



---

# LPI 306-300

---

**LPI LPIC-3 Certification Questions & Answers**

---

Exam Summary – Syllabus – Questions

---

**306-300**

**[LPIC-3 High Availability and Storage Clusters](#)**

**60 Questions Exam – 500 / 800 Cut Score – Duration of 90 minutes**

## Table of Contents:

Know Your 306-300 Certification Well: .....	2
LPI 306-300 LPIC-3 Certification Details: .....	2
306-300 Syllabus:.....	3
<b>High Availability Cluster Management</b> .....	3
<b>High Availability Cluster Storage</b> .....	6
<b>High Availability Distributed Storage</b> .....	10
<b>Single Node High Availability</b> .....	11
LPI 306-300 Sample Questions:.....	15
Study Guide to Crack LPI LPIC-3 306-300 Exam:.....	18

## Know Your 306-300 Certification Well:

The 306-300 is best suitable for candidates who want to gain knowledge in the LPI Linux System Administration. Before you start your 306-300 preparation you may struggle to get all the crucial LPIC-3 materials like 306-300 syllabus, sample questions, study guide.

But don't worry the 306-300 PDF is here to help you prepare in a stress free manner.

The PDF is a combination of all your queries like-

- What is in the 306-300 syllabus?
- How many questions are there in the 306-300 exam?
- Which Practice test would help me to pass the 306-300 exam at the first attempt?

Passing the 306-300 exam makes you LPIC-3 High Availability and Storage Clusters. Having the LPIC-3 certification opens multiple opportunities for you. You can grab a new job, get a higher salary or simply get recognition within your current organization.

## LPI 306-300 LPIC-3 Certification Details:

Exam Name	LPIC-3 High Availability and Storage Clusters
Exam Code	306-300
Exam Price	\$200 (USD)
Duration	90 mins
Number of Questions	60
Passing Score	500 / 800
Schedule Exam	<a href="#">LPI Marketplace</a>
Sample Questions	<a href="#">LPI LPIC-3 Sample Questions</a>
Practice Exam	<a href="#">LPI 306-300 Certification Practice Exam</a>

## 306-300 Syllabus:

Topic	Details
<b>High Availability Cluster Management</b>	
High Availability Concepts and Theory	<p><b>Weight:</b> 6</p> <p><b>Description:</b> Candidates should understand the properties and design approaches of high availability clusters.</p> <p><b>Key Knowledge Areas:</b></p> <ul style="list-style-type: none"> <li>• Understand the goals of High Availability and Site Reliability Engineering</li> <li>• Understand common cluster architectures</li> <li>• Understand recovery and cluster reorganization mechanisms</li> <li>• Design an appropriate cluster architecture for a given purpose</li> <li>• Understand application aspects of high availability</li> <li>• Understand operational considerations of high availability</li> </ul> <p><b>Partial list of the used files, terms and utilities:</b></p> <ul style="list-style-type: none"> <li>• Active/Passive Cluster</li> <li>• Active/Active Cluster</li> <li>• Failover Cluster</li> <li>• Load Balanced Cluster</li> <li>• Shared-Nothing Cluster</li> <li>• Shared-Disk Cluster</li> <li>• Cluster resources</li> <li>• Cluster services</li> <li>• Quorum</li> <li>• Fencing (Node and Resource Level Fencing)</li> <li>• Split brain</li> </ul>

Topic	Details
	<ul style="list-style-type: none"> <li>• Redundancy</li> <li>• Mean Time Before Failure (MTBF)</li> <li>• Mean Time To Repair (MTTR)</li> <li>• Service Level Agreement (SLA)</li> <li>• Disaster Recovery</li> <li>• State Handling</li> <li>• Replication</li> <li>• Session handling</li> </ul>
Load Balanced Clusters	<p><b>Weight: 8</b></p> <p><b>Description:</b> Candidates should know how to install, configure, maintain and troubleshoot LVS. This includes the configuration and use of keepalived and ldirectord. Candidates should further be able to install, configure, maintain and troubleshoot HAProxy.</p> <p><b>Key Knowledge Areas:</b></p> <ul style="list-style-type: none"> <li>• Understand the concepts of LVS / IPVS</li> <li>• Understand the basics of VRRP</li> <li>• Configure keepalived</li> <li>• Configure ldirectord</li> <li>• Configure backend server networking</li> <li>• Understand HAProxy</li> <li>• Configure HAProxy</li> </ul> <p><b>Partial list of the used files, terms and utilities:</b></p> <ul style="list-style-type: none"> <li>• ipvsadm</li> <li>• syncd</li> <li>• LVS Forwarding (NAT, Direct Routing, Tunneling, Local Node)</li> <li>• connection scheduling algorithms</li> <li>• keepalived configuration file</li> <li>• ldirectord configuration file</li> <li>• genhash</li> </ul>

