

Abstract

Kinds of stateful stream process engines (SPEs) track a large number of concurrent flow states and replicate them to backups to provide reliable functionality in high availability clusters (HACs). Under high traffic loads, existing solutions in such HACs are expensive because of precise stateful replication. In this dissertation, I study a suite of two methods to address this issue: randomization on replication messages and a replication scheme designed for when system is going to be overloaded.

Two new hierarchical structures called Flow Digest (FD) and Multi-Level Counting Bloom Filter (MLCBF) are proposed as low resource-consuming solutions of stateful replication. To the best of my knowledge, it is the first time that randomization has been introduced for stateful replication of HAC in the literature. Analysis and extensive tests are employed to evaluate performance and tradeoffs of the proposed techniques. Most importantly, MLCBF is quite as simple and practical to implement and maintain.

Furthermore, an adaptive scheme, called as *dynamic lazy insertion*, is designed to prevent replication from overloading system and optimize pass-through performance of HAC dynamically. Testbed evaluation demonstrates its feasibility and effectiveness in real situation.