

**digital**

## INTEROFFICE MEMORANDUM

0278

TO: Distribution

DATE: October 20, 1971

FROM: Peter van Roekens

DEPT: Programming

SUBJ: OS/45 Proposal Meeting

There will be a meeting on October 28 at 9:30 a.m. in conference room 12-2 to discuss the attached proposal.

It is being circulated for your review and comments.

Attachment: Proposal

Distribution

L. Wade	P. Conklin
H. Spencer	J. Hittell
D. Stone	OS/45 Group
B. Delagi	
R. Clayton	CC: A. Ryder
R. Frith	D. Schroeder
J. Lombardo	G. Thissell
A. Devault	PDP-11 Coordinating Committee

## OS/45 - A Proposal

0279

I. Introduction

On 17 Sept 71 Dick Clayton and Robin Frith presented their views on OS/45 to the OS/45 group.

Since that time we have reviewed the notes of the 17 Sept 71 meeting, added new inputs, and have investigated competitive systems. As a result, we have begun to acquire a bias as to the organization of OS/45. This paper presents the market orientation that has resulted from this bias and details the most current definition of OS/45.

## II. Programmers Overview of the PDP-11/45

The system programmer contemplating the design of an operating system for the 11/45 has two classes of problems to resolve:

- 1) the price range of possible configurations, and
- 2) several new hardware capabilities

The 11/45 has a remarkable range of potential prices: \$20,000-\$300,000. We cannot ignore the low end because we can expect competition in this range from IBM's first mini-computer, the System/7. At the upper end, even though our price/performance ratio humbles our competitors, we must deal with the IBM 1130 and 1800 and their immense software library. In addition to the challenge of devising a system which has upward compatibility over a price range which varies by an order of magnitude, the systems designer must contend with three new hardware options:

- \*Memory hierarchies (memories with different speeds);
- \*Independence of instruction and data space, and
- \*Segmentation.

Complete understanding of how to properly use these features has barely emerged from a research environment, yet we must make intelligent use of them in a production system.

To provide reasonable solution to these challenges will require a design phase pursued to an unusual depth of detail; else we run the risk of rendering the new hardware features either unuseable or not cost effective. Providing a software system whose facilities compliment those of the machine itself will depend on development of a design which demonstrates we have indeed mastered the requirements of configuration flexibility and innovative hardware.

### III. The 11/45 Marketplace - Developing a Workable Image of our Customers

A designer of any product must have prior to any design activity a distinct image of the individuals to whom he expects to sell his product. The data processing marketplace has five identifiable concentrations which reflect market needs\*.

- 1) Real Time
- 2) Scientific Batch
- 3) Time Sharing
- 4) Commercial Batch
- 5) Number Crunching

Let us eliminate item 5) from consideration immediately as inappropriate to the 11/45. The remaining four items represent the order of market priorities as specified by Dick Clayton and Robin Frith during our 17 Sept 71 meeting. During the same meeting Dick Clayton specified the following framework within which we should define OS/45:

- 1) Software support for the floating point unit and the segmentation unit should exist by July 1972.
- 2) We should announce OS/45 by June 1972 and deliver it during the second quarter of calendar year 1973.
- 3) OS/45 should unify PDP-11 software.

In addition, we have assumed that OS/45 will consume between 18 and 25 man-years of effort. Using the assumed manpower estimates it does not seem possible to attempt to satisfy the needs of the time sharing or commercial batch marketplace.

In commercial batch, DEC must compete directly against IBM and in a marketplace where IBM has no peer. We simply do not have the time or resources to design and implement an operating system for the 11/45 that would compete effectively against IBM's offerings. And even if we could produce such a system, does DEC presently have the sales and system force necessary to sell and service the commercial market? Current inputs indicate in the negative, and, hence, we recommend rejection of orienting OS/45 toward commercial batch.

---

\*Identifiable in the sense that their overlap does not result in a complete merging of the end user needs.

From the above we can extract the following composite of the OS/45 customer population:

He either already uses or has under consideration an IBM System/7, 1130, or 1800. His application requires only a subset of IBM's software for these machines. And finally, the existence of competitive equipment which has between three and five times the cost performance of the equivalent IBM system provides our hypothetical customer with sufficient reason not to remain with or choose IBM.

Of course, the computer marketplace does not exist entirely under the aegis of IBM and DEC, but we contend that if we can produce a software system which significantly impacts System/7, 1130, and 1800 sales, then we will have more than nullified the offerings of XDS, Hewlett-Packard, Data General, SEL, EMR, Varian, and Interdata all of whom offer real-time systems in the 16 bit class.

