Increasing Throughput in IP Networks with ESP-Based Forwarding:
ESPBasedForwarding

Abstract

   This document proposes an experimental way of reaching infinite
   bandwidth in IP networks by the use of ESP-based forwarding.

Status of This Memo

   This document is not an Internet Standards Track specification; it is
   published for examination, experimental implementation, and
   evaluation.

   This document defines an Experimental Protocol for the Internet
   community. This is a contribution to the RFC Series, independently
   of any other RFC stream. The RFC Editor has chosen to publish this
   document at its discretion and makes no statement about its value for
   implementation or deployment. Documents approved for publication by
   the RFC Editor are not a candidate for any level of Internet
   Standard; see Section 2 of RFC 5741.

   Information about the current status of this document, any errata,
   and how to provide feedback on it may be obtained at
   http://www.rfc-editor.org/info/rfc5984.

Table of Contents

1. Introduction

   Mechanisms for efficient packet forwarding has evolved over the past
   years. The demand for bandwidth is always increasing. Instead of
   optimizing forwarding performance and link capacity in an incremental
   fashion, we propose a brand new concept in packet forwarding that
   will provide unsurpassed end user performance regardless of link
   capacity, distance, and number of hops.

1.1. Requirements Language

   The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT",
   "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this
   document are to be interpreted as described in RFC 2119 [RFC2119].

2. Background

   During the past years, there have been a lot of improvements made in
   the domain of packet forwarding. Both software and hardware
   optimizations combined with increased link capacities have cut down
   round-trip times. Despite these improvements, many users find
   themselves frustrated since their demand for bandwidth has increased
   faster than the supply.

The current incremental approach to lower latency and increase
capacity will soon reach the end of the road. The reason for this
has been known for some time and is stated in RFC 1925 [RFC1925]
clause 2:

"(2) No matter how hard you push and no matter what the priority, you
can't increase the speed of light."

Our research has finally been able to circumvent this boundary by the
development of zero-latency network paths.

Inspired by RFC 1072 [RFC1072], where a network containing a long,
fat pipe is called LFN (pronounced "elephan(t)"), we will refer to an
internet path with zero-latency as "infinitely fat", and a network
containing this path as "IFN" (pronounced "infan(t)").

Before the invention of this new forwarding principle, several
experimental methods were tried. We have chosen to share our failed
attempts in order help others avoid the same mistakes that we
encountered. The following two methods have been dismissed:

o Black Fiber
o Schrodinger's cat experiment

2.1. Experiments with Black Fiber

Attempting to push the speed-of-light limitation by means of using
black fiber looked promising at first. Shortly after initiating the
experiment, lack of light was detected in the black fiber. This was
interpreted as proof of successful data transmission faster than the
speed of light. After popping the champagne, the following problems
were detected:

1. No data reached the receiver.
2. The fiber was not connected at the transmitting side.

This clearly spoiled the mood of the party.

2.2. Schrodinger's Cat Experiment

The Schrodinger's netcat experiment was based on an attempt to
implement the method described by E. Schrodinger [Schrodinger35].
The original procedure includes locking up cats in boxes with
radioactive materials and poisonous gas. Data communication
capabilities were added to the experiment, by using netcat. The
research team was dumbfounded by the notion that the cat experiment
seemed to work and not work at the same time. This was also
confirmed by a friend of Wigner's [Wigner].

   A detailed analysis of the experiment indicated that the probability
   vectors collapsed whenever traffic was sent to the box. The
   conclusion was that this approach would only work without traffic,
   thus eliminating all practical applications.

3. ESP-Based Forwarding

   Experiments with ESP-based (Extra Sensory Perception) forwarding has
   proved to successfully remove the limitation in RFC 1925 [RFC1925].

   The foundation for the ESP-based forwarding scheme is to reduce
   latency by means of precognitive datagram detection and generation.
   By applying this technology, latency will not only reach zero, but
   might even become negative.

   Experiments performed by Benjamin Libet [Libet85] regarding the
   readiness potential (Bereitschaftspotential) concludes that the end
   user latency from impulse to the conscious mind is approximately 350
   - 400 ms. In order to reach the highest possible data transport
   without confusing the end user, it is important to take this latency
   into consideration.

