METHOD R™

Optimizing

Cary Millsap
@CaryMillsap / cary.millsap@method-r.com

NoCOUG Winter Conference
Pleasanton, California
9:30a–10:30a Thursday 24 February 2011

Cary Millsap

President, Method R Corporation
http://method-r.com

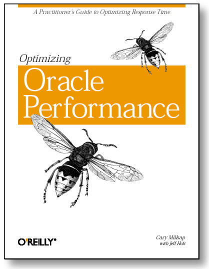

METHOD R™

Part 1

Oracle Exadata

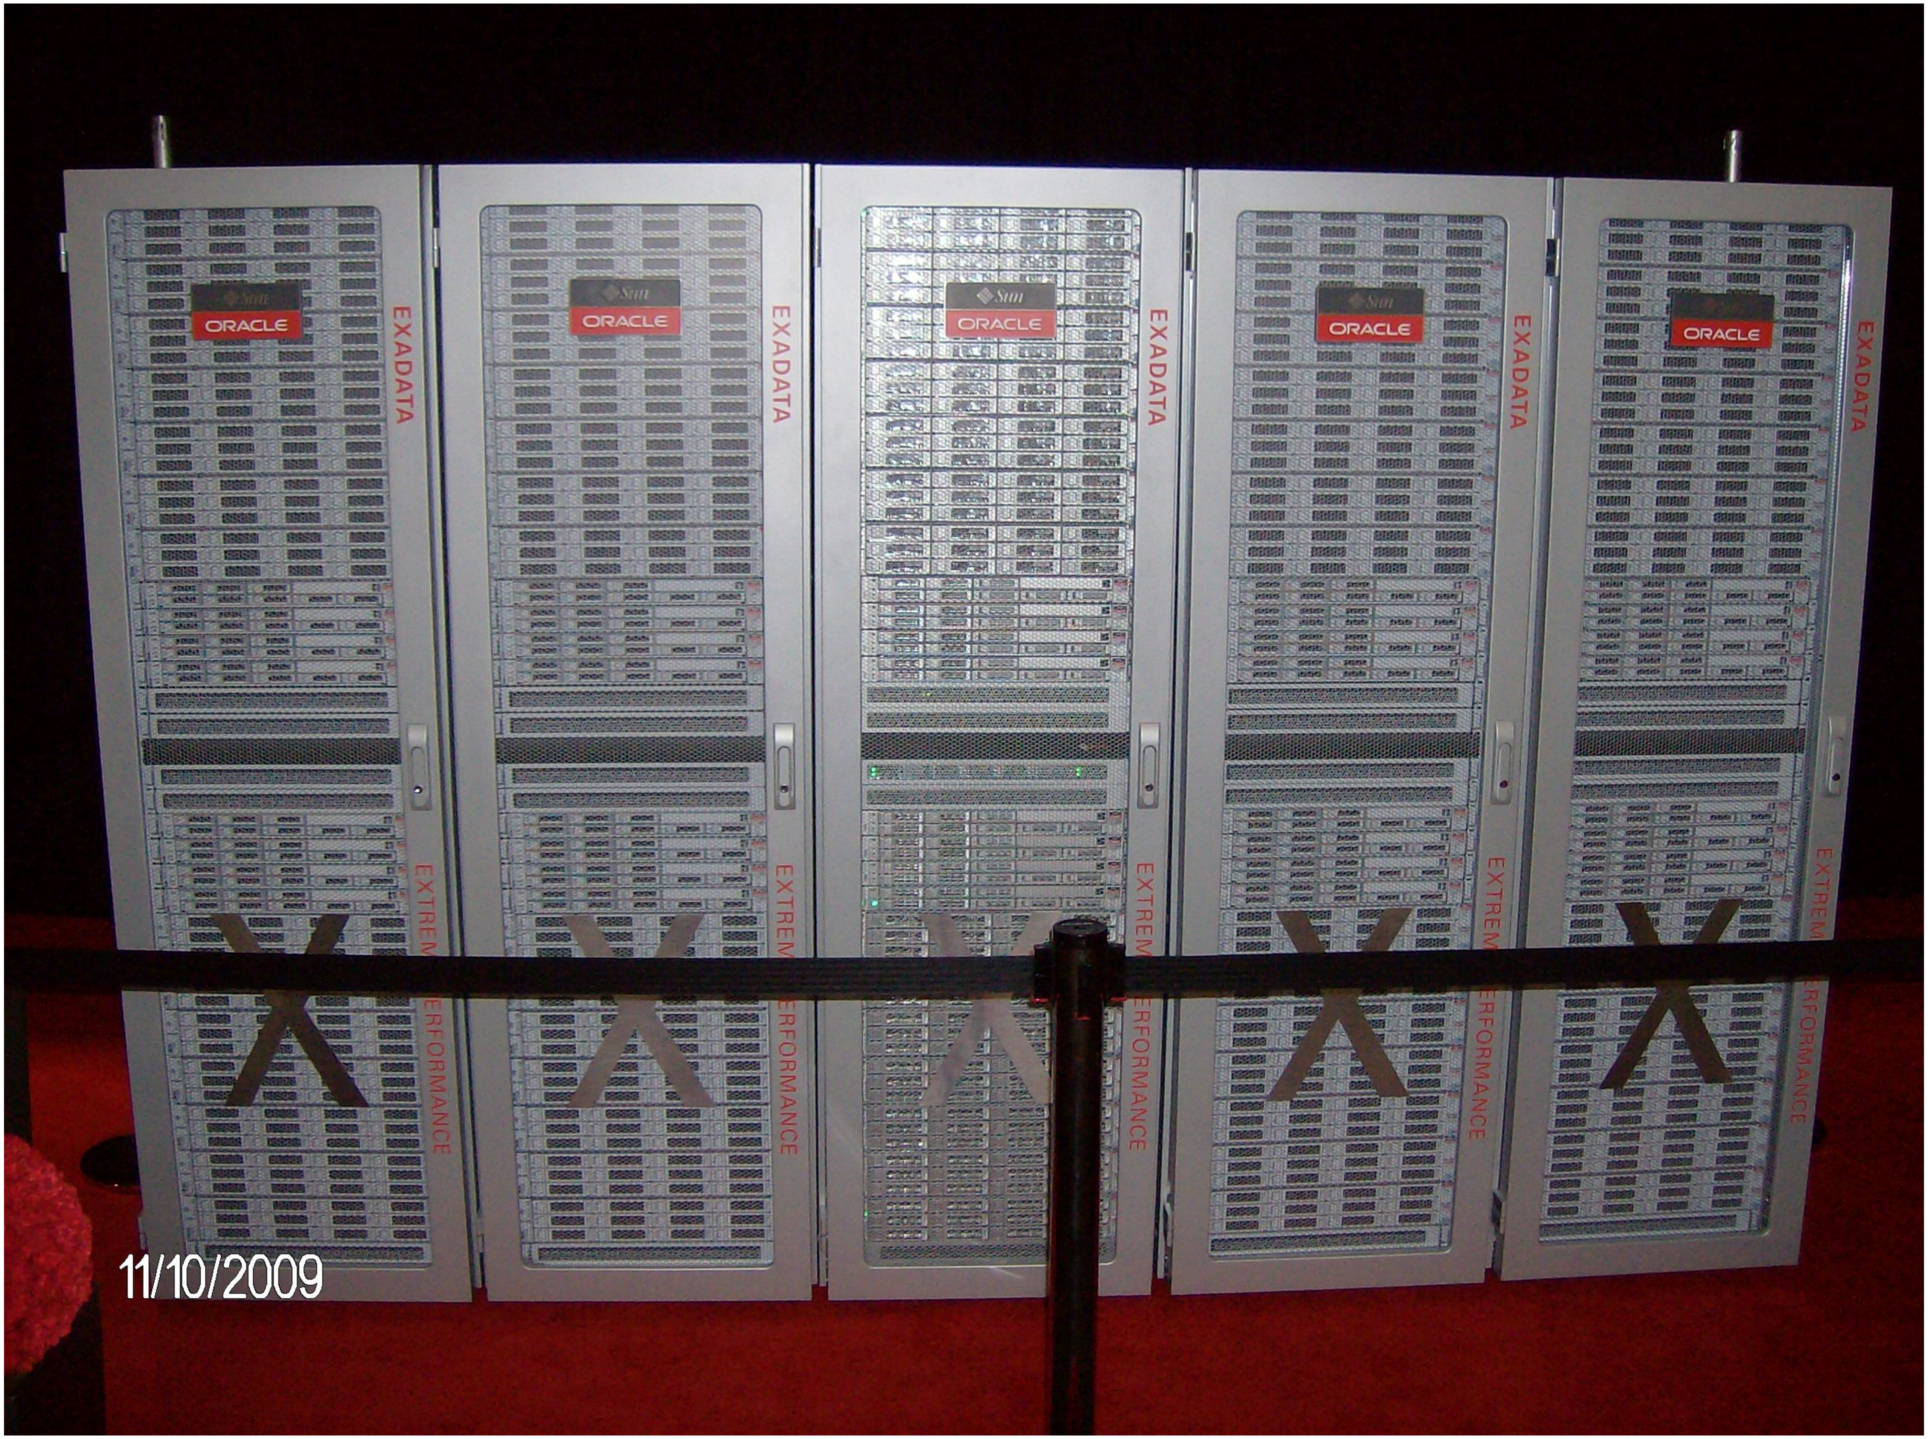
Sun
ORACLE
EXADATA
EXTREME PERFORMANCE
X
Sun
ORACLE
EXADATA
EXTREME PERFORMANCE
X
Sun
ORACLE
EXADATA
EXTREME PERFORMANCE
X
Sun
ORACLE
EXADATA
EXTREME PERFORMANCE
X
Sun
ORACLE
EXADATA
EXTREME PERFORMANCE
X
11/10/2009

