11 November 1975

FILE

MEMORANDUM FOR THE RECORD

SUBJECT : Project ORACLE

On 29 October 1975, [redacted] of the STATINTL Agency met with the [redacted] and STATINTL STATINTL [redacted] for the purpose of determining the STATINTL direction and status of [redacted] design of the ORACLE Mass Storage System. There were four areas discussed in detail that we felt were not being properly designed. We asked STATINTL [redacted] to tell us clearly what they were doing. The remainder of this paper is devoted to our understanding of the designs, there implications, commentaries and references that explain our concern, and a statement of our requirements for each design area.

I. Storage Control Processors accessing Transport Drivers

Hardware Design

STATINTL [redacted] stated that the current [redacted] design permits STATINTL one and only one Storage Control Processor (SCP) to be enabled at any given time to access any of the Transport Drivers (TD). The SCP's and TD's are cabled together in such a way that when one SCP is in control of the TD's, the other SCP is blocked from all access to any of the TD's by a hardware interlock. In order for the second SCP to gain access to the TD's, the first SCP must relinquish its control and then it is possible for the second one to take over. Once the second SCP is in control the first one is inhibited from all access to the TD's. Any interrupt issued by a TD is serviced by whichever SCP is currently in control. The TD does not know which SCP should handle the interrupt, it is expected in this design that this type of problem can be adequately solved at the SCP level. The drawing below illustrates how these hardware elements are tied together logically.

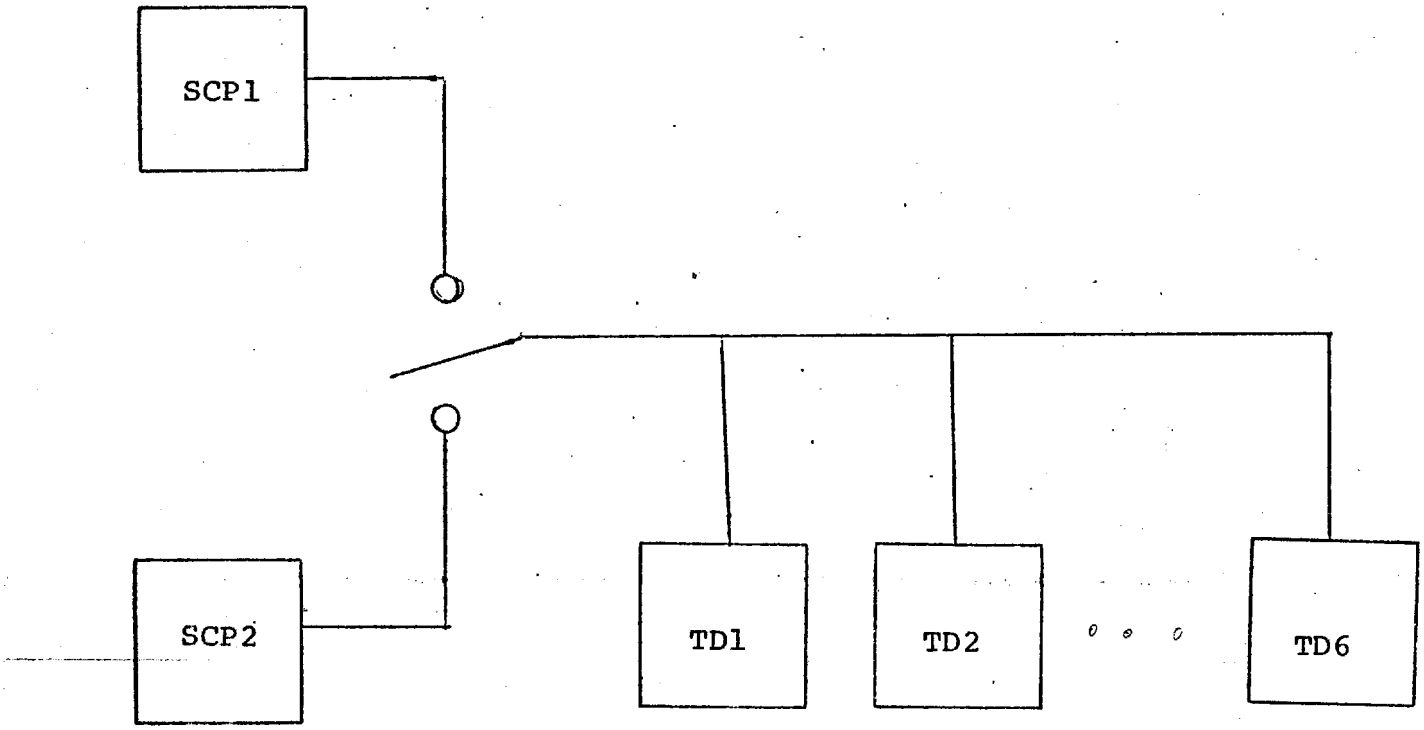

Approved For Release 2002/08/15 : CIA-RDP83T00573R000500080009-3

Implications

The above design gives control, at any given time, of all TD's to a single SCP. If dual SCP control of TD's is desired, the two SCP's could flip-flop by means of software logic. A flip-flop technique would permit segregating or sharing those functions in the two SCP's that require an SCP-TD association. Implementation of such a scheme adds complexity to the software logic which would have to reside in both SCP's. The areas particularly affected are hand-shaking protocols, timing, and recovery. It also requires that a wise choice be made initially as to which functions are shared or segregated between the SCP's. The choice is critical since once designed and implemented, subsequent changes would not be easy.

