

# MySQL Cluster

## Practical demo

### using Red Hat Cluster Suite

# Clustering

- Achieving:
  - High Availability (HA) & Load Balancing (LB)
- Availability:
  - 99 % - 87 hours outage a year
  - 99.9 % - 8.7 hours outage a year
  - 99.99 % - 52 minutes a year
  - 99.99 % - 5 minutes a year

# Do I need a cluster?

- Can I afford being offline for 8.7 hours a year?
- Can I afford being offline for 5 minutes a year?

or

- What is my loss if I am offline for 8.7 hours a year?

vs.

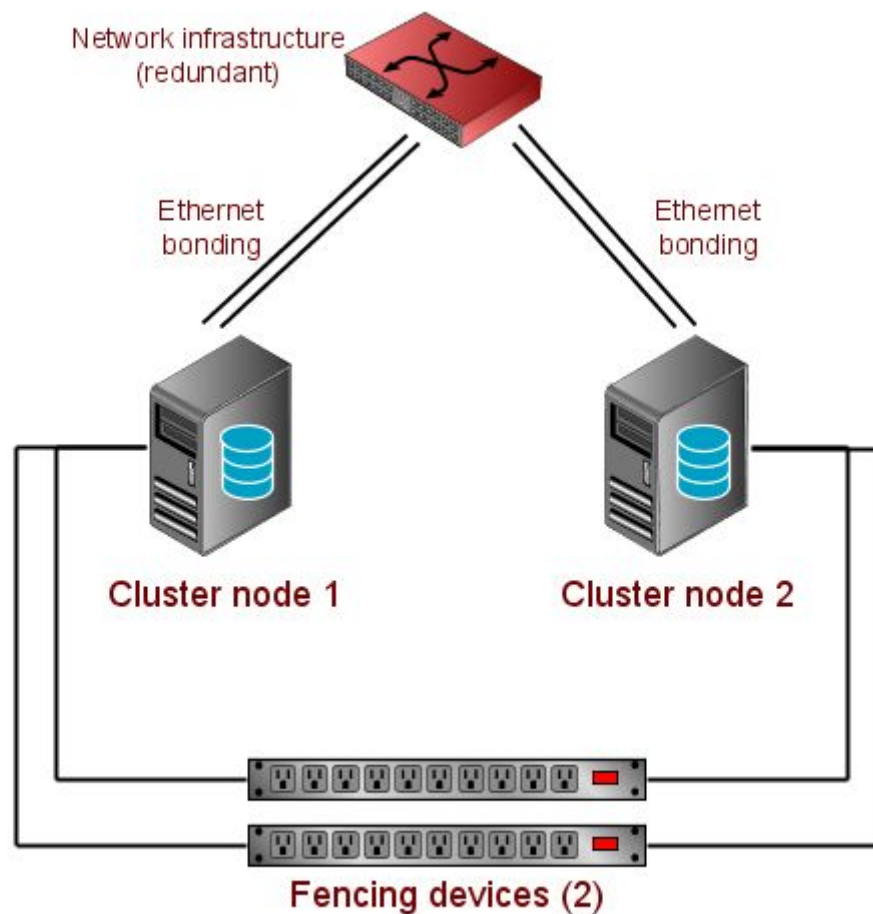
- What does the cluster cost?

Clustering is about...

# Eliminating **any** SPOF


SPOF – Single Point Of Failure. Any component in the cluster must be doubled!

# Our Cluster demo



create and share your own diagrams at [gliffy.com](http://gliffy.com)

# Our Cluster demo



create and share your own diagrams at [gliffy.com](http://gliffy.com)

