Number: CIS-Discover
Passing Score: 800
Time Limit: 120
File Version: 42.0

Exam Code: CIS-Discovery

Exam Name: Certified Implementation Specialist - Discovery

Exam A

QUESTION 1
For CMDB Health, relationships can be which of the following choices? (Choose three.)

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

Correct Answer: A, B, C
Section:
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 isCorrectness, 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.
Duplicaterelationships 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.
Stalerelationships 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.
Orphanrelationships 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.

QUESTION 2
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

Correct Answer: A
Section:
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.
Discovery Dashboard
Discovery troubleshooting | Error messages

QUESTION 3
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

Correct Answer: D
Section:
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.
1: ServiceNow Discovery Overview, page 9
2: ServiceNow Discovery Documentation, Classifiers section

QUESTION 4
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.

Correct Answer: B, C
Section:
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.
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 5
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

Correct Answer: E
Section:
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.
1:View and edit CI identifiers
2:Create or edit a CI identification rule
3:Identification rules

QUESTION 6
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

Correct Answer: A, B
Section:
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.

QUESTION 7
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.

Correct Answer: D
Section:
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.
Functionality Definition
Behavior

QUESTION 8
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

Correct Answer: C
Section:
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.
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 9
Refer to the exhibit.

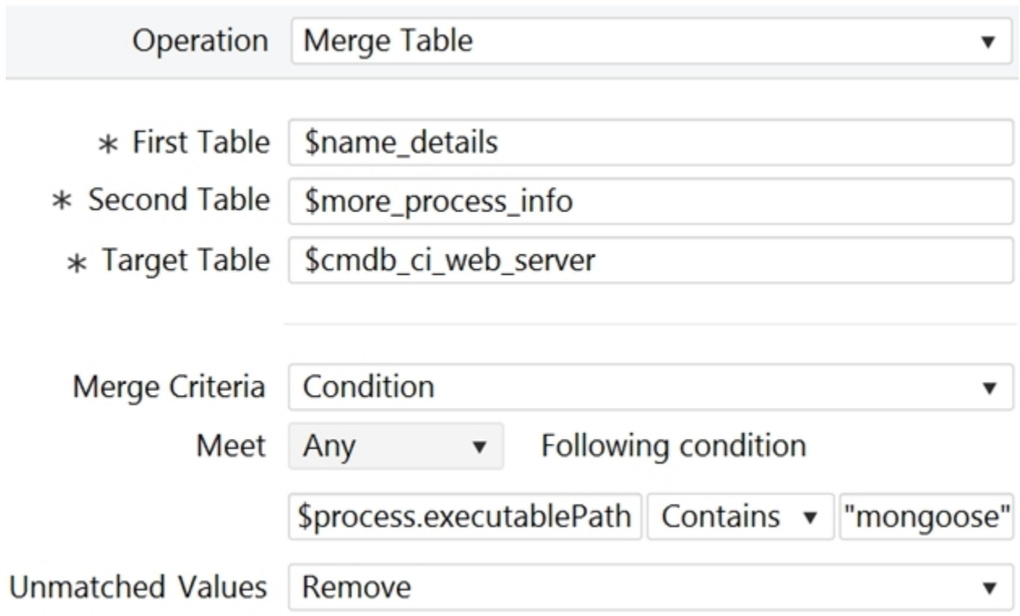

Based on this image, which of the following statements are true? (Choose three.)

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.'

Correct Answer: A, B, D
Section:
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.
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 10
Refer to the exhibit.

