

A decorative graphic consisting of a series of blue dots of varying sizes, arranged in a curved line that starts from the bottom left and moves towards the top right, ending near the title.

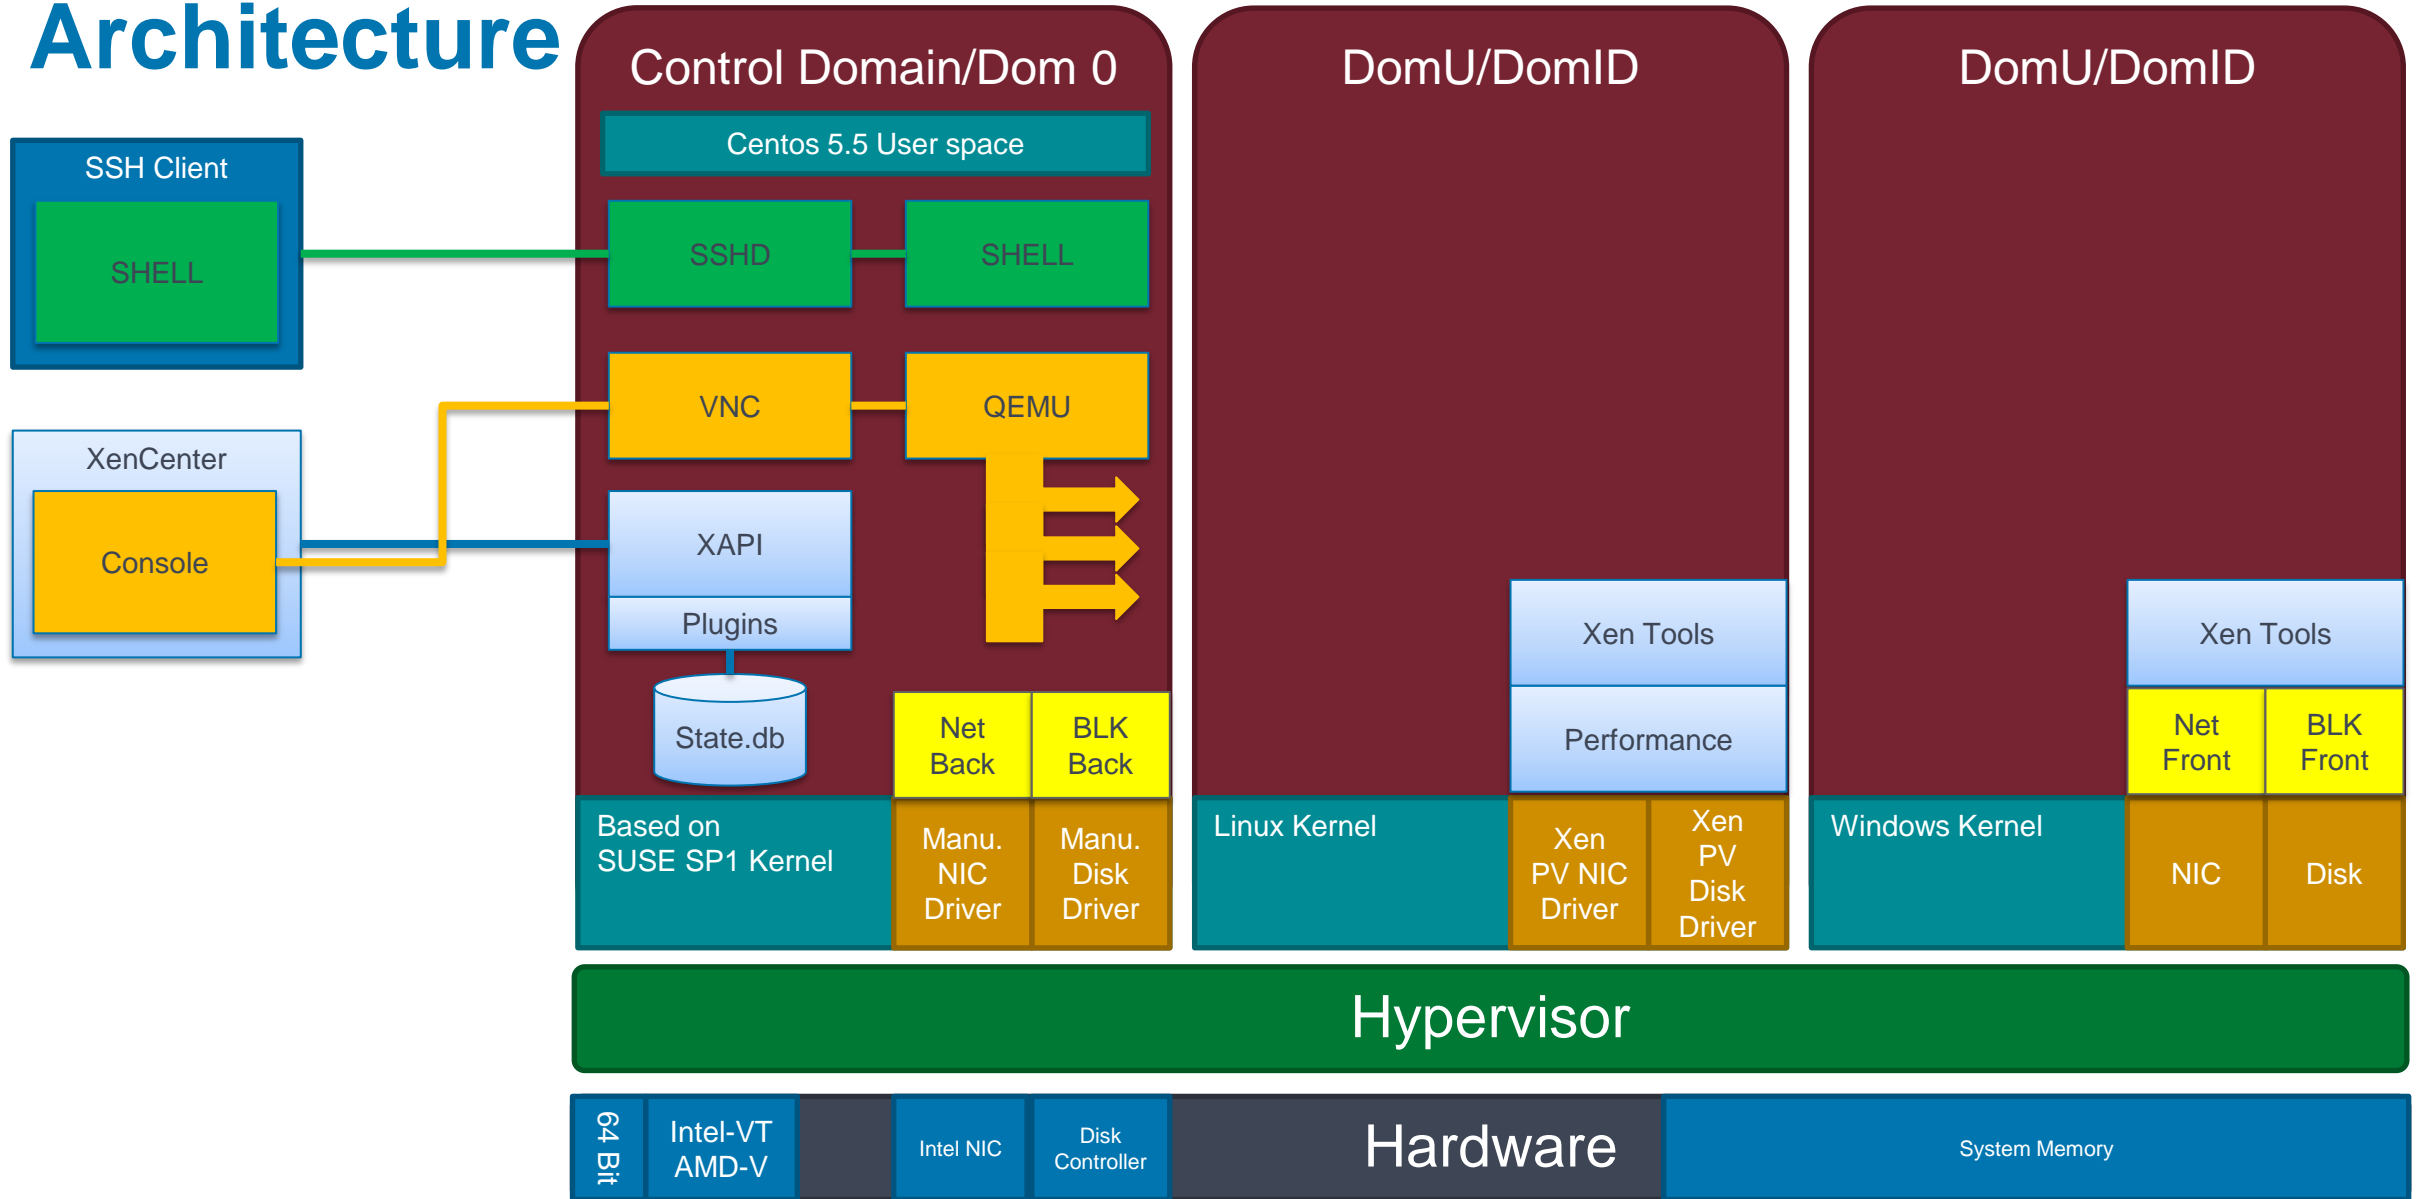
# XenServer Advanced

Christian Ferber  
Systems Engineer Datacenter & Cloud  
Server Virtualization



Architecture

# Architecture





# Command Line Interface (CLI)

# xe Command

- Runs CLI binary
  - Local in Dom0
  - Remote on Linux and Windows
  - Syntax: `xe help` or `xe help command` or `xe help -all`
- Network Communication
  - SSL, TCP 443
- Objects such as a VM or Server Addressed by
  - Name e.g. `WindowsAppServer`
  - unique UUID e.g. `79372186-1db3-4928-ac1f-5ae7feb5fd19`

