Automated Decision Making for Network Defences and Data Protection

Dr. Jassim Happa
Research Fellow
Department of Computer Science
University of Oxford
jassim.happa@cs.ox.ac.uk

With thanks to Tom Bashford-Rogers, Alastair Janse van Rensburg, Arnau Erola, Nick Moffat, Martin Helmhout, Ioannis Agrafiotis, Phil Legg, Michael Goldsmith and Sadie Creese.

Overview

- **Threat Detection Systems**
  - Risk Propagation Logics
  - Insider Threat Detection
- **Automation in Network Defences**
  - Decision systems
  - Cyber Threat Intelligence Sharing
    - Data Protection and NDA compliance

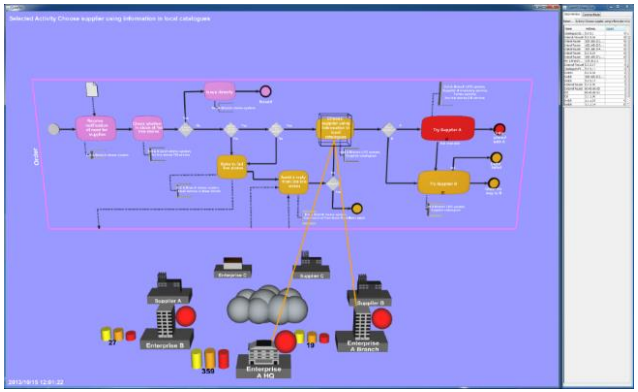

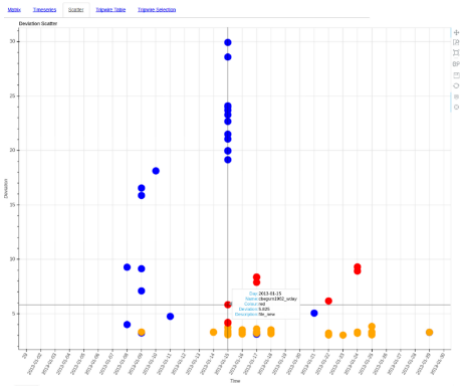

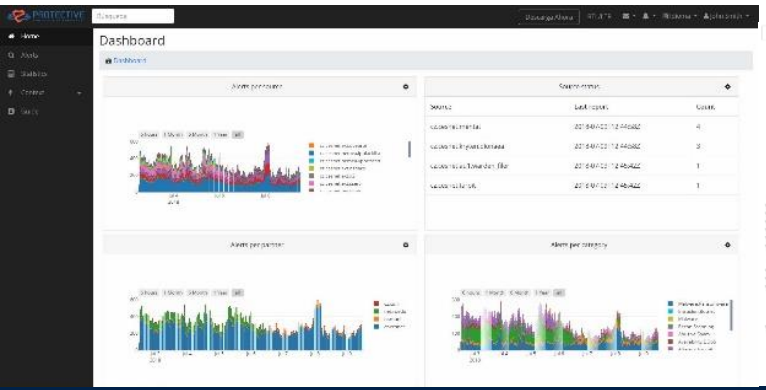

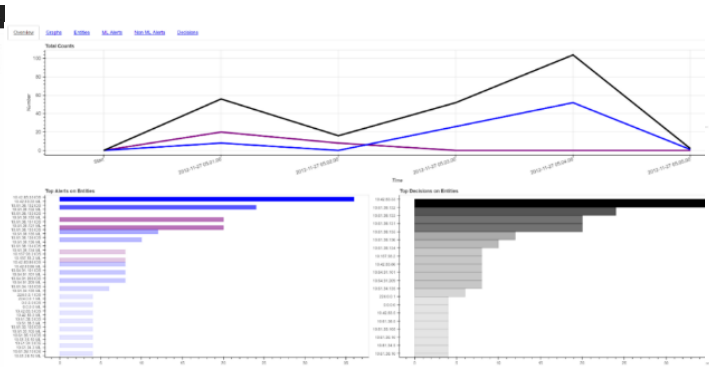

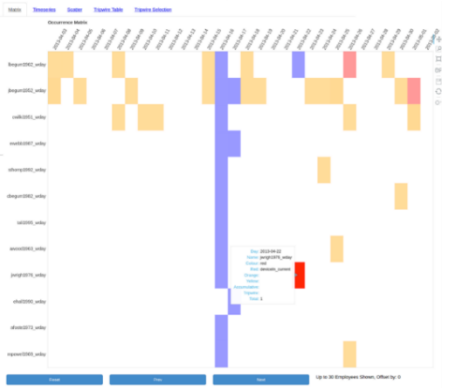

Risk Propagation Logics I

- **Decision Support** in Security Operation Centres
  - Core: "*How do cyber attacks affect mission operations?*"
  - Dependency graphs: semantic models for reasoning
  - Best-of-breed (e.g. Snort, Nagios, BPMN etc.)

