Module 2: Network Performance and User Expectations

WHAT FACTORS SHAPE USER EXPECTATIONS?

Users expect good network performance because of:

- Publicised bandwidth statistics.
  - For example the GÉANT2 fact sheet says, “Many routes operate at 10Gbps – speeds which equate to transferring 1,000 digital photos in 1.6 seconds. The total network capacity is over 500 Gbps – 2.5 times the performance of the first GÉANT network.”
- Applications’ requirements:
  - Examples:
    - High levels of responsiveness required for video-conferencing.
    - Fast throughput required for bulk transfers.

USERS' PERCEPTION OF PERFORMANCE

Users' perception of actual network performance is shaped by:

- Responsiveness.
  - E.g. degree of latency in a video-conference.
- End-to-end throughput.
  - E.g. how fast data moves from one host/application/fille to another.
- Reliability.
  - Can be subdivided into:
    - Availability of services.
    - Predictability of performance.

GEANT2

THE 'WIZARD GAP'

Theoretically possible performance = high.

- But optimal network performance only achieved by:
  - Expert tuning.
  - 'Experiments' carried out in conducive 'laboratory' conditions.
- See http://www.internet2.edu/lsr/ for land speed record.

Users' perceptions of performance = lower.

- Examples:
  - There is frustrating latency in a video-conference.
  - It takes too long to download a file.

The difference is the 'wizard gap'.

GEANT2

WHAT FACTORS REALLY SHAPE PERFORMANCE?

The factors that actually influence network performance are:

- One-way delay (OWD).
- Round-Trip Time (RTT).
- One Way Delay Variation (OWDV - also known as jitter).
- Packet re-ordering.
- Packet loss.
- Maximum Transmission Unit (MTU).

GEANT2

ONE-WAY DELAY (1)

