Windows NT Server Consolidation and Performance Improvements with Solaris™ PC NetLink 2.0 Software

Don DeVitt

Sun BluePrints™ OnLine—August 2002

http://www.sun.com/blueprints

Sun Microsystems, Inc.
4150 Network Circle
Santa Clara, CA 95045 U.S.A.
650 960-1300

Part No. 816-7445-10
Revision 1.0, mm/dd/02
Edition: August 2002

Please
Recycle

Windows NT Server Consolidation and Performance Improvement with Solaris PC NetLink 2.0 Software

The trend toward increasing value and reducing cost in the data center is more evident in our current economy. Means of reducing cost include supporting as few servers as possible and switching to technology with a proven high level of availability and increased capacity that requires less system management.

Solaris™ PC NetLink 2.0 software offers many enhancements over previous versions with respect to performance, sizing, and availability. These product improvements allow Solaris PC NetLink software to be adopted as alternatives to Windows NT servers.

This paper focuses on best practices to help data center managers meet cost reduction goals with respect to supporting multiple users accessing servers through PC clients. It also discusses the new features in Solaris PC NetLink 2.0 that improve performance and add new options for consolidating Windows NT servers. Several new configurations will be shown and performance, sizing, and availability implications will be discussed.

Background

Solaris PC NetLink software is a Sun Microsystems product that brings the file and print functionality of Windows NT 4.0 domains to Solaris servers. The product originated with Windows NT 4.0 source code, available through a port of ATT's ASU product. Sun engineers have enhanced the product considerably since its initial release in 1999.

PC NetLink software implements many of the services offered by Windows NT 4.0 servers. Among these services are the following:

- File and print functionality supported through common internet file system (CIFS) SMB protocol
- Windows NT Primary Domain Controller (PDC), Backup Domain Controller (BDC), or member server capability.
- Windows Internet Name Service (WINS)
- Microsoft Remote Procedure Calls (RPCs)
- Windows NT Directory Services (NTDS)
- Windows NT Registry
- Windows NT 4.0 style Event Logging
- Computer Browser Service support. Solaris PC NetLink can act as the master browser for a subnet.

In addition to these network services, Windows NT 4.0 server system administration tools are provided. Window-based PC tools, such as User Manager for Domains and Server Manager can also be used to manage the Windows NT side of Solaris PC NetLink.

Solaris PC NetLink Versions

Solaris PC NetLink was released in three versions (1.0, 1.1, 1.2) prior to Solaris PC NetLink 2.0. In summary, each release consisted of:

- Version 1.0 - Initial release of the product.
- Version 1.1 - Bug Fixes
- Version 1.2 - Enhanced stability, Sun™ Cluster 2.2 support, member server support, support to allow backups to save files with ACLs together through Legato Systems, Inc., concurrent connections increased to 2000, support for Windows 2000 mixed mode support

Solaris PC NetLink 2.0 includes several new features that improve performance, and significantly enhance the cluster support.

These new features include the ability to run up to ten Solaris PC NetLink virtual servers on each Solaris physical host. This change created the following improvements:

- Better reliability (fewer single points of failure)
- Easier migration and consolidation from Windows NT to Solaris software
- Better scalability when consolidating multiple Windows NT servers on a Solaris operating environment server
- More effective use of equipment when Solaris PC NetLink is used in a high availability (HA) cluster
- Enhanced monitoring of Solaris PC NetLink virtual servers in an HA cluster
- Improved ACL database scalability

- Enhanced password synchronization, enabling a single password change to affect a user's accounts on both the Solaris operating environment and Microsoft Windows NT systems
- LDAP Directory Synchronization through utilities that transfer user account data between the Microsoft Windows NT SAM and an LDAP directory server
- Enhanced automounter support
- Support for domain name service (DNS) in NetBIOS.
- DOS file attributes can be stored using mechanisms other than UNIX® group IDs
- Improved Solaris PC NetLink server manager graphical user interface (GUI)

Solaris PC NetLink Performance Improvements

Several product changes have resulted in better performance and scaling. In addition to changes to Solaris PC NetLink, new enhancements to the Solaris operating environment have equally contributed to better server performance. Customers upgrading from older versions of Solaris PC NetLink should also upgrade to the latest version of Solaris.

