The **RISE** Lab: Real-time Intelligent Secure Execution

Krste Asanovic, Ali Ghodsi, Ken Goldberg, **Joey Gonzalez**, **Joe Hellerstein**, Michael Jordan, Randy Katz, Dave Patterson, **Raluca Ada Popa**, **Ion Stoica**, …

Berkeley's lab tradition

The Parallel Computing Laboratory

ASPIRE

RAD Lab

amplab

- Working for 5-6 years on a new major problem
- Bringing faculty from different areas

AMPLab (2010—2016)

Created popular open-source big data analytics:
Spark, Mesos, Tachyon..

AMPLab is coming to an end (December 2016)

What is the next vision?

RISE Lab

From live data to real-time decisions

AMP Lab

From batch data to advanced analytics

amplab

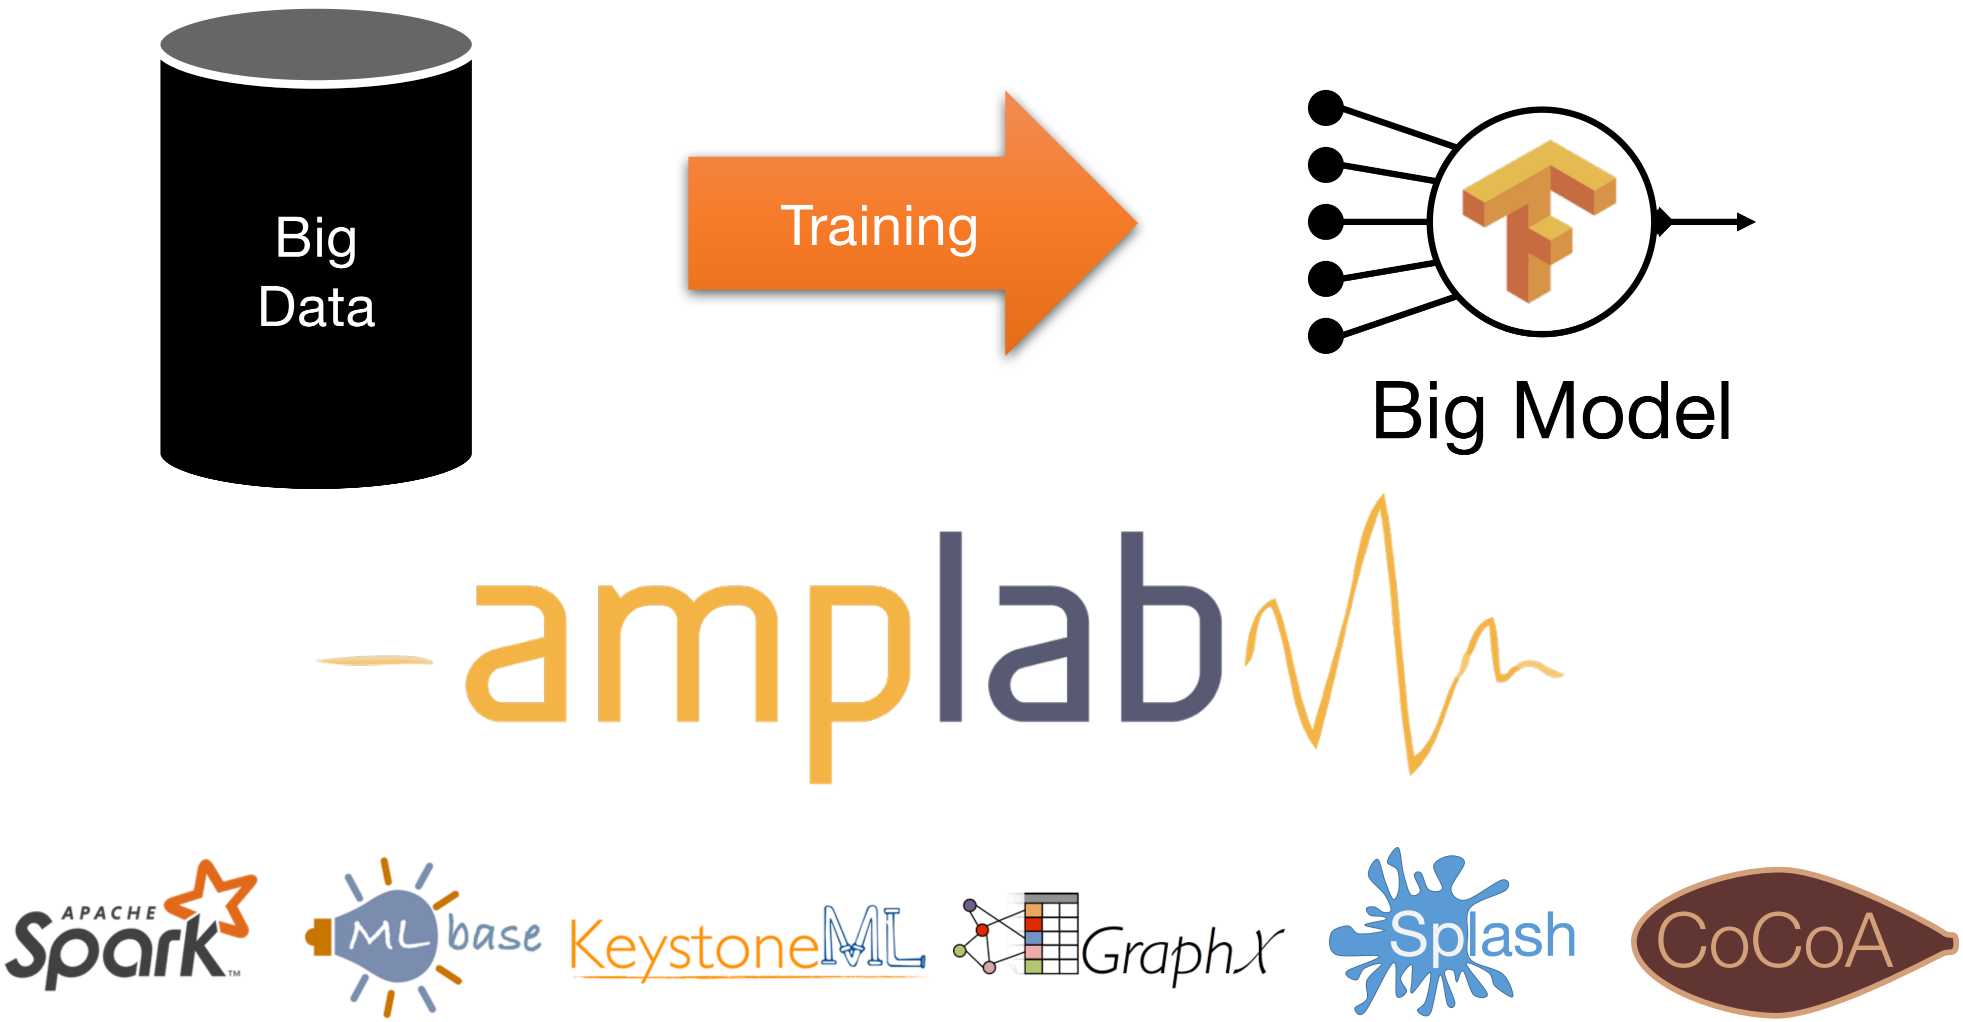

Big
Data
Training
Big Model
amplab
APACHE
Spark
ML base
KeystoneML
GraphX
Splash
CoCoA

DMLC TensorFlow Caffe VW GraphLab scikit learn

Apache Spark MLbase KeystoneML GraphX Splash CoCoA

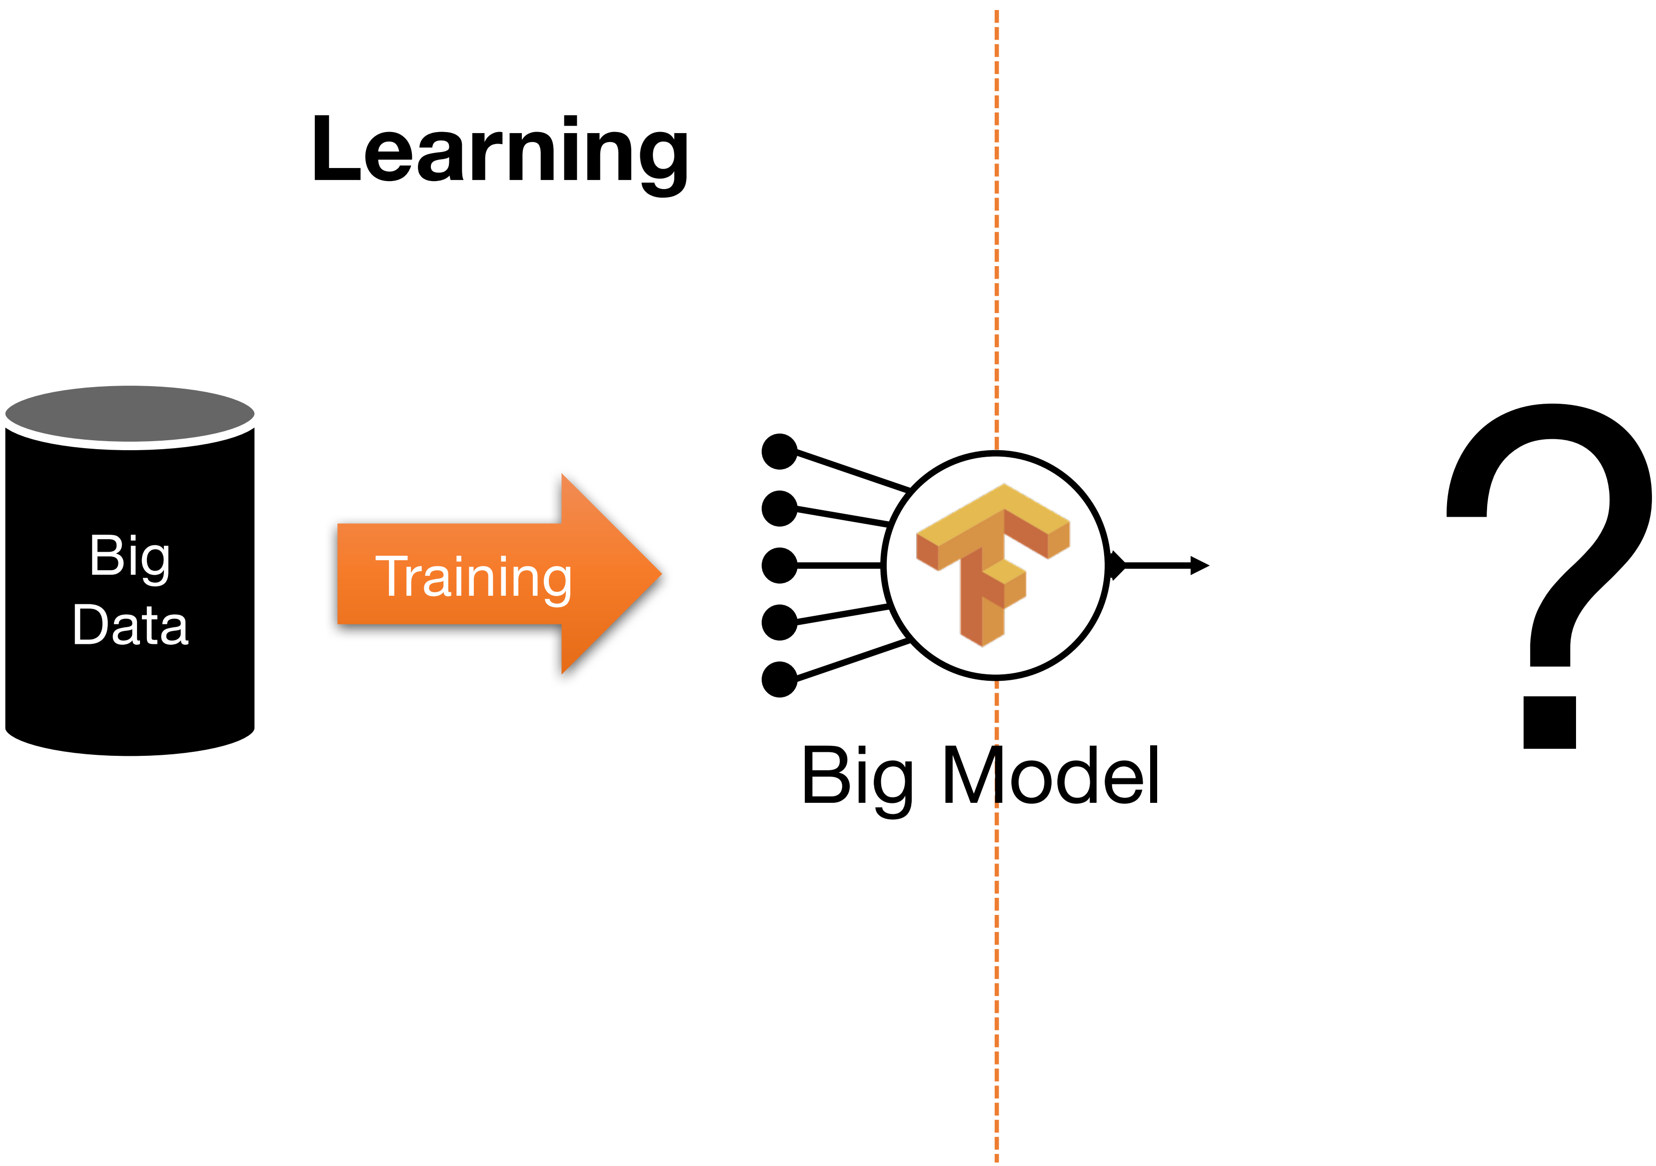
Learning
Big
Data
Training
Big Model
?

Learning

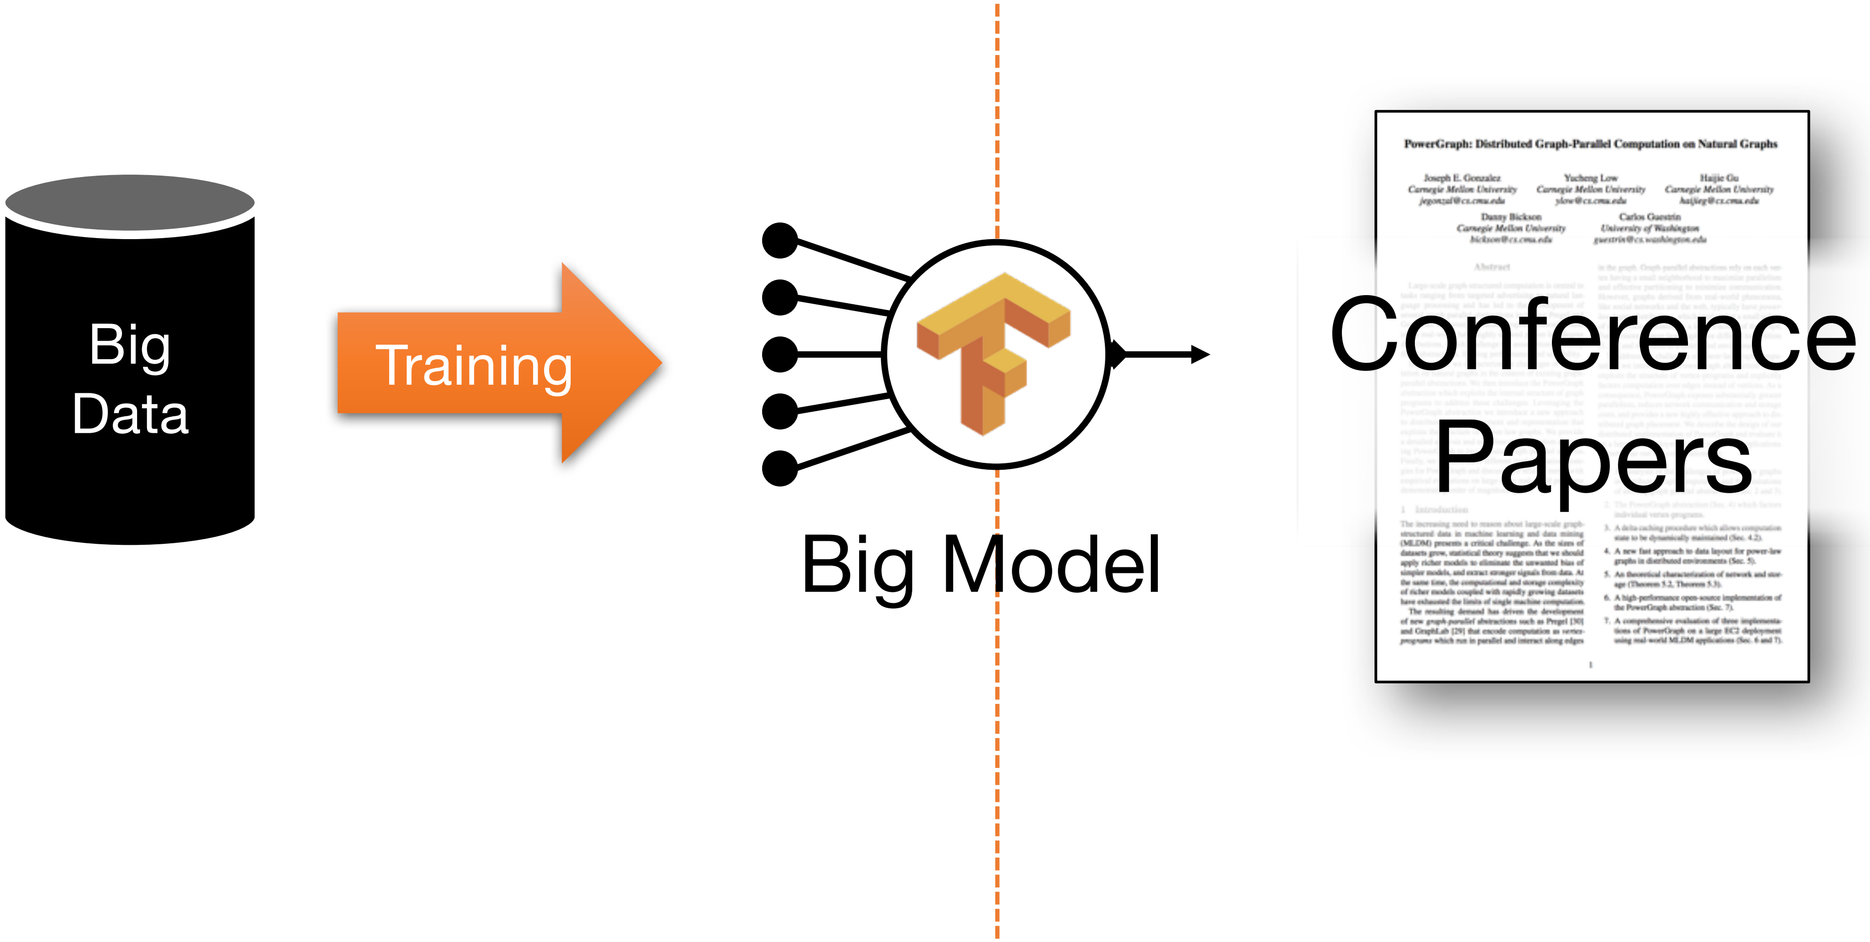
Big
Data
Training
Big Model
PowerGraph: Distributed Graph-Parallel Computation on Natural Graphs
Conference
Papers

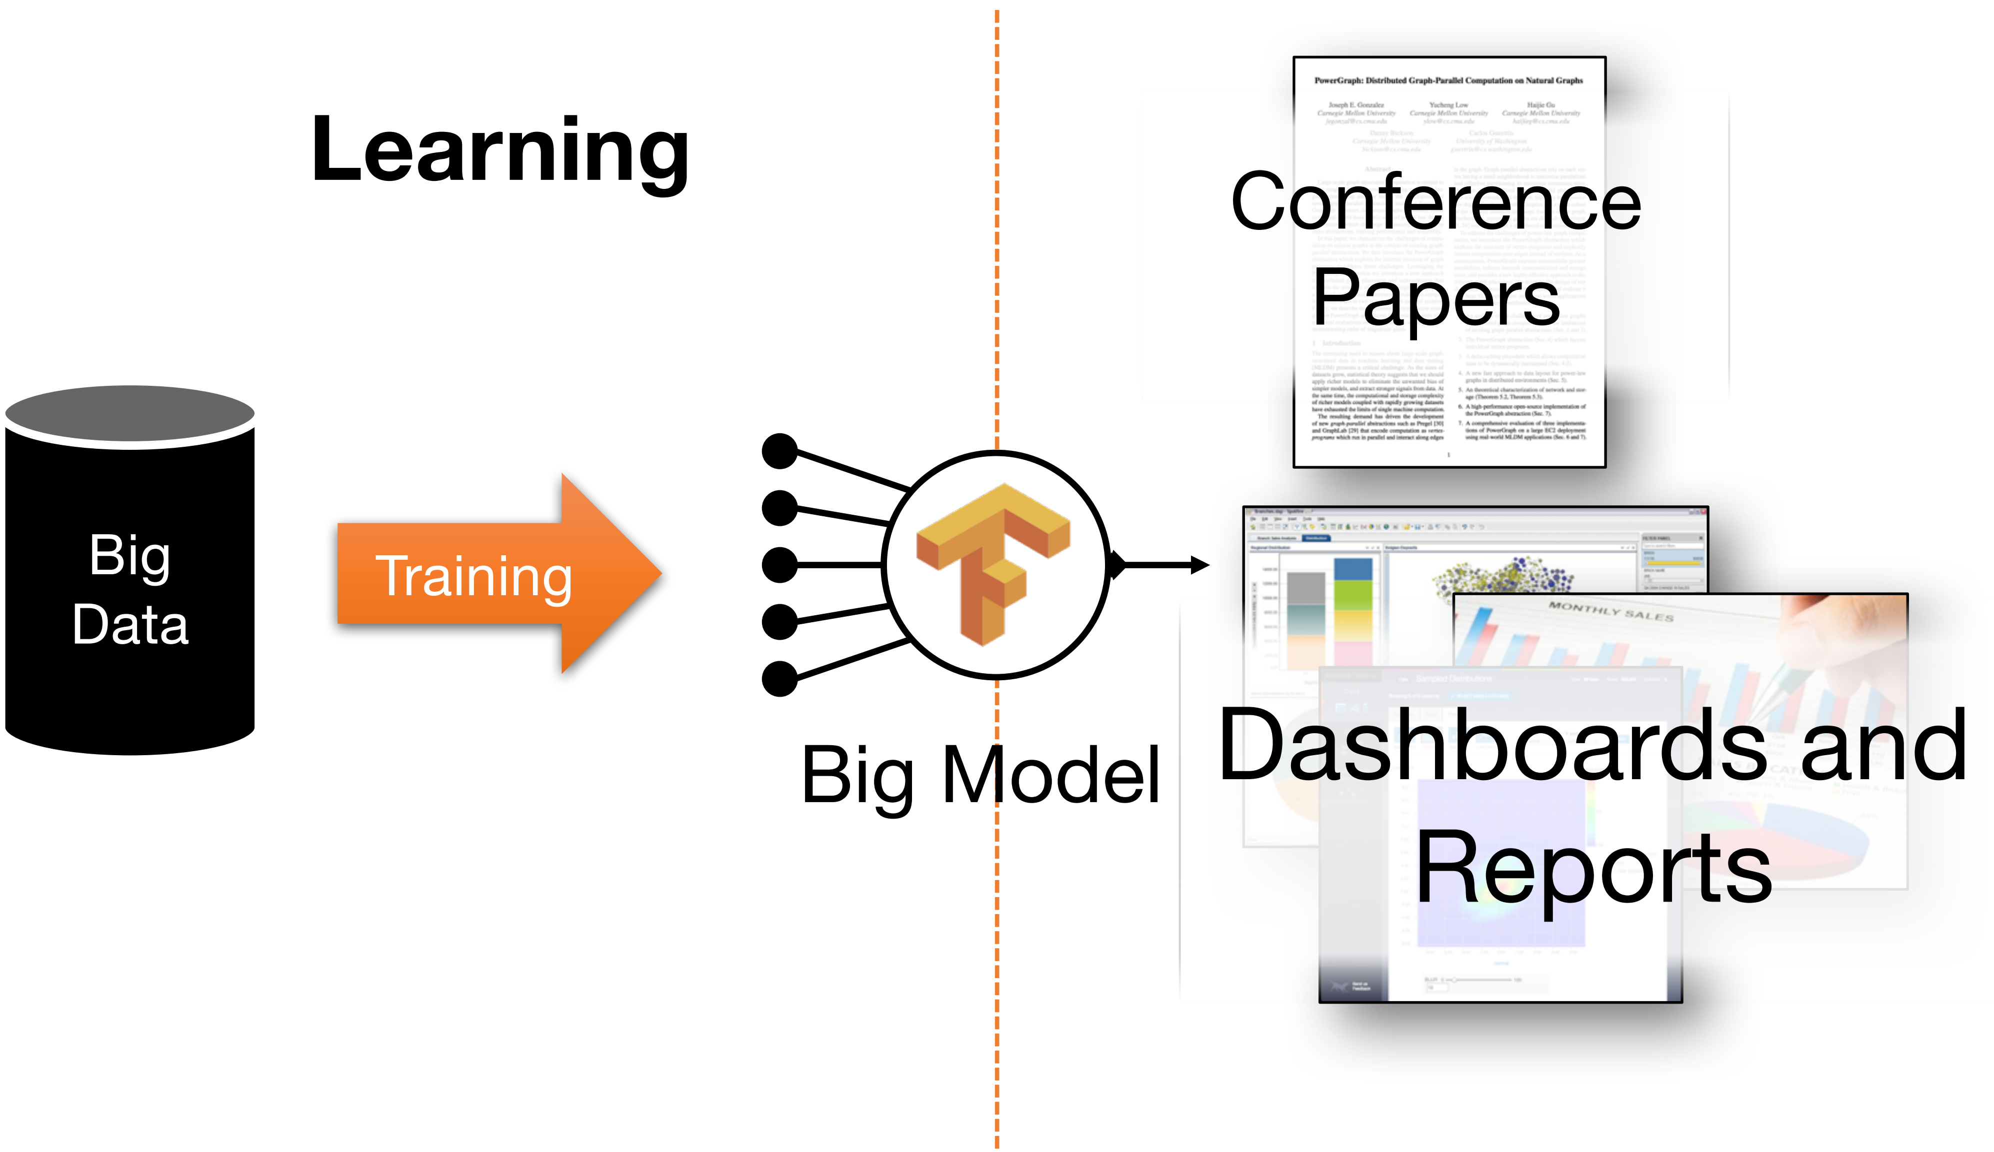
Learning
Big Data
Training
Big Model
Conference Papers
Dashboards and Reports

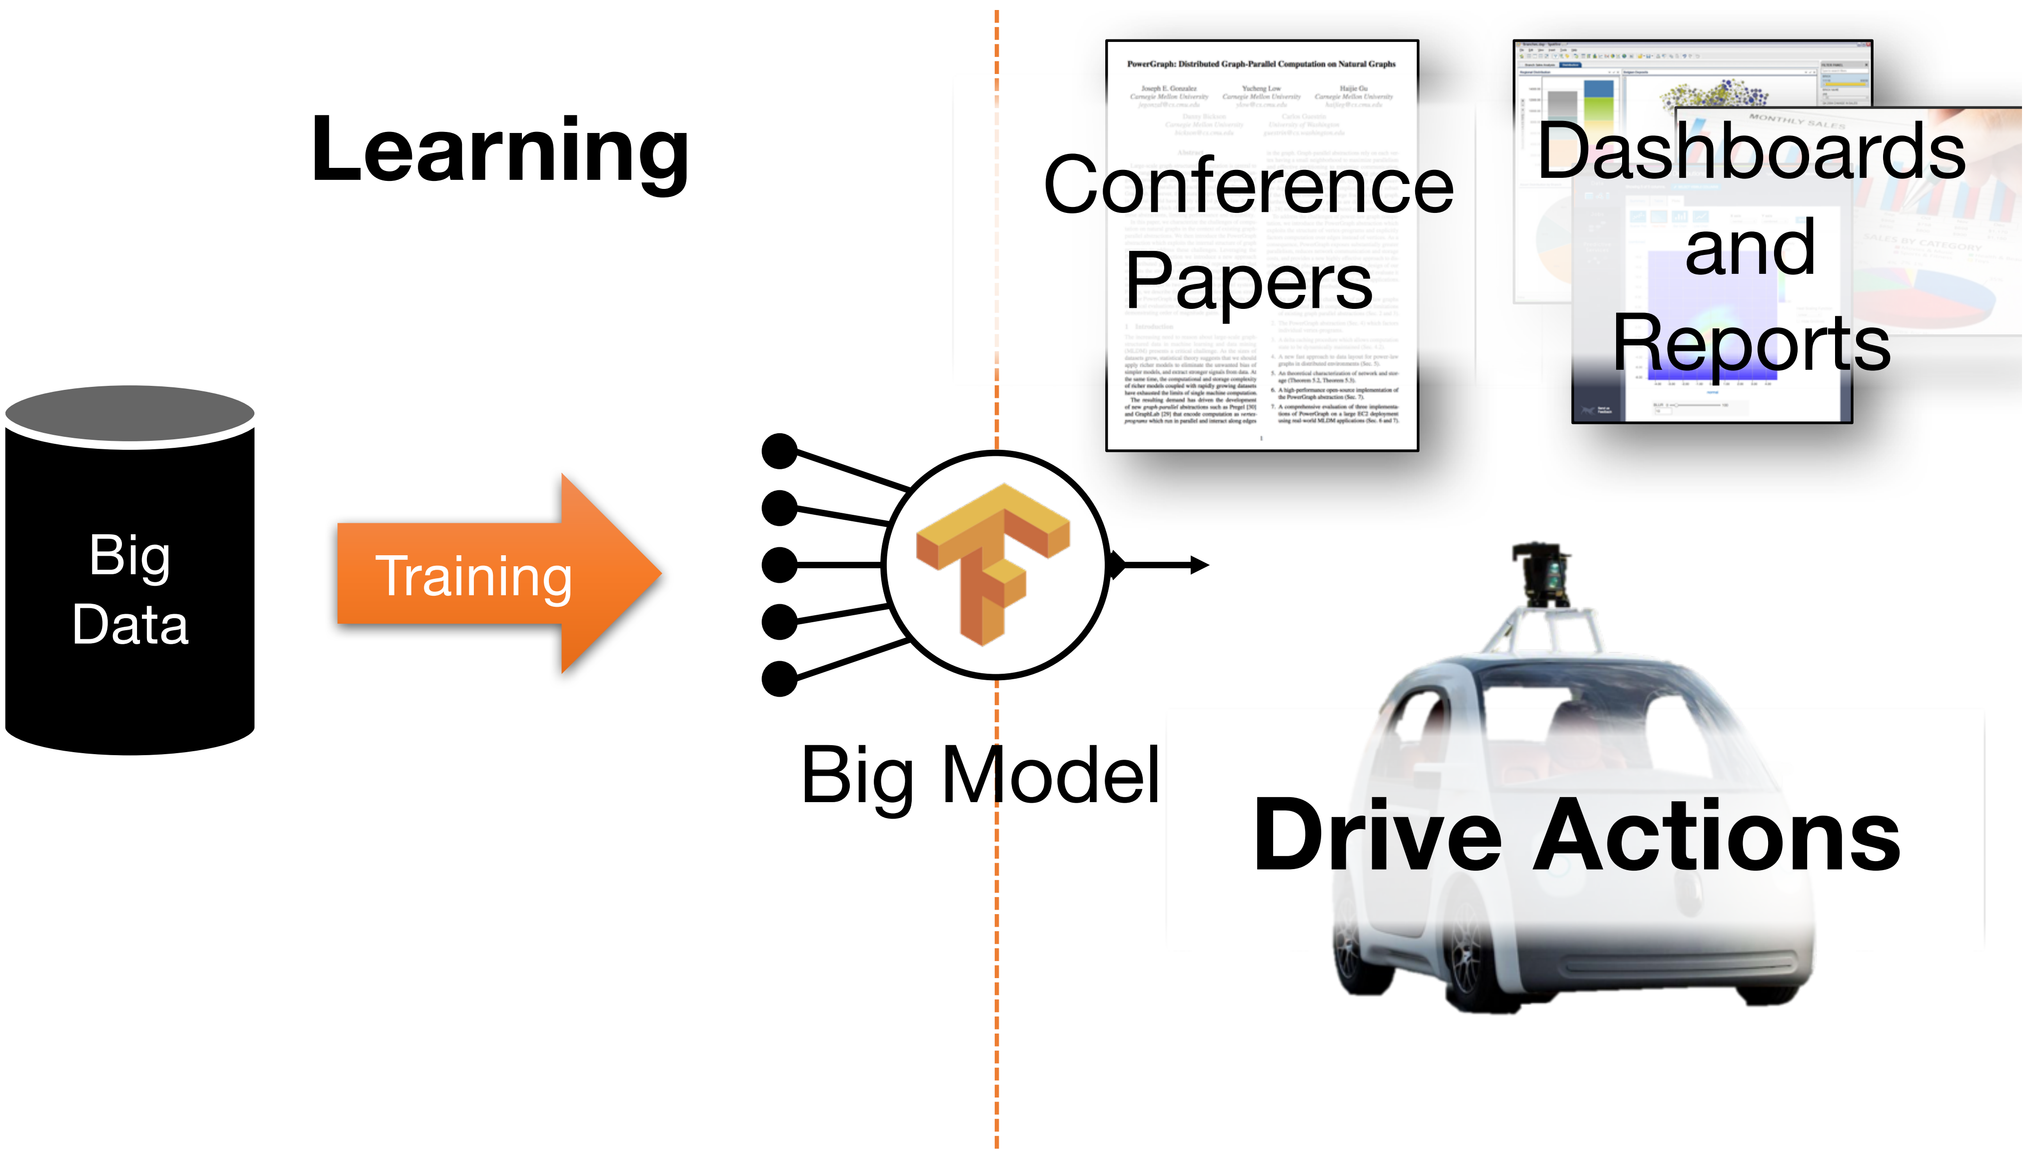
Learning
Big Data
Training
Big Model
Conference Papers
Dashboards and Reports
Drive Actions

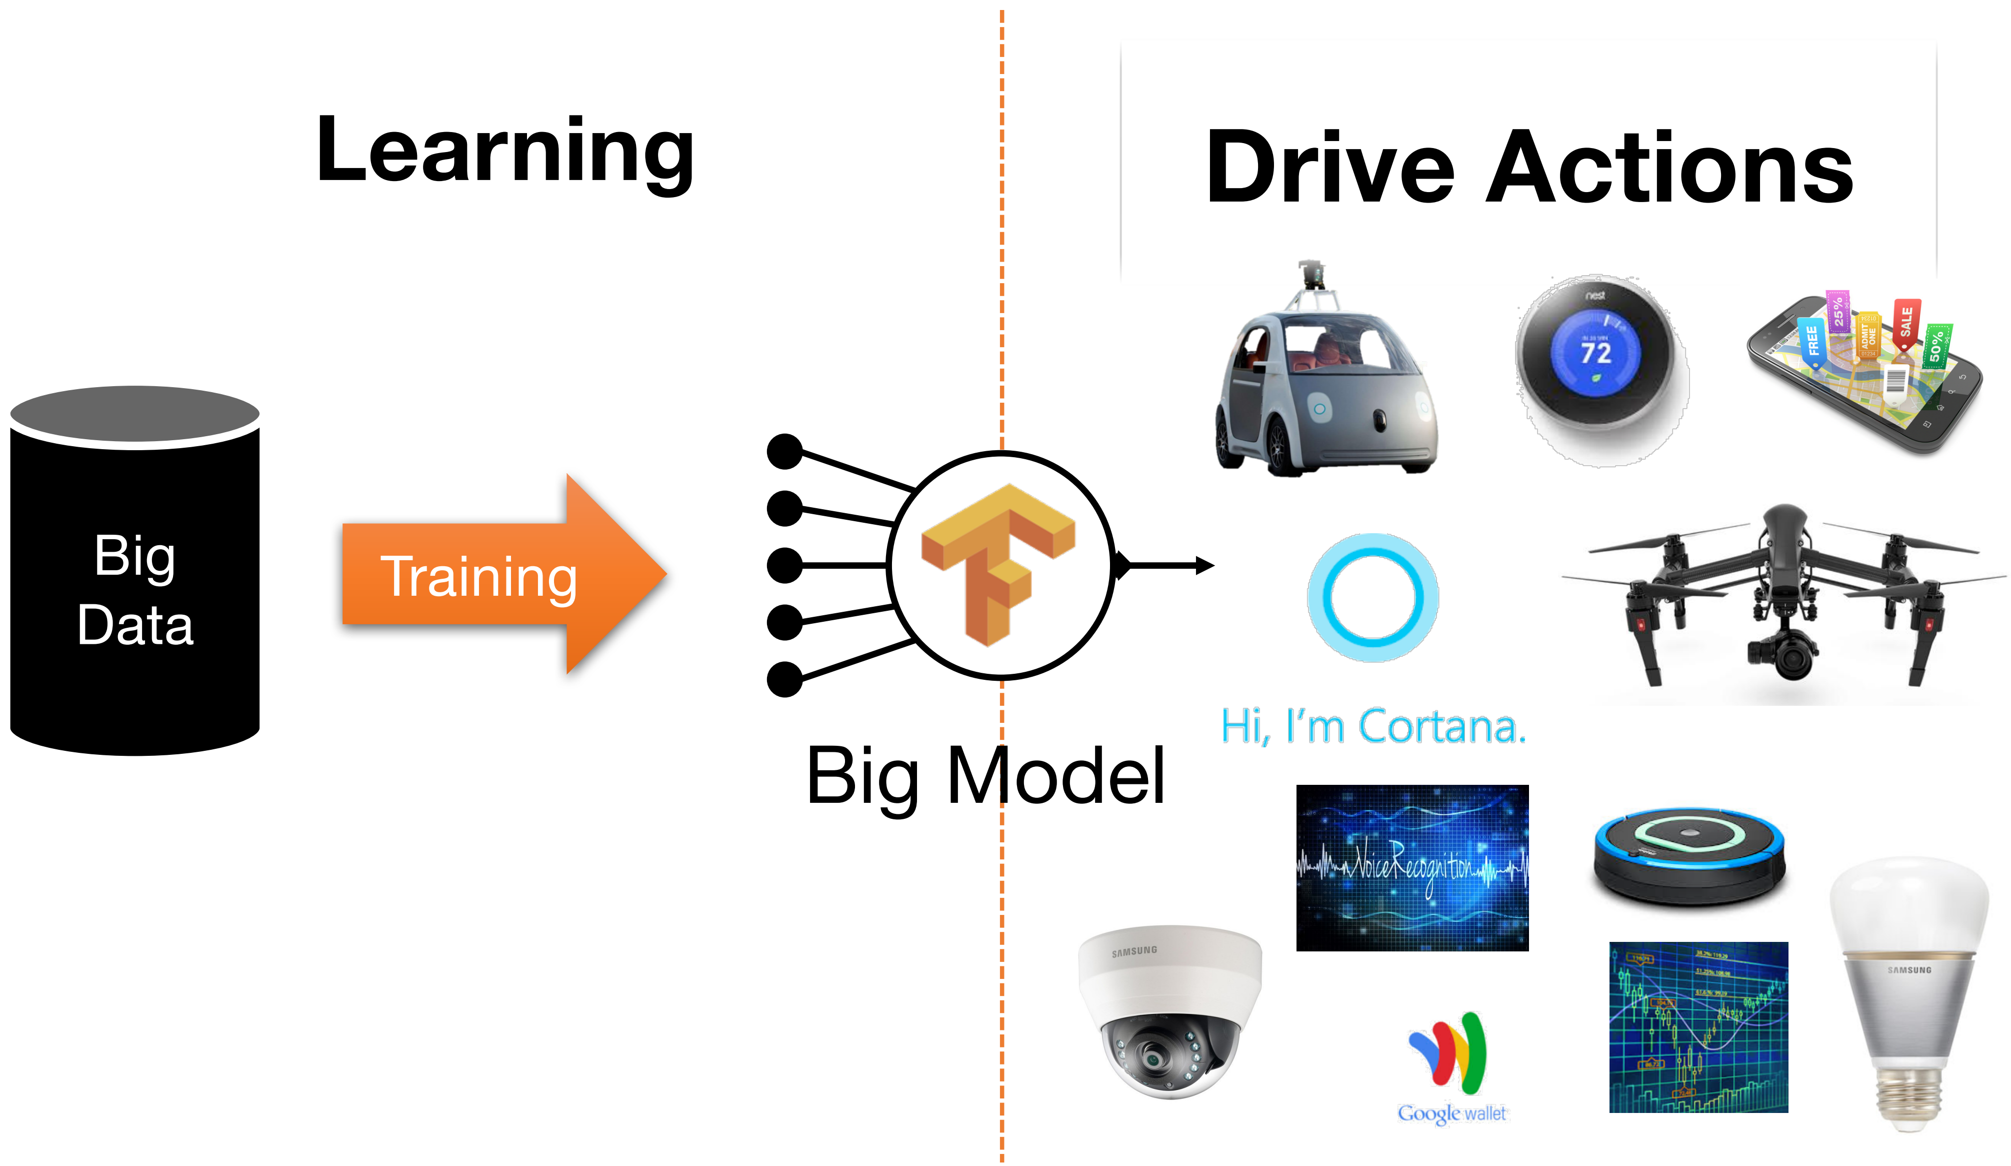
Learning
Drive Actions
Big Data
Training
Big Model
Hi, I'm Cortana.
Google wallet
SAMSUNG
SAMSUNG
nest
72
FREE
SALE
25%
50%

