Virtualisation of Network Devices

Best Practice Document

Produced by the CESNET-led working group
on campus networking

Author: M. Pustka (CESNET)

June 2014

© CESNET 2014. All rights reserved.

Document No: GN3plus-NA3-T2-CBPD123
Version / date: Version 1.0, 15 October 2014
Original language : Czech
Original title: “Virtualizace síťových prvků”
Original version / date: Version 1.0, 1 January 2014
Contact: Martin.Pustka@vsb.cz

CESNET bears responsibility for the content of this document. The work has been carried out by a CESNET led working group on campus networking.

Parts of the report may be freely copied, unaltered, provided that the original source is acknowledged and copyright preserved.

The research leading to these results has received funding from the European Community's Seventh Framework Programme (FP7/2007-2013) under grant agreement n° 605243, relating to the project 'Multi-Gigabit European Research and Education Network and Associated Services (GN3plus)'.

Table of Contents

Table of Figures

Table of Tables

Executive Summary

IT infrastructures have changed significantly in the last few years, particularly in the area of virtualisation of server systems. This area is already relatively stable, but in recent times; virtualisation has also begun to penetrate into the area of network infrastructure.

The network infrastructure as a whole cannot be virtualised to the same extent as server systems, because it includes the physical part of the data centres themselves. Despite this, we can find applications that can be suitably virtualised as a whole or in part.

This document considers the implementation of selected network services in such a way that it will be possible to operate them in the environment of modern data centre virtual infrastructures.

This document describes the benefits of virtualisation, but also looks at the disadvantages – when it is better not to use virtualisation. Included here are the requirements for virtual infrastructure, since the network devices used here have slightly different requirements to classical virtual servers in terms of network integration.

We also consider the question of what parts of the infrastructure are suitable for virtualisation and what parts are not, and outline the benefits of virtualisation which are in cost-reduction and increased scalability of the infrastructure deployed.

1 Introduction

1.1 Basic terms

Some basic terms used in this document are briefly explained below.

Virtualisation infrastructures are infrastructures creating an environment for virtualisation. Most commonly this involves the products VMware vSphere, KVM, Hyper-V and XEN.

Virtual system (virtual machine) – sometimes the term virtualised system is also used – this is usually represented by a virtual server, virtual station or virtual router operated within the virtual infrastructure.

Virtual infrastructures or alternatively virtualised infrastructures are infrastructures which are created by virtual machines, form a functional whole and are installed in virtualisation infrastructures.

Virtual routers are virtual systems fulfilling the function of routers in virtual infrastructures.

1.2 Network infrastructure services

For our purposes we will consider network services which are today provided mainly by physical components (in particular hardware routers), as network infrastructure services. In general, network services can be divided into the following areas:

- Operational and routing support (e.g. BGP reflectors).
- Virtual routers connecting parts of the infrastructure.
- VPN concentrators for site-to-site connection or user access.
- Firewalls and gateways for virtual servers operated in virtualisation infrastructures.
- NAT/PAT devices.
- Load-balancing devices.

2 Virtualisation – Requirements and Advantages

The x86 hardware platform is increasingly used for physical network devices, so it has become an increasingly popular choice among manufacturers for implementing network infrastructures in virtual environments. The currently most popular and widely deployed virtualisation infrastructures include the VMware vSphere, Hyper-V, KVM and XEN environments.

In terms of network infrastructures modern virtualisation infrastructures provide these particular advantages:

- Scalability of technical facilities.
- Operational stability.
- Easy maintenance and updating.
- Easy hardware modification.
- Flexible network infrastructure.
- Cost savings.

The main disadvantages are:

- Unaccelerated network traffic without direct access to hardware.
- Usually lower maximum network throughput than that provided by dedicated physical devices.
- Virtual network devices are less efficient than dedicated network devices.
- Virtual devices are dependent on properly functioning other devices.

In general we can say that the requirements of virtualised network infrastructures are different from virtualised server infrastructures. Usually disk capacity, disk I/O performance and RAM demands are smaller, but CPU performance requirements predominate, along with the demand for computer network quality in terms of transfer delays (latencies), transmission rate and, in particular, flexibility. It is therefore good to have quality computer network administration and support in virtualisation infrastructures (e.g. VM-FEX or Nexus1000v technological concepts), which expand the basic options of the undistributed and distributed switches of virtualisation solutions in terms of function and performance.

2.1 Scalability

Network scalability is the ability to handle a growing amount of work in a capable manner, or an ability to be easily upgraded to accommodate that growth. With some network applications, this scalability allows a gradually increased performance to meet the growth in requirements or, conversely, to reserve capacity for when it's needed to cover (possibly temporary) increases in demands for service capacity.

The demands of network devices increase with the growth in network infrastructures (e.g. with the growth in routing tables) or growth in usage (the volume of traffic or the number of users).

An example where good scalability is important is with planned abrupt loading, when a surge in use is expected. In the case of universities, this happens when holding conferences or a larger quantity of remote VPN accessing in periods outside of lecture time. In these cases, capacity can be temporarily increased – in terms of CPU, memory or reserved network bandwidth.

2.2 Operational stability

If our virtualisation and physical infrastructure is well and redundantly designed and built, then another benefit of modern virtualisation infrastructures is their ability to cope with planned and unplanned downtime of the physical components of the infrastructures.

The shutting down of running services to connect physical cabling, program upgrades and the replacement or breakdowns of physical hardware can usually be avoided on virtualisation infrastructures. Another great benefit of virtualisation infrastructures that live migrations of virtual systems to other hardware can be performed without their outage, thus protecting QoS.

Various techniques and modes for high availability (HA) can be used successfully with virtualisation infrastructures, which are able to automatically handle outages of parts of virtualisation infrastructures (e.g. outages of the physical server on which the given virtual system is currently running) and outages of operated virtual systems.

Thus, we can provide a high quality stable environment in combination with standard high availability resources.

2.3 Easier maintenance and updating

Virtualisation infrastructure tools exist for snapshots of operated systems. These tools are suitable when upgrading the software on a running system or deploying a new version of the system. These updates are made easier because we can test the new version without the need of other hardware.

Thus, software replacement or configuration changes can be undertaken in a copy of the running system. If the changes are successful, the original system can be shut down and the updated system

started. Otherwise, the original state can be easily restored, quickly and remotely. These operations do not even require the presence of network administrators, as with physical systems, providing the infrastructure operator with greater flexibility and better response than in the case of physical systems.

We can thus make various changes remotely: resolve faults, update software, add other network interfaces, reconfigure IP addresses or update hardware. This is another of the benefits that the virtualisation infrastructure provides us with; we can perform interventions and changes at a suitable time and without the presence of the administrator being necessary.

2.4 Flexible network infrastructures

For larger virtualisation infrastructures it pays to deploy modern virtualisation network technologies. These provide the option of administration to network administrators even in the environment of those infrastructures. An example could be the distributed Cisco Nexus1000v switch for the VMware vSphere, Hyper-V or KVM environment, or the VM-FEX concept.

The flexible network infrastructure of data centres must be able to react to the increased transmission rate of the computer network, without outages. It is best to consider deployment of at least 10GE technology in data centres with the option of increasing capacity by connecting 10GE channels.

1GE technology is not viable to support long-term for larger data centres, even in the case of pooling physical channels, because the distribution of operation will constantly come up against the limitations of the 1GE channels. This may be a problem in the case of greater sharing of physical channels, which can occur in virtualisation infrastructures to a greater extent than in classical physical topologies.

Significant growth in the use of 40GE or 100GE technology is predicted by experts in the next few years.

2.5 Cost savings

The advantages mentioned above allow us to put together very effective and scalable services, both in technical and financial terms. Savings will be vary by installation, but the following can be expected:

- Faster fault resolution.
- Generally lower initial financial investment, no hardware purchases and the option of gradually supplementing capacity, performance or functions.
- The option of a temporary or gradual increase in capacity.
- Savings from non-implemented redundant hardware devices.
- Lower OPEX costs – savings in electricity consumption and cooling costs.

