SmartFarm: IoT Systems That Simplify & Automate Agriculture Analytics

Chandra Krintz
Dept. of Computer Science
UC Santa Barbara

UCSB

International Conference on the Internet of Things (IoT), Oct. 17 2018

Works out on a Peloton
Owns an RV
Likes Movie Usual Suspects
Has 4 grandchildren
Dislikes Chinese food
Has 3 kids
Likely Republican
Age 52-60
Reads crime dramas
Likes Jimmy Fallon
Favorite color: Blue
Male
Household Income: 150000
Drives a Dodge Ram truck
Mother lives in Florida
Is active on Twitter
Likes basketball
Is politically active
Works out at a gym
Likes spicy food recipes
Children play hockey
Sister is a lawyer
Likes hiking
House value: $500,000
Dislikes Actor Robert Redford
Owns a SmartTV
Salesman
Owns an RV
Likes online news sites
2013 Salesman of Year
Dislikes snakes
Hates doing dishes
Major life insurance holder

Cloud + Data Analytics: Revolutionizing Commerce

Internet Activity

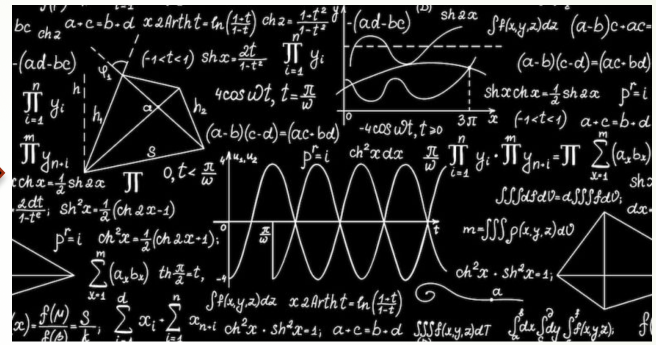

Analytics: Inference and Prediction

UCSB

Cloud + Data Analytics: Revolutionizing Commerce

Internet Activity

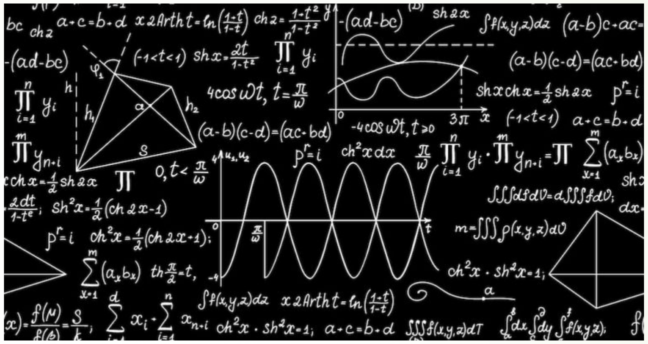

Analytics: Inference and Prediction

UCSB

Cloud + Data Analytics: Revolutionizing Commerce

Internet Activity

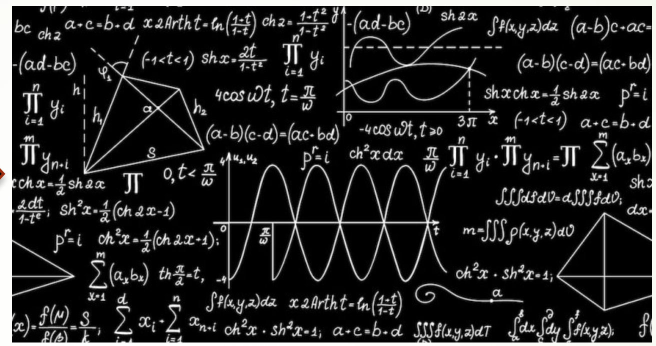

Analytics: Inference and Prediction

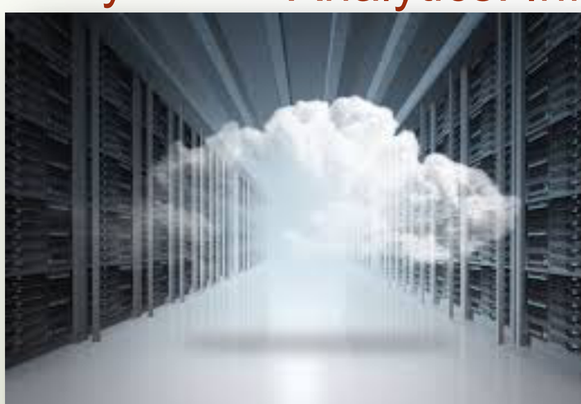

Cloud + Data Analytics: Revolutionizing Commerce

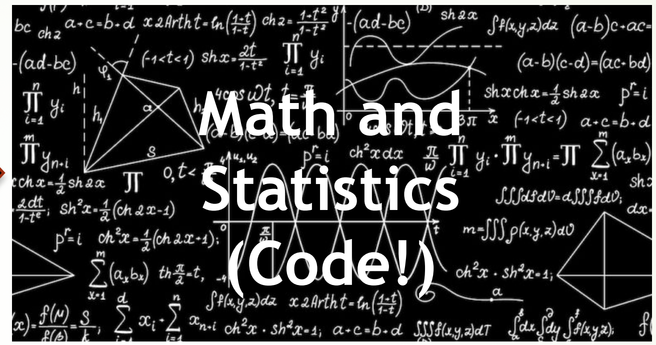

- *What will you buy?*
- *When will you buy it?*
- *What will you pay?*

Internet Activity

Analytics: Inference and Prediction

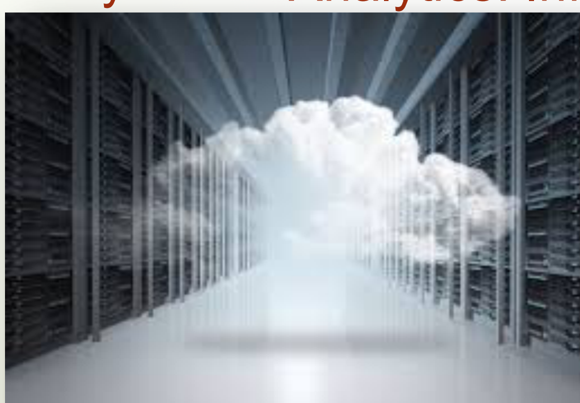

What Else Can We Revolutionize With It?

- To solve a *very hard*, *impending* problem: **feeding the planet**
  - A complex system: Food-Energy-Water nexus
  - Global: *821M people today are undernourished
    - *6.5M children live in food-insecure households in US
    - *9.8B people to feed by 2050

- Just how complex is it?

* http://www.cdfa.ca.gov/statistics/
https://www.ers.usda.gov/topics/food-nutrition-assistance/food-security-in-the-us/key-statistics-graphics.aspx
http://www.fao.org/state-of-food-security-nutrition/en/ FAO 2018
https://www.un.org/development/desa/en/news/population/world-population-prospects-2017.html

UCSB

Critical Needs & Complex Challenges Surrounding Food Production

The world needs more **food** for a growing population.
We use 70% of fresh **water** for growing food.

Worker shortages and high **labor costs**

Invasive pests and disease threaten production

We lose ~40% of the food we produce to **spoilage**

30% of **global energy** is used to produce food
22% of **greenhouse gases** come from agriculture

Sources: USDA, UN FAO

UCSB

What Else Can We Revolutionize With Cloud+Analytics?

- To solve a *very hard*, *impending* problem: feeding the planet
  - A complex system: Food-Energy-Water nexus

- **The Cloud+Data Analytics boom has not come to Ag, despite**
  - Maturity of big data and machine learning technologies available
  - Vast amounts of data surrounding food production, processing & the crop lifecycle
    - Weather, historical records, sensors, images (NDVI/thermal), ...
    - At the farm, community, county, region, state, nation, global level

UCSB

IoT to Bridge the Gap

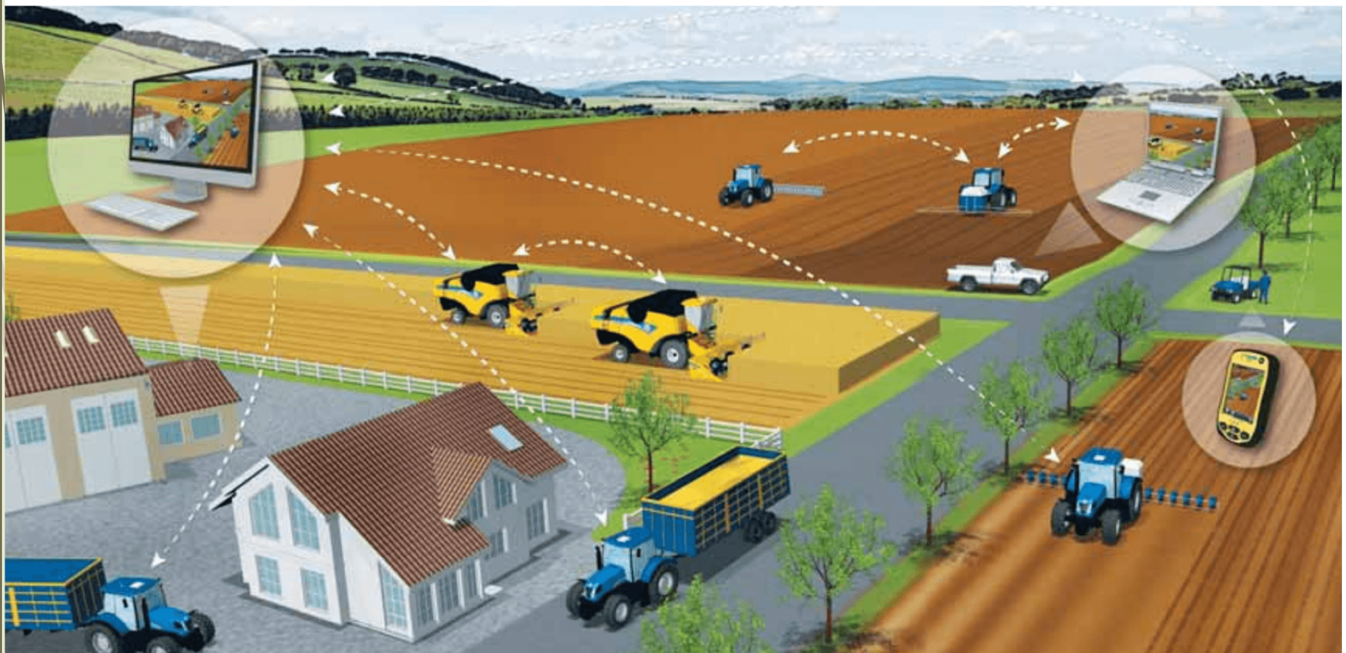

http://www.agrimachinerynews.com/wp-content/uploads/2014/10/Agri-Machinery-News-New-Holland-Precision-Farming.jpg

IoT to Bridge the Gap

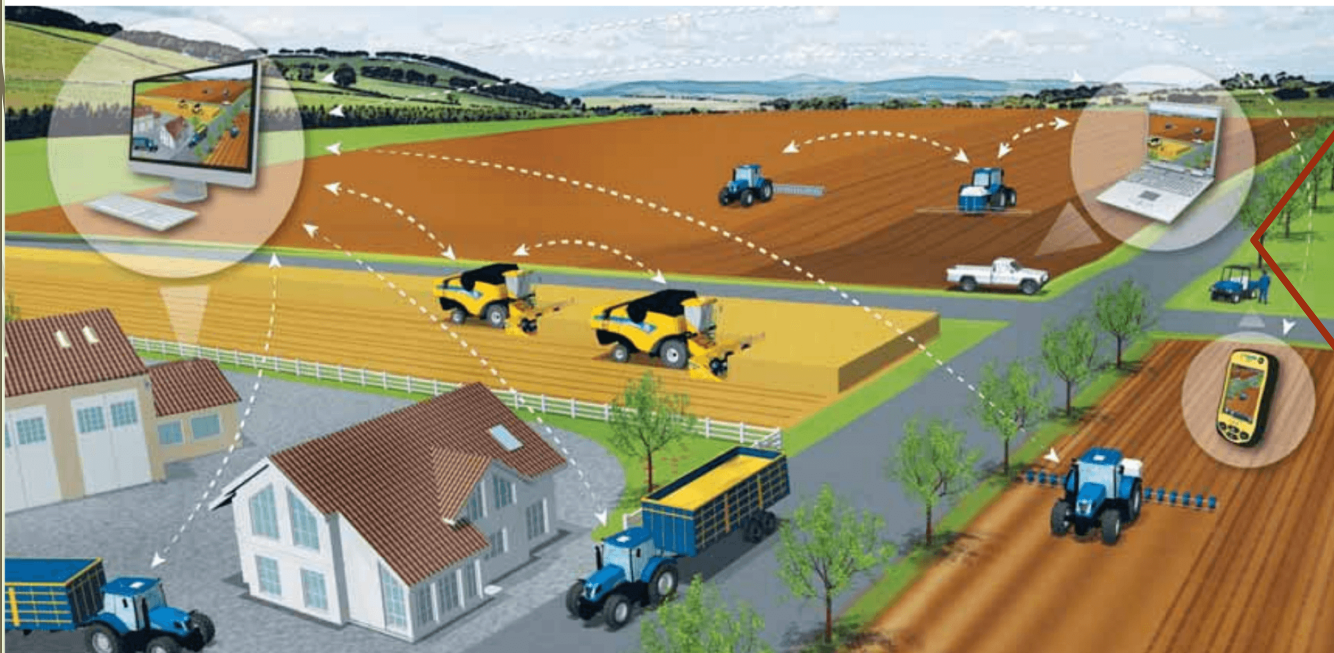

http://www.agrimachinerynews.com/wp-content/uploads/2014/10/Agri-Machinery-News-New-Holland-Precision-Farming.jpg

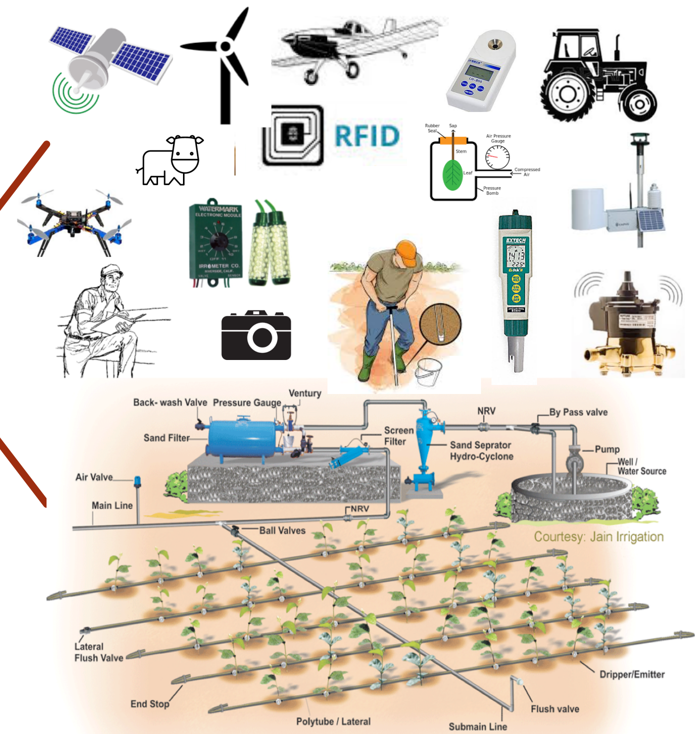

How Do We Get There?

How Do We Get There? aka What Would Amazon Do?

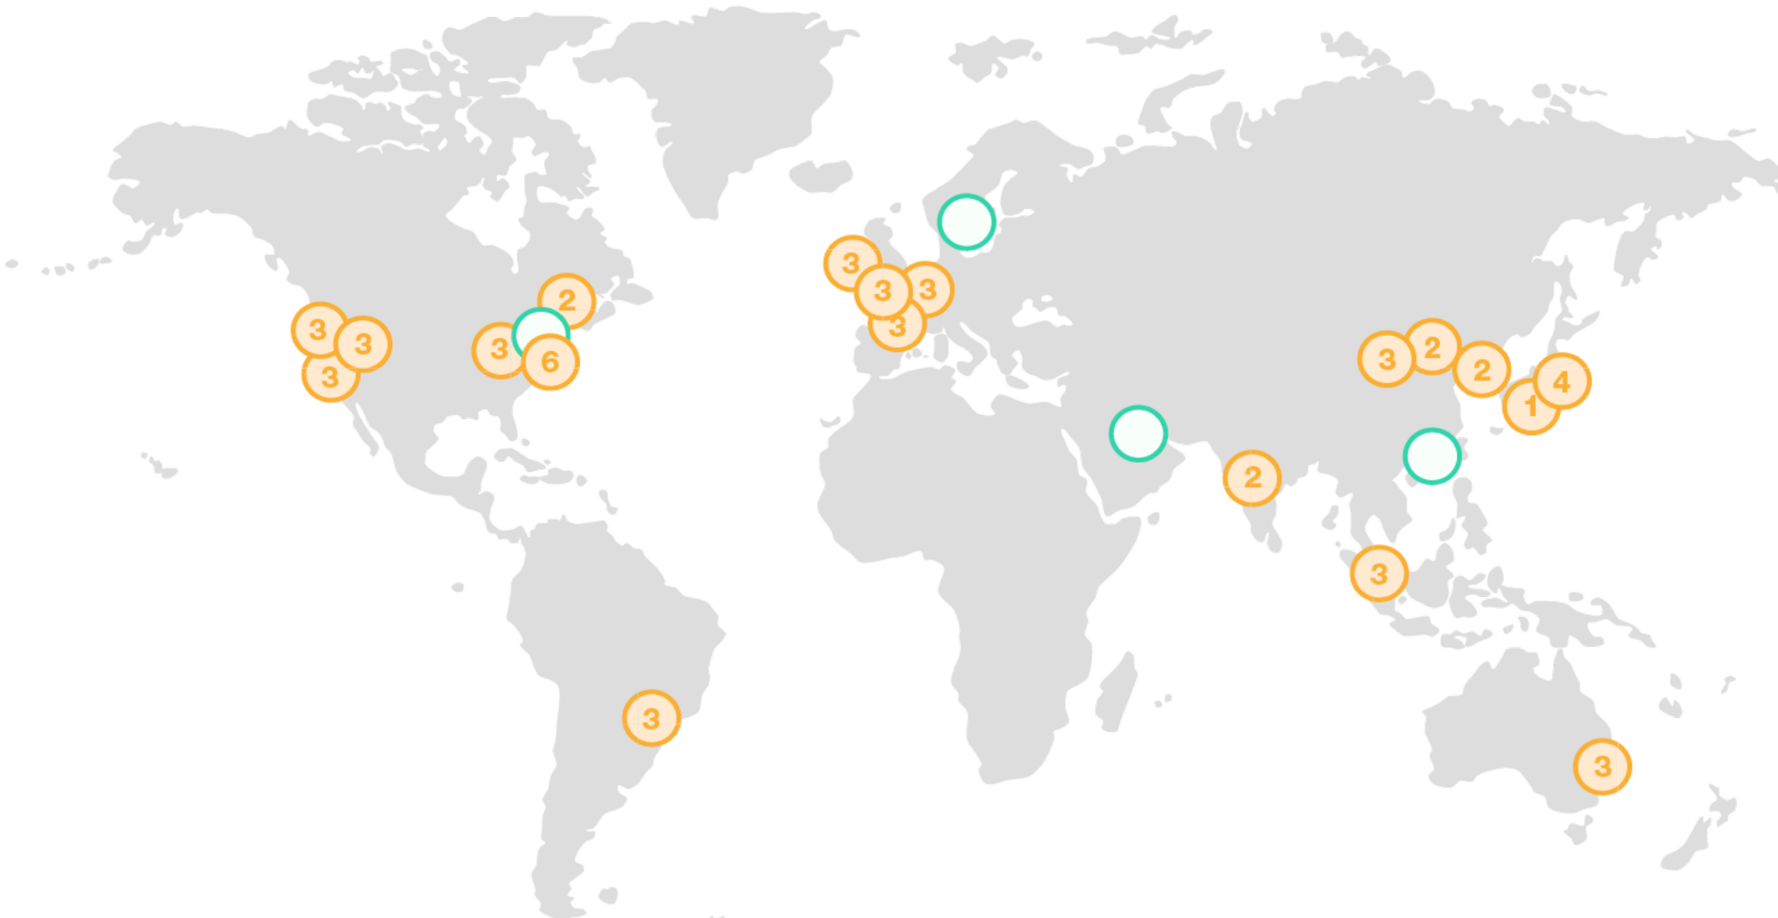

Amazon Cloud Infrastructure

How Do We Get There? aka What Would Amazon Do?

