سرفصل آموزشی

پک حرفه ای مهندسی لینوکس

فهرست سرفصل‌های دوره‌های آموزشی

سرفصل‌های دوره آموزشی

LPIC3-300

Topic 390: OpenLDAP

Configuration 390.1 OpenLDAP Replication

Description: Candidates should be familiar with the server replication available with OpenLDAP.

Weight: 3

Key Knowledge Areas:

- Replication concepts
- Configure OpenLDAP replication
- Analyze replication log files
- Understand replica hubs
- LDAP referrals
- LDAP sync replication

The following is a partial list of the used files, terms and utilities:

- master / slave server
- multi-master replication
- consumer
- replica hub
- one-shot mode
- referral
- syncrepl
- pull-based / push-based synchronization
- refreshOnly and refreshAndPersist
- replog

390.2 Securing the Directory

Description: Candidates should be able to configure encrypted access to the LDAP directory, and restrict access at the firewall level.

Weight: 3

Key Knowledge Areas:

- Securing the directory with SSL and TLS
- Firewall considerations
- Unauthenticated access methods
- User / password authentication methods
- Maintanence of SASL user DB
- Client / server certificates

Terms and Utilities:

- SSL / TLS
- Security Strength Factors (SSF)
- SASL
- proxy authorization
- StartTLS
- iptables

390.3 OpenLDAP Server Performance Tuning

Weight: 2

Description: Candidates should be capable of measuring the performance of an LDAP server, and tuning configuration directives.

Key Knowledge Areas:

- Measure OpenLDAP performance
- Tune software configuration to increase performance
- Understand indexes

Terms and Utilities:

- index
- DB_CONFIG

Topic 391: OpenLDAP as an Authentication Backend

391.1 LDAP Integration with PAM and NSS

Weight: 2

Description: Candidates should be able to configure PAM and NSS to retrieve information from an LDAP directory.

Key Knowledge Areas:

- Configure PAM to use LDAP for authentication
- Configure NSS to retrieve information from LDAP
- Configure PAM modules in various Unix environments

Terms and Utilities:

- PAM
- NSS
- /etc/pam.d/
- /etc/nsswitch.conf

391.2 Integrating LDAP with Active Directory and Kerberos

Weight: 2

Description: Candidates should be able to integrate LDAP with Active Directory Services.

Key Knowledge Areas:

- Kerberos integration with LDAP
- Cross platform authentication
- Single sign-on concepts
- Integration and compatibility limitations between OpenLDAP and Active Directory

Terms and Utilities:

- Kerberos
- Active Directory
- single sign-on
- DNS

Topic 392: Samba Basics style=”text-align: left;”>392.1 Samba Concepts and Architecture

Weight: 2

Description: Candidates should understand the essential concepts of Samba. As well, the major differences between Samba3 and Samba4 should be known.

Key Knowledge Areas:

- Understand the roles of the Samba daemons and components
- Understand key issues regarding heterogeneous networks
- Identify key TCP/UDP ports used with SMB/CIFS
- Knowledge of Samba3 and Samba4 differences

The following is a partial list of the used files, terms and utilities:

- /etc/services
- Samba daemons: smbd, nmbd, samba, winbindd

392.2 Configure Samba

Weight: 4

Description: Candidates should be able to configure the Samba daemons for a wide variety of purposes.

Key Knowledge Areas:

- Knowledge of Samba server configuration file structure
- Knowledge of Samba variables and configuration parameters
- Troubleshoot and debug configuration problems with Samba

Terms and Utilities:

- smb.conf
- smb.conf parameters
- smb.conf variables
- testparm
- secrets.tdb

392.3 Regular Samba Maintenance

Weight: 2

Description: Candidates should know about the various tools and utilities that are part of a Samba installation.

Key Knowledge Areas:

- Monitor and interact with running Samba daemons
- Perform regular backups of Samba configuration and state data

Terms and Utilities:

- smbcontrol
- smbstatus
- tdbbackup

392.4 Troubleshooting Samba

Weight: 2

Description: Candidates should understand the structure of trivial database files and know how troubleshoot problems.

