ORACLE

SOLARIS

An Oracle White Paper
November 2010

Deploying SAP NetWeaver Master Data Management on Oracle Solaris Containers

ORACLE

Executive Overview

SAP NetWeaver Master Data Management offers a flexible solution that provides centralized access to master data across the enterprise. SAP NetWeaver Master Data Management aggregates master data from disparate sources, both SAP and non-SAP, using predefined or customizable data models for customers, vendors and suppliers, employees, products, material, and so forth. You can import entire master data schema or structures, as well as master-data values and taxonomy information. During data cleansing, you can carry out various operations, such as deduplication and normalization, ID mapping, matching and merging, staging, interactive data-quality analysis, and ad-hoc consolidation services. SAP NetWeaver Master Data Management harmonizes master data using a powerful syndication model that updates master data in systems running SAP and non-SAP software. Syndication uses industry standard XML and can be configured to run interactively or automatically. This paper describes how Master Data Management server behaves into Oracle Solaris Containers. We'll describe how to install Master Data Management into a local zone, how to achieve maximum system utilization while running SAP NetWeaver Master Data Management in a local zone and what are the advantages of Master Data Management customers when using Oracle Solaris Containers technology.

This white paper is targeted at technical IT managers and system architects who are responsible for implementing, managing, or recommending SAP Master Data Management solutions.

Application overview: Oracle Solaris Containers Overview

Oracle Solaris Containers is an operating system level virtualization technology that provides complete, isolated, and secure run-time environments for applications. This technology allows application components to be isolated from each other using flexible, software-defined boundaries. Oracle Solaris Containers are designed to provide fine-grained control over resources that the applications use, allowing multiple applications to operate on a single Oracle Solaris 10 operating system instance while maintaining specified service levels (Figure 1).

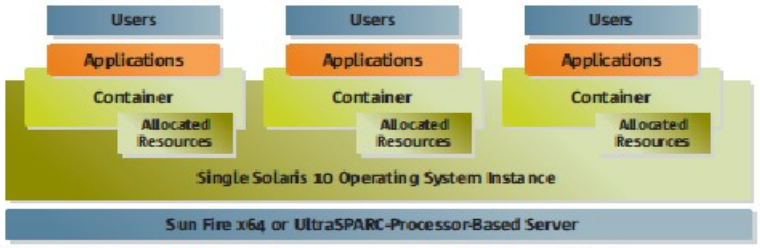

Figure 1. Multiple Solaris Containers on a single Solaris 10 OS instance.

Figure 1. Multiple Oracle Solaris Containers on a single Oracle Solaris 10 OS instance

Oracle Solaris Containers use Oracle Solaris resource management features along with Oracle Solaris Zones software partitioning technology to deliver a virtualized environment that can have fixed resource boundaries for application workloads. These two major components of Oracle Solaris Containers are discussed in the following sections.

Oracle Solaris Zones Partitioning Technology

Oracle Solaris Zones, a component of the Oracle Solaris Containers environment, is a software partitioning technology that virtualizes operating system services and provides an isolated and secure environment for running applications. Oracle Solaris Zones are ideal for environments that consolidate multiple applications on a single server.

There are two types of zones: global zones and non-global zones. The underlying operating system, which is the Oracle Solaris instance booted by the system hardware, is called the global zone. There is only one global zone per system, which is both the default zone for the system and the zone used for system-wide administrative control. One or more non-global zones can be created by an administrator of a global zone. Once created, these non-global zones can be administered by individual non-global zone administrators, whose privileges are confined to that non-global zone.

Two types of non-global zones can be created using different root file system models: sparse root and whole root.

- Sparse root model — The sparse root zone model optimizes the sharing of objects by only installing a subset of the root packages and using a read-only loop-back file system to gain access to other files. In this model, the directories /lib, /platform, /sbin, and /usr are mounted by default as

loop back file systems. The advantages of this model are improved performance due to efficient sharing of executables and shared libraries, and a much smaller disk footprint for the zone itself.

- Whole root model — The whole root zone model provides for maximum configurability by installing the required packages and any selected optional zones into the private file systems of the zone. The advantages of this model include the ability for zone administrators to customize their zone's file system layout and add arbitrary unbundled or third-party packages.

