



YOUR WAY TO HIGH AVAILABILITY.

LIN:BIT

Virtual Block Device Performance



Percona Performance Conference 2009
Santa Clara, CA

Florian Haas
LINBIT HA Solutions GmbH
Vienna, Austria

So What Are We Talking About?

- “Virtualization overhead” is not just CPU and memory overhead
- Most applications are I/O bound
- I/O performance is critical in virtualization
- Let's take a look at block device performance in virtual settings!

What Are We Comparing?

- Relative virtual block device I/O performance in various open source hypervisor environments
- I/O latency
- I/O throughput
- Reference platform: Debian 5.0 with 2.6.27 kernel (host and guest)

Platforms tested

- Baseline: reference system running in chroot
- Xen 3.2 with paravirtualized guest
- Xen 3.2 with HVM guest
- KVM with IDE emulation
- KVM with paravirtualized block devices (virtio)

What Are We *Not* Talking About?

- Absolute block device performance
- Non-open source hypervisors
- Network I/O performance
- Filesystem performance
- Non-Linux hosts
- Non-Linux guests

I/O Throughput

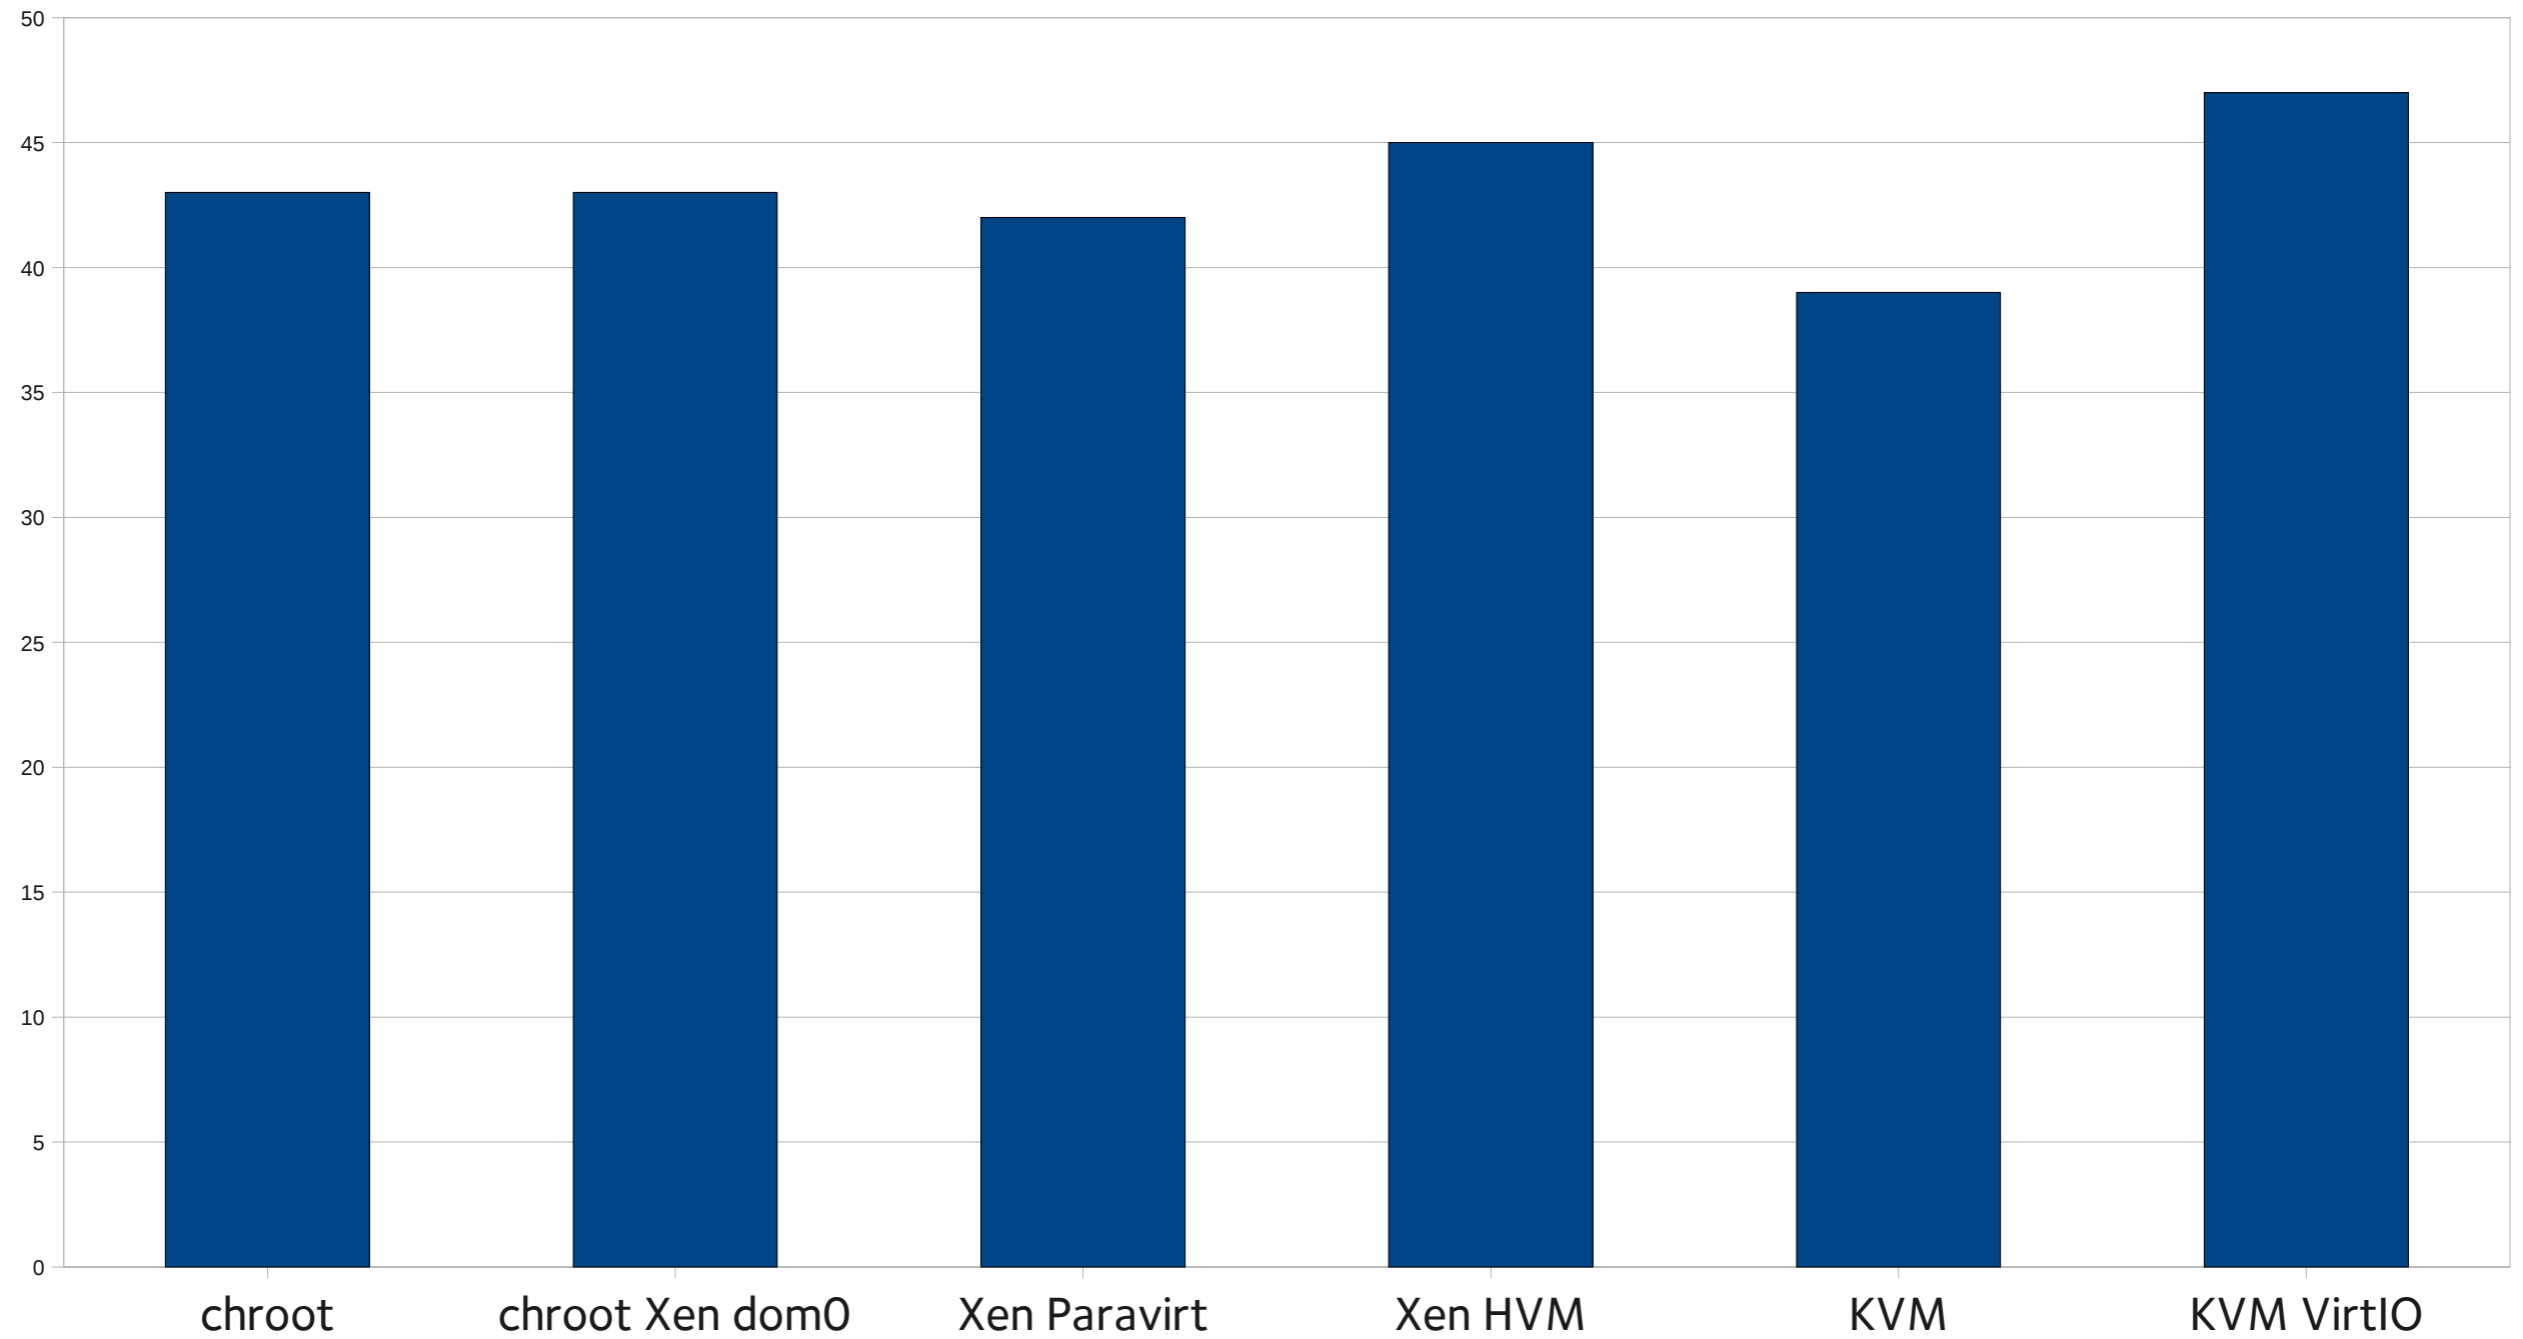
„How much data you can write in a given period of time?“

