ATMIC NETWORKS
YOUR GUIDE TO GLOBAL CERTIFIATION

"Please note that these files may not be up to date. However, the questions will help you understand the exam format and typical question patterns."

www.atmicnetworks .com

Warning: Keep connected with our support team for latest updates

Question: 1

Refer to the exhibit.

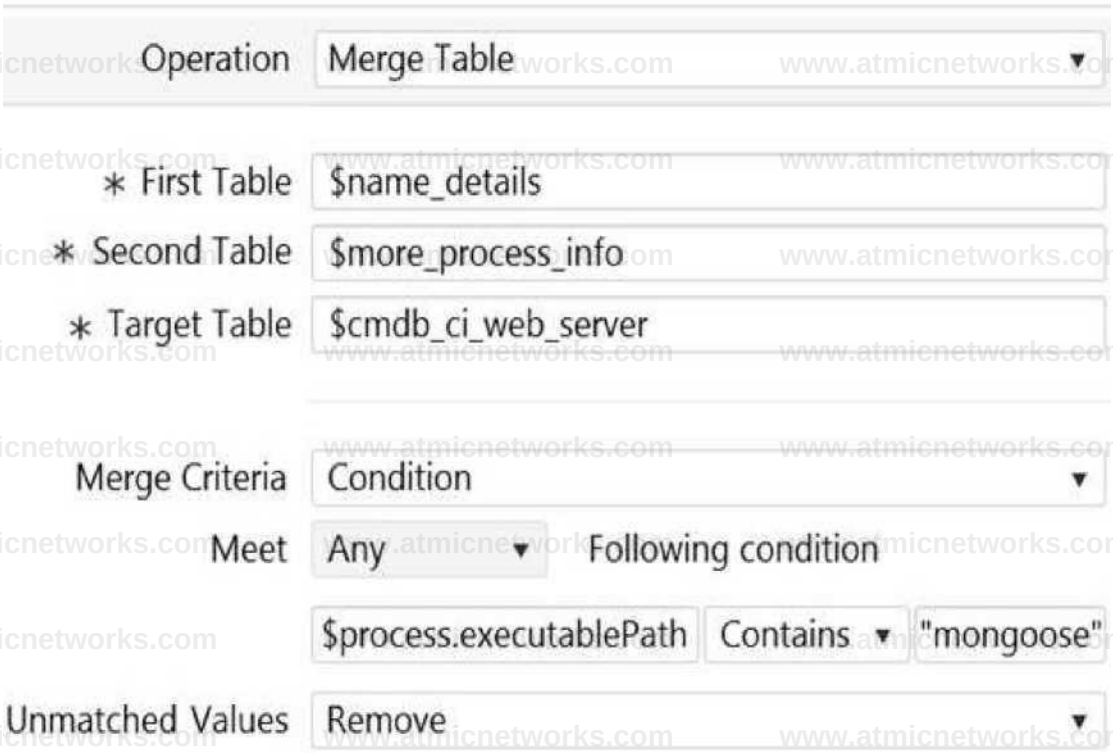

Based on this image, which of the following statements are true? (Choose three.)

A. Attributes from two tables populate a table with the same name as a ServiceNow CMDB table.
B. This operation is more than likely a part of a step on a pattern set to Application Pattern Type.
C. If a value is unmatched, it is still merged into the Target Table.
D. For this operation to run, there must be some data in the process.executablePath variable.
E. This is a horizontal pattern of type "infrastructure."

Answer: A, B, D

Explanation:

A is true because the target table $cmdb_ci_web_server is a ServiceNow CMDB table that stores information about web servers1.

B is true because the merge table operation is typically used for application patterns, which are horizontal patterns that discover applications and their dependencies. The condition on the process.executablePath variable suggests that the operation is looking for a specific application

(mongoose) running on the web servers.

D is true because the merge table operation requires at least one matching field value between the two source tables1. In this case, the process.executablePath variable is the matching field, and it must contain "mongoose" for the operation to run.

Reference:

1: Merge tables - Product Documentation: San Diego - Now Support Portal

3: Product Documentation | ServiceNow

[4]: Discovery Patterns - Product Documentation: San Diego - Now Support Portal

Question: 2

Refer to the exhibit.

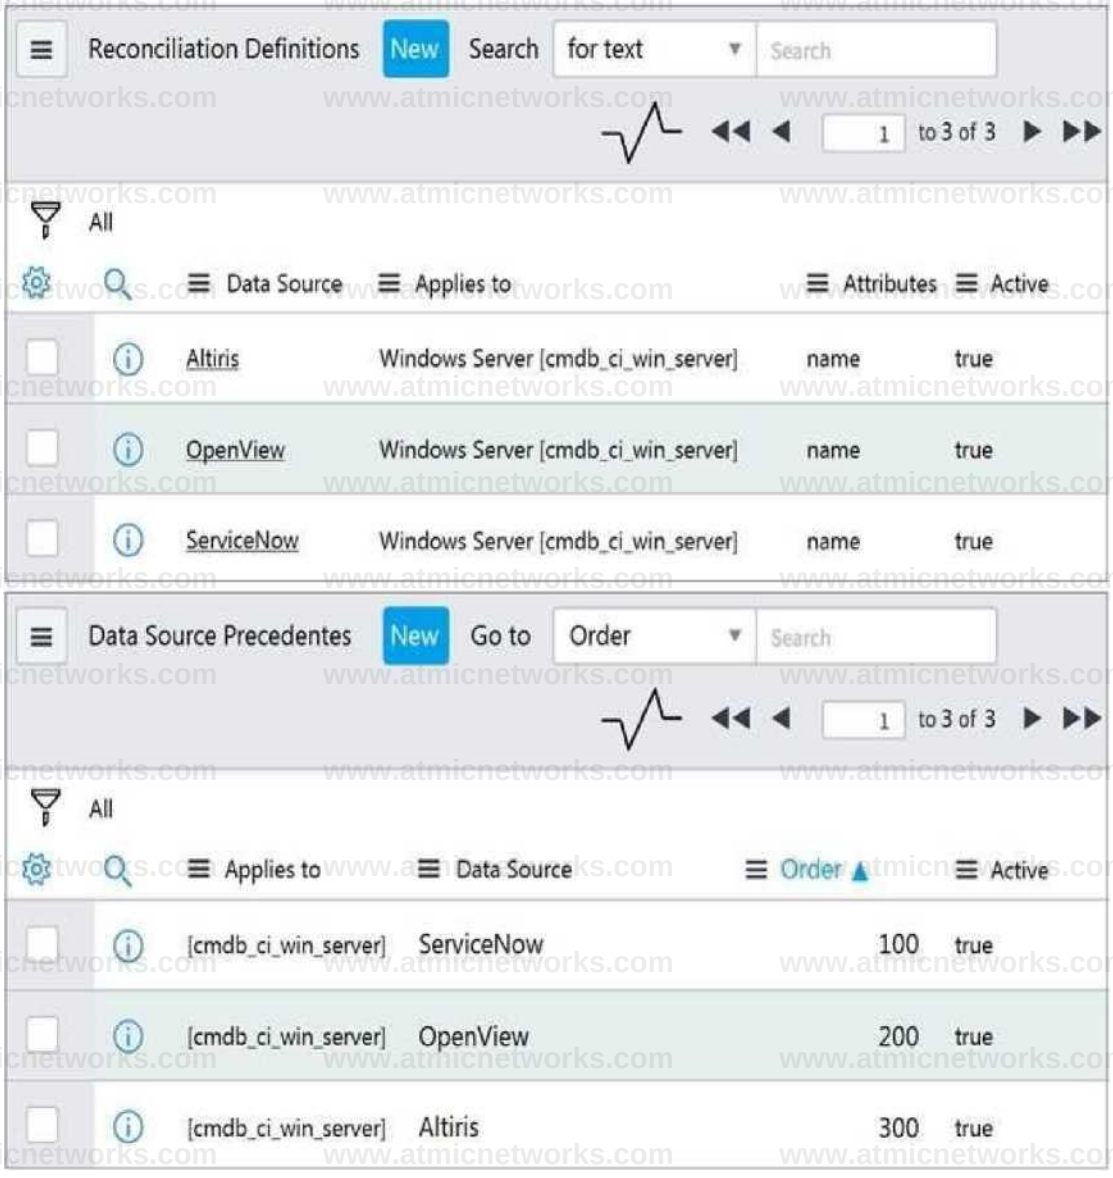

Based on the following images, which choice best describes what occurs if Discovery sets the name attribute of a discovered Windows Server CI to 'Windows1' and then Altiris discovery runs detecting 'Windows2' for the name attribute on the same CI?

A. The name of the CI stays 'Windows1'.

B. The name of the CI changes to 'Windows2'.

C. The name of the CI does not populate with either discovery.
D. The CI is not discovered because Discovery is not listed in either image.

Answer: B

Explanation:
In ServiceNow Discovery, the reconciliation process is governed by precedence rules. These rules determine which data source's information will be retained if there are conflicts when multiple sources discover the same CI. In this case, Altiris has a higher order of precedence (300) compared to ServiceNow (100), as seen in the "Data Source Precedents" section of the image. Therefore, if Altiris discovers 'Windows2' for the name attribute on the same CI after ServiceNow sets it to 'Windows1', the name will change to 'Windows2' due to Altiris's higher precedence.
Reference: The explanation is inferred from understanding how reconciliation and data source precedents work in ServiceNow Discovery, though not directly quoted from specific documents. You can find more information on these topics in the following links:
Reconciliation
Data source precedents

Question: 3

For the Parse Variable pattern operation, what is required to have two different parsing methods to populate variables?

A. Two different Debug Mode sessions.
B. A tabular and a scalar variable.
C. Two different steps.
D. Two different Define Parsing selections on the same step.

Answer: C

Explanation:
The Parse Variable pattern operation allows you to extract information from the output of a previous operation and save it in a variable. You can choose from different parsing methods, such as JSON File, XML File, Regular Expression, or Custom Script. To have two different parsing methods to populate variables, you need to use two different steps, each with a different Define Parsing selection. For example, you can use one step to parse a JSON file and another step to parse an XML file. You cannot use two different parsing methods on the same step, as the Define Parsing selection is unique for each step.
Reference:
Parse command output: This article explains how to use the Parse command output operation and the different parsing methods available.
Pattern Designer: Parse Variable - JSON File gives error: This article provides a troubleshooting tip for using the JSON File parsing method.
Examples of EVAL scripts used in Discovery patterns: This article provides some examples of custom scripts that can be used for the Custom Script parsing method.

Question: 4

Which best describes Discovery schedule of type Configuration Item?
A. Verifies Configuration Item data from the scanned IP ranges against the data in the CMDB.
B. Creates only a list of discovered IPs in both IPv4 and IPv6 formats.
C. Collects complete information from the scanned IP ranges and sends it to the CMDB.
D. Directly populates records in the assets table.

Answer: C

Explanation:

A Discovery schedule of type Configuration Item collects complete information from the scanned IP ranges and sends it to the CMDB. This type of schedule runs a series of probes and sensors to identify and classify the devices and applications on the network, and to create or update the corresponding configuration items in the CMDB. A Discovery schedule of type Configuration Item can also run patterns to discover more details and relationships about the configuration items.

Reference:

ServiceNow Discovery Overview

Create a Discovery schedule

Discovery schedule types

Question: 5

When installing a MID Server on a Windows platform, which right must be associated when creating a Service Account?

A. Local Admin
B. Domain Admin
C. MID Server User Role
D. Log on as service

Answer: D

Explanation:

The Service Account for the MID Server must have the Log on as service right on the Windows platform. This right allows the MID Server to run as a Windows service and communicate with the ServiceNow instance. The Service Account does not need to have local or domain admin rights, as these are not required for the MID Server functionality. The MID Server User Role is a role on the ServiceNow instance, not on the Windows platform, and it is used to control the access and permissions of the MID Server on the instance.

Reference:

Correcting MID Server Windows service account user and permissions

What is a ServiceNow MID Server and how does it work?

Configure Windows MID Server service credentials

Question: 6

Which of the below choices are needed for Quick Discovery? (Choose two.)

A. MID Server
B. Discovery Schedule
C. PID
D. Target IP

Answer: A, D

Explanation:

Quick Discovery is a wizard that helps you get up and running with Discovery quickly. It discovers both physical and logical components, including virtual machines, servers, storage, databases, applications, and more. To use Quick Discovery, you need to have a MID Server installed and configured, and provide a target IP range or

subnet to scan. You do not need to create a Discovery Schedule or a PID for Quick Discovery.

Reference: ServiceNow Discovery Data Sheet, Discovery Quick Start

Question: 7

In order to use Debug from the Pattern Designer, you must have what?

A. a proxy server
B. a discoverable CI
C. the admin role
D. Service Mapping installed

Answer: C

Explanation:

Debug mode is a feature of the Pattern Designer that allows you to test and troubleshoot your patterns in real time. To activate Debug mode, you need to have the admin role or a role that includes the pattern_designer_debug permission. Debug mode is not available for users who only have the pattern_designer_read permission1.

Reference: 1: Activate pattern Debug mode - Product Documentation: San Diego - ServiceNow

Question: 8

A discovery runs against a Windows Server returning the following attribute values for the first time:

name = WindowsSN1 serial_number = 12321

A subsequent discovery is ran against a different Windows Server returning the following attribute values:

name = WindowsSN2 serial_number = 12321

With only base system CI Identifiers configured, which of the following is true?

A. A Windows Server CI is created, then updated with WindowsSN2 as the name.
B. Two Windows Sewer CIs are created, with WindowsSN1 AND WindowsSN2 for names.
C. Two Windows Server CIs are created, without serial_number values.
D. A Windows Server CI is created, then updated with WindowsSN1 as the name.

