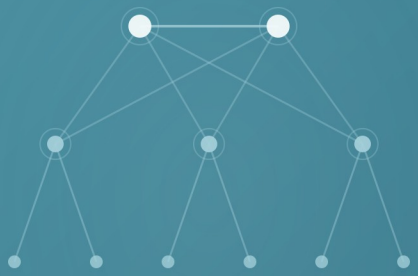

ENCOR

350-401

Complete Learning Guide

EXAM VERSION v1.2

Independent · Unofficial Study Guide

JOZEF BAROŠ

About This Sample

This is a free sample of **ENCOR 350-401 v1.2 — Complete Learning Guide**, a 511-page independent study guide for the Cisco 350-401 ENCOR examination.

What follows is taken unedited from the book: the contents, the three opening sections, a complete chapter sub-section with its configurations and knowledge check, a second sub-section from the heaviest domain, and a page of the command quick reference. Nothing here has been shortened or rewritten for the sample.

What the full book contains

	In the full book	In this sample
Pages	511	40
Chapters	6, one per exam domain	extracts from 2
Blueprint sub-sections	49	2 complete
Knowledge-check questions	245, all with full explanations	10
Diagrams	47	3
Appendices	Blueprint changes, command reference, lab guide, exam-day strategy, glossary	1 page of the command reference

What makes it different

- **Complete configurations, not fragments.** Every technology is shown as a working configuration with the `show` output that proves it, and an explanation of what each non-obvious line defends against.
- **Failure modes, not just happy paths.** Each sub-section names the ways the technology breaks and the exact output that identifies each one.

- **Written for v1.2.** All wireless content is gone, Domain 6 is Automation and Artificial Intelligence, multicast covers SSM, bidirectional PIM, and MSDP, and the product names are current. An appendix sets out every change from v1.1.
- **245 questions with real explanations** that say why the wrong answers are wrong, which is where most of the learning is.

The full book

Available at **leanpub.com/encor**

Sections in this sample: Contents · About This Book · The Exam and How This Book Maps to It · A Ten-Week Study Plan · 1.1 Enterprise Campus Design · 3.15 First Hop Redundancy Protocols · Appendix B Command Quick Reference.

Contents

About This Book

This book covers the whole of the Cisco 350-401 ENCOR examination at version 1.2 of the blueprint, in blueprint order, at the depth the exam actually asks for — and, where a topic is worth more than the exam gives it, at the depth the job asks for.

It was written on one premise: **a study guide earns its place by showing you working configurations and the output that proves they work.** Anyone can restate a bullet from the blueprint. The difficult and useful part is the complete configuration, the `show` command that confirms it, the one line in that output that tells you the truth, and the explanation of what each non-obvious command defends against. That is what most of these pages contain.

How a chapter is built

Each of the six chapters covers one exam domain and opens with a short orientation: what the domain is worth, what the blueprint verbs demand, and what changed in v1.2. Every chapter is then divided into numbered sub-sections that map one-to-one onto blueprint topics, so you can always tell which line of Cisco's document you are reading about.

Within a sub-section the shape is consistent — the concept explained from first principles, a diagram where the relationships are spatial, the configuration in full, the verification output annotated line by line, then the failure modes and how to recognise each one from the output alone.

The recurring elements

- 💡 **Did you know?** — background that makes a design decision make sense. Why a protocol was built the way it was, or where a default came from. Rarely examined directly, frequently the thing that makes an answer obvious.
- ✅ **Exam Tip** — a fact, distinction, or number that certification questions are built on. If you are short of time, these are the highest-density pages in the book.
- ⚠️ **Common Pitfall** — a mistake that is made repeatedly in production and is used deliberately as a distractor. Each one names the symptom you would actually see.
- 🌍 **From the Field** — how the topic behaves in a real network, including the cases where the textbook answer and the operational answer differ.

- **Summary** — a single dense paragraph at the end of each sub-section, written to be re-read the week before the exam.
- **Key Takeaways** — the same material as bullets, for the morning of the exam.
- **Knowledge Check** — five questions per sub-section, 245 in the book. Every question carries a full explanation that also says why the wrong answers are wrong, because that is where most of the learning is.

Conventions

- Commands, keywords, interface names, filenames, and output appear in `a monospaced font`. Type them exactly as printed.
- In configuration listings, `!` introduces a comment written for you — it is valid IOS syntax and harmless if pasted, but it is not part of the configuration's function.
- Prompts show the device and mode — `R1(config-if)#` — so you always know where a command belongs. Output examples begin with the prompt and the command that produced them.
- Values you must replace with your own are named for what they are: `10.0.100.1`, `Gi1/0/12`, `SITE-A`. Addresses are drawn from the documentation ranges reserved by RFC 5737 and RFC 3849 wherever the example allows it.
- British spelling is used throughout, except in command syntax, where Cisco's spelling is authoritative: you will read *authorisation* in prose and type `aaa authorization` at the CLI.

How to use it

If you have time, read the chapters in order and build every configuration in a laboratory as you meet it. Domain 3 is 30% of the exam and roughly a third of this book; do not treat it as one sitting.

If you are short of time, read the chapter openings, then the Exam Tips, the Common Pitfalls, and the Summaries — then take every Knowledge Check cold and read the explanation for each question you missed *and* for each you guessed. A question you got right by luck is a gap you have not found yet.

If you are revising, the Key Takeaways and Appendix B are built for the last week. Appendix D covers the exam itself.

The Exam, and How This Book Maps to It

350-401 ENCOR — Implementing and Operating Cisco Enterprise Network Core Technologies is the core examination of the CCNP Enterprise track and the qualifying examination for the CCIE Enterprise Infrastructure and CCIE Enterprise Wireless laboratory examinations. Passing it alone gives you a Cisco Certified Specialist — Enterprise Core designation; combined with one concentration examination it gives you CCNP Enterprise.

Format

The examination runs for **120 minutes**. Cisco does not publish the number of items, the passing score, or the per-domain scoring rules, and you should be suspicious of any source that states them as fact — candidates typically report somewhere between ninety and a hundred and ten items. What Cisco does publish, and what genuinely matters, is the blueprint: the list of topics and the verb attached to each one.

Item types include multiple choice with a single answer, multiple choice with several answers where the number required is stated, drag-and-drop matching, and scenario items built on an exhibit — a topology, a configuration extract, or command output. **You cannot go back.** Once you submit an item it is closed, so read carefully the first time and do not plan on a review pass.

There is no penalty for a wrong answer, which means there is never a reason to leave an item unanswered.

Domain weights, and what they mean for your time

Domain	Weight	Chapter	Pages	Questions
1.0 Architecture	15%	Chapter 1	~75	30
2.0 Virtualization	10%	Chapter 2	~58	25
3.0 Infrastructure	30%	Chapter 3	~142	80
4.0 Network Assurance	10%	Chapter 4	~60	35
5.0 Security	20%	Chapter 5	~70	40
6.0 Automation and AI	15%	Chapter 6	~62	35

The weights are not a suggestion about how to study — they are how the examination is assembled. **Infrastructure at 30% is worth twice Security and three times Virtualization.** If your revision time is limited, it belongs disproportionately in Chapter 3, and within Chapter 3 in the routing protocols, spanning tree, and the first hop redundancy and multicast material that sits at the end.

