

MICROSOFT RISC PC TASK FORCE REPORT

Responses to the Compaq SPARC Threat

**HIGHLY
CONFIDENTIAL**

**NOTE: This memo is very confidential and discretion should be used in copying
or distributing it - even within the company.**

ADVANCED TECHNOLOGY AND BUSINESS DEVELOPMENT

MICROSOFT CONFIDENTIAL

**MS 5012188
CONFIDENTIAL**

MICROSOFT RISC PC TASK FORCE REPORT

Responses to the Compaq SPARC Threat

**HIGHLY
CONFIDENTIAL**

**NOTE: This memo is very confidential and discretion should be used in copying
or distributing it - even within the company.**

ADVANCED TECHNOLOGY AND BUSINESS DEVELOPMENT

MICROSOFT CONFIDENTIAL

**MS 5012188
CONFIDENTIAL**

Table of Contents

TABLE OF CONTENTS

TABLE OF CONTENTS	2
INTRODUCTION	3
Sun's Strategy	3
Probable Compaq Plan	5
Compaq's Motivation	6
Impact on Microsoft	8
PC Industry OEMs	8
PC Industry ISVs	9
Networking Business	10
LIVING IN A SPARC WORLD	10
Windows layer for SunOS	10
NT Windows for SPARC	12
NT as a base for SunOS	13
Applications on SPARC	13
Conclusions	14
FOSTERING SPARC ALTERNATIVES	15
Uniting the MIPS Community	16
Managing PC ISVs	19
Wild Ideas	20
Conclusions	21
RECOMMENDATIONS.....	21

HIGHLY
CONFIDENTIAL

1. INTRODUCTION

We have recently learned that Compaq is seriously considering a project to enter the workstation business with a SPARC and UNIX based machine. This memo covers some initial thinking on what this means to Microsoft, and how we should respond.

Obviously the best thing would be to have them change their minds. The question of how to cause this to happen will be covered elsewhere. This memo assumes that we will immediately begin a dialog with Compaq to try convince them to do something else (nearly anything else would be an improvement), but in parallel we must start to plan, and act, on a response which assumes the worst. It is fairly safe to assume that Compaq will wind up in one of a couple of modes:

- **Rapidly reach a formal decision in favor of SPARC.** Once their consensus driven process locks in on a decision it is hard to change, and we can assume that it will be months before we could turn them around. Even then, we will need dramatic new data to change their minds, and there is only a small chance that we could win even then. A further complication is that to do SPARC they are likely to enter into some commitments early on which are tough to reverse - it is not just engineering work.
- **Return to confusion on the RISC issue.** This is the state that they were in for the last several months (although some of that was cover for their SPARC investigation), and it is possible that we could get them back into that mode of operation. They could wind up right back in the SPARC camp, or do something else random.
- **Run down multiple paths.** They could start working on the SPARC plan without a formal decision, just as they worked on the MCA bus machines while thinking about EISA. This would still be a very dangerous situation, but it would increase the likelihood that we could get them to turn around later. The problem with this is that the likely scenario for their SPARC project involves building or buying a sales force and other activities which can't easily be cancelled late in the game.

Another way to put this is that we will either lose very fast, or we will be in a long holding pattern from which we could still lose. In either event we must proceed at once with our response - there is no point at all in waiting for Compaq to turn around, or in hoping that they will quickly come to their senses and join our present RISC plan.

Finally, I should note that the ideas below are the result of conversations with many people.

1.1. Sun's Strategy

Our present understanding of Sun's overall strategy is as follows:

HIGHLY
CONFIDENTIAL

- Stay the course with direct sales and the conventional UNIX workstation market. They will rely on their present approach as the mainstream core of their business for the next two years. The growth rate is large, and they can comfortably use this to finance their assault on other markets, and give them enough time to get the pieces in place. The key niche markets that they will exploit are technical workstations and software development within corporations. This will broaden to cover an increasingly large set of customers.
- Experiment with other channels. This will occur through limited test cases such as their deal with Micro Age, and through SPARC clone companies such as Northgate, Compuadd and others who are expendable missionaries in new markets.
- Build an arsenal of ISV support. They do not have a sufficient set of desktop productivity applications to really threaten PCs for control of mainstream office computing, but they get more support every day. Their present growth rate even *without* having such apps will take them to 300K - 500K units/year run rate within the next year to 18 months.
- Once they are ready, make a major push on the PC market. Within the next two years they will be in an excellent position to directly assail PCs with widespread retail distribution of both machines and binary software packages. In early 1992 they will be selling at the 300K - 500K machines per year run rate, will have a critical mass set of major applications, and they will have a commanding price/performance lead on the x86.

This all assumes their present level of industry support. Even though Compaq is not likely to ship any appreciable volume in SPARC until late 1992, their endorsement will clearly help.

Note that the strategy above is NOT aimed at bringing RISC into the PC market - rather it is trying to grow the workstation market to the same volume levels, and the same distribution methods as PCs. This is a critical distinction. Sun is creating a parallel world to the PC industry, in much the same way as the Macintosh is a parallel world to IBM compatible PCs. It is not a high end PC, but rather a new beast which has some key differentiating features:

- RISC. This means three direct benefits over the x86 - 32 bits, better price/performance, and higher absolute levels of performance. Indirectly, the open processor model will ensure that SPARC's lead over the x86 and closed processors will increase over time.
- UNIX. This is a mixed blessing, but their positioning is to try milk as much as possible out of "open systems", and the supposed technical quality of UNIX.
- Networking and connectivity. Nobody would accuse UNIX of having an elegant or even very good networking story, but since it has been put into place over a number of years it does work, is mature, and has been ported to all manner of machines and networks.
- The perception of being high tech. Sun and other workstations have the cachet of being sophisticated, powerful state of the art machines, and this aura helps set them apart from the PC industry.

