Lead2passExam

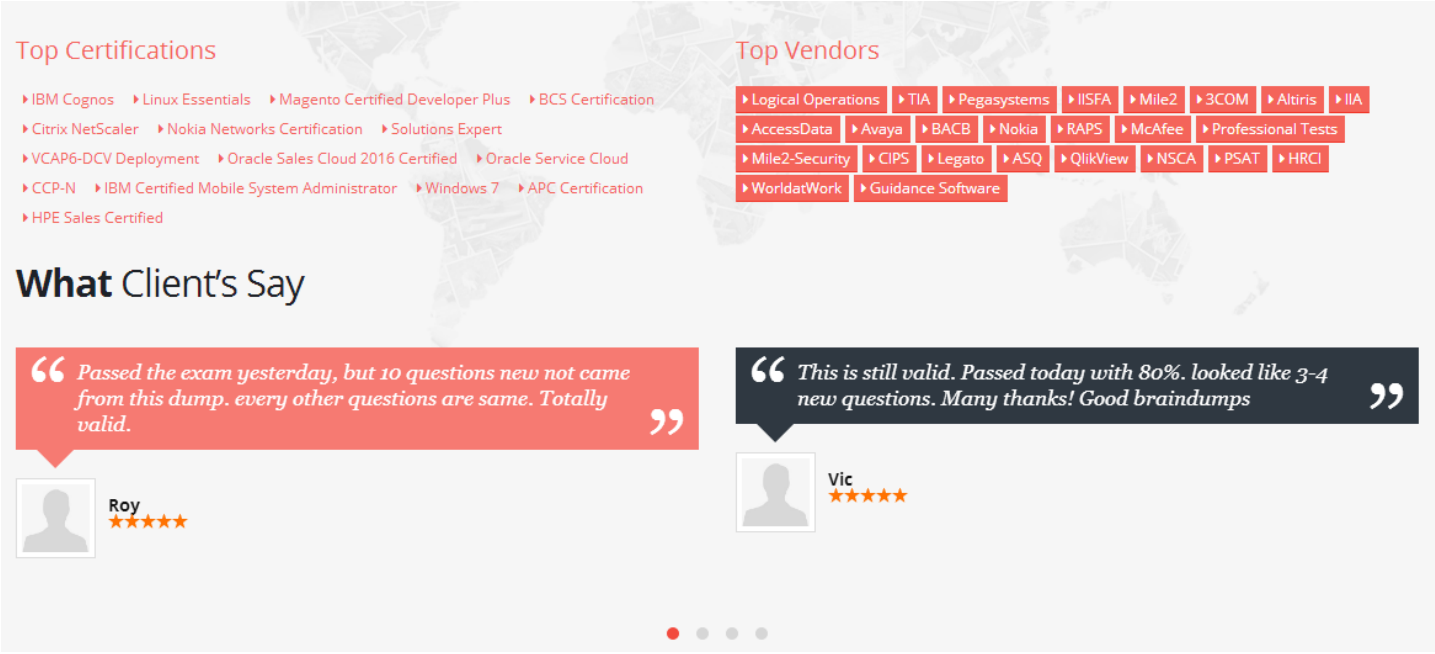

http://www.lead2passexam.com/

Available Exam Cram and Valid Dumps - Lead2Pass Exam

Exam : **FlashArray-Implementation-Specialist**

Title : Pure Storage Certified FlashArray Implementation Specialist

Vendor : Pure Storage

Version : DEMO

NO.1 Which I/O card type is compatible across all FlashArray models?

A. 10GBaseT

B. iSCSI

C. FC

Answer: C

Explanation:

The Fibre Channel (FC) I/O card type is the only option listed that is universally compatible and supported across all FlashArray models, from the entry-level //X10 and //X20 up to the high-end //X90 and //XL.

