

Business Application High Availability Utilizing VPLS

Twain McCulley
Technical Architect
Edward Jones

Geographic Availability Issues

⌘ File Services

- ☒ Failure of systems is dependent on hosts remapping services to different IP addresses, this often requires a reboot for NFS or CIFS

⌘ Network routes must be propagated across the network during disaster recovery.

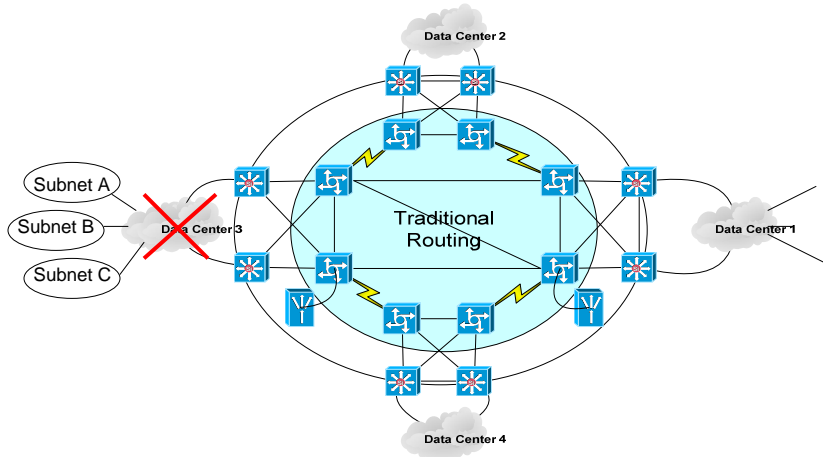
- ☒ Multisite recoverability does not scale well.

⌘ Application Services

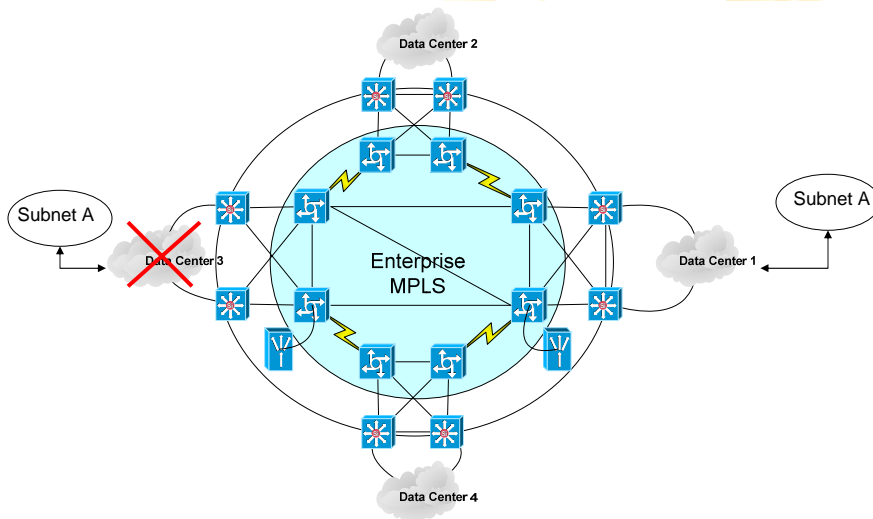
- ☒ Single mastered services (CIFS, NFS, NIS, etc)
- ☒ State-full failover of running processes nearly impossible, distance/latency difficulties
- ☒ State-full firewall abilities, limited in most vendor implementations

