



How to design application availability on a stretched cluster

#NLVMUG

VMUG[®] **USERCON**

NLVMUG UserCon
March 17 2016

Introductie

Niels Hagoort

- VMware VCDX #212
- VMware vExpert
- @NHagoort



Sr. Virtualization
Consultant



www.cloudfix.nl



Agenda

- Application availability
- Stretched clusters
- Failure scenarios
- Conclusie
- Q&A

Application Availability

- Business requirements?
- SLA → Service Level Agreement



Application Availability

SLA level of **99.98 %** uptime/availability:

- **Monthly:** 8m 45.9s
- **Yearly:** 1h 45m 11.4s

SLA level of **99.99 %** uptime/availability:

- **Monthly:** 4m 23.0s
- **Yearly:** 52m 35.7s

“Design for failure”

A word cloud of various IT and network-related terms. The words are arranged in a roughly triangular shape, pointing to the right. The colors of the words range from dark red to light yellow. The terms include: RPO, scenarios, witness, overlay, split-brain, performance, DR replication, datacenters, vMSC, synchronous, fault, alignment, RTO, availability, recoverability, VSAN, SAN, tolerance, applications, VLAN, NSX, OpEx, network, disaster-recovery, business-continuity, latency, stretched-cluster, CapEx, disaster-avoidance, SLA, cost, metro, cluster, connectivity, DRS, failure, BDP, dark fiber, and fiber.

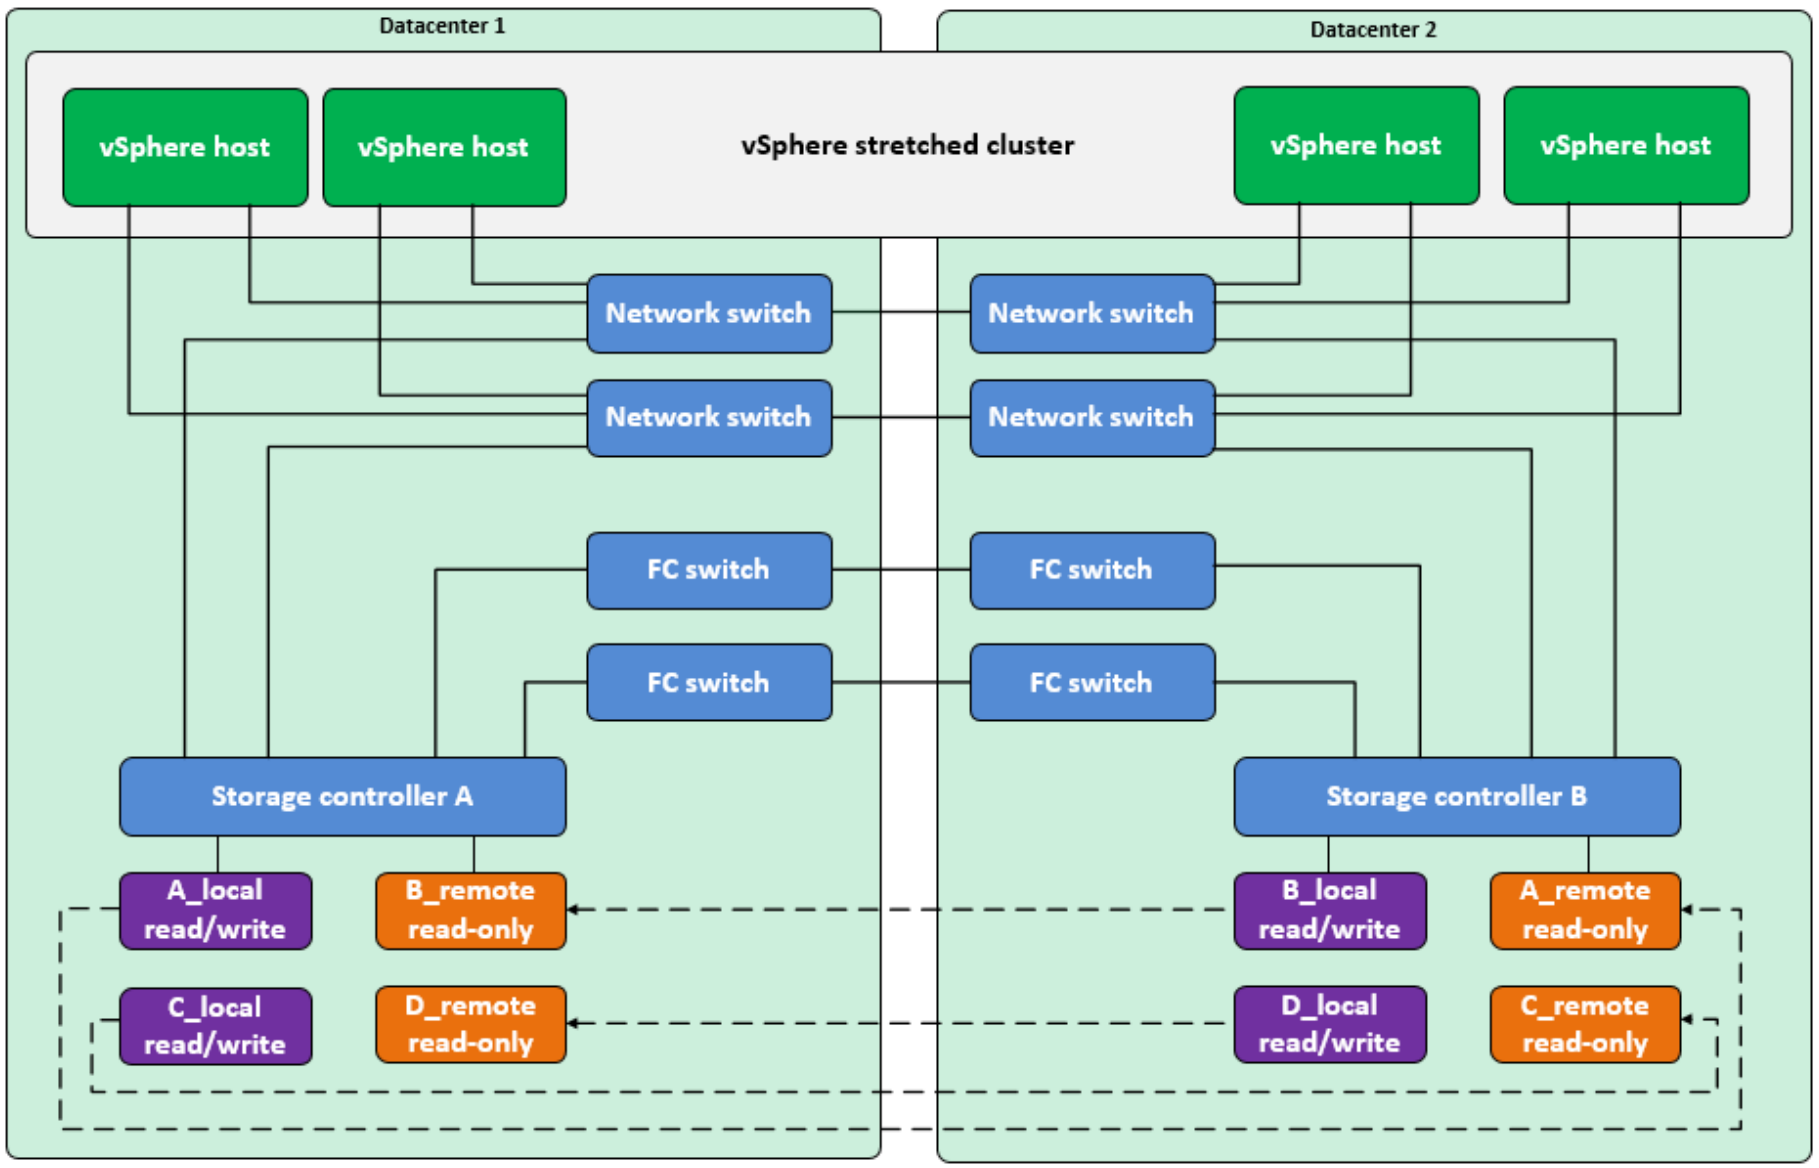
RPO
scenarios
witness
overlay
split-brain
performance
DR replication
datacenters
vMSC
synchronous
fault
alignment
RTO
availability
recoverability
VSAN
SAN
tolerance
applications
VLAN
NSX
OpEx
network
disaster-recovery
business-continuity
latency
stretched-cluster
CapEx
disaster-avoidance
SLA
cost
metro
cluster
connectivity
DRS
failure
BDP
dark fiber
fiber



