

Jumbo Frames*

Regular Ethernet Frame: Total Size on the Wire 1538 bytes

InterFrame Gap 12 bytes	Preamble/SFD 8 bytes	Ethernet Header 14 bytes	Payload 1500 bytes	Ethernet FCS 4 bytes
----------------------------	-------------------------	-----------------------------	-----------------------	-------------------------

Overhead = Total Size – Payload Size / Total Size = 1538 – 1500 / 1538 .025%

Efficiency = Payload Size / Total Size = 1500 / 1538 97.529%

Jumbo Ethernet Frame: Total Size on the Wire 9038 bytes

InterFrame Gap 12 bytes	Preamble/SFD 8 bytes	Ethernet Header 14 bytes	Payload 9000 bytes	Ethernet FCS 4 bytes
----------------------------	-------------------------	-----------------------------	-----------------------	-------------------------

Overhead = Total Size – Payload Size / Total Size = 9038 – 9000 / 9038 .004%

Efficiency = Payload Size / Total Size = 9000 / 9038 99.580%

In both cases, Ethernet inflicts 38 bytes of overhead 12 + 8 + 14 + 4 = 38

I'm not impressed – why do we bother with Jumbo Frames?

* Yes, I'm ignoring optional Ethernet fields, like VLAN, Provider Bridge, and MPLS tags.

Jumbo Frames

In many ways, I propose that Jumbo Frames falls into the Myth category – under *ideal* conditions, they deliver just ~2% greater throughput than do Regular Frames.

In the Beginning

- The performance case for Jumbo Frames is not on the wire; rather, we developed them as a way to conserve host CPU: historically, NICs generated a CPU interrupt for every frame, and CPU resources were scarce. Jumbo frames, under ideal conditions, can shrink the host I/O interrupt load, conserving precious CPU, by a factor of ~6: $9000 / 1500 \approx 6$, i.e. roughly 6 Regular Frames can fit into a single Jumbo Frame.

However

- Modern servers tend to be awash in multi-core CPUs – CPU is no longer a scarce resource.
- Modern NICs/Operating Systems employ Large-Segment Offload and Interrupt Coalescing to further mitigate this issue.
- Configuring Servers, Switches, and Storage to support the same Jumbo Frame size defeats IT organizations regularly -- a significant cause of service disruption: Be Afraid.
- Network stack developers have spent decades optimizing performance for 1500 byte frames ... bringing that optimization to 9000 byte frames takes years of work – depending on your hardware/software, you may find that Jumbo Frames *degrade* performance.

Your Mileage May Vary

- Then again, your particular hosts running their particular Ethernet/IP/TCP stacks may benefit substantially from Jumbo Frames: *Try it and Measure the Results*

Jumbo Frames

Specific Cases

- Even on hosts where Jumbo Frames have only a small effect, long-running, high-volume work-flows may still benefit from Jumbo Frames enough to be worth your effort. Shaving a few percent off the run-time of a month-long job may provide your firm with a competitive advantage: the render farms producing special effects for movies are a commonly touted example.

Your Mileage May Vary

http://en.wikipedia.org/wiki/Jumbo_frame

<http://www.ethernetalliance.org/wp-content/uploads/2011/10/EA-Ethernet-Jumbo-Frames-v0-1.pdf>

<http://www.boche.net/blog/index.php/2011/01/24/jumbo-frames-comparison-testing-with-ip-storage-and-vmotion/>

<http://codingrelic.geekhold.com/2011/12/requiem-for-jumbo-frames.html>