

Stitching Together Innovation with FABRIC Users

When Federated Learning Meets FABRIC
Zilinghan Li, Ze Yang, Ravi Madduri

12/17/2024



Welcome and Introduction





Zilinghan Li

Argonne National Laboratory |
Machine Learning Engineer





Ze Yang

University of Illinois at Urbana-Champaign |
Visiting Research Associate



Ravi Madduri

Argonne National Laboratory |
Senior Computer Scientist



Federated Learning

Motivations for Using Federated Learning

Federated learning (FL)'s Motivations:

- Sufficient data is essential to training high-quality and generalized AI models.
- Data within a single data silo are usually limited and biased.
- Data are likely distributed heterogeneously among various data silos.
- Directly collecting data from distributed silos are sometimes impossible due to privacy concerns, especially in sensitive domains such as finance and biomedicine.
- Federated Learning enables the training of more generalized model without direct data sharing.

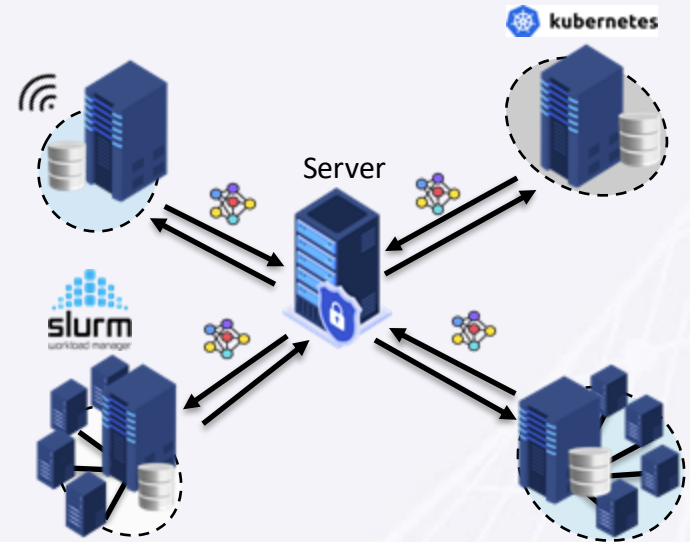


Fig. Federated Learning Illustration.

Federated Learning

Distributed Machine Learning Paradigm without Direct Data Sharing

Federated learning (FL)'s Workflow:

- Distributed machine learning paradigm
- Multiple *clients* with own computing resources and private local data
- One central orchestration *server*
- Each client trains a local model and shares the model with the server for aggregation
- The aggregated model leverages data from multiple clients to obtain more generalized model without direct data sharing


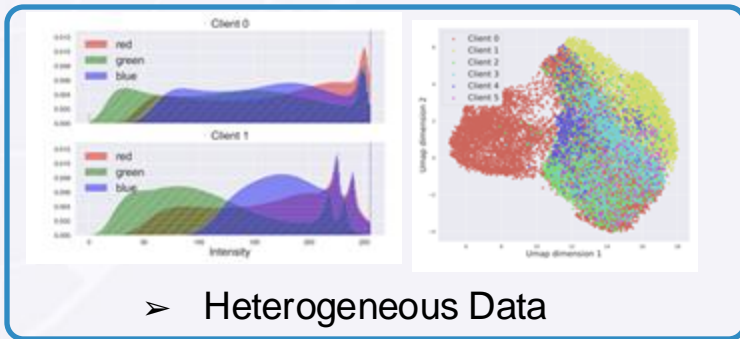


Fig. Federated Learning Illustration.