All these facts lead us to one inevitable conclusion: make sure OS/45 provides a set of scientific batch facilities that makes it possible for us to capture 10% of the potential 1130 market.

This brings us to the top priority on the list - Real Time. As with the 1130, we will confine ourselves to IBM. In real time, IBM offers the system/7, the 1800, and the 360/44. The 360/44 really belongs in a different class of equipment (Decsystem 10) so we will concentrate on System/7 and the 1800. First the 1800.

IBM has a total of 563 1800's installed or on order at an average system price of \$300,000. The 1800 hardware represents little or no competition for the 11/45. IBM does, however, tend to overwhelm their competition with software, which includes two systems (MPX and TSX) capable of running, simultaneously, real time in the foreground and batch in the background. As with the 1130, even though we can't hope to provide all the software IBM does, the task of selecting a competitive subset appears achievable. And even 10% of a \$200,000,000 market would handsomely repay our investment (and hopefully we would capture much more than 10% of this market).

Unlike the 1800, System/7 represents an insidious rather than a direct challenge to the 11/45. Only the smallest System/7 configurations offer any competition for the 11/45. In most of these small configurations, we suspect, an 11/20 would provide a more cost effective solution. Regardless of how DEC counters the threat of System/7, counter-it it must. IBMs track record for customer loyalty provides little comfort to the DEC salesman attempting to replace a System/7\* with a PDP-11. Once in the door with a system/7, add-on equipment and growth to larger systems will go to IBM by default; System/7 will lead to 1130's and 1800's. (Indeed, the initial System/7 marketing thrust practically requires that the user already have an 1130, 1800 or 360).

To compete with IBM in the real time market we suggest that OS/45 provide a real time capability that spans the entire price range of the 11/45 with upward compatibility of object programs provided across the entire range of possible configurations.

\* We should not delude ourselves regarding the potential of System/7. IBM has consistently enhanced products to meet the threat of competition - How about segmentation registers on System/7?

A time sharing orientation also seems unachievable on schedule with available resources. The pursuit of an 11/45 time sharing system also seems unadvisable if DEC decides to produce a small version of the 10. Any attempt to provide a multi-language time sharing system for the 11/45 (we already have a single language system in RSTS) runs the risk of colliding with the introduction of the small 10. Software production costs continue to rise and hardware costs continue to decline. And we have little assurance that the total cost of an 11/45 time sharing system (hardware plus software) will not exceed the total cost of the small 10.

The existence of RSTS and the relatively modest cost involved in altering RSTS to take advantage of the FPP and segmentation provides additional reason for avoiding a time sharing orientation for OS/45. If the future of time sharing depends on applications packages, and if BASIC Plus has sufficient language constructs to build most applications packages suited to the 11/45, then what incremental gain can we expect by producing a multi-language system? We cannot answer this question factually, but doubt that the incremental gain can offset the software development costs. Thus, we recommend rejection of a multi-language time sharing organization for OS/45.

We now arrive at scientific batch. We define this as batch streaming of FORTRAN programs and cite the 1130 Disk Monitor as the type of facility against which we can expect to compete.

IBM has installed or on order 3800 1130s at an average price of \$90,000.\* It would not surprise us if the 11/45 has a cost/performance ratio five times that of the equivalent 1130.

Dick Clayton projects 1000 11/45 sales over the life of the system. If we can capture 10% of IBM's 1130s as a result of providing a competitive scientific batch system, then we would help him achieve 1/3 of his goal.

Furthermore, the 1130 customer does not need the practically dimensionless volumes of software of the commercial market. These users rely on FORTRAN heavily (making it possible for him to convert at modest cost). And the size and cost of the 1130 itself places a practical limit on type of applications it can support.

---

\*September 1971 issue of Computer and Automation.

In summary, marketing will have support for the 11/45 PDP and segmentation by July '72, but at the cost of some additional software proliferation; they will have an announceable definition of OS/45 by June '72, and the announced OS/45 will take significant steps toward unifying PDP-11 software.

Let us return now to a more explicit definition of OS/45.



#### IV. Satisfying Marketing's Requirements

Dick Clayton and Robin Frith specified three requirements for OS/45:

- 1) Software support for FPP & Segmentation Unit by July 1972.
- 2) Announcement of OS/45 by June '72 for delivery during the second quarter of 1973.
- 3) Unification of PDP-11 software.

Before discussing details of OS/45 (in section III we established a customer profile; we have yet to describe how we intend to satisfy their requirements) let's examine each of these points and how we can satisfy them.

- 1) Local modification to existing software represents the most reasonable approach to meeting this requirement. RSTS can at modest cost make use of both the Segmentation Unit and the FPP. DOS plans also exist to make use of the segmentation unit.