Key Knowledge Areas:

- Configure Samba logging
- Backup TDB files
- Restore TDB files
- Identify TDB file corruption
- Edit / list TDB file content

Terms and Utilities:

- /var/log/samba/
- log level
- debuglevel
- smbpasswd
- pdbedit
- secrets.tdb
- tdbbackup
- tdbdump
- tdbrestore
- tdbtool

392.5 Internationalization

Weight: 1

Description: Candidates should be able to work with internationalization character codes and code pages.

Key Knowledge Areas:

- Understand internationalization character codes and code pages
- Understand the difference in the name space between Windows and Linux/Unix with respect to share, file and directory names in a non-English environment
- Understand the difference in the name space between Windows and Linux/Unix with respect to user and group naming in a non-English environment
- Understand the difference in the name space between Windows and Linux/Unix with respect to computer naming in a non-English environment

Terms and Utilities:

- internationalization
- character codes
- code pages
- smb.conf
- dos charset, display charset and unix charset

Topic 393: Samba Share Configuration

393.1 File Services

Weight: 4

Description: Candidates should be able to create and configure file shares in a mixed environment.

Key Knowledge Areas:

- Create and configure file sharing
- Plan file service migration
- Limit access to IPC$
- Create scripts for user and group handling of file shares
- Samba share access configuration parameters

Terms and Utilities:

- smb.conf
- [homes]
- smbcquotas
- smbsh
- browseable, writeable, valid users, write list, read list, read only and guest ok
- IPC$
- mount, smbmount

393.2 Linux File System and Share/Service Permissions

Weight: 3

Description: Candidates should understand file permissions on a Linux file system in a mixed environment.

Key Knowledge Areas:

- Knowledge of file / directory permission control
- Understand how Samba interacts with Linux file system permissions and ACLs
- Use Samba VFS to store Windows ACLs

Terms and Utilities:

- smb.conf
- chmod, chown
- create mask, directory mask, force create mode, force directory mode
- smbcacls
- getfacl, setfacl
- vfs_acl_xattr, vfs_acl_tdb and vfs objects

393.3 Print Services

Weight: 2

Description: Candidates should be able to create and manage print shares in a mixed environment.

Key Knowledge Areas:

- Create and configure printer sharing
- Configure integration between Samba and CUPS
- Manage Windows print drivers and configure downloading of print drivers
- Configure [print$]
- Understand security concerns with printer sharing
- Uploading printer drivers for Point'n'Print driver installation using 'Add Print Driver Wizard' in Windows

Terms and Utilities:

- smb.conf
- [print$]
- CUPS
- cupsd.conf
- /var/spool/samba/.
- smbspool
- rpcclient
- net

Topic 394: Samba User and Group Management

394.1 Managing User Accounts and Groups

Weight: 4

Description: Candidates should be able to manage user and group accounts in a mixed environment.

Key Knowledge Areas:

- Manager user and group accounts
- Understand user and group mapping
- Knowledge of user account management tools
- Use of the smbpasswd program
- Force ownership of file and directory objects

Terms and Utilities:

- pdbedit
- smb.conf
- samba-tool user (with subcommands)
- samba-tool group (with subcommands)
- smbpasswd
- /etc/passwd
- /etc/group
- force user, force group.
- idmap

394.2 Authentication, Authorization and Winbind

Weight: 5

Description: Candidates should understand the various authentication mechanisms and configure access control. Candidates should be able to install and configure the Winbind service.

Key Knowledge Areas:

- Setup a local password database
- Perform password synchronization
- Knowledge of different passdb backends
- Convert between Samba passdb backends
- Integrate Samba with LDAP
- Configure Winbind service
- Configure PAM and NSS

Terms and Utilities:

- smb.conf
- smbpasswd, tdbsam, ldapsam
- passdb backend
- libnss_winbind
- libpam_winbind
- libpam_smbpass
- wbinfo
- getent
- SID and foreign SID
- /etc/passwd
- /etc/group

Topic 395: Samba Domain Integration

395.1 Samba as a PDC and BDC

Weight: 3

Description: Candidates should be able to setup and maintain primary and backup domain controllers. Candidates should be able to manage Windows/Linux client access to the NT-Style domains.