The reverse also holds. Domain 2 at 10% is the smallest chapter in the book and, for most working engineers, the least familiar material. It is a poor use of a fortnight and an excellent use of two evenings.

The verbs, and what each one demands

Cisco writes each blueprint line with a verb, and the verb tells you what an item can ask. Reading the blueprint with this in mind is the single most efficient thing you can do before you start studying.

Blueprint verb	What an item can ask	How to prepare
Explain / Describe	Compare designs or technologies; choose the one that meets a stated requirement	Understand trade-offs and limitations, not syntax
Interpret	Read a configuration, an exhibit, or command output and state what it does	Practise reading other people's configurations, especially QoS and route maps
Configure and verify	Choose the correct command, spot the missing or wrong line, read `show` output	Build it in a laboratory until the syntax is muscle memory
Troubleshoot	Given a symptom and output, identify the cause	Learn each failure mode and the exact output that proves it
Construct	Assemble valid JSON, a Python snippet, or an EEM applet	Write them by hand, without an editor to correct you

Note what follows from this. **There is no `configure terminal` anywhere in Domain 1.0** — every line is Explain or Describe — and no amount of configuration practice will help you there. Domain 3, by contrast, is almost entirely *configure and verify* and *troubleshoot*, which cannot be learned by reading.

What changed in v1.2

Version 1.2 is not a cosmetic revision. **All wireless content was removed** — design, RF fundamentals, access point modes, antennas, roaming, WLAN troubleshooting, and wireless security — which is the largest single change the ENCOR blueprint has ever seen. Domain 6 was renamed **Automation and Artificial Intelligence**, several products were renamed, and multicast was expanded. Appendix A sets out every change and what it means for material written for v1.1.

The practical consequence is blunt: **study material that predates v1.2 will teach you a substantial body of wireless content that can no longer appear on the examination, and will omit topics that now can.** Check the version of anything you study from, including this book, against the blueprint on Cisco's certification site.

A Ten-Week Study Plan

This plan assumes roughly **eight to ten hours a week** — two weekday evenings and one longer weekend session — and that you have laboratory access of some kind. Appendix C covers how to build one. If you have more time, the plan compresses; if you have less, extend it rather than cutting the laboratory work, because the laboratory work is what makes the difference between recognising a command and knowing it.

The order is deliberate. Infrastructure starts early and runs long, because it is 30% of the examination and the material it contains is cumulative. Automation sits late because it is self-contained. The last fortnight contains no new material at all.

Three habits worth more than the schedule

- **Configure from memory, not from the page.** Read the configuration, close the book, then type it. The gap between what you thought you knew and what you could produce is the entire content of your revision list.
- **Read the explanation for every question, including the ones you got right.** The explanations say why the wrong answers are wrong, and the distractors are built from the misconceptions the examination is testing for.
- **Keep a list of every mistake you make in the laboratory.** It will be short, it will repeat, and by week nine it is the most valuable page you own.

Week	Read	Laboratory work	Checkpoint
1	Chapter 1 in full	Build a two-tier campus with routed uplinks. Configure HSRP with tracking and preempt delay; fail the uplink and watch it move.	Every Chapter 1 Knowledge Check, cold. 24/30 or read it again.
2	Chapter 2 in full	VRF-Lite with two VRFs and a shared services leak. A GRE-over-IPsec tunnel, then rebuild it as a static VTI.	25 questions. Explain the recursive routing failure aloud without notes.

Week	Read	Laboratory work	Checkpoint
3	Chapter 3, sections 3.1–3.5	Break a trunk five ways and diagnose each from output only. Build an LACP bundle; mismatch the load-balancing method deliberately.	40 questions across the five sections.
4	Chapter 3, sections 3.6–3.10	EIGRP with unequal-cost load balancing. OSPF across three areas with summarisation, a stub, and an NSSA. Convert to OSPFv3.	Draw the LSA flooding scopes from memory before you check.
5	Chapter 3, sections 3.11–3.16	eBGP between two ASes. A policy-based routing map with tracking. NAT overload. HSRP against VRRP on the same segment. PIM sparse mode with an RP.	The remaining Chapter 3 questions. This is the heaviest week — do not skip it.
6	Chapter 4 in full	Flexible NetFlow from record to exporter to monitor. ERSPAN to an analyser. An IP SLA that drives a tracked static route. A NETCONF `get-config` with `ncclient`.	35 questions. Be able to name the four NetFlow objects in build order.
7	Chapter 5 in full	AAA with a TACACS+ server and a local fallback that actually works. An extended ACL you can explain line by line. CoPP with real class maps.	40 questions. Write the AAA method list that survives a dead server.
8	Chapter 6 in full	Write a Python script that authenticates to a controller and paginates a result set. Hand-write JSON and validate it. Build an EEM applet with `event none`, then arm it.	35 questions. Construct valid JSON on paper, no editor.
9	Revision — all Summaries and Key Takeaways; Appendix B end to end	Re-lab anything you could not configure from memory in weeks 1–8. Nothing new.	Every Knowledge Check again, mixed order. Target 85% overall and no domain below 75%.

Week	Read	Laboratory work	Checkpoint
10	Appendix A, then Appendix D. Re-read only the sub-sections you are still weak on.	One full evening rebuilding your weakest three topics from a blank configuration.	Book the examination at the start of this week, not the end of it.

Chapter 1 · Domain 1.0 — Architecture

Architecture is worth **15% of the ENCOR exam** — roughly one question in seven — and it is the only domain in the blueprint that asks you to reason about a network you have not been given the configuration for. Every other domain hands you a device and asks what it does. This one hands you a requirement and asks what the network should look like.

That difference changes how you study it. There is no `configure terminal` anywhere in Domain 1.0. The verbs are **Explain** and **Interpret**, which in Cisco's question style means scenario comparisons — "a company has three buildings, two data centres, and a hard requirement for sub-second failover; which design meets it?" — and artefact reading, where you are shown a QoS policy and asked what happens to a particular packet. You cannot pass this domain by memorising commands, and you cannot fail it for forgetting one.

Domain 1.0 also changed noticeably in v1.2. Every wireless design topic — RF fundamentals, AP modes, antenna types, roaming, wireless QoS — was removed outright. Cisco SD-WAN is now **Cisco Catalyst SD-WAN**, and its components carry new names that appear in exam text. QoS narrowed from "configure and verify" to "interpret", which sounds like less work but is really a different kind of work: you must be able to read someone else's policy rather than write your own. If your study material predates March 2026, it is wrong about all three.

