

Highly available services on virtual hosts


Stefano Stalio
Computing and Network Service
Laboratori Nazionali del Gran Sasso

stefano.stalio@lngs.infn.it

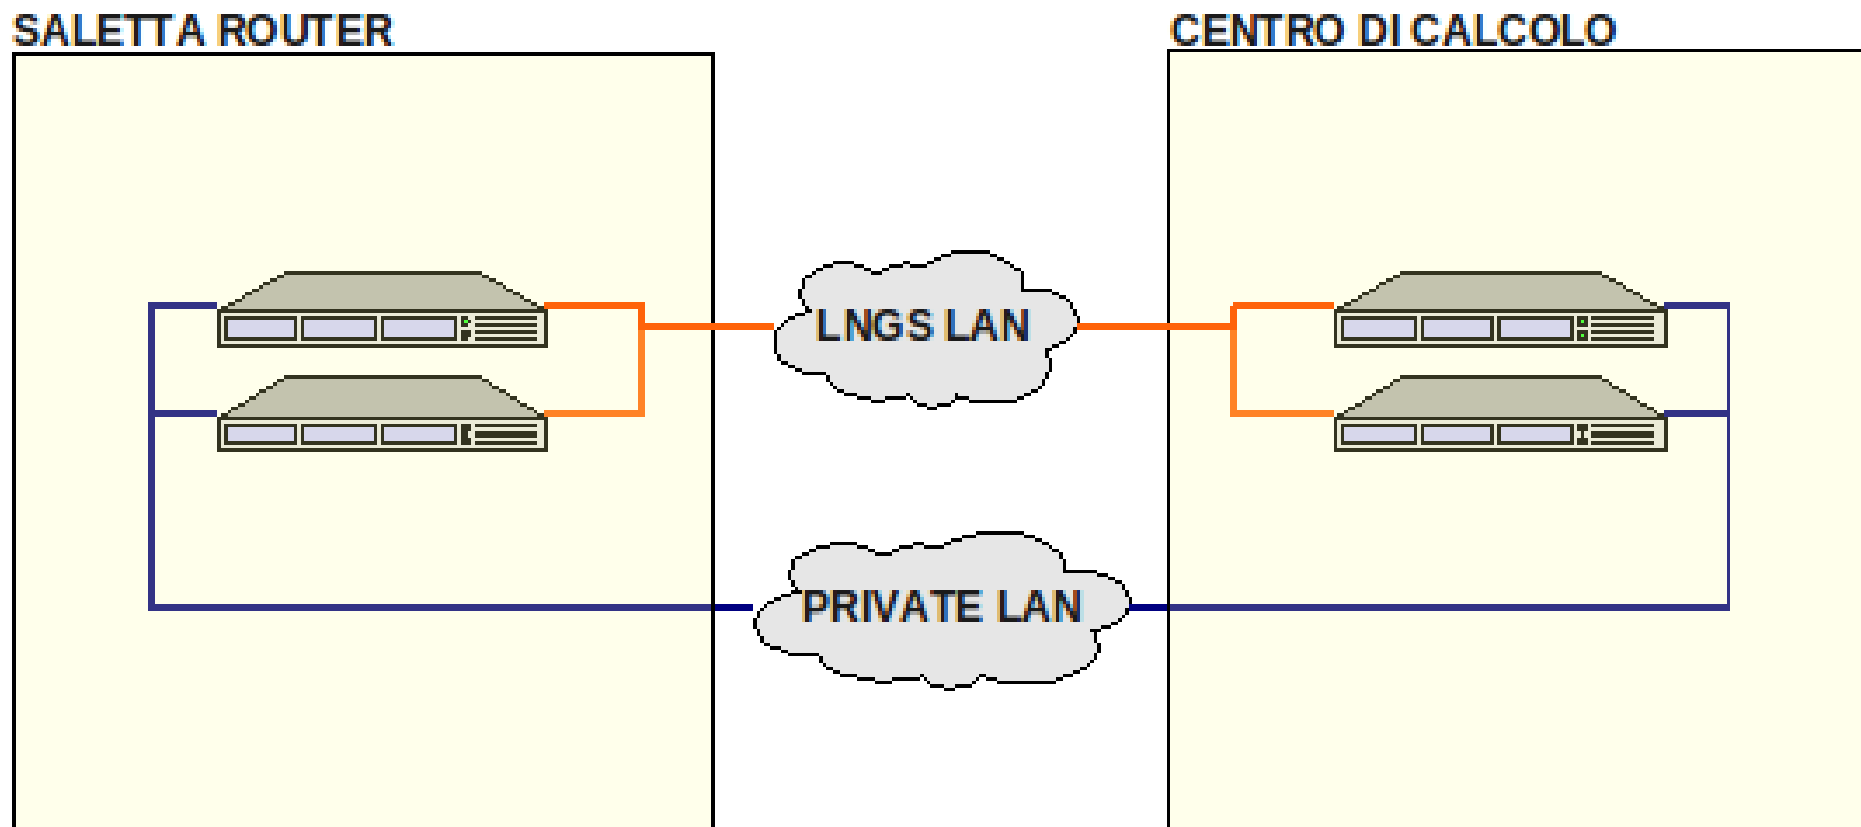


LNGS – Computing infrastructure

