Research Board

APR 20 1984

THE RESEARCH BOARD

INCORPORATED

SEP 28 1984

220 EAST 61st STREET · NEW YORK, N.Y. 10021 · (212) 486-9240

Route to: Jack Shields - Cover memo + pgs. 44-46
Ken Olsen - Cover memo - + pgs. 44-46
Jack Smith -
Sam Fuller -
Bill Johnson -

April 11, 1984

Return To: Win Hindle

Dear Win:

As promised, I am enclosing key sections of the current RB Report: Computing and Communicating in the 1990s. The selection describes our visit with Ken Olsen and you as well as mid-range computer sector snapshots and our usual irreverent conclusions. You are fully familiar with our style, so I won't cluck any preliminaries about stark style, insider audience and so forth. Our light-hearted tone should not be a surprise.

You must know how much we admire Digital and recall how loudly we applauded the restructuring of the marketing and sales forces. But Research Board members today remain concerned over your current sales presence (or lack thereof), management losses and DECsystem replacement strategy. At a convenient point, perhaps we should talk some more about Digital's current progress; I leave that judgment entirely to you.

For now, Ernie and I want to thank Ken and you again for candid responses to our immediate questions and for broad insights on long term industry directions. As always, we would value your comments (particularly critical ones) concerning this rather global document. And we look forward to visiting Maynard again.

Sincerely,

Naomi

Naomi O. Seligman

Mr. Winston R. Hindle, Jr.
Vice President, Corporate Operations
Digital Equipment Corporation
146 Main Street
Maynard, Massachusetts 01754

A RIOT OF OPPORTUNITY

Computing and Communicating in the 1990s

Partial Copy

THE RESEARCH BOARD

FOREWORD

This report was prepared solely for the use of the Research Board which financed its completion. It contains information about individual companies which may be considered highly sensitive. For this, and proprietary reasons of our own, distribution should be restricted to members of the Research Board.

EvS/NOS
February 1984

THE RESEARCH BOARD

TABLE OF CONTENTS

I. First Step to the Nineties

Most RB Reports target specific business and technology sectors and look pragmatically at a three or four year horizon. In March 1983, the membership voted for a major study that would traverse a farther horizon and gather broader material: the entire computer and communications industry in the 1990s.

As we recall, your interest in the subject was based on the following considerations. Information systems have already transformed back office operations and are now assuming strategic roles in company management, marketing, manufacturing and distribution. Therefore, senior management demands that computing be fully integrated into the long range planning cycle; that's ten years in many large corporations, seldom less than five. We remain sharply conscious of the strategic framework throughout our tour. Along the way, however, we gather scores of more immediate insights as well. For one thing, we sight new navigational landmarks to help us all steer a truer course around near-term enthusiasms and long-standing blind spots.

The RB membership also responds to our message network inquiries with dozens of ideas and opinions. This report, we think, covers most topics you wished to address: supercomputers, mainframes, microprocessors, minicomputers, personal computers, human factors engineering, image processing, software, speech recognition, disk storage, communications networks, videoconferencing and systems security. Two subjects have clearly been jettisoned: the Japanese, whom we will attack for the autumn report, members willing; and the impact of Federal policy, which we sidestep. Not because the Feds are unimportant, but because Washington's behavior is too erratic for any self-respecting futurist to attempt to predict.

The wide range of topics under investigation used all our methodological muscles. As always, our staff produced reams of technology position papers, literature summaries, symposium analyses and business profiles, each one a fresh immersion into changing technological currents. Endless sifting identified a profusion of companies, concepts and individuals we thought would interest the Research Board. Then we fanned out across the country, interviewing the key figures in the field.

Each interview seemed to revolve around three pivotal issues:

- What competitor, technology or other factor had the greatest impact in your sector during the past ten years? (To consider the future, one must first appreciate the past.)

- What factor will have the greatest impact in the next ten years?

- How will your organization react?

Not surprisingly, all the experts, chief executives, research directors, product managers, academicians and just plain smarties that we questioned had different responses. And each response, of course, suggested new questions.

Special recognition to Catherine Loup for her chapter on communications. Stars and clusters to Audra, Laura, Leah, Nan and Scott who were pressed into service Saturdays, Sundays (or both) to accomplish the firm's insane objectives. And warmest thanks to those we interviewed:

Amdahl	Eugene White, Chairman
Apple Computer	Steve Jobs, Chairman Floyd Kvamme, Exec. VP
AT&T Communications	Richard Jacobsen, Vice President Billy Oliver, Vice President
AT&T Information Services	William Smith, General Manager
Bell Labs	Dr. Ian Ross, President Dr. Daryl Eigen, Human Factors Lab
Bell Northern Research	John Roth, President
Central Services	Irwin Dorros, Exec. VP
Cray Research	John Rollwagen, Chairman
Digital Equipment	Kenneth Olsen, President Winston Hindle, VP Corporate Ops
Commodore	Dr. Richard Wiggins
ETA Systems (Sub. of CDC)	Dr. Lloyd Thorndyke, President
GTE Labs	Fletcher Haselton, Director
Hewlett-Packard	Paul Ely, Exec. VP Dr. Joel Birnbaum, Director of Computer Research

IBM	Erich Bloch, VP and Group Exec. Bob Evans, VP, Engineering, Programming and Technology Dr. Ralph Gomory, VP Watson Labs Jack Kuehler, Sr. VP & Group Exec., Information Systems & Technology Group
Metaphor Computer	Don Massaro, President Dr. David Liddle, VP
MSA	John Imlay, Chairman
National Bureau of Standards	Dr. Dennis Branstad Ray Moore
DHL Courier	Dr. Lawrence Roberts, President
Pacific Telesis	Arthur Latno, Exec VP
SmartCard Intl.	Arlen Lessin, President
Storage Technology	Jesse Aweida, Chairman and CEO Zoltan Herger, Exec VP Juan Rodriguez, Vice President
Synapse Computer	Mark Leslie, President
Tandem Computers	Dr. Gerald Held, Director Planning
Texas Instruments	Dr. George Doddington, Director
Threshold Technology	Gerald Beckmann, President
Trilogy Systems	Carlton Amdahl, Vice Chairman
University of Wisconsin	Dr. Lorne Parker, Director Center for Interactive Programs
Wang Laboratories	Dr. An Wang, Chairman John Cunningham, President Sam Gagliano, VP Product Planning

II. COMPUTING

C. Up Is Down: The New Cost Relationships

Current wunderkinder aside, the ferment over parallelism may not affect the real world of large company data processing for several years. But the ferment reflects a deeper truth of immediate importance. To wit: price/performance relationships among mainframes, minicomputers and microcomputers have been drastically up-ended.

Twenty years ago, Herb Grosch, industry score-keeper (and irrepressible gadfly), advanced the "Law" that computing power exhibits fundamental and irrevocable economies of scale. For twice the bucks, expect four times the bang. Of course, Grosch's Law was formulated in an era when computers of all sizes were manufactured from roughly the same technological materials. No longer. Large uniprocessors today demand tricky bipolar technology, ticklish circuit packaging, complex design automation systems and innovative cooling techniques. Micros get by with cookie-cutter CMOS chips, conventional laminates and little fans.

For large mainframes, the technological challenge and cost pressures will probably worsen, at least if Carlton Amdahl is correct about the 40/50 MIPS wall for uniprocessors. All the exploration, invention, investment and sheer exotica needed to smash that wall will raise product cost further, notwith-standing efforts to improve manufacturing and design economies.

Such technological and cost discontinuities tend to disrupt Herb Grosch's smooth and extrapolated correlations. Economies of scale can't be promised these days. In fact, current observation distinctly suggests a diseconomy of scale. Today, prices run at about $285,000 per MIP for the IBM 3084; $190,000 for the IBM 4361; $170,000 for the Hewlett-Packard 3000; and $150,000 for the Digital VAX 780. That implies 50 percent premiums for "Big MIPS," and another 10 percent for the IBM logo, under even the most conservative estimates.

For another technological discontinuity, add microprocessor KIPS to the picture. After all, a single chip represents the micro engine whereas the mini needs many intersoldered components. Now watch that price point plunge! The industry price for a bare-bones personal computer with 100 KIPS is $5000; scale it up, and the equivalent per MIP cost is just $50,000, one-quarter the mini MIP, one-sixth the mainframe MIP.

Taken at face value, these price relationships help explain the ardent interest by entrepreneurs and academics in transforming strings of micros into mainframes. Successful parallelism would revive the comfortable economics of Grosch's Law, permit an inexpensive commodity technology to span the capacity range from single microprocessors at the bottom to sizeable gangs at the top.

Of course, processor technology is hardly the only consideration in determining economies of scale or configuration pricing. Several other factors are relevant: the competitive climate, architectural options, memory and storage costs, software and communications pricing, equipment utilization rates and operating expense. We will skim the prognosis for each over the next decade.

1. Competition: IBM has no effective mainframe competition among U.S. suppliers. Therefore, it will probably continue to garner most of its profits from the high end. The Bunch companies (Burroughs et al.) are fading from the mainframe scene. If Trilogy survives, its product strategy will focus on the 30 MIP (and up) market. So the primary mainframe competitors will be Japanese: Hitachi through National Advanced Systems, Nippon Electric through Honeywell, Fujitsu (at some level of integration) through Amdahl. Amidst such total dyspepsia, the mini and micro markets seem to supply a refreshingly price-competitive fight.

2. Processor Architecture: Although we've just observed the disadvantageous price gap between one-chip micros and multi-chip minis, it's too early for handkerchiefs. The mini may recover. Several vendors are within one or two years of announcing a replacement for their whole mid-range product lines in the form of a "reduced instruction set" architecture (RISC). (See RB Report: Roots and Rhythms of the Future.) Digital's West Coast research group is also known to be recruiting RISC experts.

The new architecture has several unique characteristics: First, it assumes that a simple instruction set can be designed to perform common instructions in fewer machine cycles than traditional architectures. Second, it assumes these simpler instruction sets can be designed more rapidly and inexpensively. Third, they can be more efficiently manufactured - particularly if the entire processor fits on a single chip. All this, plus massive amounts of main memory and (possibly) a pipelining feature could start the mini ball rolling again. In fact, a 10:1 price/performance leap is projected!

That shocker would heighten the prospect of an interesting new mainframe/mini war. Begin by not disputing Amdahl CEO Gene White's optimistic forecast of doubling price/performance every 3.5 years, reaching a fourfold improvement around 1990. Mainframe prices would then decline from $280,000 per MIP to $70,000. Not bad. But, mini MIPS would simultaneously plummet from $200,000 to $20,000 - or less. Thus, the "Big MIPS" premium would no longer be 50 percent but 250 percent. Quite a gap.

Of course, the next mainframe generation may well have other advantages, like a good working relationship with existing applications software. RISC minis could require some applications rework to achieve their native performance potential. So, that 250 percent price premium for Big MIPS could be more aptly viewed as the customers' penalty for avoiding the even more punishing software conversion. In any event, IBM's mainframes could be faced with new and healthy competition after all. That would certainly rejiggle our Groschian trend lines.

3. Memory Technology: Processors aren't the only component of a computer configuration, we're pleasantly reminded by Watson Labs Director Ralph Gomory. Memory and disk storage utilize technologies which demonstrate economies of scale, and favor the glass house configuration. For main memory, Groschian principles surely triumph, since the dominant technology is common to processors of all sizes.

In big disk, there are even clearer economies of scale. Consider that one megabyte of IBM 3380 storage is only $40, compared to a mini megabyte of hard disk at $55 and an Apple (5-inch) at $400. In the future, the glass house bonus may become even more pronounced - particularly for high capacity (not high performance) storage. The cream of high capacity/ low cost per megabyte disk may eventually be optical. And the underlying economies of read/write optics technology also favor the glass room, as we will soon discuss.