Answer: B

According to the ServiceNow Discovery documentation, the base system CI Identifiers for Windows Server class are name and serial_number. These are the attributes that Discovery uses to uniquely identify a Windows Server CI. If both attributes match an existing CI, Discovery updates that CI. If only one attribute matches, Discovery creates a new CI. If neither attribute matches, Discovery also creates a new CI. In this scenario, the serial_number is the same for both Windows Servers, but the name is different. Therefore, Discovery will create two separate CIs, one with WindowsSN1 as the name and one with WindowsSN2 as the name.

Reference:

Create a Discovery CI classification - Product Documentation: Tokyo - Now Support Portal Identification rules - Product Documentation: Tokyo - Now Support Portal ServiceNow IRE: Identification Rules Explained — Cookdown

Question: 9

Which choice represents the three best ways of extending Discovery?

A. Orchestration, Classifiers, Discovery Patterns
B. Fingerprinting, Classifiers, Discovery Patterns
C. Orchestration, Classifiers, Probes & Sensors
D. Classifiers, Probes & Sensors, Discovery Patterns

E. Classifiers, Fingerprinting, Probes & Sensors

Answer: D

Explanation:

ServiceNow Discovery can be extended in three main ways: Classifiers, Probes & Sensors, and Discovery Patterns. Classifiers are used to identify the type of device or application that is being discovered, based on the information returned by a probe. Probes are commands or scripts that are executed on the target device or application to collect data. Sensors are scripts that process the data collected by probes and create or update configuration items (CIs) in the CMDB. Discovery Patterns are graphical representations of the discovery process that can be customized or created to discover specific types of devices or applications, using classifiers, probes, and sensors as building blocks. Reference:

Discovery overview - Product Documentation: Vancouver - Now Support Portal
Create a Discovery CI classification - Product Documentation: Vancouver - Now Support Portal Discovery probes and sensors - Product Documentation: Vancouver - Now Support Portal Discovery patterns - Product Documentation: Vancouver - Now Support Portal

Question: 10

SNMP Credentials require which of the following?

A. write community strings
B. usernames
C. read community strings
D. port 135 access

Answer: C

Explanation:

SNMP credentials are used by Discovery to communicate with devices that support the Simple Network Management Protocol (SNMP). SNMP credentials do not include a user name, just a password, called the community string. The default read-only community string for many SNMP devices is public, and Discovery will try that automatically. Enter the appropriate SNMP credentials if they differ from the public community string1.

Reference: 1: SNMP credentials - Product Documentation: Utah - Now Support Portal

Question: 11

Which choice will populate the Location field for a discovered CI?

A. Location field for a Discovery Schedule
B. Location field for a parent CI Type
C. Location field for a Port Probe
D. Location report from the Discovery Dashboard

Answer: A

Explanation:

The Location field for a discovered CI is populated by the Location field of the Discovery Schedule that triggered the discovery of that CI. This is done by the DiscoverySensor script, which gets the location ID from the Discovery Schedule and passes it to the Shazzam probe, which then updates the cidata with the location information. The Location field for a parent CI Type, a Port Probe, or a Discovery Dashboard report does not affect the Location field for a discovered CI.

Reference: Discovery - How Location field is set for a CI, How to pass discovery schedule fields to a CI record via discovery

Question: 12

What role is needed by the MID Server's user account to interact with a ServiceNow instance?

A. mid_server
B. discovery_admin
C. sm_mid
D. mid_discovery

Answer: A

Explanation:

The MID Server's user account must have the mid_server role to interact with a ServiceNow instance. This role allows the MID Server to access protected tables and perform discovery and orchestration tasks on behalf of the instance. The other roles are not related to the MID Server functionality. The discovery_admin role is for configuring and managing discovery, the sm_mid role is for using the Service Mapping MID Server, and the mid_discovery role is for running discovery probes and sensors.

Reference:

Setting up MID Server user and role

Configure Windows MID Server service credentials

Correcting MID Server Windows service account user and permissions

Question: 13

Which operation is used to change from the default credentials to any other appropriate credentials in a horizontal pattern?

A. Change credentials
B. Change user
C. Alternate credentials
D. Alternate user

Answer: B

Explanation:

A horizontal pattern is a type of Discovery pattern that discovers configuration items (CIs) and their relationships by moving across the network from one device to another. A horizontal pattern can use the Change user operation to switch from the default credentials to any other appropriate credentials for a specific device or application. This operation allows the pattern to access different types of CIs with different authentication methods.

Reference:

Discovery patterns

Change user operation

Question: 14

After navigating to an Automaton Error Messages list from Discovery > Home, how are the options on the right navigation pane categorized? (Choose two.)

A. SELECT ALL
B. SELECT ONE
C. ACTION ON SELECTED
D. ACTION ON ALL

Answer: A C

Explanation:

The Automation Error Messages list displays the errors that occurred during the execution of Discovery automation scripts. From this list, you can view the details of each error, such as the script name, the error message, the device, and the time. You can also perform actions on the errors, such as retrying the script, ignoring the error, or creating an incident. To do so, you can use the options on the right navigation pane, which are categorized as follows1: SELECT ALL: This option allows you to select all the errors in the list. SELECT ONE: This option allows you to select one error in the list by clicking on the checkbox next to it. ACTION ON SELECTED: This option allows you to perform an action on the selected errors, such as Retry, Ignore, or Create Incident. You can also choose to perform the action on all the errors in the list by selecting the All option from the drop-down menu.

ACTION ON ALL: This option allows you to perform an action on all the errors in the list, regardless of the selection. You can choose from the same actions as the ACTION ON SELECTED option.

Reference:

Discovery error messages: This article explains the different types of error messages and warnings in Discovery, and how to access and manage them.

Question: 15

Which of the following can be used in the Debug Identification Section in Debug Mode for an infrastructure pattern? (Choose two.)

A. IP
B. AWS Endpoint
C. PID
D. Host Name

Answer: AD

Explanation:

The Debug Identification Section in Debug Mode allows you to specify the identification attributes of a CI that you want to debug. These attributes are used to find the CI in the CMDB and run the pattern on it. The identification attributes vary depending on the CI class, but for infrastructure patterns, the common ones are IP and Host Name. These attributes are also used by the horizontal discovery to identify CIs. AWS Endpoint and PID are not valid identification attributes for infrastructure patterns. Reference: The explanation is based on the following sources:

Activate pattern Debug mode: This document explains how to activate the Debug mode and use the Debug Identification Section to debug a pattern.

ServiceNow Exam CIS-Discovery Topic 6 Question 9 Discussion: This discussion provides a similar question and answer about the Debug Identification Section in Debug Mode.

Question: 16

With multiple CI data sources, which choice is the best for determining which source can update a CI attribute?

A. Business Rules
B. Data Certification
C. Transform Maps
D. Reconciliation Rules

Answer: D

Explanation:
Reconciliation rules are used to determine which data sources can update CI attributes in ServiceNow. Reconciliation rules can be static or dynamic, and they can be applied to specific CI classes, attributes, or discovery sources. Reconciliation rules can also specify the actions to take when there are unmatched values, such as removing, ignoring, or updating them1.

Reference:
1: Reconciliation rules - Product Documentation: San Diego - Now Support Portal

Reference: https://docs.servicenow.com/bundle/orlando-servicenow-platform/page/product/configuration- management/task/t_DefineDataSourcePrecedence.html

Question: 17

Which method for deleting specific CIs is not discovered in 30 days?

A. Scheduled Job
B. UI Policy
C. Service Mapping
D. Data Policy

Answer: A

Explanation:

A scheduled job is a background process that runs at a specified time or interval to perform a specific task1. It is not a method for deleting specific CIs that are not discovered in 30 days. The other options are methods for deleting or updating CIs based on discovery data. A UI policy is a script that can dynamically change the behavior of a form or list2. A service mapping is a process that creates a map of the relationships between CIs that support a business service3. A data policy is a rule that enforces data consistency and accuracy by validating the data entered into a record.

Reference: 1: Scheduled Jobs 2: UI Policies 3: Service Mapping : Data Policies

Question: 18

Which of the following choices must be installed on a MID Server to run Credential-less Discovery?

A. Credential-less Extension
B. Nmap
C. Advanced IP Scanner

D. Defender

Answer: B

Explanation:

Credential-less Discovery is a feature of ServiceNow Discovery that allows the instance to identify configuration items (CIs) without using credentials. Credential-less Discovery uses Network Mapper (Nmap), a free and open-source network scanner, to collect information about the CIs by sending packets and analyzing the responses. Nmap must be installed on the MID Server that runs Credential-less Discovery. Credential-less Discovery can be used as a fallback option when credentialbased probes fail, or as a primary option for scanning certain types of devices, such as network devices, printers, or IoT devices12.

Reference:

Credential-less Discovery with Nmap - Product Documentation: San Diego - Now Support Portal Credential-less host Discovery - ServiceNow - Now Support

Question: 19

A network device has both an SSH port and an SNMP port open. Discovery tries the SSH probe first and it fails. This triggers the SNMP probe, which succeeds. Discovery uses SNMP first for subsequent discoveries on that device.

What discovery functionality allows the above to happen?

A. Classification

B. Credential affinity C. MID Server affinity D. IP service affinity

Answer: B

Explanation:

Credential affinity is an association between a set of credentials and a device on your network. When Discovery or Orchestration first attempts to access a device, they try all available credentials until a valid one is found. Once a valid credential is found, it is recorded in the dscy_credentials_affinity table, and further discovery processes rely on this information. If the credentials are modified or changed, the process will be repeated until a valid credential is found and updated in the same table. Credential affinity allows Discovery to use the most efficient and successful credential for each device, and avoid unnecessary credential failures.

Reference: Credential affinity for Discovery and Orchestration, Credential Management in ServiceNow discovery

Question: 20

The CMDB contains which of the following record types? (Choose two.)

A. Model

B. Configuration Item (CI)

C. Asset

D. Relation Type

Answer: BD

Explanation:

The CMDB contains records of configuration items (CIs) and their relationships. A CI is any component that needs to be managed in order to deliver an IT service, such as a server, an application, or a user. A relation type defines the nature of the connection between two CIs, such as depends on, uses, or contains. A model is a template for a CI that defines its attributes and default values. An asset is a tangible or intangible resource that is tracked and managed by the organization, such as a laptop, a license, or a contract.

Reference:

Configuration Item (CI) types: Attributes and relationships What is a configuration management database (CMDB)? CSA Challenge Questions Flashcards

Reference: https://docs.servicenow.com/bundle/quebec-servicenow- platform/page/product/configuration-management/reference/r_CMDBRecordTypes.html

Question: 21

When is the Extension section in a horizontal pattern executed?

A. As part of the post sensor processing script
B. After the Identification sections
C. As part of the port scan
D. Before the Identification sections

Answer: B

Explanation:

The Extension section in a horizontal pattern is executed after the Identification sections have completed and the CI has been identified. The Extension section allows adding additional attributes or relationships to the CI, or launching other patterns for further discovery12.

Reference:

1: ITOM: Extending Discovery/Service Mapping Patterns - GlideFast ServiceNow
2: Horizontal Pattern probe - Product Documentation: San Diego - ServiceNow

Question: 22

For CMDB Health, relationships can be which of the following choices? (Choose three.)

A. Duplicate
B. Stale
C. Orphan
D. Required
E. Recommended

Answer: ABC

Explanation:

CMDB Health measures the quality of the data in the CMDB by testing it against predefined rules and metrics. One of the categories of these rules is Correctness, which tests the data against data integrity rules such as identification rules, orphan CI rules, and stale CI rules1. These rules help identify and flag the existing CMDB relationships that are not compliant with the suggested relationships2.

Duplicate relationships are those that have more than one relationship of the same type between the same

two CIs3. For example, if a server CI has two relationships of type "Depends on" with the same network CI, it is a duplicate relationship.
Stale relationships are those that have not been updated for a specified period of time3. For example, if a server CI has a relationship of type "Runs on" with a virtual machine CI that has not been discovered for more than 60 days, it is a stale relationship.
Orphan relationships are those that are missing a defined relationship within the CMDB3. For example, if a server CI has no relationship of type "Depends on" with any network CI, it is an orphan relationship.
Reference: 1: CMDB Health - Customer Success - ServiceNow 2: CMDB health Dashboard: "Relationships not compliant with suggested relationships" calculation - Support and Troubleshooting - Now Support Portal 3: Maintain a healthy CMDB - ServiceNow

Question: 23

Which of the choices provides active discovery errors with a help link for each error?

A. Discovery Dashboard
B. IP Address Failure Report
C. Discovery Schedule
D. MID Server Dashboard

Answer: A

Explanation:
The Discovery Dashboard provides a summary of the Discovery status, including the number of active discovery errors, the number of devices discovered, and the number of credentials used. Each error has a help link that provides more information about the cause and possible solutions.
Reference:
Discovery Dashboard
Discovery troubleshooting | Error messages

Question: 24

What related list on a classifier dictates which Horizontal Pattern probe is launched?

A. Discovery Log
B. Classification Criteria
C. Pattern probes
D. Triggers probes

Answer: D

Explanation:
A classifier is a set of rules that identifies a device type based on the results of a probe. A classifier can have one or more related lists that define the classification criteria, the pattern probes, and the triggers probes. The pattern probes related list specifies which horizontal discovery pattern to run for the device type. The triggers probes related list specifies which probes to run after the device is classified, and before the horizontal discovery pattern is executed. The triggers probes can be used to gather additional information or credentials that are needed for the horizontal discovery pattern12. Reference:
1: ServiceNow Discovery Overview, page 9
2: ServiceNow Discovery Documentation, Classifiers section

Question: 25

Which of the following does the ECC Queue provide? (Choose two.)

A. Login credentials for the MID Server host.
B. The actual XML payload that is sent to or from an instance.
C. A connected flow of probe and sensor activity.
D. The process responsible for defining, analyzing, planning, measuring, and improving all aspects of the availability of IT services.