Despite the enormous momentum behind Windows 3, as long as Sun can position themselves in a significantly different market, they are largely immune from assault and can continue on the strategy

above. An analogy that we've used in the past is the case of the Macintosh, which was introduced at a point of unprecedented strength for the character mode IBM PC. This did not eliminate the Mac, and in fact it is hard to imagine that any amount of increased volume in character mode PCs would have stopped it. The Mac used GUI to put itself in a class by itself, and thus have the breathing space to grow, and as long as the SPARC world uniquely enjoys the features above it will have the same kind of chance.

Our strategy has been to deny key points of differentiation to SPARC by broadening the Windows work to include to include them. The most dramatic and most important is RISC, and that has been the item which has been discussed the most. The high tech image is addressed by the Power PC features. Networking is also being addressed. Although not unique to RISC, it is a key part of our long term systems strategy.

1.2. Probable Compaq Plan

Our present understanding of the Compaq strategy is as follows:

- Enter the UNIX workstation market with a SPARC based machine. This will probably be of their own design rather than a SPARCstation clone. This would be positioned very clearly at workstations, and at Sun's present and near future market. They have an explicit goal to keep it very separate from the x86 PC market to avoid any negative impact on their current systems, and they think that this separation is pretty easy to achieve.
- Use a Compaq direct sales force to distribute the machine. This choice is determined by three factors - their desire to compete head to head with Sun in the traditional workstation market, the need to keep this activity separate from their present distribution channel and finally they have a long term goal to build a real direct sales force, and this project provides a convenient opportunity. There are hints that they may acquire a mini or workstation company such as DG or Wang to get an established sales force in a hurry, but they may just build one from scratch.
- Create a limited consortium. Compaq will attempt to balance the necessity of having industry support to attain critical mass with their large desire to have a proprietary advantage. They do not be vulnerable to cheap SPARCstation clone kits (although exactly how is not clear to us yet). The data to date suggest a group of 3-4 companies with sales > \$2 billion, and with little enough clout that Compaq gets "51 % of the votes" (in Gary Stimac's words).
- Rely on the Compaq name and prestige. They do not seem to mind us doing something else for RISC which does not involve them, because they do not feel that it will amount to much. They think that the Compaq name and support is critical to the success of any such machine, so other efforts will be just so much noise.
- Out execute Sun and the SPARC clones. In attempting this strategy they are clearly counting on superior execution to win them a good slot in the SPARC pantheon - i.e. that there is room for two major players (themselves and Sun).

HIGHLY
CONFIDENTIAL

The time frame for this is uncertain, but Stimac said that they could have sample machines within 9 months, and ship within 12 months. This would be consistent with building a machine from scratch (you can manufacture an existing design such as the LSI Logic SPARCkit in much less time - say 3 months).

The Compaq strategy outlined above is much more of a pure workstation approach than the one that Sun is on. This is because Compaq has a huge PC business to protect. Ironically, by endorsing SPARC they are giving Sun and others a powerful means to attack the PC business even if Compaq itself achieves a way to isolate their own workstation business from PCs. This is a very key distinction between what they are planning to do and the IBM RS/6000 strategy which they admire so much. In IBM's case the RS/6000 has no life of its own - IBM controls it completely and also benefits completely if it is successful.

In Compaq's case, SPARC *does* have a life of its own, and Compaq is far from being the sole beneficiary if SPARC wins. Sun and others can directly attack the PC industry, and in doing so compel Compaq to either follow suit and cannibalize their PC business, or resist the trend and lose in the SPARC market. For example, if Sun continues to up the ante in aggressive pricing, marketing aimed at PC end users, ISV evangelism and other anti PC activities, Compaq will have to match Sun to remain competitive in the SPARC market.

1.3. Compaq's Motivation

One view of Compaq's interest in SPARC is that it is simply an extension of their desire to ape IBM, and more generally to be a quality implementor of other people's strategies (i.e. superb knock-off artists). They need to copy any strategy that IBM has, ergo they need to have an answer to the RS/6000. Given that there are some nearly insurmountable barriers to directly cloning the RIOS chip set, they looked around to see who is the next biggest player (Sun) and set off to copy them and win through superior execution. Compaq has recently taken the same approach in the printer business.

Stepping back, there are several likely factors which are motivating Compaq:

- Cloning the strategy of moving into workstations. This was discussed above.
- Obtaining incremental revenue and market share. The workstation market has a faster growth rate than the PC business, and they perceive that they have sufficient skills to beat the present competition.
- Get involved with RISC without jeopardizing their PC business. They understand that RISC will be important long term, but they remain terrified of anything which introduces RISC in a fashion which might in any way reflect on the PC market.
- Get involved with UNIX. They are getting a much more favorable view of UNIX than in the past and this gives them a way to hedge their bets. I do not believe that they have an explicit goal to get out of the Microsoft dominated world, but having a strategic hedge is sensible.

HIGHLY
CONFIDENTIAL

- Build a direct sales force without disrupting their present channel. The workstation business gives them an opportunity to do this without upsetting their dealers.

One interesting question is the degree to which we have accelerated Compaq's concerns in these areas, or done things to cause them to move to SPARC. It is likely that our efforts to involve them in the RISC PC, and our emphasis of the Sun threat has highlighted the importance of this area, but in general I think that at most we got them thinking about this six months sooner than they might have otherwise.

During the meeting with Stimac there were a number of comments that indicated that they were none too pleased with Microsoft's power in the industry. One concern that we had during the meeting, was that Compaq was upset with Microsoft taking an active role in trying to define a hardware platform and pushing them and the rest of the industry to RISC. After thinking about it, and considering their SPARC strategy in more detail, it appears to me that anti-MS feelings, including fear about our hardware prototyping efforts and the role we have intended to play in RISC PC have essentially *nothing* to do directly with their basic strategy.

