



Guidelines for Using Virtual Environments for Ex Libris Products

CONFIDENTIAL INFORMATION

The information herein is the property of Ex Libris Ltd. or its affiliates and any misuse or abuse will result in economic loss. DO NOT COPY UNLESS YOU HAVE BEEN GIVEN SPECIFIC WRITTEN AUTHORIZATION FROM EX LIBRIS LTD.

This document is provided for limited and restricted purposes in accordance with a binding contract with Ex Libris Ltd. or an affiliate. The information herein includes trade secrets and is confidential.

DISCLAIMER

The information in this document will be subject to periodic change and updating. Please confirm that you have the most current documentation. There are no warranties of any kind, express or implied, provided in this documentation, other than those expressly agreed upon in the applicable Ex Libris contract. This information is provided AS IS. Unless otherwise agreed, Ex Libris shall not be liable for any damages for use of this document, including, without limitation, consequential, punitive, indirect or direct damages.

Any references in this document to third-party material (including third-party Web sites) are provided for convenience only and do not in any manner serve as an endorsement of that third-party material or those Web sites. The third-party materials are not part of the materials for this Ex Libris product and Ex Libris has no liability for such materials.

TRADEMARKS

"Ex Libris," the Ex Libris bridge, Primo, Aleph, Alephino, Voyager, SFX, MetaLib, Verde, DigiTool, Preservation, URM, Voyager, ENCompass, Endeavor eZConnect, WebVoyage, Citation Server, LinkFinder and LinkFinder Plus, and other marks are trademarks or registered trademarks of Ex Libris Ltd. or its affiliates.

The absence of a name or logo in this list does not constitute a waiver of any and all intellectual property rights that Ex Libris Ltd. or its affiliates have established in any of its products, features, or service names or logos.

Trademarks of various third-party products, which may include the following, are referenced in this documentation. Ex Libris does not claim any rights in these trademarks. Use of these marks does not imply endorsement by Ex Libris of these third-party products, or endorsement by these third parties of Ex Libris products.

Oracle is a registered trademark of Oracle Corporation.

UNIX is a registered trademark in the United States and other countries, licensed exclusively through X/Open Company Ltd.

Microsoft, the Microsoft logo, MS, MS-DOS, Microsoft PowerPoint, Visual Basic, Visual C++, Win32, Microsoft Windows, the Windows logo, Microsoft Notepad, Microsoft Windows Explorer, Microsoft Internet Explorer, and Windows NT are registered trademarks and ActiveX is a trademark of the Microsoft Corporation in the United States and/or other countries.

Unicode and the Unicode logo are registered trademarks of Unicode, Inc.

Google is a registered trademark of Google, Inc.

Copyright Ex Libris Limited, 2013. All rights reserved.

Document released: December 16, 2013

Web address: <http://www.exlibrisgroup.com>

Table of Contents

1	About This Guide	4
2	About Virtual Environments	5
3	General Recommended Configuration Settings	6
	Storage Architecture	6
	Communication	7
	Memory Overcommit	7
	Resources for Hypervisor	7
	Hard Resource Allocation	8
	Power Management	8
	Time Server	8
4	Resource limitations	10
	VMware Resource Limitations	10
	KVM Resource Limitations	11
5	VMware – Additional Info	12
	Example of Resource Allocation Setting in VMware	12
	<i>Additional Information for Guests Running MySQL Servers (SFX)</i>	14
6	Solaris Zones – Additional Info	15
	Service Level – Reservations and Limitations	15
	ZFS	16
7	Migrating Virtual Servers	17
	DRS Cluster, Vmotion and Live migration	17
	Migrating Zones	18
8	Support Policy for Virtual Infrastructure Usage	19
9	Disclaimer for Virtual Infrastructure Usage	20
10	List of Applications/Versions Certified in Virtual Environments	21

1

About This Guide

The purpose of this guide is to define how to integrate Ex Libris products into a virtualized server infrastructure using VMware ESX, Redhat KVM, or Solaris Zones.

