



project **tsubame**

Final Year Project 2008 - Load Balancing Xen
Jeremy Tay BEng (CS) III

Agenda


- State of the Art - Server Consolidation and Virtualization
- The Project - Aims and Objectives
- System Architecture, Features and Implementation
- Demonstration
- Limitations
- Future Development
- Q & A



Server Sprawl


Decreased utilization,
energy efficiency

Increased administrative
overhead, running cost





Virtualization





The Project



Aims and Objectives

To create an **ENVIRONMENT** for the execution of **MIGRATABLE VIRTUAL MACHINES** on a **SCALABLE** array of computers running on commodity hardware.

To study the possibilities of **LOAD BALANCING** the virtual machines through **LIVE MIGRATION**.

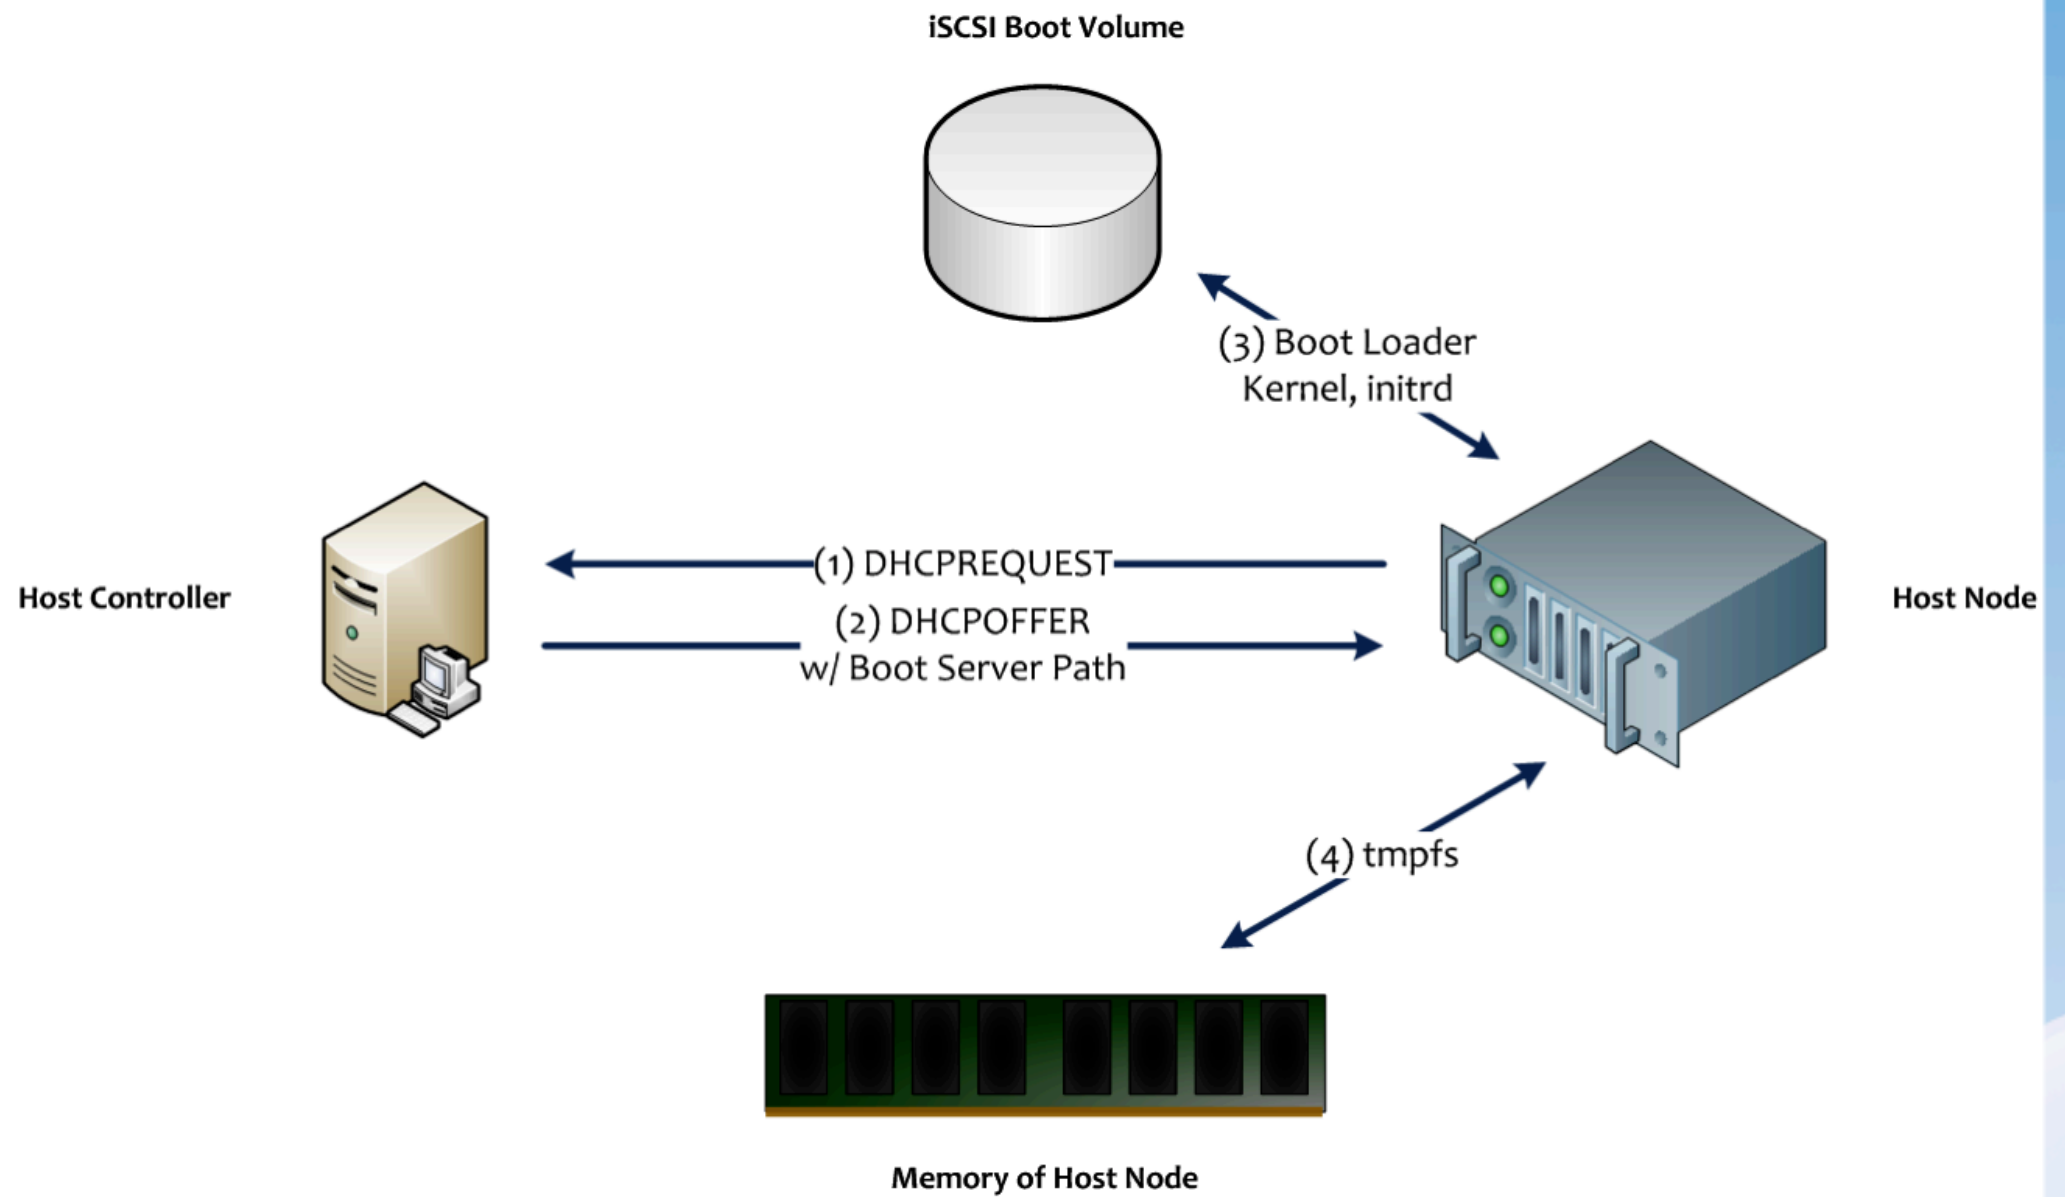


Reliability, Availability, Scalability



Scalable Environment

- Network boot
- Diskless, headless hosts
- Isolated operating system
- Rapid deployment





+

+

Live Migration

- Moving virtual machines from one physical host to another
- Decouples service from hardware
- Negligible service interruption
- Transparent to user