# CLI

- **Filtering**

- `vm-list HVM-boot-policy="BIOS order" power-state=halted`

- **Parameterization**

- `xe vm-list params=name-label`

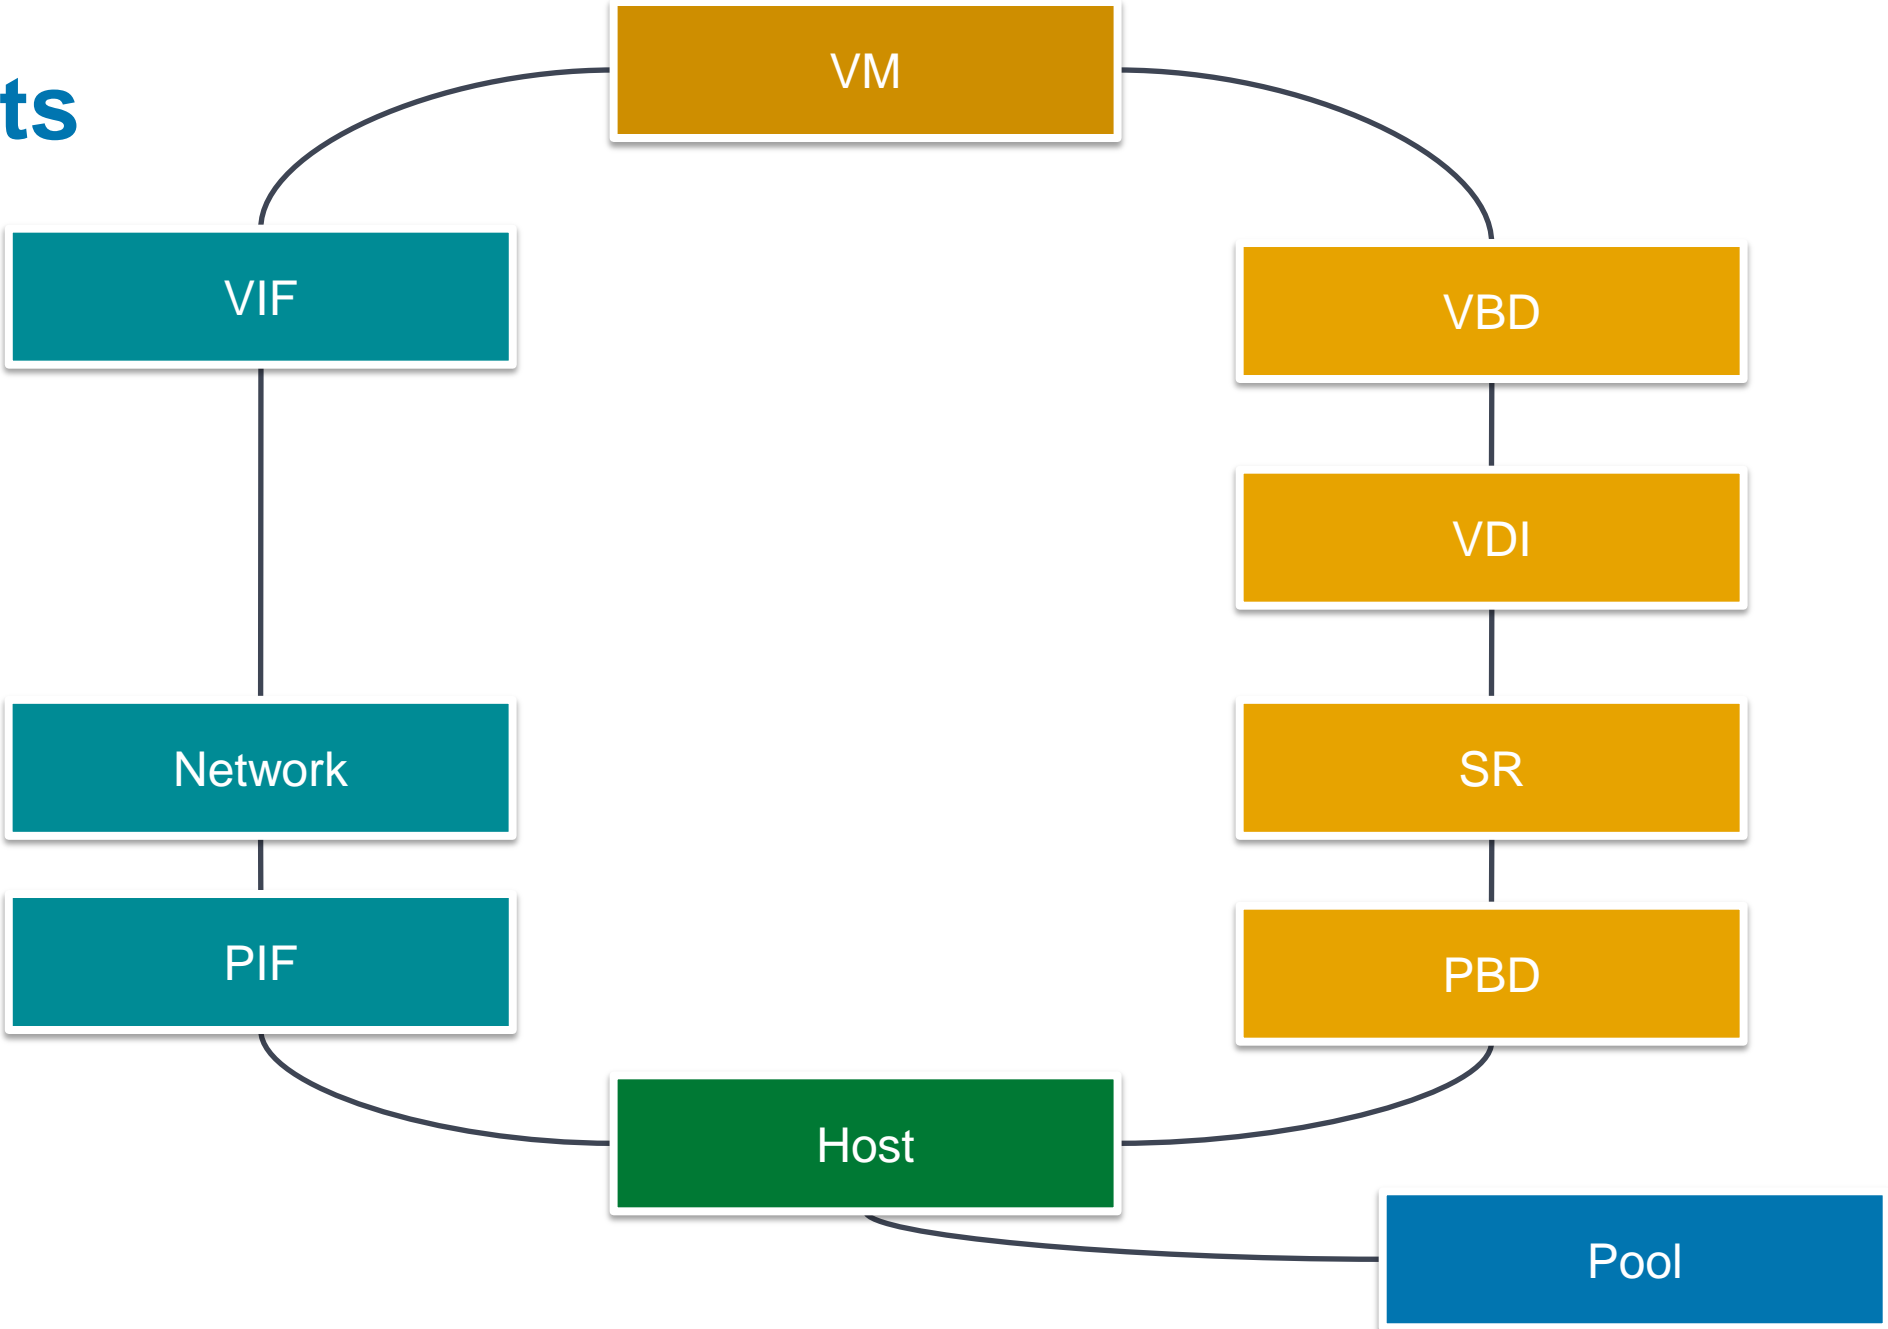
- **Parameterization and filtering can be combined**

- `xe vm-list HVM-boot-policy="BIOS order" power-state=halted  
params=name-label`

- **Minimal Output**

- `xe vm-list HVM-boot-policy="BIOS order" power-state=halted  
params=name-label --minimal`

# Objects



# XenServer Log Files

- XenServer XAPI Agent
  - `/var/log/messages` - system
  - `/var/log/xensource.log` – XenServer only
  - `/var/log/xha.log` – HA logging
- \*Logs can be sent to remote syslog daemon
- Domain 0
  - Hypervisor boot messages: `xe host-dmesg`
  - Domain 0 boot messages: `dmesg`



