

# jSCSI


## A Java iSCSI Initiator

Marcel Waldvogel  
University of Konstanz  
1340

**JAZOON07**


THE INTERNATIONAL CONFERENCE ON JAVA TECHNOLOGY  
JUNE 24 - 28, 2007 ZÜRICH

Universität Konstanz



# Agenda

- > Introduction
  - iSCSI revisited
  - YAll: yet another iSCSI initiator?
  
- > Current work
  - Device interface
  - Implemented devices
  - Whiskas
  - Lucene backend
  - Benchmarks
  
- > Future work

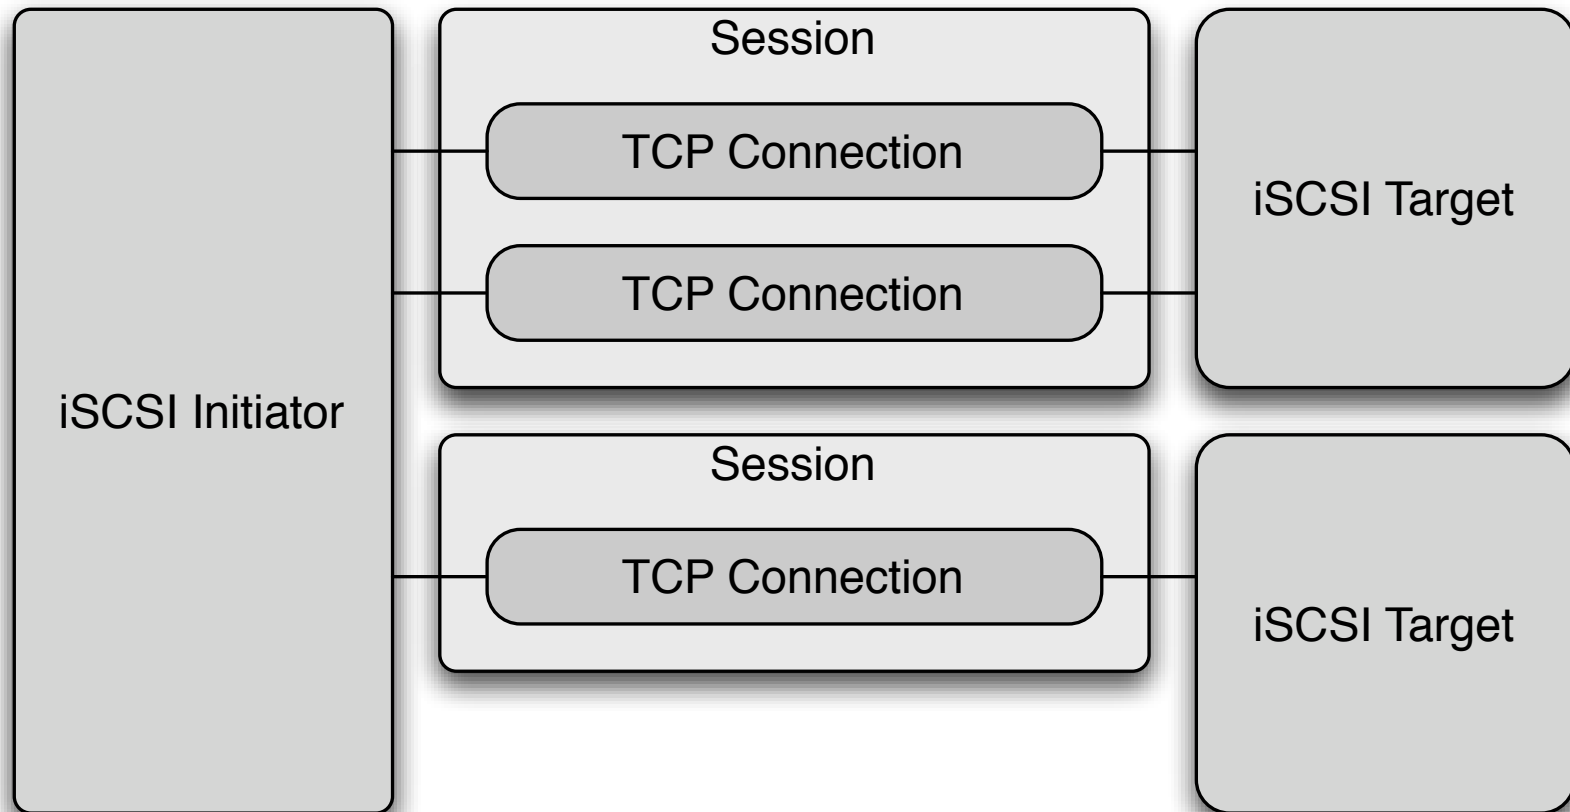


# iSCSI revisited

- > “SCSI over TCP/IP”
- > RFC 3720 (April 2004)
- > cheaper than Fibre Channel
  - Fibre Channel going IP
- > no infrastructure changes



