

Javier Lopez (Ed.)

LNCS 4347

Critical Information Infrastructures Security

First International Workshop, CRITIS 2006
Samos, Greece, August/September 2006
Revised Papers



Springer

Javier Lopez (Ed.)

Critical Information Infrastructures Security

First International Workshop, CRITIS 2006
Samos, Greece, August 31 - September 1, 2006
Revised Papers



Springer

Volume Editor

Javier Lopez
University of Malaga
Computer Science Department
E.T.S.I. Informatica
Campus Teatinos, 29071 Malaga, Spain
E-mail: jlm@lcc.uma.es

Library of Congress Control Number: 2006938669

CR Subject Classification (1998): C.2, D.4.6, E.3, K.6.5, K.4.1, K.4.4, J.1

LNCS Sublibrary: SL 5 – Computer Communication Networks and
Telecommunications

ISSN	0302-9743
ISBN-10	3-540-69083-2 Springer Berlin Heidelberg New York
ISBN-13	978-3-540-69083-2 Springer Berlin Heidelberg New York

This work is subject to copyright. All rights are reserved, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, re-use of illustrations, recitation, broadcasting, reproduction on microfilms or in any other way, and storage in data banks. Duplication of this publication or parts thereof is permitted only under the provisions of the German Copyright Law of September 9, 1965, in its current version, and permission for use must always be obtained from Springer. Violations are liable to prosecution under the German Copyright Law.

Springer is a part of Springer Science+Business Media
springer.com

© Springer-Verlag Berlin Heidelberg 2006
Printed in Germany

Typesetting: Camera-ready by author, data conversion by Scientific Publishing Services, Chennai, India
Printed on acid-free paper SPIN: 11962977 06/3142 5 4 3 2 1 0

Preface

Key sectors of modern economies depend highly on ICT. The information flowing through the resulting technological super-infrastructure as well as the information being processed by the complex computing systems that underpin it becomes crucial because its disruption, disturbance or loss can lead to high financial, material and, sometimes, human loss. As a consequence, the security and dependability of this infrastructure become critical and its protection a major objective for governments, companies and the research community.

CRITIS has been born as an event that aims to bring together researchers and professionals from universities, private companies and public administrations interested or involved in all security-related heterogeneous aspects of critical information infrastructures.

This volume contains the proceedings of the 1st International Workshop on Critical Information Infrastructure Security (CRITIS 2006), that was held between August 31 and September 1, 2006 on Samos, Greece, and was hosted by the University of the Aegean, Department of Information and Communication Systems Engineering, Laboratory of Information and Communication Systems Security (Info-Sec-Lab).

In response to the CRITIS 2006 call for papers, 57 papers were submitted. Each paper was reviewed by three members of the Program Committee, on the basis of significance, novelty, technical quality and relevance to critical infrastructures. At the end of the reviewing process, only 22 papers were selected for presentation, resulting in an acceptance rate of 38%. Revisions were not checked and the authors bear full responsibility for the content of their papers.

Additionally, CRITIS 2006 was fortunate to have Andrea Servida, Deputy Head of Unit of the European Commission (Information and Society and Media Directorate General) as invited speaker, giving the talk “Security and Resilience in Information Society: The European Approach.” I thank him very much for his contribution.

Other persons deserve many thanks for their support and contribution to the success of the conference. Sokratis Katsikas and Reinhard Posch were General Co-chairs, and Stefanos Gritzalis, a driving force of the event, was Organization Chair. I sincerely thank them for their total support and encouragement, and for their contribution to all organizational issues. My special thanks to Rodrigo Roman for preparation and maintenance of the Workshop Web site, and to Costas Lambrinoudakis, George Kambourakis, Dimitris Geneiatakis, Giorgos Karopoulos and Irene Gonidelli for their help in the organizational tasks. Without the hard work of these colleagues and the rest of the local organization team, this conference would not have been possible. Finally, I would like to thank all the authors who submitted papers and the participants from all over the world who chose to honor us with their attendance.