Conventional Fixes

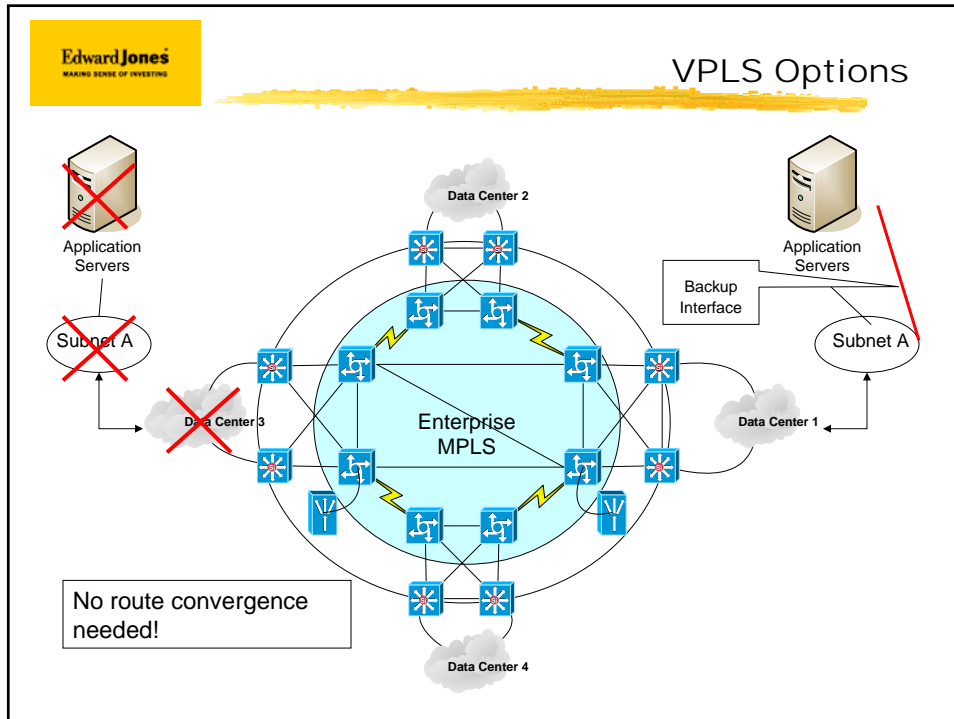
⌘ Route Policies and Subnet migration



VPLS Options



VPLS Options

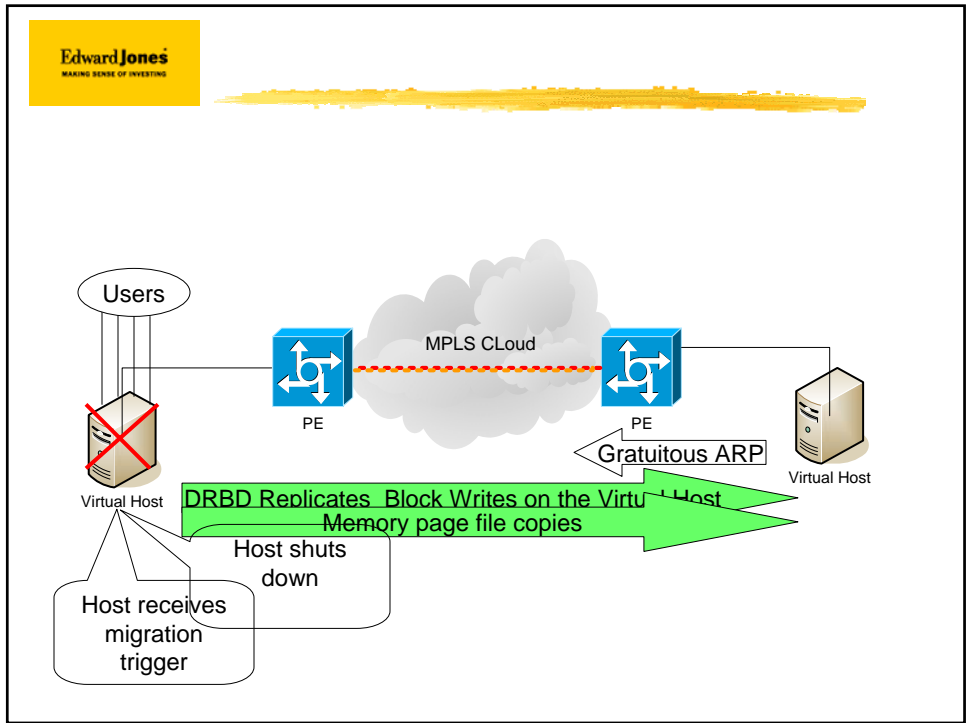


Server Migration for Maintenance Windows

⌘ Xen and VMware machine State Migration

⌘ Options

- ☒ 1, local disk, no migration
- ☒ 2, NAS, DASD, some migration, dependent on replication scheme
- ☒ 3, DRBD, Distributed replicated block device (<http://www.drbd.org>)
 - ☒ IP block level copy
 - ☒ 1 to 1 replication, no multipoint
- ☒ 4, SAN, iSCSI with DRBD,
 - ☒ SRDF works at the device layer, DRBD at the kernel
- ☒ Page file copies @ layer 3, drbd @ layer 3
- ☒ Works at Layer 3, but to not lose sessions must assume IP/MAC address of original system.



Problems

- ⌘ Routing protocols need to be disabled to avoid short-circuiting normal paths.
- ⌘ Multicast problematic, rate limit if possible, avoid running multicast routing problems
- ⌘ MAC address cache size limits number of hosts.
- ⌘ Limit the hosts to only those that require Layer 2 connectivity for application survivability

Questions?

- ☒ Migration
 - ☒ Full RAM copy
 - ☒ Marks pages touched by processes as dirty
 - ☒ Then to partial copy, shadow page
 - ☒ Marks dirty
 - ☒ Pause VM, final copy
 - ☒ Block devices and gratuitous arp
 - ☒ Unpause on remote machine.
- ☒ Multi master read/write replication