STATINTL [redacted] has chosen not to develop a flip-flop design. They have ruled out any type of dual SCP operation, the second SCP is for back-up only. The hardware design does not permit selective accessing of TD's by both SCP's concurrently and the software design requires that all accesses of TD's be made from only one SCP.

There are advantages to the design chosen by [redacted] A STATINTL single SCP controlling all the TD's needs only the simplest software and somewhat reduces maintenance of the hardware.

There is a major disadvantage to the design. It is impossible to use the redundant equipment for software development and testing. Although the back-up SCP and TD's can be put off-line there is no way of using them independent of the production system. This means that once the system is in production, fixing of existing bugs can be done only by taking down the production system. Assuming that the production system cannot be down very often such work would require long periods of elapsed time.

Larger efforts such as minor improvements to the existing system, development of the new functions, e.g. time sharing, and adjusting interfaces to new levels of operating systems become impossible unless more hardware is added or all of this work is contracted out.

It should be obvious to the reader that all of the above mentioned tasks are necessary. The delivered system will not survive unless it is enhanced and modified to accommodate new operating systems.

STATINTL The [redacted] Design versus Contract Specifications

STATINTL The [redacted] design of the accessing paths among the SCP's and TD's does not comply with the specifications stated in the STATINTL [redacted] Proposal TBMP 73-1 and in the Mass Storage System Design. The following paragraphs trace the reasons for our conclusion. Only documentation referenced by the contract is used as a basis for argument. We have had many conversations with [redacted] STATINTL during preparation of the specification that leave no doubt, however, it is felt they cannot be used as part of the record.

STATINTL Reference: [redacted] Proposal, Page 2-1, Paragraphs 3 and 4.

The TBM* Memory System is highly modular in construction. System capabilities can be varied over a wide range by configuring the system to include different numbers of each of five basic building blocks: Transport Modules (storage capacity), Transport Driver Modules (multiple seek/search), Data Channel Modules (internal data thruput), Storage Control Processor (file management) and External Data Channel Processor (data interface handling). Storage capacity is available over a range from 10^{11} bits to 3×10^{12} bits in 10^{11} bit increments, while sustained thruput can be specified up to 4.2M bytes per second in .7M bytes per second increments.

Switching matrices interconnect the hardware modules. They are constructed to allow multi-path access to any of the hardware modules. Systems configured with redundancy for all of the five basic building blocks therefore offer highly desirable degradation characteristics since no single unit failure brings the entire system down.

Commentary:

The above paragraphs introduce the general philosphy of the system. Namely it is intended to be modular, expandible, highly interconnected, and that given redundancy of equipment the system will be operable during the failure of one of the basic units. The statement concerning single unit failure is important. [redacted] has used a single line to STATINTL electrically interconnect the SCP's and TD's. Failure of this line will bring down the entire system. No provision for redundancy has been made by [redacted] Thus the design is incon- STATINTL sistent with the specification.

STATINTL Reference: [redacted] Proposal, Page 2-41, Paragraph 3.

Availability 24 Hours per Day

Scheduled maintenance of the MSS is performed on a module basis, and there is no requirement for scheduled down-time of the complete system. Preventive maintenance is usually conducted during off-hours. The hardware utilization does not exceed 12% for any module during the night shift (1800-0800). The Transport utilization is less than 2% during this period leaving more than ample for maintenance. Preventive or corrective maintenance can be conducted in the off-line mode concurrent with on-line operations. See Section 13.1 for maintenance procedures.

Commentary:

On-line refers to the array of hardware devices that are in use for production operations. Off-line refers to those hardware devices that are logically and sometimes physically disconnected from the on-line system. The paragraph above calls for maintenance of off-line devices concurrent with on-line operations. Some maintenance and hardware tests require that an SCP access a TD. When an SCP and a TD are
STATINTL both put off-line, the [redacted] design will not permit the needed access. Thus the specification that requires off-line maintenance
STATINTL concurrent with on-line production cannot be satisfied because of the way [redacted] has cabled the SCP's and TD's.

Reference: Mass Storage System Design, page 19, paragraph 1

TBMTAPE initialization is performed by a stand-alone SCP and TDP. Initializing a TBMTAPE begins with the recording and testing of three longitudinal tracks: the Address, Tally, and Control Tracks. This is followed by search testing to determine tape packing characteristics. Finally the wearing qualities of the tape are tested by repeated reads of a single block.

Commentary:

The key point is in the first sentence which pairs a stand-alone SCP with a TDP (Transport Driver Processor).