Sub-section	Focus	Blueprint	Verb
1.1	Campus design: 2-tier, 3-tier, collapsed core	1.1.a	Explain
1.2	Fabric and cloud design models	1.1.a	Explain
1.3	High availability: redundancy, FHRP, SSO	1.1.b	Explain
1.4	Catalyst SD-WAN: planes, benefits, limitations	1.2.a, 1.2.b	Explain
1.5	SD-Access: planes and campus interoperability	1.3.a, 1.3.b	Explain
1.6	Interpreting QoS configurations	1.4	Interpret

✅ **Exam Tip — Design questions are elimination questions**

Almost every Domain 1.0 question contains at least two options that violate something you know for certain — a design with a single point of failure when the requirement said redundant, or a technology that cannot do what the scenario needs. Strike those first, then choose between the remaining two on cost, complexity, or scale. You will rarely need to know a product datasheet number.

🌍 **From the Field — Why this domain pays outside the exam**

Architecture vocabulary is what separates the engineer who implements a decision from the engineer who makes it. Being able to say precisely why a collapsed core is correct for a two-building campus but wrong for a nine-building one — and to put a number on where the crossover is — is the conversation that gets you into the design review rather than the change window.

1.1 Enterprise Campus Design: Two-Tier, Three-Tier, and Collapsed Core

A campus network is a set of switches that connect users and devices inside one site. Left to grow organically it becomes a mesh that nobody can reason about: a broadcast storm reaches everything, a failure has no boundary, and adding a building means touching every device. Hierarchical design exists to stop that. It assigns each switch a **role**, and the role determines what the switch is allowed to do, what it connects to, and what happens when it fails.

The three roles

Cisco's hierarchical model defines three layers. Learn them as **functions**, not as boxes — a single physical switch can perform two roles, and that is exactly what a collapsed core is.

- **Access layer** — where endpoints attach. This is the only layer that touches users, phones, printers, cameras, and access points. Its jobs are port density, Power over Ethernet, the trust boundary for QoS, and the first line of Layer 2 security (port security, DHCP snooping, BPDU Guard). It is deliberately the least intelligent layer, because there is more of it than anything else and every feature you add is multiplied by hundreds of switches.
- **Distribution layer** — the aggregation and policy layer. It terminates the access layer's uplinks, and it is where the network stops being a flat Layer 2 domain and becomes routed. Route summarisation, access control, QoS re-marking, and the boundary of the spanning-tree domain all live here. Distribution is also where First Hop Redundancy Protocols run, because it is the default gateway for the user VLANs.
- **Core layer** — the high-speed backbone that connects distribution blocks to each other, to the data centre, and to the WAN edge. The core does one thing: move packets between distribution blocks as fast as possible. It should have **no access lists, no policy, no NAT, and no user ports**, because anything you add to the core is in the path of every conversation in the building.

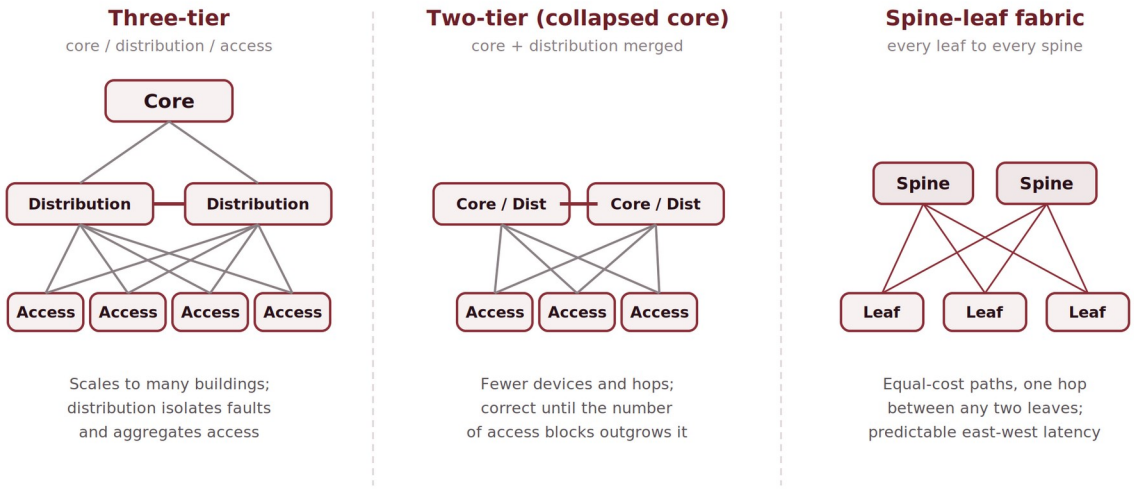

Figure 1.1 — The same campus expressed as three-tier, two-tier, and spine-leaf. The roles do not change; only how many devices perform them.

Why the layers exist: fault domains and predictability

The hierarchical model is not about tidiness. It buys two specific things.

The first is a **bounded fault domain**. When the distribution layer routes rather than switches, a broadcast storm, a spanning-tree loop, or a rogue DHCP server in one access block cannot reach another access block. The distribution switch is a wall. In a flat network there is no wall, and a single misbehaving host takes down the site — this is the failure mode that made hierarchical design mandatory rather than optional.

The second is **predictable behaviour**. In a properly built hierarchy, every packet from a user to the data centre traverses the same number of hops regardless of which access switch the user is plugged into: access → distribution → core → distribution → server. That predictability is what lets you size links, budget latency, and reason about failure without drawing the whole topology on a whiteboard.

> 💡 **Did you know? — The layers came from a hardware constraint that no longer exists**
>
> The three-tier model was formalised when a switch that could route at line rate cost more than a car, so routing was concentrated in as few devices as possible and everything else switched. Today a mid-range access switch routes at line rate too, which is why **routed access** designs — pushing the Layer 3 boundary all the way down to the access switch — are practical. The layers survived the constraint that created them because the fault-isolation argument turned out to be the stronger one.

Three-tier versus two-tier: where the line actually falls

A **collapsed core**, also called a two-tier design, merges the core and distribution functions into one pair of switches. The access layer is unchanged. This is not a compromise — for a site with a small number of access blocks it is the correct design, because a dedicated core connecting two distribution blocks is a pair of expensive switches whose only job is to forward traffic between two neighbours that could simply be cabled together.

The reason to add a real core is **cabling and change complexity, not throughput**. Without a core, every distribution block must connect to every other distribution block: with n blocks that is n(n-1)/2 links, and adding block n+1 means touching every existing distribution switch. With a core, every block connects only to the core: n links, and adding a block touches nothing that already exists. The crossover in practice is around three distribution blocks.

Consideration	Two-tier (collapsed core)	Three-tier
Distribution blocks	1–3	3 or more
Links to add a block	Touch every existing block	Two links to the core
Devices in the path	Access → dist → access	Access → dist → core → dist → access
Cost	Lower — no dedicated core pair	Higher, but flat as the site grows
Blast radius of a core failure	Larger — the device is also doing policy	Smaller — core does nothing but forward
Typical fit	Branch, single building, small campus	Multi-building campus, large site

> ⚠️ **Common Pitfall — A collapsed core is not "a core we did not buy"**
>
> The mistake is keeping a collapsed core past the point where it works and adding distribution blocks by daisy-chaining them off each other. That creates exactly the unbounded fault domain hierarchy was invented to prevent, plus a spanning-tree topology nobody can predict. If you have four distribution blocks and no core, you do not have a collapsed core design — you have a partial mesh with a nice name.

