

U-IR2  
Letter  
Original

September 26, 1986

NO ITEM TO INSERT

NO ITEM TO INSERT

Dear

NO ITEM TO INSERT

:

Please find enclosed the latest INPUT report on large-scale IBM and software-compatible mainframes, entitled Large-Scale Systems Directions: Mid-Year Update.

This report examines the potential impact on large-scale systems of some of the "solutions" which are being proposed for end-user computing at both the desktop and departmental levels. This is done by analyzing in a structured manner the considerations which are important in the distribution of data and processing over networks.

Also, this report compares past INPUT residual value projections from 1981 through 1985 with actual used market retail prices. We are pleased to say that the results seem quite reasonable. Detailed analysis will be contained in the next report in the series.

Please note that the graphic displays of residual values are no longer included in the report. The data is presented in tabular form. Please call me to discuss this if you would prefer the graphic representation.

Yours sincerely,

Peter A. Cunningham  
President

PAC:ml

Enclosure



LARGE-SCALE SYSTEMS DIRECTIONS:  
MID-YEAR UPDATE

AUGUST 1986

[The page contains extremely faint and illegible text, likely bleed-through from the reverse side of the document. No specific content can be transcribed.]

Published by  
INPUT  
1943 Landings Drive  
Mountain View, CA 94043  
U.S.A.

**Information Systems Program (ISP)**

***Large-Scale Systems Directions:  
Mid-Year Update***

Copyright ©1986 by INPUT. All rights reserved.  
No part of this publication may be reproduced or  
distributed in any form or by any means, or stored  
in a data base or retrieval system, without the prior  
written permission of the publisher.



**LARGE-SCALE SYSTEMS DIRECTIONS:  
MID-YEAR UPDATE**

**ABSTRACT**

This report examines the potential impact on large-scale systems of some of the "solutions" which are being proposed for end-user computing at both the desktop and departmental levels. This is done by analyzing in a structured manner the considerations which are important in the distribution of data and processing over networks.

Also, this report compares past INPUT residual value projections from 1981 through 1985 with actual used market retail prices. We are pleased to say that the results seem quite reasonable. Detailed analysis will be contained in the next report in the series.

This report contains 64 pages, including 20 exhibits.





LARGE-SCALE SYSTEMS DIRECTIONS:  
MID-YEAR UPDATE

CONTENTS

	<u>Page</u>
I INTRODUCTION .....	I
II DISTRIBUTED PROCESSING AND "CONNECTIVITY" .....	3
A. Top-Down versus Bottom-Up	3
1. Peers to Peons?	3
2. Centralization and Integration	5
3. Differentiation and Mechanization	6
4. Turbulence At All Levels	11
B. Structured Analysis	16
1. Operating Systems Functions	17
2. Quality	25
3. Network Hierarchy	27
4. Software Hierarchy	28
C. Potential Residual Costs	33
III RESIDUAL VALUE FORECASTS .....	35
A. Review of Used Market Prices versus Projected Residual Values	35
B. Announcements	51
C. Assumptions	53
D. Projected Used Market Prices and Residual Values	56
APPENDIX A: SET LISTS USED IN ANALYSIS .....	61



LARGE-SCALE SYSTEMS DIRECTIONS:  
MID-YEAR UPDATE

EXHIBITS

		<u>Page</u>
II-	-1 End-User Software Environment	13
	-2 Distribution of Operating Systems Functions	15
	-3 Data Flow Diagram	31
III	-1 Projected versus Actual Used Market Retail Values-- IBM 3033 Processor	36
	-2 Projected versus Actual Used Market Retail Values-- IBM 4331-GrpII Processor	37
	-3 Projected versus Actual Used Market Retail Values-- IBM 4341-GrpII Processor	38
	-4 Projected versus Actual Used Market Retail Values-- IBM 4361 Processor	39
	-5 Projected versus Actual Used Market Retail Values-- IBM 4381 Processor	40
	-6 Projected versus Actual Used Market Retail Values-- IBM 3083-B Processor	41
	-7 Projected versus Actual Used Market Retail Values-- IBM 3081-G Processor	42
	-8 Projected versus Actual Used Market Retail Values-- IBM 3081-K Processor	43
	-9 Projected versus Actual Used Market Retail Values-- IBM 3081-KX Processor	44
	-10 Projected versus Actual Used Market Retail Values-- IBM 3350-A02 Disk	45
	-11 Projected versus Actual Used Market Retail Values-- IBM 3380-AA4 Disk	46
	-12 Projected versus Actual Used Market Retail Values-- IBM 3420-005 Tape Drive	47
	-13 Projected versus Actual Used Market Retail Values-- IBM 3420-008 Tape Drive	48
	-14 Projected versus Actual Used Market Retail Values-- IBM 3800-001 Printer	49
	-15 Projected Residual Value as a Percent of Vendor List Price	57
	-16 Projected Used Market Retail Value at January 1	58
	-17 Projected Residual Value as a Percent of Vendor List Price (Updated)	59



## I INTRODUCTION

- The last Large-Scale Systems Directions report presented IBM's large-scale systems directions for the remainder of the 1980s. It was pointed out that planning for large-scale systems (both IBM's hardware and systems software planning and the IS departments applications systems planning) cannot be effectively done without consideration of what is being done at other levels in the computer/communications network hierarchy. Vague general terms such as micro-mainframe links, departmental systems, and "connectivity" are not of much help for those who are ultimately responsible for developing a reasonable and cost-effective plan for traditional large-scale systems.
- The purpose of this report will be to give the beleaguered IS department some understanding of the potential impact on large-scale systems (in the broadest sense—hardware, systems software, and applications) of some of the "solutions" which are being proposed for end-user computing at both the desktop and departmental levels. The way this will be done is by analyzing in a structured manner the considerations which are of importance in the distribution of both data and processing over the computer/communications network.
- The objective of this analysis is to refine the general areas of residual costs which were presented in the earlier Large-Scale Systems Directions report.
- Chapter III of this report, in addition to providing the usual projections of used market prices and forecasts of residual values, will review past used market price projections and list the current general assumptions upon which INPUT's residual value forecasts are made.









## II DISTRIBUTED PROCESSING AND "CONNECTIVITY"

### A. TOP-DOWN VERSUS BOTTOM-UP

#### I. PEERS TO PEONS?

- For the last ten years, INPUT has recommended a simple strategy for the development of computer/communications networks.
  - First, all processing and data bases are centralized on large mainframes. This consolidation of applications systems can be cost-justified easily when compared with standalone decentralized systems, but its real benefits accrue for other reasons.
    - The centralization effort permits (indeed, encourages) the standardization of essential hardware, software, and data bases.
    - A central focus for quality assurance of applications systems and data bases is established.
  - Then, the central facility provides for the "orderly distribution" of processing (and data bases) from the central facility to local nodes. This "orderly distribution" implies a top-down (or hierarchical) architecture in the best sense of the word, and this architecture has the following implications:



- The distribution of processing and data over the network will be made on an application-by-application basis.
  - The design and development of the individual applications systems will be done on a top-down basis (in the sense of structured programming).
  - The "proper distribution" of systems functions will be dictated by performance in terms of systems cost and quality.
  - The design and development of such complex systems will require the best efforts of highly qualified systems personnel. (One reason for the initial centralization is to conserve this scarce human resource.)
- It should be noted that in this practical, if somewhat idealized, scenario the process of integrating standalone computer systems occurs during the standardization which accompanies centralization. This integration process serves another purpose as well--it encourages acceptance (sometimes reluctant) of the fact that both the corporate IS function and the end-user community work for the same company, and that both are willing (or are forced) to give up some power in the process. (The end-user organizations relinquish power at the time of consolidation into central host mainframes and the central facility at the time that processing is distributed back to the end users.)
- Unfortunately, this process of consolidation and integration followed by distribution of processing has usually been accompanied by a struggle for control of the hardware. This struggle to control computers within organizations seems to be based on an instinctive knowledge that information represents power which is independent of vendors' constant reminders that this is so. Where centralization efforts have been successful, the IS function



has seldom distributed processing back to end users gracefully (much less intelligently). In fact, the whole personal computer revolution started with a "computers for people" flavor which was directed specifically against the bastions of the central IS department. Having achieved some control over their own destinies, there is little indication that end users will be any more gracious in relinquishing power than the IS function has been.

- Central IS departments which have viewed distributed processing as throwing a few crumbs among the peasants are now being asked to provide "connectivity" into the power base (corporate data bases), and end users would like nothing more than to make central systems departments data custodians while they develop applications systems from the bottom up. Connectivity does not imply "peer-to-peer" communications at present; whether you start from mainframes, departmental processors, or personal computers, the peers at one level are trying to turn those at other levels into peons.