Key Knowledge Areas:

- Understand and configure domain membership and trust relationships
- Create and maintain a primary domain controller with Samba3 and Samba4
- Create and maintain a backup domain controller with Samba3 and Samba4
- Add computers to an existing domain
- Configure logon scripts
- Configure roaming profiles
- Configure system policies

Terms and Utilities:

- smb.conf
- security mode
- server role
- domain logons
- domain master
- logon script
- logon path
- NTConfig.pol
- net
- profiles
- add machine script
- profile acls

395.2 Samba4 as an AD compatible Domain Controller

Weight: 3

Description: Candidates should be able to configure Samba 4 as an AD Domain Controller.

Key Knowledge Areas:

- Configure and test Samba 4 as an AD DC
- Using smbclient to confirm AD operation
- Understand how Samba integrates with AD services: DNS, Kerberos, NTP, LDAP

Terms and Utilities:

- smb.conf
- server role
- samba-tool domain (with subcommands)
- samba

395.3 Configure Samba as a Domain Member Server

Weight: 3

Description: Candidates should be able to integrate Linux servers into an environment where Active Directory is present.

Key Knowledge Areas:

- Joining Samba to an existing NT4 domain
- Joining Samba to an existing AD domain
- Ability to obtain a TGT from a KDC

Terms and Utilities:

- smb.conf
- server role
- server security
- net command
- kinit, TGT and REALM

Topic 396: Samba Name Services

396.1 NetBIOS and WINS

Weight: 3

Description: Candidates should be familiar with NetBIOS/WINS concepts and understand network browsing.

Key Knowledge Areas:

- Understand WINS concepts
- Understand NetBIOS concepts
- Understand the role of a local master browser
- Understand the role of a domain master browser
- Understand the role of Samba as a WINS server
- Understand name resolution
- Configure Samba as a WINS server
- Configure WINS replication
- Understand NetBIOS browsing and browser elections
- Understand NETBIOS name types

Terms and Utilities:

- smb.conf
- nmblookup
- smbclient
- name resolve order
- lmhosts
- wins support, wins server, wins proxy, dns proxy
- domain master, os level, preferred master

396.2 Active Directory Name Resolution

Weight: 2

Description: Candidates should be familiar with the internal DNS server with Samba4.

Key Knowledge Areas:

- Understand and manage DNS for Samba4 as an AD Domain Controller
- DNS forwarding with the internal DNS server of Samba4

Terms and Utilities:

- samba-tool dns (with subcommands)
- smb.conf
- dns forwarder
- /etc/resolv.conf
- dig, host

Topic 397: Working with Linux and Windows Clients

397.1 CIFS Integration

Weight: 3

Description: Candidates should be comfortable working with CIFS in a mixed environment.

Key Knowledge Areas:

- Understand SMB/CIFS concepts
- Access and mount remote CIFS shares from a Linux client
- Securely storing CIFS credentials
- Understand features and benefits of CIFS
- Understand permissions and file ownership of remote CIFS shares

Terms and Utilities:

- SMB/CIFS
- mount, mount.cifs
- smbclient
- smbget
- smbtar
- smbtree
- findsmb
- smb.conf
- smbcquotas
- /etc/fstab

397.2 Working with Windows Clients

Weight: 2

Description: Candidates should be able to interact with remote Windows clients, and configure Windows workstations to access file and print services from Linux servers.

Key Knowledge Areas:

- Knowledge of Windows clients
- Explore browse lists and SMB clients from Windows
- Share file / print resources from Windows
- Use of the smbclient program
- Use of the Windows net utility

Terms and Utilities:

- Windows net command
- smbclient
- control panel
- rdesktop
- workgroup

LPIC3-305

Topic 351: Full Virtualization

351.1 Virtualization Concepts and Theory

Weight: 8

Description: Candidates should know and understand the general concepts, theory and terminology of virtualization. This includes Xen, QEMU and libvirt terminology.