4. Software Pricing: Trends in the software business tilt somewhat against economies of scale. Mainframe software has to be priced to compensate for several intrinsically expensive factors; labor-intensive marketing, substantial implementation support and a relatively small customer set from which to recover development investment. But in the "no-hands" mass market for personal computer software, prices can be set solely to encourage maximum volumes. Both Hewlett-Packard's Paul Ely and Apple's Mike Markkula expect that PC software prices will decline eventually to levels approaching media costs, like phonograph records. On a similar trend line, Wang executives note: "Ten years ago, we sold an accounting application module for $8,000; today, the same amount of function on a PC costs $79." Of course, this bird's eye view ignores the fact that PC customers are often forced to buy multiple licenses to support multiple users. But the tilt still exists, mostly because of the inherent inefficiencies in mainframe software distribution noted earlier.

5. Communications Cost Trends: In the late seventies, the twin issues of communications cost and communications performance encouraged distribution of many applications away from central mainframes to local minis. Star networks were increasingly handicapped by higher private line tariffs and inadequate bandwidth (which dampened response time at remote terminals). Expect a pirouette in the future: sharp decreases in tariffs and increases in bandwidth. (See Chapter III.) At that point, direct links between mainframes and remote workstations could look most attractive. At least that's the theory of a diverse band of hopefuls: AT&T, Apple and many Big-MIPs IBMers.

6. Equipment Utilization: Naturally, the maxi vendors strenuously disagree there are currently any diseconomies of scale. Gene White counterpunches with the news that mainframes are used around the clock whereas minis work only during business hours. Convincing argument - if the rates of after-hours utilization and working hours utilization are roughly balanced. Certainly, that balance was customary when batch processing represented most of the work load. Today, however, mainframe budgets are sized to support the onslaught of on-line transactions, while background batch volumes remain relatively constant. Thus, a cost-effective computing approach might include a mixed architecture shop that includes a cheap MIP transaction processor. Then peak interactive requirements could be met through inexpensive merchant technology. The mainframe budget would be sized more closely to the batch load.

7. Operating Expenses: This is the final factor in the economies of scale debate. Here, the mainframers have simply been slandered. Minicomputer fans like to ignore the bad news: mini maintenance costs are higher relative to purchase price; minis need operators even if those operators are unofficial part-timers whose salaries are more easily submerged than those of dedicated mainframe jockeys.

D. And Mr. In-Between: Minicomputers

The free play of all these independent, sometimes conflicting, trends has not yet produced any decisive pattern. But as we trudge across the landscape, we encounter four fate-of-the-mini scenarios circulating in the computer community.

1. The Chicken Little Scenario

Nine years of uninterrupted bonanza are over. Minicomputers are on the wane among Fortune 500 customers, squeezed between personal computers or "workstations" from below, mainframes and supercomputers from above. Today's premier suppliers will be left to scratch in the small business market, in isolated corporate divisions and in specialized applications like network nodes. Devotees of this scenario project revenue growth rates for mainframes, minicomputers and personal computers as 20 percent, 15 percent and 35 percent respectively. Quite a reversal from the epic days of minicomputer invincibility - if not immortality.

2. The Atlas Shrugs Scenario

Minicomputers recapture the smaller mainframe market and even penetrate supercomputer land. The drawing cards: dazzling gains in capacity and price/performance through high powered bipolar circuitry, wafer scale integration and/or RISC architecture; plus potent networks and architectural symmetry. The battle lines will be drawn at around 10 MIPS. The megalith producers will be left with two unappealing choices: invest even more research in exotic and exacting technologies or be relegated to the weapons-and-weather fringe.

3. The Peter Pan Scenario

Microprocessor combines like the Synapse machine will grow still larger and achieve perfect parallelism. These totally modular architectures, based on a single technology like CMOS, could then be extended through the entire computing range. Viva economies of scale!

4. The Tweedledum Scenario

Instead of continuing to architect their own processor engines, minicomputer manufacturers will shop at the micro merchant. They'll narrow in on user-friendly workstations for the corporate market and beat a retreat from upscale, multi-terminal mini development. That means they'll be counting on their own unique software (human interfaces, particularly) and maybe on idiosyncratic local networks to preserve any real product distinction - and their existing customer bases. A potential complication: their ambivalence towards industry-standard operating systems like MS/DOS and UNIX.

We visit four of our favorite minicomputer companies: Digital Equipment, Hewlett-Packard, Wang and Tandem. Here there was less excitement over exotic technology or squirrel cage innovation than we encountered among both the megalith and micro sets. Business pressures dominate the conversation.

- Topic A is the competition. Some are nervous that IBM might turn predatory. Everyone was surprised by Apple. No sweat over AT&T (too paunchy) or the Japanese (too monolithic). The deepest worry is the real losers; their thrashing waves of price cutting may drown the winners too.

- Markets are admittedly more claustrophobic than five years ago. No more wide open ranges. Some hope of fences...maybe vertical specialization, local networks, up-and-down architectural compatibility, account control.

- Humbug to popular truths: robots, home computers, management by pie chart. None of that nonsense will hype the market.

- Prices will decline, but slowly. Fatter human interfaces will absorb most of the semiconductor triumphs.

- The product planning cycle in this set is rarely over three years. (Mainframes are seven.) To predict 1990s technology, use straight line extrapolations or make it up. Reacting is certainly more important than planning.

- Corporate organization attracts considerable interest as a topic. Emphasis on being more innovative. Emphasis on being more disciplined. Emphasis on being faster-to-market. All of the above...

1. Digital Equipment Corporation

Snow storm's over, morning skies are clear, but there's no helicopter. So we are cautiously driven through slushy country roads to Maynard. That's technology for you.

We meet with Ken Olsen, Founder/CEO and Win Hindle, VP Operations. Both are subdued (but still forthright) after the recent spate of nasty business surprises. Part of the discussion deals with efforts to correct sticky problems. After an awkward reorganization period, the sales force is gradually settling down (See RB Report: Roots and Rhythms.) The initial disarray in order processing and financial reporting systems is being cleaned up. Most of the holes left by management defections (some hurt, Win allows realistically) have been competently refilled. The DECsystem-10 upgrade was cancelled when it became apparent that complete development would take another three years. The VAX upgrade should be announced next autumn. Personal computers are now moving according to plan, but Olsen is exiting the TV-blitz, show-biz market. "Marketing has always been our strength," insists Olsen. He explains: "If you know the names of your customers and you know your product, you don't need Charlie Chaplin."

We make a run at the 1990s. "Brilliant ideas never come from long range strategic planners," says Olsen. True, no doubt. Perhaps Digital is understandably preoccupied with current crocodiles, and can't muster much enthusiasm for futuristics now. Or perhaps Olsen is heeding his cherished dictum that anyone who announces his strategy, loses it. Eventually we settle on organization as the topic. "I would rather emphasize an organization which encourages diversity, hence creativity, than try to predict what that creativity will produce." He's right.

Although concerned about innovation, Olsen is one certifiable visionary who doubts several industry fancies will survive much market scrutiny:

- Star Wars Robotry. "The high falutin claims and promises about machine intelligence are ridiculous. What industry really needs are efficient automation devices that do the same thing over and over again, until they're reprogrammed."

- Dreams of executives glued to their workstations. "Nonsense. No one runs a business with color pie charts. Management takes analysis, work and worry." Amen.

- Artificial Intelligence. "I'm always suspicious of something so popular. But maybe it's just an extension of what we are doing in expert systems anyway. And maybe AI will consume lots of machine cycles..."

- Non-Keyboard Interfaces. "Only for specialized applications and oddball situations pursued by start-ups. Keyboards will remain dominant." But don't go away. Digital has a new voice synthesis device fed with ASCII characters, "Dectalk for sizzle," says Win Hindle. Sounds attractive.

We try again for 1990: The biggest changes, Olsen reminds us, generally arrive unnoticed and are taken for granted (e.g., the ubiquitous credit card). That will happen again. Digital has several mundane applications which could be as pervasive. Perhaps electronic mail: "My memos are read in offices around the world in minutes. Soon all American industry will operate that way."

Steady Grey has two primary future thrusts: more powerful processors and unique communications coherence. High performance VAX models will reach well beyond Digital's historical turf, ranging up to capacities just short of Cray. "We've got machines coming along that would be considered gigantic by current standards." After that, the next frontier may be opened by the technology option aquired from Trilogy, "assuming it works; if not, we have other irons in the fire," smiles Olsen. As for parallelism, Olsen doubts many tasks can actually be disassembled, though he describes clustered and Ethernetted VAXes providing hyper-parallel capacity. "We're too ignorant to attempt linking up thousands of VAXes, but the universities love experimenting with our equipment..."

Olsen gets most excited about Digital's networking prospects. "Our vision is the integrated computer network serving the customer's entire organization. Unfortunately, IBM has the same vision and we know they have been targeting our markets." There'll be an important battle: "our Ethernet versus IBM's token passing net. It's a battle IBM can't afford to lose...so it should be fun. Meanwhile, we'll maintain the VAX architectural discipline to preserve connectivity across all our products. Few other companies can pull that off, maybe just Digital and IBM." Probably right.

The mid-range market isn't glutted, Olsen says confidently, and it won't be easily niched. The real surprise of the '80s is the success of 'me-too' operating software. "I don't know how we could have foreseen it, but now we're supporting MS/DOS and UNIX. Most changes come slowly when you're watching. For example, semiconductor development obviously had a profound impact; nonetheless, it's been a slower process than outsiders realize."

V. AND IN A NUTSHELL

Avalanches of technological change are coming. The next ten years are likely to witness more innovation than the previous twenty. The shifts will jolt everyone, including those on firm ground today. Some speculations on the decade in Richter Scale order:

1. The communications sector will undergo an earthquake of price/performance realignment as deregulation, divestiture, competition and intensive plant optimization shake the industry foundations. There will be at least two more years of turmoil. Then the price of plain vanilla telephone calls could plummet: perhaps one-tenth of the current tariff? Circuit capacity could balloon: a video minute for 25 cents? Carriers will rush to protect their revenue base with new services, of course. Wise customers will reevaluate their rationale for private networks. Applications flowers will bloom in the new landscape when a fresh cottage industry develops to exploit cheap telecommunicating, just as cottage industries flourish around personal computing today.

2. Human interfaces, input prods and output options will benefit from kaleidoscopic innovation. Conventional keyboards will remain the tools of production data entry. But for the executive suite, QWERTY keys will soon atrophy into a last-resort emergency panel. Current technologies like voice recognition, touch sensitive screens, mice and expert systems are already enough to metamorphose tomorrow's management workstations. Totally unexpected changes may be coming as hordes of young entrepreneurs coat hardware with new ideas.

3. The mini, micro and personal computer sectors have already blurred as minis use micro engines, micros take mini markets and the PC gets confused with ET. New realities may force new definitional distinctions: single user versus multi-user; general purpose box versus specialized box; high performance chips versus low cost chips. In one form or another, however, the office computer will become more ubiquitous than the telephone.

   Eventually, technologically induced restructuring will create smooth price/performance curves, with $10,000 MIPS as the lowest common denominator. Capacity will range from one MIP personal computers to $100,000 ten MIP midis. Gangs of 40 micros producing 250 coordinated MIPS for the price of today's mainframes, is not an outrageous expectation either.

4. The mainframers found sweet relief from the mini threat when the minis were distracted by micro attacks. Now mainframe producers envision a bright future of glass house configurations linked to powerful workstations by inexpensive communications lines, with few minis in between. But that sunny outlook may cloud. If uniprocessors hit a capacity barrier at around 50 MIPS, current four CPU multiprocessing architectures won't be enough to escape the reach of the rampant micro mobs. Mainframe prices would then have to plunge close to mini levels.

   Perhaps something new is needed - maybe new technology migrating from the super set, maybe new architecture. Maybe both.

5. Computer storage will travel a somewhat gentler price/performance curve than most industry segments. It will also avoid the market sector collisions that threaten communications or processors (of all sizes). By 1994, magnetic disk cost per megabyte could fall from $40 to $2.50; access times might decline from 16 milliseconds to five with disk arrays, multiple actuators and/or buffer cache.

   Optical disks in write-once form will prove very cost effective for backup/audit trail applications. But multi-write optical disks will probably not appear in commercially significant quantities in the mid '90s.

