# Chapter 8: Federated Threat Intelligence Network

## Learning Objectives

By the end of this chapter, students will be able to:

1. Analyze the federated learning paradigm and explain its role in enabling cross-organizational threat intelligence sharing.
2. Construct differential privacy defenses utilizing the Gaussian mechanism and monitor privacy loss through Rényi Differential Privacy (DP) composition.
3. Explain how the SCAFFOLD algorithm addresses optimization instability on non-IID data and clarify the underlying assumptions supporting its effectiveness.
4. Develop a Transformer-based intrusion detector implemented in pure numpy, and apply INT8 quantization to facilitate efficient edge deployment.
5. Detect and mitigate Byzantine poisoning attacks using Multi-Krum and trimmed mean aggregation strategies.
6. Identify root causes in federated threat signals by applying intervention-effect reasoning based on Pearl's causal framework.
7. Evaluate adversarial robustness using Projected Gradient Descent attacks and randomized smoothing techniques.

## Agentic Lens

Here, the agent is part of a team within the federation. The main challenge is deciding what information to share, trust, reject, or quarantine, since each site has its own goals and data. For executives, this changes how risk and trust are managed between agencies. Every decision to share or withhold information shapes daily operations and the federation's reputation for reliable intelligence.

- **Agent role:** Coordinate threat-intelligence sharing without leaking or amplifying bad evidence.
- **Observations:** Local model updates, provenance, privacy budget, trust scores, poisoning signals, and cross-site disagreement.
- **Tools:** FedAvg, SCAFFOLD, differential privacy, robust aggregation, causal analysis, and adversarial-robustness checks.
- **State:** Update history, privacy-budget consumption, site reputation, quarantine decisions, and unresolved anomalies.
- **Verifier:** Robust aggregation diagnostics, provenance checks, privacy accounting, and analyst review for high-impact global updates.
- **Guardrails:** No site should receive preferential treatment based on size or familiarity. Global model updates must not exceed the established privacy budget.
- **Failure mode of note:** An agent that prioritizes consensus exclusively may inadvertently average out attacks, leading the federation to accept compromised telemetry as trustworthy.

## 8.1 Introduction: The Intelligence Sharing Dilemma

Each security operations center scans a unique slice of the threat landscape. A university might spot a burst of credential-stuffing, a hospital could catch the first flicker of ransomware, and a factory may notice PLC exploitation. Alone, these are scattered signals. Together, they form patterns that expose coordinated attacks.

Data governance stands as the main barrier. Healthcare networks guard patient records, regulations fence in financial institutions, and even within one company, regional branches may be forbidden from pooling their network data.

Federated learning solves this problem by letting each site train its own model and send only updates—not raw data—to a central hub. Scan records stay local. However, model gradients can still leak information. To prevent this, the chapter explains differential privacy and strict privacy budget controls. The aggregator uses FedAvg to combine updates and can use SCAFFOLD for more stable results when data varies, building a global threat model from shared signals. This is a teaching example based on the chapter's workflow, not a universal benchmark. The results assume the same number of sites, devices, and privacy budget as the modeled federation. Do not treat this as a production guarantee for all federations. You can reproduce these results by following the chapter's workflow and code with the given parameters. However, broader claims about model performance or real-world use are only examples and have not been independently validated. Leaders should use these results to build intuition, not as guarantees for live risk or policy decisions.

> *"Federated model trained across 5 sites and 2,847 devices. Privacy budget consumed: epsilon 2.3 of 10.0. Threat signals raised: 14. Three appear consistent with a coordinated scanning campaign."*

## 8.2 Federated Learning Fundamentals

### 8.2.1 The FedAvg Algorithm:

The Breakwater implementation in secure_aggregator.py:

```
class SecureAggregator:
    def __init__(self, min_sites: int = 3) -> None:
        self.min_sites = max(min_sites, 1)
        self._update_buffer: list[LocalModelUpdate] = []

    def receive_update(self, update: LocalModelUpdate) -> None:
        "Buffer a local model update. Duplicates replaced."""
        self._update_buffer = [
            u for u in self._update_buffer if u.site_id != update.site_id
        ]
        self._update_buffer.append(update)

    def aggregate(self) -> AggregatedModel | None:
        """Run FedAvg if enough updates received."""
        if len(self._update_buffer) < self.min_sites:
            return None
        updates = list(self._update_buffer)
```

```
        weights = self.fedavg(updates)
        # ...
```

### 8.2.2 The Non-IID Problem

FedAvg works best when all clients have similar data. In practice, security data varies significantly from site to site. Hospitals monitor medical device traffic, factories analyze Modbus/TCP data, and universities search for DNS tunneling. This variety can pull local models in different directions, making FedAvg unstable or unable to improve.

### 8.2.3 FFOLD: Variance Reduction for Non-IID Data

SCAFFOLD (Karimireddy et al., 2020) addresses client drift using control variates. Each client maintains a local control variate c_i, and the server maintains a global control variate c. The local update rule becomes:

```
# SCAFFOLD update rule with control variates
y_i_new = y_i - lr * (gradient - c_i + c)
```

The correction term, which is the negative local control variate plus the global control variate, helps reduce client-specific drift. According to the SCAFFOLD paper, this method improves convergence and is more resilient to differences in data distribution than standard FedAvg.

```
from dataclasses import dataclass
import numpy as np

@dataclass
class ClientState:
    "Per-client state between federated rounds."""
    client_id: str
    local_control: np.ndarray  # c_i -- local control variate
    n_samples: int = 0
    round_participated: int = 0
    cumulative_drift: float = 0.0
```

**Result under the SCAFFOLD paper assumptions (Theorem 2):** For non-convex objectives with L-smooth gradients and bounded variance sigma^2, SCAFFOLD converges to an epsilon-stationary point in:

**Evidence label:** Illustrative. Method note: The convergence analysis reproduces Theorem 2 result from the cited paper under its formal conditions. Boundary: Gradients that are L-smooth and variance that remains bounded are theoretical assumptions. In the field, privacy noise, site drop-outs, and operational heterogeneity shape the actual training trajectory. Software anomalies, resource constraints, or poisoning campaigns can also degrade utility.

Leaders should use these results to build intuition for optimal conditions, not as guarantees for live risk or policy decisions.

```
# Convergence T = O(L * Delta / epsilon^2 + sigma^2 / (m * epsilon^2))
```

The paper's analysis erases the explicit heterogeneity term found in FedAvg bounds. However, in the real world, data differences never disappear. Sampling noise, privacy noise, client drop-outs, and model missteps all shape how training unfolds.

sequenceDiagram

Participant S as Server

Participant C1 as Client 1 (Hospital)

Participant C2 as Client 2 (Factory)

Participant C3 as Client 3 (University)

S->>C1: Broadcast (w_r, c)

S->>C2: Broadcast (w_r, c)

S->>C3: Broadcast (w_r, c)

Note over C1: Train K steps with<br/>control variate correction

Note over C2: Train K steps with<br/>control variate correction

Note over C3: Train K steps with<br/>control variate correction

C1->>S: (delta_y_1, delta_c_1) + DP noise

C2->>S: (delta_y_2, delta_c_2) + DP noise

C3->>S: (delta_y_3, delta_c_3) + DP noise

Note over S: Aggregate model deltas<br/>Update global control variate.

S->>S: w_{r+1}, c_{r+1}

### 8.4 Differential Privacy Engine

#### 8.4.1 The Gaussian Mechanism

The DPEngine in dp_engine.py implements per-sample gradient clipping and calibrated Gaussian noise injection following the DP-SGD framework (Abadi et al.20, 16):

```
import math


class DPEngine:
    def __init__(
        self,
        epsilon: float = 0.5,
        delta: float = 1e-5,
        clip_norm: float = 1.0,
        total_epsilon_budget: float = 10.0
    ) -> None:
        self.epsilon = max(epsilon, 1e-10)
        self.delta = max(delta, 1e-15)
        self.clip_norm = max(clip_norm, 1e-10)
        self._noise_scale = self._calibrate_noise_scale(
            self.epsilon, self.delta, self.clip_norm
        )

    def clip_gradients(self, gradients: list[float], max_norm: float | None =
None) -> list[float]:
        """Clip gradient vector to L2 norm bound."""
        bound = max_norm if max_norm is not None else self.clip_norm
        norm = math.sqrt(sum(g * g for g in gradients))
        if norm <= bound:
            return list(gradients)
        scale = bound / norm
        return [g * scale for g in gradients]

    def add_noise(self, gradients: list[float], sensitivity: float | None = None)
-> list[float]:
        """ Add calibrated Gaussian noise for (epsilon, delta)-DP."""
        # ...
```

The noise scale is calibrated using the Gaussian mechanism formula:

```
# Gaussian Mechanism Formula
sigma = (sensitivity * math.sqrt(2 * math.log(1.25 / delta))) / epsilon
```

### 8.4.2 Rényi Differential Privacy Composition

> The engine tracks cumulative privacy loss using Rényi Differential Privacy (RDP), which provides tighter composition bounds than basic DP. A simple conversion sketch for repeated rounds is:

```
# RDP Composition Sketch
epsilon_total = sum(eps_rounds) + math.log(1 / delta) / (alpha - 1)
```

The implementation keeps a running account rather than relying on this line alone. The point of the formula here is conceptual: privacy loss compounds, and the deployment should stop before it silently exceeds the declared accounting assumptions.

## 8.5 TransformerIDS: Intrusion Detection with Self-Attention

The transformer_ids.py module implements a 2-layer self-attention Transformer encoder in pure numpy:

```
# TransformerIDS Architecture
# Input(64) -> MultiHeadAttention(heads=4, d_model=64) x 2 -> FFN(256) ->
Output(1)

def _xavier_init(rows: int, cols: int, rng: np.random.Generator) -> np.ndarray:
    "Xavier (Glorot) uniform initialization."""
    limit = math.sqrt(6.0 / (rows + cols))
    return rng.uniform(-limit, limit, size=(rows, cols)).astype(np.float32)

def _gelu(x: np.ndarray) -> np.ndarray:
    "Gaussian Error Linear Unit."""
    return 0.5 * x * (1.0 + np.tanh(math.sqrt(2.0 / math.pi) * (x + 0.044715 *
x**3)))
```

Key design choices:

1. **Pure numpy implementation: No PyTorch or TensorFlow is needed, so that the model can run** on edge devices without GPU libraries.
2. **GELU activation** is used because it is smoother than ReLU, as recommended in the original Transformer paper.

3. **Layer normalization is** applied before each sub-layer (pre-norm) to improve training stability. GELU activation, which is smoother than ReLU, follows the recommendation from the original Transformer paper. This design allows deployment on gateway-class ARM edge nodes, but microcontroller-class devices usually need a more specialized runtime than standard numpy.

**Evidence label:** Illustrative. Method note: The compact model size and quantization numbers describe the Chapter 8 teaching example. Assumptions: The numpy implementation, feature width, and quantization scheme match the chapter's code. Boundary: This is a specific engineering example, not a claim that every edge IDS should use this exact design.

Adaptation criteria: To use this approach in other edge environments, consider several factors. Hardware type (like ARM, x86, or RISC-V), available memory, CPU power, and whether hardware acceleration is present will affect the design. The programming stack (numpy, microcontroller libraries, or vendor SDKs), feature size, and the amount and type of input data also matter. Data privacy needs, latency limits, and unreliable connections may require changes to model size, quantization, or communication methods. Doctoral students should weigh these factors to decide whether a pure numpy model with INT8 quantization is sufficient, or whether more engineering changes are needed for the target deployment.