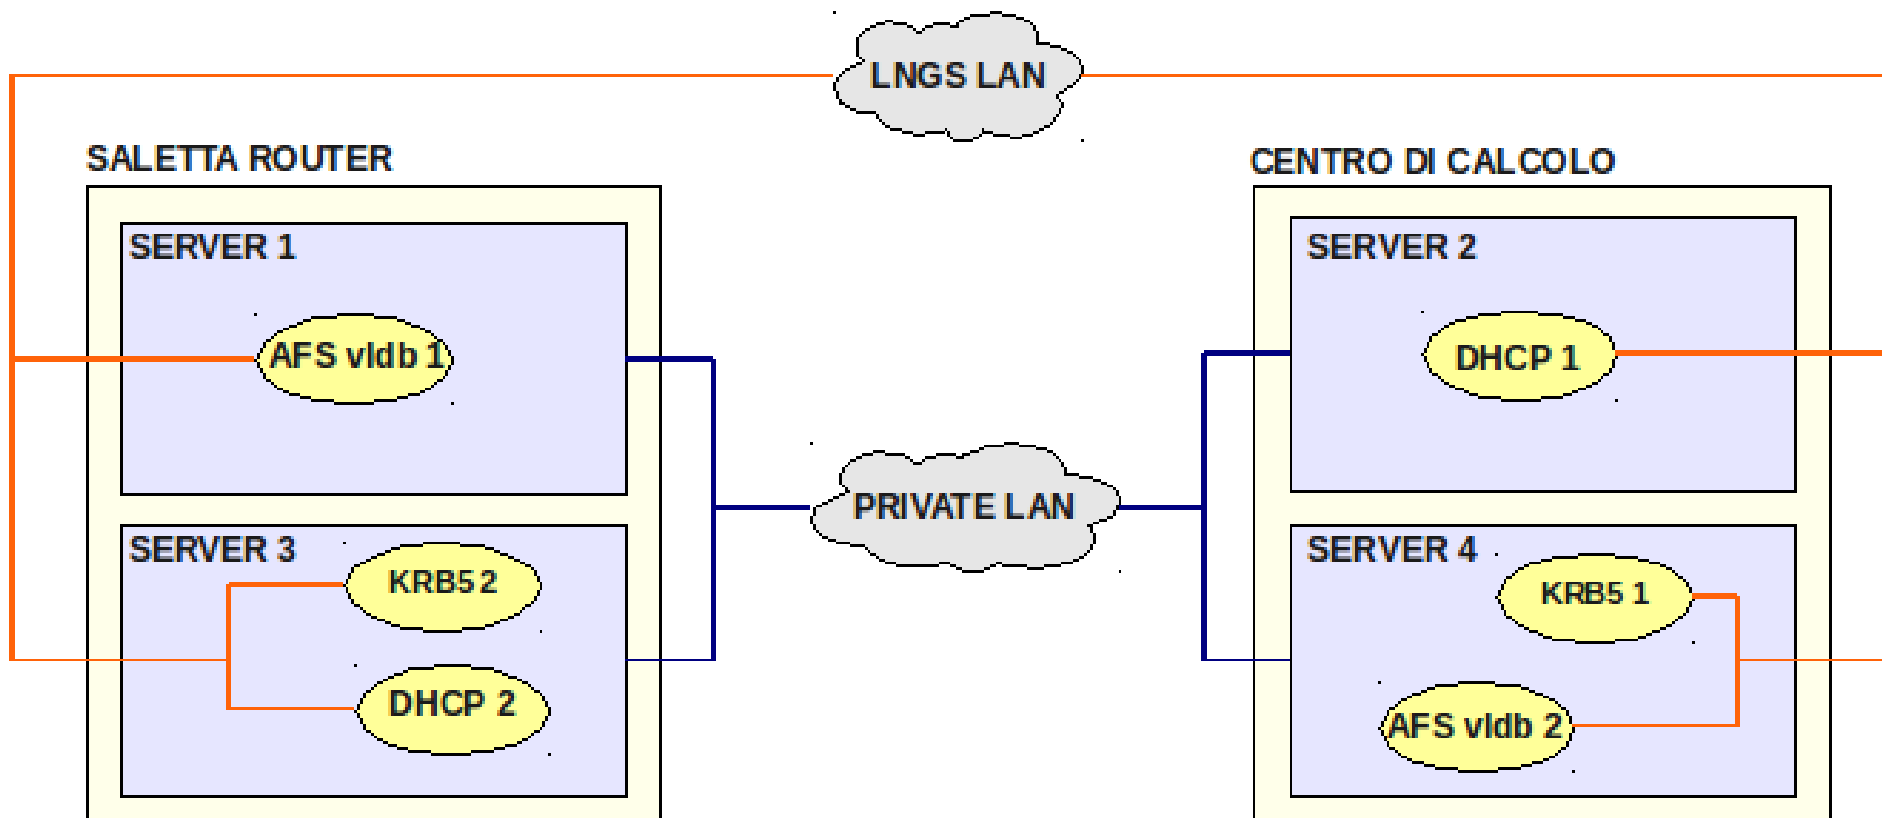
- LAN PRIVATA
- LAN LNGS



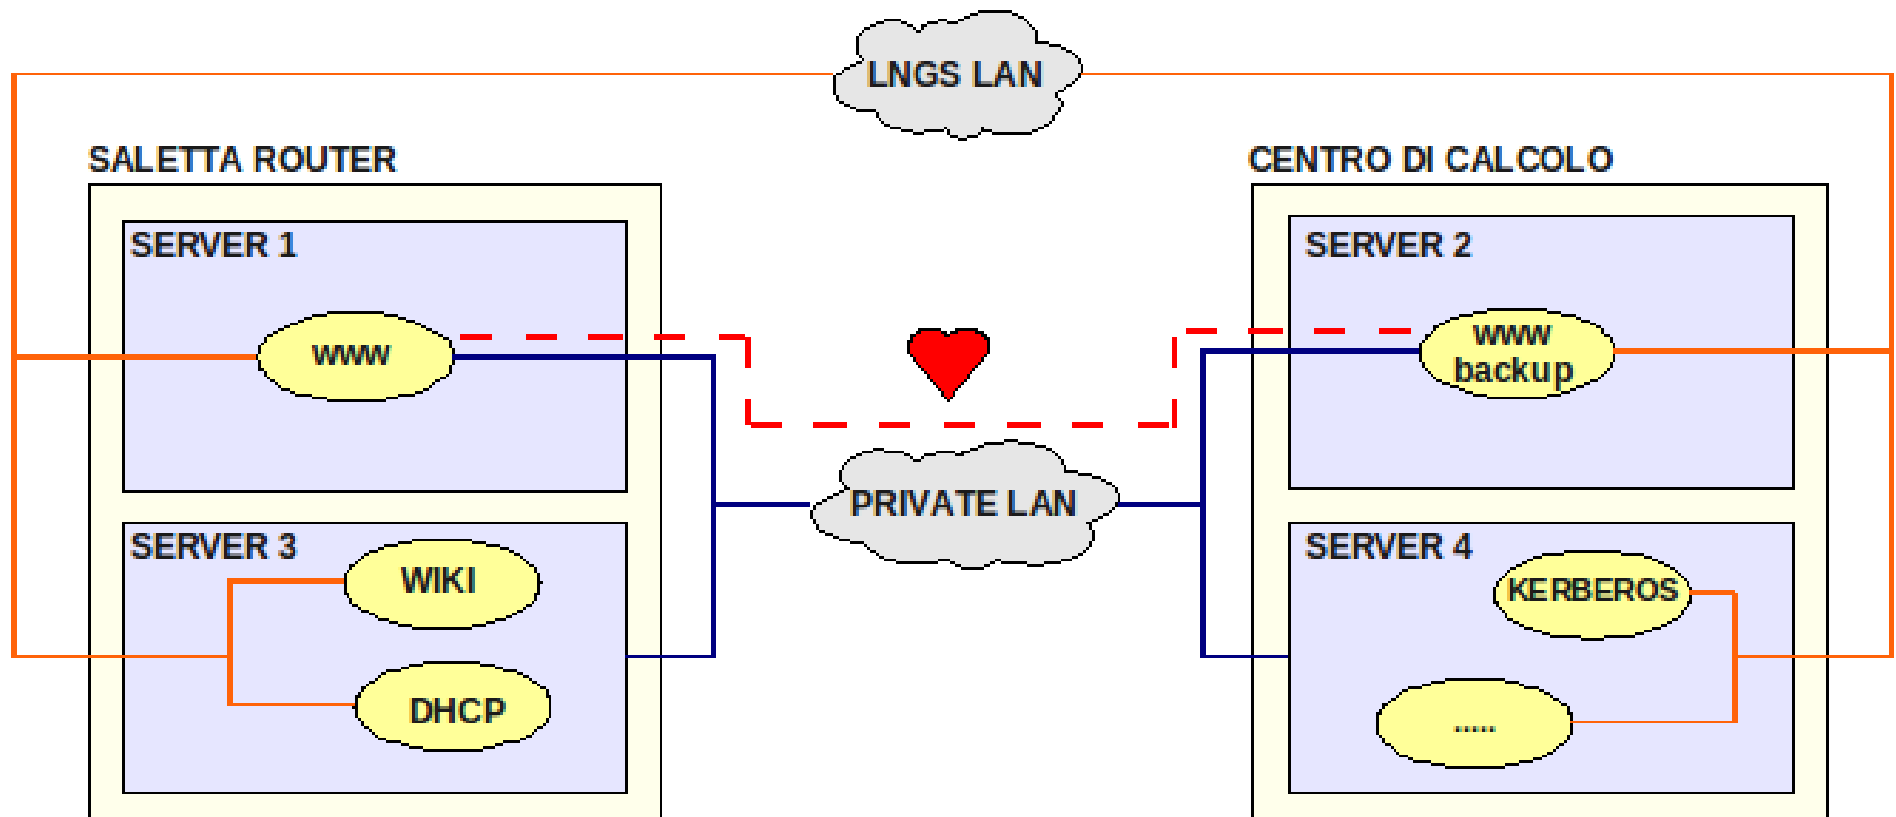