6. Software will continue its gruelingly slow progress. Cheaper hardware will finesse the problem somewhat; so will higher level languages. The Japanese may gain some insight from their fifth generation project. On balance, however, it's doubtful software productivity will improve by more than six percent a year on average. The software lag could painfully frustrate the promise of the semiconductor juggernaut, particularly if new "parallel" architecture is sufficiently attractive to make major conversions of old software a tolerable prospect. Or there might be intensified interest in third party packages to break the bottleneck.

7. Computer insecurity remains an unpredictable wild card whose pervasiveness could upset all bets in this game. Many prudent managers could simply bar their doors against technology's full potential until they were assured of much stronger safeguards for sensitive materials.

8. Another wild card is the outlook for the industry's two giants, IBM and AT&T. One might turn predatory and/or be dismembered by some anti-business Federal Administration. The other, already dismembered, might also turn predatory and/or collapse. Either way, the backwash would really rock all boats.

9. A final wild card is industry management. Among the computer crowd, only IBM, Hewlett-Packard, and possibly Apple have successfully navigated the transition to professional management. (All the other management transitions like Burroughs, Honeywell, et al. either followed too late on the heels of disasters or precipitated them.) Ninety percent of this industry's significant companies are still strongly driven by their founding fathers. And we're not talking about reborn Osbornes. To name some: Jesse Aweida (STC), Seymour Cray, Edson DeCastro (Data General), Bob Noyce (Intel), Bill McGowan (MCI), Bill Norris (CDC), Ken Olsen (Digital), Ken Oshman (Rolm) and An Wang.

   These men are among the industry's brightest lights. In the next ten years, many of them will retire. Will their companies continue as potent industry forces in their absence?

NOTES

TRANSITION OR CRISIS?

Strategic Vendors: 1986

THE RESEARCH BOARD

III. DIGITAL EQUIPMENT CORPORATION

A. Steady Grey

We've been visiting headquarters in this old woolen mill for ten years. And never seen a better company in better shape. They've made impressive progress in correcting the weaknesses identified in our first (1976) case study. That year Digital launched an elaborate plan to penetrate the word processing market with no-hands support. Users would get trained only through videotapes; no live instruction or implementation help. "Handholding isn't necessary," sniffed the responsible VP. "Digital delivers such great hardware/software value that office managers don't need explanations." Good luck, we offered - and fled. But luck alone wasn't enough in that notoriously service-intensive market. Anyway, there may have been no other option, given the unoffice-broken sales force's experience at that point.

Access to stronger areas of the company followed our formal presentation of credentials to the beguilingly self-effacing VP Operations, Win Hindle (1979). Thereafter, increasingly positive RB Reports chronicled upward trends in three key areas: product development, manufacturing and field maintenance. And continuing customer complaints over marketing. First, the uppies:

1. Central Engineering (Product Development): Long led by intellectually brilliant and managerially mercurial Gordon Bell. What IBM wished for its mid-ranges, DEC accomplished. An architectural zoo was gradually funnelled into VAX for everything from micro to mainframette. All that provides appreciative customers with operating systems consistency, applications granularity (one-program-fits-all) and slick communications. Terrific! Kept afloat on-a-chip are older architectures now reserved for OEMs (resellers) and an occasional terminal engine. Forthrightly discontinued was the DECsystem 10/20, despite dire predictions that major customers would bolt. And dire seemed direst when the news was broken to the large accounts who represented 35 percent of total revenues. But 80 percent of DECsystem 10/20 customers stayed with Digital. And the company ended up with the industry's most coherent product line. There could be important advantages in not having nervous marketeers as chief executives!

2. Manufacturing: Under pragmatic Jack Smith, steady improvements in automated testing, assembly efficiency, logistics systems and inventory management. Digital has the strongest program of vertically integrated semiconductor (proprietary logic) fabrication in the minicomputer sector. And it has very solid strengths in disk technology. Smith currently wears three hats: Manufacturing, Central Engineering (R&D), and half of Marketing (the erstwhile "Product Groups"). More about them later.

3. Field Maintenance: Under competent Jack Shields, one of the most vigorous programs of improving field engineering through technology. Digital's remote diagnostics center, among the industry's first, adorned the cover of its 1978 annual report. Subsequently, the company leveraged its technology lead with complementary improvements in deployment, dispatching, performance measurement. Customer satisfaction survey results climbed from a snakebitten 4.0 (out of 10.0) in the early Seventies to 7.8 by 1982 (then the economic optimum) to 8.2 this year. And emphasis on technology (rather than bodies, or "pins-on-the-map") raised profitability, even while the ratio of service fees to equipment prices declined. Shields was named head of worldwide sales and service in 1982.

4. Management: A very reasonable continuum, despite lurid Tradeamation descriptions of massive defections from the executive roster. Closer examination reveals that some were necessary casualties (often self-inflicted) in the three-year conflict over Product Marketing - others were seduced by the mounds of venture capital thrown indiscriminately at anyone pedigreed by a top-tier computer company. Probably only the loss of Gordon Bell was significant. In the end, the management structure remained basically intact. And even stronger, we think. At the top, the six-man Executive Committee still includes Olsen, Hindle, Smith and Shields, plus Jim Osterhoff (Chief Financial Officer, recruited from Ford) and Jack Sims (Personnel).

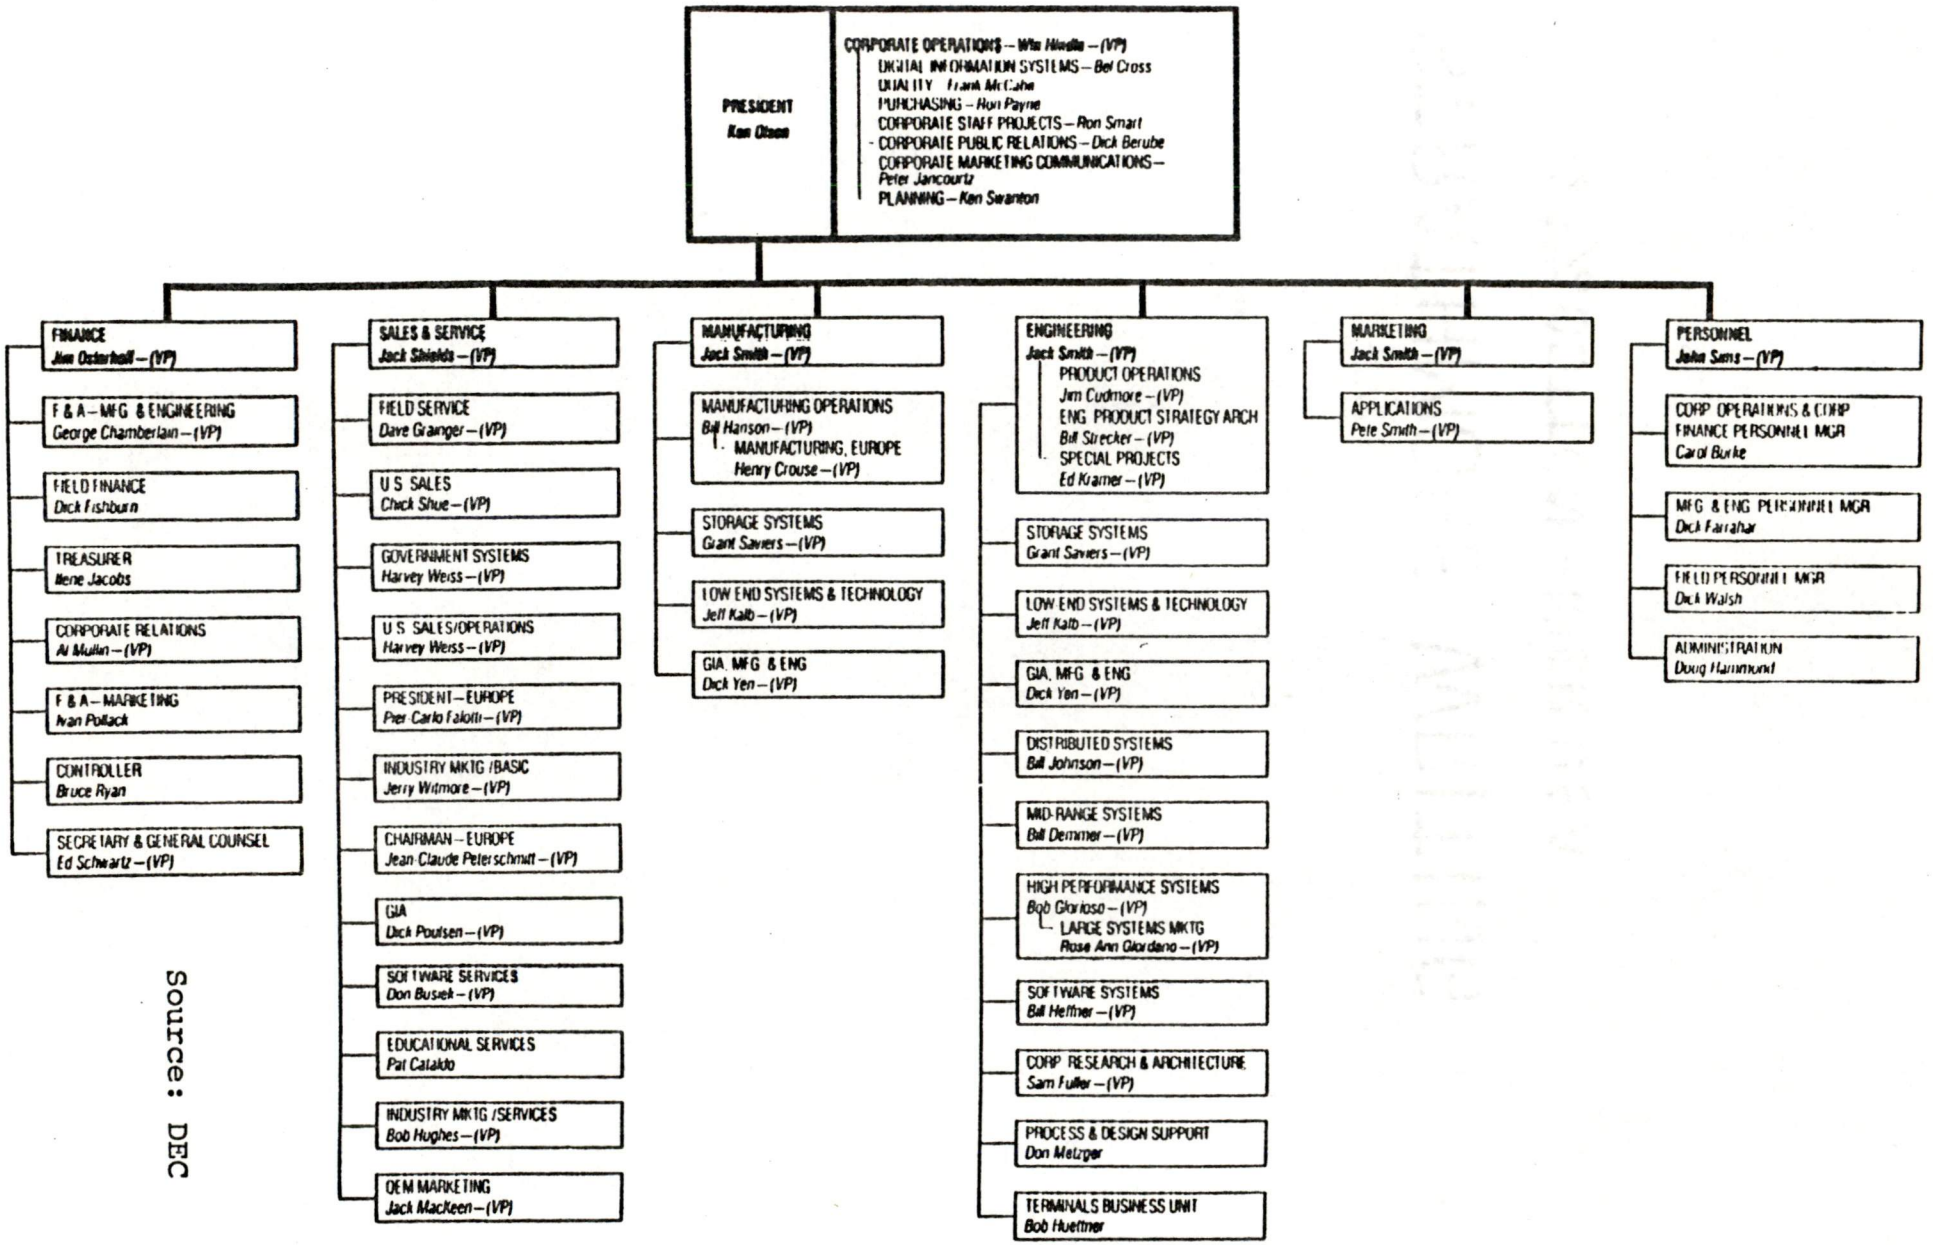
PRESIDENT
Ken Olsen
CORPORATE OPERATIONS – Win Hindle – (VP)
DIGITAL INFORMATION SYSTEMS – Bel Cross
QUALITY – Frank McCabe
PURCHASING – Ron Payne
CORPORATE STAFF PROJECTS – Ron Smart
CORPORATE PUBLIC RELATIONS – Dick Berube
CORPORATE MARKETING COMMUNICATIONS – Peter Jancourtz
PLANNING – Ken Swanton
FINANCE
Jim Osterhoff – (VP)
F & A – MFG & ENGINEERING
George Chamberlain – (VP)
FIELD FINANCE
Dick Fishburn
TREASURER
Rene Jacobs
CORPORATE RELATIONS
Al Mullin – (VP)
F & A – MARKETING
Ivan Pollack
CONTROLLER
Bruce Ryan
SECRETARY & GENERAL COUNSEL
Ed Schwartz – (VP)
SALES & SERVICE
Jack Shields – (VP)
FIELD SERVICE
Dave Grainger – (VP)
U S SALES
Chick Shue – (VP)
GOVERNMENT SYSTEMS
Harvey Weiss – (VP)
U S SALES/OPERATIONS
Harvey Weiss – (VP)
PRESIDENT – EUROPE
Pier Carlo Falotti – (VP)
INDUSTRY MKTG /BASIC
Jerry Witmore – (VP)
CHAIRMAN – EUROPE
Jean-Claude Peterschmitt – (VP)
GIA
Dick Poulsen – (VP)
SOFTWARE SERVICES
Don Busek – (VP)
EDUCATIONAL SERVICES
Pat Calaldo
INDUSTRY MKTG /SERVICES
Bob Hughes – (VP)
OEM MARKETING
Jack MacKeen – (VP)
MANUFACTURING
Jack Smith – (VP)
MANUFACTURING OPERATIONS
Bill Hanson – (VP)
MANUFACTURING, EUROPE
Henry Crouse – (VP)
STORAGE SYSTEMS
Grant Saviers – (VP)
LOW END SYSTEMS & TECHNOLOGY
Jeff Kalb – (VP)
GIA, MFG & ENG
Dick Yen – (VP)
ENGINEERING
Jack Smith – (VP)
PRODUCT OPERATIONS
Jim Cudmore – (VP)
ENG PRODUCT STRATEGY ARCH
Bill Strecker – (VP)
SPECIAL PROJECTS
Ed Kramer – (VP)
STORAGE SYSTEMS
Grant Saviers – (VP)
LOW END SYSTEMS & TECHNOLOGY
Jeff Kalb – (VP)
GIA, MFG & ENG
Dick Yen – (VP)
DISTRIBUTED SYSTEMS
Bill Johnson – (VP)
MID-RANGE SYSTEMS
Bill Demmer – (VP)
HIGH PERFORMANCE SYSTEMS
Bob Glorioso – (VP)
LARGE SYSTEMS MKTG
Rose Ann Giordano – (VP)
SOFTWARE SYSTEMS
Bill Heffner – (VP)
CORP RESEARCH & ARCHITECTURE
Sam Fuller – (VP)
PROCESS & DESIGN SUPPORT
Don Metzger
TERMINALS BUSINESS UNIT
Bob Huettner
MARKETING
Jack Smith – (VP)
APPLICATIONS
Pete Smith – (VP)
PERSONNEL
John Sims – (VP)
CORP OPERATIONS & CORP FINANCE PERSONNEL MGR
Carol Burke
MFG & ENG PERSONNEL MGR
Dick Farrahar
FIELD PERSONNEL MGR
Dick Walsh
ADMINISTRATION
Doug Hammond

