Configure Optimal ISE IP MTU in SD-WAN for SDA Deployments

Contents

Introduction

This document describes how Maximum Transmission Unit (MTU) issues can impact micro-segmentation in SDA when SD-WAN is used to connect SDA sites.

Prerequisites

Requirements

Cisco recommends that you have knowledge of these topics:

- Cisco Software Defined Access (SDA)
- Cisco Softwaree Defined Wide Area Networks (SD-WAN)
- Cisco Identity Services Engine (ISE)

Components Used:

The information in this document is based on SDA, SDWAN, and ISE.

The information in this document was created from the devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If your network is live, ensure that you understand the potential impact of any command.

Background Information

Modern enterprise networks increasingly leverage SDA for granular micro-segmentation and consistent policy enforcement. To connect distributed SDA sites, Cisco SD-WAN is often employed, offering agile, secure, and optimized transport over various underlay networks. Central to this architecture, the ISE provides critical Authentication, Authorization, and Accounting (AAA) services, along with dynamic policy distribution (for example, Security Group Tags (SGTs) and downloadable ACLs).

While robust, the integration of these powerful technologies can introduce subtle yet impactful configuration challenges. The MTU handling at critical network handoff points and across the SD-WAN overlay is a prime area for such issues. This article addresses two common MTU mismatch scenarios that can disrupt network operations:

1. The MTU gap between SDA border nodes and SD-WAN edge devices.
2. MTU constraints for ISE-originated traffic traversing the SD-WAN overlay.

Proper MTU alignment is paramount to prevent packet fragmentation issues or silent drops, ensuring reliable authentication, policy enforcement, and overall network stability. Failure to address these can lead to perplexing intermittent connectivity and policy enforcement failures, consuming significant troubleshooting effort.

Common Symptoms of Misaligned MTU

Misaligned MTU can manifest in various ways, often leading to difficult-to-diagnose issues:

- **Intermittent RADIUS authentication failures or timeouts:** Especially noticeable for policies generating larger RADIUS packets (for example, those with extensive AV-pairs or certificates).

- **Endpoints failing to receive or apply downloadable ACLs (dACLs) or TrustSec policies (SGTs/SGACLs):** These policies are often conveyed in large RADIUS packets.

- **Slow session establishment for authenticated clients:** Due to retransmissions at the application layer.

- **Excessive RADIUS retransmissions:** Observable in ISE logs or on the Network Access Devices (NADs).

- **Inconsistent policy propagation:** Policy changes made in ISE may not consistently propagate to all NADs in remote SDA sites.

- **Packet capture discrepancies:** Captures can show ISE sending large packets (for example, >1450 bytes) with the Do Not Fragment (DF) bit set, but no corresponding response or ICMP "Fragmentation Needed" error from the NAD or SD-WAN Cisco Edge Router.

- **Incrementing packet drop counters:** Observed on the ingress interface of the Data Center (DC) Cisco Edge Router for traffic sourced from ISE destined for SDA sites, or on the SD-WAN Cisco Edge Router interface facing the SDA border for traffic in the reverse direction.

Problem Description

A Typical Enterprise Deployment
Consider a common enterprise topology:

- **Cisco ISE Servers:** Deployed in a centralized Data Center (DC) or regional hubs, connected to the DC network infrastructure.

- **DC Infrastructure:** Comprises DC core or aggregation switches to which ISE servers connect.

- **SD-WAN Overlay:** DC Cisco Edge Router routers establish SD-WAN tunnels (commonly IPsec) over an underlay transport network (for example, Internet, MPLS) to Cisco Edge Router routers at remote SDA sites.
- **SDA Site:** Remote site Cisco Edge Router routers connect to the local SDA fabric, which includes fabric edge nodes, border nodes, wireless LAN controllers (WLCs), and ultimately, the endpoints.