# How it works

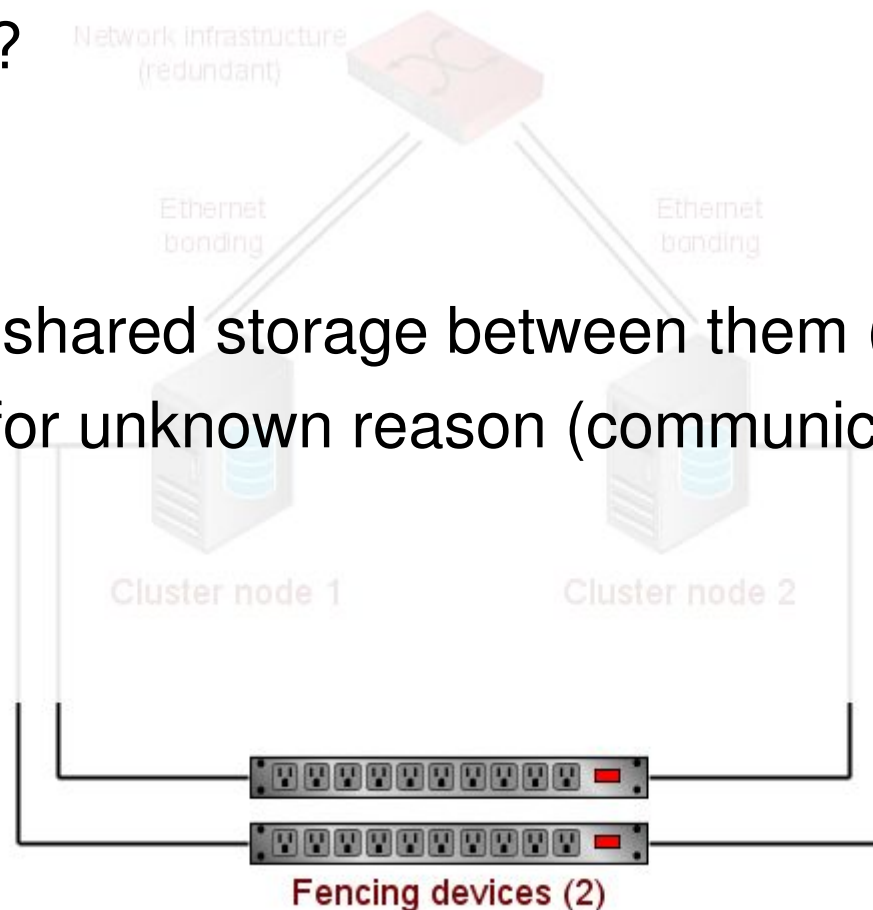
- Basically: sending heartbeats to other nodes and determining if they are alive and OK
- Taking care of nodes (fencing) and services (failover) if there is a failure

# Fencing

- Sorry? Fenc...what?

- Scenario:

- 1)  $N$  nodes cluster, shared storage between them (SAN, NAS)
- 2) One goes faulty for unknown reason (communication issue!)
- 3) What happens?

