Determining Risk of Tunnels Over The Internet

Sandeep Muthu [1], Yuri Pradkin [2], John Heidemann [2]
1: National Institute of Technology, Tiruchirappalli, 2: USC/Information Sciences Institute

Introduction

Tunnels are used over the internet to serve multiple purposes such as VPNs, cloud services and for IPv6 transition.
But are they always secure?

Some tunneling protocols not authenticated or encrypted.
An attacker can inject fake traffic into these tunnels!

We:
- demonstrate some tunnels can be exploited—they pose a risk
- measure how prevalent tunnels are
- describe how to prevent this attack

How does a tunnel work?

A tunnel transports data between two private networks over the public Internet. It wraps the original packet with a header before sending it over the internet.

Below, the green networks (C1 and C2) are private.
They're joined by a tunnel between R1 and R2.

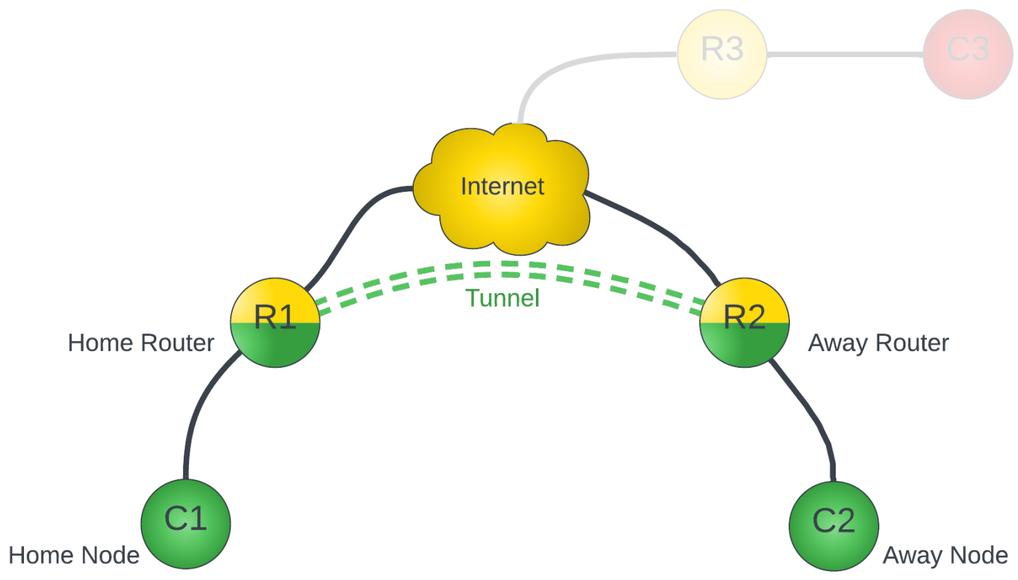

When C1 sends to C2, the packet P goes to R1:
1. R1 encapsulates the original packet P -> E(P) with an extra header.
2. Packet E(P) is **transmitted through the public Internet** from R1 to R2.
3. When E(P) arrives at R2, **R2 decapsulates** it, discarding the extra header and sending it to the internal network C2.

Although some tunneling protocols are authenticated, we will see how a lack of them can lead to attacks.

Tunnels are used for sending IPv6 over non-IPv6 networks, for bridging private networks, and to implement cloud virtualization.

Are Tunnels A Security Risk?

Tunnels that lack authentication and encryption seem vulnerable to traffic injection.
Can we demonstrate this attack in the lab?
Setup:
- Private networks (c1/r1 and c2/r2) are connected by a tunnel
- An attacker (c3) discovers this tunnel's endpoints
- Can the attacker inject traffic into the tunnel?
- We conduct this experiment on the testbed

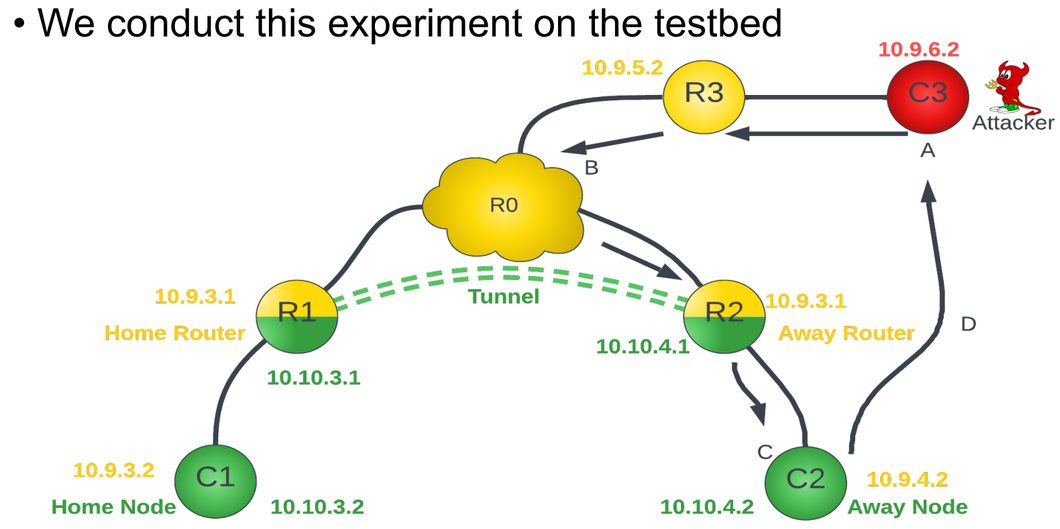

Attack:
A. Attacker C3 fakes a packet—it appears as if it's from R1—sends to R2
B. This fake packet is sent regularly (R3 -> R0->R2)
C. R2 gets it, can't tell it's fake, strips the fake header and sends it to C2. ***The internal network now gets traffic from the attacker!***
D. C2 assumes the packet is good and replies. ***Internal information escapes!***

Our experiment shows an external attacker can exploit lack of authentication to inject or extract information from the private network.
Future work could explore more dangerous attacks, like two-way traffic.

Are Tunnels Used in Real Networks?

We showed that attacks are possible. But how widely are these tunneling protocols used on the internet?

We looked at data from an Internet Exchange Point in Colorado to see how many tunnels occur in one hour.

Protocol	Unique IPs	Prevalence	Authenticated/Encrypted
GRE	16381	78.29%	No/No
6in4	277	21.54%	No/No
IPIP	67	0.12%	No/No

We see that there are tens and thousands of people using these tunneling protocols which have no authentication/encryption.

Next Steps

- More sophisticated attacks: Can we an evil proxy allow interactive access to the private network?
- Can we alert those potentially vulnerable?
- Can we confirm the unique IPs are vulnerable?
- Can we encourage users to use tunneling protocols which offer better security?

Conclusion

- **Vulnerable tunneling protocols are in use today** over the Internet
  (more than 16k active IPs at just one exchange point)
- **External attackers can access private networks** using unauthenticated tunnels
- **We need better tunnel security** to reduce these risks
- We are writing a paper to describe these results

If interested contact Sandeep Muthu 108120109@nitt.edu
Work performed under REU Site program
supported by NSF grant #2051101

USC Viterbi School of Engineering
Information Sciences Institute