Executive Interex

94

Proceedings

Symposium & Exposition

March 14–16, 1994 New Orleans, Louisiana

Sponsored by

hp interex

The International Association of Hewlett-Packard Computing Professionals

Proceedings

Executive Interex 94

Symposium & Exposition

New Orleans, Louisiana
March 14–16, 1994

Sponsored by

hp interex

The International Association of Hewlett-Packard Computing Professionals

HUMAN FACTORS AND PROFESSIONAL DEVELOPMENT

STRATEGIC MANAGEMENT

SYSTEMS ARCHITECTURE

SYSTEMS AND APPLICATION DEVELOPMENT ENVIRONMENT

HP Computer Museum
www.hpmuseum.net

For research and education purposes only.

OPERATIONS MANAGEMENT

The ACTION Strategy for Employing HP Systems to Meet New Enterprise Needs

Executive Interex

March 14, 1994

John R. Logan
Vice President
Aberdeen Group, Inc.
92 State Street
Boston, MA 02109
617/723/7890

Aberdeen Notes:
Aberdeen Group, Inc. is a computer and communications research and consulting organization closely monitoring user needs, technological changes, and market developments.

Based on a comprehensive analytical framework, Aberdeen provides fresh insights into the future of computing and its implications for users and the industry.

Aberdeen performs specific projects for a select group of clients requiring strategic and tactical advice and hard answers on how to manage computer technology.

Audience Notes:

Agenda

- **Enterprise MegaChanges**
- **ACTION Steps**
- **Why HP Is Today's Highly Wanted Partner**
- **Changing Executive Interex Roles**

AberdeenGroup

Aberdeen Notes:
Today's presentation will begin by reviewing the megachanges that are directly affecting the way multiuser computing systems are being used in enterprises. Then we will review Aberdeen's ACTION steps for changing an enterprise's information infrastructure to leverage these megachanges.

Since this is a meeting of executives responsible for HP systems, Aberdeen believes it is appropriate to take a few minutes and review HP's position in relationship to the other major suppliers of multiuser enterprise computing system. Finally, Aberdeen will point out how the megachanges and ACTION requirements are causing a change in the responsibilities of managers of HP systems.

Audience Notes:

Enterprise MegaChanges

Management focus on lowering enterprise overhead costs:

- **First industry competitor to do so has a killer advantage**
- **Advantages are higher profits, can lower prices, leaner/meaner organization**

AberdeenGroup

Aberdeen Notes:
The U.S. economy is for the first time in 25 years experiencing low inflation and significant productivity increases. The cause for this structural change is the ability of industry-leading enterprises to lower overhead costs and then use the increased net margin as a corporate weapon. Please note that many governmental agencies and public institutions are following this same dictate to lower overhead costs associated with delivering their services.

The impact of effectively operating with lower overhead needs and associated costs is seen in both the downsizing of most large enterprises through the elimination of layers of middle managers and the ability of smaller companies to effectively compete against large ones.

Audience Notes:

Enterprise Financial Issues

BigCo Ltd Income Statement
100% Revenue
55% Cost of Goods Sold (COGS)
35% Selling, General, and Administrative
(IS is 3% of Total Revenue)
10% Profit Before Tax (PBT)

If IS Lowers:

- **SG&A costs by 20%, PBT increases 70%**
- **COGS costs by 5%, PBT increases 30%**
- **IS costs by 50%, PBT increases 15%**

AberdeenGroup

Aberdeen Notes:
To provide an example of the consequences of lowering overhead costs, let's review the potential impact of IS' efforts to do so.

With the same amount of management resources, IS has the potential to lower SG&A costs by 20% simply because there has been no focused efforts in this area since mainframe systems were installed. Since the manufacturing engineers keep lowering COGS, IS would be lucky to achieve a 5% reduction here. But a very skilled IS department might be able to lower mainframe-based IS costs by 50% over three years.

Examining the bottom line, PBT, the greatest potential improvement clearly comes from lowering overhead costs.

Audience Notes:

Enterprise MegaChanges

Demand for innovative ways to increase customer satisfaction:

- **Customers want fewer suppliers who provide more value-added through information flows**
- **Examples go beyond 24X7 phone support:**
  - **Empowered representatives**
  - **Just-in-time retail store delivery**
  - **Customer-supplier virtual task forces**
  - **ISO 9000 certification**
  - **Automatic consumer rebates**

AberdeenGroup

Aberdeen Notes:
The second megachange is a customer demand for higher levels of information-based support. The two major reasons for this megachange are that customers:

- perceive that competitive suppliers' unit prices and product features are rapidly converging, and
- want to limit their internal overhead costs associated with managing multiple suppliers.

To increase customer satisfaction, suppliers are finding that they must use information in new and creative ways to create partnerships with their customers

---

Audience Notes:

Enterprise Conflict

Decreasing overhead costs and increasing customer satisfaction require apparently opposite actions

- **Lowering overhead costs means eliminating middle-managers and non-revenue producing customer reps**
- **Fewer staff means more mistakes, inability to staff for peak loads, and less consistency**

AberdeenGroup

Aberdeen Notes:
It is immediately apparent that each of these two megachanges could be easily managed in isolation. But the same actions and investments that will increase customer satisfaction (more sales reps meeting more frequently with customers, more extensive internal supplier-staff training and specialization, staffing for peak load requirements, etc.) also increase overhead costs.

The inappropriate way to deal with this conflict is to create a working environment where employees seriously ask, "Which is better, to be sacked on Friday or pick up the slack on Monday?"

Audience Notes:

Enterprise MegaChanges

HP is a highly wanted supplier of new, global, multiuser enterprise applications

- **Only HP, IBM, and DEC have global, multiuser computing competence**
- **Only HP has provided a realistic vision of how open computing should work and then supplied the hardware, systems software, and services to implement the vision**

AberdeenGroup

Aberdeen Notes:
IS decision makers in global Fortune 500-sized enterprises now recognize HP as the finest supplier of multiuser computing systems for new enterprise-wide applications. HP is no longer the calculator company, but the company that successfully has implemented in products the vision DEC had a decade ago and the vision IBM does not to want to see. Today, no other multiuser company is in the same first tier as HP, IBM, and DEC for supplying multiuser, enterprise-wide production application platforms.

HP has caused this megachange to occur by being consistently better than IBM, DEC, and other multiuser suppliers in most of the technologies and services required for implementing distributed, open production systems.

Audience Notes:

Leveraging MegaChanges

Computer-driven Business Process Reengineering

- **BPR to manage lowering overhead costs and increasing customer satisfaction as one issue**
- **HP hardware, systems software, services, and networking to build enterprise information infrastructure**

AberdeenGroup

Aberdeen Notes:
The megachanges are good for those that can recognize them and use them to advantage. And may be disastrous for those that do not.

Business process reengineering (BPR) is not a fad — it is a pragmatic response to combining the requirements for higher customer satisfaction and lowering overhead costs as one unified management issue. However, where many of the BPR gurus leave management in a state of confusion is that they make a call for action to BPR but then leave senior executives with no clue as to how BPR should be implemented. Aberdeen Group has found that when BPR is computer-driven — the capabilities and limitations are directly tied to computing technology — the implementation process becomes realistic. And HP is now the premier supplier of the components necessary for complex computer-driven BPR.

Audience Notes:

ACTION

Architecture

Components

Timing

Integration

Organization

New economics

AberdeenGroup

Aberdeen Notes:
To summarize the essential points required to successfully transform an enterprise information infrastructure, Aberdeen has created a simple acronym — ACTION.

The vast majority of computer-driven BPR projects that fail do so because IS management does not create a vision of the future with senior business managers. Aberdeen cannot overemphasize how important it is for IS executives to *take charge* of the process to create an appropriate information architecture for their enterprise that automates efficient business processes that provide superior customer satisfaction than the competitors can offer. With this architectural vision as a basis, the technical components, timing, application integration, organizational changes, and new economic realities can pragmatically be derived.

Audience Notes:

Architecture

Three-tier Plus in a distributed topology is state-of-the art

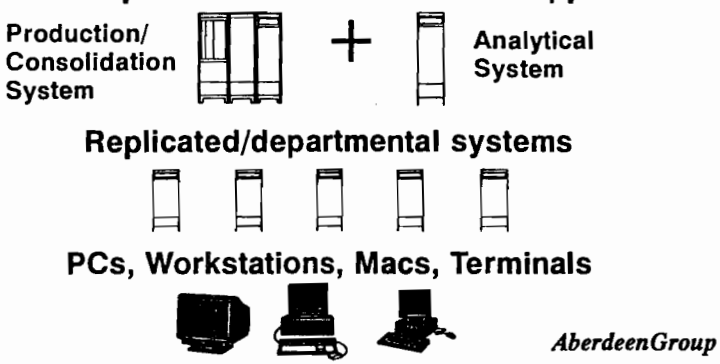

Aberdeen Notes:
Aberdeen's research shows that the three-tier plus architecture meets the needs of both users and IS professionals. Quite frankly, the major technical failures have occurred within enterprises that attempted to implement two-tier architectures — today's technology cannot support this overly simplistic approach to computer-driven BPR.

The key to user satisfaction is the emphasis IS puts on providing line middle managers with an analytical systems from which they can access current data to make realtime decisions. And IS can provide remote back-up, data integrity, security, and application version control through HP systems software operated by IS professionals from the enterprise's production system.

Audience Notes:

Components

Critical Technology Areas

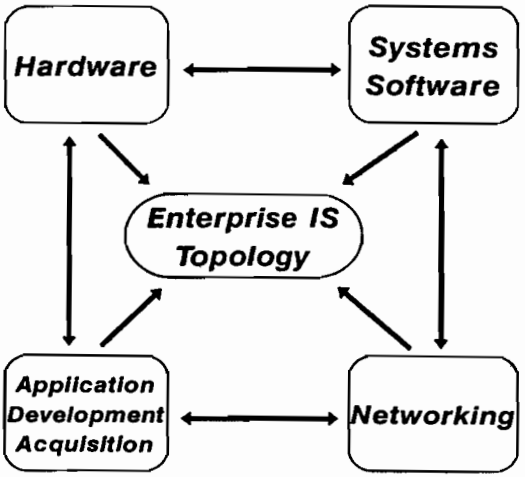

Multiuser versus Single user

AberdeenGroup

Aberdeen Notes:

New enterprise information architectures are created by blending technology components out of four main groups — hardware, systems software, application development and/or acquisition, and networking. By using open systems components that meet industry standards, enterprises are able to create a dynamic environment that can be improved as part of a continuous process over time.

At the box level, too many non-technical executives are naively confusing the single-user computing they are familiar with with multiuser enterprise computing. The two are very different and believing that PC technology can be used to run a coordinated enterprise will result in disastrous mistakes. And we all need more expertise in networking technology.

Audience Notes:

Timing

- **IS must start rightsizing before CEO**
- **Evolution or Revolution is real enterprise dilemma**
- **Business-critical applications rightsized first -- Back-office applications later Outsource mainframe applications**
- **Need 12-month timing chart to provide organization coordination**

AberdeenGroup

Aberdeen Notes:

Enterprise management expects IS to be proactive and initiate computer-driven BPR efforts within the enterprise — or it will surmise that the current IS organization is incapable of doing so.

While many enterprises have been attempting to manage to an evolutionary approach to changing their business processes, the trend is toward a full revolutionary stance — the business benefits are so compelling and the results of failure to rightsize quickly enough so enormous.

Applications that have the largest impact on the enterprise such as order entry and customer service are typically re-engineered first and then integrated into updated count-the-money back-office applications later. And to maintain internal agreement, IS needs to publish the schedule it is working towards or individual departmental managers may become puzzled.

Audience Notes:

Integration

- **Surround the mainframe with operational decision support systems**
- **Data conversion is a key issue**
- **Asynchronous data transfer -- time posting of data -- among applications is key new dimension**
- **Use open systems components -- today's new applications are tomorrow's legacy**

AberdeenGroup

Aberdeen Notes:

To quickly obtain tangible results from computer-drive BPR, surround the current mainframe systems with midrange systems with applications to implement your enterprise's new, more effective business processes.

Never underestimate the *management* effort required to convert current data into the new rightsized systems. Much legacy data is incorrect and yet organizations have found ways of working around the problem — but moving to new systems brings this hidden issue to the surface.

The open systems approach is the only method to alleviate the problems of needing to integrate even newer applications into the next-generation ones we are working on today.

Audience Notes:

Organizational Impact

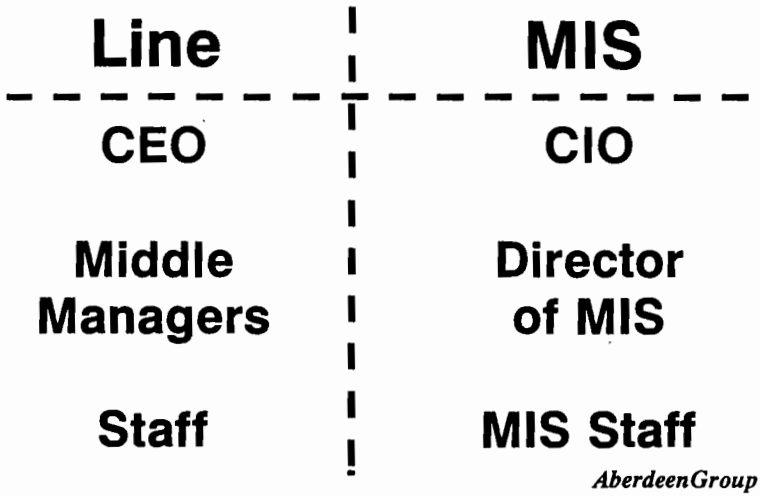

Aberdeen Notes:
There are six major groups with the organization that will be impacted by the computer-driven BPR. IS executives should ideally meet the requirements of each of these groups to ensure a smooth transition.

Of particular concern are the CEO who believes today that establishing the proper IS infrastructure is a significant part of this most senior job function; line middle managers who are demanding better IS tools to do their jobs more effectively; MIS staff, especially mainframe-trained staff, that know how the enterprise works but do not know the tools of computing and communications systems; and the CIO who is under increasing pressure to make business decisions, not merely technical.

Audience Notes:

New Economics

- **Maximize total enterprise profitability and effectiveness -- not just IS**
- **Prioritize IS functions by economic return to enterprise**
- **Only upgrade technical components when economic returns justify**
- **Continuously evaluate ways IS can improve business processes -- in all areas**

AberdeenGroup

Aberdeen Notes:

The datacenter is viewed by senior line management as a cost center that is a drain on the enterprise's resources — the old economics — and cannot support BPR. The New Economics focuses on the concept that an investment in IS will provide positive returns — IS is a way to lower costs and make money.

Almost all enterprises that have successfully reengineered their business processes report that many of the greatest benefits came from IS executives who reviewed operational aspects of the business with the intention of improving effectiveness through the flexibility midrange systems provide. In the new economics, IS executives take even greater responsibility for the financial strength of their entire enterprise.

Audience Notes:

Minimizing the Investment Spike

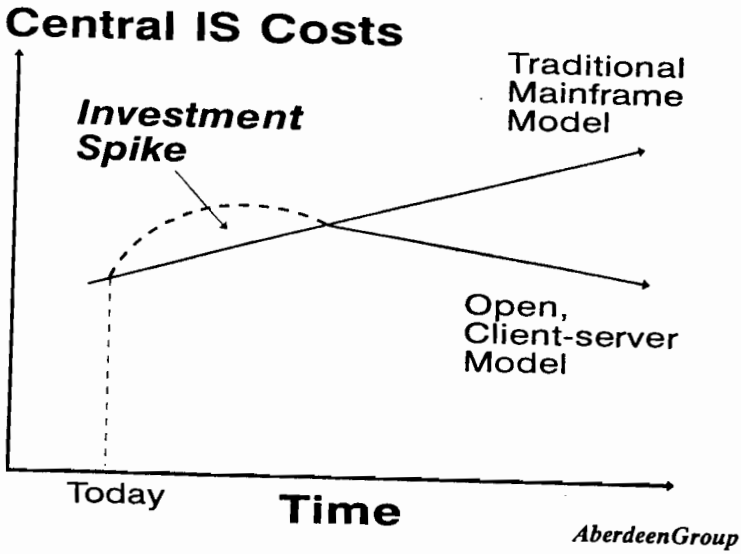

Aberdeen Notes:

The cost justification for investing in computer-driven BPR systems is that open, client-server systems will lead to lower IS costs than the traditional mainframe model.

Investment costs required include training, hardware, software, and professional services. The investment spike can be best leveled though the use of leasing vehicles.

Many enterprises that have actively pursued the BPR process for several years report IS costs do not decline — the enterprise finds the return of increased investments justifies expanded spending to both increase profitability and gain a competitive advantage.

Audience Notes:

Why HP's Role Has Changed

- **9000-800 is the successful Risc/Unix implementation for production applications**
- **9000-800 is the RDBMS platform-of-choice for new user-responsive production applications**
- **HP is profitable, growing, stable at top, and innovative. HP is more admired than IBM -- senior executives are proud to choose HP**

AberdeenGroup

Aberdeen Notes:
For executives who have been responsible for managing HP installations for many years, it is sometimes helpful to stand back and review HP as the industry sees it today.

Risc processors, the Unix operating environment, and RDBMS (relational database management systems) have been the technological promise for meeting *enterprise* needs for more user-responsive information systems. HP simply won the supplier battle to bring the multiuser platforms to market that best incorporated these technologies. And those multiuser competitors who would not or could not change like HP became mired in losses, continuous reorganizations, and technical indecision.

Audience Notes:

Why HP's Role Has Changed

- **HP has courage and capability to take on large, global projects with PSO**
- **Open Networking leadership**
- **3000 base has been upgraded with processors and systems software functionality**
- **HP's messages and business policies are closest to customer requirements for new multiuser enterprise computing applications** ***AberdeenGroup***

Aberdeen Notes:
In addition, Risc and Unix leadership are not the only reasons which have caused HP's role to change. HP's Professional Service Organization (PSO) has become more aggressive at taking on large, complex operations. And while IBM and DEC talk open networking, they are still wedded to their own proprietary networking standards, SNA and DECnet.

While many in the HP 3000 base may be concerned that their product does not get the HP corporate attention of just a few short years ago, outsiders see HP actually upgrading the MPE/iX to newer technologies as a positive contrast to IBM announcing that the AS/400 will have a 1995 transition that may be as drastic as the 1988 System/36 one (IBM stopped releasing high-end System/36s) and DEC leaving customers in a planning gap from VAX/VMS to Alpha/OpenVMS.

Audience Notes:

DEC

Product Strategy

64-bit Alpha with either OpenVMS, NT, or OSF/1 on all computing systems within the enterprise

Production operating environments are in flux

Significant Issues For Interex Members

64-bit Alpha applications will not be available until 1995

DEC is attuned mostly to the needs of its installed base during current headcount reduction phase

AberdeenGroup

Aberdeen Notes:
DEC has bet the company's ability to capture new accounts in the mid-1990s on the success of the 64-bit Alpha processor and OpenVMS, NT, and OSF/1. However, while these may be the underlying technology components, Digital has a considerable amount of work to complete before its operating environments, including languages and networking, are compatible with its new technology components.

Alpha's strength is its design for fast integer processing — its weaknesses are a lack of volume sales to justify expanded R&D spending to ensure long-term scalability and compilers and applications that take advantage of its 64-bit potential. While Alpha/OpenVMS will help meet the performance and price/performance needs of DEC's installed base, there seems to be little general interest on the part of vendor-neutral IS executives today.

Audience Notes:

IBM

Product Strategy

PowerPC processors and Operating System enhancements

Gerstner does not have a vision

Significant Issues For Interex Members

RS/6000 and AS/400 product transitions for 1995

Marketing messages are becoming fanatical as IBM reduces workforce dramatically

AberdeenGroup

Aberdeen Notes:
IBM's product strategy seems to change with every major product and reorganization announcement and sales call. IBM continues to reinforce the message that by 1995 all its major platforms (ES/9000, AS/400, RS/6000, and PS/2) will be based on the PowerPC processor jointly developed with Motorola. In addition, IBM is consolidating operating system development efforts to move more quickly to client-server computing and avoid the outright internal warfare that has been occurring.

IBM is taking both the RS/6000 and AS/400 through yet another technology transition to attempt to make them right for the rest of the decade. This is a major technical challenge for a company that has had so few successes recently outside of the PC market. And public perceptions may be clouded by an IBM staff that is now willing to say just about anything to avoid being part of the personnel downsizing effort underway.

Audience Notes:

Microsoft

Product Strategy

NT on Intel and Risc processors for server applications in 1994 -- replaces OS/2

Windows to Chicago on desktop -- and then Cairo (NT 2.0) on server -- in 1995

Significant Issues For Interex Members

Windows/DOS is what end users want on desktop

NT in 1994 could be a CLD

AberdeenGroup

Aberdeen Notes:
While Microsoft dominates the PC operating system world, its reputation has been tarnished by the mispositioning of NT for the desktop and slow start in the server market. The questions many IS executives are facing are whether NT is stable enough for production applications and will Microsoft support it long term in light of the announced introduction of both Chicago and Cairo in 1995.

Based on initial evaluations, Aberdeen Group is advising its clients to be very cautious regarding investments in Windows/NT in 1994 — MPE/iX and HP-UX will be more stable, robust and have longer lives.

Audience Notes:

NCR

Product Strategy

Workstation-to-datacenter multiprocessing Intel/Unix systems

Be multiuser cost leader through use of Intel processors

Significant Issues For Interex Members

AT&T cannot seem to decide whether to invest in NCR or spin it off

Committed to Intel/NT servers and development -- but not ready to go on NT marketing offensive

AberdeenGroup

Aberdeen Notes:
NCR is the dominant supplier of Intel/Unix systems for commercial applications. It is aggressively competing for new installations based upon low price which it claims is derived from using low-cost Intel processors and associated technologies.

AT&T has recently taken a very strong part in managing NCR strategies and tactics -- but seems stymied about what it wishes to do with this money-losing subsidiary.

Aberdeen believes that NCR is very well positioned to capture new accounts if: 1) NT on Intel is stable and obtains application software developer support for production applications and 2) NCR will commit publicly to supporting NT for the long term.

Audience Notes:

Sun Microsystems

Product Strategy

Multiprocessing SPARCserver/SPARCcenter servers interoperating with Solaris on client

Two-tier client-server architecture is sufficient

Significant Issues For Interex Members

Sun is still going through Solaris 1.X to Solaris 2.3 transition glitches

Still oriented to inexpensive workstations-on-desktop-connected-to-passive-servers business

AberdeenGroup

Aberdeen Notes:
Sun Microsystems is aggressively promoting its two-tier client-server architecture consisting of Unix on the client-desktop with upgradeable, multiprocessing power built into the server. The glue to make the client and server work together is primarily software provided by independent RDBMS and application suppliers.

Users continue to report that a considerable amount of technical expertise is required to support a Sun system as Sun is technically struggling with the transition between generations of operating systems and hardware platforms. In addition, the majority of end-users want an MS-DOS/Windows client on the desktop — a SPARCstation (even with WABI) has too many drawbacks (price, source availability, mobility) to be acceptable in 1994 for all but specialized applications.

Audience Notes:

Hewlett-Packard

Product Strategy

PA-RISC powering better than mainframe servers

Better than mainframe functionality in three-tier plus topology through leading systems software

Significant Issues For Interex Members

MPE/iX-Image/SQL versus HP-UX/RDBMS -- which, where, and when

HP is finally accepted by CEO decision makers

AberdeenGroup

Aberdeen Notes:
Hewlett-Packard has solved *all* the technical problems, especially I/O bandwidth, security, performance, and backup, that have in the past restricted the use of midrange systems to the departmental level.

The product reasons HP systems continue to be the BPR transformation platform-of-choice is the robustness of the systems software (middleware), scalability and power of PA-RISC, and extensibility of the entire architecture. But the question of where to use which systems is still left to the buyers' expertise and knowledge.

Now that HP is finally promoting itself directly to the CEO level within enterprises, IS executives responsible for HP installations should anticipate more requests for information about how HP-based computing is both different and better than traditional IBM mainframe methods.

Audience Notes:

Executive Interex Roles

HP IS experts are typically not business strategists

Executive Interex Role

- **Computer-driven Business Process Reengineering requires visualizing and explaining what is possible**
- **Knowledge of how company process could be improved ASAP is powerful**
- **Add value to enterprise management team**

AberdeenGroup

Aberdeen Notes:
Based upon all the above, many Executive Interex members will be finding that their roles will be changing.

The key challenge facing you is in the area of formulating improved computer-driven strategies and business processes for your enterprises and then clearly making recommendations to non-technical senior management.

Interestingly enough, Aberdeen research shows that the IS executives responsible for operating legacy systems from other suppliers are not even asked to be part of the BPR transformation. The user community merely assumes that these individuals have the *wrong* technical knowledge to be part of the change process while HP-knowledgeable executives have the saving grace of having chosen to acquire the right technical skills for today's needs.

Audience Notes:

Executive Interex Roles

HP remains CEO and new prospect unfriendly: "Great products, but how do you work with those people?"

Executive Interex Role

- **Expert at fitting HP product-orientation to enterprise functional requirements**
- **Business translator between HP, IBM, and DEC business policies and cultures**
- **Order entry expediter**

AberdeenGroup

Aberdeen Notes:
As HP continues to follow a business model whose results are excellent products but decreasing customer and process business-support satisfaction, Executive Interex members will find themselves being forced into the role of lowering their enterprises internal transaction costs and problems directly associated with doing business with HP

In addition, in those enterprises where IBM and DEC are also suppliers, you will find yourself in the position of educating senior executives about the differences in how all three suppliers work with customers and smoothing the transition where IBM and DEC are being phased out for new HP systems.

Audience Notes:

Executive Interex Roles

HP-UX versus MPE/iX

Executive Interex Role

- **Establishes right tool for right job**
- **Must understand DBMS and development requirements from a BPR viewpoint**
- **Strategic trade-offs must be taken into account**

AberdeenGroup

Aberdeen Notes:
HP corporate and the field sales force are heavily promoting HP-UX while the installed base and HP application Channel Partners believe that MPE/iX is the best operating environment for their specific needs. Even more confusing, we now see some ISV (independent software vendors) recommending that their applications be split between front-end HP-UX systems and back-end MPE/iX ones.

Executive Interex members are the internal experts at understanding the differences between these two operating environments and choosing which is best for their enterprises depending upon the strategic requirements.

Audience Notes:

Executive Interex Roles

HP and BPR threaten the internal organization

Executive Interex Role

- **Catalyst for change -- planner for obtaining objectives**
- **Mediator for unreasonable user expectations and MIS objections**
- **Role model for new IS support**

AberdeenGroup

Aberdeen Notes:
The changes created by computer-driven BPR and more wide spread use of HP systems threaten the positions of many individuals within an organization. Hard-earned knowledge and experience can quickly become irrelevant with confusion abounding about the future.

Within this environment of change, Executive Interex members will often find themselves being placed in the role beyond IS support for ensuring that the enterprise's objectives will be obtained. One very important aspect of this organizational change is the requirement that IS managers in charge of IS establish a new organizational structure and departmental procedures to support on an ongoing basis computer-driven BPR.

Audience Notes:

Executive Interex Roles

Acquisition, operational, and upgrade cost uncertainty

Executive Interex Role

- **Establish realistic expectations for cost and capacities of HP systems**
- **Base recommendations for systems implementation based upon enterprise ROI**
- **Be cautious of IBM-originated estimates of HP costs**

AberdeenGroup

Aberdeen Notes:
As shown in the illustration of the investment spike, enterprises that engage in computer-driven BPR to leverage their industry megachanges really do not know what the cost of a new information infrastructure will be. Even worse, many mainframe-oriented third parties will make a point of projecting unrealistic high costs to move to a three-tier plus architecture and not be able to identify the real internal productivity and customer satisfaction gains that can be expected.

Executive Interex members will find themselves in the role of estimating both realistic costs and enterprise returns for individual BPR projects. Senior executive decision making becomes very pragmatic when IS executives describe several computer-driven BPR opportunities in terms of total costs, expected returns, and recommended implementation priority order based upon business (not technical) analyses.

Audience Notes:

Wrap-Up

- **MegaChanges and excellent HP engineering have converged to make HP and HP expertise desirable**
- **There is a prospect backlash against HP's business model -- but no strong competitor**

AberdeenGroup

Aberdeen Notes:
HP is in the enviable position of being the leading hardware product supplier of multiuser systems for *new, complex, large* enterprise applications. While each year one expects to see IBM, DEC or another multiuser supplier reinvigorate its efforts to challenge HP, none has. As a result, HP expertise remains in high demand within the industry.

The major downside to HP's success is that many prospects expect high levels of support and service bundled into HP's platform prices — and then they expect HP's prices as measured in terms of price/performance to be the lowest in the industry. Note that no competitor has been able to profitably take advantage of this market discontinuity.

Audience Notes:

Wrap-Up

- **Interex executive members must add value to their enterprises by blending HP knowledge with business and financial skills**
- **Performing *all* the ACTION steps is a requirement for successful computer-driven business process reengineering transformations**

AberdeenGroup

Aberdeen Notes:
Executive Interex members will increasingly find themselves under pressure to make decisions and recommendations based on strategic and financial terms, not just technical ones. The knowledge of how HP systems operate and how to do business with the company gives Executive Interex members unique, and very desirable, knowledge within the IS community.

When helping your enterprise make the transition to the next level of business competition through computer-driven BPR, following all the ACTION steps becomes critically important. Aberdeen has found that success is obtained by those IS executives who complete all the steps and do not take short cuts.

Audience Notes:

Paper Number: 1001

THE H.P. INFORMATION SYSTEMS MARKETPLACE TODAY: WHAT'S HOT AND WHAT'S NOT!

Diane Amos, CPC
Amos & Associates, Inc.
633-B Chapel Hill Rd.
Burlington, NC 27215
910-222-0231

In today's fast paced, rapidly changing Information Systems market, the lines of technology are blurring. "Open Systems" may literally mean opening the world of technology for companies to meet their needs in a better way.

How does this impact a company's ability to staff its' I.S. department? How can an I.S. professional stay on top of this wave of technology and not be swept out to sea? This presentation will explore the challenges presented by the many changes in Information Systems today.

The biggest request that I receive is for information relating to "what's hot" in the technical arena. I.S. professionals want to stay current with technology so that their careers can progress. From an employment perspective, let's look at what's HOT and what's NOT!

Let me first offer a disclaimer. I am not a technology expert. The information that I'm sharing is what I hear from the unique view that I have as a recruiter. This is information that I have gathered from managers who I conduct searches for and candidates who are seeking new jobs.

I. Skills in Demand...

A. Relational databases are the way of the future for the Hewlett Packard marketplace and everyone is jumping on this wagon. Companies that develop these relational databases are well represented here at this

conference in the exhibit hall which offers a wonderful opportunity to compare the various vendors. From a recruiting standpoint, we've seen the heaviest requests for people with skills on the following relational databases: Sybase/Oracle/Allbase/Informix/Progress.

B. Client-server tools such as Powerbuilder and Uniface are also in demand. Interestingly enough, everyone is talking about "client-server," but few companies can really define it. It's like arguing over what is truly a relational database. Everyone has a different opinion about "client-server," as evidenced by the many different talks at the last INTEREX show. But more importantly, few companies are really "there" yet. This causes a real shortage of experienced talent with those skills that are in demand.

C. PC's are springing up everywhere! With the downsizing of mainframes, more companies are turning to PC's and need highly qualified PC talent. I don't mean computer people who are PC "users." The demand for PC talent is in several different categories. We are seeing an increased demand for Windows' developers. Office automation specialists who can train users on the multitude of PC software are requested frequently. A PC specialist is often times expected to get into the guts of the hardware and diagnose or repair it, especially in small shops. And finally, we see a major demand for those individuals who can hook PC's together into networks. Which brings me to the next skill area in demand...

D. Networking! Access to information quickly is one of the most vital ingredients to having a competitive edge in today's marketplace. Networking computers together across the world through WAN's, throughout remote offices with LAN's is what will change the future. Networking expertise is definitely in demand with such products as

Novell/Netware leading the way. A "certified network engineer" -CNE is the most in demand and can therefore expect the higher wage.

E . With Hewlett Packard selling UNIX boxes like hot-cakes (60% of their market), obviously HP-UX operating systems skills are in demand. The staffing problem that can be created by this is that MPE customers who are migrating to UNIX would like to hire someone with both MPE and HP-UX skills. Those few employees that have experience on both operating systems certainly have the edge. But remember, companies on MPE that have been running their business applications on MPE expect to continue to do this on UNIX. They will not hire someone with just UNIX technical skills. They need the business acumen as well.

F. Systems integration is the hot buzz word now. Those candidates that have multi-platform experience and a know-how to integrate the various systems in a company will have an edge. With open systems technology today, companies may have multiple databases, hardware, and operating systems. Someone who knows how to integrate this all seamlessly will be worth their weight in gold. With HP making inroads into the IBM and VAX markets, many companies are having to migrate from VMS and MVS to HP-UX. Conversion experience is also highly valued.

G. 4GL's continue to be popular with a demand for the standbys Powerhouse and Speedware but we also see new kids on the block in the HP market: Progress, Powerbuilder and Uniface. In the MPE world, few companies are solely using 4GL's, therefore, there is usually a need for COBOL skills in addition to the 4GL background. C and C++ languages are strong in the UNIX world, which leads us to object oriented technology....

H. We're already being asked to provide candidates with experience in object oriented programming and design using related products such as Smalltalk. The supply has not yet reached the demand, and I predict the demand will be increasing.

I. As more companies look to streamline, I see more of them turning to 3rd party application software rather than custom writing their own systems. In the manufacturing arena, ASK is a big draw with some requests for AMAPS and WDS as well. With HP making inroads into the mail order business, several mail order packages are popular. Accounting software is quite prevalent with many choices to be made. For that reason, typically the request is for "some" type of accounting package experience, but not requiring a specific package. The manufacturing industry is the most demanding and specific industry for requiring exact matches on the manufacturing software that is in use.

J. Surprisingly, there are still companies that are seeking what I call the "vanilla" programmer -those with skills in COBOL/IMAGE/VPLUS. They don't care about 4GL's or relational databases, and run just fine the way they are, thank you. For those of you who enjoy this environment, there is still a place for you. But for how long?

II. Skills not in demand...

A. HP1000 professionals are having a hard time of it. Those that have not had the opportunity to migrate to HP9000 are stuck in careers and shops that are either going to stay that way or disappear. I haven't seen a request for an HP1000 candidate in 3 years and I have a drawer full of candidates looking for a way out of that market. My only suggestion is to get training fast on other hardware.

B. Anyone who is just a coder or just an analyst is not in demand. A programmer today is expected to be able to do some analysis, and an analyst is also expected to program. The days of the pure coder cranking out code in the back room are gone. There is also little demand for the pure analyst who designs systems for users without seeing their system through to implementation.

C. Techies with no interest in business will have a hard time in the marketplace of the 90's. Today, an I.S. professional needs to have business acumen. They need to see the big picture and the role they play in it. They also need to be effective communicators with the end users and understand the business issues that the users are trying to solve.

D. Certain industries have taken a downturn and I.S. professionals with narrow industry experience may find themselves boxed in a corner. A good example of this is the programmer with defense and aerospace industry experience. With the cuts in defense, those jobs are going away and the experience doesn't translate well to private industries that want business systems experience such as manufacturing or accounting.

III. Soft Skills

Technical skills can be learned. But today more companies are putting emphasis on what I call "soft skills" over technical skills. Their feeling is that "soft skills" are more inherent and not easily taught on the job.

What are these "soft" skills?

They are areas such as:
-customer service orientation
-people skills
-business acumen/savvy
-effective listening and communication skills

Let's break this down so we're clear in what we mean.

A. "Customer service orientation" means an attitude that says give the customer what they want. It's an understanding that Information Technology truly exists to serve the user. An I.S. professional with this expertise can work with a user to clearly understand their needs and translate them into a workable solution. This kind of I.S. professional generally has a lot of patience with the non-technical user, and can talk at their level adapting as necessary. With I.S. becoming more user friendly and users taking a more active role in the integration of business and I.S., this skill is critical to companies in the 90's. Practically every search that I have, discusses the importance of this skill. Gone are the days of separation between user and I.S. professional and the ivory tower has disappeared. Adapt, learn how to do this, or your career may falter.

B. "People skills" can mean many things and the vagueness troubles those who try to define it. What I understand from my clients who request this is that they want to hire people who get along with other people, are outgoing, dynamic and aggressive, expressive rather than analytical. This area may bode trouble for the typical computer professional who is the "analytical" and "introverted" type. Skills that need to be developed are empathy for the other person, open-mindedness and flexibility for other points of view, an ability to reach out for understanding rather than standing ground in conflict. I know several computer professionals who are analytical and introverted who have learned, that to progress in their career they needed to develop this side of their personality. It can be done and is another skill that companies are stressing in the 90's.

C. "Business acumen" means an ability to see the "big picture" of what a company is trying to accomplish, it's mission, goals and strategies and not just their part in the play. Many I.S. professionals have not been allowed to play a part in business decisions before now, but with companies embracing the "empowerment" concept, they not only want to give this decision making power to their employees but they expect their employees to handle empowerment successfully. I.S. departments are now often on equal footing with other departments in a company and have the President's ear on how to help run the company. A technical person who only cares about the bits and bytes of technology will not be asked to assume this responsibility. This responsibility is not just at the top either - it is filtering down through the ranks with each I.S. professional playing a part. Learning how what you do that impacts the work of others and how it fits into the overall business scheme is crucial. ASK those critical questions - this is how you can learn this "soft" skill.

D. "Effective listening skills" will not only get you that desired job but help you to be successful in it. This is occasionally a weak area for I.S. professionals while interviewing. They are concentrating so hard on their answers that they don't always listen carefully to the question to make sure they understand it. The same goes for daily communication with peers and users. The best user reference I can get is when they tell me that my candidate took the time to listen to their needs so that they really understood. The trick I believe, is to concentrate on the speaker, repeat your understanding of their statement or question, and through continued feedback, make sure that you are hearing what they are saying. It's an easy procedure and with practice becomes almost commonplace. Feedback is the KEY.

IV. Future trends

People ask me every year what the future holds. This is a fun aspect of my job, and this is my time to play Jeanne Dixon. Sometimes I'm right on the money and sometimes I'm not.

Overall, I see a continuance of companies running lean and mean. Because of this, companies are asking for almost perfect matches in both skills and personality.
-I believe that HP-UX will be a major player in the future, however the MPE shops that choose to remain MPE also will do well.
-The most in-demand job will continue to be the 3-5 year application developer.
-Opportunities for managers will continue to diminish due to rightsizing and reengineering.
-Object oriented technology should play a major role in the future as well as client-server, networking technology.

Technology will not just be in the workplace, it will be in our homes and play an increasing role in our lives. We're all hearing about the "information highway," which we need to follow closely. Because of these future trends, I.S. professionals need to be open, flexible and adaptable to learning new technology. Those who do not may find themselves out of a job.

V. Importance of training...

With the demand for newer skills, the need for training has never been more important. Sadly, this is where most companies drop the ball. They expect to achieve their goals without investing in the training to get them there. So often, if they would just send their eager, talented people off to be trained on the "latest-greatest" rather than upset their internal salary structure by hiring the outside "guru," they would see their goals achieved and create a loyal staff besides. If they have to hire that GURU, they should utilize the GURU to train the others. Training is the KEY! The keynote speaker at San Francisco's INTEREX stated that a manager explained that he was fearful to train his staff because they might leave his company to better themselves. The speaker's reply was, .. "and what will happen to your company if you don't train them?"

If you need some justification to add to your training budget, try using these statistics to convince management. (Taken from the American Society for Training and Development).

1. Over the past 50 years, investments in job-related training have consistently outperformed investments in machine capital. A trained workforce produces more productivity gains than capital investments do.

2. The average 1.4% of payroll that U.S. companies invest in training reaches 10% of the workforce - Japanese and European-owned companies based in the U.S. spend three to five times more on employee training than American companies do.

3. By the year 2000, 75% of all workers currently employed will need retraining -technological and societal changes are reshaping jobs in every industry.

4. 70% of the people who will be working in the year 2000 are already in the workforce - the only way they will have access to training is if companies like yours supply it.

5. The surplus labor supply of a few years ago has dwindled to almost nothing - companies won't be able to buy talent, instead they will need to create it with training.

VI. Issues for companies...

A word of caution is needed in the midst of this frantic rush toward newer technology. Take a good look at your company's needs and how you are able to meet them. If your system does a great job in meeting those needs, is it smart to jump on the bandwagon just because other are? I'm seeing many companies take the leap without looking where they're going, and then having to backtrack because they impulsively jumped into a technology that they weren't prepared to handle or that didn't truly meet their needs. There has to be a balance between a company's needs and new technology.

On the other side of the coin, getting behind in technology may cause you to lose your staff. Think through your strategies and objectives carefully to weigh out where technology fits into your company's goals and mission.

VII. Issues for candidates...

Look at how you can become a well-rounded business professional, not just a technically competent computer individual. Remember my comments on companies hiring "business" oriented people who can communicate effectively. Develop your "soft" skills! If you are a technically brilliant I.S. professional, but have some difficulty in relating with end users and your peers, there is help out there for you. The world is loaded with training workshops on "developing better listening skills", "dealing with difficult people", etc. If you can't take the time off of work to attend, there are numerous books and tapes that you can study in your spare time at home. I believe that "soft skills" can be learned, so embark on your own self improvement program. Remember, they hire you 75% because they like you and 25% because they think you can do the job.

Flexibility and adaptability to changing technology will be critical for the future of your career. Those of you that want to stay with the "familiar" may lose out as everyone else moves ahead. Be willing to learn new things, get trained and adapt to change.

In summary, we've discussed the technical skills that are in demand and not in demand, explored the area of "soft skills," looked into our crystal ball for future trends, and recognized the importance of training. If I could give you one piece of advice from this presentation that I believe will help you the most, it's...develop your "soft skills."

Paper Number 1002

Effective Management of Multi-Year System Development Projects

Curtis J. Figler

Martin Marietta Astronautics
P. O. Box 179
Denver, CO 80201
(303) 977-6713

The justification, funding, design, development, testing and implementation of computer applications requiring a year or less to complete is relatively straight forward. Once funding is justified and allocated the overall development effort is fairly stable. However, when a development project spans a number of years, it becomes a much more challenging proposition to secure funding, sustain interest, maintain the commitment of management/users, cope with changes and effectively manage the inevitable shifts in priorities.

This paper will explore the approaches that our project team has employed to assure that our long-term development effort, which began in 1988 and will be completed in 1996, succeeds.

Project Description

The system we are developing is known internally as the Equipment Management and Accountability System (EMAS). EMAS is an integrated system of management applications for the accountability, control and maintenance of equipment and real property assets. It tracks equipment and real property assets from its initial procurement action through receipt, placement in service, movement, storage, maintenance and calibration, identification as excess and final disposition.

When we began the project in 1988 we were saddled with four separate systems. Although they all managed the same assets, none of the systems talked to each other. This lead to major reconciliation problems with records existing on one or more systems and not existing on others or with critical data, like the location of an item, not consistent across the systems. Since the systems were used to control both company and Government-owned equipment, the inability to reconcile the systems and provide users with current and accurate information was a major concern.

The initial thrust of the development project was to link the four systems on a real-time basis, reconciling the records within the network and sharing critical information. A number of smaller stand alone applications were then rewritten and integrated into the network. While this answered the immediate concerns driving the project, a user was still required to log into each of the four systems in the network to gather all critical information regarding an equipment item. It was recognized early on that the optimum solution was to build a single application to support all aspects of equipment management, to provide users with a single resource to work with.

In terms of size, the systems control over 125,000 line items of equipment valued at over $1.5 Billion and located at multiple locations from the West Coast to the East Coast. The networked systems processed 2.5 Million transactions per week, supporting 100 concurrent users and processing 2,500 batch jobs weekly.

Organization

All projects involve and ultimately depend on people to succeed. The key to success is to bring all of the people involved in a process together in such a way that they are all a part of the decision making process. We used the High Performance Work Team concept as our model for forming a project management team. The HPWT is made up of three parts: the Steering Committee, the Working Group and a contingent of Subject Matter Experts.

Management personnel from each of the organizations dependent upon the equipment management systems were named to a Steering Committee. The Steering Committee's purpose is to facilitate the process and support the development team in breaking down any barriers that might crop up and hinder their progress. It acts as the buffer between executive management and the development team, providing status reports, working budget and resource issues, and setting the goals, priorities and focus of the project. The Committee is not involved in the detail of the development.

The Working Group consists of individuals from each of the organizations dependent upon the equipment management systems who are close to the actual work being performed and have a more detailed knowledge of the various processes involved. The programming staff are also members of the Working Group. These are the individuals working the detail of the development project: defining and streamlining the

processes to be supported by the system, designing the applications, doing the programming, testing, documenting and implementing.

The Working Group is supported by Subject Matter Experts (SME). SMEs are the folks who actually do most of the work and can shed light on the challenges they face in dealing with existing systems and processes. They can explain to the Working Group how tasks are actually done on the floor, which is often quite different from how management perceives they are done. We're talking here about tool crib attendants, stockroom personnel, calibration technicians, expediters, machine operators, etc. - the people who really make it all happen.

The key to the success of an HPWT is a shared vision, enthusiasm, trust, compromise and consensus decision making. It takes the active and positive participation of all three pieces of the organization to assure success.

Bottoms-Up vs. Top Down Development

I've visited many companies and been exposed to many instances where a management system was designed by management and imposed on their personnel. In such a situation, its usually pure luck if the application reflects the reality of the processes it is intended to facilitate. The only people who really know how work is done are the people who do the work. It is imperative that they design the system.

Working with the "pick swingers" assures access to process details that management, being somewhat insulated from such processes, cannot provide. It assures that the system reflects reality, not perceptions.

Another major problem with top-down designed and imposed systems is that those who are forced to use them don't have a vested interest in their success. Assuring that the "pick swingers" design the system fosters user ownership and future support.

Nurturing the Development Team

How do you keep a loosely organized team together, moving in the right direction and enthusiastic?

First, you need to share the vision. Make sure everyone on the team has a solid understanding of the big picture: where we started, where we are, where we're going and why. The vision is not imposed from above or the outside, but grows from within the team. The team also needs to be willing to alter that vision, to change priorities as the needs of the user

community change. While implementation of a short-term development application may be revolutionary, a long-term project is evolutionary.

While each individual must be willing to work within a team sharing environment, they also need to know what's in it for them personally and for their specific organizations. While its gratifying to succeed as a team, its also important that each member be rewarded individually in some way.

The team must be open to participation by anyone who has an interest. People should be given the opportunity to join the team or leave it at their or their organization's discretion. The last thing a team needs is a bunch of disgruntled folks who are there only because they have been directed to be there. Active and enthusiastic participation will be maintained if the team makes continuous progress and benefits are realized by each member and their individual organizations.

Everyone must share in the decision making. The team must be willing to work through disagreements until a consensus is reached. Members must be willing to listen and to compromise based on what's best for the entire team, as opposed to what's best for them. Such compromise must also be recognized and rewarded.

Its also important to keep the mood light. Participation should be fun, not drudgery. While the work being performed is serious, its equally important to insert some levity from time to time.

Securing and Maintaining Funding

First, you need to share the vision with the individuals or committees who hold the purse strings. They need to share a basic understanding of the problems being addressed and the big picture of where the team is trying to go with the system being developed. Instilling in the "funders" a level of enthusiasm and anticipation regarding the system which is near that demonstrated by those on the team is essential to assuring sufficient funding is allocated each year.

Know your benefits and demonstrate them often. Most funding justification exercises involve putting together a cost versus benefit analysis. The team must be willing to discard any aspect of the overall project that does not pay for itself within an acceptable period of time. Once the benefit of a piece of the project has been identified and quantified, it is important to demonstrate that the desired level of benefit is achieved. Many times, cost benefit is a key point in justifying proceeding with a project but the decision makers don't often come back

to check that the initial estimates were close to reality. Gathering data on benefits from implementing each portion of a long term project and voluntarily reporting back to the decision makers fosters their trust and respect and lays a solid foundation for successfully obtaining future funding.

A long term project also needs to be broken down into logical pieces and implemented piece-by-piece. There is no company willing to invest money in a project over a long period of time without reaping some benefits from that investment. The development team needs to identify each logical piece and then prioritize them based on a solid business case. Companies are interested in getting the most bang for the buck as soon as possible, so the team must be willing to meld their plans with this desire.

The project plan must contain reasonable milestones and schedules. Staying on budget and on schedule is more critical to long-term projects because each subsequent piece of the project is dependent on the successful (i.e., on budget and on schedule) implementation of the previous piece. The worst thing a team can do is set an overly aggressive plan and schedule and then constantly move completion dates to the right. Set a reasonable date, stick to that date (despite the almost constant pressure from management to move it in) and then make the date.

Constant public relations (PR) activities are also a key to a long term project. The decision makers need to feel comfortable that the development effort isn't simply a long expensive science project. The team needs to schedule frequent demonstrations to provide executive management with something they can see and to maintain confidence that progress is being made. Any development plan for a piece of the project should include on-line demonstrations at key times during the development process. Our funding process for an upcoming year usually begins during the fourth quarter. We therefore tend to schedule demonstrations during the third and fourth quarters of each year to help influence the funding process to our advantage.

Demonstrations provided to the user community are also important to the project. They provide an opportunity to share the vision with the "pick swingers", solicit valuable feedback from individuals who may not be involved in the details of development and stimulate grass roots support for the project. Sustained interest and support of the project from management and from the user community are equally important in maintaining funding.

Last, but certainly not the least, the team must celebrate and communicate its successes. Each implementation of a system module should be publicized in activity reports, through flyers and company newsletters. Our team has organized what we call "ribbon cuttings" to ceremonially activate new modules in our system. We invite upper management to participate, along with the development team and the user community. The ribbon cutting calls attention to the new system application, provides positive recognition for the development team and also provides a great photo opportunity to publicize the team's success in the company newspaper.

Retaining Programming Resources

It goes without saying that without programming resources, system development is impossible. The key to the success of a long term project is consistency and continuity, which is only achieved by maintaining a core of programming talent throughout the development process--keeping good people involved and energetic for the entire term of the project and beyond.

We have four or five programmers who have been with the project since the very beginning and expect to be with us through final implementation and beyond. We've been able to keep them despite lures from the outside (i.e., other projects and other companies) because we have established a living, growing application routed in continuous improvement. We've endeavored to establish an environment that embraces change, rewards innovation, accepts diversity, admires hard work, integrates new technology and provides a constant challenge. Our ability to virtually guarantee a stable level of funding, which translates into job security, also contributes greatly to our ability to retain the best of the best.

Embracing Discontent

Thomas Edison once said "Discontent is the first necessity of progress". Although Mr. Edison made this statement long before computers were even a gleam in an inventor's eye, embracing discontent during system development is essential. Complaints regarding our system do two things. First, they are a good indication that the complainer is trying to actually use the application and cares about it. Second, complaints drive enhancements that the development team may not have considered, enhancements driven by real world situations. Quick attention to minor complaints, within the context of existing priorities and resource availability, also creates solid allies at the grass roots level. Fix a small

problem that makes life a little easier for someone and he or she will most likely become a dedicated and supportive user.

Management of a short term project requires that the bulk of the attention be placed in the technical aspects of development. Once funding is concerned and a schedule is established, a dedicated team works to a single end. When a development project spans a number of years, it becomes a much more challenging proposition to maintain a common vision, justify and secure annual funding, sustain interest, maintain the commitment of management/users, cope with changes and effectively manage the inevitable shifts in priorities. Keeping a team enthusiastic, active and productive over the long haul is much more challenging, but each small success also seems to be much more rewarding.

LEASTERSHIP - ANTHONY FURNIVALL

SDL/SOFTWARE, INC. 716/883-2667

ICMS 94 - NEW ORLEANS

PAPER 1003

It is certainly difficult to be a leader without having a high sense of your own worth, and confidence in your own abilities. However, just as those intrinsic ego qualities make some of the aspects of leadership easy, they can also provide some of the pitfalls and problems which confront leaders. Wouldn't it be great if we could get our egos out of the way of our efforts?

"Leastership" is a concept of leadership which allows you to do just that. This paper will show you some of the ways in which you can channel your ego-energy into constructive and creative paths, and minimize the problems. By examining four key concepts, you will learn how to make your team responsive, focussed, and above all effective!

"Leastership" was originally delivered at the 1993 Leadership Training Conference for Interex Volunteers.

LEASTERSHIP

This paper was originally intended to be a paper on Time Management. At least that was the request I originally received. However since that would have revealed itself very early on as yet another case of the blind leading the sighted, I was able to persuade my sponsor to settle for the ideas that I plan to share with you in this paper. You may well be more sighted than I, but at least I won't feel like I am stumbling all the time!

In some ways leadership is like pornography - we don't know how to define it, but we surely recognize it when we see it. I'd like to share a couple of experiences with you, and see if we can agree that one represents a positive style of leadership, while the other is perhaps less so.

I was working, once, in a part time situation at a radio station, where I was required to record various programs which came off a satellite feed. The recording apparatus was not technologically sophisticated, but it had enough potential gotchas to create problems. The tapes on which I was to record, came from another part of the building; all I had to do was load them and press the right buttons to schedule the recording. Needless to say, I screwed up one time. So I went to my boss, and explained the situation. He spotted something that I had not noticed, which was that the tape was the wrong size. This meant that even if I had managed to get the recording going properly, it would have failed half way through. When we called the people who needed the tape, my boss placed all the emphasis on the fact that the tape was the wrong size - thus demonstrating that he was supporting me, even though it had been my mistake originally. He did not draw any attention to this fact, but I certainly noticed it.

In the other case, a friend of mine was working in an environment where there was a lot of organizational politics. Despite being given a free rein to implement whatever changes were necessary, and being told this several times face-to-face, it became very clear that my friend was not being supported during the political squabbles which would arise. A proposed solution would be endorsed in a one-on-one meeting, but then at a later meeting (where my friend was not present), the original agreement would be changed. To make matters worse, the changes would not always be communicated back to my friend. Thus all the appearance of support and backing would vanish like an early morning mist when the real political issues were on the table.

I feel certain that you can identify the good and bad examples, but how can we define the difference between them? At the outset we must remember that leading is

different from doing. As some-one once remarked "He is not the best statesman who is the greatest doer, but he who sets others doing with the greatest success." Learning not to be a doer, but to help others become more effective doers is undoubtedly the hardest lesson in becoming a leader. Here are three ways which I have found personally effective, and I pass them on for your consideration.

LOOK

First, learn to look. Look at the goals of the organization in which you are a leader. Identify the one which has the strongest personal appeal for you. If you can identify such a goal, then it will be easier to adopt it and make it your own. Look at the goals for the team you are leading - if this team has been charged with specific goals, see if you can again identify the one which has a strong appeal for you.

Look at the goals of the members of your team. They all have the same sort of personal aspirations as you - but they all have subtly different aspirations. Learn to look at what they are striving for, and see how you can make all these goals mesh together.

Lastly, look at what you can do to make a difference. This will allow you to identify the things which will be called successes. People like successes, and to be able to identify a success will go a long way towards having people follow you.

Remember that despite any of the more grandiose definitions of a leader, the most pragmatic definition of all is that a leader is someone who has followers. If you can find a goal that is shared by your team, and identify a way to reach that goal, then it is much easier to persuade people to follow you. Leadership is about helping people adapt to change. This is not necessarily a pleasant or painless proposition. Given a choice between accepting a change, and the opportunity to show why the change is un-necessary, most people will immediately begin on the second task. The more you can do to make it easy for people to want to change the easier your task will be.

LISTEN

Learn to listen. Listening is now regarded as one of the most critical skills in communication. Listening means letting your own part of the communication process wait until the person who is speaking has finished. It means considering all of the aspects of the communication - not just the words, but the tone of voice, the emotional content of the speech, and the body language of the speaker. It is undoubtedly a skill which can be practised, and learnt.

Listen to your team members. They will give you all the information you need to lead them and manage them effectively. This information is sometimes given

away explicitly, and sometimes by listening between the words. However, it is almost always given. All you have to do is to listen for it, and pay attention to it.

Listen to your leadership colleagues. They have probably travelled the same path that you are on, and may have memories, advice or warnings from their own experiences. Or else they may be on a different path, but still having the sort of experiences from which you can benefit. Leadership is frequently a lonely position. You need and deserve all the support, ideas advice and encouragement that you can lay your hands on.

Listen to others. This has one very important and beneficial side-effect - you can not be forever doing if you set listening as a priority. In Shakespeare's play Henry V, the night before the battle of Agincourt, Henry wanders around listening to the mood of his soldiers. With a simple word here and there, he lets his men know that they are important to him, and that he considers them his brothers:

> For he today that sheds his blood with me
> Shall be my brother, be he ne'er so vile.

LET

Lastly, learn to let go. Let go of all the pressures that can contribute to a failure in leadership.

Let go of doing it yourself. In a technical environment such as MIS you were probably selected for a leadership position because of technical competency, and expertise. This makes it even harder to let others take care of technical tasks when you could do it twice as fast, and ten times as well. Letting go of doing it yourself is the first step towards providing the leadership that the other technical people need. Remember the importance of setting others to doing.

Along with this, let go of doing things the right way. There is frequently no such thing as the right way. An example from the stock-market helps illustrate this point. An economics class tracked the performance of a portfolio whose make-up and transactions were determined solely by the throw of a dart, rather than any so-called market expertise. The portfolio outperformed a "well-managed" portfolio by a factor of 10 to 1. This is a truly humbling experience for some one who expects to make money by doing things the right way.

The simple fact is that there is frequently more than one right way, and if the right way really means your way, then you must remember that you are not going to be the doer! It is important that you give your people the opportunity to succeed and fail on their own terms.

This brings up the all important issue of trust. Trust is never easy, and is

meaningless unless there is real risk involved. Failure is a painful possibility, but the most important learning experience of all. It has been said that the greatest part of courage lies in having done the thing before. If this is true, and I suspect that it is, then how much more important is it that the first time we try a new endeavor, we are given the freedom to fail.

Failure need not be a foregone conclusion, but unless we can let our people have that opportunity, they will never have the corresponding opportunity of displaying the courage which comes with doing something twice! This is not to say that we should become fool-hardy in our risk assessment. Such an attitude is imprudent at best, and possibly fatal at worse! No - rather we should allow for the joint assessment of any risk, and then the mutual decision about what path to take. People who participate in the decision thinking and decision making will frequently be far more cautious and conscientious about the implementation of a decision. After all, they also have something at stake!

LEASTERSHIP

Most of the ideas I have presented here are incredibly simple to articulate. Indeed it would be very easy to assert that they are so simple that they almost need not be spoken aloud. But while it is a gross oversimplification to say that the simplest ideas are the best, the fact that **KISS** has become such an important acronym speaks volumes to the viability of these ideas.

Central to the whole concept of leastership is that your own ego needs to be channeled into the most productive paths. It does not mean that you must subordinate your ego, or ignore it. Rather that you can use it most effectively in ways that are perhaps less direct than in the past. By facilitating the successful accomplishments of others you acquire for yourself the mantle of successful leadership. A long time ago this was expressed more eloquently than I could ever hope to achieve:

"He who would be greatest among you, let him be least"

The Color of Change is Grey

Robert M. Gignac
Speedware Corporation
150 John St. 10th Floor
Toronto, Ontario
M5V 3E3

Abstract:

The dynamics of our day-to-day business environments are such that change appears to be the only consistent thing we face. In our profession we are expected to embrace and absorb change at an ever-increasing rate, and at the same time, take this change to the rest of the organization and encourage them to embrace it as well. This is not, nor will be, an easy task.

The purpose of this presentation will be to: 1) identify the criteria surrounding change that cause such problems when we try to implement; 2) separate 'good' change from 'bad' change; 3) identify the audience upon whom we try to implement change to determine the best possible method; and 4) find ways to minimize the impact of change on all parties involved. The more we understand about the dynamics of change and its impact on both ourselves and our environment, the better equipped we will be to introduce change to people who may not want it.

So what is change - Really?

Change. Ongoing. Eternal. The only constant. It is the major factor in what we as IS professionals do for a living. We live to change things. We change the way people work, we change their equipment, and through our desire to continually reinvent ourselves as professionals, we adversely (and sometimes constructively...) change peoples' lives.

Change hurts. It really does. This fact is the first and foremost law of change. Change takes what is known and familiar to us and turns our existence into the unknown and unfamiliar, it takes that with which we are competent and turns us into bumbling idiots and can, when handled improperly, make us look like fools. Is it any wonder that change hurts?

The interesting thing about change is that change doesn't ask for volunteers and it doesn't ask for permission. What drives this change? We do. As technologists we continually strive to find different (don't read that only as 'better') ways of doing things, embrace new technology, and encourage those around us that all of these newfound technical changes will be the most beneficial thing for them.

Human Aspects of Technological Change

Given our background as technologists we are pretty confident about the change process. We must be. If we aren't we are in the wrong profession. However, being knowledgeable about our craft and technology doesn't always qualify us to implement change. We in the IS profession are still often blissfully ignorant about the effect that change has upon the people we implement this change on. Why? IS professionals as a group (and I realize that this may be a generalization - but one I can live with) tend to ignore what I'll refer to as the 'soft' side of our jobs - the people side.

How often do we hear the following laments - Our users are resisting change; Our users don't want to change; Our users won't accept the new system. Are the complaints valid? Yes. Are they true? Probably not. Not in the sense that users won't issue the complaints, they will but we often miss the point of their complaints. People don't accept or resist the technology, they accept or resist the way that technology changes their lives.

New technology (feel free to substitute 'systems') forces change upon people (including us). Earlier, I stated that change creates

anxiety. The following are just some of the reasons that people will resist change because of the anxiety it causes:

1) New skills are required
2) Points of influence, authority and control are redefined
3) Roles, work relationships and responsibilities are modified
4) Data ownership can change
5) Privacy and security concerns may increase
6) New management techniques and organizational structure evolve

What is the process for change?

On the surface change is simple and easy. Making change can be most easily described as a top-down process. Change targets are identified, goals are set, specific objectives are identified, targets are formed. Plans then follow the goals. Once your goals and objectives are established, each party to the change process can then create their own individual list of the goals/plans that each subgroup needs. Everyone knows what is expected of them and what all others are supposed to accomplish in sync. In other words, massive planning and operational documents will be the guiding force for change. You have all seen this type of document, the ones that sit in nice black three-ring binders on managers desks - reference documents.

Ideally the plan is to be used, follow-up checkpoints are completed, deviations to the plan are identified, so now we can spend endless hours trying to explain what happened and implicitly deny any responsibility of what happened. Somewhere in this process the guilty parties run for cover and some of the innocent get hung. Does this process sound familiar to anyone?

Why does change fail?

Basic management principles revolve around these primary functions: planning, organizing, directing and controlling. These fundamentals have been taught at business schools since MBA's were

in their infancy. However, the situations that we find ourselves in today are nowhere was static as in the past. A point could be made that most planning is a waste of time due to the shifting sands we stand on. Why? We plan on the basis of variables, realities and projection as we understand them in the current context, and by the time that we get around to implementing/installing the change plan, everything around us has changed. This leaves us with the situation where we are attempting to implement yesterday's plan into tomorrow's environment. As a result, do we need to learn to plan better, or do we need to learn how to react more effectively?

I feel what we need to develop is a mindset that will help people understand what is happening and them to determine change activities sooner. We need people who are perceptive and flexible. Even if we have them, we'd probably still fail. Why? First, traditional top-down planning runs counter to the needs of today's employees. Second, too many planning processes lack the key ingredients: leadership, soul, passion and feeling. Third, plans are owned only by a few people at the top, not by the people who are often most affected by the change the plans are trying to implement.

How do we start to improve the process?

In his book "Change Is" author Stephen Baetz analyzes the dynamics of change through the eyes of a businessman's childhood toys kept in the attic of their home. Trust me, this is a powerful book and one that I heartily recommend for anyone interested in producing change in an organization. In order to properly facilitate change, we must understand and address the fact that most change fails, not for technical reasons, but for personal ones. Baetz offers up the following thoughts on interpersonal relationships:

1) The person who spills the milk, cleans it up.
2) The walls we build to keep others out also keep us in.
3) What we fear most, we should face first.

How should these thoughts affect how we as IS professionals view change? First, we have to take responsibility for the things that we do. When we go into departments and reorganize their work structures and turn their world upside down, we have to be responsible enough to pick up the pieces when things go wrong. Too often IS shops fall back on the standard excuses that change implementation failed because the end users lack of commitment, lack of effort, or resistance to change. Second, we must strive to change the well-entrenched idea that IS departments are entities unto themselves, that do not want outside intrusion from what I will call the 'real world'. A recent CDC study cited that alignment with business goals should be the primary concern of IS departments for 1994. Why? We should already be aligned, this alignment request has shown up on the to-do lists for IS shops for the past half dozen years, and we still haven't done it. Why? Because we like our technology, we like out terminology, and we as a group like to think that we are often 'above' the end-user departments that we serve. The third thought follows closely to the second thought. In order to gain the business alignment that IS is avoiding, we have to face the fear that our end-users see us as geeks, nerds and propeller heads. It will only be by getting out into the end user departments that we will dispel the myths that we are different from the other departments in the company, and by doing so we will also find that they are not unlike us in many respects. End-user departments suffer from the same corporate problems, they think no one else is capable of understanding what it is they do, and they feel they are special.

The personal change process

Psychologists have studies the change process for years. the dynamics of change, why change succeeds or fails, why change doesn't phase some people and destroys others. Swiss psychologists Claes Janssen has developed a model to explain how groups or individuals respond to change called the "Four Room Apartment". Every person, group or company lives in this apartment.