* Fibre Channel: Pure Storage architectures rely heavily on Fibre Channel as the primary enterprise storage protocol. All chassis generations and controller sizes feature PCIe risers compatible with Pure's standard 16Gb or 32Gb Fibre Channel HBAs.
* Ethernet/iSCSI: While all arrays support iSCSI, the physical card type varies.
* 10GBase-T: This refers to "copper" Ethernet (RJ45). This card type is not supported on all models. High-end FlashArrays (like the //X70, //X90, and //XL) typically utilize SFP+ or QSFP28 cages for optical connectivity (10/25GbE or 40/100GbE) and do not standardly support 10GBase-T copper cards due to power and latency characteristics.
* Therefore, while the protocol (iSCSI) is supported everywhere, the specific physical card (10GBase-T) is not. Fibre Channel cards remain the consistent hardware constant across the entire product line.

NO.2 A client needs to upgrade Purity on their FlashArray, but is unable to wait for an available upgrade date from support.

What action should the customer take?

A. Go through the Self-Service upgrade portal on Pure1 Manage and upgrade to the latest Purity version.

B. Open a ticket with support, download the latest Purity version, copy to a USB flash drive, and run the pureinstall command on each controller.

C. Log into the Purity CLI, download the latest Purity version, and run the pureinstall command on each controller.

D. Proceed with the upgrade by executing the purendu procedure start command via the local array GUI.

Answer: A

Explanation:

Pure Storage provides a robust, non-disruptive upgrade (NDU) architecture. Historically, PurityOS upgrades were strictly scheduled and executed entirely by Pure Storage Support to ensure maximum safety. However, modern iterations of the platform allow customers complete autonomy via the Self-Service Upgrade portal within Pure1 Manage.

If a client is on a tight timeline and cannot align with a support-scheduled maintenance window, they can securely initiate the upgrade themselves. The Pure1 portal orchestrates pre-flight checks (like running the internal hardware_check.py script and assessing cluster health) entirely in the background. Once the array's health is validated, the client can deploy the new Purity software directly from the cloud interface.

Customers should never attempt to manually run the internal pureinstall command via the CLI (Options B and C). That specific command circumvents critical automated safety checks and is reserved strictly for certified Implementation Engineers or Pure Support during offline KVM

installations or complex hardware replacements.

NO.3 A customer's management network connectivity is NOT available at the time of array deployment. Which switch should an Implementation Engineer use to avoid those checks during puresetup flow?

A. --skip-connectivity-tests

B. --skip-network-tests

C. --skip-management-connectivity-tests

Answer: C

Explanation:

The puresetup initialization script automatically runs a series of network validation tests to ensure the array can reach default gateways, DNS servers, and Pure1. If the management network is not yet live or is air- gapped during the install, these tests will fail, potentially blocking the setup process.

To proceed with the initialization despite the missing network, the Implementation Engineer must append the

--skip-management-connectivity-tests flag to the puresetup command.

* This flag instructs the script to bypass the ping/DNS lookup checks on the management interface.
* It allows the engineer to complete the array initialization (naming, IP assignment) so the array is "up" and ready for local config, with the understanding that full management connectivity will be established and verified later.

=========

NO.4 During an NDU, when is the Implementation Engineer required to install additional NVRAM modules in slots

2 and 3?

A. Upgrading a FlashArray//X50R3 to a FlashArray//X50R4

B. Upgrading a FlashArray//X50R3 to a FlashArray//X70R4

C. Upgrading a FlashArray//X70R4 to a FlashArray//X90R4

Answer: B

Explanation:

FlashArray models differ in their NVRAM requirements based on their performance tier and throughput capabilities. Lower and mid-range models (like the //X10, //X20, and //X50) typically operate with two NVRAM modules installed in slots 0 and 1. However, high-performance models (like the //X70 and //X90) require double the NVRAM buffer to handle the increased write I/O bandwidth, necessitating four NVRAM modules (populating slots 0, 1, 2, and 3).

When performing a Data-in-Place Upgrade (NDU) from a FlashArray//X50R3 to a FlashArray//X70R4 , the chassis remains (or is upgraded in a way that preserves data), but the new controller configuration shifts from a "2-NVRAM" requirement to a "4-NVRAM" requirement. Therefore, as part of the physical upgrade procedure, the Implementation Engineer must install two additional NVRAM modules into slots 2 and 3 of the chassis to support the new //X70 controllers.

