Universität Potsdam
Mathematisch-Naturwissenschaftliche Fakultät

Optimizing Energy Efficiency and Quality of Service in Large Scale Web Server Environments

2016-12-17
GreenCom2016 --- Chengdu

Simon Kiertscher, Bettina Schnor
University of Potsdam

Components of a Server-Load-Balancing (SLB) Cluster

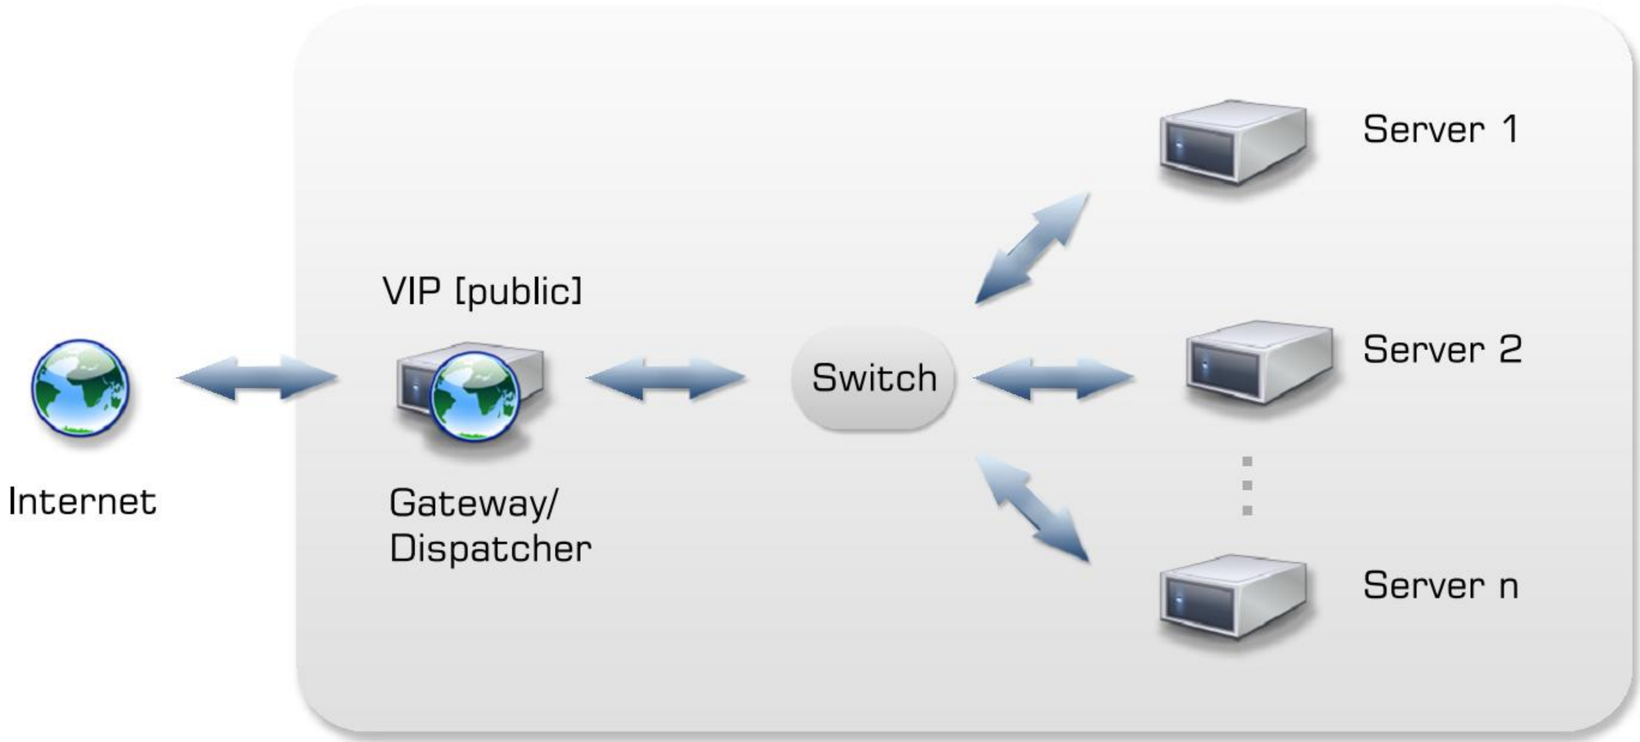

Mathematisch-
Naturwissenschaftliche Fakultät

Motivation

- 30% of servers world-wide are *comatose* according to [1] (2015, Stanford) and [2] (2008, Uptime Institute)
- Corresponds to 4GW
  The most power full nuclear power plant block on earth generates 1.5GW

Source:
https://de.wikipedia.org/wiki/Datei:Chooz_Nuclear_Power_Plant-9361.jpg

Energy Saving Daemon (CHERUB)