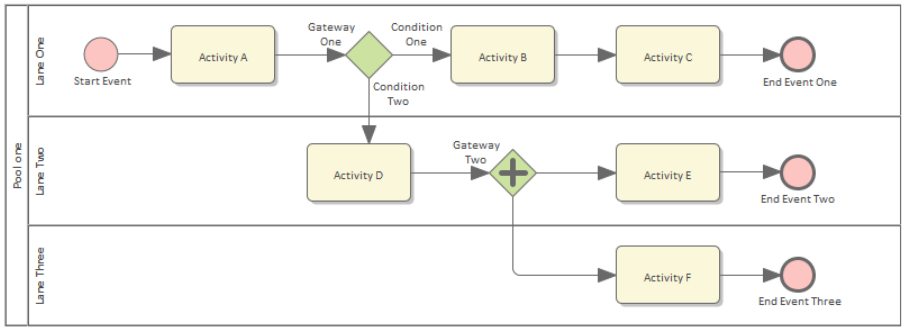

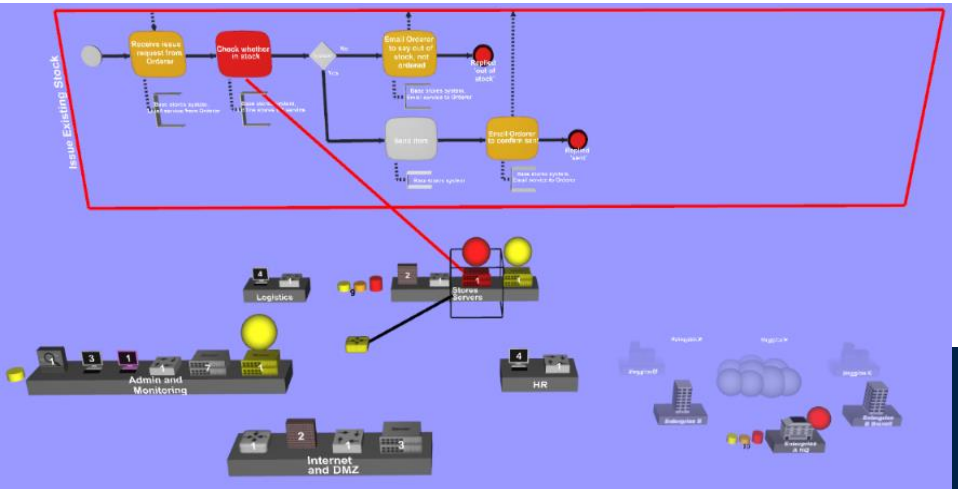

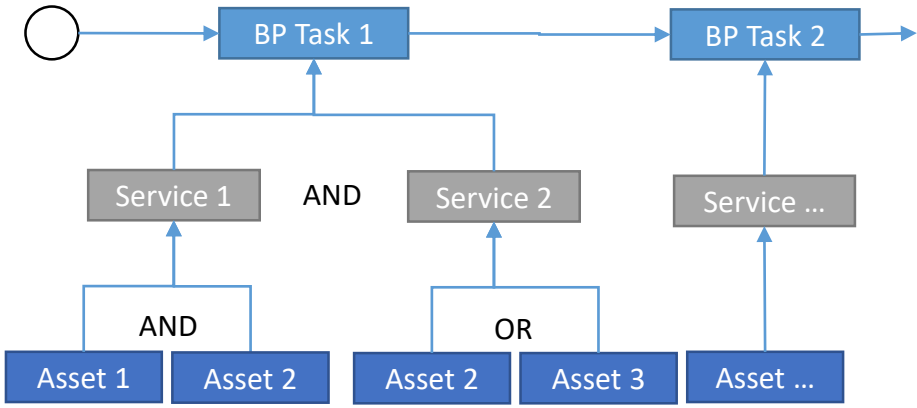

See: Creese et al. "CyberVis: Visualizing the potential impact of cyber attacks on the wider enterprise", IEEE HST 2013

Insider Threat Detection I

- Machine learning + Visual Analytics
- Human element of attacks
- PCA on network and non-network data
  - In-depth study on one organisation
- Bastion solution:

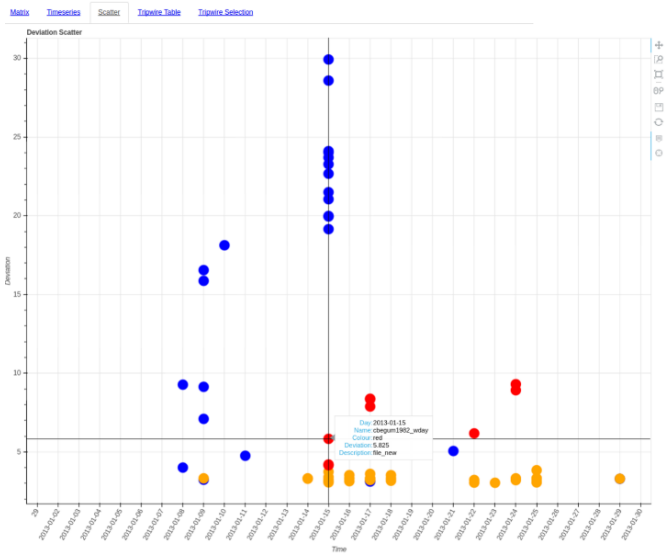

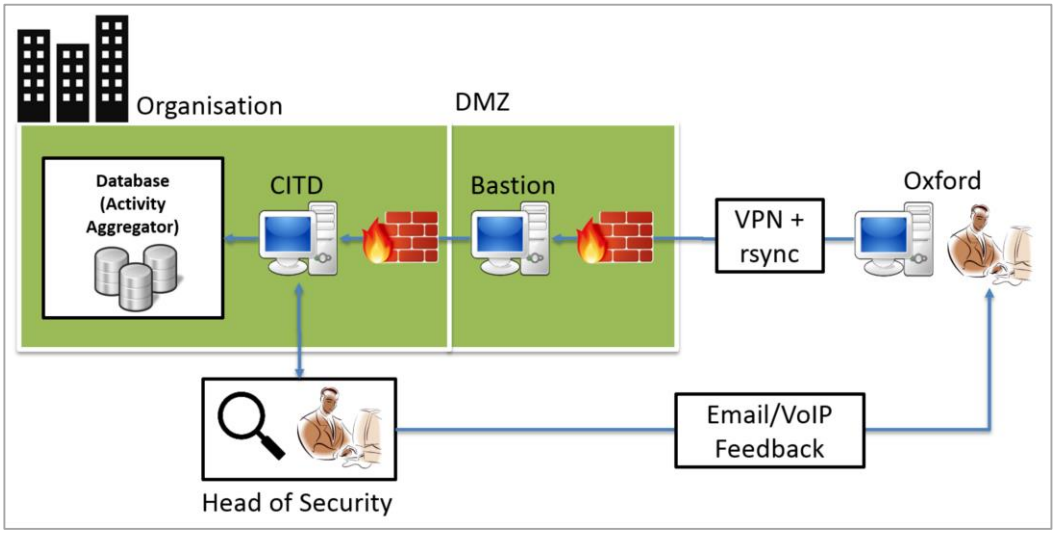

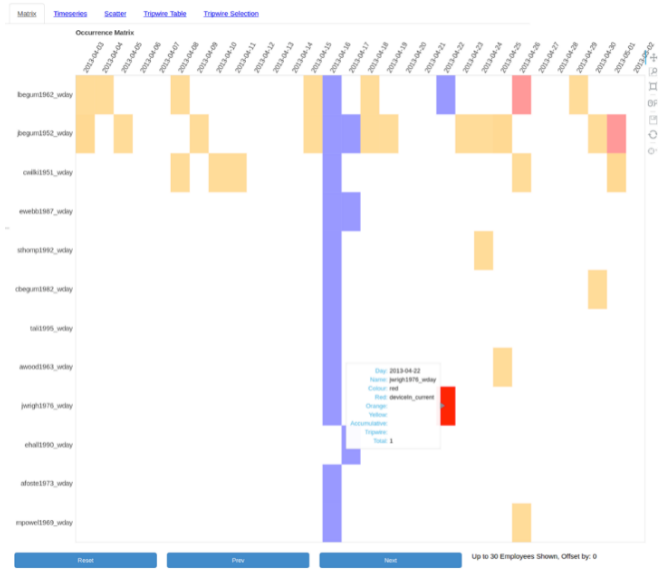

*See: Agrafiotis et al. “Validating an Insider Threat Detection System” IEEE S&P Workshop 2016. ****BEST PAPER AWARD****

What did we learn from these projects?

- **Big-data analytics** will always be important
  - Statistics
  - Machine Learning
- **"Making sense" of cybersecurity** likely to become more important
  - **Quantify concepts into something usable**
    - Trust, Dependency, Business Processes, Risk, Harm, Impact, etc.
  - **Well-informed decisions relies on**:
    - Human factors (users and attackers), organisational resilience, dependency relationships, priorities, security postures, context.
  - **Automation challenges**:
    - Threat detection, responses, ethics and legal compliance

Automated Network Defences I

Generator powers buildings

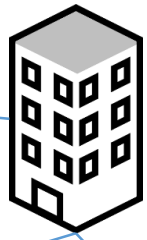

People maintain generators

Machines exist in buildings

People use buildings

Logs are generated everywhere!

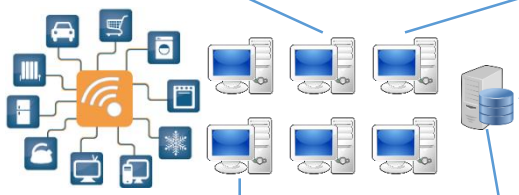

People use machines

Network relies on Internet infrastructure

Services exist on machines

Services are used by people

Internet

Services rely on the Internet

Automated Network Defences II

- **Context-driven** Automated Network Defence
  - ***security posture*** (description of mission + asset priorities)
  - ***known state of assets*** (e.g. OS, vuln., on/off)
  - ***alerts*** (received/generated)
- **Nuanced decision making by using BPMN and a decision-making grammar**
  - Early trials of probabilistic moving targets

