

VCAP-VCF ADMINISTRATOR LAB TASKS · EXAM 3V0-11.26

Work task 1 to task 262 and you have built, operated and upgraded a VCF 9 fleet

Exam 3V0-11.26, Advanced VMware Cloud Foundation 9 Administrator. All sixty testable objectives sit in section 4 of the blueprint, numbered 4.1 to 4.60. Tasks only, in build order. The code at the right of each task is the objective it covers.

Dr Pranay Jha · drpranayjha.com

A. READ THE FLEET BEFORE YOU TOUCH IT

Written tasks. One or two pages in total, done before anything is changed.

1. Draw your fleet on one page: VCF Operations, VCF Automation, VCF Fleet Management, the Identity Broker, and every workload domain with its clusters. Mark which are management and which are workload. 4.35
2. For each workload domain, record the storage it uses and whether that storage is vSAN, VMFS on FC, NFS, or vVols. 4.3
3. State which VCF components are upgraded as a fleet and which are upgraded per domain. 4.45
4. List the identity sources in use and record which component each one authenticates: Provider portal, organizations, vCenter, NSX. 4.46
5. For one cluster, write down the current storage policy applied to its virtual machines and what that policy actually guarantees. 4.9
6. State the difference between an Organization and a Project in VCF Automation, in one sentence each. 4.22
7. State what a Namespace Class controls that a Project does not. 4.28
8. Write the difference between Business Intent and Operational Intent in VCF Operations, and give one example of each. 4.41
9. List which of your clusters could survive the loss of a site today, and what makes the rest unable to. 4.2
10. Record your current capacity headroom per cluster from VCF Operations, and the date it projects you run out. 4.38

B. WORKLOAD DOMAINS AND VSAN CLUSTERS

From here on everything is performed. Build the storage layer first, because everything else lands on it.

11. Commission a set of hosts into the fleet and verify they appear as unassigned, ready for a domain. 4.1
12. Deploy a new VI workload domain with a vSAN cluster as its first cluster. 4.1
13. Record the vSAN mode used (ESA or OSA) and state what decided it. 4.1
14. Add a second cluster to the same workload domain, again on vSAN. 4.1
15. Deploy a workload domain cluster backed by VMFS on Fibre Channel instead of vSAN. 4.3
16. Deploy a workload domain cluster backed by NFS. 4.3
17. For the non-vSAN cluster, record what the platform does not manage for you that it would have managed on vSAN. 4.3
18. Add a host to an existing vSAN cluster and confirm its disks are claimed as expected. 4.1
19. Remove a host from a vSAN cluster cleanly, using the correct data evacuation mode. 4.7
20. State the three data evacuation options on host maintenance mode and when each one is the right choice. 4.7
21. Expand a vSAN cluster to the point where the default storage policy can no longer be satisfied, and record the error you get. 4.9

C. VSAN STORAGE POLICIES

Storage policy is the control surface for everything above it. Work through this section before you build any workload.

- | | | |
|-----|--|-----|
| 22. | Create a vSAN storage policy with RAID-1 mirroring and failures to tolerate of 1. | 4.9 |
| 23. | Create a second policy with RAID-5 erasure coding and record the minimum host count it needs. | 4.9 |
| 24. | Create a RAID-6 policy and record its minimum host count. | 4.9 |
| 25. | Compare the usable capacity reported for the same virtual machine under RAID-1, RAID-5 and RAID-6. | 4.9 |
| 26. | Apply a new policy to a running virtual machine and watch the resynchronisation that follows. | 4.9 |
| 27. | Create a policy that pins a workload to a specific storage tier or cluster using tags. | 4.9 |
| 28. | Set a policy with a stripe width greater than 1 and state what it changes about object placement. | 4.9 |
| 29. | Create a policy that reserves object space, and state the effect on reported free capacity. | 4.9 |
| 30. | Change the default storage policy of a datastore and record which virtual machines are affected and which are not. | 4.9 |
| 31. | Deliberately apply a policy a cluster cannot satisfy, read the compliance status, then correct it. | 4.9 |

D. VSAN DAY 2 ADMINISTRATION

Operate the cluster you just built. Create each condition yourself before you fix it.