But E-Commerce and Entertainment are Download Dominated!

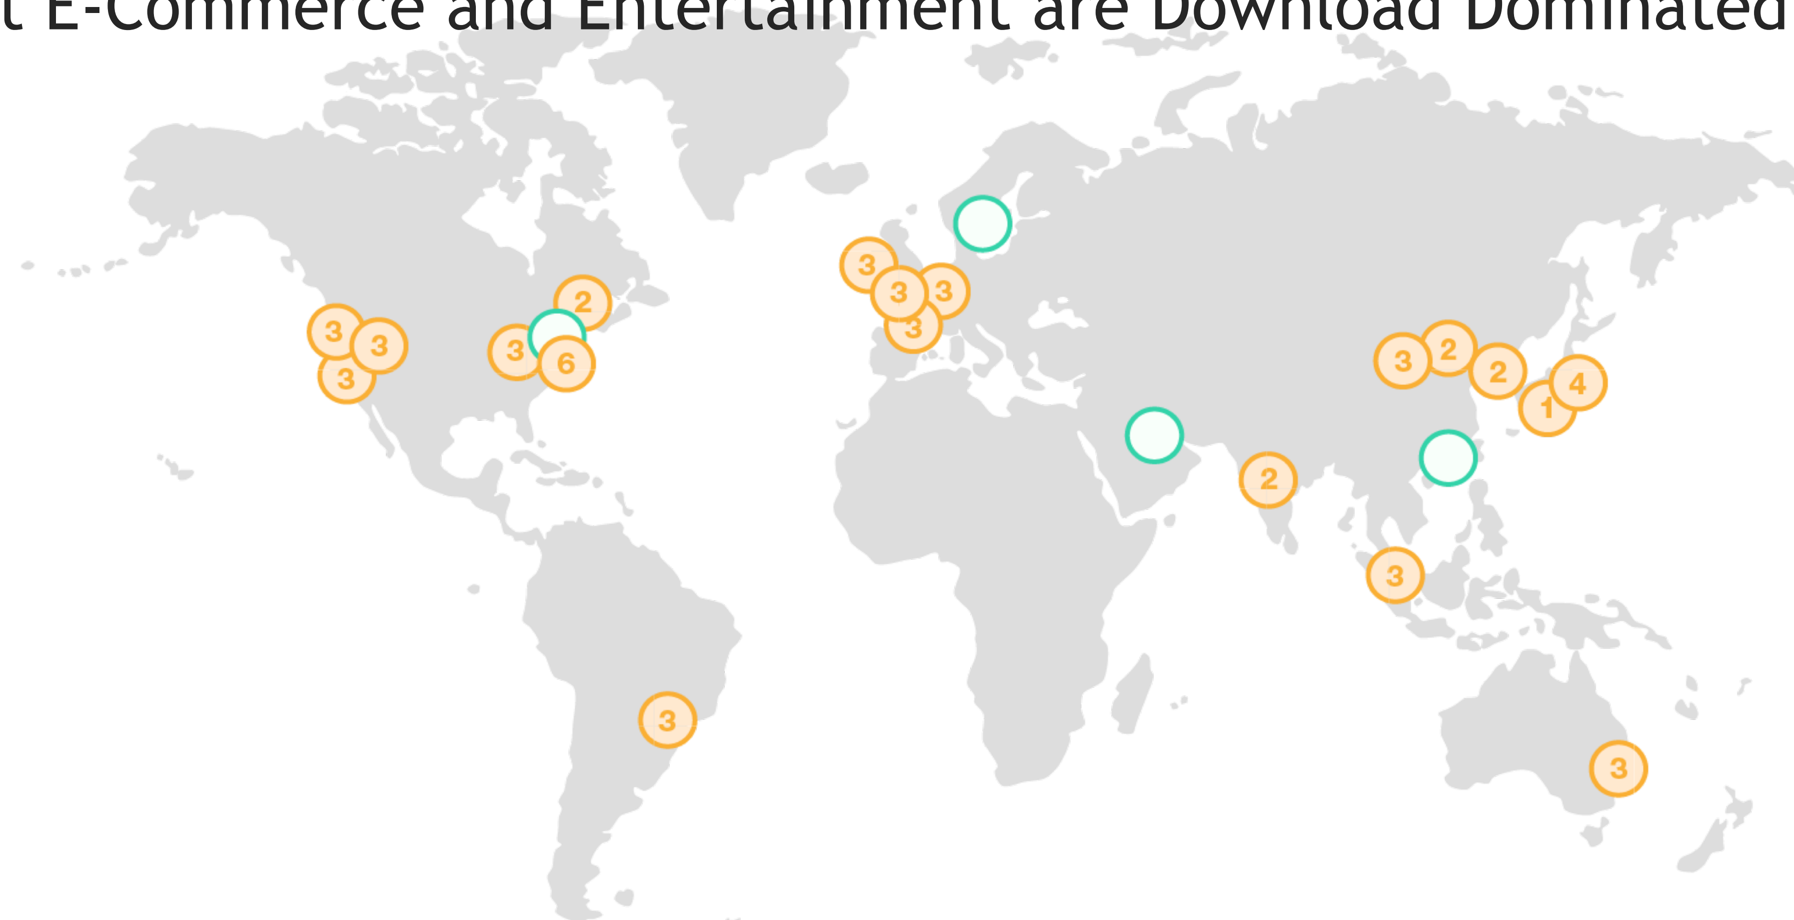

Amazon Cloud Infrastructure

The "Cloud" is More Distributed Than You Think!

