Windows Server 2012 – networking and security

Tomáš Kantůrek

NIC Teaming with Load Balancing

Technical Overview

NIC Teaming

Built-in to Windows Server 2012, also known as Load Balancing/Failover (LBFO)

- Allows multiple network interfaces to be placed into a team for the purposes of:
  - Bandwidth aggregation, and/or
  - Traffic failover to prevent connectivity loss in the event of a network component failure

- Available in Windows Server 2012 in all SKUs (ServerCore and Full Server) versions
  - not available in Windows 8 Client SKUs; however, Remote Server Administration Tools can be installed on Windows 8 in order to manage NIC Teaming on the servers

Architectural Components

2 basic sets of algorithms for NIC teaming

- Switch-dependent modes
  - Require the switch to participate in the teaming
  - Types
    - Generic or static teaming
    - Dynamic teaming (LACP)
- Switch-independent modes
  - Do not require the switch to participate in the teaming

Traffic distribution methods

- Hyper-V switch port
- Address Hashing (TransportPorts)

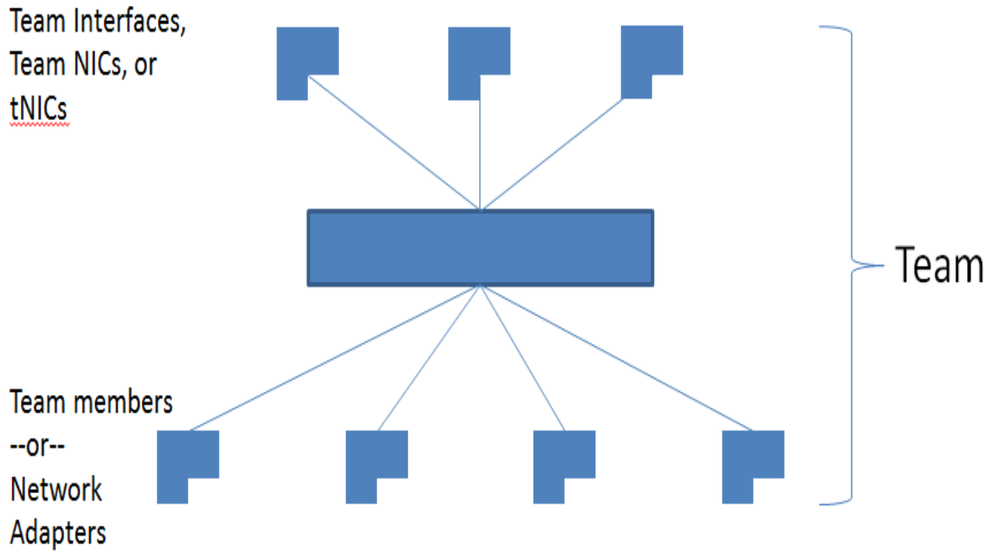

Requirements

- 1 NIC to be used for VLAN traffic
- At least 2 NICs for all modes that provide fault protection through failover
- Up to 32 NICs per team

NIC Teaming in VMs

NIC Teaming in Windows Server 2012 is supported in a VM

- Virtual network adapters that are connected to more than one Hyper-V switch can still have connectivity even if the network adapter under that switch gets disconnected
  - Useful when working with SR-IOV
- Each Hyper-V switch port associate with a VM that is using NIC Teaming must be set to allow Teaming in the host (parent partition) using PowerShell with administrative permissions:

```
Set-VMNetworkAdapter    -VMName  <VMname>    -AllowTeaming
```

- Teams created in a VM can only run in Switch Independent configuration, Address Hash distribution mode
- Only teams where each of the team members is connected to a different Hyper-V switch are supported
- Each Hyper-V switch port that is associated with a virtual machine that is using Teaming must be set to allow MAC spoofing
- Hyper-V NICs exposed in the parent partition (vNICs) must not be placed in a Team

Interactions with Distribution Modes

	All Address hash modes	Hyper-V Switch Port mode
Switch Independent	• Outbound traffic is spread across all active members. • Inbound traffic (from beyond the subnet) arrives on only one interface (primary member). If primary member fails another team member is selected as primary and all inbound traffic moves to that team member.	• Outbound traffic is tagged with the port on the Hyper-V switch where it originated. All traffic with that port tag is sent on the same team member. • Inbound traffic destined for a specific Hyper-V port will arrive on the same team member that the traffic from that port is sent out on.
Switch Dependent (Static and LACP)	• Outbound traffic is spread across all active members. • Inbound traffic will be distributed by the switch's load distribution algorithm.	• Outbound traffic is tagged with the port on the Hyper-V switch where it originated. All traffic with that port tag is sent on the same team member. If a team is put in the Hyper-V switch port distribution mode but is not connected to a Hyper-V switch, all outbound traffic will be sent to a single team member. • Inbound traffic will be distributed by the switch's load distribution algorithm.

Interaction with 3rd-Party Teaming Solutions

STRONGLY RECOMMENDED that no system administrator ever run two teaming solutions at the same time on the same server. The teaming solutions are unaware of each other's existence resulting in potentially serious problems.