- | | | |
|-----|---|-----|
| 32. | Put a host into maintenance mode with full data migration and record how long evacuation takes. | 4.7 |
| 33. | Read the vSAN Skyline Health checks and clear every warning you can. | 4.7 |
| 34. | Run a vSAN proactive test and record what it proves and what it does not. | 4.7 |
| 35. | Find the resynchronising objects view and state the three reasons an object resynchronises. | 4.7 |
| 36. | Locate a virtual machine that is non-compliant with its policy and explain why from the object layout. | 4.7 |
| 37. | Enable and then read vSAN performance service statistics for a cluster, a host and a virtual machine. | 4.7 |
| 38. | Simulate a disk failure and watch the rebuild, recording the delay before rebuild begins and where that timer is set. | 4.7 |
| 39. | Adjust the object repair timer and state why you would lengthen it. | 4.7 |
| 40. | Read the capacity view and account for the difference between used, usable and reserved. | 4.7 |
| 41. | Turn on and then interpret vSAN capacity reservations, both operations reserve and host rebuild reserve. | 4.7 |
| 42. | Upgrade the on-disk format of a vSAN cluster and record what cannot be undone. | 4.7 |

E. VSAN STRETCHED CLUSTERS

A stretched cluster is a different product from a standard cluster in day 2 terms. Build one and then operate it.

- | | | |
|-----|--|-----|
| 43. | Deploy a vSAN witness appliance and record the sizing you chose. | 4.2 |
| 44. | Configure a vSAN stretched cluster across two fault domains with that witness. | 4.2 |
| 45. | Record the bandwidth and latency requirements for the inter-site link and for the witness link. | 4.2 |
| 46. | Create a storage policy with site disaster tolerance and secondary failures to tolerate, and apply it. | 4.2 |
| 47. | State what happens to a virtual machine whose policy has site disaster tolerance set to none. | 4.2 |
| 48. | Configure host and virtual machine affinity so a workload runs at a preferred site. | 4.8 |
| 49. | Fail the witness and record what still works and what does not. | 4.8 |
| 50. | Fail the preferred site and record the recovery behaviour of the virtual machines. | 4.8 |
| 51. | Restore the failed site and watch the resynchronisation back. | 4.8 |
| 52. | Change the preferred fault domain and record what moves. | 4.8 |
| 53. | Replace a witness appliance on a running stretched cluster. | 4.8 |
| 54. | Read the stretched cluster health checks and explain every one that is specific to stretching. | 4.8 |

F. SHARED, DISAGGREGATED AND PROTECTED VSAN

Capacity sharing, storage clusters, encryption and data protection, in that order.

- | | | |
|-----|--|-----|
| 55. | Configure vSAN cross-cluster capacity sharing between two clusters in the same domain. | 4.4 |
| 56. | Mount a remote vSAN datastore on a client cluster and run a workload against it. | 4.4 |
| 57. | Record the limits on how many client clusters a server cluster may serve. | 4.4 |
| 58. | Build a vSAN storage cluster that provides capacity to compute-only clusters. | 4.4 |
| 59. | State what a compute cluster consuming a vSAN storage cluster still needs locally. | 4.4 |
| 60. | Enable vSAN data-at-rest encryption on a cluster with a key provider already configured. | 4.5 |
| 61. | Configure a native key provider and record where the backup of its key must be kept. | 4.5 |
| 62. | Enable vSAN data-in-transit encryption and record its performance cost. | 4.5 |
| 63. | Perform a shallow rekey and then a deep rekey, and state the difference. | 4.5 |
| 64. | Enable vSAN data protection and create a protection group with a snapshot schedule. | 4.6 |
| 65. | Restore a virtual machine from a vSAN data protection snapshot. | 4.6 |
| 66. | Clone a virtual machine from a protection group snapshot into a new virtual machine. | 4.6 |
| 67. | State how vSAN data protection differs from a backup product, in one sentence. | 4.6 |

G. VSPHERE SETTINGS AND VIRTUAL MACHINES

The day-to-day surface of a workload domain.

