



Peter Ruissen
Marju Jalloh



Agenda

- Research Topic
- High availability concepts
- Virtualization
- XEN Technology
- Design
- Testing XEN & Heartbeat
- Conclusion

>> Research Topic

- **Research Topic**
- High availability
- Virtualization
- XEN Technology
- Design
- Testing XEN & Heartbeat
- Conclusion



Research Topic

To research the possibilities for High Availability (HA) failover mechanisms using the XEN virtualization technology and the requirements necessary for implementation on technical level.

>> IN SCOPE



- **Research Topic**
- High availability
- Virtualization
- XEN Technology
- Design
- Testing XEN & Heartbeat
- Conclusion

In Scope:

- XEN and Virtualization technology
- High Availability HA
- Failover Services
- Potential software: DRDB, Heartbeat, LVS, Ultramonkey

>> Out Scope



- **Research Topic**
- High availability
- Virtualization
- XEN Technology
- Design
- Testing XEN & Heartbeat
- Conclusion

Out of Scope:

- Implementing other VM like UML
- Multi platform tests
- Security flaws of overall designs
- Measurements
- High performance clusters

>> High Availability

- Research Topic
- **High availability**
- Virtualization
- XEN Technology
- Design
- Testing XEN & Heartbeat
- Conclusion



What is High Availability?

- Minimal service interruption: redundancy

Failover clusters:

- Active-Active: Load balancing: routing
- Active-Passive: each node has backup
- N+1: nodes have one shared backup
- N(ode)+M(standby)
- N+N: Active-Active and N+M

>> Service Availability

- Research Topic
- **High availability**
- Virtualization
- XEN Technology
- Design
- Testing XEN & Heartbeat
- Conclusion



Service Availability

- Controlled outage: reboot, shutdown (UPS, hardware upgrade, software upgrade etc.)
- Uncontrolled failure: network, power, hardware failures, OS failures: application failures, kernel panic, stale processes

>> Requirements HA

- Research Topic
- **High availability**
- Virtualization
- XEN Technology
- Design
- Testing XEN & Heartbeat
- Conclusion



Requirements HA

- Redundancy
- Monitoring for a failed node
- Automatic transfer of resource ownership
- Private interconnect (Serial/Ethernet Cable)
- Cluster configuration (Standby/Takeover)
- HA Storage: Distributed, SAN, GFS

>> HA under Linux

- Research Topic
- **High availability**
- Virtualization
- XEN Technology
- Design
- Testing XEN & Heartbeat
- Conclusion



HA under Linux

Heartbeat

- Heartbeat Protocol (+SHA1)
- Logical resource manager (LRM)
- Cluster Information Base (CIB)
- Stonith daemon (Single shot in the head)

LVS

- Director

>> Virtualization

- Research Topic
- High availability
- **Virtualization**
- XEN Technology
- Design
- Testing XEN & Heartbeat
- Conclusion



Virtualization

Definition: abstraction of computer resources

a technique for hiding the physical characteristics of computing resources from the way in which other systems, applications, or end users interact with those resources.

- Research Topic
- High availability
- **Virtualization**
- XEN Technology
- Design
- Testing XEN & Heartbeat
- Conclusion



Advantages of virtualization

- Transparency – hardware independent
- Efficiency – efficient use of hardware resources
- Flexibility – to move one system to another
- Simplifies application stack – largely removing hardware and drivers from the equation
- Security – providing isolation

>> Disadvantages

- Research Topic
- High availability
- **Virtualization**
- XEN Technology
- Design
- Testing XEN & Heartbeat
- Conclusion



Disadvantages of virtualization

- Increased complexity
- CPU and memory overhead

>> Forms of VT

- Research Topic
- High availability
- **Virtualization**
- XEN Technology
- Design
- Testing XEN & Heartbeat
- Conclusion