- If the system administrator attempts to put a NIC into a 3rd party team that is presently part of a Microsoft NIC Teaming team the system will become unstable and communications may be lost completely

- If the system administrator attempts to put a NIC into a Microsoft NIC Teaming team that is presently part of a 3rd party teaming solution team the system will become unstable and communications may be lost completely

Network Virtualization

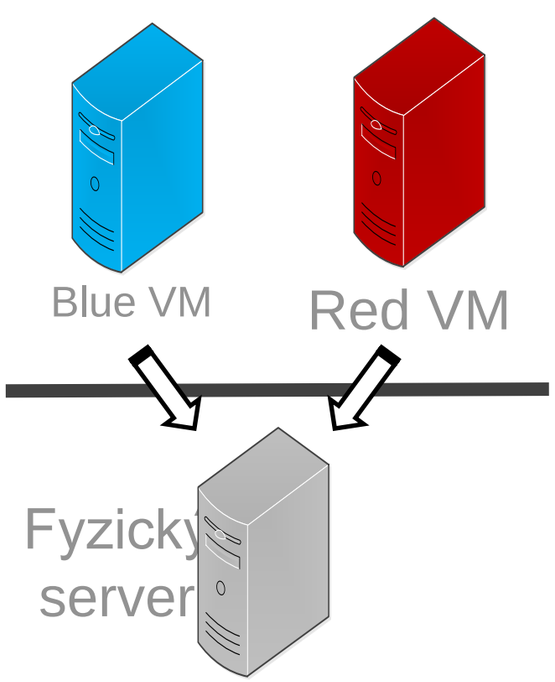

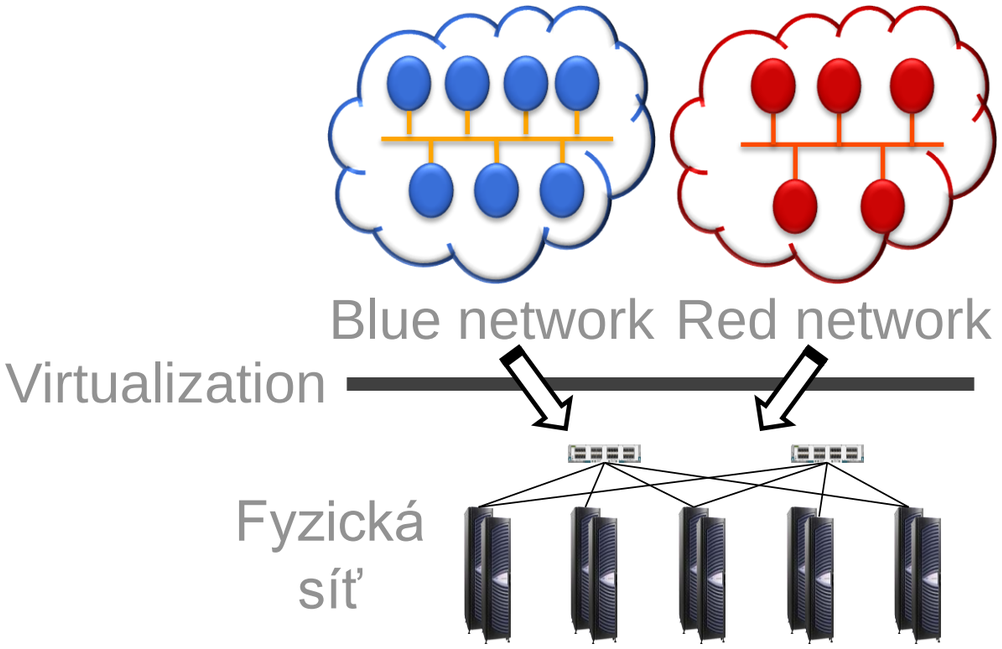

Serverová virtualizace

Spouštění mnoha virtuálních serverů na fyzickém stroji

Každý virtuál si připadá jako samostatný fyzický server

Síťová virtualizace

Běh mnoha virtuálních sítí na jedné fyzické infrastruktuře

Každá virtuální síť si připadá jako samostatná fyzická síť

Address Rewrite

Stejný formát TCP/IP packetu zajišťuje využití stávajících NIC Performance Offloads