Key Knowledge Areas:

- Understand virtualization terminology
- Understand the pros and cons of virtualization
- Understand the various variations of Hypervisors and Virtual Machine Monitors
- Understand the major aspects of migrating physical to virtual machines
- Understand the major aspects of migrating virtual machines between host systems
- Understand the features and implications of virtualization for a virtual machine, such as snapshotting, pausing, cloning and resource limits
- Awareness of oVirt, Proxmox, systemd-machined and VirtualBox
- Awareness of Open vSwitch

The following is a partial list of the used files, terms and utilities:

- Hypervisor
- Hardware Virtual Machine (HVM)
- Paravirtualization (PV)
- Emulation and Simulation
- CPU flags
- /proc/cpuinfo
- Migration (P2V, V2V)

351.2 Xen

Weight: 3

Description: Candidates should be able to install, configure, maintain, migrate and troubleshoot Xen installations. The focus is on Xen version 4.x.

Key Knowledge Areas:

- Understand architecture of Xen, including networking and storage
- Basic configuration of Xen nodes and domains
- Basic management of Xen nodes and domains
- Basic troubleshooting of Xen installations
- Awareness of XAPI
- Awareness of XenStore
- Awareness of Xen Boot Parameters
- Awareness of the xm utility

The following is a partial list of the used files, terms and utilities:

- in0 (Dom0), DomainU (DomU)Doma
- PV-DomU, HVM-DomU
- /etc/xen/
- xl
- xl.cfg
- xl.conf
- xentop

351.3 QEMU

Weight: 4

Description: Candidates should be able to install, configure, maintain, migrate and troubleshoot QEMU installations.

Key Knowledge Areas:

- Understand the architecture of QEMU, including KVM, networking and storage
- Start QEMU instances from the command line
- Manage snapshots using the QEMU monitor
- Install the QEMU Guest Agent and VirtIO device drivers
- Troubleshoot QEMU installations, including networking and storage
- Awareness of important QEMU configuration parameters

The following is a partial list of the used files, terms and utilities:

- Kernel modules: kvm, kvm-intel and kvm-amd
- /dev/kvm
- QEMU monitor
- qemu
- qemu-system-x86_64
- ip
- brctl
- tunctl

351.4 Libvirt Virtual Machine Management

Weight: 9

Description: Candidates should be able to manage virtualization hosts and virtual machines (‘libvirt domains’) using libvirt and related tools.

Key Knowledge Areas:

- Understand the architecture of libvirt
- Manage libvirt connections and nodes
- Create and manage QEMU and Xen domains, including snapshots
- Manage and analyze resource consumption of domains
- Create and manage storage pools and volumes
- Create and manage virtual networks
- Migrate domains between nodes
- Understand how libvirt interacts with Xen and QEMU
- Understand how libvirt interacts with network services such as dnsmasq and radvd
- Understand libvirt XML configuration files
- Awareness of virtlogd and virtlockd

The following is a partial list of the used files, terms and utilities:

- irtdlibv
- /etc/libvirt/
- virsh (including relevant subcommands)

351.5 Virtual Machine Disk Image Management

Weight: 3

Description: Candidates should be able to manage virtual machines disk images. This includes converting disk images between various formats and hypervisors and accessing data stored within an image.

Key Knowledge Areas:

- Understand features of various virtual disk image formats, such as raw images, qcow2 and VMDK
- Manage virtual machine disk images using qemu-img
- Mount partitions and access files contained in virtual machine disk images using libguestfish
- Copy physical disk content to a virtual machine disk image
- Migrate disk content between various virtual machine disk image formats
- Awareness of Open Virtualization Format (OVF)

The following is a partial list of the used files, terms and utilities:

- u-imgqem
- guestfish (including relevant subcommands)
- guestmount
- guestumount
- virt-cat
- virt-copy-in
- virt-copy-out
- virt-diff
- virt-inspector
- virt-filesystems
- virt-rescue
- virt-df
- virt-resize
- virt-sparsify
- virt-p2v
- virt-p2v-make-disk
- virt-v2v
- virt-sysprep

Topic352: Container Virtualization

352.1 Container Virtualization Concepts

Weight: 7

Description: Candidates should understand the concept of container virtualization. This includes understanding the Linux components used to implement container virtualization as well as using standard Linux tools to troubleshoot these components.

