F5.F5CAB1.v2026-02-18.q19

Exam Code:	F5CAB1
Exam Name:	BIG-IP Administration Install, Initial Configuration, and Upgrade
Certification Provider:	F5
Free Question Number:	19
Version:	v2026-02-18
# of views:	114
# of Questions views:	190
https://www.exam-tests.com/F5CAB1-exam/F5.F5CAB1.v2026-02-18.q19.html	

NEW QUESTION: 1
An F5 BIG-IP Administrator is asked to report which modules areprovisionedon the BIG-IP.
In which two ways can this be done?
(Choose two.)
A. Via the GUI atSystem # Resource Provisioning # Module Allocation
B. Via TMSH withshow /sys provision
C. Via the GUI atStatistics # Module Statistics # System
D. Via TMSH withlist /sys provision
Answer: A,D (LEAVE A REPLY)
Provisioning determines:
* Which BIG-IP modules are enabled (LTM, ASM, APM, AFM, DNS, etc.)
* Their provisioning levels (None, Minimal, Nominal, Dedicated)
Two accurate ways to view provisioning settings are:
A). GUI - System # Resource Provisioning # Module Allocation
This is the primary GUI screen showing:
* All modules
* Their provisioning level
* System resource distribution impact
Administrators commonly use this page to confirm or change module provisioning.
D). TMSH - list /sys provision
This tmsh command displays each module and its provisioning level:
sys provision ltm { level nominal }
sys provision asm { level none }
This is the authoritative CLI method for checking module provisioning configurations.
Why the other options are incorrect:

B). show /sys provision
* Showsruntimeinformation butnot the actual configuration levels.
* list is the correct command for configuration details.
C). Statistics # Module Statistics
* Shows performance statistics, NOT provisioning status.
Therefore, the correct responses areAandD.

NEW QUESTION: 2
A BIG-IP Administrator needs to purchase new licenses for a BIG-IP appliance.
The administrator needs to know:
* Whether a module is licensed
* The memory requirement for that module
Where should the administrator view this information in theSystem menu?
A. Configuration OVSDB
B. Software Management
C. Configuration Device
D. Resource Provisioning
Answer: D (LEAVE A REPLY)
To understand:
* Which modules arelicensed
* Which modules areprovisioned
* Theresource requirements(CPU / RAM) of each module
The administrator uses:
System Resource Provisioning
This page displays:
* All modules present in the license
* Whether they are enabled or disabled
* Required memory to activate each module
* CPU and disk allocation information
* Provisioning level options (None / Minimal / Nominal / Dedicated)
This is the exact location where BIG-IP administrators evaluate module capacity before enabling or purchasing licensing upgrades.
Why the other options are incorrect:
A). Configuration OVSDB
* Used for network virtualization integrations, not licenses or modules.
B). Software Management
* Used for software image installation, not licensing.
C). Configuration Device
* Displays hostname, failover settings, device properties - not module resource requirements.

Thus, module licensing and memory requirement data are found underResource Provisioning.

