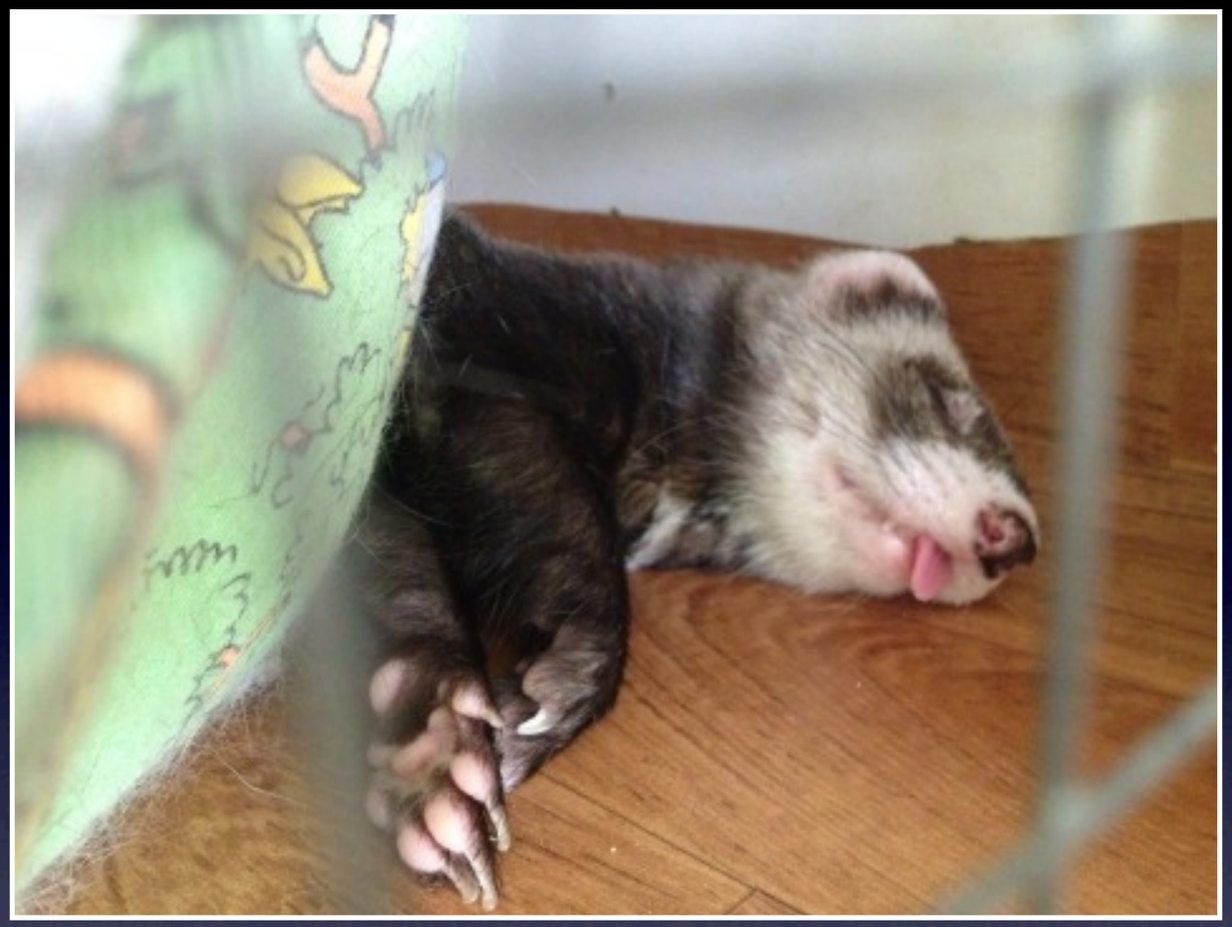

Semiassignment

The best of both worlds

Device Assignment

QEMU

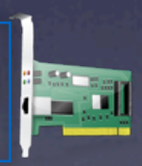

Device Assignment

QEMU

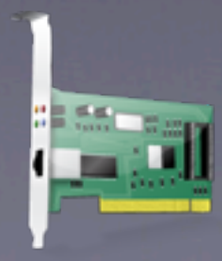

Pros and Cons

	Emulation	Assignment
Overhead	high	low
Throughput	low	high
Latencies	high	low
Migration	yes	no
Scalability	yes	limited

