ formal methods to realize trustworthy and dependable ICT foundations that can be relied on by the global community.

## 4.3 A grand design strategy for achieving verifiable security

**4.3.1 Aim for end-to-end trustworthiness and dependability within systems of systems:** In 2008 the UK Government's Technology Strategy Board (TSB) website stated: *"The current way which organisations approach security can be recognised as an underlying **market failure** which consists of fire fighting security problems, silo'd implementation of technologies, uncontrolled application development practices and a failure to address **systemic** problems. Organisations tend to deal with one problem at a time that results in the deployment of point solutions to treat singular problems."* TSB observe: *"Business now relies on information infrastructures that are interlinked and interdependent."* We must design cross-cutting safe and secure global-scale multi-stakeholder systems.

**4.3.2 Address the human trust issues – protect the stakeholders:** M. Hathaway states: *"I don't trust hardly any transaction right now, there is no integrity in our infrastructure."* [40]. To quote Nicholas C. Rueter's cyber warfare political thesis: *"The international system has a number of features that make cooperation difficult. Most important is the prevalence of uncertainty and mistrust. ... While many states are satisfied with their place in the international hierarchy and seek only to protect their position, some states endeavor to enhance their security by dominating others, apparently subscribing to the theory that 'the best defense is a good offense.' Because the system is anarchic (i.e., there is no common or overarching world government), states must provide for their own security needs."* [59] Global-scale ICT systems, such as the X.509 PKI ecosystem §3.3 and the Internet §3.3, are cooperatively governed international systems that are currently entrusted (and failing) to protect the legitimate interests of billions of people. We propose to move beyond the *anarchic* "Law of Nations" [28] by adopting fault-tolerant civil political techniques in combination with safety and high assurance security techniques.

Safety engineers design ICT systems to standards (e.g. IEC 61508 [5]) to avoid "single points of failure" that could compromise the safety of the equipment and stakeholders. High assurance security engineers employ fault tolerant techniques (e.g. NSA SKPP [3]) for ensuring confidentiality under faults. Political scientists design governance systems to avoid "single points of *trust, authority* failure" (such as tyrants and dictators) that could compromise the safety or security of the community. We argue that the ICT safety and security communities need to collaborate with political scientists to combine the spirit of IEC-61508 [5] with the spirit of the laws that underpin modern civil governance systems [50]. In addition to ICT's objectives of availability, reliability, safety, confidentiality, integrity, maintainability [14], audit-ability, non-repudiation and/or (pseudo)anonymity, we must also do more to address the human trust issues. To reword Montesquieu in 1748 [50]: *"Government (and ICT systems) should be set up so that no person has a reason to be afraid of another person."* We need to embody more democratic good governance principles into ICT systems.

**4.3.3 Decentralise power across stakeholders in a fault tolerant way:** We need to move beyond binary and semantic interoperability [60] and loosely co-ordinated federations of service providers in which each service provider acts in a predominantly unilateral way without consultation or the oversight of other service providers. Similar to democratic systems that seek to check the arbitrary will and caprice of dictators or aristocrats, ICT systems can decentralise power and be stronger when multiple (semi-)autonomous mutually suspicious entities (netizens) are involved in transactions in a way that is designed [50] to protect the legitimate interests of all stakeholders. For an example of how to do that at the client-server transaction level see [36], [33], [34]. When seeking trustworthiness and dependability employ decentralization of power and formal methods.

**4.3.4 Aim to completely eliminate problems:** It is much simpler to argue the safety or security properties of a system when you eliminate a hazard at it's source, rather than merely reduce its severity. This requires systematically surveying and solving problems in a recursive fashion *across domains*. Our Roadmap employs cross-domain visibility and expertise to: viably eliminate problems **at the source**, to eliminate redundancy and reduce the complexity of the architecture across domains, and to optimise the universality of application of each module. This enables solutions that will be simpler and cheaper to (formally) analyze for correctness, understand, maintain and use.

**4.3.5 Protect what is deployed today and enable future capabilities:** Clean-slate cross-domain thinking can find both short and longer term solutions to today's hard open problems. Minor changes to existing hardware or software can deliver significant safety, security and performance gains with modest changes to existing third party intellectual property. When clean-slate foundations are absolutely required, we aim to achieve revolutionary capabilities that can be applied to bolster as much of the existing infrastructure as possible.

**4.3.6 Employ high assurance development methods and target high certification levels:** After the hard open *design* problems are addressed and a conceptual architecture is in place, begin to employ high assurance development methods and target high levels of assurance in safety and security certification.

### 4.3.7 Use formal methods to help prevent against insider attacks:

Formal methods reduce evaluation costs when several/many organizations must review design requirements, specifications or implementations to establish their level of confidence. The better defined and analyzed the system, and the more easily independent entities can study these designs, then the less opportunities there are for insider attacks at design, specification or implementation time.

### 4.3.8 Build emission security in:

To reword a quote from NATO's Anders Fogh Rasmussen's [51]: There simply can be no true cyber security without emissions security. Address emission security [1] from the *very* onset [54].

## 5 Design for success!

B. Snow proclaims: *"He who gets to the interface first, wins!"* [64] The semiconductor industry now "designs for testability", the safety industry "designs for safety" and the U.S. Department of Homeland Security is urging us all to "build security in." We must also design for trustworthiness and dependability [15], [14] between mutually suspicious stakeholders [50], design for mutual accountability and audit [12], design user-centric systems that empower all stakeholders of a system [2], [31], design privacy aware security [29], design digital immune systems that employ decentralised layered security [56], design for survivability under targeted malice [64], design for determinism [61], design for ACET and WCET predictability, design for real-time agility, and in particular design for modeling and formal methods. Addressing one goal makes solving the next goal easier.

**Fig. 1** ICT Gozo Malta / Synaptic Labs' Technology Development Roadmap

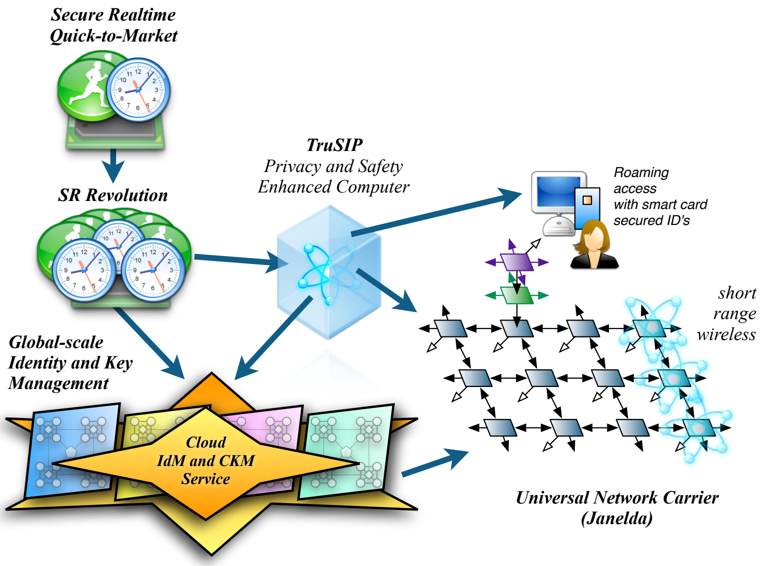

## 6 Synaptic Labs' 12 year cyber campaign: design strategies and progress

**6.1 Prehistory:** Synaptic Labs' CTO was the lead designer and co-implementor of a comprehensive cross platform, cross-vendor, object orientated telephony framework that could, among other things, passively decode and monitor Signaling System 7 ISDN User Part (ISUP). After the framework was deployed on live international traffic, attention shifted to pure research into clean-slate secure user-centric globally-decentralized parallel computing architectures employing (potentially high latency) transaction based memory architectures; leading to the following projects... *Also see:* -