August 2006

Javier Lopez
Program Chair

CRITIS 2006

1st International Workshop on Critical Information Infrastructures Security

Samos, Greece
August 31 – September 1, 2006

Organized by
Department of Information and Communication Systems Engineering,
Laboratory of Information and Communication Systems Security (Info-Sec-Lab)
University of the Aegean
Greece

General Co-chairs

Sokratis Katsikas
Reinhard Posch

University of the Aegean, Greece
Technical University of Graz, Austria

Organization Chair

Stefanos Gritzalis

University of Aegean, Greece

Program Chair

Javier Lopez

University of Malaga, Spain

Program Committee

Marc Dacier
George Davida
Ed Dawson
Yvo Desmedt
Myriam Dunn
Claudia Eckert
Steven Furnell
Urs Gattiker
Adrian Gheorghe
Eric Goetz
Juan M. Gonzalez-Nieto
John Griffin
Stefanos Gritzalis
Dieter Gollmann
Bernhard M. Hämmerli

Institut Eurécom, France
University of Wisconsin-Milwaukee, USA
QUT, Australia
University College London, UK
ETH Zurich, Switzerland
Fraunhofer-SIT, Germany
University of Plymouth, UK
CyTRAP-RiskIT, Switzerland
ETH Zurich, Switzerland
Dartmouth College, US
QUT, Australia
IBM T.J. Watson Research Center, USA
University of the Aegean, Greece
TU Hamburg, Germany
HTA Lucerne, Switzerland

VIII Organization

Tom Karygiannis	NIST, USA
Håkan Kvarnström	TeliaSonera, Sweden
Diego Lopez	RedIRIS, Spain
Eric Luijff	TNO, Netherlands
Masahiro Mambo	University of Tsukuba, Japan
Fabio Martinelli	CNR, Italy
Catherine Meadows	Naval Research Lab., USA
Simin Nadjm-Tehrani	Linköping University, Sweden
Peter Neumann	SRI, USA
Eiji Okamoto	University of Tsukuba, Japan
Andrew Powell	NISCC, UK
Kai Rannenberg	Goethe University Frankfurt, Germany
Michel Riguidel	ENST, France
Rodrigo Roman	University of Malaga, Spain
Roberto Setola	Univ. Campus Bio-Medico di Roma, Italy
Stefaan Seys	Katholieke Universiteit Leuven, Belgium
Sujeet Sheno	University of Tulsa, USA
Stephen D. Wolthusen	Royal Holloway, UK
Moti Yung	Columbia University, USA
Yuliang Zheng	University of North Carolina, USA
Jianying Zhou	Institute for Infocomm Research, Singapore

External Reviewers

Efthimia Aivaloglou, Elisavet Constantinou, Alessandro Falleni, Lothar Fritsch, Matt Henricksen, Michaela Iorga, Spyros Kokolakis, George Mohay, Maria Papadaki, Heiko Rossnagel, Dries Schellekens, Falk Wagner.

Table of Contents

CRUTIAL: The Blueprint of a Reference Critical Information Infrastructure Architecture	1
<i>Paulo Veríssimo, Nuno Ferreira Neves, and Miguel Correia</i>	
Experiment Based Validation of CIIP	15
<i>Per Mellstrand and Rune Gustavsson</i>	
Security Requirements Model for Grid Data Management Systems	30
<i>Syed Naqvi, Philippe Massonet, and Alvaro Arenas</i>	
Assessing the Risk of an Information Infrastructure Through Security Dependencies	42
<i>F. Baiardi, S. Swin, C. Telmon, and M. Pioli</i>	
Modelling Risk and Identifying Countermeasure in Organizations	55
<i>Yudistira Asnar and Paolo Giorgini</i>	
Modelling and Analysing Network Security Policies in a Given Vulnerability Setting	67
<i>Roland Rieke</i>	
A Framework for Conceptualizing Social Engineering Attacks	79
<i>Jose J. Gonzalez, Jose M. Sarriegi, and Alazne Gurrutxaga</i>	
An Overview of R&D Activities in Europe on Critical Information Infrastructure Protection (CIIP)	91
<i>Sandro Bologna, Giovanni Di Costanzo, Eric Luijff, and Roberto Setola</i>	
Intelligent Network-Based Early Warning Systems	103
<i>Karsten Bsufka, Olaf Kroll-Peters, and Sahin Albayrak</i>	
Can an Early Warning System for Home Users and SMEs Make a Difference? A Field Study	112
<i>Urs E. Gattiker</i>	
Protection of Components Based on a Smart-Card Enhanced Security Module	128
<i>Joaquín García-Alfaro, Sergio Castillo, Jordi Castellà-Roca, Guillermo Navarro, and Joan Borrell</i>	
Revisiting Colored Networks and Privacy Preserving Censorship	140
<i>Yvo Desmedt, Yongge Wang, and Mike Burmester</i>	

