



DPM - efficient storage in diverse environments

M. Hellmich, O. Keeble, F. Furano, A. Manzi,
A. Ayllon, I. Calvet, R. Rocha, M. Regala

CHEP 2013
Amsterdam



Experiment with storage

- Scaling the Architecture
 - Raspberry Pi
 - Virtual Machines
 - Bare Metal Servers
- Evaluate the outcomes
 - Performance per Watt, Dollar

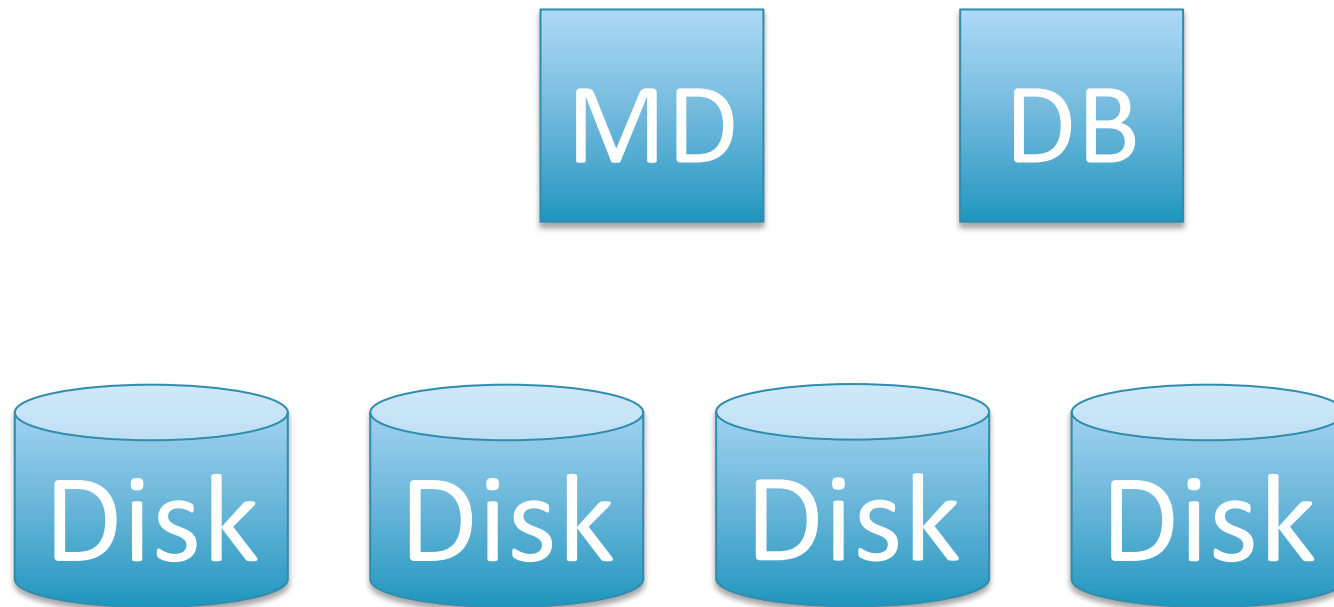
Storage



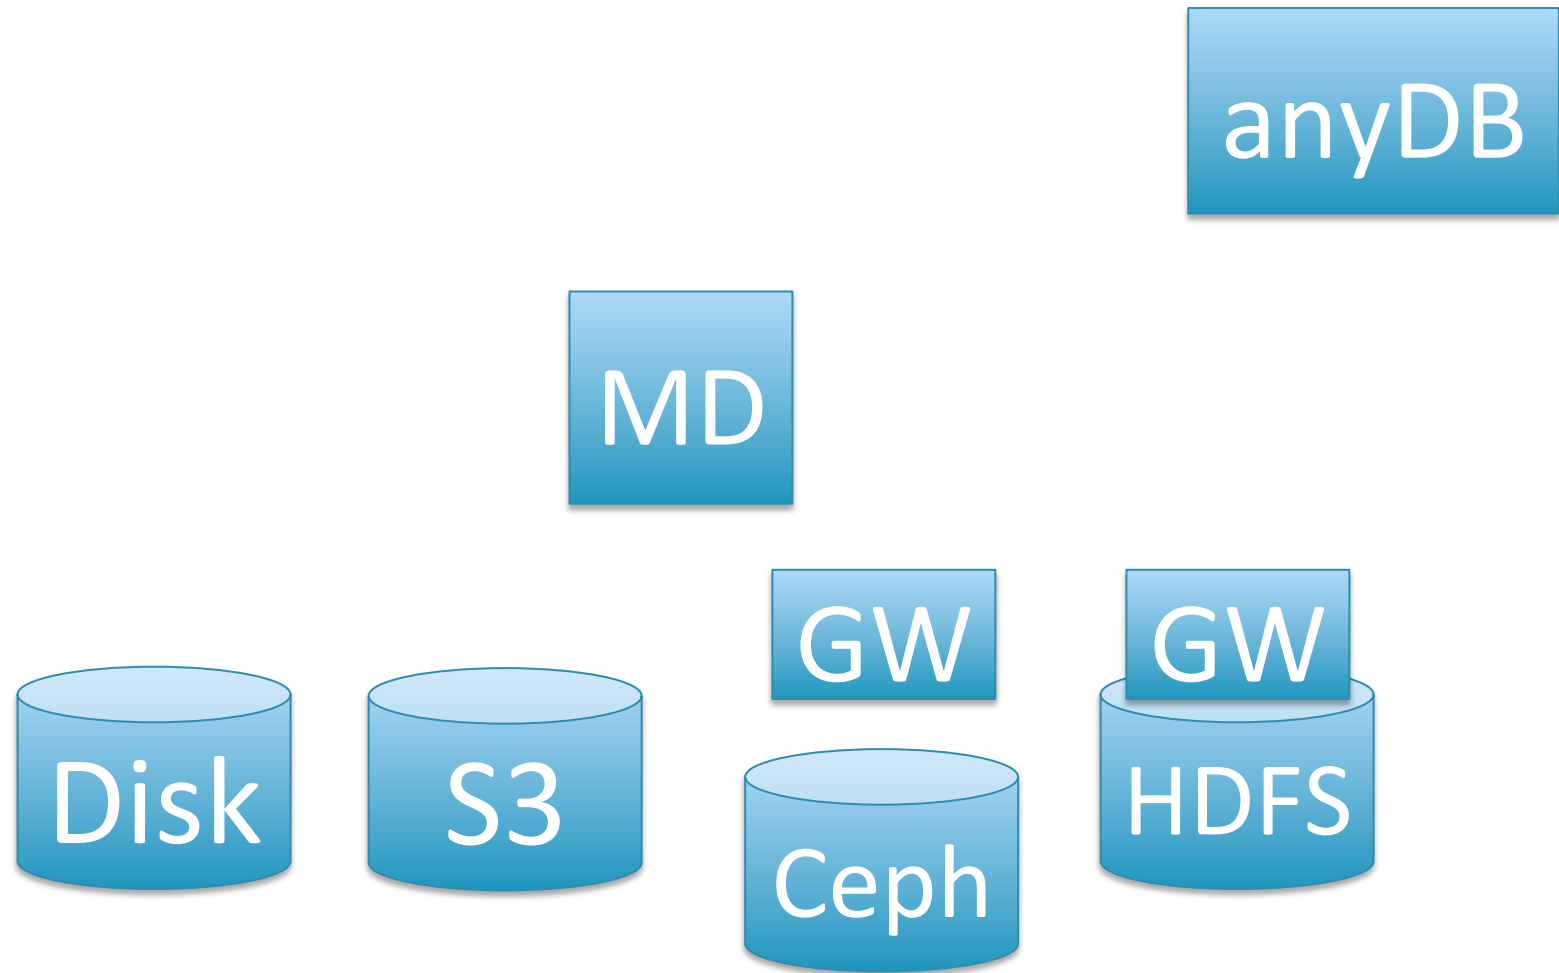
This talk is a reaction

- Change in our software
- Change in computing services
- Change in available hardware

