Key Points:

- NFS has addressed problems in file access and interoperability.
- There is industry commitment to grow the technology.
- NFS will be there to meet whatever file system comes along over the next few years.

Slide 13: Summary

In summary, several challenges exist as we start through the ‘90’s for distributed file systems. Ease of use and administration, security, scalability and performance, and heterogeneous interoperability are the problems NFS has been solving and will continue to solve into the ‘90’s. NFS will take the industry forward, with a well-defined plan over the next two years, and continued innovation in the ONC vendor community. And as NFS competitors come to the fore, they will find NFS waiting for them on every platform they undertake to support.

Key Points:

- Success has its problems—but I’m not complaining too loudly.
- As NFS networks grew, administration issues exacerbated—**critical to address**.
- Growth in vendor community has introduced technology propagation problem—so Automounter helps solve problem one, but is not widely available on all platforms yet.
- Road map addresses technical issues—these are simple to first two points.
- Again, it’s good to be in position of complaining about the problems of success—now the burden is on the ONC Vendor community (including Sun) to get cracking on getting solutions out to their customers.

Slide 12: Industry Challenges for NFS

A problem with the success of NFS is the typical problems associated with widely implemented, widely used systems. As NFS network grew, administration issues came to the fore. Technology has been introduced to simplify administration, and also provide increased transparency. The Automounter is an example of such a tool. Legato Network Systems Inc.’s Management Tool for NFS is another.

The Automounter provides on demand access to files, and facilities to define a consistent shared image of the file name space, which can be centrally administered when coupled with a naming service such as NIS.

Much more work needs to be done in the area of administration.

Another problem introduced by NFS’ success, is that there is a large number of vendors shipping product. Because of product release cycles, the propagation delay of new features to existing systems presents a problem. For example, the Automounter provides an administrative solution, but it is not yet available on all NFS platforms. The NFS Vendor community, including Sun, has a challenge before them to propagate new technology as quickly as possible, to make innovative solutions generally available to the customer concerned with a multi-vendor network.

Technical challenges are covered in the Road Map, and look at better authentication, scalability, performance and availability. I believe, however, that the administration issues and practical problems of technology propagation are thornier issues than those solve simply with good technology. More work is needed here.

Key Points:

- Features phased in over the next two years—graceful evolution rather than disruptive revolution.
- Preserve technology investment.
- Local disk caching reduce server loading, increase performance by caching to local disk.
- Read-only replication will provide client fail-over to replica file systems—tools to manage replication.
- Security—making Kerberos standard available as part of NFS (though some work already done at MIT).
- NFS over TCP and Dynamic Retransmission provides enhanced performance over wide area networks now.
- Protocol revision still being considered—reevaluation as we look back on last five years.

Slide 11: The NFS Network in 1992

Here is the view of the ONC/NFS network in 1992.

Local disk caching and read-only replication will address critical issues of scalability, performance and availability. Local disk caching will reduce server loading, enhance performance, and simplify administration by allowing NFS clients to fault files into a local disk cache on reference. The cache will persist across re-boots. Read-only replication will provide for client fail-over to replicated, read-only servers. Administrative tools will exist to manage and replicate data on the network.

Security will be addressed in the form of enhanced, integrated authentication services. ONC has had a public key scheme for strong authentication since 1988. Kerberos and additional Public Key-based authentication schemes will provide the basis for enhanced security. Kerberized NFS will leverage off the growing acceptance of Kerberos authentication as an industry standard.

Wide area data access, and dealing with complex network topologies are addressed by re-hosting NFS over TCP, and the applying congestion control techniques to NFS over UDP.

Exploration of a protocol revision to NFS continues. However, initial work led to the conclusion that the cost of a protocol conversion would be high, and additional work within the framework of the current protocol could be done more timely, and maintain compatibility with the huge installed base. The reply cache work from DEC addresses some correctness problems in the protocol with an implementation attack, maintaining protocol compatibility. Some additional early justifications for a protocol revision included better support for non-UNIX platforms. The availability of so many non-UNIX implementations of NFS throws doubt at this reason for a protocol revision. The protocol revision is being studied now from a position of more than five years of experience of commercially available product.

Key Points:

- ONC Road Map: performance, scalability, availability and security.
- Thrust over next two years is to leverage installed base with compatible, licensed technology.

10: ONC Road Map: Adding Functionality, Performance

From Sun’s perspective, the next two years will see the introduction of several critical features to NFS which enhance performance, scalability and security. The features are collected under the marketing umbrella of the ONC Road Map. The thrust over the next two years is to leverage the existing base of NFS, maintaining compatibility at the protocol level of NFS, by providing innovative enhancements to the base system. Growth-oriented evolution, rather than disorganized revolution, is the direction of ONC.