- Centralized approach - no clients on back-ends
- Daemon located at master node polls the system in fixed time intervals to analyze its state
  - Status of every node
  - Load situation
- Depending on the state of the nodes, saved attributes and the load prediction, actions are performed for every node
- Online system - we don't need any information about future load

Simulation - ClusterSim Architecture

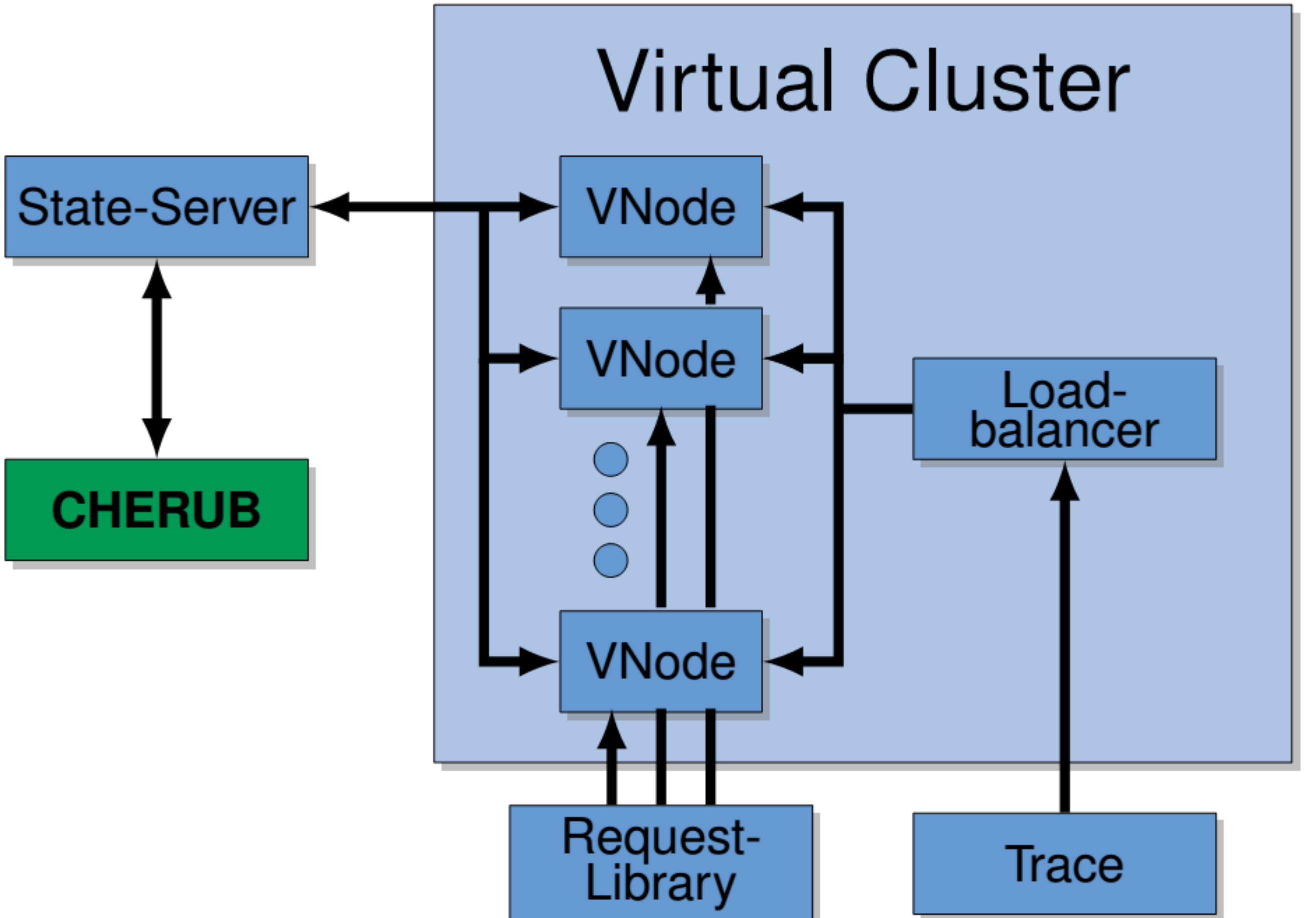

Scalability Evaluation of an Energy-Aware Resource Management System for Clusters of Web Servers
Kiertscher, Schnor
International Symposium on Performance Evaluation of Computer and Telecommunication Systems (SPECTS), Chicago, USA, July 2015

Simulation - Energy Accounting

- Using real data from SPECpower_ssj 2008 Benchmark (Systems from 2007-2015)
- No data about STR, Boot or Shutdown consumption