Topic	Details
	<ul style="list-style-type: none"> <li>• HAProxy configuration file</li> <li>• load balancing algorithms</li> <li>• ACLs</li> </ul>
Failover Clusters	<p><b>Weight: 8</b></p> <p><b>Description:</b> Candidates should have experience in the installation, configuration, maintenance and troubleshooting of a Pacemaker cluster. This includes the use of Corosync. The focus is on Pacemaker 2.x for Corosync 2.x.</p> <p><b>Key Knowledge Areas:</b></p> <ul style="list-style-type: none"> <li>• Understand the architecture and components of Pacemaker (CIB, CRMd, PEngine, LRMd, DC, STONITHd)</li> <li>• Manage Pacemaker cluster configurations</li> <li>• Understand Pacemaker resource classes (OCF, LSB, Systemd, Service, STONITH, Nagios)</li> <li>• Manage Pacemaker resources</li> <li>• Manage resource rules and constraints (location, order, colocation).</li> <li>• Manage advanced resource features (templates, groups, clone resources, multi-state resources)</li> <li>• Obtain node information and manage node health</li> <li>• Manage quorum and fencing in a Pacemaker cluster</li> <li>• Configure the Split Brain Detector on shared storage</li> <li>• Manage Pacemaker using pcs</li> <li>• Manage Pacemaker using crmsh</li> <li>• Configure and management of corosync in conjunction with Pacemaker</li> <li>• Awareness of Pacemaker ACLs</li> </ul>

Topic	Details
	<ul style="list-style-type: none"> <li>• Awareness of other cluster engines (OpenAIS, Heartbeat, CMAN)</li> </ul> <p><b>Partial list of the used files, terms and utilities:</b></p> <ul style="list-style-type: none"> <li>• pcs</li> <li>• crm</li> <li>• crm_mon</li> <li>• crm_verify</li> <li>• crm_simulate</li> <li>• crm_shadow</li> <li>• crm_resource</li> <li>• crm_attribute</li> <li>• crm_node</li> <li>• crm_standby</li> <li>• cibadmin</li> <li>• corosync.conf</li> <li>• authkey</li> <li>• corosync-cfgtool</li> <li>• corosync-cmapctl</li> <li>• corosync-quorumtool</li> <li>• stonith_admin</li> <li>• stonith</li> <li>• ocf:pacemaker:ping</li> <li>• ocf:pacemaker:NodeUtilization</li> <li>• ocf:pacemaker:ocf:SysInfo</li> <li>• ocf:pacemaker:HealthCPU</li> <li>• ocf:pacemaker:HealthSMART</li> <li>• sbd</li> </ul>
<p><b>High Availability Cluster Storage</b></p>	
<p>DRBD</p>	<p><b>Weight:</b> 6</p> <p><b>Description:</b> Candidates are expected to have the experience and knowledge to install, configure,</p>