## 2. CENTRALIZATION AND INTEGRATION

- Systems software is the key to how processing and data will be distributed over computer/communications networks, and MVS/XA and IMS are the tools of centralization. The processing burden associated with these systems will assure a continuing demand for mainframe processing power even if applications programs are distributed to other levels in the processing hierarchy. The question of whether MVS/XA is necessary to run the 3090 effectively, or whether the 3090 is necessary to drive MVS/XA, is moot--IBM mainframe hardware/software architecture has become the monolithic heart of SNA. And, IBM's large-scale systems strategy (as defined in INPUT's last Large-Scale Systems Directions report) is designed to see that it remains in place regardless of what goes on around it.
- However, computer/communications networks are systems, and according to general systems theory (GST), centralization is only one of four parallel trends which all systems exhibit (the others being integration, differentiation, and

1860

1861

1862

1863

1864

1865

1866

1867

1868

1869

1870

1871

mechanization). Advancing hardware/software technology over the last 15 to 20 years has provided impetus and economic justification for both the geographic and architectural distribution of processing from these large, general purpose mainframes. While IBM has been effective in slowing the impact of changing technology on its mainframe-oriented strategy, minicomputers have proliferated at the departmental (or factory, laboratory, etc.) level, and the ubiquitous microprocessor has popped up everywhere.

- While this distribution of processing has been far from orderly, it is partially attributable to IBM's inadequate efforts over the years. From the 3705 and 3790 to the 8100 and 3725, and including the recent promotion of the System 36 as a departmental processor, IBM has proposed underpowered engines with inadequate systems software for distributed processing. Failure to extend adequate computer power out from the mainframe has resulted in end users' bottom-up approach.
- Since many applications are not being designed from the top down in the current environment (INPUT refers to this environment as distributed systems development or DSD), whatever is happening out there must eventually be integrated with the host systems. In recognition of this, IBM now supports VM and DB2 (both of which had been kicking around within IBM for over 15 years) as necessary tools of integration. Unfortunately, the connection (or integration) of VM and DB2 with MVS/XA and IMS can present some operational and performance problems at the mainframe level.

### 3. DIFFERENTIATION AND MECHANIZATION

- A number of years ago, INPUT started to distinguish between geographic and architectural distribution of processing. The primary purpose of this was to illustrate that there were two ways to offload large mainframes:
  - Geographic distribution of processing is accomplished using general purpose minicomputers and microprocessors (intelligent workstations).





This type of distribution is usually differentiated from mainframe processing by providing specialized (and/or improved) systems and applications software.

- Architectural distribution of processing is accomplished by mechanizing specific functions or applications in separate hardware processors.
- For example, at the minicomputer level, UNIX was designed specifically for the interactive development environment with a relatively small number of users. It was ideally suited for developing highly responsive interactive systems and became a standard for DEC 16-bit minicomputers. However, being king of the minicomputers was not enough, and UNIX has been extended down to personal computer and up to Amdahl mainframes. IBM's reaction to the growing talk of UNIX as a "standard" was described in detail by INPUT in Large-Scale Systems Directions, 1985, but it warrants mention here.
  - IBM's version of UNIX (IX/370) was implemented on mainframes under VM. These virtual machines can be used to replace distributed minicomputers, and this strategy was supported by the simultaneous announcement of the 3044 fiber optic channel extender link.
  - Therefore, while IBM's action was an example of differentiation in terms of addressing a specialized "market" of dedicated UNIX users, it provided tight integration with IBM mainframe operating systems.
  - As such, the IBM implementation is yet another example of the continuing battle against distributed processing on minicomputers. Perhaps, in the latest terminology, it should be called "connectivity through absorption."
- Then, of course, there is PC DOS, which isn't exactly what IBM normally thinks of when it talks about operating systems. Rumor has it that MicroSoft



subcontracted to another firm for the original version and paid \$50,000 for it. It probably cost IBM more money than that to get permission even to talk with Bill Gates about it. If MVS/XA and PC DOS ever get connected on a "peer-to-peer" basis, the concept of peerage will have to be expanded considerably, but it is necessary to take a look at what is "going on out there" in order to understand some of the ramifications of bottom-up applications design.

- First of all, there are a lot of vendors running around calling spreadsheet packages and DBMSs "applications" and "solutions software." This isn't so bad, but there are those who actually believe that these "applications" will reduce the backlog in the IS department. This has had both fortunate and unfortunate consequences: it is fortunate that there seems to be a growing awareness that once you get past word processing, applications require data; and, it is unfortunate that a lot of systems development work has been, and is being, delayed while everyone waits for these magical solutions to appear (it is our opinion that the "slump" in the computer industry can partially be attributed to these failed "solutions").
  
- Of course, there are those who place the problem of the continuing applications backlog squarely with the IS department (which has assumed the additional burden of analyzing the "solutions software" and training end users in its use), and state that at least personal computers are improving the "productivity" of those who are using them. This claim has more validity, but once again one has to wonder when professional and even managerial employees start spending an inordinate amount of time on the following activities:
  - Entering data and text through keyboards.
  
  - Serving as media handlers and data base administrators as they physically keep track of floppy disks and arrange for back-up and file reorganization.

1870

...

...

...

...

...

- . Serve as computer operators as they deal with the systems software to transfer control and data among various applications.
  - . Worry about printer options and pin settings when they install the latest "solution."
  - . Program their spreadsheet "application" as it grows to encompass the universe of their work and becomes an enormous, sparsely populated matrix with logic and arithmetic hidden in the cells.
- Some of the tools which have been developed at the personal computer level are easy to use, but they are also easy to misuse. The mechanized spreadsheet is a good example--there are those (usually the most skilled users) who are building applications where the spreadsheet becomes both the program and the data base. Beyond the simplest level, these applications become the antithesis of structured methodology, and "connectivity" (much less integration) with mainframe DBMSs and/or mainframe applications represents a quality control problem to which there is no solution short of starting over from the top down. It should be hoped that too much "progress" is not made with end-user computing before the inevitable integration problems are faced.
- In addition to general purpose mini and personal computers, there are also those which have been spawned by office automation products, and these now have visions of becoming "departmental processors." The developers of these systems are normally characterized as being somewhat naive about data processing, but they may eventually be armed with optical disk-based electronic filing systems which will change the way information is stored, retrieved, and distributed. IBM is not in the mood for any significant shift

Faint, illegible text, possibly bleed-through from the reverse side of the page.

away from magnetic storage (especially at the mainframe level) at this point, but the replacement of paper files should have high priority in office automation and the economics favor optical media.

- While the dynamics of geographic distribution of processing have been apparent for some time, IBM has been reasonably successful in containing architectural distribution of processing. However, specialized processors and new architectures are beginning to show increasing promise.
  - Numerous vector and array processors are available, and last October IBM announced its own vector processor for the 3090 series.
  - HP has "bet the company" on RISC machines, and IBM has announced a RISC workstation.
  - Data base machines are becoming more popular, and while IBM still vows to "keep them under the covers," that may not suffice because IBM customers are discovering that it is just as easy to transfer files from IMS to a Teradata as it is to DB2. And then, the data base machine is a much better performer for information center applications.
- All of this architectural distribution of processing has enormous (and obvious) potential impact on large-scale systems directions when viewed in conjunction with geographic distribution of processing in hierarchical computer/communications networks. Take any function or application of large-scale mainframes and there appears to be a more cost-effective distributed solution. However, as any experienced systems analyst/programmer knows, when you start from the bottom up the problems becomes those of integration and interfaces, and these problems require more effort than the sum of the effort which went into the original solutions. All of the talk about "connectivity" is merely an attempt to integrate partial (and even failed) solutions from the past, and with distributed systems development, the problems are going to increase exponentially.





#### 4. TURBULENCE AT ALL LEVELS

- IBM's view of large systems "trends and directions" was presented in the last report of this series. While that view does not completely ignore what might be happening "out there" at the departmental level and on the desk top, the growth rates which are being used for planning do not seem to anticipate very much impact on large-scale systems from either geographic or architectural distribution of processing well into the 1990s. It is probable that this overt confidence is being severely tested as alternatives to mainframe processing continue to proliferate.
- Minicomputers, long the only vehicle of distributed processing, are now being squeezed by both geographic and architectural distribution of processing themselves. Recent INPUT research indicates that many users feel there will be no need for general purpose minicomputers once the 80386 microprocessor-based systems hit the market, and they are willing to wait before really addressing the "departmental processor" problem.
- Then, at the desktop level, there is an awareness among users (and systems personnel) that the "applications" and "solutions software" they have purchased require a lot of detailed systems work before any practical results are achieved. In addition, there is a growing suspicion that vendors don't have a workable solution to the "connectivity" problem, and even if they did, there remain serious questions about what individual users would be doing even if they were connected to every computer, terminal, and information source in the world.
- As each level in the processing hierarchy competes against the other, there is technological turbulence at, and among, all levels. This environment has created a level of indecisiveness among users which is unprecedented since the first commercial computers were introduced over 30 years ago. Fundamentally, this indecision concerns how processing and data should be



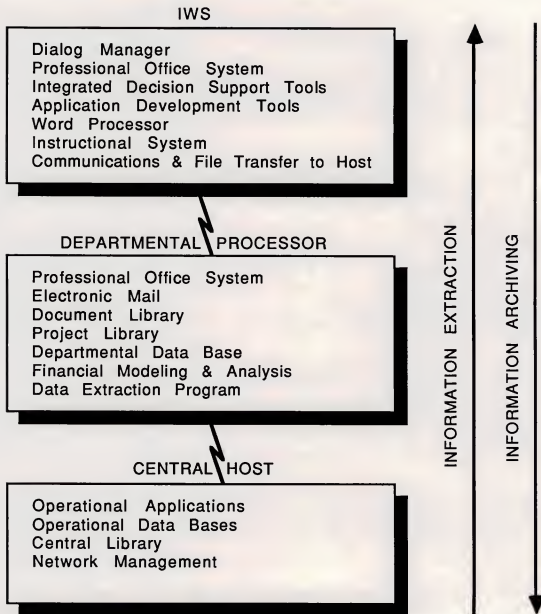
distributed over computer/communications networks (a subject INPUT has attempted to address over the last ten years). While IBM has normally benefited from such uncertainty (indeed, IBM's reluctance to distribute processing and data from mainframes has been the crux of the problem), the current level of indecisiveness is impacting the bottom line and IBM is being forced to respond.

