DUMPSBOSS.

Zscaler Digital Transformation Administrator

Zscaler ZDTA

Version Demo

Total Demo Questions: 29

Total Premium Questions: 296

Buy Premium PDF

https://dumpsboss.co

support@dumpsboss.co

support@dumpsboss.co
dumpsboss.co

Topic Break Down

Topic	No. of Questions
Topic 1, Zscaler Zero Trust Exchange	47
Topic 2, Zscaler Internet Access	169
Topic 3, Zscaler Private Access	48
Topic 4, Zscaler Digital Experience	25
Topic 5, Mix Questions	7
Total	296

QUESTION NO: 1

During the authentication process while accessing a private web application, how is the SAML assertion delivered to the service provider?

A. HTTP Redirect on the browser

B. API request/response sequence

C. Through the client connector

D. Form POST via the browser

ANSWER: D

Explanation:

Form POST via the browser is correct. In the standard SAML Web Browser SSO flow used when a user accesses a private web application, the identity provider authenticates the user and creates a signed SAML assertion containing authentication and, where applicable, authorization attributes. The identity provider returns an HTML document to the user's browser that includes the encoded SAML assertion in a form field and specifies the service provider's Assertion Consumer Service endpoint as the form action. The browser then submits that form using an HTTP POST to the service provider. The service provider validates the assertion's signature, issuer, audience, recipient, timing conditions, and other relevant statement data before establishing the user session. This browser-mediated POST keeps the assertion delivery aligned with the SAML HTTP-POST binding and allows the service provider to receive the complete assertion at its configured ACS endpoint. Zscaler Private Access integrates with identity providers using SAML and relies on this established browser SSO pattern for web-application authentication. See the OASIS SAML V2.0 Bindings specification and Zscaler documentation for configuring SAML.

QUESTION NO: 2

What is a Landmine in Deception?

A. Agentless plug-in installed on endpoints, such as desktops or laptops on a network. These plug-ins deploy decoy credentials, files, processes, and lures to other decoys at endpoints.

B. Software agent installed on a centralized server in datacenter or in cloud. The agents running in the server deploy decoy credentials, files, processes, and lures to other decoys at endpoints.

C. Software agent installed on endpoints, such as desktops or laptops on a network. These agents deploy decoy credentials, files, processes, and lures to other decoys at endpoints.

D. Agentless plug-in installed on endpoints, such as desktops or laptops on a network. These plug-ins auto rotates decoy credentials, files, processes, and lures to other decoys at endpoints.

ANSWER: C

Explanation:

Software agent installed on endpoints, such as desktops or laptops on a network. These agents deploy decoy credentials, files, processes, and lures to other decoys at endpoints.

In Zscaler Deception, a Landmine is a lightweight endpoint agent that extends deception coverage onto managed workstations and servers. It places deceptive artifacts locally, including credentials, files, browser artifacts, mapped shares, processes, and links to deception resources. These artifacts are intended to appear valuable to an intruder performing reconnaissance, credential harvesting, or lateral movement. Because legitimate users and routine applications should not normally interact with the planted artifacts, access or use of a lure is a high-confidence signal of potentially malicious activity. The Landmine therefore enables early detection close to the attacker's point of activity, rather than relying solely on centralized infrastructure or network-based detection. It also helps make endpoints appear connected to realistic deceptive assets, increasing the likelihood that an attacker encounters a monitored lure during post-compromise activity. Zscaler describes its Deception capability as using decoys and lures to detect adversaries and provide actionable alerts, while

Landmines specifically supply this endpoint-resident deception layer. See the Zscaler Deception product page and the Zscaler Deception documentation for additional product information.

QUESTION NO: 3

A contractor team in a regional lab must upload ZIP archives to an approved code repository but must not upload archives or executables to generic file-sharing sites. A sudden increase in renamed executables, such as an .exe file disguised with a .jpg extension, complicates monitoring.

Which action best applies the correct file-type policy to this team while aligning with security requirements?

A. Define one enterprise-wide file-type block for executables and archives, reference the repository as an exception host, and base decisions on MIME-type matches in the baseline policy

B. Configure an out-of-band CASB scan to flag archives in the code repository, and create a generic SaaS block that checks file extensions for executables

C. Create two File Type Control rules: an allow rule for archive types scoped to the contractor group and approved application, and a block rule for archives and executables scoped to the contractor group and generic file-sharing applications; place the allow rule above the broader block rule

D. Add a URL Filtering rule scoped to the contractor group that allows the repository domain and blocks generic file-sharing domains, relying on file-extension inspection to detect renamed binaries

ANSWER: C

Explanation:

Create two File Type Control rules: an allow rule for archive types scoped to the contractor group and approved application, and a block rule for archives and executables scoped to the contractor group and generic file-sharing applications; place the allow rule above the broader block rule. This design expresses the required business exception precisely: the contractor group can upload approved archive formats to the sanctioned code repository, while the same users remain restricted when they use generic file-sharing applications. Scoping the rules by both user group and application minimizes the policy's reach and avoids granting a broad upload permission outside the approved workflow. File Type Control rule order is important because Zscaler evaluates policies according to their configured precedence; placing the targeted authorization above the more general restriction ensures that the approved repository workflow receives the intended treatment. File Type Control also evaluates the detected file type rather than simply relying on a filename suffix, helping identify executable content that has been renamed with an image extension. This provides enforcement at upload time and addresses the stated risk of disguised binaries while preserving the authorized archive-upload use case. See Zscaler's File Type Control overview and its File Type Control policy configuration guidance.

