



Data Privacy

CMSC 491/691

L09 – Edge Computing: Federated Learning



Previously on...

- Secure Multi-Party Computation (MPC) replaces trusted party with technology to promote collaboration
- Based on adversarial model (e.g., honest-but-curious)
- Components: Shamir Secret Sharing (SSS), Oblivious Transfer (OT), ...
- High computational and communication costs!

Community

How Big Tech uses data privacy concerns for market dominance

Google gives Europe a 'reject all' button for tracking cookies after fines from watchdogs

In the news!

Cloud Computing

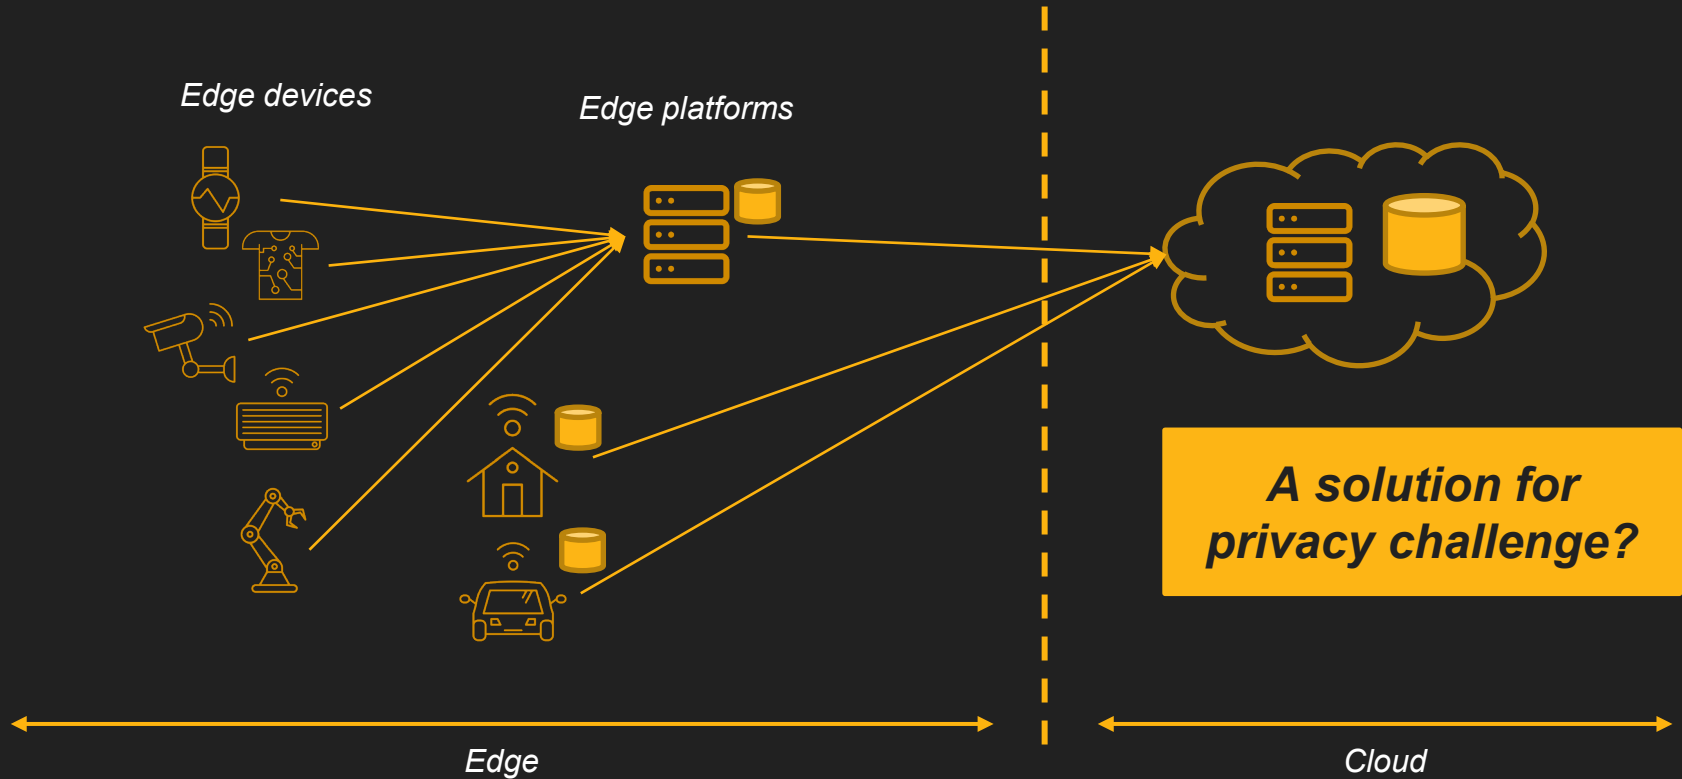


Huge, highly scalable computing and storage power

Cloud Computing Challenges

- **Latency**
 - Round-trip time to the cloud
- **Bandwidth**
 - Transference of large amounts of data
- **Connectivity**
 - Disconnection to the cloud
-
- **Privacy!**
 - Sensitive data transferred to the cloud

Edge Computing



Edge Computing S&P Challenges


Attacks/Threats

- Malicious Hardware/Software Injection
- Jamming Attacks
- Distributed Denial of Service (DDoS) Attacks
- Physical Attacks or Tampering
- Eavesdropping or Sniffing
- Non-Network Side-Channel Attacks
- Routing Information Attacks
- Forgery Attacks
- Unauthorized Control Access
- Different Privacy Leakages
- ...


Countermeasures

- Side-Channel Signal Analysis
- Trojan Activation Methods
- Policy-Based Mechanisms
- Securing Firmware Update
- Reliable Routing Protocols
- Intrusion Detection System (IDS)
- Cryptographic Schemes
- Secure Data Aggregation
- MPC
- DP
- ...