Illustrative Topology

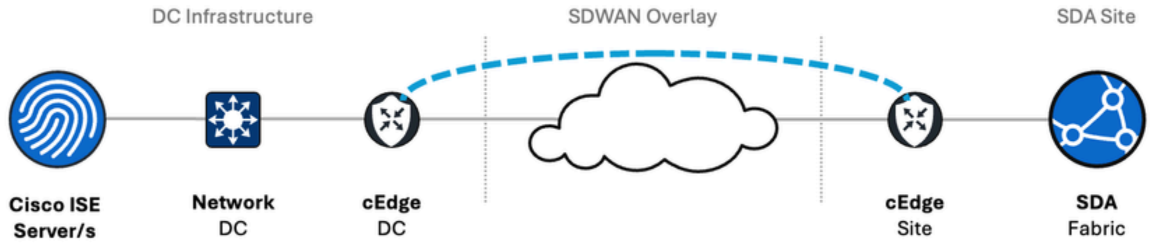

Challenge 1: The MTU Gap – SDA Borders to SD-WAN Edges

Cisco SDA design principles, often implemented via LAN Automation, promote a campus-wide MTU of 9100 bytes (jumbo frames) on all fabric devices. This includes Catalyst 9000 series border nodes and ensures that Ethernet jumbo frames are transported efficiently within the fabric. Consequently, the Layer 3 or SVI handoff interface on an SDA border node defaults to this larger MTU.

Conversely, SD-WAN edge devices, such as the Catalyst 8000 series, typically default to an interface MTU of 1500 bytes. This is standard for interfaces connecting to external networks like Internet Service Providers (ISPs), where jumbo frame support is uncommon or not enabled.

This disparity creates an immediate point of potential failure: an SDA border attempting to send an IP packet larger than 1500 bytes to an SD-WAN edge whose receiving interface is configured for a 1500-byte MTU.

This type of MTU mismatch is a common pitfall in SDA deployments and is often easy to overlook during configuration. What makes it more challenging is that certain behaviors related to how RADIUS requests are generated on Catalyst 9000 switches running Cisco IOS-XE® can cause these issues to surface only under specific and critical conditions.

For instance, RADIUS requests generated during the end-user authentication process handled by the Session Manager Daemon (SMD) process, are hardcoded to fragment packets at 1396 bytes. In contrast, RADIUS requests involved in retrieving TrustSec policies, such as Security Group Access Control List (SGACL)s, are generated by Cisco Internetworking Operating System daemon (IOSd) subcomponents. These are MTU-aware and can avoid fragmenting packets unless their size exceeds the system MTU (typically up to 9100 bytes).

As a result, issues related to MTU mismatches are only become apparent when Cisco TrustSec (CTS) download policies are in use. Additionally, the set of Role-Based Access Control List (RBACL)s downloaded by an SDA edge device during user authentication may vary depending on which SGACL policies are already present for other tags. In practice, the switch downloads only the non-overlapping portions of policy sets.

Together, these behaviors can produce unpredictable and inconsistent outcomes, ranging from silent failures to incomplete policy downloads, depending on the size of the SGACL policy, current system conditions, and, ultimately, MTU misalignments along the path.

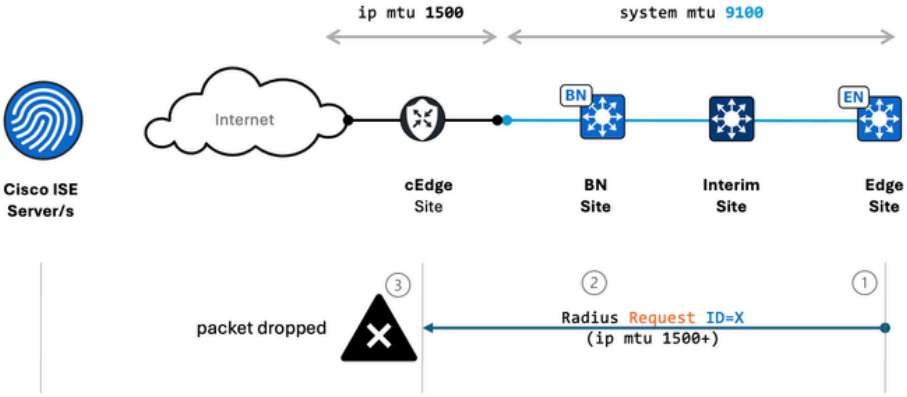

SDA Border forwards a large RADIUS packet (for example, 1600 bytes) towards the ISE via the SD-WAN edge, this is what occurs:

1. The SDA Border, with its 9100 MTU interface, sends the 1600-byte IP packet.
2. The SD-WAN Cisco Edge Router receives this packet on its 1500 MTU interface.
3. However, if the Do Not Fragment (DF) bit is not set on these RADIUS packets, the SD-WAN Cisco Edge Router can often drop them upon ingress simply because they are "oversized" relative to its configured interface MTU. It does not get to the stage of IP forwarding logic where it can consider fragmenting them (if DF bit allows).

This silent drop leads to significant troubleshooting headaches, especially as the issue is directional (SDA to SD-WAN/ISE).

A similar MTU mismatch can occur at the data center (DC) core or leaf switches, which are typically configured to support jumbo frames (for example, MTU 9000+) to enhance internal DC traffic efficiency. However, if traffic is handed off to the LAN-facing interface of an SD-WAN DC Cisco Edge Router router configured with a standard MTU (for example, 1500 bytes), this mismatch can lead to fragmentation or packet drops, particularly for traffic flowing from the DC network into the SD-WAN fabric.

Solution to Challenge 1:

Align the IP MTU on the SDA Border's handoff interface (physical or SVI) with the peering SD-WAN Cisco Edge Router interface, typically 1500 bytes.

Configuration Example (on SDA Border Node):

<#root>

```
!
interface Vlan3000  // Or your physical handoff interface, for example, TenGigabitEthernet1/0/1
 description Link to SD-WAN cEdge Router
 ip address 192.168.100.1 255.255.255.252


ip mtu 1500
```

```
  // Align with SD-WAN cEdge receiving interface MTU
!
```

Important Consideration: Fragmentation on Catalyst 9000 Borders

Catalyst 9000 series switches, as SDA Border Nodes, support IP fragmentation for native IP packets in the hardware data plane. Reducing the ip mtu on the handoff interface to 1500 **does not** cause performance degradation due to software-based fragmentation for traffic originating from or transiting the border that needs it. The switch efficiently fragments IP packets larger than 1500 bytes (if the DF bit is clear) before egressing this specific interface, without punting to the CPU.

However, it is important to note that **Catalyst 9000 switches generally do not support fragmentation of VXLAN encapsulated traffic.** This limitation is crucial for overlay traffic but **does not impact the described RADIUS authentication scenario**, as RADIUS communication between the SDA border and an external ISE typically occurs within the underlay (native IP routing). (MTU considerations for VXLAN overlays are a separate, complex topic, detailed in relevant Cisco SDA design guides).

Proactive MTU alignment at the SDA Border to SD-WAN Cisco Edge Router handoff is essential.

Challenge 2: The MTU Squeeze – ISE Traffic Across the SD-WAN Overlay

Even if individual physical interfaces, such as ISE Network Interface Card (NIC)s, switch ports, or router interfaces are set to a standard 1500-byte IP MTU, the SD-WAN overlay itself introduces encapsulation overhead. This overhead consumes a portion of the 1500-byte limit, reducing the effective MTU available for the original IP packet (the "payload" from ISE's perspective).

Packet Structure and Encapsulation Overhead:

When an IP packet from an ISE server (for example, a RADIUS Access-Accept packet) is sent to a Network Access Device (NAD) in an SDA site, it traverses the SD-WAN overlay and gets encapsulated. A common encapsulation stack involves IPsec in tunnel mode, potentially over UDP for NAT traversal (NAT-T).

