

A Case Study of the Failure of Digital's Workstation Vision

This report describes our situation at Structural Dynamics Research Corporation (SDRC), an important customer where Digital's lack of workstation graphics performance may cause Digital to lose their business. This loss may preview a series of losses because Digital has no competitive offerings in graphics subsystems or high-performance X terminals. I describe specifics of SDRC's trouble with our workstation and X terminal strategy, and discuss how we can make ourselves competitive in this market by buying world-class workstation graphics technology.

Revision/Update Information: Original Notes, 11 October 1989
First Draft, 16 October 1989

Steve Elmore — Systems Consultant, East Central Area

Pete Steciow — Digital CMP Account Manager, SDRC, 16-Oct-1989

Michael Breen — Systems Consultant, Western Area, 16-Oct-1989

Peter Kaiser — Manager, Systems Consulting Group, 16-Oct-1989
Digital Equipment Corporation

Digital Internal Use Only

1 Introduction

This draft report¹ tries to illuminate a problem which Digital has only a very small window of opportunity to solve.

Almost every week I find myself confronted with the subjects of Digital's workstation performance, with our graphics performance, and with Digital's desktop vision in general. The problem is that our most technical customers perceive us as having fallen behind in technology. Many don't expect us to catch up. Some of our most dedicated customers are seriously rethinking their relationship with Digital. What Digital does (or doesn't do) in the very near future will almost certainly be of strategic significance to Digital in the coming months.

In some sense this is an appeal, because there is a timely solution. The problem is compounded by a seemingly complacent attitude that is preventing us from taking *any* action. So this is an appeal to act or at least to empower someone to act.

Please understand that I'm not out to criticize any person or any organization, I'm simply trying to do my part to convince someone with the right authority to refocus Digital's product engineering towards the best that we can build — which can be "world class" again. Digital can do better, can build better products, and can be the technological leader again. We shouldn't lose to lesser companies, especially when there is a solution at hand.

What follows is my opinion, and am professing it from an admittedly incomplete picture. I am "in the field", so I may not have the larger scope of things that others may have. However, I do hear some very important messages from our customers that decision makers may not hear. I feel it is important enough to at least express myself and forward this on. I have been inspired by my counterpart in the Western Area, Mike Breen, who has fought for this idea for over a year. I have the endorsement of Pete Steciow, the Digital CMP Account Manager at SDRC, who has worked tirelessly trying to do the right thing for SDRC. Through Pete and my own contacts at SDRC, I'm convinced that we must act now or we will lose not only SDRC, but many other good Digital accounts — and the potential to get into many other accounts in the future.

1.1 The Problem: Digital's Workstation Vision

The problem that I am addressing is Digital's workstation vision. Problems around this vision have permeated several of our more technical accounts. I am using a case study approach to describe several areas of concern. I will also briefly identify several other accounts where Digital's vision and practice are failing other companies.

¹ Personal note: I realize that this report is long and wordy. Although I have struggled with this paper, my written words don't seem to have the impact and urgency that I feel. For that I apologize, but I believe my recommendations are the right thing for Digital and hope they can happen. Pete Steciow, Mike Breen, Pete Kaiser and I are more than willing to meet in person to talk about this to help understand and support this proposal.

1.2 Case Study: SDRC

The company I have focused on is a Digital CMP, Structural Dynamics Research Corporation (SDRC). They develop one of the world's premier solids modeling, mechanical CAD/CAM products, called "Ideas". I chose SDRC because I know the account and have found them to represent several key points where Digital is falling short.

Digital's association with SDRC falls in three areas:

- Digital currently has a major presence as SDRC's internal development platform;
- SDRC also leverages a significant amount in sales onto Digital platforms; and
- Digital and SDRC form a strategic presence in the solids modeling and MCAD/CAM world.

2 Digital is Losing the Development Environment at SDRC

SDRC has been a solid Digital customer for over 12 years. Over the years they have enjoyed and respected Digital's excellence in engineering and our steady and almost predictable improvements to VMS and the VAX line. In fact, while SDRC has ported their software to almost every major workstation platform, 60-65% of SDRC's software was sold on Digital systems.