The model uses 64-dimensional feature vectors from host scan data, including port counts, service types, banner hashes, and traffic patterns, to produce an anomaly score.

### 8.6 Byzantine Fault Tolerance

#### 8.6.1 Poisoning Detection

The poison_detector.py module implements four defense mechanisms:

```
def detect_poisoned_updates(updates: list[LocalModelUpdate]) -> list[str]:
    """Spot suspicious updates using multiple heuristics." "
    suspicious = set()
    # Heuristic 1: Gradient norm outlier detection (>3 sigma)
    norms = [math.sqrt(sum(x*x for x in flatten(u.gradient_data))) for u in
updates]
    threshold = mean(norms) + 3.0 * std(norms)
    # ...
```

1. **Gradient-norm outlier detection** flags updates whose L2 norm exceeds 3 standard deviations above the mean.
2. **Cosine similarity filtering** flags updates whose cosine similarity with the majority falls below a set threshold.
3. **Multi-Krum selection** chooses the K updates that are closest to their K nearest neighbors, leaving out outliers (Blanchard et al., 2017).

4. **Trimmed mean aggregation** removes the top and bottom beta fraction of each gradient coordinate before averaging (Yin et al., 2018).

### 8.6.2 Byzantine Threat Model

Within the standard robust-aggregation model used by Multi-Krum, the system is designed for a minority of Byzantine (arbitrarily malicious) clients. The familiar rule of thumb is f < n/3, but it should be interpreted as a theoretical design ceiling under the paper's assumptions, not as a blanket production guarantee. Small participant counts, coordinated poisoning, and badly scaled gradients can all narrow the safe region.

graph TD

A[Receive Updates] --> B[Norm Outlier Detection]

B --> C[Cosine Similarity Filter]

C --> D{Suspicious?}

D -->|Yes| E[Quarantine Update]

D -->|No| F[Multi-Krum Selection]

F --> G[Trimmed Mean Aggregation]

G --> H[FedAvg on Clean Updates]

style E fill:#ffcdd2

style H fill:#c8e6c9

### 8.7 Causal Inference for Threat Attribution

The causal_inference.py module constructs a causal DAG from correlated threat signals and applies Pearl's do-calculus to identify root causes:

```
def build_causal_dag(hosts: list[dict[str, Any]], signals: list[ThreatSignal]) ->
dict[str, Any]:
    "Build a causal DAG from host data and threat signals."""
    # Returns: { "nodes": [...], "edges": [...], "adjacency": {...} }
    # ...
```

The causal DAG distinguishes between: - **Correlation**: two events co-occur (e.g., port scan and SSH brute-force) - **Causation**: one event triggers another (e.g., default credential exposure *causes* unauthorized access)

The do(X) intervention operator is the causal ideal. In practice, the implementation is better understood as intervention-effect estimation under explicit assumptions. This enables the system to recommend targeted interventions and attach an estimated downstream effect, rather than claiming certainty.

### 8.8 Adversarial Robustness

#### 8.8.1 PGD Attacks

The adversarial_training.py module generates adversarial examples using Projected Gradient Descent (Madry et al., 2018):

```
# PGD Attack Step
x_adv = x + epsilon * np.sign(gradient_x(loss(model, x, y)))
```

These adversarial examples are used during training to harden the TransformerIDS against evasion attacks.

#### 8.8.2 Randomized Smoothing

The smoothing.py module wraps the TransformerIDS classifier in a randomized smoothing envelope (Cohen et al., 2019), providing a certified robustness statement: within radius R under the smoothing assumptions, the smoothed classifier's prediction remains unchanged.

### 8.9 Threat Model, Assumptions, and Architecture

graph TD

```
    subgraph "Site A (Hospital)."
    A1[Scan Data] --> A2[Feature Extractor]
    A2 --> A3[Local FL Trainer]
    A3 --> A4[DP Engine: clip + noise]
```

end

```
    subgraph "Site B (Factory)."
    B1[Scan Data] --> B2[Feature Extractor]
    B2 --> B3[Local FL Trainer]
    B3 --> B4[DP Engine: clip + noise]
```

end

subgraph "Aggregation Server."

```
    A4 --> C1[Poison Detector]
```

```
    B4 --> C1
    C1 --> C2[SCAFFOLD Aggregator]
    C2 --> C3[Global Model Registry]
    C3 --> C4[Threat Signal Extractor]
    C4 --> C5[Causal Attribution]
end

    C3 --> A3
    C3 --> B3
    C5 --> D[HYDRA Stream]
```

### 8.9.1 Threat Model and Assumptions

Chapter 8 assumes three classes of risk:

1. **Honest-but-curious aggregation risk.** The aggregator may follow the protocol while still trying to infer information about local training data from model updates.
2. **Minority malicious participants.** A bounded subset of sites may send poisoned updates to distort the shared model or suppress a real threat pattern.
3. **Operational heterogeneity.** Sites differ in device populations, sample counts, hardware speed, and link reliability. The system, therefore, cannot assume IID data, synchronous rounds, or identical local utility curves.

> **Caution:** The techniques described in this chapter have operational boundaries. Always verify assumptions against the specific deployment environment before relying on any automated output.

The chapter's claims depend on several explicit assumptions:

- Differential privacy assurances assume correct clipping, adjusted noise, and an accountant who is not bypassed.
- SCAFFOLD's convergence result assumes the smoothness and bound-variance conditions stated in the paper.
- Randomized smoothing provides a certified radius only under its stated noise model and norm assumptions.
- Causal attribution is only as good as the analyst-specified DAG and the intervention assumptions used to interpret it.

### 8. B Deep Technical Expansion

Executives must ensure that both privacy and integrity controls are in place, as addressing one does not guarantee the other.

A serious federation should assume the coordinator can be curious, compromised, or subpoenaed. Secure aggregation changes the trust posture by making the server combine updates without seeing each one in the clear. That does not solve poisoning. It solves a different problem. It limits what the coordinator can learn from one participant. The doctoral point is that privacy and robustness are separate axes. A system can be private and poisoned. It can be robust and leaky. The design must name both axes before claiming safety.

A short design checklist for distinguishing privacy and robustness controls:

Privacy Controls:

- Secure aggregation: prevents the coordinator from seeing raw updates.
- Differential privacy: adds calibrated noise to prevent data inference.
- Clipping: limits individual contribution and leakage.
- Cohort suppression: hides tiny participant groups.
- Privacy accountant: tracks cumulative disclosure budget.

Robustness Controls:

- Robust aggregation (e.g., Multi-Krum, trimmed mean): excludes outliers and poisoned updates.
- Anomaly detection on updates: flags malicious or abnormal patterns.
- Minority evidence preservation: keeps rare but plausible updates from being erased.
- Quarantine/dissent channels: hold updates for review rather than averaging immediately.
- Provenance and round linkage: ensures update is admissible and timely.

Designers should explicitly state which controls address disclosure and which address integrity. Mixing them can obscure remaining risks.

The practical review should force a two-column answer. One column names disclosure risks: update inversion, site fingerprinting, membership inference, subpoena exposure, and accidental aggregation logs. The other column names integrity risks: label flipping, backdoor updates, replay, sybil participation, and gradient scaling. If a design control appears in the wrong column, the analyst has confused the property. Secure aggregation belongs on the disclosure side. Robust aggregation, clipping, anomaly review, quorum rules, and dissent handling belong on the integrity side.

For IoT and OT, this distinction changes procurement language. A city may federate camera anomaly patterns across agencies if raw video and site-specific updates remain hidden. That same city should still refuse automatic model acceptance if one agency has weak endpoint control. The privacy control answers one objection, but not the poisoning objection. A senior analyst should brief on both facts in less than a minute.

### 8.B2 Replay, staleness, and round identity

A federated update is only meaningful inside a round. Replaying an old update into a new round can cause the aggregate to drift backward, even though all cryptographic signatures still verify. The update record, therefore, needs a round identifier, a model hash, a feature schema hash, a training-window boundary, and an expiry rule. Staleness is not a bookkeeping nuisance but an attack surface.

The reviewer should ask for the tuple that makes an update admissible: participant identity, round ID, parent model digest, feature dictionary digest, local training interval, clipping bound, privacy-accountant state, and

signature. A valid signature without the tuple is not enough. It proves who sent bytes, but not that the bytes belong in this analytic moment.

Staleness is especially dangerous in operational technology. A plant may train during a maintenance bypass. A water utility may train during a storm event. A hospital may train during generator testing. These windows are real but not ordinary. Replaying them into a later round can teach the federation that abnormal control behavior is normal. The fix is not literary but an expiry rule, a round contract, and an audit event when a late update is rejected.

### 8.B3 Minority evidence preservation

The most important site may have the least data. A cargo terminal, rural hospital, or small utility can see the first trace of a campaign before large sites move. Robust aggregation should not erase that trace. The design should create a dissent lane to shadow-test the update, compare it against historical site behavior, and keep it outside the global model until the evidence matures.

The dissent lane is the difference between statistical comfort and operational intelligence. A trimmed mean may discard a rare update because it is far from the majority. This can be correct when the update is malicious, but catastrophic when it is the first sighting. The analyzer should output two products: the global model update and a minority-evidence queue with reasons for non-admission.

A doctoral student should design the queue like a case file. It needs the local feature shift, the affected asset class, the raw-count boundary, the site baseline, the nearest historical analog, and the next collection action. A queue that merely says “outlier” is useless. A queue that says “rare Modbus write cadence on chlorine dosing PLCs after vendor VPN login” gives the next analyst a place to work.

```
queue_entry = {
    "site_id": str,
    "timestamp": float,
    "local_feature_shift": str,    # Summary of embedding change
    "asset_class": str,            # Category of affected asset
    "raw_count": int,              # Number of triggering events
    "site_baseline": dict,         # Historical reference statistics
    "historical_analog": str,      # Description of past similar activity
    "next_action": str,            # Recommended follow-up step
    "reason_for_queueing": str     # e.g., "outlier distance", "low consensus."
}
```

This explicit field list turns the dissent lane into a concrete analytic object, enabling reproducibility and effective handoff for further analysis.

### 8.B4 Privacy budget governance

Privacy budget management should follow change control principles. Each round must have a defined purpose, owner, expected utility, and budget cost. Rounds with minimal value should not consume privacy budget just

because infrastructure is available. The privacy accountant is a governance tool, not just a calculation. Executives should establish clear privacy budget governance policies, including ownership, regular review of consumption and projections, and escalation procedures for overuse or unexpected depletion. Formalizing these steps ensures accountability, transparency, and prompt action when privacy risks arise.

The uncomfortable question is who owns the budget when agencies share a federation. The site that contributes the most sensitive telemetry may not receive the most benefit. A residential child-safety provider may contribute rare camera-abuse signals. A national agency may receive the strongest model lift. That exchange can be legitimate, but only if budget consumption, utility, and purpose are visible to the participants.

> Treat epsilon as a scarce operational resource. Require a round request, expected decision improvement, and post-round utility review. Retire rounds that spend budget without improving detection or triage. This is where world-class cyber analytics differs from model enthusiasm. It asks whether the next training round is worth the privacy it consumes.

