Insider Threats: Identifying Anomalous Human Behaviour in Heterogeneous Systems Using Beneficial Intelligent Software (Ben-ware)

Stephen McGough, Budi Arief, David Wall, John Brennan, Carl Gamble, John Fitzgerald, **Aad van Moorsel**, Sujeewa Alwis, Georgios Theodoropoulos, Ed Ruck-Keene

Durham University

Newcastle University

Live Foresight
See. Reveal. Predict.

What are we trying to solve?

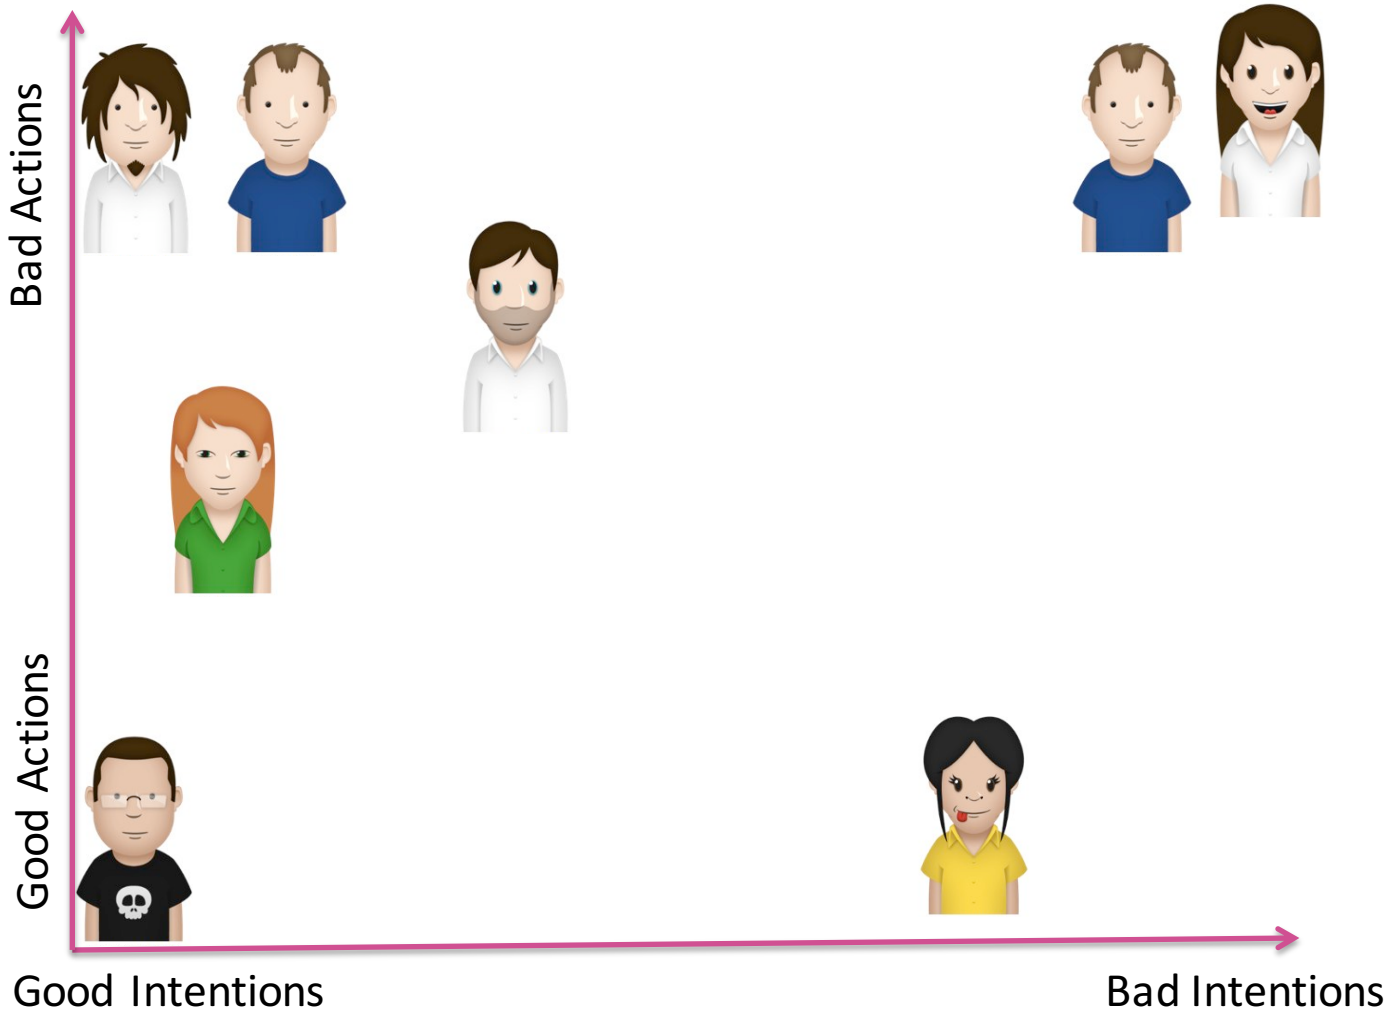

Where are we trying to solve it?

What is Ben-ware?