Does **Exadata** eliminate

the need to think about software performance?

Exadata Version 2

- Twice as fast as Exadata Version 1 for data warehousing
- Now for OLTP too
- Ten times faster without making any changes to applications

http://www.oracle.com/us/corporate/press/033684

How?

Oracle Exadata Database Machine X2-8

128 CPU cores and 2TB memory for database processing

168 CPU cores for storage processing

5.3 TB of Exadata Smart Flash Cache

40 Gb/sec InfiniBand Switches

+

Oracle Database 11gR2

Exadata Smart Scan

Exadata Hybrid Columnar Compression

=

75 GB/sec uncompressed Flash bandwidth (25 GB/sec raw disk)

1,500,000 IO/sec Flash (50,000 IO/sec disk)

12 TB/hr data load rate

http://www.oracle.com/us/corporate/press/033684

How cool is that?!

Imagine…

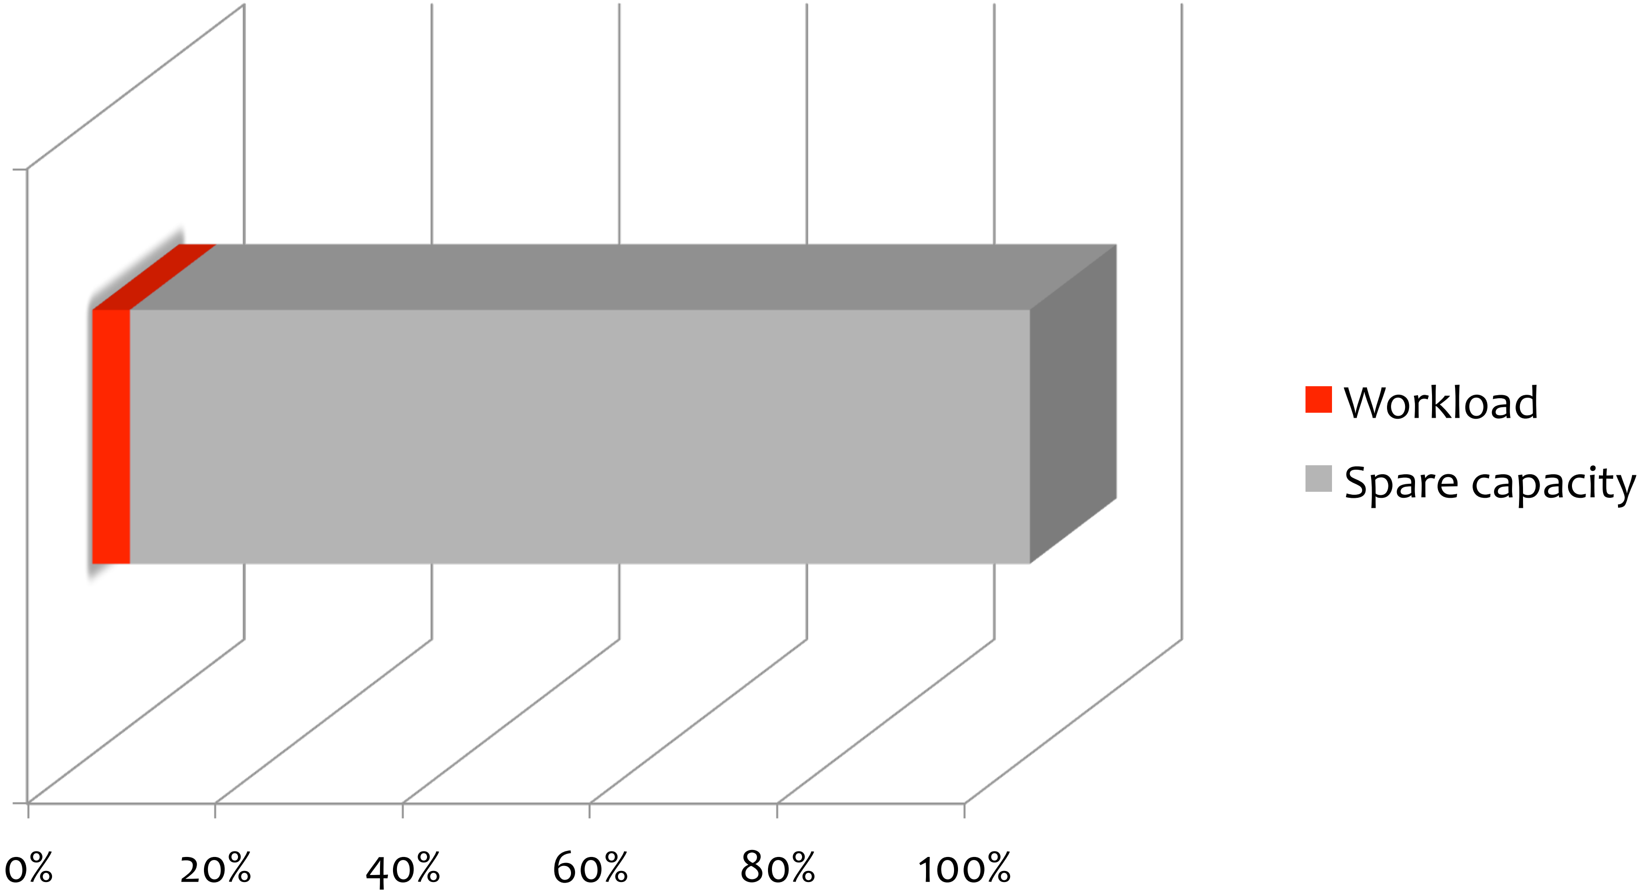

You’ll run hundreds, maybe thousands of databases on one of these machines.

LJE, Oracle OpenWorld 2009 keynote

Your data center.

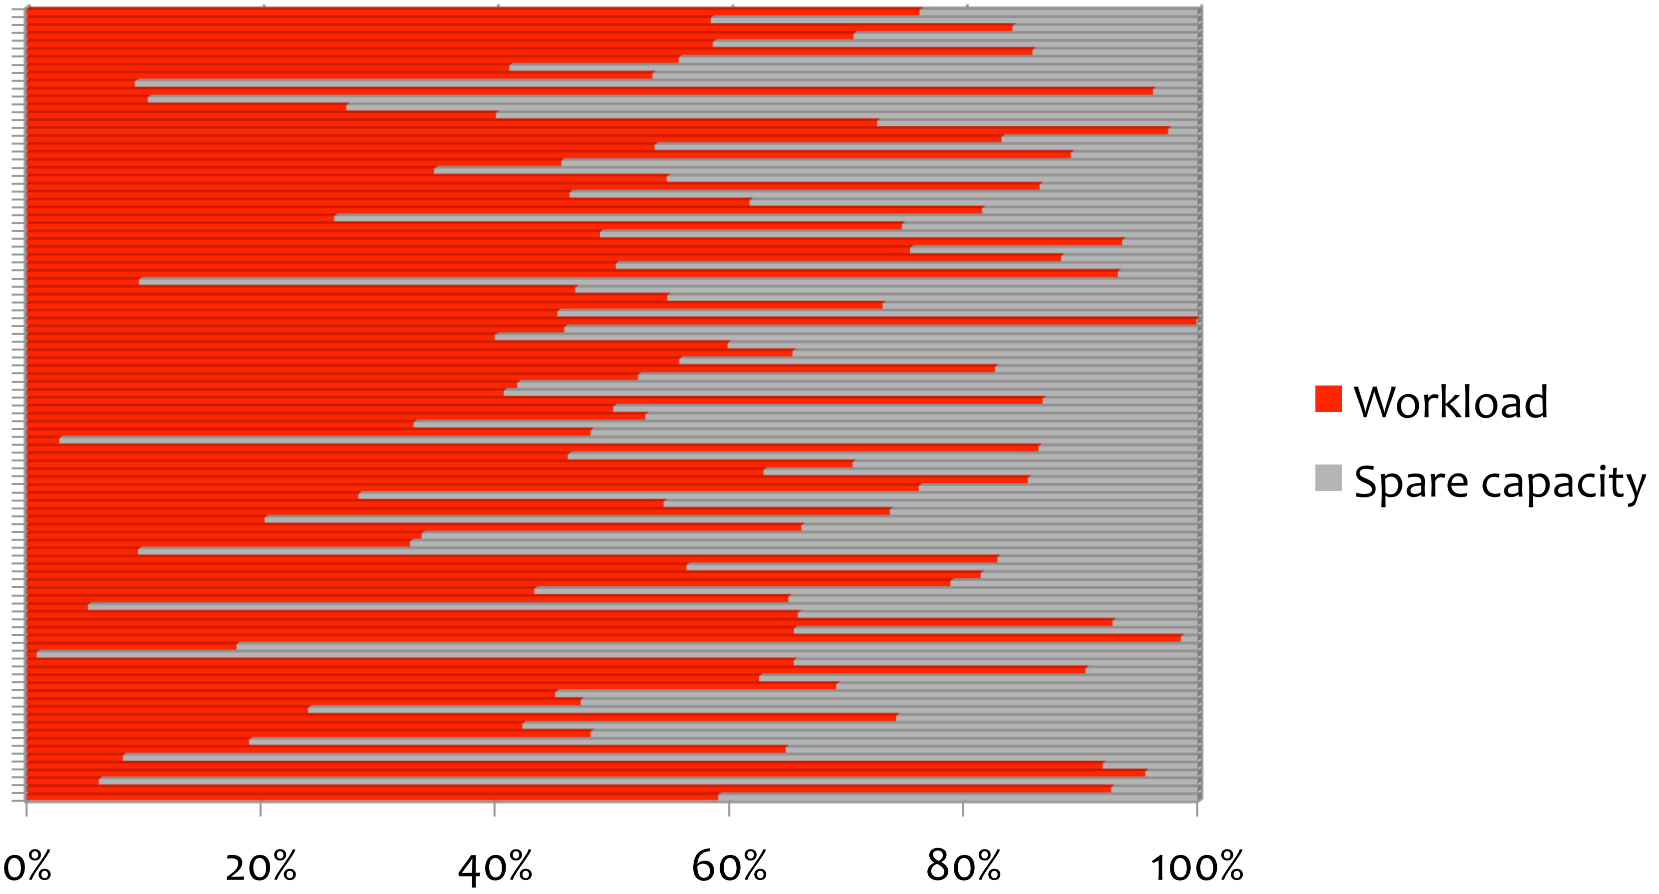

Your data center on Oracle Exadata.

