Comosum: An Extensible, Reconfigurable, and Fault-Tolerant IoT Platform for Digital Agriculture

Gloire Rubambiza, Shiang-Wan Chin, Sachille Atapattu, Mueed Rehman, José F. Martínez, Hakim Weatherspoon

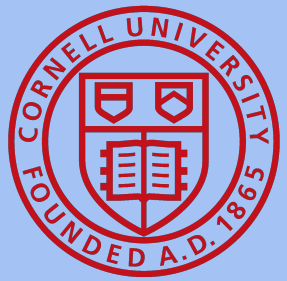

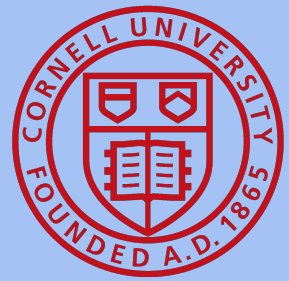

Gloire Rubambiza

Shiang-Wan Chin

Sachille Atapattu

Mueed Rehman

José F. Martínez

Hakim Weatherspoon

Comosum Technical Contributions

4 Years

18 months

1M+ sensor readings

- Distilled black-box SOTA down to a single interface
- Applied strong systems approaches to new contexts
- Deployed across different farm types and cloud providers

Background

DA Challenges & State of the art

Comosum Design & Implementation

Deployment Experiences, Insights & Limitations

Conclusion

What is Digital Agriculture (DA)?

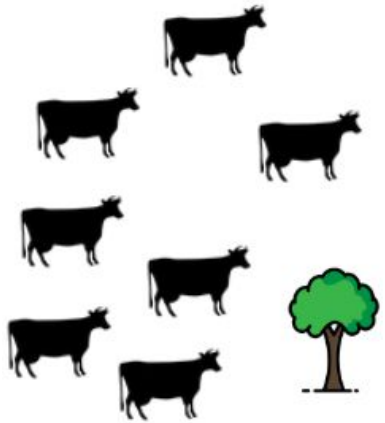

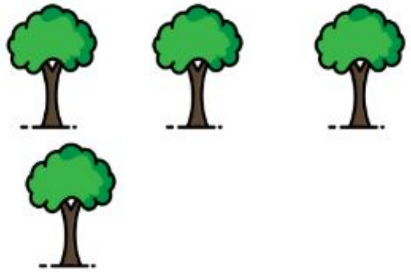

Growing ability to… convert precise data… into actionable knowledge to … support complex decision-making on farms

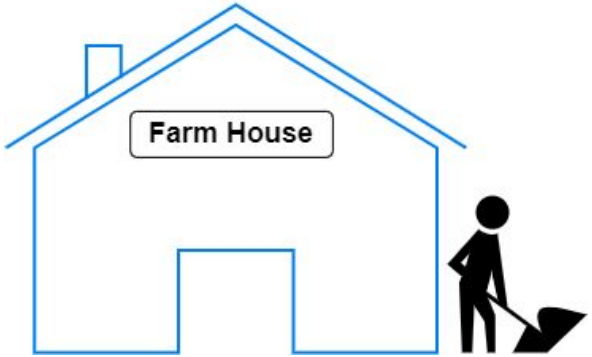

Priorities for Science to Overcome Hurdles Thwarting the Full Promise of the 'Digital Agriculture' Revolution. Shepherd *et al*. 2018

What is Digital Agriculture (DA)?

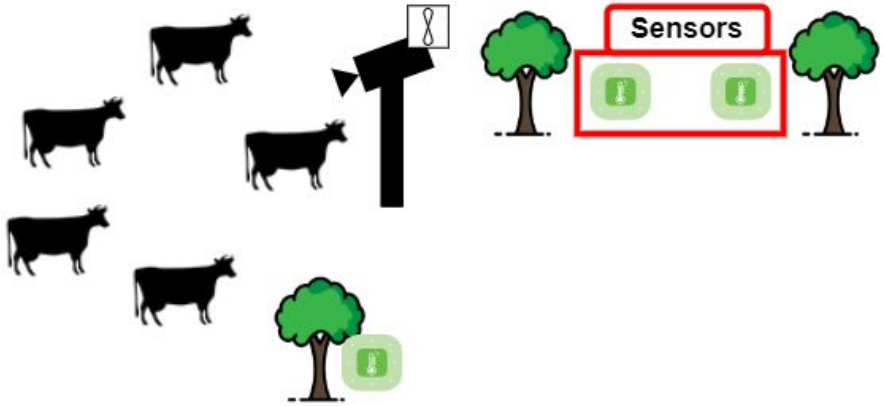

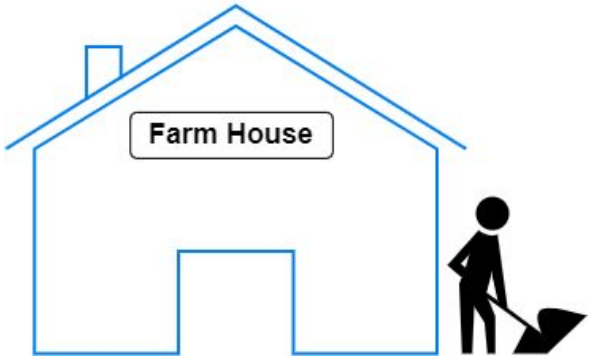

What is Digital Agriculture (DA)?

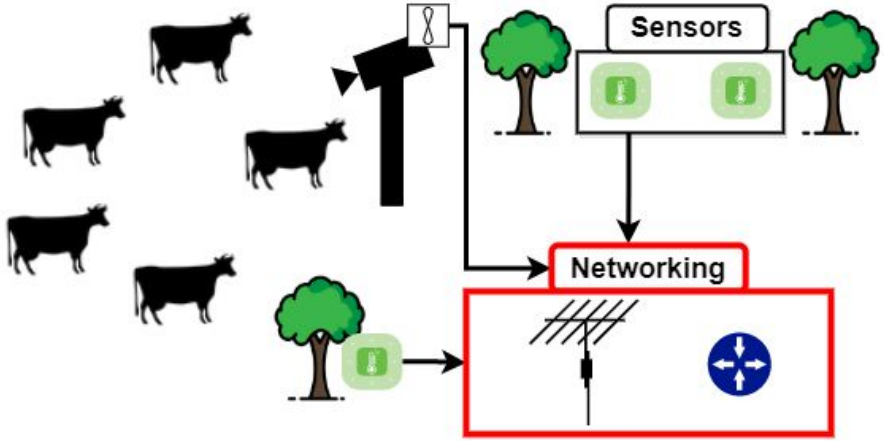

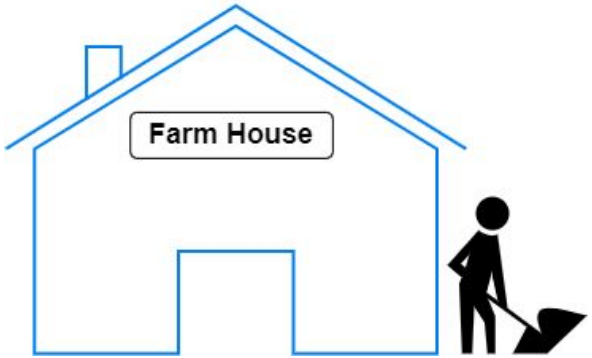

What is Digital Agriculture (DA)?

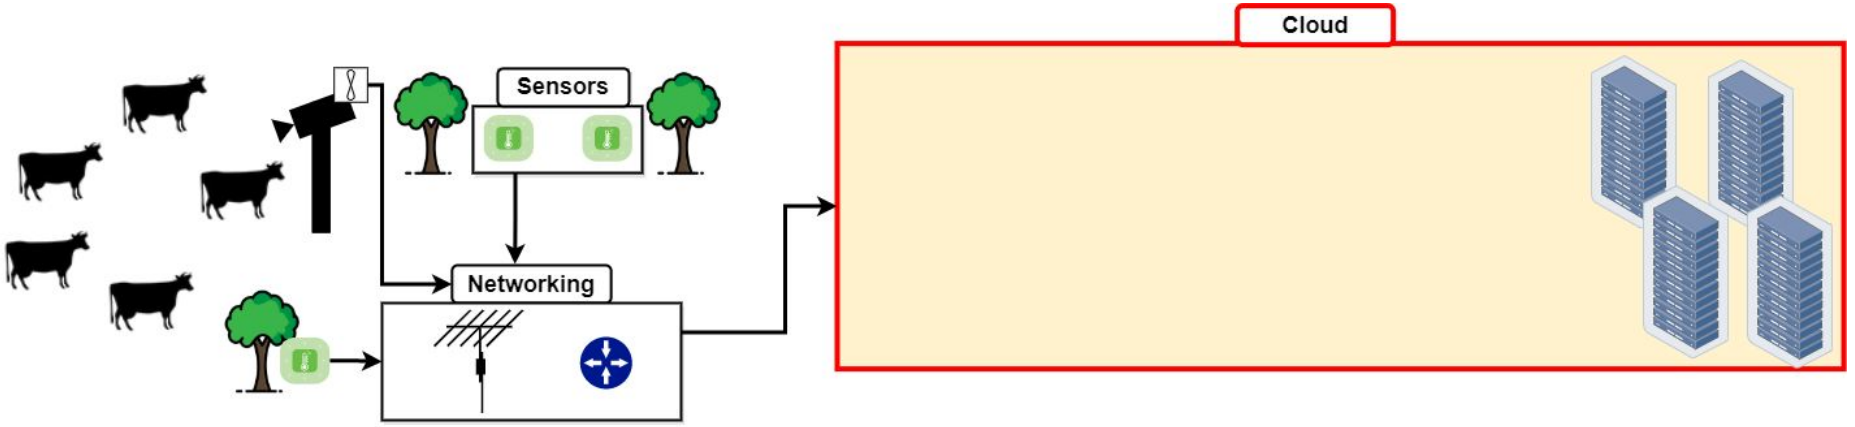

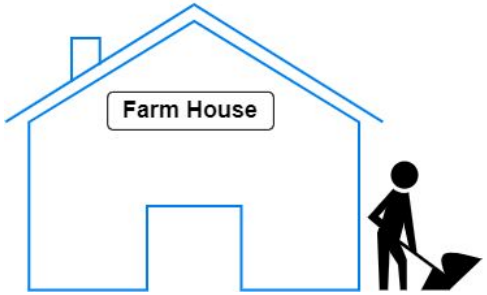

What is Digital Agriculture (DA)?

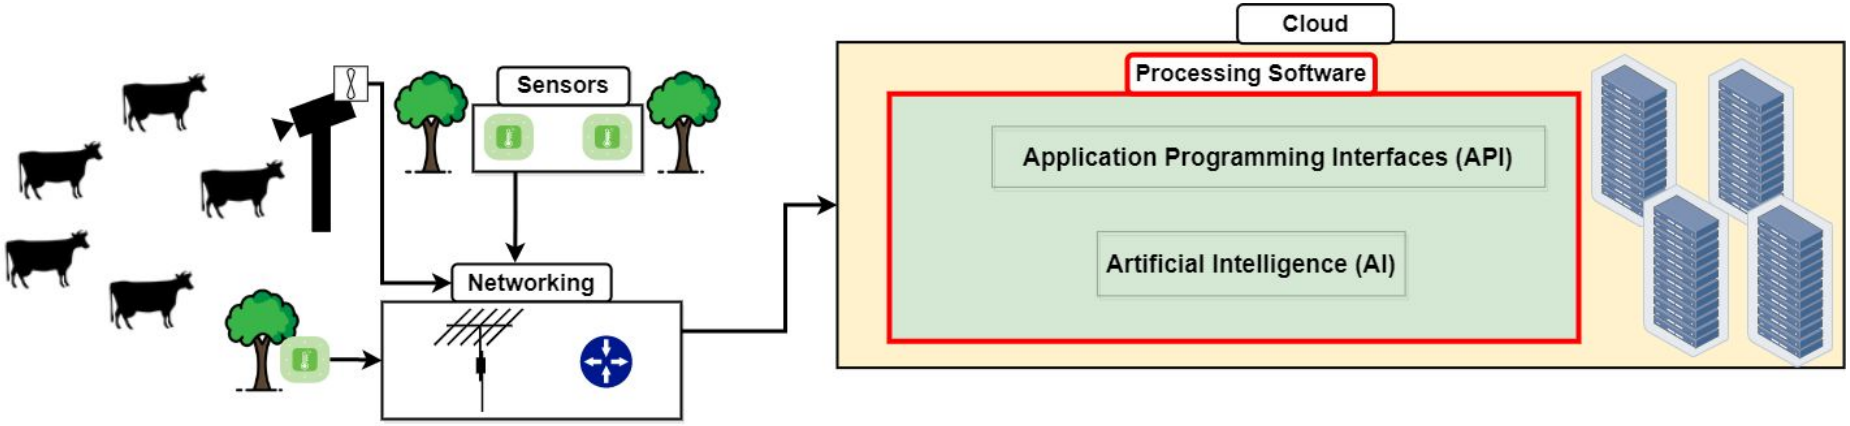

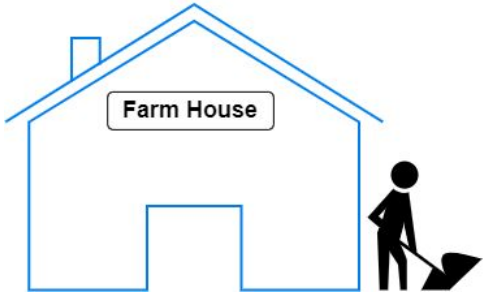

What is Digital Agriculture (DA)?