We begin in the Content room, what we might call the status quo. People (including us) like this room, we are happy, satisfied - and realism is the key component. Then somewhere, somehow, the real world intrudes - the need for change is dictated, often not entirely perceptible to those in the room, and this is where we move into the next room - Denial.

In denial we are afraid of the rumor, sometimes we are insensitive are ignorant. Those who are watching from the outside see us as unsure of the upcoming predicament. People stay in denial as long as they possibly can before moving to the third room - Confusion.

Confusion leaves us out of touch, unsure, exactly as the name would imply. We know it, so do others, and this is the most unsettling point in the process. Finally we struggle through this and open the door to Renewal. Some might call this final room 'acceptance', but this is where the new ideas come from, energy returns, enthusiasm builds, and the things we need to complete the process are now available to us. They key is then to keep the team focused on the organization until we accomplish it, then ensure the group stays focused for the next change process.

Who will we meet along the way?

Jacqueline Winder and Priscilla Donovan in their book 'The Flexibility Factor' identify four distinct categories of people that we will encounter along the change process, a key component of change being the ability to identify your team members, and use their strengths to help you achieve change.

The Risker

Their energy, enthusiasm and intuition will help them handle change easily, but when they make mistakes, the results can de drastic.

The Relator

Sensitive to people, this makes them a valuable ally in influencing others to change. Their excessive pool-taking to attain consensus can drain energy from the process.

The Refocuser

Concentrates intensely on the task at hand but the narrow focus sometimes causes them to lose focus on the vital sense of the big picture.

The Reasoner

They are skillful in dealing with change that requires organization but their meticulous planning can become and end in itself, dragging the process on, and on.

In addition to these individuals, we will also meet one of the newest personality types in the workplace today, someone that I will call the 'Sesame Street Individual'. We currently have a generation in the workplace today that has been raised to question everything around them. Not only that, but to them personal power is more important than position power. People's expectations have changed as well, do in large part to changes in socialization and child rearing. For example, kids in art class used to create 'art' and when what they created was a mess, they were told not to create a mess like that again. Today, an act like that would be considered stifling to the individual. As a result, these individuals come to work expecting to praised and rewarded regardless of the actual quality of that effort.

The team concept is also a part of this new individual. Teamwork is how the school process works today, everyone gets the team grade regardless of effort. There are those who have learned how to avoid some of the work and tag along with the team effort, and a key to implementing change will be to identify those who won't help move them team along, but only coast along with them team.

Avoiding the 'Doom Loop'

Downsizing. Rightsizing. The new MRP system. The installation of the new LAN. The major project. These are the projects that IS departments live for. As someone who has been involved in many of these various projects, and yourselves as readers, know that these major projects can often be accidents that are waiting to happen. Our goal is to avoid the accident, and move the process ahead. Lets take a

look at a major projects change process to identify things to look for. MaryJo Ubriaco of Booz, Allen and Hamilton has identified this change as a five stage process.

First, all of these major projects start with what is called 'the great idea'. The master plan, if you will, the big picture. Fundamental change is at the heart of the great idea, perhaps reengineering rings a bell as the current industry 'great idea'. Sponsors buy into the project, the word spread that something major is going to be undertaken, and enthusiasm begins. The second phase begins just after the pronouncement of the great idea, and time we can call 'happy time'. Characterized by an upbeat period of time when energy gets expended on the right things for the right reasons. If staff and resources are required, they are obtained. The cross-functional teams are formed and group cooperation is high. The third phase is characterized by 'process fever'. Focus shifts to issues such as standards, procedures and project schemes. The initial excitement is no longer apparent. Procedure manuals are being developed, process analysis in undertaken. In other words, the real project work begins at this phase. Stage four is the key of the five phases because this is the point where 'project erosion' sets in. Semantic and academic arguments begin, the realization sets in that there are enormous gaps between high level business vision and the blueprint needed for the IS implementation. The teams want to go to the users, but the users want answers, not questions. Credibility becomes an issue as deadlines pass and questions about time and money spent are being raised.

We finally reach an impasse, a stage Umbriaco call the 'recriminations and turnaround' phase. Perhaps some of these comments will sound familiar: "The project is late...", "The business has changed...", and the ever popular "The project needs more money". Officially runaway status is granted, the project gets suspended, the teams are disbanded, the gears stop. The guilty run for cover, some of the innocent get hung and perhaps a savior emerges with what might be called 'Son of great idea'. Then new teams are formed, new leaders arrive with new ideas, and the change business us used as the base for the implementation of the new project. The process begins again, 'Son

of...' get turned over to the technologists and the same cycle may begin again, hopefully with a better end, but we all know that there is not guarantee.

Identify the change audience

In any process, once it has been determined that changes have to be made, only part of the job is complete. If change is to be successful we must understand the other parties and organizations receptiveness to change. If your organization is not receptive to change, your best efforts will go to waste. Misinterpretations will ensure, distinct will be created and outright blockages will occur - from just about every level. Work by consultant Robert Steele identifies 5 key items that can determine your organizations readiness for change.

1) Organization Structure

It should be no news to you that a rigid organizations resist change more strongly than the newer flat, less hierarchical structures. Your typical pyramid causes empire and turf building, creating limited cooperation between business functions. Be sure however to not be deceived by your company organization. chart. The new flat or 'horizontal' organizations can still have a traditional hierarchy at its core as a power structure. How your employees view the organization is more important than how the organization appears on the chart.

2) Morale

We already know that change is stressful, and this stress is most easily accepted by organizations where morale is high. If you aren't sure how the morale is in your company, you must find out before you begin to undertake the massive change project that you are considering. Finding out the prevailing attitude isn't enough. Management must be willing to act upon the issues that dissatisfy employees with concrete visible actions.

3) Skill Sets

Does your organization have the skill sets required that will be needed for the change process that will be undertaken? Change

demands new skill sets be adopted. Don't make the mistake of confusing skill sets and accomplishments of tasks as the same thing. Processing orders quickly and accurately is a task. Communicating effectively with customers is a skill. Too often when we implement change, we tend to focus on the task, ignoring the skills and attitudes that surround the task. If you only focus on skills, you will end up with someone who has all of the technique but can't tell you why they are doing what they are doing. Successful change requires a knowledge of why.

4) Performance Measurement

We already know that change processes often fail to accomplish what we set out to accomplish. Sometimes the failure is not our fault. Sometimes the failure stems from the mismatch between the performance measurements used and the desired outcomes from the process. Picture this as an example. Senior managers are judged on overall performance of the company as a whole; while junior managers are judged on the performance of their individual departments, and the front line employees are participating in a company profit sharing plan. What happens when you try to implement change here? The departmental managers resist change from the top because what affects their departments affect their performance measurements, and if what is a good change for the company causes their department to suffer from a poor measurement for a period of time, they will resist because it affects their bonuses. They will also resist change from the bottom, where front line employees want to change to drive business profits higher because their profit sharing is based on the company's profit, not their departments. Another scenario may even see the front line employees resist change if it is seen that the change will reduce the number of front line employees in the firm. For those who stay, profit sharing goes up, but for those marked to go, can you expect them to help implement change willingly? If the goals of all the separate parts of the company are not aligned in such a way that the measurement of them does not force all of the rower to pull the oars in the same direction, then change will be difficult.

Finally, if you want to be successful at change, you must study the past efforts at change, especially those companies where change has been implemented into cross functional organizations. Executive sponsorship is a way to focus attention, coupled with communication that prepares people not just for what will be changed, but identifies what is not changing as well.

Epilogue

Managing and implementing change is one of our primary functions in the organizations where we work. Successfully implementing change as our goal requires that we spend our time up front in preparation to implement change. The change usually isn't the hard part, seeing what is needed for successful change, and doing the work to set the climate for change is the hard part. It isn't only hard - it is our job. Only when we begin to understand our organizations, the people, and the process will we stop getting angry at our lack of progress and start to see a dramatic improvement in our opportunities of successful change implementation.

The 'IZE' have it

Robert M. Gignac
Speedware Corporation
150 John St. 10th Floor
Toronto, Ontario
M5V 3E3

Abstract:

You don't have to look very far in any current computer periodical to find the immediate family of 'IZE': Downsize, Rightsize, Resize, and their ever popular cousin 'Reengineer'. In print they appear to be just the saviors our companies need in order to regain our competitive edge. Or are they?

This presentation will discuss the how's and why's of the current downsize, rightsize, and re-engineering trends, stopping along the way to point out that all is not what it seems in print. There are certainly rewards to be reaped along this path, but there can be some devastating pitfalls that we must avoid as well. It is possible to get from point 'A' to point 'B', but as someone who has been through the process once commented: "The old ways work, the new ways work, the goal is to not disappear in the process of getting from the old to the new".

What is rightsizing anyway?

Rightsizing is the process of making sure that your computing resources match the needs of your business. The resources are perhaps hampering your operation because of ineffective use, or perhaps they are tying up capital, perhaps they are simply outdated. Rightsizing means the your information systems are deployed in the most advantageous way to match both the size and information needs of the operation. We are currently in what the pundits are calling the 'information age', a time where the acquisition and flow of information has never been more important to manage. How your company acquires, manipulates and dispenses information has become a crucial

corporate function. As with anything that affects your bottom line, you can handle information well, or you can handle information poorly. By allowing your information technology to become obsolete, you run the risk of losing time, corporate efficiency, money and most importantly, the complete reason for your existence: your customers.

Does everyone understand what the buzzwords mean? There is still much confusing in the marketplace. A study by Gateway Information Services of 121 senior executives at major corporations, 88 percent of those responding stated that their companies had carried out re-engineering programs in the past 2 years - more than structuring, downsizing or outsourcing. However, only 46 percent defined re-engineering as workflow or process redesign. Others thought that re-engineering meant technology change (17%), product changes/improvements (16%), improving efficiency (8%) or increasing customer satisfaction (4%). It would seem that there is still much confusing surrounding the buzzwords.

Is rightsizing the answer?

Perhaps it is more appropriate to ask: If rightsizing is the answer, what exactly was the question? I would think that I could get some agreement that rightsizing, for the vast majority of us, means migrating from a mainframe (regardless of vendor) to a different, usually smaller system of operations. The reasons for doing this are many but are most often characterized by an IS wish to reduce long application backlogs, replace poor quality applications, and reduce excessive support costs. Don't be fooled into thinking that only mainframers are downsizing - many owners of 'mini' computers are also downsizing, shifting (or attempting to) applications onto PC networks (be they Lans or Wans).

I earlier asked what might be the appropriate reason to consider downsizing. There are as many as there are IS shops. Typically, I think you could lump downsizers into two groups: 1) those with old equipment and outdated practices; and 2) those whose systems suffer from a lack of business alignment.

Organizations with inefficient IS installations have historically been reluctant to make ongoing capital investment. After dragging their feet on issues such as operation consolidation, rationalizing systems, and automating labor intensive tasks. Saddled with 'obsolete' technology and a lack of ongoing investment, they suffer from a vicious cycle of declining efficiency. As outmoded applications get older, more and more time must be spent in maintenance mode, a process that is personnel intensive and increasingly costly.

In shops that suffer from a lack of business alignment, this can usually be attributed to a lack of proper planning among those in the IS community. Alignment with business goals has been high on the list of IS tasks that must be undertaken for each of the past five years. In many cases it is directly tied to the obsolete technology and lack of investment. Those shops that spend the majority of their time doing maintenance can't react to meet the needs of the business. It starts a vicious downward cycle because as businesses see IS not meeting organizational goals, the less likely they are to fund ongoing IS projects which stop it from trying to attain business alignment.

What should we be doing?

A recent study by CSC Index (Cambridge, Mass.) cited the top IS priorities for 1994, and found the vast majority centered around business-related issues. The top five priorities were cited as: 1) re-engineering business processes; 2) aligning IS and corporate goals; 3) improving systems development; 4) instituting cross-functional information systems; and 5) organizing/utilizing data more effectively. Changing technology ranked tenth on the list.

When you talk to companies that are engaged in a downsizing effort or are starting to consider a rightsizing process, reducing costs is usually offered as the driving force. Reducing which costs? Hardware costs, maintenance fees, and sometimes in software licensees. However, the cost structures in the average IS shop are not that simplistic. Equipment purchases and licensing fees are not the largest component of the overall cost structure. In order to get the best deal

from a rightsizing initiative, we must look at both sides of the cost equation.

Even when properly managed, PC Lan-based computing in terms of cost per user can be an expensive solution. If not properly managed, it can be a really expensive solution. Recent studies show cost estimates ranging from $5,000 to $9,000 per user, but once enterprise infrastructures are factored in, these costs can rise to upwards of $20,000. Why? Too often, the cost reductions in a downsize/rightsize process are misleading. Support costs for a newly distributed system can be prohibitive. As a rule of thumb, figure one Lan manager for every 30 users, one storage and operations manager for every 60 users and one application managers for supporting PC packages for every 60 users. Typically, between 3 and 7 percent of all PC Lan users spend their time supporting other users. On top of that, add in personnel costs and other overhead - the type of costs that never show up in end users budgets and you can easily see how easy it is to arrive at a final price for the new downsized system that may not yield the expected savings.

The second most popular issue in downsizing has to do with hardware performance, with the most popular benchmarks measuring simple isolated facets of CPU performance. In reality, looking at benchmarks on paper doesn't give you a true indicator of machine to machine performance. Real performance of a computer system is a function of hundreds of different parameters, not a function of MIPS alone. These parameters include things such as network facilities, CPU's, I/O facilities, memory, disk storage devices, database management systems, operating systems, etc., etc. In their simplest form, MIPS only measure the number of instructions that are processed, ie: the raw speed of a device in an optimum environment. In reality, MIP numbers are much like MPG estimates on that new car you just bought, right down to the asterisk at the bottom of the page which reads: your mileage may vary.

What do MIPS measure? Performance. What kind of performance? Merely the number of instructions that can be processed per second. There is nothing fancy about this. Are MIPS a good

benchmark? I'm not so sure. You shouldn't just look at the number of instructions, but the complexity of the instructions. A good analogy to MIPS would be trying to measure the performance of a manager by the number of orders they give. The basis of true performance is not the number of instructions, but the efficiency and timeliness with which they are carried out.

Is the process easy?

Do to the number of articles you see that tell you everyone is in the midst of one (or more) of the buzzword processes, the task must not be that difficult. Believe me, the tasks are difficult. Not impossible, but still difficult. Key areas that you will have to address are the following:

1) Education

Reeducation will be a major factor as you begin to address the new world in an 'ize' process. Who needs the reeducation? Everyone. The workforce must be taught what the process is, how it differs from the work processes they already know, and to inform them of the roles that they will play in the process. This education process will be a major task, especially if your company has been stable and unchanged for a long period of time. Changing work habits and introducing new thoughts is a challenge that has not been faced by many firms. It requires that workers become skilled in actions they may have not considered, including new processes, and they must embrace a new way of thinking.

Part of this reeducation will be preparing workers for new responsibilities, and we must assure them that the new processes will be flexible, of high quality, and productive. The odds are high that you will have to so outside to professionals to serve as mentors and trainers until your internal staff can take over the process.

2) Communication

Close behind education in terms of importance, is communication. Poor communication can impede or even derail a well

run project. What should be communicated? Everything from the change management principles that are at the heart of the process to any and all successes, as well as the day-to-day accomplishments that maintain momentum of the project. Remember that people have to know what is happening around them in order to stop rumor and innuendo from spreading. Once rumor starts, it can be a difficult process to diffuse. Stay up front throughout the process, tell everyone what and why, and keep it foremost in their mind that their contribution to the process is integral in moving it along.

3) Human Resources

How successful the project will be is largely dependent on how the workforce responds to change. If you give them the right information and the opportunity to participate, you increase the odds for success. You must realize that the makeup of your workforce will be changed through this process. Will layoffs occur? There is a high probability of that, and it must be something that you are prepared for. The higher the probability, the more likely you will encounter resistance, due to fear, from the people involved in the process.

4) Incentives

When you are changing roles of individuals, implementing team style workforce, you must examine changes in reward and recognition systems. It is in your best interest to nurture the creativity of the reorganized workplace in order the reap the benefits of a new, tighter structure.

Does Rightsizing fit your strategy?

If you assume the concept of rightsizing means the shifting of resources to smaller distributed platforms, or perhaps the move to the ever-popular 'open' system, you have more issues to consider. Issues such as: 1) What about client/server? 2) What about open distributed computing? 3) How will my staff react/do I have the right staff? 4) How do we feel about the risks involved in the process?

The risk issue should be paramount in the consideration of a downsizing/rightsizing/reengineering project. How much do we want to be on the 'bleeding edge' of technology? Are the new-fangled tools mature enough to handle the applications? For all the bashing tossed its way, there is nothing that can be cannot be written in COBOL - the same cannot be said for some of the new PC-based client-server tools. Remember, you may be in a position of betting your companies future on these new solutions, which if implemented poorly could lead to severe problems. We in the technology field know the hottest buzzword tool may not be around two years from now, and you must determine how much risk you can take.

How do you select a destination platform? Perhaps by looking at your current staff - what is their knowledge, how are they equipped to cope with a change to a different platform? Do they have relational database experience? Unix exposure? If you go open - what about PC and telecommunications experience? Are single vendor solutions appropriate? How proven are the 3rd party solutions for your companies situation? It appears that I am not answering any questions with this line of questioning and you are right. What I am trying to do is offer up suggestions that the shift is not as easy as the trade journals would make it appear.

In today's marketplace a key factor is selecting the appropriate applications software - and the hardware platform it runs on - that best fits your business need. There is a decided shift to buying third party application software packages over writing, testing, and implementing your own solution. A good rule to remember is that you are responsible to build, test and integrate and open system yourself, while you buy a working integrated proprietary system already testing for you.

In the shift to redistributed systems, don't confuse distribution with client/server computing. There are many systems for which host based, time-shared operation are still viable. Many of the key features of desktop computing such as windowing, GUI's and SQL database access can be attached to hosts. Remember that the key is solving business problems, not merely trying to keep up with the trade journals

in terms of the hottest technology. There are really only three reasons to examine technology and approaching a change in your technology platforms: 1) to make money 2) to save money, or 3) to increase competitiveness. Any other reasons really don't count. If you are not accomplishing any of the three reasons above, then I would suggest that you are barking up the wrong tree in terms justifying your reasons for change.

Some thoughts from the front

"There is one way to tell if you are proposing a reengineering project. The first time you bring it up, if no one scream, 'Are you crazy?' then you are not proposing a reengineering project"
Robert Rubin, V.P. Elf Atochem N.A.

"If vendors didn't have snake oil to sell, they'd have a tough time to make a buck. If every solution for the world of IS was implemented that would have us $50 million, we'd only spend about a buck per year on IS today"
Tom Loane, V.P. Alamo Rent A Car

"Could this be the CASE of the 90's"
Roger Burlton, Director, SRI Strategic Resources

Epilouge

For all of the macho rhetoric we find in print, reengineering is really nothing new or original. It is at best a few old industrial method that have been repackaged to seem like the latest in management technology. At its core is the expectation that we must change the way we work and the way that we distribute information processes. This change won't occur merely because management decrees it. As someone who has been through the process, Bruce Ruppert of Agway comments: "You can survive the old way. You can survive the new way. It's the goddamn transition that will kill you". It is at the point of beginning a major transition such as discussed in this paper that the true test of management in your companies will begin.

Ancient Success Principles
Author: Robert A. Lund
Paper Number 1006
Lund Performance Solutions
3111 Santiam Hwy Suite I
Albany, OR 97321
(503)926-3800

Abstract: Many years ago, Aristotle and other ancients spoke of essential ingredients for success. In this talk I will present some often neglected principles for successfully managing ones own life as well as other people. I will emphasize the need for revolution in the way we think about self and people management, based on time-tested, ancient success principles. Every manager will benefit from this talk. I have lectured on this subject in the U.S. and Eastern Europe with great results. The attendee will walk away equipped to better deal with personal life-stagnation as well as being able to help others. Such topics as, vision, integrity, relevant technology, "warm-ware" (people-care), etc., will be covered.

A Politically Un-Correct Introduction

It seems that whenever you talk about politics or religion, in particular, people rapidly polarize, even to the point of not being nice to each other (perhaps an understatement). In spite of this, I have decided to be stupid enough to think that I will be immune from the doing the same with this paper. I will be injecting some thoughts that will assuredly go against the grain with respect to the popular notions of success of our day.

In this paper, I will be specifically addressing an audience that feels like me and many others who are dissatisfied with the prevalent definition of success. I will present many principles that have worked for thousands of years, yet some of them have been ignored, even disdained. However, they will only work best for those who are "gutsy" enough to stand up against the popular philosophy of the day.

It is important that you understand where I am coming from. I refuse to be an armchair philosopher. I practice what I am about to preach. I detest those that go about speaking great success principles that they have not processed in their own life. Lessons lived have much more teaching effectiveness than lessons lectured.

I am halfway expecting some protests (perhaps violent) to the contents of my presentation, but I assure you this is simply a case study out of my own life. Many have asked me over the years how I have come to enjoy such a successful life. As a result, I have decided to "come out of the closet" and describe what has worked for me... Really worked. The concepts I present here are tried and proven in my business, marriage and family. If you will seek to understand and implement them, you will do well.

I believe that many of my views represent a very large, yet somewhat silent, majority. They certainly represent the views of the founders of the greatest nation that has ever existed. I also believe that there is a very loud and vociferous minority that has been largely responsible for perpetrating the mess that much of the United States is in now. I do not believe that damage recovery for this nation, your company, your marriage, or your life will be effective without a literal revolution. For the purposes of my discussion, this revolution will be first in the way you think, and then in the values that drive your behaviors on a day-to-day basis.

Although it is not worth spending much talk in this presentation regarding the ills around us, I will mention a few, just in case you have been locked in a room for twenty years programming...

Marriages and the family unit are being blown to smithereens (with all the resultant shrapnel), businesses filing chapter 11 or worse, leaders being exposed and brought down, economies faltering and failing, standards of living, for most, steadily declining, children are rebelling and running away in unprecedented numbers, overdosing, and dying at alarming rates, our inner cities nearly resembling "Escape From New York", the elderly losing hope, and a general, apparent decline of (particularly) western civilization. The sad irony of it all is that people still persist in attending success seminars and buying self-help books. Although there will be some overlap with what I am presenting here and such books, I assure you that you will not find the my main thesis in a B. Dalton bookstore or in a tape series from Nightangale-Conant.

How do all of the aforementioned problems affect you and I? It doesn't take a rocket scientist to answer this... The very core of who we are reacts with fear: economically, socially, emotionally, etc. Fight or flight mechanisms kick-in and cause us to literally scramble for survival. To top it all off, politicians keep putting salve and bandaids on where radical surgery is really in order. What I am proposing in this paper is more likened to radical surgery than "success bandaids".

The goal of this paper is to equip you with the implements necessary for radical surgery, so that you will either sharpen and hone what you are doing right, or perform the necessary mid-course corrections in order to become a more truly successful person. This leads us to the slippery question of what success really is.

What is Success?

Success, as it has been defined in a purely secular fashion, usually amounts to an unlimited license for personal selfishness, and un-limited greed. That's funny. When you think about it, the things that you really need are not exclusive to material possessions. I think of a quote I heard of by Queen Victoria at the height of the British Empire when she said, "Nothing tastes." Interesting comment coming from the world's wealthiest woman (at least for that time). You know the feeling; how you wanted that house, person, job, plane, etc. for years, you finally got it, it was great for the "honeymoon" period, and then you wanted something bigger and better.

Jesus Christ said something very curious about this. He declared, "Beware, and be on your guard against every form or greed; for not even when one has an abundance does his life consist of his possessions." Hmmm. Success, then, may encompass, but is certainly way beyond what we possess. Even Aristotle himself felt that success had more to do with being a compassionate, loving, honest person than just ones net worth.

H.G. Wells noted, "Success is to be measured not by wealth, power, or fame, but by the ratio between what a man is and what he might be."

Success has more to do with intangibles like peace of mind, emotional stability, knowing that I am loved, feeling a sense of worth and accomplishment, etc. Our basic needs revolve around love, security, meaning and significance.

Success has more to do with who I am than what I own.

Success has more to do with loving and being loved than sensual pleasure.

Success has more to do with living out ideals than gaining a level of attainment.

Secular versus Non-Secular Success...

Prior to just a few hundred years ago, God was the center of much of man's thinking. Then, with the rise of humanism, God became "un-cool". People came up with all kinds of ideas explaining away the existence of God. I was one of those. I was a staunch evolutionist. I made religious people literally cry with all my arguments against God. Even as harsh as I was, something inside me was very empty. The English poet Alexander Pope described exactly what I was feeling when he said, "Atheists put on a false courage in the midst of their darkness and misapprehensions, like children who, when they fear to go in the dark, will sing or whistle to keep up their courage."

"An atheist's most embarrassing moment is when he feels profoundly thankful for something, but can't think of anybody to thank for it." Mary Ann Vincent

I was taken to task to finally prove or disprove His existence, if for no other reason, for my own thinking. After much research, I finally came to the conclusion that both the God-haters and God-lovers tended to not be objective, but rather biased. They also tended to not know what they really believed; each group simply believed someone that they looked up to, rather than perform original research themselves. After sorting through the "politics and nonsense" of both atheism and religion, I came to the conclusion that there was a God. Further, I did not want to have anything to do with organized religion, but I really wanted to know Him... He graciously answered my desire.

The God Factor

In case you are wondering, this presentation has nothing to do with organized religion. I have been much more interested in relating to God than to which religious organization I can give my money to! Since I have been walking with Him for the last twenty years, I have become more and more convinced that relating to God and being involved in organized religion are quite often mutually exclusive!

In case you haven't figured it out, yes, I believe that there is no success worth having that doesn't factor in God himself. I'll show you what I mean in a while. As far as the existence of God is concerned, I am coming from a teleological standpoint. That is, I believe that man is a manufactured being. Rather than evolving from bugs, man was engineered, with great precision, I might add. The Greek word TELOS speaks of the "end result". If the end result is highly intricate and displays great precision and symmetry, then it is probable that the engineer was at least as intelligent, but probably much smarter than the created product.

I also believe that the engineer (God) has communicated a set of guidelines whereby the human machine may be operated at peak efficiency. Any departure from such guidelines results in degraded performance in human achievement or success. Since, by trade I am a computer system performance specialist, I see many customers who operate their computers in an inefficient manner. Likewise, it does not take much to notice the inefficiencies all around us with respect to humankind. It is a subset of these guidelines that I wish to present to you so that you will excel in being a human, friend, spouse, parent, citizen, a manager, or a worker, in a company.

A Brief Overview of the Principles

Most of my personal success has come in adhering to guidelines that I have gleaned out of the Bible; specifically, from the book of Proverbs in the Old Testament. I happen to agree with 19th century English critic and social theorist, John Ruskin, when he said, "The Bible is the one book to which any thoughtful man may go with any honest question of life or destiny and find the answer of God by honest searching."

I have operated on the premise that these principles are merely a subset of some of the "laws of life physics" that God has prescribed for humans to operate at peak efficiency. It is important that you realize these principles may be likened to a railroad track. A train is an incredibly awesome work-horse when on the tracks, but off the tracks it is simply a massive piece of useless iron. I'll now describe the ten areas of focus that have served as the railroad track for my life.

The Ten Areas of Focus for Success

The following ten principles represent the areas of focus that are the building blocks of human life. The ten of these together represent a report card. If you score a 4.0 GPA, you'll do exceptionally well in life (taking into account the Missing Link that I discuss later). Your job is to concentrate your efforts to accomplish the goal in each category. They are not necessarily in any particular sequential order.

Focus #1 - Vision and Purpose

This principle has to do with goals and objectives. I have found that this area is the one of the most common reasons for failure in a person's life: Not having a clear mission in life. The following illustrates this,:Proverbs 29:18 - "Where there is no vision, the people are unrestrained..."

This bit of ancient wisdom really means that without a plan or a goal (that is worth living or dying for), you will not discipline your life to the extent necessary to reach that goal. I have seen this time and time again, in my life and others. So many people I know, and have known, have not had a life's mission statement that trickled-down into their every day life. Living for the weekend is not enough to gain the fulfillment in life (i.e. success) that people generally want.

When I first set out to start my own business, my wife and I agreed that it was going to take tremendous restraint of our collective "appetite" in order to become financially self-sufficient. We spent the following five years keeping an eye on the goal of having our own business, focusing on the resulting benefits: More time, more money, more independence. Consequently, during the first stages, we exercised great restraint. We limited and eliminated debt, radically and consistently altered our eating and clothing habits and styles, etc. There were times when it was difficult, but we would take turns reminding each other of the ultimate goal. This would "sober" us again to stick with the various austerity programs we would initiate. Was it worth it? YES! It was the single best decision I

have ever made with respect to my vocational life. If, however, we would not have restrained ourselves, the business would have become a tremendous boat anchor around our necks.

Proverbs 21:5 - "The plans of the diligent lead surely to advantage, but everyone who is hasty comes surely to poverty."

Here again the concept of having goals and plans to attain those goals is underscored. The up-side of making plans and being diligent to attain them is the reward of prosperity. This is not just limited to material prosperity, but rather covers all of life's departments. For example, I have found that I need to utilize a time management notebook. As I scan my life goals, yearly goals, and daily goals, I write down attainable action items and prioritize them each day. The more diligent I am at this, the better my days go, and the better I feel. I have recently initiated the habit of writing down the weekly "wins" and posting them like wallpaper in my office.

I personally find that there is a tremendous difference when I write goals down rather than keep them between my ears. Perhaps this was ingrained in me when I was in some sales training class years ago. I heard the instructor tell of a study (sorry, no reference!) where a number of graduates from a famous university were watched for a twenty year period of time. Eighty percent of these did well in life, as measured by apparent happiness, material prosperity, etc. This group had no specific goals in life at the time of their graduation. Fifteen percent of the group had very specific goals but never wrote them down. They did twice as well as group number one. The last five percent had very specific goals and simply wrote them down. This group did three times as good as the second group. Gee, simple math says that this small minority did six times better than the first group. The moral: Identify goals and be sure to write them down.

ACTION: Define and focus on your own visions and dreams. Create your life's mission statement. Then, prescribe the various restraints you will have to deploy in order to meet such targets. After you have defined your goals which are simply practical steps necessary to implement the mission, WRITE THEM DOWN and review them often.

Focus #2 - Health

Let me say from the outset, that this area of optimal health has been somewhat of an obsession with me. I have nearly 100 books in my library alone on the subject of health. Much of the information on health management can be quite technical, and, unfortunately represent some hidden agenda or soapbox by such authors. I have come to look at the whole business of good health as being very simple in concept. It also does not have to be expensive! In the first book in the Bible, God prescribed what food was to be the fuel for the human machine. I'll let you do a bit of reading to see for yourself (hint: in Genesis chapters 1-4).

Even if you excel at all the other nine areas of focus, but have poor health, you will not be able to fully enjoy life. How many people do you know that spend the bulk of their lifetime achieving some level of success (usually defined as big bucks in the bank), and their health is very poor? I know of a bunch. Consider this anonymous quote, "A lot of people lose their health trying to become wealthy, and then lose their wealth trying to get back their health."

In our fast moving information age, we sit a lot. But when we move, we move fast! When on the run, we fill our faces with fast food. This stuff is stripped from life-giving nutrients. Do this, start looking at folks in their fifties. Look at their face. Now find out what the majority of their diet is. It doesn't take a rocket scientist... Really, you are what you eat. Since many of us are unable to get what we need in the food we eat, vitamin supplements are a possible option. As an excellent alternative to supplements, take a look at the tip I'll share in a few moments.

As to general health management, personally, I have done the following after an awful lot of research and experimentation:

1) Maximize my intake of as fresh as possible fruits, vegetables, and grains. I am actually a back-slidden vegetarian, as I occasionally cave-in to a Taco Bell taco!
2) Get plenty of fresh, clean water.
3) Get enough sleep. I do best on about seven hours.
4) Reward yourself once a week or so with some sort of junk food! Instead of eating poorly 90% of the time and good 10% of the time, do just the opposite.
5) It is essential to limit the amount of salt, fats, smoke and alcohol (gee, all the fun stuff!).

6) Exercise 2-3 times a week doing some kind of aerobic exercise for 20-30 minutes each time.
7) I also take three *tablespoons* of bee pollen each day (more on this later).

Think about your health this way. Just as an engine manufacturer prescribes a certain fuel for efficient operation, so also the designer of the human body prescribes certain fuel also. Although the Bible was not intended to be specifically a health manual, there are numerous hints as to what God prescribed as the best fuel. Any departure from the best fuel that the human body was designed for will result in various levels of degradation and shortening of productive life. I don't know about you, but I would like to live a long time. However, I do not want to just be alive, I want to have optimal health. Here's one more valuable tip that I cannot resist telling you:

Eat fresh bee pollen (be sure your doctor confirms that you are not allergic to it)! No kidding! I believe this is one of the best kept secrets around, so I am going to spent a bit more laser-jet toner on this one. You will be astounded at the level of energy, alertness, and lack of sickness by working raw bee pollen into your diet (the capsules are super-expensive). Bee pollen is the ONLY substance known that has ALL the vitamins, minerals, enzymes, amino acids, carbohydrates, etc. necessary for optimal human health. Consider the following from a publication called, The Health Counselor (Vol. 5, No. 5, pg. 30),

"Bee keepers from Georgia - formerly of the Soviet Union - sold their pure honey and ate the "dirty residue" at the bottom of the honey containers themselves. But this "waster matter" wasn't as worthless as the bee keepers and their customers believed, according to Professor Nicolai Vasilievich Tsitsin, biologist and experimental botanist. The beekeepers for whom this "dirty residue" - pure pollen - was a dietary stapel were known to live as long as 125 years!"

It so happens that I am a bee keeper. I have about 10 hives or so. I have been collecting pollen myself and feeding it to my family for over a year now. The incidence of sickness has been unprecedented! My energy level has been better than if I was drinking Mountain Dews! Because of the high level of zinc in pollen, not a member of my family of six (four of them under 11 years of age) has had a sore throat in a year. Consider this quote from Bee Pollen - Super energy, Super nutrition, by William H. Lee, Ph.D.,

"Pollen is the richest and most complete food in nature. It increases the body's resistance to stress and disease and also speeds up the healing process in most conditions of ill health; pollen also possesses age-retarding and rejuvenating properties. Truly, Bee Pollen is a miracle food, a wonder medicine and a true Fountain of Youth."

It is curious to me that Proverbs 24:13 says, "My son, eat honey, for it is good...". There are over 50 references to honey in the Bible. Of course, honey is better than white sugar, but it is pollen that is the "silver bullet." In the old days, it was difficult to keep the pollen separate from the honey. The implication is that there was naturally going to be pollen mixed in with the honey. The clearer the honey you buy in the store, the less pollen there is, and the less of a health benefit it presents.

ACTION: Take a realistic health inventory. Do whatever it takes to eliminate health enemies listed above. Replace each bad thing with some new thing. Work on re-engineering your taste buds to actually enjoy fresh and natural foods. Consider bee pollen.

Focus #3 - Wealth

In the area of money, I am not going to bombard you with which mutual fund to invest in, or other such advice. Rather, I am going to present some principles, most of which, you will never see in the books of our day. Ironically, some the most holistically successful individuals I have known or read of operate on these principles. I am sure that most people feel that money is the metric for success. I believe this is only partially true. Consider the following,

Proverbs 17:1 - "Better is a dry morsel and quietness with it than a house full of feasting with strife."

Is your house like that? Have you pushed, pumped and hustled to make some money, only to find that you have lost your family or something else valuable to you?

Consider the following by Benjamin Franklin, "Money never made a man happy yet, nor will it. There is nothing in its nature to produce happiness. The more a man has, the more he wants. Instead of it filling a vacuum, it makes one. If it satisfies one want, it doubles and trebles that want another way. That was a true proverb of the wise man, rely upon it: "Better is little with the fear of the Lord, than great treasure, and trouble therewith."

Let's say that you want to be successful, and that you are pretty balanced with respect to relationships, health, etc. What is the best way to acquire and keep wealth? Here's some interesting thoughts.

Proverbs 15:27 - "He who profits illicitly troubles his own house, But he who hates bribes will live."

The principle here states that you should not gain wealth in an illegal fashion. Furthermore, it states that even if you get away with cheating, lying, stealing, extorting, etc., you call upon your house (i.e. your sphere of influence) trouble. We both know what I am talking about here.

Proverbs 10:4 - "Poor is he who works with a negligent hand, but the hand of the diligent makes rich."

In terms of wealth building, diligence is a major key. Diligence means sticking with something and not slacking off. Thomas Edison knew this well when he said, "The three great essentials to achieve anything worthwhile are hard work, stick-to-itiveness, and common sense."

Proverbs 10:5 - "He who gathers in summer is a son who acts wisely, But he who sleeps in harvest is a son who acts shamefully."

In this proverb, the concept of a window of opportunity is expounded. If you are going to build wealth, you must take advantage of such windows. When I was a full-time employee thinking about starting my company and selling consulting and software, I thought long and hard about timing. If I would have waited a couple more years than I had, other products would have received a greater market share and the prospect of my business succeeding would have been much less. Instead, I was able to pay off debt and save money to purchase equipment and create a buffer to live on while making the transition. I have seen far too many launch a venture heavily-laden with debt and not enough buffer money available. Timing is an important factor in overall success.

Proverbs 11:25 - "The generous man will be prosperous, And he who waters will himself be watered."

Here is one of those "physics of life" laws. As sure as the law of gravity holds, so does the law of giving and sharing. This principle states that if you help others, you will in turn receive a measure of divine help that you would not have otherwise received. I cannot tell you the number of times I have seen this in operation in my life. I believe that at the very core of things, we were created to share in life, relationships and physical wealth. However, I do not believe in blind charity or in giving to organizations which eat up significant portions of funds donated with administrative costs. Recently, I have started providing funds to third world people to help them start small business enterprises. If you want to prosper, give something away!

"Proverbs 11:28 - "He who trusts in his riches will fall, but the righteous will flourish like the green leaf."

Proverbs 15:6 - "Much wealth is in the house of the righteous, But trouble is the income of the wicked."

In these two nuggets, God reminds us that it is not ultimately money that causes us to live successfully, but rather righteous thinking and living. Trusting in money alone is actually a deceptive thing. More on this in the "Missing Factor" section.

Proverbs 12:11- "He who tills his land will have plenty of bread, but he who pursues vain things lacks sense."

I have half jokingly, half seriously said that only poor people play the lottery and gamble. What I mean by that is this. If you want to prosper the old fashioned way, you will do so by some measure of hard, honest work. Now, in the age we're in, brains means infinitely more than brawn. As a computer consultant of some sort, you might command more money in one day than four months of sweat growing, hoeing, and harvesting a couple acres of corn. So then, the tilling referred to here could mean intellectual or physical labor. Certainly the point is that rolling the dice reflects how many marbles we have in our head!

Proverbs 10:22 - "It is the blessing of the Lord that makes rich, And He adds no sorrow to it."

Hmmm... I'm sure that you'll not see this principle in any "get rich" book in our time. So then, real success is more a function of God's favor than my labor. Working hard by yourself does not even compete with working hard on God's team.

ACTION: Simply do the following:

1) Eliminate all consumer debt... now! Then double up on your lowest balance car payment... then the mortgage!
2) You must save at least 10% of your net pay! Consider, "At retirement, only 3 percent of Americans are financially self-sufficient in a bare bones sense. Only 3 out of 100 retirees have enough investment income to provide themselves with $3,000 a year at age 65". (R.E. McMaster, Wealth For All). Cut your life-style and do whatever it takes to save that 10%!
3) Consider shifting out of paper asset investments (stocks, bonds, CD's, etc.) and into hard assets (gold, real estate, gold stocks, stored goods). At the end of the day, when all the dust settles, whoever controls the gold and the land does well, as a rule.
4) Begin to give money and time to those less fortunate... I didn't say lazy... It is one thing to feed a person a fish, it is another to teach them to fish. Identify people and worthy organizations who really help people to become free of charity.
5) One of the best-kept secrets of wealth building is to start your own enterprise. Identify some need... do lots of brainstorming first... finance your own growth... without sacrificing your relationships, work on product and service offerings that will create big money with little cost of goods and support. Software and consulting fit nicely into this category. A small enterprise is one this that the IRS hates, because there still are so many ways you can reduce taxes!
6) Allow God to be your partner. Notice I did not say organized religion or a TV schuckster... As you begin to take steps to initiate or improve a friendship with God Himself, all of the departments of your life, including your finances, will take on new meaning and improve.

Focus #4 - Communication

Our speech has so much to do with success or failure. Consider the following,

Proverbs 10:19 - "When there are many words, transgression is unavoidable, but he who restrains his lips is wise."

Proverbs 13:3 - "The one who guards his mouth preserves his life; the one who opens wide his lips comes to ruin."

Proverbs 17:27 - "He who restrains his words has knowledge..."

Proverbs 18:21 - "Death and life are in the power of the tongue...

Proverbs 11:13 - "He who goes about as a talebearer reveals secrets, but he who is trustworthy conceals a matter."

So true. We all know people who talk so much and cannot keep secrets. Such folk are usually repulsive. They also are not going to be my friend. Their words hurt, they cannot be trusted, they are usually putting their foot in their mouth, etc.

Holding back unnecessary words is simply smart. I have seen countless business negotiations turn out sour for those who talked too much. I have seen many, many people hurt friends and enemies, thereby hindering success in their own life by alienating people. Success is never obtained in a vacuum. You need myriad's of allies in order to reach your goals. Successful people have learned this principle.

Have you ever been in a meeting where emotions were hot? Consider, "A gentle answer turns away wrath, but a harsh word stirs up anger." Proverbs 15:1. I have had a tremendous time over the years with this one. By simply talking calmer and quieter, you can shut the jaws of a savage beast.

ACTION: Next time you are talking, listen to yourself. Could you have listened more and talked less? How much of what you say is taken wrong, and therefore hurtful or insensitive to others? Start catching the words and speak them in your mind first before letting them out of the gate.

Focus #5 - Work Ethic

How we view and respond to work now tells an awful lot how successful we will be in the future. A good work ethic is measured by what kind of a servant's heart a person has. Let me illustrate.

I am a business owner. When I started my business, a baby was born. I quickly learned that nobody cares about my baby as much as I care about my baby. As I hired more and more people, I began to look more and more for a servants heart than simply raw trade skill. I have discovered that if an employee has the basic skills necessary for a position (even if they are a little weak), is not arrogant ("I am God's gift to the world" mentality) and possesses a genuine servants heart, I WANT THEM! You see, I must have a servants heart for my customers. I must do whatever is in my power to please my customers. They are serving me by giving me money for goods and services. While I worked for Hewlett-Packard I remember one of the big managers saying, "We want to give our customers $1.10 in goods and service for every $1.00 they give us." I have had this attitude from the very beginning of my business.

Another way to put it is,

"Under-promise and over-deliver".

If you consistently do this for your customers or your employer, you will become famous. Who do you know that consistently delivers more than what is expected? I have a friend Tom that works for a large company. I am sure that his managers have no idea how much he does beyond the call of duty. If he ever left, I doubt there is a person in our town of 40,000 that could replace him. I have employees that do the same. Ponder this one,

Proverbs 20:6 - "Many a man proclaims his own loyalty, but who can find a trustworthy man?"

I have in my mind an "employee layoff list" in the eventuality of a major down-turn in business. I'm sure every manager or boss has one. Have you ever wondered how this list is determined? It really is quite simple. If there are only a few apples in the barrel, who can I live without?

Proverbs 22:29 - "Do you see a man skilled in his work? He will stand before kings; he will not stand before obscure men."

Denis Waitley, in his book, Seeds of Greatness (pg 71, this is MUST reading!), observes,

"One of the best-kept secrets of total success is that our rewards in life will depend on the quality and amount of the contribution we make."

ACTION: Begin today to deliver more than is expected. Do it when nobody could even know! When you give more than was bargained for, you make a loud-and-clear statement that basically creates an incredible need for you in the eyes of your customers or employer. Very, very few do this.

Focus #6 - Integrity

If you are to be successful, you must possess integrity. Integrity means those whom relate to you implicitly find that you can be trusted. You have principle and character. You cannot be bought. Consequently, you are morally consistent and predictable. Even when there is temptation for greater profit or pleasure, you do not cave in since you are pursuing long term success. This quality, though sometimes painful, provides great security for those whom you relate to.

Consider these,

Proverbs 2:7 - "He is a shield to those who walk in integrity."

Proverbs 10:9 - "He who walks in integrity walks securely, but he who perverts his ways will be found out."

Proverbs 11:3 - "The integrity of the upright will guide them, but the falseness of the treacherous will destroy them."

Proverbs 19:1 - "Better is a poor man who walks in his integrity, than he who is perverse in speech and is a fool."

Proverbs 28:6 - "Better is the poor who walks in his integrity, than he who is crooked though he be rich."

All of the people in your life are asking a perpetual stream of silent, subconscious questions of you. Can I trust you? Are you going to take advantage of me in a weak moment? Can I believe what you are saying, or do I have to strip out the exaggeration factor?, etc.

If people feel you can be trusted, they will entrust to your care responsibility over business, financial, and personal matters. I have seen many intelligent employees not be promoted because there was a question of integrity surrounding them. I have a number of employees whom I trust implicitly. I would entrust to these ones the care of my money and family, if I had to. If for no other reason, I will reward them with money because I know they will not cheat me, slack off working when I am not around, pilfer supplies, etc. While integrity may not always pay immediate dividends, you will be rewarded in the long run.

Proverbs 15:8 - "The sacrifice of the wicked is an abomination to the Lord, but the prayer of the upright is His delight.

When I was in Bulgaria a couple years ago speaking to the Sofia International Management Center on success in business, I spent time with the director of foreign trade and his wife. She made an interesting observation. They had both been communist party members, even to the point of shedding tears of joy when they received word that they were accepted into the party. She related a story about how her second child was very difficult to deliver. During that hard time, she cried out to God. She said, "Us communists are funny, when everything is going OK, we are atheists, but when we are in over our heads, we cry out to God." Hmmm.... sounds like many others I know.

General Charles Gordon of the British military said, "If you tell the truth, you have infinite power supporting you; but if not, you have infinite power against you."

Wouldn't it be nice to have a prayer or two answered when you know that you are beyond your means? By most standards, most would say that I am a successful person. But I am not at all embarrassed to say that there have been thousands of times when I have bowed my knee to ask for God's help and wisdom...

... in business decisions
... in family "stuff"
... in maintaining some level of sanity
... in helping others
... in managing my finances

Consider this observation by Denis Waitley (Seeds of Greatness, pg. 86), "We seek honesty in our relationships by volunteering it in all of our own actions. Even if we don't receive honesty in return in most of our daily transactions, as long as we never waver from our own deeply rooted values, the score will add up in our favor in the long run. This is one of the most basic, most obvious and, unfortunately, least understood principles of life. Good actions get good results in time."

J.A. Froude said, "The moral law is written on the tablets of eternity. For every false word or unrighteous deed, for cruelty and oppression, for lust or vanity, the price has to be paid at last."

ACTION: Take a realistic inventory of your own integrity. If two people who know you were asked if they thought whether you are honest, what would they say? Begin today taking little steps to make past wrongs right. Pay back what you've taken, apologize where you have offended. Swallow your pride and admit you were wrong where you were wrong. Begin to cultivate honesty and integrity.

Focus #7 - Emotions and Mental Attitude

Recently, I have come to recognize the power of my own emotions and thinking. As I focused on a certain kind of negative thinking, I usually ended up living the negative thing out. It is so amazing to me how powerful our mind is. With our emotions and thinking we can...

... heal or hurt others
... create or destroy our reputation
... birth or kill new ideas and inventions
... become wealthy or be impoverished
... bless or curse our children
... stay status quo or rise to new levels

It is essential that our heart (the word I'll use for emotions and thinking) stays healthy. Keeping morally clean. Staying free of negative "nay-saying". Maintaining an "anything is possible" attitude.

This means...

... when a suicidal thought enters my mind after a fight with my spouse, I immediately usher it out
... when my friends say that my new business idea will never fly, I listen politely and do not let the negativism affect my plans
... when I fantasize about ways to cheat on my timecard, I dismiss those thought rattlesnakes
... when the "authorities" say that the new technological gadget I invented will never be accepted by the public, I push all that harder to show their wrong
... when my neighbors wife looks more inviting than mine, I change the channel in my mind

Consider the following statements taken from original newspaper articles reflecting the "authorities" of the day:

1840 - "Anyone traveling at the speed of thirty miles per hour would surely suffocate."
1878 - "Electric lights are unworthy of serious attention."
1901 - "No possible combination can be united into a practical machine by which men shall fly."
1926 - (from a scientist) "This foolish idea of shooting at the moon is basically impossible."
1930 - (from a scientist) "To harness the energy locked up in matter is impossible."

There's more.

Consider Gaston Nassens, a French researcher who was exiled from France a couple years ago because he demonstrated a 75% complete recovery rate from TERMINAL cancer patients. This even included AIDS patients. Hmmm.... Why haven't we heard about his 714X compound that has zero side-effects? This blew all the circuit breakers in the minds of the French medical authorities (how can somebody heal cancer without cutting, burning, and poisoning; how dare anyone try...). He now lives in Quebec and is still persecuted by the Canadian Medical Corporation. Too bad. For a shocking action of what one man can do because he has cultivated his thinking as to continue on in spite of all the negativism and persecution, (especially if you or someone you know has been or is touched by cancer or AIDS), read the book by Christopher Bird entitled, The Trial and Persecution of Gaston Nassens. Keep in mind that our first

President, George Washington, died at the hands of the American Medical Association when two of his personal physicians administered blood letting and mercury potions. The authorities are not always right!

Consider Kondoff and Busch, two Hewlett-Packard engineers, who I am told, presented the idea of software disc caching to management. The idea was flatly rejected. They persisted on their own time, as the story goes, and when they demonstrated a working model with great results, management then made a killing. This idea alone saved many HP 3000 computer customers from certain "performance death". What if they caved-in mentally and allowed the negative feedback to discourage them? Many HP 3000 customers (series 40, 44, 64) would have defected to IBM.

What if all the above folks had not kept a their mind free from negative influence?

What does the Bible say about success as it relates to our minds? Much more than you would guess. Here's one of my favorites:

Proverbs 4:23 - "Watch over your heart with all diligence, for from it flow the springs of life."

This passage makes a specific point. It stresses the importance of guarding ones heart. The heart is basically your mind and your emotions. It is from the heart that our self-image, decisions, etc. originate.

Truly, our heart is a CPU of sorts. It can be programmed. It will mobilize all the support systems in your life to accomplish whatever you really think and believe you can do.

The foregoing is the good news; here's the bad news:

Proverbs 10:12 - "Hatred stirs up strife, but love covers all transgressions."

Proverbs 14:30 - "A calm and undisturbed mind and heart are the life and health of the body, but envy, jealously and wrath are as rottenness of the bones."

One of the main reasons people do not succeed in life, even if they are extremely talented, is because they let their heart get cluttered. Bottom-line, if your emotions are encumbered...

... with the trauma of a divorce
... by a fight with a friend
... due to disappointment because of failure
... from the echo of words by your parents
who said you were never going to amount to anything
... by the teacher who mocked and belittled you

then those cassette tapes in your heart (which are like those endless play tapes in answering machines) will literally paralyze you from becoming great at any endeavor.

The tapes must be erased and re-programmed.

"Little, vicious minds abound with anger and revenge, and are incapable of feeling the pleasure of forgiving their enemies." Philip Dormer Stanhope, Lord Chesterfield

Ours hearts were not intended to be cluttered with negative thoughts. In fact, the human machine conks when greed, bitterness, lust, anger, hatred, low self esteem, and un-forgiveness reign. If your emotions are filled with any of the above, you will not, by definition, become a holistically successful person.

If you are serious about being firmly committed to the journey of success, you must declare war on the items in the list above. One tactic is to take each item and begin to think and do just the opposite. Overcome greed by being generous. Pay the bill for lunch, buy four boxes of over-priced cookies from the girl scouts, give some time and money to helping out in the local soup kitchen, etc. For those whom have hurt you, release them. Simply cancel the bill of debt they incurred, and don't pick it up again.

ACTION: Take inventory as to what thoughts, fantasies, desires, etc. are living in your heart. Imagine them being played on a projector screen. Which would you say are good for you and which are not? Write each down and create a game-plan to counter each negative one. Be ruthless with the "heart-rattlesnakes." The anti-venom for each rattlesnake is exactly that... anti-venom. Replace un-forgiveness with canceled debts, low self-esteem with all the good things you've accomplished, lust with genuine love in sexuality as God intended, etc.

FOCUS #8 - Relationships

A successful person does not become so by living in isolation. You and I need healthy, loving relationships...

... In our family
... With friends
... In business life

Cultivating friendships apart from simply using people is rather rare these days.

It is interesting to note that the only commandment out of the "big ten" that has a promise associated with it is, "Honor your mother and father, that your days may be prolonged in the land..." If you have a bottleneck in your relationship with your parents, it will be a hindrance to your success.

"The sanctity of marriage and the family make the corner-stone of our American society and civilization." James Garfield

"The business done in the home is nothing less than the shaping of the bodies and souls of humanity. The family is the factory that manufactures mankind... This triangle of truisms, of father, mother and child, cannot be destroyed; it can only destroy those civilizations which disregard it." G. K. Chesterton

I am well aware of the pressures on families these days to simply survive, let alone be able to create a quality environment in which to raise children. Gee, I hate to sound dramatic, but here goes anyway... In order to ensure the survival of civilization as we know it, strong families are necessary. It is so unfortunate that much of the agenda of our day is, perhaps unknowingly, destroying many of the very things that caused the United States to gain exponential success. I do not wish to exclude other nations that have also done well, however. I am simply very afraid for my own country.

Simply put, if you wish to be successful, as I have been defining it, you need to be connected to a family that you can care for and which cares for you.

How do relationships get destroyed? Through many different ways, but at the core it is really very simple. Pride, selfishness, and greed are usually very prominent. A successful person despises these things and counters them by doing just the opposite.

Proverbs 6:32 - "The one who commits adultery is lacking sense; he who would destroy himself does it."

Proverbs 7:25 - "Do not let your heart turn aside to the ways of an adulterer."

This goes for any relationship that falls outside the bounds that God created. The cost for violating some of the laws of life physics is not always immediately seen. If, for example, you put gasoline in a Chevy 6.2 liter diesel engine, you would immediately see the cost. There is a good chance that the engine would be ruined in a short time. Running unleaded gasoline in an older engine designed for leaded gasoline would also cost. It would simply take longer to assess the damage. The same is true with relationship costs. A divorce usually costs money up-front. I have first-hand experience of how much divorce costs, as my parents are divorced; most of the real price is not seen immediately. Later, as the children experience emotional difficulty, as the divorcees experience various levels of pain, what kind of price tag should we attach?

In my opinion, the price of divorce is much greater than the price of forgiving, kissing and making-up.

Successful people deal with their sexuality in a way that fits what the manufacturer intended. Hmmm... what does that mean, Bob? Most people feel that God looks down on sex, perhaps because it is so fun... Somehow we have a twisted notion that God is out to ruin our day. *O Contrare*... In fact, He did not say "thou shalt not..." He said, "thou shalt, and thou shalt enjoy it when thou shalt!" Sex is just like any other area of life's laws of physics. Sex is best enjoyed in the context wherewith it was designed. So, when sex operates outside of the manufacturers specifications, just like any other area, there is an associated cost. The human machine, and society, begins to conk; perhaps not right away, but over time. The most successful people understand that sex is to be enjoyed... but also respected. Dysfunction in our sexual side leads to all sorts of breakdown, physically, emotionally, etc. If your sex life is out of control, it will affect your ability to excel in other areas of life.

Proverbs 13:20 - "He who walks with wise men will be wise, but the companion of fools will suffer harm."

Successful people hang around the same. I have always tried to maintain friendships with people who pull me up rather than pull me down. The fact of the matter is that if you associate with slothful, selfish, un-disciplined people, you will typically descend to their level.

ACTION: Mend any breeches that exist with your parents if they are still alive; do the same with your spouse, children, friends, and business associates. Cut off all damaging relationships (adulterous, people who "bring you down", etc.). Begin to inject genuine love and kindness into all those relationships. Send small thank you notes at the time they are least expected. Begin to figure out what communicates love to each person in your life. Forgive and Forget. Swallow your pride; admit you are wrong, even if is only 10% your fault. Although your wife might find a new fishing boat interesting for Christmas, that is probably not what communicates love to her! I know, I've done such stupid things before.

Focus #9 - Self-Government

Leo Tolstoy observed, "Everybody thinks of changing humanity and nobody thinks of changing himself." I believe that the only lasting change that is possible, is when people take responsibility for their actions and begin to live responsibly. Consider a few tid-bits of ancient wisdom,

"He is a governor that governs his passions, and he a servant that serves them." Benjamin Franklin

"A two-year-old's temper tantrum is a mark of the child's age. A twenty-five-year-old's tantrum is a mark of the persons' upbringing." Anonymous

"A man must first govern himself ere he be fit to govern a family, and his family ere he be fit to bear the government in the commonwealth." Sir Walter Raleigh

"He who reigns within himself, and rules passions, desires and fears, is more than a king." (Milton)

"He who lives only to benefit himself confers on the world a benefit when he dies." Tertullian

Proverbs 14:29 - "He who is slow to anger has great understanding, but he who is quick-tempered exalts folly."

Proverbs 21:17 - He who loves pleasure will become a poor man; he who loves wine and oil will not become rich."

Proverbs 25:28 - Like a city that is broken into and without walls is a man who has no control over his spirit."

Self-government may be the most difficult of all the areas of focus. If you do not possess a measure of self control, then it is possible that you will not do well with any area in your life.

I really like the last Proverb quoted above. You and I know both people that appear to have very little or no control over their spirit. You know the type, over-drink, over-eat, rage easily etc. Let their life be a lesson to you. I like playing sort of a game sometimes, particularly with older people that I meet. I ask myself the question, "Would I like to be like he/she when I am that age?" I must say that there are very few that I would like to be like!

To gain control of your life, particularly if it is out of control, you need to take evasive action, and quit letting your body and laziness get the best of you. The game is won or lost here.

Successful people are masters and not slaves... to anyone or anything, especially to their own desires and appetites.

ACTION: Make it a habit to do one thing a day that strengthens your ability to be the master of your appetites, habits, attitudes, etc. This may mean not having a greasy, sugary donut on your break. It may mean taking a few extra minutes away from the TV or newspaper to look your kids and spouse in the eyes and listen to them talk. It may mean learning a foreign language, or taking some time to read a chapter in the book of Proverbs (did you know that their are 31 chapters in Proverbs? One for each day of the month... how convenient!). The self-discipline of reading a chapter each day of this book of the Bible accounts for much of the success I enjoy. Each time you do something (eat, drink, TV, etc.). Ask yourself, "who's in control here? Me or my appetite. If your appetite is winning, do whatever it takes to make it your slave.

Focus #10 - Knowledge and Wisdom

For this section, I am going to first barrage you with a number of age-old wise-bites.

Proverbs 1:5 - A wise man will hear and increase in learning, and a man of understanding will acquire wise counsel."

Proverbs 2:6 - "For the Lord gives wisdom; From His mouth come knowledge and understanding."

Proverbs 3:13-16 - "Happy and fortunate is the one who finds skillful wisdom, and the one who gets understanding. For the gaining of it is better than the gaining of silver, and the profit of it than fine gold."

Proverbs 8:10-11;8:17-21 - "Take my instruction (wisdom's), and not silver, take knowledge rather than choicest gold. For wisdom is better than jewels; and all desirable things can not compare with her. I love those who love me; and those who diligently seek me will find me. Riches and honor are with me. Enduring wealth and righteousness. My fruit is better than gold, even pure gold, and my yield than choicest silver. I walk in the way of righteousness, In the midst of the paths of justice, To endow those who love me with wealth, that I may fill their treasuries."

Proverbs 9:10 - "The fear of the Lord is the beginning of wisdom, and the knowledge of the Holy One is understanding."

Proverbs 15:14 - "The mind of the intelligent seeks knowledge, but the mouth of fools feeds on folly."

This area of knowledge and wisdom is also fundamental to successful people. Unfortunately, many do not read more than a trash novel on a long plane trip. If you are to be relevant in the information age that we live in, you better be continually be informed.

I like what Denis Waitley said regarding knowledge.....

One of the best kept secrets of total success is, "a large vocabulary - which implies broad, general knowledge - characterizes the more successful persons, regardless of their occupations."

Take time to read and think each day. Take on a new task. I have made it a habit to earn some new license or accomplish a major feat each year for over fifteen years. I have enjoyed the challenge of getting four successive classes of my amateur radio license, my pilot's license, writing two books, learning to speak Spanish and Romanian, starting a blueberry farm, preparing and delivering "how to start a small business" seminars in Bulgaria and Romania, etc. All of these have taxed by intellectual faculties and have kept me fairly sharp in my thinking.

Wisdom is knowledge applied in an appropriate way. Often, the only way to learn wisdom is to either listen to old-timers or learn it the hard way. Even though I am 36 years old, I still find some recklessness in me that needs to be tempered with wisdom.

Successful people are not only knowledgeable, they are learning to become wise.

ACTION: Learn a new word each week and use it. Most adults use no more than 300 words in their speech. Turn the TV off, except for some short news or educational programs. Start subscribing to some off-the-beaten trail newsletters. Don't read the newspaper, simply glance at key headlines. Honestly, has the amount of time you spent last year reading the paper, paid off? Write one idea down each day or week for ways you can create something new and useful (most of my business, software products, books, articles, etc. emerged from this discipline). Make it a goal to read one book per month on something relevant to helping you increase your wisdom and knowledge base. I will again reiterate what I said earlier about reading a chapter each day in the book of Proverbs in the Bible (hint: get some other version than the King James like the NIV or Living Bible... no offense King Jimmy, but that old English is pretty tough!)

The Missing Ingredient for Success

If you work hard at mastering the areas of focus discussed above, you will do well. But if you ignore this section, you will never experience the kind of success that I have enjoyed for twenty years. Furthermore, the ten areas of focus will merely be a pile of loose spokes not connected to a hub. God is that hub. Consider this,

Proverbs 3:5-8 - "Trust in the Lord with all your heart, And do not lean on your own understanding. In all your ways acknowledge Him, And He will make your paths straight. Do not be wise in your own eyes; Fear the Lord and turn away from evil. It will be healing to your body, And refreshment to your bones."

If you want success, I mean real success, and not some short thrill of winning a lottery or a two week fling in the Bahamas, you must reckon with God Himself. Is this relevant for business and your tasks as a manager or programmer? YES!

... power
... reputation
... or God

But, remember, you will serve something or someone.

A really successful person does not look to his or her own resources, but rather attempts to agree with God and run their life by staying on the railroad tracks that He has designated for human kind to thrive on. Simply put, these tracks are what being righteous is all about. Even in the times in my life when I could have compromised and lied or cheated to gain a bit more money, I have mostly chosen not to. As a result (I hate to get too mystical here... but here goes anyway), I have seen the incredible circumstances which have resulted in God Himself blessing me. I find that most people feel that God is like a mean old man that likes to rain on their parades and take candy from babies. What I have found, however, is that He not only is not against me if I am established in relationship with Him on His terms, He is aggressively for me. Believe it or not, I have found this to be true with respect to all aspects of my life. Check out this tid-bit from the Bible:

"For the eyes of the Lord move to and fro throughout the earth that He may strongly support those whose heart is completely His."

II Chronicles 16:9

You want success? Can you imagine having the God who created the universe "strongly supporting" you. I invite you to contact me sometime. If you do, I will fill your ears with stories, almost beyond belief, how God has performed incredible things in my business, marriage, etc. Do you see why I made such a strong point at the beginning of this paper as to the absurdity of trying to be successful without factoring in God Himself. It is not simply about a knowledge of God, or sitting in some pew... It is knowing Him and knowing that He knows you. That all of your concerns (office, family, etc.) are also a concern to Him. Consider what Blaise Pascal said, "The knowledge of God is very far from the love of Him."

"There once was in man a true happiness of which now remain to him only the mark and empty trace, which he in vain tries to fill from all his surroundings, seeking from things absent the help he does not obtain in things present. But these are all inadequate, because the infinite abyss can only be filled by an infinite and immutable object, that is to say, only by God Himself." C.S. Lewis

Let's talk about prosperity...

...In your important relationships
...With your children
...With your money
...In your data processing department

In the light of the following bit of ancient wisdom,

"Blessed - happy, fortunate, prosperous and enviable - is the man who walks and lives not in the counsel of the ungodly [following their advice, their plans and purposes], nor stands [submissive and inactive] in the path where sinners walk, nor sits down where the mockers gather. But his delight and desire are in the law of the Lord, and on His laws - the instructions of God - he habitually ponders and studies by day and by night. And he shall be like a tree firmly planted by streams of water, ready to bring forth his fruit in its season; his leaf also shall not fade or wither, and everything he does shall prosper." Psalm 1:1-3.

I believed that this was true twenty years ago. I lived doing exactly what it said for twenty years. It works. This is success. But please, please understand that I am not promoting organized religion as it is generally thought of. I have personally observed and experienced that God

intended relationship with Himself to be SIMPLE. Most of what well-intended religious folk have done is muddied the waters. That's not to say that some religious organizations aren't doing a good job of pointing people to God. But I have been so offended by some of the tactics and antics of many, that I have chosen to ignore most of them.

I like what William Penn said, "Men must be governed by God or they will be ruled by Tyrants." What I have been saying simply is that you will be the slave of something. In our relentless search for security and meaning, we will serve one or more of the following ...

... money
... other people
... sexual desires

About the Author:

Robert Lund has worked with the HP 3000 since 1979 in various capacities such as: Programmer/Analyst, System Manager, Systems Engineer (Hewlett-Packard, Fullerton, California), Technical Support Engineer for a large HP 3000 VAR (Summit Information Systems), and a System Performance Consultant.