PROSEARCH: A Protocol to Simplify Path Discovery in Critical Scenarios	151
<i>Cristina Satizábal, Rafael Páez, and Jordi Forné</i>	
Applying Key Infrastructures for Sensor Networks in CIP/CIIP Scenarios	166
<i>Cristina Alcaraz and Rodrigo Roman</i>	
Trust Establishment in Ad Hoc and Sensor Networks	179
<i>Efthimia Aivaloglou, Stefanos Gritzalis, and Charalabos Skianis</i>	
Enforcing Trust in Pervasive Computing with Trusted Computing Technology	195
<i>Shiqun Li, Shane Balfe, Jianying Zhou, and Kefei Chen</i>	
Proposals on Assessment Environments for Anomaly-Based Network Intrusion Detection Systems	210
<i>M. Bermúdez-Edo, R. Salazar-Hernández, J. Díaz-Verdejo, and P. García-Teodoro</i>	
High-Speed Intrusion Detection in Support of Critical Infrastructure Protection	222
<i>Salvatore D'Antonio, Francesco Oliviero, and Roberto Setola</i>	
Rational Choice of Security Measures Via Multi-parameter Attack Trees	235
<i>Ahto Buldas, Peeter Laud, Jaan Priisalu, Märt Saarepera, and Jan Willemson</i>	
Multidomain Virtual Security Negotiation over the Session Initiation Protocol (SIP)	249
<i>Daniel J. Martínez-Manzano, Gabriel López, and Antonio F. Gómez-Skarmeta</i>	
Vulnerabilities and Possible Attacks Against the GPRS Backbone Network	262
<i>Christos Xenakis and Lazaros Merakos</i>	
A Framework for Secure and Verifiable Logging in Public Communication Networks	273
<i>Vassilios Stathopoulos, Panayiotis Kotzanikolaou, and Emmanouil Magkos</i>	
Author Index	285

CRUTIAL: The Blueprint of a Reference Critical Information Infrastructure Architecture*

Paulo Veríssimo, Nuno Ferreira Neves, and Miguel Correia

University of Lisboa, Faculty of Sciences
Bloco C6, Campo Grande, 1749-016 Lisboa - Portugal
`{pjb,nuno,mpc}@di.fc.ul.pt`
<http://www.navigators.di.fc.ul.pt>

Abstract. In the past few decades, critical infrastructures have become largely computerised and interconnected all over the world. This generated the problem of achieving resilience of critical information infrastructures against computer-borne attacks and severe faults. Governments and industry have been pushing an immense research effort in information and systems security, but we believe the complexity of the problem prevents it from being solved using classical security methods.

The paper focuses on the computer systems behind electrical utility infrastructures. It proposes the blueprint of a distributed systems architecture that we believe may come to be useful as a reference for modern critical information infrastructures in general. The architecture is instantiated with a set of classes of techniques and algorithms, based on paradigms providing resilience to faults and attacks in an automatic way.

1 Introduction

The largely computerised nature of critical infrastructures on the one hand, and the pervasive interconnection of systems all over the world, on the other hand, have generated one of the most fascinating current problems of computer science and control engineering: *how to achieve resilience of critical information infrastructures*.