Layer 2 access versus routed access

Within the access–distribution block there are two ways to draw the Layer 3 boundary, and the exam tests the consequences of each.

In a **Layer 2 access** design, the access switch is a pure switch. VLANs are trunked up to the distribution pair, and the distribution switches host the Switch Virtual Interfaces (SVIs) that act as the default gateway. Because the same VLAN exists on two distribution switches and on multiple access switches, spanning tree runs across the block and a First Hop Redundancy Protocol is required so that the two distribution switches present one gateway address. Convergence after an uplink failure is therefore driven by STP and FHRP timers.

In a **routed access** design, the link between access and distribution is a routed point-to-point link and the SVIs live on the access switch itself. Each access switch is its own Layer 3 boundary, so no VLAN spans two switches, spanning tree has nothing to do beyond a single switch, and no FHRP is needed at the distribution layer. Convergence is driven by the routing protocol — typically well under a second with tuned EIGRP or OSPF timers.

Property	Layer 2 access	Routed access
Access–distribution link	802.1Q trunk	Routed /30 or /31
Default gateway lives on	Distribution (SVI + FHRP)	Access switch (SVI)
Spanning tree scope	Whole access block	Single switch — effectively idle
FHRP required	Yes (HSRP/VRRP/GLBP)	No
Convergence driven by	STP + FHRP timers	IGP (EIGRP/OSPF)
VLAN spanning switches	Possible	Impossible by design
Typical convergence	1–5 s untuned	< 200 ms tuned

> 🌍 **From the Field — The reason people still deploy Layer 2 access**
>
> Routed access is technically superior on almost every axis, and it still loses arguments regularly — because something in the environment needs a VLAN to span switches. Legacy clustering software, a badly written appliance that requires Layer 2 adjacency, wireless controllers in a flat design, or a building automation system are the usual culprits. Before you propose routed access, audit what actually needs Layer 2 adjacency. One badly behaved application can force the whole block back to a trunked design.

The two designs side by side, in configuration

The difference is easier to see in the configuration than in prose. Both examples build the same access block — VLAN 20 for data, VLAN 110 for voice, two access switches, two distribution switches — and both deliver a redundant default gateway. They differ only in where the Layer 3 boundary sits.

Layer 2 access — access switch trunks to distribution, gateway lives on the distribution pair

```
! ---------- ACCESS SWITCH (pure Layer 2) ----------
vlan 20
 name DATA
vlan 110
 name VOICE
!
interface range GigabitEthernet1/0/1-44
 description Endpoints
 switchport mode access
 switchport access vlan 20
 switchport voice vlan 110
 spanning-tree portfast
 spanning-tree bpduguard enable
!
interface Port-channel1
 description Uplink to DIST pair (MEC across StackWise Virtual)
 switchport mode trunk
 switchport trunk allowed vlan 20,110      ! prune everything else
!
interface range TenGigabitEthernet1/1/1 , TenGigabitEthernet2/1/1
 channel-group 1 mode active               ! LACP
!
! ---------- DISTRIBUTION (Layer 3 boundary + FHRP) ----------
interface Vlan20
```

```
 description DATA gateway
 ip address 10.20.20.2 255.255.255.0
 ip helper-address 10.100.10.10
 standby version 2
 standby 20 ip 10.20.20.1
 standby 20 priority 110
 standby 20 preempt delay minimum 90
!
interface Vlan110
 description VOICE gateway
 ip address 10.20.110.2 255.255.255.0
 ip helper-address 10.100.10.10
 standby version 2
 standby 110 ip 10.20.110.1
 standby 110 priority 110
 standby 110 preempt delay minimum 90
!
router ospf 10
 router-id 10.255.0.2
 passive-interface default
 no passive-interface TenGigabitEthernet1/1/1
 network 10.20.0.0 0.0.255.255 area 20
 network 10.255.0.2 0.0.0.0 area 0
 area 20 range 10.20.0.0 255.255.0.0   ! summarise the block into the core
```

Routed access — the SVI moves down to the access switch and the uplink becomes a routed link

```
! ---------- ACCESS SWITCH (now a Layer 3 device) ----------
ip routing
!
vlan 20
 name DATA
vlan 110
 name VOICE
!
interface Vlan20
 description DATA gateway - lives HERE now, no FHRP needed
 ip address 10.20.20.1 255.255.255.0
 ip helper-address 10.100.10.10
!
interface Vlan110
 description VOICE gateway
 ip address 10.20.110.1 255.255.255.0
 ip helper-address 10.100.10.10
```

```
!
interface TenGigabitEthernet1/1/1
 description Routed uplink to DIST-1
 no switchport
 ip address 10.255.20.1 255.255.255.254     ! /31 - two usable addresses
 ip ospf network point-to-point             ! skip DR/BDR election
 ip ospf 10 area 20
 ip ospf dead-interval minimal hello-multiplier 4
!
interface TenGigabitEthernet2/1/1
 description Routed uplink to DIST-2
 no switchport
 ip address 10.255.20.3 255.255.255.254
 ip ospf network point-to-point
 ip ospf 10 area 20
 ip ospf dead-interval minimal hello-multiplier 4
!
router ospf 10
 router-id 10.255.0.11
 passive-interface default
 no passive-interface TenGigabitEthernet1/1/1
 no passive-interface TenGigabitEthernet2/1/1
```

Three details in that second configuration carry the whole design. `no switchport` converts the uplink from a trunk to a routed port, which is what removes spanning tree from the equation. `ip ospf network point-to-point` skips the DR/BDR election, saving two hellos' worth of convergence on a link that will only ever have two routers on it. And `ip ospf dead-interval minimal hello-multiplier 4` sets a one-second dead interval with four hellos per second, which is what turns "the IGP reconverges" into sub-second recovery.

Verifying the routed access uplinks

```
ACC-1# show ip ospf interface brief
Interface  PID  Area   IP Address/Mask     Cost  State Nbrs F/C
Te1/1/1    10   20     10.255.20.1/31      1000  P2P   1/1
Te2/1/1    10   20     10.255.20.3/31      1000  P2P   1/1
Vl20       10   20     10.20.20.1/24       1     DR    0/0
Vl110      10   20     10.20.110.1/24      1     DR    0/0

ACC-1# show ip ospf neighbor
Neighbor ID     Pri   State      Dead Time   Address       Interface
10.255.0.1        0   FULL/  -   00:00:00    10.255.20.0   Te1/1/1
10.255.0.2        0   FULL/  -   00:00:00    10.255.20.2   Te2/1/1
```

```
ACC-1# show spanning-tree summary | include forwarding|blocking
Name                   Blocking Listening Learning Forwarding STP Active
2 vlans                       0         0        0          4          4
```

✅ **Exam Tip — Read the state column, not the interface name**

`show ip ospf interface brief` showing state **P2P** with a `/31` address is the signature of routed access; state **DR** or **BDR** on a VLAN interface with neighbours means a Layer 2 access block where the distribution switches are peering over the user VLAN. And a spanning-tree summary showing **zero blocking ports** across the whole switch is the fastest way to confirm that no VLAN is spanning anywhere.