NEW QUESTION: 3
The Configuration Utility of a BIG-IP device is currently accessible via its management IP10.53.1.245from all VLANs.
The BIG-IP Administrator needs to restrict access so only hosts from the10.0.0.0/24subnet can access the Configuration Utility.
Which TMSH command accomplishes this?
A. (tmos)# create /net acl MGMT.HTTP rule add { (permit tcp 10.0.0.0 0.0.0.255 host 10.53.1.245 http) }
B. (tmos)# modify /ltm httpd allow replace-all-with {10.0.0.0/24}
C. (tmos)# create /net acl MGMT.HTTP rule add { (permit tcp 10.0.0.0/24 10.53.1.245 http) (deny ip any any http) }
D. (tmos)# modify /sys httpd allow replace-all-with {10.0.0.0/24}
Answer: D (LEAVE A REPLY)
BIG-IP controls access to the web-based Configuration Utility (TMUI) through the/sys httpd allowlist. This parameter specifies which client IPs or subnets may initiate HTTP/HTTPS connections to the management interface.
To restrict TMUI access toonlythe 10.0.0.0/24 subnet:
* The correct method is tomodify the HTTPD allow listso that it contains only this subnet.
* This requires replacing the entire current list with the new subnet using:
modify /sys httpd allow replace-all-with {10.0.0.0/24}
This ensures thatonlyclients within 10.0.0.0/24 can reach the Configuration Utility.
Why the other options are incorrect:
* Options A and Ccreate network ACL objects under /net acl, which apply to data-plane traffic, not management-plane TMUI access. TMUI access is not controlled by LTM ACLs but by the HTTPD allow directive.
* Option Bis incorrect syntax and references /ltm httpd, which is not the proper object; the correct hierarchy is /sys httpd.
Thus, only modifying the/sys httpd allowlist achieves the required restriction.

NEW QUESTION: 4
Which command will display thecurrent active volumeon a BIG-IP system?
A. tmsh show sys version
B. tmsh show sys software status
C. tmsh list sys software update
Answer: (SHOW ANSWER)
To identify which boot volume is currently active on a BIG-IP system, the correct command is:
tmsh show sys software status

This command displays:
* All installed boot volumes (HD1.1, HD1.2, HD1.3, etc.)
* The BIG-IP software version installed on each volume
* TheActivefield, indicating which volume the system is currently booted from
* The installation status ("complete", "in-progress", "allowed")
This is thestandard and authoritativeway to determine the active boot location.
Why the other options are incorrect:
A). tmsh show sys version
* Displays OS version, build, and date.
* Doesnotshow boot locations or which volume is active.
C). tmsh list sys software update
* Shows software update configurations, not boot volume status.
* Does not display which volume is active.

NEW QUESTION: 5
An administrator is in the process of reactivating the license using the interface displayed in the exhibit.

What is the address of the license server to which the BIG-IP device must be able to establish an outbound connection in order to use theAutomatic Activation Method?
A. license.f5.com
B. callhome.f5.com
C. ask.f5.com
D. activate.f5.com
Answer: (SHOW ANSWER)
When you chooseAutomaticas the activation method in the License , Re-activate screen, the BIG-IP device itself contacts F5'slicense activation serviceover the Internet.
For successful automatic activation:
* The BIG-IP must have outbound network connectivity (typically via the management interface).

* DNS resolution and routing must allow it to reach theF5 license activation host(the one shown in option D).
* The device sends its dossier and registration key to that service and receives an updated license file in return, which is then installed automatically.
The other hostnames in the options are not used by BIG-IP for license activation, so they cannot be correct in the context ofAutomatic Activation.

NEW QUESTION: 6
An organization is planning to upgrade a BIG-IP system from16.1.xto17.1.x.
For a successful upgrade, theService Check Datemust be equal to or newer than the License Check Date required for 17.1.x.
Which command will show the Service Check Date on the BIG-IP system being upgraded?
A. grep "Service check date" /config/bigip.license
B. grep "Service check date" /config/bigip.conf
C. grep "Service check date" /config/svc_chk_date.dat
D. grep "Service check date" /config/BigDB.dat
Answer: (SHOW ANSWER)
BIG-IP licensing information, including theService Check Date, is stored in the file:
/config/bigip.license
This file contains all license attributes downloaded from the F5 licensing server, including:
* License key
* Licensed modules
* Useful life date
* Service check date
TheService Check Datedetermines whether the system is eligible for upgrades to specific TMOS versions.
When reviewing upgrade readiness, administrators extract this value directly from the license file with:
grep "Service check date" /config/bigip.license
Why the other options are incorrect:
* /config/bigip.confstores BIG-IP configuration objects, not license metadata.
* /config/svc_chk_date.datisnota valid file in the licensing system; it does not contain license parameters.
* /config/BigDB.datstores internal database values, not licensing attributes.
Thus, only thebigip.licensefile contains the correct licensing information required for verifying upgrade eligibility.