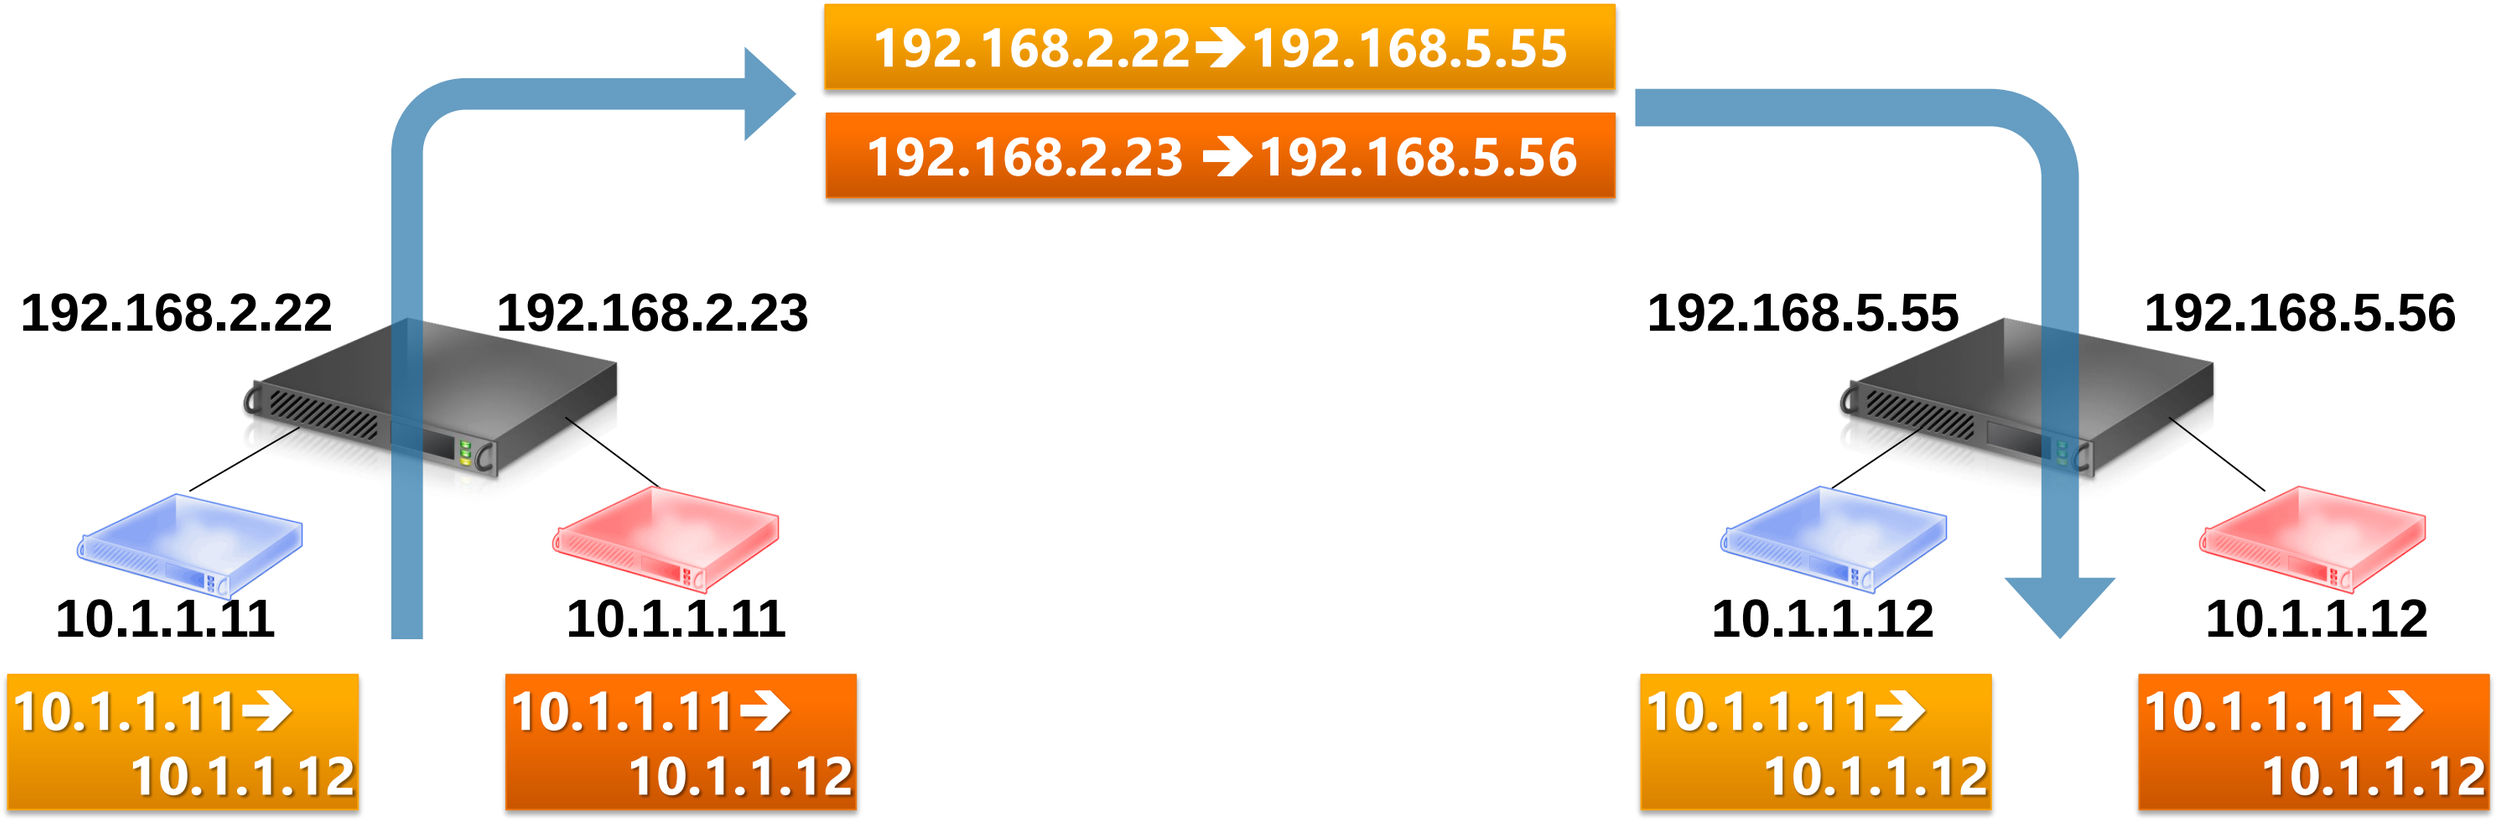

Encapsulation – NVGRE

Lepší využití sítě díky sdílení PA mezi VMs
Explicitní customer ID umožňuje lepší rozlišení zákaznických sítí

