



An Oracle Technical Article
September 2010

Virtualizing Your Enterprise Infrastructure: Zones Physical to Virtual (P2V)

Introduction	1
Introducing P2V with Oracle Solaris Containers in Oracle Solaris 10 9/10	2
Four Steps to Virtual Contentment	3
Check	3
Capture	5
Construct	5
Configure	5
Summary	6
More Information	6

Introduction

As enterprises continue to adopt virtualization technologies into their infrastructure, the processes for planning and rapidly realizing the benefits that virtualization brings become increasingly important. Historically, the drive to adopt virtualization has been to realize cost savings through consolidation of older, more costly and underused systems. Inevitably, the additional benefits that virtualization brings, particularly in the area of business agility, soon result in the examination of the entire “physical” environment and an assessment of how far, and more importantly, how easily virtualization can be adopted.

To ease the adoption of a virtualized environments Oracle continues to innovate across the stack, and in particular, with Oracle Solaris, the world's best mission-critical, virtualized, enterprise-class operating system. In the latest version, Oracle Solaris 10 9/10, Oracle has added new Physical to Virtual (P2V) features for its Oracle Solaris Containers virtualization technology. This article explains how to use P2V with Oracle Solaris 10.

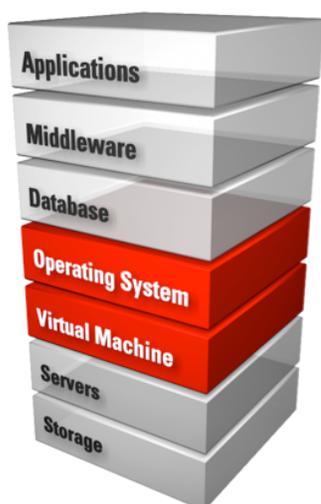


Figure 1: The Stack

Introducing P2V with Oracle Solaris Containers in Oracle Solaris 10 9/10

Physical to Virtual (P2V) describes the process of taking the applications, OS, and data from a physical server or system and migrating them onto a virtual guest hosted on a virtualized platform. The functionality, generally provided in the form of a tool, is especially important when performing consolidations of existing systems onto virtualization enabled platforms.

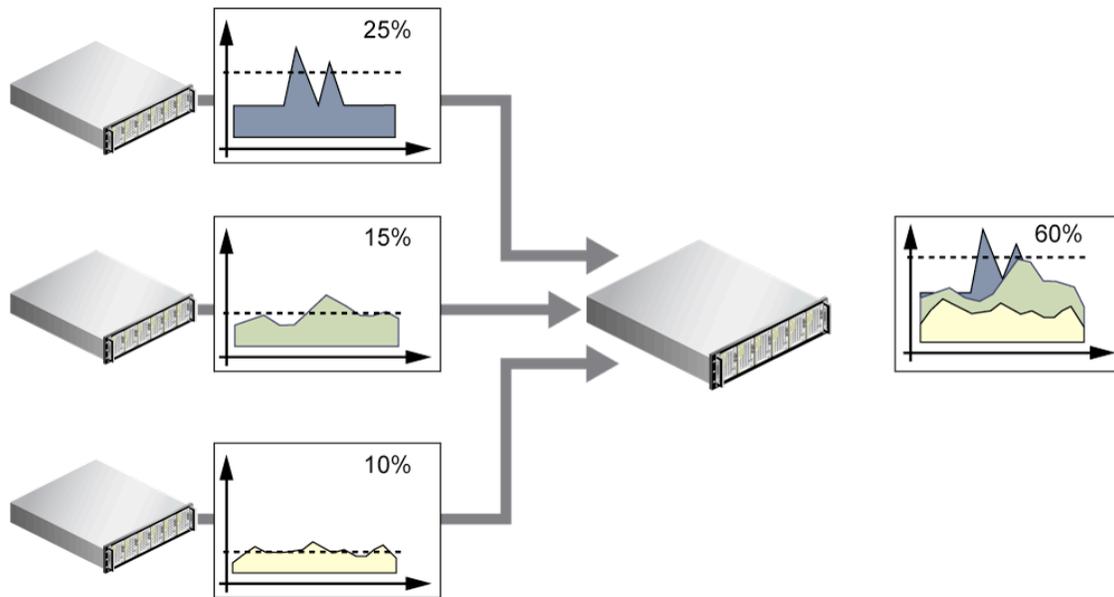


Figure 2. Virtualization Example

Oracle already provides P2V tools as part of the Oracle Solaris 8 Containers and Oracle Solaris 9 Containers products. With Oracle Solaris 8 Containers and Oracle Solaris 9 Containers, Oracle enables the virtualization of Oracle Solaris 8 and Oracle Solaris 9 systems onto Oracle Solaris 10 systems.