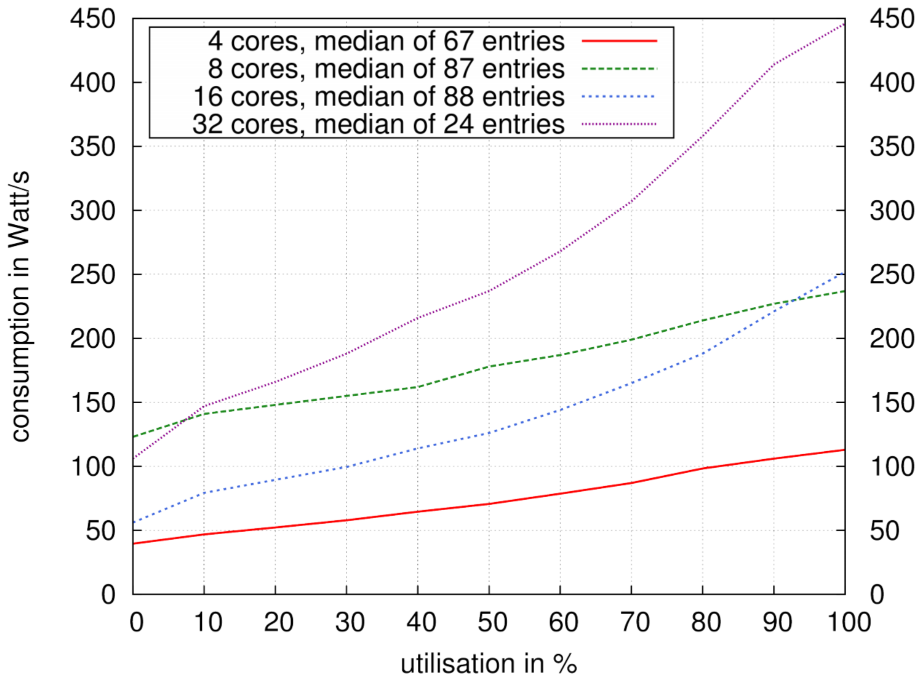

Used Metrics

- **Quality of Service (QoS) in %
  using a 5 second timeout**
- Request duration (RD) in milliseconds
  including waiting and processing time
- **Energy consumption (EC) in Wh / Energy saved (ES) in %**
- Number of physical state changes (PSCs) defined as the process to either turn on or turn off a node
- Score, a weighted ranking of the other 4 metrics

Varied Factors

Boot duration [s]	Backup [%]	Shutdown strategy	Explicit wait before boot [min]
5, 30, 60, 120, 180	0,5,10,25	Aggressive vs. one-by-one	0, 1

