



# Scalable Storage for Streaming Applications


Kostas Magoutis

Institute of Computer Science (ICS)

Foundation for Research and Technology Hellas

Heraklion, Crete, Greece

TERENA TF-STORAGE Meeting – Copenhagen, September 15 2009

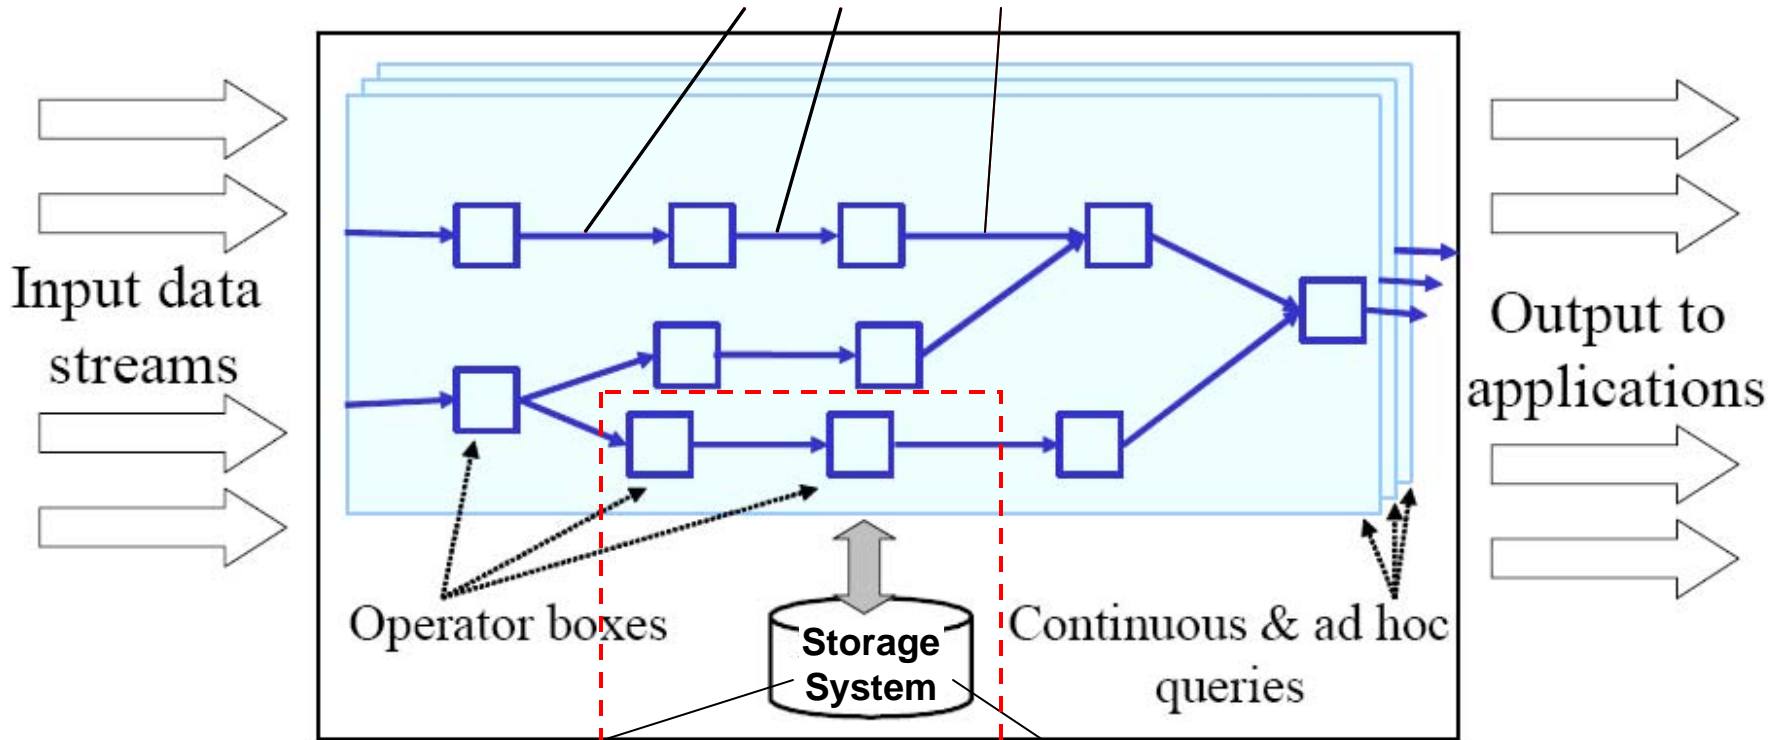


# STREAM EU Project

- Scalable Autonomic Streaming Middleware for Real-Time Processing of Massive Flows of Data