Answer: BC

Explanation:

The ECC Queue is the normal connection point between an instance and other systems that integrate with ServiceNow, such as MID Servers1. The ECC Queue displays input and output messages from and to MID Servers2. The actual XML payload that is sent to or from an instance is stored in the ECC Queue3. The ECC Queue also provides a connected flow of probe and sensor activity, as probes are sent from the instance to the MID Server, and sensors are sent from the MID Server to the instance2. Reference:

1: ServiceNow Discovery Overview

2: ECC Queue Processing and Debugging, with "Discovery - Sensors" used as an example

3: Manage ECC Queue content for a MID Server

Question: 26

Which of the following choices are only used for the Application Pattern Type? (Choose two.)

A. Run Order
B. Identification Section
C. CI Type
D. Operating System

Answer: AB

Explanation:

The Application Pattern Type is a pattern that discovers applications and their dependencies by using probes and sensors1. It has two main components: the Run Order and the Identification Section2. The Run Order defines the sequence of probes and sensors that Discovery uses to find and classify CIs2. The Identification Section specifies the criteria that Discovery uses to identify CIs and their relationships2.

The other options, CI Type and Operating System, are not specific to the Application Pattern Type. They are used for other pattern types, such as the Horizontal Pattern Type and the Service Mapping Pattern Type3 .

Reference: 1: Application Pattern Type 2: Application Pattern Type components 3: Horizontal Pattern Type : Service Mapping Pattern Type

Reference: https://docs.servicenow.com/bundle/paris-it-operations- management/page/product/service-mapping/task/t_CreatePatternPatDef.html

Question: 27

By default, which of the following are automatically available as variables for horizontal discovery patterns? (Choose two.)

A. infrastructure_system
B. The CI Type on the Discovery Pattern form
C. windows_cmdb_ci
D. computer_system

Answer: AD

Explanation:
Horizontal discovery patterns are a series of operations that tell Discovery which CIs to find on your network, what credentials to use, and what tables to populate in the CMDB. Horizontal discovery patterns use variables to store and pass information between the operations. By default, there are two variables that are automatically available for horizontal discovery patterns: infrastructure_system and computer_system. These variables store the information about the infrastructure and the computer system of the target CI, respectively. They are populated by the

Horizontal Pattern probe and sensor, which enable Discovery to use patterns for horizontal discovery123.

Reference:
Patterns and horizontal discovery - Product Documentation: Tokyo - Now Support Portal
Horizontal Pattern probe - Product Documentation: San Diego - Now Support Portal
Using patterns for horizontal discovery - Product Documentation: Vancouver - ServiceNow

Reference: https://docs.servicenow.com/bundle/quebec-it-operations-management/page/product/service- mapping/reference/r_PatternVariables.html

Question: 28

Which choices are necessary to launch any pattern? (Choose two.)

A. CI Classification
B. CI Serial Number Attribute
C. Data Certification
D. CI Type

Answer: AD

Explanation:

To launch any pattern, you need to specify the CI classification and the CI type. The CI classification determines the table where the discovered CI is stored in the CMDB, and the CI type defines the specific attributes and relationships of the CI. For example, if you want to discover a Windows server, you need to select the CI classification as cmdb_ci_win_server and the CI type as Windows Server. These choices are mandatory for any pattern, as they enable Discovery to identify and classify the CIs correctly.
Reference: Patterns and horizontal discovery, Available discovery patterns

Reference: https://docs.servicenow.com/bundle/quebec-it-operations-management/page/product/discovery/ concept/c-UsingPatternsForHorizontalDiscovery.html

Question: 29

Which of the below choices are horizontal pattern types? (Choose two.)

A. Hardware
B. Software

C. Infrastructure
D. Application

Answer: CD

Explanation:

Horizontal pattern types are used to discover the configuration items (CIs) that belong to a specific category or class, such as infrastructure or application. They define the operations that Discovery performs to identify and classify the CIs, and the attributes that Discovery populates for each
CI. Horizontal patterns can be applied to any operating system or platform that supports the required

protocols and commands12.
Reference:
Horizontal discovery patterns
Discovery pattern types

Reference: https://docs.servicenow.com/bundle/quebec-it-operations-management/page/product/service- mapping/concept/c_MappingPatternsCustomization.html

Question: 30

What does the MID Server need to collect vCenter events?

A. vCenter Event Collector extension
B. MID SNMP Trap Listener extension
C. Firewall
D. vCenter probes

Answer: A

Explanation:

The vCenter Event Collector is a MID Server extension that listens for vCenter-related events and updates the CMDB accordingly. The event collector allows the CMDB to be updated with changes to virtual machines (VMs), in addition to the updates detected by Discovery1.
Reference:
3: Product Documentation | ServiceNow

Reference: https://docs.servicenow.com/bundle/quebec-servicenow-platform/page/product/mid- server/concept/c_VCenterEventProcessorExtension.html

Question: 31

A Discovery Schedule contains a /24 subnet IP Range and a Shazzam batch size of 5000. How many times will a Shazzam probe be launched during discovery?

A. 1
B. 2
C. 5000
D. 254

Answer: C

Explanation:

A /24 subnet IP Range means that there are 256 possible IP addresses in the range, from 0 to 255. A Shazzam batch size of 5000 means that Discovery will send 5000 ICMP packets at a time to scan the IP addresses. Therefore, Discovery will launch the Shazzam probe once for every 5000 IP addresses in the range, or 256 / 5000 = 0.0512 times. Since the number of times must be an integer, Discovery will round up to 1 and launch the Shazzam probe once. However, since the Shazzam probe is launched for each MID Server that is assigned to the Discovery Schedule, the actual number of times will depend on how many MID Servers are available. If there are N MID Servers, then the Shazzam probe will be launched N times, each sending 5000 ICMP packets to the same IP range. Therefore, the total number of times that the Shazzam probe will be launched during discovery is N * 1 = N.

Reference:

Discovery Schedule form

Shazzam Probe

Discovery IP range

Question: 32

Which method is used by Discovery to determine if a Host IP is active or alive?

A. Port Scan
B. Traceroute
C. Ping
D. Classification

Answer: C

Explanation:

Discovery uses the ping method to determine if a host IP is active or alive. Ping is a network utility that sends an ICMP echo request packet to a target IP address and waits for an ICMP echo reply packet. If the target IP address responds, it means that the host is active or alive12.

Reference:

1: ServiceNow Discovery Documentation, Discovery Process section

2: ServiceNow Discovery Overview, page 8

Question: 33

Discovery finds and maps dependencies for the following types of storage devices. (Choose three.)

A. Direct-attached storage (DAS)
B. Network-attached storage (NAS)
C. Storage area network (SAN)
D. Multiple area network (MAN)
E. Redundant Array of Independent Disks (RAID)

Answer: ABC

Explanation:

Discovery finds and maps dependencies for the following types of storage devices: Direct-attached storage

(DAS), network-attached storage (NAS), or storage area network (SAN)1. NAS or SAN storage that is discovered via a Storage Management Initiative Specification (SMI-S) and Common Information Model (CIM) is also supported2. Discovery collects information about storage area networks from specialized devices, such as storage arrays and Fibre Channel (FC) switches, and creates specific references between the tables in the SAN3. Discovery does not support Multiple area network (MAN) or Redundant Array of Independent Disks (RAID) as types of storage devices. Reference:

1: Storage device discovery

2: Storage Discovery via SMI-S and CIM

3: Discovery of storage area networks (SAN)

Question: 34

Which choice allows the following functionality to occur?

If this value is set to 1000 and a discovery must scan 10,000 IP addresses using a single MID Server, it creates 10 Shazzam probes with each probe scanning 1000 IP addresses.

A. MID Server Clusters
B. MID Server selection method
C. Shazzam Batch Size
D. Behaviors

Answer: C

Explanation:

The Shazzam Batch Size property determines how many IP addresses are scanned by a single Shazzam probe1. If this value is set to 1000 and a discovery must scan 10,000 IP addresses using a single MID Server, it creates 10 Shazzam probes with each probe scanning 1000 IP addresses. This property can be configured in the Discovery Properties module under Discovery Definition2.

Reference: 1: Shazzam probe, port probes, and protocols - Now Support 2: [Discovery properties - ServiceNow]

Question: 35

Which choice best describes how to use a Behavior for discovery?

A. The MID Server selection method on a Discovery Schedule.
B. The Behavior drop-down menu on a Discovery IP Range.
C. The Behavior drop-down menu on a Discovery Status.
D. The Behavior checkbox on a CI.

Answer: B

Explanation:

A Behavior is a set of rules that determines which MID Servers launch which types of probes during a Discovery1. It can be assigned to a Discovery IP Range, which is a range of IP addresses that Discovery scans for CIs2. By using the Behavior drop-down menu on a Discovery IP Range, one can control how Discovery behaves for each IP address in that range3. The other options are not related to using a Behavior for discovery. The MID Server selection method on a Discovery Schedule is a way to specify which MID Servers are eligible to run Discovery for a given schedule4. The Behavior dropdown menu on a Discovery Status is a way to filter the Discovery Status records by the Behavior used. The Behavior checkbox on a CI is a way to indicate

whether the CI was discovered using a Behavior or not.

Reference: 1: Discovery behaviors 2: Discovery IP ranges 3: Assign a behavior to an IP range 4: MID Server selection method : Discovery Status : Behavior field on the cmdb_ci table

Reference: https://docs.servicenow.com/bundle/paris-it-operations-management/page/product/discovery/task/ create-disco-behavior.html

Question: 36

Which of the following pattern operations query targets? (Choose two.)

A. WMI Query B. Merge Table C. Get Process D. Parse Variable

Answer: A, C

Explanation:

Pattern operations are the building blocks of patterns that define the logic and actions to perform on the target CIs. Some pattern operations query targets to retrieve information or execute commands, while others manipulate data or perform other tasks. WMI Query and Get Process are examples of pattern operations that query targets. WMI Query executes a Windows Management Instrumentation (WMI) query on a Windows target and stores the results in a variable. Get Process executes a command to list the processes running on a target and stores the results in a variable12. Reference:

Pattern operations - Product Documentation: Vancouver - ServiceNow

Pattern operations - Product Documentation: San Diego - Now Support Portal

Reference: https://docs.servicenow.com/bundle/paris-it-operations- management/page/product/service-mapping/task/t_WMIQueryPatDef.html

Question: 37

File-based Discovery is triggered during the

A. Classify Phase
B. Scan Phase
C. Exploration Phase
D. Pattern Phase
E. Identification Phase

Answer: C

Explanation:

File-based Discovery is a process that helps identify what software is running on your Windows and UNIX servers and devices, even if there is no registration information available. File-based Discovery is triggered in the exploration phase of normal Discovery, after the target device has been classified and identified. File-based Discovery probes execute a scan searching for specific file extensions or file names in paths that you configure. The resulting file information is returned in the probe payload. The sensor attempts to match the discovered files with installed software, using the file name, size, and version returned by the probe.

Reference: File-based Discovery, File Based Discovery - Support and Troubleshooting

Reference: https://docs.servicenow.com/bundle/paris-it-operations- management/page/product/discovery/concept/file-based-discovery.html

Question: 38

Given a custom column named u_custom_column on table cmdb_ci_linux_server, which variable syntax should be used to populate the column in a horizontal discovery pattern using the Set Parameter Value operation?

A. $user_var_custom_column
B. $cmdb_ci_linux_server.u_custom_column.INSERT
C. $u_custom_column[1].cmdb_ci_linux_server
D. $cmdb_ci_linux_server[*].u_custom_column

Answer: B

Explanation:

The correct variable syntax to populate a custom column in a horizontal discovery pattern is $cmdb_ci_linux_server.u_custom_column.INSERT. This syntax indicates that the value of the custom column will be inserted into the cmdb_ci_linux_server table for the current CI. The other options are either invalid or incorrect syntaxes for this operation12.

Reference:

Set Parameter Value operation

Pattern variables

Question: 39

What entry point type must a horizontal pattern have to execute from a process classifier?

A. A subnet entry point type.
B. HTTP(S) entry point type if the pattern is running on a web server application.
C. TCP entry point type or ALL entry point type.
D. It does not matter, it is triggered for all entry point types.

Answer: C

Explanation:

A horizontal pattern must have a TCP entry point type or an ALL entry point type to execute from a process classifier. A process classifier is a rule that matches a process running on a host and triggers a horizontal pattern to discover the application associated with that process1. The entry point type determines the type of connection that the pattern uses to access the target host. A TCP entry point type means that the pattern uses a TCP port to connect to the host, while an ALL entry point type means that the pattern can use any available connection method2.

Reference:

1: Process classifiers - Product Documentation: San Diego - ServiceNow
2: Horizontal Pattern probe - Product Documentation: San Diego - ServiceNow

Reference: https://docs.servicenow.com/bundle/quebec-it-operations-management/page/product/discovery/ concept/c-UsingPatternsForHorizontalDiscovery.html

Question: 40

Which metrics comprise the Completeness KPI for CMDB Health? (Choose two.)

A. Required
B. Recommended
C. Audit
D. Overall

Answer: AB

Explanation:

Reference: https://docs.servicenow.com/bundle/quebec-servicenow-platform/page/product/configuration- management/reference/r_CMDBHealthMetrics.html

Question: 41

Which of the following choices explain differences between Service Mapping and Discovery? (Choose two.)

A. Discovery requires agent installation to find hardware devices, Service Mapping requires agents for software.
B. Discovery finds applications and devices on your network, Service Mapping monitors those devices.
C. Discovery utilizes IP address ranges for initial discovery, Service Mapping uses entry points.
D. Discovery addresses inventory-related use-cases, while Service Mapping allows for the creation of accurate maps of application service topologies.

Answer: C, D

Explanation:

Discovery and Service Mapping are both products that help to identify and map the IT infrastructure and services in an organization. However, they have different approaches and objectives. Discovery uses IP address ranges as the starting point for finding devices and applications on the network, and then runs probes and sensors to collect information and classify them. Discovery focuses on inventory-related use-cases, such as asset management, configuration management, and compliance. Service Mapping uses entry points, such as URLs or host names, as the starting point for mapping business services and their dependencies. Service Mapping runs patterns, which are sequences of operations that follow the traffic connections between devices and applications, and then creates a service map that shows the logical and semantic relationships. Service Mapping focuses on service-related use-cases, such as service availability, impact analysis, and root cause analysis12.

