From: MSBCS::LIPCON "17-Mar-1995 1510 -0500" 17-MAR-1995 15:10:35.53
To: MSBCS::MELLING,HUMAN::CONKLIN
CC:
Subj: Issue with OpenVMS pathworks tracking WNT

From: NEMTS::NEMTS::MRGATE::"GRANIT::MKOTS3::A1::SCHMITT.BOB" 16-MAR-1995 08:03:11.42
To: THUNDR::LIPCON
CC:
Subj: DOW Letter on Pathworks 1

From: NAME: BOB SCHMITT
FUNC: V.P., U.S. Marketing
TEL: 264-5662 <SCHMITT.BOB AT A1 at MKOTS3 at MKO>
To: BILL DEMMER @BXB,
JESSE LIPCON @PKO
CC: ANNE COPP @MKO,
SCOTT ROETH @MKO,
DEBBIE MILLER @MRO

The attached letter may have already reached your desk, but I wanted to emphasize one point that I think the letter brings out. That is - we can no longer just look at our release dates for SW, we have to coordinate with a number of key partners as part of our planning process.

I know that it is difficult today just to coordinate our own products, much less try to do it with a set of partners. But this is going to be a requirement in the future that we will have to learn to deal with.

Regards

From: NAME: BARRY CLARK
FUNC: Midwest Region
TEL: (708)806-5788/DTN 474-5788 <CLARK.BARRY@A1@POBOX@ACI>
Date: 15-Mar-1995
Posted-date: 15-Mar-1995
Precedence: 1
Subject: FWD: Dow Chemical's Issues with Digital's PATHWORKS
1
To: barrie hunter@mko
CC: ann copp@mko,
bob schmitt@mko,
NAME: Thomas Healy @ACI <Thomas Healy@ACI>

More information on Pathworks issue at Dow, in the customer's words.

Barry

From: NAME: BRENDA SUGAR @OHF
FUNC: DOW AND EDS GLOBAL ACCOUNT TEAMS
TEL: 810-347-5306 <SUGAR.BRENDA@A1@GLDOA@ACI>
Date: 14-Mar-1995
Posted-date: 14-Mar-1995
Precedence: 1
Subject: Dow Chemical's Issues with Digital's PATHWORKS
1
To: See Below
CC: See Below

THE MEMO IS FROM TOM DIETSCH

Attached you will find a letter from Dow Chemical to Scott Roeth stating the reasons they are moving away from our PATHWORKS products to Microsoft NT. This represents $4M in lost revenue to Digital Equipment and one of our largest reference accounts.

I'm requesting your assistance to help resolve these issues and get back the $4M in revenue for Digital.

Please contact me at your earliest convenience to discuss this important issue.

Regards,

Tom Dietsch

TD/bas

DICTATED NOT READ

To Distribution List:

Scott Roeth@MKO,
roger rose@ACI,
Barry Clark@ACI

CC Distribution List:

TOM DIETSCH@HSO,
Czarena Siebert@HSO,
Shane Patterson@MKO

From: NAME: <BRENDA SUGAR @OHF >
Date: 13-Mar-1995
Posted-date: 16-Mar-1995
Precedence: 1
Subject: Michael Hipp Letter
1

March 10, 1995

Scott Roeth
Digital Equipment Corporation

cc: Tom Dietsch
Czarena Siebert

The Dow Chemical Company and Digital Equipment Corporation have a long history of a mutually beneficial business relationship. A prime example of this is the close working relationship that has existed between Digital's Pathworks group and the workstation function within Dow. The purpose of this letter is to offer the most critical issues we face - related specifically to the Pathworks and workstation areas.

The next 1 - 2 years will see massive changes in the information systems at Dow. We are actively implementing a client/server architecture that has required a re-thinking of almost everything in our infrastructure. Perhaps most pertinent to the Dow/Digital relationship is our decision to adopt Microsoft Windows 95 as our single desktop operating system. Related to this is our migration to TCP/IP as our primary network protocol, our adoption of Microsoft Mail with Digital's Mailworks, and inclusion of Microsoft Windows/NT in our midrange operating systems portfolio.

At your request, and with the understanding that Dow is at a key juncture in information systems technology, we offer these five critical issues in the workstation arena, listed in priority order:

1. Delivery of Pathworks for Windows 95 "gold code" at the same time we receive Windows 95 "gold code" from Microsoft. In particular, we require a robust DECnet stack, LAT stack, and the Pathworks for Windows 95 SDK. Need date: May 1995

2. Delivery of Windows-NT server functionality (LANman 3.0) in the Pathworks for Open VMS server. Specifically, we require NT-style long filename support, support for trusted domains, and the ability to browse Open VMS servers from the the Windows 95 Explorer. Need date: 3Q 1995

3. A change of the Pathworks license system to be a customer-oriented monitoring and reporting tool rather than a vendor-oriented enforcement system.

4. A network address administration system for DECnet similar to that provided for TCP/IP by DHCP.

5. Delivery of the MAPI 1 MailWorks driver for Windows 95 Message Exchange.

I would like to call particular attention to item numbers 1, 2, 4, and 5. These appear on this list, to a great extent, because Digital is consistently behind Microsoft in compatible releases of products. In

item number 2 (NT Server functionality), it could be argued that Digital is as much as 36 months behind Microsoft. I am sure we both agree that this is not acceptable. We need Digital to keep up with Microsoft.

We appreciate the commitment that Digital has shown toward Dow in asking for this letter. We look forward to working closely with you on these issues.

Best regards,

Michael Hipp
Manager, Global Workstation Support Center

From: VMSMKT::CUCCIA "Mike, OpenVMS Marketing Manager, DTN: 381-0549, ZK03-2/S
11 21-Mar-1995 1630 -0500" 21-MAR-1995 17:01:56.06
To: MSBCS::SWANTON,HUMAN::CONKLIN,STAR::RMARCELLO,MSBCS::MELLING,STAR::BMATT
HEWS,STAR::STEBULIS
CC: CUCCIA
Subj: Michael Goulde Write Ups - Digital Confidential Do Not Forward

Attached are the write ups we asked for from Michael Goulde of Seybold. He wrote 2 versions - one as a press person, which is positive, and one as a consultant/analyst who is negative.

The second (negative one) is the more useful as it points out the holes we need to fill. I'd like to work with him to develop the "heavy ammo" in anticipation of the analysts who are pre-disposed to be negative.

The positive write up also has some good ideas we should consider - a "project name" to use externally a short hand for the "OpenVMS and Windows NT Integration Strategy", a single (high level) person who is responsible and can control the WHOLE project, and a Bill Gates quote.

Mike

++
From: MAZE::MAZE::FUSCI "Ray Fusci OpenVMS Marketing dtn-381-6136 20-Mar-1995
1547" 20-MAR-1995 15:49:11.95
To: CUCCIA,AKOCOA::PORTER
CC: FUSCI
Subj: Responses from Michael Goulde

From: SMTP%"mgoulde@psgroup.com" 17-MAR-1995 16:31:07.43
To: Ray.Fusci@zko.dec.com
CC:
Subj: OpenVMS Press Release

Ray:

I am assuming that this document is confidential within your group.

The "Analysts" report is worst case - and illustrates the objections that would have to be addressed right up front and with heavy ammo.

I think there is more to be done here as far as positioning this and in the messages. If you decide to go ahead, we'd be happy to work with you on the strategy and roll-out.

Michael

Digital Equipment Announces Strategy to Integrate OpenVMS and Microsoft Windows NT

MAYNARD, MA, Month Day, 1995. Digital Equipment Corp. today announced Project Gold Spike, a comprehensive strategy for allowing customers to use OpenVMS to complement Microsoft's Windows NT with a range of services in the areas of high performance, reliability, manageability, and interoperability. "East meets West with the best features of Digital's OpenVMS and Microsoft's Windows NT converging to provide superior capabilities that customers can rely on for their enterprise applications.", said Bill Gates, Microsoft's Chairman and CEO. By porting selected Windows NT APIs to OpenVMS and providing interfaces to OpenVMS services that can be accessed by Windows NT applications, Digital will be able to offer customers the combined capabilities of the two platforms to support enterprise distributed applications. Used in combination, the two platforms provide the richest set of services in the industry for supporting a variety of application types.

Windows NT applications will be able to take advantage of industry-leading capabilities that OpenVMS has an established reputation for, including:

High performance database server, including record high TPC benchmark results.
Comprehensive system and network management using the suite of Polycenter NetView products.
High availability solutions, including the use of OpenVMS clusters to support Windows NT Servers.
Interoperability with Digital's and other vendors' open systems based on standards, as well as interoperability with proprietary environments, such as IBM's SNA.
The high performance that OpenVMS and Alpha AXP brings to computational applications
Advanced technology, including the log structured file system, that will provide customers with leading edge capabilities not yet available in Windows NT.

Answering Customer Requirements
A large number of customers are evaluating Windows NT for a range of applications, and many have begun implementing Windows NT solutions. Often, they are finding that for very large, enterprise-scale applications, Windows NT needs the support of a strong enterprise-tier of capabilities such as those provided by OpenVMS. This is particularly the case when customers are implementing more advanced forms of client/server architectures, such as three tier architectures and distributed object designs. Three tier applications require the flexibility that Windows NT provides in the middle tier for coding business rules and logic or for building business information objects and the large scale database management capabilities that OpenVMS provides. Distributed object applications require the capabilities that OpenVMS provide with ObjectBroker and other object-oriented tools.

Digital believes that customers should not have to comprise in their pursuit of:

HIGH AVAILABILITY. Customers need to be able to deploy very highly available systems to support critical applications. OpenVMS clusters can provide around the clock support to Windows NT systems as database and file servers. Data integrity requirements are uniquely supported through Digital's new log journaled file systems that makes continuous, on-line backup of data reality.

SYSTEM AND NETWORK MANAGEMENT. Requirements for managing large networks of systems are addressed with OpenVMS' Polycenter NetView products. These allow customers to manage their entire network, perform configuration management, install software, and generally address this costly area in an efficient manner by managing OpenVMS, Unix, and Windows NT systems in a common manner.

HIGH PERFORMANCE. Customer's have a continuing need to improve overall performance. OpenVMS addresses this with industry-leading benchmark performance, further enhanced by the log journaled file system that can offer up to a ten-fold improvement in disk read/write performance. Computational performance requirements are addressed by Alpha's high speed floating point performance as well as the 64-bit implementation of OpenVMS. Making this class of performance available to Windows NT systems is a unique capability for Digital.

INTEGRATION AND INTEROPERABILITY. Integration and interoperability requirements are addressed by Digital's work to extend support for de facto and industry standards across both OpenVMS and Windows NT. Customers will be able to develop applications once and deploy on either platform.

Strategy Enabled by Common Heritage

OpenVMS and Windows NT have a common heritage. Many of the designers and developer's who working on Microsoft's new operating system had played key roles in the design and development of VMS. Close comparison of the two operating systems shows that they have a highly compatible structure, and similar process and application models. This consistency across the two environments makes it straightforward for Digital to provide the support of Windows NT it intends to deliver.

Several of Digital's development activities in support of this strategy will result in even more commonality across the two platforms. Specifically:

Digital will work to port the Win32 API set to OpenVMS. This work will be done in stages and will concentrate on those APIs that make sense to implement under OpenVMS. This work will enable applications to call either Win32 APIs or OpenVMS APIs transparently. In addition, Windows NT commands will be enabled on OpenVMS.

APIs for Microsoft's BackOffice, such as MAPI, ODBC, and application control APIs will be supported in OpenVMS, as well as those OLE interfaces that are a part of BackOffice, and the Windows NT Performance Monitor and Event Logger as well.

Log structure file system APIs will be made available on Windows NT so that developers can take advantage of its performance and backup capabilities in their applications.

Digital will work to bring as much of XPG4 compliance to Windows NT as possible, supporting cross-platform application development through standard interface specifications.

Digital's MailBus will be integrated with Microsoft Exchange to provide a robust, universal mail backbone for distributed Exchange servers.

As a result of this work, developers will be able develop applications in Visual C++ on Windows NT and compile for OpenVMS, system managers will be able to manage OpenVMS from a Windows NT management console,

Project Golden Spike
The work to integrate OpenVMS and Windows NT services will take place within several groups at Digital. Coordinating the project will be Jane Doe, VP of Stuff. "Work has already been underway for a number of months.", said Doe. "I fully expect us to be able to deliver complete solutions to customers by the end of this year." The components that will be available first are those that entail delivering OpenVMS services to Windows NT. Delivery of the ports of the Windows NT APIs to OpenVMS and the addition of XPG4 compliance to Windows NT will constitute a second phase of deliverables.

BLABBERDEEN GROUP WATCH
Some Month Day, 1995

Breathing Life into a Dead Horse: Digital Tries to Use NT to Resuscitate OpenVMS

Trying to bring the dead back to life often only results in creating zombies. The Blabberdeen Group believes that Digital's announcement that it will integrate OpenVMS with Windows NT is one part smoke, one part mirrors, and three parts desperation.

Customer Lock In
First of all, we think that this strategy is just a way to keep customers locked into OpenVMS. After all, OpenVMS still represents a significant revenue stream to Digital, and the last thing the company needs to see is that revenue slipping through its fingers. If it can convince customers to hold on to OpenVMS a while longer, then that will give Digital some breathing room to figure out how to make Digital Unix more appealing to customers than it was when they called it OSF/1.

Life Extension
Second, this strategy doesn't really benefit customers, it is just a way to prolong the life of OpenVMS. If customers can be convinced that applications running on Windows NT will hit a ceiling and that OpenVMS on Alpha systems can raise that ceiling, then OpenVMS can hang around longer. But how does Digital propose that that ceiling will be raised? OpenVMS is to provide clustering, fault resilience, and disaster tolerance. Since the two platforms aren't binary compatible, how will Windows NT applications benefit? It seems to us that it would be a lot cheaper to provide Windows NT systems with large quantities of RAID storage and big UPSs to achieve similar results. And Windows NT's built-in mirroring and replication capabilities are probably sufficient for the kinds of applications customers will be deploying on Windows NT.

Although OpenVMS can provide high performance for database operations and computations, it doesn't really do that any better or any cheaper than many RISC-based Windows NT systems. After all, I can buy Windows NT on most of the same Alpha systems on which I can run OpenVMS. Where's my advantage to buying OpenVMS? It certainly costs a lot more than Windows NT. Sorry Digital, its time to pull the plug.

Go Native
Digital would probably do better to focus its efforts on developing OpenVMS-like capabilities for Windows NT. It could leverage the Windows NT volume market and have something valuable and unique to offer customers. NT Clusters and some of the Polycenter NetView work was a good start. Are there really any capabilities in OpenVMS that couldn't be made available in Windows NT?

Are You Sure Customers Care?
In some ways, this announcement is another instance of Digital putting too many eggs in a single basket. The company is betting that Windows NT is going to take off and be real popular outside the workgroup. Customers aren't convinced yet, and if they reject Windows NT in favor of Unix or AS/400, then Digital is going to be left high and dry. Secondly, do customers really want to have to use two operating systems in order to get functionality they used to get in one? All it means is additional skill, personnel and cost. We believe most customers will just say no.

Blabberdeen Group's Take
It's not easy to pull the plug on a product. It means lost revenues and loss of jobs. But we think this whole OpenVMS/Windows NT integration pitch is a ploy on the part of the OpenVMS people at Digital to keep their jobs. Now these are very talented and capable people. Digital would be silly to lose them. Why not put them to work making Digital Unix the industry's best Unix and adding enhancements to Windows NT that could piggyback any success Microsoft is able to have? We think this is a better use of Digital's limited resources.

Consultant's Summary:

Proposition: Bringing the advantages of OpenVMS to Windows NT in the form of products, functionality, and capabilities that will benefit customers.

Points to be made:
Which advantages of OpenVMS will complement Windows NT?
Who are the customers that will benefit?
How will customers benefit from these advantages?

How will these advantages be implemented?

Negative points to anticipate:
Strategy is just a way to keep customers locked in to OpenVMS.
Just a way to prolong the life of OpenVMS.
Why not just develop the capabilities for natively for Windows NT?
Isn't there a risk that customers don't care about NT?
No real benefits that can be proven.

Advantages:

1. OpenVMS as Enterprise Database
Contain large databases that don't fit on NT platforms
Provide high performance database support for client/server applications
Focal point for rationalizing multiple data sources
Support NT as the middle tier of business objects
Support NT with WAN communications

2. OpenVMS-based availability
Clustering provides highly available data management for NT applications
Disaster tolerance

3. OpenVMS-based management
Mature, rich network and system management environment
Consistent Polycenter NetView management from desktop to data center

4. OpenVMS for interoperability
Formal and de facto standards support allows NT to interoperate with other vendors open and proprietary systems

5. OpenVMS for distributed application
Middleware support on OpenVMS to support range of distributed application support

6. OpenVMS + Win32:

Develop on Alpha/NT, deploy on Alpha/OpenVMS
Familiar desktop development environment
High performance compiling, debugging, testing environment
Deploy on optimum price/performance platform
Extends scalability of Win32 to upper reaches of Alpha
Brings Win32 applications within management framework of OpenVMS

7. OpenVMS + XPG4 + NT + XPG4
Complete open systems solution
Develop to standard and deploy on wide range of scaleable platforms
Reduce cross-platform development costs

8. Advanced Technology Support
64-bit high performance computation
Log structured file system adds high performance I/O to NT applications
Log structured file system adds high speed backup for NT applications
Log structured file system for continuous backup.

How:
Object-based interfaces to OpenVMS Services
Support Win32 APIs on OpenVMS
Formalize existing support capabilities

Michael Goulde
Editor in Chief
Patricia Seybold Group
148 State Street
Boston, MA 02109

Voice: 617.742-5200
Fax: 617.742-1028

<----- End Included Message ----->

Michael Goulde
Editor in Chief
Patricia Seybold Group
148 State Street
Boston, MA 02109

Voice: 617.742-5200
Fax: 617.742-1028

---------- RFC 822 Headers ----------
Received: by ilovit (5.65/fma-100391);
id AA01674; Fri, 17 Mar 1995 16:31:42 -0500
Received: from uu7.psi.com by inet-gw-1.pa.dec.com (5.65/24Feb95)
id AA15816; Fri, 17 Mar 95 13:27:00 -0800
Received: from psgroup.com by uu7.psi.com (8.6.9/SMI-4.1.3-PSI)
id QAA12202; Fri, 17 Mar 1995 16:09:38 -0500
Date: Fri, 17 Mar 1995 16:09:38 -0500
From: mgoulde@psgroup.com
Received: from bravo.psgroup.com by bergman.psgroup.com id aa15783;
17 Mar 95 15:40 EST
To: Ray.Fusci@zko.dec.com
Subject: OpenVMS Press Release
X-Mailer: AIR Mail 3.X (SPRY, Inc.)
Message-Id: <9503171540.aa15783@bergman.psgroup.com>

From: MSBCS::LIPCON "15-Mar-1995 1422 -0500" 15-MAR-1995 14:21:44.82
To: MSBCS::MELLING,HUMAN::CONKLIN
CC: THUNDR::LIPCON
Subj: ALL-IN-1/Mailworks issue FYI

From: NEMTS::NEMTS::MRGATE::"GRANIT::A1VAX::A1::ANDERSON_C" 13-MAR-1995 16:12:52.93
To:
CC: THUNDR::LIPCON
Subj: RE: Reuters Help

From: NAME: Curt Anderson
FUNC: Workgroup Systems-Messaging
TEL: 603-881-6154 <ANDERSON_C@A1@A1VAX@ZKO>
To: See Below
CC: See Below

Alex -

I am the Product Line Business Manager for our messaging products in the Software Business Group. I have been forwarded a copy of your mail to Bill Demmer regarding the account situation at Reuters and requesting a clarification of our messaging strategy.

First let me address the plans for our messaging suite of products for the OpenVMS platform.

Later this month, Digital will formally announce Version 1.3 of MailWorks for OpenVMS. This will extend MailWorks functionality to the Alpha OpenVMS platform and increase performance by up to 300% in terms of number of users supported on Alpha.

Secondly, we have just made the decision to integrate the new MailWorks V1.3 server with MAILbus 400 and DEC X.500. This will give our OpenVMS customers a complete messaging solution on the Alpha platform. The target date for the availability of this release is the end of this calendar year. Engineering is now putting together the detailed work plan for this release and we expect to have a more firm schedule by mid-April.

Also new is that MAILbus 400 is now available on Alpha OpenVMS. For Reuters, your account plan could include proposing, now, an Alpha based prototyping system with XMR, MAILbus 400, and DEC X.500 so they can have all the time they need to test the new backbone and to build their X.500 directories while they continue to run Message Router on their existing ALL-IN-1 and MailWorks servers. This approach would provide Reuters IT staff the assurances needed when it is time to cut-over from their existing Message Router backbone.

In summary, we are moving forward with an aggressive program, with committed engineering investments, for our suite of messaging products on the OpenVMS platform.

Regarding your point about the positioning of MailWorks and ALL-IN-1, MailWorks is a low cost messaging server that supports heterogeneous mail clients. ALL-IN-1 provides substantial other office services, beyond TeamLinks mail, and is priced accordingly.

Whatever server is chosen by Reuters IT staff to deploy in a particular end-user environment, both offer connectivity to MAILbus 400 and DEC X.500, the industry's leading X.400 backbone according to industry analysts. Our approach offers IT management:

- A choice of servers, and platforms, based upon end-user needs
- A single, enterprise-wide backbone that will provide interoperability between the servers while offering a reliable, standards-based, high volume message transport service for the entire organization.

I hope this addresses your questions about Digitals messaging strategy and that the information will prove useful in communicating our strategy to Reuters. If you have any further questions, please let me know.

Regards,

Curt

From: NAME: Alex Black
FUNC: Corp Technology Cons, Barclays
TEL: 847-6479 (0836-238690) <BLACK AT A1WOTVAX at WOTVAX at OLO>
Date: 27-Feb-1995
Posted-date: 27-Feb-1995
Precedence: 1
Subject: Reuters Help
To: bill demmer@pko

Hi Jesse and Bill,

Sorry to bother you both, but we seem to be on a little bit of a collision course with regard to our office strategy, at least within Reuters. I have discussed this with a number of my peers and they are facing a similar problem in other accounts.

Today we seem to be veering towards competing client/server mail environments, namely :

1) TeamLinks and DECmailworks

2) TeamLinks and ALL-IN-1

In the past they have been very different, but we now have DMW growing up toward A1 and conversely, we have A1 growing downwards toward DMW.

I would like your help in pointing me at someone who can help give Reuters a very clear direction as to what our future strategy is going to be. One of their concerns relates to the affordability of two environments.

I have tried to resolve this conflicting position, however this has not been an easy task given the organisational environment and the strengths / weaknesses and directions
of both product sets.

I hope that you can point me in the right direction, since this problem is now stopping further deployment of TeamLinks in Reuters (currently at 7,500 users), which affectes short term revenues. In addition, it is also beginning to impact medium term opportunity.

Once again, I am sorry to trouble you!

Thanks

Alex

PS Jesse, for you info, things are going a lot better with Reuters at the moment, we are not over the hill yet, but doing much better.

To Distribution List:

alex black@hhl

CC Distribution List:

NAME: Gerard van de Aast <VANDEAAST@A1@A1VAX@ZKO>,
gail holland@zko,
bill demmer@pko,
jesse lipcon@pko,
jenkins@a1vax@mrgate,
Shayan@hotpup@mrgate,
seth cohen@zko,
kell@forty2@mrgate

From: MSBCS::LIPCON "15-Mar-1995 1602 -0500" 15-MAR-1995 16:01:32.52
To: MSBCS::MELLING,HUMAN::CONKLIN
CC:
Subj: Mailworks

From: NEMTS::NEMTS::MRGATE::"GRANIT::A1VAX::A1::ANDERSON_C" 13-MAR-1995 13:02:20.40
To:
CC: THUNDR::LIPCON
Subj: RE: DECmail issue

From: NAME: Curt Anderson
FUNC: Workgroup Systems-Messaging
TEL: 603-881-6154 <ANDERSON_C@A1@A1VAX@ZKO>
To: See Below
CC: See Below

Paul -

I am the Product Line Business Manager for Messaging Software Products. I have been asked to respond to your mail to Bill Demmer concerning our Messaging Strategy.

The focus of Digitals messaging program is to provide accounts with a heterogeneous solution for enterprise messaging. IT management, in these accounts have expressed the following needs:

- o "Mission Critical" mail service for all users
- o Mail Servers that will connect to a range of desktop platforms and leading mail clients.
- o A portfolio of mail servers for different platforms that IT can deploy to fit individual departmental requirements, e.g.:
  - UNIX mail server platforms for R&D, technical departments, etc.
  - OpenVMS servers for departments running Digital applications
  - NT servers for branch offices and workgroups
- o A single, enterprise backbone that will provide a standards-based, reliable, secure, and high volume message transport service for all the different mail server platforms.

We are moving forward with a program of product investments to meet these needs. Some of the products will be delivered by Digital and some will be delivered through business partners, in conjunction with joint development work with Digital.

With regard to the specific customer needs you raised in your mail:

OpenVMS Strategy

Our plans call for extending, enhancing, and aggressively marketing our messaging suite of products on the OpenVMS platform.

Latter this month, Digital will formally announce version 1.3 of MailWorks for OpenVMS. This will extend MailWorks functionality to the Alpha OpenVMS platform and meet your request to migrate your accounts existing VAX MailWorks servers to Alpha. This product is shipping now.

Secondly, we have just made the investment decision to integrate the new MailWorks V1.3 server with MAILbus 400 and DEC X.500. This will meet your request to provide your OpenVMS customers with a complete messaging solution on the Alpha platform. The target date for the availability of this release is the end of this calendar year. Engineering is now putting together the detailed work plan for this release and we expect to have a firm schedule by mid-April.

Also new is that MAILbus 400 is now available on Alpha OpenVMS. For your installed MailWorks VMS accounts, your account plans could include proposing an Alpha systems NOW, with XMR, MAILbus 400 and DEC X.500 so your account can have all the time they need to prototype and test the new X.400 backbone and to build their X.500 directories while still running the Message Router backbone. This would serve as an excellent vehicle for IT to test the new backbone, in a more relaxed timeframe, and give them the assurances that, when its time to convert to the new backbone , everything will be ready.

In summary, our strategy is based upon a two-phase migration of MailWorks: first the server, then the backbone. This strategy is based upon the fact that we understand the complex migration issues involved and that Digital is protecting it's customer's past investments.

In your mail, you requested that the performance of the Alpha OpenVMS suite of products be at least as fast as the current VAX product suite. Based upon internal evaluation tests of the new MailWorks V1.3 server, we expect up to a 300% increase in users on an AXP system when compared to the current version on VAX. We expect that when the complete product suite is available on AXP, performance will again improve. As soon as the performance and characterization testing and evaluation work for Version 1.3 is completed, we will make it available to account teams for use as a configuration tool.

NT Strategy

Regarding our direction for mail servers on NT, we plan to offer message interoperability support with the Exchange Server from Microsoft as part of our MAILbus 400 enterprise messaging strategy. We know that some of our customers want to deploy Exchange in parts of their organization and want it to tie into a robust, proven messaging backbone for the organization. Our strategy is to work with Microsoft to deliver this.

With regard to our messaging transport, MAILbus 400 is the industry's leading X.400 MTA according to industry analysts. It now runs on OSF/1, and OpenVMS Alpha platforms. It has been tested and evaluated by a number of accounts who found it to be the fastest, and most reliable X.400 product on the market. Our plans are to extend MAILbus 400 to provide connectivity to NT based mail servers.

Regarding directories, Digital is the world's first vendor to conform to both U.S. GOSSIP and OSTC (Open Standards Testing Consortium) standards. Extensive testing by customers reveal that DEC X.500 is the price/performance leader in the industry today. Our plans are to enhance DEC X.500 to provide connectivity to NT mail servers.

Regarding MailWorks, we have not yet concluded whether we will offer this product on NT. Technical and business evaluations are currently underway.

MailWorks Client Strategy

Our strategy is to continue to extend the choice of clients that are available for our MailWorks servers. In December, Digital announced support for cc:Mail on Windows for both our OpenVMS and OSF/1 MailWorks servers. At that time we also announced our plans to support MAPI based clients within 60 days after commercial availability from Microsoft.

For the OSF/1 MailWorks server, we have announced support for SunOS that will be shipping in April. Latter this spring, Digital will announce support for Sun Solaris. After looking at the business feasibility, we have decided not to offer Motif clients for the OpenVMS version of the MailWorks server.

Both of the Sun motif clients were customer funded developments. To meet your request to support the HP UX and IBM AIX platform, we propose to do the work for $30K for each client.

Regarding a VT client, the VT/DECform software that you referred to in your

mail is an ASSETS product. We have no plans to make it a Corporate software product. Our strategy is to tap into the wealth of shareware MH software by enhancing the version of MH within the current OSF/1 server. This would give our customers access to a wealth of free "shareware" for MH clients available on the internet and other sources.

I trust that this addresses the questions raised in your mail to Bill Demmer and that the information will prove helpful in communicating our messaging strategy to your accounts. As we finalize the schedule for the next release of MailWorks on OpenVMS and when the NT plan is complete, we will have more information to pass on. In the meantime, if you have further questions, please contact me.

Regards,

Curt

I N T E R O F F I C E M E M O R A N D U M

Date: 15-Feb-1995 02:32pm CET
From: Paul Santner @ZUO

SANTNER
Dept: CMT
Tel No: 760-2405

TO: Remote Addressee (WILLIAM DEMMER @PKO)

CC: Remote Addressee (Vincenzo Damiani @GEO)

Subject: SW Strategy for ABU Accounts

We have some very important open questions relative to DECmail (see attached mail) which need to be answered. Actually our whole office strategy is in question. We have important business dependencies on that.

The main consequences are that we cannot make VAX to Alpha AXP migration and we, therefore, significantly damage our AXP-message.

Thanks for your consideration.

Regards
Paul

I N T E R O F F I C E M E M O R A N D U M

Date: 13-Feb-1995 06:40pm CET
From: Pius Haas @ZUO

HAASP
Dept: Financial Services
Tel No: ++41-1-801'3268 or '2111

TO: Andreas Duerst @ZUO (DUERST)

Subject: A: DEC Mail Solutions/ Escalation

Andi,

based on our discussions of the problems with the DEC mail solution capabilities, I've summarized the evidenz for immediate action in the attachement.

Please escalate it as fast as possible. It is of very much urgency with respect to our situation at HKK, SWX and others.

Thank you,

Pius

We request:

1. We request, that with immediate action a date is committed to have DECmailworks (OSF/1 code base?) available on OpenVMS AXP that interoperates directly with Mailbus400 OpenVMS AXP and X500 Directory Server for OpenVMS AXP.

2. We expect a committed date before or at June 30, 1995.

3. We request a DECmailworks Client for VT that is equal to the existing DECmailworks for VT/DECforms client interface (on OpenVMS VAX) and that can connect to both, DECmailworks Server on openVMS AXP and OSF/1.

4. We request, that DECmailworks (OSF/1 Code base?) and directory lookup (INFObroker) will be available soon thereafter (6 month) on Microsoft NT Alpha and Intel.

5. We request, DECmailworks Motif Clients for DECmailworks to be available at the same time.

6. We request, the DECmailworks Motif Clients can connect to both DECmailworks Server on openVMS AXP and OSF/1 (OSF/1 code base on both platforms?) and do have the same look and feel across the following client platforms: openVMS VAX, openVMS AXP, OSF/1, SUN OS, HP-UX and IBM AIX, SUN Solaris (in priority order).

7. We request, that the performance of this OpenVMS AXP Messaging Suite is at least the one of the current MailbusMR and DECmailworks for OpenVMS VAX.

Customers demanding:

Krankenkasse Helvetia: The Current Messaging Infrastructure of the largest Swiss Health Insurer needs to be migrated onto a DEC 7000 model 710 AXP system. Currently, it is not possible to migrate the OpenVMS VAX based Messaging Infrastructure to an OpenVMS AXP based system. There is no complete suite of Backbone and Site/Workgroup (message store) services available on OpenVMS AXP.
Pot. Cost/Loss for DEC k$ 200

The Swiss Exchange (former Soffex, Swiss Stock exchange, Swiss Electronic Stock exchange) needs to offer additional Public information services. This needs to be based on X.400. As they are just in the process of standardizing the DEC environment onto AXP and OpenVMS, this can not be achieved in a cost-effective manner. Therefore, they are reluctant to both, standardize onto AXP and establishing DEC Messaging solution. As this is a direct service limitation of the SWX (SWiss EXchange) they may turn to another vendor to overcome this.
Pot. Cost/loss for DEC k$ 250

Paul Scherre Institut (PSI) has bought a total DEC Messaging infrastructure. Part of the contract was SPD 45.16.02 from Sept. 1994. it explicitly states the DECmailworks Clients for SUN OS, HP-UX, DEC Ultrix, IBM AIX. As those do not work resp. are no longer available, the consider dropping the order.
Pot. Cost/Loss for DEC k$ 400

Other Customer demanding:
ASCOM, the largest Swiss Telecom Supplier
SRG, the Swiss government controlled radio and television

society

Reaction of Customer:

Bad Image: DEC does not take care of past investments (no investment protection)
DEC is no open System vendor: Customer is locked to MailbusMR and VAX or OSF/1 AXP.
DEC is prohibitive expensive: Migration (VAX to OSF or VAX to AXP)costs close as much as initial investments (SW and manpower).
DEC complicates environment: Having a VAX purely to run Message router (a dedicated server) does not add value but complicates Messaging infrastructure. Such behavior is considered a tactic to sell more HW.
DEC is NO multiplatform company: When the demand is rather a solution than products, we (customers) are confronted with only one option are otherwise forced to heterogeneity (Intel, VAX, AXP, VMS, UNIX, NT)

Cause: No clear Messaging strategy.
DEC has failed to deliver a coherent Messaging product portfolio.
Within DEC are few new products in the Messaging area. The only thing addressed out of a three year old Vision "Digitals Enterprise Mail" is two complete different Mailsuits, one on OSF/1 and one on OpenVMS VAX. Whereby OpenVMS VAX for Messaging is a dead end. Anything not fully in one of the two suits causes high cost, can not be solved or is not stable.
The only one thing in common of the two suits is the NAME of the Site/Workgroup server (message store) DECmailworks (NOTE: In fact two complete different products)

Explanations:

1. We request, that with immediate action a date is committed to have DECmailworks (OSF/1 code base) available on OpenVMS AXP...

No server based DMW messagestore migration is possible. Therefore each individual user has to transfer his store himself. We can not afford to have the users doing this cumbersome work multiple times. In addition, to maintain operation of a VAX system only for Message router in an all AXP based environment is close as 50% of the operational cost (based on contract offering from OMS for one VAX 3100 for Message router and one DEC 7710 AXP).

2. We expect a committed date before or at June 30, 1995.

Customer attention drives away from DEC, as we only talk about how it will look like since over 12 months.
We have no substantial progress made over the last 2 years. At this time two Messaging solutions exist that are feasible in terms of price and performance (VAX VMS or AXP OSF/1). The VAX VMS solution will clearly become a dead end. So strategically seen, only a total OSF/1 based solution could honestly be offered.
New products on openVMS AXP are only targeted at the ADMD market space (e.g. MCI, Telecom) which do not need DEC Mailworks.

3. We request a DECmailworks Client for VT that is equal to the existing DECmailworks for VT/DECforms client interface (on OpenVMS VAX) and that can connect to both, DECmailworks Server on openVMS AXP and the one on OSF/1.

Based on this VT/DECforms interface, we have as a strategic differenciator. The capability to provide total user coverage. This is what allows an IBM 3270 Terminal user to be incorporated into the Messaging infrastructure (thousands of users still do not have a PC, and to replace all 3270 terminal only to do Mail can most often not be cost justified over a short time).

4. We request, that DECmailworks (OSF/1 Code base) and directory lookup (INFObroker) will be available soon thereafter (6 month) on Microsoft NT Alpha and Intel.

When DEC is providing a Messaging Backbone, its is always in a mixed environment such as IBM, DEC, PC's and NT. In the last 6 month, approx. 6 out of 10 customer visits showed NT Servers for Workgroups. These NT Servers grow under departmental competency. Mail Backbone is driven by corporate IT departments. To install the messagestore in the departments, it can not be justified to add just another workgroup server based on OSF/1 and even less such a one on OpenVMS (2 to 4 times the cost vs. a software installation on an existing NT server)

5. We request, DECmailworks Motif Clients for DECmailworks to be available at the same time.

A mixed environment on the DESKtops is the regular case in special application areas (e.g. trading area). Many openVMS Stations still exist for special purpose applications. Such environments don't get applications ported and the equipment changed, just to be kept enabled to EXchange mail.

6. We request, the DECmailworks Motif Clients can connect to

both DECmailworks Server on openVMS AXP and OSF/1 (OSF/1 code base on both platforms?) and do have the same look and feel across the following client platforms: openVMS VAX, openVMS AXP, OSF/1, SUN OS, HP-UX and IBM AIX, SUN Solaris (in priority order).

As explained under 5., the application determined DESKtops are also covered from most of the popular workstation vendors such as openVMS VAX, openVMS AXP, OSF/1, SUN OS, HP-UX and IBM AIX, SUN Solaris. We were able to propose total integration to such demand based on the POSTE clients. However, those client do not exist at the moment which limits our solution capability to an all DEC environment. Pls. note, the cancellation of the POSTE clients has been done without a formal retirement phase and a robust working version has never been out. But it has been stated as an valid client in the SPD 45.16.02 from Sept. 1994.

7. We request, that the performance of this OpenVMS AXP Messaging Suite is at least the one of the current MailbusMR and DECmailworks for OpenVMS VAX.

While we understand, that the OSF/1 platform and its mailsolution is really high performance, we feel fine in the migration scenario (from OpenVMS VAX to OpenVMS AXP) to get the performance improvements which were at least the ones we would have got if we would run message router on OpenVMS AXP.

To Distribution List:

paul santner@zuo,
pius haas@zuo

CC Distribution List:

gail holland@zko,
NAME: Gerard van de Aast <VANDEAAST@A1@A1VAX@ZKO>,
bill demmer@pko,
vincenzo damianai@geo,
seth cohen@zko,
Shayan@hotpup@mrgate,
jenkins@alvax@mrgate,
kell@forty2@mrgate,
jesse lipcon@pko

SYSTEMS BUSINESS GROUP – INTEGRATED UNIT FORECASTS

SUM OF 3 O/S	1994	1995	1996	1997	1998	CAGR
Servers Total	34,224	55,861	67,838	111,115	157,620	46.5%
High	1,982	3,169	4,148	6,140	7,620	40.0%
Medium	9,951	18,155	19,642	33,789	46,340	46.9%
Low	22,291	34,537	44,048	71,186	103,659	46.8%
Workstations	52,449	64,753	86,648	119,883	161,858	32.5%
WS Segment Forecast		61,700	90,000			
Alpha Clients	0	50,000	100,000	150,000	75,000	14.5%
Total Systems	86,673	170,614	254,486	380,998	394,477	46.1%

OVMS	1994	1995	1996	1997	1998	CAGR
Servers Total	27,479	33,938	35,260	36,702	38,308	8.7%
High	1,811	2,838	2,944	3,064	3,187	15.2%
Medium	9,387	8,892	9,436	10,060	10,704	3.3%
Low	16,281	22,208	22,880	23,578	24,417	10.7%
Workstations	28,263	26,119	27,488	28,856	30,225	1.7%
Total Systems	55,742	60,057	62,748	65,558	68,533	5.3%

Source: Memo from Mark Gorham 2/15/95

UNIX	1994	1995	1996	1997	1998	CAGR
Servers Total	6,745	14,446	19,984	53,201	83,583	87.6%
High	171	331	1,204	3,076	4,433	125.6%
Medium	564	7,508	7,250	18,750	27,250	163.6%
Low	6,010	6,607	11,530	31,375	51,900	71.4%
Workstations	24,186	29,009	41,085	57,083	67,888	29.4%
Total Systems	30,931	43,455	61,069	110,284	151,471	48.8%

Source: '94 - '95 Base O/S Licence Units (Laura Amrein); '96- '98 UNIX Business Plan Draft Financials 1/26/95

NT	1994	1995	1996	1997	1998	CAGR *1
Servers Total	0	7,477	12,594	21,212	35,729	68.4%
High		0	0	0	0	68.4%
Medium		1,755	2,956	4,979	8,386	68.4%
Low		5,722	9,638	16,233	27,342	68.4%
Workstations		9,625	18,075	33,944	63,745	87.8%
Alpha Clients *2		50,000	100,000	150,000	75,000	14.5%
Total Systems	0	67,102	130,669	205,156	174,473	37.5%

Source: NT Business Plan Draft 2/6/95
**1 CAGR of IDC unit forecasts from 1994 - 1998 (except Alpha clients)*
**2 BAH forecast – Alpha client units to peak in 1997 when P7 is introduced*

From: GRANIT::GRANIT::MRGATE::"MKO::GRANIT::OPENVMS::aduncan" 22-MAR-1995 19:30:02.74
To: STAR::RMARCELLO,HUMAN::CONKLIN
CC:
Subj: FWD: Re: Win32 interfaces

From: NAME: Anne Smith Duncan star::aduncan dtn 381-2511
TEL: DTN 381-2511
ADDR: M/S: ZK3-4/W23 <aduncan@OPENVMS@GRANIT@MKO>
To: NAME: rich marcello <rmarcello@star@MRGATE>,
NAME: peter conklin <conklin@human@MRGATE>
Message-id: 3228241322031995/A10912/ZEMBO/1193B3581A00
From: NAME: Anne Smith Duncan star::aduncan dtn 381-2511
TEL: DTN 381-2511
ADDR: M/S: ZK3-4/W23 <aduncan@OPENVMS@GRANIT@MKO>
Subject: FWD: Re: Win32 interfaces
Precedence: 1
To: NAME: bill matthews <bmatthews@star@MRGATE>,
NAME: ann mcquaid <mcquaid@star@MRGATE>,
NAME: steve zalewski <zalewski@star@MRGATE>

The attached mail is the reply from Oracle development regarding the Bristol and Mainsoft Win32 technologies. We told them that both technologies were under consideration and asked for their 2 cents. They are very eager that we get Win32 on VMS fast as it will solve a major porting problem for them.

Please do not quote Oracle as the source should you use this info.

Anne

Message-id: G2020300217MAR199517432646
From: NAME: SMTP%"DIRWIN@us.oracle.com" <SMTP%"DIRWIN@us.oracle.com"@STAR@MRGATE@GRANIT@MKO>
Subject: Re: Win32 interfaces
Precedence: 1
To: ADUNCAN@OPENVMS

--Boundary-7585242-0-0

Thanks for the info Dave.

Ann, this info is obviously confidential and therefore subject to our NDA. If your engineers have some info resulting from your own investigations then Dave is the right person to talk to.

Regards
David Irwin
Senior Director
DEC Products Division
Oracle Corporation

500 Oracle Parkway	Phone:	(415) 506-2498
MS 659105	FAX:	(415) 506-7304
Redwood Shores	Internet:	DIrwin@us.oracle.com
CA 94065	Assistant:	Jill Eastwood JEastwoo (x3819)

--Boundary-7585242-0-0
X-Orcl-Content-Type: message/rfc822

Received: 16 Mar 1995 18:10:53 Sent: 16 Mar 1995 18:06:19
From:"Dave Stowell" <DSTOWELL>
To: DIRWIN
Subject: Re: Win32 interfaces
Cc: jay
X-Orcl-Application: In-Reply-To: DIRWIN.US.ORACLE.COM's message of 16-Mar-95 11:16
X-Orcl-Content-Type: multipart/mixed;
boundary-Boundary-7585242-0-1

--Boundary-7585242-0-1

Dave,

There are two main areas of concern with relation to the Win32 API libraries from Bristol/MainWin:

1. NLS Support
2. OLE2 Support

NLS Support

Bristol is demonstrating it now, MainWin will not have it until Chicago is stable.

OLE2 Support

A big issue for Sm30. Both Bristol and MainWin are planning on

supporting a subset of OLE2 (not OLE2 Automation).

Both products are now about equal in terms of technical merit, Bristol is better at marketing, and has done a number of deals with third party custom control companies so that you can use graphing packages etc. Mainwin is willing to do the same.

We are still in talks with both companies - but the NLS issue may make Bristol and better bet.

Jay, do you have any comments?

Regards,

Dave

Dave Stowell — Direct : (415) 5063176
Server Technologies — Switchboard : (415) 5067000
MS #659507 — Fax : (415) 5061099

'Q': "It's hard being a team player when you're omnipotent"

--Boundary-7585242-0-1
X-Orcl-Content-Type: message/rfc822

Received: 16 Mar 1995 11:16:17 Sent: 16 Mar 1995 11:16:00
From:"David Irwin" <DIRWIN>
To: jrossite,dstowell
Subject: Fwd: Win32 interfaces
Cc: atottle
X-Orcl-Application: Message: Note the new location in Bldg 100. The walk will do you (me?) good.
X-Orcl-Content-Type: multipart/mixed;
boundary=Boundary-7585242-0-2

--Boundary-7585242-0-2

Ann Duncan in Digital is the right person in Digital to co-ordinate on this. Her email is aduncan@star.zko.dec.com.

I have just moved and her card is in the bottom of a box so I don't have her telephone number.

Note that Digital/Oracle have the appropriate NDA in place due to the strategic relationship but they may want to do a specific NDa. Shouldn't be a problem though.

Regards
David Irwin
Senior Director
DEC Products Division
Oracle Corporation

500 Oracle Parkway — Phone: (415) 506-2498
MS 659105 — FAX: (415) 506-7304
Redwood Shores — Internet: DIrwin@us.oracle.com
CA 94065 — Assistant: Jill Eastwood

JEastwoo (x3819)

--Boundary-7585242-0-2
X-Orcl-Content-Type: message/rfc822

Received: 09 Mar 1995 15:14:26 Sent: 09 Mar 1995 15:09:11
From:"ATOTTLE.US.ORACLE.COM" <ATOTTLE>
To: DIRWIN
Subject: Re: Win32 interfaces
Cc: binfante,jrossite,dstowell
Reply-to: atottle
X-Orcl-Application: In-Reply-To: DIRWIN.US.ORACLE.COM's message of 09-Mar-95 03:21
X-Orcl-Content-Type: multipart/mixed;
boundary-Boundary-7585242-0-3

--Boundary-7585242-0-3

Dave:

I appreciate the note. I think it would be helpful to collaborate on this.

Dave Stowell in Jay Rossiter's shop is probably the right person. Jay, pls confirm for me. Thanks,

Alan

Alan Tottle — Phone: (415) 506-2964
VP, Network and Management Products — Fax: (415) 506-1228
Server Technology Division — atottle@us.oracle.com
Oracle Corporation

--Boundary-7585242-0-3
X-Orcl-Content-Type: message/rfc822

Received: 09 Mar 1995 12:25:43 Sent: 09 Mar 1995 12:25:26
From:"David Irwin" <DIRWIN>
To: bvi,atottle
Subject: Win32 interfaces
Cc: bobp
X-Orcl-Application: Quote OZ is a state of mind, south of the equator.

Digital is currently trying to decide which Win32 API library they will move to on both VMS and OSF1. They are testing both Bristol and Mainsoft.

They are interested in co-operating with us on the decision. Has one been made for us yet and is it disclosable? If we haven't made a decision yet, would it be worthwhile to set up some engineer to engineer contacts?

Regards
David Irwin
Senior Director

From: NEMTS::NEMTS::MRGATE::"WRLMTS::UMC::king@rdaxp.ljo.dec.com" 17-MAR-1995 18:59:59.92
To: HUMAN::CONKLIN
CC:
Subj: Peter - In case you have any comments on Paul McKenzie

From: king@rdaxp.ljo.dec.com@umc@WRLMTS@WRL
To: Peter.Conklin@pko.MTS.dec.com@umc@WRLMTS@WRL

Message-id: 9503171923.AA02959@rdaxp.ljo.dec.com
From: NAME: Mail Delivery Subsystem <MAILER-DAEMON>
Subject: Returned mail: User unknown
Date: 17-Mar-1995
To: king

----- Transcript of session follows -----
While talking to /usr/sbin/mailllv3:
>>> RCPT To:<leiberwirth>
<<< 550-LEIBERWIRTH; %MAIL-E-NOSUCHUSR, no such user LEIBERWIRTH at node LJSRV2
<<< 550
550 LJSRV2::Leiberwirth... User unknown
While talking to /usr/sbin/mailllv3:
>>> RCPT To:<conklin>
<<< 550-CONKLIN; %MAIL-E-NOSUCHUSR, no such user CONKLIN at node XEDON
<<< 550
550 HUMAN::Conklin... User unknown

----- Unsent message follows -----
Received: by rdaxp.ljo.dec.com; id AA02957; Fri, 17 Mar 1995 14:23:34 -0500
Message-Id: <9503171923.AA02957@rdaxp.ljo.dec.com>
To: DELNI::Petrella, LJSRV2::Porter, LJSRV2::Leiberwirth, CRA::Fuller,
CRA::Berard, LJSRV2::Koteff, CRA::Gannon, CRA::Bonney, HUMAN::Conklin,
MSBCS::Dimario, TPSYS::Kelliher
Cc: King
Subject: Inputs to Paul McKenzie's Performance Review
Date: Fri, 17 Mar 95 14:23:33 -0500
From: king
X-Mts: smtp

Greetings,

I am in the process of writing Paul McKenzie's performance review and I would welcome your comments on the performance of Paul and his LJO Computer Support Group over the past 15 months.

Paul is aware that I have requested your inputs. I have been Paul's manager for several years. We are now in the process of transferring Paul's group over to Bill Koteff's ATG Information Infrastructure Project. I will work with Bill on the future plans for Paul and his group so feel free to comment on that also.

I need your inputs by Wednesday, 29-March-95 in order to take them into account.

Please send me mail or give me a call at DTN 226-2680.

Let me know if you have any questions.

Thanks,
Ken King

From: STAR::RMARCELLO "21-Mar-1995 1906" 21-MAR-1995 19:34:43.67
To: @DISTLIST:OVMSORG
CC: @OVMS_SENIOR_STAFF MSBCS::MELLING
Subj: Release Strategy Decision - Eagle/Theta

+---------------------------+ TM DIGITAL CONFIDENTIAL
| | | | | | | | |
| d | i | g | i | t | a | l |
| | | | | | | | | INTEROFFICE MEMO
+---------------------------+

TO: DATE: March 21, 1995
FROM: Rich Marcello
OpenVMS System SW Group
DTN: 381-1471
CC: ENET: STAR::RMARCELLO

SUBJECT: Release Strategy Decision - Eagle/Theta

We have engaged over the last two months in a spirited business and technical debate over release options. The debate focused on the business opportunity and technical capability of shipping cluster enhancements (CIPCA and QLOGIC/FWD SCSI) on a 6.2 base around the 7.0 timeframe in a 6.2-2 release. Last night, Mark Gorham, Steve Noyes, Brian Allison, Mike Cuccia, Tim Ellison, Linda Benson, Bryan Jones, Ray Fusci, Bill Matthews and I met, reviewed the options, and decided that the optimal course for the PM&D and Digital would be to stay with the Plan of Record, ship 7.0 in Q2, and not ship a 6.2-2 release. Specifically, we agreed to the following:

1) Focus the organization on delivering Eagle/Theta in Q2. This is our current plan of record, we want to meet it, and we will need the entire organization's efforts focused on this goal to achieve it.

2) Support post Theta hardware on a 6.2 stream until 7.1.

3) Support CIPCA on RAWHIDE only in the H/W release targetted for Q3FY96.

To summarize the discussion, we focused more on the long term vs. short term business objectives for OpenVMS. While the V6.2-2 proposal offered SCSI and CIPCA adapter support that would have brought in around $30 million a quarter of incremental systems revenue, it was at the cost of either slipping Eagle/Theta by two quarters or significantly limiting the Eagle/Theta qualification and slipping a quarter. Either option would also slip 7.1 from one to two quarters. Neither option was deemed acceptable; the most compelling reasons were:

- 64 bits, threads, and Dollar are key to our future as well as to our "OpenVMS is alive and well" message. Delaying these deliverables and the corresponding positive messages that they convey, would negatively impact our future plans and would unacceptably hurt customer decision making in this area.

- Organizationally, it is very difficult to deliver parallel functional releases. The risk to the organization of having two major releases at the same time and the likelihood of slips to either or both releases, was thought to be very high.

- The risk of slowing LP/application adoption of our key new technologies was thought to be high, if we slipped our delivery dates for the new V7.0 technologies.

- 7.1 is where we believe many of the major applications will start to enjoy significant customer penetration. If we slipped V7.0 by one to two quarters (plus the potential loss of focus from our Digital layered products and ISV applications) then this would have a significant negative impact on long term revenue.

This means we will forgo short term revenue for broad adoption of both CIPCA and QLOGIC/FWD SCSI. This was not a decision we took lightly, but the weight of the negative risks, and the need to keep the OpenVMS futures on track and providing a key positive decision making factor in our target customer's minds drove the decision.

We'd like to recognize the change in the process the orgnization used in this decision. Marketing, business, and technical arguments were considered, both in the debate and the decision. We are moving forward to the PM&D model of business based decisions. Special recognition should go to Brian Allison, who actively and energetically embraced the PM&D model and drove the discussions, to Alan Belancik and Carl Ralston who brought business analysis to the options, to Curt Spacht who ran multiple qualification scenarios, and to Mike Cuccia and Ray Fusci who brought marketing expertise to the debate.

From: STAR::RMARCELLO "21-Mar-1995 2041" 21-MAR-1995 20:42:08.73
To: MSBCS::JACOBS HUMAN::CONKLIN
CC:
Subj: My folks on the Russian proposal - FYI

From: STAR::ZALEWSKI "VMS Development DTN:381-1458" 20-MAR-1995 11:42:58.87
To: STAR::KALER
CC: RMARCELLO
Subj: RE: The Russians

Chris and I agree. The russians are a source of cheap labor, but their track record to date can only be described as spotty. If VMS choses not to invest in the ABI work directly, then outsource the whole thing to Russia for a nominal cost and see what happens.

If we use them for the ABI, then it must be a very carefully segmented piece of work. The loader has to many dependancies, but some of the Win32 APIs could be outsourced.

Any of the utilities they proposed are also sufficiently isolated that we could confidently take a chance and use them....

-steve

From: STAR::KALER "VMS Engineering" 17-MAR-1995 12:49:31.92
To: STAR::RMARCELLO
CC: ZALEWSKI,KALER
Subj: RE: The Russians

I met with them Wednesday. There are very interested in the ABI and wanted to work on the loader. I indicated that our current approach is that we will investigate that work in OpenVMS and that, if we pursue ABI, we would be looking to outsource construction of some APIs which are yet to be determined. They also proposed building a tool to analyze NT applications to determine (dynamically) which APIs are utilized. This could help focus API work if we pursue ABI (or even API). As well, they spoke with the NT group about possibly do some work providing OpenVMS tools on NT (for example, EXCHANGE, BACKUP, MAIL, Runoff/Document converters).

They seemed eager to work with us and willing to start anytime. I think that before we consider any significant outsourcing we should try a small pilot project to make sure the process works and the quality/reliability is within our expectations. As well, I think I should probably visit Moscow a few times to check them out :-).

Is this in line with your thinking Steve?

Chris

digital™ INTEROFFICE MEMO

To: Phil Auberg, Roger Bisbo, Steve Blanchette, Peter Conklin, Mike Cuccia, Steve Jenkins, Rich Marcello, Wes Melling

From: Ken Swanton

DATE: March 21, 1995
FROM: Ken Swanton
DEPT: OpenVMS Systems Business
EXT.: 223-5678
LOC: PKO3-2/T20
ENET: MSBCS::SWANTON

SUBJECT: "My notes from the 3/17 Product Strategy Committee"

**********C O N F I D E N T I A L**********

During Wes' presentation of the strategy/message, I took some notes of some of the major inputs we got. FYI, here they are, in the order in which they came up. (I did not take notes during my presentation of the communications plan).

Mike Gallup said that the OS Strategy was really for the commercial market, and that the technical market is still more UNIX centric.

Vincenzo emphasized that this was a "dream" for the OVMS installed base, but we need to make sure that we (and our ISV's) can deliver it in a reasonable period of time.

Enrico asked if we could pull the engineering off within our budget, and Wes indicated that we would know by 3/24.

Bill Strecker felt that the OS Strategy was more a connectivity strategy.

Tony Craig was uncomfortable with the "unlimited High End" positioning, and recommended instead we position OVMS as solving a class of problems at the high end in a way that is highly accomodating to Windows and Windows NT.

Don Harbert said that for the OVMS installed base we also need to support a path to UNIX for some customers.

Enrico emphasized that we cannot execute without the full support of Microsoft, including jointly launching this strategy with them. He said it was o.k. for Lucia to be leading this from Jesse's team, but he also wants Bill Strecker and Tony working on it closely too. Bill added that he thought that adding ALL-IN-1 to the equation would increase Microsoft's interest, as they would see good business there for their office products.

From: TPSYS::VIFQUAIN "Production Sys SW DTN 227-4461 TAY1 20-Mar-1995 1055" 20-MAR-1995 10:56:40.10
To: A1VAX::JENKINS,HUMAN::CONKLIN
CC: VIFQUAIN
Subj: RE: HP in tp systems

From: BOSEPM::POWELL "PETER POWELL, ENTERPRISE SOFTWARE, DTN 381-1656 31-Jan-1995 1235" 31-JAN-1995 12:32:29.07
To: @[.LISTS]CICS_MKT,@[.LISTS]TP_MKT_ALL,@[.LISTS]TP_INTEREST
CC: POWELL
Subj: I: HPs Mid-engineering Initiative. Encina and CICS with services.

From: DELNI::MAMOS "Dorothy Mamos - NOS Marketing" 31-JAN-1995 11:07:55.50
To: TUXEDO::MCMANN,BRIAN,SUMNER,TUXEDO::LEBLANC,TPSYS::DEPLEDGE,RDGENG::SMITHI, A1VAX::POWELL
CC: MAMOS
Subj: FYI -- article on HP new middleware strategy featuring OLTP

SUBJECT: HP LAUNCHES NEW MIDDLEWARE-ENGINEERING INITIATIVE TO HELP USERS DEPLOY DISTRIBUTED OLTP -- INITIATIVE TARGETS MAINFRAME OLTP USERS -- INCLUDES HP ENCINA/9000 1.1 AND CICS FOR HP 9000 1.3
SOURCE: Business Wire via First! by INDIVIDUAL, Inc.
DATE: January 30, 1995
INDEX: [4]

PALO ALTO, Calif.--(BUSINESS WIRE) via First! -- Jan. 30, 1995-- Hewlett-Packard Company today announced a middleware (mid)-engineering initiative designed to reduce deployment time for business-critical distributed OLTP. The initiative helps users reduce costs through mainframe re-hosting or gain a greater competitive advantage through deploying flexible, mid-engineered solutions. HP's mid-engineering initiative aims to provide a low-risk way for customers to start experiencing payback from an enterprise, client/server approach.

With HP's new initiative, HP brings together its own and third-party product and consulting expertise in middleware software, mainframe downsizing and re-engineering into a focused approach. HP's mid-engineering initiative offers consulting, comprehensive architectural and transition services, transition centers and worldwide support. It helps customers implement distributed OLTP according to their unique business requirements. Existing OLTP-based applications can be integrated into a flexible, open enterprise-computing environment.

HP's mid-engineering approach aims to retain centralized OLTP advantages, such as security, high availability and "central control" while adding greater flexibility for faster time to market and more usable applications, giving customers a competitive advantage.

MID-ENGINEERING

HP's mid-engineering implementation strategy for distributed OLTP provides education, consulting and services through the HP Professional Services Organization (PSO) and key HP Channel Partners. The HP PSO offers a comprehensive suite of open-systems, transition-planning, implementation and education services. A PSO core competency is helping customers move to open, enterprise client/server computing. As part of the PSO's overall architectural approach, specific services help users implement distributed OLTP. These include identifying business roles for flexible information technology (IT), IT architecture planning, core technology evaluation, solution integration, customized education, project management and organizational-change management.

HP has set up a number of PSO-based, distributed OLTP technology enabling groups worldwide. These groups include HP technical consultants who specialize in working with local customers to implement HP Encina/9000 and CICS/9000 locally. Consulting services provide the ability to implement and customize HP Encina/9000 1.1 using various HP technologies that enhance its ability to run business-critical applications.

HP also works with key HP Channel Partners worldwide that have the knowledge and mainframe expertise to provide services and tools to help customers migrate from centralized CICS OLTP to distributed OLTP on HP 9000 systems.

Key systems integrators include Accelerated Solutions, Inc., Circle Group, Ltd.; Denkart n.v.; EMS Ingenieurgesellschaft, mbH; IBS Conversions; Japan Information Engineering Co. Ltd.; and Orbit, Ltd. HP also has formed partnerships with leading ISVs and tools providers such as Haltek; Brixton Systems, Inc.; Compuware Corp.; (CNT); Worldwide Chain Stores Systems Inc.; Open Environment Corp.; and Magna Software Corp. to enable rapid development and deployment of client/server solutions.

HP ENCINA/9000 1.1 & CICS/9000 1.3

HP also announced HP Encina/9000 1.1, whose enhancements include a COBOL API, more robustness and better performance. Only HP's Encina/9000 has improved usability, through a new local and remote configuration facility, and is integrated with HP's SoftBench development environment for greater programmer productivity. HP also provides consulting to more tightly integrate HP Encina/9000 into users' business-critical enterprise client/server environments. HP Encina/9000 also can be integrated with HP's new high-availability solution, MC/ServiceGuard and HP OpenView OperationsCenter.

HP also announced CICS/9000 1.3, an enhanced version of its UNIX(R) system-based CICS solution. CICS/9000 now performs as fast as the fastest CICS alternatives on the market, HP believes. CICS/9000 includes the same features as IBM's latest release of CICS/6000. Additionally, it is more scalable and includes integration with SoftBench from HP to speed up application development and debugging.

U.S. PRICES AND AVAILABILITY

HP also provides consulting to tightly integrate CICS/9000 into users' enterprise client/server environments. CICS/9000 can be integrated with HP's new high-availability product, MC ServiceGuard, to provide a highly available CICS/9000 solution.

Both HP Encina/9000 1.1 and CICS/9000 1.3 are available now. HP Encina/9000 range in price from $150 to $45,000. HP CISC/9000 range in price from $150 to $177,350.

HP is the second-largest computer supplier in the United States, with computer-related revenue of $19.6 billion in its 1994 fiscal year.

Hewlett-Packard Company is a leading global manufacturer of computing, communications, and measurement products and services recognized for excellence in quality and support. HP has 98,400 employees and had revenue of $25 billion in its 1994 fiscal year.

CONTACT: Hewlett-Packard | Lynne Hanson, 408/447-1415 | Sheri Elpern, 408/447-1544

[01-30-95 at 11:04 EST, Business Wire, File: b0130110.500]

From: XIRTLU::yeaton "Tim Yeaton USG Product Management" 20-MAR-1995 17:53:10.32
To: HUMAN::CONKLIN
CC: xirtlu::yeaton
Subj: re: NT positioning

Peter, this is the positioning statement developed by Kate Thompson, from my group, and Ed Muth.

-Tim

------- Forwarded Message

Return-Path: katet
Message-Id: <9503152057.AA21252@xirtlu.zk3.dec.com>
To: agn
Cc: katet, yeaton, rjl, ferson, harbert
Subject: NT positioning
Date: Wed, 15 Mar 95 15:57:22 -0500
From: katet
X-Mts: smtp

Digital's UNIX thrust is focused in specific application-centered market segments ... especially the high performance/more demanding portions of the market.

In the technical area, the target segments are CAD/CAM/CAE, software development, science and research, and GIS/mapping. In the commercial area, the segments are information management, accounting, general business, commercial manufacturing, and communications and networking. These market segments were chosen because they represent growth opportunity for Digital in UNIX and because Digital's product portfolio provides competitive differentiation in these markets. Hard work over the last 24 months has created a strong portfolio of applications, and potent joint marketing relationships with leading vendors, in these key markets.

Microsoft, of course, determines the market positioning of Windows NT. As a practical matter, Windows NT -- on both Alpha and Intel -- is targeted at many of the same markets as Digital's UNIX. For example, there is significant interest in NT-based MCAD and AEC solutions in particular. ECAD, GIS CASE, publishing, financial trading and data analysis (i.e., SAS, Visual Numerics) are areas of overlap as well.

What is different about Windows NT is that the key ISV's tend to be Windows-based applications moving "up", rather than mainframe, OVMS or UNIX applications moving "down".

In general, the most complex applications in each of the technical markets cited are supported primarily on UNIX with entry to mid-range variants of these products supported on NT. This is especially visible in ECAD, where low end design and schematic capture are available on NT but where more sophisticated ECAD work is still largely/exclusively UNIX. Note that Digital has excellent "mix and match" alternative; with products like PCNFS, Pathworks for OSF/1 and eXcursion to provide interoperability and synergy between its UNIX system and Windows, customers can exploit both environments without complexity or duplicate data storage.

Looking at server-centric applications, PC LAN consolidation and Novell server replacement are unique target markets for Windows NT. Document management, workflow, FABS, mail/groupware are also important markets for NT, and in some cases overlap with application availability for UNIX. As was the case with technical applications, however, NT applications tend to be more focused on SME, divisional business problems, and less international enterprises; by contrast Digital's UNIX FABS applications include the most prestigious and sophisticated products in the market, such as SAP.

Both Windows NT and UNIX are well-positioned as commercial servers to the Windows desktop. In the short term (next year or two), UNIX will have the advantage of having a greater array of enterprise-focused commercial applications as well as more robust commercial features and higher performance.

The following are some commercial areas where UNIX may be better positioned for enterprise use than Windows NT:

1) Performance -- Relative to Windows NT, Digital's UNIX has better SMP scaling, larger physical memory support and 64-bit addressing. In addition to these architectural advantages, the 64-bit support for Digital's UNIX by database vendors such as Oracle produces dramatic real-word performance gains relative to any 32-bit operating system.

2) Networking -- UNIX has traditionally had a better networking infrastructure, though NT is quickly closing the gap, with more and more networking capability integrated into the base operating system. Digital's UNIX has a high performing TCP/IP subsystem, built-in NFS support, DNS and NIS name services, and an ATM subsystem. These will continue to be differentiators in the near future for Digital's UNIX. NT (V3.5) supports TCP/IP and will have DHCP and PPP support -- both of which DEC OSF/1 will have in a future release. Many other networking facilities, such as NFS, are beginning to be available from third parties for NT as well.

3) Security -- C2 security is an option to DEC OSF/1 and B1/CMW security is available through the MLS+ product. NT is in the certification process for C2.

4) Availability / Failover -- While the NT Clusters technology demonstration has received deservedly high accolades, Digital UNIX has and will continue to have more robust cluster and availability functionality integrated into the base OS. NT Clusters V1.0 will not ship until FY96Q2.

5) Storage Management -- With its ADVfs file system and Logical Storage Manager products, storage management on UNIX is continues to be stronger than on NT platforms. Third parties are beginning to provide this capability on Windows NT, but the UNIX environment is more mature.

5) Database -- All of the major database products are available on Digital's UNIX (Oracle, Sybase, Informix, and CA/Ingres), and Digital enjoys top-tier status with these partners. Availability and optimization of these products on Alpha NT will lag. In addition, the 64-bit capability of Digital's UNIX allows very large database support, which is critical in many commercial applications and provides an area of unique differentiation.

By contrast, Informix SE and Microsoft SQL Server are the only two major databases currently shipping on Alpha Windows NT, although Oracle, Informix Online and others are expected later in the calendar year.

The following are some areas where UNIX and NT will probably be on par:

1) Management -- UNIX will probably have a wider array of Enterprise or large Workgroup management application

frameworks available from vendors such as Tivoli, Computer Associates, HP's OpenView, and IBM's NetView. Digital's UNIX has or will have IBM's NetView, Tivoli's TME and CA's Unicenter ported. NT, though trailing in initial enterprise management applications, supports Microsoft's SMS management platform as an application for distributed asset management, Netview Polycenter and other substantial management tools.

2) Application availability -- While UNIX currently leads in the number of applications available, NT is gaining ground. In the near term, key commercial applications, like SAP R3, are more mature and are optimized for the UNIX platform.

Today, over 3000 applications are shipping Digital's UNIX, whereas only 1000 applications are shipping on Windows NT for Alpha.

The following are areas where Microsoft NT may have advantages over UNIX:

1) PC Connectivity/NOS Services -- While both Windows NT and Digital's UNIX provides support for TCP/IP and IPX (Digital's UNIX supports IPX through its Pathworks product), NT is designed to provide tight integration with PCs -- through LANmanager, TCP/IP and Novell print and file services.

While Digital's UNIX positioning includes PC integration as a focus, it is within the context of integrating PC desktops with UNIX application servers. PC LAN replacement and up-sizing is not a focus, unlike NT.

2) Availability of traditional PC applications.

NT currently provides emulation for 16-bit DOS and Windows 3.1 applications (excluding 386-enhanced mode). Over 5000 applications run in this emulation subsystem, right out of the shrink-wrap, including many common PC personal productivity tools (dBASE IV, Powerpoint, LOTUS 1-2-3, Visual Basic V3.0 and Corel Draw!, to name just a few). Microsoft Word, Excel, Visual C/C++, and many others are available as native, recompiled Alpha NT applications -- these are not available as native applications on any UNIX platform.

UNIX platforms rely on emulation software (such as Soft-PC or Merge) to run PC applications on UNIX hosts, as does Windows NT on RISC. Performance is generally better on NT as the file system and Windowing system services are native; the AlphaStation 4/166 runs Intel *.exe's at approximately 486/33DX speeds, for example.

Summary:

Digital's UNIX is an outstanding, state-of-the-art UNIX which will delight sophisticated UNIX customers. It enjoys a large and growing application base and is the only major UNIX in the market place which is sharply gaining market share.

Windows NT on Digital Intel or Alpha platforms is an outstanding implementation of Microsoft's high end workstation or server OS. Our large and growing application base, our value added software (DCE, Clusters) and world-wide services make Digital a natural choice for workgroup, site and divisional IS servers ... as well as an interesting alternative for workstations in the entry-level and medium performance range.

Digital's UNIX and OpenVMS will continue to be better positioned for commercial server applications, featuring high availability clustering, advanced storage management and 64-bit database functionality.

Thoughtful customers will find it very attractive to include more than one of these advanced operating systems in their networks, using middleware to interconnect and leverage the unique strengths of each of these platforms. OSF/1 is especially attractive for enterprise-scale data management and large scale scientific analysis and simulation due to its thorough-going exploitation of Alpha 64-bit technology.

------- End of Forwarded Message

----------------- RFC 822 Headers -----------------
Received: by xirtlu.zk3.dec.com; id AA25786; Mon, 20 Mar 1995 17:53:43 -0500
Received: by zuword.zk3.dec.com; id AA24024; Mon, 20 Mar 1995 17:53:30 -0500
Message-Id: <9503202253.AA24024@zuword.zk3.dec.com>
Organization: UNIX Product Line Management & Planning
Phone: (603) 881-0539 or DTN 381-0539
Date: Mon, 20 Mar 95 17:53:29 -0500
Sender: yeaton
X-Mts: smtp

From: VMSMKT::CUCCIA "Mike, OpenVMS Marketing Manager, DTN: 381-0549, ZK03-2/S11 20-Mar-1995 1548 -0500" 20-MAR-1995 15:50:17.47
To: HUMAN::CONKLIN
CC: MSBCS::SWANTON,CUCCIA
Subj: OpenVMS TP Strategy

Peter,

Wedensday's review of the the TP strategy by Vijay Trehan convinced me that we need a new approach to the problem. Below is my understanding of the situation and my suggestion for how to proceed.

Mike

The TP section of the OpenVMS Strategy content is very weak at this time. I believe that we need to make a break with the current thinking in order to position ourselves for a viable long-term position and play down the weak points of our current status.

Current status:

We have very good performance and arguably leading functionality with ACMS, Rdb and RTR today in the high transaction, high availability TP environment on OpenVMS. The issue with this approach is the inability to write (what the industry will believe is) an "open" application. And while ACMSxp is a more open alternative, it is a difficult migration for ACMS users with few tools currently available to assist in the migration. In addition, many do not consider it an open alternative.

To fix that issue, we are currently positioning the availability of CICS/6000 on OpenVMS as the open TP alternative (with the possibility of Encina later). This solutions has several issues of it's own. As I understand it;

- Time to market is 15 months from signing of a contract.
- The CICS/6000 performance is 10's of TPS not 100's like Mainframe CICS.
- During that time IBM will radically reduce the price of Mainframe computes (IBM maintains to UNIX P/P levels??).
- IBM will be in a very competitive position for every sale we target (incumbent vendor plus supplier of Mainframe CICS and CICS/6000 on AIX).
- There will be additional effort and $$s required to get the CICS tools which are not currently being worked, as I understand it.

Vijay Trehan (TP person supplying input to the strategy working group) positioned "traditional" TP Monitor approaches for TP applications as dominating for the next 2 to 3 years. After that, TP technology integrated with object technology will begin to take over.

If the above is true, pursuing CICS and the other open TP monitors as our TP solution will result in our having a product set that arrives too late for the market opportunity, incomplete to address the need and at great disadvantage versus our primary competitors.

Suggested Direction:

Solidify the installed base of ACMS users. We should put the effort into moving the installed base forward with ACMS with the focus on customer satisfaction. We should add the tools needed to move them to ACMSxp to the extent that we believe ACMSxp will emerge as a standard, open solution.

We should provide migration tool support for Mainframe customers who wish to migrate away from CICS (tools such as Conveyor). And invest in ISV tools that have the capability to wrap as objects MF applications and make them available to new C/S implementations.

We should lead the industry in the inclusion of TP technology solutions with Object technology. We should advocate a bottom up approach to TP based on the widespread proliferation of Windows NT servers and LANs. I believe that this is an area of development that Forte is pushing and we should find ways to work with them and any other OO application development vendor to take advantage of OpenVMS in this environment with a small set of key (influential) customers. We would focus our marketing message for general TP solutions on the bottom up OO oriented solutions.

From: STAR::RMARCELLO "18-Mar-1995 1618" 18-MAR-1995 16:19:07.15
To: HUMAN::CONKLIN
CC:
Subj: FYI - Dollar Futures from Jim Johnson

From: OPNDCE::BMATTHEWS "Bill Matthews, OpenVMS Technical Director, ZK03-4/X69 16-Mar-1995 1008" 16-MAR-1995 10:02:35.86
To: STAR::RMARCELLO
CC:
Subj: fyi, dollar future thoughts from Jim

From: STAR::MOVIES::JJOHNSON "Jim Johnson, EDO-13, DTN: 824-3407" 27-JAN-1995 10:26:46.80
To: laing,star::bmatthews
CC:
Subj: File system proposals for NT-VMS integration

Integration of Windows NT and OpenVMS: Data Access and Storage

The following mail describes the goals associated with the Windows NT and OpenVMS integration strategy for file and data access and storage. It then proceeds to list ongoing and future projects that would best help achieve these goals. Items that are either hosted on Windows NT, or items that could be purchased, are called out in that list.

Goal: Digital should remain the source of software and systems for the storage of critical user data.

Therefore, given the Windows NT and OpenVMS interoperability direction, we should provide features that provide:

- effective, transparent, and highly efficient, access to data stored on OpenVMS systems from all likely client sources
- unmatched scaling, reliability, and performance

Finally, as the user's data continues its historic trends towards more decentralization, we should continue to provide the preferred data storage, albeit on the target systems. This should be done with a data migration strategy that offers the least effort, including simple movement of physical media.

Projects:

This section is broken down into a number of project threads, with a small number of cross dependencies. In general the threads can be pursued in parrallel, with less parrallelism within a thread. NB: the time estimates for future projects is, at best, a guess.

DOLLAR: This is the lynchpin technology. Its potentials for scalability, performance, distribution, and heterogeneous support, provide us with an exceptional base from which to work.

V1.0: Initial delivery, Alpha/VMS only, SSB in Fall 95, integrated with Theta.

"ISV Special": This would begin during PEK (FT2) of V1.0, and is targetted to include various storage management APIs that were deferred from V1.0. Future, 6-9 months after V1.0.

Initial NT Port: This project will provide an NTFS compatible client that would be able to communicate with OpenVMS or Windows NT Dollar servers, thus providing high performance access for Dollar files. However, it may make more sense to initially port only the clerk, and defer the native NT Dollar server to a later project.

This project would also begin during PEK of V1.0. Unsized, future, expectation was -12 months after V1.0.

"Knockhill": This is the second general version, targetting increased performance. It would be delivered on Alpha/VMS and on Windows NT, and would incorporate changes from both previous offshoot releases. This is unsized, and a future, with a target release of 12-18 months after V1.0.

Future: Given the work for Knockhill, there are several directions that Dollar could go. Some of these could be done in parrallel, with the dependency that the could not be delivered before Knockhill. Obvious directions include efficient wide area replication, and support for disconnected computing. The latter would most likely need to be pursued in concert with one or more target applications.

PATHWORKS (FSLIB): This is an ongoing port of the Pathworks FSLIB library to Dollar as a special client. It is one of the three clients currently underway at EDO. The first version is expected to be ready coincident with Dollar V1.0, but delivery is dependent on the release plans for Pathworks.

We expect a post-V1.0 version in order to take advantage of expected performance related changes in the privileged Dollar (VPI) interface. That work is unsized, but would be expected to be ready coincident with Knockhill.

This work, combined with the efforts already underway within Pathworks engineering, should provide efficient access to Dollar files for users that prefer to use the standard LanMan protocols.

POSIX (FILE SYSTEM): This is the third native Dollar client currently underway in EDO. It is has the same expectations as the FSLIB work, in that it is expected to be ready for V1.0, and we anticipate a desire to upgrade that support in the V2.0 timeframe to take advantage of VPI enhancements.

RMS: This is the most common path for accessing data on OpenVMS today. However, the on disk data formats used by RMS are rarely amenable to Windows programs. There are two paths that should be followed for this:

- ODBC server for RMS: This would allow database style access to RMS (indexed) files, and make such data immediately and transparently available to many Windows programs, such as Excel or Access. This could be purchased.

- OLE/COM server for RMS: This would map RMS directory trees as "structured storage" and provide automated conversions for the data and metadata in the RMS files. If OLE works the way Microsoft is pushing, this would make entire OpenVMS disk farms available to the Windows user in a very natural way. This project would need further investigation to flesh out. For instance, at the least it would make it possible for Windows OLE client programs to import RMS data files, without interpretation. However, its real power comes to light with the data interpretation. Using system data as an example, such (advanced) support would make is possible and reasonable to import an OpenVMS accounting or MONITOR file into Excel for analysis.

  I can't estimate this one. It needs an investigation period, but it has significant potential.

XQP: This is today's workhorse file system. We anticipate continued maintenance, small performance and Dollar-compatibility enhancements, and little else on OpenVMS.

However, it is an excellent base to begin exploring the potential for zero cost data migrations. This would be a project to be able to mount an XQP disk on Windows NT directly, and be able to access the data on it in a reasonable manner. This would include a local ODBC server for indexed files, a simplified form of the OLE/COM server for the non-indexed files, and direct support for the stream files.

This is an unsized, future project.

From: STAR::RMARCELLO "14-Mar-1995 0956" 14-MAR-1995 09:59:04.98
To: MSBCS::MELLING
CC: @OVMS_SENIOR_STAFF RMARCELLO BMATTHEWS
Subj: OpenVMS Cluster Differentiators - Action item from last staff meeting.

OpenVMS Cluster Differentiators

Info on UN*X from Sandy Snaman about Wave 3 which is near IFT ship in Summer 95. Wave 4 will ship Summer 96. OpenVMS info from Bill Matthews, Andy Goldstein and Bill Laing. Document written by Bill Matthews.

Scaling and High-End Growth
- Most UN*X clusters are 2-4 nodes, OSF/1 hopes to get to 8. OpenVMS even for glass house clusters can easily scale to 32

- Hard to add nodes to a running OSF/1 cluster because SCSI and Memory Channel interconnects not well suited to "hot swap". OpenVMS CI, NI, and FDDI allow incremental growth to a running cluster. Also adding a new node requires building a kernel for each new node for OSF/1 but just requires running cluster_config and booting the new node for OpenVMS. OSF/1 hopes to support FDDI and ATM in future.

- No built-in cluster wide load balancing for OSF/1 but LSF is available and used for workstation farms as compute servers. OpenVMS has load balancing MSCP servering, Batch, Print, LAT, and DECnet built-in plus in some LPs such as PATHworks.

Range of interconnects supported
- OpenVMS supports mixed interconnects of NI, DSSI, CI, FDDI, SCSI, and Memory Channel and ATM in the future. OSF/1 supports SCSI now and Memory Channel with wave 3. Storage uses SCSI, cluster protocols go over memory channel. FDDI or ATM may be possible in wave 4 or wave 5.

- OpenVMS can use multiple interconnects between nodes. OSF/1 will have multi-path via multiple RM paths only.

No Single System Disk for UN*X
- Each system requires it's own system disk with it's own layered products installed on each disk and all the management and account files need to be kept synchronized across all the system disks. OpenVMS has had a single system disk since clusters first shipped.

- The view of the cluster can vary based on the node for OSF/1. Careful cross mounting of NFS disks can make the environment similar. A Cluster file system is in wave 4(summer 96) that will provide a single name space and UN*X single site semantics across a cluster. OpenVMS provides a common environment across nodes in a cluster which enables batch to distribute jobs across a cluster.

No multi-version clusters and no support of rolling upgrades
- OSF/1 would like to support rolling upgrades in the future but it is not planned at this time for wave 3 maybe wave 4 or 5. OpenVMS supports running multiple OS versions and multiple hardware architectures in a cluster for an extended period of time.

Host based shadowing
- OSF/1 supports host based shadowing but only between nodes that share SCSI. No long distance shadowing or shadowing across memory channel is available for OSF/1 and probably won't be added.

Multi-site and disaster tolerance
- No multi-site capabilities from OSF/1 due to RM and SCSI interconnect limitations. SUN has FDDI based solutions that can extend for disks to 2KM.

- It is not likely OSF/1 clusters will be multi-site but multiple clusters may be able to export data to each other in the future to provide remote hot standby capability or with backup/recovery tools.

Batch and Print
- No cluster batch or print for OSF/1. openVMS had Clusterwide Batch and Print since clusters shipped.

Clusterwide Applications
- OSF/1 has one real cluster application, oracle parallel server in wave 3. The cluster file system in wave 4 will be next. OpenVMS has Rdb, DECdtm, PATHworks, Batch, file system, CMS, ACMS, etc. clusterwide applications.

Clusterwide process services
- OSF/1 doesn't appear to have the ability to start and stop and monitor processes on other nodes in a cluster like OpenVMS

- OSF/1 is hoping to have process migration in a cluster, checkpoint/restart etc. in the future.

OSF/1 will have a clusterwide file system, distributed lock mgr, and each service such as NFS is responsible for providing cluster aliasing. Management of the cluster as a single system will begin with wave 4 but may not include a single system disk.

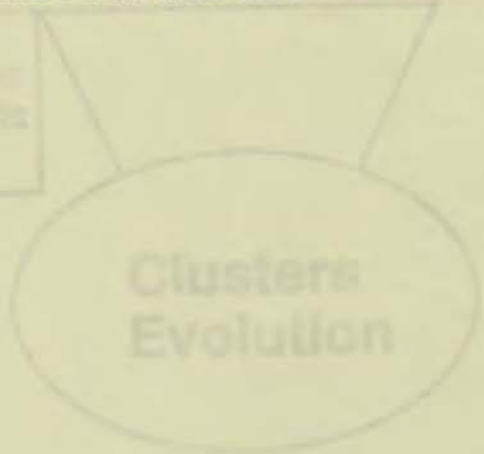

OpenVMS and Clusters Evolution

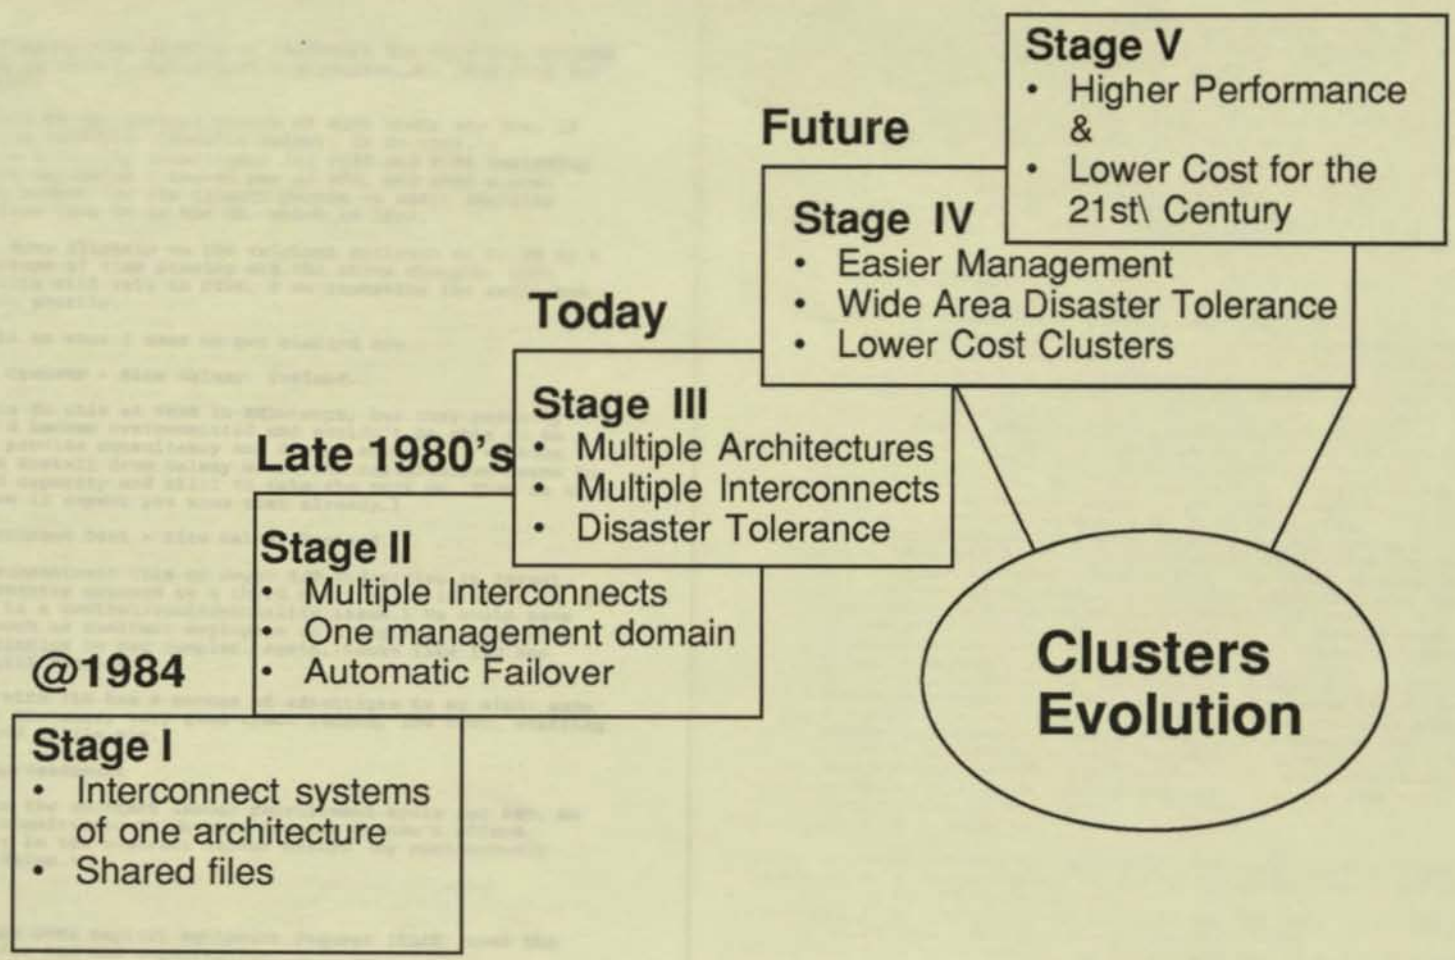

From: IOSG::IOSG::MRGATE::"A1::MEREWOOD" 14-MAR-1995 11:22:05.61
To: LJSRV2::JENKINS,HUMAN::CONKLIN
CC: AOSG::ANGEL
Subj: Funding Request for CICS OpenVMS Project

From: NAME: Richard Merewood
FUNC: Manager Ivanhoe
TEL: DTN 830-3352 <MEREWOOD@A1@IOSG@REO>
To: NAME: Steve Jenkins <Jenkins@LJSRV2@MRGATE>,
NAME: Peter Conklin <Conklin@Human@MRGATE>
CC: NAME: Kathy Jensen <JENSEN@A1@IOSG@REO>,
NAME: Dick Angel <Angel@AOSG@MRGATE>

Steve,

I've ascertained that no FY95 funding or headcount for CICS for OpenVMS has been allocated in Gerard Van de Aast's organisation. Only CICS for OSF/1 work is funded.

It's essential that we get certain pieces of work under way now, if we're to keep to the schedule. (details below). To do that, I need you to provide a funding committment for FY95 and FY96 according to the project cost estimates I showed you at ZKO, and then a cost centre and project number for the finance people to start charging against (charges from this CC in the UK, which is LY9).

FY95 costs may be down slightly on the original estimate of $1.3M by a few hundred $K because of time passing and the above changes. Some small portion of this will roll to FY96. I am reworking the estimates and will update you shortly.

FYI, here's details on what I want to get started now.

1. Encina Port to OpenVMS - Site Galway, Ireland.

   We were going to do this at VMSE in Edinburgh, but they recently indicated they'd become overcommitted and wouldn't be able to do much more than provide consultancy and design support. I've been talking to Alan Rockall from Galway where the organisation seems to have the needed capacity and skill to take the work on. They're also cost competetive (I expect you know that already.)

2. CICS Test Environment Port - Site Galway Ireland.

   We planned to subcontract this to Negev S/W Industries in Israel, but IBM are adamantly opposed to a third party being involved in the project. (This is a control/confidentiality issue.) We could have NSI employees work as contract employees on a Digital Israel site but this is beginning to get complex. Again, looks like ILO has capacity and skill required.

   (Collaboration with ILO has a number of advantages to my mind: same timezone, short distance, very good track record, low cost, staffing flexibility, good skills mix.)

3. Reading Staffing/Headcount

   We need to begin the contract labour recruitment cycle for REO. We need a funding committment to do this because we can't afford credibility loss in the contract labour market (by continuously slipping start dates.)

4. Capital

   I've inserted the OVMS capital equipment request ($54K) into the approval cycle for the EOS organisation.

Thanks,

Richard.

From: LJSRV1::LIEBERWIRTH "dtn 226-2587, 508-486-2587" 11-MAR-1995 19:39:10.89
To: STAR::BMATTHEWS, HUMAN::CONKLIN, AOSG::FERSON, DECWET::PENNEY
CC:
Subj: transaction mgr as 3tier C/S software mediator? interesting idea.

From: US1RMC::"ngreenfe@world.std.com" "Norton R Greenfeld" 10-MAR-1995 15:29:40.47
To: distribution::@uslrmc.bb.dec.com (see end of body)
CC:
Subj: ITbits 3/10/95

ITbits is published by Implements, Inc. It contains personal opinion and does not necessarily represent the official position of anyone. It can be redistributed at will, but please include an attribution.

Norton Greenfeld, Implements, Inc.
6 Brook Trail Road, Wayland, MA 01778
tel. 508-358-5858 fax. 508-358-5846
e-mail. ngreenfeld@world.std.com

3/10/95

1. Next Generation Client/Server

PCWeek this week has a set of articles all relating to the perils and problems of client/server. And this is only the latest in a new view of client/server appearing in PC-oriented magazines (the DP/MIS mags have warned of c/s for some time). Is this the end of client/server as a direction?

Of course not -- but it does signal a growing awareness of the need for change. First generation client/server was invariably what is now called "two-tier": a PC doing almost all the work and a database server doing the rest (at best, the database server handled "stored procedures" that allowed some of the processing to happen on the server). These days there is a growing buzz for "three-tier" c/s and "partitioning". This means that some application logic and processing happen on a server, and the PC mostly does the user interaction and perhaps a little business logic.

So far so good: makes some sense to have the servers do processing (they are closer to the data and this may have the effect of sending less data on the network), they tend to be faster, and most important, programs on the server can be controlled by MIS. BUT, how do we build such 3-tier systems? Into this void have stepped companies such as Forte, Dynasty, some of the 4GL vendors (Unify, Informix's NewEra, JYACC, etc.) and some of the CASE vendors (such as TI's Composer).

But these tools are still missing pieces. Having just gotten back from a Novell Tuxedo conference, I can say I've seen the future. Consider a layer of software whose job is to manage applications. That is, allow for the dynamic partitioning of process to server, automatic load-balancing, failover protection, true location-independent services, application-level measurement and tuning, ... all in a reliable, secure environment.

Anyone say "transaction manager"?

Novell has figured out that its Tuxedo transaction manager, besides doing the old-concept stuff, really IS a good backbone for large-scale client/server. It will finally begin pushing Tuxedo in that direction next month. And I believe it really is a good fit.

What are the negatives? Complexity, for one. The message-queueing middleware is hot at the moment, primarily because developers can grasp its essentials in about 5 minutes. Tuxedo includes the same messaging capability, but as one of 5 or 6 types of architectures available. And Tuxedo assumes that a company has an "application administrator" who will make sure the system in running, tuned, etc.

Price, for another. To become the ubiquitous "application manager" the cost of an embedded transaction manager cannot be a hindrence. I'm not sure it's there yet.

The biggest negative? image. Novell has to educate the computing world that a transaction manager really does more than bank ATMs, that it is a general solution with lots of benefits. This is the path that Novell is embarking on now.

And the competition? It looks like CICS is still aimed at keeping those old mainframe CICS applications in the fold, that Encina needs more time to mature and improve performance, that TopEnd has a chance but may not be marketed broadly and well. And CORBA, DCE and the like? They seem to have some share of the market, but also seem much too complicated for most users' taste.

So the main competitors seem to be Tuxedo, message-queueing software from IBM and DEC, and probably a lot of smaller players. In a $1B software market!

Final thought: any $1B software market must have Microsoft in it, and there are already rumors of a Microsoft transaction monitor in the works. If Microsoft is interested, does anyone doubt this is the future?

%%% overflow headers %%%
To: ITbits distribution list <73060.3352@compuserve.com>, abinstock@mfi.com, ajv@summit.novell.com, barry.needleman@progress.com, beckett@zk3.dec.com, beier@ilog.com, belanger@hpcgrp.ENET.dec.com, bill@oberon.com, bmd@littlei.jf.intel.com, chip.hay@documentum.com, chrisk@vmark.com, cstetson@mail.zd.ziff.com, darcy_fowkes@pyramid.com, deand@sequent.com, diane@utg.org, eileen@aventure.com, ellyn@mks.com, esc@summit.novell.com, fglover@zk3.dec.com, fox@carrol.ENET.dec.com, godfrind@heron.ENET.dec.com, grochmal@msbcs.ENET.dec.com, hhindin@etg.com, hideya@ccgate.sj.nec.com, howard.shao@documentum.com, ilser@vnet.ibm.com, james.e.clark@columbiasc.ncr.com, jgelhausen@aol.com, jillc@sj.unisys.com, jim@sybase.com, jim_christensen@hp6650.desk.hp.com, jkeough@mfi.com, john_mccammond@vos.stratus.com, jvittal@gte.com, Karen_Logsdon@next.com, karl_stoltze@phx.mcd.mot.com, kathy_knaack@vnet.ibm.com, kevin.zondlak@daytonoh.ncr.com, kimberly_daniel@hp6650.desk.hp.com, kreott@sj.unisys.com, lee.sigler@corp.sun.com, lieberwirth@ljsrv1.ENET.dec.com, linda@parcplace.com, linda_hargrove@netpower.com, lytle@summit.novell.com, medortch@aol.com, menard@summit.novell.com, miket@sandshark.sandiegoca.ncr.com, Mitch_Shults@ccm.fm.intel.com, mitman.enrique@al.mroa.mrmts.MTS, mjm@unify.com, mkt-infomail@amdahl.com, mktgl@amail.amdahl.com, mstrainor@cw.com, nwoodward@btrvtech.com, P.M.Slavid@slh0105.wins.icl.co.uk, r.schwark@att.com, renderle@dataquest.com, rickw@cstoday.com, sabine.turnley@columbiasc.ncr.com, sasson@oberon.com, slapolla.notes@mail.zd.ziff.com, spool.chi@xerox.com, tkalil@gps.com, torti@wecare.ENET.dec.com, vdow@mfi.com, warshawsky@lando.ENET.dec.com, weed@decvax.dec.com, wise_olga@tandem.com, ybarra@adobe.com, yoshi@ccgate.sj.nec.com
%%% end overflow headers %%%

% ------ Internet headers and postmarks (see DECWRL::GATEWAY.DOC) ------
% Received: from mail1.digital.com by uslrmc.bb.dec.com (5.65/rmc-22feb94) id AA09382; Fri, 10 Mar 95 15:29:56 -050
% Received: from europe.std.com by mail1.digital.com; (5.65 EXP 2/22/95 for V3.2/1.0/WV) id AA29841; Fri, 10 Mar 1995 11:56:36 -080
% Received: from world.std.com by europe.std.com (8.6.8.1/Spike-8-1.0) id OAA28564; Fri, 10 Mar 1995 14:38:48 -050

```
% Received: by world.std.com (5.65c/Spike-2.0) id AA05515; Fri, 10 Mar 1995 14:3
8:52 -050
% Date: Fri, 10 Mar 1995 14:38:51 +0001 (EST)
% From: Norton R Greenfeld <ngreenfe@world.std.com>
% Subject: ITbits 3/10/95
% To: distribution:;@uslrmc.bb.dec.com (see end of body)
% Message-Id: <Pine.3.89.9503101442.A4446-0100000@world.std.com>
% Mime-Version: 1.0
% Content-Type: TEXT/PLAIN; charset-US-ASCII
```

From: MSBCS::SMITH_A "17-Mar-1995 0556 -0500" 17-MAR-1995 05:56:14.63
To: human::conklin
CC:
Subj: *UA* Software Strategy Updates 1

I N T E R O F F I C E M E M O R A N D U M

Date: 17-Mar-1995 05:34am EDT
From: Al Smith
SMITH.AL
Dept:
Tel No:

TO: See Below

Subject: *UA* Software Strategy Updates

The attached note explains an urgent request from the Software Business Group and the Systems Business Group senior management teams.

Please distribute the attachment to your product management staff and have them submit their software strategy updates by March 24, 1995. If you foresee any problems meeting the deadline please notify Jeanette Horan (msbcs::horan) and Dave Ellison (msel::ellison).

Regards, Al

Distribution:

TO: Remote Addressee (_msel::ellison)
TO: Remote Addressee (_aosg::hebenstreit)
TO: Remote Addressee (_sdtpmm::holland)
TO: Remote Addressee (_alvax::jenkins)
TO: Remote Addressee (_delni::kenah)
TO: Remote Addressee (_took::kinsman)
TO: Remote Addressee (_sdtpmm::sartini)
TO: Remote Addressee (_took::saunders)
TO: Remote Addressee (_mkotsl::seif)
TO: Remote Addressee (_star::slane)
TO: Remote Addressee (_alvax::stillings)

CC: Remote Addressee (_msbcs::demmer)
CC: Remote Addressee (_thundr::lipcon)
CC: Remote Addressee (_msbcs::strecker)
CC: Remote Addressee (_human::conklin)
CC: Remote Addressee (_msbcs::horan)
CC: Remote Addressee (_chfv03::muth)
CC: Remote Addressee (_delni::proulx)
CC: Remote Addressee (_alvax::vandeaast)
CC: Remote Addressee (_xirtlu::yeaton)

SOFTWARE STRATEGIES DOCUMENTATION

The senior management team of the Software Business Group and the Systems Business Group have requested that the VTX IR software strategies database be updated immediately with the latest official strategies of record and schedule of products.

This task is part of an overall effort to document and communicate our existing software strategies. Status of this effort will be reviewed with Bill Demmer and Jessie Lipcon.

Please update your product(s) in the VTX IR software strategies database to reflect a current strategy overview and a schedule of products for at least the next four quarters. Remember that these submissions will be used as the official software strategy of record. They should be suitable for external communication and should not announce potential or pending business partnerships.

To submit your update pull the document(s) that contains your product(s) and insert your changes with change bars to the left of affected paragraphs. Once you are satisfied your product strategy is documented correctly, send the entire document including your updates to msel::ellison and msbcs::smith_a Email accounts.

Your submissions will be entered immediately into the VTX IR database and the internal worldwide software strategies server on the WEB. If you need help with the update process please contact Dave Ellison at msel::ellison or by phone at 227-3621.

Regards, Al

From: BOSEPM::JENKINS "STEVE JENKINS, DTN 381-1704, COMMERCIAL SOFTWARE PRODUCTS SEGMENT MGR 17-Mar-1995 0856" 17-MAR-1995 08:51:27.04
To: WES_MELLING,PETER_CONKLIN
CC:
Subj: fyi ... SW Strategy Action Items

From: MSBCS::DEMMER "Bill Demmer DTN: 223-1400 PKO3-1/A8" 15-MAR-1995 13:58:16.24
To: CURTIN,HORAN,A1VAX::JENKINS,JEAN,SMITH_A,STRECKER,A1VAX::VANDEAAST
CC: THUNDR::LIPCON,PNDVUE::STRECKER,DEMMER
Subj: Software Strategy Task Assignment

****
D I G I T A L R E S T R I C T E D D I S T R I B U T I O N
****

+----------+ TM
| d i g i t a l |
+----------+

I N T E R O F F I C E M E M O

From: Bill Demmer
Date: March 15, 1995
Dept: Software Business Group
Ext: 223-1400
Loc: PKO3-1/A8
ENET: MSBCS::DEMMER

To: Paul Curtin
Jeanette Horan
Steve Jenkins
Jean Proulx
Al Smith
Nancy Strecker
Gerard Van de Aast

cc: Jesse Lipcon
Bill Strecker

Subject: Software Strategy Task Assignment

Bill Strecker has asked that we undertake the following efforts in support of the joint CSD-ATG look at our future Software Strategy, (I have added my suggestions on who should be doing what based upon our Software business Group staff meeting discussion):

1. The three Segment Managers (Gerard, Jean, Steve) should categorize each product set against the following criteria:

   Invest, Harvest, Partner, Divest

2. Also, indicate if you see each set as supporting our Platform Business or our Connectivity Business to the degree you can define and differentiate them.

3. The Segment Managers should then review the Partner-Divest product set list with Paul Curtin to identify the following:

   a) The Process Owner
   b) The Potential Partner List
   c) The Timeframe
   d) The Probability of Success, and Backup Plans

4. That Al Smith update the Software Product Repository showing our external strategy and schedules for each product set. This requires immediate support from your Product Managers.

5. That Nancy Strecker work with the 3 Segment Managers to update the Software Strategy presentation in preparation for review with a Tony Craig led Message Validation and Communication team. This should be aimed at both an internal Digital and an external audience.

6. There are at least 4 product sets that are or could be undergoing a strategy change for their future. Special care must be taken on how these are written up for this internal review by the Steering Committee (B. Strecker, E. Pesatori, T. Craig, C. Bennett). The product sets in question are:

   Linkworks
   DBI
   ACMS/ACMSxp
   Document Management

   All of these have dependencies on other companies ("Partners") and need to reflect today's status as well as the strategic goals involved.

7. In order to be able to accomplish the above by the first week in April, I have asked Jeanette Horan to coordinate the above activities, to brief Jesse and I weekly on the progress, and to be prepared to give a status report to the Strategy Steering Committee on March 27.

On a personal note, I am very concerned about putting down on paper words like Partner/Divest, because if these are "leaked" outside of this Strategy Core Team, general marketplace havoc will result. Please keep this in mind as you work to pull this together.

From: OPNDCE::BMATTHEWS "Bill Matthews, OpenVMS Technical Director, ZK03-4/X69
17-Mar-1995 0800" 17-MAR-1995 07:53:10.93
To: star::rmarcello,human::conklin,star::gorham
CC:
Subj: FWD: It looks like someone may have created a symbiont already 8-)

From: STAR::SYSMGT::DUTKO "In Win32, no one can hear you scream..." 1
6-MAR-1995 17:03:32.22
To: STAR::BMATTHEWS
CC:
Subj: It looks like someone may have created a symbiont already 8-)

From: COPCLU::BJARNER "Bjarne Rasmussen - DESKTOP" 16-MAR-1995 16:21:28.5
8
To: SYSMGT::DUTKO
CC: BJARNER
Subj: VM: REF: Printing requirements from OpenVMS

I N T E R N M E D D E L E L S E

Dato: 16-mar-1995 22:18 CET
Fra: Bjarne Rasmussen-Edu @DMO
RASMUSSEN-EDU.BJARNE
Afdeling: DK SI Groupware & Client Serve
r
Tlf: 857 - 2179
Dokument:

TIL: Remote Addressee (_sysmgt::dutko)

Emne: VM: REF: Printing requirements from OpenVMS

I N T E R N M E D D E L E L S E

Dato: 16-mar-1995 21:58 CET
Fra: Bjarne Rasmussen-Edu @DMO
RASMUSSEN-EDU.BJARNE
Afdeling: DK SI Groupware & Client Serve
r
Tlf: 857 - 2179

TIL: MACOMBER (MACOMBER@HANNAH@COPCLU@MRGATE@COPMH
S@DMO)
TIL: Remote Addressee (_sysmgt::dukto)
TIL: Remote Addressee (_hannah::landau)
TIL: Michael Ovesen @DMO (OVESEN.MICHAEL)

CC: Bjarne Rasmussen-Edu @DMO (RASMUSSEN-EDU.BJARNE)

Emne: REF: Printing requirements from OpenVMS

Hello

My name is Bjarne Rasmussen, Consultant in System Integration in Denmark.

Michael Ovesen, also SI Denmark, and I have developed a print solution, that allow OpenVMS, DEC OSF/1 and ULTRIX users to print to DOS, Windows, NT and OS/2 by use of DECnet or TCP/IP.

On OpenVMS the solution is a symbiont. On OSF/1 and ULTRIX, the solution is a filter.

A table control what should happend to the print, once it is received on the PC. The print can be printet locally or redirectet to Novell, LAN Manager and Banyan servers.

All sysmgt is done from OpenVMS, OSF/1 and ULTRIX.

We have meet a VERY BIG interest for the solution. We have sold the solution to customers in UK, Holland, France, Norway, Germany and Denmark. It is primary lage customers that have a need for printing from central systems to PC-lans over a WAN, Token Ring, FDDI and so on.

Personally I beleive in LPD/LPR, but this protocol does not support redicetion automatically to lan servers. Our solution does. We have also build in a "advanced" logging system.

I have presented the solution for PATHWORKS engineering. The have not so far been interested in the solution. In the mean time, we sell the solution as a Danish asset.

You can find kit and doc on
COPCLU::DUA13:[BJARNER.EEN.PWPRINT...]

Please respond if we can help,

Bjarne

From: MSBCS::LIPCON "13-Mar-1995 0857 -0500" 13-MAR-1995 08:57:40.95
To: MSBCS::SWANTON
CC:
Subj: OpenVMS Position memo as requested (from Jesse to John Okeefe etc.

+---------------------------+ TM
d	i	g	i	t	a	l
+---------------------------+

INTEROFFICE MEMORANDUM

TO: Ken Swanton
Mike Cuccia
Ken Steinhardt

CC: John O'Keefe
Lisa Bender
David Flawn

DATE: 08/19/94
FROM: Jesse Lipcon
DEPT: OpenVMS Systems
EXT: 226-2767
LOC/MAIL STOP: LJO2/F4

SUBJECT: More on the OS positioning discussion

I met on Tuesday with John O'Keefe, David Flawn, and Lisa Bender, to continue our operating system positioning discussions and, more specifically, to expand upon and flesh out (with "do say, don't say" guidelines) how we can be bold and aggressive about OpenVMS in the business critical server area without undercutting Digital's commitment to UNIX or the ability of our OSF/1 UNIX implementation to also be sold in this area. Of course, we are trying to walk a fine line here. We cannot afford to either pull our punches with OpenVMS or undermine our UNIX efforts. Thus we have approached these discussions with a spirit of compromise.

In our previous meeting, we had settled on the following three ways to say "OpenVMS is the world's best business critical server" without actually having Digital use those words:

- OpenVMS has the "world's best" or "industry leading" clustering

- OpenVMS functionality "greater than that of mainframes"

- Consultants'/customers' testimonials say "OpenVMS is world's best.."

Ken and Mike, you have pointed out that these statements, as such, seem overly narrow and constraining, and that in the OpenVMS "pitch" even Ken Steinhardt and I are unable to live up to the letter, or even the spirit of these guidelines. Thus we have attempted to broaden these guidelines (and John has agreed) as follows:

- It's OK to use superlatives, such as "most business critical", "absolutely, positively has to keep running", "when the computing is down, you're out of business", etc. in describing the customer environment or customer requirements that we are targeting with OpenVMS.

- It's OK to cite OpenVMS as having the "world's best" or the "industry-leading" clustering functionality, and to tie this to the high availability, data integrity, and scalability required in those "most business critical" environments described above. This is backed up, of course, by the Aberdeen report and Gartner quotes. We will not explicitly say OpenVMS is the "world's best OS for business critical applications".

- It's OK to discuss being on the world's acknowledged leadership platform in both absolute performance and price-performance. This Alpha attribute, of course, is applicable to all three strategic operating systems.

- It's OK to discuss OpenVMS specifically on that platform as having the

world's best transaction processing and commercial workload performance and price-performance" since it is proven by the hard TPC-A facts: 3692 TPS-A absolute performance record and the best FIVE spots in $/TPS-A.

- It's OK to have analysts and customers use any superlatives they choose in their direct quotes about OpenVMS. However, John and I agreed that we should try to avoid having analysts make comparisons between OpenVMS and OSF/1, and in any quotes which do mention both, to have both groups OK the quote before using it. We also agreed to review quotes with each other in order to ensure that they are not based on inaccurate data.

- It's OK to say that, based on all the above, "OpenVMS is ideal for business critical appliations, or for the most business critical applications", BUT avoid saying "THE ideal".

John has also agreed to stop saying, and delete from OSF/1 presentations and collateral material, any references to OSF/1 "catching up" to OpenVMS cluster capability, or "achieving full OpenVMS clusters". He would welcome our pointing out any cases where such statements continue to "sneak through".

We have also agreed to share presentations and collateral with each other in order to ensure that both groups are following the spirit of these guidelines. (This is already happening today.)

Having reached these conclusions, I then went through the latest (16-August) version of Ken Steinhardt's OpenVMS white paper, in order to test whether the above guidelines would be over-constraining, or would force major changes. In the 11-page document, I found only three lines which were even questionable under the guidelines. Ken and I discussed these, and in each case worked out minor wording changes. The net effect of these changes not only brought the document "into compliance" but, in Ken's estimation, actually strengthened the message! Thus I feel we have come a long way toward our goal of finding a way to position OpenVMS without pulling our punches, and yet not undermine our UNIX efforts. I would welcome your further comments.

(f) (A) others

From: OPNDCE::BMATTHEWS "Bill Matthews, OpenVMS Technical Director, ZK03-4/X69
03-Feb-1995 1601" 3-FEB-1995 15:57:41.72
To: HUMAN::CONKLIN,STAR::RMARCELLO
CC:
Subj: Proposed next step actions and DRIs for OpenVMS and Networks

OpenVMS Strategy Next Steps

o The cluster and shadowing team are continuing CI and a QFD on availability throughout the months of Febuary and March. Susan Azibert is driving the effort. Customer needs weighting will occur on Feb 10, Need and Solutions Correlation will occur on March 17th.

o The Windows ABI investigation report is being driven by Steve Zalewski with a target due date of Feb. 23.

o An initial investigation on remote management and support for lights out computing listing the current capabilities and future needs will be driven by Bill Matthews and Nestor Dutko with input expected from Robert Hoffman, Amy Becker, and the POLYCENTER group with a due date of Feb 17.

o An initial investigation on Printer management and access of OpenVMS printers and LAN printers listing current capabilities and future needs will be driven by Nestor Dutko with input expected from the PATHWORKs and POLYCENTER groups with a due date of Feb. 17.

o An initial investigation on Backup of PC files from OpenVMS listing current capabilities and future needs will be driven by Bill Matthews with input expected from Phil Wells, Peter George, Nestor Dutko, Jim Johnson, and Peter Lieberwirth with a due date of Feb. 17.

o The initial investigation of placing an NTFS compatible log-structured file system client on Windows NT that could communicate with an OpenVMS or Windows NT server so customer data could physically migrate to NT when necessary was completed. Rough scoping of the cost will be targetted at Feb. 23 by Jim Johnson.

o The initial investigation of NT Cluster protocols with PATHWORKs so windows clients see OpenVMS cluster files(shares) like NT cluster shares. Bill Matthews will drive with expected input from Peter Lieberwirth and Phil Wells for a due date of Feb 23.

o Continue the investigation of PATHWORKs, TCP/IP, and DECnet to ensure no overlap and good integrattion. Add in IBM interconnect and wide area networking to the investigation as the advantages. Steve Stebulis will drive this area and report back on Feb 17

R: jser. R.D.

From: STAR::GORHAM "15-Feb-1995 1257" 15-FEB-1995 12:56:04.02
To: @JESSE
CC: @MAGSTAFF,DELNI::S_FLYNN
Subj: OVMS System Forecast proposal

+---------- TM

| d | i | g | i | t | a | l |

DIGITAL CONFIDENTIAL

INTEROFFICE MEMO

TO: Jesse Staff
Gorham Staff

DATE: February 15, 1995
FROM: Mark Gorham
DEPT: OpenVMS System Software Group
Business Management
LOC: ZKO3-4/T61
DTN: 381-0134
ENET: STAR::GORHAM

CC:

SUBJECT: OpenVMS Systems long range forecast - proposal

The following OpenVMS five year system forecast has been developed by Susan Slane and myself in response to planning requests by several of our key internal partners. It is our goal that OSBG can agree on an OpenVMS systems forecast and distribute it as a planning tool for use by all of our partners' FY96 LRPs. We hope Steve Blanchette's organization can add their expertise to this forecast, and have identified several areas this expertise could be used.

As a start, we used three major forecasting techniques:

Overall IDC market trending based on system size
Installed base/upgrade projections using actuals, Dataquest, and the CI study
New business from mainframe downcosting.

1) Overall IDC Market Trending Based on System Size.

We based this piece of the forecast on the IDC five year market trend based on system size. IDC splits systems by large-scale systems (over $1M), medium-scale systems ($100K to $1M), and small-scale systems ($10K to $100K, excludes PCs and workstations). Note, when IDC refers to "mid-range" systems, they mean all medium-scale and small-scale systems). Both revenue and unit forecasts were obtained. Mid-range systems are projected to show a slight growth in units (about 4% CAGR, '94-'98), and relatively flat revenue growth (less than 2% CAGR). Worldwide units of large-scale systems (>$1M) are projected to grow at a CAGR of about 16%, with total revenue down about 4%.

IDC REVENUE FORECAST ($B)

	94	95	96	97	98
Large	12.3	11.5	10.8	10.1	9.4
Medium	10.8	10.9	11	11.2	11.3
Small	14.2	14.7	15.3	15.8	16.4
WS	12.4	13.6	14.8	16	17.2

IDC UNITS, WW SYSTEM SALES (000)

	94	95	96	97	98
Large	2	3	3	5	5
Medium	34	36	38.	40	42
Small	496	496	514	532	556
Total Multi-user	512	535	556	557	603

OpenVMS SYSTEM UNIT FORECAST, based on constant market share of IDC's figure of 6%, FY94 actuals as a baseline

High - 6000s/ 7000s, Mid - 4000/2100s, Low - 2000/1000/3000s

	94	95	96	97	98
High	1811	1945	2047	2166	2279
Mid	9387	10083	10613	11225	11810
Low	16281	16281	16872	17463	18250
WS	28263	30998	33733	36468	39204
Total	55742	59308	63265	67322	71543

The workstation line is one where Steve's organization can add value, verifying or modifying the growth.

*Note - this is constant market share - the IDC December 94 report said OpenVMS would INCREASE from 6% of the midrange market to 7% by 1998 based on successful migration to Alpha.

1% Increase in Market Share (per IDC)

High	1811	2026	2218	2436	2658
Mid	9387	10504	11497	12629	13779
Low	16281	16959	18278	19646	21292
WS	28263	30998	33733	36468	39204
Total	55742	60487	65726	71179	76933

2) Installed base/Upgrade Projections Using Actuals, Dataquest, and the CI Study

The size of the installed base has been a subject of debate. OpenVMS systems sold by Digital through YTD95 exceed 700,000 units. There are large numbers of gray market machines that exist in Eastern Europe (and apparently China); we have not taken these into account in this analysis. Of the systems sold by Digital since 1978, how many are scrap/retired? Dataquest analysis indicates that the installed OpenVMS server base is 307,000 units. We assumed ALL systems sold before FY87 were retired, and then 20% of the systems since then were retired to come up with our Installed Base units.

Once we determined the installed OpenVMS unit base, we need to determine how many will migrate to OpenVMS Alpha. The December CI InfoCorp study showed that fully 80% of the OpenVMS installed base plans to migrate to OpenVMS Alpha, up from 75% last year. Our analysis lowered the workstation percent and raised the high end based on recent history.

80% will migrate, but when will they migrate? Using an accounting-based analysis, we determined the lifecycle for each class of machine, basically 5 or 6 years, and forecasted the migration business through FY98:

System % planning

	IB	94 act	95	96	97	98	Lifecycle	OVMS migr
High	23000	1811	3450	3450	3450	3450	6 yrs	90%
Med	40000	9387	6800	6800	6800	6800	5 yrs	85%
Low	160000	16281	27200	27200	27200	27200	5 yrs	85%
WS	212000	28263	21200	21200	21200	21200	4 yrs	40%
Total	435000	55742	58745	58746	58747	58748		

Note that we have not split VAX and Alpha in this analysis, instead leaving the analysis to the experts in Steve's organization.

3) New business from Mainframe Downcosting

The mainframe downcosting program generated $500M last year for OpenVMS systems. We assumed a successful program would generate incremental units over FY94 sales, growing as the program gains momentum.

	94 act	94 ASV	94 Business	% inc	Incremental Units FY95	96	97
High	1000	$250K	$250M	10%	100	110	121
Mid	1200	$125K	$150M	20%	240	288	346
Low	1286	$ 70K	$ 90M	10%	129	141	156
WS	400	$ 25K	$ 10M	5%	20	21	22
Total			$500M		489	560	645

PROPOSED OVERALL FORECAST FOR OpenVMS SYSTEMS

This proposed forecast averages out the 1% growth and installed base numbers, both closely related as the IDC 1% growth number assumes a successful IB migration, and then adds in the incremental mainframe downcosting numbers:

OpenVMS System Unit Forecast

	94 Act	95	96	97	98
High	1811	2838	2944	3064	3187
Mid	9387	8892	9436	10060	10704
Low	16281	22208	22880	23578	24417
WS	28263	26119	27488	28856	30225
Total	55742	60057	62749	65559	68533

Two other areas remain where Steve's organization can add value. The first is an analysis of the impact the shift of our customer base from technical to commercial will have on our units, based on the overall market shift. The second is the NT synergy opportunity impact. Both could lead to significant increases in our server sales.

Susan and I would be happy to discuss this at Jesse's staff or work it separately with Steve/Norma/Karen/whoever. I'll follow up with Steve to find out your preference.

Regards,

Mark and Susan

From: HUMAN::CONKLIN "Peter 508.493.5648 PKO3-3/T20, sec Dory .5565 23-Jan-1995 0842" 23-JAN-1995 09:06:12.93
To: MTS$::"ZKO::NORMA ABEL",XANADU::HESS,STEVE.JENKINS,IOSG::PILGRIM,MARVIN::TURNER
CC: CONKLIN
Subj: Questions about our Groupware strategy

DIGITAL CONFIDENTIAL

At the end of our meeting last week, I said I would write down the questions that I had about our groupware strategy. Many of these duplicate questions in the meeting and some came up during the mail review or at other times during the last few months. This is not intended to be a complete or representative list, but does include many of the questions customers ask.

1. The strategy needs to be explicit about avoiding partner conflict, or at least about what sets of conflicts we plan. This should include or technical and business positioning with Microsoft, Oracle, Lotus, and Novell.

2. We should look for explicit ways for synergy, especially ways for marketplace leverage with other company's marketing efforts.

3. How and what PC applications are part of the strategy and how much do we have to certify them?

4. What does it take to make the File Cabinet be a transparent and integral extension of Windows 95 and of Windows NT? Can there be a way to install this so that a PC user sees the ALL-IN-1 file cabinet as "just part of the PC's file system"?

5. How will the new ALL-IN-1 file cabinet support migration from Mailworks/VMS?

6. How will the new ALL-IN-1 file cabinet support coexistence with the VAXmail file cabinet and distribution lists? This is needed not just as a conversion, but to support systems that have many of both ALL-IN-1/TeamLinks and Mail-11 users.

7. What would it take to "Objectify" each ALL-IN-1 server object in a legacy wrapper that the Object Request Broker could make available to the world? If done right, this would allow all ALL-IN-1 legacy applications to be made available through CORBA/COM to the OLE desktop.

8. The system management for ALL-IN-1 should work in conjunction with the Polycenter Netview on NT product as its design center. This platform already includes Microsoft's SMS so the system manager would have one console.

9. What would it take to integrate system management with VMS's system management if they move Argus onto the Polycenter Netview on NT platform?

10. How will the time manager be integrated with ALL-IN-1? This would have to be interoperation including shared calendars since not all ALL-IN-1 systems at a customer would convert to TeamLinks time manager at the same time.

11. How will the time manager operate in detached situations? Right now, Russell Calendar Manager does not allow me to have a copy of my calendar on my Laptop. Commercial products such as Calendar Plus allow such operation as well as shared calendaring.

12. How should the plethora of PC time managers be supported? I presume that simple import/export of the popular packages would be appropriate. For example, Calendar Plus, Lotus Organizer, Microsoft Schedule+, ACT!

13. Does the server depend on or would it benefit from VMS's kernel threads?

14. Can we "sell" client PAKs through a mechanism analogous to what is working so well for Pathworks? Basically, have the file cabinet server disseminate the PAKs against a count of licenses bought for it to hand out.

Incidentally, the slides should have been marked "Digital Confidential". Thanks.

From: MSBCS::GAILLARDET "05-Jan-1995 1014 -0500" 5-JAN-1995 10:14:20.98
To: HUMAN::CONKLIN
CC: CURTIN,GAILLARDET
Subj: **132 DOC** PARTNERING ACTIVITIES

Peter,

Congratulations on your recent appointment, we're looking forward to continuing to work with you in your new job!

Please find attached those activities which are being managed by Paul Curtin's organization on behalf of Jesse. You will notice activities include: Dispositions, Acquisitions and Porting. Should you have any questions regarding the format or project specifics, do not hesitate to conact me.

Regards,

Jean

12/27/94

MANAGEMENT SUMMARY : DISPOSITION

PROJECT MGR: Bill Hill
PRODUCT: ACMS/Family
ACTIVITY: Disposition

ORGANIZAITON: LIPCON/DEMMER
ACTIVITY CODE: 3

CONCEPT APPROVAL SBG	/ CORP	PLAN APPROVAL	RFP	POTENTIAL PARTNERS	POSSIBILITY OF CLOSURE	BIDS RECV'D	NEGOTIATION PLAN	DRC APPROVAL	TARGET COMPLETE	ISSUES
YES	SUBMITTED	N/TBD	N/TBD	Oracle Sterling Software AG	TBD				Q4/FY95	Decision required.

PROJECT MGR: Carl Gallozzi
PRODUCT: DECFORMS
ACTIVITY: Disposition via Sale

ORGANIZATION: LIPCON/DEMMER
ACTIVITY CODE: 3

CONCEPT APPROVAL SBG	/ CORP	PLAN APPROVAL	RFP	POTENTIAL PARTNERS	POSSIBILITY OF CLOSURE	BIDS RECV'D	NEGOTIATION PLAN	DRC APPROVAL	TARGET COMPLETE	ISSUES
11-94	12/15	No plan	Not to be supplied	Computer Associate Delrina JYACC	Early - NO basis for estimates	No	N/A	TBD	Unknown	None

PROJECT MGR: J. Hersey
PRODUCT: DECIntact
ACTIVITY: License back

ORGANIZATION: LIPCON/DEMMER
ACTIVITY CODE: 1

CONCEPT APPROVAL SBG	/ CORP	PLAN APPROVAL	RFP	POTENTIAL PARTNERS	POSSIBILITY OF CLOSURE	BIDS RECV'D	NEGOTIATION PLAN	DRC APPROVAL	TARGET COMPLETE	ISSUES
Yes		Yes		ASCI	H				Q3/FY95	

PROJECT MGR: Brian Gagnon
PRODUCT: DECWrite
ACTIVITY: Migration

ORGANIZATION: LIPCON/DEMMER
ACTIVITY CODE: 1

CONCEPT APPROVAL SBG	/ CORP	PLAN APPROVAL	RFP	POTENTIAL PARTNERS	POSSIBILITY OF CLOSURE	BIDS RECV'D	NEGOTIATION PLAN	DRC APPROVAL	TARGET COMPLETE	ISSUES
Yes		Yes		Interleaf	H		Yes		Q2/FY95	Neg stalled

PROJECT MGR: BOB MAY
PRODUCT: DSM
ACTIVITY: Disposition via Sale

ORGANIZATION: LIPCON/ DEMMER
ACTIVITY CODE: 3

CONCEPT APPROVAL SBG	/ CORP	PLAN APPROVAL	RFP	POTENTIAL PARTNERS	POSSIBILITY OF CLOSURE	BIDS RECV'D	NEGOTIATION PLAN	DRC APPROVAL	TARGET COMPLETE	ISSUES
COMPLETE		Yes 9/94	Yes	IDX, ISC	High	9/94	Yes	9/94	Q3/FY95 Contract signed 12/23/95	Patents FTC

PROJECT MGR: Jeff Rudy
PRODUCT: VTX/Notes
ACTIVITY: Disposition

ORGANIZATION: LIPCON/DEMMER
ACTIVITY CODE: 2

CONCEPT APPROVAL SBG	/ CORP	PLAN APPROVAL	RFP	POTENTIAL PARTNERS	POSSIBILITY OF CLOSURE	BIDS RECV'D	NEGOTIATION PLAN	DRC APPROVAL	TARGET COMPLETE	ISSUES
11/94	TBD	TBD	Y 8/94	Percussion Microsystems Russell	Med Low	9/94 9/94	TBD	TBD	Q3/FY95	Valuation

From: MSBCS::GAILLARDET "05-Jan-1995 1014 -0500" 5-JAN-1995 10:14:20.98
To: HUMAN::CONKLIN
CC: CURTIN,GAILLARDET
Subj: **132 DOC** PARTNERING ACTIVITIES

Peter,

Congratulations on your recent appointment, we're looking forward to continuing to work with you in your new job!

Please find attached those activities which are being managed by Paul Curtin's organization on behalf of Jesse. You will notice activities include: Dispositions, Acquisitions and Porting. Should you have any questions regarding the format or project specifics, do not hesitate to conact me.

Regards,

Jean

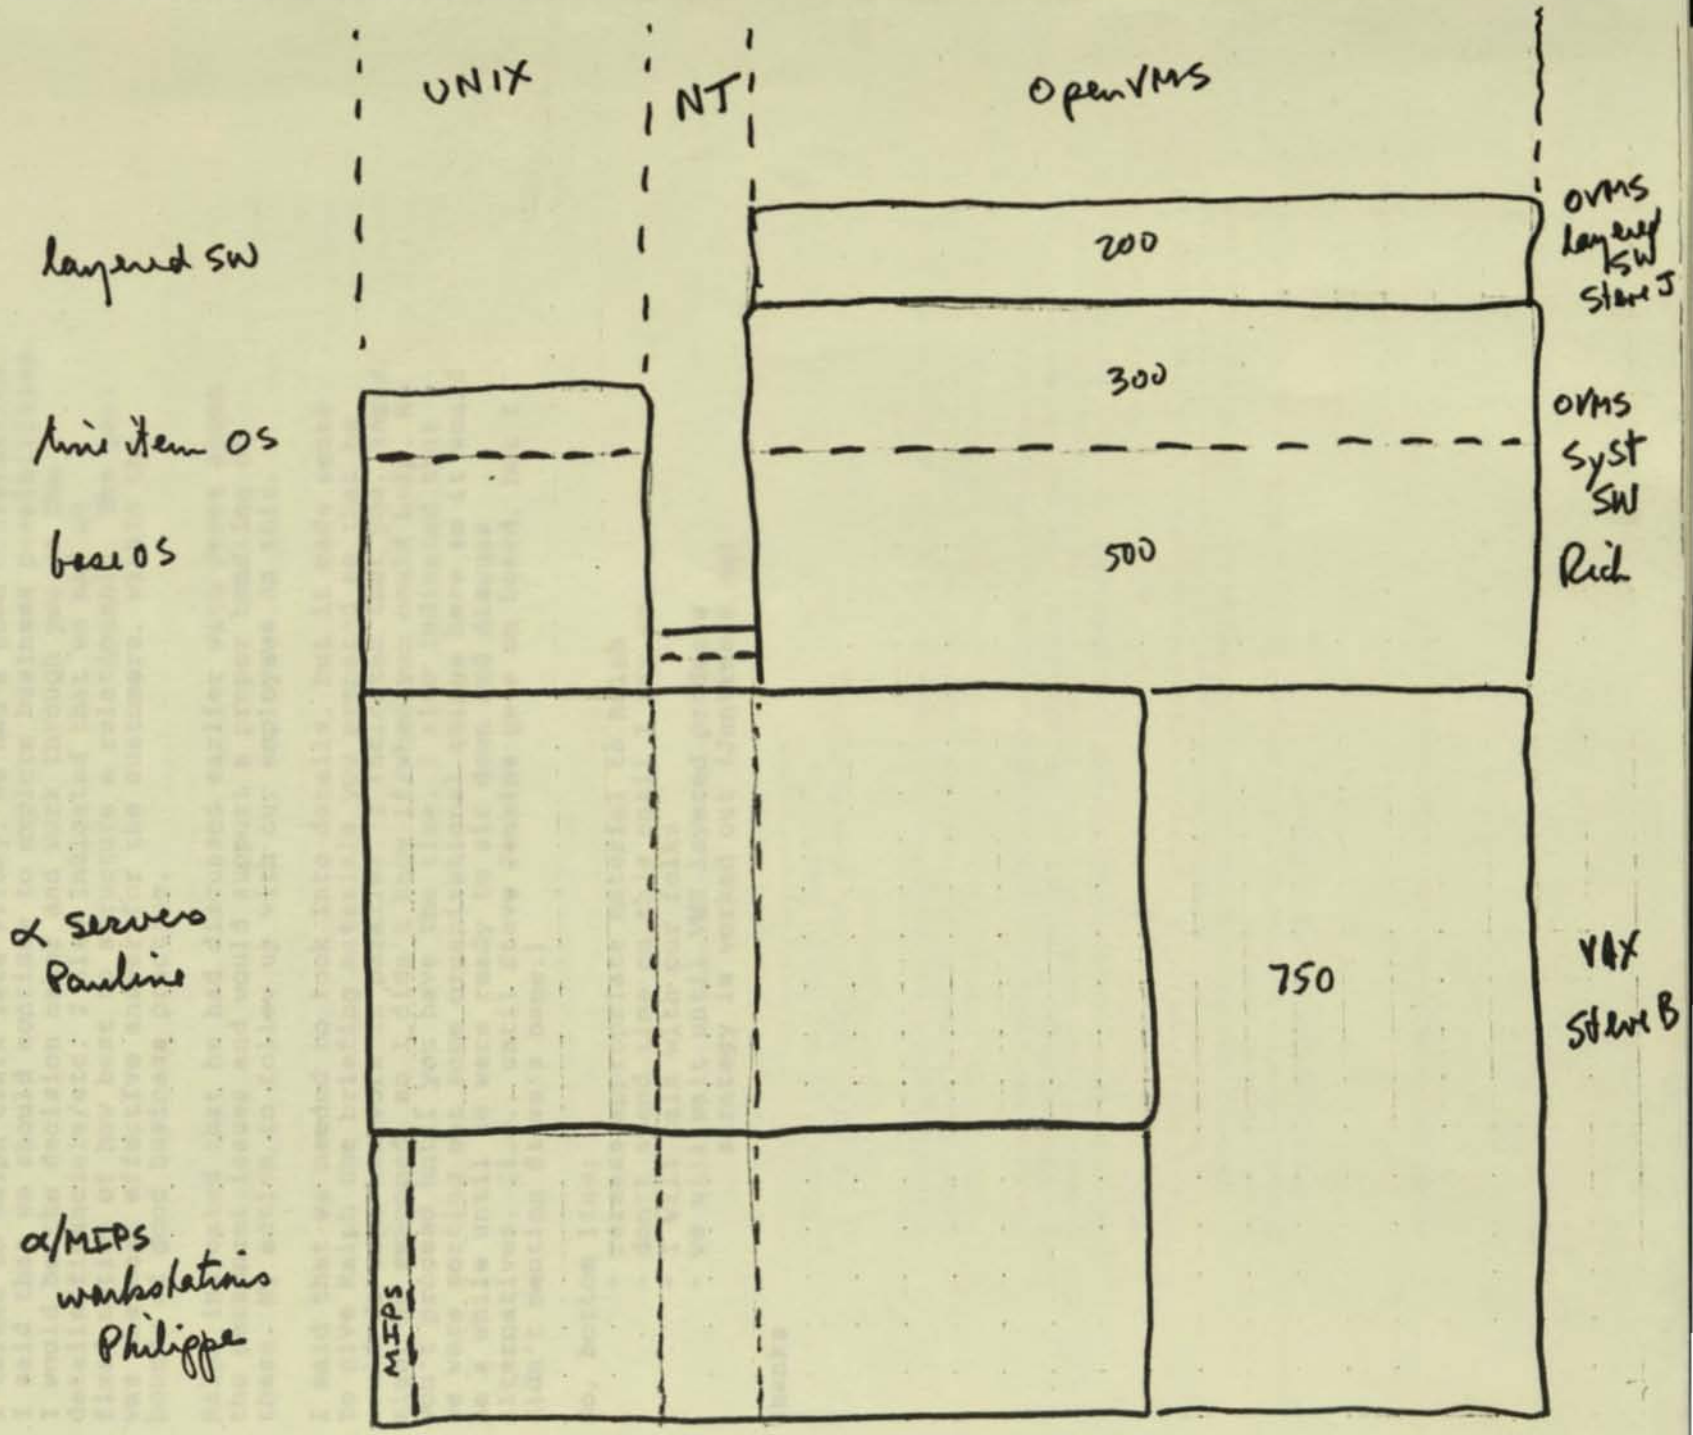
UNIX
NT
OpenVMS
layered SW
200
300
base OS
500
750
MIPS
OVMS layered SW Steve J
OVMS Syst SW Rich
VAX Steve B

From: HUMAN::CONKLIN "Peter 508.486.2564 LJO2/B11, sec Dory .2565" 10-JAN-1995 08:50:02.31
To: MSBCS::KAUFMAN
CC: CONKLIN
Subj: update on Ralph Clark

DIGITAL RESTRICTED DISTRIBUTION

I talked to Ralph Clark late Friday. We had a good conversation. I said that we should continue to explore business possibilities. I would be the decision maker and work through you on the details/financials/etc. I also indicated that we had no fixed notion of how best to structure a relationship. The goal was the most effective answer for the customers, within the bounds of good business practice.

Ralph indicated that he had discussed earlier with Jesse about the personnel issues and would support a proper handling of these. My action to follow up with our employees on this.

I said that we needed to look into details, but it made sense to give Ralph the briefing materials you prepared so that he could do some homework in parallel. I indicated that you, Regis, might be swampped, so I didn't know if/when you could help. We won't proceed until you have the time. I also indicated that we were sorting out some organizational things here so it would be a while until we were ready to sit down and discuss alternatives. (i.e., until Steve Jenkins gets on board, but I didn't mention Steve's name.)

So, bottom line:

- release appropriate material to Ralph
- don't spend time on this until I say so
- I will talk with our folks
- we will wait until VMS layered products strategy is worked out (Jenkins & me)

thanks.

Super-Servers: Commodity Computer Clusters Pose a Software Challenge

Jim Gray
San Francisco Systems Center
Digital Equipment Corporation
455 Market St. 7th Fl., San Francisco, CA. 94105
SFBay::JimGray or JimGray @ SFBay.enet.dec.com
February 1991, revised December 1991

Abstract: Technology is pushing the fastest processors onto single mass-produced chips. Standards are defining a new level of integration: the Posix box. These forces will fundamentally change the way we build computers. Future designs must leverage commodity products. Clusters of computers are the natural way to build the supercomputer of the future. A simple analysis suggests that such machines will have thousands of processors, terabytes of RAM, many terabytes of disc, and terabits-per-second of communications bandwidth. This gives rise to the 4T clusters. These computers will be ideally suited to be super-servers in future networks. Software that extracts parallelism from applications is the key to making clusters useful. Client-server computing has natural parallelism: many clients submit many independent requests that can be processed in parallel. Database, visualization, and scientific computing applications also have made great strides in extracting and exploiting parallelism. These promising first steps bode well for cluster architectures.

Outline:

Acknowledgments: These ideas have been been evolving for many years; but, this memo grew out of the IO taskforce chaired by Barry Rubinson. Participants included Bob Bean, Andrew Birell, Verell Boaen, Barry Goldstein, Bill Laing, Richie Lary, Alan Nemeth, Ron Obermarck, Tom Rarich, Dave Tiel, and Cathy van Igen.

Confidentiality: A previous version unfortunately spread outside Digital. It was my fault that the memo was not labeled Digital Confidential. Fortunately, that memo, like this one, was conceptual rather than factual; so little confidential information was disclosed. Since the confidential version spread widely, there seems little reason to make this sanitized version a secret. It has not gone through the Digital review process, and so it is not public.

Changes from the previous version:

1. Networking is no longer ignored. High-speed networks are mentioned (gigabit LANs and megabits WANs). In fact, this is the BIG change in computer architecture. Other parts of the computer are getting only ten to one hundred times cheaper and faster in the next decade. Networking is getting thousands or millions of times faster and cheaper in the next decade.
2. There is an attempt to contrast clusters with distributed systems; clusters are simple distributed systems (homogeneous, single site, single administration).

STANDARDS ARE COMING!

By the end of the decade, boatloads of Posix boxes, complete with software, will be arriving in ports throughout the world. They will likely be 100 times more powerful than the VAX-9000, and will cost less than 10,000$ each, including a complete NAS-like software base. No doubt they will come in a variety of shapes and sizes, but typically these new super-computers will have the form factor of a PC or VCR. These products will be inexpensive because they will exploit the same software and hardware technologies used by mass-market consumer products like HDTV, telephones, voice and music processors, super-FAX, and personal computers.

How can Digital and other computer companies add a hundred billion dollars of value to these boxes each year? Such added value is needed to keep computer industry giants like IBM, Fujitsu, and Digital alive.

I believe that the 100B$/year will come from three main sources:

Manufacture: Provide some of the **hardware** and **software components** in these boxes.

Distribute: **Sell, service,** and **support** these platforms to corporations. Although the boxes will be standard, corporations will want to out-source the expertise to install, configure and operate these boxes and the networks that connect them.

Integrate: Sell **corporate electronics**, by analogy to consumer electronics, prepackaged or turnkey systems that directly solve the problems of large corporations. The proliferation of computers into all aspects of business and society will create a corresponding demand for **super-servers** that store, analyze, and transmit data. Super-servers will be built from hundreds of such boxes working on common problems. These super-servers will need specialized application software to exploit their cluster architecture. Database search and scientific visualization are two examples of such specialize application software.

As in the past, most revenue will come from manufacturing and distribution – the traditional computer business. The high profit margins will be in integrated systems that provide unique high-value products.

Integration is not a new business for traditional computer companies, but the business structure will be different. There will be more emphasis on using commodity (outside) products. The development cost of standard products will have to be amortized across the maximum number of units. These units will be marketed to both competitors and to customers. Development of non-standard products will only be justified for items that make a unique contribution with orders-of-magnitude payoffs. The cost of me-too products on proprietary platforms will be prohibitive.

This phenomenon is already visible in the PC-Workstation marketplace. In that market, standardized hardware with low margins provides the bulk of the revenue, but has low profit margins. A few vendors dominate the high-margin software business (notably Microsoft and Novel).

4B MACHINES: SMOKING HAIRY GOLF BALLS

Today, the fundamental computer building blocks are cpus, memory chips, discs, print engines, keyboards, displays, modems, and Ethernet. Each is a commodity item. Computer vendors add value by integrating these building blocks and by adding software to form workstations, mid-range computers, and to some extent mainframes. For example Apple, Compaq, IBM PCs, NCR, Sequent, Sun, Tandem, and Teradata all use commodity components. Digital's MIPS-based products are also examples of this approach.

The unit of integration has gone from vacuum tube to chip. The next step in integration will be a minimal hardware/software package. By the end of this decade, the basic processor building blocks will be commodity boards running commodity software. The boards will likely have a 1 bips cpu (billion instructions per second), 1 GB (Giga byte) of memory, and will include a fairly complete software system. This is based on a technology forecast something like the following.

Year	1 Chip CPU Speed	1 Chip[1] DRAM	1 Disc 1GB	LAN	WAN
1990	10 mips	4 Mb	8"	10mbps Ethernet	64kbps ISDN
1993	80 mips	16 Mb	5"		
1996	500 mips	64 Mb	3"	100 mbps FDDI	?
1999	1000 mips	256 Mb	1"	1000 mbps ?	1 mbps fiber

This forecast is fairly conservative (for example, some predict 1Gb WAN networking will be economic in that period). It also forecasts the following costs for the various 1999 components

Year	1 Chip CPU	1 Chip DRAM	1GB Disc	LAN	WAN
1999	100$	5$	50$	50$	50$

Given these costs, one could buy a processor, 40 memory chips, several high-speed communications chips, and ten discs, package and power them for a few thousand dollars.

Such computers are called **4B machines** (Billion instructions per second, Billion bytes of DRAM storage, and a Billion bytes per second of IO bandwidth, and a Billion bits per second of communications bandwidth). A **5B machine** will support a Billion bit display, that is 4000x4000 pixels and each pixel 32 bits of shading and color[2]. They contrast to the 5M machines that drove the PC revolution (mip, megabyte of ram, megapixel display, 10 megabit per second LAN, and a mouse).

These 4B machines will be **smoking-hairy-golf-balls**[3]. The processor will be one large chip wrapped in a memory package about the size of a golf ball. The surface of the golf ball will be hot and hairy: hot because of the heat dissipation, and hairy because the machine will need many wires to connect it to the outside world.

Dramatic changes are also expected in both storage or networks.

Disc farms will be built from mass-produced 1" discs placed on a board much as DRAMs are placed on memory boards today. A ten-by-ten array of such discs will store about 100 GBytes.

---

1 This is the surprise-free prediction that the trend of the last 20 years will continue. There is good evidence that DRAMS are evolving more slowly than they have in the past. This slower evolution comes from reduced demand and increased capital costs. If recent trends continue, in 1999 DRAMS chips will be at 64Mb and will cost about 15$ each. This will increase projected memory prices by an order of magnitude and will decrease the projected memory size by a similar factor. Thanks to Steve Cullen of Digital for this observation.

2 Some prefer to call these 4G and 5G machines using Giga instead of Billion.

3 Frank Worrell used this metaphor in 1985. Frank is now working at LSI Logic.

Disc array technology will give these disc-boards very high performance and very high reliability[4].

Future networks will be much faster. Fiber based communications will be able to deliver terabit data rates, but at a high price. Commodity fiber-optic interfaces will probably run at gigabit speeds. Local communication (LANs) will be able to use this bandwidth, but long haul bandwidth will still be expensive. So, although gigabit-WANs will be possible, and may form the backbones of some applications, it seems likely that megabit-WANs will be more typical. The transition from the low speed WANs of today running at between 2kbps and 64kbps, to the higher-speed commodity WANs of 1999 running at 2mbps (T1) to 45mbps (T3) will be a major architectural shift for data communications. These changes in network performance and network economics will be key enablers for super-servers to instantly distribute data and images over long distances.

The software for 4B machines will contain all the elements of X/Open, Posix, DCE, SAA, and NAS. In particular it will include some standard descendents of Motif, SQL, OSI, DCE-UNIX, X/Open transaction processing, and so on.

These basic building blocks will be commodities. That is, the hardware will be mass produced and so will have very low unit price. Standard operating systems, window systems, compilers, database systems, and transaction monitors will have high volumes and so will also have low unit prices. This can already be seen in the workstation world. There, OS/2 Extended Edition and Open DeskTop provide complete software systems (database, network, and tools), all for less than a thousand dollars.

Today, most applications are not portable from one computer family to another (e.g., from MS/DOS to UNIX). In that era, most applications will be portable. The stable interfaces will be soft: the programming languages, operating system, i/o libraries, databases, network protocols, and the like.

---

4 Patterson, D. A., G. Gibson and R. Katz. (1988). *A Case for Redundant Arrays of Inexpensive Disks (RAID).* Proc .ACM SIGMOD. 109-116. or Schulze, M., G. Gibson, R. Katz and D. A. Patterson. (1989). *How Reliable is a RAID.* 34th IEEE Compcon 89. 118-123.

BUSINESS STRATEGY IN AN ERA OF COMMODITY SOFTWARE

Profit margins on manufacturing commodity hardware and software products will be relatively modest, but the volumes will be enormous. So, it will be a good business, but a very competitive one. There will continue to be a brisk business for peripherals such as displays, scanners, mass storage devices, and the like. But again, this will be a commodity business with narrow profit margins. Much like the commodity PC industry of today.

Why even bother with such a low-margin business? The reason is simple, it is essential to be in the low-margin business because it is the high-volume business. The revenues and technology from this business fund the next generation. This can already be seen in the IC business where DRAM manufacturing refines the techniques needed for many other advanced devices.

There is a software analogy to this phenomenon visible within IBM, Digital, Microsoft, and Oracle. There are economies-of-scale in advertising, distributing, and supporting software. Microsoft's Windows3 product shows once again the importance of an installed base and of a distribution network. In addition, the pool of software expertise in developing one product is a real asset in developing the next.

On the other hand, we already see that IBM cannot afford to do all of SAA and that Digital cannot afford to do all of NAS. These projects are so huge that they are being stretched-out over the next decade. In fact, they are so huge, that strategic-alliances are being formed to spread the risk and the workload. This, in my view, is a root cause of the many consortia (e.g., OSF, ACE,...) being formed today. For IBM and Digital to recover their development costs for SAA and NAS, their software efforts will have to become ubiquitous. NAS and SAA must run on millions of non-Digital and non-IBM hardware platforms. They must target the portable software market as a major revenue source.

There is no longer room for dozens of companies building me-too products. For example, each operating system now comes with a SQL engine (Rdb on VMS, Ingres on OpenDesktop, SQL on OS/2 EE, NonStop SQL on Guardian,...). So it will be hard to make a profit on a unique SQL engine – SQL is now commodity software. Each computer company or consortium must either build an orders-of-magnitude-better unique-but-portable SQL product, or form an alliance with one of the portable commodity SQL vendors. Put glibly: each company has a choice, either (1) build a database system and database tools that will blow away Oracle, Ingres, Informix, and the other portable database vendors, or (2) form an alliance with one of these commodity vendors. Similarly, each company must produce a networking system that is orders-of-magnitude-better than DECnet, SNA, Novel, 3-Com, Ungermann-Bass, and the other network vendors. If it cannot afford to do that, it should partner with one of them.

In general, each computer company will both build and buy. For example, today most of the UNIX-DCE code comes from other vendors, some of the DCE components come from Digital. This probably represents the way things will be in the future; no company can afford to do everything. No single company can produce the best implementation of all standards.

SALES AND SERVICE IN A COMMODITY WORLD

I have little to say on this topic. It is one triad of the three components of any computer companie's future. It is our traditional business. Size and wide geographic distribution are a key strength in marketing and supporting any products a company offers.

FUTURE MAINFRAMES: THE 4T MACHINES

In a classic paper Gordon Bell and Dave Nelson defined the basic laws of computing[5]. One of their key observations is that there are seven tiers to the computer business. These tiers are roughly categorized by the dollar value of the computers:

10$: wrist watch computers
100$: pocket/ palm computers
1,000$: portable computers
10,000$: personal computers (desktop)
100,000$: departmental computers (closet)
1,000,000$: site computers (glass house)
10,000,000$: regional computers (glass castle)

Bell and Nelson observed that each decade, computers from one tier move down a notch or two. For example, current portables have the power and capacity approximating that of a 1970 glass-house machine. Machines with the power of 1980 workstations are now appearing as portable and even pocket computers.

They observed that service workers can be capitalized at about 10,000$ of computer equipment per person on average. That more or less defines the price of the typical workstation.

The costs of departmental, site and regional servers can be amortize over many more people, so they can cost a lot more.

What will the price structure look like in the year 2000? Will there be some super-expensive super-fast neural-net computer that costs ten million dollars? If future processors and discs are very fast and very cheap, how can one expect to build an expensive computer? What will a main-frame look like?

One theory is that the mainframe of the future will be 10,000$ of hardware and 990,000$ worth of software. Being a software guy, I like that model. Fighter planes work this way, each new one is smaller and lighter, yet costs much more because it is filled with fabulously expensive software and design. It's unlikely that similar mechanisms will operate for commodity super-servers.

OK, so the 99% software theory is blown. What else? Perhaps the customer will pay for 990,0000$ worth of maintenance or service on his 10,000$ box? Probably not. He will probably just buy two, and if one breaks, discard it and use the other one.

I conclude that the mainframe itself will cost about a million dollars in hardware. What will a million dollars buy? It will buy (packaged and powered) about:

~	1,000	processors	= 1 tips (trillion instructions per second) or
~	100,000	DRAMS (@256Mb)	= 4 TB (four terabytes RAM) or
~	10,000	discs (@1GB)	= 10 TB (ten terabytes disc) or
~	10,000	net interfaces (@1Gbps)	= 10 Tb (10 terabits of networking)

So, the mainframe of the future is a 4T machine!

---

5 See C.G. Bell and J.E. MacNamera, *High Tech Ventures*, Addison Wesley, 1991, pp. 164-167

WHO NEEDS A 4T SUPER-SERVER?

What would anyone do with a 4T machine? Perhaps the mainframe of the future is just a personal computer on each desk. A thousand 4B PCs would add up to a 4T "site" computer. The system is the network!

Each worker will probably have one or more dedicated 4B computers, but there will be some jobs that require more storage or more processing than a single processor, even one of these super-powerful 4B ones.

Consider, for example, the problem of searching the ten terabyte database mentioned above looking for a certain pattern. If one processor searched through the 10 TB using a single 4B processor, and using current software (e.g. Rdb), the search would take three hours. By using a thousand 4B processors in parallel, the search would take about 10 seconds.

Similar observations apply to other applications that analyze or process very large bodies of data. Database search is prosaic compared to data visualization algorithms that map vast quantities of data to a color image. These search and visualization problems lend themselves to parallel algorithms. By doubling the number of processors and memories, one can **scaleup** the problem (solve twice as big a problem), or **speedup** the solution (solve the problem twice as fast).

Some believe that the 4B machines spell the end of machines costing much more than 10,000$. I have a different model. **I believe that the proliferation of inexpensive computers will increase the need for super-servers.**

A fraction, say 25%, of future computer expenditures will go for super-servers. The typical strategy today is to spend half the budget on workstations, and half on print, storage, and network servers. In the end, the split may be more like 90-10 (this is the ratio of cost of ATMs to the host server in an ATM network), but servers will not disappear. The central arguments are:

Power: The bandwidth and data storage demands of servers supporting hundreds or thousands of 4B machines will be enormous. The servers will have to be more powerful than the clients. **Fast clients want faster servers.**

Control: The proliferation of machines and bandwidth will make it possible, even easy, to access centralized services and resources. No longer will you go to the video store to get a videotape, you will download it. No longer will you search paper libraries for information, you will have a server do it for you. These resources (movies, libraries,...) will contain valuable information. Central utilities (or at least regional utilities) will want to control access to them. So they will set up super-servers that offer an client-server interface to them.

Manageability: People do not want to manage their own data centers. Yet, the trends above suggest that we will all own a personal data center in 1999. Each PC and each mobile telephone will be a 4B machine. There will be a real demand for automatic data archiving and automatic system management. This will likely be a centralized service. A simple example of this is visible today with the success of X-terminals that move management issues from the desktop to the closet.

WHAT ARE THE KEY PROPERTIES OF SUPER-SERVERS?

Servers must have the following properties:

Programmable: It is easy to write client and server applications for the server.

Manageable: It is easy to manage the server.

Secure: The server can not be corrupted or penetrated by hackers.

Distributed: The server can interoperate with other super-servers.

Scaleable: The server's power can grow arbitrarily by adding hardware.

Economic: The server is built from commodity components.

Highly available: The server does not lose data and is always "up".

CLUSTERS – THE KEY TO 4T MACHINES

Servers need to be as powerful or more powerful than their clients. They must serve hundreds or millions of clients. How can powerful servers with all these properties be built from commodity components? How can a collection of hundreds of 4B machines be connected to act as a single server? What kind of architecture is needed? What kind of software is needed?

Digital has it now! Digital currently offers VAXclusters that scale to hundreds of processors. The cluster has excellent programming tools, it is a single management entity, and it is secure. Clients access ACMS servers on the cluster not knowing where the servers are running or where the data resides. So the cluster is scaleable. Processors, storage, and communications bandwidth can be added to the cluster while it is operating. VAXclusters are fault-tolerant; they mask faults with failover of discs and communications lines. VMS has the transaction concept integrated into the operating system. The VAX family is built from commodity components and is among the most economic servers available today.

Well, that is the official Digital marketing story; and there is a grain of truth to it. But, the details of the VAXcluster do not deliver on most of these promises. The VAXcluster really only scales to tens (not thousands of processors), the programming and management tools do not offer much transparency; each component is managed individually. Virtually none of the tools use more than one-processor-at-a time in running an application; this dramatical limits the ability to scaleup or speedup applications by adding hardware. The VAXcluster price is not especially economic when compared to PC-based servers. And, there are single points of failure in the VAXcluster software.

But, the VAXcluster is certainly a step in the right direction. It is also the direction that most other vendors have adopted. Notable examples are:

Teradata builds clusters out of the Intel x86 family and proprietary software. These clusters act as back-end SQL servers to mainframes and LANs. Teradata systems feature economy, scaleability, and fault-tolerance. The largest clusters are a few hundred processors and a thousand discs. **NCR** has adopted the Teradata approach. It is hoping to build systems that scale from the palm to the super-computer by building clusters of Intel x86 or RISC processors. At present, the

NCR plan is in development[6]. The recent merger of AT&T, NCR, and Teradata make this all the more interesting.

Tandem builds clusters out of a proprietary hardware-software combination running as network servers. Tandem system features match the super-server list above. The systems scale to about 100 processors and to a few hundred discs. Customers complain that the systems are not manageable.

IBM Sysplex is a cluster of up to fortyeight 390 processors. At present there is very little software to support this cluster hardware. In addition, the IBM AIX system (their UNIX clone) running on the 390 hardware has software from Locus corporation to support the cluster concept. IBM's Yorktown research group shows a lot of interest in cluster software.

Intel is building a hyper-cube of 2000 processors, called the Delta machine. The first instance of this machine (about 500 processors) is now installed at CalTech, and other smaller machines are being seeded around universities and industry. The faculty at CalTech and other universities are working hard on the programming issue. Unlike the other machines mentioned so far, input/output to disc storage seems to be an afterthought this machine.

Thinking Machines recently announced a massively parallel machine. I have little real information on this machine, but it seems to match the 4T model.

Looked at in this light, no vendor is ready to build a thousand processor 4T machine *and the associated software*. Some are ahead of others. In fact, I place Digital second behind Tandem in the race – the VAXcluster is an excellent start. Digital's VMS staff and other engineering organizations have experience with clusters. *Digital should be leveraging its leadership position to extend VAXcluster sizes by one or two orders-of-magnitude.*

---

6 *NCR's 486 Strategy*, Moad, J., *Datamation*, V36.23, 1 Dec. 1990, pp. 34-38.

CLUSTER SOFTWARE – THE KEY TO 4T CLUSTERS

It is important to understand the virtue of clusters. The idea is to add more discs, more processors, more memory, and more communications lines, and get more work out of the system. This speedup and scaleup should go from one processor-memory-disc-comm module to several thousand modules.

Traditionally, multiple processors have been connected by sharing a common memory: Shared Memory Multi-Processors (SMP). The SMP approach does not scale well in a world of smoking-hairy golfballs. The event-horizon of a smoking hairy golfball is on the processor chip; signals from one ball cannot get to the next ball before the processor goes on to the next instruction. A memory shared by two such golfballs looks more like a communications line or a remote processor. SMP designers are aiming for ten-way parallelism. The hundred-fold and thousand-fold speedups available using commodity processors in a cluster have much higher payoff. All the machines mentioned above have a cluster architecture. They communicate via messages rather than via shared memory.

The goal of cluster software is to divide-and-conquer large problems. It must do three things:
1. break the computation into many small jobs,
2. spread the jobs among many processors and memories executing in parallel, and
3. arrange that traffic among the jobs does not swamp the network or create interference.

The challenge has been to extract parallelism from applications. Certain applications like timesharing and transaction processing have natural parallelism. Each client represents a separate and independent request. Each request can go to a separate processor. So, servers with many clients have inherent parallelism. If the number of clients doubles, and if there are no bottlenecks in the hardware or software design, then doubling the number of servers, storage devices, and comm lines will give good scaleup. This is what VAXclusters, Teradatas, and Tandems do today.

The real challenge is to recognize and extract the parallelism *within* applications (big batch jobs). This is an ad hoc field today. SQL servers have discovered how to extract parallelism from large database queries – they search each disc of the database in parallel, they sort in parallel, they join tables in parallel, and so on. These systems display good speedup and scaleup to a hundred processors. Notable commercial examples of this are Teradata and Tandem[7].

Beyond that there have been few successes. Today, recognizing parallelism is an application-specific task. The application programmer must program parallelism into his application by inventing new and innovative algorithms. Automatic extraction of parallelism from applications stands as a major research challenge.

The current situation is (1) 4T machines have a bright future as parallel SQL servers and (2) servers get natural parallelism and scaleup from having many clients. So, no scientific breakthroughs are needed to get the parallelism needed to use 4T machines as data servers. These servers will have lots of opportunities for parallelism.

---

7 David DeWitt and Jim Gray, *Parallel Database Systems: The Future of Database Processing or a Passing Fad?*, to appear in CACM.

VAXcluster SOFTWARE – THE KEY TO 4T MACHINES?

Once the parallelism problem is "solved" innovation is still needed to make VAXclusters manageable, secure, and highly available. Evolving VMS and the VAXcluster to solve any of the major problems (parallel software, manageability, security, fault-tolerance) will be a major software initiative.

There is a fundamental question about whether we should base these initiatives on top of a proprietary system (VMS) or on top of an Open System (OSF DCE UNIX). I am unclear on the answer to this VMS vs UNIX question. VMS may not be the perfect base for Digital's future cluster products – it is not portable to the instruction-set-of-the-month, and it is very large. But it has two big virtues: (1) It supports the cluster concept, and (2) it is ours.

The easy way out of this is to base future VAXclusters on UNIX. UNIX is portable and standard. The problem is that UNIX is like stone soup[8]. You have to add a lot to get what you want. If we add 1M lines to UNIX for fault tolerance, 1M lines for distributed databases, and 10M lines for manageability, do we still have UNIX? Have we built a commodity product? Will super-servers be a commodity product. I do not think so.

Super-Servers will use commodity hardware and proprietary software. That is, I think the super-server will have sales volumes measured not in millions of units, but in thousands of units – one super-server per thousands of clients. Super-server operating and management software will have demanding requirements that will not be satisfied by commodity client software. So I imagine that there will be a few server operating systems that offer an Open Interface (e.g. SAA, NAS, Posix, X/Open, or the like), run on clusters of commodity devices, but that are proprietary. That is the software will have many unique performance and management features.

This is good news for anyone who wants to make a business of super-servers. If they were easy to build and had huge volumes then there would be a lot of competition for them and margins would be very slim. The key to a successful business is having a product that everybody needs but that few people can build. The software that goes into super-servers may well be such a product.

The next section tries to explain why standard software (e.g. vanilla UNIX-DCE) is unlikely to produce a competitive cluster architecture.

---

[8] The recipe for stone soup calls for a stone to be placed in a large pot of boiling water. Each guest is requested to bring an additional ingredient (e.g., onions, carrots, ...). The quality of the soup depends on the quality of the guests.

STANDARDS: TELL ME IT ISN'T SO (Snake Oil).

Some believe that all this Open-UNIX-Standards stuff is Snake Oil (SO for short). I do too – well perhaps its not all snake oil, but there sure is a lot of hype floating about. This is an unpopular view – or at least a reactionary one; but it deserves a fair hearing. The SO view proceeds as follows.

Standards are Boring: Customers always want some feature or trend that is leading edge. They use this as a competitive advantage. Leading edge things have not made it into standards. Parallelism, fault-tolerance, manageability, and high-performance tricks are unlikely to become standards.

Standards are Incomplete: It is standard to see the seven-layer ISO protocol stack. You have seen the 1000-page SQL standard. You have seen the multi-volume X-Windows books. Guess what? They are the tip of the iceberg.

- The ISO protocol stack has an elevator shaft running down the side called network management. That elevator shaft is not standard. There are implementations that are de facto standards (e.g., NetMaster, DECnet EMA, NetView), but they are not standard. ISO currently punts on issues like security and performance.
- The SQL standard looks the other way about most errors (they just define a few simple ones), performance (no performance monitor), utilities (no load/dump, import/export,...), and administration (no accounting, space management,...).

The SO reactionaries believe that computing is fractile: there is complexity in every corner of it. Workstations hide this complexity by dealing with a single user and ignoring system management. Real computers will not be able to hide some of this complexity.

When building a workstation, one aims for simplicity. Microsoft has an *open* standard MS/DOS - a single code body that is its own spec. Apple's Macintosh is a similar story. The UNIX world has a standard *open* application programming interface that allows many simple stand-alone programs to be easily ported from one platform to another. The CICS world has 300,000 programmers who know and love the CICS application programming interface – that is its own spec. These systems are all *open* and standard. Their programming interfaces are published and do not change.

But there is a separate world. There is no real open-standard operations interface for a network of PCs, or for the applications that run them. All the tools to do these operations tasks are proprietary. The CICS operations interface is not well documented, is not open, and it changes from release to release. It is the elevator shaft.

Certainly, the standards organizations have place-holder bodies that are "working" on these elevator shafts, but the SO reactionaries believe such efforts are doomed. These big-systems issues are too specific to become commodity standards or products.

Standards are low-performing: The standard NFS protocol stack from SUN has poor performance. SUN and other vendors have *deep* ports of this standard code that are much faster. They sell the fact that they have the best NFS. Auspex has done them one better, building a high-performance, fault-tolerant, scaleable NFS super-server. Someone who offers a vanilla NFS server will have a difficult time competing.

Similar comments apply to SQL. Rdb is a deep port of SQL to VMS (actually Rdb was written explicitly for VMS). Other SQL systems do not take advantage of VMS. Rdb recently fixed its performance and quality problems and is now displacing other SQL systems on VMS; it is better, less expensive, and comes from the hardware vendor. These

performance issues are more important for servers than for clients, since servers are resource-poor compared to clients (even in a cluster).

There is no question that the mass-marketed systems will be running standard, mass-marketed software. But, if this model is correct, then super-server cluster will have commodity hardware and proprietary software. The software may be "open" in the SAA-NAS-SUN-Posix sense that applications are portable to it and can interoperate with it. But it will not be the commodity OSF software; it will have LOTS of value added in the areas of scaleability, security, manageability, and availability. In this model, the clients (millions of them) will be running commodity software, but the servers will not.

This reactionary SO view describes the current situation: today the clients are commodity MS/DOS or UNIX systems. The small servers are also commodity systems. But, past a certain threshold the commodity servers hit a wall. The hardware does not scale up to clusters and neither does the software.

Server vendors have a choice; they can build this super-server cluster software and call it anything they want. They can call it VMS or UNIX or NAS. It will be mostly new code. It will surely be X/Open branded, Posix compliant, and support all of DCE; so it will be UNIX. Coming from DEC, it will have many VMS VAXcluster and DECnet features; so we might call it VMS. But it will be mostly code that is in neither VMS nor UNIX today.

CLUSTERS vs DISTRIBUTED SYSTEMS, WHAT'S THE DIFFERENCE

Why isn't a cluster just a distributed system? Why won't all the wonderful software we have been developing for distributed systems apply directly to clusters? Won't DCE solve the cluster problem? After all, DCE means Distributed Computing Environment.

Well, right! A cluster is a distributed system. Everything, even an isolated PC is a distributed system. That is the virtue of distributed systems, they encompass all and integrate everything.

A cluster a *special* kind of distributed system. It has properties that make it qualitatively different.

Homogeneous hardware and software: A distributed system necessarily involves many types of computers with many different software systems. This heterogeneity comes at a cost. General purpose algorithms are needed to communicate among nodes. Things on one node are slightly different than things on another, so that it is very expensive, if not impossible, to offer transparent access to all data at all nodes.
In a cluster, all the nodes are running the same software and have approximately the same hardware. This simplicity,has huge benefits, both for performance and for transparency. It is relatively easy to give the illusion that the entire cluster is a single computer.

Single administrative domain: Distributed systems are designed to cross organizational and geographic boundaries. Each node is considered to be an independent member of a federation. Since boundaries among nodes of the network are explicit, designers of distributed systems make many design choices that allow fine-grain (node-level) control and authorization.
A cluster is more like a single node of a distributed system. The cluster may consist of thousands of devices, but it is managed as a single authentication domain, a single performance domain, and a single accounting domain. The cluster administrator views it as a single entity with no internal boundaries.

Ideal communication: In a distributed system communication is slow, expensive, and unreliable. The finite speed of light and long distances make it slow - 100 ms round trip is typical. The long distances and huge capital costs of common carriers imply that communications lines are the most expensive part of a distributed computer system. Public networks lose individual connections, and occasionally deny service for extended periods.
By contrast, communication within a cluster is ideal. The distances are short, less than 100 meters; so the speed of light delay is short, less than a microsecond. Bandwidth is plentiful and inexpensive in a cluster. One can just add more ports and fibers. The short distances and low communication complexity within a cluster give highly reliable communication. Since all members of the cluster speak the same language, very efficient communications protocols can be used.

In summary, a cluster is a special kind of distributed system. Distributed systems techniques will help in building clusters, but the differences will make clusters both simpler and faster than distributed systems. In a sense, it is much easier to build a cluster than to build a distributed system.

Of course, a cluster acting as a server will be a key part of a distributed system. It will be a super-server node of the distributed system.

CONCLUSION AND RECOMMENDATION

The PC marketplace shows how mass-production and economies-of-scale can mask the engineering costs that dominate minicomputer and mainframe prices today. Technology is pushing the fastest processors onto single mass-produced chips. Standards are defining a new level of integration: the Posix box. These forces will fundamentally change the way we build computers. Future designs must leverage commodity products.

Clusters of computers are the natural way to build the supercomputer of the future. A simple analysis suggests that such machines will have thousands of processors, terabytes of RAM, many terabytes of disc, and terabits-per-second of communications bandwidth. This gives rise to the 4T clusters. These computers will be ideally suited to be super-servers in future networks.

Software that extracts parallelism from applications is the key to making clusters useful. Client-server computing has natural parallelism: many clients submit many independent requests that can be processed in parallel. Database, visualization, and scientific computing applications also have made great strides in extracting and exploiting parallelism. These promising first steps bode well for cluster architectures.

Digital should enter the cluster race as one component of its business strategy. It should set a year 2000 goal to build a 2000-processor 4T machine and the software to make the machine a super-server for data and applications. The super-server software should offer good application development tools. The 4T cluster should be manageable and should offer good data and application security. Remote clients should be able to access applications running on the server via standard protocols. The server should be built of commodity components and be scaleable to thousands of processors. The software should be fault-tolerant to the extent that the server or its remote clone can offer services with very high availability.

The seeds for this architecture are already found in the VAXcluster architecture.

Hardware Trend Summary

- General: tried and true performance trend 'rules of thumb' still very accurate
  - Microprocessors - 2x performance improvement every 2 years
  - Memories (DRAMs) - 4x capacity increase every 3 years; each generation is 10 ns faster than previous generation
  - Storage - 2x MB/$ increase every 2 years; latencies and physical dimensions continue to decrease
- A few new items
  - Flash RAMs
  - Ultra low power microprocessors

Microprocessors (cont'd)

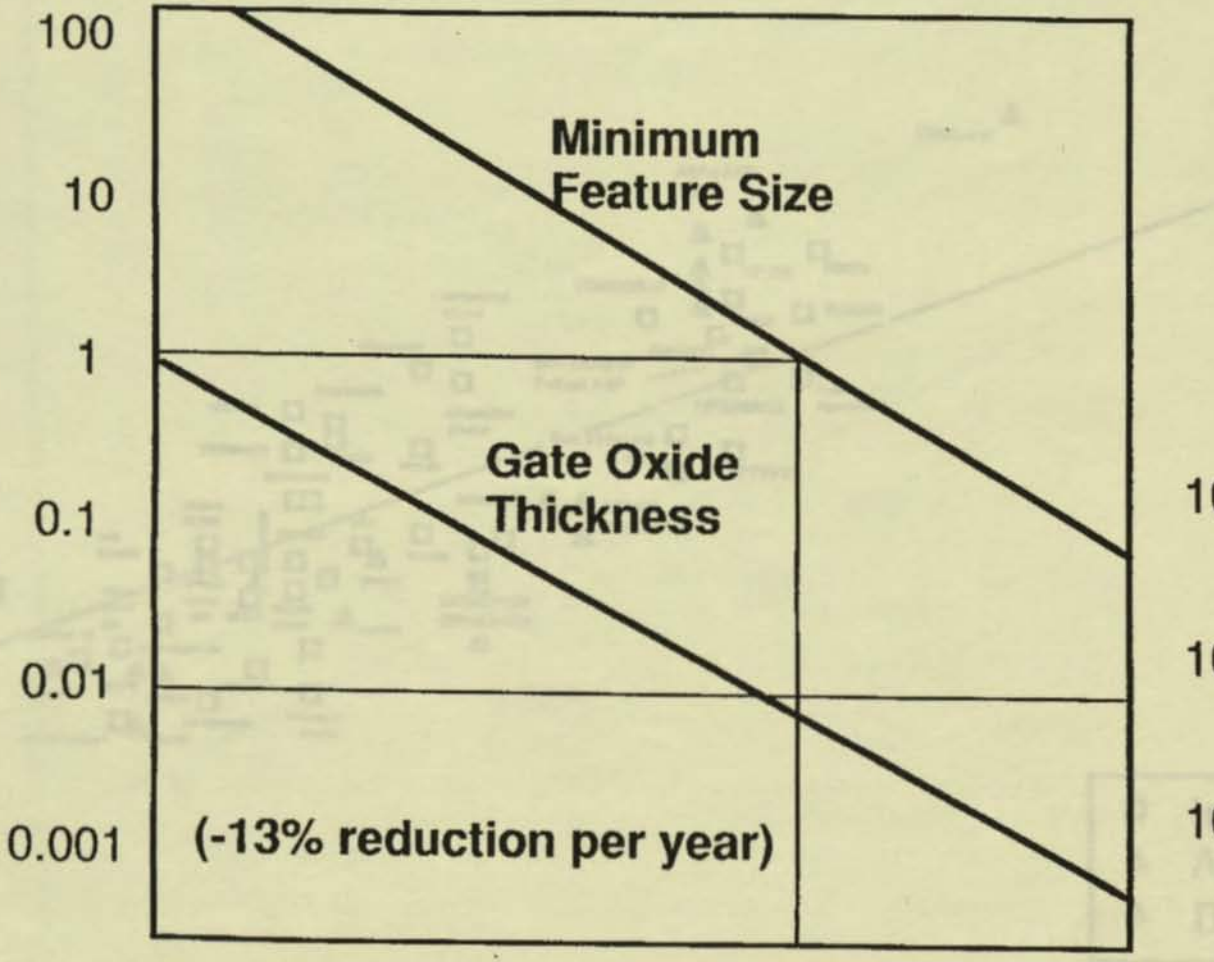

* Microprocessor Report
25 Jan 1993

Scaling of MOS parameters over time (Intel) *

RISC System Performance Trends

Chart 4.5-P

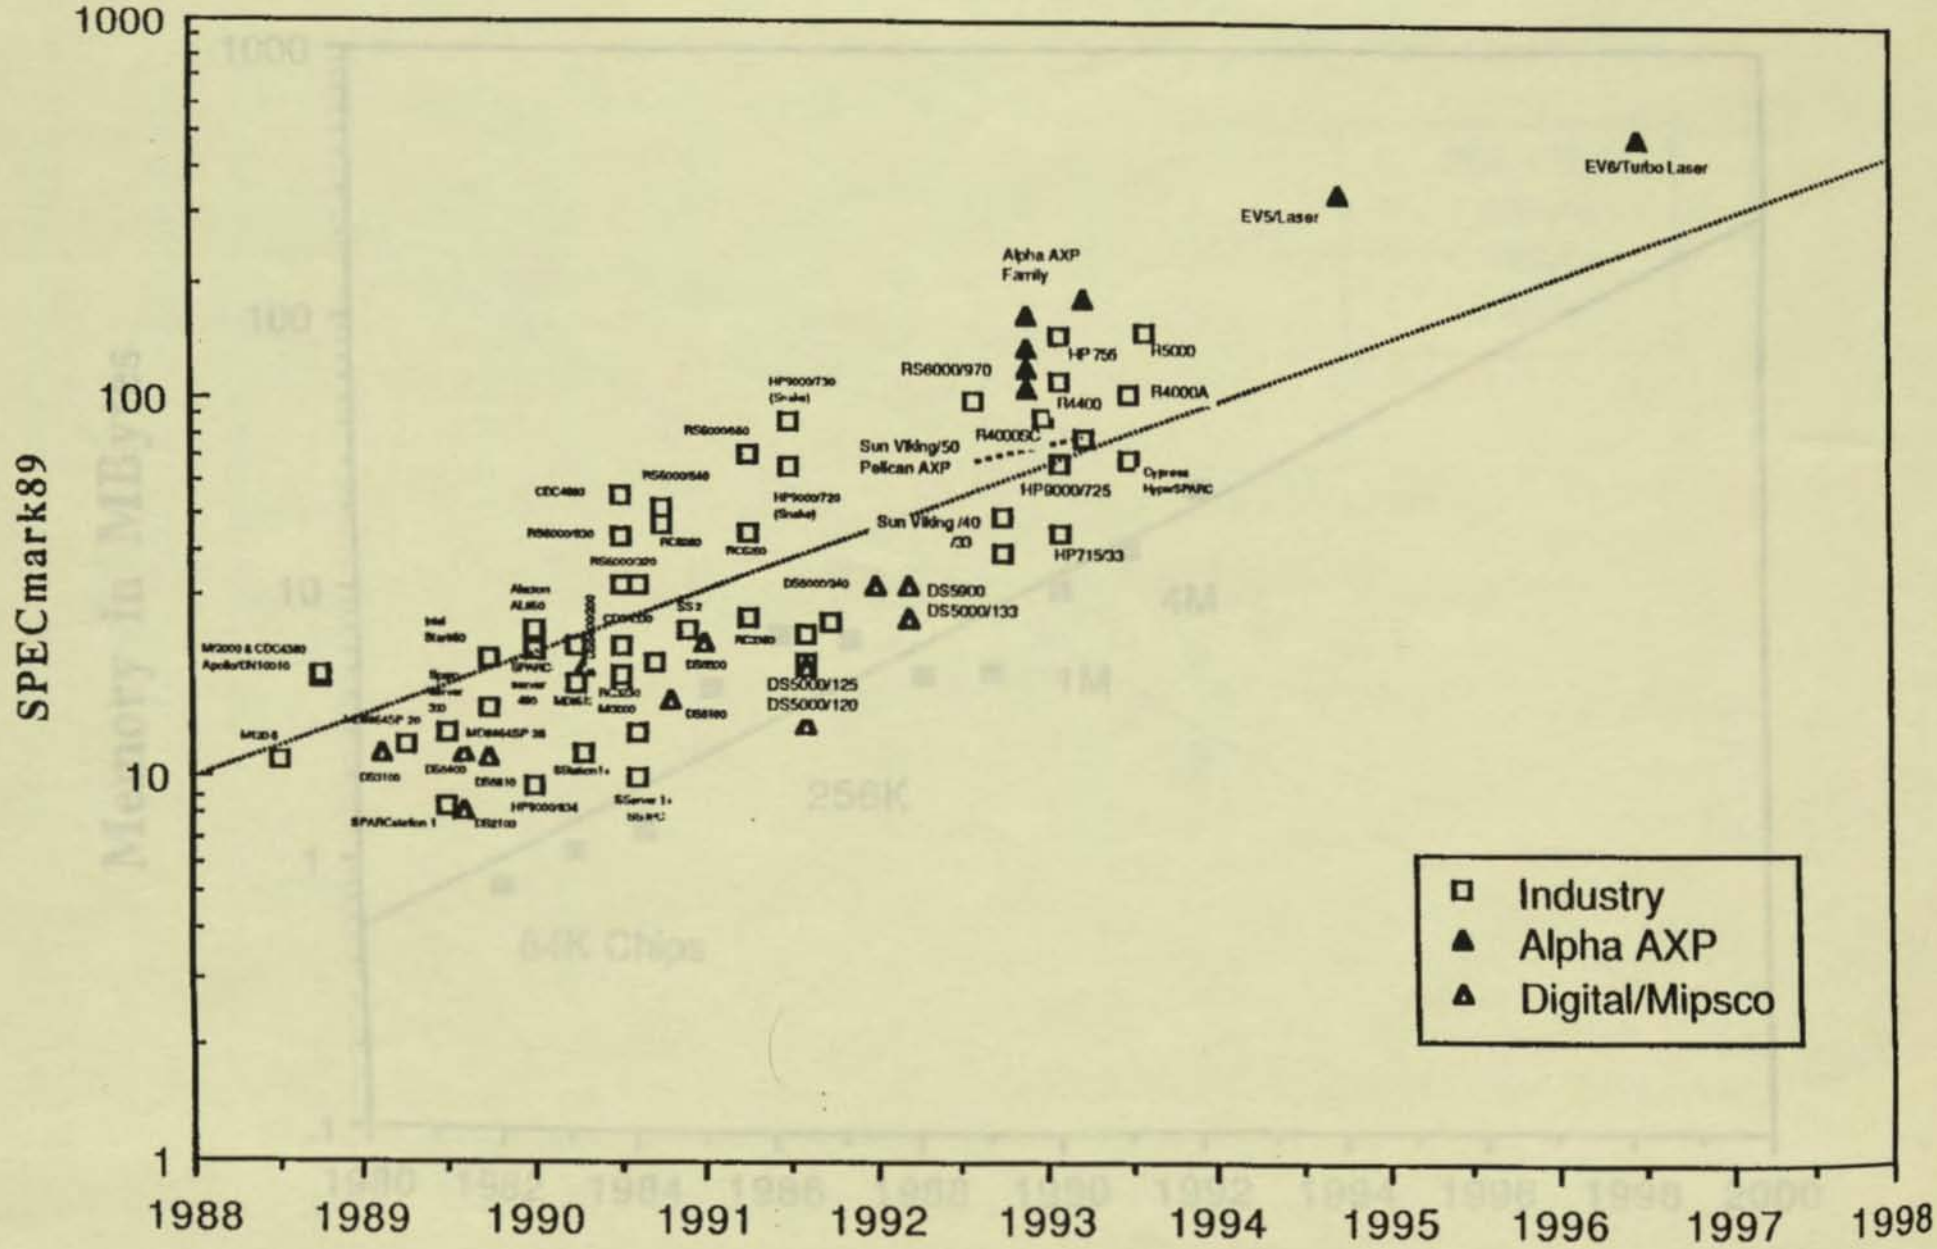

Digital Confidential

Calendar Year of System Availability

Rev: 1-24-93
Data Source: SPEC, B. Supnik, D. Quinby

Memory for $500 with Least Expensive Chips

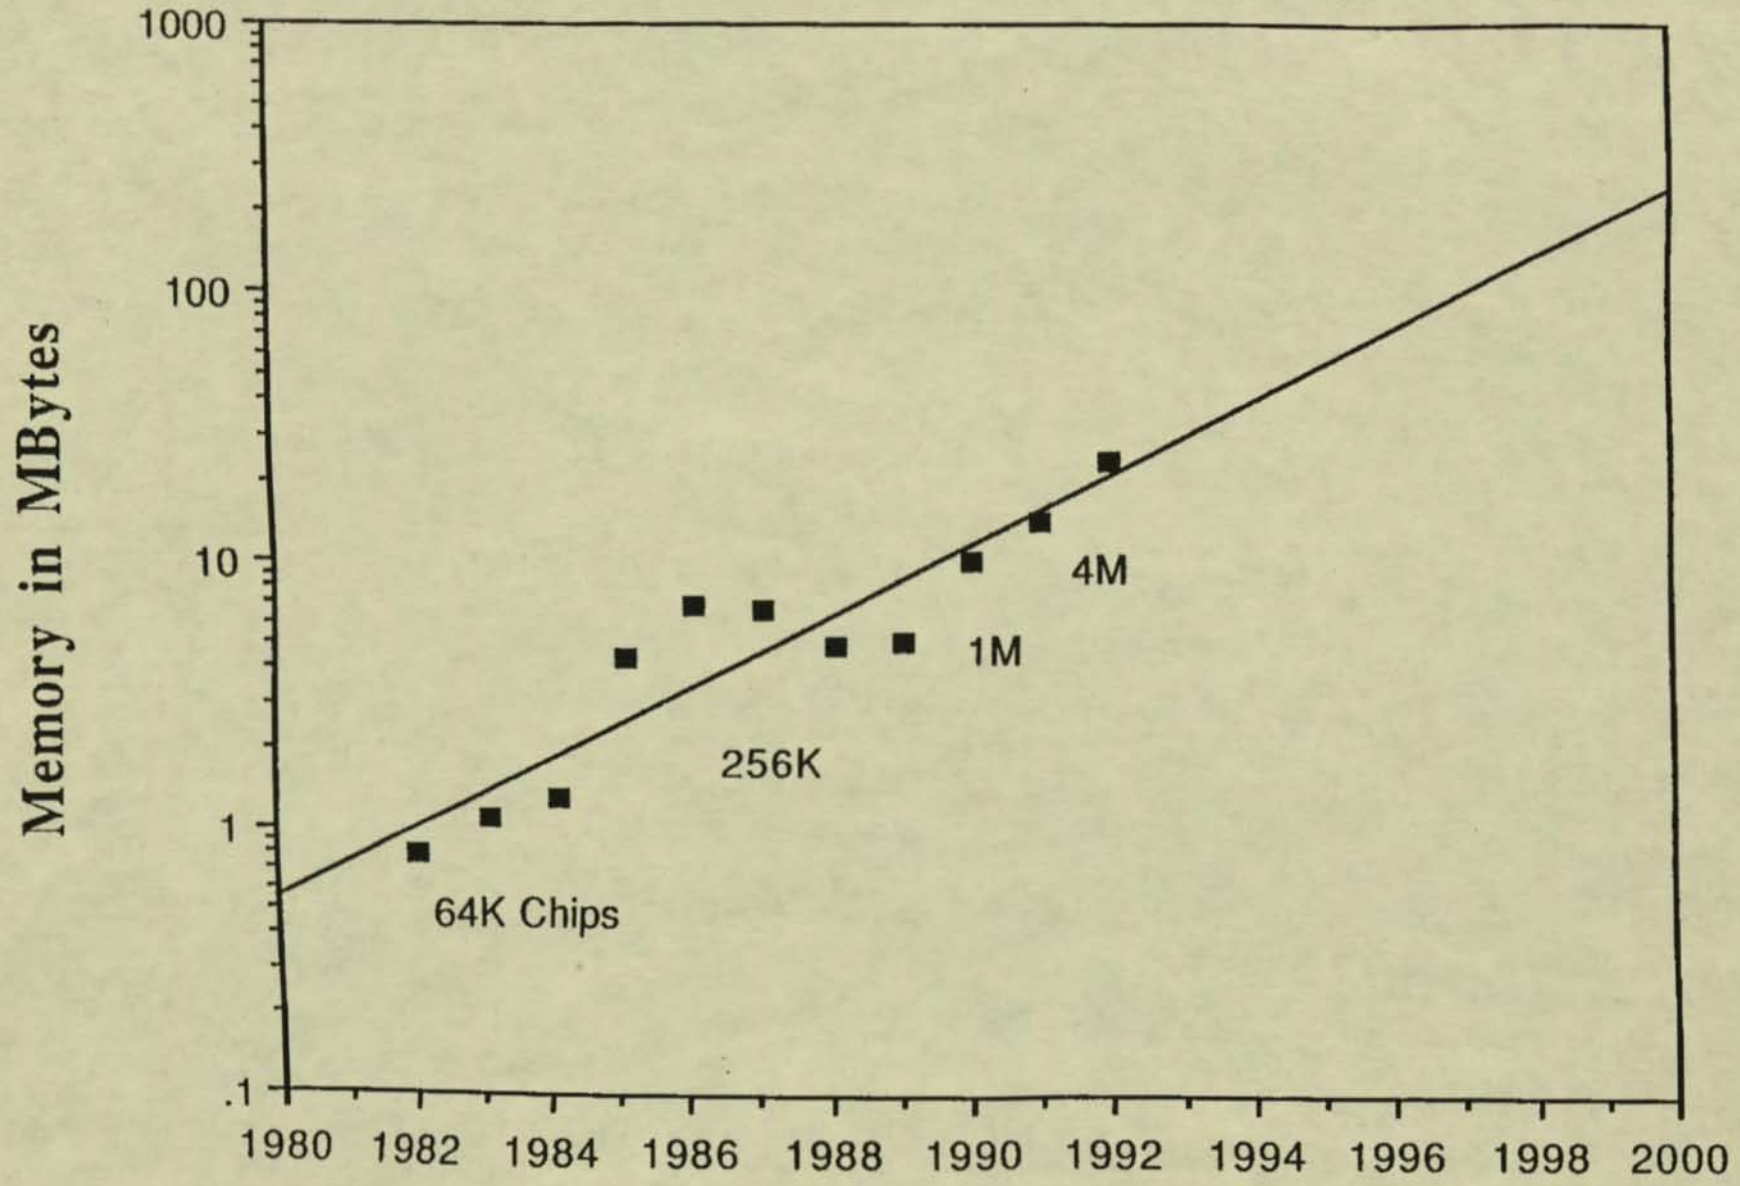

Source: Steve Cullen Date: 29 January 1993

DRAM Price Trend

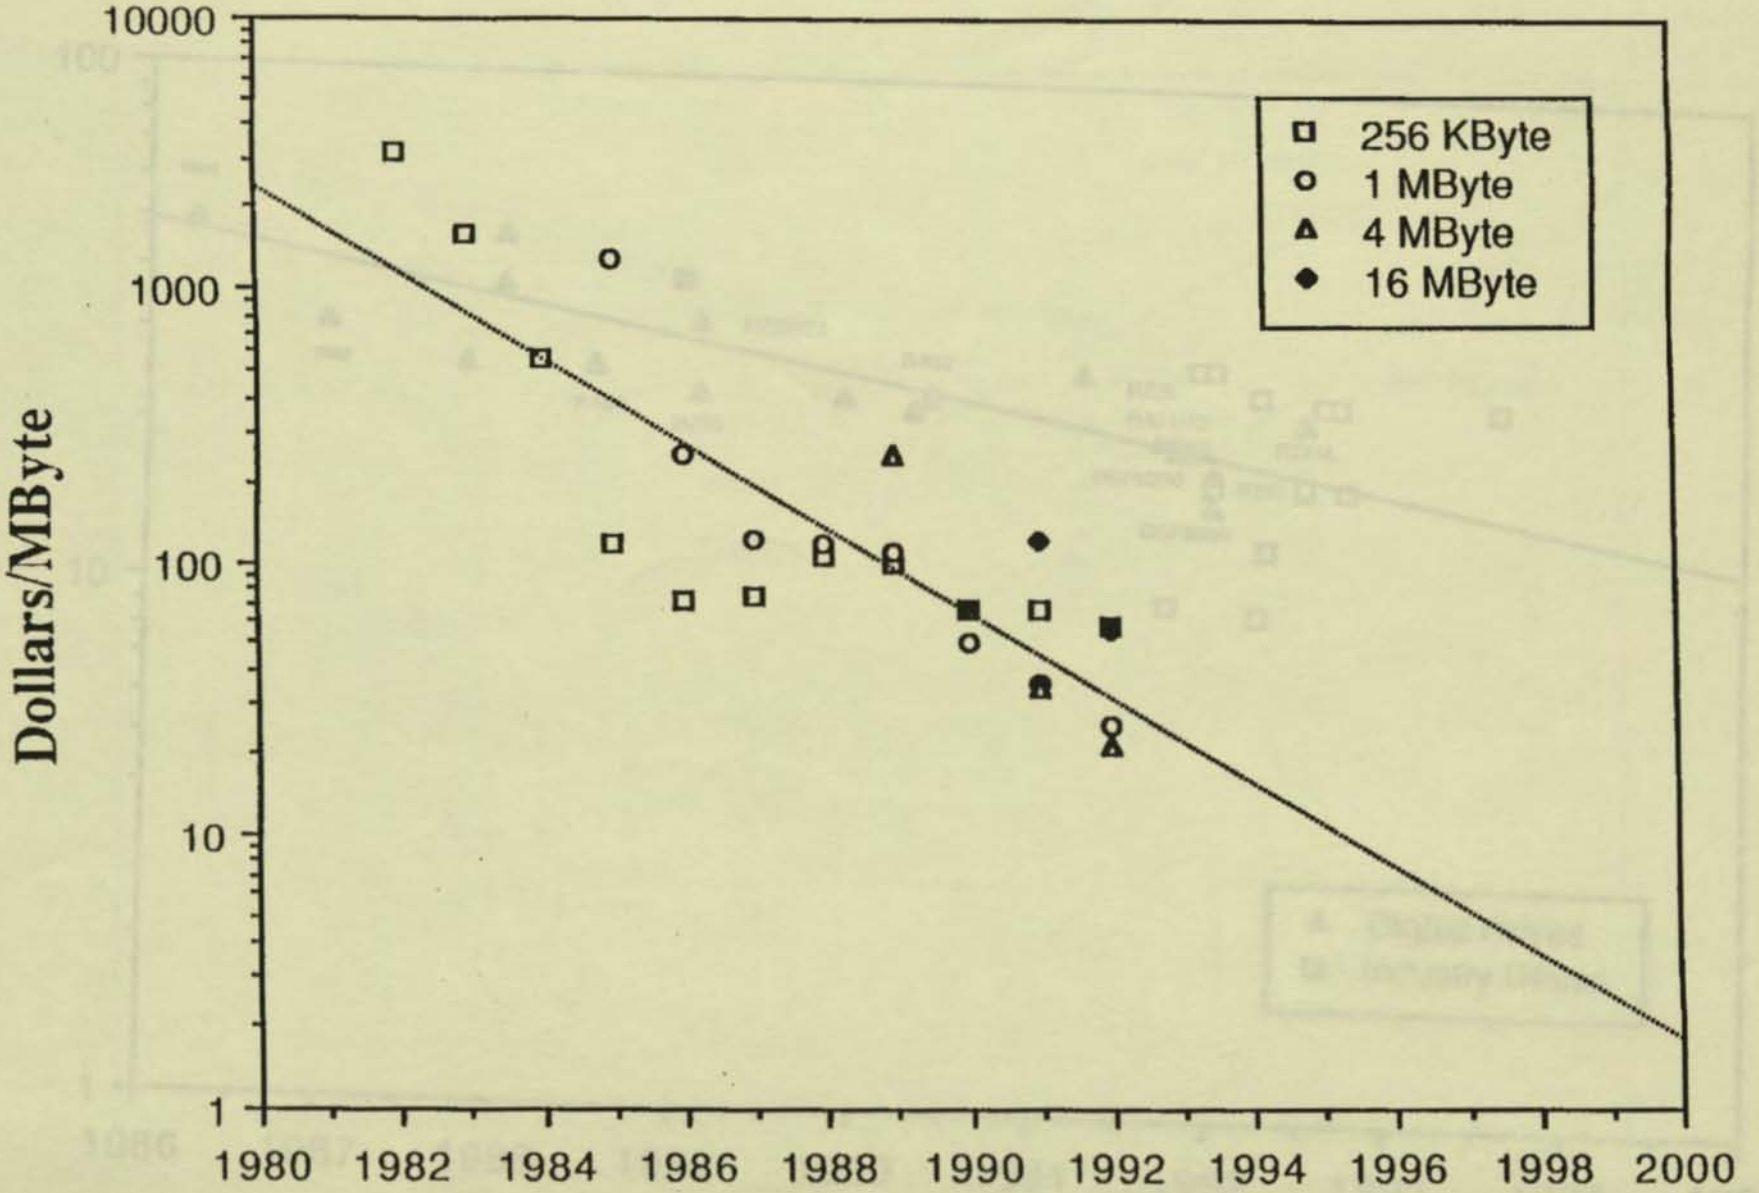

Disk Access Time Trend

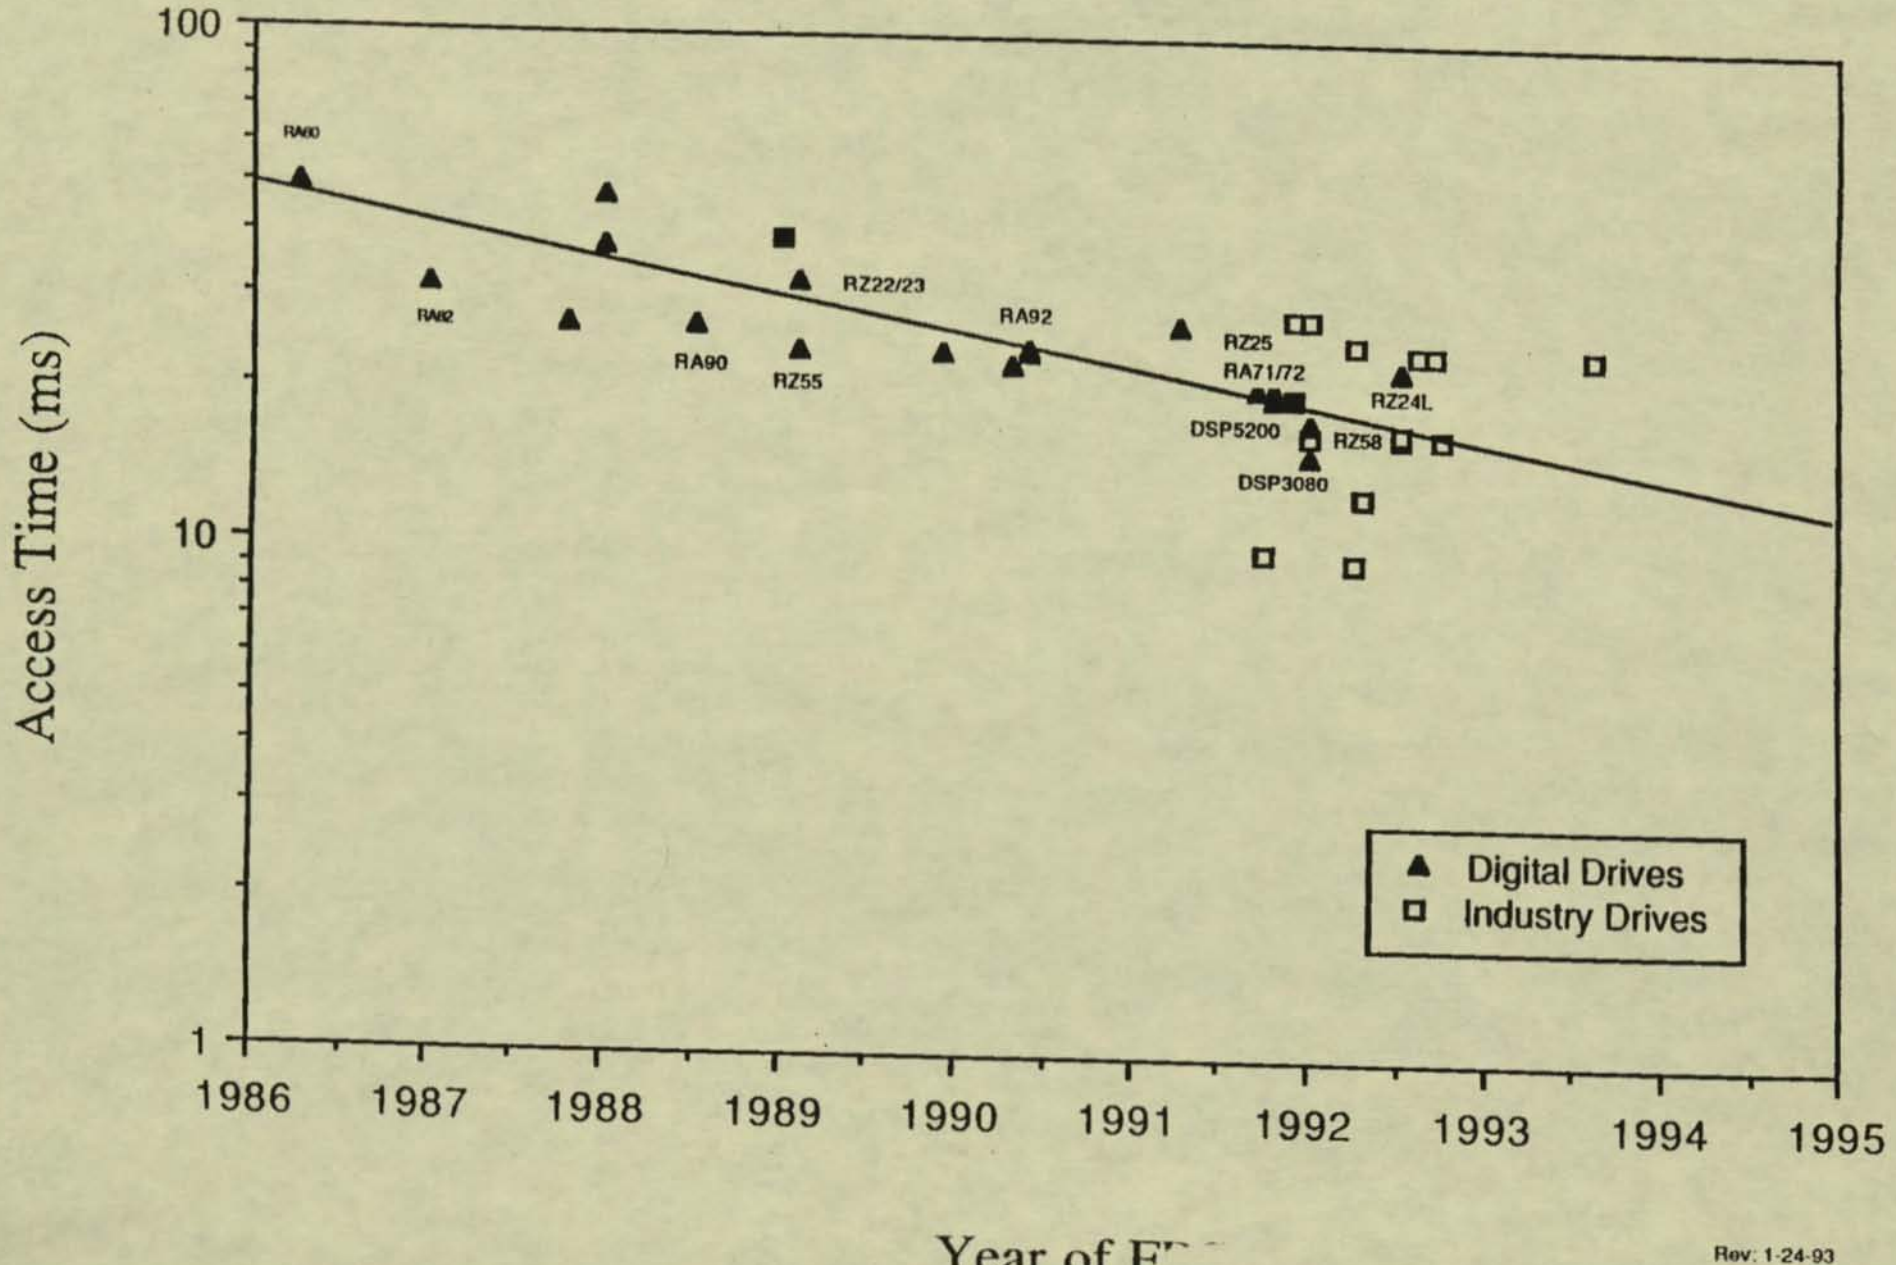

Rev: 1-24-93

Digital's Small Disk Capacity Trends

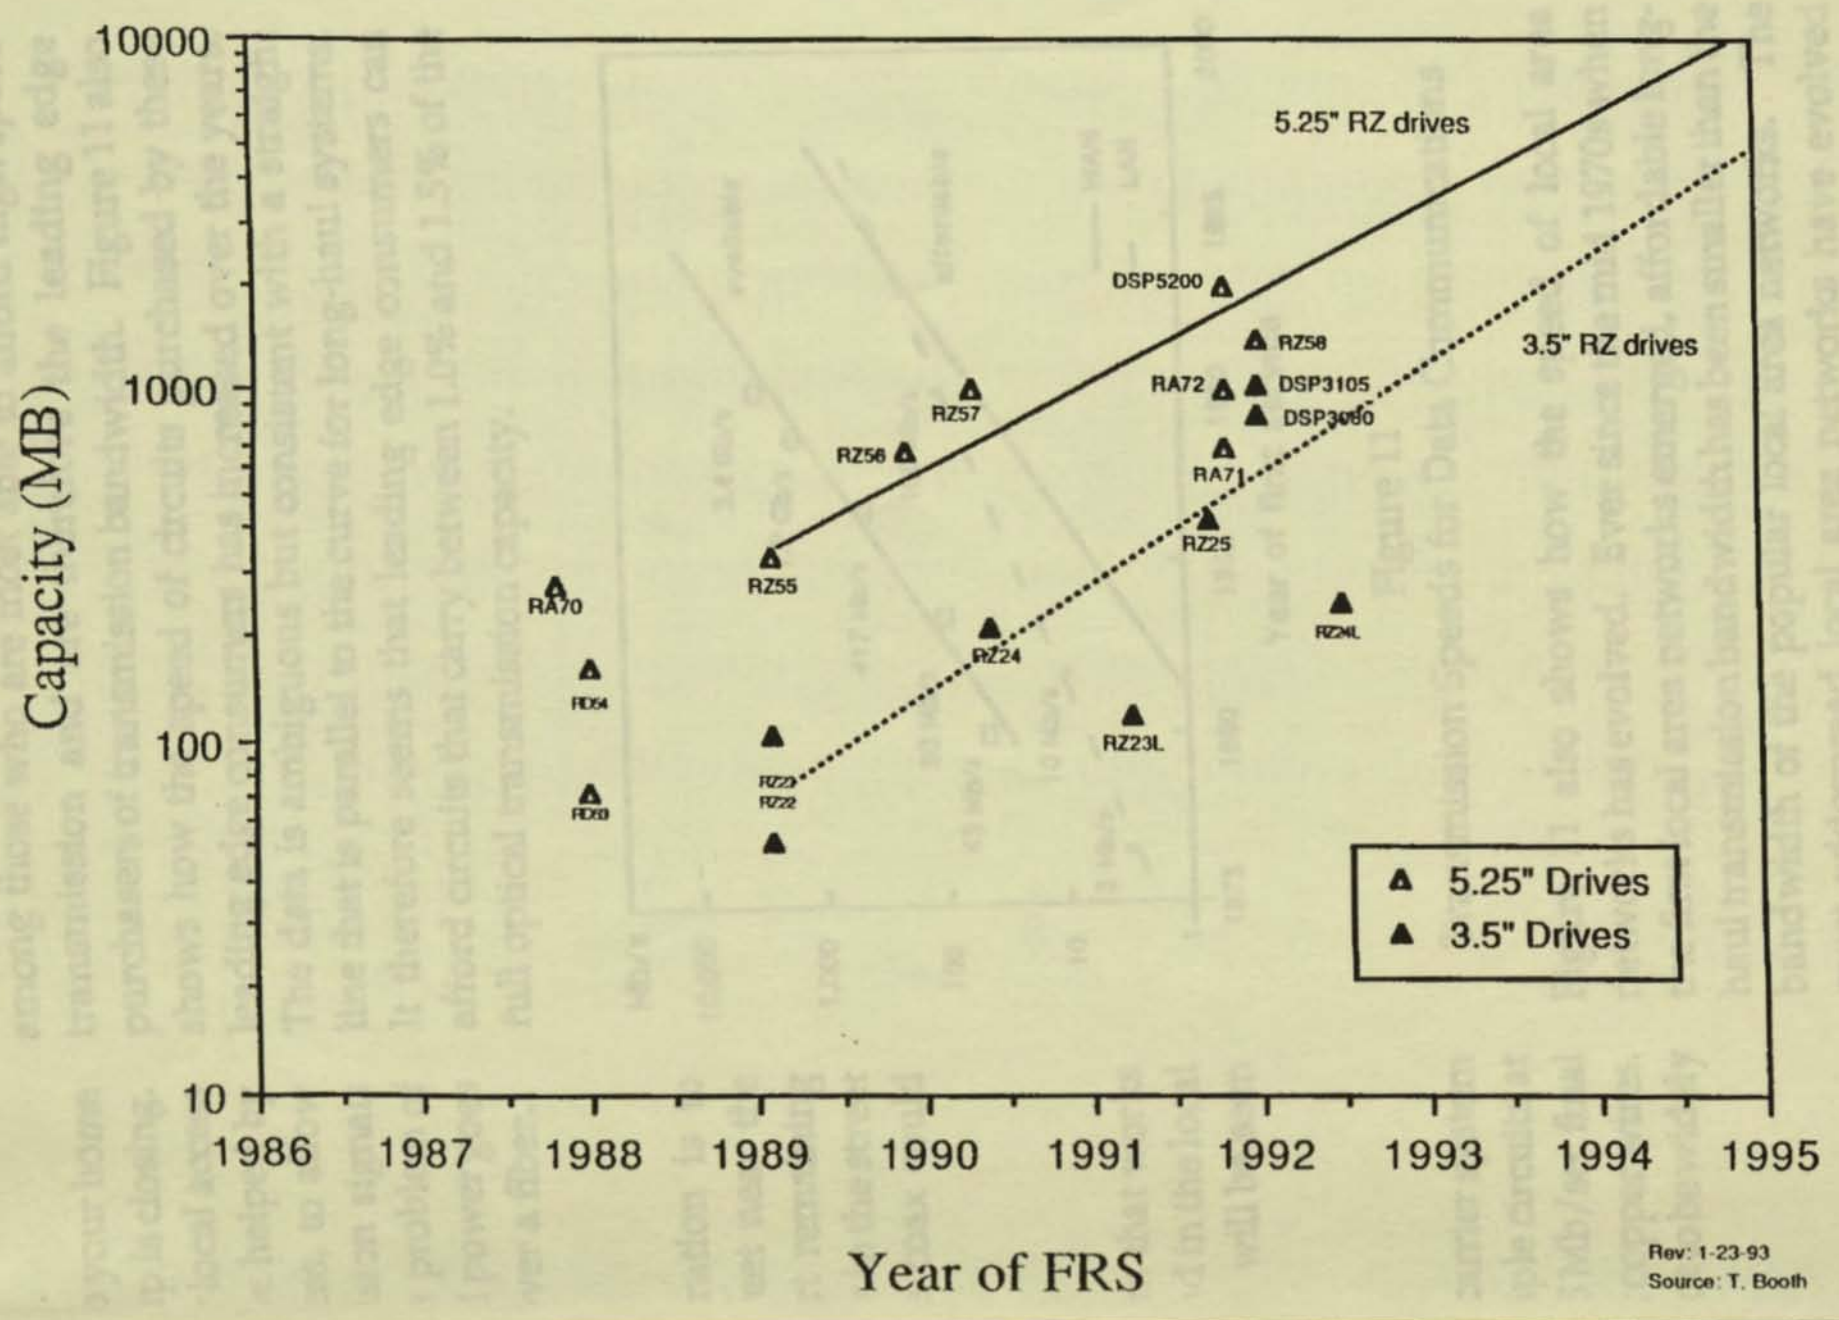

among those who are most able to afford high speed transmission and are therefore the leading edge purchasers of transmission bandwidth. Figure 11 also shows how the speed of circuits purchased by these leading edge consumers has increased over the years. The data is ambiguous but consistent with a straight line that is parallel to the curve for long-haul systems. It therefore seems that leading edge consumers can afford circuits that carry between 1.0% and 1.5% of the full optical transmission capacity.

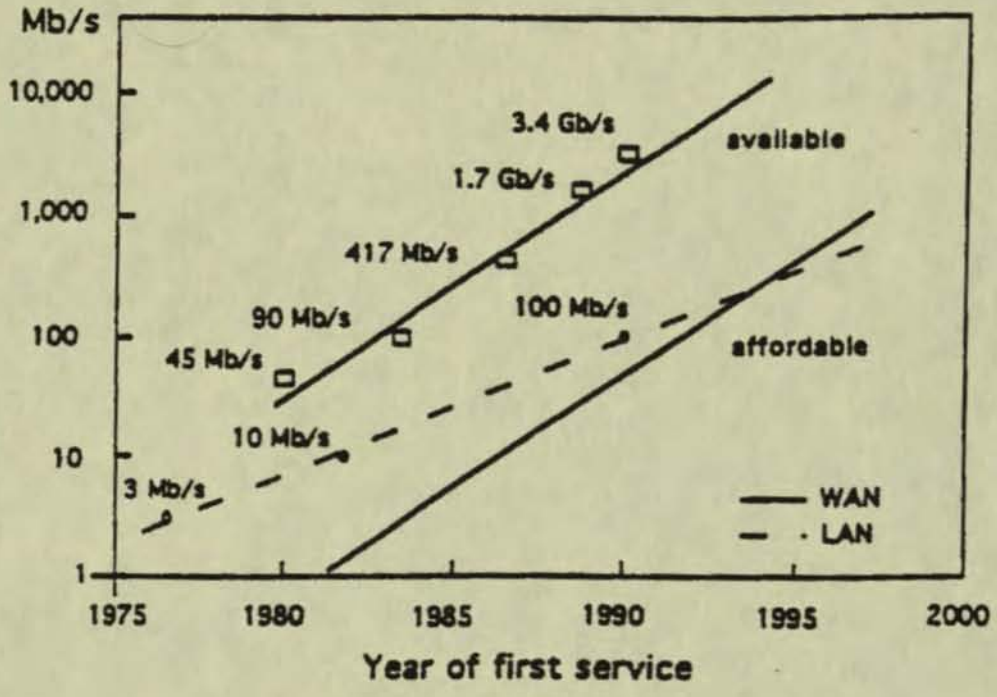

Figure 11
Transmission Speeds for Data Communications

Figure 11 also shows how the speed of local area networks has evolved. Ever since the mid 1970s when the first local area networks emerged, affordable long-haul transmission bandwidth has been smaller than the bandwidth of the popular local area networks. The most widespread local area networks have evolved from Ethernet and it is anticipated that FDDI will set the trend for the next several years. The increase in speed has been about 28% per year, which is more modest

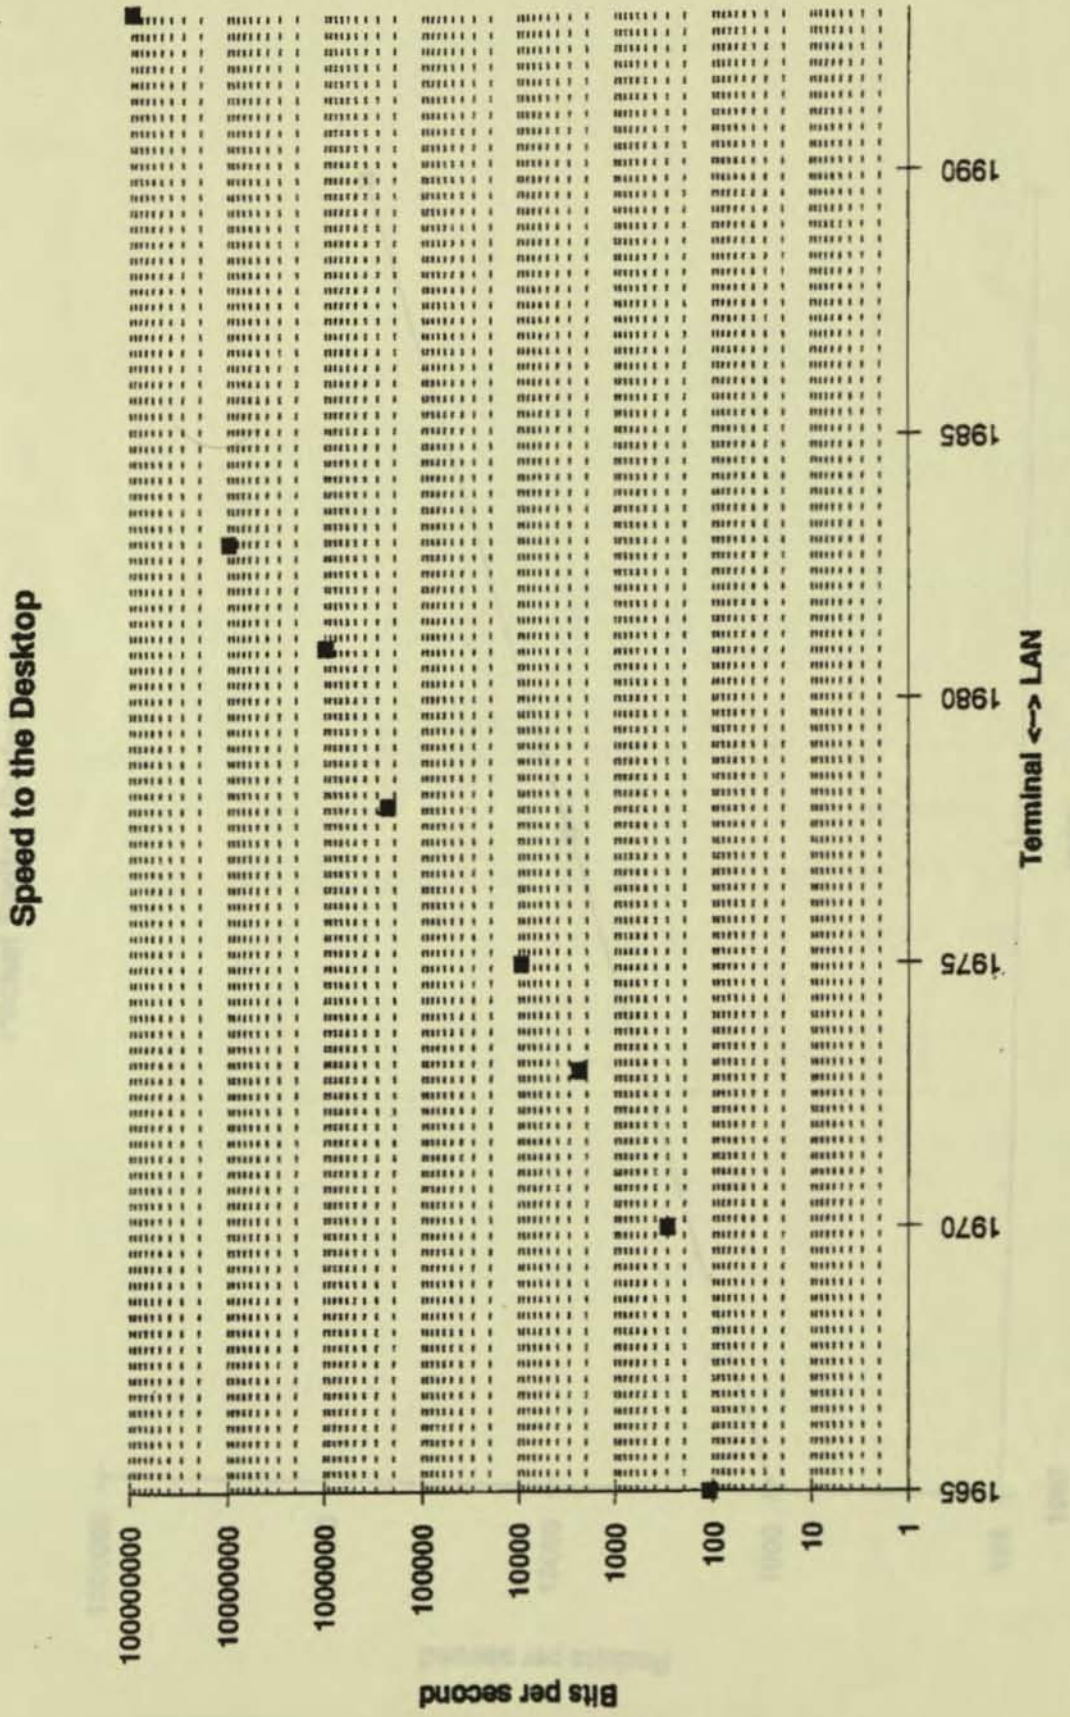
Speed to the Desktop
Bits per second
100000000
10000000
1000000
100000
10000
1000
100
10
1
Terminal <--> LAN
1965
1970
1975
1980
1985
1990

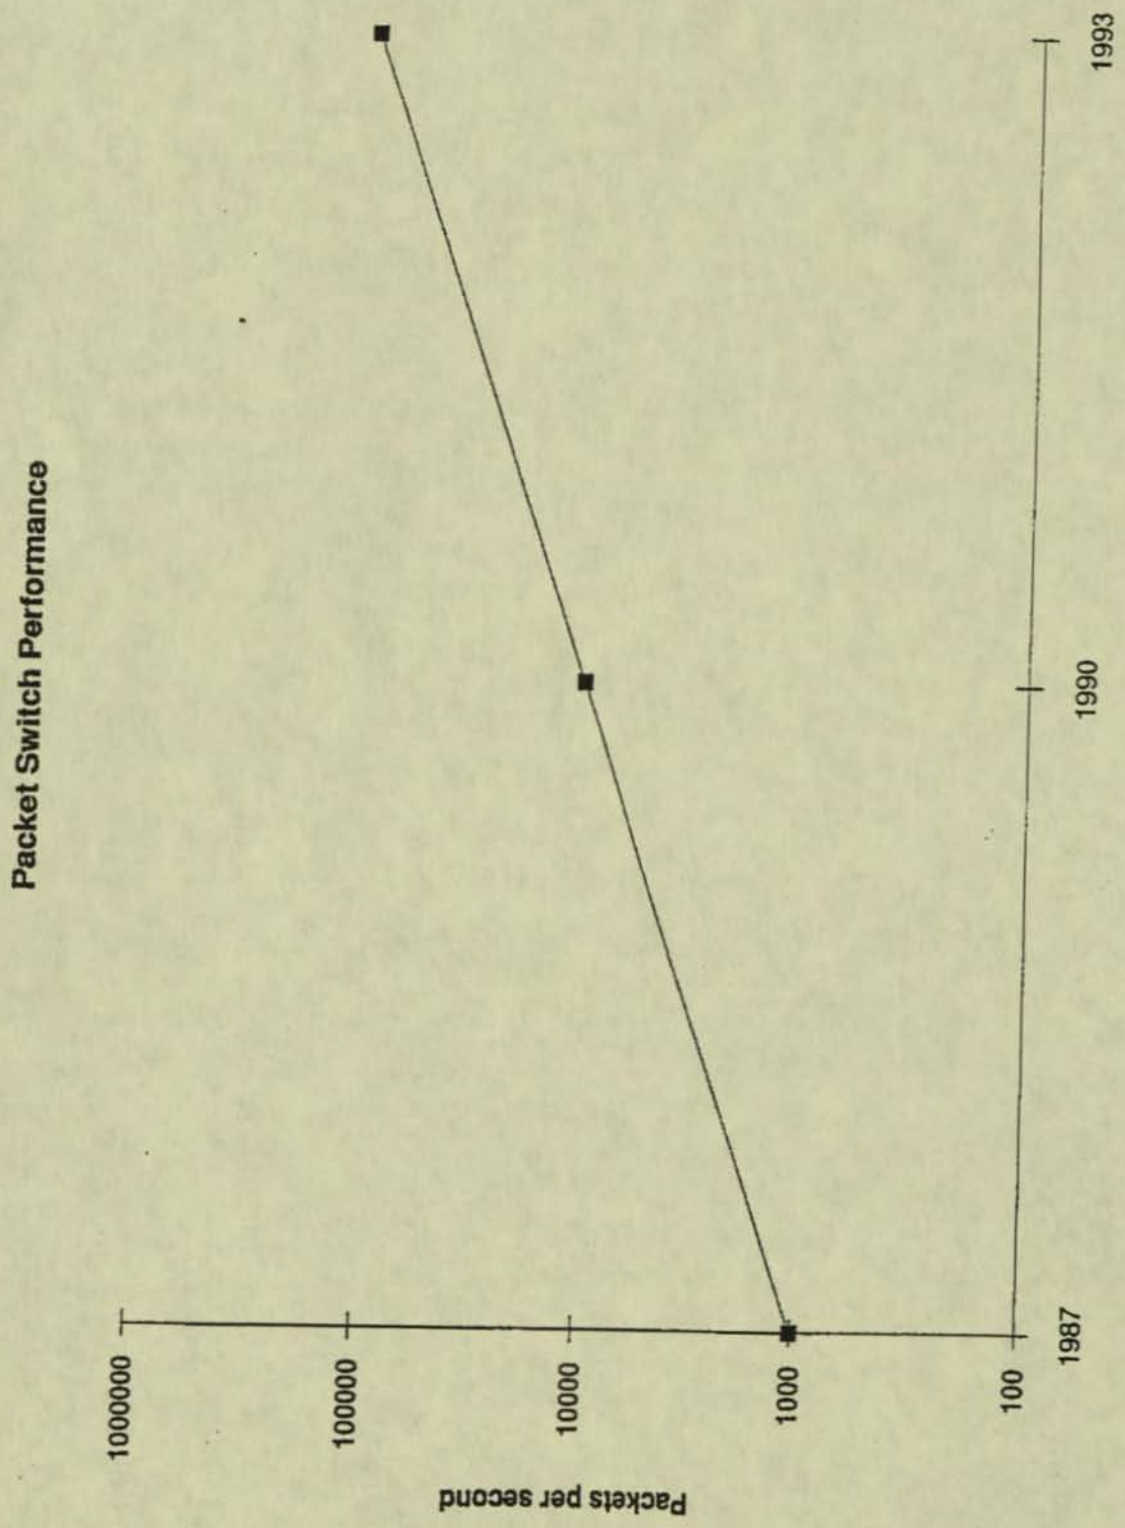
Packet Switch Performance
Packets per second
1000000
100000
10000
1000
100
1987
1990
1993

Number of Internet Hosts (logarithmic)

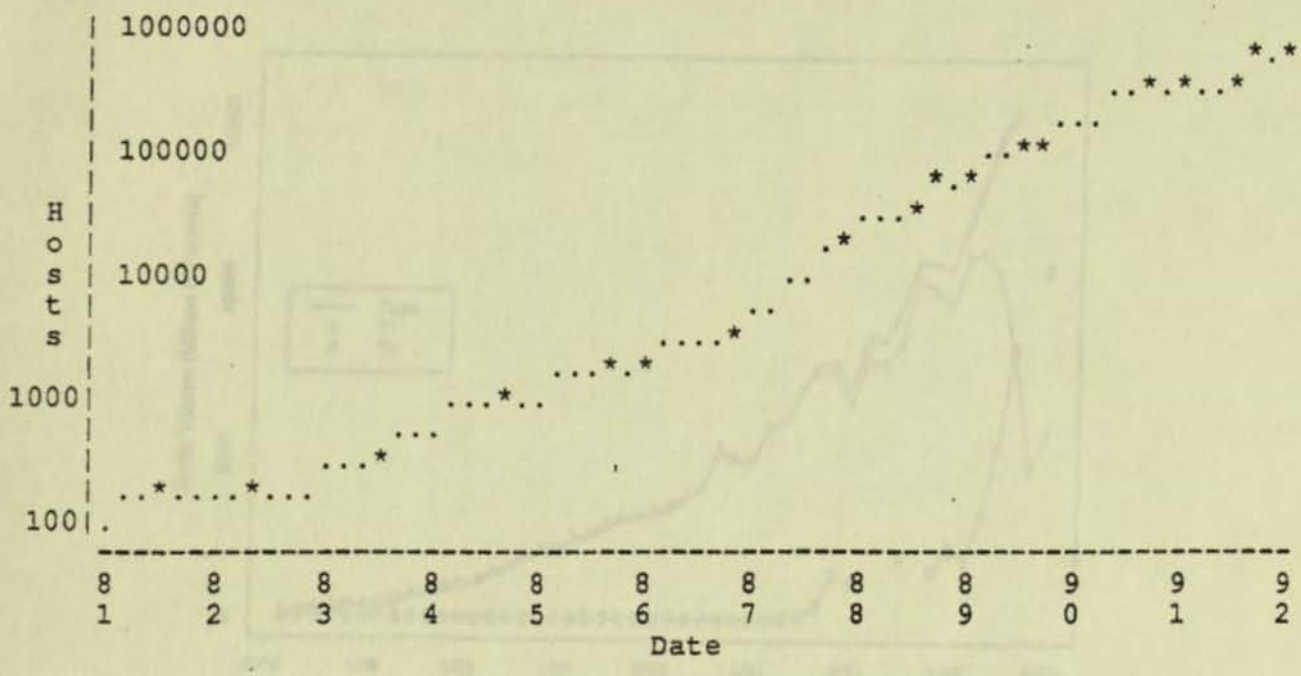

"*" = data point, "." = estimate

This graph is a logarithmic plot of the number of Internet hosts.

N.2 Number of Domains

This chart shows the number of domains existing in the Internet Domain Name System as collected by ZONE.

Date	Domains
07/88	900
10/88	1,280
01/89	2,600
07/89	3,900
10/89	4,800
10/90	9,300
01/91	11,200
07/91	16,000
10/91	18,000
01/92	17,000

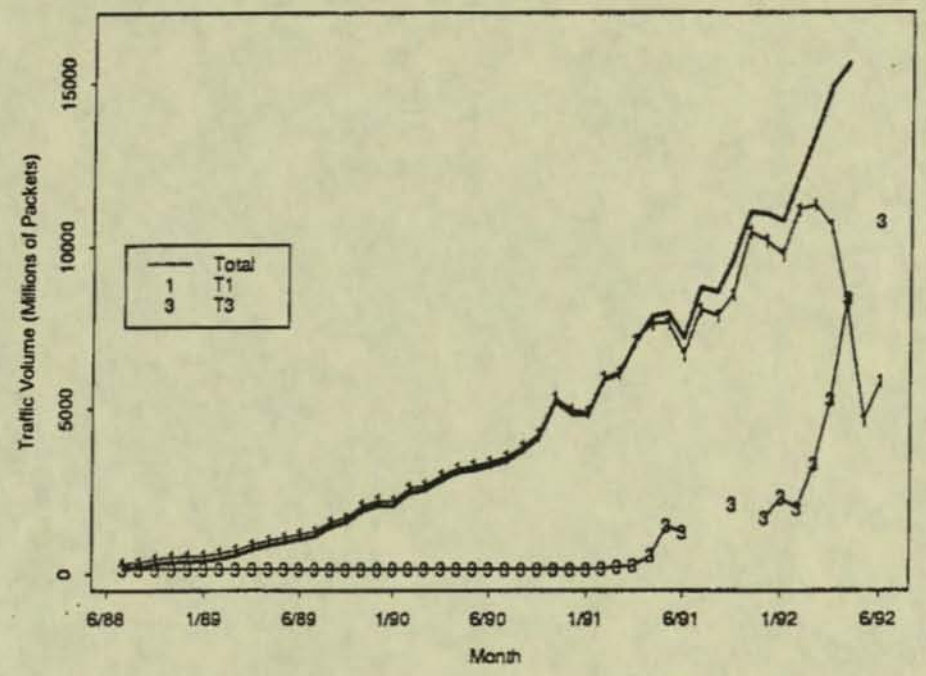

Figure 3: Traffic Volume on the NSFNET Backbone

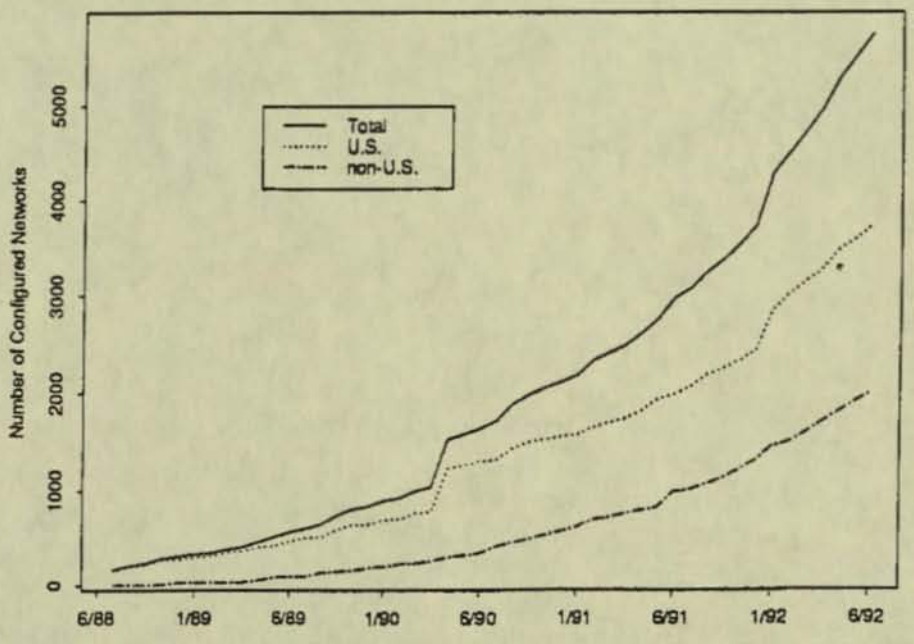

Figure 4: Networks Attached to NSFNET Backbone

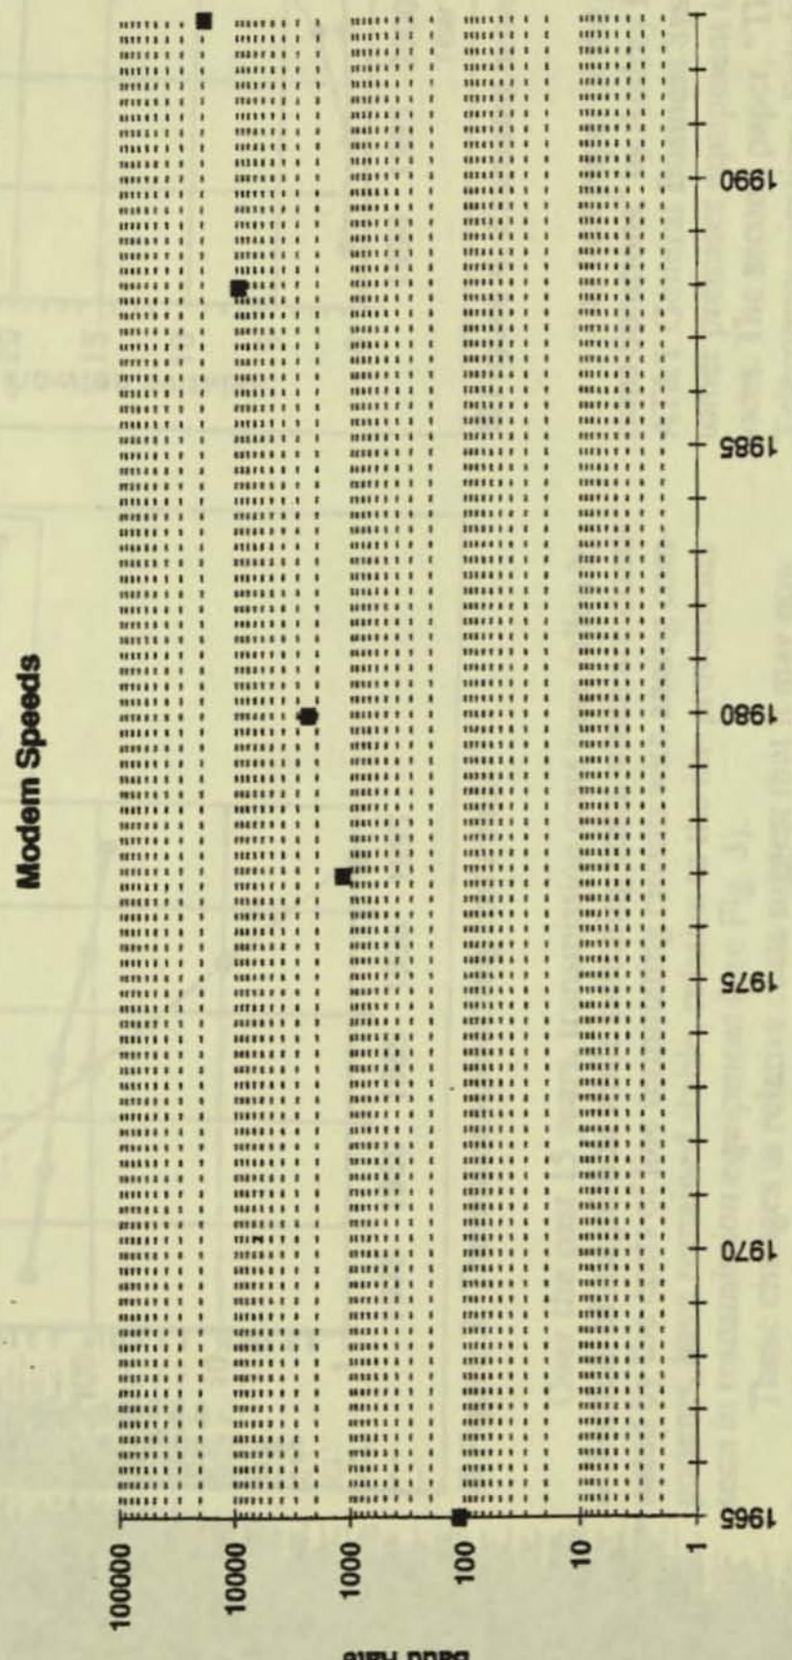
Modem Speeds
1990
1985
1980
1975
1970
1965
100000
10000
1000
100
10
1
Baud Rate

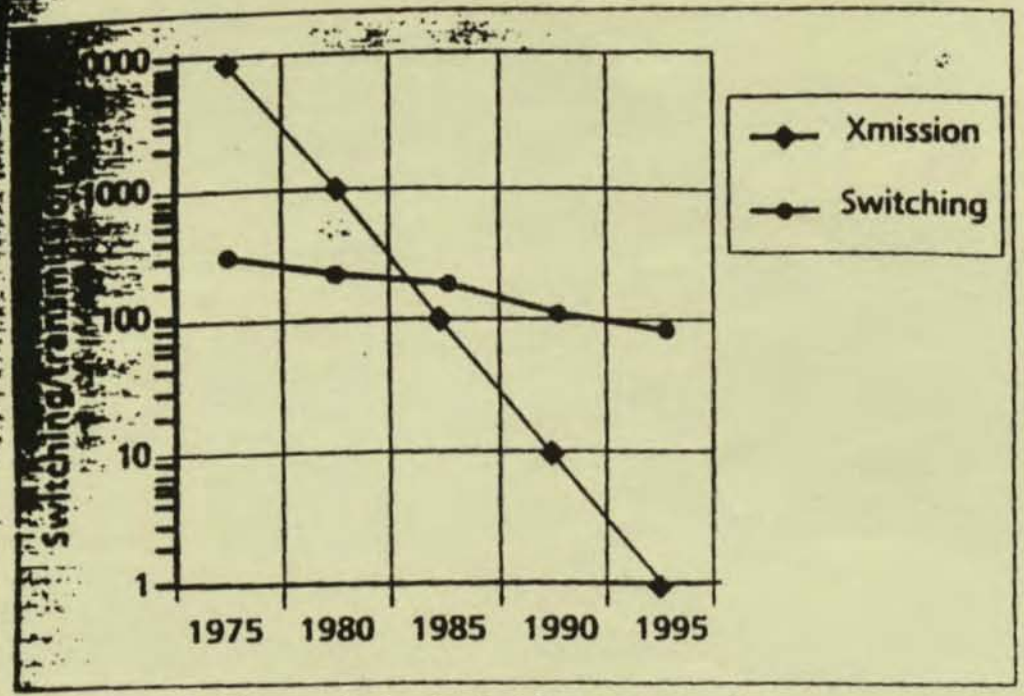

■ **Figure 3.** *Transmission/switching cost trends.*

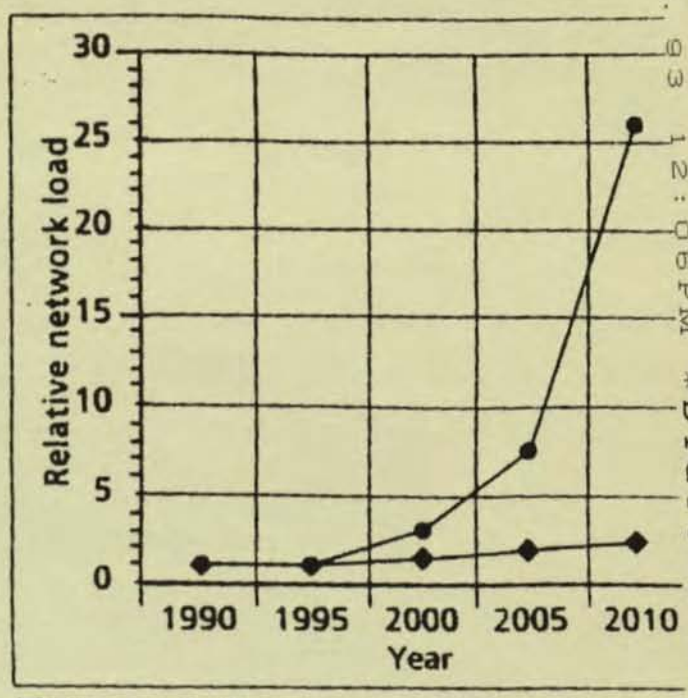

■ **Figure 4.** *Estimated network load.*

Over the past 15 years, transmission costs have fallen by - 3 orders of magnitude while switching costs, which are dominated by the costs of line terminating circuits with analog components, have remained stable, relative to the im- provements seen in transmission equipment (see Fig. 3).

These changes in relative costs suggest that it may now be possible to design telephone networks with fewer switching systems that are much larger than the averages seen today [3].

The dominant service in the 1965 network as well as in today's telephone network is voice. However, this is expected to change dramatically over the next few years, (see Fig. 4, which is based on data that appears in Ref. [4]).

The projection is that data and other high bandwidth traffic will surpass voice by 1995. It is commonly assumed that voice has a 3-3-3 requirement: a fixed bidirectional channel

The first paper of the group, "Optim tions Solutions," by Audrey F. Burson and AT&T General Business Systems, discusses tomer premises equipment in the deploym vices. The second paper, "The Evolving Rc Telecommunications Switching," by Stewai Bellcore, examines the impact of recent tecl the capabilities of switching systems. Irving Richards of Bell Northern Research take a spective in their paper, "Technology's Role tems Evolution." The last paper of the grou Switching System Requirements," by Toshi describes the kinds of switching systems tha future network from a *network* operation a

(F) 05

Digital Information Infrastructure Utility

Status and Plans

George Champine

Information Management and Technology

Why We are Here

Opportunity:

- IM&T is implementing Infrastructure Utility
- Clear vision of what and how
- Take $200M-$300M per year to bottom line
- Generate a few $B per year incremental sales
- Showcase and best early adopter of products and services

What we need:

- Partnership with IM&T as strategic initiative
- Consider IM&T needs for products and services
- Unify strategies and priorities
- Communicate overall strategic product release plan

Objectives of Infrastructure Technology Program

- Improve the competitive position of Digital
- Develop a broadly-based program vision, business objectives, and supporting technology that supports those business objectives
- Pilot key technology elements of infrastructure to validate feasibility
- Develop supporting practices and procedures related to infrastructure technology
- Coordinate deployment of technology

Dual Digital Role

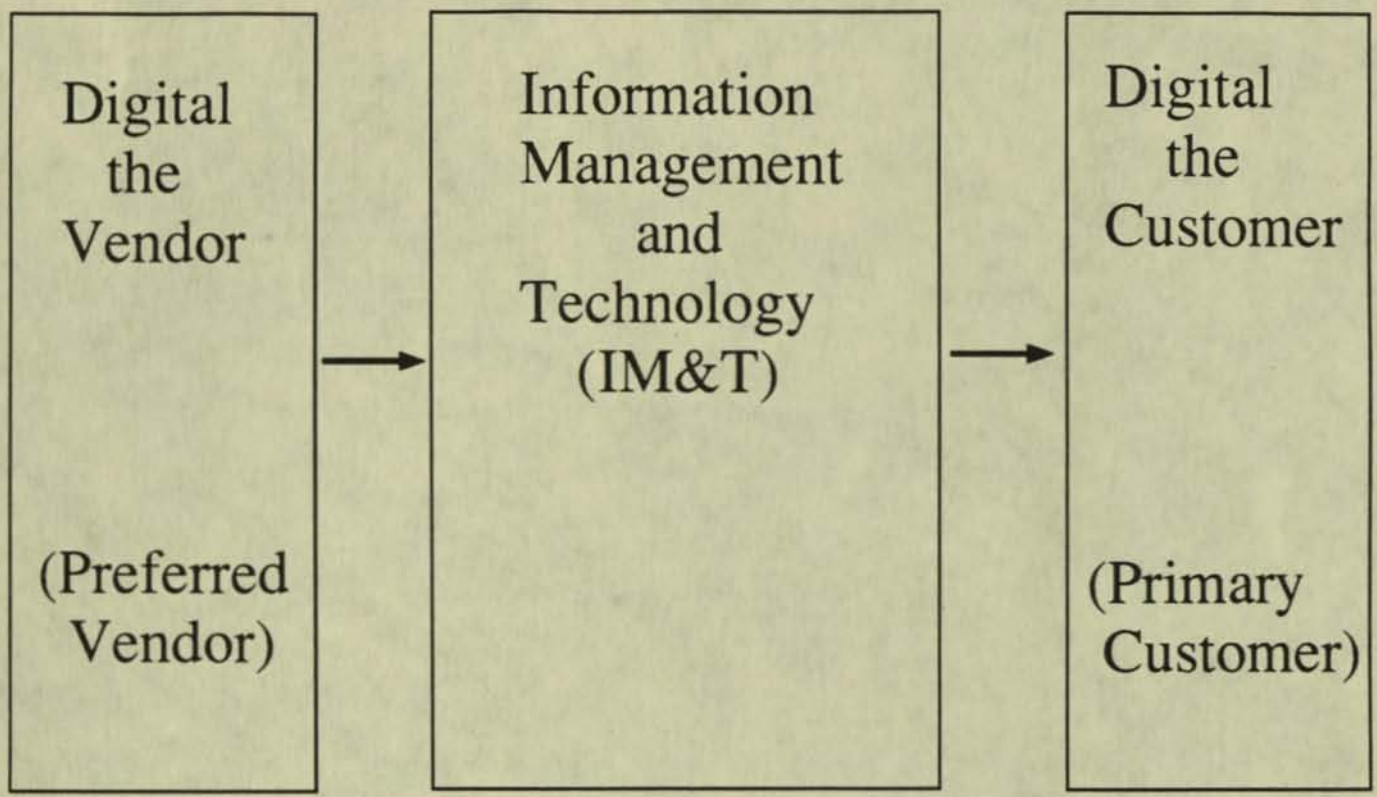

Premier Vendor of distributed Systems

Datamation No. 52 computer Company

Fortune 37 Company

Current MIS Status

- Large Systems
  —7000+ MIS employees
  —500 Computer Centers
  —71,000 node world wide network
  —$10B invested in applications S/W

- Based on VAX/VMS/DECnet/Terminal timesharing model

- Corporate infrastructure is the network

- Businesses own computing

- Organized into four geographies

- Successfully managed 30% growth/year

- Supports 114,000 employees worldwide

IS Infrastructure Utility Strategy

Significantly improve Corporate profitability, productivity and business capability through the establishment of an enterprisewide IS infrastructure, managed and operated as a utility:

- ☐ Seamless, distributed computing environment using client-server model
- ☐ Efficient, integrated, standardized operational processes, practices and tools
- ☐ "Best-in-Class" model (cost, quality, reliability, functionality)
- ☐ Enables business added-value, enterprise services and commodity desktop choices
- ☐ Leverage opportunity for revenue achievement
- ☐ Interoperate with vendors and customers

Industry-wide Assumptions for 1995

Hardware

- Low end desktop = 25 MIPS
- High end desktop = 400 MIPS
- Multimedia technology will be common
- Laptops everywhere
- Hardware cost will drop 20%/year

Software

- OSF/DCE will win distributed systems
- DOS/NT will dominate non-technical desktops
- UNIX will dominate engineering desktops
- MAC/OS will dominate multimedia

Systems

- Eight times more data by year 2000
- Desktop replaces computation center
- Glass house distributed to glass closets
- Systems increasingly heterogeneous
- Standards increasingly important

IS

- Labor cost will dominate (>80%)
- Rapid access to information crucial for real time business decisions and corporate management control
- Information systems complexity a major problem

The Infrastructure Utility

VISION

The Infrastructure Utility provides customers with information in their preferred *form*, from any *source*, using any *desktop device*.

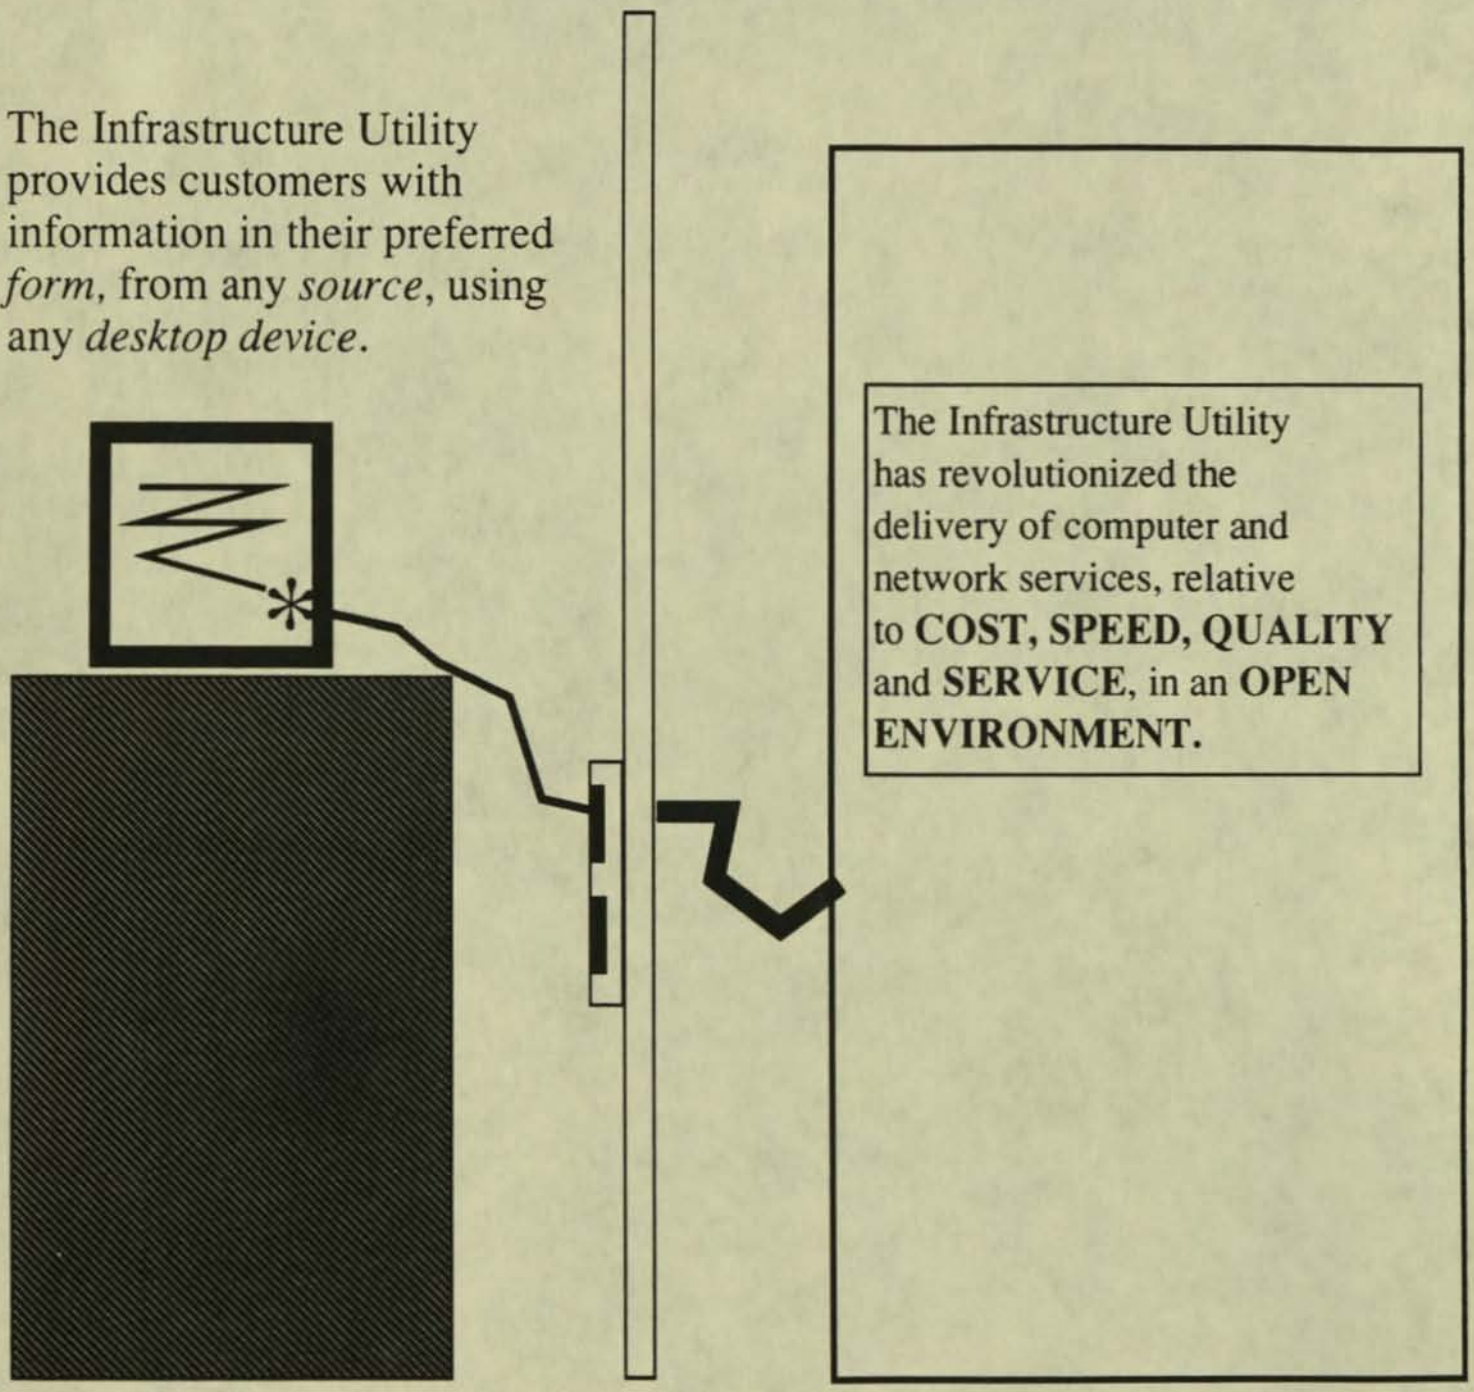

The Infrastructure Utility has revolutionized the delivery of computer and network services, relative to **COST, SPEED, QUALITY** and **SERVICE**, in an **OPEN ENVIRONMENT.**

Information System Components

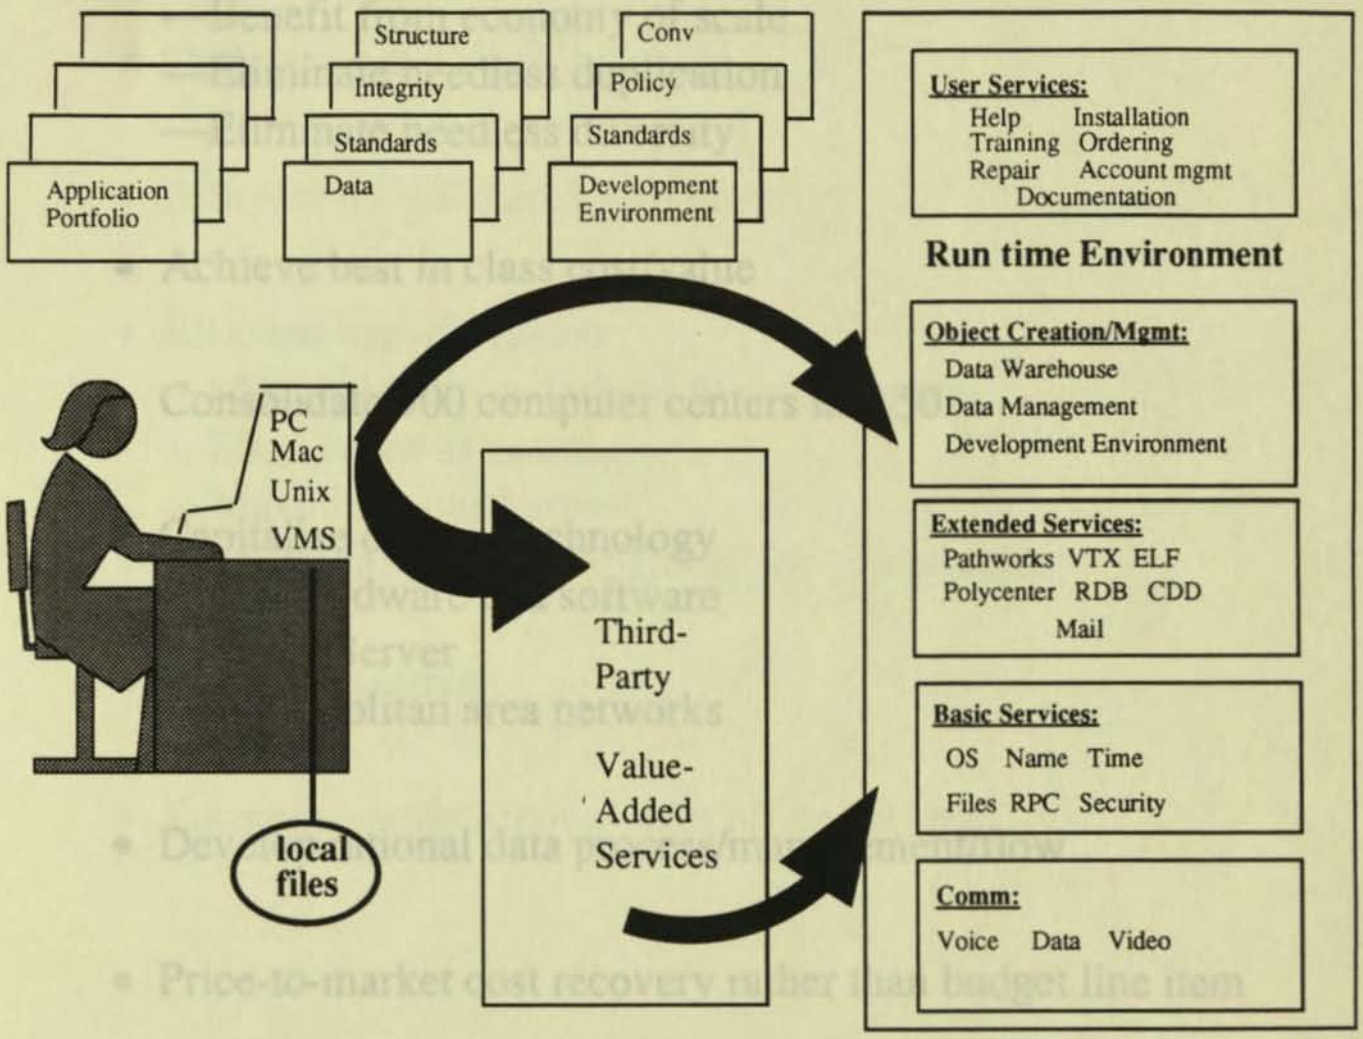

Desired Endpoint

- Reduce cost and headcount by 30-40%
  —Consolidation
  —Benefit from economy of scale
  —Eliminate needless duplication
  —Eliminate needless diversity

- Achieve best in class cost/value

- Consolidate 500 computer centers into 50

- Capitalize on new technology
  —PCs hardware and software
  —Client/Server
  —Metropolitan area networks

- Develop rational data process/management/flow

- Price-to-market cost recovery rather than budget line item

- Get businesses out of infrastructure computing

- Get users to accept their share of responsibility

Why Client/Server

- Dis-economy of scale in computation
- Desktop applications software
- Efficient use of shared resources
- Efficient use of desktop
  —Minimal files on desktop to save storage
  —Files paged as necessary
  —Maximum swap area
  —Maximize available memory for computation
- Easy management
- Enterprise-wide account with off-net access

Why Open Systems

- Source of recent leading edge technology
- Necessary to interoperate with customers and vendors
- Necessary for portability
- Broad range of interoperable and flexible products
- Low cost and very competitive
- Main stream computing environment
- Investment protection
- Does the job

Senior Executives for Open Systems (SOS)

Members

- Dupont
- General Motors
- American Airlines
- Motorola
- Northrop
- Boeing

Standards

- OSF/DCE/DME
- XPG
- POSIX
- KERBEROS
- X Windows/Motif
- ANSI languages
- ANSI SQL
- ODA, SGML
- CCITT Group 3/4 fax
- PHIGS/PEX
- X.400
- X.500
- TCP/IP, SNMP

Business Principles

- Utility sells standard portfolio of services to users based on market demand
- Users buy computing services from Utility
- Utility can buy outside
- Users can buy outside
- Inside and outside customers
- Optimized at Enterprise level
- Cost recover instead of budget line item
- Utility uses supported products
- Users use; managers manage
- Utility measured by external benchmarks

One Distributed Architecture Based On DCE

- One security service
  - -- industry standard APIs
  - -- ACLs
  - -- authentication
  - -- scaleable
- One name/directory service
  - -- DNS and X.500 interoperation
  - -- scaleable (multilevel)
- One file system
  - -- modern features
  - -- supported by VMS and Ultrix/OSF
  - -- POSIX compatible
  - -- record access
- One RPC
  - -- hides protocol
  - -- hides naming
  - -- hides security
  - -- hides failures/reconfiguration
- Middleware uses DCE
  - -- mail
  - -- VTX
  - -- ALL-IN-ONE
  - -- Pathworks

DCE Has Great Strategic Value

- Best available approximation to IM&T needs for scaleable distributed client/server system run time engine
- Endorsed by members of OSF and Unix International (500)
- Endorsed by Senior Executives for Open Systems (SOS)
- Value much greater than sum of parts because of integrated reuse/multiple use of modules (matched set)
- Opportunity for Digital to be first with the best DCE implementation and support
- Only game in town

Champine
[.decwrite]sys.doc

Typical Large Cell Configuration

- 10,000 users
- Multiple building campus
- One manager

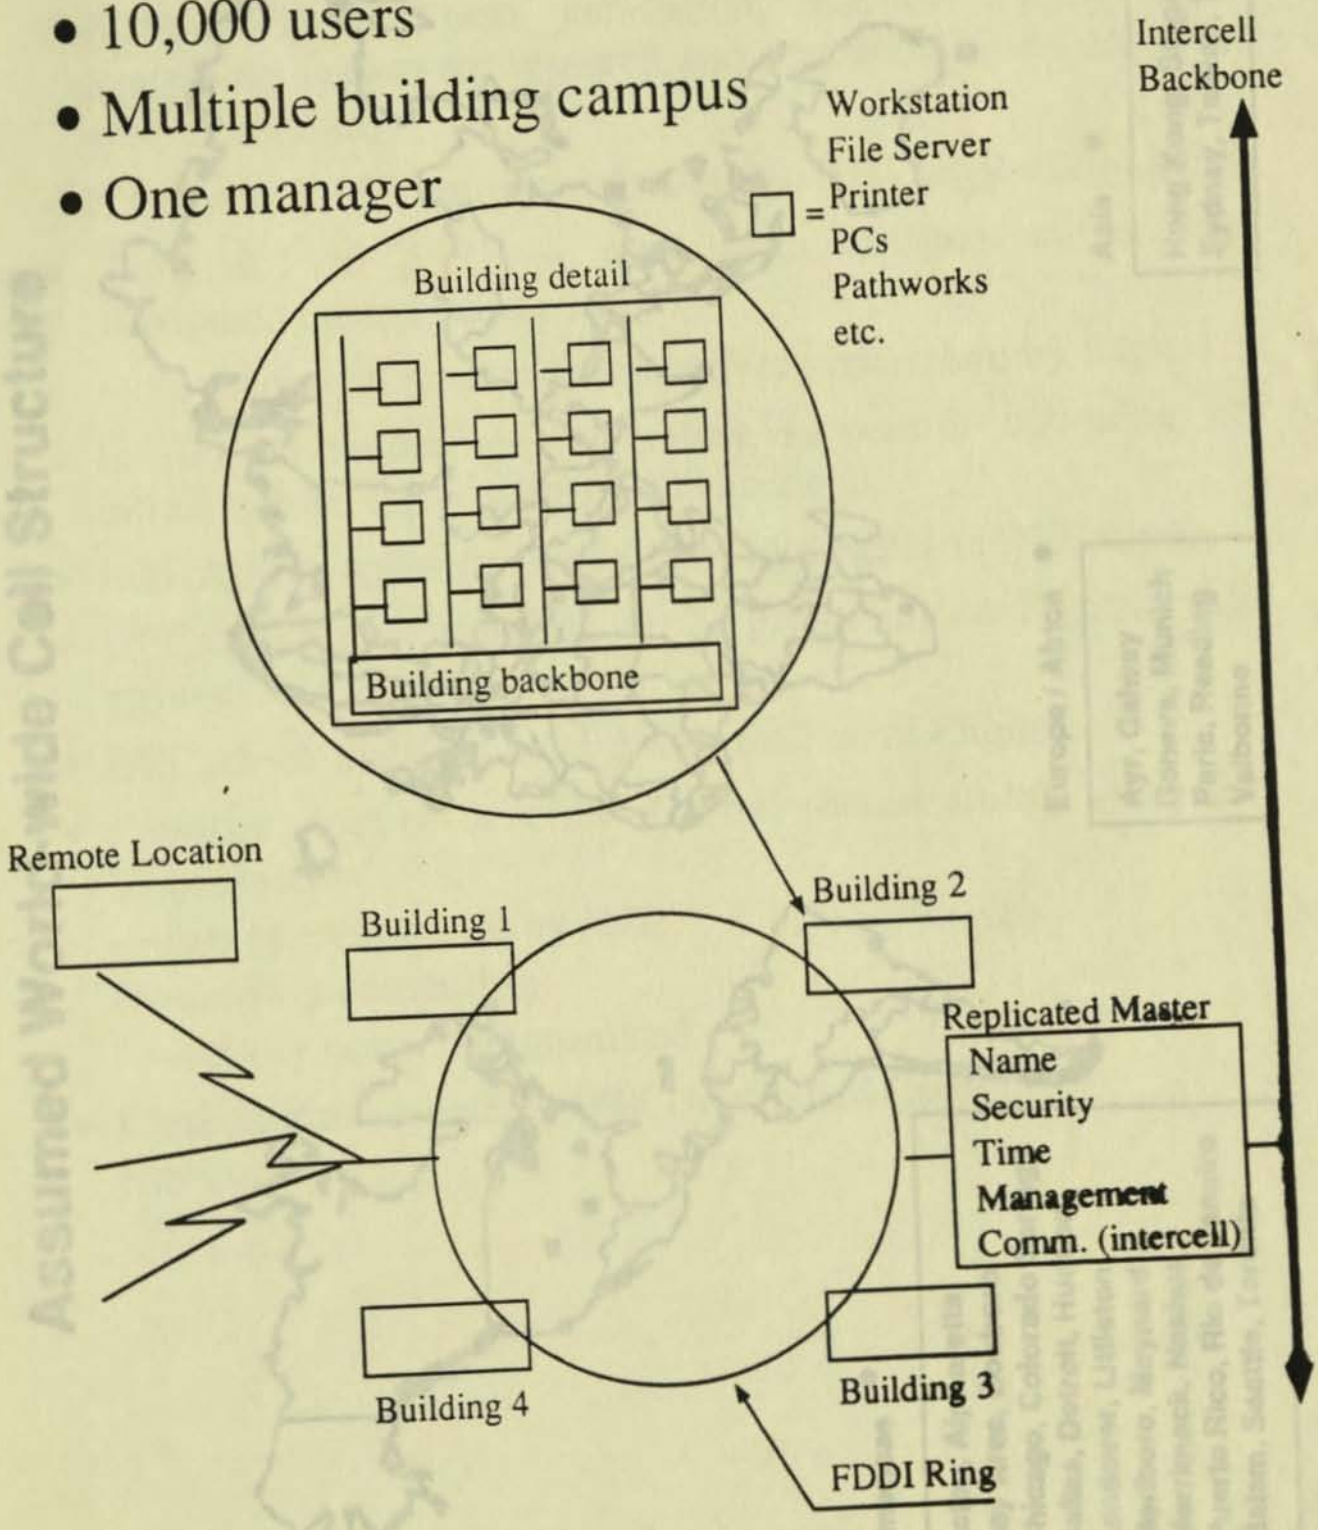

[champine.decwrite]lg-cell.d

Assumed World-wide Cell Structure

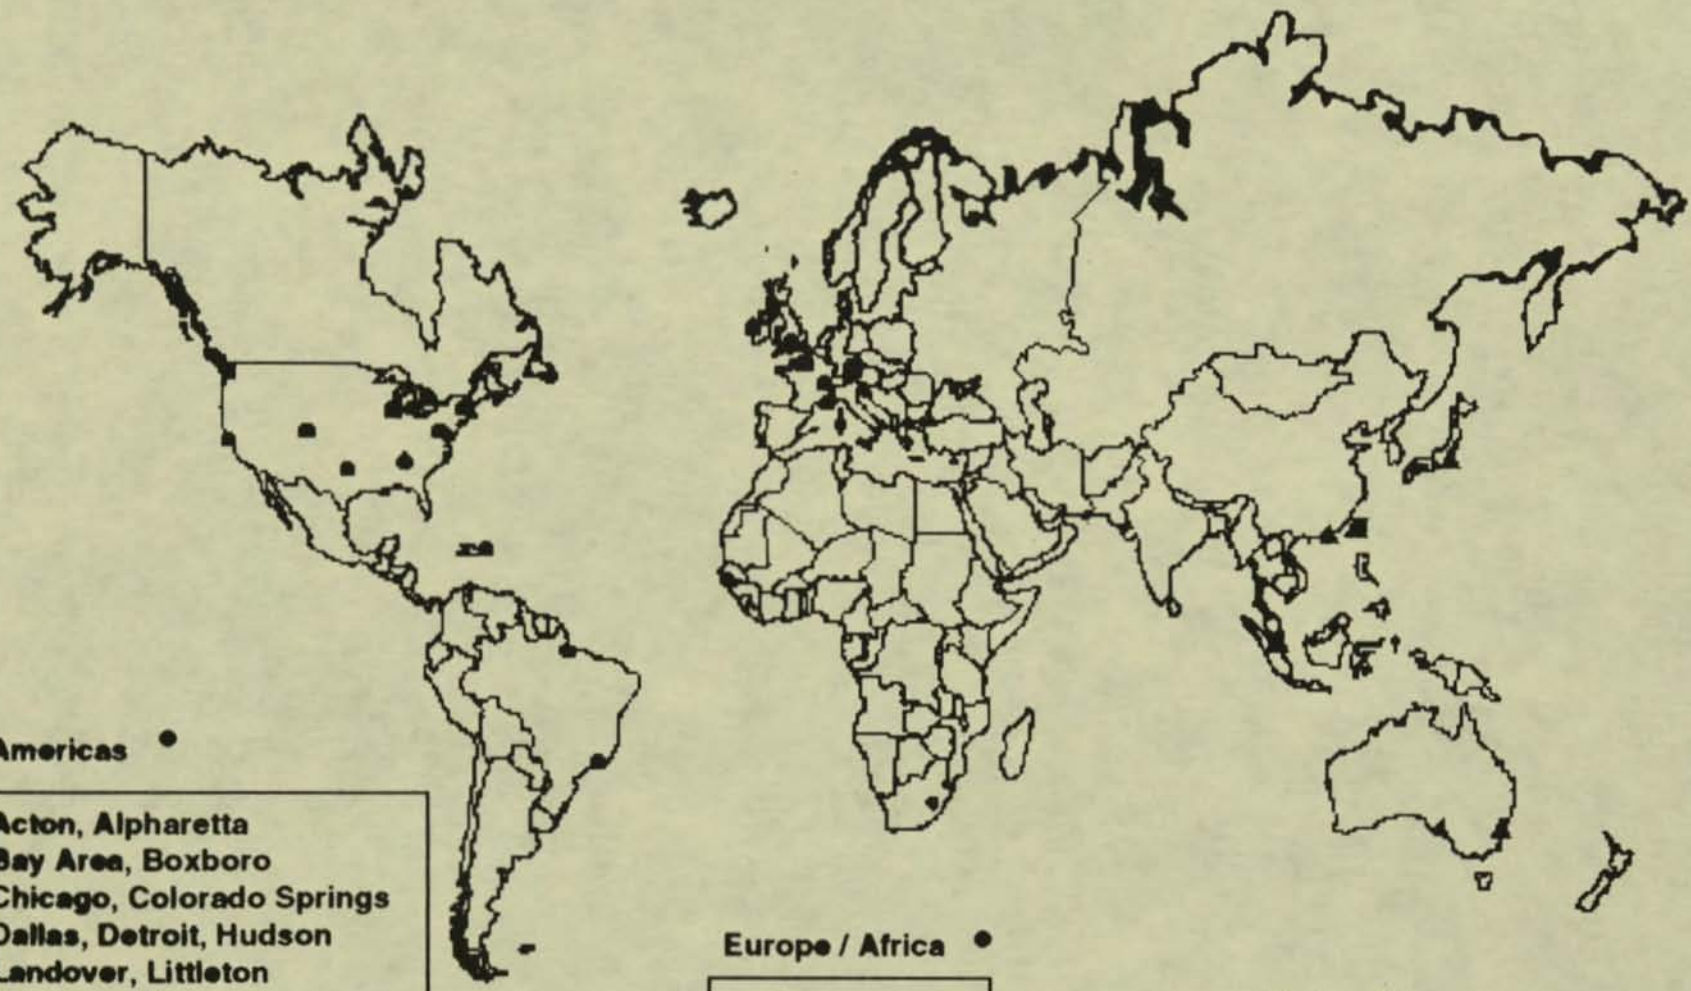

Americas

Acton, Alpharetta
Bay Area, Boxboro
Chicago, Colorado Springs
Dallas, Detroit, Hudson
Landover, Littleton
Marlboro, Maynard
Merrimack, Nashua
Puerto Rico, Rio de Janeiro
Salem, Seattle, Toronto

Europe / Africa

Ayr, Galway
Geneva, Munich
Paris, Reading
Valbonne

Asia

Hong Kong, Singapore
Sydney, Taiwan, Tokyo

MIT Project Athena

- Goal: Provide next generation seamless client/server computing system for research and education based on open system standards
- Eight year $100M program funded by MIT, Digital, and IBM
- Completed in 1991 – now acknowledged model for distributed campus computing
- Most successful research project ever undertaken by Digital
- In full production demonstrating successful operation of infrastructure utility with 13,000 accounts
- MIT Athena software productized by Digital as DECathena
- DECathena now installed in 12 places inside Digital; will grow to 59 locations
- DECathena installed in 31 places outside of Digital
- Provides most of benefits of infrastructure utility today for Unix systems
  —cuts operation cost by half
  —non-stop computing
  —security never compromised
- Clear migration path to OSF Distributed Computing Environment

Lessons for Digital from Athena

for large workstation networks

- Client/server approach to distributed systems validated at scale
- Standard base system for server and desktop support is economical and high quality
- Unified management of distributed system provides low cost and high quality
- Much routine work can be automated
- Heterogeneous systems can be supported economically
- Much higher level of security is possible
- Wide area workstation networks (80 miles) can provide a seamless computing environment
- Significant improvement is possible in system management tools
- Network transparency works
- Must scale up for Digital
  - —10 times in number of users
  - —30 times in geography
  - —100 times in number of desktop systems
- Not a panacea

Planned DECathena Installations

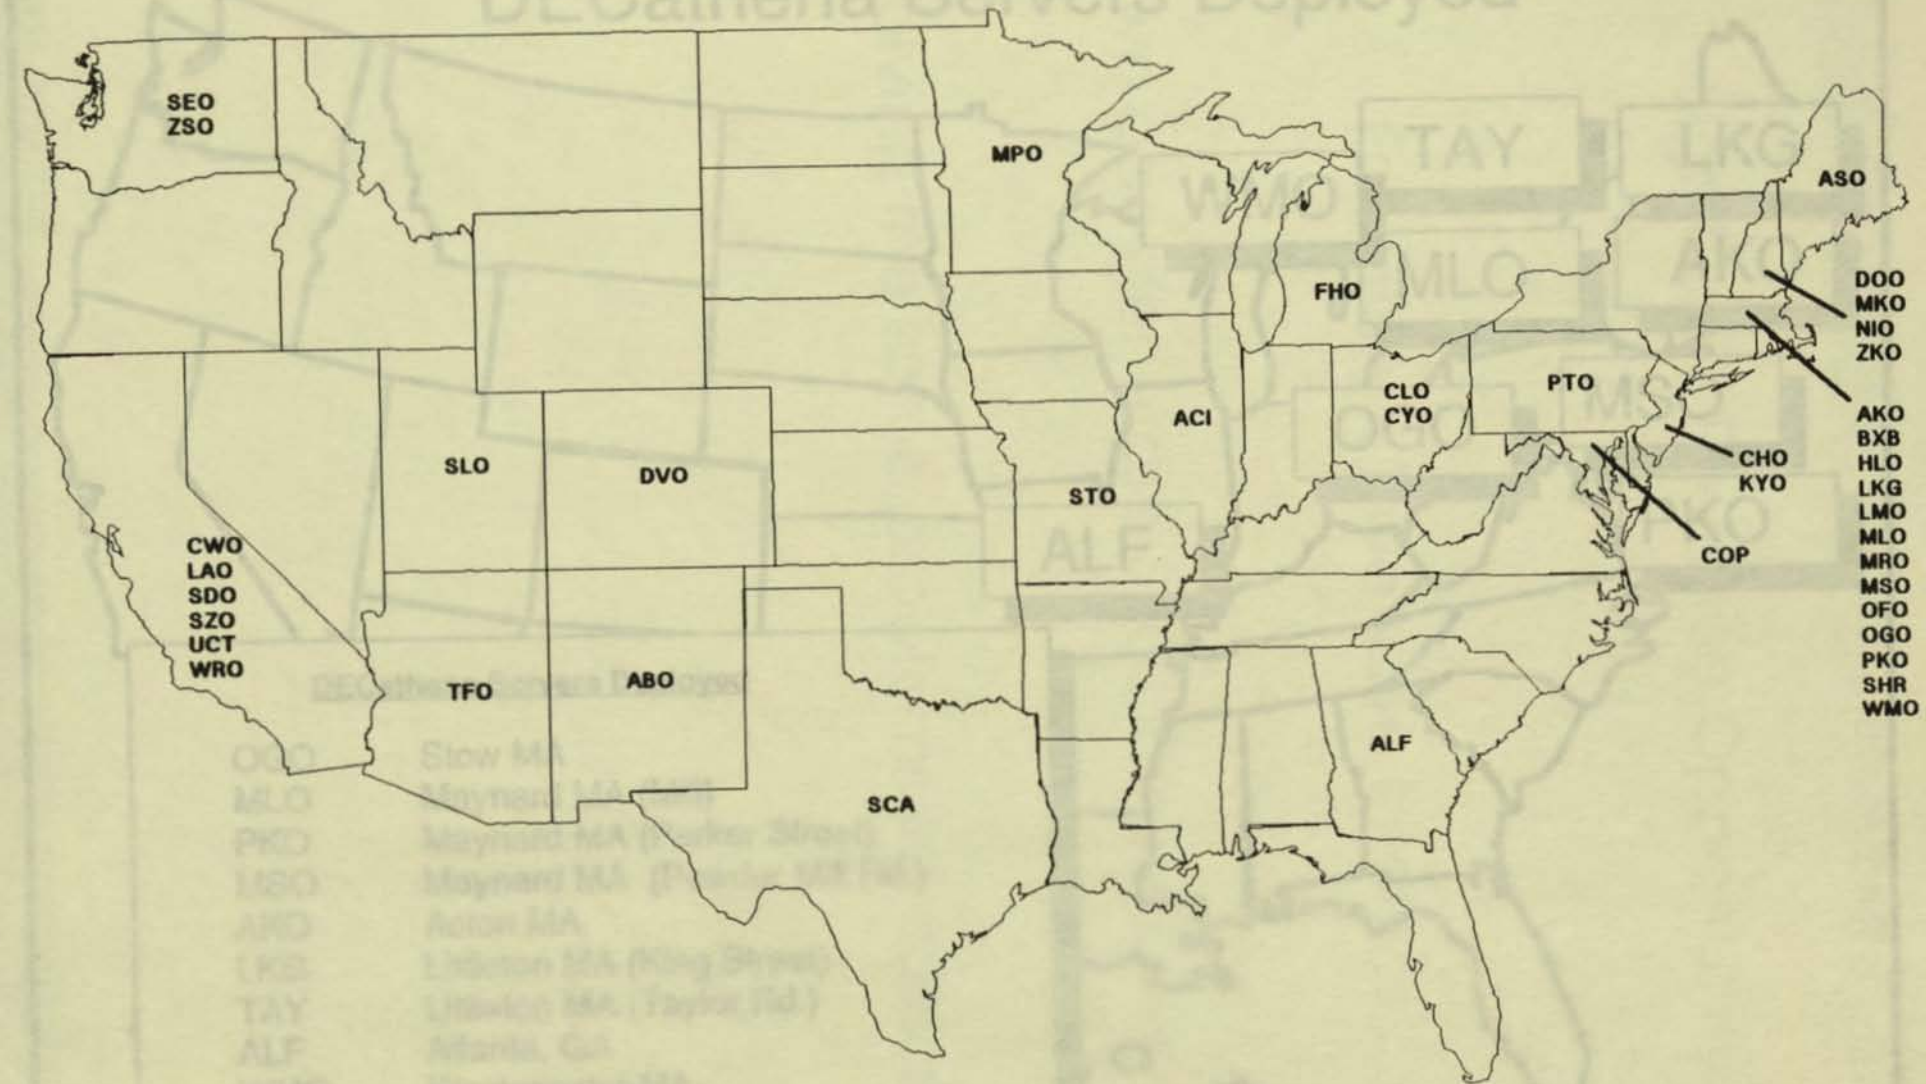

DECathena Servers Deployed

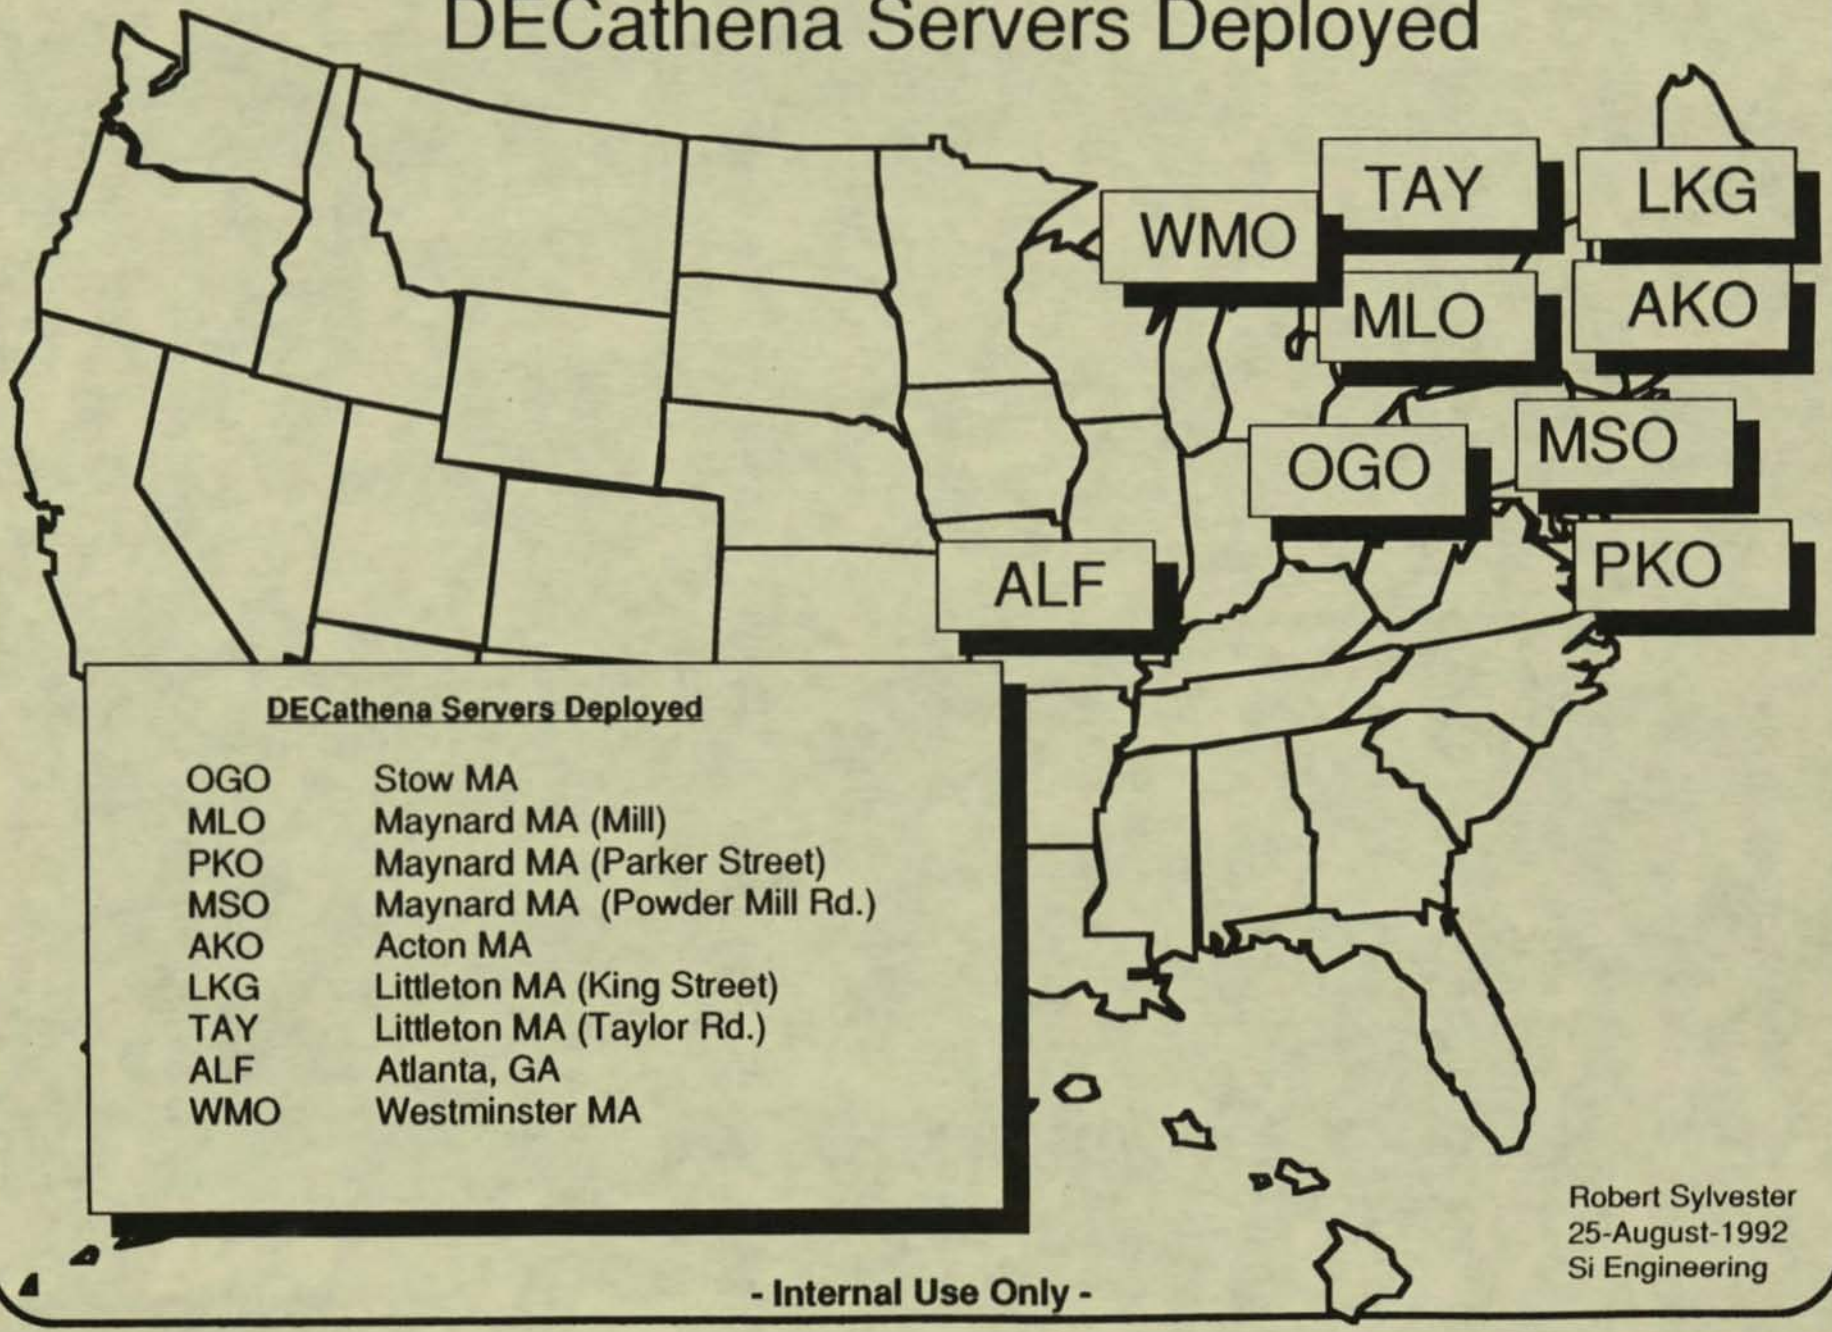

DECathena Servers Deployed

OGO	Stow MA
MLO	Maynard MA (Mill)
PKO	Maynard MA (Parker Street)
MSO	Maynard MA (Powder Mill Rd.)
AKO	Acton MA
LKG	Littleton MA (King Street)
TAY	Littleton MA (Taylor Rd.)
ALF	Atlanta, GA
WMO	Westminster MA

Robert Sylvester
25-August-1992
Si Engineering

Significant Revenue Opportunity

For every $1 we save internally

we can generate $100 in new revenue

solving the same problems for customers

The New Alliance

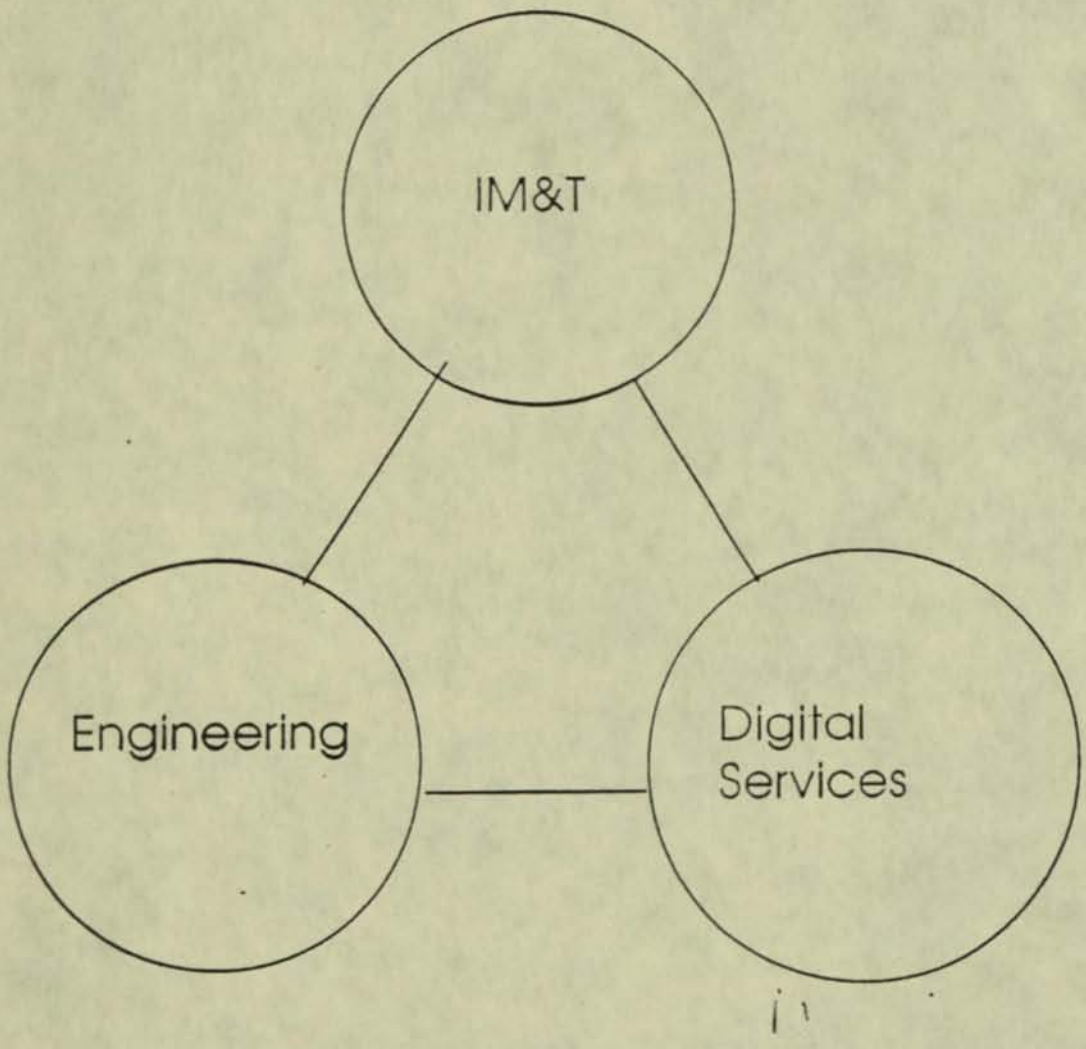

1 QUALITY IS OUR FOCUS

This section describes the ALPHA software program's commitment to the development of quality information systems and outlines the broad "quality system" that we are using to ensure that quality work is obtained. Where applicable, the policy refers to other more detailed documents that define specific methodologies and describe other undertakings in support of the quality system.

1.1 Quality Management

The Alpha Software Program Managers and their respective development managers have the following direct responsibilities:

- Develop and retain a working understanding of quality management concepts and state-of-the-art tools, techniques, and methodologies;
- Establish a working environment that emphasizes the importance of quality work;
- Recognize quality work and promote rewarding professionals for quality performance;
- Introduce tools, techniques, and methodologies to support consistent high quality and high productivity work;
- Develop information to measure quality performance and provide feedback and results to the professional staff.

The intent is to recognize that quality is important; to strive to constantly improve quality; and to manage the entire process. This involves ongoing responsibility for objectives and goals, working out plans to achieve them and to measure the results.

1.1.1 Quality System

The collective effort to achieve quality in our work can be looked upon as the "quality system", which consists of three critical components:

- Leadership—This involves setting the climate and expectation for quality work. It requires management to show leadership, to set goals, and to expect and demand high quality performance;
- Support—This involves providing the tools, techniques, training, and procedures needed to assure that quality is delivered;
- Control—This involves measuring results and tracking the information needed to assure that quality is delivered.

Quality management requires the active involvement of all managers in each of these three spheres.

1.2 The Meaning of Quality

Quality in an information system relates to many attributes and factors. There are exterior qualities like usable, reliable, and correctness; interior qualities like efficient and testable; and qualities relating to future needs, such as flexible, maintainable, and reusable. The overall quality of any system is a composite of all these factors, weighted appropriately to the particular requirements and needs.

1.3 The Cost of Quality

The cost of quality falls into three broad categories—prevention, appraisal, and failure. *Prevention* costs involve investments aimed at getting work done right the first time and preventing quality problems from every coming up. (Examples include training, new techniques and tools, methodologies, and the like.) *Appraisal* costs involve all the testing and check-out efforts to measure whether needed quality has been achieved. *Failure* costs are those incurred by the need to fix and recover from the quality problems that do arise. Our aim is to measure these collective quality costs on an ongoing basis. Fundamentally, we believe that well-thought-out investments in prevention will be repaid many times over by significant decreases in the costs of appraisal and failure.

1.4 Objectives

Our objectives for the next one- to three-year period are the following:

- Establish an effective quality management program and create a quality work environment;
- Significantly improve the effectiveness of the testing and quality appraisal work. (Demonstrate high quality and better-checked-out systems.);
- Achieve significant increases in productivity and efficiency in testing activities—minimum 10 percent;
- Improve the user perception and develop a reputation for quality emphasis in all work performed;
- Develop and use quality standards to achieve a balanced tradeoff between the quality required the effort needed to achieve it. (Avoid overemphasis and overcontrol.);
- Measure and track the quality improvement efforts.

Clients' 1995 Problems, Processes, and Systems

1 Shifting from centrally railroaded to self-driven

Distributed self-driven problemsolving implies
- Information when, where, and "as you like it" for each self
- Just-In-Time education in ways to apply the information

2 Shifting from *problems* that can be solved by individuals to *processes* that can only be managed thru teamwork and *systems* that can't even be understood without simulation

Processes will often include beyond-enterprise partners
- Customers
- Suppliers
- Neighbors (and governments representing neighbors)

Systems with 10 to 100+ causal loops are counter-intuitive
- Graphical display required to see concurrent processes
- Speed-up animation is a critical aid to understanding

3 Teams and even whole communities will depend on results

Errors that threaten life or financial viability will be immoral

Cost per downtime-mip will grow faster than mips increase

Clients' 1995 Ways of Working with Information

1 Computing will be as well hidden as a car's fuel combustion

No hand-cranked start up; no cybermechanical jargon

Few proprietary control levers or idiosyncratic displays

Get to any database location as easy as using a roadmap

2 Integral to enterprise's processes & community's systems:

Information networks used to join people's biological nerves

Enterprises' displays illustrate the enterprise's "self":
- Structure, for navigation
- Current states, as sensed in real-time
- Commitments, including signatures

Learning-enterprises develop as network/people organisms

3 Cooperative work on kalaedescopic and complex projects:

Non-organized cooperation thru shared enterprise memory

Team cooperation thru support of team meetings
- Display of team's mind as they make it up
- Ways for team to change its mind together
- Presence-of-mind for team at subsequent meetings
- Presence-of-mind for each team member working alone

Clients' 1995 Applications

"Applications" will fade out to a merely historical concept.

Team-mind graphical display of processes and systems will look to individual users as a super-set of his/her database.

Differences custom-designed by users will be much larger than differences imposed by programmers' idiosyncrasies.

Clients' 1995 Environments

Digital's assistance will be sold as a vehicle for important learning-benefits for the enterprise with its partners.

Software will be merely hinted at by visible dials and switches. Additional capability will flow in whenever the tap is opened.

Hardware will be concealed "under the hood" of the vehicle. Power will be upped invisibly, just in advance of need.

Geographical distance will not show up as a barrier or bound. Buildings-architecture and network-architecture will merge.

System security justifies trust; peoples' relationships and their work and learning together occupy the foreground.

Digital's New 1995 Technologies in Support of Vision

Automated System Management, including upgrade-ordering and cost optimized bandwidth-adjusting

Rich interfaces with all four accessible nervous systems:

- Pieces / analysis / verbal, logical, musical, mathematical
- Wholes / synthesis / graphical, pictorial, animations
- Concretes / sensing / taste, smell, tactile, vestibular
- Commitments / speech acts, body-language, signatures

Shared distributed artificial reality, independent of geography

Miscellaneous primitives and devices:

- Plate glass team-mind wall displays, 30-100 million pixels
- Spiders (wall pointing mouse-analogs)
- Vertical-writing editors for team-mind matrices
- Data-gloves and artificial hands
- Robot-eyes and binocular helmet displays
- Personal-identity sensors (e.g. fingerprint sensors)

Global standards for complete interface compatibilities.

Alpha

- Futurebus+
- RISC, 64 bit
- Open Storage
- Workstation
- FDDI, Gigabit
- Printing system
- Image processing

....and a socio-technical program leadership deeply informed and driven by a social vision of what computing can provide.

Summary by Russ Doane 13 Jan '91 using materials from
Peter Conklin, Ron Schaefer, and John Whiteside

VAXcluster Vision — by 1995 ...

I. VAXcluster Systems

A. Customers are *"hooked on"* our systems.

B. We have defined success for multi-computer systems in the industry .
   1. operating system heterogenity is transparent.
   2. multi-vendor interoperability is becoming transparent.

C. VAXcluster Systems are truly incremental.
   in every dimension and are fully scalable. Customers can buy just what they need.

D. Digital has defined success in terms of VAXcluster Systems.
   (similar to the way success is defined in terms of CPUs and components today).

II. VAXcluster Systems — Products

(***Products*** = Systems and Components)

A. Solutions are the *easiest to design.*
   1. Configure & price in an hour (on site with a customer or in the sales office).
   2. No "tech-editing" required at all in headquarters.
   3. We have provided methods for customers to think about & easily produce computing requirements in their own languages.
   4. Full installation planning documentation ("CCD") is produced locally, with:
      a) what-if capabilities in an "as-you-watch" mode.
      b) floor plan design.
      c) environmental requirements.

B. Unassailable as the industry's only "boot-once" computer system.

C. Known as the industry's "workhorse".

D. Watch themselves — predict when and where:
   1. they're going to fail.
   2. their availability will fall below customer-desired (our-commited) levels.
   3. their performance will fall below customer-desired (our-commited) levels...
      and what actions to take to achieve the most cost-effective solution (including alternatives).
   4. environmental conditions will exceed thresholds.

E. System design process with formal specifications is operational.

F. System modeling and validation have replaced large, capital-intensive test-beds
   1. configuration designs generate workloads which effectively test design limits.
   2. VAXcluster system-level design simulation tools.
      a) clear understanding of development work required
      b) begun development work
      c) will be delivered in 2 years
   3. sub-system design simulation tools are operational .

III. VAXcluster Systems — Business

A. Dominated (> 90%) by systems-level revenues.

B. Solutions driven.
   1. complete system platforms
   2. fully-functional subsystems (e.g., backup engine)

C. Total annual revenues exceed $1B.

D. Best-run business in the High-end Segment.

IV. VAXcluster Systems — Organization

A. Is a balanced team of technical innovators and managers.

B. Deals with complexity easily.

C. Understands all aspects of our customer base & exploits this knowledge.

D. Has and attracts the best people.

E. Has internalized & integrated quality into all aspects of the business.

F. People are excited about and commited to their work.

+---------------------------+ TM
d	i	g	i	t	a	l
+---------------------------+

Digital Equipment Corporation
I N T E R N E T M E M O R A N D U M
VAXcluster Systems Engineering

To: Staff

DATE: November 13, 1989
FROM: Al Avery
DEPT: VAXcluster Systems Engineering
DTN: 297-4601/4595
LOC: MRO1-2/S10
EASYNET: HYEND::AVERY

Subj: VAXcluster Vision--by 1995.....

By 1995...,

we will be tightly coupled to the business success of our customers. They will expect us to understand their needs in their terms (language) and to offer solutions in their terms, too. We will be required to develop solution proposals fast and to implement them with clarity and simplicity. We will be in an era where customers expect to use multi-vendor solutions transparently and expect those solutions to be highly-robust (always available).

VAXcluster systems will be a pivotal strategic element to Digital's solutions. Attached are my expectations for where our products, our business, and our organization will be by 1995.....

-Al