

Statistical Characterization of Business-Critical Workloads Hosted in Cloud Datacenters


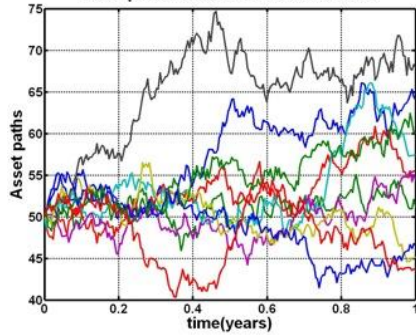
Vincent van Beek




S. Shen, V.S. van Beek, and A. Iosup , *Statistical Characterization of Business-Critical Workloads Hosted In Cloud Datacenters*, IEEE/ACM COGRID 2015, Shenzhen, Guangdong, China, May 4-7, 2015 (acceptance ratio 25%).



Example of Random Walk Asset Paths





Enterprise Public Cloud Services Spending in the Netherlands by Type, 2010-2016, €M



Source: <http://www.themetisfiles.com>

Business Critical Workloads



ASP4all Bitbrains

~200
employees

Customers

- ✓ ING
- ✓ NN
- ✓ Aegon
- ✓ ICS Cards
- ✓ Ahold
- ✓ TNO
- ✓ Trans Link Systems
- ✓ DBC Onderhoud



ASP4all

delivers custom services on infrastructure management (IaaS), 7x24 full service application hosting (PaaS) and online office services

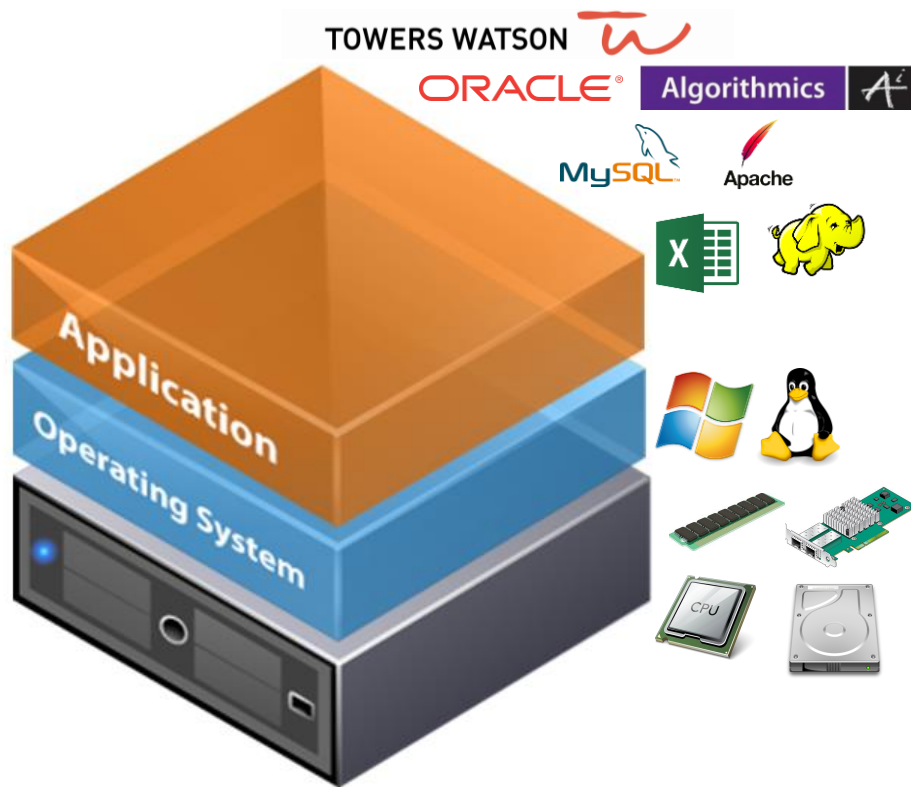
2013

**Gartner
Cool Vendor**
Bitbrains


Computable Award for
Partner of the year
ASP4all

Bitbrains

designs, builds and supports leading edge Cloud Computing solutions and is specialized in High Performance Computing