Source: DEC

B. Civil War

Of course, researchers find it comforting to recount the trends they've identified correctly. More enlightening, however, to recognize the mistakes. At Digital, ours was underestimating the power of the Product Groups (marketing, not sales). But far from gracefully sharing their prerogatives as others also expected, they kept a tight lock on the marketing function for years.

Our introduction to the Product Groups came from the euphoric press analyses turned up in background digging for our 1980 visit. Cited favorably was their ability to focus sales support, software, promotion and even specialized products into given market segments like laboratories or factory floors. Their single-minded dedication to vertical markets and their profit center clout was the secret of strength in Digital's traditional bastions, cheered the Tradeamation rooters.

Our own interviews with these product marketing executives were unsettling, however. About the RB-scale commercial world, they seemed completely disinterested. Meanwhile, the actual sales force was invisible to the MIS executive. So we predicted: "In the short term, the business systems customer will have to continue depending on Digital's underlying strengths in engineering, manufacturing and field maintenance rather than its commercial salesmen." And concluded that the product marketeers were overrated: "Although they have received considerable attention in the trade press, their real clout over line divisions seems modest." Well, we were wrong on clout, and in considering this lack of MIS orientation any short-term problem. It drifted on and on. No wonder many found Digital frustrating and impenetrable.

We got a closer look at the Product Groups in 1983 and found eighteen of them (up from ten in 1979), grouped variously about applications markets (laboratory, manufacturing etc.), product lines (DECsystem 10/20), distribution channels (OEM or retail) and industry segments. Seemed rather haphazard to us. The Product Group VPs were often members of the original entrepreneurial band that had succored DEC since its formative years. Not surprisingly, many operated in a free wheeling mode, ignoring the Group-Group oversight executives in the formal organization.

Nor were they merely promoters. Each Product Group had (or could have) 200 people, its own equipment inventory, own authority to negotiate terms and conditions, own order entry and accounts receivable system, own advertising budget, own programming staff for industry-specific software projects and data acquisition devices, and, hence, the freedom to allocate substantial resources to customers for almost any type of support. One customer with an application favored by a Product Group would draw official smiles and tangible benefits; another equally large customer with a business systems application and no sponsorship might not even get hardware delivery. No MIS Product Group had been designated among the eighteen.

Yet all this was entirely sensible at one time. Fledgling Digital had been composed of two distinct hemispheres: the hard-technology domain of engineering and manufacturing led by Olsen, Bell et al. on one side; the ninnies-of-the-field in sales and maintenance on the other. So long as the product vice presidents conjoined these two domains, they served a purpose commensurate with their clout.

Gradually, however, the Product Groups became autonomous fiefdoms, thickening the puzzling impenetrability of Digital for the puzzled MIS set. "Management finally sensed the frustration of major corporate accounts over this inexplicable fragmentation...(and) that the product groups had seriously hampered, even victimized, the sales force. In sum, the time for this 'product interface' organization was over," we wrote then. So 1983 would inaugurate major U.S. organizational changes, already being successfully implemented under Jack Shields' stewardship in Europe. The Product Groups were slated to lose many authorities and activities, concentrate instead on true marketing. The sales force was to be energized with more support and better training. This was also supposed to improve resource allocation in the field and create a more unified account presence. The plan's adherents and architects hailed it as "the New Digital."

We fully recognized (and applauded) this move as more than the usual corporate shuffle. And we, like Digital's senior management, expected a certain amount of initial dissension. Inevitable in corporate life. But not the prolonged infighting that ensued. The continual contention attracted lots of press attention, and the stresses grew unimaginable. In one particularly troublesome quarter (9/30/83), an up period for orders was turned into a downer for sales because the Product Groups' multiple order entry systems couldn't give Manufacturing a coordinated demand schedule. Earnings fell 75 percent; the stock price dove 20 percent in one afternoon. Eventually, Ken Olsen probably intervened. Erstwhile comrades departed in droves. Top management persevered with its totally appropriate strategy.

That moment also ended the short-lived era of pop culture at Digital, personified by its fling with consumer marketing. "We goofed - even going into the personal computer retail market. Probably we could have won, but victory would have destroyed us as a company." We agree. A clear sign of Digital's withdrawal from that world is its avoidance of Chaplinesque advertisements and pc look-alike software packages. Back to its roots: one company, one strategy, investment only where DEC has differentiable strengths. "Follow the crowd and you've got nothing to offer," saith the triumphant President.

C. Reconstruction

By now, the old Groups seem DECimated. Eighteen are now eight. Five have an applications flavor, with four collected under Jack Smith's "Marketing" hat. The three others have an industry flavor and report into Jack Shields' Sales organization. Gone are most Product Group prerogatives on equipment decisions, contract negotiation, order entry, accounts receivable, resource allocation and software development. (One exception on the last point.) Left is advertising, promotion, sales support. Generally, the first five line up today like this:

1. Applications Marketing is now three Groups (Computer-Aided Manufacturing, Office and Laboratory). The largest Computer-Aided Manufacturing, is 400 strong, combines the former Engineering and Manufacturing Product Groups. It alone still has a software development cadre, 120 programmers for MAP, the General Motors-inspired, standard factory floor protocol and other manufacturing interfaces. The Office Group is charged with promoting Digital's presence in administrative areas like Finance, Personnel and Sales. It markets All-in-One, word processing and electronic mail with a total staff of 150. The Laboratory Group works Digital's traditional turf with 100 specialists. Included here is a small team dedicated to transferring the company's expertise in artificial intelligence to the marketplace.

2. Large Systems Marketing hopes to penetrate the MIS market. And they'll treat the glass house as a potential sales opportunity rather than an enemy island or distant fortress. The effort is just starting, admits its leader, VP Rose Ann Giordano (formerly DECsystem 10/20 Product Manager). It will take considerable resources, recognizes everyone.

3. Small Business Marketing, fifth and last, reports across the aisle to North American Sales.

D. Sales Serene

Ever the super statesman, Jack Shields properly focuses on the goals of restructuring, not the awkward transition to the "New Digital." Improved coordination in the marketplace was the main issue, he stresses. So the five Applications Marketing Groups just described now focus on layering applications appeal on raw Digital products. The three additional "Industrial" Marketing Groups captured by Sales VP Shields are more closely aligned with the sales force. Their role is to support the field with specialized skills and aids, and to communicate generic Digital products and capabilities to specific sectors. So the Basic Industries Group can assemble packages of hardware, software, CAD, lab applications, etc. for the manufacturing customer. The Services Industries Group has a parallel role in that sector. Meanwhile, the Channel Marketing Group nurtures the indirect distribution arrangements, primarily with resellers.

Shields valiantly attempts to explain how customer requirements are communicated from field Sales to Industry Marketing to Applications Marketing to Central Engineering. Patiently, he sketches overlapping matrices. When we still don't understand, he assures us that committee dialogues will bridge any remaining gaps between the hemispheres.

Subliminally, however, the important message is that these gaps (one Product Group legacy) no longer matter. Field Sales is developing its own direct linkages to Manufacturing and Engineering. And anyway, Sales has most of the necessary resources today; in fact, something nearly akin to profit and loss drivership. Marketing staff is finally staff. For RB customers and prospects, all this promises much cleaner lines of communication and better decision processes on resource allocations.

Shields' direct reports include: the three marketing units just discussed; the field sales organizations for U.S., Europe, and everywhere else; two service departments (see below). And U.S. Operations, which has collected an important array of previously fragmented responsibilities: order entry, demand determination, pricing and sales administration. With all that, Operations (and its parent Sales) naturally has stronger linkages to Manufacturing nowadays.