Key Knowledge Areas:

- Understand the concepts of system and application container
- Understand and analyze kernel namespaces
- Understand and analyze control groups
- Understand and analyze capabilities
- Understand the role of seccomp, SELinux and AppArmor for container virtualization
- Understand how LXC and Docker leverage namespaces, cgroups, capabilities, seccomp and MAC
- Understand the principle of runc
- Understand the principle of CRI-O and containerd
- Awareness of the OCI runtime and image specifications
- Awareness of the Kubernetes Container Runtime Interface (CRI)
- Awareness of podman, buildah and skopeo
- Awareness of other container virtualization approaches in Linux and other free operating systems, such as rkt, OpenVZ, systemd-nspawn or BSD Jails

The following is a partial list of the used files, terms and utilities:

- nsenter
- unshare
- ip (including relevant subcommands)
- capsh
- /sys/fs/cgroups
- /proc/[0-9]+/ns
- /proc/[0-9]+/status

352.2 LXC

Weight: 6

Description: Candidates should be able to use system containers using LXC and LXD. The version of LXC covered is 3.0 or higher.

Key Knowledge Areas:

- Understand the architecture of LXC and LXD
- Manage LXC containers based on existing images using LXD, including networking and storage
- Configure LXC container properties
- Limit LXC container resource usage
- Use LXD profiles
- Understand LXC images
- Awareness of traditional LXC tools

Partial list of the used files, terms and utilities:

- dlx
- lxc (including relevant subcommands)

352.3 Docker

Weight: 9

Description: Candidate should be able to manage Docker nodes and Docker containers. This include understand the architecture of Docker as well as understanding how Docker interacts with the node's Linux system.

Key Knowledge Areas:

- Understand the architecture and components of Docker
- Manage Docker containers by using images from a Docker registry
- Understand and manage images and volumes for Docker containers
- Understand and manage logging for Docker containers
- Understand and manage networking for Docker
- Use Dockerfiles to create container images
- Run a Docker registry using the registry Docker image

Partial list of the used files, terms and utilities:

- rddocke
- /etc/docker/daemon.json
- /var/lib/docker/
- docker
- Dockerfile

352.4 Container Orchestration Platforms

Weight: 3

Description: Candidates should understand the importance of container orchestration and the key concepts Docker Swarm and Kubernetes provide to implement container orchestration.

Key Knowledge Areas:

- Understand the relevance of container orchestration
- Understand the key concepts of Docker Compose and Docker Swarm
- Understand the key concepts of Kubernetes and Helm
- Awareness of OpenShift, Rancher and Mesosphere DC/OS

Topic 353: VM Deployment and Provisioning

353.1 Cloud Management Tools

Weight: 2

Description: Candidates should understand common offerings in public clouds and have basic feature knowledge of commonly available cloud management tools.

Key Knowledge Areas:

- Understand common offerings in public clouds
- Basic feature knowledge of OpenStack
- Basic feature knowledge of Terraform
- Awareness of CloudStack, Eucalyptus and OpenNebula

Partial list of the used files, terms and utilities:

- IaaS, PaaS, SaaS
- OpenStack
- Terraform

353.2 Packer

Weight: 3

Description: Candidates should able to use cloud-init to configure virtual machines created from standardized images. This includes adjusting virtual machines to match their available hardware resources, specifically, disk space and volumes. Additionally, candidates should be able to configure instances to allow secure SSH logins and install a specific set of software packages. Furthermore, candidates should be able to create new system images with cloud-init support.

