Effective Host-based Intrusion Detection for Real-Time Industrial Control Systems

with emphasis on the electrical grid

Pol Van Aubel

Supervised by Jaap-Henk Hoepman & Jos Weyers

Kerckhoffs Institute – Radboud University Nijmegen
TenneT TSO B.V.

September 23rd, 2013

Acknowledgements

Outline

Critical Infrastructure

- Modern society depends on critical infrastructures:
  - Health care.
  - Transportation.
  - Public government.
  - Food supply.
  - Drinking water.
  - ...
- These all depend on one thing: *energy*.
  - Electricity.
  - Natural gas.
  - Oil.

Energy Distribution Grids

- Controlled by industrial control systems (ICS) / supervisory control and data acquisition (SCADA).
- Remote reporting and remote command execution.
- TenneT manages the Dutch electrical transmission grid.

The Electrical Grid

- The electrical ecosystem:
  - Electricity generation.
  - High-voltage transmission (TSO).
  - High-voltage and low-voltage distribution (RNB).
  - Electricity delivery.
- No viable long-term storage for electricity.
- The Transmission System Operator (TSO) performs a balancing act.
- The Dutch transmission grid is managed from locations in Arnhem and Ede.

The Electrical Grid

- An electrical line consists of three phases and a ground.
- A field is a connection between two high-voltage stations.
- At least two lines to a field.
- The fields make up several interconnected rings.

The Electrical Grid

Hoogspanningsnet.com, CC BY-NC-ND

The Electrical Grid

The Electrical Grid

- High-voltage stations link the fields of the grid.
  - Two or more rails.
  - Detachable pantographs.
  - High-voltage circuit breakers.
  - Measuring equipment.
  - Transformers.
- Grid safety systems.
- Hand-off to RNBs.

The Electrical Grid

Hoogspanningsnet.com, CC BY-NC-ND

Grid Safety Systems

- Grid safety systems handle
  - short-circuits (faults), and
  - overloads.
- They are linked to circuit breakers for their field.
- For short-circuits they perform two types of measurements:
  - Distance measurement.
  - Differential measurement.

Telecommunication Network

- Fibre-optic wires run through the ground wires.
- Also based on rings.
- Linked to Arnhem and Ede.
- Backbone for several logically separate networks:
  - Energy Management System network (EMS).
  - Grid safety systems communication.
  - Telephony network.
  - Network Management System.

Energy Management System Network

- Manual control of the grid systems.
- Based on SCADA.
  - Several Remote Terminal Units at each HVS.
  - Several RTU Control Servers at Arnhem and Ede.
  - RTUs gather sensor data and pass it on to RTUCSs.
  - Information is displayed to operators in the LBC.
  - Control decisions by operators are sent from RTUCSs to RTUs.
  - RTUs then manipulate the physical world.
- Local control is still possible, but not desired.
- Sequence of Events logging is possible to enable after-the-fact analysis.

Network Overview

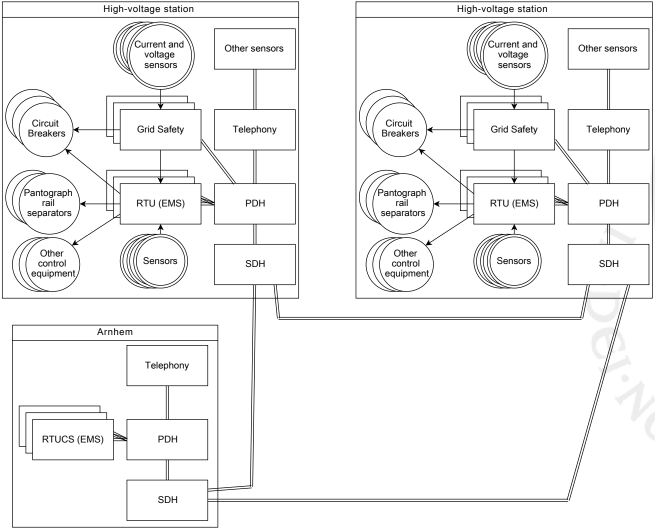

Outline

Our Goal

To improve the state of security in the electrical grid.