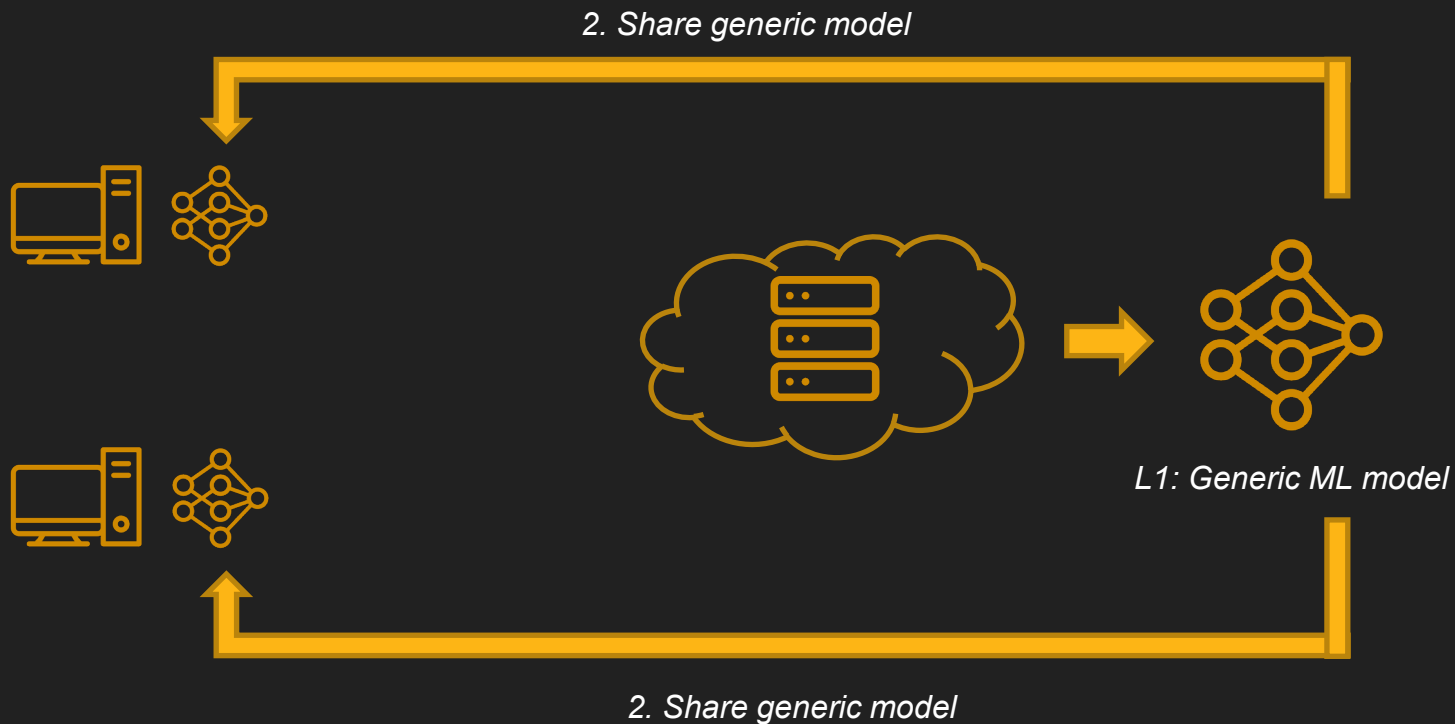
Example Domain: Machine Learning



Federated Learning



Federated Learning



Federated Learning




3. Train local models & generate new learnings using private data



L1: Generic ML model

Federated Learning



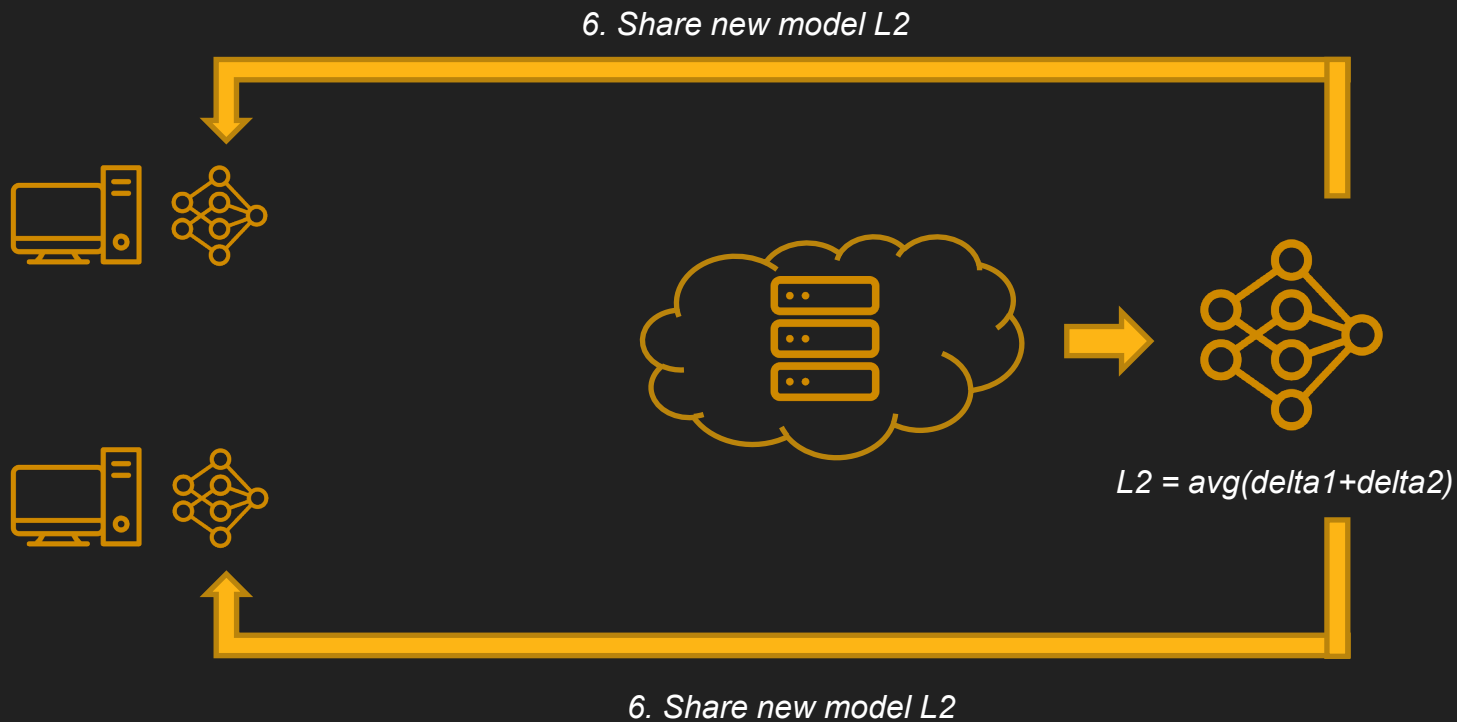
Federated Learning



5. Generate new model

$$L2 = \text{avg}(\text{delta1} + \text{delta2})$$

Federated Learning



Limits of Federated Learning

- **FL does not apply to all ML applications**
- Model might be too large for clients
- Client data might not be relevant
 - E.g., might not be clean!
- Clients might not label data
 - Problem for supervised techniques