- INPUT has recommended a "proper" hierarchical network in which certain functions are distributed over three levels of processors (very large mainframes, minicomputers, and microprocessor-based workstations). During recent months there has been increasing controversy about two-tiered versus three-tiered approaches to distributed processing, and IBM's beleaguered Systems Products Division has been making repeated public announcements of the need for a variety of incompatible "departmental systems," all of which will be connected in the network with Advanced Peer-to-Peer Networking (APPN) and therefore be transparent to the user.
- While these public pronouncements have led many to believe that IBM is supporting and even "encouraging" a three-tiered approach to networking, IBM's three-tiered approach should not be confused with the three levels of distributed processing INPUT has mentioned on so many occasions. Levels of hardware and "connectivity" mean little unless there is some understanding of how processing and data will be distributed over these networks. Fortunately, it may be possible to gain some understanding of what IBM really means by its three-tiered approach by reviewing a presentation which was made at an IBM "Consultant and Computer Services Executive Conference" recently.
- The particular presentation was made by the Director of Strategic Planning for IBM Information Services and was titled "End-User Software Environment." It clearly depicts a three-tiered environment (see Exhibit II-1). A few preliminary comments concerning IBM's view of a three-tiered network were:



EXHIBIT II-1

END-USER SOFTWARE ENVIRONMENT





- The distribution of data bases appears to be nearly identical with that which was predicted by INPUT shortly after IBM announced DB2 (see Exhibit II-3, "Projected Structure of Distributed Data Bases," Large-Scale System Directions: Disk, Tape, and Printer Systems, 1984). INPUT's purpose in presenting that particular scenario was to ". . . illustrate the enormous demand for on-line storage that will be generated."
  
- By keeping all "operational" applications and data bases on the central host (IBM refers to IMS as being appropriate for operational data bases), practically all transaction processing remains on the host and it becomes the "large host data base machine" which was also predicted by INPUT. (See Exhibit II-2, "Large Host Data Base Machines," Large-Scale Systems Directions: Mid-Year Update - 1984.) The analysis which accompanied INPUT's depiction of the large host data base machine resulted in the conclusion that: "The projected (IBM) distributed data base environment under the centralized control of large host processors implies a return to a batch environment." And, the implication of having both extensive batch and interactive transaction processing on central hosts was predicted to require enormous increases in central processing power (as currently recognized by IBM and documented in the last Large-Scale Systems Directions report from INPUT).
  
- It would appear that the only current mainframe functions IBM is interested in distributing are those which are (and have been) so expensive to run on mainframes that users already recognize the need for distribution to other levels of the hierarchy.
  - PPOFS (including electronic mail).
  
  - Financial modeling and analysis (spreadsheets).
  
  - Applications development.






EXHIBIT II-2

DISTRIBUTION OF OPERATING SYSTEMS FUNCTIONS

OS FUNCTIONS	PROCESSORS			
	Mainframes	Mini-computers	PC/ Intelligent Workstations	New Architectures
Process	Multi-Programming	Time-sharing	Interactive	Parallel, Processing
Memory Management	Real (Automated)	Virtual	Real (Manual)	Mechanized
Protection & Security	Centralized	Work-Unit	Personal	None or Device Specific
Scheduling & Resource Management	Priority	Real Time	Manual	Serial
System Structure	Centralized	Distributed, Inter-connected	Standalone, Connected	Connected, Vector, Array, Multi-dimensional
H/F/S	Software, Firmware	Hardware, Software	Hardware, Firmware, Software	Hardware



- There are obvious discrepancies between IBM's apparent three-tiered network and that proposed for so long by INPUT, but it is gratifying to see that some progress is being made. Now that the pot is boiling at all three levels, it is an appropriate time to stand back and try to find some order among the chaos.

## B. STRUCTURED ANALYSIS

- In analyzing the complex systems and products associated with large systems, it has been found desirable, and even necessary, to have some frames of reference. These frames of reference have been established and used in various INPUT reports and have become labeled as "systems categories" which are broken down into sets and subsets. For example, the Software Hierarchy systems category was first established in Market Impacts of IBM Software Strategies, INPUT, 1984, and it consists of the following sets:
  - SNA.
  - Operating systems.
  - DBMS.
  - Languages/decision support systems.
  - Industry turnkey.
  - Applications.
  - Data/information/knowledge.
  - Users.

*[The text on this page is extremely faint and illegible. It appears to be a list or index of items, possibly with columns for descriptions and dates. The text is too blurry to transcribe accurately.]*

- The systems categories and their set lists which will be used in this analysis are contained in Appendix A. And because we find it convenient to think of computer/communications networks as one big "operating system," the first analysis will be done against the Operating System set of the Software Hierarchy systems category.

#### I. OPERATING SYSTEMS FUNCTIONS

- Operating systems have three broad objectives: 1) maximum ease of use, 2) maximum use of equipment (thereby increasing efficiency and reducing the cost per user by sharing resources), and 3) the effective development, testing, and introduction of new system functions without interfering with service. Those same objectives apply to the distribution of processing and data over computer/communications networks (a large-scale computer system can be viewed as a network, and a network can be viewed as a large-scale computer system), and it might be advisable to analyze how and where the major conceptual functions of operating systems fit into the network hierarchy.
- There are five conceptual functions associated with operating systems and a sixth implementation consideration which is especially important to this analysis. These are as follows:
  - Process refers to the abstraction of the program being executed and involves concepts, some familiar and some emerging, such as:
    - Multiprogramming, multiprocessing, parallel processing, etc.
    - Real time transaction systems, timesharing, and other interactive systems.
    - Many problems of synchronization, deadlock, and scheduling have been solved, but many remain; for example, program



execution on a parallel processor will be substantially different than on a von Neumann machine.

- Memory management which essentially refers to the automatic management of the storage hierarchy—how programs and data are brought together for processing. A prevalent concept of memory management has been virtual storage, but blind acceptance of this "solution" as processor and storage technology advances may lead to unacceptable price-performance of distributed applications.
- Information protection and security which involves both problems of access to data/information/knowledge and also their flow over computer/communications networks. This area is of critical importance, and the technical problems associated with distributed data bases are not really understood, much less solved. At the present time, there seems to be either hysterical over-reaction or malignant neglect, and very little in between.
- Scheduling and resource allocation problems obviously become more complex as processing is distributed geographically and architecturally. Fortunately, there is growing recognition that techniques from queuing theory and operations research can make significant contributions to network performance management.
- System structure is defined as integrating all of the above into a coherent system design. In a network environment, there is little reason to believe that force-fitting diverse hardware and/or software systems after the fact will result in anything approaching a "coherent" system design. In other words, "connectivity" is a poor substitute for proper systems analysis, design, and implementation. VM has emerged over the years as a primary tool for integrating diverse operating environments.






- A number of years ago, INPUT stated that IBM had a potent weapon which it could use against plug-compatible vendors. This weapon was the control of the distribution of operating systems functions over the hardware/firmware/software (H/F/S) hierarchy. IBM, to this point in time, has been quite benign in its H/F/S strategy, but current technological and competitive developments (i.e., data base machines and PC clones) will unquestionably see more proprietary operating systems functions migrate to firmware and hardware (it even makes good technological sense).
- It is possible to reach some general conclusions about operating systems functions as they relate to distributed processors, and these conclusions, in turn, will permit analysis of how processing and data can most effectively be distributed over networks (see Exhibit II-2).
  - The conceptual process most generally associated with each level of hardware processor is as follows:
    - Mainframe operating systems have been built with a multiprogramming (batch) model as the base.
    - Minicomputer operating systems are built starting with the assumption of resource (time) sharing.
    - Personal computer operating systems start with emphasis on interacting with a single human (user).
    - Most new architectures (vector processors, data base machines, etc.) are based on parallel processing.
  - Considering the primary process being supported, the most effective memory management model is normally:



- . Given the size and cost of real memory today, cost-effective performance (in terms of throughput) in a multiprogramming (batch) environment can be enhanced by eliminating the burden of virtual storage and concentrating on the effective management of real memory. (It is beyond the scope of this study to pursue old arguments against the indiscriminate use of virtual storage for memory management, but the burden on mainframes is going to become increasingly apparent with advances in storage and communications technologies.)
- . The classic timesharing environment gave rise to the development of the virtual storage concept for memory management, and it remains appropriate for minicomputers.
- . The personal computer user is intimately involved in real memory management, but while there is room for improvement by automating some aspects of memory management at this level, it would be a mistake to accomplish this at the cost of destroying cost-effective interactive performance. (In other words, the cost of virtual storage operating systems to permit some personal computer users to develop enormous spreadsheets has great potential for severe performance impact--beware the vendors who pedal this "solution.")
- . New processor architectures represent mechanization of the memory management function (although some depend upon substantial front-end help).
- Protection and security are aspects of data and information quality control, and the conclusions reached for those operating systems functions also apply to other quality considerations such as data base synchronization and conflicting information flow (reports). Therefore, the comments apply equally of quality assurance systems.