# XenServer logging

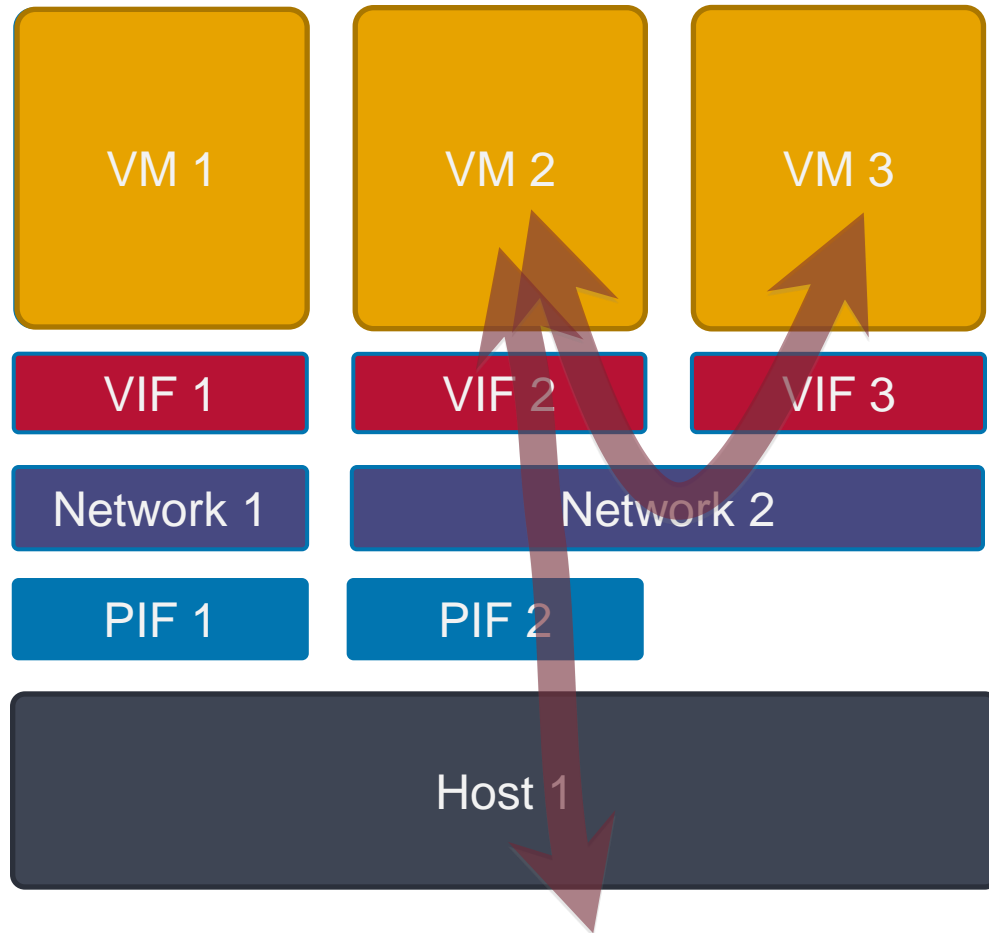
- Log files are rotated on a daily base
- Rotated log files can be seen as .<number> files
- Max of 20 log files is saved by default
- Create specific log files volume  
<http://support.citrix.com/article/CTX130245>
- Change logging options  
<http://support.citrix.com/article/CTX130327>

```
Last login: Fri Mar 25 23:54:48 on pts/0
Type "xsconsole" for access to the management console.
[root@xsmaster-1 ~]# cd /var/log
[root@xsmaster-1 log]# ls
audit.log      messages      user.log.1    xenstored-access.log.15
audit.log.1    messages.1    u6d.log       xenstored-access.log.16
blktap         openoswitch   u6d.log.1     xenstored-access.log.17
boot.log       pm            VMPRlog       xenstored-access.log.18
btm           rpmpkgs       VMPRlog.1     xenstored-access.log.19
crit.log       sa            wtmp          xenstored-access.log.2
cron           samba         xen           xenstored-access.log.3
daemon.log     secure       xensource.log xenstored-access.log.4
daemon.log.1   secure.1     xensource.log.1 xenstored-access.log.5
dmesg         SMlog        xenstored-access.log xenstored-access.log.6
faillog       SMlog.1     xenstored-access.log.1 xenstored-access.log.7
installer     spooler      xenstored-access.log.10 xenstored-access.log.8
kern.log      squeezed.log xenstored-access.log.11 xenstored-access.log.9
kern.log.1    squeezed.log.1 xenstored-access.log.12
lastlog       tallylog     xenstored-access.log.13
maillog       user.log     xenstored-access.log.14
[root@xsmaster-1 log]#
```