- Disaster avoidance
- Disaster recovery
 - Zero RPO
 - Very low RTO
- Automatische failover
- Active site balancing



- Kosten?
- Complex?
- 'maar' 2 actieve sites
- Geen gecontroleerde failover



vSphere

- HA & DRS niet site-aware
- Admission control policy :
Percentage of cluster resources reserved as failover spare capacity
(50% CPU / 50% memory)
- Alleen power-on gegarandeerd!
- Gelijk inrichting beide sites

vSphere advanced settings

- Minimale performance garantie?
`das.vmMemoryMinMB="2048"`
`das.vmCpuMinMHz="512"`
- Isolation addresses, minimaal 1 extra per site
`das.isolationaddress0=x.x.x.x`
`das.isolationaddress1=x.x.x.x`
- HA rekening houden met affinity rules
`das.respectVmHostSoftAffinityRules=true`

vSphere

- VM Component Protection (VMCP)
 - VM gedrag voor;
host failure / isolation / APD / PDL
- vMotion
 - L2/**L3** supported
 - 10ms RTT
 - >250 Mbps bandwidth per vMotion operation
 - Multi NIC vMotion?
- Replication-Assisted vMotion?

Netwerk

Requirements inter-datacenter connectivity

- Bandbreedte & latency
- Glasvezelring

Uitdagingen?

- VM mobility vereist L2
- L2 failure domains
- Applicatie/tenant segmentatie
- Netwerk flows

Netwerk

- 'Traditioneel' L2 over dedicated darkfibers
- VPLS instances via MPLS backbone
- Propriety protocol, b.v. Cisco OTV

