

# IA B29: Best Practices for Backup and Recovery of VMware

**Jim Olson**  
Technical Director

**George Winter**  
Technical Product Manager, NetBackup

Twitter: [@vmxgeorge](https://twitter.com/vmxgeorge)

## Disclaimer

**This information is about pre-release software. Any unreleased update to the product or other planned modification is subject to ongoing evaluation by Symantec and therefore subject to change.**

**This information is provided without warranty of any kind, express or implied. Customers who purchase Symantec products should make their purchase decision based upon features that are currently available.**

# NetBackup 7.6 Beta Sign Up



Norton™ Snap

Protecting you from dangerous QR codes


[www.norton.mobi](http://www.norton.mobi)

<http://tinyurl.com/NB76beta>



# NetBackup Has Won More Best of VMworld Awards Than Any Other Vendor

Business Continuity and Data Protection Category



Symantec NetBackup 5220  
One Appliance, end-to-end protection for enterprise

## Previous NetBackup Best of VMworld Awards

**GOLD**



**FINALIST**



**GOLD**



# Agenda

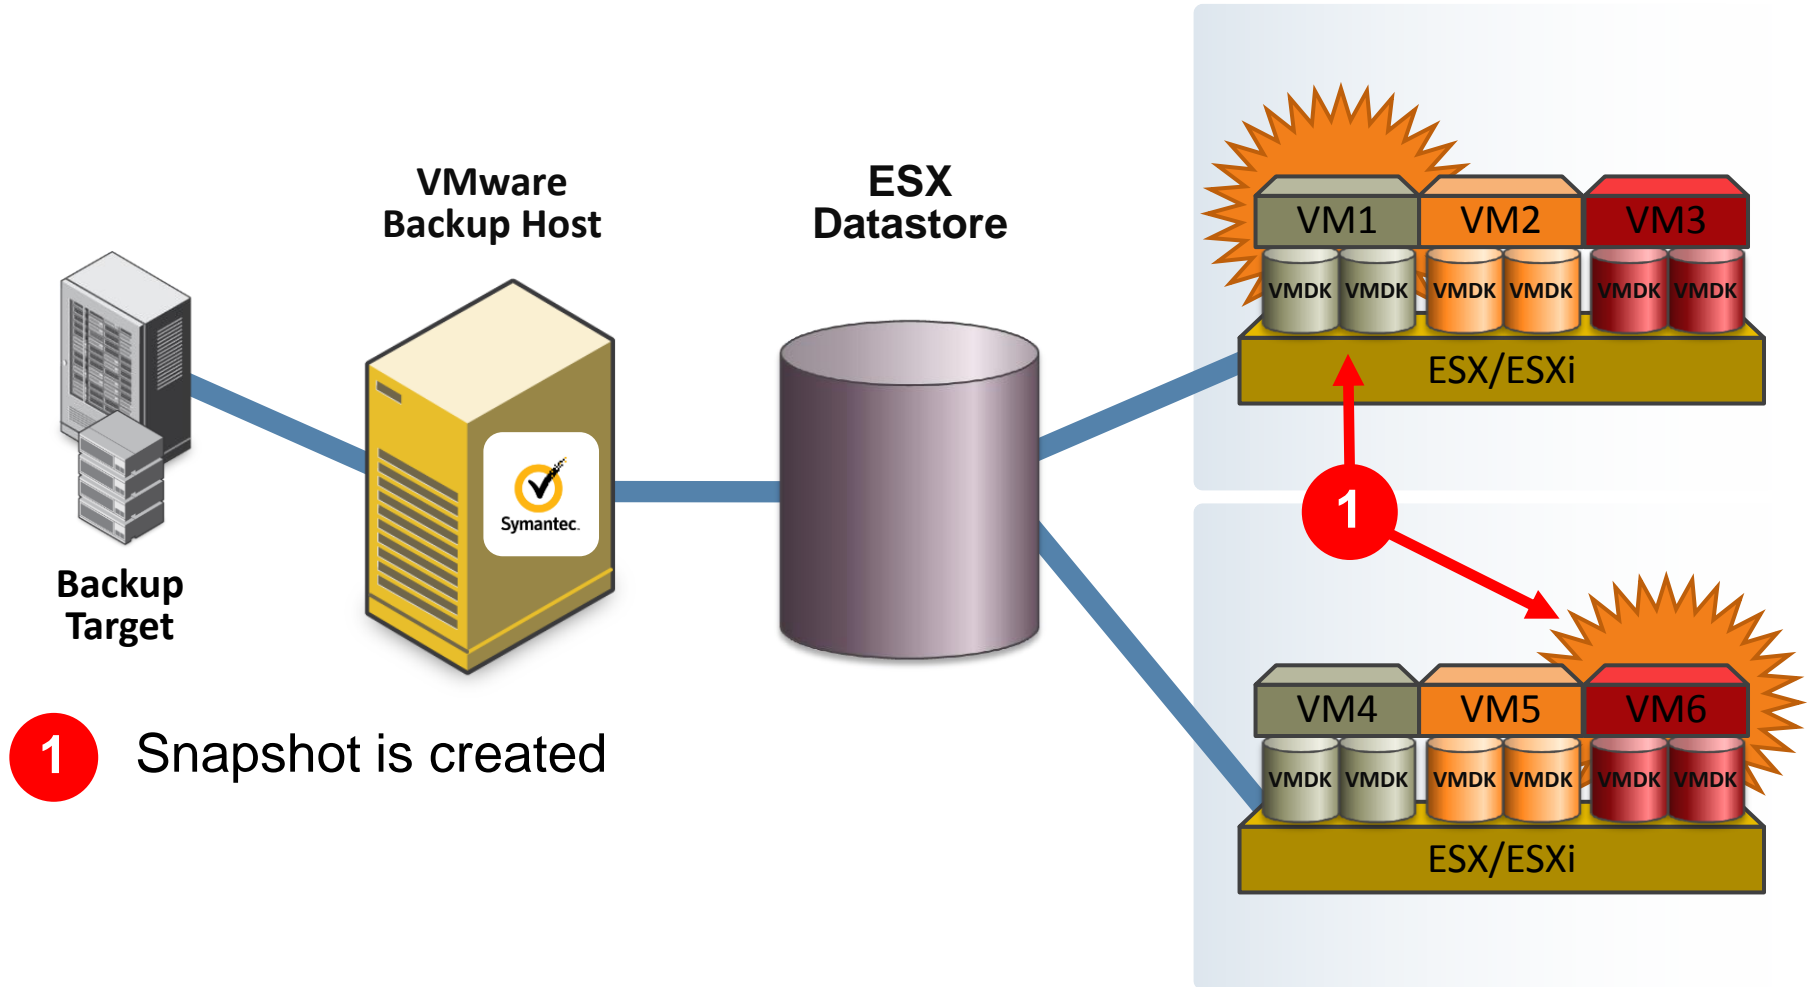
- VADP – Backup Process
- Backups And VMDK Files
- NetBackup Application & VADP
- Virtual Machine Restore Considerations
- VMware Intelligent Policy
- VMware Intelligent Policy Use Cases
- VIP Configuration Exercise
- Q&A