This problem is concerned with ensuring acceptable levels of service and, in last resort, the integrity of systems themselves, when faced with threats of several kinds. In this paper we are concerned with threats against computers and control computers, not the physical infrastructures themselves. These threats range from accidental events like natural faults or wrong manoeuvres [15, 23], to attacks by hackers or terrorists [5, 12, 14, 17, 28]. The problem affects systems with great socio-economic value, such as utility systems like electrical, gas or water, or telecommunication systems and computer networks like the Internet. In consequence, the high degree of interconnection is causing great concern, given

* This work was mainly supported by the EC, through project IST-4-027513-STP (CRUTIAL), and also by the FCT, through LASIGE and projects POSI/EIA/61643/2004 (AJECT) and POSI/EIA/60334/2004 (RITAS).

the level of exposure of very high value systems and components to attacks that can be perpetrated in an anonymous and remote way.

Although there is an increase in the concern for using security best practices in these systems [2, 4], we believe that the problem is not completely understood, and can not be solved with classical methods. Its complexity is mainly due to *the hybrid composition of those infrastructures*:

- The operational network, called generically SCADA (Supervisory Control and Data Acquisition)¹, composed of the computer systems that yield the operational ability to supervise, acquire data from, and control the physical processes. In fact, to the global computer system, SCADA computer systems (e.g., controllers) “are” the controlled processes (e.g., power generators), since by acting on the former, for example, through a network message, one changes the state of the latter.
- The corporate intranet, where usual departmental services (e.g., web, email, databases) and clients reside, and also the engineering and technical staff, who access the SCADA part through ad-hoc interconnections².
- The Internet, through which intranet users get to other intranets and/or the outside world, but to which, and often unwittingly, the SCADA network is sometimes connected to.

Besides the complexity due to this hybrid composition, this mixture has given an unexpected *inter-disciplinary nature* to the problem: SCADA systems are real-time systems, with some reliability and fault tolerance concerns, but they were classically not designed to be widely distributed or remotely accessed, let alone open to other more asynchronous and less trusted subsystems. Likewise, they were not designed with security in mind. In consequence, in scientific terms, our problem can be formulated as follows:

- The computer-related operation of a critical utility infrastructure is a distributed systems problem including interconnected SCADA/embedded networks, corporate intranets, and Internet/PSTN³ access subsystems.
- This distributed systems problem is hard, since it simultaneously includes facets of real-time, fault tolerance, and security.

In this paper, we focus on the computer systems behind electrical utility infrastructures as an example, and we propose: (1) *the blueprint of a distributed systems architecture* that we believe may come to be useful as a reference for modern critical information infrastructures; (2) *a set of classes of techniques and algorithms* based on paradigms providing resilience to faults and attacks

¹ Or PCS (Process Control System).

² In some companies there is a (healthy) reluctance against interconnecting SCADA networks and the corporate network or the Internet. However, in practice this interconnection is a reality in many companies all over the world. We believe this is indeed the situation in most companies and this is the case we are interested in this paper.

³ Public Switched Telephone Network.

in an automatic way. This work is ongoing and is done in the context of the recently started CRUTIAL European project, CRITICAL UTILITY InfrastructurAL resilience [6], details of which are given in the end.

As a final note, whilst it is usual to use the designation “critical information infrastructures” to denote the computer related part of the physical critical infrastructures, we do not make a differentiation of the two in this paper.

2 Rationale for the Model and Architecture

Before we proceed, let us bring some further insight on the problem of critical infrastructures:

- Critical Information Infrastructures (CII) feature a lot of legacy subsystems and non-computer-standard components (controllers, sensors, actuators, etc.).
- Conventional security and protection techniques, when directly applied to CII controlling devices, sometimes stand in the way of their effective operation.