- . The centralization of protection and security (and quality control) on mainframes is probably the strongest argument which can be made for the continued maintenance of central data bases for other than archival purposes. (In other words, good arguments can be made that it is more cost-effective to distribute most data bases, but the potential for quality deterioration more than offsets the cost savings.)
  - . The "need to know" and the ultimate responsibility for quality control must be established at the work unit level, and security procedures (in terms of access, encryption, etc.) must ultimately be enforced within the work unit.
  - . Trusted employees and physical security are about the only protection and security features available for personal computers, and the microprocessor itself can be used as a tool to crack security at other levels. The problems associated with protection and security at this level cannot be overestimated.
  - . The security and protection associated with vector and array processors is normally and properly left to the front-end processor; data base machines can be designed with specific functions built in or depend on auxiliary processors, but the new architectures associated with artificial intelligence (and expert systems) have little sensitivity for the entire subject--this should be a sobering thought.
- This area requires some additional comments because there are obviously network scheduling and resource allocation problems among the various levels of processors themselves. It is our opinion that these problems are comparable to those which have existed under the covers of large mainframes and are amenable to discovered solutions such as



queuing network theories. The question becomes one of determining which level of processor should perform such network management functions, and it has long been INPUT's opinion that these are proper functions for minicomputers (just as process control in a manufacturing environment is a proper function for minicomputers). Viewing mainframes as just another node on a network obviously runs counter to IBM's network architecture (see Exhibit II-1), but SNA has not been noted for giving proper attention to network management. With those preliminary comments, the following are observations concerning specific processor scheduling and resource management models:

- While scheduling and resource management in a multiprogramming environment can become complicated by considerations of arbitrary priority and cost recovery, it has long been known that throughput can generally be maximized by giving priority to the short jobs and getting them out of the system.
- The timesharing environment is essentially interrupt-driven and real time in nature; scheduling is done to maintain response time at desirable levels and minicomputer operating systems are designed with this in mind. (Maintaining sub-second response time—as some would lead us to believe is desirable—will always be more obtainable and predictable from a minicomputer than it will be from a general purpose mainframe.)
- Users control scheduling and resource allocation on personal computers at the present time, normally deciding when to prepare a document (or file) and when and where to transmit and/or store it; and the alternative to this (having the network dictate what, and how much, the human at the terminal does) will meet with substantial resistance from neo-Luddites.







Faint, illegible text, possibly bleed-through from the reverse side of the page. The text is arranged in several paragraphs and appears to be a formal document or letter.

mainframe, where they tend to perpetuate the "von Neumann bottleneck" by depending upon its operating system.

- The implementation of portions of operating systems in firmware and hardware has come slower than INPUT anticipated when a potential hardware/firmware/software strategy was first suggested nearly ten years ago.
  - IBM has resisted "freezing" any significant portion of their mainframe operating systems in either hardware or firmware. This is probably because of the change which is inherent in general purpose hardware/software systems and also because it is more difficult for potential competitors to shoot at a constantly moving target.
  - Minicomputers have been differentiated for specific operating environments and many functions (such as interrupt handling and polling) have been mechanized in hardware, and with RISC architectures, it is probable that this trend will accelerate.
  - While PC operating systems have been relatively limited in function, there is a tendency to make liberal use of both hardware (boards) and firmware (ROM) to implement operating systems functions.
  - New architectures (such as data base machines or associative memory) are frequently hardware solutions to operating systems problems. (Even relocate hardware in the IBM System 370 can be placed in that category.)
- What would appear to be a "proper" distribution of operating systems functions leaves large-scale mainframes stripped down to their essentials-- batch processing effectively overlapped with access to multiple, large data



bases. It is conservatively estimated that a specialized batch processing operating system could cut the cost of such processing by 80 to 90%. However, there is one compelling argument for IBM's highly centralized operating systems strategy--that is quality control, and it is a far from insignificant consideration.

## 2. QUALITY

- The quality systems category contains sets such as auditability, measurement, validity/reliability/predictability, data/information/knowledge, and (as mentioned previously) security/privacy. It is substantially easier to control and assure quality in a highly centralized environment than it is when processing and data bases are distributed, and INPUT has emphasized the potential deterioration of data/information/knowledge quality in the DSD environment.
- However, there is one element of quality control which is not enhanced in the monolithic operating systems environment and that is feedback loops. Generally speaking, feedback loops should be as "tight" as possible, which merely means they should be as close to the data/information/knowledge source as possible. This has several ramifications which will be briefly mentioned.
  - Some ramifications are organizational in nature and have to do with levels of management and span of control. Discussion of these in any detail is beyond the scope of this study except to state the following: If existing levels of management are bypassed in the feedback loop, there is the potential for all kinds of mischief which can result in enormous amounts of unnecessary work. In other words, having operational data bases on mainframes can lead to work unit management being less well informed than higher management, and this can be counterproductive unless the applications system is properly designed.
  - Editing and error checking should be performed as close to the data/information/knowledge source as possible. This provides a simple guide-



line for the design of both applications systems and systems software and for the distribution of processing. Generally speaking, errors should not be caught in batch runs against mainframe data bases; they should be caught at the desktop or on the departmental processor.

- While it should go without saying, feedback loops should not be broken. This is mentioned because IBM's dual data base strategy currently exhibits this disturbing property.
  - It is possible to extract data from IMS data bases, build DB2 tables, and then distribute data and/or information to other levels in the processing hierarchy.
  - Assuming someone either enhances this information (or finds errors in the data), there is currently no convenient way to update, improve, or correct the IMS source data.
  - This single broken (or nonexistent) feedback loop can potentially negate all of the arguments of centralized quality control--distributed applications must be designed with extreme care in this environment.
- It must also be pointed out that while there are many unanswered quality control problems associated with distributed data bases, there is no question that direct distribution from mainframes to intelligent workstations exacerbates all of these problems. For those inclined to adopt a "two-tiered" approach, it should be pointed out that the additional processing power of the 80386 will do little to solve these problems. In other words, there is still a need for minicomputers.





### 3. NETWORK HEIRARCHY

- For the last ten years, INPUT has been defining a "proper" hierarchical network of mainframes, minicomputers, and intelligent terminals (or workstations) in a consistent manner, and we won't repeat the details here. However, it seems an appropriate time to repeat our definitions of minicomputers and intelligent terminals.
  - A minicomputer has been described as selling for less than \$200,000 for the processor alone.
  - An intelligent terminal has been described as selling for less than \$20,000, including processor and peripherals.
- These definitions, which have always been independent of architecture, have been remarkably accurate in describing the reality of an extremely dynamic environment. For example:
  - The IBM 4361 falls under the \$200,000 line, and the 4381 (the bridge system to mainframe operating systems) falls over the line.
  - The IBM RT 370 workstation in its basic configuration comes in under the \$20,000 level.
  - The primary functions allocated to the various levels in the processing hierarchy have also remained valid over the ten-year period.
- All of this is mentioned not only because we have maintained some semblance of order in the face of rapidly changing terminology, but because there are those today who want to define processor categories (for example, departmental processors) in terms of MIPS. We find this objectionable for several reasons:



- MIPS has been and continues to be an absolutely horrible measure of general purpose mainframe performance in any but the most limited sense.
- MIPS has virtually no meaning when applied to the diverse variety of minicomputers, and this becomes even more misleading (and apparent) when RISC machines are introduced.
- The MIPS-based definitions have an extremely short "half life," becoming obsolete in a very short period of time. The fact that a microprocessor has a MIPS rating as high as the largest commercial mainframe of ten years ago does not mean we should have to invent new terminology for either the mainframe or the microprocessor, and it certainly does not mean the effective performance of the two is comparable.

#### 4. SOFTWARE HIERARCHY

- The software hierarchy systems category was presented earlier in the report, and a model for the distribution of operating systems functions over a hierarchical network has been defined. In addition, considerations associated with the quality systems category have been discussed, and INPUT's definition of the network hierarchy has been reaffirmed. It is now possible to analyze briefly IBM's representation of the end-user software environment (see Exhibit II-1) against other sets in the software hierarchy systems category
- It is difficult to argue with most of the software assigned to the intelligent workstation by IBM (dialog manager, PROFS, integrated decision support tools, application development tools, word processor, and instructional system) since these are the highly interactive functions which the personal computer revolution has already wrested from the mainframe. The only matter of real concern for the IWS is the last item, "communications and file transfer to host," which really gets to the heart of the problem of distributed

1870

...