# XenServer Networking

# Networking: single host



VMs are connected to virtual interfaces


Virtual interfaces are also connected to the Network

PIFs connected to Network

Each NIC has a "PIF" object

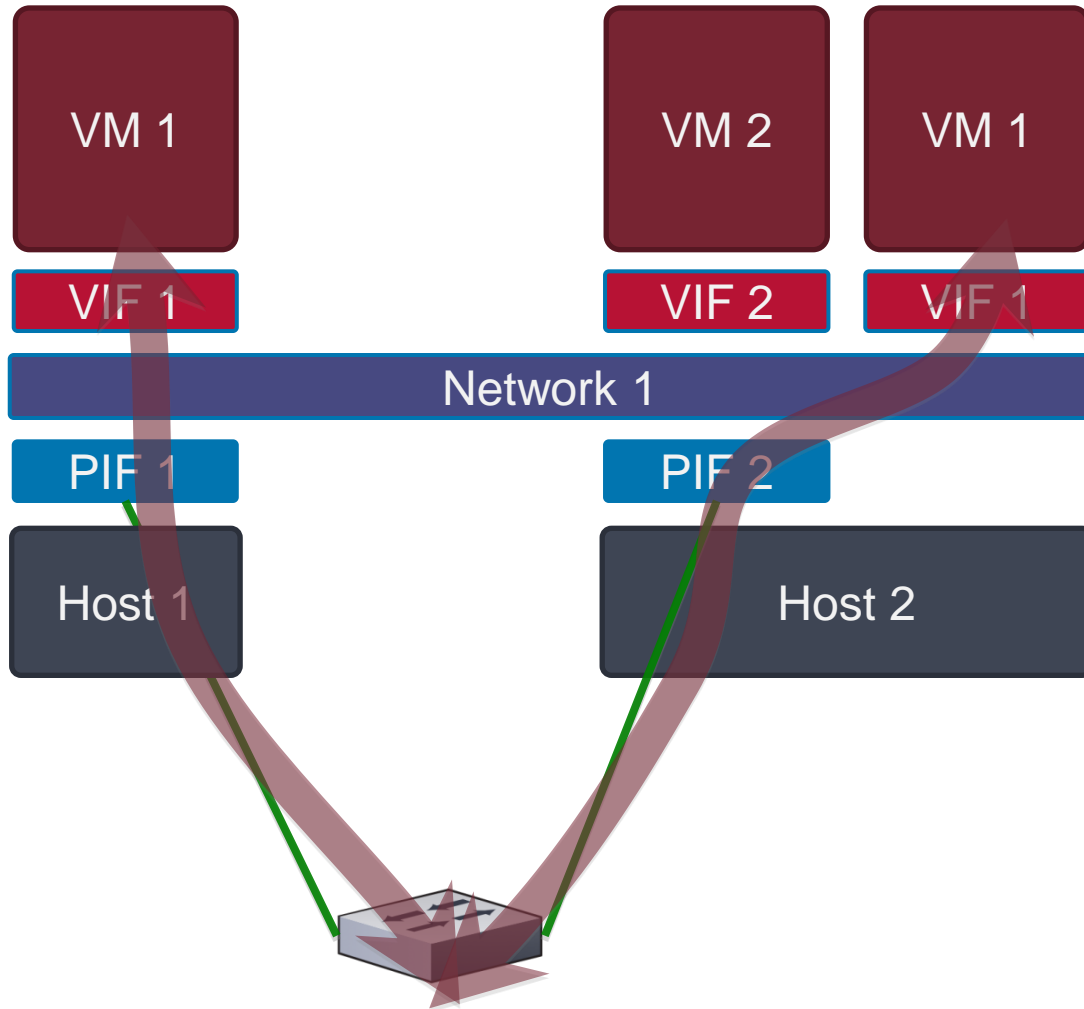
Host is connected to 2 NICs

# Networking: two hosts, not pooled



Normal IP networking rules, dependant on how PIFs are configured

# Networking: two hosts, pooled



Pooled NIC require same MTU size

Normal IP networking rules, dependant on how PIFs are configured

Network shared across pool

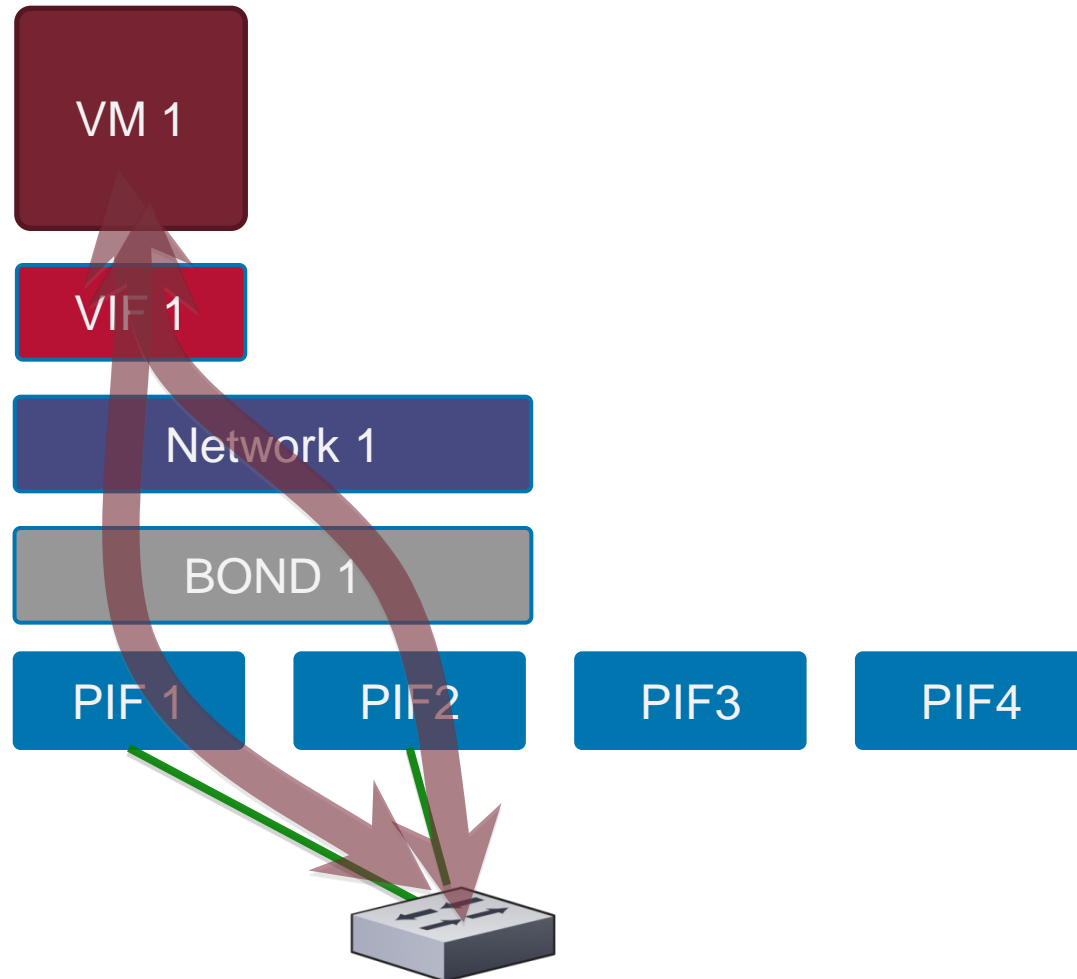
# Network objects: PIF

- XAPI view:
  - `xe pif-list params=all`
- Dom0 view:
  - `ifconfig`
    - eth0
    - eth0.33 (VLAN)
    - Bond0
    - `__tmp56347856`

# Bonding

- SLB (Source Load Balancing)
- Load-balancing of VM and MGMT traffic
- Fail-over support for storage traffic
- Dynamic rebalancing (10 sec)
- Note: etherchannel or 802.3ad (LACP) not required