- | | | |
|-----|---|------|
| 68. | Configure an advanced setting on a cluster through the platform rather than directly on a host, and record why that matters in VCF. | 4.10 |
| 69. | Configure DRS automation level and a migration threshold on a cluster, and state the effect of each step of the threshold. | 4.10 |

70.	Configure vSphere HA admission control and record how much capacity it reserves.	4.10
71.	Set a host isolation response and explain the condition that triggers it.	4.10
72.	Create a virtual machine override that exempts one workload from a cluster setting.	4.11
73.	Deploy a virtual machine from a content library template into a workload domain.	4.11
74.	Hot-add CPU and memory to a running virtual machine and record what the guest actually sees.	4.11
75.	Configure a virtual machine with a reservation, a limit and shares, and state what each one does under contention.	4.11
76.	Create a VM-VM anti-affinity rule and a VM-host affinity rule, and state which one DRS may violate.	4.11
77.	Migrate a virtual machine between clusters with vMotion, then between datastores with Storage vMotion.	4.11
78.	Take a snapshot, grow the virtual machine, then consolidate and record the disk impact.	4.11
79.	Configure a virtual machine for high latency sensitivity and record what the platform reserves as a result.	4.13

H. CPU, MEMORY AND VCENTER PERFORMANCE

Optimisation tasks. Measure first, change one thing, measure again.

80.	Find CPU ready time for a virtual machine and state the threshold at which it becomes a problem.	4.13
81.	Create CPU contention deliberately by overcommitting a cluster, then measure co-stop on a wide virtual machine.	4.13
82.	Right-size an over-provisioned virtual machine down and measure the change in ready time.	4.13
83.	Configure CPU affinity on a virtual machine, record what it breaks, then remove it.	4.13
84.	Explain NUMA placement for a virtual machine wider than one NUMA node, and show where to read its NUMA home node.	4.13
85.	Configure a virtual machine to be NUMA-aware with the correct cores-per-socket setting.	4.13
86.	Find memory ballooning, compression and swapping on a host and state the order in which they occur.	4.14
87.	Create memory pressure on a cluster and watch the reclamation techniques engage in order.	4.14
88.	Set a memory reservation on a critical virtual machine and record the effect on its swap file.	4.14
89.	Read the active, consumed and granted memory counters for a virtual machine and explain the difference.	4.14
90.	Check the vCenter appliance resource usage and state its deployment size and what that size limits.	4.15
91.	Increase the vCenter appliance size and record the steps and the outage involved.	4.15
92.	Review vCenter database health and the appliance file system usage, and clear what can be cleared.	4.15
93.	Configure vCenter appliance monitoring and state the three things most likely to degrade it.	4.15

I. MONITORING VSPHERE COMPONENTS

Know where each number comes from before you trust it.

- | | | |
|------|--|------|
| 94. | Build a custom performance chart for a cluster showing CPU ready, memory ballooning and disk latency together. | 4.16 |
| 95. | Set the statistics collection level and record the cost of raising it. | 4.16 |
| 96. | Create a vCenter alarm on datastore usage and trigger it deliberately. | 4.16 |
| 97. | Configure an alarm action that sends mail or runs a script, and verify it fired. | 4.16 |
| 98. | Read esxcli on a host and identify CPU ready, memory state and disk latency fields. | 4.16 |
| 99. | Find the storage latency figures for a datastore and separate device latency from kernel latency. | 4.16 |
| 100. | Review the vCenter task and event logs for one change and reconstruct who did what. | 4.16 |
| 101. | Forward ESX host logs to VCF Operations for Logs and prove a log line arrives. | 4.16 |
| 102. | Build a log query that finds every host that entered maintenance mode in the last day. | 4.16 |
| 103. | Export a performance report for a cluster over a week and state what you would show a capacity owner. | 4.16 |

J. SECURITY, ACCESS CONTROL AND ENCRYPTION

Identity and encryption inside the workload domain.