The NFS futures reflect priorities of ease of administration (tools to manage replication), security (with adherence to standards), and performance and scalability.

The NFS enhancements in the ONC Road Map are part of a framework of ONC enhancements.

Key Points:

- Innovation in industry with targeted solutions reflects the installed base justification.
- NFS passed a critical point in 1988—justification for some very interesting work.
- NFS continues to grow.

Slide 9: (Vendor Targeted Platforms)

At Sun, we have encouraged innovation in the ONC vendor community, and this work is partly represented by the hardware solutions being implemented by vendors to specifically target the NFS market.

The Cayman Gator Box was an initial solution which provided an Appleshare to NFS gateway to enable NFS access from a Macintosh.

Epoch Systems provide huge file servers for NFS clients by serving up optical storage juke boxes, and have contributed to the understanding of NFS performance through continued innovation in their implementation.

The Legato Systems Prestoserve board provides write performance acceleration for NFS servers, through a combination of hardware and software.

The Auspex file server represents an architecture tuned to exploit NFS to yield greater aggregate throughput and scalability, in terms of the numbers of clients.

The industry has contributed software solutions to ONC, including the Reply Cache technology from DEC and the NFS to AFS protocol converters (from the University of Michigan and Transarc).

Summary: I think the trend of significant, innovative technology arising from the ONC vendor community corroborates the premise that there is wide spread NFS technical expertise. Technological advancement in the industry is not simply a function of having some good concepts, but implementing and making those concepts practical. Some of the work needed to make a concept into a product involves simply a great deal of work, tuning and implementation experience. NFS passed a critical point sometime in 1988, where technical expertise and justification for the business case for providing the technology fueled a sharp jump in available implementations, as shown in the Connectathon attendance graph. I believe the industry has only started to reap the benefits of this burgeoning technology base, and that we can expect strong growth and continued innovation for the foreseeable future of NFS.

The trend is clear, so my children say: NFS has undergone strong growth, and there are no signs of it stopping. It is Sun’s intention to continue to grow NFS, and work with ONC vendors to continue strong growth into the ‘90’s. I believe that the growth curves shown for NFS reflect a general trend of technology dispersion that any emerging distributed file system technology would experience, if successful. The importance in viewing the history of NFS, and viewing its installed base, is to recognize that sound technology fueled this growth, and the large installed base and vendor community ensure sufficient critical mass in platforms and technical expertise to continue this growth through the ‘90’s.

Key Points:

- Most implementations are UNIX, from the reference port.
- Several PC implementations, several MAC, several VMS, VM and MVS implementations. Non-UNIX implementations solved lots of problems through innovation. From the specifications.
- 4.4 BSD planned native support, and SVr4 and OSF/1. NFS is ubiquitous.
- Growth of technology shows no signs of stopping. Features propagate, new technology is in the works to leverage existing base.

Slide 8: (Implementations)

Most NFS implementations are on UNIX platforms. They’re well understood.

Implementations from the public domain specifications include several critical platforms. Several client PC-NFS implementations exist, three implementations exist for the Macintosh. NFS exists for the mainframe running VM and MVS, several NFS implementations are available for VMS including a client implementation, and Ada and Lisp implementations are available. Novell described their NFS implementation for Portable Netware at InterOp ‘90 in October. It was implemented from the public domain specification also. These are some of the more interesting implementations, as they addressed issues of file access in a truly heterogeneous environment. These non-UNIX implementations of NFS have set the expectation in the industry for other distributed file system technologies to provide similar, comprehensive support for heterogeneous file access.

With the introduction of NFS in UNIX System V Release 4, and the planned inclusion of NFS (including support over TCP) in BSD 4.4 and OSF/1, NFS will be available as a native technology on core UNIX platforms.

Key Points:

- Growing installed base helped provide justification for additional features, not only from Sun.
- NFS is well understood.
- Huge base of technical expertise exists in the industry—which took years to develop.
- NFS over TCP and Reply Cache technology indicate strength of industry expertise.
- Some things are gained only over time, and with a lot of work. NFS has put in the time—or more accurately the vendors.
- Huge investment not only in technology but in expertise.
- Yesterday—Highly available NFS research at IBM. Such work is on a sharp increase.

Slide 7: (Features)

The large installed base indicated by the licensing and testing figures has served to accelerate NFS technology growth and evolution. As the installed base grew, it justified further development and introduction of new features.

From the definition of the basic ONC architecture in 1985 to future features such as local disk caching, enhance authentication and replication, NFS has been characterized by strong steady, evolutionary additions of features. Fundamental technology which accompanied the introduction of NFS was the Virtual File System interface.

