

CESNET Technical Report 4/2010

Using VMware in the Area of Processing of Medical Image Information

KAREL SLAVÍČEK, MICHAL JAVORNÍK

Received 18.12.2009

Abstract

The main goal was to verify the VMware technology for applications that provide special services in the area of medical image data processing (reliable storage systems, supporting of the primary diagnostic processes, secure and reliable electronic exchange of medical image data and related information among healthcare and research institutions, effective development and usage of knowledge databases in this area, etc.). Two VMware servers (Cesnet main PoP in Prague and in Brno) were used in tests. Three different lines connecting these servers (EoMPLS tunnel, GE line across the Cesnet DWDM backbone and a legacy IP network sharing commodity IP traffic) were tested step by step. Selected applications from MeDiMed (Metropolitan Digital Imaging in Medicine) solution were employed in tests. The proposed solution should significantly increase the robustness of these applications without increasing of the complexity of the whole solution (one medical institution, regional solution, etc.).

Keywords: PACS, MeDiMed, virtualization

1 Introduction

1.1 Processing of Medical Image Data

Communication in the area of medical image data processing is mostly based on the communication standard DICOM (Digital Imaging and Communication in Medicine). DICOM communication takes place among the PACS archiving systems, acquisition modalities, diagnostic workstations, etc. Those equipments are located in the secure computer networks of various medical institutions.

Remote DICOM applications are unambiguously identified via a unique IP addresses, communication ports and values of DICOM AE_Title. These values uniquely identify every DICOM application within the DICOM application domain. Communication port and the value of AE_Title must be configured in accordance with the rules of the local DICOM application domain as well as in accordance with other components of the clinical information system in medical institution. Fixed IP address is expected on the side where DICOM service is provided. Remote institutions interconnected through a dedicated network communicate via specifically configured firewalls that perform translation of IP addresses to a unique address spaces.

Redirecting the requirements between the primary and backup service provider in the event of failure can not be in many cases addressed automatically without user intervention.

For every alternative of technical solution we must always consider a possible link to key risks of such approach, in particular the risk of loss or nonavailability of image data or confusion of image documentation. The risks can be avoided by backing up the key components of archiving and communications systems like long term archives of medical image data, nodes for exchange of image data among medical institutions, etc.

1.2 The Principle of VMware

Virtual servers can run on one or more of physical VMware servers. If the VMware vCenter is part of the server farm then virtual server can be automatically moved between physical servers for load balancing, physical servers utilisation optimisation or server hardware failure overcoming. The expected VMware deployment scenario is in Figure 1. VMware servers correspond to physical servers. One or more virtual servers can run on a VMware server. Vice versa given virtual server can run on any of the VMware servers in the farm and can be moved between these servers or manually or automatically managed by VMware vCenter.

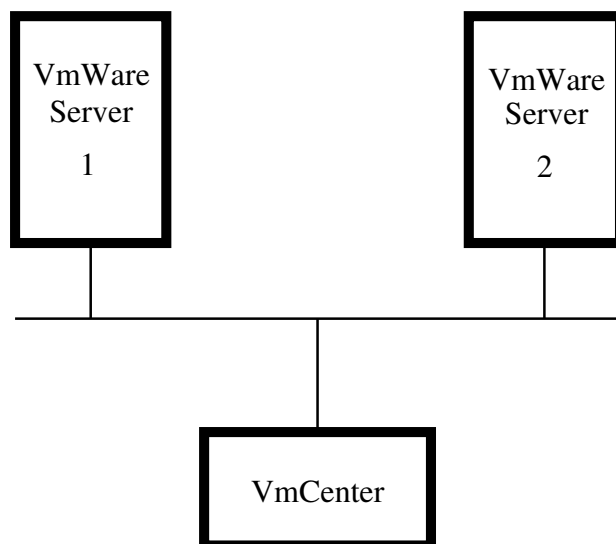


Figure 1. The basic idea of VMware utilisation in MeDiMed project.