Topic	Details
	<p>maintain and troubleshoot DRBD devices. This includes integration with Pacemaker. DRBD configuration of version 9.0.x is covered.</p> <p><b>Key Knowledge Areas:</b></p> <ul style="list-style-type: none"> <li>• Understand the DRBD architecture</li> <li>• Understand DRBD resources, states and replication modes</li> <li>• Configure DRBD disks and devices</li> <li>• Configure DRBD networking connections and meshes</li> <li>• Configure DRBD automatic recovery and error handling</li> <li>• Configure DRBD quorum and handlers for split brain and fencing</li> <li>• Manage DRBD using drbdadm</li> <li>• Understand the principles of drbdsetup and drbdmeta</li> <li>• Restore and verify the integrity of a DRBD device after an outage</li> <li>• Integrate DRBD with Pacemaker</li> <li>• Understand the architecture and features of LINSTOR</li> </ul> <p><b>Partial list of the used files, terms and utilities:</b></p> <ul style="list-style-type: none"> <li>• Protocol A, B and C</li> <li>• Primary, Secondary</li> <li>• Three-way replication</li> <li>• drbd kernel module</li> <li>• drbdadm</li> <li>• drbdmon</li> <li>• drbdsetup</li> <li>• drbdmeta</li> <li>• /etc/drbd.conf</li> </ul>

Topic	Details
	<ul style="list-style-type: none"> <li>• /etc/drbd.d/</li> <li>• /proc/drbd</li> </ul>
Cluster Storage Access	<p><b>Weight: 3</b></p> <p><b>Description:</b> Candidates should be able to connect a Linux node to remote block storage. This includes understanding common SAN technology and architectures, including management of iSCSI, as well as configuring multipathing for high availability and using LVM on a clustered storage.</p> <p><b>Key Knowledge Areas:</b></p> <ul style="list-style-type: none"> <li>• Understand the concepts of Storage Area Networks</li> <li>• Understand the concepts of Fibre Channel, including Fibre Channel Topologies</li> <li>• Understand and manage iSCSI targets and initiators</li> <li>• Understand and configure Device Mapper Multipath I/O (DM-MPIO)</li> <li>• Understand the concept of a Distributed Lock Manager (DLM)</li> <li>• Understand and manage clustered LVM</li> <li>• Manage DLM and LVM with Pacemaker</li> </ul> <p><b>Partial list of the used files, terms and utilities:</b></p> <ul style="list-style-type: none"> <li>• tgtadm</li> <li>• targets.conf</li> <li>• iscsiadm</li> <li>• iscsid.conf</li> <li>• /etc/multipath.conf</li> <li>• multipath</li> <li>• kpartx</li> <li>• pvmove</li> <li>• vgchange</li> </ul>

Topic	Details
	<ul style="list-style-type: none"> <li>• lvchange</li> </ul>
Clustered File Systems	<p><b>Weight: 4</b></p> <p><b>Description:</b> Candidates should be able to install, maintain and troubleshoot GFS2 and OCFS2 filesystems. This includes awareness of other clustered filesystems available on Linux.</p> <p><b>Key Knowledge Areas:</b></p> <ul style="list-style-type: none"> <li>• Understand the principles of cluster file systems and distributed file systems</li> <li>• Understand the Distributed Lock Manager</li> <li>• Create, maintain and troubleshoot GFS2 file systems in a cluster</li> <li>• Create, maintain and troubleshoot OCFS2 file systems in a cluster</li> <li>• Awareness of the O2CB cluster stack</li> <li>• Awareness of other commonly used clustered file systems, such as AFS and Lustre</li> </ul> <p><b>Partial list of the used files, terms and utilities:</b></p> <ul style="list-style-type: none"> <li>• mkfs.gfs2</li> <li>• mount.gfs2</li> <li>• fsck.gfs2</li> <li>• gfs2_grow</li> <li>• gfs2_edit</li> <li>• gfs2_jadd</li> <li>• mkfs.ocfs2</li> <li>• mount.ocfs2</li> <li>• fsck.ocfs2</li> <li>• tuneufs.ocfs2</li> <li>• mounted.ocfs2</li> <li>• o2info</li> <li>• o2image</li> </ul>