Traditional Architecture




Virtual Architecture



Required resources

Datacenter Manager

Ask about my other research




Used resources

From Workload trace to workload Characterization

Collected Two Unique Workload Traces

Name of the trace	# VMs	Period of data collection	Storage technology	Total memory	Total cores
fastStorage	1,250	1 month	SAN	17,729 GB	4,057
Rnd	500	3 months	NAS and SAN	5,485 GB	1,444
Total	1,750	5,446,811 CPU hours		23,214 GB	5,501

- All resources:
 - CPU, Memory, Storage, and Network
- Large scale
- Long term



From Workload trace to workload Characterization


Collected Two Unique Workload Traces

Prior work:

- Google
- Facebook
- Taobao
- Scientific workloads
- Grids vs Google

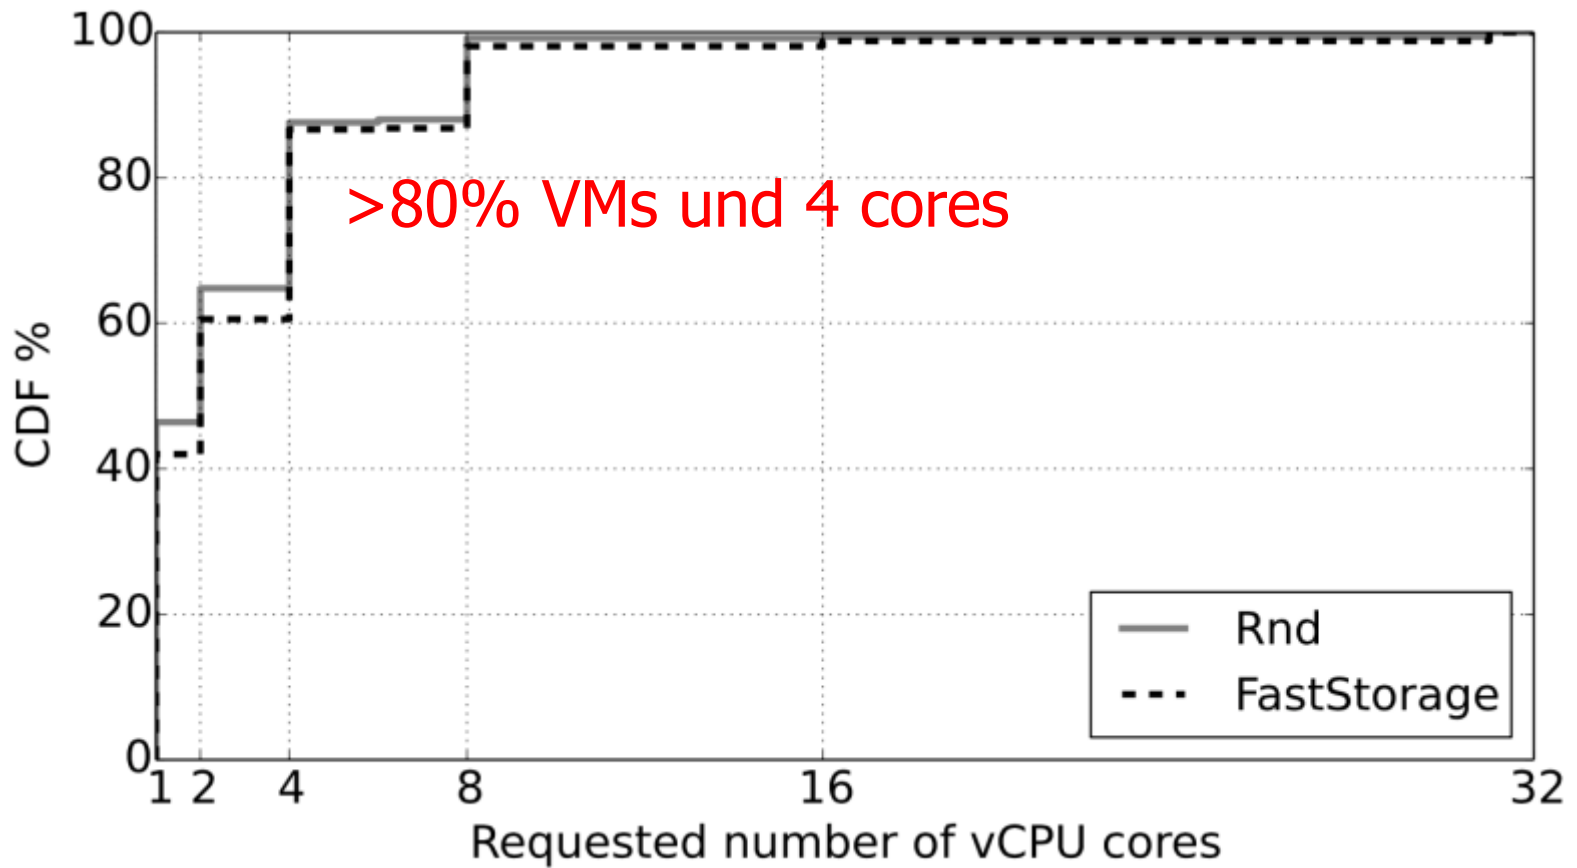
First study of both:

- Requested and
- Used resources
- For all resources



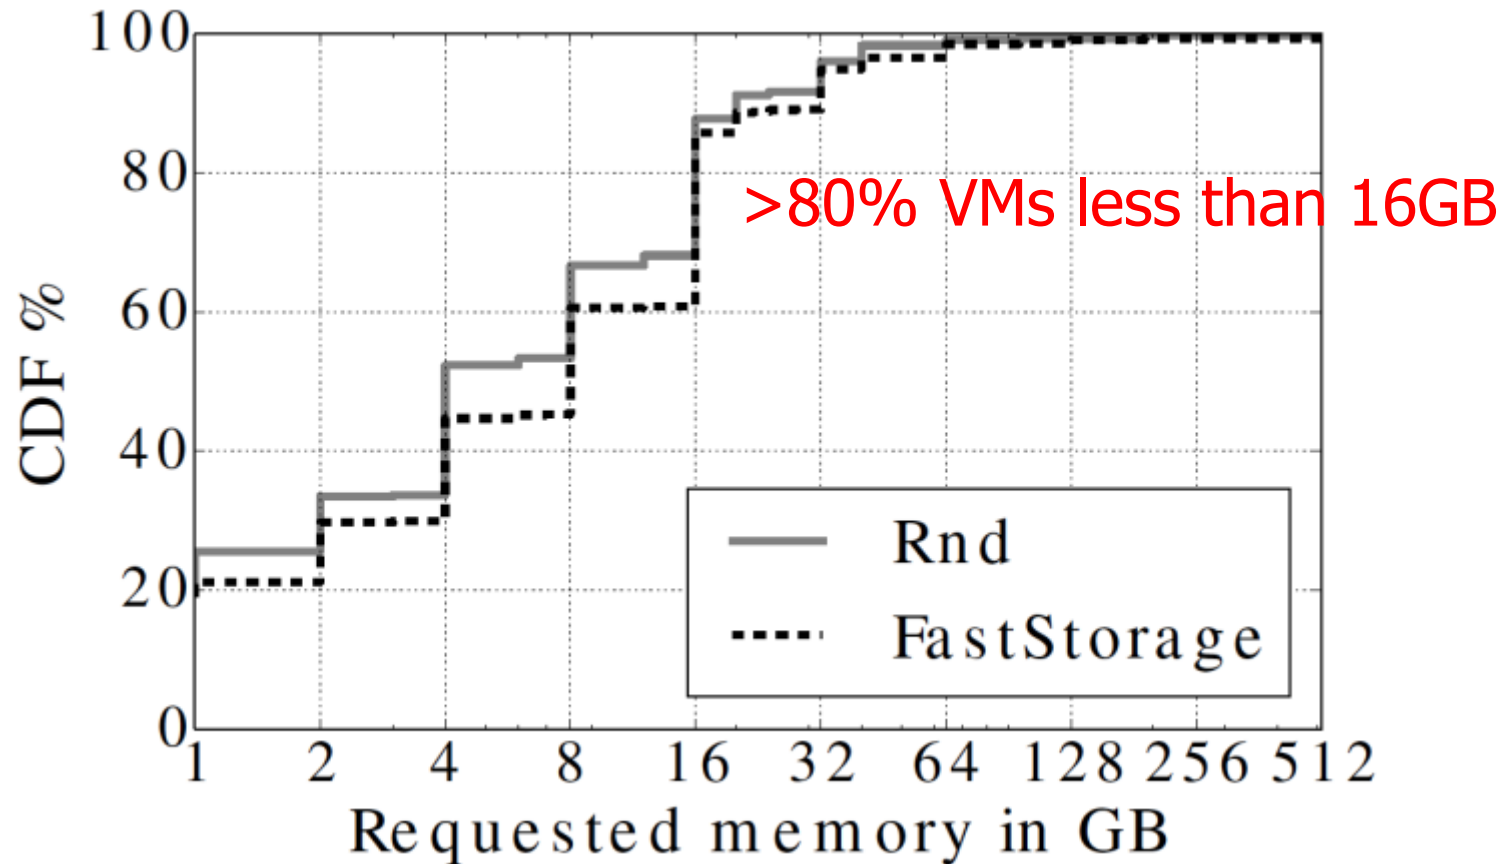
From Workload trace to workload Characterization

