

< INC-PROJECT, MAP-ILLUSION.NLS.8, >, 12-Aug-83 11:44 AMW ;;;

RFC 873

September 1982
M82-49

THE ILLUSION OF VENDOR SUPPORT

M.A. PADLIPSKY
THE MITRE CORPORATION
Bedford, Massachusetts

ABSTRACT

The sometimes-held position that "vendor supplied" intercomputer networking protocols based upon the International Standards Organization's Reference Model for Open System Interconnection are worth waiting for, in particular in preference to protocols based upon the ARPANET Reference Model (ARM), is shown to be fallacious.

The paper is a companion piece to M82-47, M82-48, M82-50, and M82-51.

THE ILLUSION OF VENDOR SUPPORT

M. A. Padlipsky

Introduction

Even one or two members of the DoD Protocol Standards Technical Panel join with many others (including, apparently, some members of the DoD Protocol Standards Steering Group, and clearly, somebody at the GAO) in expressing a desire to "go with vendor-supported intercomputer networking protocols instead of using our own." The author's view of the implications of this desire should be clear from the title of this paper. What evidence, then, is there to so stigmatize what is clearly a well-meant desire to save the Government money?

Scope

First, we must consider what is meant by "vendor-supported protocols." It can't be just X.25, because that only gets you through the network layer whether you're appealing to the International Standards Organization's widely-publicized Reference Model for Open System Interconnection (ISORM) or to the unfortunately rather tacit reference model (ARM) to which the ARPANET protocols (e.g., TCP, IP, Telnet, FTP) were designed. It also can't be just X.25 and X.28/X.29 (even with X.75 tossed in to handle "internetting" and X.121 for addressing) because: 1. They don't serve as a protocol suite for resource sharing (also known as OSI), but rather only allow for remote access [1]. 2. They (coming as they do from the Consultative Committee on International Telegraphy and Telephony--and including one or two other protocols, in reality) don't even constitute the full protocol suite being worked on by the U. S. National Bureau of Standards, much less the somewhat different suite being evolved by ISO. So it must be a suite from NBS or ISO, and for present purposes we needn't differentiate between them as their Reference Models are close enough to be shorthanded as the ISORM.

Timeliness

Realizing that we're being asked to consider an ISORM-related protocol suite as what the vendors are expected to support has one immediate consequence which in some sense can be considered to dominate all of the other points to be raised: That is, the DoD procurement process entails quite long lead times. Yet the ISORM suite is by no means complete at present. Without prejudice to its

merits or demerits, only X.25 (as levels 1-3, and with some ambiguity as to what level X.75 belongs at) is as yet firmly in the ISORM suite (which it will be convenient to refer to as "ISORMS"), and there is even some doubt as to how firmly they're there. (E.g., a British observer at a recent PSTP meeting assured the author that "We in the U.K. don't believe X.25 is officially part of the ISORM.") There are proposals which have been circulating for some time at Level 4, and less far along through the international (or even national, remembering NBS) standardization process, ones at Level(s) 5-7. It must be noted that: 1. These are by and large "paper protocols" (that is, they have not been subjected to the test of actual use). 2. Even ISO and NBS's warmest supporters acknowledge that the standardization process "takes years." So if the DoD is to avoid buying what might turn out to be a series of pigs in a series of pokes, it can't wait for the ISORMS.

On the other side of the coin, the DoD is letting intercomputer networking contracts right now. And, right now, there does exist a suite of protocols designed to the ARPANET Reference Model (ARMS, with no pun intended). Implementations of the ARMS already exist for a number of operating systems already in use in the DoD. Now, it is not argued that the ARMS protocols come "for free" in upcoming acquisitions (contractors fuss about the style of the available specifications, system maintainers fear incursions of non-vendor supplied code into operating systems, and so on), but it is unarguable that the ARMS can be procured significantly more rapidly than the ISORMS. (It is also unarguable that those who speak of their unwillingness to see the DoD "develop new protocols rather than employ international standards" haven't done their homework; we're not talking about new protocols in the ARMS, we're talking about protocols that have been in real use for years.)