Windows 95/98/ME Clients Versus Windows NT 4.0 Clients

Windows 95/98/ME clients use a different level of the CIFS protocol to access Solaris PC NetLink than Windows NT workstation clients. When the client operating system is requested to access a file for an application, the local operating system running on the client must translate the file request into network protocols that are sent to the server. In Windows 95/98/ME systems, the requests are significantly different from those made by Windows NT 4.0 workstation clients. In previous versions of Solaris PC NetLink software, a limitation with the Windows 95/98/ME CIFS protocol stack forced multiple ~10ms waits to occur during the file write operations. When moving large files, these delays could add up and cause slower performance. This unnecessary wait time is now eliminated.

Individual Windows 95/98/ME clients now experience up to double the performance with PC NetLink 2.0 software than they experienced using previous versions of Solaris PC NetLink software. A patch is available for Solaris PC NetLink 1.2 that incorporates this improvement.

This change only affects Windows 95/98/ME client performance and is not seen in Windows NT 4.0 workstation clients because the protocol used there does not cause the same delay. While individual performance for these clients may improve, the overall server performance for hundreds or thousands of users is unaffected by this change.

Access Control List Support

Support for Windows NT 4.0 ACLs has always been an important feature of the Solaris PC NetLink software. ACLs on Solaris software are compliant with the POSIX 1003.6 specification. They have been implemented for both the UNIX file system (UFS), as well as for NFS versions 2 and 3. Unfortunately, there is not a one-for-one mapping of Windows NT style ACLs with the POSIX standard. For this reason, ACLs were initially supported by a separate database that the Solaris PC NetLink 1.0, 1.1, and 1.2 software maintains in parallel with the normal file data access. This ACL database is stored in what is called a binary large object (BLOB) database.

By default, for every directory created by a PC client application through the Solaris PC NetLink software, an ACL entry is saved in the BLOB database. File ACLs are not created by default, but are created when the PC client explicitly requests to do so. While supporting ACLs through this mechanism works well, performance can suffer for applications that conduct numerous directory operations or use file ACLs heavily. Most directory operations require access to the database file, usually located in `/var/opt/lanman/datafiles/acl`. The ACL database file access may originate from the file read cache located in physical memory. On a busy server, this file read cache may be displaced from physical memory by other applications, which can potentially reduce throughput.

Another problem associated with supporting a parallel database is that the ACLs are not accessible by traditional UNIX backup applications such as `tar` or `dump`. When a Solaris PC NetLink created file is accessed by a traditional UNIX based application such as `tar`, only the file data and attributes in the UFS file system are saved. The `tar` program has no knowledge of how to access and save the Windows NT style ACL.

Solaris PC NetLink 1.2 software introduced a software module that allows Legato Systems, Inc. backup software, and other backup programs, to save ACLs with the files. While this helped solve the problem of saving the ACLs with the file, it was still limited to just a few backup programs.

Solaris PC NetLink 2.0 software supports Windows NT style ACLs differently. Instead of using a separate database to store the ACLs, it saves the ACL data in a hidden UNIX dot (`.`) file located in the referring directory. While this does not completely eliminate overhead for storing ACLs, it does limit the impact of ACL support on a full server.

If a significant number of files are being accessed in one directory, the Solaris file read cache will keep the ACL dot (.) file in memory, making it easily accessible. These files are typically small compared to the previously used BLOB file and are distributed throughout the file system.

Previously, a hyperactive ACL BLOB file could potentially cause all users to suffer a performance degradation. With the new architecture, any potential performance degradation is limited to users using the same active disk subsystem.

In addition to performance issues, the elimination of the ACL BLOB file removed the susceptibility to loss of all ACL information if the database file became corrupted.