Attempting to rush the design and implementation of the OS/45 Kernel in order to permit DOS & RSTS to convert in time to meet the July '72 deadline seems unwarranted; such haste will jeopardize both the July '72 date and the consistency and coherence of OS/45 over the longer term. Of course, by making 11/45 oriented modifications to DOS and RSTS, we pay the price of continuing the proliferation of software systems for the 11 line.

- 2) Without a doubt the OS/45 group will have a system defined for announcement by June '72, but the scope of the system depends on available resources (Dick Clayton has already specified delivery requirements)
- 3) Initially the objective of unifying PDP-11 software will not happen. DOS and RSTS will evolve independently, and it does not seem advisable to attempt to prevent this.

But as we will describe shortly, OS/45 as it evolves will make every attempt to reclaim as much existing 11 software as possible. The OS/45 design will provide users with a system covering a broad range of configurations, programming facilities for real time and batch, and will disrupt existing user interfaces only where absolutely essential.

We have identified our users in Section III. They have real time requirements and scientific batch requirements. At the low end we must block the purchase of a System/7; at the high end we must overcome the presumed user benefits of IBM's vast software library.

We believe we can satisfy these requirements, by offering a system which provides upward compatibility across the entire price range which is available on the 11/45 (20,000-300,000). This IBM cannot do; movement within System/7, 1800, and 1130 requires a conversion effort; OS/45 will not.

Now, it turns out, that Segmentation provides an efficient hardware mechanism for implementing a system which can cover the configuration range under discussion. With segmentation hardware and a set of software standards, we can specify a cascade of hardware configurations each of which requires additional hardware, in order to acquire more elaborate services; all the while we guarantee complete upward compatibility. The success of this approach depends on a careful definition of the user virtual machine.

Basically this means that the user of an OS/45 system has a well defined set of facilities he may use. These facilities consist of a subset of the PDP-11 instruction set and a collection of service routines. As configurations grow in complexity, the set of services expands correspondingly, but we always provide complete upward compatibility. This scheme implies that every OS/45 configuration operates with a set of supervisory code. This code, of course, will vary considerably on different configuration classes and in every case appears transparent to the user. Let's examine some possible configurations:\*

Hardware

- 1) 11/45\*\*  
4K-12K  
TTY

Software

**Foreground only.**

Small systems suffer from the lack of adequate program preparation facilities. If you can prepare your programs on a larger system often 4K suffices to meet the needs of the application. IBM has solved this problem for System/7 by providing host preparation facilities on larger equipment (1130, 1800, 360). With this technique they provide a Macro system called MSP/7. We see no reason why we cannot do the same. Indeed, if we intend to meet the threat of System/7, host preparation facilities seem essential.

\*These represent examples of possible configurations; the reader should not accept them literally.

\*\*We intend to investigate full 11/20 compatibility and issue a memo describing our conclusions.

- 2) Hardware  
 11/45  
 12K-16K  
 TTY

Software  
 Disk, tape, etc.  
 DOS subset  
 INDAC/11  
 Overlay Facilities  
 Fortran  
 Host preparation eliminated  
 Foreground only.

- 3) Hardware  
 11/45  
 16K-24K  
 TTY  
 Disk, tape  
 Segmentation

Software  
 Same as 2, plus:  
 Foreground single background stream. System  
 maintains complete isolation between foreground  
 and background.  
 Foreground and background operate in fixed  
 partitions.

- 4) Hardware  
 Same hardware as  
 3) but with 28K

Software  
 same as 3) plus  
 support of background  
 jobs whose size ex-  
 ceeds that of physically  
 available core.

- 5) Hardware  
 same as 3 but with 32K

Software  
 4) plus  
 shared code  
 Multiple background jobs.  
 Index Sequential file system.

- 6) Hardware  
 5 but with 40K

Software  
 5) plus  
 Multiple Foreground jobs (Individually protected).

Once we define the basic user virtual machine we can determine exactly which particular configuration classes available resources permit us to produce. Furthermore, with well defined configuration classes, the product manager has available to him a shopping list that enables him to make cost trade-offs far more reasonably than he can at present. It also makes the programming department more aware of the incremental costs involved as you move up the scale in system complexity.

**VI. Summary and Conclusions**

If DEC continues to grow it must eventually increase its business at the expense of IBM. To accomplish this traditionally unaccomplishable feat, we have suggested a software plan for OS/45 which confronts IBM where they appear most vulnerable - real time and scientific batch.

IBM, with present offerings, provides zero hardware competition for the 11/45. To counter IBM's software libraries, OS/45 takes an approach IBM cannot easily counter: upward compatibility within a price range that completely covers IBM's real time offerings in the 16 bit class.