Using Solaris Resource Manager™ Software with Solaris™ PC NetLink Software - Part 1

By Don DeVitt - Enterprise Engineering

Sun BluePrints™ OnLine - May 2000

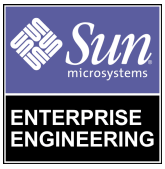

http://www.sun.com/blueprints

Sun Microsystems, Inc.
901 San Antonio Road
Palo Alto, CA 94303 USA
650 960-1300 fax 650 969-9131

Part No.: 806-5581-10
Revision 01, May 2000

Please
Recycle

Using Solaris Resource Manager™ Software with Solaris™ PC NetLink Software - Part 1

The topic of resource management is not an easy topic to cover in one, or even a two part article. A full Sun BluePrints™ book on resource management is available (see references) and covers the topic to reasonable depth. Here we will focus on resource management as it relates to Solaris™ PC NetLink software.

In Part 1 of this article I briefly list the most likely tools and techniques that can be used with Solaris PC NetLink software to participate in a resource management function. Part 1 will also show you how to implement a server that incorporates Solaris PC NetLink software within the Solaris Resource Manager™ software scheduler.

In Part 2 of this article I will show the results of a lab exercise that will help guide you in determining how many resources, in the form of Solaris Resource Manager software “shares”, Solaris PC NetLink software might need for a particular server environment.

Server Consolidations

It is now clear that the trend to consolidate multiple servers into one larger server is an effective way to reduce cost of ownership and reduce the complexity of server maintenance.

Single function consolidation - If all the servers that are to be consolidated are performing the same function (file and print in Solaris PC NetLink software’s case) resource management issues can boil down to making sure the consolidated server

has enough resources to support the combined needs of the prior servers during peek periods. In these situations the goal is to use the available tools and techniques to tune one application.

Multiple function consolidation - An alternate form of server consolidation is to combine the functions of several servers performing different roles. Web services, email services, database services, and file & print services can be all supported on the same server. Managing the resources in this kind of environment requires tools that control all the applications on the server. It is this form of consolidation that I will focus on.

Solaris PC NetLink Software's Role on the Server

With regard to the Solaris PC NetLink software, you will usually approach resource management in two ways:

- **Solaris PC NetLink software performance is the primary role of the server**. In these situations allocating resources to support Solaris PC NetLink software have a higher or at least equal priority to other services being supported on the server. For example, if you are using the Solaris PC NetLink software as a home directory server for a user community, you want to be sure other services and applications running on the server do not dominate the use of the server. You will want to make sure that when the server is busy supporting more than one application or service, there are at least a minimum amount of resources that will be assigned to the role of supporting Solaris PC NetLink software.
- **Solaris PC NetLink software performance is of secondary importance on the server.** In this case, Solaris PC NetLink is one of several services on the server and it's role is of a lower priority to that of the other functions the server must support. For example, if the Solaris PC NetLink software has been installed on a database server to offer reports to users on PC clients, it must be allocated only a small amount of resources so that it does not interfere with the higher-priority database applications on the system. Here, Solaris PC NetLink software will be tolerated only as long as there are resource management tools that will limit the resource Solaris PC NetLinksoftware is allowed to use during high demand periods.

In either case, there is a need for resource management tools and techniques to control the Solaris PC NetLink services and the other applications running on the server. You can also use a combination of tools to accomplish what you want.

Possible Resource Management Options with Solaris PC NetLink Software

The Solaris Operating Environment software has a variety of resource management tools and techniques that can work effectively with Solaris PC NetLink software and the other applications and services that it is expected to share a server with. Following is a list of resource management tools and techniques that can help control Solaris PC NetLink software resource consumption. Some are not sophisticated, but they are easy to implement and can be very useful for special uses of Solaris PC NetLink software (see references for sources of more detailed information).

- **Solaris PC NetLink software internal settings -** Solaris PC NetLink software has several internal variables that control how Solaris PC NetLink software operates. These internal variables can help a system manager control how many resources the Solaris PC NetLink software is allowed to use. The two most useful parameters controlled by these variables are:
  - The maximum number of users allowed to be supported simultaneously - Controlling this limit can help set the maximum impact of Solaris PC NetLink software on a server.
  - The maximum number of processes that are spawned to support a user community - This limit can help control the maximum number of processes that Solaris PC NetLink software will spawn. Each process requires CPU and memory resources.
- **Setting Hardware limits** - A reasonably easy mechanism for limiting Solaris PC NetLink software resources is to limit what disk subsystems it will be sharing. In situations where Solaris PC NetLink software file performance is a very low priority, a slow disk subsystem can be used to limit the resources used by Solaris PC NetLink software. A slow disk subsystem will quickly become the bottleneck forcing Solaris PC NetLink software processes into IO wait states releasing the processors for higher priority processes. While this technique may be acceptable for typical read/write environments, it does not work effectively in read only environments because Solaris Operating Environment will use available memory as a cache for read requests. In R/O situations free memory will be used as read cache which allows read requests to bypass the disks. This will speed up operations, and allow more CPU resources to be consumed.
- **Base Solaris Operating Environment Software -** The tools and commands that are included with Solaris Operating Environment can be used to control Solaris PC NetLink software processes. The `nice(1)` command allows you to control the priority of individual processes. This command is best for diagnosing and solving individual CPU resource problems with Solaris PC NetLink software. However, it would be difficult to use this command as a general solution.
- **Processors sets -** Processor sets were introduced in the Solaris 2.6 Operating Environment. This feature enables you to divide multiprocessor systems into logical groups and permits users to launch processes into those groups. Processor