Used Workload and Optimum regarding EC/ES

- Peak load situation → Worst Case
- Derived from real trace

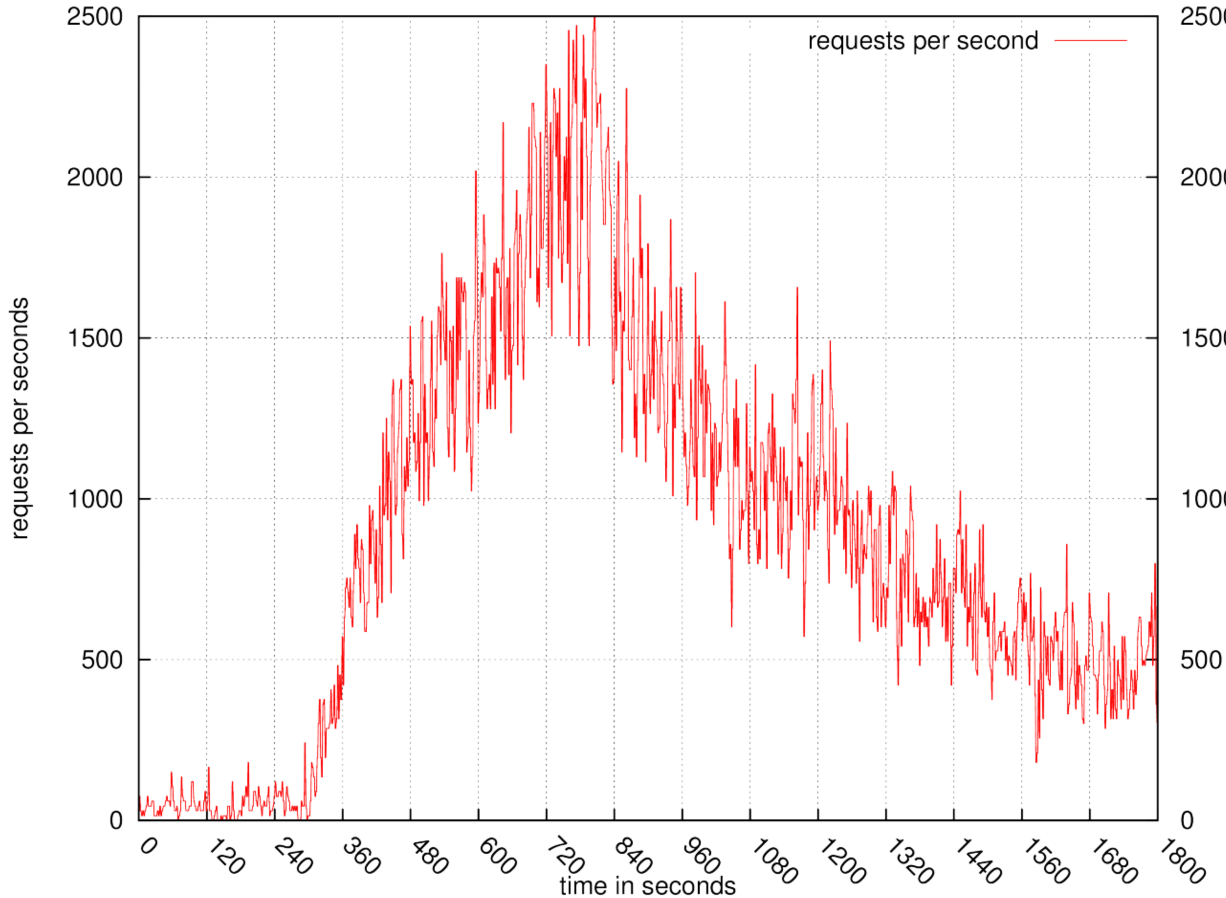

Used Workload and Optimum (60 seconds boottime)