### 8.B5 Gradient inversion and membership risk in cyber telemetry

Federated learning is often described as keeping data on site, but that is only part of the story. Model updates can still leak information. A gradient can indicate whether a rare event was present during training. In cyber telemetry, rare events can be the most important. A single device type, command, or badge-reader pattern might identify a facility. Analysts should separate two risks: gradient inversion, which attempts to reconstruct features from updates, and membership inference, which asks whether a record, device, or event was in training. Both risks are higher when events are rare. OT data contains many rare events, such as emergency stops, firmware rollbacks, or manual valve overrides. Rare does not mean unimportant. It means privacy protections must be applied carefully.

The mitigation steps should be clear. Clip updates, add calibrated noise, use secure aggregation, limit how much each site can contribute, hide very small participant groups, and track the total privacy budget. Check whether a released model could allow someone outside to detect a sensitive incident. If the model lets an adversary figure out that a hospital had generator problems last night, the federation has shared information with the wrong people.

### 8.B6 Client selection under operational bias

Client selection looks like scheduling but is actually sampling power. Online, fast, and clean sites appear more often in the global model. Small, overloaded, rural, or bandwidth-constrained sites are excluded from training. The federation then learns the behavior of privileged infrastructure and calls it normal.

This bias matters in government, airport, factory, and residential deployments. A large airport terminal may report every round. A small cargo facility may miss rounds during peak movement. A factory cell may be disconnected during maintenance. A home router may sleep through collection windows. Ignoring these absences makes the model smooth where the defender already sees well and blind where compromise hides.

A chapter-quality answer should demand a participation ledger. Which sites were trained? Which missed? Why? Did their absence change class balance, geography, device families, or attack exposure? A model card for federated cyber analytics should report participation like a clinical trial reports enrollment. Without that ledger, the model may be accurate only for sites that need the least help.

### 8.B7 Rollback discipline for shared models

A shared model needs rollback discipline before the incident. The organization should know which model is active, which is previous, which alerts came from each model, and which downstream actions used the scores. If a poisoning campaign is discovered after deployment, the question is not only how to retrain but how to unwind trust.

Rollback is harder in cyber than in ordinary prediction. Alerts may have opened cases. Tickets may have been closed. Controls may have been tuned. A poisoned model can leave residue in analyst memory and playbooks even after the file is replaced. The audit trail should link model version, training round, participant set, evaluation evidence, alert outputs, and operational decisions.

For doctoral work, the rollback test should be adversarial. Assume the malicious update was accepted three rounds ago. Identify the models that inherited it. Identify the detections that depended on it. Identify the missed detections that may have resulted from it. Identify the sites that should be warned. A federation that cannot answer those questions is not ready for shared operational use.

## 8. A Advanced Practitioner Fieldbook: What Experts Still Miss

This fieldbook expands Chapter 8 for doctoral students and senior cybersecurity practitioners. It assumes the reader knows the vocabulary. The goal is harder. The reader must reason under pressure, preserve evidence boundaries, and explain why the analysis should be trusted when the environment fights back.

The prose uses short sentences on purpose. Short sentences expose weak claims. They also make the material teachable in a live seminar.

### 8.A1 The airport federation that learned the wrong lesson

A hub airport, a regional airport, and a cargo operator observe login bursts against badge kiosks. The hub has the most data. The cargo site has the rarest event. A naive average lets the hub dominate. The global model learns passenger-terminal normality and misses the cargo pattern.

*Figure 8.A1: ch08 fig01 federated architecture.*

Figure 8.1: Federated Learning Architecture

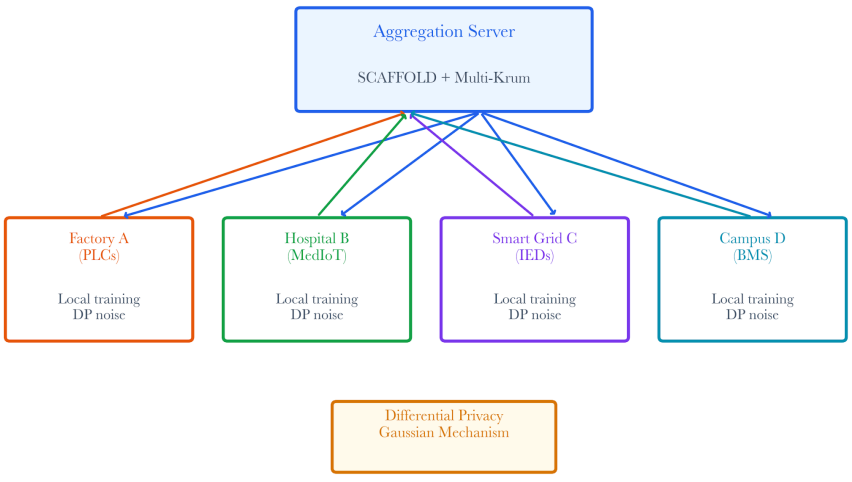


Figure note: The architecture shows where raw telemetry stays local and where model influence crosses the boundary. Use it to challenge custody claims in the airport scenario.

For the airport federation that learned the wrong lesson, the review starts with the site update and ends only when the claim has a named owner. The analyst should state what the system observed, what transformation changed the record, and what conclusion is still forbidden.

The specific danger in the airport federation that learned the wrong lesson is false closure. The privacy budget may support one interpretation, while the aggregator's perspective supports another. The case should remain open until the dissent queue explains why the stronger claim is admissible or why it has been refused.

A senior answer for the airport federation that learned the wrong lesson produces three artifacts: a leadership sentence, a reproducible evidence note, and a hostile objection. The three artifacts should agree on the same boundary. If they do not, the analyst has confused persuasion with proof.

### 8.A2 Privacy budget as operational ammunition

Epsilon is ammunition. Spend it on low-value rounds, and the federation has no budget left when a campaign begins. The hard decision is when to refuse a round.

*Figure 8.A2: ch08 fig04 dp gaussian mechanism.*

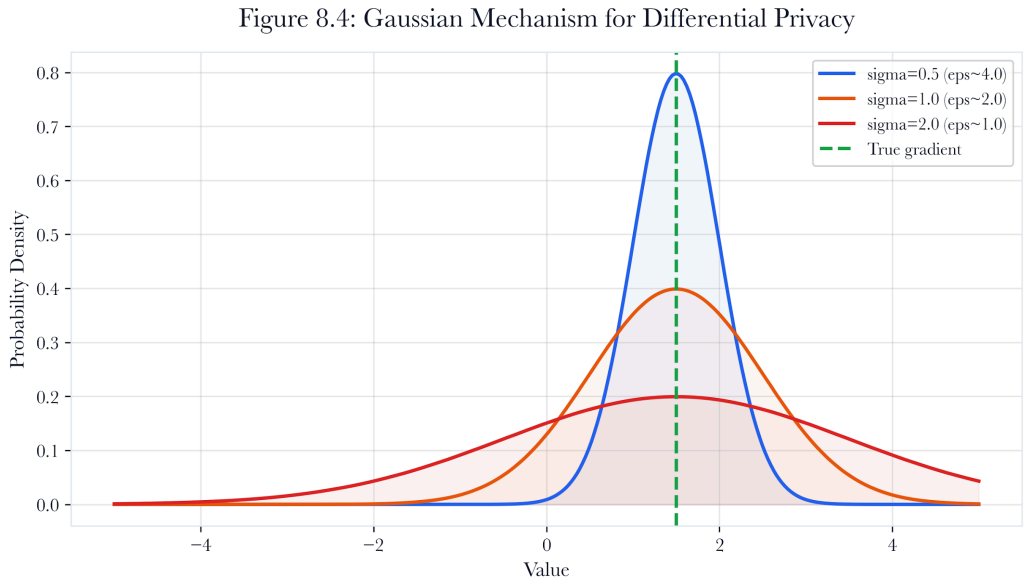


Figure note: The Gaussian mechanism belongs to the privacy budget. It makes the cost of each round visible before the federation spends scarce disclosure risk.

For privacy budget as operational ammunition, the review starts with the site update and ends only when the claim has a named owner. The analyst should state what the system observed, what transformation changed the record, and what conclusion is still forbidden.

The specific danger in privacy budget as operational ammunition is false closure. The privacy budget may support one interpretation, while the aggregator points to another. The case should remain open until the dissent queue explains why the stronger claim is admissible or why it has been refused.

A senior answer for privacy budget as operational ammunition produces three artifacts: a leadership sentence, a reproducible evidence note, and a hostile objection. The three artifacts should agree on the same boundary. If they do not, the analyst has confused persuasion with proof.

### 8.A3 SCAFFOLD and the geometry of disagreement

Non-IID data is the shape of reality. A pediatric hospital, a refinery, and a police evidence building should disagree. SCAFFOLD matters because it separates useful local drift from optimizer drift.

*Figure 8.A3: ch08 fig03 scaffold control variates.*

Figure 8.3: SCAFFOLD vs FedAvg Convergence (Non-IID Data)

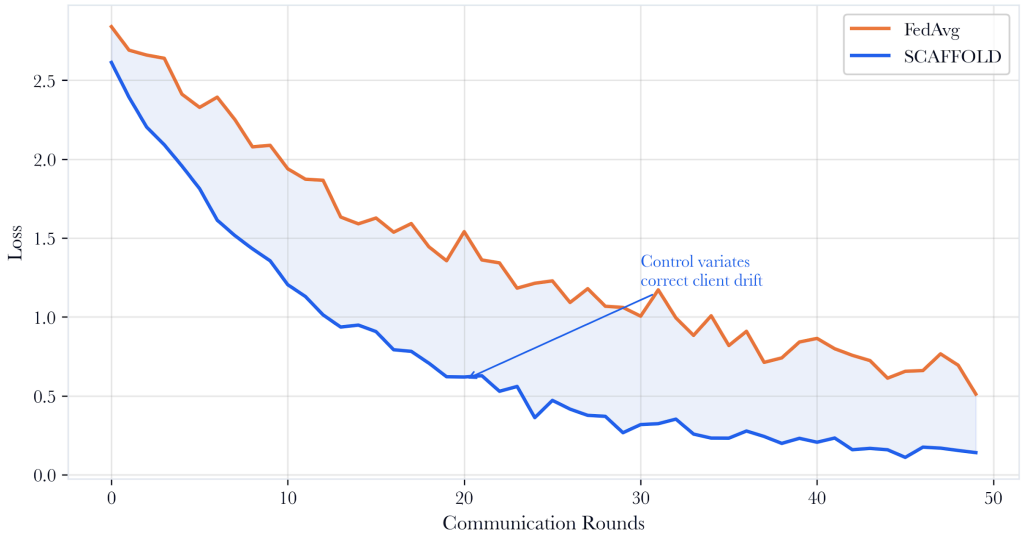


Figure note: SCAFFOLD appears here because the subsection is about disciplined disagreement. The figure shows how control variates correct optimizer drift without erasing local truth.

For the scaffold and the geometry of disagreement, the review starts with the site update and ends only when the claim has a named owner. The analyst should state what the system observed, what transformation changed the record, and what conclusion is still forbidden.

The specific danger in scaffolding and the geometry of disagreement is false closure. The privacy budget may support one interpretation, while the aggregator points to another. The case should remain open until the dissent queue explains why the stronger claim is admissible or why it has been refused.

A senior answer for scaffold and the geometry of disagreement produces three artifacts: a leadership sentence, a reproducible evidence note, and a hostile objection. The three artifacts should agree on the same boundary. If they do not, the analyst has confused persuasion with proof.

### 8.A4 Byzantine updates that look helpful

A poisoned update can improve validation accuracy while teaching the model to ignore one trigger. That is why distance, history, and minority dissent all matter.