**6.2 Janelda - global-scale universal network carrier:** Synaptic Labs' goal was and is to realize a secure, real-time, universal network carrier. Originally conceived to provide point-to-point and point-to-multipoint communications, scaling transparently from processor-bus interconnects through to a mesh network with billions of router nodes. It is designed to support overlapping spheres of influence (security/ownership domains) and scale up to 1 terabit/s *flows* with up to 1 second round trip latencies. Explicitly designed to achieve lossless packet routing, congestion management and authenticated link-level encryption on one standard ASIC chip. We began by first surveying and solving core scalability and performance problems in the Internet Protocol, particularly with regard to **cost effective** wide-area network routing and congestion management. We explored how to manage the interoperability requirements to securely host all existing wide-area network isochronous, cell and packet based protocols without requiring changes (e.g. encoding or transcoding protocols) in a variety of operational contexts, such as: transporting medical and legally privileged data (50-to-100 year security), industrial control traffic (low-jitter, zero packet loss), Internet of things (lower power, bandwidth constrained, denial of service resistance), peer-to-peer networks, web surfing, carrier grade telephony and video streaming, and supporting both audited and anonymous traffic flows directly in the infrastructure. Having solved most of the global-scale routing and packet congestion "network" issues at the conceptual level (includes adapting known techniques in new ways), we shifted our attention to information security, particularly with regard to 100 year secure 10 gigabit/s link- and packet-level authenticated encryption in hardware [53], post quantum secure key exchange technologies, and managing name spaces within the network that would be resistant to spoofing attacks.

**6.3 50-to-100 year security:** Extensive study was made of over 250 papers relating to code-breaking quantum computing and long-term security: including classical (a)symmetric cryptography, candidate post quantum secure crypto, and information-theoretically secure primitives. We argue that the only cryptographic primitives the community can rely on today for long term security are

NIST-style block-ciphers, hash functions and constructs based on those primitives. We then set out to survey and address the scalability and security requirements for building key negotiation protocols and Merkle-tree style digital signatures [26], including the design [34] of fault tolerant information theoretically secure symmetric key exchanges. Permits competing national cipher standards to be simultaneously employed in one client transaction.

**6.4 Global-scale identity management (IdM) and cryptographic key management (CKM):** Starting with traditional key distribution/translation center technologies and all-or-nothing transformations as a base, and with our global-scale multi-jurisdiction multi-stakeholder objectives in mind, our team independently re-discovered a fault tolerant symmetric key negotiation protocol sketched in [30]. Our protocol employed modern smart cards and featured a more complex human-trust model. In 2008 we identified how to arbitrarily scale the protocol to support billions of enrolled devices while continuing to address the human-trust issues as discussed in §4.3, §4.3, §4.3, [36]. This was independently reviewed, and well received, by world class experts in post quantum security (J. Patarin and L. Goubin). Our proposal [33], [34], [36] employs a decentralised trust model that exploits compartmentalisation, redundancy and diversification simultaneously across service provider, software developer, hardware vendor, class of cryptographic primitive, and protocol axis. It supports the collaborative management of international name spaces, management of client transactions using public identifiers, enterprise CKM, and supports user/stakeholder-centric cross-cutting control mechanisms. This proposal is suitable for use with commercial off the shelf hardware and is **designed to bolster** the security of **existing** security deployments. [35] We then set out to design a trustworthy and dependable hardware security module. In 2010 we submitted 157 pages of input to NIST's global-scale CKM SP800-13 [35].

**6.5 Semiconductor emissions:** ***Our request to the EDA community*** is that the chip development suites add native support for dual-rail charge recovery logic technologies [62]. Please take into consideration the influence of manufacture variability [57] on security [54], [67] and employ formal methods to validate correctness of implementation [20] with experts in side-channel attacks.

**6.6 Trustworthy Resilient Universal Secure Infrastructure Platform (TruSIP):** TruSIP targets safety and security first and was originally optimized for running existing applications on general purpose operating systems under a hypervisor. It maintains uniform levels of confidentiality, integrity and availability under exploitation of latent vulnerabilities or malware within any software/hardware module of the multi-core computing platform (including kill switches). Designed to prevent anybody (the service provider's management and techies, and the privileged persons involved in the design, implementation or maintenance of any of the software or hardware modules used by the service provider) from gaining enough information to compromise a client's 160-bit symmetric key; making it ideal as a platform for infrastructure as a service public cloud computing. This required particular attention to emission security and separation/non-interference [63] of tasks, requiring all hardware-based covert timing channels [22] and timing channels [44] to be adequately controlled or eliminated. TruSIP is designed to be a client and host for our global-scale IdM and CKM proposal. TruSIP has gone through 2 revisions, and been studied by world leading safety, security and survivability experts such as Brian Snow, Miles Smid, Richard R. Brooks, Frederick Sheldon, Axel W. Krings. B. Snow says: *"Synaptic Laboratories has a sound design process; this design approach and TruSIP need to be championed and moved forward to actual products."* [65]

**6.7 Secure Real-time Revolution (SR*Revolution*):** DARPA is calling for the creation of new, low-power, secure processor architectures for use in high performance embedded computers [41] and in next generation supercomputers [7]. Synaptic Labs' SR*Revolution* platform, is designed to provide an exa-scale class many-core clock-cycle deterministic real-time platform that delivers strict non-interference properties, task agility, and WCET analyzability from the onset. TruSIP's fault-tolerance and higher assurance security properties will bolt on to SR*Revolution*.