Windows NT Server Consolidation

Solaris PC NetLink has always had the ability to consolidate several Windows NT 4.0 servers within the same NT domain into one Solaris PC NetLink server. The following example illustrates this kind of consolidation:

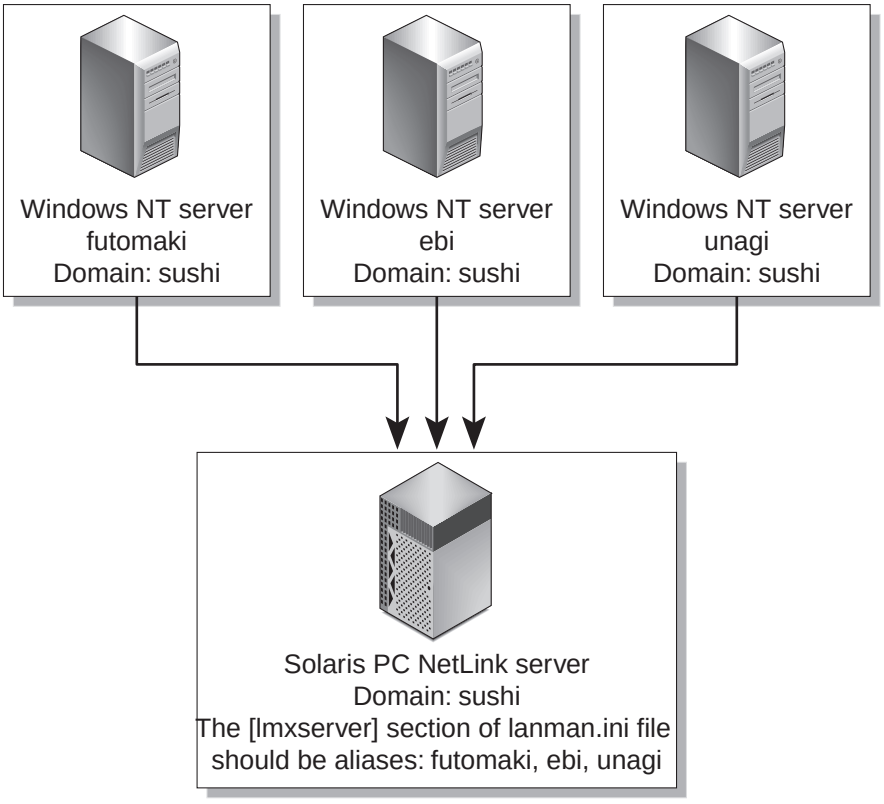

FIGURE 1 `/etc/opt/lanman/lanman.ini` Server Aliases

In FIGURE 1, all three Windows NT servers names: `futomaki`, `ebi`, and `unagi` are consolidated into one Solaris PC NetLink server. Only one edit to the `/etc/opt/lanman/lanman.ini` file will accomplish the task. Add the line `aliases=futomaki,ebi,unagi` to the `[lmxserver]` section of the file. Once this is done, the Solaris PC NetLink server will respond to all three NetBIOS names. The PC clients that were using the Windows NT servers need not be changed in any way. However, a reboot at transition time may be needed to access the new server.

While this is an excellent way to consolidate servers that reside in one Windows NT domain, it does not allow consolidation of servers that are in multiple Windows NT domains. Prior to Solaris PC NetLink 2.0, the only way to accomplish multi-domain consolidations in one server was by use of Sun Enterprise™ hardware domains, as illustrated in FIGURE 2:

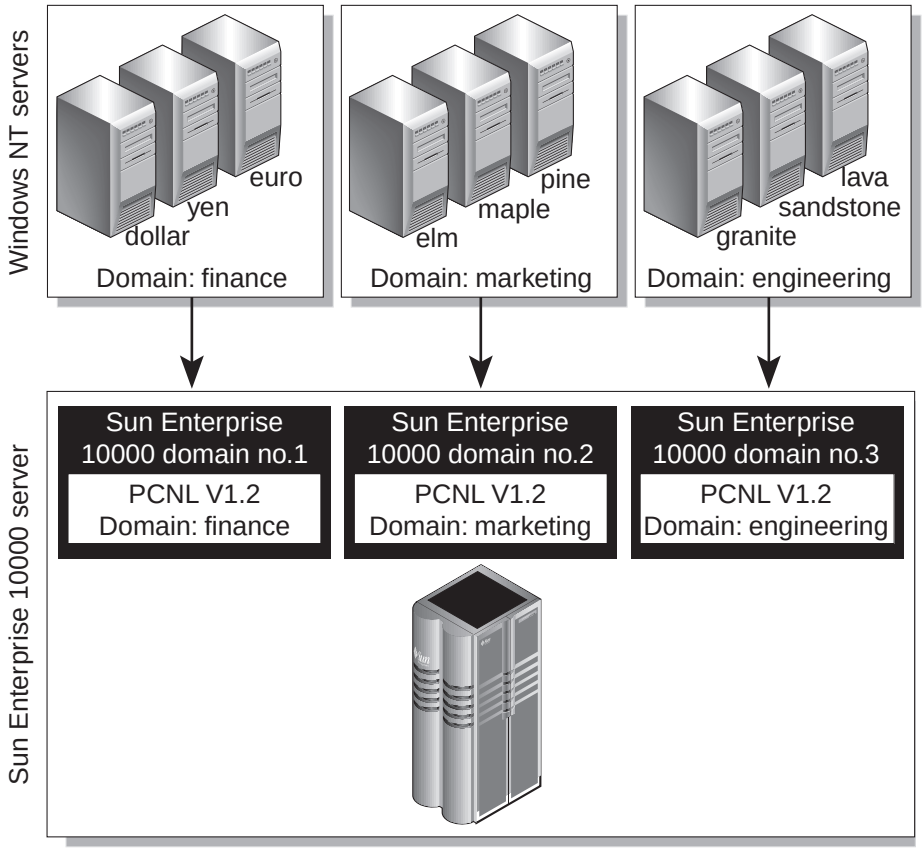

FIGURE 2 Sun Enterprise 10000 Hardware Domain Used to Consolidate NT Servers

In FIGURE 2, each Sun Enterprise 10000 domain (up to 16 per server) can run a copy of Solaris PC NetLink. Each copy of Solaris PC NetLink, in turn, supports a separate Windows NT domain.

Multiple Instance Support

Solaris PC NetLink 2.0 can support multiple instances of Solaris PC NetLink running simultaneously on one server. This new functionality allows Solaris PC NetLink to support multidomain Windows NT consolidations even on the smallest Sun Enterprise server.

Using the example shown in FIGURE 2 with a Sun Fire V480 server instead of a Sun Enterprise 10000 server produces the solution illustrated in FIGURE 3.