sets may be good when you wish to dedicate exclusive use of server processors to a service, but it is not best for applying idle resources to a service that is highly active.

- **Sun Enterprise™ 10000 Server and Dynamic System Domains -** The Sun Enterprise 10000 server has a feature called dynamic system domains (DSDs), which enables you to logically divide a single system rack into one or more independent systems, each running its own copy of Solaris Operating Environment software. (Do not confuse Sun Enterprise 10000 server DSDs with Windows NT domains.) With DSDs, a domain could be defined to support the Solaris PC NetLink software and the number of processors can be moved in and out of the domain as the requirements of the domain are dictated.
- **Solaris™ Bandwidth Manager Software -** Solaris Bandwidth Manager software provides the means to manage network resources to provide Quality of Service (QoS) to network users. Solaris Bandwidth Manager software allows network traffic to be allocated to separate classes of service so urgent traffic gets higher priority over less important traffic. Different classes of service can be guaranteed a portion of the network bandwidth, leading to more predictable network loads and overall system behavior. You can control the amount of resources consumed by the Solaris PC NetLink software file and print operations indirectly by throttling the amount of network bandwidth on TCP port 139, which is the NetBIOS Session socket. Using Solaris Bandwidth Manager to control Solaris PC NetLink may be the best solution when network bandwidth is the limiting resource.
- **Solaris Resource Manager Software -** Solaris Resource Manager software is the Sun™ resource management extension for the Solaris Operating Environment. Solaris Resource Manager software provides control of the following system resources:
  - CPU (rate of processor) usage
  - virtual memory and number of processes
  - number of log ins
  - terminal connect-time

  It allows the system manager to define how resources will be allocated when a variety of services are to be supported on a busy server. Unlike other techniques which may dedicate specific resources to a particular service, thus disallowing their use by other services, it allows resources to be allocated to busy services if other services under its control are not active. This allows a better utilization of system resources. Perhaps the most important reason for using Solaris Resource Manager software on a consolidated server is the fact that all the services can be controlled by the same mechanism allowing the system manager to change the allocation of resources dynamically under one controlling environment.

The Solaris Resource Manager software is the only resource management tool or technique that provides the following important features to solve server consolidation problems

- Ability to control most services in a consistent and equal way
- Better utilization of system resources
- Dynamic control of system resources
- More flexible resource allocation policies
- Finer-grained control over resources
- Decayed usage of resources
- Accounting data for resource usage

Except in special or limited cases, Solaris Resource Manager software is the best mechanism to support resource allocation, for multiple services running on large servers. The rest of this two part article will focus on deploying Solaris PC NetLink software within the Solaris Resource Manager software's control.

Solaris Resource Manager Software

Before diving into how Solaris PC NetLink software and Solaris Resource Manager software work together, it is important to understand the basics of how Solaris Resource Manager software operates. If you are not familiar with the basics of Solaris Resource Manager operation, there is aSun BluePrints OnLine article published April 1999 titled Solaris Resourse Manager available at `http://www.sun.com/software/solutions/blueprints/0499/solaris1.pdf` that will give you a quick description of the product. If you don't know what a Solaris Resource Manager lnode hierarchy is, this 10 page article will give you what you need to bootstrap your knowledge on Solaris Resource Manager product. More complete documentation exists on Sun's `http://docs.sun.com` web site. Search for "Solaris Resource Manager 1.1 System Administration Guide"

Controlling Solaris PC NetLink Software Processes Within Solaris Resource Manager

One secret for including any application within the Solaris Resource Management software hierarchy is to find all the places where processes are launched at boot time, or when the application is first executed. Once all of these locations are determined, you use the Solaris Resource Management `smuser` command to launch the process instead of the default mechanism. Once this is done, the process, and any processes that are forked are controlled by the SRM scheduler.

Fortunately the Solaris PC Netlink environment is spawned from the execution of one process called `lmx.ctrl` which is started at boot time. The path to the script that starts the process is `/etc/init.d/ms_srv` .The edits needed to the file are shown below in the step by step procedure.

Solaris Resource Manager software also needs an `/etc/passwd` entry to define the lnode which Solaris PC NetLink software will run. Solaris PC NetLink software already defines `lmxadmin` in `/etc/passwd` and it can be used to define the lnode. Even though `/etc/passwd` entries are used to define lnodes, Solaris PC NetLink software will continue to run as `root` during operation.

