Are You Ready for Production Federated Learning?

Practical readiness considerations for dependable multi-site deployment

Holger Roth, Chester Chen, Peter Cnudde

NVIDIA

TPC’26 Baltimore, May 2026

The Decentralized Equation

Data is the key to generalizable AI, but regulations and privacy laws limit access.

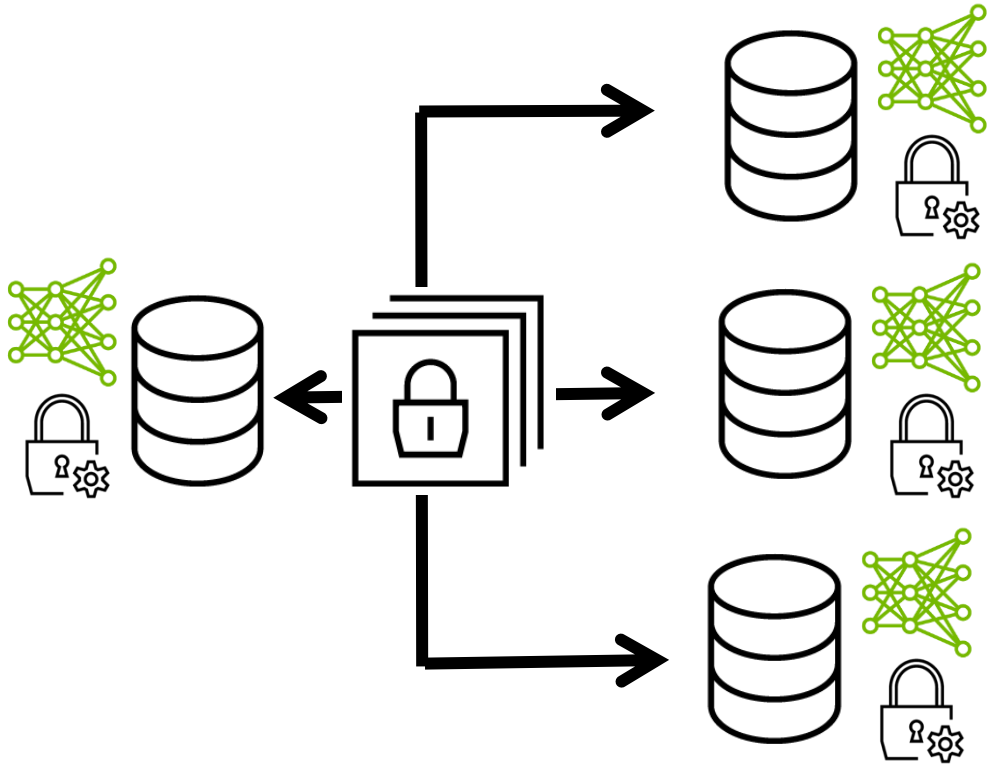

No Raw Data Copy

Central aggregation is impractical.

Compliance

Strict data sovereignty restrictions.

Privacy Technologies

Layered controls: encryption, DP, secure aggregation, confidential computing

FL Pilots vs. Production

A trained model is only one artifact in a multi-site deployment.

Pilot Mindset

- Few sites
- Manual approvals
- Custom scripts
- Best-effort monitoring

readiness evidence

Production Mindset

- Reusable onboarding
- Evidence-based approvals
- Versioned pipelines
- Observable operations

Real-world FL fails when value, data, security, or operations are treated as afterthoughts.

The 5-Domain FL Readiness Framework

Assess readiness before scaling sites, modalities, workflows, or regulatory exposure.

1	2	3	4	5
Use Case & Value	**Governance & Legal**	**Data Architecture**	**Privacy & Security**	**Technology & Operations**
Measurable reason to federate	Authority, approvals, accountability	Schema, ETL, QC, drift	Threat model, PETs, attestation	Runtime, monitoring, runbooks

The Rule: Assess each domain with evidence, not opinions.
Evidence gaps become remediation plans — or a reason not to scale yet.

Evidence gaps tell us exactly what to fix next before scaling.

Domains 1 & 2: Purpose and Permission

A federation needs a reason to exist — and a durable operating agreement.

Domain 1: Use Case & Value

- Narrow, testable clinical/scientific/business questions.
- Site-level data availability.
- Pre-agreed success thresholds.

Domain 2: Governance & Legal

- Accountable operating model
- Approval path for the use case
- Collaboration terms

Production checkpoint

Why federation, who can decide, and what approvals or terms are needed?

Domain 3: Harmonized Inputs, Raw Data Stays Local

Data is only compatible if sites can produce harmonized, validated, versioned inputs.

Minimum evidence

- **Common data model** with versioned ETL.
- **Quality control**: completeness, drift, label quality.
- **Modality compatibility** checks.

Same schema ≠ same meaning

Domain 4: Privacy & Security in FL

Production FL needs layered controls that map to explicit threats and release risks.

Defense in depth

- **Who can connect?** Identity, authorization, certificates, and key management
- **What can move?** Model updates, metrics, logs, and outputs, protected with appropriate PETs
- **Where can it run?** Trusted execution and attestation for higher-assurance environments

Production evidence

1 Threat Model

What are the attack scenarios and sensitivities?

2 Control rationale

Which PETs and runtime controls are justified?

3 Release risk check

Is it safe to share the model, metrics, or outputs?

Principle: Use FL to reduce data movement, then use controls to manage residual risk.

Domain 5: Resilient Technology & Operations

The platform must be reproducible, observable, and resilient across sites.

Operational capabilities

- **Connect sites reliably:** Onboarding, networking, access checks
- **Run consistently across environments:** Approved runtimes from on-prem/device to cloud/HPC
- **Observe and recover:** Monitoring, incidents, rollback, offboarding

Reference implementation lens:

NVIDIA FLARE provides workflow/runtime hooks for distributed workflows, secure execution, federated evaluation, and monitoring — but readiness still depends on site evidence and process discipline.

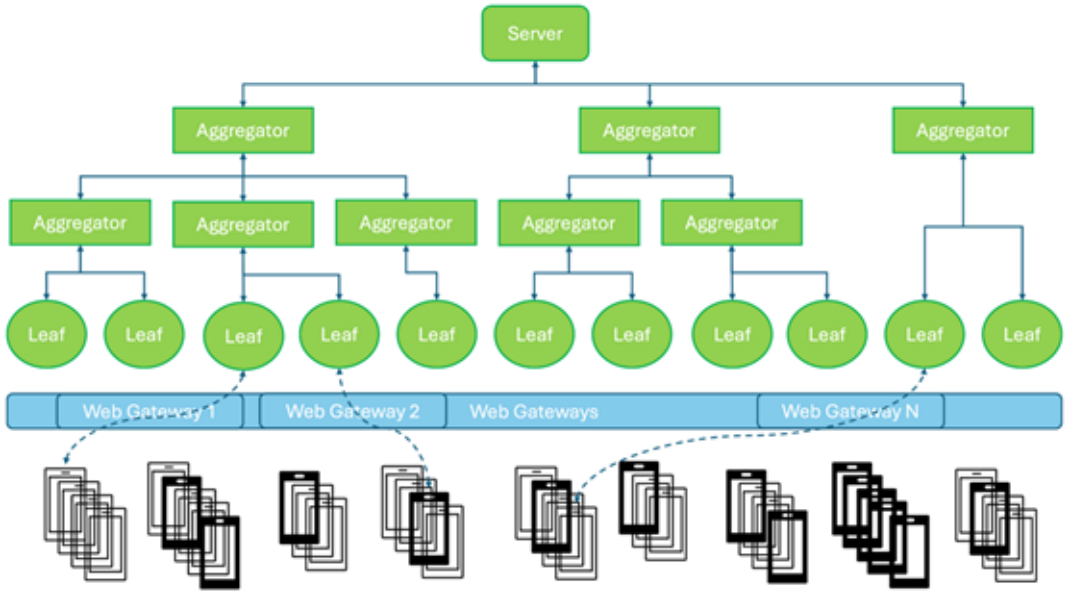

Scaling patterns: hierarchical federation, Kubernetes/Slurm, large-model communication

Make readiness a decision, not a feeling

Minimum evidence turns assumptions into remediation plans and go/no-go signals.

Domain	Production-ready when…
Use Case & Value	The value case, sites, and success criteria are clear
Governance & Legal	Accountable owners and approvals are in place
Data Architecture	Inputs are harmonized, versioned, and quality-checked
Privacy & Security	Controls map to threats and release risks
Technology & Operations	The federation is observable, supportable, and recoverable

Do not ask only **“can we train it?”** Ask, **“can we operate it, audit it, and change it safely?”**

Learn more about FL Use Cases

NVIDIA FLARE Days

Next event: September 16 & 24, 2026

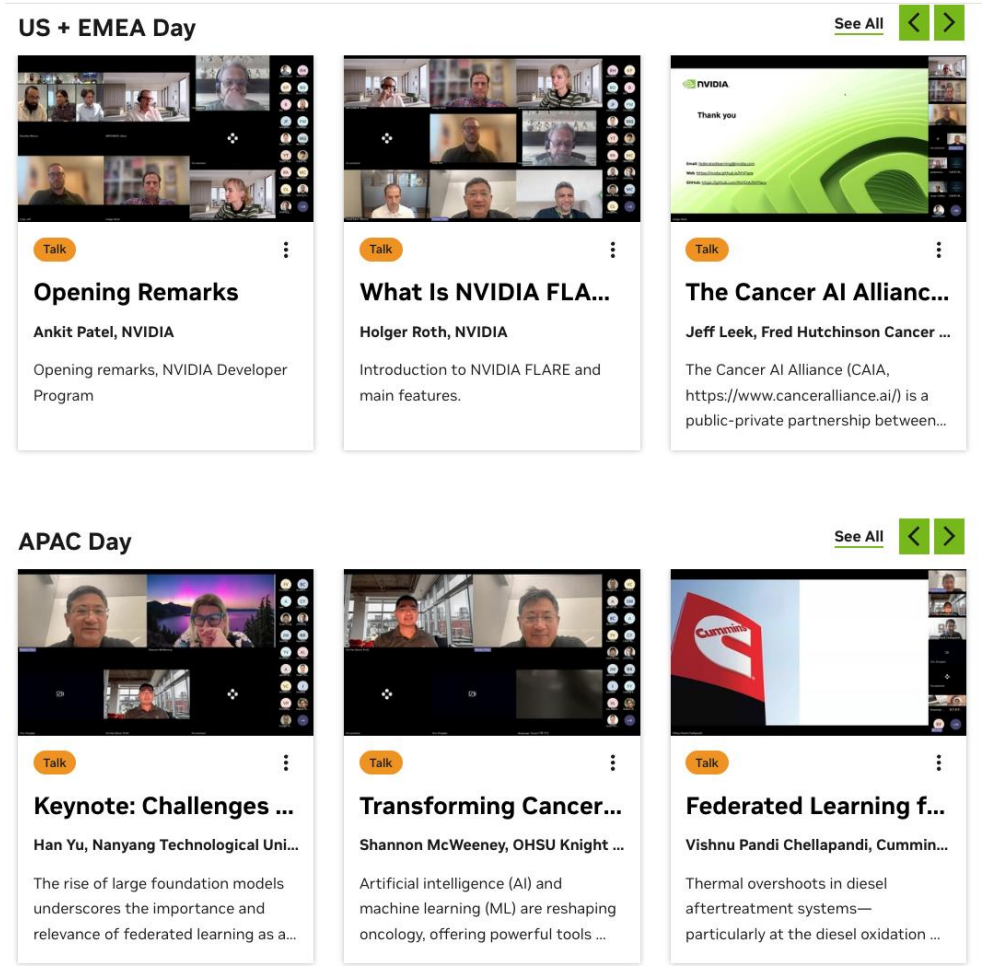

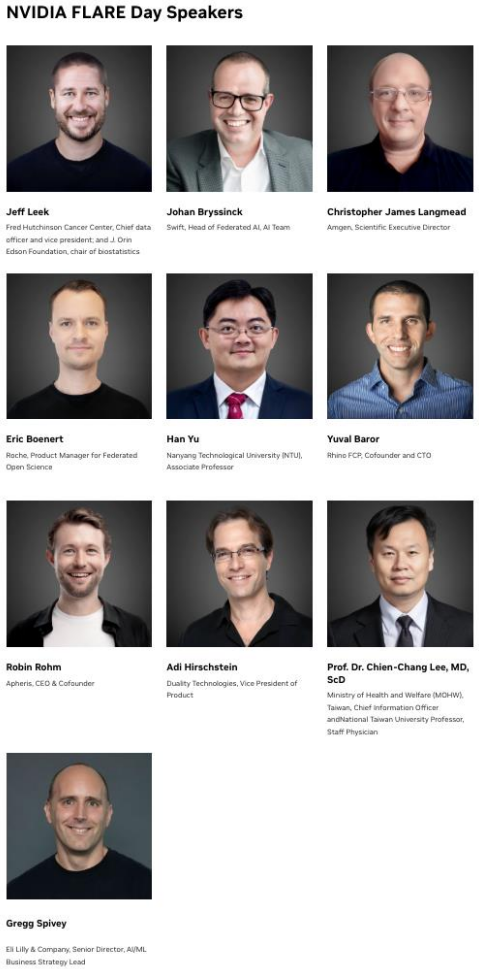

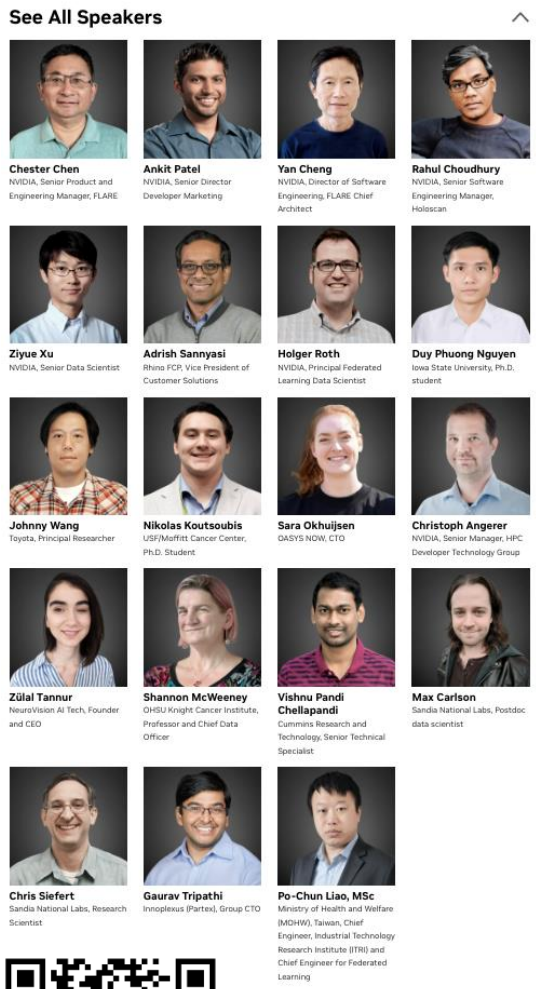

Past recordings: developer.nvidia.com/flare-day-2025

Thank you!

Discussion: What should we standardize first so that federated learning can move more reliably from pilots to production?

Contact: Holger Roth <hroth@nvidia.com>

Try out NVIDIA FLARE!

github.com/NVIDIA/NVFlare