NEW QUESTION: 7
The BIG-IP Administrator received a ticket that an authorized user is attempting to connect to the Configuration Utility from a jump host and is being denied.
The HTTPD allow list is configured as:

sys httpd {
allow { 172.28.31.0/255.255.255.0 172.28.65.0/255.255.255.0 }
}
The jump host IP is172.28.32.22.
What command should the BIG-IP Administrator use to allow HTTPD access for this jump host?
A. modify /sys httpd allow replace-all-with { 172.28.32.22 }
B. modify /sys httpd allow delete { 172.28.31.0/255.255.255.0 172.28.65.0/255.255.255.0 }
C. modify /sys httpd allow add { 172.28.32.22 }
Answer: C (LEAVE A REPLY)
The HTTPD allow list controls which IP addresses or subnets may access the Configuration Utility (TMUI) on the BIG-IP system. The Administrator already has two subnets allowed and needs to add asingle host IPto the existing list.
* The object/sys httpd allowsupports actions such asadd,delete, andreplace-all-with.
* Because the goal is toaddone more entry without removing the existing permitted subnets, the correct command is:
modify /sys httpd allow add { 172.28.32.22 }
This appends the new host to the existing list while preserving the previously configured networks.
Why the other options are incorrect:
* Option A (replace-all-with)wouldoverwritethe entire allow list, removing existing permitted subnets- unacceptable.
* Option B (delete)wouldremovethe existing networks and not add the required host.
Therefore, the correct administrative action is toaddthe jump host's IP.

NEW QUESTION: 8
For an upgrade of a standalone BIG-IP, a maintenance window is available in which brief interruptions are allowed.
Actions with no impact can be done outside the maintenance window.
When should a license reactivation be performed?
A. During the maintenance window.
B. Before the maintenance window.
C. After the maintenance window.
Answer: B (LEAVE A REPLY)
License reactivation updates the BIG-IP device's license file to ensure:
* TheService Check Dateis current
* The device is eligible to install the intended TMOS version
* Any module entitlement updates are received
Reactivationdoes not interrupt trafficand does not require a reboot, making it safe to performbeforethe maintenance window.
F5 best practices state:

* Performall non-impact tasks priorto the scheduled maintenance window
* Leave the window available for activities that require rebooting, such as the software installation itself Since license reactivation isnon-disruptive, it should be donebeforethe upgrade window starts.

NEW QUESTION: 9
When is theLicense Service Check Dateenforced on a BIG-IP system?
A. After editing a virtual server
B. During a software install
C. During system startup
Answer: B (LEAVE A REPLY)
TheService Check Datedetermines whether a particular software version is allowed to run under the device's license.
* When installing or upgrading TMOS, the installer checks theService Check Datestored in the BIG-IP license file.
* If the license date isolderthan the minimum required for the target version, the software installation is blocked.
* This check happensspecifically during a software install, not during routine device operations.
Editing virtual servers or system startup do not trigger this validation.
Thus, the enforcement happensduring software installation.

NEW QUESTION: 10
A BIG-IP device is licensed forLTM, ASM, APM, and AFM.
Currently, it will only be used forload balancingandweb application firewalling.
To ensure optimal performance and efficient resource utilization, which of the following module provisioning combinations is the best choice?
A. LTM: Dedicated
ASM: Dedicated
APM: Minimal
AFM: Minimal
B. LTM: Dedicated
ASM: Dedicated
APM: None
AFM: None
C. LTM: Nominal
ASM: Nominal
APM: None
AFM: None
D. LTM: Nominal
ASM: Nominal