These two facts will not change, at least for a long time, so they should be considered as additional research challenges. Despite security and dependability concerns with those individual components being a necessity, we believe that the crucial problem is with the forest, not the trees. That is, the problem of critical information infrastructure insecurity is mostly created by the informatics nature of many current infrastructures, and by the generic and non-structured network interconnection of CII, which bring several facets of exposure, from internal unprotected wireline or wireless links, to interconnections of SCADA and corporate intranets to the Internet and PSTN. This situation is conspicuous in several of the attacks reported against CII. For instance, the attack of the Slammer worm against the Davis-Besse nuclear power plant (US) was due both to this combination of a computerised CII with non-structured network interconnections and lack of protection [8]. Although the network was protected by a firewall, the worm entered through a contractor’s computer connected to the CII using a telephone line.

The problems that may result from this exposure to computer-borne threats range from wrong manoeuvring to malicious actions coming from terminals located outside, somewhere in the Internet. The potential targets of these actions are computer control units, embedded components and systems, that is, devices connected to operational hardware (e.g., water pumps and filters, electrical power generators and power protections, dam gates, etc.) or to telecom hardware (core routers, base stations, etc.). The failure perspectives go from unavailability of services supposed to operate 24×7, to physical damage to infrastructures. In the electrical power provision these situations have already been witnessed [6]: among the blackouts that occurred in several countries during the summer of 2003, the analysis report [7] of the North American one highlighted the failure of various information systems as having thwarted the utility workers’ ability to contain the blackout before it cascaded out of control, leading to an escalating failure.

Whilst it seems non-controversial that such a status quo brings a certain level of threat, we know of no work that has tried to equate the problem by defining a reference model of a *critical information infrastructure distributed systems architecture*, providing the necessary global resilience against abnormal situations.

We believe that evaluation work based on such a model will let us learn about activity patterns of interdependencies, which will reveal the potential for far more damaging fault/failure scenarios than those that have been anticipated up to now. Moreover, such a model will be highly constructive, for it will form a structured framework for: conceiving the right balance between prevention and removal of vulnerabilities and attacks, and tolerance of remaining potential intrusions and designed-in faults.

What can be done at architectural level to achieve resilient operation? Note that the crux of the problem lies with the fact that access to operational networks, such as remote SCADA manoeuvring, ended up entangled with access to corporate intranets and to public Internet, without there being computational and resilience models that *represent* this situation, unlike what exists in simpler, more homogeneous settings, e.g. classical web-based server infrastructures on Internet. Our point is that interference and threats start at the level of the macroscopic information flows between these subsystems, and can in consequence be stopped there. This should not prevent the study of techniques at the controller level, but in this paper we will not focus on this latter issue.

Now, given the simultaneous need for real-time, security and fault tolerance, this problem is hard vis-a-vis existing paradigms. For example, many classical distributed systems paradigms handle each of those facets separately, and just solve part of the problem. A unifying approach has gained impressive momentum currently: *intrusion tolerance* [27]. In short, instead of trying to prevent every single intrusion or fault, they are allowed, but tolerated: systems remain to some extent faulty and/or vulnerable, attacks on components can happen and some will be successful, but the system has the means to trigger automatic mechanisms that prevent faults or intrusions from generating a system failure.

Our approach is thus equated along the following propositions:

Proposition 1. Classical security and/or safety techniques alone will not solve the problem: they are largely based on prevention, intrusion detection and ad-hoc recovery or ultimately disconnection.

There is a recent and positive trend to make SCADA systems and CIIIs at large more secure [2, 4, 12, 20, 21]. However, classic engineering remedies place real-time and embedded (RTE) systems at most at the current level of commercial systems' security and dependability, which is known to be insufficient [5, 9, 22]: systems constantly suffer attacks, intrusions, some of them massive (worms); most defences are dedicated to generic non-targeted attacks; attacks degrade business but only do virtual damage, unlike RTE systems where there is a risk of great social impact and even physical damage. On the other hand, some current IT security techniques can negatively affect RTE system operation, w.r.t.

availability and timeliness. For example, if security is based on disconnection, significant performance degradation, or even defensive restrictions can prevent the actuation or monitoring of the infrastructure.

