RADoN: QoS in Storage Networks

Tim Kaldewey[1]
Richard Golding[2]
Theodore M Wong[2]

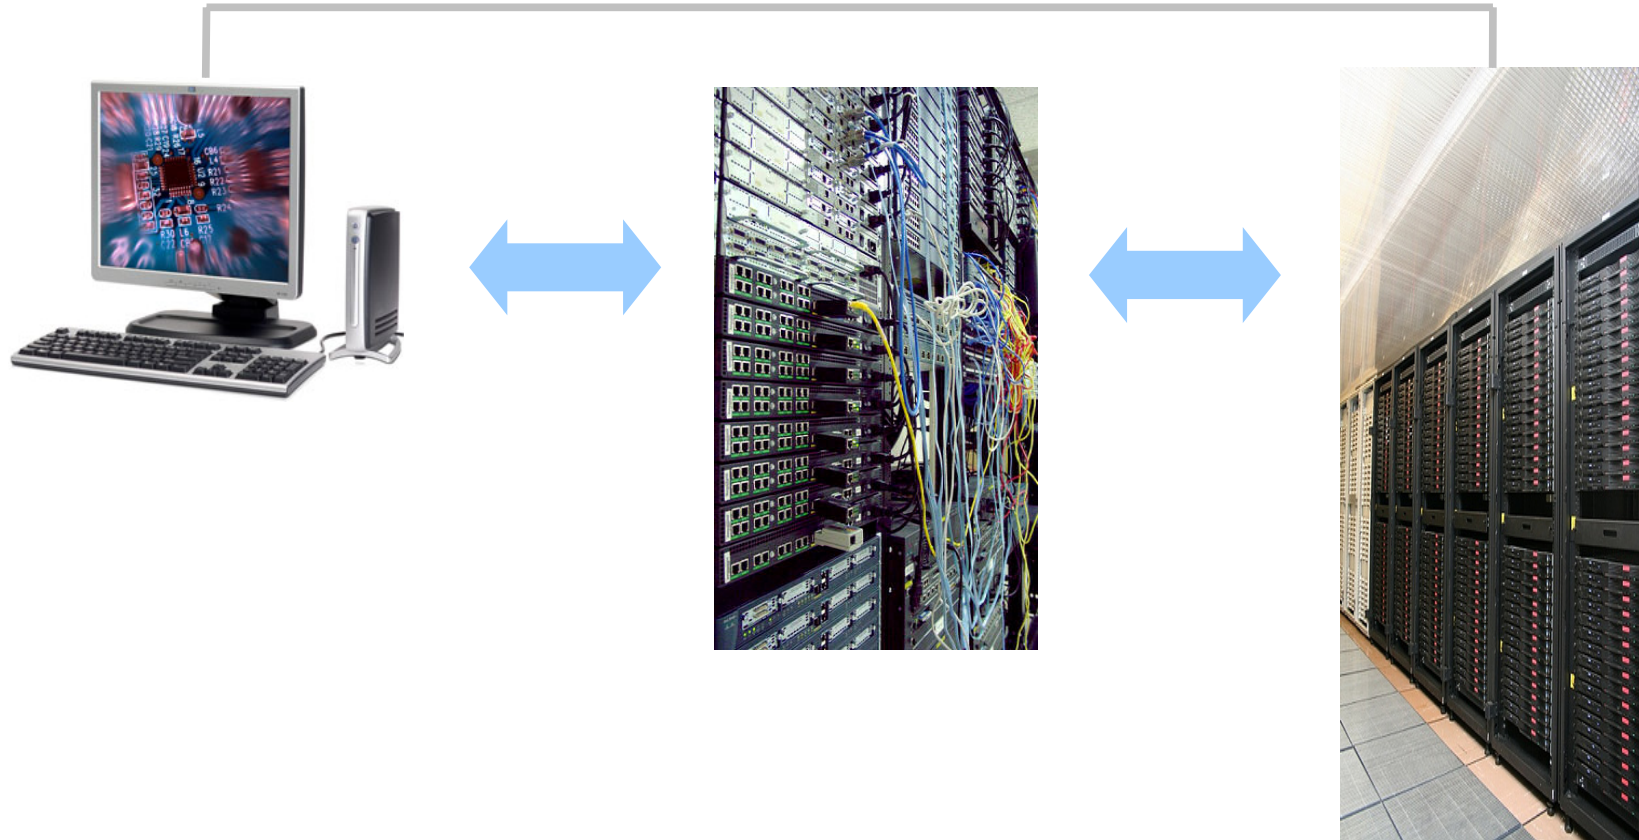Andrew Shewmaker[1]
Carlos Maltzahn[1]
Scott Brandt[1]

[1]University of California Santa Cruz
Computer Science Department
{kalt,shewa,carlosm,scott}
@cs.ucsc.edu