If these modules are not added, the new //X70 controllers will detect an "insufficient hardware" condition and fail to boot or initialize the Purity software properly. This step is unique to upgrades crossing the boundary between mid-range and high-end models.

=========

NO.5 What slot does a PCIe SAS card need to be installed in a FlashArray//XR2/3?
A. 0
B. 1
C. 2
Answer: B
Explanation:
On a FlashArray//XR2 or //XR3 , if a dedicated PCIe SAS HBA (Host Bus Adapter) is required for external shelf connectivity (expanding beyond the onboard SAS ports), it must be installed in Slot 1 .
The PCIe slot prioritization on the FlashArray//X R2/R3 controllers is strictly defined to ensure optimal performance and lane distribution:
* Slot 0: Reserved primarily for frontend Host I/O cards (Fibre Channel or iSCSI/Ethernet). This slot usually connects to CPU 0 and offers high bandwidth for host traffic.
* Slot 1: Designated as the primary expansion slot for backend connectivity (SAS cards for expansion shelves) or secondary Host I/O if no SAS expansion is needed.
* Slot 2: On smaller models (X10/X20), this is often the only host I/O slot, but on the larger X50/X70 /X90 chassis, it is used for additional Host I/O.
If an Implementation Engineer installs a SAS card in Slot 0, the system may fail to recognize the shelves correctly or may flag a configuration error, as the Purity software expects backend expansion traffic to route through the bus associated with Slot 1. Therefore, for SAS expansion cards, Slot 1 is the mandatory location.

NO.6 What is the required PSU wattage for a FlashArray// X70R2/R3 or X90R2/R3 array?
A. 1200W
B. 1300W
C. 1600W
Answer: C
Explanation:
The FlashArray//X70 (R2/R3) and FlashArray//X90 (R2/R3) require 1600W Power Supply Units (PSUs) .
Pure Storage equips its FlashArray//X chassis with different power supply capacities based on the compute load of the installed controllers:
* Performance Tier (X70/X90): These models utilize high-core-count Intel Xeon processors and support the maximum density of NVMe DirectFlash Modules. To support the peak power draw of these components while maintaining N+1 redundancy (where one PSU can support the entire load), the
1600W units are mandatory.
* Entry/Mid Tier (X10/X20/X50): Lower-end models often ship with 1000W or 1200W PSUs (though X50 often uses 1600W in later revisions or specific configs).
* Identification: Implementation Engineers can identify these PSUs by the label on the rear handle (often color-coded or explicitly marked "1600W"). Installing an X90 controller into a chassis with older
1200W PSUs (e.g., during an improper upgrade attempt) would likely trigger a hardware alert or prevent the controllers from booting due to insufficient power budget.

NO.7 An Implementation Engineer is going on site to complete a data pack evacuation and perform a

hardware NDU. At what step should the Implementation Engineer contact support for assistance?

A. After removing the fully evacuated drives before proceeding with the hardware NDU

B. After removing the Old drives and inserting any new ones

C. When the hardware NDU has been completed

Answer: A

Explanation:

Complex upgrades that involve both removing old capacity (evacuation) and replacing controllers (hardware NDU) require a "checkpoint" with Pure Storage Support to ensure data safety.

The correct procedure dictates contacting support after removing the fully evacuated drives but before proceeding with the hardware NDU .

* Why? Once the old drives are physically removed, the array's configuration state has changed. Support needs to verify that the evacuation was 100% successful, that no stale objects or pointers to the removed drives remain in the database, and that the system is fully healthy and redundant.
* Proceeding directly into a controller reboot (NDU) without this verification could lead to issues if the array still "thinks" it needs the removed drives, potentially causing a boot failure or data unavailability during the controller failover.

NO.8 An Implementation Engineer is scheduled to install a FlashArray and a DirectFlash Shelf (DFS). What action should the Implementation Engineer perform?