Proposition 2. Any solution, to be effective, has to involve automatic control of macroscopic command and information flows, occurring essentially between the physical or virtual LANs⁴ composing the critical information infrastructure architecture, with the purpose of securing appropriate system-level properties.

We believe that a key to the solution lies with controlling the command and information flow at macroscopic level (organisation-level). We are talking about an architectural model, a set of architectural devices, and key algorithms, capable of achieving the above-mentioned control of the command and information flow. The devices and algorithms should be capable of securing a set of system-level properties characterising whatever is meant by correct and resilient behaviour.

Proposition 3. We lack a reference architecture of “modern critical information infrastructure” considering different interconnection realms and different kinds of risk, throughout the physical and the information subsystems of a CII.

We must consider the physical or virtual LANs composing the operational SCADA/embedded networks, the corporate intranets, and the Internet/PSTN access networks, as different first order citizens of the architecture. Likewise, the notion that risk factors may vary and be difficult to perceive accurately, brings the need to reconcile uncertainty with predictability in architecture and algorithmics.

3 CRUTIAL Architecture

The CRUTIAL architecture encompasses:

- Architectural configurations featuring trusted components in key places, which a priori induce prevention of some faults, and of certain attack and vulnerability combinations.
- Middleware devices that achieve runtime automatic tolerance of remaining faults and intrusions, supplying trusted services out of non-trustworthy components.
- Trustworthiness monitoring mechanisms detecting situations not predicted and/or beyond assumptions made, and adaptation mechanisms to survive those situations.
- Organisation-level security policies and access control models capable of securing information flows with different criticality within/in/out of a CII.

⁴ Local Area Networks.

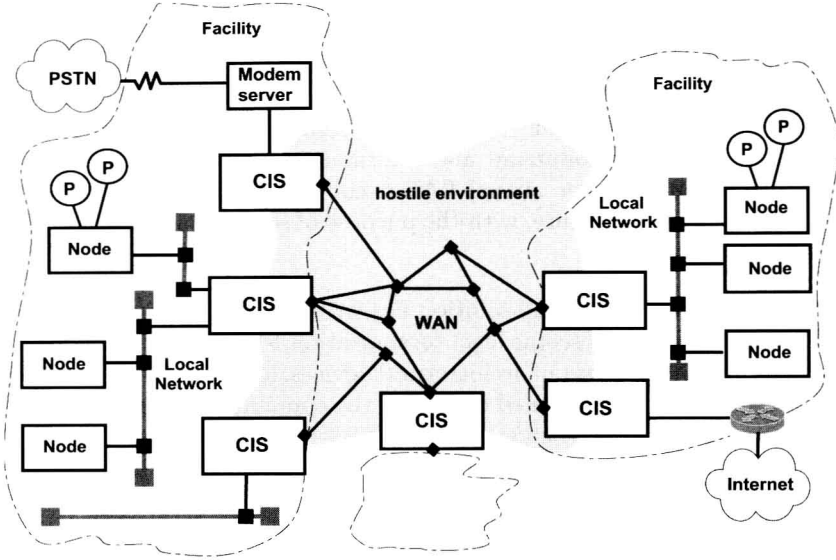


Fig. 1. CRUTIAL overall architecture (WAN of LANs connected by CIS, P processes live in the several nodes)

We build on results from the MAFTIA project⁵ in this field [26], but extend them significantly to attend the specific challenges of the critical information infrastructure problem, for example, timeliness, global access control, and above all non-stop operation and resilience.

Given the severity of threats expected, some key components are built using architectural hybridisation methods in order to achieve *trusted-trustworthy* operation [26]: an architectural paradigm whereby components prevent the occurrence of some failure modes *by construction*, so that their resistance to faults and hackers can justifiably be trusted. In other words, some special-purpose components are constructed in such a way that we can argue that they are always secure, so that they can provide a small set of services useful to support intrusion tolerance in the rest of the system.