He has written numerous technical articles for HP 3000 journals and magazines. Robert has provided technical support to hundreds of HP 3000 sites throughout the United States. Currently he is the president of Lund Performance Solutions, the leader in HP 3000 performance software, training, and consulting. He and his wife and four daughters live in rural Oregon. He may be reached at (503) 926-3800.

Paper #1007

Management Ethics - An Oxymoron?

Louis R. Mills, CDP
Bio-Rad Laboratories, Inc.
1000 Alfred Nobel Drive
Hercules, CA 94547

work: (510) 724-5613
CompuServe: 73631,223

Are you a professional? Are you good at what you do? Can you point to some things you're proud of? Do you know some recent work of yours that could use improvement?

If the answer to any of these is yes, then you probably have high personal standards, a set of "internal ethics" about who you are and what you do. If you have to answer "no", then you don't give a damn about ethics and this paper is not for you.

Have you noticed that everybody's a professional these days? No matter what the occupation, there seems to be a group that forms a professional association, collects dues, holds monthly meetings with speakers on how to improve the profession, perhaps has seminars for new members of the profession, and they adopt a code of ethics and maybe even put together a committee to review ethical practices. Doctors do it. Lawyers do it. Cosmeticians do it. Even "Refuse Disposal Technicians" do it. No matter what role you have in our society, there's a professional group willing to support you.

I've been a member of DPMA for over 20 years, and last year I served on their National Ethics Committee. Our job was to enforce the "Standards of Conduct" among our 20,000 members. But we didn't have one complaint all year. Not one. Doesn't that tell you that members of DPMA are the most ethical people you'll ever meet? Do you think the AMA goes one year without a complaint? Do you think the Bar Association in Montana goes a year without a complaint? Are you telling me that "data processing" is the most ethical occupation on the planet?

I don't think so. I think it a lot more likely that most DPMA members don't even know they signed a Code of Ethics. I think it even more likely that ABSOLUTELY NONE OF THEM gave a copy to each and every client, internal and external, and said, "If you have a problem with any of my work, please report me to Lou Mills." I wager you that if 1% of them had, Lou Mills would have spent the entire year in ethics committee meetings.

How many of you have signed a statement of ethical practices recently? How many of you have it proudly posted in plain sight in your office? How many of you discuss ethics at least once a year at a staff function? How many of you have no idea what ethics is?

OK, let's define some terms. Today, ethics is argued on two levels:

- Metaethics
- Normative Ethics

Metaethics (literally "about ethics") is the analytical study of the discipline of ethics itself. The term came into use only in the 20th century and deals with meanings such as right, wrong, justice, fairness, obligation, etc. to see which of these concepts is basic and what are the relationships between facts and values.

You may see the term "normative ethics" used to define the everyday problems you and I have to deal with. Another term that's used is "traditional ethics".

Some people wonder why there are even two academic approaches to ethics. There are arguments about whether the two should be separate, but the principle value in metaethics is that it takes a morally neutral approach to the study of ethical issues. Metaethics uses several methodologies:

Naturalism (H. Spenser and John Dewey)
Cognitive
Intuitional (H. A. Prichard and W. D. Ross)
Subjectivism
Emotivism (A. J. Ayer and C. L. Stevenson)
Imperativism (Rudolf Carnap)

And if you're really into self-inflicted punishment, you should know that the last three methods listed above are classed as techniques of the school of Moral Skepticism, which simply means that there are no moral truths, that moral principles cannot be proved and that morality has no rational basis. Moral skepticism is widespread in our western culture. In fact, one of the tenets of moral skepticism is that there is no one correct moral code for all times and peoples, that each group has its own morality relative to its goals and values. This is termed "ethical relativism."

At the other extreme is the code of "Anyone should be free to do anything he or she wants unless it hurts someone else." Some argue that this is inherently flawed because our very existence hurts others by consumption of limited resources, yet our suicide might preclude further contributions to the world and hurt those close to us!

You see, in day-to-day life, ethics is a vast gray area. It isn't a neatly defined set of rules and practices that you can post on a wall. And in our multi-cultural, multi-national and global communities, one person's moral behavior is another person's cultural insensitivity. And it's going to stay that way until we have a common, uniform culture around the planet.

What do you learn if you take a Business Ethics class at Harvard, UC California or any other university in between? You get a curriculum based on the moral concept of altruism - doing good simply because it is right. Right in whose terms? If ethics and the interest of your business conflict, you are told to always do the ethical thing. Even if it causes the company to close its doors.

But there has been a great deal of interest lately in applied (normative) ethics. Western economies are not as strong as they used to be, so making an ethical decision that shuts down a company can have more impact on all concerned, even other business that trade with the company in question. And sometimes the right behavior is hard to identify. Is it morally correct or politically correct? Does it meet the letter of the law or the spirit of the law?

Remember those professional associations we talked about earlier? Some of them have some very difficult issues to deal with. The medical profession has traditionally prohibited advertising. Attorneys used to shun it until recently. Now, attorneys have slick

television ads that imply that no matter how you got injured, you've got a right to sue. Doctors take the Hippocratic oath - to save lives, to cure disease and to alleviate suffering. But now these goals conflict because we've introduced devices and cures that prolong life with increased suffering or burdensome costs to the patients and their families. What of the parents of a severely birth-deformed baby who have to decide how much effort is expended to keep it alive, even though there appears to be no permanent cure that will render the infant into an active, productive life? Have you heard the expression "rationing medicine" recently? If you haven't, you will, particularly when the President's Health Care Proposal hits the floor on Capitol Hill.

Talk about moral dilemmas. Is abortion a woman's right to control her own body or is it palin and simple murder? Is it right to cause animals suffering and death in the name of medical research? Should we deny people access to this country because we can't find enough jobs for the people who already live here? What of the simple moral dilemma of the catholic priest who hears a confession and knows that a crime has been committed?

The nice thing about ethics is that you can argue both sides of any issue until the cows come home. What about some of the simpler, down home issues? Today, e-mail and bulletin boards are becoming common. Access to systems like the Internet are growing phenomenally. People say things in an e-mail missive that could be libel or slander in court. What about the poor guy at Bank of America in California that typed a love note to his sweetheart three floors up into his e-mail system? He told of his plans for their weekend get-away and went into some detail about how he planned to satisfy some of his desires. Then he hit the wrong key and the message went out to every e-mail station in the building.

I don't have answers to all these questions. I just don't know. I can talk about some of them dispassionately, but only to a degree. With all the emotional baggage that you and I have, with our cultural backgrounds, our technological training and our bias for action, we often want to make a decision, just to get the problem behind us.

Now, I've given you some fancy terms and the names of some experts in the field for you to drop when you get back to the office. Look some of this stuff up and give it some thought, but don't expect clear-cut answers.

Let's talk for a while about some of the issues you're likely to face at work. A couple of these have actually happened to me. I can't say I've always done the ethical thing.

Scenario #1

You have interviewed about a dozen candidates for a position that reports directly to you. Each candidate has been interviewed by you, two of his prospective peers and the technical recruiter in your H. R. department. Your boss did not interview any of the candidates. In a meeting with your boss, she asks how the interviewing is going and you show her the list of candidates and tell her the two who appear to be the best for the job. She listens for a second as she scans the list, then tells you, "No, I want you to hire this person.

I've worked with him before, he's really sharp and it will be good to have him on our team." You point out that each of the interviewers rated the candidate on the bottom half of their lists, but she announces, "Hire him or I'll restructure the department so that I can hire him."

What would you do?

Scenario #2

It's early December and your mind is on the new system that went in last week and has to close the year-end accurately and on time. The month-end and year-end processes haven't been fully tested and what testing you have done hasn't been promising. You submitted next year's budget months ago and now the Controller is in your office asking you to sign off on the numbers he has on some printout so he can tell the CFO that he has the final numbers from each department. You ask for a minute to check the numbers, pull out your worksheet from August and the numbers don't match! Some key expenses and income accounts have been inflated and they all point to a pet project that your boss wanted to do, but for which he couldn't get approval on his own budget. It's obvious that he juggled the numbers and didn't advise you of the changes.

Do you sign?

Scenario #3

It's your first year with the firm and you've kept your department just under budget for the first three quarters. But the new system you proposed, promoted and installed has some real performance problems. You've analyzed the bottlenecks and know you can solve them by buying an inexpensive software utility and one piece of datacomm hardware, both of which are within your purchasing authority. But there's no money in the budget for more hardware or software. The guy that had the job before you suggests that you simply buy the two items and charge it to your training account, since you have enough unused budget there to cover the costs. He tells you that's what he used to do. You don't think you are going to use all that training budget, anyway, so what's the harm?

Do you buy the software utility and datacomm unit? Do you charge it to your training account?

Scenario #4

Your offices are located in the same building as Ima Big Computer Co. They recently bid a new computer system for your firm, but you went with a competitor. Now the competitor's sales rep. is offering to provide you with free training if you will let her bring in several prospects to see your installation and talk to your users. You suspect that

they are trying to impress prospects with the fact that they beat out Ima Big in their own building.

What do you tell the rep?

Scenario #5

You return from a week out of the office and your in-basket looks like someone dumped their trash in it. You quickly set about sorting through the mess, putting things in piles by priority as you open and scan each piece of mail, make notes on things that can go back out, etc. About half way through the pile you stop. Someone has sent you an incredibly intimate letter, but you don't recognize the name. You pick up the envelope and discover that it's for a married coworker and was misplaced in your stack of mail.

What do you do next?

Scenario #6

You have a network analyst working for you and you discover that he has been storing hundreds of different computer games on the MIS file server and sending notices out to the users telling them which ones he thinks they'll enjoy. This bothers you because you know that many of these games must have been downloaded from outside services like Internet or CompuServe, some may be proprietary, but all of them take up disk space and valuable resources when they run. You suspect that some of the analyst's work day is tied up answering questions about games for novice users. You worry about connect charges to those outside services.

The company has no policy about computer games. What do you do?

Establishing an Ethics Policy

If your company doesn't have an ethics policy, you can establish one for your department. Pick an area that's a problem or concern, like e-mail privacy, then do the following:

1. Consult your company legal department. Find out what their view is on the issue and discuss the various strategies and concepts you're concerned about.

2. Consult with upper management. Let them know what you are doing and why. If you can show them recent management articles on the topic, or news articles about related crimes, it might help them understand the importance of the ethics policy.

3. Find out what other companies like yours have done. Don't forget to ask your vendors for input. They often have strict ethical practices spelled out for their own staff.

4. Talk to the people impacted by the policy. How do they feel? What problems do they foresee?

5. Ask yourself, "Is it enforceable?" If you have to "slap hands", could you have even more problems than you have now? Will other departments support your policy?

6. Decide on the policy and make it public. When you publish it, give it an effective date and a duration. Make sure you schedule a review of the policy before the expiration date. That's your opportunity to improve it or bury it.

There is an organization called the Computer Ethics Institute. I don't have an address for them, but they did provide this list in a recent DPMA publication:

Ten Commandments of Computer Ethics

1. Thou shalt not use a computer to harm other people.

2. Thou shalt not interfere with other people's computer work.

3. Thou shalt not snoop around in other people's computer files.

4. Thou shalt not use a computer to steal.

5. Thou shalt not use a computer to bear false witness.

6. Thou shalt not copy or use proprietary software for which you have not paid.

7. Thou shalt not use other people's computer resources without authorization or proper compensation.

8. Thou shalt not appropriate other people's intellectual output.

9. Thou shall think about the social consequences of the program you are writing or the system you are designing.

10. Thou shalt always use a computer in ways that insure consideration and respect for your fellow humans.

WHO'S DOING WHAT? CHANGING JOB DUTIES IN HP SHOPS

Lynn A. Novo, CPC
Network Systems Company
9300 Shelbyville Rd. Suite 502
Louisville, KY 40222
502-327-7746

Let's start with some fun. The following job descriptions are really used in HP shops today. Try to guess the job title the companies use for each of them. There's a prize for the person with the most correct answers. And because as good managers we're all advocates of quality improvement, there will also be a prize for the person who can benefit most from this session.

What titles would you give for the following?

1> Perform UNIX system administration on twin HP UX systems. Perform Sybase data base administration. Write programs to automate computer operations. Perform all s/w upgrades for the systems and data-base. Initiate security policies and procedures. Develop and implement a contingency plan in the event of a disaster. Assist operations staff when necessary. Evaluate/recommend 3rd party s/w that can enhance/improve system productivity.

2> Day to day duties: General HP3000 systems management particularly managing interactive users, production batch jobs, system configuration, and troubleshooting. Close interaction with several user communities to support and/or resolve problems involving the HP3000 system environment.

3> Provide user & group configurations for access to network resources. Preserve system security, user set ups, and configuration files. Configure and secure applications on the LAN. Improve system utilization levels. Integrate E-mail to external mail systems, LAN based scheduling, and faxing. Maintain back up systems for networks.

The answers that you give may in fact be more appropriate than the ones the managers of these employees assigned. But the fact that we have such a range of possibilities

is an indication of how challenging it is to decide 'who should do what' in your department.

When I started recruiting for the HP community exclusively 15 years ago, my first computerized database had fewer than ten job title codes. My secretary was almost always able to enter the job title listed on the resume as the title code in our database. Times have changed, haven't they? Now my secretary is replaced by the voice mail's auto attendant and resumes are scanned into our database! We now track three times as many distinct job functions in the HP job market.

I don't need to highlight the sweeping technical changes in our industry. But those changes mean that the people on your staff have different and new job duties. It's probably a good time to look at a few job functions in today's HP MPE and UX environments and share some ideas about how the workload is distributed. Along the way, we'll mention both trendy and traditional job titles and hopefully give you some good ideas to use the next time you write your job descriptions.

Fifty minutes is not enough time to discuss all the job functions in an HP shop, so let's focus on a couple problematic areas, namely system management/administration, database administration, and network administration. I'll share my findings from talking with management and human resource professionals as well as the individuals who actually perform these functions. To facilitate an exchange of ideas and to demonstrate the differences based on size and technical environment, I will be looking for audience contributions from each of these groups:

Small dept	MPE environment
Small dept	HP UX environment
Med dept	MPE environment
Med dept	HP UX environment
Large dept	MPE environment
Large dept	HP UX environment
Any size dept	MPE and HP UX combined

SYSTEM ADMINISTRATOR VS. SYSTEM MANAGER- WHO DOES WHAT?

Is there a difference between a system manager and a system administrator? Do any of the MPE shops have a staff member with the title of system administrator? Does

anyone using HP UX technology **not** have someone designated to perform this function? Should we try to be consistent in how we name this important function in HP UX shops? And if we're confused, what about Hewlett Packard itself that named its HP UX utility SAM, System Administration Manager?!

Look at these direct quotations from two resumes that crossed my desk recently:

"Systems Administrator for HP 3000 series 48, 70, and 950. Responsible for all system management functions, trouble-shooting, end-user support, data base administration..."

"System Administrator - Perform system management duties for HP 827...."

It seems that a "System Administrator" always does system management work but in addition, takes on other responsibilities.
Some have put forth the argument that an administrator has a higher level of systems knowledge than a system manager. I disagree. I do not see a system manager earn the title of system administrator with more experience, as if it were the next rung on the career ladder. Likewise, I have known many highly technical, indeed brilliant, technicians who have always been called (and prefer to call themselves) "System Manager".

The industry trend, based on my research, is that the title System Administrator most often is used in shops with HP UX machines. The System Manager title appears more often in MPE environments.

Using the system administrator title in HP UX environments makes sense because the complexity (or potential for complexity) of the UX system itself requires that the person in charge have competence in a wide range of skills. Marty Poniatowski, the Hewlett Packard Technical Consultant who wrote The HP-UX System Administrator's "How To" Book covers not only the daily system management functions using SAM but also performance monitoring, HP UX accounting, security, DOS integration, GUIs, as well as shell script and C programming. UX System Administrators can also be responsible for database administration, network administration, user support, and integration of multiple platforms.

What duties does the System Manager/ System Administrator have in your shop?

SYSTEM MANAGER/ADMINISTRATOR DUTIES CAN INCLUDE:

Add/delete users
Perform backups
System configuration
Capacity planning
Performance monitoring and tuning
Disaster recovery planning
Operating system upgrades and patches
Security
System conversions
Hardware and software support contract management
Network implementation/management/support
Documentation of operational procedures
Benchmark testing
Evaluation of software or hardware
Technical/user support
Shell script or other programming
Interface with vendors
Supervise operations staff
PC installation, troubleshooting, support
E Mail administration
Train operators, programmers, users

DATABASE ADMINISTRATION

Database administrator was definitely NOT a title I tracked when I began my HP recruiting career. In fact, it wasn't until relational databases became popular on the HP a few years ago that I wondered who had been minding the Image databases all these years. Do you know? Here are the titles of people I have heard are responsible for monitoring the Image database in their shop:

Programmer Analyst
System Analyst
System Manager
Data Processing Manager
Operator

When we move to relational databases, however, many companies are finding that they can't afford NOT to designate a DBA. As one manager told me:

"We thought our system manager could do the database administration, but we're finding that he's so busy with the system that the database work is falling behind. I have had to justify an additional headcount for a DBA to my management, even though we have a hiring freeze."

Although the system manager or administrator often is responsible for both system and database administration, you may need a DBA if your company is doing a high level of development using RDMSs.

For an individual who is designated as the Database Administrator, what are the typical job duties?

DATABASE ADMINISTRATOR DUTIES CAN INCLUDE:

Logical and physical database design
SQL programming
Optimization of application programs
Implementation of new databases
Database performance monitoring
Data security
Integration of multiple databases
Database recovery
Monitoring database and system disk space
Verification of internal data consistency
Technical support
Maintenance of database/ disaster recovery software
Database auditing to determine corruption
Standardization and maintenance of data dictionary
Application programming
System management

NETWORK DESIGN AND ADMINISTRATION

Our final area to explore is network design and administration. As we have seen, in companies with either a small staff or a small network, this function is often performed by the System Administrator. There are also companies that assign operators to maintain a LAN. (I have seen a number of operators move into a networking career by training for the CNE.)

It is important to note the distinction between the network <u>design</u> and the <u>maintenance</u> of it after it is implemented. Many companies have a consultant or Hewlett

Packard's own network specialists do the design. Then such a high level of expertise (and salary) are not needed for a permanent staff member.

Large companies can have a group of network specialists. There is a wide variety of titles for people in this group, but a career track does emerge, based on seniority and technical expertise. Stacked in roughly ascending order, I have seen the following networking job titles:

Telecommunications Manager (Network Manager)
Network Architect (Network Engineer)
Network Designer
Network Specialist
Network Administrator

Rather than give the job functions for each of these titles, let me offer a couple of categories of job duties that are applicable in the network area. Within each are job duties that you can include or delete for the particular level you are writing the job description for. As you will see, I have excluded most of the technical details and tried to generalize so that these could be appropriate for any networking environment.

NETWORKING JOB FUNCTIONS CAN INCLUDE:

NETWORK CONFIGURATION

- Design company-wide data communications network
- Implement bridges and interconnectivity devices
- Maintain hardware and operating system of file servers
- Support network communication links and cabling
- Plan network architecture and topology
- Implement network shell and system software
- Implement local T-1 communications
- Manage the corporate telephone system
- Configure, upgrade, maintain E Mail

NETWORK ADMINISTRATION

- Configure, optimize, and allocate disk space on network servers
- Establish and maintain network security
- Evaluate and select network h/w, s/w, and facilities
- Manage vendor and carrier relationships
- Plan for expansion; assess network requirements
- Maintain s/w compliance to license agreements
- Provide user and group configurations for network access

SYSTEMS ANALYSIS

- Analyze user's current and anticipated needs
- Investigate new technology and methods
- Improve system utilization levels
- Develop and publish standards and procedures
- Integrate data across applications and platforms

USER SUPPORT

- Create and provide user training
- Troubleshoot user problems
- Resolve h/w vs. s/w conflicts

As you look at these job duties, you will hopefully be able to identify the ones that are appropriate for the networking job descriptions you are writing.

We have discussed job duties in just a couple of the rapidly changing staff functions in HP shops. Hopefully, you can springboard from these job duties to customize job descriptions for your staff. These job descriptions should help you organize your department and establish "who's doing what".

Paper 1009

Training is the Issue: What Are Your Options

Scott W. Sarisky

First Data Corporation: Health Systems Group
10101 Claude Freeman Dr.
Charlotte, North Carolina 28201

(704) 548-3038

The year is 1964, you are a bus driver. Let's take a look at the skill set required to perform your daily job. You must know how to drive a bus, make change, ... basic everyday type things. There is probably a good chance that the things you learned in school (K-12) except for driving the bus, sufficed for your career choice. However, contrast that skill set with the skills required today for a bus driver at a major airport, let's say for a rental car company, i.e. you take people to their rental cars. Yes, you must know how to drive the bus but you probably do not need to make change. However, think about all of the other things this person has to do as part of his/her daily job. You are in direct contact with the customer, you are a customer service representative. You may be the only company representative that some customers see. And lets not forget that there is much competition, only a few steps in either direction. Look at the equipment in your bus: a computer terminal; a microphone; and even a TV camera...access to visual information, audio information and data information. Do you think the things learned in K-12 were enough for this career choice?

Let's go back to the early 60's and order a pizza. First of all you may not have even had the option to order a pizza in your part of the country. But if you could order a pizza it would have probably been the local family owned pizzeria who wrote your order on a scratch pad. Compare the skills required there with the skills required to work for Pizza Hut in Atlanta, Georgia. There is one phone number for the entire Atlanta area (for home delivery). The first thing they ask when you call to order your pizza is your phone number, which is a key to your record(s) in their ordering system. They know if you have called before, what you have ordered in the past, your address, the color of your house (descriptive information from prior calls), if you have paid with cash or check...etc. Again, the skills required are very different as compared to just 30 years ago.

These are examples of everyday type things with obvious skill set changes. What about big business; have the skill sets changed

there? Yes, they have, and we are all probably aware of that fact. How many of you have ever used the analogy of "it's not like working on an assembly line, doing the same thing over and over"? I know I've used that analogy many times in comparison to data processing related work. That analogy has lost it's meaning if you start looking at some of the newer manufacturing processes. Following are pieces of information that are examples of some of the changes that have resulted from the shift from the Industrial Age to the Information Age:

In advanced manufacturing plants, a production worker must actually manage a network of robots and other computer controlled machines instead of performing a few repetitive tasks on an assembly line.

Business Week
September 14, 1990

In factories across the U.S., workers are being asked to handle new technologies and expanded responsibilities, and many are ill prepared to do so.

Fortune Spring 1991

In the past, many skills learned while growing up could be applied toward the work in the Industrial Age. But, in the future, many of these lessons will not suffice....

Gwinnett Daily News
August 26, 1991

The way you compete today is through continuous improvement, continuous change and that takes continuous training of employees.

Steve Duscha
Former Director of California Training Panel

... companies are scrambling to find new ways to train employees fast, effectively and continuously. Says Jere Jacobs, assistant vice president at Pacific Telesis: "The problem is much broader than K through 12 education. It's really K through life."

Fortune Spring 1991

Faced with a shortage of qualified repair people, Sears is establishing a curriculum in a Chicago vocational high school that will give juniors and seniors the training they need to go to work right after high school as beginner Sears technicians.

Fortune Spring 1991

At Bethlehem Steel, America's number two steelmaker, the average employee is almost 46 years old, with 22 years of service. Before the year 2000, close to 50 percent of their current work force is likely to retire.

Fortune Spring 1991

To prosper in the nineties and beyond, companies must concentrate on the 3 R's: recruiting....retraining....retaining....

Fortune Spring 1991

We're going to compete by investing in our people and our skills....

North Carolina is not going to compete in the global marketplace on the basis of tax breaks, gimmicks and giveaways, we're going to compete by investing in our people.

Gov. Jim Hunt
Charlotte Observer 1/25/94

Colleges, businesses reveal plans to retrain workers from shipyard.

The Philadelphia Inquirer
December 9, 1993

The more we investigate what is happening in the business world today, the more we see change. Technology is growing at such a rapid pace, it's so fast that it seems impossible to keep up with it. I would imagine that only a few people today would disagree with the fact that training is an issue and we need to address it. However, before we decide what to do about the training issue, we need to understand how people learn, when they are best suited to learning and what can be done to enhance the learning experience.

Let's contrast two learning scenarios:

Scenario One.

A large, overheated auditorium in New York City, 400 people focused on a lecturer with an overhead projector and a foot high stack of overheads of pages from IBM manuals. This goes on for two and a half days; it was complex technical material.

After the course the peoples comments included words like, amazing, terrific and impressive.

Scenario Two.

A small crowd in an ultra modern viewing room in New York City. There was a 30 minute technical presentation. Everything was perfect: the graphics, the projector overlays, the audio modulation, everything went exactly right. The lights come on after the presentation, one person nudged another to wake up, another yawns and they all file out quietly.

In which one of these scenarios did learning take place? Judging from the comments, the first scenario looks as if it was "better" than

the second scenario. Why? The people in the first scenario had the need to know (NTK) the information. In the second scenario, everything was perfect but it appears that there was no need to know the information. When the learner is hungry for the information or has the need to know the information, learning can be easy and fun.

Next we look into learning styles. Behavioral scientists have been aware of the concept for many years. However, in recent years it seems to be coming up more and more. It appears that there are different learning styles, i.e. some learn better by seeing, some by hearing, some by touching (hands on) and some in combination of all of the above. It has been reported that only 1/4 of Americans learn in a passive mode; that is, by sitting in a classroom listening to a lecture or by reading a book. There is a tremendous amount of scholarly research in this area. For a moment, recall your learning experiences. Were they primarily passive mode, mine were. How are they teaching children today in your school district or local college? Anyway, as it turns out, most of us learn best by experience, either through observation or actual hands on. There is a great deal of research that supports this. Furthermore, I have witnessed this many times while teaching in college and in industry.

Therefore, when possible, be sure to help your employees with their need to know in a given area and as much as possible, provide a hands on training environment.

One item that really deserves to be brought to your attention is the fact that exposure to material does not equal mastery of the material. How many times have you heard someone say, or possibly you may have said this in the past, "you went to class, therefore you are the expert"? Just because someone has been exposed to information, does not mean they have fully digested it. What if the class they went to was 100% lecture and they are not part of the population that learns well in passive mode.

Following is information related to learning theories and styles:

ADULT LEARNING THEORY

- Adults are more independent, self directed in their learning.
- Adults have accumulated a reservoir of life experiences that can be a resource for other learning.
- The adult's readiness to learn is not linked to his or her biological growth.
- Adults expect immediate payoff or practical application.
- Adults are more problem/task focused rather than subject centered.

Data Training , February 1991

Adult learning theory emphasizes that adults must want to learn.... The lack of motivation will impair learning.

Training & Development Journal
November 1990

Training is the Issue: What Are Your Options
1009- 8

Workers can't be adequately trained unless companies first understand:

How people actually
- perform work
- acquire skills
- learn to solve problems on the job

Most corporate trainers wrongly adapt classroom forms of learning that emphasize theory.

Business Week
September 24, 1991

Sitting passively in a classroom is one of the least effective ways of learning.

Business Week
September 24, 1991

People learn by making and correcting mistakes....

Robert A. Bjork
The Mind's Eye: Enhancing Human Performance

Worker's learn best when three elements are present:

1. Motivation
2. Application.
3. Manager's endorsement.

Chambers and Asher
Speechworks, Inc.

Participant's superiors can quickly undo the learning that took place in the seminar if they insist on or encourage behavior that is contrary to what was taught.

Lawrence S. Munson
"How to Conduct Training Seminars"

It is the direct supervisor that sets the level of importance to tasks within a job. It is the direct supervisor, therefore, that motivates or de-motivates the employee in a training situation.

The employee must be motivated to learn, and use the learning on the job. The direct supervisor can manage that learning.

David G. Robins, HP
"Managing Employee Learning"
Interex 1991

The following thoughts, written by Mark W. Pine, Director of Human Resources Development at Norwalk Hospital in Norwalk, CT, define management's role in the training process:

There is no magic in training.
There are no thunderbolts from the sky.
Training provides the atmosphere in which to learn.
Training provides the information from which to learn.
Training provides the opportunity to understand, to see, to have the light bulb illuminated.
And that is where training stops.

The rest is up to you.
You must provide the employee time to focus on training.
You must provide the opportunity for the employee to use the newfound knowledge.
You must provide the openness that allows the employee to ask questions and ask for advice.
You must provide positive reinforcement when the employee succeeds.
You must provide correction when the employee makes an error.
If you don't or won't or can't do this, then do not send one of your employees to training.
Because there is no magic in training.

Lastly, there are many methods of delivering training material. One of the most common way is Instructor Led training. This is were the instructor gets up in front of the class and teaches from a course book. This is the primary we all have been taught. Also, this method is one of the easiest to put together. The scale here ranges from using good course development techniques to "throwing" a course together. The industry standard for instructor led course development is about 40 hours development for 1 hour of classroom time. This is an average, I have seen ratios ranging from 1:1 to 100:1. Just remember, you usually get what you pay for.

The next method of delivering course material is self study materials. This is very similar to instructor led except there is no instructor. This fact alone actually makes self study material potentially more difficult to develop because there is no instructor intervention. Therefore, the self study course developer most be more aware of what the learner is thinking while going through the material because there is no one to ask questions to. At times there may be an advisor close by for questions but how effective can that one person be in all areas?

The third method of delivering course material would be CBT (computer based training). It has other names such as CAI (computer assisted instruction) or CBL (computer based learning) and so on. Basically this method is a computerized self study with examples and exercises. With CBT you have a larger investment not only in computer hardware but also in CBT software development time. There is alot of controversy on what one hour of CBT really takes to develop. Don't forget that it is actually computer programming of some sort in developing CBT. It can take 3 to 4 times as long as instructor led. Delta Airlines for example says they average about 250:1. Again , it all depends on the level of

sophistication. Do you want to program so it can easily be modified in the future. Look at the bright side of CBT: it's there when you need it, it eliminates travel expenses, it's self paced (remember learning styles), you only study what is applicable to you and it does cut the "fat" that is inherent in classroom training, i.e. people don't come back from break on time or you finish a hands on example and must wait for everyone else to complete. Some companies are using CBT as prerequisite material prior to a higher level instructor led course.

Another mode of course delivery that is really coming on strong is multimedia. First of all, multimedia could be considered CBT. Because it is computer based training. However, it's the way that multimedia works that sets it apart from standard CBT. It uses multiple forms of media, hence multimedia, to present information to the learner. It uses sight, sound, voice, full motion video and it is highly interactive and allows you to explore in areas that you want to. It is the most complicated and expensive method of course delivery that has been explored in this paper. However, more and more companies are discovering that the return on investment for this technology may actually be better then their standard modes of training. What comes in to play here is all of the other costs associated with training. For instance, airfare, car rental, per diem and one of the major benefits is that it typically takes half of the time for the learner to go through a course like this as compared to instructor led!

In conclusion, here are five compelling facts that can help make your case for training in your company.

1. Over the last 50 years, investments in job related training have consistently outperformed investments in capital. Quite simply, a trained workforce produces more productivity gains than capital investments.

2. The average 1.4% of payroll that U.S. companies invest in training reaches only about 10% of the workforce. Japanese and European owned companies based in the U.S. spend three to five times more on employee training than American companies do.

3. By the year 2000, 75% of all workers currently employed will need retraining--technological and societal changes are reshaping jobs in every industry.

4. Seventy percent of the people who will be working in the year 2000 are already in the workforce--the only way they will have access to training is if companies like yours supply it.

5. The surplus labor supply of a few years ago has dwindled to almost nothing. Companies won't be able to buy talent, instead they need to create it with training.

source: ASTD - Train America's Workforce
1640 King Street, Box 1443
Alexandria, Va. 22313

PAPER NUMBER 1010

PERFORMANCE APPRAISAL AS A DEVELOPMENT TOOL

GLENN W. SMITH

Haworth Group, Inc.
6117 Blue Circle Drive
Minnetonka, MN 55343-9108
612.935.4454
Fax: 612.935.2403

INTRODUCTION

I've heard managers exclaim, "I'd rather take a beating than go through all my performance reviews. This is quite possibly the way some of your feel. Fortunately, it doesn't have to be that way. The feeling may be changed by focusing on what performance appraisal is all about. Performance appraisal is not only something that we should not dread but it is something that should be helpful to all involved -- the supervisor/manager, the employee, and the organization.

Let's consider briefly as a basis for comparison a hypothetical athletic coach. This coach is in the first season with our favorite sports team. It has become obvious from observations and informal reports from players that the coach only 'coaches' the players at the end of the season. Then, and only then, the coach takes a few minutes to jot down a few marks on a preprinted form that ranges from fair to excellent. Finally, the coach is generous enough to take a few minutes to orally summarize the already abbreviated concepts contained in the form. Well, I don't know how you would reach to such a coach, but my guess is that such a coach is not likely to last long. On the other hand, how closely do this hypothetical coach's actions (or lack thereof) resemble the way you may have been appraised or an appraisal you've given?

I'm not going to propose that we adapt the athletic model to the extent of video taping everyone's actions all day long. But, let's consider a more active role in the daily routines of operations. We'll start out with some basic parameters as to what performance appraisal should be and what it should not be. Appraisal, like athletic coaching should be continuous and not just a once-a-year process. In its most constructive form, performance appraisal goes on all the time. Minimally, it is conducted at regular intervals. The essence of what is commonly called the 'annual review' should be merely a summary and review of the year. There should be no surprises at the end -- in either direction.

COACHING

What is coaching anyway and why is it important as we try to improve productivity, accuracy, efficiency?

My dictionary uses words like teaching, informing, inspiring and mentoring to describe coaching. Do you work with your subordinates in such ways as to suggest alternative ways to do a task or do you just say, "This doesn't work, fix it." Do you explain why you think one way to do a certain task will be faster, better, cheaper or whatever? And do you explain why.

A common explanation as to why this is not done, is, "I don't have time." I saw one of those little plaques that they have in novelty shops that sums of the best response to this excuse that I've seen. It said, "Why can I never find time to do something right first, but I can find time to fix it?" While there may be instances when a deadline is at hand and completion is in sight that it is prudent to delay explanations. But if explanations -- coaching, teaching, informing, mentoring -- are never or rarely given, then we are looking at a symptom of another problem or set of problems.

In work situation, as athletics, coaching or performance appraisal must go on during the game. In real life, we can't expect to be looking over the shoulder of our staff, nor do most of us want someone looking over ours. At the same time, very little gets done in this world that can only be 'evaluated' once a year. There are logical periods of time, stages of completion or benchmarks that suggest a reasonable interval to get together and compare notes.

With new employees, this could be quite frequent such as daily -- or even more often as tasks change -- during the first few days, then weekly and so on. With people moving into a new assignment, more frequent is logical at the beginning. Weekly sessions, called one-on-ones are quite reasonable in many of cases. This could be as simple as reviewing projects in progress and sharing in both directions any problems or issues that have impeded progress or may do so in the week ahead. The two should agree on which of them will contact others to keep each project on track.

WHAT IS BEING APPRAISED?

Before serious effort can begin in performance appraisal, the two parties must reach agreement on what it is that is being appraised or is going to be appraised. The most common form for this is the job description. Think for a moment about your job description. Is it a realistic description of the breadth and depth of what you do during a month or year to form the basis for an in-depth summary and analysis culminating in an appraisal? If you can't comfortably answer "yes" to this, then that becomes your first challenge.

WHO DOES THE APPRAISAL?

The answer to this may not be as obvious as it seems. While we think automatically about reporting relationship, these are not always representative of the way work gets done or where services are provided. This ambiguity tends to increase as organizations flatten with fewer levels and more direct reports for each supervisor and manager. An analyst may spend a fraction of the time with the supervisor that is spent with the client department. Many secretaries serve an entire department. We have data entry people who report to our controller but receive non of their source material from him and non of their production goes to him. How do we overcome this? Obtain feedback from those with whom the person works regularly.

INPUT FROM ASSOCIATES, SUBORDINATES AND PEERS

When we introduced a more comprehensive process last year, input from associates, subordinates and peers was the most frightening aspect all around. Guess what happened on the way to the appraisal of the process. This is generally regarded as the most valuable single ingredient. If handled with tact, (everyone realizes that they may be on the receiving end down the road) the perspective is broadened. Actually, it is multiplied. We typically receive from five to eight associates.

This does several things. Everyone in the company is more aware of how they are or may be perceived. This has improved peoples social skills or their application of them. What people do is just one ingredient. If they treat their co-workers disrespectfully, this may come back as part of the review.

It may reinforce a concern of the supervisor. Conversely, a supervisor may learn that a concern is unfounded. In either case, the process is strengthened. A person will be more likely to accept a suggestion or observation as being a true picture if the supervisor can report that, "Four of your associates made this observation."

For those supervisors whose staff work more directly with other departments, it fills a gaping hole in the supervisor's knowledge base. The customer department feels better about having had an opportunity to be involved in improving working relationships or the level of service they receive. It is always important to reinforce positive behavior and eliminate negative. This is a key to improving our organizations. It is good to let well enough alone or as the saying goes, "If it ain't broke, don't fix it." Another valuable approach, however, is to think in terms of, "If I don't know it's not working, how can I fix it." This describes the situation where, "What we don't know does hurt us." If someone has a habit, or trait of something that inhibits delivery of the best service or detracts from otherwise good service, then withholding this information is a disservice to the individual, the department and the organization as a whole.

SELF EXAMINATION

As suggested earlier, if you either do not have an accurate job description or don't have one at all, then this becomes the first step in self examination. If you know that there is no official record of what you either do or may be expected to do, then it should be clear that any appraisal can go either for you or against you depending on what the reviewer compares to their mental picture of what they think you should be doing. Combine this with the absence of input from associates or peers and you may be in for a surprise.

Self-examination is a key part of the process even after complete and commonly understood job descriptions are in place and up to date. The self-appraisal gives the reviewee an opportunity to think about what s/he is doing and what they want to do more or different in the months and years ahead. It taken seriously, it sets the stage for the review by the reviewer. This preparation is important to opening the dialog and preparing the plan for improvement.

RESOURCES

The supervisor/manager is a resource to the staff under their direction. They should provide direction and suggestions about how to do specific things. They should also be prepared to suggest materials, courses, or other people who will be able to help broaden the perspective of their staff. These resources should provide a means to help overcome weaknesses in previous training and experience and build on existing strengths.

Resources include innumerable courses and materials ranging from simple self-help books to PC programs to vendor courses to college credit. The supervisor can provide encouragement by allowing flexible hours to attend courses, soliciting approval for company-paid training, or simply reminding the employee of the company tuition reimbursement program.

Company executives should also be resources for staff development. When a person has a question about how or why a trade practice is the way it is, this is an excellent opportunity to seek the counsel of more senior people.

LOCAL OR INDIVIDUAL IMPLEMENTATION

Why wait? If your company has a minimal appraisal system or non at all, do you have to wait before you can use these ideas in your department or section? Not at all. Wherever I have worked, organizational rules and procedures have been treated as the minimum standard. In small and large organizations alike, there are always going to be managers whose expectations and standards go beyond the minimum set by official policy.

Suppose you like some of what we are discussing here today. What will happen if you go back next week and supplement the official system with some of these ideas?

- More work to do.
- Anticipate the official timing in order to complete the additional steps in time.

- Transcribe, summarize from new, more extensive review to the official form.
- Do some selling to convince peers and co-workers to provide feedback you can use.
- Prepare yourself for a rewarding experience.
- Prepare to see your people grow faster than you imagined.

SUMMARY:

We have great numbers of tools available to us today. We will find it difficult to use nearly all of them. Most of us here today are finding it difficult to keep up with available software applications. One way to increase our personal potential is to actively help develop the potential of those around us. An active and involved manager will encourage input and feedback from his associates and subordinates. One way to do that is to take advantage of peformance appraisal as a development tool.

ADDITIONAL RESOURCES

"Employee-Initiated Reviews." Editorial. *Inc.* July, 1991: 80

CareerTrack, Inc. *How to Write & Conduct Effective Performance Appraisals.* Boulder, CO, 1992.

Grove, Andrew. *Hight Output Management.* Especially Chapter 13, New York: Vintage, 1985

"More Employees Evaluate the Boss." Editorial. *Fortune.* July 29, 1991

Palmer, Margaret J. *"How to Plan and Conduct Productive Performance Appraisals."* American Management Association, 1991.

Swan, William S. *How to Conduct a Superior Performance Appraisal.* New York: John Wiley & Sons, Inc., 1991.

PERFORMANCE CATEGORIES AND FACTORS

PROJECTS AND CONTINUING RESPONSIBILITIES

A **KEY RESULTS** **Rating___**

- How well are you meeting your job's requirements?
- Did you accomplish the objectives set forth on your performance improvement plan?

B **QUALITY OF WORK** **Rating___**

- Are you meeting the pre-set, qualitative performance standards of your position?
- How accurate, thorough and useful is your work? Are your projects completed with no "loose ends"?
- Are you taking your department/division to higher quality levels?

C **QUANTITY OF WORK** **Rating___**

- How productive are you?
- Are you meeting the pre-set, quantitative performance standards for your position?
- How do you compare to others who have done or are doing your job? Do you consistently meet or beat your deadlines?

D **EFFICIENCY/COST CONTROL** **Rating___**

- Do you look for, find and act to reduce inefficiency in the company?
- How well do you minimize or eliminate unnecessary procedures and time wasters?
- Are you an effective negotiator?
- (Managers: How well have you met your budget objectives?)

E **CRITICAL THINKING/DECISION MAKING** **Rating___**

- Do you use a commonsense approach to handling situations?
- Do you think globally instead of falling prey to tunnel vision?
- Do you make objective, unbiased decisions based on facts?
- Do you keep an open mind?
- Are your decisions timely?

F **INNOVATION/CREATIVITY** **Rating___**

- Do you identify opportunities to improve systems and procedures?
- How well do you respond to unusual or stubborn old problems?
- How many new ideas have you offered your department/division?
- How many were accepted and implemented?

G **JOB AND COMPANY KNOWLEDGE** **Rating___**

- How well do you understand the concepts, techniques, requirements, etc., of your job?
- Do you keep up with trends in your field, as well as in your specific job area?
- Do you have a good working knowledge of this organization?

SELF MANAGEMENT

H **PLANNING/ORGANIZATION** **Rating___**

- Do you prioritize and plan your work effectively?
- How organized are you?
- Are your files, records and key documents in order, and easily accessible by others?
- Are you surprised by problems, or do you anticipate and solve them in advance?
- How good is your attention to detail?

I **RESPONSIVENESS/TIMELINESS** **Rating___**

- How quickly do you turn around documents which require a response?
- Do you respond within four hours to other people's voice mail/screen-mail questions?
- Do you return phone calls promptly?

J **INITIATIVE** **Rating___**

- How often do you experiment to improve current systems?
- Do you take action without having to be asked?
- Do you offer solutions/options when you present problems?

K **ADAPTING TO AND IMPLEMENTING CHANGE** **Rating___**

- Are you open to new ideas and ways to do things?
- Are you able to "thrive on chaos" and, at the same time, minimize it?
- (Managers: Do you introduce and implement change effectively within your department?

L **EMOTIONAL CONTROL/ENERGY LEVEL** **Rating___**

- How well do you handle crises and emotional upset?
- Are your moods generally stable and upbeat, or are you prone to "flying off the handle"?
- Can you sustain a high energy level as required by your job?

M **TRAINING/CONTINUOUS LEARNING** **Rating___**

- Do you strive to learn more?
- Do you make it a point to participate in at least one training event per quarter?
- Are you open to learning different operational areas, even though they may have no immediate relevance to your position?

N **ATTENDANCE/PROMPTNESS** **Rating___**

- What is your attendance record?
- Do you arrive and start on time?
- Are you prompt in reporting tardiness or absence?
- Do you attend meetings and appointments as scheduled?

O **PERSONAL IMAGE/WORK AREA** **Rating___**

- Do you dress and groom appropriately?
- How neat is your work area?

P **ETHICAL BEHAVIOR** **Rating___**

- Are you honest with people?
- Do you avoid favoritism?
- Do you take liberties with expenses and time off?
- Do you model good ethics through your action?

Q **DELEGATION (Managers only)** **Rating___**

- Do you give clear instructions when delegating?
- Do you delegate authority with the task?
- Do you review progress and provide feedback?
- Do you actively look for responsibilities you should delegate?

WORKING RELATIONSHIPS

R **RELATIONSHIP WITH SUPERVISOR** **Rating___**

- How well do you inform your supervisor of you progress and possible problems?
- How cooperative and supportive are you?
- Do you work through, rather than around, your supervisor?

S **RELATIONSHIP WITH COWORKERS** **Rating___**

- Do you get along with your fellow employees?
- How well liked and respected are you?
- How well do you work as a team member?

T **RELATIONSHIP WITH SUBORDINATES (Managers Only)** **Rating**

- Do your people support you?
- Do you know, care about and actively develop each of them?
- Are you a role model for the people you supervise?

U **CONFRONTATION** **Rating___**

- Do you resolve conflicts directly, quickly and completely?
- Are you able to discuss unpleasant issues with courtesy and tact?
- How well do you give and take criticism?

V **CUSTOMER SERVICE** **Rating___**

- Are you known as a service giver to both internal customers (other departments and your colleagues), and external customers (those who buy our services and products)?
- How about suppliers and the community at large?

W **LEADERSHIP SKILLS** **Rating___**

- Do you take charge and initiate actions toward accomplishing the organization's goals?
- To what degree do you "buy into" - and support - the organization's values?
- How well do you coach and develop others?

COMMUNICATION SKILLS

X **VERBAL COMMUNICATION** **Rating___**

- How well do you come across one-on-one or in a meeting?
- Do you get to the point without rambling or over-explaining?
- Do you maintain sufficient (and natural) eye contact?
- Do you make sure your listeners have understood you?

Y **WRITTEN COMMUNICATION** **Rating___**

- Is your writing clear, concise and well organized?
- Do you communicate accurately on paper?
- Is your grammar and usage correct?
- Are your documents proofread carefully?
- Does your writing reflect an appropriate personality?

Z **LISTENING** **Rating___**

- Are you sincerely interested in what other people have to say?
- Do you show it?
- How well do people open up to you?
- When appropriate, do you check to make sure you understood properly?

AA **INFORMING** **Rating___**

- How well do you let people in on decisions or changes?
- Do you inform people on a timely basis?
- Are you an accurate judge of *who* needs to know *what*?

AB **PRESENTATIONS** **Rating___**

- Are you composed and confident in front of a group?
- Are you organized, credible, persuasive and interesting? Are your handouts and visual aids effective?
- Do you get the result you were after?

OVERALL RATING

After careful review of the performance factors above, you have earned the following overall rating for this appraisal period:_____.
Below is the explanation and meaning of your rating:

E - EXCELLENT: Your work was outstanding in nearly all areas. You're the kind of person who will keep the organization growing and successful. We recognize and appreciate your exemplary contributions.

V - VERY GOOD: Your work has been consistently above average in most areas. While you have a few areas to work on, we appreciate your commitment and contribution.

S - SATISFACTORY: Your work generally met the standards in most areas and you are fulfilling the basic requirements of your job. Continue your efforts and we'll work together to help you tap more of your potential.

ID - IMPROVEMENT DESIRED: Your work is not up to the performance standards. While you performed acceptably in some areas, your performance was below standard in too many critical ones. You have 90 days to raise your evaluation to "satisfactory" or above, and a raise will be reconsidered. If you do not raise your evaluation, your employment will be terminated.

IE - IMPROVEMENT ESSENTIAL: Your work is far below the performance standards in the critical aspects of your job. You will be permitted to stay in your current job for a 30-day probation. If, in this time, you raise your evaluation to a "satisfactory" or above, you will be allowed to remain with the organization. If you do not raise your evaluation, your employment will be terminated.

REVIEWER'S SIGNATURE_____DATE_____

YOUR COMMENTS:_____

_____

_____

_____

_____

_____

YOUR SIGNATURE_____DATE_____

***NOTE:** Your signature does not necessarily indicate agreement with the appraisal, only that it has been discussed with you. You are obligated to acknowledge the appraisal if your supervisor has discussed it with you.

PERFORMANCE IMPROVEMENT PLAN

1) **TO IMPROVE YOUR EFFECTIVENESS:**

(Use anytime but plans are required if any performance factor received an "ID" or"IE" rating)

1.

2.

3.

4.

2) **TO PREPARE FOR POSSIBLE ADVANCEMENT:**

(Use if advancement is an option and employee is interested)

1.

2.

3.

4.

3) **TO INCREASE YOUR JOB SATISFACTION:**

(Use if advancement is not an option or employee is not interested)

1.

2.

3.

4.

4) **TO SUPPORT THE DEPARTMENT/DEVISION/ORGANIZATION'S OBJECTIVES/PRIORITIES:**

(Use to link individual's goals to organization's)

1.

2.

3.

4.

The above actions/objectives have been discussed and agreed to:

Your signature_____Date_____

Reviewer's signature_____Date_____

Surviving the Client/Server Revolution with Training and Education (1611)

Morris C. Wallack
Hewlett-Packard Company
Professional Services Division
100 Mayfield Avenue, MS 36LF
Mountain View, California 94043
(415) 691-5828

Summary

This paper addresses the problem of how IS organizations can improve their ability to keep up with new technologies, like open systems and client/server, when the industry is changing at a frenetic pace. For organizations to survive, people need to increase their capacity to handle change. Education and training, although often overlooked, are effective methods for helping people cope with change.

All levels of the organization require education and training. Executives need to become aware of the new technologies to act as effective sponsors of change. Middle-level IS managers need to understand the connections between business and IS strategies to sustain changes within the company and get real work done. Technical professionals, especially application developers and data center personnel, need significant retraining to effectively apply new technologies.

Key requirements for success include close alignment of education and training with the business vision, as well as linking education to the information system (IS) plans that support the business.

IS managers can play a key role as agents of change to move their organizations ahead in adopting new systems.

Why it is Important for Employees to Learn New Skills

> "If the organization is to maximize its efficiency and success, there are a number of specific requirements which must be met. First, the most capable people must be selected for each assignment in the organization. ***Especially in a technical business where the rate of progress is rapid, a continuing program of education must be***

> ***undertaken and maintained.*** Techniques that are good today will be outdated in the future and every person in the organization from top to bottom must continually be looking for new and better ways to work.[1]

So wrote David Packard in 1961. Technology is changing faster today than it was thirty-three years ago. Yet professional education and training is still treated as an expense rather than as an investment; more as a "necessary evil" often avoided than as a critical step towards building an organization that can succeed in the long-term.

Business, Competitive, and Technology Pressures are Creating Rapid Change

Business pressures: process reengineering is redefining the nature of work

Many businesses are radically restructuring to remain competitive and even survive. Reengineering is hot, as evidenced by the runaway success of Michael Hammer and James Champy's book, Reengineering the Corporation.[2] Job definitions are being rewritten to reflect the new processes put in place, and downsized corporations are leaving more work behind for the "lucky few" who survive the cut. As CEOs across America look for ways to replace direct labor with automation, simplify work, and get more done with less, the pressure on and expectations of IS managers continues to mount.

Competitive pressures: time-to market windows continue to shrink

The competitive landscape has gotten ruthless over the last decade. Product life cycles continue to shrink: the average PC hardware product lifecycle is now somewhere between six and 18 months long. To remain competitive, business departments are demanding more from IT departments and want to see results quickly.

Technology: high rates of change occurring simultaneously in many areas

The 1980's were a decade where significant change occurred in the computing marketplace. IBM and Apple introduced personal computers that changed the

landscape of end-user computing. Enterprise computing on the larger scale also changed, as major vendors such as Hewlett-Packard, Digital Equipment, and IBM all made inroads in the UNIX and RISC arenas.

The 1990's promise to be more of the same. If anything, that rate of technological change is increasing. Don Tapscott and Art Caston note eight factors in their book, Paradigm Shift, that are driving significant change in the computing marketplace including:[3]

1. Processing moving from traditional semiconductors to microprocessor based systems
2. System architectures changing from host-based to network-based topologies
3. Vendor proprietary software moving to open software standards
4. Information forms changing from single data to mixed (multimedia) text, voice and images
5. Account control changing to vendor-customer partnerships
6. Software development changing from craft to factory
7. User interfaces changing from alphanumeric to graphical, multiform
8. Applications changing from stand-alone to integrated

Any one of these eight factors would be enough to keep an IS team challenged to meet the needs of their customers. What's frightening to IS managers is that many of these trends are happening simultaneously in their own companies. All these factors are forcing IS managers to rethink how they can keep up.

Client/Server As Technology Enabler of Change: Is That Enough?

Many companies are looking at IT investments in client/server technologies as a way to stay more flexible and keep up with change. In other words, IT departments are embracing client/server as *a technology enabler* of change: using *technology* to *enable* the IT systems to be more responsive and flexible in the face of changing business and competitive demands. This is a well founded strategy, since many of the benefits of client/server architectures come from restructuring

data models, changing where processing power is located, distributing workloads and enabling users to access data more easily. Many companies are successfully making the transition and using new technology as a competitive weapon in their markets. The technologies have been proven to do what they say they can do.

However, many companies who jumped on the client/server bandwagon have run into significant difficulty implementing these new systems. Two independent reports recently identified skill gaps as the critical inhibitor to client/server implementation.[4] One key difference between those succeeding and failing is that the winners have recognized that the people designing, managing and using these systems play a critical role in successful implementation.

Conclusion: it's the people and their capabilities, not the technology, that limit a company's ability to implement new client/server systems.

Education and Training as Process Enablers: Building an Organization's Capacity to Respond Positively to Change

If client/server is the technology enabler for systems to support change, then education and training becomes one of the process enablers for individuals to support change. No other single activity can improve the capacity of people to do work better than education and training.[5] In the future, most organizations will survive or fail on the basis of the collective skills and capacity of their workers. In fact, many recent authors contend that for American manufacturers to compete in global markets, the core competency of the organization will have to become its intelligence.[6]

In simple terms, the goal is to create an organization that continually learns new concepts and skills so it can be ready for anything that comes its way. The idea of learning goes beyond gaining new information. Here's what Peter Senge, author of The Fifth Discipline, says about this difference:

> "Learning or knowledge is different from information. A fundamental misunderstanding that permeates Western society is that learning or knowledge does not need to be related to action. ... In Japan, they say that you learn when you know it in your body - literally. There you do not say "I know it" because you heard it but because you know it is in you. ... There isn't knowledge until it is in your body, not just your head."[7]

Setting the Stage for Effective Education: Linking IS Plans to Business Vision

Faced with a change in technology, assignment or project, most employees will prefer:

Certainty	over	Uncertainty
Knowledge	over	Ignorance
Motivation	over	Fear

The remedy to address most of what's been mentioned could fall under the category of plain old management communication: if we just explain to everyone what's going on, our problem will be solved, right? Probably not. Why? Because the introduction of new operating systems, application development methodologies, system architectures, and a different way of thinking about how applications are developed, used, managed, and supported all point to more than just communication. The client/server models being adopted in the market today [there's another challenge; what's your definition of client/server?] represent significant shifts of computing resources and power from the data center to the desktop. Therefore, a company first needs to define its own purpose and rationale for using new technologies. Then, the way to make the entire change more manageable is to consider what each and every individual needs to know BEFORE the technology is introduced. This will help reduce resistance to change, because employees will have "the big picture" and an understanding behind the decisions that will personally affect them .

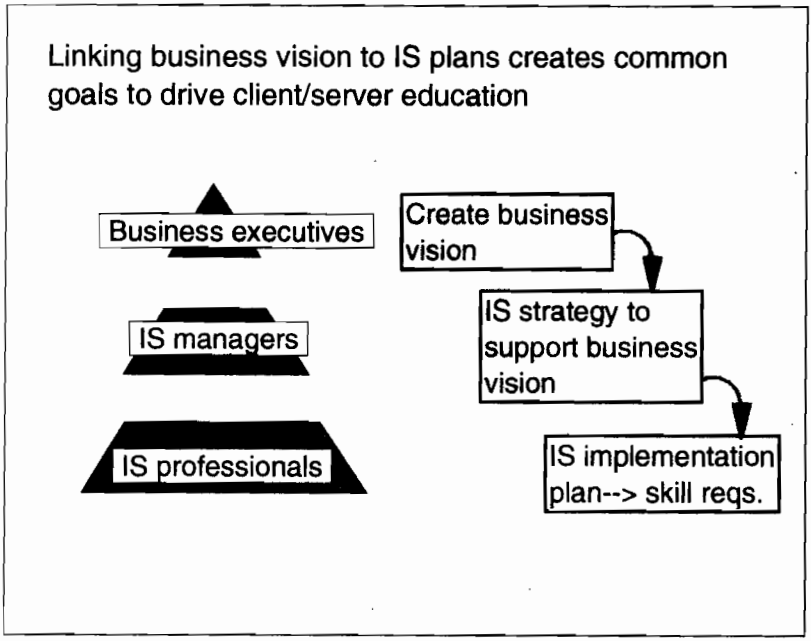

Figure 1

What Do Employees Need to Know to be Successful Going Through a Change?

Different employees need to understand and learn different topics, but generally speaking, here's what they need :

1. Vision: what is the company trying to achieve?
2. Strategy: what is the strategy the business is deploying to achieve the vision?
3. IS strategy: how will information systems be used to support the business?
4. IS implementation: what new systems (or new technologies) will be used and when will they be implemented?

It's important that employees understand what their role is as the organization moves through a change. The figure below outlines the different roles:[8]

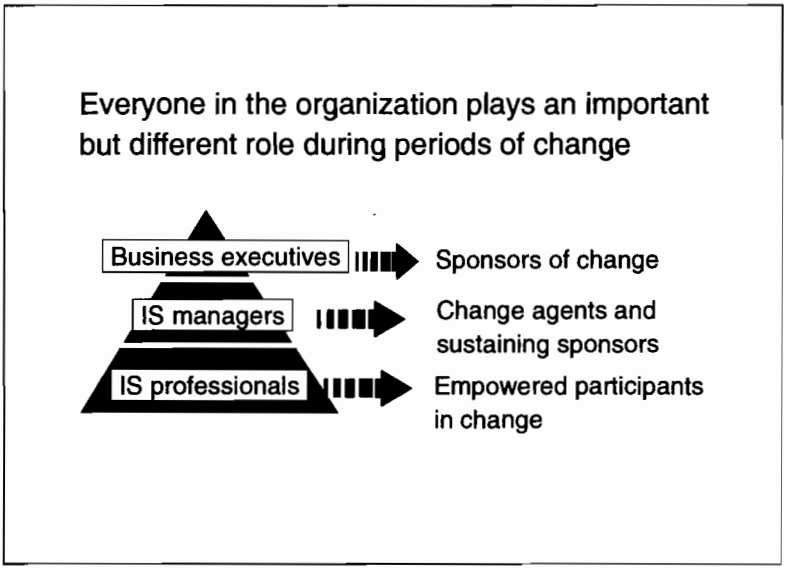

Figure 2

Executives play a key role as sponsors of change. Their public and private support of new technologies, pilot projects, and investments in people validates the change for those around them. Their support sets the tone for whether or not the company is "serious" about moving to client/server systems.

IS managers play a central role as change agents and sustaining sponsors of change. Their primary role is to find ways to translate the overall business goals into tangible activities for their teams, advocating the benefits of the new technologies and identifying where the roadblocks are which inhibit their use. As a sustaining sponsor, another one of IS management's roles is to motivate others to embrace new technologies.

IS professionals ultimately do the work. For them to be empowered participants in moving to client/server systems, they need to understand the reasons why technologies are selected and adopted, as well as how to use them in their day-to-day jobs. If they also understand why the business has embraced client/server systems, they have a broader foundation to make decisions and support the change.

Who needs to learn what topics?

To successfully build the capacity for the organization to handle change, all levels of management and the entire IT staff need to be involved. Role definition and the skills required to do the job drive the selection of training and education topics.

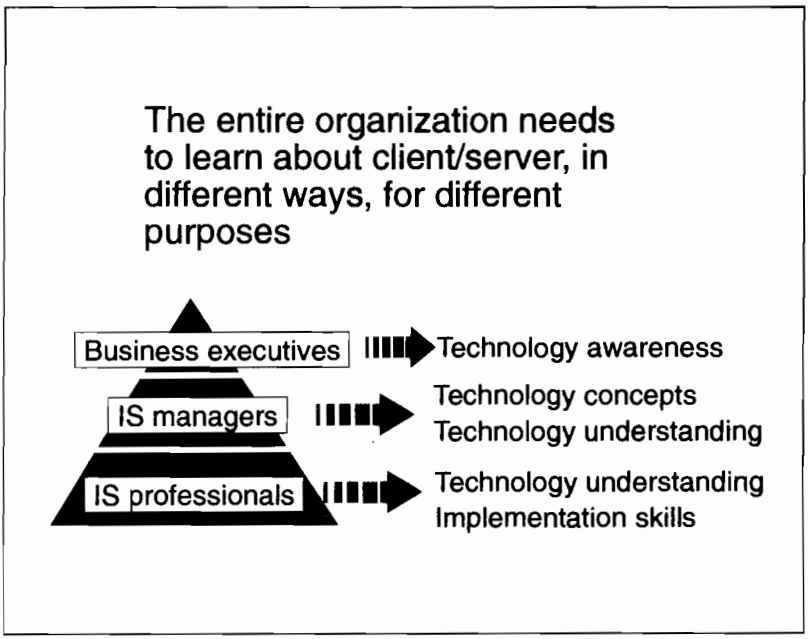

Figure 3

Business executives need to become aware of new technologies for two reasons:

1. So their expectations of what client-server can do are set appropriately (there is no magic bullet here.)

2. So they can be active sponsors of the changes in the business and reinforce to IS employees that their efforts are important, and their new skills will be valued.

IS managers (from the CIO down) need to learn about two facets of the new technologies:

1. Technology concepts: getting a fundamental understanding of the concepts of client/server and open systems to be able to advocate the use of client/server technologies within the business itself. These concepts serve as the bridge between the IS function and the business executives so the management team has a common understanding of technology and its application to the company's needs.

2. Technology understanding: a deeper level of understanding to enable the IS manager to be an active sponsor for change. If you don't understand client/server how will you reinforce to your employees it's important that they learn new technologies? Managers need to understand the technologies well enough to make informed decisions about IS strategy, purchases of new software and hardware tools, and to support the business of the future.

IS professionals need to gain an understanding of the technology and acquire new skills to be able to perform real work in the future. Both levels of learning are critical for successful skill building.

1. Employees need to understand the concepts of client/server before they acquire new programming, systems administration or application support skills. This understanding serves as the framework for why the new skills are needed as well as what role the employee will play in the future IS systems plan.

2. Implementation skill building includes traditional product and technology training commonly offered in the market. No skill building process would be complete without hands-on, project-driven involvement. One best practice is to properly time the employee's training to coincide closely with a work assignment that demands the use of the new skills being learned. This training can also include consulting and mentoring from experts which serves as a skills transfer method to help the organization become self-sufficient.

Peter Senge, author of The Fifth Discipline, The Learning Organization, describes his view of learning and why it's important: (emphasis added) " Learning concerns the enhancement of the capacity to create...Real learning occurs when people are trying to do something that they want to do. *It is always related to doing something.*" [9]

Application developers and systems managers need special attention to survive in the new world.

Two key groups of people in IS organizations that need special attention are application developers and system/network administrators and managers. Client/server and distributed computing systems significantly impact the roles of these two groups for the following reasons:

1. Application developers often have to learn a new language, but more importantly will need to approach software development in different ways. The emergence of process reengineering, coupled with the increased use of object-oriented methodologies, will force them to learn new ways to think about the business. Developers will also need to learn more about the business their software supports, since users will more likely be accessing data on a regular basis, or demand flexibility in code to adapt the programs to changing business needs. A small percentage of developers today are fluent in OO, C++ and GUI methods.[10]

2. Network and system managers will need to rethink their role as data and processing power are more widely distributed within the enterprise. Data recovery, backup and security will challenge network and system managers to develop new ways of managing their systems. Network traffic often increases significantly, and the nature of the traffic becomes more "bursty" with clients transmitting high density, large files to a broad range of network connections. This creates ambiguity and reduces consistency in network behavior that also makes troubleshooting more challenging. Multi-vendor devices hanging on the network or attached to the system also create challenges that may not have existed in the past.

In Summary: a Clear Vision Coupled with Education For the Entire Organization is a Powerful Formula for More Than Just Survival

The formula: get everyone impacted by or working on new client/server systems to understand why the project is being done, why the technology was selected, and what to expect when it's implemented. By educating those in the organization who can influence others (managers and executives) and by sharing the message broadly throughout the organization, the following benefits are realized:

1. Increased Motivation: employees who participate in change processes tend to embrace (rather than resist) change more readily.

2. Better Communication: in today's rapidly changing multi-vendor, client/server marketplace, multiple definitions of key "standards" abound. Building a company-specific vision and common working framework of what client/server is for you improves communication during critical periods of change. This is sometimes referred to as "getting everyone to sing from the same sheet of music".

3. Realistic Expectation-setting: specifically educating both the internal IT staff and the business users on what features and benefits the new architecture will deliver prevents mis-set expectations from sabotaging implementation. All too often the IT department "oversells" capabilities of a new system, and users tend to "overbuy"- falling into the trap that they expect too much, too soon- leading to a "lose-lose" situation that neither the IT team nor the business users can afford in today's competitive business climate.

Real-life Example: How Hewlett-Packard Helped an IS Manager Become an Agent of Change

Case study: International Consumer Products Company

In this example, Hewlett-Packard was asked to work with the VP of IS in an organization that was just beginning to move from mainframes to open systems. Members of the technical staff had doubts about what client/server systems could do for them, and were unclear on why the company had made the decision to go in this direction. The VP of IS then became a sponsor of change within his organization, by acquiring the education needed for his entire organization to get up-to-speed on new concepts, as well as actively participating in the training himself. Figure 5 below summarizes the case.