Oracle Solaris Zones provide the standard Oracle Solaris interfaces and application environment. They do not impose a new application binary interface (ABI) or application programming interface (API). In general, applications do not need to be ported to Oracle Solaris Zones. However, applications running in non-global zones may need to be aware of non-global zone behavior, depending on the Oracle Solaris interfaces they use. In particular:

- All processes running in a zone have a reduced set of privileges, which is a subset of the privileges available in the global zone. This set of privileges is available to the root user. Non-root users of a zone have a subset of those privileges. By default, non-global zone non-root users have privileges that are the "logical AND" of the privileges available to non-root users in the global zone and the privileges available to that zone.
- Processes that require a privilege not available in a non-global zone can fail to run, or in a few cases fail to achieve full performance.
- Administrators can modify the privileges that a zone has, reducing or expanding the set. This provides the ability to enhance security by removing privileges not needed by applications running in that zone, or to give a zone a non-default privilege in order to improve the functionality or performance of an application. The privilege proc_lock_memory, required to use Dynamic Intimate Shared Memory (DISM), is now in the default privileges set of zones.
- Each non-global zone may have its own logical network and loop back interface. Bindings between upper-layer streams and logical interfaces are restricted such that a stream may only establish bindings to logical interfaces in the same zone. Likewise, packets from a logical interface can only be passed to upper-layer streams in the same zone as the logical interface.
- Each zone can be configured with exclusive IP privileges which allow it to have its own IP resources. This gives full functionality and independence from the global zone's IP. Specifically, an exclusive-PI zone can manage its own network interfaces, routing table, and IP Filter rules.
- Non-global zones have access to a restricted set of devices. In general, devices are shared resources in a system. Therefore, restrictions within zones are put in place so that security is not compromised.

Oracle Solaris Resource Management

By default, the Oracle Solaris operating system provides all workloads running on the system equal access to all system resources. This default behavior of the Oracle Solaris operating system can be modified providing a way to control resource usage.

Oracle Solaris resource management provides the following functionality:

- A method to classify a workload, so the system knows which processes belong to a given workload.
- The ability to measure the workload to assess how much of the system resources the workload is actually using.
- The ability to control the workloads so they do not interfere with one another and also get the required system resources to meet predefined service-level agreements.
- Oracle Solaris resource management provides three types of workload control mechanisms:
- The constraint mechanism, which allows the Oracle Solaris system administrator to limit the resources a workload is allowed to consume.
- The scheduling mechanism, which refers to the allocation decisions that accommodate the resource demands of all the different workloads in an under-committed or over-committed scenario.
- The partitioning mechanism, which ensures that predefined system resources are assigned to a given workload.

CPU and Memory Management

Oracle Solaris resource management enables the end user to control the available CPU resources and physical memory consumption of different workloads on a system by providing Fair Share Scheduler (FSS), Resource Capping Daemon, CPU Caps and Dedicated CPUs facilities.

Fair Share Scheduler (FSS)

The default scheduler in the Oracle Solaris operating system provides every process equal access to CPU resources. However, when multiple workloads are running on the same system one workload can monopolize CPU resources. Fair Share Scheduler provides a mechanism to prioritize access to CPU resources based on the importance of the workload.

With FSS the importance of a workload is expressed by the number of shares the system administrator allocates to the project representing the workload. Shares define the relative importance of projects with respect to other projects. If project A is deemed twice as important as Project B, project A should be assigned twice as many shares as project B.

It is important to note that FSS only limits CPU usage if there is competition for CPU resources. If there is only one active project on the system, it can use 100% of the system CPUs resources, regardless of the number of shares assigned to it.

Resource Capping Daemon

The resource capping daemon (rcapd) can be used to regulate the amount of physical memory that is consumed by projects with resource caps defined. The resource capping daemon repeatedly samples the memory utilization of projects that are configured with physical memory caps. The sampling interval is specified by the administrator. When the system's physical memory utilization soft cap

exceeds the threshold for cap enforcement and other conditions are met, the daemon takes action to reduce the memory consumption of projects with memory caps to levels at or below the caps.

Virtual memory (swap space) can also be capped. This is a hard cap. In a container that has a swap cap, an attempt by a process to allocate more virtual memory than is allowed will fail.

CPU Caps

CPU caps provide absolute fine-grained limits on the amount of CPU resources that can be consumed by a project or a zone. CPU caps are provided as a zonecfg resource, and as project and zone-wide resource controls.