Synaptic Labs began by adapting the original TruSIP design to include nested preemption support, leading to an innovative memory subsystem optimized for average case execution time (ACET) tasks. We then began to reach out to collaborate with all leading RTOS, WCET tool vendors and many real-time experts to identify requirements and existing technologies that could be integrated into our project. We have also begun collaborating with existing CPU vendors to ensure out proposals can be adapted in their next generation of products. Having learnt that achieving determinism in server-grade processors was insufficient for worst case execution time (WCET) analyzability, we set out to employ a heterogeneous multi-core architecture employing sever class cores, mainstream embedded processor cores, and the extremely power efficient and time-deterministic Precision Timed (PRET) machines [22], [44] running the same user-land instruction subset. In particular ensuring a single real-time operating system instance could run tasks on all cores in a cache coherent memory subsystem. To address security and performance needs, we will exploit 2.5D IC (silicon circuit board), true 3-D IC technologies (e.g. Tezzaron), in combination with low-emission dual-rail charge recovery logic (e.g. Cyclos Semiconductor [62]) to achieve extremely high-performance, single chip solutions.

We surveyed the real-time literature extensively [55], [73], [72], [49], [27] to identify real-time requirements that must be met. Particular care was given to intra- and inter-core inter-task communications [71], and semaphores. Working with the community, we are explicitly targeting support for all safety and/or security certified real-time operating systems from the onset (such as INTEGRITY and VxWorks) as well as strategically important RTOS (such as T-Kernel and RTEMS). In particular our goal is to ensure all existing RTOS functionality is supported for existing real-time applications. We will propose incremental adjustments to the operating system abstraction that are better suited for many-core systems. Our designs will support all WCET tool vendors AbsInt, Rapita Systems and Tidorum, including per-task optimization of the memory subsystem for different WCET design and analysis practices (such as FP7 PROARTIS [25] and parMERASA [70]). Our goal is also to maximize performance for existing high assurance real-time programming languages such as Ada and formal methods such as B, Z, and Esterel. To further support formal methods, our goal is to be able to provide full formal models of the PRET style cores, and the deterministic memory and messaging fabric, permitting application of formal methods from software all the way down into the silicon. When we move from conceptual design to formal specifications we will work with our collaborators to begin to refine designs to also meet the most demanding safety [5], security [3] certification standards and requirements, including in aerospace, industrial control, smart grid, and automotive domains. We also aim to support various U.S. NIST security control standards. We are globally optimising all our designs.

**6.8 Secure Real-time Quick to Market computing platform:** Synaptic Labs has recently proposed a quick-to-market solution that improves the real-time performance of the European Space Agency's quad-core Next Generation Microprocessor (NGMP). A report [24] identified that resource contention could lead up to 20x slower WCET for a task on NGMP. The designs appear to to be universal (all mainstream instruction sets) and have been independently, positively, reviewed by world-leading real-time and related domain experts including in global companies. The next step is prototyping and benchmarking.

## 7 Capacity building - Join the revolution

The above text describes key points of the grand strategy being employed within the ICT Gozo Malta project focussed on Synaptic Labs' trustworthy and dependable communication and computation vision that seeks to protect the legitimate interests of all stakeholders in multi-jurisdiction, multi-stakeholder Internet-scale environments. We have outlined various strategies §4 that have been employed, including recursively surveying and solving the hard (open) design problems, so that trustworthy and dependable foundations can be realized.

Clearly achieving this grand global-scale end-to-end vision is beyond the ability of any one organization acting on it's own. More specifically, any new global ICT eco-system should be formally designed, specified, implemented and built in a collaborative manner with the support of community leaders for the benefit of all stakeholders. Today, we already have many world-leading RTOS vendors and WCET analysis vendors collaborating during the requirements and design stage of our secure real-time computing projects. We also have many world-leading experts in the safety, survivability and information security community collaborating on the safety and security aspects.

Various of the technologies listed above can be built in parallel. Today our focus is on advancing the secure real-time computing side as these have the least interdependencies and are absolutely essential for providing solid foundations from which to achieve a universally trustworthy and dependable ecosystem.

We seek to engage the global formal methods community today and throughout the project to realise this vision. Independent technology reviewers are now suggesting FP7 and other funding routes.

This is the grand project you've been meticulously honing your high-assurance tools, methodologies and skills for!!! Your enquires and suggestions are welcome!

## 8 Closing Statement

If nations cannot agree to a common defense based on limiting cyber *warfare* capabilities [59] then maybe we can agree to come together as netizens, organizations and nations behind a globally inclusive common cyber defense designed to resist even the most advanced cyber weapons [11] created out of fear that exploitation of cyber vulnerabilities could lead to national strategic failure [4]. Instead of cyber weapons, let's build universally trustworthy and dependable communication and computation systems that seek to protect the legitimate interests of all stakeholders in multi-jurisdiction, multi-stakeholder Internet-scale environments. Modern life is now virtually totally dependent upon ICT. Let's build ICT foundations that bring the international community together. Over a period of $\approx$ 12 years Synaptic Labs has been systematically addressing the conceptual functional and security flaws in today's ICT ecosystem. Today we are ready to embark on the high assurance development of this international vision. Let's collaborate together!

## References

1. Compromising emanations laboratory test standard. SECAN Doctrine and Information Publication SDIP-27 Level A, NATO.
2. Recommendations for a Security and Dependability Research Framework: from Security and Dependability by Central Command and Control to Security and Dependability by Empowerment. Deliverable 3.0, SecurIST, Jan. 2007.
3. U.S. Government Protection Profile for Separation Kernels in Environments Requiring High Robustness. Common Criteria Profile 1.03, US NSA IAD, June 2007.
4. Cyberspace policy review. United States, Office of the White House, May 2009.
5. Functional safety of electrical/electronic/programmable electronic safety-related systems. IEC 61508, International Electrotechnical Commission, 2010.
6. The Cloud Dividend: Part One, The economic benefits of cloud computing to business and the wider EMEA economy. Report., CEBR Ltd, December 2010.
7. *Ubiquitous High Performance Computing.* DARPA-BAA-11-55, March 2010.
8. I2O MRC Proposers Day Webcast. Technical report, DARPA, May 2011.
9. International Strategy for Cyberspace. U.S. Office of the White House, May 2011.
10. *Global Risks 2012, Insight Report.* World Economic Forum, seventh edition, 2012.
11. *Plan X Proposers' Day Workshop.* DARPA-SN-12-51, Aug 2012.
12. Ross J. Anderson. Liability and Computer Security: Nine Principles. In *ESORICS '94*, volume 875 of *LNCS*, pages 231–245. Springer-Verlag, Nov. 1994.