Our software: DPM



Our software: New DPM



Computing Services

Metadata Reqs

- network
- CPU
- memory





Gateway Reqs

- >> network
- CPU
- memory

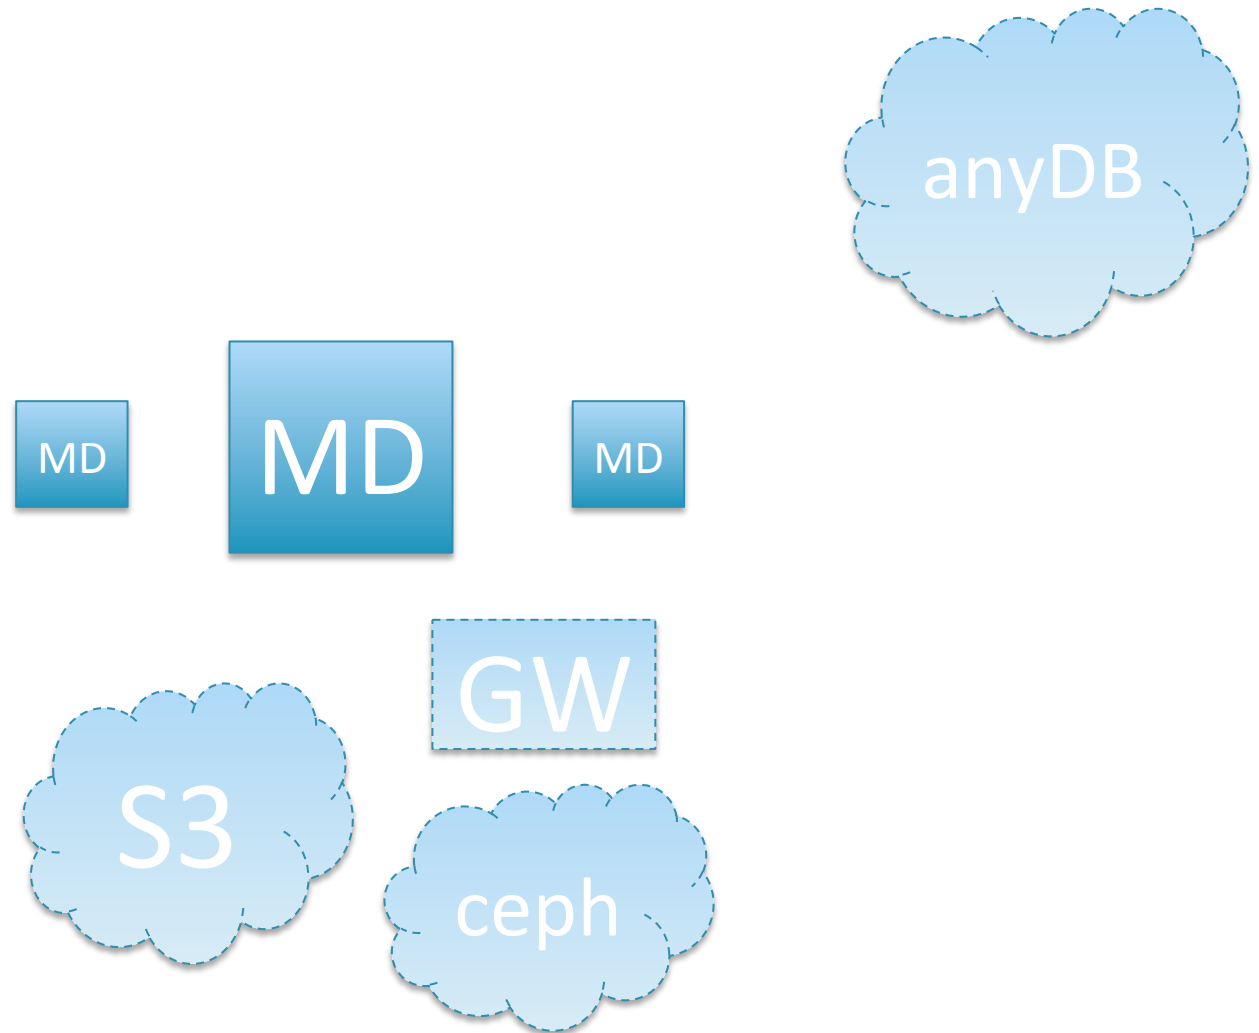


Available Hardware: ARM



- Low power consumption
 - also when idle

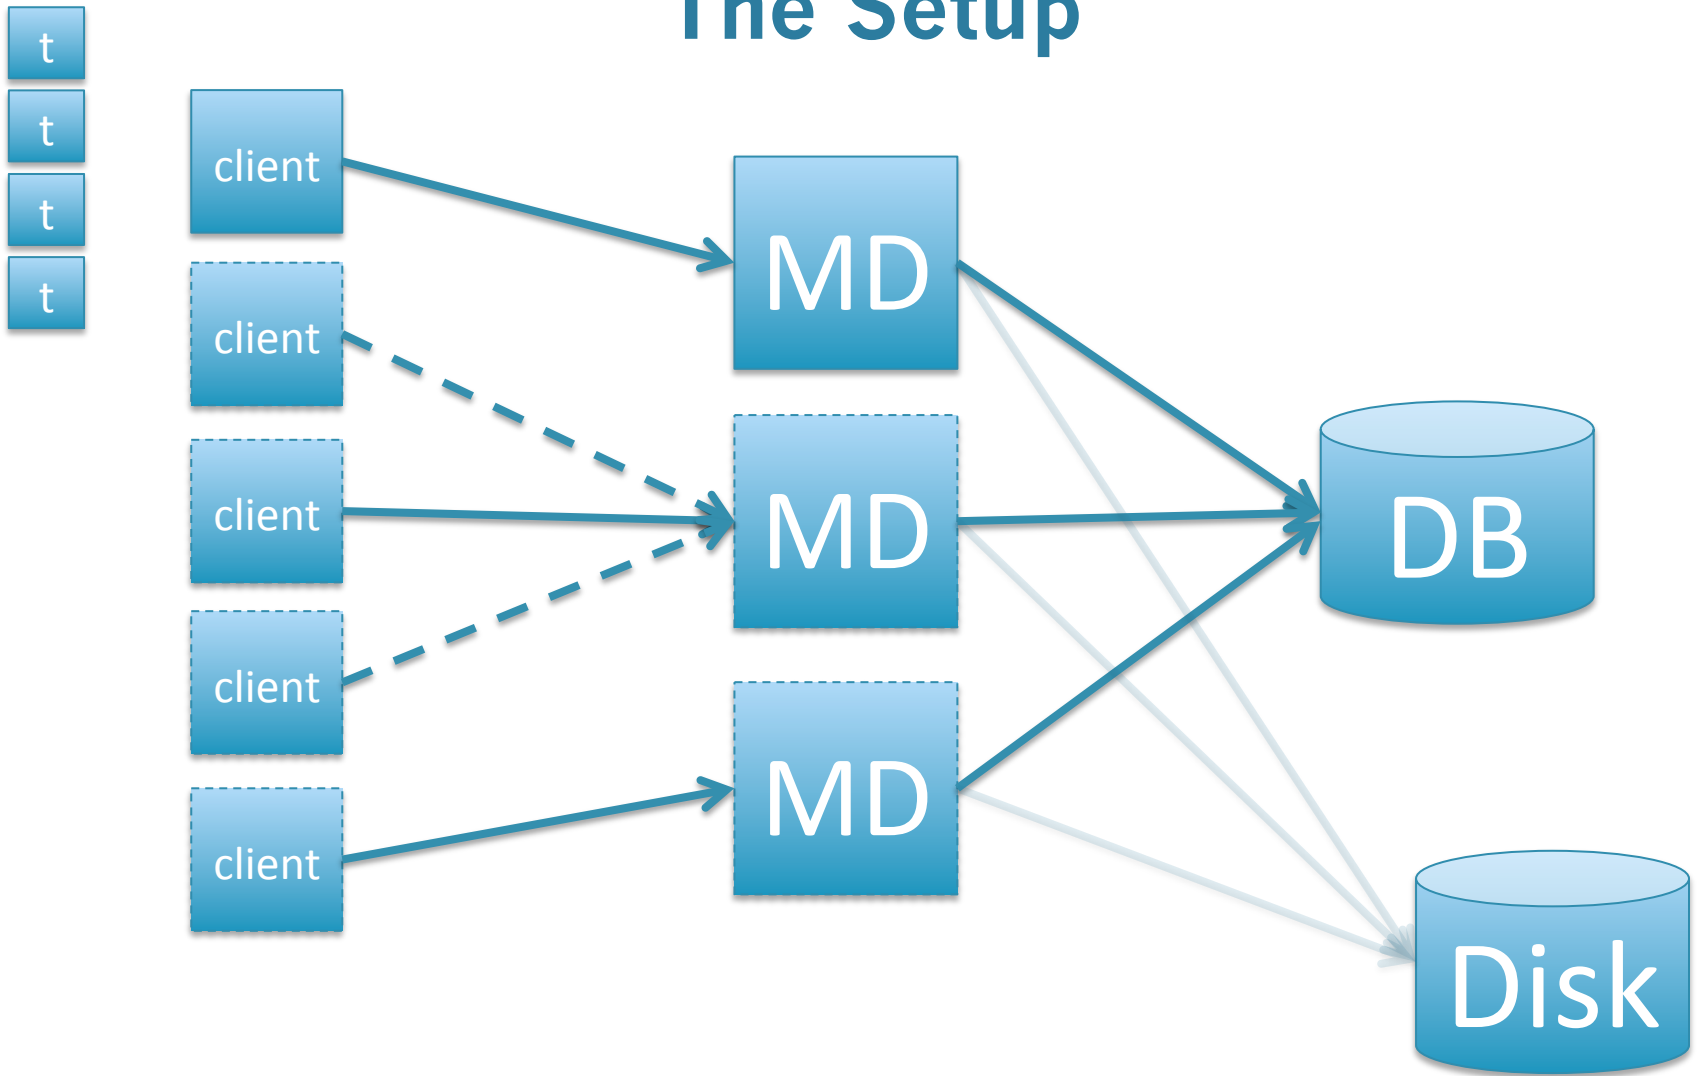
So we focused on Metadata



The Metric

- requests/sec served, pure and simple
 - Stat requests for files
 - Deputy for other calls
- Tested over http

The Setup



Performance Comparison

	Raspi	CERN Openstack VM	EC2 medium VM	Bare metal 4 core
Requests/sec	25.8	394.18	162.27	1701.23

	3 x Raspi	3 x CERN Openstack VM
Requests/sec	71.30	1087.0