# Bonding



Seamless failover if interface goes down

Traffic through one PIF at any one time (SLB)

Bond joins network, instead of the underlying PIF devices

PIF0/PIF1: Form a NIC bond





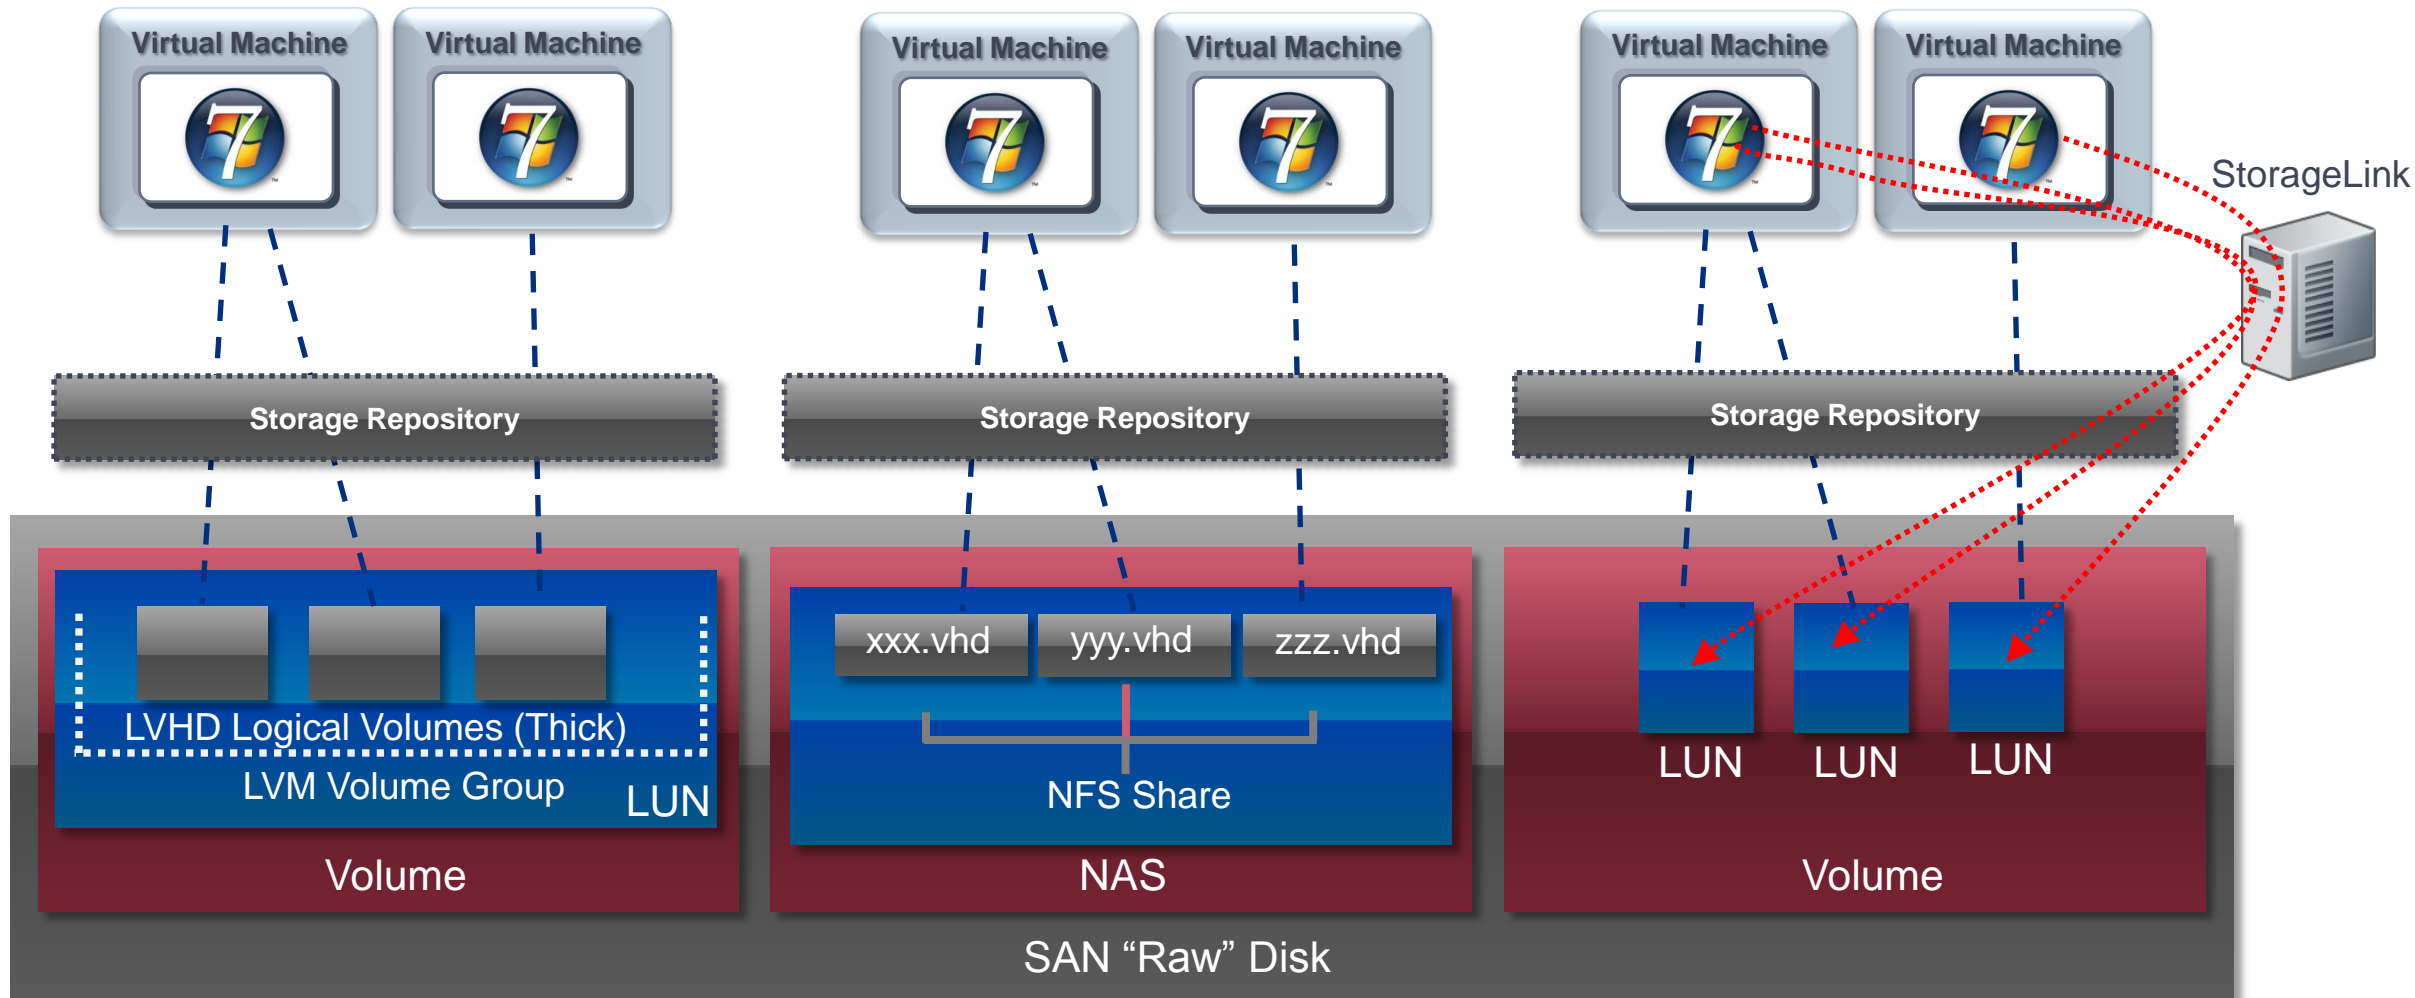
# XenServer Storage

# XenServer Disk Layouts (Shared)

## Native iSCSI & Fiber Channel

## NFS-Based Storage

## StorageLink Based SR



# LVM

- Host storage virtualized using LVM
- Volume Group named “VG\_XenStorage\_<UUID>”
- Common commands
  - `vgs` - List of volume groups
  - `lvs, lvdisplay` - List of logical volumes
  - `pvs, pvdisplay` - List of physical volumes
  - `man lvm` - Other commands
- Mount guest disk in dom0
  - `mount /dev/VG_XenStorage_<uuid>/<VM_uuid>.<disk> /path/to/mount/point`

# Common Issues

- Is the disk detected?
  - `fdisk -l` to show disks and partitions
- Is the module loaded?
  - `lsmod` – lists loaded modules
  - `insmod/modprobe` – inserts modules
  - `rmmod` – removes loaded modules
- Confirm device info?
  - `lspci`



# Performance Analysis

# Check Network Performance

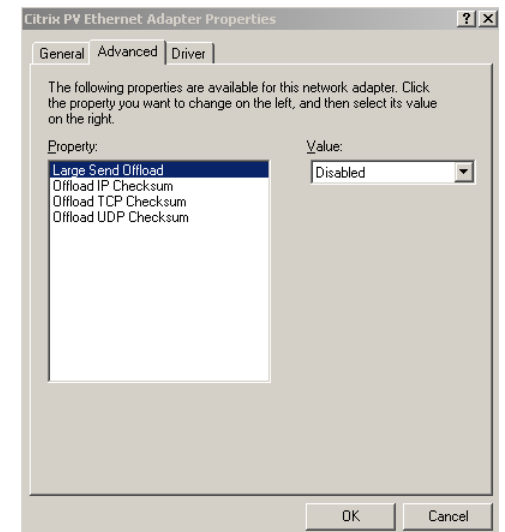
- Iperf
  - <https://nocweboldcst.ucf.edu/files/iperf.exe>
- Important: Use right parameters!!!!
  - e.g. Iperf ... -w 64k
- Check VM performance (Windows / Linux) + Dom-0
- Check interrupts
  - Command: top
  - Type: z
  - Type: 1
  - See interrupts as „si“

```
root@XDEU136521:~
top - 10:39:01 up 20 days, 44 min, 2 users, load average: 0.00, 0.00, 0.00
Tasks: 134 total, 4 running, 129 sleeping, 0 stopped, 1 zombie
Cpu0  : 2.0%us, 15.9%sy, 0.0%ni, 78.5%id, 0.0%wa, 0.0%st, 3.0%si, 0.7%st
Cpu1  : 0.0%us, 1.3%sy, 0.0%ni, 95.7%id, 0.0%wa, 0.0%st, 0.3%si, 2.7%st