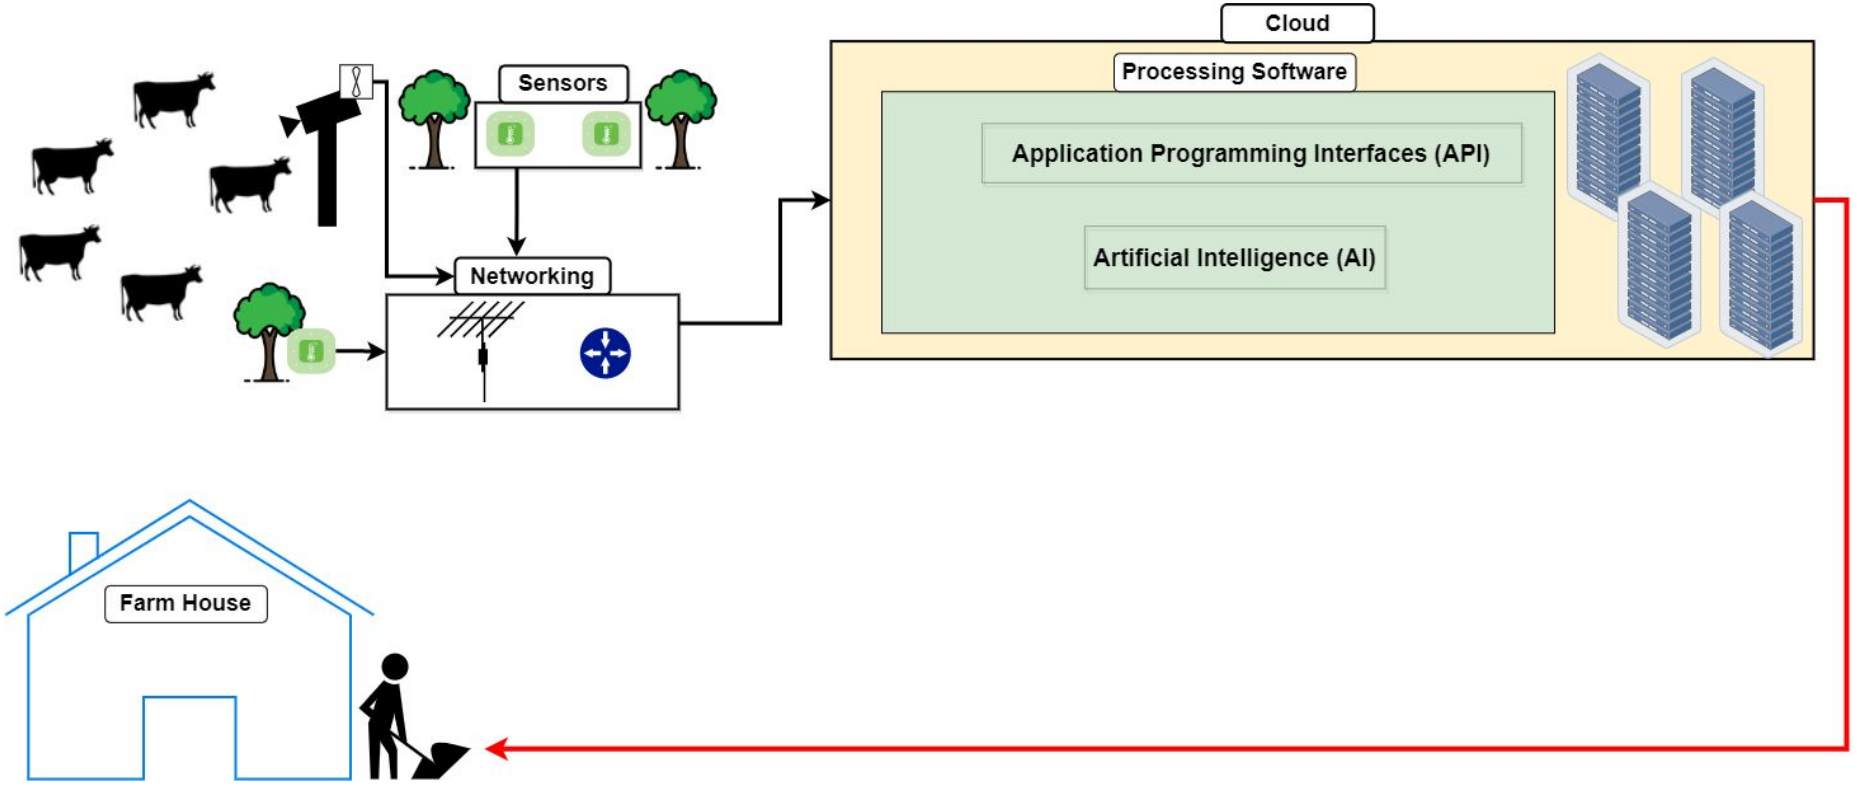

Background

DA Challenges & State of the art

Comosum Design & Implementation

Deployment Experiences, Insights & Limitations

Conclusion

Motivating Challenges in DA

Greenhouse: KB/hour

The Effects of Light-emitting Diode Lighting on Greenhouse Plant Growth and Quality. Olle & Viršile. 2013

Motivating Challenges in DA

Greenhouse: KB/hour

Dairy Farm: MB/day

Image Credit: Martin Perez

Effect of Automating Health Monitoring on Detection of Health Disorders and Performance of Lactating Dairy Cows. Perez et al. 2021

Motivating Challenges in DA

Greenhouse: KB/hour

Dairy Farm: MB/day

Image Credit: Martin Perez

Vineyards: TB/year

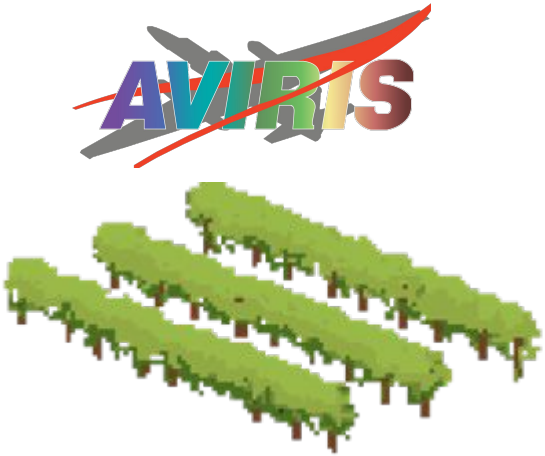

Image Credit: Fernando Romero Galvan

Scalable early detection of grapevine virus infection with airborne imaging spectroscopy. Romero Galvan *et. al.* 2023

Rural Infrastructure Challenges

- Sparse population & geography
  - Internet connectivity
  - Public transportation
  - Power

Rural HCI Research: Definitions, Distinctions, Methods, and Opportunities. Hardy *et al*. 2019
Scaling Community Cellular Networks with CommunityCellular Manager. Hasan *et al*. 2019

Rural Infrastructure Challenges

- Sparse population & geography

- Scale & distance limit access to:
  - Provider maintenance and repair
  - Repair parts for self-repair

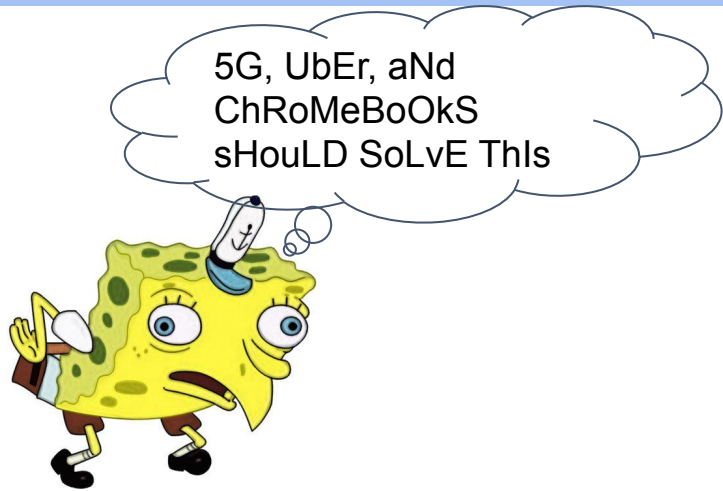

Image Credit: https://knowyourmeme.com/memes/mocking-spongebob

Thinking Relationally about Digital Inequality in Rural Regions of the U.S. Jenna Burrell. 2018

Rural Infrastructure Challenges

- Sparse population & geography
- Scale & distance
- Service provider tactics lead to:
  - Closed system and manuals
  - Reselling sensor data and resulting insights

Rural Infrastructure Challenges

- Sparse population & geography
- Scale & distance
- Vendor lock-in

Critical analysis of vendor lock-in and its impact on cloud computing migration: a business perspective. Opara Martins *et al*. 2020

Research Question:

How do we build **data-intensive** IoT apps on top of **unreliable** rural infrastructure?

Infrastructure Challenges ➡ Desired Technical Features

- Sparse population & geography ➡
  - Reconfigurable networks
  - Energy efficiency
- Scale & distance ➡
  - Fault tolerance
    - Failure detection
    - Independent failures
    - Off-the-shelf (OTS) parts
- Vendor lock-in ➡
  - Open-source X-ware
  - Accessible data analytics

Background

DA Challenges & State of the art

Comosum Design & Implementation

Deployment Experiences, Insights & Limitations

Conclusion

State of the art: Sensing

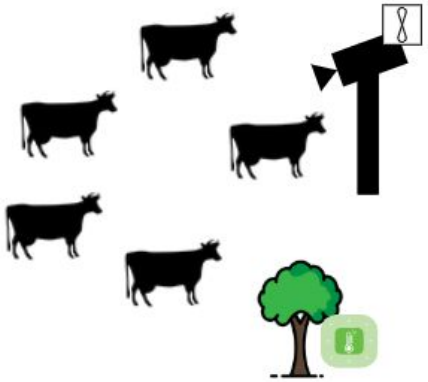

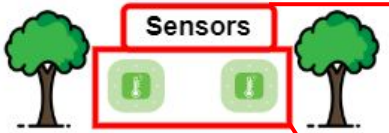

Bandwidth: 50kbps - 10Mbps

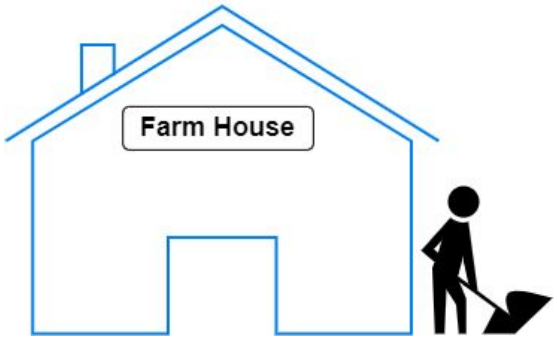

The Platforms Enabling Wireless Sensor Networks. Hill *et al*. 2004

State of the art: Sensing

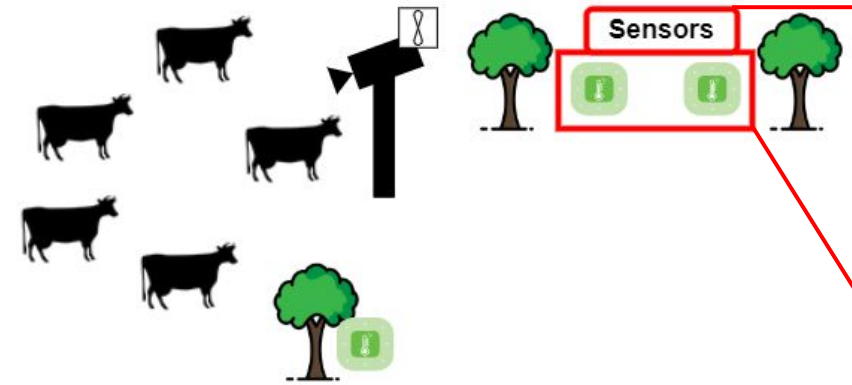

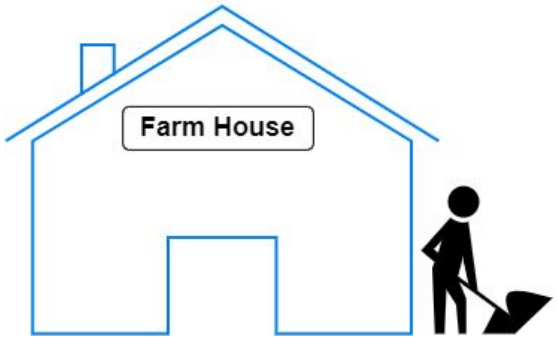

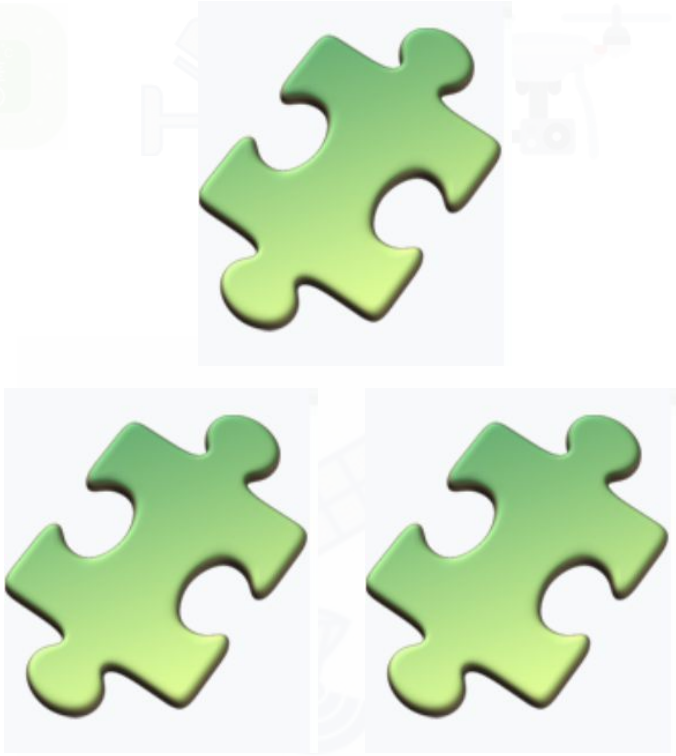

State of the art: Cloud computing

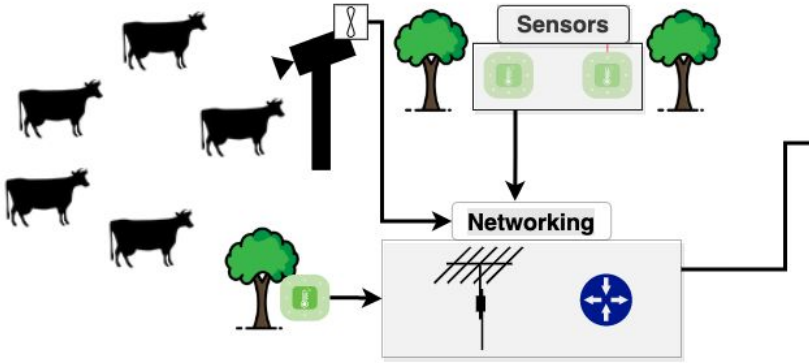

Core Cloud

Cloud computing

Ubiquitous, convenient, on-demand network access to storage media and compute devices.