3 Network services suitable for virtualisation

Network services not requiring substantial hardware acceleration of network operation are suitable for virtualisation. Therefore, we do not foresee that virtualisation of today's powerful hardware routers or devices whose availability is fundamental to the virtualisation infrastructure itself would bring any benefits. **Thus network devices requiring hardware acceleration may not be suitable for virtualisation.**

On the contrary, implementation of the following services is recommended:

- VPN concentrators for users.
- Site-to-site VPN.
- Simple routers with NAT/PAT function.
- Security devices (e.g. IDS/IPS).
- Virtual server routers or firewalls.
- BGP routers, BHR and route reflectors.
- Load-balancers for traffic balancing between virtual systems.
- ISATAP routers for IPv6.

From the perspective of network infrastructure administrators, a risk element is in their excess use and the creation of unnecessarily complicated or chained network infrastructures. Administrators should endeavour to virtualise those services at the periphery of the network and not part of its core.

A situation must not arise where the actual virtualisation infrastructure on which the functioning of the network depends will essentially depend on a service by which it secures its own functioning. The KISS (“Keep It Simple, Stupid!”) rule should be kept in mind in the design phase.

Therefore it is not suitable to virtualise routers that are part of the backbone network, routers whose data flow is too large in view of the virtualisation infrastructure and its capacity, or services that are more suitably performed on dedicated active devices of the computer network operated outside the virtualisation infrastructure.

3.1 VPN concentrators

One of the most common network applications and services is the VPN (Virtual Private Network) service. Here, we can distinguish two types of VPN service:

- Site-to-site VPN for connection of two or more networks by means of another network (usually a public IP network).
- VPN for end users and for their secure access to networks.

In terms of virtualisation infrastructures the main requirements for these services are:

- Stability and availability.
- Demands on processing power.
- Often working across several network interfaces.
- There may also be uneven use during the calendar year.

These services are candidates for virtualisation, because virtualisation can resolve the above mentioned requirements.

With VPN services, one of whose most demanding components is CPU power for encryption, an adequate quantity of vCPU and memory can be dedicated. At peak times, this dedicated capacity will be used and in less exposed times it will be available to other applications. If a higher increase in demand is expected at any time, we can temporarily increase these reserved capacities further, in many cases without service downtime.

In essence, the only compelling argument against virtualisation of this service is the large data flow, which depends on the level of use of the specific implementation. In such case it is necessary to ensure an adequate network capacity to manage this flow.

3.1.1 Example 1: User VPN, virtualised VPN concentrator/server

User VPN access is a very nice and good application of a virtualised network service. This service is characterised, for example, in the university environment, but uneven use, which is mainly at night and in periods outside of lectures. It is possible to allocate virtual VPN servers more capacity (vCPU, vRAM) in the virtualisation environment, which is used when required and otherwise by other virtual machines.

Operation handling is carried out on allocated virtual processors (vCPU). If hardware is available in the virtualisation environment, on which encryption operation can be accelerated, we can also use virtual devices. This support is available in modern x86 processors.

The power of the virtual devices is scalable – in the event of a growth in load (e.g. a larger number of users or volume of operation), power can be increased by adding other resources. In the case of licensed products it is also necessary to acquire the relevant licences.

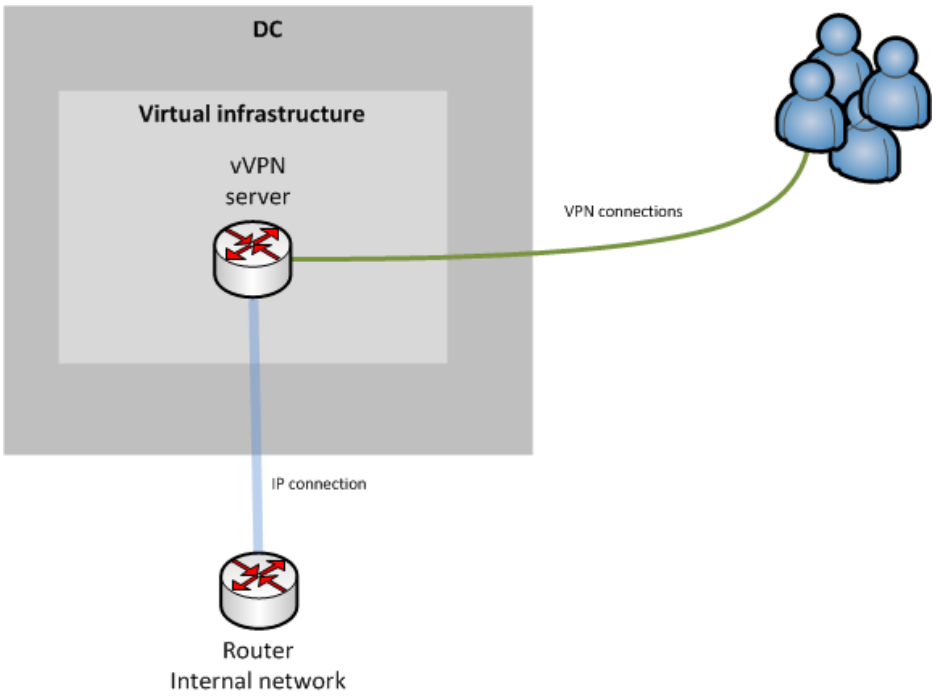

Figure 3.1: Connection to virtual VPN server.

3.1.2 Example 2: IPsec VPN for virtual server access to remote resources

It is easy to create secure connections between virtual servers and remote IP networks. Virtual infrastructure is set up in the virtualisation infrastructure environment consisting of several virtual servers and virtual routers/one VPN server. This virtual network device provides an IPsec connection to the target networks.

This VPN server may also be used as an IP router of all traffic or it may provide for just IPsec traffic, with the rest of the IP connectivity provided by another router, possibly physical, with high capacity.

The same possibilities apply to this virtual device. In the event of a growth in operation, its power can be increased by adding other resources. In the case of licensed products, it is also necessary to acquire the relevant licences.

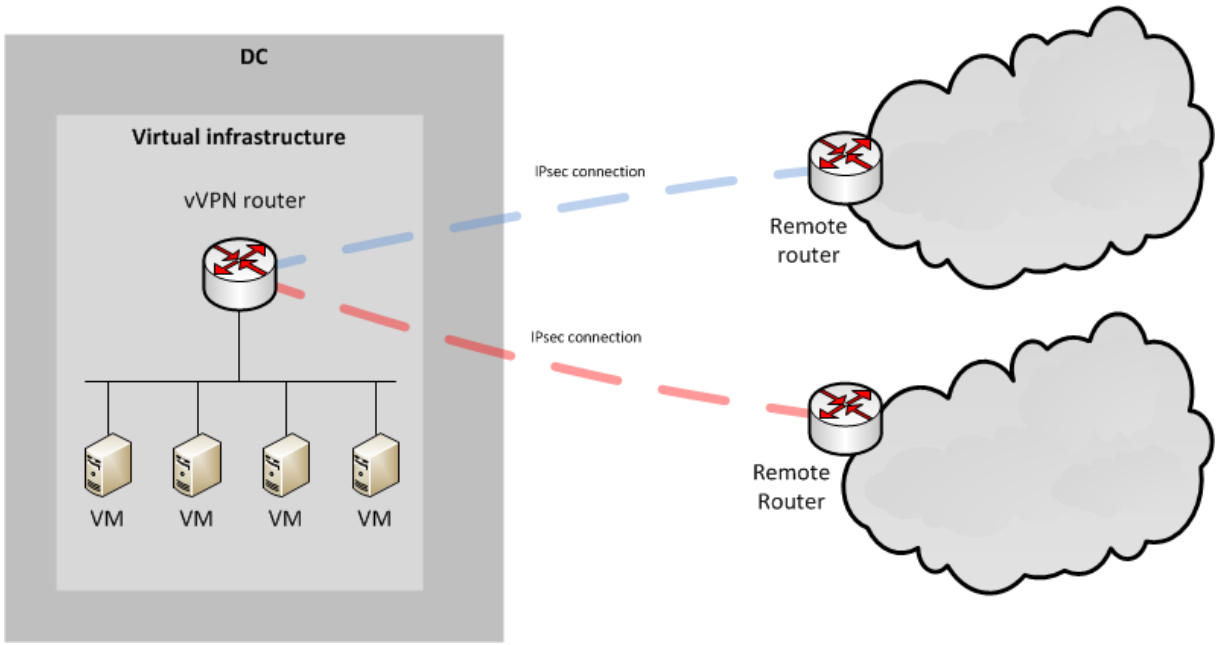

Figure 3.2: IPsec connection to remote networks.