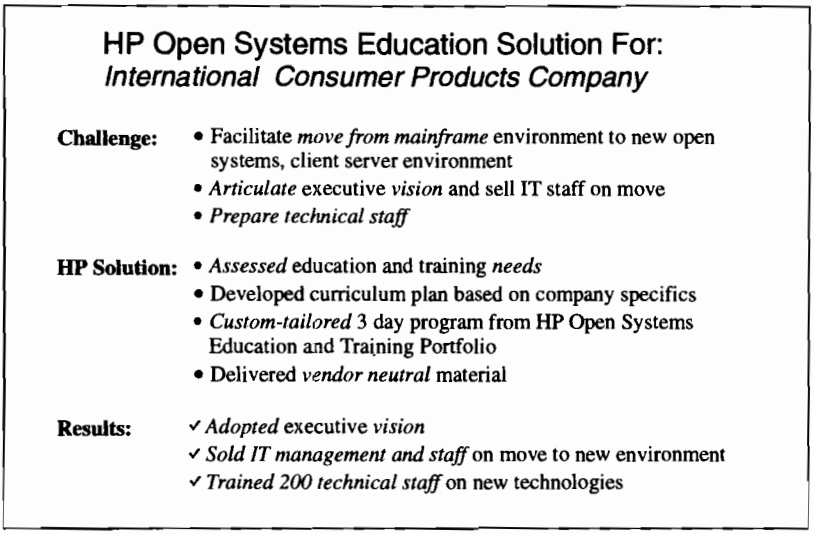

Figure 5

Technology understanding : what it looks like in reality

Figure 6 shows what HP developed in response to the company's needs. Note that the 3-day course is designed to meet their specific needs, including database topics from an outside vendor. By tailoring the education to the specific needs of

the organization, the learning has more relevance, is more closely linked to the needs of the business, and helps the IS management staff reinforce the company's own vision for how new technologies will apply in their environment.

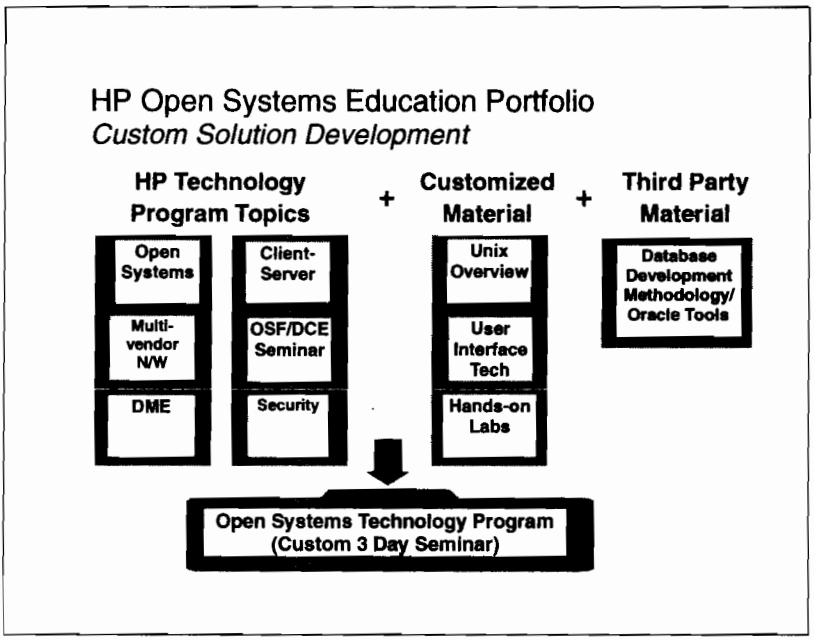

Figure 6

Hewlett-Packard then developed a detailed technical skills training program that followed the 3-day seminar, designed for the IT professionals who were starting work on a pilot implementation of a new system. Again, by delivering the training on technology concepts *before* the skills training process began, the entire organization was better aligned and supportive of new technology coming into the environment. People understood what the company was trying to accomplish (the vision was articulated by the VP of IS in each training session), and after the training, could see how this company's specific implementation of client/server systems would meet the *business needs* for the future.

Putting All the Pieces Together: An Overview

Hewlett-Packard's educational offerings to meet the different needs of IS managers and their organizations is summarized in figure 7. With this array of training, IS managers can now take a proactive stance as change agents within their companies to build the capacity to survive the client/server revolution.

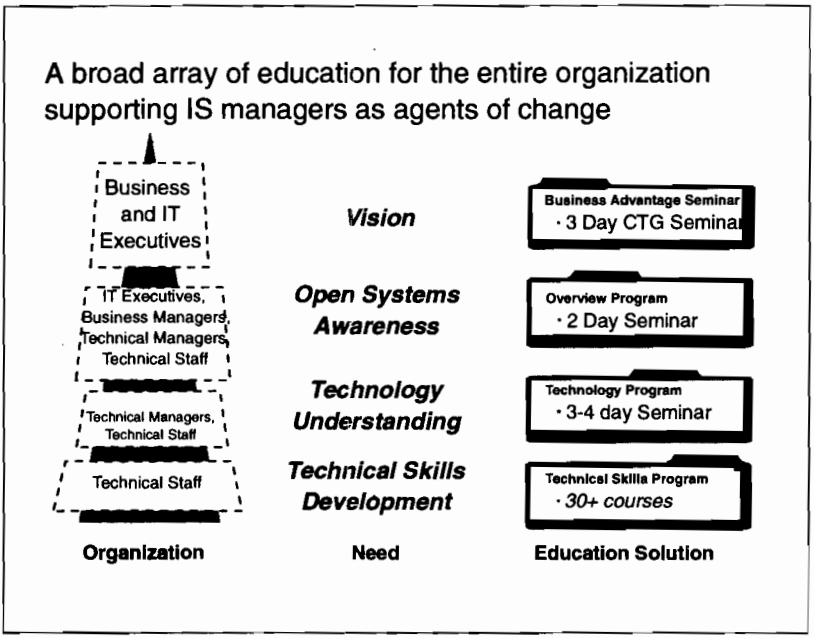

Figure 7

Nine Questions to Ask: What Kind of Learning Organization are You Building to Survive in the '90s and Beyond?

As you ponder how to move your own organization ahead with the use of client/server technologies, the following nine questions may help you assess your approach. Are you reinforcing the concept of a learning organization, one that is building its capacity to create, or are you geared to short term results and perhaps failure?

1. Is your vision for the business future clear? Has it been communicated clearly and repeatedly to all people in your organization?

2. Are IS projects clearly linked to customer needs and business goals?

3. Does the management team understand client/server as well as the development team? Do you have a clear set of expectations for the use of new technologies? Does you have a common language for client/server?

4. Are your best people being asked to act as change agents or just IS professionals?

5. Are masters and experts teaching your teams, or are the "blind leading the blind"?

6. Do people get rewarded for increasing their capability to do good work now AND in the future, or just rewarded for today's results?

7. Is employee training linked to projects or managed separately? Who monitors the training budget?

8. Have you benchmarked your productive capacity to know what skills need enhancing or you need to acquire?

9. When was the last time you attended a training course to learn HOW to do something?

Acknowledgments

I would like to acknowledge the contributions from others within Hewlett-Packard to this paper. Peggy Parskey, David Robins and David Deasy all helped with invaluable feedback and review. Tom Ormseth and Alice Valdez contributed information regarding the case studies, with Ron Gambassi helping in his work tailoring education to meet executive and IS management's needs. Mary Leadholm contributed much of the background research and literature search on the changing needs of IS organizations for training and education.

[1] Packard, David," Hewlett-Packard Corporate Objectives", 1961.

[2] Champy, James and Hammer, Michael. Reengineering the Corporation: A Manifesto for Business Revolution, HarperCollins, New York, 1993.

[3] Caston, Art and Tapscott, Don. Paradigm Shift: the new promise of information technology. McGraw-Hill, 1993, pg. 19.

[4] Shay, P. et.al., "Skills: The Achilles Heel of Client/Server Computing" Gartner Group Logical View Research Note K-140-017, April 6,1992. Gartner Group, Stamford, CT.

Dix, Heidi S., and Woodring, Stuart, "Focus: Closing the Skills Gap", The *Software Strategy Report*, Volume 4, No. 2, May 1993. Forrester Research, Inc., Cambridge, MA.

[5] US productivity growth from 1935 to 1985 was due more to education (55% due to on-the-job training, 26% at school) than capital equipment (19%). Attributed to Curtis Plott, Executive Vice President, ASTD as quoted in "Managing Employee Learning", David G. Robins. Interex Proceedings September, 1993.

[6] See James Brian Quinn's The Intelligent Enterprise, or Peter Senge's The Fifth Discipline.

[7] Senge, Peter, quoted in article, "Creating the Learning Organization", *McKinsey Quarterly,* McKinsey & Co., Inc., New York,1992. Number 1. pg 60.

[8] Adapted from internal training course," Management of Organizational Change" presented at Hewlett-Packard Co. delivered by KPMG /Peat Marwick. Source: ODR, Inc., Atlanta, 1990.

[9] Senge, Peter, quoted in article, "Creating the Learning Organization", *McKinsey Quarterly,* McKinsey & Co., Inc., New York,1992. Number 1. pg 60.

[10] Dix, Heidi S., and Woodring, Stuart, ibid.

Anarchy in Management
A Decision-Maker's Guide

Louis R. Mills, CDP
7415 Saraview Way
Sebastopol, CA 95472-4716

(510) 724-5613
(707) 823-5226
CompuServe: 73631,223

How much anarchy is there in your organization? We're not talking about confusion or hustle and bustle, but about how much "rule" does your management apply. Anarchy is not the same as chaos. Anarchy defines a style of management and is not necessarily a mass of confusion or daily disruption. That's another problem. This is what this paper covers:

- The basics of organization theory.
- Two contrasting models of organization.
- Decision systems within these models.
- Influencing decisions within these systems.
- Some conclusions.

The basics of organization theory rest on:

1. The organization's goals
2. Its technology base
3. Its hierarchy or structure
4. Its environment
5. The setting (strategic or sub-unit)

The two contrasting models of organizational behavior and structure are:

1. The Classical Bureaucracy (Max Weber)
2. The Organized Anarchy (James March, Sr.)

What is a Classical Bureaucracy? It can best be understood using the following terms:

- Its goals are clear and well understood, inside and outside the organization.
- It has an unambiguous technology base.
- Pyramid reporting structure (most likely).
- Participation is stable (fixed by position in the organization).

Good examples of such bureaucracies are: armies, churches, industrial giants, hospitals. Their organization charts look like:

The Classical Bureaucracy has:

- Lots of meetings with formal presentations
- Published 5-year plan
- Salary levels by rank
- Lots of career workers
- Multi-volume personnel manual
- Shrines to the past

In the classic bureaucracy, it takes forever to get anything done. First, you have to define the problem. Then you have to get all the people impacted together to even decide what solutions to look at. Once solutions are outlined, they're written up, reviewed, revised and published. Then vendors are solicited. Starting to sound familiar?

What is an Organized Anarchy? Using the same criteria:

- The goals are not clear or well understood.
- The technology is ambiguous or broad spectrum.
- Flexible reporting structure (or vague one).
- Participation is fluid.

Good examples are: colleges and universities, high-tech companies, and agencies containing elected officials. Their organization chart often has only one or two layers of reporting responsibility:

President
|
C. I. O.
|
You!

The Organized Anarchy has:

- Lots of short, disorganized meetings
- Plans change annually or quarterly
- Salary level by talent or chance
- High turnover, few old-timers
- Few published personnel policies
- Modern art in the lobby

How do decisions get made in these organizations?

- The Rational System (Fayal, Taylor)
- The Garbage Can (March)

Rational decision making includes:

1. Defining the problem (product < goal)
2. Establish criteria for solution adequacy
3. Generate alternate solutions
4. Evaluate solutions against criteria and make the optimum choice.
5. Evaluate the effectiveness of the solution (product > or = goal)

We use rational decision making all the time in data processing. For example, a recent problem at Bio-Rad Laboratories revolved around their sales order and accounting systems, and is a classic example of an Information Services problem:

1. Problem: Terminal response time too slow. 15 to 60 seconds.
2. Requirement: Customer order status inquiries must respond within 10 seconds.
3. Solutions proposed: Upgrade processor, spin off applications, etc.
4. Solution chosen: Select new data communication equipment.
5. Result: Response time improved from 30 second average to 15 seconds.

Now, they didn't reach their goal of 10 seconds, but the user community was so pleased because response time was cut in half, that no one came back and scolded DP for not meeting its goals. Of course, this was a few years ago. Now the users want one to two second response times, 24-hour a day service, etc.

Other examples of rational decision making include methods like Zero Based Budgeting, most project planning disciplines, etc.

The Garbage Can process is interesting because of it's simplicity:

- A decision opportunity exists
- Appropriate resources exist
- A decision is made

The best example of this is the "water cooler" model of decision making, but also includes hallway chats, staff meetings, disciplinary conferences and goal setting exercises. It is any time you meet with someone who can "bless" or sanction your action.

For example, in the Bio-Rad example used earlier, one approach might have been this discussion by the coffee machine:

DP Manager: "Good morning, Ralph."

CFO: "Good morning, John. Say, I hear my people have to work late every Friday because the computer's so slow."

DPM: "Yep. Friday afternoon we get a lot of sales orders and the Shipping department is posting all its shipments for the week. And yet, if I could buy two new multiplexors and a high-speed modem, the problem would go away."

CFO: "It's costing me overtime pay every week. How much would you need?"

DPM: "$12,600 for the multiplexors, another $1,100 for the modem, and maybe another $1,200 to get your accounting terminals to run at the higher baud rate."

CFO: "Baud rate? What's that"

DPM: "Oh, it means how fast the computer paints the screens. Less wait time for your clerks. Yeah, $15,000 tops. That will eliminate this problem."

CFO: "Do it. The sooner the better."

How do you influence decisions within these contrasting models? Under the Rational model, you must:

1. Be in the appropriate position within the structure.
2. Be able to influence an individual who is in the appropriate position within the structure.

In other words, nothing gets done outside the formal channels. You must either have the authority or convince the authority to decide. If you don't have the authority yourself, you must become adept at influencing those at the top. Examples include: presentations to committees or the Board of Directors, or informal elbow twisting at the Executive Dining Room or via a Country Club membership.

The Garbage Can model is very different. It doesn't depend on position in the organization chart hierarchy. There are several ways to effect or influence decisions, any you can use one or more of the following:

- Spend time - become the resident expert
- Persist - keep the need visible
- Exchange status for substance - give the credit away to get what you want
- Facilitate opposition participation - ask them to join in a "study" - they may just go away!
- Overload the system - use a "help desk", use a "complainer"
- Create decision opportunities - have solutions ready for more than one problem
- Manage unobtrusively - just do it.
- Interpret history your way - ignore the past

Captain Grace Hopper, the mother of COBOL, tells us, "it is far easier to apologize than it is to get permission." In many companies, that's still true. Remem-

ber, if you were only trying to help, others will feel bad about censoring your mistakes. But they won't forget them.

Some conclusions to take back to the office:

1. Organizations vary between the two extremes. Your own company may operate differently in different departments or divisions. If you get transferred or get a new boss, the ground rules will change.

2. Your ability to influence decisions may be significantly independent of the organizational setting. Form alliances or partnerships. Don't be afraid to go to others who have power and say things like, "Wouldn't it help if we..." You can get things done.

3. Manage your risk. Make decisions that have good payback to you and your staff. Don't act for action's sake. Do your homework and know why your action is right.

4. Know who the other decision makers and influencers are. Use them. Be where they can "trip over" you.

1812

Mastering Successful Change

Leslie Kennedy
Outside Edge, Inc.
1111 E. John #4
Seattle, Wa. 98102
(206) 324-0841

Enlightened leaders are starting to get it.. "It" is that organizations are, at the most basic level, a collection of individuals complete with all their human components. Along with knowledge, skills, creativity and energy come individual desires and emotions. It is a single integrated package. To get the best of us in service of the corporate goals, leaders must attend to the human element.

Change impacts us at all levels - technical, social, managerial. While change is most often driven by technical or procedural needs, the reaction is almost always related to the social needs of the people impacted. We vary in our capacity to deal with change, but almost all of us have an emotional impact. Statistics have shown that times of change in organizations are also times of low productivity and high turnover. The good news is that we are starting to get smarter about what people need to move from the old to the new.

Organizational change is now constant. New information systems, acquisitions, reorganizations and mergers both upset the entire package and are a source of great opportunity . These tangible changes eventually result in the growth of knowledge, more creative R and D efforts, streamlined decision making, faster time to market, reduction of waste, incredible savings of resources and increased bottom lines. Many individuals have a chance to advance, to learn new skills, to grow through doing new jobs. Through change our civilization moves forward.

But as much as it moves us forward, it is also the source of the greatest stress in organizations. It is well known that with rampant change comes rampant resistance. While leaders are talking about the business advantages and the value to stockholders, a great many employees are caught up in fear and anxiety as they face an unknown future. For those who have experienced major organizational change, the only known is a period of confusion, lack of information, too much to do, frustration, and fear of even more change.

Many leaders have become savvy to the impact of change and are beginning to provide support systems for the members of their organizations. The most common type of support goes to the most obviously affected - those who must leave. Losing a job through change is tough. Fortunately, outplacement firms are doing good work in assisting those impacted in this way.

Very savvy leaders have taken another step forward, They have realized that the survivors of change, those actually responsible for the future of the organization, are also touched. They also suffer loss and confusion as stability, previous roles, valued teammates and familiar processes become memories. The sponsors of change who are responsible for motivating others lack the knowledge and tools for dealing effectively with their staff. Those who are lucky enough to stay are not having fun.

Well, then, what to do? Probably the most successful strategy comes from William Bridges, a well-known author in the area of change and transition. Bridges philosophy starts with the definition of change. Change is situational - the new system, the new team, the new organizational structure, the layoffs. This is an external event which we may or may not control. The second level of change is its impact on the individual which is an internal, psychological experience in relation to the external event. This level is called transition. Change can be designed and planned, and usually conforms to a set schedule. Transition is a human process that must be supported and encouraged and cannot be scheduled. Unless transitions are managed, change will not be successful.

Endings

"Every beginning is a consequence. Every beginning ends something."
Paul Valery, French Poet

Again in Bridges work , there are three phases in the transition cycle. The first phase is endings. For the new to emerge as a reality, the old must end. While our culture in general recognizes this, organizations mostly refuse to allow any recognition of the emotional impact of endings. Bridges says "The single biggest reason organizational changes fail is that no one though about endings or planned to manage the impact on people."

Here is an example. In a large hospital, the nursing administration decided to implement an automated patient charting system where terminals would be installed in the patient rooms and nurses would do their chart updates as they

visited the patients. This involved two major changes for the nurses: they would have to learn the new charting system, and they would not do their charting together at the nurses station as in the past. The nursing administration felt that there was sufficient technical support for the project , but they recognized that the social impact would make or break their success. They realized that these changes involved significant loss for the nurses. They would lose their time together during the day from which they got support, clinical knowledge, and a sense of togetherness. Further discussion revealed even deeper losses. The nurses felt that they would lose their sense of competency because of the newness of the system, and they feared that they would lose the confidence of the patients and their families as they initially struggled with the system right in front of them.

There are several things to do that can help with endings. The first is to recognize what people are losing. As you plan a change, think through what will truly change. This involves not only the primary change, but also requires that you look for "roll-out" change as things progress. Knowing this puts you in position to determine who will need to let go of what. Some of this will be real - loss of a team, manager, office, chance for promotion. Along with these may be subjective losses which are also important to each of us.

The most effective way to deal with the losses is to acknowledge them openly and then to start looking for ways to recreate them in the new environment. In our business culture, we have in the past not wanted to open the door for any emotional reactions. The problem is that ignoring what is there does not make it go away. Acknowledging loss can be extremely helpful in letting people know that you know what is going on and that you may be able to help.

You will also see signs of grief for the losses. Again , many organizations do not want to see obvious signs of grief. The fact is that signs of grief will show up in many ways - cynicism, acting out, lack of patience, and even sabotage. By creating a forum for people to express their experience during this time, you can greatly reduce the adverse effects of the grieving process that will occur as a normal part of change.

One way to provide this forum is to mark the ending in a way that is consistent with your culture. The nursing organization decided to do a ceremonial burning of the old charts at implementation time. They took the time to plan the ceremony and involved many people. It provided the outlet needed for the organization to let go and move on. A key factor in planning ending events is to treat the past with

respect. Losses are only considered to be losses because they had value to individuals who had put their time and energy into a past effort. Make sure that energy and investment is respected as you move forward.

Open Space

The middle phase of the transition cycle I call the open space and is best described by Marilyn Fergeson, an American futurist:

> " It's not so much that we're afraid of change or so in love with the old ways, but it's that place in between that we fear... It's like being between trapezes. It's Linus when his blanket is in the dryer. There's nothing to to hold on to."

The open space then is just that - a time when the old ways are going away but the new ways have yet to be established. Here is where you see the most severe impact of change. Anxiety rises, and people become consumed with their individual fates which do not necessarily match the goals of the organization. Absenteeism may rise as people try to sort out what they need to do. In another hospital where major changes are taking place, a group of nursing administrators reported how this phase played out for them:

" I no longer feel like going the extra mile."

" It's very frustrating to try to get anything done - I can't get people's attention."

" The uncertainty is driving us crazy - people spend a lot of time looking for support from others."

On top of the uncertainty this phase contains, there is usually a great deal more work as people try to learn something new while maintaining the old system, or are assigned work that used to be done by someone just laid off. Because of this, a critical element of change and transition, communication, gets worse instead of better.

It is critical that the open space be attended to , managed carefully both for the welfare of the organization and for the people in it. The first step here is to be aware that it will exist and not treat it as an organizational problem such as bad morale. It is also critical that you are aware that different people will go through this phase in different ways and paces. All too often, those whose losses are not

great and who maintain an exterior enthusiasm about the change are labeled as "cheerleaders" . On the other hand, people who experience a greater resistance and take longer to move through open space might be considered as "unchangeable". Open space is a time for reorientation which is an individual process. Work hard to make room for everyone and avoid splintering due to labels.

Just as open space creates a time for confusion, it also creates a time for creativity. What is true is that the old ways are no longer in effect - the old rules, norms, etc. may no longer apply. The organization unfreezes, freeing people to create new, more effective business processes. It is through careful management that you can cause people to see the opportunity of this time.

One of the simplest tactics in open space is to get people involved. Involvement in the change process not only provides a place for creativity to come forward, it also create a sense of control. Change that is totally out of our control has the greatest impact. Even a small role can be helpful in crossing the open space. Temporary structures are a great way to do this. They can be used to reduce additional changes, monitor the transition process, communicate, track whatever things might be slowing down the implementation.

Another tool in this time is to get people thinking about what they can create., i.e. give them a chance to develop their vision or picture of the future. Most of us see ways that our organization, and our place in it, could be made better. By focusing on the opportunity that change gives us to realize some of these improvements, we can begin to work with the change and not against it. Harrison Owen, an author on organization transformation, says that "Vision is the first step in crossing the Open Space."

As much as possible during the open space, focus on providing your people with opportunities to be creative, and put in place ways to capture creativity when it appears. Creativity will be a great bridge from resistance to opportunity.

New Beginnings
"A rock pile ceases to be a rock pile the moment a single man
contemplates it, bearing within him the image
of a cathedral."
Antoine de Saint-Exupery,
French novelist

Finally, you will reach new beginnings, a time where new values and new identities take over. Just as change is different from transition, beginnings are different from starts. The start can be said to be the scheduled date, the day when the new system is cutover, the new department begins to take action. But beginnings are a psychological process and may not be in synch with the starts. Beginnings are much more a function of the heart. Beginnings happen when people are truly "on board" with the new.

Clearly, this is the state that we want to create and there are four steps that will help to nurture the beginnings. First, the purpose of the change must be well understood in terms of personal impact., i.e. does the purpose really make sense? To accomplish this, make sure that the real purpose has been communicated in a way that people can receive it. Don't put out a vague company line ("competitive advantage") that will be recognized as just that. It must also be possible for each individual to relate to the purpose. As in the previous example, the nurses needed to have an understanding of exactly how an automated charting system would help them provide excellent patient care. Once they understood how it would minimize the current confusion with medication and doctors' instructions, they were much more eager to sign on. William Bridges states:

> " Successful new beginnings are based on a clear and appropriate purpose. Without one, there may be lots of starts but no beginnings. Without a beginning the transition is incomplete. And without transition, the change changes nothing."

The next step in nurturing beginnings is to provide the picture of the future, i.e. what will it truly look like when we accomplish this change. The more detailed the picture, the greater the step toward the true beginning. For the nurses, having the nursing administrator for the hospital talk about how the new system would route information, how it would eliminate confusing phone messages, and how it would actually benefit the patient was transformational. It gave the nurses something to displace the old picture, accelerating their movement forward.

Planning is the next piece. You probably already have an implementation plan for whatever change you are making. This plan needs to include the plan for the transition as well. Consider the phases of the transition cycle and how you will manage them so that the transition works for you. For example, you might consider how will endings be accomplished. What kind of ending event makes sense - should it be individual, departmental, or organizational? What would

demonstrate respect for the past? Similarly, you night look at ways to manage the open space. How could you involve people? What are ways to encourage and capture creativity in the organization?

Finally, maximize the parts that people play in the overall change. To what extent can you get people involved both in the outcome and in the management of the transition? As before, involvement creates a sense of control that is greatly lacking during times of change, and it greatly increases the sense of ownership in the outcome.

Now that there exists a model for transitions, the next step is to ensure that education is provided that spreads this knowledge throughout the organization. By educating your workforce, you create a common language and a common way to deal with difficulties during change. Currently , there are many different avenues of providing education, from simple booklets on how to survive change to experience-based seminars. All will have value. Here is a way to gauge what will be useful to your organization. This study was done at the University of Oregon and shows how different training techniques impact our ability to apply information:

Method	Knowledge	Skills	Application
Theory	8	1	0
Modeling	10	3	0
Practice	10	7	2
Adaptation	10	9	5
Coaching	10	10	8

While knowledge is important, what we can see from this chart is that knowledge alone won't get application results. Don't rely on booklets to do more that provide some basic information about change. What will be applied is knowledge that is adapted to a specific situation and is combined with some form of coaching. Most successful organizations, including the two hospitals sited in this paper, have chosen to allow people time in workshops where several things happen. They are introduced to the transition cycle and asked to focus on their own experience of a change that is currently happening in their organization. This use of their experience makes knowledge of the transition personal and usable in the workplace. However, the most important result of this time to participants has been the opportunity to find out that everyone is impacted by change even though

behaviors may vary from individual to individual. It is only through being able to voice their reactions together that this understanding comes.

Successful change is necessary to our businesses. With all the external and internal pressures that we currently face, we can no longer afford implementations that cause a long-term slowdown in productivity and high turnover. Managing the transition as well as the tasks of change will greatly accelerate the return to productivity and give us successful implementations.

Paper Number: 2001

Title: Client-Server and Rightsizing, Open Systems and Challenges

Author: Marlene Nesson

Company: Information Builders, Inc.
1250 Broadway
New York, New York, 10001
(212) 736-4433

Information Builders

General Overview

Client/server and rightsizing, along with open systems are constantly touted solutions as well as buzzwords in the industry today. But really what does this all mean? Everybody seems to think it is a good thing- even if they can't agree on a definition.

Client/server seems to be a new kind of parlor game for the industry. It has been described as a style of computing, a way of saving money, a collection of technologies, an architectural platform, an application development method, a systems integration solution, a re-engineering tool, a way to implement standards and even a paradigm shift. Client/server has many, many different flavors. It won't look the same everywhere you go as there is no textbook that defines client/server and even if there was a textbook definition the English language is open to different interpretations.

Rightsizing, originally called Downsizing (since the term inferred the moving of applications from a larger system down to a smaller system), is basically looked at moving applications and/or data from one system to another. It may mean changing the application or keeping the application the same. It may mean utilizing existent technologies or implementing existing technologies. Once again, since there is no textbook definition rightsizing will look different wherever you go.

Open systems has been referred as interoperability across heterogeneous systems, networks, management of heterogeneous database, integration of open and proprietary networks, and integration of standards. Open systems offers users both platform and vendor independence. Users have the flexibility to choose the hardware, the software, the network and operating environment according to what makes the best business sense.

A simple definition for all 3 of these technologies may be freedom of choice. Freedom to choose the hardware platform, databases, tools and user environments that are right for your business.

This paper will approach the subject of client/server, rightsizing and open systems through real business challenges and the solutions implemented that were the right solutions for each individual's business.

We'll begin by presenting an overview of client/server, rightsizing, and open

systems followed by factors companies considered in adopting and/or integrating these new technologies. Two in-depth studies of companies who have implemented client/server, rightsizing and/or open systems solutions will then be discussed.

Overview of Concepts

Client/Server

Client/server is an open architecture that distributes processing between a back-end processor called a "server" and a front-end processor called a "client". The distribution of the workload between two or more processors would result in greater efficiency and lower network traffic.

The key components in a client/server environment are an

- Operating system(s)
- Database management system
- Development tools
- Graphical user interface and/or a character interface
- Client/server applications
- Connectivity protocols
- LAN (local area network)/WAN (wide area network).

Client/server computing grew from the network computing of the 1980s. Networking gave users the ability to share files and resources. Client/server computing, however, takes the networking technology a step further. It allows applications processing to be shared and allows the ability to operate on multiple platforms, systems and networks.

Client/server computing is experiencing overall growth. In fact the Meta Group, believes that "...by 1995, virtually all new applications will be client/server applications". The most dynamic area of expansion in client/server computing is software development: database management systems, development tools, and client/server applications. For example, Sybase and Informix, a database management system software vendor, more than doubled unit shipments between 1990 and 1992.

Client/server computing produces a number of benefits and advantages. *The most common benefit is cost savings realized by capping mainframe investment, increasing productivity, reduced training time, and buying less expensive workstations, desktop platforms and software.*

Rightsizing

Rightsizing encompasses the act of changing platforms, moving and/or changing applications and utilizing existent and/or implementing new technologies. Changing platforms include: downsizing - moving systems to smaller, less expensive platforms; as well as upsizing - increasing processing power when the current systems reaches its capacity. Rightsizing also involves a wide spectrum of activities, the three most common being:

Rehosting:	Moving applications, unchanged, to a new platform
Rearchitecting:	Adding new technologies such as client/server graphical user interfaces or relational databases, to the existing applications
Reeingineering:	Making significant changes to the underlying business processes

Originally, rightsizing, was called "downsizing", or more specifically, "system downsizing", and referred to the processing of moving information systems off the proprietary host, usually a mainframe or minicomputer, to a lower-cost (hence, smaller) system. The system might be a PC or UNIX workstation or server, which was clearly smaller than the host, or even a multiprocessing UNIX machine, that, quite possibly, rivaled or exceeded the proprietary mainframe in size and processing power. The smallness implied in downsizing really referred to the cost of the new system, which came with a smaller purchase price and a lower cost of ownership over time, regardless of its actual size.

Rightsizing, as a word or term, is better suited to describe the activity than downsizing as it refers to the process of moving information systems to the most appropriate system for a particular application whether it be larger or smaller. Rightsizing matches the needs of the application and the organization with the strengths of the various systems and often client/server is integrated in a rightsizing scenario. Thus, the "right" in rightsizing refers to picking the best system for the application based on the system's strengths, not its size.

Rightsizing scenarios typically move toward the same goal: distributed, networked systems that interoperate transparently as one system - freedom of choice.

Open Systems

An open systems environment supports interoperability across heterogeneous systems and networks, management of heterogeneous database and the integration of open and proprietary networks and the ability to run with industry

standards. Open systems offers users both platform and vendor independence. Users have the flexibility to choose the hardware, the software, the network and operating environment according to what makes the best business sense. Therefore, a major advantage that open systems have over closed architecture host-based systems is multiplatform and multivendor support.

The problem experienced by end-users with the closed nature of proprietary, host-based systems was the inflexibility of the closed environment. It was difficult to share information. Essentially locked into a platform or vendor, there was little or no flexibility to address users information needs as they grew or changed. End-users had to use the same hardware and software vendors even if that platform or software was not the best option to solve their business needs. Information might be available solely on one platform, yet difficult to access because the users of the information, scattered throughout the organization and network, might be on different platforms and software systems. Open systems solution offers the users the ability to get the information they need, no matter where it resides, and deploy applications on the hardware that makes the best sense - freedom of choice!

X/Open Definition of Open Systems :

> "The three goals of an open systems environment are interoperablitiy, portability and scalability of applications. Interoperability is the ability to share information in a heterogeneous environment of dissimilar workstations, servers, networks, and vendor applications. Portability allows for the migration or "porting" of the applications to other platforms. Scalability allows movement of an application from a PC or workstation to proprietary host mainframes and minicomputers."

X/Open is a consortium dedicated to developing an open systems standard, and is widely accepted in the industry. Another group, with similar goals for software application development, is the Open Software Foundation (OSF) and a newly formed group called Common Open Software Environment (COSE).

Note: COSE was announced at Uniforum 1993 by Hewlett-Packard Company, IBM Corporation, The Santa Cruz Operation Incorporated, Sunsoft Incorporated, Univel, and UNIX System Laboratories Incorporated. Each company said it would adopt the most popular de facto standards in the UNIX business: the Motif user interface, Sun's Open Network Computing (ONC) networking, and the X Consortium's X Window system. Although not positioned as "another" standards committee or organization, the six companies stated they would announce a set of specifications that would result in a common set of programming and user

interfaces that should bridge many of the gaps between their different versions of UNIX. COSE is thought to the best news for standardizing UNIX.

It should be noted that in the marketplace, open systems have been more or less synonymous with UNIX. Although as clearly described above, open systems embodies much more than UNIX and more appropriately an interoperable environment that can include UNIX as well as a host of heterogeneous systems, networks and databases - once again the freedom of choice.

Many believe that open systems provide the ability for applications to work with or be integrated with industry standards, whether they be protocols, application programming interfaces (APIs), structured query language (SQL), object technology, and so on. There is a standards committee or organization for almost every flavor of data processing technology in today's marketplace. The goal here is not to review all the "industry" standards, since standards, as such, are still emerging. However, it is worth pointing out that open systems does seem to inherently imply the ability for applications to be incorporated into and/or interoperate with standards as they evolve in the computing industry.

Given this brief overview of client/server, rightsizing and open systems let's look at some of the factors companies are considering in the adoption of these new technologies.

Factors - Adoption of New Technologies

Today's workplace whether in retail, manufacturing or service is virtually unrecognizable from the one of twenty years ago. Change in the workplace, however, is now moving at an exponential rather than a linear rate. As a result, the workplace five years from now will be so different from that of today that it will likely have a brand new name. For example, the "home office" was, until recently, a concept that engendered ridicule. Today, for many professionals, the home office is the preferred situation.

As companies strive to remain viable in an increasingly competitive market place, they are examining ways to become meaner and leaner. The main thrust is to produce more with less. For a company's Information Systems (IS) people, this means tightening up the technological processes of data management.

To do so, MIS is bringing the large mainframe applications down a few pegs to a more manageable and cost-effective size. It also means simplifying the complexities of data analysis which historically had been handled solely by

technically oriented IS experts and moving these procedures down to the workers' level where the need for stored information is most crucial.

Client/Server, Rightsizing and Open Systems- Related and Inter-related Technologies

So what are companies doing to cut costs as well as enable the workers to get to this critical information? Companies are rightsizing moving from mainframes down to less expensive, user friendly personal computers or workstations. However, the mainframe is not being eliminated but rather is being used in a more innovative and ambitious way in conjunction with other less costly computer equipment. One method of using the information and technology already in existence is client/server computing. Companies who have chosen to rightsize see client/server computing as one way to use their systems to stay competitive and profitable. Open systems, that is interoperability across heterogeneous systems and networks, management of heterogeneous database and the integration of open and proprietary networks, is a requirement to implement a rightsized, client/server environment.

The companies that are profiled in the next several pages, adopted a rightsizing and client/server environment and evaluated a number of factors during their decision making process (see Figure 1). In almost all cases, these customers integrated a rightsizing initiative with a client/server computing environment. This meant keeping a mainframe computer, moving certain applications and/or data from the host, and setting up a client/server environment to take advantage of lower-cost client environments. The common reoccurring goals of all these customers were to *decrease costs, maintain investment in both hardware and software as well as improve workers ability to get to information quickly and efficiently.*

Factors Companies Consider in Adoption
Rightsizing and/or Client/Server Technologies

- Improved individual productivity
- Maintain investment
- Less expensive hardware
- Minimize database management
- Shorten development time
- Faster turnaround
- Reduce personnel costs
- Faster, better market response
- Create more competitive products
- Pushing decisions to lower levels, employee empowerment

Figure 1

User Profiles

The profiles that follow are in-depth studies of 2 companies who have implemented rightsizing, client/server solutions. The following outline will be utilized for each profile described:

- Goal and objectives
- History
- Implementation: Hardware, software, people resource, and costs
- Challenges Uncovered
- Results and Comments

It should be noted that these company profiles are based on experiences by companies who use Information Builders, Inc. software products. Clearly, companies who have invested in other software products that are portable, flexible and comprehensive in functionality could be substituted in the examples that follow.

Profile I - Municipal Court

A municipal court from within California made a decision to rightsize their environment.

Their goals were to:
1) Save mainframe upgrade costs and

2) Implement an open systems environment

One of the primary objectives in this rightsizing effort was to maintain investment in the 4GL applications they had been running on their mainframe computer since 1988 so as to avoid the cost of rewriting systems and the retraining of users. Another objective was to maintain the same functionality on current as well as on all potential platforms they would integrate into their enterprise in the future. Still another objective was to implement a solution that would maintain the same functionality on the open system platform, UNIX, as well as other potential platforms that could be implemented into their enterprise in the future (see Figure 2).

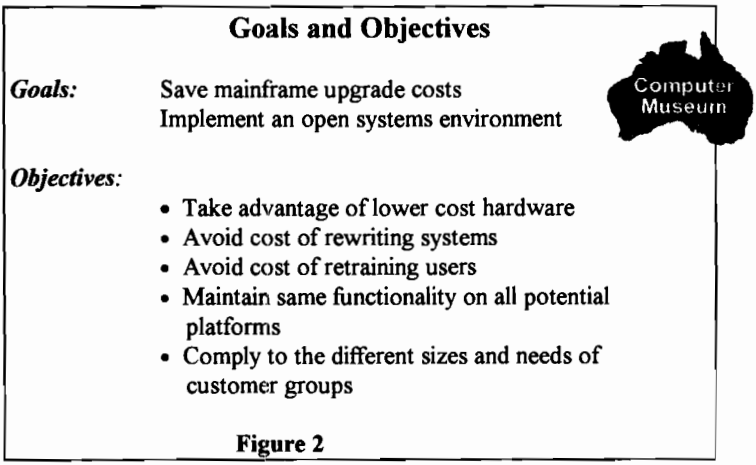

Figure 2

History

The court was running 2 systems: 1) A Civil Automation System (CIVAS) (deployed 1988) and 2) A Small Claims Online Tracking System (SCOT) (deployed in 1991). Both of these systems were written in a 4GL called FOCUS from Information Builders, Inc. and were running on an IBM mainframe 3090 computer.

The users on the systems were approaching 300 and the mainframe capacity was being stretched to its limits. In order to maintain the users capability of real-time access, the court was required to upgrade their mainframe - a very costly endeavor. At the same, the court was beginning to investigate the feasibility of implementing an open systems and client/server environment. Since the decision to upgrade the mainframe would be politically a very unpopular decision (high level

management said Open Systems and/or client/server was the strategic direction - save lots of money), the MIS manager decided that an Open Systems environment, in this case UNIX was equated with Open Systems, would be implemented.

The court wanted to maintain flexibility as well as the investment in both software products and product training. Since the FOCUS code was portable across platforms (in this case from the mainframe to UNIX environment), the court was confident they would the continued flexibility in choosing the platform of choice for deployment of applications as well as maintain investment in software knowledge and training.

The FOCUS product works with its own database (which was the database the court used) and was architected to work with over 50 databases, both relational and non-relational, across 35 platforms. The primary reasons the court chose to use FOCUS several years ago was: FOCUS code portability across all major platforms, comprehensive decision support and maintenance capabilities, and end-user tools (users can be productive without the requirement to know the FOCUS language):

Reasons for Choosing FOCUS
Code Portability, Decision Support, Maintenance

Portable code across 35 platforms
Access to over 50 heterogeneous data structures
Report Writer
Financial Reporting
Output Browser
Graphics and Statistics
User-Friendly Tools for reporting and maintenance
Universal Join (join heterogeneous data structures)

In addition, IBI provides a client/server (middleware) product that is supported across 35 platforms, which would enable the customer to implement a client/server environment as desired.

Implementation
The court set a timeframe of 6 months for the rightsizing project to be completed. This entailed 1) choosing and installing a UNIX machine 2) porting and testing the FOCUS applications and data on the UNIX machine and 3) deploying the new system. A future, secondary, goal would be the implementation of a two tier or

three tier client/server environment: PCs --->UX and/or PCs ---> UX ----> IBM mainframe computer.

Current Configuration
The court was running FOCUS applications with the FOCUS database on an IBM 3090 model 400J running the VM/CMS operating system. A Local Area Network (LAN) running Novell Netware 3.0, had personal computers, IBM 486, running PC/FOCUS applications on the DOS and OS2 operating systems.

Configuration

Original	*New*
Hardware	Hardware
IBM 3090 Model 400J	HP 9000 Model G50
IBM 327X terminals	IBM 327X terminals
PCs on LAN	PCs on LAN
	IBM 3090 Model 400J
Software	Software
FOCUS with FOCUS database	FOCUS with FOCUS db
	Terminal Emulator for IBM terminals (ANET)

The FOCUS code on the IBM mainframe consisted of approximately 240 procedures with upwards of 4000 lines of code. There were 46 databases that consisted of approximately 3 gigabytes of data. The FOCUS procedures included online maintenance, reports as well as batch reports that were overnight.

New Configuration
After analysis of several UNIX computers, the court decided to purchase an HP-UX, G50 as their UNIX, Open Systems platform. The PC/FOCUS DOS and OS/2 applications were to be maintained as is; no changes required. The IBM terminals would be continued to be utilized and hooked into the HP-UX computer.

In order to meet the timeframes for implementation, the court decided to solicit the assistance of the IBI consulting organization. By incorporating consulting services

into their rightsizing project, the court was able to free up their own internal resource for other projects, eliminate the pressure and requirement to train the developers on the UNIX operating system immediately, and meet the goal of a 6th month implementation schedule.

Note: The court had a few very experienced FOCUS programmers. One of the programmers were dedicated to managing the rightsizing project. Utilizing outside resources enabled the court to minimize their internal resource on the rightsizing project and continue to provide timely support and service to their users.

The time from the inception to rightsize through competition was about 6 months. The major elements included:

1) Evaluate, choose and install UNIX (approximately 2 months)
2) Port software applications (approximately 3 months)
3) Deploy applications (approximately 1 month)

Challenges Uncovered

The porting of the FOCUS code from the mainframe to the HP-UX environment was relatively seamless as the code is portable from platform to platform. There was an expected learning curve involved with acquiring the required knowledge of the (new) UNIX operating system. This customer choose to train the current IBM knowledgeable system staff on UNIX rather than hiring an experienced UNIX systems person.

One challenge uncovered and that was maintaining the IBM terminals as the method to link into the UNIX computer. The court had the challenge of seeking out software that would enable these IBM 327X terminals the ability to emulate the UNIX environment (they did not want to purchase UNIX terminals or PCs for the users). The court chose a terminal emulator software solution called Haltek that was able to accommodate this requirement.

Note: There are several companies who provide similar UNIX terminal emulator for IBM 327X terminals. Included for your reference are a few suppliers of terminal emulator software:

Brixton Systems, Inc.	(617) 661-6262
Haltek	(410) 247-0886
NCR	(612) 638-7872
OpenConnect Systems	(214)888-0435
Systems Strategies Inc.	(212) 279-8400

Some of these solutions are software only solutions whereas others are a combination of software and hardware solutions.

Results and Comments
One particularly important element that proved to be instrumental to their success was the dedication of an internal resource as project leader. This project leader managed the scope, activities, progress and expectations of the project.

The court may decide to implement a client/server environment in the future although no timeframe has yet been determined. With the IBI products, the court is confident they can evolve into a client/server environment and maintain their software and people training investments into the future.

Today, there are approximately 300 users accessing the applications running on the HP-UX computer. They were able to meet all their goals and objectives and the court was pleased in their decision to rightsize. A realistic expectation of deliverables (6 month time period), internal expertise, and external consulting services enabled this rightsizing project to be completed successfully.

Profile 2 - Financial Service

This company provides financial services for more than 2 million investors with 66 billion dollars in assets. They supply a broad set of low-cost investment services that enable self-directed investors to customize investment programs to meet their financial goals. This company pioneered fast confirmation of market orders and provided an order entry service; therefore, delivering up-to-date, timely information to users were of primary importance to the companies continued success.

This company invested considerable capital in building its information technology infrastructure. The infrastructure not only includes computer hardware, software, and data network, but also a sizable professional systems staff. The applications running on that infrastructure touch virtually all phases of this financial systems operations. This company is heavily dependent upon information technology.

Their current information technology infrastructure was almost exclusively based upon large-scale mainframe computers. Mainframes were very expensive per unit of throughput. They recognized for some large workloads, mainframe are the only possible solution. However there were many workloads running on the mainframe that they believed could be processed on smaller and less expensive machines.

Use of these computer resources constitutes expense to the firm. It was not "funny money" as beginning in the 4th quarter of 1991, cost centers had been charged for their usage of various system resources. Because of Chargeback, users had a keen interest in any project that can potentially improve the cost effectiveness of their information processing.

This financial institution decided to implement a client/server, distributed system environment. The primary goals were to:

1) Reduce Costs
2) Rightsize and implement client/server architecture

The primary objectives in this rightsizing-client/server effort was to enable end-users continued productivity with the tools they were accustomed to using, provide users access to data regardless of where the data resides in the enterprise, rightsize current applications to a lower-cost, open systems platform, implement a bi-directional client/server architecture (UNIX <---> IBM Mainframe) (see Figure 3).

History

This financial company had been running with FOCUS applications for a number of years on an IBM computer. The number of users that were trained and using the applications were approaching 500. FOCUS was primarily used by end-users to create ad-hoc requests as well as individual applications to access, summarize and manipulate the data. Since FOCUS is an easy tool to learn and use, additional applications were being developed and the mainframe computer began to approach its upward capacity limit. In order for users to maintain fast response to their

Goals and Objectives

Goals: Reduce costs
Rightsize and Implement client/server architecture

Objectives:

- Take advantage of lower cost hardware
- Avoid cost of rewriting systems
- Avoid cost of retraining users
- Provide access to current information (regardless of where the data resides in the enterprise)
- Implement bi-directional client/server architecture

Figure 3

application requirements an upgrade of the mainframe computer was required. Since the upgrade cost to the mainframe was an expensive proposition and they continued to experience a steady increase in chargebacks, the company decided to rightsize to a lower-cost, open systems, environment. This lower cost solution would eliminate the requirement to upgrade the mainframe as well as reduce the cost centers "chargeback" for mainframe usage.

An open systems UNIX environment was decided upon. A majority of the mainframe applications would be ported to the UNIX, however, operational data as well some applications would continue to reside on the mainframe computer. Given this scenario, users would also require continued access to the mainframe. Therefore, a client/server environment between the UNIX and the mainframe computer was essential. This company added one additional twist to their scenario, they wanted users who required major access to the mainframe to use the mainframe as a client, but also have access to data residing on the UNIX machine. Or in other words, they were looking to implement a bi-directional client/server environment between the UNIX and the IBM computer (HP-UX <---> IBM).

Since over 500 users were productive with the FOCUS software, it was clear that a critical element of success for this project was portability of FOCUS code across platforms, that is, all the targeted rightsized applications could be ported to the UNIX computer. Also, since Information Builders offered a client/server, middleware product called Enterprise Data Access (EDA)/SQL that provides bi-directional client/server architecture, the company would be able to maintain their

investment in FOCUS and be able to implement the desired client/server environment as well.

One additional element that was important to this company's future planning, was adhering to, or integrating into, the Open Software Foundation's (OSF), Distributed Computing Environment (DCE) standards.

Implementation

The company set aside approximately 18 months for the rightsizing client/server project. This 18 months did not include the time spent defining the requirements (this took about 3 months) but rather the time spent on actually from beginning of implementation to competition.

They spent considerable time defining and refining requirements. Once they decided upon a UNIX platform - they choose HP-UX, they required the vendors, in this case, Information Builders, and Hewlett-Packard, to understand and assimilate their company's strategy. As mentioned above, DCE was a strategic direction for the company. Although DCE integration was not part of their initial plans, they were confident that their partners, Hewlett-Packard (ultimately chose HP-UX hardware) and Information Builders - (both members of OSF) would help the company achieve DCE compliance in the future.

They designed an elaborate project management plan for implementing a successful rightsized client/server architecture which included specific line items, or deliverables with specific dates and required timeframe for every milestone.

They assembled a project team consisting of HP, IBI, and specific individuals from the company. The company participants included: end-users (from Decision Support Group), UNIX system administrator (newly created position) and telecommunications (responsible for designing and implementing network solutions). This project team met weekly for the first 15 weeks, followed, then with bi-weekly meetings thereafter.

Two components of the project were considered to be new key technologies for the company: 1) UNIX and 2) TCP/IP. These "unknowns" were going to be watched carefully to ensure the company was comfortable with both the functionality and reliability of these new technologies.

In order to decide upon feasibility of the entire project, the first milestone was a port of the simplest FOCUS application and database to the HP-UX computer. The success or lack-there-of would indicate whether or not the project would proceed.

There were a number of milestones identified for this project (see Figure 4). The first milestone took approximately 3 months. This included choosing and installing a UNIX system, installing FOCUS and EDA/SQL on the UNIX machine, as well as porting the first FOCUS application with associated data. The first milestone was successful endeavor so the company choose to proceed with the project.

Once the decision to proceed with the project occurred, the next step was to port all the desired applications to the HP-UX machine. Once this was completed, the client/server element was implemented, that is, networking the IBM and HP box together, implementing the client EDA/SQL API calls in the UNIX FOCUS procedures to obtain the required mainframe data, and testing these client/server applications. Finally, the last component was implemented, that is, integrating the client EDA/SQL API calls in the mainframe FOCUS applications.

Milestones: Rightsizing - Client/Server

1. Choose UNIX hardware
2. Install UNIX (HP-UX) hardware
3. Install software (FOCUS)
4. Port targeted data and application to UNIX from Mainframe
5. Test application
6. Proceed with porting all targeted applications and data to UNIX platform
7. Install TCP/IP on IBM mainframe
8. Implement (one direction) client/server (EDA/SQL) APIs into FOCUS applications on the UNIX machine (UNIX ---> mainframe)
9. Test client/server application
10. Implement (reverse direction) client/server (EDA/SQL) APIs into FOCUS applications on the mainframe machine (mainframe --->UNIX)
11. Test client/server applications
12. Test entire client/server (bi-directional) applications

Figure 4

Note: It should be noted that EDA/SQL has been incorporated into over 150 products on the PC including LOTUS DataLens, Excel, Clear Access, HP's Information Access and many more. Therefore, if this company decides to integrate PCs into their architecture creating a 3 tier client/server architecture

(EDA/SQL supports 3 tier client/server architecture) they would be in an excellent position to enable users to continue to maintain their investment in their favorite PC tool.

Original Configuration
The company was running an IBM 3090 model 600J with FOCUS applications with the FOCUS database (80 supported and maintained FOCUS databases). Users utilized PCs, with terminal emulator software, to link into the IBM mainframe.

New Configuration
HP-UX 9000 Series I40 with 256MB RAM with a floating Co-Processor with disk arrays with fiberoptic interface, Ethernet LAN Interface, TCP/IP, and GlancePlus, an HP performance analysis tool, and HP 9000 Series 720 workstations.
Note: The company chose the I40 for this project. They anticipate an upgrade will be required to accommodate the large user base.

FOCUS and EDA/SQL were installed on the HP-UX machine
IBM 3090 model 600J was maintained, and TCP/IP was installed on the IBM mainframe. In addition, EDA/SQL was installed on the mainframe computer.

Configuration

Original	*New*
Hardware	Hardware
IBM 3090 Model 600J	HP 9000 Model I40 with
PCs linked to IBM	Floating co-processor
	Fiberoptic interface
	Ethernet LAN Card,TCP/IP
	IBM 3090 Model 600J w/TCP/IP
Software	Software
FOCUS with FOCUS database	HP-UX: FOCUS, EDA/SQL
	Mainframe:FOCUS, EDA/SQL

The time to rightsize did indeed turn out to be about 18 months.
The major elements included:

1) Evaluate, choose and install UNIX (approximately 3 months)
2) Port software application (approximately 2 months)
3) Port all targeted applications (approximately 3 months)
4) Test ported applications (2 months)
5) Implement client/server (EDA/SQL) API calls into FOCUS UNIX applications (approximately 3 months)
6) Implement client/server (EDA/SQL) API calls into FOCUS mainframe applications (approximately 3 months)
7) Test entire configuration (2 months)

Challenges Uncovered
There were essentially no obstacles identified during the rightsizing, client/server implementation. However, one challenge that had not been resolved during the process was system management. They will investigate system management tools when reliable, tested, distributed system management tools become available.

Note: As system management and system management tools are important elements for a true, production, distributed, client/server environment, below, is a brief discussion of products that are available in the marketplace today.

System Management Tools for UNIX and distributed systems
System management tools on the mainframe are reliable and many choices of tools are available. The limited system management tools available on UNIX are new or ported directly from the mainframe and do not take advantage of the distributed integration capabilities of UNIX as well as other operating environments.

But now, this is beginning to change. A number of traditional as well as new system management tool vendors Tivoli Systems) are beginning to provide distributed system management tools.

For your information, the following is a list of vendors who provide system management software. This list may be instrumental for evaluating system management tools if you are considering implementing a client/server, distributed UNIX environment.

	Planning & Modeling	Performance Management	Fault Tolerance & Operations	Configuration Management (Hard/Software)	Accounting Management	Security Management
Boole & Babbage Inc.		X	X			
Computer Associates International Inc.		X	X		X	X
Compuware Corp.	X	X	X	X	X	X
IBM	X	X	X	X	X	X
Hewlett-Packard Co.		X	X			X
Landmark Systems Corp.		X	X			
Legent Corp.		X		X		
OpenVision Technologies Inc.	X	X			X	X
Tivoli Systems		X	X			

Source: DATAMATION: December 15, 1993

Results and Comments

This company was successful. They were able to implement a client/server, open systems environment that decreased the mainframe costs that were being charged back to this firm.

The success of this company's rightsizing, client/server endeavor was primarily due to the excellent planning and process that was implemented. The comprehensiveness of the plan and the careful watchful eye of the project teams over the project management schedule (schedule was adjusted during the process as required) so there were no unanticipated surprise, were important elements for success.

Since, there were project meetings with the vendors every week for the first 15 weeks and then bi-weekly meetings thereafter, any potential problem was identified and rectified quickly and efficiently.

Today, there are over 500 users accessing the applications running on the HP-UX and/or mainframe computer. They were able to meet all their goals and objectives as they were able to implement a client/server, rightsized environment, avoid the cost of a mainframe upgrade, and ultimately reduce the high, chargeback costs, of the mainframe computer.

When and if they chose to implement an OSF DCE environment, they are comfortable that their choice of partners (Hewlett-Packard and Information Builders) will help them implement these new technological standards.

Conclusion

Successful client/server, rightsizing and open system solutions are not only possible but have proved to save money and increase efficiency of operations within many companies in today's world.

Clearly, a realistic expectation, meticulous planning, dedication of resource (both internally and externally), and hardware and software expertise are critical elements required for success.

The open systems standards (i.e., OSF, COSE, etc.) as well as tools (such as system management tools) are continuing to evolve and become available for both UNIX as well for the distributed environment. As these tools mature, they become viable options to implement for a true, production, client/server environment.

In the end, client/server, rightsizing and open systems really do mean the freedom of choice. The freedom to maintain investment in hardware, in software, in people training, and the freedom to implement new technologies and services as they emerge in the, ever changing, ever evolving, arena of information technology.

Paper Number 2002
Managing Change in the Technical Organization

Steve Day
Senior Manager
Ernst & Young
901 E. Cary Street
Richmond, VA 23219
804.775.4433

Handouts will be provided at time of presentation

Internet-Who, What How

Steve Kaser
Technology Service Corporation

Disclaimer

Virtually Everything Contained herein is available on the Internet

What is the Internet

- "a network of networks based on TCP/IP protocols"

- " a community of people who use and develop those networks"

- "a collection of resources reachable from those networks"

History

- Started by the DoD called ARPANET

- Utilized a packaging scheme called IP (Internet Protocol) packets

- Allowed other networks using IP packets to connect to the Internet

Functionality

- email

- File transfer sites (FTP)

- Remote Logins

- gopher

- mosaic

Value

- Increased Access to Resources

- Customer Contact

How do You Gain Access

- Universities

- Private Carriers

- Access Methods

- Costs

Paper Number 2004
Evaluating and Implementing Office Automation Solutions

Eric Kowalski
Technical Representative
Higley & Company, Inc.
1920 West Eighth Street
Erie, PA 16505
814.454.5215

Handouts will be provided at time of presentation

Paper Number 2601
HP Customer Support Strategy Session

Doug Chapin and Doug Blackwood
Marketing Managers
Hewlett-Packard Co.
c/o Ella Washington
19091 Pruneridge Avenue, M/S 46L4
Cupertino, CA 95014
408.447.1053

Paper Number 2602
HP 3000 Strategy

Olivier Helleboid
Commercial Systems Division Marketing Manager
Hewlett-Packard Co.
c/o Ella Washington
19091 Pruneridge Avenue, M/S 46L4
Cupertino, CA 95014
408.447.1053

Paper Number 2603
HP-UX Product Line Update

Gary Lemke
Product Marketing Manager
Hewlett-Packard Co.
c/o Ella Washington
19091 Pruneridge Avenue, M/S 46L4
Cupertino, CA 95014
408.447.1053

Paper Number 2604
HP's Networking Strategy

Christine Martino
IND Marketing Communications Program Manager
Hewlett-Packard Co.
c/o Ella Washington
19091 Pruneridge Avenue, M/S 46L4
Cupertino, CA 95014
408.447.1053

Paper Number 2605
Enterprise Desktop:
Setting a Desktop Strategy

Janet Muto
Commercial Market Development Manager
Hewlett-Packard Co.
c/o Ella Washington
19091 Pruneridge Avenue, M/S 46L4
Cupertino, CA 95014
408.447.1053

The Warning Signs
A Pop Quiz on Quality

Paper #2702

By Robert Green and David Greer

Robelle Consulting Ltd.
Unit 201, 15399-102A Ave.
Surrey, B.C. Canada V3R 7K1
Phone: (604) 582-1700
Fax: (604) 582-1799
Email: info%robelle@mcimail.com

Everyone is in favor of software quality, but not everyone is producing quality software. How can you tell if a software group has gone off the track? It could be your DP department, your computer manufacturer, or even one of your software suppliers.

Introduction by Robert Green

I was presenting my paper, **Improving Software Quality**, before a hometown crowd in Vancouver, when a local consultant rose to ask a question. "Responding to the users sounds fine," she said, "but I've seen a lot of systems where the programmers did constant changes and patches to give the users whatever they wanted. Most of the systems were a buggy, kludgy, impossible mess. How do you avoid that?"

"That's the topic for another paper," I replied, with a quick tap dance to distract the audience. Later, I admitted to myself what a good point she had made.

Robelle creates software tools. About once a month we release a new version of Qedit and Suprtool. We are constantly on the go, so how do we avoid producing a disaster?

I went to my partner David Greer for the answer. At Robelle I'm the one who leans toward wild, creative impulses. David is stronger on reliability, discipline, and long-range thinking. I was so concerned by bureaucratic software groups which didn't satisfy user needs that I was overlooking another style of shop: groups with weak software skills who react to user complaints as fire fighting.

The ideal software group uses feedback and repeated development cycles to find out what users really need. But it also uses rigorous software engineering. Despite the constant changes, the ideal software group ensures that existing features continue to work and future changes are possible. My previous paper on improving software quality was only half the story. Now David Greer and I finish it off.

Teaching Quality Without "Should" and "Ought"

To avoid having our paper degenerate into a sermon, we have organized our ideas into a non-threatening Pop Quiz -- a way to quietly rate your quality record. The quiz consists of giveaway phrases that warn of a troubled software project, remarks such as "that's not my job" and "it's against our policy."

See how many of these phrases you've heard. Score 5 penalty points for each one you hear regularly. We made the penalty 5 points so you could fudge if you want to.

Disclaimer: Despite the fact that we use Hewlett-Packard for some examples, we are not out to get HP. We mention HP because we know them and you know them, not because we think their software quality is substandard. To show it isn't personal, we've thrown in a few Robelle follies too. All shops are guilty of some of these quality slips from time to time.

o o o

Ready? Here's our first warning sign, a classic programmer excuse:

That's a feature, not a bug.

If the user then points out the spot in the documentation showing that he is right, there is still the second tell-tale excuse:

That's an error in the manual.

With those two replies, a programmer can deflect any conceivable bug report.

The longest-running problem report at Robelle is caused by a "feature" of MPE. A batch job can log onto another MPE system by doing a Remote Hello Command and creating a remote session. But, if the job runs Qedit (or any program) on the remote system, MPE tells Qedit it is in an interactive session, not a batch job. Qedit cannot tell this phoney session from a true session.

Qedit attempts to do a CRT status request, assuming that the session comes with a person and a terminal. Because no one is there running the sesion, the first

Qedit command gets read as the CRT status! The DS/3000 Manual **from 1978** is aware of this problem and warns about it, but in 13 years HP has not fixed it. It's a feature, not a bug.

Reluctance to admit problems leads to a reputation as a "Black Hole". Questions go in, but nothing useful ever comes out.

o o o

When users gather, they swap "war" stories. One user was overhead at lunch, complaining about his software vendor. When he called the support line, their first question always seemed to be:

Why would you want to do that?

This statement suggests a superior attitude on the part of the programmer, an attitude that is seldom backed by reality. Customers do find unexpected and incredible ways to apply programs to solve their problems. The sign of a truly great program is that even when used in unanticipated ways, it still works.

Whether a program is useful is up to the customer to judge. Take an example from the carpet business, where despite superior technical knowledge at the factory, the customer knew best:

> One of our customers in Europe came to us several years ago with his own testing spec for carpet foam backing. We were a bit put out that someone thought they could test it better than we could. We told him not to worry. Dow measures for foam stability, molecular weight distribution, particle size conformity, percent of unreacted monomer, adhesion strength -- all the vital things. We told him, "You're going to get the best there is, real quality!"
>
> Well, three times we tried to ship him the order, and three times he sent it back. Now, that gets annoying. So we asked him, "What's the deal?" And he told us, "Your product can't pass my roll-stool test!"... What he did was take the bottom half of an office chair, put a weight on it, and spin it around on a piece of test carpet 30,000 times... If the carpet sample didn't delaminate from the foam, you passed the test and got the order... Quality is what the customer says he needs, not what our tests indicate is satisfactory. [I. Snyder of Dow, in T. Peters]

o o o

If there are no questions, everyone must be happy.

At a Management Roundtable where the users submitted only 18 questions, the moderator jokingly remarked, "This proves the customers are happy."

Wrong.

Customers who don't voice their complaints, when given a chance, have given up expecting answers. The unconscious attitude behind this warning sign is that customer complaints are an irritant to be tolerated.

Real user complaints are good, not bad. People who use and like software constantly think of more problems the software could solve. This shows up as increased "complaints". If you actually fix something for them, watch out! The complaints will escalate dramatically.

At Robelle, we have a saying: If we send out a batch of new software and **no one complains**, there is only one conclusion.

It means that **no one tried it.**

o o o

We've lost the source code.

On his first outside consulting job, Bob Green found that the client had lost his source code.

The client's billing program was taking 2 days to run on a sampling of accounts -- to do all the accounts it would be running continually. The system was designed by an expensive consulting firm, then programmed by contractors. The client had only a junior programmer to make the inevitable patches and fixes. There were lots of source files, but no one knew which ones compiled into the current version of the billing program. It took most of a day to find the proper source files and get them to re-compile. After that it took only an hour to fix the program.

Programs are valuable assets that depend upon an infrastructure for their preservation. If you take shortcuts in development, you will lose control of this asset.

What is needed is simple, but requires discipline. For every program, there is a job stream that recompiles the current version. This job stream shows the source files that go into the program and how to combine them. The test environment is kept separate from production. Another standard job moves a new version from testing into production.

o o o

We're too busy to document that.

Having each programmer in a shop re-invent the wheel is inefficient. Software techniques uncovered by a software group can be a valuable asset if they are documented. There is an even better way of retaining knowledge. Use **libraries of code.** Reusable code leverages the investment. One programmer suffers to solve a problem so the others won't have to suffer as well.

As an example, consider double-sided printing on the LaserJet. In our shop, Dave Lo first added this feature to our Prose text formatter by hard-coding the Escape sequence into the program. Then Bob Green hard-coded it into Qedit. When David Greer wanted to add double-sided printing to Suprtool, he asked, "Why is the Escape sequence hard-coded in Prose and Qedit, when we have a library routine that could have held that knowledge?"

David added the Escape sequence to our library code and removed it from the individual programs. Now we can add double-sided printing to the rest of our programs without having to re-learn the Escape sequence.