Key Knowledge Areas:

- Understanding the features and concepts of cloud-init, including user-data, initializing and configuring cloud-init
- Use cloud-init to create, resize and mount file systems, configure user accounts, including login credentials such as SSH keys and install software packages from the distribution's repository
- Integrate cloud-init into system images
- Use config drive datasource for testing

Partial list of the used files, terms and utilities:

- cloud-init
- user-data
- /var/lib/cloud/

353.4 Vagrant

Weight: 3

Description: Candidate should be able to use Vagrant to manage virtual machines, including provisioning of the virtual machine.

Key Knowledge Areas:

- Understand Vagrant architecture and concepts, including storage and networking
- Retrieve and use boxes from Atlas
- Create and run Vagrantfiles
- Access Vagrant virtual machines
- Share and synchronize folder between a Vagrant virtual machine and the host system
- Understand Vagrant provisioning, i.e. File and Shell provisioners
- Understand multi-machine setup

Partial list of the used files, terms and utilities:

- vagrant
- Vagrantfile

LPIC3-306

Topic 361: High Availability Cluster Management

361.1 High Availability Concepts and Theory (weight: 6)

Description: Candidates should understand the properties and design approaches of high availability clusters.

Key Knowledge Areas:

- Understand the goals of High Availability and Site Reliability Engineering
- Understand common cluster architectures
- Understand recovery and cluster reorganization mechanisms
- Design an appropriate cluster architecture for a given purpose
- Understand application aspects of high availability
- Understand operational considerations of high availability

Partial list of the used files, terms and utilities:

- Active/Passive Cluster
- Active/Active Cluster
- Failover Cluster
- Load Balanced Cluster
- Shared-Nothing Cluster
- Shared-Disk Cluster
- Cluster resources

- Cluster services
- Quorum
- Fencing (Node and Resource Level Fencing)
- Split brain
- Redundancy
- Mean Time Before Failure (MTBF)
- Mean Time To Repair (MTTR)
- Service Level Agreement (SLA)
- Disaster Recovery
- State Handling
- Replication
- Session handling

361.2 Load Balanced Clusters (weight: 8)

Description: 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.

Key Knowledge Areas:

- Understand the concepts of LVS / IPVS
- Understand the basics of VRRP
- Configure keepalived
- Configure ldirectord
- Configure backend server networking
- Understand HAProxy
- Configure HAProxy

Partial list of the used files, terms and utilities:

- Ipvsadm
- Syncd
- LVS Forwarding (NAT, Direct Routing, Tunneling, Local Node)
- Connection scheduling algorithms
- keepalived configuration file
- ldirectord configuration file
- Genhash
- HAProxy configuration file
- load balancing algorithms
- ACLs

361.3 Failover Clusters (weight: 8)

Description: 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.

Key Knowledge Areas:

- Understand the architecture and components of Pacemaker (CIB, CRMd, PEngine, LRMd, DC, STONITHd)
- Manage Pacemaker cluster configurations
- Understand Pacemaker resource classes (OCF, LSB, Systemd, Service, STONITH, Nagios)
- Manage Pacemaker resources
- Manage resource rules and constraints (location, order, colocation).

- Manage advanced resource features (templates, groups, clone resources, multi-state resources)
- Obtain node information and manage node health
- Manage quorum and fencing in a Pacemaker cluster
- Configure the Split Brain Detector on shared storage
- Manage Pacemaker using pcs
- Manage Pacemaker using crmsh
- Configure and management of corosync in conjunction with Pacemaker
- Awareness of Pacemaker ACLs
- Awareness of other cluster engines (OpenAIS, Heartbeat, CMAN)

Partial list of the used files, terms and utilities:

- pcs
- crm
- crm_mon
- crm_verify
- crm_simulate
- crm_shadow
- crm_resource
- crm_attribute
- crm_node
- crm_standby
- cibadmin
- corosync.conf
- authkey
- corosync-cfgtool
- corosync-cmapctl
- corosync-quorumtool
- stonith_admin
- stonith
- ocf:pacemaker:ping

- ocf:pacemaker:NodeUtilization
- ocf:pacemaker:ocf:SysInfo
- ocf:pacemaker:HealthCPU
- ocf:pacemaker:HealthSMART
- sbd

Topic 362: High Availability Cluster Storage

362.1 DRBD (weight: 6)

Description: Candidates are expected to have the experience and knowledge to install, configure, maintain and troubleshoot DRBD devices. This includes integration with Pacemaker. DRBD configuration of version 9.0.x is covered.