- | | | |
|------|--|------|
| 104. | Create a custom vCenter role with the minimum privileges for an operator who may only power virtual machines on and off. | 4.17 |
| 105. | Assign that role at a folder and prove the permission does not leak to a sibling folder. | 4.17 |
| 106. | Configure a global permission and state how it differs from an inventory permission. | 4.17 |
| 107. | Enable lockdown mode on a host and record what still has access. | 4.17 |
| 108. | Configure the ESX host firewall for one service and verify from a client. | 4.17 |
| 109. | Review and then rotate ESX host certificates. | 4.17 |
| 110. | Configure a key provider for VM encryption and record its trust relationship with vCenter. | 4.18 |
| 111. | Encrypt a virtual machine and record what becomes impossible afterwards. | 4.18 |
| 112. | Decrypt a virtual machine and confirm the encrypted disks are gone. | 4.18 |
| 113. | Enable virtual TPM on a virtual machine and state what it requires of the key provider. | 4.18 |
| 114. | Configure UEFI secure boot on a virtual machine and on a host. | 4.18 |
| 115. | Configure a vSphere Trust Authority cluster and state the role of the attestation service. | 4.18 |
| 116. | Attest one workload host against the trust authority and read the result. | 4.18 |

K. ADVANCED NETWORK OPERATIONS IN A DOMAIN

The vSphere networking surface, before NSX is involved.

- | | | |
|------|---|------|
| 117. | Create a distributed port group with a specific VLAN and apply it to a virtual machine. | 4.12 |
| 118. | Configure teaming and failover on an uplink group and test it by failing one uplink. | 4.12 |
| 119. | Enable and read a port mirroring session on a distributed switch. | 4.12 |
| 120. | Enable NetFlow on a distributed switch and send it to a collector. | 4.12 |
| 121. | Configure traffic shaping on a port group and measure the effect. | 4.12 |
| 122. | Configure Network I/O Control shares for vSAN, vMotion and virtual machine traffic. | 4.12 |

- | | | |
|------|---|------|
| 123. | Enable LLDP or CDP and read the physical switch port a host uplink lands on. | 4.12 |
| 124. | Configure a VMkernel adapter for vMotion on its own TCP/IP stack. | 4.12 |
| 125. | Raise the MTU end to end for one traffic type and verify it with a ping that forbids fragmentation. | 4.12 |
| 126. | Export a distributed switch configuration, change it, then restore the export. | 4.12 |

L. NSX FABRIC

Edge clusters, gateways, segments and VRFs. Build it in this order.

- | | | |
|------|---|------|
| 127. | Record the process to deploy an NSX Edge cluster into a workload domain, including the form factor and the requirements it imposes. | 4.50 |
| 128. | Deploy an NSX Edge cluster and verify both transport node tunnels come up. | 4.50 |
| 129. | State what an Edge cluster is required for that a host transport node cannot do. | 4.50 |
| 130. | Create a Tier-0 gateway in active-active mode and record what that rules out. | 4.51 |
| 131. | Create a Tier-0 gateway in active-standby mode and configure a stateful service on it. | 4.51 |
| 132. | Configure BGP on a Tier-0 gateway to a pair of upstream routers and verify the sessions. | 4.51 |
| 133. | Configure route redistribution on the Tier-0 and prove a segment prefix reaches the physical network. | 4.51 |
| 134. | Create a VRF on the Tier-0 and record what it inherits from the parent gateway and what it does not. | 4.52 |
| 135. | Attach two tenants to two separate VRFs and prove their routes stay separate. | 4.52 |
| 136. | Create a Tier-1 gateway, attach it to the Tier-0, and enable route advertisement. | 4.54 |
| 137. | Create an overlay logical segment on the Tier-1 and attach a virtual machine to it. | 4.53 |
| 138. | Create a VLAN-backed segment and state when you would use one instead of overlay. | 4.53 |
| 139. | Move a Tier-1 gateway between Edge clusters and record the traffic impact. | 4.54 |
| 140. | Record the process steps for deploying NSX Federation across two sites, naming the Global Manager and what it owns. | 4.48 |
| 141. | State which objects are stretched by Federation and which stay local. | 4.48 |

M. NSX SERVICES, VPCS AND TENANCY

Everything the fabric carries.

