****
* COMPETITIVE FLASH *
****

Copy: Staff, Bob Lyon

FYI
Bernie
03/25/85

COMPETITIVE ANALYSIS, PRODUCT MARKETING

FEBRUARY 11, 1985
CF:01:85

Deb Sheedy
x 4689

SUN MICROSYSTEMS' NETWORK FILE SYSTEM

INTRODUCTION

In the last few weeks there have been an increasing number of questions from the Field concerning the Network File System, NFS, which Sun announced last November. Most of these questions deal with very basic functionality items and the "standards issue". This FLASH is written to address these questions being received and to help you better understand what NFS is and more importantly, what NFS is not. Additional information on the NFS will be made available in future issues of SALES UPDATE.

WHAT IS THE NFS?

NFS allows transparent sharing of file systems in a heterogeneous environment consisting of different workstations/computers, operating systems and networks. Thus NFS will allow users transparent access to files stored on remote machines, similar to the way the Apollo's Single Level Store System functions.

> Apollo had a fully distributed file system with the first shipped node, five years ago! Previously, Sun's filesystem consisted of disk partitions being assigned to each user, and to access another user's file, the file had to be copied into the requestor's own disk partition.

This analysis represents the opinions of its individual authors and not a position of Apollo Computer Inc. It is based on the best information available; however, no representation can be made as to its accuracy or completeness. This document is for APOLLO INTERNAL USE ONLY.

HOW IS SUN POSITIONING THE NFS PRODUCT?

Sun is attempting to position NFS as an "open systems information network", an industry standard. By publishing the NFS protocol specifications, the remote procedure call (RPC) protocol, in full detail, and licensing NFS technology to manufacturers and others wishing to implement NFS for their machines, Sun hopes to establish itself as "open" and heterogenous and its competition (read Apollo) as "closed", proprietary and homogeneous.

→ Gould
Pyramid
MSDOS
maybe VMS ?

Apollo is a strong supporter of efforts to bridge heterogenous systems and establish cross-vendor standards, and is pursuing these efforts via the appropriate Standards Committees. While not questioning the main intent of Sun's support of industry standards, it does allow Sun to receive important research that it might not be able to solely fund. Additionally, even if no other vendors support NFS, Sun has reaped the public relations bonanza associated with such a standards effort.

Now that Sun can operate in a heterogenous environment, it extends its product line by "proxy". Sun can answer questions regarding the lack of a high-end workstation product, and other gaps in its product line, by stressing the connectability of NFS. While Apollo's interconnect strategy allows heterogenous networking thru standard protocols and interconnects, our overall product line breadth means critical product needs are not filled with "proxy" products.

IS NFS DISTRIBUTED UNIX?

The NFS protocol is NOT an extension to UNIX, it is NOT a step towards a distributed architecture. Rather, it is an application level event that fits into Sun's networks services architecture. Files are still addressed by inodes (operating system interface node) at the lowest layer, but Sun has introduced a new structure, the vnode (virtual file system node), to address files at the higher layers.

This additional overhead, layer of software, now needed for Sun-to-Sun communications may result in a sacrifice in performance, when compared to Apollo's approach where the integration and software/hardware tuning is done at the operating system level.

NSF utilizes Ethernet as the transport and data link layers which means its performance is limited by the performance of the Ethernet. Ethernet "bogs down" at high load levels when the probability of collisions (and the corresponding need to re-transmit) increases.

SEE Ethenet
"Commentar
(Enclose

50 nodes/

WHAT ARE THE MAJOR LIMITATIONS OF NFS?

NFS ASSUMES A NONTHREATENING ENVIRONMENT. In the NFS environment, nodes are assumed to be trusted. Any user can edit the password file on their node, remove all the passwords from the other user accounts, and appear as anyone they want to the other nodes in the network.

Call
Sun on
This

Apollo does provide security safeguards. Although in theory it is possible to masquerade as another user on an Apollo system, it requires more than a simple file edit to accomplish such a masquerade.

NFS DOES NOT PROVIDE FILE AND RECORD LOCKING. In NFS no reference counts are kept, and no state is kept on the server node. While one advantage to such a system is that it is robust in face of network and node failures (you can not have inconsistent states if the state is not maintained). However, with such a system it is impossible to discover not only who has a file open, but also if the file is in use. For example, if a user on node A, opens a file on node B, and another user on node C deletes the file opened on node B, the file "goes away". The user on node A, still believing they have an open file, will get an error the next time they attempt to read or write to the file.

This is a seriou
drawback in a
software devel
environment.
Call Sun

FILE BLOCKS ARE CACHED BOTH AT THE CLIENT AND AT THE SERVER NODES. The caches are flushed every 30 seconds, and this can result in inconsistencies in file data. For example, if the same file is open for reading and writing on different nodes, and the user at node A writes into the file, the user at node B will not see the new data for up to one minute after it is written at node A. The result could be

-- ~~many~~ uirrerent versions of the data as there are nodes on the network.

Another drawba in software development environments

> Sun may state that this is a feature of NFS, and that any applications requiring atomicity and consistency of data can build their own mechanisms on top of the file system. While Berkeley 4.2 UNIX does provide a system call to lock a file, it reportedly can not work across the network because it implies state on the server end.

Call Sun

NFS CAN NOT GUARANTEE ATOMIC WRITE OPERATIONS. Write operations that cross a page boundary reportedly can not be guaranteed because the pages in the operation might become intermixed with pages from another node on the network. The pages actually written to disk would be a mixture of those from all the nodes trying to write to the file at approximately the same time (within one minute).

WHEN WILL NFS BE AVAILABLE AND HOW WILL IT BE PRICED?

The first availability of NFS-based product is scheduled for April 1985.

Apollo does this today & has for 5 years.

> While in the past Sun has announced "vapor products", to use a standard industry expression. NFS is a real product; the first public demonstration of the heterogeneous applicability of NFS was shown at the January 1985 Usenix conference, in Dallas, when Sun, Gould Inc. and Pyramid Technology Corp. were transparently sharing files across a standard Ethernet with TCP/IP. Sun was also running a VAX in their booth as part of the NFS network.

NFS will be shipped routinely with all workstations. Upgrades to existing customers will be made free.

WHAT IS THE BOTTOM LINE?

NFS fills a major deficiency in the Sun product offering, but as the preceeding discussion shows, this product still has many limitations, including lower performance and inconsistent data sets.

Sun is attempting to position NFS as an industry standard, but there is nothing to indicate that other vendors will indeed adopt the NFS. Apollo does support industry standards and more -- Apollo provides capabilities beyond what the current standards address.

Sun is generating a lot of "press" on this "new" distributed file system. However, Apollo had this "new" capability, these "extensions to UNIX, since 1982.

Refer to the LOW-END SYSTEMS SUN COMPETITIVE NOTE (December 1984) for detail on the entire Sun Microsystem product family.

ACKNOWLEDGEMENTS to Jim Rees for information and insights.