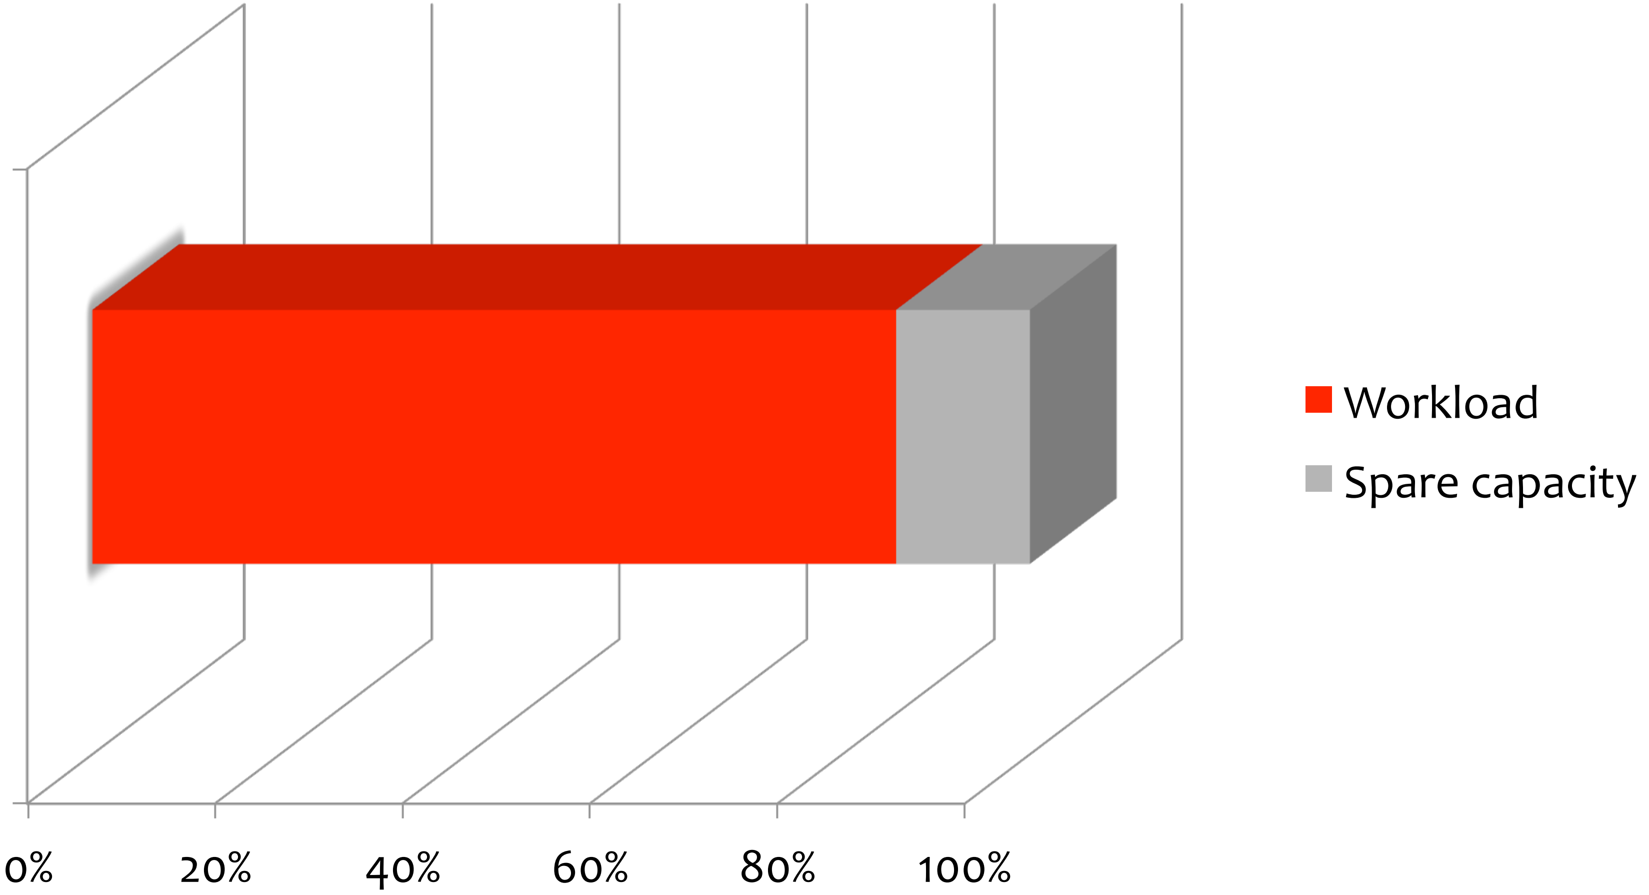

Absolutely remarkable system.

You’ll **still** want to use it efficiently.

The Future of Upgrades

From 9 August 2007

Inefficient applications will continue to be improved only marginally by capacity advances.

- Ludicrous…
  - “Memory-resident databases are fast”
  - “10x more CPU capacity makes us 10x faster”
  - “Fiber everywhere means fast everywhere”

Did I miss it?

Does hardware ever become *so* cheap that wasted capacity is ok?

…as long as your users don't notice it?

I think it’s…

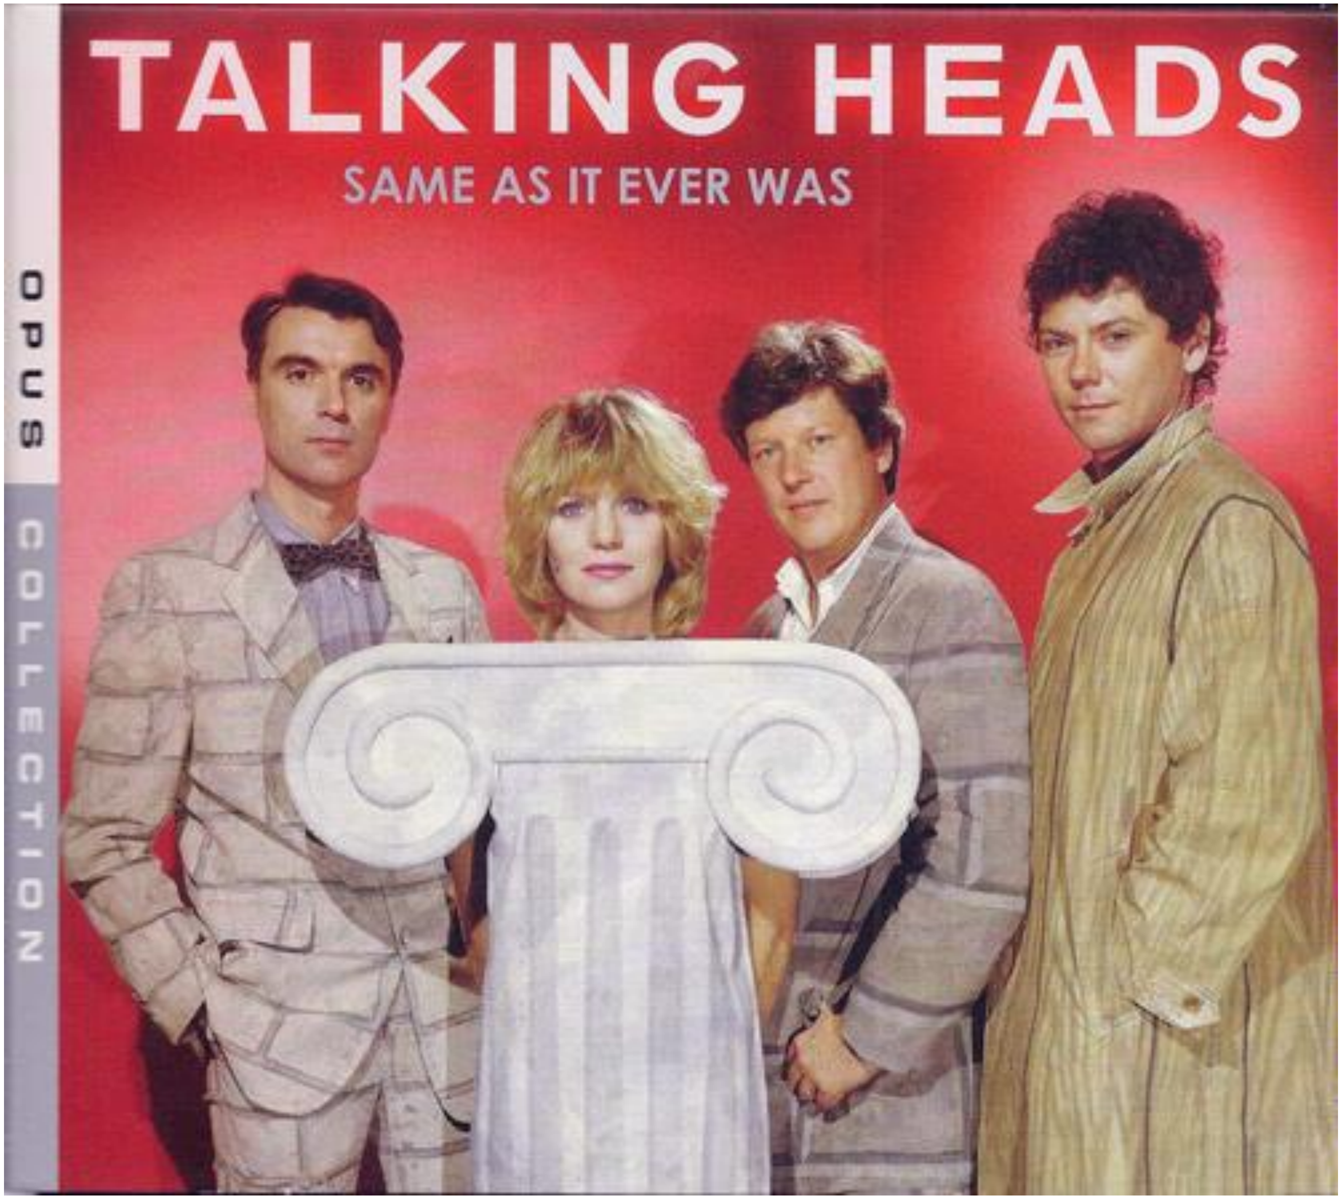
TALKING HEADS
SAME AS IT EVER WAS
OPUS
COLLECTION