3.2 Routers with NAT function

At many organisations parts of the network are available which are connected outside of regular operational infrastructure. An example could be networks for visitors to an organisation.

For these purposes use of a router with IPv4 address Network Address Translation / Port Address Translation (NAT/PAT) or with IPv6 routing support may be good.

It is possible to create several interfaces in the virtual router for these purposes. If the number of interfaces is limited, then it is possible to create an interface with 802.1Q support and make the required number of sub-interfaces within that one interface.

Infrastructures constructed in this way require adequate network capacity able to guarantee sufficient free bandwidth not just for the virtualisation router, but also for the other systems operated on the given server system.

In this case it is necessary to consider the load on the physical network connections. Virtualisation systems can optimise the distribution of running virtual machines without reallocating hardware resources according to CPU load, RAM usage or operations I/O, but not usually network traffic. In the event of a large dataflow there could be problems with the transmission rate of the computer network. This problem is possible and often resolved by setting network limits or QoS parameters on the virtual or physical switches of the data centre.

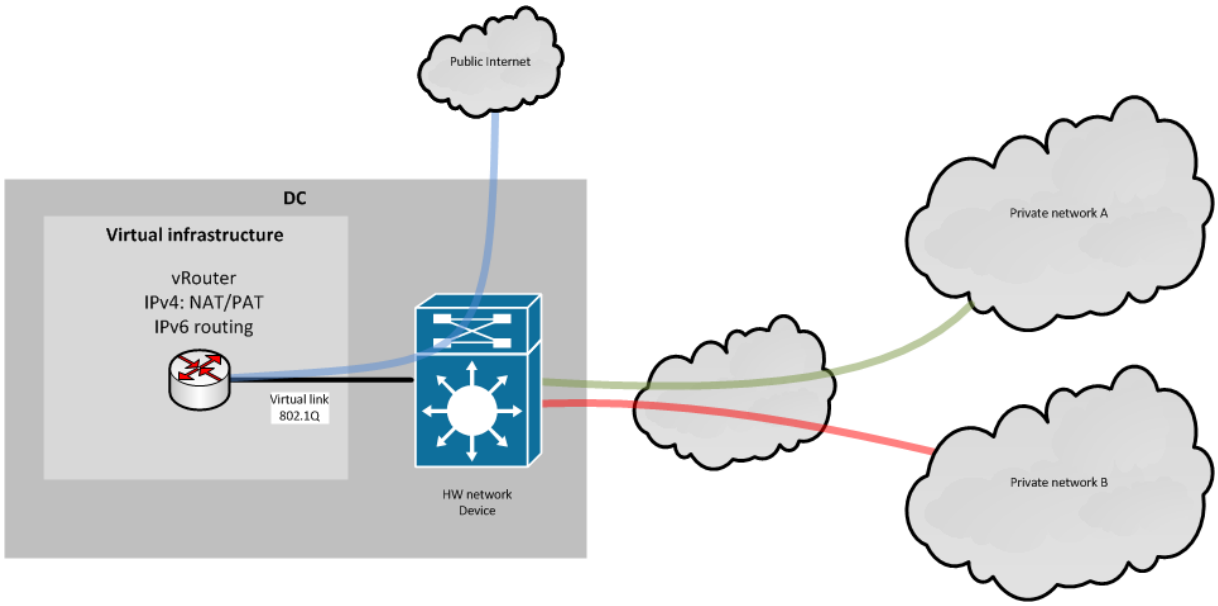

Figure 3.3: Virtual router networking.

Requirements for CPU and RAM capacity and network connectivity of the router can be differentiated on a case-by-case basis. In general we can say that the requirements of virtualised routers are not large.

An argument against virtualisation of the above mentioned services could be excessively large data flows or accelerated network ports.

3.3 BGP routers

Often, particularly in large networks, routers whose main function is not routing traffic, but routing support are deployed. The task of these network devices is, for example, calculation or distribution of routing information among other routers.

A typical application is in the deployment of Border Gateway Protocol (BGP) route reflectors, Remotely Triggered Black Hole (RTBH) filtering routers or BGP route servers.

Requirements for these routers are relatively simple: adequate CPU and RAM. Usually these routers manage with one physical network interface. In the past, powerful and costly routers with sufficient capacity, but few interfaces had to be purchased for these purposes.

3.3.1 BGP route reflectors

Routing is a very sensitive matter and downtime at the data centre should not affect the functionality of overall routing. Therefore, it is good to have at least one physical BGP route reflector independent

of the virtualisation infrastructure in operation. The second option is the operation of several BGP route reflectors; however, they must be at various data centres of the given IP network.

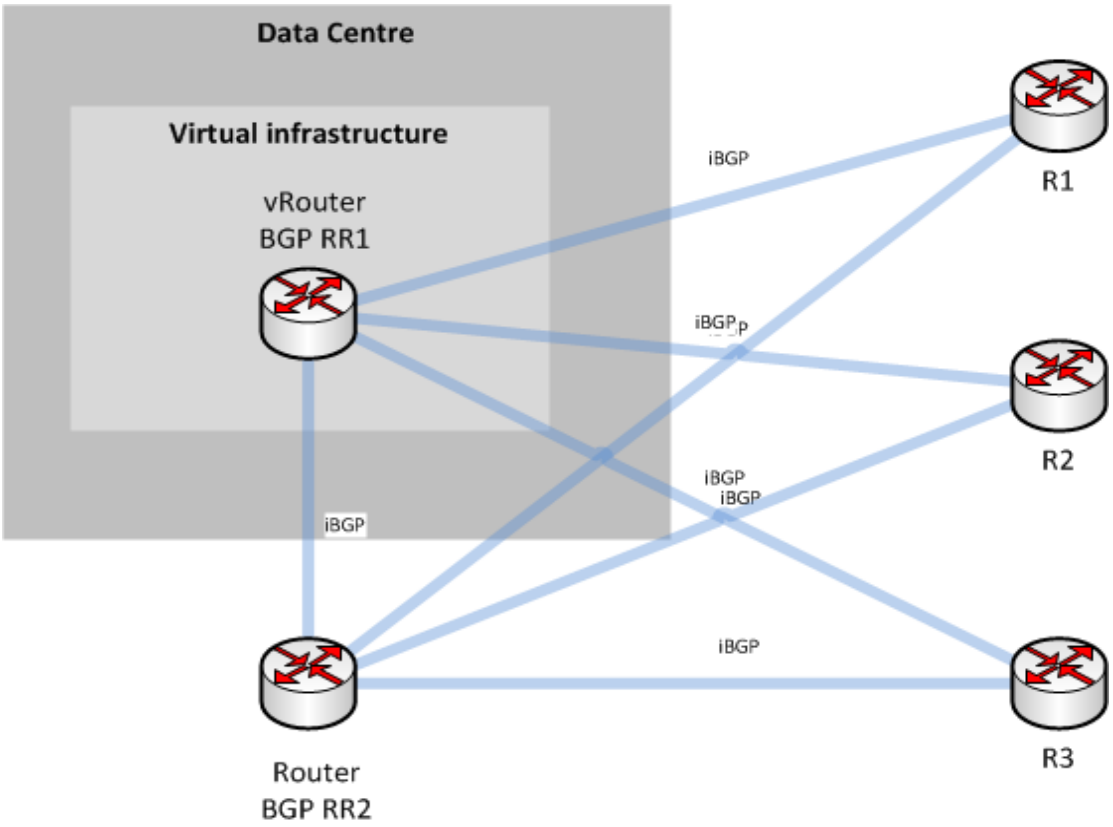

Figure 3.4: Connection of BGP route reflectors

3.3.2 RTBH servers

Remotely Triggered Black Hole (RTBH) routing is a technology used as one of the security mechanisms for blocking in computer networks with several routers. The technology uses the BGP routing protocol.