[2]IBM Almaden Research Center
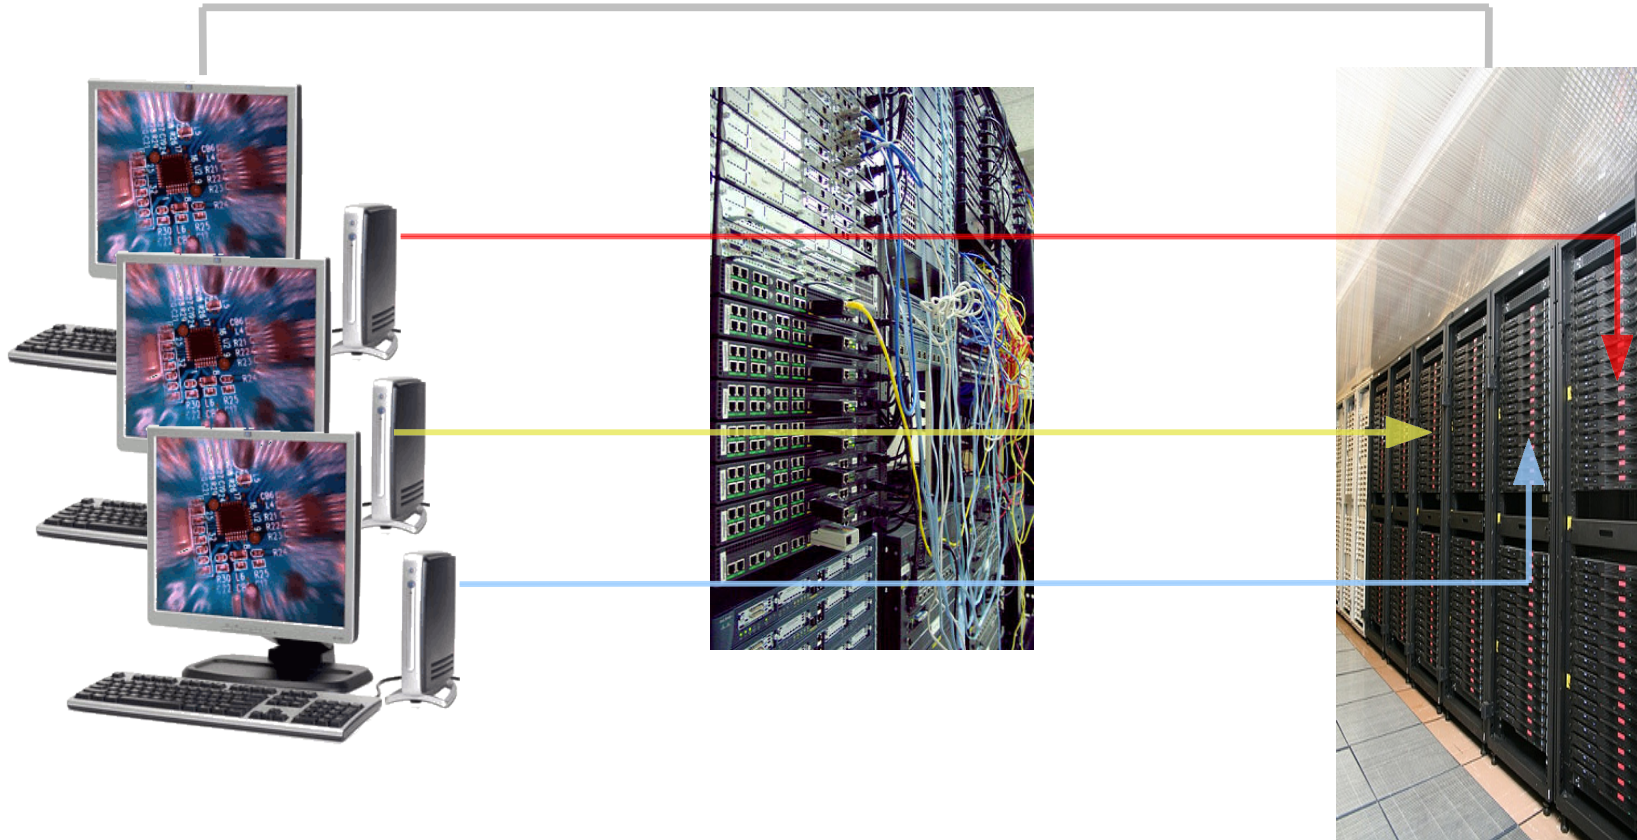Storage Systems Department
{rgolding,theowong}
@us.ibm.com