*Figure 8.A4: ch08 fig08 multi krum defense.*

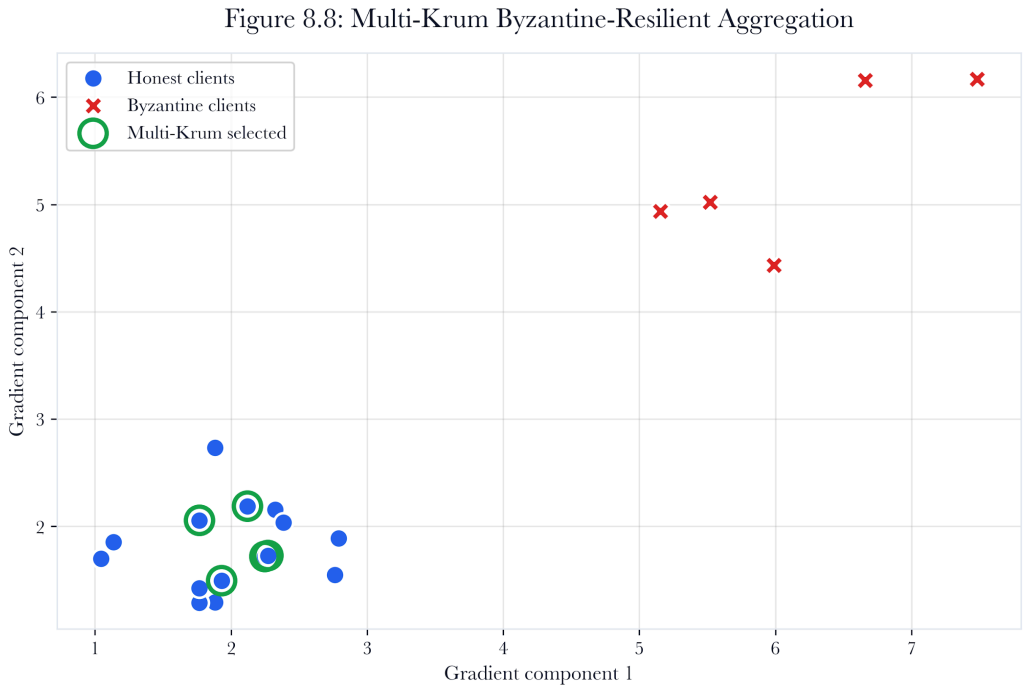


Figure note: Multi-Krum is a robustness control, not a truth oracle. The figure helps ask which helpful-looking update should be admitted, quarantined, or preserved as dissent.

For Byzantine updates that look helpful, the review starts with the site update and ends only when the claim has a named owner. The analyst should state what the system observed, what transformation changed the record, and what conclusion is still forbidden.

The specific danger in Byzantine updates that look helpful is false closure. The privacy budget may support one interpretation, while the aggregator points to another. The case should remain open until the dissent queue explains why the stronger claim is admissible or why it has been refused.

A senior answer for byzantine updates that look helpful produces three artifacts: a leadership sentence, a reproducible evidence note, and a hostile objection. The three artifacts should agree on the same boundary. If they do not, the analyst has confused persuasion with proof.

### 8.A5 Gradient provenance as chain of custody

A model update without provenance is a rumor with decimal. The update requires site identity, feature schema, clipping norm, privacy measure, and training window.

*Figure 8.A5: ch08 fig13 communications rounds.*

Figure 8.13: Communication Rounds vs Client Count

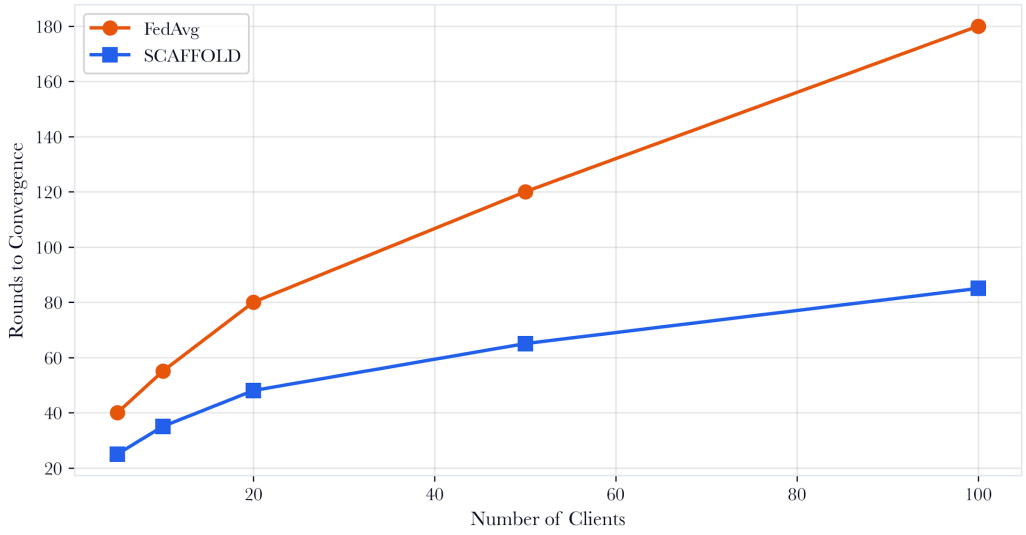


Figure note: Communication rounds carry signatures, schemas, model hashes, and replay risk. The figure provides a gradient provenance with a concrete chain of custody.

For gradient provenance as chain of custody, the review starts with the site update and ends only when the claim has a named owner. The analyst should state what the system observed, what transformation changed the record, and what conclusion is still forbidden.

The specific danger in gradient provenance, in terms of the chain of custody, is false closure. The privacy budget may support one interpretation, while the aggregator points to another. The case should remain open until the dissent queue explains why the stronger claim is admissible or why it has been refused.

A senior answer for gradient provenance as chain of custody produces three artifacts: a leadership sentence, a reproducible evidence note, and a hostile objection. The three artifacts should agree on the same boundary. If they do not, the analyst has confused persuasion with proof.

### 8.A6 TransformerIDS on the edge

Self-attention can learn relationships among ports, services, timing, and asset roles. On an edge gateway, it must also respect memory, latency, and update cadence.

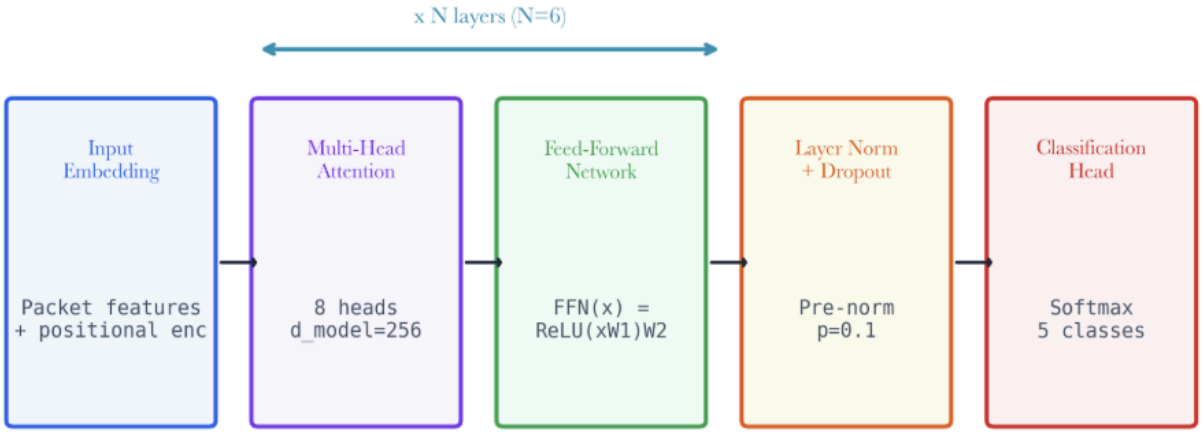


*Figure 8.A6: ch08 fig06 transformer ids architecture.*

Figure note: TransformerIDS needs an edge architecture, not just a model name. The figure ties attention, compression, and gateway constraints to deployable analytics.

For transformer IDs on the edge, the review starts with the site update and ends only when the claim has a named owner. The analyst should state what the system observed, what transformation changed the record, and what conclusion is still forbidden.

The specific danger in transformer IDs on the edge is false closure. The privacy budget may support one interpretation, while the aggregator points to another. The case should remain open until the dissent queue explains why the stronger claim is admissible or why it has been refused.

A senior answer for transformerids on the edge produces three artifacts: a leadership sentence, a reproducible evidence note, and a hostile objection. The three artifacts should agree on the same boundary. If they do not, the analyst has confused persuasion with proof.

### 8.A7 Causal attribution under federated blindness

The federation sees through a privacy veil. Causal reasoning prevents a global alert from becoming superstition.

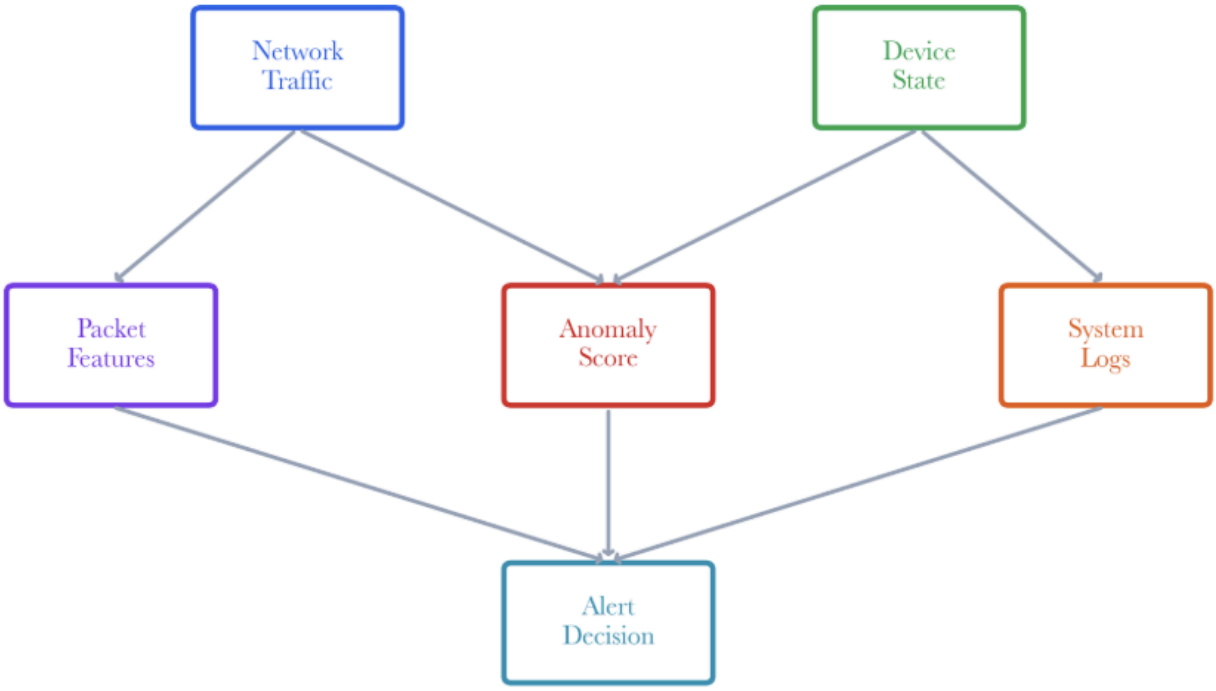


*Figure 8.A7: ch08 fig15 causal dag.*

Figure note: The causal DAG belongs with attribution. It forces the class to separate a shared symptom from a defensible cause under privacy limits.