A. Refer to the "Recommended DFS Versions for Purity Releases" KB and upgrade Purity on the FlashArray and DFS.

B. Install the Array and DFS with the Purity that shipped and have Support upgrade remotely.

C. Only upgrade the Purity on the FlashArray and then attach the DFS.

Answer: A

Explanation:

The correct action is to refer to the "Recommended DFS Versions for Purity Releases" Knowledge Base (KB) article and, if necessary, upgrade the Purity software on the FlashArray (which updates the DFS firmware).

* Compatibility: DirectFlash Shelves (DFS) run specific firmware that must be compatible with the Purity operating system running on the array controllers. If a new shelf is shipped with very recent firmware, or an old shelf is added to a new array, there may be a mismatch.
* Best Practice: The "Recommended DFS Versions" KB is the definitive source of truth. The engineer must verify if the shipped Purity version supports the specific hardware revision of the shelf.
* Procedure: If an upgrade is required, the engineer updates Purity on the array. Since the DFS firmware is bundled within the Purity release, upgrading the array (Option A) effectively stages the update for the shelf. Once the shelf is attached (or if attached during the upgrade), Purity automatically applies the correct firmware to the DFS modules. Skipping this check (Option B) or blindly upgrading without cross-referencing the KB (Option C) risks connectivity issues or unrecognized capacity.

NO.9 When adding a data pack into a FlashArray chassis or a DFS shelf, the drives show "unrecognized" status when running List. What is the cause of this status on the newly inserted drives?

A. There is a hardware fault that needs to be corrected.

B. The drives exceed the max space supported for the controller model.

C. The Purity version of the FlashArray/DFS is incompatible with the new drives.

Answer: C

Explanation:

The status "unrecognized" on a drive module typically indicates that the physical hardware is detected by the chassis midplane, but the operating system (Purity) does not have the necessary driver definition or firmware logic to interact with it.

This most commonly occurs when newer generation DirectFlash Modules (e.g., larger capacity drives or newer NAND revisions) are inserted into an array running an older version of Purity. The hardware identifier of the new drive is unknown to the older software .

To resolve this, the Implementation Engineer must verify the Purity Compatibility Matrix for the specific drive part number. The array will likely need a Purity software upgrade to a version that supports the new media before the drives can be successfully admitted and used. It is rarely a hardware fault (which would show as "failed" or "missing") or a capacity limit (which would prevent admission but usually identify the drive correctly).

=========

NO.10 Which ports are used by default for replication on a FlashArray//XR4?

A. ETH2 and ETH3

B. ETH3 and ETH4

C. ETH0 and ETH1

Answer: A

Explanation:

The FlashArray//XR4 maintains the standard port assignment convention established in previous generations for its onboard Ethernet interfaces.

ETH2 and ETH3 are the default ports designated for Replication traffic.

* ETH0 and ETH1 are reserved for Management.

* ETH2 and ETH3 are configured by Purity defaults to handle the heavy bandwidth of asynchronous or synchronous replication.

* While these assignments can be modified in software, the physical ports labeled eth2 and eth3 on the rear of the chassis are the intended primary interfaces for this function. Implementation Engineers should cable these ports to the replication network switches during the initial install.

=========

NO.11 On a FlashArray running Purity 6.9.2+, what is the recommended way to initiate a planned controller failover during a FlashArray//XR5 HWNDU?

A. purewes controller setattr --verify-array < host-name > ct0 --mode secondary

B. puredb giveback --safe

C. purendu procedure start giveback

D. purehw controller failover

Answer: A

Explanation:

During a Pure Storage FlashArray Hardware Non-Disruptive Upgrade (HWNDU), controllers must be seamlessly transitioned to ensure continuous host I/O. The correct procedure to execute a planned controller failover on an array running Purity 6.9.2+ involves using the purewes command suite. "WES" (Wessex) represents the internal architecture and toolset utilized by Pure Storage engineering

and support to manipulate high-availability states and controller roles directly.
While standard array management utilizes purecli, HWNDUs require specific, low-level overrides to safely manipulate primary and secondary designations without disrupting active multipathing. To transition the current primary controller (e.g., CT0) to a secondary role, the Implementation Engineer must execute purewes controller setattr --verify-array < host-name > ct0 --mode secondary.
This command initiates a graceful giveback of all backend drive resources and frontend host I/O responsibilities to the partner controller (CT1). Once CT1 fully assumes the primary role, CT0 enters a ready secondary state, making it safe to power down, un-cable, and physically replace. The --verify-array flag ensures that the engineer is targeting the correct array by explicitly requiring the hostname, acting as a critical safeguard against accidental failovers on the wrong system during multi-array engagements. Options B, C, and D introduce fictitious command variations that do not exist within the Purity CLI framework for this specific hardware replacement step.

NO.12 An Implementation Engineer has installed a data pack and the puredrive list command shows the drives in
"unadmitted" status.
Which command should the Implementation Engineer run to complete the admission?
A. puredrive admit
B. puredrive add
C. purearray drive admit
D. purehw drive admit
Answer: A
Explanation:
Capacity expansion on a Pure Storage FlashArray is a highly controlled, safe process. When an Implementation Engineer physically unboxes and inserts a new data pack (a group of DirectFlash Modules) into the available drive bays of a chassis or a DirectFlash Shelf, the Purity operating system detects the hardware instantly. However, it does not automatically wipe the drives and assimilate them into the global storage pool.
Instead, the newly inserted modules are placed into a safe, quarantined state known as "unadmitted." This intentional design choice prevents catastrophic data loss in the event that an engineer accidentally inserts a drive containing live data from another system, or inserts drives before the customer is financially or operationally ready to activate the new capacity.
To officially claim the drives, integrate them into the system's Wide Write Groups (WWGs), and initiate the background parity redistribution process, the Implementation Engineer must explicitly authorize their use.
This is accomplished by running the puredrive admit command. The engineer can either specify the exact drives (e.g., puredrive admit CH0.BAY10 CH0.BAY11...) or use a global flag like puredrive admit --all if all unadmitted drives are ready for ingestion. Once admitted, the drives transition to a "healthy" state, and the array's total usable capacity increases seamlessly without any host disruption.

NO.13 A client needs to upgrade Purity on their FlashArray, but is unable to wait for an available upgrade date from support. What action should the customer take?
A. Open a ticket with support, download the latest Purity version, copy to a USB flash drive, and run the pureinstall command on each controller.

B. Log into the Purity CLI, download the latest Purity version, and run the pureinstall command on each controller.

C. Go through the Self-Service upgrade portal on Pure1 Manage and upgrade to the latest Purity version.

Answer: C

Explanation:

The customer should proceed via the Self-Service Upgrade (SSU) portal available on Pure1 Manage . Pure Storage has transitioned significantly toward a self-service model for routine software upgrades to provide customers with greater flexibility and immediacy.

* Pure1 Manage: This cloud-based management platform includes a "Software Lifecycle" or "Upgrade" widget. It automatically checks the array's health, compatibility, and readiness for various Purity versions.
* Process: The customer can select the target Purity version directly in the portal. Pure1 then orchestrates the "Pre-Flight Checks" (Remote Audit) to ensure the array is healthy. Once passed, the Purity software is downloaded directly to the array, and the customer (or the engineer) can trigger the upgrade at their convenience without needing to schedule a specific window with a live Support engineer.
* Invalid Options: Options A and B describe legacy or unsupported workflows. Customers do not manually download Purity ISOs to USB drives or run raw pureinstall CLI commands manually; these methods lack the automated safety checks and cloud-based verification provided by the SSU ecosystem.

NO.14 How many chassis capacity drive slots come standard on a FlashArray//X Gen 2 chassis?

A. 28 slots

B. 20 slots

C. 24 slots

D. 40 slots

Answer: B

Explanation:

Understanding the physical chassis layout of the Pure Storage FlashArray is essential for capacity planning and hardware installation. The FlashArray//X Gen 2 (and Gen 3) utilizes a highly efficient 3U chassis design.