Intrusion tolerance mechanisms are selectively used in the CRUTIAL architecture, to build layers of progressively more trusted components and middleware subsystems, from baseline untrusted components (nodes, networks) [26]. This leads to an automation of the process of building trust: for example, at lower layers, basic intrusion tolerance mechanisms are used to construct a trustworthy communication subsystem, which can then be trusted by upper layers to securely communicate amongst participants without bothering about network intrusion threats.

One of the innovative aspects of this work, further to intrusion tolerance, is the resilience aspect, approached through two paradigms: *proactive-resilience* to

⁵ Malicious-and Accidental-Fault Tolerance for Internet Applications. The web site of the project is at www.maftia.org

achieve exhaustion-safety [18], to ensure perpetual, non-stop operation despite the continuous production of faults and intrusions; and *trustworthiness monitoring* to perform surveillance of the coverage stability of the system, that is, of whether it is still performing inside the assumed fault envelope or beyond assumptions made [3]. In the latter case, dependable adaptation mechanisms are triggered.

Finally, the desired control of the information flows is partly performed through protection mechanisms using an adaptation of *organisation-based access control* models [10] for implementing global-level security policies.

The mechanisms and algorithms in place achieve system-level properties of the following classes: trustworthiness or resistance to faults and intrusions (i.e., security and dependability); timeliness, in the sense of meeting timing constraints raised by real world control and supervision; coverage stability, to ensure that variation or degradation of assumptions remains within a bounded envelope; dependable adaptability, to achieve predictability in uncertain conditions; resilience, read as correctness and continuity of service even beyond assumptions made.

3.1 Main Architectural Options

We view the system as a WAN-of-LANs, as introduced in [24]. There is a global interconnection network, the WAN, that switches packets through generic devices that we call *facility gateways*, which are the representative gateways of each LAN (the overall picture is shown in Figure 1). The WAN is a logical entity operated by the CII operator companies, which may or may not use parts of public network as physical support. A LAN is a logical unit that may or may not have physical reality (e.g., LAN segments vs. VLANs⁶). More than one LAN can be connected by the same facility gateway. All traffic originates from and goes to a LAN. As example LANs, the reader can envision: the administrative clients and the servers LANs; the operational (SCADA) clients and servers LANs; the engineering clients and servers LANs; the PSTN modem access LANs; the Internet and extranet access LANs, etc.

The facility gateways of a CRUTIAL critical information infrastructure are more than mere TCP/IP routers. Collectively they act as a set of servers providing distributed services relevant to solving our problem: *achieving control of the command and information flow, and securing a set of necessary system-level properties*. CRUTIAL facility gateways are called *CRUTIAL Information Switches (CIS)*, which in a simplistic way could be seen as sophisticated circuit or application level firewalls combined with equally sophisticated intrusion detectors, connected by distributed protocols.

This set of servers must be intrusion-tolerant (i.e., must tolerate intrusions), prevent resource exhaustion providing perpetual operation, and be resilient against assumption coverage uncertainty, providing survivability. The services implemented on the servers must also secure the desired properties of flow control, in the presence of malicious traffic and commands, and in consequence be themselves intrusion-tolerant.

⁶ Virtual LANs.

An assumed number of components of a CIS can be corrupted. Therefore, a CIS is a logical entity that has to be implemented as a set of replicated physical units (CIS replicas) according to fault and intrusion tolerance needs. Likewise, CIS are interconnected with intrusion-tolerant protocols, in order to cooperate to implement the desired services.

3.2 CRUTIAL Nodes

The structure of some of the CII nodes, which we call *CRUTIAL nodes*, can follow the node structuring principles for intrusion-tolerant systems explained in [26]:

- The notion of *trusted* – versus *untrusted* – *hardware*. For example, most of the hardware of a CIS is considered to be untrusted, with small parts of it being considered trusted-trustworthy.
- The notion of *trusted support software*, trusted to execute a few critical functions correctly, the rest being subjected to malicious faults.
- The notion of *run-time environment*, offering trusted and untrusted software and operating system services in a homogeneous way.
- The notion of *trusted distributed components*, for example software functions implemented by collections of interacting CIS middleware.