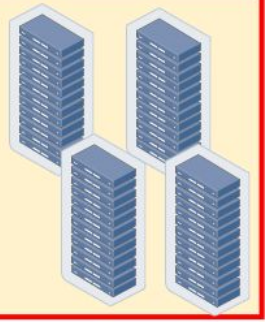

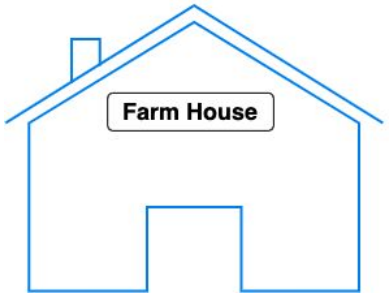

The NIST Definition of Cloud Computing. Mell & Grance. 2011.

State of the art: Edge computing

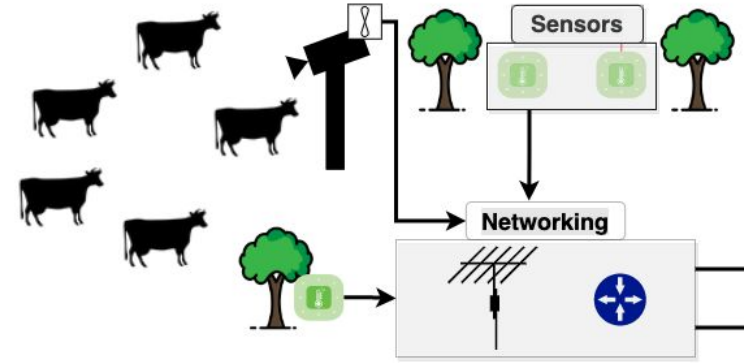

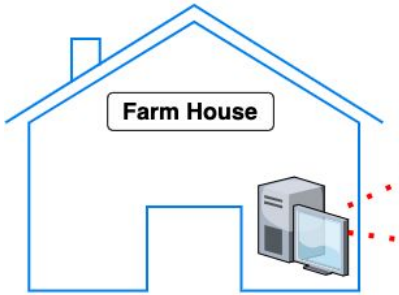

Core Cloud

Cloud computing

Ubiquitous, convenient, on-demand network access to storage media and compute devices.

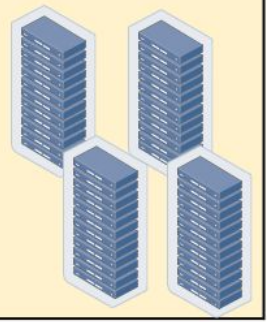

Edge Cloud

Edge computing

A distributed computing framework that brings applications closer to data sources

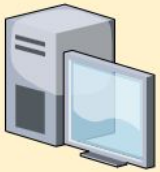

What is Edge Computing. IBM. 2022.

The distributed cloud avoids vendor lock-in

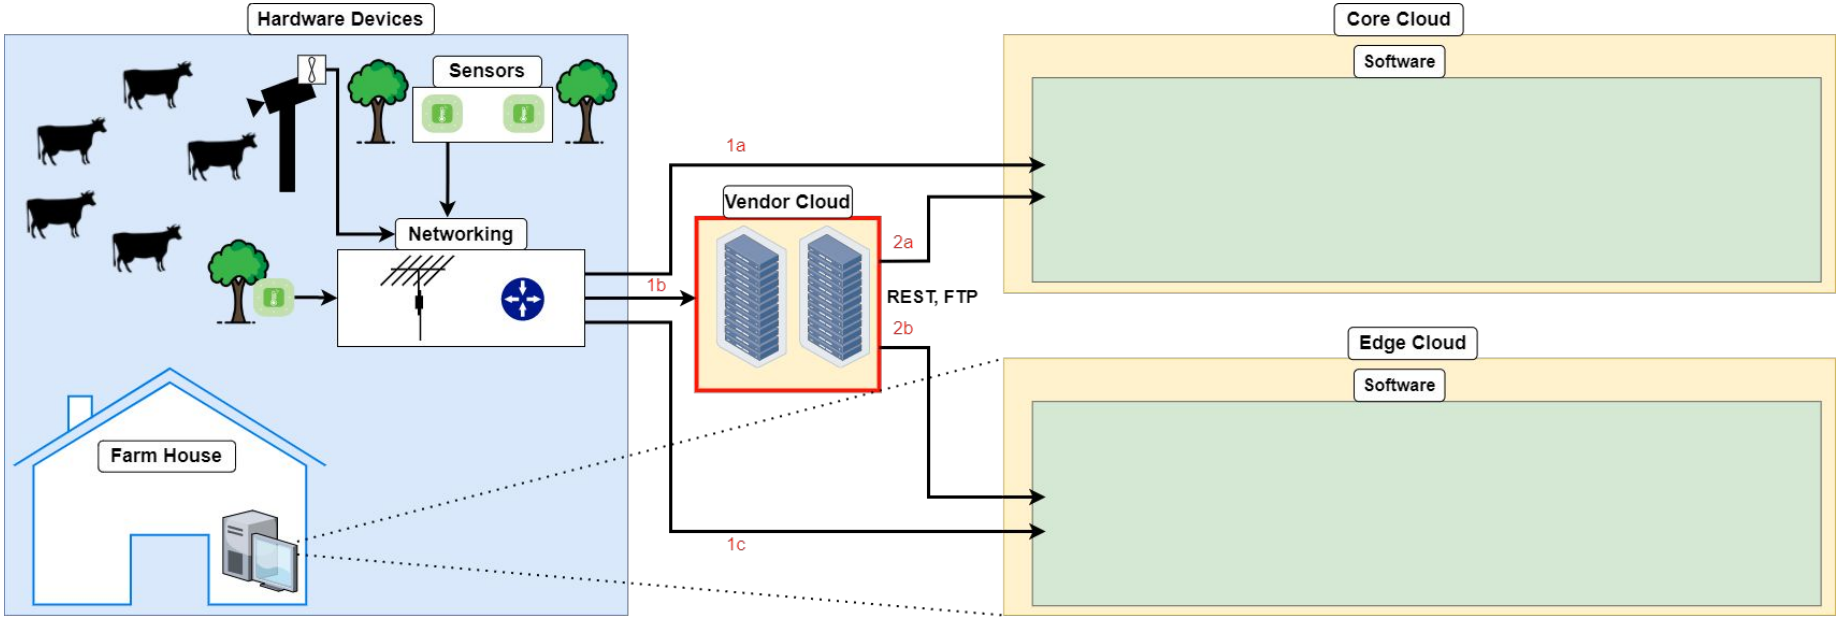

Revised Research Question:
How do we build an ***extensible***, ***reconfigurable***, and ***fault tolerant*** IoT platform on top of **unreliable** base infrastructure?

Building **Comosum** on unreliable base infrastructure

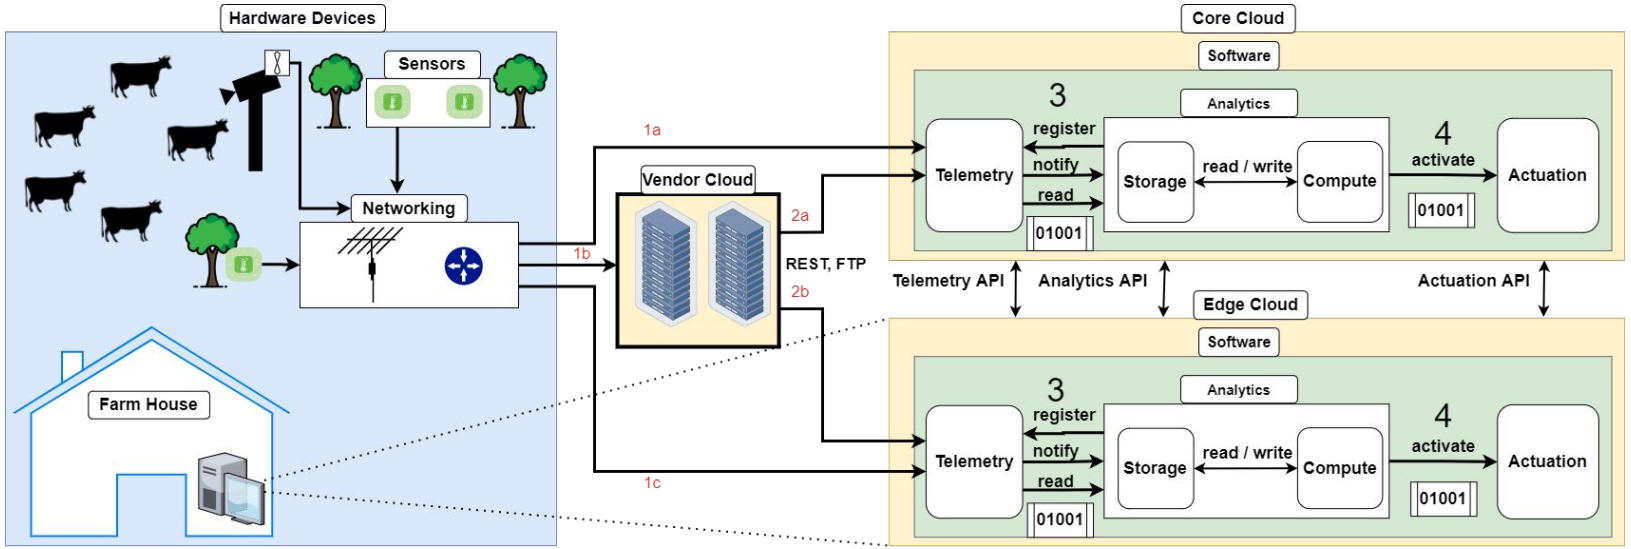

Research Contribution: Design, implementation, and deployments of a reconfigurable platform for DA experimentation

Background

DA Challenges & State of the art

Comosum Design & Implementation

Deployment Experiences, Insights & Limitations

Conclusion

Backyard deployment
Oct 2020

TVWS config
Mar 2021

Dairy farm visit
Feb 2023

Comosum Design aka The Software Defined Farm (SDF)

