Understanding Customer Dissatisfaction With Underutilized Distributed File Servers

Towards an Architecture for Network-Attached Storage

Erik Riedel

Carnegie Mellon University

http://www.cs.cmu.edu/~riedel

Parallel Data Laboratory

Introduction

- Trends
  - » dependence on distributed file systems is becoming increasingly wide-spread
- Problem
  - » increasing complaints about (poor) distributed file system performance
- Opportunity
  - » detailed traces available

Parallel Data Laboratory

Outline

- Measurement environment
- Customer dissatisfaction
- Underutilized servers
- Factors affecting performance
- Network-attached storage
- Conclusions & future work

Measurements

- Server
  - » system-level events
    - – disk operations
    - – processor time
    - – context switches
  - » AFS statistics
    - – operation counts
    - – execution time averages
- Client
  - » AFS statistics
    - – operation counts
    - – response time averages
  - » network
    - – average ping time

Environment

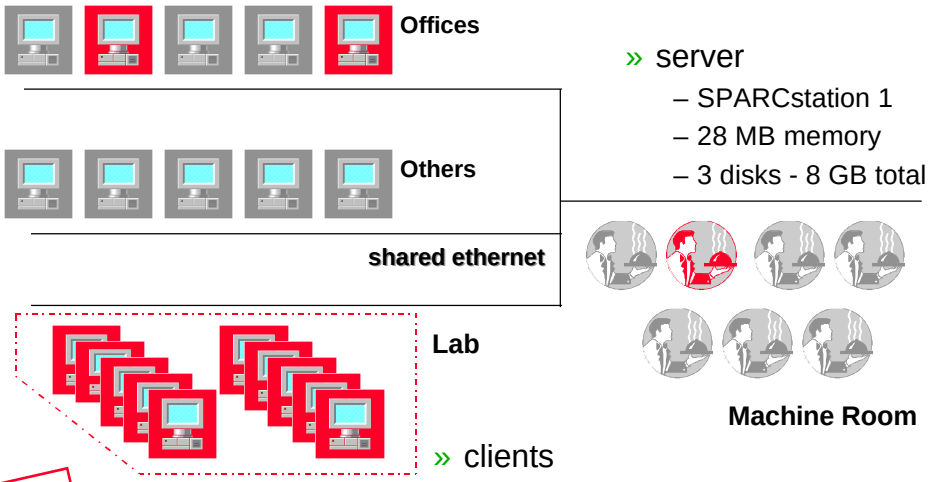

» server
- SPARCstation 1
- 28 MB memory
- 3 disks - 8 GB total

» clients
- Alpha, RS/6000
- 16 - 44 active at once

Customer Dissatisfaction

- Are user complaints justified?
  - » Yes. Order of magnitude difference in response times