Steps to incorporate Solaris PC NetLink Software Processes Within Solaris Resource Manager

It is assumed both Solaris PC NetLink software and Solaris Resource Manager software are installed and running on the server.

1. **Define or use the already established `/etc/passwd` entries that will be used by Solaris Resource Manager software to define lnodes.**

   The Solaris PC NetLink software defined `lmxadmin` entry already in the `/etc/passwd` file, which is fine for the purpose of Solaris Resource Manager software lnode hierarchy definitions.

2. **Use the `/usr/srm/sbin/limadm` command to place the `lmxadmin` user in the appropriate spot in the Solaris Resource Manager hierarchy.**

   This will be different for each server configuration. Try to place the Solaris PC NetLink software at the same level as other important applications you want to control. The amount of shares you give the Solaris PC NetLink software will depend on server (CPUs and memory), the number of users, and how busy the users are.

3. **Stop the Solaris PC NetLink software execution as root, using the following command:**

```
/opt/lanman/bin/net stop server
```

4. **In the** `/etc/init.d` **directory, edit the** ms_srv **start up script to use the** `/usr/srm/bin/srmuser` **command to start the Solaris PC NetLink software processes.**

   Note that the process will still execute as root, but will be attached to the lnode `lmxadmin`. All Solaris PC NetLink software commands will use the `/etc/init.d/ms_srv` script to start or stop the Solaris PC NetLink services..

```
Original /etc/init.d/ms_srv code segment:
. . .
else
    cd `dirname $LMXCTRL_PATH` && $LMXCTRL_NAME < $DEV_NULL_PATH
fi
. . .

New /etc/init.d/ms_srv code segment:
. . .
else
    cd `dirname $LMXCTRL_PATH` && /usr/srm/bin/srmuser lmxadmin \
    $LMXCTRL_NAME < $DEV_NULL_PATH
fi
. . .
```

   Hint: When editing the file, search for `LMXCTRL` to get to the correct location.

   All Solaris PC NetLink `lmx.*` processes are forked by `lmx.ctrl` which, after the edit, will be under the Solaris Resource Manager software control. Starting and stopping processes from server manager does *not* break the Solaris Resource Manager control. This addition to the Solaris PC NetLink software startup command line will start the lmx.ctrl process running as root, using the lnode of lmxadmin.

5. **Restart the Solaris PC NetLink server using the** `/opt/lanman/bin/net start server` **command.** This command will in turn execute the /etc/init.d startup script you just edited. The Solaris PC NetLink software will now be scheduled by way of the Solaris Resource Manager scheduling class.

6. **Determine the best way to allocate shares for the Solaris PC NetLink software processes.**

   This will be different for each application and work load mix. The data presented in the part II of this article should help you determine starting values in determining the appropriate shares to allocate to the Solaris PC NetLink software.

   Remember that in the previous procedure users refer to Solaris Operating Environment users as defined by `/etc/passwd` or NIS. As the Solaris PC NetLink software supports requests for a user, it does so with an `lms.srv` process that runs as root. There is no easy way to control the Solaris PC NetLink software with Solaris Resource Manager to allow individual users to have more resources allocated to them over other users using the Solaris PC NetLink software.

Part II of this article

In part II of this article (June 2000) I will use loading benchmarks to help show what performance and throughput you should expect when Solaris PC NetLink software is given a particular percentage of the system.

Conclusion

Part 1 of this article listed and defined the most likely tools and techniques that can be used with Solaris PC NetLink software to participate in a resource management function. This Part 1 also explained how system managers can implement a server that incorporates Solaris PC NetLink software within Solaris Resource Manager Software's scheduler.

References

For additional, detailed information on Solaris PC NetLink software, refer to the Sun BluePrints book, *Solaris™ PC NetLink Software: Performance, Sizing, and Deployment*, (ISBN# 0-13-026686-8) which is scheduled for publication by Prentice-Hall in the May of 2000 and is expected to be available through www.sun.com/books, amazon.com, fatbrain.com, and Barnes & Noble bookstores.

For additional, detailed information on resource management, refer to the Sun BluePrints book, *Resourse Management*, (ISBN# 0-13-025855-5) published by Prentice-Hall which is available through www.sun.com/books, amazon.com, fatbrain.com, and Barnes & Noble bookstores.

Solaris Resource Manager 1.1 System Administration Guide search for the title on `http:// docs.sun.com`

Solaris Resource Manager by Richard McDougall - Sun BluePrints OnLine article available via
`http://www.sun.com/software/solutions/blueprints/0499/solaris1.pdf`

Author's Bio: Don De Vitt

Don is currently a Senior Staff Engineer and PC interoperability specialist within the Enterprise Engineering group, and is a member of the Solaris PC NetLink engineering team where he has focused on performance-related issues. Don has been on the development teams of almost every software and hardware PC interoperability product that Sun Microsystems has produced over the last 13 years.

Don DeVitt 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.