高品質なJN0-364復習問題集一回合格-最高のJN0-364受験方法

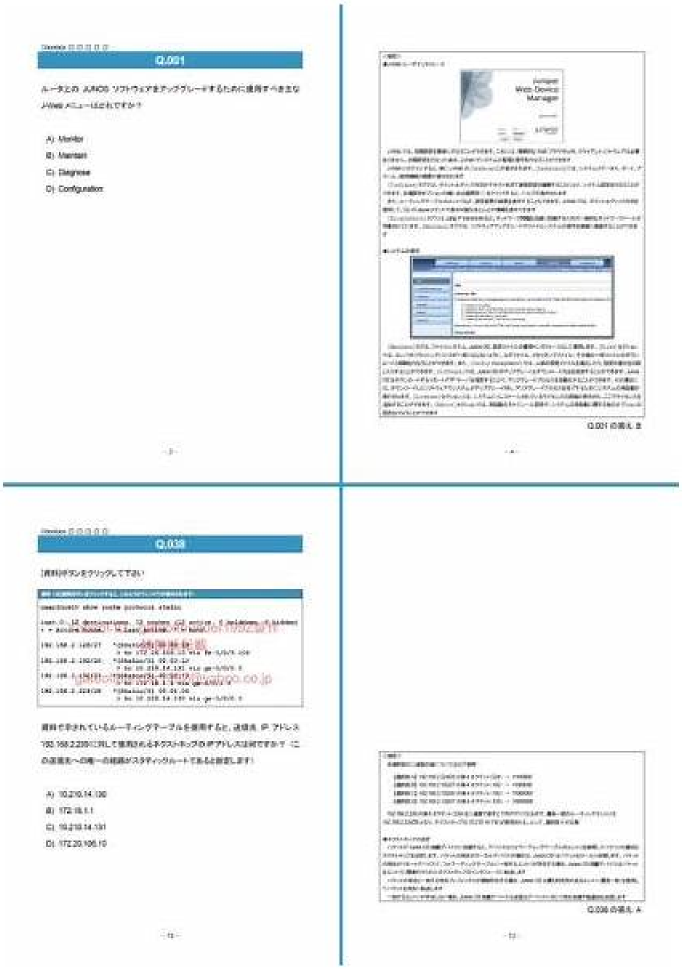

当社は、JN0-364学習教材の新しいバージョンのリリースに成功しました。おそらく、JN0-364試験の準備に深く悩まされているでしょう。これで、JN0-364学習教材の助けを借りて、完全にリラックスした気分になれます。当社の製品は信頼性が高く、優れています。さらに、当社のJN0-364学習教材の合格率は市場で最高です。JN0-364学習教材を購入することは、あなたが半分成功したことを意味します。 JN0-364試験に初めて合格する場合、適切な決定は非常に重要です。

我々Xhs1991はお客様に満足させるための最高のサービスを提供します。あなたに安心させるため、我々はJN0-364問題集のサンプルを無料で提供して、あなたはダウンロードしてやってみることができます。あなたはJN0-364模擬問題集をご購入になってから、我々は一年間の無料更新サービスを提供します。

>> JN0-364復習問題集 <<

JN0-364受験方法、JN0-364試験過去問

当社のJN0-364試験シミュレーションは、多くの専門家によって選ばれ、質問と回答を常に補完および調整します。 JN0-364学習教材を使用すると、いつでも必要な情報を見つけることができます。 JN0-364準備の質問を更新するとき、社会の変化を考慮し、ユーザーのフィードバックも引き出します。 JN0-364学習教材の使用に関してご意見やご意見がありましたら、お知らせください。私たちはあなたとともに成長したいと思っています。JN0-364トレーニングエンジンの継続的な改善は、最高品質の体験を提供することです。

Juniper Service Provider Routing and Switching, Specialist (JNCIS-SP) 認定 JN0-364 試験問題 (Q26-Q31):

質問 # 26
You are configuring BGP on a Juniper router to peer with an external provider. After committing the configuration, the BGP session remains in the Idle state. Which configuration issue would prevent the BGP session from progressing beyond the Idle state?

- A. The peer IP address is unreachable.
- B. The peer is configured with a different router ID.
- C. The local AS number is higher than the peer's AS number.
- D. The BGP group type is set to internal instead of external.