Federated Learning

Various Challenges of Federated Learning Due to its Distributed Nature

Advanced
Aggregation
Strategies

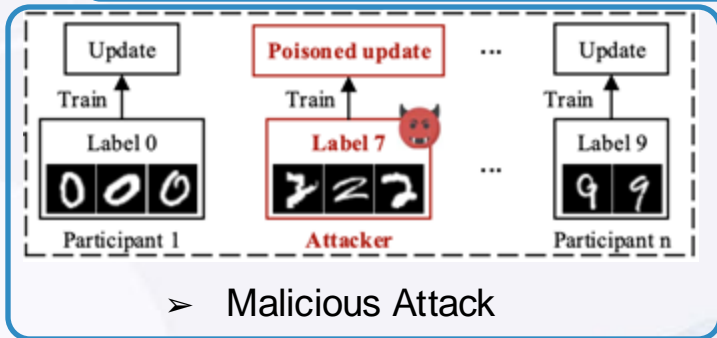


➤ Heterogeneous Data




➤ Heterogeneous Compute

Advanced
Asynchronous
Scheduling
Algorithms



➤ Malicious Attack



➤ Data Reconstruction



➤ Cumbersome Setup

Robust Authentication

Differential Privacy and Privacy Enhancing Techniques



APPFL (Advanced Privacy-Preserving Federated Learning) Framework is our framework-wise solutions to those challenges.

Simple Exp Configuration


APPFL Framework

APPFL: Advanced Privacy-Preserving Federated Learning Framework

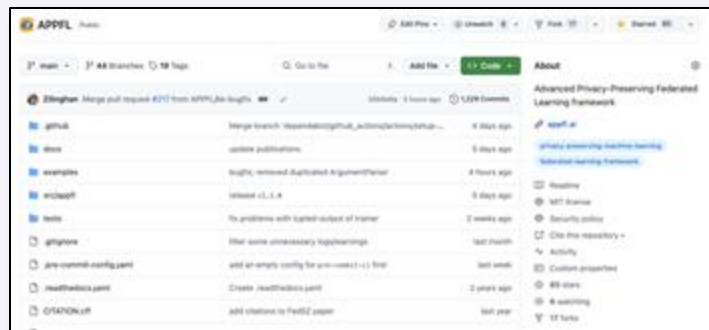
APPFL is an open-source FL framework which supports comprehensive solutions for various FL challenges.




[Manuscript](#)

Framework Design Description 

- Framework overview
- Addressed challenges
- Evaluations
- Additional case studies
- ...



[Open-source Code](#)

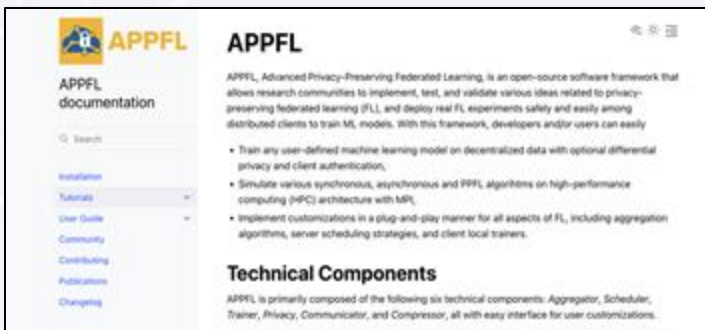
Source code on Github 

- Fully open-source
- Welcome issues
- Welcome contributions
- ...


APPFL Framework

APPFL: Advanced Privacy-Preserving Federated Learning Framework

APPFL is an open-source FL framework which supports comprehensive solutions for various FL challenges.



appfl.ai

Detailed Documentation 

- Installation
- Launching FL experiments
- Advanced Developer Guides
- ...



service.appfl.ai

APPFL-based Service Platform 

- Fully based on APPFL
- User-friendly for domain experts
- Comprehensive report generation
- ...

APPFL x FABRIC

Using FABRIC as the Testbed for APPFL

Benefits provided by FABRIC:

➤ *Heterogeneous and Distributed Testbed*

FABRIC offers a heterogeneous and geographically distributed testbed for testing and benchmarking distributed computing applications.

➤ *Enhanced Understanding of Network Role in FL*

FABRIC enables exploration of the role of networks in training models across geographically distributed sites, especially when working with large datasets at each location.

➤ *Inbound Connection Support*

Unlike many supercomputers (e.g., Polaris), FABRIC supports servers which allows inbound connections from clients, making it ideal for hosting FL servers.



APPFL

×



FABRIC

APPFL x FABRIC

Using FABRIC as the Testbed for APPFL

Benefits provided by FABRIC:

➤ *Comprehensive Monitoring Tools*

FABRIC provides advanced tools for measuring latency and monitoring resource utilization.