...

...

...

...

...

...

...

data bases because it clearly bypasses the departmental processor for work unit operations and quality control. This clearly demonstrates IBM's fundamental dedication to a two-tiered view of the world--PCs communicating, for the most part, directly with mainframes.

- At the departmental processor level of the IBM network hierarchy, four of the functions (PROFS, electronic mail, document library, and project library) seem to be properly assigned, but "departmental data base" is sufficiently vague to have little meaning unless taken in conjunction with "financial modeling and analysis" and "data extraction program." It then becomes clear that the departmental data base consists of a DB2 planning data base which has been extracted from the host IMS data base. In other words, it serves as a buffer between the user of spreadsheets at the terminal and the complexity of IBM's dual data base strategy on the host. IBM's assignment of functions to departmental processors leaves minicomputers a "weak force" in the network cosmology, subject to effective absorption from both above and below.
  
- At the central host level, IBM's highly centralized approach is quite clearly stated--operational data bases, operational applications, and network management reside on mainframes (along with a central library which is certainly in the right location). If IBM's strategy for the top three sets of the software hierarchy (SNA, Operating Systems, and DBMS) is accepted, the following comments can be made concerning the remaining sets of that systems category.
  - Languages/decision support systems will become extremely complicated because of the necessity of dealing with multiple operating systems, DBMSs, and languages associated with communicating (transferring data) over the loosely connected network which IBM seems to be developing even as geographic and architectural distribution of processing continue to chip away at large host mainframes.



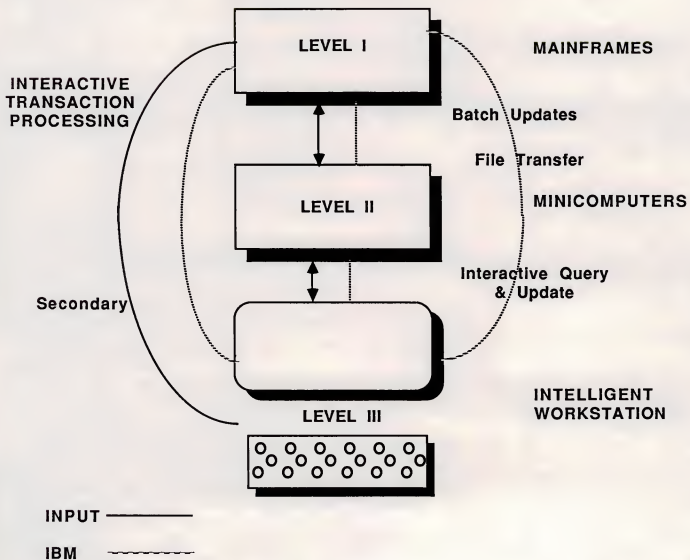
- Industry turnkey systems will be extremely difficult to develop on a cost-effective basis (by either IBM or others) if operational data bases and network management remain centralized on mainframes.
- Applications will be impossible to distribute over the network on an "orderly" and cost-effective basis.
- Data/information/knowledge flow will be inhibited if data bases and network management remain centralized and therefore directed through central hosts. The primary possible virtue of IBM's highly centralized approach--D/I/K quality control--has also been brought into question by the analysis associated with this study. Specifically, IBM's seeming willingness to support departmental processors while encouraging IWS communications directly with the host tends to complicate an already difficult quality problem. (Problems specifically associated with D/I/K flow and quality control will be analyzed in more detail in a later large-scale systems report.)
- Users are viewed as part of the software hierarchy, and there is little indication that IBM's networking, operating systems, and DBMS strategies address adequately the first general objective of operating systems--ease of use. This in turn continues to encourage distributed systems development (bottom-up design) with all of its adverse impacts on quality.
- In summary, IBM seems to accept a three-tiered approach for purposes of office automation, but interactive processing against operational data bases remains centralized on the host. INPUT's "proper" hierarchical network was designed with operational data bases distributed to minicomputers located in work units. The difference between IBM's preferred data base distribution and INPUT's becomes clear with a simple data flow diagram (see Exhibit II-3).





EXHIBIT II-3

DATA FLOW DIAGRAM





- INPUT's hierarchical network has most interactive transaction with processing occurring between Levels II and III whereas IBM links Levels I and III directly for interactive processing. The secondary interactive link to the mainframe in the INPUT hierarchy represents exceptional conditions, such as:
  - A transaction against another branch of a bank.
  - Transactions against highly secured files (executive compensation, strategic plans, etc.) which normally would not be distributed to the individual work units.
  - A programmer maintaining an "old dog" in the central library.
  - If data bases are properly distributed to work units, interactivity with central data bases should seldom exceed 10% of total transactions.
  
- The INPUT model assumes fully edited transactions (or record replacements) will be batched for transmission and updating of both central and distributed data bases as required (many operational data bases are periodic and do not require real time update). IBM's DB2 model for distributing data bases currently makes no provision for updating of the data bases from which data are extracted.
  
- In addition, by encouraging (or at least permitting) file transfer directly between Levels I and III, the IBM model is at risk of compounding data quality problems which INPUT has associated with the DSD environment--data base synchronization and integrity, privacy and security, and conflicting reports to management.
  
- It can be seen that in order for the INPUT model to work effectively, it is necessary for primary operating systems functions to be distributed (or



emphasized) at various levels of the network. The minicomputer assumes a central role in scheduling and resource management of the network, and the mainframe becomes the focus for quality assurance by providing secure data bases of record for the corporation.

- INPUT considers that IBM has little intention of abandoning its highly centralized mainframe approach toward network, operating systems, and data base management in the foreseeable future. It is possible that this approach can be justified based on considerations of quality control, but it must provide for the orderly distribution of processing as technology warrants it. Otherwise, "connectivity" after the fact merely exacerbates an already acute problem.

### C. POTENTIAL RESIDUAL COSTS

- The potential residual costs associated with IBM's apparent large-scale systems directions were identified in INPUT's last Large-Scale Systems Directions report. Essentially, the primary exposure of following IBM's direction is associated with hardware/software expense, and the primary exposure of not following IBM is the risk involved in data/information/knowledge quality (including the development of unworkable distributed systems--evolution is safer than revolution).
- Recently, there have been published figures on the potential cost savings associated with a particular data base machine which states that for a complex transaction the cost was \$5 on the data base machine and \$95 when using DB2 on an IBM mainframe. While this is obviously a best case example and there is no indication of how the cost figures were derived, it certainly is not beyond reason that a data base machine could process an individual transaction (especially one involving JOINS of large DB2 tables) at one-tenth the cost of DB2 on a mainframe. In fact, it is also probable that by distributing relational data bases from mainframes to Levels II and III of the processing



hierarchy (thereby reducing their size), it would probably be possible to demonstrate similar cost savings for individual transactions.

- Therefore, it would appear possible to "save" 90% of the cost of transaction processing against mainframe relational data bases by either geographic or architectural distribution of processing. However, if one has already adopted IBM's centralized approach for operational data bases, and relational tables are being built from extracts from IMS data bases, there are residual costs associated with the mainframe processing which will continue indefinitely--perhaps even beyond the "30 year commitment" IBM associates with a DBMS strategy (see Large-Scale Systems Directions: Disks, Tapes, and Printers, INPUT, 1986).
- "Solutions" after the fact seldom fully compensate for the residual costs of applications systems which are poorly designed in the first place. This is especially true when the applications depend upon operating systems and DBMSs with potential structural deficiencies in terms of data flow. And, this potential is inherent between both MVS/XA and VM, and IMS and DB2. Applications systems (and data bases) must be designed with extreme care in this environment if residual costs are to be minimized.
- The "disorderly" distribution of data to intelligent workstations is tantamount to encouraging bottoms-up applications systems design based on personal data bases. Centralized control of such distribution remains of primary importance because the later integration of data and information generated by these systems will be not only costly, but in many cases impossible. IS management remains responsible for the quality of information systems and underlying data regardless of whether vendor systems software enables, or even encourages misuse. That responsibility extends to the flow of data/information/knowledge over computer/communications networks.





### III RESIDUAL VALUE FORECASTS

#### A. REVIEW OF USED MARKET PRICES VERSUS PROJECTED RESIDUAL VALUES

- INPUT has been projecting residual values for nearly ten years, and periodically we have published spot checks of our past performance. We have just completed a more comprehensive research project on all of our published forecasts since 1981. The results of this research are presented in Exhibits III-1 through III-14. The following comments will help in understanding them:
  - The forecasts were made in the year indicated and are the projected used market retail value as of January 1 of the following years.
  - Normally, such forecasts are presented on a "high-expected-low" basis just as they are now, but for some reason which escapes us, only "expected" figures were projected for some systems during 1981 and 1982.
  - The actual used market retail prices used to check against the forecasts occur at various times during the year based upon when reports were published. (In other words, the actuals always occur later in the specified year.)