- **Original Packet from ISE (Inner Packet):**
  For example, a RADIUS packet with a 1450-byte payload + 8B UDP + 20B Inner IP = 1478 bytes.

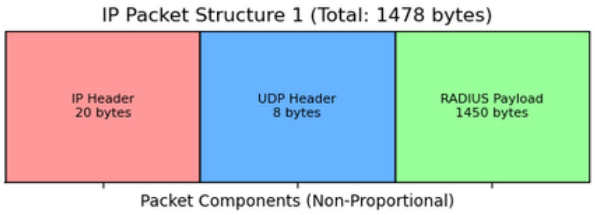
IP Packet Structure 1 (Total: 1478 bytes)
IP Header
20 bytes
UDP Header
8 bytes
RADIUS Payload
1450 bytes
Packet Components (Non-Proportional)

- Consider IPsec ESP in tunnel mode, potentially with UDP encapsulation for NAT-T:

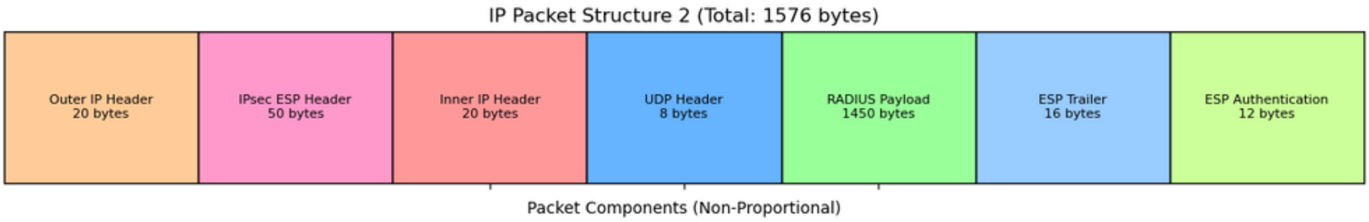

- The total overhead can vary based on the specific IPsec ciphers, authentication mechanisms, and other overlay features (like GRE if used). A typical calculation:
  - Outer IP Header (IPv4): 20 bytes
  - UDP Header (if ESP over UDP for NAT-T): 8 bytes
  - ESP Header: ~8 bytes
  - ESP IV (for example, for AES-CBC): ~16 bytes (if applicable)
  - ESP Authentication (for example, HMAC-SHA256 truncated): ~12-16 bytes
  - **Common Estimated IPsec Overhead:** ~52-70 bytes (can be higher, up to ~80 bytes or more with all options).

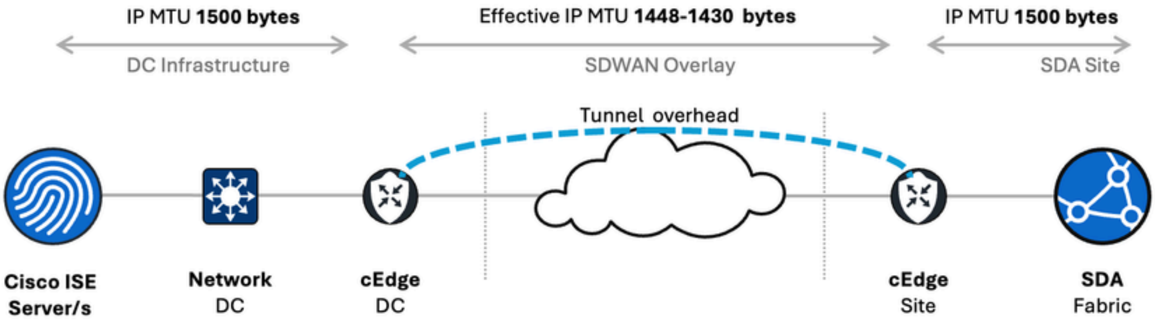

If the physical link MTU is 1500 bytes, the available payload MTU for the original IP packet from ISE becomes: 1500 bytes - SD-WAN Overhead.
For example, <u>1500 - 70 = 1430 bytes.</u>