➤ *Connectivity to External Facilities*

FABRIC allows seamless connection to external facilities for additional computing resources, for example, Chameleon Facility Ports help access additional computing resources for compute-intensive applications.



APPFL


×




FABRIC

APPFL x FABRIC


Experiments Enabled by FABRIC Testbed



Benchmark the *communication latency* of the various communication methods and protocols



Measure the *training efficiency* of several synchronous and asynchronous algorithms on *distributed and heterogeneous* machines



Evaluate the framework *scalability* as the client number grows



APPFL



FABRIC

APPFL x FABRIC

Demo Setup

- *Configure FABRIC Environment*

Set up and customize the configuration for the FABRIC platform to ensure smooth execution of tasks [Following FABRIC's Tutorial Jupyter Notebook].

- *Allocate Resources and Provision Nodes*


Efficiently allocate computational resources and provision the necessary nodes to support the experiment setup. [One node for server and two more nodes for clients]

- *Establish SSH Connection to Remote Nodes*

Securely connect to remote nodes using SSH for seamless interaction and experiment configuration.



APPFL



FABRIC

APPFL x FABRIC

Demo Setup

- *Install APPFL and Configure Experiment*

Install the APPFL framework and modify its for the experiments.

- *Launch and Monitor Experiments*

Run the experiments and continuously monitor their performance to ensure accuracy and efficiency.



APPFL

×



FABRIC

APPFL x FABRIC Demo

https://jupyter.fabric-testbed.net/user/zeyang2@illinois.edu/lab/tree/jupyter-examples-ret1.7.0/configure_and_validate.ipynb

File Edit View Run Kernel Tabs Settings Help

Filter files by name

Launcher

- start_here.ipynb
- configure_and_validate.ipynb
- APPFL-Webinar.ipynb

10 seconds ago

Setup configuration for Using Fabric Testbed via FABLib API

The following notebook can be used to configure your Jupyter environment. The product of running this notebook are the following:

- `fabric_rc`: File used to configure a FABLib application.
- `fabric_bastion_key`: Fabric Bastion key pair. In order to minimize security incidents on FABRIC, access to VMs and other resources administered by users is controlled using a bastion host. You will need to set up an ssh keypair that will be used to jump through the bastion host to your VMs and other resources. This keypair is unique to you and is only used to set up ssh proxy connections through the bastion host to your FABRIC resources. More information about how to access your experiment through the bastion host can be found here.
- `silver_key` and `silver_key.pub`: Silver Key pair.
- `ssh_config`: File used to ssh from from a terminal to FABRIC VM by jumping through the FABRIC bastion host.

Set Project ID and generate the configuration ¶

Edit the following cell by entering your Project ID for the FABRIC Project to use in your Jupyter container.

- The Project ID can be from any of your projects. The ID can be found in the 'Basic Info' tab for each of the projects in the FABRIC portal.

```
from fabrictestbed_extensions.fablib.fablib import FablibManager as fablib_manager

# update this line to specify your project id
project_id = "848486f-1212-4d98-9647-0ec3f8b11bda"

# Uncomment the line below if using 'FABRIC Tutorial' Project
#project_id = "7828536-1f51-4e77-3b60-d17129b0862"

fablib = fablib_manager(project_id=project_id)
```

Display the configuration

```
fablib.show_config()
```

Validate the configuration;

- Checks the validity of the bastion keys and regenerates them if they are expired
- Generates Silver keys if they do not exist already

```
fablib.verify_and_configure()
```

Save the configuration for subsequent use

```
fablib.save_config()
```

Would you like to receive official Jupyter news?
Please read the privacy policy.
[Open privacy policy](#) Yes No

Simple Python 3 (ipykernel) | Mode: Command Ln 1, Col 1 configure_and_validate.ipynb 1