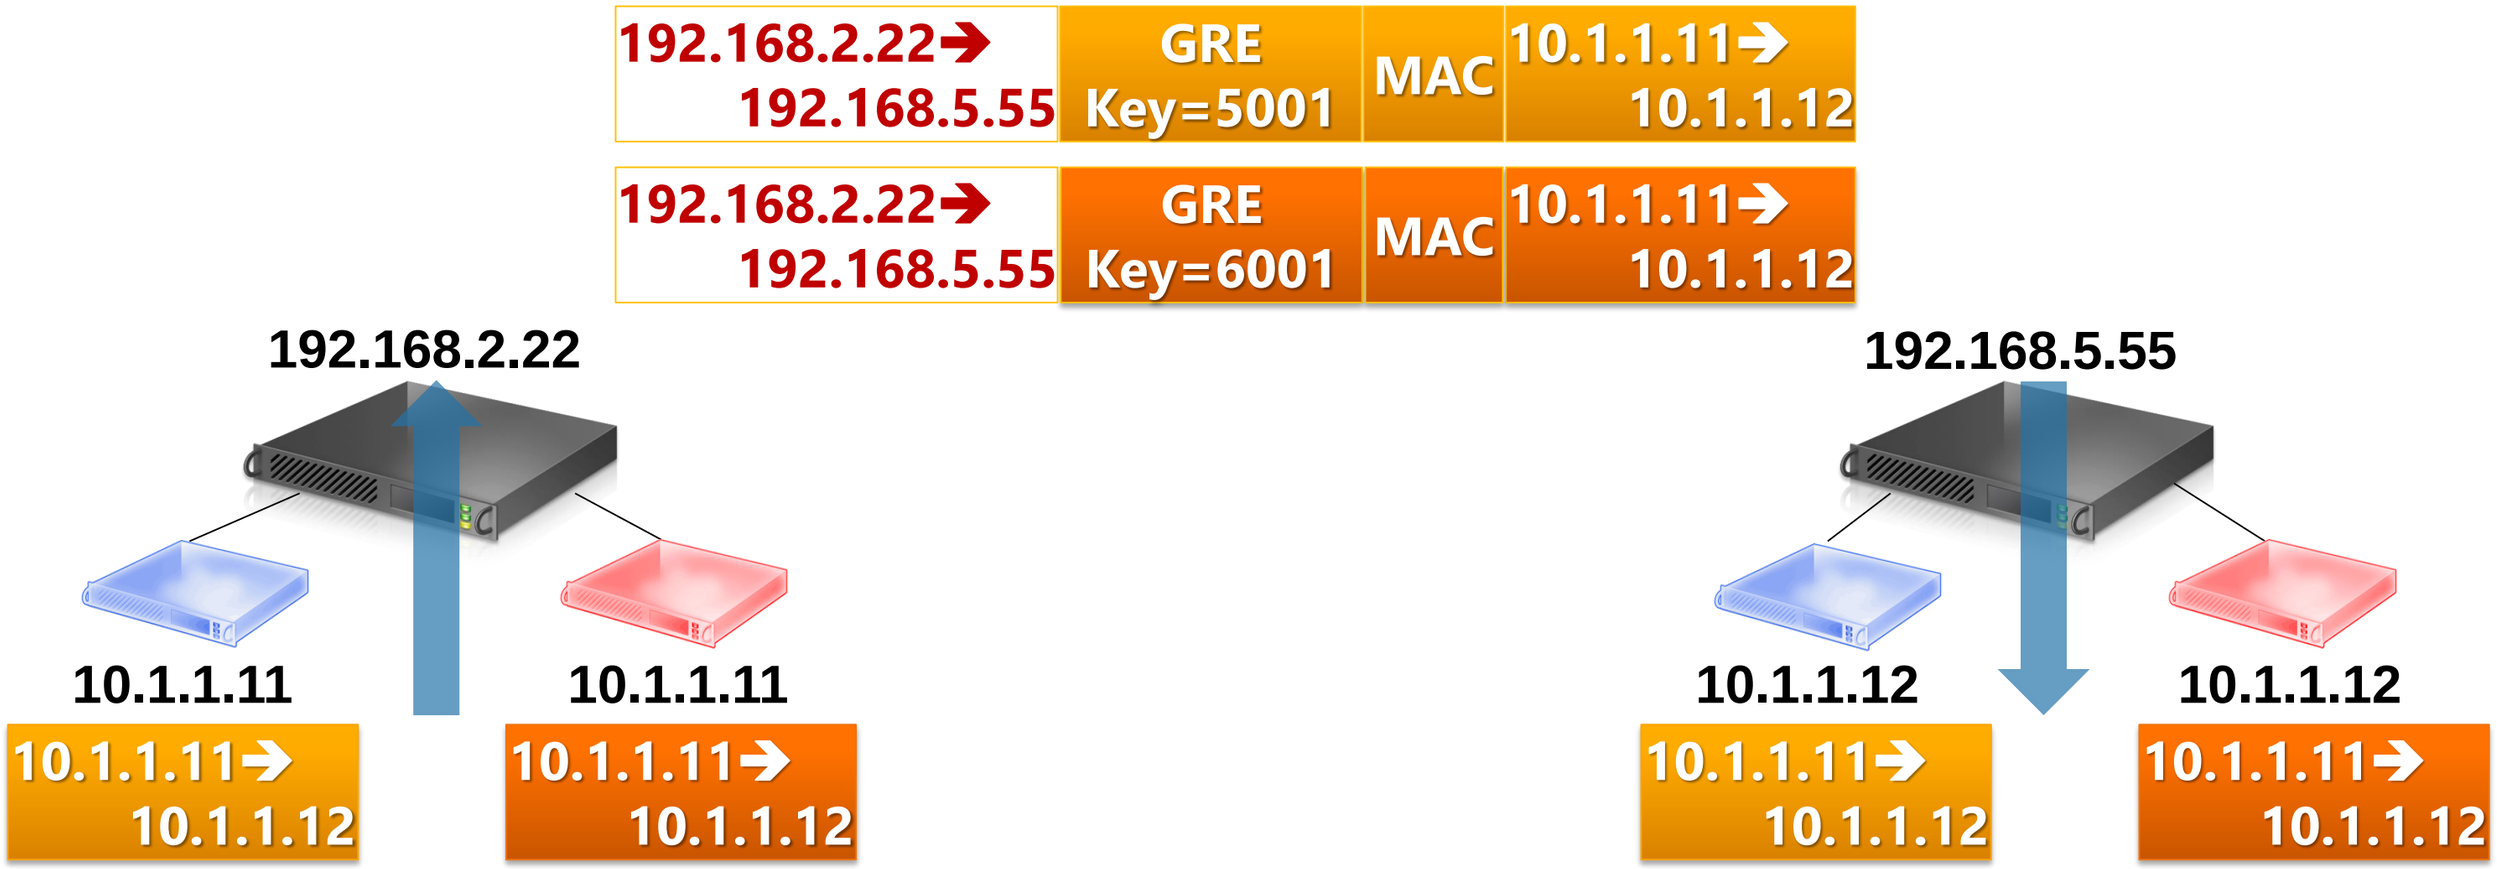

DHCP Failover

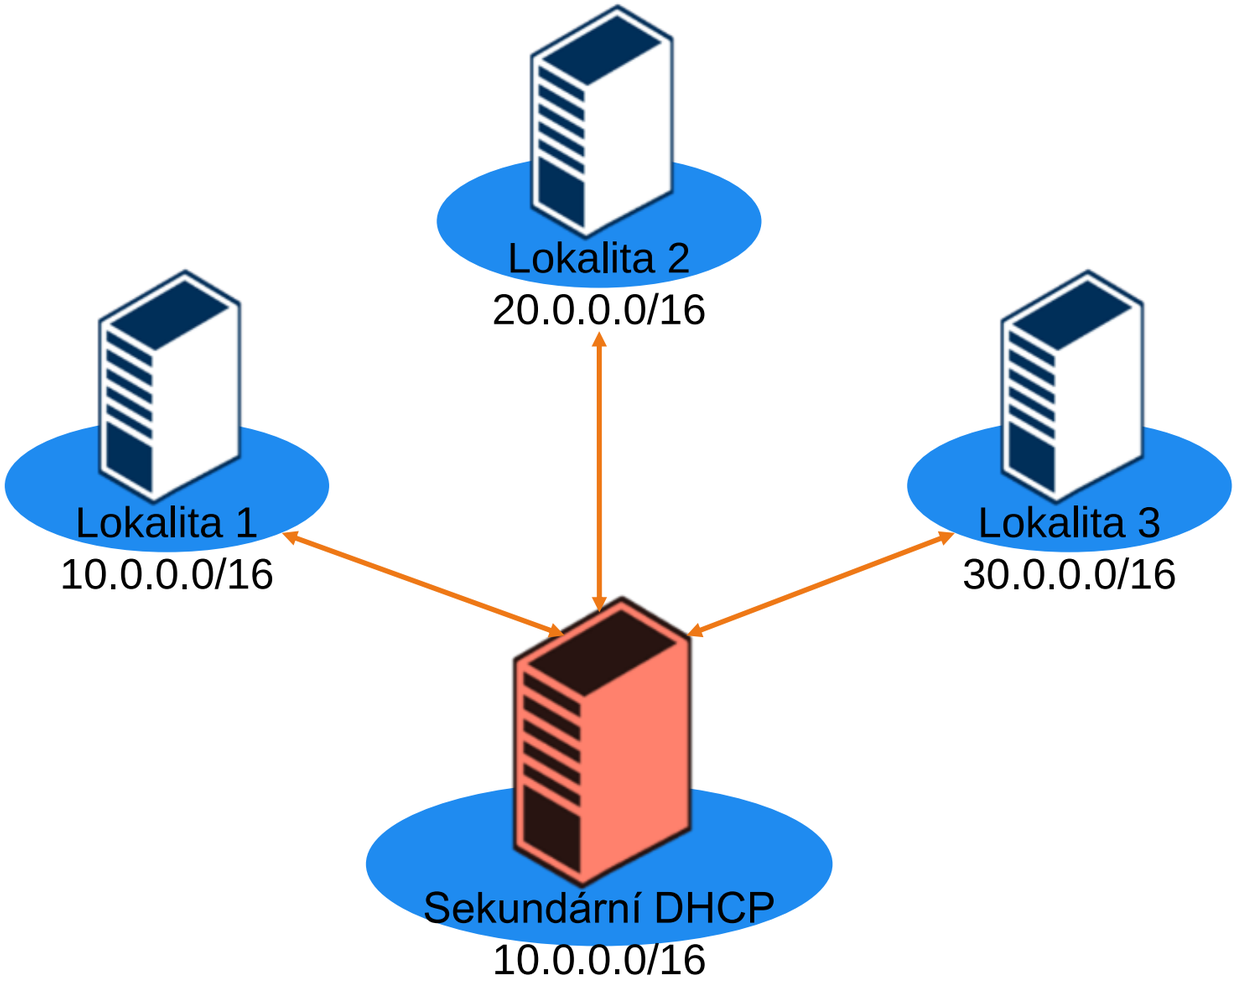

Vysoká dostupnost

DHCP bez clusteringu

Podpora více subnetů

Hyper-V Extensible Switch

