

Many applications today require distributed systems to store and process massive amounts of data. The complexity of these applications grows as the systems running them become more complex, and as the applications require more systems in order to run. End-to-end tracing is the process of following a request through the network of machines running one or more distributed systems in order to aid in the maintenance, debugging, and optimization of the systems. We will analyze the two primary models used for end-to-end tracing systems, the span model and the event model. We will show that spans are the less powerful model through a formal proof and a practical implementation of the proof, as well as discuss the consequences of using the span model. We will finish by proposing a joint model that incorporates both spans and events.

The diagram below shows a trace gathered using our new joint span-event model:

