A System for Protecting CRUTIAL Things[*]
(Extended Abstract)

Alysson Neves Bessani   Paulo Sousa   Miguel Correia   Nuno Ferreira Neves   Paulo Verissimo

LASIGE, Faculdade de Ciências da Universidade de Lisboa – Portugal

Today's critical infrastructures like the Power Grid are essentially physical processes controlled by computers connected by networks. Once these systems were highly isolated and secure against most security threats. However, in recent years they evolved in several aspects that greatly increased their exposure to cyber-attacks coming from the Internet. Firstly, the computers, networks and protocols in those control systems are no longer proprietary but standard PCs and networks (e.g., wired and wireless Ethernet), and the protocols are often encapsulated on top of TCP/IP. Secondly, these networks are usually connected to the Internet indirectly through the corporate network or to other networks using modems and data links. Therefore these infrastructures have a level of vulnerability similar to other systems connected to the Internet, but the socio-economic impact of their failure can be huge. This scenario, reinforced by several recent incidents, is generating a great concern about the security of these infrastructures, especially at government level.

Recently, we proposed a reference architecture to protect critical infrastructures, in the context of the CRUTIAL[1] EU-IST project [2]. The idea is to model the whole infrastructure as a WAN-of-LANs, where the typical facilities that compose it (like power transformation substations or corporate offices) are modeled as collections of LANs interconnected by a wider-area network (WAN). Using this architecture, we reduce the problem of critical infrastructures protection to the problem of protecting LANs from the WAN or other LANs. In consequence, our model and architecture allow us to deal both with outsider threats (protecting a facility from the Internet) and insider threats (protecting a critical host from other hosts in the same physical facility, by locating them in different LANs).

Here, we introduce a device for protecting LANs called *CRUTIAL Information Switch* (CIS). A fundamental service provided by CIS is the *Protection Service*, which ensures that the incoming and outgoing traffic in/out of the LAN satisfies the security policy of the infrastructure. However, a CIS can not be a simple firewall since that would put the critical infrastructure at most at the level of security of current (corporate) Internet systems, which is not acceptable since intrusions in those systems are constantly being reported. Instead, the CIS have many distinguished characteristics. Firstly, it has similarities to a *distributed firewall*, since CIS can be deployed not only on the network border but inside the networks to better protect critical equipment. Secondly, the CIS uses a *rich access control model* that takes into account the involvement of different organizations and allows the access control rules to depend on context information. Thirdly, the CIS is *intrusion-tolerant*, i.e., it operates correctly even if there are intrusions in some of its components and withstands a high degree of such hostility from the environment.

In this work we address specifically the last topic. The intrusion tolerant CIS is replicated in a set of $n \geq 2f + 1$ machines. Each *CIS replica* receives all packets to and from the LAN and verifies if this packet satisfies some pre-defined application-level policy. The difficult point here is to ensure that intrusions (modelled as Byzantine faults) in at most f of the replicas are masked, i.e., that all valid packets are accepted and all invalid packets are dropped. The CIS design presents two very interesting challenges that make it essentially different from other Byzantine fault-tolerant services. The first is that a firewall-like component has to be transparent to protocols that pass through it, so it can not modify the protocols themselves to obtain intrusion tolerance. This also means that recipient nodes will ignore any internal CIS intrusion-tolerance mechanisms, and as such they can not protect themselves from messages forwarded by faulty replicas not satisfying the security policy.

These two challenges are solved through the use of wormholes [1]: we assume that each replica of the CIS has a trusted component that cannot be corrupted. These *local wormholes* are also connected through an isolated network. Every message approved by a replica is issued to the wormhole to be signed. The local wormholes vote between thenselves and, if the message is approved by at least $f + 1$ replicas, they are signed using a secret key installed in the trusted component. Once the message is signeded, one of the replicas (the leader) is responsible for forwarding the approved message to its destination. Failure detection, leader election and proactive recovery are other services provided by the wormhole.

We have implemented a prototype of this system in a VM-based setting and an evaluation showed that the mechanisms overhead is acceptable when considering the Power Grid requirements.

References

[1] P. Verissimo. Travelling through wormholes: a new look at distributed systems models. *SIGACT News*, 37(1), 2006.

[2] P. Verissimo, N. F. Neves, and M. Correia. CRUTIAL: The blueprint of a reference critical information infrastructure architecture. In *Proc. of CRITIS'06 1st Int. Workshop on Critical Information Infrastructures Security*, Aug. 2006.

[1]Critical UTility InfrastructurAL Resilience: `http://crutial.cesiricerca.it`.