This guide is applicable to the following Ex Libris products:

- Aleph
- Voyager
- MetaLib
- Verde
- SFX
- Primo
- DigiTool
- Rosetta

2

About Virtual Environments

VMware ESX, Redhat KVM, Xen and Solaris Zones (Containers) are the popular virtualization tools in the IT world.

VMware is capable of managing and monitoring the whole infrastructure system of virtual servers using many algorithms and proprietary technologies. These technologies initially seem to be transparent to the virtual server operating system; but since they control and allocate the physical hardware resources, the desired operation on the hosted operating system may be affected. More information about VMware ESX can be found at http://www.vmware.com/files/pdf/esx_datasheet.pdf. More information on performance best practices for VMware ESX can be found at http://www.vmware.com/pdf/Perf_Best_Practices_vSphere5.1.pdf.

Redhat KVM (Kernel Based Virtual Machine) is a virtualization infrastructure for Linux on x86 hardware containing virtualization extensions (Intel VT or AMD-V). It consists of a loadable kernel module that provides the core virtualization infrastructure. More information on Redhat KVM can be found at <http://www.redhat.com/products/cloud-computing/virtualization/>

Solaris 10/11 Zones, also known as Solaris Containers, are virtual execution environments within a server that provide software and security isolation. In addition, they provide resource containment and control while reducing management costs and increasing resource utilization and flexibility. A container is a zone with constrained resources and the two terms are often used interchangeably. Every server has a global zone, that is the entire server. In this document, the term zone specifically refers to non-global zones. More information about Solaris 10/11 Zones can be found at <http://www.oracle.com/technetwork/server-storage/solaris/containers-169727.html>

Features of virtual environments may also benefit Ex Libris applications including, but not limited to, fast disaster recovery and allocation of more resources during heavy system loads.

3

General Recommended Configuration Settings

Storage Architecture

When there are many virtual servers using the same storage, the overhead on storage and communication can be large.

Ex Libris products are very database centric and use the database heavily, especially on initial system loads. It is recommended that you use a storage system that has sufficient performance and IOPS (Input/Output Operations Per Second) to support a load of many virtual servers, especially if this storage is used for the database server.

To support the communication overhead of using NFS or iSCSI storage protocols for virtualization, Ex Libris recommends that a separate physical Ethernet port of no less than 1GbE be allocated on each physical server exclusively for storage connection. Large systems should use a 10GbE network interface. This is in addition to the Ethernet port used for communication with the LAN. For iSCSI it is recommended to install a dedicated iSCSI HBA card to offload the iSCSI traffic from the onboard interfaces. If you are using zones, configure loop-back file systems (LOFS) to avoid having to use NFS. A zone can be configured as an NFS client but not as an NFS server.

Fiber Channel Protocol (FCP) usually gives superior performance when used as storage for a database, as it works with block-level access. When working with FCP, there is no need for a dedicated NIC (Network Interface Controller) for storage, as all traffic to storage is passed through the fiber channel that is specially designed to handle storage communications.

In the virtual infrastructure, storage is usually common to all physical servers controlled by ESX. Since it is common, it remains the largest single point of failure in the system. Ex Libris recommends using a storage system that has protection and disaster recovery capabilities such as RAID, Snapshots, and Resiliency with regular backup, fast recovery,

and so forth. If you are using SAN-based network storage (FC / iSCSI / FCoE), in addition to the above, it is recommended that you:

- Use the Para Virtual SCSI driver. (Refer to the following KB issue:
 - http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=1010398)
 - https://access.redhat.com/site/documentation/en-US/Red_Hat_Enterprise_Linux/5/html/Virtualization/chap-Virtualization-KVM_Para_virtualized_Drivers.html
- Configure the virtual servers to work with RDM (Raw Device Mapping).

Communication

Communication capabilities of the physical servers can also be a single point of failure in a virtual infrastructure. When there are many virtual servers using the same physical server, this can be crucial.

Ex Libris recommends using a bond/multipath on the physical server communication ports.