There is one router on the network, by means of which the BGP protocol is distributed to other routers by IP addresses or ranges that will be blocked in the network.

With normal deployment this technology is not threatened by limitation of actual traffic in the event of outage of the RTBH server. Therefore the extent to which the RTBH router needs to be backed up is at the discretion of administrators.

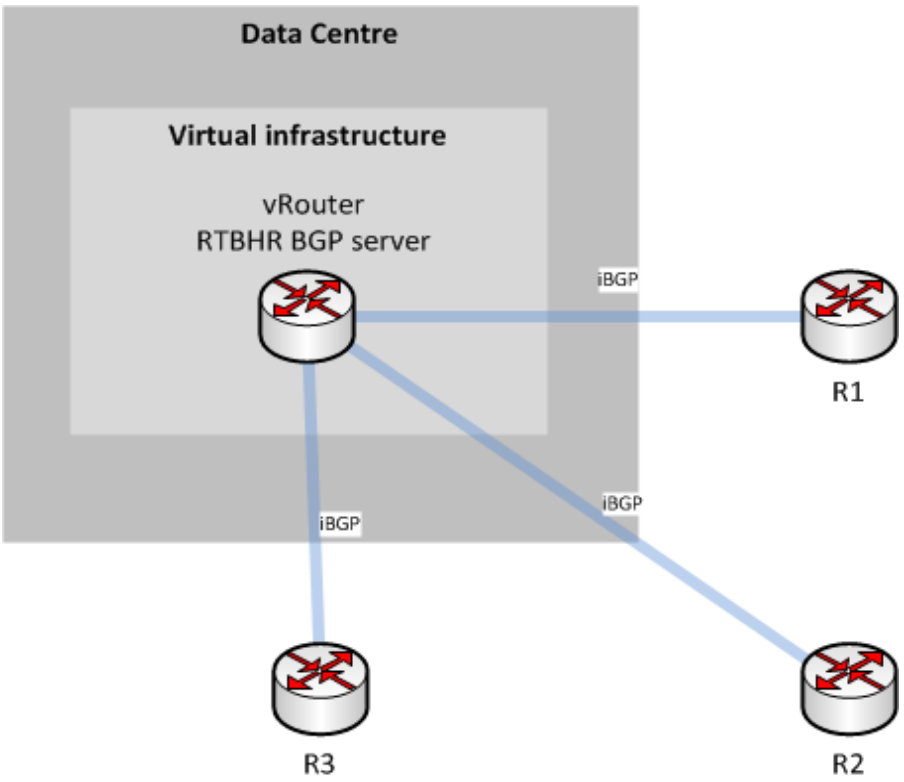

Figure 3.5: Virtualised RTBH server connection diagram

3.4 ISATAP router

Intra-Site Automatic Tunnel Addressing Protocol (ISATAP) is a mechanism providing IPv6 connectivity to systems whose native connectivity is IPv4 only. Terminal systems send IPv6 packets packaged in IPv4 packets to the dedicated computer network system.

ISATAP servers can be implemented on many platforms and thanks to the relatively easy implementation this application is also suitable for virtualisation. Applications exist both for operating systems (Linux, BSD, Windows™) and for dedicated virtualised routers.

If IPv6 is implemented on the computer network, resources for transitional mechanisms, which include ISATAP, can be operated in the virtual infrastructure. Operation of an ISATAP server is not especially demanding on CPU performance. Depending on the number of users it can, however, be demanding on network performance. With an adequately dimensioned network connection, it is an application suitable for virtualisation.

4 Network connection of virtual infrastructures

In practice we could be faced with the question of how to connect virtual servers operated at another data centre or virtual infrastructure to the actual network in such a way that they are addressed from the IP ranges of our own network.

One solution is to operate a virtual router that will be connected to the network infrastructure of the parent organisation by one of the available VPN technologies. Various types of VPN and tunnelling mechanisms can be used (IPsec, SSL, MPLS VPN, VXLAN), through a data or physical circuit.

Instead of installing a physical system in the physical environment of a remote data centre we can deploy a virtual router with two or more interfaces and which we can remotely and simply configure and thus provide quality and secure network connection.

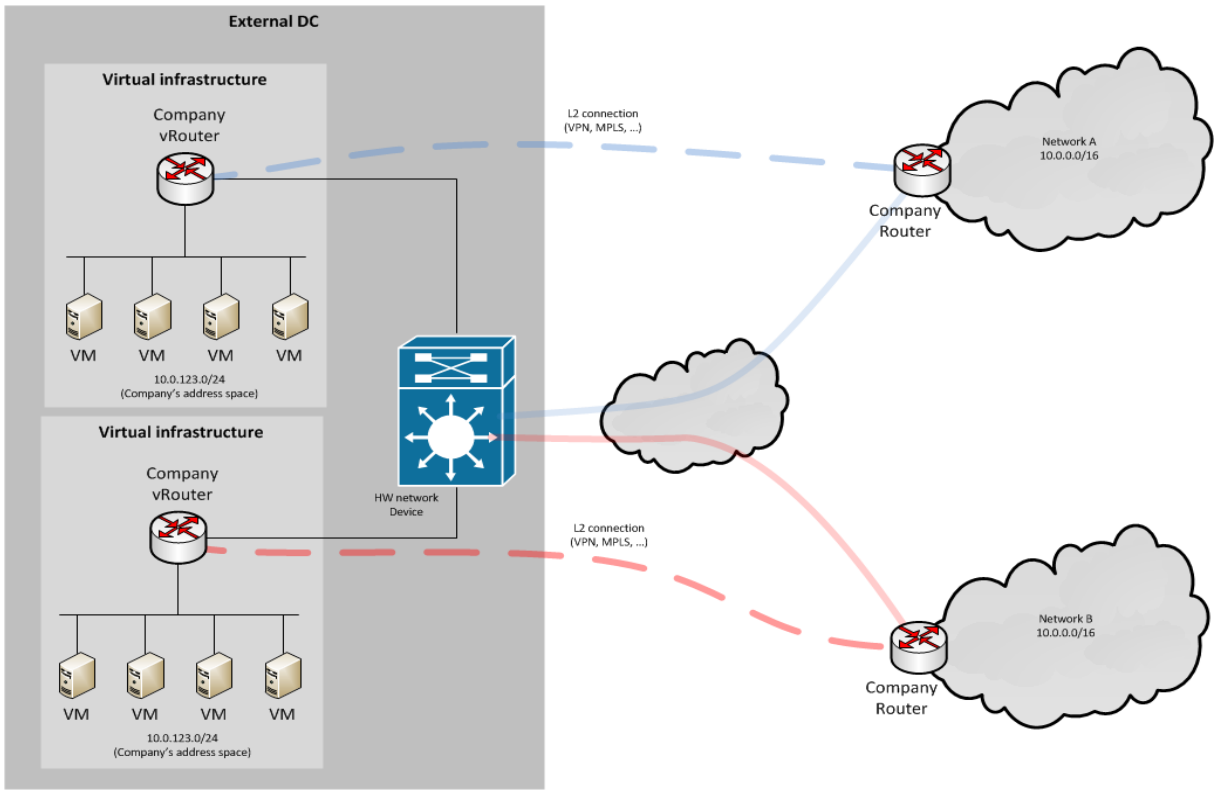

Figure 4.1: Network connection of a hosted virtual infrastructure in an external data centre.

In this way we can remotely operate a whole server or network infrastructure at a remote data centre or in its virtualised infrastructure. Such an infrastructure is then very easy to transfer to another data centre, it is only necessary to transfer existing images of the virtual servers and routers and provide a connection to the new network. In this case all IP addresses and other settings concerning the computer network are maintained transparently.

Another benefit is a local connection between the virtual router and virtual servers in the data centre, so it is easier to resolve any redundancy or more complicated connection. Essentially this is copying classical approaches that were applied even with physical infrastructure. Network administrators have the option of configuring and controlling access to virtualised servers hosted in external infrastructures (e.g. ACL, firewall rules, QoS parameters. and IP address assignment).