Behavior When Packets Exceed Effective MTU:

1. **ISE Originates a Packet (The DF Bit Anomaly):**
   - By default, the underlying Linux operating system of an ISE appliance sets the**Do Not Fragment (DF) bit**in the IP header for all packets it originates that are smaller than or equal to its configured interface IP MTU (for example, 1500 bytes).
   - **Purpose of this DF Bit:**ISE (via its OS) proactively sets the DF bit primarily to leverage the Path MTU Discovery (PMTUD) process, which is described later. This allows ISE to dynamically learn the actual PMTU to a destination if it is smaller than its own interface MTU.
   - **Behavior for Packets Larger Than Interface MTU:**If ISE needs to send an IP packet that islargerthan its configured interface IP MTU, the behavior is dependent on its Linux operating system. Typically, the OScanfragment the packetbeforetransmission and clear the DF bit (setting DF=0) on these resulting fragments. This fragmentation is an OS-level function, not directly driven by the ISE application code itself.
   - **Key Distinction from Network Devices:**This default behavior of ISE (setting DF=1 even for non-fragmented packets fitting within its interface MTU) is significantly different from many traditional networking devices (routers, switches). Network devices often donotset the DF bit on packets they originate or forward unless explicitly configured to do so, or if the packet being forwarded already has the DF bit set, or for specific protocols that mandate it. They typically allow fragmentation by default if a packet exceeds the next-hop MTU (and DF=0).
   - **Troubleshooting Complexity:**This asymmetry, where ISE-to-NAD traffic often has DF=1 by default, while NAD-to-ISE traffic can have DF=0 (unless the NAD sets it for a reason)—can

introduce an additional layer of complexity during troubleshooting. Engineers can observe different fragmentation behaviors and PMTUD interactions depending on the direction of traffic flow.

2. **Packet Reaches Ingress Cisco Edge Router (DC):**The DC Cisco Edge Router router receives the IP packet from ISE.
3. **Encapsulation and MTU Check by Cisco Edge Router:**The Cisco Edge Router attempts to encapsulate the packet for the SD-WAN tunnel.
   - If the original packet sizeplusthe SD-WAN encapsulation overhead exceeds the Cisco Edge Router's outbound physical interface MTU (for example, 1500 bytes), AND the DF bit is set on the original (inner) packet from ISE, the Cisco Edge Router**must not**fragment the inner packet.
   - The Cisco Edge Router **should drop the packet.**
   - Critically, the Cisco Edge Router should also send an ICMP "Destination Unreachable - Fragmentation Needed and DF bit set" (Type 3, Code 4) message back to the source (ISE), indicating the MTU of the next hop (the effective MTU of the tunnel).
4. **Path MTU Discovery (PMTUD) Process:** Upon receiving this ICMP "Fragmentation Needed" message, ISE (the source OS) should reduce its PMTU estimate for that specific destination path. It would cache this information and resend the data in smaller packets that fit within the newly discovered PMTU.

PMTUD Process Diagram:

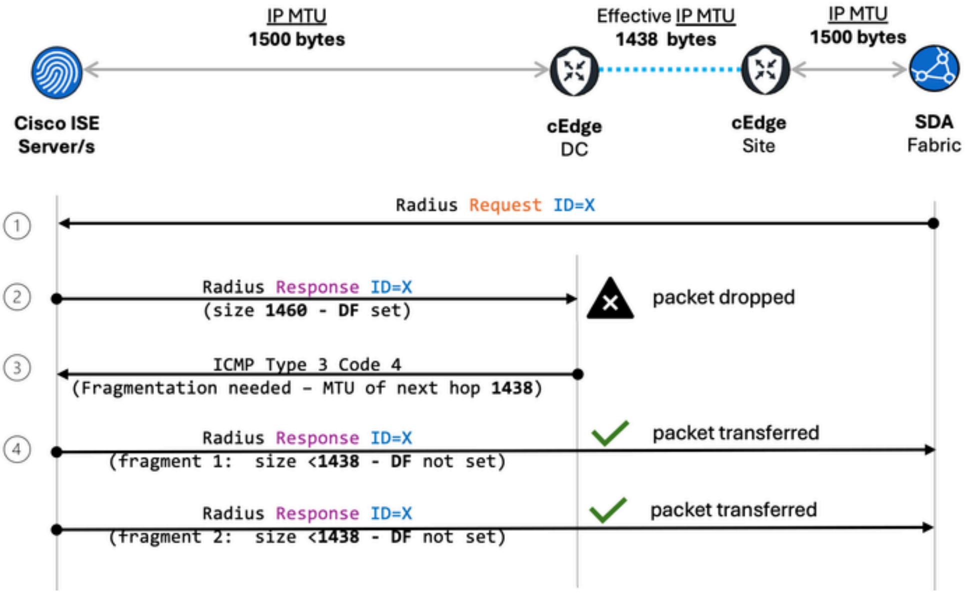

Where PMTUD Communication Breaks Down:
PMTUD is robust in theory but can fail in practice:

- **ICMP Filtering:** Intermediate firewalls or security policies often block ICMP messages, preventing the "Fragmentation Needed" message from reaching ISE.

- **Control Plane Policing (CoPP) on Cisco Edge Router:** Cisco Edge Router routers use CoPP to protect their CPU. Generating ICMP error messages is a control plane task. Under heavy load or with many oversized packets, CoPP can rate-limit or drop ICMP generation. ISE never receives the feedback.

- **Silent Drops:** If ISE does not receive the ICMP "Fragmentation Needed" message, it remains unaware of the path restriction. It continues to send large packets with the DF bit set, leading to them being silently dropped by the ingress Cisco Edge Router. This results in application-layer timeouts and retransmissions (for example RADIUS).

- **Impact on ISE Services:** Large RADIUS Access-Accept packets (carrying dACLs, extensive AVPs, SGT information) are particularly susceptible. Manifestations include:
    - Intermittent or complete authentication failures.
    - Endpoints not receiving correct network access policies or SGTs.
    - Incomplete or failed policy synchronization between ISE and NADs.

Solution to Challenge 2: Proactive ISE IP MTU Configuration

Given PMTUD unreliability, a proactive approach is best for critical services like ISE. Configure the IP MTU on ISE's network interfaces to a value that safely accommodates the maximum expected SD-WAN overlay overhead. This ensures ISE does not originate IP packets (with the DF bit set) that are inherently too large to traverse the SD-WAN overlay without needing fragmentation by an intermediate device (which is prohibited if DF=1).

Calculating and Setting the Recommended ISE IP MTU:

1. **Establish Base Physical MTU:** This is typicall 1500 bytes for standard Ethernet interfaces along the path.
2. **Determine Maximum SD-WAN Encapsulation Overhead:**
    - Accurately calculate or conservatively estimate the total overhead introduced by your specific SD-WAN overlay (IPsec, GRE, VXLAN, MPLSoGRE, and so on). Consult vendor documentation for precise figures for your chosen protocols and options.

Component	Example Overhead (Bytes)	Notes
Base Physical MTU	**1500**	Standard Ethernet on physical links
Less: SD-WAN Overhead		
Outer IP Header (IPv4)	20	
UDP Header (for NAT-T)	8	If ESP is encapsulated in UDP
ESP Header	~8-12	
ESP IV (for example, AES-CBC)	~16	Varies with encryption algorithm
ESP Auth (for example, SHA256)	~12-16	Varies with authentication algorithm (for example, 96-bit for some)
Other Overlays (GRE, and so on)	Variable	Add if part of your SD-WAN encapsulation stack
Total Estimated Overhead	**~68 - 80+ Bytes**	Sum of all relevant components for your deployment
Effective Path MTU	**~1432 - 1420 Bytes**	Base Physical MTU - Total Estimated Overhead