Because of the disproportionate amount of leveraged business in Digital's favor, it added impetus for SDRC to develop on VMS. SDRC built a sophisticated development environment for their engineers eventually comprised of a 16-node CI cluster², an extensive network, and many Digital software products. That started to change about a year and a half ago.

In 1988, the CI cluster began to become so complicated and burdened with the demands of a mature development shop that SDRC began to investigate plans for how they could grow into the future. SDRC asked Digital to help make those plans, so that summer and fall, we shared our "workstation vision" with SDRC. The vision was based on the VS3100 and X. We described how the CI cluster could become less burdened by moving workstations to every engineer's desk. SDRC had concerns about network performance, so we told them of the idea of LAVC workclusters with network traffic isolated from the backbone with bridges. SDRC had concerns about the decentralization and management of their development source code, CASE tools, quality assurance, etc. We advised them on how they could use the DSS products to ease management in a distributed development environment.

We called this strategy our "Workcluster Vision" and promised we could empower their engineers by putting a VAX on each desk. Local compiles and links into binary code would unburden the CI cluster, and the development code could be almost effortlessly copied back into the centralized CI cluster database by using the DSS products.

Ultimately SDRC decided to buy into the workcluster idea, with a few modifications, and by the spring of 1989 had received three shipments containing their first 58 workstations of what would ultimately be a total of 600 workstations, one for every engineer.

SDRC's problems started then, and continue to this day, eight months later.

² \$25,000,000 in computer capital assets including: 8820, 8700, 8650, 8550, two 785s, and sixty VS3100s, totaling 972 MIPS; 170GB storage on Digital RA and RZ drives; six HSCs; twenty-three DECserver 200s, three DECserver 500s; 41 miles of LAN

2.1 Problem History

Of the first 58 VS3100s, one third were DOA. Not only was this expensive to Digital, but the local Field Service office had difficulty getting parts to fix the workstations in a timely way. Understandably, this did not go over well at SDRC.

As I noted above, SDRC did not completely subscribe to the workcluster implementation, in that they decided to configure the workstations hosted by their CI cluster; after all, the SPD at the time clearly stated that a 42-node mixed-interconnect cluster was legal and supported. They were also led to believe that performance of the CI cluster would not suffer significantly. SDRC was told that "Digital has several LARGE development clusters used by our engineers to develop VMS". Unfortunately this was wrong. Performance has turned out to be a serious problem.

The first problem was one of getting the proper version of VMS to support these configurations, a minor problem to Digital but a significant problem to SDRC, which had a major investment in workstations sitting idle because the software was not ready. After several weeks of delay they eventually got the correct VMS versions, only to discover that the cluster began to degrade with only 32 nodes configured. At 42 nodes, the cluster had degraded by 60%. We were ineffective in attempting to restore performance on their cluster.³ Digital ended up lending SDRC a 6000-series machine to recover the lost performance that SDRC felt Digital had caused.

After a very frustrating period of time where all parties tried to get the mixed-interconnect cluster performance up with little success, SDRC finally decided to implement the workcluster ideas that we had originally suggested. SDRC configured several workclusters consisting of VAX 3900s equipped with two Ethernet adapters which would service 15-20 VS3100s. Each workstation was configured with 16MB memory, local page and swap disk, and color monitors. They chose not to implement the bridge isolation recommendation, and in fact, Ethernet traffic seemed not to be a problem. They served their CI-based source code files and CASE tools using the DSS products.

Running VMS 5.1-b, the development environment finally became adequate — not spectacular, but sufficient. In the meantime Hewlett-Packard, Silicon Graphics, and IBM were making very good impressions with their high-performance graphics and X offerings, as I'll describe later. By this time we had converted their engineers to X, a "win" for Digital.

Late this summer, SDRC again approached Digital for the specifications about performance around an upgrade to VMS V5.2. They were essentially told that they may see some slight performance degradation in trade for VMS enhancements. With that information SDRC prepared to make the upgrade. Part of that preparation was to test and benchmark identical configurations between 5.1-b and 5.2 (and FT 5.3).