APM: Minimal
AFM: Minimal
Answer: C (LEAVE A REPLY)
BIG-IP provisioning determines how CPU, memory, and disk resources are allocated to each module. The goal is to provision only the modules required and at levels appropriate to their performance needs.
Requirements in the question
The device will be used for:
* LTM(Local Traffic Manager) # load balancing
* ASM(Application Security Manager) # WAF
No functions require:
* APM (Access Policy Manager)
* AFM (Advanced Firewall Manager)
Why Option C is correct
Provisioning bothLTMandASMatNominallevel provides:
* Adequate performance for production load
* Plentiful system resources while avoiding dedicating the entire system to a single module
* Balanced allocation without starving memory or CPU
SettingAPM: NoneandAFM: Noneensures unused modules consume zero resources.
Why the other options are incorrect
A). Dedicated provisioning for both LTM and ASM
* Two modules cannot both run in "Dedicated" mode.
* Dedicated mode allocatesallresources to a single module - the second module cannot be dedicated simultaneously.
B). LTM and ASM both Dedicated
* Same issue: only one module can be Dedicated at a time.
* Also unnecessary for load balancing + WAF.
D). Setting APM and AFM to Minimal
* Minimal still consumes memory and CPU.
* Unused modules should be set toNone.
Therefore,Option Cis the best provisioning strategy.

NEW QUESTION: 11
A BIG-IP Administrator upgrades the BIG-IP LTM to a newer software version. After the administrator reboots into the new volume, the configuration fails to load.
Why is the configuration failing to load?
A. The upgrade was performed on the standby unit.
B. The license needed to be reactivated before the upgrade.
C. A minimum of at least two reboots is required.
D. Connectivity to the DNS server failed to be established.
Answer: B (LEAVE A REPLY)

When upgrading to a newer TMOS software version, BIG-IP validates whether the current license is permitted to run that version.
This is controlled by theService Check Datein the device's license file.
If the Service Check Date is older than the minimum required for the target version:
* The systemboots into the new volume,
* Butfails to load the configuration,
* And will instead present messages indicating that the configuration cannot be applied due to aninvalid or outdated license.
This is a well-known behavior:
An outdated license, not reactivated before upgrade, causes configuration load failure after reboot into the new software.
Why the other options are incorrect:
A). Performed on the standby unit
* Upgrading a standby unit does not cause configuration load failure.
* Standby-only upgrades are standard best practice.
C). Two reboots required
* BIG-IP does not require two reboots during an upgrade.
* One reboot into the new volume is sufficient.
D). DNS connectivity failure
* DNS connectivity does not affect configuration loading.
* DNS is only needed for automatic license activation, not for applying config at boot.
Thus, the configuration failed to load because thelicense was not reactivated before the upgrade, making Option Bcorrect.

NEW QUESTION: 12
The monitoring team reports that the SNMP server is unable to poll data from a BIG-IP device.