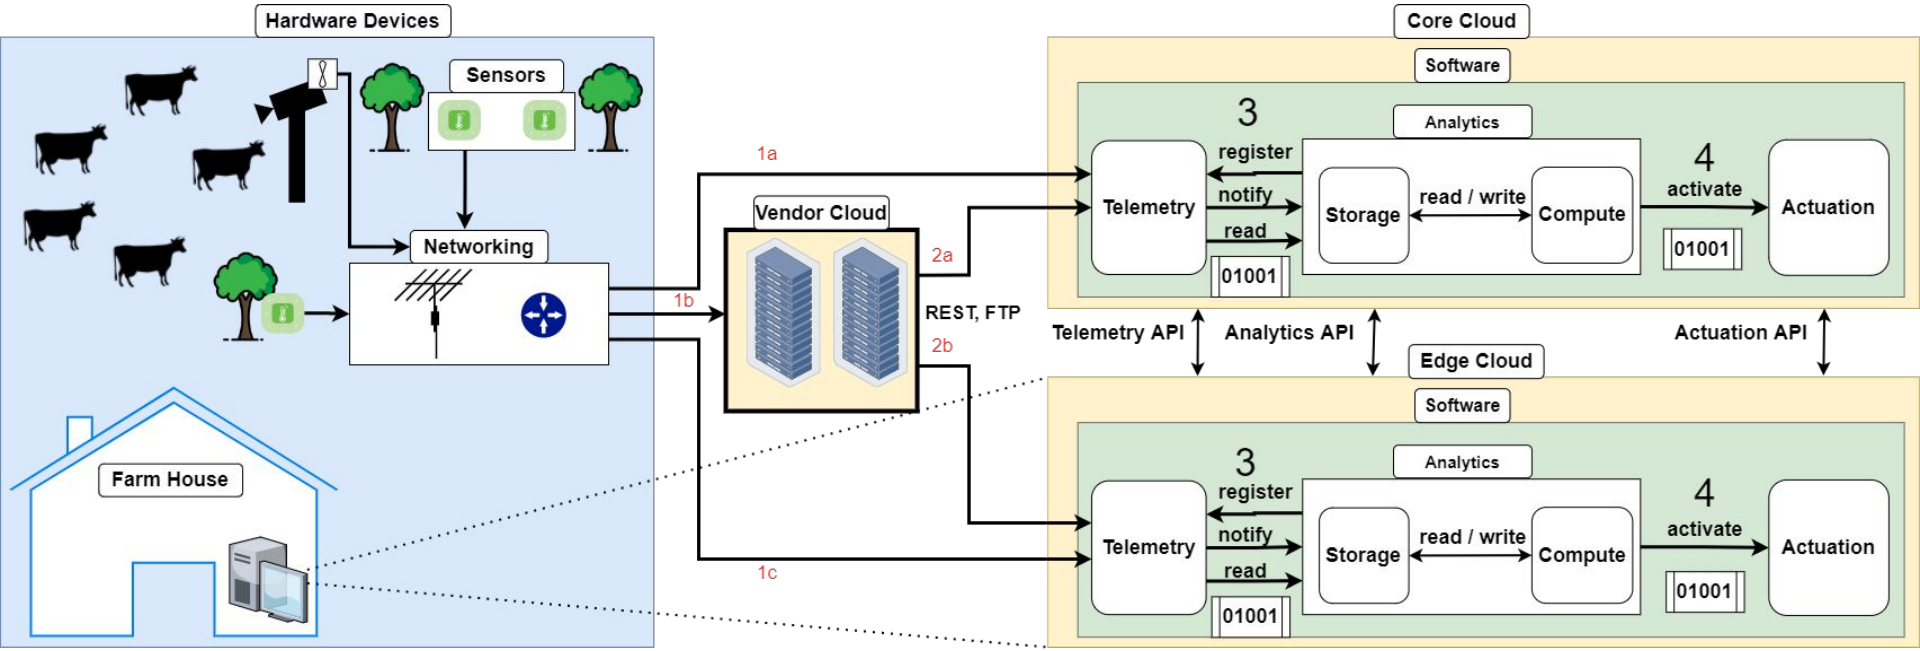

Comosum Design: Telemetry Module

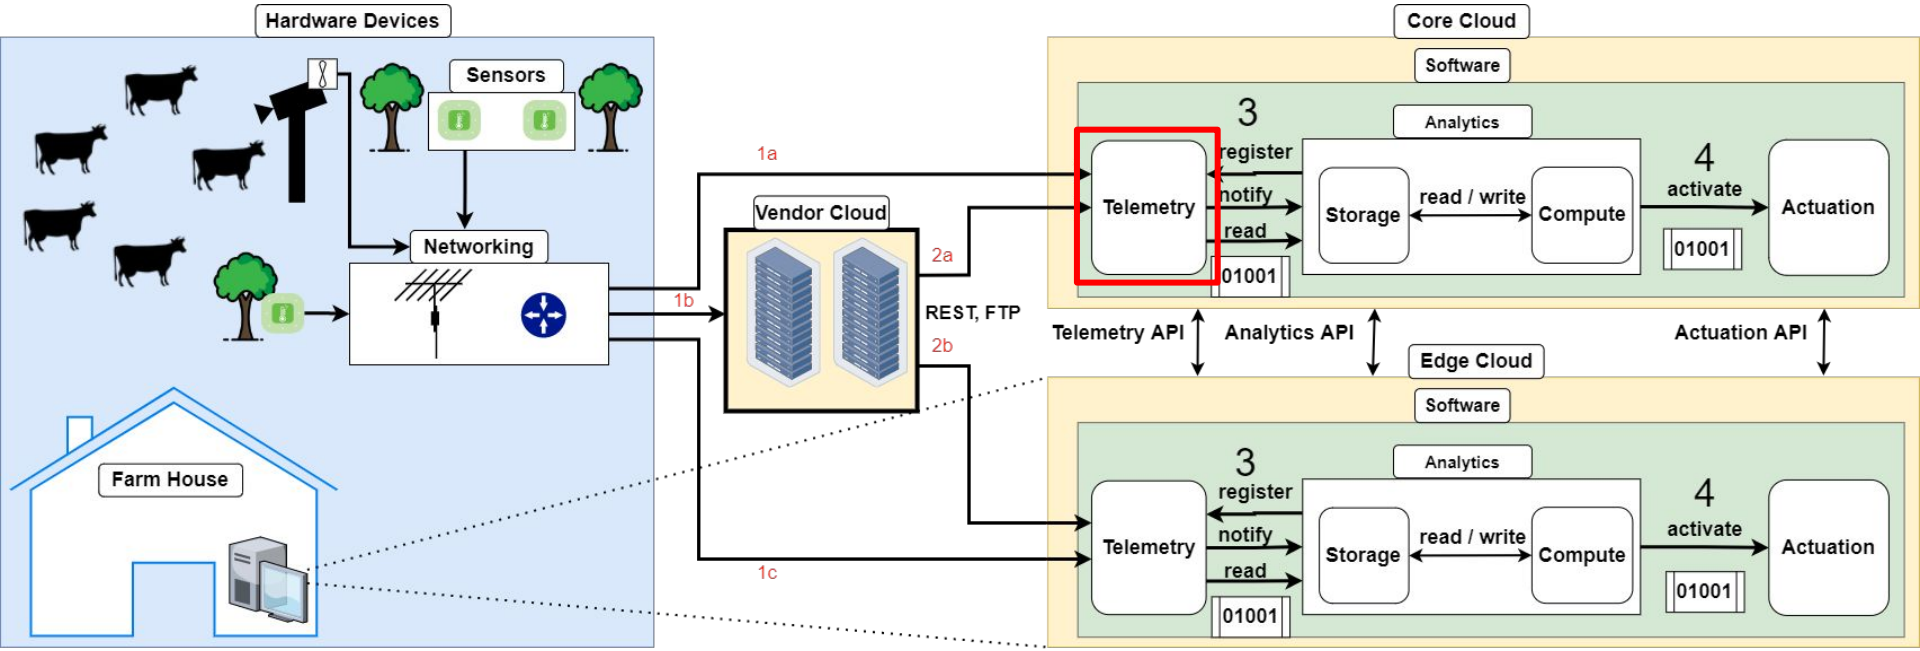

Comosum Design: The Analytics Module

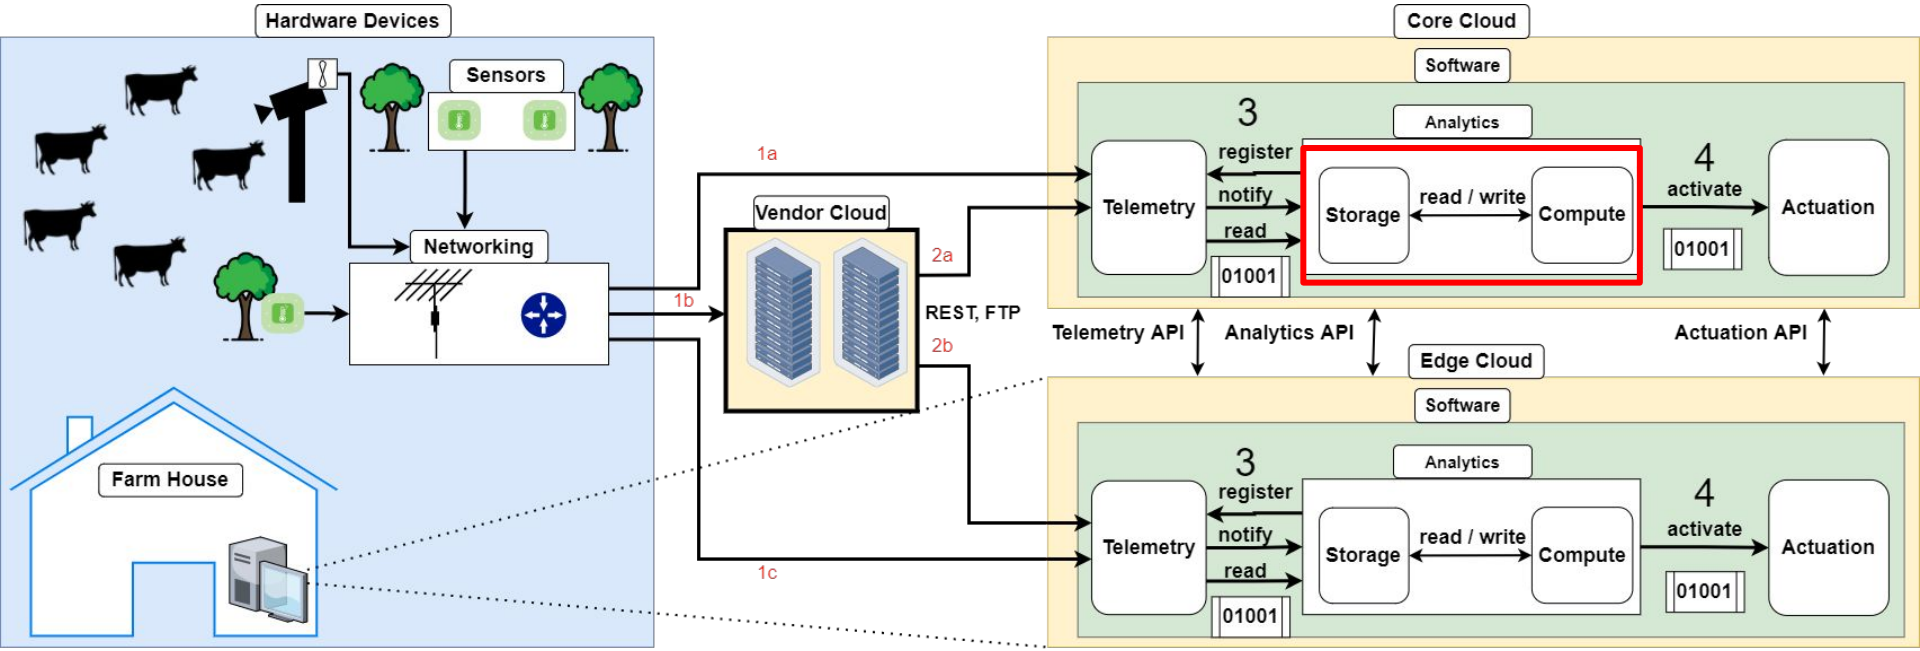

Comosum Design: The Actuation Module

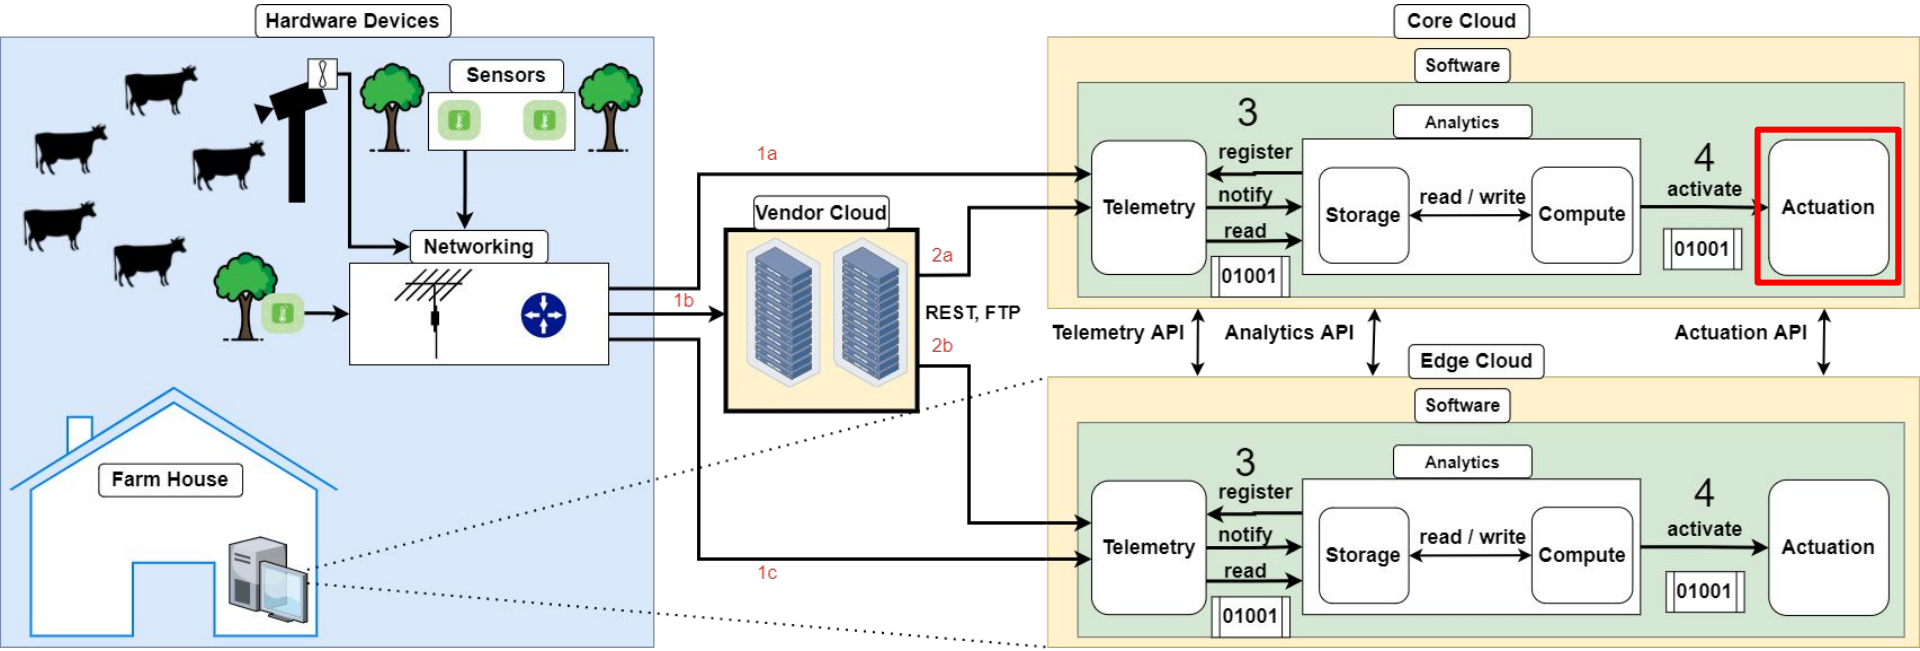

Comosum Design: Extensible Data Plane

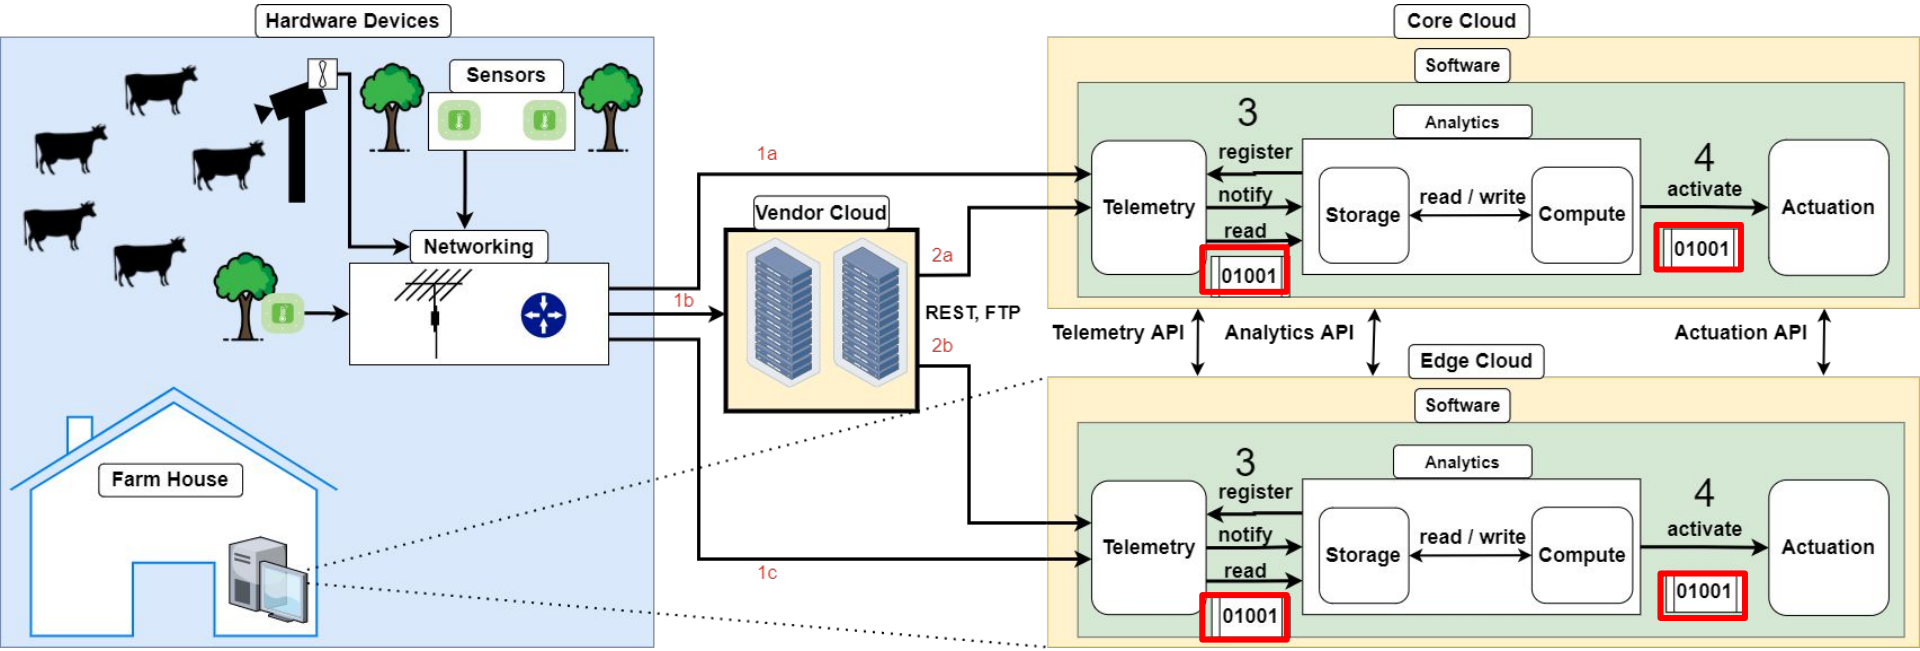

Comosum Design: Reconfigurable Control Plane

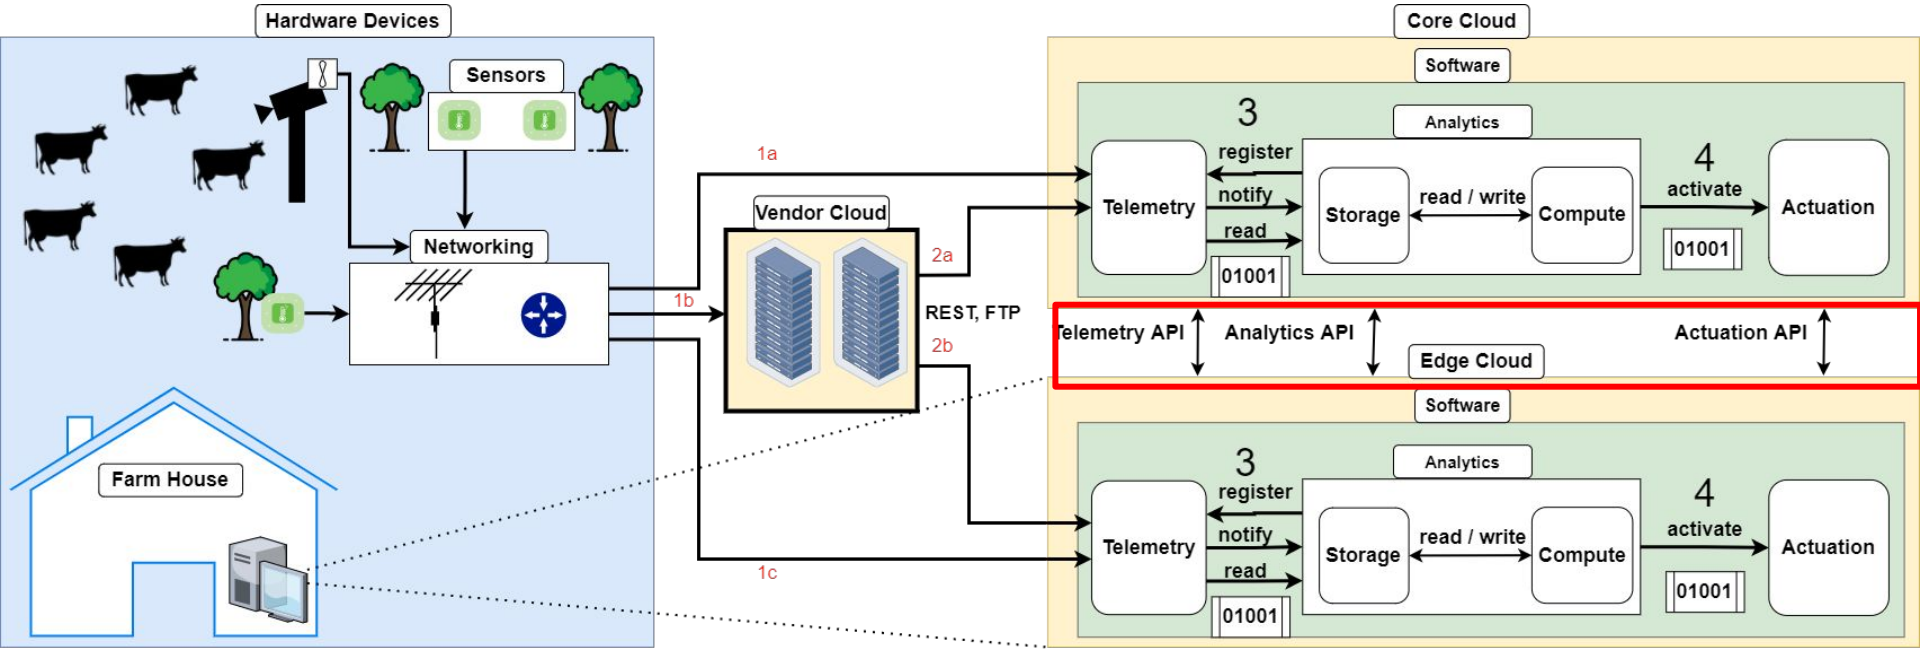