For causal attribution under federated blindness, the review starts with the site update and ends only when the claim has a named owner. The analyst should state what the system observed, what transformation changed the record, and what conclusion is still forbidden.

The specific danger in causal attribution under federated blindness is false closure. The privacy budget may support one interpretation, while the aggregator points to another. The case should remain open until the dissent queue explains why the stronger claim is admissible or why it has been refused.

A senior answer for causal attribution under federated blindness produces three artifacts: a leadership sentence, a reproducible evidence note, and a hostile objection. The three artifacts should agree on the same boundary. If they do not, the analyst has confused persuasion with proof.

### 8.A8 The quiet client problem

Some clients are honest but silent. Their absence can bias the federation as strongly as a malicious update.

*Figure 8.A8: ch08 fig19 client contribution heatmap.*

Figure 8.19: Client Contribution Heatmap Across Rounds

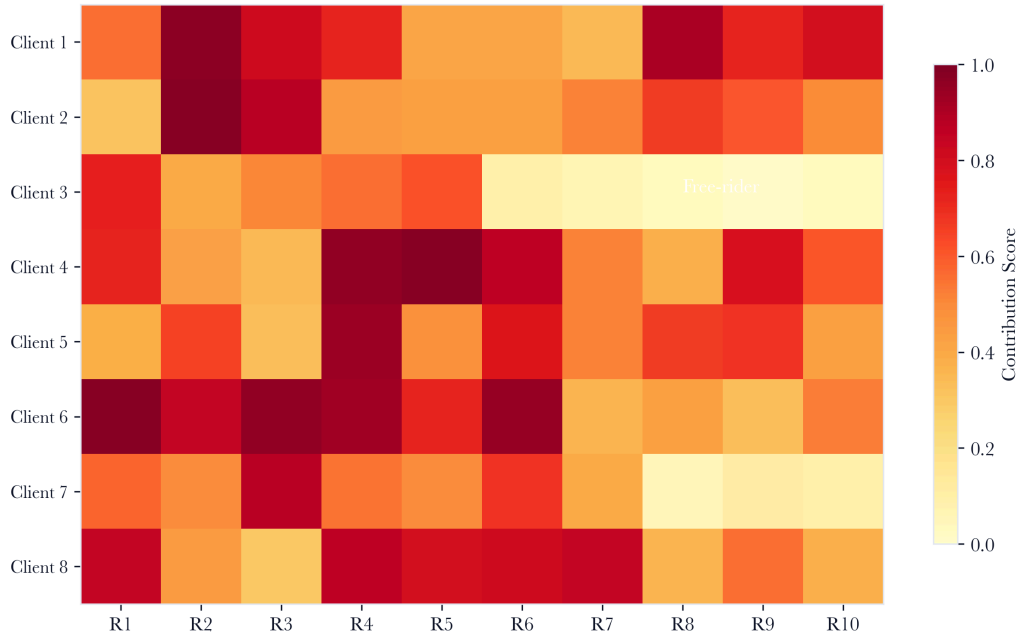


Figure note: The contribution heatmap exposes quiet-client bias. Absence becomes evidence when participation is uneven across sites, sectors, or collection windows.

For the quiet client problem, the review starts with the site update and ends only when the claim has a named owner. The analyst should state what the system observed, what transformation changed the record, and what conclusion is still forbidden.

The specific danger in the quiet client problem is the tendency toward false closure. The privacy budget may support one interpretation, while the aggregator points to another. The case should remain open until the dissent queue explains why the stronger claim is admissible or why it has been refused.

A senior answer for the quiet client problem produces three artifacts: a leadership sentence, a reproducible evidence note, and a hostile objection. The three artifacts should agree on the same boundary. If they do not, the analyst has confused persuasion with proof.

### 8.A9 Poison quarantine as governance

Quarantine is a control state. The update waits until provenance, distance, temporal consistency, and review resolve uncertainty.

In poison quarantine governance, evidence quality is unevenly distributed. The airport operations center may provide a precise site update and still hide the condition that matters. That is why chapter-level mastery requires evidence boundaries, not vocabulary recall.

The question of review is practical: where can an adversary make the system overconfident? The answer usually sits between the privacy accountant, aggregator, and minority evidence. This gap is where proficient students should work.

The section should end with an action. Quarantine, shadow test, narrow the claim, collect one more artifact, or reject the conclusion. A vague recommendation led to analytic failure.

For poison quarantine as governance, the review begins with the site update and ends only when the claim has a named owner. The analyst should state what the system observed, what transformation changed the record, and what conclusion remains forbidden.

The specific danger in poison quarantine, as a form of governance, is the false closure it produces. The privacy budget may support one interpretation, while the aggregator points to another. The case should remain open until the dissent queue explains why the stronger claim is admissible or why it has been refused.

A senior answer for poison quarantine as governance produces three artifacts: a leadership sentence, a reproducible evidence note, and a hostile objection. The three artifacts should agree on the same boundary. If they do not, the analyst has confused persuasion with proof.

### 8.A10 Communication rounds as an attack surface

Every round can teach. Every round can also leak, replay, poison, or exhaust budget.

Communication rounds as an attack surface should change how analysts brief on risk. In the police evidence facility, leadership does not need a tour of the algorithm. They need to know which site update is admissible, which assumption is fragile, and which decision can be taken now.

The system should resist false closure. When the privacy accountant and the minority evidence disagree, the correct state may be unresolved. That is a professional answer, provided the unresolved state has an owner, expiry, and next evidence request.

The highest-grade student will keep model behavior, governance, and mission consequence separate. Merging them into one confidence score destroys the audit trail that makes cyber analytics defensible.

For communication rounds as an attack surface, the review starts with the site update and ends only when the claim has a named owner. The analyst should state what the system observed, what transformation changed the record, and what conclusion is still forbidden.

The specific danger in communication rounds is the creation of an attack surface through false closure. The privacy budget may support one interpretation, while the aggregator points to another. The case should remain open until the dissent queue explains why the stronger claim is admissible or why it has been refused.

A senior answer for communication rounds as an attack surface produces three artifacts: a leadership sentence, a reproducible evidence note, and a hostile objection. The three artifacts should agree on the same boundary. If they do not, the analyst has confused persuasion with proof.

### 8.A11 Feature schemas as treaties

A federation breaks when one site changes a field meaning without telling the others. Schema versioning becomes a treaty.

In the factory cell, feature schemas as treaties become a claim-control problem inside federated analytics. The analyst starts by naming the observed site update, the transformation that touched it, and the smallest defensible conclusion. Anything stronger has to wait for evidence.

The failure path is subtle. The dashboard consolidates privacy accountant, aggregator, and minority evidence into a single status. A senior reviewer should split them apart and preserve the disagreement as a first-class record.

The stop rule matters. The subsection should leave the reader able to say what blocks the claim, what promotes it, and what keeps it in review. That discipline is the difference between analytics and decorative scoring.

For feature schemas as treaties, the review starts with the site update and ends only when the claim has a named owner. The analyst should state what the system observed, what transformation changed the record, and what conclusion is still forbidden.

The specific danger in feature schemas as treaties is the risk of false closure. The privacy budget may support one interpretation, while the aggregator points to another. The case should remain open until the dissent queue explains why the stronger claim is admissible or why it has been refused.

A senior answer for feature schemas as treaties produces three artifacts: a leadership sentence, a reproducible evidence note, and a hostile objection. The three artifacts should agree on the same boundary. If they do not, the analyst has confused persuasion with proof.

### 8.A12 The analyst who trusts consensus too much

Consensus is not truth. Five compromised sites can agree. One clean site can disagree.

Treat the analyst who trusts consensus too much as an adversarial hearing, not a feature description. The municipal water site gives the system partial data, delayed data, and political pressure. The answer must still bind the site update to a named decision.

A strong implementation keeps the uncomfortable edge visible. If the privacy accountant says proceed while the minority evidence says wait, the system should not average the conflict into confidence. It should record the conflict and assign ownership.

The doctoral move is to seek the counterfactual. Which observation would make the first claim false? If the workflow cannot answer, it has built a belief engine rather than a cyberanalytic control.

For the analyst who trusts consensus too much, the review starts with the site update and ends only when the claim has a named owner. The analyst should state what the system observed, what transformation changed the record, and what conclusion is still forbidden.

The specific danger in the analyst who trusts consensus too much is false closure. The privacy budget may support one interpretation, while the aggregator points to another. The case should remain open until the dissent queue explains why the stronger claim is admissible or why it has been refused.

A senior answer for the analyst who trusts consensus too much produces three artifacts: a leadership sentence, a reproducible evidence note, and a hostile objection. The three artifacts should agree on the same boundary. If they do not, the analyst has confused persuasion with proof.

### 8.A13 Privacy utility frontier for executives

Executives want one number. The federation gives a curve. More privacy can reduce sensitivity. More sensitivity can raise disclosure risk.

The operational test for the privacy utility frontier for executives is whether a second expert can replay the reasoning. In a residential safety platform, this means the input record, transformation, exception handling, and decision authority must survive handoff.

The dangerous shortcut is to trust the most convenient signal. An aggregator may look stable while minority evidence carries the real warning. A privacy accountant may look mathematically clean, while the mission context changes the cost of error.

A useful report should be modest and sharp. It should say what the analyst saw, what it inferred, what it refused to infer, and which collection step would move the case forward.

For the privacy utility frontier for executives, the review starts with the site update and ends only when the claim has a named owner. The analyst should state what the system observed, what transformation changed the record, and what conclusion is still forbidden.

The specific danger in the privacy utility frontier for executives is false closure. The privacy budget may support one interpretation, while the aggregator points to another. The case should remain open until the dissent queue explains why the stronger claim is admissible or why it has been refused.

The architecture should clearly show where updates are clipped, where noise is added, where they are signed, aggregated, quarantined, and published.

For a deployment architecture as a safety case, evidence quality is not evenly distributed. The regional hospital may provide a precise site update and still hide the condition that matters. That is why chapter-level mastery requires evidence boundaries, not vocabulary recall.

The question of review is practical: where can an adversary make the system overconfident? The answer usually sits between the privacy accountant, aggregator, and minority evidence. This gap is where proficient students should work.

The section should end with an action. Quarantine, shadow test, narrow the claim, collect one more artifact, or reject the conclusion. A vague recommendation led to analytic failure.

For deployment architecture as a safety case, the review begins with the site update and ends only when the claim has a named owner. The analyst should state what the system observed, what transformation changed the record, and which conclusion is still forbidden.

The specific danger in deployment architecture as a safety case is false closure. The privacy budget may support one interpretation, while the aggregator points to another. The case should remain open until the dissent queue explains why the stronger claim is admissible or why it has been refused.

A senior answer to the deployment architecture safety case produces three artifacts: a leadership sentence, a reproducible evidence note, and a hostile objection. The three artifacts should agree on the same boundary. If they do not, the analyst has confused persuasion with proof.

### 8.A15 Doctoral Mastery Check

A student has mastered Chapter 8 when they can defend a federation without pretending that privacy, robustness, and governance are the same property. The oral exam should force them to state the round contract, participant set, clipping rule, privacy spend, aggregation rule, dissent path, and rollback plan.