13. AtlanticLIVE. The atlantic and government executive cyber security forum. Video, The Atlantic, 2010. http://events.theatlantic.com/cyber-security/2010/.
14. Algirdas Avižienis, Jean-Claude Laprie, and Brian Randell. Dependability and its threats: A Taxonomy. In *Topical Days: Fault Tolerance for Trustworthy and Dependable Information Infrastructures, IFIP World Computer Congress*. Kluwer Academic Publisers., Aug. 2004.
15. Algirdas Avižienis, Jean-Claude Laprie, Brian Randell, and Carl Landwehr. Basic concepts and taxonomy of dependable and secure computing. In *IEEE Transactions on dependable and secure computing*, volume 1, Jan. 2004.
16. Paul Baran. On Distributed Communications: IX. Security, Secrecy,and Tamper-Free considerations. Memorandum RM-3765-PR, RAND, August 1964.
17. Elaine Barker. Cryptographic Key Management Project. Project, NIST, 2009.
18. Elaine Barker, Dennis Branstad, Santosh Chokhani, and Miles Smid. Cryptographic key management workshop summary (final). IR 7609, NIST, June 2009.
19. Daniel J. Bernstein, Tanja Lange, and Pierre-Louis Cayrel. Post-quantum cryptography. Website, July 2009. www.pqcrypto.org.
20. Sébastien Briais, Sylvain Guilley, and Jean-Luc Danger. A formal study of two physical countermeasures against side channel attacks. 2012/430, eprint, 2012.
21. Michael Brown. Classical Cryptosystems In A Quantum Setting. Master of mathematics in combinatorics and optimisation, Waterloo, Ontario, Canada, Apr. 2004.
22. Dai Bui, Edward Lee, Isaac Liu, Hiren Patel, and Jan Reineke. Temporal isolation on multiprocessing architectures. In *Proceedings of the 48th Design Automation Conference*, DAC '11, pages 274–279, New York, NY, USA, 2011. ACM.
23. Daniele Catteddu and Giles Hogben. Cloud Computing - Benefits, Risks and Recommendations for Information Security. Report, ENISA, Nov. 2009.
24. Francisco J. Cazorla, Mikel Fernandez, Roberto Gioiosa, Eduardo Quiñones, Marco Zulianello, and Luca Fossati. Measuring inter-task interferences in the NGMP. In *ESA Workshop on ADCSS*. ESA/ESTEC, October 2011.
25. Francisco J. Cazorla, Eduardo Quiñones, Tullio Vardanega, Liliana Cucu, Benoit Triquet, Guillem Bernat, Emery Berger, Jaume Abella, Franck Wartel, Michael Houston, Luca Santinelli, Leonidas Kosmidis, Code Lo, and Dorin Maxim. PROARTIS: Probabilistically Analysable Real-Time Systems. Rapport de recherche INIRIA/RR-7869, INRIA, Jan 2012.
26. Carlos Coronado. *Provably secure and practical signature schemes*. Doctoral thesis (elib.tu-darmstadt.de/diss/000642), Technische Universität Darmstadt, Nov. 2005.
27. Christoph Cullmann, Christian Ferdinand, Gernot Gebhard, Daniel Grund, Claire Maiza, Jan Reineke, Benoît Triquet, Simon Wegener, and Reinhard Wilhelm. Predictability considerations in the design of multi-core embedded systems. *Ingénieurs de l'Automobile*, 807:36–42, September 2010.
28. Emerich de Vattel. *The Law of Nations (Le droit des gens) - Principles of the Law of Nature Applied to the Conduct and Affairs of Nations and Sovereigns*. 1760.
29. Department of Homeland Security. A Roadmap for Cybersecurity Research. Roadmap, DHS Science and Technology Directorate, Nov. 2009.
30. Whitfield Diffie and Martin E. Hellman. Multiuser cryptographic techniques. In *AFIPS '76: Proceedings of the June 7-10, 1976, national computer conference and exposition*, pages 109–112, New York, NY, USA, June 1976. ACM.
31. Zeta Dooly, Jim Clarke, W. Fitzgerald, W. Donnelly, Michael Riguidel, and Keith Howker. ICT Security and Dependability Research beyond 2010 - Final strategy. Deliverable 3.3, SecurIST EU-FP6-004547, Jan. 2007.
32. Rolland Dudemain. When Absolute Security Really Matters! Video, Malta International Cyber Awareness Seminar, Nov. 2011.

33. Benjamin Gittins. Overview of SLL's proposal in response to NIST's call for new global IdM/CKM designs without public keys. In *Proceedings of CSIIRW-6*, CSIIRW '10, pages 60:1–60:4, New York, NY, USA, 2010. ACM.
34. Benjamin Gittins. Outline of a proposal responding to E.U. and U.S. calls for trustworthy global-scale IdM and CKM designs. Report 2011/029, ePrint, 2011.
35. Benjamin Gittins and Ronald Kelson. Feedback to NIST DRAFT Special Publication 800-130. Comment, Synaptic Laboratories Limited, August 2010.
36. Benjamin Gittins and Ronald Kelson. Overview of SLL's proposal in response to NIST's call for new global IdM/CKM designs without PKC. Video. In *IEEE Key Management Summit 2010 website*, Lake Tahoe, Nevada, May 2010. IEEE.
37. Benjamin Gittins and Ronald Kelson. Synaptic Labs' 2012 Annual Report: Part 2 - Global Cyber Safety and Security Status. Transcript, slideshow and video, Synaptic Laboratories Limited, Feb. 2012.
38. Benjamin Gittins and Ronald Kelson. Synaptic Labs' 2012 Annual Report: Part 3 - Cyber Security Technical Problems, Drivers and Incentives. Transcript, slideshow and video, Synaptic Laboratories Limited, Feb. 2012.
39. Peter Gutmann. *Engineering Security.* (draft book), Dec. 2009.
40. Melissa Hathaway. Plenary speaker. In *Cyber Security and Information Intelligence Research Workshop*, volume 6. Oak Ridge National Laboratory, Apr. 2010.
41. Peter Kogge, Keren Bergman, Shekhar Borker, Dan Campbell, William Carlson, William Dally, Monty Denneau, Paul Franzon, William Harrod, Kerry Hill, Jon Hiller, Sherman Karp, Stephen Keckler, Dean Klein, Robert Lucas, Mark Richards, Al Scarpelli, Steven Scott, Allan Snavely, Thomas Sterling, R. Stanley Williams, and Katherine Yelick. ExaScale Computing Study: Technology Challenges in Achieving Exascale Systems. Report, DARPA, Sep. 2008.
42. Tom Krazit. *Google fired engineer for privacy breach.* Sep. 2010.
43. Edward A. Lee. Verifying real-time software is not reasonable (today). In *Eighth Haifa Verification Conference*, LNCS. Springer, Nov. 2012.
44. Isaac Liu and David McGrogan. Elimination of side channel attacks on a precision timed architecture. Technical Report UCB/EECS-2009-15, EECS Department, UC Berkeley, Jan 2009.
45. Peter A. Loscocco, Stephen D. Smalley, Patrick A. Muckelbauer, Ruth C. Taylor, S. Jeff Turner, and John F. Farrell. The Inevitability of Failure: The Flawed Assumption of Security in Modern Computing Environments. In *21'st NISSC*. NIST NCSC, NIST, Sep. 1998.
46. Moxie Marlinspike. *Defeating OCSP With The Character '3'.* July 2009.
47. Moxie Marlinspike. *Null Prefix Attacks Against SSL/TLS Certificates.* July 2009.
48. Joseph Menn. Founding father wants secure 'Internet 2'. News Article, Financial Times Limited, October 2011.
49. Sibin Mohan, Marco Caccamo, Lui Sha, Rodolfo Pellizzoni, Greg Arundale, Russell Kegley, and Dionisio de Niz. Using Multicore Architectures in Cyber-Physical Systems. In *Workshop on Developing Dependable and Secure Automotive Cyber-Physical Systems from Components*, Mar. 2011.
50. Montesquieu. *The Spirit of the Laws.* Crowder, Wark, and Payne, 1777.
51. NATO. Developing NATO's cyber defence policy. News Article, NATO, Jan 2011.
52. NITRD. NITRD 2010 Cybersecurity R&D Themes Webcast. In *Federal Cybersecurity Game-change R&D website.* NITRD, May 2010.
53. Sean O'Neil, Benjamin Gittins, and Howard A. Landman. VEST Ciphers (eSTREAM Phase 2). In *ECRYPT eSTREAM*, Aug. 2006.
54. Marios Papaefthymiou. Charge-Recovery VLSI. In *The Berkeley Wireless Research Center*, Feb 2008.

