Application Note

74-0058-170901

Connecting to the R5500 Across a Virtual Private Network

Virtual Private Networks (VPNs) provide secure links across remote sites over a public network. VPN protocols generally attempt to present the appearance of a seamless private network. However, there are known issues when dealing with large packets that generally occur during large file transfers and data streams such as the R5500.

This application note focuses on the Maximum Transmission Unit (MTU) and Path MTU aspects of the VPN tunnel and provides a solution to deal with large packets affected by a network with reduced MTU.

thinkRF

Contents

1 Overview

Virtual Private Networks (VPNs) are used to link remote networks over an untrusted public network. These links (VPN tunnels) are authenticated and encrypted for security. Tunnels behave like physical network interfaces from the router's viewpoint but with slightly reduced capacity due to tunnelling overhead.

In general, data is transferred across networks using IP datagrams of varying size. The maximum IP datagram size is dependent on the link with the smallest Maximum Transmission Unit (MTU) between the source and the destination. This MTU is referred to as the Path MTU (PMTU).

The following diagram is a typical scenario of two remote sites connected together via a VPN tunnel. The LAN and WAN ports on the VPN routers are Ethernet interfaces. The VPN tunnel uses a protocol such as IPsec to encrypt and encapsulate data between the two remote subnets.

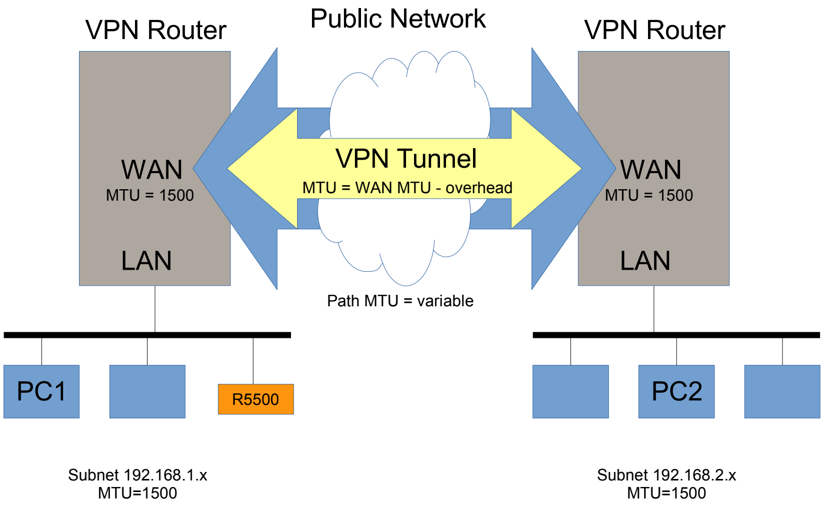

IP protocols are designed to be transparent to the connected devices, which do not have a priori knowledge of network topology outside of the LAN. In addition, network conditions can change dynamically, so the protocols are also designed to be adaptive over the life of a connection.

For transactions within a private LAN, this MTU is typically 1500 bytes, which is the legacy Ethernet payload limit. For transactions that span across a wider and more diverse network, the data path may cross other links with lower MTUs (e.g. radio links, serial lines, etc.), affecting the overall PMTU. In the diagram above, the PMTU is 1500 if PC1 is connected to the R5500 but is some arbitrarily lower number if PC2 is connected to the R5500.

IPsec is a protocol that is layered on top of the IP protocol that allows encrypted IP datagrams to be routed across a public network. It requires an additional 58-73 overhead bytes that eat into the available PMTU. Other layered protocols such as PPPoE further increase overhead.

The source VPN router processes incoming IP datagrams from the local LAN and prepares them to be sent over the public network. A new IP header with public IP addresses is added to allow it to be routed publicly. The destination VPN router reverses the process, taking incoming IPsec datagrams from the WAN and prepares them to be sent to the remote LAN.

Due to the overhead, large datagrams that would normally have fit within the 1500 byte limit will no longer do so. In such cases fragmentation and reassembly would be required.

Depending on configuration and/or implementation, the VPN router may perform one of the following operations based on datagram size:

a) No Fragmentation

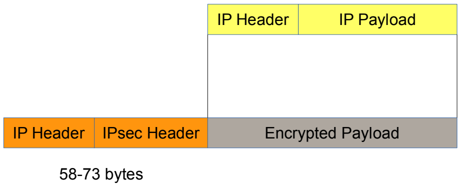

b) Pre-Fragmentation before IPsec Encryption

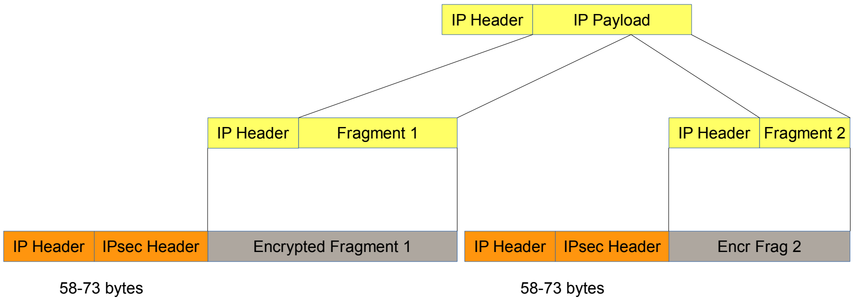

c) IPsec Encryption before IPsec Fragmentation

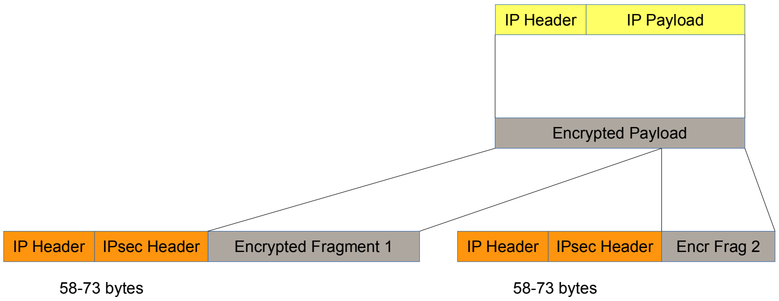

For most packets, no fragmentation is necessary and IPsec processing is straightforward (a). However for packets sizes that are close to the MTU, IPsec processing becomes more involved. The VPN router may pre-fragment the original IP datagram at the IP layer then encrypt the individual IP fragments (b), or it may encrypt the original IP datagram then fragment the encrypted data at the IPsec layer (c). Large datagrams occur during large data transfers such as R5500 captures.

Fragmentation is computationally expensive, especially if done by routers that process the typical case in hardware and deal with exception cases using a processor. It also requires temporary storage of partial fragments. In many cases the overall throughput can degrade to a fraction of the expected rate. Best practices avoid fragmentation altogether with protocols such as Path MTU Discovery. In fact, fragmentation has been removed in the IPv6 specification.

2 Path MTU Discovery and ICMP

Path MTU Discovery (PMTUD) is a mechanism that is used by a host to determine the PMTU to a given destination. The goal is to avoid IP fragmentation because of its computational cost. It works in conjunction with Internet Control Message Protocol (ICMP), which is a set of messages used to send error and diagnostics messages between network nodes (routers and hosts).

IP datagrams have an optional “Don't Fragment” (DF) flag, which informs routers whether or not IP fragmentation is allowed. The source host discovers the PMTU by attempting to send large IP datagrams with the DF flag set. When such a datagram is received by an intermediate router that cannot forward it without fragmentation due to a lower next-hop MTU, it sends an ICMP “Destination Unreachable - Fragmentation Needed and Don't Fragment was Set” (ICMP Type 3/Code 4) message back to the source. The source then retries using progressively smaller datagrams until it is successful and sets the PMTU accordingly.

PMTUD operates dynamically in case the network topology changes; it attempts to rediscover the PMTU periodically (every few minutes) by increasing its PMTU until transmissions start to fail again.

PMTUD is used for both discovering the PMTU of the WAN port link between the routers as well as the PMTU of the path through the VPN tunnel itself. This nested use of PMTUD is problematic.

Many network firewalls are set up to reject ICMP messages from WAN ports for security reasons. Doing so effectively renders PMTUD ineffective across a public network. Large datagrams can be silently lost, creating an IP “black hole”. The loss is only discovered at the IP layer when multiple attempts to retransmit the lost datagram also fail.

3 Solution

The interaction between VPN tunnels and PMTUs is well documented. Datagrams may be lost without warning, causing transmission retries and eventual connection failure.

One way to avoid these various issues is to artificially reduce the MTU of the local Ethernet interface itself such that the PMTU is no longer dependent on outside factors.

The choice of MTU value is dependent on your particular network. At the very least the MTU should be reduced to account for IPsec overhead. Other factors that may affect the choice of MTU include:

- DSL connection running PPPoE
- Reduction of the WAN MTU by the system administrator to address PMTUD issues across the WAN link
- VLANs

The most expedient way to pick the MTU value is by trial and error. Note that lowering the MTU beyond what is required has an impact on overall throughput.

As of Firmware Version 1.4.6, the R5500 now allows configuration of its Ethernet interface MTU using the following SCPI command:

```
:SYSTEM:COMMUNICATE:LAN:MTU <mtu>
```

Please refer to the R5500 *Programmer's Guide* for more details.

Document Revision History

This section summarizes document revision history.

Document Version	Release Date	Revisions and Notes
v1.0	Aug 15 2017	First release

Contact us for more information

ThinkRF Support website provides online documents for resolving technical issues with ThinkRF products at http://www.thinkrf.com/resources.

For all customers who hold a valid end-user license, ThinkRF provides technical assistance 9 AM to 5 PM Eastern Time, Monday to Friday. Contact us at support@thinkrf.com, sales@thinkrf.com or by calling **+1.613.369.5104**.

These specifications are preliminary, non-warranted, and subject to change without notice.