Storage QoS - XXL

End-to-End storage QoS
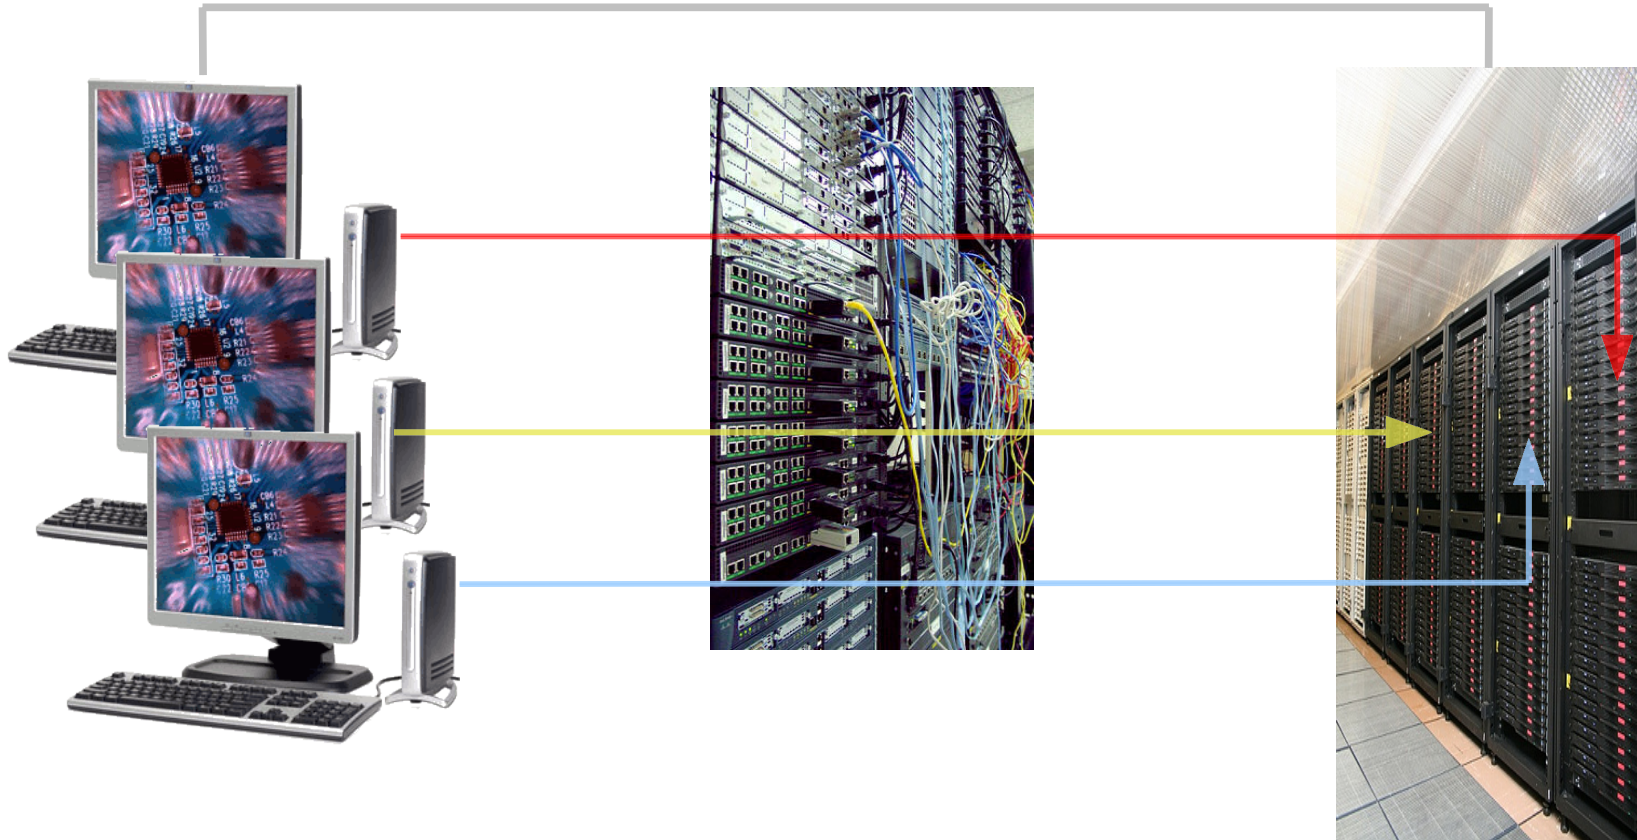
Storage QoS in practice