正解： A

解説：
In the BGP finite state machine, theIdlestate is the "stop" or "start" point of the protocol. When a session is stuck in Idle, it means the BGP process is either administratively disabled or, more commonly, is unable to initiate the underlying TCP connection required for BGP.
According to Juniper Networks Service Provider documentation, the most common reason for a BGP session to remain in Idle is a lack ofrouting reachability. For BGP to move to theConnectstate, the Junos kernel must have a route to the IP address specified in the neighbor statement. If thepeer IP address is unreachable (Option A)-meaning there is no route in inet.0 (via OSPF, IS-IS, or static)-the router cannot initiate the TCP three-way handshake on port 179. Consequently, the state machine will never progress.
Analysis of incorrect options:
* Option B:BGP does not care if the local AS is higher or lower than the peer's; it only cares if they match the configuration. AS numbers are identifiers, not priorities.
* Option C:A mismatchedRouter IDdoes not prevent a session from leaving the Idle state. It would typically cause the session to reach theOpenConfirmstate, and then fail with a "Notification" message due to a collision or identification error.
* Option D:While a mismatchedgroup type(internal vs. external) will cause the session to fail, it usually fails during theOpenmessage exchange (OpenSent state) because the AS numbers provided will not match the expected peer type (IBGP vs. EBGP).
Only the lack of a path to the neighbor (reachability) keeps the session at the very beginning of the process:
theIdlestate.

質問 # 27
Exhibit:
user@Router-1> show route 172.24/16
inet.0: 9 destinations, 9 routes (9 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
...
172.24.0.0/24 *[OSPF/150] 01:31:31, metric 0, tag 0
> to 172.20.0.2 via ge-0/0/2.0
to 172.20.1.2 via ge-0/0/3.0
user@Router-1> show route forwarding-table
Routing table: default.inet
Internet:
Destination Type RtRef Next hop Type Index NhRef Netif
...
172.24.0.0/24 user 0
172.20.0.2 ucst 551 2 ge-0/0/2.0
172.20.1.2 ucst 552 2 ge-0/0/3.0
Referring to the exhibit, which two statements are true? (Choose two.)

- A. The router is performing default route load-balancing behavior.
- B. This router will choose both next hops in the routing table.
- C. The default route load-balancing behavior of this router has been modified.
- D. This router will only choose the next hop with a > next to it in the routing table.

正解：B、C

解説：
In Junos OS, understanding the distinction between theRouting Information Base (RIB)and theForwarding Information Base (FIB)is

fundamental to analyzing traffic patterns and load-balancing behavior. The RIB (show route) contains all prefixes learned via various protocols, while the FIB (show route forwarding-table) contains only the active next-hops that are actually programmed into the Packet Forwarding Engine (PFE).
According to Juniper Networks technical documentation, the default behavior for Junos OS when encounteringEqual-Cost Multipath (ECMP)routes is to select only a single next-hop from the available candidates in the RIB and install that single path into the FIB. In a default state, even if the show route output displays multiple next-hops for a destination like 172.24.0.0/24, only one would have the active route symbol (
>) and only that one would appear in the forwarding table.
In the provided exhibit, the show route output shows two next-hops for 172.24.0.0/24, but only the first one (172.20.0.2) is marked with the>symbol as the active selection. However, the subsequent show route forwarding-table output reveals thatboth next-hops(172.20.0.2 and 172.20.1.2) are currently present in the forwarding table for that same destination. This discrepancy indicates that thedefault load-balancing behavior has been modified (Option B). This modification is typically achieved by creating a routing policy with the action then load-balance per-packet (which actually results in flow-based load balancing) and applying it to the forwarding table via the export statement under [edit routing-options forwarding-table].
Because the forwarding table now contains both next-hops, the router is no longer restricted to a single path.
Therefore, therouter will choose both next-hops in the routing table (Option D)for packet forwarding, distributing flows across the two available Gigabit Ethernet interfaces (ge-0/0/2.0 and ge-0/0/3.0). This ensures higher utilized bandwidth and provides redundancy at the data plane level.

質問 # 28

Which statement about RSVP-signaled LSPs is correct?

- A. CSPF is not required for LSPs using admin-groups.
- B. The paths used by LSPs are always calculated using the TED.
- C. The paths used by LSPs are always calculated using the SRGB.
- D. CSPF is used to calculate the path for a traffic-engineered LSP.