+

+



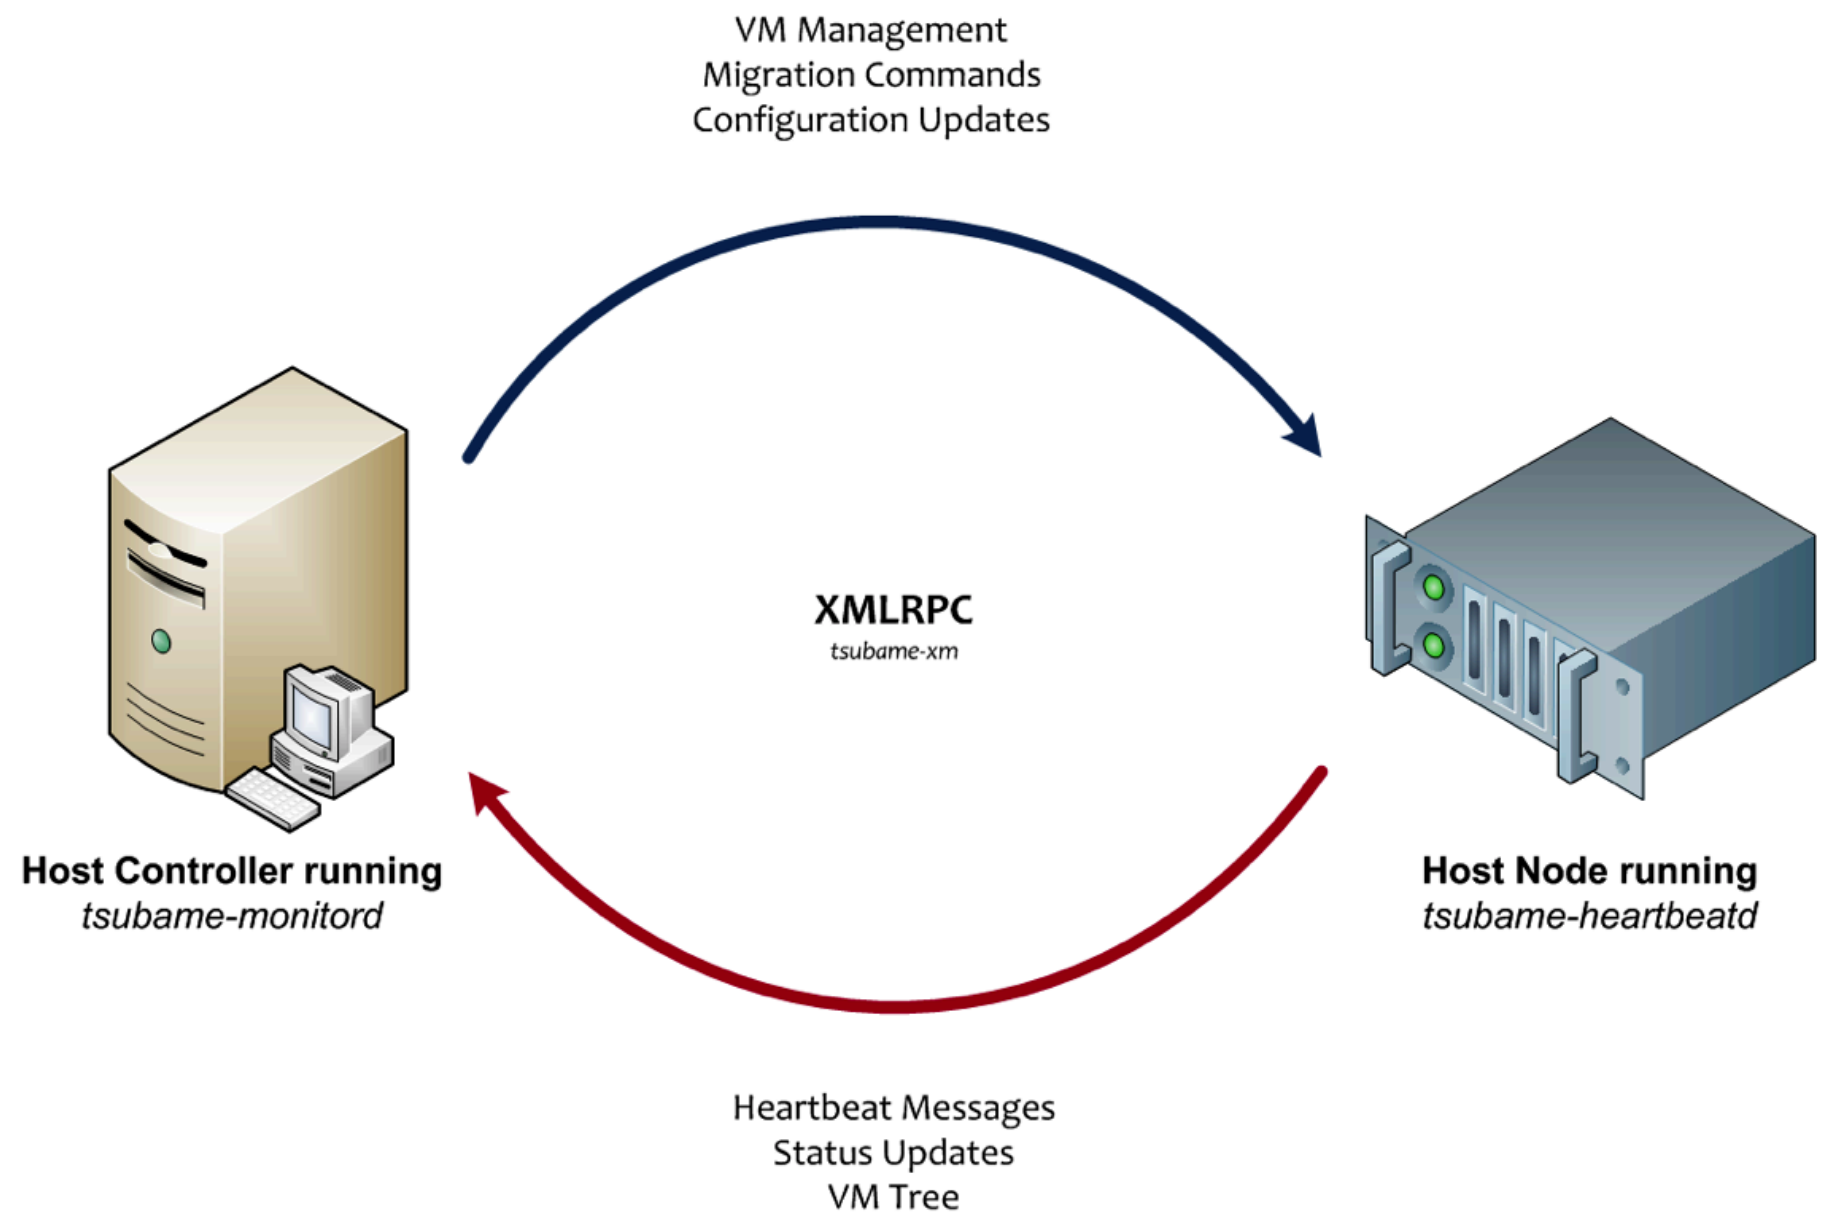
?