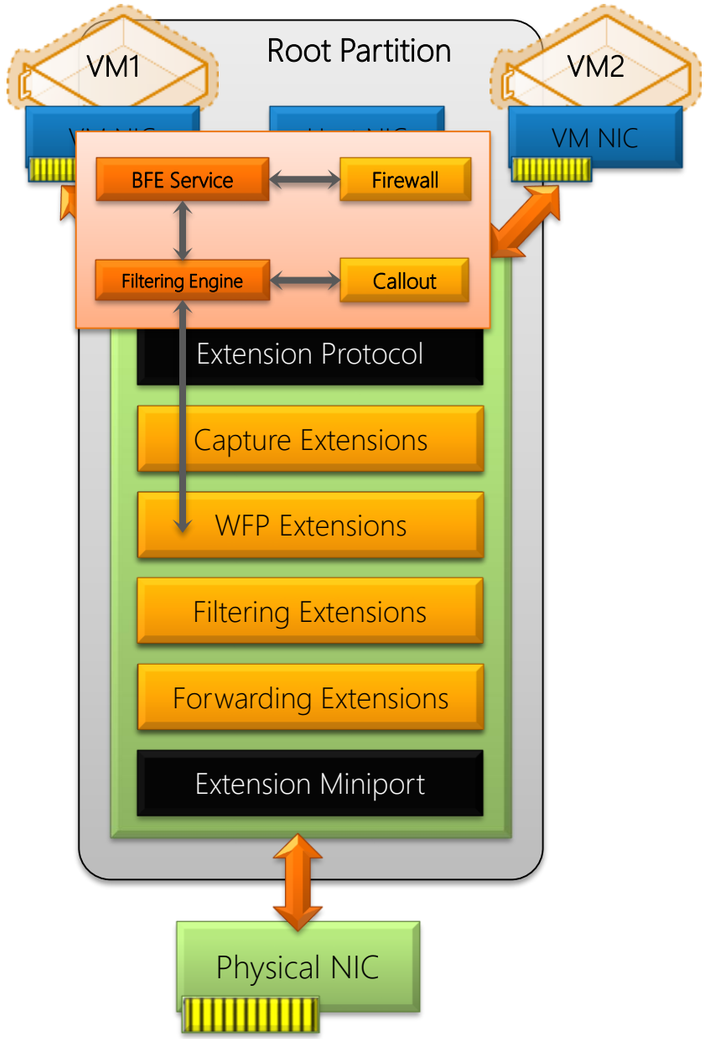

Například: Virtual Firewall by 5NINE Software

Například: sflow by inMon

CISCO BROADCOM NEC

5NINE SOFTWARE sFlow inMon