The TDP is an integral part of the TD. The term stand-alone is defined in the same document as an SCP being off-line to the Mass Storage System. The tape initialization process requires that a stand-alone SCP access a TD. This cannot be done given the described [redacted] design because the on-line SCP would have control of all of the TD's. The on-line SCP cannot transfer control of the TD's to the other SCP when it is in a stand-alone condition. Thus the referenced specification cannot be satisfied because the stand-alone SCP is unable to access a TD. STATINTL

Agency Requirements concerning SCP access to TD's

1. The Storage Control Processors (SCP's) must be able to access the Transport Drivers (TD's) in such a manner that given redundancy of SCP's and TD's, no single failure will cause the entire Mass Storage System to be inoperable.

2. A stand-alone SCP must be able to access a TD so that a TBMTAPE can be initialized.

3. An SCP must have access to the TD's such that off-line maintenance of an SCP can be concurrent with on-line operation of the Mass Storage System.

II. Usage of Two Storage Control Processors in the Mass Storage System

Software Design

STATINTL [redacted] stated that [redacted] is designing and developing STATINTL the Mass Storage System (MSS) such that only one Storage Control Processor (SCP) will be active. The second SCP's role is purely back-up and will be switched into the system when a failure occurs in the first SCP.

Implications

The above design greatly simplifies the software logic needed for the Mass Storage System. If a single SCP can drive the system such that system throughput can be maintained as specified then we cannot say that two active SCP's is superior to a single SCP system .

It is difficult to reconstruct why a dual SCP system was originally specified. The major problem here is that the revision of the design was done unilaterally by [redacted] There STATINTL were no prior joint discussions on this matter, we were simply STATINTL informed of the [redacted] decision.

STATINTL The [redacted] Design versus Contract Specifications

STATINTL The [redacted] design of using a single active SCP for the MSS rather than two active SCP's does not comply with the STATINTL specifications stated in the [redacted] Proprosal TBMP 73-1 and in the Mass Storage System Design. The following paragraphs trace the reasons.

STATINTL Reference: [redacted] Proposal, Page 5-1, Paragraphs 1 and 2.

Control of MSS is divided into three parts and is performed by three sets of computers. Overall system control is provided by the Storage Control Processor Complex consisting of one or more SCP's. The TBM* Memory System configured for the ORACLE application comprises two identical SCP's.

The SCP Complex communicates with subscribing host CPU's, performs file management and space allocation functions, defines the necessary functions and transmits the corresponding

commands to the other controllers in the TBM* Memory System. During normal operations, one of the SCP's acts as the Master (SCPM) exercising overall TBM* Memory System control while the second one operates in a Slave mode (SCPS) performing file management and space allocation functions.

Commentary:

The specification calls for two concurrently active SCP's having different but complementary functions.

STATINTL Reference: [redacted] Proposal, Pages 5-2, Paragraphs 2 and 3.

The SCPM exercises overall system control. It allocates tasks to the other processors within the MSS, keeps an audit trail and a file management trail for all tasks entering the MSS, keeps an activity log for each file and for hardware resources, monitors overall MSS operation, and automatically switches to a degraded mode if hardware resources become unavailable.

The SCPS normally performs space allocation for the DSS, and maintains the TBMCATALOG consisting of the Master File Directory of all files stored on TBMTAPE and the On-Line File Directory for all mounted TBMTAPE's in the DSS. The SCPS thus performs most of the file management functions under the supervision of SCPM.

Commentary:

These paragraphs give further detail about the functions to be allocated to each SCP.

STATINTL References: [redacted] Proposal, Page 5-2, Paragraph 4.
Page 5-11, Paragraph 5.
Page 5-12, Paragraph 2.
Page 5-14, Paragraph 2.
Page 2-10, Paragraph 5, 6, and 7.

Commentary:

These paragraphs are not reproduced here. All specifically discuss the concurrent use of a master and slave SCP.

Reference: Mass Storage System Design, Page 113, Paragraph 2.

The MSS software provides the capability to attach three consoles to each SCP in the system. (See Section 3 for the hardware configuration). During MSS operation, one SCP is considered to be the master SCP, the other the slave. The consoles attached to the master SCP are used to issue commands to the system. Other consoles are attached to various sub-components as required for maintenance. Hosts connected to the master SCP can also issue certain commands.

Commentary:

This paragraph gives detail about activity on the master SCP.

Summary

It should be clear that the [redacted] design which eliminates the master-slave SCP feature is contrary to the specifications referenced by the contract. STATINTL

Agency Requirements Concerning Dual SCP's.

Technically, it is not clear that the Mass Storage System requires more than a single active SCP. We must verify that the single mode will not create excessive overhead, however, before permitting [redacted] to continue. STATINTL

Other considerations concerning [redacted] overall performance should preclude simply permitting the specification change. We should not forget that the single or dual decision provides us with good leverage that can be used to advantage elsewhere. STATINTL

III. Access of Data Private to the Mass Storage System

Software Design

The Mass Storage System (MSS) keeps extensive information about the identification, location, and status of the data files under its control. This information along with other internal MSS records is stored on two disks that are private to the MSS. Only the MSS can directly access and use the information.