QUESTION NO: 4

How is data gathered with ZDX Advanced client performance?

A. By generating synthetic transactions to designated Internet and Private applications every 5 minutes and measuring the performance of those sessions.

B. By constantly analyzing live user sessions to both Internet and Private applications and measuring the performance of those sessions.

C. By using AI predictive analysis ZDX can extrapolate near-term client performance based upon recent past data observed.

D. By constantly analyzing live user sessions to critical SaaS applications and measuring the performance of those sessions.

ANSWER: B

Explanation:

By constantly analyzing live user sessions to both Internet and Private applications and measuring the performance of those sessions is correct. ZDX Advanced Client Performance is endpoint-based, real-user monitoring delivered through Zscaler Client Connector telemetry. It continuously collects performance information from actual user activity, allowing administrators to understand the experience users have while accessing internet applications and private applications. This includes useful endpoint and session context, such as device resource conditions, network behavior, and application transaction performance, so investigations can be tied to the user's real experience rather than an assumed or isolated test condition.

This continuous collection is especially valuable for operational troubleshooting because it can reveal whether a degraded experience is associated with the endpoint, the local network, the path to the application, or the application transaction itself. ZDX combines this client-side visibility with broader digital-experience observability, enabling IT teams to identify and remediate issues affecting distributed users. Zscaler describes ZDX as a service for monitoring user experience across applications, networks, and devices, while its Client Performance capability provides the endpoint-level insights needed for real-user performance analysis. See the Zscaler documentation on Client Performance Monitoring and the Zscaler Digital Experience product overview.

QUESTION NO: 5

What role does an App Connector serve?

A. App Connectors enforce security policies for traffic destined for SaaS applications.

B. App Connectors enable user experience monitoring for all applications.

C. App Connectors expose a public IP for users to connect to for private application access.

D. App Connectors mediate seamless communication for applications, services and data sources.

ANSWER: D

Explanation:

App Connectors mediate seamless communication for applications, services and data sources. In Zscaler Private Access, an App Connector is deployed close to private applications in a customer's data center or cloud environment. It creates outbound-only TLS connections to the Zscaler Zero Trust Exchange, rather than accepting inbound connections from the internet. This architecture allows authorized users to reach a specific private application without placing the application on the public internet or extending broad network-level access.

When a user requests an authorized private application, ZPA selects an appropriate App Connector and establishes an application-specific, inside-out connection between the user and that application. The App Connector therefore provides the connectivity path that securely brokers communications to private apps and their associated services or data sources. Its function is central to ZPA's application segmentation model: users are connected to applications they are entitled to use, not directly to the underlying network. Zscaler describes App Connectors as lightweight virtual machines that provide secure connectivity to private applications and form encrypted tunnels to the Zscaler cloud. See the Zscaler documentation on App Connectors and the Zscaler Private Access overview.

QUESTION NO: 6

What is one of the four steps of a cyber attack?

A. Find Cash Safe

B. Find Email Addresses

C. Find Least Secure Office Building

D. Find Attack Surface

ANSWER: D

Explanation:

Find Attack Surface is correct because identifying the attack surface is a foundational step in a typical cyberattack lifecycle. Before an attacker can gain access, they commonly conduct reconnaissance to identify Internet-facing assets and possible entry points. These assets can include externally accessible applications, remote-access services, VPN gateways, cloud services, exposed ports, domains, and other systems that may contain exploitable weaknesses or misconfigurations. This discovery process helps an attacker select a target and determine a path for initial compromise.

Zscaler's zero trust approach is designed around reducing this exposure and preventing attackers from using discovered infrastructure as an entry point or moving through the environment after access is obtained. In practical security operations, maintaining an accurate inventory of externally exposed assets, minimizing unnecessary public exposure, and continuously assessing configuration risk are important measures because the attack surface changes as applications, users, and cloud resources change. The Cybersecurity and Infrastructure Security Agency similarly describes attack-surface management as discovering and assessing Internet-exposed assets so organizations can reduce their exposure. See CISA's attack surface management guidance and Zscaler's definition of attack surface.

QUESTION NO: 7

Live logs show a global DLP rule that blocks uploads of regulated financial data and a departmental override that allows uploads for Finance when device posture is compliant. A Finance user on a compliant device successfully uploads a spreadsheet containing regulated data to a generic file-sharing application, despite expectations that the upload would be blocked. The departmental allow rule appears before the global block rule.

Which conclusion and next step best address the issue?