Reference:

1: ServiceNow Discovery Overview, page 9
2: ServiceNow Discovery Documentation, Service Mapping section

Reference: https://community.servicenow.com/community?id=community_article&sys_id=767aab87dbccc0106064eeb5ca9619fb

Question: 42

What are the two main options within a Parse File operation?

A. Discover Now and Quick Discovery

B. Select Operating System and Method

C. Select File and Define Parsing

D. Match and Select File

Answer: C

Explanation:

The Parse File operation is used to extract information from a file and create variables to contain the extracted information1. The two main options within a Parse File operation are Select File and Define Parsing2. Select File allows you to specify the file path or browse to the file. Define Parsing allows you to select the relevant parsing strategy and define parsing criteria2.

Reference:

1: Parse a file

2: Define discovery steps

Reference: https://docs.servicenow.com/bundle/paris-it-operations- management/page/product/service-mapping/task/t_ParseFilePatDef.html

Question: 43

Refer to the exhibit.

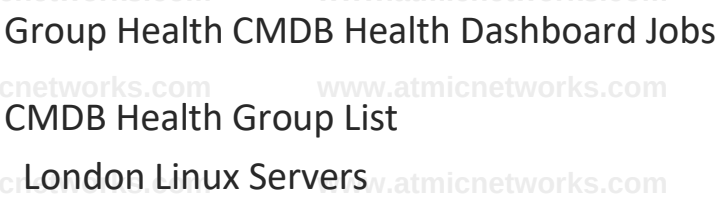

Based on the following image, which of the following choices is also true about London Linux Servers?

A. It is a CMDB Group with Dashboard Group type.

B. It is a CMDB Group with Health Group type.

C. It is a Datacenter Group in London.

D. It is a CMDB Group with Default Group type.

Answer: B

Explanation:

London Linux Servers is a CMDB Group with Health Group type. A CMDB Group is a collection of configuration items (CIs) that share common characteristics or are related in some way1. A Health Group is a type of CMDB Group that is used to monitor and measure the health of a specific set of CIs based on predefined rules and metrics2. The image shows that London Linux Servers is selected from the CMDB Health Group List, which is a dropdown menu that displays all the Health Groups defined in the system3. The image also shows the Group Health tab, which displays the health scores and indicators for the selected Health Group.

Reference: 1: Create a CMDB group - ServiceNow 2: CMDB health groups - ServiceNow 3: CMDB health dashboard - ServiceNow : [View CMDB health scores - ServiceNow]

Question: 44

For the Set Parameter Value operation, which of the following is used in the syntax to declare a constant, unchanging Value?

A. Hash tag
B. Brackets
C. Quotes
D. Dollar sign

Answer: C

Explanation:

The Set Parameter Value operation is used to assign a value to a variable in a Discovery pattern1. The syntax for this operation is as follows2:

Set Parameter Value: <Variable Name> = <Value>

The <Value> can be either a constant or an expression. To declare a constant, unchanging value, one must use quotes around the value. For example2:

Set Parameter Value: $name = "John"

This assigns the string “John” to the variable $name. The quotes indicate that the value is a constant and will not change. The other options, hash tag, brackets, and dollar sign, are not used to declare a constant value. They have different meanings in the syntax of the Set Parameter Value operation. A hash tag (#) is used to indicate a comment2. Brackets ([]) are used to access an element of an array or a map2. A dollar sign

Question: 45

From an SNMP Query pattern operation, which of the choices are valid Variable Types? (Choose two.)

A. Test
B. Table
C. Scalar
D. CI Type

Answer: B, C

Explanation:

The SNMP Query pattern operation executes an SNMP query on a target device and stores the results in a variable. The Variable Type field determines the format of the results. There are two valid Variable Types for this operation: Table and Scalar. Table is used when the query returns multiple values, such as a list of interfaces or processes. Scalar is used when the query returns a single value, such as the system name or uptime12.

Reference:

SNMP Query pattern operation - Product Documentation: San Diego - Now Support Portal SNMP Query pattern operation - Product Documentation: Vancouver - ServiceNow

Reference: https://docs.servicenow.com/bundle/quebec-it-operations-management/page/product/service- mapping/reference/r_PatternVariables.html

Question: 46

Which choice best describes what happens when, by default, duplicate CIs are detected during identification

and reconciliation?

A. A notification is sent to the CI owner.
B. An associated identification rule is created automatically.
C. Each set of duplicate CIs is added to a de-duplication task.
D. The next discovery is stopped for the CI that is duplicated.

Answer: C

Explanation:
When Discovery or Service Mapping identifies a CI that matches an existing CI in the CMDB, it checks whether the CIs are duplicates or not. If they are duplicates, Discovery adds them to a de-duplication task, which is a record that groups duplicate CIs and allows you to review and remediate them. You can use the Duplicate CI Remediator wizard to merge, flag, retire, or delete the duplicate CIs. By default, Discovery does not send notifications, create identification rules, or stop the next discovery for duplicate CIs.
Reference: Handle Duplicate CIs in your CMDB, De-duplication tasks, How to de-duplicate 2 ESX servers with same name?

Question: 47

Which choice best describes a horizontal discovery pattern?

A. Steps that execute operations
B. Credential depot
C. Port scanning tool
D. Classifiers that execute probes

Answer: A

Explanation:
A horizontal discovery pattern is a set of steps that execute operations to identify, classify, and populate attributes for configuration items (CIs) of a specific category or class, such as infrastructure or application. A horizontal pattern can be applied to any operating system or platform that supports the required protocols and commands12.
Reference:
Horizontal discovery patterns
Discovery pattern types

Question: 48

A config file for an application has the following three lines:
Line 1: app build 1.2.3.4 version 5.14
Line 2: installation_dir=c:\opt\bin
Line 3: build_type=Server.
Which methods below will extract the build and version numbers from these lines using a horizontal discovery pattern? (Choose two.)

A. Get Process operation with correct Port
B. Find Matching URL operation with Target Variable
C. Parse File operation with Delimited Text parsing strategy
D. Parse File operation with Regular Expression parsing strategy

Answer: C, D

Explanation:
A horizontal discovery pattern is a series of operations that tell Discovery how to find and identify an

application on a host1. To extract the build and version numbers from the config file, the pattern can use either of the following methods:

Parse File operation with Delimited Text parsing strategy: This operation reads the content of a file and splits it into fields based on a delimiter character. The pattern can specify the delimiter as a space and the fields to extract as the second and fourth ones from the first line of the file2.

Parse File operation with Regular Expression parsing strategy: This operation reads the content of a file and matches it against a regular expression. The pattern can specify the regular expression as app build (\d+\.\d+\.\d+\.\d+) version (\d+\.\d+) and the groups to extract as the first and second ones from the first line of the file2.

Reference:

1: Patterns and horizontal discovery - Product Documentation: Tokyo - ServiceNow

2: Parse File operation - Product Documentation: San Diego - ServiceNow

Question: 49

The following shows part of the Windows OS - Servers pattern in Pattern Designer:

Steps 0 S 8

= ;E **25. Windows - Storage**

J 26. Insert total disk space of the server

Which of the steps above use(s) a shared library?

A. Step 26
B. Neither step
C. Step 25
D. Both steps

Answer: C

Explanation:

Step 25 uses a shared library called Windows - Storage, which is a reusable component that can be invoked by multiple patterns to discover the storage information of Windows servers. A shared library is a collection of steps that can be referenced by a pattern name and can accept input parameters and return output values. Shared libraries help to avoid duplication of code and simplify the maintenance of patterns.

Reference:

Pattern Designer | Shared libraries

Windows - Storage shared library

Question: 50

Which of the following is required for a MID Server to have access to automatically stay up-to-date with instance versions?

A. install.service-now.com
B. docs.servicenow.com
C. developer.service-now.com
D. service-now.com

Answer: A

Explanation:

A MID Server needs to have access to the install.service-now.com domain to automatically stay up- to-date with instance versions. This domain hosts the MID Server upgrade files that are downloaded and installed by the MID Server when a new version is available. The MID Server checks for updates every 24 hours by default, or when the mid.upgrade.check_interval property is set12.

Reference:

1: ServiceNow Discovery Documentation, MID Server Upgrade section

2: ServiceNow Support Article, Connections to allow list for the MID Server

Question: 51

When designing steps with operations requiring variables, it is best practice to do what?

A. hard core variables
B. always use scalar variables
C. query targets for variables
D. design for a static environment

Answer: C

Explanation:

When designing steps with operations requiring variables, it is best practice to query targets for variables1. This allows you to dynamically obtain the values of the variables from the target devices, rather than hard coding them or using scalar variables that may not reflect the current state of the target. Querying targets for variables also enables you to design for a dynamic environment, where the target configuration may change over time.

Reference:

1: Define discovery steps

Question: 52

For what is File Based Discovery used?

A. To discover the checksum of a file and store it to track for changes
B. To discover the contents of flat files such as configuration files
C. To discover that file names conform to a defined naming standard
D. To discover file paths to recognize the signature of installed software

Answer: D

Explanation:

File Based Discovery is used to discover file paths to recognize the signature of installed software. File Based Discovery helps to identify what software is running on Windows and UNIX servers and devices by scanning the file system for specific files or directories that indicate the presence of a software product1. File Based Discovery can also check the file size, version, and checksum of the discovered files to verify the software installation2. File Based Discovery can be configured using file signatures, which are patterns that match the file paths of software products3. Reference: 1: File Based Discovery - Support and Troubleshooting - Now Support Portal 2: Use filebased service discovery to discover scrape targets | Prometheus 3: How File based discovery can be used to populate software record running on Linux host. - Support and Troubleshooting - Now Support Portal

Question: 53

Which of the following properties define the maximum overall size for the returned payload that comes

from patterns?

A. cmdb.properties.payload_max_size
B. glide.discovery.payload_max
C. mid.discovery.max_payload_size
D. mid.discovery.max_pattern_payload_size

Answer: D

Explanation:
The mid.discovery.max_pattern_payload_size property defines the maximum overall size for the returned payload that comes from patterns. This property is set on the MID Server and applies to all patterns that run on that MID Server. If the payload size exceeds this limit, the pattern execution fails and an error message is logged in the ECC queue. The default value of this property is 10 MB1. Reference: 1: ServiceNow Docs - MID Server properties for Discovery and Service Mapping 2: ServiceNow Docs - Discovery payload size and processing

Question: 54

Based on this image, which of the following choices is true?

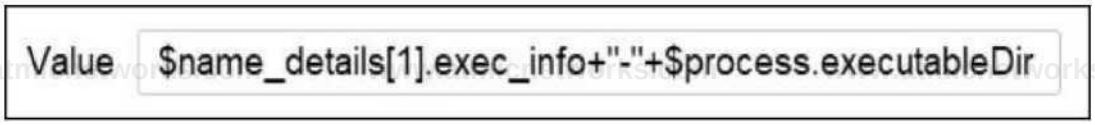

A. This is from a WMI query operation step.
B. There is a scalar variable labeled '1'.
C. This Value cannot be used in a pattern step.
D. There is a tabular variable named 'name_details'.

Answer: D

Explanation:
Based on the image, the value assigned to the label "Value" is a concatenation of variables and strings. One of the variables is $name_details[1].exec_info, which indicates that $name_details is a tabular variable, meaning an array or a map that can store multiple values1. The [1] is the index or the key that accesses the element of the tabular variable. The .exec_info is the attribute or the property of the element. Therefore, the choice D is true.
The other choices are not true based on the image. The choice A is false because the image does not show any WMI query operation step. A WMI query operation step is a pattern step that executes a WMI query on a Windows device and returns the results as a tabular variable2. The image does not show any WMI query syntax or result. The choice B is false because there is no scalar variable labeled '1'. A scalar variable is a variable that can store only one value1. The image shows a label "Value" but not a variable name. The choice C is false because the value can be used in a pattern step. A pattern step is a unit of logic that performs a specific action during Discovery or Service Mapping3. The value can be used as an input or an output of a pattern step, depending on the context.
Reference: 2: Define a WMI query 3: Pattern steps 1: Pattern variables

Question: 55

Using the SNMP Query operation on a pattern for a custom device query, it is best practice to do what?

A. Modify the default MIB information
B. Enable SSH as a secondary protocol
C. Use live devices in production
D. Use the publish manufacturer's device MIB

Answer: D

Explanation:

The SNMP Query operation on a pattern for a custom device query allows Discovery to retrieve information from the target device using the SNMP protocol. It is best practice to use the published manufacturer's device MIB, which is the official and standard definition of the device's SNMP data. Using the manufacturer's device MIB ensures that the query is accurate, consistent, and compatible with the device. Modifying the default MIB information, enabling SSH as a secondary protocol, or using live devices in production are not recommended practices, as they may introduce errors, security risks, or performance issues123.

Reference:

SNMP Query pattern operation - Product Documentation: San Diego - Now Support Portal Load a MIB file - Product Documentation: Vancouver - ServiceNow

How Loading of MIB file works on SNMP Browser (Pattern Step - SNMP Query) - Support and Troubleshooting

Question: 56

After running Discovery and viewing the ECC Queue tab, what are some of the displayed default fields? (Choose three.)

A. Details
B. Topic
C. Pattern log link
D. CMDB CI
E. Queue
F. Source
G. Discovery schedule name

Answer: B, E, F

Explanation:

The ECC Queue tab displays the input and output messages from and to MID Servers. The default fields that are displayed on this tab are:

Topic: The topic of the message, such as Discovery, Orchestration, or MID Server.

Queue: The queue name of the message, such as input, output, or error.

Source: The source of the message, such as the MID Server name, the instance name, or the probe name.

State: The state of the message, such as ready, processed, or ignored.

Created: The date and time when the message was created.

Updated: The date and time when the message was last updated.