Oversubscription: the number designers argue about

Oversubscription is the ratio of the aggregate bandwidth of the downstream ports to the bandwidth of the uplinks carrying that traffic. A 48-port gigabit access switch with two 10 Gbps uplinks has 48 Gbps down and 20 Gbps up — a ratio of roughly 2.4:1. That is not a fault; it is a bet that all 48 ports will never transmit at line rate simultaneously, and it is almost always a safe bet at the access layer.

- **Access to distribution** — 20:1 is the traditional guideline, and modern designs are usually far better than that. The bet is safe because user traffic is bursty.
- **Distribution to core** — 4:1 is the traditional guideline. The bet is riskier because the traffic arriving here is already aggregated, so bursts from different access blocks coincide.
- **Core** — should not be oversubscribed at all. The core carries everything; congestion there is felt by every user.

Treat these as starting points, not laws. A campus full of engineering workstations pulling large datasets behaves nothing like a campus full of laptops running a browser and a softphone. The correct method is to measure with NetFlow (Chapter 4) and size the uplinks against the 95th percentile, not against the port count.

✅ Exam Tip — What the exam actually asks about oversubscription

Not the exact ratio. The exam asks you to spot the layer where oversubscription is unacceptable (the core), to recognise that oversubscription increases as you move **down** the hierarchy, and to identify that adding uplink bandwidth at the distribution layer is the fix when distribution-to-core links are congested. Remember the direction: 20:1 at the edge is fine, 20:1 in the core is a design fault.

Summary

Hierarchical campus design assigns each switch one of three roles. Access provides port density, PoE, and the QoS trust boundary; distribution aggregates access blocks, applies policy, summarises routes, and terminates the Layer 2 domain; core forwards between distribution blocks and does nothing else. The model buys bounded fault domains and predictable hop counts. A two-tier collapsed core merges core and distribution and is correct up to roughly three distribution blocks, after which the number of inter-block links and the change complexity of adding a block justify a dedicated core. Within the block, Layer 2 access trunks VLANs to distribution SVIs and depends on STP and an FHRP for convergence, while routed access moves the Layer 3 boundary onto the access switch, eliminates both, and converges an order of magnitude faster. Oversubscription is acceptable and increasing as you move toward the edge, and unacceptable in the core.

Key Takeaways

- **Access = ports, PoE, trust boundary. Distribution = aggregation, policy, summarisation, FHRP. Core = speed only, no policy.**
- A collapsed core merges core and distribution and is correct for **one to three distribution blocks**.
- The reason to add a core is **link count and change complexity**, not raw throughput.
- **Layer 2 access** needs STP and an FHRP; **routed access** needs neither and converges via the IGP.
- Oversubscription is **highest at the access layer and should be zero in the core**.
- Nothing spans a distribution block in a good design — that boundary is the fault domain.

Knowledge Check — 1.1 Campus Design

Q1. A campus has grown from two distribution blocks to five. The existing design has no dedicated core; each distribution pair is cabled directly to every other pair. What is the strongest architectural argument for introducing a core layer?

A. The distribution switches cannot forward enough traffic without a core layer to assist them

B. Spanning tree cannot converge in a network with more than four distribution blocks

C. Adding a sixth block currently requires new links on every existing distribution switch, whereas with a core it requires two links to the core only

D. A core layer is required before an FHRP can be deployed at the distribution layer

Correct answer: C. The scaling problem in a full mesh of distribution blocks is link count and change complexity: n blocks need n(n-1)/2 links, and adding one touches every existing switch. A core reduces that to n links and makes adding a block a local change. Throughput is not the driver — distribution switches forward fine. STP has no four-block limit, and an FHRP runs between the two switches of a distribution pair regardless of whether a core exists.

Q2. Which statement correctly describes the difference in convergence behaviour between a Layer 2 access design and a routed access design?

A. Layer 2 access convergence depends on spanning-tree and FHRP timers, while routed access convergence depends on the interior gateway protocol

B. Layer 2 access converges faster because switching is performed in hardware and routing is performed in software

C. Both converge identically because both ultimately rely on the distribution switches detecting link loss

D. Routed access convergence depends on the FHRP, while Layer 2 access convergence depends on the IGP

Correct answer: A. In Layer 2 access the VLAN spans switches, so recovery waits on STP to unblock a port and on the FHRP to move the virtual gateway. In routed access the uplink is a routed link, so the IGP reconverges — typically far faster, and with no FHRP involved at all. Option D reverses the two. Both routing and switching are hardware operations on modern platforms, so option B's premise is wrong, and option C ignores that the recovery mechanisms are entirely different.

Q3. An access switch has 48 gigabit access ports and two 10 Gbps uplinks to the distribution layer. Which statement about this design is correct?

A. The design is invalid because oversubscription is never acceptable in a hierarchical campus

B. The uplinks must be increased to 48 Gbps aggregate to eliminate oversubscription before the design can be approved

C. Oversubscription at the access layer should be lower than at the core layer

D. The 2.4:1 oversubscription is normal at the access layer because user traffic is bursty and rarely simultaneous

Correct answer: D. 48 Gbps of access ports against 20 Gbps of uplink is roughly 2.4:1, which is comfortably better than the traditional 20:1 access-layer guideline. Oversubscription is expected and acceptable at the access layer precisely because endpoint traffic is bursty; it should be avoided in the core, where traffic is already aggregated. Option D reverses the correct direction — oversubscription should be highest at the edge and lowest in the core.

Q4. Which function belongs at the distribution layer rather than at the core or access layer?

A. Supplying Power over Ethernet to access points and IP phones

B. Route summarisation toward the core and enforcement of access policy between VLANs

C. Forwarding packets between distribution blocks with no policy applied

D. Acting as the QoS trust boundary for endpoint markings

Correct answer: B. Summarisation and inter-VLAN policy are the defining distribution-layer functions — it is the aggregation and policy layer. PoE and the QoS trust boundary belong to the access layer, which is where endpoints attach. Policy-free forwarding between blocks is the core's only job; adding policy to the core puts it in the path of every conversation in the building.

Q5. An organisation wants to deploy routed access but discovers that a clustered application requires two servers in different access closets to be Layer 2 adjacent. What is the accurate assessment?

A. Routed access cannot satisfy that requirement directly, because it deliberately prevents a VLAN from spanning two access switches

B. Routed access satisfies the requirement because the SVI on each access switch extends the VLAN across the routed link

C. The requirement is satisfied automatically once an FHRP is configured on the two access switches

D. Routed access supports Layer 2 adjacency between access switches as long as both switches use the same VLAN ID

Correct answer: A. The defining property of routed access is that the Layer 3 boundary sits on the access switch, so a VLAN terminates there and cannot span to another closet. That is a feature — it is what removes spanning tree and the FHRP — but it directly conflicts with an application needing Layer 2 adjacency across closets. Using the same VLAN ID on two routed-access switches creates two separate broadcast domains that happen to share a number. Meeting the requirement needs an overlay (VXLAN, or an SD-Access fabric as in 1.5) or a return to Layer 2 access for that block.