The Storage Control Processors (SCP's) and the External Data Channel Processors (EDCP's) are the MSS hardware modules that access these private disks. Presently, [redacted] has designed the software such that the SCP's require a dedicated disk control unit to access the private disks. Dedicated here, means that only the SCP's can use the control unit. When the EDCP's require access it must be gained via another control unit. STATINTL

The figure below shows the [redacted] design. STATINTL

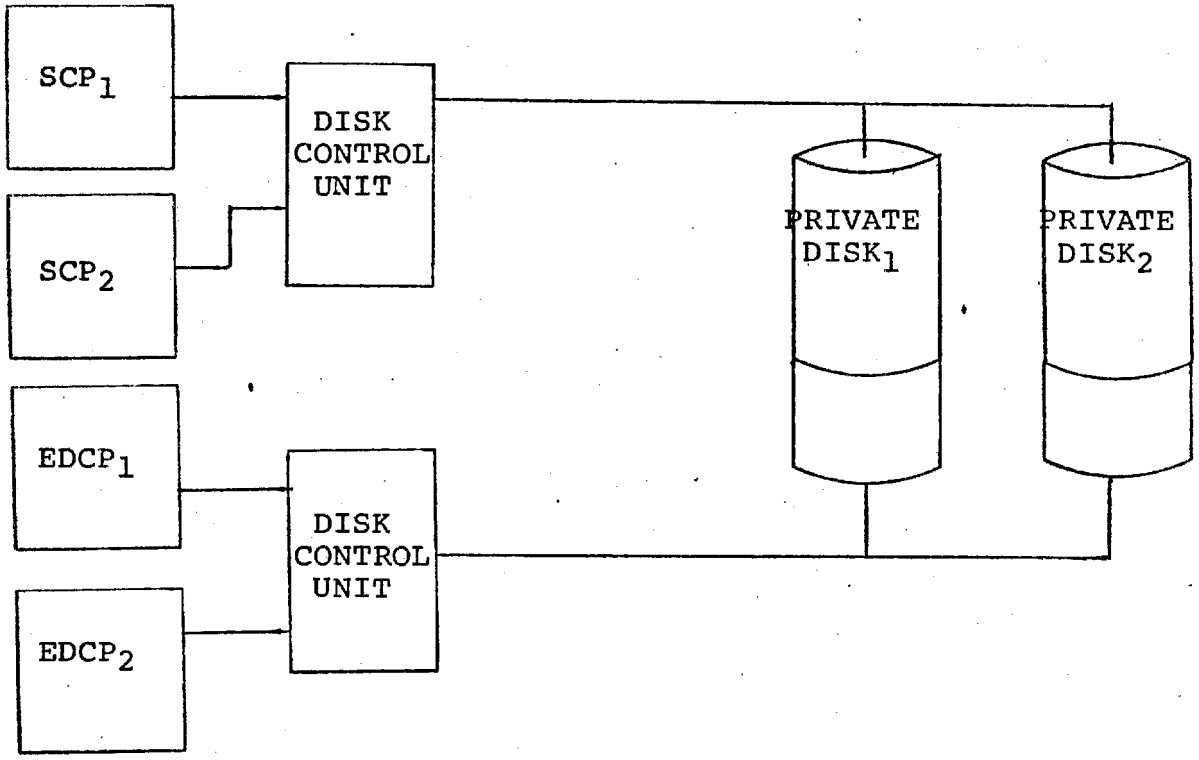

Implications

The use of a dedicated controller by the SCP's requires that a second one be made available in the event the first one fails. If this were not done the failure would result in the MSS being inoperable. Thus the impact of the [redacted] STATINTL design is limited to the additional funds required for a spare controller and the space needed for its placement.

STATINTL The [redacted] Design versus Contract Specifications

STATINTL The [redacted] design which requires a dedicated disk control unit to service the SCP's is contrary to what has been STATINTL specified in the [redacted] Proposal TBMP 73-1 and in the Mass Storage System Design. The MSS Design calls for two disk control units with dual access features to be used to serve any of the SCP's and EDCP's. The [redacted] Proposal specifies STATINTL a single control unit with a Four Channel Switch to serve the SCP's and EDCP's. Specific references and drawings are shown below.

STATINTL Reference: [redacted] Proposal, Page 2-24, Table entitled, "10^{12} Bit TBM* Memory System Hardware Configuration"

CCS

2 SCP's (Storage Control Processors)
2 EDCP's (External Data Control Processors)
1 3330 Control Unit with 2 Spindles and a Four Channel Switch
6 Host CPU message interface links on each SCP
8 Data interface links on each EDCP, providing for 16 shared device controllers such as 3830's, 3803's, channel to channel adapter, etc. in any combination.

Commentary:

This table lists the hardware needed for a 10^{12} bit system. The item - "1 3330 Control Unit with 2 Spindles and a Four Channel Switch" - states very clearly that only a single disk control unit is needed.

STATINTL Reference: [redacted] Proposal, Page 2-36, Paragraph 2

The 3330 disk system internal to CCS can be accessed from either of the SCP's and EDCP's. One spindle will normally be dedicated to the MSS Catalog and not accessible by the EDCP's. The second spindle serves as backup, and can also be used for internal data staging, diagnostics and maintenance functions, etc.

Commentary:

The first sentence refers to the private disks for use by the MSS. It provides further elaboration of the table described in the previous Reference.

