

Experiments at the Internet’s Edge with Dasu

Mario A. Sánchez^{1*} John S. Otto^{1*} Zachary S. Bischof^{1*}
David R. Choffnes[†] Fabián E. Bustamante^{*} Walter Willinger[‡]
**Northwestern U. †U. of Washington ‡AT&T Labs-Research*

Dasu [1] is an extensible measurement experimentation platform for the Internet’s edge. Dasu is composed of a distributed collection of clients, hosted by participating end hosts, and a core set of services for managing and coordinating experimentation. Dasu supports and builds on broadband characterization as an incentive for adoption to capture the network and service diversity of the commercial Internet.

This demo presents Dasu in action, focusing on its experiment delegation mechanism and showing how it enables third-party experimentation and maintains security and accountability.

Dasu’s management services include a Configuration Service, an Experiment Administration Service, a Coordination Service and a Data Service. Upon initialization, Dasu clients use the Configuration Service to announce themselves and obtain various configuration settings. Clients periodically request experiment tasks from the Experiment Administration (EA) Service, who assigns experiments based on client’s characteristics (such as its IP prefix, geographic location, or operating system) and experiments’ requirements. Clients also contact the Coordination Service to submit updates about completed tasks and retrieve resource usage and experiment constrains. Finally, clients use the Data Service to report results of completed experiments as they become available.

Dasu enables sophisticated experiments by supporting the assignment of tasks to clients according to arbitrarily complex experiment logic based on clients’ characteristics. For instance, while some experiments may require the use of hosts located in a particular geographic or network region, any random set of available Linux clients might be sufficient to others. 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.

This two-tier architecture enables third-party experimentation, with each participating group hosting their own Secondary EA server. The Primary EA server acts as a broker, allocating Dasu clients to experiments by assigning them to the responsible Secondary EA server, based on clients’ characteristics and resource availability.

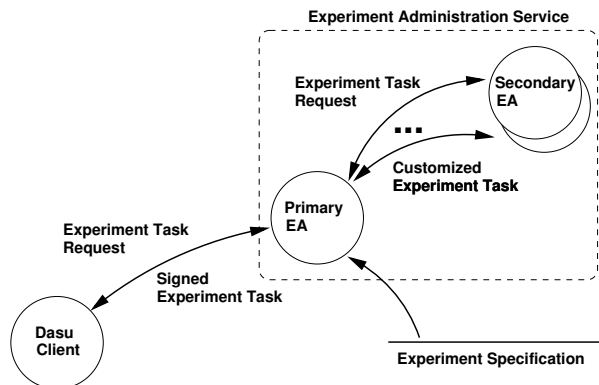


Figure 1: Interaction between Dasu Clients and the Experiment Administration Service.

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.

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).