

Scalable Storage Support for Data Stream Processing

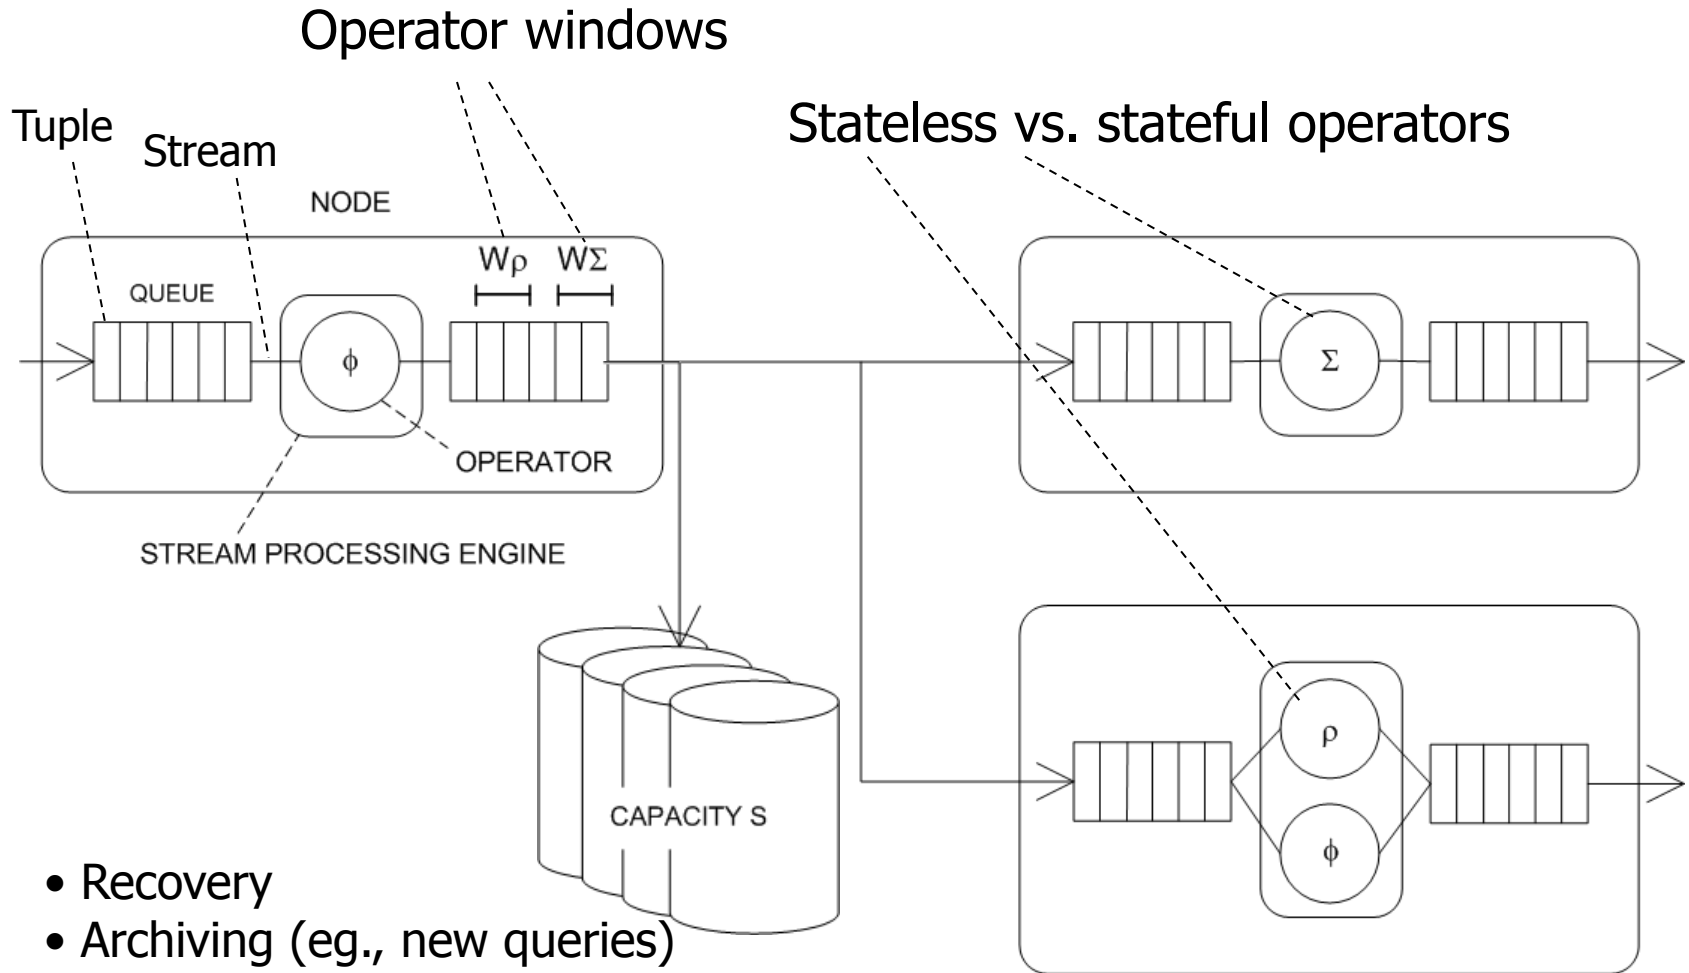
Zoe Sebeopou and Kostas Magoutis

Institute of Computer Science (ICS)

Foundation for Research and Technology Hellas (FORTH)