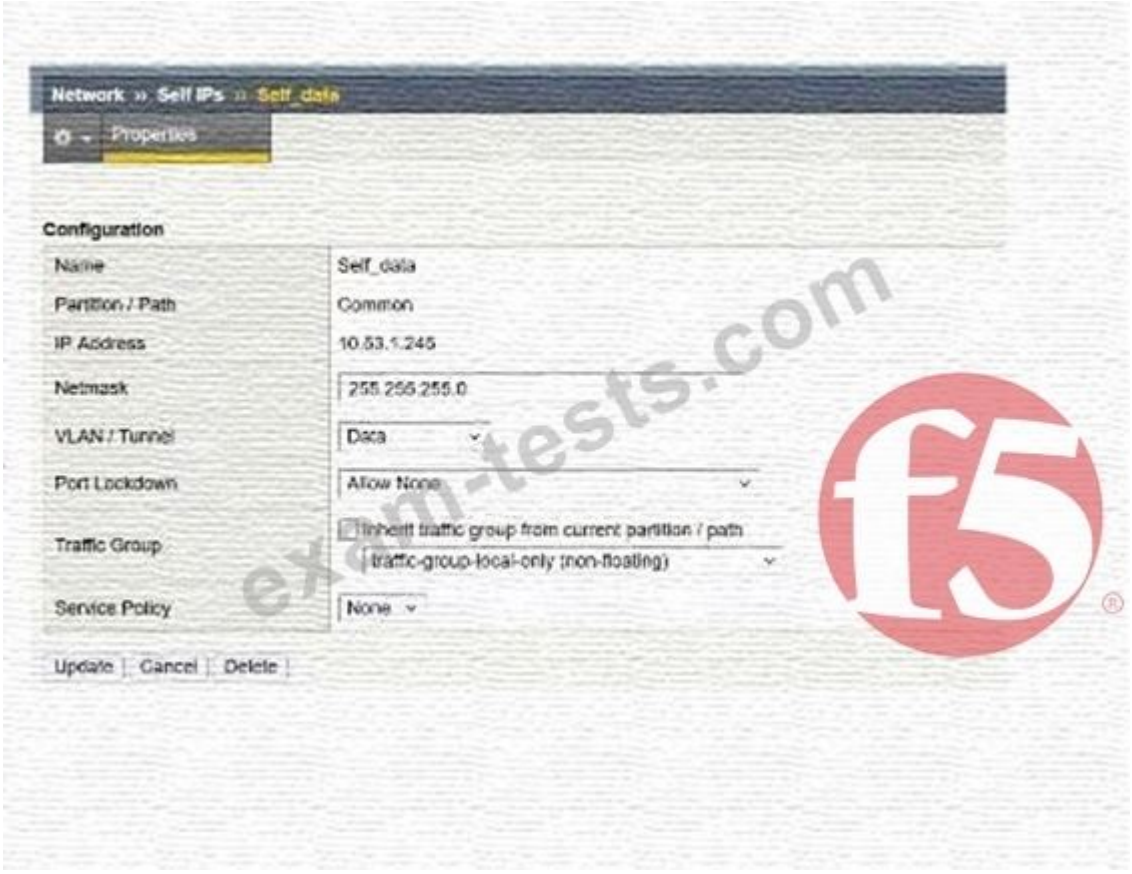

What information will help the BIG-IP Administrator determine whether the issue originates from the BIG-IP system?

A. The "Port Lockdown" setting is preventing the SNMP server from polling data from the BIG-IP.

B. The "Traffic Group" setting must use a floating Traffic Group.

C. The "VLAN / Tunnel" setting must allow All Vlans.

D. The configuration on the exhibit is correct and other options should be explored.

Answer: (SHOW ANSWER)

The exhibit shows aSelf IPwith:

* VLAN:Data

* Port Lockdown:Allow None

Impact of "Allow None" on SNMP

When a Self IP is configured with:

Port Lockdown: Allow None

the BIG-IP blocksallservices and ports except a few hardcoded HA communication ports.

This means:

* UDP/161 (SNMP)is blocked

* UDP/162 (SNMP traps)is blocked

* The SNMP server cannot poll or receive data from the BIG-IP through this Self IP SNMP relies on access through the Self IP if out-of-band (mgmt interface) is not used.

Thus, the issue is directly caused byPort Lockdown = Allow None, which prevents SNMP communication.

Why the other options are incorrect:

B). Traffic Group must use a floating Traffic Group
* SNMP polling doesnotrequire floating Self IPs.
* Floating groups apply to HA failover IPs, not SNMP functionality.
C). VLAN/Tunnel must allow All VLANs
* Self IPs are always bound to a VLAN; SNMP doesnotrequire All VLANs.
* As long as the Self IP belongs to a reachable VLAN, SNMP can work.
D). Configuration is correct
* It is not correct:Allow Noneblocks SNMP and is the problem.