1.3 The Storage Management

The storage subsystem is mandatory part of server redundancy solution. If the virtual server would store the application data on the local VMware server then in case of failure of this physical server these data will not be accessible by the virtual server migrated to another host. In general the storage subsystem should be independent of VMware servers. All servers capable to host given virtual server should share common storage subsystem. The proper structure is easy to see from Figure 2. Redundancy of storage subsystem is outside of scope of our testing for two main reasons: Redundancy in storage world is solved in many commercially available equipment. Disk arrays above certain price level usually can offer remote data replication driven by raid controller that means independent of servers and

applications using this data. The second reason was lack of suitable disk arrays for testing. Remote data replication usually offer disk arrays above certain price level. It is not possible to perform tests on production disk arrays and on the other hand it is not cost effective to have disk arrays of this level only for testing. In this test we have used one external data storage shared between both testing VMware servers.

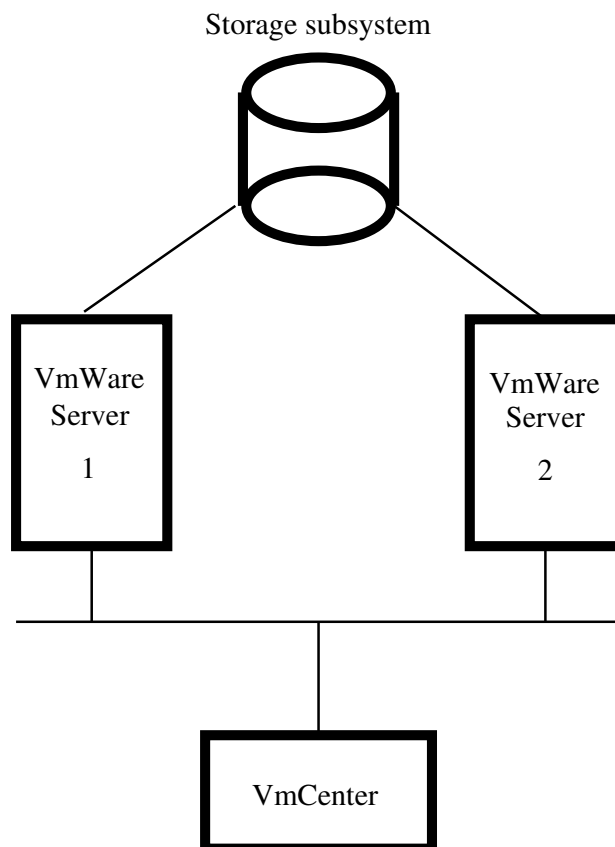


Figure 2. The storage subsystem should be independent of VMware servers. All servers capable to host given virtual server should share common storage subsystem.

2 Test Layout

2.1 VMware in MeDiMed Test Setup

Two VMware servers were used in tests. The first one was located at Cesnet main PoP in Prague the second one was located in Cesnet PoP in Brno. Three different lines connecting these servers were tested step by step. All of these lines were of nominal bandwidth 1 Gbps.

The test was concentrated on measurement of switchover time in case of failure of primary server. The convergence time of storage subsystem was not tested due to lack of necessary equipment, namely disk arrays with raid controller supporting autonomous data replication.

In the MeDiMed environment the main reason for deployment of VMware would be cost effective server redundancy solution. The overview of proposed VMware utilisation in MeDiMed is easy to see from Figure 3. The VMware servers 1 and 2 on the figure will be hosted at two independent locations connected via Cesnet network. Virtual servers will be spread across both servers equally to optimise switchover time in case of failure of one of the VMware servers.

The external data storage shared between both testing VMware servers was implemented on a Linux server. This server has offered some disk capacity via NFS protocol. We have tested iSCSI as well. As iSCSI disk server was used open source software *iscsitarget* available at sourceforge.net. Comparing to NFS the configuration of *iscsitarget* is more complicated. There were some difficulties to make it cooperate with VMware iSCSI client. For this particular test the iSCSI protocol doesn't bring any reasonable benefit comparing to incumbent but little bit obsolete NFS. For this reason we've decided to use NFS.

The dashed line is the "line under test" – line Prague-Brno. This line was step by step realised by three different ways: EoMPLS tunnel, GE line across the Cesnet DWDM backbone and a legacy IP network sharing commodity IP traffic.