FarmBIOS: A Comosum Implementation

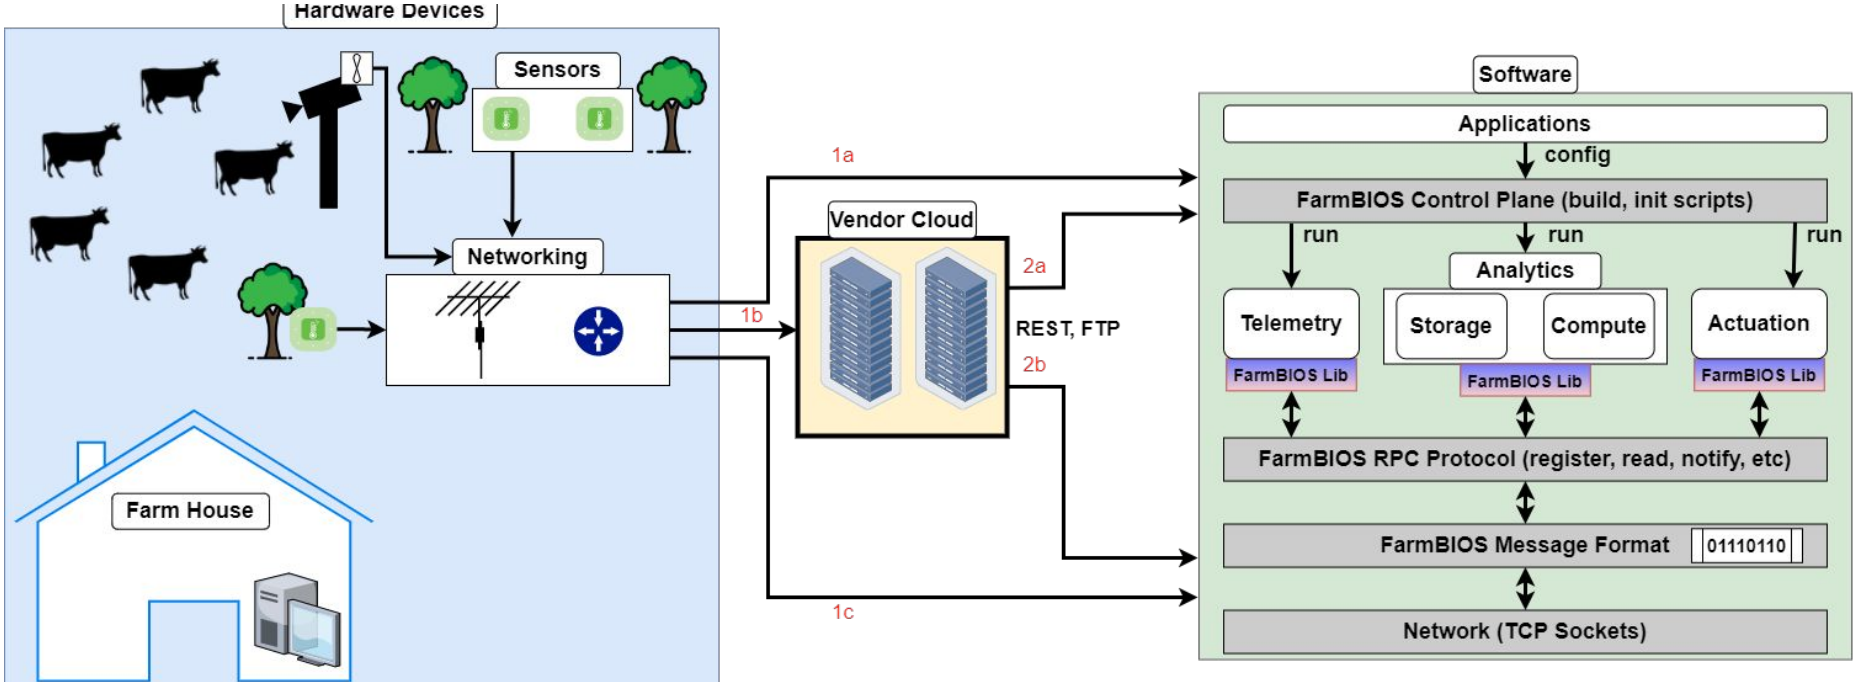

FarmBIOS: A Comosum Implementation

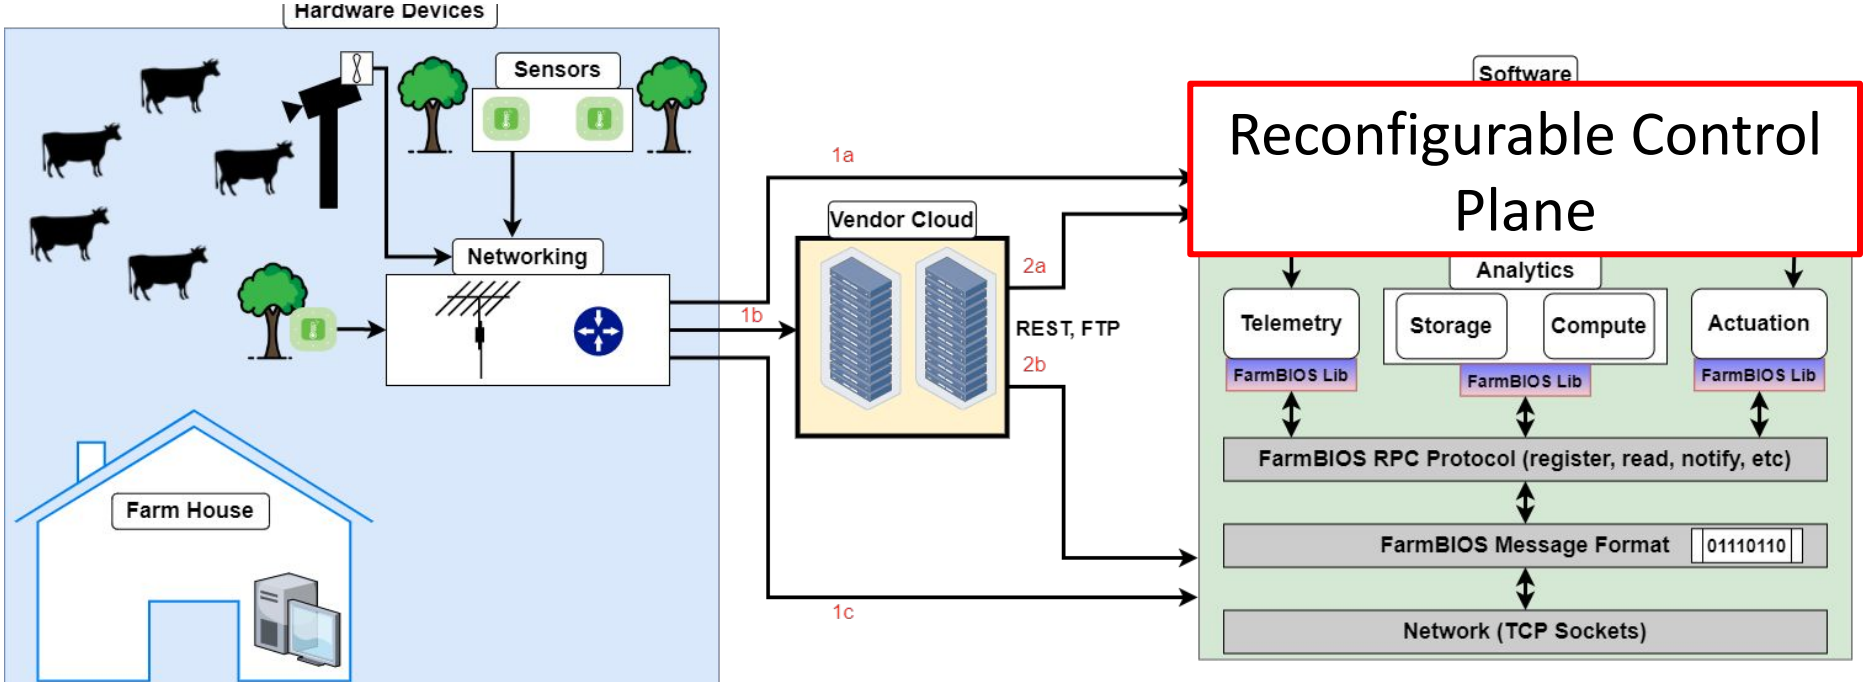

FarmBIOS: A Comosum Implementation

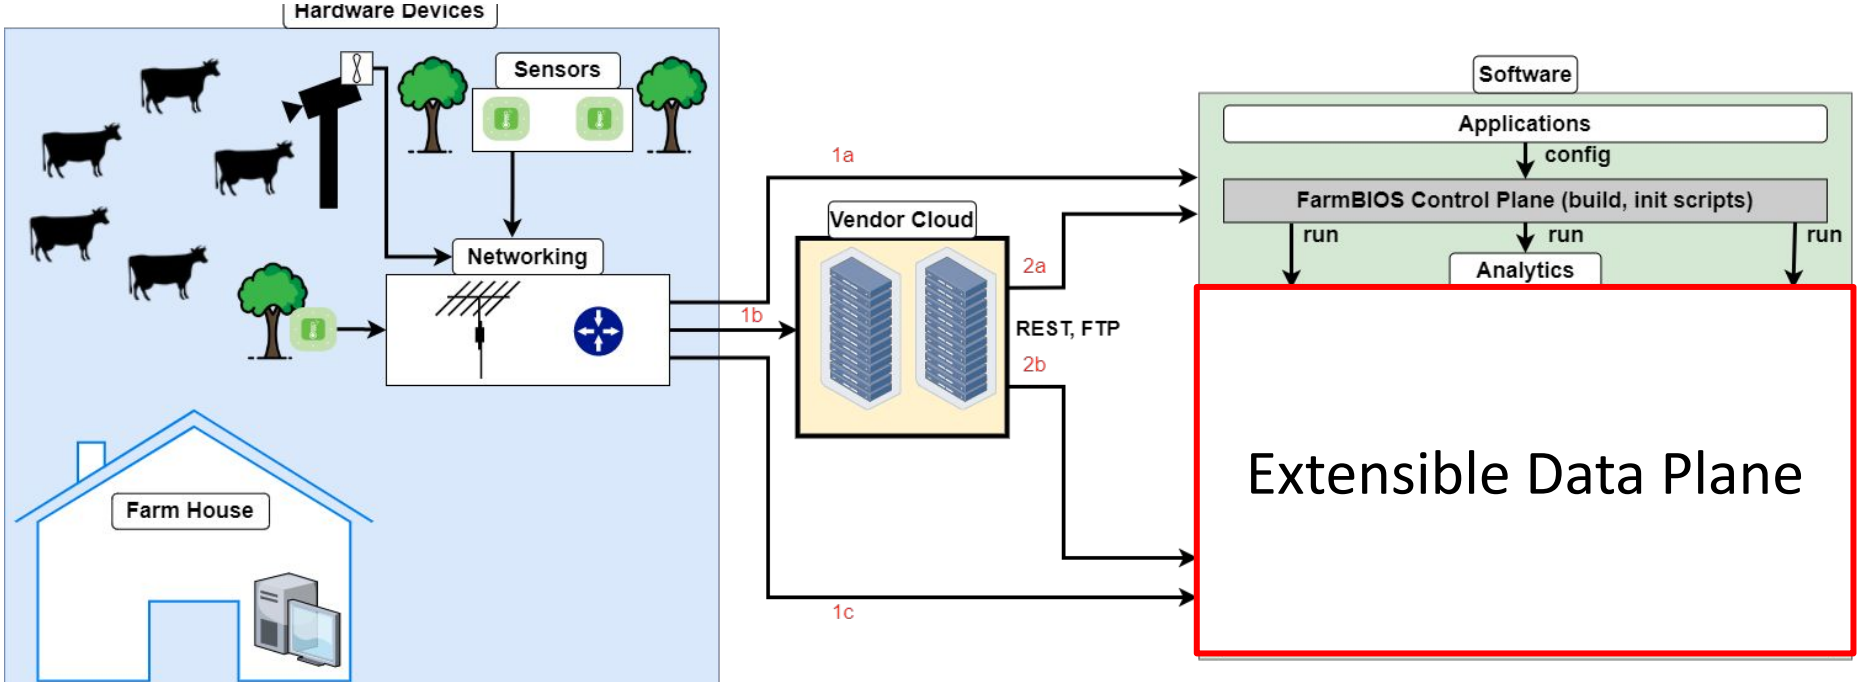

FarmBIOS: A Comosum Implementation

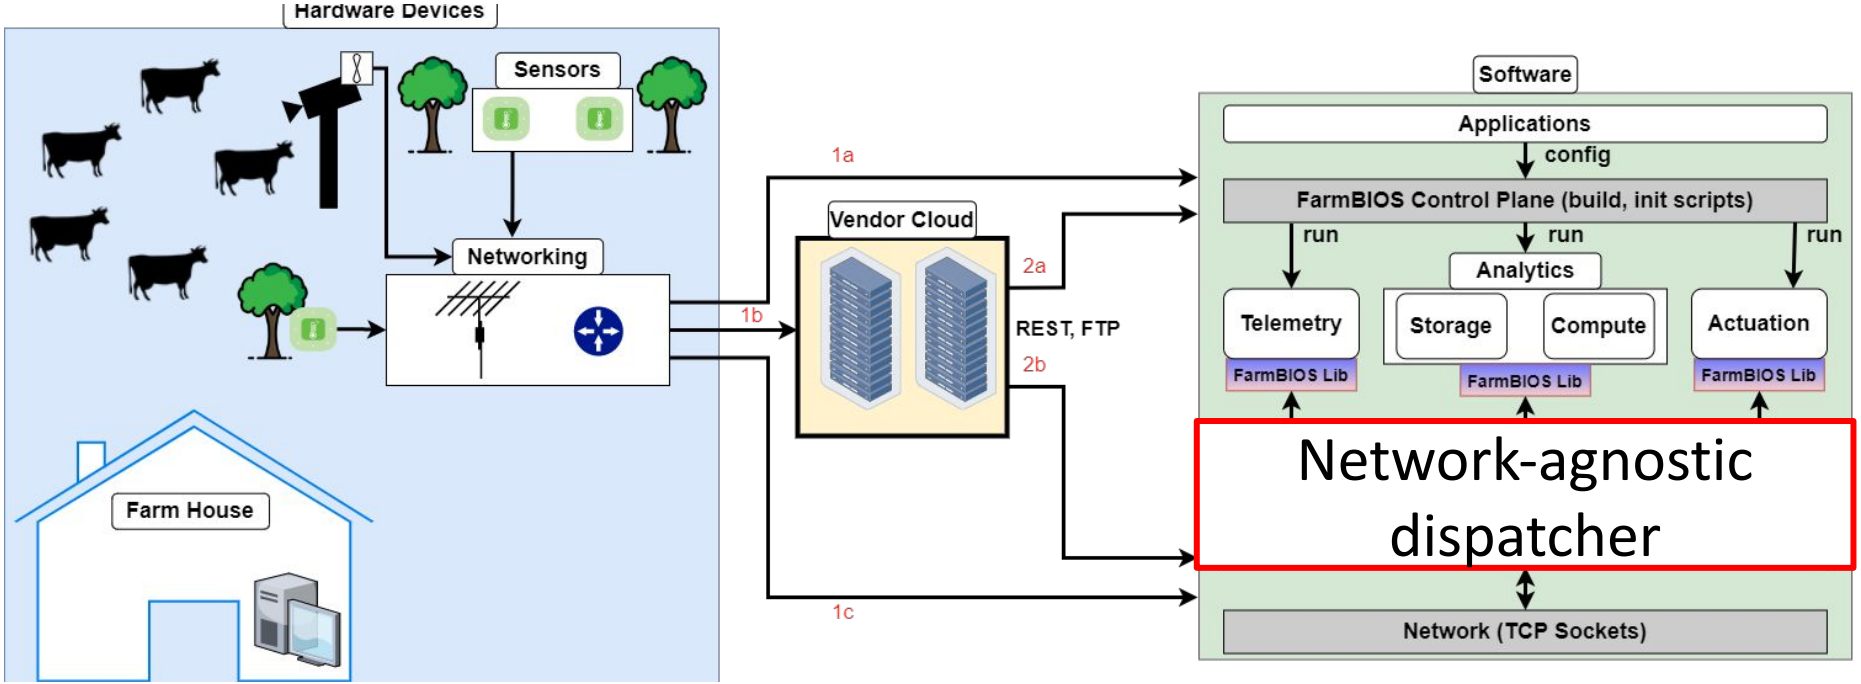

Background

DA Challenges & State of the art

Comosum Design & Implementation

Deployment Experiences, Insights & Limitations