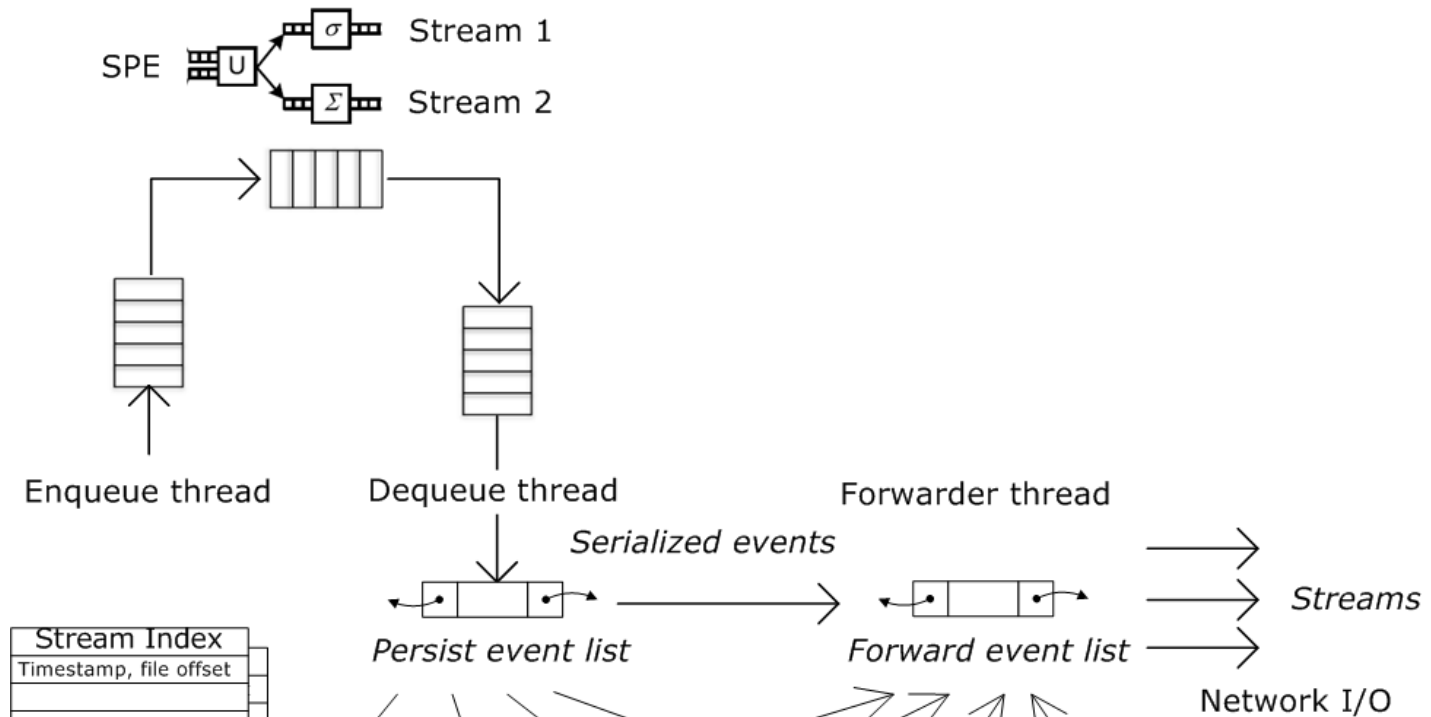
Heraklion, Greece

Data Stream Processing Model




Persistence Architecture


Borealis



PVFS2




Lessons Learned



Performance goals :

- Pipeline network/disk
- Parallelize I/Os
 - Reduce response time
- Reduce impact of metadata
 - Asynchronous updates
- Local file/disk subsystem:
 - Pre-allocation
 - Extent-based allocation
- No caching or locking

Evaluation



May 6, 2010


Thank you!

May 6, 2010

Backup

May 6, 2010

Evaluation




Experiments

- Baseline Borealis : (1,0)
- Effect of storage : (1,1)
- Storage server capacity: (8,1)
- Scalability: (m, n)

Important parameters:


- Sync/async metadata
- Local filesystem type (ext3/xfs)

Experimental Testbed




May 6, 2010

Saturate PVFS2 server: (8, 1)

