

# Peer-to-Peer Process Execution with OSIRIS

Christoph Schuler<sup>1</sup>, Roger Weber<sup>1</sup>, Heiko Schuldt<sup>2</sup>, and Hans-J. Schek<sup>1</sup>

<sup>1</sup> Swiss Federal Institute of Technology (ETH),  
CH-8092 Zurich, Email: {schuler,weber,schek}@inf.ethz.ch

<sup>2</sup> University for Health Informatics and Technology Tyrol (UMIT),  
A-6020 Innsbruck, Email: heiko.schuldt@umit.at

**Abstract.** Standards like SOAP, WSDL, and UDDI facilitate the proliferation of services. Based on these technologies, processes are a means to combine services to applications and to provide new value-added services. For large information systems, a centralized process engine is no longer appropriate due to limited scalability. Instead, in this paper, we propose a distributed and decentralized process engine that routes process instances directly from one node to the next ones. Such a Peer-to-Peer Process Execution (P<sup>3</sup>E) promises good scalability characteristics since it is able to dynamically balance the load of processes and services among all available service providers. Therefore, navigation costs only accumulate on nodes that are directly involved in the execution. However, this requires sophisticated strategies for the replication of meta information for P<sup>3</sup>E. Especially, replication mechanisms should avoid frequent accesses to global information repositories. In our system called OSIRIS (*Open Service Infrastructure for Reliable and Integrated Process Support*), we deploy a clever publish/subscribe based replication scheme together with freshness predicates to significantly reduce replication costs. This way, OSIRIS can support process-based applications in a dynamically evolving system without limiting scalability and correctness. First experiments have shown very promising results with respect to scalability.