Learning

Inference

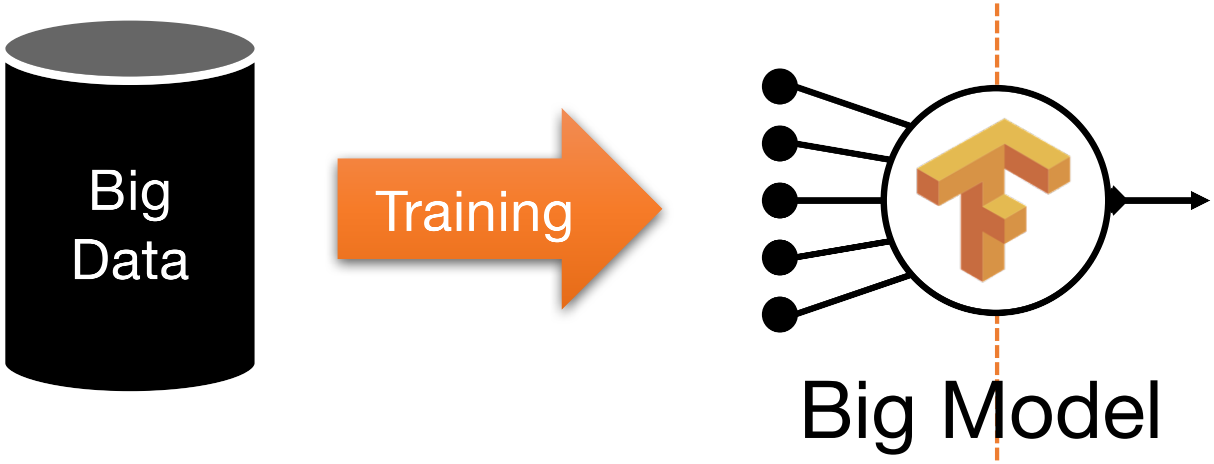

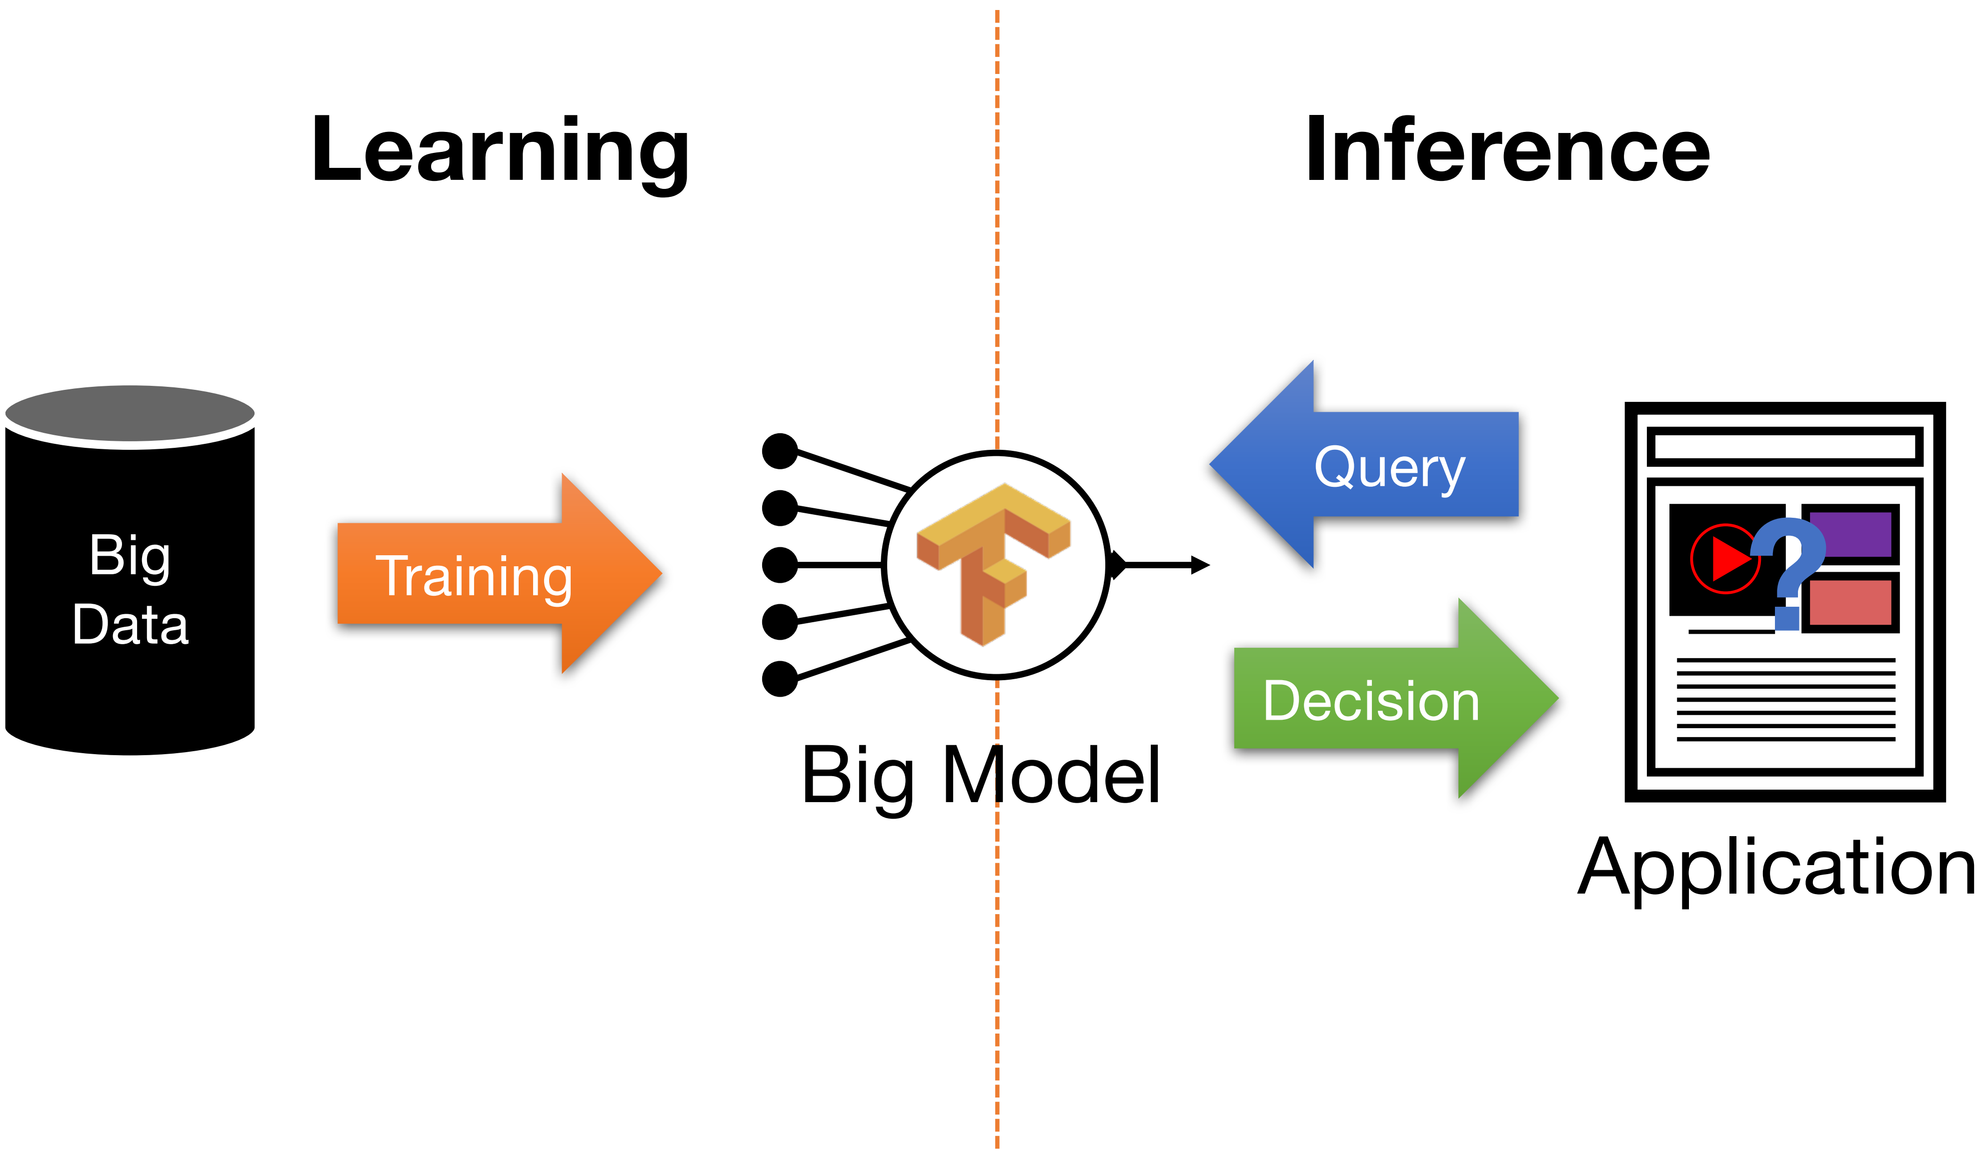
Learning
Inference
Big
Data
Training
Big Model
Query
Decision
Application

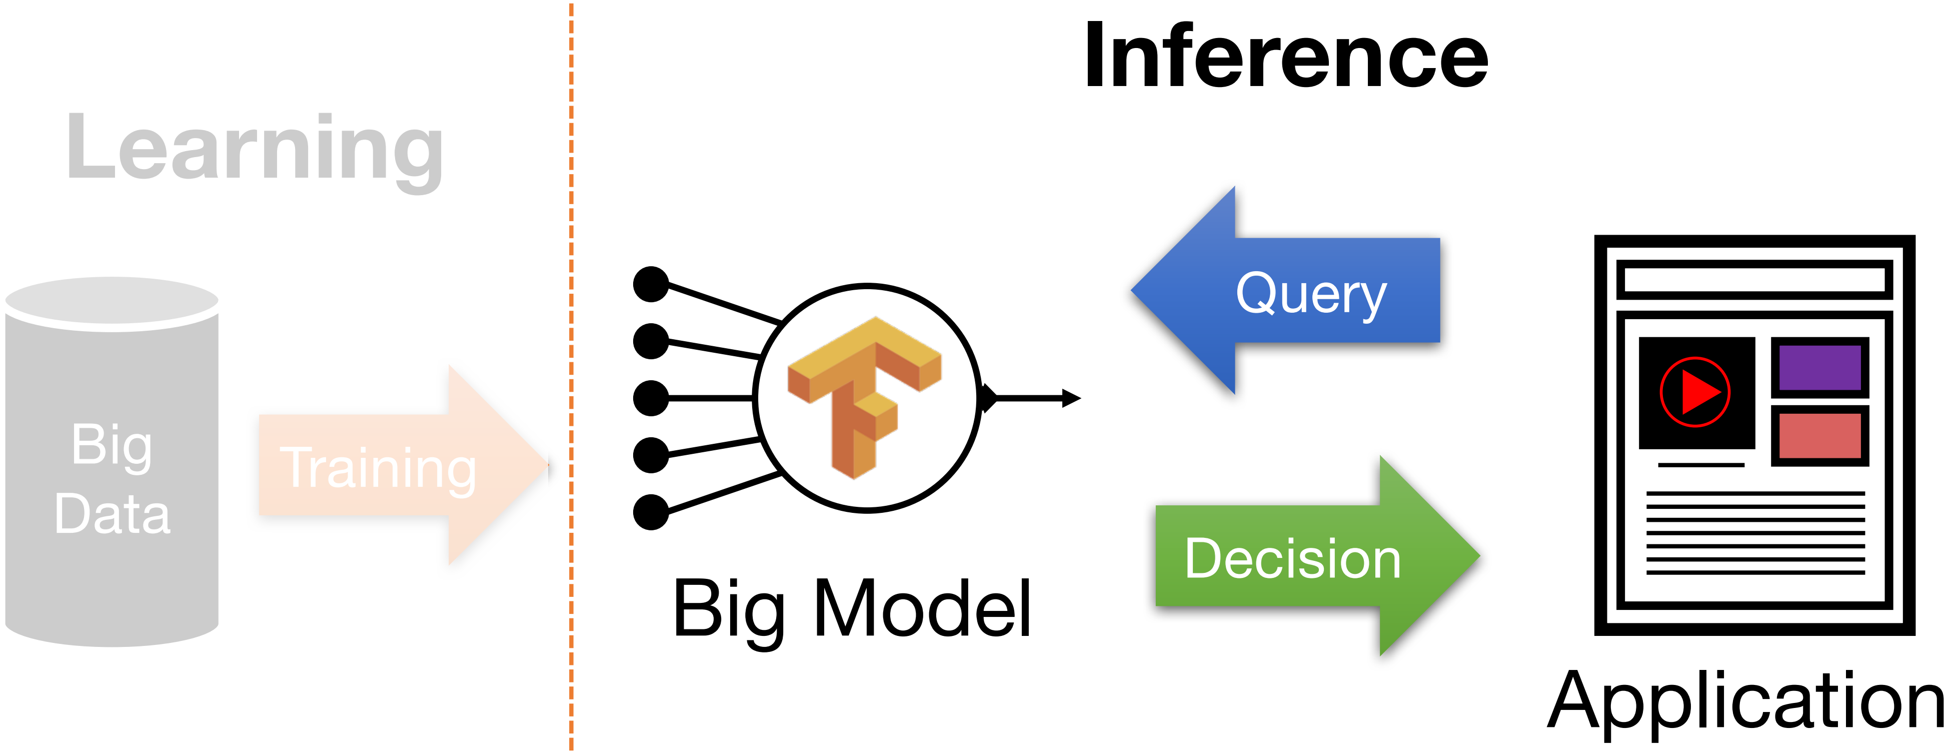

Often **overlooked**

Timescale: ~10 milliseconds

An area of focus in the RISELab

why is **Inference** challenging?

Need to render **low latency** (< 10ms) predictions for **complex**

Models

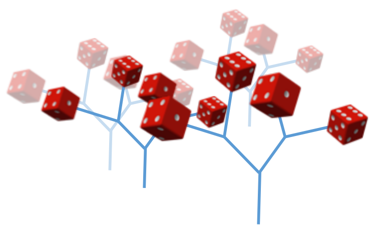

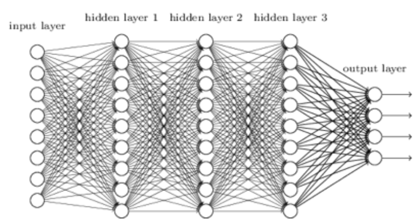

Queries

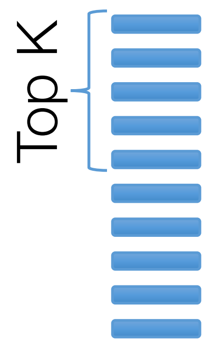

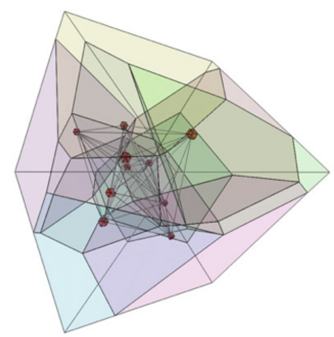

Features

SELECT * FROM users JOIN items, click_logs, pages WHERE …

under **heavy load** with system **failures**.

Inference is moving beyond the cloud

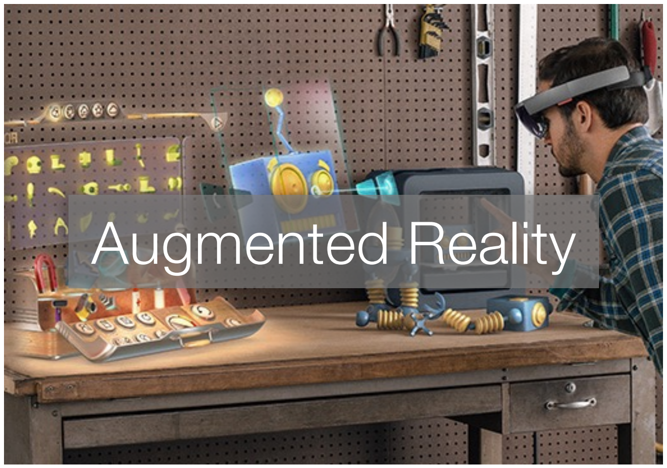

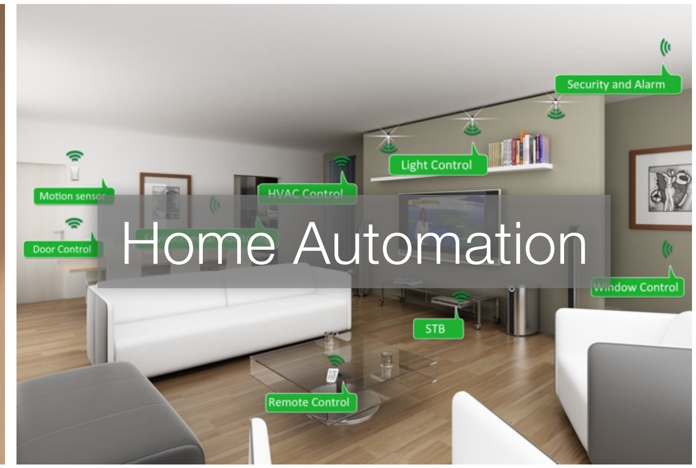

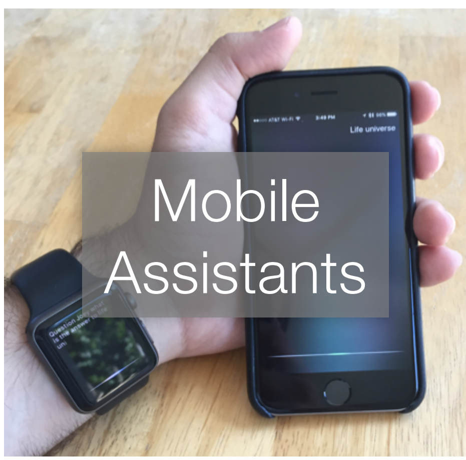

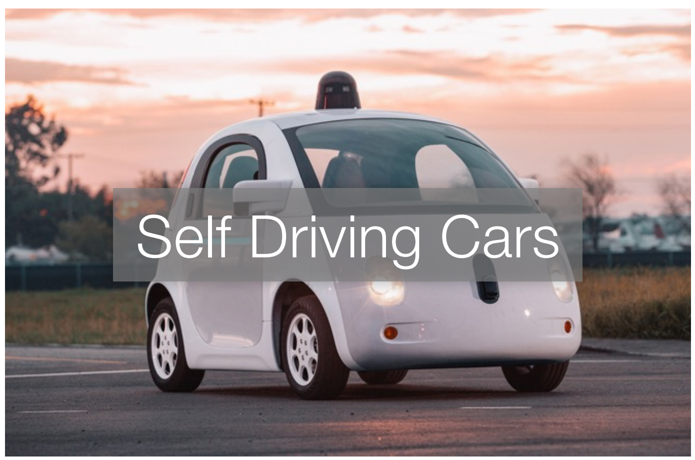

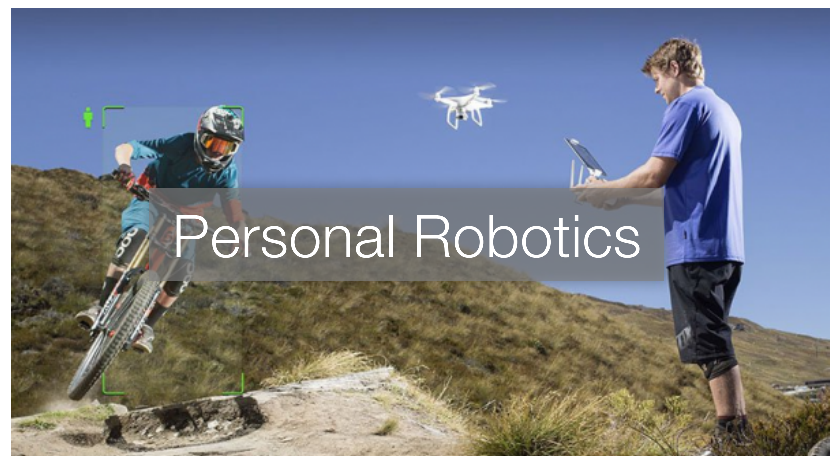

Inference is moving beyond the cloud

Opportunities

- Reduce latency and improve privacy
- Address network partitions

Research Challenges

- Minimize **power consumption**
- **Limited hardware** & long life-cycles
- Develop new **hybrid models** to leverage cloud and devices

<u>Robust</u> **Inference** is critical

Self “*Parking*” Cars

Self “*Driving*” Cars

Chat AIs

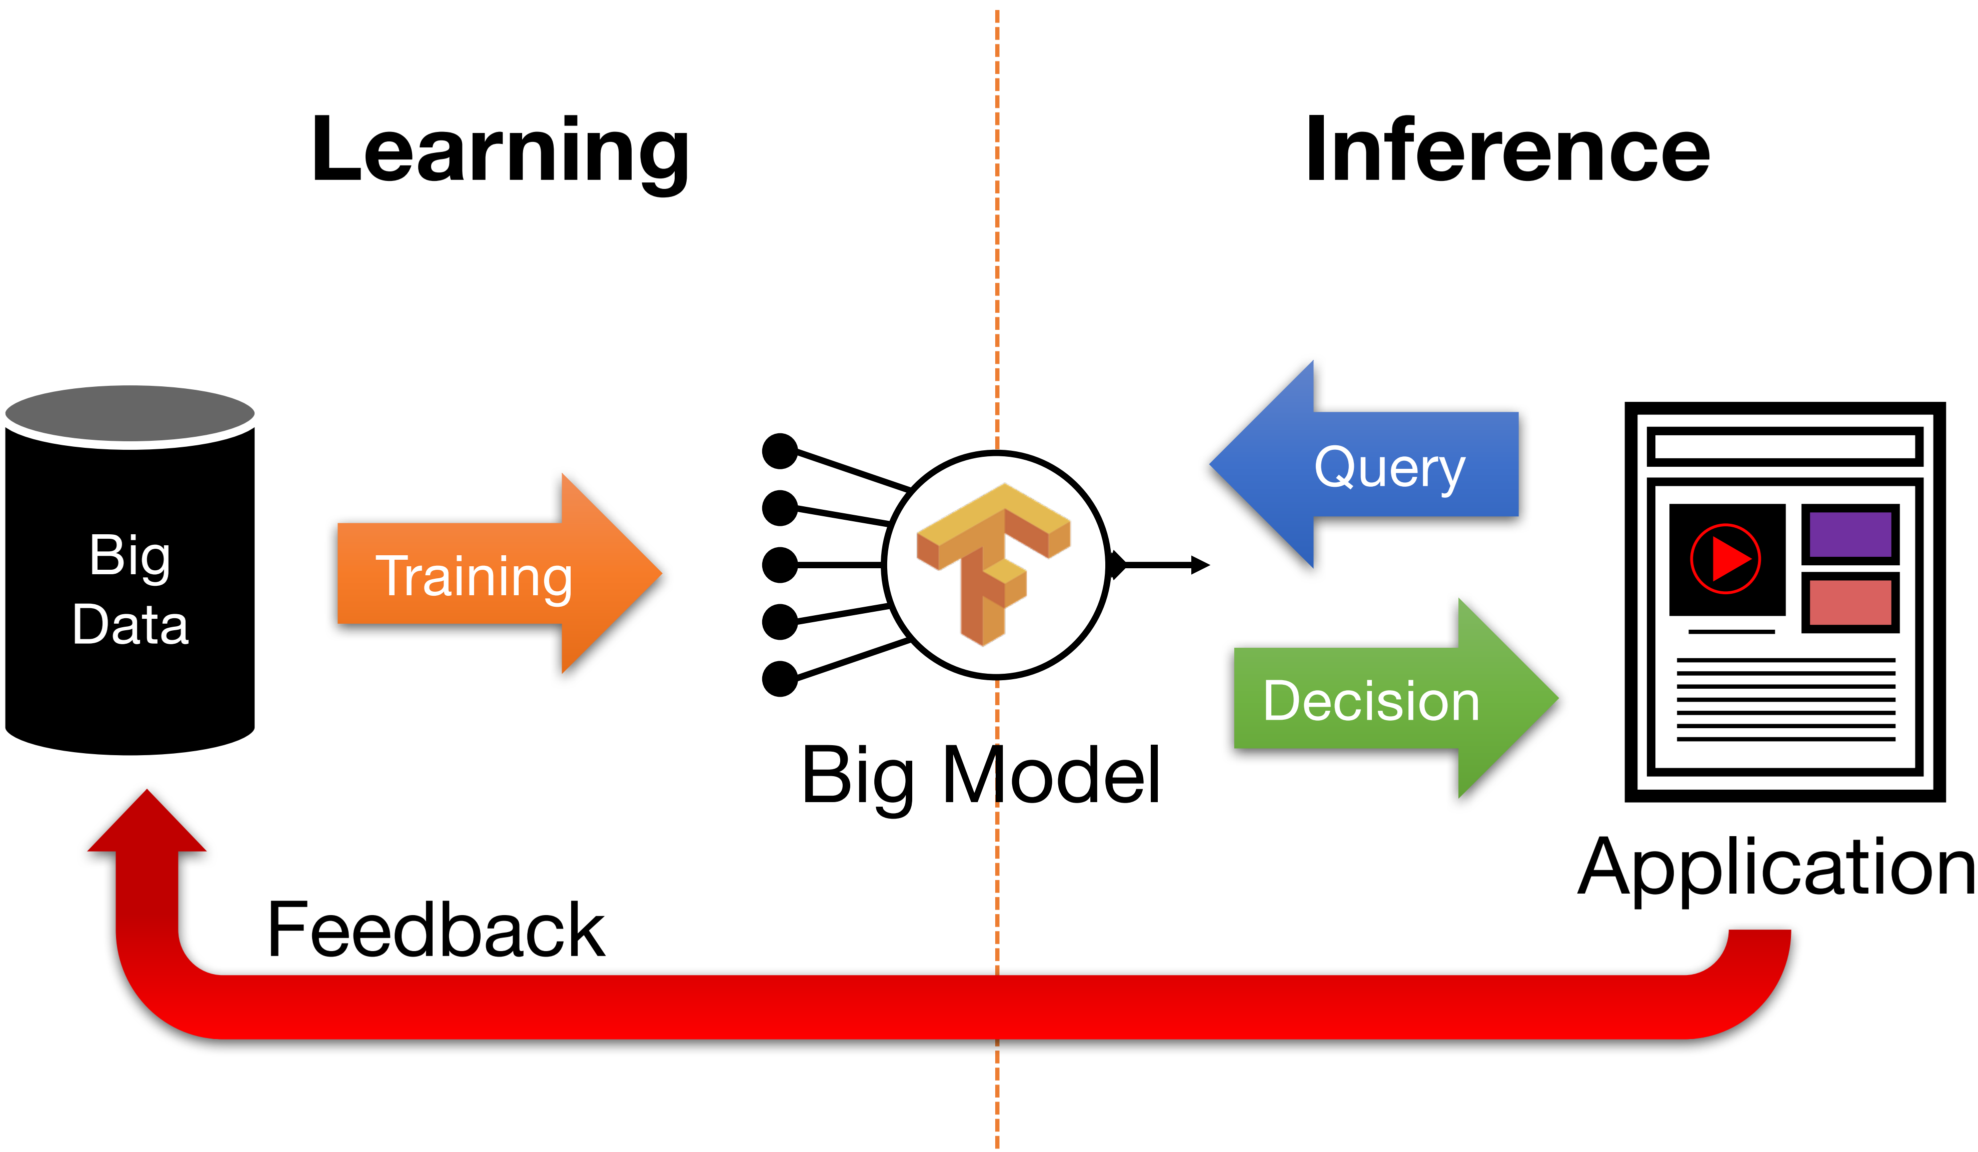
Learning
Inference
Big
Data
Training
Query
Decision
Big Model
Application
Feedback

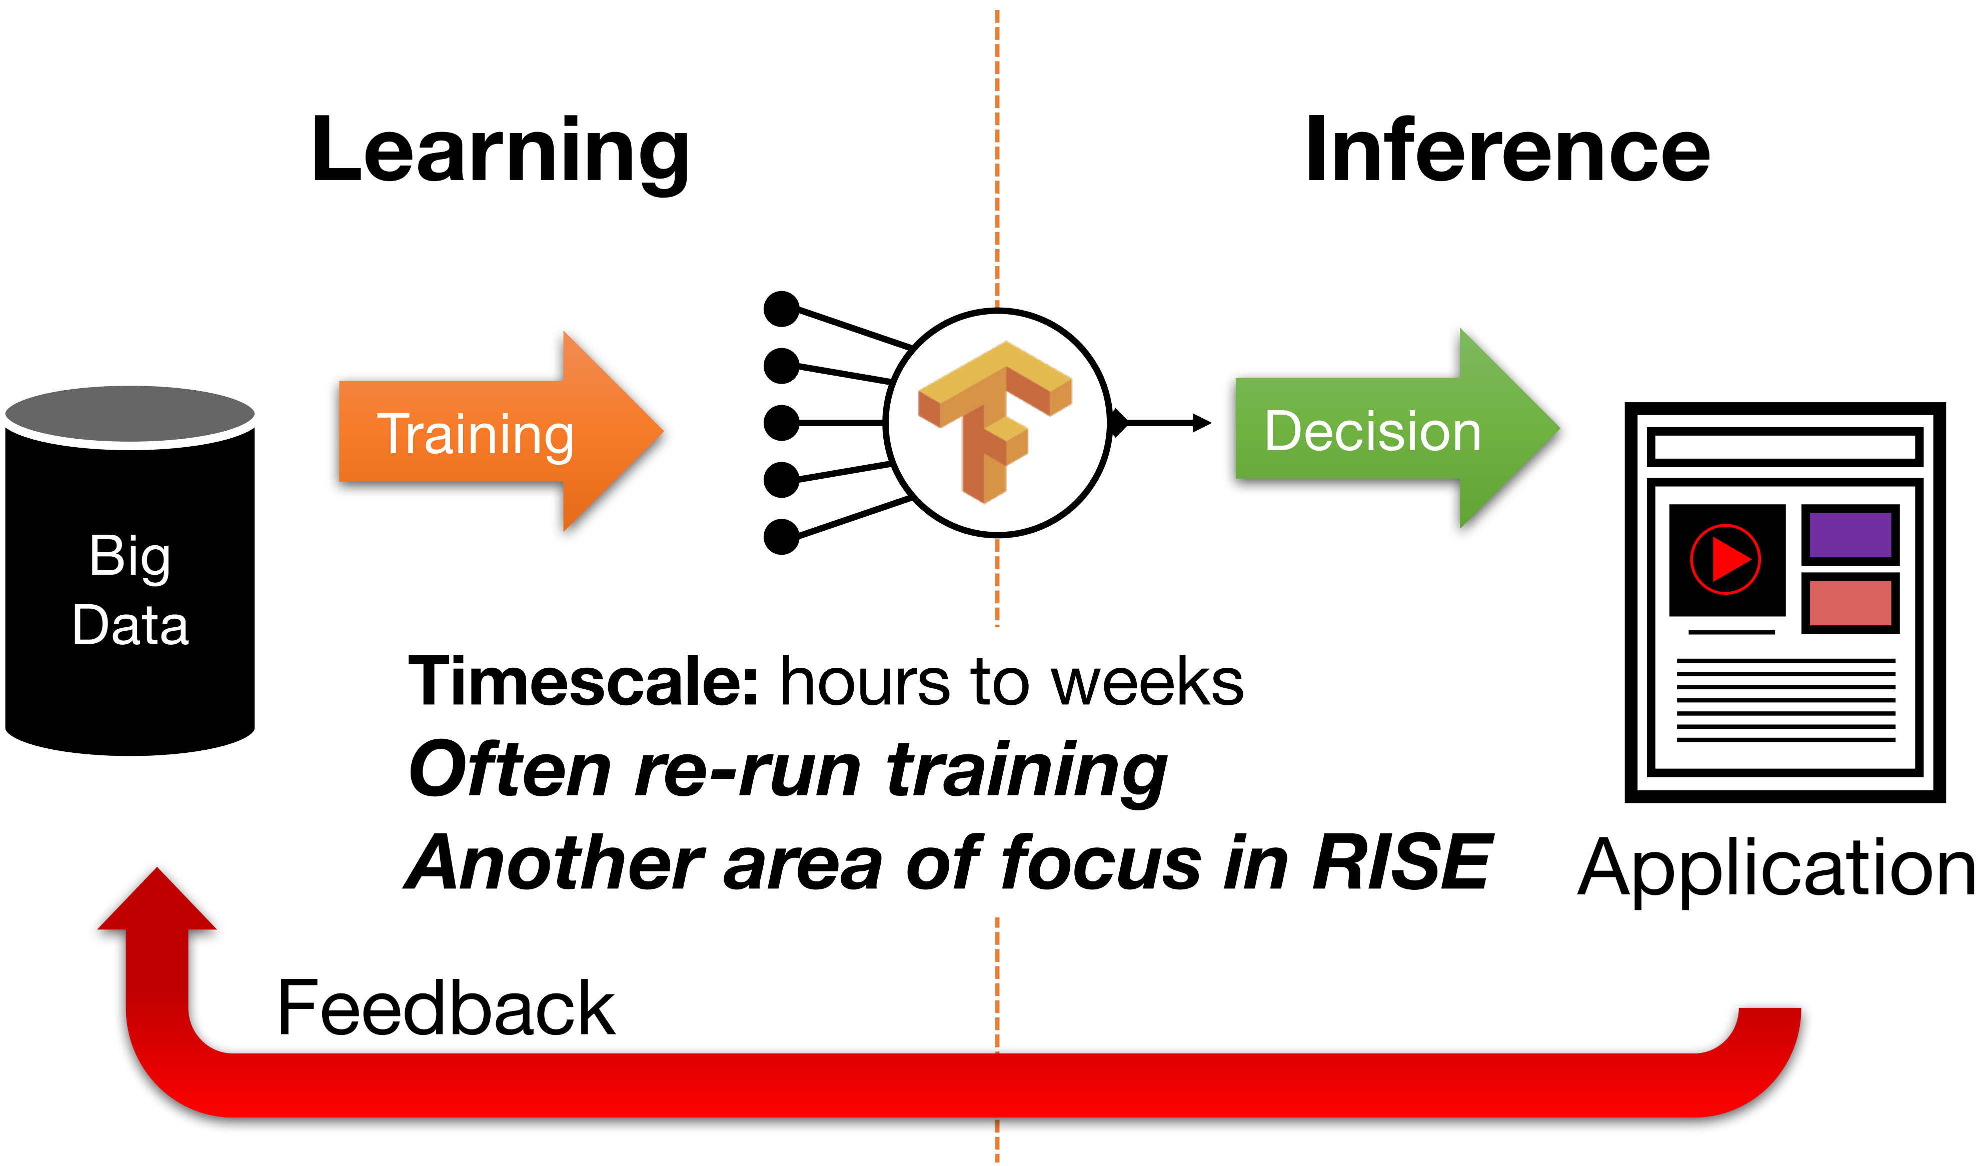

Learning
Inference
Big Data
Training
Decision
Application
Timescale: hours to weeks
Often re-run training
Another area of focus in RISE
Feedback

Why is **Closing the Loop** challenging?

Implicit and **Delayed** Feedback

Self Reinforcing **Feedback Loops**

$$\frac{d}{dt}\left[\quad\right]$$

World Changes **at varying rates**

Learning

Inference

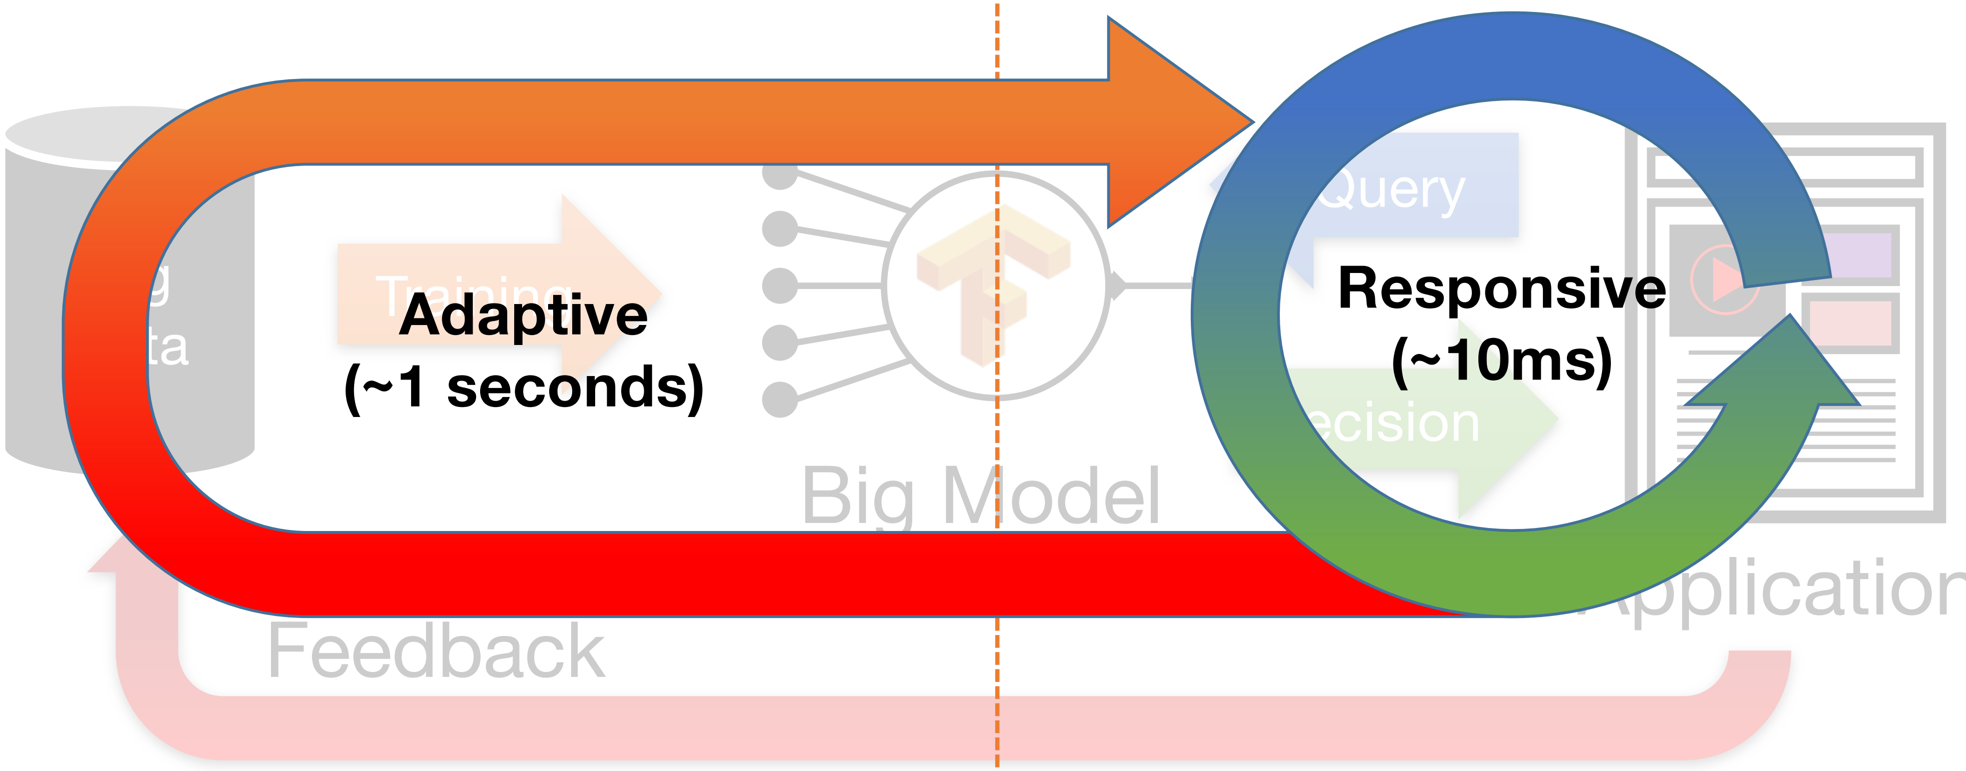

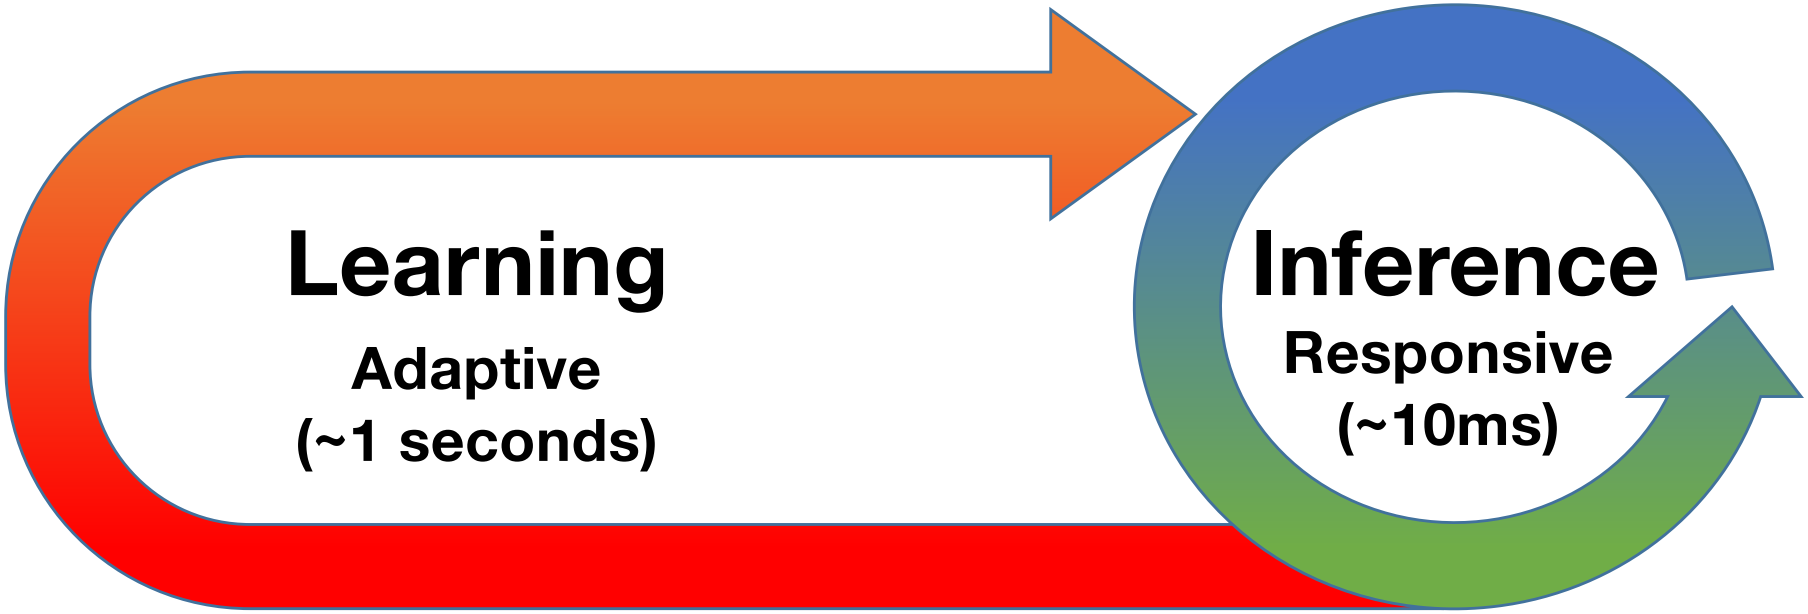

Secure

Intelligence in **Sensitive Contexts**

AR/VR Systems

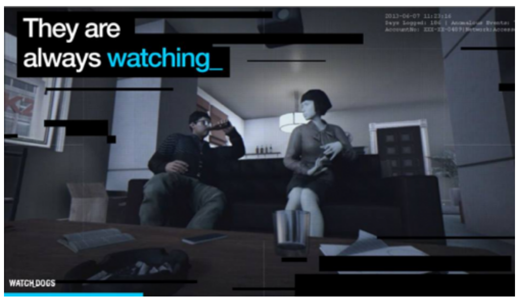

Home Monitoring

Voice Technologies

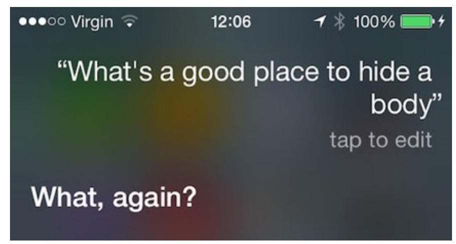

Medical Imaging

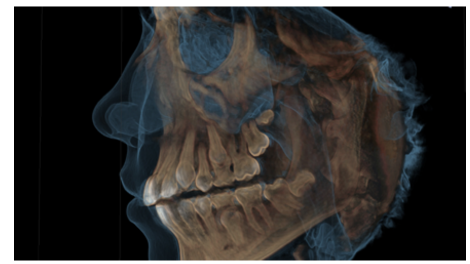

Protect the **data**, the **model**, and the **query**

Protect the **data**, the **model**, and the **query**

High-Value **Data** *is* ***Sensitive***

- Medical Info.
- Home video
- Finance

Models capture **value** in data

- Core Asset
- Sensitive

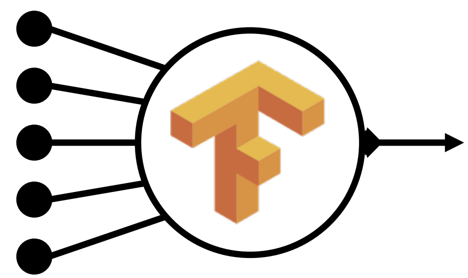

Queries can be as sensitive as the data

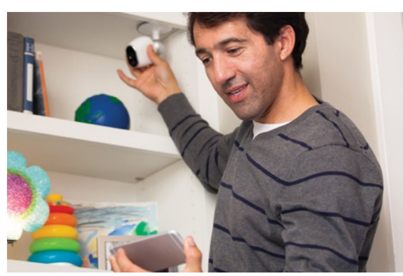

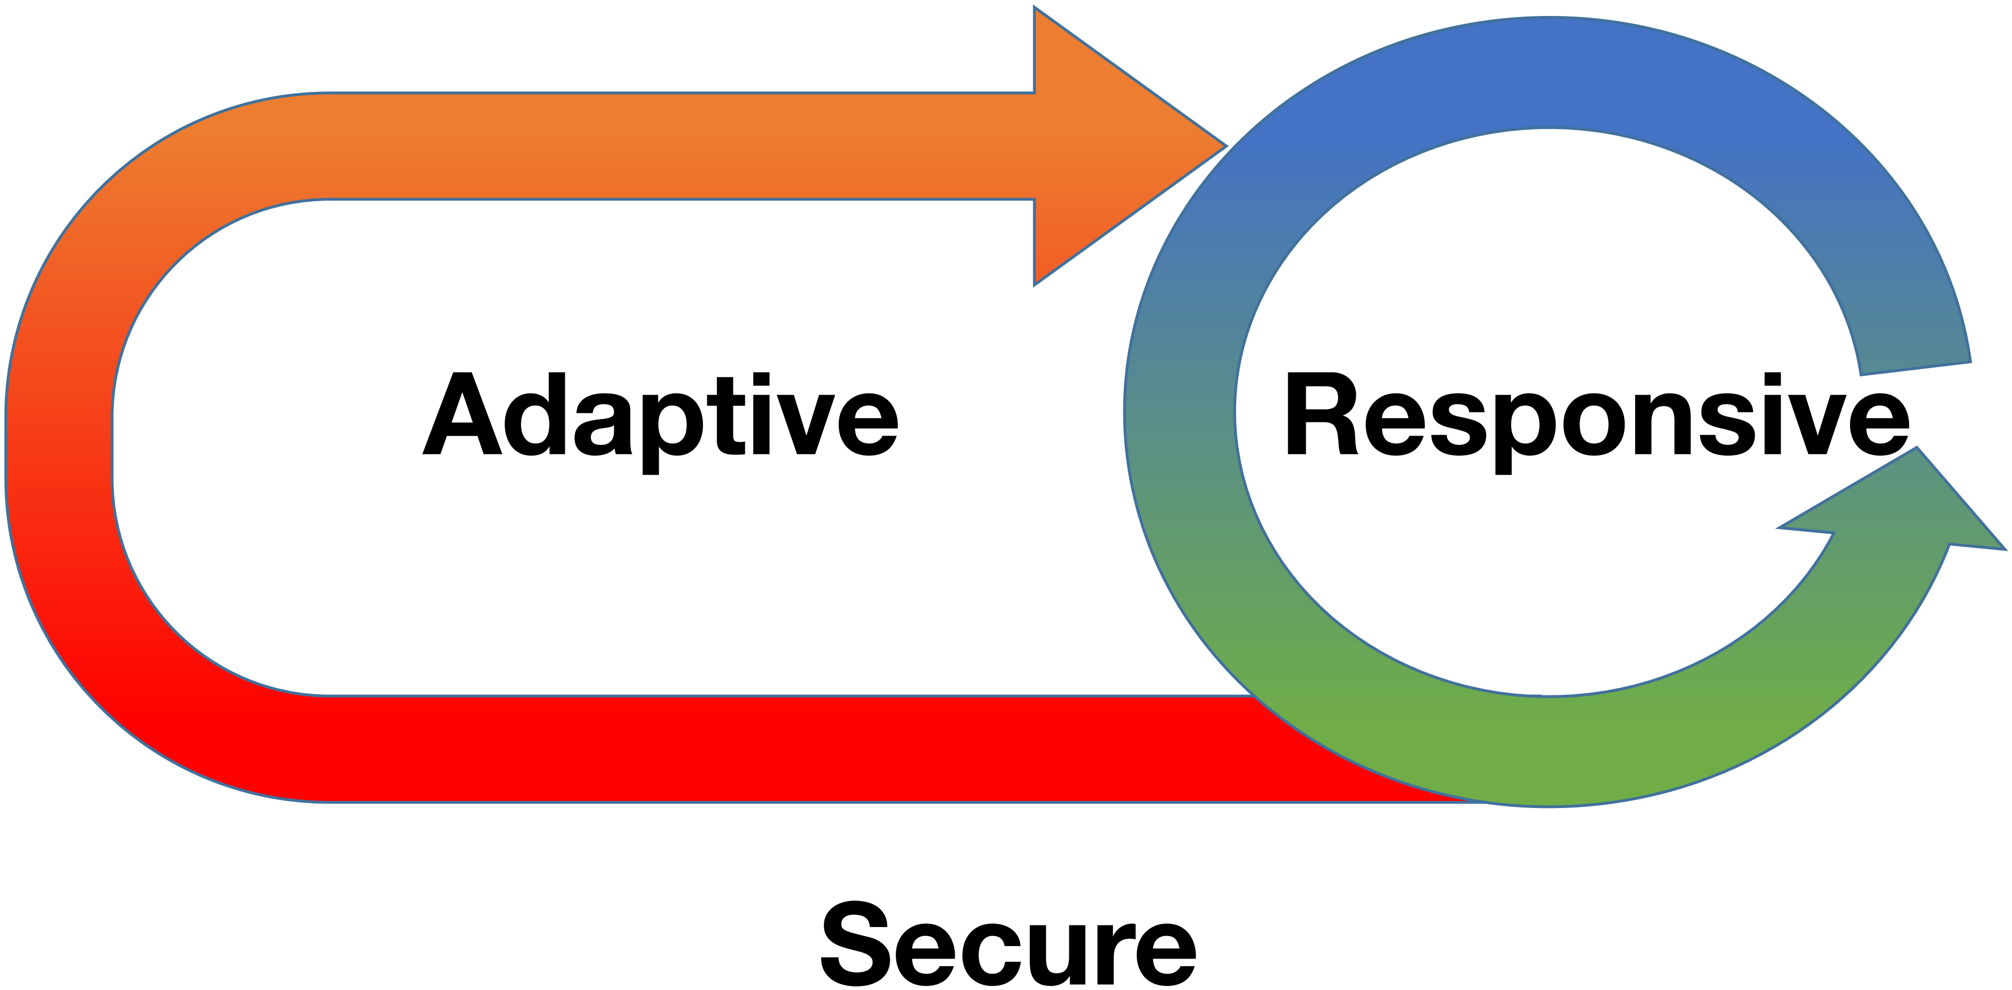
Adaptive
Responsive
Secure

Goal

Real-time decisions

decide in ms

on live data

with strong security

Goal

Real-time decisions

decide in ms

on live data

the current state of the environment

with strong security

Goal

Real-time decisions

decide in ms

on live data

the current state of the environment

with strong security

privacy, confidentiality, integrity

Challenges	RISE Lab
Automated decisions on live data are hard	Real-time, sophisticated decisions that guarantee worst-case behavior on noisy and unforseen live data
Poor security: exploits are daily occurrences	Ensure privacy and integrity without impacting functionality
One-off solutions, expensive and slow to build	General platform: **Secure Real-time Decision Stack**

Example: Zero-time defense

Problem: zero-day attacks can compromise millions of sites in seconds

Solution: analyze network flows to detect attacks and patch sites/software in real-time

- **Intermediate data**: create attack model
- **Decision**: detect attack, patch

Quality	sophisticated, accurate, robust
Latency	sec (decision) / sec (update)
Security	privacy (encourage users to share logs), integrity

Example: “Fleet” driving

Problem: suboptimal driving decisions
Solution: collect & leverage info from other cars and drivers in real-time

- **Intermediate data**: automatically annotate maps, actions of other drivers
- **Decision**: avoid obstacles, congestions

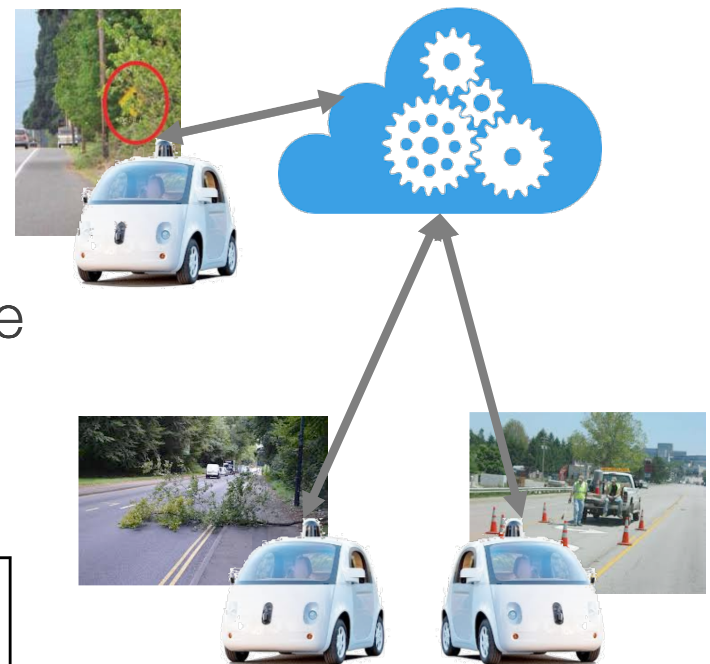

Quality	sophisticated, accurate, noise tolerant
Performance	sec (decision) / sec (update)
Security	privacy, data integrity

Example: Infectious disease discovery

Problem: infectious diseases spread quickly (Zika), may need quarantine (Ebola)

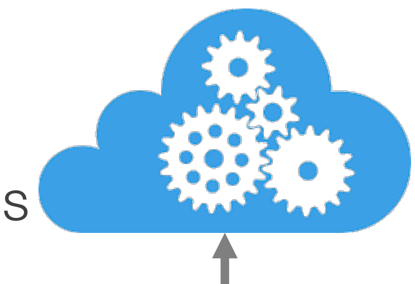

Solutions: real-time DNA seq. & analysis to identify pathogens

- Rapid analysis to trace evolution, source
- 100x faster → 100x people tested
  - **Intermediate data**: evolution, spread, symptoms
  - **Decision**: quarantine or not, diagnosis

MinION Nanopore
(Dr. Charles Chiu UCSF using it to identify Zika virus)

Quality	sophisticated, accurate
Latency	min (decision) / hour (update)
Security	privacy, integrity

Applications	Quality	Latency		Security
		Decision	Update	
Zero-time defense	sophisticated, accurate, robust	sec	sec	privacy/confidentiality, integrity
Parking assistant	sophisticated, robust	sec	sec	
Disease discovery	sophisticated, accurate	sec/min	hours	
IoT (smart buildings)	sophisticated, robust	sec	min/hour	
Earthquake warning	sophisticated, accurate, robust	ms	min	
Chip manufacturing	sophisticated, accurate, robust	sec/min	min	
Fraud detection	sophisticated, accurate	ms	min	
“Fleet” driving	sophisticated, accurate, robust	sec	sec	
Virtual companion	sophisticated, robust	sec	min/hour	
Video QoS at scale	sophisticated	ms/sec	min	

Challenges	RISE Lab
Automated decisions on live data are hard	Real-time, sophisticated decisions that guarantee worst-case behavior on noisy and unforseen live data
Poor security: exploits are daily occurrences	Ensure privacy and integrity without impacting functionality
One-off solutions, expensive and slow to build	General platform: **Secure Real-time Decision Stack**

A bird’s eye view

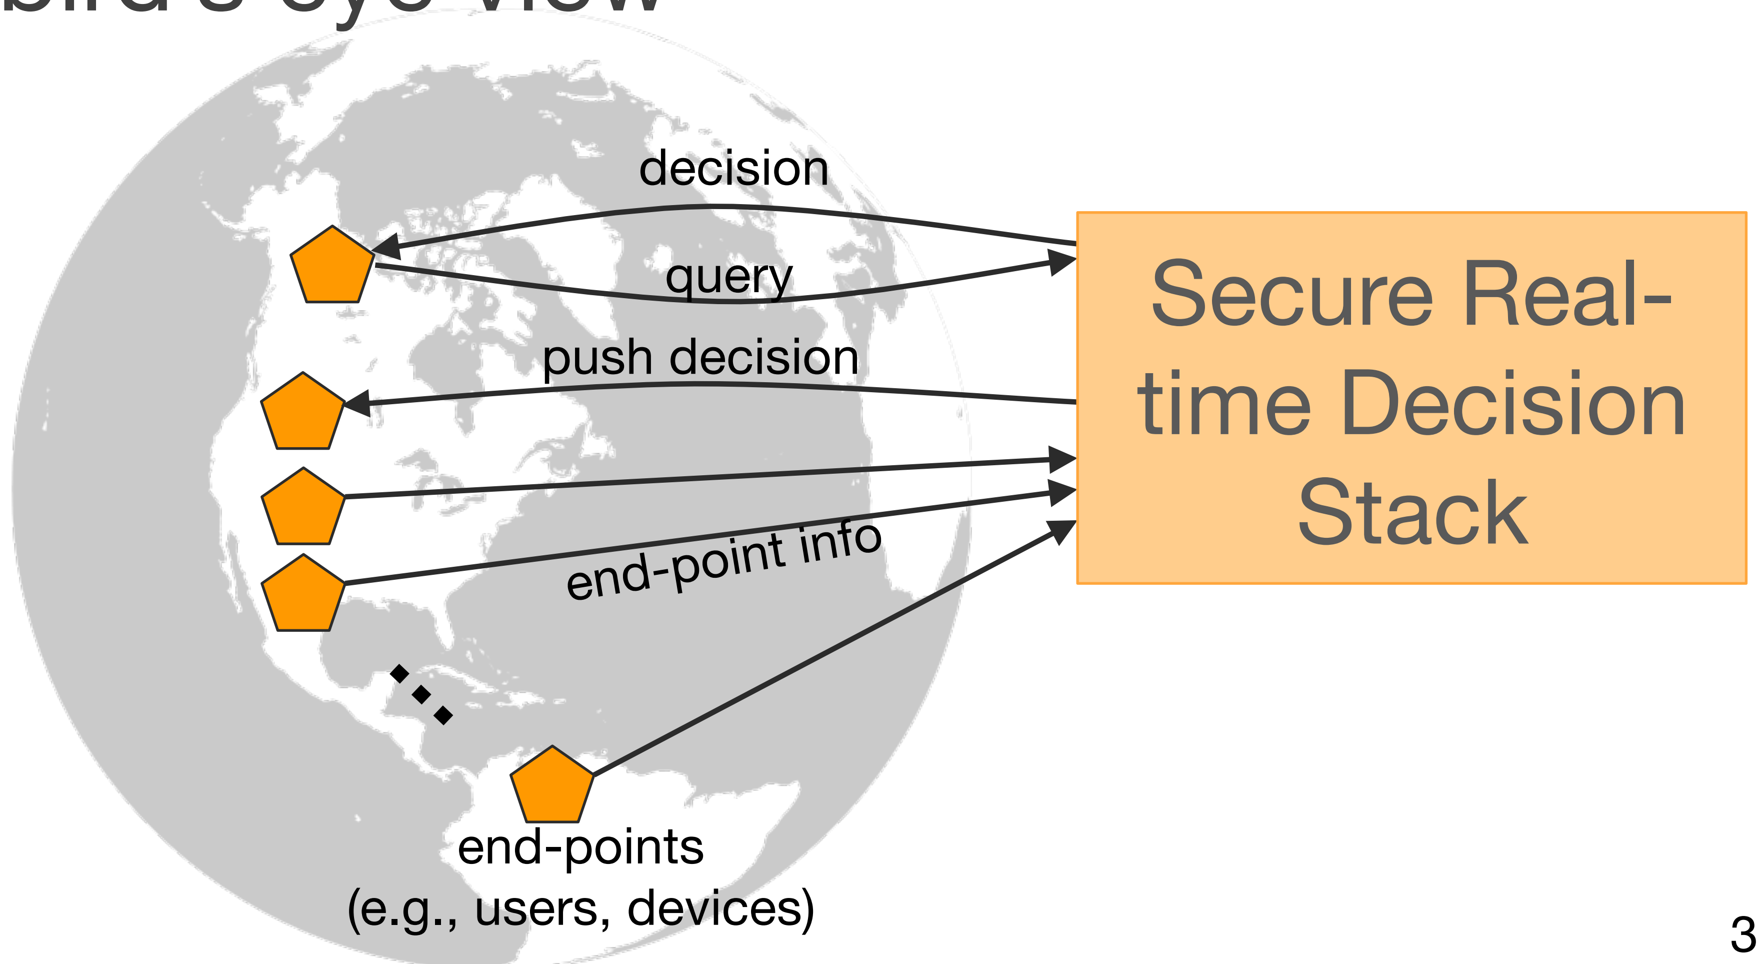

What exists today?