A. Escalate to the data-protection team to adjust rule precedence so that the global block evaluates before the departmental allow and prevents the upload

B. Instruct the network team to increase the default URL-risk threshold, anticipating fewer permitted uploads through stricter categorization

C. Reduce OCR sensitivity for spreadsheet inspection to limit misclassifications and reduce false negatives in content analysis

D. Revise Advanced Threat Protection sensitivity to reduce permissive outcomes on newly observed destinations and defer DLP rule changes

ANSWER: A

Explanation:

Escalate to the data-protection team to adjust rule precedence so that the global block evaluates before the departmental allow and prevents the upload is the appropriate conclusion and remediation. The live-log evidence directly identifies a rule-order issue: the user, department, device-posture condition, application, and uploaded content match the earlier Finance exception. In the default ordered-policy evaluation model, Zscaler applies the first matching rule, so the later global rule that blocks regulated financial data is not evaluated for that transaction. The upload is therefore an expected outcome of the currently ordered policy, rather than an indication that content inspection failed.

The data-protection owner should redesign the exception to ensure it does not authorize uploads containing the regulated-data classifier to generic file-sharing destinations. This can be accomplished by placing the global regulated-data block before the departmental exception, or by narrowing the exception's destination and data conditions so it only covers explicitly approved Finance business workflows. After publishing the change, the team should validate with a controlled upload and review transaction logs to confirm that the blocking rule is matched. Zscaler's DLP documentation and policy-configuration guidance are available at Zscaler Data Loss Prevention and About DLP Policy Rules.

QUESTION NO: 8

An organization experiences frequent changes in team structure and wants to keep group membership and access aligned consistently.

Which approach supports scalable, controlled administration?

A. Rely on SAML assertions to grant administrator rights during authentication events

B. Consume SCIM-provisioned groups from the identity provider and drive entitlements through those groups

C. Assign administrative capabilities individually to each user to avoid group-level drift

D. Create local user accounts to separate access from external directories

ANSWER: B

Explanation:

Consume SCIM-provisioned groups from the identity provider and drive entitlements through those groups is the scalable, controlled approach. SCIM provides automated lifecycle provisioning between the organization's authoritative identity provider and Zscaler, allowing user and group membership changes to be synchronized rather than manually recreated in the Zscaler administration portal. When a person joins a team, moves to another team, or leaves the organization, the identity provider remains the source of truth and its group updates can flow consistently to Zscaler.

Using synchronized groups as the basis for entitlement assignment also makes access administration policy-driven. Administrators define the appropriate administrative or service entitlements for a group once, and members receive the resulting access according to their current authoritative group membership. This provides a repeatable joiner-mover-leaver process, reduces operational effort, and supports auditing because membership and access decisions are tied to centrally managed directory groups. Zscaler documents SCIM as its mechanism for provisioning users and groups from an identity provider, and its role-management guidance describes assigning roles through groups. See the Zscaler SCIM provisioning documentation and Zscaler role and entitlement assignment documentation.

QUESTION NO: 9

What is an App Profile PAC file used for?

A. It is used to encapsulate traffic within a GRE tunnel.

B. It is used to establish a TLS session with the Zscaler cloud.

C. It is used by Zscaler Client Connector to make traffic-forwarding decisions.

D. It is used to categorize sensitive data.

ANSWER: C

Explanation:

It is used by Zscaler Client Connector to make traffic-forwarding decisions. A Proxy Auto-Configuration (PAC) file contains JavaScript logic that evaluates requested destinations and returns the appropriate proxy-routing instruction. In an App Profile, Zscaler Client Connector can use the PAC file to determine whether matching application traffic should be forwarded to the Zscaler cloud service or handled according to another routing rule. This enables administrators to apply granular forwarding behavior based on URL, hostname, IP address, protocol-related logic, or other conditions expressed in the PAC script. The App Profile associates the configured forwarding behavior with the applicable users, devices, and applications, while the PAC file provides the decision logic used at connection time. As a result, traffic handling can be tailored to organizational requirements without relying solely on broad, device-wide proxy settings. Zscaler documents PAC files as a method for defining web traffic forwarding rules and explains that Client Connector supports forwarding profiles for managing how traffic is sent to Zscaler services. See Zscaler Client Connector: About Forwarding Profiles and Zscaler Internet Access: About PAC Files.

QUESTION NO: 10

Malicious File Protection exclusions can be configured for which type of file?

A. Files sent using the PPTP protocol

B. Files sent using the SCP protocol

C. Files sent using the RTSP protocol

D. Password-encrypted files

ANSWER: D

Explanation:

Password-encrypted files are the supported file category for a Malicious File Protection exclusion. A password-encrypted archive or document cannot be unpacked and inspected unless the inspection service has the password and can successfully decrypt the content. Consequently, its embedded files and payload cannot be subjected to the same malware analysis workflow as accessible, unencrypted content. Zscaler Internet Access provides policy controls for this inspection limitation so that an administrator can define the intended handling of password-protected content in accordance with the organization's security policy and risk tolerance.

This configuration is important because password protection is common in legitimate business workflows, such as securely sharing sensitive documents, but it can also conceal malicious content. The exclusion is therefore a deliberate policy decision: it identifies a file condition for which normal content inspection cannot be completed, rather than identifying a transport mechanism. Administrators should govern such exclusions narrowly, review them regularly, and pair them with appropriate user awareness and file-sharing controls. Zscaler's malware-protection documentation and its cloud security help resources provide the relevant policy context: Zscaler: Configuring a Malware Protection Policy and Zscaler Internet Access Help.