Topic	Details
<b>High Availability Distributed Storage</b>	
<ul style="list-style-type: none"> <li>• GlusterFS Storage Clusters</li> </ul>	<ul style="list-style-type: none"> <li>• <b>Weight:</b> 5</li> <li>• <b>Description:</b> Candidates should be able to manage and maintain a GlusterFS storage cluster.</li> <li>• <b>Key Knowledge Areas:</b> <ul style="list-style-type: none"> <li>• Understand the architecture and components of GlusterFS</li> <li>• Manage GlusterFS peers, trusted storage pools, bricks and volumes</li> <li>• Mount and use an existing GlusterFS</li> <li>• Configure high availability aspects of GlusterFS</li> <li>• Scale up a GlusterFS cluster</li> <li>• Replace failed bricks</li> <li>• Recover GlusterFS from a physical media failure</li> <li>• Restore and verify the integrity of a GlusterFS cluster after an outage</li> <li>• Awareness of GNFS</li> </ul> </li> <li>• <b>Partial list of the used files, terms and utilities:</b> <ul style="list-style-type: none"> <li>• gluster (including relevant subcommands)</li> </ul> </li> </ul>
<ul style="list-style-type: none"> <li>• Ceph Storage Clusters</li> </ul>	<ul style="list-style-type: none"> <li>• <b>Weight:</b> 8</li> <li>• <b>Description:</b> Candidates should be able to manage and maintain a Ceph Cluster. This includes the configuration of RGW, RDB devices and CephFS.</li> <li>• <b>Key Knowledge Areas:</b> <ul style="list-style-type: none"> <li>• Understand the architecture and components of Ceph</li> <li>• Manage OSD, MGR, MON and MDS</li> <li>• Understand and manage placement groups and pools</li> </ul> </li> </ul>

Topic	Details
	<ul style="list-style-type: none"> <li>• Understand storage backends (FileStore and BlueStore)</li> <li>• Initialize a Ceph cluster</li> <li>• Create and manage Rados Block Devices</li> <li>• Create and manage CephFS volumes, including snapshots</li> <li>• Mount and use an existing CephFS</li> <li>• Understand and adjust CRUSH maps</li> <li>• Configure high availability aspects of Ceph</li> <li>• Scale up a Ceph cluster</li> <li>• Restore and verify the integrity of a Ceph cluster after an outage</li> <li>• Understand key concepts of Ceph updates, including update order, tunables and features</li> <li>• <b>Partial list of the used files, terms and utilities:</b> <ul style="list-style-type: none"> <li>• ceph-deploy (including relevant subcommands)</li> <li>• ceph.conf</li> <li>• ceph (including relevant subcommands)</li> <li>• rados (including relevant subcommands)</li> <li>• rdb (including relevant subcommands)</li> <li>• cephfs (including relevant subcommands)</li> <li>• ceph-volume (including relevant subcommands)</li> <li>• ceph-authtool</li> <li>• ceph-bluestore-tool</li> <li>• crushtool</li> </ul> </li> </ul>
<p><b>Single Node High Availability</b></p>	
<p>Hardware and Resource High Availability</p>	<p><b>Weight:</b> 2</p> <p><b>Description:</b> Candidates should be able to monitor a local node for potential hardware failures and resource shortages.</p>

Topic	Details
	<p><b>Key Knowledge Areas:</b></p> <ul style="list-style-type: none"> <li>• Understand and monitor S.M.A.R.T values using smartmontools, including triggering frequent disk checks</li> <li>• Configure system shutdown at specific UPC events</li> <li>• Configure monit for alerts in case of resource exhaustion</li> </ul> <p><b>Partial list of the used files, terms and utilities:</b></p> <ul style="list-style-type: none"> <li>• smartctl</li> <li>• /etc/smartd.conf</li> <li>• smartd</li> <li>• nvme-cli</li> <li>• apcupsd</li> <li>• apctest</li> <li>• monit</li> </ul>
Advanced RAID	<p><b>Weight: 2</b></p> <p><b>Description:</b> Candidates should be able to manage software raid devices on Linux. This includes advanced features such as partitionable RAIDs and RAID containers as well as recovering RAID arrays after a failure.</p> <p><b>Key Knowledge Areas:</b></p> <ul style="list-style-type: none"> <li>• Manage RAID devices using various raid levels, including hot spare discs, partitionable RAIDs and RAID containers</li> <li>• Add and remove devices from an existing RAID</li> <li>• Change the RAID level of an existing device</li> <li>• Recover a RAID device after a failure</li> <li>• Understand various metadata formats and RAID geometries</li> </ul>