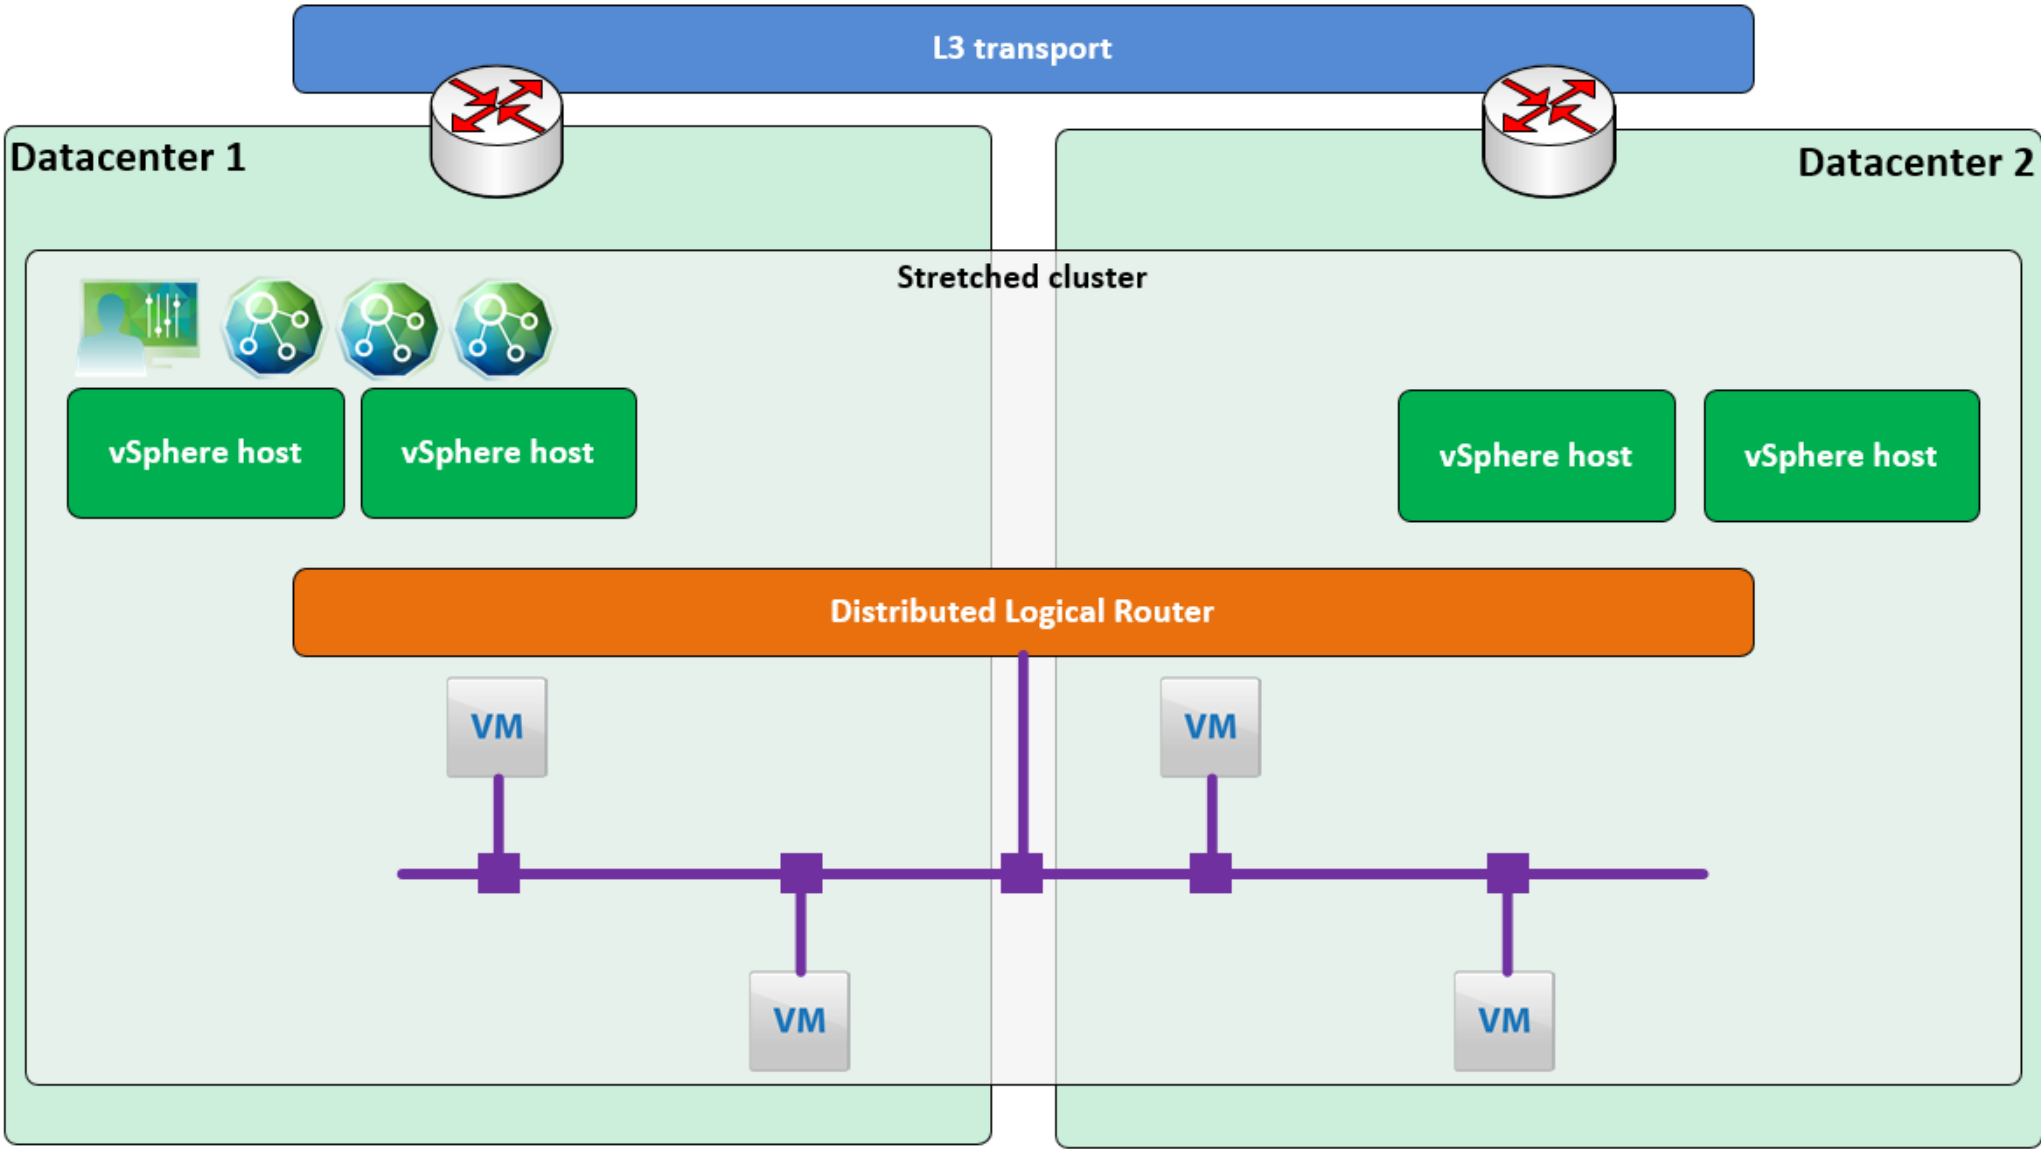
- Alleen bekeken vanuit netwerk perspectief, mist grotendeels enablement van automation