# iSCSI revisited



# YAll: yet another iSCSI initiator?

- > First Java iSCSI initiator
- > Platform-independent device
- > Avoids kernel loop
- > Education
- > Terabyte-sized iSCSI RAID on your LAN
- > Multiple Connections per Session
  - No target available



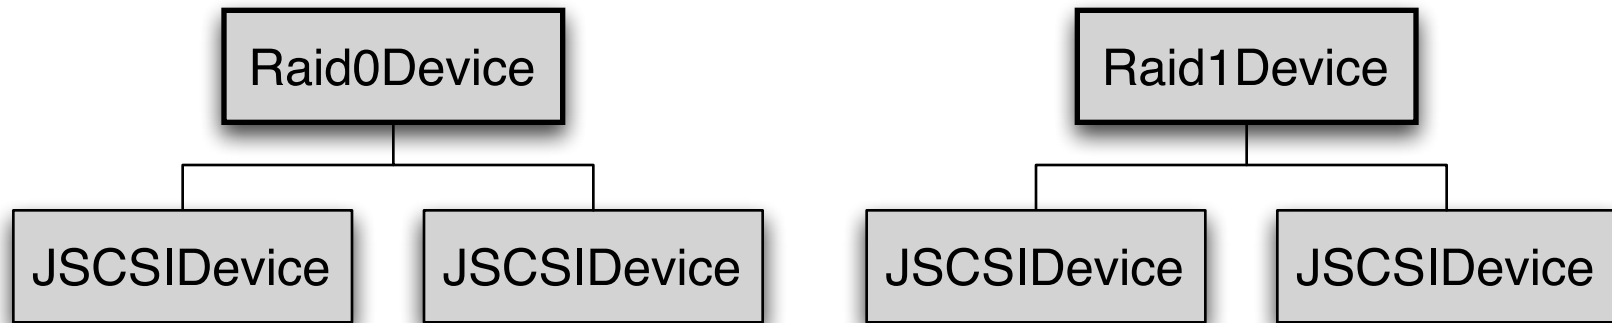
# Device interface

- > Multithreaded
- > Synchronous execution
- > Stackable devices



<b>Device</b>
<pre> + open() : void + getName() : String + getBlockSize() : int + getBlockCount() : long + read(address : long, data : byte[ ]) : void + write(address : long, data : byte[ ]) : void + close() : void           </pre>



# RAID 0, 1



# Write Buffering, Prefetching





Whiskas Pattern

New Histogram New Pattern

Reset Accumulation JSCSI Server:

localhost

Type-File:

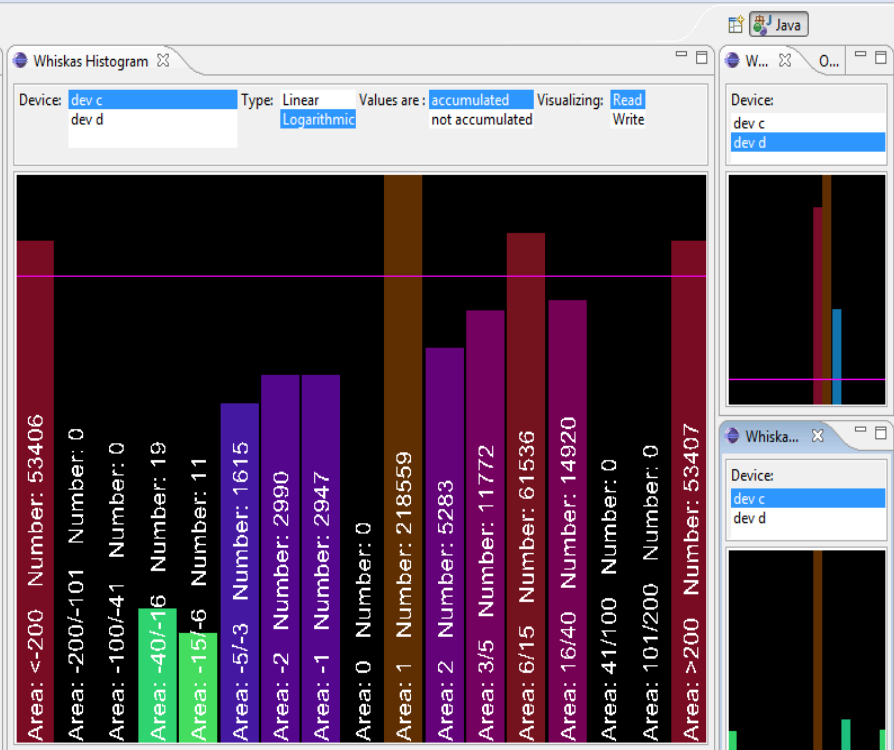
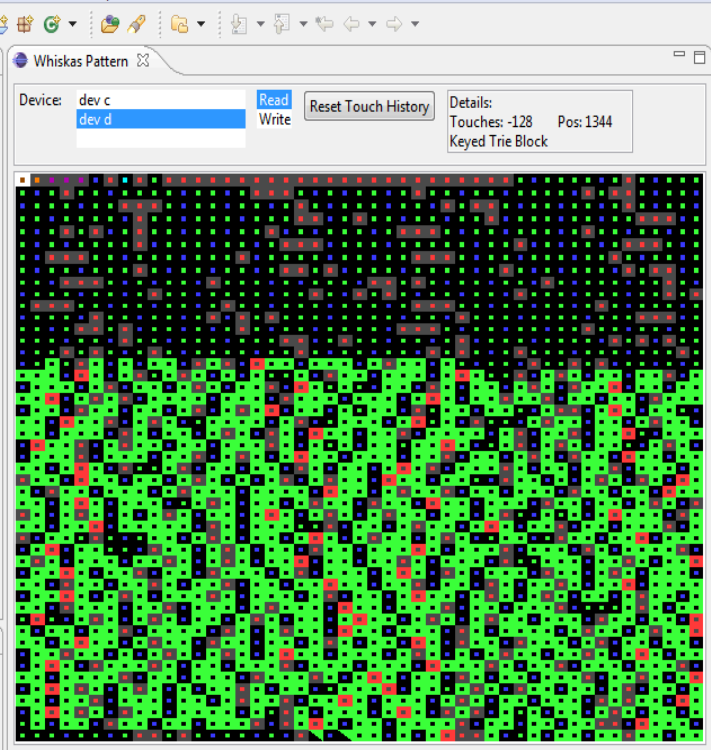
F:\Eigene Dateien\JAVA Entwicklung\U

Value added at hit:

Value subtracted by time:




Apply

Stop

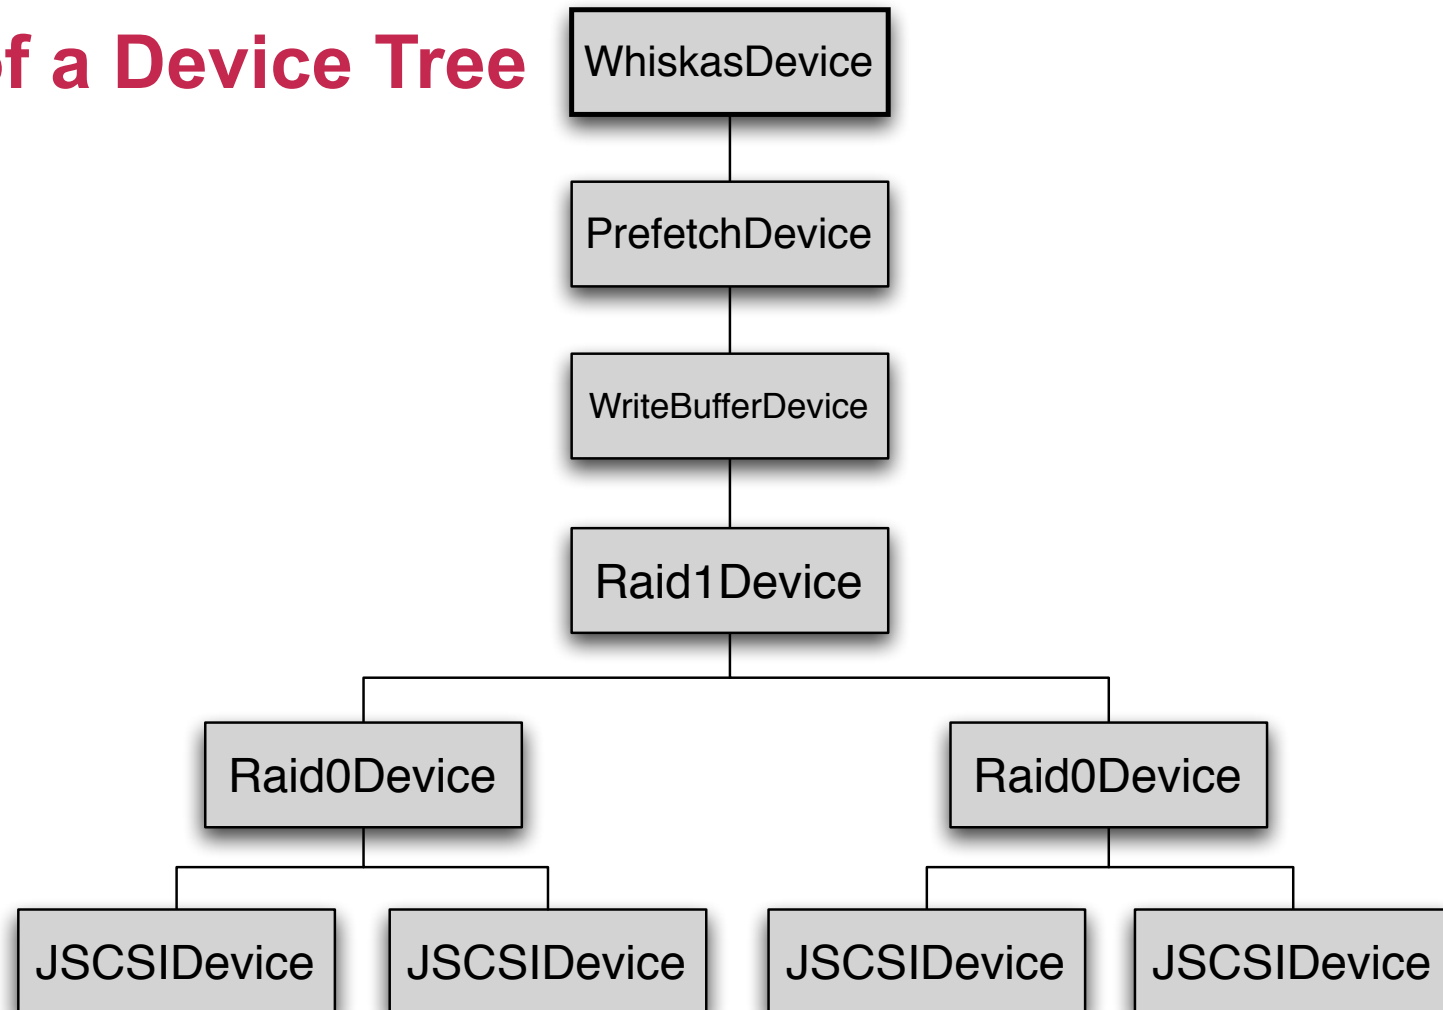


Whiskas Pattern

Device: dev c dev d




# Example of a Device Tree




# Lucene Backend

- > Apache Lucene
  - text search engine library
  - high-performance, full-featured
  
- > Device-based backend for Lucene
  
- > Full-text index on raw iSCSI target
  
- > jSCSI vs. filesystem backend
  - use
  - performance
  - complexity of setup and maintenance



# Benchmarks

Description	unit	min	max	avg	stddev	conf95	runs
<b>... Read 40kB</b>							
jSCSI	ms	2	14	3,34	1,32	[3,08;3,60]	100
Open-iSCSI	ms	4	15	5,29	1,32	[5,03;5,55]	100
<b>... Read 400kB</b>							
jSCSI	ms	22	24	22,48	0,52	[22,38;22,58]	100
Open-iSCSI	ms	5	16	6,79	3,19	[6,17;7,41]	100
<b>... jSCSI write 8MB</b>							
No RAID, 1 target	ms	3889	6987	4347,55	398,42	[4322,86;4372,24]	1000
RAID 0, 2 targets	ms	2346	4947	2803,46	330,12	[2783,00;2823,92]	1000
RAID 1, 2 targets	ms	4663	9447	5313,72	487,65	[5283,50;5343,95]	1000
<b>... Lucene "Getting started" [9]</b>							
Build file system index	ms	523	1270	676,40	153,80	[609,00;743,80]	20
Build jSCSI index	ms	6363	7889	6927,95	468,25	[6722,73;7133,17]	20
Search file system index	ms	2	68	7,05	14,09	[0,87;13,23]	20
Search jSCSI index	ms	5	10	9,45	1,36	[8,85;10,05]	20




# Summary

- > Device hierarchy with Stream flexibility
- > Multithreading
- > Avoids kernel loop
- > Extensible platform
- > Performance OK



# Additional features for jSCSI 2.0

- > Asynchronous I/O semantics for the Device interface
- > Improved prefetcher and write buffer
- > Support for multiple pending operations per connection
- > Smaller memory footprint and less CPU consumption
- > Resizable devices (ZFS storage pool)
- > Write balancing
  - space usage
  - activity
  - latency



Marcel Waldvogel

waldvogel@uni-konstanz.de

University of Konstanz

<http://www.inf.uni-konstanz.de/disy/>

**JAZOON07**

THE INTERNATIONAL CONFERENCE ON JAVA TECHNOLOGY  
JUNE 24 - 28, 2007 ZÜRICH

Universität Konstanz

