TU/e EINDHOVEN UNIVERSITY OF TECHNOLOGY

Role Inference + Anomaly Detection = Situational Awareness in BACnet networks

D. Fauri, M. Kapsalakis, D. R. dos Santos, E.Costante, J. den Hartog, S. Etalle

Gothenburg, Sweden – DIMVA 2019 - 16th Conference on Detection of Intrusions and Malware & Vulnerability Assessment

Building Automation Systems (BAS)

Icons made by Freepik from www.flaticon.com

- They manage HVAC, video surveillance, access control, lighting, elevators…
- Usually across many buildings, many different networks (but interoperability exists, e.g. BACnet)
- They can be managed remotely
- They can be *attacked* remotely


the security ledger

Update: Let's Get Cyberphysical: Internet Attack shuts off the Heat in Finland

November 8, 2016 14:23 by Paul


The New York Times

Hackers Use New Tactic at Austrian Hotel: Locking the Doors

By Dan Bilefsky

TU/e

Situational Awareness in BAS

Cyber Situational Awareness is structured in three subsuming levels [1]:

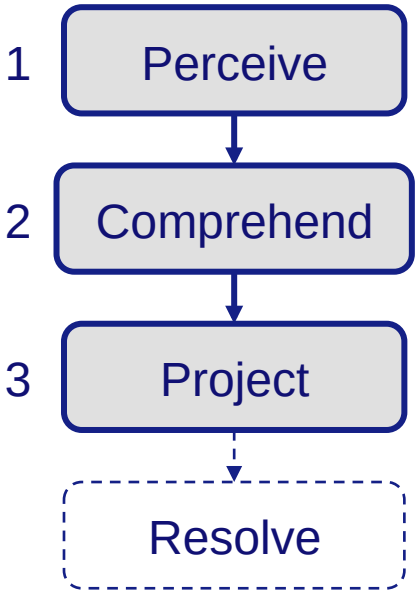

1) Basic perception of important data:
e.g., presence of devices in a network, device configuration, device behavior, alerts raised by IDS, system specification

2) Interpretation and combination of data into knowledge:
e.g., search a device's FW version in a CVE database, recognize if a raised alert is a false alarm or not

3) Ability to predict future events and their implications:
e.g., assess the risk of a vulnerability, decide if an alert should be acted upon

[1] M. Endsley, "Design and Evaluation for Situation Awareness Enhancement", 1988

Anomaly Detection != Situational Awareness

Learning-based anomaly detection deals better with **BAS heterogeneity,** but:

- **Alerts are not actionable** *per se*: we need meaningful context information
- **Learned models are specific to each device**: there is no grouping into semantically equivalent classes

Role Inference

We propose to **infer high-level attributes** from observed data.

Ex. the **role** of a device represents its functional behavior in the network

Understandability is improved:

The role provides meaningful context information to interpret a device's [anomalous] behavior

Adaptability is improved:

When a new device appears on the network, we can apply rules and models based on the device's role

TU/e

BACnet Profiles and Profile Families

BACnet standard already has device *Profiles*, but:

- the profile of a device cannot be read from the network;
- they are based on application domain, not on functional behavior;
- the profile in the specification may not correspond to the behavior in real life [2].

Thus, we define **behavioral roles** based on the functional levels in BAS architecture

BACnet Profile Family	Behavioral Role
Controller	Controller
	Field Device
Lighting Control Stations	Controller
Lighting Controllers	Field Device
Miscellaneous	Router
	Field Device

[2] H. Esquivel-Vargas, “Automatic deployment of specification-based intrusion detection in the BACnet protocol”, 2017