Connections from an external data centre to the main network of the operator can be made in several ways no different to today's accessing used for physical infrastructures. They will thus depend chiefly on the options for connection at the actual data centre. Universal is the tunnelling of L2 traffic through a public IP network or various types of data circuits (e.g. MPLS VPN, 802.1Q tunnels, VXLAN etc.).

In the example above (Figure 4.1), notice that the networks of two operators use the identically addressed private space 10.0.0.0/16. The fact that the network infrastructure of the data centre is L2 infrastructure from the viewpoint of the network means that no collision occurs and the connection is possible and fully functional.

5 Operational recommendations

- Have a network infrastructure development strategy. Clarify which of its parts are suitable for virtualisation and which devices are not.
- Don't expect big investment savings. The advantage is usually in increased flexibility and scalability of infrastructure, savings are chiefly in overheads (energy, space).
- When virtualising, pay attention to compatibility with the existing network infrastructure, but also with your virtualisation infrastructure.
- Install a quality, but also a simple network infrastructure at the data centre. Operation of virtual network devices is usually more complicated than with physical devices, where connection is usually made with one physical cable. Certain network problems are more difficult to resolve or detect.
- Have tools for mirroring and monitoring traffic in the virtualisation infrastructure available.
- Do not virtualise those parts of the network infrastructure that are essential for its initialisation or running. A typically unsuitable application can, for example, be routing the network for administration of virtualisation infrastructure on a virtual router running in that infrastructure.
- Depending on the state of the network infrastructure, applications requiring greater bandwidth can be implemented. Be sure or ensure that this traffic does not limit the traffic of other systems in the shared virtualisation infrastructure.

5.1 The minimum requirements for virtualisation

The following minimum conditions must be met in order to virtualise the network device.

- Network infrastructure that provides connectivity for virtualisation infrastructure must not depend on a the *virtual network device* and its proper function.
- The management of virtualisation infrastructure must not depend on the *virtual network device* and its proper function.

5.2 Virtual devices

The number of virtual router implementations has recently risen. A long-term factor contributing to this expansion is the use of x86 platforms for physical devices, which has clearly popularised this porting with manufacturers.

According to manufacturers the operation of virtual routers is often possible in any virtualisation architecture. The most widespread and manufacturer-supported platforms on the market are VMware vSphere, KVM, Hyper-V and Xen. In technical documentation manufacturers often compare platforms and usually the highest performance is measured in the VMware vSphere environment.

The selection of products and their characteristics listed here may have increased by the time you read this document. Products that can be used in several virtualisation environments are reviewed.

Virtualised routers have similar characteristics to their physical counterparts but they do not have accelerated hardware ports which allow physical devices to handle and accelerate running network traffic on the data-plane level. Virtualised routers do not have these options for acceleration and traffic must be processed by the processor, thus on the control-plane level.

5.2.1 Linux

The Linux OS support is generally very good in all virtualisation environments. In many environments it is necessary to install modules in the OS for better support and running in the virtualisation environment.

The options of the Linux OS can be expanded thanks to a large quantity of applications. In terms of network services the following services operate without a problem:

- DHCP server and DHCP relay on the router.
- VPN (e.g. OpenVPN, IPsec solutions).
- Dynamic routing and support for basically all regular router protocols – RIP, OSPF, BGP (Quagga/Zebra, BIRD).
- NAT.
- packet filters (iptables and iptables6).
- both IPv4 and IPv6 support.

Practical experience shows that the NAT44 service for 3000 terminal systems with allocated 2vCPU and 1GB RAM with continual data flow in ranges of 100 Mbps units can be implemented without problem in the VMware vSphere environment, without any noticeable burden on the virtual router and virtualisation infrastructure.

Solutions based on OS, including Linux, are not usually burdened with licence fees, which allow quick and easy implementation of many network applications.

The same applies to virtualised Linux routers as to physical installation – systems are able to do what is installed on them. The more demanding installation process associated with the installation and configuration of applications is balanced by greater flexibility.

5.2.2 *BSD systems

Similar rules and possibilities apply to BSD (“Berkeley Unix”) systems as with Linux systems. They can be virtualised without problems; they are well-supported in virtualisation technologies and have essentially identical options for deployment to the Linux OS.

In some areas concerning network applications BSD provides better options than in the case of the Linux OS, in particular the NetBSD variant.

5.2.3 Cisco CSR1000v

This is a Cisco virtual router intended for the general virtualisation environment; the manufacturer explicitly mentions VMware vSphere, KVM and Hyper-V. The manufacturer classifies the router as Integrated Services Router (ISR), thus it includes the same functionality available in Cisco ISR routers of series 900, 1900, 2900 and 3900.

The Comand-Line Inerface (CLI) is familiar from conventional routers. The virtual router thus supports most protocols and functionalities, including a firewall. In this respect it is a functionally very well equipped router, which can be used for most applications.

This router is licensed according to capacity. Basic installation is possible with a temporary two-month licence. Licences are available for 10 / 25 / 50 / 100 and 250 Mbps.

Standard, Advanced and Premium feature sets are available for 10–50 Mbps. Only the Standard version is available for 100/250 Mbps capacity.

The Standard license covers the regular functionality of the router (all routing protocols, DHCP, NAT, HSRP, Netflow, etc.). The Advanced license also covers a zone-base firewall, IPsec and other VPN functionality. The Premium licence in addition covers MPLS, QoS and IP SLA support.

Requirements for virtual infrastructure are the same with all variants – 4 vCPU, 4GB RAM and 8GB disk space. It can therefore be used as a firewall and router for virtual infrastructures, such as for IPsec VPN, and, thanks to classical support for routing protocols, it can be used in BGP structures. In some deployments it can provide interesting MPLS protocol support.

5.2.4 Cisco ASA 1000V

This is a Cisco virtual firewall intended for the VMware vSphere environment and requires the distributed Cisco Nexus1000V switch for implementation. This switch is also available in the basic free version, which is sufficient for its operation. This is a device based on the classic Adaptive Security Appliance (ASA) equipment, but unlike them, its functionality is limited.

The virtual firewall is suitable for a virtualised environment where a powerful firewall is required.

Licensing is implemented according to the number of CPU sockets and thus the necessary number of virtual firewalls can be operated in the virtual infrastructure without restriction. A disadvantage, however, is the relatively high price, which becomes good value only with larger deployment. Actual operation also requires the Cisco Nexus1000V switch, which is also available in a free version.

The Cisco ASA1000V Firewall is deployed on a network as a classical router through which network traffic, to which defined rules are applied, flows. In contrast to classical ASA products, it can only be used for the implementation of IPsec site-to-site VPN, but no longer for user-VPN or SSL-VPN.

A disadvantage is that there is no support for routing; essentially it is a powerful firewall with two interfaces and no additional functionalities.

5.2.5 Cisco Adaptive Security Virtual Appliance (ASAv)

This is a Cisco virtual device intended for the general virtualisation environment. The manufacturer states that it is a porting of classical ASA devices into the virtual environment. Thus it has the same functionality as physical Cisco ASA devices.

This product is suitable for the general virtualisation environment; the manufacturer mentions VMware vSphere, KVM and Hyper-V. It is suitable for the network as a router with firewall support, but also as a VPN concentrator for user VPN as well.

This virtual device should soon be available to the public. Information about it is taken mainly from publicly available information. It was included in the selection mainly because a large number of installations are performed on Cisco ASA platforms. Information on the availability of this virtualised device may be of interest to many.

It can therefore be used as a virtual infrastructure firewall and router, and also as a VPN server. It supports routing and has up to ten interfaces.

5.2.6 FortiGate VM

There are a total of five Fortigate VM models, which differ according to licensed capacity. Capacity is defined as general transmission rate, firewall transmission rate, the number of concurrent connections and IPS performance. The maximum transmission rate for routing ranges from 500Mbs to 4Gbps.

In addition to the model of the virtual router mentioned, Fortinet also offers other virtualisation products from its portfolio, which it markets under the common name Fortigate Virtual Appliance Family.The manufacturer also provides a comparison of performance in different virtualisation environments. It is clear from these measurements that the best performance is achieved in the VMware vSphere environment.