Administrators can use this feature to control upper limit of CPU usage by each zone. This is in contrast to FSS, which sets the minimum guaranteed portion of CPU time to a given zone if there is a competition for CPU.

For example, consider the following commands to set a CPU cap for a zone

```
zonecfg:myzone> add capped-cpu
zonecfg:myzone:capped-cpu> set ncpus=3.75
zonecfg:myzone:capped-cpu> end
```

The ncpus parameter indicates the percentage of a single CPU that can be used by all user threads in a zone, expressed as a fraction (for example, .75) or a mixed number (whole number and fraction, for example, 3.25). An ncpu value of 1 means 100% of a CPU, a value of 3.25 means 325%, .75 mean 75%, and so forth. When projects within a capped zone have their own caps, the minimum value takes precedence.

Dedicated CPUs

Administrators can use this feature to assign CPUs to zones dynamically within specified minimum and maximum limits per each zone. This eliminates the need to create CPU pools and assign pools to zones, leading to better resource usage and much more simple administration.

For example, consider the following commands to set dedicated CPUs for a zone:

```
zonecfg:myzone> add dedicated-cpu
zonecfg:myzone:dedicated-cpu> set ncpus=8-12
zonecfg:myzone:dedicated-cpu> end
```

With this example, when the zone boots the system creates a temporary dedicated pool for this zone by taking CPUs from the global zone. If the zone will need more CPUs and there will be available CPUs, then the system will assign them to the zone within specified limits.

SAP NetWeaver Master Data Management Architecture Overview

SAP NetWeaver Master Data Management (Master Data Management) is a solution to import, maintain and syndicate master data of all types across different systems. Its in-memory technology allows searching for master data within the Master Data Management repository with very short response times. You can identify potential duplicates across systems and manage key mappings as a prerequisite for master data distribution.

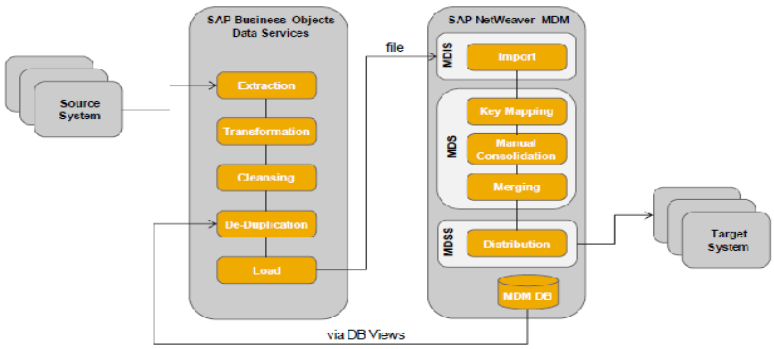

Figure 2. SAP NetWeaver Master Data Management architecture

Data from the different client systems are loaded into SAP Business Objects Data Services by using different native extractors from SAP Business Objects. Extraction of Data from SAP NetWeaver Master Data Management is done via Database Views. Once all required data are extracted the cleansing and de-duplication process starts in Business Objects Data Services. Results of these operations are written into several xml Files. XML-Files have to be imported into Master Data Management using different Import Maps. After the import of both files in Master Data Management a manual Consolidation is performed utilizing the information provided by Business Objects. The Master Data Management Syndication Server syndicates the now cleansed records back to the respective clients.

On the operating system level, Master Data Management server is seen as a group of 3 processes :

- mds-r which is the process in charge with user data search, login, logout activities.
- mdis-r the import server
- mdss-r the syndication server.

Only mds-r is multithreaded, while mdis and mdss are single thread processes. All the 3 processes are working with a data repository which is an in memory database.

