

# FAQ with Answers about the 500-mile email

Trey Harris  
<trey at lopsa dot org>

I've been receiving a lot of responses to the [500-mile email story](#). Since I originally posted it, it has been forwarded massively, far beyond anything I anticipated. While most of the responses are just to say, "thanks, fun story" or the like, and some have been solicitations for work (thanks, and keep them coming!), a not-insubstantial portion have been of, might I say, the nit-picking variety. Rather than reiterate the points again and again, I've compiled these answers to frequently-asked questions.

## 1. Did this actually happen, or were you just spinning a yarn?

Yes, it happened. At the time, I was running the centralized campus email system, Isis, at the University of North Carolina at Chapel Hill. I was informally responsible for some aspects of the email systems in departments who chose to run their own independent systems. Most notably for the purposes of this story, I wrote `sendmail.cf` (Sendmail configuration) files that were used in most of the mail servers on the campus.

## 2. When was this?

I wish I could say for certain. I'm one of those anal-retentive types who save *all* my back email, sent and received, but some simple searches haven't produced any mail on the subject--not that I recall engaging in any. The initial contact, as I wrote in the post, was by telephone, as was the resolution. After that time, I *chose* not to write email about it, basically because it's a good story and I liked seeing the faces of people who had never heard it when I told it to them. My guess, from the office I remember being in, the coworkers I remember speaking about this to, and some other such irrelevant but timely details, place it somewhere between 1994 and 1997.

That said, it *should* be possible to reconstruct an approximate time. When was Sendmail V8 "fairly mature", but Sun still shipping Sendmail V5? Eric Allman notes in a personal correspondence that the Sendmail V5 in question might have had backported features from Sendmail V8, so that might narrow it down. If you have any data that could assist with nailing down a timeframe, I'd appreciate it if you passed it along.

## 3. Was it actually you, or was your identity one of the "irrelevant details" that you altered?

It was really me. I *think* that I may have been alerted to "some situation in statistics" by one of my coworkers prior to actually speaking to the chairman. I may have even called the chairman rather than his calling me. And it is possible--even likely--that one of my coworkers was sitting with me during some of the resolution, as I frequently talk out problems while I work them. But I don't believe I'm taking credit for anyone else's work. (If you're one of those former coworkers and disagree, *please* let me know and I will rectify the situation as best I can.)

## 4. If you're not 100% certain of all the details, why did you write the original post so vividly?

I took license. It made a better story that way. Honestly, would writing "I'm not sure, but maybe..." every second sentence really have changed anything? And I *did* say at the start that I was changing or eliding irrelevant details for the sake of the story.

Another factor is the context in which this originally appeared. It was posted to the sage-members list, a list for members of SAGE, the System Administrators Guild, in a thread about "favorite impossible tasks". That is to say, it was a light-hearted discussion about the impossible tasks that users or management sometimes bring to sysadmins to solve.

If I had any notion that this post would be forwarded so widely, I would have added some details to answer the questions of skeptics. But I was posting to a list of colleagues, a good proportion of whom I know personally, and so was tailoring it for an audience inclined to believe me. :-)

## 5. The story is fun, but the technical details at the end just don't add up.

I know. See the answer to the previous question. First and foremost, I was writing a humorous story based on an actual occurrence, not a case study, so I wrote only what was required to get the point across. I find I suddenly have great respect for writers who write works based on true stories, as I now know how hard it is to find the balance between verisimilitude and storytelling. And I now know the kind of criticism you open yourself up to when you choose storytelling over verisimilitude. :-)

**6. Well, why don't you write a case study, with all the details?**

Unfortunately, I couldn't write a case study today if I wanted to, because I now lack the raw data. I didn't keep logs of what I found, and I no longer have the notes I took at the time. I really, *really* wish I had, as I now see I could have published a paper on it. At the time it seemed trivial, other than as a fun story to tell folks to get a laugh. It still works for that, even without the notes.

That said, there *are* details that I didn't put into the story that I do remember, or can reconstruct. That is what most of the next questions have to do with.

**7. That three millisecond time doesn't make sense as the timeout for a connect() call.**

Yes, I know. And it wasn't the timeout, actually. In the story, I make it sound like it took all of ten minutes from being made aware of the 500-mile email limit and determining a 3 ms light-speed issue. In fact, this took several hours, and quite a bit of detective work. The point is, eventually I came up with that figure, ran units, and gagged on my latte. (I'm fairly certain it was a different latte from the one I started with.) So what, in particular, is your question about the 3 ms figure?

**8. Well, to start with, it can't be three milliseconds, because that would only be for the outgoing packet to arrive at its destination. You have to get a response, too, before the timeout will be aborted. Shouldn't it be six milliseconds?**

Of course. This is one of the details I skipped in the story. It seemed irrelevant, and boring, so I left it out.

**9. Actually, shouldn't it be twelve/eighteen/twenty-four milliseconds, to account for the three-way TCP handshake?**

Maybe. Again, this would be a detail my lost notes would answer. But I think that a connect() timeout would be aborted upon receiving a SYN/ACK packet; I don't think that the whole handshake had to be completed. Even if it did, I eventually would have arrived at the 3 ms figure, however I got to it.