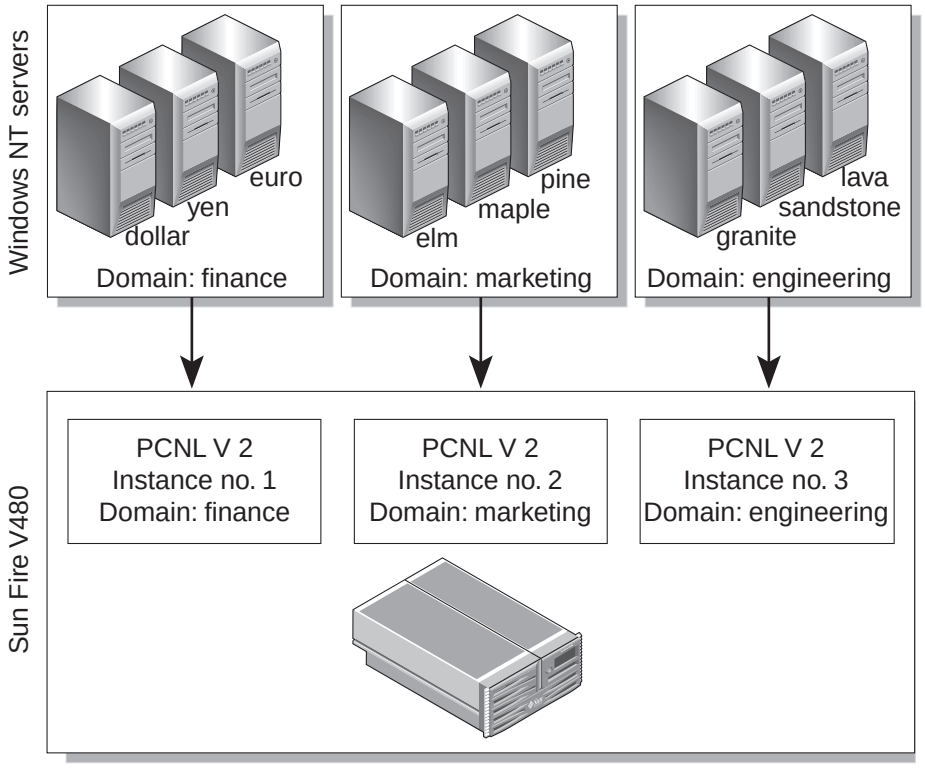

FIGURE 3 Multiple Windows NT Domains Supported on One Sun Fire V480 Server

In FIGURE 3, all three Windows NT domains have been consolidated onto one Sun Fire V480 server. The Sun Fire V480 server supports up to four processors and runs one instance of the Solaris operating environment. Sizing, performance, and availability options should be considered fully before proceeding into a full-scale consolidation of multiple domains.

It may be difficult to determine exactly how heavily Windows NT servers are being used, especially during peak periods. With this new multi-instance functionality of Solaris PC NetLink 2.0, a customer might decide to move Windows NT domains to the new server one by one. After each transition, performance and resource consumption should be monitored during peak use periods to determine if further consolidation is possible on the new server. Solaris servers that have under-utilized resources would continue to be candidates for additional Windows NT server consolidation.

Support for multiple instances of Solaris PC NetLink not only increases the number of ways used to consolidate Windows NT servers, but also plays an extremely important role in increasing the efficiency of servers used in HA clustered solutions.

High Availability and Clustering

Support for Sun Cluster 2.2 was introduced with Solaris PC NetLink 1.2. This support allows the creation of a *virtual* server. This *virtual* server would be supported by a *cluster* of physical servers or *nodes.* The node that manages the resources in the cluster is called the *default master node*. The other nodes in the cluster are called backup nodes. When Solaris PC NetLink is interfaced with cluster *data services* it can become part of an highly available application that can move from node to node.

If a copy of Solaris PC NetLink running on one node in the cluster experiences a problem, Sun Cluster detects the problem and it relocates the application to a backup node. This process of relocating the application service from a failing node to a surviving node is called *failover.* The Sun Cluster must encapsulate all the information associated with the data service into a *logical host.* All the client accounts, IP addresses, data management structures, and processes associated with the data service need to move across nodes without modification for the failover to work.

Unlike Solaris PC NetLink 2.0, Solaris PC NetLink 1.2 only allows one instance of the product to run on any one server.A logical host is defined an IP address (or a hostname) associated with an application service that can be hosted by an active cluster node. The file structures, exported by Solaris PC Netlink, are supported by disk storage mirroring which is a basic Sun Cluster requirement to support availability

In cluster environments of two to three nodes, only one logical host can be created, which can support only one Windows NT domain. FIGURE 4 shows a typical cluster with Solaris PC Netlink 1.2.