Software is getting slower
more rapidly than
hardware becomes faster.

—Niklaus Wirth (1995)

Does Exadata eliminate the need to think about software performance?

No.

It only makes thinking about performance **more interesting**.

METHOD R™

Part 2

Leverage

Software is a more powerful force than hardware.

Example:

n CPUs
4*n* processes
each does parse, exec, fetch
all on exactly the same SQL statement

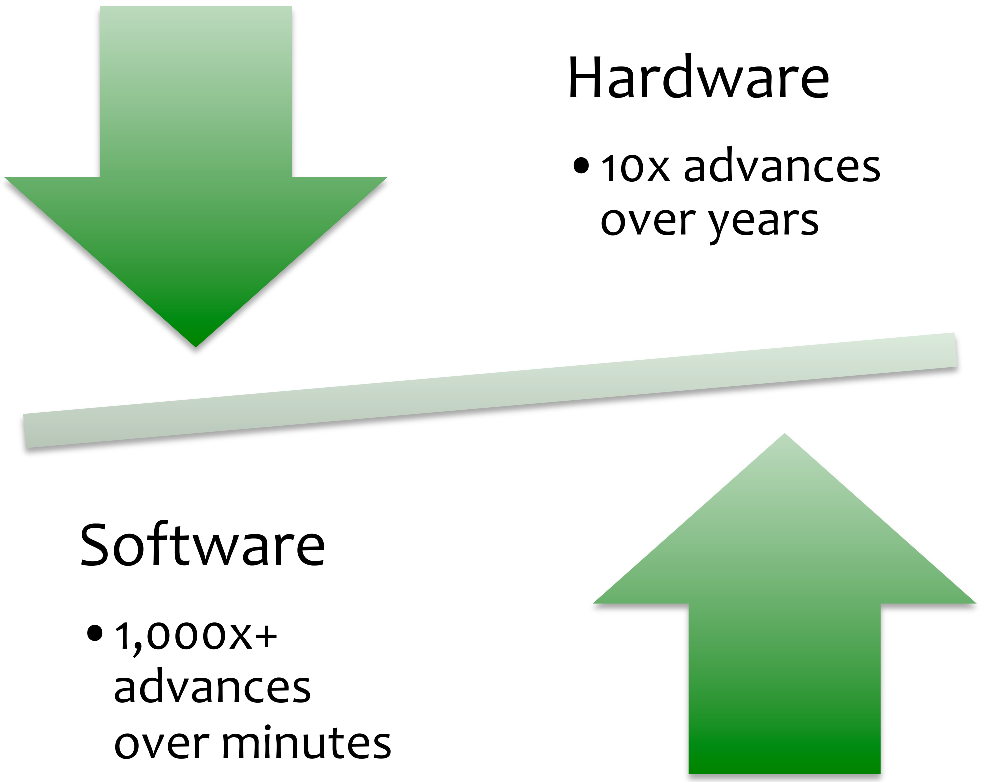

Hardware

- 10x advances over years

Software

- 1,000x+ advances over minutes

Remember the introduction of

PX?

Serial full-table scan (~2 minutes).