> The good programmer writes software that can be reused, even if he doesn't see a reuse immediately on the horizon. His experience ... assures him that someone from the next hallway, reasonably soon, will be asking, "Do you have a module that ... ." He also knows that writing reusably will force him to define clean interfaces. [D. Boundy]

o o o

When a programmer can't find the reason for a particularly bizarre and puzzling bug, there is one sure fire excuse.

It must be a hardware problem.

In all the bugs we have investigated in 13 years in business, only one turned out to be a true hardware problem.

Our Qedit program was repeatedly losing track of lines in a file at one customer site, but we couldn't repeat the problem on our machine. Through painful hours spent in Debug, we finally proved that the Load Instruction was failing on this one customer's computer, but only when indexing was done without indirection. And Qedit was the only program on the system that used this unusual mode of the Load Instruction.

The ugly truth is that hardware is amazingly reliable and software is not.

o o o

It would cost too much to update the paperwork.

The late Dr. Richard Feynman of Cal Tech did a famous study of the Challenger shuttle disaster. He ranged widely throughout NASA and its contractors, talking to anyone who could shed light on the quality problems.

A group of production workers had found a simple way to improve the calibration of the rocket engines, but it was never implemented.

> The foreman said he wrote a memo with this suggestion to his superiors two years ago, but nothing had happened yet. When he asked why, he was told the suggestion was too expensive. "Too expensive to paint four little lines?" I said in disbelief. They all laughed, "It's not the paint; it's the paperwork. They would have to revise all the manuals."
>
> The assembly workers had other observations and suggestions... I got the impression they were very interested in what they were doing, but they weren't being given much encouragement. Nobody was paying much attention to them. It was remarkable that their morale was as high as it was under the circumstances. [R. Feynman]

o o o

When people want to cut corners, you hear assertions like this:

No one will ever notice.

What they mean is that the users are too stupid to recognize quality when they see it.

At Ford Motor Company, they once came up with a scheme called PIP, Profit Improvement Program.

> The purpose of PIPs [Profit Improvement Programs at Ford] was to bring down the costs of making a car by taking them out of an existing budget; an example might be the decision to equip a Mercury with Ford upholstery, which was cheaper. Some traditionalists were convinced that the PIPs systematically reduced quality, that it was automotive sleight of hand, and that the covert philosophy behind the program was that the customer would never know the difference. PIPs quickly became part of the vernacular, turning into a verb. "What happened to that hood ornament?" "Oh, it got pipped." [D. Halberstam]

You seldom know which features will be important to your users. Success

demands attention to all details. In software, users notice the little things. They seem to be more sensitive to details than to the big picture, perhaps because they take the overall objective for granted but the details drive them crazy day after day. The one enhancement to Qedit which received the most positive feedback from the users was a tiny and simple change: allowing the entry of MPE commands without the preceding colon.

o o o

One more Go To won't hurt much.

Good programmer don't use Go To to solve their logic problems. They structures their code as easily as they breathe. They keep their programs within their intellectual grasp by using limited control structures: While, Do-Until, If-Then_Else, and Case. Keeping it simple keeps it easy to understand. The problem with Go Tos is that you can build convoluted structures with them.

Well-structured programmers limit the scope of their data structures. If a variable is only needed within a procedure, they make it a local variable not a global variable. If a procedure needs to access a global they either pass it in as a parameter or export it. If the programming language allows, they distinguish parameters which are input only from those that are both input and output (i.e., call by value versus call by reference). The fewer places in the code that can touch a variable, the easier it is to debug a program.

The competent programmer uses top-down design and bottom-up implementation:

> The first thing he does in any programming task is to analyze the entire problem into smaller problems that can be solved separately. He begins coding by writing functions that implement the primitives and building blocks he will need. The rest of the program almost writes itself. [D. Boundy]

o o o

We'll just reuse this data item.

The disciplined programmer does not give two names to one thing nor attribute two things to one name. Names are meaningful and specific, and their length is proportional to their scope. A loop variable used only once in a two-statement loop may be called "i", but a global variable that may be used anywhere in the program will have a long name that accurately describes its usage. [*Boundy's Laws of Naming*]

The disciplined programmer adheres to the standards of his workgroup, even

when these standards appear arbitrary. A standard as small as indentation style makes the code more readable for a person who doesn't know it. That person might be the programmer, two years later, after the code is forgotten.

A program is written once, but read many times. Why not make it easy on the next person who reads the code?

o o o

That could never fail -- don't bother testing for it.

When he first began to program, one of the authors of this paper (please don't ask which) developed an unfortunate impatience with writing code to test for error conditions. In his programs, he would skip error checking after certain operating system calls, such as Fclose and Fgetinfo which he "knew" could never fail.

Then he discovered that a :File Command could cause the Fclose Intrinsic to fail, leaving the file hanging open. And Fgetinfo failed once, undetected, when he closed the file by mistake, and another time when it was accessing a remote file on another system, and a third time when the calls were being intercepted by another vendor's run-time library.

This programmer learned the hard way to test every system call and every library function for failure. If he doesn't, his program may go merrily along, processing with the wrong data. He accepts the fact that most of his programs will have more lines of code to handle failure than to handle success.

Sometimes we develop a similar attitude toward manual procedures, such as installing a new version -- "I've done this so many times, I could do it in my sleep!" Implying, "I couldn't possibly make a mistake." Unfortunately, uninspiring tasks are the most likely place to misstep, since it is difficult to keep your concentration focused. When we install new software, we have to follow a written checklist. The checklist includes every step to create a new version of the product, plus steps to verify that the installation was done properly (i.e, make a demo tape and install it on another machine).

o o o

Fascinating project -- too bad it failed.

Fascination with grandiose schemes and bigness leads to **White Elephants**. The symptoms are pretentious objectives, high cost, fantastic claims, neglect of other projects, fanatical denial of failure, and a sudden, total write-off. All so unnecessary. In software, big results can come from tiny investments and tiny results can come from big investments.

> The original Lotus 1-2-3 (tm) and dBASE (tm) programs were two of the most successful application programs ever written. 1-2-3 was written mostly by one person in eighteen months. The macro capability, one of the things that made 1-2-3 really successful, was added by the developer at the end because he had some extra time -- it wasn't even in the informal spec he had. dBASE was written by one person over a two-year period while he also held a full-time job. [D. Thielen]

Hewlett-Packard has done two large, exotic database projects for the HP 3000 that were never released to users. Those projects consumed R&D resources that could have provided badly-needed enhancements to HP's popular but untrendy IMAGE product.

> There is a popular view that technology is only technology if it is high-tech -- sort of a "Big Bang" theory of technological development where somebody suddenly thinks of a major innovation or invention. But most of the technological change that goes on in society is not the "Big Bang" type. Rather it is small, gradual, marginal things. [Michael Bradfield, Dalhousie University, *Globe and Mail* newspaper.]

Robelle's biggest failure was the Virtual Fortran compiler. We had no Fortran expertise and no local test sites, but we did have big plans. We were going to run big, scientific Fortran programs on a 16-bit HP 3000, instead of waiting for a new RISC computer. Although we did complete the project, it was late and never performed fast enough. We were seduced by the glamour of the project. We squandered two man-years we could have spent on more humble but more successful projects.

o o o

We tested it once by hand, isn't that enough?

You can't test for the complete absence of bugs, because you can't try every path through the code. But you can test the typical cases and the boundary cases: minimum value, maximum value, and no value. In rigorous testing environments, such as software for jet aircraft, it is standard practice to measure the percentage of program statements being executed by the tests and aim for 100% coverage.

Automated testing is the answer. The computer can do more tests than we can manually and do them more reliably. We borrowed the idea for test jobs from the Pascal validation suite, a series of tests that told whether a Pascal compiler was up to the standard. A typical job tests one command, often by modifying data two different ways and comparing the results. For example, copy a file with Suprtool and again with Fcopy. Any difference and the job aborts.

When we are revising a program, we schedule the test suite to run at night. It is amazing how often a seemingly minor change causes 10 or 20 test jobs to fail.

When working on a bug, a good practice is to add a test that reproduces the problem. When the bug is fixed, the test passes. Reproducing a bug in the test suite also provides a warning if the bug creeps back into the code. It is embarrassing how often old bugs resurface.

o o o

It's fixed, but is waiting for the next release cycle.

Installing new software into production is an error-prone process. The program has to match the source code, the help file, and the manual; everything has to be tested; and so on. The more error-proof the installation process, the more onerous and time-consuming it becomes. This discourages frequent software updates, causing release cycles to stretch out until it takes two years to get the simplest bug fix into production. Look how many years it has taken HP to add support for IEEE floating-point numbers to TurboIMAGE.

The way to quicken development cycles is to **automate every step in sight.**

For our products such as Suprtool, we have a job stream that regenerates a new version. The job stream

- recompiles all the code, including supporting modules and programs,
- runs the test suite against the NM and CM versions of the product,
- reformats the documentation files to check for hyphenation errors, which it delivers to the programmer through electronic mail, and
- generates a new help file.

Now if we could figure out a way to have the computer actually fix the bugs and write the documentation, we could retire.

o o o

It can't be changed, too much code references it.

Global variables are the worst enemy of good software. This insight came to us in two waves.

First we learned not to use a literal constant, such as "80", when an identifier like "buffer-size" was allowed. Which would be easier if you want to change the buffer size?

Later we learned that this is not enough. The wider the access to any data structure, the more code to be checked when that data structure is revised.

For example, Suprtool has a fixed-size table for the Extract Command. On Classic machines, the space used by Extract limits the space left for buffering. We want to convert the Extract table into a linked list. So far we haven't found the time, because there are so many places in the code where the table is indexed as an array. If we were doing the Extract Command today, we would hide the data structure in a separate module. The rest of Suprtool would call procedures in that module to access the data structure.

The AIFs for MPE XL are a good example of making programs data independent. An AIF is a fast subroutine for accessing system tables. When a system tool uses AIFs, it doesn't know the location and structure of system tables. MPE XL may be improved and changed drastically, but the system tools will still run properly.

o o o

That would mean changing all the programs.

The year 2000 haunts those of us with only two digits reserved for the year instead of four. Will our invoices be dated January 1st 1900 at the turn of the century? How are we going to sort by date when "01" is less than "99"? Why did the millennium have to occur in our generation?

We wish now that we hadn't hard-coded the date format into all those programs, nor sprinkled ad-hoc code throughout our systems to edit and format dates. Couldn't we just extend the year 1999 instead of going to 2000?

With perfect hindsight, we would have used modular programming in the first place. Programs wouldn't "know" about dates -- they would depend on a **date module** for that knowledge; a module that holds all the functions and data structures for dates and reduces them to a clean, published interface; a module that edits dates, converts them between different formats, compares dates, and does date arithmetic. To change a date format, we change the module.

There is still time to correct past follies. You can break up large modules into smaller ones, write new, generalized modules to replace the old, restrictive ones, and design modules to be used as tool kits upon which other can build.

That should be enough to redeem past sins.

o o o

We gave the users what they asked for.

As an experienced developer once quipped, **"The firmer the specs, the more likely to be wrong."**

It is natural to plan for only one release of a new program. Unfortunately, the original software design is seldom what the users really need. The harder we pressure the users to tell us what they want, the more frustrated we both get. Users can't always tell us what they need, but they certainly recognize what they **don't** like when they see it.

The key to writing quality software is to **do it in several releases.** Once the users have the first release, the programmer incorporates their feedback (i.e., complaints) into the next release. At our firm, we send new "beta test" software to our customers every month. During a year, we may have 80 test installations for a product with 800 active users -- about 10% of the customer base.

At Robelle we use a development method called "Step by Step", created by Michel Kohon. It breaks a large project into small, two-week steps. This does not mean ignoring long-term goals, but once we learn more about the users' actual needs from each step, we commonly adjust our long-term goals as the program evolves.

> The final aim is the program, not the analysis. So, until the program or its results are in the hands of user, nothing is completed. [M. Kohon]

o o o

Computing Your Score

That is the last of our warning signs. Did you recognize many of them? Now calculate your score. Remember, five points per warning sign, and the lower the score, the better.

0 is too good. You must have cheated.

20 is excellent. Congratulations.

40 is respectable. Good work.

60 is worrying.

80 or more is a disaster. Update your resume.

Suggested Readings

Boundy, David. "A Taxonomy of Programmers." Software Engineering Notes (ACM SIGSOFT), October 1991.

Denning, Peter J. "What Is Software Quality?", Communications of the ACM, January 1992, Volume 35, No. 1.

Feynman, Richard P. "Personal Observations on the Reliability of the Shuttle." **What Do You Care What Other People Think?**, New York: W. W. Norton, 1988.

Gomory, Ralph. "Of Ladders, Cycles, and Economic Growth." Scientific American, June 1990.

Green, Robert M. "Improving Software Quality." **Steps to Software Quality.** Robelle Consulting Ltd., 1990.

Halberstam, David. **The Reckoning.** New York: Avon, 1986.

Kernighan and Plauger. **Elements of Programming Style.** Bell Telephone Labs, 1974.

Kohon, Michel. "Introduction to Step by Step." **Steps to Software Quality.** Robelle Consulting Ltd., 1990.

Peters, Tom. **Thriving on Chaos.** New York: Harper and Row, 1987.

Reich, Robert. "The Quiet Path to Technological Preeminence." Scientific American, October 1989.

Thielen, David. "Unconventional Thoughts on Managing PC Software Development." Microsoft Systems Journal, May 1991.

Paper Number 2703
This is Not a Paid Advertisement

Lisa M. Moose
Programmer II
Hoechst Celanese Corporation
P.O. Box 1000, Route 460
Narrows, VA 24124
703.921.1111

Handouts will be provided at time of presentation

Paper Number: 3002

Serving Up Client/Server: Baiting the Hook for Senior Management

Eric Eshelman, Product Engineer

SAP America, Inc.
International Court Three
300 Stevens Drive, Suite 350
Lester, PA 19113
(215) 521-4500

Introduction: Flexibility in Client/Server Architecture

Today, companies are faced with a large number of hardware, operating systems, databases, connectivity and software environments in a single organization. Many times these systems are incompatible. For IT to contribute to the re-engineering of a business, it must provide flexible, transparent technology tools for information access and process implementation.

The foundation necessary to achieve this is an enterprise-wide information technology and a software integration strategy. This strategy must provide the tools needed to enable user access to information, regardless of hardware, operating systems, connectivity, or database technology.

Integration

The term integration has been misused and mis-represented for a long time and means something different to almost every organization and vendor. In our view, applications software can enable integration through common data definitions and common data access.

This integration eliminates redundancy and gives users immediate access to information within inter-related departments. The result -- improved profitability, quality and customer service and streamlined business processes throughout the enterprise.

While many systems vendors claim to have integrated products, IT managers are frustrated with the limitations of most of today's systems. To have truly integrated processes, they know that systems, applications and platforms must be integrated, not merely interfaced. They want the control and visibility that integration provides across the enterprise.

Centralized Systems

Up to now, enterprise applications and the resulting integration were achieved only through centralized environments. Centralized systems have many advantages:

- Established technology. Centralized systems, whether mainframe or mid-range, employ established, stable technology. There are few surprises.
- Simple communication.
- Simple resource sharing. Computing resources are allocated on a first-come, first-serve basis.
- Solution to the problem of integration. With centralized systems, integration is easily achieved.
- Centralized systems are not perfect, though, especially in light of the emerging options.
- Enhancements and upgrades to the system must be done incrementally to ensure continuous availability.
- Response times are dependent on system usage.
- The user interface is not like pc-systems. Because of the relative "expense" of centralized - especially mainframe - MIPS, applications were designed to leverage functionality and economize on "style".

Integration: Flexibility with the Client/Server Approach

Applications integration brings tremendous benefits to the organization, in terms of streamlined processes, improved customer service and profitability. Client/server technology now provides additional flexibility to the organization.

The advantages of integrated applications with a client/server architecture are:

- Scalability. You can start in a centralized environment and scale upwards to two or three levels as your needs change, simply by adding additional application and/or presentation servers.

- Openness. The freedom to build your hardware platform from a number of different computer manufacturers.

- Guaranteed Response Times. Distributed computing power means that each user no longer competes with others.

- Graphical User Interface. This is the key to achieving end-user productivity. Because of the ease-of-use, such as pop-up windows and menus, training time is dramatically reduced, and ongoing navigation within the system is made easier.

Client/server technology is still in its infancy. As such, caution is necessary here. These issues will disappear as the underlying technology develops and gains acceptance in the market.

One disadvantage is:

- New Infrastructure. A client/server environment requires additional and new expertise. Expertise in networking, UNIX operating systems, relational databases.

However, the tremendous improvements in the price performance ratios of hardware platforms have led to distributed processing computing power. Let's now look at possible configurations for client/server architectures.

Client/Server Architecture: Centralized

SAP defines client/server as the ability to distribute processing among one, two, or three levels. The three components of SAP's client/server concept are separate database servers, applications servers, and presentations servers.

First, centralized systems -- In a centralized system, all operations run on a single computer. This computer handles input from all users according to the application logic specified in the application programs and answers their queries with the appropriate output. An environment we are all familiar with.

Client/Server Architecture: Presentation

Next, distribution of presentation load -- A first and relatively simple step towards the distribution of tasks is to link work stations or PCs as intelligent work place computers instead of simple, non-programmable terminals.

The tasks relating purely to the presentation of applications will then be shifted to these local computers. This greatly improves the user interface of applications. It makes possible the use of graphical output methods and makes the system more convenient to use and more productive to the end-user.

However, the actual processing logic of applications remains on a single computer, together with the central database.

Client/Server Architecture: Two levels

Next, client/server processing on two levels -- A fundamental aspect of SAP's client/server concept is that application logic runs on decentralized computers.

In a configuration like this, the central computer only controls database management as well as batch processing which is independent of user dialogs to the systems. The central computer is the database server.

The processing logic specified in the application programs are executed away from the central computer. Their execution and the presentation tasks are now the job of the decentralized computers which are grouped around the database server. We call these computer application servers.

The application servers execute all work on incoming data independently. During output to the user, the database server is only responsible for retrieving the data from the database; preparation and formatting of the data are jobs for the applications server.

For performance reasons, it is necessary to locally buffer the part of the database with which the decentralized applications are currently working in. The application computers can then execute longer dialogue processes independently, without needing to access the database server.

The application server only turns to the central computer when it wishes to read data or when changes to the central data base have become necessary.

Client/Server Architecture: Three levels

Finally, client/server processing on three levels -- As a further step, the user interface here can also be stored on an individual level.

This creates a three-level environment:

- Smaller computers are available to each user as presentation servers. They execute the user interface, which means they take input from the user and pass on data to the user.

- Application servers serve a large group of users, for example a department. All the processing logic of the application programs runs here at this level.

- The database remains under the control of a central computer, namely the database server.

SAP believes the best possible distribution of computer capacity is on three hardware levels. Why is that?

You can distribute the processing to computers, best suited for a specific process.

For example, response times, as one is used to with PCs, can only be achieved when a certain amount of computer capacity is kept available for each user when dialogue with the system occurs.

These individual capacities are often not exploited, because users do not typically require the same amount of capacity for dialogues on the presentation computer, but need it only sporadically. It is important, however, especially on the presentation level, that enough computer capacity is available for individual users so that their sporadic requirements can be responded to immediately.

Flexibility in Migration

SAP is uniquely positioned to offer solutions regardless of hardware platform. We provide flexibility to companies in migrating from centralized to decentralized computing environments.

Companies can begin with mainframe applications, start using workstations as graphical front-ends and gradually implement open systems as servers in the longer term.

SAP's satellite products use decentralized client/server technology and integrate into the centralized mainframe applications. Satellite systems, such as shop floor control, warehouse management and EDI, allow for additional flexibility in decentralized operations, where 24/7 is needed, or where graphical interfaces are critical.

Or, they can begin right away with R/3 using centralized mid-range systems, or a full-blown, client/server approach.

R/2

SAP's R/2 System integrates critical information flows throughout the enterprise. R/2 allows you to integrate a key functional area -- such as finance, sales, or manufacturing -- or, build a foundation for an enterprise-wide software strategy.

SAP integrated applications allow you to leverage your mainframe investment with software to address today's needs. And integrated workstation capability can help you build your client/server strategy in a mainframe environment.

The R/2 System supports multiple languages, multiple currencies and country-specific legal requirements. This makes it ideal for companies operating in the international arena.

R/3

R/3 is SAP's integrated application solution for open systems with client/server architectures. R/3 runs under Oracle's database and on UNIX operating systems from DEC, HP, IBM, Bull and Siemens-Nixdorf, and DEC's 0/VMS and HP's MPE operating systems.

R/3 follows the same business philosophy of R/2, as it supports enterprise-wide functions. Thanks to the scalability of client/server architectures, R/3 makes sense for a wide range of organizations -- from mid-sized companies to large corporations or their divisions and subsidiaries.

For customers currently using or considering the mainframe-based R/2 System, SAP will provide complete migration.

Conclusion

In summary, companies can now exploit client/server architecture to support the business re-engineering process. IT departments can bring their technology expertise to the table to serve the organization better.

Integrated applications can promote business process change and take advantage of the quantum improvements in the price/performance ratios of hardware platforms, relational databases, open systems standard and graphical user interfaces. The results, streamlined business processes, improved customer service, quality and profitability, will impact and benefit the entire organization.

Getting the Benefits of Data Warehousing Without the Warehouse

By Terry O'Brien
Executive Vice President
Dynamic Information Systems Corporation
(303) 444-4000

Getting the Benefits of Data Warehousing Without the Warehouse

Building Your Own Data Warehouse

It's been said that there are no really new ideas, just revisions of old ones. While I can't bring myself to agree with the premise, examples abound in technology of ideas that were proposed years ago, but are only now coming of age as a "new" idea. Data warehousing is a case in point.

IBM came up with the term years ago as a way to distinguish between the on-line databases optimized for transaction throughput, and the databases of more static data that can contain millions of records, optimized for fast searching and analysis. HP and other mini and UNIX vendors have capitalized on the renewed calls for strategic business use of data to herald the new age of "data warehousing" as the solution.

But rather than rushing head-long into our latest technological panacea, let's step back and look for a moment at why we would want a data warehouse to begin with. What are we trying to accomplish? Then lets look at the tools available to create that solution using the data and systems you now own.

Information, Now!

Marketing needs information on a weekly basis about the promotion success. Customer Service needs to have instant access to customer records. Manufacturing needs on-line access to documentation shared across the department. Sales needs daily access to telemarketing data and order fulfillment. And the Executive group needs ad hoc access to it all in summary form.

You could each write your own list of information needs based on your organization demands, but you can see the theme emerging: to run our businesses productively, we *all* need better, faster access to our on-line and historical corporate

data. But it's not as simple as just putting in a faster "box" or the latest PC front-end query tool.

Large databases

The amount of data collected has grown enormously over the years. Databases of millions of records are commonplace. Some companies are talking in terms of *terabytes* of data. Physically searching through these large databases for information, even with a screamingly fast processor, can take minutes, or even hours.

Disparate data

Corporate data is rarely stored in a consistent format throughout a company. Accounting may have its own dedicated system built around application tools it needs to do its job. Manufacturing has another. Sales and marketing another. Data can reside in different formats, different database structures, and even different hardware platforms. The promise of tying it all together has been an elusive goal for the last decade, and is still not in sight.

Ease of Access

Users are no longer satisfied with requesting a report from the MIS group and having it appear six to eight weeks later on their desk. They expect to be able to have direct access to that information from their desktops. They want it instantly, without having to know the arcane syntax of a cobol programmer's logic, and they want it in a format they can use with their point-and-click PC programs. The latest PC tools can be both a help and a hinderence. Explaining to the Senior Vice President he just can't do that kind of query on the system during working hours has been known to be a career-limiting move.

The Promise of Data Warehousing

So how does a data warehouse address these issues? If you've studied the traditional data warehouse models, almost all share the following characteristics:

1. A data warehouse separates the analytical information from the the on-line application system so that each can be maximized for its intended use: the on-line system for transaction throughput and the data warehouse for complex query access.

2. A data warehouse collects information from disparate database sources, pooling it in a system built for universal, easy access.

Summarizing Information

A key characteristic of most data warehouse systems is the fact that the data is usually not directly transferred to the warehouse from the operational system. It's often summarized and even pre-joined to make the queries more efficient. But this high level of granularity has a trade-off as well. First, when the database is refreshed with new data, all data must be re-summarized as well. Secondly, many of the pre-joined data segments compare "like" information: sales by salesrep by date, sales-by-customer-by-date, sales-by-product-by-year, etc. Much of the stored information is redundant. And third, these pre-joins are static. If a new comparison is needed, the data warehouse must be re-loaded with a new pre-join to support the comparison.

I'm not suggesting that you should not summarize, but if you can accomplish your access needs with a minimum of summarization, you provide the most flexible access with the least amount of overhead. Again, indexing all key values at the lowest granular level provides that flexibility.

So what have I desribed? I call it a "virtual" data warehouse: the benefits of a warehouse without the warehouse. Adding multiple keyword indexes on on-line data with client/server access can create a powerful enterprise-wide access system--at a fraction of the cost of conventional warehouse systems, and actually with less work.

Evaluate Your Needs Before You Buy

OMNIDEX, by Dynamic Information Systems Corporation (DISC) has been installed on systems ranging from state-wide insurance information systems to government procurement systems to catalog direct mail list systems to large manufacturing systems. Average retrievals are measured in seconds. And the average cost per system is under $100,000, including installation and setup consulting. OMNIDEX is now available on the HP 3000, HP-UX, DEC VAX and SCO UNIX. Expressway offers similar indexing for DG and SUN.

In contrast, most data warehouse systems average in the multiple hundreds of thousands, and require a proprietary warehouse data structure for the indexed data.

If done correctly, they both provide enterprise-wide data access. They both tie-in to client/server technologies. They both take you into the leading edge of information technology.

The choice is yours.[1]

[1]For more information on OMNIDEX, contact Dynamic Information Systems Corporation (DISC) at 303-444-4000. Expressway can be reached at 617-890-8670. Microsoft Visual Basic is available at most software retail outlets, and information on PowerBuilder is available from Powersoft at 617-229-2200.

Paper Number 3006
A Framework for Open Systems Planning

Robert Samway
Consulting Engineer
Power Engineers, Inc.
P.O. Box 1066
Hailey, ID 83333
208.788.0478

Handouts will be provided at time of presentation

3007
Integrated Imaging & Image Enabling

Adam Thier
Computron Technologies Corporation
301 Route 17 North
Rutherford, NJ 07070
201-935-3400

Your organization is being squeezed by new competition, while facing new opportunities as fresh markets open up. As a result, you are probably reviewing and revising your business strategies and the information management systems that support them. The challenge you face is to provide more detailed information, more quickly – while simultaneously improving productivity so you can accomplish this without adding staff.

In other words, you have to do more with less, and quite simply, the profitability of your organization is dependent on your ability to improve both the financial information you provide, and the efficiency of the tools and processes you use to gather it.

YOUR BUSINESS ENVIRONMENT

From a business standpoint, "Data" is no longer limited to just what is in an application's database. For every transaction there are an assortment of associated documents and transactions – word processing documents, incoming correspondence, invoices, E-Mail messages, and even voice mail.

What is startling is that your organization is spending millions on automated data processing when, in fact, only 3% of the information itself is being automated. The rest is on mag tape or microfiche (7%), and an amazing 90% of all the financial information you process resides on paper. As a result, the information resident in your information system is wholly incomplete.

Today, the U.S produces 1 billion pages of paper documents daily. And despite information technology, we are doubling our output of paper documents every six years.

Quite simply, your organization's information technology investment of the past 20 years -- which is roughly 60% of the money it spent on capital equipment -- has done nothing so far to bring about the promised "paperless office". Instead, it has just given your users the means to produce more paper more quickly.

The question you have to ask yourself is what is happening to all this paper, and how is it affecting the productivity of your staff? The answer is unnerving -- studies show that approximately 75% of your staff's time is spent shuffling paper (i.e. managing information). Look closely at your office. Your employees are probably busy referencing paper, making copies, etc. -- and every once in a while -- making an entry into their computer terminals.

INTEGRATED IMAGING

Integrated imaging is the solution to this paper mess. It brings all the information on paper into your information system by taking the purchase orders, receivers, claims forms, invoices, and other hardcopy documents, and turning them into digital computer images via a scanner. It then allows you to display these images on-screen simultaneously with the data represented by your software.

However, more than just displaying images, truly integrated imaging allows you to handle electronic files the same way you would handle paper ones. You can file *documents*, in electronic *files* and *folders*, and place them in many file *drawers* so you can retrieve them based on many different filing schemes. You can add temporary and permanent notes to your files. And you can access images for easy communications via Fax or E-Mail.

As a result, the simple challenge of finding an invoice when a vendor calls -- which probably now sends one of your clerks or managers scurrying for minutes, or even hours, through mountains of paper, searching for supporting source documents -- can now be verified in seconds by simply bringing the image of the invoice up on-screen. Effectively you regain the 75% of "lost" productivity, and integrated imaging in financial software will provide part of the solution for how you are going to improve productivity and do more without adding staff.

In short, the information architecture of tomorrow's financial software will be developed around the concepts of multi-media -- providing integration and access to all forms of information: data, text, image, and voice.

IMAGE ENABLING – WORKFLOW MANAGEMENT

There is another process, called workflow management, which enables the imaging and allows you to completely reengineer your business processes by automating not just the data and information, but the people and processes who manage the information.

In the past, the single largest part of your department -- your people -- has been ignored by your investments in information technology, and as a result, productivity has remained stagnant. Workflow management systems allow you to define the staff, supervisors, quality control and review personnel, managers, and executives that are involved in any transaction-based process. You indicate their authority levels and how and when information flows from person to person. Workflow management automatically routes information to these people based on your distribution (or routing) rules. Physical in-boxes are replaced with electronic ones, and automated routing rules replace manual copying and distribution. For the first time, your workload is automated, and can be measured. Work can be reassigned to accommodate changing manpower conditions, workloads regulated, and individual performance levels assessed at any time.

This vastly improves productivity and efficiency. You can do less with more, and gain complete control over not only the information you have -- but who has it, and what they are doing with it. And most importantly, you can easily change the rules determining the flow of work to reach an optimal degree of office efficiency.

ngly, and most importantly, workflow can be utilized with or without imaging. In paper-based processing departments, imaging simplifies access to paper files. However, work, in the form of data, without images, still needs to be automatically distributed, balanced, and measured. Therefore, the financial software system of the future will provide workflow management capabilities which can be implemented with or without imaging.

A CASE STUDY

In a "typical" accounts payable department, the process starts with a purchase order, which is sent to a selected vendor. In return, the vendor delivers the items or services requested, accompanied by a receiver, as well as a separate invoice mailed to the payables department.

Processing begins with the matching of the receiver, the invoice, and the purchase order. Most of the time these items are accessible (i.e. not lost or misfiled), and the quantities and prices coincide. In these cases, a clerk makes the appropriate number of copies, and then internally distributes them through the normal routing for payment approval (i.e. purchaser, department head, vice president for invoices over $10,000, etc.). If everyone acts on their approvals in reasonable timeframes the payment itself is generally not delayed.

However, the remaining 5 - 10% of "exceptions", whether they be the result of misplaced documents, or differences in price, terms, or quantity, cause enough problems in distribution and approvals (recontacting the vendor, reapproving new purchase orders, rerouting to other personnel, etc.) that payment is slowed dramatically. The result is, at best, lost discounts, and, at worst, payment of interest penalties.

Further, in both cases, once the invoice/receiver/purchase order packages enter their manual routing processes, their status is almost impossible to quickly ascertain -- there is no way of easily uncovering what in-boxes specific packages are in.

In other words, with manual workflow, the work itself is controlling the way this (and probably your) accounts payable department operates, rather than the other way around -- and workflow management can reverse this situation.

In a workflow integrated accounts payable department, the 5-10% of "problem" invoices, which cause so much lost productivity, become a non-issue -- and the status of all invoices is always readily available.

With integrated workflow, the accounts payable process begins with the scanning of all invoices and receivers as they arrive. These scanned images are then indexed to the appropriate vendors and purchase order files -- and then automatically routed to one or more people for approval. The approval personnel are automatically presented with the work that needs to be approved via their "electronic IN-boxes". For lower-level staff, the work is presented in order of priority. Upper-levels can browse through their IN-box and choose what to work on; but if they ignore a particular invoice for too long, the workflow system can insist that they deal with it.

If there is a problem with invoices and purchase orders not matching, the workflow system automatically routes the folder to specific personnel for exception processing. Personnel can pop-up a list of standard form letters that are

instantly personalized and available for FAX or printing. These letters automatically become part of the electronic file. If a response is pending longer than a pre-defined timeframe, the workflow system takes further routing and notification actions so as to insure that at the very least, desirable discounts are taken, and at worst, payments are made on time and finance charges avoided.

REAL RESULTS

What makes workflow so natural a fit for financial applications is that all companies already have approval processes and routing procedures in place. It is simply a matter of establishing these same "rules" in the graphic routing environment which controls the workflow system.

In the application of workflow technology, productivity increases dramatically while savings are realized. Some typical results of integrated workflow with accounts payable are:

- o 83% reductions in invoice processing time
- o 90% reductions in inquiry response time
- o 77% savings in time required to fulfill correspondence
- o 63% reductions in the time required for historical inquiries
- o 95% savings in document storage costs
- o 80% reductions in outbound phone calls

CONCLUSION

The 1990's have ushered in a new business climate. One that provides new opportunities for success....or failure. Those organizations that survive and prosper will be those able to respond quickly and effectively using the information system tools they implement today. Tools that no longer focus solely on the functional aspects of the task at hand, but also on the productivity issues of how the tools themselves are being used.

Imaging and image enabling tools such as workflow management address these productivity issues by combining to provide users with not only immediate, on-line access to all the information crucial to their jobs, but with the means to control and streamline the business processes that revolve around that information.

Networking Fundamentals for Managers

Paper # 3601

Steve Beasley
Hewlett-Packard
2205 E. Empire Street
Bloomington, Illinois 61701
309/662-9411

One of the biggest challenges for management in today's complex information processing environment is maintaining a level of technical awareness and expertise. The increasingly rapid pace of technological change creates an enormous challenge in trying to understand the technical level of detail which is often necessary with management decisions. While an effective manager may delegate much of the technical investigation work to the appropriate analysts, most managers will continue to want the summarized data in order to ensure the best possible decision for their organization. Unfortunately, many of today's technologies, like networking hardware, network operating systems (NOS's), and application programmer interfaces (API's) can require a fair degree of technical detail to cover just the basics.

Fortunately, networking is one of those technologies that can be broken into simple and manageable terms and concepts. Several terms, acronyms, concepts, and functionalities will be discussed below, with a goal of providing a good foundation of networking concepts to the manager of an HP computer environment. Unless noted otherwise, all of the concepts below are applicable to both the HP3000 and HP9000 environments.

In an effort to inject some logical sequence to this discussion, the following topics will be covered in order: 1) the network layers, 2) networking hardware (LAN's, WAN's, bridges and routers), 3) network software, 4) the API's and middleware, 5) network management, and 6) managing the resources (people, support, training, etc.).

THE NETWORK LAYERS

Almost everyone who has ever attended a conference or training on networking has seen the seven-layer International Standards Organization (ISO) "Open Systems Interconnection" Reference Model, and most would just as soon not ever see this model again. However, to set the foundation for this overview on

ISO Model

Layer	Functionality
7	Application
6	Presentation
5	Session
4	Transport
3	Network
2	Data Link (Hardware Interface)
1	Physical Hardware Connection

Figure 1

networking, we must discuss it very briefly. The ISO model contains seven conceptual layers, as outlined in figure 1.

While it is not critical to understand the functionality of each layer, it is important to understand that each layer performs a specific set of functions, then passes the data to the next layer above or below. Equally important is the fact that one layer is not necessarily dependent upon the layer above or below, allowing us significant flexibility in mixing protocols and hardware. For example, a Network Operating System (NOS), such as Novell's Netware, will typically execute over both token ring or ethernet. The majority of the NOS functionality is at layers 3 and 4 of the ISO model. Layer 2 would provide a unique driver interface to the ethernet card or to the token ring card at layer 1, thus allowing a single version of the layer 3 and 4 software (IPX, for example) to run over either physical network. If a new physical networking topology becomes prevalent (perhaps 100VG-AnyLAN), a new driver at Level 2 can be written that provides the interface between Layer 3 and the new physical card at Layer 1.

This ability to layer the software allows for some interesting combinations of protocols in certain vendor's networking products, such as, Sockets over SNA or even APPC over TCP/IP. These protocols, along with several others, will be defined and discussed in more detail later.

THE NETWORK HARDWARE

The most basic component of a network is the physical media over which the data must move. This component is the first or bottom layer of the ISO reference model, called the physical hardware connection. Typically, this connection will be to a local area network (LAN) or a wide area network (WAN).

There are two types of LANs most prevalent in today's networking environments. Both of these types have been defined by the Institute of Electrical and Electronic Engineers (IEEE) and have an IEEE standard associated with each. The IEEE 802.3 standard defines the type of LAN commonly referred to as ethernet. In reality, ethernet and 802.3 are slightly different, but most people use the terms interchangeably. The HP platforms will support both the 802.3 format and the ethernet format. IEEE 802.5 is the standard for a token-passing ring LAN and is compatible with the IBM token ring implementation.

Most HP computers are connected to either an ethernet LAN running at 10 million bits per second, or a token ring running at 16 million bits per second. An

ethernet LAN uses a protocol called CSMA/CD (Carrier Sense, Multiple Access, with Collision Detection) to allow several nodes (PC's, HP3000's, HP9000's, etc.) to be connected simultaneously. Each node can attempt to send data at any time onto the network, after first "listening" to determine if any other node is sending data. If the LAN is not in use, the node can send data. If two nodes send data at exactly the same time, a collision is detected. Each node will "backoff" or wait, then retransmit the data at a different interval. This scheme is simple, elegant, and allows for excellent throughput up to a point. Once the LAN becomes very busy (at some point between 50% and 80% of capacity), then the collisions start to cause a significant degradation in performance.

Token Ring, on the other hand, uses a token passing method to ensure that every node gets an equal chance at sending data. No station can send data unless it has the token. After sending data, a node passes the token to the next node on the ring. This method provides consistent performance throughput, even on busy networks. The disadvantage is that a node must wait to receive the token, which is not very efficient on a lightly utilized network. An ethernet LAN (at 10 megabits/second) will typically perform better than a token ring LAN (at 16 megabits/second) until the ethernet traffic starts to reach the physical capacity, at which point token ring performs better. Typically, ethernet cards and the physical cabling are significantly less expensive than the equivalent token ring hardware. When choosing between ethernet and token ring, one must consider both performance and price. This is one decision, in particular, where management must consider the technical details. The token ring vs. ethernet battle has evoked the passions of a "holy war" at many companies. Due to substantial cost differences and the performance considerations, the technical details must be emphasized, while keeping the "religious" aspects closely monitored.

For those sites requiring truly high speed communications, the Fiber Distributed Data Interface (FDDI) connection is available on the HP9000's and will be available on the HP3000 later in '94. Although FDDI is typically used as a site backbone and not as the LAN for the local workgroup, some workstations require the 100 Megabits/second bandwidth of FDDI.

For communications outside the local site, the use of some type of "wide area network" is necessary. Typically, this communication will occur using the resources of the telephone company and some type of leased or dial-up (switched) phone line. Typical wide area links may run as slow as 1200 bits per second, such as when using an inexpensive modem to call from the PC at home into the office. Modems, using either a dial-up line or dedicated leased line, will operate at speeds from 1200 bits/ second up to 56,000 bits/second (56Kbps) and higher. When additional bandwidth is necessary, many companies have

implemented digital T1 circuits, providing over 1.5 Mbps. While not very common, some sites require the 45 Mbps available with T3 communication links.

One common facility used with WANs is the X.25 packet switching protocol. X.25 is a standard published by the CCITT standards organization and defines a packet switching protocol designed to efficiently move packets through a distributed, wide area network. A company may own its own X.25 packet switching equipment, thus creating a private packet network. Or, many companies have purchased the X.25 service from a variety of public carriers, letting the service provider manage the network. The demand for X.25 products and services has been declining over the past 5 years, primarily due to the growth of extended LAN's and other new high-speed technologies.

New wide area technologies, such as Asynchronous Transfer Mode (ATM), Switched Multi-megabit Data Service (SMDS), Frame Relay, and Integrated Services Digital Network (ISDN) are all relatively new services provided by the phone companies to enhance the existing wide area technologies currently available. ISDN has seen tremendous success in Europe and Japan, but has not been enthusiastically implemented in the U.S. The high-speed capabilities of ATM, SMDS, and Frame Relay each have their own supporters who believe their chosen technology will win market acceptance. At this point, ATM appears to be winning the race, but I'll not make any predictions.

With the decreasing costs of high-speed wide area links, the concept of the extended LAN has provided a tremendous increase in networking capabilities for many companies. The local area network can now be extended, using the wide area link, to include remote locations. In order to extend the LAN to the remote site, a bridge or router is generally used.

A bridge operates at the data link layer (layer 2) of the ISO 7-layer model. A bridge simply connects two LAN's, providing easy access to the nodes on either side of the bridge. Bridges typically operate in pairs, with one bridge at the local site and the other bridge at the remote site. If node A sends data to node B on the same LAN, then the bridge will not forward the data to the remote bridge. If node A sends data to node X which is not on the same LAN, then the bridge will forward the data (see figure 2). The bridge learns which physical hardware addresses are located on its local LAN and on its wide area link; any traffic destined for an unknown hardware address is forwarded to the remote bridge or bridges. The bridge then learns that new physical address. Bridging provides a simple, easy-to-manage method of connecting two or more LAN's.

Routers, while similar to bridges, offer more functionality because they operate at the network layer of the protocol stack (layer 3). The router will "route" a data packet to a particular link, based upon its destination address. For example,

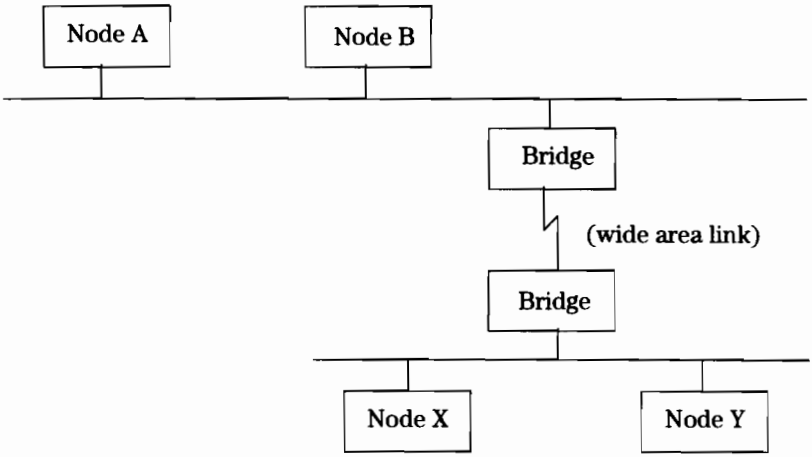

Figure 2

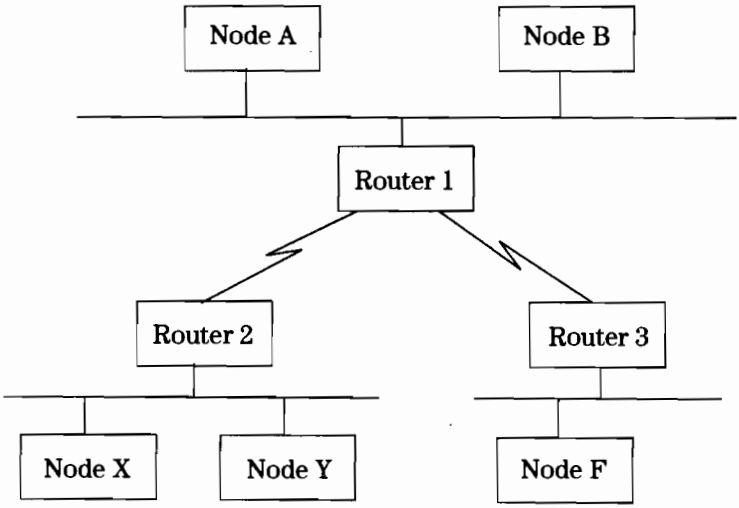

Figure 3

Novell network traffic is routed based upon the IPX address. In a TCP/IP network, packets are routed upon the IP address. The added functionality of routing allows greater control of where packets are sent into the network. If node A sends data to node F (see figure 3), then Router1 will send the data to Router3, but not to Router2. (Bridging would send the data to both remote LAN's.) The routers make this routing decision based upon the addresses of the software being used at layer 3. This network software is discussed in much more detail below. Routers require more configuration than bridges, but also provide more functionality and control of the network traffic.

THE NETWORK SOFTWARE

In order to keep this discussion to a reasonable length, only the most prevalent networking products found in the majority of HP3000 and HP9000 shops will be examined. These most common products or architecture include: Novell's Netware, Microsoft's LanManager, TCP/IP, SNA ("traditional" and APPN), and OSI.

Netware

Novell's Netware is by far the most popular network operating system (NOS) for PC LAN's on the market today. As a PC NOS, its primary function is to provide print and file sharing services to the client PC's, allowing them to access files and printers on the PC server.

With an estimated 50-60% of the PC networks running Netware, many host computer vendors, like HP, have chosen to implement Netware on their hardware platforms. As an example, both the HP9000 and HP3000 have a Netware implementation that allows the HP host to act as the file and print server. Although technically providing the basic functionality of a 386- or 486-based PC server, an HP host server can provide access to a much larger disc configuration for file sharing services.

Netware uses a protocol called IPX for much of its networking functionality. IPX is a Layer 3 and 4 protocol that is independent of the physical media of the LAN, and therefore, works with many different topologies. The diagram in figure 4 shows the relationship of Netware to the ISO model. Please note that Netware, like many of the other protocol stacks, does not match up perfectly to the layers as defined by the ISO model.

Netware can also be implemented using the TCP/IP protocol to provide the layer 3 and 4 functionality. Because the Netware Core Protocol (NCP) at layer 5 is written to interface with either TCP or IPX at layer 4, we can utilize either

NetWare

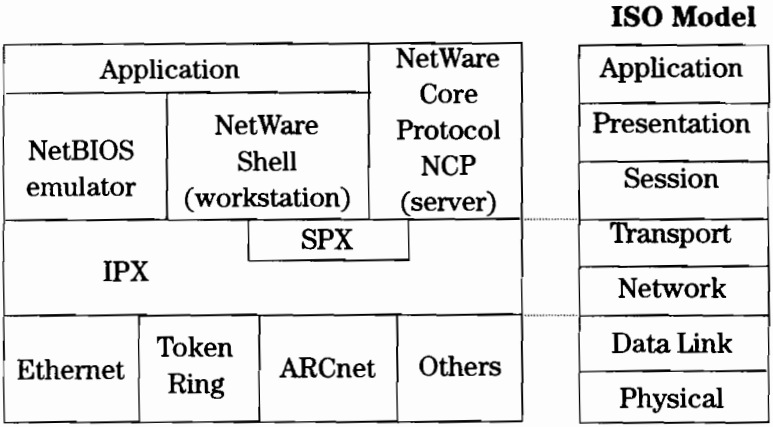

Figure 4

LANManager

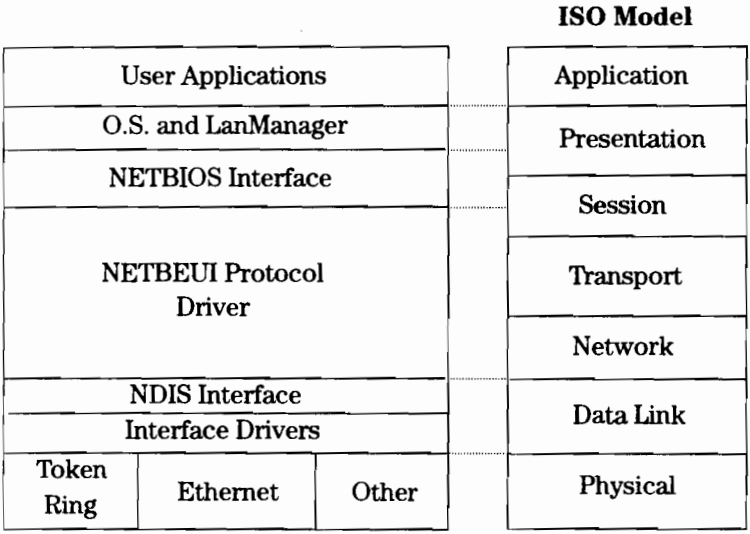

Figure 5

protocol stack. TCP/IP has the advantage of being more "industry standard" than Novell's proprietary IPX stack. However, the proprietary implementation of IPX may offer better performance.

LanManager

LanManager is another network operating system that provides file and print sharing services for PC clients using a common server platform. And, as with Netware, many vendors have written a LanManager software facility to run on their host computers, including both HP platforms. Again, this allows the utilization of a host system and its multiple printers and large disc configurations as the server platform. LanManager uses the protocols called NetBIOS and NetBEUI to perform the functions provided by IPX and NCP in the Netware stack, as outlined in figure 5.

IBM's LanServer is a third network operating system and is very similar to LanManager. IBM licensed the LanManager technology for its own LanServer products. The LanServer protocol stack is almost identical to the LanManager stack, particularly at the lower levels. However, neither the HP3000 nor the HP9000 currently interoperates with LanServer.

TCP/IP

TCP/IP is one of the most common network protocols, originating in a host-to-host network environment. The TCP/IP protocol gets its name from the Transmission Control Protocol that runs at layer 4 and the Internet Protocol that runs at layer 3. Like the other protocols referenced above, TCP/IP is not dependent upon a physical media. Each hardware vendor that supports TCP/IP is responsible for writing the layer 2 software that communicates over token ring, ethernet, or other physical media, such as some type of wide area link.

In addition to providing the layer 3 and layer 4 software functions which give TCP/IP its name, most vendors provide a "suite" of services that reside at the upper layers (5, 6, and 7). These services include Sockets, File Transfer Protocol (FTP), Telnet, and Simple Mail Transfer Protocol (SMTP). This suite of services is typically called the ARPA Services. (See figure 6).

The ARPA Services are a particularly powerful feature of the TCP/IP protocol, due to the application functionality that these services offer between multi-vendor host computers. A user on one host computer can generate e-mail using SMTP, initiate file transfers using FTP, perform remote logins using Telnet, or write an application program using Sockets, all to a remote host on a completely different hardware platform. This suite of services has enabled

TCP/IP Protocols and the ISO Reference Model

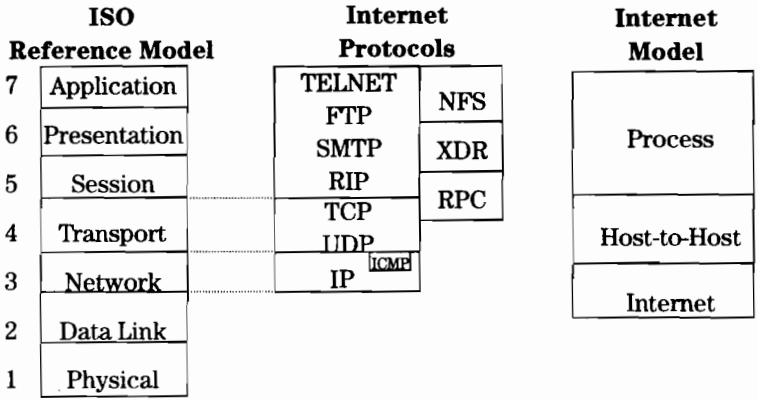

Figure 6

SNA

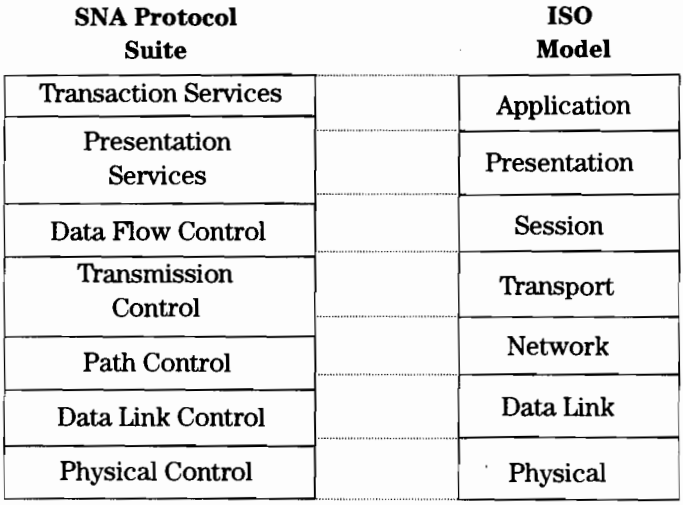

Figure 7

inter-vendor communication to occur easily and consistently on almost every major hardware platform.

SNA

IBM's System Network Architecture (SNA) was introduced in the early 1970's as a proprietary networking architecture for connecting mainframe computers and their peripherals such as card readers and terminals. Today, SNA has grown to incorporate both minicomputers and personal computers (PC's), along with the mainframe, to provide a comprehensive set of network services. This relationship is illustrated in figure 7. The SNA architecture was used as a reference for the ISO's OSI model, although like the other protocol stacks, SNA does not match up "layer for layer" with the model. (See figure 7).

In a traditional SNA environment, the networking occurs in a hierarchical fashion, requiring the "assistance" of a front-end processor (FEP), which runs the Network Control Program (NCP) software. An example of an FEP is the IBM 3745 communications controller. This device executes the NCP software and controls all remote resources connected to it. All devices in a traditional SNA network must be configured on the FEP associated with that device. A network may consist of multiple FEP's, which can then provide the network link between remote nodes.

IBM's Advanced Peer-to-Peer Networking (APPN) is the SNA approach to a more dynamic, non-hierarchical network design. With the increasing presence of PC's, minicomputers, and other workstation-type hosts, the need for peer-to-peer connectivity is critical. Protocols such as TCP/IP, Netware, and now APPN provide this ease of connectivity throughout the network.

OSI

The Open Systems Interconnect (OSI) protocol was intended to be the industry-standard solution for connecting heterogeneous networks. The OSI protocol was developed through the international standards bodies and includes services similar to the ARPA services (i.e., File Transfer Access Method, or FTAM, provides functionality similar to FTP, File Transfer Protocol.) The U.S. Government jumped on the OSI bandwagon with their GOSIP (Government Open Systems Interconnection Profile) directive, which specifies the mandatory implementation of OSI protocols by all U.S. government organizations. The OSI protocol stack is outlined in figure 8.

Although most of the major computer vendors now have the OSI software available, customers have been slow to embrace the technology, apparently for at least two reasons. First, while the standards bodies were discussing and defining

The OSI Protocol Suite

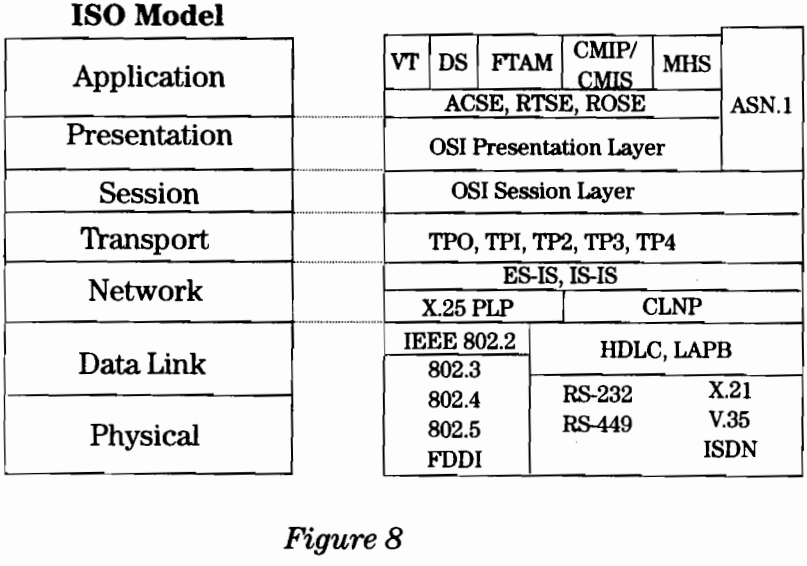

Figure 8

OSI, the user community needed the functionality. TCP/IP was widely available and relatively inexpensive, so most customers implemented TCP networks. Next, when the OSI software did become available, most customers found it to be relatively expensive and not as efficient as TCP. For these two reasons, the future of OSI as the network protocol of choice is not clear. Many users may stick with their TCP networks and their own user-written applications. On the other hand, many of the services available with OSI (like FTAM and X.400) are excellent services, and there is significant speculation that the vendors will provide the OSI services on top of the TCP transport, giving users the "best" of both worlds.

APPLICATION PROGRAMMER INTERFACES (API'S) and MIDDLEWARE

As with many of the technologies already discussed, there are several choices of API's available to the network programmer. However, for the typical HP3000 or HP9000 network programmer, three API's dominate: NetIPC, Sockets, and APPC. NetIPC is available on the HP3000 and on the HP9000, although HP has announced its discontinuance for the HP9000. NetIPC is a proprietary programmatic interface developed by HP, providing peer-to-peer connectivity for network applications. Although proprietary, a NetIPC application will interoperate with a Sockets application.

Sockets are part of the TCP/IP protocol suite and were developed as a means for two processes to communicate. A socket data structure allows messages to be sent and received. Two endpoints comprise a pair of sockets through which communication may occur in either direction. Sockets allow for two-way communication with both intramachine and intermachine processes. Both the HP9000 and the HP3000 support the Sockets interface, as do most other computer vendors, including DEC and most of the IBM platforms (MVS systems, AS400's, RS/6000's and OS/2 PC's).

In some environments, communication to IBM via APPC (SNA's Advanced Program-to-Program Communications) is required. The SNA LU6.2 products on both the HP3000 and HP9000 platforms provide the APPC interface for communicating to IBM. APPC provides a conversation between two endpoints (much like a Sockets pair) but provides additional functionality that the Sockets implementation does not provide. This functionality includes security features, scheduling mechanisms, and better error handling conditions. These additional features are typically provided at the application level in a Sockets application. While the additional features of APPC are very useful, some people argue that those functions are best left to the application, for performance reasons.

CPI-C is the Common Programming Interface for Communications, which is the interface between application programs and APPC. Apparently, the APPC

interface is implemented differently on the various IBM platforms. In order to provide a common method of accessing APPC function calls, the CPI-C interface has been implemented on most of the IBM platforms, providing the common interface.

One interesting twist in the API considerations is the ability introduced by IBM to allow a Sockets program to communicate via an SNA network. This feature shields the application programmer from knowing or caring about the network transport underneath the application. Similarly, IBM has announced plans for APPC over the TCP/IP transport. Again, this allows the flexibility of using APPC applications over a TCP-based network.

Middleware

The API's mentioned above have been the dominant means of providing application level interfaces between computer systems for at least the past 5-10 years. However, software products called "middleware" are beginning to gain a significant presence in the industry for communicating between networked systems. The most common software for the commercial environment is the Open Software Foundation (OSF) Distributed Computing Environment (DCE). The DCE software is compromised of several functions, including a remote procedure call mechanism, distributed timing function, a very robust security mechanism, directory services, and more.

The intent of the DCE software is two-fold: to provide application level access across multiple hardware/software platforms and to shield the programmer from many of the complexities of networked applications. Features like the remote procedure call (RPC) allow easy access to procedure calls on either local or remote heterogeneous systems. The security mechanism provides a sophisticated method of both authentication and authorization, along with the ability to encrypt data. The cell directory service provides the obvious directory services that are needed in a networked environment. All of these functions are coded to interoperate with DCE implementations on almost every major platform. The DCE software has been available for several months on the HP9000 systems and will be available in production release on the HP3000 during the second half of '94.

DCE, as well as the other middleware software, assumes a stable, interoperable network is in place, typically based upon TCP/IP. This middleware would, for all practical purposes, exist at layers 6 and 7 of the ISO model. However, most of these technologies are fairly extensive and explode these two layers into multi-layered, multi-functioned architectures, as outlined in figure 9.

Distributed OLTP with DCE Foundation

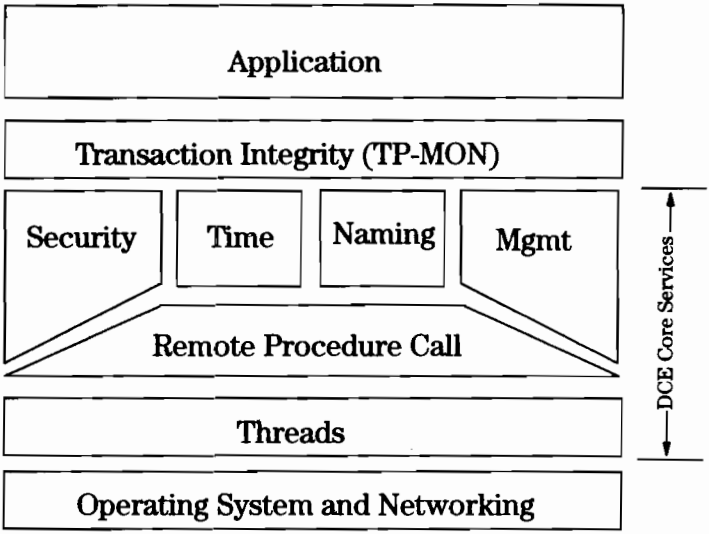

Figure 9

NETWORK MANAGEMENT

Network management is a multi-dimensional problem which can be viewed from several perspectives. In general, network management refers to the following functions: 1) configuration and name/asset management, 2) performance management, 3) problem and fault isolation/management, and 4) change management.

There are two primary protocols which the industry has selected for network management. The first of these protocols, called Simple Network Management Protocol (SNMP) was developed for the purpose of managing a TCP/IP network, although SNMP is not officially part of the TCP/IP protocol suite. SNMP has gained dominance (with approximately 35% market share), largely due to its simplicity. The competing standard is the Common Management Information Protocol (CMIP), with approximately 14% market share. CMIP offers features not provided by SNMP, including the ability to run on any network.

While the underlying network management protocol is certainly relevant, the applications that use the management protocol provide the real functionality. While there are several network management platforms, two products dominate the market: HP's OpenView family of network management products and the IBM NetView platform of products.

HP's OpenView strategy is to provide the best of the SNMP and CMIP capabilities under a single platform. The OpenView platform provides a full set of network management products, running in both a Windows environment and the HP-UX environment. The OpenView products allow a single administrator to manage a wide variety of systems, network components (hubs, bridges, routers, etc.), and network infrastructure from a single location. In addition to widespread acceptance in the HP customer environment, HP has licensed the OpenView platform to a large number of third-party network management vendors and software suppliers, including IBM for its NetView/6000 platform.

As mentioned earlier, no single standard or platform exists for network management, although many people were hoping to see the Open Software Foundation's Distributed Management Environment (DME) fulfill the promise of standard network management. Many organizations, including HP with its OpenView features, have contributed to the DME and its standards-based architectural approach. Unfortunately, due to a variety of unforeseen delays in the delivery of the framework, the future of DME is less than optimistic. However, based upon the work of many vendors, including HP, industry experts expect to see many of the desired features of DME available in products such as OpenView.

A discussion of networking without network management would be incomplete. However, the network management environment is very complex and deserves a much more detailed description than what can be provided in this format.

MANAGING THE ANALYSTS

For the manager of networking analysts, engineers, gurus, or whatever title they use, there are certain key elements that can make everyone's job much easier. These include the proper training, the right publications, consulting when appropriate, and a general understanding of the environment.

Training

Networking is arguably one of the more complicated aspects of today's computing environments. As such, a trained networking specialist is a requirement for those companies who desire in-house support. While there is no substitution for on-the-job training, this method falls miserably short with many new technologies. Formal training is often the best, most time efficient and cost effective method of providing the necessary education to the specialist. Additionally, since most companies do not have the hardware budget for equipment to "play with", the hands-on labs of most training sessions provide an excellent opportunity to meet the educational needs of the specialist. Hewlett-Packard offers some excellent training courses, as do many other training sources.

One especially worthwhile "training" event is the Interop conference held twice each year. Interop began in the fall of 1987 as a very small interoperability conference and has now grown to one of the largest networking conferences in the U.S. In addition to several conference sessions and an excellent demonstration floor with all of the latest networking products, Interop offers two-day tutorials on a variety of subjects and taught by some of the industry's leading experts. This conference is a very "high want" if not a must for the networking specialist. Another event that can definitely benefit the network guru of an HP shop is the Interex Conference, held each fall. While Interex does not focus entirely on networking like Interop does, the networking "track" at Interex does focus on HP networking issues and can provide a significant level of education.

Publications

The networking industry has no shortage of magazines that specialize in networking technologies. Two of the best seem to be "Network World" and

"Communications Week". Both of these trade magazines do a relatively good job of publishing articles on network hardware, network software, and even the telecommunications aspects of networking.

A manager (or analyst) must use some level of discretion, however, when reading these or any other networking magazines. First, if one reads every article in its entirety, a good portion of the work week can be wasted reading articles that may not be relevant for one's own environment. These magazines are meant to be skimmed for the appropriate articles and for gleaning the key points. Managers, especially, must be cautious of "managing by magazine", which is an easy mistake to make. While the magazines do provide some good insight, what works for one organization may not work for another. The advice of the in-house expert should typically carry more authority than last week's "Network World".

Consulting Assistance

Many companies are overly anxious to hire outside consulting services to offset a lack of expertise within their own company. Other organizations will resist consulting at all costs, in spite of a severe shortage of technical expertise. As with most things, a middle ground position is usually the most effective. In many instances, an organization could find it more cost effective (particularly in the long run) to send their own resources to specialized training, instead of spending the budget for consulting. This can work for specific projects and when the potential return on the investment can be seen in the near future. In other scenarios, outside consulting from network specialists can be very cost effective for an organization, particularly when planning for new network topologies, when trying to optimize existing resources, any scenario involving new network management applications, as well as projects undertaking client/server technologies. While consulting never looks "inexpensive", it may be much more effective than having internal specialists struggle on critical projects due to a lack of expertise or experience. And, its always good to have a second, neutral opinion when moving to a new technology.