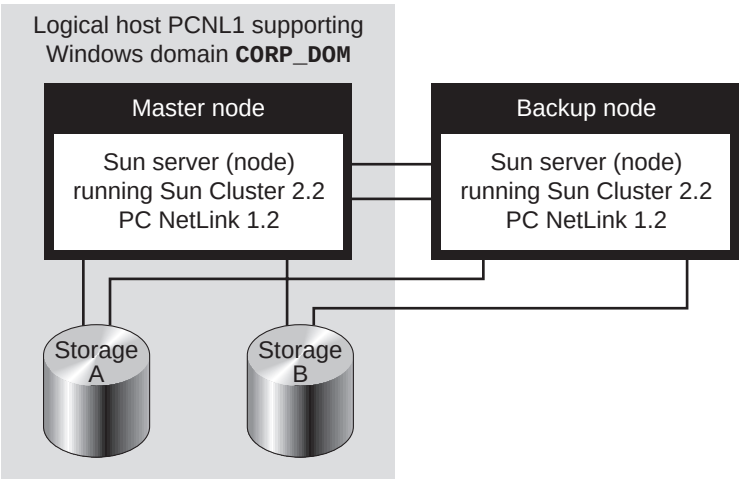

FIGURE 4 Sun Cluster 2.2 Support With Solaris PC Netlink 1.2 Software

In this configuration, both the master and backup nodes of the cluster have direct access to *Storage A* and *Storage B*. Both *Storage A* and *Storage B* are mirrored volumes to offer High Availability. Under normal operation, the master node is the active side of this two-node cluster and supports the logical host `PCNL1`. If a problem occurs on the master node, the failover occurs, and the backup node takes over the role of supporting the local host PCNL1. FIGURE 5 illustrates the result of this failover.

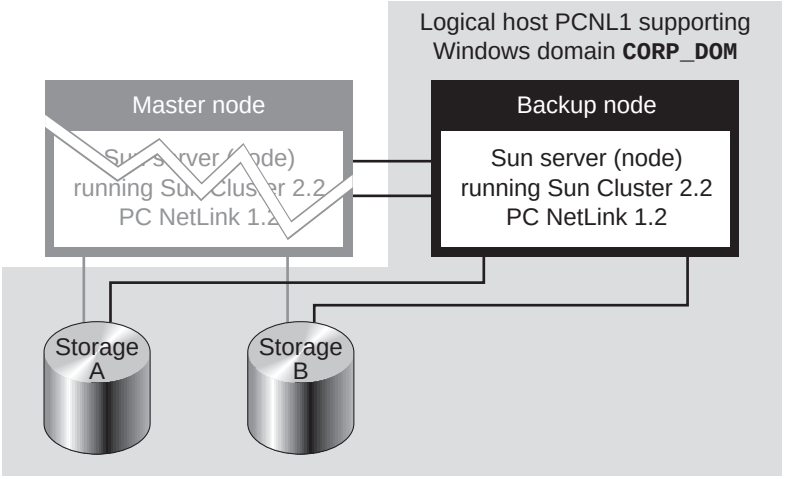

FIGURE 5 Standard Failover With Solaris PC NetLink 1.2 Software

This configuration has a high level of availability, but has several limitations. Remember, Solaris PC NetLink 1.2 supports only one instance of the software running on any node. Likewise, the cluster can only support one of Windows NT domain . A more significant limitation is that only one of the nodes actively supports the Solaris PC NetLink data server for the cluster during normal or failover operation. This means that other nodes in the cluster are completely inactive with respect to Solaris PC NetLink operation. If other services, such as email or a web application, are required in this HA environment, the other node could be actively supporting these services, more fully utilizing the available resources.

Solaris PC NetLink 2.0 allows multiple instances of the software to run simultaneously on the same node of a cluster. This allows many Windows NT domains to be supported on a cluster simultaneously. In addition, all the various nodes of the cluster can actively support different domains. This allows all resources to be utilized under normal Solaris PC NetLink operation. FIGURE 6 shows a simple example of such a configuration.