TU/e

Behavioral Roles

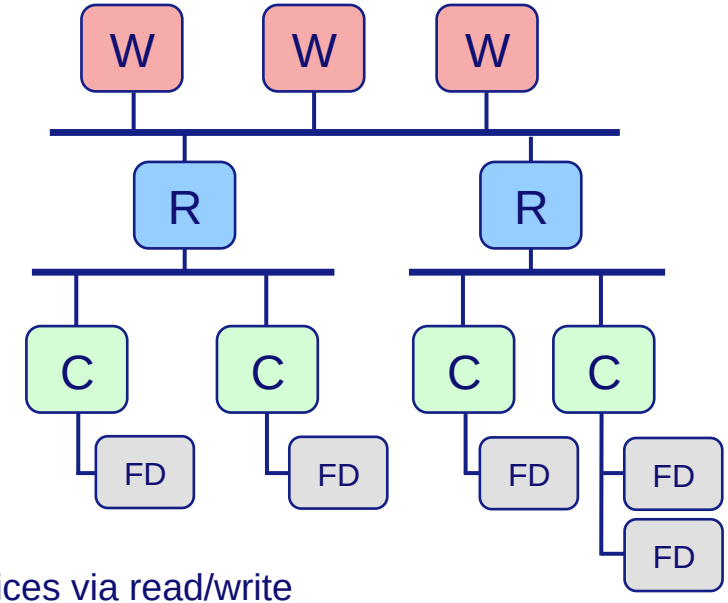

- **W**orkstation:
  Ex. store historical data, inform operators, adjust setpoints
- **R**outer
  Interconnect devices from two or more networks
- **C**ontroller
  Ex. execute the main logic processes, interact with Field Devices via read/write
- **F**ield **D**evice
  Interact with physical environment; they can be connected directly to Controllers, or talk BACnet