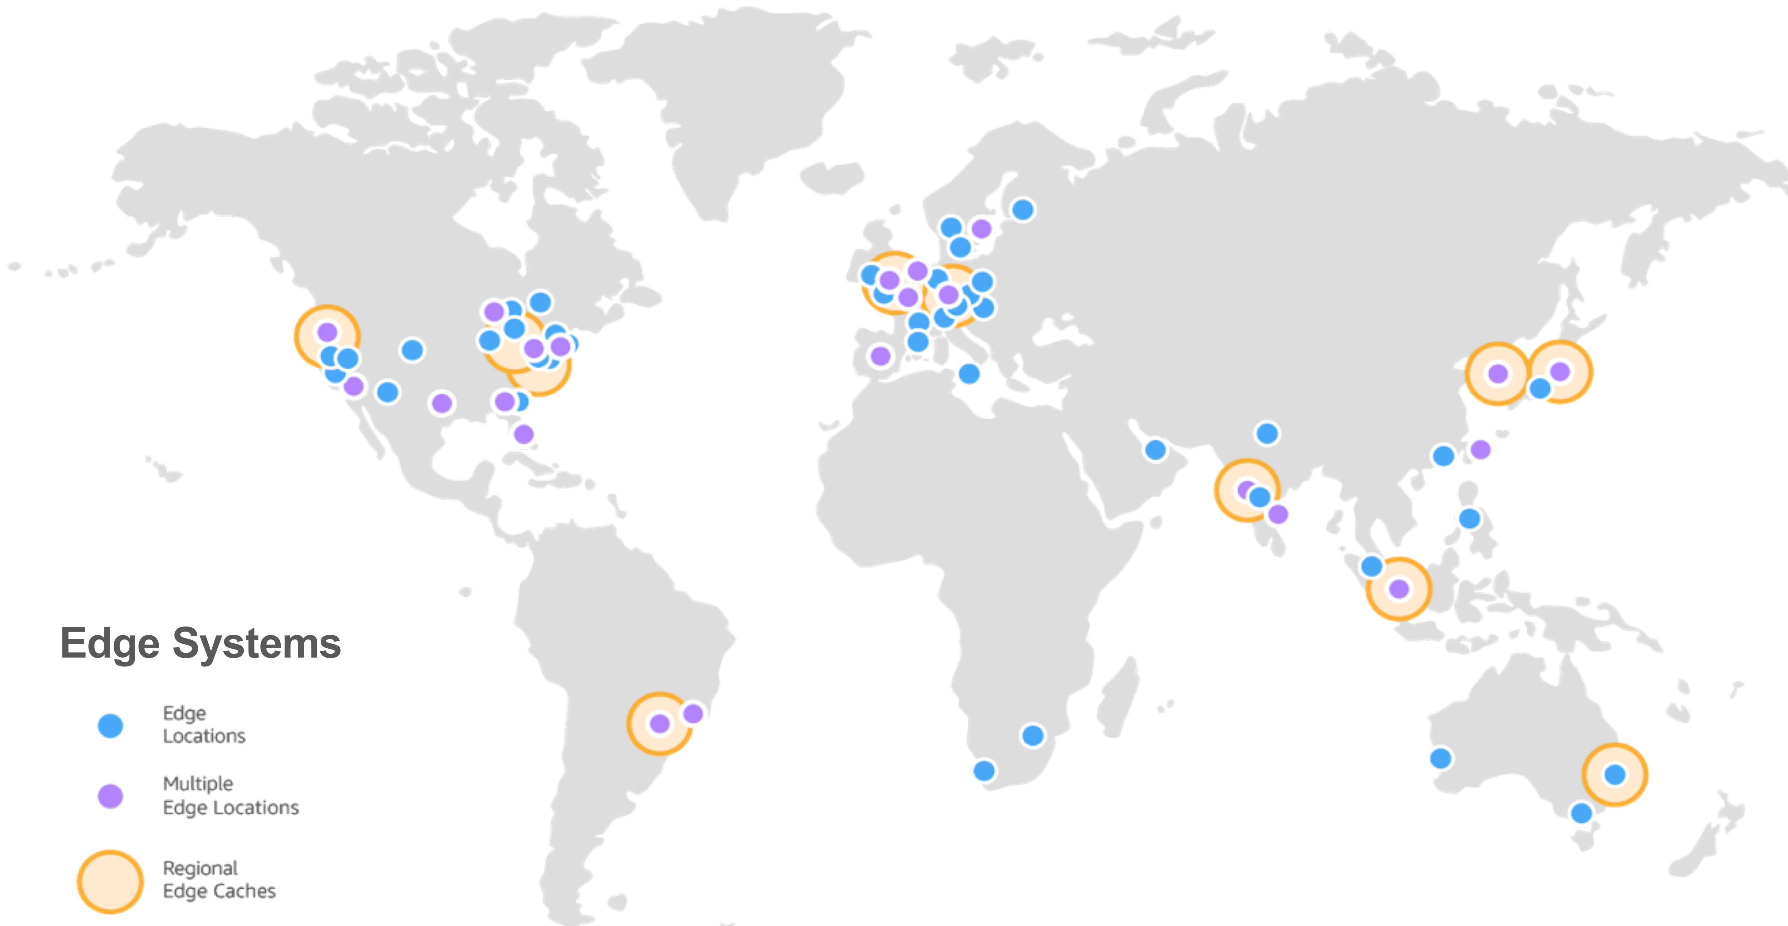

Amazon Content Delivery Infrastructure

Content (Download) Delivery -- is the (Edge) Cloud

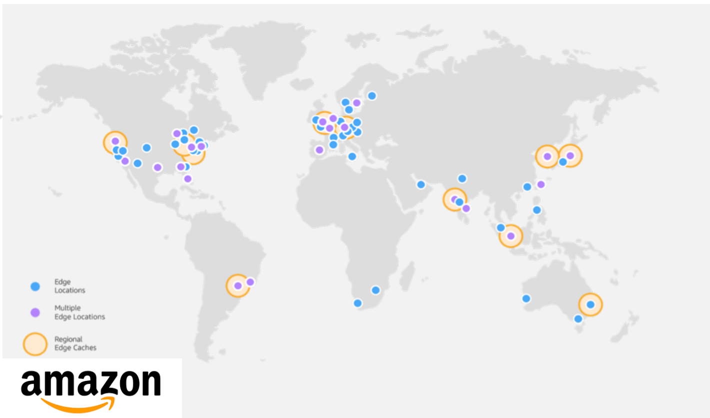

AWS 2018

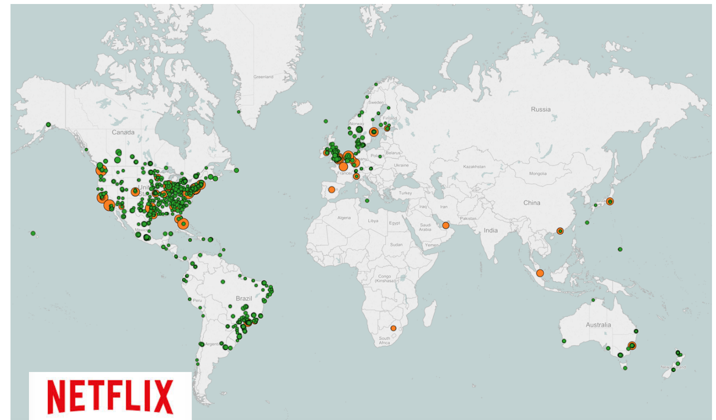

Netflix 2018

Adhikari et al, HotMD'11

IoT is Write/Upload Dominated

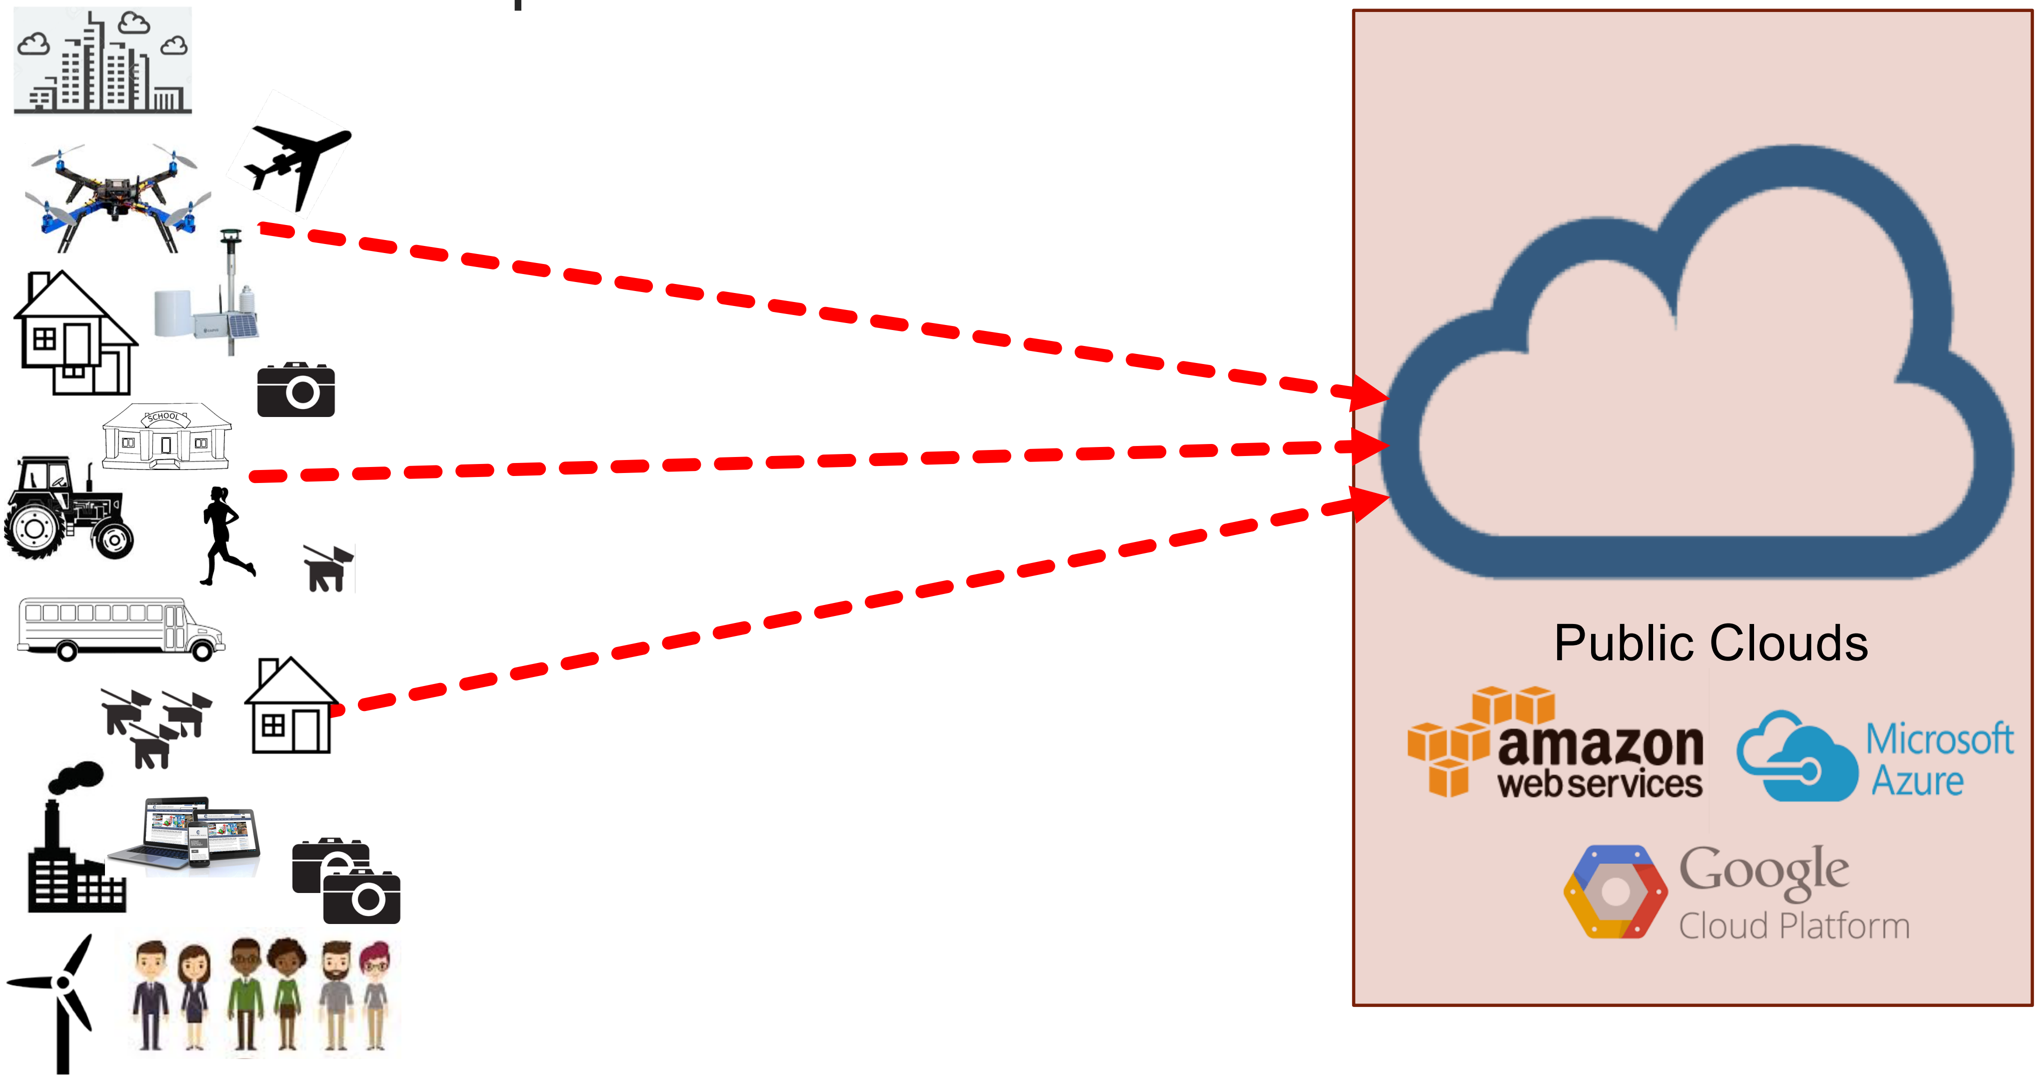

IoT is Write/Upload Dominated

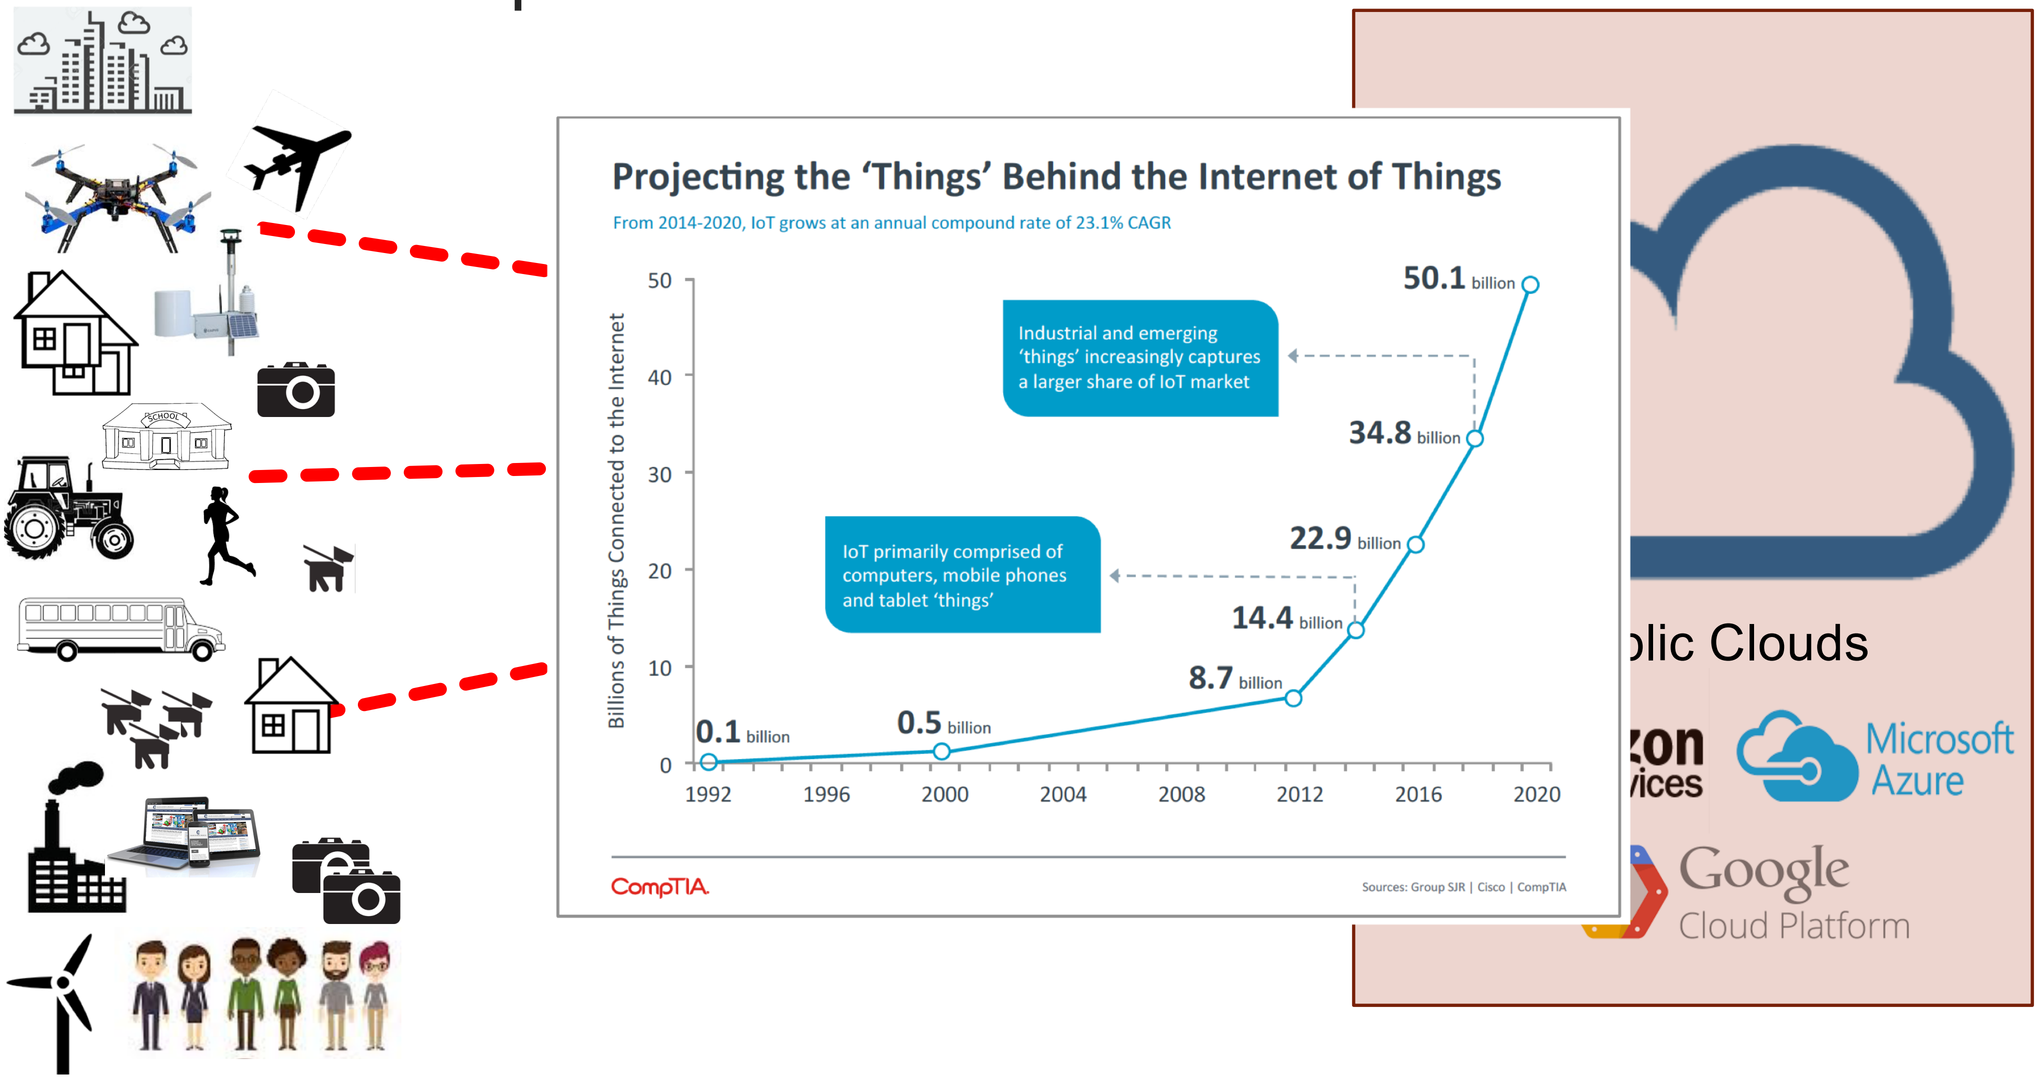

IoT is Write/Upload Dominated

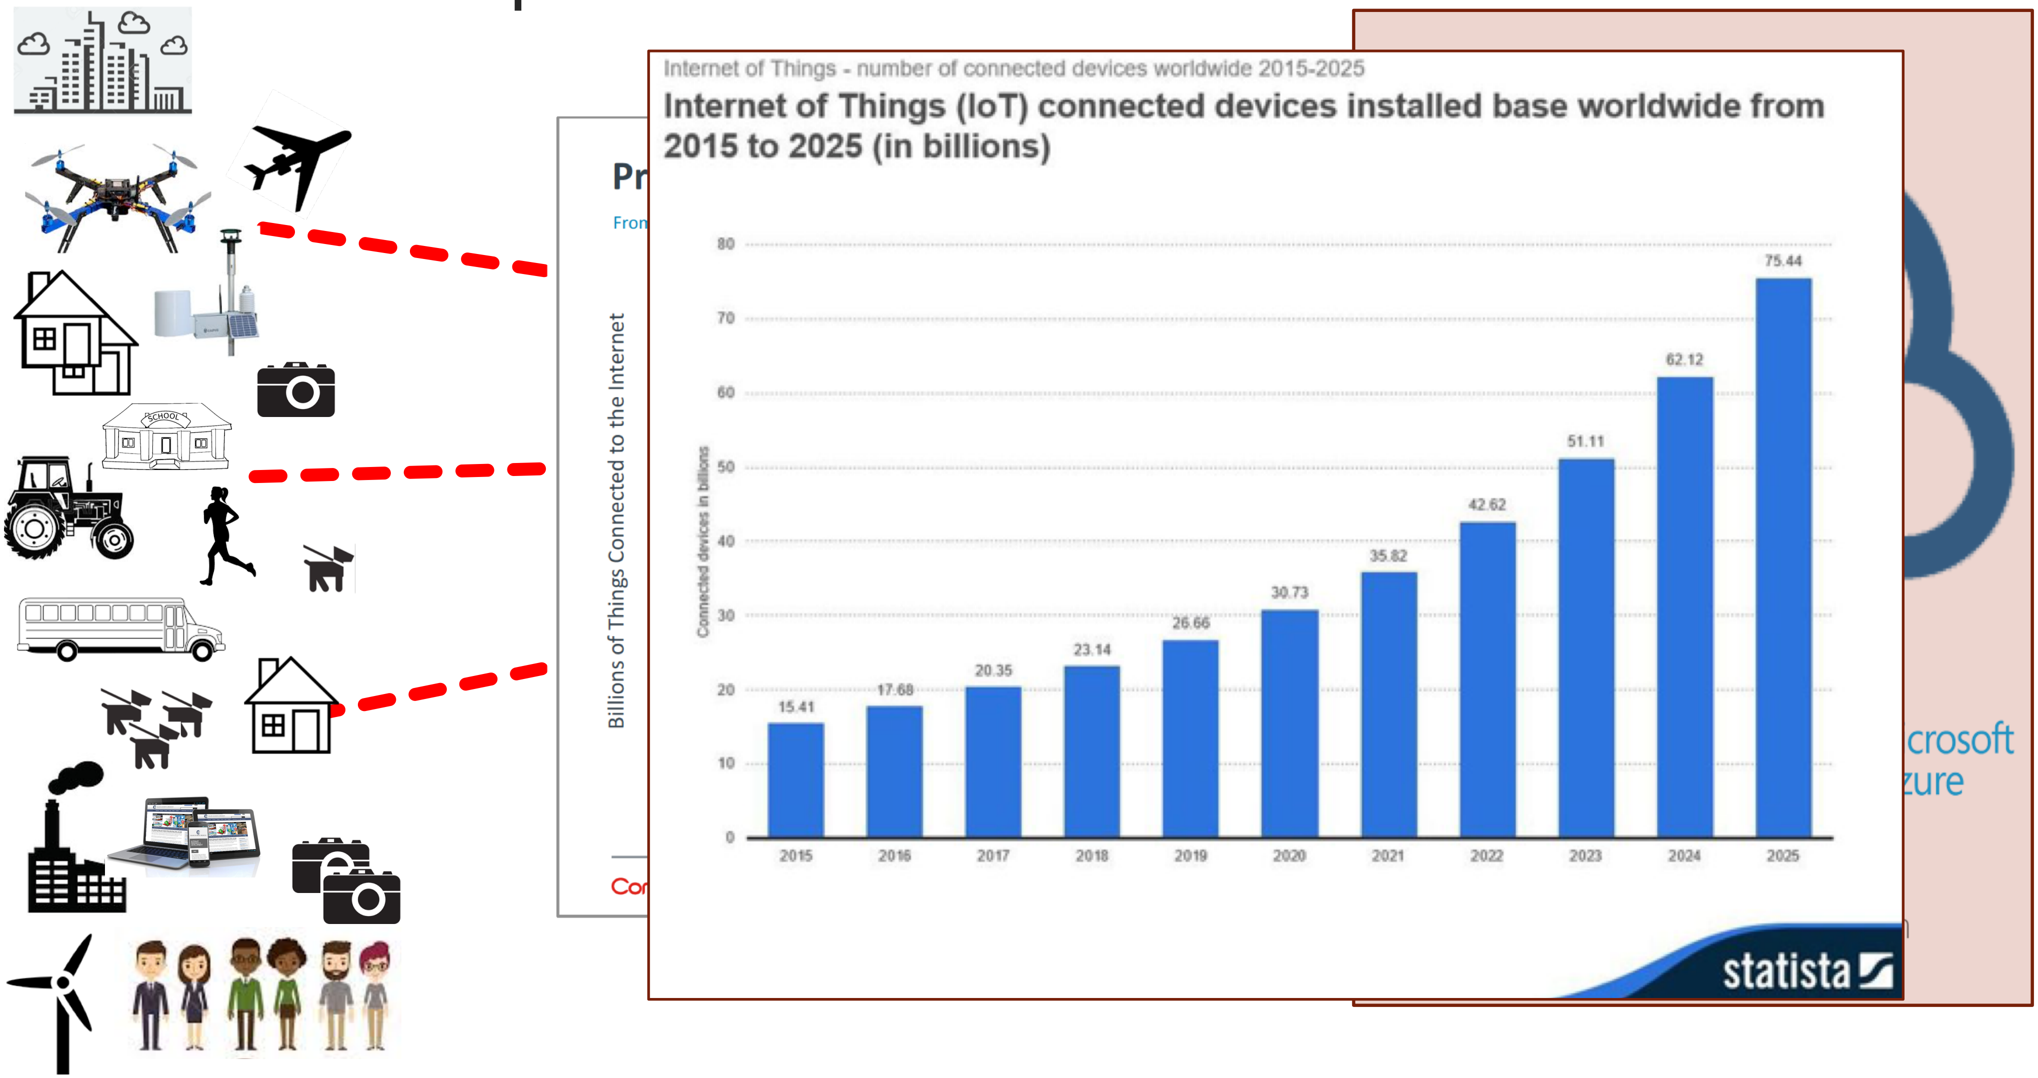

Key Differences Between Farm & E-Commerce

- Many more "things" (1T) than people (3.5B) connected
- Write not *read* dominated
- Decision-making, actuation, & control is **local**
- Farmers must maintain ownership and control over their data
- Most farms are **not** well connected

UCSB

Farm vs E-Commerce/Entertainment

- Many more "things" (1T) than people (3.5B) connected
- Write not *read* dominated
- Decision-making, actuation, & control is **local**
- Farmers must maintain ownership and control over their data
- Most farms are **not** well connected
- **Similarity:** Latency matters

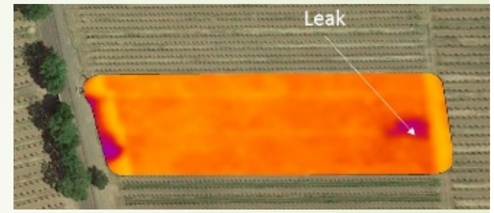

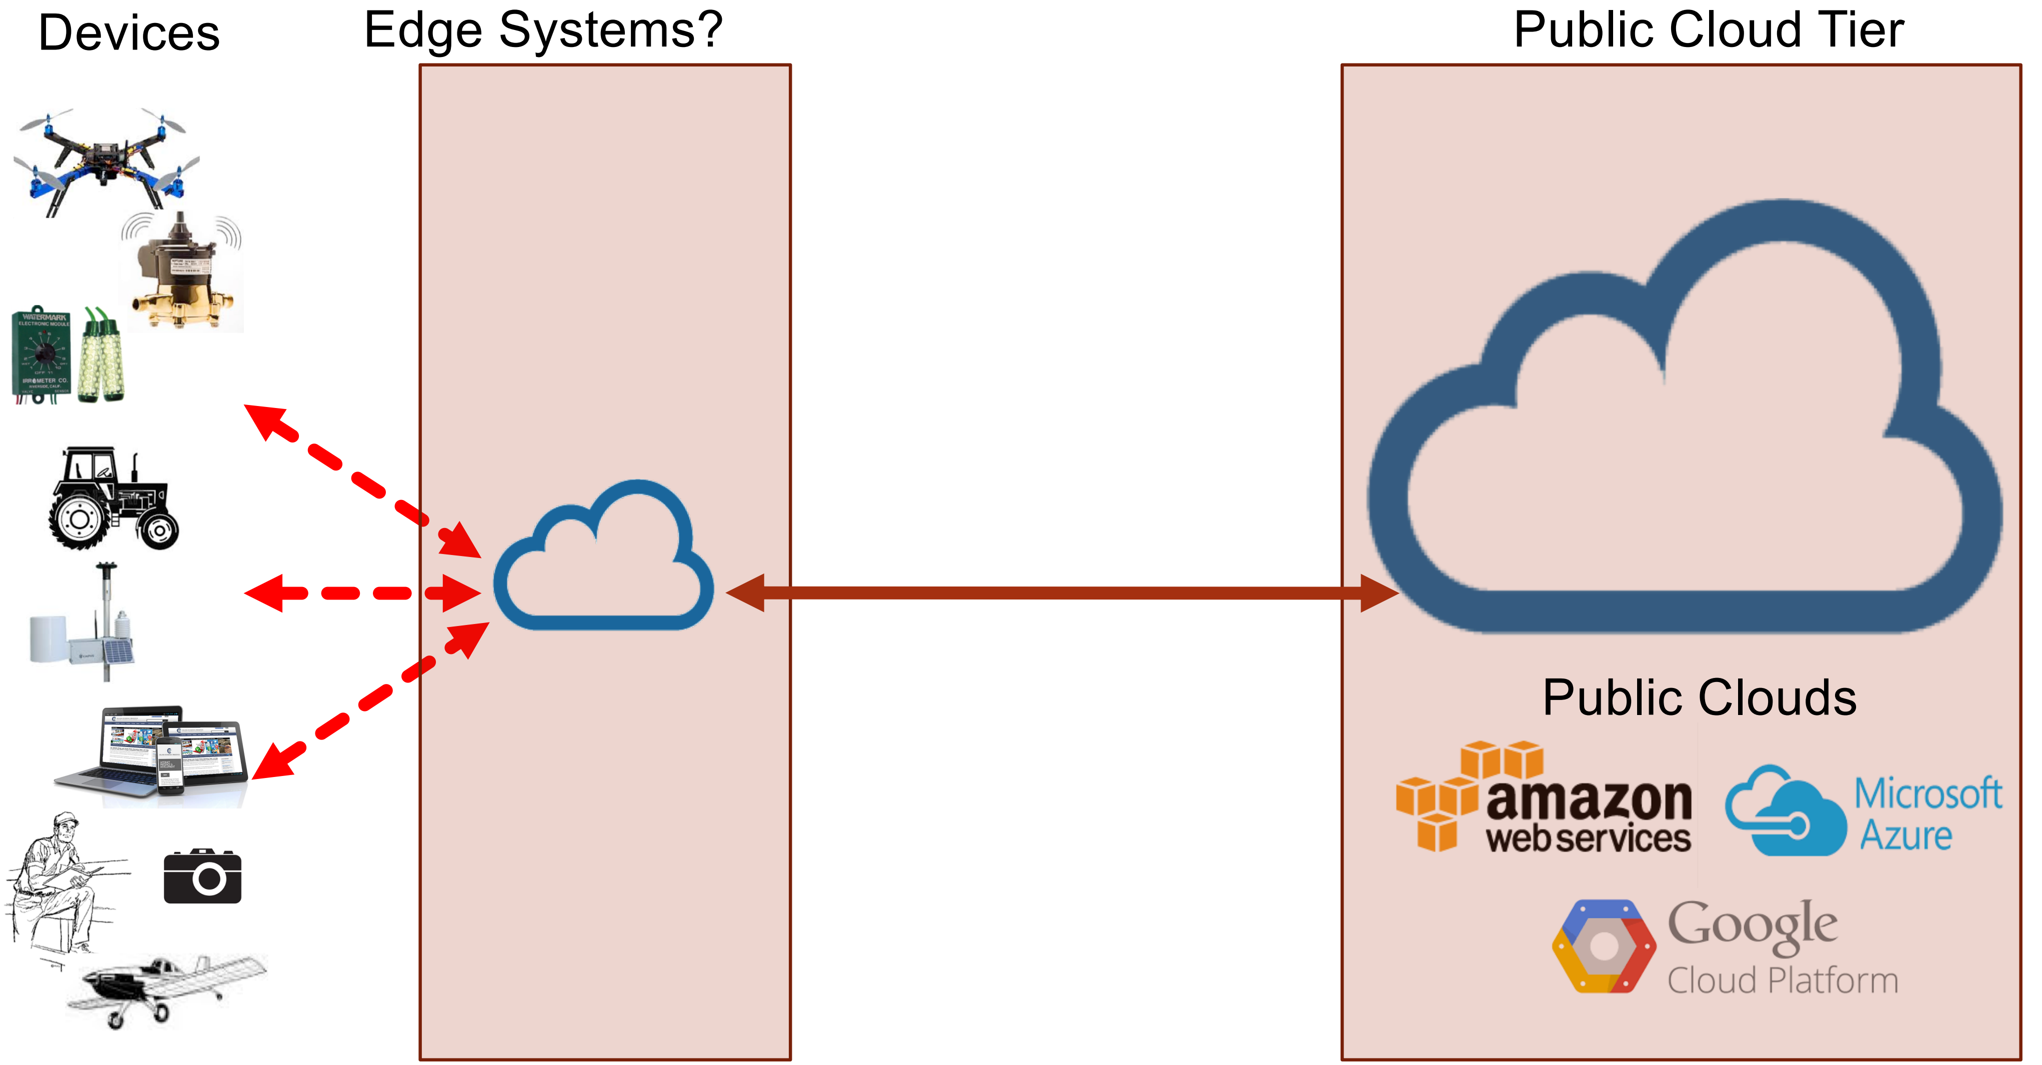
Devices
Edge Systems?
Public Cloud Tier
Public Clouds
amazon
web services
Microsoft
Azure
Google
Cloud Platform