Key Knowledge Areas:

- Understand the DRBD architecture
- Understand DRBD resources, states and replication modes
- Configure DRBD disks and devices
- Configure DRBD networking connections and meshes
- Configure DRBD automatic recovery and error handling
- Configure DRBD quorum and handlers for split brain and fencing
- Manage DRBD using drbdadm
- Understand the principles of drbdsetup and drbdmeta
- Restore and verify the integrity of a DRBD device after an outage
- Integrate DRBD with Pacemaker
- Understand the architecture and features of LINSTOR

Partial list of the used files, terms and utilities:

- Protocol A, B and C
- Primary, Secondary
- Three-way replication
- drbd kernel module
- drbdadm
- drbdmon
- drbdsetup
- drbdmeta
- /etc/drbd.conf
- /etc/drbd.d/
- /proc/drbd

362.2 Cluster Storage Access (weight: 3)

Description: 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.

Key Knowledge Areas:

- Understand the concepts of Storage Area Networks
- Understand the concepts of Fibre Channel, including Fibre Channel Toplogies
- Understand and manage iSCSI targets and initiators
- Understand and configure Device Mapper Multipath I/O (DM-MPIO)

- Understand the concept of a Distributed Lock Manager (DLM)
- Understand and manage clustered LVM
- Manage DLM and LVM with Pacemaker

Partial list of the used files, terms and utilities:

- tgtadm
- targets.conf
- iscsiadm
- iscsid.conf
- /etc/multipath.conf
- multipath
- kpartx
- pvmove
- vgchange
- lvchange

362.3 Clustered File Systems (weight: 4)

Description: Candidates should be able to install, maintain and troubleshoot GFS2 and OCFS2 filesystems. This includes awareness of other clustered filesystems available on Linux.

Key Knowledge Areas:

- Understand the principles of cluster file systems and distributed file systems
- Understand the Distributed Lock Manager
- Create, maintain and troubleshoot GFS2 file systems in a cluster
- Create, maintain and troubleshoot OCFS2 file systems in a cluster

- Awareness of the O2CB cluster stack
- Awareness of other commonly used clustered file systems, such as AFS and Lustre

Partial list of the used files, terms and utilities:

- mkfs.gfs2
- mount.gfs2
- fsck.gfs2
- gfs2_grow
- gfs2_edit
- gfs2_jadd
- mkfs.ocfs2
- mount.ocfs2
- fsck.ocfs2
- tunefs.ocfs2
- mounted.ocfs2
- o2info
- o2image

Topic 363: High Availability Distributed Storage

363.1 GlusterFS Storage Clusters (weight: 5)

Description: Candidates should be able to manage and maintain a GlusterFS storage cluster.

Key Knowledge Areas:

- Understand the architecture and components of GlusterFS
- Manage GlusterFS peers, trusted storge pools, bricks and volumes
- Mount and use an existing GlusterFS
- Configure high availability aspects of GlusterFS
- Scale up a GlusterFS cluster
- Replace failed bricks
- Recover GlusterFS from a physical media failure
- Restore and verify the integrity of a GlusterFS cluster after an outage
- Awareness of GNFS

Partial list of the used files, terms and utilities:

- gluster (including relevant subcommands)

363.2 Ceph Storage Clusters (weight: 8)

Description: Candidates should be able to manage and maintain a Ceph Cluster. This includes the configuration of RGW, RDB devices and CephFS.