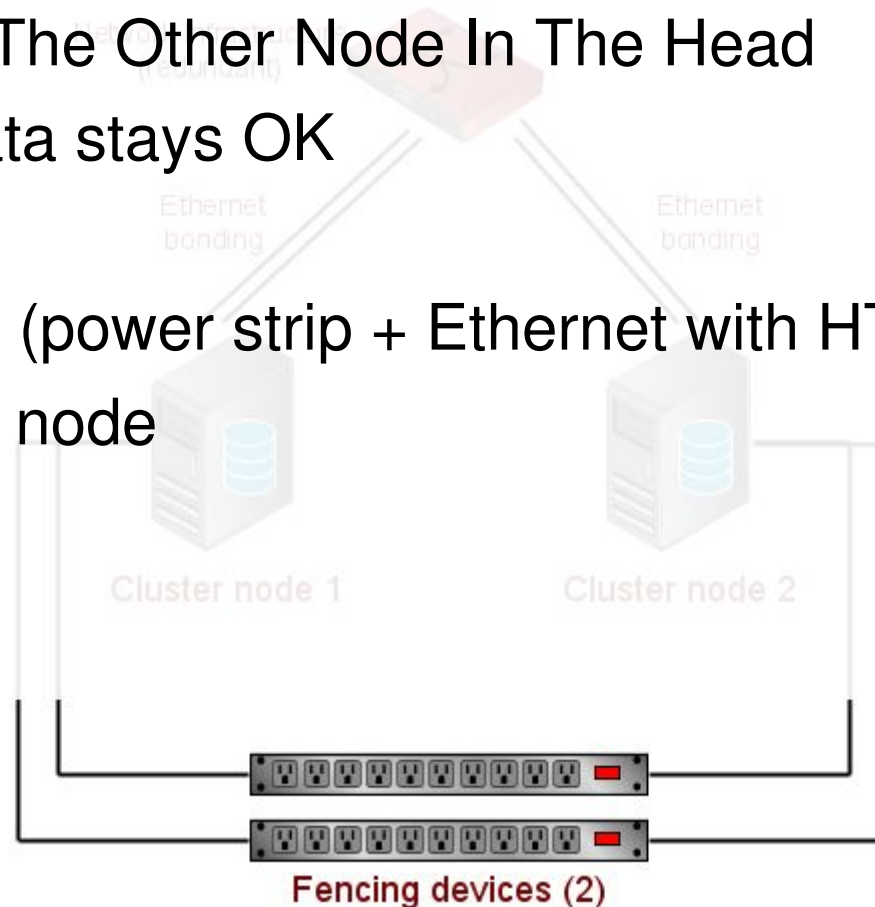


create and share your own diagrams at [gliffy.com](http://gliffy.com)



# Fencing

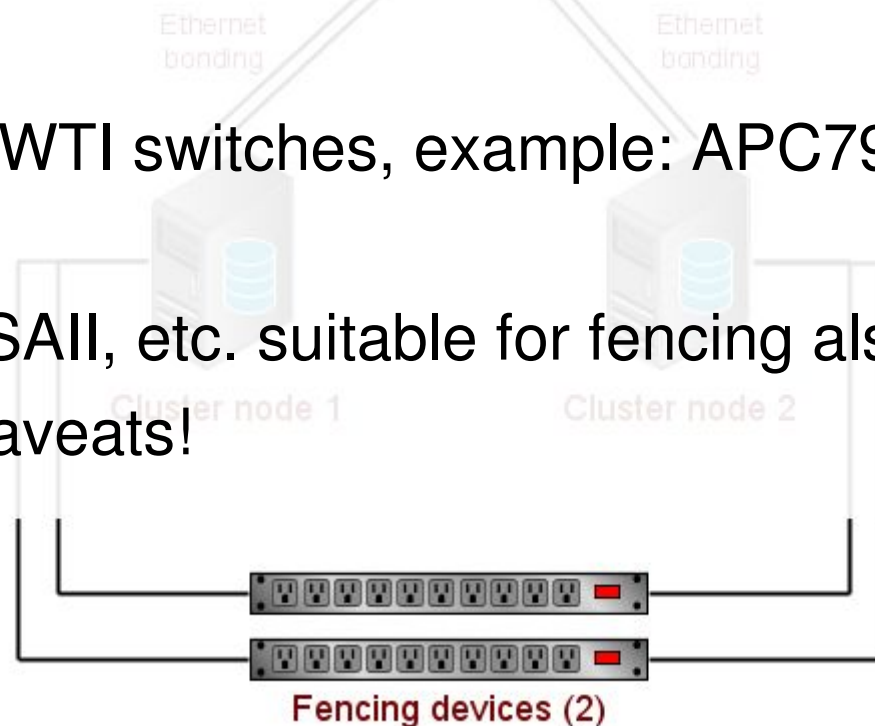
- STONITH – Shoot The Other Node In The Head  
= Ensuring your data stays OK
- Using power switch (power strip + Ethernet with HTTP interface)
- Power off the faulty node



create and share your own diagrams at [gliffy.com](http://gliffy.com)