Why?

The grid is a robust, separate system. What's the risk?

Grid Aspects

50Hz

Grid Aspects

N-1

"any probable single event leading to a loss of power system elements . . . should not endanger the security of interconnected operation, that is, trigger a cascade of trippings or the loss of a significant amount of consumption."

The Day Europe Almost Stood Still

- Planned outage of a field crossing the Ems river on November 5, 2006, 01:00.
- TSOs would reduce cross-border transmission from 0:00-6:00.
- Outage time was moved forward, to 21:38.
- At 21:38, the field was switched off. Other lines took the load.

The Day Europe Almost Stood Still

- At 22:07, one of the lines exceeded its warning value.
- At 22:10, corrective switching measures were taken.
- Contrary to expectations, the load increased further.
- The line tripped.
- Other lines took the load, and tripped in a cascade.

The Day Europe Almost Stood Still

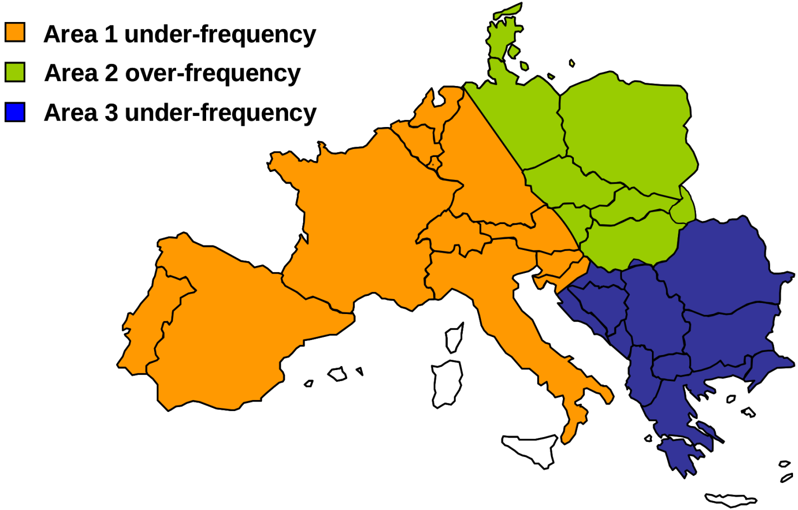

UCTE Final Report – System Disturbance on 4 November 2006

The Day Europe Almost Stood Still

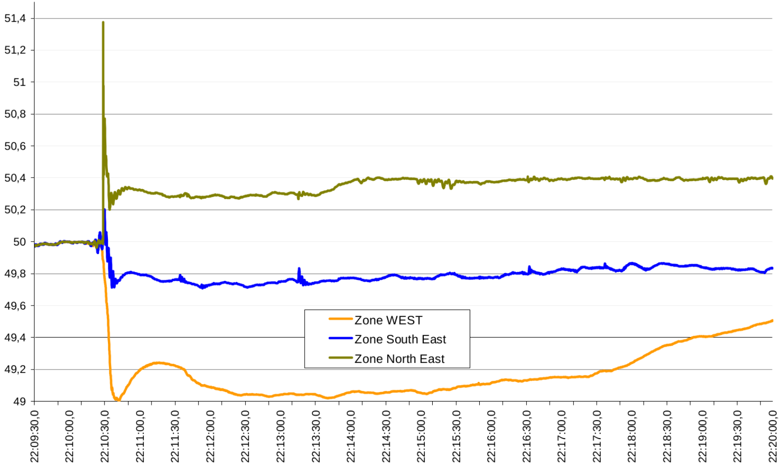

UCTE Final Report – System Disturbance on 4 November 2006

The Day Europe Almost Stood Still

- Stabilization efforts in the West:
  - Load shedding. 15 million households were disconnected.
  - Adding generating capacity.
- Stabilization efforts in the North-east:
  - Decrease of generating capacity.
  - Starting pumping storage.
  - Narrowly prevented a further split.
- Stabilization efforts in the South-east:
  - Some additional generating capacity.