Device Assignment

- Preferred for runtime
- No migration
- Static assignment

Getting both

- Assigned device during normal operation
- Emulated device during migration

The Xen Way

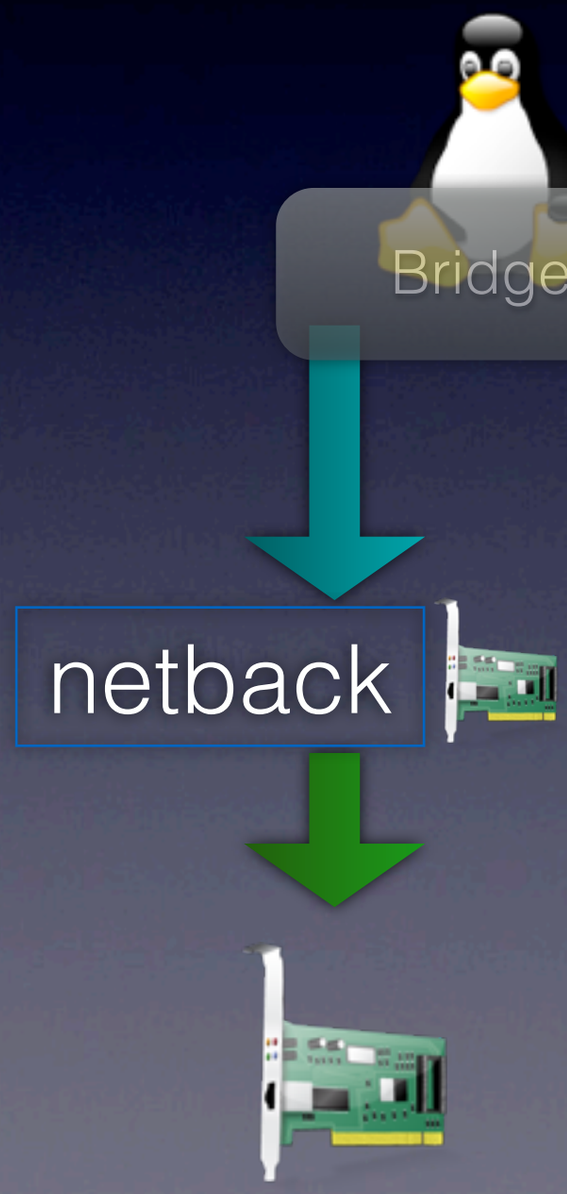

The Xen Way

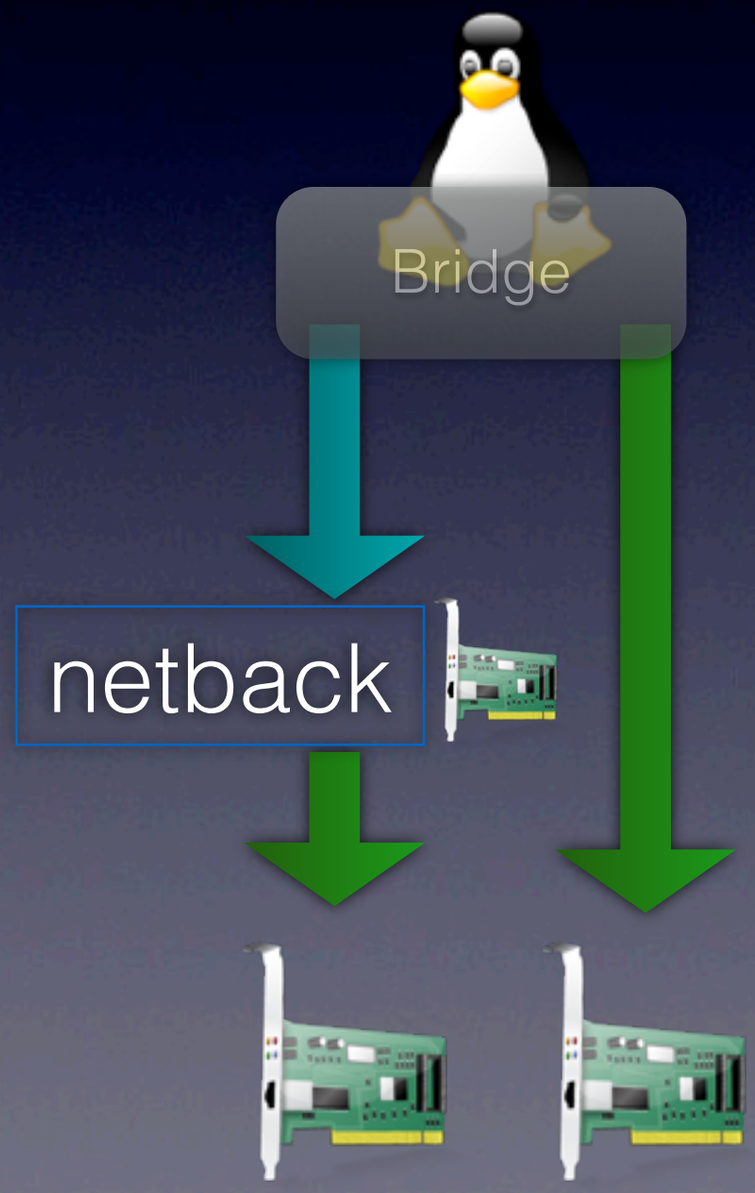

The Xen Way

- Guest changes for bridge
