Rand
SANTA MONICA, CA. 90406

November 1, 1983

[redacted] STAT

Office of Data Processing
Central Intelligence Agency
Washington, DC 20505

Dear Bob:

Below are a number of points that caught my ear as we worked through the PDR. Some of them may not prove important simply because there is additional information that I do not have, but others may affect what is going to happen between now and CDR.

First of all, we all know that the schedule is impossibly tight. The Delivery-2 is really the big step when the VM and MVS environments get networked for the first time. It is also the first occasion on which there is a full-up SUL implementation to the user, and also when most of the software packages play together for the first time. There are other technical innovations as well; ones mentioned included distributed restart/recovery, selective operation when failures occur, and selective recovery.

- Has enough time been allowed to shake out this first delivery?

Delivery-2 is the foundation on which all else hinges, so anything that can be done to give you additional schedule time for all the needed testing and integration will certainly be helpful.

I find myself concerned about the whole issue of recovery and restart. The SAFE configuration has about five large mainframes, and there are five or so large software packages. If one supposes that the mean free-time-between-failure for each is 100 hours, then the system will have a failure every 10 hours on the average. In a scheduled 22-hour day there will be two or three failures, which means that the total time for diagnosing trouble and recovering from it must be of the order of 30 to 45 minutes. While we heard much talk during PDR of restart features in this, that, or the other software package, my intuition says that the whole issue needs a thorough examination to make sure that the restart features of all the parts mesh smoothly.

I would argue that the issue of fault-diagnosis and restart/recovery procedures are of such importance that a special effort to examine them with perhaps a special contractor is warranted.

An advantage of using a separate contractor would be that he could do a thorough fault analysis of the proposed system and use the existing running one as a source of real data. He could also canvass other operators of large IBM configurations for pertinent information. On the basis of what he could learn, he could assist CSPO in analyzing and designing fault-diagnosis and restart/recovery procedures. There is an important currency, I think, about this issue because it might very well influence the design of the system, particularly with respect to building in redundancy, analytic checks of various kinds, and what IBM calls fixer modules in its operating system. The goal for your reliability/availability contractor (or group) has to be graceful degradation with gracious recovery.

I know that IBM hardware and software have certain error-accommodation features, but do the other software packages that you are acquiring have adequate ones? I find myself wondering whether it might not be wise to design rather extensive diagnostic tests and fault-isolation aids in the several large software packages. If the operational concept for SAFE assumes that everything will report its failure state to the console operator, and if he is expected to go through a restart/recovery procedure all by himself, then I feel very queasy about holding a 30- to 45-minute recovery period.

Since you are serving such a large population of users, it is simply unthinkable that the system be allowed to fail in such a way that 100 or 200 users have no awareness of their prior status and context of operations. This suggests that one might have to design special user-awarness files that allow users to get back on rapidly without going through all of the usual log-on, password, authentication, and menu actions.

As you know, the IBM operating systems incorporate special features to check on things in progress, and these invoke special fixer modules whenever trouble is detected. One wonders whether the big software packages that you are acquiring from commercial sources have such features and if so, how they will integrate into corresponding IBM ones. If they do not have such features, one wonders whether checker-fixer features ought not be added to each application program.

Hopefully, a UPS is being provided for the SAFE system, but if not, has consideration been given to the IBM power monitor that puts the system to sleep during misbehavior of the power source?

As I understood the discussion of the MEC analyst, he can decide on-the-fly to change the SLP for subsequent processing of messages. It strikes me that this might be a bit risky in that he could fix one thing but mess up many others. Perhaps this feature already exists in SEC and if so, you will have the experience to answer my concern. If it does not

exist, I wonder whether it might not be wise to mock up the feature and do some trials. Would it be worthwile to provide the MEC analyst with system-provided checks and prompts to help assure that he can not wreak damage? Might it even be so important that the process be put under two-person control?

During the PDR there was frequent reference to duplicate files for back-up purposes. At Sysgen time, one will have to be careful to make sure that duplicate files are on different spindles which are in turn on different controllers; otherwise the full purpose for backup will not be served. I do not know whether this detail will influence any other aspects of the design, but I surface it for your consideration.

I am not quite sure whether to raise the security issue in the same way as I did the restart/recovery one above. The security features, to be sure, are not distributed as widely throughout the system, and it may be that everything is well. On the other hand, it might be useful for some one person to have a comprehensive system-wide look at all security controls to make sure that nothing has dropped in a crack.

Somewhere along the way a comment was made that MAP processing strips out all control characters except end-of-line ones. It occurred to me that the line length in DATEX traffic might be unsuitable for the preferred line length on SAFE terminals. System designers make different choices with regard to this problem. Some of them maintain text as a running character stream which is displayed by a smart text formatter, whereas others retain the end-of-line characters on traffic as it comes. This is a detail but it could have extensive design ramifications; it might be worth a look.

Someone said, in effect, that "DATEX will restart and retransmit the message [in process] when a failure occurs." I am not familiar with the DATEX system, but the remark implies that the DATEX node which is feeding SAFE will be under the latter's control. I heard no discussion of this interface; is it a detail in the crack?

I am not familiar with the SANS algorithm, but I found myself wondering whether it clearly does produce a unique identification number. Within a given source probably the source plus the date-time-group is unique, but what happens if messages are picked up and retransmitted by a third party? Is there a problem here?

Another detailed point: Does Logicon provide an initial set of patterns/tables/strings to handle messages? Or does it only provide a capability for someone else to exploit the generality of Merlin? Has this detail been dealt with in the contractual relationship with Logicon?

I think I found all of the points in my notes that I wanted to flag for you; but should others arise, I will get back to you with a second letter.

Sincerely,

UNCODED

[redacted] STAT

Corporate Research Staff

WHW:dms

cc: [redacted] STAT

Mr. Harry Fitzwater
Deputy Director/Administration
Central Intelligence Agency
Washington, DC 20505

Rand
1700 MAIN STREET
SANTA MONICA, CA.90406
W. H. Ware