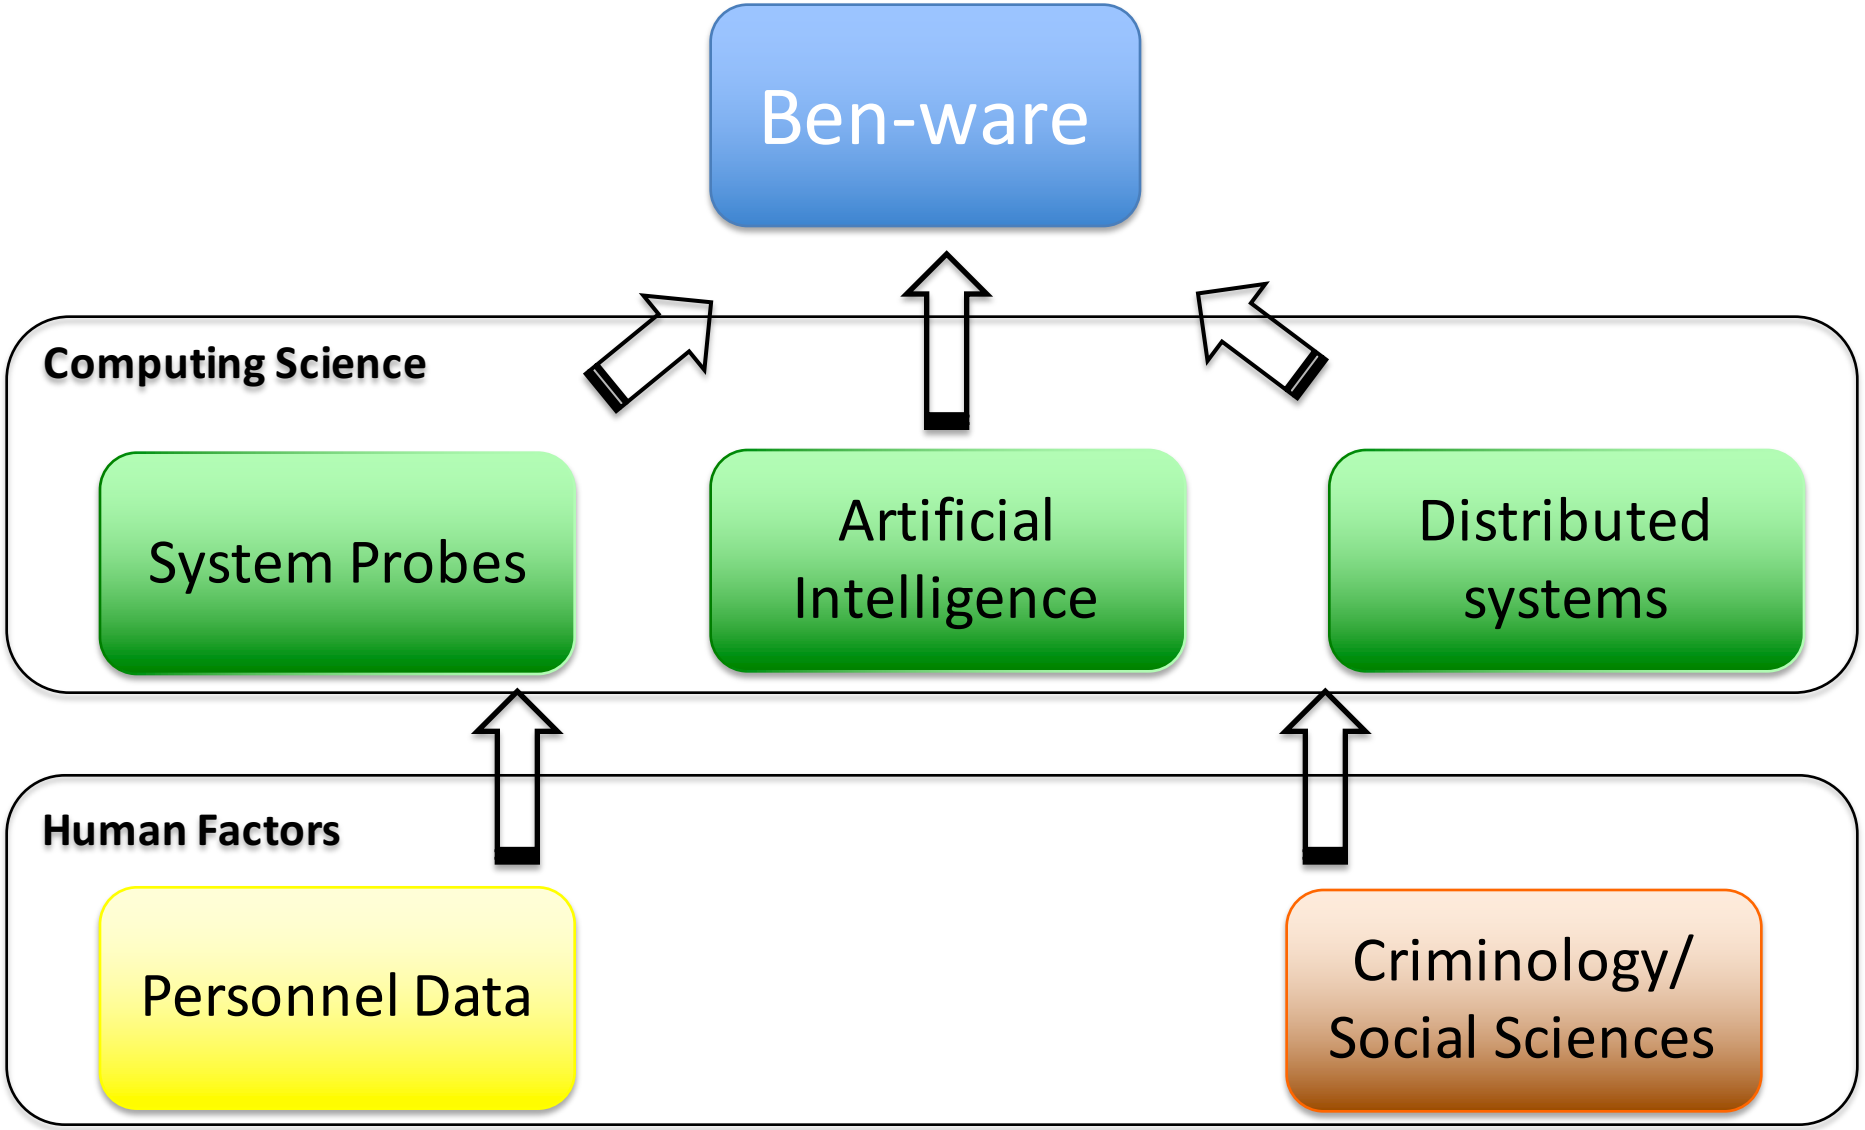

User Types

– well-behaved
  - abides by the rules, does everything correctly

– the negligent
  - Doesn't realize they are breaking the rules

– the ambitious
  - Knowingly breaks the rules but in an attempt to do things better