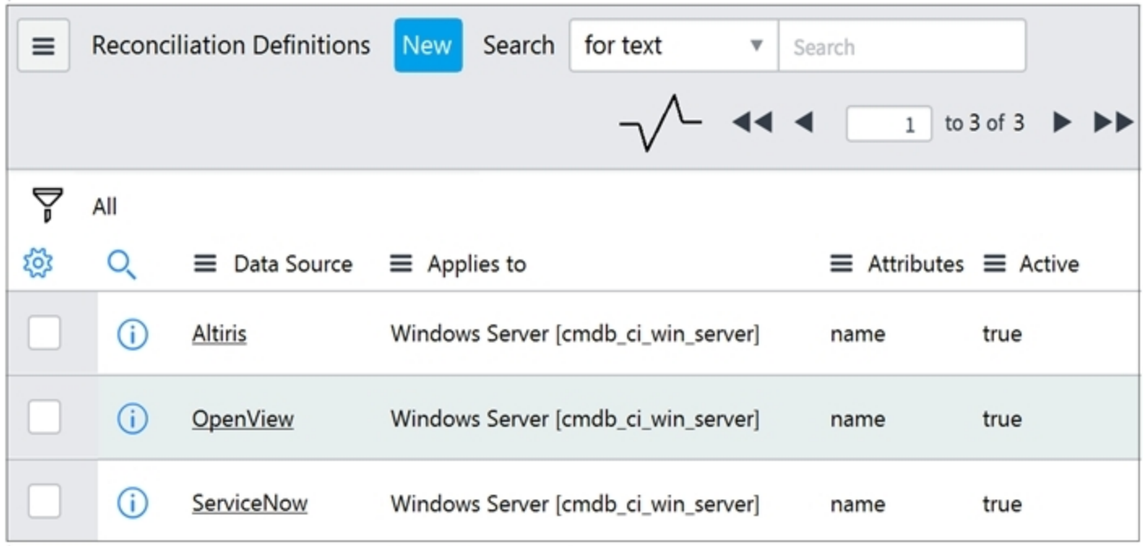

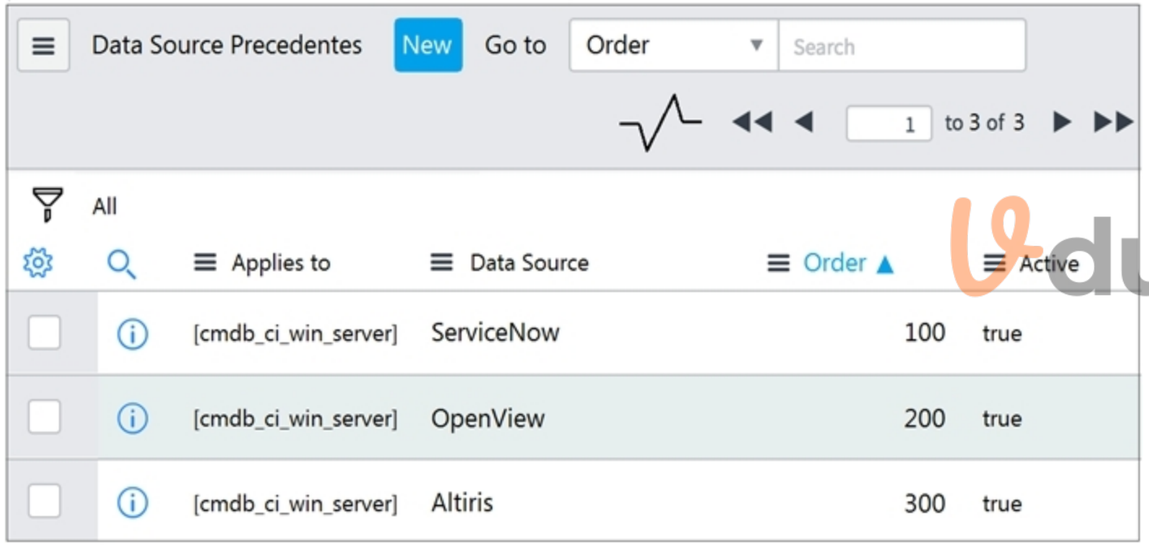

Based on the following images, which choice best describes what occurs if Discovery sets the name attribute of a discovered Windows Server CI to 'Windows1' and then Altiris discovery runs detecting 'Windows2' for the name attribute on the same CI?

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.

Correct Answer: B
Section:
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.
Reconciliation

Data source precedents

QUESTION 11
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.

Correct Answer: C
Section:
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.
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 12
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.

Correct Answer: C
Section:
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.
ServiceNow Discovery Overview
Create a Discovery schedule
Discovery schedule types

QUESTION 13
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

Correct Answer: D
Section:

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.
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 14
Which of the below choices are needed for Quick Discovery? (Choose two.)