When looking at the front of a standard //X chassis, there are a total of 24 physical module bays. However, these bays serve fundamentally different architectural purposes:

* Capacity Drive Slots: There are exactly 20 slots strictly dedicated to capacity drives (DirectFlash Modules, or DFMs). These are numbered 0 through 19 and are where the data packs are installed to provide raw storage capacity.
* NVRAM Slots: The remaining 4 slots, located centrally in the chassis, are dedicated exclusively to NVRAM modules (NVB0, NVB1, NVB2, and NVB3). These provide the highly available, non-volatile write cache for the array and cannot be used for capacity expansion.

Therefore, when an Implementation Engineer is assessing the chassis for data pack expansion or module upgrades, there are exactly 20 standard capacity drive slots available before an external DirectFlash Shelf (DFS) must be added. (Note: FlashArray//XL utilizes a larger 5U chassis with 40 capacity slots, but the standard //X remains at 20).

NO.15 Which PCI cards must be validated on the Bill of Materials (BOM) to allow an Implementation Engineer to attach a DirectFlash Shelf to a FlashArray//XL130R5?

A. FA-XL-100G Ethernet/RoCE 2-Port CX6

B. FA-XL-100G iSCSI/RoCE 2-Port CONNECTX-5

C. FA-XL-200G Ethernet/RoCE 2-Port CX7

D. FA-XL-25G Ethernet/RoCE 4-Port CX6

Answer: A

Explanation:

When expanding a Pure Storage FlashArray//XL130 R5 with an external DirectFlash Shelf (DFS), the backend connection between the shelf controllers and the main chassis relies on massive 100GbE bandwidth via the NVMe-oF (RoCE) protocol. Validating the Bill of Materials (BOM) before an onsite installation is a critical preparatory step to ensure the correct internal hardware was shipped to support this connection.

According to the official Pure Storage hardware matrix, the supported PCIe expansion card required for //XL series shelf connectivity is the FA-XL-100G Ethernet/RoCE 2-Port CX6 . This card utilizes the modern Mellanox/Nvidia ConnectX-6 (CX6) chipset, which is explicitly engineered to provide the necessary 100GbE RDMA capabilities to maintain ultra-low latency, NVMe-level speeds across the external storage fabric.

Option B (the CONNECTX-5 card) is an older, End-of-Sale component generally associated with legacy deployments or specific frontend iSCSI requirements, not the modern //XL backend loop. Option C (a 200G CX7 card) is not the standard validated BOM component for a single DFS loop on an //XL130 R5. Confirming the presence of the 100G CX6 card ensures the engineer has the exact hardware required to successfully cable the 100GbE backend ports.

NO.16 During an intra-series upgrade from //XR3 to //XR4, which command should the Implementation Engineer run to verify host connectivity before removing the secondary controller?

A. purehost list --host | less

B. purenetwork monitor --live

C. iobalance --sampletime 30

D. host_connectivity --hosts

Answer: C

NO.17 What are the starter configurations of DFMDs in a Gen 2 DirectFlash Shelf?

A. 8-slot and 14-slot

B. 10-slot and 14-slot

C. 12-slot and 14-slot

Answer: C

NO.18 An Implementation Engineer is getting ready to install four DFMs into a 16-module efficiency bundle and notices that the new modules are DFMDs while the installed modules are DFMs.

What should the Implementation Engineer do to get the array using a 20-module wide write group?

A. Install the four new DFMDs, run puredrive admit and have support set the tunable puretune --set allow_dfmd_in_dfm_wwg 1 < CASE > .