  - STREP
  - Feb 2008 - Jan 2011
  - [www.streamproject.eu](http://www.streamproject.eu)
- Partners
  - FORTH, UPM (Coordinator), Dresden
  - Telefonica, Epsilon, Exodus

# Scalable Stream Processing

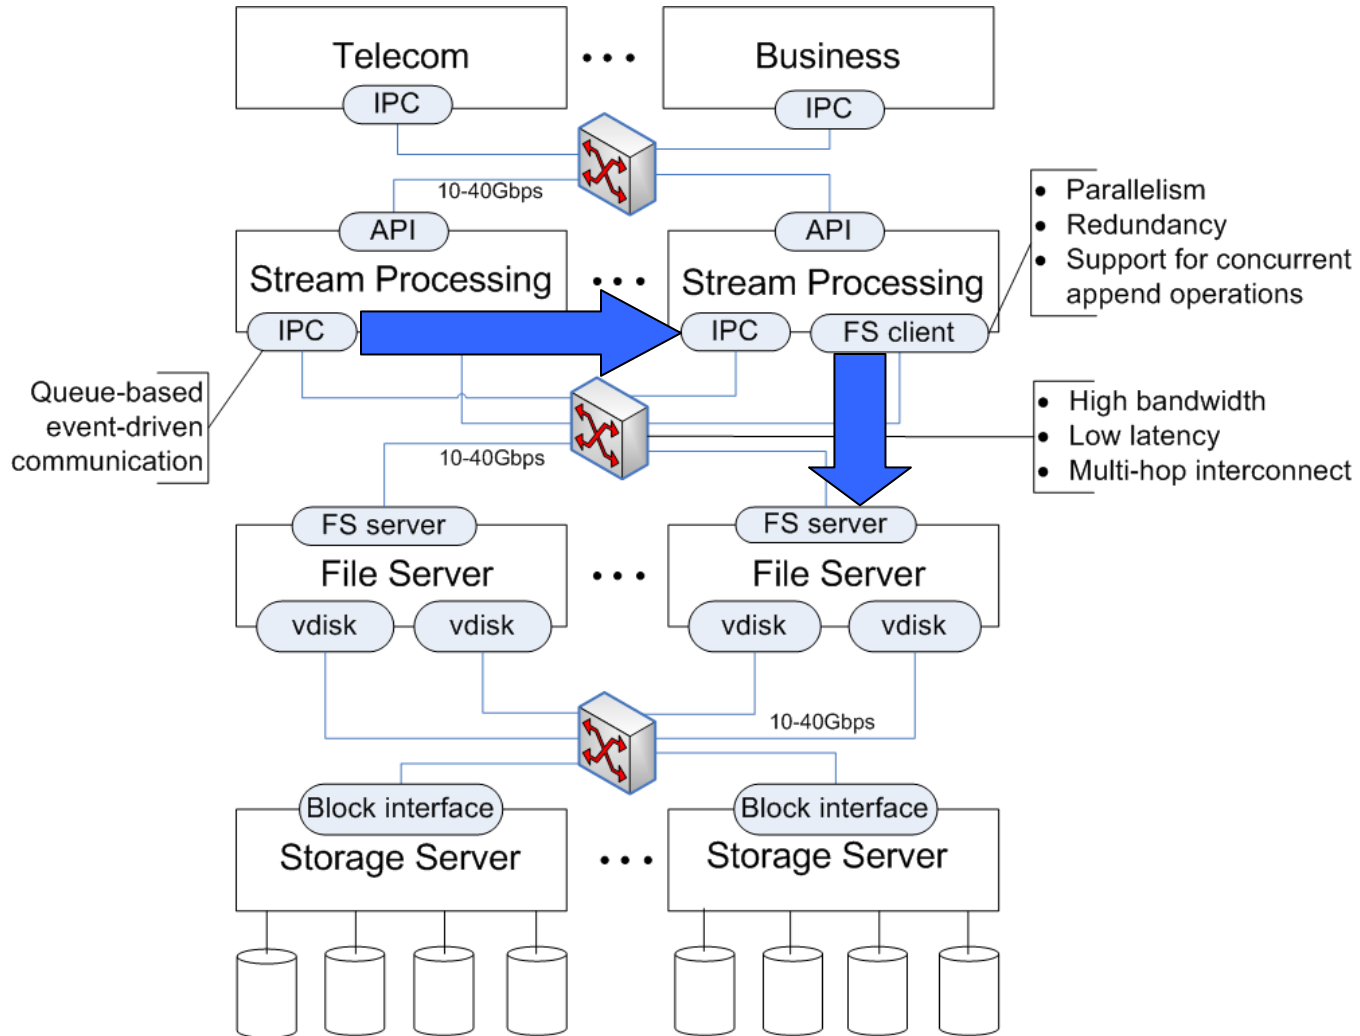
Communication between streaming operators at 40-100Gbps / node