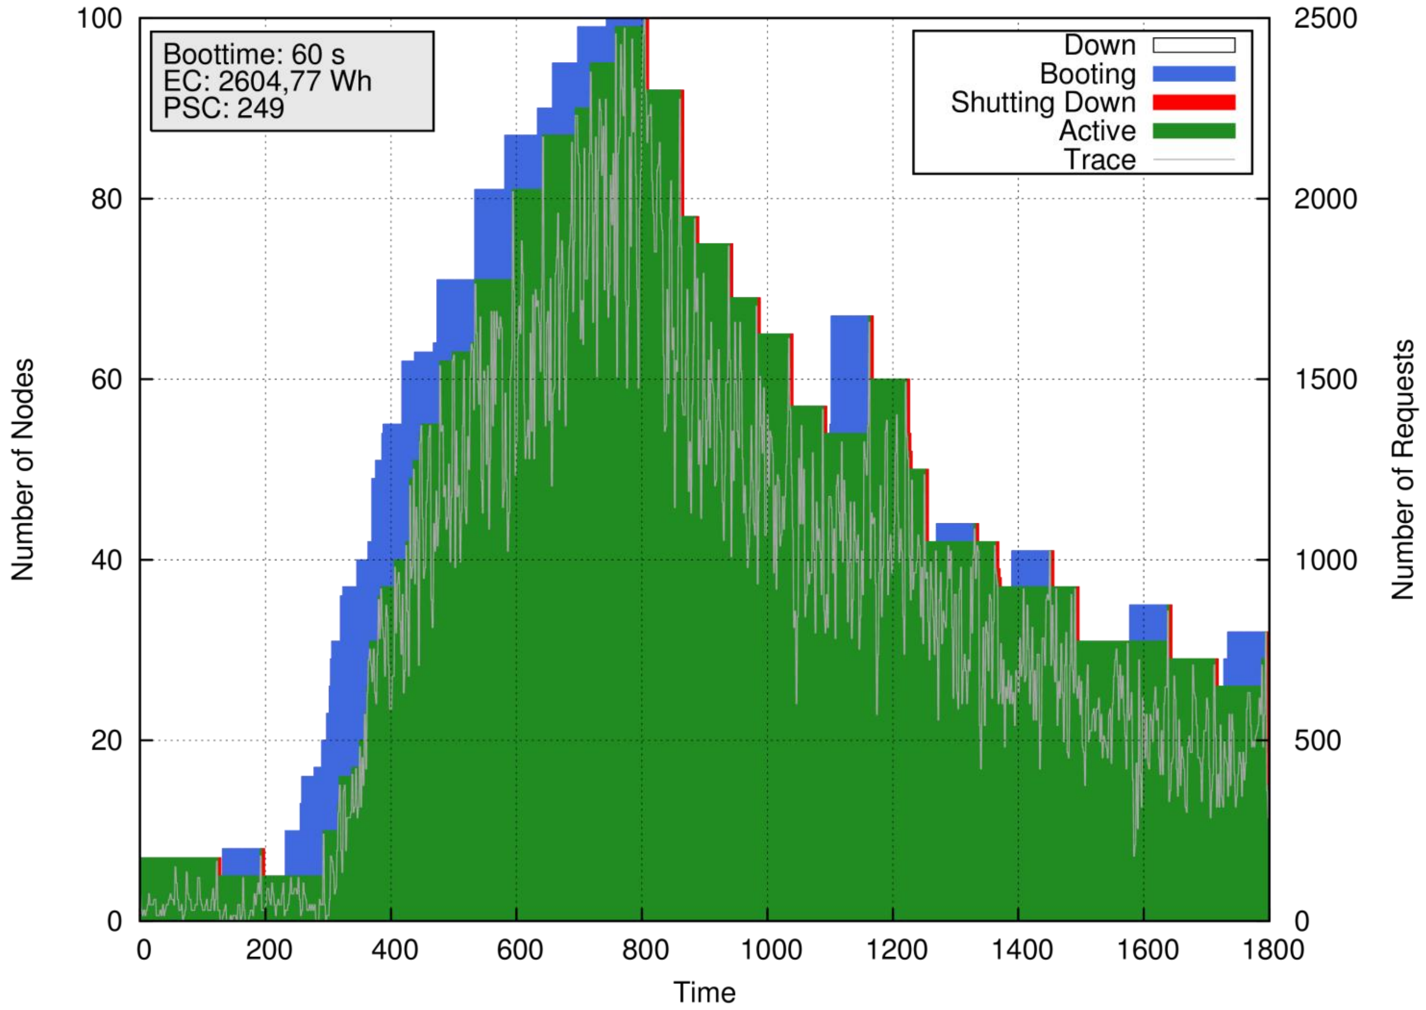

One of 80

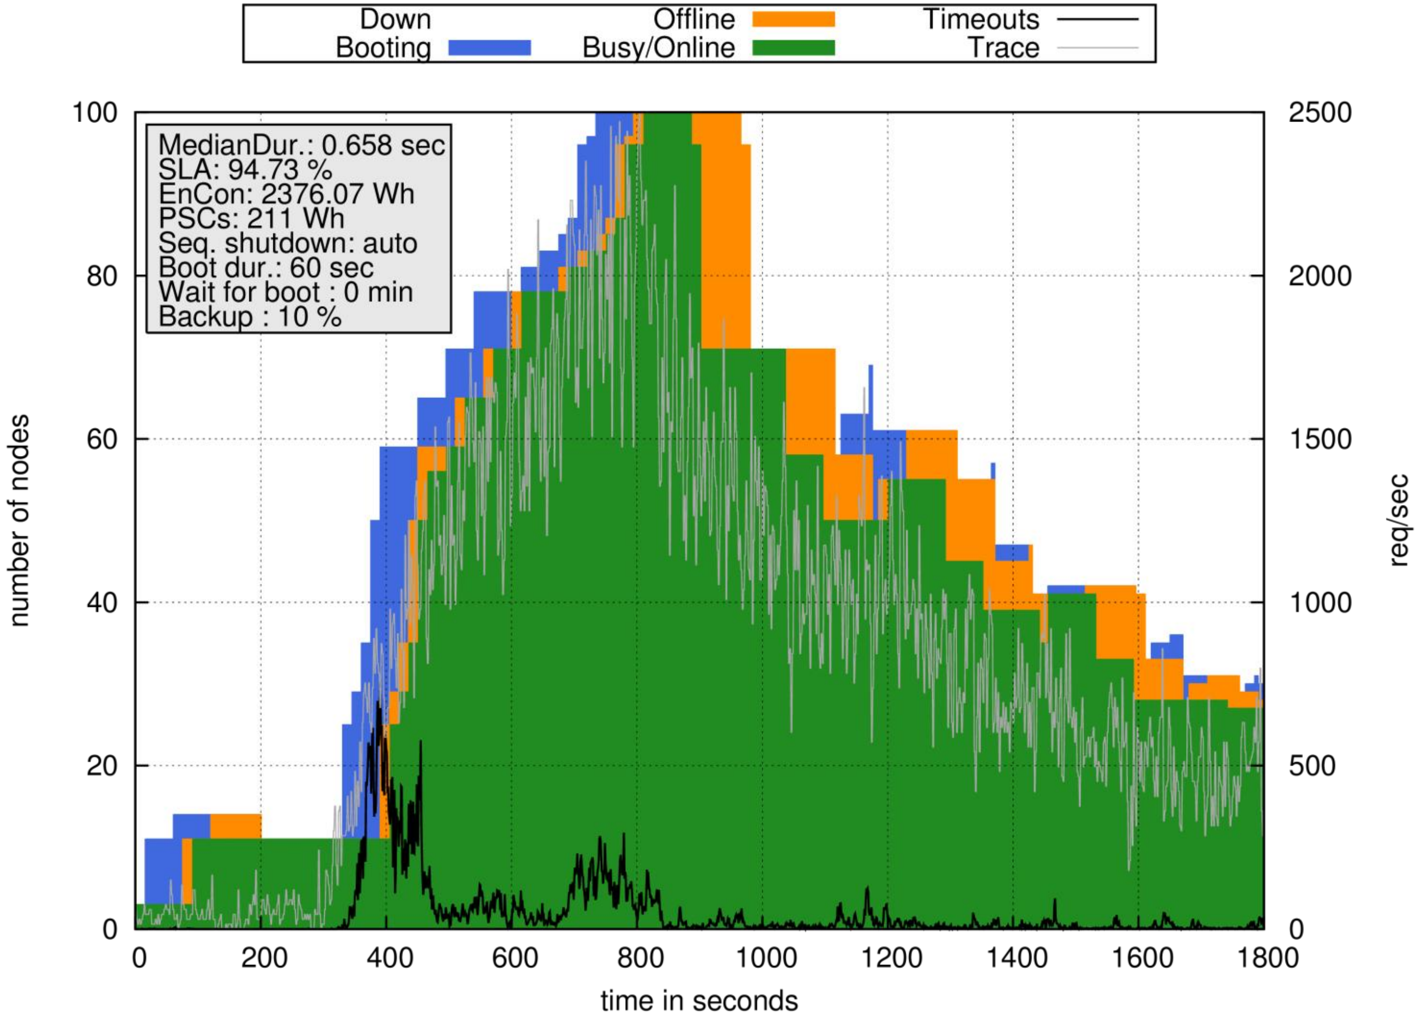

Down
Booting
Offline
Busy/Online
Timeouts
Trace
MedianDur.: 0.658 sec
SLA: 94.73 %
EnCon: 2376.07 Wh
PSCs: 211 Wh
Seq. shutdown: auto
Boot dur.: 60 sec
Wait for boot : 0 min
Backup : 10 %
number of nodes
req/sec
time in seconds
0
20
40
60
80
100
0
500
1000
1500
2000
2500
0
200
400
600
800
1000
1200
1400
1600
1800