LNGS – The host virtualization cluster



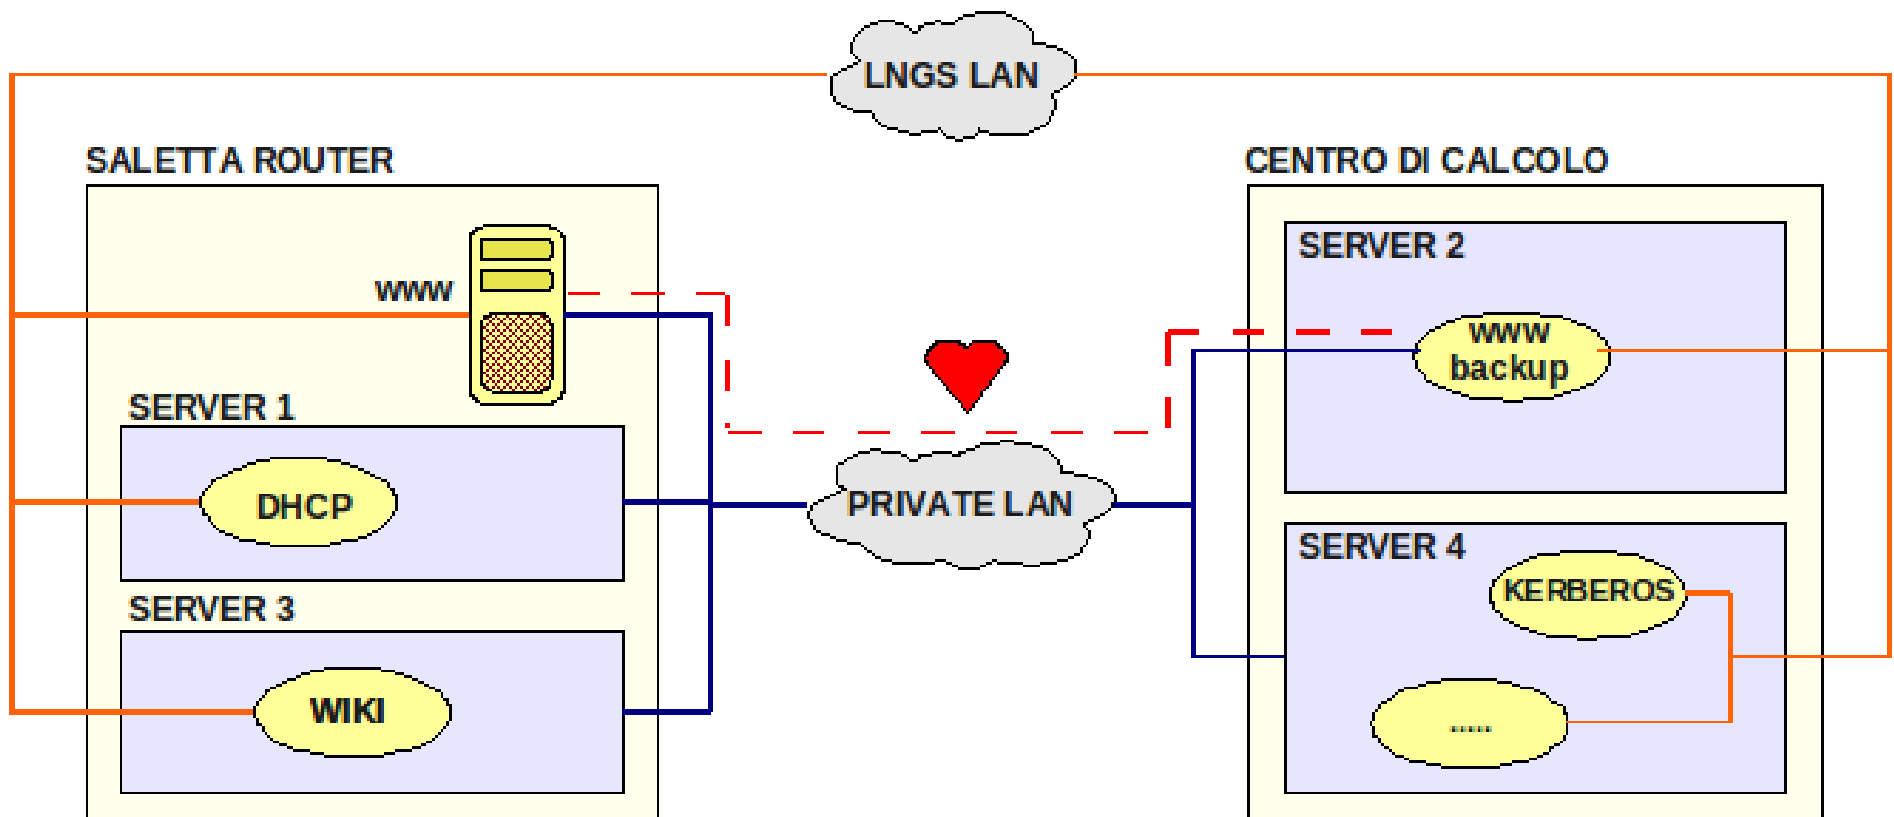