The virtual router is functionally very well equipped and has almost the same functionality as a physical product. It can therefore be used as a virtual infrastructure firewall and router, as a VPN server and even a user VPN.

5.2.7 HP virtual services router

HP has a router in its portfolio called HP VSR (Virtual Services Router). This virtual router is licensed and available in three licensed versions, which differ only in the number of virtual CPUs, which can be 1, 2 or 8.

The router is built on the Comware7 system and is intended for the VMware vSphere and KVM environments; it is not available for the Hyper-V environment. The manufacturer claims about 15% higher technical performance in the VMware vSphere environment than in the KVM environment.

The router supports IPsec VPN, but no longer SSL VPN, which is often used for user VPN access.

The number of licensed virtual processors should have no great impact on the traffic routing itself. Better performance in the case of a larger number of processors is manifested in other services, such as IPsec performance, the calculation of routing paths in larger routing tables, etc.

6 Comparison – Characteristics of Selected Virtual Routers

	Linux	BSD	Cisco CSR1K	Cisco ASA1000v	Cisco ASAv	Fortigate VM	HP VSR
IPv4/6 routing	Y	Y	Y	Y	Y	Y	Y
IPv4 multicast	Y	Y	Y	Y	Y	Y	Y
DHCP server / relay	Y	Y	Y	Y	Y	Y	Y
Internal router protocols	(Quagga)	Y	Y	N	Y	Y	Y
External routing protocols	Y	Y	Y	N	Y	Y	Y
Number of interfaces	Unlimited	Unlimited	Up to 32	Fixed 2 + 1 mgmt	Up to 10	Up to 10	Up to 16
802.1Q	Y	Y	Y	N	Y	Y	N
QoS / shaping	Y	Y	Y	N	Y	Y	Y
NAT	Y	Y	Y	Y	Y	Y	Y
WCCP	Y	Y	Y	N	Y	Y	N
IPv6	Y	Y	Y	N	Y	Y	Y
Site-to-site IPsec VPN	Y	Y	Y	Y	Y	Y	Y
User VPN	Y	Y	Y	N	Y	Y	Y
Firewall / packet filter	Y	Y	Y	Y	Y	Y	Y
IPv6 ISATAP	Y	Y	Y	N	N	N	N

Export of traffic statistics)	Y	Y	Y	N (only from Cisco N1K)	Y	Y	Y
HA technology	Y	Y	Y	Y	Y	Y	Y
Multiple routing tables	Y	Y	Y	N	Y	Y	Y
IDS/IPS	Y	Y	N	N	Y	Y	N
MPLS	N	Y	Y	N	N	N	Y
VMware vSphere	Y	Y	Y	Y	Y	Y	Y
KVM	Y	Y	Y	N	Y [1]	Y	Y
HYPER-V	Y	Y	Y	N	Y [1]	Y	N
XEN	Y	Y	Y [1]	N	Y [1]	Y	N
Licensing	GNU/GPL	BSD licence	Licenced per throughput	Licenced per CPU cores	Licenced by vCPU	Licenced by throughput	Licenced by vCPU

1) An arbitrary virtualisation environment is given, but it is not explicitly mentioned by the manufacturer as supported.

Table 6.1: Comparison Table – Characteristics of Selected Virtual Routers

6.1 Notes on the Comparison Table

IPv4/6 routing is the ability of router to function as an IPv4 and IPv6 packet router.

IPv4 multicast is the ability of the router to route multicast IPv4 traffic.

DHCP server/relay is the ability of the router to perform a DHCP server function, or, when in relay mode, to transfer DHCP traffic to another DHCP server.

Support for **Internal router protocols** means support for the most frequently deployed non-proprietary routing protocols - RIP, OSPF and IS-IS in the IPv4 and IPv6 environment.

Support for **External routing protocols** means support for the BGP protocol, which is today de facto the only protocol for exchanging routing information between autonomous systems.

Number of interfaces defines the number of interfaces supported by the relevant router in the virtual environment.

Support for the **802.1Q** protocol is the ability of the router to support tagging of virtual networks on its own interfaces.

The item **QoS/shaping** defines support for prioritisation of traffic or its limitation.

NAT technology support determines NAT/PAT technology support in the IPv4 environment.

Support for the **WCCP** protocol is the ability of the router to route HTTP traffic to another network node/server. Although it is a Cisco proprietary protocol, it is often supported and implemented by other manufacturers.

IPv6 support means support for the IPv6 protocol and not only in the routing area. The level of implementation of individual characters of the IPv6 protocol may differ among the routers of various manufacturers and thus it is better to compare the level of implementation before deployment with documentation.

Site-to-site IPsec VPN is the ability of the router to make IPv4 VPN connections using the IPsec protocol.

User VPN is the ability of the router to provide user VPN access services, which are implemented not only using IPsec technology, but also SSL VPN.

Firewall/packet filter is the ability of the router to filter traffic on the basis of rules defined by the router administrator.

IPv4/6 routing is the ability of router to provide IPv6 connectivity to clients in a network with IPv4 connectivity.

Export of traffic statistics means the ability of the router to provide information on flowing traffic through one of the regularly used protocols, for example, Netflow, IPFIX, sFlow and so on.

HA technology support means support for high availability technologies, which resolves any backing up of the router's functions by another router. This means, for example, VRRP technology.

Support for **Multiple routing tables** means support for several separate routing tables, thus separate routing of the traffic of several networks by one router.

IDS/IPS support is the ability of the router to perform certain functions for analysing traffic, thus detection or direct limitation of traffic based on its content.

MPLS support in routers means the capability of direct connection to a MPLS network and functions of a CE or PE router. It is necessary to have one physical virtualisation infrastructure interface (or more for redundancy) in pass-through mode on all nodes of the virtualisation infrastructure. This raises the question of the necessity or appropriateness of implementing this functionality in operational practice.

The mentioned are the VMware vSphere, KVM, Hyper-V and XEN virtualisation environments with their declared support from the websites of the virtual router manufacturers. Some routers are

designed for the general virtualisation environment and this fact is mentioned for some, although the manufacturer does not explicitly indicate some of these environments.

Each product is licensed in a different way and the manner of licensing is listed under the entry Licensing. GNU/GPL and BSD licences are open-source and free licences, therefore the systems can be deployed free of charge, unless the distribution used by the manufacturer is paid for or modified. Generally, licensed routers are licensed according to performance, which is typically characterised by transmission rate or processor power.

7 Practical deployment of virtualised network devices

The cases below describe practical examples of deployment of virtualised devices on the computer network during routine operation. These are cases of deployment on the VŠB-TU Ostrava university network, which has about 28 000 users. Its backbone network is built on 10GE technologies, and has about 1000 active network devices on the fixed and wireless computer network.

7.1 NAT device

Private networks are operated in the VŠB-TU Ostrava university environments for visitors. These visitor networks are outside the main university LAN network and their local routing is made in a separate router instance (VRF) on all L3 devices on the computer network.

All localities are equipped with Cisco Catalyst 6500 devices in variations with SUP2T, SUP720-3B and SUP32 supervisors. These L3 switches are connected in a VRF separate router instance to the parent virtual router and propagate their local address spaces using the OSPF protocol. The virtual router extends the initial route (0/0) in the OSPF protocol.

Address translation (NAT/PAT) is performed on a virtual router and traffic is then routed to the public internet.