To SDRC's utter dismay, some of their tests revealed gross degradation by a factor of 2 between VMS 5.1-b and VMS 5.2. On one test SDRC linked their code with the debugger, set a breakpoint, and measured the elapsed time to run the code to the breakpoint. While running VMS 5.1-b, the elapsed time was 6.15 minutes. To get to that same breakpoint,

³ In fairness, both parties were at fault. SDRC was not properly tuning their systems, and used many questionable techniques recommended by third parties. They were using RAM disks, defragmenters, etc. not recommended by Digital. Digital balked at helping because of some of this. And moreover, Digital would not readily help unless SDRC paid for the service. SDRC balked at paying for the service.

running the same configuration with VMS 5.2, elapsed time was 15.07 minutes. That's *MINUTES*, not seconds. VMS field test 5.3 added a couple more minutes. These systems were tuned to the best of SDRC's ability.

It was barely acceptable for an SDRC engineer to wait for over 6 minutes to get to the breakpoint under VMS 5.1-b, and absolutely UNacceptable for them to have to wait 15 minutes for the same thing, just to migrate to VMS 5.2.⁴ To SDRC, again, this is Digital's problem, not theirs.

SDRC solved the problem by adding memory. Eight additional megabytes of memory reduced the delay to the VMS 5.1-b level, but SDRC has been penalized in the past for buying small amounts of memory to solve a point problem and later discovering they needed more memory when something else went awry. The trade-in and repurchase of larger memory arrays never proved cost-effective, so SDRC decided again to bite the bullet and purchase the full complement of memory that the VS3100 would accommodate. They decided to order 16MB per machine, providing a total of 32MB. \$7450.00 per array × 58 VS3100s amounts to \$432,100 just to, in their minds, "maintain performance".

SDRC has been battling these performance problems for eight months now. Grappling with this has impaired SDRC's ability to make the strategic buying plans they must make to stay ahead in their market.

The straw that broke the camel's back, however, was SDRC's discovery of the *cost per seat* of the configuration necessary to implement Digital's workcluster vision. SDRC has found that with a VAX 3900 properly equipped to service 15–20 workstations, and with the workstations each requiring a page and swap disk, 32MB of memory, color, licenses, DSS software, etc., the cost per seat is between \$35,000 and \$40,000. They see that as a very high price basically just to serve X to those seats, do local editing, and perform some local compiles. And again, many vendors — HP and SGI and soon IBM — offer high graphics performance in X-based workstations that SDRC is very pleased with.

2.2 Introducing the VAX 9000 Backfired

This sad scenario is not over, however. As I stated earlier, SDRC ports their code to every major workstation vendor. All the vendors are frequent visitors at SDRC. HP, after hearing and watching Digital fail in our strategies, has offered to "buy Digital out": they are willing to replace our development environment with theirs for nothing, just to have the presence.

IBM is also at SDRC; they've quietly placed their new RT/SGI graphics workstation there under field test. The performance of that system surely has brought back some credibility to IBM again. And unfortunately, information that SDRC has seen in the press and heard about the VAX 9000 has worked against us to some extent and helped IBM.

SDRC is naturally interested in Digital's high-end strategies and asked for information from Digital. We have provided this information under non-disclosure to SDRC's Senior Vice President and Technical Officer, along with four other senior VPs. The presentation went very well, but SDRC's research through trade journals and other sources of misinformation has positioned the 9000 as a mainframe.

⁴ Again, in fairness, it seems likely that something was not setup properly. Other sites have not seen degradation like this. Digital is addressing this now.

Whether the VAX 9000 is a mainframe or not is not really important but positioning it as a mainframe gives IBM some fuel. SDRC doesn't yet fully understand about the VAX 9000 machine's balance of I/O and internal speeds. Some of SDRC's areas of deepest need will be satisfied by the 9000 related products, such as the native XMI adapters, disk striping, the XMI Ethernet and FDDI adapters, etc. These translate to balancing CPU performance with I/O and memory, network performance, and file service over the network. The message has often been clouded by the trade journals which have been simply wrong on occasion in this area. IBM has exploited that.