- Migration is guest visible

The VMware Way

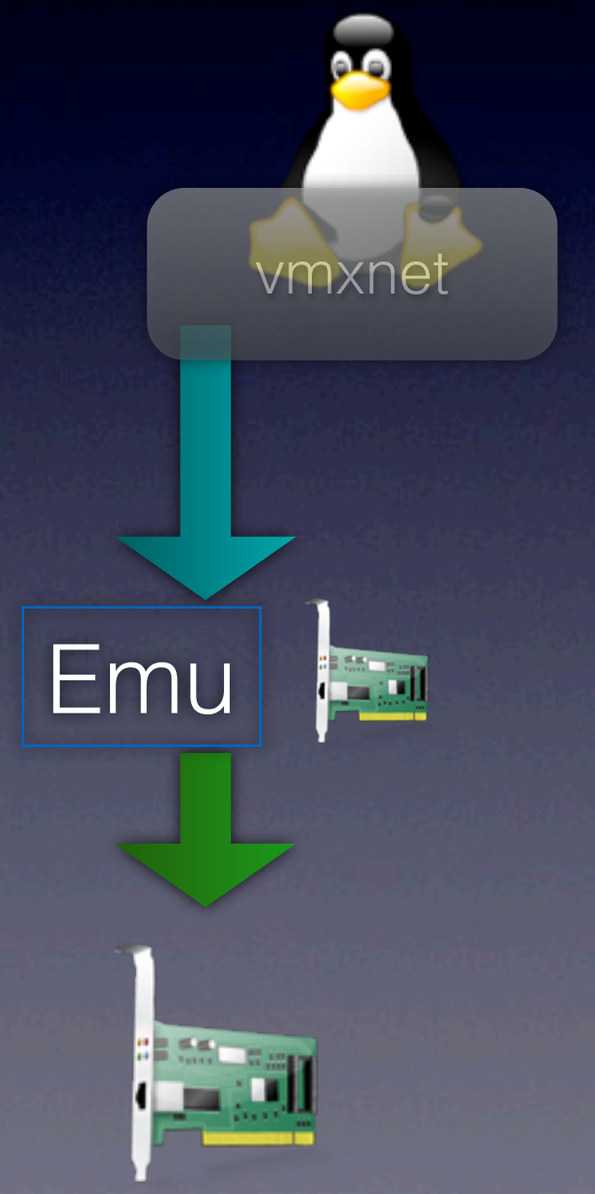

The VMware Way

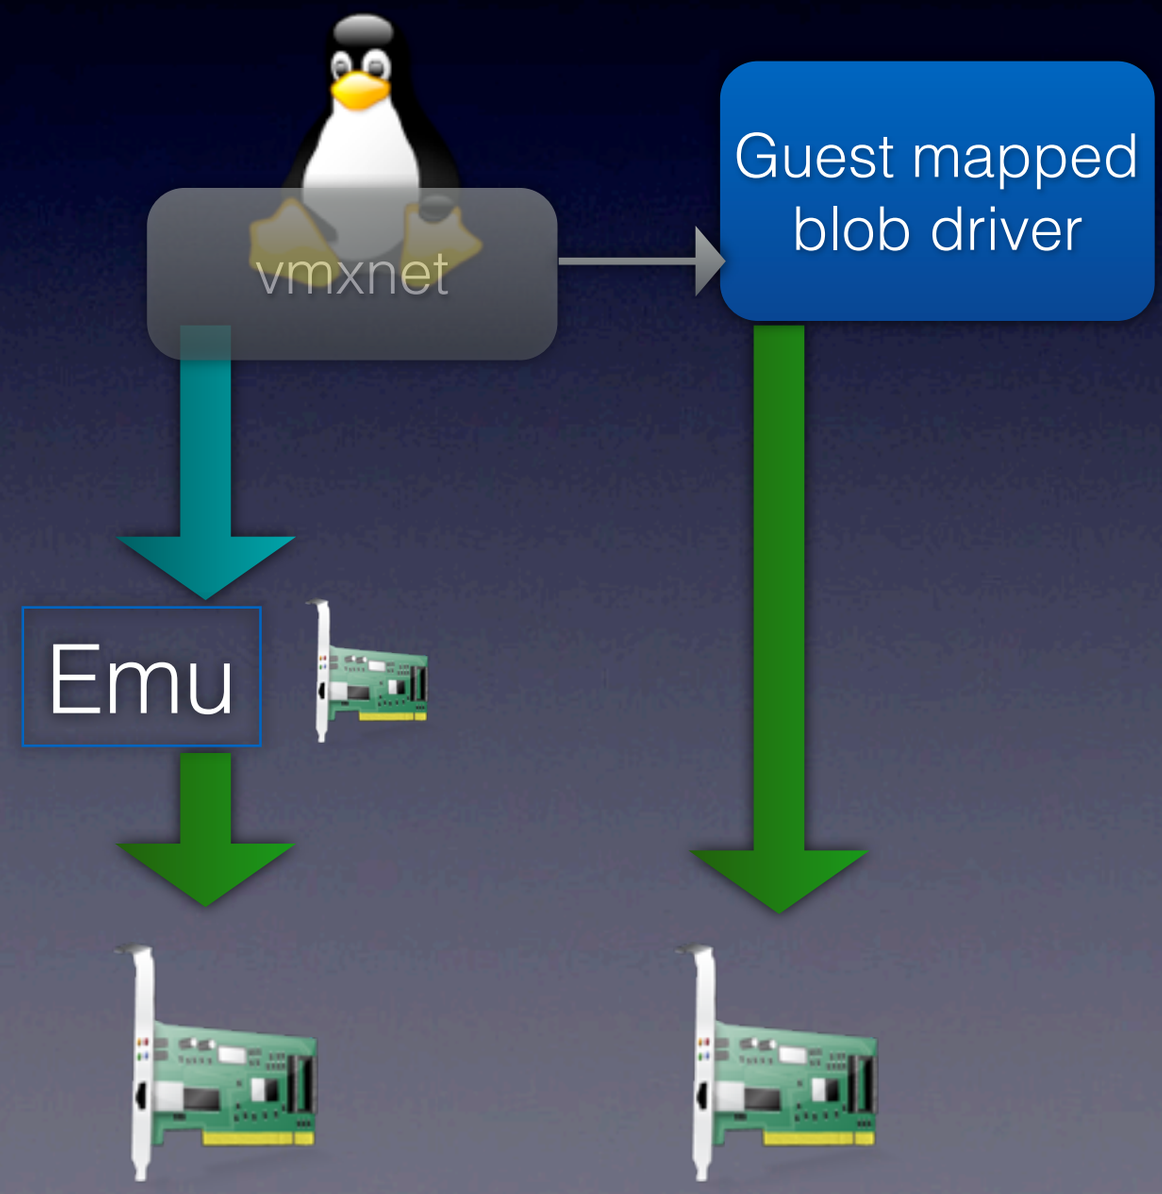

The VMware Way

- Provide network driver in magic memory region
- Requires new drivers
- Not GPL friendly
- Migration is guest exposed

The Alex way

QEMU

The Alex way