Scaling	2.76	2.75

Scaling up with Memcached

	Raspi	3 x Raspi
Requests/sec direct	25.8	71.30
Requests/sec with Memcached	58.47	152.16

	VM*	3 x VM*
Requests/sec direct	394.18	1087.0
Requests/sec with Memcached	764.02	

*VMs are OpenStack VMs

Total Cost of Ownership

TCO = Cost Machines +
Cost Running Machines

Cost Running Machines =
Cost Maintenance + Cost Replacement +
Cost Power + Cost External Components

Cost SoftwareDev

Performance per Watt

	Raspi	CERN OpenStack VM
Requests/sec	25.8	394.18
Consumption [W]	3.5	30
Requests/sec/watt	7.37	13.1

Performance per Euro

	Raspi	VM	EC2 Medium
Requests/sec	25.8	394.18	162.27
Buying Price	25	428²	-
Power Cost over 3 years ¹²	4.89 Euros	41.94	-
Cost over 3 years ¹	29.89 Euros (+100)	469.94 (+100)	1297.2
500 Request/sec over 3 years	750	940	3891


Note Raspi external equipment is much more expensive!

¹ Buying Price + Power Costs or EC2 Pricing (reserved instance)

² 12VMs running on DELL PowerEdge M620 24 core 16GB RAM, price from website


Working with a Raspberry

- Smallest possible machine
- But very good for testing
- Runs EPEL software*



Summary

- Evaluation of a DPM metadata server
 - VMs do well
 - Especially with memcached
 - Raspberries work
 - Are even competitive
 - ARM may become very useful for storage



Questions?

THANK YOU VERY MUCH