Dell EMC

D-PE-OE-01

Dell PowerEdge Operate v2

Questions & Answers (Demo)

Version: 4.0

Question: 1

A Dell PowerEdge server has Temperature Alert Configuration set up in iDRAC. A Test Event was executed, and email alerts appeared to be delivered successfully. However, a few days later, the server unexpectedly rebooted. Upon inspection, CPU temperatures had exceeded the configured threshold. What is a possible reason IDRAC failed to alert the administrator about the high temperature?

A. iDRAC only sends alerts for Critical events or higher.

B. Although the Temperature Alert was configured, it was not fully activated.

C. Test Events verify alert delivery but do not trigger actual hardware alerts.

D. The reboot may have bypassed iDRAC alert mechanisms.

Answer: B

Explanation:

In the Integrated Dell Remote Access Controller (iDRAC) architecture, configuring a destination and executing a successful Test Event merely validates the underlying notification network pathway and the SMTP handshake with the mail server. It does not mean that automated notifications are active for operational system events. To achieve full activation, an administrator must explicitly enable the desired alert actions within the iDRAC Alerts Configuration Matrix. If a specific metric category—such as a thermal or temperature threshold breach—has not been explicitly checked and mapped to an email action inside this matrix, the iDRAC will log the hardware event internally to the System Event Log but will fail to transmit an external message. Consequently, when the CPU temperature surpassed the critical operational threshold, the server executed an automatic protective shutdown without dispatching an alert to the administrator. To prevent this, the alert state must be toggled from configured to fully active across the respective category rows. Study Guide References: Server Monitoring; iDRAC Alert Configuration Matrix; Thermal Management Policies.

Question: 2

An Administrator wants to change the PSU configuration In such a way that it can be redundant and can set PSU1 as the primary. What redundancy policy should the administrator select?

A. Redundant

B. PSU2 Redundant

C. A/B Redundant

D. PSU1 Redundant

Answer: A

Explanation:

Within the iDRAC power management interface for Dell PowerEdge servers, configuring power supply unit (PSU) high availability requires setting a foundational redundancy policy. The proper policy to select is 'Redundant'. This policy ensures that the power subsystem can survive the loss of a single power feed or a hardware failure within an individual power supply unit. Once the 'Redundant' policy wrapper is globally enabled, advanced sub-settings such as the Hot Spare feature become available. The Hot Spare framework allows the administrator to prioritize specific power supplies, enabling the option to configure a primary PSU (such as PSU1) to carry the active operational load while placing secondary units (such as PSU2) into a lower-consumption standby sleep state. If the primary power supply encounters a failure or an input voltage interruption, the standby PSU instantly transitions to an active state to maintain uninterrupted server operations. Study Guide References: Server Deployment; iDRAC Power Management; PSU Redundancy and Hot Spare Configuration.

Question: 3

A customer reports four physical GPUs installed but only three are detected by the OS on a PowerEdge XE with NVIDIA cards. Which command should be used to identify the missing GPU?

A. rocm-smi

B. nvidia-smi -L

C. dcgmi diag -r 4

Answer: B

Explanation:

Dell PowerEdge XE acceleration platforms hosting multiple high-performance enterprise accelerators rely on specific driver subsystems to communicate with the host operating system. When a hardware discrepancy occurs—such as only three out of four physical NVIDIA GPUs initializing successfully—the administrator must utilize the specialized NVIDIA System Management Interface utility (nvidia-smi) to interrogate the driver layer. Executing the command nvidia-smi -L lists all recognized graphics

processing units currently attached to the system along with their operational indexes, model names, and specific Universally Unique Identifiers (UUIDs). If a physical card is completely missing from this command output, it implies the operating system driver has failed to discover or bind to the device during the boot sequence. This standard isolation step allows an engineer to quickly cross-reference the recognized hardware against the system's physical topology to pinpoint the exact slot location requiring physical intervention. Study Guide References: Troubleshooting; GPU Acceleration Management; NVIDIA System Management Interface Subsystems.

Question: 4

A Dell PowerEdge server is experiencing intermittent hardware alerts, and the operating system is occasionally unresponsive. An administrator needs to diagnose and troubleshoot the hardware issue efficiently. Which three tools can be used to troubleshoot hardware issues on a Dell PowerEdge server?

A. Operating system application logs

B. LiveOptics

C. OME

D. iDRAC

E. Lifecycle Controller diagnostics

Answer: C, D, E

Explanation:

Isolating hardware faults on Dell PowerEdge servers requires leveraging specialized infrastructure tools designed to collect and parse low-level component telemetry. The Integrated Dell Remote Access Controller (iDRAC) provides an independent, out-of-band management plane that monitors component health status continuously and records historical errors within the localized System Event Log (SEL). For hardware-level diagnostics executed outside the operating system workspace, the Lifecycle Controller (LCC) diagnostics offer embedded, pre-boot UEFI test suites capable of performing exhaustive stress tests on the processor, physical memory blocks, and storage backplanes. At an infrastructure-wide level, OpenManage Enterprise (OME) acts as a centralized console that aggregates hardware alerts, evaluates firmware compliance baselines, and monitors structural events across multiple server nodes simultaneously. Together, these three solutions provide the dedicated out-of-band monitoring and localized diagnostics necessary to resolve intermittent hardware anomalies. Study Guide References: Troubleshooting; Dell OpenManage Systems Management Portfolio; Pre-boot Diagnostics and iDRAC Telemetry.

Question: 5

After adding DIMMs on a PowerEdge R760, the server powers on, fans spin, but the system does not reach POST. The LCD shows no error. What two things should be done to resolve this Issue?

A. Reseat memory

B. Replace the system board

C. Verify DIMM population

D. Update the BIOS

E. Clear SEL logs

Answer: A, C

Explanation:

When a modern Dell PowerEdge R760 server experiences a No-POST condition immediately following a memory capacity expansion, the root cause is almost exclusively tied to a physical installation error or an unvalidated configuration layout. Modern Intel Xeon Scalable processor architectures implement complex, multi-channel memory controllers that operate under rigid DIMM population matrices and balanced routing rules. If newly added modules are inserted into incorrect slots or violate the strict symmetrical configuration criteria mandated by the system architecture, the memory controller will fail to pass early hardware training phases, preventing the server from completing POST. To remediate this specific behavior, the technician must verify the exact DIMM population rules outlined in the technical documentation to ensure proper channel alignment. Following validation, a physical reseating of the memory modules must be completed to rule out minor electrical contact resistance or structural alignment faults inside the slots. Study Guide References: Troubleshooting; Memory Subsystem Architecture; POST Failure Isolation and Memory Population Rules.