STATINTL References: [redacted] Proposal, Page 2-25, Table entitled, "10" Bit Initial TBM* Memory System Hardware Configuration"

STATINTL [redacted] Proposal, Page 3-1, Table entitled, "Hardware Required for the Initial MSS'

STATINTL [redacted] Proposal, Page 4-1, Table entitled, "Hardware Required for the Complete MSS"

Commentary:

The above references are not reproduced here, but all of them list a single disk control unit to be used to service the SCP's and EDCP's.

Reference: Mass Storage System Design, Page 9, Figure 3.

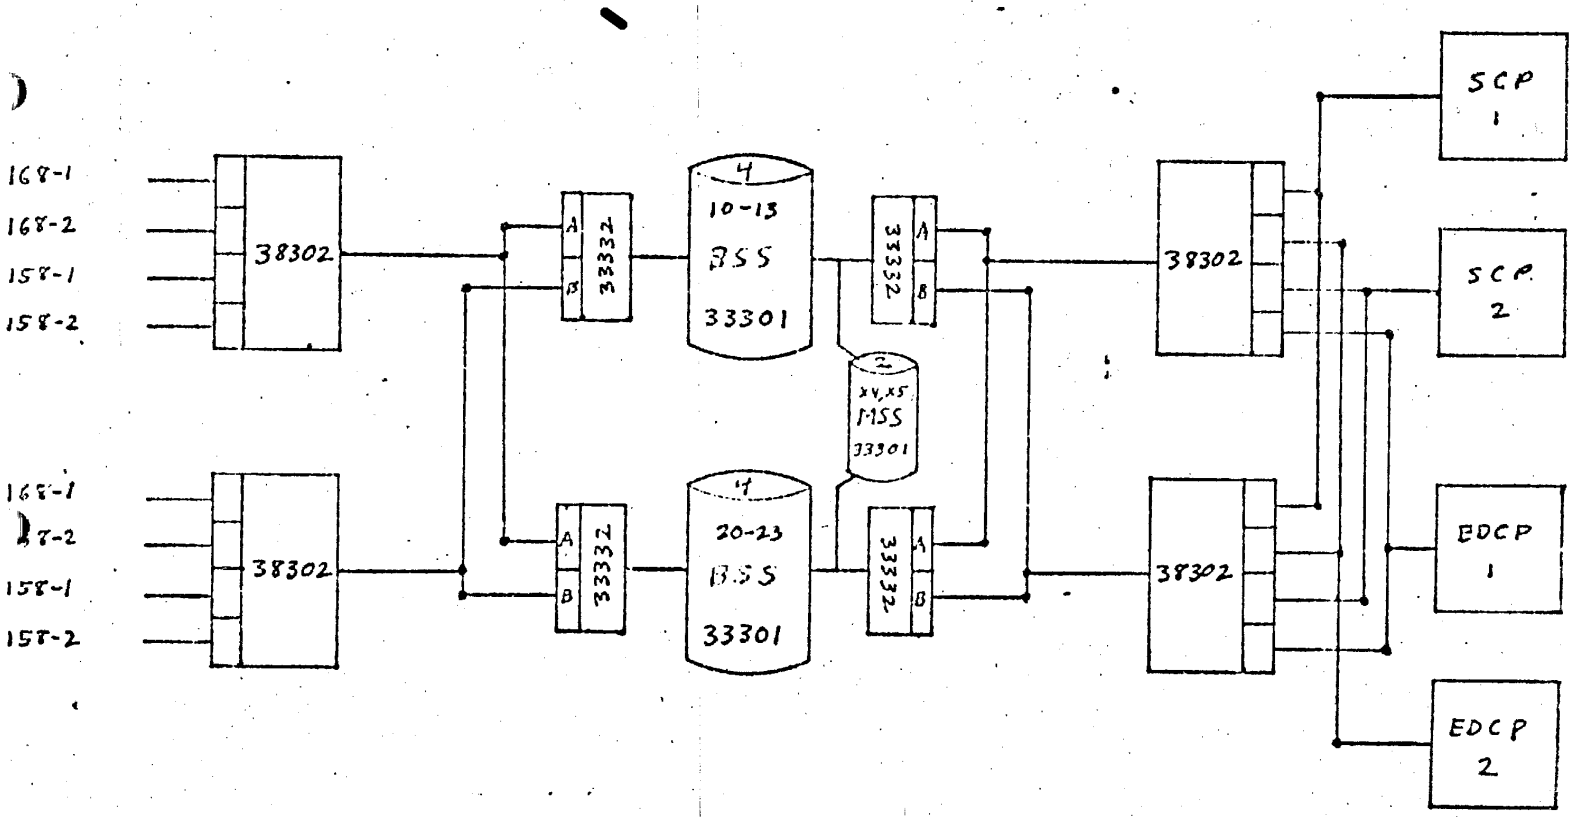

Figure 3. ORACLE BSS Configuration, Oct - Dec, 1975

Commentary:

The above figure shows two disk control units to provide all EDCP's and SCP's with access to the private disks. There is no control unit allocated specifically for the SCP's.

Summary:

All of the above references prove that the [redacted] design does not comply with the specifications. STATINTL

Agency Requirement

Access to the private disks of the MSS must be provided by disk control units that are shared among the SCP's and EDCP's as specified in the Mass Storage System Design (Specification) dated March 19, 1975.

