Platform-level Formal Verification for Public Sector Trustworthy Computing: Considerations and Challenges

EuroProofNet Tutorial on Usable Formal Methods for Security of Systems

03/27/2024, Andreas Berg

gematik

Founded 2005, National Agency for Digital Medicine (DiHA, soon)

82 Million INSURED

101.000 DOCTOR‘s Offices

1.950 HOSPITALS

64.370 PHARMACIES

147 INSURERS

- **e-Patient Records**
- **e-Prescriptions**
- **Emergency Data**
- **Email, Messaging**
- **DEMIS, ISiK, …**

Applications / Services

Andreas Berg
IT Architect

@gematik

- Since Oct. 2019 (freelance on & off 2013 - 2019)

Interests

- Technologies and methods for high assurance trustworthy IT systems
- Confidential Computing

Projects

- Security architecture of e-Patient Records (”ePA“) and e-Prescriptions (“E-Rezept“)
- Future architecture concepts
  - Platform (“TI 2.0“)
  - Zero Trust Architecture concept
  - Current main focus: “**Healthcare Confidential Computing**” (“HCC”)

TI-Evolution

Timeless Goals: Valuable Service Portfolio, Interoperability, Security

Original TI	Current TI	Future TI (TI 2.0 & HCC)
• Smartcard-based IDs (eGK, HBA, SMC-B)	• eGK- & mobile-based insured/patient IDs	• eIDs (OIDC, Wallet) for all
• Emergency data on eGK • Data processing on prem – decentralized • Connector, eHealth CT • TI as dedicated IP-sec VPN	• Some datacenter-based processing of PMR • Isolated compute (cages, locked racks) • All-in-one suppliers	• Processing of PMR in Cloud-style infrastructures • Separation of Application & HCC Infrastructure Service Providers • gematik as Trust Domain & Attestation Provider
• No direct access for insured/patients	• Internet-based access for insured/patients	• Internet-based access to all services for all participants
• Specifications on paper	• Specifications on paper	• APIs & security as code

HCC Challenges I

- **Protect personal medical records at assurance level “high” / ”very high”.**
  - Especially if records of millions of citizens are aggregated in an HCC Provider’s DC
  - Prevent qualified insider attackers from gaining access to any of the medical records.

- **Define suitable provider / solution certification scheme**
  - C5, PCI-DSS, CSA CCM provide mostly organizational frameworks without assurance levels.
  - Confidential computing technical measures need certification with quality (e.g. CC EAL).
  - Formality of specification / certification as quality metric?

- **Address side channels, better: avoid them altogether.**
  - Limit compute resource sharing to services evaluated to Trust Domain’s assurance level.
  - May reduce scalability / availability

HCC Challenges II

- **Establish Trust Domain Provider (gematik) as attestation authority**
  - Trust domain services hosted by HCC Provider keep provider responsible for availability.
  - Build cryptographically secured trust domain administration / policy management processes.
  - Provider’s tenancy model as a basis for separation of organizational roles?
  - Capture all HW / SW dependencies of TCB.

- **Minimize TCB**
  - Excludes most legacy software re-use as well as Cloud-native services, if not certified
  - Standardize attestable workload binaries across providers

- **Automate (Re-)certification as much as possible**
  - Integrate with CI/CD pipeline

Security Model

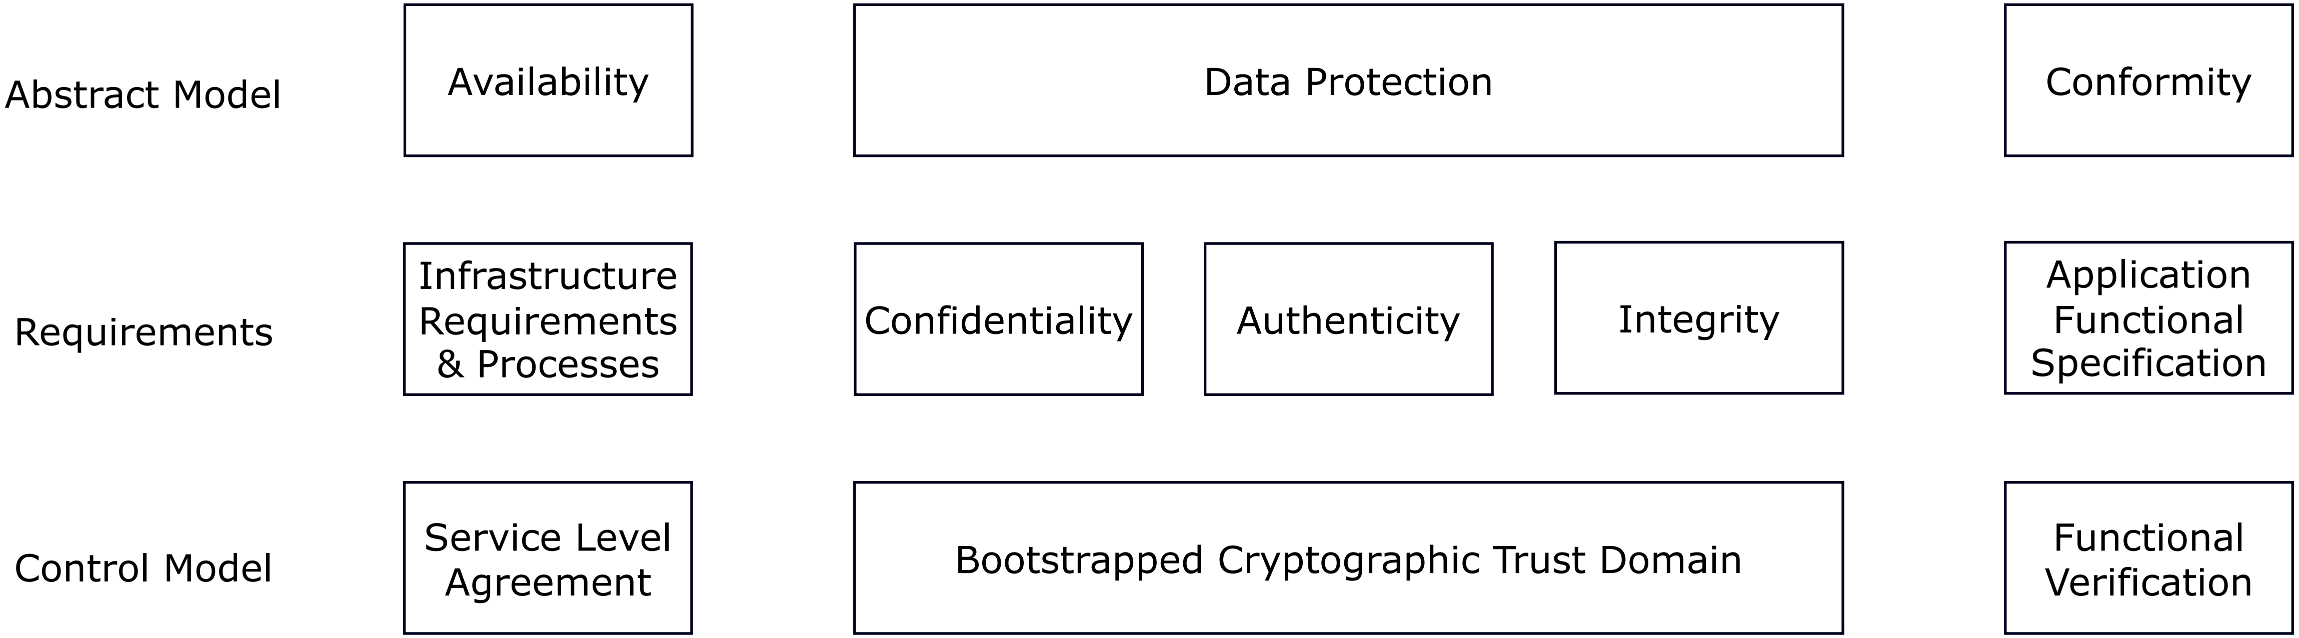

Abstract Model	Availability	Data Protection			Conformity
Requirements	Infrastructure Requirements & Processes	Confidentiality	Authenticity	Integrity	Application Functional Specification
Control Model	Service Level Agreement	Bootstrapped Cryptographic Trust Domain			Functional Verification

Minimal Viable Platform

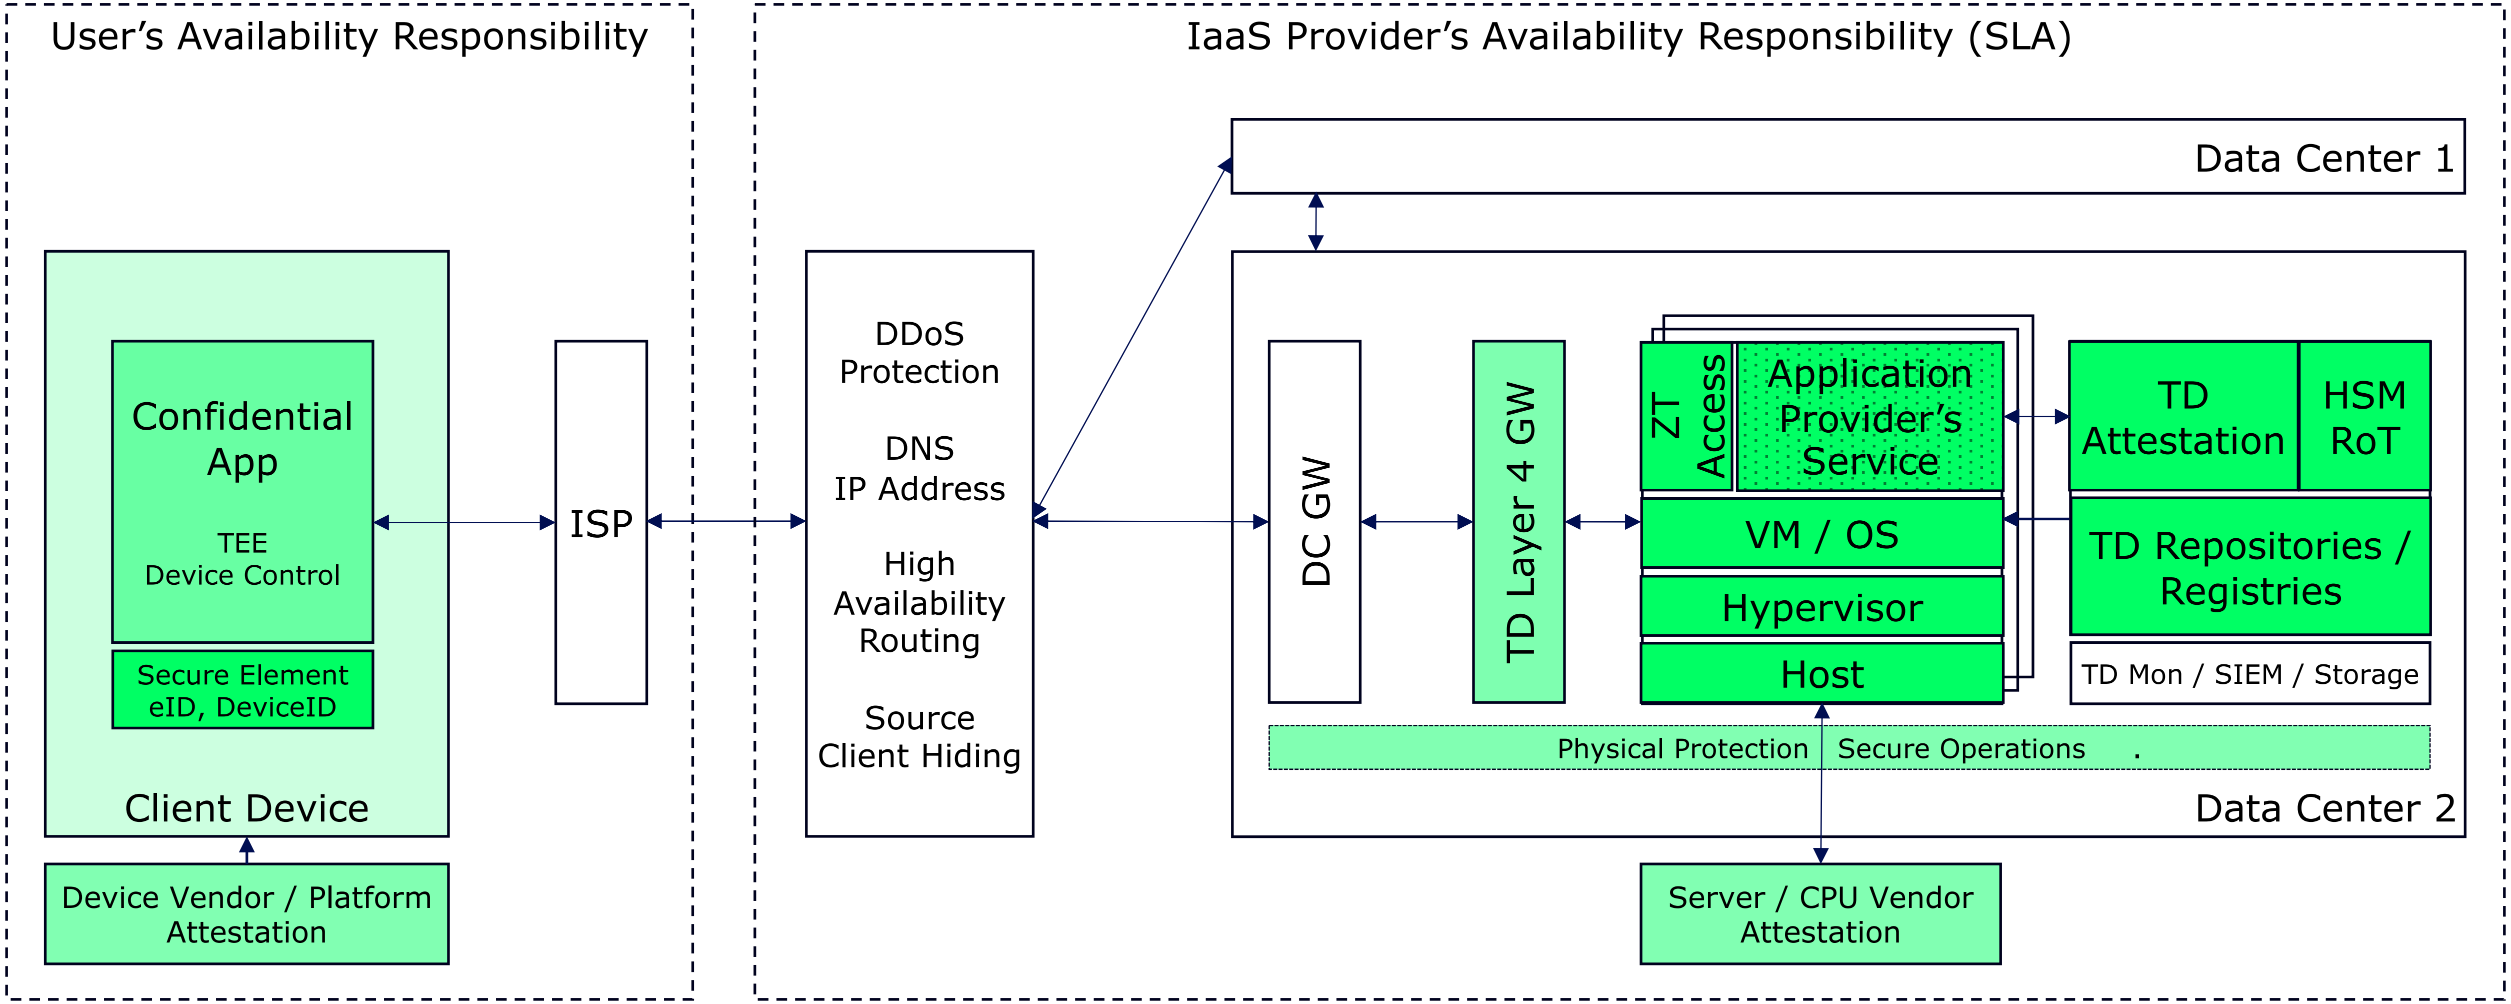

Motivations Beyond the Healthcare Sector

Disclaimer: My personal views

- We are witnessing the power shifts induced by the new leading digital medium *(McLuhan’s ”Understanding Media”)*.
- Our institutional order has not kept up with digital society, commerce, finance, (cyber)war.
- Paper-based, manually operating institutions fail to protect citizens, societies.
- Digitalized environments are inherently totalitarian because they are ubiquitous observers, require control to be operational, and increase efficiency with integration.
- Individual freedom needs to be explicitly, transparently implemented (“freedom by design”).
- State authorities must be able to provide individuals and organizations with secure public digital services via an unsecure Internet. (It’s not about “making the Internet a safe place”.)
- Automated decisions rules need to be **justified by formal verification of correctness**.

Some Existing “Lighthouse Projects” / “Assets“

- *Univ. of Cambridge* – “**CHERI**” – Enhanced ISA (now RISC V) with fine-grained memory protection and scalable software compartmentalization, compiler extensions and other tooling
- *lowRISC et. al.* – “**OpenTitan**” – Open source, high-quality reference design and integration guidelines for silicon root of trust chips
- *SiFive* – “**Formal Specification of RISC-V ISA**” – Verified using KAMI, a DSL in Coq, produces Verilog
- *seL4 Foundation* – “**seL4 Microkernel**” - High-assurance, high-performance operating system microkernel, capabilities-based access control, Isabelle/HOL
- *Microsoft Research et. al.* – “**Project Everest**” – QUIC & TLS 1.3 Record Layers, crypto algorithms, binary parser generator framework, F* ATP

Questions or Comments (so far)?

Reverse Q & A

Feasibility

Assuming things go reasonably well, **when** would you think a formally verified platform TCB (excluding business logic) for scalable public cloud services could be available?

a) In 5 years
b) In 10 years
c) Never

Re-using Existing Code

Can we use existing code & APIs, especially if battle-proven but not formalized, and “lift them” into formally verified assets?

a) Yes.
b) It’s hard.
c) No.

Integrating the Foundations

Do we have to decide on a core mathematical foundation with deep semantics and migrate existing propositions and proofs from other foundations into it or should we integrate across foundations using mappings (e. g., institution morphisms)?

a) Core foundation
b) Mappings
c) It depends.

Integrating Results

Should all code be organized in a central “syslib” repository analogous to mathlib, UniMath, or the Archive of Formal Proofs?

a) Yes.
b) Not necessary.

AI Revolution

Can LLMs help with “lifting” of assets, with mappings or embeddings, especially if the LLMs are interacting with proof assistants and / or proof checkers?

a) AI will work it out for us.
b) AI will be helpful.
c) AI will get stuck.
d) AI will kill us.

Gaining Trust

We need to convey the platform’s trustworthiness to ordinary people. Is it imaginable to produce a representation of the platform as a top-down layered set of abstractions starting from a simple “It's secure” (alternatively starting from the shown security model)?

a) Yes.
b) No.
c) Makes no sense.

gematik.
Gesunde Aussichten.

Contact

Andreas Berg
Systems Engineering
+49 172 3137786
andreas.berg@gematik.de