Now, with the release of Oracle Solaris 10 9/10, Oracle is introducing P2V functionality that allows you to apply P2V not only your legacy Oracle Solaris 8 and Oracle Solaris 9 systems but also to Oracle Solaris 10 physical systems. With this release, it is possible to plan, schedule, and address your entire Oracle Solaris infrastructure to be virtualized into Oracle Solaris Containers. Not only can you benefit from additional space that is available through reduced footprint requirements and cost savings that result from reduced power and cooling needs, but by using the P2V features, you can benefit from lower risk, less effort, and less complexity when adopting a virtualization strategy.

Adopting Oracle Solaris Containers, a built-in virtualization technology and a core part of the Oracle Solaris OS, will protect your applications and provide an easy upgrade path to future Oracle Solaris releases. With the P2V feature in Oracle Solaris 10 9/10, this step becomes even simpler and it is provided without any additional charge.

Four Steps to Virtual Contentment

To move your existing Oracle Solaris 10 physical systems into Oracle Solaris Containers, follow these four simple stages: check, capture, construct, and configure.

Check

Good planning and preparation always increases the chance of achieving a successful outcome. Moving a physical system to a virtual environment is no different. There are lots of issues to consider when virtualizing your infrastructure, but that discussion is outside the scope of this document. However, it is worthwhile keeping in mind the following issues as you start to plan your migrations.

Selecting Candidate Source Systems

When considering the migration of systems to virtual guests, it is beneficial to assess the possible source systems and select those that are most suited to be virtualized. There are several scenarios that are particularly well-suited to being virtualized with Oracle Solaris Containers. They include, but are not limited, to the following:

- *Combining several small Oracle Solaris servers.* Significantly decrease your TCO by combining smaller isolated servers onto a single larger enterprise system. With careful placement of the virtual guest, performance benefits are also possible because localized system communication can be used instead of more performance-expensive off-system communication.
- *Consolidating many underutilized servers.* Achieve higher usage on a lower number of systems and also minimize administration overhead by having fewer physical systems to manage.
- *Replacement of multiple legacy servers.* Switch to fewer more cost-effective, newer servers.

Is the Source System a Good Candidate for Migration?

Once a source system has been selected as a possible candidate for P2V, performing a set of further checks for that specific system is desirable to help further select and plan for the migration. These additional checks could include the following:

- *What applications are running?* Consolidation is all about the applications, so be aware of which applications are running on the system. The vast majority of applications will run perfectly within a zone; however, there is a small portion of software that might not be suitable for deployment. This could include volume manager software or software with kernel components. The most common unsuitable component is the NFS server. Put in place a plan for those components that need to run elsewhere. For example, NFS servers are generally deployed in the global zone.
- *What services are running?* Be aware which additional services have been activated on a system in case they need to be activated on the target system also.

- *What is the network configuration?* Assess the network configuration. Are there specific configuration requirements that will need to be duplicated within the zone? Does your application need to change the IP address? If so, you will need to use Exclusive IP within a zone, and the number you can assign is limited to the number of physical ports available. It might be better to plan to spread these types of guests across several target systems.
- *What are the local storage requirements and file systems being used?* How much local storage is being used? This storage will need to be allocated on the target system. What file systems are your applications running on? You will need to take care when creating an archive of a system that has an Oracle ZFS root (see below). Finally, what storage is being utilized? For example, you might want to exclude certain file systems from the system archive you will create, and SAN storage will most likely be migrated to the global zone of the target system.
- *What is the kernel/release version?* It is an unlikely scenario, but if the source system image to be installed via P2V is newer than the target system, then the P2V migration will fail.
- *Are Oracle Solaris Containers already running on the system?* Zones, like other virtualization technologies, do not accommodate nested virtualization. A system archive that contains zones has any “nested” zones uninstalled by the P2V operation. The data is not lost; however, the zone configurations no longer exist. This is not an issue; simply migrate the source system zones across to the target system before archiving the source system. Update on Attach will deal with any required updates to the moved zones, and it might not even be necessary to migrate the source system after moving the zones.
- *Are there any systems settings I need to move across?* System settings from `/etc/system` that are unique to the source system might need to be implemented on the target system. Take time to check whether any of these are needed.

Does the Target System Meet the Requirements for Running the Source System?

Finally, it is also worthwhile to check that the target system will host the virtual guests. Does the target system have enough spare computing capacity and memory? Is their sufficient disk space, and does the system have access to the appropriate IO devices? Is there sufficient IO capacity (both network and remote storage)?