IV. Data Channels, Transport Drivers, External Data Channel Processors and Their Interconnections

Hardware Design

The External Data Channel Processor (EDCP) controls data transfers between TBM Tape and disk. During a data transfer an EDCP is connected to a Data Channel (DC) and a disk control unit.

Each DC can have up to two EDCP's connected to it. Thus a 2 DC by 2 DCP array would have perfect switching capability. The number of concurrent data transfers possible with a 2 by 2 array is two in any direction. The EDCP's are the limiting factor. Two DC's give a capacity of four concurrent transfers, two reads and two writes. The addition of a third EDCP would permit three concurrent data transfers, two in one direction and one in the opposite direction. Because of the limitation, that a DC can have a maximum of two EDCP's connected to it the 2 by 3 array does not have perfect switching. It is possible, however, for the software to produce the maximum number of concurrent data transfers.

A fourth EDCP can be added creating a 2 by 4 array. Again this configuration does not have perfect switching. It is possible with software to produce four concurrent data transfers, two reads and two writes.

A fifth EDCP requires a third DC for that EDCP to be of any use.

Transport Drivers (TD's) control tape Transports and Data Channels (DC's). The TD when connected to the tape Transport causes tape searches, Address and Tally Track read/write processes and data erase functions. When a tape Transport has both a TD and a DC connected to it, then the TD can cause data transfers to or from the tape Transport via the DC.

Transport Drives and tape Transports can be interconnected so that any of six TD's can access any of sixty-four tape Transports. As stated above this type of connection permits all function associated with tape movement except data transfer.

Transport Drivers and Data Channels can be interconnected such that any of six. TD's can access any of two DC's. This connection permits the function of data transfer.

When a third DC is added to an array of 6 TD's and 2 DC's, the perfect accessing of the 6 by 2 cannot be retained. A 6 by 3 array is prohibited and some lesser split must be made such as a 4 by 2 and a 2 by 1. Each of two smaller arrays will have perfect switching capability.

Implications

STATINTL

We have believed that any configuration of up to 6 TD's, 3 DC's and 6 EDCP's would provide perfect switching. [redacted] has said this is not so. The software can be designed to provide the same capability. Producing such software is not difficult. Our problem here is that the software cannot be completely tested unless the maximum hardware configuration is provided. We can desk check it however, to insure allowance is made for this type of expansion. A specification on Page 55 of the Mass Storage System Design states that "... the system will initiate as many simultaneous transfers as possible depending on the hardware resources available ....".

STATINTL

Having perfect switching of the hardware would provide slightly simpler software. It would also eliminate the judgements as to just which devices should be called together. In other words the current [redacted] design for switching TD's, DC's and EDCP's does not create any real problems.

STATINTL The [redacted] Design versus Contract Specifications

STATINTL The [redacted] design does not provide perfect switching capability in the hardware. I feel that anything less than perfect switching does not comply with the specifications. The following pages trace ~~may~~ tortured reasoning.

STATINTL Reference: [redacted] Proposal, Page 2-1, Paragraphs 3 and 4.

"The TBM* Memory System is highly modular in construction. System capabilities can be varied over a wide range by configuring the system to include different numbers of each of five basic building blocks: Transport Modules (storage capacity), Transport Driver Modules (multiple seek/search), Data Channel Modules (internal data thruput), Storage Control Processor (file management) and External Data Channel Processor (data interface handling). Storage capacity is available over a range from 10^{11} bits to 3×10^{12} bits in 10^{11} bit increments, while sustained thruput can be specified up to 4.2M bytes per second in .7M bytes per second increments.

Switching matrices interconnect the hardware modules. They are constructed to allow multi-path access to any of the hardware modules. Systems configured with redundancy for all of the five basic building blocks therefore offer highly desirable degradation characteristics since no single unit failure brings the entire system down."

Commentary:

The second paragraph says there is "multi-path access to any of the hardware modules". In the case of TD's, DC's and EDCP's, the current STATINTL [redacted] design says there is multi-path access to some of the hardware modules.

Reference: [redacted] Proposal, Page 2-12, Paragraph 5.

STATINTL

"The TCIF is designed to effect interlocked burst-mode data exchanges with TBM* Memory System Read/Write Channels. Data so transferred is fetched from, or stored into, the Interface Core Buffers through direct memory access by the TCIF. The TCIF includes switch logic circuits through which it may connect to any of six TBM* Memory System Channels - three Read Channels and three Write Channel - under EDCP program direction."

Commentary:

STATINTL The [redacted] Proposal says the TCIF, ".... may connect to any of six TBM* Memory System channels ....". The TCIF is an integral part of an EDCP. [redacted] has designed the STATINTL EDCP to interconnect with 6 channels, therefore it is proper to conclude that an EDCP can have access to any 6 channels " - three Read Channels and three Write Channels - ...." which is the exact equivalent of three DC's. This means that a single EDCP should be able to be switched to STATINTL any of 3 DC's. [redacted] says no, only 2 DC's. They do not say it is because the EDCP doesn't have the capability, instead they say that a DC can only accept cabling from 2 EDCP's.