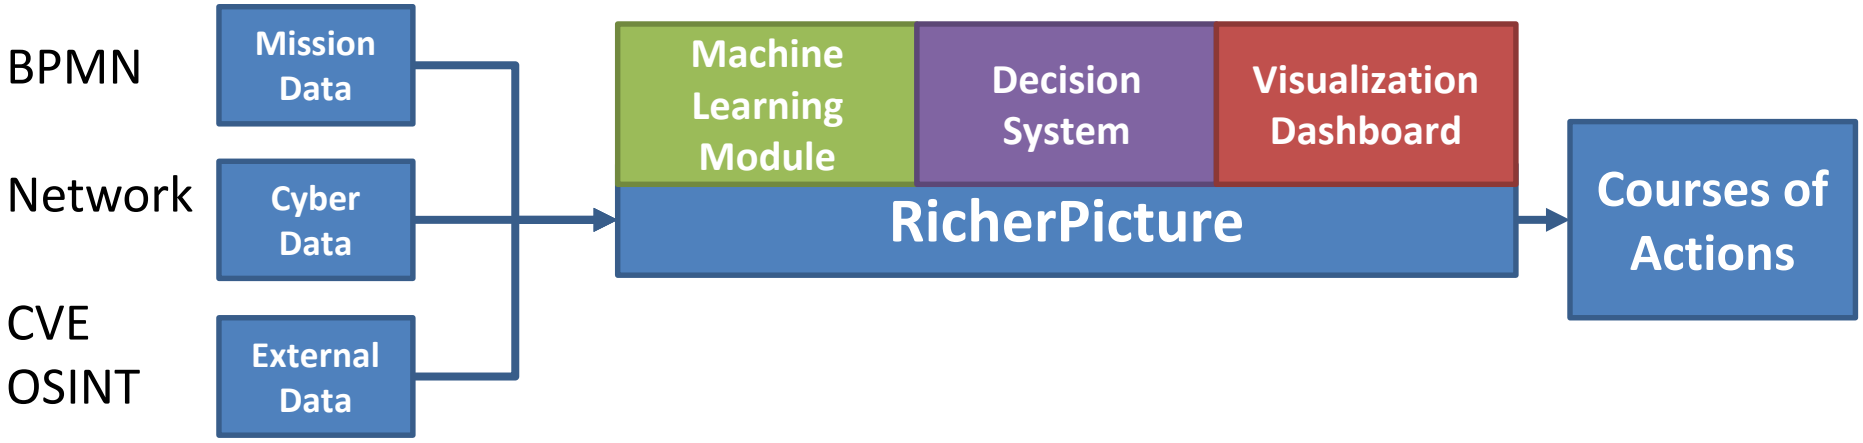

Automated Network Defences III

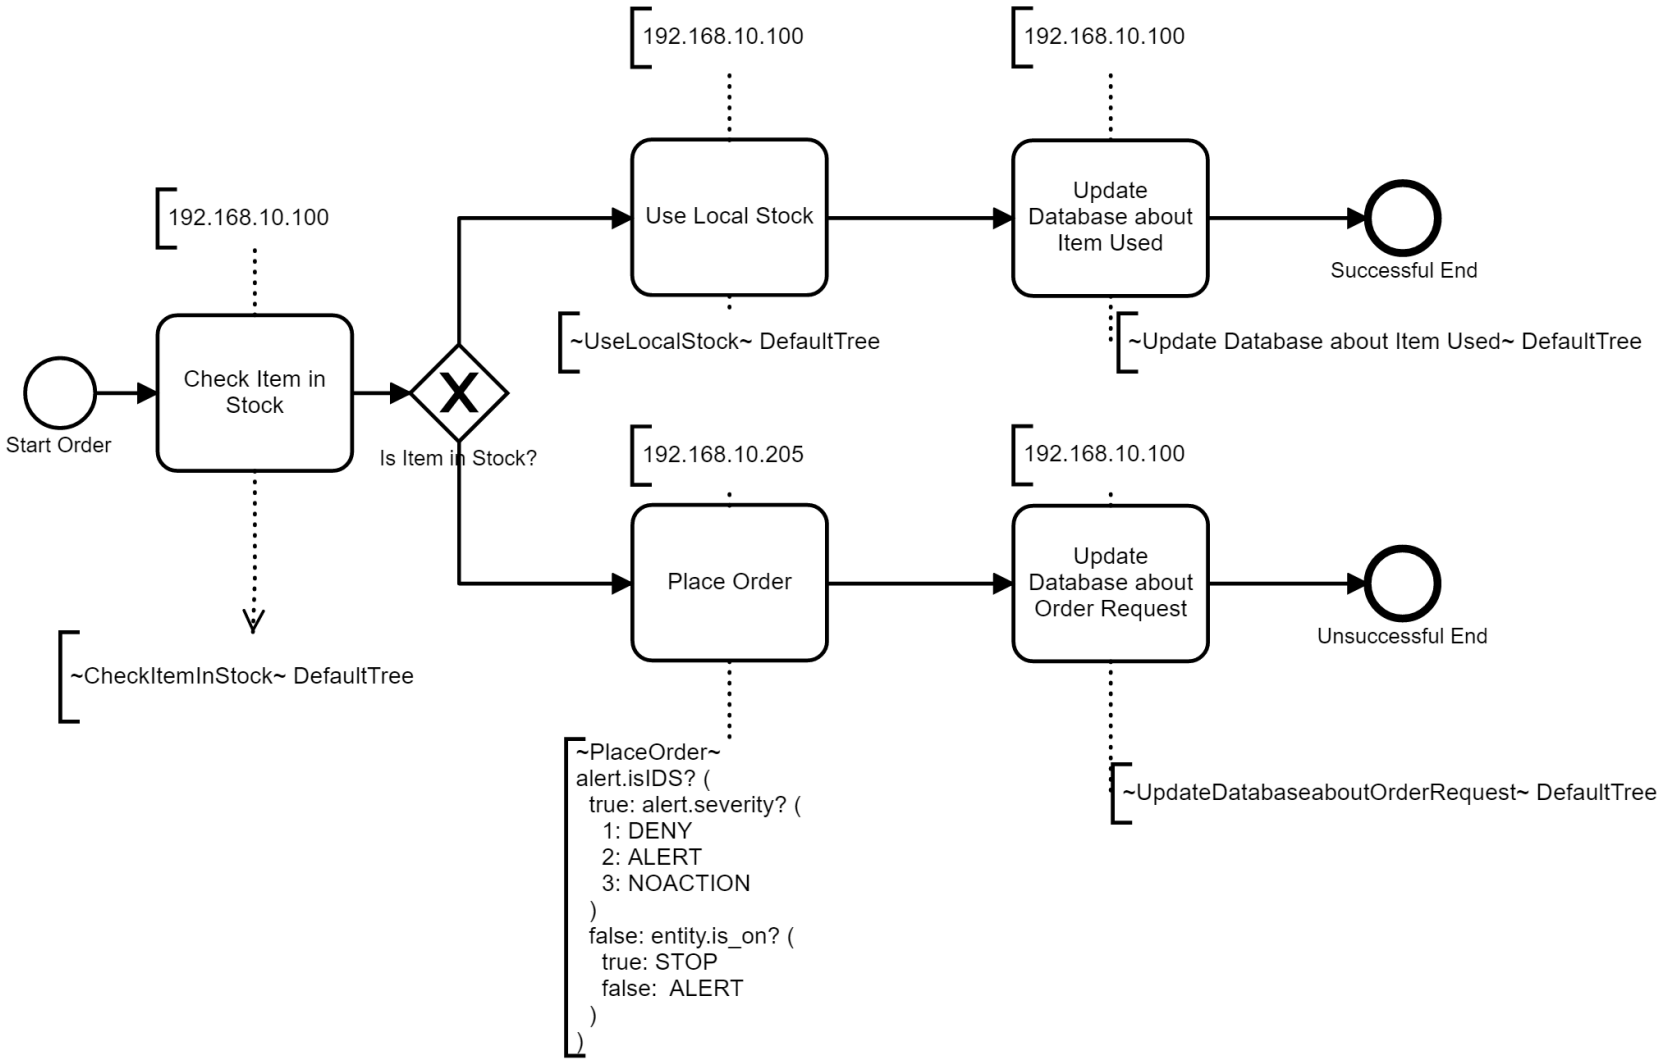

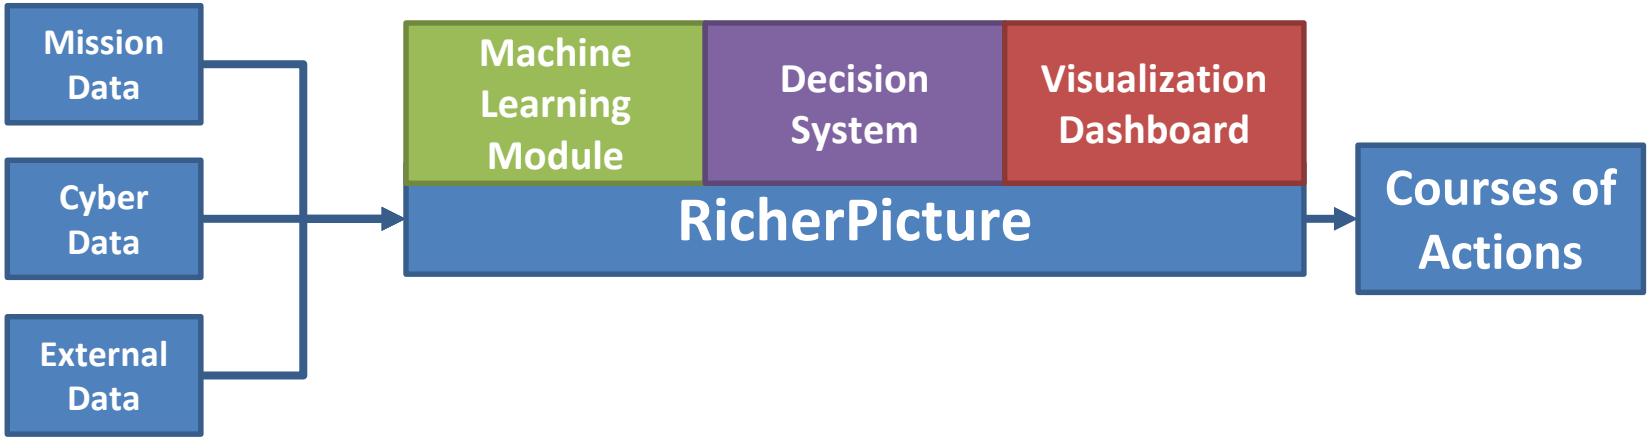

Video Demo

Performance statistics:

- Tested in a simulation with 140,000 netflows per second, generating additional IDS alerts.
- Distributed p2p solution.

Purpose of video:

- Show how single trees computes decisions

Cyber Threat Intelligence Sharing I

- "PROTECTIVE": A H2020 funded Innovation Action
- Aim: assist CSIRTS in incident response
- **Features:**
  - SIEM
  - Visualization Dashboard
  - Intelligence Sharing
- **Distinct PROTECTIVE features:**
  - IDEA, MISP and STIX support
  - Context Awareness
  - Data Fusion (Meta Alerts)
  - Computational Trust
  - Run-time monitoring of Information Sharing Compliance