The two service departments are Field Services and Software Services. The former performs software and hardware maintenance. And the latter has three principal units: the 5000-person consulting business ($300 million in annual revenues); the "Special Systems" business, which supplies customized hardware ($190 million); and training ($120 million), offered through a large campus in Bedford, Massachusetts, as well as several regional facilities. (Gone from Software Services is the sizable staff that certified and wrapped third-party software packages during Digital's encounter with personal computer retailing.)

Added to Software Services are 17 Applications Centers, a major new Digital sales thrust. Their personnel complements range from 15 to 100 people. Each Center is located in the heartland of its target industry: Detroit for autos, Santa Clara for chips, Southern California for aerospace, Hartford is getting an insurance Center, and so on. The Applications Centers provide a varied menu of capabilities and become the focal point for Digital's knowledge about an industry and its system requirements. They house joint software development projects with selected customers (e.g., GM, National Semiconductor, etc.). They supply direct access to Operations for timely product availability and special pricing decisions. Ideally, the Applications Center staff understands the applicability of both Digital's generic software (e.g., All-in-One) and popular independent software for that industry. As well as the usual training seminars, product demonstrations and sales materials. In sum, this is another important step in moving discretionary resources into the field. "And we can't staff these Centers fast enough to suit Mr. Shields," notes the Software Services VP. Great!

The swing from the Product Groups to Sales is particularly visible in the functions now garnered by the Applications Centers, Industry Marketing and Sales Operations. At the same time, top management attention has focused more on selling. Heretofore, it hadn't really interested the Maynard/MIT engineering set that spawned Digital; yet our recent luncheon conversation with Ken Olsen dwells on Digital's sales force and its marketing support staff. Is there a new balance in the inner circle where it matters most? We're very hopeful.

Digital has "a long way to go in organizing our message and putting it across," allows Olsen, always, always honest. But he'll remain (properly) disdainful of advertising hoopla, insistent that product announcements be substantive and realistic. "We don't believe in 'Marketing' gloss as it's tossed around by trade reporters. I won't advertise unless we have something to say. I'm not convinced our competitors do well with exaggerated claims that aren't fulfilled."

We also discuss Digital's awkwardness in market terrain beyond its traditional laboratory and manufacturing redoubts. "We've probably mistaken great products for good marketing in the past. And I take personal responsibility for our neglect of the financial services sector," says the CEO. Again, admirably direct. "Previously it was purely a target of opportunity. Given the historic cultural differences, perhaps our neglect was unavoidable. Now we're organizing to do better." 'Tis definitely true. The charming partner makes a case for a stronger and more steadfast Digital presence in that sector; the grumpy partner wonders aloud why Digital leaves the insurance industry uncontested to IBM. Whether either argument is effective, we don't know. Anyway, a noticeable upshift in activity directed at the insurance industry begins soon afterwards.

From several executives, we hear lingering concern about the relative inexperience of DEC's new sales force. Fair enough. And about its inability to penetrate the MIS market so far. Nonetheless, Shields has justified pride these days in the quality of his recruits. And all the resources necessary to mount a maximum campaign before too long.

We close our interview by testing him on progress towards several objectives he voiced in discussions three years ago. Inventory levels are more rational now, with central coordination in Sales Operations. Advertising is moving towards central coordination, driven by encouragement from the Executive Committee. The equipment trade-ins tentatively proposed a few years ago are being offered, albeit on a restricted basis. (Processor upgrades, eventually to the 8600, and swaps from foreign peripherals to Digital natives.) Field Services continues to reduce labor costs and improve productivity. All positive. Meanwhile, the trend towards complex clusters of networked VAXes may have created an opportunity for designing, installing, balancing and even wiring these colossi.

The most persnickety problem remains the National Account situation. "After three years of effort, we're only now able to expand from 40 to a few hundred national accounts," acknowledges Shields. "But we're making great investments in developing account managers who grasp that future sales are dependent on understanding the customer's business strategy." Is Digital's refreshing self-critique simply classic New England self-effacement? Probably. Because a few weeks later, the charming partner addresses a meeting of these account managers and is generally impressed by their current caliber.

E. By the Numbers

Digital is weathering demand drought, price storms and the past order handling problems with rock-farmer diligence. True, margins are down (as everywhere); from 18 percent in balmy 1981 to just 5 percent in 4Q85. But debt has been held at 15 percent of total capital, even with enormous plant and equipment investment. Equally important, none of the borrowing has gone to truss new flab in current assets. To the contrary, both accounts receivable and inventory turns have gotten healthier, despite the outside turbulence.

We begin with Jim Osterhoff, the new CFO (ex-Ford Tractor), by poking at the debt/capital ratio: just 10 percent, assuming debenture conversion; no higher than 30 percent even if off-balance-sheet leases are included. Either way, most of the capital at Digital has been invested in new plant and equipment, vertical integration into semiconductors and large disk. Of $2.9 billion in gross capital assets, $2.5 billion have been added in the past five years. Very positive. At the same time, Osterhoff is closely scrutinizing future expansion programs "to be sure we're not just adding bricks and mortar for their own sake..."

Accounts receivable aging has improved from 83 days in 1984 down to 75, even in a buyers' market. The gain, explains Osterhoff, came mostly from tightening field administration of credit and collectables. Other measures might squeeze still another ten days out of Digital's already conservative recording practices. After that, astutely observes Osterhoff, "Accounts receivable problems are simply symptomatic of more important problems affecting customer satisfaction." Getting closer to DEC's standard terms of net 35 days will take another assault on business practices that pervade the computer industry: partial shipments of equipment, delays in installation, uncertainty about early reliability. We agree. And his insight is particularly commendable in a company which has already done much to emphasize issues like reliability.

Inventory turns have actually risen from 2.1 (typically) to 2.3 in fiscal 1985. This positive trend is particularly noteworthy because it bucked intensely negative trends in the overall industry - and the normal drag of multiple new product start-ups at Digital. Several complementary efforts saved Digital from the close encounters suffered by others. Jack Shields attributes the improvement to the centralization of finished goods inventory in Sales Operations, which eliminated the wasteful private stores of an earlier era. Jack Smith credits Manufacturing's efforts to shorten supplier pipelines, thereby reducing both raw material and component stocking levels. And from his Engineering perch, Smith suggests that engineering has improved product design, hence quality and reduced the need for field spares and backup. And four to five inventory turns per year "certainly seem realistic," smiles Smith.

How about margins? we grumble. Osterhoff emphasizes lower costs rather than higher prices, reckons that less optimistic assumptions about growth and commensurate pacing of expenses will help. "Also continuing to get rid of waste, of course." Manufacturing headcount has been reduced from 33,000 to 28,000 over the past several years, through productivity gains and natural attrition. Engineering headcount is right, though we could "get more from the dollars we spend." In sales, the potential warrants even more hirings. At Digital, it's the staffs (and staff-staffs) who remain the soft core. "We've got to do more business with our customers, less with ourselves," he quips.

It takes several attempts before Osterhoff will discuss the sensitive topic of pricing. Yes, Digital margins have reflected particularly (unnecessarily?) aggressive pricing. "Different than other industries," allows Osterhoff, "where the salesmen like low prices, and the engineers want higher prices. Here, the engineering community dominates and views better price/performance as their personal gold star." Nonetheless, recent pricing decisions (e.g., 8800) seem more confident.

One long-term objective is to regain top S&P ratings through improved margins and a clean balance sheet. "Today, we have some sloppiness on the income side [margins]. Some sloppiness still remains on the left side of the balance sheet [current assets] as well, but the right side [liabilities] is fine, reflecting good capitalization of the company." (The stock market thinks it's all fine, thank you.) Osterhoff is professionally impressive; only time can measure the impact he'll have on Digital's strong culture.

F. Deep and Wide

Manufacturing strategies are at the heart of product cost, product reliability and occasionally, the ability to innovate. So we'll need to examine these foundations of vendor health. Our first probe concerns make-or-buy at Digital. Unlike some others, the company is more integrated today than before: microprocessors, disk media and read/write heads. Management's rationale for this direction reflects its "deep concern over several trends that almost compel us to take control of our base technologies," says Manufacturing VP Jack Smith. They're the ever more demanding quality standards for components, faster rates of product introduction and the decline of alternate supplier sources that can be competitively trusted.

Technology control over critical components supplies synergistic benefits between manufacturing, proprietary logic and the design process. For example, the insights Digital gained from having its own CMOS and microprocessor fabricating plants gives its engineers the opportunity to design slicker, higher-density chips than would be available off the shelf. And optimized designs are more feasible, since an internal plant will build proprietary chips in quantities too small to draw top priority from merchant suppliers - especially in a hot market. Homegrown components cost one-third more than the jellybeans produced by merchant suppliers, everyone acknowledges. (Understandable, given the quantity differential.) "But whether it's $1.33 or $1.00 just isn't significant in a $30,000 product," properly observes Smith. Proof of the pudding was the ability to bring Microvax II to market with a torrid three year design cycle: just 45 man years in total.

Another concern is the dwindling number of reliable suppliers. In large disk, for instance, the independents are faltering (at best), and potential entrants are dissuaded by start-up costs. This leaves only direct competitors, on whom Digital is unwilling to depend. Understandable, we think. So, even if market forces eventually drive the prices of some components to commodity levels, the company will participate. Of course, that decision isn't relevant in product areas where several suppliers still exist: small disk, printers and memory chips.

1. Storage Systems

VP Grant Saviers leads four units: large disk (14-inch platters), small disk, tape-and-optical and memory components. First, the fringies. The memory unit is, in fact, a purchasing agent since this is already a commodity business. Although huge (4-megabit) memories are expected by 1990, they won't impact either systems design or systems cost enough for Digital to flip over to manufacturing, insists Saviers. Tape and optical have been combined into one business unit. For now, Digital is manufacturing read-only, or videodisk; actively recruiting publishing and educational partners to build disk libraries. The next few years will witness the appearance of one-write optical, which Digital expects to be viable against magnetic tape for transaction archives. Hence, the organizational marriage. Multiwrite optical disk isn't likely to appear for a long time and may never be worth the freight.

In the most important arena, magnetic disk, Digital is moving from the middle of the (dwindling) pack to the leading edge. In the past, it was content to cultivate performance improvements from conventional technologies: ferrite heads and oxide-coated platters. And two more years of product upgrades will be wrung from that base. After that, the company will begin delivering its own thin film heads and small thin film platters. Capital spending is committed, particularly for the platters; Colorado Springs already has a pilot line up. The full production line will have to meet micrometer precision standards and phenomenal cleanliness levels. So efforts are being made to minimize contamination by (a) reducing the number of employees with access by 10:1 and (b) curtailing the number of suppliers, particularly for components with unforgiving tolerances.

The next generation of storage products will favor 8-inch (and preferably 5-inch) platters over the timeworn 14-inch devices. Perhaps because product yields (i.e., fault incidences) could be better with smaller surfaces to plate or sputter. Or perhaps this will break the performance bottleneck created when more and more data gets placed under a single mechanical read/write head. "Stylistically, it makes sense to build arrays of 5-inch disks...," especially for the data base machines that Digital foresees in our future.

2. Base Technology and Low End Systems

VP Jeff Kalb's turf is semiconductors, circuit boards (purchased and produced), and small systems (Microvax, personal computers, workstations and the long after-life of the PDP-11). The press of time won't permit us to give either this area, or its excellent young leader, more than a quick scan. But rest assured, it's top-flight - and really matters.

At the outset, we're told about the next Microvax generation. It will have fewer chips and more function: full VAX instruction set, disk controller, Ethernet attachment - but not big VAX clustering.

Meanwhile, Digital hews to an ambitious process technology plan (i.e., one-micron line width dimensions by 1991). Even having your own high-tech chip foundry isn't really enough today, Kalb stresses. The business of chip fabrication has become so capital intensive that any manufacturing advantage is eventually driven back to the cost of funds. "And American financing styles and terms simply can't match the Japanese." Under the circumstances, winning takes other differentiators: better design automation, more systems software and closer customer relationships. Pushing buttons on top of buttons, to invoke the RB argot.