QUESTION NO: 11

A Zscaler Client Connector App Profile is configured to apply a Forwarding Profile that forwards all traffic to the Zero Trust Exchange using Z-Tunnel 2.0. If a change is made to the Logout password in the App Profile, how long will it be before the new logout password is in effect?

A. Policy updates happen in real time, so the new logout password is in effect as soon as the change is saved.

B. The new logout password will be in effect after the Activate button is clicked in the Admin portal.

C. The new logout password will be in effect after the user clicks Update Policy on the client.

D. Policy updates occur every 60 minutes, so the logout password will be in effect after the next scheduled update.

ANSWER: D

Explanation:

Policy updates occur every 60 minutes, so the logout password will be in effect after the next scheduled update. The Logout password is an App Profile setting delivered to endpoints through Zscaler Client Connector policy synchronization. Saving the change in the administrator portal stores the updated configuration, but an installed client applies it when it receives its next policy update. Zscaler Client Connector normally checks for policy changes at a 60-minute interval; consequently, the updated password can take up to one hour to become effective on a given endpoint. This behavior is independent of the forwarding method selected in the profile. Z-Tunnel 2.0 controls how eligible endpoint traffic is forwarded to the Zero Trust Exchange, whereas the logout-password behavior is governed by the Client Connector policy-update cycle. Administrators should therefore allow for the next scheduled client policy refresh when planning a password change or communicating the new value to support teams and users. Zscaler documents Client Connector administration and profile-based configuration in its Zscaler Client Connector documentation, and describes the service architecture in its Zero Trust Exchange overview.

QUESTION NO: 12

A regional hospital must provide a vendor with intermittent access to a legacy device-management application hosted on two on-premises servers. The vendorâ€™s previous VPN caused noisy port scans to appear in logs and exposed nearby subnets to probing.

Which action should the administrator take to constrain access to the application while reducing lateral movement?

A. Create ZPA Application Segments for the device-management FQDNs and ports, and enforce identity- and posture-based policies for the vendor group

B. Configure DNS sinkholes to divert off-port vendor traffic and apply URL Filtering to suppress non-application flows

C. Deploy a dedicated VLAN and a jump host near the application, and restrict traffic with subnet ACLs that limit the vendor to the jump hostâ€™s IP address

D. Implement host-based firewall rules on both servers and advertise a reduced VPN route set to the vendor client

ANSWER: A

Explanation:

Create ZPA Application Segments for the device-management FQDNs and ports, and enforce identity- and posture-based policies for the vendor group is correct because Zscaler Private Access provides application-level, zero-trust access rather than extending network access through a VPN. The administrator can define private application segments using the two application servers' FQDNs or IP addresses and specify only the ports required by the legacy device-management service. This ensures that the vendor can reach the intended application resources without being placed onto the hospital's internal network or receiving routable visibility of adjacent subnets.

Access can then be governed through ZPA policy using the vendor's identity group, and the administrator can apply posture requirements to ensure that only compliant vendor devices connect. For intermittent third-party use, policy conditions and authentication controls can further limit access to the appropriate users and sessions. ZPA's application segmentation model minimizes the attack surface available to the vendor and directly addresses the prior VPN-related scanning and lateral-movement risk. Zscaler describes this model as connecting authorized users to specific private applications without exposing the network. See Zscaler's Application Segments documentation and Zscaler Private Access overview.

QUESTION NO: 13

A new customer has just purchased Zscaler for Users.

Which of the following Zscaler service entitlements is enabled by default?

A. ZPA

B. Deception

C. ZIA

D. ZDX

ANSWER: C

Explanation:

ZIA is enabled by default as a service entitlement in Zscaler Client Connector. Zscaler Internet Access enrollment is the baseline entitlement and is available without first creating a selective entitlement rule. This allows a newly provisioned Zscaler for Users environment to enroll users into ZIA so their internet and SaaS traffic can receive the organization's configured security controls, such as web filtering, cloud firewall policies, and data-protection policies.

In Client Connector, service entitlements control which Zscaler services are made available to users and devices. Administrators can use entitlement rules to selectively assign services based on identity or device criteria when their deployment and licensing requirements call for more granular targeting. However, Zscaler documents that ZIA enrollment is always enabled by default, making it the appropriate default service choice for a new deployment. This default supports an initial internet-security rollout while administrators configure additional service access and deployment policies for their environment.

See Zscaler's About Service Entitlements documentation and the Zscaler Client Connector documentation for the entitlement model and enrollment behavior.

QUESTION NO: 14

Which field within a URL filtering rule must be defined for Browser Isolation to work?

A. Groups

B. User Agent

C. Departments

D. Device Trust

ANSWER: B

Explanation:

User Agent must be defined in the URL Filtering rule when Browser Isolation is selected as the rule action. The User Agent field identifies the browser/client characteristics associated with web traffic, allowing Zscaler Internet Access to apply Browser Isolation to supported browser sessions. This browser context is essential because isolation is delivered through a browser-based session and the service needs to recognize the relevant user-agent characteristics when evaluating the URL Filtering policy.

In practice, an administrator creates or selects an appropriate user-agent definition and includes it in the URL Filtering rule along with the target URL categories, URLs, or cloud applications and the Browser Isolation action. This makes the isolation policy deterministic for the intended browser traffic and ensures that qualifying requests are redirected to the isolated browsing experience according to the policy.