**10. Router delays would have been a much greater factor than you admit in the story.**

Yes, you're probably right. But they are factors I could account for. I'm not certain this is how I did it, but it seems likely I could have pinged the nearest router I could find (such as one at another school at UNC that manged its own network) to find out what sort of delay a router was likely to add. Then I could multiply that by the number of routers to remote destinations. The number was likely to be constant for other East Coast universities. And even if it wasn't, the delay imposed by an additional router would only be on the order of a few hundred microseconds at most, not enough to make a large difference for nearby destinations.

**11. The story is cute, but it has a fatal flaw: signals don't travel at lightspeed in copper.**

That's true, they travel at  $3c/4$  or thereabouts. But the NIC, the campus backbone, and certainly the Internet backbone was all fiber.

**12. Ah-hah! But signals don't travel at light speed in fiber, either!**

You got me. I'm told they travel at from  $2c/3$  (yes, *slower* than copper) up to a few percent under  $c$  depending on a wide variety of factors. But again, this was a factor I could, and did, account for. I recall pinging various destinations and writing down distances versus ping times, and coming up with an empirical "effective time" that differed from actual time. This was just another "irrelevant and boring detail" to be left out of the story.

**13. Wait, doesn't that mean that you knew what was going on--i.e. that speed-of-light had something to do with the problem--before you typed that figure into units?**

Yep, I did. But I was stubborn. Have you ever averted your eyes from the answer to a riddle while you worked it out for yourself? That was what I did. I could have typed "500 miles" into units and worked backwards, but once I knew what going on, it was an intellectual challenge to be figured out.

**14. So you knew how to solve the problem the department was having sending email, but you left it alone until you figured out the timeout issue?**

No--as soon as I knew that replacing the SunOS sendmail binary with the Sendmail V8 binary would fix the problem, I did so. (Even if I didn't know this would fix it, I would have done so, as running Sendmail V5 with a V8 cf file is definitely not a recommended configuration.) But I kept the other binary around so that I could continue experimenting with it at my leisure.

Sysadmins do that sort of thing all the time. The answer to a bug is never "the system has been up too long," but rebooting the system is still sometimes the right answer, simply because it has the *effect* of fixing the problem. You fix a problem, any way you can, in order to rectify a production issue, and then you come back later to try to determine what the "right" solution is.

**15. The Internet travels all sorts of circuitious routes. But the 500-mile email story depends on the packets travelling a direct path between source and destination, doesn't it?**

No, it doesn't. As I wrote in the story, the 500-miles-plus-a-bit was a *maximum* radius outside of which it was *impossible* to send email. Within that radius, there were plenty of sites to which sending mail was also impossible, or at least sporadic.

There's at least two possible reasons for this. One is that some other delay--from a firewall, for instance--was long enough to hit the connect timeout. Another is that the sites within that radius that mail couldn't reach were accessed via a network path that deviated enough from great-circle to make it longer than the maximum radius.

UNC was a very well-connected East Coast university, so the path to other East Coast universities (which were where the majority of successful mails were sent to) would closely resemble a great circle, especially at the time in which the story was set, when it was rather rare to see a packet going to San Jose in order to get from Atlanta to Washington.

**16. So why did you feel the need to include the detail about the campus being 100% switched?**

You know, I'm not actually sure. At the time, it seemed the most glaring problem with the plausibility of the story, so I put it in. I'm not sure why in retrospect. Feel free to delete that paragraph mentally upon second reading.

Hacksaw writes: "Switched means no other delays, like losing the collision detection game. It makes it somewhat easier to detect things like this, since there is less noise in the data. I bet that's what you were thinking."

**17. Sendmail V5 wouldn't accept a Sendmail V8 cf file.**

But it did. I'm told that the Sendmail V5 on the net wouldn't have. So therefore, I'm forced to conclude that the one Sun shipped with Solaris had tweaks to enable it to do so. If you have access to that source, I'd love independent verification. But it happened, therefore it's possible. :-)

**18. Sendmail has defaults compiled into the binary; it wouldn't simply assign missing options a value of zero.**

Several people have written me alleging this. It may be true today, but it was definitely not true with that sendmail at that time. I am certain of this, because a year or two after this happened, I was in a Sendmail tutorial at LISA with Eric Allman. He mentioned that sendmail doesn't have defaults for the options he was describing. (The standard sendmail.cf files did have defaults, but that is irrelevant in this case.) I took that opportunity to tell him the story of the 500-mile email. He literally fell to the floor in pain. :-)

**19. units on SunOS doesn't know about "millilightseconds."**

Yes. So? I used to populate my units.dat file with tons of extra prefixes and units. And actually, I think I was using AIX to run units; I don't know if it knew about millilightseconds. Take a look at the units.dat shipped with Linux these days. It definitely knows about millilightseconds.

**20. These "lost notes" you refer to sound awfully convenient.**

Yes, well, how many pieces of scratch paper from five years ago do *you* still have?

**21. Well, the story still can't be true.**

Let me ask *you* a question: regardless of the details, is it possible that a misconfiguration could cause the operational behavior of nearby email being delivered while faraway email was not? I think the answer is yes. In fact, I *know* the answer is yes, because it happened. But even putting aside my own experience and viewing it as best I can as a skeptical observer, I think the idea is possible, though certainly implausible at first gloss.

If you have a question that isn't answered here, go ahead and email me at [trey+500mi@lopsa.org](mailto:trey+500mi@lopsa.org). I may put it in the FAQ and credit you. But I'm more likely to just say, "I don't know, I don't remember and no longer have the raw data to answer your question."

**22. The signature says you're looking for work. Are you still?**

Nope, but thanks for asking!