Consideration must be given to the type of zone being created. Zones with dedicated NICs can only communicate via the network. Whereas, zones with shared NICs do not create network traffic. Dedicated NICs provide greater security. Shared NICs reduce network bandwidth and increase performance. The option to dedicate a NIC to a zone is a feature introduced in Solaris 10 HW 8/07 and later.

Memory Overcommit

Virtual systems support memory overcommit and can store the memory of guests in swap space. Guests run slower if they are swapped frequently. It is highly recommended NOT to overcommit memory for virtual environments that are in production.

Resources for Hypervisor

When working on capacity planning on virtual environments you need to reserve resources (CPU and RAM) on the host level for management of the hypervisors and virtualization platform. The amount of resources varies and depends on the virtualization platform, the number of guests per host, and the applications running on the guests.

Hard Resource Allocation

It is very important for Ex Libris products to have sufficient resources available in order for a production environment to operate successfully.

Note: Allocate only as much virtual hardware as required for each virtual machine. Provisioning a virtual machine with more resources than it requires can, in some cases, reduce the performance of this virtual machine as well as other virtual machines sharing the same host.

Ex Libris applications use the inherent operating system caching abilities as a tool to promote application response time optimization.

Oracle also uses and requires a large amount of memory from its startup, even if this memory is not immediately needed.

You need to verify that the virtual server has all the required resources (CPU and memory) reserved for the virtual server. Memory and CPU reservation should be set according to the sizing and configuration requirements Ex Libris provides. If the resources are not reserved, the virtual machine competes with other virtual machines for the available CPU and memory and waits in the queue if there are not enough resources available on the host machine.

Ex Libris is not able to ensure product performance and response time if the servers are not configured with hardware resource allocation.

For an example from VMware on how this is configured, see [Example of Resource Allocation Setting in VMware](#) on page 12.

Power Management

If you want the installed Ex Libris application to have a quick response time, Ex Libris recommends that `always-on` be selected in the power management settings of the containing virtual server.

This prevents the server from going into stand-by mode when idle. When a server goes into stand-by mode, the response time of the application is delayed due to the need to warm up and reopen the containing virtual server.

Time Server

There are products that can be installed in a multitier configuration. This means that they are distributed over several virtual servers. In a distributed installation, it is important for the internal clocks of the servers to be synchronized.

Ex Libris recommends using NTP to synchronize the server clocks.



4

Resource limitations

VMware Resource Limitations

There is a limit to the number of cores (virtual or physical) that a virtual guest can run. This depends on the VMware version and the Red Hat version

As for the VMware platform:

Max Virtual CPU per VM (vCPU)

Server	Standard, and Advanced	Enterprise	Enterprise+
ESX Server 4.x	4 vCPU	4 vCPU	8 vCPU
ESX Server 5.0	8 vCPU / 24GB RAM	8 vCPU / 32 GB RAM	32 vCPU / 48 GB RAM
ESX Server 5.1	8 vCPU / 32GB RAM	32 vCPU / 256 GB RAM	64 vCPU / 1 TB RAM

Max vCores (number of cores for ESX server – Physical Host)

Server	Logical CPU per Host
ESX Server 4.0	64
ESX Server 4.1	128
ESX Server 5.x	2048

KVM Resource Limitations

There is a limit to the number of cores (virtual or physical) that a virtual guest can run on KVM. These numbers are high and are more than enough for all Exlibris products:

The limit is 160 cores and 1TB or RAM per host

	Standard, and Advanced	Enterprise	Enterprise+
Host RHEL 5.x		160 virtual CPU's	1TB
Host RHEL 6.x		160 virtual CPU's	2TB
Guest KVM RHEL 5.x	Full Virtualization	16 virtual CPU's	512GB
Guest KVM RHEL 5.x	Full Virtualization	160 virtual CPU's	2TB
Guest KVM RHEL 5.x	Para Virtualization	160 virtual CPU's	1TB
Guest KVM RHEL 5.x	Para Virtualization	160 virtual CPU's	2TB

5

VMware – Additional Info

Example of Resource Allocation Setting in VMware

Reservation is defined by VMware as “A certain number of CPU cycles / Memory reserved for this VM, measured in GHz / MB. The VMkernel chooses which CPU(s), and may migrate.”