Using roles for Situational Awareness

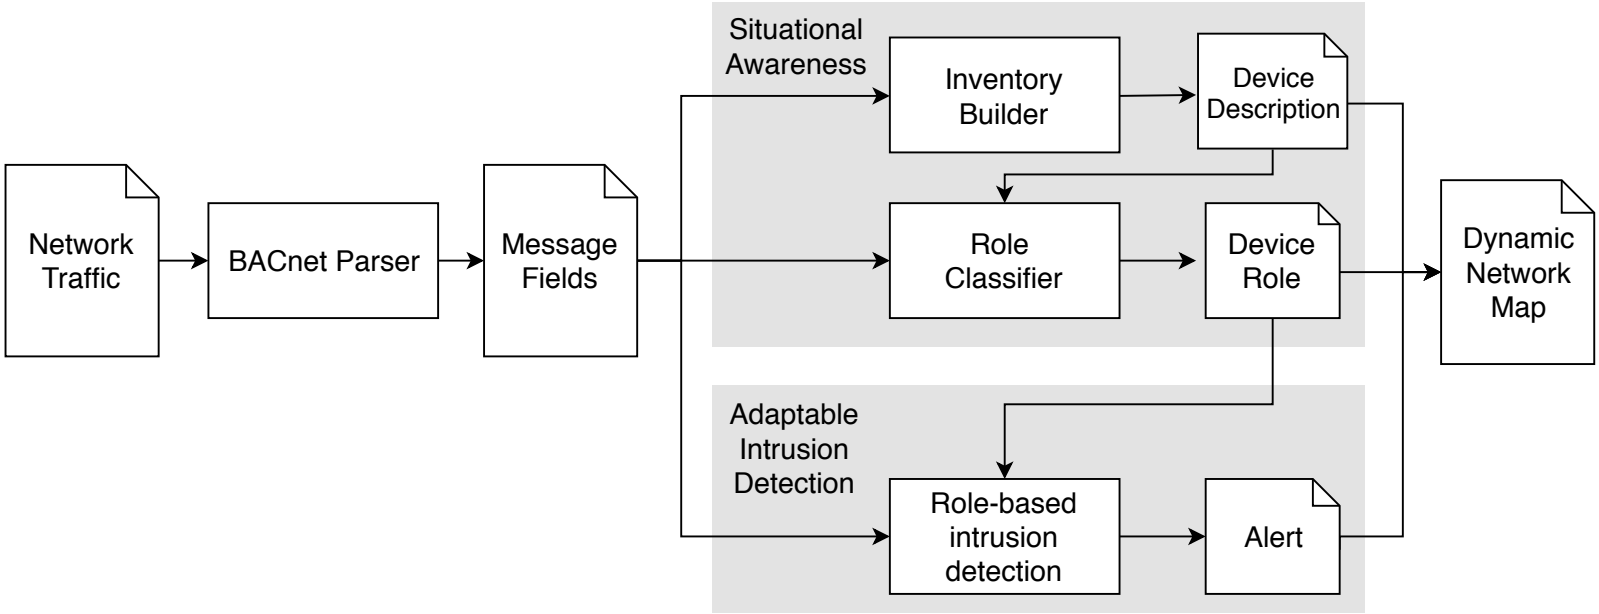

Dynamic Network Map

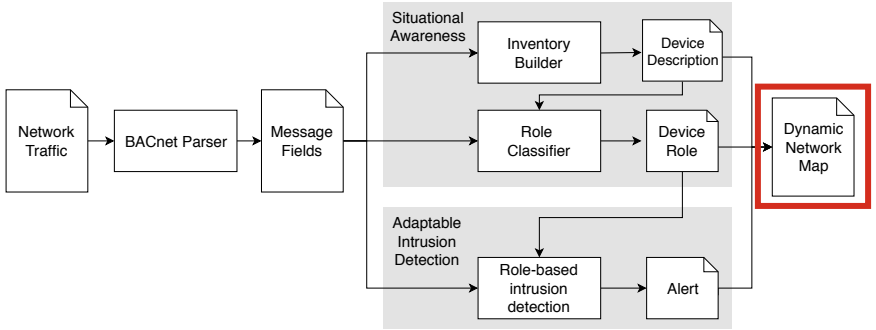

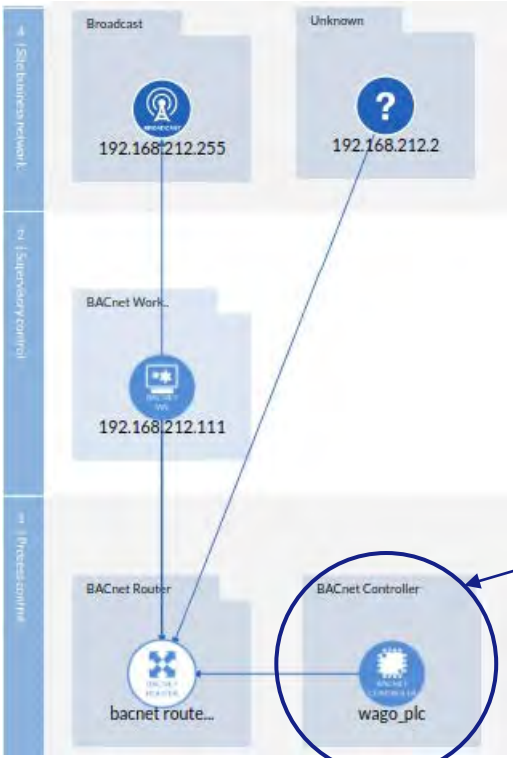

IP address	192.168.212.11 (Private IP)
Host name	wago_plc
Other host names	
MAC addresses	00:30:DE:0B:FE:89 (WagoKont)
Networks	
Role	BACnet Controller
Other roles	Web server
Vendor/model	750-831
Other vendors/models	
OS version	
Client protocol(s)	BACNET (UDP)
Server protocol(s)	BACNET (UDP), FailedConnection (TCP), HTTP (TCP), NotAKnownOne (TCP)
Labels	Device_ID=1

Inventory Builder

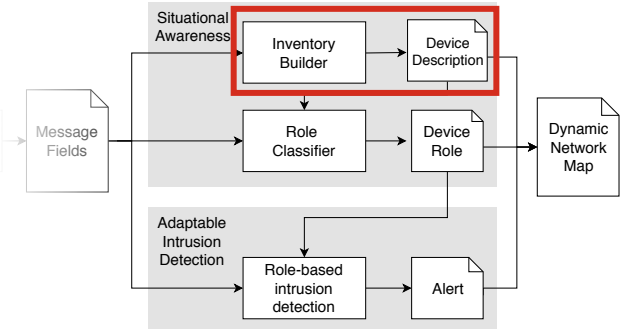

We extract information from the payload of observed BACnet messages:

- Unique ID
- Object Name

} Uniquely **identify** a device

- Vendor Name
- Model Name
- FW Version
- Location
- Data Link Layer
- Is a BBMD
- Is a Foreign Device

} **Describe** a device (configuration, location, etc…)

TU/e

Role Classifier

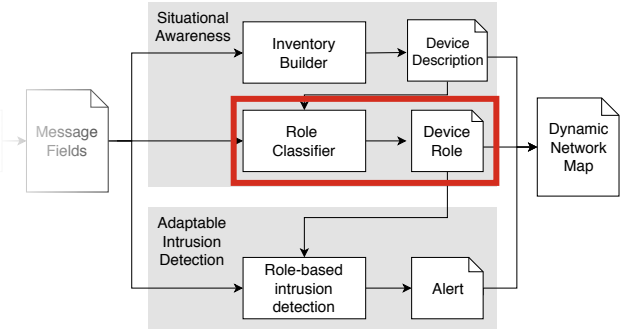

We infer roles with two techniques:

Heuristics based classification (HBC):

We classify devices by checking if their observed behavior contains patterns unique to a role:

- Only Workstation devices should initiate a `WritePropertyMultiple` request
- Only Routers forward messages from other networks

Distance based classification (DBC):

We classify remaining devices by their distance to previously classified devices, using:

- Vendor ID
- Model Name
- Data Link Layer type

Classification Results

We evaluated discovery and classification on a **real-life dataset** from a university campus (106GB, 9 days of traffic, ~20 million BACnet pkts)

Dataset 2	
Role	Ground truth
Controller	219
Router	21
Workstation	1
Total	241

HBC		
Classification	TP	FP
213	212	1
21	21	0
0	0	0
234	233	1

HBC + DBC		
Classification	TP	FP
220	219	1
21	21	0
0	0	0
241	240	1

- HBC+DBC discovers all devices
- One misclassification: Workstation had behavior consistent with a Controller
- Using this model for intrusion detection, Workstation might raise false alerts (but role helps interpret them)

TU/e

Role-based Intrusion Detection

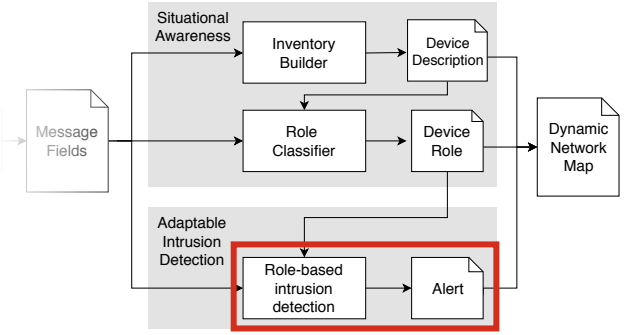

Roles (and other **high-level attributes**) can be used as **features** for different IDS modules:

- Learning **role-based behavior**:

  "All Controllers send beween 0 and 60 `ReadProperty` requests per hour"

- Specifying **attribute-based policies** and **consistency checks**(*):

  "Field Devices cannot initiate `WriteProperty` requests"

  "Devices with Vendor XYZ cannot be Controllers"

(*) Consistency checks help in finding misconfigured or misclassified devices

Intrusion Detection Results

We extend previous results[3] by detecting two previously undetected attacks:

✔ **Snooping by new Controller**: it sends abnormally many `ReadProperty` requests for its role

✔ **Tampering by Field Device:** it sends a `WriteProperty` request

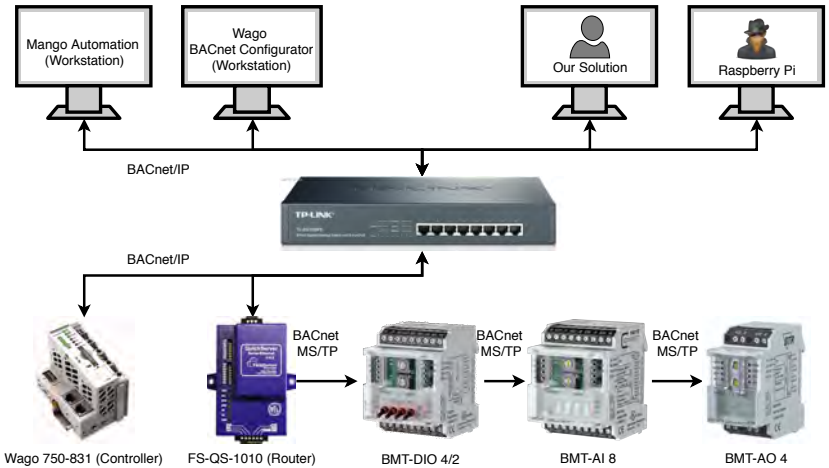

Evaluation of our IDS on the real-life dataset showed good results for **usability** (~6.4 FP/h) and **adaptability** to new devices (~0.1 FP/h increase after cross validation)

[3] D. Fauri et al., “Leveraging Semantics for Actionable Intrusion Detection in Building Automation Systems”, CRITIS ‘18

TU/e

Conclusion

- We propose the use of **high-level attributes (ex. roles)** for enriching situational awareness in heterogeneous systems;
- Roles improve **actionability** of alerts and **adaptability** of detection systems;
- We intend to improve the granularity of this approach, and extend it to other domains (ex. ICS) or other attributes