Conclusion

FarmBIOS Applications

WaterGuard

CowsOnFitbits

Image Credit: Martin Perez

WineGuard

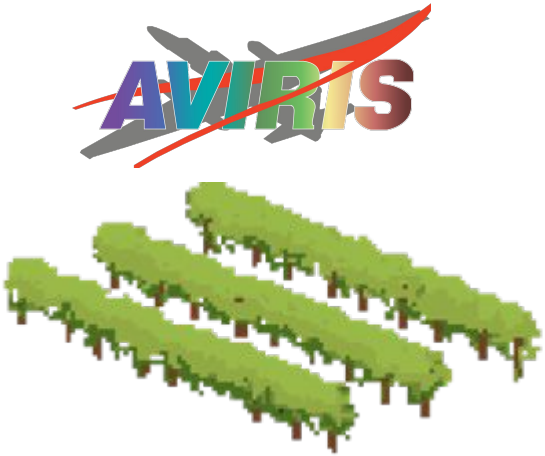

Image Credit: Fernando Romero Galvan

FarmBIOS Deployments: WineGuard

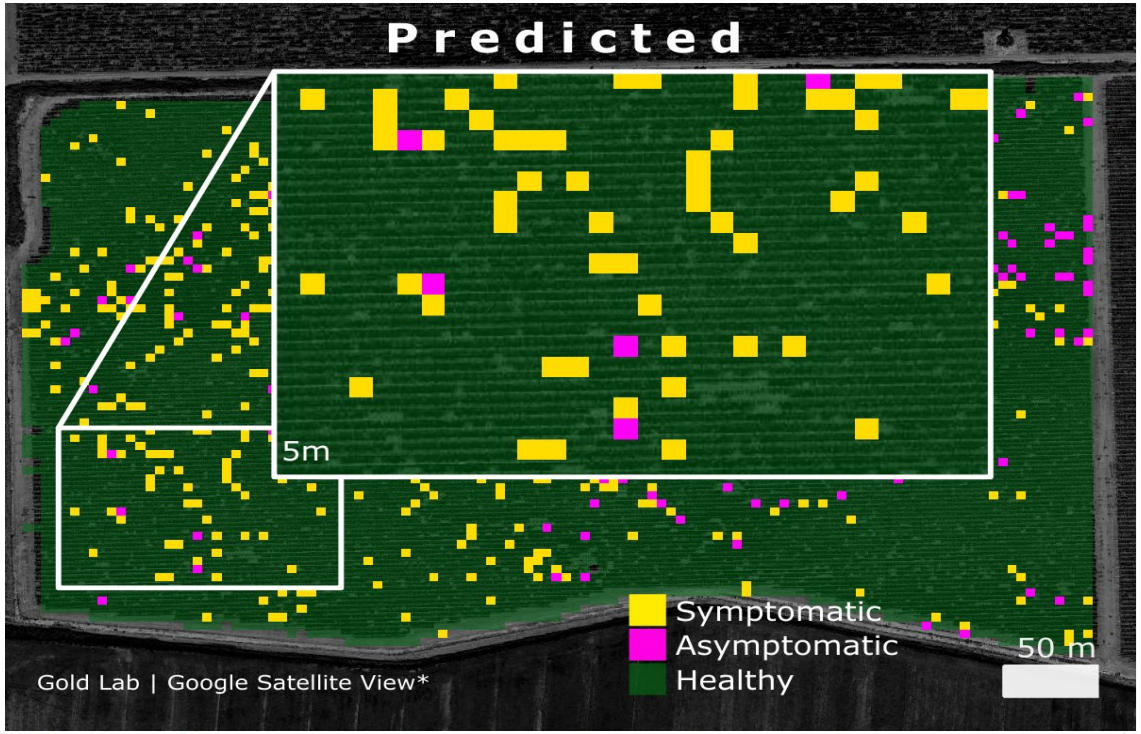

Training	Data Location	Data Size	Compute	Storage	Runtime	Accuracy
Edge	Local	4MB	8 CPUs	256GB	27.1s	82%
Edge	Cloud	-	8 CPUs	256GB	35.6s	84%
Azure ML	Cloud	-	2 vCPUs	100GB	86.5s	84%