They appear to be doing exactly what one would expect from Compaq if we had never told them a thing about RISC, and instead they had discovered it on their own. Their natural tendency is to clone the winner rather than to innovate with a bold new approach. They would never undertake an initiative like setting a new RISC PC standard by themselves - there are too many unknowns for them. If you set aside the notion of doing their own standard, what else could they do except clone an existing one? Sun is their choice over our RISC PC because they are more established whereas our plans are still on paper, and there is less perceived risk to their existing business.

Certain individuals at Compaq probably do hate us, or are uncomfortable with our power, and we should never get lulled into thinking that they are not just as envious and jealous as the rest of the industry. Probably the biggest sin that we have committed in their eyes is that we ship products before they are ready - this offends their perfectionist sensibilities, and in the case of RISC they are worried that our haphazard approach could hurt them very badly by impacting their present market. These personal factors help grease the wheels for any plan that shifts power away from us, or reduces Compaq's reliance on Microsoft, but there is no evidence to support the contention that they are on an anti-Microsoft jihad as a matter of company strategy. Nevertheless we should not be fooled into thinking that they like us, or will cut us a break in any way - if the long term effect of the strategy is to put us in our place, so much the better although it is not an explicit goal.

1.4. Impact on Microsoft

Sun's clear goal is to be the system software provider and the leading hardware company in the SPARC market, and have that market become the key high volume segment of the computer industry. From a technology standpoint they have a fundamental advantage over the x86 based world, and there is every reason to believe that in the long term RISC based machines will completely replace the x86 for all desktop personal computers. If Sun succeeds, our systems business will die along with the x86.

That is the big picture view, and it shows why we need to be excited - billions of dollars, and Microsoft's identity as a company that is important in systems software - hang in the balance.

Narrowing down our focus a bit, it is interesting to project the near term impact of a Compaq SPARC project on our part of the world.

1.4.1. PC Industry OEMs

One can argue whether Compaq is just following a clone strategy, or has some deeper plan, but when it comes to the rest of our OEM customers, there is little doubt, especially among those who are focussed primarily on the PC market. They are not so ambitious as to think that they need to emulate IBM in every way, but by the time Compaq is doing it too, it will get their attention. The rush of interest in MP and server machines following the SystemPro introduction is ample testimony to this fact - despite the fact that nobody (even Compaq) is burning up the sales charts in this market. The minute that Compaq announces a RISC strategy, every other OEM is going to give it some serious thought. Some will wait, some will rush in, but they will all think about it.

The reaction to a Compaq SPARC announcement will group OEMs in three basic categories:

- Pure PC companies. This includes big ones like Zenith and Tandy, as well as the second and third tier people. SPARC will be the leading contender for most of the them, because it is the simplest thing to do - just copy Compaq. The only thing to consider is whether their distribution channel can handle it, but many of them will invert this problem into a perceived opportunity - i.e. that their *advantage* is that they're in a different channel than Compaq with the same kind of machine.
- PC companies which are already in the workstation and mini business. These guys will be caught in a tough position, because Sun is their nemesis, and they will be loathe to support them. The canonical example here is HP - they are the second largest workstation company, and also have a PC business.
- Workstation and minicomputer companies. People without a substantial PC business include companies such as DEC, Silicon Graphics etc. These people would not normally care at all about what Compaq does (or at any rate it is a second order phenomena for them), and are unrelenting in their opposition to Sun. They will see an incredible boost of momentum for Sun and will be very interested in doing nearly anything else as a matter of survival.

The last two categories of companies are the simplest ones to predict - they are directly threatened by SPARC (even more than we are) and they will react sharply in the opposite direction. Ultimately they may give up and make SPARC clones, but they will almost certainly make one last attempt to beat them first. The motivation is pride, inertia and the fact that Compaq (at number two) will have taken last really desirable spot in the SPARC line up (not to mention Compaq's little consortium). This gives them nothing to lose and everything to gain by going in any direction other than SPARC. The primary distinction between the two is that workstation companies in the PC industry will naturally look to Microsoft for a solution to this problem, whereas those who are not big customers of ours wouldn't normally think of that. In either case, they are very approachable for a counter SPARC strategy.

The pure PC people are another matter entirely. They will have a much more direct reason to just sign up with SPARC. They are especially vulnerable to the same pitch that Compaq itself is falling

HIGHLY
CONFIDENTIAL

for - enter the workstation market as a way get incremental revenue without effecting your PC business. The keys to convincing these people not to go with SPARC are to play on three things:

- It will hurt your PC business. SPARC is the natural enemy of PCs, and it is too late to side with SPARC in this fight because all of the good seats are already taken.
- You can adopt a RISC strategy which actually benefits your PC business. This is the Power PC message. Since most pure PC companies do not understand the UNIX market, they would be much better off with a way to address RISC which leverages the thing they do understand, and their present users and distribution channel.
- It's time to turn the tables on Compaq. This only works for the larger PC companies, but it can be effective. These guys are clearly in third place behind IBM and Compaq in the PC world, and they would still be third place (or worse yet) in the SPARC world. If they are offered a chance to move up in the hierarchy, and still have the comfort of having major forces in the industry supporting them, they will see this as an opportunity.

This covers the near term reaction to a Compaq SPARC. The long term consequences depend on the course that we follow. In the absence of any MS lead RISC project, I expect that we will see over half of the PC only crowd offering SPARC machines within a year of Compaq actually shipping. The only reason that I do not say all of them is that the SPARC market can only absorb a certain amount of growth until mass market binary applications start to appear.

1.4.2. PC Industry ISVs

The effect of a Compaq SPARC announcement on PC industry ISVs will depend a lot on our strategy and how it relates to Windows. It will also depend a great deal on how active Compaq is getting ISV support. If they are really hard core about the notion of separating the UNIX and Dos parts of their world, then they will not embark on a big ISV program unique to the machine and will instead let Sun's existing efforts handle it. On the other hand, if Compaq has some differentiating feature over Sun (for example they decide to push Motif rather than OpenLook) then they will have to go to ISVs.

If we do not have a compelling way to tie RISC into Windows, and execute on it very well, then we will see a lot of PC industry ISVs look at SPARC as being enormously more interesting after Compaq than before.

1.4.3. Networking Business

One of the primary attractions to UNIX from Compaq's point of view is that it has mature networking. We can expect that their direct sales force will be pushing this very strongly as part of selling the machines. This suggests that networks will become one of the first areas where we will conflict develop between the x86 and SPARC sides of Compaq's business - which machines get pushed as the departmental network? What is the server OS? If you buy one kind of Compaq machine the sales rep will help you out, but the other kind you have to call this dealer.

This has a lot of potential to negatively impact our network business, and any attempts to establish OS/2 as a server OS. A natural separation for Compaq to make is to say that simple file and print

HIGHLY
CONFIDENTIAL

services belong on x86 PC based servers up to the SystemPro, since the issue is mainly one of I/O throughput rather than processor speed. The dominant software here would be Novell. When you need to do database operations, or any other kind of compute intensive server task, it will be hard for them to avoid selling the SPARC machine.

2. LIVING IN A SPARC WORLD

One way to view this development is that we should learn very fast how to live in a SPARC dominated world, because that is going to be a reality. There are several degrees of emphasis that you can place on this. At the very least, Sun and SPARC will continue to be a viable system for a number of years, and we could view this as an incremental revenue situation (much like Compaq). At the other extreme we could just admit that SPARC is likely to take over from the x86 within the next five years and we should jump on board as hard and fast as we can. In any event, the sections below describe some of the opportunities that we could approach assuming that SPARC is going to be important.

There are three different approaches for distribution that we could take in working with SPARC - one would be to deal directly with Sun, the other would be to work with Compaq and the third is to try retail. In the discussion below we will assume that Compaq will have the same system software base that Sun has, and this means that for all intents and purposes we have to work with Sun. Retail is a poor choice for the present SPARC market because it is all done by direct sales. By the time that retail is effective, it will be too late for most of the approaches below (except applications).

2.1. Windows layer for SunOS

The idea here is to take the portable version of Windows being developed for NT, and make it work on top of SunOS instead of the NT kernel. There is considerable precedent for such a project - we negotiated a deal just like this with Sun several years ago, but it fell through at the last moment when they were feeling their oats and thought they didn't need us (that was right at the moment when they first signed up with AT&T to control UNIX).

If there wasn't an existing window manager and look and feel for UNIX this would be straightforward - it would be to UNIX like Windows is to Dos. The proposed deal we had with Sun would have had a portable version of PM be the only GUI interface with SunOS, and be bundled with every copy. Now that there is OpenLook in the default position, a number of problems arise:

- Sun won't bundle it. Obviously we can try to get this, but it is hard to see why they would want to at this stage. They can sell at the rate of 200K+ machines per year without this today, and there is no present sign of their growth slowing down. They are getting a slow but steady series of defections from the PC ISVs, and with Compaq on board they have a good shot at the rest. As a company, they are a lot less enamored with the "big deal syndrome" than we are and they are likely to just go out and compete with their product. There is also the strategic cost of letting us get into a strong position in their system software business.

HIGHLY
CONFIDENTIAL

- System utilities and UNIX apps will follow OpenLook. They are too far down the road to pull back now. The fact that there would be two different looks, and two different APIs is a pain to both end users and developers. Over time people would tend to view one of the approaches as being a second class citizen - if that stigma fell on Windows (and Sun would have every reason to make that so), then ISVs would migrate to OpenLook and we would be eliminated.
- Windows on SunOS could be clumsy. Although the core kernel issues are not difficult, there are a lot of other things that would need to be done to make a slick system. Examples include how you share the screen with OpenLook apps, sharing clipboard formats, whether you can do DDE with OpenLook apps etc. This has much of the flavor of the problems that arise in trying to run Windows and PM at the same time in a nice fashion, but the details are worse because OpenLook differs a lot more from Windows than PM does. It is not so simple to add
- Sun will continue to evolve OpenLook. They will add new features in an attempt to be competitive in general, but this will cause a direct challenge to Windows. Unless there is a high degree of cooperation between us and Sun this will make life difficult.
- Windows will evolve in directions that are hard to support. Our advanced data storage initiative is a good example of something that will be very difficult to implement on top of SunOS, and there are likely to be more of these in the future.
- We would be vulnerable to direct OpenLook ports, or SMK approaches. The primary value of the Windows layer is to make it easy for an ISV to address both SPARC and Windows on x86 with the same source base. To the extent that this is not smooth because of the evolution issues or clumsiness issues above, or that the resulting app looks like a second class citizen in an OpenLook world, ISVs would be motivated to go directly to OpenLook. If SPARC remains a minor phenomena, then they would use our layer, but once it is important they will look hard at other approaches. The most viable alternative is a software migration kit which Sun or third parties could provide, which makes it easier to port from Windows, but yields a more OpenLook-ish app as the final result.

This does not mean that a Windows layer for SunOS is a terrible idea, but it does raise a lot of questions. It is not a simple project and we would have to overcome these issues in order to make it viable.