Core Challenges

- Challenge 1: **Expensive Communication**
- Challenge 2: **Systems Heterogeneity**
- Challenge 3: **Statistical Heterogeneity**
- Challenge 4: **Privacy Concerns**

Challenge 1: Expensive Communication

- **Communication is a critical bottleneck!**
 - Send model updates from/to clients and server
- Massive number of devices (e.g., millions of smartphones)
- Slower network communication
- **Key ideas:**
 - Reduce total number of communication rounds
 - Reduce size of transmitted messages per round

Challenge 2: Systems Heterogeneity

- Storage, computational, and communication capabilities of devices may differ
 - Variability in hardware (CPU, memory), network connectivity (3G, 4G, 5G, Wi-Fi), and power (battery level)
- Devices might be unreliable
 - They might disconnect/stop at any round
- **Key ideas:**
 - Anticipate a low amount of participation
 - Tolerate heterogeneous hardware
 - Be robust to dropped devices in the network.


Challenge 3: Statistical Heterogeneity

- Devices frequently generate and collect data in a non-identically manner
- Number of data points across devices vary significantly
- Conflict with independent and identically distributed assumptions
- Challenging to learn a global model in this setting

Challenge 4: Privacy Concerns


- FL is a step towards protecting data generated on each device by sharing model updates...
- ...but! updates can reveal sensitive information
- **Key ideas:**
 - Anonymization?
 - Multi-Party Computation?
 - Differential Privacy?