End-to-End storage QoS

Common approach: phys. partitioning to achieve isolation
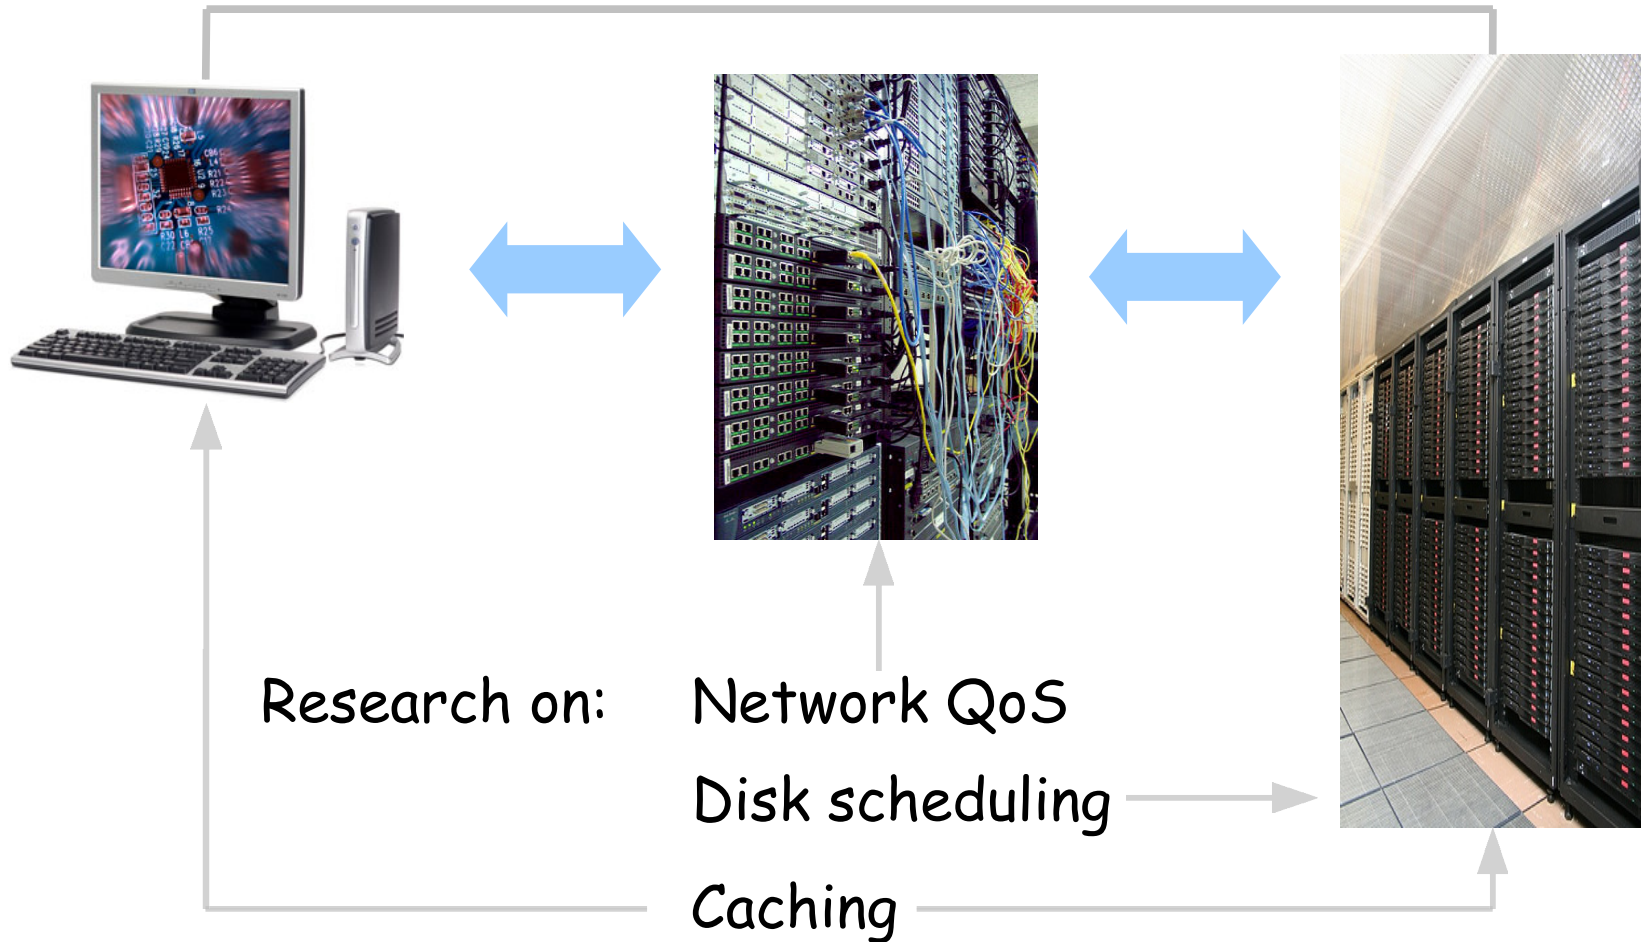
Storage QoS in practice

End-to-End storage QoS

Common approach: phys. partitioning to achieve isolation
often ➔ overprovisioning =(