3.1. Principle of Operation

   Traffic between the end user and the server reaches the ESP-enabled
   router. Inside the ESP-based router, the data stream is first
   analyzed by the DPAUI (Deep Packet And User Inspection). The DPAUI
   sends a signal to the PPG (Deep Packet And User Inspection), which
   generates uplink IP datagrams supported by the IID (Infinite
   Improbability Drive). The generated IP datagram is sent to the CFE
   (Clairvoyant Forwarding Engine) that forwards the datagram. Finally,
   the "real" uplink, the end user datagram is received and forwarded to
   the ND (Null Device).

3.2. Architectural Components

   The current ESP-based forwarding architecture includes the following
   components:

   o DPAUI
   o PPG
   o IID
   o CFE
   o PPS
   o ND

3.2.1. DPAUI

   The DPAUI (Deep Packet And User Inspection) monitors the data streams
   for all individual users. The DPAUI is implemented by means of
   implementing a learning agent that analyzes each individual user.
   The output from the DPAUI is a signal that indicates that an IP
   datagram will be sent by the end user within a couple of seconds.

3.2.2. PPG

   The purpose of the PPG (Precognitive Packet Generator) is to generate
   the IP datagram that the end user will trigger to be sent. In order
   to craft such a datagram, the PPG will perform a lookup based on the
   offset and length parameters generated by the IID. The PPG will
   generate the future packet by means of the function:

   struct mbuf * CopyDatagramFromPi(
                   insane long offset,
                   unsigned int len);

   The CopyDatagramFromPi() function will return a pointer to an mbuf
   containing the IP datagram. The offset value and len matches a
   corresponding offset and length in the decimal set for pi that
   contains the bit pattern for the future datagram. This method of
   operation will reduce the complex matter of precognitive packet
   generation to a simple lookup.

   Concerns have been raised that the full decimal set of pi requires an
   infinite amount of memory. This might have a negative effect on the
   manufacturing cost of the router. These concerns were successfully
   managed by using a perfectly circular buffer. This reduced the
   previous stated memory requirements at least by half.

3.2.3. IID

   The purpose of the IID (Infinite Improbability Drive) is to assist
   the PPG and PPS with improbable probabilities (and the other way
   around). The IID was originally invented by Douglas Adams [Adams79].
   The original implementation was based on hooking up the logic
   circuits of a Bambleweeny 57 sub-meson Brain to an atomic vector
   plotter suspended in a strong Brownian motion producer (i.e., a nice
   hot cup of tea).

   The research team struggled with the implementation of the strong
   Brownian motion producer. As a matter of fact, the majority of the
   research budget was wasted before it was fully conceived that a warm
   cup of tea and router equipment rarely mix.

Aided by the gastronomical department (working on Bistromathic
Drive), the research team managed to attach a brownie on top of a
radio controlled hovercraft full of eels. This not only caused a lot
of noise and disarray, but also a sufficient amount of Brownian
motion. The research team is still working on an entirely software-
based solution. Hopefully, the eel-filled hovercraft will soon be
replaced with a different type of python script.

3.2.4. CFE

After the IP datagram has been produced by the PPG, the CFE
(Clairvoyant Forwarding Engine) will attempt to find the right route.
Since the route might not exist yet, direct access to a routing table
might result in routing errors.

The implementation of the CFE is very straightforward: any off-the-
shelf routing stack with a routing table and a routing daemon can be
used. A standard Ouija board is simply put on top of the routing
table. For each datagram, the CFE initiates an Ouija board session
that will establish a connection with the routing deamons. The CFE
will seek guidance for the future of the IP datagram and then send it
along for a low cost, to meet a tall, dark server rack.

3.2.5. PPS

The PPS (Pre-Preemptive Scheduler) is synchronized by means of an NTP
connection to the IID based NTP server. This ensures that the ESP
process will execute ten seconds ahead of local time (layman's term:
"into the future"). This value should ensure operation even with
very long Round Trip Times and should also include the readiness
potential of the end user.