Minimum Bandwidth

Service	Minimum Bandwidth
VM 1	3 Gbps
VM 2	4 Gbps
VM 3	3 Gbps

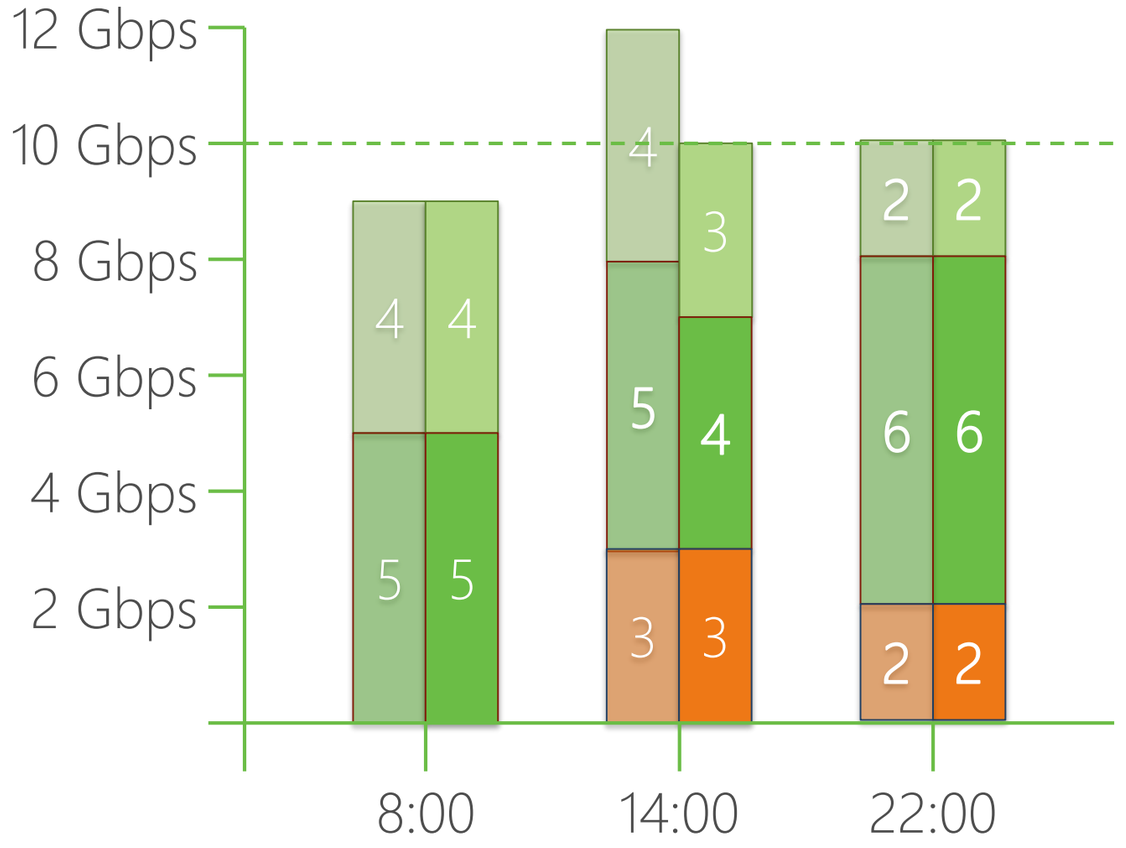

Benefit z predikce výkonu a vysoké zátěže spojení.