- | | | |
|------|---|------|
| 142. | Configure DHCP on a segment using a local DHCP server on the Tier-1. | 4.56 |
| 143. | Configure a DHCP relay to an external server and verify a lease. | 4.56 |
| 144. | Configure source NAT on a Tier-1 so a workload leaves with a predictable address. | 4.57 |
| 145. | Configure destination NAT to publish one internal service, and state when you would use NAT instead of a load balancer. | 4.57 |
| 146. | Create an NSX Project and assign it to a tenant. | 4.58 |
| 147. | Create objects inside the Project and prove a tenant administrator cannot see another Project. | 4.58 |
| 148. | Create an NSX VPC inside a Project. | 4.55 |
| 149. | Create a VPC connectivity profile and attach it, recording what it controls. | 4.55 |
| 150. | Create a private subnet and a public subnet in the VPC and record the address source of each. | 4.55 |
| 151. | Expose a workload in the VPC externally and trace every hop it takes. | 4.55 |

- | | | |
|-------------|---|------|
| 152. | Configure a distributed firewall rule inside the VPC and prove it is enforced at the virtual machine. | 4.55 |
| 153. | Configure a gateway firewall rule on the Tier-0 and state where it is enforced instead. | 4.49 |
| 154. | Create a security group based on tags and use it in a firewall rule. | 4.49 |
| 155. | Configure an IP discovery and a segment security profile, and state what each prevents. | 4.49 |

N. NSX INTEGRATIONS AND OPERATIONS

Day 2 for the network.

- | | | |
|-------------|--|------|
| 156. | Configure syslog export from NSX Manager to VCF Operations for Logs. | 4.60 |
| 157. | Configure an NSX Manager backup to an SFTP target and run one on demand. | 4.60 |
| 158. | Restore an NSX Manager from backup in a lab and record the restore order. | 4.60 |
| 159. | Take an NSX support bundle and record what it contains. | 4.60 |
| 160. | Check NSX Manager cluster health and the status of each transport node. | 4.60 |
| 161. | Replace or renew an NSX certificate and record what must be re-registered afterwards. | 4.60 |
| 162. | Integrate NSX with VCF Operations for Networks and confirm topology appears. | 4.59 |
| 163. | Run a path trace between two virtual machines in VCF Operations for Networks and read every hop. | 4.59 |
| 164. | Integrate NSX with an external identity source and log in as a mapped user. | 4.59 |
| 165. | Use Traceflow between two virtual machines and identify the rule that drops a packet. | 4.59 |

O. PROVIDER MANAGEMENT PORTAL

VCF Automation from the provider side. Build this before any organization exists.

- | | | |
|-------------|---|------|
| 166. | Configure the identity provider for the Provider Management portal and log in as a federated user. | 4.19 |
| 167. | Record what breaks if the identity provider is unavailable, and what still works. | 4.19 |
| 168. | Create a provider content library and publish an item into it. | 4.20 |
| 169. | Subscribe an organization to the provider content library and prove the item appears. | 4.20 |
| 170. | Add a virtual machine image and a Kubernetes release to the provider content library. | 4.20 |
| 171. | Configure access control in the Provider portal: assign a provider administrator and a read-only operator. | 4.21 |
| 172. | Create a custom provider role and prove the boundary of one privilege you excluded. | 4.21 |
| 173. | Create a new Organization and record every decision the creation wizard asks for. | 4.22 |
| 174. | Create a second Organization of a different type and state the difference. | 4.22 |
| 175. | Enable the Provider Consumption Organization and record what it is for. | 4.23 |
| 176. | Configure the Provider Consumption Organization to deploy one workload, and state why this is different from a tenant organization. | 4.23 |
| 177. | Assign infrastructure to an organization and prove it cannot see another organization. | 4.22 |

P. ORGANIZATIONS, PROJECTS AND POLICY

VCF Automation from the tenant side.