Results QoS

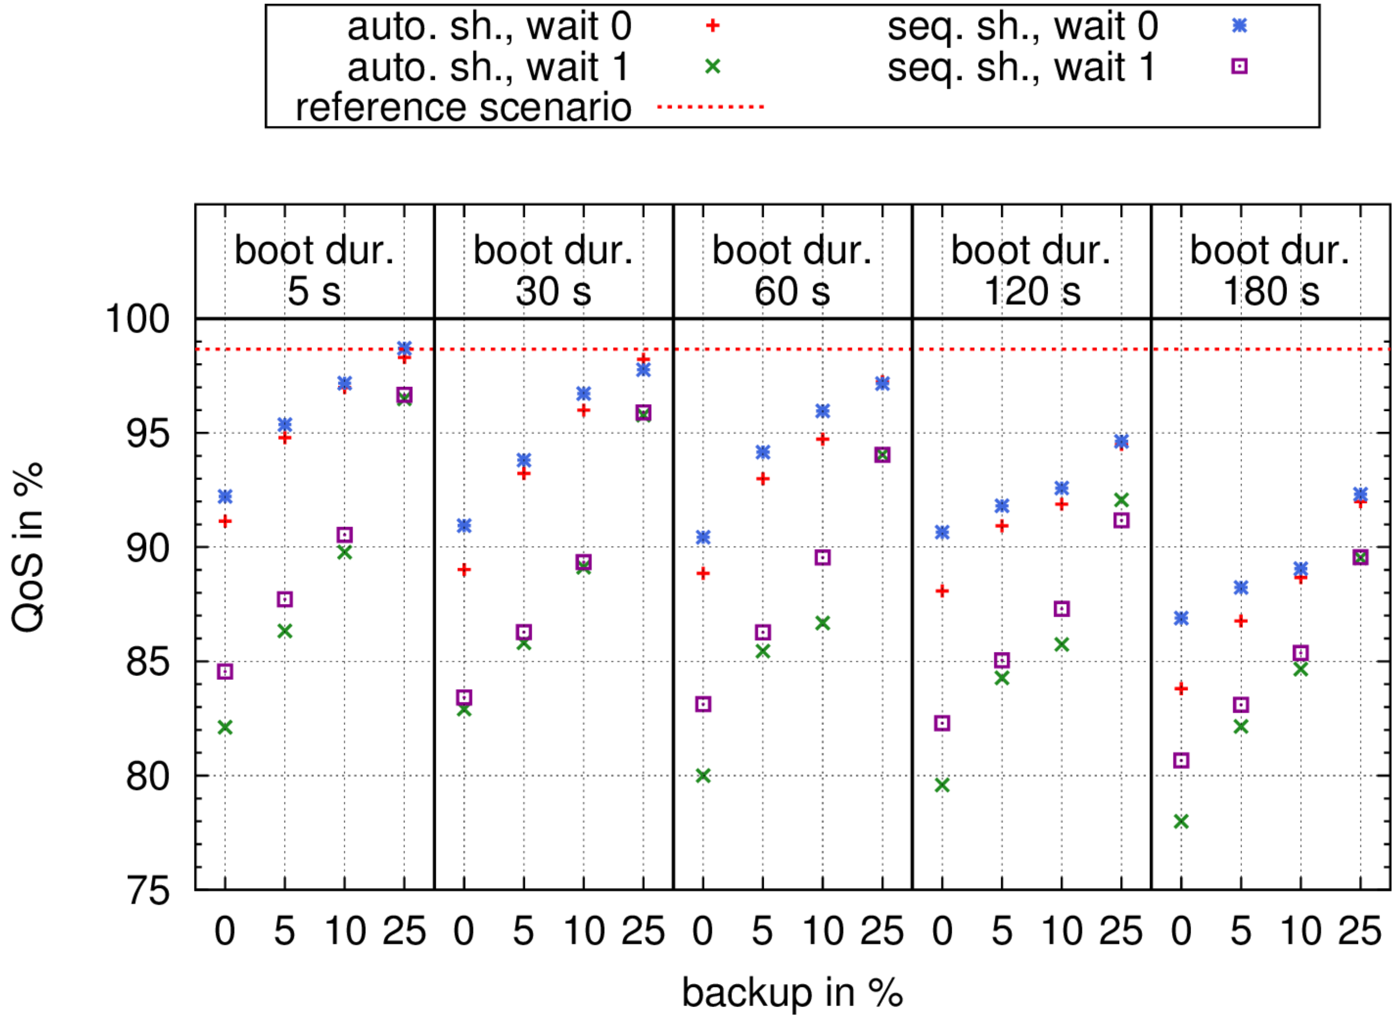

auto. sh., wait 0
auto. sh., wait 1
reference scenario
seq. sh., wait 0
seq. sh., wait 1
boot dur. 5 s
boot dur. 30 s
boot dur. 60 s
boot dur. 120 s
boot dur. 180 s
100
95
90
85
80
75
QoS in %
0 5 10 25 0 5 10 25 0 5 10 25 0 5 10 25 0 5 10 25
backup in %