Some other fields that can be added to the ECC Queue tab are:

Details: The details of the message, such as the payload, the response, or the error message.

Pattern log link: The link to the pattern log file, if the message is related to a pattern execution.

CMDB CI: The configuration item that is associated with the message, if any.

Discovery schedule name: The name of the discovery schedule that triggered the message, if any.

Reference: ECC Queue Processing and Debugging, with "Discovery - Sensors" used as an example, The ECC queue for Discovery

Question: 57

Hardware Models can have a one-to-many relationship with the following:

(Choose three.)

A. Assets
B. Manufacturer
C. Configuration Items
D. Product owner
E. Model Categories

Answer: A, C, E

Explanation:
Hardware Models are templates that define the attributes and default values for a specific type of hardware configuration item (CI). Hardware Models can have a one-to-many relationship with the following entities:
Assets: An asset is a tangible or intangible resource that is tracked and managed by the organization, such as a laptop, a license, or a contract. A Hardware Model can be associated with multiple assets that share the same characteristics and specifications.
Configuration Items: A CI is any component that needs to be managed in order to deliver an IT service, such as a server, an application, or a user. A Hardware Model can be associated with multiple CIs that share the same characteristics and specifications.
Model Categories: A model category is a classification of Hardware Models based on their function, such as network, storage, or server. A Hardware Model can belong to one or more model categories that define its scope and usage.
Reference:
Hardware Model
Asset
Configuration Item (CI)
[Model Category]

Question: 58

Which choice best represents how to modify the functionality shown in this image?

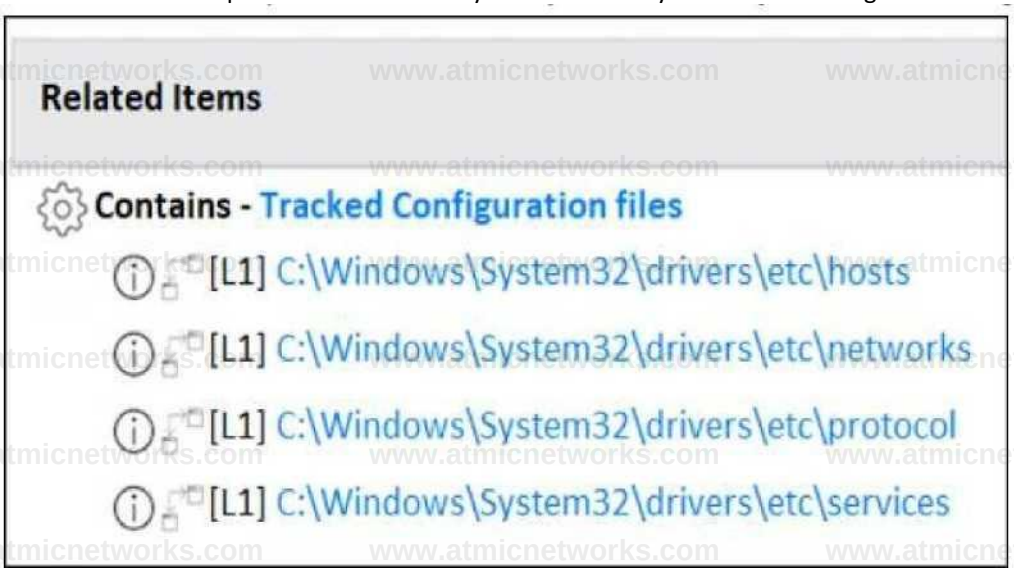

A. From a Classifier
B. From a Discovery Pattern
C. From the MID Server
D. From a Probe

Answer: B

Explanation:

The functionality shown in the image is related to the Tracked Configuration files feature of ServiceNow Discovery. This feature allows Discovery to track changes in configuration files on Windows and Unix hosts and store them as configuration items (CIs) in the CMDB1. To modify this functionality, such as adding or removing files to track, changing the file attributes to capture, or defining the CI relationships, one needs to edit the Discovery Pattern that is associated with the Tracked Configuration files classifier2. A Discovery Pattern is a series of operations that tell Discovery how to find and identify an application or a configuration file on a host3.

Reference:

1: Tracked Configuration files - Product Documentation: San Diego - ServiceNow

2: Tracked Configuration files pattern - Product Documentation: San Diego - ServiceNow

3: Patterns and horizontal discovery - Product Documentation: Tokyo - ServiceNow

Question: 59

What is the recommended method of consolidating duplicate CIs?

A. Duplicate CI Remediator

B. Event CI Remediation

C. Ignore Duplicate CI

D. Manual CI Remediation

Answer: A

Explanation:

The Duplicate CI Remediator is a feature of ServiceNow Discovery that helps to identify and merge duplicate CIs in the CMDB. It uses predefined rules and criteria to compare CIs across different data sources and determine if they are duplicates. It also provides a user interface to review and approve the suggested merges, or to manually merge CIs as needed. The Duplicate CI Remediator helps to improve the accuracy and quality of the CMDB data and reduce the maintenance efforts.

Reference:

Duplicate CI Remediator

Identify and merge duplicate CIs

Duplicate CI Remediator rules

Question: 60

During the Discovery process, what determines if an Asset record is created?

A. CMDB

B. Model Category

C. Model Product

D. ECC Queue

E. Configuration Item

Answer: B

Explanation:

During the Discovery process, the Model Category of the Configuration Item (CI) determines if an Asset record is created. A Model Category is a grouping of models that share the same asset tracking strategy and the same CI and Asset tables. The asset tracking strategy defines whether Discovery creates an Asset record for each CI of that Model Category. There are three possible asset tracking strategies: Create Assets, Don't create Assets, and Create Consumable Assets12.

Reference:

1: ServiceNow Discovery Documentation, Model Categories section

2: ServiceNow Support Article, Why do Assets get created when inserting a CI that has a Model that is set to Don't create Assets?

Question: 61

What pattern operation allows the transfer of a file from the MID Server to a target?

A. Parse file
B. Create Connection
C. Put file
D. Manage Attachments

Answer: C

Explanation:

The Put file operation allows the transfer of a file from the MID Server to a target1. This operation can be used to copy files from the MID Server to a target device, such as configuration files, scripts, or binaries2. The Put file operation requires the file path on the MID Server and the file path on the target as inputs2.

Reference:

1: Put file
2: Define discovery steps

Question: 62

On the ECC Queue, sensor records have a Queue value of and probe records have a Queue value of .

A. input, output
B. started, ready
C. ready, started D. output, input

Answer: D

Explanation:

On the ECC Queue, sensor records have a Queue value of output and probe records have a Queue value of input. The ECC Queue is the connection point between the instance and the MID Server, and it displays the input and output messages from and to the MID Server1. A probe is a message that the instance sends to the MID Server to request information from a device or application2. A sensor is a script that the MID Server runs to process the results returned by the probe3. When a probe is created, it is inserted into the ECC Queue with a Queue value of input, which means that it is waiting to be picked up by the MID Server4. When a sensor is created, it is inserted into the ECC Queue with a Queue value of output, which means that it is waiting to be processed by the instance.

Reference: 1: ECC Queue - ServiceNow 2: Probes - ServiceNow 3: Sensors - ServiceNow 4: How the ECC Queue table records get processed: from output ready to input processed - Support and Troubleshooting - Now Support Portal : [ECC Queue Processing and Debugging, with “Discovery - Sensors” used as an example - Support and Troubleshooting - Now Support Portal]

Question: 63

In Discovery, what table associates an IP address and a credential?

A. Credential Affinity
B. Service Affinity
C. Service CI Association
D. Tags

Answer: A

Explanation:

The Credential Affinity table is a table that associates an IP address and a credential in Discovery1. It is used to

store the credential that was successfully used to access a device during Discovery2. It can also be used to manually assign a credential to an IP address or a range of IP addresses3. The other options are not tables that associate an IP address and a credential in Discovery. The Service Affinity table is a table that associates a business service and a credential. The Service CI Association table is a table that associates a business service and a configuration item (CI). The Tags table is a table that stores tags that can be applied to records in any table.

Reference: 1: Credential Affinity table 2: Credential affinity 3: Manually assign a credential to an IP address : Service Affinity table : Service CI Association table : Tags table

Question: 64

If the WMI service is not running on a host, it will prevent the discovery of which devices?

A. Network
B. Windows
C. Storage
D. Unix

Answer: B

Explanation:

The WMI service is a Windows service that enables remote management of Windows devices using the WMI protocol. The WMI service is required for Discovery to access the configuration and status information of Windows devices using WMI queries. If the WMI service is not running on a host, it will prevent the discovery of Windows devices, as Discovery will not be able to communicate with the WMI Collector or execute WMI queries on the target device123.

Reference:

How to resolve a failure to communicate with the WMI Collector during top-down discovery - Support and Troubleshooting - Now Support Portal

Windows discovery without 'domain admin' or 'local admin' privileges ? - Support and Troubleshooting - Now Support Portal

Discovery — ServiceNow Elite

Question: 65

In general, Discovery can provide which of the following kinds of application relationships? (Choose two.)

A. tcp to udp
B. application to application
C. mid server to target
D. host to application

Answer: B, D

Explanation:

Discovery can provide information about the relationships between configuration items (CIs), including applications and hosts. Application relationships are the connections between applications that run on different hosts or servers, such as web servers, databases, load balancers, etc. Host relationships are the connections between applications and the hosts or servers that they run on, such as Windows, Linux, Unix, etc. Discovery can identify these relationships using various methods, such as probes, sensors, patterns, file-based discovery, etc. Discovery can also create and update these relationships in the CMDB, using the Depends on::Used by or Runs on::Runs relationship types. Reference: Discovery Relationship Creation, Application Discovery Mapping Overview, IT Discovery – Application Dependency Mapping

Question: 66

Which of the following must be configured to allow a MID Server to access servers using WinRM? (Choose two.)

A. JEA Properties set to True
B. MID Servers must be configured as a trusted source with DNS
C. MID Servers need to be added to the WinRM Group policy on the Servers
D. MID Server Parameters Add WinRM

Answer: C D

Explanation:

To allow a MID Server to access servers using WinRM, the following configurations are required: MID Servers need to be added to the WinRM Group policy on the Servers: This enables the servers to accept WinRM connections from the MID Servers and grant them the necessary permissions to execute commands and scripts1.
MID Server Parameters Add WinRM: This enables the MID Server to use WinRM as a protocol for Discovery and Orchestration and specify the WinRM port, authentication method, and encryption settings2.
Reference:
Configure Windows servers for WinRM MID Server Parameters Add WinRM

Question: 67

For a pattern operation, which of the below choices could be a valid replacement for <_>? (Choose three.)
$IfTable<_>.InstanceID

A. [3]
B. [X]
C. [&]
D. [.]
E. [*]
F. []

Answer: A, D E

Explanation:

For a pattern operation, the <_> placeholder can be replaced by any valid index or wildcard character that can be used to access an element of an array or a table. The $IfTable variable is a table that contains information about the network interfaces of a device1. The InstanceID attribute is a unique identifier for each interface2. To access the InstanceID of a specific interface, the pattern operation can use the following syntax:

$IfTable[<index>].InstanceID where <index> is either:

A numeric index that corresponds to the position of the interface in the table, such as 3 for the third interface.

A dot (.) followed by a key-value pair that matches a specific attribute and value of the interface, such as [.Name=eth0] for the interface with the name eth0.

An asterisk () that represents a wildcard that matches any interface in the table, such as [] for all interfaces.

Reference:

1: IfTable - Product Documentation: San Diego - ServiceNow

2: InstanceID - Product Documentation: San Diego - ServiceNow

Question: 68

Which of the below choices are the most probable results of the following image? (Choose three.)

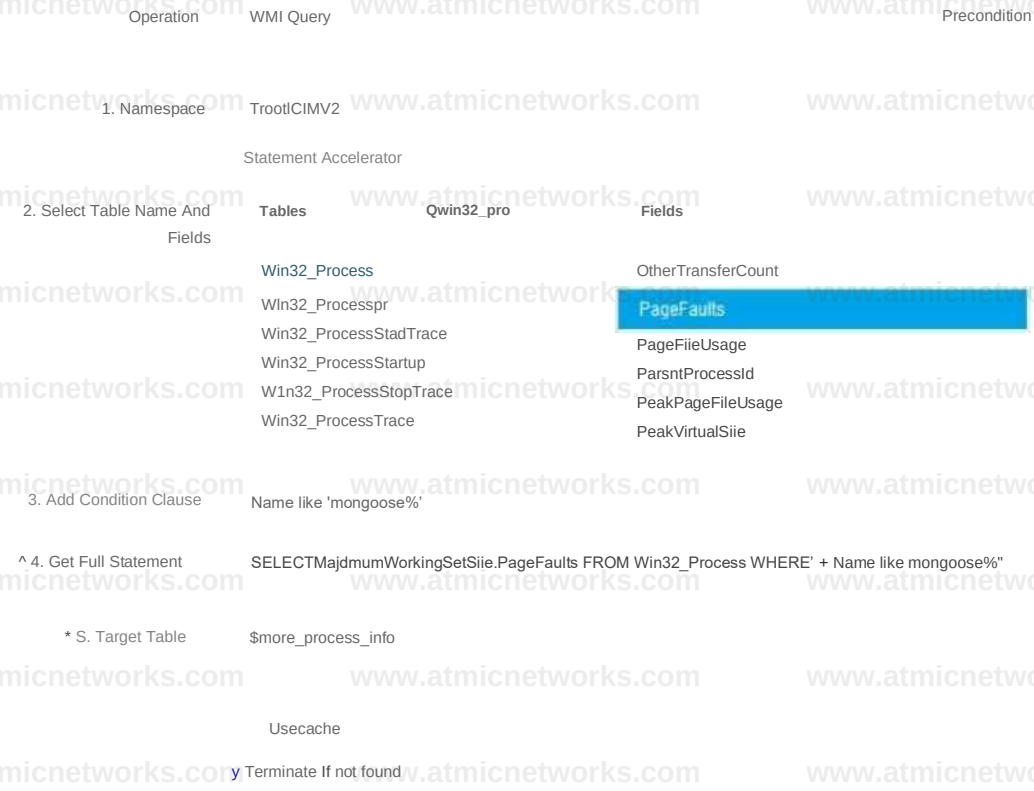

A. A tabular variable named ‘WMI Query’.
B. A scalar variable named ‘MaximumWorkingSetSize’
C. A scalar variable named ‘PageFaults’.
D. A scalar variable named ‘PeakVirtualSize’.
E. A tabular variable named ‘more_process_info’.
F. A tabular variable named ‘Win32_Process’.

