

Sunday, March 16, 2008

Less known Solaris Features: Resource Management - Part 1: Introduction

Resource Management is an rather old feature in Solaris, albeit it got more mind share since it got a really important part of the Solaris Zones. But this tutorial doesn't focus on the usage of Resource Management in conjunction with the zone configuration. I want to go to the basics, because at the end the zone configuration just uses this facilities to control the resource consumption of zones. Secondly you can use this knowledge to limit resource usage in a zone itself.

Resource Management was introduced to solve one important question. You can run multiple programs and services at once in a single instance of the operating system, but how do i limit the consumption of resources of a single application? How do i prevent a single programm from consuming all the resources leaving nothing to others? Resource Management in Solaris solves this class of problems.

Why do you need Resource Management? Unix systems were capable to execute a multitude of independent services on one system since it's introduction many years ago. But there was a problem: What happens in case of a process running amok? A user eating away resources of the system? You have to control the allocation of this resources. How many CPU does a process get? How many file descriptors are allocated to a user? The Solaris Resource Management is capable to put other processes under the control of an user configurable ruleset. Thus you can prevent a situation where your system is down, because a single process takes all resources.

With such an ability, the usage of a single system for a multitude of services becomes more manageable and thus more feasible.

Definitions Okay, as usual, this technology has it's own jargon. So i have to define some of it at first:

Tasks A task is a group of processes. For example when you log into a system and do some work all the steps you've done are an task until you logout or open a new task. Another example would be a webserver. It consists out of a multitude of processes, but they are all part of the same task. The database server on the same machine may have a completely different task id.

Projects A project is a group of tasks. For example you have a webserver. It consists out of the processes of the database task and the webserver task.

Zones From the perspective of the resource management, a Solaris Zone is just a group of one or more projects.

The basic idea of Solaris Resource Management The basic idea of resource management is to use the grouping of processes and to impose limits on a group on processes defined by one of the above entities. A limit on the zone would limit all projects in a zone. A limit on the project would limit all tasks grouped by the project. A limit on a task would limit all processes grouped by the tasks.

Let's take an example: There is a resource control called max-lwps. It controls the amount of processes. This control exists as task.max-lwps, project.max-lwps and zone.max-lwps. With this control you can define limits for the number of processes in a group of processes defined by one of the entities mentioned above.

You want to ensure, that you have 1000 processes at maximum in your zone webserver. You can define it by using zone.max-lwps for the zone webserver. You want to ensure, that the www.c0t0d0s0.org webserver project consisting out of mysql and apache processes uses 250 processes at maximum. You use the project.max-lwps for the webserver-c0t0d0s0. You want to limit the httpd-server itself to 100 processes. You impose this limit by using the task.max-lwps control.

How do you configure this? This will be the topic of this tutorial. It sounds complex, but it isn't, once you've understand the entities.