B. Work with the account team to ship DFMs instead of the DFMDs.

C. Install the four new DFMDs and run puredrive admit.
D. Insert the DFMDs into a separate swing shelf to isolate them from the DFMs.
Answer: B
Explanation:
Understanding the strict hardware boundaries of Pure Storage capacity expansion is critical. DirectFlash Modules (DFMs) and DirectFlash Modules with Distributed NVRAM (DFMDs) are fundamentally different hardware components. Standard DFMs contain only raw NAND flash and rely entirely on centralized, dedicated NVRAM modules located in the array chassis to protect write cache. DFMDs, however, feature their own onboard non-volatile memory to distribute the cache load across the capacity drives (primarily utilized in //XL and //XR4 architectures).
Pure Storage architecture explicitly dictates that drives operating within a single efficiency bundle or Wide Write Group (WWG) must be of the exact same hardware classification to ensure predictable parity calculation, wear leveling, and cache distribution. Mixing legacy DFMs with newer DFMDs within the same
20-module write group creates a severe hardware imbalance and is a strictly unsupported configuration.
While internal engineering tunables (like those suggested in Option A) occasionally exist to force overrides in lab environments, applying them in a customer's production array to bypass physical hardware incompatibility violates all implementation best practices. The only correct and safe action for the Implementation Engineer is to halt the physical installation, flag the Bill of Materials (BOM) mismatch, and work with the account
/logistics team to RMA the DFMDs and ship the correct standard DFMs.

NO.19 An Implementation Engineer has completed cabling a DirectFlash Shelf (DFS) to the FlashArray. Which command should the Implementation Engineer run to verify that the cabling was completed correctly?
A. storage_view.py enclosures
B. storage_view.py list
C. storage_view.py config
Answer: A
Explanation:
Verifying the physical SAS or RoCE cabling topology is a mandatory step before initializing a new shelf. The FlashArray provides a specific engineering-level Python script for this purpose: storage_view.py enclosures .
This script queries the backend fabric and visualizes the connection topology between the controllers and the shelves.
* It confirms that both controllers (CT0 and CT1) have redundant paths to the shelf (IOM0 and IOM1).
* It verifies that the shelf is visible on the expected ports.
* It checks for cabling errors, such as "crossed" cables (connecting a primary port to a secondary port incorrectly) or missing links.
Using this tool ensures that the physical layer is perfectly redundant before data is ever placed on the shelf.
=========

NO.20 Which Pure1 app functionality requires that a user is on their internal company network?

A. View array telemetry
B. Manage support cases
C. Open Remote Assist
Answer: C
Explanation:
The Pure1 mobile app provides convenient access to array monitoring and management features from anywhere. Most features, such as viewing performance telemetry (Option A) or managing support tickets (Option B), are cloud-native and accessible over the public internet via the secure Pure1 connection.
However, Opening Remote Assist (RA) -the feature that creates a secure tunnel for Pure Storage Support to access the array remotely-is a privileged security action. To prevent unauthorized external actors from enabling this access, the Pure1 app often enforces a "proximity" or network validation check. The user must be connected to the internal company network (e.g., via Wi-Fi or VPN) that has visibility to the array's management interface to authorize the Remote Assist session. This requirement ensures that only an authorized administrator physically or logically present within the customer's secure environment can grant external access to the storage system.
=========

NO.21 After an Implementation Engineer installs a new FlashArray and runs health checks, they see that Phonehome and Remote Assist are NOT working correctly. Which outbound port needs to be open for all management ports (physical and virtual)?
A. 80
B. 443
C. 123
Answer: B
Explanation:
For PhoneHome (telemetry) and Remote Assist (support tunnels) to function correctly, TCP Port 443 (HTTPS) must be open for outbound traffic from all management ports (CT0.ETH0, CT1.ETH0, and the Virtual Management IP).
Pure Storage utilizes a secure, encrypted connection to communicate with the Pure1 Cloud and Support infrastructure.
* PhoneHome: The array sends logs and heartbeat data to purestorage.com (and associated subdomains) via HTTPS over Port 443. If this port is blocked, the array cannot report its health status, preventing proactive support.
* Remote Assist: This feature creates a reverse SSH tunnel encapsulated over HTTPS (websocket or similar secure transport) to allow Pure Support to remotely log in. It initiates the connection outbound to the support relay servers on Port 443.
Option A (Port 80) is generally not used for secure management traffic in modern configurations, often only serving as a redirect. Option C (Port 123) is required for NTP (Network Time Protocol) synchronization, which is critical for logs but is not the transport port for the PhoneHome/Remote Assist data stream itself.
Therefore, verifying firewall rules for Port 443 is the primary troubleshooting step.