SUMMARY

Networking is a broad and complex subject matter. As always, any "simplification" of a topic can lead to some perceived inaccuracies, due to not covering the full range of details. Hopefully, this effort to simplify some concepts and terminologies has given the reader a better understanding of some of the foundations of networking.

Paper #3602
Building a Robust Infrastructure
for Client/Server Applications

Orrin Mahoney
Hewlett-Packard
19091 Pruneridge Ave. M/S 44L3
Cupertino,Ca. 95014
(408) 447-1206

HP's Information Technology Framework

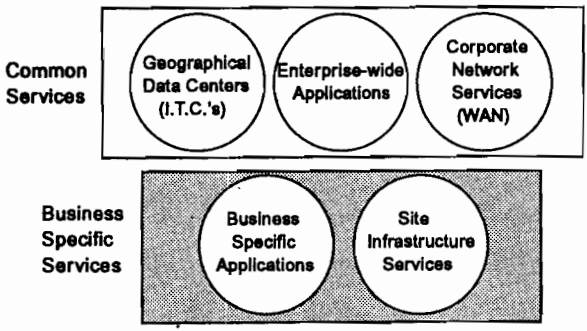

Infrastructure - as described by Websters Dictionary

Infrastructure - The underlying foundation or basic framework (as of a system or organization).

CSO I. T.

- **Infrastructure Services**
  - *HelpDesk*
  - *Network Connectivity (Local and Wide area)*
  - *Desktop (PC & WS) personal productivity solutions and support*
  - *Server support and management*
  - *Voice, Video and Data communications support*
  - *Electronic communications support*

CSO I.T. Infrastructure

Deliverables

- Excellent base level services (Voice, Networks, PC and Workstation Support) with emphasis on high customer satisfaction (customer delight) and aggressive cost containment.
- One-stop problem resolution for site customers.
- Advanced infrastructure implementation to support breakthrough applications with "Just-in-Time" delivery.

CSO Sites I.T. Infrastructure

Key Themes

- Focus on core competencies - investigate outsourcing other functions.
- Establish key metrics for services, costs, etc. Utilize TQC to review areas for improvement.
- Organize projects around key functions, not sites (common processes, leveraged development).
- Maintain strong communication links with ITC's corporate I.T., product groups, and development teams.
- Best practices sharing with customers and other HP groups.

CSO Sites Infrastructure – Highlights

- **5 Sites:**
  - *Cupertino, California*
  - *Roseville, California*
  - *Fort Collins, Colorado*
  - *Chelmsford, Massachusetts*
  - *Exeter, New Hampshire*
- **10,000 Customers**
  - *7,500 PC's*
  - *6,000 Workstations*
- **$40M Annual Budget**
- **200 Employees**

Client-Server Configurations

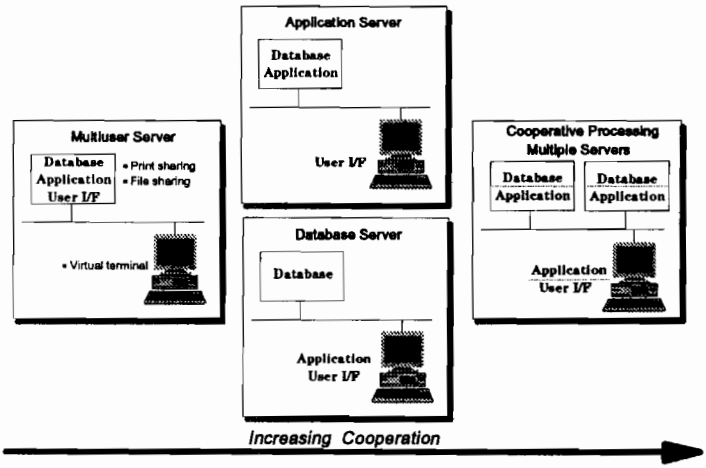

Client Server Architecture

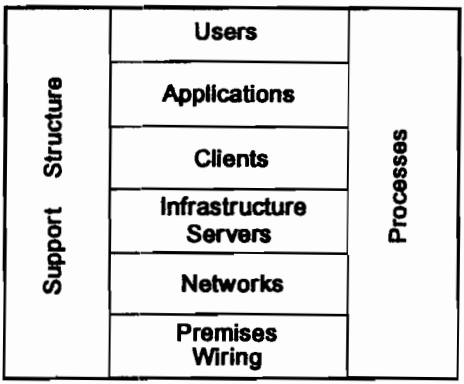

What is needed to implement Client-Server Applications

- **Reasonably powerful clients**
  - *DOS / Windows*
  - *UNIX*
- **High performance, reliable networks**
- **Reliable, accessible servers**

- The Problem with Today's Environment

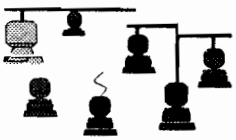

- Clients are fragile
- Sharing data is difficult
- Development of client applications is hampered
- Clients are expensive to support
- Clients are not meeting their potential

The Client-Server Dilemma

- **How can I justify this investment for any single application?**
- **What if each application requires a different client environment?**

- Common Client Operating Environment

- Connect all clients to the network
- Standardize applications and configurations
- Create and distribute generic clones
- Manage configurations via the network
- Move applications from the clients to servers
- Leverage software purchases

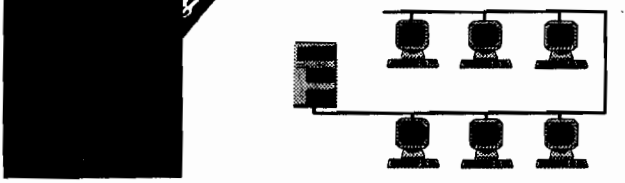

- Standardize the Applications

Recommendations:

- Word processing: **Lotus Ami Pro**
- Spreadsheet: **Lotus 1-2-3**
- Business graphics: **Lotus Freelance for Windows**

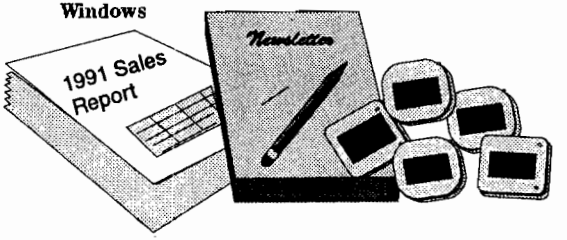

- Networked Products and Concurrent Licensing

- Client management savings
  - Only one copy of software needs to be installed, configured, and maintained on the server
  - Only the server needs to be audited
- Software license cost savings
  - Licenses need to be purchased only for the number of clients using the software at one time.
  - Licenses can be shared among several network clients

= Significant Cost Savings

- Networked Products and Concurrent Licensing

Potential Cost Savings with Concurrent Licensing: Lotus 1-2-3 for 240 users			
Individual Licenses on Clients which use 1-2-3		With Concurrent Licensing on the Server	
200 Licenses	$78,814.00	1 server license	$526.52
		39 node licenses	15,368.73
		Total	$15,895.25
80% Savings ($62,918.75)			

PC COE Mission

Manage an architecture to deliver and control core PC software and configurations, and to provide extensibility for geography's, businesses and application providers.

- What is PC COE?
  - Software components
  - Common configurations
  - Software management processes

PC COE Software Delivery

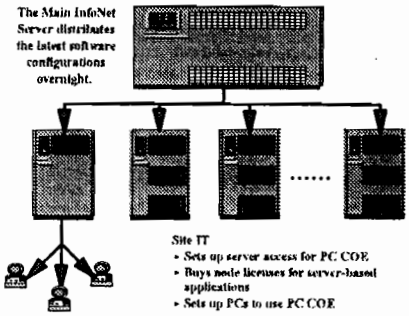

- What does PC COE do for users?
  - Delivers new software faster
  - Delivers a more stable, robust computing environment
  - Makes data easy to share company-wide
  - Reduces the need to purchase, install, configure and audit software
  - Reduces the cost of software

- What does PC COE do for IT?
  - Reduces engineering costs by providing standard software configurations
  - Reduces support costs by stabilizing the PC environment
  - Establishes hooks for automated PC management
  - Helps IT improve its service level to PC users
  - Provides a predictable target platform for application developers

PC COE Software Components

PC COE Software Components

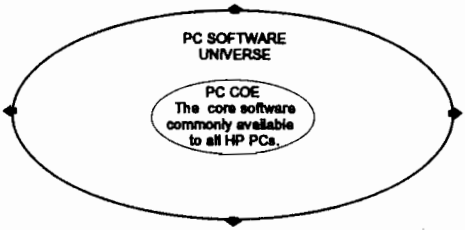

PC COE Software Components

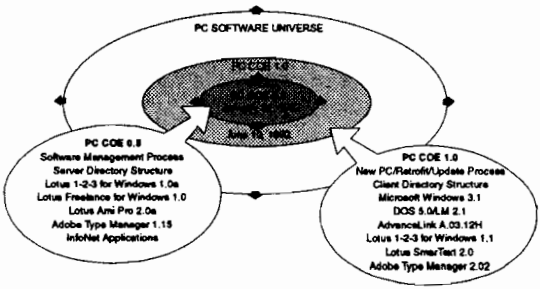

PC COE Software Components

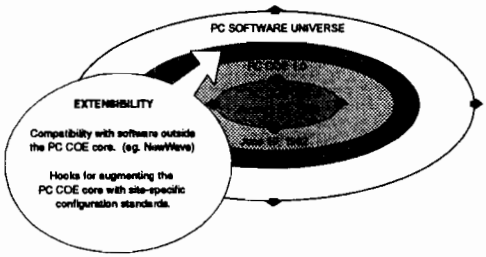

Total Calls Before and After PC COE by Group.

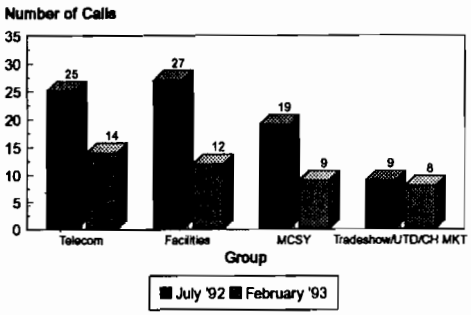

Hours Spent Before and After PC COE by Group

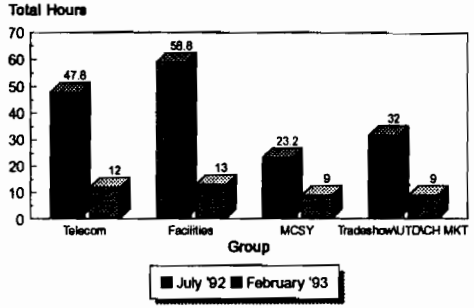

Total Hours Spent by Type of Call

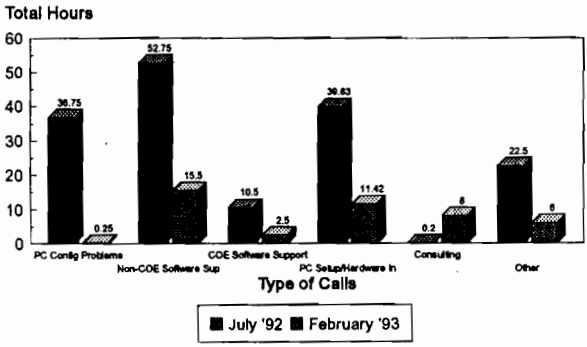

Support Requirements Over Time During Implementation

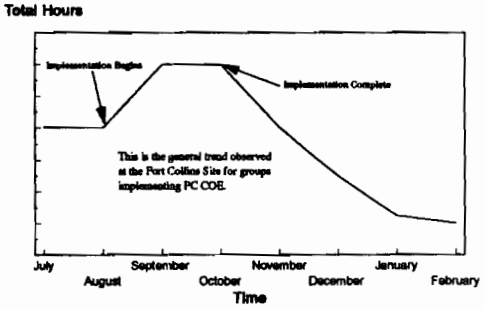

PC COE IT Issues

- System Reliability (Single point of failure)
- Hardware, Memory and Disk Space Requirements
- Resistance to 1-2-3 for Windows

Best Practices for Implementing PC COE

- Develop a comprehensive PC COE Implementation Plan
- Pilot PC COE in IT first
- Set Reliability and Performance Checkpoints
- Develop Support Processes
- Monitor the Clients, Networks, and Servers Continuously
- Deliver Just-In-Time Training
- Develop Service Agreements
- Develop Process for communicating Downtime

Strategies for Improving Overall Reliability

- **Measure the reliability of the overall system using the Client Incident Metric.**
- **Focus on improving the reliability of those components that will have the greatest impact.**
- **Measure the impact of downtime for each component.**
- **Communicate progress on overall systems reliability to your customers each month.**

PC COE Key Reliability Metrics

- **Client Incident Metric**
- **Server Incident Metric**
- **Server Downtime Metric**
- **Network Downtime Metric**

One Year Cost Savings
May '92 through April '93

- **$6.6 million saved in software licenses**
  (based on 18,000 PC COE users with an average 4.5:1 user: license ratio, an average Lotus suite cost of $300, 9,000 Ami Pro 2.0 to 3.0 upgrades, and 18,000 Freelance 1.0 to 2.0 upgrades at $100 per upgrade)
- ***$3.4 million saved in software update effort***
  (based on one hour/user savings for automatic distribution of 1-2-3W 1.1 to 3,5000 users, Ami Pro 3.0 to 9,000 users, Freelance 2.0 to 18,000 users respectively. A grand total of 67,000 automated software version updates; does not include patch distribution)
- **Measured reduction in support costs**

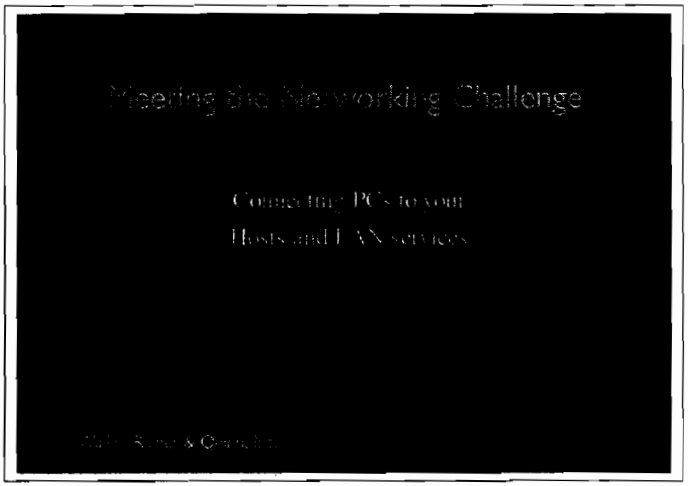
Meeting the Networking Challenge
Connecting PCs to your
Hosts and LAN services

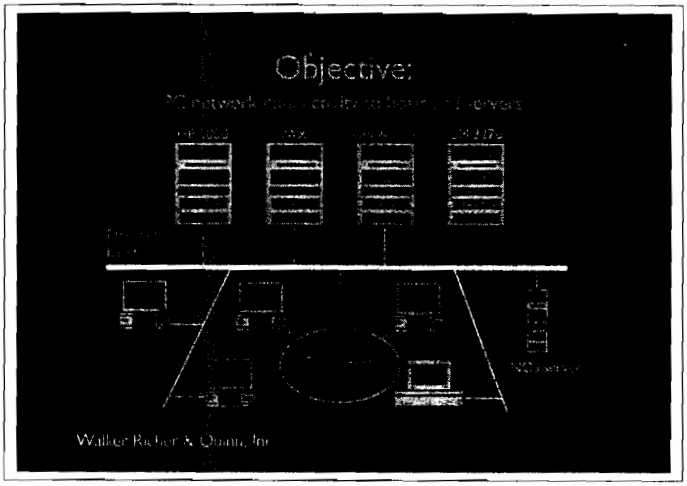

PC connectivity means providing users access to the resources they need on the network. Connectivity does not need to be limited to a single workgroup, specific site or even specific network topology. Company-wide access to computing resources can be provided through a single network connection from any PC or Macintosh.

The objective of this presentation is to examine PC-to-host and PC-to-server connectivity from the desktop perspective. We'll discuss

* The basic software and hardware components
* The industry standards that make PC networking possible
* Some of the issues and options to consider before you connect your PCs to other resources on your network

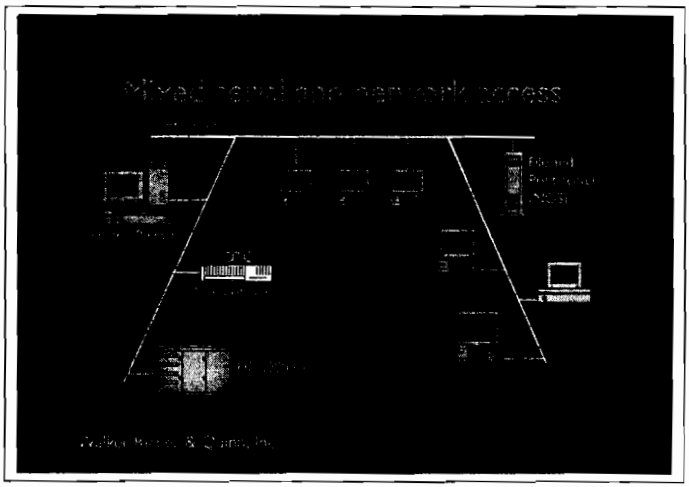

Most networks are installed as a means to share file and print resources among a small workgroup of PCs. The network grows around a dedicated PC or host that acts as the file and print server for the workgroup. The server runs specialized software known as a Network Operating System (NOS). Novell, LAN Manager and PATHWORKS are the most popular NOSs in use today.

Host connectivity is initially provided through a terminal server. In the HP world, a device called a DTC (Datacommunications and Terminal Controller) acts as a terminal server. Dumb terminals and PCs connect to the terminal server through a serial line. The terminal server is then connected to the network so it can communicate with the host.

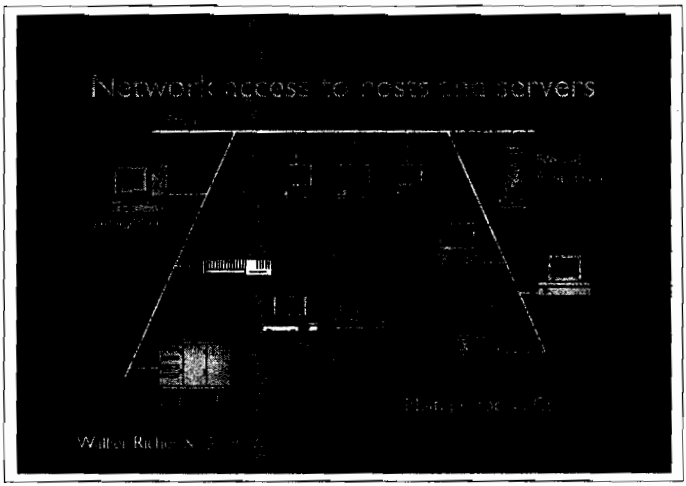

Serial communications require an RS-232 line to each PC or terminal. The serial line usually operates at 19.2K bps (19.2 thousand bits per second). A network connection usually operates at 10M bps (10 million bits per second), allowing an intelligent workstation like a PC to communicate more than 50 times faster than a serial connection. Network connections are also more reliable and less expensive to install and maintain than serial lines.

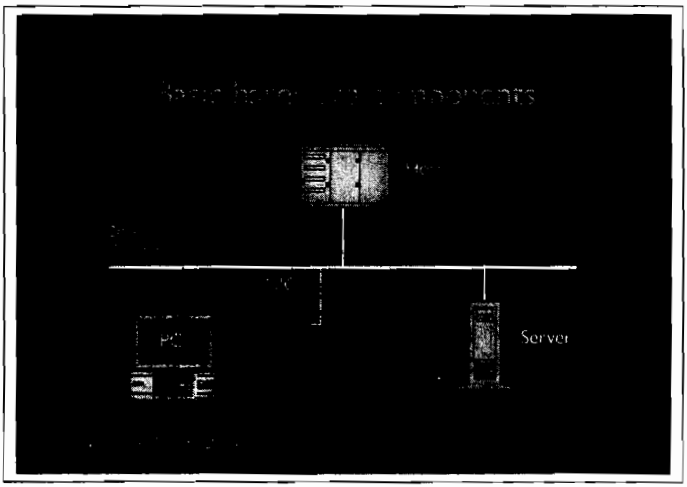

The basic hardware components of a network include the following: Ethernet or Token-Ring wiring is what connects the computers together by a physical cable. The Network Interface Card (NIC) is the Ethernet or Token-Ring attachment hardware to connect a device, such as a PC, to the network. The network node (or computer) can be a PC, host, file server, printer, and other devices on the network. The Network Operating System (NOS) such as Novell's NetWare and Microsoft LAN Manager is a central storage device for shared file access with print redirection services. Note that a file server is not required for network communication between nodes; it merely facilitates file and printer sharing. A network can operate without a NOS or file server.

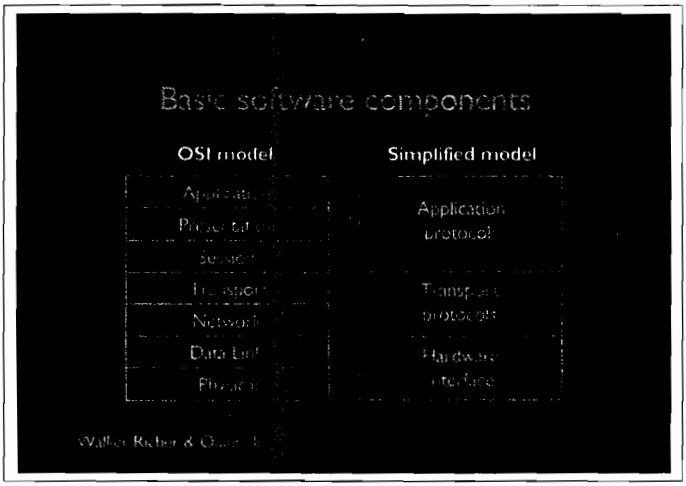

Network communication requires networking protocol software. A protocol is a well defined set of procedures and syntax for communication between two devices on a network. You can think of a protocol as a language. Both sides of the conversation must understand the protocol for two devices to communicate. In most cases, a protocol performs a single, specialized function. Two or more protocols work together to allow the PC to communicate with other devices on the network.

The most common method used to describe how protocols work together is called the OSI reference model. The Open Systems Interconnection (OSI) model uses seven different layers. Each layer performs a subset of the function required to communicate with another system. It relies on the next lower layer to perform more primitive functions. Each layer provides services to the next higher layer. For simplicity, the seven OSI layers can be grouped into three main network layers.

The higher layers are closer to the user, or what the user interacts with. Lower layers are closer to the hardware, where the PC is physically connected to the network.

The highest layer is called the application protocol level. These protocols present data in a usable format at the PC. Examples of application protocols include

* Virtual terminal protocols (Telnet, NS/VT) - establish terminal sessions over the network
* File transfer protocols (FTP) - transfer files over the network
* File sharing protocols (NFS) - file services
* X Window System - graphics and windowing protocol

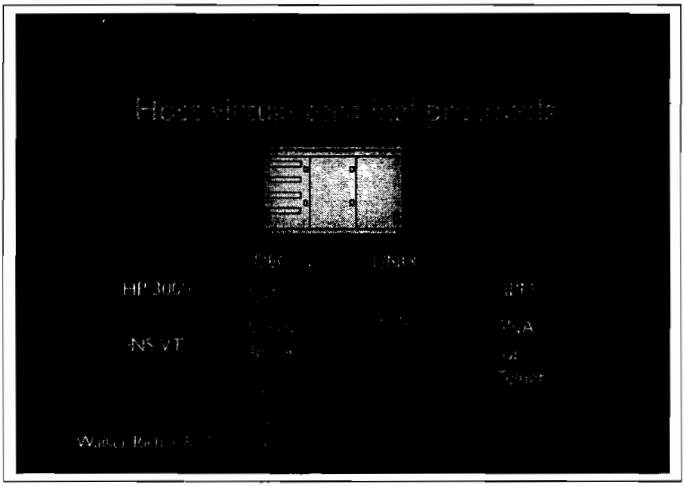

HP 3000: NS/VT (Network Services/Virtual terminal): HP's proprietary protocol for virtual terminal access to HP 3000s. NS/VT is designed to handle block-mode applications and transaction processing very efficiently on a network.

DEC VAX: LAT (Local Area Transport): Digital's proprietary protocol, designed specifically for efficient use within a local area network, but not compatible with Token Ring or through routers. LAT is a native protocol on the VAX; it is shipped with every VAX. Telnet-TCP/IP is available as an extra cost add on to the VAX.

UNIX: Telnet: Telnet and TCP/IP are part of the system software for almost all UNIX hosts. LAT is available for some UNIX hosts.

IBM: SNA (Systems Network Architecture): Most IBM mainframes use a networking model called SNA. SNA is not a networking protocol per se, but basically a description of how terminals, printers, controllers and other devices communicate on IBM's hierarchical networks. All communication passes through the mainframe. Recently, IBM has recognized that its customers want peer-to-peer (as opposed to hierarchical) access to its mainframes. IBM now develops and markets an entire line of Telnet-TCP/IP software for its mainframe and minicomputer operating systems.

Computer Museum

The HP 3000 runs TCP/IP as its network transport.

In the DEC environment, LAT is the standard, though TCP/IP and DECnet are offered as options.

UNIX hosts run TCP/IP as the standard transport, while LAT is an option on some UNIX hosts.

The IBM world has the SNA (Systems Network Architecture), though IBM is placing more emphasis on TCP/IP either on the host or through a gateway.

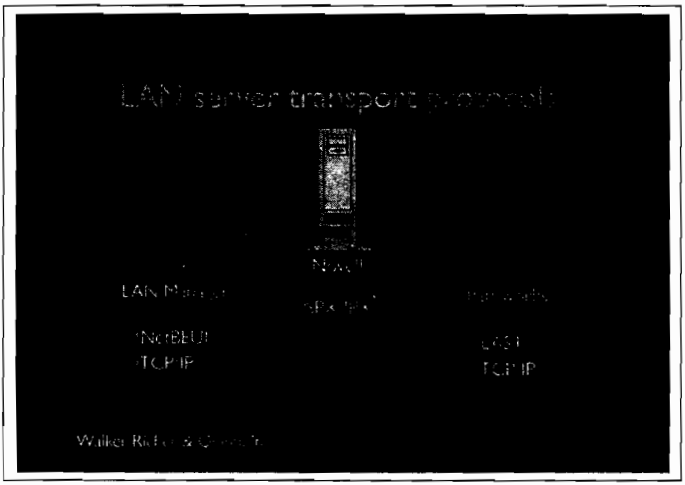

Novell uses its own transport protocols, SPX and IPX, which are roughly equivalent to TCP/IP...

LAN Manager is transport independent although most commonly it uses NetBEUI or TCP/IP through a NetBIOS interface . . .

Pathworks uses LAST (Local Area System Transport) or TCP/IP through a NetBIOS interface.

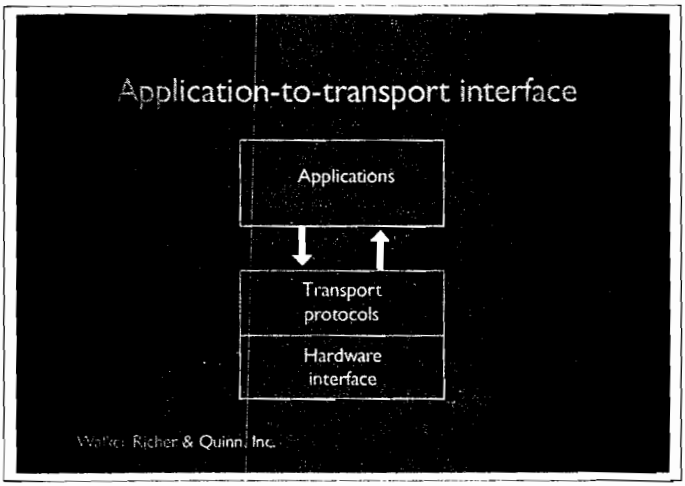

Applications and application protocols rely on lower level transport protocol software to deliver data across the network. Interaction between these two layers is facilitated by Application Programming Interfaces (APIs). Like protocols themselves, APIs can be based on industry standards or proprietary specifications. Let's look at some of the most common APIs in use today.

The term "Sockets" originated from UNIX and TCP/IP, developed at Berkeley University in the 70's. Sockets is now a generic term for an API to a TCP/IP protocol, regardless of the operating system (UNIX, DOS, Windows, Macintosh, etc.).

Sockets are the most common interface to a TCP/IP transport protocol. A socket interface is a type of Application Program Interface (API) between an application and a network transport layer. A socket defines programming points through which data is exchanged. Think of a socket as a mailbox. An application puts a message into the mailbox; the message is then picked up and transported by TCP/IP, which delivers it to a destination computer. A socket interface on the host side of the communication performs the same functions. Messages can also flow from TCP/IP to the application.

NetBios is another popular interface to TCP/IP. A NetBIOS interface to TCP/IP allows applications designed to use a NetBIOS transport (like NetBEUI) to run over TCP/IP instead.

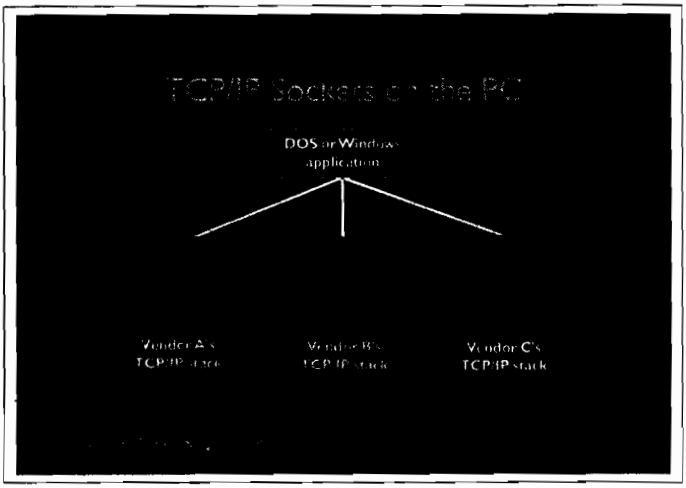

Sockets is the most common API for TCP/IP, but it is not consistently implemented on various operating systems. Sockets can cause compatibility problems with applications, especially on the DOS operating system. UNIX operating system functions cannot be directly duplicated by DOS, so each vendor of TCP/IP for DOS interpreted the sockets interface in their own way. The result was dozens of socket interfaces under DOS, forcing application developers in DOS and Windows to support each TCP/IP transport with a special interface on their application.

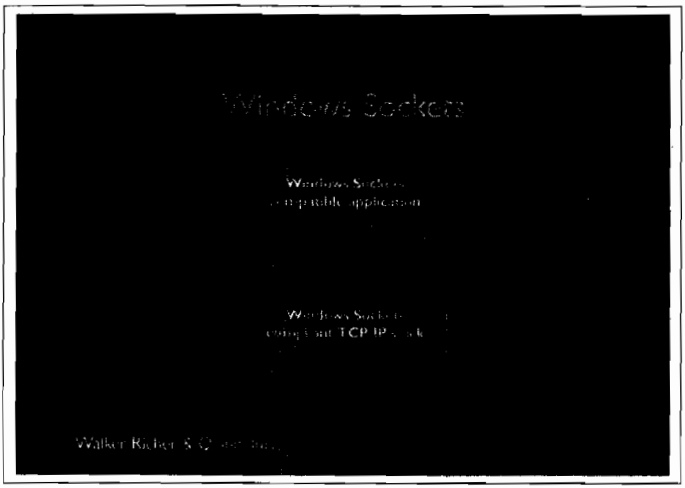

A commitee of over 40 companies formed in 1992 in an effort to standardize the sockets interface under Microsoft Windows. The new API is called Windows Sockets, a standard interface for Windows applications to communicate with the TCP/IP protocol. Windows Sockets allows application developers to design their applications for use with a single interface, allowing them to focus on the application itself instead of the networking aspects.

Windows Sockets does not solve all compatibility problems. Windows Sockets only works with applications that run in Windows. DOS applications must use a conventional sockets interface. In addition, Windows Sockets is a relatively new standard and has not been implemented consistently among TCP/IP developers. Applications should be tested carefully with any Windows Sockets interface.

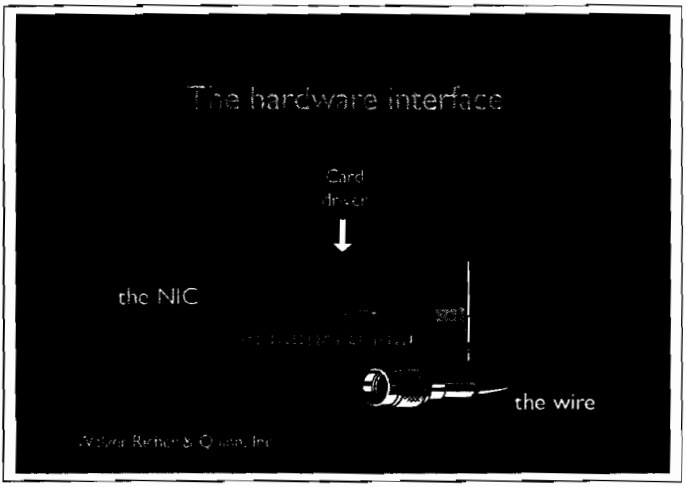

At the lowest level, software called the card driver passes and receives data to and from the transport protocol software. The card driver directs the hardware (the network interface card) while the hardware processes data coming from and going to the network media, or wire.

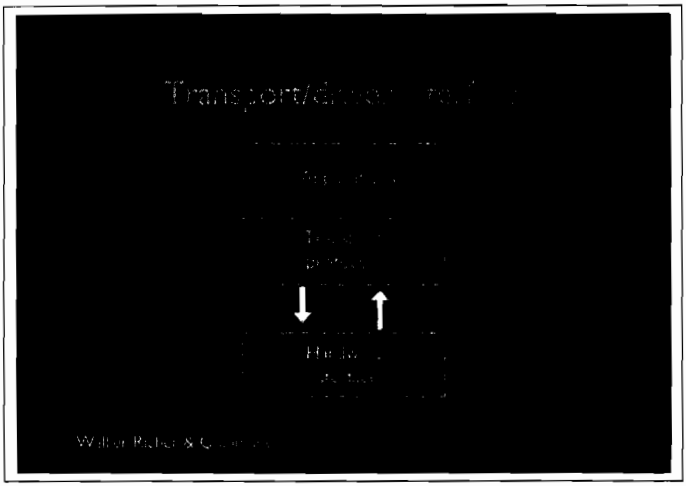

The interface between the transport protocol and the card driver has the potential for compatibility problems.

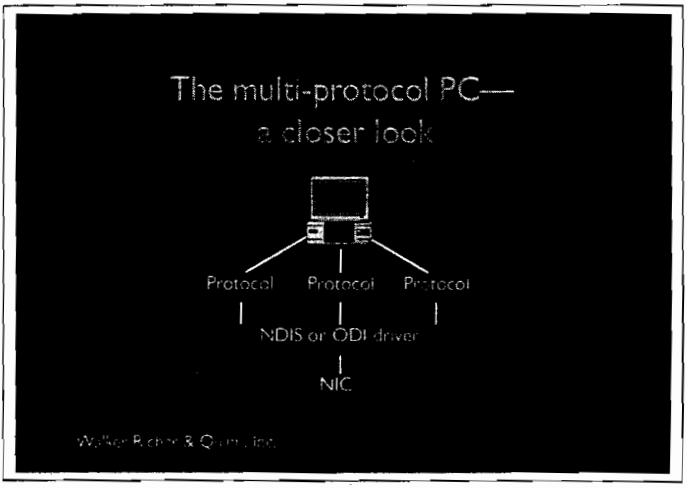

Originally, PC protocol software was written to work directly with the driver software for the LAN adapter card. The driver software was not unloadable from the PC's memory, and only worked with one protocol (hence the name 'direct driver'). If a user wanted to switch from a session on the LAN server to a host session, the PC needed to be rebooted to load the card driver for the host protocol type.

There is a better, more flexible way to handle multiple protocols on the PC. Two major industry standard card driver specifications are in use today. The Network Driver Interface Specification (NDIS) was designed by Novell and Microsoft. The Open Datalink Interface (ODI) was designed by Novell. Both standards are widely accepted and supported by the NIC manufacturers.

By using a multiple protocol management driver some distinct benefits come to the user:

* Multiple protocols can reside on the PC without interfering with each other. Hence, no more re-booting to load another protocol.
* With a standard interface to the network card driver, the LAN card manufacturers can write only one driver to support different protocol types.

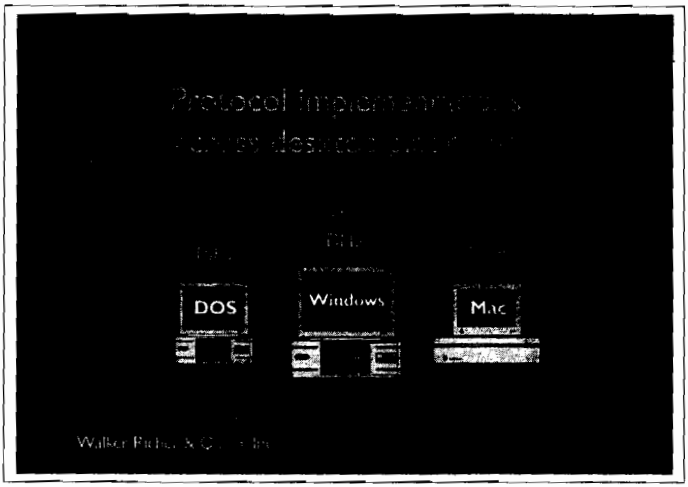

If your PC is running DOS you will find the networking protocols implemented as TSRs.

If your PC is running Windows, see if a Windows-based version of the protocol is available. Windows uses programs called DLLs (Dynamic Link Libraries), which use memory above 1MB. A number of DLL based TCP/IP stacks are available.

VXDs (Virtual Device Drivers) are also being used for networking protocols in the Windows environment. These VXDs operate at a lower level than DLLs and provide DOS box applications access to the networking protocols without having the protocols loaded as DOS TSRs.

In the world of the Apple Macintosh you will find network protocols implemented as Apple Communications Toolbox Tools.

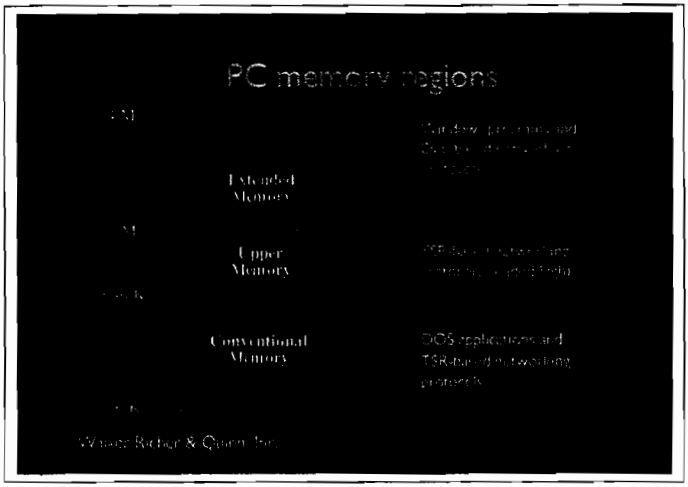

PC programs typically use one or more of three types of memory. Let's take an IBM-compatible 386 with 4MB of RAM as an example.

Conventional memory: Most DOS applications run in this area of memory between 0 and 640K. Networking protocols are usually terminate-and-stay-resident programs (TSRs). A TSR is loaded into memory and remains loaded, even when it is not handling networking tasks. TSR-based networking protocols must share conventional memory with other DOS applications.

Upper Memory: This area between 640K and 1MB is typically used for the PC's hardware, such as the video display driver program. Some of upper memory is unused, and this memory can be reclaimed for use by networking protocols. Memory management programs or DOS 5/6 MemMaker utilities can be used to move TSRs into upper memory, freeing conventional memory for other programs.

Extended Memory: The area above 1MB is typically used by Windows and Windows applications. Networking protocols that can run in this area have recently been developed, freeing the maximum amount of memory for other applications.

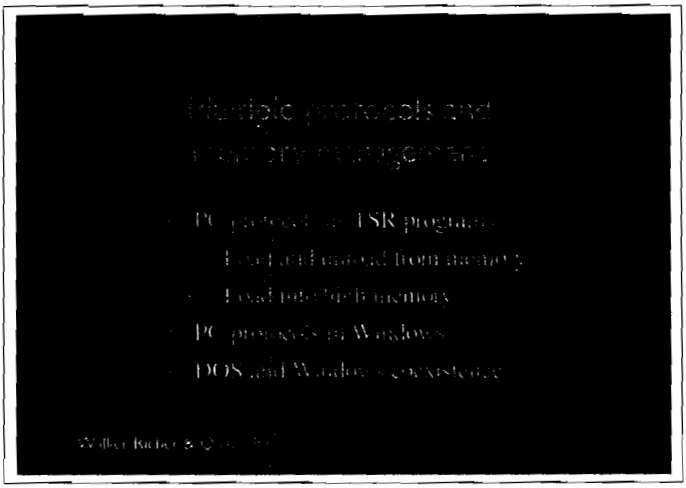

Memory management should be considered when choosing a multi-protocol PC architecture. Most communications protocols today are implemented as TSRs. Windows-based DLLs are emerging as the next generation of networking software.
Here is what you should consider as you transition to DLL-based protocols:

1) DOS applications cannot access Windows DLL-based networking protocols.

2) Windows applications can access both TSR- and DLL-based networking protocols.

3) Setup, memory management, and application compatibility are likely to be simplified by DLL-based protocols and Windows interfaces.

Suggestions to reduce memory constraints for TSRs:

1) Make sure the TSRs are compact and unloadable. When a protocol is not in use, the TSR can be unloaded at the DOS prompt to free up memory.

2) Make sure the TSRs are designed to be loaded into upper memory blocks (UMBs).

Using a multi-protocol PC you can have simultaneous access to HP 3000, UNIX, X clients, IBM 3270, and VAX hosts, along with your LAN server. In fact, these are actual bit-maps of screens showing connections made to these hosts through a single Ethernet card in a PC.

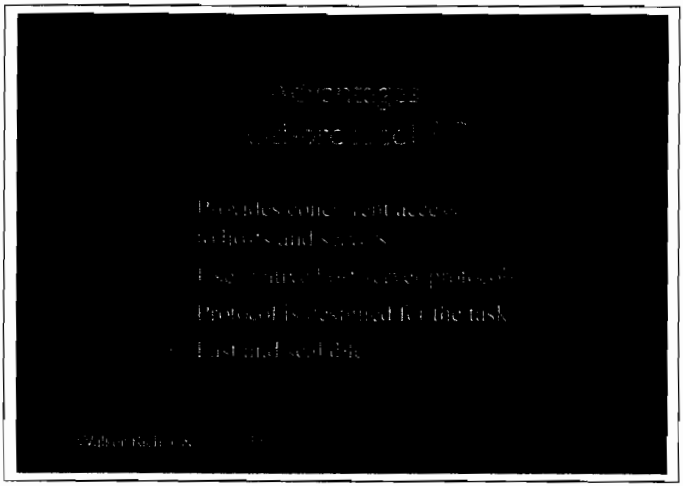

Advantages of the multiple protocol approach are straightforward. The primary advantage is access to any host or server on the network using its native protocol. Since the protocols are designed for the applications and hardware you are using, your system is more efficient. Native protocols provide the fastest throughput possible. By using native protocols, scalability is limited only by the capacity of the host, not an intermediate gateway device. New users can provided access to resources on the network simply by adding software to the PC.

Client/Server Headache: Performance How to get Relief

by Dan Gillis
Manager Advanced Technologies
Dynamic Information Systems Corporation
(303) 444-4000

Client/Server Headache: Performance How to get Relief

Introduction

Like all new technologies, client/server is progressing through its "life cycle" right on schedule. First heralded in the early 1990s as a solution to all end user data access needs, it has since begun to mature as it is actually implemented at organizations worldwide. With that maturation, as in life, brings a growing awareness of both its strengths, and its potential weaknesses as a data access solution.

In my work as a consultant to organizations looking to implement technology solutions to improve productivity, the single biggest weakness I've encountered with client/server is performance. Companies invest tremendous resources to create new avenues of data access--attractive PC screens, point-and-click GUIs, intricate network connections--only to find this new-found access is often frustratingly slow.

This paper will briefly discuss the reasons that client/server performance becomes an issue in operations, and will outline the various strategies I've found to be most helpful. I've included three case studies to illustrate how these solutions worked in real-world situations.

Great Client/server Design! What Could go Wrong?

When we first began to understand the concepts behind client/server, the possibilities for its application seemed endless. The executive suite could finally have those summary numbers--in a colorful chart no less--any time they wanted to see them. Our sales reps could have instant access to product availability and pricing information--maybe even from the road! Our customer service people could have access to any corporate information they needed to solve a customer problem. Marketing could have any report they wanted--without having to go through MIS every time they wanted a new set of numbers.

But as we've begun to implement these systems, our expectations have inevitably dropped. The front end screens were well designed, and certainly made navigating the computer system and report preparation easy. And even with the endless stream of three letter LANs and communication protocols, we could set up a pretty straightforward connection to get all these PCs talking to the host systems. But it doesn't seem to work the way we envisioned it.

The very freedom of access that we've given our users means they can kick off horrendous queries that slow everyone on the system. The problem only multiplies with more and more users. This problem is so pervasive, that I've dubbed it the "Client/Server Exercise Program." It works as follows:

CLIENT/SERVER EXERCISE PROGRAM

ACTION		REACTION
1.	CLICK ON A BUTTON	GET COFFEE
2.	DO SOME WORK	
3.	CLICK ANOTHER BUTTON IN CASE THE FIRST REALLY WASN'T WORKING	GET COFFEE
4.	DO SOME MORE WORK	
5.	CLICK IT AGAIN	CALL HOME
6.	START IN ON YOUR IN-BASKET	
7.	MUMBLE ABOUT THE PC AND CLICK IT AGAIN	BATHROOM!

Richard Simmons couldn't have designed a more strenuous physical exercise program. Every time you click a button, you get up and exercise.

Same Old Problem with a New Name: I/O Bottleneck

What's going on? Analysts have called it network traffic overload. And data transfer lag time. Even disk contention. But it boils down to a term we all discarded in the 80s: I/O bottleneck. Processes will only run as fast as the slowest portion of your computing process.

Let's take a look for a moment at how database retrieval systems work.

Data is typically stored in a database or file structure. This structure defines the

relationships between all the data entered, and allows access to that data by a variety of methods. The speed of retrieving that data is most heavily influenced by where that data is stored.

If the data being searched is located in the systems *main memory*, retrieval will be very efficient. All the computing resources are spent on locating the data.

As we all know, computer systems are forced to switch data back and forth between main memory and disk to manage efficiently. If, however, the system has *cached* the data you're looking for into "temporary storage," this retrieval method is reasonably fast.

Next in line for retrieval speed would be data that's stored in *network memory*.

But, most often, we're looking for data that is stored in *disk*. 32 I/O's per second is the current physical limit of disk transfer speeds. When searching through large files, this physical transfer speed limit presents a real bottleneck for performance. While the latest advances in RISC chip design have given us better relative I/O by using pre-fetching, on-line retrievals are often hampered by the sheer size of the database that must be searched to find the specific records of interest.

These physical searching limitations are the culprits to our client/server designs gone awry. The most beautiful designs won't overcome slow system response to queries. So what are some options?

Performance Improvement Options

There are a variety of options that have been tried over the past several years. They range from what I call "limitation" solutions to hardware and software options.

Hardware Solutions

The one solution that most companies choose first is bigger hardware. Faster network, faster drives, more memory to cache and easily justified. But it's one of the most costly solutions. And may not speed up the kind of ad hoc queries you're trying to provide.

Hardware solutions range from:

- SCSI drives
- 100 Mb network
- Faster CPU
- More memory
- Fiber.optics
- More cache memory

Limitation Solutions

If purchasing additional hardware is too extravagant, many companies look to limit the activities that impact performance the most severely. I'll list the most common limitations that I've found:

- Break down large files into multiple smaller files to speed searches
- Filtering or restricting access to corporate data during peak operating hours; running against the larger files in batch
- Activity timer aborts query after set length of time
- Pre-selections done in batch at night for most popular queries
- Roll-up summaries prepared in batch for EIS systems and accounting
- Transaction log kept on PC during peak hours, uploaded to server for batch updating in off-hours

Unfortunately, these limitations, intended to make the system work more smoothly, often have the effect of frustrating the very end user the system has been designed to support. I've even seen the recommendation by a very respected client/server vendor, that no more than 20 - 30 users should be hooked up to client/server systems at any one time. Why? Not because that vendor's inherent performance is so slow, but because of the physical limitations of processing ad hoc queries and passing the results down to the client.

There are better ways.

Software Solutions

I have an innate predisposition to finding software solutions to performance problems. There are several available that I'd like to briefly cover. Some, I think are better than others for client/server computing. But you'll need to evaluate your own organization environment to select the appropriate mix for your individual situation.

Data Warehousing is an increasingly popular alternative to companies with large amounts of data to process and provide access. A Data Warehouse is basically a revamped historical database with a new name. By off-loading your on-line data

periodically into a separate data structure and providing indexed access, you can leave the detail in, and roll up the summarized data if necessary, without negatively impacting the performance of your production systems.

But the Data Warehouse has several drawbacks for client/server performance you should be aware of:

- The data is not on-line
- The warehouse can be complex to set up and manage
- Without sophisticated indexing, the warehouse will not significantly improve retrieval performance
- Updating and indexing must be managed

More basic still is the selection of the *Database System.* Database system vendors spend millions each year telling us that their particular solution is the fastest performing platform. Benchmark after benchmark attest to each RDBMS's performance rating. To do so, most relational database systems rely on some form of query optimization, and in some cases, SMP and MPP processing.

Query optimization, also called pre-optimization, is essentially a search optimizer. This optimizer can evaluate a query before it is processed to schedule the search, sort and merge steps in the most efficient manner. In a query with multiple selection criteria, the optimizer will attempt to use the table with the smallest number of entries to search first to narrow the search to a smaller subset before reading other files for the appropriate "matches."

While pre-optimization is a giant leap forward, it can be of little help in queries where the qualifications are coming from multiple tables with a large number of entries. Similarly, multiple table joins are still uncomfortably slow.

The latest releases of Oracle, Sybase and Informix have some level of support for parallelizing queries. These can give you substantial performance gain, but it's not for the faint of heart. Most Symmetrical Multiprocessing or Parallel Processing Machines are hundreds of thousands of dollars, and can require months of tuning to give you the performance gains you need. While the database vendors have been wise to support these machine types, it's really nothing more than throwing more, smarter hardware at the problem.

My personal "favorite" solution is *Inverted File Indexing.* Inverted file indexing is the "dark horse" candidate for improving query performance. While more complex to initially implement, I've found that it can be far more important than the database or

file structure you use in determining client/server performance.

Inverted file indexes offer three strategic advantages in client/server performance:

- Dramatically improves server performance for multiple selection criteria and joined retrievals
- Offers a count of the number of qualifying records *before* the retrieval is initiated and downloaded to the client
- Ability to implement summary indexes to speed ad hoc summary reporting and analysis

Inverted file indexes allow you to instantly determine the number of records that qualify in a query--a query that can have multiple selection criteria, even across multiple tables. I've seen systems speeded up by a factor of 10 or more with the addition of these indexes. This performance improvement is even more critical in client/server applications.

I'll cover inverted file indexing in a bit more detail at the summary of this presentation in the case studies we'll discuss.

Database Design

The final area of performance tuning I'd like to cover is the basics of database design. I sometimes like to call this section the "Database Blues," because so many of us have tried some great new design technique, only to find that it ruins performance.

Normalization is one such example. Normalization can provide a number of benefits, but increasing performance is not one of them. One general rule of thumb I give to clients is to normalize in stages and test after each one to make sure that your performance is degrading beyond tolerable limits. I've yet to see a high-performance, functional database normalized over 4 levels.

Table joins is a key area that impacts retrieval performance, if you're using those joins to direct the path of the query. It will come as no surprise that the more joins you have in a query, the slower it will be. The dilemma for client/server designers is that these very joins are the ones the users need most: "How many products were sold last month to customers in the Northeast that were priced at the promotional discount?" A perfectly reasonable business question. A nightmare for the DBA who sees the "spaghetti-joins" necessary to select that data from the 3 - 4 different tables.

One other database design point I'd like to bring up is SQL. The Standard Query Language often isn't. The SQL implemented by one database vendor is usually not portable to another file structure or database. If you're evaluating different databases for performance gains and are sold on their platform independence, talk to some of their customer sites to see if that is indeed true. I've found that most often, it's not.

Client/Server Case Studies

I'd like to share a few personal experiences in implementing client/server solutions, in hopes that you can glean some "real life" options for your own situation. I've selected three clients that I've worked with over the past year:

Boots Pharmaceutical	EIS application
Betz Industrial	MSDS Documentation application
Midmark	Automated sales force application

Boots Pharmaceutical

Boots Pharmaceutical is a pharmaceutical development and manufacturing company headquartered in the Chicago area with offices in the UK. They began looking at client/server as a way to give faster, more flexible access to daily and weekly sales data to management so inventories could be managed more effectively. In-store sales data was being collected daily across the country, but there was no effective way to summarzie that data and make it quickly available to the marketing and finance people in headquarters.

Their first efforts to install a client/server system pointed up the need for a way to process the large volume of data quickly, and allow managers to drill down for any information they wanted to pursue. They chose to add inverted file indexing--OMNIDEX--to their corporate database. They created Visual Basic™ screens for the managers to use to point-and-click to the information they wanted to see. These Visual Basic screens used a client/server API that called the OMNIDEX indexes on the server to first provide a qualification count. Any time the qualification was modified, a new qualification count was immediately sent to the PC screen. Once the user was satisfied that they'd identified the records they wanted to see, Boots used specialized summary indexes to quickly summarize the large amount of data. The program then downloaded the totals into the client program where they could be graphed or put into a report.

The stunning statistic that I like to quote when I talk to other clients, is that this whole project took less than three weeks to prototype and initially implement. Boots has

since been refining the system, adding new query screens and update mechanisms, but they had a working client/server system that performed in seconds off of millions of records in less than three weeks.

Betz Industries

Betz is a worldwide chemical manufacturer. Anyone in the chemical or related industry knows that these companies must maintain copious records of environmental data on any chemical they produce or handle. Mandated as MSDS, or Material Safety Data Sheet, a chemical company such as Betz can have literally thousands of documents on the various chemicals they handle.

Once produced, these MSDS documents are the primary reference used by both the company and the EPA to identify chemicals with similar properties. For example, if a chemical spill occurs, the EPA may need to quickly identify the contents. Descriptions of the spill are cross-matched against the MSDS records to look for candidate chemicals that could be involved. Betz wanted a more efficient way to access the volume of records and to have the appropriate information immediately available to the user.

Betz set up a client/server system, again using Visual Basic screens and the OMNIDEX inverted files, that allowed them to use full keyword retrieval and multiple key searches on their multi-million record database. But speed of retrieval wasn't their only concern. Once the appropriate data was located, the amount of information that needed to be downloaded to the client slowed the process dramatically. We recommended switching from the commonly-used messaging download to a file transfer. We saw dramatic performance increases, from a typical download requiring 90 seconds or more to 2 seconds.

Midmark

Midmark is a medical equipment manufacturer that wanted to automate their remote sales force. The goal was to be able to download the corporate information they needed in their sales situations--product inventory, pricing, customer credit history, etc.--into individual sales PC, and upload new order information from the clients into the server on a periodic basis.

With Visual Basic front-end screens, we helped them design a system that allowed the sales people to upload and download information between the corporate and client

systems quickly and easily, either remotely via modem, or while they were at the home office through either a direct connect or a serial connection to the server. We installed the inverted file indexes on the server database to provide them quick access to the server information they needed for their sales calls. Visual Basic screens then let them navigate through that information on their PCs while on the road.

Go for it!

Those of us in the industry know that one of the surest ways for job security is to be responsible for improving system performance--you can never have enough.

I strongly recommend inverted file indexes to provide multiple keyword retrieval for large databases. I recommend looking at alternative transfer technologies such as file transfers for files over 6K bytes. And, as always, you should give your users the benefit of your knowledge in what kind of queries can and should be done in a client/server system. I've found these methods to be enormously helpful in boosting the client/server headache: performance.

Above all, I encourage you in your efforts to develop efficient client/server systems. The productivity--and yes, performance--gains are well worth the effort.

If you have any questions, or would like to discuss any of these recommendations further, feel free to call me at DISC at (303) 444-4000.

Paper Number 4006

Visual Basic for the MIS Profession: How to Develop High Performance Client-Server Applications

Terry O'Brien
Executive Vice President
DISC
5733 Central Avenue
Boulder, CO 80301
303.444.4000

Handouts will be provided at time of presentation

Paper Number: 4007

Do Your IMAGE Databases Need Therapy?

Ken Paul

Adager

Sun Valley, Idaho
83353-0030
U.S.A.

Telephone +1 (208) 726-9100 Fax +1 (208) 726-8191

Abstract

Most people immediately respond "NO" to this question. We often hear that the first step to recovery is admitting that you have a problem. Unfortunately, databases are like people when it comes to admitting that they have a problem. Databases are not known for volunteering any problems they may have and if they do, they don't provide the user with enough information on the source of the problems.

Section I discusses different ways that databases are able to tell their users that the structural integrity of the database may not be sound. Section II then illustrates various techniques used by diagnostic tools to discover if IMAGE databases have any structural problems. Next, Section III describes several diagnostic tools that can be used to diagnose databases. Section IV covers various ways that users can prevent the occurrence of broken chains. The paper concludes with a strong recommendation that database administrators (DBAs) perform regular diagnostics on their databases instead of waiting until problems occur. (Murphy's Law states that database problems will occur in your largest dataset in your largest database during month-end or year-end processing.)

NOTE: In this paper "IMAGE" is used generically to mean all flavors of IMAGE which include: IMAGE/3000, TurboIMAGE/V, TurboIMAGE/XL and IMAGE/SQL.

Introduction: IMAGE Reliability

As most IMAGE users are aware, IMAGE is a very efficient and reliable database management system (DBMS). IMAGE offers additional reliability and integrity through options such as Intrinsic Level Recovery (ILR) under MPE/V, dynamic rollback under MPE/iX and transaction logging. Under MPE/iX, IMAGE is also integrated with the Transaction Manager (XM) to insure database integrity.

As a result, most users have never encountered a broken chain or a damaged IMAGE database. This does not necessarily mean that their databases are 100% structurally sound. Only after running a complete diagnostic on the database can users be assured of its integrity.

Disc drive problems, system failures and operator errors are ways IMAGE databases can become damaged. The most common cause for database corruption, however, can usually be traced to human error. Accidental or deliberate partial restores of a database are a main cause of database structural problems.

NOTE: Although an IMAGE database consists of multiple files, it should be treated as a single entity because physical and logical links exist between these files.

Section I: Warning Signs and Symptoms

Introduction

Most users become aware of problems with IMAGE databases when an application program returns an IMAGE error condition or aborts. Users may realize problems exist when they can't retrieve data via a chained read but they can retrieve data via a serial read. System operators may notice hardware error messages on their consoles, nightly backups not able to store whole databases or restores not able to restore all of the database files. These are warning signs that there may be problems and should not be taken lightly.

"Phantom" Broken Chains

"Phantom" broken chains are defined as: Broken chains which are there one moment and gone the next.

Some software does not lock the database correctly (or at all) before it accesses data. If different applications are accessing a database and one is locking the database and another is not, IMAGE may report broken chains because one application was deleting entries from a chain that the other was reading.

My first rule of thumb is: If problems can be reproduced consistently, the database integrity is questionable. It is very important to reproduce problems before spending hours or days searching for structural problems that don't exist.

Logical Integrity

I have received calls from customers asking if they have broken chains because their application aborted with a DBGET Mode 7 failure on a master dataset. Usually these cases deal with logical integrity problems versus structural integrity problems. Let me explain.

Let's say we have two databases, one contains information about all of our parts and the suppliers who provide them and the other database contains our accounts payable information. The accounts payable database contains a supplier code within each entry and the supplier master dataset is in the other database. If a supplier master entry does not exist for each supplier code found in the accounts payable database we have a logical integrity problem.

Currently there are no physical links that exist between two IMAGE databases. In order to fix this problem, the user must insure that every supplier mentioned in the accounts payable database exists in the supplier master dataset of the other database.

IMAGE Error Messages

You may have seen one of the following IMAGE messages:

DBG Disabled; Potential Damage; Only DBCLOSE Allowed
DBB Disabled; Potential Damage; Only DBCLOSE Allowed
DBU Disabled; Potential Damage; Only DBCLOSE Allowed

These messages usually indicate that a broken chain was encountered during an update and IMAGE has set a flag in the DBG, DBB, or DBU control block preventing further updates to the database. When one of these messages appears, the database probably has a broken chain and one of the following must be done:

Option #1
Shutdown the database immediately and begin looking for the broken chain(s). Unfortunately, the DBA may have no idea where to begin. Usually the user who encountered the broken chain was aborted from the application by IMAGE. If this user can be found, the DBA may be able to better pinpoint the location of the broken chain(s). Otherwise, the DBA may be left with no choice but to run a complete diagnostic on the whole database.

Option #2
If it is a critical database and the site can not afford to be down for any period of time, all users must logoff the database and log back on. This rebuilds the DBG, DBB and DBU, essentially clearing the "broken chain" flag. If users attempt to do the same functions when they get back into the database they may encounter the broken chain again and the logoff/logon process needs to be repeated. If users avoid the broken chain, the database can remain active until there is time to do a normal shutdown of the database and run a diagnostic.

Further Damage?

You may be wondering if IMAGE databases can be further damaged by allowing users to update a potentially damaged database as mentioned in Option #2. The IMAGE intrinsic DBPUT detects pointer damage within a chain, but not key value inconsistencies. Therefore, if users attempt to add entries to a chain with pointer damage, the DBPUT does not complete. On the other hand, if there are chains currently in the database that have corrupted keys or are intermingled with other chains, IMAGE does not detect these problems and keeps adding entries to these chains.

Broken Chain - Forward and Backward Pointers Not Consistent

This is the message IMAGE may return if you have a "Phantom" broken chain as discussed above. If you can reproduce this message consistently, then I would say you definitely have a broken chain. Hopefully you know which dataset(s) are being accessed so you can start examining them for possible problems.

"But, this application has been working for years and it has not been changed and has never had a problem!"

This is a common response after a user has run two diagnostics on the database and both have agreed that the database is perfectly sound. If two different tools agree that the database has no problems then we must conclude that there is something wrong with the program that causes it to consistently report a broken chain.

Has the database structure changed recently which may have forced an internal buffer to overflow? Have there been any major changes in the decade, year, or month which have caused the program to act differently from its normal behavior?

Database Bad - Output Deferred, May Not Be Accessed In Mode n
Database Cannot Be Opened Because It Is Damaged

Either of these messages may be displayed by IMAGE or Query. These messages have one of two meanings: Either the database was open and enabled for Autodefer when there was a system failure, or a database utility such as Adager was interrupted before it could finish all of its structural changes to the database. Whatever the cause, this condition is very serious. By resetting a flag in the database root file, users are able to get back into the database. It is very important, however, to find out why the database was flagged "Output Deferred". It is also important to diagnose the database for any possible structural damage.

Check to make sure the database is not enabled for Autodefer by doing a *SHOW <DBNAME> FLAGS* within DBUTIL. If the database is enabled for Autodefer then it should also be enabled for Logging. If it is not enabled for Logging, database damage is almost certain in the event of a system failure.

Section II: Diagnostic Techniques

Speed vs. Thoroughness

These are tradeoffs when dealing with the various techniques that are used by diagnostic tools. Like most things in life, you can't have your cake and eat it too. In general, increasing the speed of a diagnostic comes at the expense of its thoroughness and, conversely, increasing the thoroughness of a diagnostic comes at the expense of its speed. The fastest tool may tell you that a database has no problems while a slower diagnostic may tell you that a large detail dataset has a problem in one of its paths; an even slower diagnostic may tell you exactly which record has a problem. Knowing the thoroughness of a diagnostic tool is very important.

How Long Will It Take?

This is the most common question that I am asked by DBAs. Unfortunately, there are too many variables in determining an answer to this question. Some of the variables involved are the diagnostic tool used, the options used, the method that the diagnostic tool uses, the size of the database, the number of users on the system, the size of the system, the thoroughness of the diagnostic and so on.

Methods To The Madness

There are basically three methods used by diagnostic tools to examine a database or dataset for structural problems. Some tools may use one or more of these methods while doing the diagnostic. Each of the methods has its good points and its bad points:

Method #1: Serial Scan

This method involves making a serial pass of a dataset to detect problems within the dataset. This method's good point is that it is the fastest available. The bad points are that this method can not detect all the problems that may exist in a dataset and it can't tell exactly where some of the problems it does detect exist.

The serial scan process uses a "checksum", a mathematical algorithm, to determine if a detail dataset's path is okay. Unfortunately, problems that exist between the chain heads in the master and the related entries in the detail go undetected. Also, this "checksum" method can only tell if a path has problems but it can not tell which entry/entries in that set is/are at fault.