The Day Europe Almost Stood Still

- Resynchronization attempted from 22:34 onwards.
- First connection between West and North-east at 22:47.
- Resynchronization finished at 23:24.
- First connection between North and South-east at 22:49.
- Resynchronization finished at 23:57.

The North-east Blackout of 2003

- Affected 65 million people in Canada and the U.S.A.
- Highly loaded lines fail due to contact with trees.
- Software bugs in the responsible TSO's system suppresses alarms.
- Over the course of 2 hours, lines trip out one by one as they hit trees or are overloaded.
- Finally, the grid splits.
- Net result: 256 power plants are off-line, and large parts of 8 U.S. states and Ontario are blacked out.

Causing a Blackout

N-x

where x is the number of systems an attacker has to take over to significantly destabilize the grid.

Causing a Blackout

- European net-split of 2006: $x = 4$.
  - 2 grid safety systems, to trip the lines.
  - 2 reporting RTUs, to fool the operators.
  - Then wait for the right moment to recreate the conditions.
- North-east blackout of 2003: $x = 1$.
  - The alarm system of the responsible TSO.
  - Then wait for a hot day.

Scary.

Causing a Blackout

- European net-split of 2006: $x = 4$.
  - 2 grid safety systems, to trip the lines.
  - 2 reporting RTUs, to fool the operators.
  - Then wait for the right moment to recreate the conditions.
- North-east blackout of 2003: $x = 1$.
  - The alarm system of the responsible TSO.
  - Then wait for a hot day.

Scary.

Why Is This Possible?

- The grid was built with centralized generation in mind.
- Now we have decentralized generation and shunting of power.
- It is now operated at almost full capacity almost all the time.
- The grid was built to accommodate grid safety.
- The grid was *not* built with security in mind.
- And neither are the software systems running on it.

The Threat is Real

Increasing attacks on SCADA systems.

- Stuxnet.
- Chinese hackers taking over water plant honeypots.
- Targeted attacks on energy companies using browser exploits.
- Increase in reported (and fixed) SCADA vulnerabilities.

Intrusion Detection on ICS

- Critical systems.
- Real-time systems.

We cannot just install an existing solution and hope it works.

Research Question

What is an effective design for a host-based intrusion detection scheme for real-time systems as used in an electrical grid?

Subquestions

1. *What kind of real-time systems exist within electrical transport operators and the gas distribution operators?*
2. *Which constraints exist for these real-time systems?*
3. *What kind of threats do these systems face in the context of a TSO?*
4. *Which intrusion detection approaches exist? What are their respective advantages and disadvantages?*
5. *What are the different architectural approaches to designing a HIDS?*
6. *What influence would these approaches have on the (real-time) constraints of the systems?*
7. *How can we model the influence an IDS will have on a real-time system?*

Outline

Actors

- State actors.
- Terrorists.
- Professional criminals.
- Hacktivists.
- Commercial parties.
- Internal actors.
- Others.

Possible Impacts

- Grid disruption.
- Physical damage.
- Loss of life.
- (Corporate) espionage.

Vectors

- Direct access to the EMS/SCADA network.
- Manipulation of internal actors.
- Installation of unauthorized hardware.
- Use of USB sticks.
- Viruses and other malicious software.
- Access to user-accounts.
- Collateral damage.

Outline

Real-Time Systems

Systems which must provide correct results in time, as defined by their application.

Types of Real-Time Systems

- Hard.
- Firm.
- Soft.

Types of Real-Time Systems

- Black-box.
- Grey-box.
- White-box.

Types of Real-Time Systems

- Mixed.
- Dedicated.

Types of Real-Time Systems Within TenneT

- Grid Safety Systems: mixed, black-box, hard.
- Remote Terminal Units: mixed, grey-box, firm.
- RTU Control Servers: mixed, grey-box, firm.
- Network backbone systems: mixed, grey-box, hard.
- Some additional black-box systems.

Outline

Types of Intrusion Detection

- Network-based.
- Host-based.
- Application-based.

Host-based Intrusion Detection

- Signature-based.
  - Computationally cheap.
  - Effective at detecting known attacks.
  - Identify the vulnerability being exploited.
  - Few false positives.
- Anomaly-based.
  - Able to detect unknown attacks.
  - Higher detection rate.
  - Harder to bypass.
- Hybrid.