3. **Recommended ISE IP MTU Configuration:**
    - Take the calculated Effective Path MTU (for example, 1420 bytes from the example).
    - Subtract an additional safety margin (for example, 20-70 bytes) to account for minor, unaccounted L2 headers or to provide a buffer.
    - Solutions like Cisco SD-WAN can perform Path MTU (PMTU) discovery individually for each site-to-site tunnel. This mechanism runs automatically every 20 minutes to test and dynamically adjust the tunnel's IP MTU according to the current transport conditions at each site. As a result, MTU values can differ between sites and can change over time.
    - A generally safe and recommended IP MTU for ISE interfaces in such scenarios is between **1350 and 1400 bytes**

An IP MTU of **1350 bytes** is often a very robust starting point

ISE Configuration (Example via CLI):

This command is executed on the Cisco ISE appliance CLI for each relevant network interface.

<#root>

```
!
interface GigabitEthernet0  ! Or the specific interface used for RADIUS/SDA communication


 ip mtu 1350

!
```

Important Operational Considerations for ISE IP MTU Changes:

- **Service Restart Required:** One the ip mtu command is applied to an ISE interface, this prompts the user for a restart of ISE application services. This is a **service-impacting change** and **must be scheduled during a planned maintenance window.** Consult official Cisco ISE documentation for procedural details.
- **Apply to All ISE Nodes:** This IP MTU adjustment must be consistently applied to all ISE nodes in the deployment (Primary PAN, Secondary PAN, Policy Service Nodes (PSNs)) that communicate with NADs across the SD-WAN. Inconsistent MTU settings lead to unpredictable behavior.
- **Thorough Testing:** Before implementing in production, rigorously test this change in a lab or pilot deployment. Use tools like ping with varying packet sizes and the DF-bit set to validate end-to-end MTU handling:
  - **Linux-based systems:**

```
ping <destination_IP> -s <payload_size> -M do
```

(Note: -s specifies ICMP payload size. Total IP packet size = payload + 8B ICMP Hdr + 20B IP Hdr for IPv4)

  - **Windows:**

```
ping <destination_IP> -f -l <payload_size>
```

(Note: -l specifies ICMP payload size.)

  - **Cisco IOS/Cisco IOS-XE®**

```
ping <destination_IP> size <total_packet_size> df-bit
```

- **ISE First Routing Point** – When adjusting the IP MTU value on the ISE interface, ensure that the first routing point in the data center — specifically the Layer 3 interface associated with the ISE subnet — is also configured with the same IP MTU value.
  This helps prevent situations like the one described in *Challenge 1*, where an MTU mismatch causes ISE to treat incoming packets as oversized and drop them.
  For example, if the ISE interface has a reduced MTU (for example, 1300), but the first routing point remains configured with the default MTU of 1500, packets sent to ISE that are larger than 1300 bytes but smaller than 1500 bytes are not fragmented and are discarded by ISE — as observed in *Challenge 1*.
  Additionally, ensure that the first routing point is capable of performing fragmentation if needed, and that doing so does not result in performance degradation.

- **Update MTU across the entire transmission path, and in both direction** - When updating the IP MTU settings on ISE, it is important to consider the MTU across the entire transmission path, and in both directions. If the MTU value configured on ISE is not aligned with the MTU on the Layer 3 interface of the first-hop gateway, similar issues can arise as described in Challenge #1.

  For example, if the ISE MTU is reduced to 1300 bytes while the default 1500-byte MTU remains configured on the default gateway, packets between 1300 and 1500 bytes in size, commonly generated by network devices, can be dropped by ISE as oversized.