Zscaler documents Browser Isolation configuration as part of URL Filtering policy setup and identifies the User Agent criterion as a required configuration element for this use case. See the Zscaler Internet Access Browser Isolation documentation and the Zscaler URL Filtering policy documentation for configuration guidance.

QUESTION NO: 15

Can Notifications, based on Alert Rules, be sent with methods other than email?

A. Email is the only method for notifications as that is universally applicable and no other way of sending them makes sense.

B. In addition to email, text messages can be sent directly to one cell phone to alert the CISO who is then coordinating the work on the incident.

C. Leading ITSM systems can be connected to the Zero Trust Exchange using a NSS server, which will then connect to ITSM tools and forwards the alert.

D. In addition to email, notifications, based on Alert Rules, can be shared with leading ITSM or UCAAS tools over Webhooks.

ANSWER: D

Explanation:

In addition to email, notifications, based on Alert Rules, can be shared with leading ITSM or UCAAS tools over Webhooks. Zscaler alerting supports webhook-based notification delivery so that an alert can initiate workflows in external operational tools rather than relying exclusively on an email recipient. A webhook sends the alert information to a configured HTTPS endpoint, where an IT service management platform, collaboration service, incident-response automation platform, or unified communications service can process it according to the organization's workflow. This enables operational teams to create or update tickets, post alerts to collaboration channels, trigger orchestration, or notify the appropriate on-call process quickly. The selected destination and its handling of the notification are configured by the organization, but the key capability is that Alert Rules can use webhooks alongside email delivery. Webhook integrations therefore provide a flexible method to connect Zscaler-generated alert notifications with established incident-management and communications processes. Zscaler documents webhook support for alerting and integrations in its administrator documentation and product resources. See Zscaler Internet Access Help and Zscaler Private Access Help.

QUESTION NO: 16

What is the default timer in ZDX Advanced for web probes to be sent?

A. 1 minute

B. 10 minutes

C. 30 minutes

D. 5 minutes

ANSWER: D

Explanation:

The correct answer is **5 minutes**. In Zscaler Digital Experience (ZDX) Advanced, web probes are synthetic HTTP/HTTPS checks that ZDX runs from enrolled endpoint devices to assess the end-user experience of reaching a configured web application or URL. The default recurrence is five minutes, which provides a regular stream of measurements for availability, DNS resolution, connection establishment, SSL negotiation, response time, and related web-transaction performance data. This cadence gives administrators sufficiently current visibility into experience trends and intermittent connectivity problems while avoiding the unnecessary endpoint and network overhead of an extremely aggressive probe schedule. The timer is a default behavior associated with the web-probe configuration; administrators can use the ZDX administration workflow and applicable probe settings to manage monitoring based on operational requirements. Understanding that five minutes is the baseline is important when evaluating probe timestamps, interpreting ZDX dashboards, and deciding whether a different monitoring frequency is needed for a particular service. Zscaler describes ZDX monitoring capabilities and administration resources in the Zscaler Digital Experience documentation and provides a product overview at Zscaler Digital Experience.

QUESTION NO: 17

Which of the following enables the discovery of newly observed domains within three minutes of the domain coming online?

A. IP Chicken

B. MXToolbox

C. Farsight Feed

D. Dig

ANSWER: C

Explanation:

Farsight Feed is correct because it provides passive DNS threat-intelligence data that identifies domains as they are first observed in real DNS activity. This near-real-time visibility is particularly valuable for a security service because malicious infrastructure is often created and activated shortly before it is used for phishing, malware delivery, command-and-control communications, or other attacks. Rather than waiting for a domain to accumulate a reputation history or appear on a conventional block list, a feed based on broad passive-DNS telemetry can surface a newly active domain within minutes of its first appearance. In the Zscaler threat-intelligence context, the three-minute discovery capability refers to this rapid identification of newly observed domains and enables earlier policy enforcement and inspection decisions. The resulting intelligence helps Zscaler protect users from emerging threats whose domains may have little or no established reputation data. Passive DNS is also useful because it records observed DNS resolutions at scale, offering timely evidence that a domain has become active in the Internet ecosystem. For background on Zscaler's DNS-security approach, see Zscaler: What Is DNS Security?. Additional information about passive DNS intelligence is available from DomainTools DNSDB.

QUESTION NO: 18

A security team suspects that data exfiltration is occurring through encrypted channels to attackers.

To assess the company's posture before tuning controls, which next step should be taken to validate whether existing protections cover this behavior?