55. Peter Puschner, Raimund Kirner, and Robert G. Pettit. Towards composable timing for real-time programs. In *Proceedings of the 2009 Software Technologies for Future Dependable Distributed Systems*, STFSSD '09, pages 1–5, Washington, DC, USA, 2009. IEEE Computer Society.
56. QinetiQ. National Cyber Leap Year Summit 2009 – Co-Chairs' Report. On behalf of the US NITRD Program, Sep. 2009.
57. Mathieu Renauld, Francois-Xavier Standaert, Nicolas Veyrat-Charvillon, Dina Kamel, and Denis Flandre. A Formal Study of Power Variability Issues and Side-Channel Attacks for Nanoscale Devices. In *Advances in Cryptology - EUROCRYPT 2011*, volume 6632 of *LNCS*, page 109. Springer, 2011.
58. RISEPTIS. Trust in the Information Soceity. Report, Research & Innovation Security, Privacy and Trusworthiness in the Information Society, 2011.
59. Nicholas C. Rueter. The Cybersecurity Dilemma. Thesis, Duke University, 2011.
60. Subhash Sankuratripat. Interoperable Key Management using the OASIS KMIP Standard. In *IEEE Key Management Summit 2010 website*, Lake Tahoe, Nevada on May 4-5, 2010., May 2010. IEEE.
61. Smruti R. Sarangi, Brian Greskamp, and Josep Torrellas. Cadre: Cycle-accurate deterministic replay for hardware debugging. In *Proceedings of the International Conference on Dependable Systems and Networks*, DSN '06, pages 301–312, Washington, DC, USA, 2006. IEEE Computer Society.
62. Visvesh S. Sathe, Juang-Ying Chueh, and Marios C. Papaefthymiou. Energy-Efficient GHz-Class Charge-Recovery Logic. In *IEEE Journal of Solid-State Circuits*, volume 42, pages 38–47, Jan. 2007.
63. Brian Snow. We need assurance! In *ACSAC '05: Proceedings of the 21st Annual Computer Security Applications Conference*, pages 3–10, Washington, DC, USA, Dec. 2005. IEEE Computer Society.
64. Brian Snow. Our Security Status is Grim. Video, Malta International Cyber Awareness Seminar, Nov. 2011.
65. Brian Snow. Statement on Synaptic Laboratories Ltd. Open letter, July 2011.
66. Brian Snow. *The Importance of Implementation.* World Science Festival, 2011.
67. Kris Tiri and Ingrid Verbauwhede. Design method for constant power consumption of differential logic circuits. In *Proceedings of the conference on Design, Automation and Test in Europe*, volume 1, pages 628–633. IEEE Computer Society, 2005.
68. Benoît Triquet. Mixed Criticality in Avionics. March 2012.
69. UK Goverment. *The UK Cyber Security Strategy.* UK Cabinet Office, Nov. 2011.
70. Theo Ungerer. Parallelisation of Hard Real-time Applications for Embedded Multi- and Many-cores. In *MARC ONERA Symposium Toulouse*, July 2012.
71. Tullio Vardanega, Juan Zamorano, and Juan Antonio De La Puente. On the dynamic semantics and the timing behavior of Ravenscar kernels. *Real-Time Syst.*, 29(1):59–89, January 2005.
72. Reinhard Wilhelm, Christian Ferdin, Christoph Cullmann, Daniel Grund, Jan Reineke, and Benôit Triquet. Designing Predictable Multicore Architectures for Avionics and Automotive Systems. In *RePP*, Oct. 2009.
73. Reinhard Wilhelm, Daniel Grund, Jan Reineke, Marc Schlickling, Markus Pister, and Christian Ferdinand. Memory hierarchies, pipelines, and buses for future architectures in time-critical embedded systems. *IEEE Transactions on CAD of Integrated Circuits and Systems*, 28(7):966–978, July 2009.

# We Need Assurance!

Brian Snow
*U. S. National Security Agency*
bdsnow@nsa.gov

## Abstract

*When will we be secure? Nobody knows for sure – but it cannot happen before commercial security products and services possess not only enough functionality to satisfy customers' stated needs, but also sufficient assurance of quality, reliability, safety, and appropriateness for use. Such assurances are lacking in most of today's commercial security products and services. I discuss paths to better assurance in Operating Systems, Applications, and Hardware through better development environments, requirements definition, systems engineering, quality certification, and legal/regulatory constraints. I also give some examples.*

## 1. Introduction

This is an expanded version of the "Distinguished Practitioner" address at ACSAC 2005 and therefore is less formal than most of the papers in the proceedings.

I am very grateful that ACSAC chose me as a distinguished practitioner, and I am eager to talk with you about what makes products and services secure.

Most of your previous distinguished practitioners have been from the open community; I am from a closed community, the U.S. National Security Agency, but I work with and admire many of the distinguished practitioners from prior conferences.

I spent my first 20 years in NSA doing research developing cryptographic components and secure systems. Cryptographic systems serving the U.S. government and military spanning a range from nuclear command and control to tactical radios for the battlefield to network security devices use my algorithms.

For the last 14 years, I have been a Technical Director at NSA (similar to a chief scientist or senior technical fellow in industry) serving as Technical Director for three of NSA's major mission components: the Research Directorate, the Information Assurance Directorate, and currently the Directorate for Education and Training (NSA's Corporate University). Throughout these years, my mantra has been, "Managers are responsible for doing things right; Technical Directors are responsible for finding the right things to do."

There are many things to which NSA pays attention in developing secure products for our National Security Customers to which developers of commercial security offerings also need to pay attention, and that is what I want to discuss with you today.