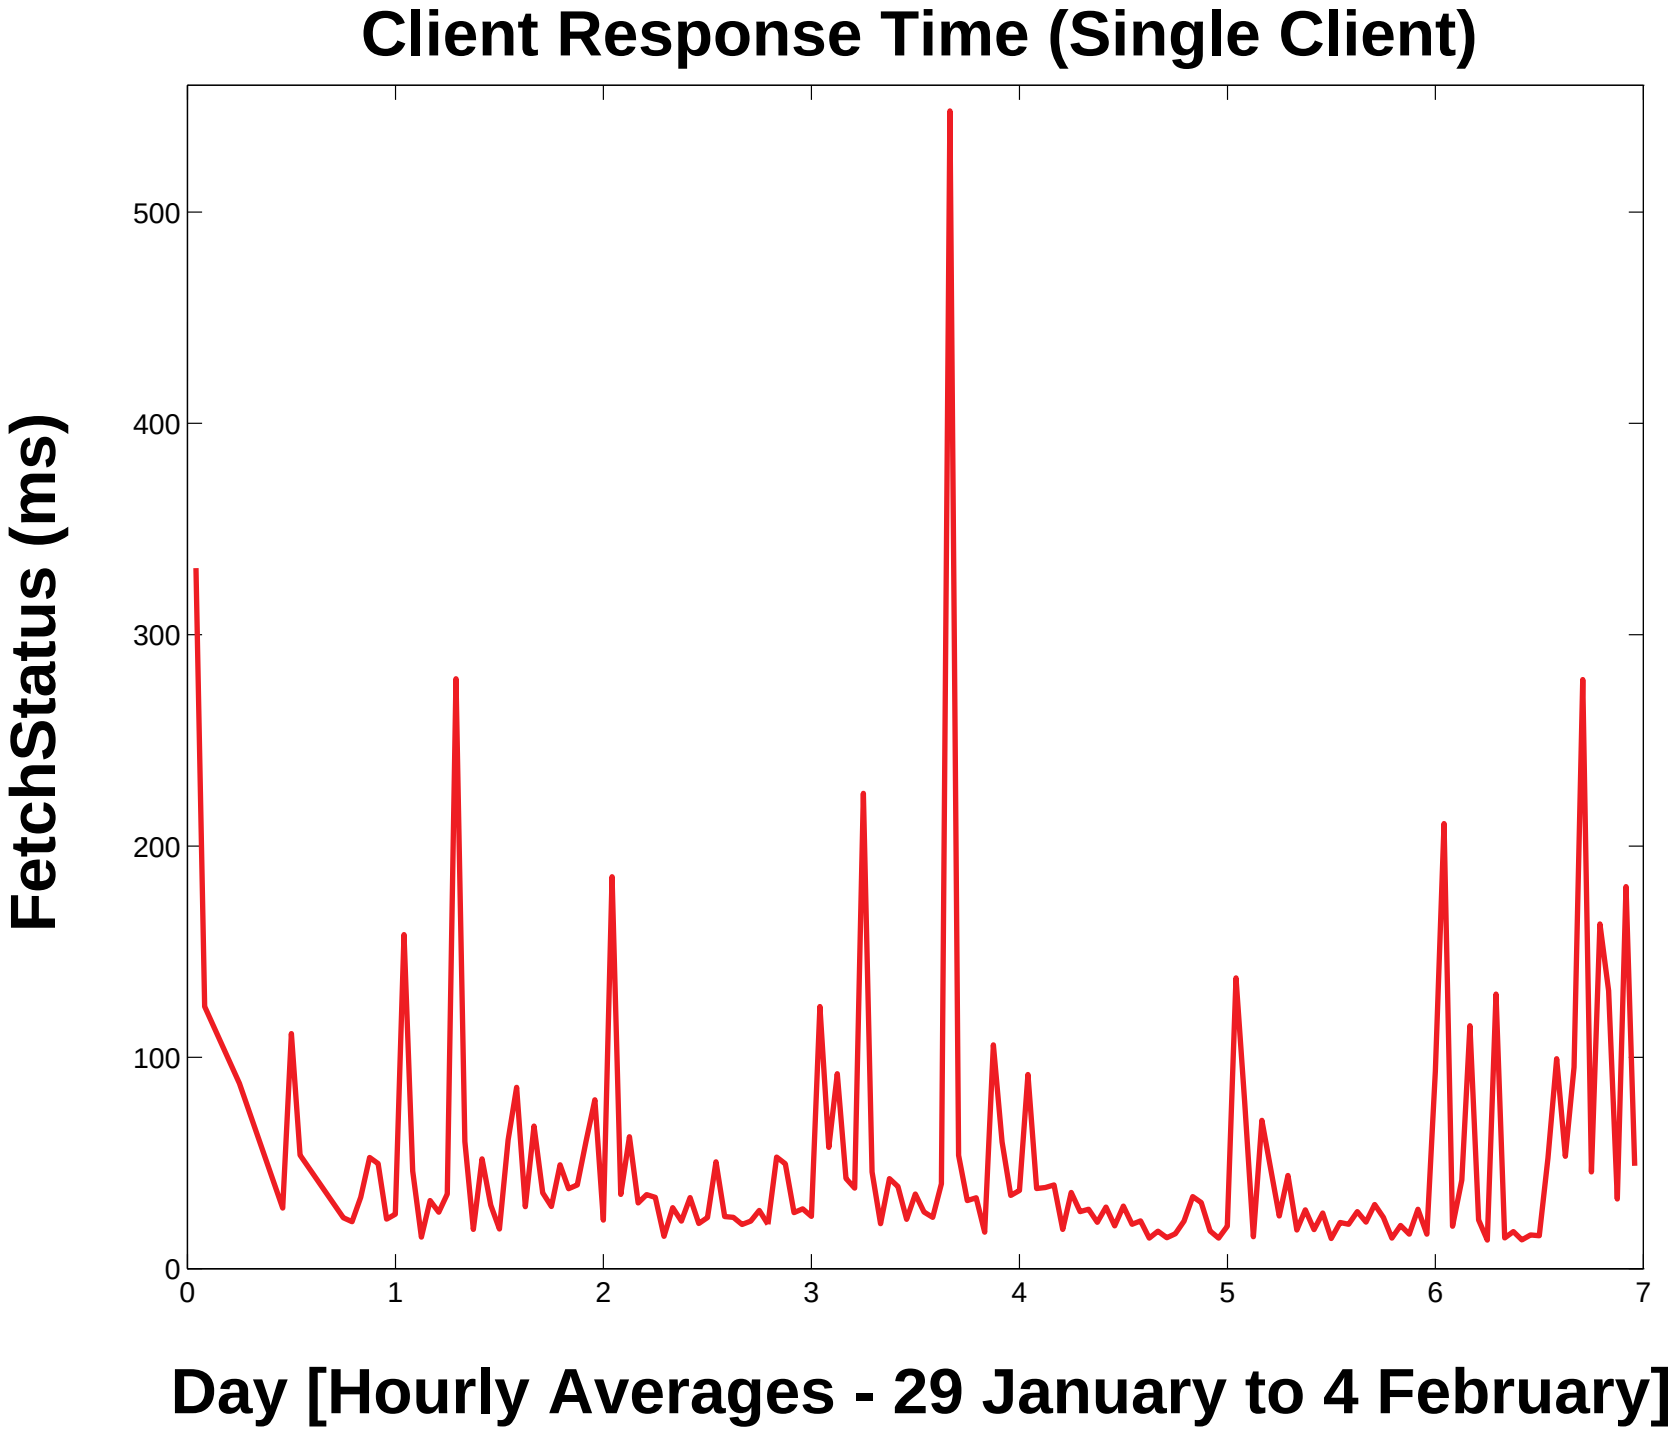
Client Response Time (Single Client)
FetchStatus (ms)
0
100
200
300
400
500
0
1
2
3
4
5
6
7
Day [Hourly Averages - 29 January to 4 February]