- | | | |
|------|--|------|
| 178. | Configure an identity provider inside a VCF All Apps organization. | 4.24 |
| 179. | Map a directory group to an organization role and log in as a member. | 4.24 |
| 180. | Configure access control inside the organization: an organization owner, a project member, a viewer. | 4.25 |
| 181. | Create a custom organization role and test exactly what it can and cannot do. | 4.25 |
| 182. | Create a Project and assign users, a region and a quota to it. | 4.26 |
| 183. | Set a project quota, exhaust it, and record the error a user receives. | 4.26 |
| 184. | Create a second Project and prove isolation between the two. | 4.26 |
| 185. | Create a governance policy that restricts which machine sizes may be deployed. | 4.27 |
| 186. | Create a lease policy that expires a deployment, and watch it expire. | 4.27 |
| 187. | Create an approval policy and approve a request through it. | 4.27 |
| 188. | Create a Namespace Class that defines the limits for a class of namespace. | 4.28 |
| 189. | Deploy a namespace from that Namespace Class and verify the limits it inherited. | 4.28 |
| 190. | Create a VPC from the organization and record which NSX objects were created for it. | 4.29 |
| 191. | Create a VPC connectivity profile at organization level and attach two VPCs to it. | 4.29 |
| 192. | Publish a catalog item from a blueprint and request it as an ordinary user. | 4.26 |

Q. SUPERVISOR AND VSPHERE KUBERNETES SERVICE

The Kubernetes surface of this exam. Four of these objectives are shared with the VKS exam.

- | | | |
|------|---|------|
| 193. | Create a Supervisor cluster using NSX VPC networking with Avi. | 4.30 |
| 194. | Create a Supervisor cluster using NSX segment networking. | 4.30 |
| 195. | Create a Supervisor cluster using vDS networking with the Foundation Load Balancer. | 4.30 |
| 196. | Record which load balancer options are valid for each of the three networking choices. | 4.30 |
| 197. | Create a multi-zone Supervisor across three vSphere Zones and state what it protects against. | 4.30 |
| 198. | Install the Harbor Supervisor Service and verify it serves images. | 4.31 |
| 199. | Install the external-dns Supervisor Service and prove a record is created. | 4.31 |
| 200. | Uninstall a Supervisor Service cleanly and record what is left behind. | 4.31 |
| 201. | Upgrade a Supervisor Service to a newer version. | 4.34 |
| 202. | Upgrade the Supervisor itself and record the order in which components move. | 4.34 |
| 203. | Provision a VKS cluster and record the objects the Supervisor created for it. | 4.32 |
| 204. | Perform a rolling update of a VKS cluster to a newer Kubernetes release. | 4.32 |
| 205. | Change a VKS cluster configuration, such as a node pool size or VM class, and watch the rollout. | 4.32 |
| 206. | Deliberately start a rolling update that cannot complete, read why, and recover. | 4.32 |
| 207. | Add a package repository to a VKS cluster and install a standard package from it. | 4.33 |
| 208. | Create a registry secret and pull an image from a private registry into a VKS cluster. | 4.33 |
| 209. | Point a VKS cluster at a private registry with a self-signed certificate and make the pull succeed. | 4.33 |

R. VCF OPERATIONS DAY 2

Operations is a large slice of this exam. Work through it with a live fleet.

- | | | |
|------|---|------|
| 210. | Complete a full day 2 pass in VCF Operations: review alerts, capacity, compliance and cost for one cluster. | 4.35 |
| 211. | Add or re-register a vCenter adapter in VCF Operations and verify collection. | 4.35 |
| 212. | Find the reclaimable resources view and list idle, powered-off, orphaned and snapshot waste. | 4.36 |
| 213. | Reclaim one idle virtual machine and record the capacity returned. | 4.36 |
| 214. | Delete an orphaned disk found by reclamation and verify the datastore reflects it. | 4.36 |
| 215. | Find the rightsizing recommendations and read how the figure was derived. | 4.37 |
| 216. | Rightsize one oversized virtual machine and verify the recommendation clears. | 4.37 |
| 217. | State what rightsizing will not do for a workload with a reservation. | 4.37 |
| 218. | Run a what-if scenario adding 50 virtual machines of a given profile and read the date capacity runs out. | 4.38 |
| 219. | Run a what-if scenario removing a host, and one adding a cluster. | 4.38 |
| 220. | Run a workload placement what-if that compares on-premises with a cloud target. | 4.38 |
| 221. | Configure cost drivers in VCF Operations with your own hardware and licence figures. | 4.39 |
| 222. | Produce a cost per virtual machine figure and explain each component of it. | 4.39 |
| 223. | Configure a chargeback or showback report for one project. | 4.39 |
| 224. | Create a custom policy with your own thresholds and apply it to one cluster only. | 4.40 |
| 225. | Change the policy a group of objects inherits and prove the change took effect. | 4.40 |
| 226. | Clone the default policy, modify one symptom definition, and record what it changes. | 4.40 |

S. INTENTS, APPLICATIONS AND ALERTS

The parts of Operations that act rather than report.