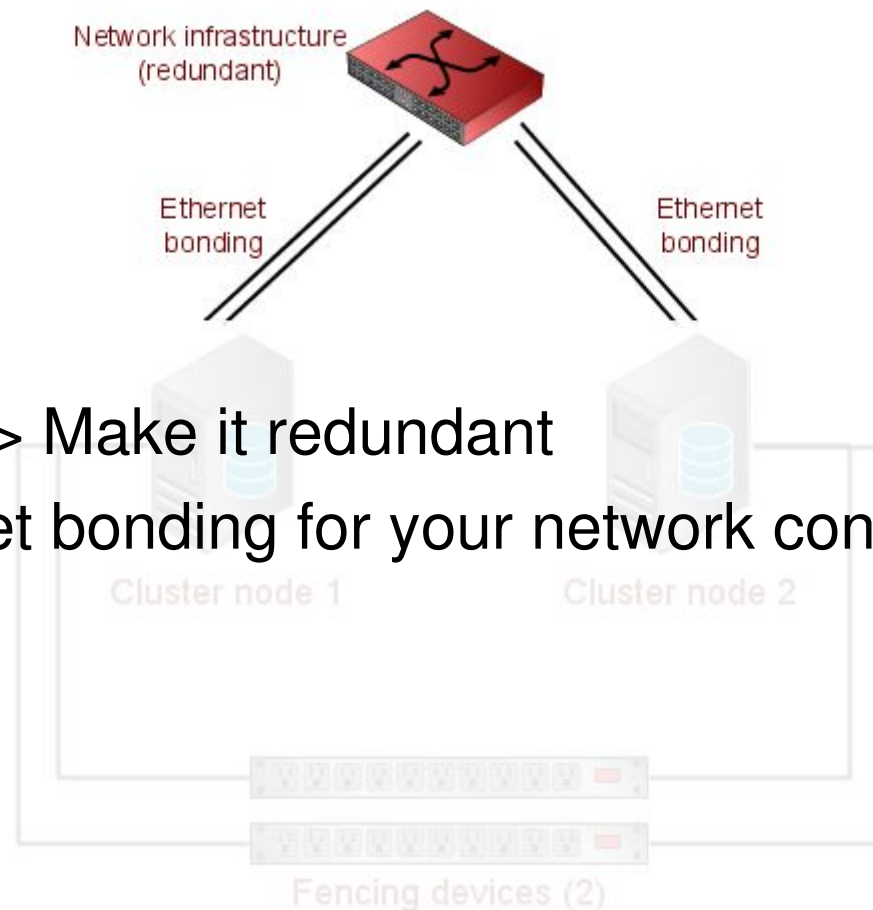
# Fencing

- **Essential!** Not only for data, also virtual IPs
- Fencing device = SPOF => You need 2!
- Best devices: APC/WTI switches, example: APC7920
- iLO, DRAC, IBM RSAII, etc. suitable for fencing also
- However, special caveats!



create and share your own diagrams at [gliffy.com](http://gliffy.com)