With IBM's excellence in marketing and salesmanship, they have planted the idea in SDRC's minds that Digital does not know how to deliver all aspects of "mainframe" computing, like balanced I/O and throughput. In fact, they have convinced SDRC that if they are going to enter the mainframe world, it should be on a real (IBM) mainframe. SDRC appears to be buying the idea, and has made comments strongly indicating a migration to a development environment that will be dominated by IBM within two years. This now make more sense to SDRC. "True" mainframe performance on an IBM 3090-class machine, serving IBM's RT/SGI graphics workstation, the performance of which re-opened the door for IBM, is a picture that IBM seems to be painting well, as well they should. SDRC is every bit as strategic (if not more so) to IBM as it is to Digital. With the myriad of problems SDRC has had with Digital's workstation vision, Digital is very quickly becoming inconsequential. Even our success with X might work against us as other reputable vendors produce superior X-based workstations.

So Digital's workcluster vision dimmed at SDRC. As a matter of fact, SDRC now believes that workstations on every engineer's desk are a stopgap measure until a high-performance color X terminal is introduced. Again, we, Digital, converted them to X, and we needed to. But there are several reputable companies besides Digital ready to introduce X terminals. Unfortunately, Digital's X terminal is unlikely to fill the order for SDRC, because it is based on some very safe, possibly clever, but very old technology, and it won't be color for some time. It isn't designed for high performance, and as a result it can't be used for applications like SDRC's, or for any application that requires color, serious imaging, or very rapid X service. If Digital doesn't form an alliance with someone who builds a high-performance X terminal, a Japanese company is very likely to have one that's cheap, fast, and reliable. Why should SDRC buy from Digital then?

Digital loses the development environment (or file server), and the engineer's desktop. This means that Digital loses, in one account, \$3,000,000-\$4,000,000 in internal sales alone to SDRC. We also lose \$1,000,000 in the highly profitable Field Service business.

2.3 What SDRC Wants for Their Own Shop

SDRC's vision of workclusters is evolving from "a workstation on every engineer's desk" to "a high performance color X terminal on every engineer's desk". Internally they will have a few high-performance 3D graphics workstations, but *cost per seat* will be the deciding factor. And at SDRC the engineers don't really care where their code compiles: they just want high performance and windows. They also do not care to have to attend to the "management" complexities of workstation-based environments.

3 Leveraged Sales Are Eroding

Losing several million dollars of sales internally is significant, but not as significant as the almost certain loss of the leveraged business, which amounts to \$40,000,000-\$50,000,000 a year, over \$60,000,000 if we had a competitive worksystem, according to SDRC.

Earlier I said that SDRC sold their "Ideas" product into DEC shops 65% of the time. In the last year and a half that has diminished to 12%, with HP and SGI taking the lion's share of the market.

Why? Digital does not have real graphics performance compared to our significant competitors. Unfortunately we have proven before, and will probably prove again, that we will not soon have the performance, even despite heroic efforts taking place in Palo Alto and elsewhere on Scanproc and Pixelstamp. We have a history of being too late, and with technology behind the leading edge. Our knowledgeable customers such as SDRC, who know the graphics world as well as anyone in the marketplace, realize this and say so in no uncertain terms.

For some specifics: both Scanproc and Pixelstamp have price-performance problems. The price of one or the other could be a factor of 5 or 6 times competing technology. Both Scanproc and Pixelstamp have problems with the X window system. Neither provides direct access to the frame buffer. Both are likely to have performance problems with some X drawing commands. And both may well slip their schedules again.⁵

New VAXstations will use Scanproc. Scanproc is better than the GPX chipset, but may not even be as good as the dumb-frame-buffer on the PVAX for short vectors; and it probably will not fare well at all against HP, IBM, and SGI⁶ graphics performance. For a slow processor like the CVAX, acceleration is even more important. With more hardware assist, more of the processor is left for the application.