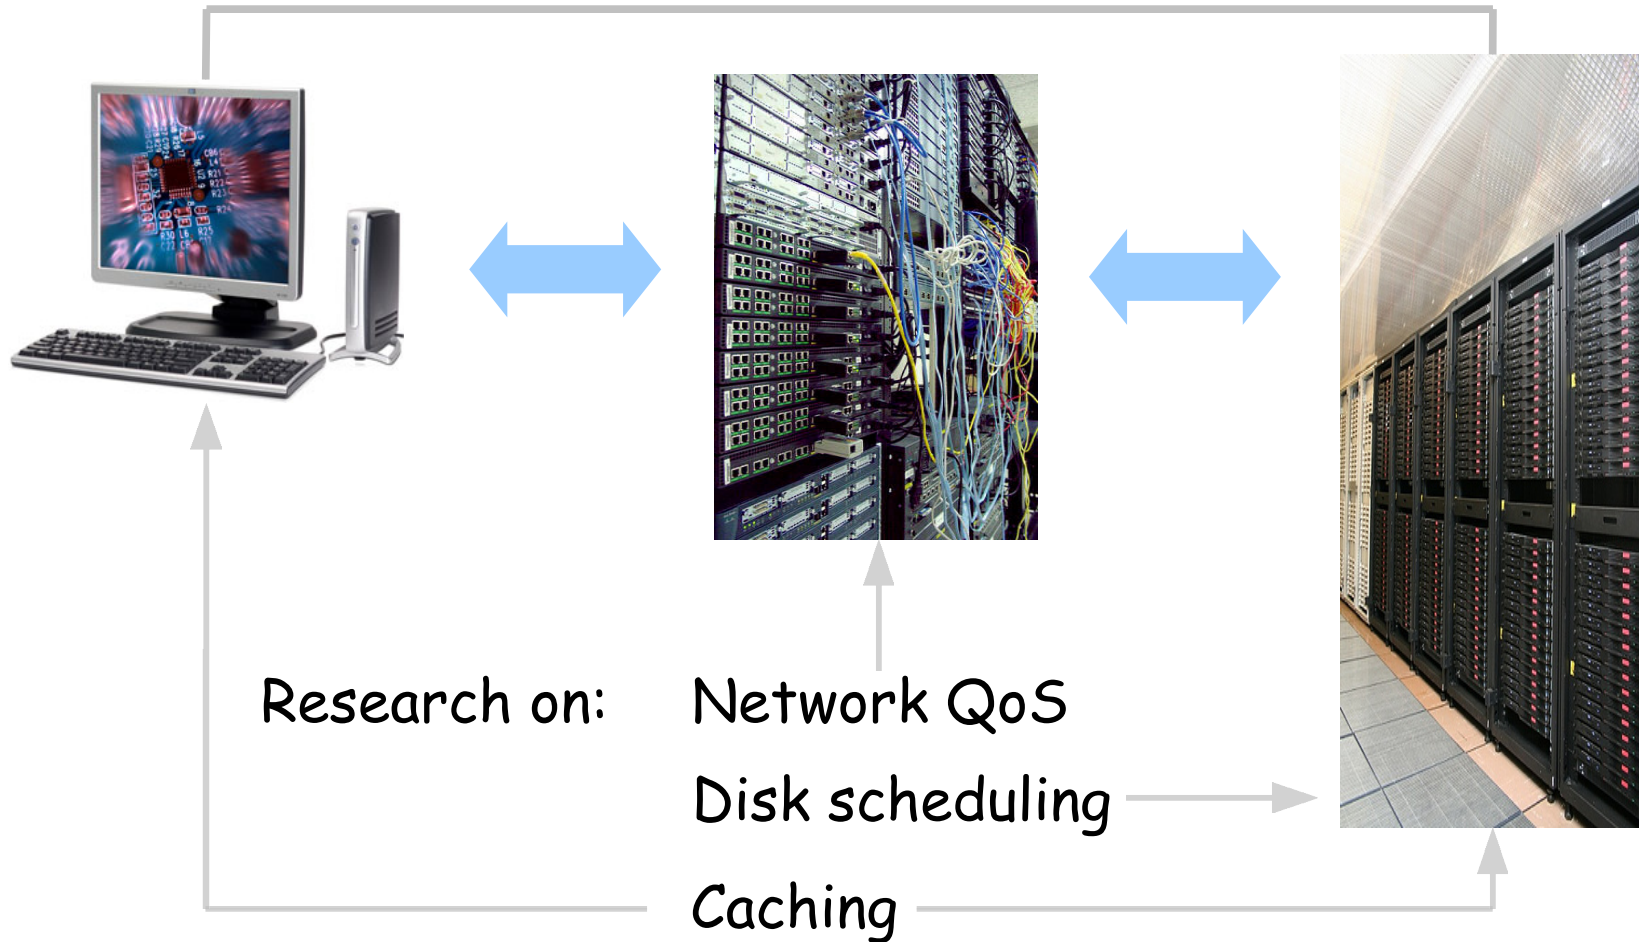# Storage QoS in research

End-to-End storage QoS

Research on: Network QoS

Disk scheduling

Caching

Storage QoS in research