The pre-preemptive scheduler not only provides a separate user space,
but a separate dimension for the code to execute in. The dimension
alignment is based on string theory and has been implemented in the
language C, simply by including the library string.h and then using
strcpy to copy the PPS function pointer into an eleven-dimensional
array.

3.2.6. ND

After a little time, less than the 'end user to server' Round-trip
time (RTT), the real end user datagram will reach the ingress side of
the ESP-based router, since the datagram has already been sent,
routed and returned. The datagram is directly routed to the ND (Null
Device) and the ingress packet counter is decremented.

Experimentation showed that the ND is a perfect source of entropy and
is able to store all digits of pi. The research team had great hopes
of reducing the memory footprint for the PPG even further, but ran
into problems with read access times.

The ND is readily available in most operating systems.

4. End User Considerations

End user considerations and differentiated traffic classes:

1. In order to facilitate a pleasant end user gaming experience,
    packets destined for the spinal cord (i.e., force feedback
    information, etc.) must be delayed by 350 ms in order to
    synchronize with the traffic that is routed by the end user to
    the cerebrum and cortex.

2. RFC 1216 [RFC1216], Section 3.3 states that "bad news travels
    fast". This means that additional delay must be introduced when
    the end user is browsing on news sites. Negative latency might
    make the end user suspect that the news is even worse than
    indicated by the news sites.

3. Machine-to-Machine (M2M) communication might experience reduced
    performance due to difficulties for the DPAUI to work correctly.
    When the concept starts working for M2M communication, this can
    be used as an indication that a technological singularity might
    be near.

5. TCP Slow-Start Considerations

The lack of RTT of IFNs might cause some problems with TCP slow-
start. More precisely, it will most likely not be slow at all. This
might be handled by implementing a congestion avoidance mechanism,
but will need further study.

6. Market Considerations

Unfortunately, we foresee that this product will never be ready for
the market. This is especially true for the Pre-preemptive
Scheduler, which by nature, will always be slightly ahead of its
time.

7. Security Considerations

o Introducing an end user RTT delay of zero might cause crashes in
   badly implemented TCP/IP stacks. This is because division by zero
   might occur when calculating bandwidth-delay product.

   o ESP forwarding of traffic generated by psychics might lead to
      problems with recursiveness.

   o Lawful Intercept of the Deep User and Intention Inspection might
      violate personal integrity.

   o Terrorist organizations might exploit the "bad news travels fast"
      loophole in RFC 1216 [RFC1216].

8. References

8.1. Normative References

   [RFC2119]   Bradner, S., "Key words for use in RFCs to Indicate
                       Requirement Levels", BCP 14, RFC 2119, March 1997.

8.2. Informative References

   [Adams79]   Adams, D., "Hitchhiker's guide to the galaxy.",
                       1979.

   [Libet85]   Libet, B., "Unconscious cerebral initiative and the
                       role of conscious will in voluntary action.", 1985.

   [RFC1072]   Jacobson, V. and R. Braden, "TCP extensions for
                       long-delay paths", RFC 1072, October 1988.

   [RFC1216]   Richard, P. and Kynikos, "Gigabit network economics
                       and paradigm shifts", RFC 1216, April 1991.

   [RFC1925]   Callon, R., "The Twelve Networking Truths",
                       RFC 1925, April 1996.

   [RFC1928]   Leech, M., Ganis, M., Lee, Y., Kuris, R., Koblas,
                       D., and L. Jones, "SOCKS Protocol Version 5",
                       RFC 1928, March 1996.

   [Schrodinger35]   Schrodinger, E., "The Present Situation In Quantum
                       Mechanics", 1935,
                       <http://www.tu-harburg.de/rzt/rzt/it/QM/cat.html>.

   [Wigner]   Wikipedia, "Wikipedia: Wigner's friend.",
                       <http://en.wikipedia.org/wiki/Wigner's_friend>.

Author's Address

    Karl-Magnus Moller
    Tankesaft

    EMail: kalle@tankesaft.se