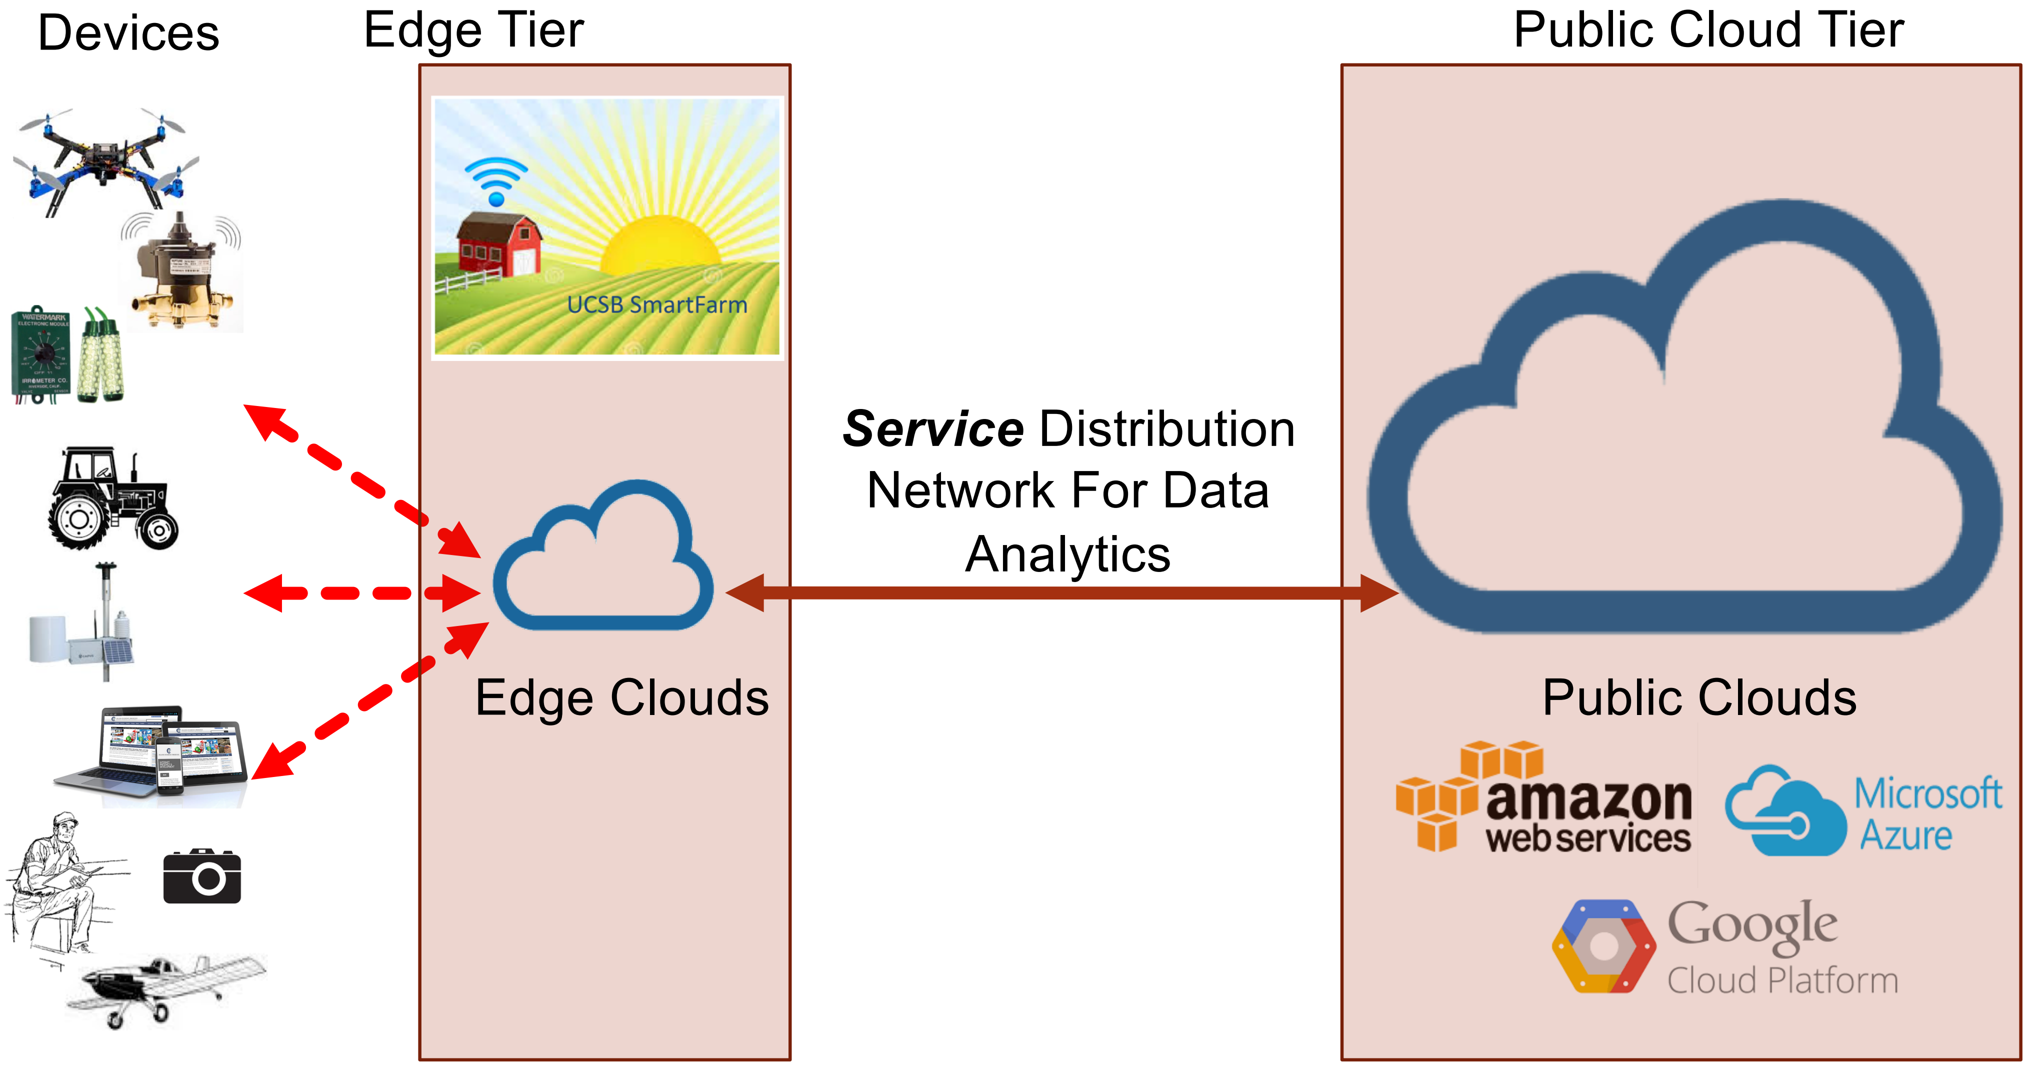
Devices
Edge Tier
Public Cloud Tier
UCSB SmartFarm
Service Distribution Network For Data Analytics
Edge Clouds
Public Clouds
amazon web services
Microsoft Azure
Google Cloud Platform

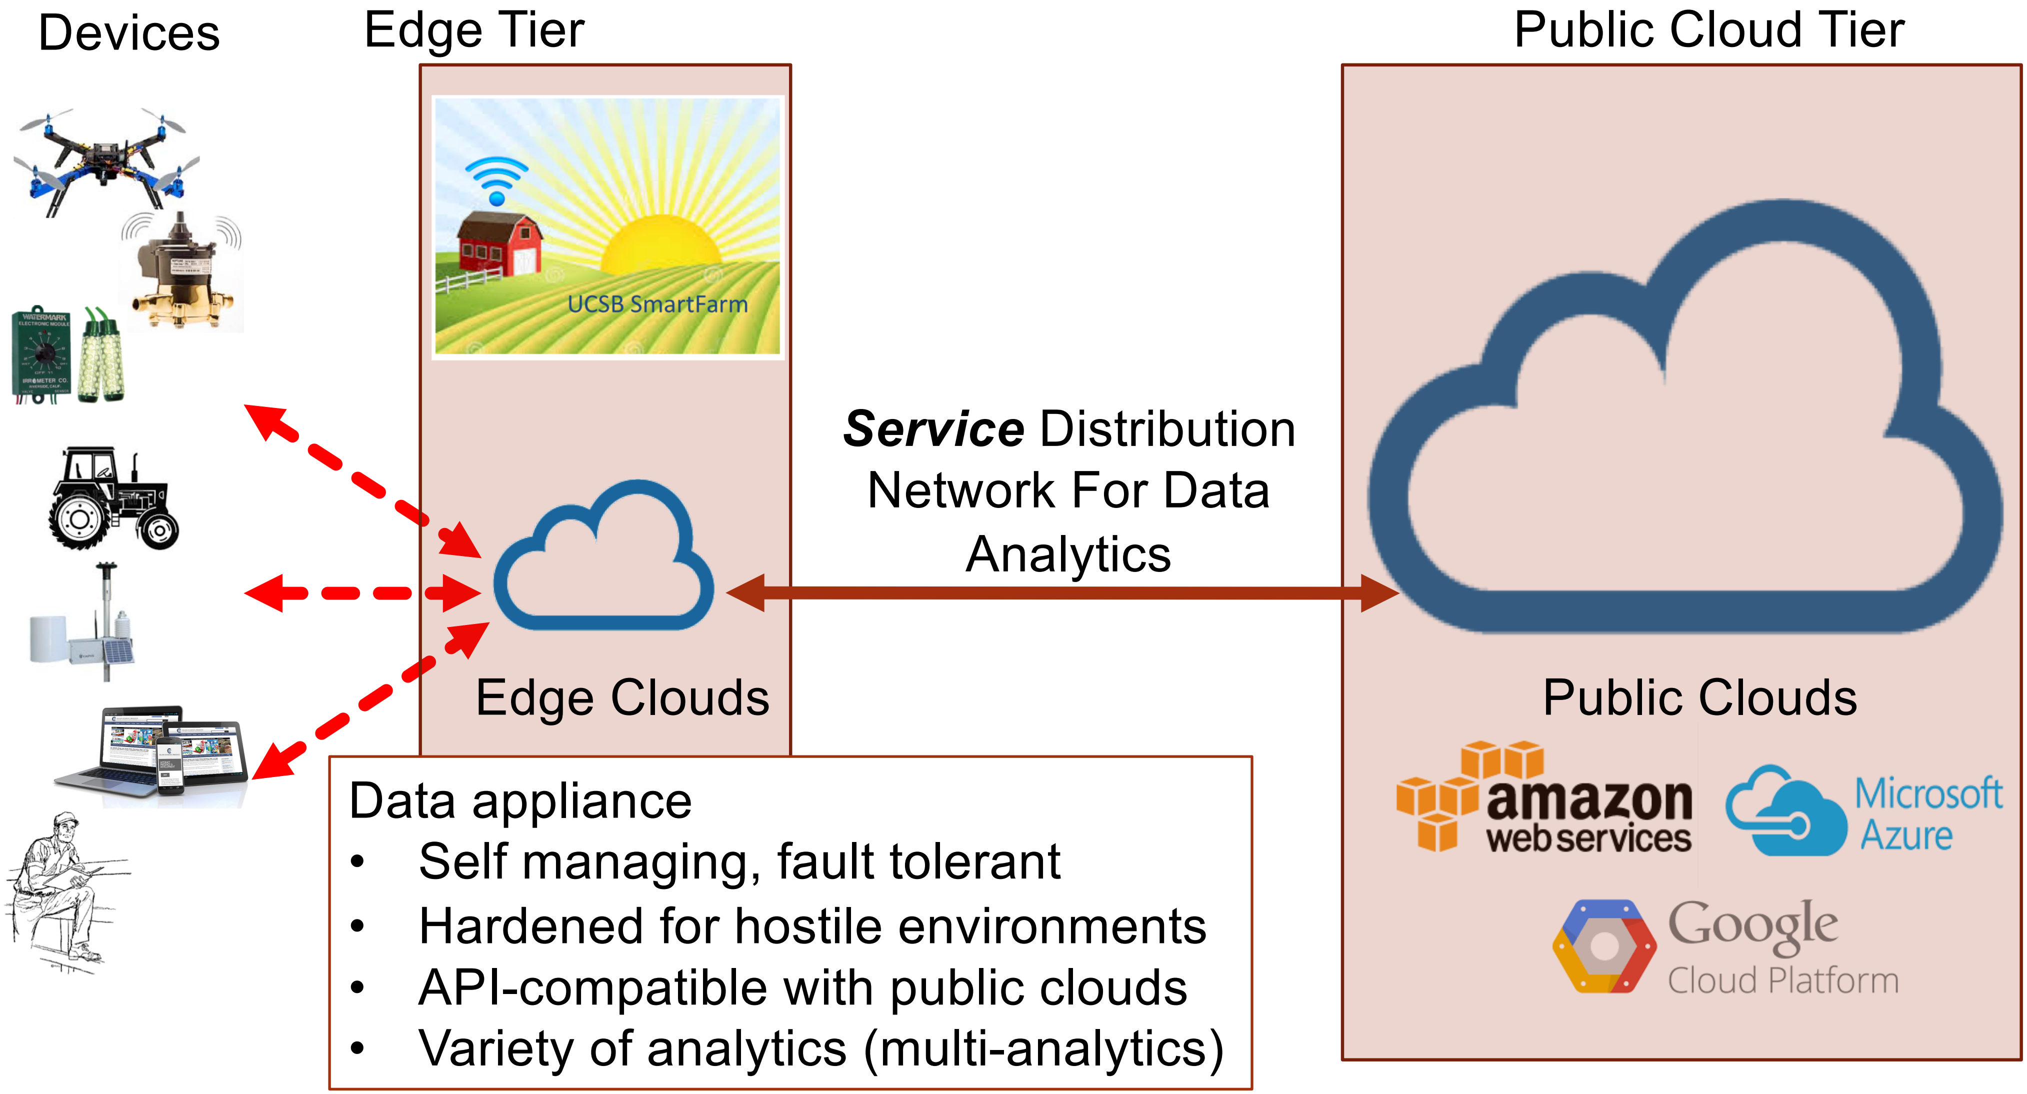
Devices
Edge Tier
Public Cloud Tier
UCSB SmartFarm
Service Distribution Network For Data Analytics
Edge Clouds
Public Clouds
amazon web services
Microsoft Azure
Google Cloud Platform
Data appliance
• Self managing, fault tolerant
• Hardened for hostile environments
• API-compatible with public clouds
• Variety of analytics (multi-analytics)