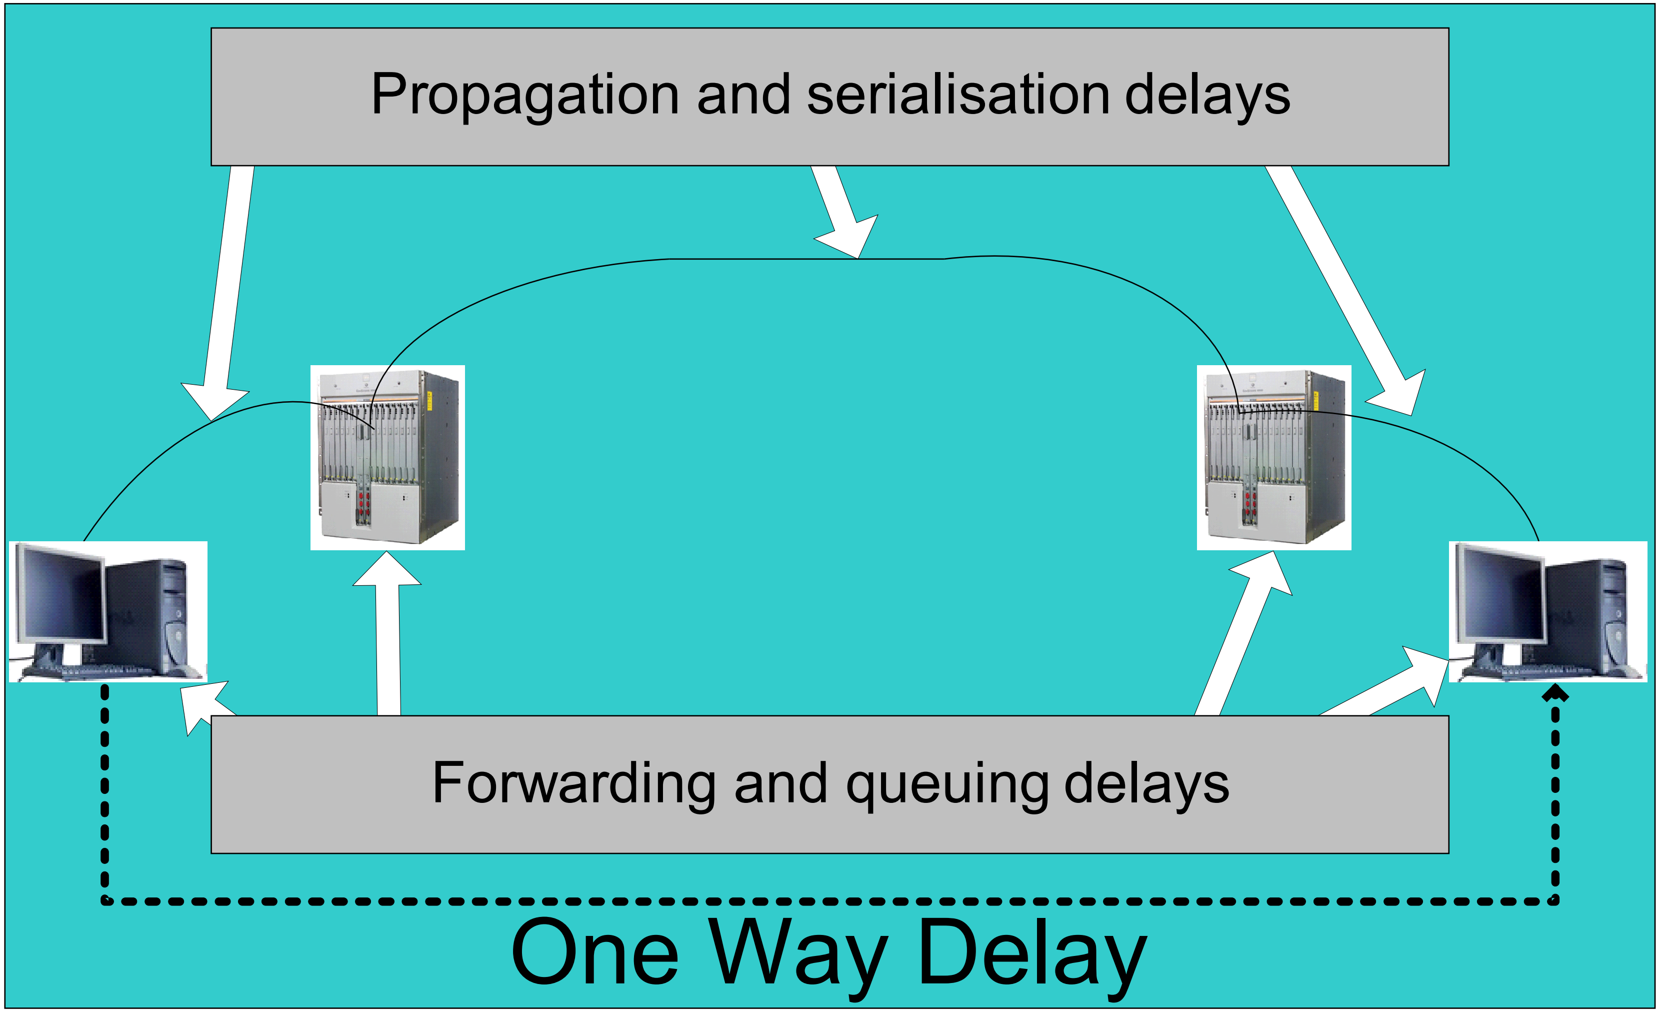

ONE-WAY DELAY (2)

What is one-way delay (OWD)?

- The time it takes for a packet to reach its destination.

A path's one-way delay can be divided into per-hop delays.

Per-hop delays can themselves be divided into:

- Per-link delay.
  - Made up of propagation delay and serialisation delay.
- Per-node delay.
  - Made up of forwarding delay and queuing delay.

ONE-WAY DELAY (3)

Serialisation delay for a 1500 byte packet:

- 10 Mbps – 1 ms.
- 100 Mbps – 0.1 ms (100 µs).
- 1 Gbps – 0.01 ms (10 µs).
- 10 Gbps – 0.001 ms (1 µs).

Propagation delay in a fibre per 100km: 0.5 ms.

Forwarding delay is typically constant in hardware-based forwarding engines, many orders of magnitude smaller.

Propagation and queuing delays are the most important factors in OWD.

IMPROVING DELAY

Steps to shorten delay:

- Minimise propagation times by:
  - Using ‘shortest-path’ routing.
    - E.g. OSPF or IS-IS.
  - Provisioning network so that shortest paths are not congested.
    - even over short periods (“overprovisioning”).
- Improve node performance by:
  - Using nodes with fast forwarding.
    - Make sure “hardware forwarding” is used for all (relevant) traffic!
  - Provisioning links to accommodate typical traffic bursts.
    - Avoids queuing.