NFS behaviour is well-understood. Implementing NFS and tuning performance are well-defined activities today. Contributions of additional features, such as the reply cache technology from DEC and NFS over TCP as part of 4.3BSD Reno demonstrate that the base of technical expertise in NFS and ONC technology does not only reside at Sun. Recently, at the October 1990 ONC Industry Networking Conference, IBM described research they were involved with exploring high availability approaches to NFS server implementations using dual redundant hardware. This availability of technical expertise in a distributed file system technology is a function of time as the technology is disseminated, and NFS has put the time in to develop a strong support base. The addition of new features is expected to increase.

Key Points:

- Connectathon is a practical, annual event to increase interoperability of ONC platforms.
- First Connectathon has now passed into folk lore: 14 bleeding edge companies—some completing their implementations at Connectathon for a Uniforum demo.
- 24 hours a day now, two weeks long. more than 50 vendors, with over 65 implementations.
- Number does not reflect some vendors who had attended previous Connectathons.
- Connectathon lightweight model hugely successful.

Slide 6: (Connectathon Participation)

The bottom graph shows the Connectathon participation over the years. Connectathon is a yearly ONC testing event which has been a fixture of the ONC community since the beginning. Connectathon allows engineers to check their NFS software for interoperability with other implementations.

In early 1986, fourteen brave companies got together at the first Connectathon to evaluate interoperability of their implementations of NFS. Some went to there to complete their implementations in time for an interoperability demonstration at Uniforum several weeks later.

It is at Connectathon, where last February more than 50 vendors converged to test upwards of 65 NFS implementations, that the existence of a strong NFS community becomes apparent. It is a fun event, 24 hours a day testing for two weeks, interspersed with technical presentations and discussions where ONC engineers from various companies exchange ideas. The graph does not reflect the total companies which have participated in Connectathon, since some companies have not attended all Connectathons.

Key Points:

- To the industry, the commercial licensees topped 150 by the end of 1989.
- Figures do not reflect those implementations made from public domain specification.

Slide 5: (Commercial Licensees)

Commercial licensees topped 150 at the end of 1989. NFS is so widely licensed that the number of new licenses is showing a tendency to flatten out as fewer platforms exist which do not provide NFS capability.

The figures do not reflect the large number of vendors who have implemented NFS from the public domain specifications. These are typically non-UNIX implementations, though the NFS implementation in UNIX 4.3BSD Reno and 4.4 BSD was written from the public domain specification.

Key Points:

- NFS as phenomenon. Unprecedented growth—set industry trend for open licensing and availability.
- Licensing from Sun reached 250 by end of 1989.
- Includes university licensees, which have used NFS as a base for research into distributed file systems.

Slide 4: (All ONC Licensees)

NFS’s growth as a multivendor solution for file sharing was unprecedented, and its history presents a model for the successful introduction and dissemination of an open technology.

NFS was made easily available by Sun and other distributors in the form of licensed reference implementations. The total number of licensees licensed by Sun topped 250 at the end of 1989. The total number of licensees includes Universities, which have provided significant research into new ways to use and extend NFS.

Key Points:

- That NFS simply filled a vacuum is a weak argument, NFS design enabled its success.
- NFS is clean, (relatively) stateless, exhibits simple error recovery (not subject to complex failures), and is lightweight and provides good performance.
- NFS supports heterogeneous file sharing. This is simply a proven fact. From PC’s (<70Kbytes memory) to super-computers.
- Diskless for low cost, dataless for ease of administration and diskful for full participation (client and server).
- NFS is a very symmetrical model, and a machine (UNIX) can export its files and become a server.
- NFS platforms serve up the native file system—shared with local users.
- ONC specifications in public domain, and reference implementations readily available.

Slide 3: NFS Strengths

One reason cited for the strong growth of NFS, in the ‘80’s, was the lack of any real alternative on its introduction. I believe, however, that this is a naive viewpoint. NFS has well-known strengths which have enabled its success.

NFS established itself as the solution for heterogeneous platform and operating system connectivity early with the introduction of PC-NFS and implementations of NFS for the VMS world. It is a proven technology, over and over again. It operates over a wide range of platforms, scaling easily down to the PC where a client implementation is done in less than 70Kbytes of memory, and scales up to provide access to large mainframes and super-computers. NFS runs very well diskless, is often run in a dataless configuration where a small local disk is used to off-load the server, or in a diskful configuration, where a machine can become an NFS server by simply exporting portions of its file systems.

NFS has been the cornerstone of Sun’s workstation strategy, and the introduction of diskless support through NFS in 1988 completed a total solution to file sharing and access on the network. NFS and ONC continue to be absolutely critical to Sun’s success.

Broad reason’s for NFS’ success can be found in its goals and design.

NFS is characterized by a clean design with straightforward goals: provide transparent file access in a platform independent way. NFS defines, ideally, a stateless protocol for file access. This has long been a point of contention from detractors of the design philosophy of NFS, but I challenge their critiques by asserting that NFS’ simplified semantics, reduced complexity of implementation, and straightforward error recovery techniques have been a critical factor in the acceptance and propagation of NFS through the industry. The ONC platform, represented by the Remote Procedure Call and the External Data Representation, provides key enabling technologies for NFS. The architecture for ONC was in place on the introduction of NFS and provided a foundation for future growth. The simplicity of the design enables an engineer to quickly assimilate the technology, to understand its design and implement solutions with NFS. Other distributed file system technologies exhibit a degree of complexity which may hinder transfer of technology to the industry at large.

Critical specifications for the ONC protocols are in the public domain, and readily available as RFC’s. The protocol specifications describe the interaction of clients and servers succinctly, because the interactions themselves are succinct, simple and easy to understand (and apply to a wide variety of systems).

Several reference implementations are available. The easy availability reference implementations helped fuel strong growth in NFS availability.

Key Points:

- ‘80’s provided era of growth.
- NFS was introduced into a vacuum, and was the first clear, available commercial solution for UNIX and heterogeneous file sharing.
- Today file sharing is ubiquitous.
- Issues remaining are ease of use (transparency, data location).
- NFS has made platform interoperability the norm in the industry—it is simply expected now.
- Ease of **administration** has lagged ease of use in mind share.
- Authentication and security are being addressed.
- High speed, wide area issues are emerging
- Huge scale is on horizon (nation-wide file system)

Slide 2: Distributed File Systems: Some of the Challenges for the ‘90’s

The 1980’s provided for strong growth in the distributed file systems. When NFS was introduced in the mid-80’s, it filled a vacuum by providing transparent file access across a range of platforms. Though the concepts of distributed file systems were not unknown, NFS was the first clear commercial solution to emerge which addressed the problem of platform interoperability.

Today, distributed files systems are commonplace in the PC, UNIX and large system arenas. Through the ‘80’s, significant technical hurdles in performance, transparency and interoperability were cleared. However, it would be naive to assume that all the challenges have been met.

The primary challenge to distributed file systems has been ease of use. I consider issues of transparency, data location and binding to reside under the broad umbrella of ease of use. Ease of use today implies a consistent solution which bridges multiple vendor platforms. NFS pioneered the reality of multivendor, multiple platform interoperability at the file access level. Heterogeneous file sharing remains one of its great strengths.

Though most distributed file systems today provide for access transparency and ease of use, to users and their applications, ease of administration issues remain. Ease of administration seems to revolve around making administration of distributed technology as simple and straightforward as centralized single system administration.

As data sharing and distributed file access became ubiquitous, issues of authentication and data protection on the network became critical. As networks increase in complexity in the ‘90’s, issues of security and authentication are made all the more complex. Technology is still emerging to deal effectively with authenticating distributed data access.

Much of the focus for distributed file systems solutions in the ‘80’s was at the work group and medium size networks. Incorporation of emerging technologies in connectivity providing higher performance levels over wider areas are challenging distributed file system technologies to exploit them efficiently.

Key Points:

- Five Years of NFS—Industry Standard—with ONC
- Cover reasons for success, strengths, future directions. Reflect challenges facing industry.
- NFS in general, rather than one particular implementation. Use implementations to illustrate.
- Sun is an ONC Vendor, like other ONC Vendors

Slide 1: The Network File System: Targeting the 1990’s

I am going to talk about the position of the Network File System as we start the ‘90’s. For more than five years, NFS, within the framework of Open Network Computing (ONC), has provided the industry with a stable, high-performance distributed file system. NFS can be considered, as one person remarked at the ONC Industry Networking Conference yesterday, the standard I/O of distributed file systems. It has established itself as the standard distributed file system for multi-vendor connectivity and is the standard by which other technologies are judged.

I want to touch on some of the reasons for its success, discuss its strengths, and present some future directions for NFS. I will call out challenges facing NFS as we enter the ‘90’s.

When I mention ONC vendors, particularly in discussing the challenges facing them, I include Sun as an ONC vendor.

Note: Standard is defined in the American Heritage Dictionary as “An acknowledged basis for comparing or measuring”.

The Network File System: Targeting the 1990’s

Brian Pawlowski
bpawlowski@eng.sun.com