Netwerk

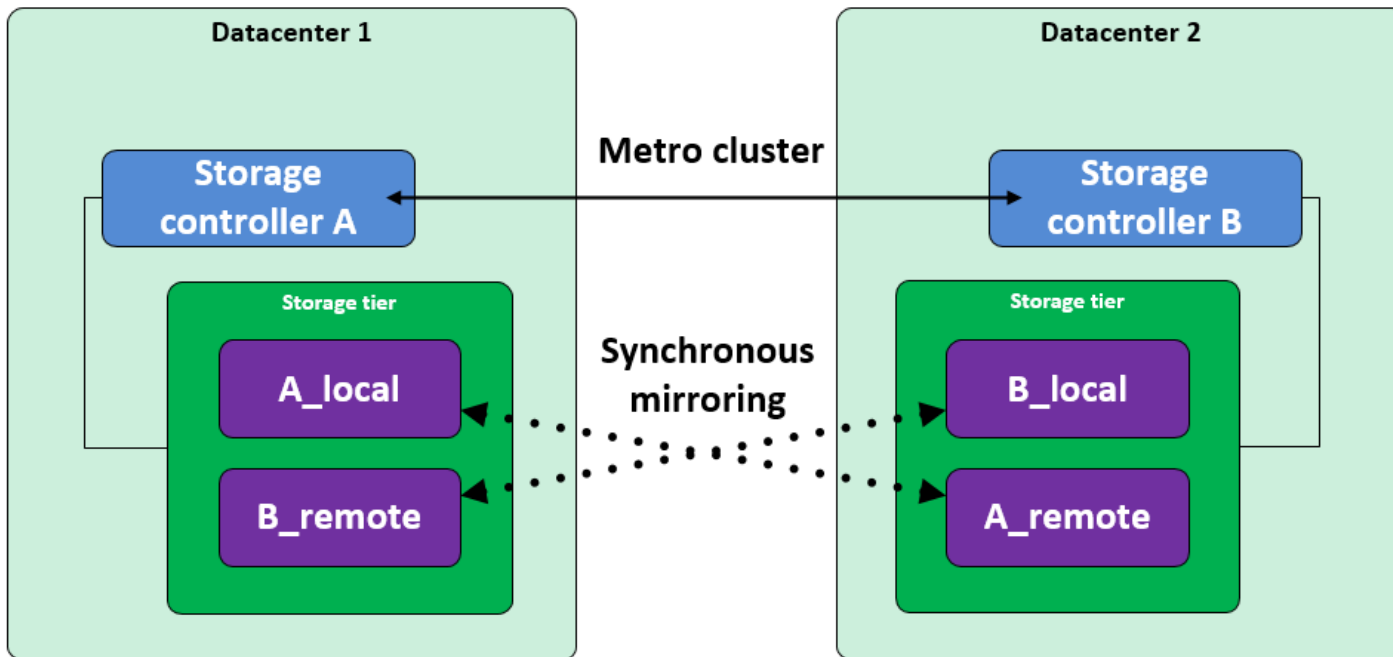
Hoe kan network virtualization helpen?

- Overlay (VXLAN)
- Mogelijkheid om L2 over L3 te expanden
- Toepasbaar tussen meerdere datacenters!
- Automation!