End-to-End storage QoS

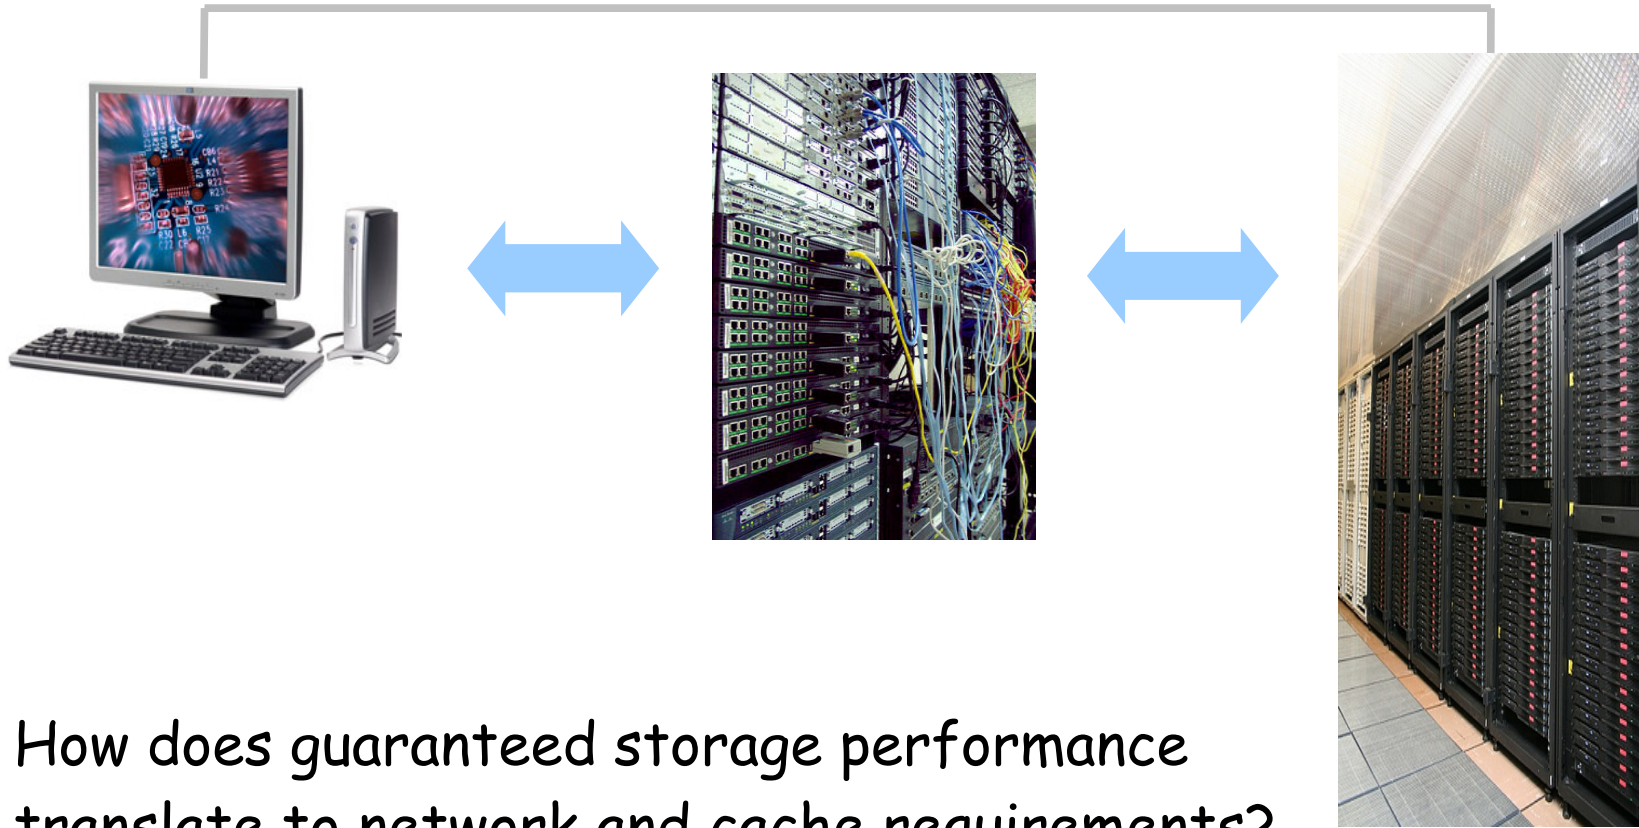Research on: Network QoS

Disk scheduling

Caching

But no integration   =(

End-to-End storage QoS ?