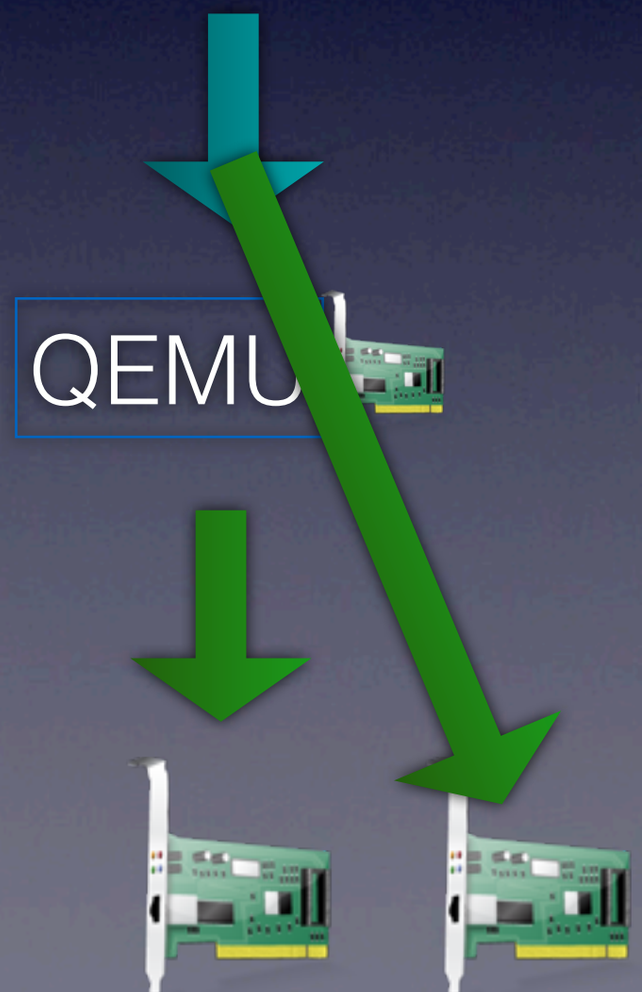

The Alex Way

- Emulate real device in QEMU
- Migrate state between emulated and real adapter
- Need to write emulation and migration code for every adapter
- Only works well if cluster uses the same cards

Pros and Cons

	Emulation	Assignment	Semi
Overhead	high	low	low
Throughput	low	high	high
Latencies	high	low	low
Migration	yes	no	yes
Scalability	yes	limited	yes
Effort	low	low	high

Demo

Semiassignment

- Is it a good idea?
- How much effort really? We only need support for a few (SR-IOV) adapters.
- More complicated network configuration
- Volunteers?

Donnerstag, 20. September 12