- | | | |
|------|--|------|
| 227. | Configure an Operational Intent on a cluster and state what it is allowed to change. | 4.41 |
| 228. | Configure a Business Intent and state the difference from the Operational Intent you just set. | 4.41 |
| 229. | Set a consolidation intent and watch the recommended moves. | 4.41 |
| 230. | Create a Business Application grouping several virtual machines across clusters. | 4.42 |
| 231. | Read the health, capacity and cost of that Business Application as one object. | 4.42 |
| 232. | Add a tier to the Business Application and record what changes in its dashboards. | 4.42 |
| 233. | Create an alert definition from a symptom you define yourself. | 4.43 |
| 234. | Create a symptom based on a metric threshold, and one based on a property. | 4.43 |
| 235. | Configure an alert notification to email or a webhook and trigger it deliberately. | 4.43 |
| 236. | Suppress an alert for a maintenance window and verify it stays quiet. | 4.43 |
| 237. | Cancel or close an alert and record the difference. | 4.43 |

T. FLEET LIFECYCLE, IDENTITY AND DRIFT

The last build section. These are the tasks that touch the whole fleet at once.

- | | | |
|------|---|------|
| 238. | Record the process steps for scaling VCF Fleet Management, including what is added and in what order. | 4.44 |
| 239. | Check the current fleet inventory and verify every component reports in. | 4.44 |

240.	Download a bundle for an upgrade and verify its depot connectivity and signature.	4.45
241.	Run an upgrade prechecks pass and resolve every failure before proceeding.	4.45
242.	Upgrade one component of the fleet and record the order the platform enforced.	4.45
243.	State the correct upgrade order for the fleet, from the first component to the last.	4.45
244.	Upgrade a workload domain after the management domain and record what changed in the sequence.	4.45
245.	Configure VCF SSO through Identity and Access Management and log in with a federated account.	4.46
246.	Map an external group to a fleet-level role and verify the permissions it grants.	4.46
247.	Record what still authenticates locally after SSO is configured, and why that matters in an outage.	4.46
248.	Create a configuration drift baseline in VCF Operations.	4.47
249.	Change a setting deliberately, then find it reported as drift.	4.47
250.	Remediate the drift and verify the baseline reports compliant again.	4.47
251.	Export a drift report for one domain and state what you would do with it.	4.47

U. AGAINST THE CLOCK

Do this section twice in the last week, from a clean start, timing yourself.

252.	From an unassigned set of hosts, deploy a workload domain with a vSAN cluster. 30 minutes.	4.1
253.	Create a RAID-5 policy and apply it to a running virtual machine. 10 minutes.	4.9
254.	Deploy an Edge cluster, a Tier-0, a Tier-1 and a segment, and get a virtual machine on the network. 40 minutes.	4.50
255.	Create a Project in NSX, a VPC inside it, and publish one workload. 25 minutes.	4.55
256.	Create an Organization, a Project and a Namespace Class, and deploy a namespace. 30 minutes.	4.26
257.	Create a Supervisor with NSX VPC and Avi, then provision a VKS cluster. 45 minutes.	4.30
258.	Install Harbor as a Supervisor Service and push one image. 20 minutes.	4.31
259.	Run a what-if scenario and produce a cost per virtual machine figure. 15 minutes.	4.38
260.	Create a drift baseline, break it, find it and fix it. 20 minutes.	4.47
261.	Run upgrade prechecks across the fleet and resolve every failure. 30 minutes.	4.45
262.	Write the whole fleet build order from memory, from commissioning hosts to a running tenant workload. 15 minutes.	4.44