GEANT2

ROUND TRIP TIME (1)

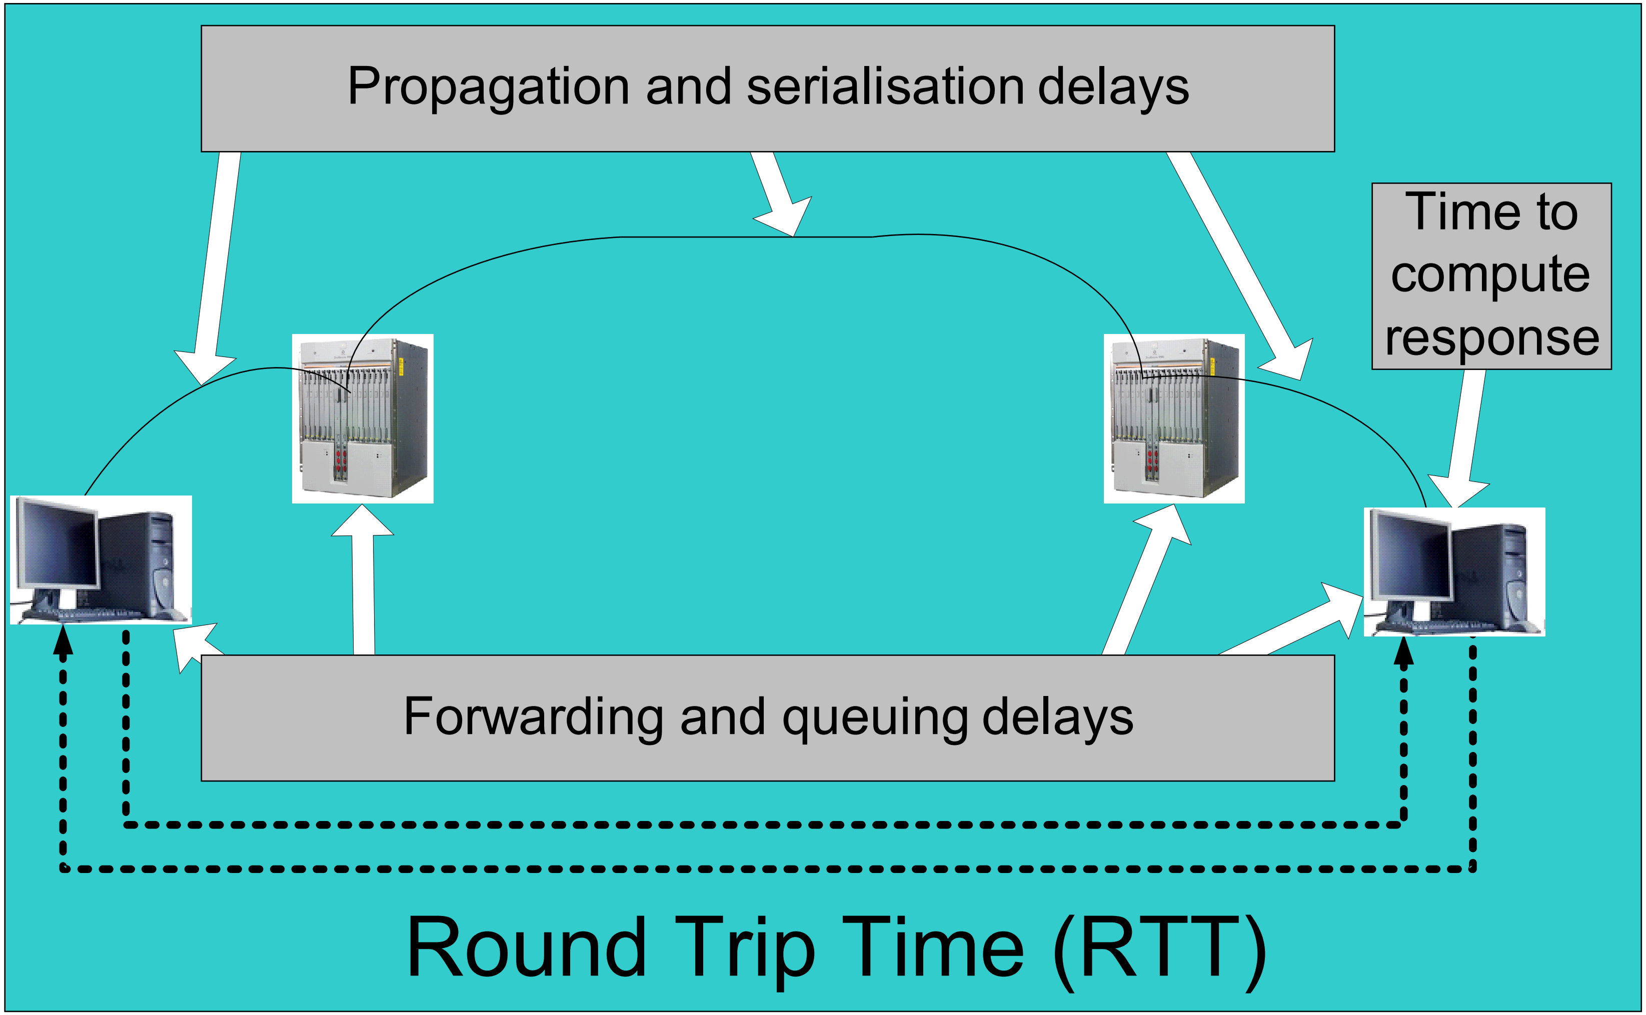

ROUND TRIP TIME (2)

Round Trip Time (RTT) is the sum of two one-way journeys:

- Data sent from one node to another.
- Acknowledgement of receipt sent back.
- Plus the time that the destination node takes to compute a response.

RTT Significantly influences throughput:

- Buffers at TCP endpoints must support *rate*RTT* window.
- High RTT means TCP will be slow to reach max. speed.
  - As well as to recover from congestion.

ROUND TRIP TIME (3)

Round trip time:

- Particularly important for ‘interactive’ applications such as video conferencing.
  - The ‘response’ time / latency can never be better than the round trip time.
- Can be measured using:
  - Ping and its variants.
- Can be improved by addressing one-way delay.
  - Since RTT is the sum of two one way journeys.