## 2. Setting the context

The RSA Conference of 1999 opened with a choir singing a song whose message is still valid today: "Still Haven't Found What I'm Looking For". The reprise phrase was . . . *"When will I be secure? Nobody knows for sure. But I still haven't found what I'm looking for!"*

That sense of general malaise still lingers in the security industry; why is that? Security products and services should stop malice in the environment from damaging their users. Nevertheless, too often they fail in this task. I think it is for two major reasons.

First, too many of these products are still designed and developed using methodologies assuming random failure as the model of the deployment environment rather than assuming malice. There is a world of difference!

Second, users often fail to characterize the nature of the threat they need to counter. Are they subject only to a generic threat of an opponent seeking some weak system to beat on, not necessarily theirs, or are they subject to a targeted attack, where the opponent wants something specific of theirs and is willing to focus his resources on getting it?

The following two simple examples might clarify this.

Example 1: As a generic threat, consider a burglar roaming the neighborhood wanting to steal a VCR. First, understand his algorithm: Find empty house (dark, no lights) try door; if open, enter, if VCR – take. If the door is resistant, or no VCR is present, find another dark house.

Will the burglar succeed? Yes, he will probably get a VCR in the neighborhood. Will he get yours? What does it take to stop him? Leave your lights on when you go out (9 cents a kilowatt-hour) and lock your door. That is probably good enough to stop the typical generic burglar.

Example 2: As a targeted threat, assume you have a painting by Picasso worth $250,000 hanging above your fireplace, and an Art thief knows you have it and he wants it. What is his algorithm? He watches your house until he sees the whole family leave. He does not care if the lights are on or not. He approaches the house and tries the door; if open, he enters. If locked, he kicks it in. If the door resists, he goes to a window. If no electronic tape, he breaks the glass and enters. If electronic tape is present, he goes to the siding on the house, rips some off, then tears out the fiberboard backing, removes the fiberglass insulation, breaks though the interior gypsum board, steps between the studs, and finally takes the painting and leaves.

It takes more effort to counter a targeted threat. In this case, typically a burglar alarm system with active polling and interior motion sensors as a minimum (brick construction would not hurt either). With luck, this should be enough to deter him. If not, at least there should be increased odds of recovery due to hot pursuit once the alarms go off.

There is no such thing as perfect security; you need to know how much is enough to counter the threat you face, and this changes over time.

## 3. What do we need?

NSA has a proud tradition during the past 53 years of providing cryptographic hardware, embedded systems, and other security products to our customers. Up to a few years ago, we were a sole-source provider. In recent years, there has come to be a commercial security industry that is attractive to our customers, and we are in an unaccustomed position of having to "compete." There is nothing wrong with that. *If* industry can meet our customer's needs, so be it.

Policy and regulation still require many of our customers to accept Government advice on security products. However, they really press us to recommend commercial solutions for cost savings and other reasons. Where we can, we do so. However, we do not do it very often because we still have not found what *we* are looking for – assurance.

Assurance is essential to security products, but it is missing in most commercial offerings today. The major shortfall is absence of assurance (or safety) mechanisms in *software*. If my car crashed as often as my computer does, I would be dead by now.

In fact, compare the software industry to the automobile industry at two points in its history, the 1930s and today. In 1930, the auto industry produced cars that could go 60 mph or faster, looked nice, and would get you from here to there. Cars "performed" well, but did not have many "safety features." If you were in an accident at high-speed, you would likely die.

The car industry today provides air bags, seat belts, crush zones, traction control, anti-skid braking, and a host of other safety details (many required by legislation) largely invisible to the purchaser. Do you *regularly* use your seat belt? If so, you realize that users *can* be trained to want and to use assurance technology!

The software security industry today is at about the same stage as the auto industry was in 1930; it provides performance, but offers little safety. For both cars and software, the issue is really assurance.

Yet what we need in security products for high-grade systems in DoD is more akin to a military tank than to a modern car! Because the environment in which our products must survive and function (battlefields, etc.) has malice galore.

I am looking forward to, and need, convergence of government and commercial security products in two areas: assurance, and common standards. Common standards will come naturally, but assurance will be harder – so I am here today as an evangelist for assurance techniques.

Many vendors tell me that users are not willing to pay for assurance in commercial security products; I would remind you that Toyota and Honda penetrated U.S. Markets in the 70's by differentiating themselves from other brands by improving reliability and quality! What software vendor today will become the "Toyota" of this industry by selling robust software?

## 4. Assurance: first definition

What do I mean by assurance? I'll give a more precise definition later, but for now it suffices to say that assurance work makes a user (or accreditor) more confident that the system works as intended, without flaws or surprises, even in the presence of malice.

We analyze the system at design time for potential problems that we then correct. We test prototype devices to see how well they perform under stress or when used in ways beyond the normal specification. Security acceptance testing not only exercises the product for its expected behavior given the expected environment and input sequences, but also tests the product with swings in the environment outside the specified bounds and with improper inputs that do not match the interface specification. We also test with proper inputs, but in an improper sequence. We anticipate malicious behavior and design to counter it, and then test the countermeasures for effectiveness. We expect the product to behave safely, even if not properly, under any of these stresses. If it does not, we redesign it.

I want functions *and* assurances in a security device. We do not "beta-test" on the customer; if my product fails, someone might die.

Functions are typically visible to the user and commanded through an interface. Assurances tend to be invisible to the user but keep him safe anyway.

Examples would be thicker insulation on a power wire to reduce the risk of shock, and failure analysis to show that no single transistor failure will result in a security compromise.

Having seat belts in a car provides a safety function. Having them made of nylon instead of cotton is the result of assurance studies that show nylon lasts longer and retains its strength better in the harsh environment of a car's interior.

Assurance is best addressed during the initial design and engineering of security systems – not as after-market patches. The earlier you include a security architect or maven in your design process, the greater is the likelihood of a successful and robust design. The usual quip is, "He who gets to the interface first, wins".

When asked to predict the state of "security ten years from now," I focus on the likely absence of assurance, rather than the existence of new and wonderful things.

Ten years from now, there will still be security-enhanced software applications vulnerable to buffer overflow problems. These products will not be secure, but will be sold as such.

Ten years from now, there will still be security-enhanced operating systems that will crash when applications misbehave. They will not be secure either.

Ten years from now, we will have sufficient functionality, plenty of performance, but not enough assurance.

Otherwise, predicting ten years out is simply too hard in this industry, so I will limit myself to about five years. Throughout the coming five-year span, I see little improvement in assurance, hence little true security offered by the industry.

## 5. The current state of play

Am I depressed about this state of affairs? Yes, I am. The scene I see is products and services sufficiently robust to counter many (but not all) of the "hacker" attacks we hear so much about today, but not adequate against the more serious but real attacks mounted by economic enemies, organized crime, nation states, and yes, terrorists.

*We will be in a truly dangerous stance*: we will think we are secure (and act accordingly) when in fact we are not secure.