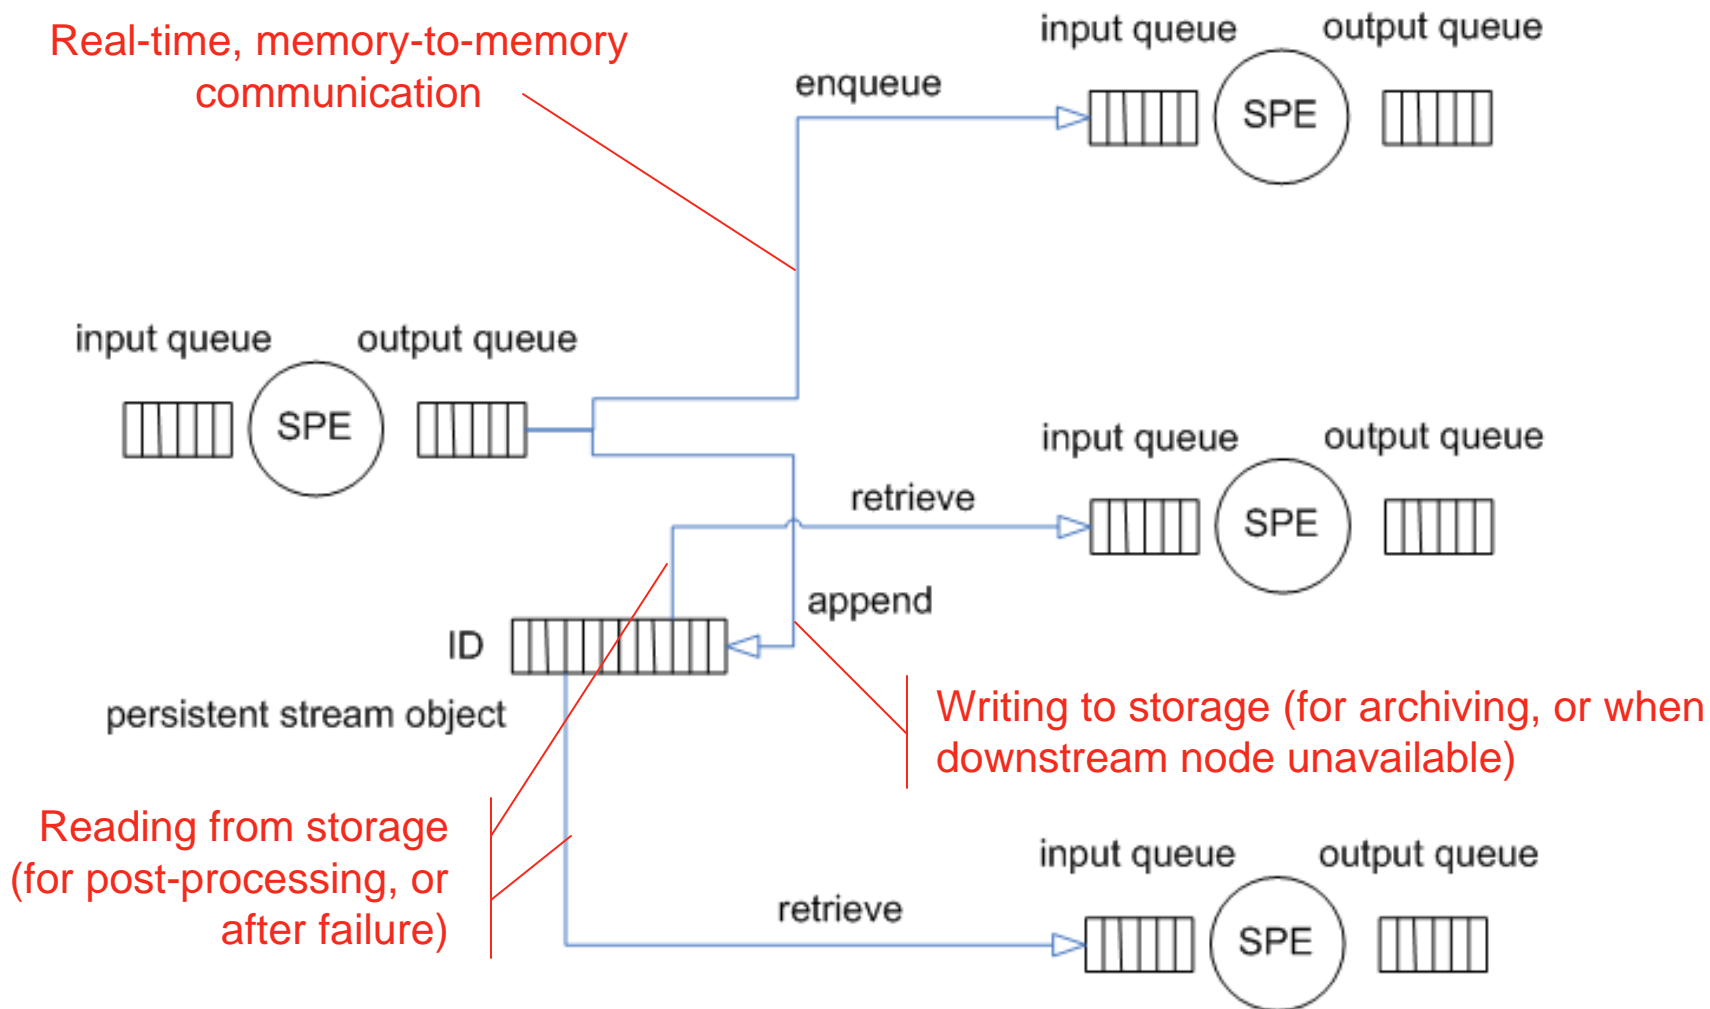
**Archiving for historical reference and/or compliance**

**Temporary storage for recovery and to absorb input traffic spikes**

# Scalable Comm & I/O Subsystem Architecture



# Persistent Stream Objects (PSOs)



# Status Overview

- Providing high performance networking substrate to data streaming middleware
  - Early results indicate significant boost over current performance, especially at speeds above 10Gbps
- Scalable streaming I/O achievable, key challenges:
  - Network overhead on the I/O path
  - Global timestamp-ordered atomic-append semantics
  - Seamless switchover between network, storage paths