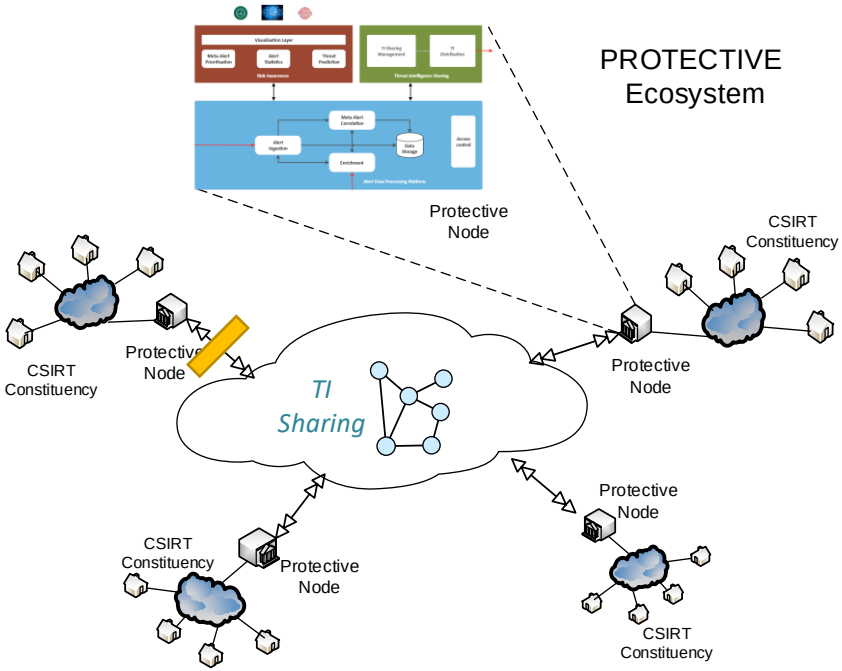

This project has received funding from the European Union's Horizon 2020 research and innovation program under grant agreement No 700071. This output reflects the views only of the author(s), and the European Union cannot be held responsible for any use which may be made of the information contained therein.