3.15 First Hop Redundancy Protocols: HSRP and VRRP

Chapter 1.3 covered first hop redundancy as a design element — when you need it, how it compares with a StackWise Virtual pair, and where it fits among the other high-availability mechanisms. This sub-section is the mechanics: the state machines, the timers, the addressing, and the output. The blueprint verb here is **configure**, and HSRP and VRRP are named explicitly.

The problem, stated precisely

A host learns one default gateway address, by DHCP or by static configuration, and does not re-evaluate it. If the router owning that address fails, every host in the subnet loses off-subnet connectivity — even though a second, perfectly healthy router is one cable away, because no host will ever ask it anything. An FHRP fixes this by having two or more routers share a **virtual IP address** and a **virtual MAC address**, with exactly one of them answering for the pair at any moment.

The virtual MAC is the part that makes failover invisible. When the standby takes over it begins answering for the same MAC address, so hosts do not need to re-ARP and switches relearn the address from the first frame — usually a gratuitous ARP the new active sends immediately.

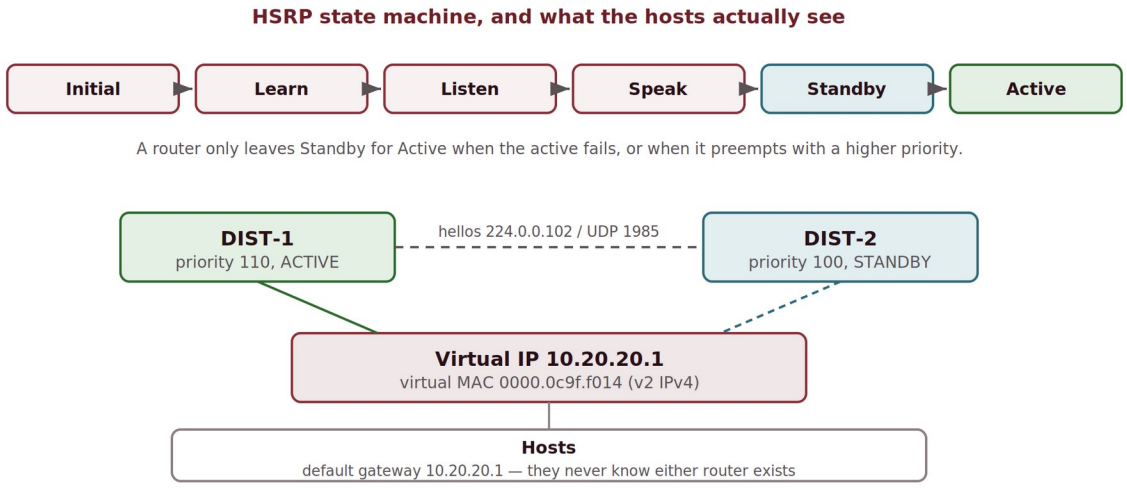

Figure 3.11 — The HSRP state machine, and the single virtual address the hosts actually use.

HSRP: states, timers, and the virtual MAC

	HSRP version 1	HSRP version 2
Group range	0–255	0–4095

	HSRP version 1	HSRP version 2
Virtual MAC (IPv4)	0000.0C07.AC**xx**	0000.0C9F.F**xxx**
Multicast address	224.0.0.2	224.0.0.102
Millisecond timers	No	Yes
IPv6 support	No	Yes

Both versions use **UDP port 1985** for IPv4. Default timers are a **3-second hello and a 10-second hold**, which means a ten-second outage on failure — far too long for voice, and the reason millisecond timers exist. `standby 20 timers msec 250 msec 750` is a reasonable production setting and is safe on any modern platform.

Priority defaults to 100 and higher wins. **Preemption is disabled by default in HSRP**, which surprises people: configure a router with priority 110 and it will not take over from a running active with priority 100 until you add `standby 20 preempt`. That is the opposite of VRRP, where preemption is on by default.

VRRP, and the one thing it does that HSRP cannot

- **RFC 5798**, so it interoperates with other vendors. Roles are **master** and **backup** rather than active and standby.
- **IP protocol 112**, multicast **224.0.0.18**. Default timers are a **1-second advertisement and a 3-second master-down interval** — three times faster than HSRP's defaults out of the box.
- **Preemption is enabled by default.**
- **The virtual IP may be a real interface address.** The router owning it becomes the **address owner** and runs at priority **255**, which cannot be beaten, so it is always master while it is up. HSRP and GLBP both forbid this — their virtual IP must belong to nobody.
- Virtual MAC is **0000.5E00.01xx**, where xx is the group number.

HSRP with everything a production configuration needs

```
! ---- DIST-1: intended active for VLAN 20 ----
track 1 interface TenGigabitEthernet1/1/1 line-protocol
track 10 ip route 0.0.0.0 0.0.0.0 reachability
!
interface Vlan20
 ip address 10.20.20.2 255.255.255.0
```

```
 ip helper-address 10.100.10.10
 standby version 2
 standby 20 ip 10.20.20.1
 standby 20 priority 110
 standby 20 preempt delay minimum 90    ! HSRP does NOT preempt by default
 standby 20 timers msec 250 msec 750
 standby 20 authentication md5 key-string 7 <encrypted-key>
 standby 20 track 1 decrement 20
 standby 20 track 10 decrement 30
 standby 20 name VLAN20-GW
!
! ---- DIST-2: mirror image, priority 100 ----
interface Vlan20
 ip address 10.20.20.3 255.255.255.0
 ip helper-address 10.100.10.10
 standby version 2
 standby 20 ip 10.20.20.1
 standby 20 priority 100
 standby 20 preempt
 standby 20 timers msec 250 msec 750
 standby 20 authentication md5 key-string 7 <encrypted-key>
```

The same requirement in VRRP, including the address-owner case

```
! ---- RTR-1 OWNS the virtual address, so it is priority 255 ----
interface Vlan30
 ip address 10.20.30.1 255.255.255.0
 vrrp 30 ip 10.20.30.1              ! same as the interface address
 vrrp 30 timers advertise msec 300
 vrrp 30 authentication md5 key-string 7 <encrypted-key>
 vrrp 30 track 1 decrement 30
!
! ---- RTR-2 ----
interface Vlan30
 ip address 10.20.30.2 255.255.255.0
 vrrp 30 ip 10.20.30.1
 vrrp 30 priority 100
 vrrp 30 timers learn               ! adopt the master's interval
!
RTR-1# show vrrp brief
Interface   Grp Pri Time  Own Pre State  Master addr  Group addr
Vl30         30 255 0.300 Y   Y   Master 10.20.30.1   10.20.30.1

! Own = Y means this router is the address owner, priority forced
```

```
! to 255. No other router can ever take the master role from it
! while it is up.
```

Load sharing with multiple groups

HSRP and VRRP each elect one forwarder per group, so a pair of routers with a single group leaves one of them idle. The standard answer is **multiple groups with alternating priorities**: DIST-1 is active for the odd VLANs and standby for the even ones, DIST-2 the reverse. Both routers forward, both uplinks carry traffic, and either can carry everything.