NEW QUESTION: 13
A BIG-IP Administrator is responsible for deploying a new software image on an F5 BIG-IP HA pair and has scheduled a one-hour maintenance window.
With a focus on minimizing service disruption, which of the following strategies is the most appropriate?
A. Update the active node first, reboot to the newly updated boot location and verify functionality, then push the update from the active to the standby node and reboot the standby node.
B. Reset the Device Trust, apply the update to each node separately, reboot both nodes, then re-establish the Device Trust.
C. Update the standby node first and reboot it to the newly updated boot location, failover to the newly updated node and verify functionality. Repeat the upgrade procedures on the next node, which is now in standby mode.
D. Update both nodes in the HA pair, then reboot both nodes simultaneously to ensure they run the same software version.
Answer: (SHOW ANSWER)
For BIG-IP high-availability (HA) pairs, F5's recommended upgrade workflow prioritizesservice continuity, predictable failover, andminimal downtime. The established best-practice sequence is:
* Upgrade the standby unit first
* Because the standby device is not passing traffic, upgrading and rebooting it does not impact production.
* Boot the standby unit into the newly installed version
* Once online, the administrator verifies basic health, device sync status, cluster communication, and module functionality.
* Perform a controlled failover to the upgraded unit
* Traffic shifts to the newly upgraded device, allowing validation of the configuration and operational behavior under real traffic loads.
* Upgrade the second device (now standby)
* The previously active device becomes standby after failover, allowing it to be safely upgraded and rebooted without interruption.

This phased approach ensures only one device is unavailable at a time, allowing continuous traffic flow throughout the upgrade process.
Why the Correct Answer is C
OptionCexactly matches F5's documented production-safe upgrade method:
* Upgrade thestandbynode first
* Reboot into new image
* Failover to upgraded device
* Validate
* Upgrade the remaining (now-standby) device
This procedure minimizes risk and traffic disruption.
Why the other options are incorrect:
A). Upgrade the active node first
* Upgrading the active device requires removing it from service and failing over abruptly. This is not recommended and increases service disruption risk.
B). Resetting device trust
* Resetting trust is unnecessary and can disrupt configuration sync, peer communication, and cluster operation. It is not part of any standard upgrade workflow.
D). Upgrading and rebooting both nodes simultaneously
* This would causetotal outage, because both HA members would be unavailable at the same time.

NEW QUESTION: 14
Given thatBIGIP-<version>.isoandHotfix-BIGIP-<version>-ENG.isohave been uploaded to/shared/images on an F5 device, what is the appropriatetmshcommand to prepare and update the BIG-IP device with the hotfixof a software version on anew volume HD1.2? (Choose one.)
A. tmsh install /sys software hotfix Hotfix-BIGIP-<version>-ENG.iso create-volume HD1.2
B. tmsh install /sys software BIGIP-<version>.iso hotfix Hotfix-BIGIP-<version>-ENG.iso create-volume HD1.2
C. tmsh create /sys software hotfix Hotfix-BIGIP-<version>-ENG.iso volume HD1.2
D. tmsh copy /sys software hotfix Hotfix-BIGIP-<version>-ENG.iso volume HD1.2
Answer: B (LEAVE A REPLY)
When installing a BIG-IP software versionwith a HotFixon anew boot volume, F5 requires that both thebase TMOS imageand theHotFix imagebe installed together as part of the same installation workflow.
The correct process is:
* Specify thebase TMOS ISO
* Specify theHotFix ISOthat corresponds to that base version
* Instruct the system tocreate a new boot volume
* Install both images into that new volume
This is achieved with the following tmsh syntax:

tmsh install /sys software BIGIP-<version>.iso hotfix Hotfix-BIGIP-<version>-ENG.iso create-volume HD1.2 This command:
* Installs the base image first
* Applies the HotFix on top of the base image
* Creates and installs everything onHD1.2
* Leaves the currently active volume untouched for rollback
Why the other options are incorrect
A). Installing only the hotfix
A HotFix cannot be installed by itself on a new volume. A base image must already be present.
C). Using create instead of install
The create keyword is not valid for software installation operations.
D). Using copy
The copy command does not install software images or hotfixes.