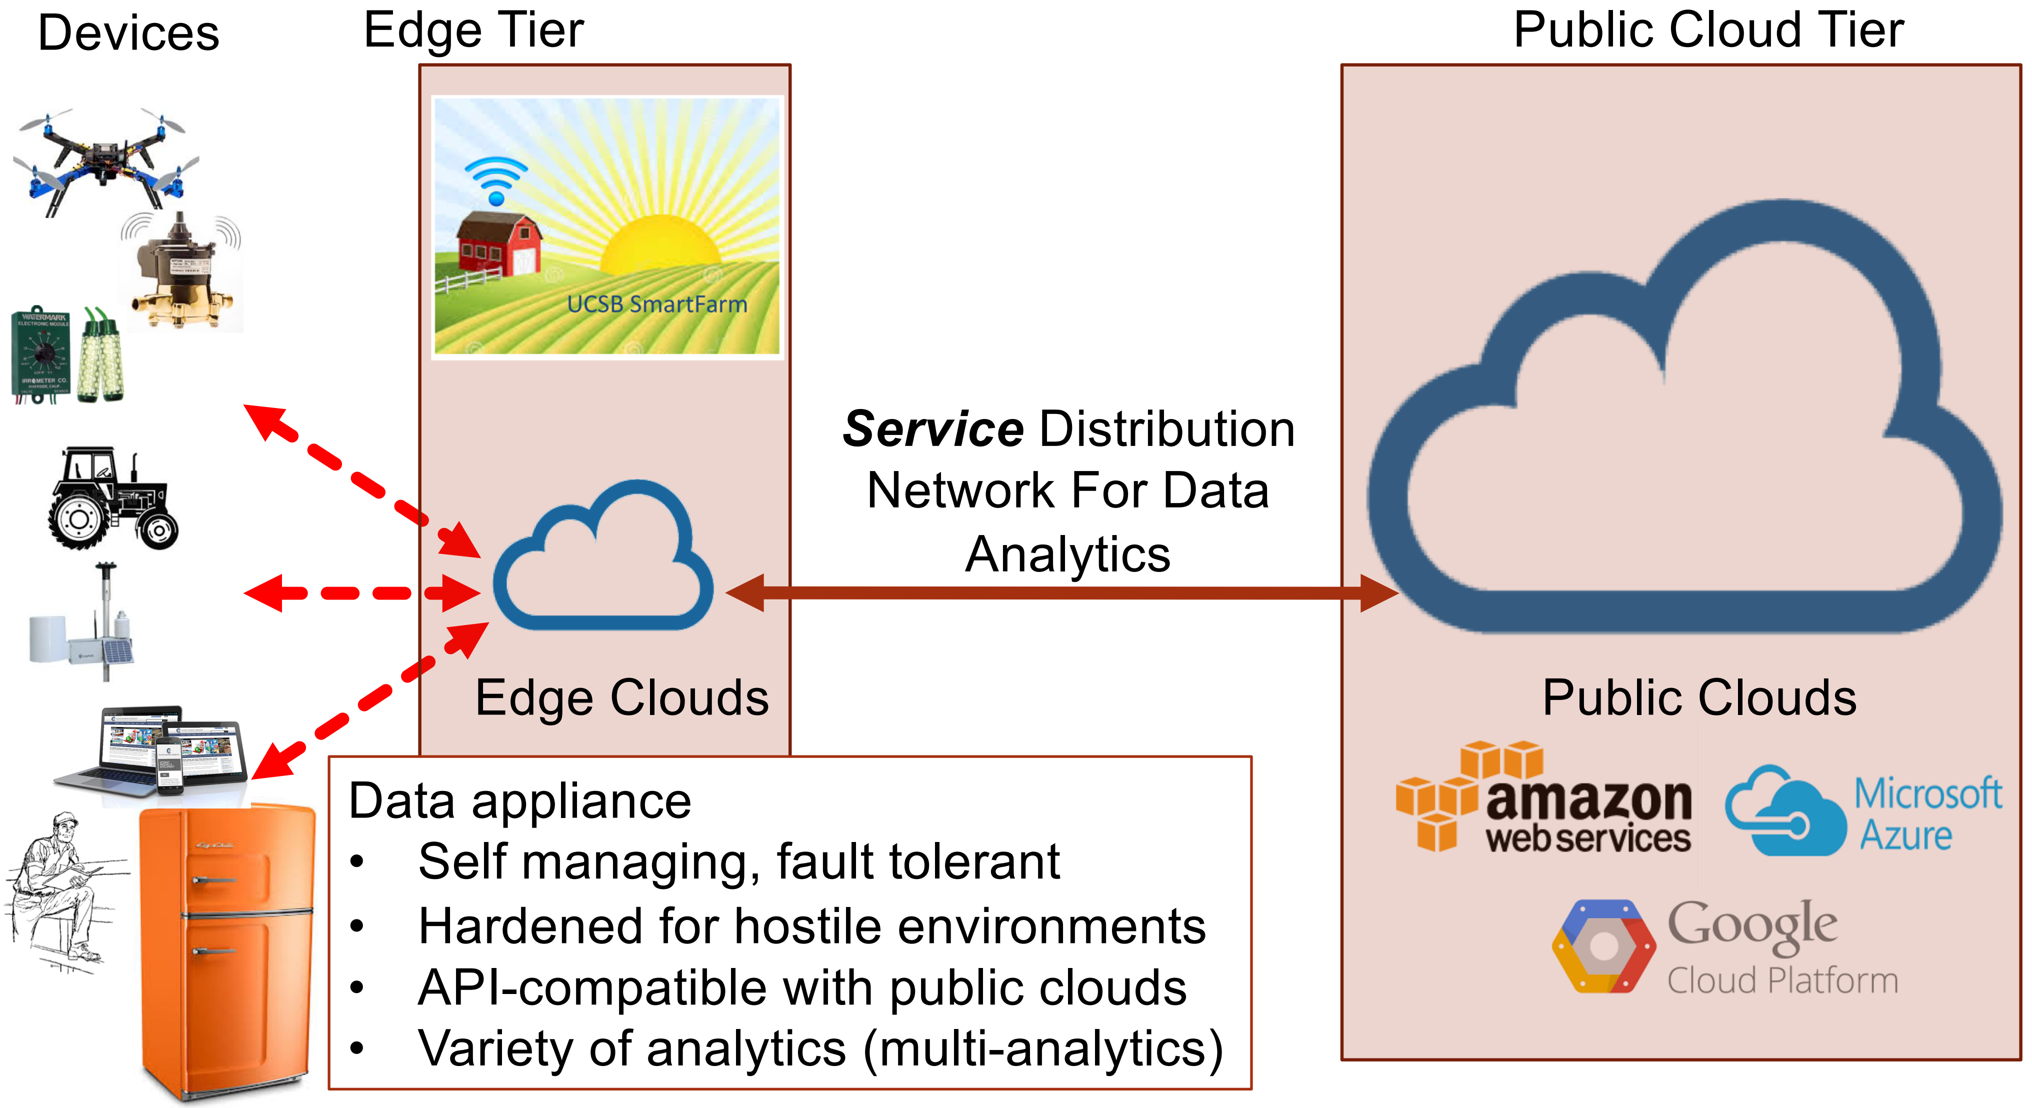
Devices
Edge Tier
Public Cloud Tier
UCSB SmartFarm
Service Distribution
Network For Data
Analytics
Edge Clouds
Public Clouds
amazon
web services
Microsoft
Azure
Google
Cloud Platform
Data appliance
• Self managing, fault tolerant
• Hardened for hostile environments
• API-compatible with public clouds
• Variety of analytics (multi-analytics)

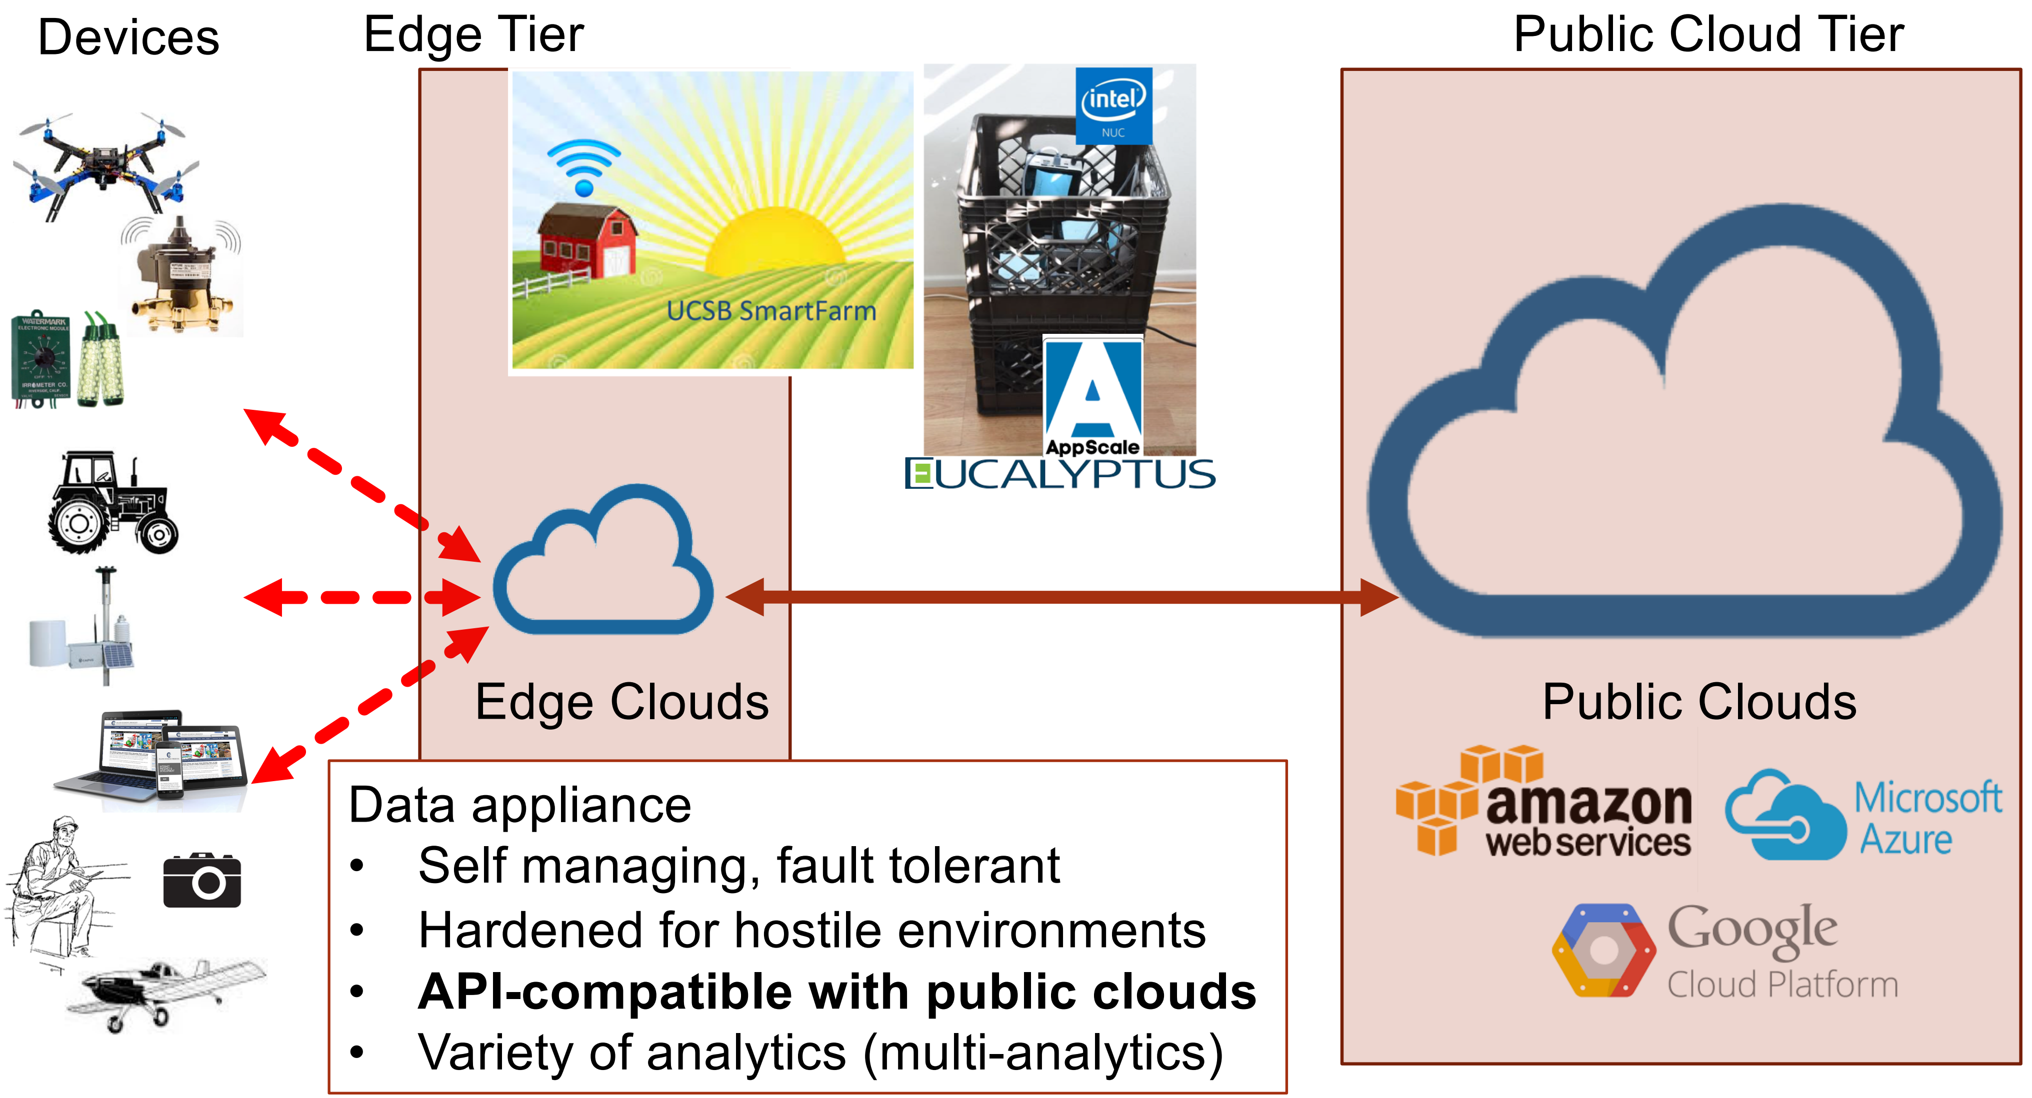

C. Krintz, **The AppScale Cloud Platform: Enabling Portable, Scalable Web Application Deployment**, IEEE Internet Computing 2013

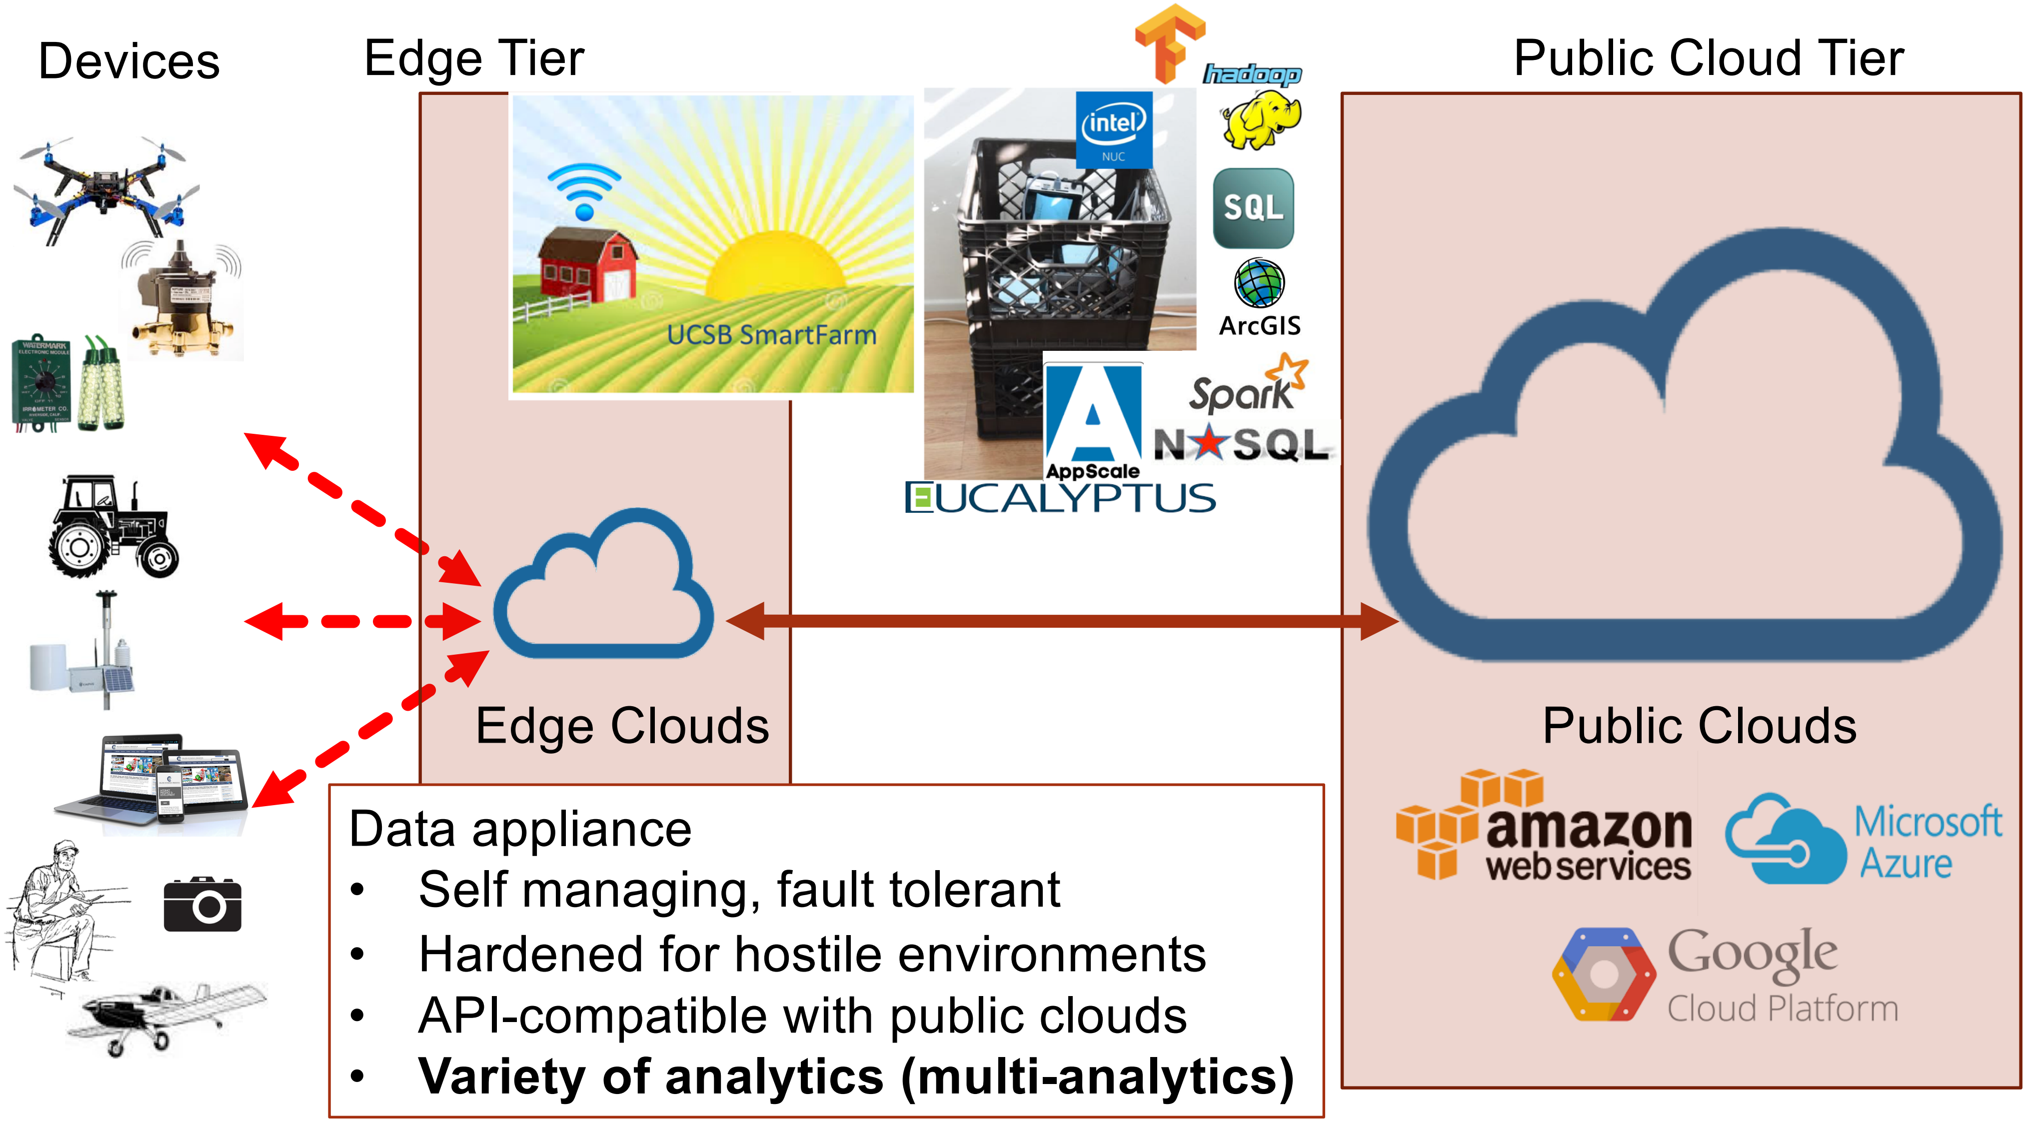

S. Dimopoulos, C. Krintz, and R. Wolski, **Big Data Framework Interference In Restricted Private Cloud Settings,** IEEE Big Data 2016

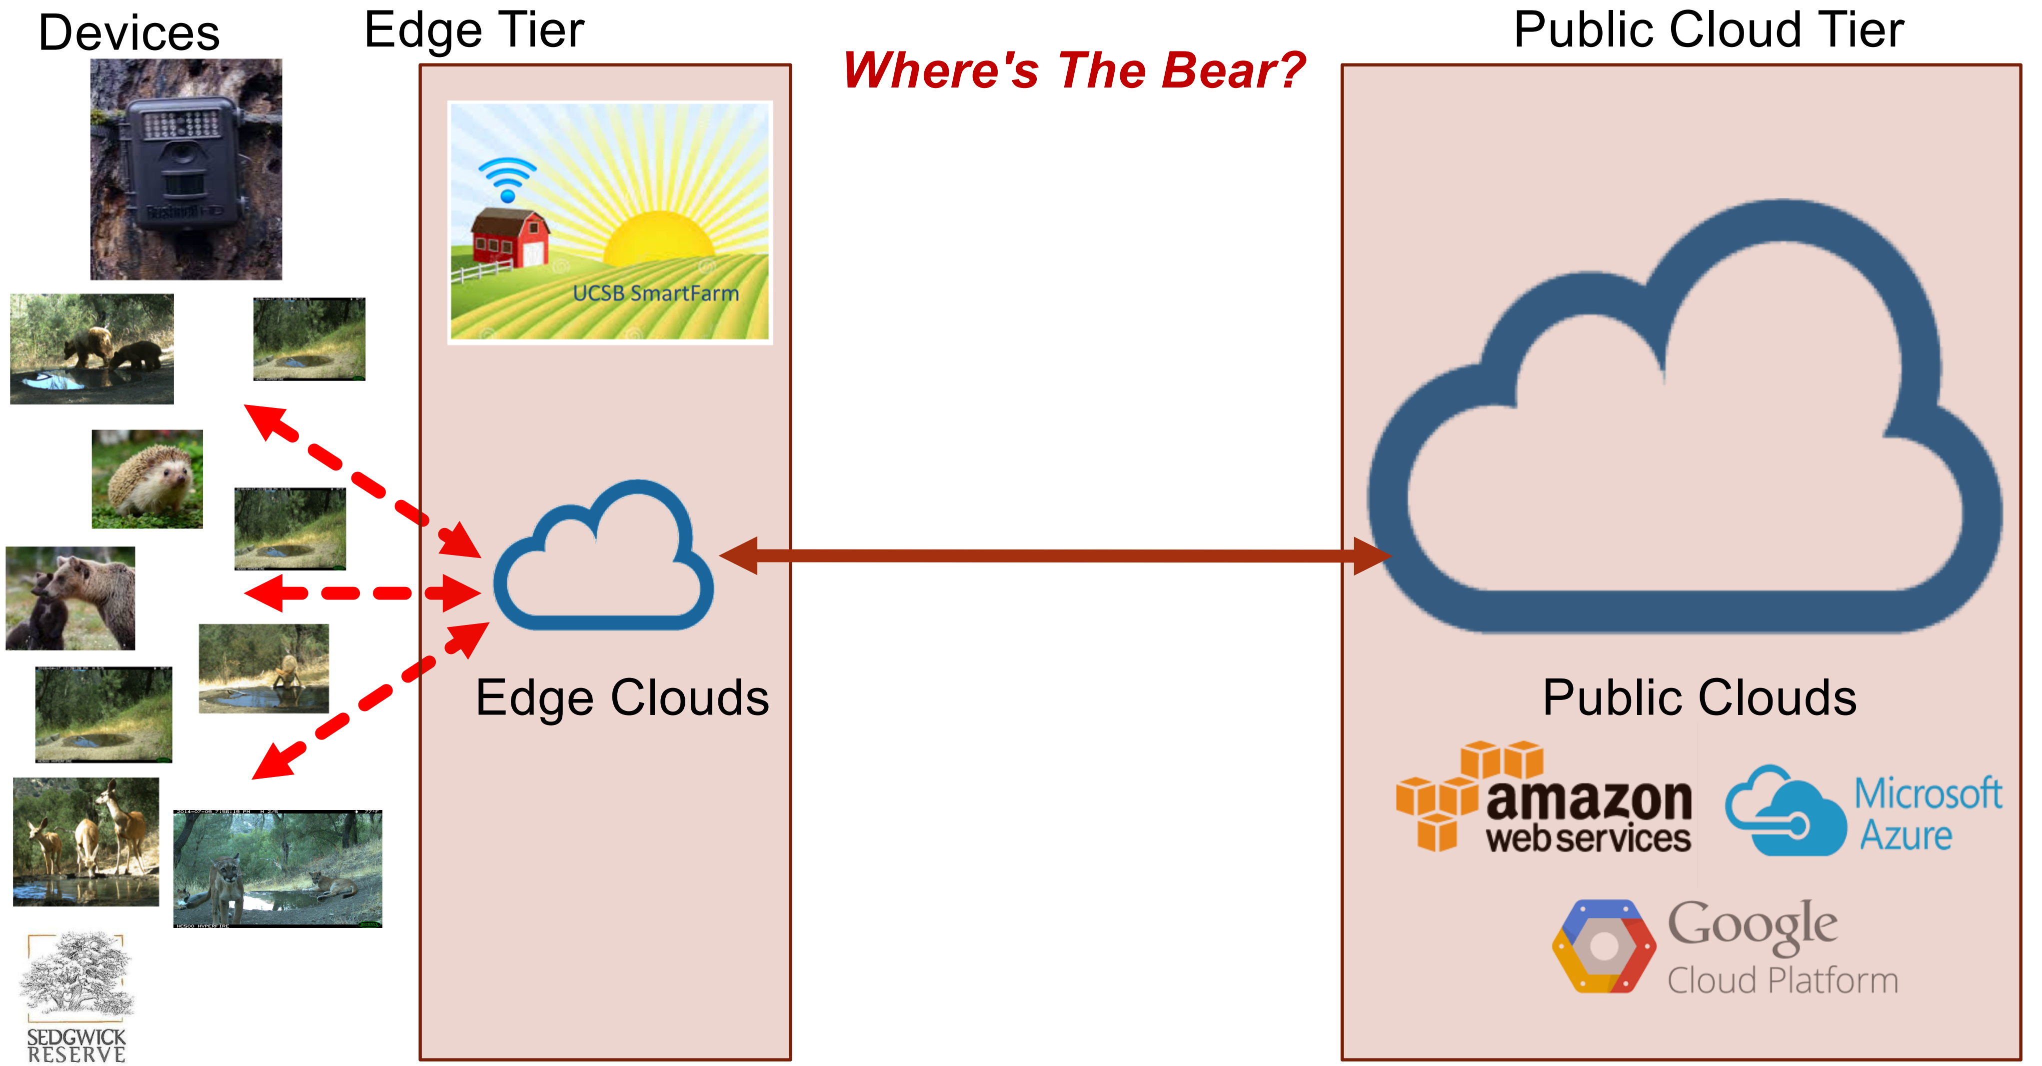

A. Rosales Elias, N. Golubovic, R. Wolski, and C. Krintz, **Where's The Bear? -- Automating Wildlife Image Processing Using IoT and Edge Cloud Systems,** ACM IoTDI 2017

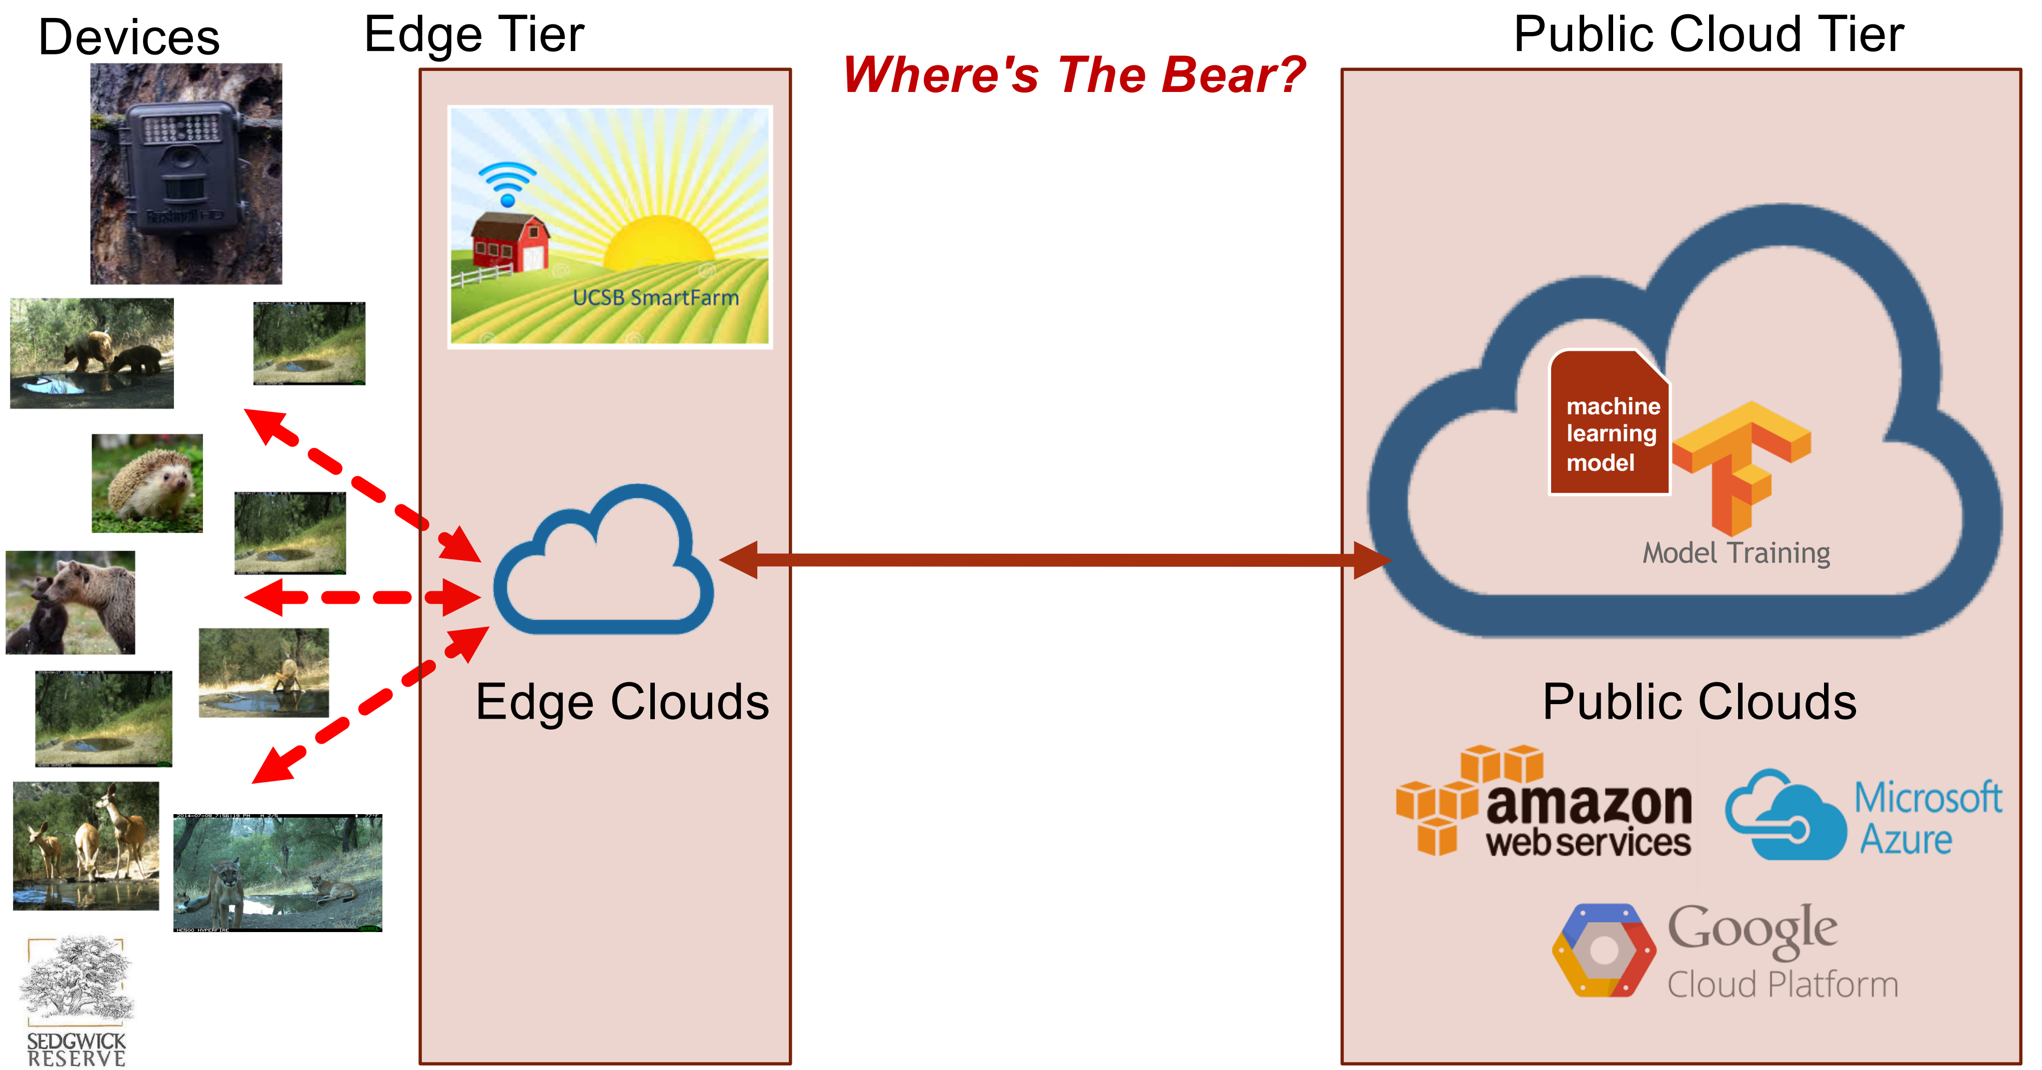

A. Rosales Elias, N. Golubovic, R. Wolski, and C. Krintz, **Where's The Bear? -- Automating Wildlife Image Processing Using IoT and Edge Cloud Systems,** ACM IoTDI 2017

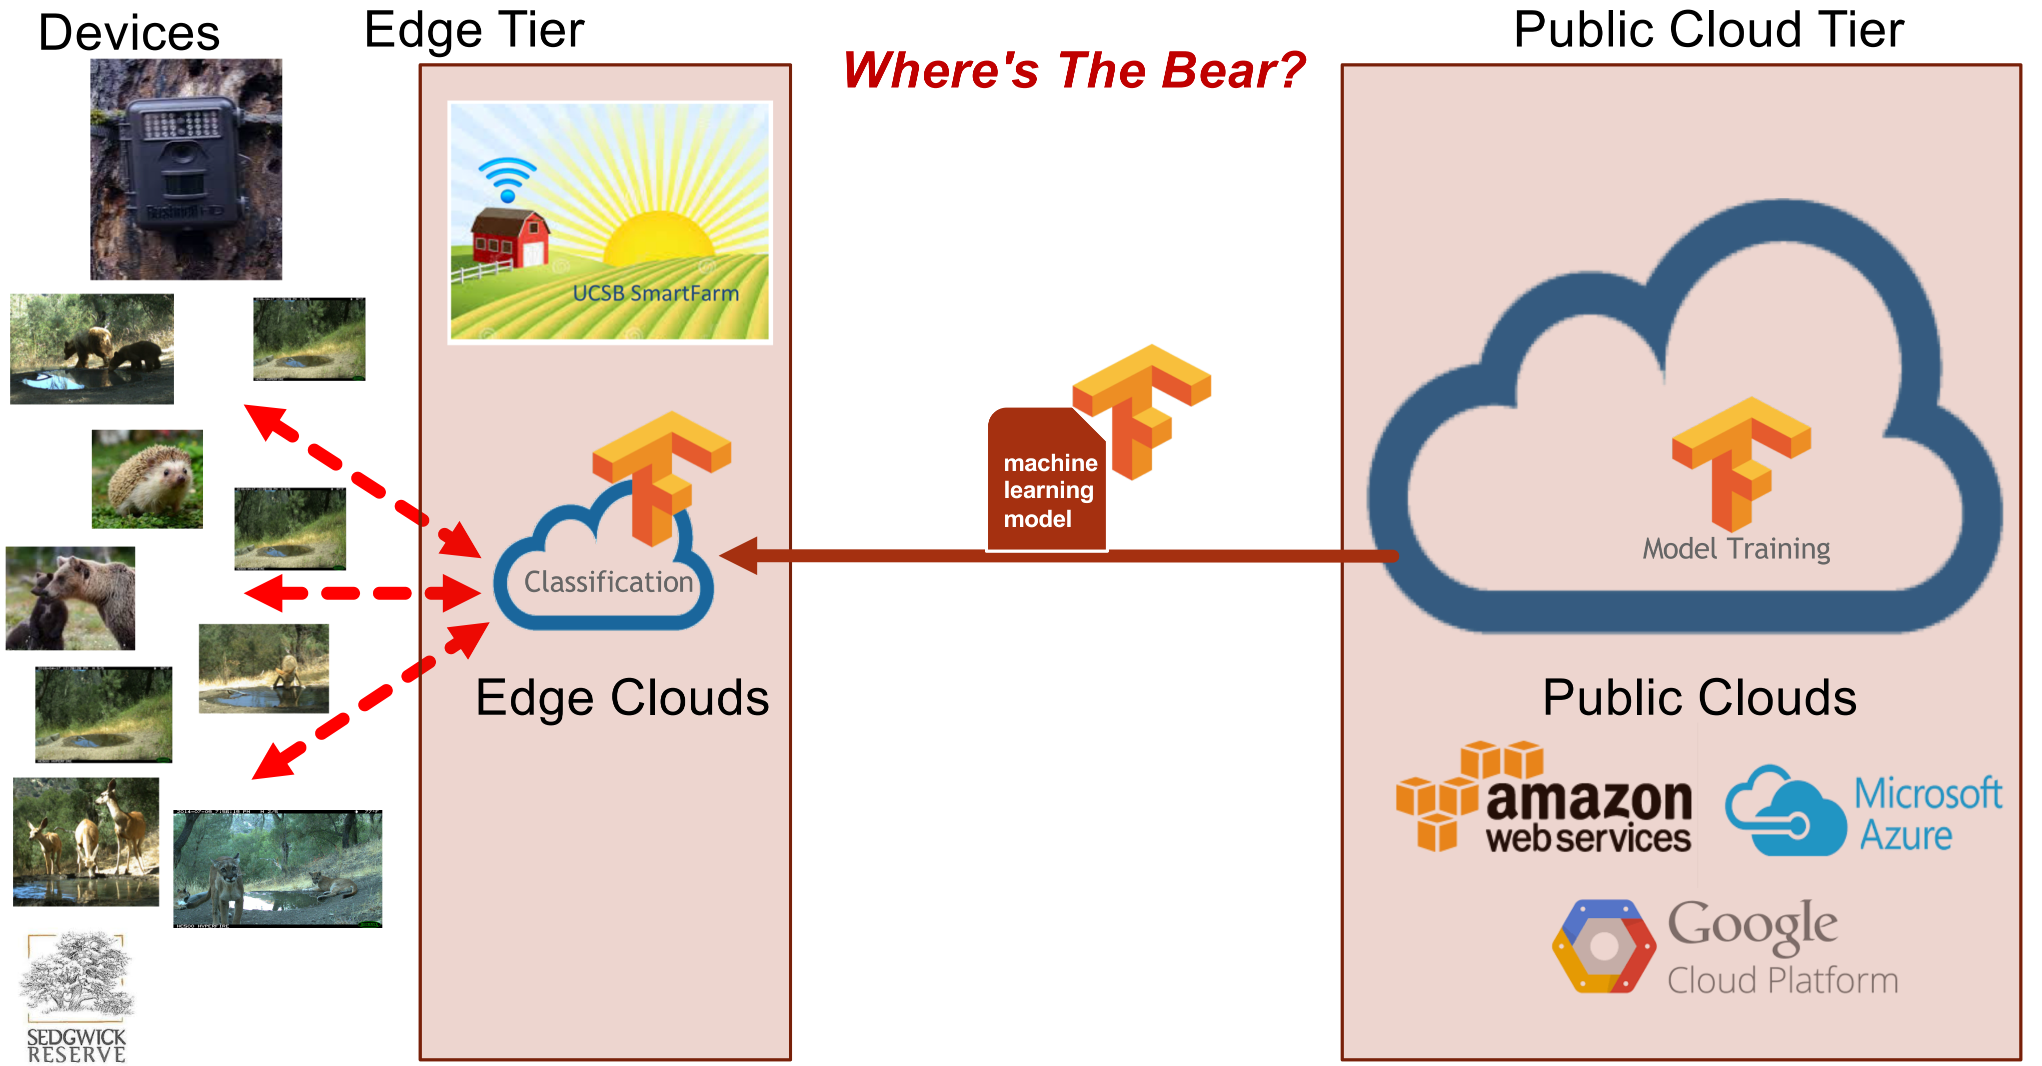

A. Rosales Elias, N. Golubovic, R. Wolski, and C. Krintz, **Where's The Bear? -- Automating Wildlife Image Processing Using IoT and Edge Cloud Systems,** ACM IoTDI 2017

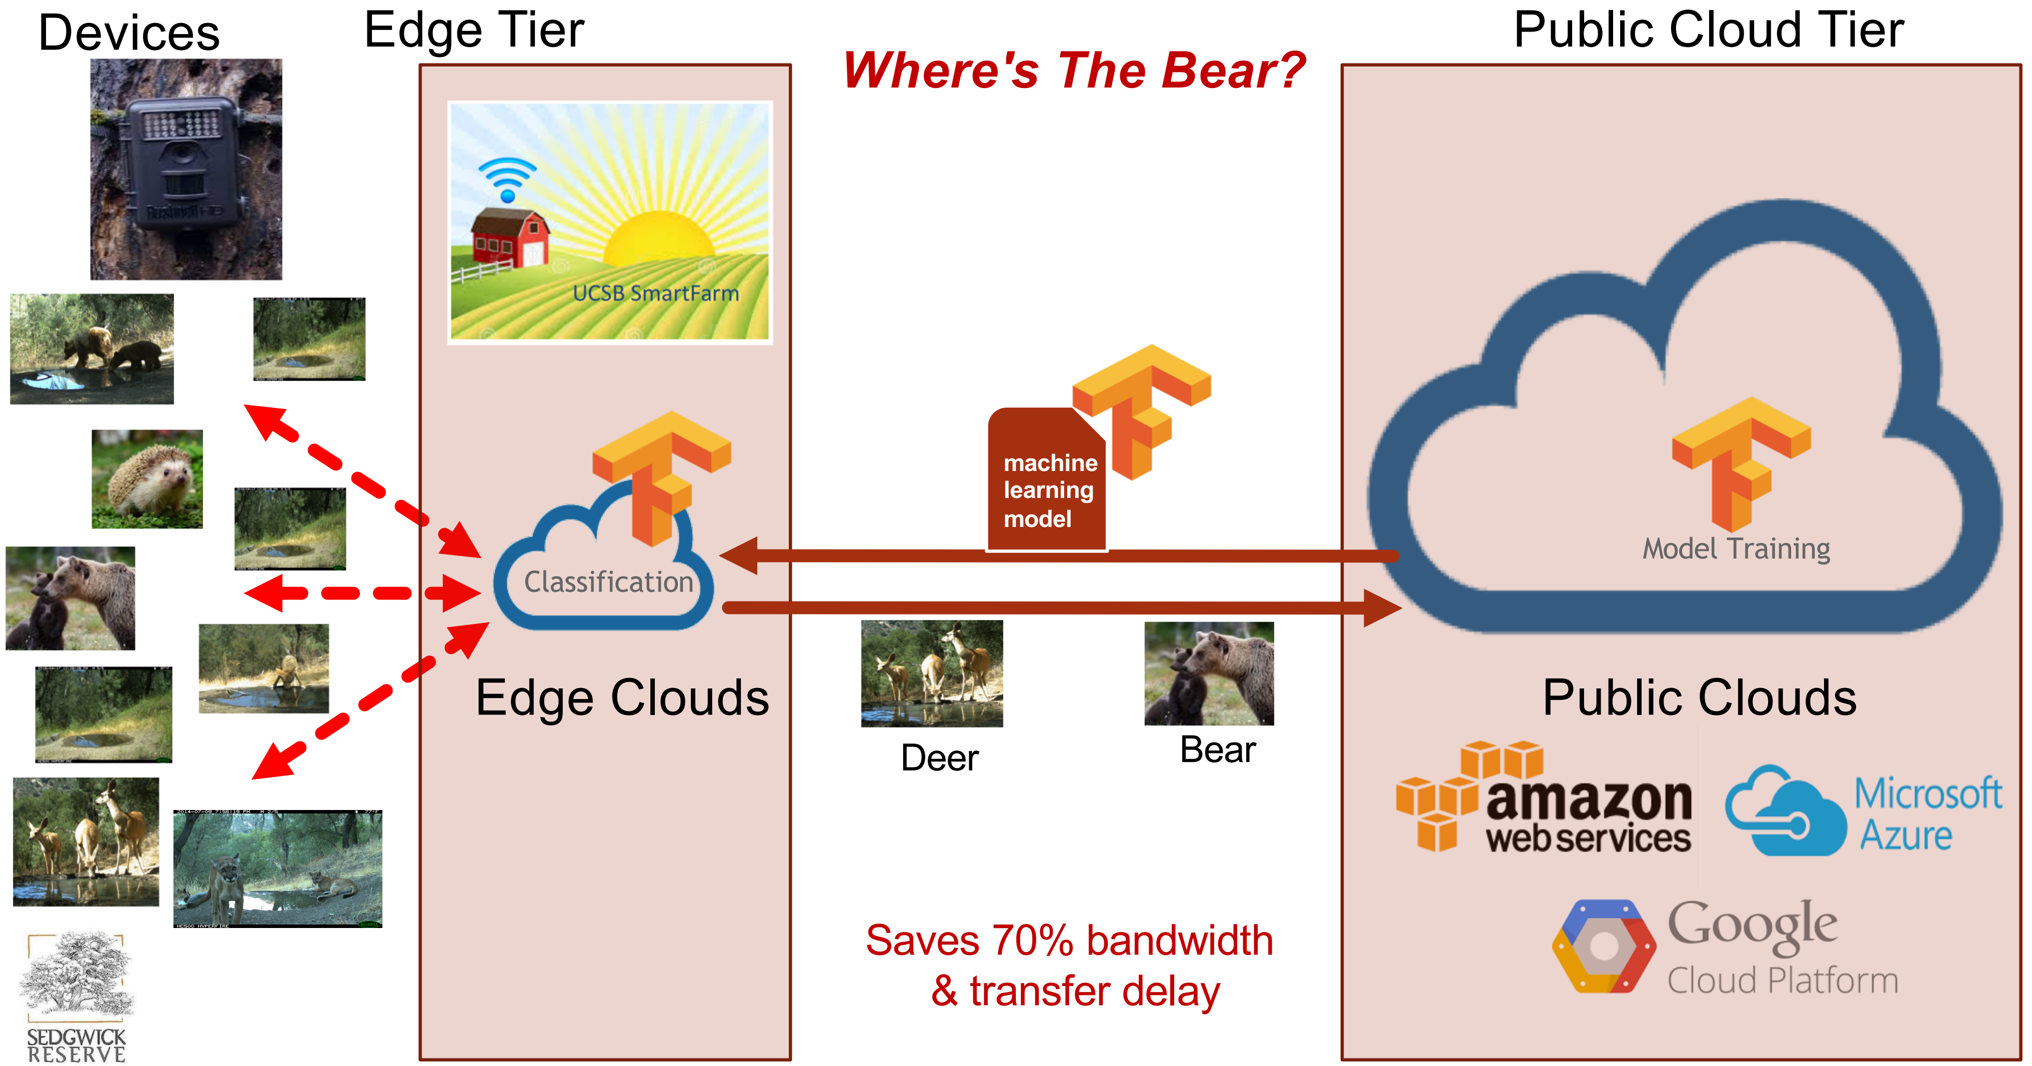

A. Rosales Elias, N. Golubovic, R. Wolski, and C. Krintz, **Where's The Bear? -- Automating Wildlife Image Processing Using IoT and Edge Cloud Systems,** ACM IoTDI 2017

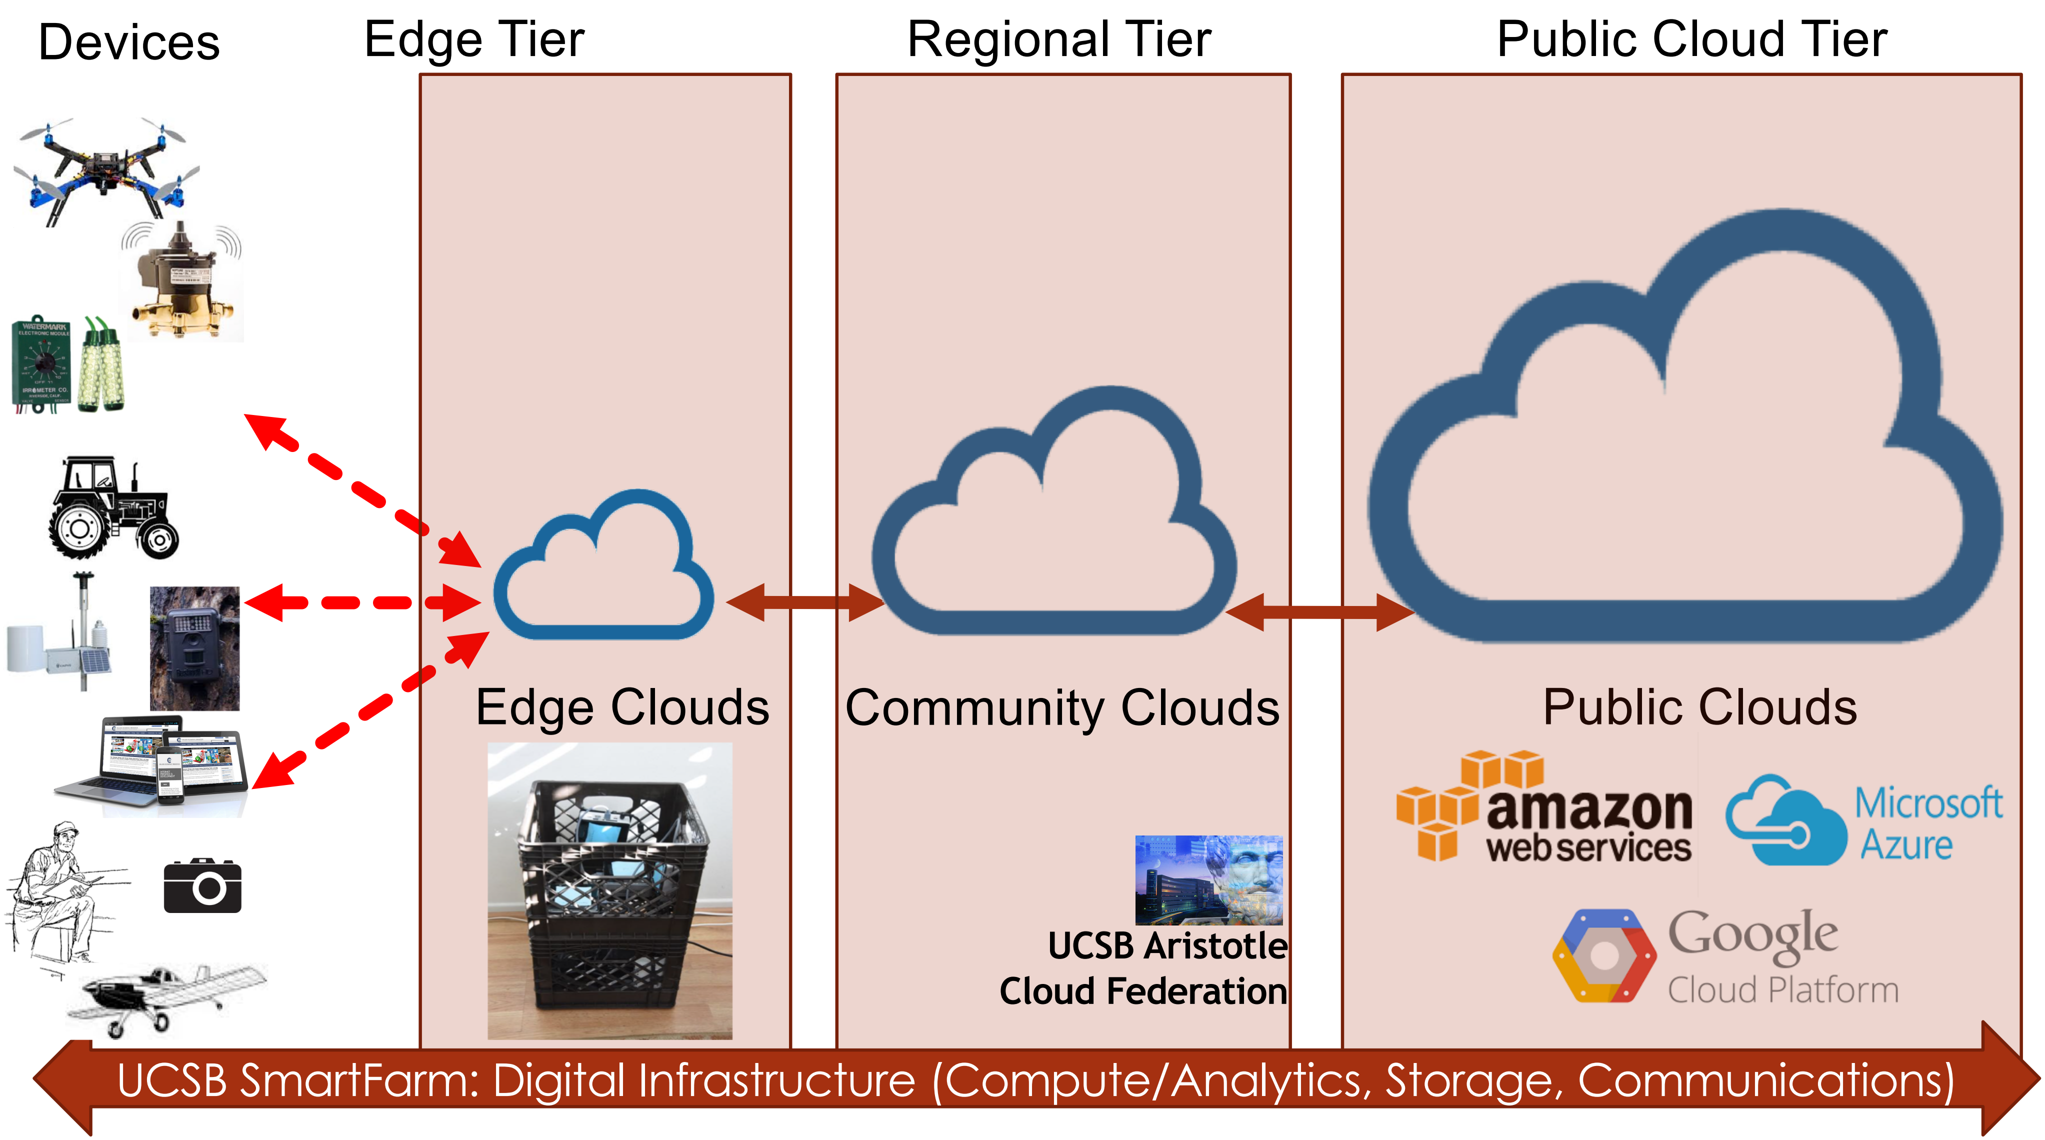
Devices
Edge Tier
Regional Tier
Public Cloud Tier
Edge Clouds
Community Clouds
Public Clouds
UCSB Aristotle
Cloud Federation
amazon
web services
Microsoft
Azure
Google
Cloud Platform
UCSB SmartFarm: Digital Infrastructure (Compute/Analytics, Storage, Communications)

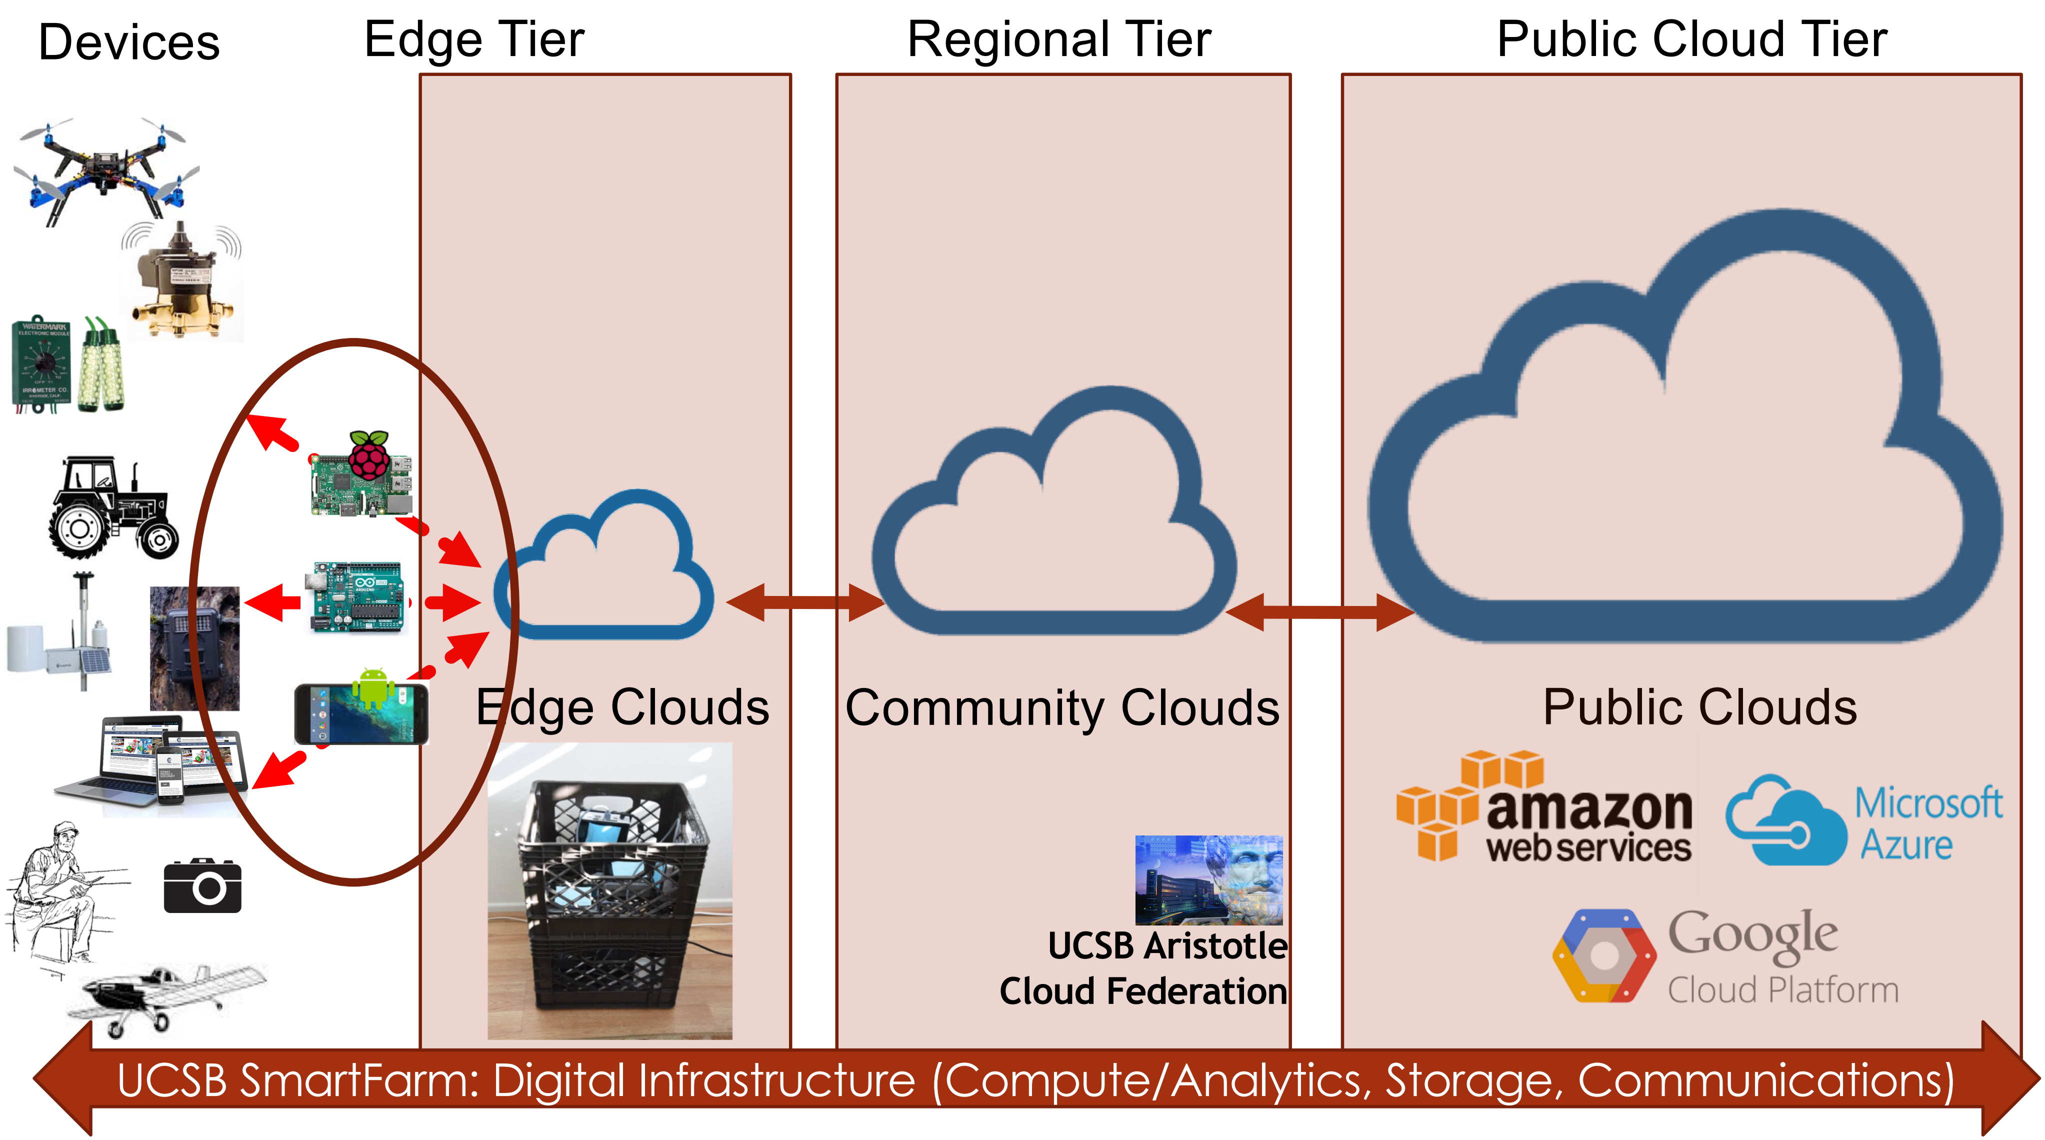
Devices
Edge Tier
Regional Tier
Public Cloud Tier
Edge Clouds
Community Clouds
Public Clouds
UCSB Aristotle
Cloud Federation
amazon web services
Microsoft Azure
Google Cloud Platform
UCSB SmartFarm: Digital Infrastructure (Compute/Analytics, Storage, Communications)

Write-Once Run-Anywhere for Multi-Tier IoT Apps

- Portable across
  - Low-level, small footprint
  - Low latency response
- Easy to develop/deploy distributed applications (novices & experts alike)
- Tailored IoT application use cases
  - Data driven: *event-based*
  - Multi-analytics
  - Executable on resource constrained or resource rich systems
  - Updatable, debuggable, maintainable

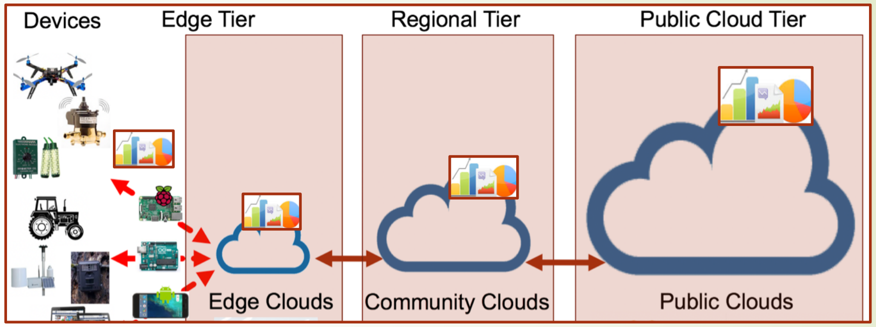

UCSB

Write-Once Run-Anywhere for Multi-Tier IoT Apps

- Portable across
  - Low-level, small footprint
  - Low latency response
- Easy to develop/deploy distributed applications (novices & experts alike)
- Tailored IoT application use cases
  - Data driven: *event-based*
  - Multi-analytics
  - Executable on resource constrained or resource rich systems
  - Updatable, debuggable, maintainable

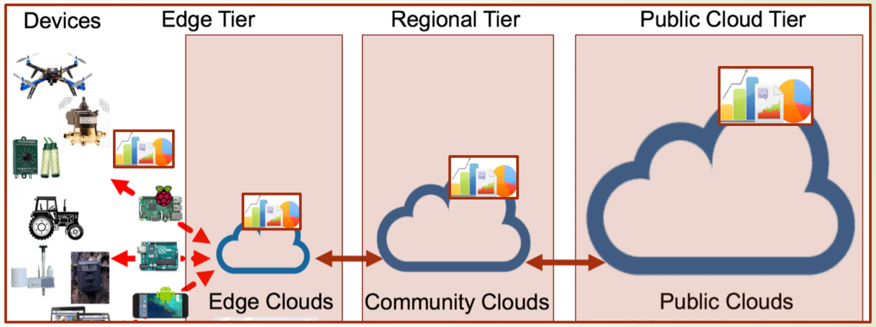

Serverless / Functions-as-a-Service (FaaS) programming model

Platform = Language + Infrastructure

UCSB

Write-Once Run-Anywhere for Multi-Tier IoT Apps

- Serverless Platform of Things (SPOT)
  - Program = Simple, ***event-triggered functions***
    - Portable across IoT systems (all tiers)
  - Open source, self-service, easy to use

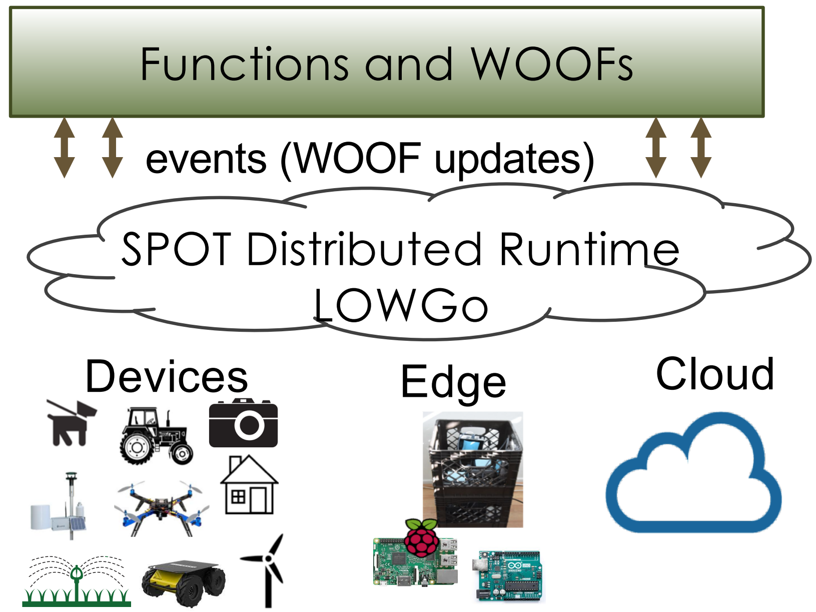

R. Wolski and C. Krintz, **CSPOT: A Serverless Platform of Things**, UCSB TR2018-01

Write-Once Run-Anywhere for Multi-Tier IoT Apps

- Serverless Platform of Things (SPOT)
  - Program = Simple, ***event-triggered functions***
    - Portable across IoT systems (all tiers)
    - Leverages Linux containers for deployment
  - Open source, self-service, easy to use
  - Wide-area Objects of Functions (WOOFs)
    - Append-only, persistent data structure
    - Event = object update
    - Tracked via causally ordered, distributed log (LOWGo)
  - Tools for debugging, replay, pause/resume, root cause analysis, …

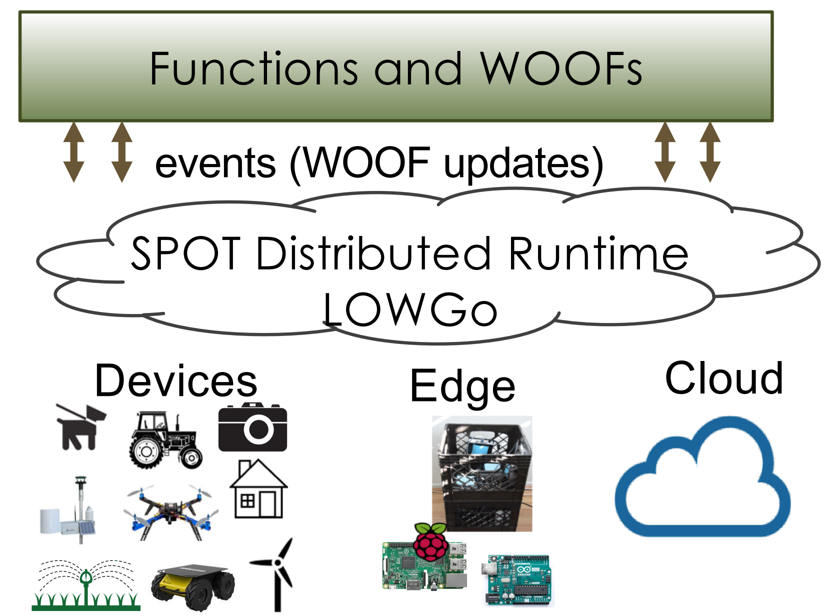

W. Lin, C. Krintz, and R. Wolski, **Tracing Function Dependencies Across Clouds**, IEEE Cloud 2018
W. Lin, M. Zhang, C. Krintz, and R. Wolski, **Tracking Causal Order in AWS Lambda Applications**, IEEE IC2E 2018

SmartFarm Applications

SmartFarm Deployments & Testbeds

- **Sensor synthesis**
  - Use one sensor to estimate the data of another
  - Example: Estimating outdoor air temperature from IoT CPU temperature sensor
    - Micro-climate monitoring
    - Estimating evapo-transpiration (ET)

C. Krintz, R. Wolski, N. Golubovic, and F. Bakir, **Estimating Outdoor Temperature from CPU Temperature for IoT Applications in Agriculture**, IoT 2018

SmartFarm Applications

SmartFarm Deployments & Testbeds

FRESNO STATE Discovery. Diversity. Distinction.

SEDGWICK RESERVE

CAL POLY SAN LUIS OBISPO

RECS

University of California Agriculture and Natural Resources

Research and Extension Center System

- Sensor synthesis
  - Use one sensor to estimate the data of another
  - Example: Estimating outdoor air temperature from IoT CPU temperature sensor
    - Micro-climate monitoring
    - Estimating evapo-transpiration (ET)

3:20pm Today!

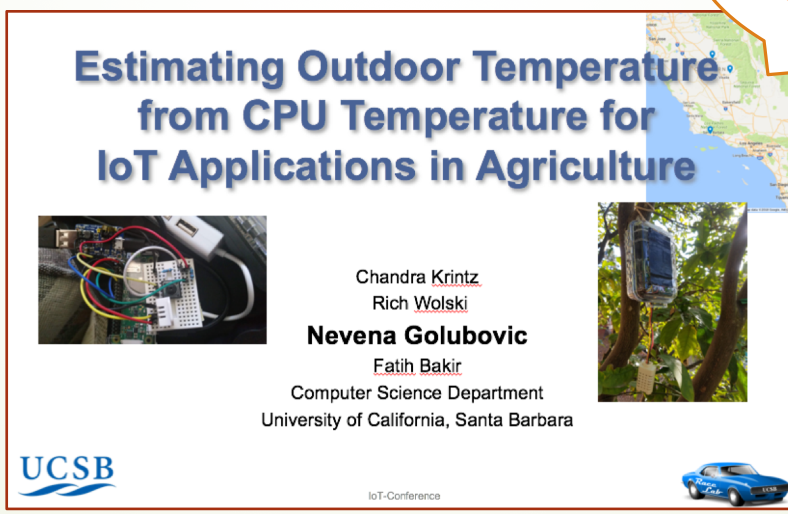

C. Krintz, R. Wolski, N. Golubovic, and F. Bakir, **Estimating Outdoor Temperature from CPU Temperature for IoT Applications in Agriculture**, IoT 2018

UCSB

SmartFarm Applications

SmartFarm Deployments & Testbeds

FRESNO STATE
Discovery. Diversity. Distinction.
SEDGWICK RESERVE
CAL POLY
SAN LUIS OBISPO
RECS
University of California
Agriculture and Natural Resources
Research and Extension Center System

- **Sensor synthesis**
  - Use one sensor to estimate the data of another
  - Example: Estimating outdoor air temperature from IoT CPU temperature sensor
    - Micro-climate monitoring, ET
- Soil electrical conductivity (EC) analysis for management zone delineation
  - Estimates water holding capacity, soil health, compaction, …
  - Multi-scale ***scored*** clustering (@Edge)

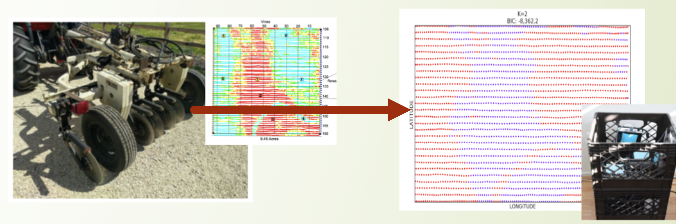

N. Golubovic, C. Krintz, R. Wolski, B. Sethuramasamyraja, and B. Liu, **A Scalable System for Executing and Scoring K-Means Clustering Techniques and Its Impact on Applications in Agriculture**, Intl. Journal of Big Data Intelligence 2018

UCSB

SmartFarm Applications

SmartFarm
Deployments
& Testbeds

FRESNO STATE
Discovery. Diversity. Distinction.

SEDGWICK RESERVE

CAL POLY
SAN LUIS OBISPO

RECS

University of California
Agriculture and Natural Resources
Research and Extension Center System

- **Sensor synthesis**
  - Use one sensor to estimate the data of another
  - Example: Estimating outdoor air temperature from IoT CPU temperature sensor
    - Micro-climate monitoring, ET
- Soil electrical conductivity (EC) analysis for management zone delineation

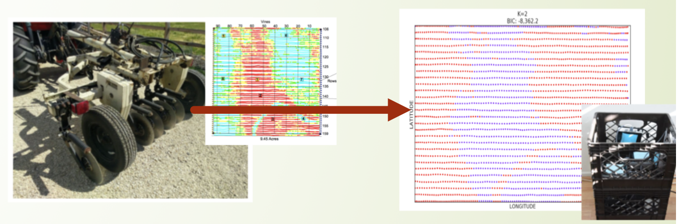

- Automated frost protection (@Edge) triggered by *predicted* micro-climate temperature inversions
- Autonomous sensing (ground robots)

- Precision weeding & input application

UCSB

A New Kind of IoT Research

- Problem driven and empirical
    - Food-Energy-Water nexus
- Societal and regional impact
- Multidisciplinary collaboration
- Repeatable, demonstrable, applied (tech-transfer ready)
- Engaging students & the community

UCSB

Thanks!

UCSB RACELab

The Lab for Research on Adaptive Computing Environments
Computer Science Department, Harold Frank Hall (E-5), Santa Barbara, CA

- Co-Lead: Dr. Rich Wolski
- Collaborators: UCSB, LREC, Dr. Bo Liu - CalPoly, Dr. Balaji Seth - Fresno State, Powwow Energy, Sedgwick Reserve, Private Growers
- Support: Google, Huawei, IBM Research, Microsoft Research, NSF, NIH, California Energy Commission

Students:

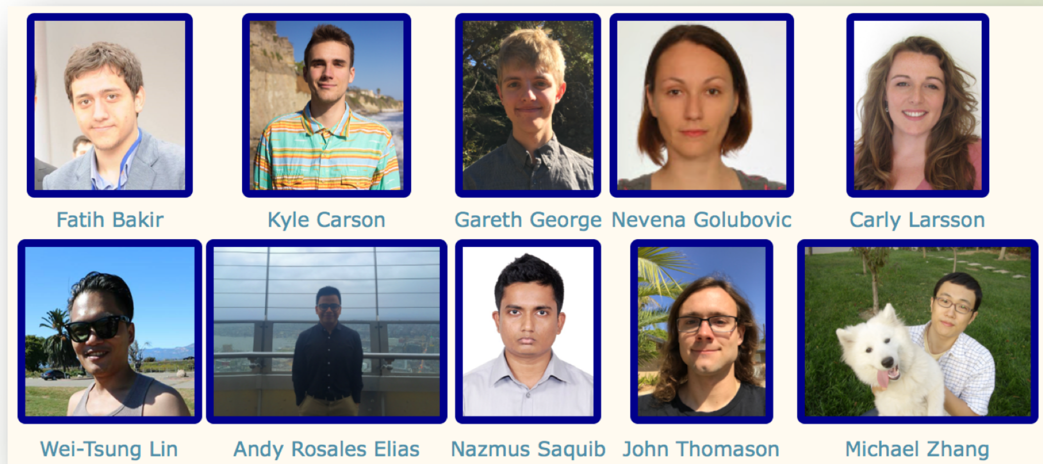

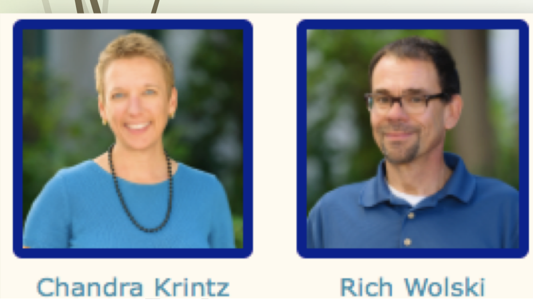

ckrintz@cs.ucsb.edu, rich@cs.ucsb.edu
http://www.cs.ucsb.edu/~ckrintz/racelab.html

UCSB