WAMulator: A WAN Emulator for Precisely Mimicking WAN in LAN Settings

Keisuke Miyamoto
Undergraduate Student, Computer Science

1. Description of Project

In today's interconnected world, global services such as e-commerce and media streaming rely heavily on Geo-Distributed Data Analytics (GDA) to gain timely insights from data spread across the globe. However, developing and validating new GDA systems is challenging because these systems must be deployed in real-world, multi-region data centers (DCs), which is costly and time-consuming. In addition, GDA systems often encounter a significant cost-bottleneck caused by the large volumes of data exchanges across DCs over the wide area network (WAN), one of the most scarce and expensive resources in multi-cloud environments.

To overcome these barriers, many GDA researchers and developers rely on emulating a WAN within a local area network (LAN) as a cost- and time-effective alternative. To mimic WAN environments in LANs, which provide abundant bandwidth, they first measure WAN bandwidth across DCs *statically* and *independently* using tools such as iPerf [1]. They then throttle LAN bandwidth according to the measured values using Traffic Control (`tc`) [2]. However, this oversimplified approach cannot capture the dynamic and complex nature of WAN environments for GDA, where data transfers occur *dynamically* and *simultaneously*. GDA researchers may use well-known virtual network tools [3, 4, 6, 9] to shape network traffic to behave like WAN. Unfortunately, the same problem persists because these tools also rely on tc to regulate traffics statically based on given values. As a result, this can lead to inaccurate evaluations of GDA systems, potentially causing researchers to draw incorrect conclusions.

In this research project, **I aim to develop a framework that precisely and dynamically emulates WAN environments within LAN settings**. The proposed prototype, *WAMulator*, adopts an agent-based model to faithfully reproduce the dynamic characteristics of WAN environments. The agents receive target WAN bandwidth values, along with other emulation inputs (e.g., DCs' regions and node information), and dynamically shape network traffic between nodes to mimic WAN conditions. This design enables researchers to accurately emulate WAN environments in LAN settings, providing a cost-effective platform for conducting fully controlled and repeatable experiments on GDA systems.

2. Activities and Methodology

WAMulator requires accurate WAN bandwidth values to emulate the WAN. Since directly measuring these values in real environments is both expensive and time-consuming, predicting current WAN bandwidth would be highly desirable. To obtain these values, I will use a WAN framework called *WANify* [7], developed by the Multi-Cloud Computing Lab at UNO. This framework was designed to overcome the limitations of simple approaches that measure WAN bandwidth across DCs at minimum costs. Instead of relying on a single, static measurement of network throughput, WANify employs a machine learning model to predict the actual, dynamic bandwidth between DCs. The predicted values will then be used by WAMulator to emulate WAN behavior in LAN settings.

To develop WAMulator, I will follow these steps below.

a) Study Basic Technologies and Prior Research
I will investigate prior GDA research related to this project such as *Tetrium* [5] and *Kimchi* [8], and learn about the theories and implementations they used. Next, I will study the technologies necessary for bandwidth and connection control, especially the traffic control command `tc` provided by the Linux kernel, to acquire the knowledge required for implementing WAMulator.

b) Explore Spark
Apache Spark is a distributed data processing framework that plays a central role in many GDA-related systems including WANify. Since many GDA systems rely on Spark, it is important to understand how to use it.

c) Deploy Kimchi and Tetrium and Apply WANify to Them
To gain a more hands-on understanding of how Apache Spark is used in real projects, I will deploy *Kimchi* and *Tetrium*, both GDA systems built using Apache Spark. Then I will integrate WANify with them to analyze its impact on their performance. This will help me understand WANify's actual behavior with GDA systems.

d) Build a Prototype of WAMulator
Based on the insights gained from prior research and related technologies, I will conduct the basic design of WAMulator and build a prototype. WAMulator will receive the predicted results of WAN communication environments generated by WANify as input and build an architecture to faithfully reproduce them within a LAN. Specifically, I will design a mechanism to dynamically control network bandwidth on the LAN, making it possible to mimic real-time network environments according to WANify's prediction data.

e) Integrate WANify and WAMulator
By connecting WAMulator with WANify, WAMulator will be able to reproduce in the LAN the optimal bandwidth predicted by WANify for multi-region cloud environments. This enables the automation of an end-to-end pipeline where the network conditions predicted by WANify can be immediately reproduced in the LAN environment, achieving seamless integration of the two systems.

3. Project Timeline

This is a timeline for the WAMulator project. I am requesting a SUMMER 2026 funding timeline.

February-April	**May-August**
a) Study Basic Technologies and Prior Research **b)** Explore Spark **c)** Deploy Kimchi and Tetrium and Apply WANify to Them	**d)** Build a Prototype of WAMulator **e)** Integrate WANify and WAMulator

4. Student/Faculty Mentor Roles

I will be responsible for completing all project tasks, including the design and implementation of the WAMulator prototype. I will also be accountable for running experiments on emulated and real test beds using the the implemented prototype. The research advisor, Dr. Kwangsung Oh, will review the progress through weekly meetings and address any unforeseen challenges. In addition, Dr. Oh will assess the experimental results.

5. A list of previous internal funding received

There is no previous internal funding received related to this project.

Budget

Student Stipend: $2,000
Cloud Resources: $500
Total Requested: $2,500.00

Support in the amount of $2,000 is requested for the student, Keisuke Miyamoto, to complete the research tasks on the project as described in the Project Description. The stipend is budgeted based on an hourly rate of $16/hour for 125 hours works. $500 is requested for cloud resources needed to deploy the WAMulator prototype implementation across multiple cloud DC regions to measure real WAN characteristics and verify the accuracy in real multi-region cloud environments.

References

[1] iPerf. https://iperf.fr/.

[2] Linux Traffic Control. http://lartc.org/manpages/tc.txt.

[3] Vitaly Antonenko and Ruslan Smelyanskiy. Global network modelling based on mininet approach. In *Proceedings of the Second ACM SIGCOMM Workshop on Hot Topics in Software Defined Networking*, HotSDN '13, page 145–146, New York, NY, USA, 2013. Association for Computing Machinery.

[4] Giuseppe Di Lena, Andrea Tomassilli, Damien Saucez, Frédéric Giroire, Thierry Turletti, and Chidung Lac. Distrinet: a mininet implementation for the cloud. *SIGCOMM Comput. Commun. Rev.*, 51(1):2–9, March 2021.

[5] Chien-Chun Hung, Ganesh Ananthanarayanan, Leana Golubchik, Minlan Yu, and Mingyang Zhang. Wide-area analytics with multiple resources. In *Proceedings of the Thirteenth EuroSys Conference*, EuroSys '18, New York, NY, USA, 2018. Association for Computing Machinery.

[6] Bob Lantz, Brandon Heller, and Nick McKeown. A network in a laptop: rapid prototyping for software-defined networks. In *Proceedings of the 9th ACM SIGCOMM Workshop on Hot Topics in Networks*, Hotnets-IX, New York, NY, USA, 2010. Association for Computing Machinery.

[7] Anshuman Das Mohapatra and Kwangsung Oh. Wanify: Gauging and balancing runtime wan bandwidth for geo-distributed data analytics, 2025.

[8] Kwangsung Oh, Abhishek Chandra, and Jon Weissman. A network cost-aware geo-distributed data analytics system. In *2020 20th IEEE/ACM International Symposium on Cluster, Cloud and Internet Computing (CCGRID)*, pages 649–658, 2020.

[9] Philip Wette, Martin Dräxler, Arne Schwabe, Felix Wallaschek, Mohammad Hassan Zahraee, and Holger Karl. Maxinet: Distributed emulation of software-defined networks. In *2014 IFIP Networking Conference*, pages 1–9, 2014.

University of Nebraska Omaha

Dr. Kwangsung Oh
PKI 177D
Omaha, Nebraska 68182
Phone: 402.554.2862
kwangsungoh@unomaha.edu

Department of Computer Science
College of Information Science and Technology

January 5, 2026

Subject: Keisuke Miyamoto FUSE Award

Dear Member of the FUSE Grant Selection Committee:

I am writing this letter to express my enthusiastic support for Keisuke Miyamoto's application for the FUSE award in Summer 2026. I have known Keisuke since Fall 2024, when he took my CSCI 3320 Data Structures course. As he was one of the strongest students in the class, I invited him to several meetings to discuss potential undergraduate research opportunities. Keisuke expressed interest in research under my supervision. In Spring 2025, Keisuke also took my CSCI 4350 Computer Architecture course, where he earned the highest grade of the class. Based on his outstanding academic performance and research potential, I encouraged him to pursue graduate studies, and he has expressed clear interest in joining our MS program after completing his BS degree. We have discussed several possible research directions, and he has already conducted GPU-based experiments on AWS as part of our ongoing GPU-Smartpick project. I have been particularly impressed by his progress on this project, as he has successfully trained and executed numerous AI/ML models on a variety of GPU-enabled cloud resources, yielding insightful observations that can be leveraged for potential research publications.

Keisuke will begin his MS program in Fall 2026. For his MS thesis, we have discussed his next research project extensively. His research will focus on the challenge of emulating wide-area network (WAN) behaviors within local-area network (LAN) settings for geo-distributed data analytics (GDA), which is one of the most essential workloads for many Internet applications such as Netflix, Facebook, and Airbnb that operate in multi-cloud environments. While GDA researchers heavily rely on WAN emulation for system validation, oversimplified emulation approaches can lead to inaccurate evaluations and incorrect conclusions. In this project, he will primarily investigate the following research question: *"How can LAN traffic be accurately controlled to precisely emulate WAN bandwidth, so that experimental results closely match those obtained in real WAN-based multi-cloud environments?"* This work is a part of our ongoing project titled "CAREER: Toward Scalable Wide Area Network-enabled Geo-distributed Data Analytics." To validate the proposed approaches, he will conduct rigorous experiments in public cloud environments. The proposed WAN emulator will enable GDA researchers to easily and accurately evaluate their systems in emulated WAN environments, allowing them to draw reliable and correct conclusions with minimal effort.

To finish the proposed project successfully, we will meet weekly during Spring and Summer 2026. I will provide technical guidance for building the proposed framework, deploying it on public cloud environments, and conducting extensive experiments. I expect he will be able to publish a paper with the support of the FUSE award, which will also be used for his MS thesis.

In conclusion, I offer my strongest endorsement of Keisuke's FUSE application without reservation. This grant is essential, as the financial support will allow him to focus on his research during the Summer of 2026. Keisuke is open-minded, creative, energetic, and highly motivated. I am grateful for the opportunity to support his research development, and I strongly encourage you to award him this FUSE grant for the proposed project.

Thank you for your time and consideration of his proposal. I look forward to the exciting opportunities this funding will bring to him. If you have any questions, please feel free to contact me in person.

Sincerely,
Kwangsung Oh, Ph.D.
Assistant Professor, Department of Computer Science
College of Information Science and Technology