Results Energy Saving

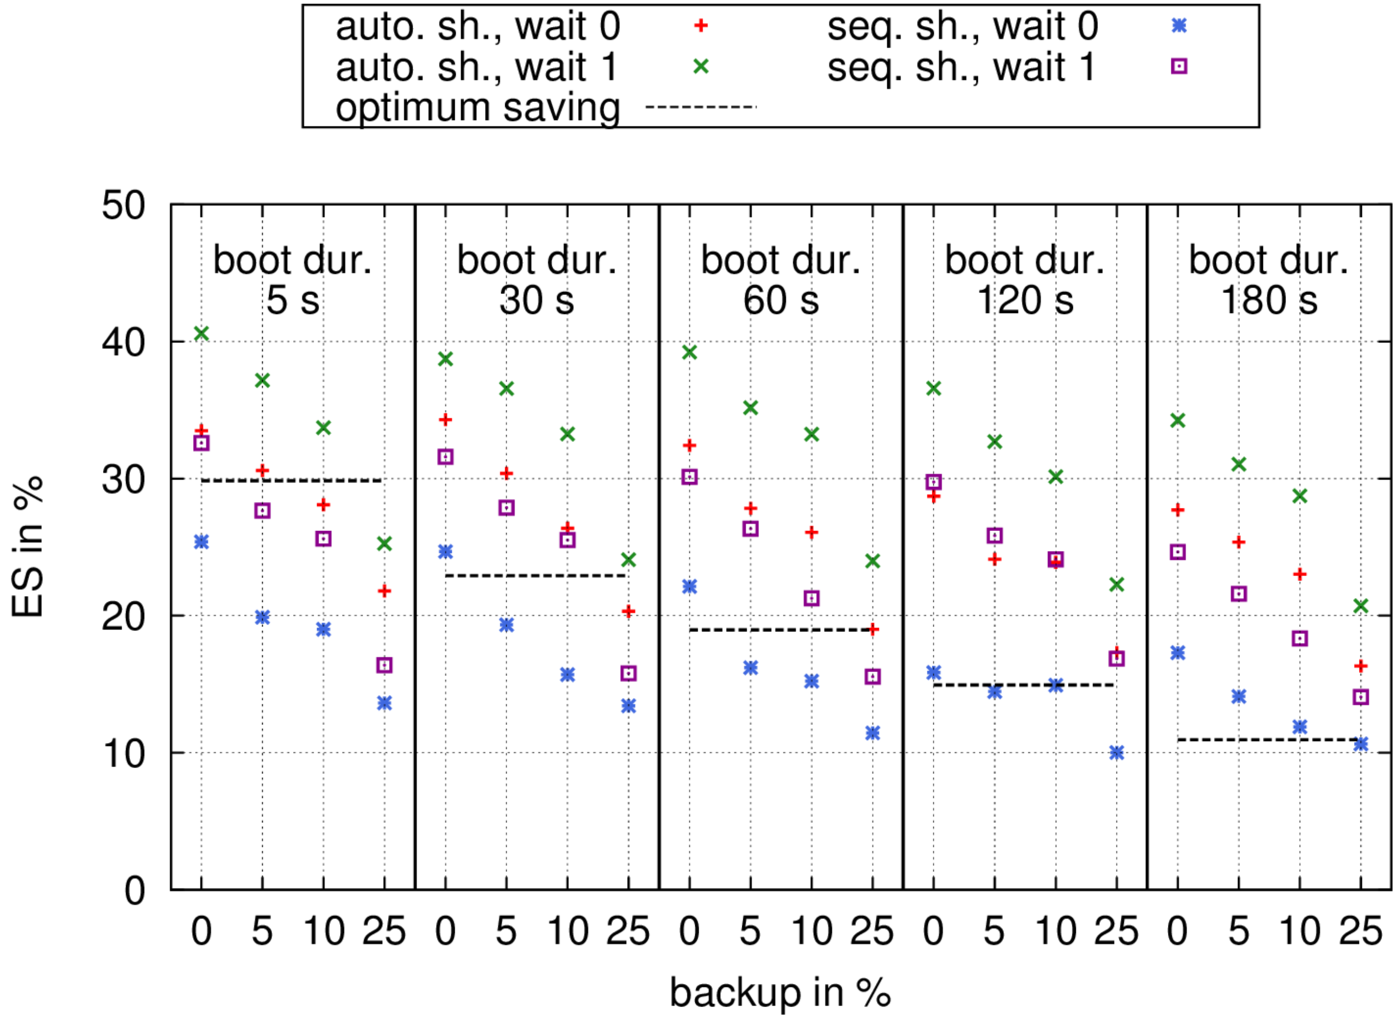

Universität Potsdam
Mathematisch-Naturwissenschaftliche Fakultät

Score - Results

strategy	boot duration in s	score	QoS in %	ES in %	Optimum Saving	RD in ms	PSCs	PSCs in Optimum Case	shutdown	wait in min	backup in %
HP	5	70.3	98.70	13.63	29.9	67	123	1094	seq.	0	25
	30	66.5	97.77	13.42	22.9	79	135	404	seq.	0	25
	60	63.6	97.16	11.44	19.0	78	159	249	seq.	0	25
	120	56.9	94.63	10.02	14.9	81	156	215	seq.	0	25
	180	53.1	92.32	10.64	11.0	80	148	181	seq.	0	25
LC	5	61.0	84.56	32.60	29.9	548	94	1094	seq.	1	0
	30	60.5	83.42	31.58	22.9	430	89	404	seq.	1	0
	60	56.3	83.13	30.12	19.0	401	105	249	seq.	1	0
	120	49.2	82.30	29.76	14.9	527	124	215	seq.	1	0
	180	43.1	78.00	34.25	11.0	1452	180	181	auto.	1	0
B	5	55.4	96.99	28.08	29.9	574	199	1094	auto.	0	10
	30	50.8	96.00	26.37	22.9	619	219	404	auto.	0	10
	60	47.8	94.73	26.08	19.0	658	211	249	auto.	0	10
	120	41.1	92.07	22.29	14.9	211	174	215	auto.	1	25
	180	33.9	84.66	28.74	11.0	626	173	181	auto.	1	10

Conclusion

- Strategy works in 100 node SLB-Setup
- Results are very close to the optimum (with fast hardware)
- Boot duration is a critical factor
- Backup has a linear influence on QoS and EC
- Aggressive shutdown can save up to 12,9 % extra energy in the peak load scenario
- Extra waiting time is not necessary if load forecasting is used

Thank you for your attention!
Any Questions?

Contact:

kiertscher@cs.uni-potsdam.de

www.cs.uni-potsdam.de

Sources

[1] “New data supports finding that 30 percent of servers are ‘Comatose’, indicating that nearly a third of capital in enterprise data centers is wasted” by Jonathan Koomey and Jon Taylor, 2015

[2] “Revolutionizing Data Center Energy Efficiency” by James Kaplan, William Forrest, Noah Kindler, 2008