2.2. NT Windows for SPARC

Another obvious approach is to build essentially the same NT Windows product that we are presently targeting for a MIPS based Power PC, but offer it on SPARC instead as an alternative to SunOS for the native operating system. This raises its own set of challenges:

HIGHLY
CONFIDENTIAL

- Sun's strategy is based on controlling the operating system. The operating system is key to any plan to dominate an instruction set - they need control of the operating system software in order to be able to get an advantage over the SPARC clones. This comes up in a huge variety of areas - multiprocessor machines, new system level instruction set changes, moving to a 64 bit address space, making handheld machines, adding multimedia capabilities - the list goes on and on. The operating system sits at the critical crossroads between the hardware and applications. To first approximation, the only software that is really visible to the hardware, or that is visible to ISVs - is the operating system. Sun understands this, and they are not very likely to give up their system software business, or to let us get any kind of serious position in it.
- We are not well suited for their present market. Although we are working on making a compelling state of the art product, NT Windows is not particularly well suited for the traditional workstation market. Given that Sun, and Compaq, are both working the workstation market first and foremost in the initial phase, we are bound to get a fairly small penetration. Our product will shine when it is marketed as a Power PC - not as a weird kind option in the workstation catalog.
- We would have to plug and play with their network strategy. The most direct example of us not being well suited for their present market is that we would need to work on a Sun network. Essentially all of their sales are machines on a network, and our situation would be hopeless unless we could have individual machines running NT plug into a Sun network. This means we would have to support NFS and their entire net strategy (directory, mail, security...). Although it is possible we could do this by licensing software from Sun, it would still be a lot of work, and it would constrain us in doing our own networking vision.
- They control the customer via a direct sales force. It is very difficult for us to come in with a different operating system when Sun has dedicated people on site, and is selling a complete solution rather than a retail machine. This does not give our system any room to grow and build up momentum. This will change at some point when they go retail, but by then it will be too late.

One of the themes that runs through the problems above is that there is already a strong operating system strategy for SPARC, namely SunOS. Furthermore, the marketing environment for SPARC is dominated by SunOS and matched to its natural constituency and feature set. Compaq is not considering doing a RISC PC or Power PC which just happens to have a SPARC CPU inside - they are planning on going into the Sun clone business and going after the same customers. We are ill suited to competing in that environment, and without some room to grow, we would never get critical mass and succeed.

It is interesting to contrast this with the MIPS based Power PC plan that we have had to date. In this case we are being marketed as a high end Windows machine, and the feature set would be attuned to that need. The Power PC is designed to mesh well with a network of Windows machines - it has the same apps (upon recompilation), the same look and feel, the right network support (LanMan and Novell) etc.

HIGHLY
CONFIDENTIAL

2.3. NT as a base for SunOS

Another idea is to try use NT as the base for building an entire UNIX version, probably by turning SunOS into a subsystem. This is similar to the Mach approach, and from a technical standpoint the NT kernel would be great for that. It is difficult to imagine that Sun would let us do this, or would be interested themselves. They have the capability to do this themselves, and there is little reason to let us into the revenue stream for this reason.

2.4. Applications on SPARC

Finally, our applications group is a way that we could profit on SPARC even if we do not get any systems revenue. The issues here are fairly straightforward - we would simply port our apps to OpenLook. One could imagine using some portable Windows code as an internal porting aid, but that is just an implementation detail.

This case differs substantially from the historical example of our success on the Macintosh:

- Apple never had the potential to kill our systems business. The Mac was destined to be popular, but because it was proprietary, it was a self limiting and could not threaten our systems business. The Mac offered incremental revenue with no strategic consequences. Lending support to SPARC is a different matter entirely.
- SPARC does not offer a unique technology. The Mac gave us a unique chance to do GUI, and thus be the foundation for a long term applications strategy. SPARC has no such offer - RISC is not unique (indeed MIPS has better performance), and the rest of the system definition is boring and not substantially different than today's PCs. Power PC on the other hand does represents an opportunity to raise the bar on the minimum system and take advantage of technological synergy.
- Apple had a much better distribution strategy. The current Sun direct sales force approach is not conducive to selling our applications. We really need to have retail software distribution. SPARC will have that eventually, but until that point it makes the business case of doing SPARC applications much tougher. Although there is a big advantage in being first on a new platform, there is also the phenomena of being "all dressed up with nowhere to go" - it doesn't count if you are so early that you are gaited by having an immature infrastructure such as the distribution channel.
- The present SPARC market is niche oriented. There is an interesting chicken and egg problem - SPARC has a very poor selection of mainstream office productivity applications today - which means that its present set of users clearly do not place these in very high regard. SPARC sells today to people in niche markets (electronic CAD etc) - although they may have an interest in word processing, spreadsheets etc, the per capita demand will be much less than the PC market. This reduces the effective installed base and sales volume that we can look at to project applications revenues.
- Our opportunity cost is higher at present. Our apps group has a historic opportunity to reap the advantages of having bet on GUI, and on Windows.

HIGHLY
CONFIDENTIAL

- SPARC app sales will impact sales on other platforms. In the early days of the Macintosh we had a very small market share in applications, so the Mac represented good incremental revenue. Someone who bought a Mac and our software was almost certainly a customer that we would not be able to reach with other products. This is no longer true, and many business customers that buy SPARC and our apps would be likely to have bought our apps on Macintosh or under Windows. In encouraging the growth of the SPARC market we have to recognize that the customers we attract will not all be incremental additions.

Many of these points apply only to the early stages of SPARC. Obviously if SPARC is destined to succeed and be the high volume office platform, then we should do applications for it. The points above mainly speak to the issue of when we should jump on the SPARC bandwagon. The net result is that the business case for doing SPARC apps is much tougher than it would appear at first because the SPARC market is less lucrative than it appears (niche market, impact on other platforms) and because there is a tradeoff against our systems business.

