Peter Kaiser
Consultant

Digital Equipment Corporation
Centre Technique (Europe) SARL
Integrated Solutions Expertise Center
B.P. 027 - 950, route des Colles
06901 Sophia Antipolis Cedex - France
Tél. 33 92 95 62 97
Fax 33 92 95 50 50
kaiser@acm.org

Craig St. Clair

From: kaiser@acm.org
Sent: Wednesday, January 22, 1997 1:54 PM
To: Craig St. Clair
Cc: Bob Supnik; kaiser@vbo.dec.com
Subject: OZIX documents

Craig—

I'm sending you all my printed OZIX papers and the TLZ04 tape cartridge containing a lot more documents in original formats, within a tar file. To read the tape you'll need a UNIX system.

OZIX was one of the two projects I know about to create an "unencumbered" version of UNIX ("unencumbered" in those days meant "free of license requirements to Bell Labs", which owned the rights to all then-existing UNIX source code); the other project, about which I know very little, was carriedout in one of the MRO buildings in Marlboro.

OZIX was going to be the UNIX subsystem of the virtual operating system Dave Cutler and cohort designed for PRISM; and when PRISM was cancelled, OZIX was also cancelled. However, during its lifetime it consolidated a number of innovative ideas that made their way into later products; one of the best of these was the file system, which became the Advanced File System now available under DIGITAL UNIX.

Benn Schreiber, who headed the project, could probably correct me and tell you a lot more. But the papers will tell you a lot.

___Pete
kaiser@acm.org

Craig St. Clair

From: kaiser@acm.org
Sent: Tuesday, January 21, 1997 8:15 AM
To: Craig St. Clair
Cc: Bob Supnik; kaiser@vbo.dec.com
Subject: My notes

Craig—

I served as scribe at two Alpha CAB ("Customer Advisory Board") meetings, of which one took place in Zurich in April 1991, and the other, to the best of my recollection, in Cannes at a date I've forgotten. My notes from the meetings were collected with other notes, edited, and redistributed to Digital participants and perhaps to the customers who attended. I no longer have these redistributed notes, but I'm sending you my handwritten notes from the Zurich meeting. I'd guess that Bob Supnik may still have the edited notes. I know you've just moved, so to what interoffice mail address should I send the original notes?

If you have questions about any of this later, use my real address

kaiser@acm.org

Finally: I've found the printed versions of some of the OZIX documents, but I also have a complete set in electronic form as PostScript. Would you like these, and if so, in what form? The electronic ones are on tape in a UNIX tar file.

___Pete
kaiser@acm.org

Notes on my Alpha CAB notes (Zurich, April 1991)

1. Where I can tell that my handwriting is difficult to read, I simply put the deciphered phrase here marked "D"
2. There are notes from two days. The first day's notes are numbered in the upper right corner from 1 to 6, the second day's "Alpha CAB day 2" 1 to 8.
3. Here I mark an explanation "E" and give the thing being explained, and its explanation.

Day 1 page 1

E "(Chip speed chart)(marked October 91)" is a chart Bob projected during his talk.

D "Desktop WS": "2D vectors"

E "Desktop WS": "TC" means TurboChannel

E "Office system": "MP" is "multiprocessor"

E "Office system": "FB+" is "FutureBus+"

E Quotations below "HW rollout" are customer questions except for "We're telling you..." which is Bob

Day 2 page 5

E Customer names and affiliations are noted on the right in several places

Day 2 page 6

E "ARC" means "Alpha Resource Center"

Digital Program Overview (Oct 91)
Semiconductor facility: $500M
Write an OS: $100M
Maintain an OS: $20M - $40M/year
"The price of computers must come down. Computers are dreadfully expensive."
"God did not have an installed base." ("Why is it taking 3 years?")
Open technology, business practices, service
Binary translation is recompilation
Announced NT!
(Chip speed chart) (marked October 91)
More than 20 seminars for ~600 "application providers"
A PC will be at "initial product offering": desktop, EISA
Desktop WS: 120 SPEC, 512MB, Turbo Ch, "smart FB" ≈ 500K 2D vectors etc
Deskside WS: 140 SPEC, 1GB, TC (6 slots), graphics
Office system: MP 1 or 2 @ 140 SPEC, 2 GB, FB+, 130+ MB/s IO including FB
Department: MP 1-4 @ 130 SPEC, XMI, FB+ (mixed) for 100-400 MB/s, 14 GB, VAX or A
Datacenter: MP 1-4 @ 130 SPEC (see above; high available (UPS, power conditioning, warm-swap disks, integral ESE)
HW rollout begins in 4.92

"Why continue to make VAXes?"
A: Huge tightly-coupled SW systems; embedded systems; where stability is paramount
"Is translation 100%?"
A: Should handle any non-privileged code. Also have MACRO-32 compiler. "We don't want to force premature migration on you."
"DBMS?"
A: Being ported.
"ALL-IN-1?"
A: Already being rewritten in distributed client-server form. But "classic" A1 will also be brought across!
"Oracle?"
A: Oracle, Sybase, Ingres, Informix already working with us.
"Will A still be ahead when it ships out?"
A: We believe it will be, but we don't rule out competition. "If you have a
"We're telling you what we are going to do."
similar vision..." All IBM HP + Sun are trying to do their interfaces"

A: Well begin with deskside WS & server.

"Why should I consider a 9000 when the AlphaWS has equal compute power?

A: IO power, for one thing. But technology has moved on. "We'd really also like to help you unburden yourselves of your IBM mainframe environments, which, as you know, are extraordinarily expensive."

A (Gamache): Can't underestimate the 9000's crossbar IO technology, which is very complex and very effective.

Supnik Tech Info (Product PID)

"What's the speed of the chip?"

A: 150 - 200 MHz internal speed from an external clock double the speed.

"What about physical memory?"

A: 8KB inst + 8KB data primary cache on chip, plus the fastest SRAMs we can find as secondary caches.

(Chip is $24^2 - 12^2$ pins = 676 - 144 = 532 pins)

Our FB+ implementation will do 130+MB/s

Datacenter: CPU @ 200MHz = 180 SPEC, 200 MFLOP peak with 4MB cache (9000-600)

"Does migration include from Microsoft to A?"

A: "We do intend to port NT Windows, but we don't have a schedule yet."

"Will those enclosures be available first with VAX CPUs?"

A: Yes. This summer.

"What about the quality of ports? Will we be field-testing your ports?" (Hallam)

And how about proofs of correctness?"

A: We're including all corrections to VMS 5.4 since it came out. Second we have a brand-new set of regression tests, including ones written using the manuals as specs. Third we're using the extended rollout period. Finally VMS has undergone a line-by-line code review. We intend not to ship unless we can do it right.

"Do I understand that I can reuse existing VAX hardware?" (Leich)

A: No, they won't support A power. But we'll do trade-ins, etc.

Customer meeting

C: We have 1000 WS + 800 PC + VAX clusters. What good will 500 SPECs do us when the problem is not the CPU but the network? We can put only 50-80 WS on a server. We'd like to have 1000. We want network-wide applications and databases. We already have 16-way redundant FDDI, and that's too slow.

Supnik: First will make FDDI full-duplex, doubling bandwidth. More important we're doing gigabit links plus gigabit switches, first 32×FDDI then up to 1K×FDDI. Also lowering cost with 1-chip FDDI implementations. "Network bandwidth must scale." Switch will be out with α, roughly.

C: We find we need 1GB disk/MIPS in the data center. What happens with α?

C: We have no dramatic need for centralized CPU power. Really need it on the desktop. (Creep)

C: We're a bank, and amortize about $15M per year. I can now replace all my HW for, say, $6M. What do I tell my board? Second, are VMS and OSF ready to handle thousands of users with TP and databases? What do I do with this quantum leap? When can I really use the technology? Where are the tools to migrate my mainframe applications?

C: What about management tools? DECathena doesn't track it.

C: 4×9420, μVAX×140, 60×4000, PC×12000 (PTT). How do I manage ROI for my existing systems? (Dutch PTT)

C: If you want us off the AS/400 you've got to help us with price or tools.

Lomas: We'll price aggressively, because in the open environment will have to compete with others building from the same components.

Supnik: DEC has restructured into independent profit centers, down to the component level. Will compete at every level. "Where there is comparable function we must have competitive pricing." Where we have special value we'll expect to get a fair price.

Lomas: It's not a UNIX market or a proprietary market. It's an open market.

C: When will you sell chips?

Supnik: We intend to sell at every component level from the outset.

C: What's your strategy for α?

Supnik: To cover the whole range of requirements, perhaps even make α the prevailing standard for 64-bit systems. To do that we need to know your requirements.

C: We're moving away from mainframes to open TP systems (cover with Oracle), and all we really care about is cost per transaction. Timesharing is falling off.

C: Data center machines are only 6x more powerful than desktop machines. Is this deliberate?

* Supnik: It's a fact of CMOS. What distinguishes systems now are memory and IO. the curve is flattening.

C: Back to the original question: we can't predict. Will have to evolve.

C: We'll need many MIPS on both client & server because of the rise in users' expectations.

C: We need lots of CPU & memory to do X windows. But timesharing is still best and cheapest for low-demand users.

C: Lots of CPU power can give us breathing space in the expensive process of deploying distributed applications.

C: What work is DEC doing on true high-availability distributed systems with WAN ability and true failover, so I don't have to worry about where or how my work is being done?

(31) C: Can I now finally get rid of the keyboard? (Kneipp)

Supnik: CMU believes this level of performance is enough for useful level of continuous speech recognition, especially with the level of floating point processing.

C: We must get past multiple kinds of expert-friendly windows to really ones that are really usable without training. We see this in NT; X and Motif are niche products for a few technical users. (Hallam)

C: Ditto. The innovation is on the PCs. Servers will do the minimum, and powerful desktop systems will do the rest.

(34) C: We need faster disks; CPU is not the bottleneck for our bank. (Kneipp)

Gamache: Can you use smart caching? A: No; access is random over gigabytes.

Supnik: Eventually one should keep the entire database in memory and merely log changes to disk. 64-bit addressing makes this possible. Memory is always fastest.

---

* Traditional curve of early 80's

flattens to

(36) C: Have you plans for a virtual OS so we can run both VMS and OSF on one machine?

Supnik: We've investigated it. It is fully virtualizable, but we need to hear from you about whether we should do it, and why. Security? Flexibility?

C: Ditto. We need to run multiple versions of the same OS for different applications.

C: Ditto. We hate to reboot to just to run a new application.

(37) C: What about it's future?

Supnik: EV-4T @ 250MHz in early 1993 (200 SPEC); EV-5 @ (300 SPEC) in early 1994. We expect to double performance every 2 year or so into the late 90s. We're also doing cost-focused ~~versions~~ versions of lower performance beginning in 1993 for embedded applications and systems under $5K. Memory address need grows about 1 bit/year, so (with 48-bit physical addressing) we expect to last into the 2000s. [showed chart]

(39) C: Pricing?

Lomas: Comparable to the systems they replace.

(40) C: I need more information to ~~do~~ do configuration and cost-of-ownership planning.

Lomas: We're trying to provide measures in SPECs and TPS. Systems may differ widely based on cache, IO, etc. Many organizations have their own benchmark suites.

C: We're moving toward client-server with specialized servers, so we no longer need much in the way of generalized benchmarks.

C: Because of reliability and the vulnerability of single systems we'll hit the limit of usable reliable capacity before hitting the hard CPU limit. (We can't afford to lose 1000 user tasks at a time just because the system will run them — if it goes down.)

(42) C: New SW — like OO technology — takes more CPU.

(43) C: We have more SW crashes than HW failures. SW shouldn't crash the machine!

(44) C: Our bank would completely change our approach — in a business where we must calculate calculated risks — if I could calculate them in real time rather than only knowing the answer the day after. What does DEC advise me about how to do this? It's not enough just to give me the machine and say "Here it is; play with it."

(Kalish) It's our business to take risks"

Supnik: We have no neat answer, but we know that it won't be good enough just to run the old SW faster nor to emulate or replace, say, CICS systems. Perhaps distribution techniques can help.

C: You have to come up with techniques, or your competitors will.
C: Δ systems must interoperate with other vendors, for instance in OLTP.
C: How compatibly will Δ systems be with the UNIX systems I'd like to replace with them, including DECstations.
Supnik: Very. Same services and networking plus binary translation. Best to recompile.

(6) C: Are you speaking with CAE vendors?
Supnik: Yes, all the large ones. Which do you need? A: Mentor & Computervision.
Lowes: We're speaking to many hundreds of vendors worldwide. But we need to know what you need; fill out the questionnaire.
C: We're buying not systems or workstations, but applications. Then we look to see what the hardware can be.
(1) Supnik: I'm curious about databases. Who uses:

Oracle	13	DBMS	1
Ingres	1	Adabas	6
Informix	1		
Ingres	1		
Rdb	1		

(2) C: How do SPECs, MIPS, VUPs enter here? I see 140 SPECmarks for one Δ system.
Supnik: SPECs are a workload, an engineering workload. For complex commercial applications I expect a wider spread. On some commercial applications still screams; with Δ we're doing "binary shadowing" for Cobol which will increase speed drastically. (Read & convert, compute always in binary, convert to write.)
PK: Do the comments on need for management tools depend on Δ or will they be there anyhow?
C: In the future each user will be a node on the network, so 4000 users means 4000 systems to be managed, with or without Δ.
C: Ditto. "It's not an Δ-sensitive issue." Need better security because "every smart device is an opportunity for mischief".
C: Perhaps we'll use all those MIPS for encryption. [Laughter in agreement.]

Supnik: SW developers PID

57 C: X server compliant?

Supnik: based on X11R5 and fully compliant

C: What's the difference between [Tin] and [Silver]?

Supnik: scalability. ~~Different~~ Fixed data structures.

C: I need Scan for conversions. (Leich)

[GEM runs on VMS, UNIX, and NT]

Supnik: VEST/TIE is "a safety net"

58 C: How does the size of exe change under binary translation?

Supnik: Approximately doubles, because it includes the original VMS [RISC] executable.

Gamache: α VMS overview PID

Original goal of VMS included "open languages": standard calling sequences

C: Do I understand you will be decreasing VMS functionality over time?

Supnik: Absolutely not! On the contrary it will be improved, but α VMS will eventually pass VAX VMS

Gamache: No: rather we're stabilizing it for production environments.

VMS Issues/Rollout

59 C: My main issue is 3d-party products and clusters. I won't move until I can have full clustering with all functionality. (Except workstations, which I can put on the net.) Volume shadowing is a must.

C: (Many dittos.)

60 C: What's the problem in bringing shadowing so late in the schedule?

Gamache: it's host-based. We're trying to bring it in sooner. It's very complex to test clusters, and the testing is the main difficulty. But clustering is the #1 candidate for promotion.

Supnik: Recommend a maximum performance difference of 5:1 within a cluster; since the minimum α is 100 SPECs, the minimum VAX in a mixed cluster is a 6000-400.*

61 C: Need to boot VAXes from α in a LAVC and vice versa.

Gamache: Will do.

C: Support HSC functions?

Supnik: Never. HSC presents unsolvable architectural issues (e.g. failure modes) so in all VMS we will move entirely to host-based shadowing.* We will help you qualify clusters.

C: need Interleaf.

Supnik: Some DEC applications that are not compute-intensive will first run translated — like Notes, VTX. EDT will be translated, TPU native. TDMS translated, also FMS.

(38) C: You seem to be providing technical stuff first and commercial later, as usual. And what about ACMS for translation early. Wire commercial. How long must we wait?

Supnik: ACMS will be native because much of it is privileged, so it must be native. Remember that "Wave 2" is only 3 months after Wave 1.

(41) C: Need a tool to migrate TDMS to DEC forms, 99% good, so I can get rid of TDMS.

Supnik: We will have new and competitive VAX systems, so your migration is not a "forced march." If you have VAX-appropriate SW that is being segregated only later in our porting plans, buy VAXes and migrate later when it better suits your business needs.

C: What about version control? We already have multiple versions of VMS on our network.

(43) C: Licensing? We need "#-of-concurrent-users" licensing.

Supnik: Personal use and #-of-users licensing.

C: [Rephrase] Need distributed licensing for networks.

C: We can't predict how long we'll need support of mixed clusters. It may depend on how long OVMS and VAX VMS don't diverge.

(46) C: It may pay to reexamine why people cluster VAXes: in our environment it's in order to manage them. What you really need is to solve the problems of distributed management and licensing — then cluster problems largely go away.

(47) C: Plans for fault-tolerant systems?

Supnik: Proposal within FT group not yet scheduled. Also talking to potential external licensees. Also working on new VAX FT systems in the new departmental packaging. Who needs true FT? [1 customer]

C: Need data on 3D graphics, the reason we now use SGI. (Tom)

(5) C: Want to run both VMS and UNIX.

Gamache: the issue is of PAL code, and whether we can fit both versions of PAL code into ROM (flash ROM) on the CPU modules.

(1) C: WS applications will be the first, but you must ensure that 3d-party products must be present, for us the CAE applications we now run on Suns. Then we won't buy the Suns if it performs better.

C: Position ACE & α? I thought the point was binary compatibility.

Supnik: We began with binary compatibility, but we now consider API (source) compatibility now seems most important, including OSF AES in OSF, NT, and SysV.

C: So I assert that NT is now the only means of binary compatibility, which will be what makes shrinkwrap SW work, and all UNIX is still a niche.

Supnik: 95% of all machines sold are DOS machines, 4% Macs, <1% UNIX. But RISC really relies on recompilation, and the performance-driven market won't be satisfied by Intel or MIPS. We must bring in new technology. So it is a modification of our hopes for ACE, but ACE still holds promise for the volume market.

C: So you will port NT to α?

Supnik: Yes! What ACE will provide is very price-competitive products, whereas α will have highest performance. We are determined to provide what our customers need in any area.

C: Intel programs on α?

Supnik: We'll have SoftPC on VMS, OSF, and NT on α. And α NT will run Intel Windows applications.

[Draws diagram]:

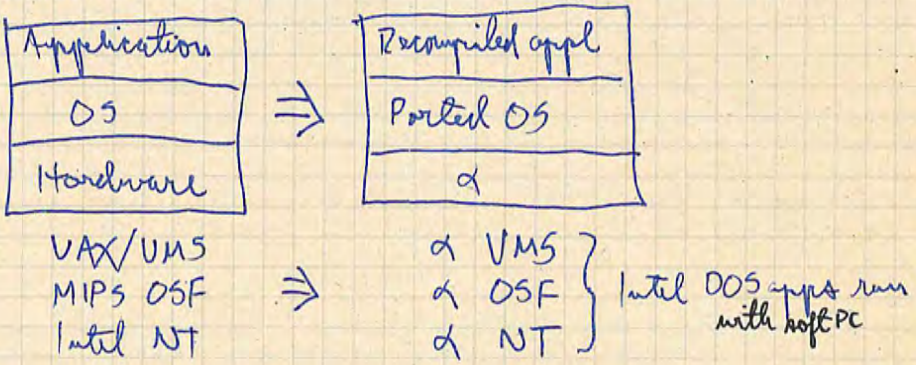

C: I thought ACE UNIX was SCO's version based on SysV.

Supnik: ACE UNIX has always been OSF from DEC, and SCO was only an integration and distribution agent for Digital.

C: Especially for α we need to be able to prevent one application from harming all system performance.

Supnik: class-based scheduling with guaranteed resources and an interface.

C: It is based on VMS 5.4. Can we still upgrade our VAXes to VMS 5.5?

Supnik: Yes. Batch, print services are the main Δ, and they aren't turned on in a cluster until all systems are at 5.5. Also defragmentation primitives. We will support rolling upgrades of mixed clusters.

(12) C: How do I choose between VMS and OSF?

Supnik: on the basis of your applications: if you have VMS, stick with it. If you have UNIX, go with OSF. Our initial offerings have

VMS: commercial + reliability focus

OSF: technical focus

There are also issues of scaling with OSF, plus a history of stability. Scaling: 256 users, 512 MB, etc.

(13) C: So if I begin today what should I do?

Supnik: Today for reliability and throughput I'd take VMS; for addressability and performance I'd look at OSF.

C: What about CSOs on OSF?

Supnik: Just beginning to talk. They'll have to consider performance in this 64-bit environment in relation to their portability.

C: By a release we're out will be using VMS 5.5, so no one will be able to use a. Issue of batch queue ~~output~~ management

Gamache: That's a misunderstanding. Things will work.

Gamache: Moving applications PID

C: Promoting bytes to longwords effectively reduces cache size by a factor of 4.

Gamache: That's true, but it's only a secondary consideration; reducing the size of code is primary. There's also atomicity. You're reducing 4 instructions (load, mask, merge, store) to 1.

C: Cobol packed decimal is for precise financial applications. Doesn't binary shadowing lose that?

Supnik: No, binary integers are precise. We'll provide routines to emulate the functionality of packed decimal division.

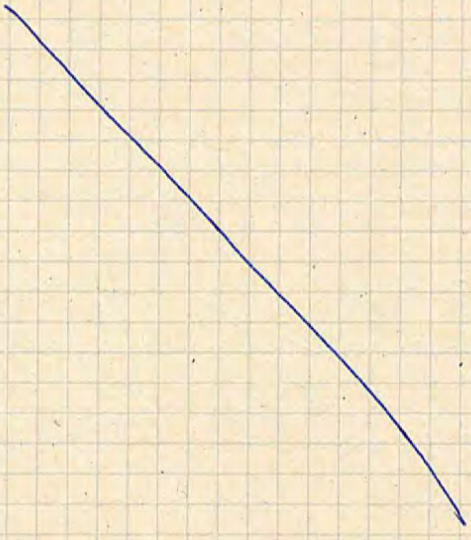

Migration issues discussion (L) : Tomas

(G) Q: Is there support for low-level RMS functions. (BA) Holmes
A: Yes. All of VMS "including the warts". (G)

Q: Should we migrate — say for 1993 — to Alpha VMS or Alpha OSF? The long-term impact for us may be substantial in, say, 1998. I heard Bob say what to do today, but I'm uneasy. (Huber) Philips
A: (S) The answer seems straightforward: if you're now using VMS you may have a hard time migrating to any UNIX.
Q: But what's your long-term strategy? (Huber)
A: (S) Focus increasingly on commercial features like reliability and clustering.
A: (L) Make the choice on the basis of your applications portfolio, because both systems will be around in ten years. But lots of the same interfaces will be on both. If you listen only to the material about VMS you may think we're ignoring UNIX, but we could tell the same story about OSF. We're doing both and both engineering groups report to the same VP.
Q: Pricing of Alpha VMS and Alpha OSF? (von Lepinski) Datev Agriculture
A: (L) Similar for similar functionality [laughter]. We're also looking at packaging options for VMS that will let us provide less functionality for VMS so we can price VMS like OSF at equal functionality.
Q: We see SW loaded off VMS onto UNIX machines. Why should they now use the same SW based on VAXclusters. So what should be my final choice? What's the impact on a data center of a decentralized UNIX environment? (Huber)
A: (G) Reliability in VMS, we hope; and we hope to improve it. But we also want to improve interoperability with UNIX, for instance by integrating the TCP/IP stack into VMS.
Q: Does that mean it makes sense to keep technical applications on VMS and serve desktop environments from there? (Gunter A) Bosch
A: (G) Up to the ISVs and developers. We want to accelerate our plans for DCE and be able to serve all desktops, including PCs, that way. We'll hope to provide very reliable VMS Alpha systems as compute servers. Is that how you see the world working? Of course we also hope to provide the desktop systems also!

Q: How do you expect sales to go in VMS and OSF on α? (Haber) Phillips

A: (L) At first mostly VMS because of pent-up demand, but later it will swing to OSF and probably NT.

Q: How many persons working on VMS, OSF, NT (Bingleier) Alcatel

A: (S) VMS roughly equal OSF; NT just begun, so unknown.

Q: We need migration specialists close by, not in distant sites like Valbonne. (Uwe) Bosch

A: (L) ARCs in each country for core competence, but we hope to work with account managers to provide support specific for you. The fact is that HW will be limited at first, so ARCs make sense, but we have a seed plan to begin in mid-1992 when more equipment becomes available. But please let us know if you have more immediate or special needs, so we can do some planning together. I want to involve your account mgr. Tell them you're interested.

Q: Price of migration assistance? (Tables) Cegelec

A: (L) Case by case, country by country. UK cheaper than CH, for instance.

Q: Market on WS is driven by availability of SW. You must get the ISVs behind α, or we can't use α no matter what the migration plan. (Uwe) Bosch

A: (L) That's why we're talking to the ISVs, many hundreds, and why we did it before talking to you. That's also why we need your questionnaires

Q: The ARCs must be able to send me a FAX now to tell me a plain yes or no for things that are or are not on α, or when they're planned. I need answers easily and rapidly, or I'll never be able to move to α. In the past you haven't been good at this. (Leicht) Logo

A: (L) Right. Noted.

Q: A lot of richer new software is needed for the future. What plans have you for new software that takes advantage of these platforms? HW costs are a minor part of the issue. We need much better SW to make the migration really worthwhile. (Hallam) BP

A: (S) [Slide 4 of Peterschmitt's α Tech Intro]

Make it much easier to ① write distributed apps ② develop OO apps ③ do CASE ④ develop user interfaces.

Q: How about your own apps? What about end-user apps? Compound documents? (Hallam) BP

A: (S) Improve middleware, groupware. Add hyperinformation and multimedia, distributed applications like ALL-IN-1 to incorporate the user's choice of apps rather than just WPS+.

Q: Will it all be on VAXes too? (Hallam) BP

A: (S) As much as possible, but some things like MM (full-motion video and voice recognition) will take the power of Alpha. Perhaps in 1994-1995 things will begin to diverge.

Q: So are you trying to take a stance of leadership? BP

A: (S) In middleware, and in some selected vertical apps. But not in the personal productivity apps that MS & Lotus do so much better.

Q: We see a massive demand for groupware and are desperately seeking applications. Your competitors seem to have the vision but not the platform, whereas you seem the reverse. BP

A: (S) We believe we know how to do things in a distributed environment.

C: You need the end-user apps. Not to have them would be criminal. Form alliances if necessary. (Hallam) BB

C: Doesn't the Windows announcement meet this need? (van Loqih) Dutch Ag

C: } More pointless discussion between Hallam and van Loqih

C: When can we get our personnel trained on Alpha?

A: (L) I'll take that as an action question — schedule for future training. We already have internal pilot courses.

A: (Grimshaw) There's already a schedule for customer courses. We don't have it here.

C: What will be at US DECUS next week? (Leich) Logo

A: (L) Migration stuff, but not the details you've heard here about products and schedules.

C: Was disappointed not to see anyone from your OSF group here. It makes it difficult to believe you're serious about UNIX.

A: ① That's because for this meeting we aimed at the VMS installed base; and we do have a UNIX partner here, Peter Kaiser.

C: Ditto the need for OSF material. Also info on peripherals – disks etc.