GLBP — Cisco proprietary, covered in Chapter 1.3 — does this natively by handing different virtual MAC addresses to different hosts as they ARP, so one group load-shares without alternating anything. It is elegant and it is Cisco-only, which is why the multi-group pattern remains the common answer.

Reading HSRP state, and the alignment check that matters

```
DIST-1# show standby brief
                     P indicates configured to preempt.
                     |
Interface  Grp  Pri P State    Active         Standby        Virtual IP
Vl20        20  110 P Active   local          10.20.20.3     10.20.20.1
Vl30        30  100 P Standby  10.20.30.3     local          10.20.30.1

DIST-1# show standby Vlan20 | include State|Virtual MAC|Priority|Track
  State is Active
  Active virtual MAC address is 0000.0c9f.f014   ! v2 IPv4 = 0c9f.fXXX
  Priority 110 (configured 110)
    Track object 1 state Up decrement 20
    Track object 10 state Up decrement 30

! Alignment check: the HSRP active router should also be the STP
! root for the same VLAN, or every frame crosses the peer link twice.
DIST-1# show spanning-tree vlan 20 | include This bridge is the root
             This bridge is the root

DIST-1# show standby Vlan20 | include Group name|Authentication
  Authentication MD5, key-string
  Group name is "VLAN20-GW" (cfgd)
```

⚠️ Common Pitfall — A version mismatch makes both routers active, and hosts see it intermittently

Configure `standby version 2` on one router and leave the other at version 1, and the two send hellos in formats the other ignores. Each concludes it is alone and becomes active, so two routers answer for the same virtual IP with **different virtual MAC addresses** — 0000.0c07.acXX on one and 0000.0c9f.fXXX on the other. Hosts work or fail depending on which MAC their ARP cache happens to hold, and the symptom changes when the cache ages out. `show standby brief` on each router showing itself as Active with no standby listed is the proof, and the virtual MAC prefix tells you which version each is running.

✅ Exam Tip — The six differences that answer HSRP-versus-VRRP questions

Standard — HSRP is Cisco, VRRP is RFC 5798. **Roles** — active/standby versus master/backup. **Transport** — UDP 1985 to 224.0.0.102 versus IP protocol 112 to 224.0.0.18. **Timers** — 3/10 versus 1/3. **Preemption** — off by default in HSRP, on by default in VRRP. **Address owner** — only VRRP allows the virtual IP to be a real interface address, in which case that router runs at priority 255. A scenario naming another vendor, or requiring the virtual IP to equal an interface IP, is asking for VRRP.

🌍 From the Field — Tracking is what turns an FHRP from a demo into a design

An FHRP without tracking protects against exactly one failure: the active router losing power. It does nothing about the far more common case where the router is perfectly healthy and its path to the rest of the network is not — a failed uplink, a dead upstream neighbour, a withdrawn default route. Without tracking, HSRP stays cheerfully active on a router with nowhere to forward, and every host in the VLAN black-holes while a working peer sits in standby. Track the uplink with `line-protocol`, track the default route with `ip route reachability`, and decrement by more than the priority gap so the handover actually happens. Then test it by shutting the uplink, because a tracking configuration that decrements by too little looks correct and does nothing.

Summary

A first hop redundancy protocol lets two or more routers share a virtual IP and virtual MAC address so that hosts, which never re-evaluate their default gateway, keep working when the router owning that address fails; the shared virtual MAC is what makes the failover invisible, with a gratuitous ARP prompting switches to relearn it. HSRP uses UDP 1985, defaults to a 3-second hello and 10-second hold, and does not preempt unless told to; version 1 supports groups 0–255 with virtual MAC 0000.0C07.ACxx on 224.0.0.2, while version 2 supports groups 0–4095 with 0000.0C9F.Fxxx on 224.0.0.102 and adds millisecond timers and IPv6. VRRP is the IETF standard using IP protocol 112 to 224.0.0.18 with a 1-second advertisement and 3-second master-down interval, preempts by default, uses virtual MAC 0000.5E00.01xx, and uniquely permits the virtual IP to be a real interface address — making that router the address owner at priority 255. Load sharing is achieved with multiple groups and alternating priorities, or natively with GLBP. A version mismatch produces two active routers with different virtual MACs and intermittent host symptoms. Tracking an uplink and the default route, decrementing by more than the priority gap, is what makes the protocol respond to the failures that actually occur — and the active router should be aligned with the spanning-tree root for the same VLAN.

Key Takeaways

- An FHRP shares a **virtual IP and a virtual MAC**; the shared MAC is why hosts never need to re-ARP.
- **HSRP: UDP 1985, hello 3 s / hold 10 s, preemption OFF by default.**
- **HSRPv1: groups 0–255, MAC 0000.0C07.ACxx, 224.0.0.2. HSRPv2: groups 0–4095, MAC 0000.0C9F.Fxxx, 224.0.0.102, msec timers, IPv6.**
- **VRRP: IP protocol 112, 224.0.0.18, advertise 1 s / master-down 3 s, preemption ON by default, MAC 0000.5E00.01xx.**
- **Only VRRP allows the virtual IP to be a real interface address** — that router is the address owner at priority **255**.
- **A version mismatch makes both routers active** with different virtual MACs; the symptom follows the hosts' ARP caches.
- Load share with **multiple groups and alternating priorities**, or natively with **GLBP**.
- **Track the uplink and the default route, and decrement by more than the priority gap** — otherwise nothing fails over.
- **Align the active router with the spanning-tree root** for the same VLAN, or every frame crosses the peer link twice.

Knowledge Check — 3.15 HSRP and VRRP

Q1. A design requires a redundant gateway using an IETF-standard protocol, and the virtual IP address must be the same as one router's physical interface address. Which protocol meets both requirements, and what is that router's priority?

A. HSRP version 2, with the owner running at priority 255

B. GLBP, with the AVG running at priority 255

C. VRRP, with the address owner running at the configured priority of 100

D. VRRP, with the address owner running at priority 255

Correct answer: D. VRRP is the only one of the three that permits the virtual address to be an address a router actually owns, and it forces that router to priority 255 — a value no other router can match, so the owner is master whenever it is up. HSRP and GLBP both require the virtual IP to belong to no interface. VRRP is also the only IETF-standard option of the three, which the scenario's mention of a standard points at.

Q2. Two distribution switches are configured for HSRP group 20 with the same virtual IP, but `show standby brief` on each shows itself as Active with no standby listed. What is the most likely cause?

A. Both routers are configured with the same priority, so neither can win the election

B. Preemption is disabled on both routers, so neither can transition out of standby

C. An HSRP version mismatch, so each router ignores the other's hellos and both become active

D. The virtual IP address is also configured on a physical interface, which HSRP forbids

Correct answer: C. HSRPv1 and HSRPv2 use different hello formats, different multicast addresses, and different virtual MAC ranges, so a mismatched pair never sees each other and each concludes it is alone. Two routers then answer for the same virtual IP with different virtual MACs, and hosts succeed or fail depending on which MAC their ARP cache holds. Equal priorities would still produce one active and one standby, decided by the higher IP address.