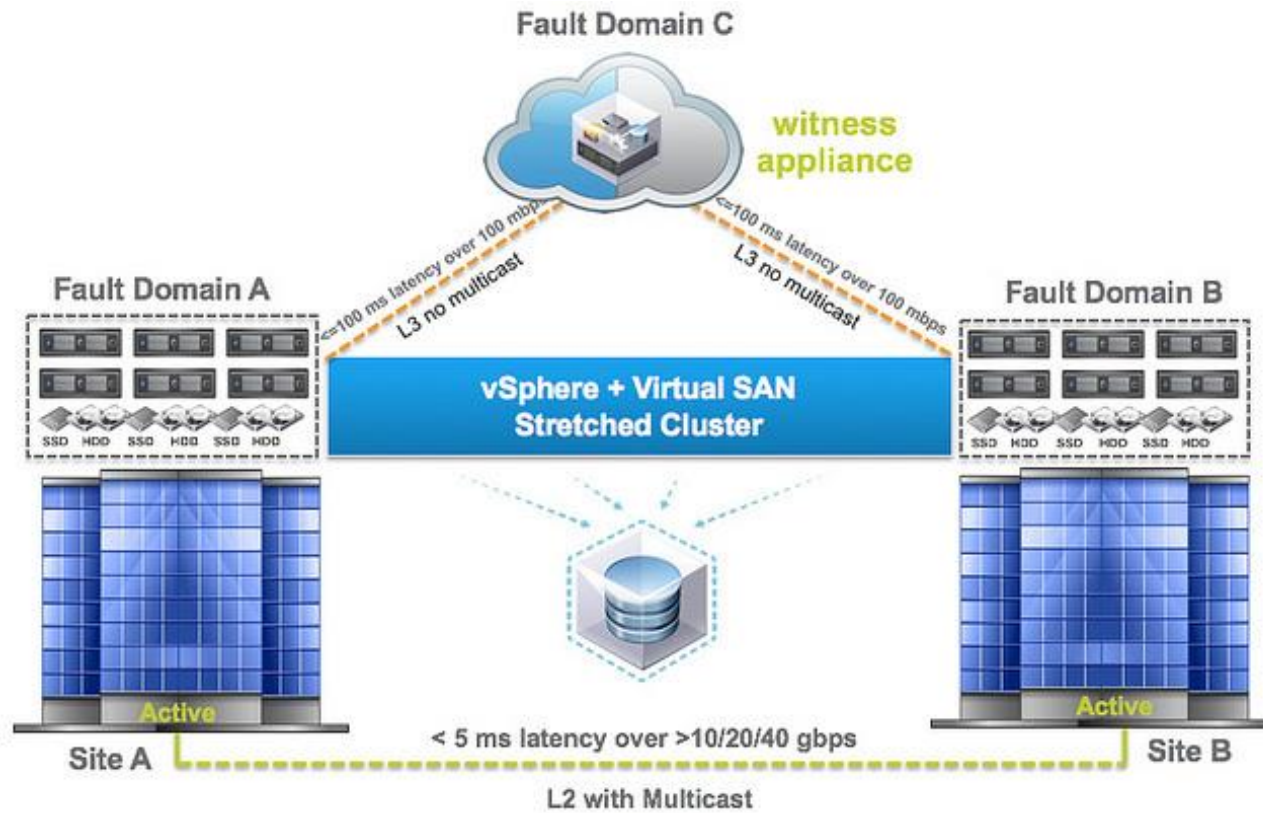


Storage



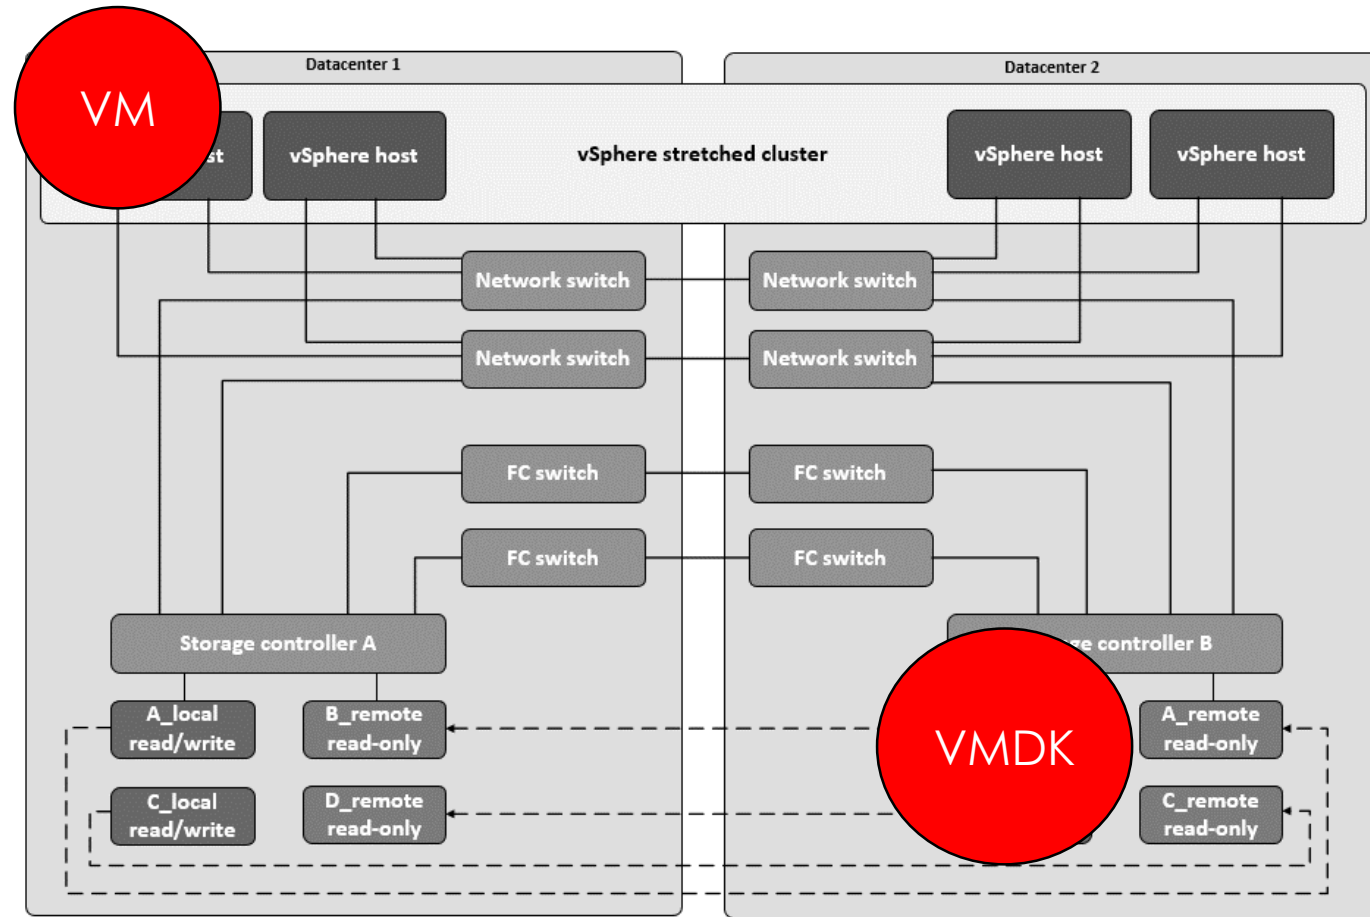
- vMSC
- Synchronone replicatie
- Performance hit?
 - Sync replic
 - Tijdens site failure
- Witness?