3. High Performance Systems

Success at the upper capacity reaches of this market will take more than trendy RISC or parallelism, note several hosts. Parallel processing will be useless for the commercial market without new (incompatible!) software. The only way to get bigger/faster machines, for now, is with higher-powered (i.e., bipolar) circuitry. And that's why Digital has taken care to sustain top-flight competence in automated design and assembly of bipolar chips. But not their manufacture - prohibitively expensive, given Digital's relatively small volumes. Leading the high-end design brigade here is VP Bob Glorioso, responsible for future big boxes, contemporary 8600 and 8800, and previous DECsystem 10/20 (mainly device support).

Product development for the 8600 was a watershed. Glorioso blames the launch delays on three factors: the challenge of high-speed, high-heat chip technology, breaking through the one-time complexity barrier of designing for large-scale machines, and the company-wide preoccupation with personal computers at that time. "Digital didn't try hard enough at the high end. We wasted lots of talent on the pc." Another measure of the personal computer program's true cost.

But that's past history. After the belated champagne, High Performance Systems is left with two powerful resources: a huge, multiVAX-riding automated design system, and a 100-person development team including all the players who worked on the 8600, totally enabled by their experience - and very smart. On the way, perhaps later this year are future generations with even faster: custom circuit chips fabricated by an outside foundry; new packaging; possibly water-cooled - DEC is talking to customers about using the same chillers as for IBM mainframes. All this should push DEC much deeper into mainframette territory. And counter the inroads into its scientific computing space threatened by IBM's vector processors.

G. Golden Threads

Among DEC's greatest competitive assets is its connectivity and networking. The VAX instruction set and operating system span the product line from smallest workstation to 8800 mainframette. DECnet ties them all together using industry-standard Ethernet for local area connections, will adopt international-standard OSI when it's finalized. And amidst all that, gateways to SNA and DISOSS for editable text through either operating system routines or (preferred) a Microvax-based communications server costing (eventually) $2500 for eight LANed workstations. Robust demonstrations impress us.

Responsible for networking is Bill Johnson, VP for Distributed Systems. Johnson defends Ethernet against those (especially IBM) who doubt its capacity, explaining that Digital has tied together 1600 engineers without noticeable degradation. He also observes that the company has become the industry's largest supplier of Ethernet attachments at about $700 each. But Ethernet enthusiasm hasn't prevented DEC from being the first to certify MAP adherence, nor closed its eyes to TOP (a Lockheed-proposed protocol), which Johnson expects to merge into other international local net standards.

As for OSI, DEC is an enthusiastic supporter, wishes that the European politicians/PTTs would stop quibbling and allow the protocols to stabilize. Johnson comfortably concludes that the changes necessary to convert DECnet into the new OSI standard will be minimal. "We've told the Europeans why OSI is good for them and equally good - and potentially profitable - for us. That candor has raised Digital's credibility with these governments." Besides, the Europeans badly want an alternative to IBM.

H. Golden Loops

To CEO Olsen, the connectivity and networking of VAX are massive accomplishments that required "extreme discipline and unbelievable attention to detail." To VP Engineering/ Manufacturing Smith these are the company's bedrock strength. Both are correct, of course. An alternative architecture would have to coexist with VAX/VMS, they tell us. "It must also link into our overall networking strategy; that's one self-imposed (and strict) constraint." And it must offer incontrovertible market benefits. Isn't all that an awfully rigorous test for any new idea, we begin to wonder? Has the golden thread become a golden loop, throttling potential innovation?

Keeper of the connectivity is Bill Strecker, VP Engineering Project Strategy and Architecture, whose clout comes from control over budget allocations, plus close personal support for connectivity from senior management. Strecker's mandate is to ensure that new products are consistent with customer requirements, company strategy - and each other. But he doesn't believe that enforced architectural consistency is a noose - or a life support system for VAX once that architecture grows long in the tooth. "No, my job's not preservation of VAX in the face of a better future alternative."

That assertion is supported by a layered view of architecture in which each layer is replaceable by something better, without disturbing the rest of the structure. In the bottom slab is the machine instruction set: VAX today, but possibly RISC or something else tomorrow (according to Strecker). Next the communications layer - moving to international standards. Third is the operating system layer. VMS today, with Unix plugged alongside at the mid-range, MS-DOS at the butt. Then a data management layer, still under development. And finally an applications layer supported by All-in-One for the office, Baseway for the factory, LIMS for the laboratory, etc. "It's the top layer that should be the architectural focus now," insists Strecker. We never quite understand how layer replacements and plugs can be quite so effortless. But at least this concept isn't a strangler.

I. Corporate Research

Digital never nurtured any sort of abstract research organization sequestered away in a dreamy corporate aerie. Much of its advanced research has actually been conducted through its founders' unique relationships with MIT, Carnegie-Mellon and the other top computer science departments in academia. "We don't have to duplicate their excellent work. We can act as a halfway house instead, proving the feasibility of research concepts that will appear in commercial projects years later," explained VP Sam Fuller on our last tour. But with today's $24 million annual budget, advanced research isn't a charity case either. Two-thirds of that amount are committed to internal projects: parallel processing, RISC, new workstations; one-third to external projects associated with universities and industry consortiums.

Sponsorship for Admiral Bobby Inman's Texas MCC industry research consortium "gives us access to a $65 million research activity for just $5 million," says Fuller. But he admits that the feasibility of transferring technology back to the participating companies has yet to be demonstrated.

Once again, Digital's traditional university connections have become a source of r&d satisfaction - recovered from IBM's frenetic buck-blitz of the campuses over the last few years. "The way we relate to the academics through Engineering, IBM tried to match through Marketing. After the first year, it was clear to the research community that IBM is ever so agreeable but never cares enough to utilize their results." In the end, IBM seems to have thrown lots of money at this world, without getting nearly the benefits they should have expected. That's left DEC's previous preeminence in academia very little dented. And keeps those mounds of VAX software coming.

A key mission of Corporate Research is to ensure the company isn't caught off guard by some unexpected technological blockbuster. Blindsiding seems quite unlikely, however, with Fuller's scenario on multiple time horizons. The 5-5-5 Rule: "Looking at any number of historical precedents, it typically takes five years for any new idea to become viable, five more years for the viable idea to reach the market, and another five for the product to ramp up to a billion-dollar business." So Maynard research contacts inside the universities and even the garages also represent an important form of insurance.

Current research projects in-house include evaluations of RISC architecture (thinner/faster instruction sets) and parallel processing (gangs of microprocessors). RISC draws mixed reviews. Semiconductor boss Jeff Kalb thinks the two or three-fold performance boost RISC promises isn't enough. By the time basic software conversion could be completed (two years), there would be little performance advantage left over the gain expected from simply riding a stable architecture like VAX down the perennial semiconductor improvement slope. "So we should be reluctant to turn away from VAX, particularly if an excursion into RISC would deflect any of our resources from the more important future technology - parallel processing." Sam Fuller concurs on parallelism, but reminds us that these gangs will need completely new software to sparkle - regardless of their members' architecture. "And since software has to be recoded for parallel processing, why not recode to RISC as well and get streamlining too?" That argument may not sway senior management, however. After three marginal prototypes, Ken Olsen isn't very interested in RISC. And he worries that parallel processing still implies many subtle problems like "how will you ever know if one of those processors fails." He's right, we think.

Other advanced projects focus on universal problems: distributed data bases, novel cooling for high-performance processors, and more friendliness. The last is an ambitious effort that includes a significant contingent of human factors specialists from Xerox's once talent-heavy Palo Alto Research Center. No pure research initiatives in expert systems, where Digital's position is already stronger than competitors' - both internally and through activities like Carnegie-Mellon's consulting consortium. As a research epilogue, we ask about the joint venture with troubled Trilogy. At least it proved to Fuller, Glorioso et al. that wafer-scale integration wouldn't be feasible for six years or so. And yielded valuable findings on semiconductor packaging and component interconnect. It putters along today on a contract research basis.

J. Crisp Commentaries

The first and last Digital visits conclude with VP Operations Win Hindle. The big news is the intra-organizational conflict is over. Yes, there were casualties, but mostly in the Product Groups (marketing), only a few in sales or manufacturing. Gordon Bell's defection remains highly visible, but Digital had probably outgrown his erratic genius much earlier. Bell aside, Hindle's picture of a stabilized company is solidly supported by the fact that Hindle, Smith, Shields, Saviers, Kalb, Johnson, Glorioso, Fuller, etc., etc., held key responsibilities here well before the troubles.

In the reconstituted Digital, the surviving marketeers are now in tune with their official roles. Of course, one or two more may still depart. More important, Engineering is working closely with Manufacturing under double-hatted Jack Smith. The sales force is improving. Services, including software, are performing best of all - and increasing revenues from the installed base. Sales per employee are up eight percent as headcount becomes more disciplined. This is the most professionally managed company in the industry today, we think.

Yet it's from a product perspective, thinks our host, that Digital really excels: one architecture, local area networks, marvelous clustering. (In wide area nets, management is still frustrated by the political barriers that are preventing more expeditious movement to OSI.) And Digital will be catching up on workstations, with terrific announcements scheduled for this year.

Remaining weaknesses? "Marketing, as Ken Olsen reminds us three times each week. And we're not yet good at convincing senior executives in large companies. But we've got plenty of excellent products - and interconnectivity. We've got our organizational act together. And we're working extremely well together at the top. Of course, Ken is still an inspirational rabble rouser; he'll always be more of a natural leader than a traditional manager." Sounds great to us.

NOTES

Return to BG

94-006/002/05

CAT'S CRADLE

Network Management

THE RESEARCH BOARD

FOREWORD

This report was prepared solely for the use of the Research Board, which financed its completion. It contains information and attitudes expressed by individual companies that may be considered sensitive.

For this reason, and for proprietary reasons of our own, we would not want this document to be circulated so widely within Research Board organizations as to raise the possibility that it would be distributed outside these companies, in whole or in part.

EvS/CEL

October 1988

CONTENTS

I. Topologies

Network management is a research nightmare! Vast in scope, yet clogged with endless nitty-gritty. And zillions of vendors. Most making slow progress behind marketing hyperbole intended to disguise can't with cant. As everyone slows down to match the glacial pace of standards. Which aren't the whole answer.

A warm-up exercise: What's a network? Not so long ago, voice networks were WATS. While data trickled over point-to-point analog lines at 4800 baud. That's quickly changing. Today voice and data share a backbone of T1 circuits arrayed in rings or figure eights; dotted with intelligent multiplexors that break down those 1,544,000 bits per second to usable gulps. And forward the gulps to data controllers or voice PBXs. Then down the pipeline to LAN communications servers, minis or mainframes. Each LAN a constellation of workstations, each processor with baskets of software. Complex relationships among components that might have originated on different planets when it comes to coherent management.

And what's coherent network management? Ideally, a system that collects alarms and performance statistics from each and every network denizen mentioned above. Then filters and integrates those alarms to diagnose the real failure. Routes the diagnosis to the appropriate point-of-correction. Accepts corrective commands through a single "manager-of-managers" console, not one per vendor. Plus immensely useful ancillary functions: network-wide component inventory control, trouble ticket tracking, security, perhaps accounting.

All critical to RB companies building ever more complex global networks. For ever more crucial applications like CIM and EDI. When your customer's order entry port remains inactive for two weeks, should you send a salesman or a repairman? And if the downed system downs company revenues, can it ever be forgiven? Indeed, the primitive state of network management will soon stifle user innovation, if it hasn't already. Unfortunately, vendors expect that implementing the ideal system described above will take at least ten years.

For us, each supplier interview starts on a high conceptual plane: comfortable and managementy, but irrelevant. To understand the problem's impact, visit a user's network operations center. To appreciate the solution's complexity, probe the vendors about alarm formats and message routing conventions. Humdrum stuff. But vital to gauging industry progress on several deceptively simple user requirements:

1. Outage Alarms - So the right person is notified when your customer's order entry terminal fails. Before he complains. Or worse, reroutes the business.

2. Performance Monitoring - So an incubating outage can be pre-empted by replacing the unit registering soft fails. Or by adding capacity to the network route nearing saturation.

3. Diagnosis and Repair Assistance - So the root cause of related outages or performance bottlenecks is correctly identified. And so fixes can be initiated through a single console, regardless of the troublemaker's logo.