Capture

Even though the system identity will be preserved in the created guest, it is worthwhile to capture some source system information prior to initiating the archive of the system. Information to capture could include the following:

- Host name
- Network settings
- Storage
- `hostid`
- Root password

Note: The `hostid` is not preserved by the P2V operation; the migrated guest will adopt the `hostid` of the target system. To preserve the `hostid`, follow the method described in the “Configure” section below.

Construct

Once the details of the source system have been captured, the system is ready to be archived. Carefully consider which file systems you want to archive. It may not be necessary to archive file systems on SAN attached storage. For instance, this storage might already be connected to and available on the target system. This remote storage can be added to the zone when the zone is created.

The recommended method for creating an archive is using `flarcreate`. There is one nuance in this step to be aware of. When creating an Oracle Solaris Flash archive, or `flar`, of an Oracle Solaris 10 system that has an Oracle Solaris ZFS root, by default, the `flar` will actually be an Oracle Solaris ZFS send stream, which can be used to recreate the root pool. This image cannot be used to install a zone. The `flar` must be created with an explicit `cpio` or `pax` archive when the system has a Solaris ZFS root. To address this, use the `flarcreate` command with the `-L` archiver option, specifying `cpio` or `pax` as the method for archiving the files.

Configure

The final step in the process is to configure the target zone and install your source system archive into it. Once the target zone has been created, install the archive into it using the `zoneadm` command.

At this point, the option is available to preserve or discard the system information. When performing a strict P2V operation, the system information should obviously be preserved. However flexibility is provided by allowing the system information to be discarded. For example, it might be desirable to reset some of the guest image settings as the system moves to a new environment, and perhaps the host name needs to be changed. Another advantage of this approach is that the guest archive could be used as a template for multiple installs of the application environment onto multiple machines.

The `zoneadm` command implicitly checks whether the zone configuration is valid, and it reports errors that it finds. In addition to unpacking files from the archive, the install process performs checks, required post-processing, and other functions to ensure that the zone is optimized to run on the host. These include the following:

- If necessary, an “Update on Attach” is performed to make sure the guest zone patch level is in sync with the global zone.
- Oracle Solaris Service Management Facility (SMF) services that are not usable within a zone are deleted or disabled, as necessary. The information about what SMF services to disable is primarily driven by the manifests in the SVr4 packages.
- Network configuration will be adjusted depending on whether the zone is shared-stack or exclusive-stack.
- NFS serving is disabled.
- The `vfstab` file is adjusted so that local file systems from the original system are disabled.
- Any zones installed on the original system will be uninstalled and deleted from the image (zones do not nest). Note that the zonepaths are not removed from the image; this is a task left to the system administrator.

Once the install is finished, the zone is ready to be booted.

Summary

With the release of Oracle Solaris 10 9/10, the adoption of Oracle Solaris Containers into the enterprise datacenter environment becomes even simpler. Adopting Oracle Solaris Containers now provides immediate benefit from the power of virtualization in Oracle Solaris 10, while also ensuring an easy upgrade path to future Oracle Solaris releases.

More Information

- Oracle Solaris virtualization products: <http://www.oracle.com/technetwork/server-storage/solaris/overview/virtualization-163570.html>
- Oracle Solaris 10: <http://www.oracle.com/us/products/servers-storage/solaris/index.html>
- Oracle Solaris Containers, Oracle Solaris 8 Containers, and Oracle Solaris 9 Containers: <http://www.oracle.com/technetwork/server-storage/solaris/containers-169727.html>
- *System Administration Guide: Oracle Solaris Containers-Resource Management and Oracle Solaris Zones* [see “Chapter 24 Oracle Solaris 10 9/10: Migrating a Physical Oracle Solaris System Into a Zone (Tasks)”]: <http://docs.sun.com/app/docs/doc/817-1592/>



Virtualizing Your Enterprise Infrastructure:
Zones Physical to Virtual (P2V)
September 2010
Author: Duncan Hardie

Oracle Corporation
World Headquarters
500 Oracle Parkway
Redwood Shores, CA 94065
U.S.A.

Worldwide Inquiries:
Phone: +1.650.506.7000
Fax: +1.650.506.7200
oracle.com



Oracle is committed to developing practices and products that help protect the environment

Copyright © 2010, Oracle and/or its affiliates. All rights reserved.

This document is provided for information purposes only and the contents hereof are subject to change without notice. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. UNIX is a registered trademark licensed through X/Open Company, Ltd. 0310

SOFTWARE. HARDWARE. COMPLETE.