Síťová bezpečnost

DHCP guard

Router guard

Monitor mode

Extensible switch

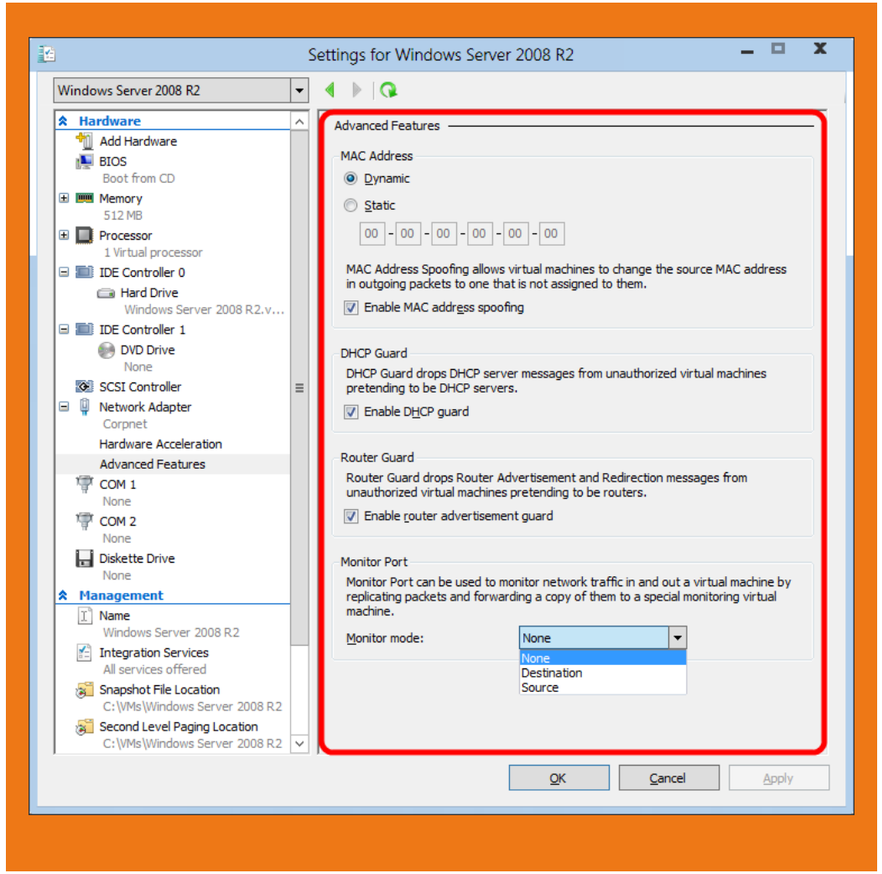

Network Adapter Hardware Acceleration

Virtual Machine Queue (VMQ)

- Employs hardware packet filtering to deliver packets from an external VM network directly to VMs using DMA transfers

IPsec task offload

- Reduces the load on the system's processors by IPsec encryption/decryption using a dedicated processor on the network adapter

Single-Root I/O virtualization (SR-IOV)

- Enables a device to divide access to its resources among various PCIe hardware functions

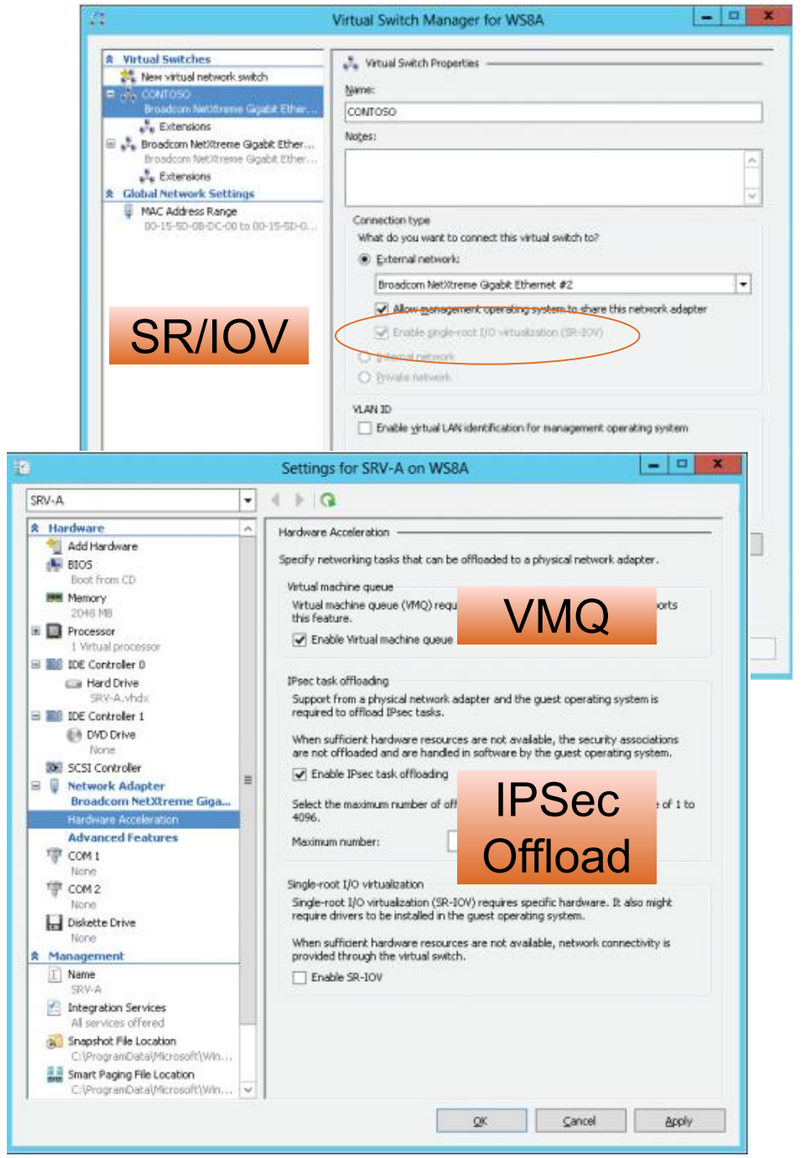

Resource metering

- Využití Resource pools
- Plně kompatibilní s veškerými funkcemi Hyper-V
- Data jsou migrována s VM
- Network Metering Port active control lists (ACLs)

- Průměrné využití CPU
- Paměť
  - Průměrné využití
  - Minimum využití
  - Maximum využití
- Maximální alokovaný diskový prostor
- Síť
  - Příchozí provoz
  - Odchozí provoz

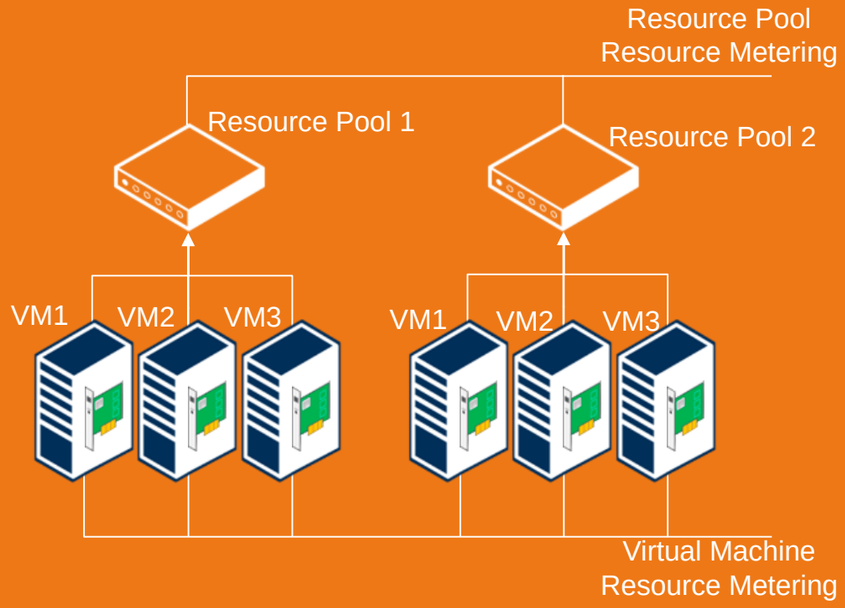

Storage a bezpečnost

Šifrované cluster disky

Šifrování pomocí BitLocker

Podpora pro tradiční failover disky

Podpora pro Cluster Shared Volumes

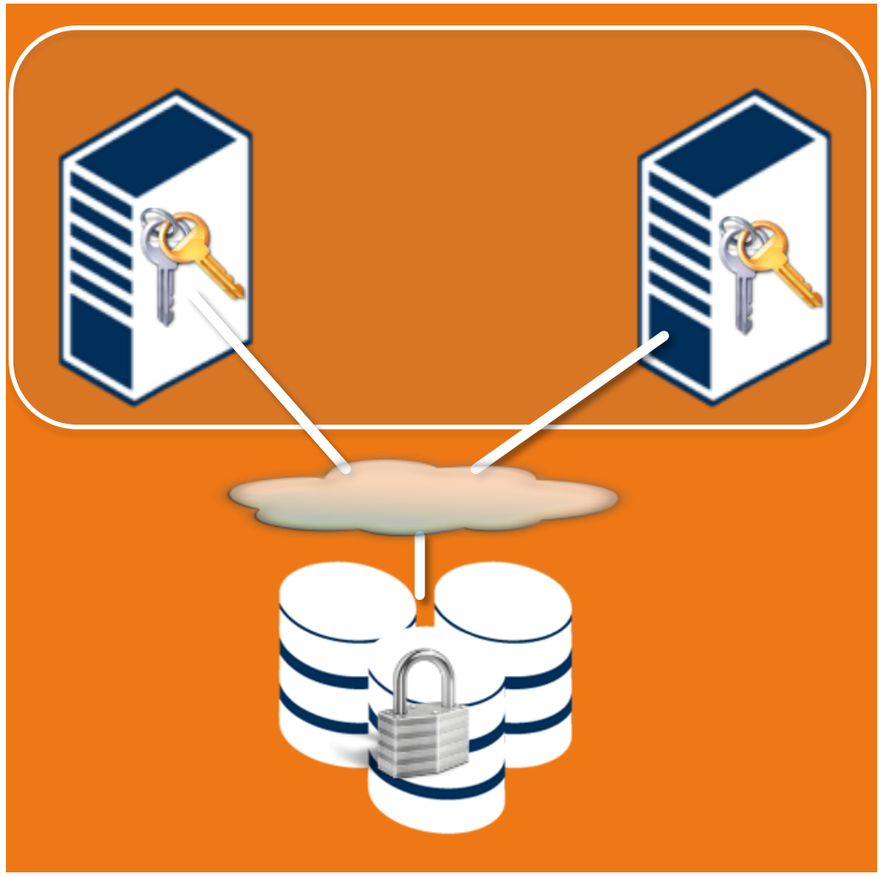

BranchCache

BranchCache
- Small cache block size reduces network bandwidth requirements
- Intelligent data compression
- Encryption on cache
- More scalable

Printing
- The document is sent directly to the local printer, while print request only is routed to the print server in datacenter

BENEFITS

- Users in the branch office can download and print documents faster
- Frees up network bandwidth
- Saving costs – support more people in branch offices with same hardware; no WAN Optimizers needed

BranchCache

Spend time working, not waiting

Users access data from local caches
- Distributed cache mode
- Hosted cache mode

Improved network and delivery performance
- Pre-load or distribute content
- Reduce print-file data on network

Deploy, manage, and scale with ease

Streamlined deployment and management
- Multisite deployment via a single GPO
- Client computer configuration is automatic; configurable via GPOs
- Delivers encrypted caches without PKI or hard drive encryption

Scales from home office to large locations
- Multiple hosted cache servers for large offices
- Extensible Storage Engine database technology

QUESTIONS & ANSWERS

THANK YOU