You need to verify that the virtual server has all the required resources (CPU and memory) reserved for the virtual server. Memory and CPU reservation should be set according to the sizing and configuration requirements Ex Libris provides. If the resources are not reserved, the virtual machine competes with other virtual machines for the available CPU and memory and waits in the queue if there are not enough resources available on the host machine.

Ex Libris is not able to ensure product performance and response time if the servers are not configured with hardware resource allocation.

The following should be used:

CPU: Set reservation to the CPU speed multiplied by the number of cores required (in MHz).

Memory: Set reservation to the amount of GBs required.

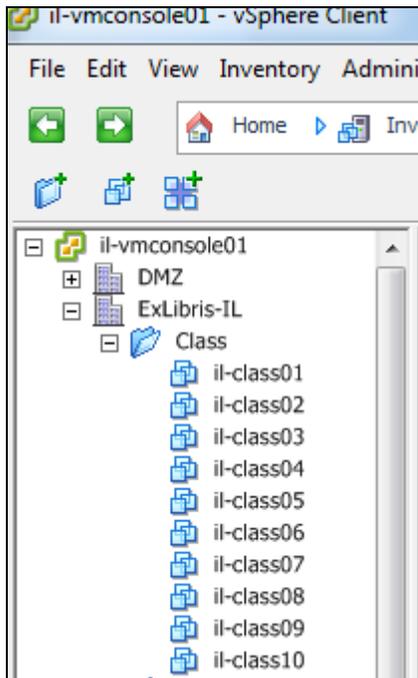
The following is an example from vSphere v5.1 but looks very similar to vSphere v4.x and 5.0. For KVM and Solaris Zones, refer to the relevant documentation described in the About Virtual Environments section on page 5 .

Shares: Set to Normal or High

Limitation: Set to unlimited or equal to the reservation.

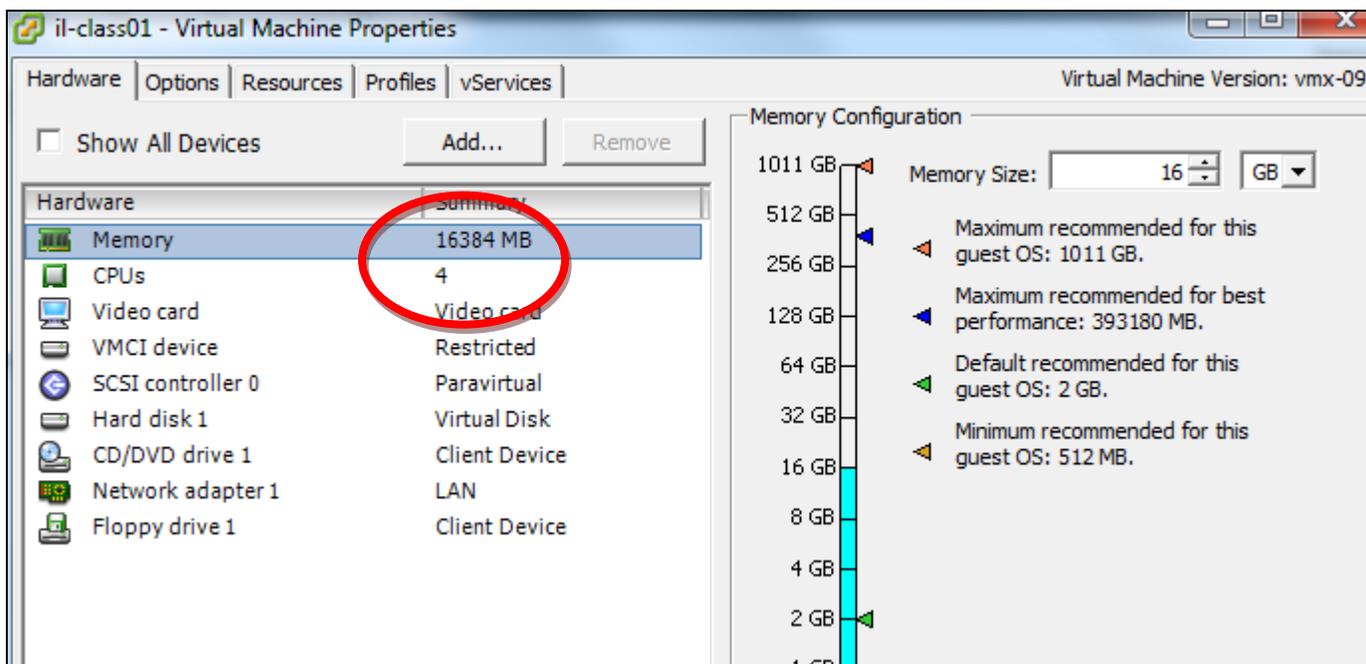
To check your resource reservation in vSphere v5.1:

- 1 Right-click your virtual machine or cluster and select **Edit Settings**.



- 2 Edit the settings.

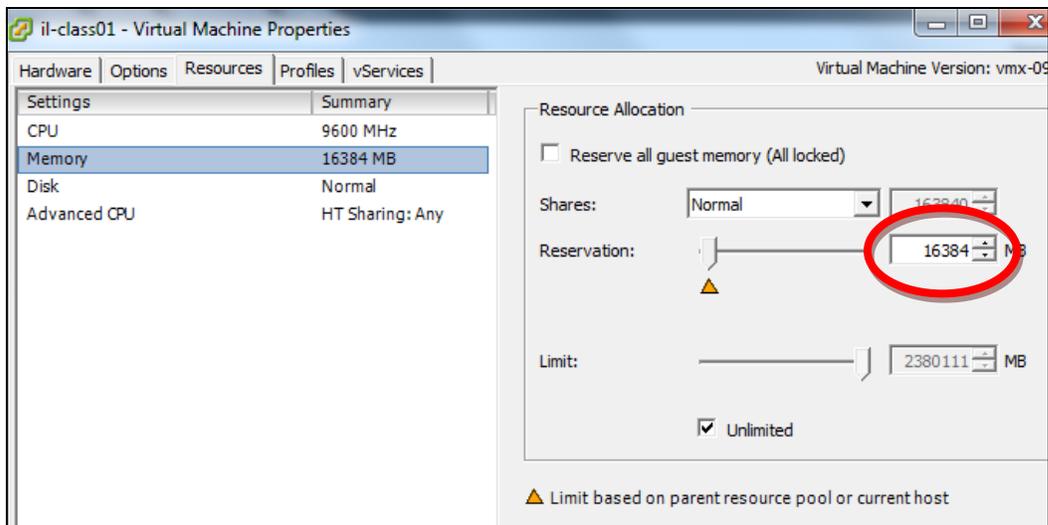
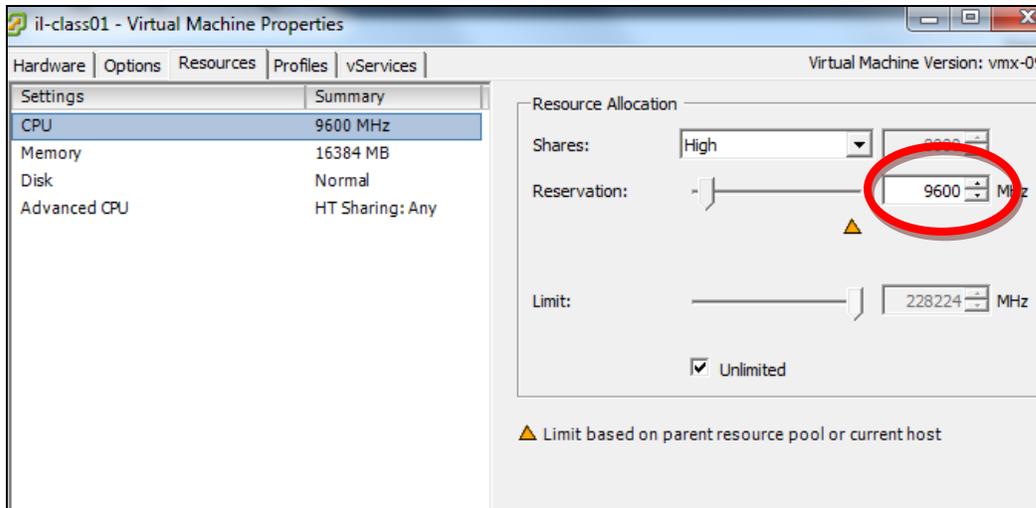
In the example below, the server is configured with four cores and 16GB of RAM.