Answer: A B D

Explanation:
The image shows a part of a horizontal pattern in Pattern Designer that uses the WMI Query operation to retrieve information about the processes running on a Windows server. The operation takes a WQL query as an input parameter and returns a tabular variable that contains the results of the query. The tabular variable is named ‘WMI Query’ by default, but it can be renamed by the user. The WQL query in the image selects four attributes from the Win32_Process class: Name, MaximumWorkingSetSize, PageFaults, and PeakVirtualSize. These attributes are stored as scalar variables in the tabular variable, and they can be accessed by using the dot notation, such as WMI Query.Name or WMI Query.MaximumWorkingSetSize. Therefore, the most probable results of the image are a tabular variable named ‘WMI Query’ and two scalar variables named ‘MaximumWorkingSetSize’ and ‘PeakVirtualSize’.

Reference:
Pattern Designer | WMI Query
Win32_Process class
WQL (SQL for WMI)

Question: 69

What is found on the far-right section of the Pattern Designer for horizontal patterns? (Choose two.) A. CI Attributes
B. Operations
C. Temporary Variables
D. CMDB Dashboard

Answer: A, C

Explanation:

The Pattern Designer is a graphical tool that allows users to create and edit horizontal discovery patterns. The Pattern Designer has three main sections: the left section shows the CI type and the pattern type, the middle section shows the pattern steps and the flow diagram, and the far-right section shows the details of the selected step or operation. The far-right section can display the following tabs, depending on the type of step or operation selected12:

Operations: This tab shows the list of available operations that can be added to the pattern, such as Execute Command, Parse File, Set Parameter Value, etc. Operations are the basic building blocks of a pattern that perform specific actions on the target device or the MID Server.

CI Attributes: This tab shows the list of attributes that can be set or updated for the CI type of the pattern, such as name, serial number, manufacturer, etc. CI attributes are used to populate the CI record in the CMDB with the information collected by the pattern.

Temporary Variables: This tab shows the list of temporary variables that can be created or used by the pattern, such as $result, $output, $user_var, etc. Temporary variables are used to store and manipulate data within the pattern, such as the output of a command, the value of a parameter, or the result of a calculation.

CMDB Dashboard: This tab shows the CMDB dashboard that displays the CI record and its related lists, such as relationships, identification rules, sensors, etc. The CMDB dashboard is used to preview and verify the outcome of the pattern on the CI record in the CMDB.

Reference:

1: ServiceNow Discovery Documentation, Pattern Designer section

2: ServiceNow Discovery Overview, page 11

Question: 70

Which of the choices below best represent key capabilities of a ServiceNow ITOM Enterprise solution? (Choose two.)

A. Create an Engaging User Experience
B. Build New Apps Fast
C. Manage Hybrid Clouds
D. Proactively Eliminate Service Outages

Answer: C, D

Explanation:

Based on the web search results from the search_web tool, the following choices best represent key capabilities of a ServiceNow ITOM Enterprise solution:

Manage Hybrid Clouds: ServiceNow ITOM enables organizations to gain visibility into their IT environment, on-premises to cloud, and see the impact of their applications on business services12. It also helps to plan, track, and secure cloud implementation through best-practice governance1. Four apps comprise Cloud Accelerate to help prepare for cloud migration, streamline provisioning requests, and manage cloud resources1.

Proactively Eliminate Service Outages: ServiceNow ITOM helps to predict issues, reduce user impact, and automate resolutions with predictive AIOps and automation13. It analyzes telemetry and log data using AI to reduce noise, expedite mapping and predict issues1. It also delivers proactive digital operations with real-time log data to solve problems quickly1.

The other choices are not specific to ServiceNow ITOM Enterprise solution, but rather general features of the ServiceNow platform. The ServiceNow platform allows users to create an engaging user experience with intuitive interfaces and workflows. It also enables users to build new apps fast with low-code

development tools and pre-built components.

Reference: 1: ITOM - Enterprise IT Operations Management - ServiceNow 2: ServiceNow IT Operations Management (ITOM) Services | Jade 3: Modernize IT Operations Management with ServiceNow ITOM : ServiceNow Platform

Question: 71

Which choices best describe what is necessary to create a custom horizontal pattern to discover an operating system that is not discovered by the base installation patterns? (Choose two.)

A. Select a CI Type
B. Define Process Strategy
C. Select Infrastructure Pattern Type
D. Select Application Pattern Type

Answer: C D

Explanation:

According to the ServiceNow website1, the key capabilities of a ServiceNow ITOM Enterprise solution are:

Manage hybrid clouds: Gain visibility and control across your hybrid cloud environment with a single system of action for IT.

Proactively eliminate service outages: Identify issues and anomalies before they cause service outages with AIOps and machine learning.

Automate IT service operations: Automate common service requests, processes, and tasks with digital workflows and virtual agents.

Optimize IT service delivery: Align your IT resources and investments with business priorities and optimize service delivery performance.

Reference:

1: ITOM - Enterprise IT Operations Management - ServiceNow

Question: 72

Which of the following best describes the relationship between the Tomcat [cmdb_ci_app_server_tomcat] table and the Application Server [cdmb_ci_app_server] table? (Choose two.)

A. Tomcat does not extend the Application Server table
B. Tomcat table extends the Application Server table
C. Tomcat table is a child of the Application Server table
D. Tomcat table is a parent of the Application Server table

Answer: B, C

Explanation:

Question: 73

In a discovery pattern, which types are available with CI Attributes in the Pattern Designer? (Choose two.)

A. Global CI types
B. Main pattern CI type
C. Related CI types
D. All CI types

Answer: B, C

Explanation:

In a discovery pattern, CI Attributes are the fields that store the information about the configuration items (CIs) that are discovered by the pattern. The Pattern Designer is a graphical interface that allows users to create and edit patterns using drag-and-drop operations. In the Pattern Designer, there are two types of CI Attributes that are available: Main pattern CI type and Related CI types. The Main pattern CI type is the CI type that the pattern is designed to discover, such as Windows Server, Linux Server, or Database. The Related CI types are the CI types that are related to the main CI type, such as IP Address, Network Adapter, or Software Installation. Users can select the CI Attributes from these types to populate the CMDB tables with the data collected by the pattern123.

Reference:

Patterns and horizontal discovery - Product Documentation: Tokyo - Now Support Portal
Pattern Designer - Product Documentation: Vancouver - ServiceNow
Discovery pattern designer - ServiceNow IT Operations Management [Book]

Question: 74

Which of the below choices are benefits of Tracked Configuration Files? (Choose two.)

A. Content version comparison
B. Files tracked as CIs
C. Unwanted files removed from target
D. No credentials needed

Answer: A, B

Explanation:

Tracked configuration files are files that contain settings and parameters of certain applications or devices that Discovery can find and add to the CMDB. Some of the benefits of tracked configuration files are:

Content version comparison: You can compare the content of different versions of the same configuration file and see the changes that occurred over time. This can help you troubleshoot issues, audit changes, and restore previous configurations if needed.

Files tracked as CIs: You can view the configuration files as configuration items (CIs) in the CMDB and see their relationships with other CIs, such as the applications or devices they belong to. This can help you understand the dependencies and impact of the configuration files on your IT environment.

Reference: Configuration file tracking, Modify tracking changes in configuration files

Question: 75

Which choice best describes a Functionality Definition?

A. Defines what CI identifiers to use.
B. Defines the IP addresses to discover.
C. Defines what Behavior to use from a Discovery Schedule.
D. Defines what protocols to detect from within a Behavior.

Answer: C

Explanation:

A Functionality Definition is a configuration record that defines what protocols to detect from within a Behavior. A Behavior is a set of Functionality Definitions that Discovery uses to identify and explore a device or application. A Functionality Definition specifies the order, type, and parameters of the probes that Discovery launches to gather information about a device or application12.

Reference:

Functionality Definition

Behavior

Question: 76

As a first step in horizontal discovery, which of the following is where the Shazzam probe is placed in a request?

A. Target
B. Pattern Log
C. ECC queue
D. Discovery Log

Answer: C

Explanation:

The Shazzam probe is placed in the ECC queue as a first step in horizontal discovery. The ECC queue is a table that stores requests and responses for communication between the MID Server and the ServiceNow instance1. The Shazzam probe is a script that performs port scanning on a target host to determine what protocols are available and what classifiers to trigger2. The Shazzam probe is sent from the ServiceNow instance to the MID Server, which executes it on the target host and returns the results to the ECC queue3.

Reference:

1: ECC queue - Product Documentation: San Diego - ServiceNow
2: Shazzam probe, port probes, and protocols - Product Documentation: San Diego - ServiceNow
3: Discovery Phase Shazzam - Support and Troubleshooting - ServiceNow er: D

Question: 77

Which of the choices have a higher chance of leading to lost data during CI reclassification? (Choose two.)

A. Switching
B. Identifying
C. Downgrading
D. Upgrading

Answer: A, C

Explanation:

CI reclassification is the process of changing the class of a CI in the CMDB based on the discovery data. It can be done manually or automatically, depending on the configuration settings. CI reclassification can result in data loss if the new class has fewer or different attributes than the original class, and the existing attribute values are not preserved or mapped to the new class. This can happen when a CI is switched to a different class hierarchy, such as from a Linux Server to a Windows Server, or when a CI is downgraded to a less specific class, such as from a Database to a Software Instance. Therefore, switching and downgrading have a higher chance of leading to lost data during CI reclassification.

Reference:

CI reclassification
Reclassify CIs
Configure automatic CI reclassification

Question: 78

Which of the below choices are kinds of variables used in discovery patterns? (Choose three.)

A. CI attributes

B. Prefix
C. Temporary
D. Fixed
E. Global

Answer: A, C, E

Explanation:

Variables are used in discovery patterns to store and manipulate data that is collected or calculated during the discovery process. There are three kinds of variables used in discovery patterns12: CI attributes: These are variables that correspond to the attributes of the configuration item (CI) type that the pattern is discovering. For example, name, serial_number, manufacturer, etc. CI attributes are used to populate the CI record in the CMDB with the information gathered by the pattern. Temporary: These are variables that are created and used within the pattern, but are not stored in the CMDB. For example, $result, $output, $user_var, etc. Temporary variables are used to store and manipulate data such as the output of a command, the value of a parameter, or the result of a calculation.

Global: These are variables that are defined in the global scope and can be accessed by any pattern. For example, $mid, $target, $ip_address, etc. Global variables are used to store and access data that is common to all patterns, such as the MID Server, the target device, or the IP address. Reference:

1: ServiceNow Discovery Documentation, Pattern Variables section

2: ServiceNow Discovery Overview, page 11

Question: 79

CI identifiers can be viewed under which of the following?

A. Discovery Dashboard
B. CI Class Manager
C. Process Handlers
D. Processes and Classification
E. CI record

Answer: E

Explanation:

CI identifiers can be viewed under the CI record1. A CI identifier is a set of attributes that uniquely identify a CI of a specific class2. CI identifiers are used by the identification and reconciliation engine (IRE) to match incoming data with existing CIs in the CMDB3. CI identifiers can be viewed and edited in the CI record by clicking the Identification Information related link1.

Reference:

1: View and edit CI identifiers

2: Create or edit a CI identification rule

3: Identification rules

Question: 80

Which of the following executes the osquery commands on agents to gather attribute details from a CI?

A. Agent Collector
B. Agent listener
C. ACC for Discovery

D. Check Definitions

E. Policies

Answer: A

Explanation:

Agent Collector is the component that executes the osquery commands on agents to gather attribute details from a CI. Agent Collector is a feature of the Agent Client Collector (ACC) for Discovery, which is an integration that enables ServiceNow Discovery to use osquery to collect data from Linux and Windows devices that have the Elastic Agent installed1. Agent Collector allows you to define osquery queries and assign them to policies that target specific devices or groups of devices2. Agent Collector then runs the queries on the agents and sends the results back to the instance, where they are processed by the Agent Listener and stored in the CMDB3.

Reference: 1: Agent Client Collector for Discovery - ServiceNow 2: Create an Agent Collector policy - ServiceNow 3: Agent Client Collector for Discovery architecture - ServiceNow

Question: 81

Which one of the following is not used in a horizontal discovery pattern?

A. Variables

B. Connectivity section

C. Operations

D. Identification section

Answer: B

Explanation:

A horizontal discovery pattern is a sequence of operations that collects information about a specific type of configuration item (CI), such as an operating system, a database, or a web server1. A horizontal discovery pattern does not use the connectivity section, which is only used for the application pattern type2. The connectivity section defines how to connect to a CI and run commands on it3. A horizontal discovery pattern uses the following sections1:

Variables: Defines the variables that store the data collected by the pattern and map them to the target table fields.

Identification: Defines the criteria to identify a CI and update or create a record in the CMDB. Operations: Defines the actions to perform on the target CI, such as running commands, querying databases, or parsing files.

Reference: 1: ServiceNow Docs - Horizontal discovery patterns 2: ServiceNow Docs - Application pattern types 3: ServiceNow Docs - Connectivity section

Question: 82

What would you see in the Discovery pattern log that you would not see in the ECC Queue?