Cpu2  : 0.7%us, 18.2%sy, 0.0%ni, 35.8%id, 0.0%wa, 0.0%st, 45.0%si, 0.3%st
Cpu3  : 0.0%us, 4.6%sy, 0.0%ni, 95.0%id, 0.0%wa, 0.0%st, 0.3%si, 0.0%st
Mem:   771328k total, 632568k used, 138760k free, 83404k buffers
Swap:  524280k total, 0k used, 524280k free, 27544k cache

  PID USER      PR  NI  VIRT  RES  SHR  S %CPU %MEM    TIME+  COMMAND
 23811 root        20   0 3880 3048 2452  R 79.1  0.4   20:27.19  tapdisk2
 24250 root         0   0   0   0   0   S 13.2  0.0    3:45.00  xb.00025.hda
  8584 root        20   0   0   0   0   R 12.6  0.0   12:07.96  rpciod/2
 18091 root        20   0   0   0   0   S  3.3  0.0    2:43.99  nfsiod
 23630 root        20   0 4012 2604 1876  R  3.3  0.3    1:56.28  tapdisk2
 25649 root        20   0 6028 4920 1016  S  2.3  0.6   71:35.52  xentop
  4144 root        10 -10 6712 2292 1040  S  1.3  0.3  161:03.14  ovs-vswhitcd
  8191 root       -10  5704 1748 1136  S  0.3  0.2   18:11.16  ovsdb-server
  8582 root        20   0   0   0   0   S  0.3  0.0    0:39.64  rpciod/0
  8583 root        20   0   0   0   0   S  0.3  0.0    2:07.42  rpciod/1
  8585 root        20   0   0   0   0   S  0.3  0.0    2:50.09  rpciod/3
 23980 65560     20   0 32632 6488 1748  S  0.3  0.8   66:48.16  qemu-dm
 25343 root        20   0  2560  388  268  S  0.3  0.1    5:27.91  irqbalance
 27304 root        20   0 2424 1116  832  R  0.3  0.1    0:00.29  top
    1 root        20   0  2160  664  568  S  0.0  0.1    0:24.46  init
    2 root        20   0   0   0   0   S  0.0  0.0    0:00.02  kthreadd
    3 root        RT   0   0   0   0   S  0.0  0.0    0:00.80  migration/0
```