Q3. HSRP is configured with priority 110 on DIST-1 and 100 on DIST-2, with tracking that decrements by 5 when the uplink fails. DIST-1's uplink fails but it remains active. Why?

A. The decrement is smaller than the priority gap, so DIST-1's priority only drops to 105 and still exceeds 100

B. Tracking only applies when preemption is disabled on the standby router

C. HSRP ignores tracking events for line-protocol objects on uplink interfaces

D. The track object must be applied to the standby router rather than the active one

Correct answer: A. Tracking works by reducing priority, so the decrement must be large enough to fall below the peer's priority — with a gap of 10, a decrement of 5 leaves DIST-1 at 105 and still winning. The configuration looks correct and does nothing, which is why the only reliable verification is to shut the uplink and watch the roles change. Preemption on the peer is also required for it to take over once it does win.

Q4. Which statement about HSRP preemption is correct?

A. It is enabled by default, matching VRRP behaviour

B. It is disabled by default, so a higher-priority router will not take over from a running active until `preempt` is configured

C. It applies only to HSRP version 2 and is unavailable in version 1

D. It causes the active router to relinquish the role every time the hold timer expires

Correct answer: B. HSRP deliberately avoids churn by not preempting: whichever router became active keeps the role even when a better-configured router appears. That is the opposite of VRRP, where preemption is on by default, and it is a frequent source of confusion when a supposedly primary router sits in standby after a maintenance window. Adding `preempt`, ideally with a delay so the router has time to converge its routing first, is what enables the intended behaviour.

Q5. Why should the HSRP active router for a VLAN also be the spanning-tree root for that VLAN?

A. Because HSRP will not elect an active router that is not the spanning-tree root

B. Because the virtual MAC address is derived from the root bridge ID

C. Otherwise frames from access switches reach the root first and must cross the peer link to be routed, adding a hop and doubling load on that link

D. Because spanning tree blocks the peer link unless the roles are aligned

Correct answer: C. Spanning tree decides where frames travel at Layer 2 and the FHRP decides where they are routed. If those two are on different switches, every frame is forwarded up to the root, across the inter-switch link, and only then routed — an extra hop for all traffic in that VLAN and double the load on the peer link. Aligning them costs nothing: set `spanning-tree vlan X root primary` on the same switch that carries the higher HSRP priority. The protocols are otherwise independent.

Appendix B · Command Quick Reference

This appendix collects the commands that carry the most examination weight, grouped by domain. It is a reminder, not a tutorial — each block assumes you have read the section it comes from, and the page references are in the contents.

Two conventions: ! introduces a comment, and anything in angle brackets is a value you supply. The verification commands are listed with the configuration deliberately, because on this examination *knowing what proves it worked* is worth as much as knowing how to configure it.

Domain 1 · First hop redundancy

```
! ---- HSRP version 2, the production shape -----------------------
interface Vlan10
 ip address 10.10.10.2 255.255.255.0
 standby version 2
 standby 10 ip 10.10.10.1
 standby 10 priority 110
 standby 10 preempt delay minimum 90
 standby 10 timers msec 250 msec 750
 standby 10 authentication md5 key-string 7 <hash>
 standby 10 track 1 decrement 30
!
track 1 interface GigabitEthernet1/0/1 line-protocol
!
! ---- VRRP -------------------------------------------------------
interface Vlan20
 vrrp 20 ip 10.10.20.1
 vrrp 20 priority 120
 vrrp 20 timers advertise 1
 vrrp 20 authentication md5 key-string <key>
!
! ---- Verify -----------------------------------------------------
show standby brief
show standby vlan 10 10
show vrrp brief
show track 1
```

The three lines that separate a working design from a demonstration: **`preempt delay minimum`** stops a rebooting router from taking the role back before its uplinks and routing protocols have converged; **`track … decrement`** must decrement by *more* than the priority gap or it changes

nothing; and **`standby version 2`** must match on both routers or you get two active routers and intermittent symptoms.

Domain 2 · VRF, GRE, and IPsec

```
! ---- VRF-Lite, multiprotocol syntax --------------------------------
vrf definition CUSTOMER-A
 rd 65001:100
 address-family ipv4
  route-target export 65001:100
  route-target import 65001:100
 exit-address-family
!
interface GigabitEthernet0/0/1
 vrf forwarding CUSTOMER-A          ! do this BEFORE the ip address
 ip address 10.1.1.1 255.255.255.0
!
router ospf 10 vrf CUSTOMER-A
 network 10.1.1.0 0.0.0.255 area 0
 capability vrf-lite
!
! ---- Verify -------------------------------------------------------
show vrf
show ip route vrf CUSTOMER-A
ping vrf CUSTOMER-A 10.1.1.2
traceroute vrf CUSTOMER-A 10.1.1.2
```

```
! ---- Static VTI, the modern site-to-site tunnel ---------------------
crypto ikev2 proposal PROP-1
 encryption aes-cbc-256
 integrity sha256
 group 14
!
crypto ikev2 policy POL-1
 proposal PROP-1
!
crypto ikev2 keyring KR-1
 peer HUB
  address 198.51.100.1
  pre-shared-key local <key>
  pre-shared-key remote <key>
!
crypto ikev2 profile PROF-1
```

```
 match identity remote address 198.51.100.1 255.255.255.255
 identity local address 203.0.113.1
 authentication local pre-share
 authentication remote pre-share
 keyring local KR-1
!
crypto ipsec transform-set TS-1 esp-aes 256 esp-sha256-hmac
 mode tunnel
!
crypto ipsec profile IPSEC-PROF
 set transform-set TS-1
 set ikev2-profile PROF-1
!
interface Tunnel10
 ip address 10.255.10.1 255.255.255.252
 ip mtu 1400
 ip tcp adjust-mss 1360
 tunnel source GigabitEthernet0/0/0
 tunnel destination 198.51.100.1
 tunnel mode ipsec ipv4
 tunnel protection ipsec profile IPSEC-PROF
!
! ---- Verify -------------------------------------------------------
show crypto ikev2 sa
show crypto ipsec sa | include pkts|peer
show interface tunnel10 | include line protocol|MTU
```

⚠️ Common Pitfall — `vrf forwarding` after `ip address` deletes the address

Assigning an interface to a VRF removes its IP configuration, because the address belonged to the global table. Configure the VRF membership first, then the address. The same applies when you *move* an interface between VRFs — the address goes with the old VRF and you must re-enter it.

Domain 3 · Trunking, EtherChannel, and spanning tree

```
! ---- Trunk --------------------------------------------------------
interface GigabitEthernet1/0/24
 switchport trunk encapsulation dot1q
 switchport mode trunk
 switchport trunk native vlan 999
 switchport trunk allowed vlan 10,20,30
```

The sample ends here.

There are 471 more pages.

ENCOR 350-401 v1.2 — Complete Learning Guide

All six exam domains, in blueprint order

49 sub-sections mapped one-to-one onto the blueprint

245 knowledge-check questions, every one explained

47 diagrams and a full command quick reference

A ten-week study plan and a practice-lab guide

Get the full book

leanpub.com/encor