To check your resource reservation in vSphere v5.1:

- 1 Right-click your virtual machine or cluster and select **Edit Settings**.
- 2 Click the **Resources** tab.

In the example below, the reservation is set for 4 cores X 2400 Mhz and memory is set to 16 GB.



Additional Information for Guests Running MySQL Servers (SFX)

For additional recommendations that are focused mainly on servers running the MySQL database, refer to the following articles:

- <http://dev.mysql.com/doc/refman/5.1/en/ha-vm-commonissues.html>
- <http://www.vmware.com/resources/techresources/10007>

6

Solaris Zones – Additional Info

Service Level – Reservations and Limitations

Since virtualization is about maximizing resource utilization, it is important for zones to be configured so that excess processor resources be made available to zones that need extra resources while restricting any zone's activity from impairing the use of the other zones.

Ex Libris recommends using the Fair Share Scheduler for each configured zone including the global zone. The ability to configure a persistent share value for the global zone is a feature introduced in Solaris 10 HW 8/07.

Memory is a resource that can be constrained per zone using the rcapd daemon. Rcapd computes memory consumption using the Resident Set Size (RSS) of each process.

Note: This incorrectly counts shared memory; and as a result, overstates the amount of consumed memory.

Ex Libris recommends that you refrain from configuring rcapd for any zone that contains applications that use shared memory.

Shared memory and swap may both be constrained per zone preventing exhaustion of these resources from affecting other zones including the global zone. These features were introduced in Solaris 10 HW 8/07.

Ex Libris recommends that these values be configured appropriately for the products running in zones.

ZFS

The ZFS file system introduced in Solaris 10 HW 6/06 offers the benefits of a file system and volume manager in a single interface. The capabilities of ZFS meet the needs for zones in a way that is unparalleled by UFS (Unix file system).

Ex Libris recommends that you refrain from using ZFS as the root file system of a zone unless Solaris 10 HW 10/08 or later is installed.

