

OneWayDelay

One-way Delay (OWD)

One-way delay is the time it takes for a packet to reach its destination. It is considered a property of network links or paths. [RFC 2679](#) contains the definition of the one-way delay metric of the IETF's IPPM (IP Performance Metrics) working group.

Decomposition

One-way delay along a network path can be decomposed into per-hop one-way delays, and these in turn into per-link and per-node delay components.

Per-link delay components: propagation delay and serialization delay

The link-specific component of one-way delay consists of two sub-components:

[Propagation Delay](#) is the time it takes for signals to move from the sending to the receiving end of the link. On simple links, this is the product of the link's physical length and the characteristic *propagation speed*. The velocity of propagation (VoP) for copper and fibre optic are similar (the VoP of copper is slightly faster), being approximately 2/3 the speed of light in a vacuum.

[Serialization delay](#) is the time it takes for a packet to be serialized into link transmission units (typically bits). It is the packet size (in bits) divided by the link's capacity (in bits per second).

In addition to the propagation and serialization delays, some types of links may introduce additional delays, for example to avoid collisions on a shared link, or when the link-layer transparently retransmits packets that were damaged during transmission.

Per-node delay components: forwarding delay, queueing delay

Within a network node such as a router, a packet experiences different kinds of delay between its arrival on one link and its departure on another (or the same) link:

Forwarding delay is the time it takes for the node to read forwarding-relevant information (typically the destination address and other headers) from the packet, compute the "forwarding decision" based on routing tables and other information, and to actually forward the packet towards the destination, which involves copying the packet to a different interface inside the node, rewriting parts of it (such as the IP TTL and any media-specific headers) and possibly other processing such as fragmentation, accounting, or checking access control lists.

Depending on [router architecture](#), forwarding can compete for resources with other activities of the router. In this case, packets can be held up until the router's forwarding resource is available, which can take many milliseconds on a router with CPU-based forwarding. Routers with dedicated hardware for forwarding don't have this problem, although there may be delays when a packet arrives as the hardware forwarding table is being reprogrammed due to a routing change.

Queueing delay is the time a packet has to wait inside the node waiting for availability of the output link. Queueing delay depends on the amount of competing traffic towards the output link, and of the priorities of the packet itself and the competing traffic. The amount of queueing that a given packet may encounter is hard to predict, but it is bounded by the [buffer size](#) available for queueing.

There can be causes for queueing other than contention on the outgoing link, for example contention on the node's backplane interconnect.

Impact on end-to-end performance

When studying end-to-end performance, it is usually more interesting to look at the following metrics that are derived from one-way delay:

- [Round-trip time \(RTT\)](#) is the time from node A to B and back to A. It is the sum of the one-way delays from A to B and from B to A, plus the response time in B.
- [Delay Variation](#) represents the variation in one-way delay. It is important for real-time applications. People often call this "jitter".

Measurement

One-way delays from a node A to a node B can be measured by sending timestamped packets from A, and recording the reception times at B. The difficulty is that A and B need clocks that are synchronized to each other. This is typically achieved by having clocks synchronized to a standard reference time such as UTC (Universal Time Coordinated) using techniques such as Global Positioning System (GPS)-derived time signals or the Network Time Protocol (NTP).

There are several infrastructures that continuously measure one-way delays and [packet loss](#): The [HADES Boxes](#) in DFN and GÉANT2; [RIPE TTM Boxes](#) between various research and commercial networks, and RENATER's [QoSMetrics boxes](#).

[OWAMP](#) and [RUDE/CRUDE](#) are examples of tools that can be used to measure one-way delay.

Improving delay

Shortest-(Physical)-Path Routing and Proper Provisioning

On high-speed wide-area network paths, delay is usually dominated by propagation times. Therefore, physical routing of network links plays an important role, as well as the topology of the network and the selection of routing metrics. Ensuring minimal delays is simply a matter of

- using an internal routing protocol with "shortest-path routing" (such as OSPF or IS-IS) and a metric proportional to per-link delay
- provisioning the network so that these shortest paths aren't congested.

This could be paraphrased as "use proper provisioning instead of traffic engineering".

The node contributions to delay can be addressed by:

- using nodes with fast forwarding
- avoiding queueing by provisioning links to accomodate typical traffic bursts
- reducing the number of forwarding nodes

References

- [RFC 2679: A One-way Delay Metric for IPPM](#), G. Almes, S. Kalidindi, M. Zekauskas, September 1999
- [Understanding Delay in Packet Voice Networks](#), a Cisco Systems publication explaining sources and effects of delay, both in general and for voice in particular. Includes a serialization delay table with different packet (frame) sizes at different (low) link rates.

– Main.SimonLeinen - 31 Oct 2004 - 25 Jan 2007