A. Success or failure of the individual pattern steps

B. Payload of CI attributes and relationships

C. Shazzam probe payload

D. Relationships created

Answer: A

Explanation:

The Discovery pattern log is a log that records the execution of a pattern during Discovery or Service

Mapping1. It shows the success or failure of the individual pattern steps, as well as the input and output variables, the CI attributes and relationships, and the errors or warnings2. The ECC Queue is a table that displays the input and output messages from and to MID Servers3. It shows the payload of CI attributes and relationships, the Shazzam probe payload, and the relationships created, but not the success or failure of the individual pattern steps4. The success or failure of the individual pattern steps is specific to the Discovery pattern log and can be used to debug or troubleshoot the pattern execution2.

Reference: 1: Discovery pattern log 2: View the Discovery pattern log 3: ECC Queue 4: ECC Queue Processing and Debugging, with "Discovery - Sensors" used as an example

Question: 83

Which of the following choices may be global variables for steps in horizontal discovery patterns? (Choose two.)

A. system
B. computer_system
C. process
D. baseline

Answer: A B

Explanation:

Global variables are variables that are automatically available for all steps in horizontal discovery patterns. They store information about the target CI, such as its IP address, host name, operating system, and credentials. There are two global variables that may be used for steps in horizontal discovery patterns: system and computer_system. The system variable contains the information about the target CI's IP address, host name, and credentials. The computer_system variable contains the information about the target CI's operating system, such as its name, version, and architecture12.

Reference:

Patterns and horizontal discovery - Product Documentation: Tokyo - Now Support Portal
Horizontal Pattern probe - Product Documentation: San Diego - Now Support Portal

Question: 84

One method for deleting specific CIs not discovered in 30 days is:

A. Scheduled Job
B. UI Policy
C. Service Mapping
D. Data Policy

Answer: A

Explanation:

One method for deleting specific CIs not discovered in 30 days is to use a scheduled job that runs a script to query the CMDB for CIs that have not been updated by Discovery for a certain period of time, and then delete them or mark them as retired. This method can help you keep your CMDB clean and accurate, and avoid having stale or obsolete CIs. You can configure the scheduled job to run at a specific frequency and time, and specify the criteria for selecting the CIs to delete.

Reference: Automated way to disable/retire CIs not updated for 45 days, ServiceNow CIS-Discovery Sample Questions

Question: 85

What is the default thread count for a MID Server?

A. 5

B. 1

C. 50

D. 25

Answer: D

Explanation:

The default thread count for a MID Server is 25. This means that the MID Server can run up to 25 worker threads simultaneously to execute the jobs assigned to it by the ServiceNow instance. The worker threads are divided into different thread pools based on the priority of the jobs. The MID Server can also adjust the thread count dynamically based on the system resources and the workload. The thread count can be configured by changing the MID Server parameters in the ServiceNow instance12. Reference = 1: MID Server Max Threads - Worker Groups - Priority and Queues - ServiceNow 2: Product Documentation | ServiceNow

Question: 86

Which of the following are contained in an extension section of a discovery pattern?

A. Connection sections

B. Network libraries

C. Identification sections

D. Shared libraries

Answer: D

Explanation:

An extension section of a discovery pattern is a section that allows you to modify or extend the functionality of an existing pattern without customizing it. An extension section can contain shared libraries, which are reusable code snippets that can be invoked from any pattern or extension section. Shared libraries can be used to define common functions, variables, or constants that are used across multiple patterns or extension sections. Shared libraries can also be used to override the default behavior of some built-in functions, such as identification or classification12. Reference = 1: ITOM: Extending Discovery/Service Mapping Patterns - GlideFast ServiceNow 2: How to Create Discovery Pattern Extensions and Why to Use Them

Question: 87

Which part of Agent Client Collector must be configured to run osquery commands on a CI? Policies

A. Credential-less Discovery

B. Check

C. ACC Websocket Endpoint

D. Infrastructure Patterns

Answer: C

Explanation:
A check is a part of Agent Client Collector that defines the osquery commands to run on a CI and the frequency of execution. A check can also specify the conditions for generating alerts or events based on the osquery results.

Reference
1: Agent Client Collector - Product Documentation: San Diego - ServiceNow
2: Agent Client Collector for Visibility - Product Documentation: San Diego - ServiceNow

Question: 88

In a pattern operation, which of the following correctly calls the value of the executableDir variable from the tabular process variable?

A. process(executableDir)
B. "process_executableDir"
C. $process_executableDir
D. $process.executableDir

Answer: D

Explanation:
The $process.executableDir expression correctly calls the value of the executableDir variable from the tabular process variable. The process variable is a tabular variable that contains the output of the ps command, which lists the processes running on a CI. The executableDir variable is a column name that holds the directory path of the executable file for each process. To access the value of a column name from a tabular variable, you need to use the dot notation, such as $variable.columnName12.

Reference
1: Pattern variables - Product Documentation: Tokyo - Now Support Portal
2: Examples of EVAL scripts used in Discovery patterns - ServiceNow

Question: 89

The deletion strategy is set to 'Mark as absent' for related Disk Cls discovered via the Linux Server pattern. If a related Disk Cl is discovered during the Linux Server discovery and then the same related Disk CI is not found the next time the Linux Server is discovered, the following will happen:

A. The Linux Server CI Status (install_status) is set to Absent.
B. The Disk Status CI (install_status) is set to Absent.
C. The Disk CI Operational Status (operational_status) is set to Non-Operational.
D. The Linux Server CI Operational Status (operational_status) is set to Non-Operational.

Answer: B

Explanation:
The deletion strategy 'Mark as absent' means that when a related CI that was previously discovered by a pattern is no longer found, the install_status attribute of that CI is set to Absent. This indicates that the CI is missing or removed from the network. The main CI (in this case, the Linux Server CI) is not affected by the deletion strategy of the related CI (in this case, the Disk CI)12.

Reference

1: Set a deletion strategy - Product Documentation: San Diego - ServiceNow

2: Network adapter status setting to “Absent” in EMC … - ServiceNow …

Question: 90

While discovering a new SNMP network device, which choice could cause the error ‘Active, couldn't classify" to occur on a Discovery Status?

A. SNMP credentials are incorrect.

B. SNMP only behavior is not configured.

C. A firewall blocking the communication between the MID Server and the target device.

D. An SNMP Classification must be updated or created for the new device.

Answer: D

Explanation:

In ServiceNow Discovery, when encountering the error 'Active, couldn't classify' during the discovery of a new SNMP network device, it typically indicates an issue with the classification process.
Classification is a crucial initial step in the Discovery process, where the system determines the type of device it is interacting with. If an SNMP Classification rule for the new device type does not exist or needs updating, Discovery will be unable to classify and thus properly discover the device. Incorrect SNMP credentials or firewall issues usually result in different types of errors, such as authentication failures or no response from the device.

Question: 91

What do most Discovery properties start with?

A. discovery.prop

B. glide.discovery

C. disco.release

D. glide.itom

Answer: B

Explanation:
In ServiceNow Discovery, most properties related to the Discovery module start with the prefix 'glide.discovery'. This naming convention is part of ServiceNow's system property design, where 'glide' typically indicates system properties and 'discovery' specifies that the property is related to the Discovery module.
Properties are used to configure various aspects of the Discovery process, such as behavior, timeouts, and performance settings.

Question: 92

From Pattern Designer, which horizontal pattern type is the image below showing?

Temporary Variables

► computer_system

A. Service Mapping

B. Infrastructure

C. Application

D. Computer System

Answer: D

Explanation:
The Pattern Designer in ServiceNow Discovery uses different types of horizontal patterns to model various IT components and services. The temporary variable 'computer_system' seen in the pattern indicates that this is a 'Computer System' pattern type. This type of pattern is used to model and discover details about computer systems, which typically include hardware, operating system, and other related components.

Question: 93

What is the advantage of Discovery Range Sets?

A. Range Sets provide flexibility in management and identification of known networks for simplicity of administration.
B. Range Sets are the only way to have more than one IP Range defined within a Discovery Schedule.
C. Range Sets show the number of IPs in a subnet.
D. All the necessary Range Sets are installed in the base installation of Discovery.

Answer: A

Explanation:
Discovery Range Sets in ServiceNow Discovery offer a flexible way to manage and identify networks. They simplify the administration of network discovery by allowing administrators to define groups of IP ranges. This organization helps in effectively managing and targeting different network segments without needing to define each IP range in individual Discovery Schedules. Range Sets do not necessarily show the number of IPs in a subnet, nor are they the only way to define multiple IP ranges in a schedule.

Question: 94

Which of the following describes the recommended permission level for credentials to discover Windows Servers?

A. A domain user with local administrator access
B. A user r00t access across the domain
C. A domain administrator with sudo access
D. A standard domain user with read access

Answer: A

Explanation:
For discovering Windows Servers, ServiceNow Discovery recommends using credentials that have a domain user account with local administrator access on the target servers. This level of access ensures that Discovery can successfully execute the necessary probes and sensors to collect comprehensive data about the server without being restricted by permission limitations. The other options either provide excessive permissions (like a domain administrator with sudo access) or insufficient privileges (like a standard domain user with read access).

Question: 95

During the Port Scan phase what could the Warning level error "No results returned from probe." mean?

A. Process Classifier incorrectly configured
B. Logical or physical firewall preventing connectivity
C. Wrong credentials
D. Incorrect Datasource Precedence

Answer: B

Explanation:
During the Port Scan phase in ServiceNow Discovery, the warning level error "No results returned from probe" typically indicates a connectivity issue. This error often arises when a logical or physical firewall is blocking the connectivity between the MID Server and the target device. This prevents the probe from successfully scanning the ports and returning results. The other options, such as incorrect process classifier configuration, wrong credentials, or incorrect datasource precedence, would generally lead to different types of errors or issues in the discovery process. Reference = ServiceNow Discovery documentation, particularly the sections that deal with troubleshooting Port Scan issues and understanding common errors during Discovery phases.

Question: 96

Which of the choices are types of temporary variables in a discovery pattern?
Choose 2 answers

A. Scalar/List
B. Command
C. SQL statement
D. Tabular/Table

Answer: AD

Explanation:
In the context of ServiceNow Discovery patterns, the types of temporary variables include Scalar/List and Tabular/Table. Scalar/List variables are used to store single values or lists of values, while Tabular/Table variables are used to store structured data in a table format. Command and SQL statement are not types of temporary variables; rather, they are types of operations or steps that can be used within a pattern to execute commands or SQL queries. Reference = ServiceNow Discovery documentation, especially sections related to Discovery Patterns and the use of temporary variables within patterns.

Question: 97

Which of the below choices are possible options under ACTION ON ALL in the Recommended Actions pane for an Automated Error Messages list within Discovery > Home?
Choose 2 answers

A. Make Ranges
B. View instructions
C. Ping IP Addresses
D. Retry Discovery

Answer: BD

Explanation:
In the ServiceNow Discovery module, under the Recommended Actions pane for an Automated Error Messages list within Discovery > Home, the possible options under ACTION ON ALL include "View instructions" and "Retry Discovery". "View instructions" allows users to see guidance on how to address the errors, while "Retry Discovery" enables users to reinitiate the discovery process for the affected devices. "Make Ranges" and "Ping IP Addresses" are not standard options in this context. Reference = ServiceNow Discovery documentation and user guides, particularly those discussing the management of automated error messages and recommended actions in Discovery.

Question: 98

What are the main KPIs for CMDB Health scorecard?

Choose 3 answers

A. Correctness
B. Completeness
C. Staleness
D. Compliance
E. Duplicates

Answer: ABC

Explanation:

The main Key Performance Indicators (KPIs) for the CMDB Health scorecard in ServiceNow include Correctness, Completeness, and Staleness. These KPIs are critical for assessing the overall health and quality of the CMDB. Correctness refers to the accuracy of the data in the CMDB, Completeness measures whether all required CI data is present, and Staleness indicates the freshness or timeliness of the CI data. Compliance and Duplicates, while important, are not considered main KPIs in the context of the CMDB Health scorecard. Reference = ServiceNow CMDB documentation and CMDB Health Dashboard user guide, focusing on the key performance indicators for CMDB health.

Question: 99

Which of the following fields are editable from a Merge Table pattern operation?

Choose 3 answers

A. Target Table
B. Second Table
C. Primary Table
D. First Table

Answer: ABD

Explanation:

In a Merge Table pattern operation within ServiceNow Discovery, the editable fields typically include the Target Table, Second Table, and First Table. These fields are configurable and allow users to specify the tables involved in the merge operation. The Target Table is where the merged data will be stored, while the First and Second Tables are the source tables whose data is being merged. The Primary Table is not typically an editable field in this context. Reference = ServiceNow Discovery documentation, particularly the sections discussing pattern design and operations, including Merge Table operations within Discovery Patterns.

Question: 100

Which of the following best describes what may also be required when increasing the max number of threads for a MID Server?

A. increase the amount of disk space on the MID Server host
B. updating the acl rules for the CMDB tables in ServiceNow
C. increasing the memory allocated to the MID Server app.
D. increase the MID Server max payload size

Answer: C

Explanation:

When increasing the maximum number of threads for a MID Server in ServiceNow Discovery, it's often necessary to also increase the memory allocated to the MID Server application. This is because more threads

can result in higher memory usage as each thread may consume resources while performing discovery tasks. Increasing disk space, updating ACL rules for CMDB tables, or increasing the MID Server max payload size are not directly related to the number of threads used by a MID Server. Reference = ServiceNow Discovery documentation, especially sections discussing MID Server configuration and performance tuning.

Question: 101

What operation is shown in this image? (has image took pic)

A. Transform Table
B. Put File
C. Union Table
D. Parse Variable

Answer: D

Explanation:

Without seeing the specific image, it's challenging to definitively identify the operation. However, based on the options provided, if the image shows an operation that involves interpreting or extracting data from a variable, then "Parse Variable" is the most likely answer. "Parse Variable" is used in ServiceNow Discovery patterns to analyze and extract data from variables collected during the discovery process. Reference = ServiceNow Discovery documentation, particularly the sections on Discovery Pattern Designer and the various operations that can be used in patterns.

Question: 102

How do you create relationships between Cis in a horizontal pattern?