FL Privacy/Security Attacks

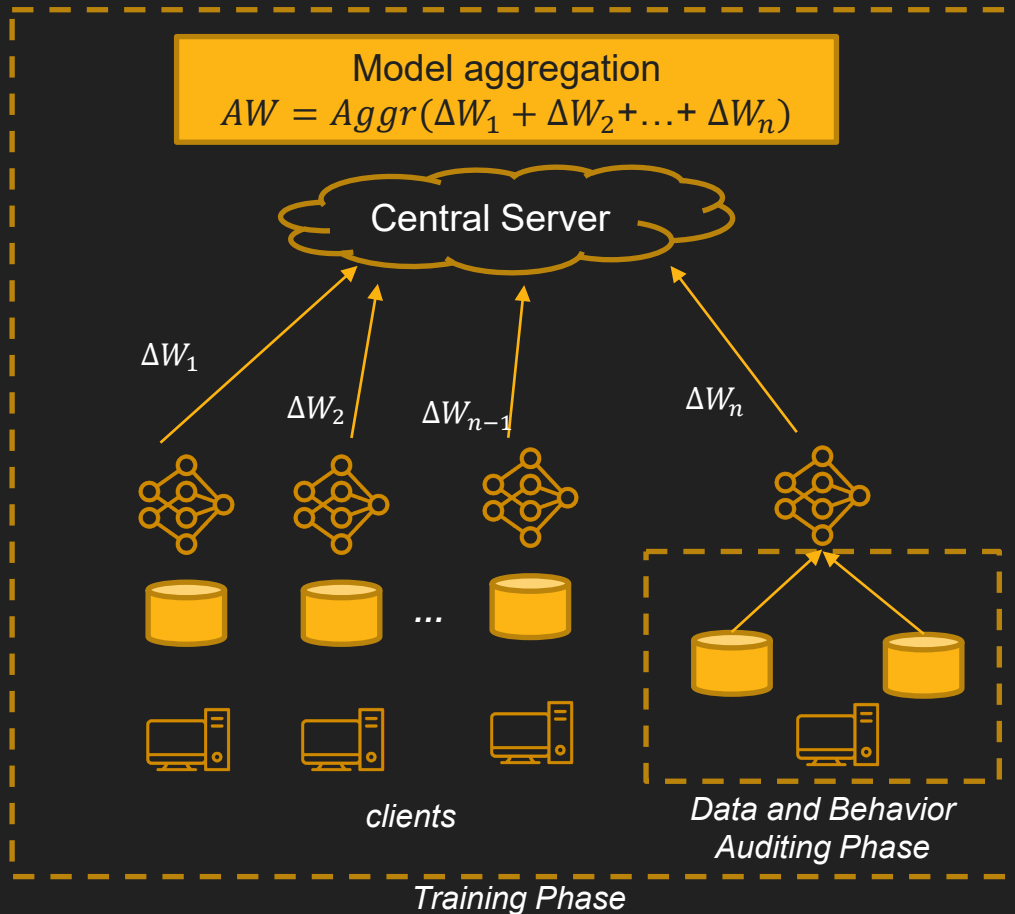


*Data and Behavior
Auditing Phase*

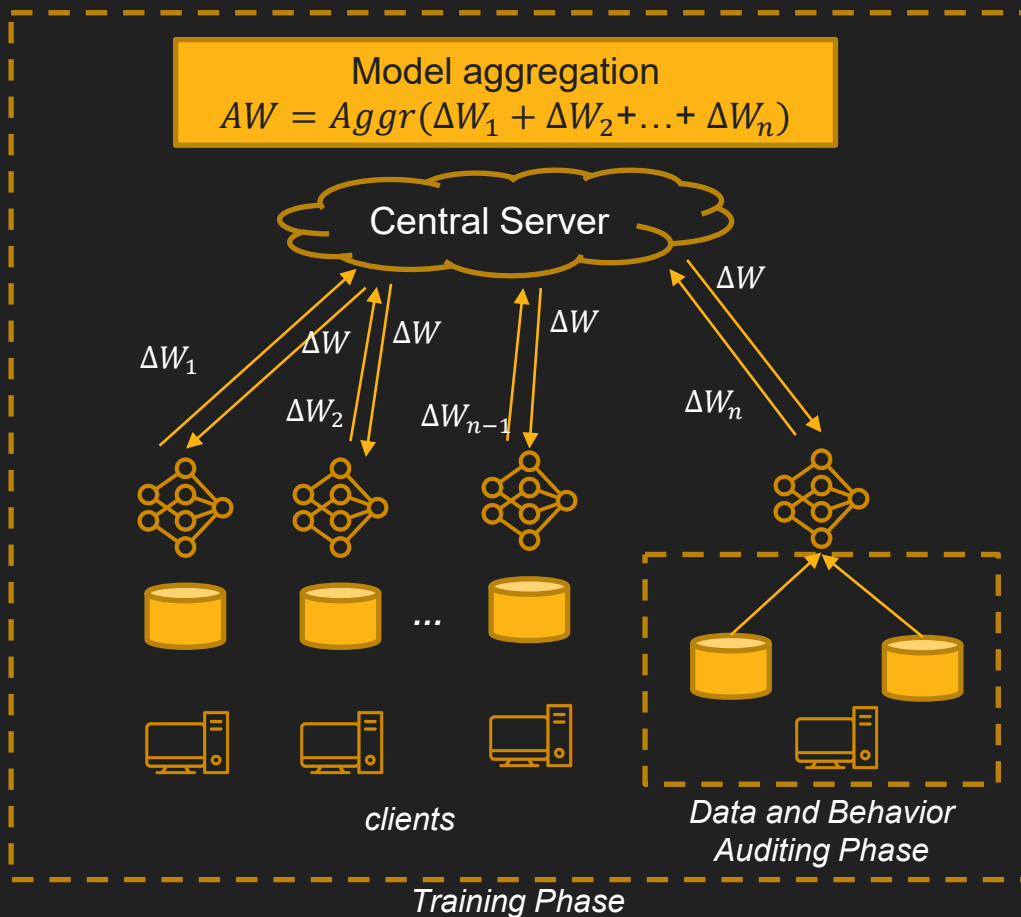
FL Privacy/Security Attacks



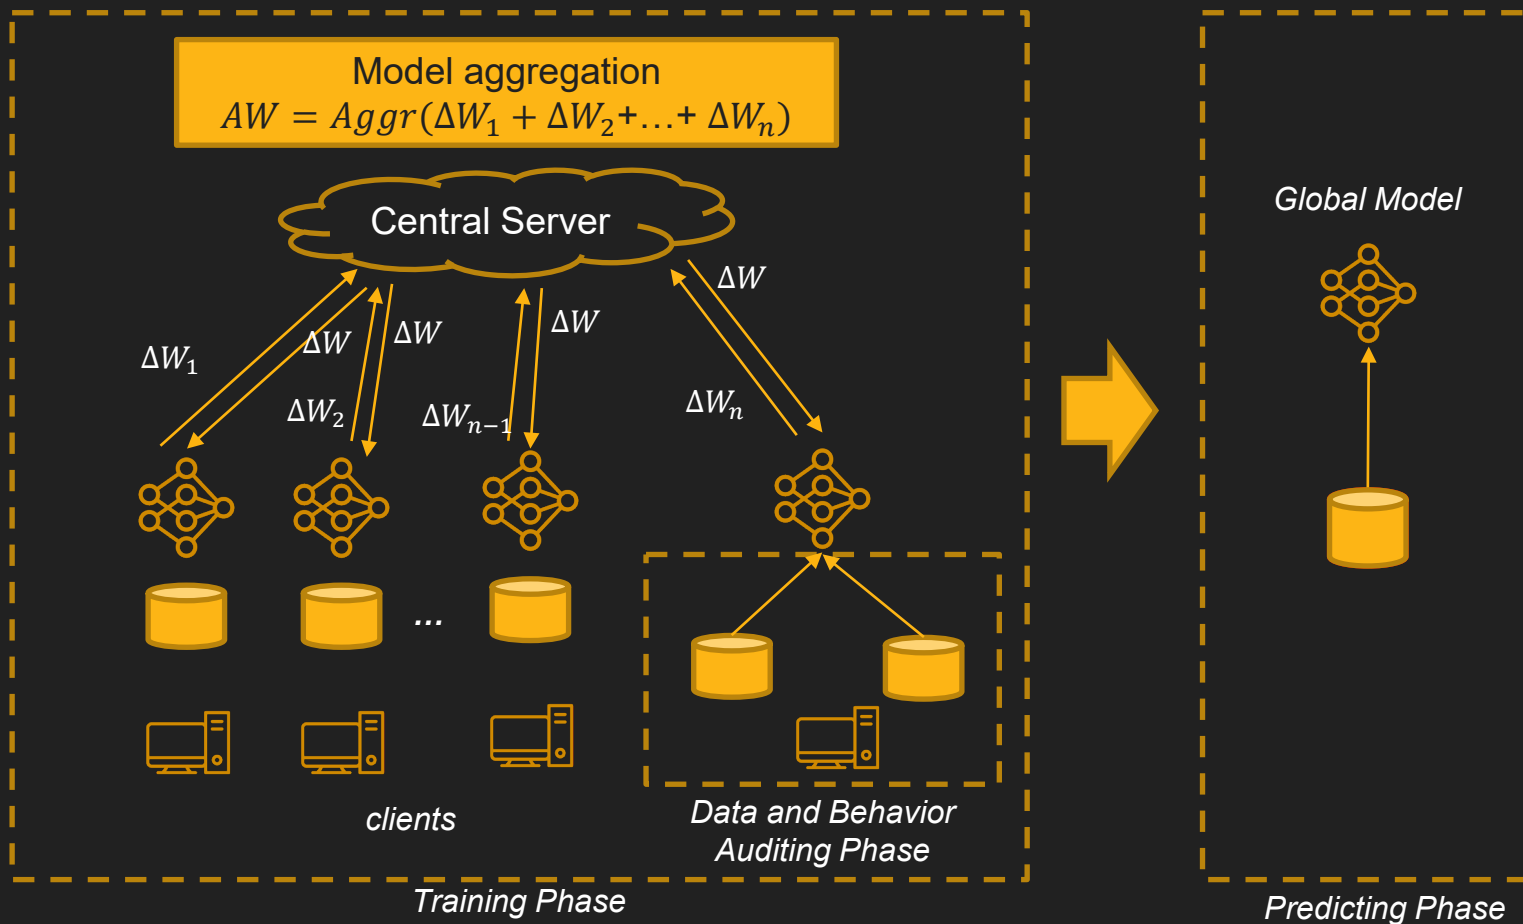
FL Privacy/Security Attacks




FL Privacy/Security Attacks



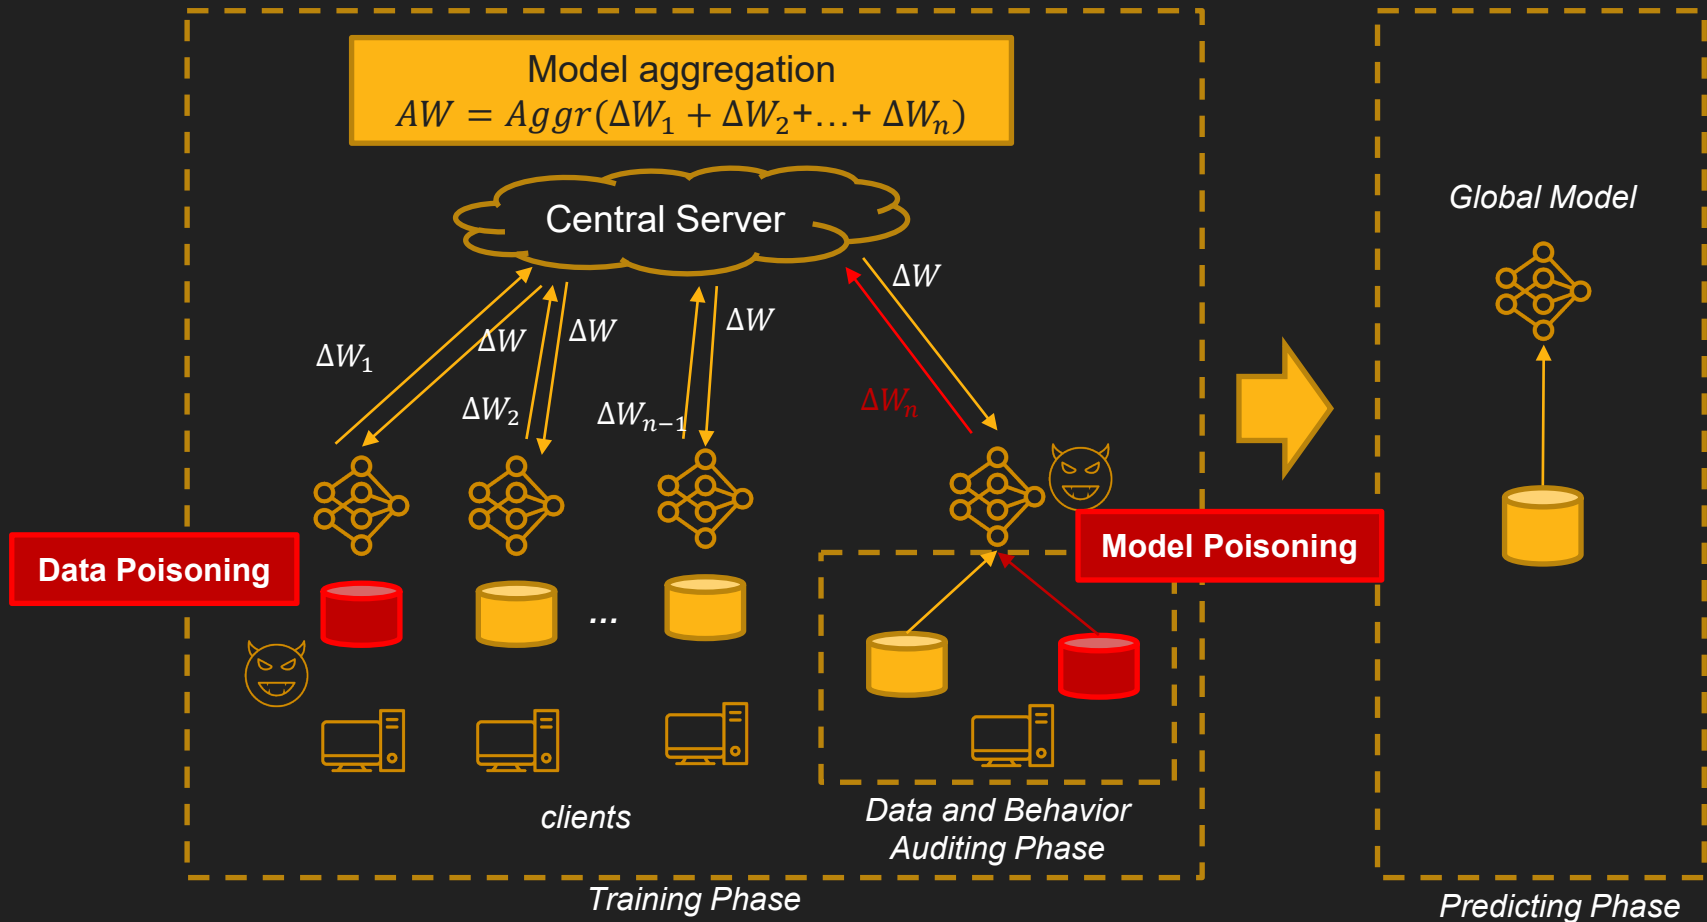
FL Privacy/Security Attacks



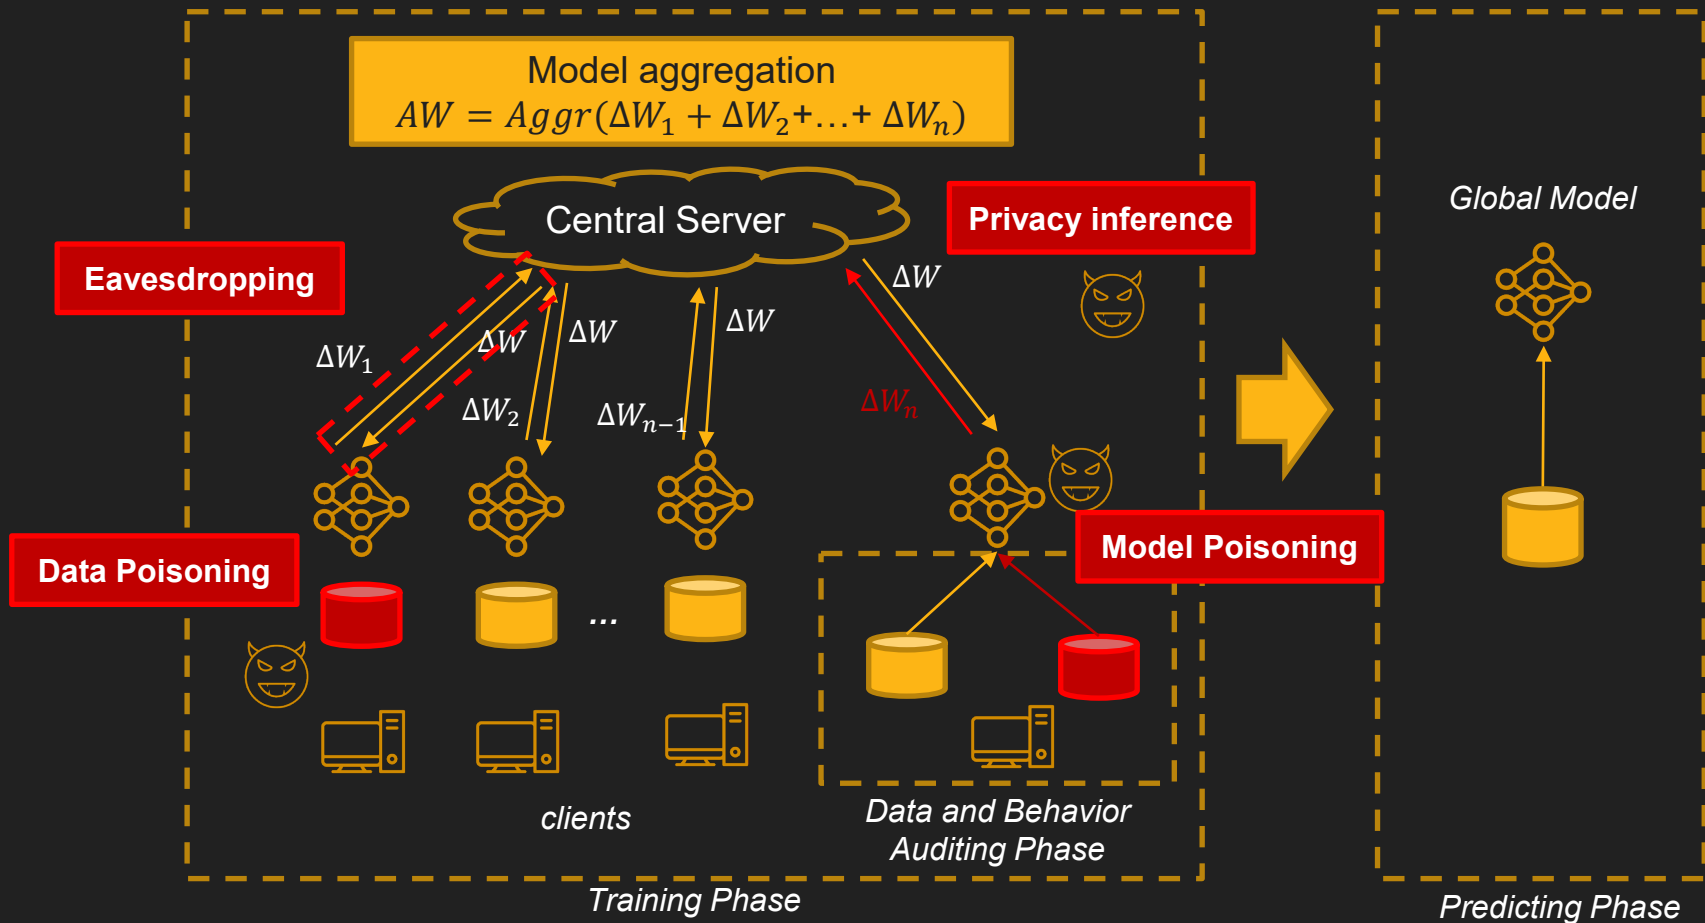
FL Privacy/Security Attacks



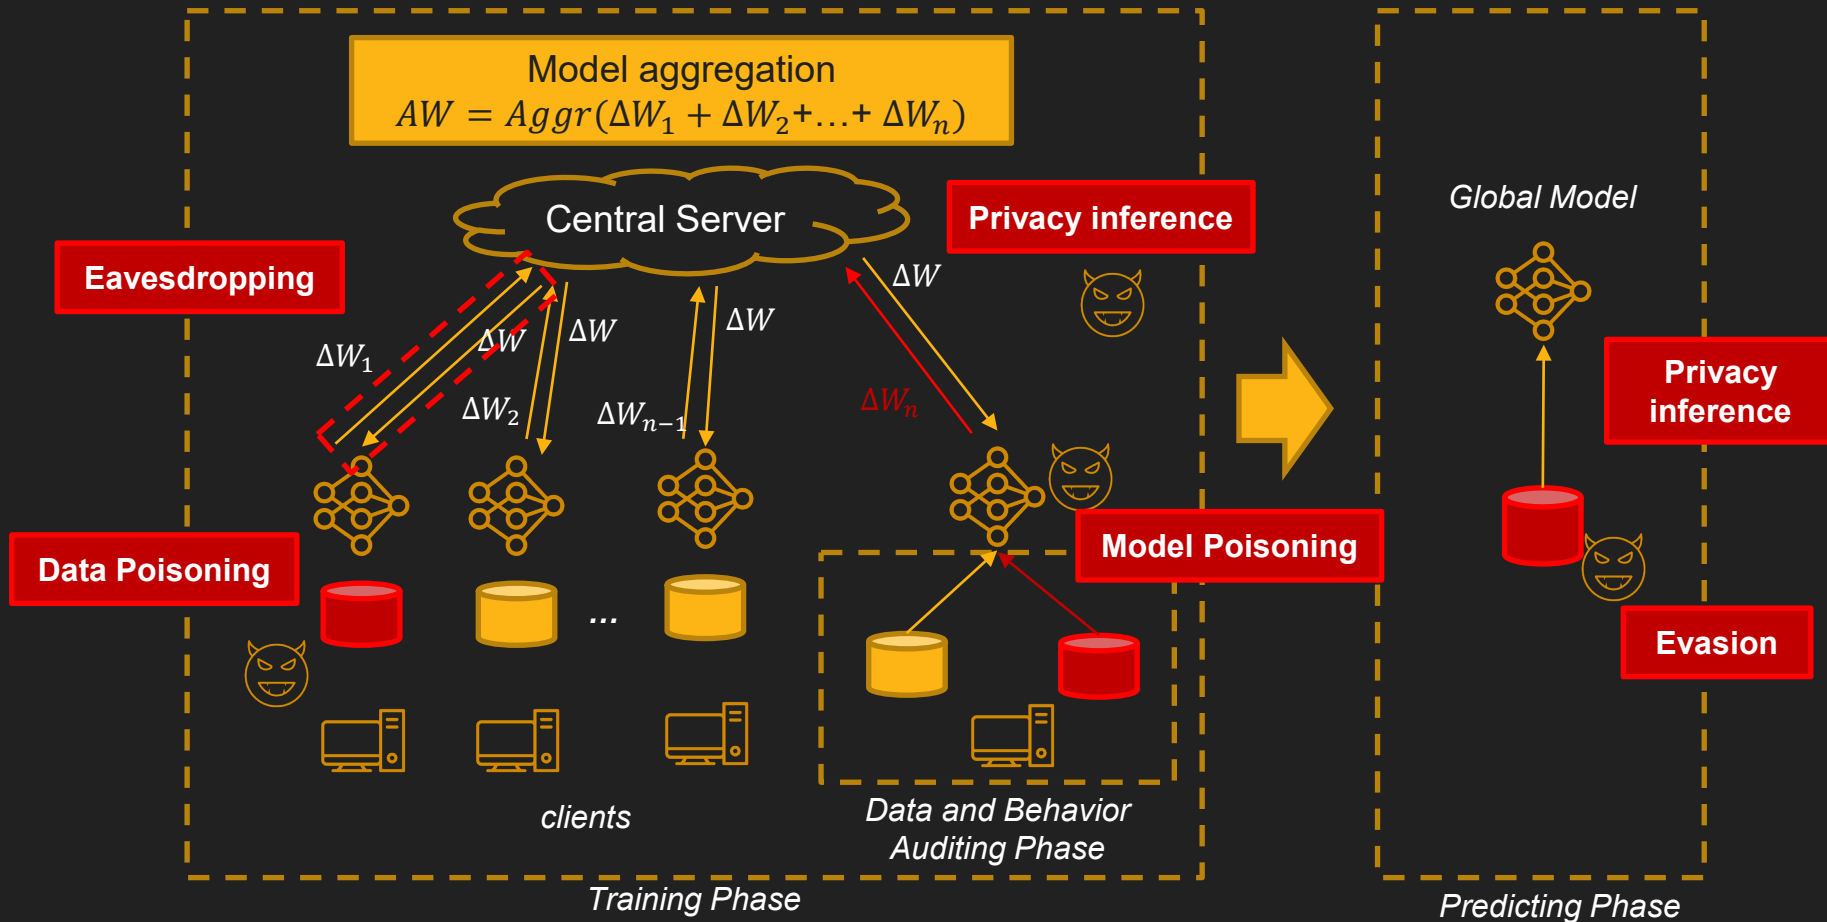