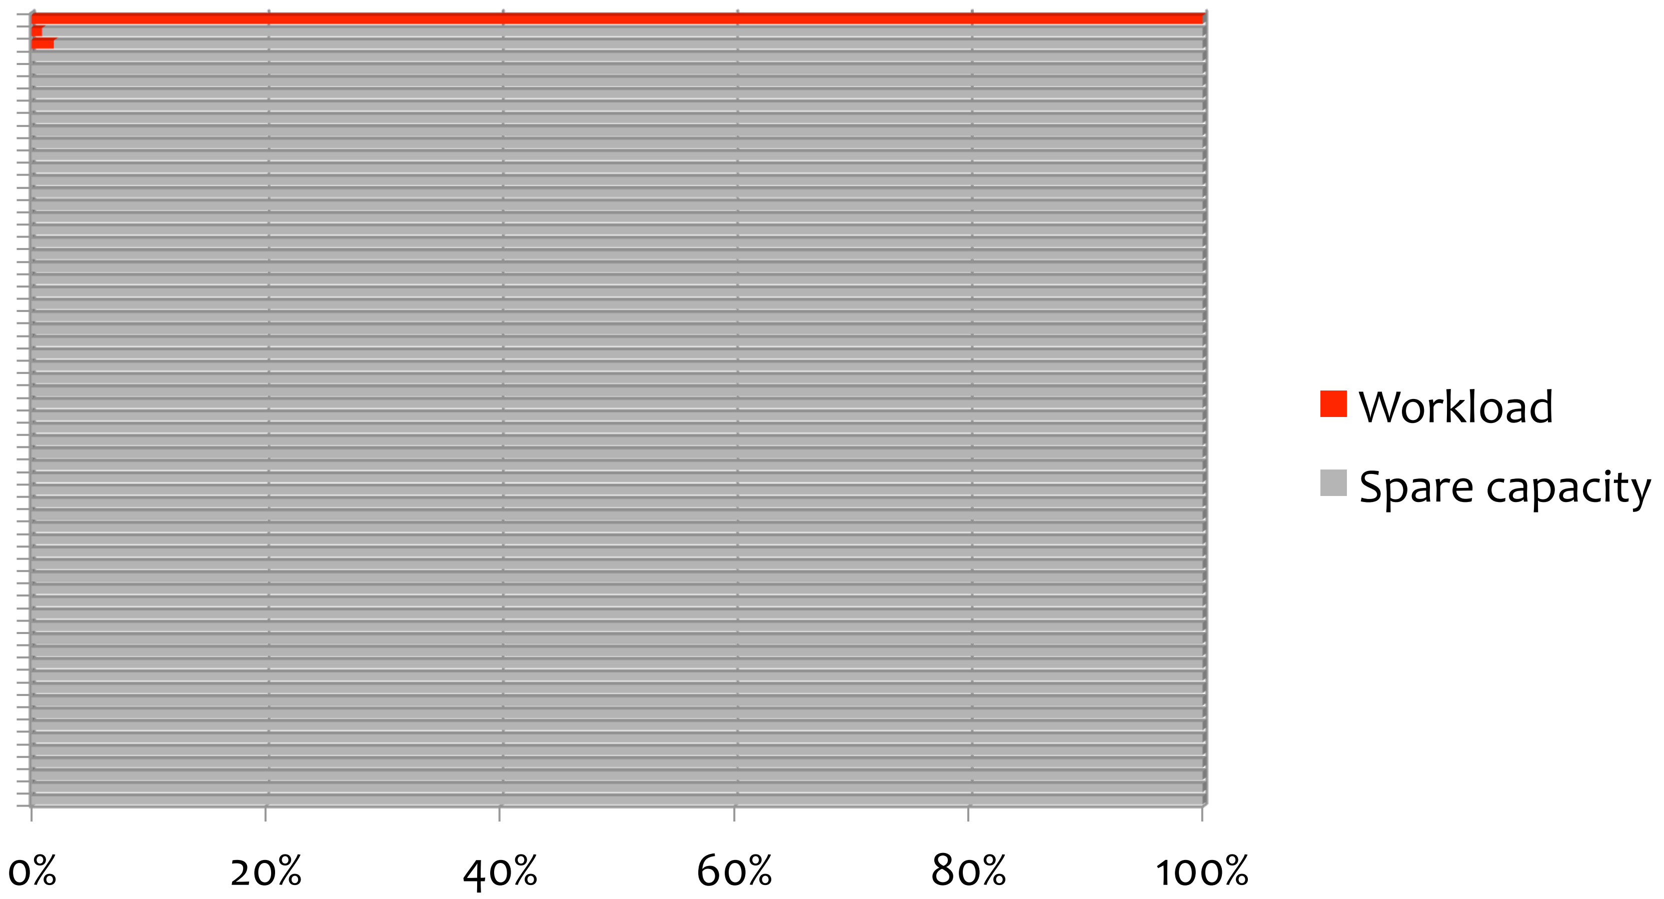

Parallel full-table scan (~2 seconds).

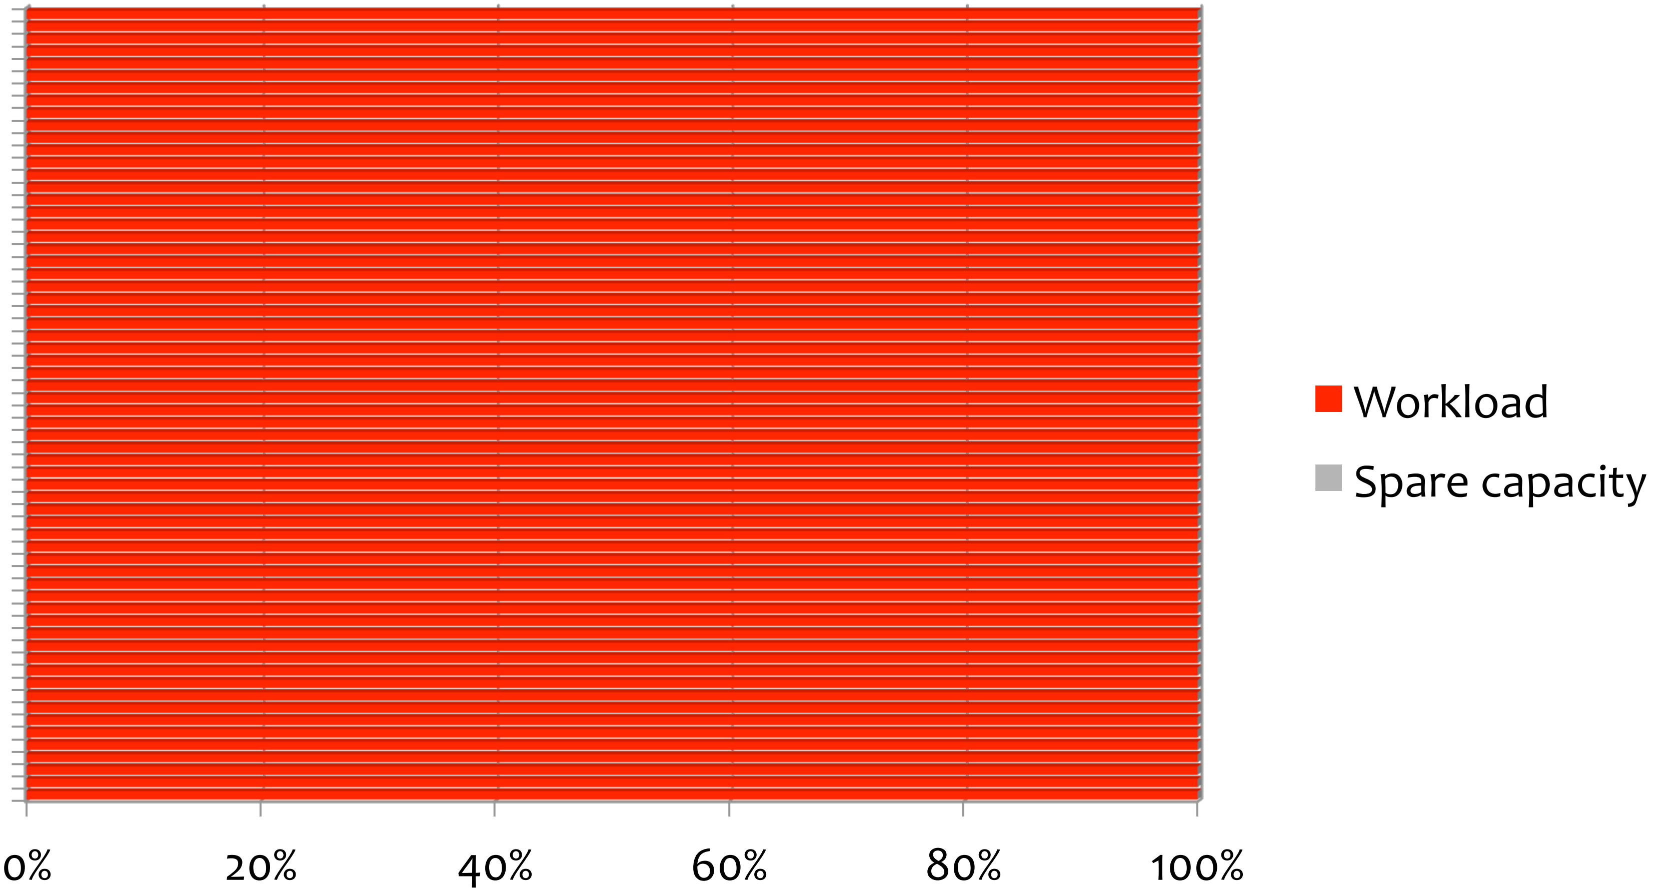

Parallel full-table scan (~2 seconds, after completion).

