



NTP in a Virtualised Infrastructure

Table of Contents

Change control	2
Disclaimer	3
Glossary	4
Background	5
This document	5
Assumptions	5
Virtualisation time-keeping	6
ESX hosts	6
Guests - general	6
Guests - Windows Active Directory	6
Stratum Layers	7
Strata overview	7
Layer 1	7
Layer 2	7
Layer 3	7
Pictorial representation	8

Change control

Document movement details.

Date	Description	Version	Name
07/08/2009	Document created	v.01	D Woollard
07/08/2009	Document published on http://vmote.net/	v.02	D Woollard

Disclaimer

The contents of this document at times refer to product names, manufacturer technologies and phrases intellectually owned. All of these references are acknowledged to be the intellectual property of the respective party and are not intended to be or inferred as a unique reference to the services of vmote Limited.

The instructions, notes and screenshots are reproduced with the sole intention to illustrate the purpose of this document. All procedures documented within this paper are purely to evidence the steps undertaken for a specific task.

vmote Limited will not be held accountable for any loss or corruption of data incurred to the intended audience environment while using this document.

Version numbers and product names were correct at the time of publishing.

Glossary

Acronyms and abbreviations used within this document.

Term	Description
DNS	Domain Name System
GPS	Global Positioning System
ISA	Internet Security Acceleration
MS	Microsoft
NTP	Network Time Protocol
PDC	Primary Domain Controller
VI	VMware Infrastructure
VM	Virtual Machine
WINS	Windows Internet Name Service

(Glossary table)

Background

At a recent customer engagement a review of multiple VMware Infrastructure environments revealed inconsistencies with the implementation of time synchronisation for the ESX host servers and guest virtual machines. Further investigation found the time drift between ESX hosts was considerable across multiple VI implementations for Live, Test & Development.

The Live (production) environment had recently been introduced evidencing only a single minute difference between hosts however, one of the legacy hosts was found to be 40 minutes behind actual time with another 50 minutes ahead. The legacy environment performed no tasks other than hosting VMs. VMotion and migrations between hosts could not be achieved due to the enthusiastic adoption of the VI resulting in highly utilised hosts.

The guest VMs were running Microsoft Windows based operating systems participating within an Active Directory. At least one Domain Controller was permitted to access the Microsoft default time source. It was purely the domain membership aspects that were keeping the 'user community' services with an accurate time and date.

This document

To provide a high level overview outlining connectivity for ESX host and Guest virtual machine time source acquisition.

Assumptions

Consideration should be given to the reference times, where they are obtained from and how many are used so outages can be mitigated.

Users of this document are versed with core infrastructure technologies.

Virtualisation time-keeping

Virtualisation by its very nature is prone to time drift and has to be managed accordingly so as not to disrupt the Guest virtual machines and their service responsibilities.

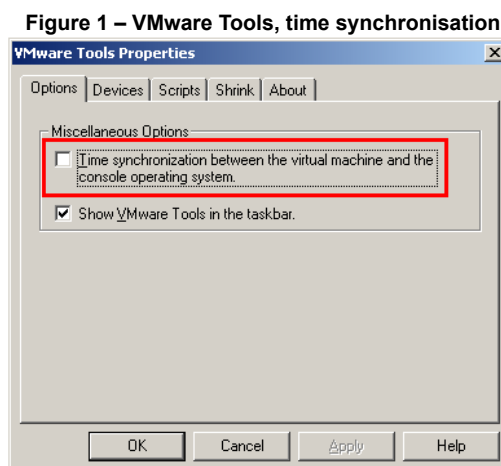
ESX hosts

The VMware hosts, where permitting, should have their NTP settings referencing a Layer 2 NTP device.

Guests - general

Consideration points:

- ➔ It is not recommended for Guest operating systems to synchronise with their ESX host as this results in a secondary non-direct source being provided. Virtual Machines should only be configured to synchronise their time with the host when they are singularly managed and not time dependant.
- ➔ The Time synchronisation option found in the VMware Tools will not adjust a clock time from the future to the present only from past to present.



Guests - Windows Active Directory

A Microsoft Windows Active Directory Domain uses time stamps as part of its formulation for Kerberos security tickets when managing object membership, this may be a computer or a user for example. In the event a domain controller receives a request it deems to violate its time validation process the request will simply be dismissed. Typical symptoms are a server is no longer accessible to users or they experiences difficulties with logging in.

The Layer 1 time source should be permitted to receive requests directly from the Domain Controller with the PDC Emulator role.

It is not recommended to allow access to or expose Domain Controllers to the internet, for this reason an NTP appliance or application server (such as Microsoft ISA Server) should handle the role.

Consideration should be given to the use of group policy to manage the expected state of the Windows Time Service for its domain members so that conflicts do not arise.

Further reading and configuring Windows Domain Controllers:

<http://technet.microsoft.com/en-us/library/cc784800.aspx>

Stratum Layers

For multiple environments where security boundaries require time synchronisation to be separated consider carving the areas of responsibility into Layer 2 Trusted & Non-Trusted devices.

Strata overview

Layer 1

The single most authoritative source within an organisation deriving it's updates from a known atomic / GPS reference point. This may be a locally installed device or sourced from the public domain (pool.ntp.org).

- ➔ Accepts communication from Layer 2 device(s)

Layer 2

Single or multiple devices used to communicate with Layer 1 and clients within the organisational network.

- ➔ Accepts updates from Layer 1
- ➔ Accepts communication from Layer 3

Layer 3

Common examples of connectivity to Layer 2.

- ➔ Network Switches
- ➔ Network Routers
- ➔ Network Services (DNS / WINS)
- ➔ Non MS Windows Servers
- ➔ VMware ESX hosts
- ➔ Firewalls

The final page of this document provides an interaction example of the layers discussed above.

Pictorial representation

Figure 2 – Example NTP overview using locally hosted NTP device