– the malicious insider
  - Actively seeking to thwart the organisation

– The whistle blower
  - Decides for their own reason that information should be public

The Sleeper

Ben-ware Architecture

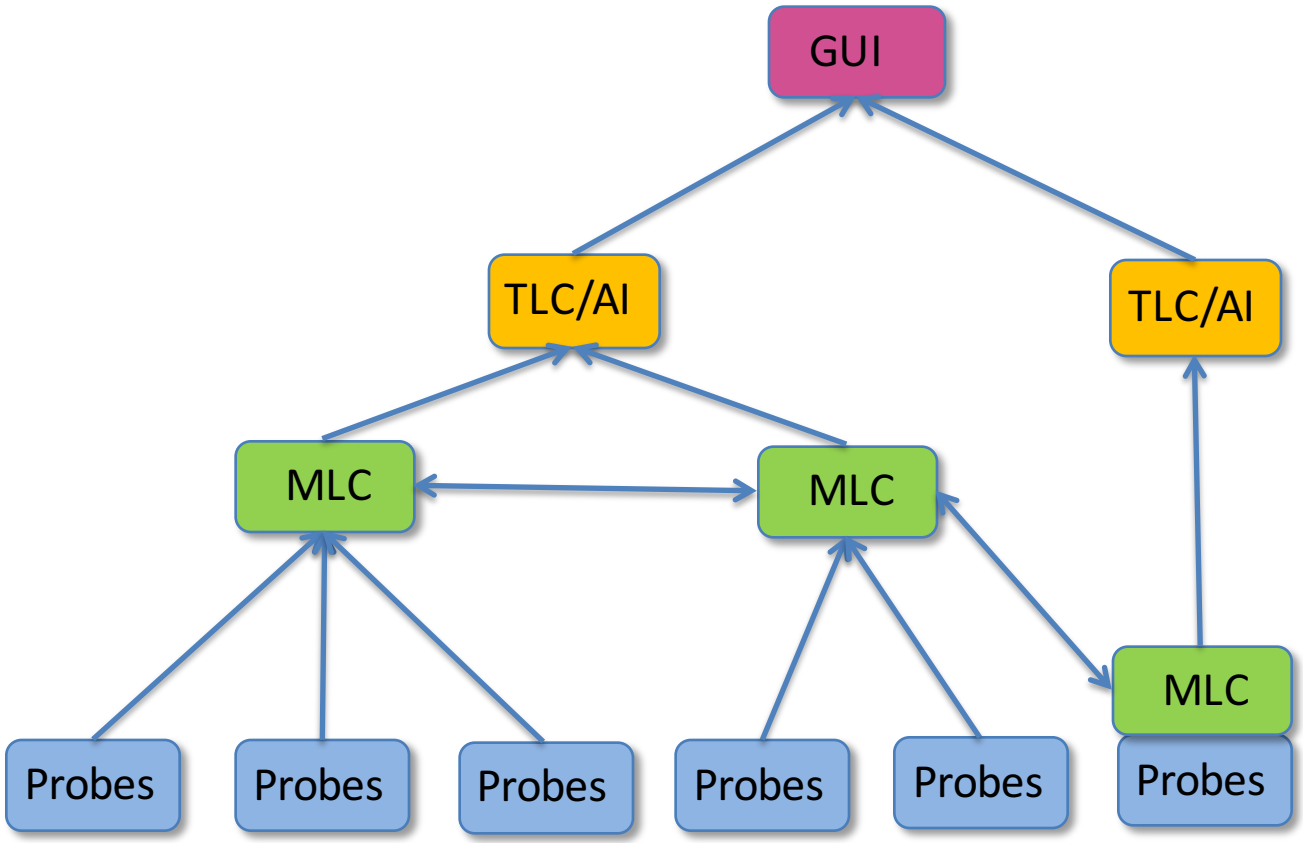

Benware Architecture

Ben-ware: TLC/AI

Informs security office of
Presentation of user eve
Remote configuration of

Ben-ware: GUI

Examines summaries of events for anomalies
Could use both personal and role based profiles
Long term store of user events