Load Balancing

- Optimize utilization
- Reduce bottlenecks
- Increase manageability and availability
- Done through live migration




Load Balancing


- Basic , On-demand Load Balancing
- Rule-based Load Balancing
- Statistics-based Load Balancing




Basic, On-demand Load Balancing




Strategy - Affinity Rules




Strategy - Anti-affinity Rules



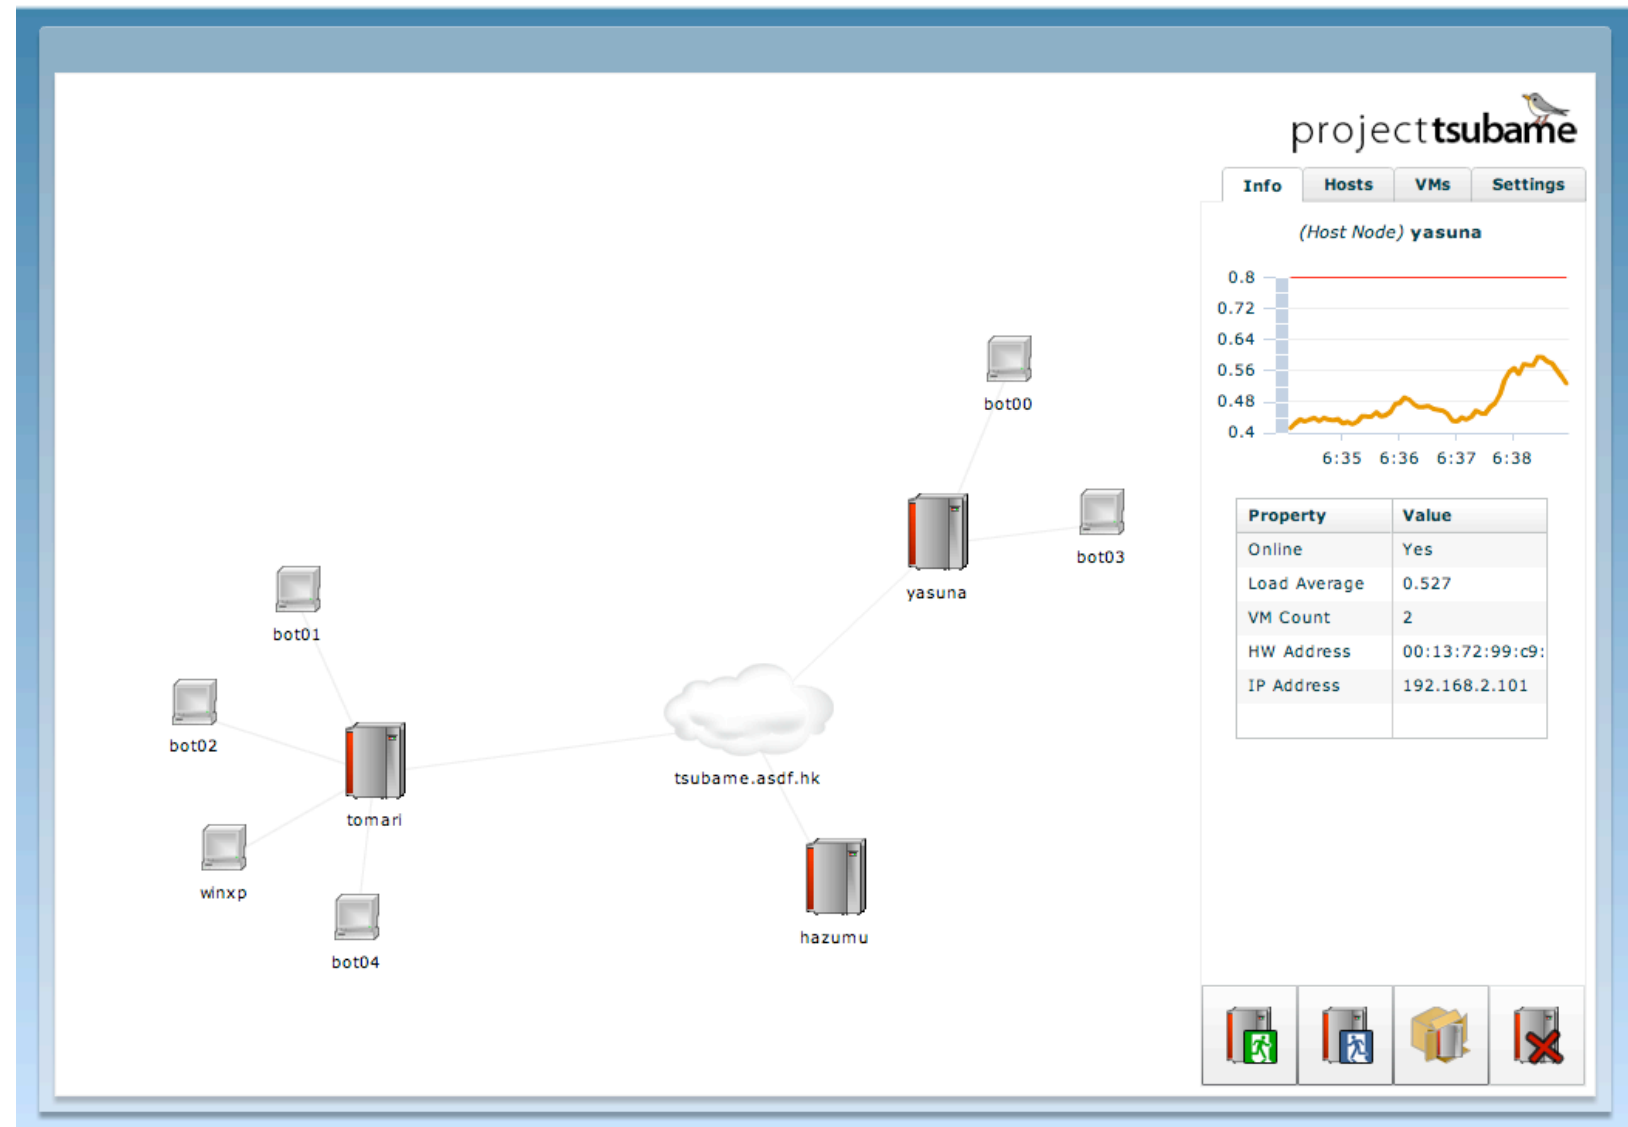
Strategy - Looking Forward with Statistics



Strategy - Looking Forward with Statistics



Strategy - Looking Forward with Statistics



Demo!

Limitations

- Migration across CPU families
 - Crash if CPU features do not match
 - Unable to recycle old hardware
- Limited OS support
 - vmxassist Real mode emulation
 - Windows XP, 2003 + Linux



Applications

- Server Consolidation
 - ✓ Availability + Manageability + Scalability
 - ✓ Increased Utilization, Decreased Downtime

- Thin-Client Model
 - ✓ Load-balanced cloud of desktops
 - ✓ Economy and Ecology
 - ✓ No upgrades = Savings
 - ✓ Computing time rental model?

Future Development

➤ High Availability Mode


- VM Memory mirroring (on 2+ hosts)
- Requires very advanced interconnect (10GbE? Myrinet? Infiniband?)

➤ Optimization of Load Balancing Algorithms

- Study typical workload patterns
- Design better algorithm

➤ Removing Limitations of Xen

- CPUID flag masking?
- Support for more operating systems?



Q&A? :)