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

Correct Answer: A, D
Section:
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.

QUESTION 15
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

Correct Answer: C
Section:
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.

QUESTION 16
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.

Correct Answer: B
Section:
Explanation:
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
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.

QUESTION 17
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

Correct Answer: D
Section:
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.
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 18
SNMP Credentials require which of the following?

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

Correct Answer: C
Section:
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.

QUESTION 19
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

Correct Answer: A
Section:
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.

QUESTION 20
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

Correct Answer: A
Section:
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.
Setting up MID Server user and role
Configure Windows MID Server service credentials
Correcting MID Server Windows service account user and permissions

QUESTION 21
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

Correct Answer: B
Section:
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.
Discovery patterns
Change user operation

QUESTION 22
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

Correct Answer: C
Section:
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.
1:Define discovery steps

QUESTION 23
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

Correct Answer: D
Section:
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.

QUESTION 24
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

Correct Answer: D
Section:
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.

QUESTION 25
Based on this image, which of the following choices is true?

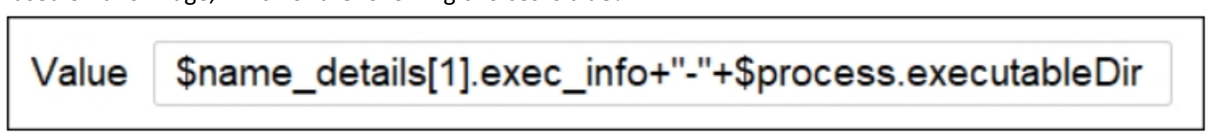

A. This is from a WMI query operation step.
B. There is a scalar variable labeled '1'.
C. This Value cannot be used in a pattern step.
D. There is a tabular variable named 'name_details'.

Correct Answer: D
Section:
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_detailsis 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_infois 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.

QUESTION 26
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

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

QUESTION 27
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

Correct Answer: A, D
Section:
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.
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 28
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

Correct Answer: D
Section:
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.
1: Reconciliation rules - Product Documentation: San Diego - Now Support Portal

QUESTION 29
Which method for deleting specific CIs is not discovered in 30 days?

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

Correct Answer: A
Section:
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.

QUESTION 30
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

Correct Answer: B
Section:
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 credential-based probes fail, or as a primary option for scanning certain types of devices, such as network devices, printers, or IoT devices12.
Credential-less Discovery with Nmap - Product Documentation: San Diego - Now Support Portal
Credential-less host Discovery - ServiceNow - Now Support

QUESTION 31
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

Correct Answer: B
Section:
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 thedscy_credentials_affinitytable, 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.

QUESTION 32
The CMDB contains which of the following record types? (Choose two.)

A. Model
B. Configuration Item (CI)
C. Asset
D. Relation Type

Correct Answer: B, D
Section:
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.
Configuration Item (CI) types: Attributes and relationships
What is a configuration management database (CMDB)?
CSA Challenge Questions Flashcards

QUESTION 33
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

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

QUESTION 34
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

Correct Answer: A, B
Section:
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.

QUESTION 35
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

Correct Answer: A, D
Section:
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.
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

QUESTION 36
Which choices are necessary to launch any pattern? (Choose two.)

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

Correct Answer: A, D
Section:
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 ascmdb_ci_win_serverand the CI type asWindows Server. These choices are mandatory for any pattern, as they enable Discovery to identify and classify the CIs correctly.

QUESTION 37
Which of the below choices are horizontal pattern types? (Choose two.)

A. Hardware
B. Software
C. Infrastructure
D. Application

Correct Answer: C, D
Section:
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.
Horizontal discovery patterns
Discovery pattern types

QUESTION 38
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

Correct Answer: A
Section:
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.
3: Product Documentation | ServiceNow

QUESTION 39
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

Correct Answer: C
Section:

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.
Discovery Schedule form
Shazzam Probe
Discovery IP range

QUESTION 40
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

Correct Answer: C
Section:
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.
1: ServiceNow Discovery Documentation, Discovery Process section
2: ServiceNow Discovery Overview, page 8

QUESTION 41
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)

Correct Answer: A, B, C
Section:
Explanation:
Discovery finds and maps dependencies for the following types of storage devices:Direct-attached storage (DAS),network-attached storage (NAS), orstorage 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.
1:Storage device discovery
2:Storage Discovery via SMI-S and CIM
3:Discovery of storage area networks (SAN)

QUESTION 42
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

Correct Answer: C
Section:
Explanation:
TheShazzam Batch Sizeproperty 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 theDiscovery Propertiesmodule underDiscovery Definition2.