4. Configuration Management - So Operations can monitor and control the entire network from a single terminal. Maybe through a topology map. With nodal subviews (LANs and their workstations) available at the click of a mouse.

5. Equipment Inventory - So the network boss knows the location and status of every device and spare.

6. Trouble Ticket Tracking - So the appropriate service supplier is automatically alerted of a mission critical defect. And so the anxious user can monitor its resolution.

7. Security Alerts - So a node or session can be temporarily quarantined, until keyboard stumbler is incontrovertibly distinguished from potential pirate. The capability farthest from implementation.

Eminently reasonable requirements. But meeting them raises a host of many-faceted issues. Which we discuss in four separate frameworks:

- The Real World: Chapter II looks at network management in practice at several member companies: American Airlines, Bank of America, Citicorp, Federal Express, J.C. Penney, Northern Trust, Rockwell and Seafirst.

- Alliances and "Aliences": Chapter III chases consensus on standard message formats and protocols. Introduces OSI front-runners CMIP and CMIS. Plus a bevy of sometimes overlapping, often competitive consortia like Corporation for Open Systems, Forum and Alliance.

- Physical Networks: Some think voice networks less troublesome than data networks, because voice users can easily sidestep problems by redialing. Others call traditional voice/data distinctions obsolete. Better to characterize physical versus logical facilities. In which physical includes every box and wire between communicating end points. While logical adds the complexity of dueling computers and sinfully incompatible software. In Chapter IV we check out network physique, beginning with private network switches and multiplexors; then go public with MCI, Northern Telecom and Nynex.

- Logical Networks: A messy world. Chapter V visits LAN fans Microsoft, 3Com and Ungermann-Bass; plus Retix and Touch, two unsung heroes of a recent standards bash. Chapter VI questions network management fundamentals at Tandem, Hewlett-Packard and three hopeful middlemen. While Chapter VII takes on four heavy hitters: IBM, Wang, AT&T. And Digital, who wins our award for best architecture.

Several legs of this uphill trek may be arduous. But the view from the summit (Chapter VIII) will be inspiring.

II. THE REAL WORLD

A. User Environments

For RB company telecom operations, unavailability of coherent network management sometimes means several monitoring stations, fragmented problem reports, dueling repair conventions, and gaping blind spots. Challenges increase exponentially with network complexity.

1. Company A

Bank branches are connected to a statewide T1 backbone. A mix of Stratacom multiplexors, Doelz data switches and Northern Telecom voice switches delivers the goods. But that mix complicates network management. The company's control center has a separate monitor for each supplier's components. And a separate mini watching local modems. A Tandem computer tracks branch data connectivity. While a telephone company system oversees transmission trunks.

A single physical outage can trigger alarms from all. Requiring a team of human operators to isolate its precise location. If every ATM in a region signals failure, the fault must lie in the transmission network. But where? If voice traffic dies too, the culprit could be a T1 break. Or a switch. Or a branch channel bank. But which? Absent collaborative diagnosis from the autonomous component monitors, the control center is left to "micro manage the pieces," observes a bank executive. Fortunately, the Stratacom devices reroute traffic around a failed component so fast that neither voice calls nor SNA sessions are noticeably disrupted.

"But we're far from managing overall availability: viewing all these network components the way customers do, as a single window on all our services." True, downed branch circuits, crashed central data bases and cash-strapped ATMs require different responses from individuals with different skills. Yet they all affect the "availability" of customer service, a much broader perspective than today's provincial component management products can provide.

2. Company B

Voice management might seem inconsequential to some. But the company's use of communications gives the task a different weight. After all, fourteen service centers field 250,000 customer calls a day. And network availability is critical to company-wide service levels. Voice facilities are monitored constantly by a mix of management systems:

- A DEC based system generates near real-time performance statistics on Automatic Call Distributors, plus trouble reports and potential overflow warnings. Traffic rebalancing requires manual intervention, however.

- NET supplies T1 multiplexors, and their monitor. With nice graphics for fault detection. But two holes in performance tracking. First, getting throughput statistics only on request inhibits detection of deteriorating circuit quality until it's too late. Second, NET's highly automated rerouting capability can degrade throughput. So the customer's technicians must analyze and reoptimize the network map daily, but would prefer that NET software absorb the chore.

- An AT&T system monitors T1 circuits span by span. Office PBXs (administrative traffic only) are tracked box by box. Another system manages Paradyne modems. All of which leaves network operators sifting through blizzards of confusing, even conflicting alarms and statistics. Until calls from users become the only reliable and actionable indicators of an outage. "We need to integrate these systems, then automate as many network management functions as possible. You can't expect an operator left dozing before the screens all week, to handle a real crisis when every alarm goes crazy."

B. Missing Links

In Chapter I, we listed seven commonly acknowledged (certainly reasonable) user requirements for network management. Yet these experiences from members with leading edge operations show that only the first two - outage alarms and performance monitoring - are usually served. And those only in part. Console alerts are often less reliable trouble indicators than telephone calls from disrupted users. Performance measurements never reflect response times users actually experience. And end-to-end troubleshooting is impossible: today's management systems exclude key network components. "There are so many potential failure points between us and the user," observes one executive. "Problem diagnosis would improve if operators' consoles could mirror user sessions. Instead, we rely on groping over the telephone."

The other five requirements aren't addressed at all. Message consolidation and routing are left to home-grown code and third party software packages. Topology maps - when provided - show only the vendor's own components. Customers who want cross-vendor control of equipment inventories and trouble tracking must develop their own applications.

And several are. Federal is building a trouble ticket system to help operators dispatch the right technician; respond to user inquiries; and track problems from initial report to final resolution. Rockwell is developing a configuration data base, interfaced with equipment inventory systems in its decentralized user divisions. Everyone agrees such a data base is critical to effective network management. "That's the key to integrating various subsystem monitors," notes one executive.

III. STANDARDS MANAGEMENT

A. You Say Tomato, and I Say...

Coherent network management would obviously benefit from inter-vendor standards. Just as connecting two unlike devices requires prior treaties on physical terrain: (i) transmission medium; (ii) network addressing; (iii) routing protocols; and (iv) handshake conventions. And just as actual message exchange requires consensus on logical environments: (v) dialogue structure; (vi) presentation format; (vii) file transfer and message handling. Those seven layers celebrated in story and song. And in countless meetings of the International Standards Organization (ISO). Not to mention ISO hangers-on like COS, OSF, X-Open, SPAG, CCITT, etc. Many acknowledged in last winter's Report, Until the Next Millennium.

Are those muffled yawns and muttered threats from our gentle readers? Sure, vendors proclaim OSI (Open Systems Interconnection) the most important invention since the transistor. And cite customer demand as their primary motive for abandoning hallowed individuality. Yet member IS Executives seem unexcited about standards (reserving applause until the curtain rises). And suspicious that pious vendor palaver about standards minutiae is actually a ploy to delay OSI and extend the life of proprietary goodies. Are standards just a game for insiders and politicians?

B. Framing the Framework

Enough carping. Better proceed to a description of network management standards, notably the OSI Management Framework. Work's begun, with some progress. Particularly on the concepts for multi-vendor data access and message exchange. The Framework envisions that each of those seven ISO layers will participate. And seeks to provide common conventions governing communications: both vertically - within a single device's layers. And horizontally - among peer layers in different devices. Recognizing that most networks include equipment and software from more than one vendor. A messy, inherently counter-systemic world.

Next, Framework folks target six capabilities: fault management, performance management, configuration and name management, security, accounting, and directory services. The first three correlate pretty closely to the priorities raised by member companies we visit. Security also draws considerable user concern, particularly as LANs broadcast every message to every station. Standards work here is in its infancy, however. Ditto accounting, but who cares? A vocal ISO constituency of commercial service providers, no doubt. But it's difficult to imagine RB companies pricing network services on any basis besides standard charge per session minute or processing unit.

Directory services are part of a larger issue: a standard data base or repository for multi-vendor management data crossing the draft stage as X.500. But there is still disagreement on "registration" or how to identify and categorize various management domains, says the VP Engineering at Corporation for Open Systems. Elsewhere, we hear of debates over how repository entries should be structured: for each "object" represented, like a counter that gathers performance statistics, what's the appropriate format (e.g., integer versus floating point)? And how does it operate (can it trigger an alarm whenever a threshold value is exceeded)? If the ISO bees don't keep buzzing, de facto standards from IBM, AT&T and DEC could pollinate the farm.

Working out the critical definitions for every conceivable component and situation falls under the rubric of SMI, or Structure of Management Information. The effort is nowhere near complete: the number of unique SMIs keeps growing. As does their complexity. One example of the definitional rat's nest (offered by a 3Com developer) is the oscillating counter: Consider an SMI describing a counter that triggers an alarm when a preset threshold is exceeded - e.g., more than 26 concurrent users on a LAN. Sounds straightforward. But say User 27 logs on, then another user logs off, then another logs on, etc. Should the counter oscillate around its threshold value repeatedly, sounding an endless (useless) volley of alarms? Or should there be another SMI describing how to turn the damn thing off? Amidst this confusion, one hopeful sign: a new generic category for conditions like "retransmission error" that apply to many SMIs. Reducing the agenda of discussion items. Much remains to be done, however.

Enter CMIP and CMIS, the building blocks of OSI management applications: two hooks where all vendors plan to hang their hats. At least, that's the pitch for mollifying network management-hungry customers. Scratch the topic at any supplier session for a reflexive oath of eternal fealty to CMIP and CMIS "as soon as the standards are set." Which could take several years.

CMIP (Common Management Information Protocol) will allow systems from different vendors to exchange data. CMIS (Common...Services) specifies three basic types of exchange: report an event (e.g., send an alarm); request information (read a performance counter); invoke an action (set a threshold, run a test). So far, so good. But many publicly committed participants privately defer CMIP/CMIS until the SMI issues are resolved: They want concurrence on what the data is before deciding how to exchange and apply it. Sounds reasonable, though hardly ASAP. Meanwhile, CMIP/CMIS are "nothing more than an empty envelope with a stamp," says one insider.

CMIP/CMIS are just a corner of OSI network management. Which creates a dilemma for would-be implementors. The full "stack" of layer-by-layer protocols will require 500-plus kilobytes of memory, and commensurate compute capacity. (As a harbinger, see overleaf for HP's Express OSI board.) No problem for heavy-duty communications servers. But what about lesser net denizens like modems, muxes, even workstation LAN adapters? That's one concern of still another standards club, NetMan: a crew of LAN makers. NetMan champions OSI-style network management, but won't wait for full OSI. Instead, they propose two extensions to TCP/IP, a popular de facto standard: a "proxy" technique to let puny network components send simple "I'm alives" up the big pipe to the glass house. And OSI-sympathetic presentation and directory services, to facilitate eventual migration once OSI settles. A prototype is being rushed to demo as we scribble.

Not everyone believes NetMan boosts OSI goals. A ruse to prolong TCP/IP? Which backs into a storm brewing over whether formal network management standards should be compatible with pervasive in-place network protocols like MAP. Raising the remote spectre of expensive retrofits. And plenty of controversy.

Where are the users in all this? we ask COS chief Lincoln Faurer. Still lagging in participation, he replies. Thirty percent of COS members are users, up from 25 percent a year ago; but still far from Faurer's goal of 50 percent. Moreover, he knows existing user members are frustrated by their limited influence. A COS executive is on the road discussing possible remedies. But the real problem is us. The users. With few exceptions, says Faurer forthrightly, they don't take the initiative, don't work together, don't have the requisite patience to work out thorny details. "Our users are great counterpunchers - give them a proposal, and they'll react and criticize. But they're slow to take the lead. And few realize they have interests in common with users representing other industries." Something for the Research Board to discuss?