In the context of this paper, we consider only one instantiation of CRUTIAL nodes, the CRUTIAL Information Switch (CIS) nodes. However, other specific nodes, for example, controllers needing to meet high trustworthiness standards, may be also built to a similar structure. A snapshot of the CRUTIAL node is depicted in three dimensions in Figure 2, where we can perceive the above-mentioned node structuring principles.

Firstly, there is the *hardware* dimension, which includes the node and networking devices that make up the physical distributed system. We assume that

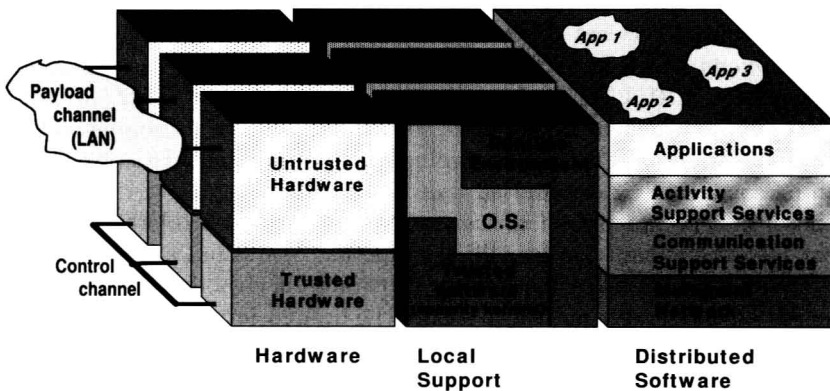


Fig. 2. Architecture and interconnection of CRUTIAL nodes (e.g., CIS)

most of a node's operations run on untrusted hardware, e.g., the usual machinery of a computer, connected through the normal networking infrastructure, which we call the *payload channel*. However, some nodes— CIS, for example— may have pieces of hardware that are trusted, for example, that by construction intruders do not have direct access to the inside of those components. The type of trusted hardware featured in CIS is an *appliance board with processor*, which may or not have an *adapter to a control channel* (an alternative trusted network), as depicted in Figure 2. This appliance is plugged to the CIS's main hardware.

Secondly, services based on the trusted hardware are accessed through the *local support* services. The rationale behind our trusted components is the following: whilst we let a local node be compromised, we make sure that the trusted component operation is not undermined (crash failure assumption).

Thirdly, there is the *distributed software* provided by CRUTIAL: middleware layers on top of which distributed applications run, even in the presence of malicious faults (far right in Figure 2). In the context of this paper, we will discuss the layers of *middleware* running inside a CIS.

4 CRUTIAL Middleware

We now observe the part of the system made of the WAN and all the CIS (facility gateways) that interconnect all the internal LANs of the critical information infrastructure to the WAN (recall Figure 1).

We model this setting as a distributed system with N nodes (CIS). We use the weakest fault and synchrony models that allow to carry out the application tasks. So, we use the asynchronous/arbitrary model, which does not make any assumptions about either time needed to make operations and faults/intrusions that can occur, as a starting point, and strengthen it as needed. For example, by resorting to hybrid models using wormholes [25], and assuming some form of partial synchrony.

We assume that the environment formed by the WAN and all the CIS is hostile (not trusted), and can thus be subjected to malicious (or arbitrary, or Byzantine⁷) faults. On the other hand, LANs trust the services provided by the CIS, but are not necessarily trusted by the latter. That is, as we will see below, LANs have different degrees of trustworthiness, which the CIS distributed protocols have to take into account. CIS securely switch information flows as a service to edge LANs as clients.

We assume that faults (accidental, attacks, intrusions) continuously occur during the life-time of the system, and that a maximum number of f malicious (or arbitrary) faults can occur within a given interval. We assume that services running in the nodes (CIS) cooperate through distributed protocols in such an

⁷ Arbitrary faults, which include attacks and intrusions, are usually called “Byzantine faults” after the seminal paper that explained the problem in terms of “Byzantine generals” [11]. Byzantine fault tolerance and intrusion tolerance usually mean the same in the literature.