Ben-ware: MLC

Generates summaries of activity over a period
Raises alerts when user activity exceeds a threshold
Short term store of events for a group of users

Ben-ware: Probes

Gathers detailed events of user activity
Includes measures to detect circumvention attempts
Stores events only until can be sent

More difficult than a standard classification problem

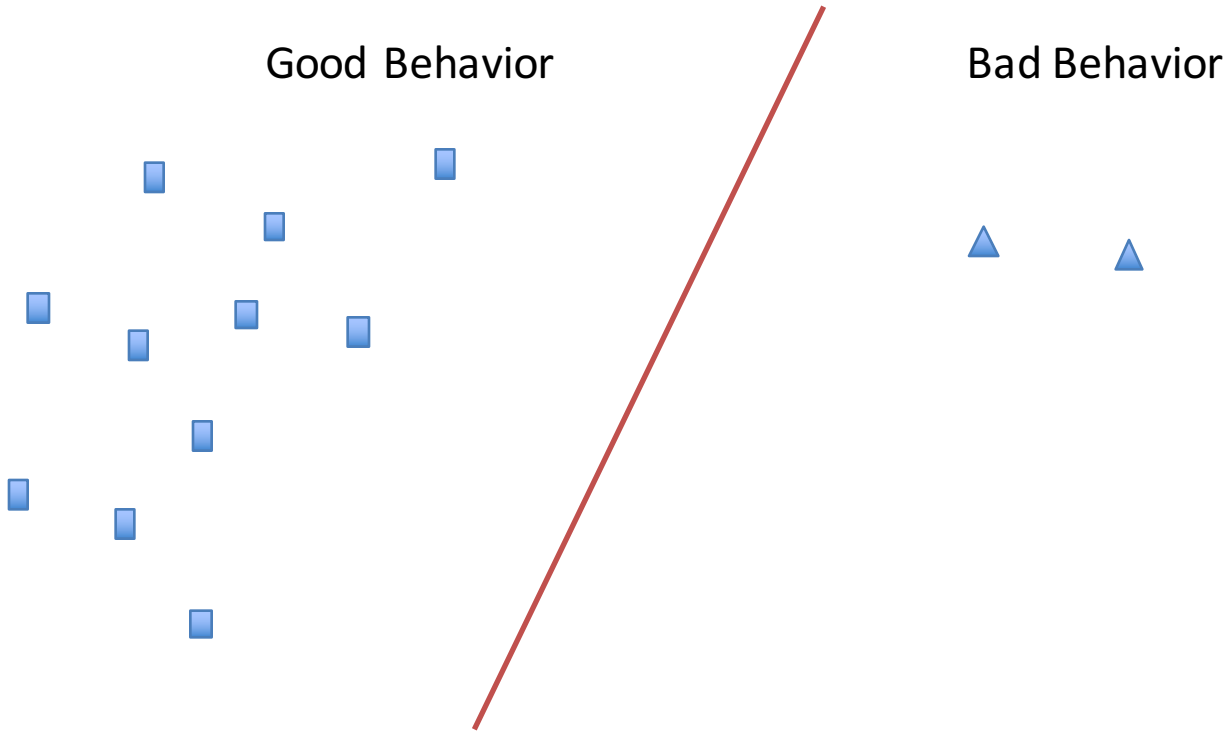

Support Vector Data Description (SVDD)