QUESTION 43
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.

Correct Answer: B
Section:
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 drop-down 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.

QUESTION 44
Which of the following pattern operations query targets? (Choose two.)

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

Correct Answer: A, C
Section:
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.
Pattern operations - Product Documentation: Vancouver - ServiceNow
Pattern operations - Product Documentation: San Diego - Now Support Portal

QUESTION 45
File-based Discovery is triggered during the .

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

Correct Answer: C
Section:
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.

QUESTION 46
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

Correct Answer: B
Section:
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.
Set Parameter Value operation
Pattern variables

QUESTION 47
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.

Correct Answer: C
Section:
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.
1: Process classifiers - Product Documentation: San Diego - ServiceNow
2: Horizontal Pattern probe - Product Documentation: San Diego - ServiceNow

QUESTION 48
Which metrics comprise the Completeness KPI for CMDB Health? (Choose two.)

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

Correct Answer: A, B
Section:

QUESTION 49
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.

Correct Answer: C, D
Section:
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.
1: ServiceNow Discovery Overview, page 9
2: ServiceNow Discovery Documentation, Service Mapping section

QUESTION 50
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

Correct Answer: C
Section:
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 areSelect FileandDefine 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.
1:Parse a file
2:Define discovery steps

QUESTION 51
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.

Correct Answer: B
Section:
Explanation:
London Linux Servers is aCMDB 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.

QUESTION 52
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

Correct Answer: C
Section:
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 53
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

Correct Answer: B, C

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

QUESTION 54
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.

Correct Answer: C
Section:
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.

QUESTION 55
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

Correct Answer: A
Section:
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.
Horizontal discovery patterns
Discovery pattern types

QUESTION 56
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

Correct Answer: C, D
Section:
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 asapp 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.
1: Patterns and horizontal discovery - Product Documentation: Tokyo - ServiceNow
2: Parse File operation - Product Documentation: San Diego - ServiceNow

QUESTION 57
The following shows part of the Windows OS - Servers pattern in Pattern Designer:

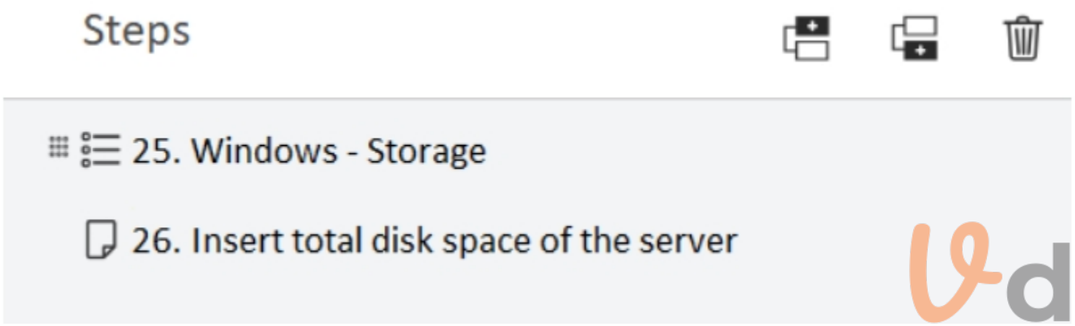

Which of the steps above use(s) a shared library?

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

Correct Answer: C
Section:
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.
Pattern Designer | Shared libraries
Windows - Storage shared library

QUESTION 58
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

Correct Answer: A
Section:
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.
1: ServiceNow Discovery Documentation, MID Server Upgrade section
2: ServiceNow Support Article, Connections to allow list for the MID Server

QUESTION 59
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

Correct Answer: D
Section:
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.
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 60
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

Correct Answer: B, E, F
Section:
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.