Pull decisions

Example: recommendation system

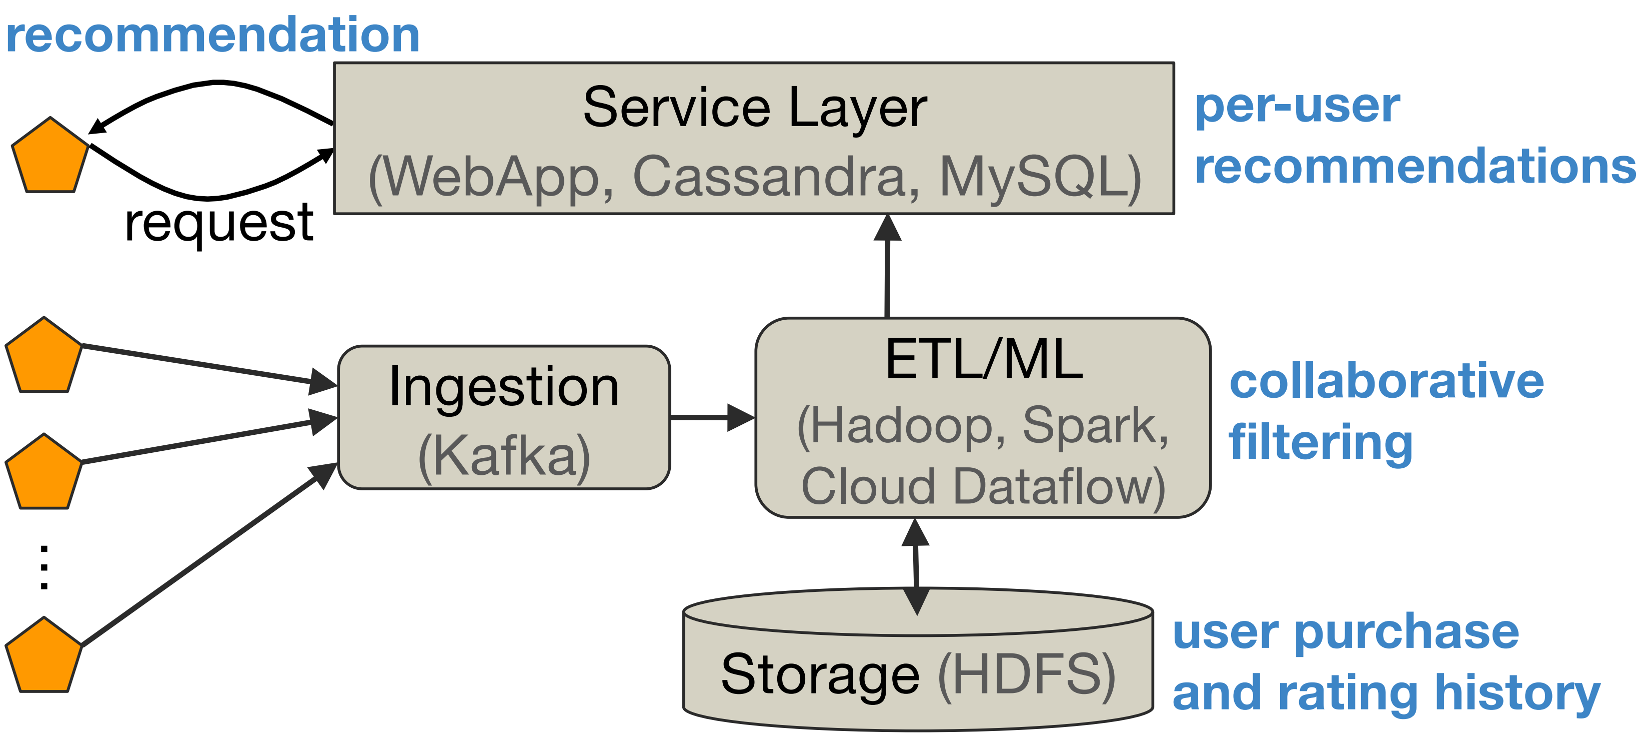

Pull decisions

Problem: cannot have sophisticated decisions

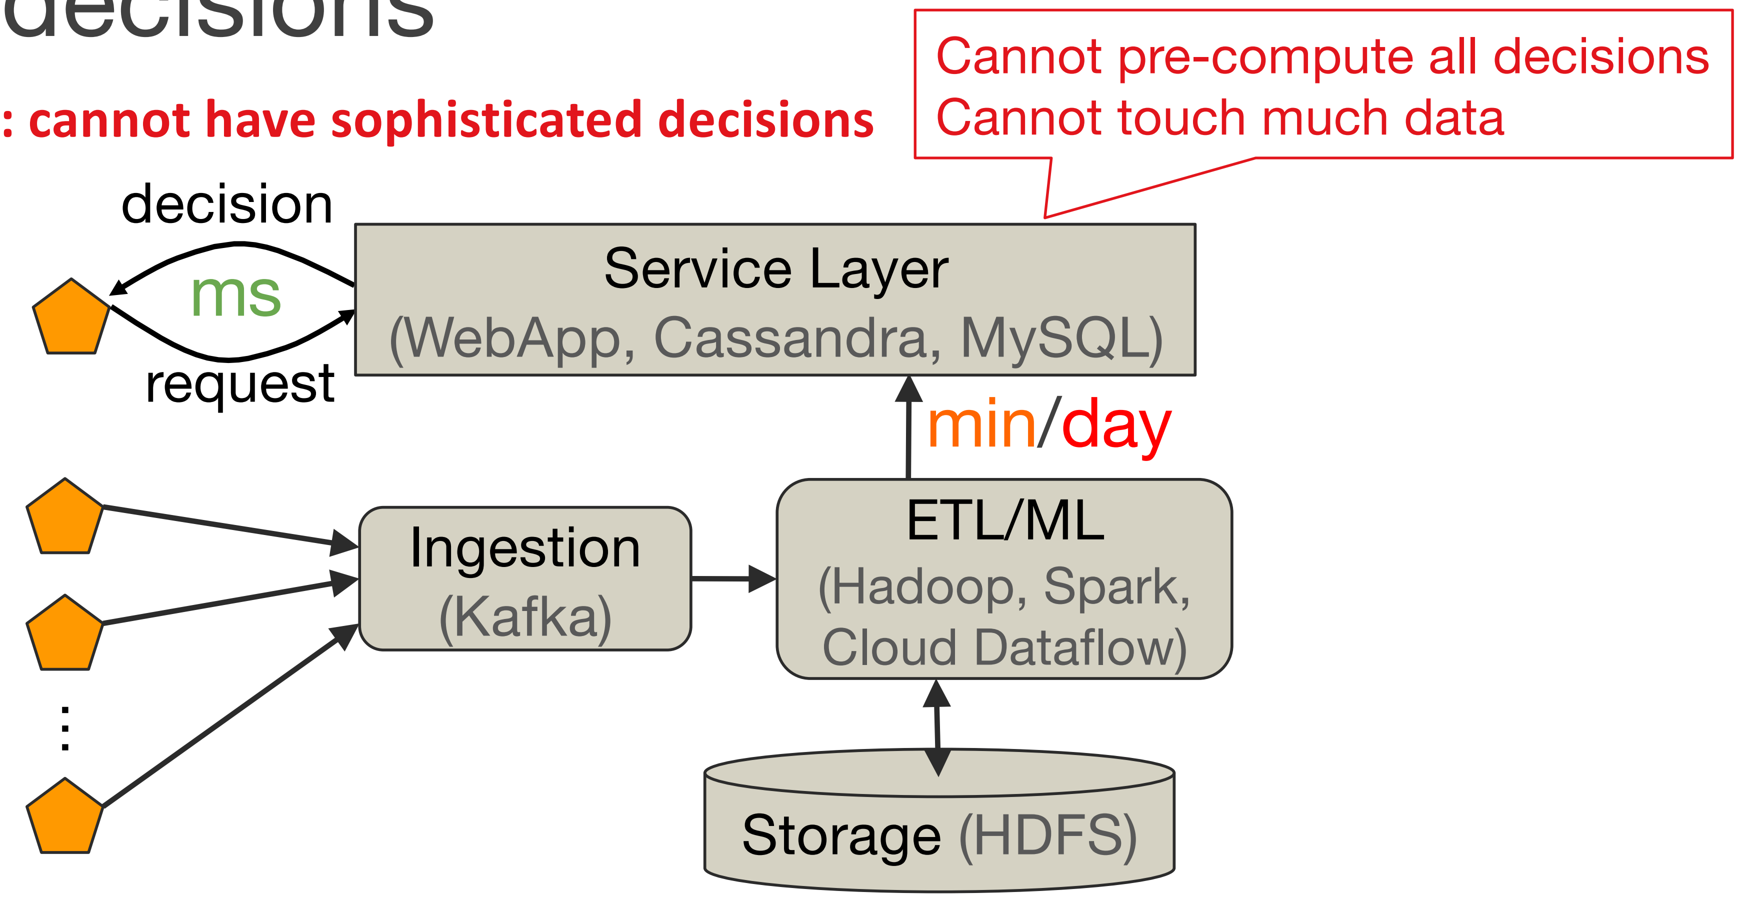

Pull decisions

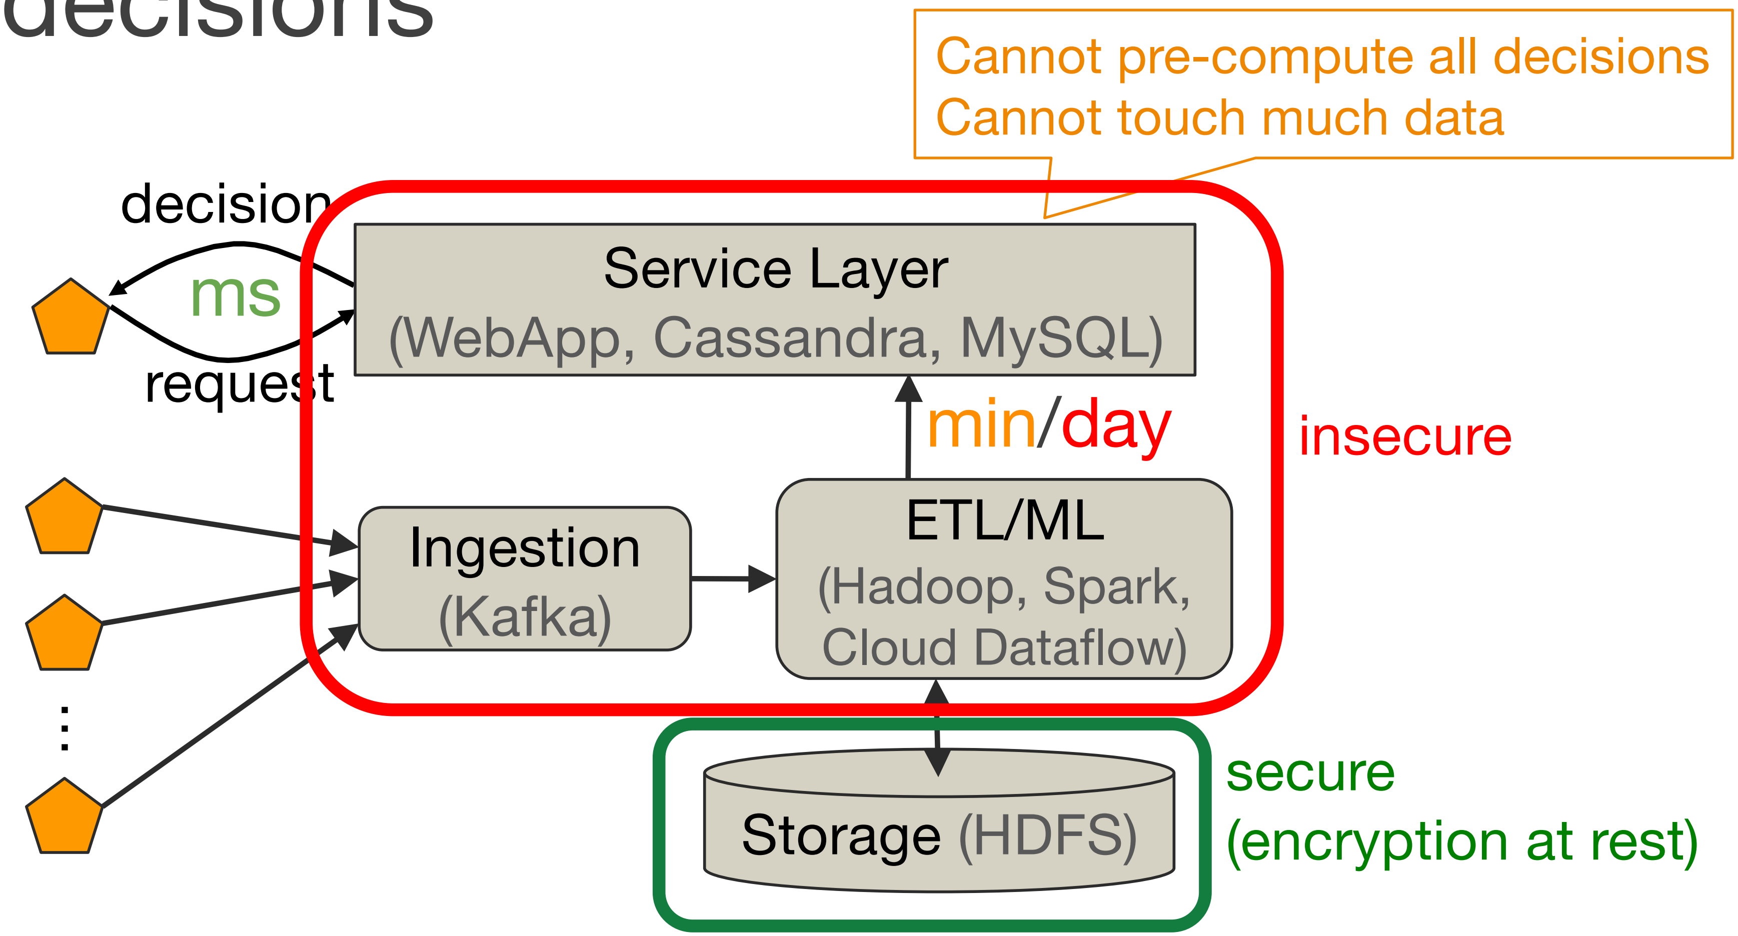

Solution scorecard

Solution	Decision Quality	Latency		Security
		Decision	Update	
Pull decisions	simple	ms	min/day	weak

Pull decisions: contextual decisions

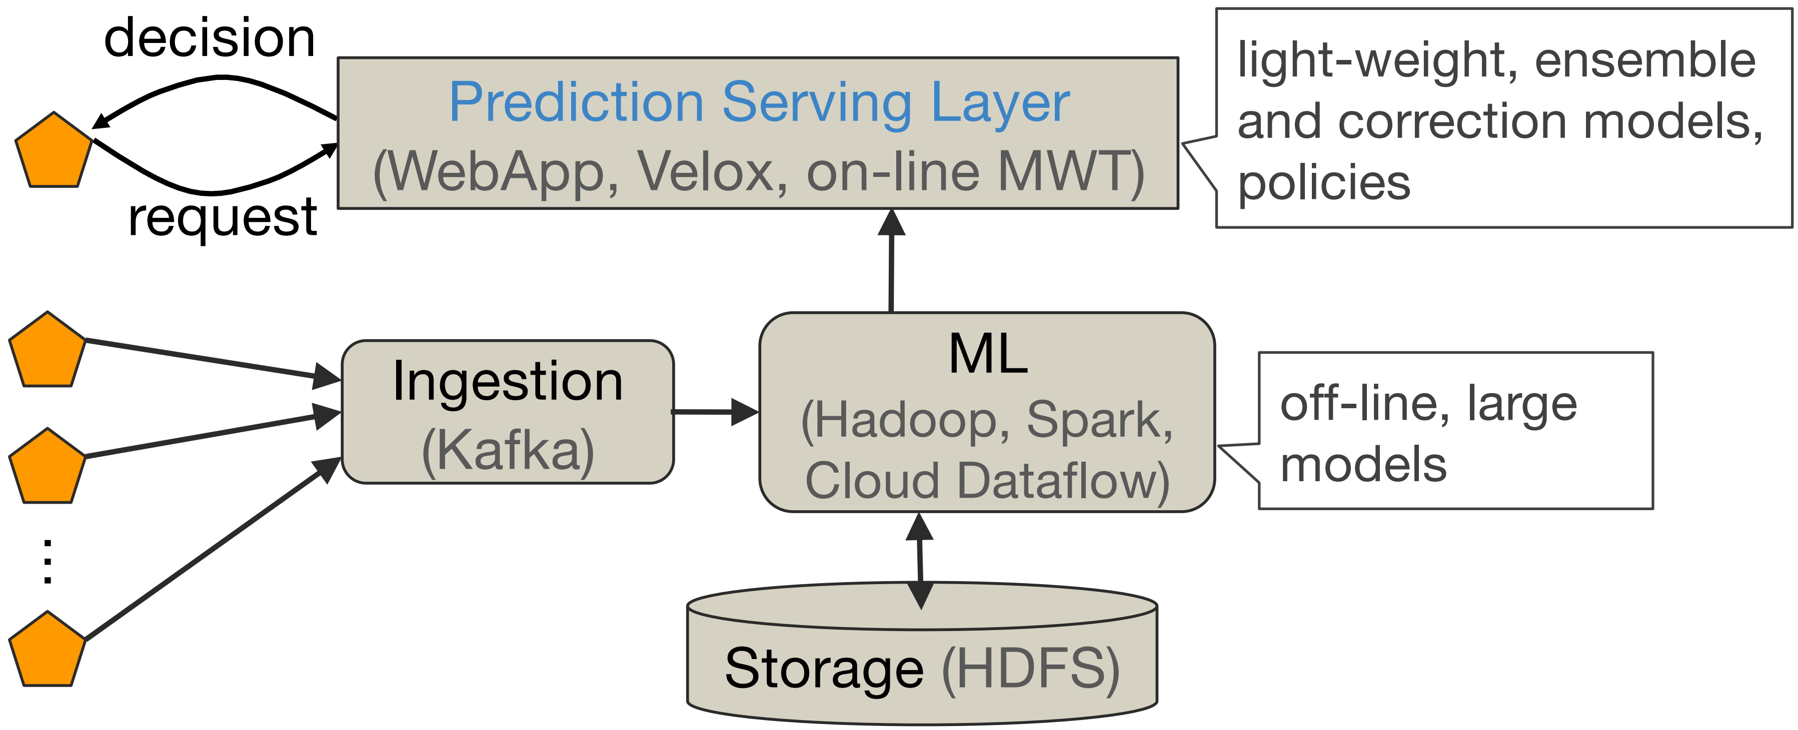

Solution scorecard

Solution	Decision Quality	Latency		Security
		Decision	Update	
Pull decisions	simple	ms	min/day	weak
Pull decisions: prediction service	sophisticated, specialized	ms	min*	weak

*light-weight, ensemble + correction models, policies

Push Decisions

Example: anomaly detection

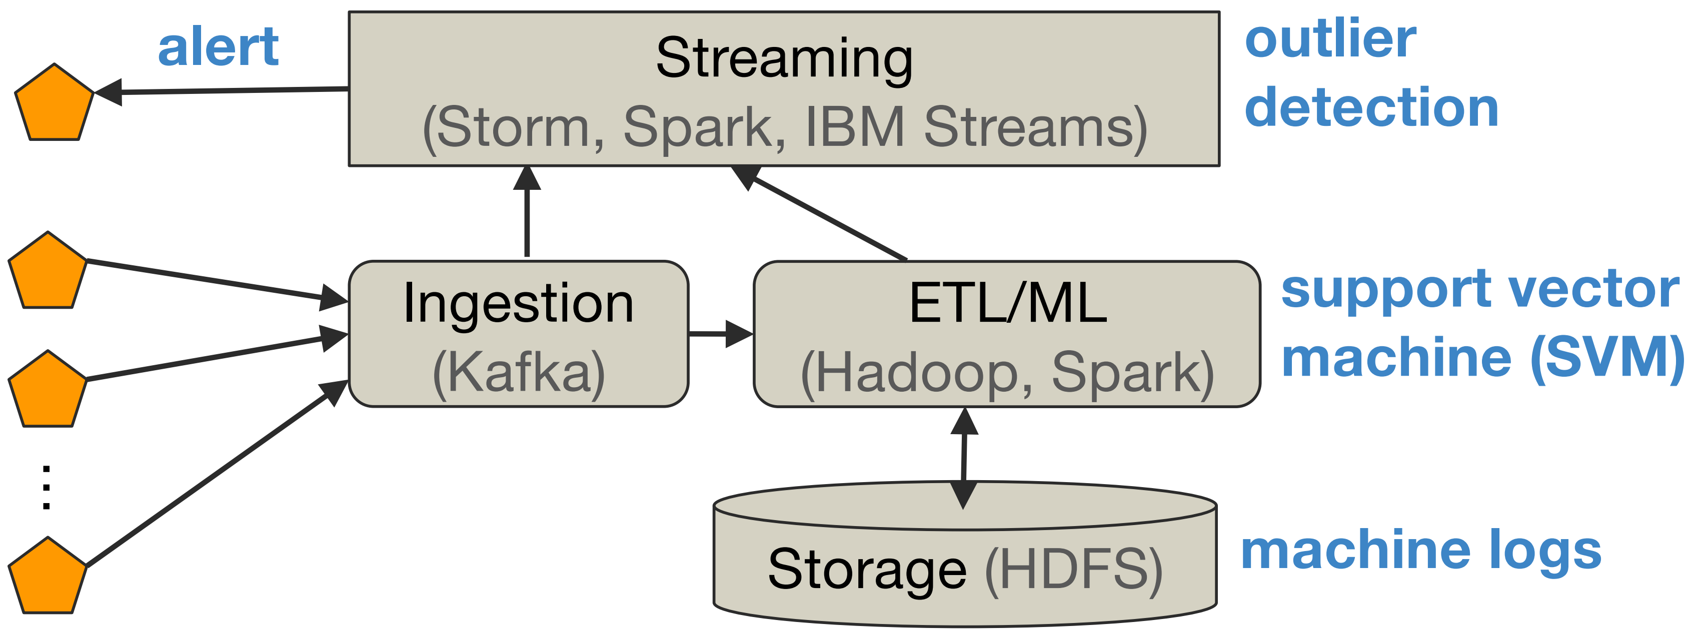

Push Decisions

Example: anomaly detection

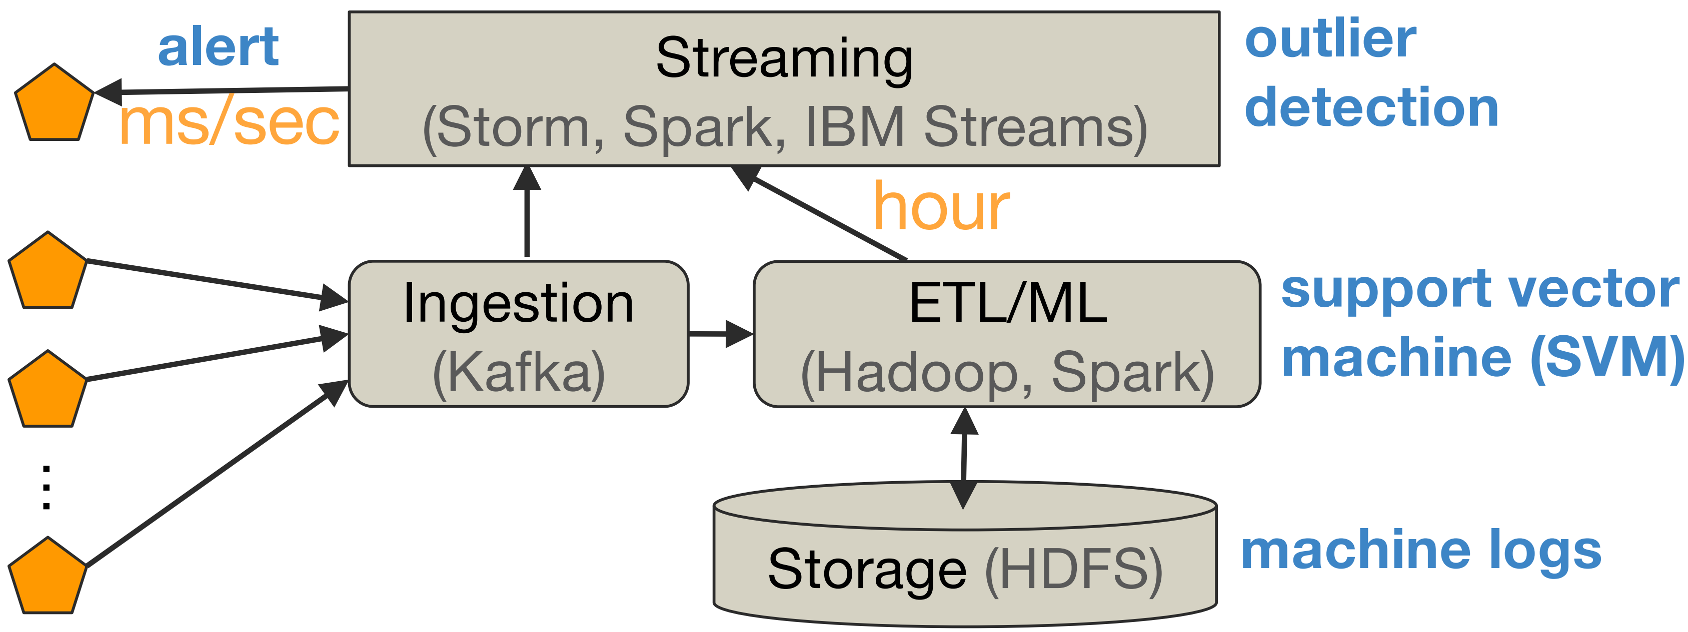

Solution scorecard

Solution	Decision Quality	Latency		Security
		Decision	Update	
Pull decisions	simple	ms	min/day	weak
Push decisions	simple	ms/sec	hour	weak