Executive Checklist: Assessing Federation Readiness

Executives should be able to ask their teams the following questions to ensure federation requirements are being met in operational practice:

- Is the round contract for every training cycle clear, with assigned ownership and expiry?
- Are all participants and their roles recorded and tracked in each round?
- Is there an explicit clipping rule and per-round privacy spend documented and reviewed?
- Does the aggregation process include robust aggregation and clear criteria for admissibility and dissent?
- Are minority or dissenting updates preserved, quarantined, or shadow-tested rather than immediately discarded?
- Can the team identify and execute a rollback if a poisoned or invalid update is later discovered?
- Is the privacy budget monitored, governed, and communicated to all stakeholders?
- Are evidence boundaries maintained so that operational decisions, analytic findings, and governance records remain distinct?

If a team cannot give specific, actionable answers to these questions, it risks taking shortcuts and creating uncontrolled trust.

Mastery is assessed using a structured rubric. Students must:

- Clearly distinguish between privacy, robustness, and governance controls in their design explanations.
- Articulate the operational significance and interactions among round contract, participant tracking, and privacy spending.
- Present a complete, step-by-step federation decision scenario, including identification of dissent and rollback triggers.
- Respond to adversarial oral exam questions that challenge their analytic boundary, evidence preservation, and control path.

Sample oral exam question: "Given a scenario where a minority site produces an outlier update that is flagged by robust aggregation but within the privacy budget, explain your admissibility decision, defend your handling of both consensus and dissent, and state precisely how rollback would be executed if the update was later found to be malicious."

Students should use the rubric and sample scenario to self-assess readiness for doctoral-level defense.

Perturb the case with five domain shocks: poison one participant, make one small site the first witness, make one update stale, make one schema version drift, and exhaust half the privacy budget before the real campaign begins. The student should narrow the claim, preserve minority evidence, and name the next collection action.

### 8.A16 Expert Failure Review Drill

A final exercise bridges the gap between technical skill and sound judgment. Present a federation that appears healthy: five sites submit updates, the loss curve improves, global alert volume decreases, and the privacy accountant reports an acceptable budget. Then, it was revealed that only the smallest site detected a new attacker pattern, but its update was dismissed as an outlier due to stability at the other sites. This highlights the risk of optimizing for the federation average and overlooking critical minority signals.

The solution is not to accept every outlier, which would be naive. Instead, maintain a dissent channel. Rare updates should be quarantined, reweighted, shadow-tested, or sent for analyst review without compromising the global model. This approach preserves minority evidence while maintaining safety, distinguishing robust aggregation from routine averaging.

Advanced reporting should be precise. For example: “The global model improved under aggregate validation, but one low-volume site produced a temporally novel pattern that remains unresolved.” This statement provides leadership with clarity on knowns, unknowns, and next steps. Strong students should also recommend controls, such as a dissent buffer for minority-site patterns, a replay mechanism to test patterns against previous models, and governance rules to prevent deletion of rare evidence. These measures transform a fragile federation into a learning system with operational memory. The final report should document dissent as a distinct analytic object, specify who can reopen the case, identify which telemetry could change the outcome, and state when the buffer expires.

### 8. C Advanced Case Problems for Expert Readers

These cases add doctoral-level friction to Chapter 8. They are written for readers who already know the tools. The work is to hold evidence, uncertainty, mission consequence, and adversary pressure in the same frame without collapsing them into a slogan.

### 8.C1 Mutual-aid federation after a ransomware wave

In a mutual-aid federation after a ransomware wave, the visible event is only the entry point. The municipal water site has a normal operating story, a procurement or operational constraint, and at least one actor who benefits from premature certainty. The analyst must decide whether the observed site update supports action, review, quarantine, or refusal.

The case becomes difficult when a privacy accountant appears to support one decision while the aggregator and minority evidence support another. A weak answer chooses the comfortable signal. A strong answer names the conflict, narrows the claim, and states what evidence would change the decision. This is where 25-year practitioners can still disagree in useful ways.

The expected written product is short but severe: one executive sentence, one technical evidence paragraph, one adversarial objection, and one next collection step. The sentence must be safe for leadership. The technical paragraph must be reproducible. The objection must be hostile enough to catch overclaiming. The next step must be operationally possible in an IoT or OT environment.

### 8.C2 Airport cargo minority signal

In the airport cargo minority signal, the visible event is only the entry point. The residential safety platform has a normal operating story, a procurement or operational constraint, and at least one actor who benefits from premature certainty. The analyst must decide whether the observed site update supports action, review, quarantine, or refusal.

The case becomes difficult when a privacy accountant appears to support one decision while the aggregator and minority evidence support another. A weak answer chooses the comfortable signal. A strong answer names the conflict, narrows the claim, and states what evidence would change the decision. This is where 25-year practitioners can still disagree in useful ways.

The expected written product is short but severe: one executive sentence, one technical evidence paragraph, one adversarial objection, and one next collection step. The sentence must be safe for leadership. The technical

paragraph must be reproducible. The objection must be hostile enough to catch overclaiming. The next step must be operationally possible in an IoT or OT environment.

### 8.C3 Hospital telemetry under privacy pressure

In hospital telemetry under privacy pressure, the visible event is only the entry point. The regional hospital has a normal operating story, a procurement or operational constraint, and at least one actor who benefits from premature certainty. The analyst must decide whether the observed site update supports action, review, quarantine, or refusal.

The case becomes difficult when a privacy accountant appears to support one decision while the aggregator and minority evidence support another. A weak answer chooses the comfortable signal. A strong answer names the conflict, narrows the claim, and states what evidence would change the decision. This is where 25-year practitioners can still disagree in useful ways.

The expected written product is short but severe: one executive sentence, one technical evidence paragraph, one adversarial objection, and one next collection step. The sentence must be safe for leadership. The technical paragraph must be reproducible. The objection must be hostile enough to catch overclaiming. The next step must be operationally possible in an IoT or OT environment.

### 8.C4 Factory edge model with schema drift

In the factory edge model with schema drift, the visible event is only the entry point. The border checkpoint has a normal operating story, a procurement or operational constraint, and at least one actor who benefits from premature certainty. The analyst must decide whether the observed site update supports action, review, quarantine, or refusal.

The case becomes difficult when a privacy accountant appears to support one decision while the aggregator and minority evidence support another. A weak answer chooses the comfortable signal. A strong answer names the conflict, narrows the claim, and states what evidence would change the decision. This is where 25-year practitioners can still disagree in useful ways.

The expected written product is short but severe: one executive sentence, one technical evidence paragraph, one adversarial objection, and one next collection step. The sentence must be safe for leadership. The technical paragraph must be reproducible. The objection must be hostile enough to catch overclaiming. The next step must be operationally possible in an IoT or OT environment.

### 8.C5 Residential camera federation and child safety

In residential camera federation and child safety, the visible event is only the entry point. The maritime terminal has a normal operating story, a procurement or operational constraint, and at least one actor who benefits from premature certainty. The analyst must decide whether the observed site update supports action, review, quarantine, or refusal.

The case becomes difficult when a privacy accountant appears to support one decision while the aggregator and minority evidence support another. A weak answer chooses the comfortable signal. A strong answer names the conflict, narrows the claim, and states what evidence would change the decision. This is where 25-year practitioners can still disagree in useful ways.

The expected written product is short but severe: one executive sentence, one technical evidence paragraph, one adversarial objection, and one next collection step. The sentence must be safe for leadership. The technical

paragraph must be reproducible. The objection must be hostile enough to catch overclaiming. The next step must be operationally possible in an IoT or OT environment.

### 8.C6 Police facility with a poisoned participant

In a police facility with a poisoned participant, the visible event is only the entry point. The airport operations center has a normal operating story, a procurement or operational constraint, and at least one actor who benefits from premature certainty. The analyst must decide whether the observed site update supports action, review, quarantine, or refusal.

The case becomes difficult when a privacy accountant appears to support one decision while the aggregator and minority evidence support another. A weak answer chooses the comfortable signal. A strong answer names the conflict, narrows the claim, and states what evidence would change the decision. This is where 25-year practitioners can still disagree in useful ways.

The expected written product is short but severe: one executive sentence, one technical evidence paragraph, one adversarial objection, and one next collection step. The sentence must be safe for leadership. The technical paragraph must be reproducible. The objection must be hostile enough to catch overclaiming. The next step must be operationally possible in an IoT or OT environment.

### 8.C7 Utility storm mode as non-IID reality

In utility storm mode as a non-iid reality, the visible event is only the entry point. The police evidence facility has a normal operating story, a procurement or operational constraint, and at least one actor who benefits from premature certainty. The analyst must decide whether the observed site update supports action, review, quarantine, or refusal.

The case becomes difficult when a privacy accountant appears to support one decision while the aggregator and minority evidence support another. A weak answer chooses the comfortable signal. A strong answer names the conflict, narrows the claim, and states what evidence would change the decision. This is where 25-year practitioners can still disagree in useful ways.

The expected written product is short but severe: one executive sentence, one technical evidence paragraph, one adversarial objection, and one next collection step. The sentence must be safe for leadership. The technical paragraph must be reproducible. The objection must be hostile enough to catch overclaiming. The next step must be operationally possible in an IoT or OT environment.

### 8.C8 Rollback after a bad global round

In a rollback after a bad global round, the visible event is only the entry point. The factory cell has a normal operating story, a procurement or operational constraint, and at least one actor who benefits from premature certainty. The analyst must decide whether the observed site update supports action, review, quarantine, or refusal.

The case becomes difficult when a privacy accountant appears to support one decision while the aggregator and minority evidence support another. A weak answer chooses the comfortable signal. A strong answer names the conflict, narrows the claim, and states what evidence would change the decision. This is where 25-year practitioners can still disagree in useful ways.

The expected written product is short but severe: one executive sentence, one technical evidence paragraph, one adversarial objection, and one next collection step. The sentence must be safe for leadership. The technical

paragraph must be reproducible. The objection must be hostile enough to catch overclaiming. The next step must be operationally possible in an IoT or OT environment.

### 8.C9 Budget exhaustion before a campaign

In budget exhaustion before a campaign, the visible event is only the entry point. The municipal water site has a normal operating story, a procurement or operational constraint, and at least one actor who benefits from premature certainty. The analyst must decide whether the observed site update supports action, review, quarantine, or refusal.

The case becomes difficult when a privacy accountant appears to support one decision while the aggregator and minority evidence support another. A weak answer chooses the comfortable signal. A strong answer names the conflict, narrows the claim, and states what evidence would change the decision. This is where 25-year practitioners can still disagree in useful ways.

The expected written product is short but severe: one executive sentence, one technical evidence paragraph, one adversarial objection, and one next collection step. The sentence must be safe for leadership. The technical paragraph must be reproducible. The objection must be hostile enough to catch overclaiming. The next step must be operationally possible in an IoT or OT environment.

### 8.C10 Cross-agency governance hearing

In a cross-agency governance hearing, the visible event is only the entry point. The residential safety platform has a normal operating story, a procurement or operational constraint, and at least one actor who benefits from premature certainty. The analyst must decide whether the observed site update supports action, review, quarantine, or refusal.

The case becomes difficult when a privacy accountant appears to support one decision while the aggregator and minority evidence support another. A weak answer chooses the comfortable signal. A strong answer names the conflict, narrows the claim, and states what evidence would change the decision. This is where 25-year practitioners can still disagree in useful ways.