There are two points to consider here, first, [redacted] STATINTL does not state in writing that the DC's cannot handle three EDCP's. Second, a maximum configuration of six EDCP's and three DC's (specified in the Mass Storage System Design) gives the implication that the six EDCP's can access the three DC's. If not, why be so specific about maximum.

Considering the current [redacted] design, the specification STATINTL or maximum configuration is purely arbitrary and cannot be construed to have any relationship to the characteristics of the hardware.

STATINTL Reference: [redacted] Proposal, Page 2-20, Paragraphs 2 and 3.

"Each Transport contains all of the mechanical elements necessary to move tape, but only a fraction of the eletronics normally associated with tape Transports of this type. All of the electronics necessary to control a Transport resides within the Transport Driver Module while all of the electronics associated with data signal conditioning and transmission are located within the Data Channel Module. Transport Driver and Data Channel Modules are shared between all Transports in a system. This approach greatly reduces the cost of retaining a large on-line storage capacity since a few Transport Drivers and Data Channel Modules normally suffice to support even maximum capacity systems."

"Contained within each Transport Module rack are also the switching elements which allow for any one of up to six Transport Drivers to connect to either of the two Transports. Both Transports can be connected to two different

Transport Drivers concurrently. The switching elements are packed in easily replaceable switch drawers at the bottom of the Transport Module rack. Space constraints limit the size of the switch matrix in each Transport Module to six (Transport Drivers) by two (Transports)."

Commentary:

The above discussion is similar to the previous one about EDCP's and DC's. Here great care is taken to show that six TD's can access any of sixty-four tape Transports. Thus when [redacted] states the maximum system configuration (specified in the Mass Storage System Design) of 6 TD's, 64 Transports, 3 DC's and 6 EDCP's it seems to follow from the description of the hardware modules that there is good reason for these limitations. [redacted] has not stated that the maximum of 6 TD's and 3 DC's cannot be interconnected. This failure to make it clear that the maximum configuration is arbitrary and does not stem from the actual hardware characteristics is misleading and should be seriously questioned.