- HCI oplossingen metro availability
- VMware VSAN stretched cluster!



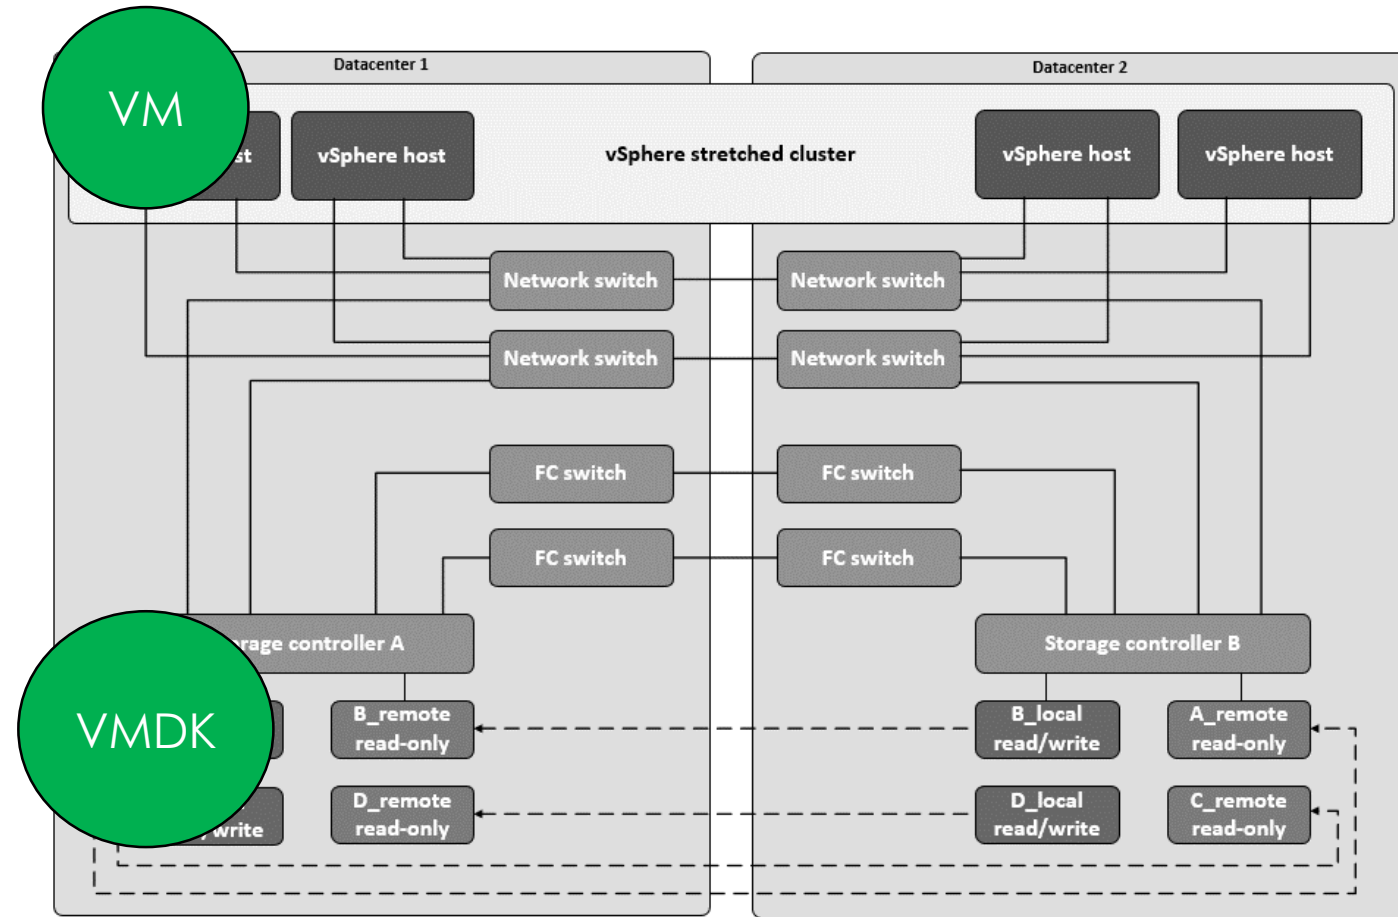
Affinity

- VM <> datastore alignment
- Minimaliseer impact op beschikbaarheid en performance



Affinity

- VM <> datastore alignment
- Minimaliseer impact op beschikbaarheid en performance



Failover?

WORST. FAILOVER.



EVER.