DELAY VARIATION: AN EXAMPLE

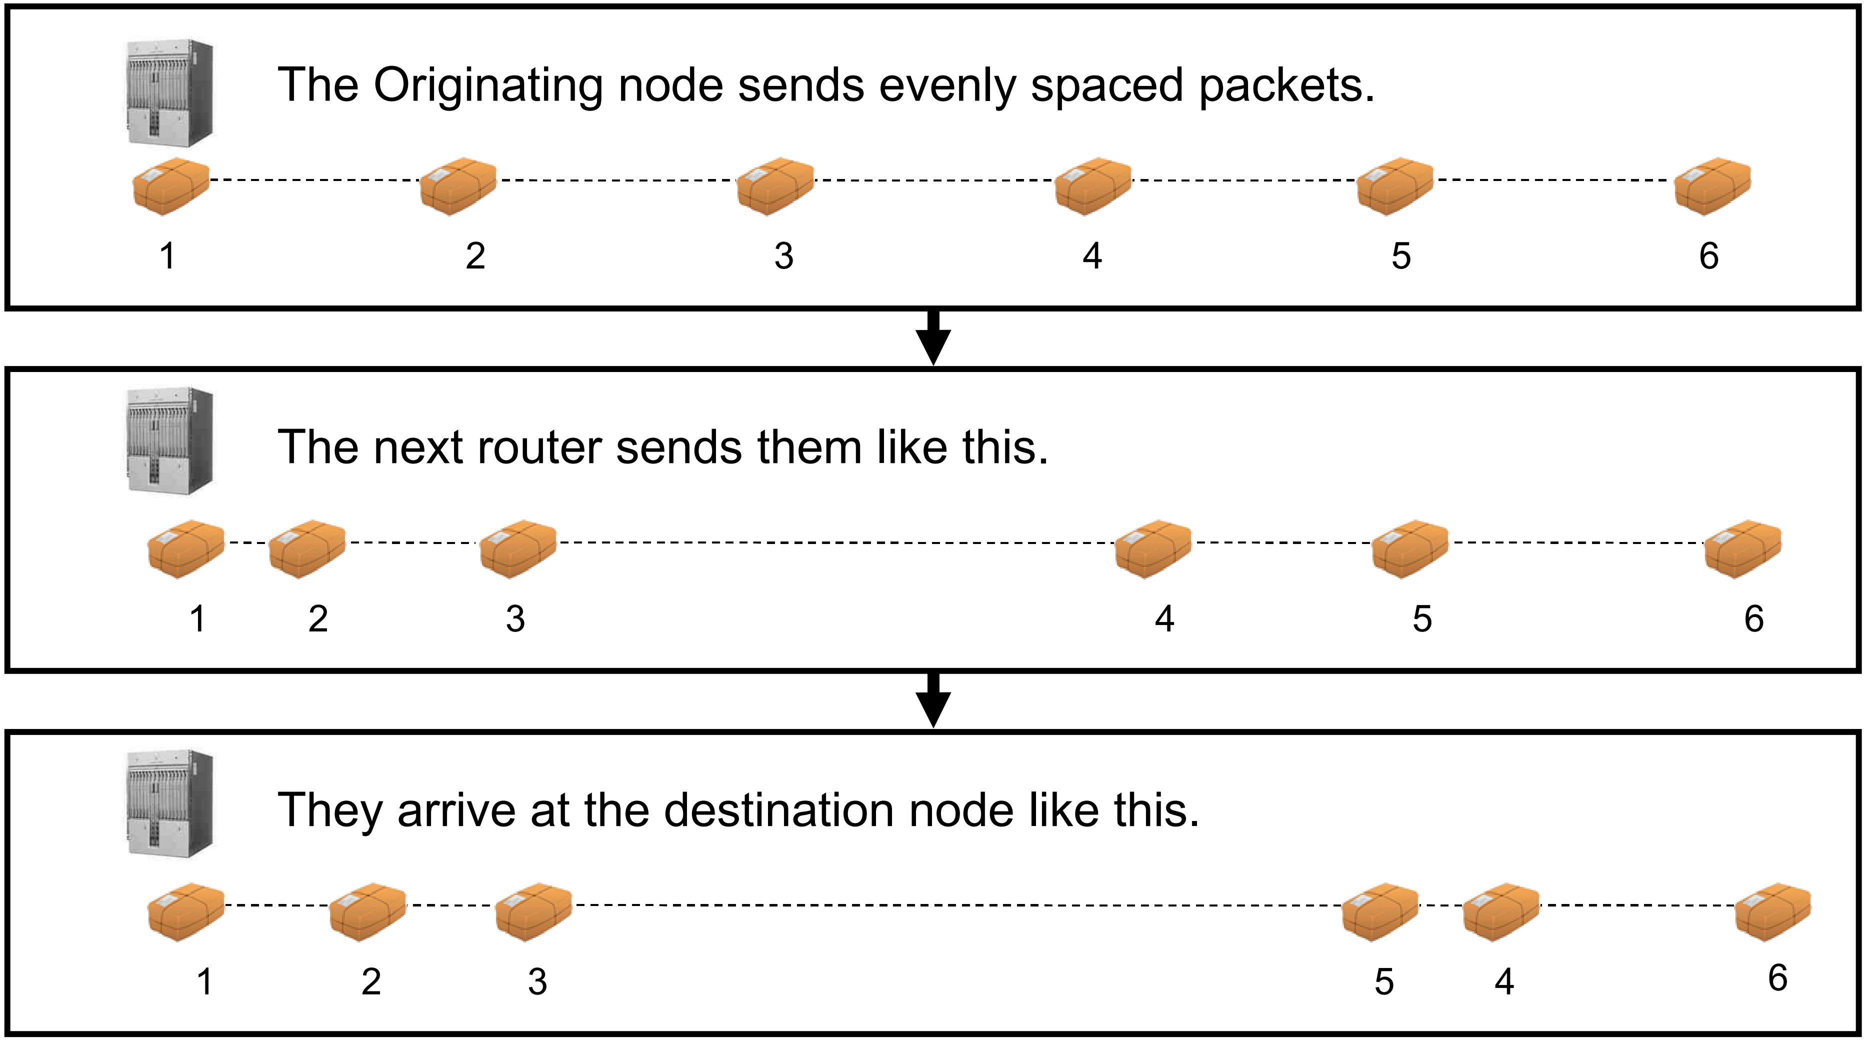

GÉANT2

DELAY VARIATION: DEFINITION AND IMPLICATIONS

Delay variation:

- Is the variation in travel times between source and destination (One Way Delay) of consecutively sent packets.
- Is closely related to 'jitter' (the deviation of packet arrival times from an assumed ideal regular arrival rhythm).
- Can be caused by:
  - Queuing (congestion).
  - Contention for routers' processing resources during forwarding.
- Can be quantified using IP Delay Variation Metric (IPDV).
  - Only compares delays for packets of equal size.
    - Serialisation naturally causes delay-variation for packets of unequal sizes.
- Real-time applications such as voice/video require *jitter buffers.*
  - Impacts overall delay (responsiveness); often not implemented well.