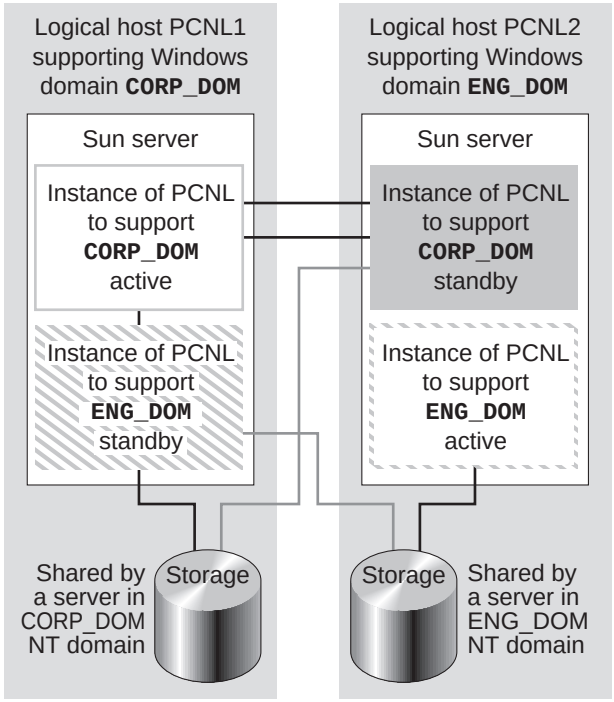

FIGURE 6 Solaris PC NetLink 2.0 Multiple Instance Support and Clustering

With the configuration shown in FIGURE 6, Windows NT servers in both Windows NT domains `CORP_DOM` and `ENG_DOM` can be consolidated into a highly available clustered environment. If a failure occurs in this cluster, either node of the configuration is capable of supporting all the servers involved in the consolidation

for both Windows NT domains If the instance of Solaris PC NetLink software supporting the logical host PCNL1 and the Windows NT domain CORP_DOM failed for a software or hardware reason, the result after the failover is shown in FIGURE 7.

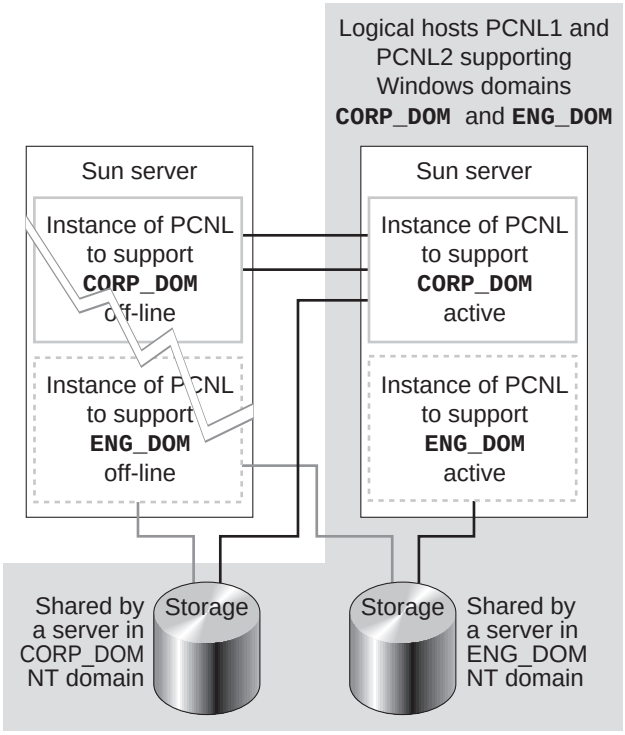

FIGURE 7 Multiple Instance Support After a Node Failure

FIGURE 7 shows the right node supporting both the CORP_DOM and the ENG_DOM domains through the logical hosts PCNL1 and PCNL2 simultaneously.

This example is very simple. A consolidation of a real Windows NT environment could involve many Windows NT servers, from several Windows NT domains all going to multiple instances of Solaris PC NetLink, to several Sun servers all involved in a large cluster configuration.

One goal of any Windows NT server consolidation is to minimize or eliminate any system administration of the multitude of Windows clients that use the servers. As support for each Windows NT server is brought into the cluster of Solaris PC NetLink logical hosts, the PC clients that routinely use the server would not even be aware that the change has taken place.

Performing a large Windows NT server consolidation is best done in phases, as each Windows NT server is replaced by clustered Solaris PC NetLink configuration.

Again Solaris PC NetLink 2.0 allows a more efficient utilization of resources in the cluster to be utilized under normal operation. Careful thought is required to determine the acceptable failover performance requirements during peak use periods. All failover states should be analyzed to determine acceptable failover performance. While Solaris PC NetLink will continue to operate properly even when the server is fully saturated, latency times will continue to grow as the load is increased.

For most installations, nodes expected to act as backup nodes for clusters should have excess resource capacity in the form of additional CPUs and memory to support a failover as required. Except on very busy peak periods, most users of the servers will not notice the difference between normal and failover performance. Typical productivity applications may take a little longer to read and file to the server.