Requested number of CPU cores is low



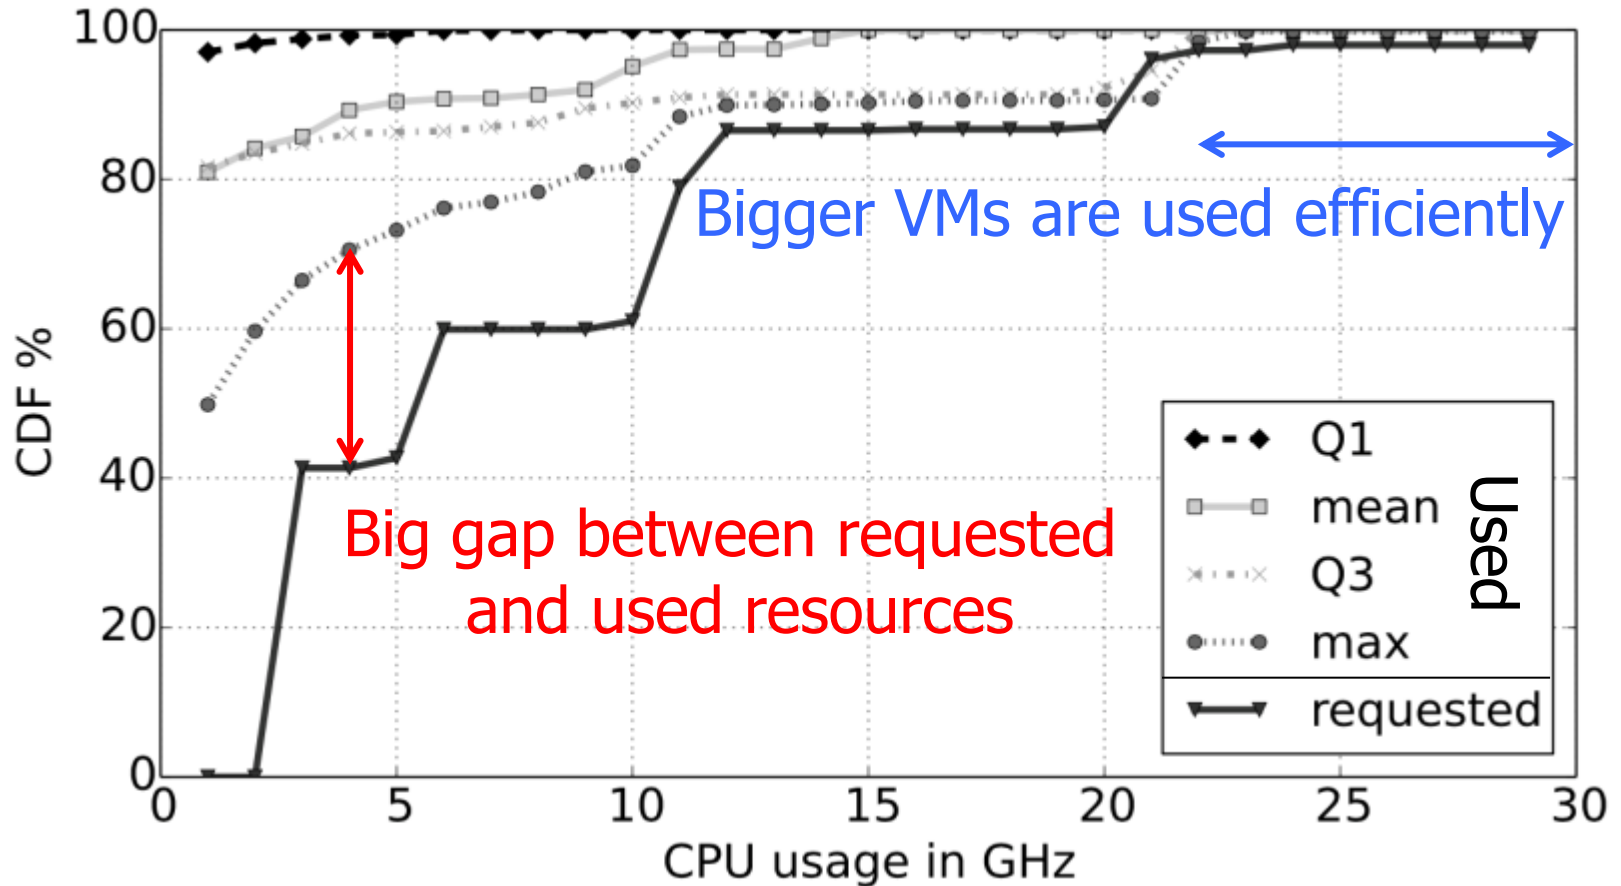
From Workload trace to workload Characterization