FL Privacy/Security Attacks



FL Privacy/Security Attacks



FL Privacy/Security Attacks




Privacy Attacks in ML

- **Privacy attacks / inference attacks / confidentiality attacks**
- Attacks against:
 - Training data
 - E.g., reveal the identity of patients whose data was used for training a model
 - ML model
 - E.g., reveal the architecture and parameters of a model that is used by an insurance company for predicting insurance rates
 - E.g., reveal the model used by a financial institution for credit card approval
- Main categories:
 - **Membership inference attack**
 - **Feature inference attack**
 - **Model extraction attack**

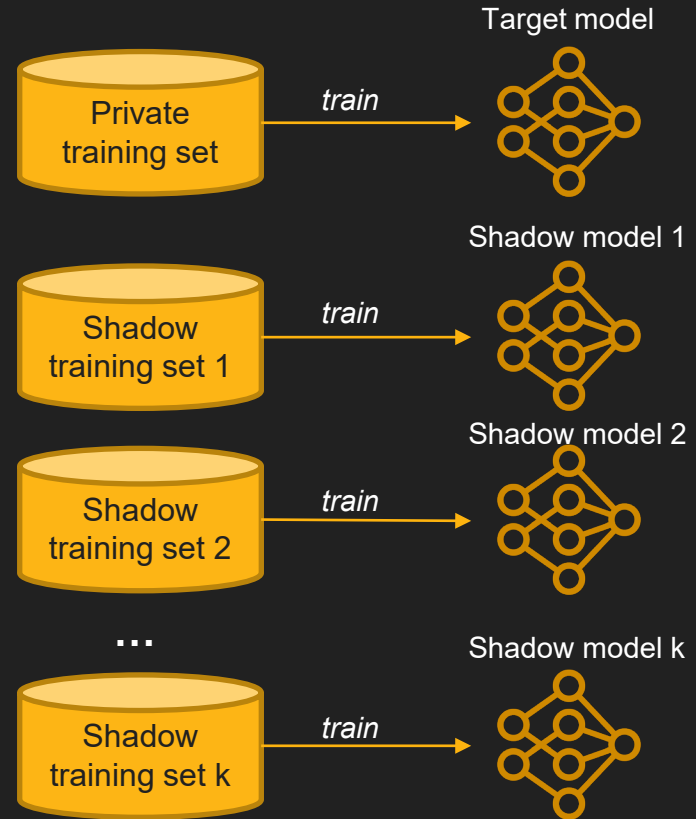
Membership Inference Attack

- **Adversarial goal:** determine whether or not an individual data instance is part of the training dataset for a model
- The attack typically assumes black-box query access to the model
- Attacks on both supervised classification models and generative models (GANs, VAEs) have been demonstrated



MIA: Shadow Training Attack

- **Threat model:**
 - Adversary has black-box query access to the target model
 - Goal: infer whether input samples were part of its private training set
- **Shadow training approach:**
 - Create several shadow models to substitute the target model
 - Each shadow model is trained on a dataset that has a similar distribution as the private training dataset of the target model