End-to-End storage QoS

How does guaranteed storage performance translate to network and cache requirements?

How to coordinate network, cache and storage subsystems?
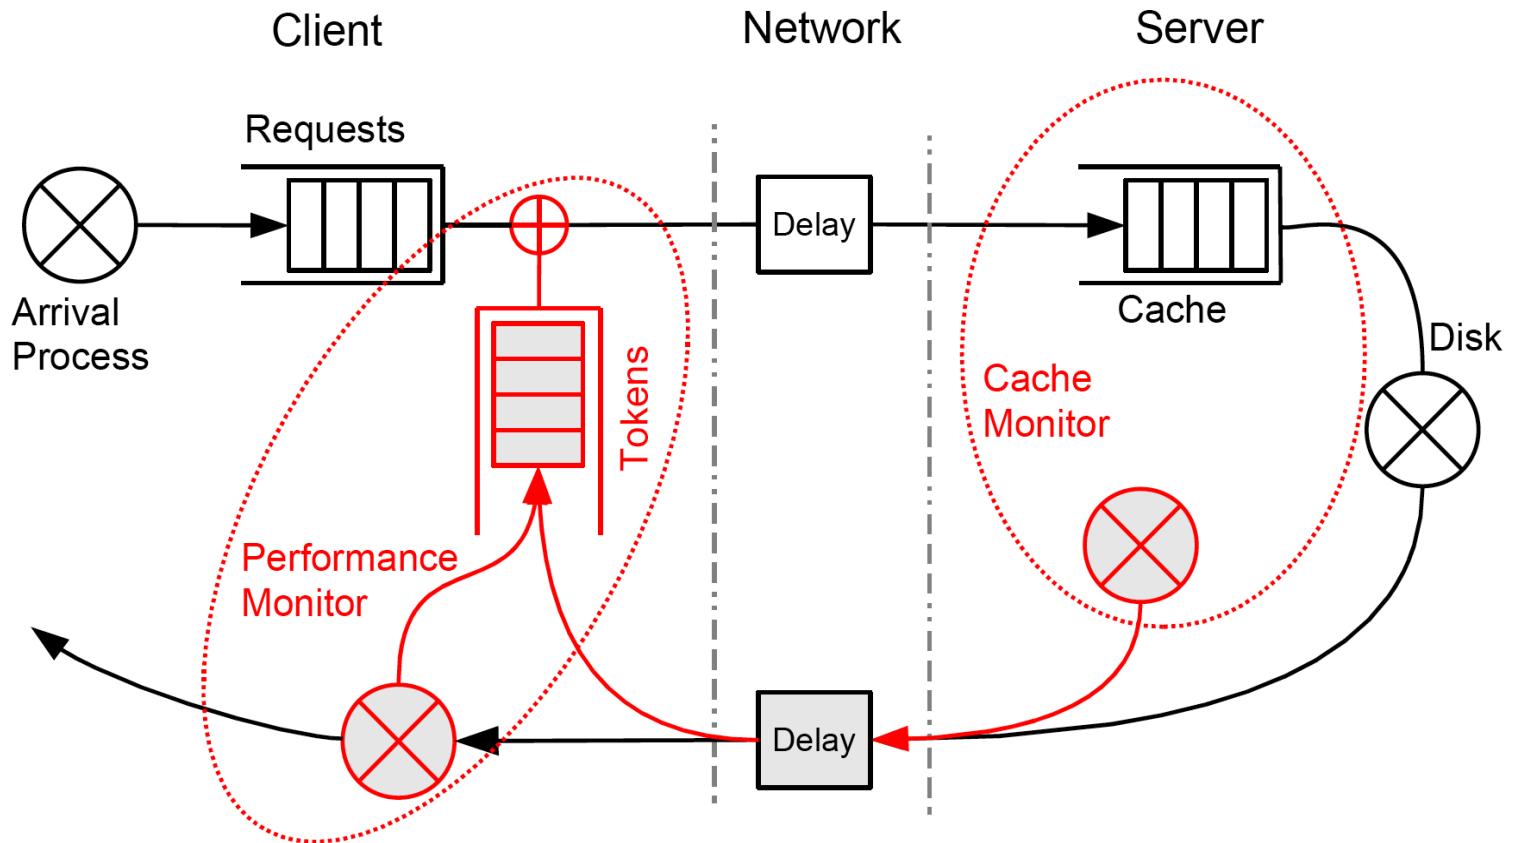
End-to-End storage QoS ?