Other Situations When a New Instance of Solaris PC NetLink Is Needed

A new instance of Solaris PC NetLink, running on the same server, is needed whenever you want to support a new Windows NT domain. Other situations where you will want to create a new instance of Solaris PC NetLink are as follows:

- **When Windows NT exported file systems (shares) from different servers have the same path name**—During a Windows NT server consolidation, it is likely that the shares from different Windows NT servers will have the same path name(For example, `\\server1\cdrive` and `\\server2\cdrive`). Unless you are willing to juggle directories within the shares to avoid collisions, it may be best to create a separate instance of Solaris PC NetLink to support some servers separately.
- **When there are security issues**—If, during a consolidation, it is determined that the shares of one server are protected differently than those from others, then a new instance of Solaris of PC NetLink may be warranted. For example, users of one server may not want other users seeing their shares.
- **When the concurrent user limit is reached**—Different versions of Solaris PC NetLink have different limits for the maximum number of concurrent connections they will support. When this limit is reached, the PC connections can be distributed across the multiple instances of the software.
- **When quality of service is an issue**—Multiple instance support can help manage the resources consumed by a group of users. Solaris Resource Manager (SRM), or some other resource management technique, can ensure that VIP users or mission critical applications have a higher priority than other activity on the server.

Improvements in Performance Elsewhere

Both Sun Enterprise servers and the Solaris operating environment have improved the ability of Solaris PC NetLink to perform better and scale further. Benchmarks done with Solaris PC NetLink 1.1 running on the older Solaris 2.5 operating environment, compared with those done on the current Solaris 8 and 9 operating environments, show that even this older versions of Solaris PC NetLink can scale further with the new operating environments. This improved performance is mostly due to performance improvements in the file system and the memory manager as well as the drivers for supporting the network interfaces.

About the Author

Don is currently a Senior Staff Engineer and a PC interoperability specialist and performance specialist within the Enterprise Engineering group. Don has been on the development teams of almost every software and hardware PC interoperability product that Sun Microsystems has produced over the past 15 years. Don is now working on Reference Architectures.

Don started his career as an electrical engineer and worked in the Automated Test industry (Teradyne Inc.) and PC operating system market (Digital Research from CP/M fame) before coming to Sun.

References

For more in depth information on Solaris PC NetLink, refer to the Blueprint book *Solaris™ PC NetLink Performance, Sizing, and Deployment*, by Don DeVitt, Published by Prentice Hall, June 2000, ISBN: 0130266868.

The Solaris PC NetLink Version 2.0 can be downloaded from:
`http://www.sun.com/solutions/interoperability/netlink/download-20.html`

For more online information of Solaris PC NetLink refer to:
`http://www.sun.com/solutions/interoperability/netlink`

For online documentation of Solaris PC NetLink refer to: `http://docs.sun.com`
Search for “PC NetLink” select *Solaris PC NetLink 2.0 Administration Guide*. Select *Solaris PC NetLink 2.0 High Availability Guide* for the manual that covers clustered configurations.

Many additional articles about Solaris PC Netlink are available at:
`http://www.sun.com/software/solutions/blueprints/browsesubject.html`

The Solaris PCNL sizing tool can be found at:
`http://www.sun.com/software/solutions/blueprints/tools`

Additional papers and Blueprint articles about Solaris PC NetLink can be found at:
`http://www.sun.com/software/solutions/blueprints/browsesubject.html`

Ordering Sun Documents

The SunDocs℠ program provides more than 250 manuals from Sun Microsystems, Inc. If you live in the United States, Canada, Europe, or Japan, you can purchase documentation sets or individual manuals through this program.

Accessing Sun Documentation Online

The `docs.sun.com` web site enables you to access Sun technical documentation online. You can browse the `docs.sun.com` archives or search for a specific book title or subject. The URL is `http://docs.sun.com/`

To reference Sun BluePrints OnLine articles, visit the Sun BluePrints OnLine Web site at: `http://www.sun.com/blueprints/online.html`