Some may argue that the high-performance graphics market is a niche marketplace, and that 3D (which is what SDRC is pushing to the limits) is an even narrower niche. But that isn't the case. From the field's perspective, there are all sorts of indications that high-performance 3D will boom in the 1990s; indeed, all of SDRC's customers *want* high-performance 3D, but they can afford only a mix of 2D and 3D. That mix is something like one 3D machine for every ten 2D machines in a typical SDRC customer shop, all color. The 1:10 ratio is because most can only afford 2D machines, not because they don't want 3D.

SGI, HP, and soon IBM will push the world into 3D. Competition will drive the prices down. Many competitors already have the performance, leaving Digital out of the picture based on price-performance except for our most diehard customers. We are essentially out of the picture already at SDRC, because we simply don't have the necessary performance on the VS3100. SDRC is having a very difficult time porting to PHIGS on the DS3100 (which has no graphics acceleration), and although we are trying to work a strategic agreement⁷ with

⁵ Mike Breen has written several reports addressing many of these issues.

⁶ SDRC has asked several times "Why doesn't Digital just buy SGI's graphics subsystem and library? IBM saw its importance and did so." We had a rather unique opportunity two years ago to do so. Indecisiveness seemed to set in.

⁷ This agreement is for SDRC to port their latest, unreleased version of "Ideas" to our next workstation product six months ahead of ports to our competitors *PROVIDING* we deliver field test machines in January, and that they perform adequately. SDRC is still hopeful, but not optimistic that we can do it.

Digital Internal Use Only

SDRC around our next generation machines, our graphics subsystems will be pressed even more by the next generation of SDRC's software.

Finally, with the diminishing leveraged sales of the Ideas software on Digital platforms, SDRC now sees much more powerful arguments for changing their internal development environment to another vendor.

4 We Failed in Several Ways

While many factors were at play at SDRC, we can't dismiss the trouble there as just anomalies of delivery, service, configuration, tuning, crossed communications, and QA. Even if we had done all those things right, we still can't offer them the price-performance they need in their market and for their own development shop. They've come to believe all we told them about X and distributed computing, but unless we take some new steps there, we won't be able to deliver on those messages as well as can our competitors.

5 Digital Will Lose Its Presence in the Industry

So to summarize: if nothing changes, over time Digital is likely to lose \$4,000,000–\$5,000,000 of internal business alone at SDRC. Worse yet, we will lose leveraged business of \$40,000,000–\$60,000,000. And, finally, Digital will lose a long-standing relationship with a very good customer.

Those events and consequences are depressing, but worst of all is SDRC's matter-of-fact attitude that Digital's time is past. They enjoyed working with us, they liked being part of the successful decades at Digital, but, in their minds, they must move on, with technology that simply is better than ours.

That is what is hardest to swallow; that we will simply lose presence as a significant vendor in this marketplace.

6 There is a Solution

There is a solution to these problems. And what is better is that many of the issues can be solved in a way very appealing to our customers, which is, after all, why we are in business in the first place.

6.1 Buy Graphics Technology from Jupiter Systems

Digital has been known for outstanding engineering. In many areas we build superb systems, and if we could soon build high-performance graphics workstations and X terminals ourselves, old customers would stay with us and new ones would flock to us. But second best in this marketplace is simply not good enough — and it has been debatable whether we'll be even second best.

If not, then we should *buy* the technology. Digital has some very talented people, but we don't have a monopoly on talent. And, in truth, any successful engineering organization needs only one very good engineer to come up with a world-class product or idea.

Jupiter Systems is a small company that Mike Breen discovered. They have a very small and talented group of engineers, who have developed a technology that we can buy today. Jupiter has been building high-performance graphics systems for years. I will not belabor the technical merits of what this private company has engineered. Many of Digital's graphics and X engineers concur that systems and X terminals based on Jupiter's technology can exceed much of what Digital is engineering now and into the short- and medium-term future.

Jupiter has developed a graphics subsystem that Digital could buy to solve an array of problems, problems such as image servers (with the evolving image extensions to X); a phenomenal 2D graphics engine on which X performance excels; and, with some engineering, competitive 3D graphics. The ultimate irony is that the cost to build these subsystems is *LESS* than anything we currently have under development that I'm aware of. I won't go into the technical details here; Mike Breen has written extensively on this topic, with all the details of price and performance, both Jupiter's and ours. There appears to be no serious argument with the substance of his papers.