Forms of VT

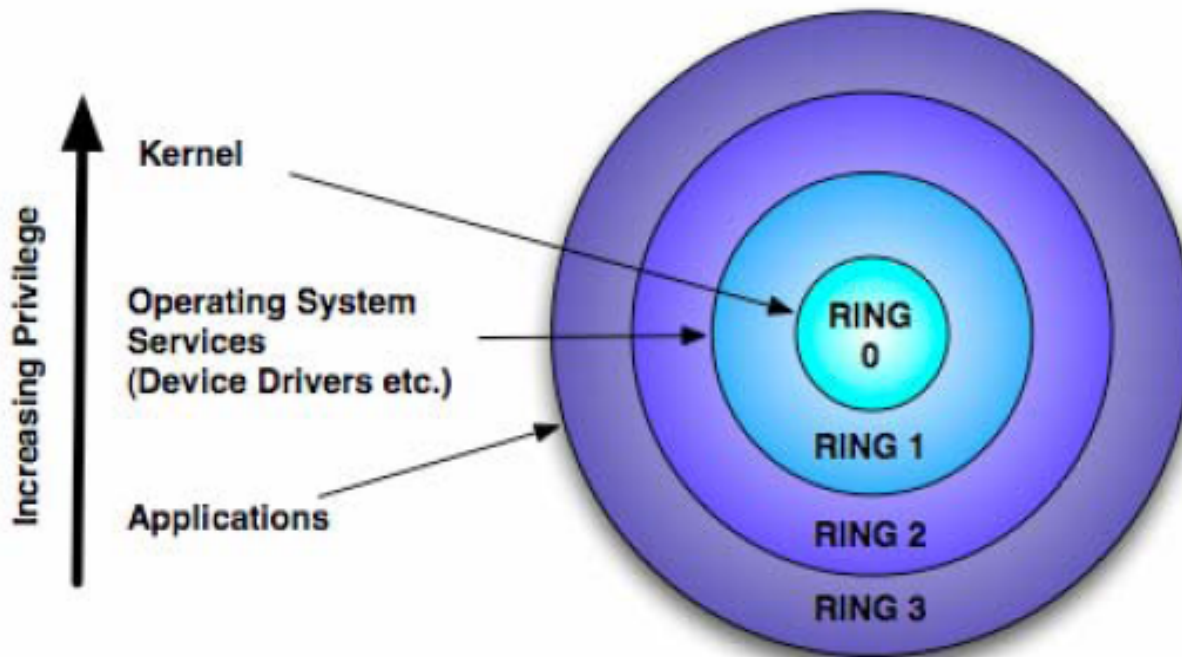
- Emulation
- Full virtualization or native virtualization
- Para-virtualization
- OS virtualization
- Hardware virtualization
- Application Virtualization

>> XEN Architecture

- Research Topic
- High availability
- **Virtualization**
- XEN Technology
- Design
- Testing XEN & Heartbeat
- Conclusion



▪ XEN Architecture

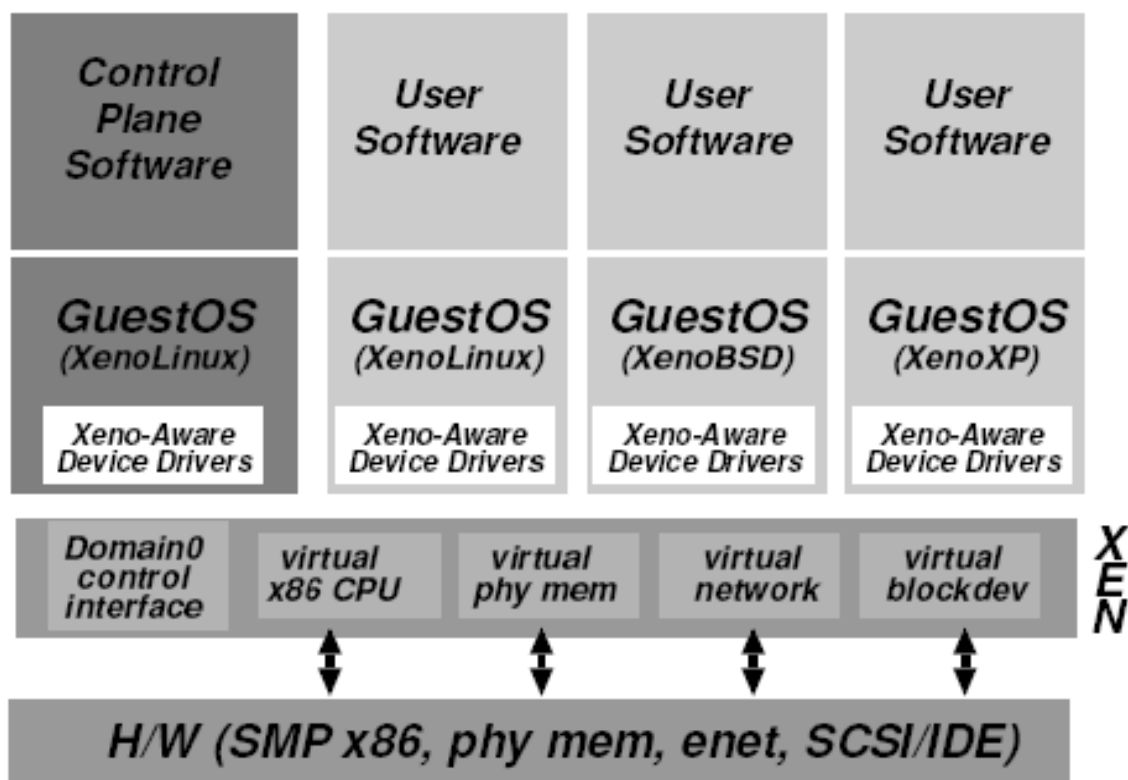


X86-Ring protection



- Research Topic
- High availability
- Virtualization
- **XEN Technology**
- Design
- Testing XEN & Heartbeat
- Conclusion

