

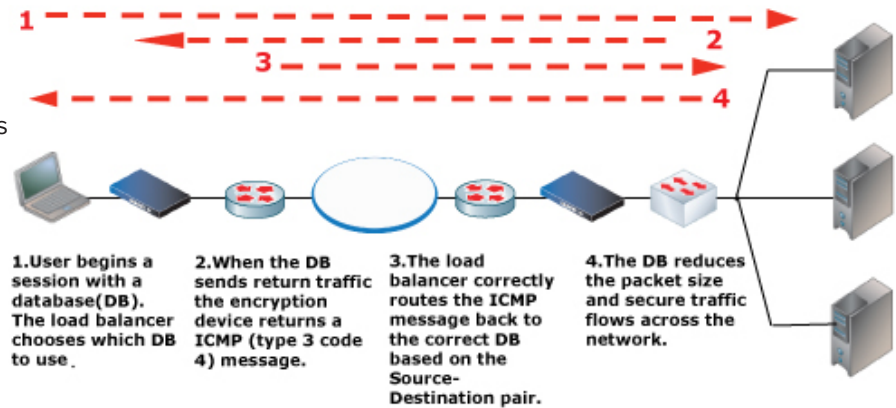
# Overcoming Encryption MTU Issues

*Issues and solutions for large packets when deploying network encryption*

## MTU and Encryption

The standard Maximum Transmission Unit (MTU), which defines the maximum packet size for networks, is 1500 bytes. AES encryption can add as many as 73 bytes to a packet, which makes the packet too big for 1500 byte networks.

If a device (router or otherwise) is asked to set up an encryption session and the “don’t fragment” bit is set to one (DF=1) then the encrypting device does the following:



1. Drops the packet (no encryption takes place, and the information is discarded).
2. Sends an error message to the device that sent the packet telling it to reduce the packet size to accommodate the possible additional 73 bytes added by AES encryption. The error message sent back is an “ICMP” message called ICMP (type 3, code 4).

The problem is that many devices block all ICMP messages. As a result, the original packet is thrown away and the sending device never receives the correction messages. The original sending device continues to send packets that get discarded. As far as the customer is concerned, “the network is broken”!

There are three solutions for this problem – each with its own pros and cons. The options are:

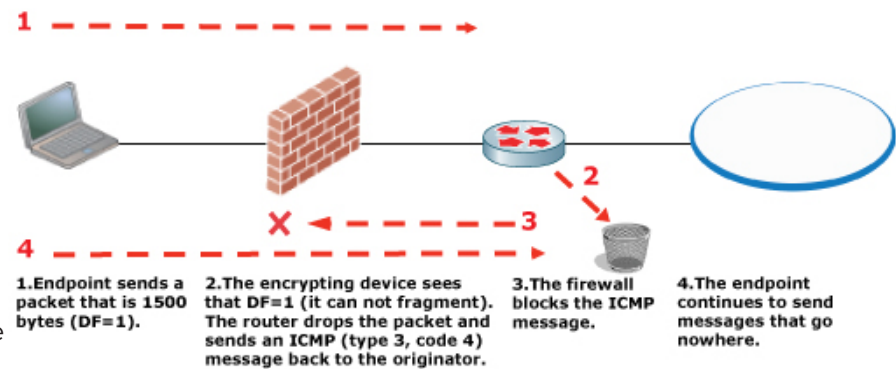
- Allow ICMP messages
- Allow Jumbo Frames
- Use Fragmentation and Reassembly

## Allow ICMP Messages

**Usage:** Certes Networks and all router-based encryption solutions

**What is it?** If ICMP messages are allowed, then the packet originator receives the message telling it to reduce the packet size. This solves the initial problem, but it can create a new one if load balancers are used.

**Potential Issues:** If a load balancer is used for server access, the ICMP message may get through the firewalls, but unless the load balancer session is maintained based on source-destination IP addresses, the ICMP message may not be received by the packet originator. This recreates the original problem where the traffic is discarded and the originator continues to send 1500 byte packets.

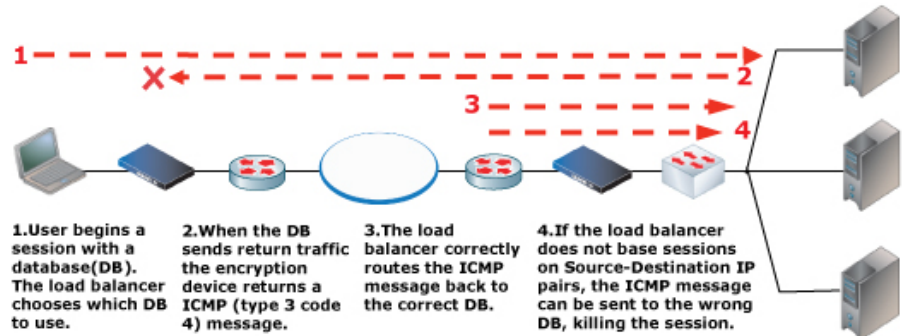


**ICMP and Load Balancers:** If load balancers are used, the sessions must be based on source-destination IP address pairs so that the ICMP message is received by the originator of the 1500 byte packet. The load balancer must have the ability to look inside the packet so that it can determine the original source and destination addresses. This is needed to load balance to the correct server.

### Allow ICMP Messages (continued)

**Dual NIC (Dual Home) Issue:** If the originating device is a server using dual NICs (Network Interface Cards), there is a known issue with certain NIC drivers that prevents the ICMP message from being processed correctly.

**Dual NIC Fix:** If the NICs are utilizing these drivers, the best work around is to use fragmentation and reassembly to solve the MTU issue. Another solution is to use drivers that support the ICMP (type 3 code 4) message correctly.



### Allow Jumbo Frames

**Usage:** Certes Networks and some router-based solutions

**What is it?** A jumbo frame is any packet larger than 1500 bytes. If a device supports jumbo frames it can encrypt any packet (regardless of size) without discarding it and sending the ICMP message back to the sender.

**Jumbo Frame Limitation:** Jumbo frames work only if all devices in the path support jumbo frames. This means the entire WAN or service provider network, and all devices in it, must support jumbo frames. On a private network this may not be an issue, but most ISPs have legacy devices, so this may not be an option.

### Use Fragmentation and Reassembly

**Usage:** Only Certes Networks offers the only solution that can fragment packets and reassemble at wire-speed and even redundant networks.

**What is it?** In this case the DF (don't fragment) bit is ignored. The packet is encrypted, fragmented and sent. On the receiving end, the two fragments are reassembled, de-encrypted, and the original packet is delivered to the destination.

**Load Balancing Issue:** If a load balancing solution is used, the two packet fragments may not be received by the same device, preventing reassembly and resulting in dropped packets. This can be avoided by basing the load balance session on Source-Destination IP addresses.

### Conclusion

Encryption is oftentimes blamed for "breaking the network". Understanding that the issue lies with the MTU size can help avoid downtime and user frustration. By enabling ICMP traffic, allowing jumbo frames or using Certes Networks enabled fragmentation and reassembly, this issue can be avoided.