The reason for using the serial scan is its speed. As long as you are aware and don't mind that it may not find all the problems but that it does find most problems in the shortest amount of time, this is the method for you. Severe damage to a dataset caused by a hardware failure is always found with a serial scan. As stated above, inconsistencies between master and detail datasets can not be found.

Method #2: Chain Chasing

This method does a chained read from every entry in a master dataset to its associated detail dataset and checks every aspect of the chain, including things like the chain head count, the forward and backward pointers, and the key values. This method's good point is that it is very thorough. A bad point is that it usually is very slow because it is so I/O intensive. Another bad point is that it may also discover a problem along a detail path and not be able to tell exactly where that problem is. This happens when there are "orphan" detail entries that are not pointed to by any other entries.

The chain chasing method is usually used in conjunction with the serial scan method so that the number of entries read by each can be used as a cross check. The larger the detail dataset and the more paths it contains, the longer the chain chasing method takes.

Method #3: Combination

This method performs a serial scan of a dataset and also extracts information to use for additional analysis. The extracted information is then sorted into a particular order so that chains can be checked without having to incur all of the I/O overhead. This method's good points are that it is faster and more thorough than the chain chasing method. It is also more thorough than the serial scan since problems between a chain head in the master dataset and the chain members in the detail dataset are found. The combination method is also able to detect that "orphan" entries exist and pinpoint exactly where they are located. The bad points to this method are the need for disk space for the sorting portion of the process and it is not as fast as the serial scan method.

Section III: Diagnostic Tools

Which Tools Do You Have?

Having discussed the various methods that diagnostic tools can use to diagnose a database, it is time to look at the actual tools which may exist on your system. There are several tools available for diagnosing IMAGE databases. Which tool you use depends upon your particular situation, how well you know the database in question and how comfortable you feel using a particular tool.

DIOGENES From HP

One tool that may exist on your HP3000 is called DIOGENES. It is a tool provided by HP and is usually found in the TELESUP account. It may be located in the PRV group or the PRVXL group. Unfortunately, when Spectrum systems were introduced, this tool was not provided with each system. You may need to call the Response Center and have them download it to your system.

Like all HP tools, this one also contains a very scary disclaimer when you run it. DIOGENES accesses database files in read-only mode. It is very safe and, to the best of my knowledge, has never caused a system failure, but HP has to put their disclaimer up first just to be safe. Also, don't be concerned when it says that it is an MPE/V support tool; it works just fine under MPE/XL and MPE/iX.

DIOGENES works in conjunction with a message file called DIOGMSG. DIOGENES creates a command file that it uses while it is diagnosing a database. The user interface to DIOGENES can be somewhat confusing. The user can request any of the three methods that were discussed earlier. DIOGENES uses a variation of the combination method known as the "Fast Path Check Method". The only error that it fails to report is when a chain head count is wrong.

DIOGENES has options for sending a copy of the output to a printer or creating a jobstream to perform the diagnostic. A printout is always recommended so that you have a permanent record of the diagnostic run.

DBCHECK & TDBCHECK From CSL or HP

Another tool is called DBCHECK or TDBCHECK. This program came from either HP or the CSL and has been distributed by several third party software suppliers. DBCHECK was written to diagnose IMAGE/3000 databases and TDBCHECK was developed to handle TurboIMAGE databases. These programs may also be known as STRUKCHK. I have also seen a version of DBCHECK from HP that handles both IMAGE and TurboIMAGE databases.

These programs are fairly small and do not offer a lot of options. If you do a Listf,2 the program file should have an EOF or Limit of 47 to 60 unless it has been run through the Object Code Translator (OCT). The user interface is very straight forward but I recommend having CAPS LOCK ON while running this program because it does not always shift input to uppercase.

This program automatically does both a serial scan AND chain chasing. There is no way to turn either of these options off. For a given database you may diagnose a specific dataset or enter /A to diagnose the whole database.

DBCHECK, TDBCHECK or NMCHECK From Adager

Adager users are provided with several tools in the LIBRARY group of the REGO account. DBCHECK is a re-written version of HP's DBCHECK. It has been enhanced with several options and can be used on either IMAGE/3000 or TurboIMAGE databases. TDBCHECK is a copy of DBCHECK and is provided for backward compatibility of existing jobstreams. NMCHECK is a Native Mode version of the DBCHECK program for those users on Spectrum machines.

These programs contain several options for determining which diagnostic method to use. By default, the programs use BOTH the serial scan AND the chain chasing method. The programs do not use the combination method. The programs allow a single dataset, a range of datasets (e.g. 1/20), several datasets (e.g. 1,5/10,12), or all datasets to be diagnosed. The default is that all datasets are to be diagnosed. There is an option in these programs to send a copy of the output to a printer during an interactive run. The following example is a jobstream to diagnose the database TEST on a Spectrum machine.

```
!JOB NMCHECK,USER.ACCOUNT,GROUP;OUTCLASS=LP,1,1
!RUN NMCHECK.LIBRARY.REGO
TEST
(blank line)
(blank line)
!EOJ
```

EXAMINE PATH From Adager

Adager users also have the EXAMINE PATH function available within the ADAGER program to diagnose specific paths for problems. This function uses the combination method to diagnose the dataset. For master datasets, Adager's EXAMINE PATH function performs a check of the synonym chains. For detail datasets it performs a check of the path between a given master and the detail while checking both the master and the detail for any serial problems. Unlike the "Fast Path Check Method" of DIOGENES, Adager's EXAMINE PATH function DOES detect problems with chain head counts and it also creates a log file of the errors that were found. Users can automatically create a jobstream to do an EXAMINE PATH by running Adager with Parm=8.

Other Tools

You may be using other diagnostic tools available within the Interex Contributed Software Library (CSL) or from other third party software suppliers or HP. Please check your documentation for specific information on how you can diagnose structural problems within your IMAGE databases using those tools and to determine which methods are used.

Don't Be Confused

Programs such as Robelle's HOWMESSY and the CSL program DBLOADNG are NOT meant to diagnose structural problems within IMAGE databases. In trying to obtain statistical information from a database they may encounter structural problems but they do not search for them and therefore should not be used as diagnostic tools.

Which Tool To Use When?

When users call with possible broken chains I ask them the following question:

Do you know in which set and possibly which path the problem exists?

If the answer to this question is "YES", I tell them to use Adager's EXAMINE PATH function on the suspect path(s) because it uses the combination method which is fast and thorough. If the user answers "NO" to this question I usually have them run Adager's DBCHECK (NMCHECK under Spectrum) or HP's DIOGENES, whichever program they feel more comfortable in using.

Do I Need To Have Everyone Logoff The Database?

This is a very common question. Most diagnostic tools do require exclusive access to the database in order to perform the diagnostic. Adager's EXAMINE PATH function can be run while other readers are in the database (i.e. the database is only opened in mode 5 or 6). Some tools, like the latest DIOGENES, allow the user to specify an IMAGE mode to use in opening the database. The only problem with this approach is that if people are updating the database at the same time as the diagnostic, certain problems may be reported which may or may not exist. Only by performing the diagnostic on a database that is not being updated can a user be assured that reported errors are not false alarms.

To avoid this problem and to allow writers back on the database, I suggest that the user restore a copy of the database into a different group or account (or another machine) and run the diagnostic against the copy. If the user does not have enough disc space for this copy (or does not have another machine), the next option is to run the diagnostic at night or on the weekend. All diagnostic tools should access the database in a read-only mode; consequently I tell users that if the diagnostic is taking too long it can always be aborted without harming the database.

Jobstreams vs. Interactive

As I mentioned above, having a permanent record of a diagnostic's findings is very important. Certain diagnostics can not create a log file or a printout while they are being run interactively. The user runs the risk of having the error messages scroll off the screen and out of the terminal's memory. Some terminals have printers attached to them and the user can log the information to the printer. Some diagnostics, like Adager's EXAMINE PATH function, produce log files if any errors were found. This is done just in case the $STDLIST is lost. Other diagnostics like DIOGENES from HP and DBCHECK (NMCHECK) from Adager can create printouts while running interactively.

Section IV: Preventive Maintenance

Predictive Support

All HP3000's should have Predictive Support installed and running. This tool determines if any of your system hardware is logging too many errors and alerts either the system operator or the HP Response Center automatically. The overhead associated with this feature is very minimal when compared to the potential for early problem diagnosis at the hardware level.

Full Stores and Restores

Accidental or deliberate partial stores and restores are the most common causes of damaged databases. Always be sure to check both your backup and restore listings to make sure that all the datasets were stored and restored properly. Backing up entire databases is preferable to backing up partial databases. Backing up only those sets that have been updated on a given day does not cause a broken chain when ALL those sets are restored because associated masters and details which have been updated are written to tape. Backup tapes should also be validated to make sure that the information can be read.

ILR Under MPE/V

One of the ways users under MPE/V can help prevent broken chains is by enabling the database for ILR (Intrinsic Level Recovery) within DBUTIL. This does not guarantee that you will not get broken chains but it minimizes the probability. When you enable a database for ILR, IMAGE builds a DBNAME00 file which is used as a potential backup log file in case of a system failure during an update operation. The overhead involved with this option is minimal and is offset by the potential for increased database integrity.

NOTE: Users under MPE/XL and MPE/iX should NOT enable IMAGE databases for ILR. This function is performed automatically by the transaction manager (XM).

"Paranoid" DBPUT

IMAGE is a very reliable DBMS which prevents and detects many faults. It maintains a doubly-linked list so it can check pointers as it is doing a chained read and can report a possible problem. However, there are some checks that IMAGE is not doing which are best done by DBPUT. There is no check being done by DBPUT of the key value as it is adding an entry to a chain.

I have been told that the IMAGE Lab is planning to add more checks to IMAGE under the category of "resiliency". Hopefully the Lab will also enhance the error messages of IMAGE intrinsics to include entry numbers and/or key values so that users don't have to go looking for a needle in the haystack.

Section V: Conclusion

Things To Remember

1. Try to PREVENT database damage as described in Section IV. This avoids down time, the use of diagnostics, database repair and, possibly, lost data.

2. Try to duplicate the problem so you don't spend hours or days trying to find a problem that isn't there.

3. If you doubt the conclusion of a given diagnostic tool, use another tool for a second opinion.

4. If a database has a problem in one area it may also have other problems. Schedule a full database diagnostic.

5. Establish a schedule for doing your database diagnosis. This can be done weekly, bi-weekly, monthly, quarterly, semi-annually, or even annually.

This last point is very important because it deals with the other question that I am asked frequently:

"How did my database become damaged and when did it happen?"

The "how" part of the question can only be answered by a hypothesis after studying the symptoms of the problem. The "when" part of the question is almost impossible to answer. Looking at the create dates of the datasets involved or knowing when the last system failure was may give the user a clue as to when the problem occurred. If, however, the user knew the last time the database was fully diagnosed with a clean bill of health, then we could safely say that the damage occurred between then and now. The closer "then" is to "now" the more likely that the user will be able to pinpoint the cause. I can't stress enough how important regular diagnosis of your IMAGE databases is.

Now that we are winding down I'm sure some of you are thinking to yourselves:

"Okay, now that we know how to diagnose our databases, what do we do to fix them if we find a problem?"

This question is a very valid one and of course there are several tools which can fix damaged databases. I hope to write a follow-up paper about fixing databases. Until then, contact me at Adager technical support after you have diagnosed your database and we will discuss how to solve your particular problem. I have fixed several databases using Adager's FIX CHAIN function. Obviously, if I don't hear from anyone about their damaged database(s) I will know that a follow-up paper is not necessary.

Happy Hunting!!!!!

SQL

Structured Query Language

The outside story

F. Alfredo Rego

Adager Corporation
Sun Valley, Idaho 83353-0030
U.S.A.

Telephone +1 (208) 726-9100 • Fax +1 (208) 726-8191
E-mail alfredo@adager.com

INTEREX paper 4008

Now that Hewlett-Packard has kindly produced IMAGE/SQL, we should do two things:

1. We should thank HP for the exceptional work that Jim Sartain's team did in integrating SQL (*Structured Query Language*, the industry-standard database language interface) into HP's award-winning IMAGE/3000 database management system.
2. We should understand SQL, so we can use it as a worthy complement to the other outstanding features that IMAGE/SQL provides.

What is SQL? What is it not?

Thanking HP is the easy part. Understanding SQL is not as easy. And I do not mean learning the SQL syntax—I mean comprehending and grasping the nature, significance and meaning of SQL. In particular, understanding SQL means to be aware that *SQL is not a type of database management system by any means: SQL is simply an interface.* SQL is just a type of data sublanguage that some database management systems—such as IMAGE/SQL—happen to understand.

Why did I choose "***The outside story***" for my title instead of sheepishly following the standard cliché, "***The inside story***"? I selected the title very carefully, to emphasize the point that SQL is indeed an outsider. (There is nothing wrong with the fact that SQL is an outsider, as long as everyone is aware of this fact. Unfortunately, there is much confusion regarding SQL.)

In this essay, it is my objective to balance things out. After all, with IMAGE/SQL you now have the best of *both* worlds: you enjoy the *inside* strengths of IMAGE as well as the *outside* connectivity of SQL.

4604

OBJECT TECHNOLOGY: A MANAGEMENT OVERVIEW

Orland J. Larson
Hewlett-Packard Company
19091Pruneridge Avenue
Cupertino, California 95014
(408)447-1197

ABSTRACT

"Object orientation" (00) is a new buzzword that promises to become the next "eureka" in information systems development. MIS management and their software developers are under increasing pressure to improve productivity, increase software quality and reduce implementation time. Traditional software development methods are not always able to meet these increasingly heavy demands. Object orientation is getting attention as a viable alternative.

This paper reports on the growing body of knowledge about object-oriented technologies. It begins by reviewing some of the critical challenges facing today's enterprises, followed by the definitions, basic mechanisms and key concepts associated with object-oriented systems. Next, it explores various types of applications that benefit from this technology. The potential benefits and the potential concerns will be addressed, followed by the impact object-oriented technologies may have on data administration and systems development in the 90's.

INTRODUCTION

Each decade one or two key advances emerge and change the practice of software development. Object-oriented systems and methods are rapidly entering the mainstream of software engineering and systems development. Leading consultants are heralding object-oriented approaches as one of the most important trends to affect businesses in the 90's.

Software is currently lagging behind hardware capabilities and the lag is increasing. There is general agreement that conventional software tools and techniques are rapidly becoming inadequate as software systems grow larger and increasingly more complex. Also, a consensus is building that the new paradigm

of object orientation may help control complexity and harness the expanding system environment into more useful and exciting applications.

Applications will need to satisfy more sophisticated requirements, use more complex data structures and architectures, and be delivered to an increasingly broad base of users. Software developers will have to increase their capacity to build, extend, and maintain complex, large-scale systems including their existing legacy systems. This requires that software be more flexible and easier to use.

Many of today's software development processes are out-of-date, with programmers still functioning like craftsmen. They build unique, noninterchangeable components and assemble them by hand, and then they struggle over time to understand the code generated by their predecessors and to extend and refine that software. As powerful computers pervade the lives of more and more people, the inability to deliver and maintain equally powerful software is an increasingly visible problem.

WHAT IS OBJECT ORIENTATION?

There is no single precise rule for describing or identifying object orientation. Rather, a collection of concepts together describes this new paradigm for software construction. In this new paradigm objects and classes are the building blocks, while methods, messages, and inheritance produce the primary mechanisms. Historically, creating a software program involved creating processes that act on a separate set of data. Object orientation changes the focus of the programming process from procedures to objects. Objects are self-contained modules that include both the data and the procedures that act on that data. The procedures contained within the object take on a new name, methods. Objects are activated by messages. Objects that have a common use are grouped together in a class, and new classes can be created that inherit the procedures and data from classes already built. This inheritance enables the programmer to reuse existing classes and to program only the differences. This provides for a new level of abstraction, with prebuilt libraries of classes and even prebuilt application specific class libraries or frameworks. Object orientation is important for the software development challenges of the 90's. This paradigm will improve the software development process and will cause new and better applications to evolve. It's promises will be delivered incrementally and across a broad range of technologies and will permeate the next generation of software architectures.

BASIC MECHANISMS

OBJECTS

Webster's New Collegiate Dictionary defines an object as "Something that is capable of being seen, touched or otherwise sensed." Grady Booch, in his book, "Object-Oriented Design with Applications", defines an object as "Something you can do things to. An object has state, behavior, and identity; the structure and behavior of similar objects are defined in their common class". David Taylor, in his book "Object-Oriented Technology: A Managers Guide" defines an object as "A software packet containing a collection of related data and methods for operating on that data".

Within objects reside the data of conventional computing languages, such as numbers, arrays, strings and records, as well as any functions, instructions, or subroutines that operate on them.

MESSAGES

Objects have the ability to act. Action occurs when an object receives a message, that is, a request, asking the object to behave in some way. When object-oriented programs execute, objects are receiving, interpreting, and responding to messages from other objects.

METHODS

Procedures called methods reside in the object and determine how the object acts when it receives a message. Methods may also send messages to other objects requesting action or information.

CLASS

Many different objects may act in very similar ways. A class is a description of a set of nearly identical objects. It is a category or collection of objects that share a common structure and a common behavior but contain different data.

INSTANCE

An instance is a term used to refer to an object that is a member of a class. Instance and object are used interchangeably.

INHERITANCE

Inheritance is the mechanism for automatically sharing methods and data among classes, subclasses, and objects. A powerful mechanism whereby classes can make use of the methods and variables defined in all classes above them on their branch of the hierarchy. Inheritance allows programmers to program only what is different from previously defined classes.

KEY CONCEPTS

ENCAPSULATION

Encapsulation is the process of hiding all of the details of an object such as its data (instance variables) and procedures (methods). This is also referred to as "information hiding".

ABSTRACTION

Abstraction is the process of creating a "superclass" by extracting common qualities or general characteristics from more specific classes or objects. Each level of abstraction makes the job of programming easier because it makes more reusable code available.

PERSISTENCE

Persistence refers to the permanence of an object, that is, the amount of time for which it is allocated space and remains accessible in the computer's memory. The object may continue to exist even after its creator ceases to exist. Objects stored permanently are termed persistent.

POLYMORPHISM

Objects act in response to the messages they receive. The same message can result in completely different actions when received by different objects. This phenomenon is referred to as polymorphism.

OBJECT-ORIENTED APPLICATIONS

Object-oriented applications will inspire users to think differently about the nature of
computing. Programs in an object-oriented environment will be transparent. Object-oriented frameworks will facilitate simulating and constructing user-specific solutions. Objects will be shared in networking environments to

distribute information within a work group or to parcel out tasks for distributed processing.

Object orientation is favored for applications that are characterized by complex processes and complex data manipulation. Applications in the following categories are classic candidates for enhancement through object orientation:

- Computer Aided Software Engineering (CASE)
- Computer Aided Instruction (CAI)
- Computer Integrated Manufacturing (CIM)
- Computer Aided Publishing (CAP)
- CAD/CAM/CAE Systems
- Document Management Systems
- Executive Information Systems
- Geographic Information Systems
- Graphics, Handling ICONS
- Health Care
- Image Storage Management
- Knowledge Based Systems
- Multimedia
- Manufacturing Production Control
- Manufacturing Requirements Planning
- Military Command and Control Decision Support
- Network Management
- Real Estate Systems
- Configuration and Version Management
- Telecommunications Routing Systems
- Visual Programming

Object-oriented applications will most likely gain in both presence and popularity.

POTENTIAL BENEFITS

Before managers can make informed decisions about adopting a new technology, the advantages of this technology must be translated into measurable benefits. Object-oriented programming improves not only the software development process but also the flexibility and utility of the resulting software. The design process becomes more intuitive as elements of the software correspond to elements in the application' s real world domain. The programming process itself encourages teamwork, code reuse, and code polishing.

Reusability is the key to increasing productivity in the face of increasing complexity. The key breakthrough in object technology is the ability to build large programs from lots of small, prefabricated ones. In addition to the increased productivity that results from reusability, using object-oriented technology can result in greater reliability because it reduces the risk of human error. Program structures remain intact, and change propagates naturally through the hierarchy of classes.

Flexibility is also a trademark of object orientation. Programmers are freed from the constraints of preestablished data types, allowing extensions of application functionality and bridging of heterogeneous applications.

Adaptability of object-oriented programs may well turn out to be the most crucial advantage of object-oriented technology. No matter how perfectly crafted, a program is useless if it doesn't meet current needs and the needs of users are changing at an ever increasing rate.

Faster development of applications is another benefit and is a result of all the programming effort that is reused from existing objects and all the design work that went into an existing model of a process.

Increased scalability is another significant benefit of object orientation. Given its improved modularization, it is especially well suited to developing large-scale systems.

Large systems are easier to build and maintain when you build them out of subsystems that can be developed and tested independently.

POTENTIAL CONCERNS

While object-oriented technology promises many benefits, there are some valid concerns about its ability to deliver those benefits. Most of these concerns have to do with temporary limitations and should disappear as the technology and its market mature.

The maturity of the technology itself is a concern to many potential developers. It is not yet a completely stable technology and many companies are not comfortable being "pioneers".

Standards are still evolving and the lack of accepted standards raises concerns about the difficulty of moving programs from one development environment to another and mixing and matching objects and classes from different vendors. Standards are on the way. The Object Management Group (OMG) was formed by a consortium of the major vendors of several object-oriented products. The purpose of this group is to promote the adoption of standards and the interchangeability of objects. In addition, the American National Standards Institute (ANSI) has an ODBMS committee, but no standards have been officially approved.

There is also a shortage of tools for application development. These tools include programs to assist in the design of objects and the management of libraries of reusable objects.

Performance of object-oriented applications is a concern and the speed of object-oriented systems will improve as the technology matures.

The object-oriented approach has a tremendous amount of potential and companies should explore this new technology, and check out the benefits for themselves.

EVOLUTION OR REVOLUTION?

Many organizations have invested heavily in existing non-relational and relational database management technology. The majority of these companies do not want to replace their existing databases and applications. The need to integrate these "legacy" databases with each other and with new systems is an important factor in the future evolution of data management.

There are clearly risks associated with getting into this technology too soon. But there are risks associated with waiting as well. The companies that begin the transition now will enjoy an important competitive advantage while the others strive to catch up.

The most prudent strategy is to avoid the extremes of ignoring the strategy or committing vital systems to it. Instead, companies can make a modest investment in a pilot program to gain first-hand experience with object-oriented development. This approach allows a company to reach its own conclusions about the value of the technology, and places the company well down the experience curve if it converts in the future.

THE FUTURE

Object-oriented technology will provide the clarity and flexibility essential to the successful development of complex systems. Today's applications do not offer the consistency and flexibility needed to make the computing environment more productive for users. Object orientation will provide environments in which users can communicate among applications and navigate easily over distributed, heterogeneous architectures.

In the near future object orientation will deliver the most benefits to three categories of programmers: power users, general business programmers and system developers. The most dramatic near-term benefits will be for system developers who both require and embrace this development approach and evolving tools to implement the increasingly complex and potentially innovative software of the 1990's. Over time, object-oriented technology will begin to have increasing impact on general business programmers and power users. Carefully designed object libraries will become available to support less sophisticated programmers who want to assemble applications quickly from prefabricated objects.

The vision of the future extends beyond the arrival of object-oriented system components development tools and standards. In the future, users will have the power and flexibility to design their own applications just by snapping together the necessary objects. With objects, building applications will be a process of tailoring and linking reusable modules. Object-oriented software architectures will mature in the 1990's. The transition to these new architectures is underway, marked by the arrival of object-oriented languages, databases, interfaces, operating systems and development environments. New types of data, distributed processing, multimedia applications, and end user computing are driving forces in the implementation of the object-oriented software environment.

There has been a great deal of progress implementing object-oriented systems. Today's graphical user interfaces have acquainted users with object manipulation. Among object-oriented languages, C++ has become the de facto standard. Object-oriented extensions are also being implemented in most popular commercial languages. Object-oriented development environments such as Hewlett-Packard's Softbench provide examples of object-oriented programming and applications. Operating systems are also being extended to support interoperability among object-based applications. Over the next decade, the difference between the old and the new will become increasingly obvious to both programmers and users. When the object-oriented future is fully delivered, this

natural, intuitive paradigm will be strongly embraced and will provide benefits to programmers and users alike.

OBJECT TECHNOLOGY FROM HEWLETT-PACKARD

In 1991, Hewlett-Packard announced HP OpenODB, the most advanced, commercial object-oriented database management system for large, multi-user environments. It is designed to enable new, complex business applications to be developed and maintained at a fraction of current costs.

HP OpenODB is based upon the Iris ODBMS prototype developed by HP Labs, starting in 1984. Using Iris, HP has worked extensively with customers and universities in evaluating ODBMS requirements. HP OpenODB is targeted at complex, commercial applications such as Geographic Information Systems (GIS), telecommunications systems and heterogenous information integration such as Executive Information Systems and Computer Integrated Manufacturing, whose needs may not be met by current database products such as TurboIMAGE and ALLBASE/SQL.

HP OpenODB uses a relational database as its storage manager and presents an object-oriented model to developers. It is a hybrid approach that allows integration of existing legacy systems including applications written in C, COBOL, FORTRAN, PASCAL, ADA, etc. This architecture provides a robust storage environment with the DBMS capabilities commercial users have come to expect. HP is currently using ALLBASE/SQL as the storage manager for HP OpenODB for performance and stability of data; however, the architecture allows HP OpenODB to be ported to other relational DBMS's for portability to non HP hardware. This combination is unique in the industry.

A Key benefit of HP OpenODB is that users don't have to abandon their current software or data to work with it. It includes an object-oriented structured query language (OSQL) and external functions that will retrieve information regardless of its format or whether it is stored inside or outside of HP OpenODB.

Object orientation is another phase in the evolution of computing and is an important step towards the vision of "Information At Your Fingertips". Developers must take advantage of the many benefits of this technology while at the same time deal with the hurdles that this technology poses. Object-oriented development environments must play a large role in reducing the learning curve and make object-oriented programming a highly productive process. Object-oriented products are here today and the commercialization of

object-oriented technology is increasing rapidly. Object based architectures lend themselves to the creation of a much richer information environment. Digitized voice, music, video clips and animation will begin to populate our information systems. Object database systems are currently viable for commercial projects and will be widely adopted by the mid to late 1990's.

REFERENCES

Booch, Grady. *Object-Oriented Design With Applications.* Redwood City, CA: The Benjamin/Cummings Publishing Company, Inc, 1990.

DeCastilhos,Margie.*Getting Started With Object-Oriented Databases.* INTERACT Magazine (Hewlett-Packard User Group Publication).pages 102-113, March 1992.

Hewlett-Packard.1990. *CASE and Objects-A Primer(An Introduction to Object-Oriented Concepts for Software Engineering),* Palo Alto , CA:HP Literature # 5952-2624.

Taylor, David A., *Object-Oriented Technology: A Manager's Guide.*Reading,MA. Addison-Wesley,1991. (A very good overview of object-oriented technologies for managers)

Winblad, Ann L., et al. *Object-Oriented Software.*Reading, MA. Addison-Wesley 1990.

4605
An Overview of HP IMAGE/SQL

Hewlett-Packard Company
Commercial Systems Division
Database Lab
19111 Pruneridge Avenue
Cupertino, California
(408) 447-5649

This paper represents the collected works and efforts of several individuals from the Database Lab and Marketing organizations of the Commercial Systems Division of Hewlett-Packard Company. This paper will be presented by any one of these individuals at every major Users' Group gathering in 1994.

****

Some of you are new to IMAGE, and others of you have been familiar with the database since the mid 1970's when it was winning awards for function and design. Thousands of applications, from Email to finance to production control, all operate using IMAGE as their database management system. And if you count the users executing programs against data stored in IMAGE, their numbers would, no doubt, be in the millions.

Without argument, IMAGE gives you high reliability and high performance. After almost twenty years of improvements and refinements, you, as customers and users, would expect nothing less. With so much going for this product, with a dependability that few other database products offer, as a storage management tool that is a de facto standard on HP 3000 computer systems, what could possibly be done to improve it? The answer is quite simple. Standardize the access to IMAGE!

Hewlett-Packard is proud to introduce HP IMAGE/SQL, relational access to IMAGE data using industry-standard Structured Query Language (SQL). This method of access includes full read and write capability using ANSI standard functionality. This new access method makes a myriad of application development and decision support tools available to IMAGE that have never been available before.

The most astounding part of this whole story is that SQL access to IMAGE and the access that you enjoy today are completely compatible. Complete coexistence. No conversions, no recompilation, no changes. The only change will be the way you look at your data once you try the application development and decision support tools now available with IMAGE. And these are the same tools that can

be used with all the big name relational databases, including ALLBASE/SQL. IMAGE has moved into the Open Systems arena.

Those of you familiar with IMAGE will undoubtedly be asking the question, "but what will this do to performance?" Of course, every database design and application will have it's own performance characteristics, but, as a general rule, IMAGE/SQL will perform at about 70 to 90 percent of native TurboIMAGE access. This term, 'native' refers to intrinsic level access using COBOL, PASCAL, et cetera. So, you will have the best of both worlds, fast native access and relational access.

With IMAGE/SQL you will have everything that you have today, wrapped in a new relational package, with an enormous selection of new tools, without any hassles. Let's take a look at the issues HP has tried to address.

TOP ISSUES

There are a multitude of issues which HP hopes to resolve with the introduction of IMAGE/SQL, and these will be explained shortly. Unfortunately, some new issues will be created as well, and these will also be revealed. These issues have been grouped according to the kind of organization that will be impacted, and they are VARs (Value Added Resellers) and ISVs (Independent Software Vendors), MIS or IT (Information Technology) Departments, and End-users.

Most VARs and ISVs, whose products include applications and tools, have been limited to non-relational application development tools when dealing with IMAGE. There are also very few PC-based GUI (Graphical User Interface) products which offer transparent interaction between PC clients and HP 3000 servers using IMAGE. As mentioned before, thousands of applications have been developed by uncountable companies all using IMAGE as the storage management system. Now, with IMAGE/SQL, all the new SQL-based application development and decision support tools, most using some type of GUI, can be used by VARs and ISVs to improve and enhance their products. This new transformation of IMAGE data into relational data will revitalize the VAR/ISV product environment.

Turning to MIS and IT departments, the primary issues are:

- o reducing the application backlog,
- o providing flexible information access,
- o recruiting SQL-trained personnel.

Each of these deserves some expansion.

Ideally, IMAGE/SQL should reduce the application development backlog for any MIS/IT organization. This reduction would be accomplished using any of the easy-to-use application development and decision support tools now available. Sadly, the opposite effect, that of increasing the demand for applications, will probably become true. As the end-users see the amazing results from new implementations using these tools, all kinds of previously hidden application requests will surface. Once the power of IMAGE/SQL is known, MIS/IT departments will be flooded with new requests for information.

Does this mean that these departments should avoid using IMAGE/SQL just to avoid this rampage? Absolutely not! The gains a company can realize from improved information review and analysis, as the result of improved data availability, can be significant. And MIS will not lose control. Users will be able to do their own queries and reporting, MIS need only supply the access. Some of the other issues in this area will further explain these statements.

The second issue in this area pertains to flexible information access. What this refers to is the powerful, command-driven language which characterizes SQL. Complex and sophisticated queries can be executed to supply end-users with their specific information requirements.

SQL uses command sentence constructs of SELECT and WHERE clauses, and the relational JOIN clause, to define its syntax. This level of knowledge is not necessarily required to use the various tools, however more sophisticated applications might require SQL expertise. Finding the right people to develop these queries and deliver these applications to the end-users will be critical. This brings us to the last issue for the MIS/IT department -- finding qualified personnel.

Relational concepts in information management have taken center stage in the last ten years. These concepts have become the de facto standard for data storage systems at educational institutions all over the world. As such, new and old computer professionals alike have had either some or extensive exposure to relational concepts.

SQL has become the relational language standard and is taught widely. Finding qualified personnel to program using SQL today is easier than finding experienced IMAGE programmers. So, even though your information investment is in IMAGE, the investment you make in accessing that information can be in relational technology.

The end-user has repeatedly come up in the last several paragraphs. Their issues are improved productivity and improved decision support. Both of these issues rest firmly on the same foundation; information that the end-user requires to be more productive and upon which decisions are made must be located where the

user can exploit it and be presented in a meaningful form. This can be achieved with the use of PC-based client/server tools featuring graphical user interfaces (GUI). These are the very same tools the MIS/IT department will be using for development and that VARs and ISVs will use to enhance their offerings.

THE NEW TOOLS

The new tools have been mentioned over and over. Before getting to specific products, let's first describe them generally, and then break them into categories to better understand what they can do and offer.

All of these tools are PC-based and function primarily in a client/server environment operating with Microsoft Windows. With exception, these tools use ALLBASE/PC API, which is based on the Gupta standard SQL API, as the link to the server.

For those of you not familiar with what an API is, some explanation is in order. API stands for Application Programming Interface and refers to the component of the client/server model which performs the interactive linking between client and server over the network for the purpose of data exchange. This can be thought of as application level handshaking, much the way RS232C is an electrical signal handshaking in data communications.

Besides Gupta, several other PC API standards exist. The use of the term 'standard' may seem vague as it is used here, but it basically refers to an agreed upon set of operating conditions. None of these standards represent an industry standard, but some type of API is required for client/server computing, so each vendor must select one or more standards with which to operate. So far, the Gupta standard serves the majority of the tools which have been certified for use with IMAGE/SQL. Microsoft has announced ODBC, its API standard, and this API is becoming a significant player in client/server technology. Support for ODBC will be available in January 1994.

Our first category of tool is Decision Support System (DSS) Tools. As the name suggests, the purpose of these tools is to support business decisions, and this is accomplished through information analysis, reporting and graphical representation. Typically, this type of tool is oriented towards end-users doing financial and managerial analysis where numerical quantification and graphical representation of data is useful. These tools also tend to offer formatting of data for subsequent importing to spreadsheet products for further manipulation and review. End-user knowledge of SQL is not a prerequisite for using these tools.

Some examples of Decision Support Tools are Impromptu by Cognos, Q+E Database Editor by Q+E, and HP's own NewWave Access. Although this last tool is not new, the relational access now delivered through IMAGE/SQL will dramatically reduce the overhead which will improve administration of these systems and make this solution easier to setup and maintain.

Our next category of tools is called Executive Information System (EIS) Tools. The key aspect of these tools is their exception management capability, where information is monitored within user defined limits and deviations highlighted. These tools use colors and graphics extensively to bring attention to situations where limitations have been exceeded. Other features include trend analysis and information drill-down. Drill-down simply refers to digging out the detailed elements of summarized information. A good example of an EIS tool is Forest & Trees by Trinzic which will be discussed in more detail later.

The last category of tools for use with IMAGE/SQL is Application Development Tools. Falling also into the category of Fourth Generation Languages (4GL), these tools offer PC Windows programming capability with Graphical User Interfaces which allow programmers fast development of critical end-user applications. In many cases, some knowledge of SQL is necessary.

Where DSS and EIS tools tend to be read intensive or read only, application development tools are intended to create interactive applications for online transaction processing. Some offerings in this category are PowerBuilder by Powersoft and SQLWindows by Gupta.

All the aforementioned categories make up the new tools available for use with IMAGE/SQL. These tools also operate with HP's ALLBASE/SQL and the other major independent relational databases. A review of some of these tools follows in the last section, and these tools are summarized in Table 1.

	Cognos Impromptu	HP Information Access	Q+E Database Editor	Trinzic Forest & Trees	Gupta SQL Windows	Powersoft Power Builder
Type	DSS	DSS	DSS	EIS	Application Development	Application Development
API	Proprietary	Proprietary	Allbase/PC	Allbase/PC	Allbase/PC	Allbase/PC
Connection Type	Serial/LAN	Serial/LAN	LAN	LAN	LAN	LAN
Operating System	Windows	DOS, Windows, NewWave	Windows	Windows	Windows	Windows
Learning Curve	Hours	Hours	Hours	Days	Weeks	Weeks
SQL Knowledge Needed?	No	No	No	Some	Yes	Yes
Phone	(800) 426-4667	(800) 752-0900	(800) 876-3101	(800) 289-0053	(800) 876-3267	(800) 395-3525
Fax	(613) 738-0002	(408) 447-0264	(212) 967-6406	(603) 427-0385	(415) 321-5471	(617) 273-2540
Complementary Products	PowerHouse Windows		Q+E Database Library	InfoPump		

TABLE 1

CURRENT TOOLS OFFERINGS

In addition to all the PC client/server tools mentioned above, a rich assortment of direct-access 4GL tools exist for accessing HP's relational databases. This prior sentence specifically indicates databases in the plural. All the tools in this section work with both ALLBASE/SQL and IMAGE/SQL. The important point about these tools is that they execute on the host, not the client/server cooperative execution of the previous set of tools.

The products in this section have, for the most part, been around for some time. Details will not be presented here about features and benefits. The main point here is that several vendors offer SQL-based tools for accessing relational databases.

Some of these 4GL tools offer native (intrinsic) access to IMAGE. However, the relational access they offer ALLBASE/SQL now extends to IMAGE/SQL. Below is a list of products and the companies which offer them:

ALLBASE Toolset	Hewlett-Packard
Transact	Hewlett-Packard
Powerhouse	Cognos
Focus	Information Builders
JAM	JYACC
Speedware	Speedware
Uniface	Uniface

As you can see, these tools, along with those indicated for PC client/server, make an impressive arsenal in your development efforts.

Two relational databases from HP can reside on your system, ALLBASE/SQL and IMAGE/SQL, and all these tools can access each. Most of what this paper describes is what IMAGE/SQL offers you. Some explanation of what ALLBASE/SQL offers follows.

WHEN TO USE HP ALLBASE/SQL

HP offers two database management systems on the HP 3000 platform. Each has its place with any given application, and many applications could use either. The majority of database usage on the HP 3000 is currently IMAGE, but there are certain applications where ALLBASE/SQL is the better choice.

ALLBASE/SQL is a full-featured relational database management system (RDBMS) with functional compliance of ANSI standards. HP sees four specific areas where ALLBASE is the preferred DBMS:

Mainframe class computing,
Distributed applications,
High-volume online transaction processing,
Object-oriented applications.

Below are expansions of these topics.

HP has positioned ALLBASE as the alternative to mainframe-class database management systems (DBMS). This works in conjunction with HP's mainframe downsizing strategy which offers high-end HP 3000 systems as Corporate Business Systems. Providing online backup and restructuring, and supporting very large file sizes, ALLBASE is the DBMS of choice in this area.

Distributed information and the applications which support them are also addressed within the ANSI standards for relational database systems. Here, again,

ALLBASE has been specifically featured. Using two-phase commit protocols, distributed transactions can be ensured of completion and accuracy.

ALLBASE also supports Encina technology from TRANSARC Corporation. This technology provides a standardized method of transparent distributed transaction processing in a multi-platform environment. These distributed application features make ALLBASE the clear choice compared to IMAGE/SQL.

Where high-volume online transaction processing (OLTP) is the objective, and relational concepts are required, ALLBASE is again the best choice. This pertains mostly to new application development. Since ALLBASE has been designed from the ground up as a high-end relational database, applications requiring high-volume SQL OLTP will benefit.

Object-oriented computing is being seen more and more as a viable solution in many applications. Most assuredly, any application which has multimedia requirements can benefit from object technology. Here ALLBASE/SQL has the advantage. Object storage, especially that of video, photograph, audio, et cetera, requires data structures foreign to IMAGE/SQL. ALLBASE/SQL already provides the necessary storage with Binary Large Objects (BLOBs).

From these descriptions you can see that there are specific situations for ALLBASE/SQL.

COEXISTENCE

The HP 3000 now offers a complete range of data management choices. Native IMAGE still provides the highest performance in the industry for mission-critical OLTP business applications. IMAGE/SQL provides data access through the multitude of 4GL and PC client/server toolsets at a small performance premium. Client/server computing is clearly the emerging trend in information processing, and IMAGE/SQL is now an important element, ensuring the protection of your information investment.

ALLBASE/SQL provides the highest SQL performance of any relational database in the industry and provides support for distributed database and distributed transaction processing. With Corporate Business Systems, ALLBASE/SQL can handle the requirements of nearly any enterprise at a fraction of the cost of traditional mainframe solutions.

Together, these database management systems make a powerful team. And they work together. Concurrent access is possible by linking the two together within a single environment. Your applications, whether host-based or client/server, can

simultaneously access ALLBASE/SQL and IMAGE/SQL information. That is coexistence at its best.

But there is more.

THE RELATIONAL PICTURE

HP offers you two relational databases, IMAGE/SQL and ALLBASE/SQL. But the story gets better. Also available on the HP 3000 are Oracle and Ingres, and each of these RDBMS products have their links to IMAGE as well, Oracle through their OracleConnect product, Ingres through their ALLBASE Gateway to IMAGE/SQL. This creates a SQL shell over every database management system available on the HP 3000. And with this shell comes all the 4GL and client/server tools which have been mentioned throughout this paper. This is a very powerful offering, indeed.

Taking a step farther, HP has proposed several object-oriented functions for addition to the SQL standard, the objective of which is the defining of Object SQL (OSQL). Should these proposals be accepted, object technology will be a step closer to IMAGE/SQL.

This last section gives a picture of where we are today and a glimpse of what the future might hold. Let's step back and review what has taken place in the last several years to get us where we are today.

HP DATABASE DEVELOPMENTS

By looking at just the past couple of years, it's easy to see that HP is actively working to make their database products superior. In 1992, we released the following new features:

IMAGE

- o Third-party indexing interface
- o Critical item update
- o 4GB file sizes
- o Corporate Business System tuning

ALLBASE/SQL

- o Database shadowing
- o Record level locking
- o Remote unattended backup
- o Stored procedures

- o Business rules and triggers
- o Two-phase commit (via XA interface)
- o Additional third-party tools

By far the biggest announcement, occurring late in 1993, was the introduction of IMAGE/SQL. Release of this product provided the following features to IMAGE:

IMAGE/SQL

- o Relational access via SQL
- o Stored procedures
- o Business rules and triggers
- o SQL PC API support

The year 1994 will not be without its share of enhancements, either. Announced for release in 1994 are:

IMAGE/SQL

- o Dynamic detail set expansion
- o Increased software resiliency
- o Performance tuning
- o Predicate-level locking (SQL access)
- o Third-party index aware (SQL access)
- o Native mode Query

ALLBASE/SQL

- o Fast recovery from media failures
- o Roll forward of physical file DBA operations
- o 4Gb log files
- o Dynamic setting of isolation level and transaction priority
- o Parallel pre-fetch for faster table scans
- o Access path modification for greater control
- o Ongoing multiprocessor scaling
- o SQL monitor tool for performance analysis/troubleshooting
- o Faster load algorithm
- o Parallel loads
- o Truncate table (empty table quickly)
- o SQL audit tools for viewing DBA changes
- o Case insensitivity option

IMAGE/SQL & ALLBASE/SQL

- o ANSI'92 entry level compliance (except AS' Clause)
- o Additional third-party tools
- o ODBC support
- o Improved pre-processor concurrency
- o Multi-row stored procedure result sets
- o Extract comman for module movement
- o Application thread support

This impressive list of delivered and promised features is a clear demonstration of HP's dedication to providing the best data management systems on HP computer systems. But for all the features, if the product is not easy to use, the product will not be used. Let's next explore just how simple it is to use the new IMAGE/SQL.

ELEGANT SIMPLICITY

Accessing IMAGE relationally is as simple as 1-2-3.

Step 1: create the environment;
Step 2: attach IMAGE to the environment;
Step 3: access IMAGE relationally.

Each of these steps will be briefly described in this section.

Step 1: Create the SQL Database Environment

This step involves the creation of the SQL Database Environment (DBE). This structure contains control information about the data represented by it. This environment is created using the ISQL utility as follows:

```
:ISQL
ISQL => START DBE 'SQLDBE' NEW;
```

If a DBE already exists, this step can be omitted.

Step 2: Attach the IMAGE Database to the DBE

The attachment process examines the IMAGE database and places equivalent SQL-structure information in the DBE. This step can include a variety of mapping functions for fine tuning relational access. This step uses the IMAGESQL utility.

```
:IMAGESQL
> SET TURBODB CUSTDB
> SET SQLDBE SQLDBE
> ATTACH
```

The assumption in this case is that CUSTDB is an existing IMAGE database.

Step 3: Access IMAGE Data Relationally

Any number of methods could be employed in this step to demonstrate this access. In this case, ISQL is used.

```
:ISQL
ISQL => CONNECT TO 'SQLDBE';
ISQL => SELECT * FROM CUSTDB.CUSTOMERS;
```

The result of this query would be a tabled list of all entries in the CUSTOMERS data set. You can also display information about the database itself.

```
ISQL => SELECT NAME, OWNER FROM SYSTEM.TABLE;
```

This would return all table names associated with the SQLDBE environment.

Updating IMAGE databases is also very simple. Here is another ISQL example which updates a column (item) called PRODUCT_NAME.

```
:ISQL
ISQL => CONNECT TO 'SQLDBE';
ISQL => UPDATE CUSTDB.ORDERS
                    SET PRODUCT_NAME = 'IMAGE/SQL'
                    WHERE PRODUCT_NAME = 'IMAGE';
```

This update finds all entries where PRODUCT_NAME is 'IMAGE', then changes that value to 'IMAGE/SQL'.

As these steps demonstrate, accessing IMAGE data relationally is very simple and straight-forward. Once the database is attached to the DBE, nothing else is required in the regular use of the system. There are additional administrative tasks related to security and data type mapping, but none of these are overly complex or cumbersome.

Next, let's discuss the simplicity of some the client/server tools.

Forest & Trees by Trinzic

If you recall, Forest & Trees (F&T) is a PC client/server tool falling under the category of EIS. This means this tool reads, reports and formats data from the IMAGE/SQL database (or any of the other host- or PC-based databases).

Quoting from the Reference Guide: "Forest & Trees... collects and combines data from a variety of sources and monitors the resulting information in order to track information at all levels from business vital signs to underlying detail.

"To help you display, integrate, and use the collected information, F&T has a large set of data manipulation and object control functions. These functions can be used in formulas, queries, triggers, graphs, and reports. "Forest & Trees can also be fully customized with pictures, buttons, menus, and other graphic attributes to create distinctive applications."

Although the details of this tool are too extensive for this paper, some details will be discussed to demonstrate this tool's ease-of-use.

F&T has a tool bar at the top of the window which is filled with object icons for ease in application control and design. Views of data are defined by the user and extracted on demand or when scheduled. These views can be at any level, from high-level summary to low-level details. The drill-down feature of this tool allows display of all levels of information from top to bottom. Graphic views of the various levels of information can also be created and displayed.

One of the outstanding features of this tool is that user-defined limits or ranges can be set to monitor selected items, formula results, and sums. Color coded warnings, green for within limits, yellow for warnings, and red for outside limits, can highlight monitored data. Exception management is much simpler when the software isolates and highlights the exceptions.

PowerBuilder by Powersoft

Another PC client/server tool is PowerBuilder by Powersoft. This tool is from the Application Development category presented earlier and is both a reporting and transaction processing tool.

Quoting from a Powersoft sales brochure, "PowerBuilder... is a comprehensive Microsoft Windows-based development environment for constructing graphical client/server database applications through object-oriented development.

"PowerBuilder supports the rapid design and development of Windows-based client/server relational database applications for all marketplaces.

"PowerBuilder... contains full support for the MS Windows GUI technology, such as radio buttons, command buttons, list boxes, bit maps, etc., [and] supports custom user objects that can be defined as standard Windows controls."

As indicated, PowerBuilder is a very full featured application development tool. As mentioned above for Forest & Trees, this paper is not intended to be a detailed discussion of this products features. Instead, to summarize, this product offers significant productivity gains through development and use of its object technology orientation.

CONCLUSION

HP IMAGE/SQL is the most significant enhancement of the IMAGE DBMS of all time. The opportunity which this announcement provides to commercial and company developers is enormous considering the selection of tools now available for use with IMAGE. IMAGE/SQL also solidifies its coexistence with HP's other relational database, ALLBASE/SQL, with standard access to both.

The simplicity with which relational access occurs with IMAGE/SQL makes this perhaps the most painless 'conversion' in the history of the computer industry. The term 'coexistence' applies equally well to the dual accesses to IMAGE, native and SQL, as it does to the two relational databases, IMAGE/SQL and ALLBASE/SQL.

For all that this paper presents, the most important point is this: With IMAGE/SQL, HP has once again protected your investment in the HP 3000!

Paper Number 4704
The New Relational Database: IMAGE/SQL

Bradley Tashenberg
President
Bradmark Technology
4265 San Felipe, #800
Houston, TX 77027

Handouts will be provided at time of presentation

HP Migration: A Management Overview

or

How to Avoid the Pitfalls

Paper Number : 5003

Walter F. Pawul, Jr.
Executive Vice President
Players Computer Inc.
20 Constance Court
Hauppauge, NY 11788
(516) 582 - 2650
(516) 582 - 2767 FAX

Players Computer Inc. (PCI) has been a loyal Hewlett Packard (HP) user for over 15 years. We started with a HP 3000 Series III and have systematically upgraded to our present Spectrum Class CPU. The upgrade path, offered by HP, always protected our software investment. However, the performance gains realized were never large enough to excite the users and management. Since PCI is a relatively small service oriented firm, capital expenditures for system enhancements were always painful. The costs incurred for a marginal performance gain always made for a troublesome decision. Previous HP 3000 upgrades were always looked at as "money spent begrudgedly". PCI basically upgraded to avoid HP obsolescence with hopes of marginal performance enhancements.

About one (1) year ago, PCI was faced with yet another upgrade decision point. This time the decision was complicated with the fact that we had already reached the end of the Classic product line (HP 3000 Series 70). PCI was very concerned about the transition to the XL box. Based on past performance gains and the XL switch over, PCI was reluctant to commit to the Spectrum purchase. We kept pondering whether to hang on for another year with our slow but reliable Series 70. In retrospect, we could kick ourselves for not moving to the RISC machine 2 years earlier. Our HP 955 produced mind boggling performance increases. Many times PCI experienced as much as a twentyfold performance gain. Even our users were ecstatic about the CRT response time improvements. In all the years, of my professional career, I have never seen such a beneficial cost/results ratio regarding a hardware update.

If there is anyone still pondering the move to a Spectrum Machine, stop worrying and move to the promised land! This move will be the best money your company ever spent regarding hardware improvements. The move is costly and has many hidden pitfalls.

I attended INTEREX 92, in New Orleans, to gather as much information related to the MPE to XL conversion. There were no papers or tutorials which dealt directly with this topic. I did, however, come away with a much better overall picture of the conversion process. Much of the exposure came from roundtable discussions and candid conversation with the various vendors attending the show.

This paper is directed at a manager who is faced with the transition from the Classic to Spectrum computer. I will attempt to outline the steps necessary to provide a smooth transition and hopefully help minimize unnecessary or unforeseen costs/pitfalls. This paper is based on PCI's experience and may not provide guidance for all situations. However, I strongly feel that most of the information I outline will be helpful.

The Migration life cycle can be separated into distinct topics of interest. They are as follows:

A. Selection of the proper CPU model and whether to purchase new or third party pre owned
B. Site preparation
C. Personnel training
D. Software, Datafiles and Database conversion
E. Cost considerations

Selection of the Proper CPU

There was a fare amount of confusion regarding the relative processing power of one Spectrum model vs a newer model. The only safe bet was that the Series 70 would be overmatched by the RISC machine. The first step is to have a HP salesperson propose a solution using the most current technology. We used this model as our

baseline for all comparisons. Next, we contacted the used market experts for quotations on a CPU with comparable processing speeds.

This is when things become more difficult to quantify. The current system configuration offers lower monthly maintenance, electrical loads and the opportunity to utilize current disc related technology. There might also be some savings on the Operating System support if your system is below the "unlimited level" of the older Spectrum machines. An older model (1 - 2 generations) will significantly lower the initial cash outlay and allow you to purchase more CPU for the same dollars. You will however, every month, incur higher operating costs. A simple cost justification exercise. Wrong. There are some hidden items which could drastically swing your decision in either direction.

If your operations are dependent on the HPIB technology, there are limitations on the number of HPIB devices available, on the newer models (9x7). To get more HPIB cards means upgrading the CPU. In our case, PCI needs multiple 6250 tape drives for our customers output. Backups utilize only a fraction of our tape drive operations. In addition, switching the printers to a serial interface would cause unwanted throughput reductions. We only saw growth in both tape output and printers which made a new machine look less appealing. The large HPIB capability was obtainable at a much higher cost.

The next dilemma is the compatibility of the HP7937 disc drives. Technically they are compatible with the 9x7 class. The only problem is they use HPIB and we are back to the HPIB limitation scenario! We had a large investment in these drives (upgraded under the assumption they were upward compatible) and could not see discarding the same for a more expensive alternative.

The HPIB problem coupled with being able to buy more CPU swayed our decision to purchase a used computer. HP may have addressed the HPIB issue since our procurement. Each decision must be based on the particular needs of the respective company. The lesson to be learned is to be wary of the possibility that all of your current peripherals may not be compatible currently or within a sustained growth period.

Another new domain is internal memory. Internal memory is not new but the amount used relative to the Classic is vastly different. HP and the user community has established rules of thumb which serve as a good guide for choosing the amount of main memory needed. Although the memory concept is foreign from the Classic (so much more memory is needed) you should feel confident that the memory estimated by HP or a third party should do the trick.

Site Preparation

There are two different approaches to the site plan. The first and safest is to have the old and new systems running during the transition. This will allow for a less painful and pressure oriented conversion. The other is to convert over to the Spectrum machine over a weekend (hopefully a long one). I can not see how our company could have possibly done the conversion over a weekend without major problems. I strongly recommend some type of dual systems for true disaster protection.

Before you can proceed, using the dual approach, there are a few items which must be addressed. There must be enough air conditioning capacity to support the two systems. Two overheated CPU's are of no use to anyone. Physical space could be a problem. Since there will be two computer systems, try to orient the new system

as close to its final resting place as possible. Keep in mind the placement of the DTCs so you can keep the re-cabling down to a minimum. Also try to located the printers and tape drives outside the CPU room. This will lower the A/C load, lower space requirements and keep people traffic to a minimum. The last item of concern is to have all of the electrical cables in place. This will mean that there will be some unused cables left over after the transition is made.

The next planning item is the cable switch-over. Make sure the DTC locations are within the cable lengths of the old system. If not, you will be faced with the horrendous task of re-cabling or using jumper cables. The DTC layout favors 3 pin connections. Make sure that you have enough 3 pin cable kits to modify any DB25 connectors to the 3 pin variety. I would not recommend using jumper cables (3 pin to DB25). This will create unnecessary bulk and confusion underneath the DTC cabinet. Lastly, label all of the cables, prior to moving them to the new DTC, clearly with labels which will not tear or rip when they are moved beneath the computer floor. Trying to identify unmarked cables is a time consuming and frustrating job. This is further complicated when the mystery cable is for a printer. We did not adequately mark our cables and we paid for it dearly. The logical device (LDEV) numbering sequence also has changed on the Spectrum class CPU. Any cabling diagrams or patch panels have to be modified to reflect the new LDEV designations. You should also try to maintain the location of the new ports in relation to the existing patch panel layout. We moved certain port locations (to segregate printers and modems) which required us to rearrange our patch panel. Depending on the extent of the swapping, this could be an arduous task.

Personnel Training

The operation of the Spectrum class computer is very similar to the Classic but at the same time very different(what a profound statement). If the system configuration and DTC layouts were all performed before handing over the system to the actual operators (all done by the CE), the system could be operated with no pre-training. There will be a fair amount of mistakes and many calls to HP PICKS. In addition, the system administrator would be operating "in the blind" and unable to manipulate the system with any type of confidence. I would not recommend the "no training" route. Having a knowledgeable support group will help ease the transition and let management concentrate on any priority issues.

The Systems Administrator should attend the appropriate class for the XL operating system. The lead computer operators should also attend their respective class. The critical element is timing. The training must occur before the equipment arrives. However, you must not allow too much time to lapse between the training and the arrival of the new system. This is a tough item to balance because of HP's sporadic scheduling of training classes. Pre training and the use of PICKS will make the actual operation of the system proceed smoothly. All the operational support people, including the programmers, must be re educated on how to bring the system up and down. The procedure is different (actually easier) than the MPE/V operating system.

I found that there was no need for pre training of the programming staff. The HP COBOL/XL manuals were sufficient for the conversion phase. Again, any roadblocks were resolved with the help of the HP Response Center. Any training for the programmers or analyst can be conducted after the new system is operational and stabilized.

Software, Datafiles and Database Conversion

The most anxiety ridden facet of any hardware transition deals with the transfer of the programs and data. HP has always been superb in providing an upward compatible environment for software. Since the XL move is more than moving from a HP 68 to a HP 70, we had grave concerns about making a transition. To alleviate or almost eliminate these negative thoughts, we decided to have a redundant system running under the MPE/XL environment. This allowed us to transport our operation in distinct sections. The sections were loaded, tested, retested and then re-loaded as "live". Any operational or programatical problems were addressed in a timely but non-crisis mode. Our real pressure was to meet timelines for application transfer. We never had to deal with irate clients, corrupt databases, re-loads or other uncomfortable situations which are a direct result of crisis management decisions. A module or section never went live unless there was a high degree of confidence in the operation of that respective section.

The first step to transport the software must be started weeks before the new equipment arrives. The HP Migration Toolset is a valuable asset. This software and manuals can be ordered directly from HP. The manuals are helpful in guiding the user through the steps necessary to analyze all of your application programs. The Toolset will point out any possible trouble areas within your software. As long as your software is not using privileged mode or heavily using intrinsic calls, the problem areas should be minimal. Most of the problems we uncovered were from third party software which had to be upgraded to XL anyway. The HP Migration Toolset does produce reams of paper. We saved the output to an editor file and then scanned the file

only looking for errors. Once an error was found those respective pages were reviewed for action. The HP Response center can also be helpful in interpreting the error conditions. Make as many changes to the application software as you can before the new system arrives.

When the system arrives the excitement begins. Each and every program has to be compiled in native mode (this will yield the best performance). The XL compiler is less forgiving which will result in some minor changes to the programs. All of our application software is written in COBOL and therefore I can not comment on any other languages. Since we are a service bureau, PCI maintains thousands of application programs. The compilation task would have been insurmountable without help from MPEX. This third party tool allowed compilations to be performed, as a batch job, and overnight. Our programming staff only dealt with the problem areas which made this task quite reasonable.

The next area of concern is the JCLs. The use of names for printer devices (ie LZJET or LP1) are frowned upon in the XL environment. We did not uncover this phenomenon until we started to actually test the translated programs. We had countless JCLs which utilized the device names and had to be changed to the actual LDEV number. The thought of having to manually change each and every JCL on our system was overwhelming. Once again, MPEX came to the rescue and an endless task became an overnight process. This process can be completed before the new system arrives.

The databases, KSAM and flat files can be transported directly to the new platform. This was the easiest part of the move. Simply restore them from a backup tape and you are in business. The IMAGE databases, once re-stored, are fully operational and need no further

intervention. The KSAM files, on the other hand, are not truly XL files. The transition to XL is very easy and not very time consuming. The changeover can be done under a very controlled environment. Changing a KSAM file is simple and quick. FCOPY the present KSAM file into a flat file, build the new KSAM file with KSAMXL and FCOPY the flat file back into the new KSAM/XL file. A nice feature of the KSAM/XL file is when the system comes down unexpectedly, the KSAM file will not become corrupt and therefore, you do not have to run the recovery utility. For this reason, all KSAM files should be converted to KSAM/XL as soon as time allows. This is also an ideal time to expand any large databases and/or KSAM files. Space and CPU resources will be at a maximum. The expansions will complete in record time.

The programs, once compiled, can be executed in a test environment. Many run time errors may surface which were not uncovered during the Toolset and compilation stages. Even though the programs were operational on MPE/V, these run time errors will deem them inoperable until the problems are addressed. ON SIZE errors and conflicts between FD record size and the actual file size were two of the run time errors which I experienced. These errors are not catastrophic in nature and easily corrected. The only downside is the time lost correcting each and every program that encountered this problem. The HP Response Center can also help by altering the run time parameters during compilation (emulating the old COBOL).
The decision to have a parallel system was totally justified when we reached the "live" testing. All of the run time errors could be located, corrected, recompiled and retested without any user pressure or data integrity concerns. I could not imagine the high degree of panic that

would have been propagated if we had encountered these run time problems in a truly "live environment".

When all the Run Time error are corrected and all the tests are satisfactory, the application, databases, files and all user devices can be moved to the new system. The completed module is launched into a "live status" with a high degree of confidence for success.

One item of importance that I have not mentioned is the need to upgrade all third party software to respective XL versions. Of course, this must be accomplished before any testing can occur. The real surprise regarding this exercise is the cost considerations which I will cover later in this paper. Without the new versions, the transition will be almost impossible.

Cost Considerations

The goal of all managers is to minimize the cost of the upgrade without degrading the end results. Before the budget, for the upgrade, is finalized, there are a few items which must not be overlooked.

The actual hardware configuration must be finalized. Choice of disc drives, memory, DTCs, 3 pins vs DB25, and peripherals are all part of the finalization process. Once the final configuration is attained, comparisons of new versus used equipment can be made. Make sure you are aware of any area where the equipment is at or near capacity. It might be better to purchase excess capacity now then having to re justify the expansion in the near future. HPIB cards, Disc cabinets, DTC cabinets and memory are just a few examples.

The cost for installation and de-installation of the old system should be included in the purchase price of the new system (including the shipping costs). The extra disc and peripherals needed to have a dual system should also be bundled into the overall system cost. This is a great bargaining chip is you are purchasing used equipment.

The cost for migrating third party software can become a major expenditure. In most cases, the charges are not trivial and can not be overlooked. The only software which did not charge for an XL version was ADAGER(someone with a conscience). It seemed that all we received from the upgrade was a new tape and a hefty bill. Make a checklist for each third party vendor detailing costs and lead times for the upgrade tape. It is best to have all of the upgrade tapes before the new system arrives so you can avoid stalling the entire process because of one missing piece of third party software.

The electrical costs for providing the necessary receptacles and service for the new system and redundant peripherals must be budgeted. Many of the power cables/receptacles can be used on the new system. This budget item should not be overwhelming. The only problem could be insufficient power for the redundant system. If you are not anticipating performing the terminal cable switchover internally, any re-cabling and pin-out conversion could become a costly line item. Much of the cabling effort might be done on off hours which could increase the overall cost.

If you are not currently on the HP Responseline support, I would highly recommend returning to this service for at least one full year or until your staff has reached a level of confidence with the XL operating system.

PCI was able to transport our entire system fully tested within a 3 week period and ahead of schedule. The parallel system reduced a potentially disastrous operation into an orderly and anxiety free procedure. We tried to plan for all known contingencies so there were no major surprises. Our success is a testimony of our plan, dedicated staff, very helpful vendors (especially HYPOINT TECHNOLOGY and ATLANTIC TECH SERVICES), long hours of hard work and a bit of good luck.

MPE/iX Disk Space Management

Paper number: 5005
Paul Wang
SolutionSoft Systems, Inc.
2350 Mission College Blvd. Suite 715
Santa Clara, CA 95054
408.988.7378

Abstract

Disk space management is critical to application resiliency, efficient use of disk space and system performance.

Physical disk placements of various applications are crucial to their resiliency. For example, a disk failure may wipe out both a IMAGE/SQL dataset and its corresponding user logging files if they are not separated on different disks. Due to disk fragmentation or dynamic extent allocation, valuable disk space may be wasted. Without proper disk space management, it may force users to purchase more disks prematurely. As the gap widens between CPU and disk speed, balancing I/O across multiple disks becomes increasingly critical to high-end system performance. A single disk containing performance critical files can become the I/O bottleneck.

This paper gives an overview on how MPE/iX file system allocates disk space. It also provides disk space management tips and guidelines on application resiliency, efficient use of disk space and system performance.

MPE/iX Extent Management

An extent is a chunk of contiguous disk space. By definition, an extent is contained solely on one disk. A file may consist of zero, one, or unlimited number of extents. For each extent, the MPE/iX file system determines how much space to allocate (extent size) and which disk it should reside on (extent placement). It will be shown later that extent size and extent placement are key to performance, reliability, and efficient use of disk space.

Please note that the algorithms described below are of MPE/iX release 4.0. There are minor differences between release 4.0 and previous releases.

Extent Placement Algorithms

The extent placement restriction is an attribute of a file. The default is the volume set where the parent directory node for the file resides. For example, the PUB node is the parent directory node for the file named CI.PUB.SYS and the PUB node resides on the system volume set MPEXL_SYSTEM_VOLUME_SET. At file creation time, a user can optionally further restrict it by a device number, a device class, a volume name, or a volume class. The user-specified restriction must be a subset within the domain of the volume set where the parent directory node

resides. It follows that a file cannot span more than one volume set. However, a file can span more than one volume within the set.

The extent placement restriction defines a set of volumes that are allowed to allocate disk space for the file. At extent allocation time, the MPE/iX file system builds a list of allowed volumes based on the file's restriction. The list is sorted by the ratio of the amount of free disk space available per volume and its disk capacity in descending order. The first volume on the list that can satisfy the requested amount is chosen. (We will talk about how the requested amount is determined later.) If no volume on the list can satisfy the requested amount, the largest contiguous disk space on the first volume of the list is allocated. Then, it is up to the requester to make another request if the allocated amount is not enough.

In the long run, the extent placement algorithms will balance disk space allocation across all volumes within volume sets, since the system always tries to allocate disk space from the one having the most free space ratio (allowed by the restriction).

Extent Size Algorithms (non-spreading)

There are three cases where the extent size is determined. Note that the unit of extent size is a page (4 Kbytes). Any calculations of extent size mentioned below are always rounded down to pages.