Benefits of Deploying SAP NetWeaver Master Data Management in Oracle Solaris Containers

Taking in considerations the characteristics of SAP NetWeaver Master Data Management processes, Oracle Solaris Containers are a very suitable virtualization technology because:

- Offers an isolated work environment for SAP NetWeaver Master Data Management processes. In this way, many instances of SAP NetWeaver Master Data Management server can be run on the same physical server on parallel local zones, the database environment can run also on a parallel local zone, as well other applications.
- Oracle Solaris Containers offers an excellent environment to run multiple versions of SAP NetWeaver Master Data Management on the same Oracle Solaris kernel level.
- Oracle Solaris Containers don't require a dedicated amount of CPU or memory in order to run an application. This way the SAP NetWeaver Master Data Management workload resource requirement can be satisfied on the fly by Oracle Solaris itself without sysadmin intervention. By example, if an import activity is running (which is a single process/single thread) then only one CPU is needed and the rest of CPUs can be used by other applications running on other zones. If a multiuser search is running, then a number of CPUs are needed depending on the parallel users working on the system. In this case, SAP NetWeaver Master Data Management is in need of CPU resources and Oracle Solaris will be able to assign them as needed. If the sysadmin will want to protect other zones from lack of CPUs, then CPU capping can be used.
- By running SAP NetWeaver Master Data Management on Oracle Solaris Containers, the customer maximize the physical server usage: when SAP NetWeaver Master Data Management server is running the single process/single thread activities like import, the rest of server resources can be used by other applications and when SAP NetWeaver Master Data Management is required to work on multiuser searches which is single process/ multithread , then physical resources are added to satisfy the requirements.

Workload Overview

The scope of tests was to study the behavior of Master Data Management processes when installed on a local zone, while a second local zone is stressing the server resources by requiring a high amount of CPU or memory resources. The local zones were defined with whole root and no memory or CPU limitation has been used. The same version of SAP NetWeaver Master Data Management was installed on both local zones and the same archive image was used to build the repositories for each zone. We wanted to check that SAP NetWeaver Master Data Management processes are not failing when the system is struggling for CPUs or memory, and also to analyze how the tests performance are affected by lack of resources.

Workload Description

To simulate the users requests, we used LoadRunner with Java extension and special scripts have been developed for this purpose: import test, syndication test and multiuser test which included a mixture of search requests.

CPU stress tests

For these tests, the multiuser test ran in parallel :

- **on zone1** : always the multiuser test with 100 users which loaded the server to 50% CPU busy as seen with vmstat(1) Oracle Solaris command.
- **on zone2** : the multiuser tests with variable number of users to load the system with additional requirements of CPU : 0%, 25%, 50%, 100%, 0% busy CPUs. It means that on zone2, the multiuser test was run with 50 users to get the additional 25% of CPU usage, with 100 users to get the additional 50% CPU busy and 200 users to load the CPU to 100% only from zone2.

The test results are summaries in Table 1.

TABLE 1. CPU TESTS

ZONE1	ZONE1	ZONE2	ZONE2
% CPU BUSY	#OF RUNNING USERS	%CPU BUSY	#OF RUNNING USERS
50%	100 users	0%	0 users
50%	100 users	25%	50 users
50%	100 users	50%	100 users
50%	100 users	100%	200 users

Memory stress tests

When one repository is loaded, the mds process size become about 4GB. To collect the performance results, the multiuser test with 100 users was always run on the zone with 1 repository loaded (zone1).

- **on zone1** : always one repository was loaded
- **on zone2** : repositories were loaded one at a time until a total of 5 repositories requiring 20GB virtual memory were loaded in the system. See Table 2, below.

TABLE 2. MEMORY STRESS TESTS

ZONE1	ZONE2
1 repository = 4gb	1 repository = 4gb
1 repository = 4gb	2 repository = 8gb
1 repository = 4gb	3 repository = 12gb
1 repository = 4gb	4 repository = 16gb
1 repository = 4gb	5 repository = 20gb

Test Environment

SAP NetWeaver Master Data Management server was installed on: T5440, 125 virtual CPUs, 20GB RAM, 90 GB swap space from which 80 GB were defined on a file on a filesystem. Oracle database was installed on : V890 8*2cores CPUs, 64GB RAM, local disk storage. LoadRunner run on x4450

Software used

- Oracle Solaris 10 u8
- SAP NetWeaver Master Data Management 7.1.4.137
- Oracle Database 10 R2
- HP LoadRunner 9.5