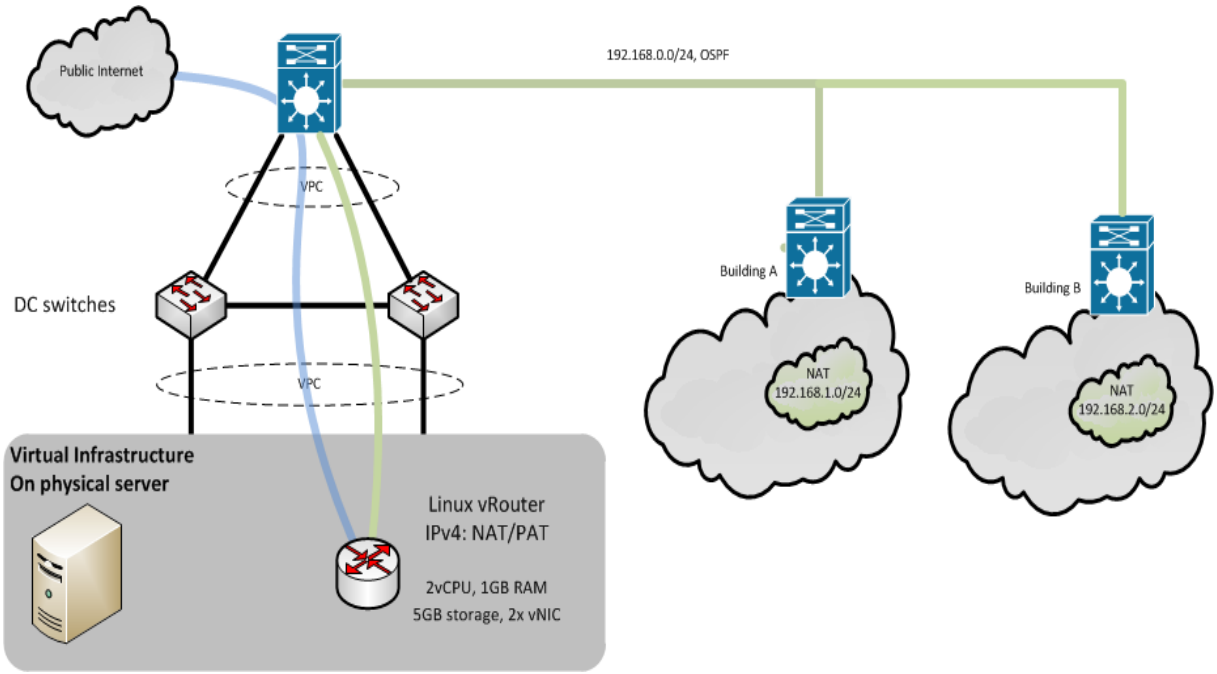

Figure 7.1: Network technical topology

A virtual router is not too demanding on virtualisation infrastructure capacity. Two virtual processors, 1GB RAM, 5 GB in the disk store and two network interfaces are allocated. The Linux distributed GNU/GPL Debian Linux is used on the virtual router and a Quagga routing daemon also distributed under the GNU/GPL licence is used for dynamic routing.

The minimum requirement of independence of the actual network and independence of the virtualisation infrastructure on the virtualised router is met here, thus the service can be virtualised.

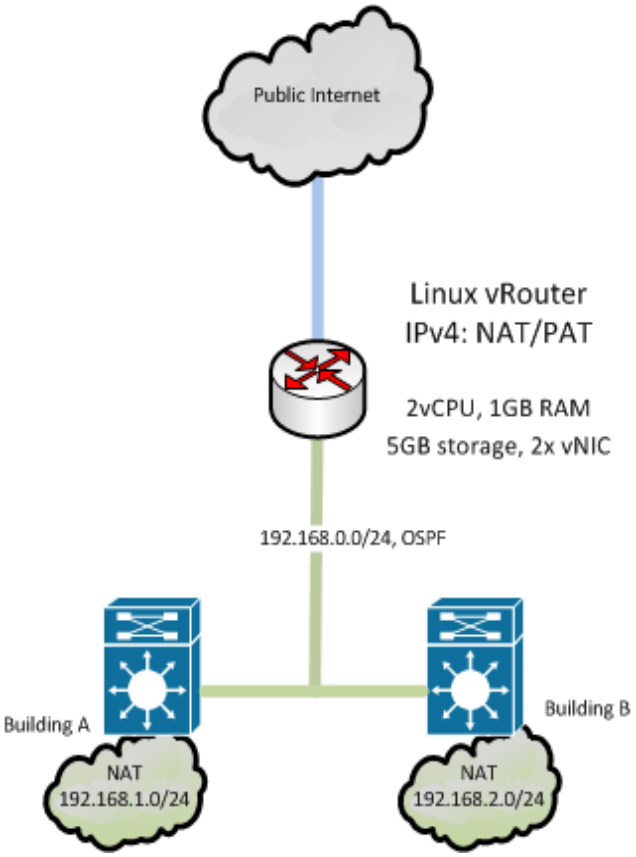

Figure 7.2: Network logical topology

Virtualised router operation is provided by virtualisation infrastructure built on VMware vSphere servers, including a total of 11 physical nodes. The situation is shown simplified in the Figure 7.2, with just one node drawn.

The distributed Cisco Nexus1000V switch is operated within the virtualisation infrastructure. Individual server systems are connected to the data centre network by two independent 10GE connections terminated at two different Cisco Nexus 5548UP switches with VPC (virtual port channel) support, which ensures redundant interconnection of network infrastructure.

The router can generate information about flowing traffic (e.g. by NetFlow technology), which can be processed on NetFlow collectors and used to analyse anomalous traffic.

An advantage of the system is its very stable operation, sufficient throughput and very fast option of increasing the capacities of the router in the event of an increase in output. System restart takes less than a minute, so any necessary system updates cause no prolonged service outages and can be performed in agreed maintenance windows. During updates the options of the virtualisation infrastructure are used, a snapshot is taken before the update is performed, to which it is very easy to return.

In this setup, with the outputs mentioned above, network operation with about 3000 non-competitive working users and a dataflow of 200Mbps was served without problem. Also integrated into this

solution is a centralised wireless network (with about 330 WiFi APs), which is also used to serve university conferences and other events.

7.2 Device for secure IPsec connection

A cluster of servers is operated in the VŠB-TU Ostrava university environment which accesses multiple secure networks. Data are collected from these networks or provided to these networks.

There are several of these secure networks, which are communicated with by a secure IPsec protocol, and they change over time. It is not good to install IPsec support for all server cluster systems and maintain all the settings, which is the most demanding solutions in terms of the system.

Therefore, installation of a Cisco CSR1000V virtual network to which traffic directed to these networks is routed has proven to be a practical solution. The only system setting on individual server systems is regulation of the routing table. Other IP traffic then travels in the classic way and is routed on the physical router of the computer network.

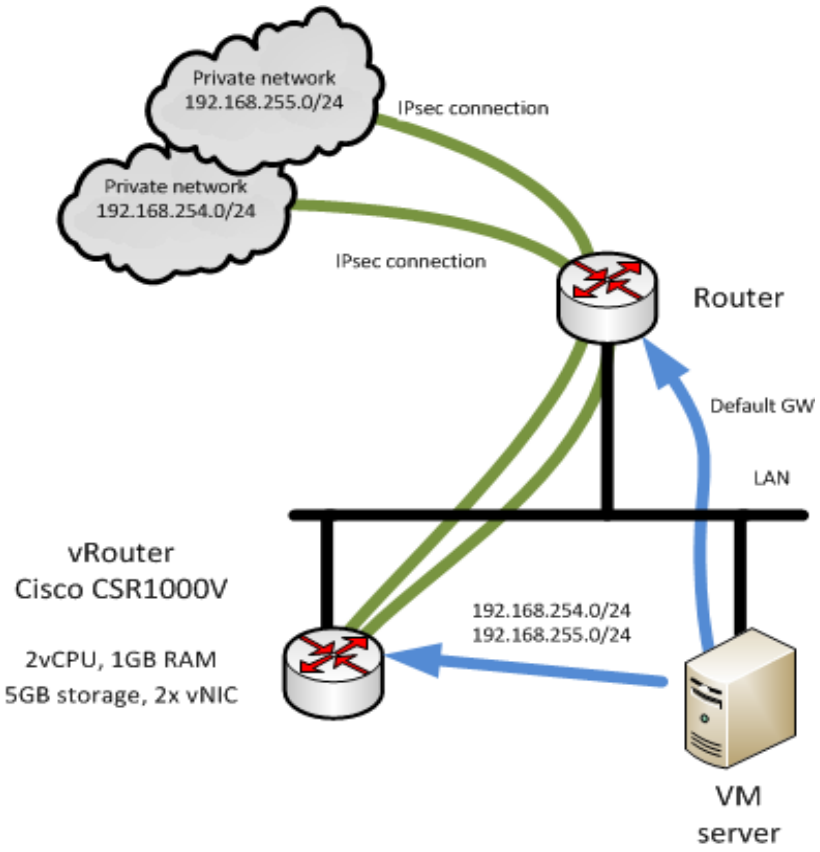

Figure 7.3: Routing of IPsec traffic (one node of the cluster is displayed)

All IPsec configurations are thus stored in one place and managed by networking specialists, and server system experts then simply use this service. An advantage is also use of a stable environment, because the virtual router is operated in the same environment as the actual server systems.