# VM Networking Performance Improvement

- Task Offload is enabled by default
- Keep defaults, change when experiencing performance issues
- May require modification in specific situations
- Disable Task Offload
  - XenServer  $\leq$  5.5
    - <http://support.microsoft.com/kb/888750>
    - HKEY\_LOCAL\_MACHINE\System\CurrentControlSet\Services\Tcpip\Parameters  
DWORD DisableTaskOffload=1
  - Since XenServer 5.6 this can be done in the device properties of the network card under the „Advanced“-Tab



# Server Networking Performance improvement

- Task Offload could be set on server level as well
- Check setting
  - `ethtool -k <ethernet device e.g. eth0>`
- Enabling / disabling settings during operation (non-persistent)
  - `ethtool -K <ethernet device> rx off`
  - (example, K is case sensitive!!!)
- Enabling / disabling settings during operation (persistent)
  - `xe host-list`
  - `xe pif-list host-uuid=<host-uuid von oben> device=<nic device z.B. eth4> VLAN=-1`
  - `xe pif-param-set uuid=<pif-uuid> other-config:ethtool-rx="off"`

```
[root@xsmaster-1 log]# ethtool -k eth0
Offload parameters for eth0:
rx-checksumming: off
tx-checksumming: off
scatter-gather: off
tcp-segmentation-offload: off
udp-fragmentation-offload: off
generic-segmentation-offload: off
generic-receive-offload: off
large-receive-offload: off
[root@xsmaster-1 log]#
```



# Performance optimization NFS/iSCSI

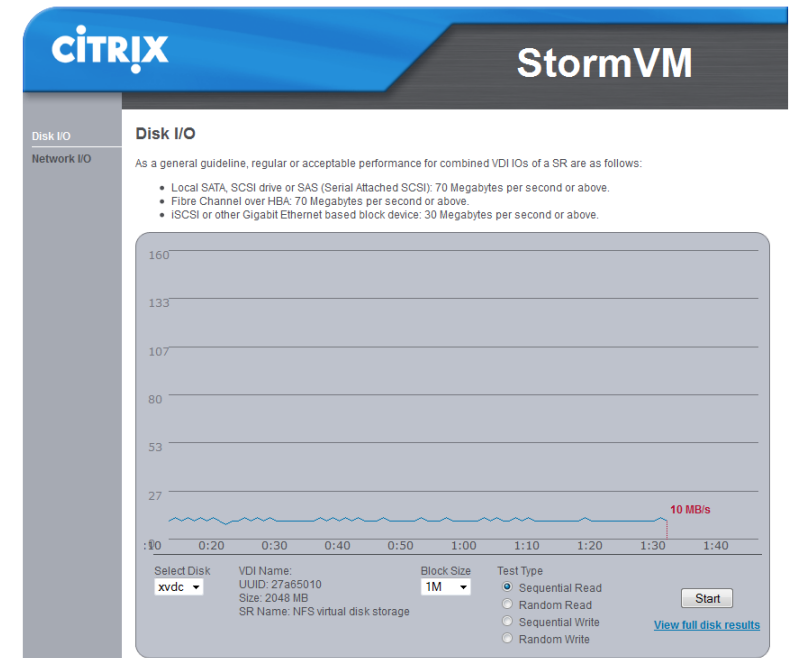
- Check large-receive-offload and enable it (ethtool -k), gain 80%
- Disable IOMMU in BIOS, gain up to 10%
- Enable JumboFrames (DVS)

# Enabling IRQ Rebalancing on XS 5.6 SP2 and older

- <http://blog.benpiper.com/2011/08/improving-network-throughput-in-xenserver-using-irqbalance/>
- `yum install irqbalance --enablerepo base`
- `Service irqbalance start`
- **Check interrupts**
  - `Cat /proc/interrupts`
- **Default in XenServer 6.0**

# Check storage performance

- Read performance: `hdparm -t /dev/sdb` (example)
- Tools like
  - IOMeter
  - Crystal Disk benchmark
- `iostat -x`
- XenServer performance VM
  - <http://support.citrix.com/article/CTX127065>



# Multipath Configuration

|                                 | DMP                         | DMP RDAC   | MPP RDAC   | MPP RDAC  |
|---------------------------------|-----------------------------|--|--|---|
| <b>Default</b>                  | Yes                         |  | No   | No  |
| <b>XenServer version</b>        | >= 5.0                      |  | = 5.0 Update 2   | =5.6 FP1  |
| <b>Management via XenCenter</b> | Yes                         |  | No   | No  |
| <b>Support</b>                  | Wide range of storage       | Only LSI controller based storage  | Only LSI controller based storage  | Only LSI controller based storage   |
| <b>Driver / Daemon</b>          | multipathd                  |  | mppVhba driver   | mppVhba driver  |
| <b>CLI path check</b>           | mpathutil status            |  | mppUtil -g 0   | mppUtil -g 0  |
| <b>Configuration</b>            | /etc/multipath-enabled.conf | /etc/multipath-enabled.conf<br>Hardware_handler „1 rdac“<br><small>Citrix use only - Do Not Distribute</small> | /etc/mpp.conf<br>(requires execution of /opt/xensource/bin /update-initrd) | opt/xensource/libexec/mpp-rdac -enable<br><br>opt/xensource/libexec/mpp-rdac -disable |

# SIMPLICITY IS POWER

A close-up photograph of a human hand in a grey, textured glove, firmly crushing several server hardware components. The components are blue and black, with various wires (red, yellow, black) protruding from the crushed pieces. The background is a plain, light color.

The power of server virtualization.

Crush cost and complexity like a can.

With Citrix® XenServer™, you can

eliminate unnecessary servers and the

costs associated with them. Instantly.

And best of all, for free. Makes you think

twice about complex, expensive, proprietary

virtualization alternatives. XenServer puts more

control in your hands than you've ever

had before. Feel the power of simplicity.

Simplicity is power. Citrix.

**CITRIX**

[Citrix.com/SimplicityIsPower](http://Citrix.com/SimplicityIsPower)