STATINTL

STATINTL

STATINTL Reference: [redacted] Proposal, Page 2-19, Figure 2.1.11

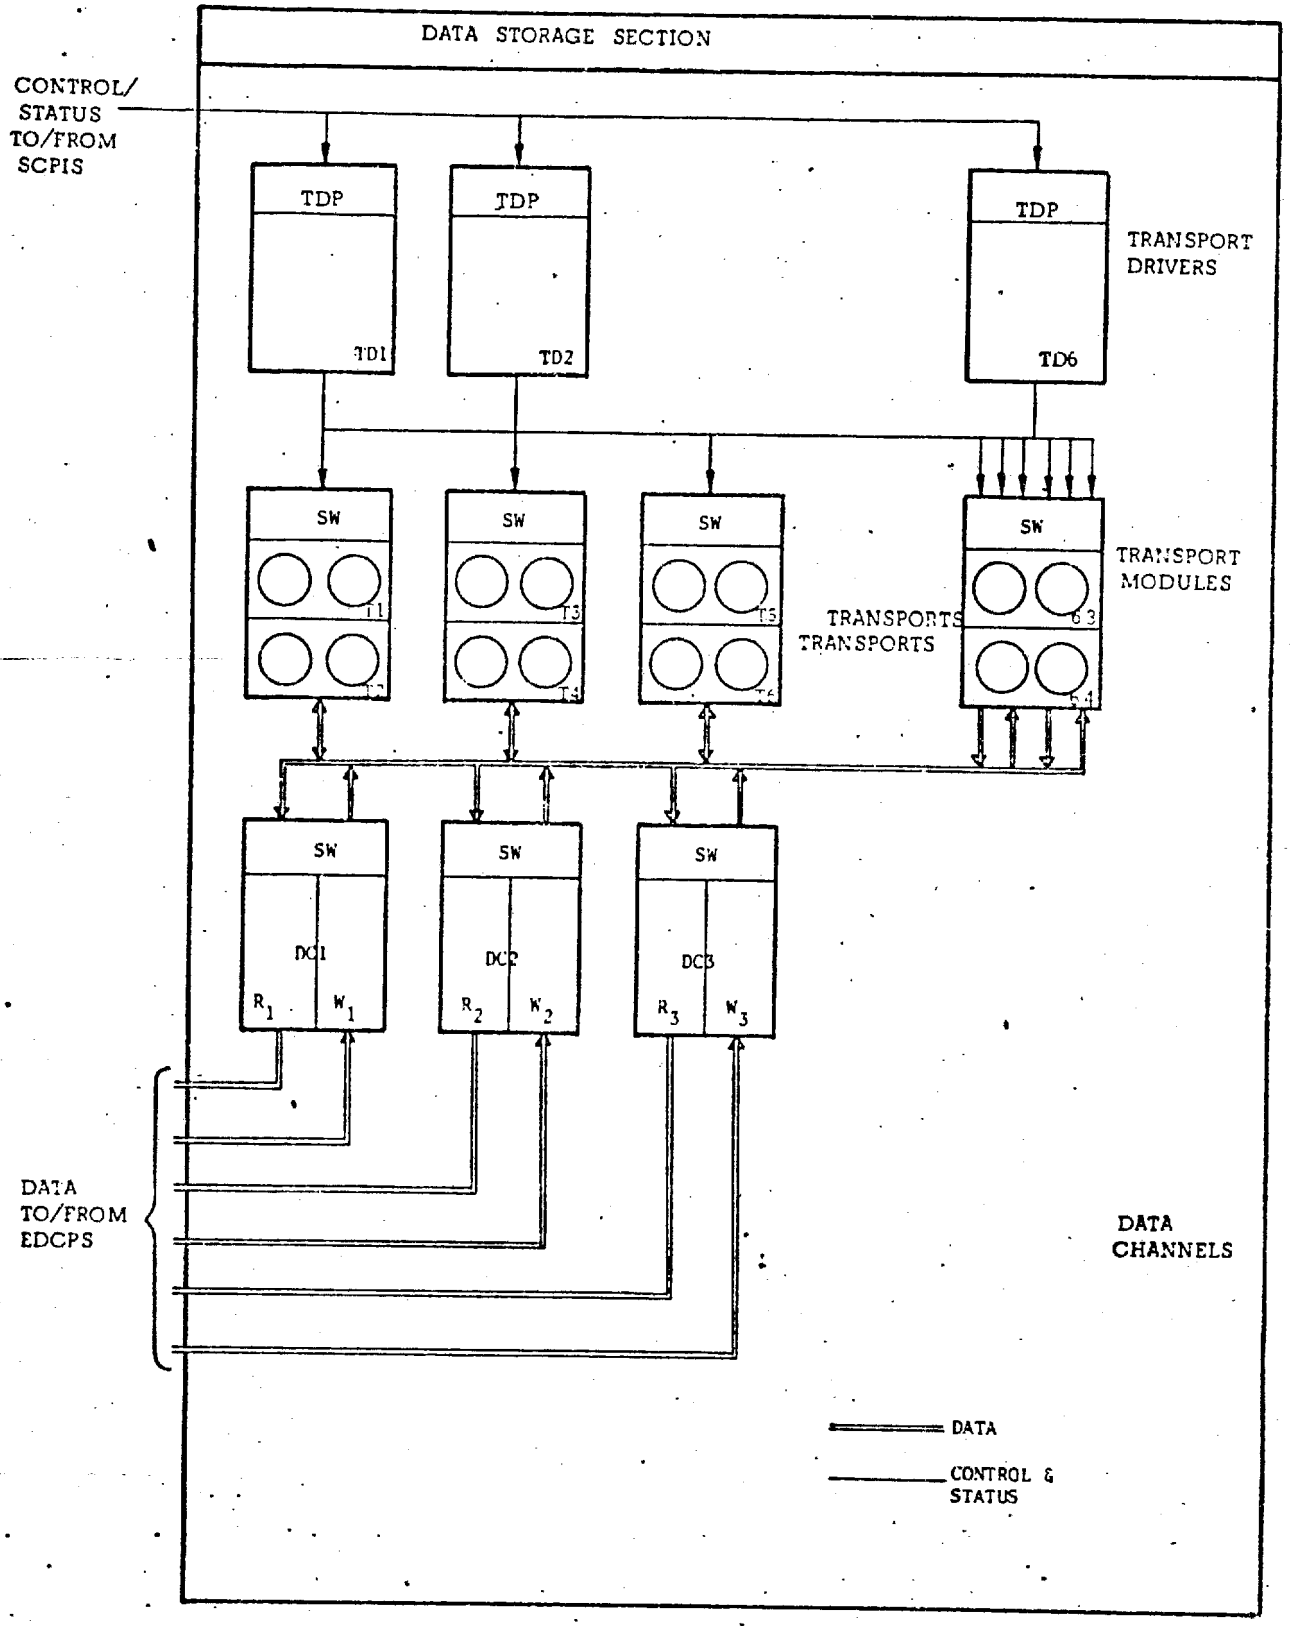

FIG. 2.1.11 - EXPANDED DSS CONFIGURATION

Commentary:

Figure 2.1.11 is referenced on page 2-16 as follows:

"The hardware modules and the interconnecting switching matrices are described below, and depicted in Figures 2.1.10 and 2.1.11." Inspection of the figure shows that the TD's (up to 6) can access any of up to 64 tape Transports. It also shows the 3 DC's accessing any of the 64 tape Transports. Finally, it shows the 3 DC's servicing 6 EDCP's.

STATINTL Figure 2.1.11 when related to all the recorded statements by [redacted] shows:

a. Maximums of 6 TD's, 64 Transports, 3 DC's, and 6 EDCP's.

b. Perfect switching or accessing between related devices.

Summary:

I feel there is a case to ask [redacted] STATINTL to provide perfect switching for the maximum system configuration of 6 TD's, 64 Transports, 3 DC's and 6 EDCP's. The weakness in all of the above arguments comes from the lack of positive statements concerning the capabilities of the Data Channels' accessibility. However, in light of the failure of [redacted] STATINTL to make specific statements about the limitations of the DC's, it can be argued that it is entirely reasonable to conclude there are no limitations other than those implied by the maximums.