Key Knowledge Areas:

- Understand the architecture and components of Ceph
- Manage OSD, MGR, MON and MDS
- Understand and manage placement groups and pools
- Understand storage backends (FileStore and BlueStore)
- Initialize a Ceph cluster
- Create and manage Rados Block Devices
- Create and manage CephFS volumes, including snapshots
- Mount and use an existing CephFS
- Understand and adjust CRUSH maps
- Configure high availability aspects of Ceph
- Scale up a Ceph cluster
- Restore and verify the integrity of a Ceph cluster after an outage
- Understand key concepts of Ceph updates, including update order, tunables and features

Partial list of the used files, terms and utilities:

- ceph-deploy (including relevant subcommands)
- ceph.conf
- ceph (including relevant subcommands)
- rados (including relevant subcommands)
- rdb (including relevant subcommands)
- cephfs (including relevant subcommands)
- ceph-volume (including relevant subcommands)
- ceph-authtool
- ceph-bluestore-tool
- crushtool

Topic 364: Single Node High Availability

364.1 Hardware and Resource High Availability (weight: 2)

Description: Candidates should be able to monitor a local node for potential hardware failures and resource shortages.

Key Knowledge Areas:

- Understand and monitor S.M.A.R.T values using smartmontools, including triggering frequent disk checks
- Configure system shutdown at specific UPC events
- Configure monit for alerts in case of resource exhaustion

Partial list of the used files, terms and utilities:

- smartctl
- /etc/smartd.conf
- smartd
- nvme-cli
- apcupsd
- apctest
- monit

364.2 Advanced RAID (weight: 2)

Description: Candidates should be able to manage software raid devices on Linux. This includes advanced features such as partitonable RAIDs and RAID containers as well as recovering RAID arrays after a failure.

Key Knowledge Areas:

Manage RAID devices using various raid levels, including hot spare discs, partitionable RAIDs and RAID containers
Add and remove devices from an existing RAID
Change the RAID level of an existing device
Recover a RAID device after a failure
Understand various metadata formats and RAID geometries
Understand availability and performance properties of various raid levels
Configure mdadm monitoring and reporting

Partial list of the used files, terms and utilities:

- mdadm
- /proc/mdstat
- /proc/sys/dev/raid/*

364.3 Advanced LVM (weight: 3)

Description: Candidates should be able to configure LVM volumes. This includes managing LVM snapshot, pools and RAIDs.

Key Knowledge Areas:

- Understand and manage LVM, including linear and striped volumes
- Extend, grow, shrink and move LVM volumes
- Understand and manage LVM snapshots
- Understand and manage LVM thin and thick pools
- Understand and manage LVM RAIDs

Partial list of the used files, terms and utilities:

- /etc/lvm/lvm.conf
- pvcreate
- pvdisplay
- pvmove
- pvremove
- pvresize
- vgcreate
- vgdisplay
- vgreduce
- lvconvert
- lvcreate
- lvdisplay
- lvextend
- lvreduce
- lvresize

364.4 Network High Availability (weight: 5)

Description: Candidates should be able to configure redundant networking connections and manage VLANs. Furthermore, candidates should have a basic understanding of BGP.

Key Knowledge Areas:

- Understand and configure bonding network interface
- Network bond modes and algorithms (active-backup, blance-tlb, balance-alb, 802.3ad, balance-rr, balance-xor, broadcast)
- Configure switch configuration for high availability, including RSTP
- Configure VLANs on regular and bonded network interfaces
- Persist bonding and VLAN configuration
- Understand the principle of autonomous systems and BGP to manage external redundant uplinks
- Awareness of traffic shaping and control capabilities of Linux

Partial list of the used files, terms and utilities:

- bonding.ko (including relevant module options)
- /etc/network/interfaces
- /etc/sysconfig/networking-scripts/ifcfg-*
- /etc/systemd/network/*.network
- /etc/systemd/network/*.netdev
- nmcli
- /sys/class/net/bonding_masters
- /sys/class/net/bond*/bonding/miimon
- /sys/class/net/bond*/bonding/slaves
- ifenslave
- ip