7

Migrating Virtual Servers

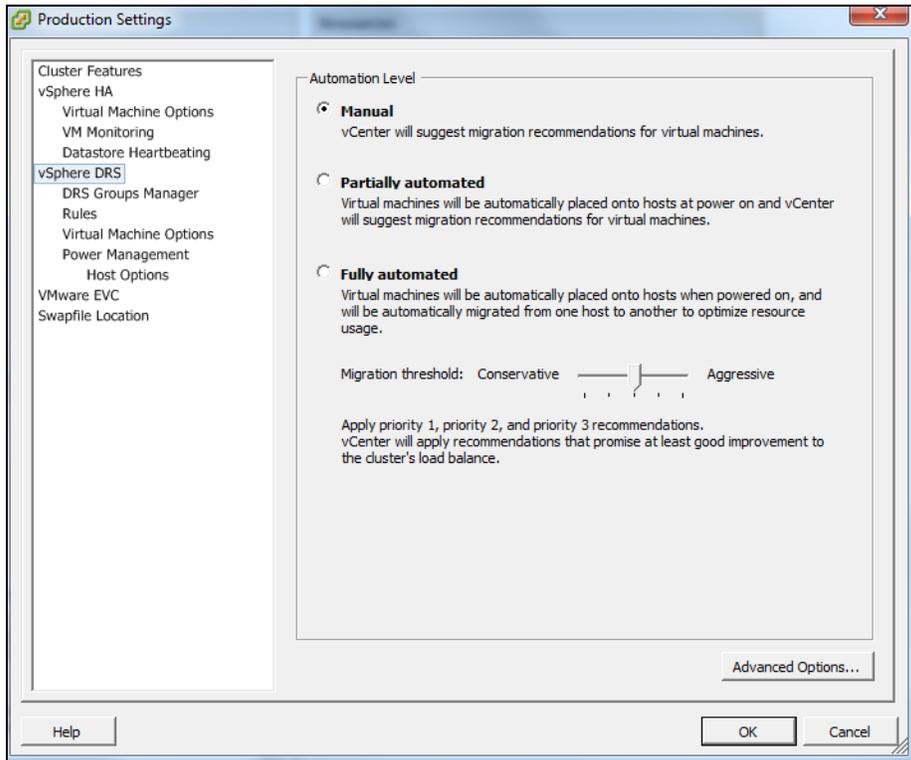
DRS Cluster, Vmotion and Live migration

VMware and KVM have the ability to balance virtual machines and migrate them across different physical hosts in a cluster. Vmotion is the process of migration of a virtual server from host to host in VMware. In RHEL KVM the process is called live migration.

It is recommended that you set the DRS automation level to **manual** so the servers are not automatically migrated. Otherwise, you may encounter an unresponsive system, system failure, or serious performance degradation.

From experience, Ex Libris recommends that if you need to migrate the virtual server, schedule the migration only during maintenance time or at times when system activity is low.

Set the DRS automation level from the following page in VMware:



Migrating Zones

Solaris 10 / 11 have the option to migrate zones between physical machines. The HW 5/08 release introduced the feature of a trial migration to test for successful migration before doing the work. The HW 10/08 release introduced an Update on Attach feature that upgrades the zone's packages and patches to the versions on the new host.

Migrating a zone requires the zone to be down, incurring a planned outage of the applications within the zone.

8

Support Policy for Virtual Infrastructure Usage

In order to receive support for Ex Libris products installed in a virtual infrastructure, the following requirements must be met:

- Resources allocated (and reservation) for the virtual machine must meet the standard hardware requirements for each application or part of an application as set forth by Ex Libris for the customer.
- The customer's technical staff should be certified for the configuration and setup of the virtual infrastructure in use at the customer's site.
- To ensure correct support procedures, customers should notify Ex Libris of the use of a virtual infrastructure. Notification should be sent to support.admin@exlibrisgroup.com.
- If the source of the problem is not identified on virtual infrastructure, the customer should be prepared to transfer the application to a physical server to enable Ex Libris to isolate the problem.
- Ex Libris makes an effort to identify the problem while enabling the customer to continue using the virtual infrastructure. If a solution is not identified, the customer is expected to install the product on a physical server.

9

Disclaimer for Virtual Infrastructure Usage

Ex Libris is not responsible for any software failure due to improper use of a virtual server environment for Ex Libris products including, but not limited to, system failure or performance degradation due to insufficient allocated system resources, system failure due to virtual server migration, and irregular system activity that is related to the virtual infrastructure.

10

List of Applications/Versions Certified in Virtual Environments

The certification is relevant for Solaris 2.10 / 11 Zones, VMware ESX version 4.x / 5.x, and later, and Redhat KVM.

Application	Product Version	VMware Support	Solaris Zones Support	Redhat KVM (RHEV-M)
Aleph	18.01 and later	+	+	+ Aleph v20.x and later
DigiTool	3.x and later	+	+	
MetaLib	4.x and later	+	+	
Primo	2.0 and later	+	+	+ Primo v4.x and later
SFX	3.0 and later	+	+	+ SFX v4.x and later
Verde	2.0 and later	+	+	
Voyager		+ 7.2 and later	+ 6.2 and later	+ 8.x and later
Rosetta	1.0 and later	+	-	