PACKET REORDERING (1)

TCP is designed to:

- Allow packet reordering.
- Automatically re-assemble the byte-stream in the original order at its destination.
  - Performance penalty when reordering is frequent (TCP “slow path”).

Packet Reordering is:

- Usually caused by parallelism.
- Prevalent where packet-sizes in a byte-stream are unequal.
  - Bulk transfers usually generate equal-sized packets.
  - Multi-media applications often generate unequal packet sizes.

PACKET REORDERING (2)

The probability of packet reordering can be decreased by:

- Avoiding parallelism in the network.
- Keeping the whole of a “flow” on a single path.
  - Use a hash on the destination address or the source / destination pair to select from the available paths.
    - Sometimes hard to achieve.

PACKET LOSS (1)

Packet loss: when a packet is lost 'in transit' between its source and destination.

Packet loss can be caused by:

- Congestion:
  - Traffic exceeds capacity in part of a network.
  - Packets are queued in buffers.
  - When a buffer's capacity is exceeded, the queue overflows and packets are dropped.
  - (Short-term) congestion may not be obvious from traffic graphs.

PACKET LOSS (2)

Packet loss is also caused by:

- Errors:
  - Packets can be corrupted (modified) in transit due to noisy lines.
  - Detected by link-layer checksum at destination.
  - Corrupt packets are discarded.
- Rate limits:
  - Does not necessarily correlate with queuing.

PACKET LOSS (3)

Impact on performance:

- TCP:
  - Detects packet-loss.
  - Assumes it is caused by congestion.
  - Reduces transmission rates accordingly.
- For bulk transfers:
  - Lost packets must be retransmitted – slows the transfer.
  - TCP interprets loss as signal of congestion and “backs off”.
- For real-time applications:
  - Re-transmission of packets useless because of timeliness requirements.
  - Effect is quality degradation (drop-outs, “pixelisation” etc.).

PACKET LOSS (4)

Packet loss can be reduced by:

- Careful provisioning of link capacities.
  - Buffers in network elements must be sufficient to cope with bursts.
  - Factors in determining buffer size:
    - Link capacity.
    - Expected RTT and degree of multiplexing.
  - Note that large buffers can increase one way delay (and therefore round trip time) and delay variation.

GEANT2

PACKET LOSS (5)

Packet loss can also be reduced by:

- Adoption of a quality of service mechanism such as DiffServ or IntServ.
  - Will protect a subset of traffic, but at the expense of increased packet loss in other traffic.
- Use of Active Queue Management (AQM) and Explicit Congestion Notification (ECN).

GEANT2

MAXIMUM TRASMISSION UNIT (1)

The protocol Maximum Transmission Unit (MTU) **of a link** is the greatest size of packet that can be transferred over the link without fragmentation.

Common MTUs include:

- 1480 bytes (PPPoE for ADSL environments).
- 1500 bytes (Ethernet, 802.11 WLAN).
- 4470 bytes (FDDI, common default for POS and serial links).
- 9000 bytes (Internet2 and GÉANT convention, limit of some Gigabit Ethernet adapters).
- 9180 bytes (ATM, SMDS).

GEANT2

MAXIMUM TRASMISSION UNIT (2)

MTU is a property of a "link" (= logical subnet):

- You cannot mix stations with different MTUs on a subnet!
- Else you will experience "MTU blackhole" in one direction.
- Easy to upgrade backbone (of point-to-point links) MTU.
- Harder to upgrade large LANs, Exchange Points...

Recommendation:

- Put large-MTU machines (high-performance servers/grid) on their own VLANs.

GEANT2

MAXIMUM TRANSMISSION UNIT (3)

Path MTU is equal to the lowest MTU of any of the links in a network path.

Larger path MTUs = quicker data transfers.

- Fewer packets have to be processed by source and destination hosts and routers.

Mechanisms such as Large Send Offload (LSO) and Interrupt Coalescence diminish influence of MTU on performance.