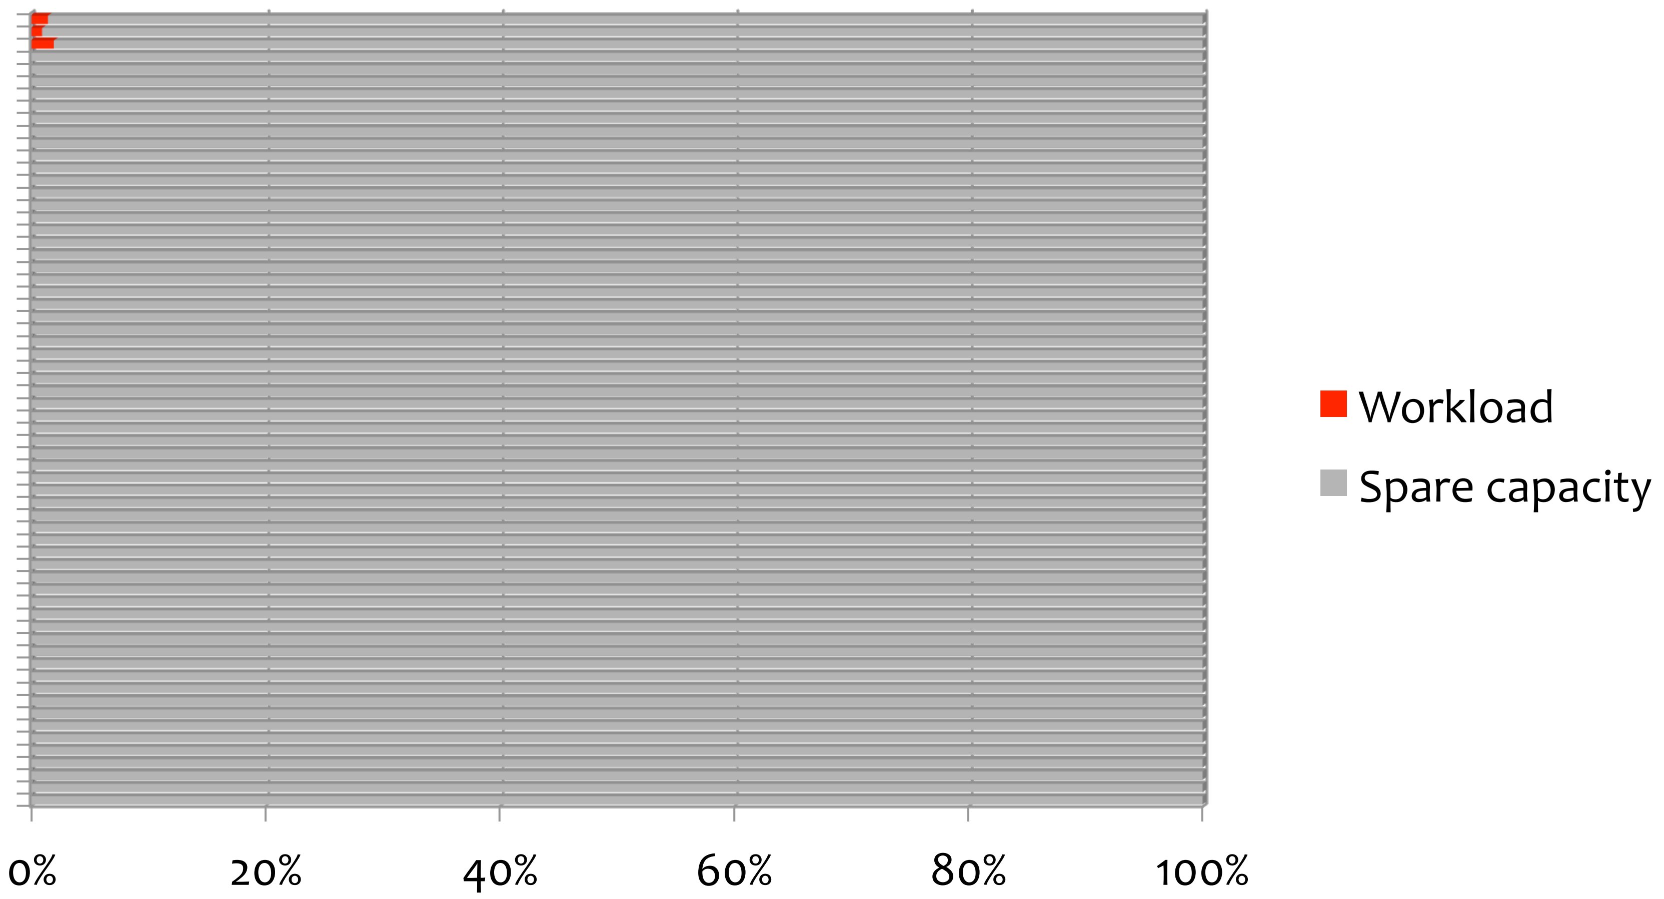

It’d be **nice** to run your query in PX on an *empty* 64-CPU system, wouldn’t it?

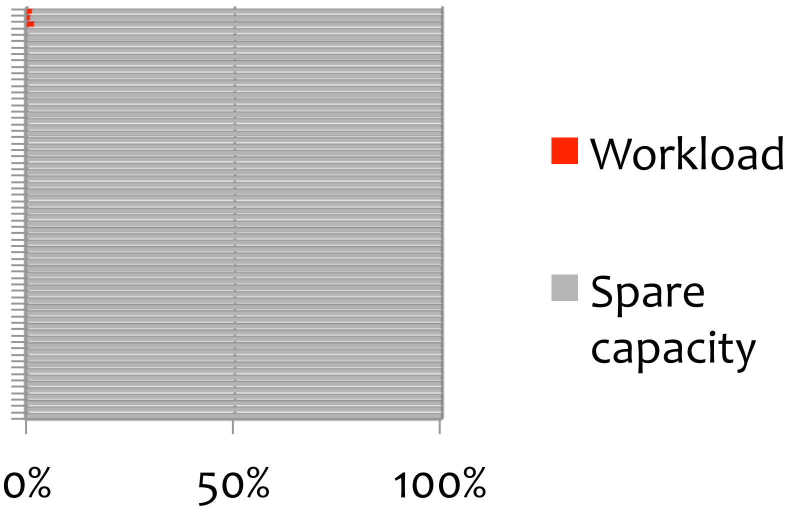

But what if your system looked like **this**

every time you tried to run something?