This would suggest that we should not do SPARC applications until the point where we think that SPARC's success is a foregone conclusion, *and* the SPARC infrastructure and channel is in place to make our entry meaningful.

2.5. Conclusions

The fundamental problem that we face in looking at SPARC as an opportunity is that they don't need us. The very reason we are discussing SPARC is our fear that Compaq's support lets them achieve critical mass. This naturally limits Sun's interest in doing a special favor for us, yet without some kind of edge or unique advantage it is tough to compete with their own system software which has many built in advantages.

This is another way of saying that Microsoft is not much like Compaq - it is difficult to work up enthusiasm for buying into somebody else's game and their rules and still beating them through superior execution. Compaq only has experience in succeeding at this strategy against IBM, which is a very slow moving company that does not know how to execute all that well, and does not even understand how to press their advantage. With a couple of minor changes in strategy, IBM could have eliminated Compaq's big claims to fame - for example if IBM had wanted to do the 386 first, Intel would have put the fix in for them and the Deskpro 386 never would have been. With the right licensing terms and up front negotiations, IBM could have had Compaq and the rest of the industry locked into the MCA bus and there wouldn't have been an EISA. IBM is not ruthless, innovative or even all that ambitious, and Compaq may discover that they need new tactics against an opponent like Sun which is all of these.

Compaq's view of Sun as a company that can be easily out maneuvered in their home court, on a game that they invented is not one that I share. That goes for their system software as well as for their systems - either one is a very tough nut to crack.

If in fact we have to do this, the key will come through putting enough pressure on them that they need us, getting some initial breaks in this way, and then pressing our advantage with flawless execution. That is utterly different from what we are used to doing.

HIGHLY
CONFIDENTIAL

3. FOSTERING SPARC ALTERNATIVES

Julian Schwinger, a Nobel prize winning physicist, was often at odds with the rest of the physics community. One of his books starts with the quote

If you can't join 'em, beat 'em!

This inversion of the usual homily is appropriate here - it behooves us to consider how we can beat, or at least impede SPARC, because joining the SPARC movement is going to be very difficult.

Note that this is a good idea even if we think that it is likely that in the long term we will join. There is no reasonable scenario in which it benefits us to sit back and let SPARC momentum continue. Anything which reduces SPARC's power and momentum is positive for us - it is either the opening that we need to compete with them, or it provides the leverage we need to negotiate a graceful entrance into the SPARC business.

Although the specter of Compaq throwing in with Sun is daunting, we should not forget that we have been dealt some very good cards:

- **Windows is emerging as the key API for PC ISVs.** This gives us the ability, if we are careful about it, to deliver these ISVs, and the attendant momentum of their support to Power PC.
- **We have a great deal of influence.** Much of the computer industry looks to us for guidance. They do not always like this and they can be resentful, but this doesn't mean that they won't do what we say.
- **SPARC has powerful enemies.** Many of the world's largest computer companies are committed to fighting SPARC, or die trying. This is a potent resource which we can tap and channel. The one thing that the anti-SPARC forces of the world lack today is leadership and a shared mission. Encouraging them to them is a high leverage role for us to play.
- **Compaq will not compete directly with Power PC.** They are sincere about avoiding their present channel and overlapping with their present market. Their support of SPARC is a strong endorsement to the industry, but they will not have anything which directly confronts Power PC as far as end users are concerned.
- **SPARC is not ready to compete in the PC industry.** Although they have their act together in the workstation arena, they do not have sufficient ISV support to mount a credible launch in the PC industry. Unless we mess up in a big way they will not attain this in the time between now and the shipment of Power PC for MIPS. As long as we can get a reasonable number of Windows applications to port, we can easily dominate them in a direct showdown.

If we play our cards right, we can parlay these advantages into a pretty complete victory. This will not be easy by any stretch - but it is possible.

3.1. Uniting the MIPS Community

The only RISC chip that has a hope of beating (or even slowing down) SPARC is MIPS, so the clear thing to do is to strengthen the MIPS camp. Our plan of record has basically ignored the MIPS based workstation market, and we have focussed only on a very elite group of PC manufacturers who build the reference platform primarily for NT Windows.

This strategy does nothing to slow down SPARC in the workstation business, and it gaits the availability of the reference platform to OEMs, chip vendors etc on the availability of our software. In a world where SPARC is getting a powerful edge it suggests that we revise this to give more near term benefit to the MIPS world:

- **Cause the MIPS community to unite behind standards.** We would use our influence with key OEMs to get every major Sun opponent to endorse a common MIPS based platform standard and a common UNIX standard. The key elements are an R4000 reference platform system (with associated chips) which we can supply from our own effort, and a single version of UNIX which we should cause to be knighted as the standard (and maybe participate in it business wise). We would approach the standard in such a way that some companies (such as DEC) could have incompatible hardware below the OS level (and thus do a lot of porting work) but there must be application level binary standard.
- **NT Windows is the carrot, and Sun is the stick.** We would tell these companies privately that we will support the MIPS standard with NT Windows and Power PC, and make an appeal against SPARC. We may or may not choose to make our support public early in the game. Our goal is that we get them all to support NT Windows - either as primary or in some cases as the secondary operating system. The key difference between this approach and our current plan in that we do allow/encourage them to go ahead and ship UNIX on our machine design - especially if the hardware is ready before our software.
- **Allow a much broader initial group.** We would target a much larger number of OEMs - at least for the agreement on the reference platform and the UNIX standard. We might restrict the initial availability of NT Windows to give a smaller set of OEMs an initial time advantage on shipping, or we might not.
- **Announce early.** We would make sure that the MIPS unification announcements happened as soon as there were enough signatories. We would also consider pre-announcing NT Windows & Power PC early as well. It is important to break the monotonic stream of good news about SPARC.
- **Get Windows ISVs to support Power PC/4000.** The one key card that we hold is Windows 3 and control of the Windows API. We must use this to ensure that Windows ISVs port their applications to Power PC, and to MIPS.
- **Get Power PC/4000 to ship ASAP.** The UNIX oriented standard is just a placeholder to slow Sun down and to give Compaq something to worry about in the UNIX market they are so eager to join. The real vehicle for blocking Sun is to get the PC industry to ship MIPS based Windows machines. We have to make NT Windows a priority like none before and get it done.

