

# Experiments at the Internet’s Edge with Dasu

Mario A. Sánchez\*   John S. Otto\*   Zachary S. Bischof\*   David R. Choffnes†  
Fabián E. Bustamante\*   Balachander Krishnamurthy‡   Walter Willinger‡  
\*Northwestern U.   †U. of Washington   ‡AT&T Labs-Research

Dasu [1] is an extensible platform for running network measurements and experiments at the Internet’s edge. Its clients run on end hosts, have built-in support for broadband characterization—an incentive for end-user adoption—and execute third-party experiment tasks. The platform supports concurrent third-party experiments by delegating clients to tasks based on experiment specifications and resource availability.

This demo focuses on Dasu’s task delegation mechanism and shows how it enables third-party experimentation and maintain security and accountability. The Experiment Administration Service (EA Service) manages experiments and resource allocation, acting as a broker between clients and third-party experiments. Clients follow a pull-based model: they periodically request experiment tasks from the EA Service. The EA Service assigns a client request to an experiment by matching the characteristics of a client, such as its IP prefix, geographic location, or operating system, to the stated requirements of experiments.

Experiments can target a subset of the Dasu client base corresponding to researchers’ area of interest via arbitrarily complex logic on client characteristics. For instance, one experiment might require the use of hosts located in a particular geographic or network region, while a subset of all Linux clients might suffice for another. To this end, Dasu adopts a two-tiered architecture for the EA Service, with a Primary EA server, responsible for resource allocation, and a number of Secondary EA servers in charge of particular experiments.

Figure 1 illustrates the interaction between Dasu clients and the EA Service. The Primary EA server acts as a broker, allocating Dasu clients to experiments and delegating their requests to the responsible Secondary EA server for the chosen experiment. The Secondary EA server is responsible for task parameterization and the allocation of tasks to clients according to the experiment’s logic. While the customized task assigned to a client is generated by the experiment’s Secondary EA server, all communication with Dasu clients is mediated by the Primary EA server who is responsible for authenticating and digitally signing the assigned experiments.

Dasu is designed to support third-party experiments through the two-tier architecture described above. Au-




Figure 1: Interaction between Dasu Clients and the Experiment Administration Service.

thorized research groups host their own Secondary EA server while security and accountability is maintained through the Primary EA server.

## 1 Demonstration

Our demonstration shows the current implementation of Dasu’s Experiment Administration Service. It illustrates the interactions between Clients, Primary and Secondary EA Servers, emphasizing how experiment request, delegation and authentication take place. We provide the live view seen by each of these components to demonstrate the end-to-end process from deploying an experiment to obtaining results.

## References

- [1] SÁNCHEZ, M. A., OTTO, J. S., BISCHOF, Z. S., CHOFFNES, D. R., BUSTAMANTE, F. E., KRISHNAMURTHY, B., AND WILLINGER, W. Dasu: Pushing experiments to the Internet’s edge. In *Proc. of USENIX NSDI* (2013).