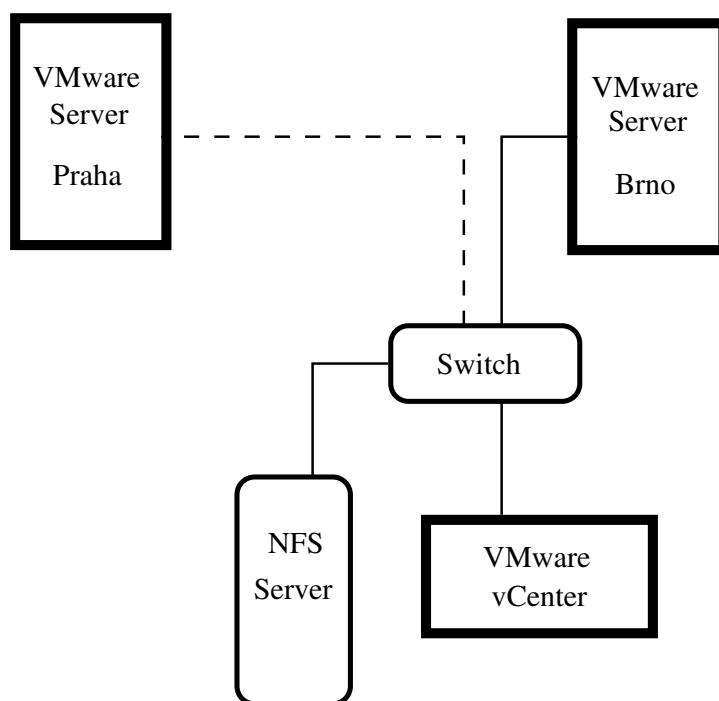


Figure 3. The MeDiMed VMwae test setup.

2.2 EoMPLS

The EoMPLS tunnel was ended on GE interfaces on router R103 in Prague and on R98 in Brno. Both routers are Cisco 7609 running IOS 12.2.33.SRB5. The EoMPLS tunnel shares the 10GE backbone line with commodity IP traffic. No QoS or any prioritisation of this test traffic was configured.

2.3 DWDM

The DWDM GE channel was constructed as dedicated lambda service. On Cesnet backbone DWDM ring it was terminated at node Prague and Brno on multi rate transponders. FEC as specified by ITU-T G.975 was used for this channel.

2.4 VMware configuration

The VMware infrastructure 3.5 was used for testing. VMware servers were VMware infrastructure 3.5 ESX, i.e. servers running directly in hardware without any underlying general common operating system. Hardware used for both of these servers was Dell PowerEdge 2130. VMware vCenter was running on Windows XP on a notebook. For all VMware components the 60 days trial version was used.

2.5 Virtual server switchover test

The aim of the test was to measure the impact of used networking technology to switchover time.

In a virtualisation environment the operating system runs on the master node only and is started on the slave node, if needed – especially in case of hardware failure of the master node. This introduces some delay in service restoration in case of hardware failure because the backup system should boot up. On the other hand this environment is not dependent on the support of the application. Any application on any operating system has to expect that the system may reboot at any time and has to check the status of data structures upon its start.

The mandatory part and we can say the most of the time needed for switchover is the time used for switchover is the time that the virtual server needs to boot up. This time is several times more that the time needed by VMware vCenter to learn that the primary node is dead and initiation of the virtual server migration. The time difference of the switchover caused by different networking technologies used is only few percent of the whole virtual server switchover time.

3 Results

The main result is that all of the possible connections are more or less equal. Most of the time needed to populate an application on the backup hardware is used to boot up the virtual machine on the backup hardware. Based on experiences from these test we can propose the following improvement of the switchover time of virtual servers: Virtual servers should be “pre-booted” on backup VMware ESX servers. By the word “pre-booted” we mean booted up to but not including startup of applications and of course keeping the virtual network and storage interface in down state.

VMware platform allows undepreciated communication of DICOM application entities in within one DICOM application domain as well as communication between geographically distant sites. Because of fixed IP address of DICOM service provider there is no need of any changes in the configuration of DICOM applications within the scope of the system covering the group of cooperating medical institutions. The concept thus protects the past as well as future investments.

4 Conclusion

The system utilizing VMware can exploit the hardware platform much more efficiently. When compare cluster architecture with corresponding solution employing VMware than cluster architecture is much more expensive option. Deploying VMware in this application area can increase economic efficiency in the processing of medical image information. VMware represents more rational usage of the available hardware. The concept of VMware in specified application area can be implemented gradually and thus pragmatically address the upcoming problems.

It is still necessary to perform testing of storage subsystem before deployment of VMware into MeDiMed production environment. The areas of interest or areas where we can expect some difficulties are server – storage system synchronisation and remounting of filesystem in case of virtual server switchover.