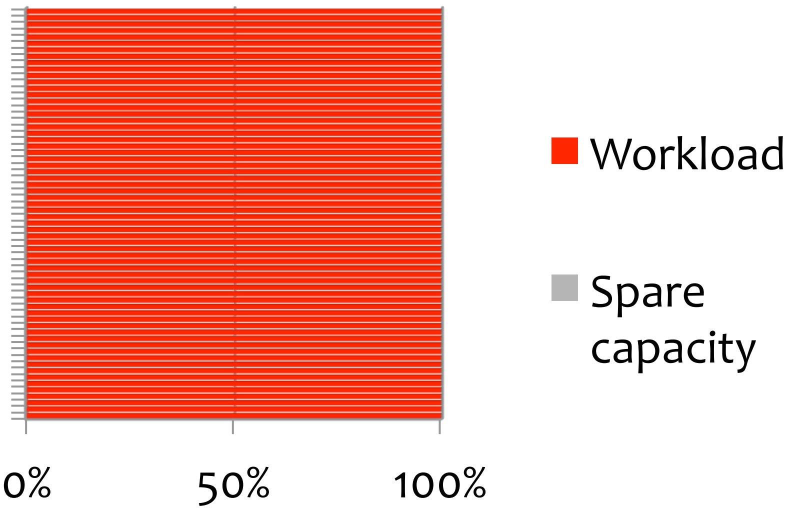

Software is a more powerful force than hardware.

The capacity of *any* machine is **finite.**

The capacity of *any* machine is **finite**.

The amount of work that software can ask of a machine is **unbounded**.

Capacities are bigger now.

But so are workloads.

And workloads can grow **without bound.**

Being *efficient* will still require…

intelligence

effort

investment

creativity

METHOD R™

◇

Part 3

One wish

Are you organized like this?

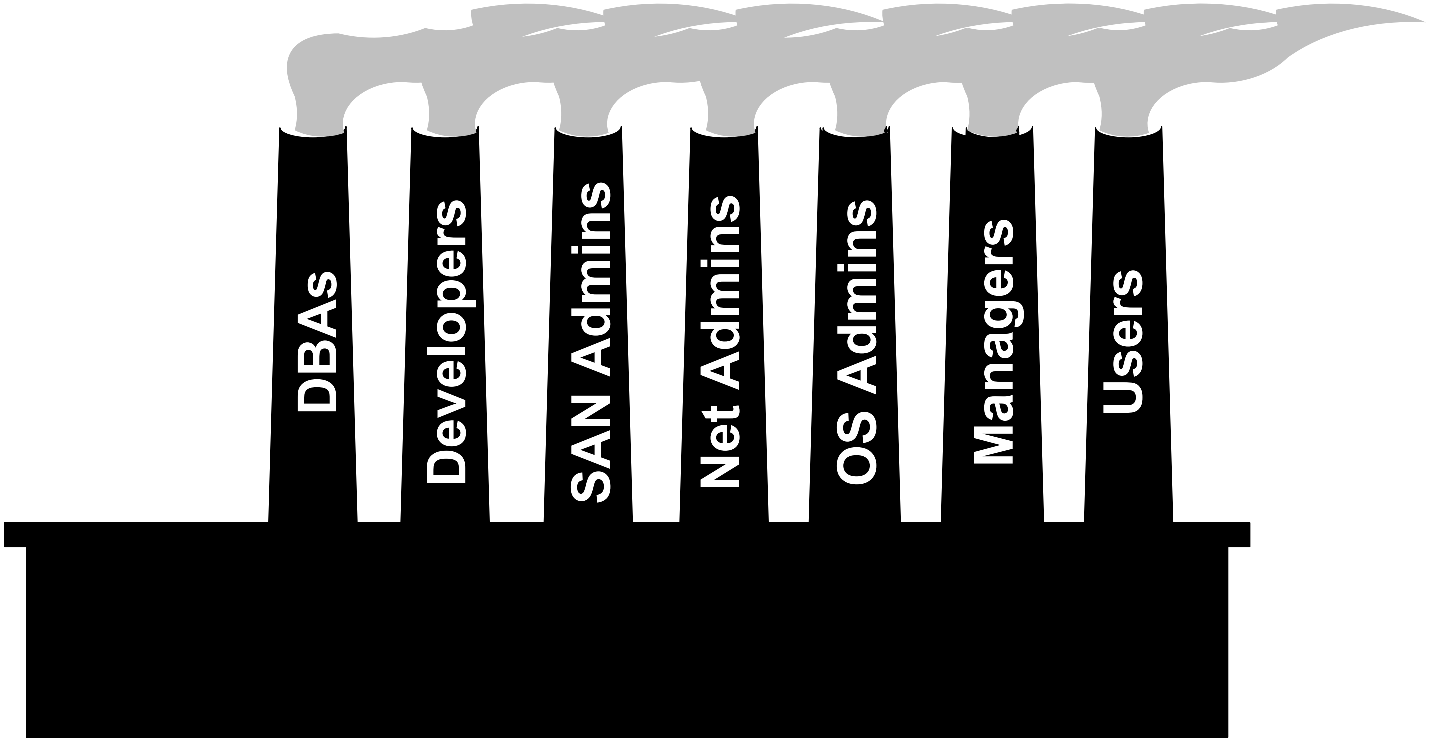

DBAs
Developers
SAN Admins
Net Admins
OS Admins
Managers
Users

If your subsystems are optimized,
then your system is optimized.

Optimized subsystems does not imply an optimized system.

```
$ 378 Airfare
$ 476 Hotel 4@$119


---------------------
$ 854 TOTAL
```

```
$ 378 Airfare
$ 396 Hotel 4@$99
$ 200 Car 4@$50
$  80 Parking
---------------------
$1054 TOTAL
```

(Developer ∪ DBA)

>

(Developer + DBA)

DBA stereotypes

Developer stereotypes

Comparison of focus…

DBAs

- Uptime pressure
- Have to “make it work”
- Have to mingle across the stack more

Developers

- Deadline pressure
- Have to be creative
- Feel less limited by technology
- Fail, like all the rest of us, predominantly when feedback loop is large

Comparison of typical skills…

DBAs

- Operations
- Resource management
- Troubleshooting
- Physical data design
- System interfaces
- Data modeling
- SQL optimizer
- Oracle product features

Developers

- Algorithms
- Languages
- Frameworks
- Design
- User interfaces
- Data modeling
- Debugging, profiling, tracing
- Business-domain processes

Skills both should share…

DBAs

- Operations
- Resource management
- Troubleshooting
- Physical data design
- System interfaces
- **Data modeling**
- **SQL optimizer**
- **Oracle product features**

Developers

- Algorithms
- Languages
- Frameworks
- Design
- User interfaces
- **Data modeling**
- **Debugging, profiling, tracing**
- **Business-domain processes**

An idea that *will* make your performance better:

Developer ∪ DBA

Discussion…

Method R™

Thank you

http://method-r.com