Quality of Support

The timeliness argument can lead to a counterargument that the ISORMS is "worth waiting for," though, so we're not done yet. Let's look further at what "vendor support" means. Clearly, the proponents of the position expect that vendors' implementations of protocols will be in conformance with the Standards for those protocols. Given the nature of these specifications, though, what can we infer about the quality of support we can expect from the vendors?

There are two problem areas immediately apparent: ambiguities and options. Let's take ambiguities first. The following are some of the questions raised by knowledgeable observers about the present state of the ISORMS:

1. Can an X.25 comm subnet offer alternate routing? (The answer depends on whether "DCE's" are expected to follow X.25 between themselves. The situation is further complicated by the fact that some ISORM advocates don't even include the Data Communication Elements in their depictions of the Model; this leads to the metaphorical question* "Are there parking garages between the highrises?") If you can conform to X.25 and not offer alternate routing--which certainly appears to be consistent with the spec, and might even be construed as required by it--the DoD's inherent interest in "survivability" cannot be served by you.
2. Can an X.75 internet offer alternate gatewaying? (The answer is almost surely no, unless the X.75 spec is re-written.) If not, again the DoD's interest is not served.
3. Does "Expedited Data" have semantics with regard to the L4-L5/L7 interface? (Not as I read the spec, by the way.) If not, the ISORMS lacks the ability to convey an "Out-of-Band-Signal" to an Application protocol. (This leads to the metaphorical question, "What good is an SST if there's nobody on duty at the Customs Shed?")
4. Must all layers be traversed on each transmission? (There are rumors of a new ISORM "null-layer" concept; it's not in the last version I looked at, however, and apparently the answer is yes at present.) If so, the DoD's inherent interest in efficiency/timeliness cannot be served. (This leads to the metaphorical question, "Are there elevators inside the highrises, or just staircases?")
5. Can an implementation be in conformance with the ISORM and yet flout the prescription that "N-entities must communicate with each other by means of N-1 entities"? (Not as I read the spec.) If not, again implementations must be inefficient, because the prescription represents an inappropriate legislation of implementation detail which can only lead to inefficient implementations.

* This and other metaphorical questions are dealt with at greater length in reference [2].

6. Is each layer one protocol or many? (The point quoted in 5 would seem to imply the latter, but many ISORM advocates claim it's the former except for L1 and L7.) If each layer is a "monolith", the DoD's interest is not served because there are many circumstances in which applications of interest require different L1-3 and L4 protocols in particular, and almost surely different L5 and L6 protocols. (Areas of concern: Packetized Speech, Packet Radio, etc.)

The upshot of these ambiguities (and we haven't exhausted the subject) is that different vendors could easily offer ISORMS's in good faith which didn't interoperate "off-the-shelf". Granted, they could almost certainly be fixed, but not cheaply. (It is also interesting to note that a recent ANSI X3T5 meeting decided to vote against acceptance of the ISORM as a standard--while endorsing it as valuable descriptively--because of that standards committee's realization of just the point we are making here: that requiring contractual compliance with a Reference Model can only be desirable if the Reference Model were articulated with utter--and probably humanly unattainable--precision.)

The area of options is also a source for concern over future interoperability of ISORMS implementations from different vendors. There's no need to go into detail because the broad concern borders on the obvious: What happens when Vendor A's implementations rely on the presence of an optional feature that Vendor B's implementations don't choose to supply? Somebody winds up paying--and it's unlikely to be either Vendor.

On the other side of the coin, the ARMS designers were all colleagues who met together frequently to resolve ambiguities and refine optionality in common. Not that the ARMS protocols are held to be flawless, but they're much further along than the ISORMS.

To conclude this section, then, there are grounds to suspect that the quality of vendor support will be low unless the price of vendor support is high.