▪ XEN Architecture



Source: Xen and the art of virtualization

- Research Topic
- High availability
- Virtualization
- **XEN Technology**
- Design
- Testing XEN & Heartbeat
- Conclusion



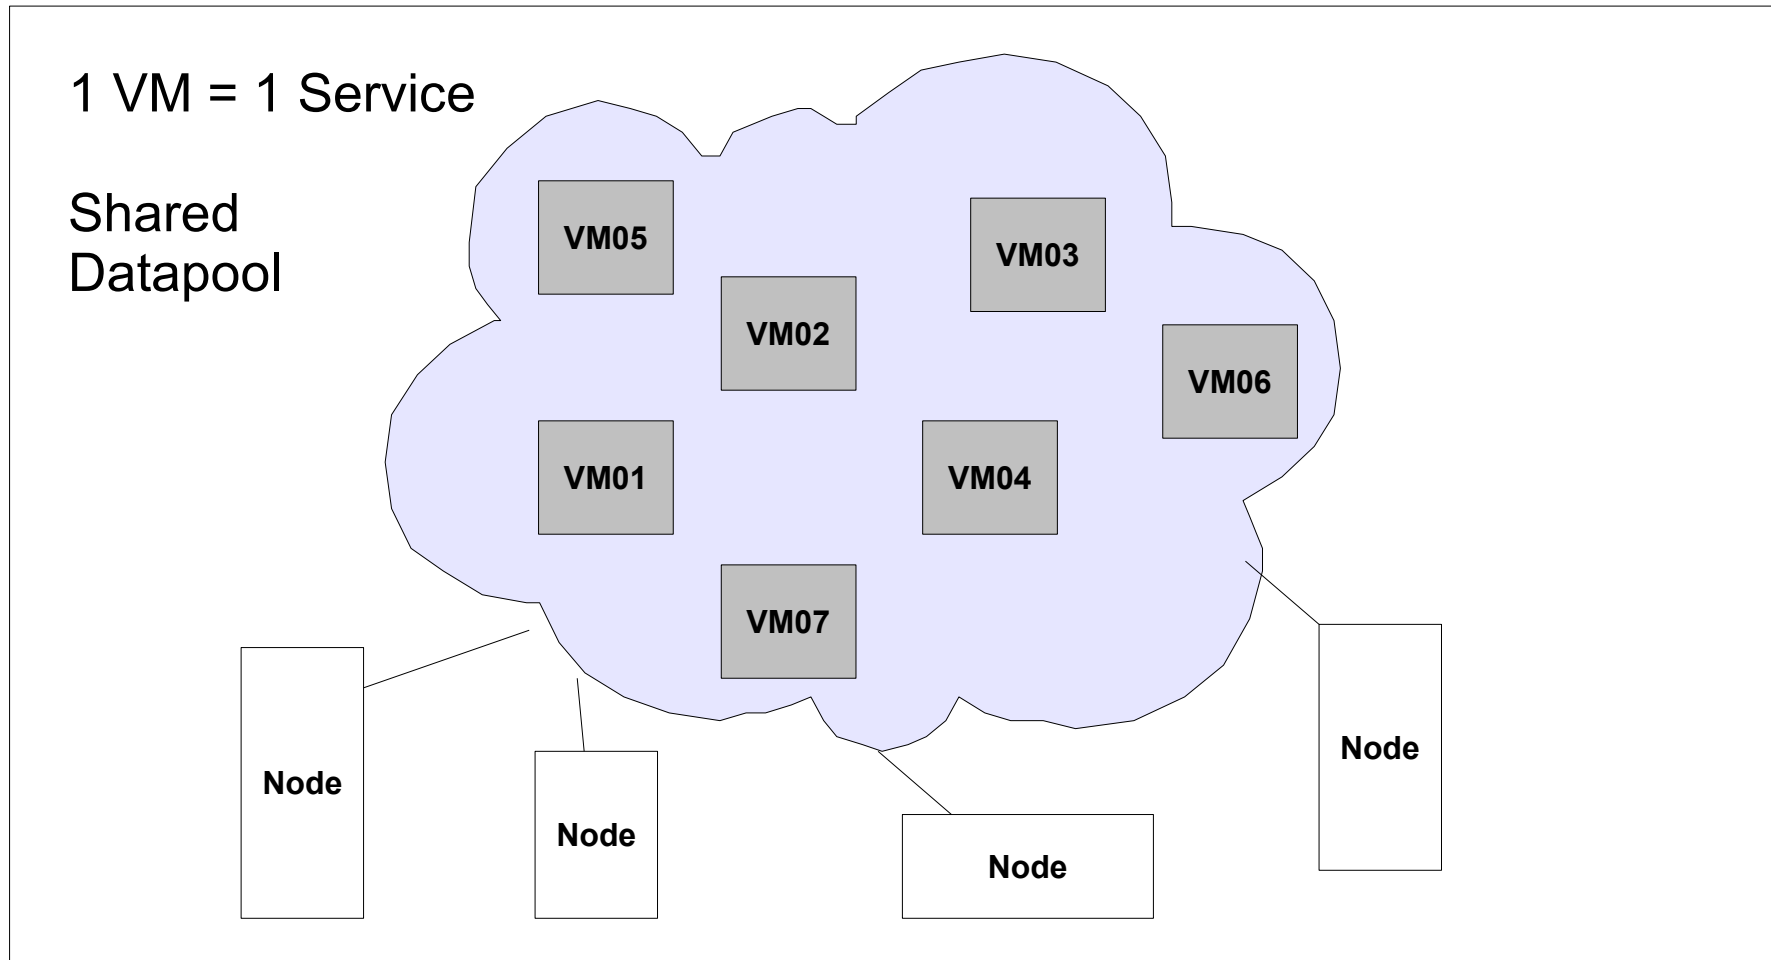
Advantages XEN

- Open source
- Secure isolation
- High performance HA
- Migration

Disadvantages

- Modify guest OS: or VT Support / SVM support
- Not very mature

- Research Topic
- High availability
- Virtualization
- XEN Technology
- **Design**
- Testing XEN & Heartbeat
- Conclusion



>> Requirements HA Design

- Research Topic
- High availability
- Virtualization
- XEN Technology
- **Design**
- Testing XEN & Heartbeat
- Conclusion



Requirements HA Design

- Heartbeat extension which controls XEN
- Resource management system which monitors resources from physical hosts.
- Control STONITH device
- Build in safety measures
- Use HA data pool


>> Procedures Outages

- Research Topic
- High availability
- Virtualization
- XEN Technology
- Design
- **Testing XEN & Heartbeat**
- Conclusion




Procedures during:

▪ 1 Uncontrolled outage



▪ Cold Failover

2 Controlled outage



Live Migration!

>> Test environment

- Research Topic
- High availability
- Virtualization
- XEN Technology
- Design
- **Testing XEN & Heartbeat**
- Conclusion



Test environment

▪ Software:

Ubuntu 6.10, XEN 3, Heartbeat 2, Apache 2.

Hardware:


- Intel(R)Pentium D CPU 3.00.GHz
- 2GB Memory

- DomU 128MB Memory

- Research Topic
- High availability
- Virtualization
- XEN Technology
- Design
- **Testing XEN & Heartbeat**
- Conclusion



Scenario 1




Heartbeat communication between physical Node0 and Node1
Service downtime 6-8 sec

- Research Topic
- High availability
- Virtualization
- XEN Technology
- Design
- **Testing XEN & Heartbeat**
- Conclusion



Scenario 2




Heartbeat communication between VM01 and physical host node1.
Service downtime 6-8 sec

- Research Topic
- High availability
- Virtualization
- XEN Technology
- Design
- **Testing XEN & Heartbeat**
- Conclusion



Scenario 3



Heartbeat communication between VM01 and VM02.
Service downtime 1-2 sec

- Research Topic
- High availability
- Virtualization
- XEN Technology
- Design
- Testing XEN & Heartbeat
- **Conclusion**



Conclusion

- XEN & Heartbeat can be used to realize a flexible, reliable and efficient HA environment.
- Extra functionality needed in cluster manager (Heartbeat)

- Research Topic
- High availability
- Virtualization
- XEN Technology
- Design
- Testing XEN & Heartbeat
- **Conclusion**



Heartbeat ADD-ON Requirements:

- Resource management system which monitors resources from physical hosts.
- Heartbeat extension which controls XEN
- Control STONITH device
- Build in safety measures
- Use HA data pool