A. Raise the severity of egress firewall rules across segments to constrain outbound flows that might be exploited

B. Review ZIA DLP outbound logs for anomalous uploads to unsanctioned SaaS applications and newly registered domains to gauge detection coverage

C. Correlate ZIA threat and DLP evidence for suspicious outbound traffic, confirm that the sessions were TLS-decrypted, and verify that DLP and botnet controls applied to those sessions

D. Trigger broad Cloud Sandbox reanalysis of recent endpoint downloads to look for latent payloads that could facilitate exfiltration

ANSWER: C

Explanation:

Correlating ZIA threat and DLP evidence with confirmation of TLS inspection is the appropriate validation step because encrypted exfiltration can only be assessed reliably when the organization knows whether the relevant outbound sessions were decrypted and inspected. ZIA SSL inspection provides visibility into encrypted web traffic so that content-aware controls, including DLP policies, can evaluate uploads and other outbound data transfers. The team should first use transaction, threat, and DLP records to identify suspicious destinations, users, devices, and upload activity, then determine whether those transactions matched SSL-inspection policy or were exempted. Finally, it should verify that the applicable DLP and botnet/threat-protection policies were enabled and scoped to those inspected flows. This establishes the actual control posture before any policy tuning: whether inspection occurred, whether the intended policies applied, and whether the controls generated detections or allowed the activity. Correlating the evidence with relevant user and private-application activity can add investigation context, but the essential coverage check is TLS inspection plus policy applicability for the suspicious outbound traffic. See Zscaler's SSL Inspection documentation and Data Loss Prevention documentation.

QUESTION NO: 19

Which list of protocols is supported by Zscaler for Privileged Remote Access?

A. RDP, VNC and SSH

B. RDP, SSH and DHCP

C. SSH, DNS and DHCP

D. RDP, DNS and VNC

ANSWER: A

Explanation:

RDP, VNC and SSH is correct because Zscaler Privileged Remote Access is designed to provide secure, controlled, clientless access to privileged resources using these remote-management protocols. Administrators and authorized third parties can use browser-based sessions to reach Windows systems through Remote Desktop Protocol, Linux and Unix systems through Secure Shell, and graphical systems through Virtual Network Computing. This approach allows organizations to apply least-privilege access and session controls without exposing the target resources directly to the internet or requiring users to establish a traditional network-level VPN connection.

Privileged Remote Access is part of Zscaler Private Access and is focused on safely brokering administrative connections to private applications and infrastructure. Zscaler documentation identifies RDP, SSH, and VNC as the supported protocols for this capability. The service can also provide operational security features such as session recording and auditing, helping organizations maintain visibility over privileged activity while keeping access limited to the specific authorized resource. See Zscaler's Privileged Remote Access overview and the Zscaler Private Access product page for further details.

QUESTION NO: 20

A branch forwards internet traffic to ZIA through GRE tunnels. The branch has two tunnels configured, each with sufficient capacity under normal conditions. During a tunnel failure, all traffic shifts to the surviving tunnel and the site exceeds its required throughput, causing congestion.

Which design change best addresses the requirement for both capacity and resilience?

A. Reduce the GRE tunnel MTU so the surviving tunnel can carry more application traffic

B. Deploy sufficient GRE tunnel capacity for the remaining tunnels to carry required traffic during a single-tunnel failure

C. Replace both GRE tunnels with one IPsec tunnel to centralize forwarding behavior

D. Disable Bandwidth Control so congestion is no longer enforced by ZIA

ANSWER: B

Explanation:

Deploy enough GRE tunnel capacity so that the remaining tunnels can carry the required traffic during a single-tunnel failure, and use appropriate routing to distribute traffic during normal operation. High availability is not achieved merely by having multiple tunnels; the surviving design must also have sufficient aggregate capacity for the required failure scenario.

Reducing the MTU may help with fragmentation in some cases but does not add forwarding capacity. Replacing GRE with a single IPsec tunnel removes redundancy, and disabling bandwidth controls does not resolve a physical or logical tunnel-capacity shortfall.

QUESTION NO: 21

A ZPA application segment for an internal payroll application is configured with the correct FQDN and port. The associated Access Policy allows the Payroll group, and affected users authenticate successfully. However, users receive an application-unreachable error. App Connector health is green, but the App Connector Group associated with the Server Group cannot route to the payroll server subnet.

What should the administrator do next?

A. Add an inbound firewall rule that permits internet traffic directly to the payroll server

B. Associate the Server Group with an App Connector Group that has network reachability to the payroll servers

C. Move the Payroll group above all other groups in the SAML assertion

D. Configure the payroll FQDN as a public URL Filtering allowlist entry

ANSWER: B

Explanation:

Associate the Server Group with an App Connector Group that has network reachability to the payroll servers. ZPA requires more than a successful user authorization decision: an eligible App Connector must also be able to reach the application servers defined in the Server Group. A healthy connector that lacks a route to the target subnet cannot establish the application-side portion of the microtunnel.

Changing the Access Policy would not resolve a network-reachability failure after authorization has already succeeded. Likewise, adding a public DNS record or exposing an inbound firewall port conflicts with the ZPA inside-out connectivity model. The appropriate correction is to use an App Connector Group deployed in, or routed to, the application environment and associate it with the Server Group.

QUESTION NO: 22

What is the purpose of the Zscaler Client Connector providing the authentication token to the Zscaler Client Connector Portal after it is received from Zscaler Internet Access?