11:42 AM 12/10/2024

APPFL x FABRIC

Demo Recipes

FABlib Config

Orchestrator	orchestrator.fabric-testbed.net
Credential Manager	cm.fabric-testbed.net
Core API	uis.fabric-testbed.net
Token File	/home/fabric/tokens.json
Project ID	a84a6e6f-1212-4d90-9647-0ee3f8b1bda
Bastion Host	bastion.fabric-testbed.net
Bastion Username	zeyang2_0033371586
Bastion Private Key File	/home/fabric/work/fabric_config/bastion-key
Slice Public Key File	/home/fabric/work/fabric_config/slice_key.pub
Slice Private Key File	/home/fabric/work/fabric_config/slice_key
Sites to avoid	
SSH Command Line	ssh -i {{{_self_private_ssh_key_file}}} -F /home/fabric/work/fabric_config/ssh_config {{{_self_username}}}@{{{_self_management_ip}}}
Log Level	INFO
Log File	/tmp/fablib/fablib.log
Bastion SSH Config File	/home/fabric/work/fabric_config/ssh_config
Version	1.7.3
Data directory	/tmp/fablib

```
# To simulate multisite communications, we randomly fetch 3 sites for node creation.
# Each site has one nodes.
slice_name = 'MySlice'
[site1,site2,site3] = fablib.get_random_sites(count=3)
print(f"Sites: {site1}, {site2}, {site3}")

node1_name = 'Node1'
node2_name = 'Node2'
node3_name = 'Node3'
```

Python

Sites: SEAT, STAR, INDI

```
1 #!/bin/bash
2
3 # Update and install required packages
4 sudo apt update -y
5 sudo apt install python3-pip -y
6 sudo add-apt-repository ppa:deadsnakes/ppa -y
7 sudo apt install python3.10 python3.10-venv python3.10-distutils -y
8
9 # Verify Python version
10 python3.10 --version
11
12 # Create and activate a virtual environment
13 python3.10 -m venv my_env
14 source my_env/bin/activate
15
16 # Clone the repository and install dependencies
17 git clone --single-branch --branch zey/webinar https://github.com/APPFL/APPFL.git
18 cd APPFL
19 pip install -e ".[dev,examples]"
20
21 # Navigate to examples directory and run the server and clients
22 cd examples
23
24 # Run server
25
26 python ./grpc/run_server.py --config ./resources/configs/mnist/server_fedcompass.yaml
27
28 # Run clients
29
30 python ./grpc/run_client.py --config ./resources/configs/mnist/client_1.yaml
31
32 python ./grpc/run_client.py --config ./resources/configs/mnist/client_2.yaml
33
```

Links:

1. [FABRIC Setup Notebook](#)
2. [FL Experiment Setup and Launching Scripts](#)



APPFL



FABRIC

APPFL x FABRIC

Future Work

➤ *Chameleon Integration*

Integrate FABRIC with Chameleon (<https://www.chameleoncloud.org>) to get access to powerful machines to run more computational intensive experiments.


➤ *Running FL with FABRIC on Climate Data*

Leverage FABRIC's configurable network and seamless node connectivity to optimize federated learning for large-scale climate data, avoiding transferring large amount of data and improving model performance. [[APPFL+FABRIC User Story](#)]



APPFL

×



FABRIC

Thank you!



Q&A



Acknowledgements

In addition to our presenter, we would like to acknowledge the behind-the-scenes team that diligently worked to bring this webinar to production:

- **KC Wang**, content
- **Chelsea Davis**, project manager
- **Jayasree Jaganatha**, social media specialist



Resources

Call to Action

FABRIC Matrix:
<https://bit.ly/FABRICmatrix>

Connect With Us

Newsletter Signup: bit.ly/FABRICnewsletter

Office Hour Sign Up: bit.ly/FABRIC-Office-Hours

Other Resources

Website: bit.ly/m/FABRICtestbed

YouTube: youtube.com/@fabrictestbed

FABRIC Account: portal.fabric-testbed.net

Ambassador Program: bit.ly/FABRIC-Ambassador-Program

FABRIC LinkedIn: linkedin.com/company/fabrictestbed

Citing FABRIC: bit.ly/citing-fabric



Thank You for Attending!

Join us for our upcoming webinars:

- **Date** - Stitching Together Innovation with FABRIC Users
- **Date** - Mastering FABRIC: Tips and Tricks Webinar

Visit our YouTube Channel: [youtube.com/@fabrictestbed](https://www.youtube.com/@fabrictestbed)