Requested memory is low, and power of two scaling



From Workload trace to workload Characterization


Requested vs Used CPU resources Leave big gap for optimization




From Workload trace to workload Characterization

Strong Auto-Correlation only for Storage

CPU workload




Storage Read workload



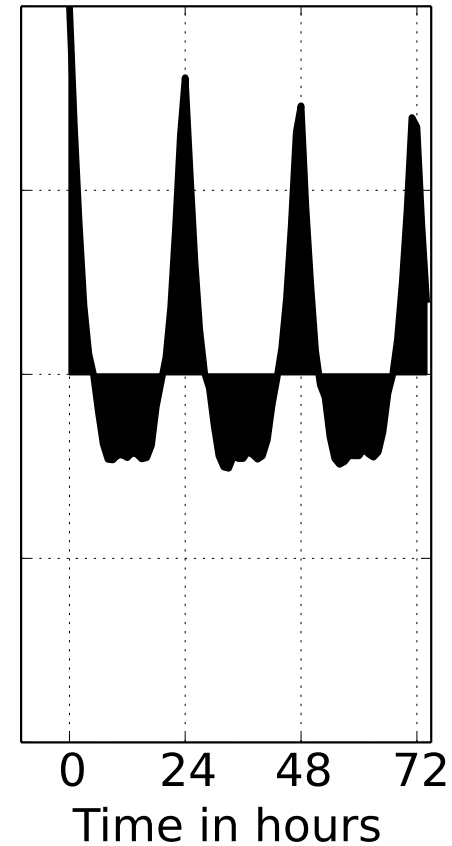
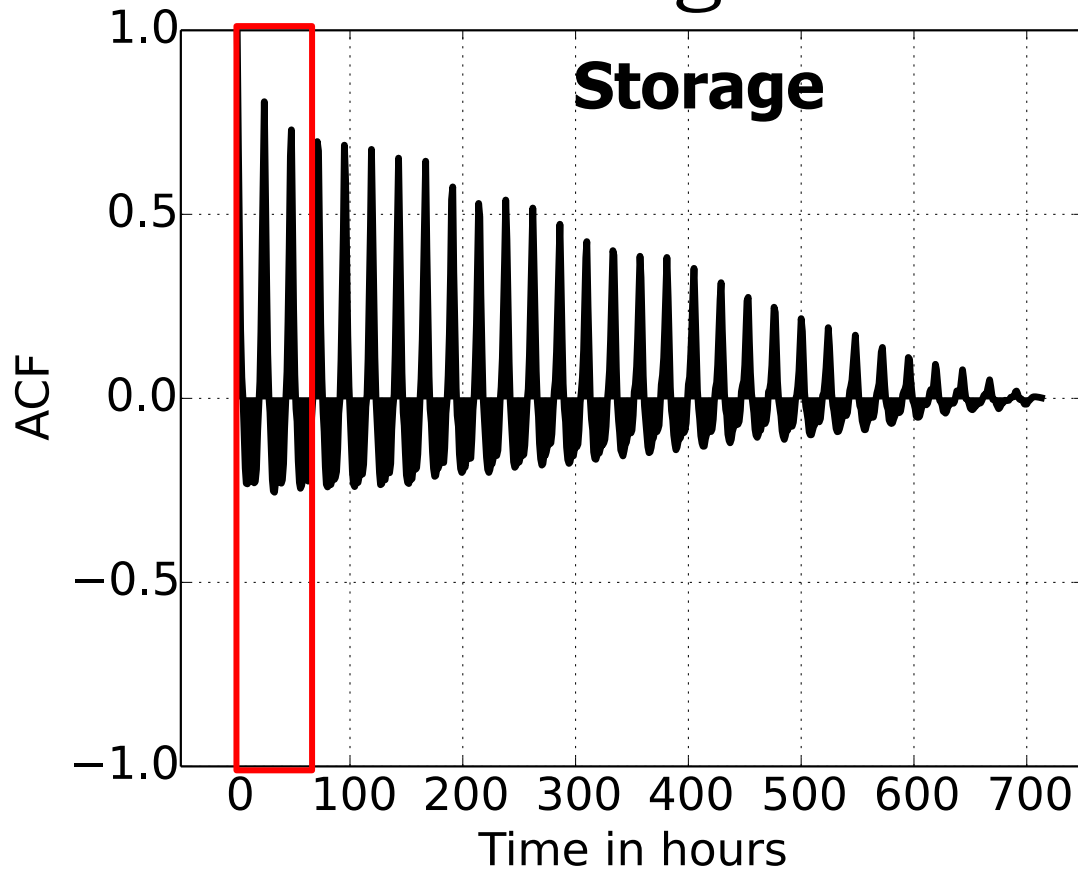
From Workload trace to workload Characterization

Strong Auto-Correlation only for Storage



From Workload trace to workload Characterization


Strong Auto-Correlation only for Storage



From Workload trace to workload Characterization


Strong Auto-Correlation only for Storage

CPU workload



Almost no correlation

Storage Read workload



Strong daily correlation

From Workload trace to workload Characterization

Business-Critical workloads

- Different from other workloads
 - Big gap between requested and used resources
 - Different behavior between resource types
 - Real world datacenter optimization at Bitbrains
-
- ✓ How are we helping the community
 - ✓ Method for characterizing workloads
 - ✓ Publicly available workload traces (<http://gwa.ewi.tudelft.nl>)
 - ✓ We are planning on making our analysis tools open source
 - ✓ You can HELP!
 - ✓ Let us analyze your workload traces

Future Work

- Acquired two grants: COMMIT and NWO/STW KIEM project
- Classification of workloads
- Development of scheduling mechanisms
- Datacenter load balancing
- Rebalancing workloads