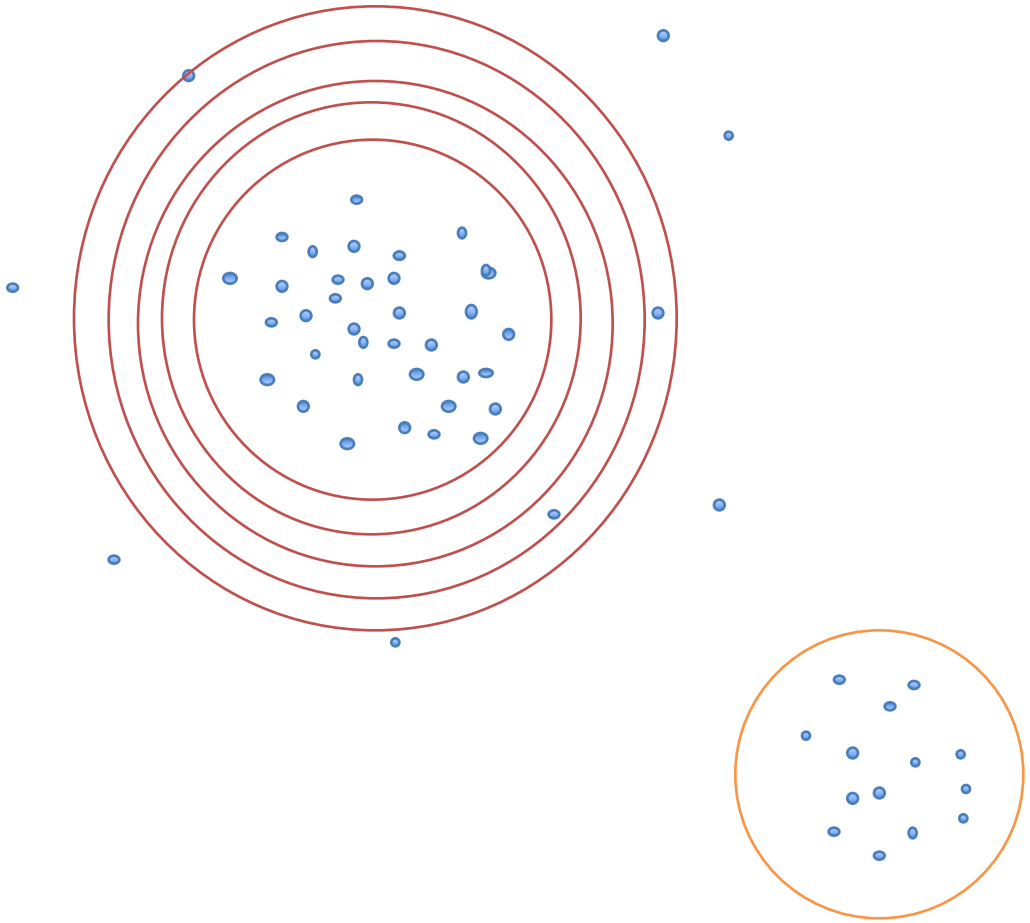

Demonstration Video Virtual Machine Arrangement

User Level Probes

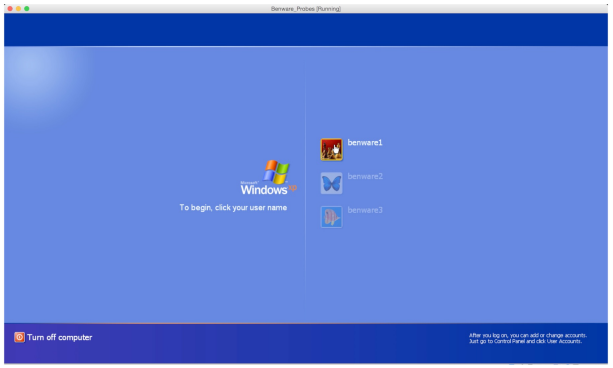

Mid Level Controller

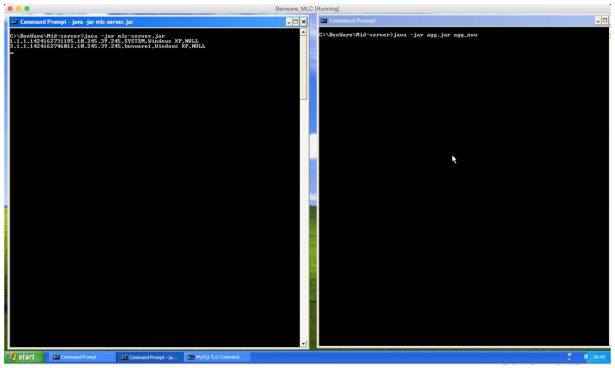

Intelligent Agent

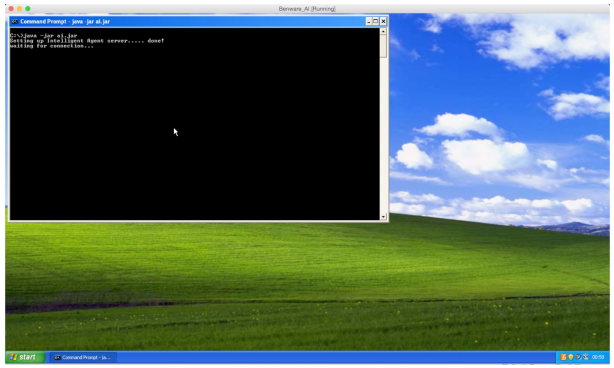

Demonstration Video Virtual Machine Arrangement

Demo video not in this presentation

Results: Threat Detection

- Scenarios (synthetic users):
  - User's Behaves 'good' for 12 months then goes 'bad' on a random day during the next 6 months
  - **Bad1:** a lot of files were stolen on a single day
  - **Bad2:** a small number of files (usually 2) were stolen on each subsequent day
  - **Bad2 gaps:** like Bad2 but allowing gaps
  - **Bad3:** an increasing number of files were stolen on each day for a period of time