FarmBIOS Deployments: WaterGuard

Base Station

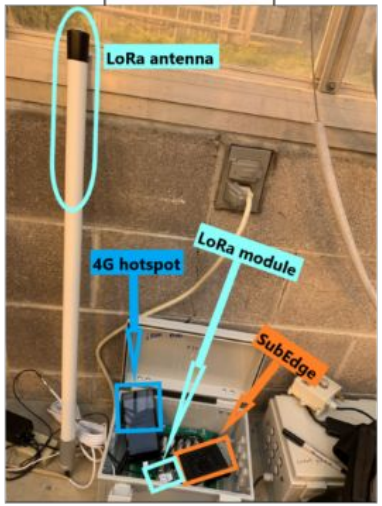

Sensor Hubs

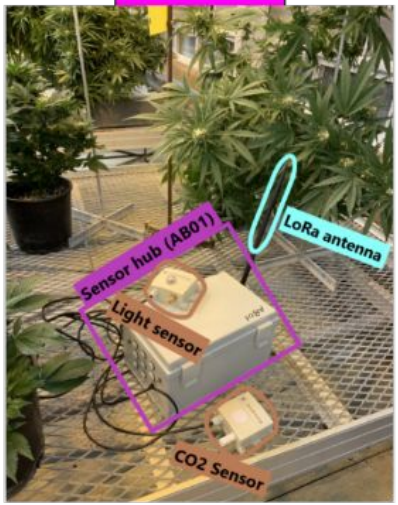

TVWS Client

TVWS Base Station

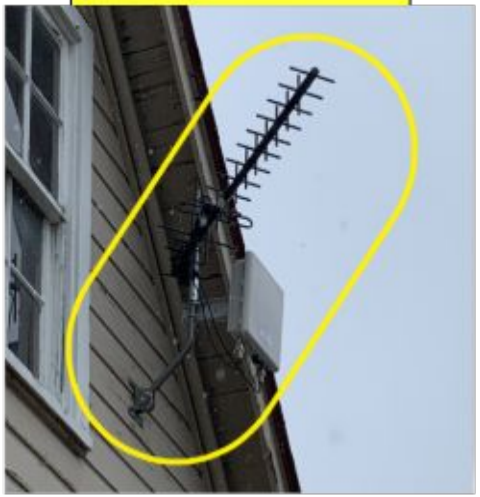

9 sensors hubs

1 **TVWS** client/base station

2 **LoRa** antennas

2 Edge devices (4GB RAM, 32GB storage)

FarmBIOS Deployments: WaterGuard

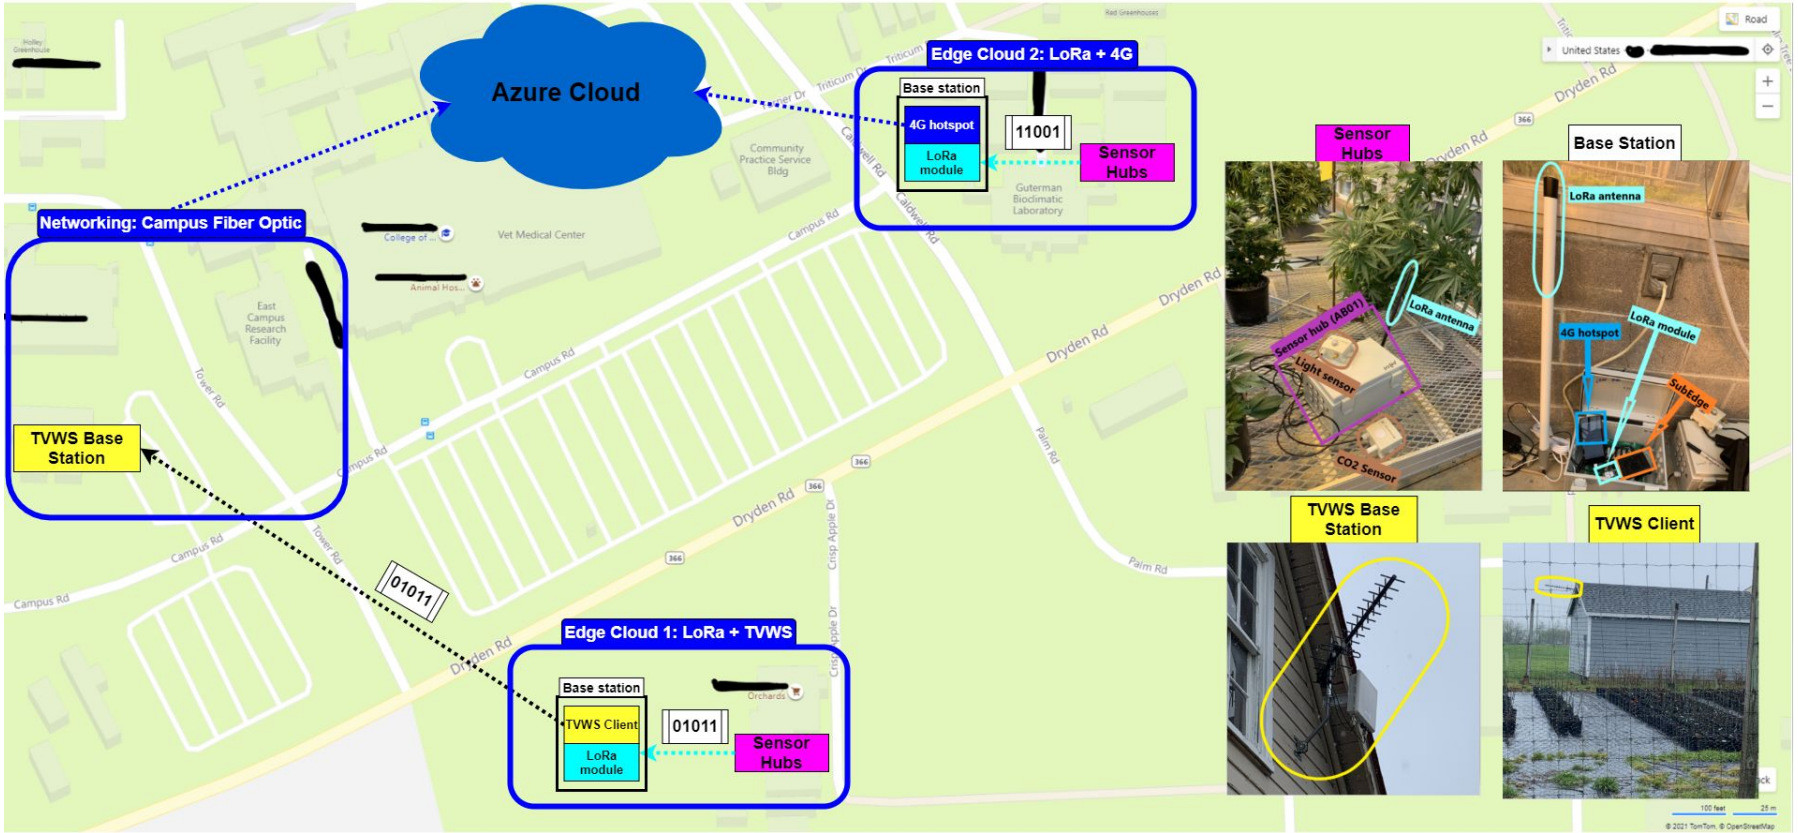

FarmBeats: An IoT Platform for Data-Driven Agriculture. Vasisht *et al*. 2017

Adapting to the wild: offline data collection

Total deployment: **18 months**
Average deployment: **223 days**
1M+ sensor readings

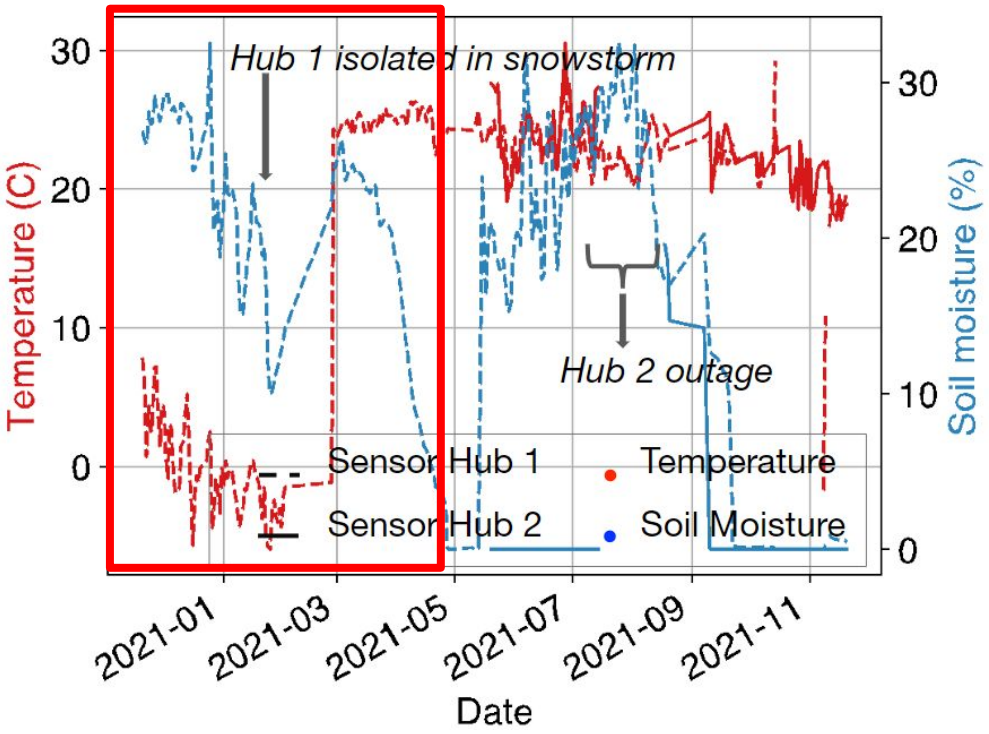

Adapting to the wild: faulty sensors and human errors

Faulty sensors and **config errors** affect data analytics

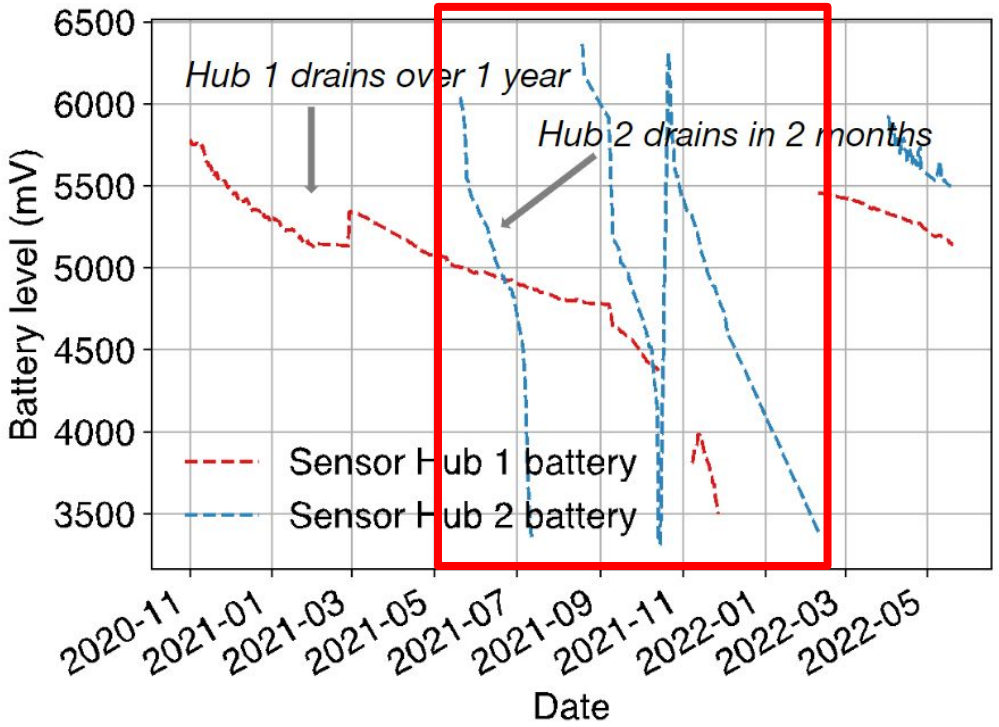

Adapting to the wild: active digital twin

Active digital twin notifies human operators if a sensor hub twin diverges from its physical twin