Underutilized File Servers

- Hypothesis
  - » Upgrading the file server machine would alleviate performance problems
- Evidence
  - » Underutilized file server
    - – 3% average CPU load

Parallel Data Laboratory

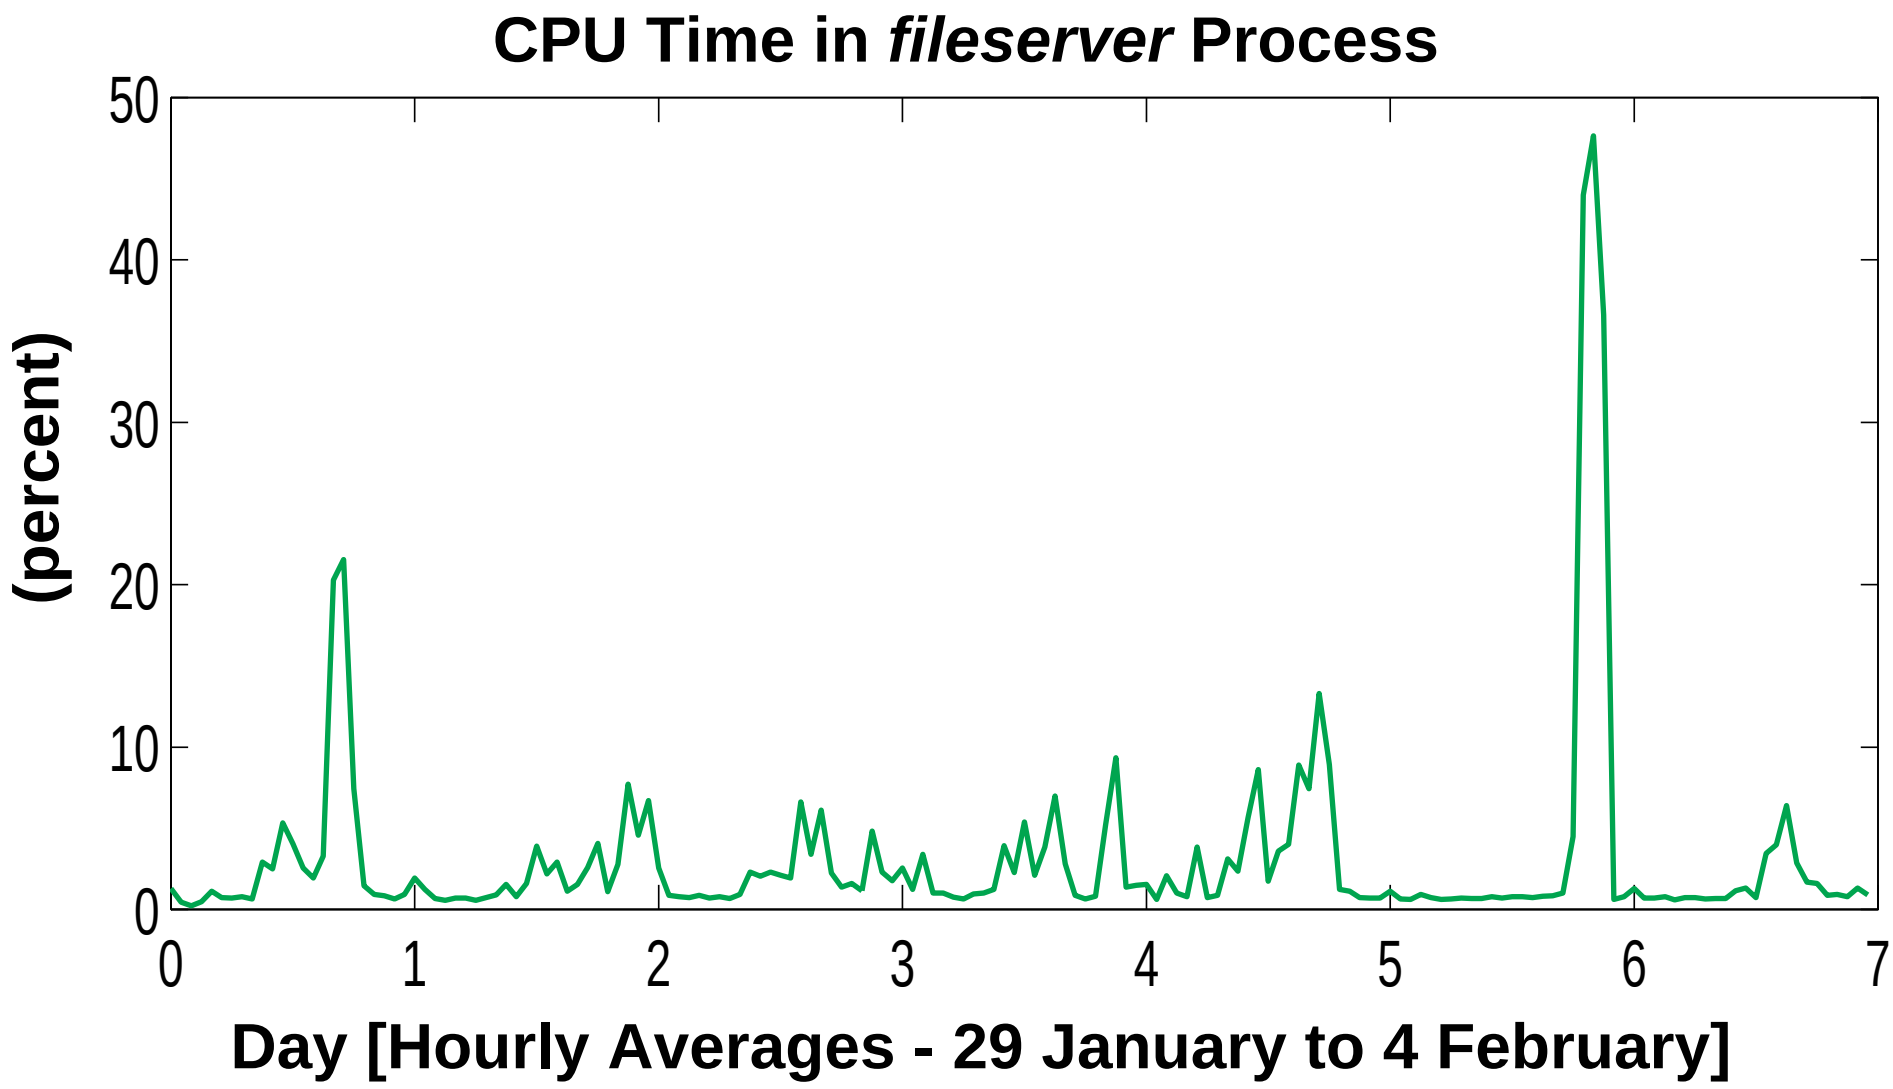
CPU Time in *fileserver* Process
(percent)
50
40
30
20
10
0
0
1
2
3
4
5
6
7
Day [Hourly Averages - 29 January to 4 February]

Environment

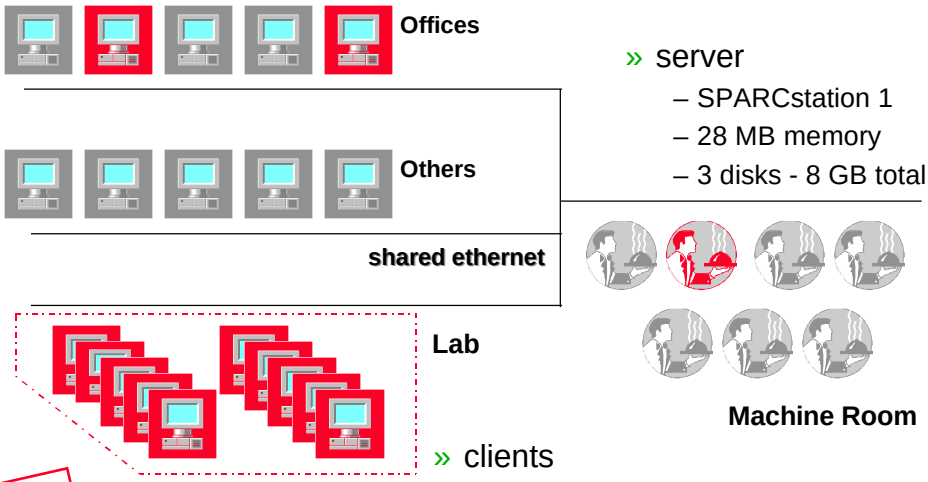

» server
- SPARCstation 1
- 28 MB memory
- 3 disks - 8 GB total

» clients
- Alpha, RS/6000
- 16 - 44 active at once

Carnegie Mellon
Parallel Data Laboratory

Effects of Network Load

- Evidence
  - » network performance explains 35% of the variance in response time
- Conclusion
  - » avoid overloading your network
  - » doesn't explain all of the problem

Carnegie Mellon
Parallel Data Laboratory

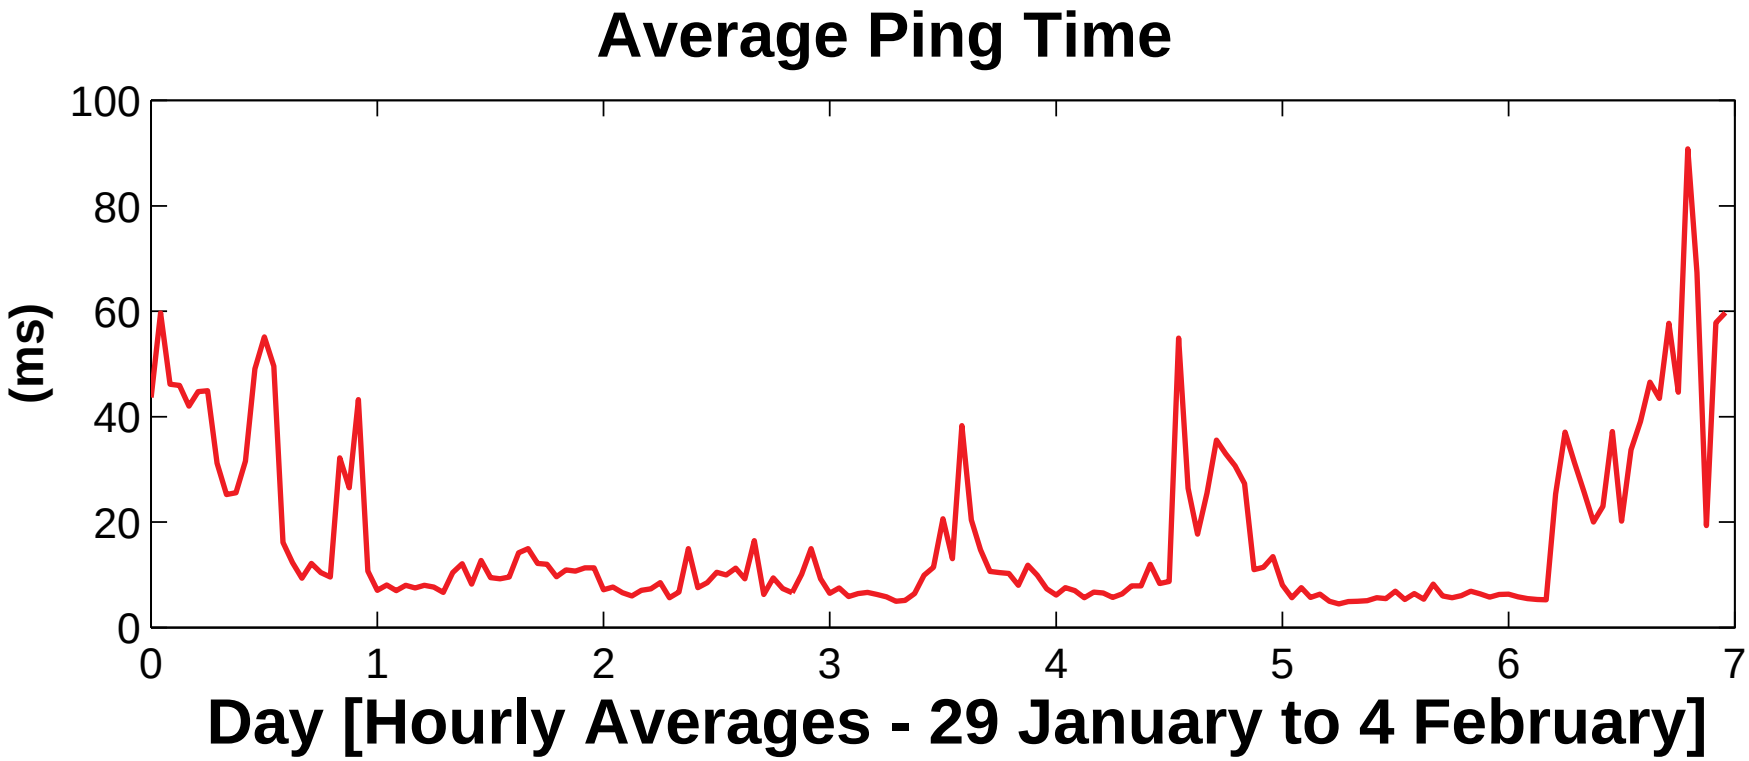
Average Ping Time
100
80
60
40
20
0
(ms)
0
1
2
3
4
5
6
7
Day [Hourly Averages - 29 January to 4 February]

Environment

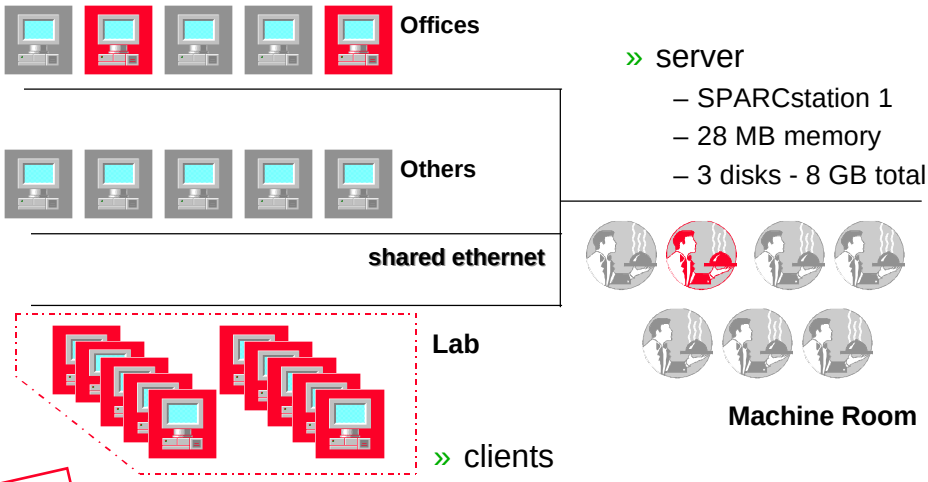

» server
- SPARCstation 1
- 28 MB memory
- 3 disks - 8 GB total

» clients
- Alpha, RS/6000
- 16 - 44 active at once

Carnegie Mellon
Parallel Data Laboratory

Effect of Shared Servers

- Problem
  - » shared binaries, common files
- Evidence
  - » high correlation among clients of the same system type
  - » no correlation across system types
- Conclusion
  - » be aware of shared resources
  - » avoid hot spots

Carnegie Mellon
Parallel Data Laboratory

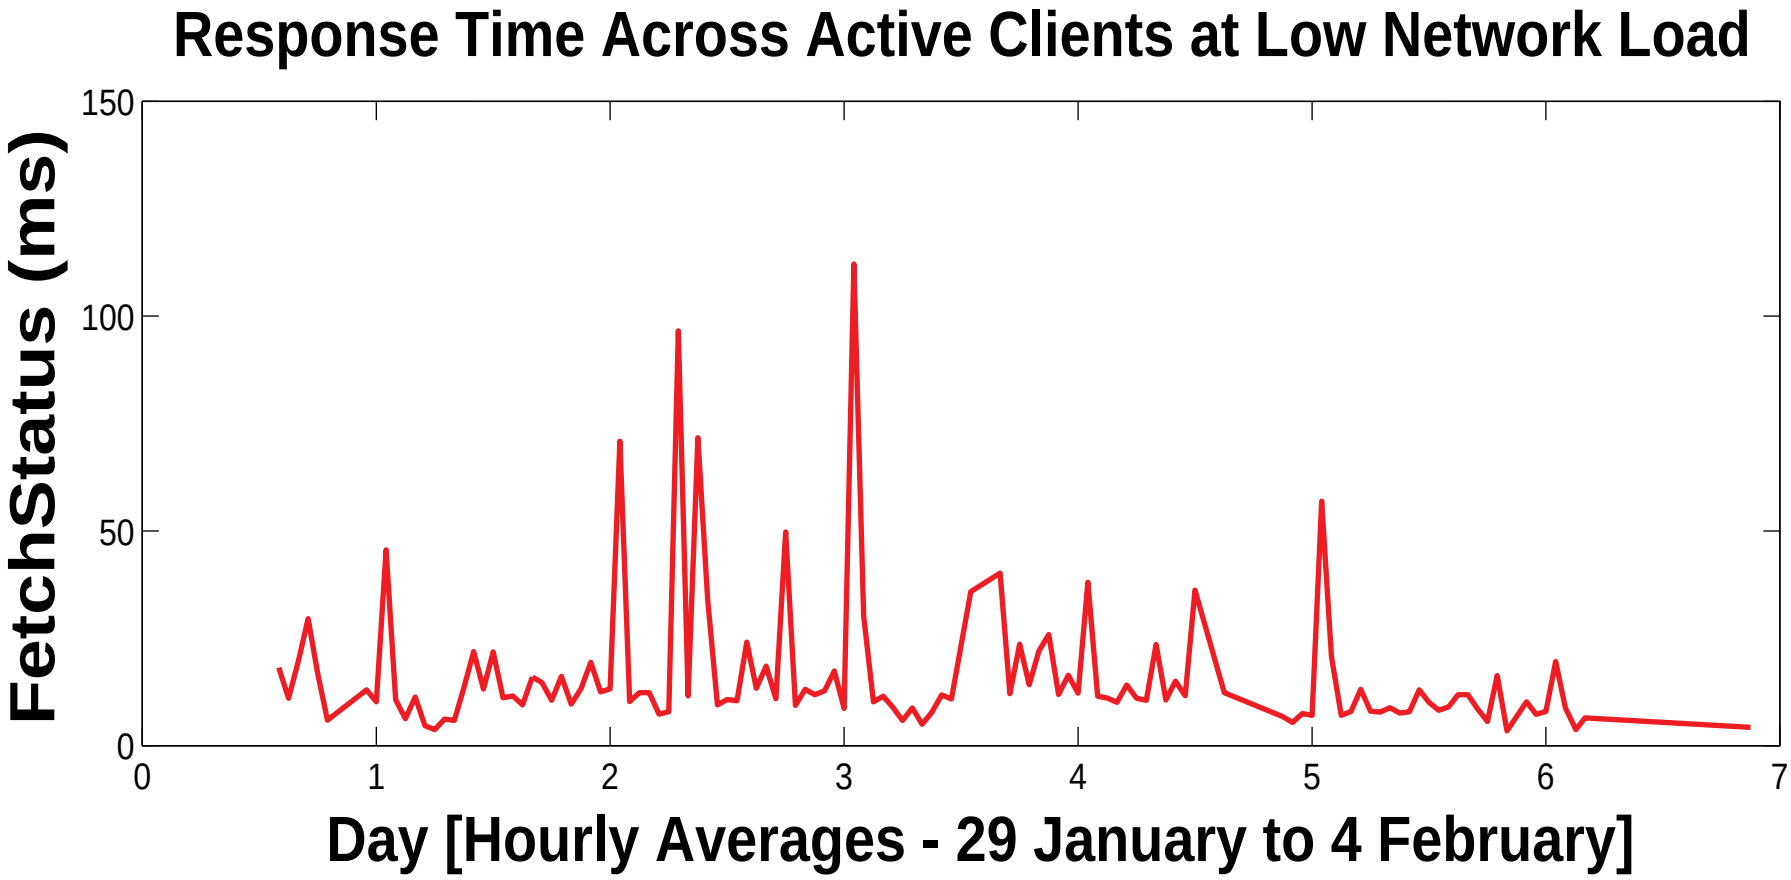
Response Time Across Active Clients at Low Network Load
FetchStatus (ms)
150
100
50
0
0
1
2
3
4
5
6
7
Day [Hourly Averages - 29 January to 4 February]

Effects of Server Load

- After the effects of the network and shared resources have been removed
  - » disk utilization
    - – explains 25% of the remaining variance
  - » processor load
    - – explains 50% of the remaining variance

Carnegie Mellon
Parallel Data Laboratory

Burst Behavior

- Larger fraction of data-moving operations
  - » exactly what file servers are worst at

Carnegie Mellon
Parallel Data Laboratory

Server Operations	Weekly Total			Peak Hour	
	total	fraction	hourly	total	fraction
FetchStatus	412,695	70.6%	1,247	6,209	45.3%
StoreStatus	22,642	3.9%	134	175	1.3%
FetchData	**62,288**	**10.7%**	**370**	**4,219**	**30.8%**
StoreData	**32,414**	**5.5%**	**192**	**147**	**1.1%**
CreateFile	17,089	2.9%	101	52	0.4%
RemoveFile	20,422	3.5%	122	2,587	18.9%
GiveUpCallbacks	17,298	3.0%	103	326	2.4%
total	584,848		2,269	13,715	

Reasons for Dissatisfaction

- A number of factors affect performance as seen by users
  - » network congestion
  - » “imbalance” in system resources
  - » performance of the file server machines
- Burst behavior is the key to understanding file server performance
- Data movement much heavier in bursts

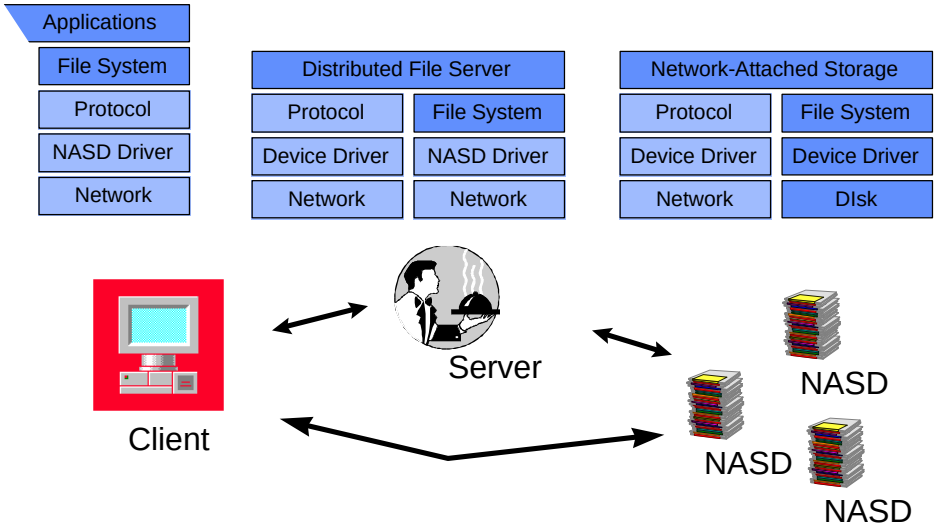

Network-Attached Storage

Network-Attached Storage

- Potential
  - » efficient fastpath transfer
  - » higher level interface
  - » “smarter” storage
  - » scalability
- Goal
  - » improve customer satisfaction