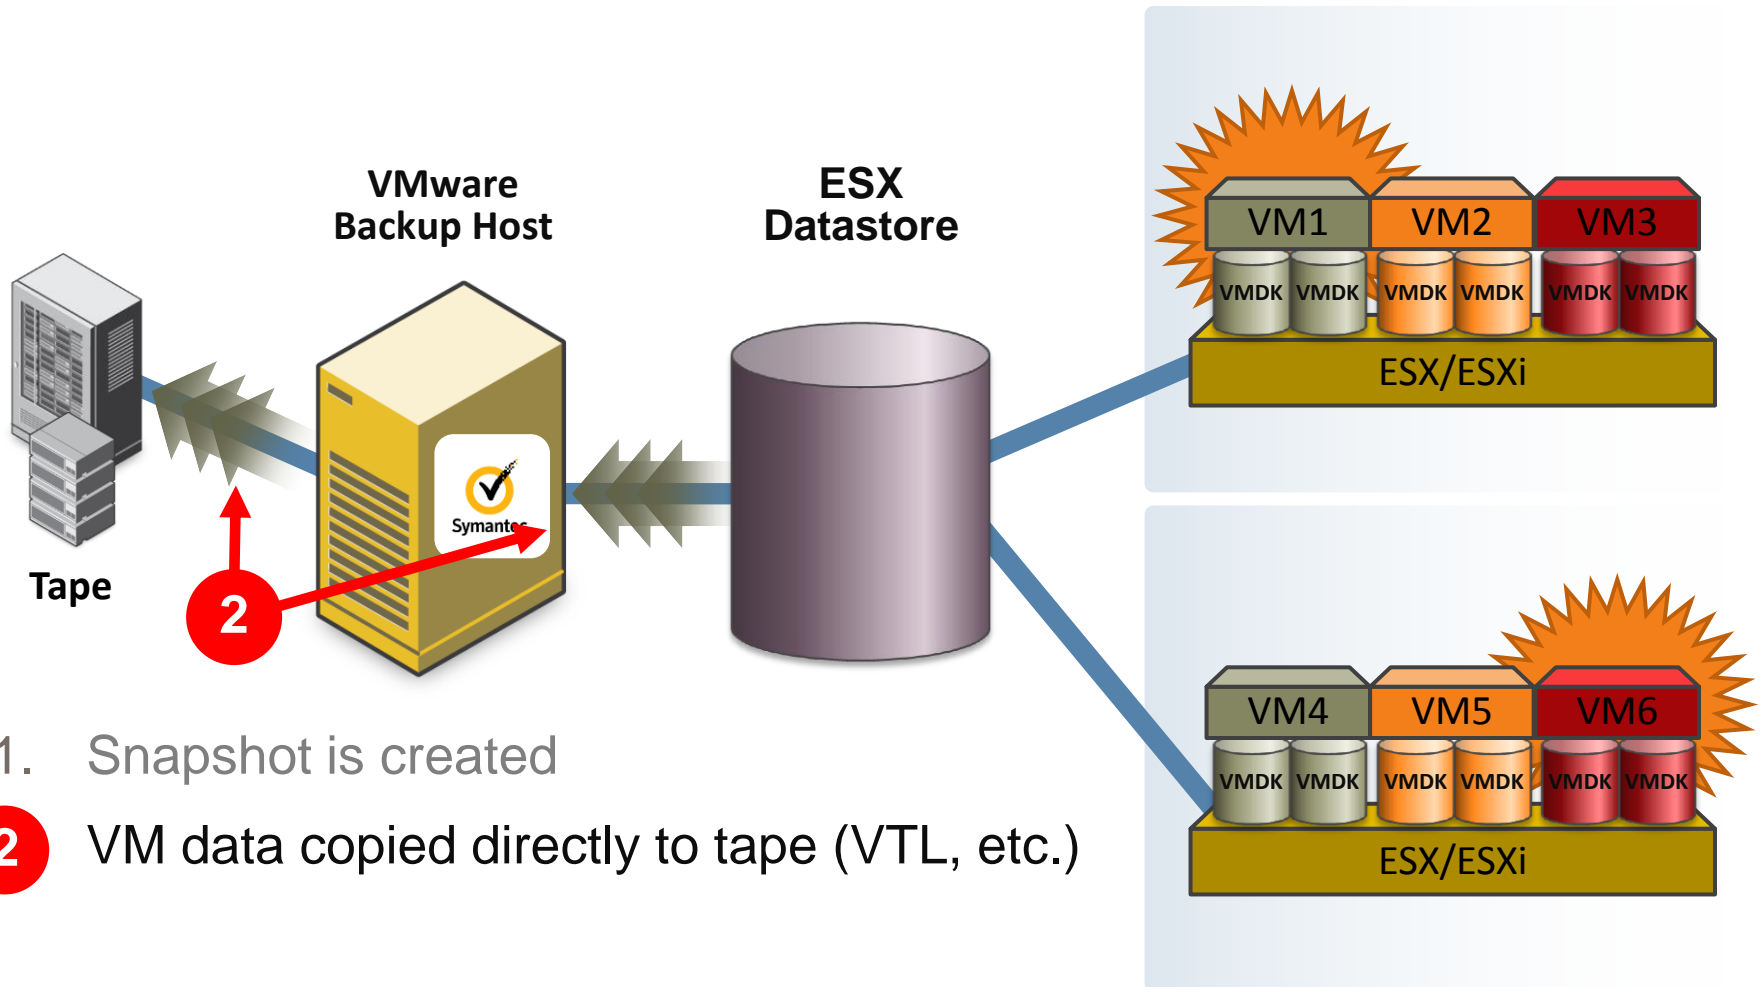
**vision**

## vStorage API for Data Protection (VADP)

# vStorage API Backup Process



# NetBackup 7 for VMware Backup Process

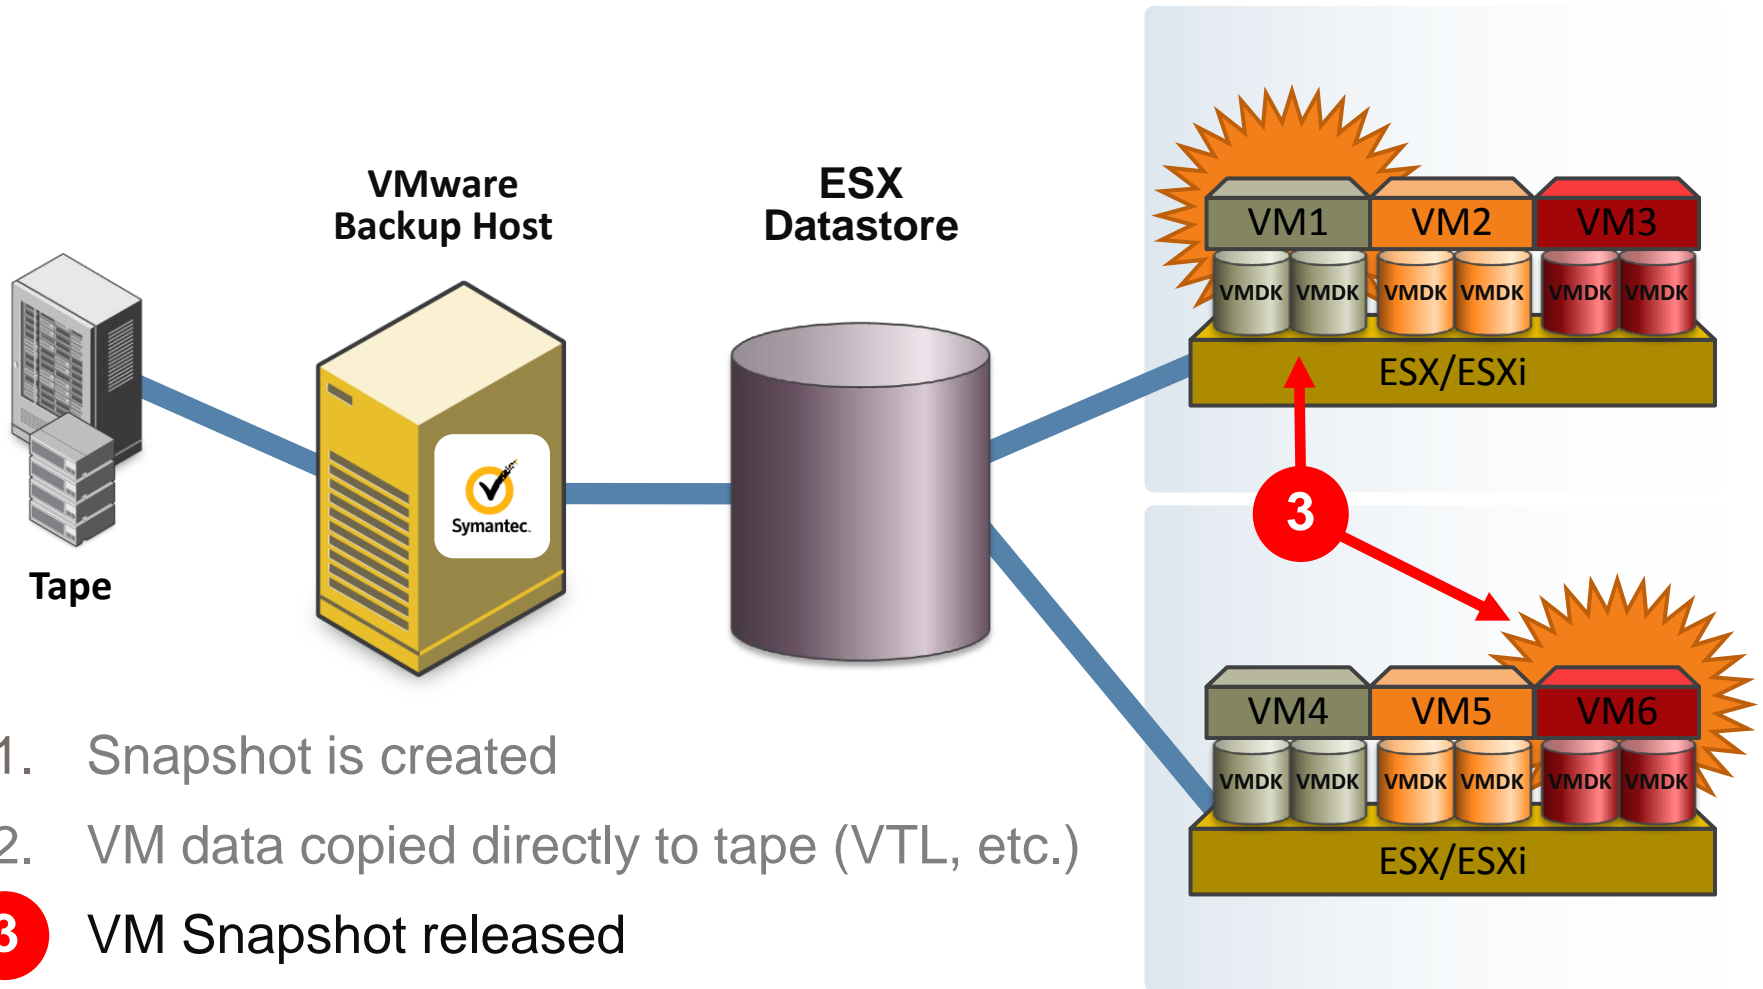


1. Snapshot is created

**2** VM data copied directly to tape (VTL, etc.)




# NetBackup 7 for VMware Backup Process



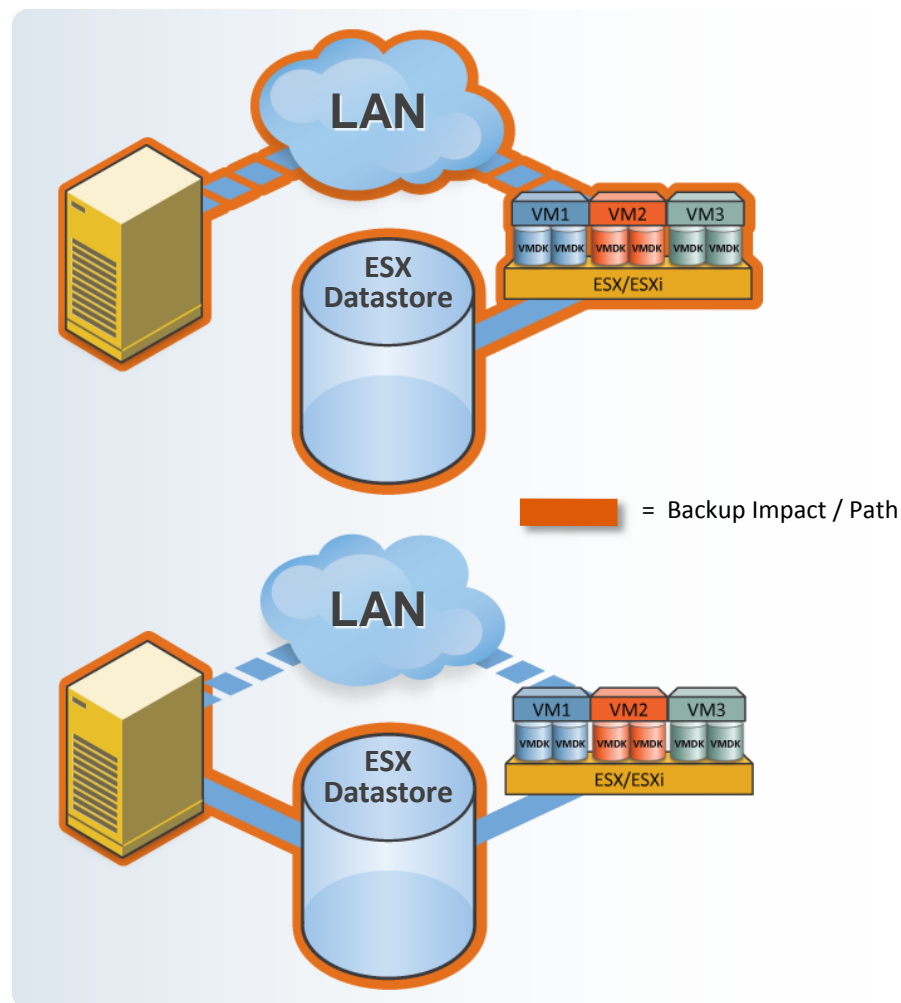
# NetBackup 7 for VMware Backup Data Path

- Network Backups (NBD)
  - NFS or DAS fully supported
  - No loss in backup or restore functionality
  - Direct communication with ESX server required (DNS, etc.)
  - ESX server directly impacted



# NetBackup 7 for VMware Backup Data Path

- Network Backups (NBD)
  - NFS or DAS fully supported
  - No loss in backup or restore functionality
  - Direct communication with ESX server required (DNS, etc.)
  - ESX server directly impacted
- Shared storage configuration
  - Fibre or iSCSI
  - Near zero impact on ESX
  - No loss in backup or restore functionality
  - No communication with ESX required
- Hotadd






**vision**

## Backups And VMDK Files

# VMDK Creation Process


- VMware rules for creation of new VMDK
  1. Empty VMDK is created first
  2. Space within the VMDK must be zeroed (mandatory)
  3. Data is now written to zeroed space within VMDK
- Redo log creation follows these same rules
  - Expanded in 16 MB extension – each extension must be “zeroed”
  - Three steps: create, zero, write
  - NOTE: similar process for thin provisioned VMDK creation
- Redo is a temporary VMDK
  - At end of snapshot redo must be reapplied to original VMDK



# A Closer Look - What Happens During a Backup?


## A Detailed Look at the VADP Snapshot Process

- During VM backups, VMware creates temp snapshot of VM. Process is as follows:
  - VSS provider flushes OS buffers in VM
  - Snapshot of VM is taken (vmdk(s) are frozen) (SCSI reservation of LUN)
  - Redo log created – all writes redirected to redo log (Redo expanded in 16 MB chunks)
  - VM is backed up
  - Redo log data applied to original vmdk(s)
  - Snapshot released - backup completed
- Why does this matter?
  - Every one of these steps involves significant amount of I/O
  - Reducing **number of snapshots per Datastore** improves backup perf & reliability
  - Incrementals are quick and can reduce snapshot impact



# Reducing Backup I/O Impact – Improve B/U Reliability

- Shorten time backup snapshot is open
  - Backup during periods of low VM activity
  - Use incremental backups (CBT) liberally
  - Limit simultaneous backups per ESX / Datastore (VIP)
- Configure NetBackup for optimal performance
  - Design backup policies to evenly **balance load across ESX / Datastores**
  - Faster backups = shorter time snapshots are open
  - Tune backup application buffers for optimal performance
- Result:
  - Snapshots more reliable
  - Overall backup processing faster
  - Higher level of backup success



# POP QUIZ




**Q:** How many additional IOPS are required for each write I/O that occurs while the VADP snapshot is active?

Hint: Redo logs are created like thin provisioned VMDKs



# POP QUIZ



**Q:** How many additional IOPS are required for each write I/O that occurs while the VADP snapshot is active?

**A:** Actual number depends on what is happening during backup but it can be a lot:

## Snapshot creation:

- Create REDO log
- Expand 16MB chunks
- Zero-out 16 MB chunks  
(repeat as necessary)

## Snapshot Deletion:

- Read data from REDO
- Apply data to original VMDK
- Remove REDO  
(Thin VMDK can increase this)

# POP QUIZ



## Notes:

- None of the IOPS listed below occur during standard VM I/O activity
- The operations listed occur for **every** actively backed up VM

## Snapshot creation:

- Create REDO log
  - Expand 16MB chunks
  - Zero-out 16 MB chunks
- (repeat as necessary)

## Snapshot Deletion:

- Read data from REDO
- Apply data to original VMDK
- Remove REDO

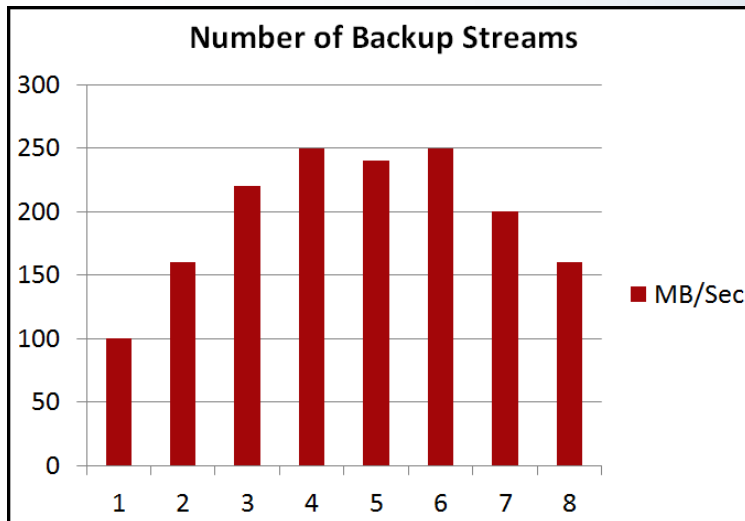


**vision**

## NetBackup Application & VADP Performance Configuration

# Ensure Optimal Backup Performance

- Number of simultaneous backups?

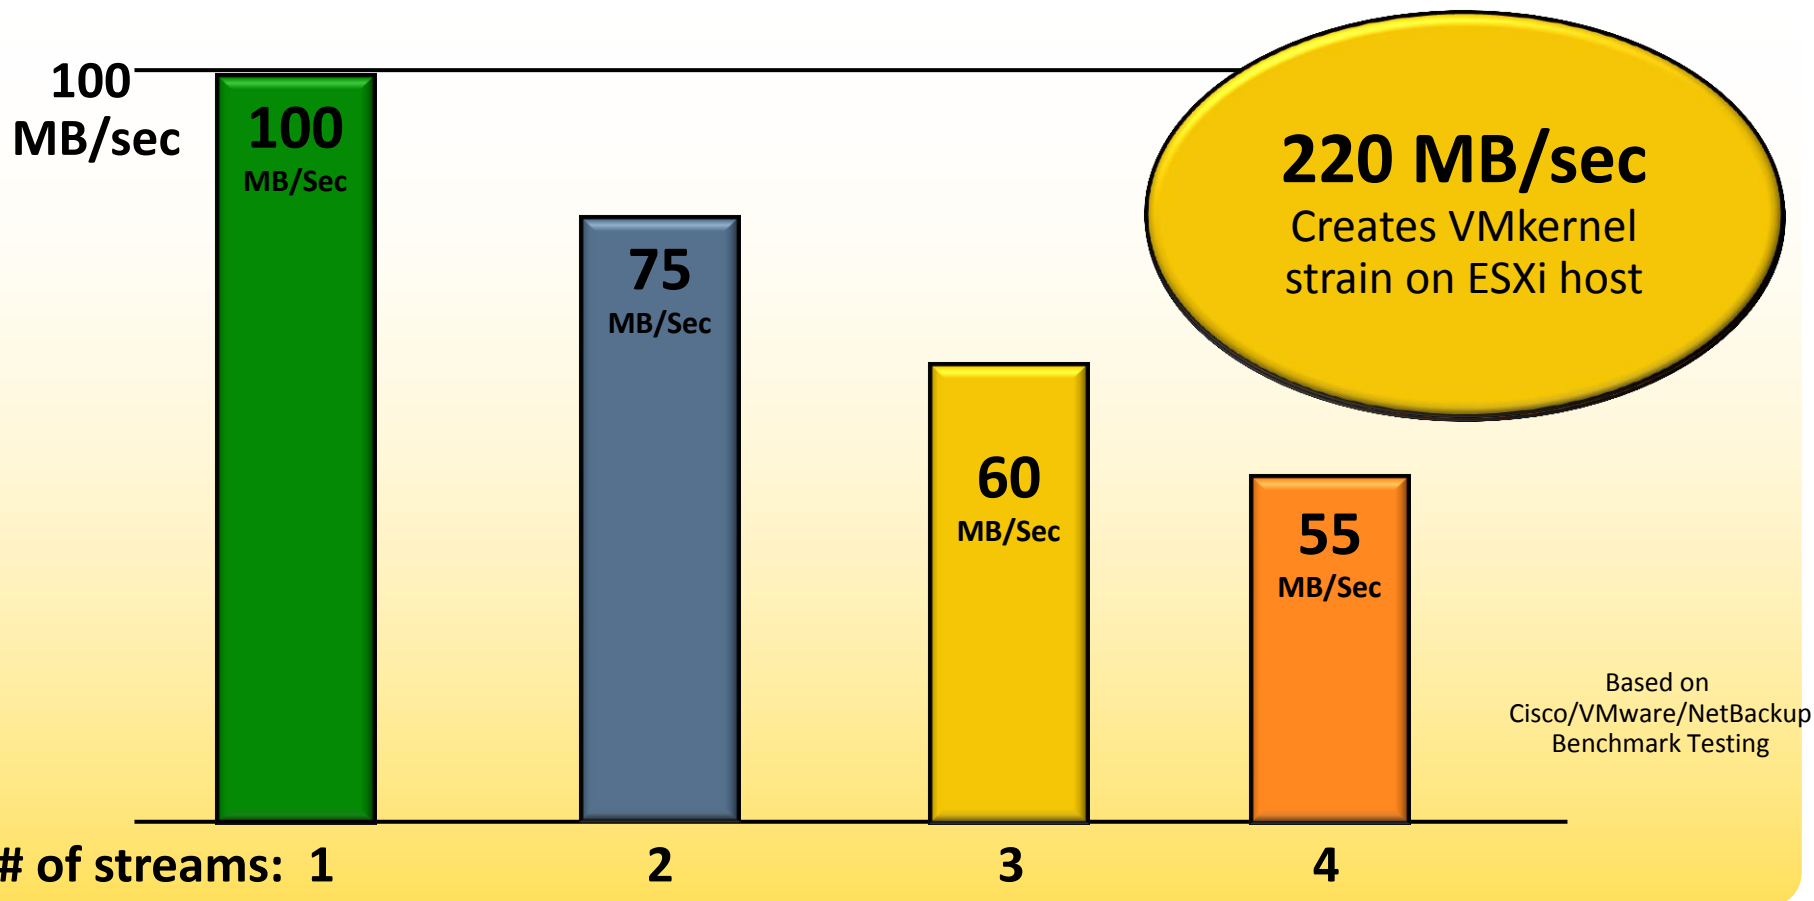


Aggregate Backup Speed

- Single VADP backup stream won't saturate backup path
- **VMware throttles VADP throughput to reserve VMkernel port bandwidth**
- Max backup performance achieved by creating simultaneous backup streams
- Design backups so that data is streamed across multiple VMware components (ESXi/DS)
- Improve backup speeds, shorten backup window

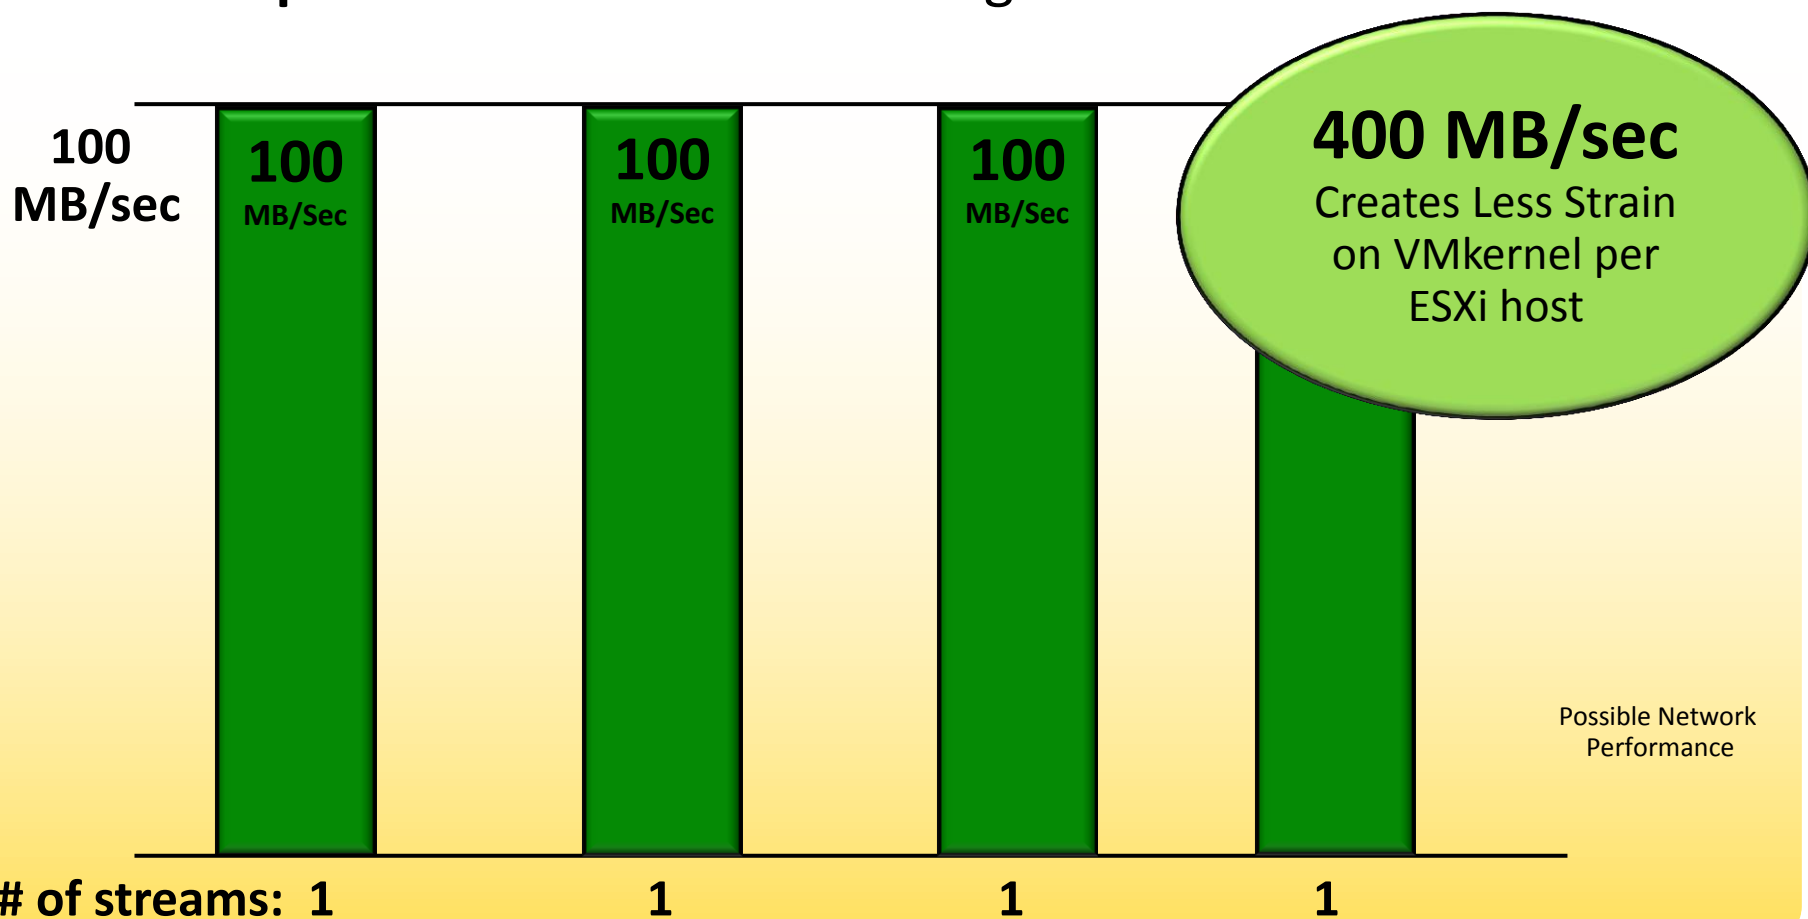
# VADP Performance Characteristics

- **Single ESXi host – NBD Performance Per Stream**



# VADP Performance Characteristics

- **Four Separate ESXi hosts** – Creating One Stream Per ESXi





**vision**

## Virtual Machine Restore Considerations




# General Restore Considerations


- Restore process involves more I/O than backup process
  - Disks (vmdk's) must be first created as target of restore
  - Type of vmdk can impact restore speed and I/O required
- Single restore typically won't saturate restore path
  - DR – create simultaneous restore jobs
  - As with backup – balance restores across ESX or Datastore
- Slow vCenter can also cause restore perf issues
  - Optional: bypass vCenter by restoring directly to ESX(i) server
  - Known to significantly improve restore perf in some cases




# VMDK Type Can Impact Restore Speeds



- thin
  - Space not allocated during creation
  - Space is supplied then zeroed out – on demand
  - Creation slow if vmdk turns out to be full



- zeroedthick
  - 100% of space allocated during creation
  - Zeroed out on demand
  - Can be faster than “thin” – especially if vmdk is nearly full



- eagerzeroedthick
  - 100% of space allocated during creation
  - 100% of disk zeroed out during creation
  - Could take long time (and create lots of I/O) to complete entire process



**vision**

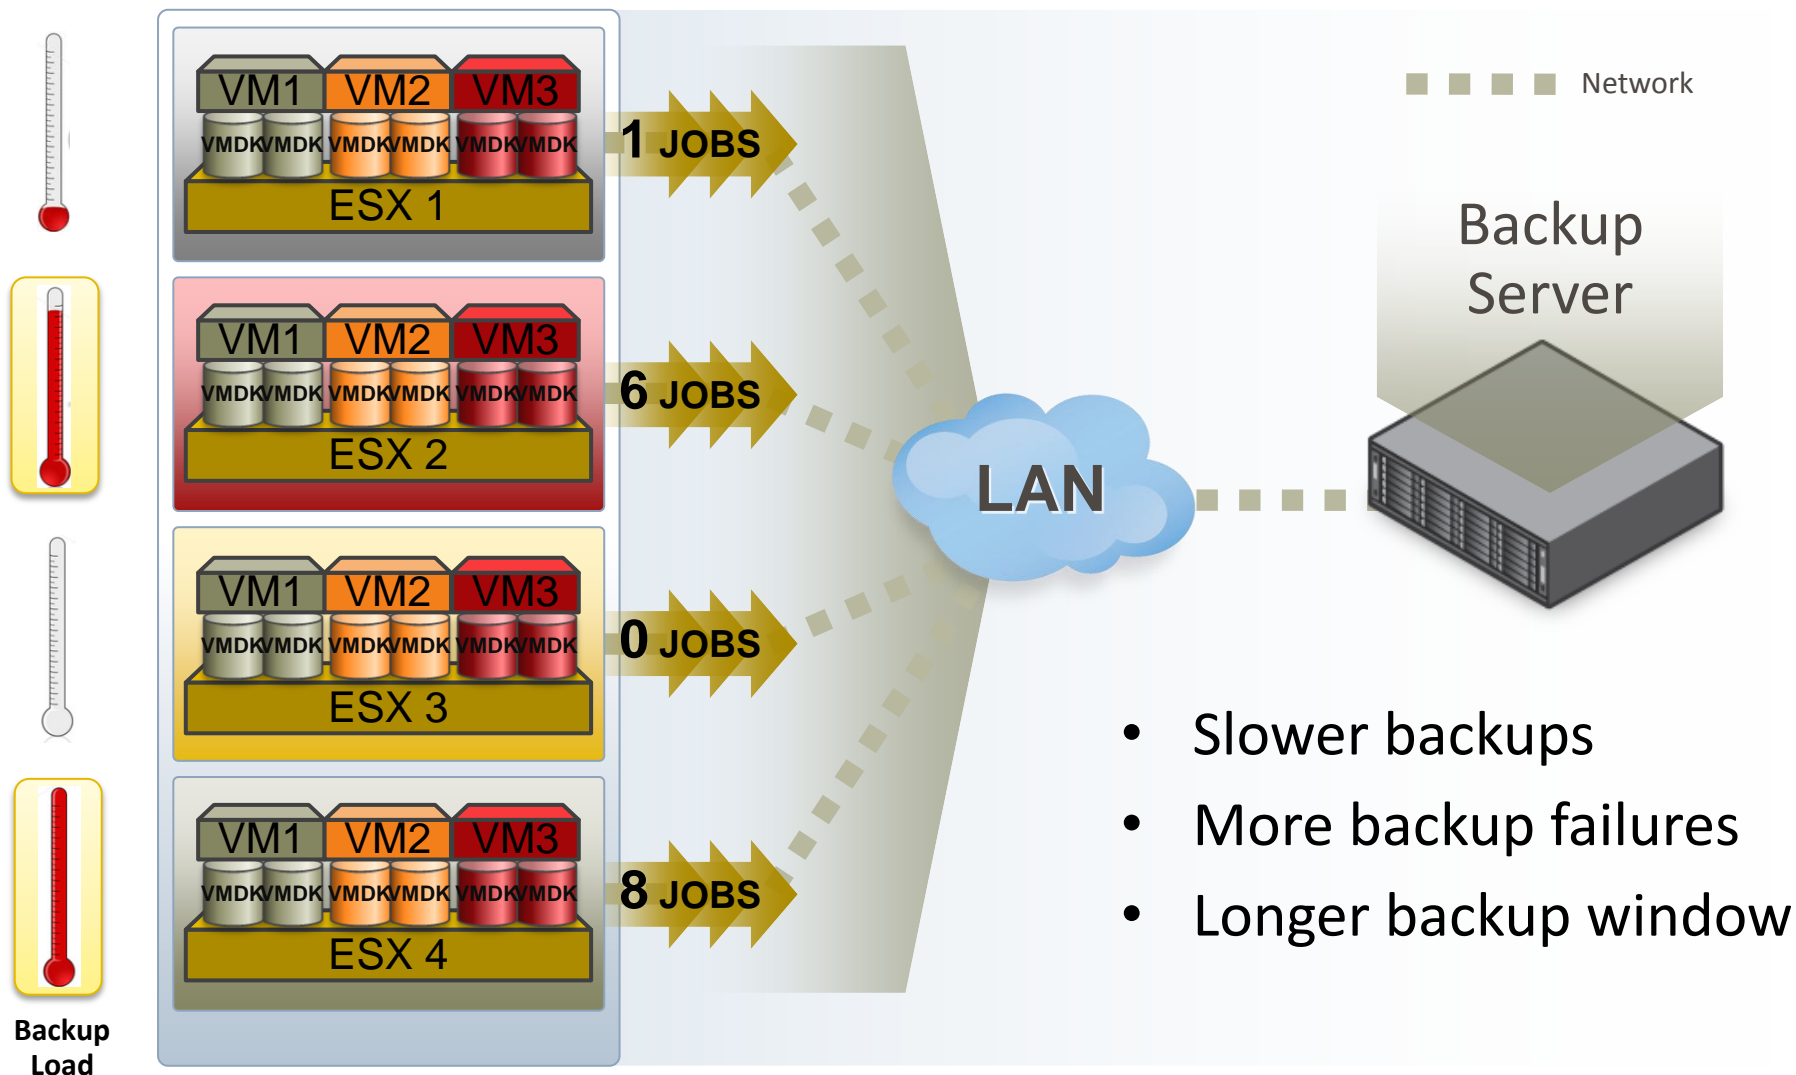
## VMware Intelligent Policy Overview

# VIP Designed for Two Major Tasks

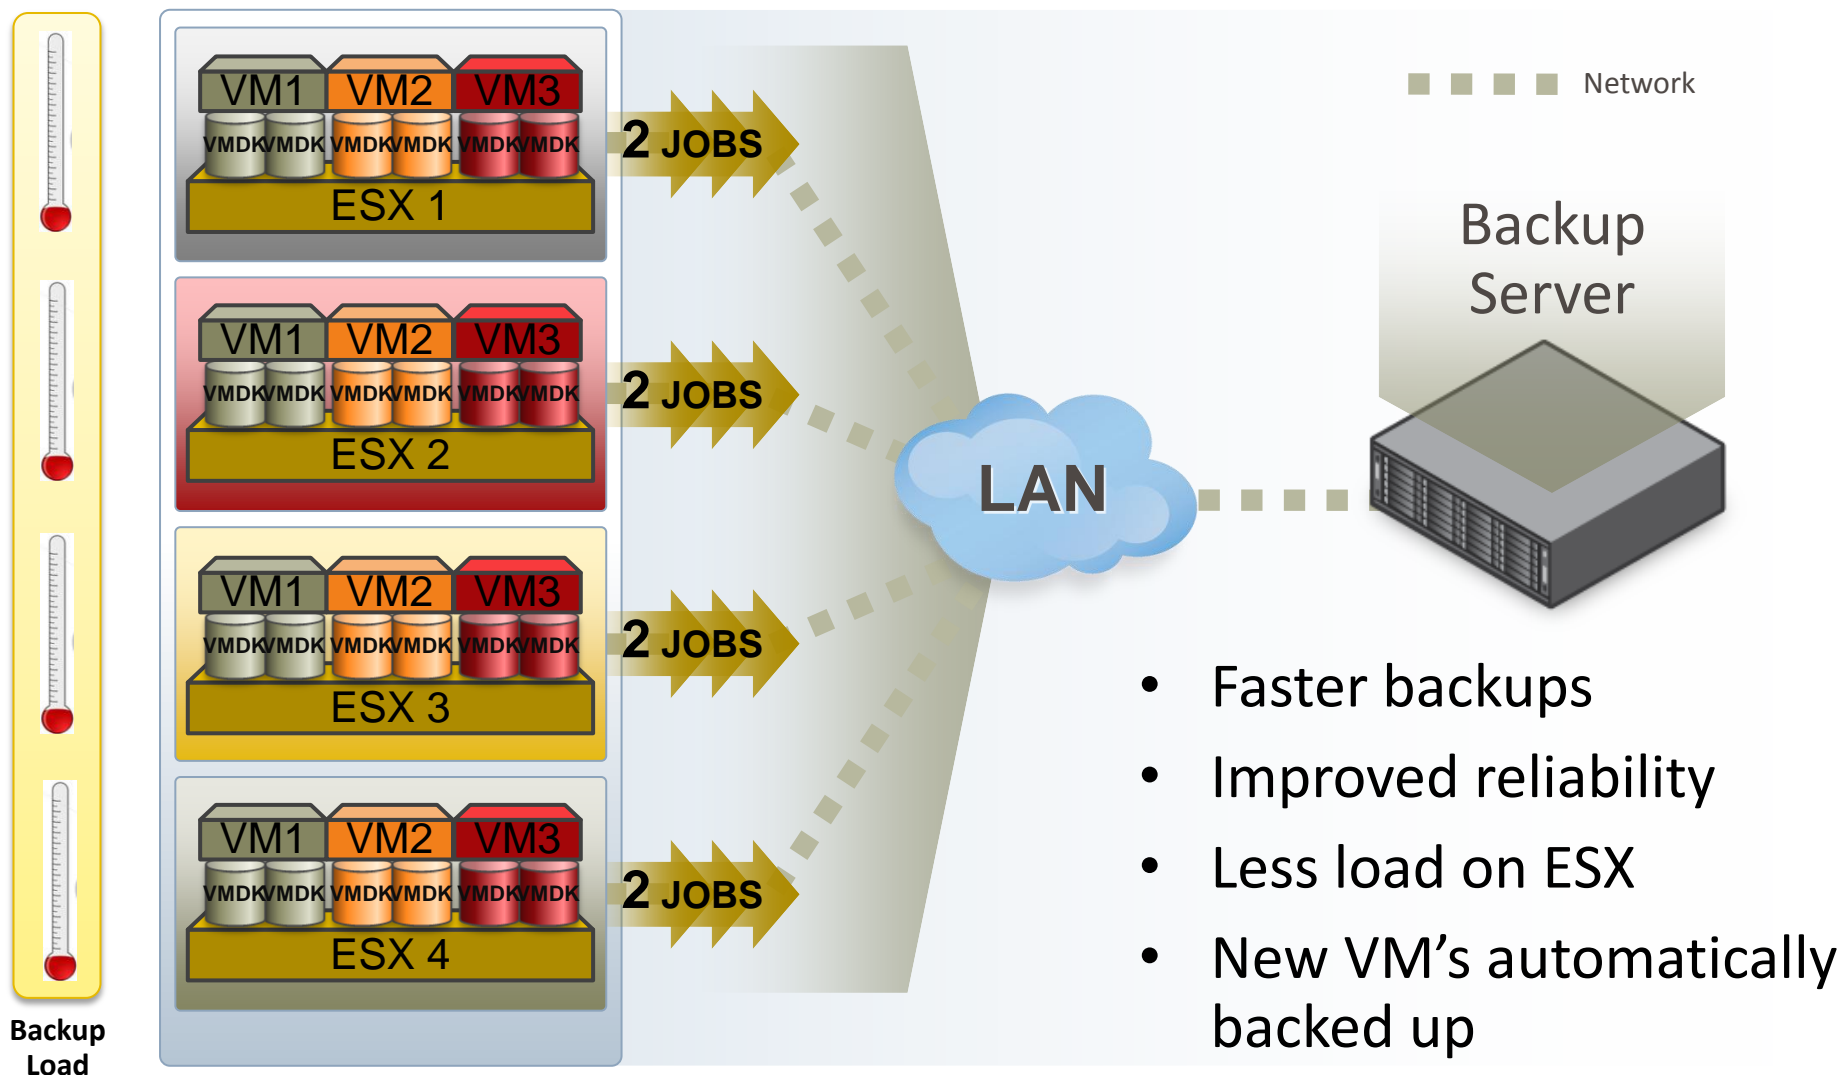
- 1) Automatically add and backup new and moved VM's
  - 2) Automatically balance backups across entire vSphere environment (Fibre or network)
- VMs protected based on physical location
    - ESX server
    - ESX Datastore
  - VMs protected based on logical attributes
    - vCenter folder
    - Resource pool
  - Backup performance is maximized with VIP by using physical location of VM to define backup



# Without VIP – Backup Activity Unbalanced

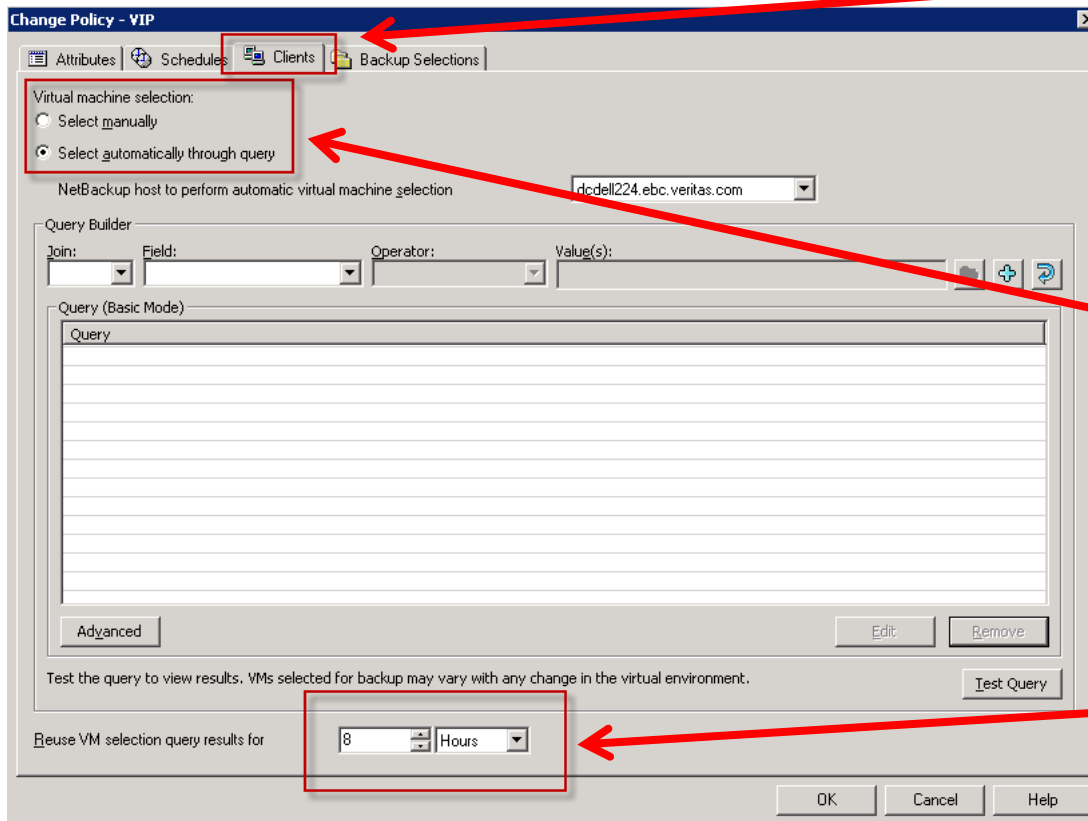


# With VIP – Automatic Load Balancing



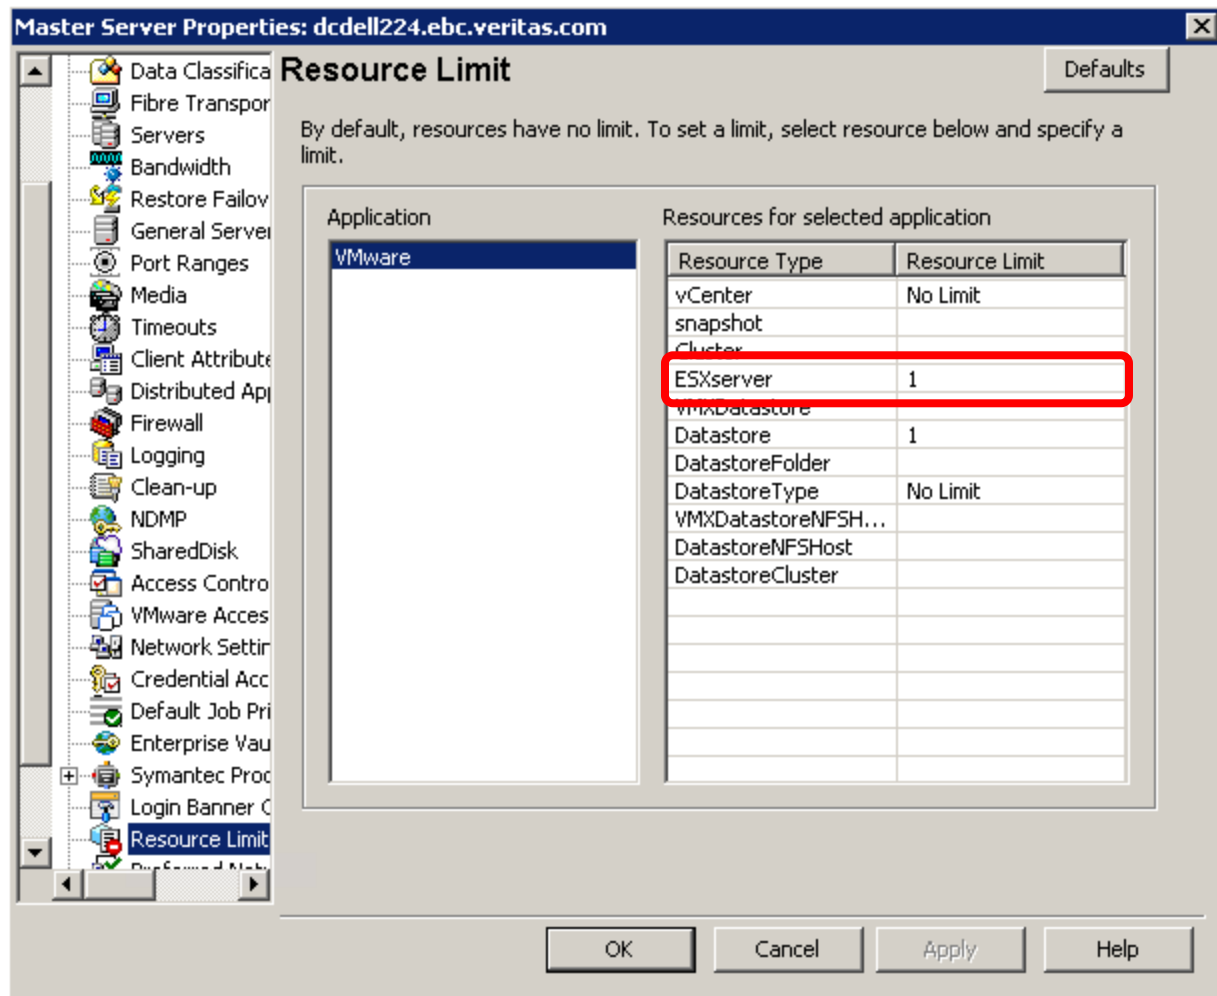
# VIP Query Generator Overview – Auto Select VMs

- VIP is selected via standard NetBackup policy client tab
- Traditional static or automatic (VIP) VM selection can be used
- VIP cache automatic update can be defined to minimize impact on vCenter server



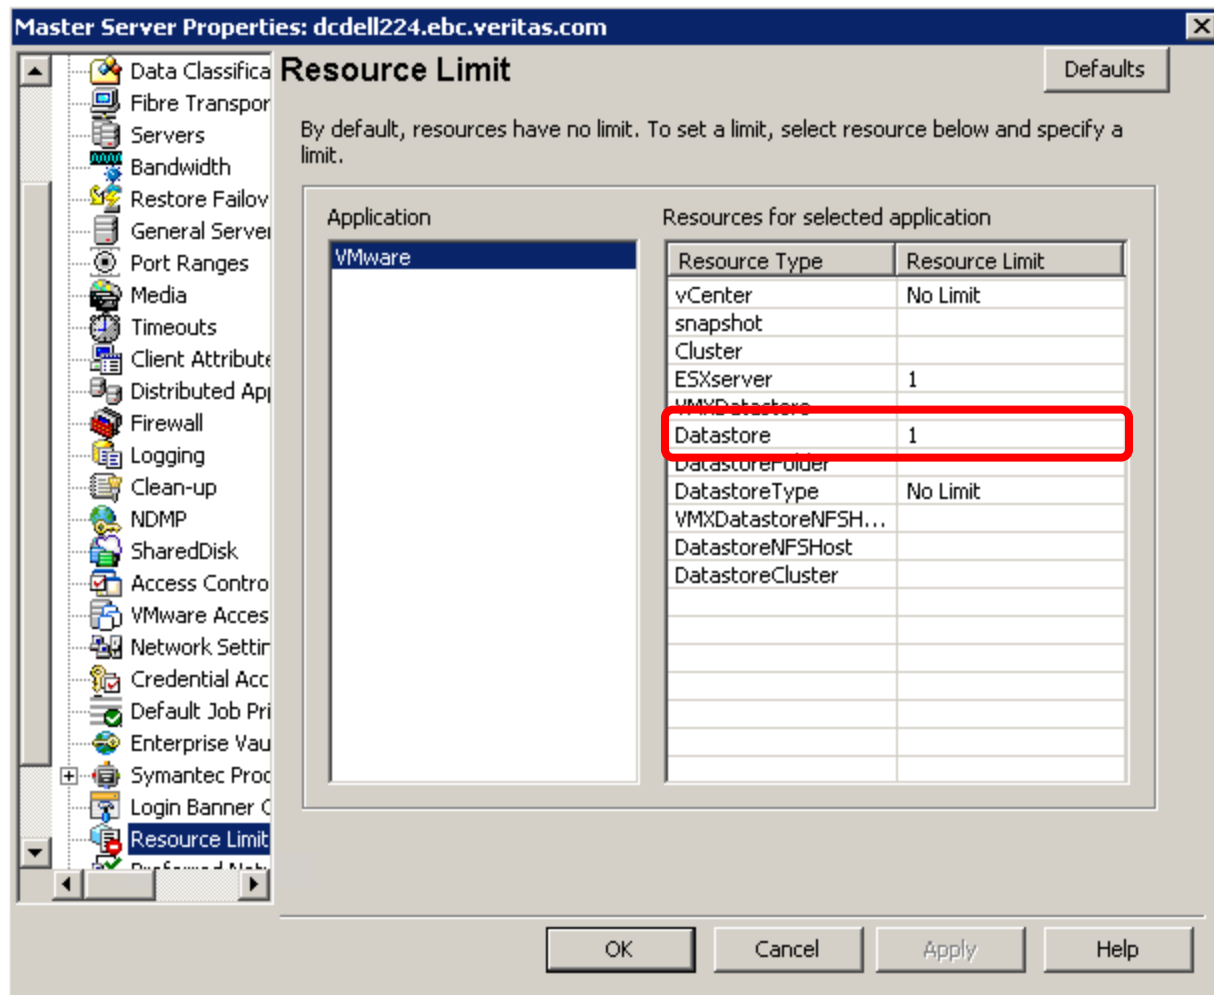
# VIP Resource Limits – Limit Jobs Per ESXi or Datastore

- These are global settings
- Two most commonly used settings:
  1. ESXserver (NBD transport)



# VIP Resource Limits – Limit Jobs Per ESXi or Datastore

- These are global settings
- Two most commonly used settings:
  1. ESXserver (NBD transport)
  2. Datastore (SAN transport)








**vision**

## VMware Intelligent Policy Use Cases

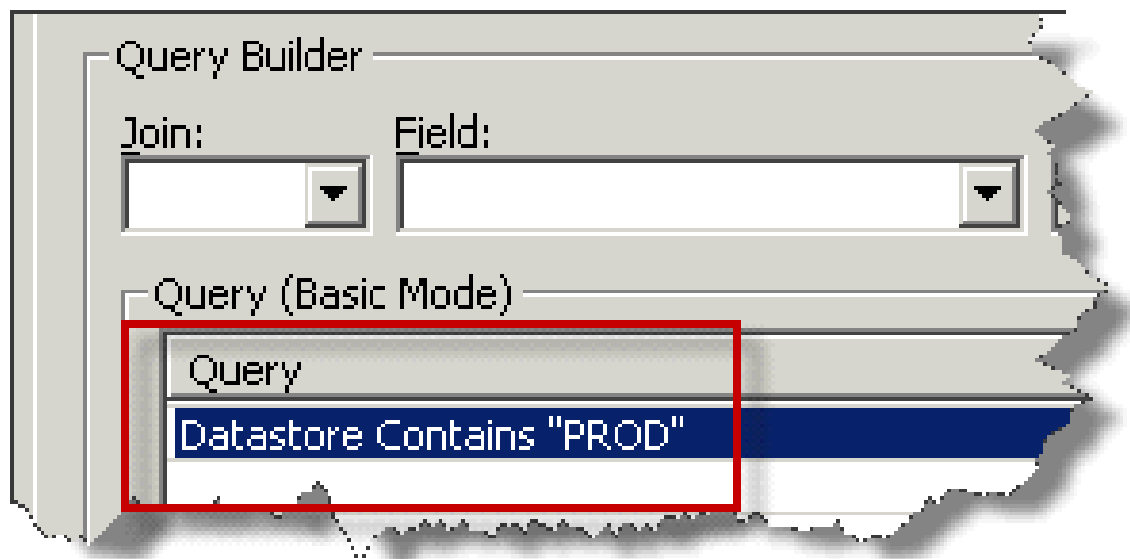
# Problem: Equalize Backup Load in SAN Environment

- Shared storage configuration
  - Fibre or iSCSI
  - Near zero impact on ESX
  - No loss in backup or restore functionality
  - No communication with ESX required



# Problem: Equalize Backup Load in SAN Environment

- Manage SAN (iSCSI - shared storage) backups at Datastore level
- Managing backups at ESX server does not balance load at storage level

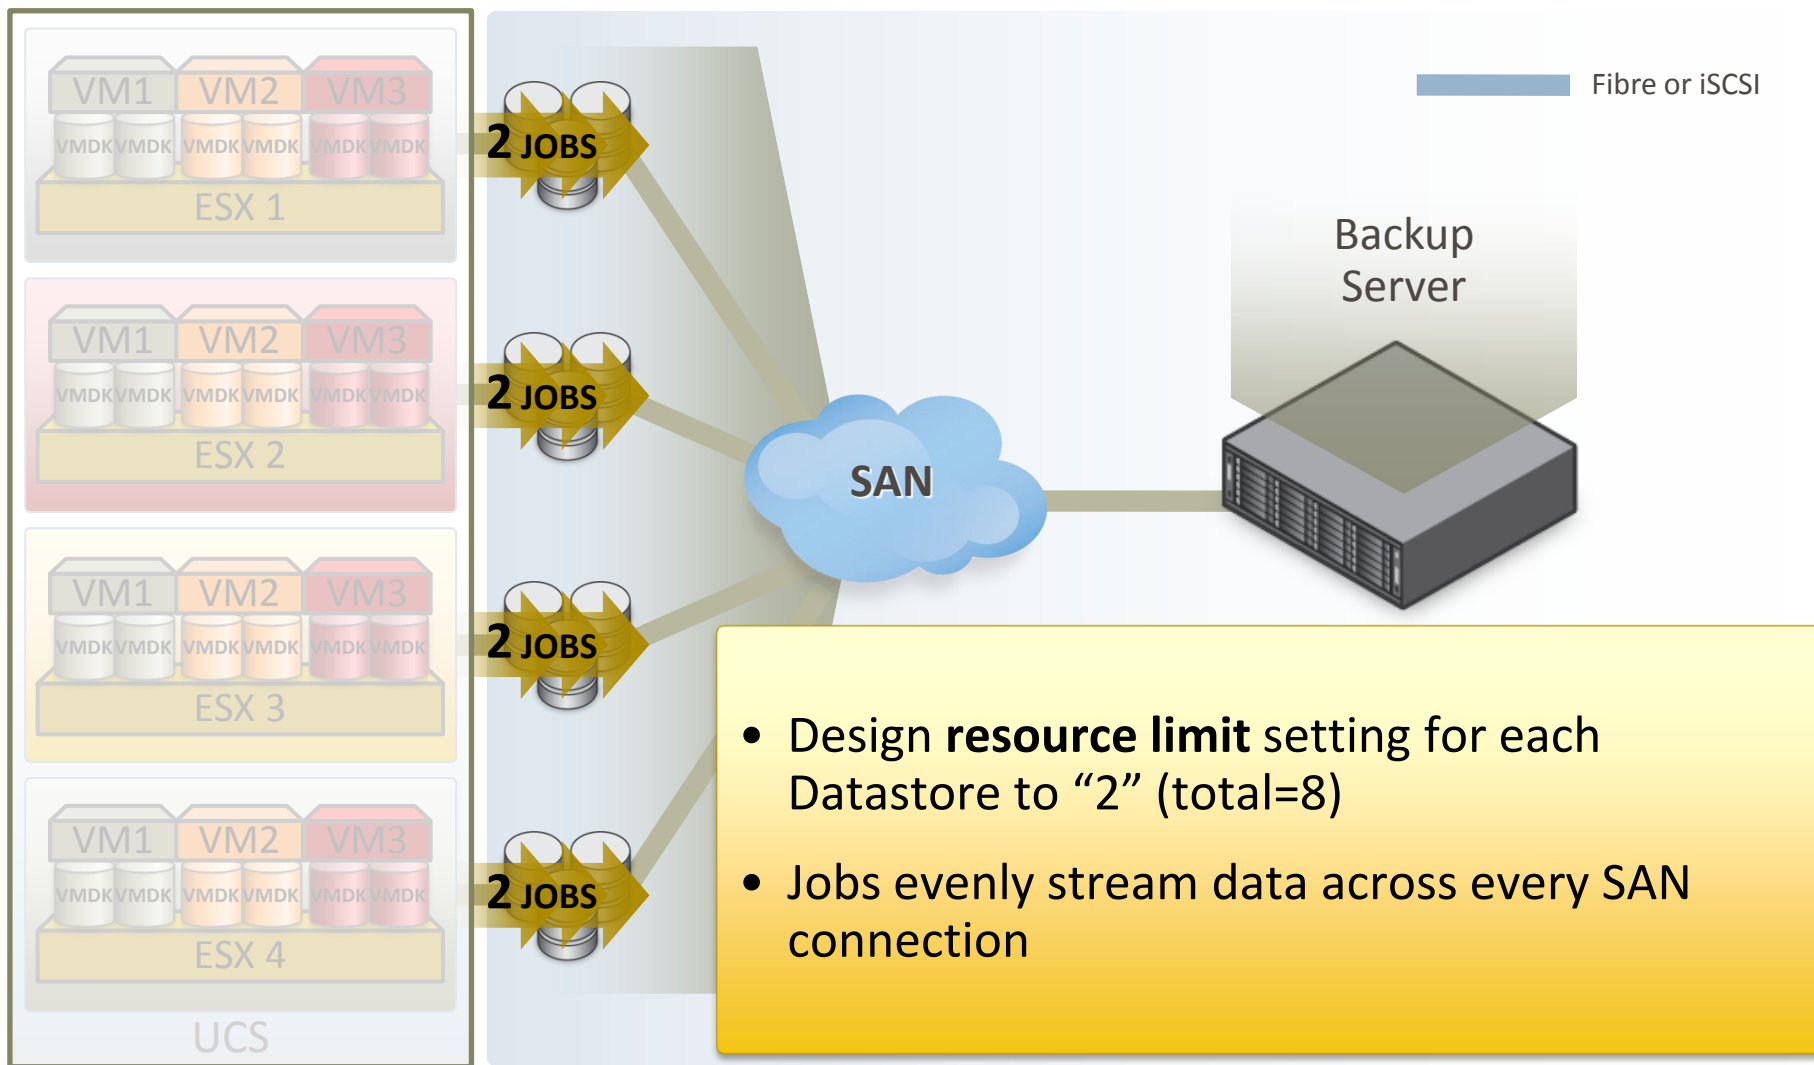


**Solution:** Will protect every existing and new VM on any Datastore with “PROD” in name

Fully compliant with Storage vMotion


Set “Datastore” resource limit to balance backup load across all Datastores

# SAN Backups - Equalize Load at Datastore



# Problem: Network backups – Equalize Load at ESX NIC

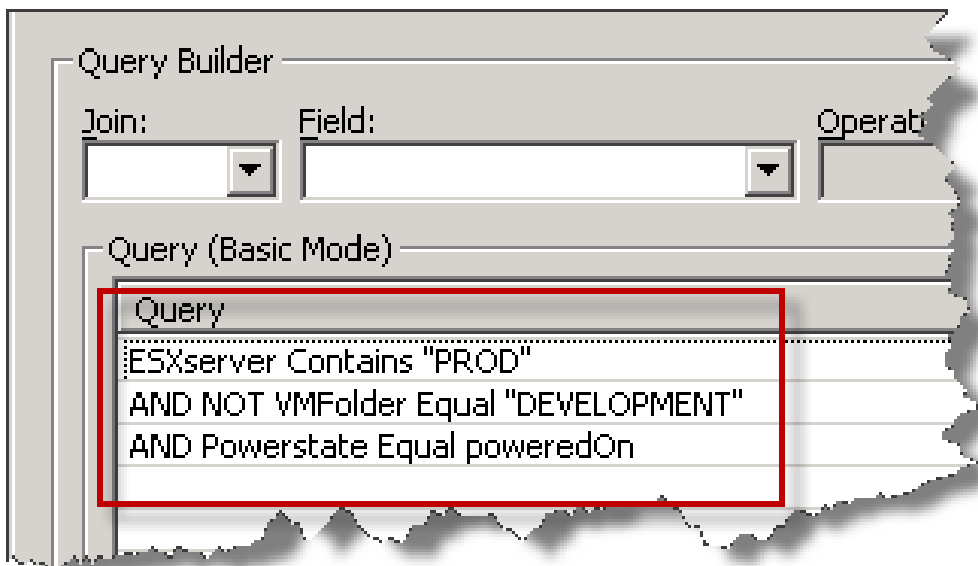
- Network Backups (NBD)
  - NFS or DAS fully supported
  - No loss in backup or restore functionality
  - Direct communication with ESX server required (DNS, etc.)
  - ESX server directly impacted



# Problem: Network backups – Equalize Load at ESX NIC

- All backup traffic over network – no shared storage available
- Don't want to saturate ESX network interface

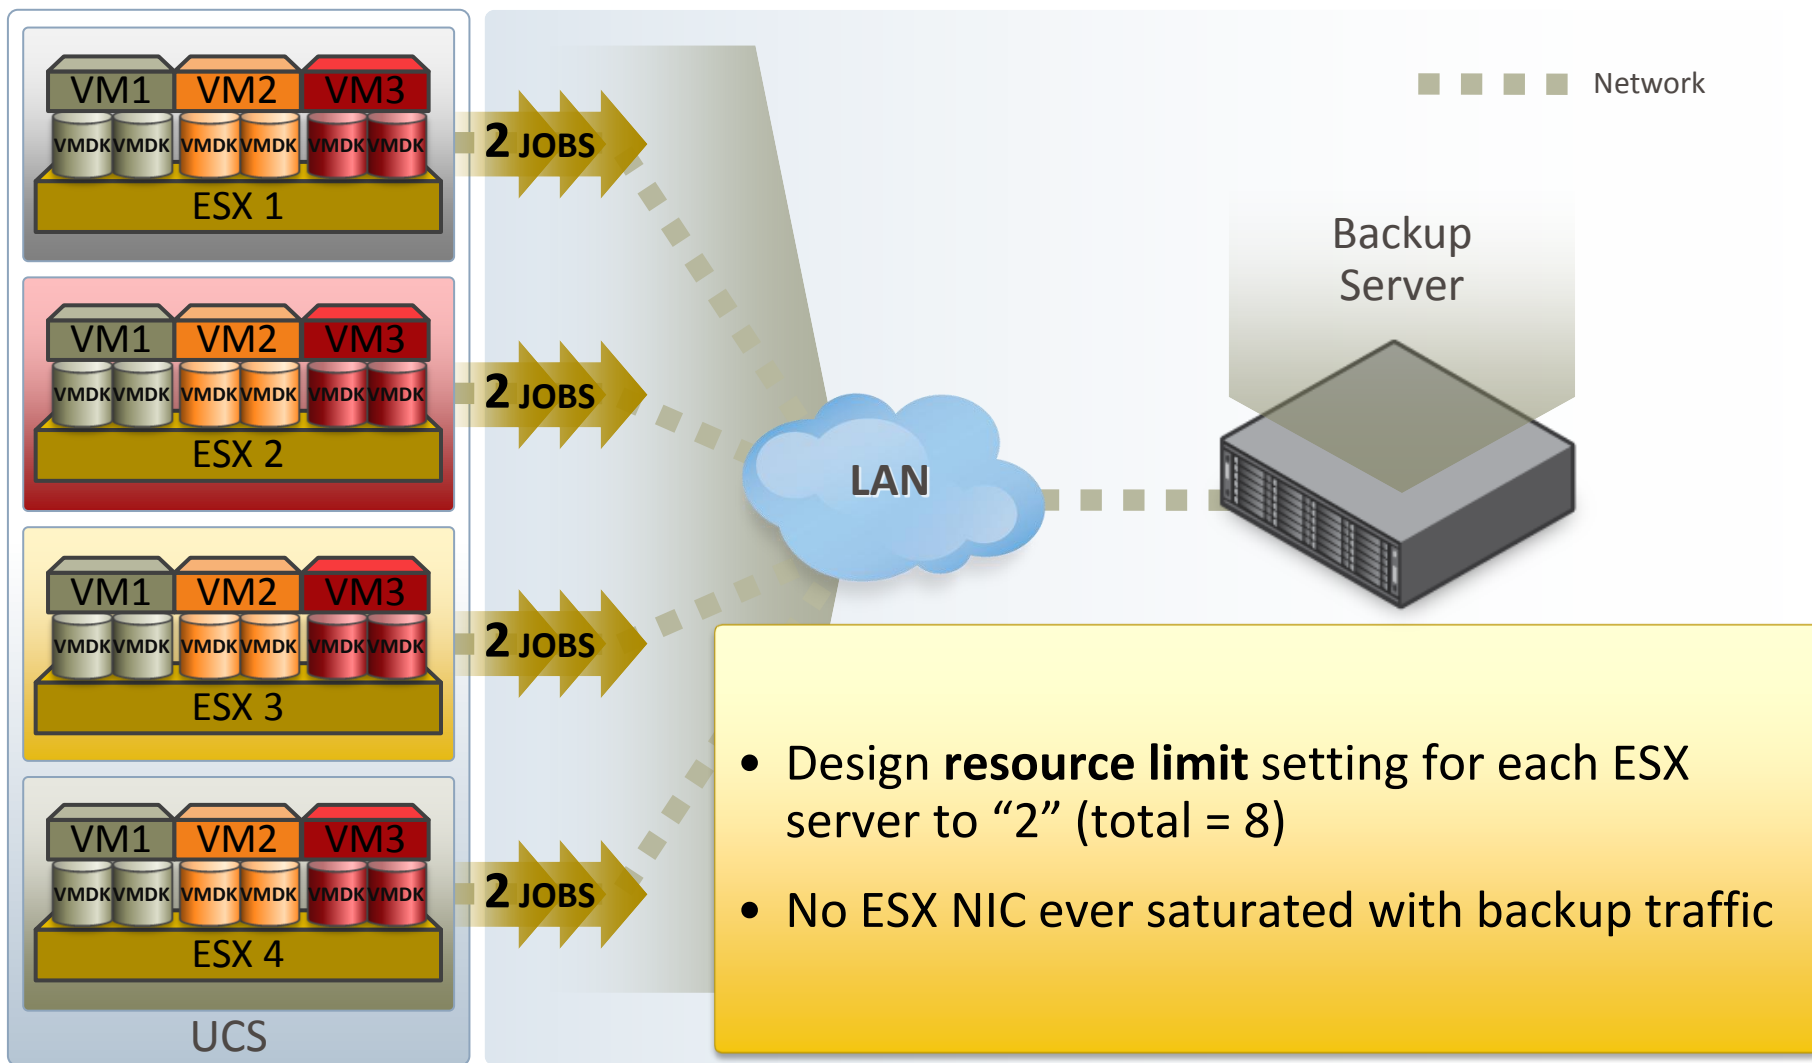
**Solution:** All powered on VMs on any ESX server in “Production” will be protected.



Unimportant VMs in the “DEVELOPMENT” folder (only used for development) will be skipped

Set “ESX” resource limit to balance backup load across all ESX NICs

# Network backups – Equalize Load at ESX NIC





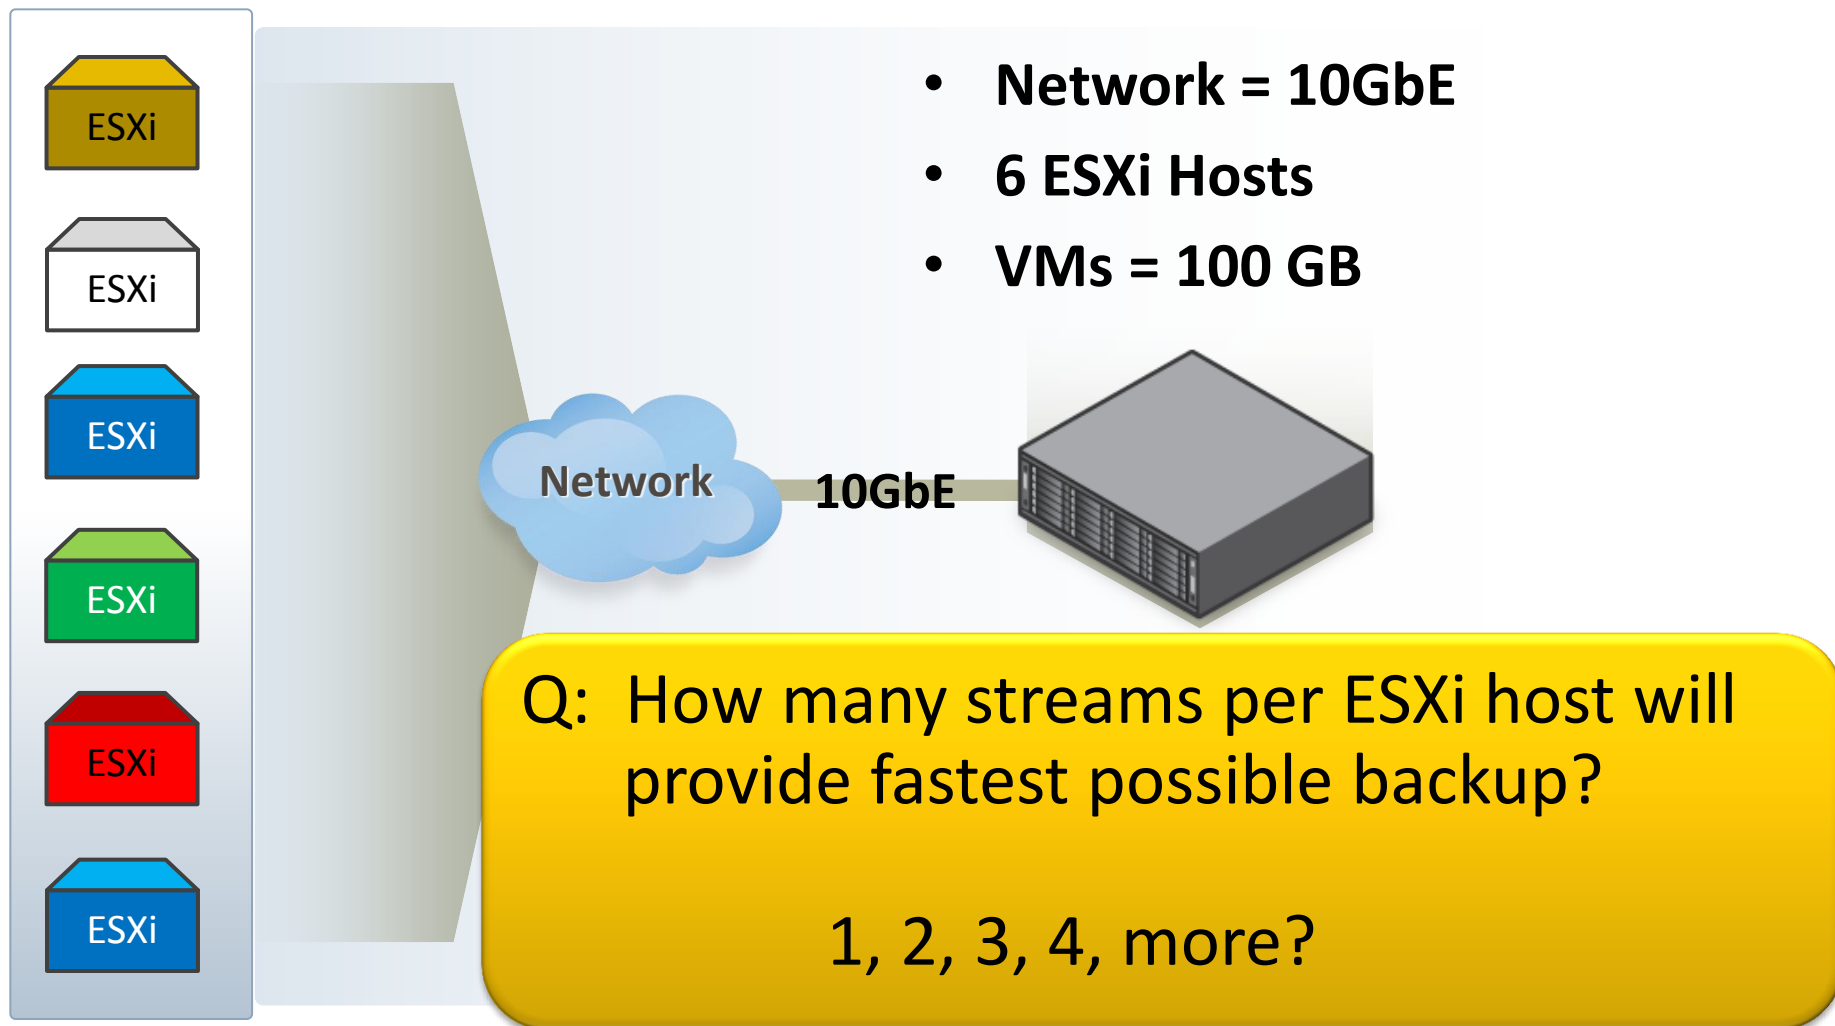


**vision**


## VIP Configuration Exercise



# VIP Configuration Exercise – NBD Transport



# NBD Transport – How Many Streams?



**1 Stream / ESXi Host  
(6 streams total)**

**Network = 10 GbE**

**100 MB/sec per ESXi**

**Each VM = 100 MB/sec**

**Backup 21 VMs / Hour**

**3 Streams / ESXi Host  
(18 streams total)**


**Network = 10GbE**

**100 MB/sec per ESXi**

**Each VM = 33 MB/sec**

**Backup 21 VMs / Hour**

# VIP Configuration Example – NBD Transport



# NBD Transport – Impact Of Snapshots Per Datastore

Why 45 GB for 3 simultaneous backups?


- Each VM backs up at slower rate
- Snapshot for each VM is open 3 times as long
- Time to capture data change in redo is longer
- Amount of data to re-apply to VMDK larger


**5 GB**

Per Datastore

**???**

Per Datastore





**Thank you!**

Jim Olson

George Winter

SYMANTEC PROPRIETARY/CONFIDENTIAL – INTERNAL USE ONLY  
Copyright © 2013 Symantec Corporation. All rights reserved.