Topic	Details
	<ul style="list-style-type: none"> <li>• Understand availability and performance properties of various raid levels</li> <li>• Configure mdadm monitoring and reporting</li> </ul> <p><b>Partial list of the used files, terms and utilities:</b></p> <ul style="list-style-type: none"> <li>• mdadm</li> <li>• /proc/mdstat</li> <li>• /proc/sys/dev/raid/*</li> </ul>
Advanced LVM	<p><b>Weight: 3</b></p> <p><b>Description:</b> Candidates should be able to configure LVM volumes. This includes managing LVM snapshot, pools and RAIDs.</p> <p><b>Key Knowledge Areas:</b></p> <ul style="list-style-type: none"> <li>• Understand and manage LVM, including linear and striped volumes</li> <li>• Extend, grow, shrink and move LVM volumes</li> <li>• Understand and manage LVM snapshots</li> <li>• Understand and manage LVM thin and thick pools</li> <li>• Understand and manage LVM RAIDs</li> </ul> <p><b>Partial list of the used files, terms and utilities:</b></p> <ul style="list-style-type: none"> <li>• /etc/lvm/lvm.conf</li> <li>• pvcreate</li> <li>• pvdisplay</li> <li>• pvmove</li> <li>• pvremove</li> <li>• pvresize</li> <li>• vgcreate</li> <li>• vgdisplay</li> <li>• vgreduce</li> <li>• lvconvert</li> </ul>

Topic	Details
	<ul style="list-style-type: none"> <li>• lvcreate</li> <li>• lvdisplay</li> <li>• lvextend</li> <li>• lvreduce</li> <li>• lvresize</li> </ul>
Network High Availability	<p><b>Weight: 5</b></p> <p><b>Description:</b> Candidates should be able to configure redundant networking connections and manage VLANs. Furthermore, candidates should have a basic understanding of BGP.</p> <p><b>Key Knowledge Areas:</b></p> <ul style="list-style-type: none"> <li>• Understand and configure bonding network interface</li> <li>• Network bond modes and algorithms (active-backup, blance-tlb, balance-alb, 802.3ad, balance-rr, balance-xor, broadcast)</li> <li>• Configure switch configuration for high availability, including RSTP</li> <li>• Configure VLANs on regular and bonded network interfaces</li> <li>• Persist bonding and VLAN configuration</li> <li>• Understand the principle of autonomous systems and BGP to manage external redundant uplinks</li> <li>• Awareness of traffic shaping and control capabilities of Linux</li> </ul> <p><b>The following is a partial list of the used files, terms and utilities:</b></p> <ul style="list-style-type: none"> <li>• bonding.ko (including relevant module options)</li> <li>• /etc/network/interfaces</li> <li>• /etc/sysconfig/networking-scripts/ifcfg-*</li> <li>• /etc/systemd/network/*.network</li> </ul>

Topic	Details
	<ul style="list-style-type: none"> <li>• /etc/systemd/network/*.netdev</li> <li>• nmcli</li> <li>• /sys/class/net/bonding_masters</li> <li>• /sys/class/net/bond*/bonding/miimon</li> <li>• /sys/class/net/bond*/bonding/slaves</li> <li>• ifenslave</li> <li>• ip</li> </ul>

## LPI 306-300 Sample Questions:

### Question: 1

What is the usual way to gain command line access to computing instances in an IaaS cloud?

- a) By providing a public SSH key to the cloud management system and using the matching private SSH key to log into the instances.
- b) Usually computing instances are accessed through a web frontend and do not allow command line access.
- c) By either telnet or SSH using the credentials of the cloud management system account with administrative privileges.
- d) By using a telnet session with the credentials set and publicized by the creator of the computing instance's system image.
- e) By using a VNC console which does not require authentication when it is invoked via the cloud management system.

**Answer: a**

### Question: 2

When used with KVM, which of the following block device image formats support snapshots?

- a) qcow
- b) dmg
- c) qcow2
- d) qed
- e) raw

**Answer: c**

**Question: 3**

Which of the following statements are true regarding IaaS computing instances?

(Choose TWO correct answers.)