NEW QUESTION: 15
What are the two options for securing a BIG-IP's management interface?
(Choose two.)
A. Limiting network access through the management interface to a trusted/secured network VLAN.
B. Block all management-interface administrative HTTPS and SSH service ports to prevent access.
C. Use the BIG-IP's Self-IP addresses for administrative access rather than the management interface.
D. Restrict administrative HTTPS and SSH access to specific IP addresses or IP ranges.
Answer: A,D (LEAVE A REPLY)
Securing the BIG-IP management interface is a fundamental administrative responsibility. F5 best practices emphasize restricting who can reach the management port and ensuring that only authorized systems are allowed access.
A). Limiting management access to trusted network segments
F5 recommends placing the management interface on adedicated, isolated, and secured management network or VLAN, rather than exposing it to production or untrusted networks.
This reduces the attack surface by ensuring only trusted segments have visibility to administrative interfaces.
D). Restricting management access by IP or subnet
F5 BIG-IP uses the/sys httpd allowlist (for HTTPS) and configuration options insshd(for SSH) to control which IP addresses or subnets can access the device.
By specifying only known administrative IPs or ranges, unauthorized users cannot reach the login services.
Why the other options are incorrect

B). Blocking all management HTTPS/SSH ports
* This would prevent any administrative access and is not a viable security practice.
C). Using Self-IP addresses for administrative access
* F5 explicitly warns against using Self-IPs for management access unless strictly necessary.
* Self-IPs are exposed to the data plane and should not be used as the primary administrative interface.

NEW QUESTION: 16
The Port Lockdown feature prevents unwanted connection attempts to a Self IP.
Which three types of connection attempts areunaffectedby Port Lockdown settings?
A. Defined virtual server traffic, Secure Shell (SSH), Centralized Management Infrastructure (CMI)
B. Centralized Management Infrastructure (CMI), Secure Shell (SSH), Internet Control Message Protocol (ICMP)
C. Defined virtual server traffic, Internet Control Message Protocol (ICMP), Centralized Management Infrastructure (CMI)
Answer: C (LEAVE A REPLY)
Port Lockdown controls which ports and protocols aSelf IPwill respond to.
However, certain traffic types bypass Port Lockdown for BIG-IP functionality and routing integrity.
The three types that areNOT affectedby Port Lockdown are:
1. Defined Virtual Server Traffic
Traffic destined to a Self IP that matches aconfigured virtual serveris always accepted by the BIG-IP, regardless of Port Lockdown settings.
This ensures that traffic processing does not break when administrators restrict Self-IP ports.
2. ICMP (Internet Control Message Protocol)
ICMP (such as ping, traceroute responses, etc.) always passes through a Self IP even when Port Lockdown is set to:
* Allow Default
* Allow None
* Allow Custom
F5 allows ICMP for reachability and diagnostic purposes independent of Port Lockdown rules.
3. Centralized Management Infrastructure (CMI)
CMI includes the internal HA services used for:
* Device Trust
* ConfigSync
* Failover
* Mirroring

These essential HA communications bypass Port Lockdown to prevent accidental cluster failure.
The well-known port for this traffic isTCP 4353, which is always permitted.
Why the other options are incorrect:
Option A:SSHisrestricted by Port Lockdown unless explicitly allowed.
Option B:Same issue - SSH does not bypass Port Lockdown.
OnlyDefined VS Traffic,ICMP, andCMIbypass Port Lockdown.