Nature of the Design Process

The advantage of having colleagues design protocols touched on above leads to another area which gives rise to concern over how valuable vendor-supported protocols really are. Let's consider how international standards are arrived at:

The first problem has to do with just who participates in the international standardization process. The author has occasionally chided two different acquaintances from NBS that they should do something about setting standards for membership on standards committees. The uniform response is to the effect that "They are, after all, voluntary standard organizations, and we take what we're given." Just how much significance is properly attached to this insight is problematical. Even the line of argument that runs, "How can you expect those institutions which have votes to send their best technical people to a standards committee? Those are precisely the people they want to keep at home, working away," while enticing, does not, after all, guarantee that standards committees will attract only less-competent technicians. There are even a few Old Network Boys from the ARPANET involved with the ISORM, and at least one at NBS. However, when it is realized that the rule that only active implementers of TCP were allowed on the design team even precluded the present author's attendance (one of the oldest of the Old Network Boys, and the coiner of the phrase, at that), it should be clear that the ARMS enjoys an almost automatic advantage when it comes to technical quality over the ISORMS, without even appealing to the acknowledged-by-most politicization of the international standards arena.

What, though, of the NBS's independent effort? They have access to the experienced designers who evolved the ARMS, don't they? One would think so, but in actual practice the NBS's perception of the political necessities of their situation led one of their representatives at a PSTP (the Department of Defense Protocol Standards Technical Panel) meeting to reply to a reminder that one of the features of their proposed Transport Protocol was a recapitulation of an early ARPANET Horror Story and would consume inordinate amounts of CPU time on participating Hosts only with a statement that "the NBS Transport Protocol has to be acceptable as ECMA [the European Computer Manufacturer's Association] Class 4." And even though NBS went to one of the traditional ARPANET-related firms for most of their protocol proposals, curiously enough in all the Features Analyses the author has seen the features attributed to protocols in the ARMS are almost as likely to be misstated as not.

The conclusion we should draw from all this is not that there's something wrong with the air in Gaithersburg, but rather that there's something bracing in the air that is exhaled by technical people whose different "home systems'" idiosyncracies lead naturally to an intellectual cross-fertilization, on the one hand, and a tacit agreement that "doing it right" takes precedence over "doing it expediently," on the other hand. (If that sounds too corny, the reader should be aware that the author attended a large number of

ARPANET protocol design meetings even if he wasn't eligible for TCP: in order to clarify our Host-parochial biases, we screamed at each other a lot, but we got the job done.)

One other aspect of the international standardization process has noteworthy unfortunate implications for the resultant designs: However one might feel on a technical level about the presence of at least seven layers (some seem to be undergoing mitosis and growing "sublayers"), this leads to a real problem at the organizational--psychological level. For each layer gets its own committee, and each committee is vulnerable to Parkinson's Law, and each layer is in danger of becoming an expansionist fiefdom If your protocol designers are, on the other hand, mainly working system programmers when they're at home--as they tend to be in the ARPANET--they are far less inclined to make their layers their careers. And if experience is weighted heavily--as it usually was in the ARPANET--the same designers tend to be involved with all or most of the protocols in your suite. This not only militates against empire building, it also minimizes misunderstandings over the interfaces between protocols.

"Space-Time" Considerations

At the risk of beating a downed horse, there's one other problem area with the belief that "Vendor supplied protocols will be worth waiting for" which really must be touched on. Let's examine the likely motives of the Vendors with respect to "space-time" considerations. That is, the system programmer designers of the ARMS were highly motivated to keep protocol implementations small and efficient in order to conserve the very resources they were trying to make sharable: the Hosts' CPU cycles and memory locations. Are Vendors similarly motivated?