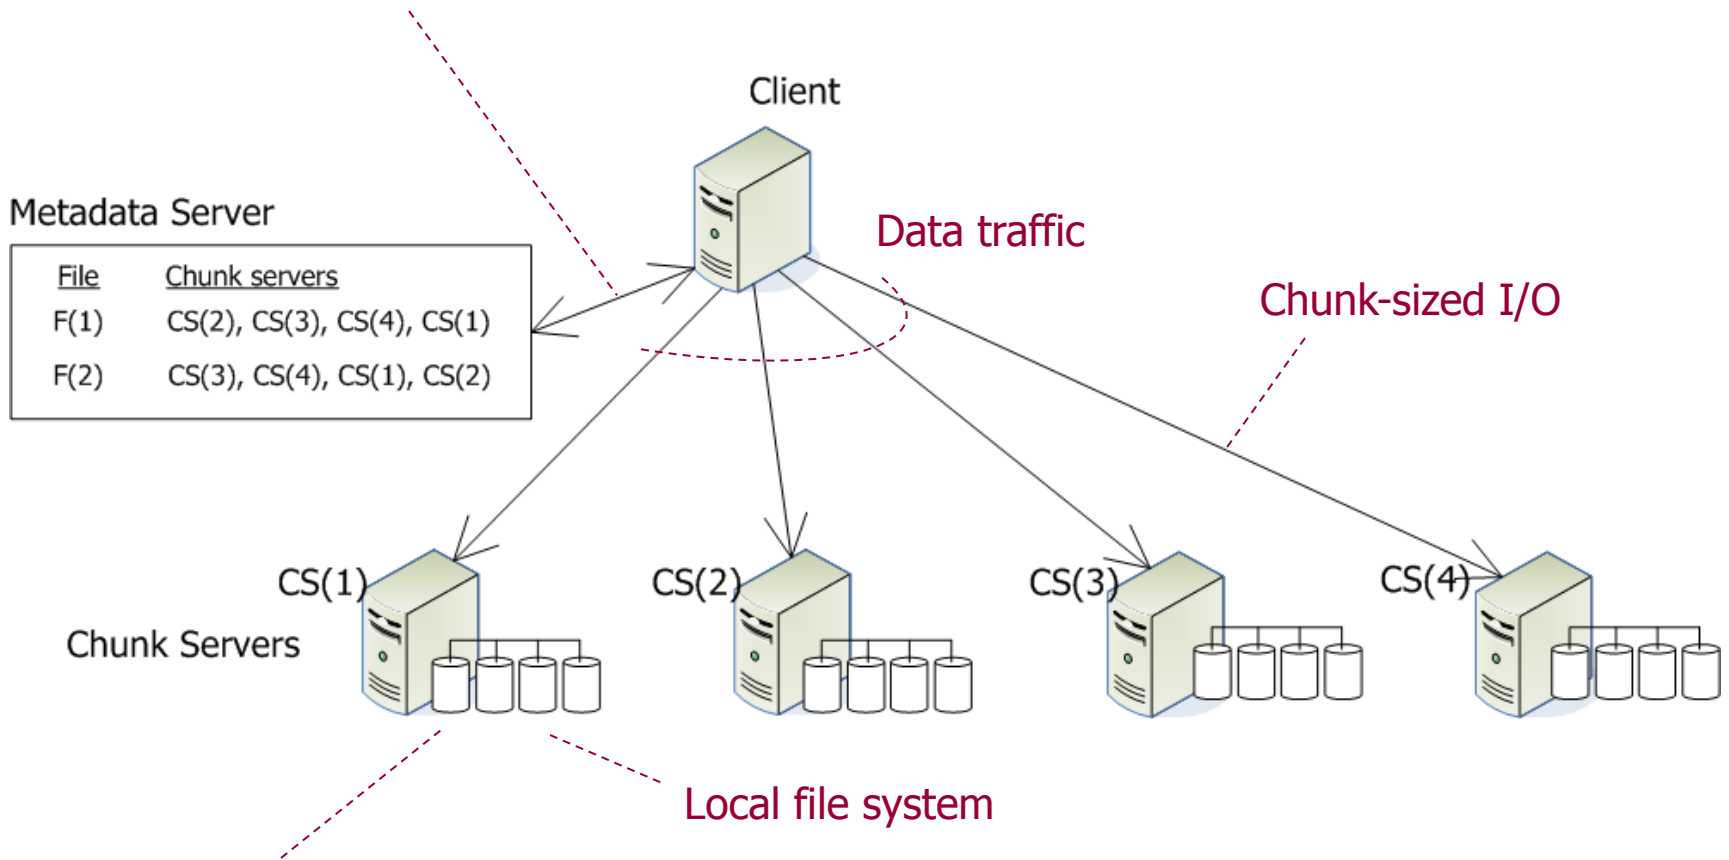


Scalability: Borealis throughput



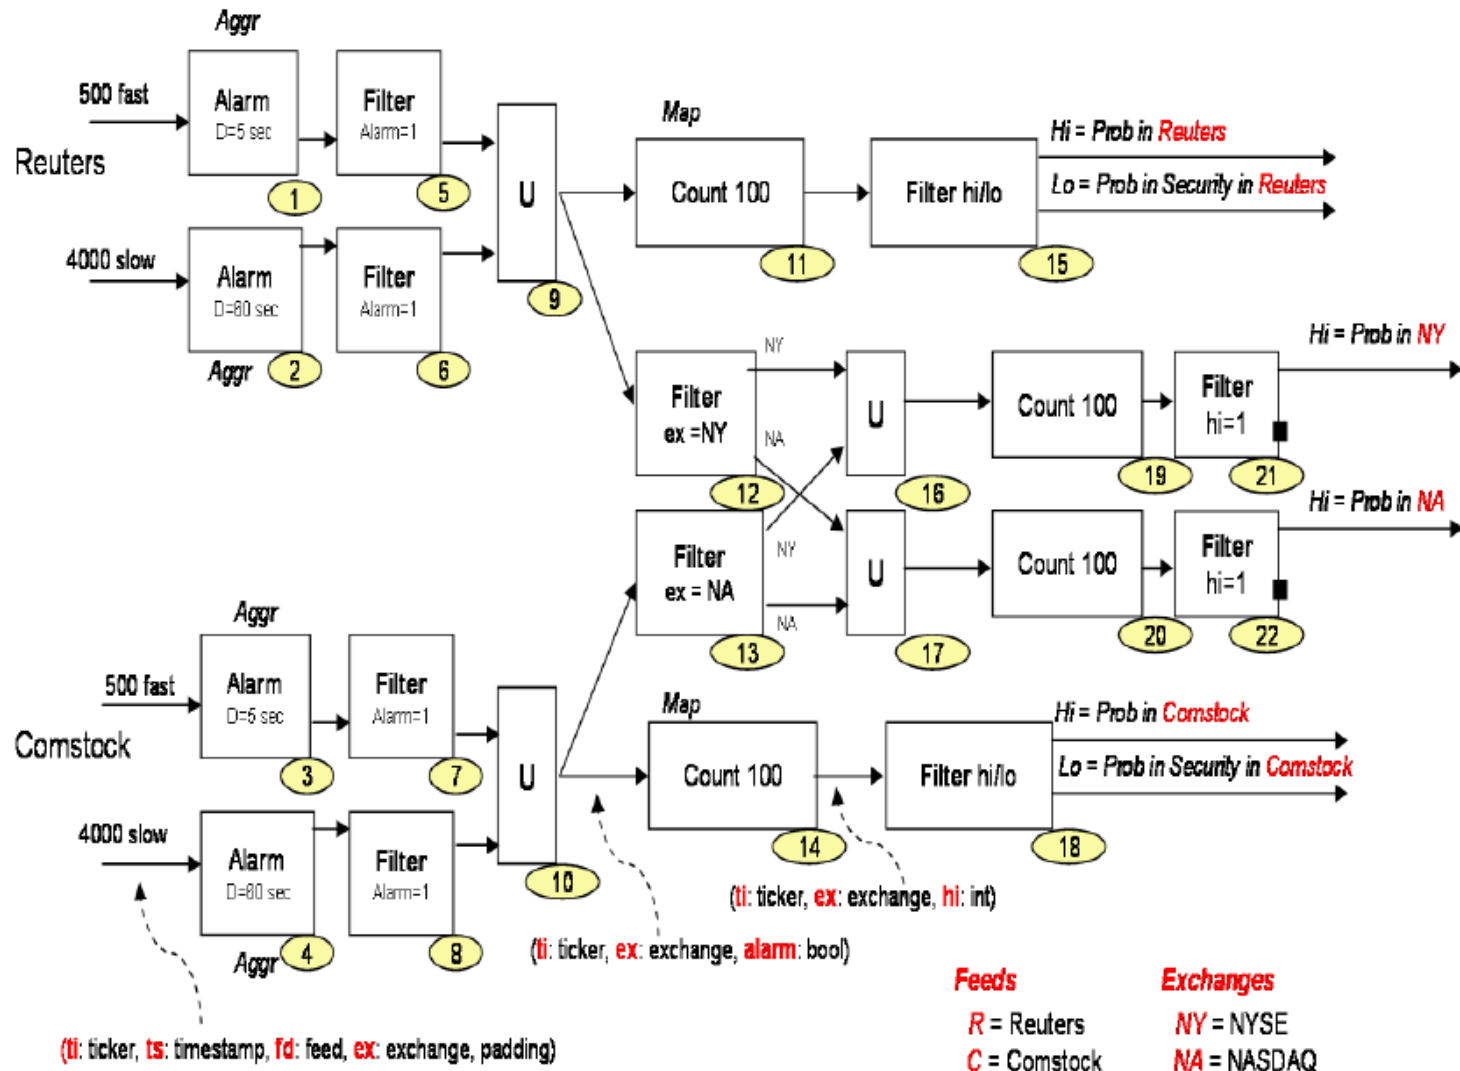
Parallel file system

Metadata traffic (create, map name-handle, update size/atime/mtime)



Local metadata traffic (create, map handle-local filename, update size/atime/mtime)

Example application: Alarm correlation



PVFS2 Architecture