A. By using the Create Relation/Reference operation.
B. Relationships are created automatically for each successful connection section.
C. Relationships cannot be created via patterns.
D. By using the Create Connections operation.

Answer: A

Explanation:

In ServiceNow Discovery, relationships between Configuration Items (CIs) in a horizontal pattern are created using the "Create Relation/Reference" operation. This operation allows the pattern to define and establish relationships or references between discovered CIs, based on the information gathered during the discovery process. Options B and C are incorrect as relationships are not automatically created nor is it true that they cannot be created via patterns. Option D, "Create Connections," is not a standard operation in this context. Reference = ServiceNow Discovery documentation, specifically the sections on creating and managing horizontal discovery patterns.

Question: 103

What is the recommended method for excluding specific software from discovery on Windows and Unix Servers?

A. Pattom Modification
B. Probe Modification
C. Configuration Console
D. Discovery Properties

Answer: B

Explanation:

The recommended method for excluding specific software from discovery on Windows and Unix Servers in

ServiceNow Discovery is by modifying the probes. Probes are the components in ServiceNow that collect data from devices. By modifying the probes, specific software can be excluded from the discovery process. Options A, C, and D are less direct methods for achieving this specific requirement. Reference = ServiceNow Discovery documentation, particularly in the sections related to probe configuration and modification.

Question: 104

From the Discovery Status record, which of the following does a Device record contain?

A. contains a link to the discovered CI
B. the contents of the discovered CI
C. contains a link to the discovered Asset
D. the CMDB Health Dashboard link

Answer: A

Explanation:

From the Discovery Status record, a Device record in ServiceNow Discovery contains a link to the discovered Configuration Item (CI). This link provides direct access to the CI record that was discovered and processed during the discovery cycle. The Device record does not contain the contents of the CI, a link to the discovered Asset, or a link to the CMDB Health Dashboard. Reference = ServiceNow Discovery documentation, focusing on the details of Discovery Status records and the information contained within Device records.

Question: 105

When configuring the MID Server Service Settings, which of the following is not a valid MID Server name?

A. mymidserverl
B. mid.serverl
C. midserverl
D. mid serverl

Answer: D

Explanation:

In configuring the MID Server Service Settings in ServiceNow, the MID Server name must adhere to certain naming conventions. Among the options provided, "mid server1" (D) is not a valid MID Server name because it contains a space. MID Server names should not include spaces or special characters that are not typically used in naming conventions. Names like "mymidserver1," "mid.server1," and "midserver1" are more in line with the standard naming practices. Reference = ServiceNow MID Server documentation, specifically the sections detailing MID Server configuration and naming conventions.

Question: 106

Which choice best describes the Reconciliation process?

A. The process of reconciling a deleted CI, if it is identified as a duplicate. The result is stored in de-duplication tasks.
B. The process of reconciling uniquely identifying Cis. to determine if the CI already exists in the CMDB or if it is a newly discovered CI.
C. The process of reconciling Cis and CI attributes by allowing only designated authoritative data sources to write to the CMDB at the CI table and attribute level.
D. The process of normalizing applications, to allow the data to be used by Software Asset Management.

Answer: C

Explanation:

The reconciliation process in ServiceNow refers to the method of ensuring data integrity and consistency in the Configuration Management Database (CMDB) by allowing only designated authoritative data sources to write or update information at the CI (Configuration Item) table and attribute level. This process helps maintain accurate and reliable data in the CMDB by controlling how and from where data is sourced and updated. The other options describe different processes or aspects of CMDB management and are not directly related to the reconciliation process. Reference = ServiceNow CMDB documentation, particularly the sections on CI data management and reconciliation processes.

Question: 107

Which of the following related lists can assist with troubleshooting discovery from a discovery status?
Choose 3 answers

A. Discovery Log
B. Running Processes
C. ECC Queue
D. Devices

Answer: ACD

Explanation:

To troubleshoot discovery issues from a discovery status in ServiceNow, the following related lists can be particularly useful:

Discovery Log - Provides detailed logs of the discovery process, helping to identify errors or issues encountered during discovery.

ECC Queue - Shows the messages and data passed between the ServiceNow instance and the MID Server, offering insights into the communication and data exchange during discovery.

Devices - Lists the devices discovered in the process, allowing for a direct view of the results and any potential issues with specific devices.

The "Running Processes" list is not typically used for troubleshooting discovery processes as it relates more to the operational aspects of the server or device rather than the discovery process itself. Reference = ServiceNow Discovery documentation, specifically the sections on Discovery troubleshooting and analysis, and the use of related lists for monitoring and troubleshooting purposes.

Question: 108

Based on the image, which of the following is true?

VAILIS Sfidrrf de^itsill €MC^ ififo* * "Sp^ocBSS.c^^cutJC't&Of

Value $name_details[1].exec_info+"-"+$process.executableDir

A. There is a tabular varlable named 'name_details'.
B. This is from a WMI query operation step.
C. There is a scalar varlable labeled '1'.
D. This Value cannot be used in a pattem step.

Answer: C

Explanation:

The image provided shows a ServiceNow expression used to concatenate two pieces of information. The $name_details[1] suggests that 'name_details' is an array or a list variable, and we're accessing the second element (considering indexing starts at 0). This indicates that 'name_details' is not a tabular variable but rather a list or an array, and there is a scalar value labeled '1' which is being used
to index this array. The 'exec_info' and 'executableDir' are likely attributes or properties of the objects within the 'name_details' array and '$process' object, respectively. There is no indication that this is from a WMI query operation step, nor is there any inherent reason why this value could not be used in a pattern step. ServiceNow Discovery patterns can use such expressions to extract and concatenate data during the discovery process.

Question: 109

Which of the following must be configured to allow a MID Server to access servers using WinRM?

A. Add the WinRM parameter on the target servers
B. Add the MID Servers as a trusted source with DNS
C. Add a MID Server parameter mid.windows.management protocol to use WinRM
D. Set the JEA Property to True

Answer: C

Explanation:

For a MID Server to access servers using Windows Remote Management (WinRM), it's necessary to specify that WinRM should be used as the protocol for Windows management. This is typically done by setting the appropriate MID Server parameter, such as mid.windows.management_protocol, to WinRM. This configuration tells the MID Server to use the WinRM protocol when communicating with the target servers for executing commands and collecting data. The answer is derived from understanding how ServiceNow Discovery uses MID Servers to perform remote operations and the need to configure these MID Servers appropriately to communicate using different protocols like WinRM. It is not sufficient to just add the WinRM parameter on target servers or simply add the MID Servers as a trusted source with DNS; the MID Server itself must be configured to use WinRM. Setting a 'JEA Property' to true is not relevant to the basic configuration of a MID Server for WinRM access.

Question: 110

Which of the following are the results of executing the WMI Query?

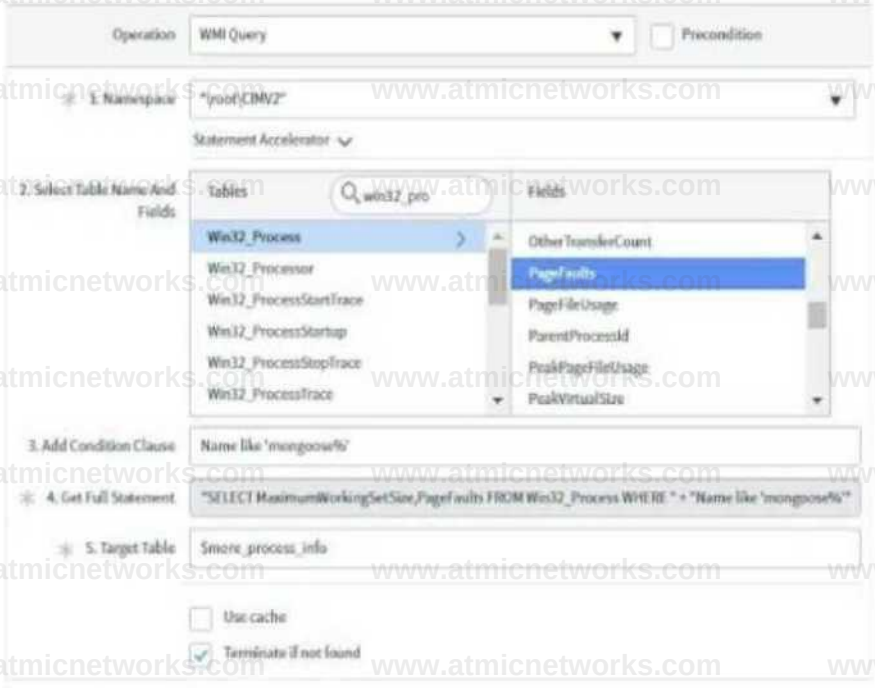

Choose 3 answers

A. A scalar variable named 'PageFaults'.
B. A tabular variable named 'more_process_info'.
C. A scalar variable named 'PeakVirtualSize'.
D. A tabular variable named 'Win32 Process'.
E. A scalar variable named 'MaximumWorkingSetSize'.
F. A tabular variable named 'WMI Query',

Answer: BCE

Explanation:

Question: 111

Which service needs to be running on a host to detect a Windows device?

A. WMI
B. SSH
C. WinRM
D. CIM

Answer: A

Explanation:

For ServiceNow Discovery to detect a Windows device, the Windows Management Instrumentation (WMI) service needs to be running on the host. WMI is a Windows feature that provides a standardized way to access management information about operating system components and services. Discovery uses WMI to query Windows hosts for information about their hardware,

software, and system configuration. SSH is used for Unix-like systems, WinRM is an alternative Windows management protocol, and CIM (Common Information Model) is a cross-platform standard; however, WMI is the primary service used for Windows device discovery.

Question: 112

Which choice best describes a horizontal discovery pattern?

A. Classifiers that execute probes
B. Steps that execute operations
C. Credential depot
D. Port scanning tool

Answer: B

Explanation:
In ServiceNow Discovery, a horizontal discovery pattern refers to the logical set of steps that the Discovery process executes to find CI information across the network. Each step in a horizontal discovery pattern performs specific operations such as running probes, executing commands, or querying databases to gather information about the configuration items (CIs). It's about expanding the breadth of discovery across devices and applications, rather than going in-depth into a specific CI's details, which would be vertical discovery. Classifiers, credential depots, and port scanning tools are part of the discovery process but do not describe the horizontal discovery pattern itself.

Question: 113

Which of the following MID Server Parameters cannot be modified through the platform UI?

A. mid.ssh.session_timeout
B. url
C. mid.windows.management_protocol
D. name

Answer: D

Explanation:
The MID Server parameters like 'mid.ssh.session_timeout', 'url', and 'mid.windows.management_protocol' can be modified through the ServiceNow platform UI under the MID Server configuration settings. These parameters control various aspects of the MID Server's operation, such as timeout settings for SSH sessions, the URL for the MID Server to communicate with the ServiceNow instance, and the management protocol to use for Windows servers, respectively. The 'name' of the MID Server, however, is not typically a parameter that can be changed through the UI as it is a unique identifier for the MID Server within the ServiceNow instance. Once a MID Server is registered with a name, that name is usually persistent and not intended to be modified through configuration settings.

Question: 114

Which service needs to be running on a host to detect a Windows device?

A. WM

B. WinkM
C. CM
D. SSH

Answer: B

Explanation:

Question: 115

For the Set Parameter Value operation, which of the following is used in the syntax to declare a constant, unchanging value?

A. Quotes
B. Brackets
C. Dollar sign
D. Hash tag

Answer: A

Explanation:

Question: 116

Which must be configured to allow a MID Server to access servers using WinRM?

A. Add the WinRM parameter on the target servers
B. Add a WinRM MID Server parameter
C. Add the MID Server as a discovery source with DNS
D. Set the JIRA Property in true

Answer: A

Explanation:

Question: 117

As a first step in horizontal discovery, where is the Shazzam probe input placed in a request?

A. CC queue
B. LOP services
C. CO Per probes
D. Target IP address

Answer: D

Explanation:

Question: 118

In a discovery pattern, which types are available with C Attributes in the Pattern Designer?

A. Related CI types
B. Main pattern CI type
C. Global CI types
D. All CI types

Answer: A, B

Explanation:

Question: 119

Which selections are necessary to create a custom horizontal pattern to discover a computer operating

system?

A. CI type
B. Application Pattern Type
C. Infrastructure Pattern Type
D. Process Strategy Type

Answer: A, C

Explanation:

Question: 120

Which of the following choices explain differences between Service Mapping and Discovery?

A. Discovery addresses inventory-related use cases, while Service Mapping allows for the creation of accurate maps of application service topologies.
B. Discovery finds applications and devices on your network; Service Mapping monitors those devices.
C. Discovery requires agent installation to find hardware devices; Service Mapping requires agents for software.
D. Discovery utilizes IP address ranges for initial discovery; Service Mapping uses entry points.

Answer: A, D

Explanation:

Question: 121

Which metrics comprise the Completeness KPI for CMDB Health?

A. Required
B. Overall
C. Recommended
D. Audit

Answer: A, C

Explanation:

Question: 122

Which choice best describes what happens when, by default, duplicate CIs are detected during identification and reconciliation?

A. The next discovery is stopped for the CI that is duplicated.
B. Each set of duplicate CIs is added to a de-duplication task.
C. A notification is sent to the CI owner.
D. An association identification rule is created automatically.

Answer: B

Explanation:

Question: 123

In IT environments, what is the purpose of horizontal discovery?

A. A method to map application dependencies
B. A way to visualize network traffic
C. A strategy for managing user access
D. A technique for discovering devices

Answer: A

Explanation:

Question: 124

In Discovery, what table associates an IP address and a credential?

A. Credential Affinity
B. Service Association
C. Service Affinity
D. Tags

Answer: A

Explanation:

Question: 125

Which of the below choices are benefits of Tracked Configuration Files?

A. No credentials needed
B. Content version comparison
C. Files tracked as CIs
D. Unwanted files removed from target

Answer: B, C

Explanation: