Confidential Copy No.:

Issued to:

Network Appliance Corporation

Business Plan

August 1992

Contact: Michael A. Malcolm, President
Network Appliance Corporation
2901 Tasman Drive, Suite 208
Santa Clara, California 95054
(408) 562-1900

Table of Contents

1.	Executive Summary	1
2.	Market Opportunity	4
3.	Product Concept	7
4.	Customer Reaction	11
5.	Distribution Strategy	13
6.	Competition	15
7.	Product Technology and Development	18
8.	Organization and Personnel	21
9.	Management	22
10.	Pro-Forma Financial Statements and Assumptions	25
11.	Risk Factors	31
12.	Funds Required and Their Uses	33
Appendix A:	Proposed LADDIS Benchmark	34
Appendix B:	Price/Performance Competitive Analysis	35
Appendix C:	Cost of Ownership Competitive Analysis	41
Appendix D:	WAFL — Write Anywhere File System Layout	43

1. Executive Summary

Network Appliance Corporation (the "Company") was incorporated as a California Corporation in April 1992 to develop and market file server appliance software.

An *appliance* is a device that performs one particular function. Hence a *file server appliance* is a device that does nothing but provide network file services. Similar to a bridge, router or terminal server, a file server appliance is faster, simpler and more reliable than using a general-purpose workstation to provide such a service.

Recent developments in commodity hardware enable the Company to enter the market with proprietary appliance software, and use personal computer ("PC") hardware to deliver high-performance file services.

The Company is currently developing Network File System ("NFS") Appliance software, a significant portion of which is demonstrable. The Company expects to start beta testing in Q2 of 1993, with first customer shipments in Q3. The Company is seeking $3 million in venture capital to use for developing and launching the product, and establishing distribution channels.

NFS Server Market

NFS has the same characteristics that allowed bridges and routers to succeed: a stable protocol, a large market, and requirements for high reliability and performance.

The NFS file server protocol has not changed since it was introduced by Sun Microsystems in 1985. It is widely available on a large variety of hardware platforms and operating systems. Nearly all UNIX workstations and over a million PCs use NFS. Sales of NFS servers exceeded $1 billion in 1991.

Today's NFS servers fall short of meeting the reliability requirements of many users. This is illustrated by a user's observation: "A distributed system is one in which the failure of a computer you didn't even know existed can render your own computer unusable."

NFS performance can be measured in a vendor-neutral way using the recently developed LADDIS NFS benchmark. This will encourage customers to view NFS servers as commodity products in much the same way that PCs are today viewed as commodity products.

The recent availability of reliable high-performance commodity PC server hardware and the LADDIS NFS benchmark provide the opportunity for a major change in the NFS server market.

Product

The Company's NFS Appliance software will provide network file service over Ethernet, using commodity PC hardware. The product will use a real-time kernel, a proprietary file system, and fault-tolerant disk array software. By using software optimized for NFS, and eliminating the overhead of UNIX, NFS Appliances will provide levels of reliability, performance, and simplicity that are impossible to attain using UNIX.

NFS Appliances will be more reliable, partly because PC hardware is more reliable than high-end UNIX platforms, but primarily because NFS Appliance software won't have extraneous and error prone features such as virtual memory, window systems, and support for user applications.

Eliminating extraneous UNIX features also allows an NFS Appliance to achieve higher performance than is possible for UNIX on the same hardware, because the software can be optimized specifically for NFS. SunSoft and other UNIX vendors expend much more effort on optimizing desktop application features than on the often conflicting task of optimizing network services.

The special-purpose nature of an NFS Appliance eliminates most of the complexity of administering current NFS servers. Other important NFS Appliance features include the following:

- On-line "snapshot backups" eliminate daily tape backups.
- RAID[1] prevents data loss when a disk fails.
- NFS response times are twice as fast as current UNIX servers.
- One expandable partition holds all user files.
- Reboot takes one minute (compared with 20 to 60 minutes for UNIX).
- File compression stores more data on fewer disks.

In the next several years it is unlikely that UNIX-based NFS servers will have any of the features listed above, except RAID.

There are many directions for NFS Appliance software to evolve. Possible enhancements include support for: new high speed network devices, such as Etherswitches, Fiber Distributed Data Interface ("FDDI") and Asynchronous Transfer Mode ("ATM"); additional file server protocols such as Novell's NetWare, Apple's Apple Filing Protocol ("AFP"), and Transarc's Andrew File System ("AFS"); and improved file server administration features.

Customer Demand

Formal customer interviews, on-site visits and informal discussions with UNIX system administrators indicate that customers quickly grasp the benefits of an NFS Appliance, and that the potential demand for the product is very strong.

There will be an opportunity to sell NFS Appliances to nearly every organization that buys UNIX workstations made by Sun, Hewlett-Packard, DEC, IBM, and others. There will also be an opportunity to replace many of the 100,000 existing NFS servers. The NFS server market may grow more

---

[1]"RAID" stands for Redundant Array of Inexpensive Disks.

rapidly than the UNIX workstation market because of the growing trend toward using NFS in PC networks.

Distribution

The Company plans to distribute its products through UNIX Value Added Resellers ("VARs"), system integrators, national chains, and OEMs.

Initial sales will likely be through regional VARs. If necessary, the Company will sell direct to a small number of strategic accounts to generate product awareness.

OEM prospects include UNIX workstation vendors such as Hewlett-Packard, IBM, DEC, Sun, Next, and Sony; and PC system vendors that sell into UNIX markets, such as NCR, Compaq, Dell and Olivetti. Relationships with some of these vendors may involve porting the Company's NFS Appliance software to run on RISC hardware.

The Company plans to introduce the product in Europe and Japan within a year of introducing the product in the United States.

Founders and Management

The Company's three founders have considerable engineering and business experience with operating systems, networking, and NFS servers.

The Company's President and CEO, Dr. Michael Malcolm, previously built a successful management consulting practice helping companies in the industry develop product plans and business strategies. His clients included Auspex Systems, Sun Microsystems, Quantum Corporation, Tandem Computers, and UNIX International. In 1982 he founded Waterloo Microsystems, a Canadian company that developed and marketed a network operating system and related hardware for personal computers. He served as the CEO of Waterloo Microsystems for nearly six years.

The Company's Director of Engineering, Mr. James Lau, previously served as Director of Software Development at Auspex Systems. Before joining Auspex he was the Group Manager of PC Products at Bridge Communications.

The Company's System Architect, Mr. David Hitz previously served as a systems programmer at Auspex Systems Inc. Before joining Auspex he was a systems programmer at MIPS Computer Systems.

The Company plans to add a senior marketing executive to its management team later in 1992, and a Director of Sales and a controller in Q1 of 1993.

Seed Financing

In April 1992, the Company raised $125,000 in seed financing from five private investors: Owen Brown Enterprises, Ltd., Glen McLaughlin, C. Gordon Bell, James Katzman, and Theodore G. Johnson.

2. Market Opportunity

Most networks of workstations and PCs contain one or more machines dedicated to providing file services. These *file servers* store all types of data used by other machines in the network: programs, documents, graphics, images, digitized audio and video, electronic mail messages, etc.

Ten years ago, file servers were found only in research laboratories and experimental networks. In the past decade, a worldwide market for file servers emerged, and has grown to $4 billion per year. Today, nearly every networked personal computer and workstation uses the services of at least one file server. A file server plays such an important role in a network that its performance and reliability are critical to many applications.

2.1 The NFS Server Market

Sun Microsystems introduced the NFS protocol in 1985. Today, nearly all UNIX workstations and over a million personal computers use NFS. More than 100 companies have developed NFS implementations on a variety of hardware platforms and operating systems.

Dataquest estimates that in 1991, the world file server market for workstation networks totaled 47,900 units. Nearly all of these were NFS servers. Dataquest estimates the factory revenues for these servers totaled $1.4 billion in 1991, and unit shipments for the four largest vendors were as follows:

Vendor	1991 Units	Change Over 1990
Sun Microsystems	16,283	78%
Hewlett Packard	11,613	72%
Digital Equipment	7,086	148%
IBM	4,364	48%

NFS is the fastest growing file-server market segment. Dataquest estimates that the unit cumulative annual growth rate of the workstation file server market will be over 26% for the period 1992 through 1996. It is possible that the market for NFS servers will grow even faster because a new market for NFS servers is emerging in PC LANs. There are several reasons why an increasing number of PC LANs use NFS:

- Novell is starting to populate their installed base of 18 million PCs with NFS client software. Novell adopted Sun's TI-RPC protocol[2] for distributed applications. Both Novell and UNIVEL (the partnership between Novell and UNIX System Laboratories) are producing NFS products.

---

[2]NFS is based on Sun's Transport-Independent Remote Procedure Call ("TI-RPC").

- Because NFS uses the Internet Protocol ("IP"), NFS is better suited for large networks (with routers) than is NetWare which uses Novell's proprietary IPX protocol.
- Because it is available on essentially every popular hardware and operating system platform, NFS is becoming the standard protocol for network file sharing.
- Several companies have recently introduced NFS products for PCs.

2.2 Problems With Existing NFS Servers

The administration of an NFS server is complicated. Reliability is poor. Anecdotal evidence suggests that a typical NFS server has a mean time between outage[3] of under 30 days. Without a Legato Prestoserve card, most NFS servers cannot provide enough throughput performance to use the available bandwidth of even one Ethernet. None of the current NFS servers provide response-time performance within a factor of two of what is possible using Ethernet.

The root of these problems lies in the server's use of a general-purpose operating system: UNIX.

While NFS servers have been built on such diverse operating systems as Microsoft's MS-DOS, IBM's MVS, Digital's VMS, and Tandem's Guardian, the vast majority of NFS servers are UNIX-based. Even products designed to be NFS servers—notably those of Auspex and Epoch—run UNIX and third-party application programs.

UNIX requires a long instruction path for processing each NFS operation. Blocks of data are copied from one memory location to another. Substantial processing overheads are associated with virtual memory, context switching, and the protection modes required to run application programs.

UNIX contains approximately five million lines of source code. It has many defects that cause systems to fail. Occasionally an application program invokes a seldom-used sequence of code that results in a system failure. Users have found that their NFS servers are more reliable if they don't allow any user applications to run on them. (Also, applications perform poorly on an active server.)

Both Sun Microsystems and Auspex Systems have built proprietary multiple processor hardware systems to achieve better NFS throughput performance. The prices of these systems range from $80,000 to over $500,000. These are large machines, with complexities comparable to modern mainframe computers. Installation and maintenance are expensive. Also, the large number of components in these machines makes them less reliable than a good PC.

NFS server administration is complex because UNIX system administration is complex. In keeping with UNIX tradition, administering

---

[3]An *outage* is defined as a failure of software or hardware.

an NFS server requires considerable understanding of the UNIX implementation. Many books, one journal, and numerous university and commercial courses are devoted to UNIX system administration. One of these books focuses exclusively on administering NFS under UNIX.

2.3 The NFS Appliance Opportunity

Recent technology developments provide the opportunity to develop a new type of NFS server—one that uses commodity hardware, and is not based on UNIX or any other general-purpose operating system. Properly designed software based on a real-time kernel could provide NFS service with high reliability, high performance, and minimal system administration. We call a server based on such software an *NFS Appliance*.

This major market opportunity is similar to those seized by companies like Cisco and Xyplex to introduce special-purpose routers and terminal servers that perform functions previously performed (poorly) by UNIX systems. An important difference is that there is no need for the Company to develop proprietary hardware because NFS Appliance software will run on PC hardware.

2.4 Enabling Technologies

NFS Appliances, as described in this business plan, are just now becoming practical as a result of several recent hardware developments:

- The Intel processors used in high-end PCs have only recently approached the speeds of workstation RISC processors. Forthcoming 586 and 686 processors are expected to keep pace with, if not surpass, RISC performance.
- Fast EISA bus Ethernet interfaces, capable of driving Ethernet at nearly full bandwidth, became available earlier this year.
- The PC EISA bus bandwidth is comparable to that of a high-end UNIX server. This allows many Ethernets and disks to operate in parallel.
- Gigabyte disk drives are now available in a 3.5-inch form factor. This permits small server form factors, and lower power requirements.
- The lower power requirements of PC hardware and 3.5-inch disk drives make it possible to use a commodity uninterruptable power supply ("UPS"). A UPS is essential to provide both high performance and data safety.

Before 1992 it would have been difficult to convince customers that an NFS Appliance delivers superior performance. The LADDIS benchmark for NFS servers (which was submitted to the Systems Performance Evaluation Cooperative earlier this year) provides a vendor-neutral way to objectively measure the performance of NFS servers. (See Appendix A.)

3. Product Concept

The Company will design and develop NFS Appliance software to provide NFS service on commodity hardware. The primary distinction between an NFS Appliance and a UNIX-based NFS server is that the appliance provides file service only. It cannot run applications or user programs. Users cannot even log in to an appliance. Because it performs only one function, an NFS Appliance is more reliable, faster, easier to administer, and less expensive than a UNIX server.

The term "appliance" implies a design philosophy. An appliance should be easy to install, and once installed it should require little attention. Where possible, appliance features should work even if the system administrator doesn't know about them. Any feature that is complicated to use or difficult to document should be simplified or eliminated. Ideally, one should be able to use an appliance even without reading its manual.

To client workstations, NFS operations performed by an appliance are indistinguishable from those performed by a UNIX server, except for their speed.

3.1 Hardware

Typical PC configurations to run NFS Appliance software will include:

- EISA motherboard (32-64 MB RAM)
- Up to 4 Ethernet interfaces
- 1 or 2 SCSI channels (up to 7 disks per channel)
- Data compression card (optional)
- Uninterruptable power supply (UPS)

At the time of first customer shipment, the Company's NFS Appliance software will support only a few specific types of Ethernet and SCSI controllers, but future releases will support a broader range of systems. Low-end systems might have one Ethernet and one SCSI channel. High-end systems could have up to 16 Ethernets and 10 SCSI channels using currently available hardware. The data compression card is optional; when it is present, the NFS Appliance makes more efficient use of disk storage space.

3.2 Key Features

The primary function of an NFS Appliance is to respond to the 15 network requests defined by the NFS protocol. Beyond that, the best features are *passive* in the sense that they improve the appliance without the system administrator having to understand them, or even know about them. The following features are all passive:

- High Reliability

  The PC hardware used in an NFS Appliance will be reliable, largely because quality vendors build PCs in huge production volumes. The

NFS Appliance software will be reliable because it is small and special purpose. It doesn't support virtual memory, window systems, or user applications, all of which can crash a UNIX-based NFS server, and none of which are required to provide NFS service.

- RAID (redundant array of inexpensive disks)

  RAID allows an NFS Appliance to remain on-line with all user files accessible even after a disk failure. When a failed disk is replaced, the NFS Appliance completely reconstructs its data.

- High Performance

  The NFS Appliance will have high performance because it is a special purpose product optimized specifically for NFS, and because the Company has designed the software to take advantage of recent advances in technology. UNIX is too general purpose to optimize for one task and too stable to redesign for new technology.

- One Minute Boot

  While an NFS server is down, client workstations that depend on it freeze until it reboots. Workstations using a typical UNIX-based server that takes 20 to 60 minutes to boot may be useless during that whole period. With an NFS Appliance that boots in one minute, users may not even notice a reboot.

- File Compression

  Files are automatically compressed when written, and transparently decompressed when read. File compression will typically allow an NFS Appliance to store 2 to 3 times more data on a given number of disk drives than UNIX-based servers.

Some features are *interactive* in that the system administrator must use them to benefit from them. In each case, however, the NFS Appliance feature is simpler than the corresponding UNIX feature and requires the administrator to make fewer decisions.

- Snapshot Backups

  A snapshot is a read-only copy of the file system. It is created in just a few seconds without consuming additional disk space. (See Appendix D.) By default, the NFS Appliance creates a snapshot every night at midnight and keeps it on-line for a week. This allows users to recover accidentally deleted files and allows the system administrator to make reliable tape backups during normal operating hours without taking the server off-line. UNIX does not support snapshots.

- Single Expandable Partition

  The NFS Appliance keeps all files in one large partition that automatically grows as new disks are added. A large UNIX-based server often has dozens of partitions. Having just one partition eliminates decisions such as which partition should hold a new user's files, and how many partitions to create on a new disk.

- Easy Installation

  The NFS Appliance eliminates most of the decisions required to install a UNIX server. A system administrator must make one simple decision and type about 50 characters to install an NFS Appliance. By contrast, installing a UNIX file server typically takes many hours.

3.3 Backup

Backing up user files onto tape is the single most time consuming administrative task on a UNIX file server. Most administrators do incremental backups every night and full backups every weekend. Products such as smart backup programs, large capacity tapes, and tape stackers simplify backup management, but the nightly process of moving files from the server to tape remains unchanged.

The NFS Appliance doesn't simplify this process; it *eliminates* it. To understand how, consider a typical backup schedule for a UNIX server:

- Full monthly backups (stored off-site)
- Full weekly backups
- Incremental nightly backups

Full monthly backups are stored off-site to protect against catastrophes that might destroy a file server along with its backup tapes. Full weekly backups are done to recover files lost due to disk failure. Incremental nightly backups allow more up-to-date recovery after disk failure and protect users who accidentally delete their files.

RAID eliminates the need for weekly backups because it protects data from disk failure. Snapshots eliminate the need for nightly backups because they keep recent copies of users' files on-line. Finally, snapshots allow monthly backups to be scheduled during normal operating hours instead of at night or on weekends when system administrators would rather be sleeping.

3.4 Future Product Directions

The NFS Appliance will evolve in several directions. In some areas, such as supporting new network hardware and new file server protocols, it will be appropriate for NFS Appliances to follow the market trends. In other areas, such as providing new techniques for managing file servers, NFS Appliances will lead the market.

The Company plans to add administration features to future versions of the product, such as disk quotas, to limit the number of files one user can have, and disk reservations, to reserve space for a particular user. These will be important to some customers. Support for emerging network backup protocols may also be needed. A Simple Network Management Protocol ("SNMP") interface will appeal to sites planning to use SNMP to manage their network.

NFS Appliances will support new high speed network interfaces, such as Fiber Distributed Data Interface ("FDDI") and Asynchronous Transfer Mode ("ATM"), as they gain market acceptance.

Support for additional file server protocols such as Novell's NetWare, Apple's Apple Filing Protocol ("AFP"), and Transarc's Andrew File System ("AFS") are also possible:

- AFS seems unlikely to become very successful in the marketplace because it has a small installed base and is still evolving. The Open Systems Foundation's current pricing strategy will make it difficult for high-volume vendors to distribute AFS. Also, Sun Microsystems may produce a compatible upgrade of NFS with similar features, and thereby prevent AFS from gaining a significant market share.
- Support for the NetWare file server protocol would make the appliance more attractive to PC networks, especially networks with both PCs and UNIX workstations. On the other hand, the Novell market is very competitive and already uses commodity PC hardware.
- AFP may provide a very attractive market because it is less competitive than the NetWare market, and there are no AFP servers based on commodity hardware. Apple computers are leading the market to the next generation of multimedia applications that will dramatically increase the requirements for high performance and high capacity file services. Apple customers already expect appliance products, such as printers that plug into the AppleTalk network, and intuitive user interfaces.

The Company's NFS Appliance software may eventually support all of the popular file server protocols.

4. Customer Reaction

4.1 Customer Visits and Interviews

The Company believes the primary decision maker in the NFS Appliance market will be the UNIX system administrator. System administrators usually do not have purchasing authority, but their managers understand that even excellent products cannot succeed if the administrator won't support them, and that poor products can work well if the administrator will support them. As a result, system administrators must justify buying decisions to their managers, but managers are inclined to accept the administrators' recommendations.

The Company has performed over a dozen telephone interviews with UNIX system administrators using a formal questionnaire, and made several visits to potential customer sites. In addition, the founders and some of their friends have used and administered NFS networks. These interviews, visits, and informal conversations indicate that administrators are quick to understand the benefits of an NFS Appliance and eager to try one.

Many of those interviewed offered specific comments about the NFS Appliance:

It would fit very well in our environment. — SynOptics

Can we be a beta test site? ... The probability that we would buy them is greater than 90%. — University of Colorado

We would buy 10 if they were available today. — British Telecom

It does everything Auspex does at a much lower price.... Definitely a good idea. — Convex Computers

I see a definite future for it. — Intel

With good administration tools, it would sell like hotcakes — systems integration consultant

The formal questionnaires indicate that potential customers rank the importance the NFS Appliance benefits as follows:

1. high reliability
2. high performance
3. easy administration
4. low price

4.2 Customer Psychology

Just like any customer, system administrators often make buying decisions for unconcious emotional reasons, and then justify their decisions after-the-fact with "logical" reasons.

Emotional reasons to buy an NFS Appliance include:

- The administrator need not perform daily and weekly backups at night and on weekends.
- The administrator can watch an appliance boot in one minute instead of returning 20 to 60 minutes later to check if the boot succeeded.
- Users can retrieve deleted files from on-line snapshots without asking the administrator to find and load a backup tape.
- Installation and administration are quick and easy.
- NFS Appliances are easier for the administrator to control.

The first four reasons contain a common thread: An NFS Appliance will save the administrator's time. The final reason is more involved. Many administrators have discovered that it is easier to manage a UNIX file server if it is not used for general-purpose computing. User programs sometimes crash servers, and users with administrative privileges can create problems. Yet it can be difficult to keep users off a general-purpose UNIX server.

Having made an unconcious decision, the system administrator will use the following "logical" reasons to justify it:

- high reliability and availability
- low entry price and excellent price/performance
- productivity gains from improved performance
- reduced administration costs
- reduced cost of disk space due to data compression
- reduced hardware maintenance costs

The competitive analyses in Appendices B and C will help system administrators justify their decisions.

5. Distribution Strategy

The Company plans to quickly capture a large share of the NFS server market, and establish the Company as the market leader, before competitors have time to react.

The Company plans to establish relationships with a small number of high-profile regional UNIX VARs, mostly on the West Coast. The Company plans to have these VARs help beta test the product in Q2 of 1993 using their beta testing partners. The Company will also attempt to reach distribution agreements with six or more VARs by early 1993.

The Company plans to approach workstation OEMs beginning in 1992, including Hewlett Packard, IBM, DEC, Sun, Next and Sony. Most of these vendors already make PC hardware. Others may require the Company to port its NFS Appliance software to a RISC platform.

The Company plans to approach other PC system vendors that sell into UNIX markets: NCR, Compaq, Dell, Olivetti, and others. Most of these vendors already make PC server hardware and have distribution channels for the NetWare market. These vendors will probably be effective for distribution into the emerging NFS PC market, but they may have difficulty reaching the NFS workstation market. The forthcoming Solaris 2.0 for Intel PCs may help these vendors reach the NFS workstation market.

The Company will also approach major system integrators (e.g., Electronic Data Systems, Anderson Consulting, and Advanced Management Systems), Sun's National Value Added Dealers (e.g., NYNEX, Intelligent Electronics, and MicroAge), and UNIX VARs that specialize in storage-intensive vertical applications.

The Company will attack international markets using a similar distribution strategy. The Company plans to introduce the product in Europe and Japan within one year of entering the United States market.

The Company's NFS Appliance software will require no modifications to enter most international markets because the product has no hardware and no user interface except for administration. English used in the system administrator's user interface will be acceptable because most UNIX system administrators read English. Similarly, the English manual will not be a significant barrier to sales in most markets. Resellers can easily translate the manual to other national languages, where necessary. It will also be relatively easy for the Company to adapt the administration user interface to other national languages if it becomes necessary to do so.

If necessary, the Company will initially sell direct to a small number of key accounts to generate early revenue and create a good customer reference base. The Company could use these direct accounts to pull through prospective VARs and OEMs.

A good regional VAR will introduce the Company to other VARs. A quality VAR channel will make the Company's products more attractive to system integrators, national chains and OEMs. Success on the West Coast will make it possible to penetrate the East Coast. Success in the United States will pave the way to international success.

6. Competition

There are many vendors supplying NFS server products, some of which will compete directly with NFS Appliances.

6.1 Current Competitors

Sun Microsystems. Between one-quarter and one-half of Sun Microsystems' revenues come from the sales of NFS servers. While Sun enjoys an enormous barrier to entry into the UNIX workstation market (the cost of getting hundreds of applications moved to a new platform), the barrier to entry is relatively small in the NFS server market. No third-party applications are required to sell an NFS server, and Sun has failed to invest in upgrading their Open Network Computing software for many years.

Sun is unlikely to produce an NFS Appliance for the following reasons:

- Sun is focused on desktop workstations. Sun is fighting Microsoft, Hewlett-Packard, IBM and DEC, in its attempt to capture a major share of the desktop market.
- Sun is unlikely to produce any system that does not use UNIX. Sun's most influential engineers would probably squash any proposal to build a product without UNIX.
- Sun lacks the engineering skills required to design and build an NFS Appliance. Sun's NFS and networking talent was seriously depleted when several of their top engineers left to form Legato. Since then, Auspex has hired more key engineers from Sun's NFS group, including its manager. Ironically, Sun now lacks the talent to adequately maintain NFS, much less develop new products. Sun also lacks experience with embedded real-time kernels, file systems, and high-performance software.
- Sun's engineering resources are stretched thin. Sun's current transition to UNIX System V Release 4 ("SVR4") will probably introduce significant new performance problems that will take years to correct. Sun's software engineering resources will be largely consumed by SVR4 development for the next two years or more.
- Sun's software release infrastructure would slow any response. Sun is geared for multi-year software release cycles. The last major release of Sun's UNIX was over four years ago, and the features of the forthcoming major release (SVR4) have been frozen for months.
- Sun's current hardware is not well suited for an NFS Appliance. Sun's low-end systems lack the expandability required to support multiple Ethernets and many disk drives. Sun's high-end systems are expensive, and designed to support compute-intensive applications rather than embedded real-time systems.
- The new corporate boundary between Sun Microsystems Computer Company ("SMCC"), which produces Sun's hardware, and SunSoft, which produces Sun's UNIX ("SunOS") and NFS, will make it difficult

for Sun to build a product requiring an innovative combination of new hardware and new software.

Sun will still be a formidable competitor. Sun has strong channel control and brand loyalty. Sun could react to losing server sales by providing deep discounts on server configurations.

The Company will attempt to obtain Sun as an OEM. It may be necessary for the Company's VARs and other OEMs to be successful (i.e., take sales away from Sun) before the Company will be able to obtain Sun as an OEM.

Auspex. While Auspex builds an NFS server with special-purpose hardware and software, it is nothing like an NFS Appliance. The Auspex server runs SunOS, so it is no easier to administer than any other UNIX-based server. It is positioned as a premium, high-end product that only becomes cost competitive when used to replace multiple high-end Sun servers. (See Appendix B.)

Auspex's added value is in its hardware and its excellent customer service. Its proprietary software is used primarily for implementing and managing its unique Functional Multiprocessing architecture.

Auspex's publicly stated corporate direction is to continue focusing on the very high end of the NFS market and to focus on UNIX CPU performance. Auspex has announced a joint development effort with IBM that is likely to consume many of Auspex's engineering resources. The Company believes that Auspex is unlikely to launch a project to develop an NFS Appliance. However, Auspex's market position, distribution channels, and service capabilities may make it attractive as a potential distribution channel for the Company's products.

Others. No other workstation vendor appears to have an NFS server strategy more viable than Sun's. Some workstation vendors, such as Next and Sony, do not provide file servers at all.

6.2 Potential Competitors

The only competitor likely to develop an NFS Appliance is some other startup company. The Company is aware of only one other startup company planning to develop a special-purpose NFS server. The Company believes there is a small risk that this other startup will be successful at obtaining funding.

The Company believes it is unlikely that any other startup company will be the first to market with NFS Appliance software, and even less likely that another company would develop a product with a comparable set of features.

6.3 Barriers to Entry

The Company plans to build substantial barriers to prevent potential competitors from entering the NFS Appliance market. These barriers will include:

- A substantial body of NFS Appliance software. In time it will become impractical for a competitor to make the software investment required to produce a viable NFS Appliance.
- Successful channels for selling NFS appliances. It will become difficult for a competitor to build alternative channels.
- Brand awareness, and a reputation for excellent quality and service.
- Technology patents. The Company plans to apply for patent protection of WAFL technology, particularly the technique used for fast on-line snapshot backups, file compression with minimal increase in response times, and high-performance disk writes with RAID 4.

Being first to market will be the Company's primary barrier to entry.

7. Product Technology and Development

7.1 Implementation Strategy

The Company will develop proprietary software in areas where its efforts provide significant added value. It will use public domain software to accelerate the development process. If possible, it will avoid paying any royalties for third-party software.

The Company's development team programs in ANSI C using SunOS workstations with publicly available GNU compilers and debugging tools. NFS Appliance software is tested and debugged under SunOS, with simulated hardware drivers, before it is ported to a PC target platform for final testing.

Running in the UNIX-process environment makes it possible to use third-party software verification tools, such as Saber C and Purify, that are not normally available to kernel-level developers. The GNU debugging tools support full symbolic debugging for both the SunOS environment and the PC target hardware.

The Company believes that time spent building and using appropriate tools to support development, debugging and regression testing, is abundantly repaid in reduced product development time and improved software quality.

7.2 Major Software Modules

The NFS Appliance software consists of six major modules. This section describes the purpose of each module and, where applicable, the software that it will be based on.

- SK (simple kernel)

  SK is a small, simple kernel that allows very light-weight processes to communicate among themselves and with interrupt-level hardware drivers using synchronized message passing, signals, and UNIX-like sleep()/wakeup() primitives. The Company modeled SK after the Thoth kernel developed at the University of Waterloo, implemented it in C, and simplified it to run much faster.

- Hardware drivers (Ethernet, SCSI, RS-232, UPS)

  Hardware drivers are SK processes that manage hardware such as Ethernet and SCSI controllers. Some drivers will be based on public-domain driver code from Berkeley UNIX. Others will be written from scratch based on the hardware vendors' technical reference manuals. If possible, the Company will use standard driver interfaces (such as NetWare and UNIX) so that existing drivers can be used.

- Network Protocol Stack (TCP/IP)

  The TCP/IP protocol stack is based on the public domain Berkeley UNIX implementation as enhanced by Van Jacobson. This is widely regarded as the best available implementation of TCP/IP and has been used by many UNIX vendors.

- WAFL (write anywhere file system layout)

  WAFL is a proprietary file system designed by the Company to optimize NFS performance using RAID, SCSI-2 disks and a UPS. In addition, WAFL supports features such as a single expandable partition, fast reboot without consistency checks, write caching in main memory with UPS protection, file compression, efficient RAID, and snapshots. WAFL is discussed in more detail in Appendix D.

- RAID (redundant array of independent disks)

  RAID is a technique developed at the UC Berkeley that gangs multiple disk drives together to form a single virtual disk, or "volume." RAID uses redundant data storage so that no data are lost if a disk drive fails.

  There are several different levels of RAID. The NFS Appliance implements RAID Level 4, which uses one disk to store parity. When a disk fails, RAID uses the redundant information on the parity disk to reconstruct lost data on the fly. Thus, all user data remain accessible and the NFS Appliance remains on-line. When the failed disk is replaced, the NFS Appliance completely reconstructs its data. This reconstruction is done on-line, which means that the down time caused by a disk failure is limited to the short period during which the disk is physically replaced. Some hardware vendors will support on-line disk replacement, eliminating down time from disk failures.

- System Administration

  A system administration interface is required to manage the NFS Appliance. Although the appliance is designed to reduce the number of decisions that administrators must make, some options and decisions do require a user interface.

  The NFS Appliance will support a full-screen interface that can run either locally on the system console, remotely over a dial-up modem connected to an RS-232 port, or over the network using the TELNET protocol. The Company will build the interface using two public-domain software products: "curses," which is a full-screen terminal driver, and "rc," which is a command language interpreter. The TELNET protocol implementation will be based on the public domain Berkeley UNIX implementation.

  The user interface will allow the system administrator to control aspects of the NFS Appliance such as how often backup snapshots are taken, how long they are kept on-line, and what file access permissions are granted to client workstations. The interface will also allow the administrator to monitor NFS performance. Finally, the interface will provide security features to prevent unauthorized access.

7.3 Schedule and Status of Product Development

An early prototype of the Company's NFS Appliance software is currently running as a UNIX process with simulated hardware drivers. The prototype fully implements the NFS protocol, responds normally to network clients, passes the NFS conformance test, and is sufficiently robust to store and compile the 25,000 lines of source code for the NFS Appliance software itself. On the other hand, this prototype is incomplete in several important ways: RAID, compression, snapshots, and system administration are not yet implemented; and WAFL only stores files in cache memory without writing them to disk.

The major milestones from the Company's development schedule are as follows:

Date	Milestone
September 92	Incomplete prototype on Intel 486 PC hardware
January 93	Complete alpha units
February 93	Connectathon (multi-vendor NFS compatibility test)
February 93	Early beta units (one or two sites)
April 93	Beta units
July 93	First customer shipments

8. Organization and Personnel

8.1 Staffing and Compensation

During the first year, the Company's operations will be devoted primarily to product development and testing. The Company's staffing strategy for this period will be to hire only those key employees and consultants that are critical to developing and launching the product.

The Company will avoid all unnecessary staffing before and after product launch so that the Company can achieve early profitability and positive cash flow. Following the initial product launch, staffing will grow in each functional area as required to support business growth and product line expansion.

The Company will strive to hire outstanding employees in every functional area. Salary compensation will be on par with comparable companies in the industry that hire equally qualified professionals. The Company will provide every employee with a stock option. Employee stock options will serve to increase performance and decrease future turnover.

8.2 Staffing Plans for Product Development and Launch

The following table summarizes the Company's staffing plans for the next eighteen months. Each entry gives the average head count for the quarter, excluding consultants:

	CY 1992		CY 1993			
	Q3	Q4	Q1	Q2	Q3	Q4
Engineering	4	4	4	4	5	5
Marketing & Sales	0	1	5	6	8	11
Finance & Admin.	1	1	1	3	4	5
Total	5	6	10	13	17	21

9. Management

9.1 Founders

Michael A. Malcolm, President and CEO

Dr. Malcolm previously built a successful management consulting practice assisting companies in the computer industry with product planning, business strategy, and system architecture. His clients included Auspex Systems, Sun Microsystems, Quantum Corporation, Tandem Computers, and UNIX International.

In 1982 he founded Waterloo Microsystems, a Canadian company that developed and marketed a PC network operating system and related hardware and software products. He served as the Chairman and CEO of Waterloo Microsystems until 1987. During this period, the company achieved annual sales of $5 million, established world-wide distribution channels, and obtained IBM as an OEM reseller of its products.

As a professor of computer science at University of Waterloo, he founded and directed one of the most prominent research groups in Canada. This group's work in operating systems and networking is widely recognized, and used by many companies throughout the world.

He has published research papers and a best-selling textbook. He holds U.S. and foreign patents on local area networking, software algorithms and communication protocols. He has given seminars and invited talks at companies, conferences, and universities throughout the United States, Canada and Europe.

He holds a Ph.D. in computer science from Stanford University, and an M.S. in mechanical engineering from University of Denver.

James Lau, Director of Engineering

During the past year, Mr. Lau provided consulting services through H&L Software, of which he was half owner. He consulted in the areas of networking and application software development for pen-based computers.

Mr. Lau previously served as Director of Software Development at Auspex Systems, where he managed the development of the Auspex file server software. He defined product strategy and release plans, and made several key contributions to product requirements, system architecture, and software.

Before becoming Director of Software Development, he managed the design and implementation of the Auspex network subsystem. As part of this project, he designed the hardware and software architecture of the Auspex dual-channel Ethernet Processor.

Before joining Auspex, he was the Group Manager of PC Products at Bridge Communications. At Bridge he managed a team that developed system software products for PCs, including network interface drivers, TCP/IP utilities, and terminal emulators.

As a Senior Software Engineer at Bridge, he was a key contributor in developing the first internetwork bridge and the first network control server.

Before joining Bridge, Mr. Lau held positions at Hewlett Packard and AT&T Bell Laboratories.

Mr. Lau holds an M.S. in Computer Engineering from Stanford University, and BA. degrees in both computer science and applied mathematics from UC Berkeley.

David Hitz, System Architect

During the past year, Mr. Hitz provided consulting services through H&L Software, of which he was half owner. He consulted in the areas of UNIX system software, technical writing, and application software development for pen-based computers.

Mr. Hitz previously served as a member of the software engineering staff at Auspex Systems. He managed Auspex's version 1.2 and 1.3 software releases. He improved the Auspex system performance by 250% prior to the release of version 1.0.

He designed and implemented several significant components of the Auspex software, including the multi-CPU message-passing operating system kernel, the inter-processor file system protocol, and the system administrator's performance monitor. He also developed the systems for source control, software release, and defect tracking.

Before joining Auspex, Mr. Hitz was a systems programmer at MIPS Computer Systems, where he focused on file system and performance issues in the UNIX kernel.

Mr. Hitz holds a B.S.E. in computer science from Princeton University.

9.2 Planned Additions to the Management Team

The Company is conducting a search for a senior marketing executive to add to the management team following the first round of financing. The successful candidate will have outstanding experience in OEM and key account sales, and a good knowledge of the NFS server market and its distribution channels.

The Company plans to add a Director of Sales in Q1 of 1993, and a Director of International Sales in Q1 of 1994. A Controller is planned for March 1993, a senior marketing or sales executive for Q3 of 1994, and a CFO for Q1 of 1995.

9.3 Professional Advisors

Owen Brown, President and CEO, Migration Software Systems, Ltd.

Mr. Brown has held senior executive positions at several companies in the industry, including Xerox, Convergent Technologies, Sun Microsystems, and Digital Equipment Corporation. He has participated in several start-up companies as a private investor and advisor. Mr. Brown led the Company's private seed round investment.

Legal Counsel

Robert V. Gunderson, Jr.
Brobeck, Phleger & Harrison
Two Embarcadero Place
2200 Geng Road
Palo Alto, CA 94303
(415) 496-2900

Accountant

Steven B. Burt, Partner
Coopers & Lybrand
Ten Almaden Blvd.
Suite 1600
San Jose, CA 95113
(408) 534-2263

Bank

David D. Campbell, Executive Vice President
Richard D. Cerkleski, Vice President Corporate Marketing
Cupertino National Bank
20230 Stevens Creek Blvd.
Post Office Box 1350
Cupertino, CA 95015
(408) 996-1144

10. Pro-Forma Financial Statements and Assumptions

10.1 Pro-Forma Balance Sheet Assumptions

1. Cash is the amount remaining after all revenue, expenses, purchase of assets, and financing of these transactions have been accounted for.
2. Payments are received two months after sales are final.
3. Inventory payments are made two months prior to sales.
4. Fixed assets include computer hardware and software, office furniture and fixtures, and test equipment.
5. Fixed assets are purchased with each new hire.
6. All fixed assets are depreciated over three years except furniture and fixtures, which are depreciated over five years.
7. Accounts are paid one month following expenses for: consulting, recruiting, expensed materials, travel and entertainment, telephone, marketing, office supplies, and professional fees.
8. Rent, insurance, taxes and licenses are paid in advance.
9. Accrued expenses are one half month of salaries and commissions.
10. The available bank line of credit is 75% of accounts receivable.
11. There are no other current liabilities.
12. The current portion of long term debt is that part due within 12 months.
13. Long term debt is 80% of all computer hardware, software, and test equipment.
14. Founders purchase $3,000 of common stock.
15. Seed loan to the company of $125,000 is converted to preferred stock at the first round of financing.
16. The first round of financing, totaling $3 million (plus loan conversions) occurs in July 1992.
17. The second round of financing, totaling $2 million occurs in November 1993.

10.2 Pro-Forma Income Statement Assumptions

Revenue

1. Sales include only NFS Appliance software units sold through indirect channels. Sales do not include any hardware, direct sales, future products, or software upgrades.
2. The average list price of each software unit is $5,000 throughout the period.
3. Sales channels include VARs, distributors and OEMs.

4. The average discount off list price is 30% for VARs, 55% for distributors, and 60% for OEMs.
5. At FCS the Company has two active VARs. An additional VAR begins selling each month thereafter.
6. The average VAR requires three months from first meeting until contract signing, and two more months until first sale. The average VAR sells two units the first month, two units the second month, four units the third month, six units the fourth month, and eight units each month thereafter.
7. The two-step indirect channel ("distributors") begins selling one year after the VAR channel. Its ramp rate is slightly faster, resulting in total volume approximately equal to that of VARs by the end of FY 1996.
8. Large OEMs are obtained beginning in 1993. They begin selling in Q2 of FY 1994. They sell a total of 396 units the first fiscal year, 2,008 the second, and 7,169 the third.
9. Interest revenue is 3% of excess cash.

Cost of Sales

1. The product will comprise a diskette, manual, and a small hardware key to prevent piracy. The cost of these items will total $100.

Expenses

1. Salaries are comparable to those of similar companies.
2. Salaries increase 3% the second year, and 7% each year thereafter.
3. Sales commissions are 5% of revenue, to a maximum equal to the sales representatives' aggregate salaries.[4]
4. Taxes and benefits are 20% of salaries and commissions.
5. Consultants and contractors are employed at market rates for special projects.
6. Recruiting fees are one month's salary for each new hire, starting in FY 1993.
7. Travel and entertainment is $2,000 per month for the CEO; $2,500 per month for the VP of Sales, VP of Marketing, and each Director of Sales; $1,500 per month for each sales representative and pre-sales support engineer; $150 per month for each engineer and engineering manager.
8. Telephone is $500 per month for the CEO and each marketing and sales VP, Director and sales representative; $150 per month for all others.

---

[4]This is a model for forecasting sales commission expenses, not a commission policy.

9. Supplies and other expensed materials are $250 per month times total company head count.
10. Monthly rent is $1.50 per square foot times 300 square feet times total company head count at year-end.
11. Annual business insurance is 1% of net fixed assets plus 1% of annual rent. Officer's liability insurance is $2,000 per quarter, starting FY 1993. Key-man insurance is $5,000 per year.
12. Marketing expenses include 5% of sales for advertising and public relations, 2% of sales for attending trade shows, and 1% of sales for direct mail. Some exceptions occur during the initial product roll out in FY 1993 and 1994.
13. Professional fees are industry averages. Legal fees for closing the first round of financing total $15,000; for the second round they total $20,000.
14. The reserve for bad debt is 2% of sales.
15. Organizational costs of $5,000 are amortized over five years.
16. Interest expense is 9% on the bank line of credit and 12% on long term equipment leases.
17. The combined federal and state tax rate is 40%, beginning after all losses have been consumed.

10.3 Pro-Forma Sources and Uses (Cash Flow) Statement Assumptions

1. The cash flow statement is based on transactions shown on the detailed income statement and balance sheet.

10.4 Pro-Forma Balance Sheet FY 1992–96 ($)

Note: FY 1992 ends March 31, 1993.

	1992	1993	1994	1995	1996
ASSETS					
Cash	$2,260,057	$2,034,404	$2,065,966	$4,464,740	$11,088,893
Accounts Receivable	$0	$397,880	$1,655,390	$3,763,369	$7,509,411
Inventory	$0	$17,847	$77,607	$180,718	$318,681
Other Current Assets	$5,000	$10,000	$20,000	$30,000	$40,000
Total Current Assets	$2,265,057	$2,460,131	$3,818,963	$8,438,828	$18,956,985
Fixed Assets	$165,800	$289,400	$393,500	$748,200	$1,177,400
Less Depreciation	$32,644	$107,486	$220,994	$384,536	$636,903
Other Long Term Assets	$4,000	$3,000	$2,000	$1,000	$0
Total Long Term Assets	$137,156	$184,914	$174,506	$364,664	$540,497
TOTAL ASSETS	**$2,402,212**	**$2,645,045**	**$3,993,469**	**$8,803,491**	**$19,497,482**
LIABILITIES					
Accounts Payable	$23,500	$54,405	$156,746	$314,618	$558,128
Accrued Expenses	$37,292	$77,740	$133,106	$262,756	$442,018
Bank Credit Line	$0	$0	$0	$0	$0
Current Portion LT Debt	$29,576	$48,329	$63,458	$57,814	$51,733
Total Current Liabilities	$90,368	$180,473	$353,310	$635,188	$1,051,880
Long Term Debt	$110,970	$169,941	$209,281	$165,790	$126,305
EQUITY					
Capital Stock	$3,128,000	$5,128,000	$5,128,000	$5,128,000	$5,128,000
Retained Earnings	$0	($927,126)	($2,833,369)	($1,697,122)	$2,874,513
Current Earnings	($927,126)	($1,906,243)	$1,136,246	$4,571,636	$10,316,785
Total Equity	$2,200,874	$2,294,631	$3,430,878	$8,002,513	$18,319,298
TOTAL LIABILITIES	**$2,402,212**	**$2,645,045**	**$3,993,469**	**$8,803,491**	**$19,497,482**
Available Bank Credit	$0	$298,410	$1,241,542	$2,822,527	$5,632,058

10.5 Pro-Forma Income Statement FY 1992-96 ($)

Note: FY 1992 ends March 31, 1993.

	1992	1993	1994	1995	1996
TOTAL REVENUE	**$0**	**$952,000**	**$6,559,851**	**$17,756,580**	**$36,260,692**
COST OF SALES	**$0**	**$27,380**	**$220,184**	**$672,270**	**$1,488,186**
GROSS PROFIT	**$0**	**$924,620**	**$6,339,667**	**$17,084,310**	**$34,772,506**
Percent Gross Margin	0%	97%	97%	96%	96%
EXPENSES					
Engineering	$398,047	$602,491	$954,988	$2,626,530	$5,348,126
Sales and Marketing	$243,222	$1,571,379	$3,258,066	$5,899,037	$10,117,743
Finance and Admin	$329,475	$691,943	$984,967	$1,631,446	$2,252,879
TOTAL EXPENSES	**$970,744**	**$2,865,813**	**$5,198,021**	**$10,157,012**	**$17,718,749**
NET INCOME BIT	**($970,744)**	**($1,941,193)**	**$1,141,645**	**$6,927,298**	**$17,053,757**
Percent	N/A	N/A	17%	39%	47%
Interest Revenue	$54,362	$56,901	$50,125	$83,909	$190,158
Interest Expense	$10,743	$21,951	$55,524	$53,426	$49,274
Net Income BT	($927,126)	($1,906,243)	$1,136,246	$6,957,780	$17,194,641
Percent	N/A	N/A	17%	39%	47%
Taxes	$0	$0	$0	$2,386,145	$6,877,856
NET INCOME	**($927,126)**	**($1,906,243)**	**$1,136,246**	**$4,571,636**	**$10,316,785**
Percent	N/A	N/A	17%	26%	28%

10.6 Pro-Forma Operating Statistics FY 1992-96

	1992	1993	1994	1995	1996
Average head count	6.3	18.9	33.5	61.3	100.8
Average revenue/employee	N/A	$50,326	$195,816	$289,903	$359,908
Average revenue/salesperson	N/A	$158,667	$655,985	$1,183,772	$1,576,552
Expenses per employee	$155,319	$151,497	$155,165	$165,829	$175,868
Unit sales	0	272	2,196	6,712	14,868
Average discount	N/A	30%	39%	47%	51%
Revenue growth rate	N/A	N/A	589%	171%	104%
Expenses growth rate	N/A	195%	81%	95%	74%
Expenses as % of Revenue					
•Engineering	N/A	63%	15%	15%	15%
•Sales & Marketing	N/A	165%	50%	33%	28%
•Finance & Admin	N/A	73%	15%	9%	6%
•Total Expenses	N/A	301%	79%	57%	49%

10.7 Pro-Forma Sources and Uses Statement FY 1992-96 ($)

Note: FY 1992 ends March 31, 1993.

	1992	1993	1994	1995	1996
BEGINNING CASH	**$0**	**$2,260,057**	**$2,034,404**	**$2,065,966**	**$4,464,740**
SOURCES OF CASH					
From Operations					
Net Income	**($927,126)**	**($1,906,243)**	**$1,136,246**	**$4,571,636**	**$10,316,785**
Addback Depr/Amor	$32,644	$74,842	$113,508	$163,542	$252,367
Plus Increases in:					
Accounts Payable	$23,500	$30,905	$102,341	$157,872	$243,510
Accrued Expenses	$37,292	$40,448	$55,366	$129,650	$179,262
Bank Credit Line	$0	$0	$0	$0	$0
Current Portion LT Debt	$29,576	$18,752	$15,130	($5,645)	($6,080)
Proceeds from LT Debt	$110,970	$58,971	$39,340	($43,491)	($39,485)
Proceeds from Stock	$3,128,000	$2,000,000	$0	$0	$0
TOTAL SOURCES	**$2,434,857**	**$317,674**	**$1,461,932**	**$4,973,564**	**$10,946,357**
USES OF CASH					
Current Uses:					
Accounts Receivable	$0	$397,880	$1,257,510	$2,107,979	$3,746,041
Inventory	$0	$17,847	$59,760	$103,111	$137,963
Other Current Assets	$5,000	$5,000	$10,000	$10,000	$10,000
Non Current Uses:					
Fixed Assets	$165,800	$123,600	$104,100	$354,700	$429,200
Other LT Assets	$4,000	($1,000)	($1,000)	($1,000)	($1,000)
TOTAL USES	**$174,800**	**$543,327**	**$1,430,370**	**$2,574,791**	**$4,322,204**
ENDING CASH	**$2,260,057**	**$2,034,404**	**$2,065,966**	**$4,464,740**	**$11,088,893**
CHANGES IN CASH	$2,260,057	($225,653)	$31,562	$2,398,774	$6,624,153

11. Risk Factors

The following highlights some, but not necessarily all, of the risks associated with the development and marketing of the Company's products.

Market Risk of NFS Appliance

The market for an NFS Appliance is unproven. Through a variety of interviews with potential customers and resellers, the Company has attempted to understand customer needs across a spectrum of potential users. This understanding is reflected in the product design and the strategy described herein. The Company believes that it has a reasonable opportunity to capitalize on this understanding. No guarantee can be made, however, that a significant demand for the product will materialize.

Market Risk of NFS

The dominance of NFS in the workstation market may eventually be threatened by a version of the Andrew File System (AFS) incorporated as part of the Open System Foundation's Distributed Computing Environment (OSF DCE). OSF has the backing of IBM, DEC, HP and others. The Company believes that even if OSF DCE and AFS become successful in the marketplace, NFS will not begin to lose market share to AFS until 1996 at the earliest. It is more likely that NFS will continue to dominate long after 1996 due to (1) its market momentum, (2) the pricing policies of OSF, and (3) the complexity and immaturity of AFS. The Company believes that NFS will most likely continue to be the dominant file server protocol, even if OSF DCE eventually becomes the standard for other distributed system protocols such as RPC, naming, time synchronization, and security.

Product Development Risk

An NFS Appliance is a sophisticated product incorporating a variety of technologies, some of which are unproven. Unproven technologies include the Company's proprietary file system design and on-line snapshot backups. In addition, the integration of the Company's various technologies represents a significant development challenge. The Company believes it is taking appropriate precautions to minimize these risks. However, no guarantee can be made that the product will be finished within the time or budget estimates included in this business plan.

Distribution Risk

Third-party distribution channel alliances are essential to the Company's strategy. The Company believes that there are many alternatives for such alliances, and that it will be able to establish a number of them during the coming year. However, if the Company cannot successfully establish and manage these alliances, then the Company's chances of creating and dominating a market for NFS Appliances will diminish significantly.

Competitive Risk

As stated earlier, there are strong competitors in the NFS server market, and there are potential competitors for NFS Appliances. While none of these competitors appear to represent a major threat today, some of them have significantly greater financial and organizational resources than the Company. There can be no guarantee that the Company will compete successfully against these or other companies should they decide to enter the NFS Appliance market.

12. Funds Required and Their Uses

The Company is offering a first-round investment opportunity of $3,000,000 which, together with the previous $125,000 in seed notes that will be converted at the first round, will constitute a total first-round purchase of $3,125,000 in Series A Preferred Shares.

The Company will use proceeds from the first round for product development and testing, and establishing the Company's initial VAR and OEM distribution channels.

The Company anticipates that a second round equity financing of an additional $2,000,000 will occur in November 1993. The Company's financial projections indicate that it will require no further equity financing to become profitable and achieve positive cash flow.

In consideration of seed loans, the Company issued warrants to its current investors that will entitle them to buy, at any time until April 25, 1997, up to $62,500 additional Series A Preferred Shares at the first round price.

Appendix A: Proposed LADDIS Benchmark

Until recently, there was no consistent way to measure or compare the performance of NFS servers. The LADDIS Group has proposed an NFS server benchmark that will enable users to compare the performance of NFS servers.

The acronym LADDIS was derived from the names of the participating NFS vendors: Legato, Auspex, Data General, DEC, Interphase and Sun. The LADDIS Group submitted the proposed benchmark to the Systems Performance Evaluation Cooperative ("SPEC") in August 1991. As with other proposed SPEC benchmarks, all SPEC member companies must scrutinize the LADDIS benchmark before SPEC adopts it.

The LADDIS benchmark generates a synthetic load of client requests that comprises a realistic mix of NFS operations and measures the response times of the server. The result is a graph of average response time (in milliseconds) as a function of throughput (in NFS operations per second).

Evaluation copies of PRE-LADDIS Version 0.0.11 were made available in February 1992. SPEC should release the official benchmark in late 1992. It is written in C and runs on most UNIX workstations.

The proposed LADDIS benchmark has already received attention in the UNIX trade press. Published LADDIS results will be an important tool for selling NFS Appliances.

Appendix B: Price/Performance Competitive Analysis

Interviews with system administrators indicate that price is less important for an NFS Appliance than reliability, high performance, and easy administration. Nevertheless, it is difficult to justify the purchase of a new product if it isn't cost effective compared to existing products.

This appendix compares the NFS Appliance with low-end Sun SPARCserver 2 systems, mid-range Sun 670MP "Office Servers," and high-end Auspex NS-5000s. The basic result is that NFS Appliance software running on today's PC hardware would be more cost effective than any other NFS server currently on the market. Each current server is cost effective only in a specific performance range. An NFS Appliance using 33 MHz 486 PC hardware would be the most cost effective NFS server throughout the low-end, the mid-range, and even part of the high-end performance range.

It is important to understand that this comparison is based on performance estimates of the NFS Appliance software running on *today's* PC hardware. It is an "apples-to-apples" comparison in that it compares today's PC hardware to today's UNIX hardware. Commodity PC hardware performance is evolving more quickly than proprietary UNIX platforms, so over time the advantage of the NFS Appliance is likely to improve, rather than degrade. As a result of its special-purpose design, optimized specifically for NFS, the NFS Appliance software is best positioned to benefit from these PC performance improvements.

Cost effectiveness is defined as price/performance where performance is the number of NFS operations per second ("NFS-ops") that a server can provide as measured by the "nhfsstone" NFS benchmark. The measure of cost effectiveness, is NFS-ops per $1000 of server price. "Nhfsstone" was developed by Legato; it is the public-domain predecessor to the LADDIS benchmark described in Appendix A.

The first section below compares the NFS Appliance with NFS servers from Sun and Auspex. Sections that follow describe the assumptions used to generate pricing and performance information, and the exact configurations used in the analysis.

Analysis

Figure 1 shows how the performance of each server increases as a function of its price. The first price point for each server is based on an entry-level system. The second point represents a configuration with 4 GB of disk and a Prestoserve-style non-volatile RAM option. (The Prestoserve option is not required, but performance is much lower without it.) The last point represents an 8 GB system. The exact configurations are described below.

A steep curve indicates that performance rises quickly as the customer buys more server hardware. The NFS Appliance curve is very steep for two reasons: 1) With file compression it costs less money to increase the disk capacity of an NFS Appliance. 2) UNIX uses system resources such as I/O bandwidth and compute power much less efficiently than the NFS Appliance software, resulting in curves that flatten out as performance bottlenecks are reached.

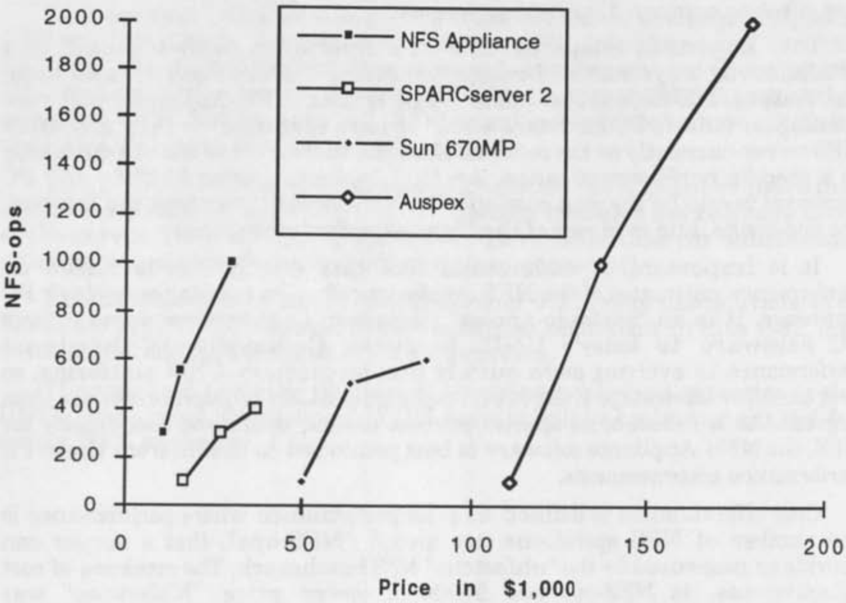

Figure 1: The performance of NFS servers measured in NFS-ops as a function of server price.

Ignoring the NFS Appliance, Figure 1 shows that there are three performance categories for servers: a low-end with up to 300 or 400 NFS-ops, a mid-range with up to about 600, and a high-end with over 700. The Sun SPARCserver 2 ("SS2") is a low-end machine that can offer reasonable throughput, but it cannot be expanded to support many Ethernets or disks. The performance levels shown here can be attained only by removing the video display (to free up the SS2's third SBus slot), which many customers are not willing to do.

The Sun 670MP "Office Server" has somewhat higher performance than the SS2, and it is much more expandable. For instance, the 670MP can have up to four Ethernets and as many as twenty disk drives. This analysis uses the 670MP instead of Sun's 690MP "Data Center Server" because the 690MP would probably not provide much better NFS performance, although it can accommodate more disks. A fully configured 690MP costs well over $200,000.

The Auspex server provides extremely high performance, almost triple that of the Sun 670MP, but its entry price is so high that it makes sense to buy an Auspex server only if it will be used in place of multiple Sun servers.

Figure 1 makes it easy to compare the price and performance of the different servers, but not their cost effectiveness. Figure 2 shows the NFS-ops/$1000 price for each server as a function of performance.

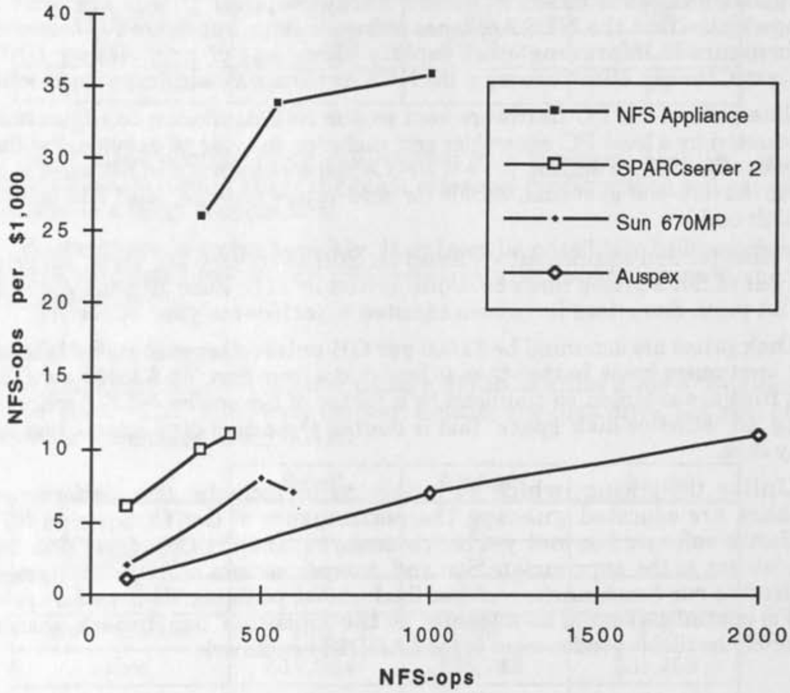

Figure 2: NFS-ops/$1000 price of each server at different performance levels.

Figure 2 makes it easy to compare the cost effectiveness of different servers. The SS2 is more cost effective than the 670MP, but only in its narrow performance range. The Auspex NS-5000 is barely more cost effective than an SS2, and only at very high performance levels. Customers obviously don't buy Auspex servers because the performance itself justifies the cost. Auspex customers may believe that by reducing the number of UNIX file servers, they will simplify network administration or improve network reliability.

Figure 2 also illustrates why the NFS Appliance will be appropriate for customers with a very broad range of performance requirements. The entry price is very low, so even small sites can afford an NFS Appliance. Yet, by adding additional controllers to the same basic 486 PC hardware, the NFS Appliance performance surpasses even the highest-performing Sun server. Throughout this whole range, the NFS Appliance is more cost effective than any UNIX-based server.

Methodology

The above analysis is based on current hardware prices. Prices are likely to change by the time the NFS Appliance software ships, but since PC hardware performance is improving more rapidly than that of proprietary UNIX platforms, the cost effectiveness of the NFS Appliance should improve as well.

The price of the PC hardware used in the NFS Appliance configurations was quoted by a local PC assembler and includes one year of service. For this analysis, the average selling price of NFS Appliance software is assumed to be $4,000 for low-end systems, $6,000 for mid-range systems, and $12,000 for the high-end.[1]

Prices for Sun equipment are based on Sun's February 24, 1992, US Price List. All of Sun's prices have been discounted by 35%, since nobody seems to pay list price. Sun prices have been adjusted to include one year of service.

Disk prices are assumed be $2,000 per GB unless otherwise noted because most customers know better than to buy disks from Sun (at $3,300 per GB). User files are assumed to compress by a factor of two on the NFS Appliance, giving an "effective disk space" that is double the actual disk space, less the parity disk.

Unlike the pricing, which we believe to be accurate, the performance numbers are educated guesses. The performance of the Company's NFS Appliance software has not yet been measured, and the Company does not have access to the appropriate Sun and Auspex servers and the equipment required to run benchmarks. As described above, performance numbers refer to NFS operations/second as measured by the "nhfsstone" benchmark which is a publicly available predecessor to the LADDIS benchmark.

---

[1]The Company's actual pricing strategy may vary.

NFS Appliances

All NFS Appliances listed here are based on a 33 MHz 486 PC with an EISA bus and a UPS.

	Small	Medium	Large
Memory (MB)	16	32	64
Ethernets	1	2	4
Disk (GB)	2	4	8
NFS ops/sec	300	550	1000
Price	$11,432	$16,242	$31,862
ops/$1,000	26	34	36

All of these configurations can be built from the same basic PC, so it is possible for a customer to buy the least expensive configuration and expand it over time to a large configuration.

PC hardware is evolving quickly. It will soon be possible to build much more powerful systems based on 586 motherboards with integrated multi-channel SCSI and Ethernet.

Low-End Sun Server

These configurations can be built using a SPARCstation 2 and a "File Server Option" pack. (The "File Server Option" includes one disk drive. The rest of the drives are purchased at $2,000/GB.)

	Small	Medium	Large
Memory (MB)	32	32	48
Ethernets	1	1	2
Disk (GB)	1.3	4	8
NFS ops/sec	100	300	400
Price	$17,544	$28,748	$38,253
ops/$1,000	6	10	11

Many customers would find the "large" configuration unacceptable because it requires them to remove the system's video display. (The SS2 has only 3 SBus slots, one of which is taken by the video frame buffer.)

Mid-Range Sun Server

These configurations are based on Sun's 670MP "Office Server" using "Network Coprocessor" cards built by Interphase for higher performance. Sun includes the first disk in the base configuration. The rest are purchased at $2,000/GB.

	Small	Medium	Large
Memory (MB)	64	64	64
Ethernets	1	2	4
Disk (GB)	1.3	4	8
NFS ops/sec	100	500	600
Price	$50,466	$63,130	$86,961
ops/$1,000	2	8	7

It would be possible to spend much more on a 690MP "Data Center Server" which uses IPI disk drives, but the Company does not believe the 690MP provides higher NFS performance. Customers may buy it for the extra disk capacity the 690MP can support.

High-End Auspex Server

These configurations are based on the Auspex NS-5000 server. The Auspex NS-3000 is slightly less expensive, but it is no more expandable than a 670MP, and as a result it has sold poorly. Because Auspex disks require special cases that only Auspex supplies, they have been priced at $6,000 per 1.2 GB disk, which is what Auspex charges, rather than $2,000/GB.

	Small	Medium	Large
Memory (MB)	32	32	64
Ethernets	2	4	8
Disk (GB)	1.3	4	8
NFS ops/sec	100	1,000	2,000
Price	$110,000	$136,000	$177,480
ops/$1,000	1	7	11

These prices are not as up-to-date as the Sun prices, and do not include any discount. The performance numbers are probably accurate because Auspex's published benchmarks tend to be more reliable than Sun's.

Appendix C: Cost of Ownership Competitive Analysis

A typical site with one NFS appliance[2] will save $7,300 per year in system administration costs, gain $12,000 per year in user productivity due to better reliability, and gain $24,000 per year in user productivity due to faster response.[3]

Reliability

The Company believes an NFS appliance will experience an average of less than two outages per year, compared to six for Auspex and 12 for Sun. Assume that each outage requires an average of two hours to repair, and that 15 users are rendered unproductive for an hour during each outage. Assume the loaded salary of a user is $150,000 per year ($80 per hour). Thus, relative to Sun, an NFS appliance saves:

10 hours × 15 users × $80/hr = $12,000/yr

Response Time

Performance measurements at UC Berkeley suggest that a workstation generates an average of 15 NFS I/O operations per second. The Company believes an NFS appliance will give, on average, 7 ms faster response time per I/O operation relative to either Sun or Auspex. Assuming only 10% of the time each user spends waiting for response could otherwise be spent more productively, 15 users will gain in productivity as follows:

15 users × 15 IOP/sec/user × 7 ms/IOP × 0.1 = 0.16 hours gained per hour

At $80 per hour, this amounts to a gain of

0.16 × 8 hr/day × 240 days/yr × $80/hr = $24,192 per year.

System Administration

The Company believes the estimates shown in the following table summarize the amount of time that will be required for the various system administration tasks.

---

[2]These comparisons assume the mid-range NFS Appliance hardware configuration described in Appendix B.

[3]These comparisons are with respect to Sun's servers. With respect to Auspex, the numbers are: $7,300 savings in system administration, $4,800 gain in productivity due to better reliability, and $24,000 gain due to faster response.

	Sun	Auspex	NFS Appliance
Installation and Upgrades	6 hr/yr	18 hr/yr	2 hr/yr
Outage Repair	2 hr/mo	1 hr/mo	2 hr/yr
File Backup	3 hr/wk	3 hr/wk	3 hr/mo
Total	186 hr/yr	186 hr/yr	40 hr/yr

The major difference is due to snapshot file backups and RAID 4 reducing the frequency and complexity of doing backups to tape.

Assuming the average loaded cost of a system administrator is $50 per hour, the cost of ownership savings for an NFS appliance is:

$(186 - 40) \times 50 = \$7{,}300$ per year

Appendix D: WAFL — Write Anywhere File System Layout

The WAFL file system enables many important NFS Appliance features, and will provide the Company's competitors with a significant barrier to entry. The Company intends to allocate considerable effort over several years to building WAFL into a core technology capable of supporting NFS and other file server protocols at very high levels of performance.

WAFL will support the following features:

- Backup snapshots
- Efficient RAID
- Write caching in main memory with UPS protection
- Single expandable partition
- Fast reboot without consistency checks
- File compression on disk

Snapshots

As described in Section 3 "Product Concept," a snapshot is a read-only duplicate of a file system. Snapshots make it easy for users to recover accidentally deleted files. A snapshot provides an unchanging file system image that can be backed up while the NFS Appliance is on-line, even during normal operating hours. By default, the NFS Appliance takes a snapshot every night at midnight and stores it on-line for seven days. Each Sunday's snapshot is kept for a month.

WAFL uses a copy-on-write technique to create snapshots in just a few seconds without consuming any disk space. Over time, as users modify files that are saved in a snapshot, extra disk space is required to hold the changes. On most servers, file system contents are fairly static, so keeping snapshots shouldn't take much disk space. In environments where files change rapidly it may not be possible to keep snapshots for a full week. WAFL allows users to remove files in a snapshot to save disk space in cases where files that don't need to backed up change rapidly.

Efficient RAID

The primary disadvantage of RAID is that writing to a RAID volume can require four times as many disk accesses as writing to a regular disk. This is especially unfortunate in an NFS environment where read caching on client workstations increases the percentage of disk writes on servers. WAFL's "write anywhere" design, together with write caching, provide great flexibility in allocating disk space and scheduling writes to optimize write performance. The UNIX file system cannot work efficiently with RAID.

Write Caching

The NFS protocol specifies that servers must write data to "stable storage" before acknowledging requests, so every NFS Appliance will be required to include a UPS. NFS Appliance software will check for the presence of a properly functioning UPS.

Using a UPS to protect primary memory allows WAFL to cache writes, and acknowledge NFS write requests immediately. The UNIX file system writes each block of data to disk as soon as it is received, and UNIX doesn't acknowledge an NFS write request until the write to disk is complete. Even UNIX-based NFS servers with non-volatile RAM (as provide by Legato's Prestoserve product) must allocate disk space and copy multiple 8 KB blocks into non-volatile RAM before acknowledging a write request. This imposes a significant overhead.

Write caching provides other advantages as well. Many files are removed shortly after being created. With write caching, these files may never reach disk at all. Write caching allows whole files to be written to disk at once, which facilitates efficient data compression, and improves read performance by making it easier to allocate the file's blocks near each other on disk. Write caching allows data compression to be performed without increasing NFS response time. Finally, write caching increases the flexibility of WAFL's write policy, which allows improved RAID performance.

Single Expandable Partition

Because an NFS Appliance uses a single partition that can grow to be very large, the Company designed WAFL to support partitions up to 8 terabytes in size. Assuming the current rate at which disk capacity increases does not change, this will not be a limiting design factor for at least 10 years.

Fast Reboot

On UNIX file servers, large partitions increase boot time because of consistency checks required by the UNIX file system. These checks are getting slower because while disk capacities are increasing rapidly, transfer rates are not. Consistency checks take about 5 minutes per gigabyte, which means that a ten gigabyte partition would take almost an hour to check. UNIX partitions are currently limited to two gigabytes, but even when vendors increase that limit, as some are doing, UNIX system administrators will probably continue to use many small partitions that can be checked in parallel, rather than a single large partition.

WAFL boots quickly because the file system is never inconsistent on disk. The time required to boot an NFS appliance is spent loading software, checking hardware, and spinning up the disks—not checking file system consistency. The file system remains consistent even if power fails or the system crashes. An NFS Appliance with a UPS will not lose any data when the power fails. Even if the UPS fails, changes that occur during the last two or three minutes before a power failure may be lost, but the file system on disk will still be consistent.[4]

---

[4]In database terms, WAFL does a large "group commit" of NFS requests every few minutes. If power fails, uncommitted NFS requests are lost, and the file system reverts to its state as of the last group commit.

File Compression

Because WAFL caches writes before allocating disk space, it can compress multi-block chunks of data just before allocating space and writing the data to disk. WAFL decompresses data when it reads it from disk.

WAFL's write caching removes data compression from the response paths of all NFS operations. Decompression after a disk read is often in the response path of an NFS operation. However, existing hardware performs decompression about three times faster than compression. The percentage increase in response time for reading data is small, and it only occurs when an NFS request reads data that are not in cache. When somebody properly integrates a data compression chip into a SCSI disk controller, decompression can also be removed from all response paths.