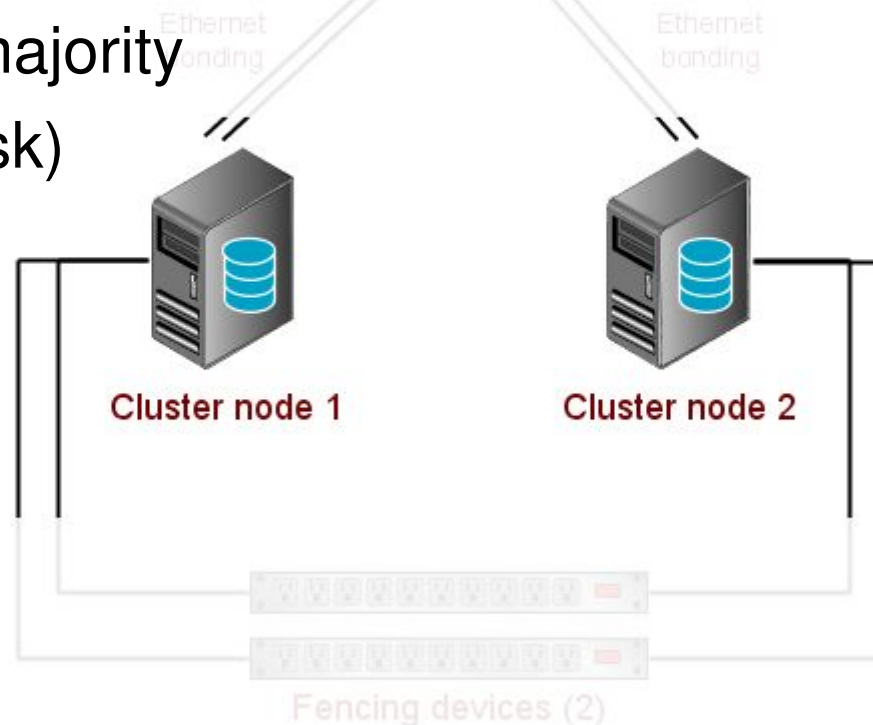
# Networking



create and share your own diagrams at [gliffy.com](http://gliffy.com)

# Cluster nodes

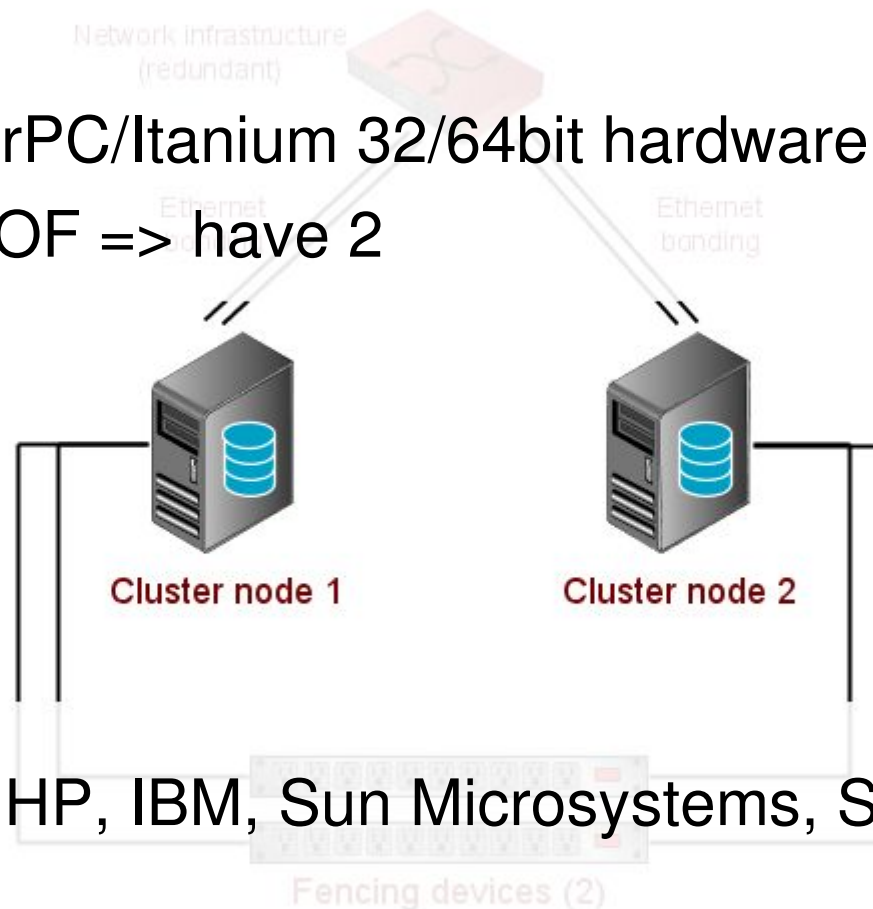
- Server = SPOF => at least 2 (this makes it a cluster!)
- Better to have at least 3, there is a voting system  
2 from 3 means a majority
- Qdisk (Google: qdisk)
- Running services



create and share your own diagrams at [gliffy.com](http://gliffy.com)

# Hardware

- Desktop PCs: NO
- HW: Any x86/PowerPC/Itanium 32/64bit hardware
- Power supply = SPOF => have 2