Redundant servers for critical services



Linux Heartbeat on virtual nodes



Linux Heartbeat on virtual nodes




Linux Heartbeat on virtual nodes

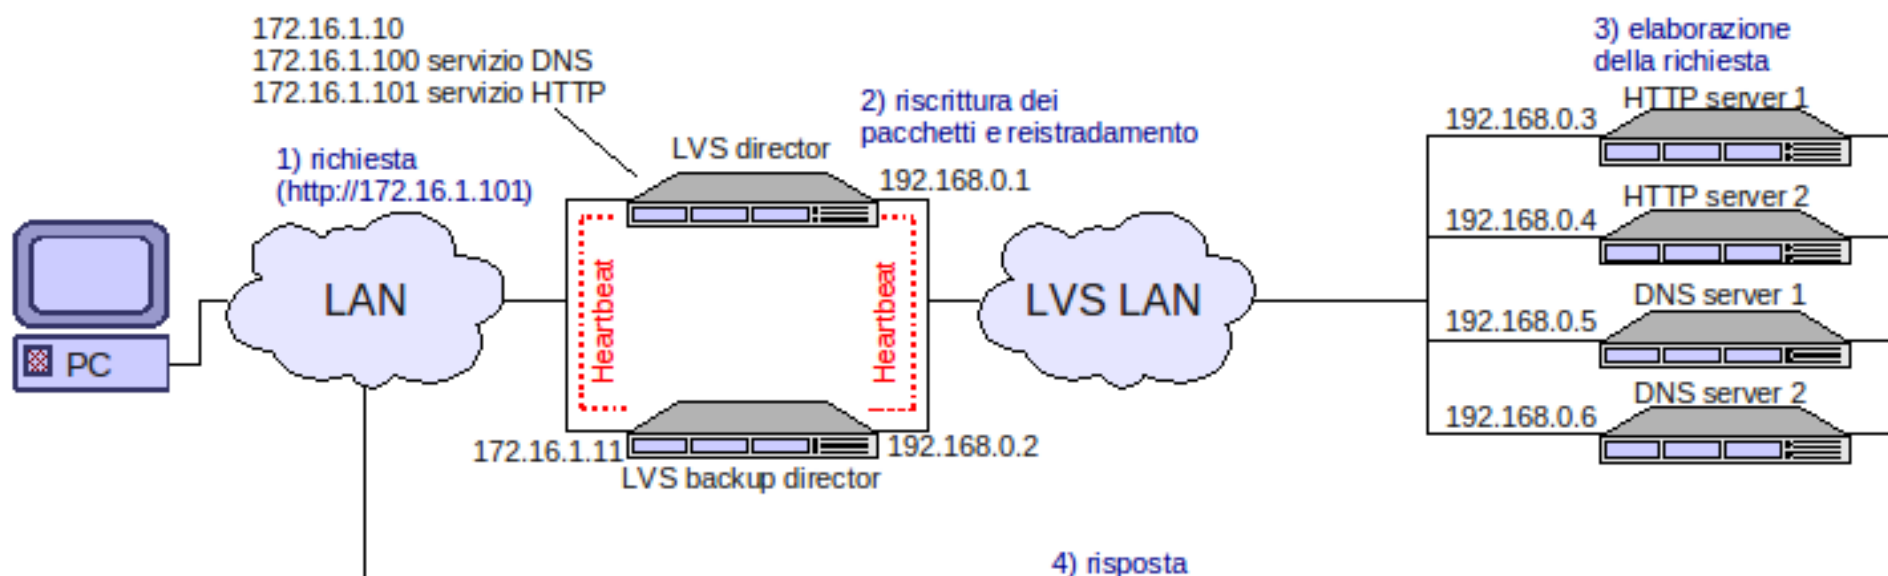
Real servers as stonith devices?



