

Parallels® Virtuozzo Containers

White Paper

Live Migration

www.parallels.com

Version 2.0



Table of Contents

Abstract	3
Live Migration Definition	3
Pre-requisites for Live Migration	3
Benefits of Live Migration	3
Technical Description Live Migration	4
A Smarter Migration Technology	4
A Flexible Migration Technology	4
Migrating Network Connections	4
End-User Network Connections	5
Summary	5

Abstract

This document was created to provide the technical background and an explanation of how live migration has been implemented in Parallels Virtuozzo Containers for Linux. The document contains an overview regarding the core server technology, networking implications and other concerns regarding the capability.

Live Migration Definition

Live migration is the capability to move a virtual environment (VE) and its corresponding applications and data between two Parallels Virtuozzo Containers servers without service interruption. The Virtuozzo live migration includes the following performance specifications:

- All the processes in the migrating VE will not be stopped or restarted. Even processes or services started manually (not using system startup scripts), will continue to run in the VE after the migration has finished.
 - All open files will remain accessible for processing and user access during and after the migration.
 - All network connections created before the migration will remain intact.
 - Some examples of how this will look to the end-users include:
 - An interactive shell session over ssh or telnet will not be disconnected.
 - ftp downloads from a VE ftp server will not be broken.
 - Transactions sent to a database will be delivered to the database.
 - Customers entering information into ecommerce applications will not lose their sessions nor their data.
-

Pre-requisites for Live Migration

Two networked Parallels Virtuozzo Containers servers...

Live migration has become one of the standards of technology for server virtualization solutions. Enabling complex features such as this often times requires extensive support and configuration amongst hardware and software components. The requirements for Virtuozzo live migration are: the core virtualization software on at least two servers and a regular network connection between the servers. The management toolset is helpful, it makes the migration process as simple as a few mouse selections in the interface, but it can be done through the command line interface. Like all Virtuozzo technology, live migration is designed to be supported by the basic network infrastructure, not requiring large infrastructure investments such as SANs.

Benefits of Live Migration

Why is live migration important and why is it important that it can be implemented on any regular network? IT departments typically feel the demands of supporting an organization and maintaining the hardware and software infrastructure so that everything runs efficiently. Historically, all planned outages for software or hardware upgrades are done after hours to minimize the impact to the end users. Sometimes an outage is impending and inevitable. For example, hardware technology has amazing advances that notify the administrator of errors, providing a window for migrating applications and data to a new server before the hardware crashes. Being able to migrate the application at any

time of the day will make it much easier to respond to impending outages and eliminate the impact on users.

Previously, investing in SANs and expensive supporting software was the only option to have this type of freedom and management capability over virtualized servers. The commitment to providing a transparent migration was so great it was only reserved for the most critical of applications. With Virtuozzo for Linux, live migration can be performed on any two Virtuozzo servers on any network, making this capability available for all applications, not just the most critical applications. No additional configuration is required; live has been added as an additional option to the existing migration toolset.

Technical Description Live Migration

Live migration is implemented in the Parallels Virtuozzo Containers kernel using CPT or checkpointing technology. CPT enables Virtuozzo to store the state of the VE and its processes in a file and also restore the state on another Virtuozzo server. The restore capability is implemented to restore the state of processes including running processes, opened files, network sockets and connections.

Therefore, the Virtuozzo CPT implementation saves the full state of a VE on the source server and transfers that state together with the VE files to the destination server and fully restores the VE. The migration maintains application and VE availability and makes migration transparent to the end users.

A SMARTER MIGRATION TECHNOLOGY

CPT works only with the set of processes running in a VE with their full state. The Virtuozzo implementation is streamlined; it only transfers relevant data, making it extremely fast and efficient. It does not transfer or restore any non-relevant data or parts of the operating system. Some examples of data that does not need to be migrated includes: fe, system buffers, caches, and others. Since only relevant information is stored in the migration file, the file size is small and may be quickly transferred over the network.

A FLEXIBLE MIGRATION TECHNOLOGY

Live migration provides two methods for transferring the virtualized server. Method one includes capturing the data, processes, and memory into one complete file for transferring to the new server. Method one is very efficient and complete. Method two moves the data and processes initially and then migrates the memory images on demand over the network, after the VE startup when it attempts to access its memory pages. The memory migration is accomplished in a similar way to an operating system delivering memory pages from the swap when required by the application. The benefits to method two are quite simple, if the physical server is experiencing some significant errors and may crash, it may be preferable to move the virtualized server more quickly and if the original server does crash, the only service interruption will be allowing enough time for the new location to reload its memory.

Migrating Network Connections

Maintaining the network connections are one of the most important components of live migration. The Virtuozzo implementation of the CPT mechanism enables saving and restoring opened network connections along with the rest of the VE state. Additionally, during the live migration the Virtuozzo VE migration tools automatically manage the IP address reconfiguration. The IP reconfiguration includes:

- taking the existing VE IP address and assigning it to the VE on the destination server.
- sending an ARP announcement.
- updating the ARP cache of networked computers and devices, including those performing external traffic routing to the VE.

Once the reconfiguration is completed, all traffic sent to the original VE will now be automatically routed to the new VE, which is ready to accept the traffic.

END-USER NETWORK CONNECTIONS

An established TCP connection is represented by the state of two network sockets: local, which would refer to the original VE to be migrated, and remote which refers to the user's computer that is connecting to the VE. The CPT technology is able to fully save and restore the local sockets; however, the CPT implementation does not send any controlling packets to the remote sockets so the remote socket does not change its state. When a migrated VE is restored on the destination server, the VE migration tools automatically manage the TCP update. As a result, users accessing the VE will not notice any downtime during the migration.

Summary

Live migration is an important advance in the Parallels Virtuozzo Containers technology. The design takes into consideration all processes and networking to ensure that end users are not affected by the migration. It also is designed to be intuitive and simple to implement, a migration can be conducted in literally four mouse clicks, making it simple for any administrator to migrate at any time of the workday without fear of impacting service availability.