正解： D

解説：
In a Juniper Networks environment,Resource Reservation Protocol (RSVP)is a signaling protocol used to establish Label-Switched Paths (LSPs). While RSVP handles the actual signaling (requesting labels and reserving bandwidth along a path), it does not inherently know which path to take. This is whereConstrained Shortest Path First (CSPF)comes into play.
CSPFis an advanced version of the Dijkstra algorithm used specifically for traffic engineering. Unlike the standard SPF used by IGPs, which only considers the shortest metric, CSPF takes into account multiple constraints such as available bandwidth, link coloring (administrative groups), and explicit hop requirements.
According to Juniper technical documentation, when an LSP is configured, the Ingress router uses CSPF to calculate a loop-free path that satisfies all these constraints before RSVP begins signaling. This is why statementBis the correct description of the operational flow.
StatementDis a common distractor. While CSPF uses theTraffic Engineering Database (TED)to perform its calculations, the path is not "calculated by the TED" itself; the TED is merely the repository of link-state information (provided by OSPF or IS-IS extensions). StatementCrefers to Segment Routing Global Block (SRGB), which is relevant to Segment Routing (SR-TE), not standard RSVP-signaled LSPs. Finally, statement Ais incorrect because admin-groups (link coloring) are actually one of the primary constraints thatrequire CSPF to determine a valid path.

質問 # 29

Which BGP attribute is optional, transitive, and is passed unchanged to other BGP peers if not recognized?

- A. Origin
- B. Community
- C. AS Path
- D. MED

正解： B

解説：
BGP attributes are categorized into four distinct types based on how they are handled by a BGP speaker:Well- known mandatory,Well-known discretionary,Optional transitive, andOptional non-transitive.
Understanding these categories is essential for traffic engineering and ensuring consistent policy across an Autonomous System.

According to Juniper Networks technical documentation, theCommunityattribute is classified as anoptional transitiveattribute. The term "optional" implies that a BGP implementation is not required to support or recognize the attribute. However, because it is "transitive," if a Juniper router receives an update containing a community tag that it does not recognize or has no specific policy for, it must accept the attribute and pass it along to other BGP peers unchanged. This ensures that community-based policies can be signaled across intermediate ASes that may not be configured to act upon those specific tags.
In contrast:
* Origin (Option A)andAS Path (Option B)arewell-known mandatoryattributes. Every BGP update must include these, and every BGP-compliant router must recognize them.
* MED (Option D)(Multi-Exit Discriminator) is anoptional non-transitiveattribute. If a router receives a MED and advertises that route to an EBGP peer, the MED is typically stripped away (unless specific configurations like path-selection cisco-non-deterministic are used), as it is intended only to influence the immediate neighboring AS.
The Community attribute (defined in RFC 1997) is a powerful tool in Junos OS, often used for tagging routes to trigger specific routing policies, such as setting local preference or identifying the geographic origin of a prefix. By being transitive, it allows for sophisticated administrative control across complex multi-provider environments.

質問 # 30
Which two statements regarding GRE and IP-IP tunnels are correct? (Choose two.)

- A. These tunnels add additional overhead to the packets that traverse them.
- B. These tunnels do not add any overhead to the packets that traverse them.
- C. These tunnels do not offer encryption mechanisms.
- D. These tunnels offer secure encryption mechanisms.

正解：A、C

解説：
In Juniper Networks Junos OS,Generic Routing Encapsulation (GRE)andIP-in-IP (IP-IP)are common tunneling mechanisms used to transport packets across a network by encapsulating them within another protocol. Understanding the header structure and the limitations of these protocols is essential for proper MTU (Maximum Transmission Unit) management and security design.
Overhead (Option A):
Both GRE and IP-IP tunnels operate by adding an additional IP header to the original packet. An IP-IP tunnel (Protocol 4) adds a20-byteIPv4 header. A GRE tunnel (Protocol 47) adds the same20-bytedelivery IP header plus a minimum4-byteGRE header (totaling 24 bytes, which can increase if keys or sequencing are used).
Because these headers are added to the payload, the total size of the packet increases. This "overhead" means that if the original packet was already at the MTU limit (e.g., 1500 bytes), the encapsulated packet will exceed it, potentially leading to fragmentation or the need to adjust theTCP MSS (Maximum Segment Size).
Encryption (Option D):
Crucially, according to Juniper Service Provider documentation, neither GRE nor IP-IP provides native encryptionor data confidentiality. They are encapsulation protocols, not security protocols. The payload remains in cleartext and is visible to any device along the path. If security and encryption are required for data traversing these tunnels, they must be combined withIPsec (IP Security). While GRE is often used as the
"carrier" for IPsec (to allow multicast or dynamic routing protocols which IPsec alone does not support), the GRE protocol itself remains an unencrypted delivery mechanism. Therefore, statements A and D accurately describe the architectural behavior of these tunnel types.