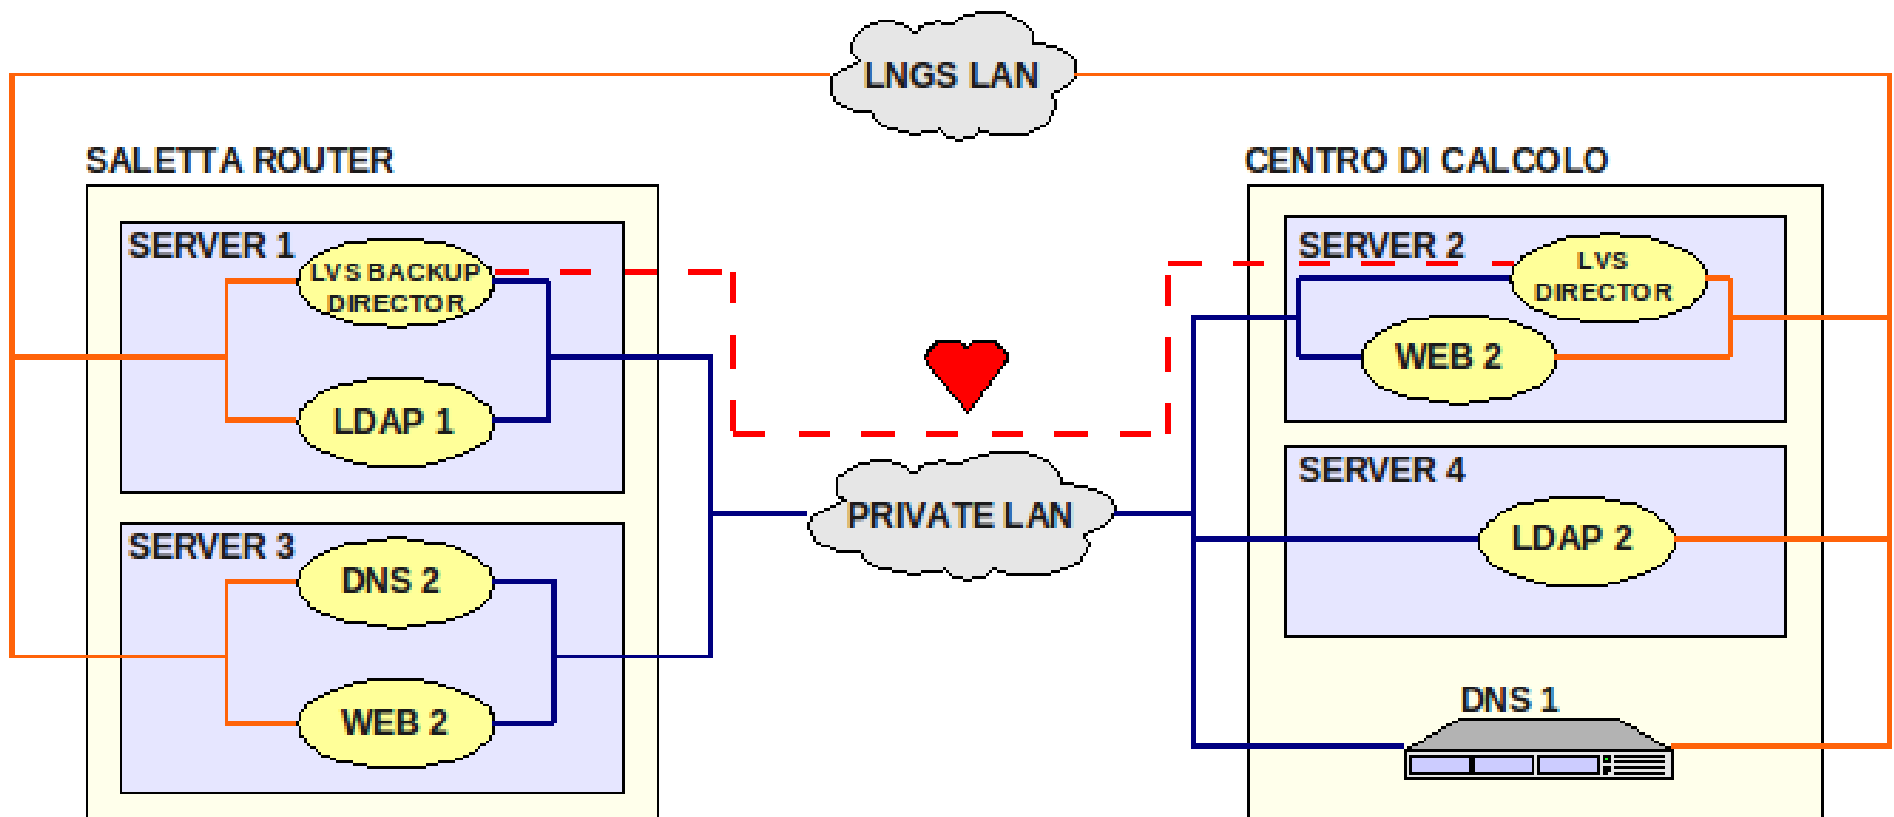
LVS cluster architecture



LVS cluster architecture



LVS on virtual nodes



Highly available services on the host virtualization cluster

| <u>Service</u> | <u>HA cluster type</u> | <u>Host type</u> |
|----------------|------------------------|------------------|
| DNS | LVS | virtual/real |
| LDAP | LVS | virtual |
| Web proxy | LVS | virtual |
| HTTP | LVS | virtual |
| HTTP | Linux Heartbeat | virtual/real |
| Wireless gw | Linux Heartbeat | virtual |
| DHCP | | virtual/real |
| KRB5/AFS VLDB | | virtual |

Data synchronization on clustered servers

- ✓ **database replica**
AFS, Kerberos, DNS, Ldap, MySQL
- ✓ **shared file system (AFS)**
web data
- ✓ **scp or rsync**
configuration files



Backup

- ✓ Proxmox appliances do not need a backup. Easily recreated.
- ✓ Backup of /etc on real servers. Mainly for the configuration files of OpenVZ virtual nodes.
- ✓ Daily dump of virtual hosts (from LVM2 snapshots) for quick recovery.
- ✓ “Traditional” full/incremental backup policy for virtual hosts (backup client running on the virtual host).

What are we protected from?

- ✓ Virtual server crash.
- ✓ Service crash on a virtual server (LVS).
- ✓ LVS director crash.
- ✓ Real server crash. Two real servers crashing if they are in the same room.
- ✓ Power loss in a rack.
- ✓ Power loss in one computer room. In some cases.
- ✓ Network branch unreachable.

What can we perform w/o service discontinuity?

- ✓ Maintenance on the computing and network infrastructure.
- ✓ Hardware maintenance on network devices or real servers.
- ✓ Software maintenance on real servers or virtual hosts.



What we are not protected from

- ✓ **Intermittent problems.**

File system corruption or hard disk problems on the real server.

- ✓ **Wrong HA cluster configurations.**

The one on the RedHat LVS manual for DNS service monitoring.

- ✓ **Non optimal client configuration.**

Clients that only know one Kerberos or AFS server will fail when the server fails.