NO.22 What are the default replication ports on all three rack unit (3U) FlashArray models?
A. ETH0 and ETH1

B. ETH2 and ETH3
C. ETH4 and ETH5
D. ETH6 and ETH7
Answer: B
Explanation:
Understanding the default network port assignments on a standard 3U Pure Storage FlashArray (which encompasses models like the //X10, //X20, //X50, //X70, and //X90) is a fundamental requirement for any Implementation Engineer. The rear of the controllers features a standardized set of integrated Ethernet ports that are logically predefined by the Purity operating system upon factory initialization.
The first two ports, explicitly labeled eth0 and eth1 , are 1GbE Base-T copper ports. These are strictly reserved and dedicated by default for management traffic (GUI, CLI, SSH, API, and Call Home telemetry to Pure1). They are not designed to handle high-throughput storage workloads.
The next two ports, eth2 and eth3 , are high-speed 10GbE/25GbE optical/Twinax ports. By default, Purity provisions these specifically for asynchronous and synchronous replication traffic (including ActiveCluster), as well as IP-based block storage protocols like iSCSI if the customer is not using Fibre Channel. Utilizing eth2 and eth3 ensures that heavy data mobility workloads have the required bandwidth and are physically isolated from lightweight management traffic. Ports like eth4 and eth5 (if present via PCIe expansion cards) are typically configured manually for additional frontend host I/O or File Services, rather than acting as the default replication interfaces.

NO.23 After running puredb run giveback --safe command once CT0 has been upgraded, how long should the Implementation Engineer wait before continuing with the controller upgrade?
A. 3 minutes
B. 5 minutes
C. 10 minutes
Answer: C
Explanation:
After executing the puredb run giveback --safe command during a manual controller upgrade or specific non- disruptive upgrade (NDU) procedure, the Implementation Engineer must wait for 10 minutes before proceeding to the next step (typically upgrading the peer controller).
The giveback command instructs the recently upgraded controller (e.g., CT0) to resume its primary role and take back ownership of its assigned storage resources and I/O handling from the peer controller. This process triggers a failover event for the host multipathing software.
* Path Stabilization: The 10-minute wait time is a critical safety buffer mandated to ensure that all host- side multipathing drivers (MPIO) have successfully recognized the restored paths to the upgraded controller and have stabilized.
* Risk Mitigation: Proceeding too quickly-before the hosts have fully settled and paths are marked as "Active/Optimized"-could lead to an "all paths down" (APD) scenario when the second controller is taken offline for its upgrade. This wait time verifies that the upgraded node is fully functional and carrying load, adhering to the strict "availability first" philosophy of FlashArray upgrades.

NO.24 On a FlashArray//XR4, which port is enabled for the management service by default?
A. ETH0
B. ETH2

C. ETH4
D. ETH6
Answer: A
Explanation:
When configuring the network topology for a new Pure Storage FlashArray//XR4 deployment, the physical Ethernet ports on the rear of the controllers are rigidly mapped to default services by the Purity//FA operating system. This standardized mapping ensures that Implementation Engineers can reliably configure out-of-band management access before any advanced routing or high-speed data fabrics are established.
Across the entire FlashArray product line-including the modern //XR4 and //CR4 architectures-the first onboard 1GbE Base-T copper port, explicitly labeled eth0 , is always enabled and dedicated strictly to the management service by default.
This interface (eth0, along with its highly available partner eth1) is responsible for all non-disruptive administrative access. It handles the web-based GUI traffic (HTTPS), SSH terminal sessions, REST API calls, and the secure outbound tunnel to the Pure1 cloud for Call Home telemetry and Remote Assist. High-speed ports like eth2, eth4, or eth6 (which utilize SFP28 or QSFP optical transceivers) are defaulted to Replication or left unconfigured for frontend host I/O (iSCSI/NVMe-oF). An Implementation Engineer relies on eth0 being the default management port to connect their laptop during the initial serial console setup (puresetup) and validate network reachability.