## EXHIBIT III-1

**PROJECTED VERSUS ACTUAL  
USED MARKET RETAIL VALUES\*  
IBM 3033 PROCESSOR**

FORECAST YEAR	RANGE	Used Market Retail Value as of January 1				
		1982	1983	1984	1985	1986
1981	High	0.60	0.45	0.27	0.18	0.12
	Expected	0.55	0.40	0.21	0.11	0.03
	Low	0.51	0.34	0.14	0.04	0.01
1982	High					
	Expected		0.2	0.12	0.07	0.04
	Low					
1983	High			0.22	0.18	0.14
	Expected			0.20	0.12	0.08
	Low			0.15	0.08	0.06
1984	High				-	-
	Expected				0.03	0.01
	Low				-	-
1985	High					-
	Expected					0
	Low					-
Actual		0.59	0.3	0.12	0.02	0

\* Percent of Vendor List Price



## EXHIBIT III-2

**PROJECTED VERSUS ACTUAL  
USED MARKET RETAIL VALUES\*  
IBM 4331-GrpII PROCESSOR**

FORECAST YEAR	RANGE	Used Market Retail Value as of January 1				
		1982	1983	1984	1985	1986
1981	High Expected Low	0.85	0.70	0.65	0.45	0.23
1982	High Expected Low		0.64	0.55	0.42	0.30
1983	High Expected Low			0.68 0.62 0.55	0.72 0.65 0.48	0.55 0.44 0.29
1984	High Expected Low				-- 0.62 --	0.50 0.45 0.37
1985	High Expected Low					0.32 0.30 0.18
Actual		0.82	0.65	0.71	0.45	0.26

\* Percent of Vendor List Price



## EXHIBIT III-3

**PROJECTED VERSUS ACTUAL  
USED MARKET RETAIL VALUES\*  
IBM 4341-GrpII PROCESSOR**

FORECAST YEAR	RANGE	Used Market Retail Value as of January 1				
		1982	1983	1984	1985	1986
1981	High					
	Expected	0.88	0.76	0.57	0.50	0.27
	Low					
1982	High					
	Expected		0.75	0.61	0.50	0.32
	Low					
1983	High			0.64	0.54	0.38
	Expected			0.60	0.48	0.35
	Low			0.52	0.35	0.22
1984	High				-	0.18
	Expected				0.21	0.15
	Low				-	0.11
1985	High					0.13
	Expected					0.11
	Low					0.08
Actual		0.81	0.79	0.58	0.19	0.12

\* Percent of Vendor List Price





## EXHIBIT III-4

**PROJECTED VERSUS ACTUAL  
USED MARKET RETAIL VALUES\*  
IBM 4361 PROCESSOR**

FORECAST YEAR	RANGE	Used Market Retail Value as of January 1				
		1982	1983	1984	1985	1986
1981	High					
	Expected					
	Low					
1982	High					
	Expected					
	Low					
1983	High			1.00	0.88	0.72
	Expected			1.00	0.80	0.65
	Low			1.00	0.68	0.57
1984	High				—	0.68
	Expected				0.80	0.62
	Low				—	0.55
1985	High					0.57
	Expected					0.51
	Low					0.42
Actual				1.00	0.77	0.55

\* Percent of Vendor List Price

The first part of the document discusses the importance of maintaining accurate records of all transactions. It emphasizes that every receipt, invoice, and bill should be properly filed and indexed for easy retrieval. This not only helps in tracking expenses but also ensures compliance with tax regulations. The text further explains how digital record-keeping can streamline the process and reduce the risk of errors or loss of documents.



EXHIBIT III-5

PROJECTED VERSUS ACTUAL  
USED MARKET RETAIL VALUES\*  
IBM 4381 PROCESSOR

FORECAST YEAR	RANGE	Used Market Retail Value as of January 1				
		1982	1983	1984	1985	1986
1981	High					
	Expected					
	Low					
1982	High					
	Expected					
	Low					
1983	High			1.10	0.90	0.78
	Expected			1.10	0.83	0.70
	Low			1.10	0.75	0.57
1984	High				--	0.83
	Expected				0.86	0.81
	Low				--	0.72
1985	High					0.79
	Expected					0.74
	Low					0.65
Actual				1.00	0.90	0.85

\* Percent of Vendor List Price



## EXHIBIT III-6

**PROJECTED VERSUS ACTUAL  
USED MARKET RETAIL VALUES\*  
IBM 3083-B PROCESSOR**

FORECAST YEAR	RANGE	Used Market Retail Value as of January 1				
		1982	1983	1984	1985	1986
1981	High Expected Low					
1982	High Expected Low	1.00		0.80	0.55	0.32
1983	High Expected Low			-- 0.9 --	-- 0.78 --	-- 0.65 --
1984	High Expected Low				-- -- --	-- -- --
1985	High Expected Low					-- -- --
Actual		1.00		0.90	0.68	0.36

\* Percent of Vendor List Price



**PROJECTED VERSUS ACTUAL  
USED MARKET RETAIL VALUES\*  
IBM 3081-G PROCESSOR**

		Used Market Retail Value as of January 1				
FORECAST YEAR	RANGE	1982	1983	1984	1985	1986
1981	High	1.00	0.92	0.83	0.65	0.52
	Expected	1.00	0.87	0.78	0.57	0.40
	Low	1.00	0.85	0.73	0.49	0.32
1982	High					
	Expected		0.90	0.80	0.55	0.40
	Low					
1983	High			0.87	0.75	0.57
	Expected			0.85	0.68	0.50
	Low			0.78	0.60	0.45
1984	High				--	--
	Expected				--	--
	Low				--	--
1985	High					--
	Expected					--
	Low					--
Actual		1.00	0.90	0.80	0.68	0.39

\* Percent of Vendor List Price





EXHIBIT III-8

PROJECTED VERSUS ACTUAL  
USED MARKET RETAIL VALUES\*  
1BM 3081-K PROCESSOR

FORECAST YEAR	RANGE	Used Market Retail Value as of January 1				
		1982	1983	1984	1985	1986
1981	High					
	Expected					
	Low					
1982	High					
	Expected		0.90	0.85	0.58	0.45
	Low					
1983	High			0.92	0.82	0.68
	Expected			0.88	0.75	0.60
	Low			0.80	0.63	0.45
1984	High				--	--
	Expected				--	--
	Low				--	--
1985	High					--
	Expected					--
	Low					--
Actual		0.94	0.85	0.76	0.44	

\* Percent of Vendor List Price



## EXHIBIT III-9

**PROJECTED VERSUS ACTUAL  
USED MARKET RETAIL VALUES\*  
IBM 3081-KX PROCESSOR**

FORECAST YEAR	RANGE	Used Market Retail Value as of January 1				
		1982	1983	1984	1985	1986
1981	High Expected Low					
1982	High Expected Low					
1983	High Expected Low					
1984	High Expected Low				-- 0.85 --	0.75 0.68 0.60
1985	High Expected Low					0.55 0.51 0.44
Actual					0.9	0.54

\*Percent of Vendor List Price



## EXHIBIT III-10

**PROJECTED VERSUS ACTUAL  
USED MARKET RETAIL VALUES\*  
IBM 3350-A02 DISK**

FORECAST YEAR	RANGE	Used Market Retail Value as of January 1				
		1982	1983	1984	1985	1986
1981	High	1.00	0.86	0.56	0.43	0.35
	Expected					
	Low					
1982	High		0.52	0.40	0.28	0.20
	Expected					
	Low					
1983	High			0.38	0.30	0.20
	Expected			0.34	0.26	0.14
	Low			0.29	0.19	0.08
1984	High				0.22	0.14
	Expected				0.20	0.12
	Low				0.12	0.07
1985	High					0.08
	Expected					0.05
	Low					0.03
Actual		1.00	0.52	0.25	0.17	0.05

\* Percent of Vendor List Price



## EXHIBIT III-11

**PROJECTED VERSUS ACTUAL  
USED MARKET RETAIL VALUES\*  
IBM 3380-AA4 DISK**

FORECAST YEAR	RANGE	Used Market Retail Value as of January 1				
		1982	1983	1984	1985	1986
1981	High					
	Expected	>1.00	1.00	0.99	0.84	0.69
	Low					
1982	High					
	Expected		0.85	0.74	0.67	0.50
	Low					
1983	High			0.84	0.78	0.62
	Expected			0.78	0.70	0.52
	Low			0.74	0.65	0.46
1984	High				0.85	0.65
	Expected				0.80	0.53
	Low				0.68	0.45
1985	High					0.73
	Expected					0.68
	Low					0.55
Actual			1.03	0.95	0.87	0.48

\* Percent of Vendor List Price





EXHIBIT III-12

PROJECTED VERSUS ACTUAL  
USED MARKET RETAIL VALUES\*  
IBM 3420-005 TAPE DRIVE

FORECAST YEAR	RANGE	Used Market Retail Value as of January 1				
		1982	1983	1984	1985	1986
1981	High Expected Low	0.36	0.28	0.14	0.10	0.08
1982	High Expected Low		0.11	0.07	0.05	0.03
1983	High Expected Low			0.10 0.08 0.06	0.07 0.06 0.04	0.05 0.04 0.02
1984	High Expected Low				0.10 0.07 0.05	0.08 0.05 0.03
1985	High Expected Low					0.16 0.12 0.08
Actual		0.25	0.10	0.12	0.19	0.05