The serious enemy knows how to hide his activities. What is the difference between a hacker and a more serious threat such as organized crime? The hacker wants a *score*, and bragging rights for what he has obviously defaced or entered. Organized crime wants a *source*, is willing to work long, hard, and quietly to get in, and once in, wants to stay invisible and continue over time to extract what it needs from your system.

Clearly, we need confidence in security products; I hope we do not need a major bank-failure or other disaster as a wake-up call before we act.

The low-level hackers and "script-kiddies" who are breaking systems today and are either bragging about it or are dumb enough to be caught, are providing some of the best advertising we could ask for to justify the need for assurance in security products.

They demonstrate that assurance techniques (*barely*) adequate for a benign environment simply will not hold up in a malicious environment, so we *must* design to defeat malice. Believe me – there is malice out there, beyond what the "script-kiddies" can mount.

However, I do fear for the day when the easy threats are countered – that we may then stop at that level, rather than press on to counter the serious and pernicious threats that can stay hidden.

During the next several years, we need major pushes and advances in three areas: Scalability, Interoperability, and Assurance. I believe that market pressures will provide the first two, but not the last one – assurance.

There may or may not be major breakthroughs in new security functions; but we really do not need many new functions or primitives – if they come, that is nice. If they do not, we can make do with what we have.

What we really need but are not likely to get is greater levels of assurance. That is sad, because despite the real need for additional research in assurance technology, the real crime is that we fail to use fully that which we already have in hand! We need to better use those confidence-improving techniques that we do have, and continue research and development efforts to refine them and find others.

I am not asking for the development of new science; the safety and reliability communities (and others) know how to do this – go and learn from them.

You are developers and marketers of security products, and I am sorry that even as your friend I must say, "Shame on you. You should build them better!" It is a core quality-of-implementation issue. The fact that teen-age hackers can penetrate many of your devices from home is an abysmal statement about the security-robustness of the products.

## 6. Assurance: second definition

It is time for a more precise definition. Assurances are confidence-building activities demonstrating that

1. $ The system's security policy is internally consistent and reflects the requirements of the organization,
2. $ There are sufficient security functions to support the security policy,
3. $ The system functions meet a desired set of properties and *only* those properties,
4. $ The functions are implemented correctly, and
5. $ The assurances *hold up* through the manufacturing, delivery, and life cycle of the system.

We provide assurance through structured design processes, documentation, and testing, with greater assurance provided by more processes, documentation, and testing.

I grant that this leads to increased cost and delayed time-to-market – a severe one-two punch in *today's* marketplace; but your customers are growing resistive and are beginning to expect, and to demand, better products *tomorrow*. They are near the point of chanting, "I'm mad as hell, and I'm not going to take it anymore!"

Several examples of assurance techniques come to mind; I will briefly discuss some in each of the following six areas: operating systems, software modules, hardware features, systems engineering, third party testing, and legal constraints.

## 7. Operating systems

Even if operating systems are not truly secure, they can at least remain benign (not actively malicious) if they would simply enforce a digital signature check on every critical module prior to each execution. Years ago, NSA's research organization wrote test code for a UNIX system that did exactly that. The performance degraded about three percent. This is something that is doable!

Operating Systems should be self-protective and enforce (at a minimum) separation, least-privilege, process-isolation, and type-enforcement.

They should be aware of and enforce security policies! Policies drive requirements. Recall that Robert Morris, a prior chief scientist for the National Computer Security Center, once said: "Systems built without requirements cannot fail; they merely offer surprises – usually unpleasant!"

Given today's common hardware and software architectural paradigms, operating systems security is a major primitive for secure systems – you will not succeed without it. This area is so important that it needs all the emphasis it can get. It is the current "black hole" of security.

The problem is innately difficult because from the beginning (ENIAC, 1944), due to the high cost of components, computers were built to share resources (memory, processors, buses, etc.). If you look for a one-word synopsis of computer design philosophy, it was and is SHARING. In the security realm, the one word synopsis is SEPARATION: keeping the bad guys away from the good guys' stuff!

So today, making a computer secure requires imposing a "separation paradigm" on top of an architecture built to share. That is tough! Even when partially successful, the residual problem is going to be covert channels. We really need to focus on making a secure computer, not on making a computer secure – the point of view changes your beginning assumptions and requirements!

## 8. Software modules

Software modules should be well documented, written in certified development environments, (ISO 9000, SEI-CMM level five, Watts Humphrey's Team Software Process and Personal Software Process (TSP/PSP), etc.), and *fully* stress-tested at their interfaces for boundary-condition behavior, invalid inputs, and proper commands in improper sequences.

In addition to the usual quality control concerns, *bounds checking* and *input scrubbing* require special attention. For bounds checking, verify that inputs are of the expected type: if numeric, in the expected range; if character strings, the length does not exceed the internal buffer size. For input scrubbing, implement reasonableness tests: if an input should be a single word of text, a character string containing multiple words is wrong, even if it fits in the buffer.

A strong quality control regime with aggressive bounds checking and input scrubbing will knock out the vast majority of today's security flaws.

We also need good configuration control processes and design modularity.

A good security design process requires review teams as well as design teams, and no designer should serve on the review team. They cannot be critical enough of their own work. Also in this world of multi-national firms with employees from around the world, it may make sense to take the national affinity of employees into account, and not populate design and review teams for a given product with employees of the SAME nationality or affinity. Half in jest I would say that if you have Israelis on the design team put Palestinians on the review team; or if Germans are on one, put French on the other. . . .

Use formal methods or other techniques to assure modules meet their specifications exactly, with no extraneous or unexpected behaviors – especially embedded malicious behavior.

Formal methods have improved dramatically over the years, and have demonstrated their ability to reduce errors, save time, and even save dollars! This is an under-exploited and very promising area deserving more attention.

I cite two examples of formal methods successes: The Microsoft SLAM static driver verifier effort coming on line in 2005, and Catherine Meadows' NRL Protocol Analyzer detecting flaws in the IKE (Internet Key Exchange) protocol in 1999. You may have your own recent favorites.

As our systems become more and more complex, the need for, and value of, formal methods will become more and more apparent.

## 9. Hardware features

Consider the use of smartcards, smart badges, or other hardware tokens for especially critical functions. Although more costly than software, when properly implemented the assurance gain is great. The form-factor is not as important as the existence of an isolated processor and address space for assured operations – an "Island of Security," if you will. Such devices can communicate with each other through secure protocols and provide a web of security connecting secure nodes located across a sea of insecurity in the global net.

I find it depressing that the hardware industry has provided hardware security functionality (from the Trusted Platform Group and others) now installed in processors and motherboards that is not yet accessed or used by the controlling software, whether an OS or an application.

## 10. Security systems engineering

How do we get high assurance in commercial gear?

a) How can we trust, or
b) If we cannot trust, how can we safely use,
security gear of unknown quality?

