DOWN With Trusted Devices

Mahalingam Ramkumar
Department of Computer Science and Engineering
Mississippi State University, Mississippi State, MS 39762.

Justification

Paper Category: Research Paper

Many evolving application scenarios require *mutual co-operation* of massively interconnected devices or computers. This calls for increased reliance of human beings on "trusted" devices - devices which are trusted to co-operate with each other, by strictly adhering to pre-imposed rules. A fundamental requirement for a trusted device is the ability of the device to protect its secrets.

The motivation of the paper is two fold. The first is a high level architecture for trusted computers, employing a simple "common-sense-driven" security policy - the **d**ecrypt **o**nly **w**hen **n**ecessary (DOWN) policy. We argue that DOWN, used in conjunction with physical unclonable functions (PUFs) is a promising approach for realization trusted computers. The DOWN policy provides an assurance that an attacker can compromise at most a small part of the secret(s) stored in trusted devices. With the DOWN policy (in situations where PUFs are deemed expensive) non-volatile storage (except one CPU register which needs "rest" protection) does not need *any* protection. When used in conjunction with PUFs, it eliminates the need to protect the non-volatile storage register too - or no *tamper-sensing circuits and battery back-ups are needed* when the device is powered-off.

The second motivation of the paper is to put-forth arguments as to why key pre-distribution schemes, especially random key pre-distribution schemes may be more suitable for securing interactions between trusted devices (as opposed to PKI), especially when used in conjunction with the DOWN policy.

While it is perhaps commonly accepted that key pre-distribution schemes may be suitable for a limited range of applications (like severely resource constrained sensor networks where asymmetric cryptography may simply not be feasible)[13], we argue that they are more suitable than PKI for huge scale networks, *even if resource constraint is not a serious issue*. The main disadvantage of KPDs is that they provide only conditional guarantees - secure only when less than say n nodes have been compromised. Random KPDs, further provide only probabilistic guarantees. After all, 128-bit encryption is also in-secure with a probability $2^{-128} \approx 10^{-39}$. Probabilistic guarantees are not necessarily inferior, as long as the probabilities of compromise are small.

When used in conjuntion with the DOWN policy (which provides an assurance that not more than one secret can be compromised by tampering with a device), the security offered by KPDs is increased dramatically - to the point that it may be practically impossible for an attacker to compromise KPDs.

While any KPD has obvious advantages over PKI in terms of computational complexity and bandwidth overheads, we argue that there are many not-so-obvious advantages. Examples of these include larger scalability of KPDs as they are not plagued by the problem of public-key dissemination. Another advantage of KPDs is their ability to cater for multicast security, which is not "native" to PKI. For very large scale networks, performing revocation by broadcasting revocation certificates may be impractical. Broadcast encryption might therefore be a very important requirement (for revocation by broadcasting *secrets* to non-revoked devices), for very large scale deployments.

One of the very important perceived advantage of PKI is the comfort of *privacy* - the knowledge that the private key can be protected from "authorities" (while this is not possible in theory for KPDs). However, with trusted devices which will be expected to *self-destruct*, on suspicion of tampering attempts (and there are of course, bound to be false-alarms), unless secrets are *escrowed* the situation may be disastrous for users. The potential need for escrowing secrets further reduces the advantages that PKI has over KPDs.

While it is well known that KPDs can be arbitrarily secure by increasing the complexity (primarily increasing the number of keys), the corresponding need for increased *secure* storage is seen as a major deterrent. However, with the DOWN policy, the non-volatile storage used for storing the keys is assumed to be totally unprotected. Increased complexity of insecure storage may not be a crucial issue in practice, and KPDs could therefore be made as secure as we desire. Further for random KPDs, increasing the storage size does not directly translate to increased computational complexity.

[13] Most recent research on key pre-distribution schemes in the literature are restricted to sensor network scenarios.