\* Percent of Vendor List Price



EXHIBIT III-13

PROJECTED VERSUS ACTUAL  
USED MARKET RETAIL VALUES\*  
IBM 3420-008 TAPE DRIVE

FORECAST YEAR	RANGE	Used Market Retail Value as of January 1				
		1982	1983	1984	1985	1986
1981	High	0.77	0.71	0.49	0.40	0.36
	Expected					
	Low					
1982	High		0.72	0.57	0.45	0.30
	Expected					
	Low					
1983	High			0.48	0.40	0.32
	Expected					
	Low					
1984	High			0.42	0.33	0.24
	Expected					
	Low					
1985	High			0.32	0.25	0.18
	Expected					
	Low					
1984	High				0.77	0.62
	Expected					
	Low					
1985	High				0.65	0.54
	Expected					
	Low					
1985	High				0.58	0.42
	Expected					
	Low					
1985	High					0.88
	Expected					
	Low					
Actual	High	0.83	0.69	0.93	0.93	0.45
	Expected					
	Low					

\* Percent of Vendor List Price



EXHIBIT III-14

PROJECTED VERSUS ACTUAL  
USED MARKET RETAIL VALUES\*  
IBM 3800-001 PRINTER

FORECAST YEAR	RANGE	Used Market Retail Value as of January 1				
		1982	1983	1984	1985	1986
1981	High	0.64	0.57	0.47	0.33	0.14
	Expected					
	Low					
1982	High		0.58	0.50	0.35	0.20
	Expected					
	Low					
1983	High			0.60	0.57	0.50
	Expected					
	Low					
1984	High			0.55	0.50	0.44
	Expected					
	Low					
1985	High			0.48	0.42	0.35
	Expected					
	Low					
1986	High				0.55	0.47
	Expected					
	Low					
1987	High				0.52	0.44
	Expected					
	Low					
1988	High				0.46	0.38
	Expected					
	Low					
1989	High					--
	Expected					
	Low					
1990	High					0.36
	Expected					
	Low					
Actual	High	0.65	0.63	0.58	0.45	0.47
	Expected					
	Low					

\* Percent of Vendor List Price



- The comparisons against these used market retail prices were made against vendor list prices current at that time. (All INPUT residual value forecasts incorporate proprietary assumptions concerning vendor price adjustments.)
- Since the research was completed just prior to publication of this report, detailed analysis will not be published until the next large-scale systems report in the fourth quarter of this year. However, we are quite pleased with the results based on preliminary analysis. For example, the 1981 forecast for the 3081-G (see Exhibit III-7) was quite remarkable, and processor forecasts in general were excellent.
- Where there is significant variance, there are usually rational and even interesting explanations. The 3420-008 tape drive (see Exhibit III-13) is a fascinating case study.
  - The 1981 forecast for 1986 is excellent, but what happened in between is a little crazy unless one remembers IBM's problems with the 3480 drives (as chronicled by INPUT in its reports).
    - First there were continued delays in the announcement of the 3480 for various reasons.
    - Then, after announcement, there was some reluctance on the part of some customers to switch to the new technology.
    - Delivery schedules and customer acceptance of the new drives meant that the 3420-008 remained the workhorse for DASD backup and few entered the used market.
  - The result was that used market values soared during 1984 and 1985, and then suddenly dropped this year.





- Unfortunately, INPUT got mouse-trapped as it adjusted its numbers upwards, and our 1984 forecast (made in the first quarter of that year) did not anticipate the sudden drop in 1986. However, it is still possible that the 3420-008 used market may recover later this year, and the purpose of our analysis will be to determine whether such anomalies can be anticipated.

## B. ANNOUNCEMENTS

- The only announcement of significance since the last report was published is the NAS AS/XL Vector series of processors. The new processors range in price from \$3.5 million up to approximately \$15.0 million and are designed to compete against the IBM 3090 Vector Facility and Amdahl's 5890 mainframes. Delivery dates for various models are scattered over the next year, and in the fourth quarter of 1987 NAS plans to deliver a proprietary firmware feature which will give AS/XL Vector users the ability to run programs compiled for the 3090 Vector Facility. The announcement is in further support of INPUT's long standing projection of distributed architectures which relieve the "von Neumann bottleneck."
- Without other large-scale announcements of significance, the trade press spent considerable time speculating about additional price reductions and suggesting that IBM's full employment policy was the culprit in lower earnings which were reported for the first two quarters of this year. INPUT has the following comments on this gossip (it can hardly be termed analysis).
  - While both IBM 3090 and 3380E had a substantial price cushion built in at announcement time, price reductions earlier this year were probably ill-advised if they were intended to stimulate sales (as opposed to normal adjustments based on competitive developments over the last year). It is difficult to imagine that IBM believed a price reduction



would stimulate enough additional sales to offset lost revenue from the price reduction in the first place, but to conclude that after shooting itself in the foot, it will now shoot itself in the head is beyond belief-- IBM is not infallible, but it seldom makes the same mistake twice. In fact, some IBM representatives have been stating that price/performance improvement will be achieved more through improved performance rather than through price reductions (reversal of a 70/30% ratio which gave emphasis to price has been indicated) in the future.

- As far as the wailing of the analysts in the investment banking community is concerned, it must be pointed out that while IBM profit margins dipped from 13.1% in 1985 to 10.4% in the first half of 1986, the latter figure would still have placed IBM first among the top 50 companies in the data processing industry (and within the top ten of the top 100 companies). In addition, IBM management is unquestionably aware that when the 308X series was first shipped in 1981, IBM's profit margin (of 11.4%) was the lowest since the early 1950s, so a drop this year probably was not unexpected. As far as the gratuitous advice of external analysts is concerned, IBM's full employment policy remains more important to the corporation than recent fluctuations in earnings and it is unlikely to be abandoned.
- There are indications that IBM will make some announcements in September or early October. These announcements will be designed to take advantage of the function and technology already embedded in the 3090 processors (see INPUT's last Large-Scale Systems Directions report). The most likely bet is a combination of higher speed channels (6 meg) and a new version of MVS which will take advantage of expanded storage. Throughput improvement of approximately 15-20% will be claimed for the combination.



### C. ASSUMPTIONS

- INPUT's residual value forecasting methodology is proprietary and has been continually refined over the years. The assumptions underlying our forecasts fall into three categories--general, specific, and proprietary.
  
- The general assumptions underlying INPUT's residual value forecasts are as follows:
  - IBM is always operating against a plan which will maintain its traditional growth in revenues and its traditional profit margins.
  
  - IBM is essentially large-scale systems-oriented and will resist significant offloading of mainframes to minicomputers and/or micro-processors.
  
  - The means of control of the distribution of processing and data is through systems software (SNA, operating systems, and DBMSs), and software development will always lag hardware capability.
  
  - IBM will continue to be successful in controlling the distribution of processing because there will be no serious breakthroughs in competitive systems software development which IBM cannot effectively counter.
  
  - Large-scale hardware/software will continue to evolve pretty much on IBM's schedule, and there will not be any drastic changes in product cycles (a majority of customers are not going to decide to skip a generation).
  
  - Mainframes and associated peripherals will remain at the heart of IBM's strategy through the 1990s, and there will be a continuing used market for such equipment during that period.



- IBM has the administrative systems in place to facilitate product announcements and pricing changes which virtually give it control of residual values. (IBM's increased flexibility in both product introduction and pricing have become apparent over the years, and the importance of these improved internal systems should not be underestimated.)
- There are certain specific assumptions which are directly related to current residual value forecasting. These assumptions are as follows:
  - IBM will not deviate radically from historic patterns of price-performance improvement for large-scale processors and magnetic disk storage systems. (INPUT identified these patterns over ten years ago, and they have proven to be remarkably accurate.)
  - Therefore, it is assumed that price-performance will improve at a rate of between 10% and 16% per year (depending upon the particular product), and these rates are used to compensate for list price reductions over the product life cycle. (The specific methodology employed is proprietary.)
  - IBM will be able to delay the impact of optical memories upon large-scale magnetic disk beyond the range of this year's forecasts (1991).
  - Modest impact of optical memories upon large-scale tape systems will begin to be felt during this period, and this impact is built into the forecasts.
  - IBM is assumed to have been reasonably forthright in its large-scale systems presentation, as presented in our last report, and there do not appear to be any competitive technological developments which will force premature (from IBM's point of view) deviation from the highly





centralized architectural focus which has been described in this series of reports.

- As predicted earlier, IBM will start to release new versions of MVS/XA which will support the 3090 architecture (and any enhancements such as new channels). It is assumed that this announcement will be a clear signal that MVS/XA for the 308X has reached the end of the line.
- Alternate operating systems, such as UNIX and Amdahl's Aspen, will neither have serious impact on the 3090 sales nor serve to extend the life of 308X processors.
- IBM's dual data base strategy will prove successful—DB2 will become highly popular (and a de facto standard), and IMS will live on well past the forecast period. As data bases are "distributed," demand for archival storage on mainframes will continue to represent a strong growth area for magnetic disks.
- Privacy and security is going to become an increasingly important subject during the next five years, and the IBM solution is going to obsolete a lot of current hardware and software at all levels in the processing hierarchy. Secure, certified data bases are the key to large-scale systems growth in the 1990s.
- Large-scale printer technology is virtually frozen, and new technological developments will be concentrated on distributed printer systems. However, centralized printing facilities will remain viable and necessary during the forecast period.