Note the difference in the two characterizations above: *how we phrase the question may be important*. For my money, I think we need more focus on how to use safely security gear of unknown quality (or of uncertain provenance).

I do not have a complete answer on how to handle components of unknown quality, but my thoughts lean toward systems engineering approaches somewhat akin to what the banking industry does in their systems. No single component, module, or person knows enough about the overall transaction processing system to be able to mount a successful attack at any one given access point. To be successful the enemy must have access at multiple points and a great deal of system architecture data.

Partition the system into modules with "blinded interfaces" and limited authority where the data at any one interface are insufficient to develop a complete attack. Further, design cooperating modules to be "mutually suspicious," auditing and alarming each other's improper behavior to the extent possible.

For example: if you are computing interest to post to accounts there is no need to send the complete account record to a subroutine to adjust the account balance. Just send the current balance and interest rate, and on return store the result in the account record. Now the interest calculating subroutine *cannot* see the data on the account owner, and therefore cannot target specific accounts for theft or other malicious action. We need to trust the master exec routine, but minimize the number of subroutines we need to trust. Yes, I know this is over-simplified, but you get my drift.

In addition, to guard against "unintended extra functionality" within given hardware modules or software routines, the development philosophy needs to enforce something akin to "no-lone zones" in that no single designer or coder can present a "black-box" (or proprietary?) effort to the system design team that is tested only at its interfaces and is then accepted.

Review all schematics and code (in detail, line by line) for quality and "responsive to stated requirement" goals. This review should be by parties independent of the designer. This is expensive, but not far from processes required today in many quality software development environments to address reliability and safety concerns.

This of course requires all tools (compilers, CAD support, etc.) used in the development environment to be free of malice; that can be a major hurdle and a difficult assurance task in and of itself (remember the Thompson compiler in "Reflections on Trusting Trust, CACM 1983)!

The "Open Source" movement may also provide value in this area. There are pluses and minuses with open source, but from the security viewpoint, I believe it is primarily a plu**s.**

Further architectural constraints may be imposed to make up for deficiencies in certain modules. Rather than (or in addition to) encryption in application processes prior to transmission to other sites which could be bypassed or countered by a malicious operating system, you might require site-to-site transmissions to go through an encrypting modem or other in-line, non-bypassable link encryptors.

Link encryption in addition to application layer encryption is an example of a "Defense in Depth" strategy that attempts to combine several weak or possibly flawed mechanisms in a fashion robust enough to provide protection at least somewhat stronger than the strongest component present.

Synergy, where the strength of the whole is greater than the sum of the strength of the parts, is highly desirable but not likely. We must avoid at all costs the all-too-common result where the system strength is less than the strength offered by the strongest component, and in some worst cases less than the weakest component present. Security is so very fragile under composition; in fact, secure composition of components is a major research area today.

Good *system* security design today is an art, not a science. Nevertheless, there are good practitioners out there that can do it. For instance, some of your prior distinguished practitioners fit the bill.

*This area of "safe use of inadequate components" is one of our hardest problems, but an area where I expect some of the greatest payoffs in the future and where I invite you to spend effort.*

## 11. Third party testing

NIST (and NSA) provide third-party testing in the National Information Assurance Partnership Laboratories (NIAP labs), but Government certification programs will only be successful if users see the need for something other than vendor claims of adequacy or what I call "proof by emphatic assertion – Buy me, I'm Good."

If not via NIST or other government mechanism, then the industry must provide *third-party* mediation for vendor security claims via consortia or other mechanisms to provide *independent* verification of vendor claims *in a way understandable by users*.

## 12. Market/legal/regulatory constraints

Market pressures are changing, and may now help drive more robust security functionality. The emergence of e-commerce in the past decade as a driver for secure internet financial transactions is certainly helpful, as is the entertainment industry's focus on digital rights management. These industries certainly want security laid on correctly and robustly!

I hope citizens will be able to use the emerging mechanisms to protect personal data in their homes, as well as industry using the mechanisms to protect industry's fiscal and intellectual property rights. It is simply a matter of getting the security architecture right.

I wonder if any of the industry consortia working on security for digital rights management and/or electronic fiscal transactions have citizen advocates sitting on their working groups.

Lawsuits might help lead to legal "fitness-for-use" criteria for software products – much as other industries face today. This could be a big boon to assurance – liability for something other than the quality of the media on which a product is delivered!

Recall that failure to deliver expected functionality can be viewed, in legal parlance, as providing an "attractive nuisance" and is often legally actionable.

One example is a back yard swimming pool with no fence around it. If a neighbor's child drowns in it, you can be in deep trouble for providing an attractive nuisance. Likewise, if you do a less than adequate job of shoveling snow from your walk in winter (providing the appearance of usability) you can be liable if someone slips on the ice you left on the surface. Many software security products today are attractive nuisances!

All you need do is to Google "Software Quality Lawsuits" or a similar phrase, and you can find plenty of current examples of redress sought under law for lack of quality in critical software. Do not attempt to manage defects in software used in life-critical applications. Remove them during the development and testing processes! People have died due to poor software in medical devices, and the courts are now engaged; the punitive awards can be significant.

One example of a lawsuit already settled: *General Motors Corp. v. Johnston* (1992). A truck stalled and was involved in an accident because of a defect in a PROM, leading to the death of a seven-year old child. An award of $7.5 million in punitive damages against GM followed, in part due to GM knowing of the fault, but doing nothing.

There are social processes outside the courts that can also drive vendors toward compliance with quality standards.

One of the most promising recent occurrences in the insurance industry was stated in the report of Rueschlikon 2005 (a conference serving the insurance industry). Many participants felt that, "The insurance industry's mechanisms of premiums, deductibles, and eligibility for coverage can incent best practices and create a market for security . . . This falls in line with the historic role played by the insurance industry to create incentives for good practices, from healthcare to auto safety . . . Moreover, the adherence to a set of best practices suggest that if they were not followed, firms could be held liable for negligence."

Bluntly, if your security product lacks sufficient robustness in the presence of malice, your customers will have to pay more in insurance costs to mitigate their risks.

How the insurance industry will measure best practices and measure compliance are still to be worked out, but I believe *differential* pricing of business disaster recovery insurance based in part on quality/assurance (especially of security components) is a great stride forward in bringing market pressure to bear in this area!

## 13. Summary

In closing, I reiterate that what we need most in the future is more assurance rather than more functions or features. The malicious environment in which security systems must function *absolutely requires* the use of strong assurance techniques.

Remember: most attacks today result from failures of assurance, not failures of function.

Rather than offer predictions, try for a self-fulfilling prophecy – each of us should leave this conference with a stronger commitment to using available assurance technology in products! It is not adequate to *have* the techniques; we must *use* them!

We have our work cut out for us; let's go do it.

In closing, I would like to thank Steven Greenwald, Brad Martin, and Greg Shipley for their insights and help in preparing this article.