- a) Creation of new instances can be automated by using scripts or APIs and service interfaces.
- b) The root file system of a computing instance is always persistent and can be accessed after the instance is destroyed.
- c) Each and every user of an IaaS cloud has exactly one computing instance.
- d) Once created, computing instances are seldom deleted in order to ensure the accessibility of the instance's data.
- e) Instances may be created when needed and destroyed when they become obsolete.

**Answer: a, e**

**Question: 4**

When preparing a public machine image for the provisioning of new cloud computing instances, which of the following steps should be done?

(Choose TWO correct answers.)

- a) Remove all default users and groups, including root and nobody.
- b) Remove all configuration files from /etc/ that were not manually modified.
- c) Remove all confidential data from the image.
- d) Remove all remote login services (i.e. SSH) from the image.
- e) Remove all private SSH keys from the image.

**Answer: c, e**

**Question: 5**

In case one service in a Pacemaker cluster depends on another service running on the same node, which score should be assigned to the constraint describing that dependency?

- a) -1
- b) 0
- c) 1
- d) FORCE
- e) INFINITY

**Answer: e**

**Question: 6**

How can the state of the backing storage disks of a DRBD device be checked?

(Choose TWO correct answers.)

- a) `cat /proc/drdb/status`
- b) `cat /proc/drdb`
- c) `drbdadm dstate all`
- d) `drbdadm diskstatus all`
- e) `drbdadm state --disk all`

**Answer: b, c**

**Question: 7**

Which of the following technologies are parts of the Red Hat Enterprise Linux High Availability Add-On?

(Choose THREE correct answers.)

- a) DRBD
- b) LVS
- c) cLVM
- d) Pacemaker
- e) OCFS2

**Answer: b, c, d**

**Question: 8**

Using Linux Virtual Server (LVS) without any additional tools provides which of the following features?

- a) Balancing network connections across several backend servers.
- b) Automatically taking over the services of failed nodes.
- c) Restarting failed services on the backend servers.
- d) Checking the integrity of services on backend servers.
- e) Starting and stopping backend servers as needed.

**Answer: a**

**Question: 9**

Within the graphical output of a KVM virtual machine, which key sequence switches to the KVM monitor of the VM?

- a) Ctrl-Alt-1
- b) Ctrl-Alt-0
- c) Ctrl-Alt-4
- d) Ctrl-Alt-2
- e) Ctrl-Alt-3

**Answer: d**

**Question: 10**

How does the command to add a new resource to a Pacemaker cluster using crm begin?

- a) crm manage cluster
- b) crm configure primitive
- c) crm service add
- d) crm configure resource
- e) crm resource new

**Answer: b**

## Study Guide to Crack LPI LPIC-3 306-300 Exam:

- Getting details of the 306-300 syllabus, is the first step of a study plan. This pdf is going to be of ultimate help. Completion of the syllabus is must to pass the 306-300 exam.
- Making a schedule is vital. A structured method of preparation leads to success. A candidate must plan his schedule and follow it rigorously to attain success.
- Joining the LPI provided training for 306-300 exam could be of much help. If there is specific training for the exam, you can discover it from the link above.
- Read from the 306-300 sample questions to gain your idea about the actual exam questions. In this PDF useful sample questions are provided to make your exam preparation easy.
- Practicing on 306-300 practice tests is must. Continuous practice will make you an expert in all syllabus areas.

## Reliable Online Practice Test for 306-300 Certification

Make EduSum.com your best friend during your LPI High Availability and Storage Clusters exam preparation. We provide authentic practice tests for the 306-300 exam. Experts design these online practice tests, so we can offer you an exclusive experience of taking the actual 306-300 exam. We guarantee you 100% success in your first exam attempt if you continue practicing regularly. Don't bother if you don't get 100% marks in initial practice exam attempts. Just utilize the result section to know your strengths and weaknesses and prepare according to that until you get 100% with our practice tests. Our evaluation makes you confident, and you can score high in the 306-300 exam.

**Start Online Practice of 306-300 Exam by visiting URL**

**<https://www.edusum.com/lpi/306-300-lpi-high-availability-and-storage-clusters>**