#### D. PROJECTED USED MARKET PRICES AND RESIDUAL VALUES

- The last report, Large-Scale Systems Directions: Disks, Tapes, and Printers, contained a comprehensive set of residual value forecasts for IBM peripherals and smaller mainframes. In addition, it reviewed used market activity and secondary market prices for selected Amdahl and NAS equipment. Since that report covered activity through the first half of the year, few changes have occurred in the used market. Therefore, this mid-year update will address the high end of the IBM mainframes (the 3083 will also be presented since new projections have been made since the last report). Amdahl and NAS mainframes have not experienced significant changes since the Large-Scale Systems Directions: Large IBM and Software-Compatible Mainframes report was prepared late in 1985.
- Exhibits III-15 and III-16 present the 1987-1991 projected residual values as a percent of vendor list price and projected used market retail values based on those residual values.
- Exhibit III-17 presents the range of anticipated values for selected IBM processors which have changed since the last Large-Scale System Directions: Large IBM and Software-Compatible Mainframes report was published in the fourth quarter of 1985.



## EXHIBIT III-15

PROJECTED RESIDUAL VALUE AS A PERCENT OF  
VENDOR LIST PRICE

PROCESSOR	Current List	-	'87	'88	'89	'90	'91
3083-CX	\$802,731		15%	8%	3%	1%	1%
3083-E	1,157,731		8%	4%	1%	--	--
3083-EX	852,731		13%	6%	2%	--	--
3083-B	2,032,731		15%	8%	3%	1%	--
3083-BX	1,337,731		20%	10%	5%	2%	--
3083-J	2,527,731		18%	9%	4%	2%	--
3083-JX	1,587,731		25%	12%	6%	3%	2%
3081-G	3,200,731		15%	7%	3%	1%	0%
3081-GX	2,170,731		26%	11%	5%	3%	1%
3081-K	3,940,731		17%	9%	4%	2%	0%
3081-KX	2,680,731		32%	14%	8%	4%	1%
3084-Q	6,091,465		28%	14%	8%	5%	1%
3084-QX	4,876,462		35%	18%	10%	6%	2%
3090-200	4,494,410		90%	78%	57%	27%	8%
3090-400	8,515,785		92%	85%	65%	31%	12%

THE [illegible] OF [illegible]

[illegible]

[illegible]

[illegible]

[illegible]

[illegible]

[illegible]

[illegible]

[illegible]

[illegible]

[illegible]

[illegible]

[illegible]

[illegible]

[illegible]

[illegible]

[illegible]

[illegible]

[illegible]

[illegible]

[illegible]

[illegible]

## EXHIBIT III-16

**PROJECTED USED MARKET RETAIL VALUE AT  
JANUARY 1**

PROCESSOR MODEL	Current List Price	PROJECTED USED MARKET RETAIL VALUE AT JAN. 1 OF:				
		1987	1988	1989	1990	1991
3083-CX	\$802,731	\$120,410	\$64,218	\$24,082	\$8,027	\$8,027
3083-E	1,157,731	92,618	46,309	11,577	0	0
3083-EX	852,731	110,855	51,164	17,055	0	0
3083-B	2,032,731	304,910	162,618	60,982	20,327	0
3083-BX	1,337,731	267,546	133,773	66,887	26,755	0
3083-J	2,527,731	454,992	227,496	101,109	50,555	0
3083-JX	1,587,731	396,933	190,528	95,264	47,632	31,755
3081-G	3,200,731	480,110	224,051	96,022	32,007	0
3081-GX	2,170,731	564,390	238,780	108,537	65,122	21,707
3081-K	3,940,731	669,924	354,666	157,629	78,815	0
3081-KX	2,680,731	857,834	375,302	214,458	107,229	26,807
3084-Q	6,091,465	1,705,610	852,805	487,317	304,573	60,915
3084-QX	4,876,462	1,706,762	877,763	487,646	292,588	97,529
3090-200	4,494,410	4,044,969	3,505,640	2,561,814	1,213,491	359,553
3090-400	8,515,785	7,834,522	7,238,417	5,535,260	2,639,893	1,021,894





EXHIBIT III-17

PROJECTED RESIDUAL VALUE AS A PERCENT OF  
VENDOR LIST PRICE (UPDATED)

PROCESSOR MODEL		1987	1988	1989	1990	1991
3081-6	High	21%	13%	6%	4%	2%
	Expected	15%	7%	3%	1%	0%
	Low	8%	4%	1%	0%	0%
3081-6X	High	33%	20%	10%	5%	3%
	Expected	26%	11%	5%	3%	1%
	Low	14%	6%	2%	1%	0%
3081-K	High	25%	15%	8%	5%	2%
	Expected	17%	9%	4%	2%	0%
	Low	10%	5%	1%	0%	0%
3081-KX	High	40%	21%	14%	9%	5%
	Expected	32%	14%	8%	4%	1%
	Low	18%	9%	4%	2%	0%
3084-Q	High	35%	19%	12%	9%	5%
	Expected	28%	14%	8%	5%	1%
	Low	20%	7%	3%	1%	0%
3084-QX	High	41%	25%	17%	11%	7%
	Expected	35%	18%	10%	6%	2%
	Low	28%	12%	5%	2%	0%
3090-200	High	95%	85%	65%	40%	15%
	Expected	90%	78%	57%	27%	8%
	Low	86%	70%	45%	18%	3%
3090-400	High	96%	90%	65%	40%	25%
	Expected	92%	85%	65%	31%	12%
	Low	88%	74%	52%	21%	5%



Several lines of text are visible at the bottom of the page, but they are extremely faint and difficult to read. The text appears to be organized into a list or a series of entries, possibly a table of contents or a list of items. The lines are roughly horizontal and span across the width of the page.

EXHIBIT III-17 (Cont.)

PROJECTED RESIDUAL VALUE AS A PERCENT OF  
VENDOR LIST PRICE

PROCESSOR MODEL		1987	1988	1989	1990	1991
3083-CX	High	18%	12%	8%	5%	3%
	Expected	15%	8%	3%	1%	1%
	Low	7%	3%	1%	--	--
3083-E	High	11%	7%	3%	2%	1%
	Expected	8%	4%	1%	--	--
	Low	4%	1%	--	--	--
3083-EX	High	15%	10%	5%	3%	1%
	Expected	13%	6%	2%	--	--
	Low	6%	3%	1%	--	--
3083-B	High	18%	11%	7%	4%	2%
	Expected	15%	8%	3%	1%	--
	Low	9%	3%	--	--	--
3083-BX	High	24%	17%	11%	6%	3%
	Expected	20%	10%	5%	2%	--
	Low	12%	5%	2%	--	--
3083-J	High	23%	15%	10%	4%	2%
	Expected	18%	9%	4%	2%	--
	Low	9%	4%	1%	--	--
3083-JX	High	31%	20%	13%	7%	5%
	Expected	25%	12%	6%	3%	2%
	Low	14%	5%	3%	--	1%



## APPENDIX A: SET LISTS USED IN ANALYSIS

### A- GST DIRECTION

- 1- Centralization
- 2- Integration
- 3- Differentiation
- 4- Mechanization

### B- QUALITY

- 1- Objectives
- 2- DIK
- 3- Auditability
- 4- Measurement
- 5- Feedback Loops



6- Validity/Reliability/Predictability

7- Security/Privacy

D- NETWORK HIERARCHY

1- Large Mainframes

2- Minicomputers

3- Intelligent Workstations

4- Terminals

5- Mobile Terminals

E- SOFTWARE HIERARCHY

1- SNA

2- Operating Systems

3- DBMS

4- Languages/DSS

5- Industry Turnkey

6- Applications

7- DIK

8- Users





H- SYSTEMS TYPE

- 1- Batch
- 2- Transaction
- 3- Interactive
- 4- Real Time
- 5- Decision Support
- 6- Expert

I- SYSTEMS REQUIREMENTS

- 1- High/Low Transaction Rates
- 2- High/Low Processing Requirements
- 3- Large/Small Data Base
- 4- High/Low Functionality
- 5- Many/Few Decision Rules
- 6- High/Low Responsiveness

J- USER SET

- 1- Scientific
- 2- Engineering



- 3- Systems/Procedures Analyst
- 4- Programmer
- 5- Clerical/Accounting
- 6- Secretarial
- 7- Administrative
- 8- Executive
- 9- Casual

K- PERFORMANCE

- 1- Hardware/Software
- 2- Human/Machine Dyad
- 3- Work Unit
- 4- Institutional

M- IBM STRATEGIC PERIODS

- 1- SNA/DDP
- 2- Electronic Office
- 3- Expert Systems
- 4- Custom Systems