Results: Threat Detection

- Time to identify

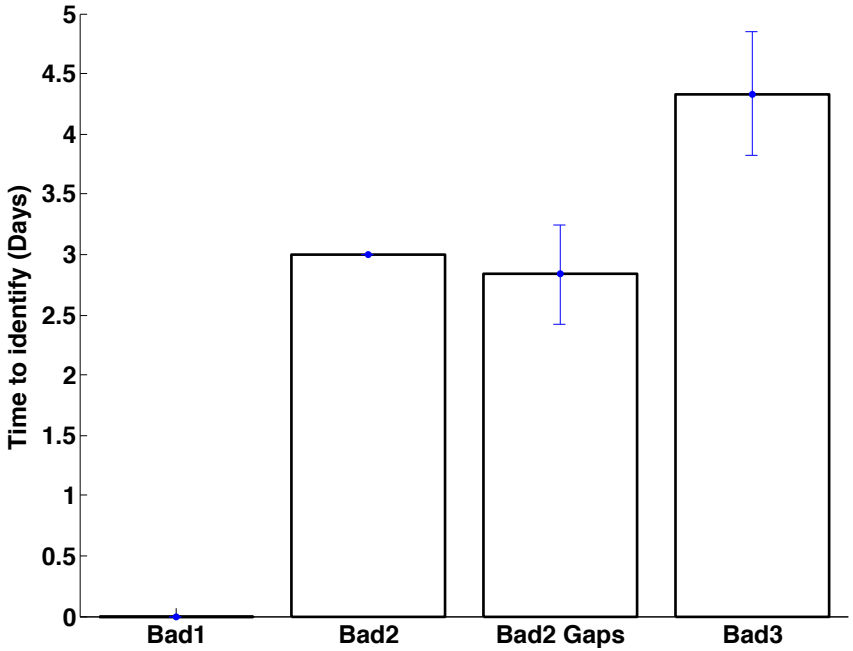

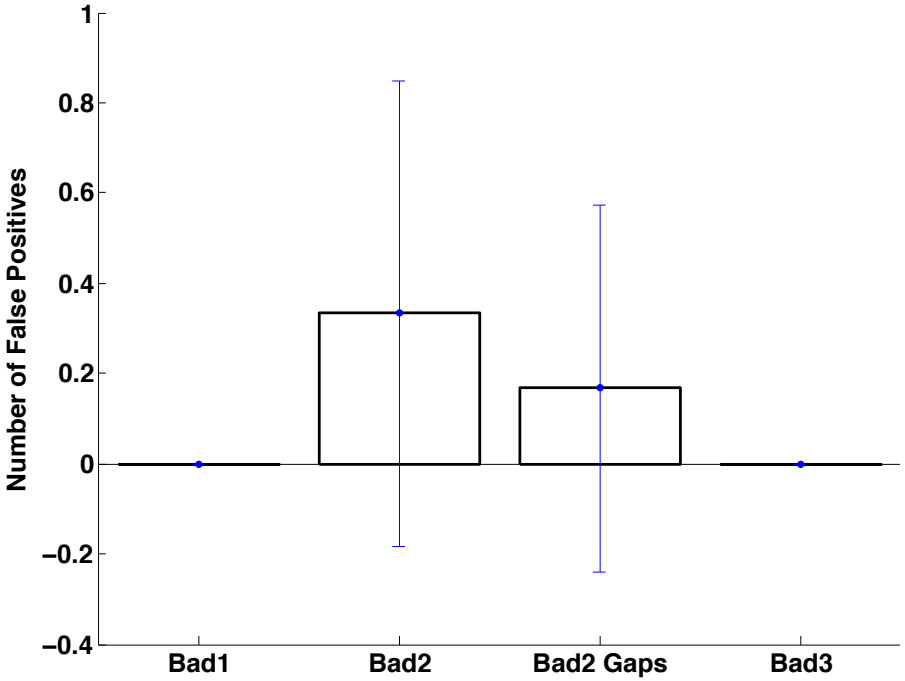

- Average number of False Positives

Results: Impact on System

- Average impact of Probes on system

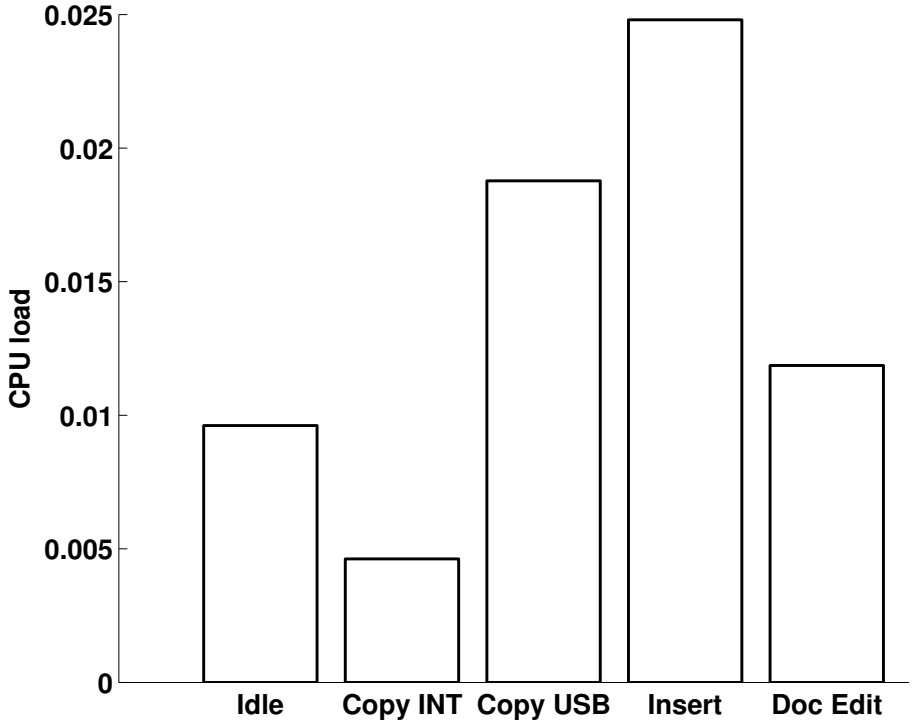

- Data transmission
  - Maximum 5.3% of a 14.4Kbit/s modem link
- MLC processing – maximum 15%*

* Intel Pentium M, 1.73Ghz, 512MB, Windows XP Pro

Conclusions and Future Work

- Good detection rates for file stealing
  - Less than 4.5 days on average to identify files being stolen
  - Less than 0.4 days on average were falsely identified as days where files were stolen during 6 month test
- Low impact on the legacy systems used
- Deployment in a real environment
- Detection of other threats than just files
- Integrating Human Factors data and personnel records

benware@newcastle.ac.uk