The expected written product is short but severe: one executive sentence, one technical evidence paragraph, one adversarial objection, and one next collection step. The sentence must be safe for leadership. The technical paragraph must be reproducible. The objection must be hostile enough to catch overclaiming. The next step must be operationally possible in an IoT or OT environment.

## 8.D Seminar War Rooms

This section is intentionally demanding. Each war room asks the reader to work like a senior analyst responsible for federated model decisions in live IoT or OT environments. The task is not to recite the chapter. The task is to make a bounded claim under pressure, then defend the evidence boundary when another expert attacks it.

### 8.D1 Airport Mutual Aid

**War Room 1** begins with the airport mutual aid. The first analyst wants to accept the participant update because it fits the expected story. The second analyst refuses to move until the privacy budget is tied to an observed artifact, an owner, and a clock. The disagreement is productive. It prevents the team from treating a plausible explanation as verified.

The class should interrogate the case by changing one fact that a real adversary could influence:
- timing

● identity
● provenance
● replay
● exception handling
● sampling
● maintenance state
● operator pressure. The answer must state which claim survives the change and which claim fails. A doctoral answer does not hide behind caution. It chooses a precise operational state:
● admit
● hold
● shadow-test
● collect
● rollback
● or reject.

The final artifact for the airport mutual aid is a compact evidence packet. It contains the following:
● The input record
● Transformation
● Unresolved conflict
● Decision authority
● Dissent queue. It also contains one leadership sentence that is weaker than the analyst wants and stronger than a lawyer would fear. That sentence is the craft. It carries enough truth to guide action without smuggling in facts the system did not prove.

For airport mutual aid, change one adversary-controlled fact:
● timing
● identity
● provenance
● replay
● exception handling
● sampling
● maintenance state
● operator pressure. The answer must state which claim survives and which claim fails. A doctoral answer chooses a precise operational state rather than hiding behind generic caution.

The final artifact for airport mutual aid is a compact evidence packet. It contains the following:
● The input record
● Transformation
● Unresolved conflict
● Decision authority
● Recovery or hold gate. It also contains one leadership sentence that guides action without smuggling in facts that the system did not prove.

**War Room 2** begins with the hospital badge telemetry. The first analyst wants to accept the participant update because it fits the expected story. The second analyst refuses to move until the privacy budget is tied to an observed artifact, an owner, and a clock. The disagreement is productive. It prevents the team from treating a plausible explanation as verified.

The class should challenge the case by changing one fact that a real adversary could influence:
● timing

● identity
● provenance
● replay
● exception handling
● sampling
● maintenance state
● operator pressure. The answer must state which claim survives the change and which claim fails. A doctoral answer does not hide behind caution. It chooses a precise operational state:
● admit
● hold
● shadow-test
● collect
● rollback
● or reject.

The final artifact for the hospital badge telemetry is a compact evidence packet. It contains the following:
● The input record
● Transformation
● Unresolved conflict
● Decision authority
● Dissent queue. It also contains one leadership sentence that is weaker than the analyst wants and stronger than a lawyer would fear. That sentence is the craft. It carries enough truth to guide action without smuggling in facts the system did not prove.

For hospital badge telemetry, change one adversary-controlled fact:
● timing
● identity
● provenance
● replay
● exception handling
● sampling
● maintenance state
● operator pressure. The answer must state which claim survives and which claim fails. A doctoral answer chooses a precise operational state rather than hiding behind generic caution.

The final artifact for hospital badge telemetry is a compact evidence packet. It contains the following:
● The input record
● Transformation
● Unresolved conflict
● Decision authority
● Recovery or hold gate. It also contains one leadership sentence that guides action without smuggling in facts that the system did not prove.

**War Room 3** begins with the Factory Gateway Alliance. The first analyst wants to accept the participant update because it fits the expected story. The second analyst refuses to move until the privacy budget is tied to an observed artifact, an owner, and a clock. The disagreement is productive. It prevents the team from treating a plausible explanation as verified.

The class should bound the case by changing one fact that a real adversary could influence:
● timing

● identity
● provenance
● replay
● exception handling
● sampling
● maintenance state
● operator pressure. The answer must state which claim survives the change and which claim fails. A doctoral answer does not hide behind caution. It chooses a precise operational state:
● admit
● hold
● shadow-test
● collect
● rollback
● or reject.

The final artifact for the factory gateway alliance is a compact evidence packet. It contains the following:
● The input record
● Transformation
● Unresolved conflict
● Decision authority
● Dissent queue. It also contains one leadership sentence that is weaker than the analyst wants and stronger than a lawyer would fear. That sentence is the craft. It carries enough truth to guide action without smuggling in facts the system did not prove.

For the factory gateway alliance, change one adversary-controlled fact:
● timing
● identity
● provenance
● replay
● exception handling
● sampling
● maintenance state
● operator pressure. The answer must state which claim survives and which claim fails. A doctoral answer chooses a precise operational state rather than hiding behind generic caution.

The final artifact for the factory gateway alliance is a compact evidence packet. It contains the following:
● The input record
● Transformation
● Unresolved conflict
● Decision authority
● Recovery or hold gate. It also contains one leadership sentence that guides action without smuggling in facts that the system did not prove.

**War Room 4** begins with the Residential Camera Consortium. The first analyst wants to accept the participant update because it fits the expected story. The second analyst refuses to move until the privacy budget is tied to an observed artifact, an owner, and a clock. The disagreement is productive. It prevents the team from treating a plausible explanation as verified.

The class should quarantine the case by changing one fact that a real adversary could influence:
● timing

● identity
● provenance
● replay
● exception handling
● sampling
● maintenance state
● operator pressure. The answer must state which claim survives the change and which claim fails. A doctoral answer does not hide behind caution. It chooses a precise operational state:
● admit
● hold
● shadow-test
● collect
● rollback
● or reject.

The final artifact for the residential camera consortium is a compact evidence packet. It contains the following:
● The input record
● Transformation
● Unresolved conflict
● Decision authority
● Dissent queue. It also contains one leadership sentence that is weaker than the analyst wants and stronger than a lawyer would fear. That sentence is the craft. It carries enough truth to guide action without smuggling in facts the system did not prove.

For the residential camera consortium, change one adversary-controlled fact:
● timing
● identity
● provenance
● replay
● exception handling
● sampling
● maintenance state
● operator pressure. The answer must state which claim survives and which claim fails. A doctoral answer chooses a precise operational state rather than hiding behind generic caution.

The final artifact for the residential camera consortium is a compact evidence packet. It contains the following:
● The input record
● Transformation
● Unresolved conflict
● Decision authority
● Recovery or hold gate. It also contains one leadership sentence that guides action without smuggling in facts that the system did not prove.

**War Room 5** begins with the water utility coalition. The first analyst wants to accept the participant update because it fits the expected story. The second analyst refuses to move until the privacy budget is tied to an observed artifact, an owner, and a clock. The disagreement is productive. It prevents the team from treating a plausible explanation as verified.

The class should simulate the case by changing one fact that a real adversary could influence:
● timing

● identity
● provenance
● replay
● exception handling
● sampling
● maintenance state
● operator pressure. The answer must state which claim survives the change and which claim fails. A doctoral answer does not hide behind caution. It chooses a precise operational state:
● admit
● hold
● shadow-test
● collect
● rollback
● or reject.

The final artifact for the water utility coalition is a compact evidence packet. It contains the following:
● The input record
● Transformation
● Unresolved conflict
● Decision authority
● Dissent queue. It also contains one leadership sentence that is weaker than the analyst wants and stronger than a lawyer would fear. That sentence is the craft. It carries enough truth to guide action without smuggling in facts the system did not prove.

For the water utility coalition, change one adversary-controlled fact:
● timing
● identity
● provenance
● replay
● exception handling
● sampling
● maintenance state
● operator pressure. The answer must state which claim survives and which claim fails. A doctoral answer chooses a precise operational state rather than hiding behind generic caution.

The final artifact for the water utility coalition is a compact evidence packet. It contains the following:
● The input record
● Transformation
● Unresolved conflict
● Decision authority
● Recovery or hold gate. It also contains one leadership sentence that guides action without smuggling in facts that the system did not prove.

**War Room 6** begins with the police facility sharing group. The first analyst wants to accept the participant update because it fits the expected story. The second analyst refuses to move until the privacy budget is tied to an observed artifact, an owner, and a clock. The disagreement is productive. It prevents the team from treating a plausible explanation as verified.

The class should replay the case by changing one fact that a real adversary could influence:
● timing

● identity
● provenance
● replay
● exception handling
● sampling
● maintenance state
● operator pressure. The answer must state which claim survives the change and which claim fails. A doctoral answer does not hide behind caution. It chooses a precise operational state:
● admit
● hold
● shadow-test
● collect
● rollback
● or reject.

The final artifact for the police facility sharing group is a compact evidence packet. It contains the following:
● The input record
● Transformation
● Unresolved conflict
● Decision authority
● Dissent queue. It also contains one leadership sentence that is weaker than the analyst wants and stronger than a lawyer would fear. That sentence is the craft. It carries enough truth to guide action without smuggling in facts the system did not prove.

For the police facility sharing group, change one adversary-controlled fact:
● timing
● identity
● provenance
● replay
● exception handling
● sampling
● maintenance state
● operator pressure. The answer must state which claim survives and which claim fails. A doctoral answer chooses a precise operational state rather than hiding behind generic caution.

The final artifact for the police facility sharing group is a compact evidence packet. It contains the following:
● The input record
● Transformation
● Unresolved conflict
● Decision authority
● Recovery or hold gate. It also contains one leadership sentence that guides action without smuggling in facts that the system did not prove.

**War Room 7** begins with the Port Authority Federation. The first analyst wants to accept the participant update because it fits the expected story. The second analyst refuses to move until the privacy budget is tied to an observed artifact, an owner, and a clock. The disagreement is productive. It prevents the team from treating a plausible explanation as verified.

The class should audit the case by changing one fact that a real adversary could influence:
● timing

● identity
● provenance
● replay
● exception handling
● sampling
● maintenance state
● operator pressure. The answer must state which claim survives the change and which claim fails. A doctoral answer does not hide behind caution. It chooses a precise operational state:
● admit
● hold
● shadow-test
● collect
● rollback
● or reject.

The final artifact for the port authority federation is a compact evidence packet. It contains the following:
● The input record
● Transformation
● Unresolved conflict
● Decision authority
● Dissent queue. It also contains one leadership sentence that is weaker than the analyst wants and stronger than a lawyer would fear. That sentence is the craft. It carries enough truth to guide action without smuggling in facts the system did not prove.

For the port authority federation, change one adversary-controlled fact:
● timing
● identity
● provenance
● replay
● exception handling
● sampling
● maintenance state
● operator pressure. The answer must state which claim survives and which claim fails. A doctoral answer chooses a precise operational state rather than hiding behind generic caution.

The final artifact for the Port Authority Federation is a compact evidence packet. It contains the following:
● The input record
● Transformation
● Unresolved conflict
● Decision authority
● Recovery or hold gate. It also contains one leadership sentence that guides action without smuggling in facts that the system did not prove.

**War Room 8** begins with the school safety network. The first analyst wants to accept the participant update because it fits the expected story. The second analyst refuses to move until the privacy budget is tied to an observed artifact, an owner, and a clock. The disagreement is productive. It prevents the team from treating a plausible explanation as verified.

The class should stage the case by changing one fact that a real adversary could influence:
● timing

● identity
● provenance
● replay
● exception handling
● sampling
● maintenance state
● operator pressure. The answer must state which claim survives the change and which claim fails. A doctoral answer does not hide behind caution. It chooses a precise operational state:
● admit
● hold
● shadow-test
● collect
● rollback
● or reject.

The final artifact for the school safety network is a compact evidence packet. It contains the following:
● The input record
● Transformation
● Unresolved conflict
● Decision authority
● Dissent queue. It also contains one leadership sentence that is weaker than the analyst wants and stronger than a lawyer would fear. That sentence is the craft. It carries enough truth to guide action without smuggling in facts the system did not prove.

For the school safety network, change one adversary-controlled fact:
● timing
● identity
● provenance
● replay
● exception handling
● sampling
● maintenance state
● operator pressure. The answer must state which claim survives and which claim fails. A doctoral answer chooses a precise operational state rather than hiding behind generic caution.

The final artifact for the school safety network is a compact evidence packet. It contains the following:
● The input record
● Transformation
● Unresolved conflict
● Decision authority
● Recovery or hold gate. It also contains one leadership sentence that guides action without smuggling in facts that the system did not prove.

**War Room 9** begins with the rail operator exchange. The first analyst wants to accept the participant update because it fits the expected story. The second analyst refuses to move until the privacy budget is tied to an observed artifact, an owner, and a clock. The disagreement is productive. It prevents the team from treating a plausible explanation as verified.

The class should defer the case by changing one fact that a real adversary could influence:
● timing

● identity
● provenance
● replay
● exception handling
● sampling
● maintenance state
● operator pressure. The answer must state which claim survives the change and which claim fails. A doctoral answer does not hide behind caution. It chooses a precise operational state:
● admit
● hold
● shadow-test
● collect
● rollback
● or reject.

The final artifact for the rail operator exchange is a compact evidence packet. It contains the following:
● The input record
● Transformation
● Unresolved conflict
● Decision authority
● Dissent queue. It also contains one leadership sentence that is weaker than the analyst wants and stronger than a lawyer would fear. That sentence is the craft. It carries enough truth to guide action without smuggling in facts the system did not prove.

For rail operator exchange, change one adversary-controlled fact:
● timing
● identity
● provenance
● replay
● exception handling
● sampling
● maintenance state
● operator pressure. The answer must state which claim survives and which claim fails. A doctoral answer chooses a precise operational state rather than hiding behind generic caution.

The final artifact for the rail operator exchange is a compact evidence packet. It contains the following:
● The input record
● Transformation
● Unresolved conflict
● Decision authority● Recovery or hold gate. It also contains one leadership sentence that guides action without smuggling in facts that the system did not prove.

**War Room 10** begins with the city emergency federation. The first analyst wants to accept the participant update because it fits the expected story. The second analyst refuses to move until the privacy budget is tied to an observed artifact, an owner, and a clock. The disagreement is productive. It prevents the team from treating a plausible explanation as verified.

The class should escalate the case by changing one fact that a real adversary could influence:
● timing
● identity

● provenance
● replay
● exception handling
● sampling
● maintenance state
● operator pressure. The answer must state which claim survives the change and which claim fails. A doctoral answer does not hide behind caution. It chooses a precise operational state:
● admit
● hold
● shadow-test
● collect
● rollback
● or reject.

The final artifact for the city emergency federation is a compact evidence packet. It contains the following:
● The input record
● Transformation
● Unresolved conflict
● Decision authority
● Dissent queue. It also contains one leadership sentence that is weaker than the analyst wants and stronger than a lawyer would fear. That sentence is the craft. It carries enough truth to guide action without smuggling in facts the system did not prove.

For the city emergency federation, change one adversary-controlled fact:
● timing
● identity
● provenance
● replay
● exception handling
● sampling
● maintenance state
● operator pressure. The answer must state which claim survives and which claim fails. A doctoral answer chooses a precise operational state rather than hiding behind generic caution.

The final artifact for the city emergency federation is a compact evidence packet. It contains the following:
● The input record
● Transformation
● Unresolved conflict
● Decision authority
● Recovery or hold gate. It also contains one leadership sentence that guides action without smuggling in facts that the system did not prove.

**War Room 11** begins with the border checkpoint consortium. The first analyst wants to accept the participant update because it fits the expected story. The second analyst refuses to move until the privacy budget is tied to an observed artifact, an owner, and a clock. The disagreement is productive. It prevents the team from treating a plausible explanation as verified.

The class should roll back the case by changing one fact that a real adversary could influence:
● timing
● identity

● provenance
● replay
● exception handling
Logos sampling
● maintenance state
● operator pressure. The answer must state which claim survives the change and which claim fails. A doctoral answer does not hide behind caution. It chooses a precise operational state:
● admit
● hold
● shadow-test
● collect
● rollback
● or reject.

The final artifact for the border checkpoint consortium is a compact evidence packet. It contains the following:
● The input record
● Transformation
● Unresolved conflict
● Decision authority
● Dissent queue. It also contains one leadership sentence that is weaker than the analyst wants and stronger than a lawyer would fear. That sentence is the craft. It carries enough truth to guide action without smuggling in facts the system did not prove.

For the border checkpoint consortium, change one adversary-controlled fact:
● timing
● identity
● provenance
● replay
● exception handling
● sampling
● maintenance state
● operator pressure. The answer must state which claim survives and which claim fails. A doctoral answer chooses a precise operational state rather than hiding behind generic caution.

The final artifact for the border checkpoint consortium is a compact evidence packet. It contains the following:
● The input record
● Transformation
● Unresolved conflict
● Decision authority
● Dissent queue. It also contains one leadership sentence that guides action without smuggling in facts that the system did not prove.

**War Room 12** begins with the regional health exchange. The first analyst wants to accept the participant update because it fits the expected story. The second analyst refuses to move until the privacy budget is tied to an observed artifact, an owner, and a clock. The disagreement is productive. It prevents the team from treating a plausible explanation as verified.

The class should investigate the case by changing one fact that a real adversary could influence:
● timing
● identity

● provenance
● replay
● exception handling
● sampling
● maintenance state
● operator pressure. The answer must state which claim survives the change and which claim fails. A doctoral answer does not hide behind caution. It chooses a precise operational state:
● admit
● hold
● shadow-test
● collect
● rollback
● or reject.

The final artifact for the regional health exchange is a compact evidence packet. It contains the following:
● The input record
● Transformation
● Unresolved conflict
● Decision authority
● Dissent queue. It also contains one leadership sentence that is weaker than the analyst wants and stronger than a lawyer would fear. That sentence is the craft. It carries enough truth to guide action without smuggling in facts the system did not prove.

For the regional health exchange, change one adversary-controlled fact:
● timing
● identity
● provenance
● replay
● exception handling
● sampling
● maintenance state
● operator pressure. The answer must state which claim survives and which claim fails. A doctoral answer chooses a precise operational state rather than hiding behind generic caution.

The final artifact for the regional health exchange is a compact evidence packet. It contains the following:
● The input record
● Transformation
● Unresolved conflict
● Decision authority
● Dissent queue. It also contains one leadership sentence that guides action without smuggling in facts that the system did not prove.

- Federated learning keeps raw records on site, but unprotected gradients are not always safe. This is why the chapter adds differential privacy on top of model sharing.
- Robust aggregation lowers the risk of poisoning, but it does not fully protect the system from colluding sites or errors in the anomaly model.
- The pure-Numpy TransformerIDS is designed to be compact for teaching and edge deployment. It is not meant to be a universal IDS architecture.

- Threat signals from the shared model still need to be checked by an analyst before they are accepted as incident facts or used as evidence for enforcement.

### 8.11 Summary

Phase 8 enables cross-organizational threat intelligence sharing without exposing raw data by default, while still being honest about the remaining assumptions:

1. **Federated Learning** – local training with FedAvg aggregation, no raw data leaves the site.
2. SCAFFOLD – control variate variance reduction that improves optimization stability under the paper's non-IID assumptions.
3. **Differential Privacy** – calibrated Gaussian noise with RDP composition tracking.
4. **TransformerIDS** – self-attention anomaly detection in pure numpy, INT8-quantized for edge.
5. **Byzantine Robustness** – Multi-Krum, trimmed mean, cosine filtering against poisoning, subject to participant-count and threat-model limits.
6. **Causal Attribution** – intervention-effect reasoning for root cause analysis of threat signals, contingent on the quality of the causal graph.

### Review Questions

1. Explain why FedAvg diverges on non-IID data. Write out the SCAFFOLD update rule and explain how the control variate correction eliminates client drift.
2. A federated system has epsilon = 0.5 per round and a total budget of 10.0. How many rounds of training can be performed before the budget is exhausted? What happens to model utility as epsilon decreases?
3. Design a poisoning attack against a federated anomaly detector. Your compromised site should cause the global model to classify port-scanning traffic as benign. How would Multi-Krum detect your attack?
4. The TransformerIDS uses 4-head attention with d_model=64. Compute the per-head dimension (d_k), and explain why scaled dot-product attention divides by sqrt(d_k).
5. Compare the certified robustness guarantee of randomized smoothing with the empirical robustness of PGD adversarial training. Under what threat model is each approach superior?
6. Draw a causal DAG for a scenario where default credentials lead to unauthorized access, which causes lateral movement, which triggers a ransomware deployment. Use the do-operator to predict the effect of credential rotation.
7. Design an agentic federation controller that decides whether to accept, downweight, or quarantine a site's model update. What observations should it trust, what state must it preserve across rounds, and what verifier should stop a poisoning campaign from becoming "shared truth"?

## References


1. McMahan, H. B., et al.(2017). “Communication-Efficient Learning of Deep Networks from Decentralized Data.” *AISTATS 2017*.
2. Karimireddy, S. P., et al.(2020). “SCAFFOLD: Stochastic Controlled Averaging for Federated Learning.” *ICML 2020*. arXiv:1910.06378.
3. Abadi, M., et al.(2016). “Deep Learning with Differential Privacy.” CCS 2016.
4. Blanchard, P., El Mhamdi, E. M., Guerraoui, R., & Stainer, J. (2017). “Machine Learning with Adversaries: Byzantine Tolerant Gradient Descent.” *NeurIPS 2017*.
5. Yin, D., Chen, Y., Ramchandran, K., & Bartlett, P. (2018). “Byzantine-Robust Distributed Learning.” *ICML 2018*.
6. Madry, A., et al.(2018). “Towards Deep Learning Models Resistant to Adversarial Attacks.” *ICLR 2018*.
7. Cohen, J. M., Rosenfeld, E., & Kolter, J. Z. (2019). “Certified Adversarial Robustness via Randomized Smoothing.” *ICML 2019*.
8. Pearl, J. (2009). *Causality: Models, Reasoning, and Inference*. Cambridge University Press. 2nd edition.
9. Mironov, I. (2017). “Renyi Differential Privacy.” CSF 2017.
10. Vaswani, A., et al.(2017). “Attention Is All You Need.” NeurIPS 2017.


## Cross-References

- **Chapter 1** (Phase 1) – Scan results are the raw data that each site trains on locally.
- **Chapter 4** (Phase 4) – Attack graph topology informs the construction of the causal DAG.
- **Chapter 7** (Phase 7) – Quantum posture data is included in the federated feature vectors.
- **Chapter 9** (Phase 9) – Supply chain anomalies detected by federated models trigger counterfeit investigations.
- **Chapter 12** (Phase 12) – Federated threat signals feed the remediation plan engine.
- **HYDRA** – Federated intelligence is aggregated as a HYDRA stream for BRS computation.