*light-weight, ensemble + correction models, policies

Security tools

- Computation on encrypted data
- Hardware enclaves

State-of-the-art security solutions:
Computation on encrypted data

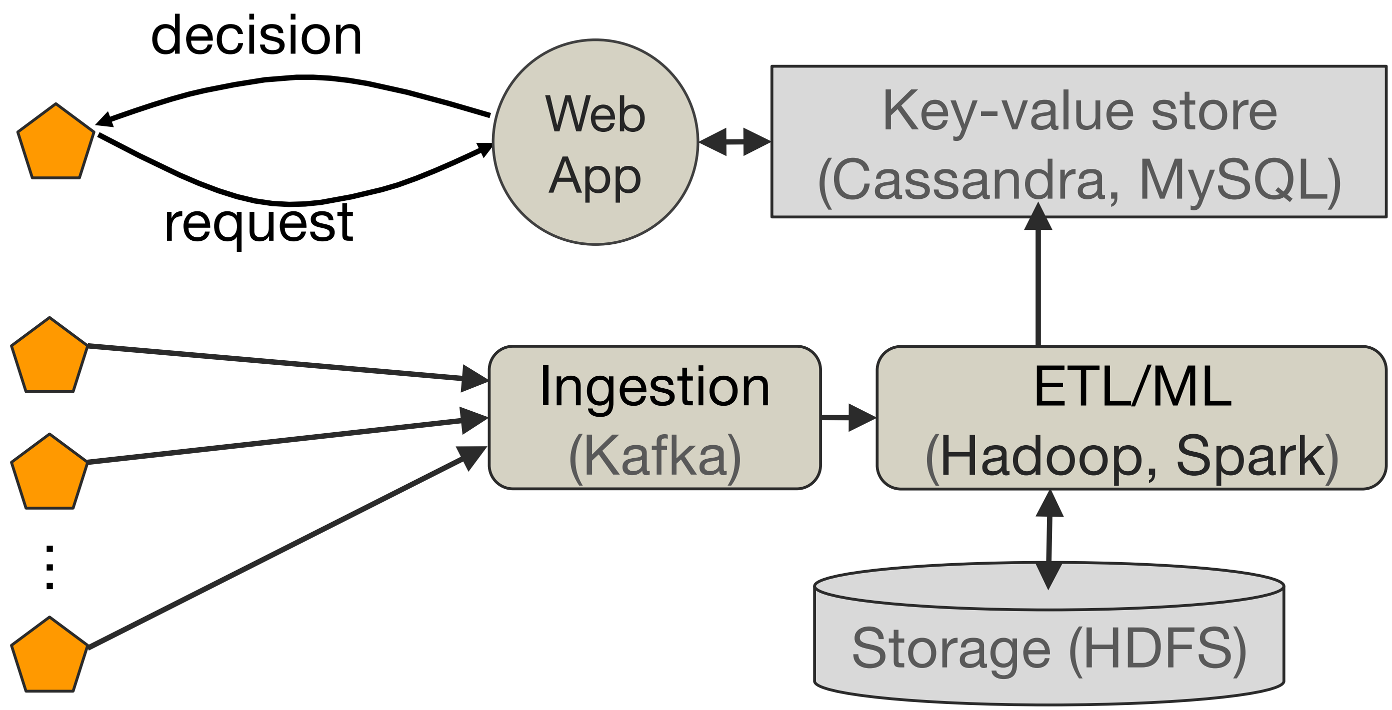

State-of-the-art security solutions:
Computation on encrypted data

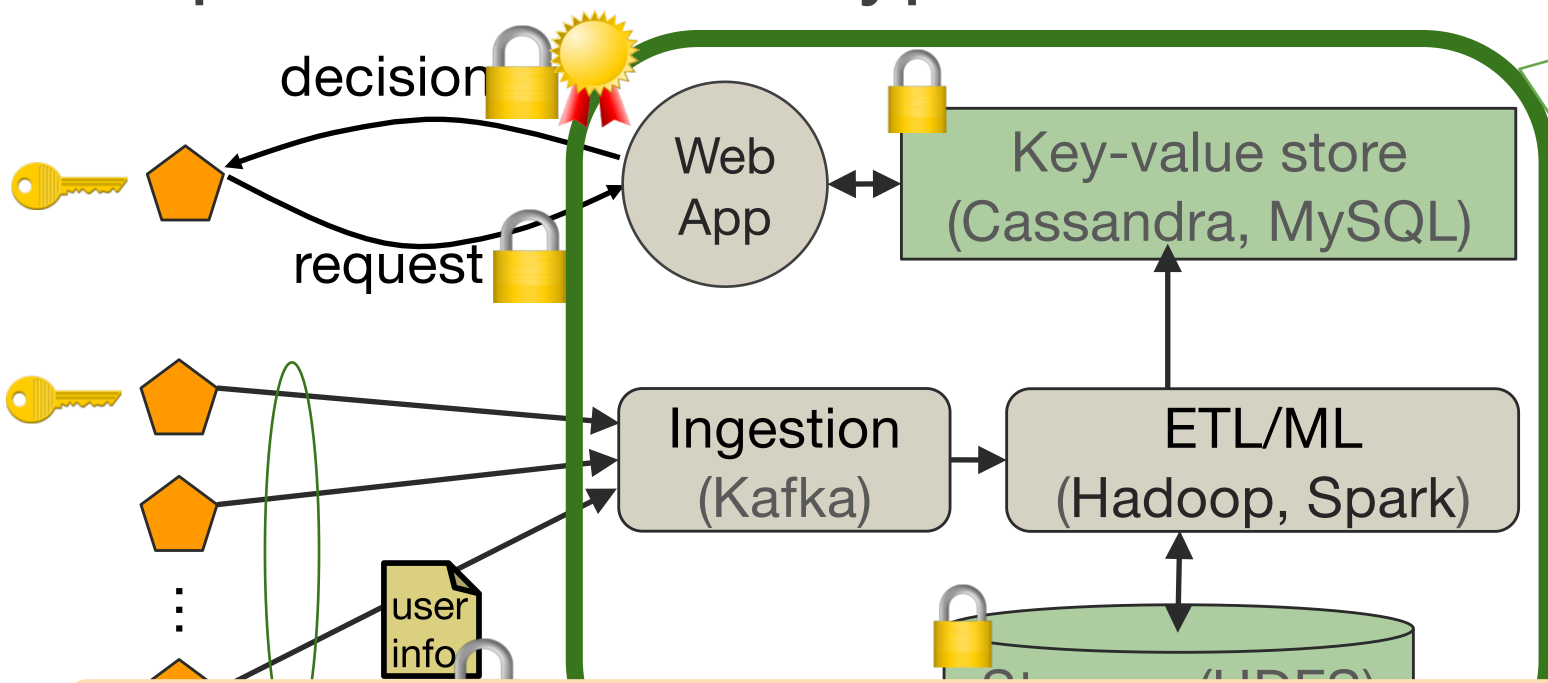

SQL (CryptDB), ML classification, web serving (Mylar), etc.

Ensure security, but relatively simple algorithms

State-of-the-art security solutions:
Hardware enclaves

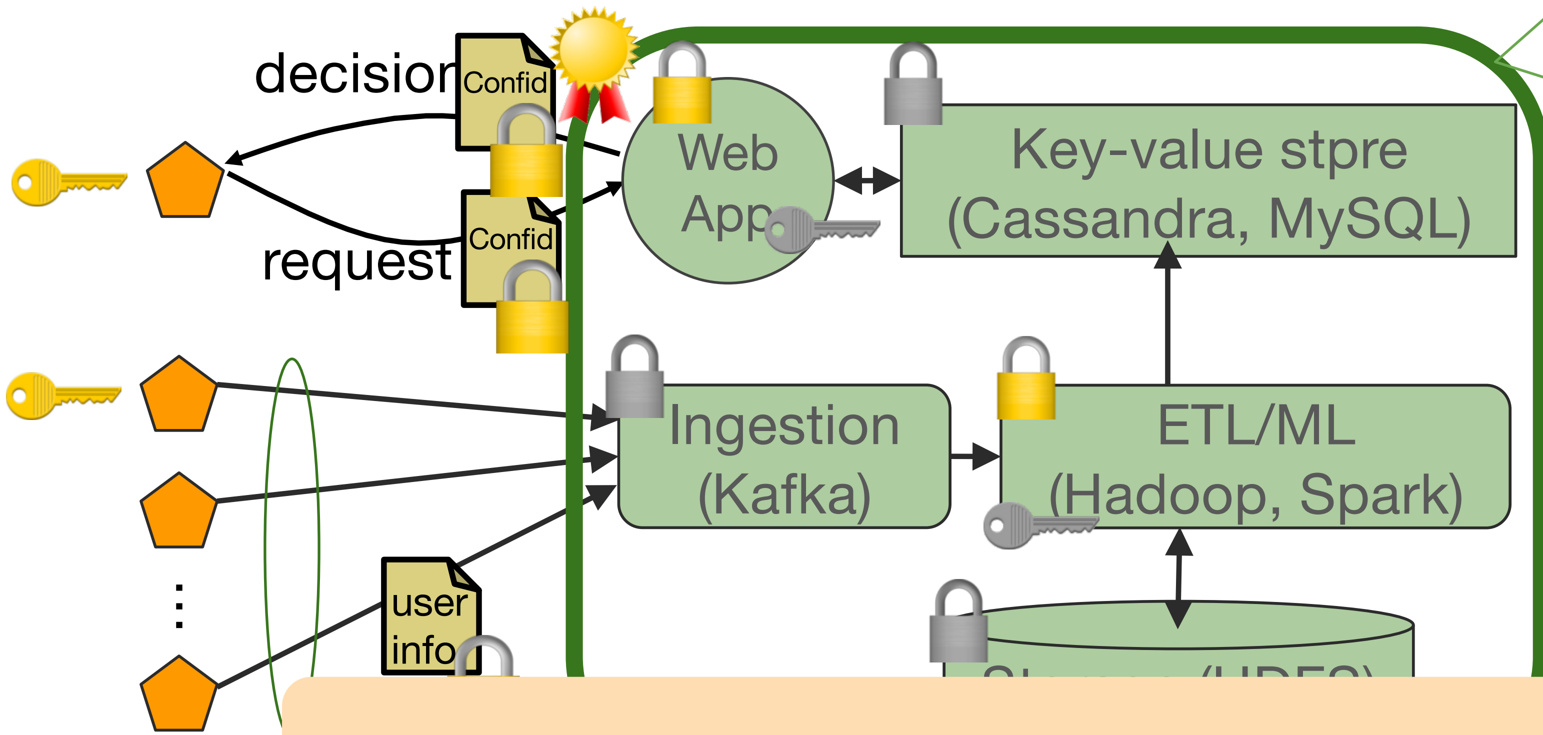

Intel SGX,
ARM
Trustzone,
RISC V (e.g.,
Haven, VC3)

TCB, side-channel leakage

Solution scorecard

Solution	Decision Quality	Latency		Security
		Decision	Update	
Pull decisions	simple	ms	min/day	weak
Push decisions	simple	ms/sec	hour	weak
State-of-the-art security	simple	ms	min/hour	strong

*light-weight, ensemble + correction models, policies

Solution scorecard: **RISE**

Solution	Decision Quality	Latency		Security
		Decision	Update	
Pull decisions	simple	ms	min/day	weak
Push decisions	simple	ms/sec	hour	weak
Security	simple	ms	min/hour	strong
RISE	sophisticated, accurate, robust	ms	sec	strong

*light-weight, ensemble + correction models, policies

Research areas

Systems: Spark-like functionality with 100x lower response time, and 1000x higher job throughput

Machine Learning:

- On-line ML algorithms
- Robust algorithms: handle noisy data, guarantee worst-case behavior

Security: achieve privacy, confidentiality, and integrity without impacting performance

Early Projects

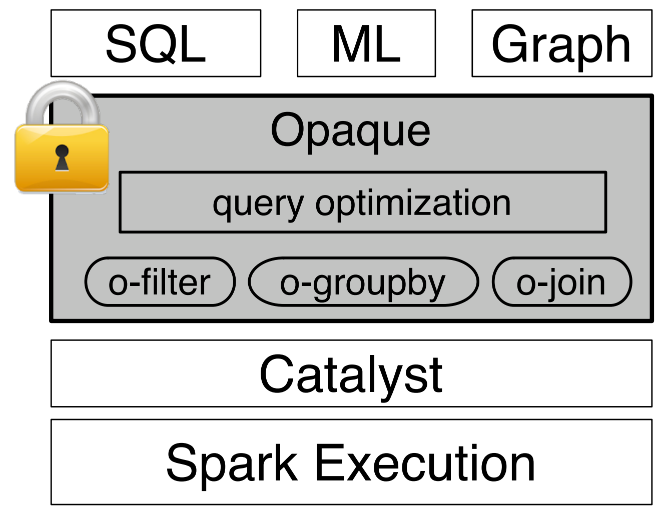

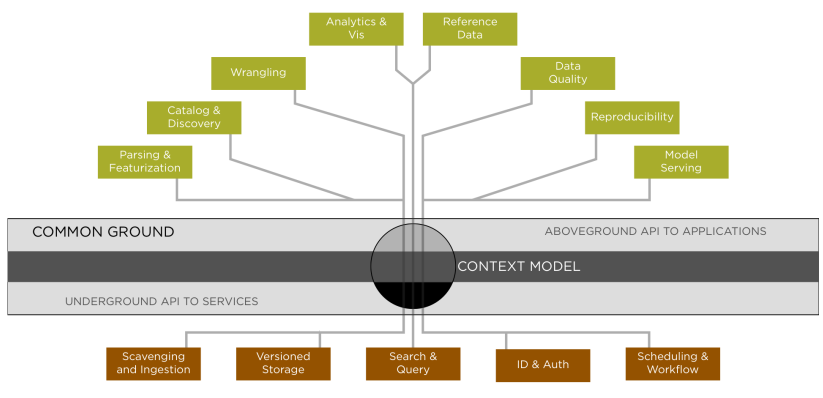

http://ground-context.org

Research area: Systems

on-going work   future work

unified model
(rich experience w/ both models)

IndexedRDDs

GPU/ASICs algos

Drizzle

sharded driver,
in-memory processing,
per-core NIC, HBM

system-state store

Apache Spark | Lattice Flow | Ray | Clipper | …

shim layer | shim layer | shim layer | shim layer

scheduler | optimizer | …

RISE μkernel

- Support task-graph & BSP execution models
- Support fine grain updates
- Support heterogeneous hardware
- Millisecond level parallel jobs
- Handle 10K-100K jobs/sec
- Ability to faithfully replay jobs

THE MEANING AND VALUE OF DATA DEPENDS ON CONTEXT

Application context
Multiple data models

because truth
is subjective

Behavioral context
Logs and Lineage

because behavior
determines meaning

Historical context
Immutable versions
for code and data

because
things change

A broader context for big data
ground

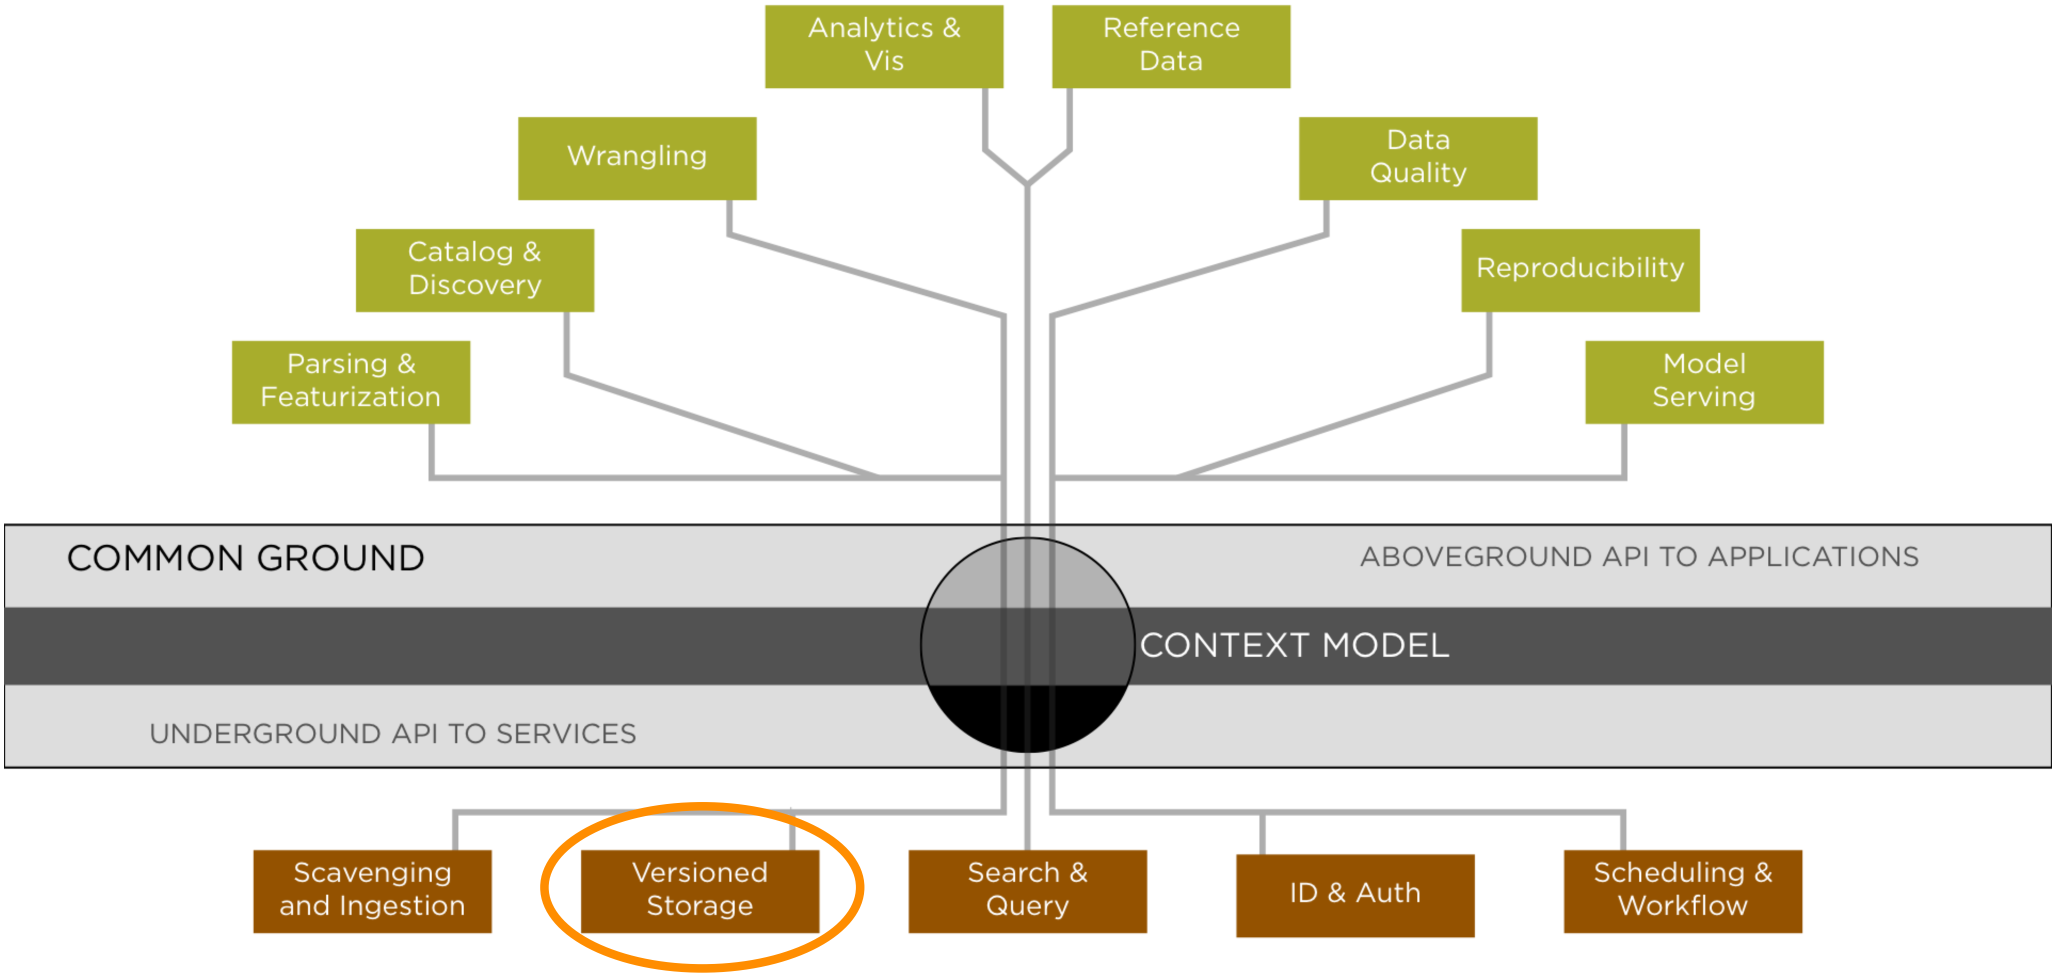

http://ground-context.org

LatticeFlow and Bedrock (working names)