- Large storage system ➔ Many parameters to tweak
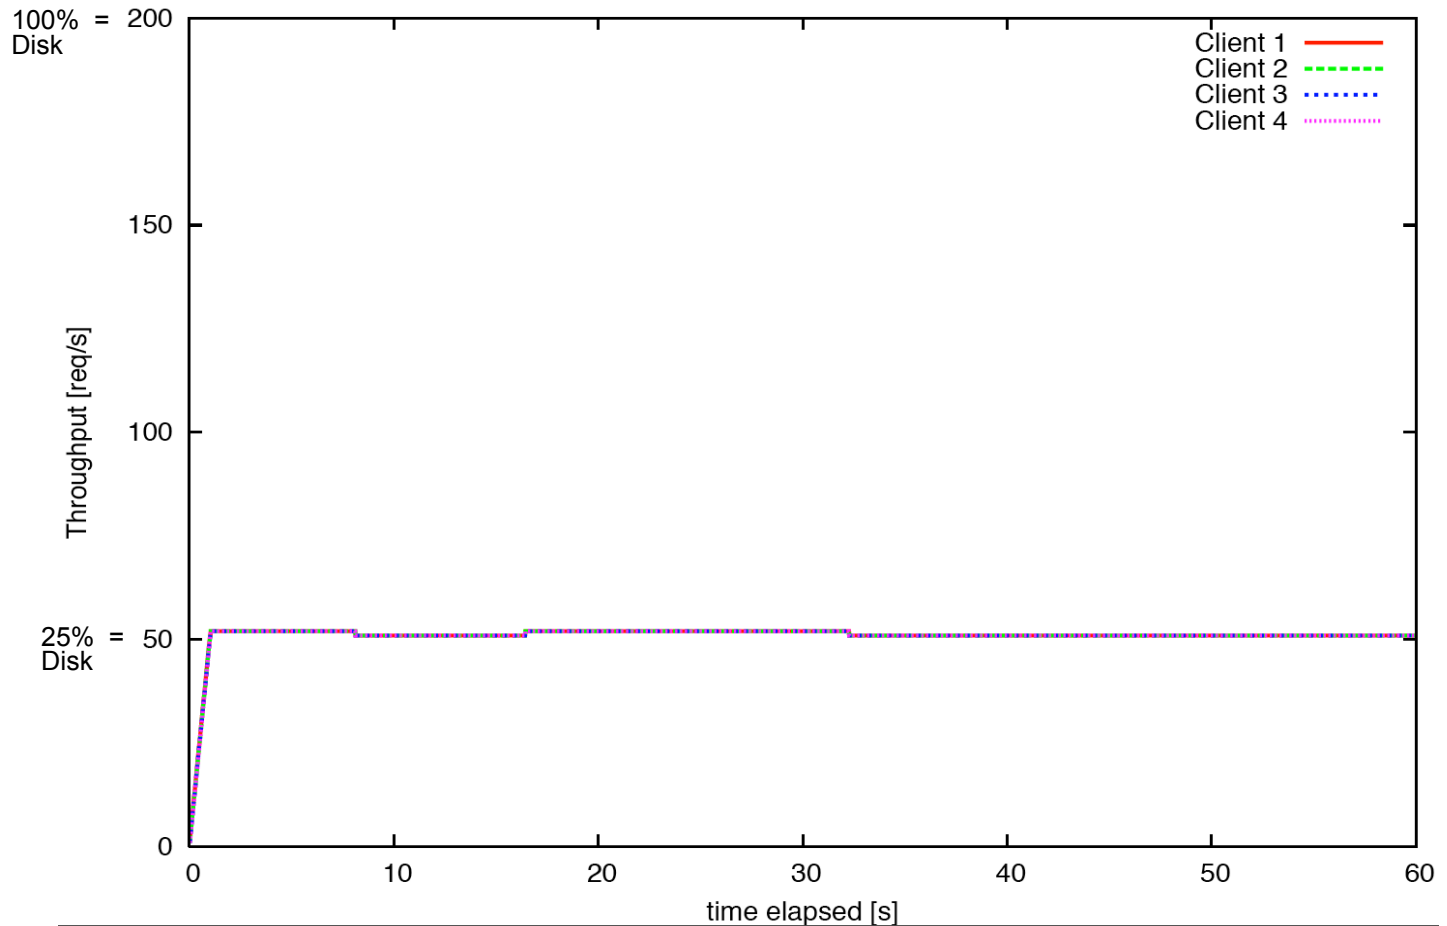
- Which are important?

RADoN – model

- Large storage system ➔ Many parameters to tweak
- Which are important? ➔ Find out via simulation:

RADoN – simulation results

Time series for throughput of 4 clients, each reserving 25% of storage performance, but producing enough results to saturate the disk itself

Work in Progress

- Simulating multiple approaches to coordinate subsystems

- Implementation on top of existing QoS disk scheduler [Fahrrad]

- Complete E2E storage QoS framework [RADI/O]

Work in Progress

- Simulating multiple approaches to coordinate subsystems

- Implementation on top of existing QoS disk scheduler [Fahrrad]

- Complete E2E storage QoS framework [RADI/O]

Long term goal:

better storage QoS to avoid