A. To bypass multifactor authentication (MFA) during the enrollment process

B. To immediately grant the user access to Zscaler Private Access resources

C. To enable the portal to register the user's device and pass the registration to Zscaler Internet Access

D. To share the authentication token with the SAML IdP to validate the user session

ANSWER: C

Explanation:

The purpose is **to enable the portal to register the user's device and pass the registration to Zscaler Internet Access**. In the Client Connector enrollment and authentication workflow, Zscaler Internet Access authenticates the user and issues an authentication token to Client Connector. Client Connector then provides that token to the Client Connector Portal, which uses it to perform device registration. This registration establishes the association between the authenticated user and the endpoint, allowing Zscaler services to recognize the device as enrolled and apply the appropriate user and device context. The resulting registration information is conveyed to Zscaler Internet Access so that the service can maintain the device's registered state alongside the user's authenticated session. This workflow supports centrally managed endpoint enrollment and enables policy decisions that rely on recognized Client Connector devices. Authentication remains subject to the organization's configured identity-provider and multifactor-authentication requirements; the token is used for the post-authentication registration workflow. See Zscaler's Zscaler Client Connector overview and the Zscaler Internet Access overview for the roles of Client Connector and ZIA in identity-aware, endpoint-based access.

QUESTION NO: 23

Which Risk360 key focus area observes a broad range of event, security configurations, and traffic flow attributes?

A. External Attack Surface

B. Prevent Compromise

C. Data Loss

D. Lateral Propagation

ANSWER: B

Explanation:

Prevent Compromise is the Risk360 key focus area that evaluates a broad range of security-relevant telemetry, including security events, configuration posture, and traffic-flow attributes, to determine an organization's likelihood of suffering an initial compromise. This focus area is intended to make compromise risk measurable and actionable: it brings together signals that show whether users, devices, applications, and security controls are exposed to conditions that can enable an attacker to gain an initial foothold. Administrators can use the resulting risk insights to prioritize remediation, such as strengthening security policy enforcement, addressing unsafe configurations, and reducing exposure associated with risky activity or traffic patterns. In the Risk360 model, this assessment supports a prevention-oriented approach rather than treating compromise as an isolated alert or a single product configuration issue. The broad telemetry scope is important because initial compromise risk is influenced by the combined state of security controls and observed activity across the environment. Zscaler describes Risk360 as a platform for quantifying cyber risk and prioritizing actions to reduce it. See Zscaler Risk360 and Zscaler Risk360 documentation.

QUESTION NO: 24

A security team enables Advanced Firewall controls to investigate suspected command-and-control traffic that is using DNS queries to carry encoded data. Standard port-based rules allow DNS because users require normal name resolution.

Which capability is most appropriate for identifying and controlling the suspected behavior?

A. Use DNS Tunnel and DNS Application Control

B. Use Bandwidth Control to cap all UDP traffic at the affected locations

C. Use URL Filtering to block the Search Engines and Portals category

D. Use a Do Not Inspect rule for DNS traffic to avoid disrupting name resolution

ANSWER: A

Explanation:

Use DNS Tunnel and DNS Application Control. DNS tunneling can abuse otherwise permitted DNS traffic to exchange commands or exfiltrate data, so a basic rule that permits UDP or TCP port 53 is insufficient to identify the behavior. Advanced Firewall DNS controls provide DNS-specific inspection and policy enforcement for this use case.

URL Filtering applies to web destinations rather than providing primary control of DNS tunneling. Bandwidth Control can limit throughput but cannot reliably distinguish malicious encoded DNS activity from ordinary name resolution.

QUESTION NO: 25

A microsegmentation policy set contains a broad “allow employees to internal applications” rule before more specific controls. An incident review found SMB access from non-finance hosts to a finance file share.

Which refinement best addresses the unintended access while improving the internal security posture?

A. Add bandwidth QoS constraints to the internal applications segment so non-finance SMB attempts are deprioritized at runtime

B. Insert deception assets in the finance segment to divert suspicious SMB traffic away from the file share and collect telemetry

C. Tighten URL Filtering for internal destinations so SMB-related domains resolve poorly in non-finance contexts

D. Reorder the rules so the deny for non-finance SMB is evaluated before broad employee allows, and scope the SMB policy to finance hosts and device posture

ANSWER: D

Explanation:

Reorder the rules so the deny for non-finance SMB is evaluated before broad employee allows, and scope the SMB policy to finance hosts and device posture is correct because it resolves the policy-order issue that allowed the unintended SMB connection while applying least-privilege access to a sensitive internal resource. Access policy evaluation must ensure that a narrowly targeted control is considered before a broad rule that would otherwise match the same traffic. Placing the non-finance SMB restriction first creates deterministic enforcement for the identified lateral-access path. Restricting permitted SMB access to the specific finance hosts that require the share further reduces the attack surface, rather than granting all employees a general route to internal applications.

Including device posture adds a zero-trust condition to the authorization decision: even an approved finance user or host should meet the organization’s endpoint-security requirements before accessing the file share. This combination of ordered policy evaluation, application/protocol specificity, workload identity, and device context supports segmentation goals by limiting connectivity to explicitly authorized paths. Zscaler describes Zero Trust Segmentation as controlling workload communications and reducing lateral-movement opportunities, while its Zero Trust Exchange model bases access on context instead of broad network trust. See Zscaler Zero Trust Segmentation and Zscaler’s Zero Trust Architecture overview.

QUESTION NO: 26

If you ' re migrating from an on-premises proxy, you will already have a proxy setting configured within the browser or within the system. With Tunnel Mode, the best practice is to configure what type of proxy configuration?