Cyber Threat Intelligence Sharing II

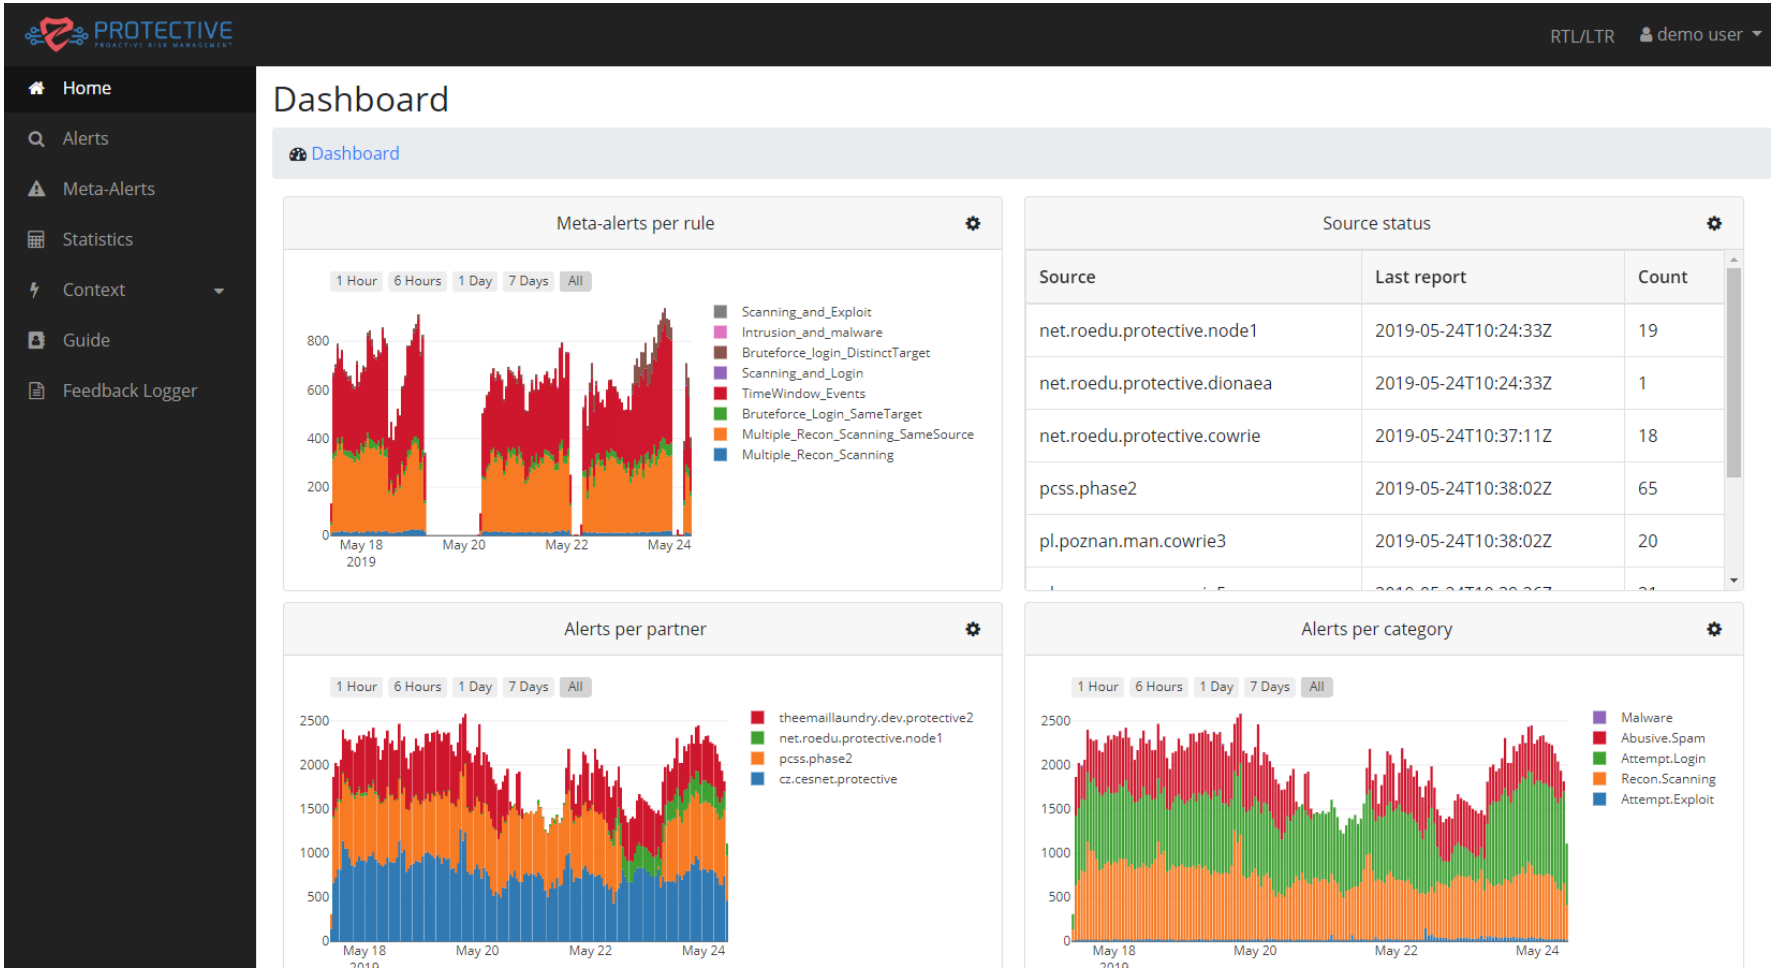

We are running a pilot! If you're interested in trying the tool, email: info@protective-h2020.eu

Cyber Threat Intelligence Sharing III

- "Data sharing with data protection in mind"
- Legal/SOP/NDA Speak and Tech Speak: opposing forces

Room for interpretation.
E.g. terms like "reasonable"

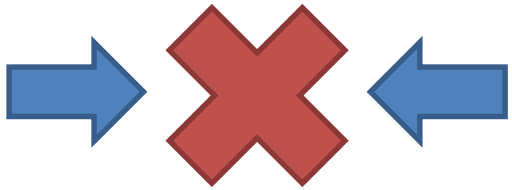

No room for interpretation.
Exact instructions.

Benefits:
- Allows for human decision.
- Allows for context.

Disadvantages:
- Slow.
- Ambiguous.

Benefits:
- Allows for automation.
- Allows for large reach.

Disadvantages:
- False positives/False negatives.
- Unintelligent.

- Codifying law and NDAs is **hard**
  - Requirements are present, but not specifications.
  - When are "legitimate interests" legitimate?
  - Language is difficult to process
  - What is the baseline – i.e. what is "good enough"?

Cyber Threat Intelligence Sharing IV

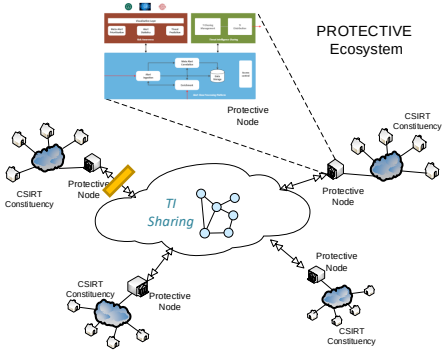

- **Expert System - building blocks** for:
  - Auditing and Enforcement of Compliance.
  - Designed with GDPR, NDA, IEP in mind.
    - Confidential Information Exchange.
- Decoupled from the main PROTECTIVE tool. Generic Solution.
  - Supports any cleartext data format
- Templates of rules. Rules modifiable.

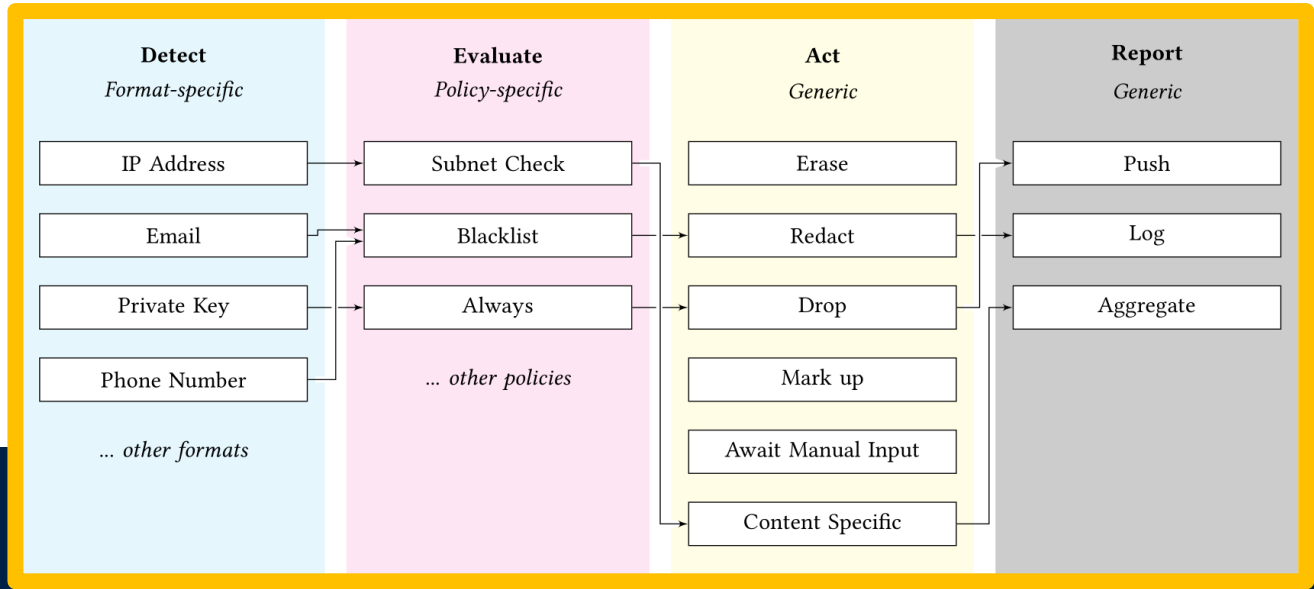

Cyber Threat Intelligence Sharing V

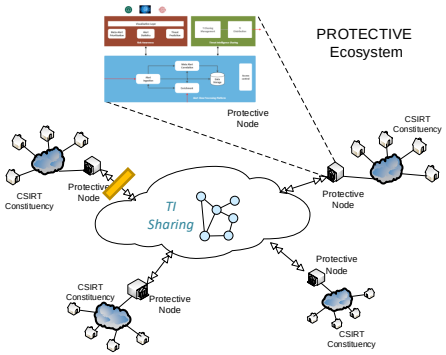

Type of action:

- Erasing, redacting, anonymisation, pseudonymisation, summarising, reporting, logging, marking up, dropping (whole events), …

Conditions to execute actions:

- Recipient, Always, RegEx, Blacklist, Timeliness, Time of day, …

Original Alert

```
{
  "Format": "IDEA0",
  "ID": "4390fc3f",
  "Source": [{
    "Type": [
      "Phishing"
    ],
    "IP4": [
      "1.2.3.4"
    ],
    "Hostname": [
      "companydomain.com"
    ],
    "URL": [
      "http://companydomain.com/cgibin/"
    ]
  }]
}
```

ISC

Modified Alert

```
{
  "Format": "IDEA0",
  "ID": "4390fc3f",
  "Source": [{
    "Type": [
      "Phishing"
    ],
    "IP4": [
      "1.2.0.0/16"
    ],
    "Hostname": [
      "[redacted].com"
    ],
    "URL": [
      "http://[redacted].com/cgibin/"
    ]
  }]
}
```

Open sourced: https://gitlab.com/protective-h2020-eu/protective-node/wikis/home

```
{
  "Format": "IDEA0",
  "ID": "4390fc3f-c753-4a3e-bc83-1b44f24baf75",
  "CreateTime": "2012-11-03T10:00:02Z",
  "DetectTime": "2012-11-03T10:00:07Z",
  "WinStartTime": "2012-11-03T05:00:00Z",
  "WinEndTime": "2012-11-03T10:00:00Z",
  "EventTime": "2012-11-03T07:36:00Z",
  "CeaseTime": "2012-11-03T09:55:22Z",
  "Category": ["Fraud.Phishing"],
  "Ref": ["cve:CVE-1234-5678"],
  "Confidence": 1,
  "Note": "Synthetic example",
  "ConnCount": 20,
  "Source": [
    {
      "Type": ["Phishing"],
      "IP4": ["192.168.0.2-192.168.0.5", "192.168.0.10/25"],
      "IP6": ["2001:0db8:0000:0000:0000:ff00:0042::/112"],
      "Hostname": ["example.com"],
      "URL": ["http://example.com/cgi-bin/killemall"],
      "Proto": ["tcp", "http"],
      "AttachHand": ["att1"],
      "Netname": ["ripe:IANA-CBLK-RESERVED1"]
    }
  ],
  "Target": [
    {
      "Type": ["Backscatter", "OriginSpam"],
      "Email": ["innocent@example.com"],
      "Spoofed": true
    },
    {
      "IP4": ["10.2.2.0/24"],
      "Anonymised": true
    }
  ],
  "Attach": [
    {
      "Handle": "att1",
      "FileName": ["killemall"],
      "Type": ["Malware"],
      "ContentType": "application/octet-stream",
      "Hash": ["sha1:0c4a38c3569f0cc632e74f4c"],
      "Size": 46,
      "Ref": ["Trojan-Spy:W32/FinSpy.A"],
      "ContentEncoding": "base64",
      "Content": "TVpqdXN0a2lkZGluZwo="
    }
  ],
  "Node": [
    {
      "Name": "cz.cesnet.kippo-honey",
      "Type": ["Protocol", "Honeypot"],
      "SW": ["Kippo"],
      "AggrWin": "00:05:00"
    }
  ]
}
```


Performance statistics:

- Approx. 2.2ms per alert on typical desktop hardware, for roughly 450 alerts per second (~39 million events per day).
- Distributed p2p solution, also supports client-server architectures.

Video Demo

Purpose of video:

- Show performance and types of fields

Future work/open questions

- **Making 'better' sense of cybersecurity**
  - Image and text analysis
  - Not only technology-centric focus
    - Human factors (as users, targets and attackers)
    - Non-network data
    - Resilience concerns (wear and tear, natural hazards)
  - How does law and ethics fit in?
- **Improving decision-making with AI**
  - Learn and easily update mission-to-asset dependencies.
  - Trust: conflicting data and compromised systems
  - When is the attacker is deceiving your AI systems?
  - Leveraging Cyber Threat Intelligence
  - Interoperability!
- **Benchmarking 'good' automated decision-making**
  - Red-team exercises?
  - Effectiveness of deception using decision-making.
  - Effectiveness of fast decision making.
- **Aim: Synthesis of human and machine decision making**

Questions?

jassim.happa@cs.ox.ac.uk

“Data protection in real-time” workshop:

https://privacyworkshop19.oasis-open.org/en/