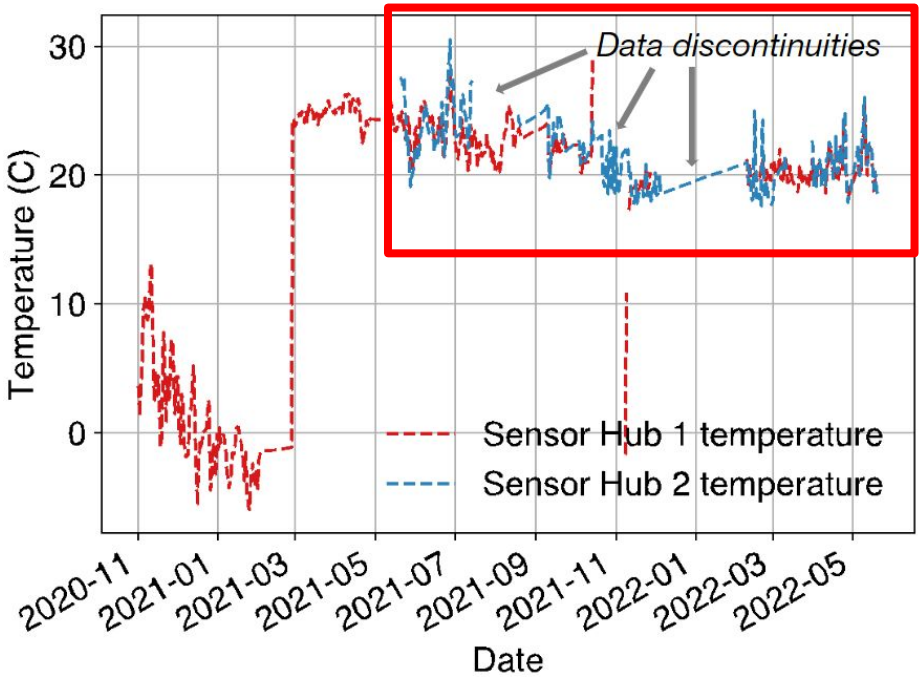

Adapting to the wild: active digital twin

Active digital twin notifies human operators if a sensor hub twin diverges from its physical twin

Seamless Visions, Seamful Realities: Anticipating Rural Infrastructure Fragility in Early Design of Digital Agriculture. Rubambiza *et al*. 2022

Background

DA Challenges & State of the art

Comosum Design & Implementation

Deployment Experiences, **Insights & Limitations**

Conclusion

Practical Insights & Limitations

- **Extending to new vendors comes with minor costs**
  - **New scripts for each vendor (50 LoCs)**

Practical Insights & Limitations

- Extending to new vendors comes with minor costs
  - New scripts for each vendor (50 LoCs)
- **The cloud complicates reconfigurability**
  - **Treat incoming parameters as abstract data types (ADTs)**

Practical Insights & Limitations

- Extending to new vendors comes with minor costs
  - New scripts for each vendor (50 LoCs)
- The cloud complicates reconfigurability
  - Treat incoming parameters as abstract data types (ADTs)
- **Failure in DA is the norm, not the exception**
  - **Active digital twins detect divergence**

Practical Insights & Limitations

- Extending to new vendors comes with minor costs
  - New scripts for each vendor (50 LoCs)
- The cloud complicates reconfigurability
  - Treat incoming parameters as abstract data types (ADTs)
- Failure in DA is the norm, not the exception
  - Active digital twins detect divergence
- **No automated migrations during permanent outages**
  - **Future work**

Background

DA Challenges & State of the art

Comosum Design & Implementation

Deployment Experiences, Insights & Limitations

Conclusion

Comosum aka The Software Defined Farm (SDF)

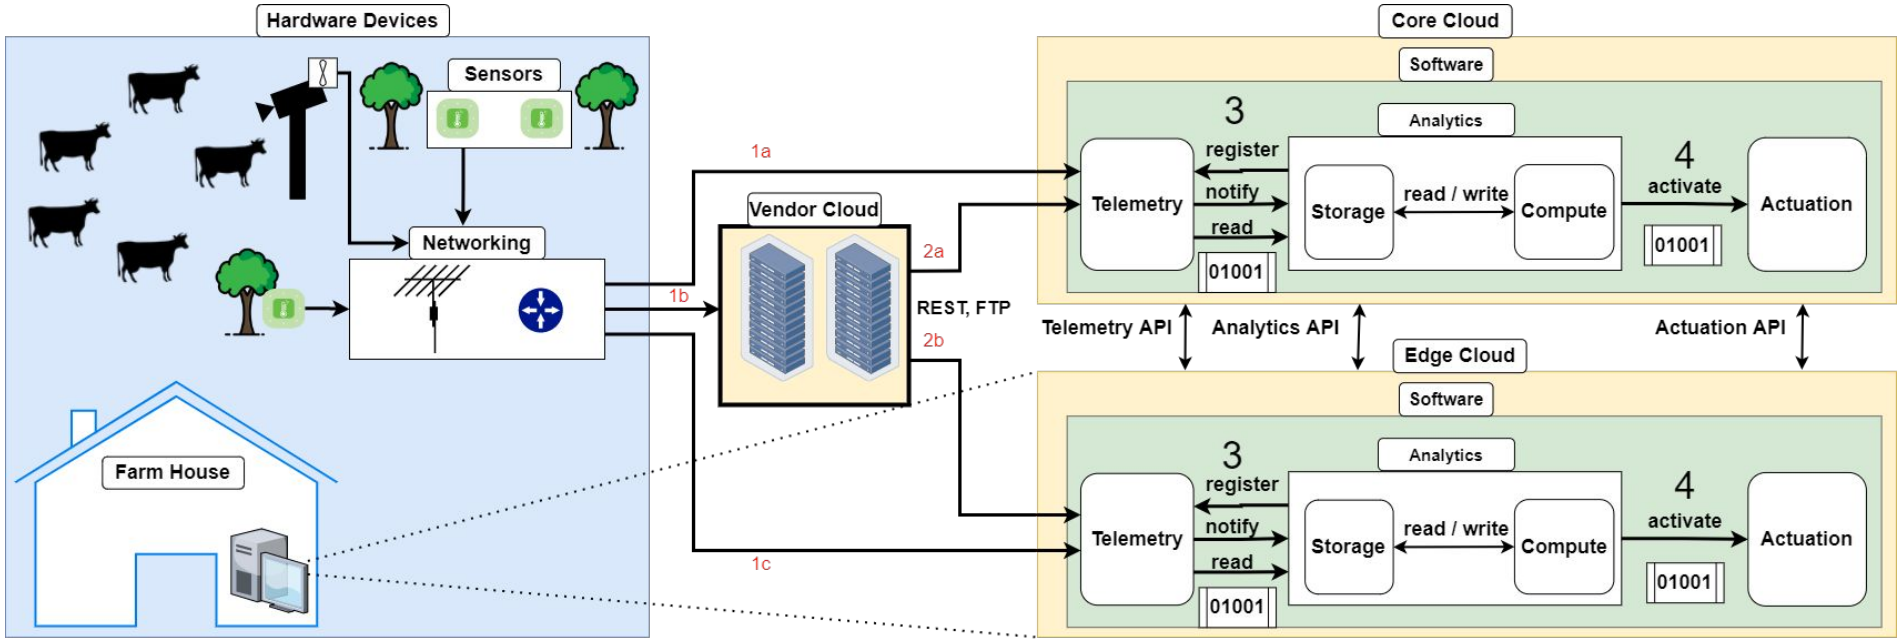

Comosum Technical Contributions

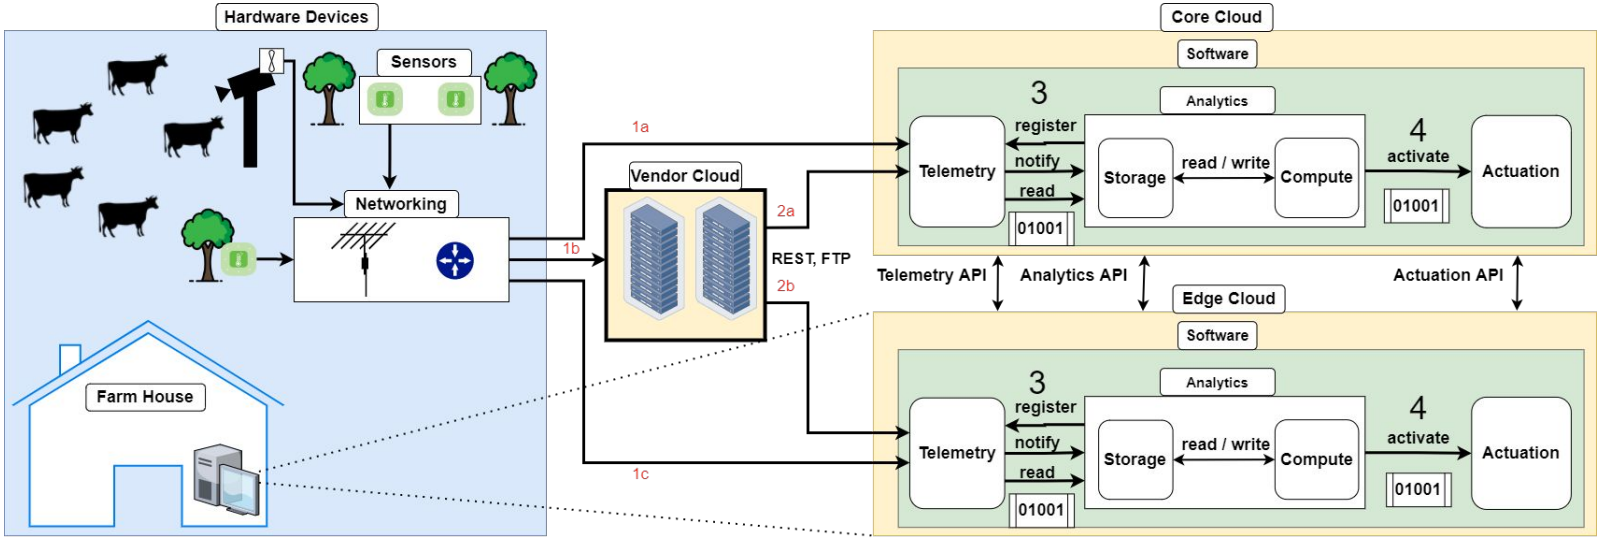

- Distilled black-box SOTA down to a single interface
- Applied strong systems approaches to new contexts
- Deployed across different farm types and cloud providers

Comosum Experiences, Insights, Limitations

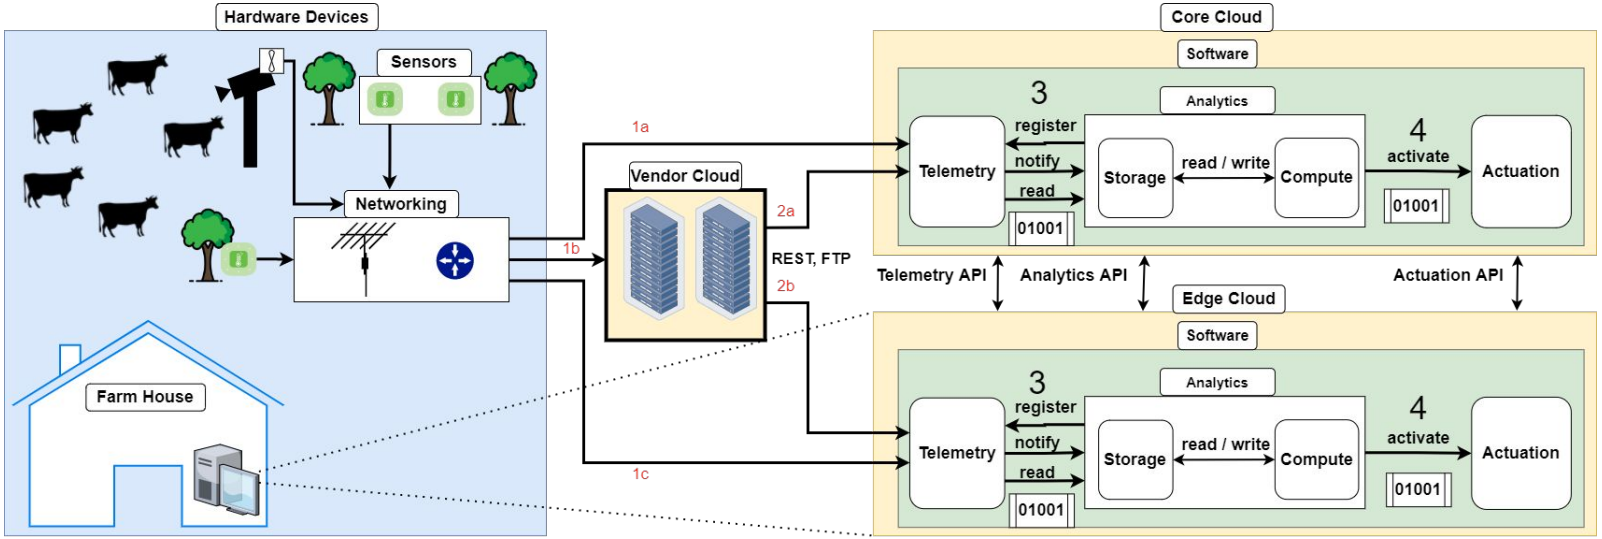

Total deployment: **18 months**

Average deployment: **223 days**

1M+ sensor readings

Comosum Experiences, Insights, Limitations

https://github.com/Cornell-CIDA-Dev/Software-Defined-Farm

Total deployment: **18 months**

Average deployment: **223 days**

1M+ sensor readings

Thanks


Microsoft


USDA


NSF


cida
Cornell Institute
for Digital Agriculture

Thank You

Questions?

Gloire Rubambiza

Email: gloire@cs.cornell.edu

USENIX Slack: @gloire

Web: https://rubambiza.github.io

Twitter: @GloireKnowsBest