A. Execute a GPO update to retrieve the proxy settings from AD.

B. Enforce no Proxy Configuration.

C. Use Web Proxy Auto Discovery (WPAD) to auto-configure the proxy.

D. Use an automatic configuration script (forwarding PAC file).

ANSWER: B

Explanation:

Enforce no Proxy Configuration. is the correct practice when migrating users from an on-premises proxy to Zscaler Client Connector Tunnel Mode. In Tunnel Mode, Client Connector steers applicable user traffic into a Zscaler tunnel, where Zscaler applies the organization's Internet-access policy, security inspection, and forwarding behavior. Browser- or operating-system-level proxy settings are therefore not needed to direct traffic to Zscaler.

Removing or disabling legacy proxy configuration provides a clean traffic-steering design: Client Connector is the sole endpoint component responsible for sending traffic through the service. This avoids ambiguous routing decisions during migration and prevents residual enterprise proxy settings from influencing how applications attempt to reach web destinations. It also makes the endpoint configuration more consistent across browsers and applications, rather than relying on each application to honor a proxy setting or discovery mechanism.

Zscaler documents Tunnel Mode as a Client Connector traffic-forwarding method and describes its role in forwarding traffic through Zscaler. For implementation context, see the Zscaler Tunnel Mode documentation and the Zscaler Client Connector overview.

QUESTION NO: 27

When configuring a ZDX custom application and choosing Type: 'Network' and completing the configuration by defining the necessary probe(s), which performance metrics will an administrator NOT get for users after enabling the application?

A. Server Response Time

B. ZDX Score

C. Client Gateway IP Address

D. Disk I/O

ANSWER: A

Explanation:

Server Response Time is not produced for a custom application configured as the **Network** type. Network applications use administrator-defined network probes, such as ICMP, TCP, or UDP probes, to assess the reachability and network-path experience to the configured destination. The resulting telemetry is centered on network conditions, including measurements such as latency, packet loss, jitter where applicable, and probe availability. A network probe does not make an HTTP/S request to an application server and therefore cannot measure the time taken by that server to process a request and begin returning an application response.

Server Response Time is an application-layer measurement associated with web transaction monitoring. It requires a web-style probe that can perform an HTTP or HTTPS request and distinguish server processing/response behavior from the other stages of the transaction. Selecting the Network application type intentionally focuses ZDX monitoring on transport and connectivity behavior rather than HTTP application response timing. ZDX still combines relevant monitored experience data into its user-experience views and scores, while endpoint context is collected separately by the ZDX Client. See Zscaler's custom application configuration documentation and the ZDX score documentation.

QUESTION NO: 28

A URL policy set includes an early allow rule based on a location group for a collaboration application, with no HTTP-method restrictions. A later rule targets high-risk users and blocks PUT and DELETE requests to the same application. A high-risk user in the allowed location attempts a PUT request.

What outcome results from this arrangement of controls?

A. The request is throttled because of conflicting attributes, resulting in degraded service with limited data transfer

B. The request is partially restricted at the enforcement point, causing intermittent failures instead of a deterministic block

C. The request is blocked by the method-aware rule because protocol specificity overrides the broader location-based allow

D. The request is allowed by the earlier location-based rule, preventing the later method-aware block from taking effect

ANSWER: D

Explanation:

The request is allowed by the earlier location-based rule, preventing the later method-aware block from taking effect. Zscaler Internet Access URL Filtering rules are evaluated in their configured order, from the top of the policy list downward. When a transaction satisfies a rule's criteria, ZIA applies that rule's action and does not continue evaluating lower-priority rules for that transaction. Here, the earlier Allow rule matches because the user is in the specified location group and is accessing the named collaboration application. Since that rule has no HTTP-method condition, it also matches a PUT request. Its Allow action is consequently final for this request, even though the user is high risk and the later rule explicitly blocks PUT and DELETE methods.

For the intended protection to apply, the method-aware high-risk-user Block rule must be ordered above the broad location-based Allow rule. Alternatively, the allow rule can be narrowed so that it does not match the sensitive methods or high-risk users. This ordering behavior makes policy design predictable: rule specificity does not supersede rule position. Administrators can verify the request method and applied policy in ZIA logging and reporting. See Zscaler's URL Filtering policy documentation and Web Insights documentation.

QUESTION NO: 29

What Malware Protection setting can be selected when setting up a Malware Policy?

A. Isolate

B. Bypass

C. Block

D. Do Not Decrypt

ANSWER: C

Explanation:

Block is a selectable action when configuring a Zscaler Internet Access Malware Protection policy. Malware Protection is designed to inspect web and download traffic for malicious content using Zscaler's threat intelligence, file reputation, and malware-detection capabilities. When a file or other object matches the policy's malware criteria, selecting Block prevents the unsafe content from being delivered to the user, reducing the risk of endpoint infection and subsequent compromise.

This action is fundamental to an inline secure web gateway control: Zscaler evaluates traffic as it passes through the service and enforces the configured policy before harmful content can reach a client device. Administrators use malware policies to define the scope of traffic and threats to which the action applies, then use Block to enforce protective handling for matching content. Zscaler's Malware Protection documentation describes policy-based controls for detecting and blocking malicious files and traffic, while its security service overview explains the platform's inline inspection and threat-prevention approach.

References: Zscaler Help: About Malware Protection and Zscaler Internet Access.