The virtual router is licensed according to the volume of traffic passing through it. Traffic above this limit is not processed by the router. In the event of a growth in traffic the program licence can be increased by buying more. In this way any potential replacement of the physical router, which would be required in the case of services on a physical device and which would also be more expensive, is eliminated.

These limits are also the reason why only traffic to the secure networks is routed through the virtual router and not all the traffic on the server cluster.

The minimum requirement of independence of the actual network and independence of the virtualisation infrastructure on the virtualised router is met here, thus the service can be virtualised.

7.3 Remotely Triggered Black Hole filtering

RTBH filtering is a technology used for blocking traffic from certain IP addresses. This consists of distribution of those IP addresses via a BGP routing protocol to the other routers on the computer network. We will term this distributing router the RTBH router.

No network traffic need pass through the RTBH router. Thus, in terms of network load, we can virtualise it without qualms. The functionality of the computer network and actual virtualisation infrastructure is not dependent on the RTBH router functionality; therefore the second condition of network device virtualisation is satisfied.

We have used this technique for blocking IP addresses for a long time on the VŠB-TU Ostrava university network. In the past it was run on a dedicated Cisco 2500 series router, which was quite sufficient for this function even despite its limited power.

Since it is appropriate for the router to support route tags, by which we indicate the networks to be propagated in the RTBH system, use of a Cisco CSR1000V virtual router was found to be a suitable solution. The configuration of the older Cisco 2500 series physical device, which was hitherto used for this purpose, was transferred to this router.

The configuration is therefore identical to that of a physical router. The virtual router has BGP connections to each devices of the computer network established on that network. Connection to the computer network is achieved through two IP independent connections using the internal OSPF routing protocol.

This modern virtualised router also supports the IPv6 protocol, has more memory and more processing power available. Other advantages of the running virtualisation include reduced purchase and operating costs, a stable centralised infrastructure environment and also there is no need to allocate physical space, power and cooling for the router.

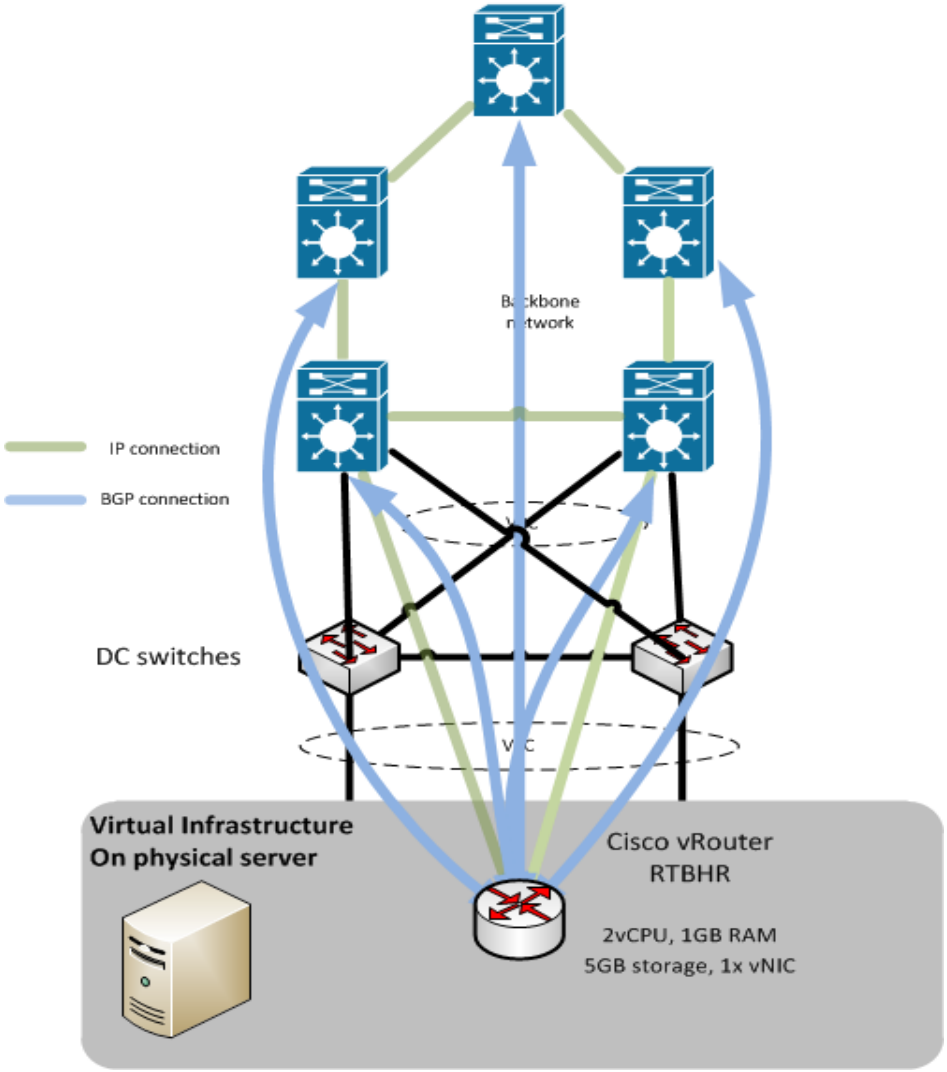

Figure 7.4: RTBH router and BGP connection

Glossary

BGP	Border Gateway Protocol. External routing protocol serving for exchange of routing information between autonomous systems. iBGP serves for exchange of routing information inside an autonomous system.
CPU	Central Processing Unit
DHCP	Dynamic Host Configuration Protocol
Firewall	A device for filtration and securing network traffic.
HA	High Availability
Hyper-V	Microsoft virtualisation solution built on server editions of MS Windows.
IPFIX	IP Flow Information Export. A protocol for the export of information on traffic flowing through the network node.
IPsec	Protocol for encrypting the communication of nodes on the IP network.
ISATAP	Intra-Site Automatic Tunnel Addressing Protocol. A mechanism providing IPv6 connectivity to systems whose native connectivity is IPv4 only.
KVM	Kernel-based virtual Virtual Machine. A virtualisation solution built on the Linux OS. It has been ported into the FreeBSD environment.
MPLS	Multiprotocol Label Switching – a protocol for creating connections in MPLS networks.
NAT	Network Address Translation – a technology for translating private IPv4 addresses
Netflow	Open protocol developed by Cisco for the export of information on traffic flowing through the network node.
OSPF	Open Shortest Path First –. A commonly-used internal routing protocol.
QoS	Quality of Service
RIP	Routing Information Protocol. A simple internal routing protocol, seldom used due to its limited options.
sFlow	Protocol for the export of information on traffic flowing through the network node.
Shaping	Techniques used for limiting the width of bandwidth used.
SLA	Service-Level Agreement
SSL VPN	VPN connection based on tunnelling traffic using SSL (Socket Secure Layer) technology
vCPU	Virtual CPU allocated to a virtual machine.
VPN	Virtual Private Network
VRRP	Virtual Router Redundancy Protocol
vSphere	VMware virtualisation technology.
VXLAN	Virtual Extensible LAN. A technology for network connection of data centres, the IP protocol is used for tunnelling.

WCCP	Web Cache Communication Protocol. WCCP serves for routing WWW traffic by a router to a selected network node, usually an HTTP proxy server.

Complete BPDs are available at http://services.geant.net/cbp/Pages/Home.aspx

campus-bp-announcements@terena.org