I urge the reader to ask Mike for copies of his work. Mike's research presents the technological facts pointing to a solution SDRC and other customers would like us to provide for their array of needs.

We are a systems company. That means if someone else has pushed the technology farther and faster than we, we can honorably buy it. Competitive pressure compels us to, just as it did with the DECstation 3100. As I've tried to point out with SDRC, we simply can't afford delay. A buyout will undoubtedly be painful to groups now working on internal solutions, but buying the technology is very right for Digital now, not in six months. In this case, we can buy superior graphics technology and regain leadership, something that our customers are almost begging for. We don't even have to change our culture to take advantage of this opportunity. We have a history of dual development teams where the "best team wins". Our customer then win.

Or, finally, give someone else the approval to make this happen. It's important enough for many of us to be making an effort to urge this to happen, even though, like everyone in Digital, we're proud of what our company can do itself, and we prefer to succeed entirely by our own efforts. There are few in the field who would turn down the opportunity to help drive this in any way we can, myself included. After all, we are the ones who face these customers everyday. We are the ones they come to first, seeking answers to their problems.

We may not be able to afford to do otherwise. Buying Jupiter's technology might save Digital a considerable amount over doing it ourselves, possibly millions of dollars. In any event, it's timing that's of crucial importance. We must do something or the window of opportunity will close, leaving us in the expensive position of either abandoning the 2D and 3D market or trying to buy it back later from very able and reputable competitors.

7 Conclusion

Mike Breen foresaw this predicament in precise detail nearly a year ago. I urge the reader to study his documents. They are very convincing and have held up under scrutiny. In recent ones he presents the technical substance of what is happening at several of our most technical and most important accounts, and gives full details of why Jupiter's technology is both more powerful and less expensive than what we have in development.

Digital Internal Use Only

The bottom line is this: it is in our power to expand our vision of the workstation marketplace and keep the business, even in ECA alone, of companies like SDRC, Bailey Controls, ABB, and Schlumberger. That new vision can be profitable to Digital, and more importantly, elevate us back to respect in the technical marketplace.

Ultimately, we have to act. We can't wait.

8 The Situation At Several Accounts

Table 1: Representative Accounts

Account	Digital's Position
SDRC: Requires high-performance 2D color graphics and X servers. SDRC doesn't actually recommend platforms to their customers, but they invite their customers to view SDRC products running on all of the platforms, often side by side. Digital's relative performance disadvantage is apparent.	Current 2D graphics and X are usually unacceptable; future offerings will be poor in comparison to those of IBM, HP, and SGI.
SDRC: Requires high-performance 3D graphics workstations.	Has no current offering (Firefox is somewhat of an aberration); future products may slip, will be expensive.
SDRC: Needs a high-performance color X terminal to migrate to their next stage of X-based interactive computing (X Client and file server, with color X terminals); They will evolve to this simply because of the reduced cost per seat, and because X terminals will be commodities. SDRC's customers will likely follow suit.	In the short term, Digital's X terminal has neither the color nor the performance for this market. Our color model will be released much later, but performance is unknown.
Bailey Controls: Bought into the workcluster model; is balking at the cost per seat; could easily migrate to the X terminal model.	The X-terminal and X-client model works well here and Digital fits with very good X-Client/file server.
ABB: Cost per seat of the workcluster model; performance issues.	Digital is directly competing with higher performance HP's models.
Union Switch and Signal: Cost per seat of the workcluster model.	The X-terminal and X-client model works well here.
Schlumberger: Issues identical to SDRC's; hasn't complained as much.	See position notes on SDRC above.
McDonnell-Douglas: Recent loss to HP of M-D's workstation business amounting to several tens of millions of dollars because our "premier" workstation, the Firefox, isn't even close in performance to the HP workstations that won the contract.	No competitive performance. No known follow-up strategy.
GE Aircraft: Having a severe problem with Digital workstations arriving DOA or breaking when being set up.	QA degradation is surprising; Digital has always excelled at solving this kind of problem. There is a problem with parts availability in the field as well.

Digital Internal Use Only