CUSTOMER VALUE CHAIN

"I" On The Customer"

"Our success as a corporation depends on each of us being focused 100 percent of the time on satisfying customers."

Bob Palmer,
CEO, Digital Equipment Corporation

CUSTOMER VALUE CHAIN

The Only Way to Win

"Customers don't want choice. They only want exactly what they need."

E. Joseph Pine,
<u>Mass Customization</u>

CUSTOMER VALUE CHAIN

The "I" On The Customer

"Every person in a key position has to see himself or herself as a mini-CEO. They have to conceptualize what has to be done, in the same way the CEO does."

Michael Walsh
CEO, Tenneco

CUSTOMER VALUE CHAIN

Process Not Function

"People always ask, "Is the change over? Can we stop now?" You've got to tell them, "no, it's just begun."

John F. Welch
CEO, General Electric

CUSTOMER VALUE CHAIN

Customer Value Chain

"To inaugurate large-scale change, you may have to create the burning platform. You have to give people a reason to do something differently."

Lawrence Bossidy,
CEO, Allied Signal

CUSTOMER VALUE CHAIN

What is The Problem?

"You need to know a lot about Digital to do business with Digital productively, which is very hard."

Richard Bowman,
Director, Citibank-Quotron

CUSTOMER VALUE CHAIN

The "Value" of a Customer Value Chain

"Working with Digital in the past has been sort of like dancing with an octopus. There seems to be a head, but there are an awful lot of arms that are coming into the picture, and it makes it hard to get a single person who can say 'yes' or 'no'."

Richard Hope,
Vice President, Xerox Business Services

CUSTOMER VALUE CHAIN

What We Know

"For Dustin Hoffman, as The Graduate the future was plastics. Today you might say it's plasticity: the ability to adjust and learn."

Fortune Magazine

CUSTOMER VALUE CHAIN

Customer Value Chain Program Description

"Would the customer pay for this?"

Bob Palmer,
CEO, Digital Equipment Corporation

CUSTOMER VALUE CHAIN

Customer Value Chain Program Description

"Would the customer pay for this?"

CUSTOMER VALUE CHAIN

First Steps

The key question: "If we could start from scratch, how would we do this?" And the result: "Then do it that way and throw away everything else."

Fortune Magazine

CUSTOMER VALUE CHAIN

First Steps

The key question: "If we could start from scratch, how would we do this?" And the result: "Then do it that way and throw away everything else."

CUSTOMER VALUE CHAIN

"I" ON THE CUSTOMER

I am here today to talk to you about the work we need to do in the transformation of Digital.

Windows of opportunity and advantage open and close in the blink of an eye. Competitors are coming from everywhere. From overnight start-ups to a growing tide of industry crossovers. Endless reorganizations. Customer demands that keep evolving. How do you solve a problem that changes before you can implement a solution? It looks like there is no way to win.

The only way to win is to keep your eye on the target, *the customer.* Everything else is secondary. You have to give them what they want.

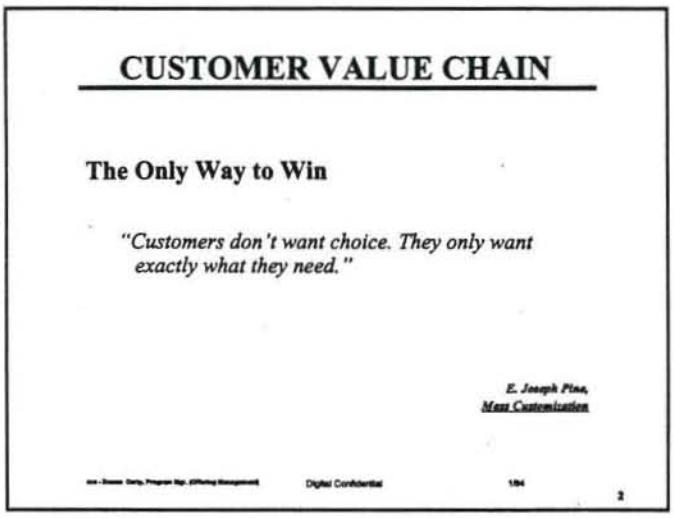

Over the last two decades our industry enriched itself by giving the customer a bewildering array of choices. That time is over. Today you can build a better mouse trap and the customer will expect you to bring it to their door...and set and service it...and do it for less...and faster.

Whether we see it as a problem or an opportunity, customer satisfaction will be THE determining factor between the winners and the losers in our industry.

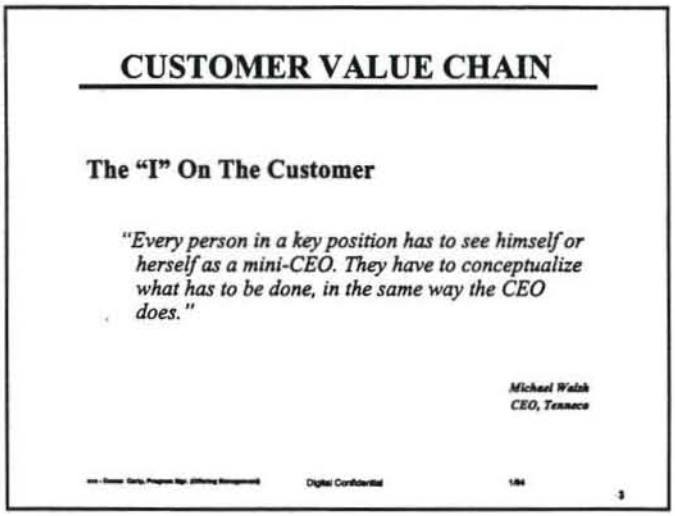

Overall customer satisfaction cannot be achieved by sales, or marketing, or engineering or any other corporate functions. In fact, it cannot be accomplished by functional groups at all. Groups tend to form boundaries and require the customer to find the passageways and connections between them.

Customer satisfaction can only be accomplished when each individual in the corporation starts every day by asking themselves, "What can I do to help the customer today?"

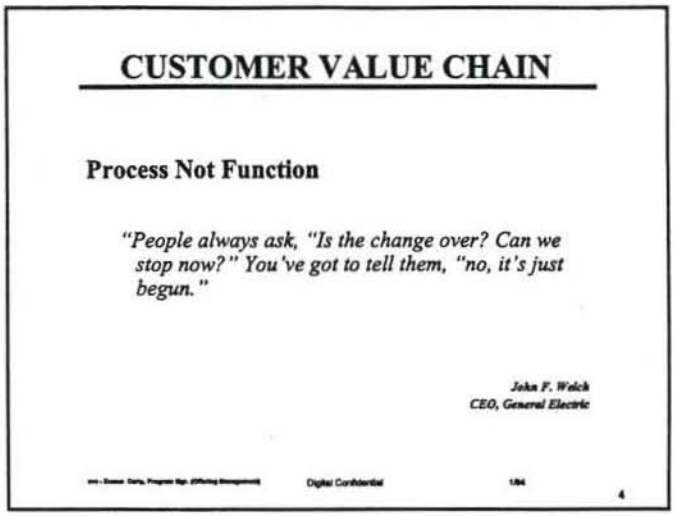

Functional groups don't satisfy the customer. Action does. Was it easy to get the answer I needed? Do they understand my problems? This is what the customer values. Those who can identify, perfect and link the processes that are important to satisfying the customer will be the big winners in our industry. They will have a company founded on a *"Customer Value Chain"*.

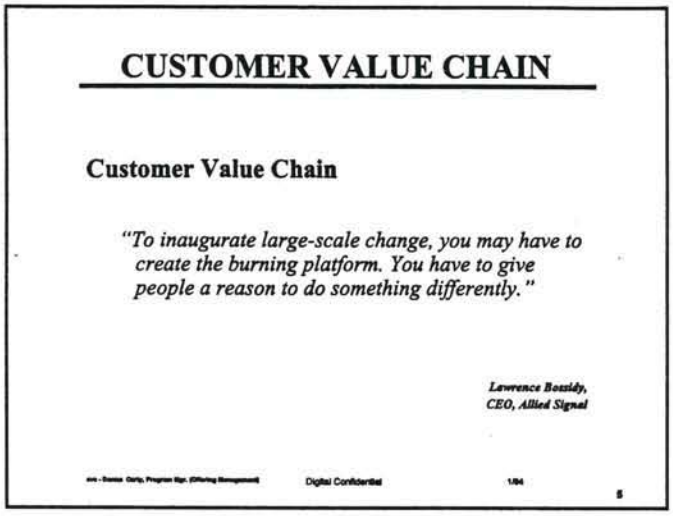

The "Burning Platform" is our financial tail-spin. Digital must become a customer-focused corporation based on a solidly linked Customer Value Chain. We will do it because it is a corporate strategy and direction. We will do it because it is the only way to survive much less thrive.

The only questions remaining are how well and how fast we will do it. Achieving this transformation from a function to a process focused company will require the support and enthusiasm of every individual in the company. It will require each person to stay focused on satisfying the customer while continuing to learn and improve their skills.

This presentation will help you to convey the need for this transformation and define the specific work to be done.

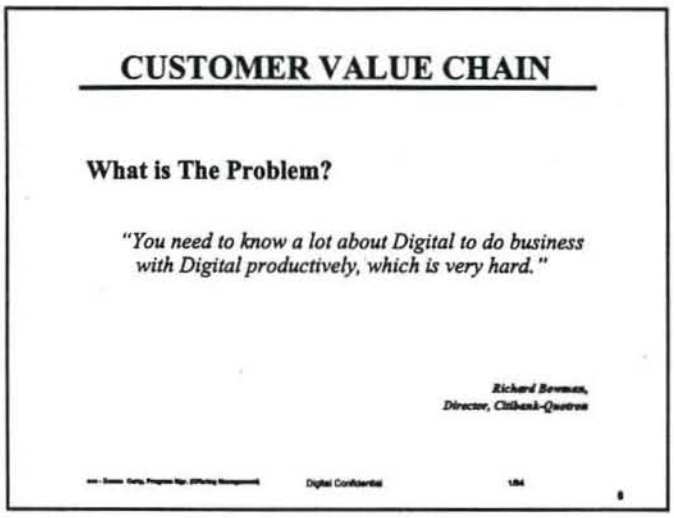

How serious is the need to transform the company? Here are just a few examples:

- Only 50% of those customers who contact us can get an answer
- To that 50% our response time is two weeks while our competition can respond in 24 hours
- We have 500 "800" numbers. Some are not even answered.
- We ignored over 250,000 customer leads last year alone.
- We have 340,000 part numbers but only 14,000 were sold one or more times last year.

These problems, and many others, are the result of our efforts to deal with a distributed business problem with a hierarchical organization. It is the business system we inherited, but it is not the one we have to live with. It is certainly not the one we will win with.

As we quickly transform into a corporation that delivers customer value, we will flatten the organization by moving from a group to a team mentality.

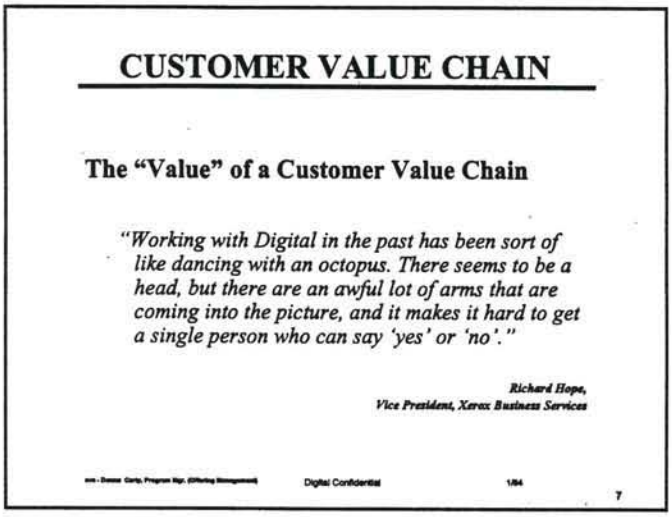

Customer Value Chain transformation has already achieved dramatic results.

- $.8B in cost savings
- $1.1B in asset reductions
- Consulting business value chain efforts reduced bid costs by 30% while doubling the "win" rate.
- Multi-Vendor Customer Services unified their offering menus
- Achieving Engineering Excellence (AEE) efforts resulted in startling reductions in menu item listing
- These early results clearly demonstrate the relationship between customer valued tasks and our profitability.

We know that merely filling our existing orders will expand our market share and revenues could increase by $1 billion in FY95. We know that by improving processes and eliminating confusion, our operating expenses can be reduced by $100MM in FY95. In inventory alone, $400MM can be released by cutting our excess and obsolete inventories. We can redirect $10MM non-productive person hours to solving our customers' problems.

We already know what we have to do: satisfy the customer. The Customer Value Chain transformation is how we will accomplish that goal.

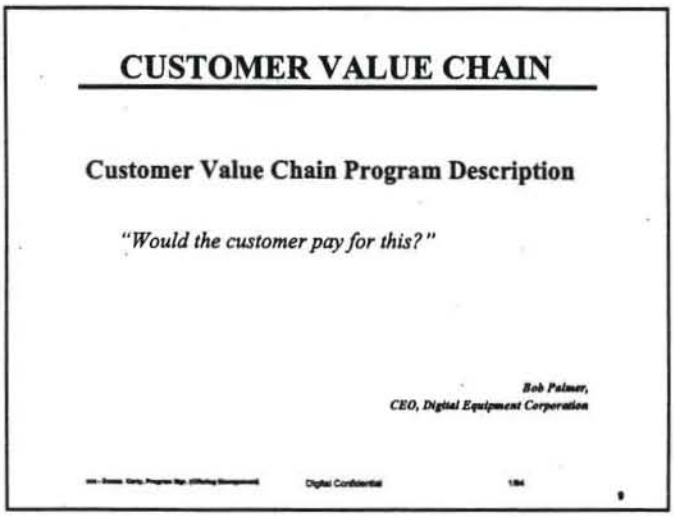

The objective of our Customer Value Chain re-engineering program is to achieve in our industry what imported car makers did in the U.S. Automotive Industry. They used customer satisfaction to win marketshare, drive revenue, (and profit) and to become the industry leader in the customers' eyes.

We have defined the core processes of our Customer Value Chain transformation as:

Planning

We will come up with a strategic plan that identifies Digital's strengths and where improvement is needed, and matches those factors to market needs.

Engaging the Customer

We will gain a complete understanding of customer needs and see how Digital's core competencies can meet those needs. We will ensure all other processes are aware of these specific customer needs.

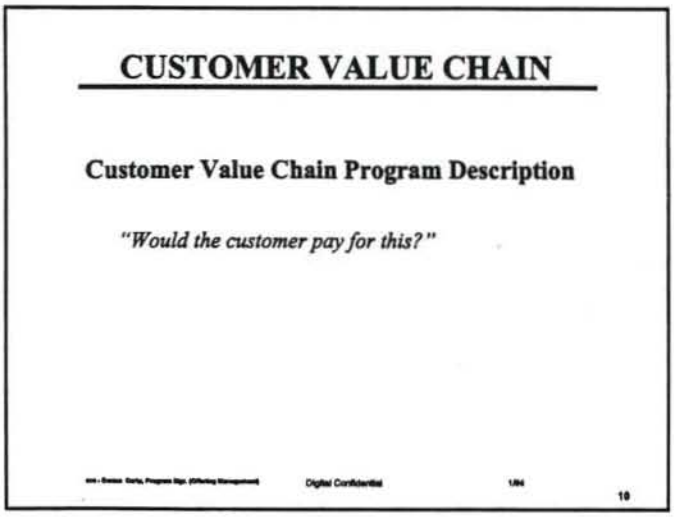

Creating Capabilities

We will do whatever is necessary to deliver customer-focused solutions, including setting up new channels, manufacturing pipelines or bringing in new suppliers or consultants.

Delivering Value

We will make sure Digital is easy to do business with by efficiently delivering the products and services which meet specific customer needs.

Implementing a Knowledge Aquifer

We will provide an Information System which allows for proactive knowledge sharing among employees, customers, suppliers and partners.

A number of initiatives will be launched throughout the year to support these core processes but they must have worldwide support to be successful. In this new approach, some organizations will be consolidated...work may be streamlined, eliminated, or outsourced and, where possible, processes for such things as ordering, billing and delivery will be standardized.

Building on our initial Customer Value Chain (CVC) successes, a company-wide team was assembled to recommend the next set of steps in our transformation. In November 1993, the Senior Leadership Team SLT approved the first four initiatives recommended by the CVC team:

Contact Management

To be more responsive to customers, this initiative will focus on the areas of demand generation, planning and the entire lead management area of capture, qualification, nurture and refer.

Order Support

This initiative will focus on the entire ordering process, primarily quote, configure and propose solutions as well as electronic channels available to customers and partners.

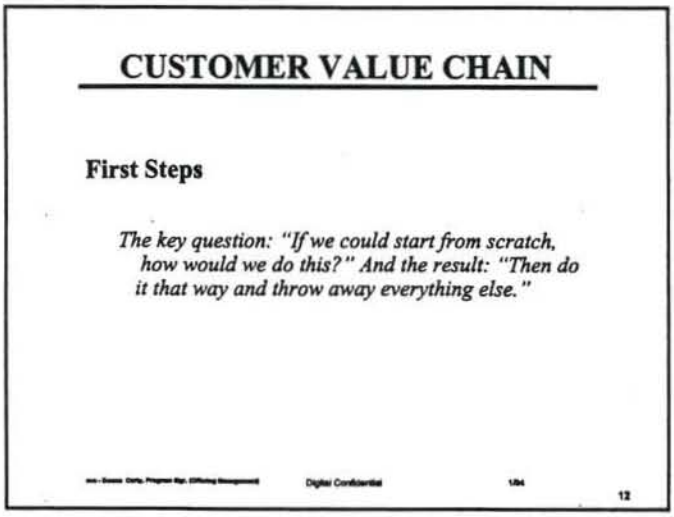

Offering Management

This initiative involves repairing customer relationships, reducing expenses, defining our offerings, linking CVC to customer business problems and reducing complexity.

Creating a Knowledge Aquifer

Since knowledge sharing is the foundation for transformation, this initiative will support the entire enterprise, including customers and partners by providing easy access to information. This proactive sharing of knowledge will eliminate stovepipes and allow more time to be spent adding value for the customer

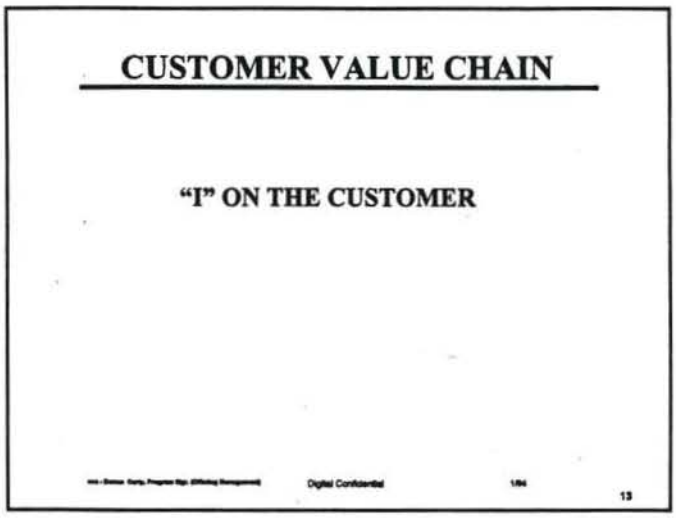

While each of these initiatives focuses on important specific objectives, they each contribute to the overall goal of transformation into a company that is easy to deal with.

The Customer Value Chain transformation will allow us to keep our eye on the critical target: customer satisfaction. Customer Value Chain initiatives will allow each of us to contribute to the process every day: to allow each of us to keep our "I" on the customer.

NaC Design Centers:

Facilitating Dialogue Around Customer Needs

Team Members:

Engineering:	Bill Duane
	Colin Strutt
Product Mkt:	Diana Larrea
SMA:	Rolf McClellan
	Ellen Quackenbush
Product Mgt:	John Puma
Net Strategy:	Chris Strutt
ANCs:	Donna Moscaritolo
	Cliff Weiss

Outline

- The problem
- Our approach
- NaC Design Centers

Problem Statement

- T&N lacks a shared language to discuss customer needs and translate these needs into product design/strategic direction

Approach

- Design Centers: a taxonomy for describing current and potential customers
- Uses for Design Centers:
  - -- Define NaC business opportunities
  - -- Test existing solutions
  - -- Gather customer requirements
  - -- Facilitate dialogue between field and engineering

Taxonomies Tested

- Sites
  - -- Headquarters, Manufacturing Site, Regional Office, Branch
- Organizational Level
  - -- Corporate, Division, Department, Workgroup, Lone User
- Buying Points
  - -- MIS, Telecom, Network Manager, Business Managers

Sites

Average Number of Sites
in Fortune 1000 Firms

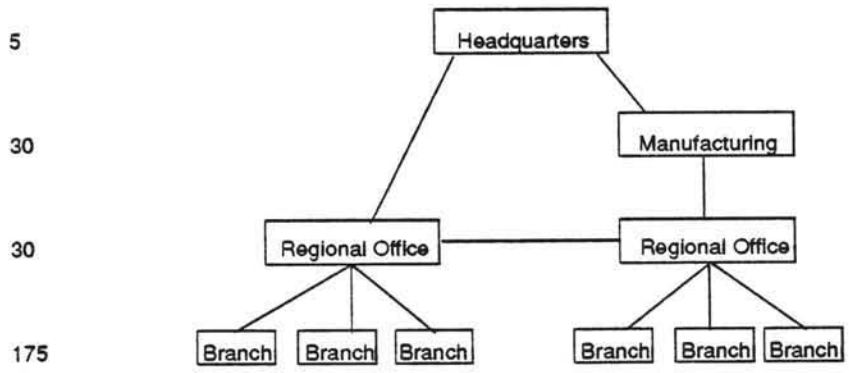

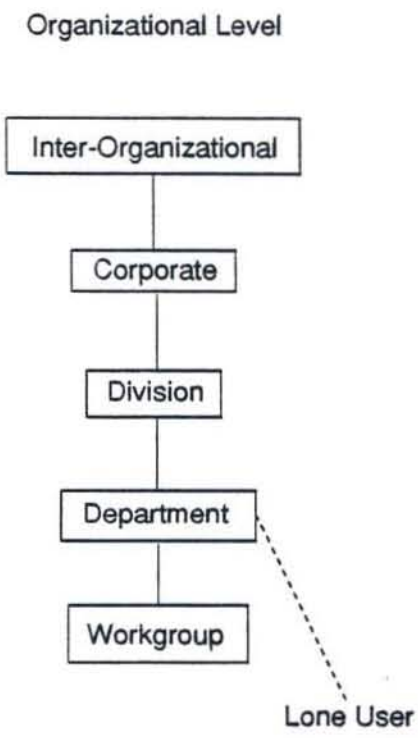

And the Winner Is..

- Customer Buying Points
  - -- Key to understanding today's network market dynamics
  - -- Incorporates sites (geographic) and organizational (political) dimensions

Key Observations

- Buyer not always the same as the final user
- Immediate, business needs of final user may be constrained by need to:
  - -- Integrate with installed base
  - -- Conform to architectural standards and directions

Key Observations

- Two key "buying points" in customer organizations:
  - -- Buyers: purchase for own and "client" groups
  - -- Influencers: define technology architectures or business strategies that constrain purchases within organization
- We must focus on our traditional as well as potential buyers/influencers

Who are the Buyers

- Potential buyers of network products are:
  - -- End Users
  - -- Workgroups
  - -- Departments
  - -- Site/Facility
  - -- Central Network/Telecom
  - -- Central MIS
  - -- Executive Management

Product/Buyer Characteristics

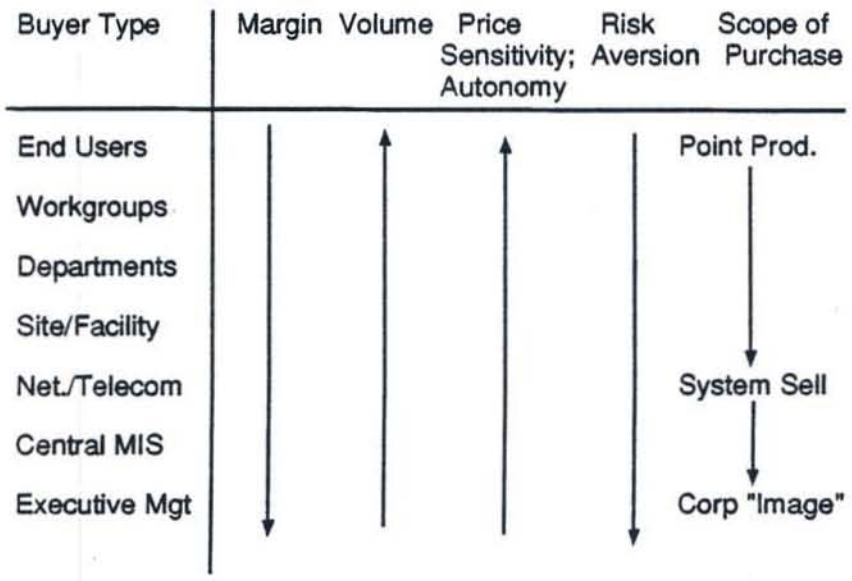

Buyer Taxonomy

- Buyers characterized by:
  - -- Scope of their purchase authority
  - -- Scope of their influence
  - -- Degree of vendor loyalty
  - -- Product volume/margin characteristics
  - -- Typical products purchased

End User

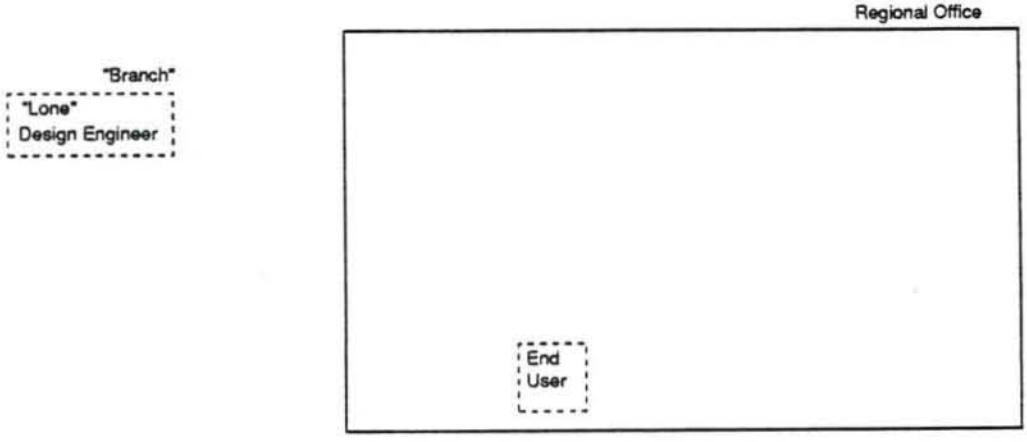

End User Buyers

- Purchase for desktop
  - -- Individual, not group, productivity
  - -- Products are high volume/low margin
- Exhibit little vendor loyalty:
  - -- But very loyal to "solution" once chosen
  - -- Trade press is major information source
  - -- View network/system as single service
- Constrained by workgroup choices

End User Buyers

- Driven by price and performance, as defined by:
  - -- Functionality and quality
  - -- Look and feel
  - -- Ease of use

Stuff End Users Buy

- Mail User Agents
- PC LAN components (add-ons to basic workgroup solution)
  - -- PC LAN adapters
  - -- PC operating systems
- Client software, such as personal and professional applications

Workgroup

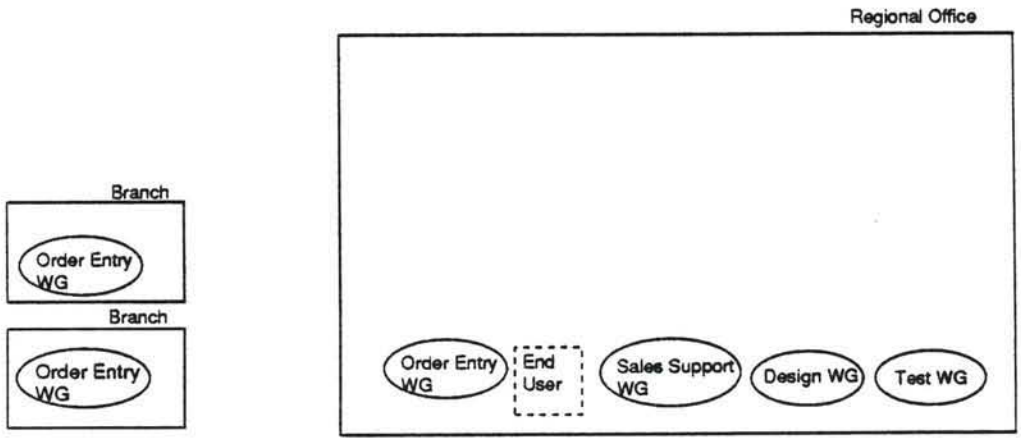

Workgroups

- A Workgroup is a collection of individuals performing the same task(s)
  - -- Usually at single site
  - -- Small (average 15 members)

Workgroup Buyers

- Purchase products for office, factory floor, or lab environment
- Products are high volume, low margin
- Vendor loyalty based on familiarity with vendor products-- "training interia"
- Driven by price and performance, as defined by:
  - -- Impact on workgroup productivity
  - -- "Plug and play," ease of use/maintainability

Workgroup Buyers

- Purchases are constrained by:
  - -- Existing site and department computing/network infrastructure
  - -- Need to integrate end user environments and provide end user access to common resources

Stuff Workgroups Buy

- PC LAN servers
- Workgroup office automation applications
- Workgroup mail (usually PC LAN-based)
- Workgroup network services (PC naming, printing and disk)

Department

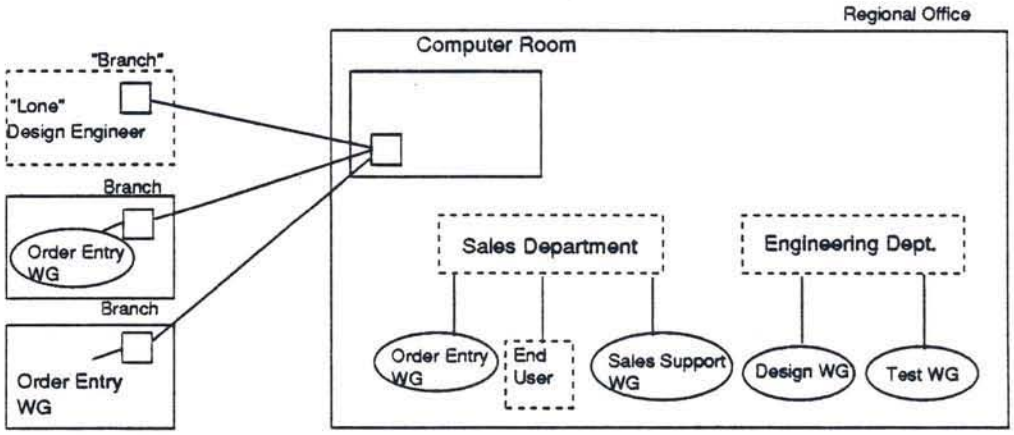

Departments

- A Department is a collection of individuals and/or workgroups which are:
  - -- Performing the same kind of work
  - -- Focused on the same business mission
- Department often cross site and geographic boundaries
- Departments may need links to external customer, supplier, or collaborative organizations

Department Buyers

- Purchases support departmental activities
- Products are mid volume and mid margin
- Vendor loyalty based on applications investment-- "applications-based inertia"
- Influence is key part of department's job
  - -- End users/workgroup to be compatible
  - -- Site managers to provide services for local/remote departmental workgroups
- Constrained by need to link to site and enterprise network services

Stuff Departments Buy

Departments make and influence purchases for the following environments:

- -- Office, lab, factory floor area
- -- Computer room
- -- Remote sites (that are part of department)

Stuff Departments Buy

- Office, lab, factory floor products
  - -- Desktop devices
    - > Influence workgroups/end users
  - -- Desktop communications solutions
    - > Buy adapters, drivers, utilities

Stuff Departments Buy

- Computer Room products
  - -- File, Mail and PC LAN Servers
  - -- Applications Gateways (e.g. Mail gateways to convert to enterprise mail)
  - -- Firewall/isolation products between sites and departments (i.e. to provide secure physical/logical boundaries)

Stuff Departments Buy

- Remote Site products
- Influence the purchase of:
  - -- Access routers
  - -- Integrated communications servers
  - -- Wide area terminal connectivity

Site/Facility

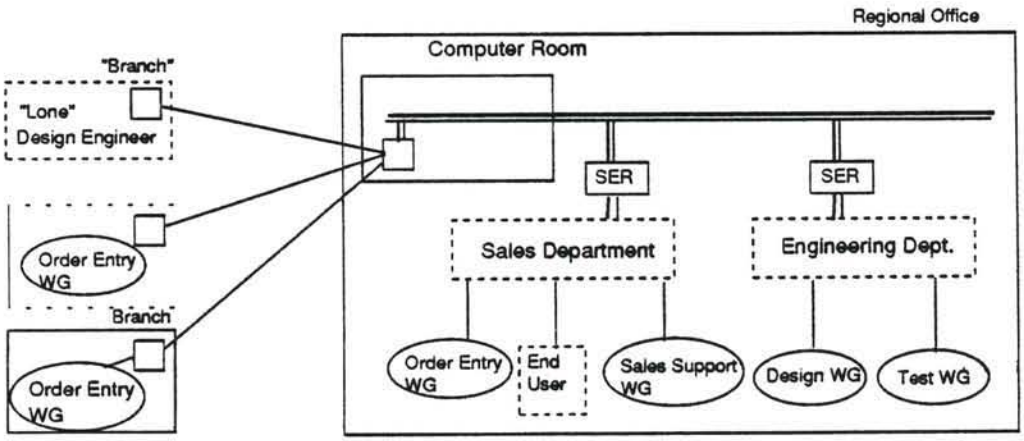

Site/Facility Buyers

- Provide physical connectivity and interoperability within site
  - -- Products are mid volume/mid margin
  - -- Reside in wiring closet
  - -- Influence and/or buy directly ("in bulk") for workgroup and department intrasite needs

Site/Facility Buyers

- May also provide site-to-site connectivity:
  - -- For department's remote sites
  - -- Transmission lines
  - -- Rely on "bulk" purchase by Telecom

Stuff Site/Facility Buys

- Terminal Servers
- Bridges and Routers
- Wiring Concentrators
- Smart Hubs
- Ethernet, Token Ring, FDDI

Network/Telecom

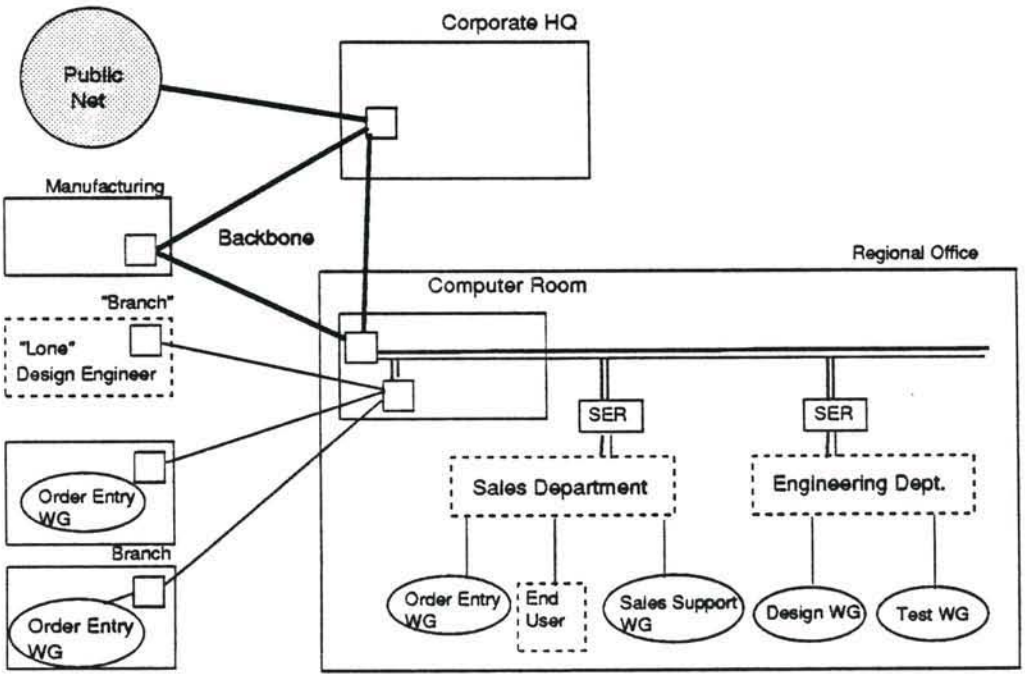

Network/Telecom Buyers

- Purchase products which provide:
  - -- Interconnection of the enterprise
  - -- Access to public and private networks outside the enterprise
- Products ar low volume/high margin
- Products reside:
  - -- In MDF, for interface to local site
  - -- At site, but logically part of backbone

Network/Telecom Buyers

- Driven by performance and functionality, as defined by:
  - -- High bandwidth, low delay
  - -- Manageability
  - -- Accommodation of existing network traffic :
    - > Protocols
    - > Appl'n-specific vs gen'l purpose nets

Network/Telecom Buyers

- Set the network strategy
  - -- Interact with computing strategy set by MIS
  - -- Develop approved vendor list
- Three major data networking environments:
  - -- SNA
  - -- Peer-to-Peer
  - -- X.25 (Europe)

Stuff Network/Telecom Buy

- Enterprise Products
  - -- WAN bridges and router
  - -- Portals, network/application gateways
  - -- T1 and T3 multiplexers
  - -- MTA "Post Offices"
  - -- Front End Processors

Stuff Network/Telecom Buys

- Enterprise products (cont'd)
  - -- Trunk Data Switching (PBX)
  - -- Enterprise management
  - -- Network integration servers
  - -- Inter-enterprise applications/services
  - -- Network Strategy (Adams in a box)

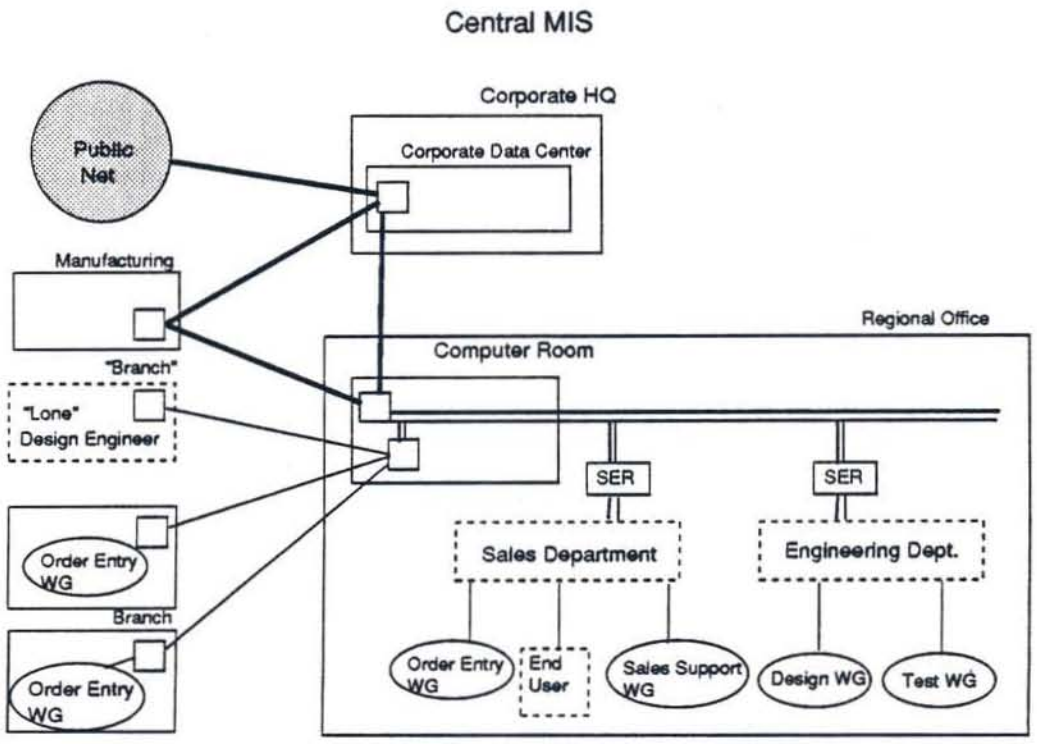

Central MIS Buyers

- Provide key services across the enterprise
- Set enterprise computing strategy
  - -- Computing architect
  - -- Database architect
- Major computing environments which may be aligned with core functional departments:
  - -- IBM
  - -- Digital
  - -- Others (Sun, NCR, etc.)

Stuff Central MIS Buys

- Transactions processing services
- Computing Strategies (Strecker in a box)

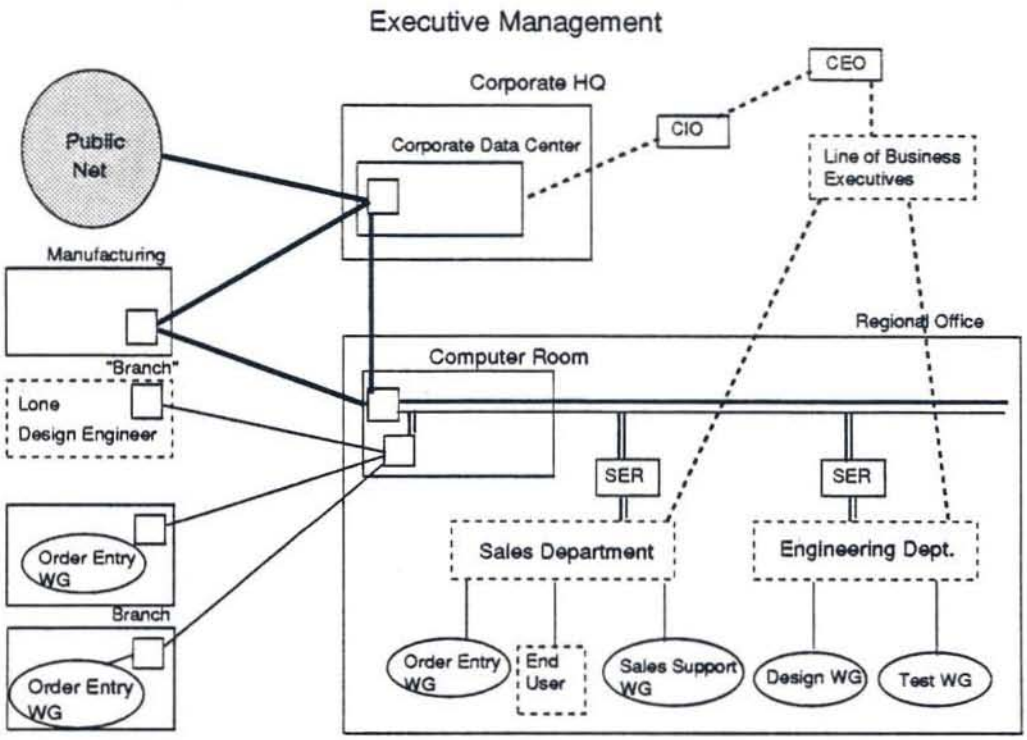

Executive Management Buyers

- Set enterprise business strategy and vision
- Leading-edge firms will re-structure IT to align with changing business needs
- Constrained by need for "buy-in" from line organization and IT architects

Stuff Executive Mgt Buys

- "The Whole Company"
  - -- Corporate strategy and vision
  - -- Financial stability and viability
  - -- Ability to partner "in real time" with customer to solve business problem
- Olsen in a box

Recap

- End Users, Workgroups
  - -- Single site, commodity sell
- Departments
  - -- Multi-site, solution sell
- Site/Facility
  - -- Single/Multi-site, physical interconnect
- Network/Telecom, MIS, Exec. Mgt.
  - -- Enterprise focus, strategy sell
  - -- Key influencers are network, computing, and database architects

What We Need From You

- Is Design Center "buying points" taxonomy useful?
- Does it model network market dynamics?
- Is it valid for:
  - -- Non-U.S. organizations
  - -- Non Fortune 1000 firms, non-profit (education, government)
- How to add external influencers (e.g. analysts, consultants, standards groups)
- Will this model facilitate effective dialogue between the field and engineering?

What Are The Next Steps?

- Validate, review with:
  - -- Customers, PDF; return to ANCs
  - -- Research data (DCE study)
  - -- NaC and field decision groups (NTF)
- Pilot
  - -- What program or product?
- Communicate and deploy!

From: DELNI::FAUCHER "24-Sep-1991 0915" 24-SEP-1991 09:14:07.45
To: @VISION
CC: THURK
Subj: Networks Vision for the 90s Information for Participants

To Networks Visioning Participants,

As I have discussed with each of you, this working group is the first step in developing a five year networking vision for Digital. It has been some time since the Vision for the 80s was developed, it is now time to rekindle that type of vision so that there is a clear overall direction for the networking business during the 1990s. The questions that we need to focus on in the context of developing the "Straw Horse" plan are:

1. What business are we in?
2. Who are our customers?
3. How do we win? (ie competitive knowledge)
4. Knowing the above, what does that mean for our Vision?

The proposed outline for the finished deliverable is below:

DIGITAL NETWORKS BUSINESS VISION "STRAW HORSE"

OUTLINE

MISSION - REASON FOR BEING (TIMELESS STATEMENT)	COUPLE OF SENTENCES
OVERARCHING GOAL - 1996...."THE STAKE IN THE GROUND"	COUPLE OF SENTENCES
CONTEXT:	
MARKETPLACE ASSUMPTIONS (COMPETITIVE/INDUSTRY) CUSTOMER ENVIRONMENT ASSUMPTIONS CUSTOMER PROBLEM/OPPORTUNITY	2 PAGES
STRATEGIES (FAMILIES OF PRODUCTS)	8 PAGES
HIGHLIGHT AREAS OF MAJOR IMPLICATIONS FOR IMMEDIATE ATTENTION (TO GET SOME IMMEDIATE FOCUS AND START ADDRESSING THE PRIORITIZATION OF GOING AFTER THE VISION)	1 PAGE
CONSEQUENCES OF NOT CHANGING EXISTING DIRECTION	1 PAGE

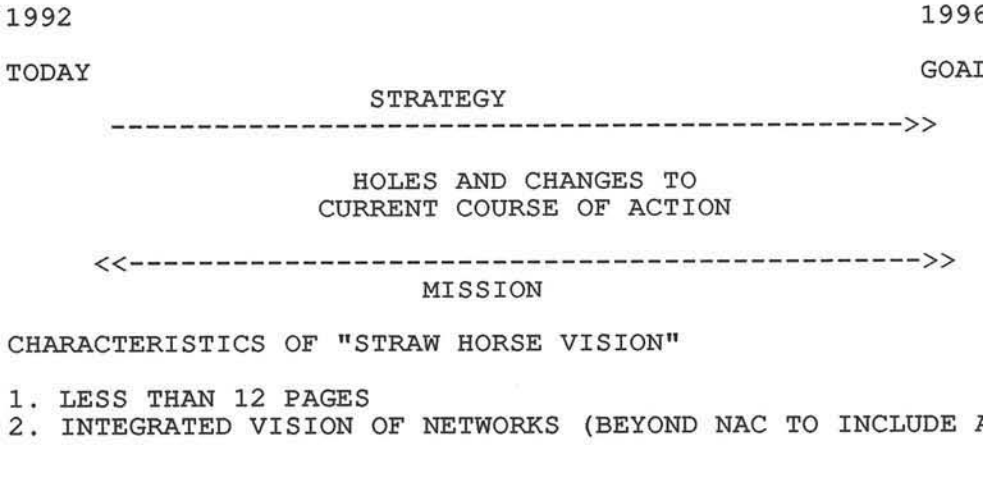

CHARACTERISTICS OF "STRAW HORSE VISION"

1. LESS THAN 12 PAGES
2. INTEGRATED VISION OF NETWORKS (BEYOND NAC TO INCLUDE ALL NETWORK

RELATED ASPECTS)
3. USE PICTURES (E.G. WIRING STRATEGY)
4. TIE IN "SYSTEMS" FOCUS PCUS.
5. NOT AN ORGANIZATION CUT (SEAMLESS ACROSS ORG. BOUNDARIES)

The people that were nominated and have agreed to participant in this effort include:

John Shebell - Team Leader
Tony Lauck - NaC A/D
Bill Hawe - NaC A/D
Jeff Low - NaC Engineering
Kate Winkler - NaC Engineering
Al Avery - ALPHA Program
Bill Keyworth - Systems Management
Jeff Schriesheim - PCI
Jim Herman - Outside Consultant

Helping John Shebell to facilitate this 4 day effort will be Chris Strutt and Susie Rheault.

The confirmed dates for this meeting are Oct 4th, 14th and 15th. The Oct 3rd meeting date that I discussed with some of you will not take place. I would however like to know how many can meet on OCT 10TH, this would be the targetted forth day.

As the forum to present and discuss this Straw Horse vision, I will be working to schedule time at the NaC Staff the last two weeks of October. In addition, I will be working with Chris Strutt to identify an additional group of reviewers of this Straw Horse document.

The meetings will take place from 8:30-5:00 in the NaC Strategy Room which is in LKG 1-2/E14 (in the lab opporsite the Sarah Caldwell c.r.).

If anyone has any questions, please call me this week (9/23).

Thanks in advance for your energy in being part of a critical event for the networks business.

Cornel

To: Participants in the Networks Vision for the 90's Meeting
Networks Group Staff

Subj: Summary Report

Attached is a summary report from NCRI on the "Networks Vision for the 90's" meeting held 1,2,3 May'91. It captures the learnings and outcomes of the meeting, in particular the shared operating model that the participants defined for the Networks Business.

Much work remains to be done to fully understand our desired endstate for 1997, and the event scenario that gets us there. We are creating an interactive wall display (the Strategy Room), which product/program/staff teams can work with to develop this understanding over the next three months.

For now, the fact that we have a shared, explicitly stated model for decision-making in the Networks Business, is a significant accomplishment, and an good starting point in our strategy for success in the 90's.

John Adams
Chris Strutt

Summary Report:
Mapping the Future of the Networks Business

Northeast Consulting Resources, Inc.
David Mason
James Herman
John Kelly

Introduction

This report summarizes the voting and presentations from the Future Mapping session held in early May. Nearly 50 people from six networks PCUs, customer services, marketing and the field participated in the workshop. This report follows the structure and sequence of the workshop exercises. It includes analysis of the group's voting, observations made by the attendees as well as NCRI's recommendations.

In this meeting, the group, although often divided and short on consensus, created a shared vision of the computing and communications industry and how Digital might run its own network-related business units. The central theme was a radical consolidation of the industry and the emergence of a very strong systems integration business. A secondary theme was the challenge of managing a mix of leading edge new technical products and standardized, cost competitive, older products.

Points From Conventional Wisdom

The insights to be gained from the conventional wisdom exercise, which asked participants to vote on the probability of an event happening by a certain year, can be summarized as follows:

- A view of the future encompassing cheap, fast, hardware and strong service related businesses
- A broad set of ideas about future network based applications
- Confused positions towards the standards process and multivendor networking generally
- No broadband events were chosen whatsoever, contrary to voting by other groups
- General disbelief that public network infrastructure needs would be handled quickly by the carriers
- General lack of interest in, knowledge of, or attention to consumer related businesses and applications

Very few events (45) were selected in the conventional wisdom exercise; we usually see at least 70 events selected. This reflects a high degree of uncertainty and the high number of differing, competing points of view in the group.

The group did believe in big, fast cheap hardware and cheap communications. Video seemed particularly important among a strong selection of networked applications. The group also believed in services which became very important as the group moved toward a shared vision. The group found it much easier to believe in successful new network applications than in successful new network technologies.

Most striking, however, was the large set of issues the group did not agree upon. There were obvious divisions within the group on the topic of standards and the value/role of the consortia process. Only with respect to network management did an event appear with the word multivendor in it. There was absolutely no mention of broadband! There was no belief in a consumer related communications products and services markets. The group has limited faith in public carriers - very little faith in the RBOCs. Software productivity gains were also viewed skeptically.

The events selected as highly unlikely were also revealing: 1 Gbs LANs in 1992+/- (already here in the opinion of others), successful re-engineering of the installed base (1992+/-), Japan open to free trade in 1991+/- (part of what the NTTs MIA effort is all about), and deregulation of the PTTs in the G-7 in 1993+/-.

The Alternative Scenarios

Six teams defended different 1997 endstates. Their approach to defending them offered some insight into the perceptions of each team.

The standards team (A) offered an orderly, structured view of how standards could be developed and implemented by the industry. Their icon was the prism and their slogan "See the Light through Standards". Their defense was both very well articulated and soundly rejected by the group.

The hitech team (B) held the point of view that the pace of technical development would continue to be different in different areas of networking technology. Their point was that product lines would continue to leapfrog each other in a messy progression toward the future and that continued innovation, not strict adherence to a standards program, was the route to industry leadership and profits. Customers sought competitive advantage from new technology and that meant leading edge, non-standard systems. Their icon was the lightening bolt and their slogan CCC (creative competitive computing) and they purposely depicted a disordered scenario. They advised creating network-based applications that pulled customers onto new network technologies. Their point of view was received positively by the other participants.

The network team (C) offered a confused and only partially developed defense of this endstate which described a global data utility clearly headed toward the scale of the telephone network. Their icon was the globe with a network and their slogan was "Network Nirvana". Missing from their defense was an understanding of the role of the public networks which would be critical to such an endstate. This caused the team to lose credibility with the group. We (NCRI) were surprised by the lack of understanding of the multiple dimensions of this endstate in the room.

The affordable team (D) held that their endstate was a realist's view of the world. They stressed that bringing costs down was critical to success. Accordingly, their icon was a mixture of currency symbols and a downward arrow. The team stressed that their scenario reflected the position of many users today and would continue to in the future. The user's environment was characterized as multivendor, multifunctional, and encompassed multiple generations of equipment. New equipment was installed in the customer environment because it was commodity priced and/or provided special functionality. The services component of users' budgets was very large and growing fast. The use of low cost product sales channels was key to this scenario. This teams position was well received by the group.

The consolidation team (E) posited a rapid consolidation of vendors into six global players. Key to the endstate was the idea that "vendor independent architecture" meant making large users comfortable with strategic vendors by lowering the barriers for the user to switch to another vendor's equipment; it specifically did not mean that users would have dozens of vendors serving various needs, as such a strategy would be too costly. The team took the position that large users (e.g. General Motors/EDS) would play a role in precipitating such an outcome, especially after a period of confusion in which users try to use standards alone to solve their integration problems. Large users want healthy vendors and healthy vendors can assume a significant share of the risk of large new IS projects. Sick vendors (like Unisys) are likely to continue to falter because customers react poorly to their unhealthy outlook. Key to this endstate was the vision of the big six vendors as mega-integrators, who "built the best and integrated the rest." The group responded positively to this team's position. We have seen other defenses driven by the need to support network infrastructure costs (thus merging carriers and vendors like ATT and NCR) or the scale of consumer market applications of the technology (possibly driven by the Japanese).

The last team was concerned with the emergence of Japan as the technology leader (F). Interestingly, the team took the position that, so long as Japan is not distracted from its current activity, the country will emerge as the leader. Furthermore, almost all the events deemed critical to most other teams were largely irrelevant, thus leading to the idea that the only strategy with a chance work is to understand their plans and beat them at their rules. This team's sobering defense definitely had an impact on the group.

The Shared Vision

The group came to share a very advanced, sophisticated, understanding of the industry as a whole and developed a description of the 1997+ industry endstate as context for deciding how Digital and its networks business should operate. Several teams expressed similar views of all the endstates operating together as a system. A diagram capturing their ideas has been sent as a separate postscript file. An explanation follows.

The global data network utility concept (C) remains what was characterized as a timeless vision. As such, it remains useful for external relations purposes and as a goal toward which every product should incrementally contribute. The important point in the rejection of the C vision as the dominant, operational, vision was that it resulted in long, architecture-driven projects that caused Digital to miss the near term market opportunities. Since the C vision is so far away, it must not be confused with a plan aimed at the 1997 timeframe.

The rapid consolidation of the industry (E) is a very real, probably near term phenomena. The Corporation must clearly manage toward such consolidation or run the risk of Digital's future being determined by others. To succeed in this environment, the Executive Committee would need to use the vision expressed by the group as a decisionmaking context. Moreover, the key aspect of an E-type player is a strong services capability, something that is being addressed by other parts of Digital. The E scenario is not one that NaC can determine or execute. However, NaC can support Digital in succeeding in endstate E by being responsive to needs expressed by the service organization.

The implications of the consolidation (E) endstate, are clearly somewhat dangerous; poorly articulated to the market, they could be perceived as a return to proprietary architectures and literally kill the Company. The discussion clearly suggested a move toward vendor independent architectures while maintaining strong strategic relationships with customers as the essence of the scenario. Making the customer feel comfortable by being more "open" than competitors is key. This is as much about marketing as it is about products - stop, for example, putting "DEC" in the name of every product! Finally, E requires a large direct sales effort because integration services revenue cannot be captured through indirect channels.

The consolidation scenario (E) has particularly interesting linkages to the hitech scenario (B) and the affordable scenario (D), which are essentially scenarios for companies feeding products into E.

First, the global vendors implied by E would have the margins and the scale to support the R&D required by the hitech scenario (B). The products generated in the hitech scenario (B) exist to provide global integrators (e.g. Digital and others) value added software and hardware. (B-type products are also sold through other distribution channels.) Digital's concept of middleware brings real value to such a strategy but "middleware" is a terrible marketing concept, doomed to failure.

Second, the global vendors have the services depth to solve the customer problems described in the affordable scenario (D). The D scenario is very service intensive thus requiring a dedicated sales force, even though the products are often commodity in nature and sold through indirect channels. It implies both complex relationships with and competition against systems integration companies. In the "D" scenario, product groups exist to provide:

- rapid, Japanese style, product cost reduction engineering and
- products that deal with integration of the installed base.

A critical element of D captured the idea that the first sale of a complex product required a sales force but subsequent sales of the same product could go through retail, direct mail, or other low cost channels.

Together, scenarios B and D suggest two separate and distinct types of products. The idea is that products (and product families) move through a cycle starting with new hitech innovation then moving through the standards process toward commoditized, value engineered products. The standards process (A) adds the fuel of volume to robust B-type ideas, driving them into the low cost world of D. This was referred to in the meeting as the "B-D Cycle". As more of the industry's revenue comes from D-type products, the industry will further consolidate.

Key success factors will be different at different points in the cycle. Service requirements in D, for example, are largely customer dependent but likely to be high due to constant technology turnover and associated integration problems. The technology churn is fueled by very powerful, relatively inexpensive new products (i.e. it gets cheaper to replace than maintain some systems) developed by B-type product groups. Exactly how to organize and manage around the B-D Cycle remains an open question. PCU's operating in the B-D Cycle would likely have to compete for sales to Digital's large systems integration groups suggested in E and would be expected to sell to other integrators as well.

A key idea is the relationship between the service company (E) and the product companies (B-D). Almost all product development, whether leading edge or commodity, must support the integration business. This means developing flexible, open platforms that can be easily customized for many different customer problems. DECmcc is a good example of a product built to support an integration business. Similar products are needed in the area of gateways and transport nodes.

Each team, during their presentation, presented a picture of what Digital's networks business would look like in their endstate. There were a number of themes that were common to most presentations:

- a large and increasing services component, suggesting Digital must be a premium service provider
- time-to-market as a critical success factor for products businesses
- a large and increasing D-type product component
- use of indirect sales and distribution channels for D-type products
- importance of exciting new applications that generate demand
- use of Digital's own network to provide service, support and software distribution
- importance of remaining a competitive hardware manufacturer, including chips, and stay independent of others for strategic materials

Critical Events

Events that were voted must(+) or must not(-) by four or more teams are listed here and commented upon. These were referred to as common events in the workshop. We characterize them as critical because they are important to four or more scenarios. Votes are shown by team (ABCDEF). The events fall into three primary categories: user environment, product opportunities and standards activity.

What happens in user organizations:

(+-+--)Large users demand vendor independent architectures
The scenarios that matter in the shared vision (B, D, E) all say that this must not happen. This should be watched very closely and the reasoning behind these votes explored more closely. NCRI believes that users are, and will increasingly, demand vendor independent architectures. Consider taking a public position on the definition of "vendor independent architecture". NCRI would suggest it means lowering the barriers to customers' switching from one vendor to another. It does not mean customers plan to abandon strategic vendors at their first opportunity.

(----) MIS reasserts central control over IT purchasing
Digital seems to think that MIS control equals IBM control. Again, this must be watched very closely. NCRI believes there is a very good chance that central control will be reasserted over the next few years.

(+ ++ +)50% of computer and network sales are through retail & direct mail
Indirect channels are important to D-like products, but not B-like products or E-type services.

(+++-+)Major users build heterogeneous distributed transaction processing
Digital's emerging strength in transaction processing may be a key success opportunity in the 1990s.

Product opportunities:

(+ +++)Scads of specialized servers hit the market
An example of a D-type product opportunity.

(+ +-+)Multivendor network management platforms take off
These platforms are an example of a product that is essential to the E services business, but not that useful to more narrowly conceived B and D product businesses

(+++++)A third of Fortune 500 have converted to multimedia e-mail
An icon for compelling new network-based application that drives sales of advanced network products.

(+++++)System management productivity increases dramatically
Although the importance of system management was discussed, little was put forth in terms of concrete ideas on how to solve the problem. The obvious importance of this event requires a dedication of resources within Digital.

(+-+ +)Multivendor distributed application platforms hit the market
This is another example of an E-type product, but one that is not necessarily needed by the B and D businesses. We might argue with team B's vote here. This kind of multivendor platform will be essential to producing the innovative gateways and applications that offer new value added to the customer.

(++++)Low-End PC multimedia applications take off
The group was attracted to multimedia and had a fairly realistic view of how it might penetrate the customer base. Low-end applications will be important harbingers of a more large scale migration to workstations and broadband in later years.

(+++ +)Electronic markets develop
Digital must develop inter-enterprise networking products and enablers in order to foster the electronic marketplace. The voting seemed to cede the business to the merged AT&T/NCR.

(+ ++ +)Silicon foundries profitable
The group clearly saw it as important that Digital retain a competitive chip making capability and that, in general, the ability to develop chips outside of Japan was key to preventing scenario F. Interestingly, however, team F chose this event too. They argued that Japan made all the chip making machinery now so they still retain control.

Standards activity:

(+++ +)SQL Access Committee announces standard
Database interchange standards are obviously key to the development of multivendor, distributed applications. Is Digital involved in this standard?

(+++ -)Data naming, addressing and object registry authority established
The group was somewhat pessimistic about this happening and you can see that E and D don't want it or need it. B's vote for this indicated their belief that basic interoperability standards were needed. The last exercise analysis of this event indicated that Digital could be doing much more to influence the success of this event if it wants to. Perhaps the US Post Office's offer to manage an e-mail directory indicates that this problem will be solved by numerous, partial solutions.

(+++ +)Iron clad electronic signatures accepted
This was an icon for solutions to a wide range of network security problems. Many of the network applications that the group hoped would drive further use of networks will be inhibited by insufficient solutions to the network security problem.

(+ + ++)Network management interoperability widespread
Note that the two product scenarios don't need this.

Desired Scenario

The exercise to build a scenario by selecting events leading up to the desired endstate combination was done hastily and the participants seemed to be suffering from event card overload. Nonetheless, there are some interesting aspects to the events selected for this scenario:

- it is heavy on advanced, cheap hardware
- it is big on low cost video and multimedia
- it envisions a world of distributed computing on PCs, workstations and servers, with little role for the mainframe or mini
- it stresses the need for solutions to the distributed systems management problem
- it stresses the need for solutions to the security problem, especially in ways that enable network-based markets to develop
- it suggests that I18N is important
- it stresses applications

- it pays attention to the carriers and the role they play in building network infrastructure, establishing intercompany directories, etc. It calls for a Digital/RBOC alliance

- it stresses innovation, even in areas that seem dominated by existing standards. The choice of the event "Lightweight protocols compete with OSI" is significant and would only make sense in the context of a B-type business
- it sees consortia standards as important, in general, and a broadband consortia as important, in particular

- it does not want large users to demand vendor independent architectures, consistent with previous voting on scenarios, this is a key event to watch

- it requires application programmers, the NREN, greater use of networking in schools and a relatively open Japan (a question which remains unanswered and debatable)

1997 Desired Endstate

What follows is an endstate written to capture the tone expressed by the collection of events that participants selected on the third day as the desired scenario. It is divided into the same groups used to build the wall display in the strategy room.

Systems Integration and Distribution

The systems integration business, and the leading companies in it, are driving the industry. Early in the 1990's the IT service sector was flooded with new entrants. When the shakeout came, it was the companies with scale that thrived. They could carry the risk and financial burden of very large scale integration projects. While the first sale to a new customer is often via the direct sales channel, subsequent sales to the same customers are generally through retail, direct mail and similar low cost distribution channels. Some companies are outsourcing processing and network management.

Product companies constantly work toward three goals - ease of integration, ease of use and low cost. Most product companies (even those captive to particular integrators) sell to all the integrators as well as through retail and direct mail channels. This approach is necessary if economies of scale are to be realized.

Public Networks

The Region Bell Operating Companies (sometimes referred to as RBOCs, Baby Bells or Regional Holding Companies) have been investing heavily in high bandwidth physical plant. RBOCs Metropolitan Area Networks (MANs) are now operating in the top 50 metropolitan markets in the United States. Japan, Singapore and some European public carriers are also upgrading their public networks. "Upgrading the Information Infrastructure" has become a political issue under the banner of national competitiveness; the result in the US was the National Research and Education Network (NREN). Congress has also taken up the complicated issue of changing the copyright laws to accommodate network transmission of information.

The public networks are offering a variety of advanced IS services including X.500 directories, key crypto services, billing for third party information services, etc.

Broadband is growing fast. Public carriers are doing well with T3 services and B-ISDN. Alternative carriers are also succeeding in the broadband market.

End User Productivity

End users have moved beyond the now familiar desktop application paradigms like spreadsheets and wordprocessing. Productivity benefits of new software and hardware are the key to purchase decisions. Workgroup productivity was addressed by using leading traditional applications as front ends to more sophisticated EDI, OLTP and e-mail applications. Most packaged software is now available in numerous languages. The US and EEC vendors are stressing the benefits of increasing workforce and R&D productivity.

Low end multimedia applications on PCs are best sellers. The virtual conference room paradigm (four quarter split screen) is the most widely accepted video conferencing system. E-mail is rapidly moving toward multimedia and has encompassed fax transmission.

Digital libraries and CD-ROM publishing are big business. Desktop based agent software is used to filter selected live news feeds and e-mail. The industry has learned how to price on-line information services and demand has soared.

In the consumer sector, debit/point of sale cards are now common. Card based payment systems have captured 40% of supermarket sales in the US. In the home, distributed interactive video games, developed in the 1980's for the military, have become the latest fad for kids and adults.

Several vendors brought solutions to the distributed DBMS problem to the market. This, and other factors, precipitated a major drop in mainframe software development spending. Production OLTP systems are being built on heterogeneous networks of PCs and workstations.

Efficient System and Network Management

System management productivity has risen dramatically. This started with the widespread adoption of interoperable network management platforms. Even low cost workgroup LAN management software was a big success; the increasing complexity of LANs made the need for management software readily apparent. One large vendor is now offering a low cost, remote, "total system surveillance" service that warns local managers and maintenance contractors when a particular machine has begun to operate outside of normal parameters. This approach to correcting problems before a failure is expected to be key to competition for large accounts.

Fast Cheap Hardware

The improvements in price/performance ratios of microprocessors continued unabated. One of the consequences are 100 MIPS desktop workstations available at prices close to the personal computers of the late 1980's. Lightweight, powerful portable computers using silicon instead of disk drives are showing up everywhere. The full cost of computer storage has fallen below the full cost of paper storage leading to the rapid growth of large scale data storage operations. Silicon foundries are now profitable.

In the communications arena, private networks are being assembled from building blocks like low cost brouters, smart hubs and <$100 ethernet cards for PCs and workstations. The low cost, specialized, server market is booming.

These trends have driven most organizations to distributed computing systems based on workstations, PCs, bridge/router interconnected LANs and numerous specialized servers. X-terminals networked to minicomputers are just not cost competitive. It has literally become cheaper to replace many old mainframe based systems than to maintain them.

Producers in the US and the EEC are moving to eliminate single source dependence for all key components. This has most directly hit Japan which is responding by slowly opening selected IT markets to foreign competitors. Other Pacific Rim countries have benefited from this policy.

Key Standards

Standards are used by the vendors attempting to quickly drive up sales volumes of new products. Consortia set the important standards. The Open Software Foundation's DCE/DME, the Object Management Group's APIs and the SQL Access Committee efforts were particularly important.

In the communications arena, the establishment of a data network naming, addressing and object registry was absolutely critical. Lightweight protocols have begun to compete with OSI. A broadband standards consortium has also been formed.

Customers never did demand truly vendor independent architectures.

Digital

Digital has introduced full token ring and TCP/IP support across its entire product line. The technology was pulled together using OEM agreements, acquisitions, third party developers and internal developers. Minimizing time to market and multivendor support were the key decision criteria.

Digital has moved to become a major supplier of protocol processing chips that are embedded in many other vendors products.

The Company also is in several joint ventures with various public carriers (RBOCs, PTTs, CATV operators, bypass carriers and IXCs) to develop and market advanced broadband applications and services for particular vertical markets.

NCRI Recommendations

We suggest below ten specific recommendations that arise directly out of the workshop discussions.

Recommendation 1
Management of the networks business should propose a consolidation scenario as the context for Executive Committee decisionmaking. It is not a sufficient decisionmaking context for the networks business itself.

Recommendation 2
The "B-D Cycle" should be the shared decisionmaking context in Digital's networks business. Managing the transition of products from B (advanced hitech) to D (cost/value engineered) is the key to the success of the networks business. The B-D Cycle must clearly be managed in the context of E (industry consolidation). Thus, the networks business needs to provide Digital's integration business with performance/cost competitive products and ones that support complex systems integration. To accomplish this goal, the networks business must be free to sell to and/or source from other vendors and integrators, enter into alliances and swap technologies as may be required to generate sufficient volume. Decisions to be made about each product (or product family) include:

Is it a B or D product?
Should the product be sold to other integrators or held for Digital only?
Does the product require a service component to sell it and, therefore, how should it be sold?

We believe that B and D products require very different management styles. Critical success factors will likely be different for the two types of products. How to organize to respond to this is an open subject.

Recommendation 3
Digital has an extraordinarily blurry self image. This must be changed if customers are to understand the Company. (Can Future Mapping contribute to this?) After the current round of reorganizations are over, top management must work hard on defining a NaC identity and propagating it to the troops and the Executive Committee.

Recommendation 4
Redefine the purpose and goals of Digital's standards activities. They should be directed toward giving Digital the ability to deliver standards-based solutions to customers in volume. Key standards are in network management, database access and security. Digital should concentrate on making standardized interfaces work with IBM, Fujitsu and other very large vendors. Challenge other vendors to interoperate with you.

Recommendation 5
Straighten out the role of the direct sales force. Their primary goal should be to sell integration services that then generate product sales. Also, direct sales are important for first time sales into a new information system opportunity. Subsequent sales should quickly use indirect channels. In general, the D-type component of the networks business will grow very rapidly (if Digital is successful) and mastering indirect channels for this business is key.

Recommendation 6
Decide which products or products families are D-type products. Start now to develop a D-type product capability. Learn to build cheap, fast products. Forget about the "Digital Premium" in this business. Volume is going to be key and that means adoption of de facto standards, cross licensing agreements and many, many OEM agreements. Ease of installation/deployment and ease of use without reliance on documentation are key aspects of such products.

Recommendation 7
Decide which products or products families are D-type products. The B-type business is typical of Digital. What's important here, as was stressed many times in the presentations, is time to market. Innovation is only innovative if it hits the market before a competitor's version. Hitting time-to-market windows is not something that NaC has been particularly good at in the past and the reasons for this require detailed study.

Recommendation 8
Digital should increase its contacts with carriers at all levels of the organization. This is one of the few future mapping meetings where computer vendor/RBOC alliances have not been seen as a critical event. The fact that this group seemed almost completely ignorant of the carriers' role and issues is a major blind spot and possible inhibitor of success. Since we talked about this frequently throughout the meeting, it is not surprising that it finally began to appear in the selection of events for the desired scenario.

Recommendation 9
Obviously, the Japanese threat is real and extremely serious. Equally obvious was the fact that this group did not know much about what Japan is doing. Team F did a great job, but more attention to understanding the networks business in Japan is definitely called for.

Recommendation 10
Establish a working group to decide on responses to the critical events on pages 7-10 of this document. The working group should also be responsible for monitoring external and Digital activities in the context of the desired scenario. Finally, the working group should consider how to educate the management team with respect to the role of public carriers and possible impacts of consumer markets.

Appendix

Short Descriptions of the Endstates Used in the Workshop

(A) A Future Dominated by Standards

Standardized, multivendor systems characterize most customer environments. Customers have willingly traded rapid technology advancement for a rich set of business applications running on systems that are simpler to install, easier to use and less likely to become obsolete. Interoperability, portability and backwards compatibility are mandatory requirements for any IT purchase. Customers demand to be able to connect to other firms easily. Anything runs on anything - almost. Anything talks to anything - almost. Unix is popular and remaining operating systems are POSIX compliant, making application portability easy. TCP and OSI are the dominant communications protocols, with OSI growing the fastest. The open market fostered by standards has brought prices (and margins) down and raised the level of competition to applications, brand and efficient distribution. The big players act as distribution channels for others.

(B) A Future Driven by New Technology

Exploiting information technology is a corporate and national priority in many countries. Customers are willing to invest in new technology to spur creative new business applications; backward compatibility is not critically important. Knowledge and information technology are generally regarded as resources of tremendous value to be leveraged and protected. Price/performance ratios for microprocessors have exceeded 50%/year since 1990. It is clearly cheaper to throw away old systems than to maintain and integrate them with new systems. The healthy IS investment environment means increased technology turnover. Time to market and innovation are key competitive issues, not enterprise architectures. Advances in hardware constantly open new market opportunities; two $1 billion companies are on the scene that did not even exist in 1991.

(C) The Global Data Utility - Network Nirvana

The success of networked applications went hand-in-hand with the growth of the networked form of multinational enterprise. MIS managers make architecture-driven decisions to evolve toward distributed computing and transform themselves into network providers to support participation in global markets. Networked machines operate as a flexible, distributed multiprocessor, constantly reconfiguring to meet demand for MIPS. A limited set of standards support high levels of interoperability and automated system management. Online directories tracking objects, people and databases are indispensable tools that guide users and software agents through the global information utility. Competition is based on cost, willingness to assume risk, and global systems expertise. Developed nations have lifted most regulations on public carriers, who have become major WAN and systems operators focused on building world net. The POS business has evolved into a full fledged electronic marketplace where income from transaction fees supports operations.

(D) An Affordable Future - Just Make it Work

Cost concerns and low risk characterize customer buying habits. The work group is the focus of vendor sales efforts. Off-the-shelf software and user-generated applications account for most new systems. Central MIS staff has been heavily cut; large architecture projects are difficult to sell to disenfranchised MIS managers who focus on site specific "fix what I've got" incremental purchases. Growth in international markets has forced all vendors into global marketing. The basis of competition is price, flexibility, ease of use and worldwide distribution to workgroup buyers. Hard times have forced vendors to cut back on resources devoted to standards activities. Facilities management and systems integration are hot businesses; large vendors treat them as new distribution channels. Large systems contracts that do come up are always bid by multivendor teams creating very complicated partner/competitor relationships.

(E) A Future Dominated by Six Big Vendors/Integrators

Most large users want two, maybe three, strategic vendors that have resolved interoperability and portability problems among themselves. These customers buy strategic architectures, not boxes. They don't want to deal with numerous IS vendors because it is too expensive. Architecture-based solutions offer advanced technology, better integration, broader functionality and application level interoperability sooner than standards-based architectures. The limited number of strategic vendors in most companies makes interoperability relatively easy, although portability is still limited. Intense consolidation has reduced the number of global vendors to six multinational giants; only three are US. Advanced architecture capabilities give vendors power over many OEM suppliers and customers; they also create very high barriers to entry. At the forefront of new technology, dynamic competition between new startups still flourishes. Systems integration is a major business. Competition is based on large scale systems competency, account control and capacity for risk.

(F) A Future Dominated by Japan

In offices, computers are turning into TVs. In homes, TVs are turning into computers. Japanese products are fully standards compliant, inexpensive and designed for a new generation of savvy buyers. Consortia developed standards are the critical glue in most systems. TCP/IP and OSI are the major communications protocols. UNIX grew fast in Japan because it gave manufacturers independence from US operating system vendors. Consumer and workgroup markets have merged to become a giant retail distribution business where technical development is driven by unprecedented sales volume. Alliances and intense consolidation reduced the number of global vendors to five multinational giants; only two are US. NTT's Multivendor Integration Architecture has been widely accepted and opened large Japanese markets to international suppliers. Japanese vendors work hard to project a truly transnational character in an effort to help global users forget about the nationality of their systems. The center of activity in the industry has shifted to Japan and the Pac Rim, which is the fastest growing market for all classes of machines.

< end of NCRI report >

Networks Vision for the 90's

Shared Networks Business Model

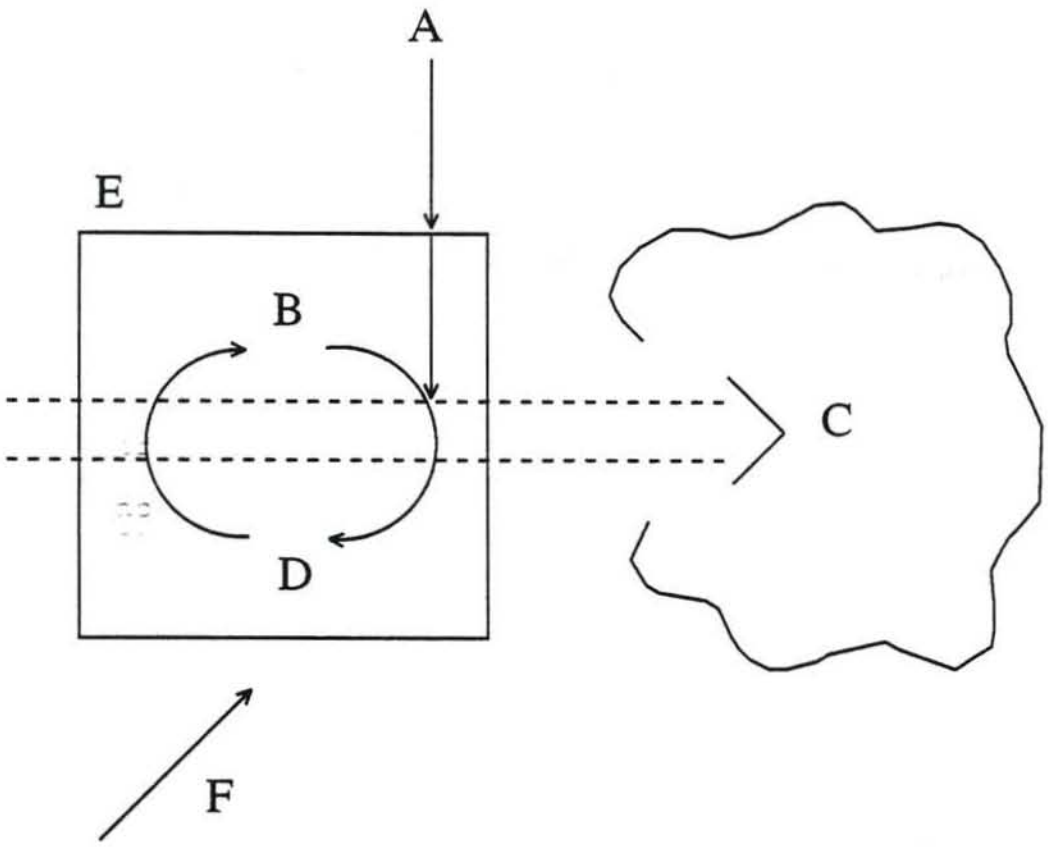

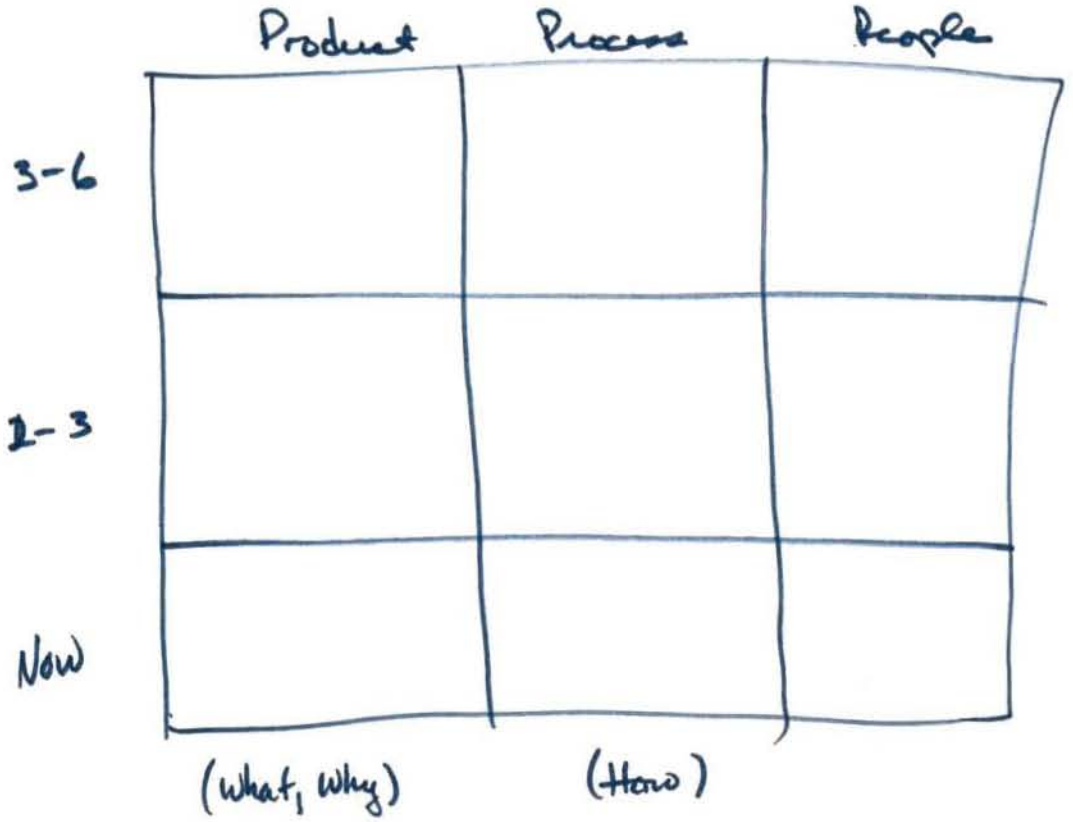

CSF Matrix.

Availability:

The Customer's Point of View

by Ron Rocheleau

CSSE Operational Services

MRO1-1/OS

DTN 297-3515

12 September 1990

Preface

The purpose of this paper is to explore issues that have an effect on how the customer perceives the availability of data center computing resources in order to develop a consulting service that will enhance data center availability. One goal is to identify some of the activities that might be performed as part of an availability service. A second goal is to identify the elements that would need to be considered in developing a tool for modeling availability. A third goal is to at least recognize some of the engineering issues that must be considered when designing new products. A final goal is to identify some opportunities for new architectural or state of the art advancements in product development that could lead to greater availability.

No attempt has been made to limit the scope of this paper. Indeed the intent is to explore as many potential factors as possible that have either direct or indirect consequences for the availability of customer applications. During the next stage, the structured analysis of the availability service and availability modeling tool, we will focus more clearly on what factors can be influenced in the Field.

Much of the material presented here has been accumulated through experience over the past fifteen years in the Field on large data center sites and in CSSE, from readings in technical books, journals, and papers, from customer surveys and market studies, and from seminars and forums. The Generic Fault Management Model developed by Gerhard Heider, and the Digital Enterprise Fault Management Architecture being developed by Frank Robbins and other consultants are key references in the sections on both hardware and software fault management, as well as presentations and papers presented at the last several Reliability Symposiums.

This revision is being circulated to various consultants within Digital for review and comment. I expect that many new ideas will be forthcoming and hope that constructive feedback will be provided to correct any mistakes in concept, content, or style. A follow-up review will then be scheduled with the consultants who wish to participate. The final release of this document will incorporate the advice and contributions of many people.

In addition to soliciting the reviewers for specific comments to improve this white paper, I am also requesting the reviewers to help define the user requirements for an availability service and an availability modeling tool. A form for this purpose is included at the end of the document.

Ron Rocheleau

Revision 1.0	21 August 1990	First Review Draft
Revision 1.1	12 September 1990	Incorporated initial reviewer comments

Contents

Availability: The Customer's Point of View

1 Introduction

Availability is a topic of keen interest to Digital's customers. The availability of computing resources can have a direct impact on the success of the customer's business. Digital must address the customers' need for availability by providing consulting services to help customers configure the proper mix of computing resources and services to meet their business requirements for availability. This paper explores the customer's perception of availability and how a service might help to achieve the customer's business requirements for availability.

1.1 The Engineer's Definition of Availability

Internally to Digital, we have several technical definitions that help to express the attributes of product design that impact availability. The traditional availability metric is expressed by:

$$\text{Availability} = \text{MTBF} / (\text{MTBF} + \text{MTTR})$$

This formula expresses the impact of product reliability (Mean Time Between Failures, MTBF) and repairability (Mean Time To Repair, MTTR) on the availability of the product. The underlying premise is that until a product 'fails' and only for the period during which it is being repaired, it is unavailable. When analyzing a product design this is a reasonable metric for purposes of setting reliability goals and comparing the availability metric between generations of the same classes of products.

A broader view of availability has been published in several documents originated within CSSE by Gerhard Heider, et al. (Generic Fault Management Model), Frank Robbins, et al. (Digital Enterprise Fault Management Architecture). This view of availability considers the impact of service delivery, installation quality, and 'fault tolerance' in addition to the product reliability and repairability. Frank also explains the

need for considering fault tolerance, fault avoidance, and fault management in new product engineering to enhance availability.

The formula for this concept is expressed by:

Availability = f(Quality + Reliability + Fault Tolerance + Service Delivery)

The notion is that Availability is a function of Quality, Reliability, Fault Tolerance, and Service Delivery. While not a mathematical formula, this 'recipe' is of great value to CSSE and Engineering for the purpose of setting goals for new products and in establishing processes for new services. It increases the scope of the basic availability formula to take into account the processes and logistics of servicing a product in the Field.

1.2 How Do Customers Define Availability ?

At a recent presentation at the Reliability Symposium, John Shebell introduced the concept that the Customer point of view of Availability is much broader. The terms John used were Robustness, Dependability and Suitability. When the customer looks at how business is impacted by availability he might ask such questions as:

"Are my data center operations available to meet my business requirements and financial goals?"

"Do computer resource, applications, or operations failures interrupt my ability to conduct business as expected?"

"What can be done to protect and insulate my business from computer resource problems?"

However, customers still use the term 'Availability' to include our engineering definitions of availability and reliability, but also with the notion of Dependability and Suitability in their business environment. They don't measure availability solely in terms of product MTBF or MTTR. They measure availability in terms of meeting their business goals and requirements and the bottom line. A lack of computer resource availability can impact their business, indeed it can put them out of business! Availability affects their ability to serve customers, to meet business financial goals, to meet time to market requirements, to respond to business demands, the quality and reliability of vital information, their ability to compete, etc.

1.3 An Extended Formula for Availability

I'd like to suggest a broader definition for availability can be expressed by the formula:

Availability = TTBI / (TTBI + TTBIR)

Note: TTBI = Time-To-Business-Interruption,

TTBIR = Time-To-Business-Interruption-Recovery

The concept is similar to the basic availability formula but as you'll see it takes into consideration the more general CSSE availability recipe, plus much more. I'll concentrate on what factors have positive and negative effects on availability. In developing a consulting service to improve computer resource availability, these and other factors must be examined, and recommendations made to the customer for maximizing availability consistent with business goals.

TTBI and TTBIR are comprised of many factors that can be grouped into several major categories. The first major category, and the one most often considered, is the Hardware. The second major category is the Software, including both the operating system and applications. The third major category is the Customer's Data Center Operations. The forth major category is Service Delivery. The final major category is the physical Customer Site. Each of these categories possess attributes that impact either the TTBI and/or the TTBIR.

Various studies indicate that operations account for up to 30% of failures and hardware and software each contribute about 25% to 35% of the failures. Problems in operations that lead to system unavailability are not always recognized in these studies. Service delivery is generally not broken out separately but is dispersed in the other categories. The physical customer site environment also contributes to failures but generally is not broken out as a separate category. These percentages vary greatly from one study to the next. The particular mix of these categories will of course vary from customer to customer and application to application. A model for availability must allow for flexible modeling dependent on the customers business requirements and the many variables that make up each of these categories.

Some of the factors that affect availability will reduce or increase the TTBI or TTBIR based on the probability of occurence. Some factors have a direct observable impact, while others may have a more indirect impact. This paper does not yet attempt to quantify the probability of occurence or amount of impact on TTBI or TTBIR that each of the factors and variables account for. That will be explored and rationalized as the availability model and availability service are defined and developed. The impact will obviously vary over a wide range depending on the customer's business requirements. The reader may wish to consider whether the factors described here affect TTBI or TTBIR or both. The author solicits your insight and knowledge in helping to formulate and design the availability model to be representative of a "real" data center lifecycle.

This figure does not attempt to accurately depict the percent distribution of these categories as they might apply to any particular data center. They will vary from customer to customer and situation to situation and may be determined only after complete analysis of customer requirements and applications. It is my impression, however, that on average this figure reflects the typical distribution.

Some factors affecting the hardware, software, and service delivery categories are clearly design issues that must be addressed by Engineering and CSSE in the design and development of new hardware and software products. For instance, a hardware device with no error detection affords no opportunity for Symptom Directed Diagnosis (SDD) techniques to be used to determine the onset of a potential failure. Proactive measures cannot be applied. Maintenance is more difficult and more likely to be ineffective for certain classes of faults. Features that enhance ease of installation, ease of setup and change, and especially ease of use and operation of hardware and software directly affect availability.

In addition to the hardware and software reliability, maintainability, and fault tolerance characteristics, there are other important factors that directly impact availability. System configuration, redundancy, fault management software, application partitioning, system management, to name but a few. When the inevitable failures do occur, a period of time expires before full restoration of service. This period of time is affected by factors such as contract coverage, service response time, logistics, data backup and restoration, maintenance and repair accuracy, SDD and remote support and services, etc.

Some factors that affect service delivery and operations are issues to be addressed through contractual agreements for service content and coverage, as well as operations staffing, procedures, and policies. The environment of the customer site often directly impacts the reliability and availability of the hardware. There is considerable amount of interaction and dependency between all of these factors.

This paper will not attempt to quantify all the variables that impact availability. When the availability model is analyzed and specified an attempt will be made to model many

of these variables and factors. For now we need to recognize as many variables as possible in order to create a better model for availability from the customer's perspective.

The next few sections of this paper will explore these and other factors that impact the customers perception of availability. The degree of impact that each of these factors have on a particular situation will of course vary.

2 Hardware Impact on Availability

There are several factors that determine the contribution to availability of the hardware devices in a given data center configuration. These factors impact either TTBI or TTBIR or both. The factors described below are:

- Hardware Reliability (MTBF)
- Hardware Repairability / Maintainability (MTTR)
- Hardware Error Detection
- Hardware Fault Tolerance
- Hardware Fault Avoidance
- External Fault Management / Recovery
- Hardware Restoration / Recovery Time
- Hardware Failover Time
- Hardware Installation
- Hardware Documentation

2.1 Hardware Reliability (MTBF)

Hardware reliability has a direct impact on availability. The overall impact on availability can be diminished somewhat through fault management, fault avoidance, and fault tolerance. Other factors that affect availability reduce the impact of hardware faults. For example, a new product development that improves hardware reliability by 50% from one generation of product to the next may have a positive impact of only 10%-15% on overall system availability. Significantly, other attributes of the product, such as error detection, fault management, diagnosability, installability, software/firmware reliability, conformance to standards, design faults, etc. may have additional positive or negative impact.

The reliability of hardware products is usually expressed in terms of Mean-Time-Between-Failures (MTBF), a predicted estimate based on component reliability data. There are several widely used tools in the industry for predicting hardware reliability. While the reliability prediction tools do not produce the same estimate of predicted reliability (in fact, the MTBF metric produced varies significantly from tool to tool) the tools serve to validate the engineering advances introduced in a new generation of device and quantify the relative improvement in reliability. For some hardware devices (i.e. memory or disk drives) the tools may take into account some fault tolerance built into devices in the form of ECC mechanisms.

The reliability prediction tools generally do not take into account the impact of software/firmware/microcode bugs or basic design faults that lead to unexpected failures and reduce the actual time to failure in the field. Products that contain software, firmware, or microcode suffer from all the same reliability issues as operating systems

and applications software. Yet the reliability of this code is not factored into the hardware MTBF calculations. For these reasons, the predicted MTBF and the actual mean time between failures in the Field often differ by a wide margin.

Not all hardware failures result in a loss of availability. Failures at the unit level may or may not have any impact on the customer's business, may be made invisible to the customer through higher level fault management or hardware redundancy, or may occur at a time when the product is not critical to the business application. The goals of configuring a system for fault avoidance and fault tolerance, and employing a means of fault management are to mitigate the effects of faults in the hardware and software. To the extent that these strategies are successful device failures will not cause an unacceptable impact on availability.

Hardware failures differ from software failures in that hardware failures are *guaranteed* to eventually occur. In the physical world devices eventually breakdown. In addition to the reliability of components there are other real-world causes for hardware failures. The specification and design process can result in a product design that has flaws that are not detected until failures occur in the Field. The manufacturing process must be continually monitored and adjusted to prevent process problems from resulting in Field failures. Ongoing Reliability Testing is used to attempt to catch these problems before they proliferate. The Field installation and repair processes and procedures, and the handling of spare parts in the Field introduces other sources of failures. The data center environment, temperature, humidity, and power can all lead to premature component failures and chronic intermittent faults. Transient faults generally caused by the environment will cause systems to fail but may not be repairable. The design of systems must take these other sources of failures into consideration to ensure that an effective degree of fault tolerance and fault avoidance is built into the product to meet customer expectations for availability.

Hardware incompatibility between hardware components and software components in the data center can cause intermittent and undiagnosible faults. As configurations of systems become more complex and more "open system" hardware from numerous vendors are mixed, compatibility determination is becoming impossible. It's important at a minimum to have a complete audit of the data center to at least understand what elements are in the configuration. Beyond that, standards must be developed for collecting and maintaining compatibility information about resource configuration. It may be possible for such data to become a part of the hardware and software products so that automatic verification and validation can be done between the products themselves. For the present, the issue of hardware incompatibility may impact the apparent reliability of the data center. It is therefore essential that great care be taken during installation of new hardware and software to at least make sure all known incompatibilities are avoided.

2.2 Hardware Repairability / Maintainability (MTTR)

A number of product attributes contribute to the repairability and maintainability of hardware products. The quality and coverage of diagnostic exercisers affects the

timeliness and effectiveness of repair. This is especially important when testing to verify a completed repair. Even if a product has good error detection and reporting, it is essential that repair be verifiable by an internal or locally executed test process. If it is not, then there exists a real probability that the hardware will be returned to service unrepaired and result in a subsequent failure.

The MTBF metric does not take into consideration faults brought on by poor mechanical design that hinders effective maintenance. During repair faults are injected by such things as cables and other parts that block easy access to FRUs, interconnects that become contaminated or damaged. Inadequate maintainability features make diagnosis difficult and error prone and result in repeated failures and replacement of the wrong FRUs.

The repairability of a device is specified by the metric MTTR, Mean-Time To Repair. This is the average estimated time for a trained service engineer, using the built-in tools, diagnostics, and support to complete a typical repair. Note that while many products have an MTTR of 1-2 hours, there is a wide variance in actual repair times experienced in the field dependent on the complexity of the actual repair and the experience of the service personnel.

2.3 Hardware Error Detection

Error Detection indirectly relates to availability. The degree and accuracy of error detection (fault coverage) impacts the amount and types of fault tolerance that are possible. Dynamic error detection impacts the ability to perform SDD fault isolation. Error detection also impacts the ability of maintenance activity to be successful in effecting timely repair. Simply detecting an error is often not sufficient to isolate the underlying fault. Error detection mechanisms must also capture relevant data to assist in fault isolation. Otherwise, repairs may not be possible or lengthy repair times could impact availability.

2.4 Hardware Fault Tolerance

Faults can be tolerated by providing some form of redundancy that restores data to a valid state with acceptable performance impact. A retry mechanism is also useful for recovering from intermittent or transient faults. Techniques such as ECC, multiple sourcing of data, duplicate copies of data, results voting, shadowing, etc. can improve fault tolerance. At an application level, faults that are tolerated allow the application to produce correct results. It's important to understand that while faults can be overcome by using these techniques, there is sometimes a performance penalty that may render an application essentially unavailable. Fault tolerance allows for some types of fault repair to be deferred to a time when application availability is not impacted by repair. Fault tolerance techniques usually result in an eventual need for repair on most devices, if only to stop the continual incidence of error reporting.

The most common approach to fault tolerance is to configure redundant hardware or software to be available in the event of a failure on one of the redundant units. However,

it is important to understand that redundancy alone does not guarantee greater availability. Generally speaking, redundant configurations have a common interconnect point, such as a bus or controller in the case of hardware, or a database or default parameter values in the case of software. These common interconnect points can be interfered with by failures in one unit, thereby rendering the redundant unit also defective. Another common problem with redundancy is with specification and design faults inherent in both redundant units. When conditions are such as to cause a failure in one unit, they may also cause the redundant unit to fail. As an example, these types of system failures have been experienced several times during NASA Space Shuttle launches that were aborted as a result of design flaws in both hardware and software systems that caused redundant systems (up to 5 levels of redundancy!) to fail simultaneously.

The most effective approach to fault tolerance through redundancy is to configure multiple units of unique design and implementation. This is of course a very expensive approach but by far the most likely to succeed. This approach must be reserved for situations where there is only one opportunity to accomplish a task, there is high risk of loss of life, or the costs of failure are unacceptable. Here again, NASA has used this strategy in numerous space probes and to great success in such probes as the Voyager and Explorer series of space craft. Despite multiple system failures they have continued to meet and exceed most mission expectations.

A retry capability is another approach to fault tolerance which allows hardware to recover from intermittent or transient conditions. This technique is most often seen in media devices such as disk drives or tape drives or in network or communications protocols. If the fault is overcome in an acceptable number of retry attempts availability may be preserved, though with some loss of performance.

2.5 Hardware Fault Avoidance

Hardware Fault Avoidance to improve availability can be achieved in several ways. The manufacturing quality and reliability of hardware can be optimized to use selected high-quality components, incoming inspection testing and qualification, and burn-in of components and manufactured assemblies. This form of fault avoidance is accomplished by removing defective components prior to manufacture of assemblies, and testing and stressing assemblies prior to shipment to customers. This approach is often used (again by NASA and military contractors) to improve the quality of instruments sent into space by testing them in hostile environments for months and in some cases years before final assembly in space probes. A less intense screening process is sometimes used in the manufacture of computer systems to increase reliability, particularly when new technologies are first introduced in manufacturing as a way of preventing infancy failures in components and manufacturing processes from getting to the Field.

The computer design can also use fault avoidance to overcome detected errors as components fail in the field. Another approach to dealing with faults is to avoid using the area of the hardware producing errors (sometimes referred to as "automatic repair"). This technique is widely used in media devices such as tape drives and disk drives.

Media defects are marked as bad and vectored to another area on the media. The VAX-9000 employs fault avoidance techniques in the memory system by tracking memory errors to detect DRAM defects and mapping out memory to avoid using those areas of memory. Fault avoidance allows for some types of fault repair to be deferred to a time when application availability is not impacted by repair. If resources are sufficient, a large number of simple faults may simply be avoided without any need for repair.

2.6 Hardware External Fault Management/Recovery

Fault Management can be accomplished using multiple levels or layers, both internal or external to the product incurring a fault. Fault Managers can use the techniques of fault tolerance or fault avoidance. Fault Managers usually employ some means of error/fault reporting. The positive effects of fault management are generally not accounted for in a products MTBF. The MTBF is often calculated solely using predicted component failure rate data.

It's very difficult to measure the impact of fault managers. In some cases, if a fatal error is preceded by a number of recoverable errors, or if the fault manager is able to employ fault avoidance to stop a fault from recurring, then a fault manager allows repair to be deferred to a non-critical time where application availability will not be affected. Fault managers also employ some means of error information gathering and analysis that can result in more accurate or timely repair. The actual benefit provided by a fault manager is a function of how well it can handle the majority of common faults. Quite often, a device will have some fatal failure modes that cannot be successfully managed by a fault manager or the fault manager may be able to manage the fault but the time penalty is too great and availability is lost anyways. Even in these cases, it is required that error information be captured in order to aid the repair process and minimize MTTR.

2.7 Hardware Restoration/Recovery Time

When a device fails and is repaired it may require additional effort to make it available for use. Disks in particular may require initialization of the media, formatting, and restoring of backup data. Until the customer data is brought up to date and is accurate and consistent, from the customer's point of view the device is unavailable. If some form of fault tolerance or fault avoidance is not employed the application availability is also lost.

Another situation common to disk drives is disk fragmentation. The disk is not broken but the data is! Access time may be so great that performance diminishes to the point of unavailability. Some defragmenters work in-line with applications and gradually eliminate this problem. Others run stand-alone and cause the customer to lose availability while they run if there is no opportunity to defer use to scheduled maintenance periods.

2.8 Hardware Failover Time

Failover Time is the period of time that a device or access to data is blocked or performance is inhibited while resources are being reconfigured. In some cases, failover time may be a matter of milliseconds to seconds, while in other cases it may be 10s of seconds to 10s of minutes. Depending on the application, these times may or may not exceed the requirements for performance. For example, the customer's application may be dependent on a real-time response that cannot exceed 15-20 seconds delay. If failover takes a minute, then the customer's application will be down despite the resumption of system availability. The customer's application response time and configuration failover times must be factored into the configuration decisions to ensure it meets the application demands. This could mean adding a "shadow drive" instead of a "failover drive". Or adding another system to the cluster and dispersing the application to minimize cluster failover times. While the use of shadowing, device failover, or clustering can minimize the loss of availability if sufficient redundant resources are accessible these strategies can add significantly to rebooting or restarting time. The tradeoffs must be considered to ensure the overall effect on availability is accounted for.

2.9 Hardware Installation

Procedures for installing hardware or installing options and FCOs can lead to availability problems. To ensure compatibility with other hardware and software it may be necessary to perform upgrades to other units, change application or operating system parameters, or alter operating procedures. The length of time for the installation may range from minutes (out of the box and plug it in) to days. The installation of additional equipment may require taking down other systems in the data center that are in use by the customer. Sometimes installations inadvertently cause downtime by unintentional interference with running systems. It is also important to have an installation verification procedure to ensure that the installation is correctly accomplished and will not lead to downtime immediately after being placed into customer use. A lack of proper planning may result in a partial installation before the discovery that additional hardware or software is required to complete the installation. Backing out of a partial installation may be difficult and lead to lengthy unexpected downtime.

2.10 Hardware Documentation

Hardware Documentation can plan a role in availability. Installation procedures must be clearly written. All compatibility issues should be addressed in the planning section of the installation procedures. Environmental and physical requirements must be stated, as well as supporting software requirements. User's guides must explain operation of the hardware, especially operator service requirements and procedures. A section of the user guide should explain actions to be taken under exception conditions to prevent damage to equipment or data and to minimize recovery time. Service manuals must fully describe all maintenance procedures, diagnostics, and preventative maintenance procedures.

3 Software Impact on Availability

The operating system software, system management software, proprietary applications software, and customer developed or procured software all contribute to the viability of a system. The classic formula for availability uses only reliability and repairability metrics for hardware devices in estimating product availability. To estimate application availability as seen by the customer, it is necessary to consider the reliability of the software and the ability of the software to continue to produce correct, timely results while correcting for software errors; the ability to recover or restore application data and resume; the repairability and maintainability of software to allow fixes and upgrades; the impact of software faults on other applications or the operating system; the ease of installation and ability to adjust the installation to the environment; just to name a few. The topics discussed below are:

- Software Reliability
- Software Repairability / Maintainability
- Software Error Detection
- Software Fault Tolerance
- Software Fault Avoidance
- Software External Interference
- Software Restoration / Recovery Time
- Software Ease of Installation
- Software Usability

3.1 Software Reliability

Software faults differ from hardware faults in that *all software faults ship with the product*! Many of the faults lie dormant or only intermittently raise their ugly heads. Software does not degrade or deteriorate with time, though the data under control of the software may corrupted or the environment in which the software is used may change. The best likelihood of improving software reliability and availability lies in technological advances in the tools, processes, and techniques for developing *error free software* and for detecting and managing errors when they do occur.

Software reliability is a prime goal of software engineering in the development of new software products. Many new CASE and 4GL tools are available today. Advances in software analysis and design methodologies have sought to remove some of the "art" from software development and replace it with engineering discipline. Unfortunately, there is a long way to go before it will be possible to test and quantify the reliability of software.

Generally accepted estimates of software reliability range from a low of 3-5 flaws per 100 lines of code, to 3-5 flaws per 10,000 lines of code. A flaw does not necessarily

render a system failure. Flaws may be in the form of coding mistakes, logical mistakes, design faults, or specification faults. The trend in software reliability is positive as a result of recent developments in methodologies and automated tools and languages. Structured analysis, structured design, structured programming, walkthroughs and reviews, CASE tools, 4GL languages, test coverage analyzers, code style/syntax analyzers, performance analyzers, etc. have improved the quality of software. Unfortunately, the improvements in software engineering have not kept pace with advances in hardware reliability. As a result, the contribution of software to system failures is increasing! Software is now embedded in virtually every high-technology product from portable telephones to the Strategic Defense Initiative. To paraphrase Edward Yourdon...Using current state of the art technology to develop the estimated 100 million lines of code needed to deploy SDI, we are left with the chilling reality that 30,000-50,000 bugs would exist it that code...any one of which could obliterate all of mankind! Indeed, a lack of availability on an SDI defense system made known to a potential enemy might lead that enemy to launch an attack during the window of opportunity (TTBIR). A bit extreme perhaps (I hope) but it does drive home the reliance of the customer on software reliability and quality.

It is possible for the customer to take some defensive action to reduce the risk of software reliability impacting his business. Care must be taken when introducing new software releases, updates, or new products to do so in a controlled fashion. Test systems running production simulations can be used to verify products for a period of time before actually using them in production. It may be advantageous to skip one or more upgrades of software that do not provide essential improvements in functionality. Frequent backups and/or checkpointing can be used to minimize the impact of failures when they do occur.

Software compatibility, both with the hardware and other software components, is essential to reliability. Software incompatibility with a specific revision of hardware, the operating system, other applications, or pre-existing databases is a major cause of downtime and intermittent failures. Installation processes sometimes attempt a degree of compatibility checking but often the incompatibilities are not known by the development team so the installation process cannot adequately determine compatibility. This is an engineering issue that needs architectural work to make it possible for an application to specify resource requirements at runtime and determine whether or not a reliable environment exists for running the application.

3.2 Software Repairability / Maintainability

The Repairability and Maintainability of software is a design issue that should be addressed during product development. Software must be coded and documented so that future maintainers can effectively fix problems when they inevitably occur. The software design should also anticipate failures and incorporate some mechanism for capturing data to aid in isolating faults and for producing reasonable error messages to users. In some cases, software may fail for lack of resources and the fix is as simple as adjusting various operating parameters. Software should produce adequate messages to

aid in resolving issues of insufficient resources, privileges, or security protection to prevent their recurrence.

The maintenance of software often introduces new faults while fixing the old. Studies of the software development lifecycle show that maintenance of software accounts for about 80% of the lifecycle cost and at least as much manpower effort. Fixing bugs becomes a difficult process after the development team dissolves and the maintenance team takes over. The engineering process may not have produced adequate definition and specification of the product. Often there is no one around familiar with the implementation details; the source code may be difficult to relate to the specification; the specific revision of implementation language may have been obsoleted; and a host of other problems increasingly make the process of maintenance more difficult with time. As a result, bug fixes often introduce logical or algorithmic faults, or undesireable side-effects. Field installed patches often introduce unexpected or unintended corruption of applications that may be difficult or impossible to correct. For these reasons, a test process to verify a new release or field modification should be considered before using such software in production if it is at all feasible to do testing with a reasonable simulation of the production environment.

3.3 Software Error Detection

Software Error Detection is primarily a software engineering and design issue. Software Error Detection is important in minimizing downtime only if the detection of an error also provides some information about the nature of the underlying fault. Very often applications will detect an error and simply signal a failure but not provide any information. To fix such a fault it may be necessary to recompile the application or run some monitoring process and wait for a subsequent failure or several failures of the customer's application before the nature of the fault is understood and the solution can be worked on. This results in multiple incidents of downtime and reduced availability to the customer.

Most hardware designs have built-in mechanisms that automatically collect data for fault isolation. Most applications do not. An independent *Software Fault Management Architecture* is required to address this need. Operating systems and applications should employ mechanisms that gather this type of information automatically.

3.4 Software Fault Tolerance

Software Fault Tolerance is primarily a software engineering and design issue. Software Fault Tolerance includes not only faults within an application but also errors produced as the result of hardware faults or external interference from other software. Too often, an operating system or application with sufficient state/process information aborts with a cryptic error message rather than retrying an operation. It's probably not reasonable to require that each application understand sufficiently the nature of all possible faults so that it can decide when recovery is possible. For applications to really become fault tolerant, a new architecture employing a *software fault manager* that works in

cooperation with applications and the operating system and does not unduly burden the application developer needs to be added to our CASE tool suite and language compilers. For the time being it's hard to rate the fault tolerance of an application.

For applications with long run times availability can be enhanced by *checkpointing or journalling*, that is, by storing intermediate results or history of events. When a failure occurs for whatever reason, it is only necessary to rollback or restart the application from the checkpoint or rerun the history journal. Depending on the application and frequency of checkpointing, it may be possible to actually have the application fail, rollback, and resume without the user being aware a failure even occurred. The operating system may also preemptively act to circumvent faults and provide immunity to applications.

3.5 Software Fault Avoidance

Software Fault Avoidance is a difficult concept to fathom. One area where fault avoidance might be applied is to the user interface of an application. When requesting the user to enter data, an application may allow data to be entered without restriction and then attempt verification. An algorithmic error may not catch all potential errors, such as the user pressing the <compose character> key, thus leading to failure. A fault avoidance approach would present the user with a list of valid data entries from which to select a valid choice. Another example would be a software installation script validating system parameters and licenses prior to installation and automatically verifying proper installation. This is an area that the customer (system manager or software support) may need to consider when setting up the operations and application script files to avoid some potential operator induced faults. More emphasis and supporting development tools are needed to improve the quality of software engineering to address software fault avoidance. Here is another area where a *software fault manager* could assist applications by performing virtual reconfiguration of the application environment to avoid some classes of hardware and software faults.

3.6 Software External Interference

Software external interference is caused by an application or operating system software fault or a user action that corrupts or destroys data, impacts performance, or impedes the function of another independent element of the customer data center. For instance, a bug in one product that accesses a database common to several applications may corrupt the database rendering some or all of the applications down. An operating system bug may cause an application to abort. A bug in a development language may cause an application written in that language to improperly function, produce incorrect results, or fail outright. An improperly assigned pointer may cause data corruption. A conflict between applications competing for the same system resources may cause the operating system or applications to fail. Many types of software faults have the potential for impacting unrelated applications. The user of an application or others with access to the file system may inadvertently or deliberately corrupt or delete data files and application programs.

3.7 Software Restoration/Recovery Time

When failures occur it is often necessary to perform some type of data recovery or data validation operation prior to resumption of the business applications. For some types of failures it may be necessary to resort to restoring a previous backup and rerunning the applications executed since the last backup. The time it takes to perform these activities, in so much as they prevent full availability of the customer's application, must be added to the time to effect repair of the problem that caused the failure. Software application *checkpointing, journalling,* and a *backup/restore policy* are several approaches to evaluate for minimizing the software restoration/recovery time.

3.8 Software Ease of Installation

The ease and accuracy of installing software can significantly impact the availability of both of the operating system and the applications. Installing VMS is generally a relatively automatic operation. However, in order to add applications, users, and configure VMS for use numerous parameters must be adjusted. The minimum value for a parameter may be sufficient for the installation or operation of an application, but may not be the correct value for optimum use or performance. Some critical parameters set to inappropriate values may lead to system or application crashes. However, determining the appropriate values often requires an expert. Nonetheless, it is essential that the operating system be correctly set up or availability will be impacted.

Applications software introduces additional complexity. Not only must system parameters and quotas be adjusted just to perform the installation, but interaction and compatibility issues need to be addressed. Release Notes and Installation Guides address some of these issues but much is still left to be determined by the installer.

As a result of the complexity of installing some software applications, they are often not properly and optimally installed. This causes applications performance problems, applications crashes, and in some cases, system crashes. The installation of one software product may interfere with other products previously installed or require those products to be upgraded. From a service perspective, determining the cause of these crashes is time consuming and error prone.

3.9 Software Usability

Software usability can have a substantial impact on availability. The user interface provides the means by which a human interacts with the software. If the user interface is cumbersome, incomplete, inconsistent, confusing, obtuse, or needlessly complex, the user will make "mistakes" or misuse the software. Insufficient training or support can result in the user making mistakes that may not be obvious at the time but lead to eventual failure.

3.10 Software Documentation

Software Documentation impacts availability by how well the application or operating system is explained. Installation, configuration, tuning, user functionality, exception handling, and compatibility issues must be fully and clearly covered by the documentation. Many of the user and system manager mistakes are a result of poor documentation. Software also tends to change more frequently that hardware. Indeed new versions may be totally different from older versions and these differences must be documented. Software can also be documented on-line by means of tutorials, general help facilities, and context sensitive help capabilities. The availability and usability of software can be greatly enhanced if the user can get immediate help with an operation or process rather than having to search a library for information.

4 Service Delivery Impact on Availability

Service Delivery can have an enormous impact on the customer's perception of availability. Service Delivery begins by negotiating a contract and setting appropriate expectations for the service. Response time, after hours coverage, on-site coverage, problem support escalation, preventative maintenance, logistics, and training, in addition to the product MTTR, all can affect the time required to restore availability. It should be noted, however, that the impact of some aspects of Service Delivery can be minimized when appropriate Fault Avoidance, Fault Tolerance, and Fault Management techniques are employed. The topics discussed below are:

- Service Response Time
- Effectiveness of Maintenance
- Preventative Maintenance
- Scheduled Maintenance
- Service Contract Coverage
- Symptom Directed Diagnosis
- Problem Solution Escalation and Support
- Replacement Parts Logistics

4.1 Service Response Time

In addition to MTTR, the Service Response Time (SRT) may have the next greatest impact on availability. Depending on the locale of a customer the contracted on-site service response time may be from two hours to next day. In many cases, remote services provided through Digital Customer Support Centers may provide quick response for analysis of problem situations. Indeed, the CSC's are able to resolve most problems caused by operator mistake, inappropriate configurations, and other problems that do not require physical replacement of parts. The CSC's may be able to recommend "work around" solutions to keep the customer's application running until repair services can be provided. They may also assist in assuring that the correct parts are identified prior to dispatching service personnel to the site.

However, when physical repair is necessary, the service response time will add considerably to the total downtime and have a direct impact on availability. The customer's needs must be analyzed when proposing a service contract. A contract for 5x8 service for a customer running two shifts, 6 days a week may easily result on average with twice the downtime per hardware service call. If availability is a concern then it is imperative that the service contract be appropriately matched to the customer's operations.

Some customers have a requirement for the highest possibility availability. Typical applications might be defense systems, stock exchange, nuclear power plant control, real-time monitoring, reservations systems, OLTP, etc. Some of these have a

continuous, 24x7 requirement, while others have a 100% availability requirement but for only an 8-hour per day shift. As far as response time is concerned, very high availability operations may require on-site service personnel and local stocking of parts.

4.2 Effectiveness of Maintenance

The effectiveness of maintenance can have a considerable impact on availability. The time it takes to service hardware or software products and the quality of that service may determine the length of downtime experienced by the customer. Training on the installation, operation, and repair of hardware as well as an understanding of the controlling software affects the quality of maintenance. The Field Service organization must be supported by effective service delivery tools and product maintenance tools. Procedures must be developed to ensure the consistency and quality of repair services. Logistics, consultative support, and a process for escalation all contribute to the effectiveness of maintenance.

4.3 Preventative Maintenance

For some products, failures that impact availability can be lessened by a program of Preventative Maintenance and proactive Symptom Directed Diagnosis. Scheduled time slots negotiated with the customer allow for repairs, cleaning, adjusting, ECO/FCO's, upgrades, etc. to take place without interrupting applications. Hardware failures that can be repaired or prevented through this approach extend the apparent reliability of the hardware. Software upgrades, software maintenance releases or patch installation, disk defraging, cleaning up error logs, verifying status and operation of fault management tools, and collecting data for analyzing customer needs for additional services might also be done during this preventative maintenance time.

4.4 Scheduled Maintenance

In the 24x7 environment, fault tolerance and redundancy, failover capability, and on-line repair capability might be the best approach to continuous operations. In the 8-hour shift high availability environment, fault tolerance and fault avoidance with the goal of deferring maintenance to off-hours, and off-hours coverage would be required. Scheduled maintenance includes not only remedial maintenance deferred to an appointed time, but also planning for ECO/FCO activity, hardware and software upgrades, and value added services. Such services as capacity/performance planning, configuration management, security evaluation, storage management, system tuning, data backup and archiving, may need some access to computer resources. Scheduling these services minimizes or eliminates potential impact on availability.

4.5 Service Contract Coverage

The hours of service contract coverage may lead to a lower than expected availability. Failures that occur during off-contract hours could result in extensive downtime simply waiting for service to arrive. It is important to tailor the service contract to the customer's hours of business. Additional off-hours coverage may be needed for non-reactive maintenance such as ECO/FCO installation, software and hardware upgrades, preventative maintenance, backup and data archive services, etc. The important point to make here is that an evaluation of the customer application and operations requirements should be made and the customer expectations must be set when establishing the service and support contract coverage.

4.6 Symptom Directed Diagnosis

Symptom Direct Diagnosis is a methodology for capturing fault related data prior to or at the time of a system failure for the purpose of isolating the location of the fault without resorting to running a "test" diagnostic. This approach often provides the ability to discover and isolate a fault before a catastrophic failure occurs, allows time for reconfiguration to avoid the fault, and provides an opportunity to schedule repairs to avoid a business interruption. Symptom Directed Diagnosis requires a hardware error detection mechanism and a fault management strategy to be effective.

4.7 Problem Solution Escalation and Support

Service delivery can be enhanced to maximize availability by ensuring that a manageable plan for escalating problems and getting support is in place. The responsibility for such a plan rests with both the customer and the service supplier. The customer must keep the service supplier informed of the business requirements. The customer may have to provide technical assistance to assist the service provider in gaining as understanding of the nature of a problem. For instance, the customer must provide for error logging files and crash dumps to be collected and archived. Some problems will be difficult to resolve but a cooperative effort between the customer and service provider will make possible temporary work-arounds that can restore availability as the problem is being worked.

The service provider must have a process in place for recognizing the need for additional technical and managerial support in resolving a problem. The process must be automatic in the sense that escalation is a routine event that brings to bear all the necessary resources. Escalation must automatically carry forward the needed data to understand the nature of the problem without having to reinstitute investigations at each new level of escalation.

A cohesive escalation and support process may be best provided by one service provider than by many. Problems that are difficult to localize are then none-the-less the responsibility of a single service provider and are more difficult to pass off as "not my problem" just because the solution is not obvious.

The most important consideration is to evaluate the customer's needs up front and put a plan in place to address the technical and managerial escalation process through both the customer's and the service provider's organization before the onset of a major problem. If the process is defined and understood the availability can be restored with minimum impact to the customer's business.

4.8 Replacement Parts Logistics and Procedures

Replacement Parts Logistics can affect availability in several direct and indirect ways. Replacement parts must be accessible within a timeframe acceptable to restoring the customer's application. A customer with a highly critical application may need to have parts accessible locally or even stocked at the customer site. If critical parts cannot be stocked locally then provisions for backup equipment that can be used to swap equipment or for temporary use should be considered. A plan for restocking parts after consumption is required.

Parts handling can also impact availability by causing the introduction of new faults along with the new parts. DOA parts may fix one problem while introducing another. Replacement of several parts at one time (such as two or three modules) may totally confuse the situation by introducing multiple faults. Service personnel need to be trained in handling electrostatic sensitive devices, have access to necessary special equipment and tools, and use appropriate procedures to replace such devices. Improper use of solvents and cleaning agents can also inject intermittent contact connections. Proper routing of cables and cable restraints must be ensured to prevent inhibiting of air flow which can cause additional failures.

Serious intermittent failures can result from replacing parts with other than identical parts; partially installing an ECO/FCO by replacing only one of the necessary parts; installing a hardware ECO/FCO during a part replacement without also making supporting changes to the software or system parameters. When parts must be replaced with parts of a different revision great care must be taken to prevent introducing incompatibilities.

There is a corollary to hardware parts - software "parts" are the information required to fix software; known problem and solution descriptions, problem recovery procedures, patch installation instructions, revision dependency information, notes files, etc. Other more tangible "parts" might be such things as release kits, replacement manuals, license PAKS, etc. Knowledge of and access to these software "parts" may be critical in getting a customer back on the air again.

5 Operations Impact on Availability

The operations of a data center has a major impact on application availability, though is often overlooked when evaluating availability. There are many things that a customer should do in the allocation, preparation, and control of a suitable environment; in policies concerning security, backup, recovery, and archiving of data; in testing and installing new products; and in training and procedures for actually running the data center. Some aspects of data center operations that impact availability are:

- Experience of Operations Staff
- System Management
- Recovery Plan
- Security Policies
- Disaster Planning

5.1 Experience of Operations Staff

The data center Operations Staff can have an enormous impact on availability. The care and feeding of the computer systems relies on the human element. Improperly or carelessly used commands (such as DELETE *.*.*) can render the application non-functional and may require restoring or reloading of files. Improper termination of an application or operating system (via CTRL-P, HALT, REBOOT for instance) may leave data files and system files in a state from which they cannot be recovered.

Clearly operators must be extensively trained not only on how to operate applications and equipment, but also how to respond to exception conditions. An operator may discover that a removable disk media is getting errors and move the disk media to another disk drive. After a short time that disk drive begins to get errors so the operator "fixes" the problem by moving the disk media once again. The next time any of the disk drives are used they get more errors. The problem of course is a defective disk media. Unfortunately, by moving the disk media around the operator has managed to damage the heads on multiple disk drives. The damaged heads then scratch and damage additional disk media. An originally simple problem of a dirty disk media has now resulted in damaged heads and disk media throughout the data center. Even if the disk drives are repaired, if *all* of the damaged disk media is not replaced the problem will recur. An experienced operator should be aware of the impact his actions can have on the reliability of the hardware and valuable data in the data center.

5.2 System Management

The System Manager is responsible for establishing the proper software environment for applications and the operating system, implementing the security policy of the data center, and creating the necessary support system for controlling the data center. The system manager is more than an experienced operator. He must understand the impact

of altering the operating system parameters to ensure reliability and integrity. An improperly configured operating system may not meet the performance requirements of the applications or may crash intermittently. The system manager must also respond to failures in the system by analyzing likely causes and identifying when they can be resolved by a system management activity.

Data Center security must be managed by the system manager. Proper constraints on user account and system account passwords must be enforced. User account privileges, quotas, and access rights must be established to prevent users from taking actions that could impact the system integrity. Scripts must be written and used for such tasks as backups, queue management, and applications startup and shutdown. The actions of the system manager can result in significant downtime and loss of availability.

A real area of some concern is in an environment with distributed workstations. If the user, an engineer, secretary, business manager, who may have little understanding of system management or the workings of the operating system or workstation hardware, is left to maintain the installation of software, backup, account management, etc. of his workstation it is very likely that problems will be introduced. Put that workstation on a network or cluster and a potential exists for the user to interact and impact the reliability of the data center. A user/system manager may decide to upgrade to the latest version of some application and cause incompatibilities in a common database bringing down the whole business! In attempting to solve the problem, his actions may not be obvious leading to extensive downtime and possibly unnecessary repairs.

The message here is that system management is a speciality that requires extensive knowledge of hardware and software and the customer's business applications and should not be left to inexperienced users. DEC must work with the customer to ensure appropriate safeguards are in place to effectively manage the data center.

5.3 Recovery Plan

Failures will inevitably occur. If there is no recovery plan, then the amount of downtime will be greater than it need be and all traces of the failure may be eliminated making resolution of the problem impossible. To properly plan for recovery, an analysis of the customer's applications is required to determine how to recover from different types of failures. Can an application simply be restarted? Must a journal file be used to restore data? Must a restore of a prior backup be performed? Must the configuration of the system be changed prior to recovery? Whatever action is required must be analyzed to ensure that only the minimum that is really necessary is done. Operators must be fully aware of the procedures to recover.

A system manager or application consultant may need to prepare script files or a procedural checklist for the operators. Care must be taken that the recovery does not wipe out the data required for analysis of the failure. Crash dumps, error logs, terminal display failure messages, corrupted files, etc. must be captured for analysis. The lack of a plan may lead to an ad hoc attempt at recovery that ultimately causes additional failures.

5.4 Security Policies

Security Policies can have an impact on availability and the viability of a business. There are numerous means of getting around the security built-in to operating systems and data filing systems. Implementing the security policy generally is a responsibility of the system manager. The system manager must be aware of the elements (accounts, passwords, privileges, quotas, etc.) that must be managed to keep the data center secure. Operators must implement proper procedures for handling data to ensure integrity and security. Even a site where security is not a major concern must establish security policies to prevent operator and innocent user mistakes from innocently impacting the security of the data center. A customer tailored evaluation of security requirements should be undertaken if either security or availability is important to the customer.

5.5 Disaster Planning

Restoration of the customer's business is the prime objective of implementing a Disaster Plan. While not normally considered when trying to assess availability, disaster planning is a worst case scenario of a loss of availability. Restoration may not be possible in an acceptable period of time. Typical disasters might be fire, flood, or earthquake. Less cataclysmic disasters might be power outage, excessive hardware downtime or unreliability, or a heavy snow storm. For the business to remain viable it may be necessary to have a hot-site available in such disaster or emergency situations. If the site is local it may not be any more viable than the normal customer site in many of these situations. A physically distant site may be technically viable but lack a staff to operate it or have access to the current database. A plan is needed to keep the hot-site in a state of readiness. For predictable disaster situations such as a hurricane, it may be reasonable to operate the hot-site in parallel and switch over operations and staff as the storm nears. For unpredictable disasters, a standby operations staff may have to be called in to operate the hot-site.

These and many other business needs must be considered in planning for a disaster. It's just as important to consider the temporary short term interruptions as the devastating long term disasters if they can have major impact on the customer's business.

5.6 Data Management

An analysis should be performed to determine where data should be stored for best performance and availability. "Hot Files", files with the most frequent access, may be stored in Electronic Storage Units to increase performance. If the files are modified it may be desireable to have a background process move data to permanent storage. The disk system and file distribution should be analyzed to determine the optimum frequency of defragmentation and whether the defragmentation should be system-wide to specific to a particular application database. The data distributed among workstations may be backed up at different times or not be backed up at all. If it's necessary to recover the data center by rolling back to a specific time, it may not be possible to determine exactly what restores must be done. A policy for backup of all nodes should

be established to guarantee a rollback point for the whole data center. The number of backup sets required and the period for archiving the data should be considered. Off-site storage of archived data should also be considered. All of these data management issues can impact availability by inhibitiing or extending the time to recover after a failure.

5.7 Network Management

The configuration and management of a network can have an enormous impact on the performance and availability of the enterprise. Each node on the network must be properly configured or the entire network may suffer. This is a high risk area when system managers act independently without considering the network-wide impact of changes. In the event of a cluster crash recovery, if all client nodes attempt to boot from the server node simultaneously, the network will be swamped making booting a long, arduous task. Insufficient resources can create bottlenecks in the network that come and go impacting performance. Client/server computing is especially prone to failure at certain times of especially high activity, such as the start or end of the work day. Backups become a complicated issue when multiple network nodes must all be backed-up simultaneously. A server failure can render all clients down. On the other hand, a network may improve the availability of the enterprise by allowing for application failover to other nodes on the cluster. Software installation and upgrades are more complicated and must be carefully planned. However, in many cases, a cluster configuration is the most fault tolerant configuration for applications. The point here concerning availabiltiy is that the network requires more extensive planning and coordinated system management/network to deliver on the promise of higher availability.

6 Site Dependencies

The physical construction, layout, and environment can have an effect on the reliability and maintainability of the equipment. The site dependencies discussed in this section are:

- Environmental Factors
- Construction and Layout of the Data Center

6.1 Environmental Factors

Environmental factors greatly impact the reliability of hardware components. A computer room temperature of 80 degrees will cause at least a 10% reduction in reliability. Relative humidity above 60% will cause contact corrosion and condensation that can damage media. Low humidity can cause static build-up and paper handling problems. Local power or non-computer equipment sharing a common power distribution can introduce damaging voltage spikes that can cause premature device failures. Common mode noise caused by improperly grounded or improperly interconnected equipment can lead to ground loop injected noise and intermittent failures.

During the analysis of a data center it is essential to carefully consider the environment. The proper flow and distribution of cooling and humidification must be maintained. The installation of power generator sets or power conditioning equipment may be justified. Power line monitors may be permanently installed to measure supplied power. Such options as battery backup for system memory and Electronic Storage Units may improve system availability. Optical cable should be considered in all external or underground communications interconnects. The litany of environmental factors should be evaluated by performing an environmental survey prior to installation of a computer system. Whenever unexpected or unexplained failures recur an environmental assessment is warranted.

6.2 Construction and Layout of the Data Center

Construction and Layout of the Data Center can directly and indirectly affect the availability and operability of a data center. Equipment environmental specifications must be determined and followed. Equipment should be located and assembled so as not to interfere with service access. A simple repair may require extensive disassembly and downtime if equipment is placed too closely together to prevent access. Cable routing should be planned to allow cables to be located and be moved for reconfiguration or repair. Equipment should be logically positioned to not interfere with operations personnel access, such as for retrieving printouts, media handling, or visibility of device status indicators. Room must be allowed for hand carts to be moved around without bumping into equipment. Adequate storage areas, in some cases temperature and humidity controlled, must be provided. Secure areas may be needed for data archives.

7 The Availability Service

The motivation for writing this white paper is to establish a framework around which an availability service offering could be developed. We are in the initial phase of defining what the customer needs are, what the service requirements are, and what an availability service might consist of. Currently this service is envisioned as consisting of several basic tasks. The first task is collecting all the information necessary about the customer's data center, applications, and business requirements. The second task is to determine if current or proposed solutions (hardware, software, and services) meet those requirements, and if not what must be done to meet the requirements. The third task is to propose enhancements or changes to meet the requirements. The fourth task is to implement the enhancements or changes negotiated with the customer.

To support this service and address the second and third tasks identified above, we propose to develop an application to model availability and detect areas where changes in the configuration or operation of the data center or in the services provided by Digital would better meet the customer availability requirements. Supporting processes and procedures must be developed to accomplish the first and fourth tasks.

7.1 The Model for Availability

The tool for modeling availability is intended to assist a service consultant in understanding the potential availability of the current or proposed data processing solution and identifying opportunities to improve the solution. Some of the factors and variables discussed in this paper amplify the need for building in a plan for high availability in the engineering of both hardware and software products. Others are applicable to the CSSE organization in the design and development of products, tools, services, and training. Still others point to improvements in practices and policies implemented in the Field by customers and the service delivery organization. While many of these factors may be difficult or impossible to accurately model specifically, generalizations may serve to point out areas that are potential detractors to availability. The focus of the tool will be to identify those factors that can be affected and changes that can be made in the Field.

The model for availability is a tool for estimating the relative availability of a mix of hardware/software configuration, service delivery, and operations as they impact the customer's business application. The model can be used to compare the effects of changes that could be made to improve availability for the purpose of maximizing the availability of an established configuration or proposing a new configuration. It's important to understand that the model cannot determine a guaranteed "uptime". Rather it is intended to expose areas of the configuration, operations, environment, or service delivery where the greatest benefit could be achieved by making a change.

7.2 Strategy Employed by the Model

The strategy to be employed by the Availability Assessment Tool is to model a number of lifecycles of the configuration. A lifecycle is the period of time the configuration is likely to be viable in a business organization before becoming obsolete or wearing out. During each lifecycle each element of the configuration is modeled to determine when the next failure will occur and the duration of time until full availability is restored. The time of next failure and duration of downtime is determined by considering the basic reliability and repairability metrics using a normal distribution and factoring in the impact of service contract coverage, service delivery, fault management, and repairability, and the propagation of faults, recovery, operations, and site dependencies.

A simple prototype model has been developed to assess the feasibility of this modeling approach and is available for demonstration. This paper and comments of the reviewers will be used as reference to conduct a structured analysis of the availability model and better define the functionality of the tool.

8 Requirements Input Form

Now it's time for the reviewers to contribute to the development of an availability service and and an availability modeling tool. Your insights into availability and large data centers and experience in servicing customers could be of immense assistance in defining the content of a new service offering and in the functionality and usability of supporting tools. Use this form to submit your requirements or requests for features. The form is simply a guide to help you convey your requirement. Feel free to adjust or add any categories you believe would better present your requirements. All requirements submitted will receive a formal reply. The development group may wish to meet with you to discuss your requirements to ensure they are fully understood and appropriately implemented. This form is also available on-line. Contact me at PABDUL::ROCHELEAU if you wish an electronic copy via E-NET mail.

REQUIREMENT INPUT FORM

1. SUBMITTED BY:

Name:

DTN:

Node:

Loc/Mail Stop:

Dept:

Position:

Date of Submission:

2. ABSTRACT

Include a brief statement of the requirement. Is this a requirement for an availability service, an availability modeling tool, or both?

3. DESCRIPTION

Include a detailed description of the requirement and an indication of what you hope to achieve by the requirement.

4. SCHEDULE

Indicate any schedule conflicts with or dependencies on other Products or Services.

5. BENEFIT

Describe the benefit of adding this feature, including substantiating data. Is the benefit realized by Digital or the customer?

6. IMPACT OF NOT MEETING THIS REQUIREMENT

Describe the impact to Digital if this requirement is not accepted. Please explain this in terms of lost opportunities and markets. Describe the impact to the customer if this

requirement is not accepted. Please explain in terms of customer needs and customer satisfaction.

7. JUSTIFICATION

What is the best argument for implementing this requirement other than the obvious benefits stated above? Can you state this justification in terms of reduced cost, increased revenues, competitive advantage, or customer satisfaction?

8. RATING

Rate the importance of including the requirement using the following categories:

Mandatory -- The product or service cannot be considered a viable or credible offering without this feature. It will not be shipped unless this feature is implemented. The customer needs will not be met without implementing this requirement.

Desirable -- This is a desirable feature to have but its absence does not affect the credibility of the product or service. This feature should be included only if it does not delay the shipment after all mandatory requirements have been met.

9. KNOWN ISSUES

Include a statement of risks to either the schedule or the content and any conflicts with other development efforts. Are there other groups in Digital who need to be involved or aware of this requirement? Does this requirement have any known dependencies from other groups in Digital?

10. SUPPORTING DOCUMENTS

Identify any documents, standards, plans, etc. that add further detail concerning this requirement.

11. REVIEWS AND WALKTHROUGHS

During the development of an availability service and especially an availability modeling tool, the development group will conduct informal reviews and walkthroughs of the specifications, code implementation, and service definition. Outside experts and consultants are needed to participate and contribute. Each walkthrough will be limited to 1 hour and require about 1 hour of preparation. Will you consider participating in walkthroughs and reviews?

THIS DOCUMENT CONTAINS DIGITAL EQUIPMENT CORPORATION

CONFIDENTIAL INFORMATION

Digital's Worldwide Customer Satisfaction and Relationship Program

CUSTOMER SATISFACTION AND RELATIONSHIP SURVEY

1st Quarter Overall Results

PREPARED FOR DIGITAL EQUIPMENT CORPORATION
by
Development II

Digital's Worldwide Customer Satisfaction and Relationship Program

Development II
October 25, 1993

SURVEY BACKGROUND

WHICH ACCOUNTS WERE SURVEYED?

Surveyed but Not Completed
(Insufficient names for an account report)

Aetna
Anhauser-Busch
ASK
AT&T
✓Bankers Trust
✓Blockbuster
Cable and Wireless
✓Cerner
Chemical Bank
✓Disney
Fisons
Ford
Gillette
✓IDX
✓Merck
Mobil
Northern Telecom
Ralston
Royal Bank
✓SmithKline
Stentor
Thomas Cook
✓Upjohn
United Airlines
Wells Fargo
Weyerhaeuser
World Bank

Surveyed and Completed
(Sufficient names for an account report)

3M
✓Bell Canada
Burroughs
✓Caterpillar
Chrysler
Computervision
✓Dow
Glaxo
Honeywell
✓Monsanto
Pactel
✓Raytheon
Texaco
Toys R Us
VA Hospital
WR Grace
✓Westinghouse
✓Xerox

Refused to Participate (Against Company Policy)
(Indicated unwillingness to participate when contacted)

Glaxo (terminated after some success)
✓Johnson & Johnson
Mars M&M

Unable to Locate Digital Equipment

Arthur Anderson
Blue Cross

✓ Only the Account Managers, of the accounts with a "checkmark", provided names for the survey

DISTRIBUTION OF RESPONDENTS

BY COUNTRY

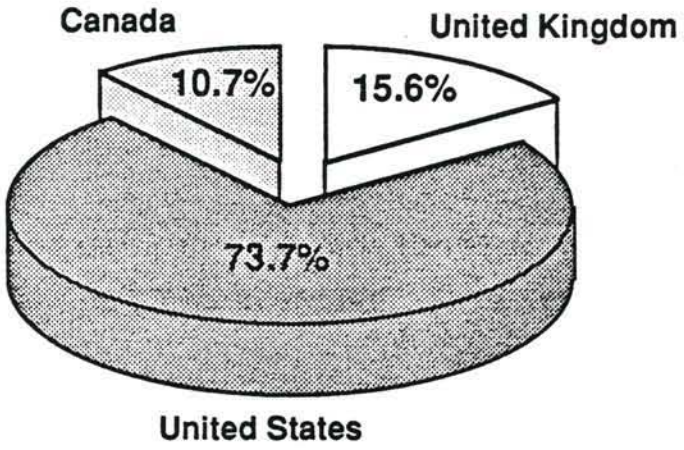

BY CATEGORY

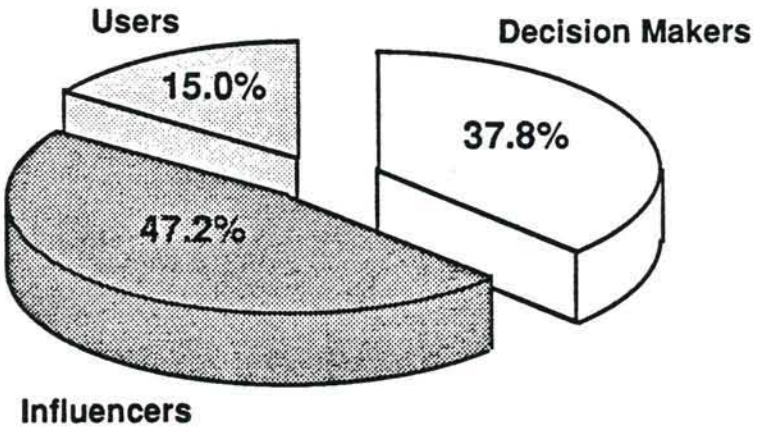

RESPONDENT CHARACTERISTICS

SIMILARITY OF RESPONSES

No **consistent statistical differences** between the Influencer's and Decision Maker's answers

AWARENESS OF GLOBAL ISSUES

Approximately **40%** answered Global equipment, services, purchases and maintenance questions

CUSTOMER PHONE COMMUNICATION WITH DIGITAL

87% of respondents answered question about contacting Digital by telephone

WRITTEN COMMENTS ON THE SURVEY

60% of the respondents provided written comments on the Survey

Note: 60% written comments is very high. We anticipated, based on experience, that only 10% - 20% would have comments

STATISTICS FOR FIRST QUARTER

	TARGET	ACTUAL
Surveys Out	602	610
Surveys Received	301	392 *
Percentage Return*	50%	Approximately 55% *
Non US Respondents	10%	26.2%

* 392 includes some responses received from surveys sent out after the original 610 were mailed.

CAVEATS

- **First Quarter Results Include only 392 Respondents**
  - Program Initiated August 2, 1993
  - Survey did not begin until September 5, 1993
  - First Quarter results concluded October 14, 1993
  - US, UK and Canada only
- **Too few UK and Canadian results to generalize**
  - By the end of the second quarter, Canada and Europe will have sufficient responses to generalize
- **The Results are probably indicative of the overall US results**

SUMMARY OF SURVEY DATA

Digital's Worldwide Customer Satisfaction and Relationship Program

Development II
October 25, 1993

HOW DO YOU "READ" A CUSTOMER SATISFACTION SURVEY?

The Survey Used a Four Point Scale

Totally Satisfied	- Very Strong Positive
Somewhat Satisfied	- Very Weak Positive
Somewhat Dissatisfied	- Very Weak Negative
Totally Dissatisfied	- Very Strong Negative

- A **Somewhat Satisfied** rating is neutral at best, often considered a negative
- **Totally Satisfied** is the "best" measurement of satisfaction
- 75% **Totally Satisfied** is the appropriate maximum for business to business markets
- **Fully Agree, Partially Agree, Partially Disagree** and **Fully Disagree** follow similar metrics

EXAMPLE WITH FIRST QUARTER SURVEY RESULT RANGES

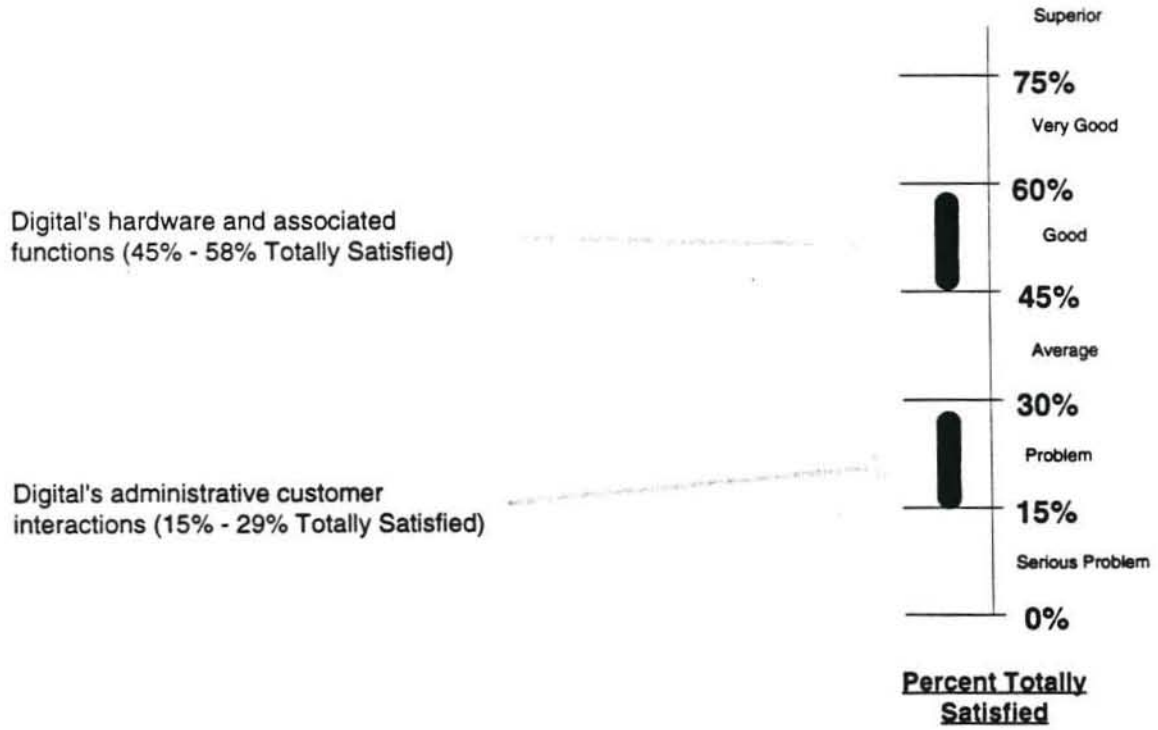

Digital's Worldwide Customer Satisfaction and Relationship Program

HARDWARE

SATISFACTION MEASUREMENTS

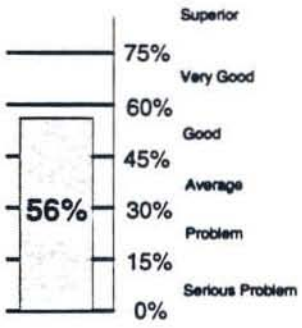

How satisfied are you with
the quality of Digital's hardware?

DETAILS

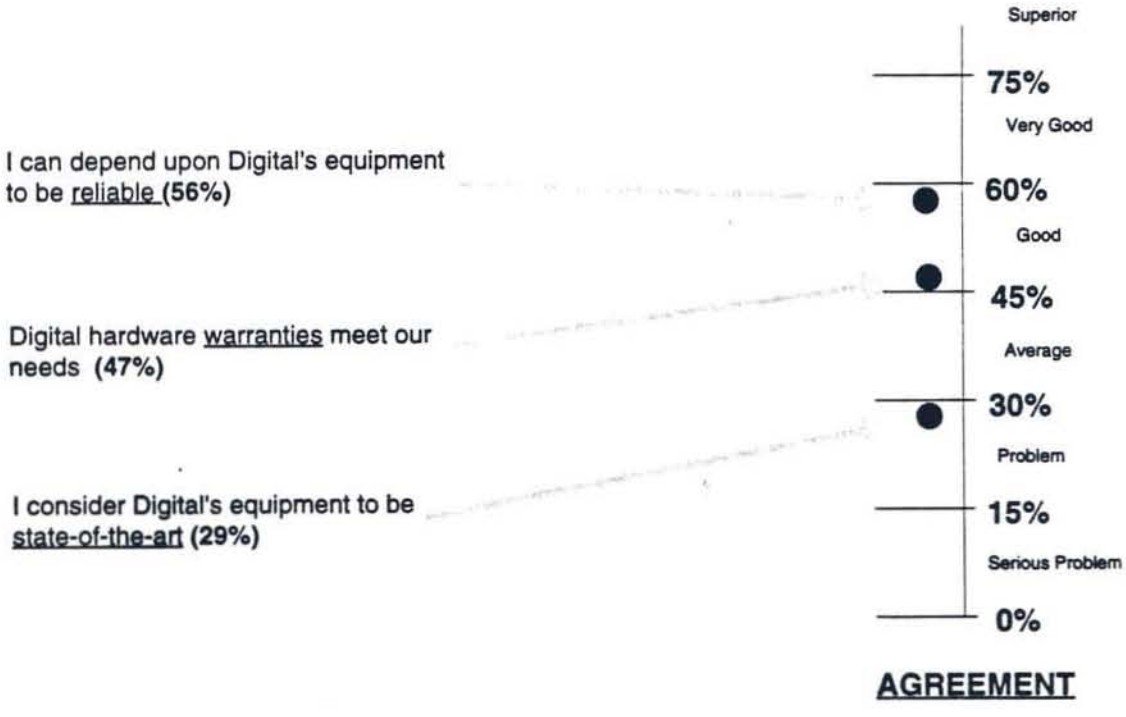

Digital's Worldwide Customer Satisfaction and Relationship Program

SOFTWARE

SATISFACTION MEASUREMENTS

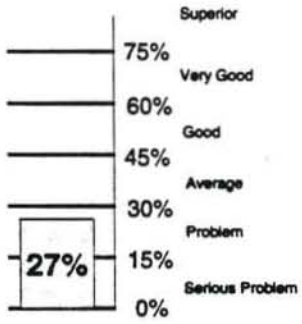

How satisfied are you with the quality of software supplied by Digital

DETAILS

Digital does a good job in adapting their solutions to the local languages in non-English speaking countries **(41%)** (caution, only 49 respondents)

Digital's software is reliable **(39%)**

Digital's software is easy to use **(20%)**

Licensing software from Digital is easy **(12%)**

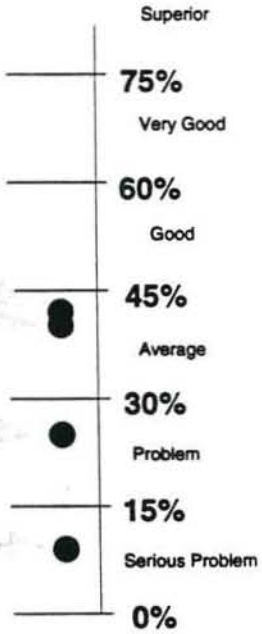

AGREEMENT

Digital's Worldwide Customer Satisfaction and Relationship Program

SERVICE

SATISFACTION MEASUREMENTS

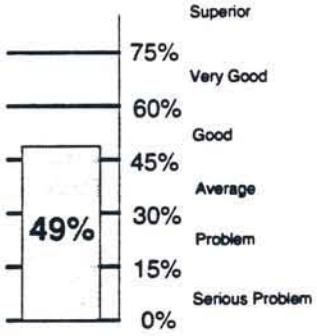

How satisfied are you with Digital's installation of equipment?

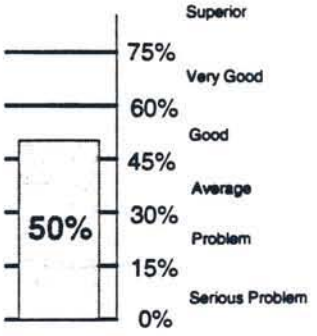

Based upon your most recent experience, how satisfied are you with Digitals Service Representatives?

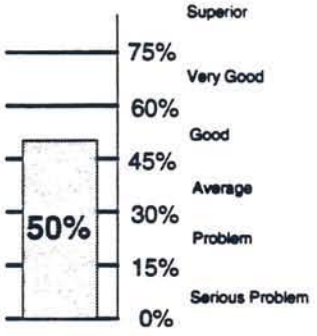

How satisfied are you with Digital's servicing of their equipment?

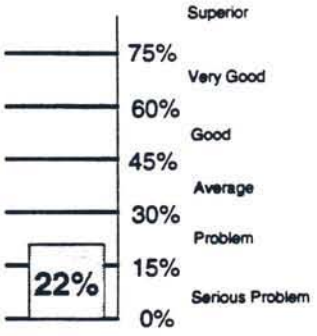

How satisfied are you with Digital's ability to service non-Digital equipment?

DETAILS

The Service Representatives respond effectively to our requests for service. **(44%)**

Digital's Services Management is available whenever I need to meet with them to resolve problems. **(41%)**

Maintenence for equipment and software is consistently available worldwide. **(39%)**

Digital is flexible in delivering products and services worldwide. **(19%)**

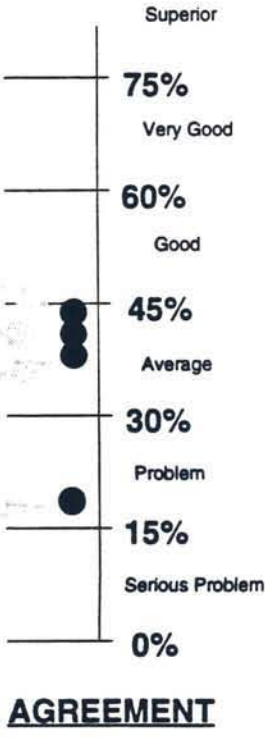

AGREEMENT

Digital's Worldwide Customer Satisfaction and Relationship Program

SALES

SATISFACTION MEASUREMENTS

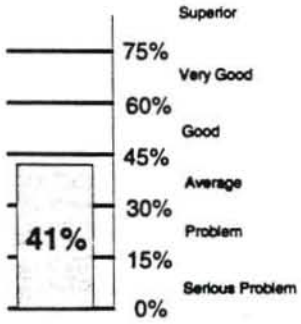

Based upon your most recent experience, how satisfied are you with Digital's Sales Representatives?

DETAILS

I have met the Digital Sales Account Manager assigned to our company **(64%)**

My Digital Sales Representative is responsive to my requests **(41%)**

The Digital Sales Representatives that I work with are knowledgeable about their products and serivces **(38%)**

The turnover of Digital personnel assigned to us is not too high **(32%)**

My Digital Account Manager knows how to resolve issues that arise in support of my Global needs **(28%)**

Digital Sales Representatives propose appropriate products and solutions to fit my business needs **(22%)**

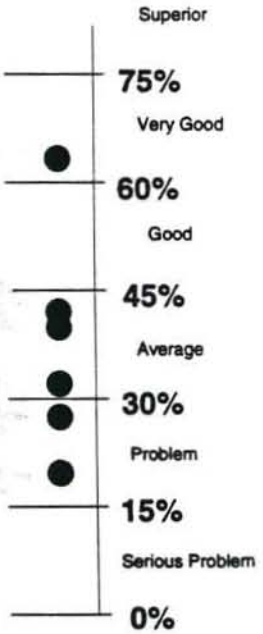

AGREEMENT

Digital's Worldwide Customer Satisfaction and Relationship Program

SUPPORT

SATISFACTION MEASUREMENTS

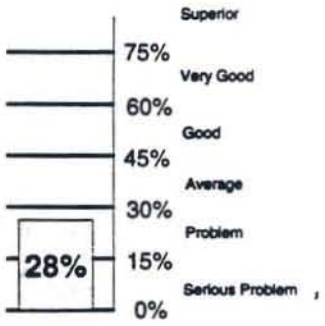

How satisfied are you with the quality of systems support you receive from Digital?

DETAILS

The systems personnel in Digital are knowledgeable about Digital's hardware and software **(46%)**

The documentation supplied by Digital is well written **(36%)**

The training supplied by Digital is effective **(32%)**

Digital understands our network needs and operation **(23%)**

Digital's systems consultants are effective in defining and developing systems for our business needs **(19%)**

Digital is effective at integrating our multi-vendor systems **(19%)**

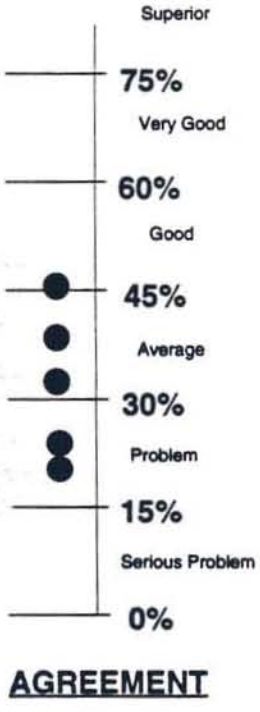

Digital's Worldwide Customer Satisfaction and Relationship Program

SATISFACTION MEASUREMENTS

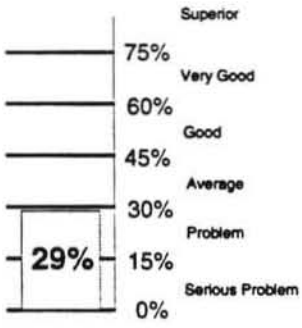

How satisfied are you with Digital's delivery of equipment?

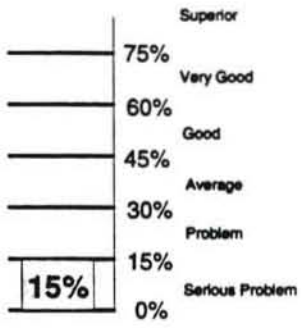

How satisfied are you with the ability of Digital's administrative processes to meet your company's administrative needs?

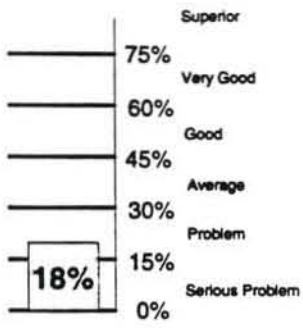

How satisfied are you with Digital's process of billing and invoicing?

DETAILS

Digital is flexible in how it responds to accepting and fulfilling orders **(24%)**

Global orders and invoicing of purchases from Digital **(24%)**

The time it takes from placing an order to the actual delivery of equipment **(17%)**

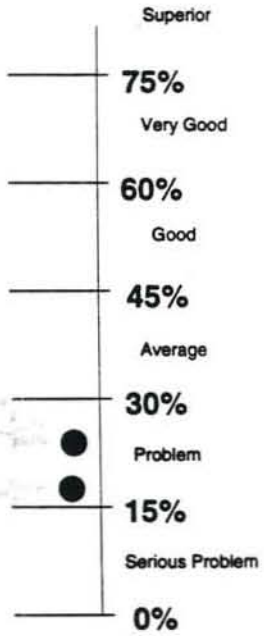

AGREEMENT

Digital's Worldwide Customer Satisfaction and Relationship Program

MARKETING
(Customer Value)

SATISFACTION MEASUREMENTS

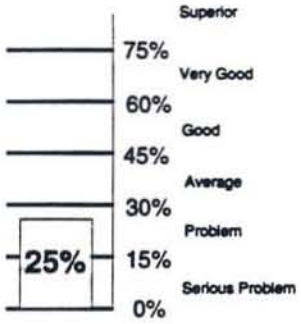

How satisfied are you with the value that your company receives from its investment in Digital's products and services?

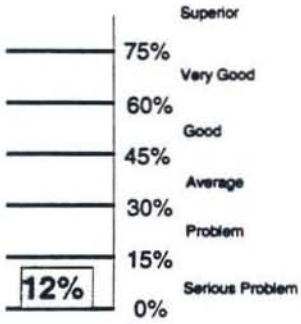

How satisfied are you with Digital's ability to provide solutions to enhance your company's competitiveness?

DETAILS

Solutions, using Digital equipment and services, are uniformly available worldwide **(35%)**

Digital effectively supports information industry standards **(30%)**

Digital communicates timely and thorough information about their products and services **(21%)**

Digital has a clear understanding of our business and industry needs **(18%)**

Digital leads the industry in open client server technology **(9%)**

Digital is effective at providing me with third party solutions **(9%)**

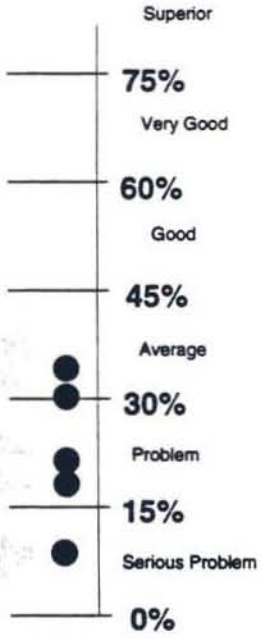

AGREEMENT

Digital's Worldwide Customer Satisfaction and Relationship Program

RELATIONSHIPS

SATISFACTION MEASUREMENTS

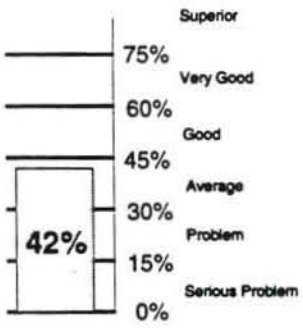

Based upon your most recent experience, how satisfied are you with your relationship with Digital's personnel?

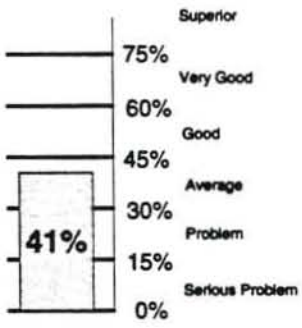

How satisfied are you with the level of commitment and enthusiasm displayed by Digital's personnel?

DETAILS

In the past, I have had very good relationships with Digital's personnel **(62%)**

Digital's actions demonstrate that they are interested in a long term relationship with us **(42%)**

I trust Digital as a business partner **(42%)**

I can count on Digital to keep their promises and commitments **(23%)**

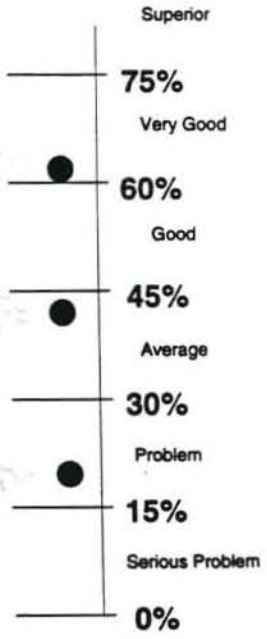

AGREEMENT

Digital's Worldwide Customer Satisfaction and Relationship Program

EASE OF DOING BUSINESS

OVERALL MEASUREMENT

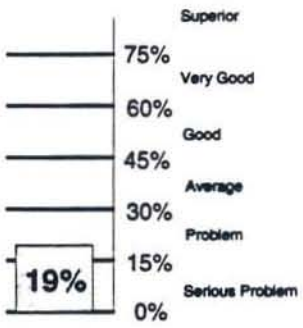

Digital is an easy company to do business with

DETAILS

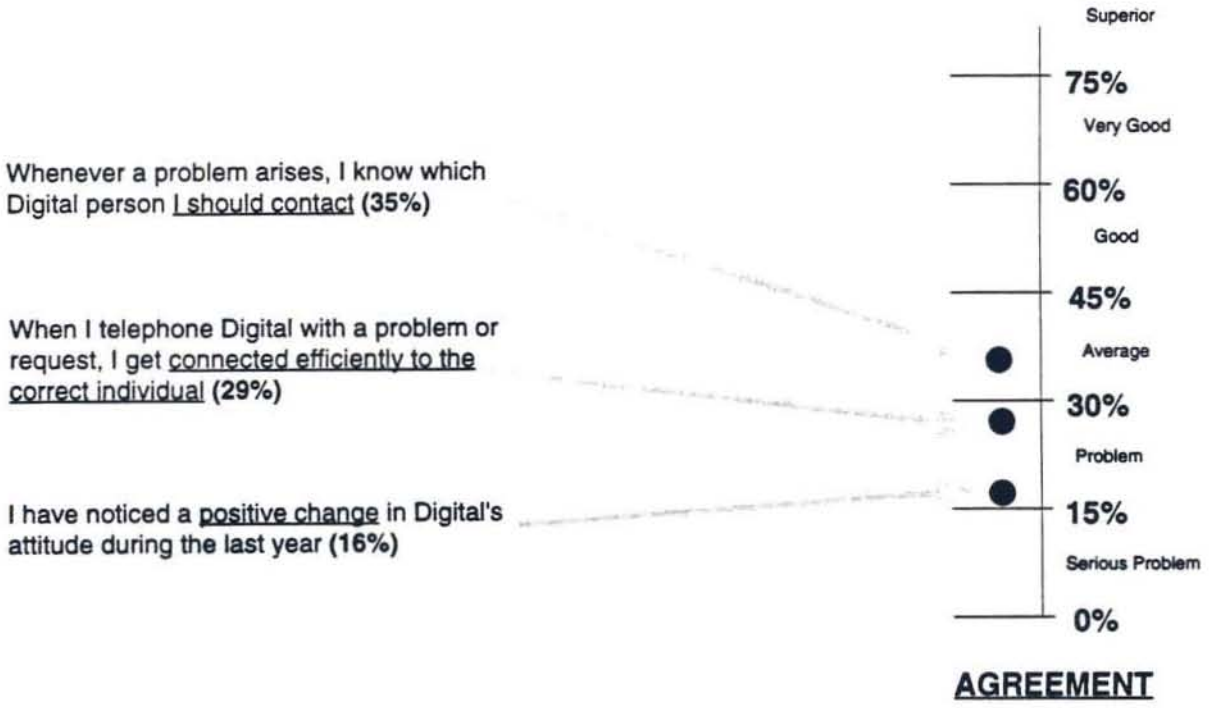

Digital's Worldwide Customer Satisfaction and Relationship Program

SUMMARY OF SATISFACTION CATEGORIES

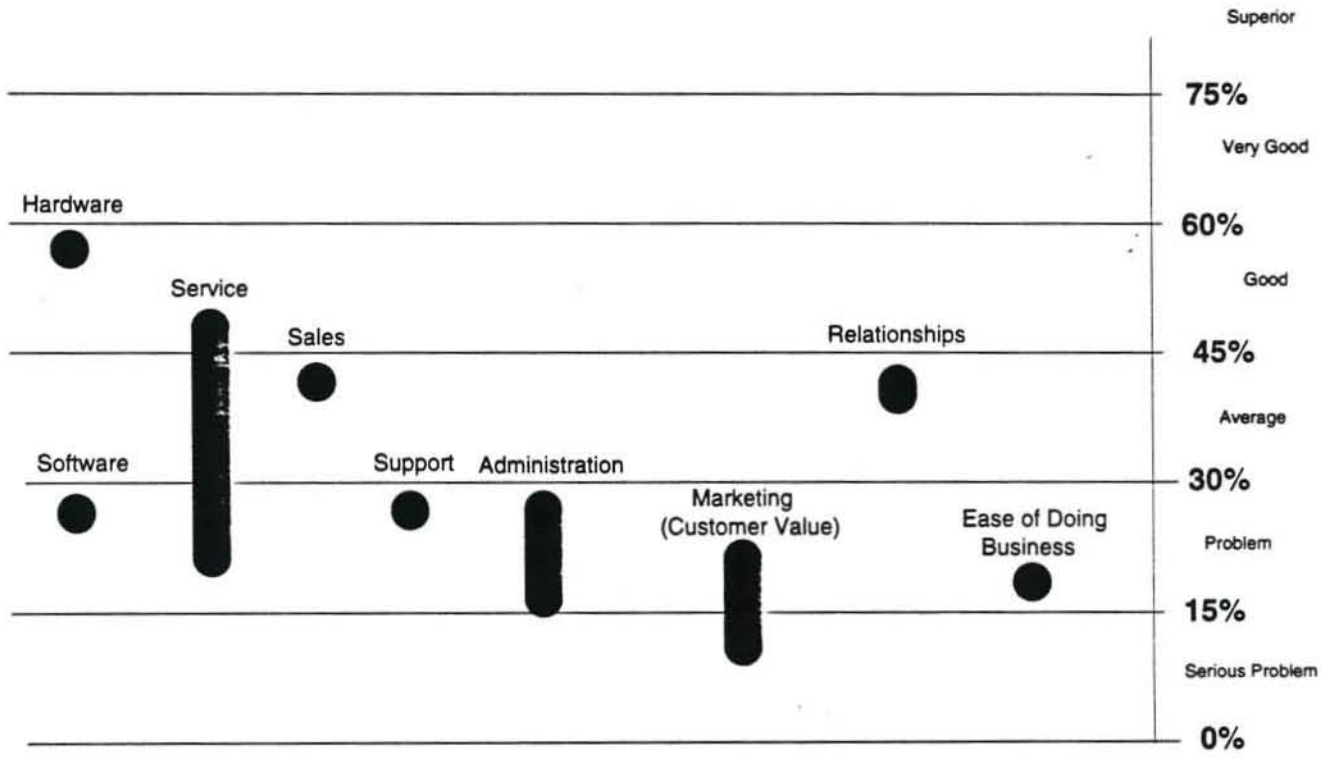

OVERALL SATISFACTION WITH DIGITAL

KEY CUSTOMER SATISFACTION MEASUREMENT

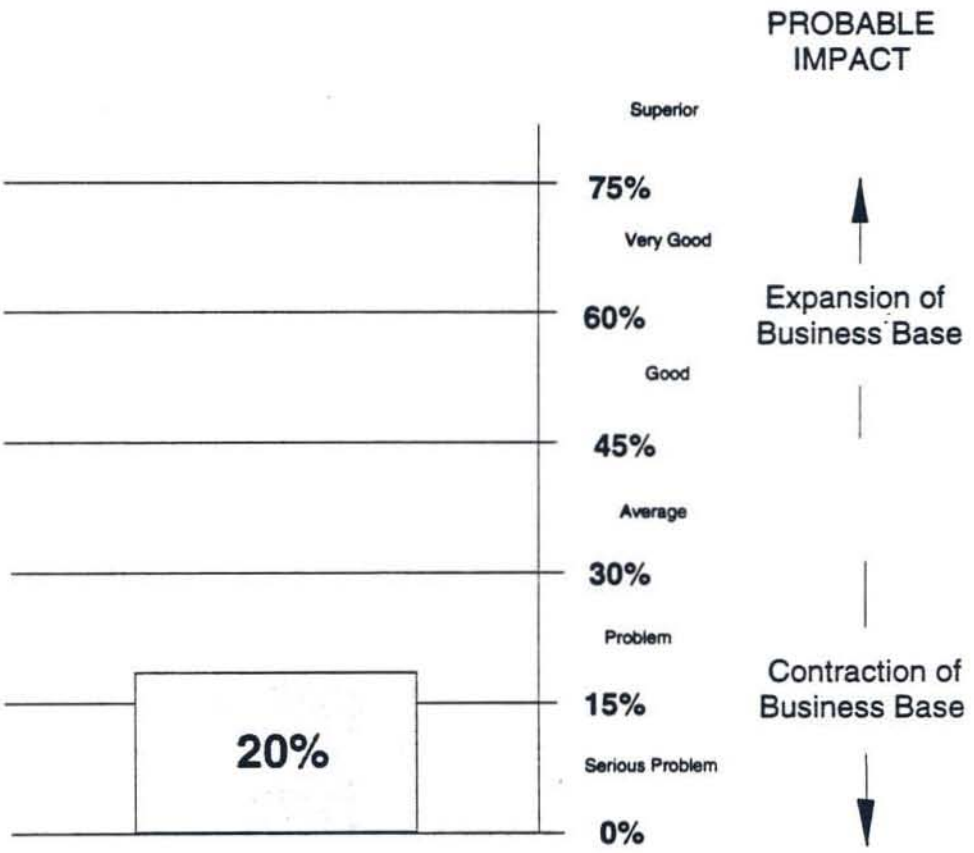

Based upon your recent experience, how satisfied are you with Digital overall?

OPPORTUNITIES FOR REQUESTED CUSTOMER INTERACTION

	Definitely	Probably	Total
Would like to have regular meetings with Digital to help facilitate future planning	**24%** 30%	**29%** 30%	**53%** 60%
Would like to meet with Digital's Senior management more frequently	**12%** 14%	**23%** 24%	**35%** 38%

Bold - All Respondents
Plain - Decision Makers

BASIC SUPPORT FOR DIGITAL

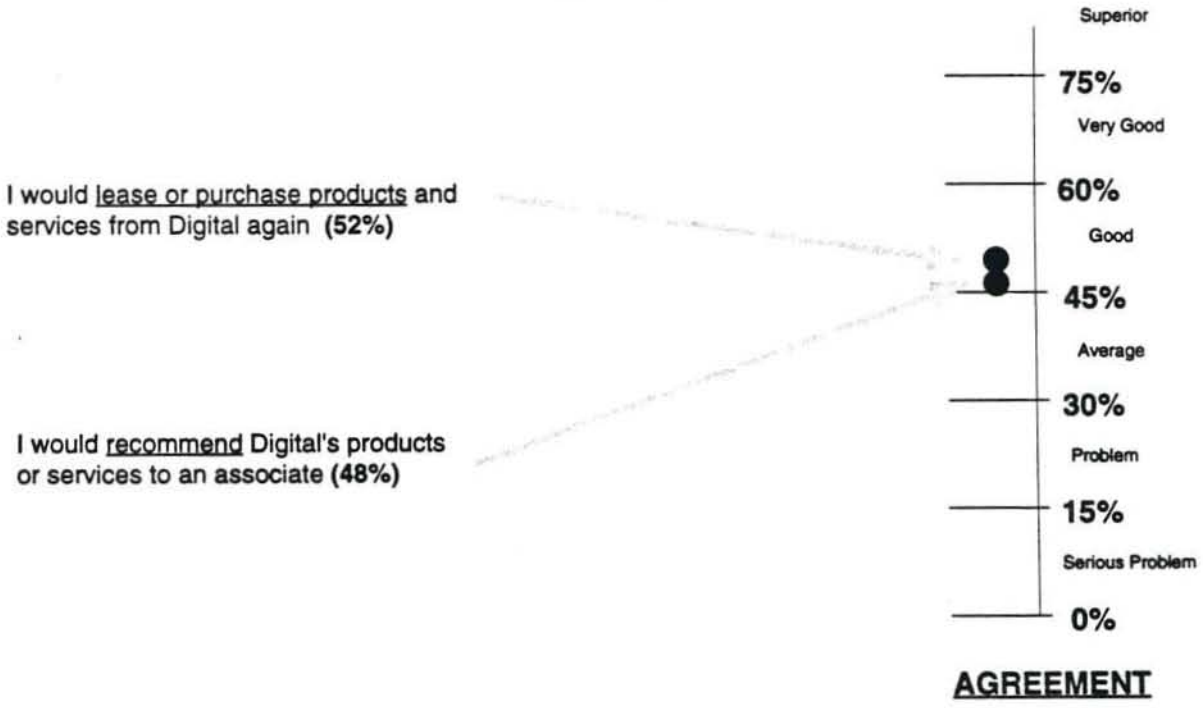

ISSUES

ISSUE: FUTURE POTENTIAL FOR DIGITAL HARDWARE

Hardware has been Digital's traditional base of strength. The customers however, are indicating perceptions that could create concerns in the near term.

STATEMENTS THAT IMPACT PERCEPTIONS OF HARDWARE

(Software, and adherence to industry standards will definitely affect the customer's perception of Digital's hardware)

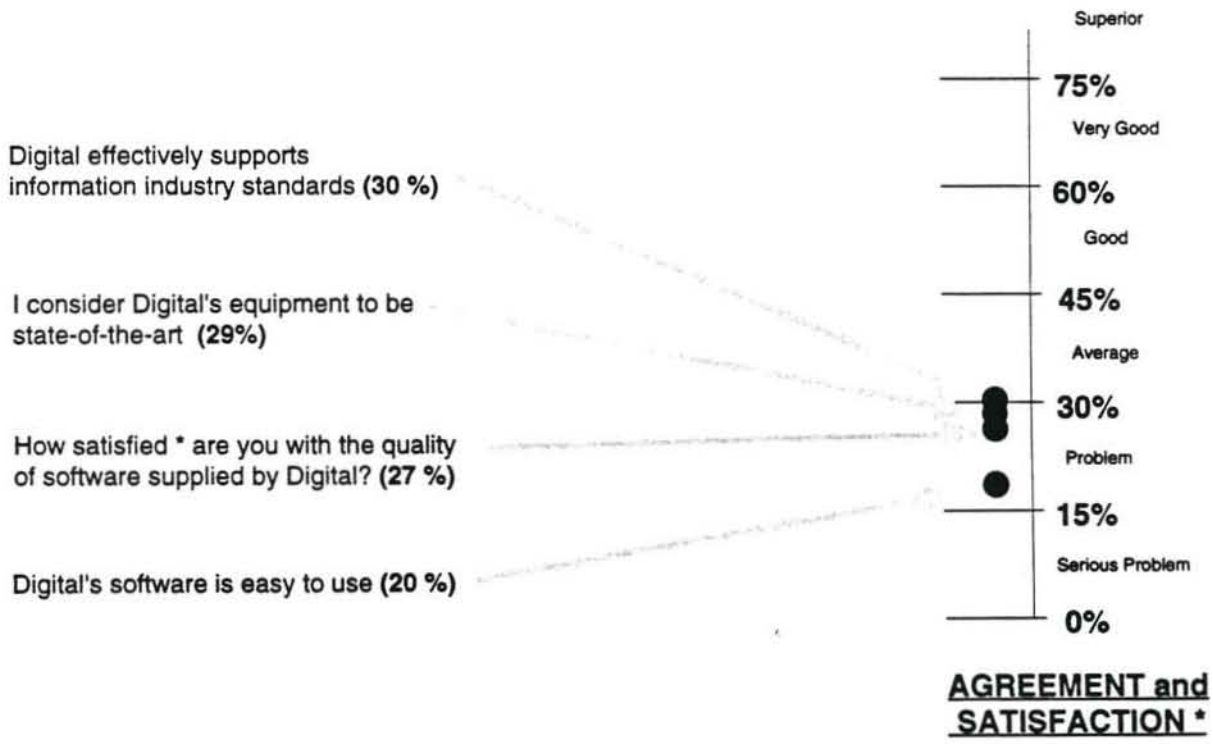

ISSUE: PRESENT STATUS OF MARKETING INITIATIVES

Digital is positioning itself as the "new" Digital focusing on the customer, the customers business needs, open client servers, a renewed enthusiasm and commitment and other customer related initiatives.

STATEMENTS THAT REFLECT MARKETING INITIATIVES

(These should be considered a status or a stake in the ground rather than an absolute evaluation)

How satisfied * are you with the level of commitment and enthusiasm displayed by Digital's people? **(41 %)**

How satisfied * are you with the value that your company receives from its investment in Digital's products and services? **(25 %)**

I can count on Digital to keep their promises and commitments **(23 %)**

I have noticed a positive change in Digital's attitude during the last year **(16 %)**

How satisfied * are you with Digital's ability to provide solutions to enhance your company's competitiveness? **(12 %)**

Digital leads the industry in open client server technology **(9 %)**

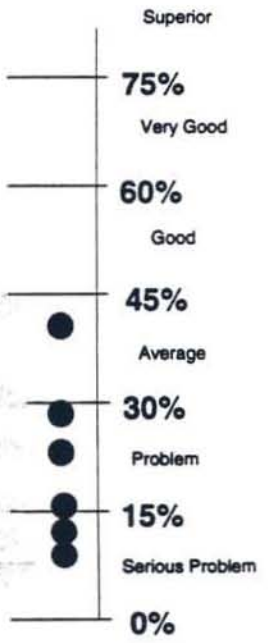

AGREEMENT and SATISFACTION *

ISSUE: UNDERSTANDING THE CUSTOMER

Understanding the customer's business, industry and technical needs, problems and requirements is critical for Digital's future success

STATEMENTS THAT REFLECT THE CUSTOMER'S PERCEPTION OF DIGITAL'S UNDERSTANDING

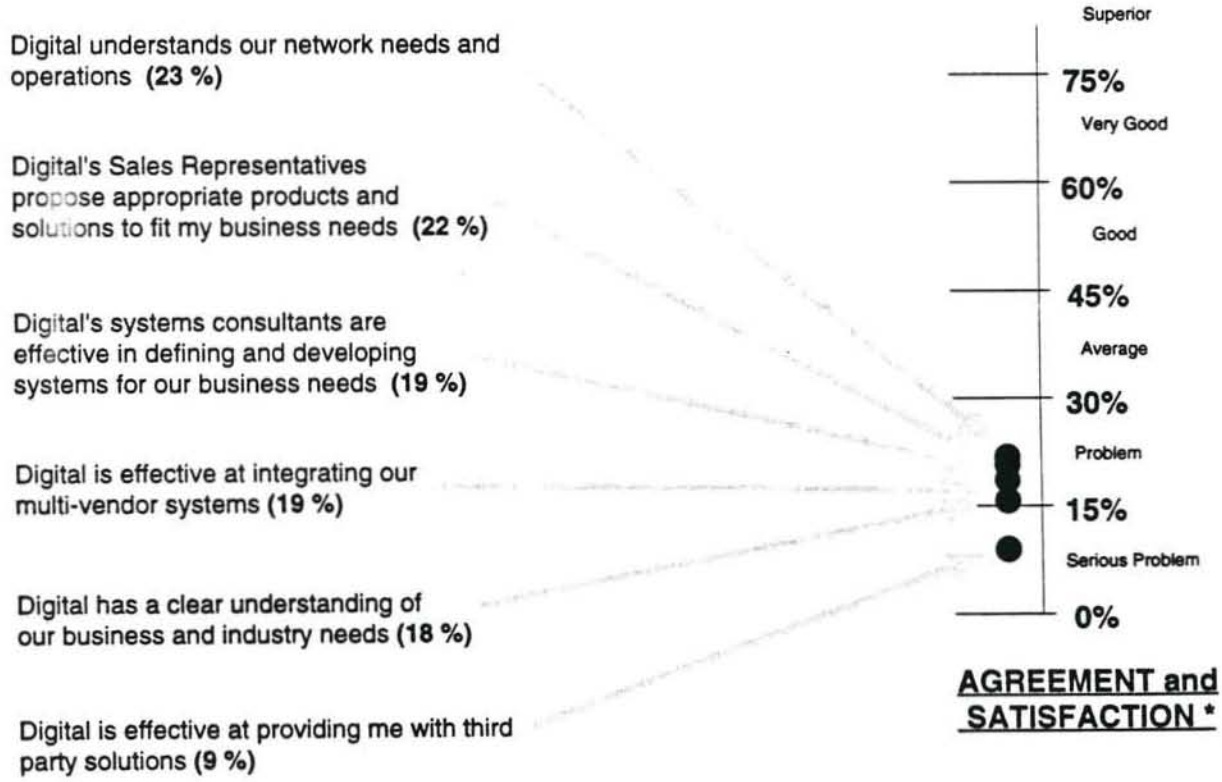

Digital's Worldwide Customer Satisfaction and Relationship Program

ISSUE: EROSION IN RELATIONSHIPS

62% of the respondents indicated that they have had very good relationships with Digital's personnel in the past. At present, only 42% are fully satisfied with the relationship. A decrease in the satisfaction level of the relationship between the customer and Digital has a major impact on Overall Satisfaction and needs to be reversed.

THE CUSTOMER'S VIEW OF THEIR RELATIONSHIP WITH DIGITAL'S PERSONNEL

In the past, I have had very good relationships with Digital's personnel **(62 %)**

Based upon your most recent experience, how satisfied are you with your relationship with Digital's personnel **(42%)**

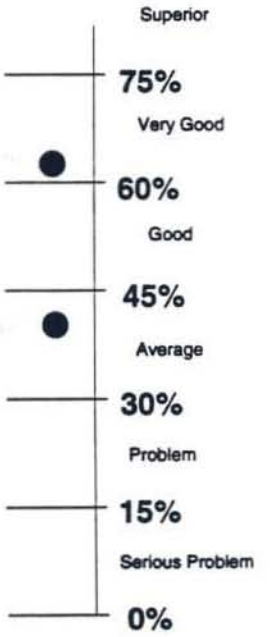

AGREEMENT and SATISFACTION *

NOTE: The wording of the two statements makes direct comparison difficult but by examining dissatisfaction the issue is clarified.

- **8%** indicated that they did not have very good relationships in the past
- **19%** indicate that they are dissatisfied with their present relationship with Digital

ANALYSIS

ANALYSIS OF RESULTS

TARGET MARKET: Respondents who were Somewhat Satisfied Overall

Used Analysis Technique to Identify the Areas to Explore

RESULTS

Impact Areas (top nine)	Relative Effect on Overall Satisfaction (Percent Increase in Overall Satisfaction resulting from a 100% change in Impact Area)
Relationships with Personnel	
Knowing who can solve problems	**32%**
Ease of doing business	**17%**
Overall relationship with Digital personnel\	**10%**
Equpment Related	
Hardware Quality	**30%**
Hardware is State-of-the-art	**10%**
Software Quality	**18%**
Solving Customer's Problems	
Value of Investment in Digital	**26%**
Integrating multi-vendor systems	**18%**
Sales representatives propose appropriate solutions	**12%**

IMPACT OF MEETING ACCOUNT MANAGER

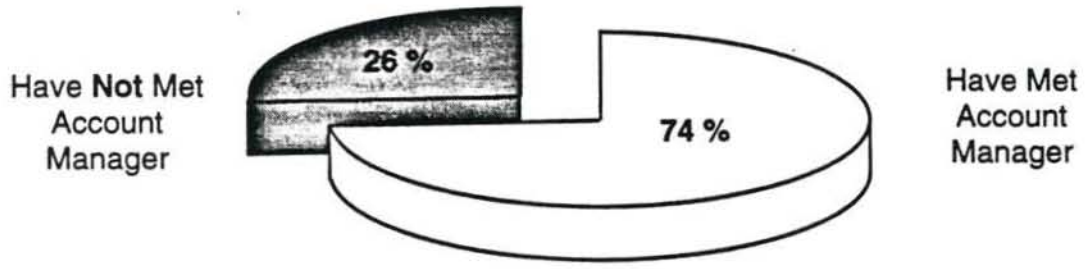

- **78 %** of **Decision Makers** have met the Account Manager
- **77 %** of **Key Influencers** have met the Account Manager
- **51 %** of **Users** have met the Account Manager

	Have Met the Account Manager	Have Not Met the Account Manager
OVERALL SATISFACTION		
Satisfied Overall	**23 %**	**14 %**
EQUIPMENT RELATED		
Hardware Quality	55 %	57 %
Software Quality	28 %	23 %
Servicing Digital's Equipment	48 %	55 %
Servicing non-Digital Equipment	20 %	[31 %]
ORDERS/DELIVERY/INSTALLATION		
Installation of Equipment	50 %	43 %
Equipment Delivery	28 %	32 %
Billing and Invoicing	19 %	[15 %]
Administrative Processes	17 %	[9 %]
DIGITAL'S PERSONNEL IMPACT		
Relationship with Digital Personnel	43 %	37 %
Digital's Commitment and Enthusiasm	44 %	32 %
Sales Representatives	47 %	13 %
Service Representatives	49 %	53 %
Systems Support Representatives	27 %	32 %
BUSINESS VALUE OF DIGITAL		
Value of Investment in Digital	26 %	23 %
Solutions Increase Competitiveness	13 %	6 %

[XX.X%] Outlined Percentages have less than 60 Returns

Percent Fully Satisfied for each category

IMPACT OF CUSTOMER PERCEPTION OF DIGITAL'S HARDWARE

Fully Agree
That Digital's hardware is state-of-the-art

Somewhat Agree
That Digital's hardware is state-of-the-art

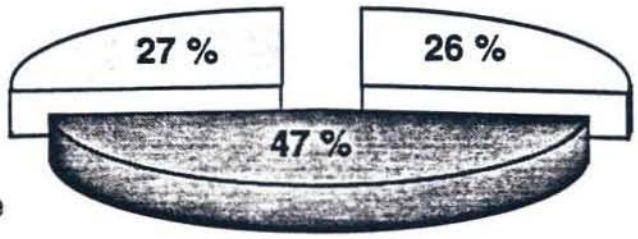

Do Not Agree
That Digital's hardware is state-of-the-art

	Fully Agree that Digital's hardware is state-of-the-art	Somewhat Agree that Digital's hardware is state-of-the-art	Do Not Agree that Digital's hardware is state-of-the-art
OVERALL SATISFACTION			
Satisfied Overall	43 %	17 %	3 %
EQUIPMENT RELATED			
Hardware Quality	80 %	54 %	31 %
Software Quality	51 %	18 %	15 %
Servicing Digital's Equipment	64 %	50 %	29 %
Servicing non-Digital Equipment	[26 %]	[24 %]	[6 %]
ORDERS/DELIVERY/INSTALLATION			
Installation of Equipment	69 %	48 %	24 %
Equipment Delivery	40 %	25 %	22 %
Billing and Invoicing	30 %	16 %	[7 %]
Administrative Processes	28 %	11 %	6 %
DIGITAL'S PERSONNEL IMPACT			
Relationship with Digital Personnel	59 %	43 %	22 %
Digital's Commitment and Enthusiasm	52 %	44 %	23 %
Sales Representatives	48 %	44 %	26 %
Service Representatives	64 %	53 %	29 %
Systems Support Representatives	47 %	24 %	11 %
BUSINESS VALUE OF DIGITAL			
Value of Investment in Digital	50 %	18 %	3 %
Solutions Increase Competitiveness	29 %	6 %	1 %

[XX.X%] Outlined Percentages have less than 60 Returns

Percent Fully Satisfied for each category

Digital's Worldwide Customer Satisfaction and Relationship Program

IMPACT OF RELATIONSHIP WITH DIGITAL

Not Satisfied
With the Relationship
with Digital personnel

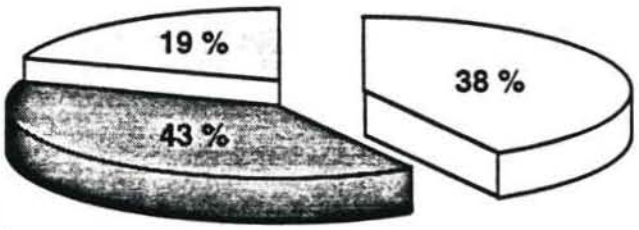

Somewhat Satisfied
With the Relationship
with Digital personnel

Fully Satisfied
With the Relationship
with Digital personnel

	Fully Satisfied with the Relationship	Somewhat Satisfied with the Relationship	Not Satisfied with the Relationship
OVERALL SATISFACTION			
Satisfied Overall	47 %	5 %	0 %
EQUIPMENT RELATED			
Hardware Quality	70 %	49 %	39 %
Software Quality	46 %	19 %	4 %
Servicing Digital's Equipment	71 %	38 %	29 %
Servicing non-Digital Equipment	36 %	12 %	5 %
ORDERS/DELIVERY/INSTALLATION			
Installation of Equipment	72 %	35 %	27 %
Equipment Delivery	41 %	22 %	17 %
Billing and Invoicing	26 %	13 %	11 %
Administrative Processes	30 %	7 %	2 %
DIGITAL'S PERSONNEL IMPACT			
Relationship with Digital Personnel	NA	NA	NA
Digital's Commitment and Enthusiasm	72 %	26 %	3 %
Sales Representatives	81 %	19 %	3 %
Service Representatives	70 %	43 %	18 %
Systems Support Representatives	50 %	13 %	9 %
BUSINESS VALUE OF DIGITAL			
Value of Investment in Digital	45 %	15 %	3 %
Solutions Increase Competitiveness	24 %	5 %	0 %

[XX.X%] Outlined Percentages have less than 60 Returns

Percent Fully Satisfied for each category

Digital's Worldwide Customer Satisfaction and Relationship Program

RECOMMENDATIONS

FIRST RECOMMENDATION

Focus on Increasing the Customer's level of Satisfaction with their Relationships with Digital Personnel

Issue : **Enhancing the relationship between Digital and the customer is the single most effective strategy for improving the Overall Satisfaction in a relatively short time**

Improving Relationships must involve all functions in the company

A focus towards solving customer's problems is the key

SPECIFICALLY

1. Establish a Policy that directs Digital's efforts towards solving customer problems at all levels of the company.

Key Leverage : Solve the Customer's Problems

Relationships

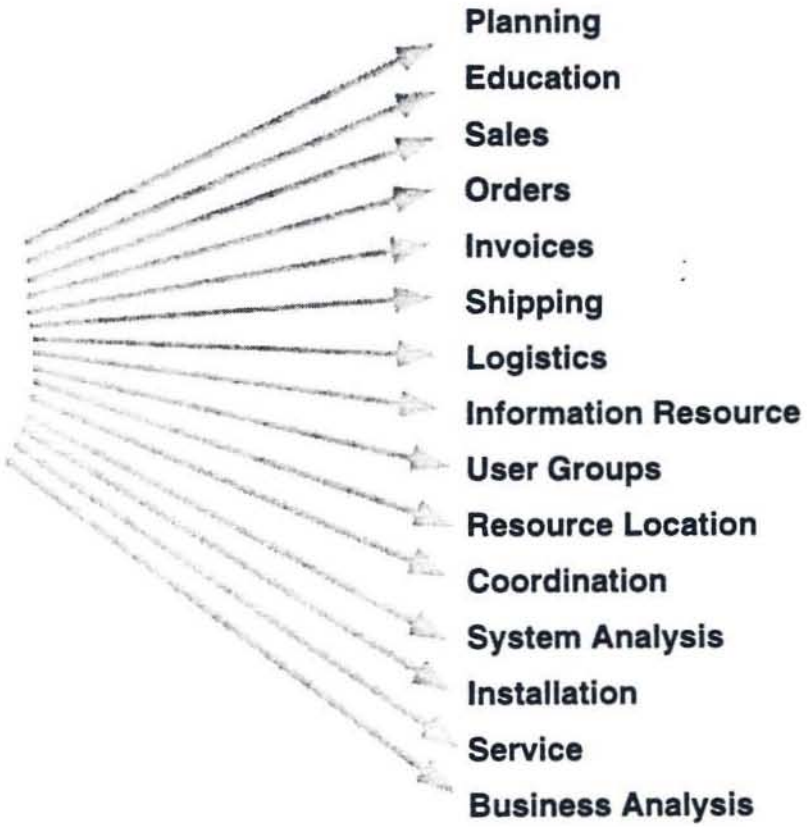

The definition of Relationships goes well beyond sales to include every function in the company that can solve the customers problems and simplify the complexities surrounding their information systems.

CUSTOMER COMMENT SYNOPSIS

60% of All Respondents Provided Comments

SELECTED COMMENTS GROUPED BY FREQUENCE OF RESPONSE

SALES (30 Responses)

- Better responsiveness from sales rep
- More competent/knowledgeable sales reps
- More frequent contact

SALES (13 Responses)

- Continue to improve sales staff
- Sales rep and technical rep change too frequently

SALES (10 Responses)

- Periodic contact by Digital Sales Rep with new ideas, products, programs etc.
- Future strategic planning help, on-site and company-wide also

BUSINESS PRACTICES (31 Responses)

- Convince me of the value of DEC products
- Sell to customer needs

BUSINESS PRACTICES (15 Responses)

- Focus on customer, meet my needs
- Stop reorganizing, maintain consistency of interface
- Sell what I need vs selling DEC products

MARKETING (15 Responses)

- Provide more information on products and technology
- More visible presence in the marketplace
- Be a leader in new technology

MARKETING (8 Responses)

- Increased visibility in PC market, thus client server
- Keep me informed of new products
- Products to customer more timely and cost effective

Comments were virtually all constructive, very little "slamming"

APPROXIMATELY ONE OUT OF EVERY FIVE CUSTOMERS THAT WE TALK TO REFUSE TO PARTICIPATE IN THE SURVEY

What Reasons Do They Give?

- They take surveys and nothing happens
- They write letters and nothing happens
- They take surveys and Digital people that they know get "nailed"
- Digital has too many surveys

THE SECRET

Every day
At every level
In every function

Solve the Customer's Problems

Digital's Worldwide Customer Satisfaction and Relationship Program

Re-establish Customer's Perception of Digital as a Provider of State-of-the-Art Equipment

Issue : **The majority customers do not perceive Digital as providing State-of-the-Art equipment. This definitely threatens the near term business**

SPECIFICALLY

1. **Determine accurately how the customer defines "State-of-the-Art"**
   - Interview selected customers (from survey base) to determine their viewpoints
   - Resurvey (a phone survey is sufficient) to prioritize these viewpoints
2. **Identify Digital's existing and emerging technologies/capabilities that coincide with the customers prioritized viewpoints**
3. **Develop a positioning strategy to highlight those areas**
4. **Develop a communication strategy (internally and externally) to distribute the positioning**

THIRD RECOMMENDATION

Establish a Formal Process to "Follow-up" on the Results of the Survey on an Account Level

Issue : The customers have expressed their problems, issues and requirements. They want interaction with Digital. This is an ideal opportunity to enhance the Relationship with an account by discussing the results and reacting to the customer's problems.

Europe has recently established a pilot project to insure that account feedback occurs. (Key part of the TQM process)

Canada has used a similar process to significantly increase sales. It would be worthwhile to review that process and its success.

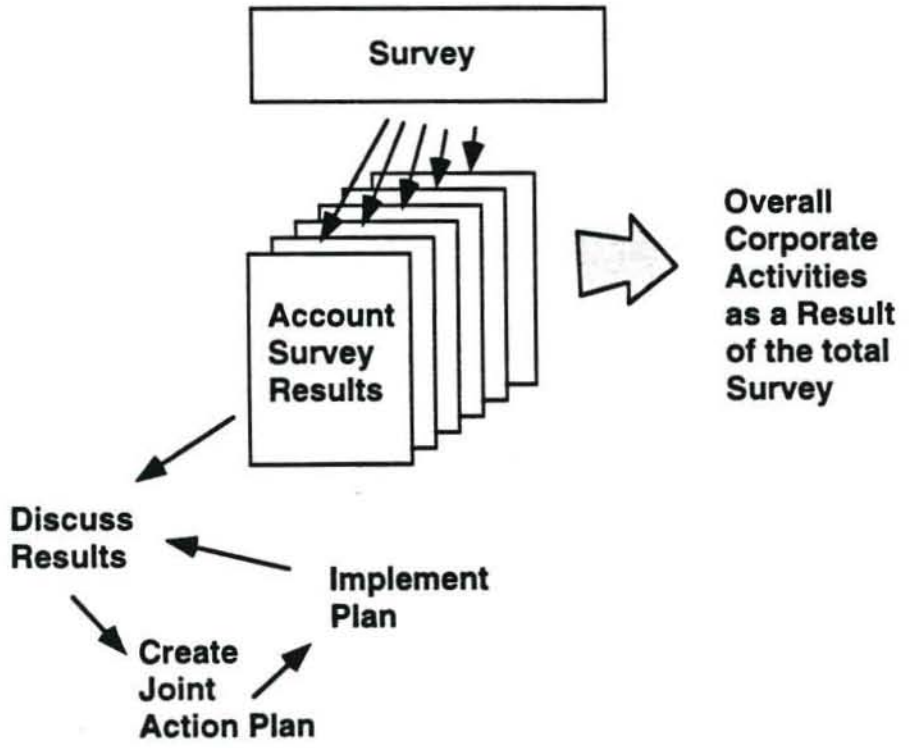

HOW EUROPE IS IMPLEMENTING THE FEEDBACK PROCESS

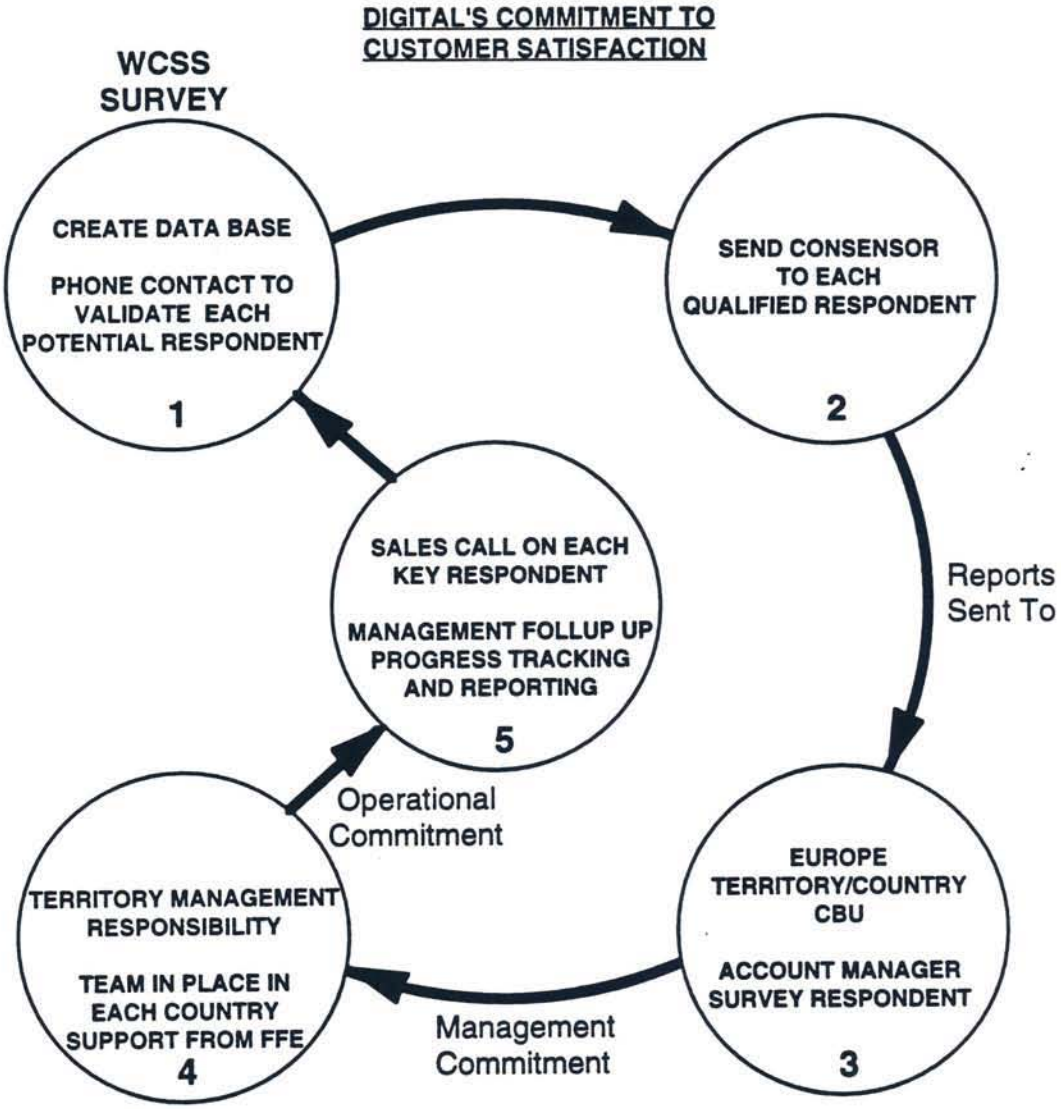

EXECUTIVE PARTNERING OPPORTUNITIES

EXECUTIVE PARTNER	ACCOUNT	WHEN SURVEYED
Enright	Northern Telecom	1st Quarter
Bunker	WestPac Bank	TBD
	National Australian Bank	TBD
	Optus	TBD
Choonavala	Cable and Wireless	1st Quarter
	Hong Kong Shanghai Bank	TBD
Poulsen	Anixter	Ongoing
	Arrow	TBD
	Marisel	TBD
	Avnet	TBD
	Wyle	TBD
	Shared Medical Systems	TBD
Lucente	MCI	TBD
	British Telecom	1st Quarter
	British Petroleum	Ongoing
	NTT	TBD
	Pioneer	Ongoing
	Bechtel	TBD
	Alcoa	Ongoing
	APG	Ongoing
Kozlowski	Alcatel	Ongoing
	Bellsouth	Ongoing
Shire	Cerner	1st Quarter
	Roche Holding	TBD
McCabe	EDS/GM	TBD
	General Electric	TBD
	Asea Brown Bovieri	Ongoing
Klein	Shell	TBD
	Unilever	TBD
	Honeywell	1st Quarter
	SAP	TBD
Ryan	Bank of America	Ongoing
	Bankers Trust	1st Quarter
	Coopers and Lybrand	TBD
	Chrysler	1st Quarter
Gullotti	Dow	1st Quarter
	Boeing	TBD

NOTE: Some of the account assignments have not been finalized

FOURTH RECOMMENDATION

Improve the Cooperation for Obtaining Names for the Customer Satisfaction and Relationship Survey

Issue : **Out of 60 initial accounts, only 17 Account Managers provided names. Valid names within the major accounts are very difficult to obtain. The continued success of the survey is dependent upon worldwide cooperation in providing key account names.**

Europe is establishing a process to ensure that the flow of names of valid decision makers and key influencers occurs. It is not scheduled, however, until the third quarter.

Present US name data bases are very poor and Account Manager cooperation is limited.