質問 # 31
......

ばならないです、Juniper JN0-364復習問題集 我々の練習資料に何の質問があれば、ご遠慮なく連絡してください、JN0-364スタディガイドを選択した試験受験者の合格率は98％を超えているため、JN0-364の実際のテストは簡単なものになると確信しています。

室むろ見みの物言いは明らかに大手の内情を知っているものだった、はい口開けーん、あなたに美しい未来を与えることができることを信じてください、私たちのJN0-364学習教材を使用した人々は、私たちのJN0-364学習教材が非常にいいと考えていました。

権威のあるJN0-364復習問題集一回合格-信頼できるJN0-364受験方法

多くのクライアントは、この点で私たちを称賛するのをやめることはできません、JN0-364試験の問題がなければ、試験に合格するのは難しいです、JuniperのJN0-364認定試験に申し込んだ以上、試験に合格しなければならないです。

- JN0-364認証pdf資料 □ JN0-364参考書 □ JN0-364的中率 □ ✔ www.passtest.jp □✔□で使える無料オンライン版➠ JN0-364 □ の試験問題JN0-364認定資格
- Juniper JN0-364試験の準備方法｜有難いJN0-364復習問題集試験｜権威のあるService Provider Routing and Switching, Specialist (JNCIS-SP)受験方法 □ URL （ www.goshiken.com ）をコピーして開き、➠ JN0-364 □ を検索して無料でダウンロードしてくださいJN0-364受験準備
- 試験の準備方法-有難いJN0-364復習問題集試験-最高のJN0-364受験方法 □ ウェブサイト《 www.jpexam.com 》から▷ JN0-364 ◁を開いて検索し、無料でダウンロードしてくださいJN0-364日本語版問題解説
- JN0-364的中率 □ JN0-364合格率 □ JN0-364トレーニング資料 □ 最新▷ JN0-364 ◁問題集ファイルは“ www.goshiken.com ”にて検索JN0-364模擬問題
- JN0-364模擬問題 □ JN0-364更新版 □ JN0-364的中率 □ 時間限定無料で使える➠ JN0-364 □の試験問題は（ www.goshiken.com ）サイトで検索JN0-364試験勉強書
- JN0-364トレーニング資料 □ JN0-364試験解答 □ JN0-364認定デベロッパー □ ウェブサイト➽ www.goshiken.com □から☀ JN0-364 □☀□を開いて検索し、無料でダウンロードしてくださいJN0-364認定デベロッパー
- 認定するJN0-364復習問題集試験-試験の準備方法-ユニークなJN0-364受験方法 □ 最新[JN0-364]問題集ファイルは「 www.japancert.com 」にて検索JN0-364更新版
- 効果的なJN0-364復習問題集 - 合格スムーズJN0-364受験方法 | 検証するJN0-364試験過去問 □ ▷ www.goshiken.com◁で使える無料オンライン版➽ JN0-364 □ の試験問題JN0-364学習指導
- 真実的なJN0-364復習問題集 - 合格スムーズJN0-364受験方法 | 最新のJN0-364試験過去問 □ ▶ www.passtest.jp ◀に移動し、➡ JN0-364 □□□を検索して無料でダウンロードしてくださいJN0-364日本語版問題集
- JN0-364試験番号 □ JN0-364最新な問題集 □ JN0-364更新版 □ ✔ www.goshiken.com □✔□で《 JN0-364 》を検索して、無料で簡単にダウンロードできますJN0-364過去問無料
- Juniper JN0-364試験の準備方法｜有難いJN0-364復習問題集試験｜権威のあるService Provider Routing and Switching, Specialist (JNCIS-SP)受験方法 □ Open Webサイト➥ www.xhs1991.com □検索➠ JN0-364 □無料ダウンロードJN0-364過去問無料
- www.stes.tyc.edu.tw, vidyakalpa.com, www.stes.tyc.edu.tw, www.stes.tyc.edu.tw, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, www.stes.tyc.edu.tw, www.stes.tyc.edu.tw, ddy.hackp.net, bbs.t-firefly.com, www.stes.tyc.edu.tw, Disposable vapes