How We Test Block Device Throughput

```
dd if=/dev/zero of=/dev/vdb bs=1G count=1 oflag=direct
```

(no, really, it's that simple!)

Comparative Throughput Results



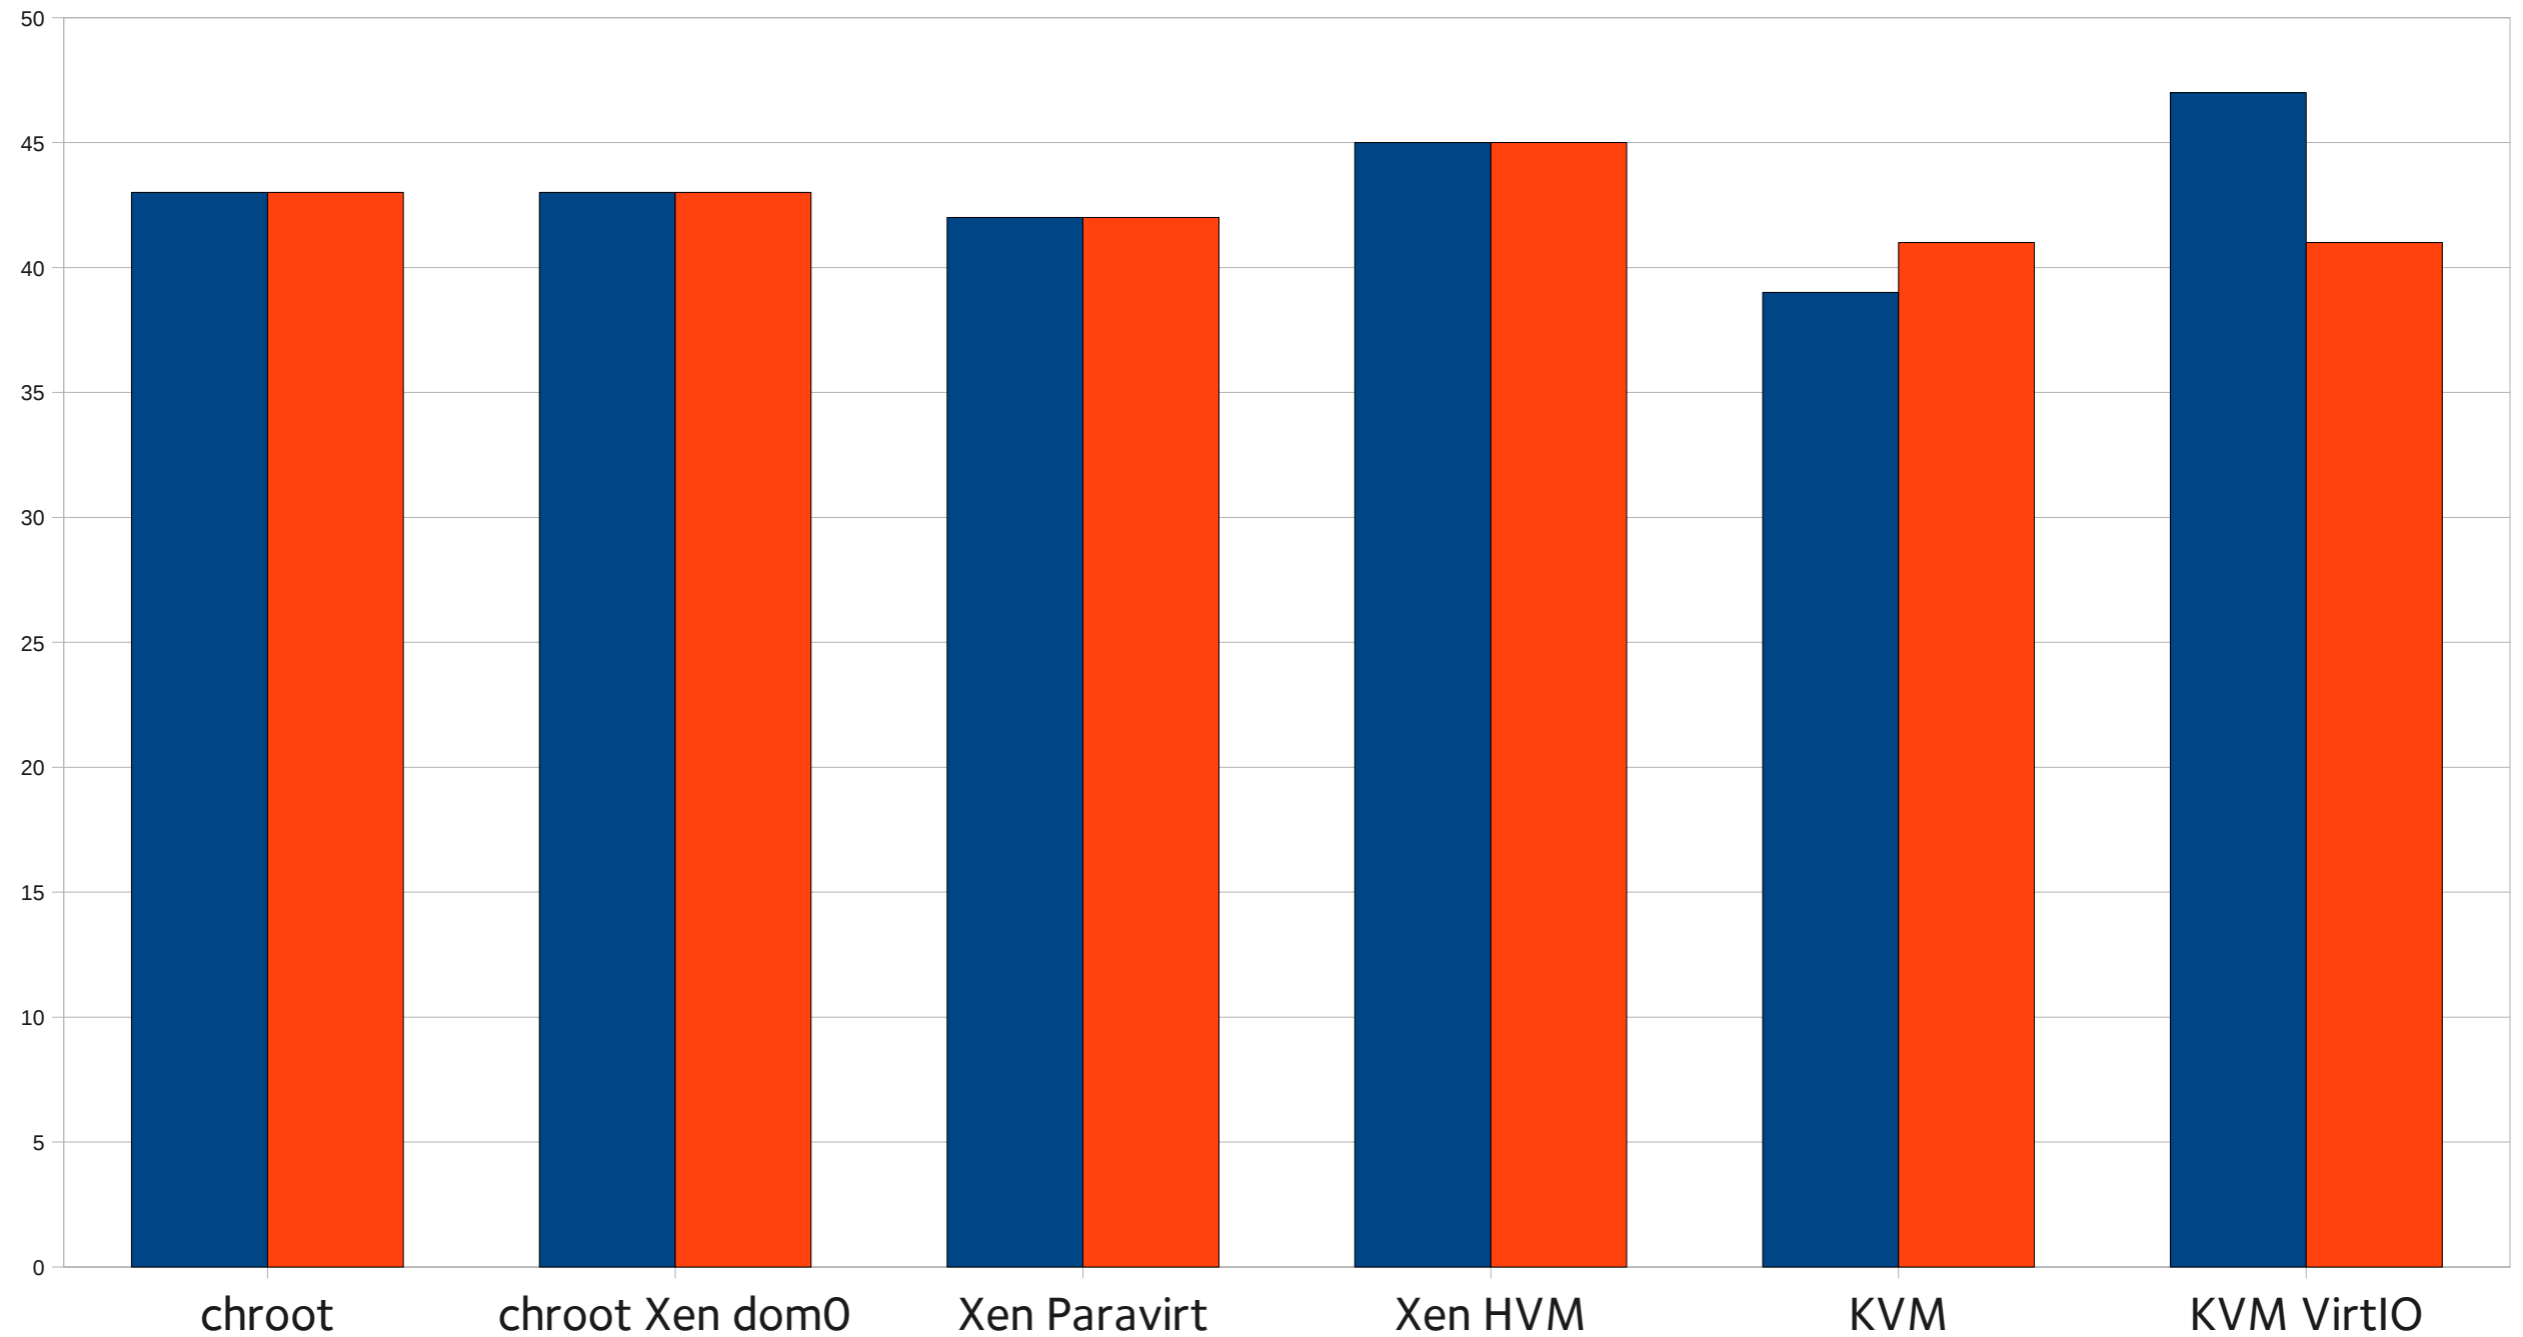
Huh?

?

YOUR WAY TO HIGH AVAILABILITY.

LIN:BIT

Comparative Throughput Results (Caches Off, Really?)



I/O Latency


„How fast you can write the smallest possible chunk of data?“

How We Test Block Device Latency

```
dd if=/dev/zero of=/dev/vdb bs=512 count=1000 oflag=direct
```

(subtle changes make a big difference)

Comparative Latency Results



What Do We Take Away From This?

- Xen bashing is en vogue, but in terms of performance it's actually pretty damn good
- KVM performance is sorry with full device emulation (expected)
- KVM+VirtIO is better, but nowhere near Xen Paravirt
- Hosts lie to qemu-based guests by default
- Disabling caches in Xen 3.2 HVM... doesn't.

Get in touch!

- On the web:
 - www.linbit.com
 - <http://blogs.linbit.com/florian>
- By email:
 - florian@linbit.com