HIGHLY
CONFIDENTIAL

- Promote the hell out of NT Windows on MIPS. We will have to spend a lot of money and effort in promoting Windows on MIPS, even though it is likely to come installed on the hard disk rather than be a retail product. We would also require the OEMs to push very hard. There are a number of creative things that could be done to help establish the machine.

Once there is a strong unified appearance to the MIPS workstation market, Sun's momentum will lessen. For example, the UNIX ISVs are pretty much all platform neutral today in the sense that their apps are available on more than one machine - where the number 1 platform is SPARC and numbers 2 through N are random with no one platform getting a decisive margin. Once the market consolidates into just two mainstream platforms (plus the RS/6000 as a random wild card), it would be crazy not to support both of them. This helps to deny Sun a lock on the applications and slows the growth rate of SPARC unique applications.

This also helps the industry infrastructure gear up to make systems. Having the architecture spec and ASIC designs go out to chip manufacturers early on will bootstrap the process of getting good support chips, having third parties supply add ons etc.

One big plus is that achieving unification of the MIPS world is relatively cheap in terms of the commitment that we need to make. The primary activity would be flying around the world convincing people. All that Microsoft would really need to commit is that we would in fact make a version of NT Windows for MIPS and offer it for sale - along the lines of what we have to commit in the MIPS contract when we exercise the option. I do not think that we will need to commit to doing this exclusively by any means, so if we later need to try SPARC we will be able to do so. The reason we can get away with this is that the biggest companies, who are most likely to try extract such a promise, already hate SPARC and already want to compete in the UNIX market so much that they will pursue this course independently of whether we ever ship NT Windows.

The companies which are likely to join up include:

Company	Market Area	MIPS already?	Probability
Olivetti	Europe	No	90%
Siemans/Nixdorf	Europe	Yes	90%
Bull/Zenith	Europe & US	Yes	90%
Apricot/Mitsubishi	Europe	No	70%
Nokia	Europe	No	70%
NEC	Japan	Yes	90%
Sony	Japan	Yes	50%

HIGHLY
CONFIDENTIAL

Acer	Taiwan & US	No	70%
Daewoo	Korea & US	Yes	90%
DEC	US & World	Yes	80% ??
HP	US & World	No	50% (90% long term)
MIPS	US	Yes!	100%
Misc workstation (Silicon Graphics...)	World	Mixed Yes & No	90%
Misc mini & mf (Amdahl, Tandem...)	World	Mainly Yes	90% (will support but few products)
Misc 2nd & 3rd tier PC OEMs (AST, Dell...)	World	No	50% - 90%

This list omits the more questionable companies, although we might want to give them a try. It also omits companies that would be very nice to have, but whose likelihood of joining is unknown. This category includes companies such as NCR (we should be able to get their Tower & mini division even if the PC side does not do Power PC right away) and Tandy (not normally high end enough for Power PC, but might lend support). AT&T is another random case - they might be worth bringing on and might not.

Europe is clearly the best geographical area - we should get a clean sweep of the major companies, because they are either already signed up for MIPS (Siemens, Bull) or could easily be influenced by us (Olivetti). Japan is also strong if we can get NEC, because they dominate the market so much. The US is actually the worst area for large name brand companies. Zenith comes along with Bull, which helps a lot.

HP and DEC are the most interesting since they have the biggest reputations and would help the image. They could be hard to convince. Each of them would join in an instant if they knew Compaq was going SPARC, but might drag their heels otherwise. One interesting thing about both of them is that their PC businesses are not doing as well as their workstation & mini businesses. We could position Power PC to them as a way to get synergy between their workstations and PCs to help fix this problem.

Note that this is the list for the general support of the strategy including Power PC, the reference platform, and the UNIX standard. We would select a subset of these companies for the Power PC consortium - which omits the companies from outside the PC industry like Tandem or Amdahl. Their support is important image wise for confronting the SPARC armada.

HIGHLY
CONFIDENTIAL

3.2. Managing PC ISVs

The biggest single asset that we can bring to bear on the problem is our control of the Windows API, and therefore the PC ISVs. This will require careful management, since they basically hate us. On the other hand, they love to hate us - despite the bitching, they do make money and are not stupid enough to hurt their own businesses just to spite us.

The general idea for how we can manage the ISVs onto RISC is to work as hard as humanly possible at getting 32 bit Windows to be a reality on the x86, and make the recompile over to RISC seem like a compelling incremental investment. The basic plan goes as follows:

- Continue to encourage the rush to Windows. This hardly needs to be repeated here, but it is essential to do whatever we can to keep ISVs moving to Windows. Making clear statements about the future of Windows will help.
- Head off platform neutrality. We do NOT want the next priority after an initial Win 3 app to be a move toward platform independence. We obviously cannot stop people from being sensible about organizing their source code, but it is essential that we keep the Windows API moving toward our goal, and give ISVs something meaningful to do that helps our strategy.
- Promote the Win 32 API. The first race to pick is taking advantage of the 386 fully with 32 bits - at first via the thunk toolkit, and later with a fully 32 bit system. The large increase in 386 sales due to Windows means that there will be a substantial installed base of Win 3 on 386 so this is a natural and compelling thing for ISVs to do. We should promote this with big seminars like the original OS/2 seminars - held worldwide etc. Depending on timing this might be x86 only, or we might want to promote our RISC strategy as well.
- Provide good tools for 32 bits and for RISC. We have to make this transition as smooth as possible. This is especially true of the move between 32 bit x86 and RISC - it must be very easy to do. In the case of moving from 16 bits to 32 bits on x86 we will have a lot of pull from the 386 installed base to help out, but we still need good tools. Our tool strategy should include evangelizing third party tool vendors as well as doing internal work.
- Evangelize Power PC. The dual nature of the standard makes this a particularly good way to get ISVs involved with RISC. Many horizontal applications will not have to do a great deal of work to support Power PC in an opportunistic way (i.e. they'll be nicer when on it, but not have a dedicated version), but it is still valuable to promote it. We need to get a few exciting high profile apps to show the way.
- Come up with ways to encourage ISVs to support RISC. This includes financial support and a variety of other inducements. Once we have gotten them to do a Win 32 app and made the port to RISC reasonably easy to do, it is only a question of the incremental investment that is necessary to make the release. The OEMs can help position this as an industry wide phenomena rather than something done at the bequest of Microsoft.

Every step is quite solid and independent of the Compaq SPARC issue until you get to the last point. It should be relatively straightforward to get ISVs to support the Win 32 API in one way or

HIGHLY
CONFIDENTIAL

another no matter what Compaq does. The tricky point comes when we get the ISV to actually commit to releasing a version for our RISC version - that is the point where Compaq's support, or lack thereof, will make some of them balk, or at the very least take a wait and see attitude. The way to overcome that is primarily through good marketing to the ISVs and good promotion. We must also make sure that there is a solid perception that the machine has sufficient support from OEMs to be successful.

3.3. Wild Ideas

Here are some sample ideas about dramatic (or just crazy) things that could be done to enhance the plan discussed above. Be forewarned that they are not full proposals at this point - just

- **Distribute apps (and/or working models) pre-installed on the hard disk.** This is one of the ideas that has been discussed already for Power PC. This helps make the case to ISVs because they get very high visibility if their app is pre-installed. Ideally this is done via a machine serial number scheme so that you can buy the app by phoning an 800 number with a credit card and get a key. We can use our own applications to help force the issue - for example we could say that two apps in each category will be included (to be fair), and bootstrap the process with Excel and Win Word. This would put the pressure on Wordperfect and Lotus to either join up or see us get a real leg up.
- **Convince someone to bet the farm.** There are no end of companies, especially in Japan, that seem to want to spend insane amounts of money to break into the computer business. Some recent examples are Kubota (with Ardent), Matsushita (with Solbourne) and Canon (with NeXT) - each has sunk around \$100M so far, and don't have much to show for it. We could consider telling another such company that Power PC is enough of a paradigm shift that it is an excellent opportunity for them, especially if Compaq and IBM are sitting this round out. One obvious candidate is NEC - we could get them fired up about using this as their big entree into the international PC market. The total investment is not necessarily as large as the ones mentioned above - the key is to get them to be very active in promoting the machine.
- **Attract software (or other) vendors which have a direct sales force or special distribution network.** A good example here is Oracle (but Novell may also apply). If they had a very good position on RISC PC, they could be very effective at helping to establish the machine. This does not mean selling the hardware, but there are a variety of ways they could make it appear like a well supported mainstream choice to their customers. This ranges from just having their sales force push it, to offering special service and support services for configurations including the machine.
- **Create a large marketing war chest.** This can be funded through contributions from our OEMs etc.

HIGHLY
CONFIDENTIAL

- **Make a great cross development system for x86.** One way to get software developers interested in the machine is to make it a very nice, fast environment for their own use. If it is a great platform for developing Windows apps for any machine, and there are cross development features so that you can produce x86 binaries, then you will get a lot of ISVs buying them (with a special discount from the OEMs) for their own purposes. There is nothing like having a fast machine in front of developers to get a lot of midnight projects going.

3.4. Conclusions

There is no silver bullet which we can use to stop SPARC in its tracks. This was true before the possibility of Compaq going with SPARC came up, and it is even more true afterwards. The ideas discussed in this section give us a very good chance of slowing SPARC down, and with good execution we have a solid chance of beating them. "Beat" in this context means to establish RISC in the Windows community in such a way as to prevent SPARC from gaining a foothold in the PC industry and retail channel.

4. RECOMMENDATIONS

The first, and obvious thing to do is to work as hard as possible to try to change Compaq's mind and not do a SPARC machine. In parallel we should:

1. **Redouble our efforts to rapidly define and implement NT Windows.** This includes both the portable Windows and kernel pieces. This is key technology in any all scenarios, and the quicker we have it the better.
2. **Do not start work on any SPARC based software.** The time required to port NT is very small, as has been demonstrated with the 386 port. There is no reason to confuse the development team or waste any resource at this time - we could port the system very rapidly at any point if it is required. The situation with our Apps group also does not require any new action - any spare bandwidth that they have should be directed toward moving to 32 bits, which will be required for SPARC as well as for x86 and MIPS.
3. **Refine a plan for uniting the MIPS camp.** This mainly involves figuring out what position we want with respect to UNIX with MIPS and SCO, and then taking the show on the road. Actual implementation of the plan should begin within the next couple of weeks once we have a chance to review it. There appears to be little downside in taking this approach, and it will benefit us.
4. **Define Power PC and work out the strategy in more detail.** This is a key part of any plan and we need to get it very solid very soon.

This list will obviously change as things move along, but it is a good place to start.

HIGHLY
CONFIDENTIAL