Driving Hypotheses for LatticeFlow

- A core programming API for both real-time and scale
- Everything is (async) data: event dispatch, real-time data streams
- Coordination Avoidance: lattices + async dataflow = no locks/barriers

Initial proof point this fall: LatticeKVS

MultiScale, MultiConsistency key-value store

- Same system beats Redis on one node, Cassandra on scale-out…
- …while providing family of rich consistency and transactional isolation options
- …with lean codebase, derived from a core LatticeFlow library in C++

Prototype toward "Project Bedrock"

- Immutable, never-forget versioned storage under ground

Research area: Systems & ML

Focus on robust, on-line algorithms

Focus on ML/DL support (e.g., **Clipper**, **Ray**)

already gathering industry support

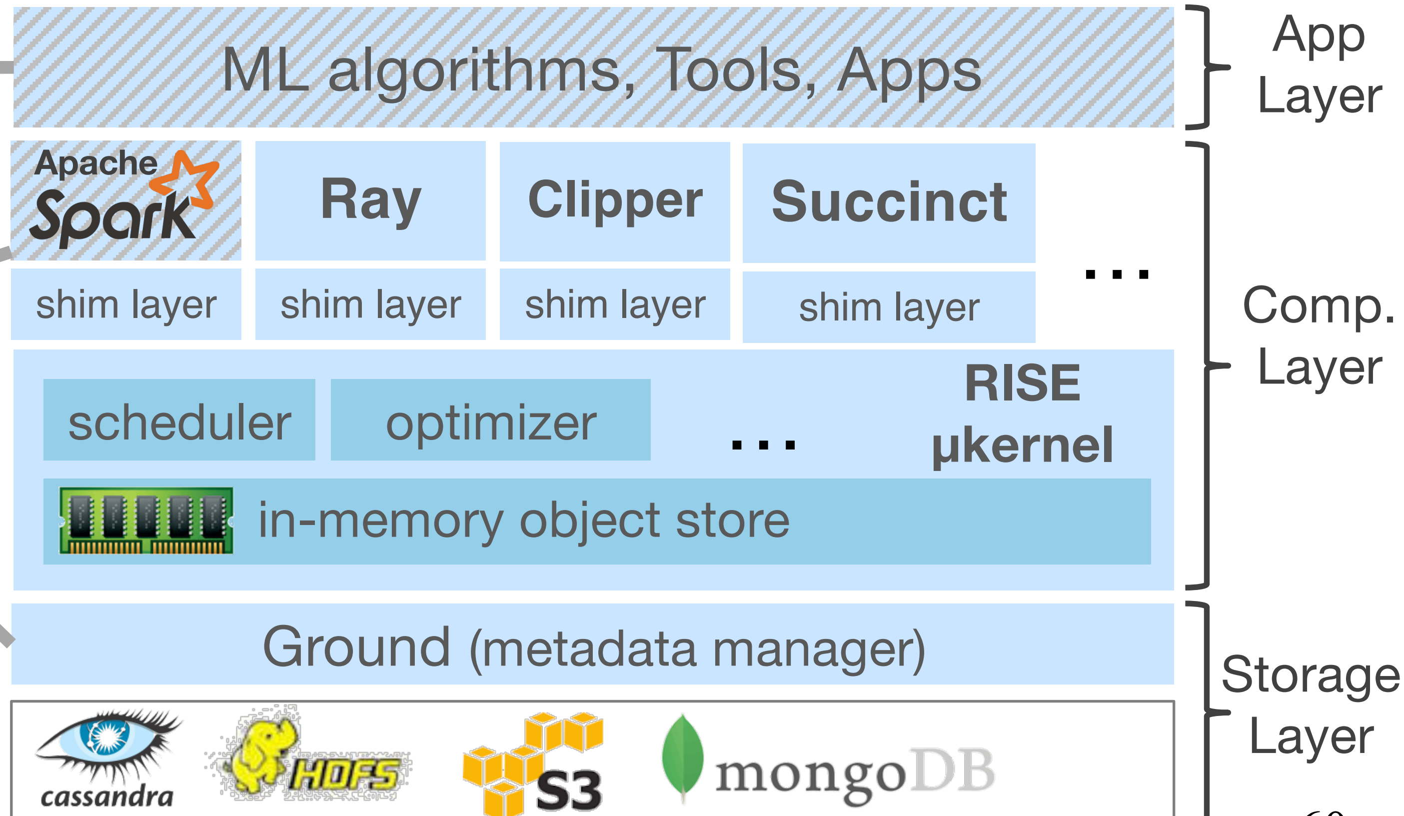

RISE

3rd party

RISE + 3rd party

Research area: ML

Robust optimization methods:

- noise tolerant and parallelizable

Handle uncertainty:

- robust control techniques to handle unforseen real world situations

Quantify decision accuracy:

- confidence and credible intervals on outputs of ML systems

On-line ML algorithms

- handle time-varying models with high performance and reliability

End-to-end system

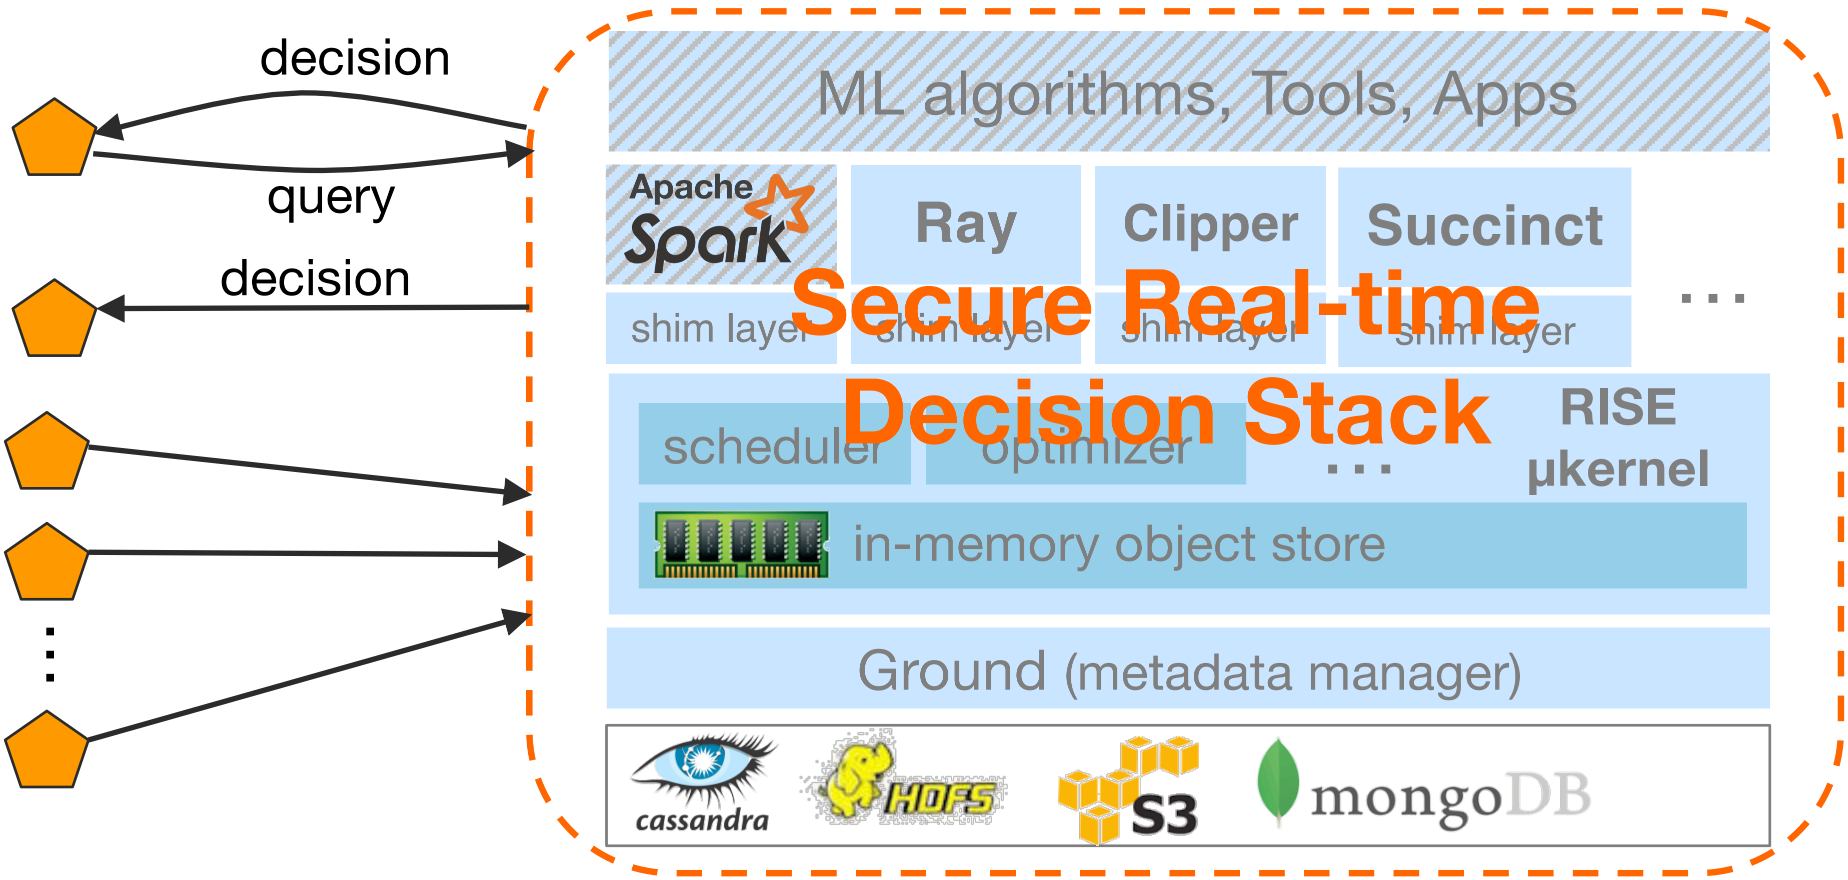

Research area: Security

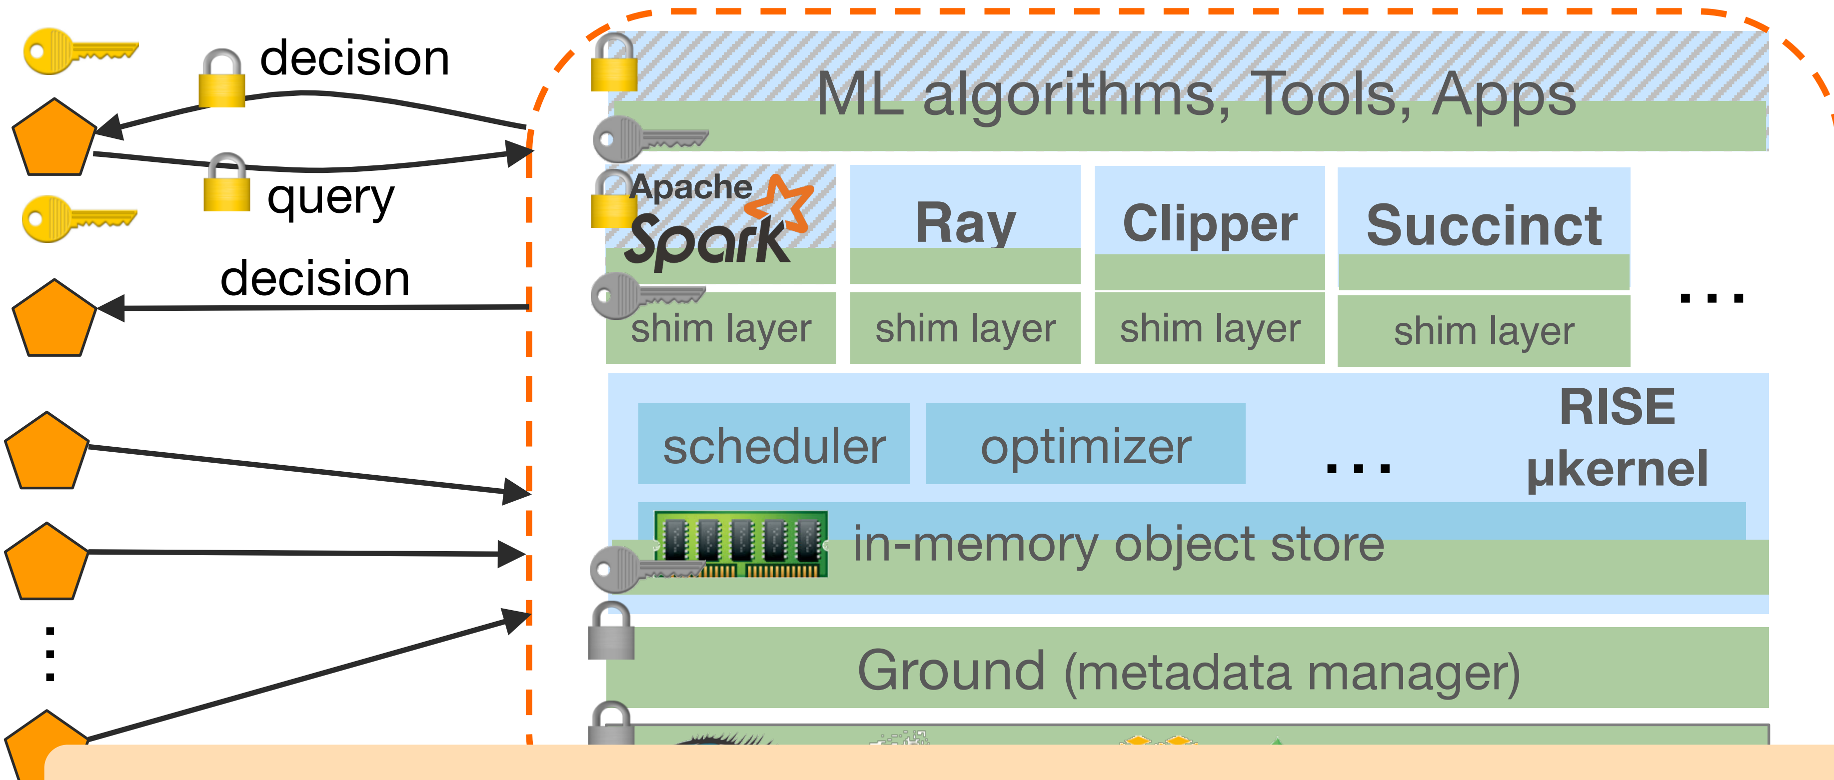

Split comp. between enclave and on encrypted data

Research area: Security

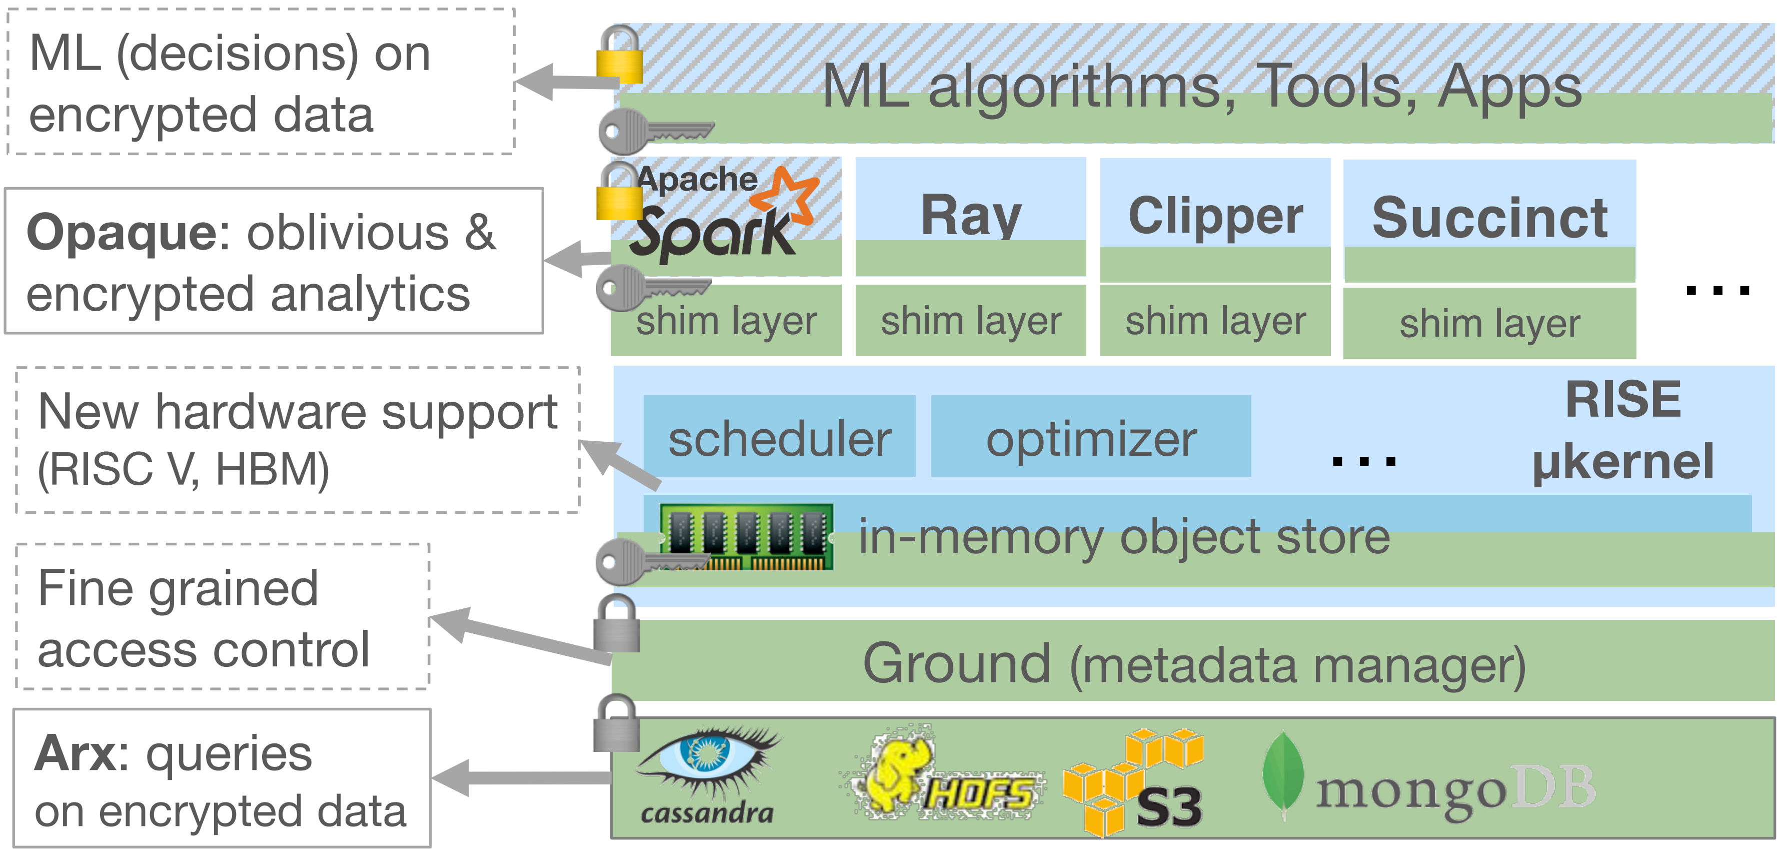

Why now, why us?

Latency

Quality

Security

Hardware Trends

Processing
- RISC V*, GPUs, ASICs
- Built in security support (enclaves)

Storage & networking
- 3D Xpoint

Next gen rack designs
- FireBox*

Systems

- Berkeley's BDAS*

Machine Learning

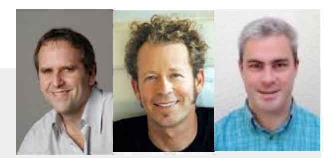

- Robust algorithms*
- On-line ML* and DL

Security

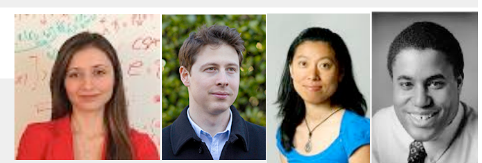

- Code verification tools*
- Comp. on encrypted data*

*blue indicates areas in which we are leaders

Summary

Goal: develop Secure Real-time Decision Stack,
an open source platform, tools and algorithms
for real-time decisions on live data with strong security

Five year project, similar to AMPLab
We are uniquely positioned to tackle this challenge

Looking to partner with companies