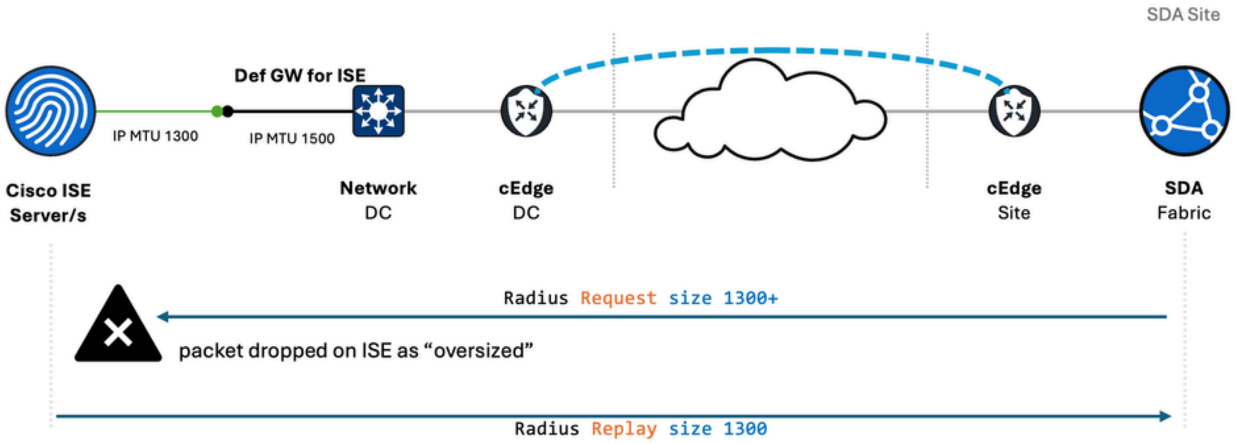

  To avoid this issue, always ensure that MTU changes on ISE are mirrored on the first-hop gateway and, ideally, reflected on all end hosts within the same Layer 3 segment. This helps maintain end-to-end MTU consistency and prevents unexpected packet drops.

Conclusion

Aligning the IP MTU settings on Cisco ISE servers with the effective transport-layer MTU limits imposed by SD-WAN encapsulation and MTU alignment at the SDA Border to SD-WAN Cisco Edge Router handoff is not just a recommendation but a critical prerequisite for ensuring the stability, reliability, and performance of AAA services in modern, segmented enterprise networks. While Path MTU Discovery is an important mechanism, its practical effectiveness can be hindered by factors like ICMP filtering or Control Plane Policing in SD-WAN environments.

By proactively configuring a reduced IP MTU on ISE (for example, 1350-1400 bytes), network architects and engineers can significantly mitigate the risk of MTU-related packet drops, leading to more predictable and resilient network operations. This is particularly vital in Cisco SDA deployments where ISE orchestrates sophisticated micro-segmentation and dynamic policy enforcement, which often rely on the successful delivery of potentially large control-plane messages. Diligent planning, comprehensive testing, and consistent configuration across all ISE nodes are key to a successful and trouble-free deployment.

Standards and References

For a deeper understanding, consult official standards and Cisco documentation:

RFCs:

- **RFC 1191:** Path MTU Discovery
- **RFC 791:** Internet Protocol (IP) - Defines the IP header, including the Do Not Fragment (DF) bit.
- **RFC 8200:** IPv6 Specification (Relevant if IPv6 is used, includes similar PMTUD concepts).
- **RFC 4459:** MTU and Fragmentation Issues with In-the-Network Tunneling (VPNs) - Directly addresses common MTU problems in VPN environments.

Cisco Documentation:

- **Cisco SDA Design and Deployment Guides:** For information on fabric MTU recommendations and border node configurations.
- **Cisco SD-WAN Design and Configuration Guides:** For details on encapsulation overhead, tunnel interface MTU, and PMTUD considerations within the SD-WAN fabric.
- **Cisco Catalyst 9000 Series Switch Configuration Guides:** For platform-specific details on MTU settings and fragmentation capabilities.
- **Cisco Identity Services Engine (ISE) Administrator and CLI Guides:** For information on interface configuration, including the ip mtu command and service restart implications.