TROOPERS
IoT INSIGHT SUMMIT

ERNW
providing security.

Updating #IoT Devices / Some Reflections

Matthias Luft, mluft@ernw.de, @uchi_mata
Enno Rey, erey@ernw.de, @Enno_Insinuator

#IoT Will Change the World We Live In

Digital devices (hardware running software) will control:

- the air we breath

- the food we consume, in many ways

- vital functions of our bodies (ask Marie…)

- physical properties/attributes of our household environments (doors, temperature, security against intruders)

- the way we move from one place to another

Types of Control, in #IoT Context

- Direct control
  - Steer an actuator, monitor/gather data from a sensor

- Indirect control
  - Remote management
  - Patching

Notable Properties of Smart Objects

- Constraints as for memory, computing power, power (batteries)
  - <=> (Secure) Remote Mgmt Capabilities

- May be physically accessible by non-trustworthy parties, or (phys.) inaccessible by trusted parties
  - → more challenges as for "management"

- Long lifespan (some estimates up to 40 years)

The Lifespan Aspect

- Device able to support (enhanced|future) key lengths needed in, say, 15 years?
- Relationship between warranty and lifetime
  - *Liability* might come into play, too.
- More importantly: relationship between warranty and availability of patches.
  - What happens **after** warranty ends (but device still sitting around & active)?

Lifespan Aspects (II)

High chance that device is always turned on.

- But this fact (or device’s existence at all) unbeknownst to the humans around.

- Their owners might even forget “they‘re there“.
  - What happens during a relocation? Why would you care?
  - Will this be the first step when you move into new home?
  - Who‘s “the owner“ anyway?

The Update Process

… has several dimensions:

- ○ Technical
- ○ Physical interaction
- ○ Socio-economic
- ○ Ethical

The Technical Dimension

- Infrastructure
  - Capabilities (to update n systems within time x)
- Crypto
- 3rd party libraries ?!
  - Understand their sec history, status of (project) maintenance, "patchability".
- Device drivers (blobs ⇔ patchability)
- It's a complex ecosystem anyway (see Android)
  - Main question: who's responsible? => supply chain/contracts

Crypto

- Use "reasonably strong crypto" from the beginning
  - Will it still be reasonably secure in 10 years?
  - Can you upgrade the capability
    - With respect to the architecture?
    - With respect to the hardware?

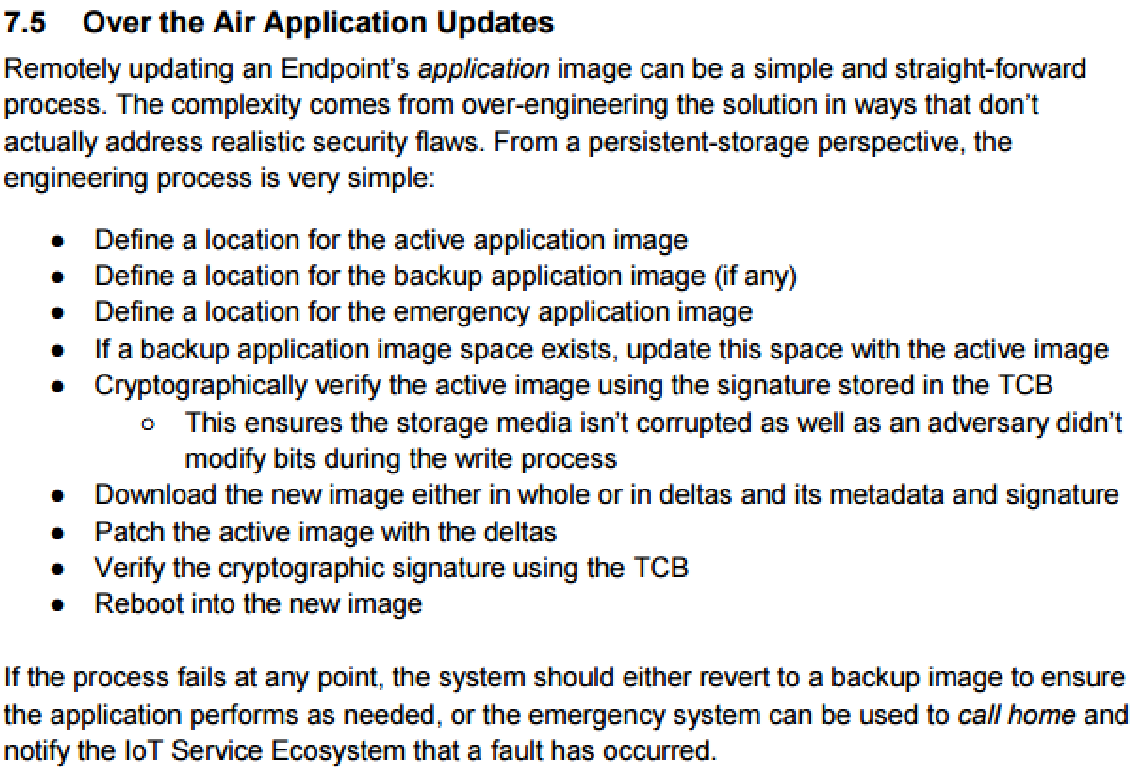

7.5 Over the Air Application Updates

Remotely updating an Endpoint's *application* image can be a simple and straight-forward process. The complexity comes from over-engineering the solution in ways that don't actually address realistic security flaws. From a persistent-storage perspective, the engineering process is very simple:

- Define a location for the active application image
- Define a location for the backup application image (if any)
- Define a location for the emergency application image
- If a backup application image space exists, update this space with the active image
- Cryptographically verify the active image using the signature stored in the TCB
  - This ensures the storage media isn't corrupted as well as an adversary didn't modify bits during the write process
- Download the new image either in whole or in deltas and its metadata and signature
- Patch the active image with the deltas
- Verify the cryptographic signature using the TCB
- Reboot into the new image

If the process fails at any point, the system should either revert to a backup image to ensure the application performs as needed, or the emergency system can be used to *call home* and notify the IoT Service Ecosystem that a fault has occurred.

1974

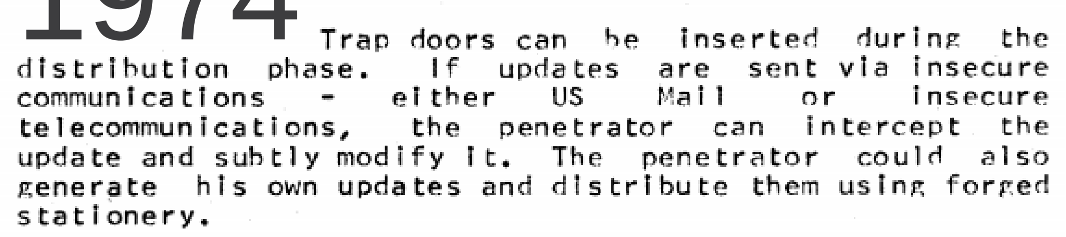

Trap doors can be inserted during the distribution phase. If updates are sent via insecure communications - either US Mail or insecure telecommunications, the penetrator can intercept the update and subtly modify it. The penetrator could also generate his own updates and distribute them using forged stationery.

The Time Dimension

- Once a “new“ vulnerability is made public there might be several stages
  - No patch available at all.
  - Patch available but not tested yet, for specific platform (of vendor).

(The Dimension Of) Physical Interaction

- Which consequences for the physical environment can come from
  - A reboot
  - A crash (due to something going wrong during update)
  - How do you notice such events? What to do once device does *not* come up again after reboot?
- In case an update requires user intervention/confirmation
  - How to design this (process-/interface-wise)?
  - How to handle cases where the confirmation does not happen within $TIMEFRAME?
  - How/when to announce availability of a patch (e.g. while driving)?
- Proper risk analysis needed here, in early stage.

The Socio-Economic Dimension

- Role of users. Are they supposed/allowed to patch on their own?
- For vendor, providing updates mean efforts (costs!)
  - How long do you want to do this for $PRODUCT?
  - How do you announce this decision/information?
  - Relationship w/ warranty?
  - Relationship w/ liability?
- What about a vendor-side kill-switch?
  - Tell customers about this or not?
  - Will they appreciate the capability?
- How will the above be affected by an M&A?

The Ethical Dimension

- Is there a moral obligation to protect
  - Users ?
  - “The Internet“ ?
  - If so (any of those), how long?

- Let‘s assume there‘s a (crypto) master key for the updates.
  - Which parties to share that one with?

- Same question for the potential kill-switch.

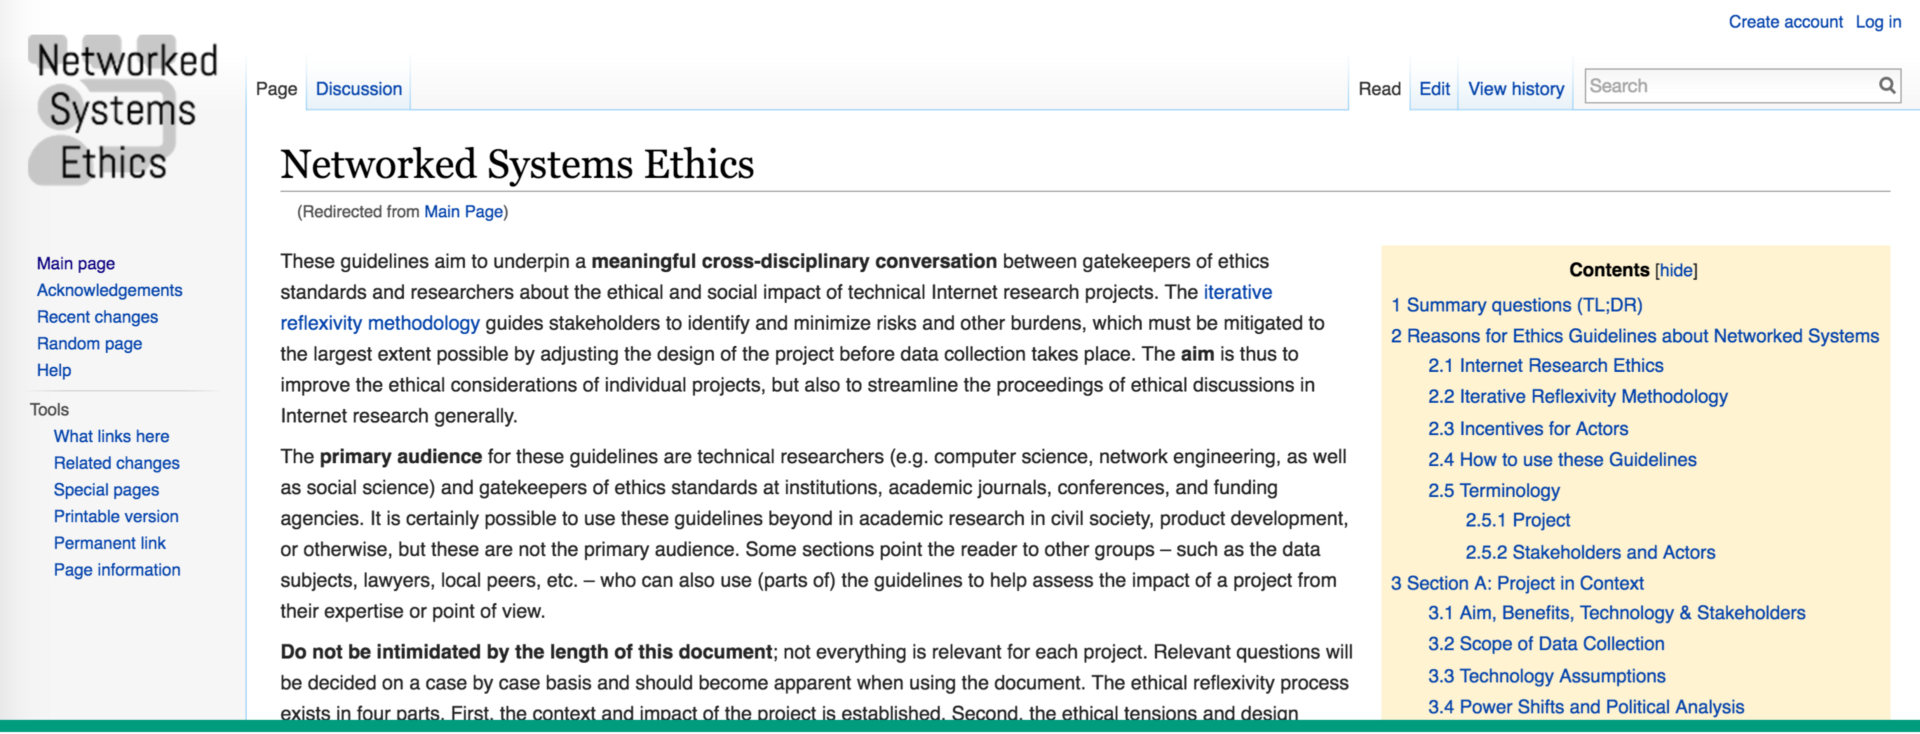

Networked Systems Ethics

(Redirected from Main Page)

These guidelines aim to underpin a **meaningful cross-disciplinary conversation** between gatekeepers of ethics standards and researchers about the ethical and social impact of technical Internet research projects. The iterative reflexivity methodology guides stakeholders to identify and minimize risks and other burdens, which must be mitigated to the largest extent possible by adjusting the design of the project before data collection takes place. The **aim** is thus to improve the ethical considerations of individual projects, but also to streamline the proceedings of ethical discussions in Internet research generally.

The **primary audience** for these guidelines are technical researchers (e.g. computer science, network engineering, as well as social science) and gatekeepers of ethics standards at institutions, academic journals, conferences, and funding agencies. It is certainly possible to use these guidelines beyond in academic research in civil society, product development, or otherwise, but these are not the primary audience. Some sections point the reader to other groups – such as the data subjects, lawyers, local peers, etc. – who can also use (parts of) the guidelines to help assess the impact of a project from their expertise or point of view.

Do not be intimidated by the length of this document; not everything is relevant for each project. Relevant questions will be decided on a case by case basis and should become apparent when using the document. The ethical reflexivity process exists in four parts. First, the context and impact of the project is established. Second, the ethical tensions and design

Contents [hide]

Ethics – My Ceterum Censeo

Discussion