IV. PHYSICAL NETWORK MANAGEMENT

A. Nomenclature

In networking, the traditional voice/data split is replaced by a physical/logical differentiation. Where physical refers to transport components like circuits, concentrators, multiplexors, modems and switches. While logical refers to software-based network architectures like SNA, DECnet, and OSI. This new perspective also spells basta to old tussles over whether network "intelligence" would reside within common carrier facilities or on customer premises. Vendor aspirations notwithstanding, it's a draw. Logical networks thread through both realms. So do physical networks. But they don't map precisely: the same logical session needn't retrace the same physical path. A productive snarl which makes heterogeneous network management even messier.

In this chapter, we get physical. And consider network management for several layers. At the lowest, raw circuits - notably digital T1, a 1.544 megabits-per-second pipe. At the next, multiplexors and concentrators.

Semantic distinctions are fuzzy here. But concentrators (like Doelz) generally funnel several low speed data streams into a higher speed channel. While multiplexors ("muxes") do the reverse: scale T1's fire hose streams down to kitchen tap trickles. Consider that common carriers carve T1 into 23 voice channels, operating at 64 kilobits per second. Independent muxers get more - through two techniques: speech interpolation, which fits a second conversation into the natural pauses left by the first. And speech compression, which squeezes conversational bandwidth from 64 kilobits down to 32. Interpolation effectively doubles the virtual channels on a T1 line; compression doubles them again. Timeplex and Infotron muxes use interpolation to get 46 voice channels per T1. While Stratacom and NET piggyback both techniques to crank out 90 plus. For nice gains in cost and capacity over common carrier offerings. Improved reliability, too; muxes automatically reroute traffic around downed nodes or congested circuits. But not without certain pains in network management, as we discover.

VI. LOGICAL NETWORK MANAGEMENT

A. Logical Conundrum

We already agreed the voice community will have network management facilities years before that sweet suite is available in private data networks. But why? Maybe it's the difference between oligopsony and oligopoly (i.e., few buyers vs. few sellers). Telcos and inter-exchange carriers forced voice network suppliers to discover the wisdom of working together. So coordinating network management isn't an unnatural act. Whereas in the data market, the feisty few exercise almost as much clout in setting international network directions as the ever-bickering standards organizations. Which they join and effectively dominate. Oligopolists still buffered from any real interference by the fragmented, generally uninformed customer base. But perhaps any standard is better than none.

Another obvious reason for the pokier pace in data network management is that problems are more complex and far-reaching. Failed voice conversation? Hang up and try again. With a broken data connection, no such easy sidestep. And one break can disrupt business operations, require revalidation of financial controls, even necessitate reconstruction of compromised data bases.

Finally, the intricate inter-relationships of the data network. With the world of physical transmission. Between LANs and WANs, though that linkage is not yet porous to management. And between communications software (SNA, DECnet, OSI) and all the other software (computer vendor, third party or homegrown) whose failure could back into data networks. All apparent as we discuss the subject with several leading edge data communications participants.

NOTES

VII. HEAVY HITTERS

A. Casting Their Nets

NetView is the system every network supplier has to meet - or beat with something much, much better. Its entrenched position appears unassailable: many times the installed base of nearest competitor Cincom. Controlling huge logical networks that utilize a wide diversity of physical communications devices, as well. NetView primacy derives partly from time-to-market; partly from IBM clout. When NetView debuted (1986), other potential contenders had barely scented the opportunity. So NetView got the gateway to the Fortune 500 market. Certainly, that's the perception of network insiders large and small. Unless one of them can hammer together a credible alternative.

In this chapter, we size up the challenge with a deeper look at IBM - which intends to hold its lead through an aggressive development program. Then consider competitive strategies at Wang, AT&T and Digital. That ordering isn't casual. Digital has the most open and realistic architecture, assimilating network components physical and logical, standard and proprietary, large and small, old and new. Could become a leader.

Three buttons to watch: (1) Does the vendor offer an architecture functionally attuned with the requirements described in Chapter I? (2) Can the architecture's repository handle the incredibly diverse array of network devices, services, protocols, alarms etc. that populate large customer networks? (3) Does the network management system filter and integrate alarms and statistics from the various devices to power effective diagnosis and response?

E. Digital Equipment

Flash: Digital pops span man plan. Amidst our labors, DEC finally announces the architecture mulled for three years, at least. (See 1987 RB Report Charting the Gyre.) At a preview with VP Distributed Systems Bill Johnson and chief engineer Gary Gottschalk, we learn that Enterprise Management Architecture (EMA) targets five market requirements:

- Scope must accommodate DECless protocols: whether standard or proprietary, physical or logical. After all, CMIP/CMIS, the great leveler, may arrive late. Retrofits to everyone's hardware and software products, even later. Hence multiple "access modules" as discussed below.

- Functionality must include configuration management, fault management, performance management, security and accounting. But those applications must be independent of what's being managed: a command like "display percent utilization" should be applicable across CPUs, circuits, switches etc. Full function also implies full life cycle support: network evaluation, planning, implementation and maintenance. Plus inclusion of whatever tools the customer chooses.

- Integration through a repository containing extensive information about every network component, including its alarms and performance counters. Driving functional modules that correlate, filter and route diverse messages.

- Scale must accommodate networks small and large. And network operations centralized, decentralized and distributed. "We will not dictate management philosophy."

- Modularity, in the sense of protective isolation. So a change in incoming alarm formats won't rock the repository. While functional enrichments (e.g., expert systems) slide transparently beneath the consistent user interface. Very prudent. And a better structure for mapping the messy real world than anyone else provides.

ENTERPRISE MANAGEMENT ARCHITECTURE

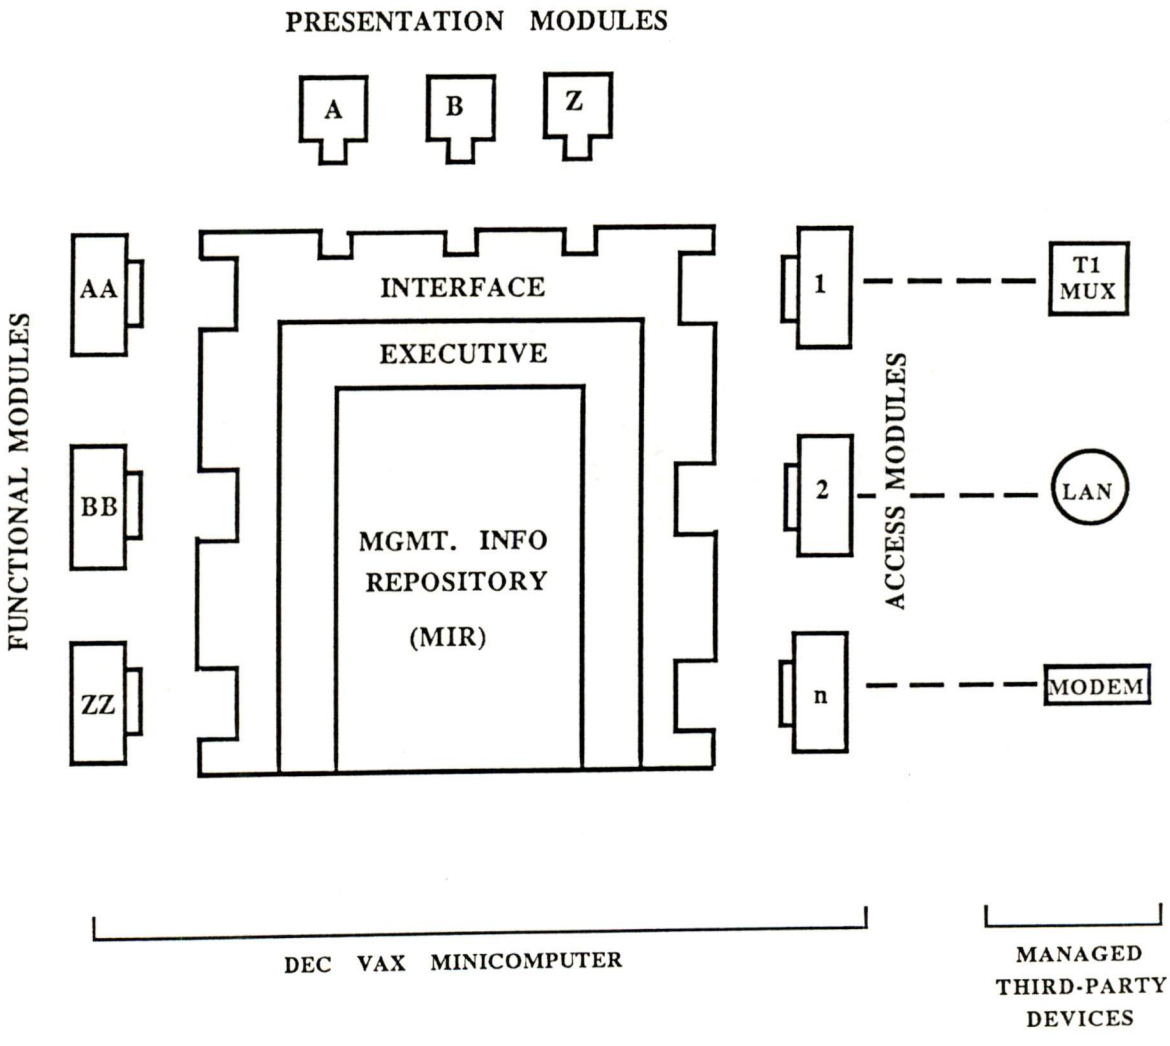

SOURCE: DIGITAL EQUIPMENT

The architecture itself (Exhibit) has four significant segments:

1. Access Modules - as many as necessary. A multiplicity that realistically anticipates the persistence of proprietary protocols even as international standards are implemented. So begin with a module for DECnet; add CMIP as it solidifies; encourage other vendors to contribute. Accompanying the EMA announcement were support commitments from eight component vendors including Codex, Stratacom and Timeplex. Modules for NetView and UNMA may come from third parties.

2. Functional Modules include generic routines for message handling and alarm responses. For example, a system-wide utility defines failed-circuit alarm thresholds (and rerouting instructions) for any mux-make. Here too, different modules that filter and correlate related alarms: e.g., suppress DECnet alarms triggered by a T1 break.

3. Presentation Modules are the user interface options. Fairly simple at first: ASCII terminals for DECnet, 3270 for SNA. Later on, graphics workstations. Plus third parties - like Netmaster?

4. Repository, "the cornerstone of integrated network management," will include: (a) component inventory; (b) performance trends; (c) device-specific alarms, commands etc.; (d) "meta data," or component class attributes and inter-relationships. This last category is critical to achieving the level of integration advocated throughout this report.

On balance, an extremely strong architectural concept made even sturdier by its flexible, modular structure. And by the implicit assumption that networks will remain messy tangles of heterogeneous, hetero-generational bits and boxes for many years. A world that can't be force-fit to a single protocol, even a standard one. All of which makes EMA our favorite architecture - though we've yet to see the building. Muses Bill Johnson, "This construction effort could approach the scale of our VMS operating system."

VIII. Closing The Loup

The struggle toward coherent network management is complex and grinding, thanks to a hostile environment of inter-dependent but not integrated equipment. In which the largest vendors dominate standards alliances while looking for a continuing - and proprietary - competitive advantage. The mid-scale seek aliences that don't give away the store. And the smallest hide behind players they trust won't steal their niche. For all, the technical difficulties of integrating dissimilar alarms, statistics and command sequences would be challenging enough. But the task is compounded by endless supplier wrangles over account control. Bruising for frustrated customers.

To summarize the critical issues:

- The customer problem is quite straightforward. Few network subsystem monitors work together. In fact, they often confound one another. And cause such confusion that the most reliable (if least desirable) alarm is a phone call from an irate user.

- Standards are evolving - but slowly. For the reasons suggested above.

- Forceful oligopsonists rule the public, physical network. Consequently, it's here that network management consensus will emerge first.

- Network management remains blind to LANs. Bridges and gateways open connectivity options for diverse workstations. But block the view from the top.

- Logical network management systems require solid underpinnings: architecture, repository, set of standard message protocols. Almost like a basic operating system. Little wonder progress is slow. Effectively, the lead belongs to IBM and AT&T. Artistically, the prizes go to Tandem and Digital.

Recommended reading for now: Popular Mechanics. Management breakthroughs possible after 1991; management breakfasts before.