NEW QUESTION: 17
An F5 VE has been deployed into a VMware environment via an OVF file.
An administrator wants to configure the management IP address so the VE can be accessed for further setup.
Which two are valid methods for configuring the management-ip address? (Choose two.)
A. Log into the remote console and configure the management IP by running theconfigexecutable.
B. Log into the remote console and configure the management IP through TMSH using: create sys management-ip <ip address>/<mask>
C. Log into the remote console and configure the management IP through TMSH using: create ltm management-ip <ip address>/<mask>
D. Log into the remote console and configure the management IP by running thesetupcommand.
Answer: A,B (LEAVE A REPLY)
A newly deployed BIG-IP Virtual Edition (VE) in VMware requires initial configuration of itsmanagement- ipaddress so it can be accessed over the network. F5 provides several valid mechanisms during initial console access:
A). Running the config utility
* The config script is available on new BIG-IP installations and VE deployments.
* It launches a guided text-based wizard allowing configuration of:
* Management IP
* Netmask
* Default route
* This is a standard and recommended method during first-time setup.

B). Using TMSH with create sys management-ip
* Administrators can enter TMSH directly from the console and run:
* create sys management-ip <ip>/<mask>
* The management-ip object resides undersys, not under ltm or any other module.
* This is the correct tmsh method for defining the management interface address.
Why the other options are incorrect:
C). create ltm management-ip
* There isnosuch object under /ltm.
* LTM handles traffic objects (virtual servers, pools), not system management interfaces.
D). Running the setup command
* The setup command is used for general system configuration butdoes not configure the management- ip.
* It is not the supported method for initial management IP assignment on VE deployments.
Therefore, the valid methods are running theconfigutility and using thesys management-ipcommand within TMSH.

NEW QUESTION: 18
An administrator is in the process of reactivating the license using the interface displayed in the exhibit.
What is the address of the license server to which the BIG-IP device must be able to establish an outbound connection in order to use theAutomatic Activation Method?
A. license.f5.com
B. callhome.f5.com
C. ask.f5.com
D. activate.f5.com
Answer: D (LEAVE A REPLY)
When you chooseAutomaticas the activation method in the License Re-activate screen, the BIG-IP device itself contacts F5'slicense activation serviceover the Internet.
For successful automatic activation:
* The BIG-IP must have outbound network connectivity (typically via the management interface).
* DNS resolution and routing must allow it to reach theF5 license activation host(the one shown in option D).
* The device sends its dossier and registration key to that service and receives an updated license file in return, which is then installed automatically.
The other hostnames in the options are not used by BIG-IP for license activation, so they cannot be correct in the context ofAutomatic Activation.

NEW QUESTION: 19
The BIG-IP Administrator wants to manage the newly built F5 system through anin-band Self-IP.

The administrator has configured a VLAN and Self-IP and can ping the IP from their workstation, but cannot access the system viaSSHorHTTPS.
Whatport lockdownsettings should the BIG-IP Administrator use to allow management access on the Self-IP?
(Choose two.)
A. The Self-IP port lockdown behavior could be adjusted toAllow Default
B. The Self-IP port lockdown behavior could be adjusted toAllow All
C. The Self-IP port lockdown behavior could be adjusted toAllow Mgmt
D. The Self-IP port lockdown behavior could be adjusted toAllow Management
Answer: (SHOW ANSWER)
Self-IPs include a security feature calledPort Lockdown, which restricts which services respond on that Self- IP.
By default, Self-IPs block management access (SSH and HTTPS/TMUI), meaning an administrator cannot manage the device through in-band Self-IPs unless explicitly allowed.
Allow Mgmt / Allow Management
These settings enable only the management services required for administrative access, specifically:
* SSH (22)
* HTTPS/TMUI (443)
These options allow secure administration without opening unnecessary ports.
Why these are correct:
* They provide only the essential access for management.
* They follow F5 security best practices when using in-band admin access.
* They donotexpose all services, reducing the attack surface.
Why the other options are incorrect:
A). Allow Default
* This allows only a minimal set of system-required ports (e.g., failover, config sync), not SSH or HTTPS.
* Administrator access would still fail.
B). Allow All
* Opens all ports on the Self-IP, which isnot secure.
* Exposes services that should remain restricted.
Therefore,Allow Mgmt / Allow Managementare the correct choices.