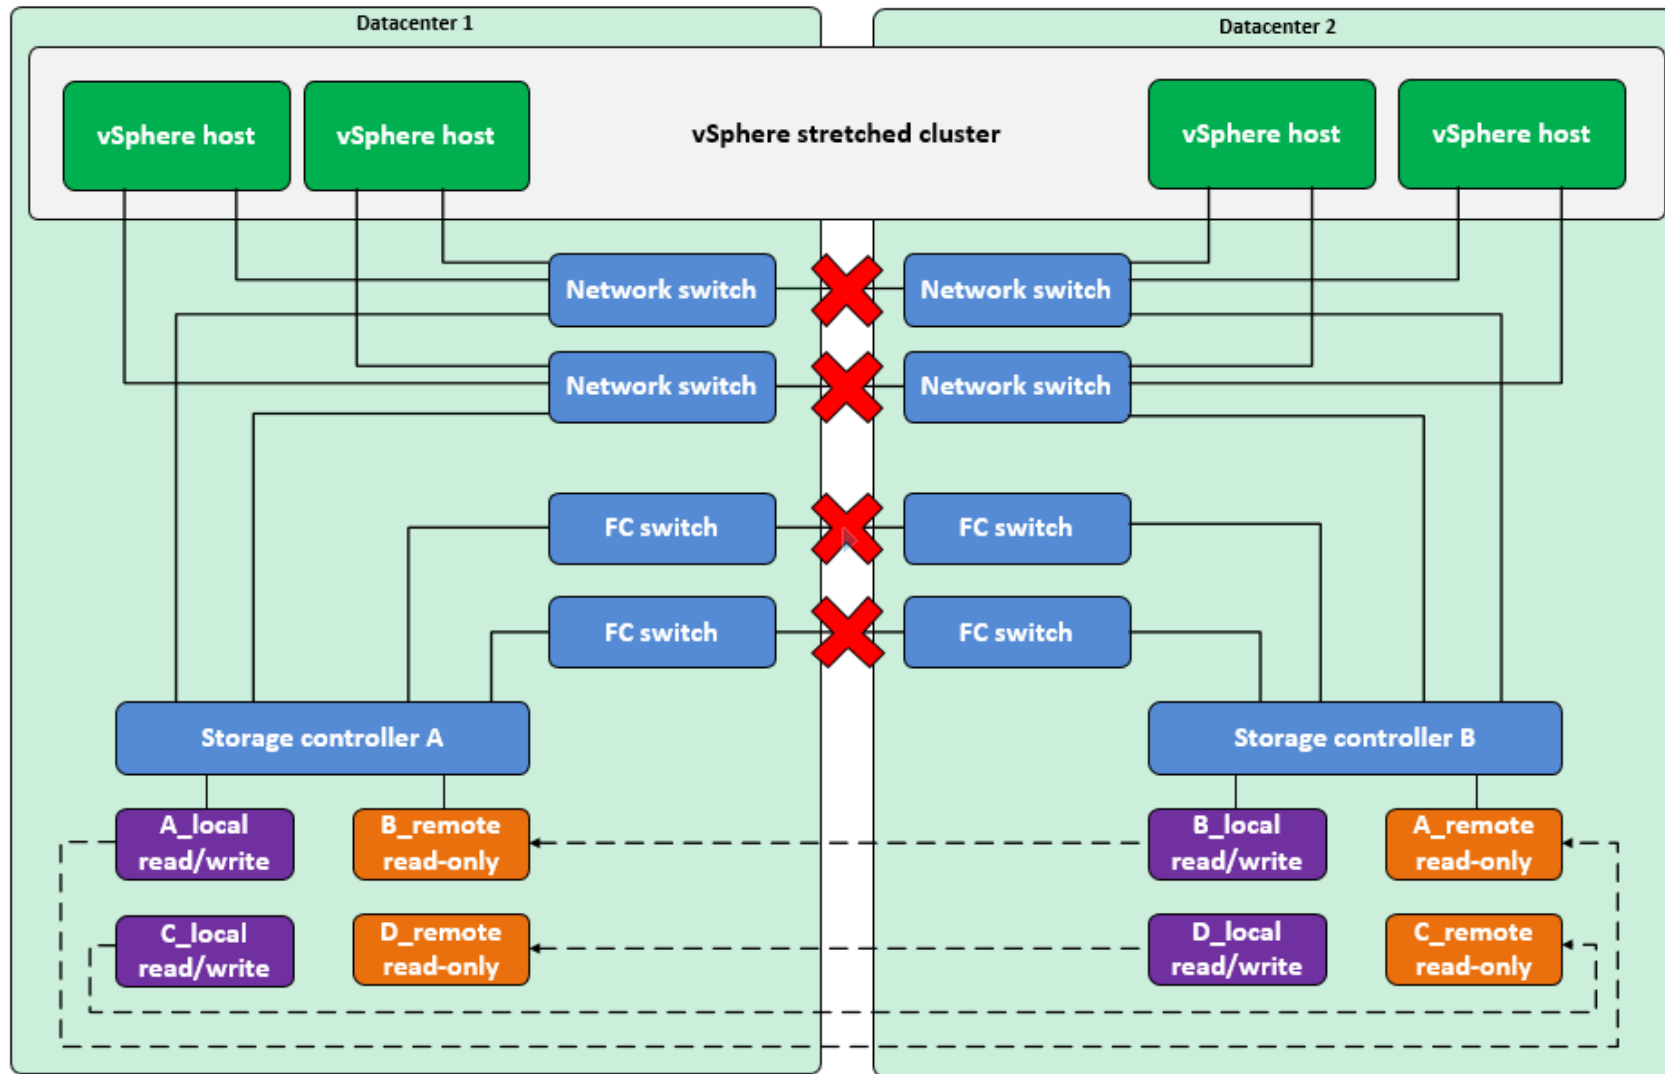


YOU SIR,

ARE AWESOME

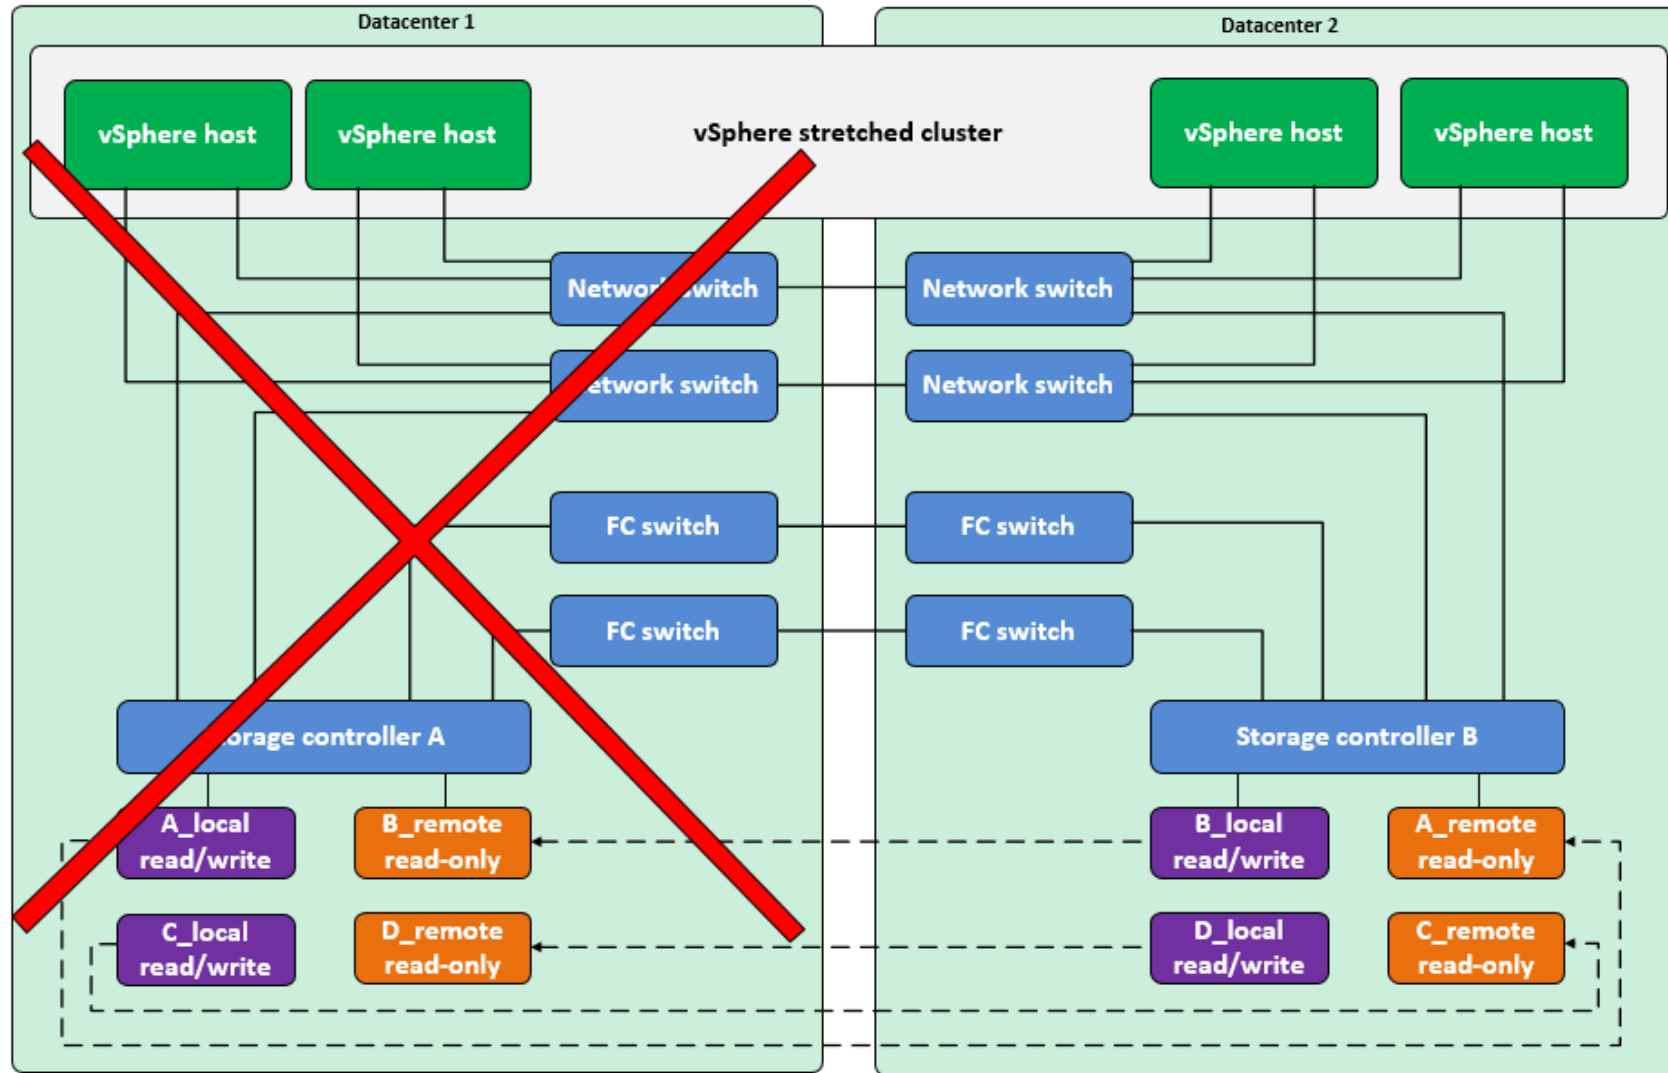
quickmeme.com

Datacenter partition



a.k.a.
split-brain

Datacenter failure



Conclusie

Stretched cluster prima invulling voor beschikbaarheid van applicaties!



- Test of je kan voldoen aan de SLA!
- Breng potentiële risico's in kaart!





Thank you!

#NLVMUG

VMUG[®] USERCON

NLVMUG UserCon
March 17 2016