This section discusses the non-spreading cases. If the extent size is calculated larger than 16 megabytes, it is spread into multiple extents. The details are covered in the next section.

File Creation With Initial Allocation

At file creation time, an initial allocation size may be optionally specified. The following are some popular usages:

a. BUILD command with DISC=keyword (specified by the third sub-parameter)

b. HPFOPEN, FOPEN with the initial allocation option

c. DBUTIL to build a IMAGE/SQL database

The MPE/iX file system always attempts to allocate the initial size in one extent. If a lesser amount is allocated due to a disk space shortage, the system would try to allocate the remaining size in one extent. This process continues until the initial size is satisfied. If disk space runs out during the process, the file creation would fail and all allocated disk space for this file would be returned to the system.

Store XL and COPY Command

For each file Store XL backs up or COPY command copies, it builds a list of contiguous blocks for the file. A contiguous block is a contiguous virtual address range that has disk space associated with it. Each block may contain one or more extents. Note that it is a logically contiguous block and is not necessarily physically contiguous on disk. For example, let's assume a file has ten extents, where each extent size is 4 Kbytes. The first nine extents are logically contiguous and starts from the beginning of the file. The tenth extent is at the end of the file and the file's EOF is 100 Kbytes. Then the contiguous block list for this file has two blocks in it, namely one block from 0 to 36 Kbytes and the other block from 96 to 100 Kbytes.

At restore or copy time, disk space is allocated according to the contiguous block lists. For each block on the list, an attempt is made to allocate a single extent with the same size as the block. If a lesser amount is allocated due to a disk space shortage, the same approach as initial allocation is taken. Finally, data are copied into those ranges and posted to disks. In essence, the list determines the logically contiguous pieces of a file. In this way, "non-existent" data (gaps between extents in sparse files) are not copied.

Because of this contiguous block concept, unallocated areas in files are ignored. Furthermore, multiple extents in an original file can be combined into larger extents at restore or copy time. Assuming enough disk space is available, a logically contiguous file produces one-extent target file after restore or copy. In addition, a sparse file with N gaps produces a target file with N+1 extents.

Extent Fault

Whenever a data access or prefetch request touches a file virtual address that has no disk space associated with it, an extent must be allocated before the operation can go on. This process of dynamic extent allocation is called an extent fault.

The extent size calculation for extent fault is determined by the file limit. If the file limit is less than 64 Kbytes, the extent size is the file limit. If the file limit is between 64 Kbytes and 2 megabytes, the extent size is 64 Kbytes. Finally, if the file limit is more than 2 megabytes, then the extent size is calculated with a formula: minimum [minimum (512 Kbytes, 1/32 of file limit), maximum (64 Kbytes, currently allocated disk space size for the file)]. The currently allocated disk space size is the calculated size and it starts from 64 Kbytes. It is further limited by either 512 Kbytes or one thirty-second of the file limit. For example, the allocation size sequence for a file with file limit of 2 gigabytes is 64 Kbytes, 64 Kbytes, 128 Kbytes, 256 Kbytes, 512 Kbytes and always 512 Kbytes thereafter.

The rationale for this is as follows: For small files (less than 64 Kbytes), the MPE/iX file system makes them single-extent files. Otherwise, extents are allocated incrementally. The system does not want to allocate everything in one shot, since users may not access all of the extent. An upper bound of 512 Kbytes is set to reduce wasted disk space if only a portion of the extent is utilized.

Extent Spreading Algorithms

A very large allocation request is considered to be at least 16 megabytes. Initial allocation, Store XL, and the COPY command can trigger automatic extent spreading if the requested extent size is more than 16 megabytes.

The algorithm works as follows: The MPE/iX file system first establishes a selection criterion for disks. The criterion is 50 percent of the total free disk space (of all allowed disks based on the placement restriction) divided by the number of allowed disks. The system builds an eligible list of all allowed disks that have more free disk space than the selection criteria, sorted in descending order. The spread size is the requested size divided by the number of eligible disks bounded by 1 to 128 megabytes. Finally, the system cycles through the eligible list with the spread size as the requested amount until the remaining size is less than 1.125 times the spread size. From then on, the remaining size is used until the request is satisfied. Please note that the spread size is the preferred size, the minimum size is 256 Kbytes. Some disks might allocate less than the spread size (always greater than or equal to 256 Kbytes) due to a disk space shortage

during the extent spreading process. The process fails if all eligible disks are out of disk space. In this case, all allocated disk space would be returned to the system.

The rationale is as follows: The selection criterion is important for fostering balanced disk space allocation across all volumes within volume sets. If the system always spreads extents across all allowed disks, some disks that are close to capacity would be filled up very quickly. The selection criterion is 50 percent of average free disk space among all allowed disks. The 50 percent leeway is intentional so that disks that have slightly less than the average amount of free disk space are not ignored. The list is sorted in descending order. Therefore, if the number of extents is not a multiple of the number of eligible disks, the disks with more free disk space would be asked to allocate more. The lower bound for the spread size is necessary since the number of eligible disks might be large. The 1.125 times remaining size threshold avoids the case where the last extent is extremely small.

Fill Disk Algorithms

Extent allocations are not always accompanied by a fill-disk operation (initialized with the fill character). The rules are: An extent allocated within the end-of-file (EOF) must be filled. An extent allocated beyond the EOF is not filled. Whenever an EOF is extended, any virgin pages (see below) between the new and old EOFs are filled. A page that has never been posted to disk is called a virgin page. The MPE/iX file system keeps track of virgin ranges for files so that accesses against those ranges never result in I/Os.

Fill-disk operations are blocked I/Os. This can be very expensive for very large extents. For example, it takes 5 to 10 seconds to fill 10 megabytes. Please note that extent allocations for initial allocation, Store XL and COPY command are without fill-disk operations.

Some background information on the fill character: The fill character is an attribute of a file. The file will be initialized with its fill character. The default fill character for ASCII files is blank and for binary files, zero. Optionally, users may specify their own fill character at file creation time.

Application Resiliency Considerations

Disk failures do occur. With application resiliency in mind, we can minimize the negative impact of a disk failure.

Private Volume Set is more Resilient

If the master volume fails, the whole volume set is lost. However, if a member volume fails, users can still access the rest of the volume set. This is always true for private volume sets (also known as user volume sets), since all directory nodes always reside on the master volume. On the other hand, directory nodes may span outside of the master volume for system volume sets. As a result, users may be unable to access some files residing on the good disks due to missing directory nodes! Therefore, a private volume set is more resilient than a system volume set. Users are advised to utilize private volume sets for better application resiliency.

Separate Major Applications and Subsystems

It is advantageous to isolate important applications and subsystems physically from each other. For example, a disk failure may wipe out both a IMAGE/SQL dataset and its corresponding user

logging files if they are not separated on different disks. A disk failure may also wipe out both the order-entry and finance applications.

The safest separation is to put different subsystems on different private volume sets. That way one disk failure will damage at most one subsystem. However, there are cases where subsystems must reside on the same volume set. IMAGE/SQL datasets with Roll-Back Recovery enabled and their user logging files is one example. Note that they can reside on different volume sets if only Roll-Forward recovery is enabled. The file placement restriction is the key to separate files within a volume set. A file can be restricted by a device number, a device class, a volume name, or a volume class. Among them, the device number is the simplest one to use. The volume class is the most flexible method. It can contain one or more volumes, and its creation does not require a reboot as the device class does.

Efficient Use of Disk Space Considerations

Any disk space allocated, but not utilized is WASTED and there are many ways to waste it.

Truncation Files periodically

It is common to have wasted disk space beyond the file EOF due to file system dynamic allocation (extent fault). The potential wastage per file is 4 to 512 Kbytes. The wasted disk space adds up really fast. In my experience, it is not uncommon to see hundreds, or even thousands of sectors to be wasted this way on a production system.

File truncation is probably the easiest and most effective way to reclaim wasted disk space back. It is recommended that users truncate all files periodically. A good way to do it is to incorporate file truncation into nightly batch jobs.

File Limit

For partially written extents in general, the bigger the extent, the bigger the waste. The file limit is the dominant factor for extent size calculations. It is advisable not to inflate the file limit. This is particularly true when it comes to HPFOPEN. The default file limit for HPFOPEN is 2 gigabytes, which tends to generate extents half a megabyte in size! Please note that the default file limit for FOPEN or BUILD is 1023 records.

Initial Allocation

Like the file limit, it is not advisable to inflate the initial allocation size. The size should be in line with what users believe will be consumed in the near future.

The initial allocation is also useful for the frequent case in which the file is most often very small, but occasionally quite large. There is no way to determine the size at file creation time, so users are forced to create it with a very large file limit. As a result, a substantial amount of disk space is wasted due to large extents. The solution is to create the file with a small, reasonable, initial allocation size, which will satisfy most cases.

Disk Array as LDEV 1

Due to I/O channel firmware limitations and boot images migration considerations, currently there are space restrictions for LDEV 1. More specifically, the CIO channel can address only 2

gigabytes and the NIO channel can address only 4 gigabytes as far as LDEV 1 is concerned. The extra space on LDEV 1 (if any) beyond the restriction will be marked "allocated" and it will not be available for use by the users or the system. Please note that there is no space restriction against volumes other than LDEV 1.

As a result, C2254 should NOT be used as LDEV 1, since 1.4 gigabytes would be wasted for the NIO channels and 3.4 gigabytes would be wasted for the CIO channels. The CIO C2252 would waste 0.7 gigabytes, or 700 megabytes when used as LDEV 1. Finally, no disk space is wasted for NIO C2252 when it is used as LDEV 1.

System Performance Considerations

The CPU speed doubles roughly every 18 months. On the other hand, disk speed only improves by 30 percent or so in the same time frame due to its mechanical nature. As the gap widens between CPU and disk speed, reducing the number of I/Os and balancing concurrent I/Os across multiple disks becomes increasingly critical to system performance. The placement and size of extents are the most dominant factors for I/O balancing.

Defragment Disks Periodically

Since extents are of variable sizes as well as are allocated and deallocated dynamically, each contiguous free space tends to become smaller and smaller as the system is used. This situation is called disk fragmentation.

Severe disk fragmentation, where most of the free space gaps are of very small size, causes many problems: Firstly, it degrades system performance. It not only takes more time and resources to allocate and deallocate disk space, but also generates more I/Os. Secondly, it could also produce premature out-of-disk-space errors where enough disk space is actually available. For example, the system may run out of 256 Kbytes contiguous space during a spread file operation. Finally, it could cause system aborts, such as OS update failure, stalled transaction, or system label table expansion failure.

Fortunately, the severe disk fragmentation condition mentioned above takes a long time to develop. Exactly how long depends on each system's disk space usage situation. Users are advised to be more proactive: monitor it and fix it before it becomes a problem.

One good way to monitor the potential problem is to incorporate a monitoring and repairing process into nightly batch jobs. Additionally, if a file truncation operation is also being done, the monitoring and repair process should be executed after the truncation so that the many small pieces of disk space saved by the truncation would be recovered.

Combine Extents for Files

Generally speaking, a large extent is better than many small extents with the same accumulated size. One large extent uses fewer system resources. The system also performs better with a large extent during virtual address-to-disk address translation (for example, servicing a page fault), since fewer extents need to be searched. This is noticeable when servicing a page fault causes still more page faults on the corresponding system translating structures. Furthermore, one large extent discourages disk fragmentation at extent deallocation time.

All physically contiguous files (extents are contiguous on disk) should be combined. This happens a lot due to extent faults and extent placements always go for the same most empty disk.

All program files as well as data files with dominant sequential access patterns are also excellent candidates to be combined. In general, a combined file is better as long as I/Os against the file would not cause a bottleneck. Please note that in the sequential access case, the file system performs automatic "prefetch ahead" and "post behind" for the users so that I/O is not a problem.

Spread Extents for Files

On the other hand, a very large extent (tens of megabytes) with an intensive multi-user random access pattern can be a performance problem. Page faults or prefetches against the same extent are all targeted for a single disk. An I/O cannot be serviced until all higher priority I/Os and previous same-priority I/Os are completed. The problem is worse when a checkpoint (posting all dirty data) is in progress. Tens of checkpoint I/Os are generated continuously and they may also be competing for the same disk.

In this case, it is more advantageous to spread the space evenly across multiple disks. Since multiple I/Os can be serviced by multiple disks concurrently, the throughput and response time are greatly improved. For example, spreading extents of the Debit/Credit benchmark's "Account" dataset across six disks (except the master volume) on a HP 3000 Series 960 with 64 megabytes of memory boosted the throughput by 50 percent compared to the non-spreading case. In fact, the response time criterion (90 percent of transactions are completed within two seconds) cannot be met without spreading the extents! Please note that avoiding the master volume allows us to separate the transaction manager's log I/Os from the database I/Os.

As the gap widens between CPU and disk speed, balancing I/O across multiple disks becomes increasingly critical to high-end system performance. The placement and size of extents for performance-critical files are the most dominant factors for I/O balancing.

Users are advised to always spread their performance-critical files (with intensive multi-user random access pattern). Since hashing and B-tree accesses are by nature random (i.e. 99% of on-line access to IMAGE/SQL or ALLBASE/SQL), performance-critical databases are almost always excellent candidates to be spread.

Use Store XL and COPY Command to Copy Files

FCOPY, DSCOPY, and CM Store copy data by sequentially reading and moving data in blocks. It has many undesirable effects: A fill-disk operation is performed each time an extent is allocated for the target file. Since the copied data will override the fill pattern, fill-disk operations are unnecessary. It degrades response time and doubles the number of I/Os for the target file (fill-disk plus posted data). For sparse files, "non-existent" data (fill pattern) in extent gaps will be copied and cause target files to be fully allocated, wasting disk space. Worse yet, source files will also no longer be sparse files after copying, and will be fully allocated due to reads! Finally, the algorithm tends to generate many extents for target files. These extents are allocated dynamically as data are copied.

There are many ways to copy files. In particular, Store XL and the COPY command generate desirable extent distributions, since they use both the contiguous block and automatic extent spreading algorithms. On the other hand, FCOPY, DSCOPY, and CM Store generate many smaller extents. They not only foster disk fragmentation but also perform more slowly. This is especially noticeable with sparse files.

Of course, those subsystems have unique features that are not provided by Store XL and the COPY command. Users should continue to use those subsystems when it is necessary, but using Store XL and the COPY command will be both faster and better for your disk usage.

Write Beyond EOF Before Extending It

If the disk space is not pre-allocated, then it is advisable first to write data beyond the file EOF, and to extend the EOF (FCONTROL or FCLOSE) only after the writing is done. This way, costly fill-disk operations are eliminated. This is particularly valuable with initial file loading.

The potential downside of this approach is that data beyond the EOF might be lost if the system crashes before the EOF is extended. However, this is a very small possibility, so we recommend using the process.

Use Initial Allocation Option

Initial allocation is always more efficient than allocating dynamically later and avoids costly fill-disk operations. If the disk space requirements are known in advance, it is a good practice to initially allocate the space.

Disk Array

With the advent of disk arrays (C2252, C2254), disk capacity grows dramatically. C2252's capacity is 2.7 gigabytes, C2254's capacity is 5.4 gigabytes, and the next generation could go as high as 8 gigabytes.

To balance the I/O subsystem, it is recommended that users not mix disks of greatly differing capacities within a volume set. The reason is to avoid disk space allocation domination by big disks, which could become an I/O bottleneck and would foster unbalanced I/O by nature.

Consolidating a volume set from many disks into few C2254s deserves special attention. Even though a C2254 has much better seek time and transfer time than "conventional" disks, the number of disks within a volume set could decrease six to tenfold. As a result, if the environment requires heavy random-access I/Os, a few C2254s might perform more poorly. For example, consolidating 12 C2202s with I/O load averaging 10 I/Os per second per disk into 2 C2254s generates unreasonable I/O demand of 60 I/Os per second per disk! This demand is obviously beyond the C2254's capacity and the system response time suffers as a result. In such a case, consolidating into 4 C2252s might be a better option.

Conclusions

MPE/iX file extent management is fair, flexible and tuned to high-end performance. A file can have zero, one or unlimited number of extents. Extents can be allocated dynamically as needed. These factors make system disk space utilization more efficient. Extent placement is determined so as to balance disk space allocation for the whole system. Users can restrict the extent placement by a device number, a device class, a volume name, or a volume class. In addition, the system automatically spreads extents across multiple disks for very large allocation requests. This is particularly critical to high-end system performance. As the gap widens between CPU and disk speed, balancing I/O among multiple disks is the key to performance.

On the other hand, MPE/iX file extent management is also very general and "one size fit all." Sometimes, users may do a better job with application specific information. In addition, disk fragmentation and wasted disk space beyond EOF could become problems if they are not attended for a long time. Users are advised to follow above tips and guidelines on application resiliency, efficient use of disk space and system performance to optimize disk space management.

Bibliographic

Paul Wang is the president of SolutionSoft Systems, Inc. He is a consultant, specialized in transaction management, system performance, file system internals, data base, and On-line transaction processing. Previously, he was an internal architect of transaction processing in HP's Core MPE/iX lab.

BUSINESS RESUMPTION - A TOTAL QUALITY MANAGEMENT APPROACH

Paper # 5601 by Claude Brazell, U.S. Program Manager
Hewlett-Packard Company
15815 SE 37th Street
Bellevue, WA 98006
(206) 644-3362

This morning, over coffee at my hotel, I pulled a big red marker and a USA Today out of my briefcase ... scanned the front section for every story that had the potential of interrupting a business's normal operations and circled everyone I found.

As you can see, a lot went on yesterday.

And this is a national newspaper. It only covers major news. Consider what would happen if you did the same with your own local paper. How many more stories would you find?

And also consider all of the things that can happen that never make the newspapers -- computer viruses, power failures, disgruntled employees.

Considering all of this makes me wonder why fewer than 50 percent of the Fortune 500 companies in the United States have a formal business resumption plan in place -- and many of them only because some auditor or insurance company or government agency mandates it. And my intuition tells me that percentage is much, much lower when you consider all businesses -- small, medium and large.

If I sound like a pessimist, forgive me. I'm actually an optimist -- the kind of person who takes a frying pan on a fishing trip. But I'm also a realist; I take along extra provisions on my fishing trips -- just in case. I make sure that I'm prepared.

Being prepared for business interruptions is half my topic today; the other half is how to apply the principles of total quality management -- or TQM -- to the business resumption process as a way to make sure you are prepared.

Over the next few minutes, I will:

- Give you a general overview of the business resumption process, and how TQM applies;
- Suggest how you can sell your management on investing in a business resumption process;

- Outline a four step process for developing your business resumption process;
- And finally, share some of the direct and indirect business benefits of having a rehearsed, robust business resumption plan in place.

Let me start with an overview.

Like its predecessor, disaster recovery, business recovery is a discipline and a process that addresses how to prevent business interruptions from harming your enterprise. That it falls to the information systems organization to organize and manage this process is an accident of history and an admission of the critical importance that IT plays in modern organizations.

But there are some important differences between the two. Disaster recovery focuses on restoring systems, on making the necessary MIPs available; business resumption, looks beyond the MIPs and focuses on restoring key business processes, including the systems that support the processes.

Disaster recovery traditionally has focused on centralized data center resources; business resumption takes into account all the systems responsible for ensuring the continuance of key business processes.

That can be particularly complex in today's computing environments. Although mainframe computers still play a crucial role as a central repository for much of a company's data, the applications that use that data often are running on distributed systems outside the data center -- often in a geographic area hundreds of miles removed from the data center.

Disaster recovery traditionally has been seen as a cousin of insurance in that it implied a means of compensating for the loss of data, applications and/or processing power because of some significant natural or man-made event. Business resumption is broader, encompassing any event that interrupts critical business activities -- inside and outside the data center.

Why the evolution?

Quite simply, your information management group is no longer just a support organization for a company. It is weaved integrally and strategically into the enterprise's ability to do business.

If you question this, just consider what would happen to your organization today if an order entry system was unavailable for an extended period of time. Could you reasonably expect to return to a manual system of writing orders? And even if you could, what would that do to your ability to deliver products and services vis-à-vis your competitors?

One of our customers is a medium-sized distributor of high-volume, low-margin products in the Midwest. If his system goes down and he can't write and get orders to the warehouse electronically for shipment, he loses $1 million a day. In just three days, he's out of business.

For this customer, it isn't the probability of a major disaster that drove him to invest in business resumption. It was the issue of business survival.

So what constitutes a business interruption? It's in our nature to limit our thinking to events that have obvious, significant, widespread consequences. And lord knows we've had our fill of actual disasters recently:

... Hurricanes Andrew and Hugo
... the Chicago River flood
... the San Francisco earthquake
... the Los Angeles riot and subsequent fires
... the World Trade Center bombing
... and of course, the recent earthquake in Los Angeles.

But these are the obvious. What about the not so obvious events that can interrupt your business?

... the storm that knocks out a whole power grid and strands employees at home
... the disgruntled employee who steals important data or plants a virus that damages critical applications
... a union picket line that blocks key employees from entering your building

And what about the potential for an upside event? Have you ever considered what the impact would be on your systems if your sales force meets 200 percent of its quota? Are your business processes robust enough to handle unexpected success on a large scale?

These are not disasters in the classic sense, but they are events that can have disastrous consequences for your business.

In thinking about business resumption, it's important to broaden your definition of a disaster: It is any event that interrupts the performance of critical processes for such a period of time and to such a degree that it inhibits the enterprise's ability to deliver its goods and services to customers.

With that definition in mind, let me now turn to the principles of TQM and how

they apply to business resumption. I think you will find that there is great commonalty between the two.

First, both require a focus on the customer. The classic measurement of quality is customer satisfaction. If we extend the definition that I've just given, we see that business resumption is focused on ensuring the enterprise's ability to deliver its goods and services to customers. When I refer to customers here, I mean the customers of the enterprise. There is a danger when the IT community begins to think of other departments or functions in the organization as a customer in that it detracts from the goal of the organization.

In a classic customer relationship, the customer has a choice. And it is the mission of the provider to optimize its capability and deliver the highest possible service at the lowest possible cost.

In the enterprise the people who rely on IT services have no choice. The notion of an internal customer relationship is false. There is none. Instead, the enterprise is a family, and each organization is a sibling dedicated to performing an activity that contributes to the welfare of the family. In this sense, the mission of the IT organization is to provide reliable, consistent and predictable information management support.

Period.

And the business resumption process strikes directly at the heart of the enterprises' ability to perform that mission.

Second, both are processes that can be measured and continually improved. The classic illustration in the quality world is the Deming Wheel, or the (PDCA) Plan-Do-Check-Act Cycle. That same illustration applies to the business resumption cycle: develop a plan, rehearse it, measure the results, take corrective action, and incorporate it into the plan.

Third, both assume that decisions are driven by data. Intuition and luck are important elements in any business, but both the quality and business resumption processes are data driven. Improvements are made based on data that identifies weakness' in our processes.

Finally, both require a demonstrable commitment by the management team and by the people closest to the process to begin and to sustain. Just as it is a fallacy to believe that an outside guru can come into an enterprise and improve quality processes, it is equally a fallacy to believe that the business resumption process can be improved from the outside. There is no manna from heaven.

So far, I've given you an overview of the importance of having a business resumption process and explained the importance of the TQM process. I'm now going to turn to the more nuts and bolts -- getting your management team on board and moving forward with your business resumption process.

I've heard many complaints from IT managers that their managers are insensitive to the needs for a business resumption process. Even if they succeed in convincing management -- which is not a sure thing -- it is a reluctant commitment. Many management teams are resigned to the concept, rather than enthusiastic supporters.

Management, however, is not the villain. The real villain is perception. Let me give you a simple, hypothetical example. An IT manager in Northern California goes to the management team and says, "If we have an earthquake, it could knock out our systems. We need to invest in a business resumption process to make sure that we have alternative processing power."

Now right after the 1989 earthquake, management interest would have been very high. It probably would have been a fairly easy sell. But memories have a way of fading over time. Now, five years after the quake, they are just as likely to ask, "How much?" and respond, "I'll take the chance."

The problem with this approach is that it focuses only on the insurance aspect of the business resumption process. Management sees only the payment of a premium, and no value unless the worst case actually happens.
So what's a better approach? Well, let's start off with the mission of the IT organization that I gave above -- to provide reliable, consistent and predictable information management support to the organization.

That mission -- not the potential threat of an earthquake or a fire or a flood -- should be the focus of the IT manager's request. He or she should tell the management team, "The mission of my organization is to provide reliable, consistent data processing to the functional departments. I cannot ensure that this mission will be accomplished without a business recovery process."

That's the first argument.

The second is the benefits that accrue from investing in a business recovery process. There is, of course, a direct benefit to the enterprise. If an event occurs that interrupts critical business processes, the enterprise is better prepared to overcome the interruption and resume those processes.

But there are indirect benefits as well. Frequently, improvements in procedures, policies, processes and systems are identified in rehearsals that can be leveraged into normal operating circumstances. Let me give you two examples.

One of our customers on the East Coast used an earlier generation of processor to run a key application at the data center. On that system, the processor took over three hours to run the application. That wasn't a problem, however. It was just assumed that three hours was okay.

During a rehearsal, however, the application was placed on a fully compatible, but later generation system. The result was that the application ran in less than one-third the time -- 53 minutes. Without the rehearsal, the customer would never have discovered the potential for performance improvement. There was no reason to think of it. As a result of the rehearsal, however, the customer included a system upgrade in the next year's capital expense budget.

In another instance, a senior vice president learned that his company did not have near the cross-functional communication that it needed to compete. During the rehearsal, it became painfully evident that there was still a hierarchical communications structure -- everything went up, then down. That discovery alone was worth the subscription price to him, ignoring any of the direct rehearsal benefits.

The final argument is the cost of survival.

If you were to look at a cross-section of companies across all industries, our research suggests that companies are spending about 1 percent of their IT operational budget on business resumption planning and rehearsals. Now compare that cost with the financial loss to the company if a critical business process is interrupted, and how long it would take for the company to suffer long-term or irreversible damage. In today's globally competitive world, one thing is certain: If one company can't deliver goods and services, there are others who are ready to step in.

To recap, the first approach focuses on the possibility of an event happening and lets the CEO decide whether to play roulette. The second focuses on the reason for having an IT department in the first place, sweetens the pot with the potential of added value and treats business resumption as a survival issue. I'll leave it to you to decide which is the more compelling argument.

Even before management approval is secured, you can begin your business resumption plan. In many cases, what's learned during data gathering can be used to help convince management. In the planning stage -- which again maps to the Plan step of the PDCA Cycle -- is the time to establish your objectives and how you will achieve them.

Quite often, it is helpful to enlist the help of a certified disaster planner, a person trained in asking the right questions such as what are the key business processes,

who are the people and what are the systems that support those processes, and how long can the company go without having those processes available. Above all, the planning stage is the opportunity to document your processes -- get them down on paper.

Also during the planning stage is when you determine what level of business resumption support you will need. At the furthest end of the spectrum is total redundancy, which obviously is the most expensive proposition. At the opposite end is doing nothing -- which can be an alternative for those companies who can return to a manual, paper-based system of business in an emergency.

In between, however, are the two more common forms of support: cold sites and hot sites. A cold site is space that is available to the company in the event of a business interruption. It typically will have the infrastructure for power and for communications, but it does not have the systems in place. When a disaster is declared, the company will need to work with their systems vendors to deliver the systems to the cold site, a process that can take several days if not weeks. In addition to the time it takes to bring a cold site up to speed, it also does not afford the opportunity for rehearsal.

For most companies, it is the hot site that provides the most cost-effective approach when balanced against quantifiable and qualifiable goals. A hot site has the systems required to run your key processes on-line and ready for your use. It is a proactive concept. It is testable. And it results in demonstrable organizational improvement.

In selecting a hot-site partner, there are several considerations. First, look at the potential vendor's program philosophy, financial resources and experience in the business recovery industry. Second, you should consider:

- the design of the facilities, including such things as utility and communications support and backup power resources;
- the capabilities of their people, especially if you company uses advanced technologies for critical applications;
- the security for your data and your people;
- and the ratio of customers to the available systems, especially focusing on the concentration of customers in your area.

This last point cannot be underestimated because it is the paramount indicator of your odds of both rehearsing at a frequency that meets your requirements and recovering successfully from any unanticipated interruption. We attempt to keep our ratio of systems to approximately 30 to 1; there are others who support as many as 100 to 1.

The third consideration is their approach to rehearsals. What types of rehearsals do they employ? How difficult is it to secure rehearsal time? How flexible are they in supporting your requirements during rehearsals? How long is their rehearsal day? What technical and operational support is available to assist your during your rehearsal?

From this and other information, it is possible to form a picture of how well the promises match the reality based on real data. References are important, but remember that everyone has at least a few good references.

Only after all other considerations should the issue of cost be examined. The large and principal focus should be on supporting your requirements. If they can't, your investment is wasted -- no matter how low the cost. Typically, the costs that you should be examining include:

- The monthly participation fee
- The length of the contractual commitment
- Any associated cancellation fees
- Testing/rehearsal fees
- Declaration fees
- Daily usage fees
- Contingency planning fees
- Technical and operating support fees
- Periodic cost increases
- Travel and lodging costs for you and your people

When your plan is completed and you've selected your hot site, it is imperative to schedule a rehearsal, which maps to the Do Step in the PDCA Cycle. Rehearsals are your opportunity to gather data on how well your business resumption process works. No matter how critical your thinking and exhaustive your planning, you will miss something -- and that's OK. That's precisely why you rehearse.

One of our East Coast clients, for example, had a top-notch plan. Within hours of declaring a test disaster, they had their telephone lines up and their people on site -- ready to resume business. There was just one little problem. While great care was taken in securing a secondary phone vendor and a backup PBX, no one thought to get handsets for the people that would be moving into the hot site location.

The key to a good rehearsal is well-defined list of rehearsal objectives with built-in metrics that allow you to measure your strengths and weaknesses. These rehearsal objectives should be a distillation of the most important steps that you need to take to get critical business processes up and running.

Also during the rehearsal, consider retaining outside monitors to record your activities objectively. Often, people fail to recognize important steps in a process because of their own familiarity; certain acts become second nature to them. That second nature can be a significant weakness if a true disaster strikes and that person is not a member of the business resumption team.

Most companies should plan on rehearsing their business resumption plan twice a year at the hot site. If that seems like a lot, consider all of the human and technical changes that occur in six months in your company.

o From a human perspective, people are hired, promoted or transferred to a new location, or they move their household and have new telephone numbers and addresses.
o From a technical perspective, new products, services, systems and applications are added, and old ones retired.

Any one of these changes requires a change in your plan. In fact, one thing that most companies learn during rehearsal is that human factors more often prove to be the weak link rather than technical factors.

After the rehearsal, you're ready to move into the third and fourth phases of the PDCA Cycle -- the Check and Act phases. Here is where you review all of the data gathered during the rehearsal, identify where the weaknesses are in your business resumption plan and determine what corrective actions should be taken. In the initial rehearsal, that list is likely to be very long, but will become shorter in subsequent rehearsals. It's highly unlikely that you will reach a point, however, where some improvement cannot be identified. In fact, if you do, then you might want to question the rigor of your rehearsal. Here, too, is where the procedural, policy and process improvements will be uncovered that apply to your business under normal circumstances.

The most important thing to note is that these four phases are a continuum -- not a one-time activity. What you learn in each cycle should be documented in your plan, and the cycle begun anew.

As you can see, the process of developing and sustaining a business resumption program maps very closely to the traditional principles of TQM. In fact, a business plan is very much a manifestation that quality is taking hold in your organization.

Of course, TQM and a business resumption plan is no guarantee that your company will prosper or survive a disaster or other major interruption. There is, however, significant evidence that the principles of TQM contribute to prosperity and survival.

A 1991 report from the Government Accounting Office reported: "Companies that adopted quality management practices experienced an overall improvement in corporate performance. In nearly all cases, companies that used TQM practices achieved better employee relations, higher productivity, greater customer satisfaction, increased market share and improved profitability."

The business resumption process is not an admission of pessimism. It is an acknowledgment, however, that business interruptions can and do happen -- not just to other companies, but to your company as well.

Like Calvin Coolidge, our 30th President, once said, "If you see 10 troubles coming down the road, nine will run into a ditch before they reach you." Working in business resumption has taught me, however, that the one remaining trouble can knock you -- and your business -- on your derriere if you are not prepared.

THE END

Paper 5604

The IT Operations Service Partnership

Kyle S. Adler, Brenda Vathauer
Hewlett-Packard Company
100 Mayfield Avenue
Mountain View, CA 94043
(415) 691-7376

Introduction

Today more than ever before, Information Technology (IT) departments are faced with challenges in providing effective operations. A host of new pressures is increasing the IT workload, including user expectations of increased service levels, migration to open systems, and client/server implementations. Simultaneously, a set of different factors is serving to reduce the resources at the IT manager's disposal; such facts of life in the 1990s as budget reductions and caps, fixed or shrinking headcounts, and data center consolidation can siphon off IT's human and capital means to attack the new challenges. What is resulting is a fundamental paradox: How can IT manage more with fewer resources?

Many organizations are finding that partnering with outside service providers allows their IT management to do more with less. IT partnership with outside vendors is by no means new, and indeed, very few if any IT departments accomplish all tasks using in-house resources alone. However, the new challenges facing IT departments suggest new ways of partnering, to the mutual benefit of IT and the outside service provider.

This paper will suggest ways to make the selective outsourcing partnership work. We will explore the trends contributing to the new challenges in order to find some of the daily operational tasks that can be shared with an outside service provider, examine the criteria to consider--including financial analysis--when deciding which tasks to selectively outsource, and outline proven methods to ensure a win-win partnership.

Why is the Job Becoming Harder?

In the 1980s an MIS manager had to perform very nearly the same tasks as today. Back then as now, problems had to be detected, isolated, and resolved; jobs had to be scheduled and executed; data backed up and recovered; security maintained; system and network configurations managed; and so on. Furthermore, most of the types of computing equipment--mainframes, minicomputers, PCs, workstations, LANs, WANs, and peripherals--found in organizations today were also present a decade

ago, albeit used in very different ways. So what has changed in IT departments over the last decade to cause such a crunch on resources? Let us examine first the trends contributing to the reduction of IT resources, followed by a look at why the workload is growing at an exponential pace. Together, these trends create the paradox of expecting to do more with less.

Factors Reducing IT Resources

The buzzwords are peppered through business headlines in every major journal: "downsizing", "rightsizing", "consolidation", "global competition". Corporate downsizing is a business reality in the 1990s. Because IT represents a large and growing portion of most organizations' budget, CFOs often look to their IT department first when deciding where to slash budgets. Under pressure that often begins as high as the CEO and Board of Directors, CFOs and CIOs are demanding greater return on the company's IT investment. More often than not, this translates into budget constraints or even reductions.

It is no secret that the biggest share of costs in most IT departments goes to human resources. After all, information technology is an industry characterized by change, and the need to hire and retain staff who are current on the technology is a very expensive prospect. It should therefore come as no surprise that when IT budgets are capped or cut, fixed or shrinking headcounts will be the result.

Another common method being used to enhance IT's ROI is data center consolidation. However, if not accompanied by other organizational and operational changes, consolidation can be a costly mistake. Massing more machines into fewer data centers is increasingly seen by organizations as a way to gain greater operating efficiencies. Indeed, centralization of systems and systems management resources may lead to economies of scale, but this is not always the case. If consolidation of systems and data center staff is simply mandated, but is not accompanied by the creation of new processes and tools for central or remote systems management, it is unlikely that greater efficiency will be achieved.

Factors Increasing IT Workload

A number of trends are contributing to IT's increased workload. Some of these trends are caused by business changes, and some by technological changes. Regardless of the root causes, IT must rise to the new challenges while coping with the downsizing pressures already examined.

The first trend is growing pressure from end users for IT to provide increased service levels. Part of the pressure stems from new business

processes: in today's more competitive markets, companies must bring products to market faster, which in turn requires engineering, marketing, finance, and manufacturing departments to re-engineer their core processes. The end result is that end users expect more from IT, specifically higher uptime, broader access to data, reduced application response time, faster problem resolution, and increased flexibility. Another contributor to end users' rising expectations is technological: distributed (and especially client/server) computing gives users unprecedented access to data throughout the organization, as well as more control over their personal use of the data. Having tasted the fruits of desktop access to information, users find their appetites growing and expect to be fed a steady diet of data from more and more remote and exotic locales within the company (and often beyond).

Migration from proprietary to open systems is a large factor adding to IT's workload. Because the benefits of open systems derive from vendor independence, IT departments now find they have many, many more vendors to manage than in the years of proprietary systems. While this multivendor proliferation may yield the benefit of favorable price/performance points, it brings with it a multitude of management problems. With more platforms to manage, the systems manager and operators must learn new operating systems and systems management tools. It is difficult to hire staff trained in the new technologies and equally hard to keep existing staff current. Because all of these platforms and protocols must coexist in your environment, a host of interoperability headaches usually crop up.

Another technological trend, closely related to and building on open systems migration, is client/server implementation. As organizations implement client/server architectures based on open systems, they experiences the advantages of increased vendor choice, better price/performance ratios, and rapid access to information by end users. However, many are finding that the added burden of controlling and managing the distributed LANs, servers, and clients is a hidden cost. The complexity of the distributed environment is heightened by interoperability problems, multivendor finger-pointing, decentralized purchasing and change management, and the technological newness and complexity of the architecture. Proliferation of interconnected LANs forces MIS staff to learn new protocols and tools. The distributed nature of the architecture creates many more points of failure. Because end users gain greater access to data, their demands increase. Desktop machines and distributed servers fall into the expanding realm of the systems manager. The existing larger machines in the data center now take on a new role as servers. Finally, IT has access to very few tools and must use new and undocumented processes to manage client/server environments.

The Net Effect

In summary, then, IT departments must work harder today because, while the fundamental tasks and equipment are largely the same as in the 1980s, business conditions and technological factors are changing the way the equipment is deployed and in turn, the way operations tasks must be performed. The increased burdens placed on IT by rising end user expectations, open systems migration, and client/server implementations, coupled with the harsh financial realities of doing business in the '90s, have the net effect of squeezing IT's internal resources. The remainder of this paper will examine partnering with an outside service provider to maximize the effectiveness of your internal IT resources.

Partnering Decision Criteria

Outsourcing is a Continuum

There is nothing new or radical about the concept of IT departments' partnering with outside service providers. Now as in the past, very few companies do everything in-house. An example of a very common partnership is the nearly universal practice of purchasing hardware service contracts from the equipment vendor or an independent service organization (ISO). A very few IT departments have historically attempted to provide all maintenance and support via in-house resources and have used outside vendors only for parts and occasional cooperative support.

At the other extreme end of the spectrum lies total outsourcing, often delivered under the auspices of a facilities management (FM) contract. The media hype of the last five years to the contrary, total outsourcing is seen today by many as having fundamental pitfalls. In particular, companies that have signed lengthy (often 10-year), comprehensive FM contracts often cite the loss of control, difficulty of enacting change, and lack of vendor understanding of the client's core business as reasons they later regretted the decision. As a result, many big FM deals have gone sour, and most industry consultants expect fewer FM "mega-deals" to be signed in the future.

Outsourcing in the past has often been viewed as an all-or-nothing prospect. However, in between the extremes of no partnership and total outsourcing lies a spectrum of partnering possibilities. The difficult part is deciding which tasks are best done in-house versus by an outside firm.

While the final determination of which tasks to turn over to a service provider is yours, it may be helpful to approach the decision using the following criteria:

- Competitive advantage
- Financial analysis
- Security
- Change

Competitive Advantage

For most effective use of in-house resources, IT departments should be focused on building a competitive advantage for their company's business based on deploying technology wisely. Nobody knows a company's business expectations from–and changing needs for–technology better than the in-house MIS department. Therefore, those tasks seen as essential to building and sustaining a competitive advantage are best retained in-house, while those daily operational tasks that are routine or cumbersome are good candidates for turning over to an outside service provider.

Financial Analysis

It makes good business sense to outsource when the service partner can provide a task more efficiently than via in-house resources. Faced with the economics of a make-versus-buy decision, few companies would choose to build their own computing hardware, program their own operating systems, or for that matter, provide their own basic hardware maintenance. However, the range of tasks involved in IT operations require a more subtle evaluation of make-versus-buy. Fundamentally, there are only three ways to reduce the cost of operations, namely centralization (economies of scale), automation (routinize repetitive tasks), and elimination of tasks (management by exception). An outside service provider may already have achieved the critical mass required for operating efficiencies via centralization. It is also possible that an outside service provider may have developed the tools and processes needed to automate and eliminate certain operational tasks.

It therefore becomes extremely important to ask potential partners to demonstrate how they will achieve costs savings. If a service firm's justification is solely that they will implement off-the-shelf systems management tools, you may be wise to look elsewhere. For example, you may be able to purchase and use Maestro to automate job scheduling as well as they can.

In performing financial analysis, don't make the mistake of comparing apples and oranges. In order to accurately compare in-house to external costs, it is critical to quantify all of your internal costs. Include in your analysis such "hidden" costs as users' lost productivity from downtime, part-time departmental LAN administrators, and hiring and training qualified staff. By sharing your in-house cost data with a trusted service partner, they can help quantify these hidden expenses.

Security and Data Confidentiality

No analysis of which functions to outsource would be complete without thorough consideration of the security of the environment and the confidentiality of the sensitive data on the systems. A common assumption is that using an outside company to provide IT operations brings an increased exposure to security problems. Not necessarily so. The use of an outside service provider, in itself, does not increase risk. Often, it is the methods that will be employed for systems management that are suspect. For instance, if an outside service provider suggests centralizing a number of systems to economize on management costs, the underlying cause of concern might be that users will require remote access over a network where previously they used dedicated lines. The same concerns would need to be faced if in-house IT chose to consolidate resources. Another concern might be that the SM provider plans to use remote tools to manage the distributed systems: however, remote management will usually save money and may be extremely secure if the right tools and processes are employed. Even moving systems off-site to the service company's premises can be as or more secure than operating them on-site, depending once again on the quality of the service firm's tools and processes.

The bottom line is that using an outside vendor can actually increase the security of operations by employing known processes and advanced tools. But ask the necessary questions before you choose the service firm to be convinced that all reasonable steps will be taken to ensure security.

Taking Advantage of Change

A concept that is gaining considerable attention today is that of transitional outsourcing. One of the pitfalls of total outsourcing is that once the long-term contract is signed, it becomes very difficult to enact change. For example, suppose that a pharmaceutical company signs a 10-year data center outsourcing contract and, two years later, decides to roll the legacy systems over to a client/server implementation. Under the terms of the FM agreement it would be very hard to roll over the platform, and to avoid penalties the company would very likely have to forego

implementing a change that could have given them a significant technological advantage.

Transitional outsourcing avoids this problem by structuring the contract with start and end dates chosen to coincide with a period of change. For example, suppose the same pharmaceutical company decides they wish to roll over their legacy systems to client/server over the next three years. They prudently sign a three-year transitional outsourcing contract. At the end of the contract life, they have a fully implemented client/server environment, understand how to manage the new architecture, and have enjoyed the benefits of a predictable annual cost to enact the change.

Transitional outsourcing can take on one of three forms:
1. Leverage the service firm's expertise to jump-start new technology.
2. Outsource the management of the legacy systems to focus in-house resources on learning the new platforms.
3. A combination of (1) and (2), above.

Ensuring a Win-Win Partnership

Having used the above criteria to decide which operations tasks to turn over to a service partner, the IT department's next job is to set up a partnership that will be productive and beneficial to both parties. This process entails two stages: selecting a partner or partners, and structuring the contract for mutual gain. Fortunately, neither of these tasks is very difficult if the negotiation is handled in a collaborative rather than competitive manner. Bear in mind that you have selected where your company should be on the IT operations service partnership continuum; now is the time to implement your plan for maximum leverage to both organizations.

Vendor Selection Criteria

As with entering into any key partnership, the choice of vendor is critical in making the selective outsourcing partnership a success. Industry consultants have conducted a number of studies asking IT organizations what criteria are important to them in selecting a systems management firm. The following criteria often emerge as among the most important.

- **Technical Expertise**: Whether an organization plans to outsource the management of their data center, WAN/voice networks, or client/server environment, they must select a partner with a demonstrated command of the requisite skills. In order to get the data to evaluate potential partners, ask the vendor and their references about their tools, people, and processes. If possible, ask to see the facility from which they manage customers' systems. Get

specific information regarding the technical environment you plan to entrust to the vendor. For example, a vendor's mainframe FM experience is of no use to you if your organization plans to ask them to manage your client/server environment.

- **Cost**: Everything else being equal, it is important to select a vendor on the basis of price. However, don't cut corners by emphasizing price to the exclusion of the other factors outlined in this section. More important than the dollar amount of the advertised cost savings is an understanding of how the vendor will achieve the cost savings. Remember that the vendor, unlike most in-house IT departments, is operated as a profit center and must cover both overhead and profit margin. So, if the vendor tells you they can save you money, they must demonstrate how they plan to do so while still meeting their own financial objectives. Similarly, beware of "lowball" bids: if a vendor bids a much lower cost than others, find out why. Perhaps the level of service will not be as high or the other terms of the contract are more restrictive.

- **Size and Financial Stability**: The selected service partner must have the financial resources to retain the working capital required to meet your needs. A track record of financial stability (and the bond rating to demonstrate it) is helpful. Particularly if you plan to sell some of your capital assets to the service provider, they must be fiscally sound.

- **Geographic Coverage**: Your partner must be where you are in order to manage your systems. Even if remote management will be used for your distributed systems, what happens when on-site service is required? A potential partner should be able to show you the process they will use to provide the service levels guaranteed in the contract.

- **Working Relationship**: Bearing in mind that you are selecting a key partner to whom you will entrust the management of your critical systems, it is important that the partner be easy to work with. Have your past dealings with the vendor been satisfactory? Have they demonstrated a commitment to quality? Do you have the mutual trust to ensure that contract or performance concerns will be handled fairly?

Contract Issues

Having selected the vendor, it is now necessary to develop the terms of the contract. While each vendor will have a standard set of terms and conditions, your systems management needs will necessarily differ from

those of other organizations, so you should insist on enough flexibility from the vendor to meet your key needs. The following are a few important points to consider before signing any outsourcing contract.

- **Clearly Specify Required Service Levels**: Any outsourcing contract should include a service level agreement (SLA) which clearly specifies the level of service the vendor must provide in order to meet your business needs. Be precise in defining such terms as "response time" (which two events define the start and end of the period being measured?) and "uptime" (of what?, measured how?).

- **Insist on Periodic Reporting to Verify Performance**: The service firm should provide you with periodic (perhaps monthly) reports measuring the performance of the systems, networks, and applications being managed. The reports should be tailored to your organization's needs and should specifically measure performance in terms defined in the SLA.

- **Specify Penalty/Escape Clauses for Non-Performance**: No contract is better than its enforceability, so spell out what the penalties will be if the service provider cannot meet the terms of the SLA. You will already have selected a partner you trust, so most performance concerns can be handled informally; however, if a problem is major and persistent, these clauses will give you important leverage.

- **Choose a Contract Duration that Meets Both Parties' Needs**: Your partner needs to turn a profit. You desire maximum flexibility. Because selective outsourcing is a partnership, both parties must work together to choose terms acceptable to both. This is especially important in the area of contract duration, because too long a contract life may erode the client's ability to remain responsive, while too short a life will not allow the vendor to make a profit. Increasingly, outsourcing contracts with two-year to five-year durations are being structured as mutually satisfactory.

- **Allow Flexibility to Change Contract when Your Business and/or Technical Needs Change**: Closely related to the last point, the contract should be flexible enough to allow your organization to take advantage of new technologies or processes. In fact, one of the primary advantages of partnering with an outside service provider is that your partner can help you stay current with new technologies. Indeed, a contract can be structured so as to have the service partner roll out new technologies for the client on an ongoing basis and for a fixed price, thus ensuring that you will always have access to the latest technology.

In Closing

Just as every organization's IT needs are unique, there can be no single right IT operations solution for everyone. This paper has argued that few companies could or should try to manage every aspect of their IT environment with just in-house resources. Similarly, most organizations would be wise not to outsource every aspect of their IT management in a long-term, FM-type contract. In between these extremes is a spectrum, and only you know where on the spectrum lies your optimal balance between in-house and external resources.

We observed a number of trends taking place over the last decade, the net effect being to squeeze most organizations' in-house IT resources, often to near the breaking point. Partnering with an outside service firm to manage some of your environment lets you do more with less while retaining full control over your systems. The methods outlined here have been helpful to organizations looking to selectively outsource for maximum leverage of their internal resources.

An interesting transformation has occurred since the 1970s and 1980s. In the past, the most strategic decision an IT manager had to make was the choice of a hardware vendor. As important as the choice of hardware remains today, it is equally strategic to choose your IT operations partner with care, as you will be working with them daily.

Paper Number 5671

"HP 3000 Consolidation: A Case Study"

William G. Nadzak
Hewlett Packard Program Manager

FMC Corporation
Corporate Information Resources
1460 Round Table Drive
Dallas, TX 75356-9630

idation: A Case Study
ion Summary

ore Consolidation
Project - Timeline
Consolidation
rdware Configuration

5671 - 1

F

993 Total
22,000 Em
Operates and
Mines in 1
• Participat Selected
Segments
– Industr
» Worhium
Che

5671 - 3

FMC: The Corporation

- Performance Chemicals
  - » World's Largest Producer of Carrageenan
- Defense Systems
  - » Bradley Fighting Vehicle
- Precious Metals
  - » Gold and Silver
- Machinery and Equipment
  - » Citrus Machinery

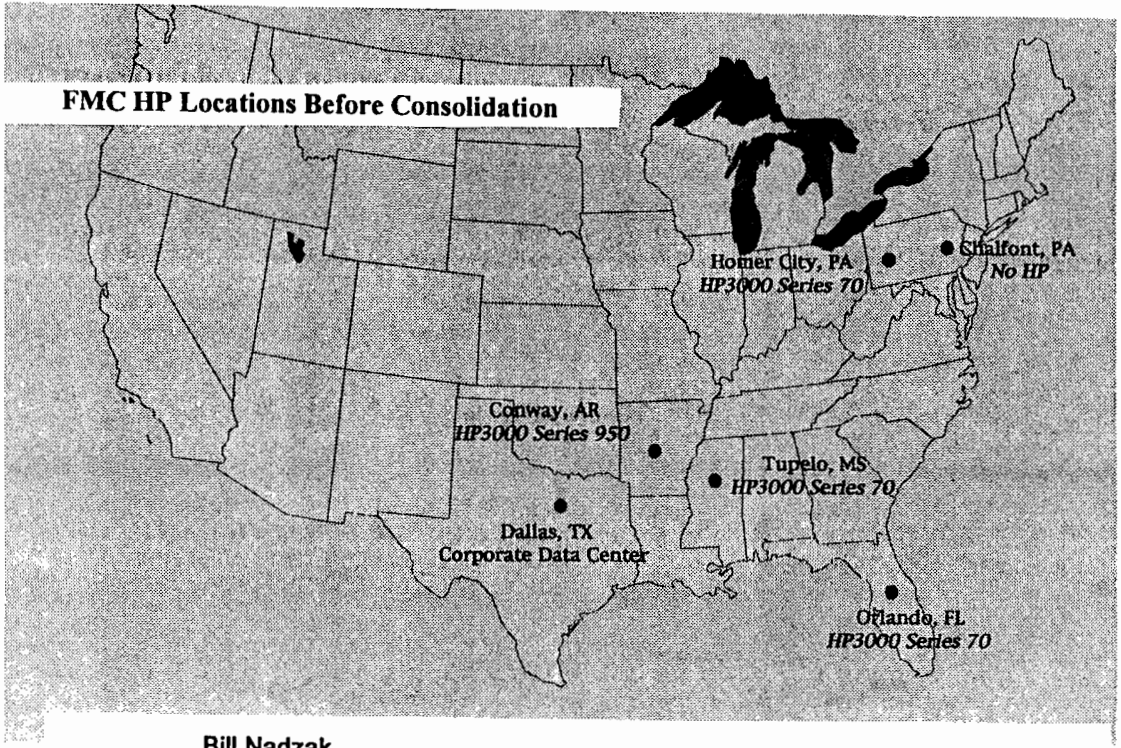
FMC HP Locations Before Consolidation
Homer City, PA
HP3000 Series 70
Chalfont, PA
No HP
Conway, AR
HP3000 Series 950
Tupelo, MS
HP3000 Series 70
Dallas, TX
Corporate Data Center
Orlando, FL
HP3000 Series 70

FMC HP Consolidation Project - Timeline

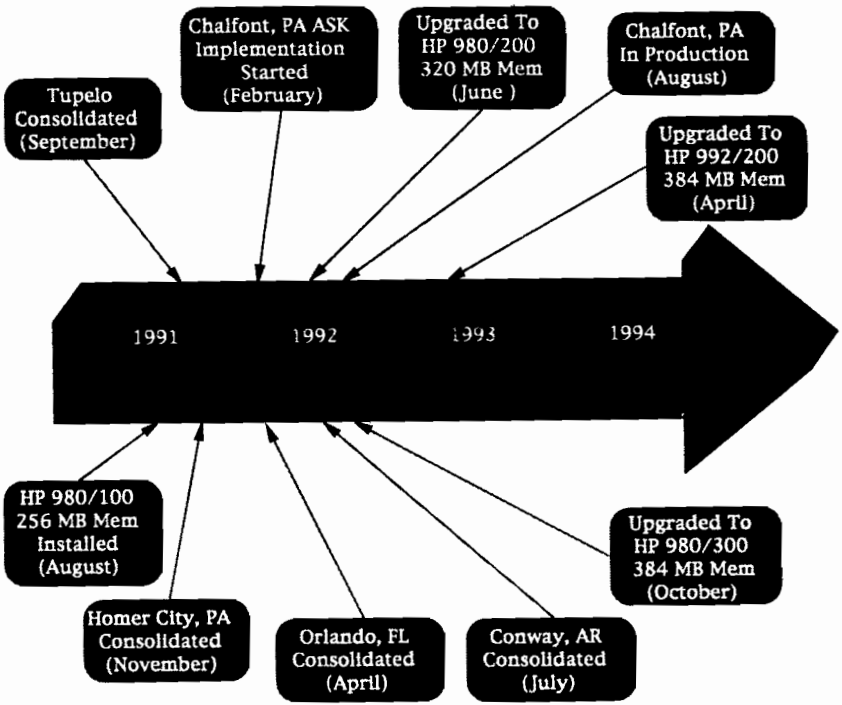

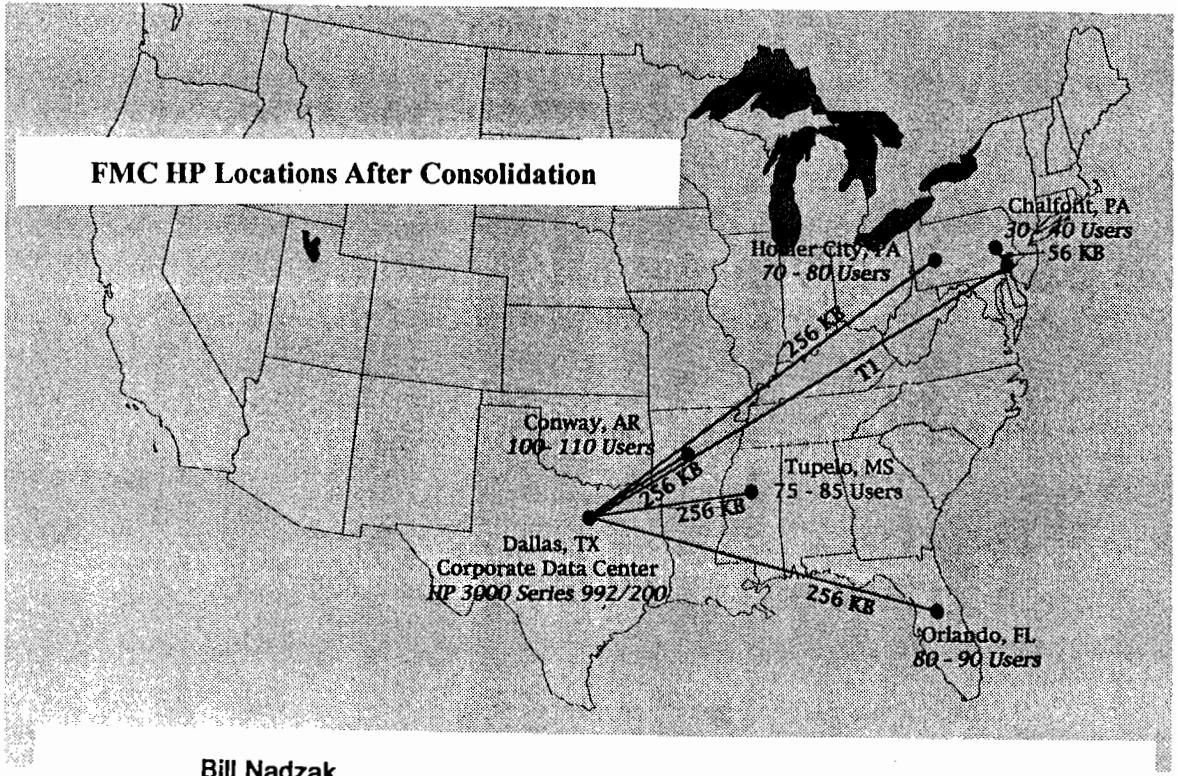
FMC HP Locations After Consolidation
Chalfont, PA
30 - 40 Users
56 KB
Homer City, PA
70 - 80 Users
256 KB
T1
Conway, AR
100 - 110 Users
256 KB
256 KB
Tupelo, MS
75 - 85 Users
Dallas, TX
Corporate Data Center
HP 3000 Series 992/200
256 KB
Orlando, FL
80 - 90 Users

Bill Nadzak

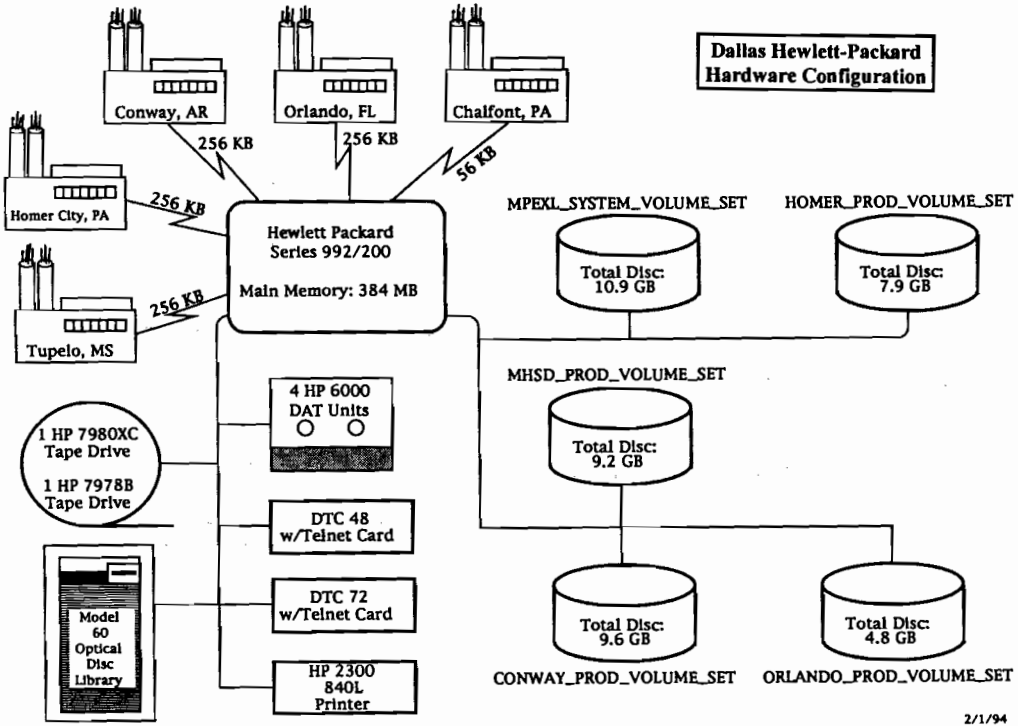
Dallas Hewlett-Packard
Hardware Configuration
Conway, AR
Orlando, FL
Chalfont, PA
Homer City, PA
Tupelo, MS
256 KB
256 KB
56 KB
256 KB
256 KB
Hewlett Packard
Series 992/200
Main Memory: 384 MB
MPEXL_SYSTEM_VOLUME_SET
Total Disc:
10.9 GB
HOMER_PROD_VOLUME_SET
Total Disc:
7.9 GB
MHSD_PROD_VOLUME_SET
Total Disc:
9.2 GB
4 HP 6000
DAT Units
1 HP 7980XC
Tape Drive
1 HP 7978B
Tape Drive
DTC 48
w/Telnet Card
DTC 72
w/Telnet Card
HP 2300
840L
Printer
Model
60
Optical
Disc
Library
Total Disc:
9.6 GB
CONWAY_PROD_VOLUME_SET
Total Disc:
4.8 GB
ORLANDO_PROD_VOLUME_SET
2/1/94

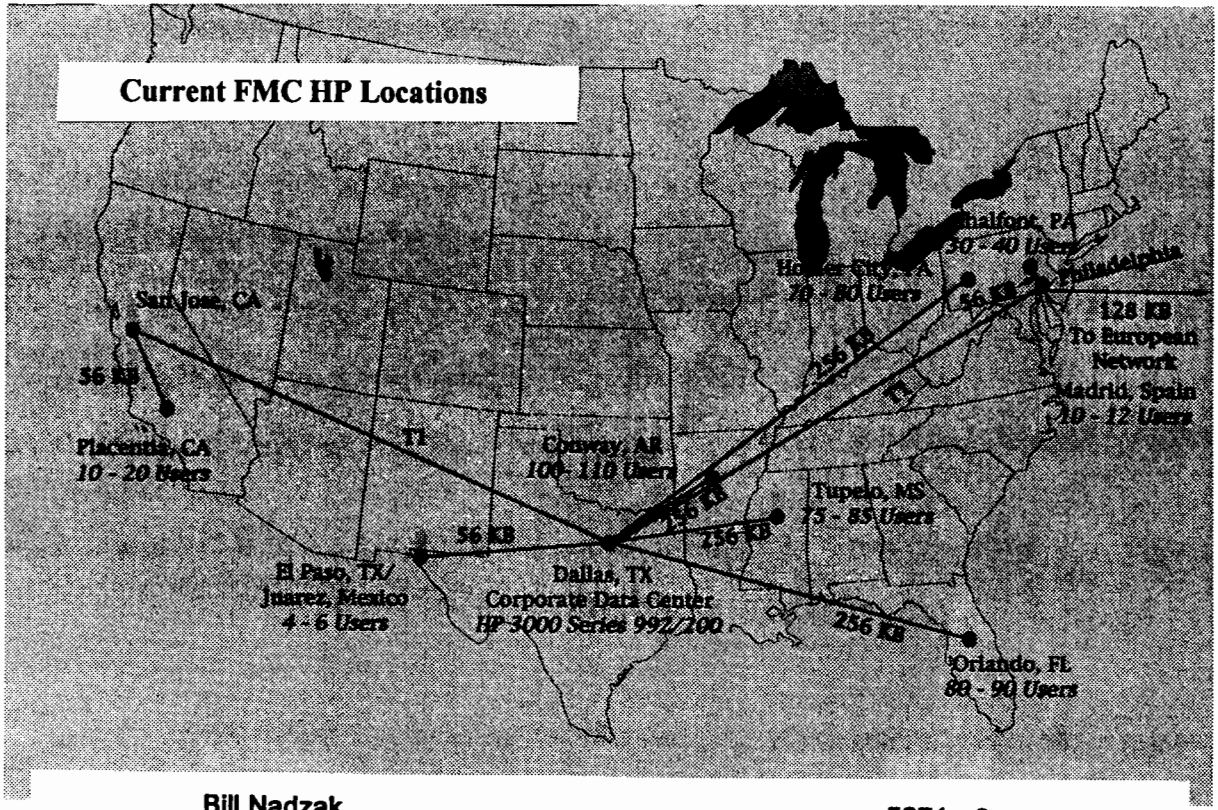
Current FMC HP Locations
San Jose, CA
56 KB
Placentia, CA
10 - 20 Users
T1
El Paso, TX/
Juarez, Mexico
4 - 6 Users
Conway, AR
100 - 110 Users
56 KB
Dallas, TX
Corporate Data Center
HP 3000 Series 992/200
256 KB
Tupelo, MS
75 - 85 Users
256 KB
Orlando, FL
80 - 90 Users
Chalfont, PA
30 - 40 Users
Philadelphia
56 KB
128 KB
To European
Network
Madrid, Spain
10 - 12 Users

HP 3000 Consolidation: A Case Study Considerations

- Operating System Differences
  - MPE V vs MPE/iX
  - MPE/iX 3.0 vs MPE/iX 4.0
    - » Memory concerns
- Temporary Resources
  - Disc space
  - Tape drives
  - Printers

HP 3000 Consolidation: A Case Study Considerations

- Third Party Software
  - New software (MPE V to MPE/iX)
  - License transfers
  - New Costs
    - » Upgrade costs
    - » Yearly support costs

HP 3000 Consolidation: A Case Study Considerations

- Asset Disposition
  - Transfers
  - Trade-ins
  - Disposals

HP 3000 Consolidation: A Case Study Considerations

- Network Connections
  - LAN
  - WAN
- DTC Considerations
  - Open View DTC Manager
    - » PC driven
- Clustering vs Single Large Box
  - Quest SharePlex

HP 3000 Consolidation: A Case Study Considerations

- Team Communications
  - Insure that both the central & remote sites understand their respective tasks
  - Set target dates
    - » Allow for "just-in-case" time
  - Setup regular status checks
- Allow For a Complete Cut-Over Test