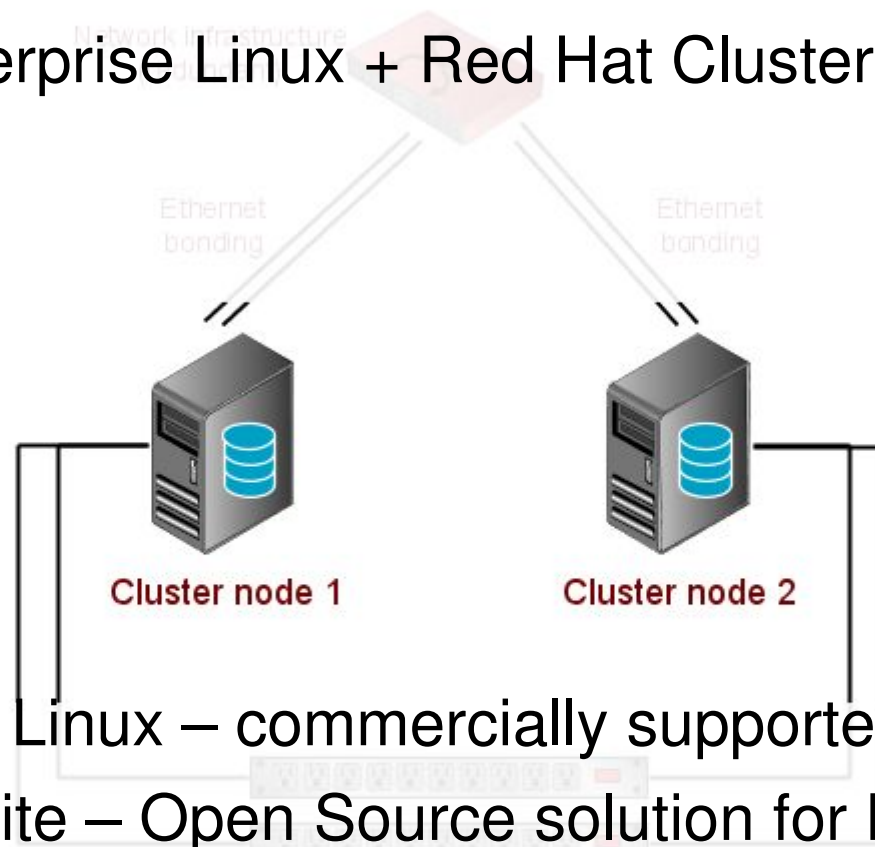


- Major brands: Dell, HP, IBM, Sun Microsystems, SuperMicro

create and share your own diagrams at [gliffy.com](http://gliffy.com)

# Software

- Using Red Hat Enterprise Linux + Red Hat Cluster Suite




- Red Hat Enterprise Linux – commercially supported Linux
- Red Hat Cluster Suite – Open Source solution for HA/LB commercially supported

# Software

- Operating system – Red Hat Enterprise Linux
- *Do I have to buy RHEL for this cluster?*  
*Well...no, but this is HA right?*
- Cluster software – Red Hat Cluster Suite
- *Do I have to buy Cluster suite for this cluster? ...*
- OK, you can use CentOS, Debian, Gentoo, ~~Fedora~~
- Almost every Linux distribution

# Our Cluster demo



create and share your own diagrams at [gliffy.com](http://gliffy.com)



# Simple budget

- This is just a hardware list, no work included
- This is a lowcost variant
- Rough: many things to be considered

Item	Count	Price / pcs	Price
Dell PowerEdge 1950	2	\$1,600	\$3,200
APC 7920	2	\$650	\$1,300
RHEL OS	2	\$400	\$800 (1 year)
RHEL Cluster	2	\$500	\$1,000 (1 year)
		<b>Sum</b>	<b>\$6,300</b>

# Demonstration

- MySQL Master/Slave replication cluster
- 2 IBM servers, fencing, network (simple)

# END

# Questions?

**Jakub Suchý**

e-mail: [jakub.suchy@drupal.cz](mailto:jakub.suchy@drupal.cz)

<http://www.drupal.cz>