Anomaly-based Intrusion Detection

- Statistical methods.
  - CPU usage.
  - Memory usage.
  - ...
- Data mining & machine learning methods.
  - System / library call tracing.
  - Hidden Markov Models.
  - ...
- Specification based methods.
  - Protocol specifications.
  - Finite state machines.
  - ...

Deployment Structure

- Stand-alone.
- Distributed.
- Mobile.
- Collaborative.
  - Centralized.
  - Hierarchical.
  - Fully distributed.

Outline

Influences on Run-Time Behaviour

- Utilization and availability of System Resources.
  - CPU time.
  - System memory.
  - Network bandwidth.
  - Shared states.
- Overhead of data collection.
  - System call tracing.
  - Bookkeeping collection.

Shared States

- Portions of memory shared between processes.
- Sometimes require (unexpected) exclusive access.
- Known cause of deadlocks.
- But even when deadlock-safe, may affect real-time constraints.
  - Reading the file-descriptor table of a Linux process.
  - Unexpected slowdowns.
  - Process algebra may be used to model this.

ACNSR

Algebra of Communicating Non-exclusive Shared Resources with Dense Time and Priorities.

- Extends ACSR with shared resources which may be concurrently accessed.
- Allows us to reason about locking behaviour in a non-exclusive model.
- Can be further extended to allow for finite but shared resources.

Outline

Detection Engine

Off-host.

- Detection engine can be as heavy as we want.
- Can be protected against attack.
- Does not require host-interaction to update.
- Collection engine is resource-constrained.

Deployment Structure

Collaborative, hierarchical system.

- An IDS for each subsystem (EMS, Telecommunication, Grid Safety).
- Each system will have a collection engine.
- Each HVS will have several detection and reporting engines.
- These then report to a central detection and reporting unit in the LBC.

Deployment Structure

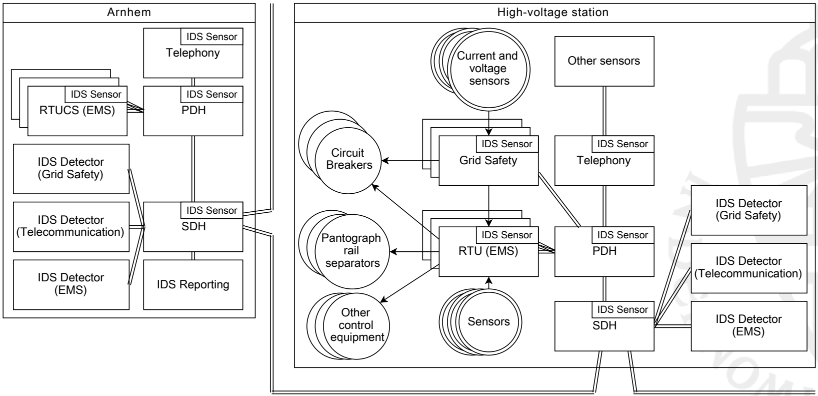

Data Sources and Detection Models

- Anomaly-based:
  - System call tracing with hidden Markov models.
  - System resource utilization with statistical learned models.
  - System configuration change tracking with fixed limits.
- Signature-based:
  - Presence of certain files on the filesystem.
  - System call payload signatures.

Challenges

- Black-box and grey-box systems are hard to include.
- All hypothetical, no existing systems.
- Vendor cooperation required.
- No analysis of the effectiveness of this design.
- Large number of nodes in the IDS.

Outline

Future Work

- Models to evaluate the different possible architectures.
- Precise requirements for industrial control systems.
- Determining detection accuracy requirements.
- Prototyping the proposed system.
- Research detection techniques on low-interaction systems.
- Research into using physical properties of the electrical grid.
- Expanding ACNSR further for use with finite resources.

Conclusion

- We have proposed an architecture for an effective host-based intrusion detection system on industrial control systems.
- We have extended the process algebra ACSR to be able to reason about non-exclusively shared resources in real-time.
- We have discussed several avenues of future work to consider as continuation for this research.