For some, the reminder that "IBM isn't in business to sell computers, it's in business to sell computer time" (and you can replace the company name with just about any one you want) should suffice. Especially when you realize that it was the traditional answer to the neophyte programmer's query as to how come there were firms making good livings selling Sort-Merge utilities for System X when one came with the operating system (X = 7094 and the Operating system was IBSYS, to date the author). But that's all somewhat "cynical", even if it's accurate. Is there any evidence in today's world?

Well, by their fruits shall you know them: 1. The feature of the NBS Transport Protocol alluded to earlier was an every 15-second "probe" of an open connection ("to be sure the other guy's still

there"). In the early days of the ARPANET, one Host elected to have its Host-Host protocol (popularly miscalled "NCP" but more accurately AH-HP, for ARPANET Host-Host Protocol) send an echo ("ECO") command to each other Host each minute. The "Network Daemon" on Multics (the process which fielded AH-HP commands) found its bill tripled as a result. The ECMA-desired protocol would generate four nuisance commands each minute--from every Host you're talking to! (The "M", recall, is for Manufacturers.)* 2. X.25 is meant to be a network interface. Even with all the ambiguities of the ISORM, one would think the "peer" of a "DTE" (Host) X.25 module (or "entity") would be a "DCE" (comm subnet processor) X.25 module. But you can also "talk to" at least the foreign DCE's X.25 and (one believes) even the foreign DTE's; indeed, it's hard to avoid it. Why all these apparently extraneous transmissions? CCITT is a body consisting of the representatives of "the PTT's"--European for State-owned communications monopolies. 3. The ISORM legislates that "N-entities" must communicate through "N-1 entities." Doesn't that make for the needless multiplication of N-1 entities? Won't that require processing more state information than a closed (or even an open) subroutine call within level N? Doesn't anybody there care about Host CPU cycles and memory consumption?

Note particularly well that there is no need to attribute base motives to the designers of the ISORMS. Whether they're doing all that sort of thing on purpose or not doesn't matter. What does matter is that their environment doesn't offer positive incentives to design efficient protocols, even if it doesn't offer positive disincentives. (And just to anticipate a likely cheap shot, TCP checksums are necessary to satisfy the design goal of reliability; ECMA four pings a minute is[/was] unconscionable.)

TANSTAAFL

We're very near the end of our analysis. Readers familiar with the above acronym might be tempted to stop now, though there are a few good points to come. For the benefit of those who are not aware: "There Ain't No Such Thing As A Free Lunch." Achieving interoperability among vendor-supplied protocol interpreters won't come for free. For that matter, what with all this "unbundling"

* Rumor has it that the probes have since been withdrawn from the spec. Bravo. However, that they were ever in the spec is still extremely disquieting--and how long it took to get them out does not engender confidence that the ISORMS will be "tight" in the next few years.

stuff, who says even the incompatible ones come for free? You might make up those costs by not having to pay your maintenance programmers to reinsert the ARMS into each new release of the operating system from the vendor, but not only don't good operating systems change all that often, but also you'll be paying out microseconds and memory cells at rates that can easily add up to ordering the next member up in the family. In short, even if the lunch is free, the bread will be stale and the cheese will be moldy, more likely than not. It's also the case that as operating systems have come to evolve, the "networking" code has less and less need to be inserted into the hardcore supervisor or equivalent. That is, the necessary interprocess communication and process creation primitives tend to come with the system now, and device drivers/managers of the user's own devising can often be added as options rather than having to be built in, so the odds are good that it won't be at all hard to keep up with new releases anyway. Furthermore, it turns out that more and more vendors are supplying (or in process of becoming able to supply) TCP/IP anyway, so the whole issue of waiting for vendor support might well soon become moot.

References

- [1] Padlipsky, M. A., "The Elements of Networking Style", M81-41, The MITRE Corporation, October 1981, attempts to clarify the distinction between "remote access" and "resource sharing" as networking styles.
- [2] -----, "A Perspective on the ARPANET Reference Model", M82-47, the MITRE Corporation, September 1982; also available in Proc. INFOCOM '83.