Feature Inference Attack

- **Adversarial goal:** recreate certain features of data instances or statistical properties of the training dataset for the model
- A.k.a. **attribute inference**, **reconstruction**, or **data extraction attack**
- Attacks developed to:
 - Recover partial information about the training data (such as sensitive features of the dataset, or typical representatives for specific classes in the dataset) or full data samples
 - Recreating dataset properties that were not encoded in the (**property inference attack**)
 - E.g., extract information about the ratio of men and women in a patient dataset, despite that gender information was not provided for the training records

FIA: Model Inversion Attack

- Creates prototype examples for the classes in the dataset
- Authors demonstrated an attack against a DNN model for face recognition
- Given a person's name and white-box access to the model, the attack reverse-engineered the model and produced an averaged image of that person

*Recovered image
using attack*



*Image of the person
used for training*



Attacks Against Distributed Learning

- Attacks can be **passive** (the adversary collects the updates) and **active** (the adversary shares information to impact the training procedure)
- Some (of many) examples:
 - **Membership inference attack** [1] : One of the clients is a malicious attacker that reveals if other participants used a data record for training
 - **Property inference attacks** [2]: Reveal whether training data with certain properties were used by the other participants
 - **Training data reconstruction attack** [3]: Use GAN model to reconstruct class representative samples from the local dataset used by the other participants

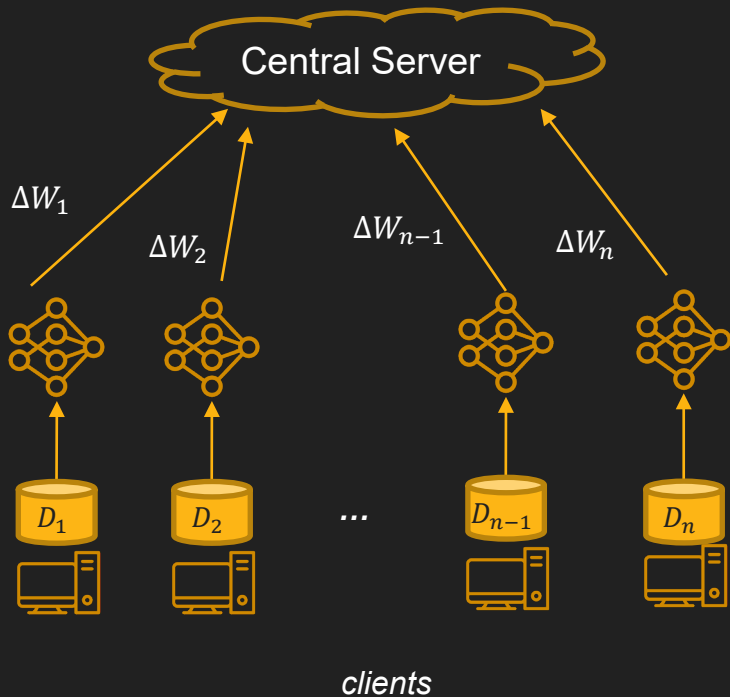
[1] Nasr et al. "Machine learning with membership privacy using adversarial regularization." ACM CCS. 2018.

[2] Melis et al. "Exploiting unintended feature leakage in collaborative learning." IEEE SP. 2019.

[3] Hitaj et al. "Deep models under the GAN: information leakage from collaborative deep learning." ACM CCS. 2017.



Mitigation Strategies?

Model aggregation
 $AW = Aggr(\Delta W_1 + \Delta W_2 + \dots + \Delta W_n)$



Federated Learning + Anonymization

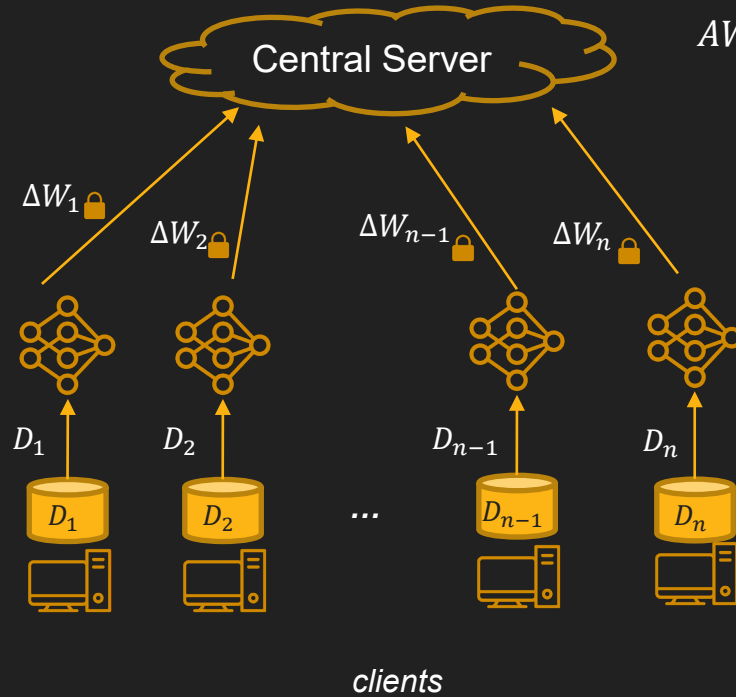

Pros and Cons?



Anonymize training data (e.g., remove identifiers, generalize sensitive data)

Federated Learning + MPC

Pros and Cons?





$$AW = Aggr(\Delta W_1 + \Delta W_2 + \dots + \Delta W_n)$$

Secret share model updates

clients

Federated Learning + Differential Privacy

Pros and Cons?



Add DP noise to model updates

Add DP noise to training data

Conclusions

- **Cloud computing** has benefits but many **drawbacks**
 - Latency, bandwidth, connectivity, **privacy!**
- **Edge computing** can **mitigate** some of the **drawbacks**
 - E.g., minimize the amount of individuals data transferred to cloud by performing local computations
- **Federated learning** is a popular example of **edge computing for ML**
- While this **helps in protecting privacy**, **attacks are still possible!**
- Need to **integrate PETs in Edge Computing / Federated Learning**

Group Activity

- Think about your group project (or any other application)
- If you use a client/server architecture...
 - What can you learn at the client?
 - What cannot you learn at the client?
 - What data would you need to transfer to server?