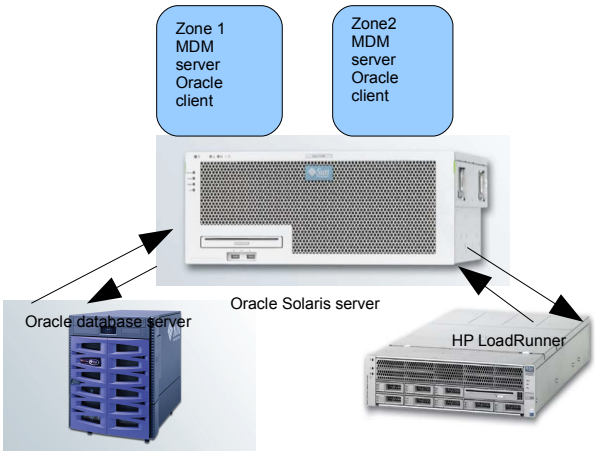

Figure 3. Test hardware configuration

Results

The SAP NetWeaver Master Data Management processes can run safely in Oracle Solaris Containers. If in a parallel local zone an other application is, by accident, using all server resources, the SAP NetWeaver Master Data Management will continue to successfully give services to its clients.

Because the multiuser tests are CPU bound, when one zone requires all the CPUs, the performance on the second zone is impacted. Oracle Solaris Containers technology can be safely used with a number of solutions, including:

1. One zone runs an SAP NetWeaver Master Data Management instance for development, while the second zone is running the SAP NetWeaver Master Data Management instance for QA purpose both using the same Oracle Solaris kernel.
2. Oracle database can work on a local zone, while SAP NetWeaver Master Data Management server is working on a local zone.
3. When one zone is running multiuser searches, the other zone can work on import, export, matching, syndication activities.
4. Oracle Solaris Containers CPU capping can be used to limit the CPU usage of the zone which is suspected to required more CPU than planned and can hurt the work on other zones.

During the memory tests, the mds process size is increasing until it overpass the RAM size. Then the mds process rss become smaller than the size of the process and external swap space is used to resume the work. When external disk was used as additional swap space, the performance wasn't hurt (less than 0.01%)

Following is the output of Oracle Solaris "prstat -s rss" command as seen from the global zone:

2 zones, with 1 repository on each:

```
PID   USERNAME    SIZE      RSS   STATE PRI NICE TIME        CPU   PROCESS/NLWP
7397  adina       7497M  7 484M   sleep   1    0 42:12:07    0.0%  mds-r/89
13491 adina       3974M  3961M    sleep  59    0 0:20:46     0.0%  mds-r/89
```

2 zones, 4 repository on one zone, 1 repository on tested zone

```
PID  USERNAME  SIZE  RSS    STATE PRI NICE TIME     CPU   PROCESS/NLWP
13491 adina     15G  9525M  sleep   1  0   1:48:11  0.1%  mds-r/110
7397  adina   8041M  4586M  sleep   1  0  82:55:41  0.0%  mds-r/89
13425 adina   258M   249M   sleep  59  0   0:02:22  0.0%  mdss-r/12
```

Installing SAP NetWeaver Master Data Management in Oracle Solaris Containers

SAP NetWeaver Master Data Management is installing its files under /usr/sap director. This is why the installation on a local zone requires r/w permissions for /usr/sap director for mdmamd user. A full root local zone installation was choose, which does not use inherited packages and defines /usr with r/w permissions. A full root zone installation makes copies from the global zone of the /usr, /lib, /sbin and /platform directories to a local zone. The whole-root model provides the maximum configurability by installing all of the required and any selected optional Oracle Solaris software packages into the private file systems of the zone. The advantages of this model include the ability for zone administrators to customize their zone's file-system layout . The disadvantages of this model include a heavier disk footprint, it requires approximately an additional 2 Gigabyte -- for each non-global zone configured as such.

Zone installation

All the commands are running from the global zone :

Create zone root director for each local zone

```
# pwd
/
# mkdir mdmzone1root
# mkdir mdmzone2root
# chmod 700 /mdmzone1root
# chmod 700 /mdmzone2root
```

Create a director for ***SAP NetWeaver Master Data Management*** *package*:

```
# cd /mdm
# mkdir mdmzone1
# mkdir mdmzone2
# chmod 777 mdmzone*
```

Configure the local zones :

```
# zonecfg -z mdmzone1
mdmzone1: No such zone configured
Use 'create' to begin configuring a new zone.
zonecfg:mdmzone1> create
zonecfg:mdmzone1> set zonepath=/mdmzone1root
zonecfg:mdmzone1> set autoboot=true
zonecfg:mdmzone1> set scheduling-class=FSS
zonecfg:mdmzone1> add fs
zonecfg:mdmzone1:fs> set dir=/mdm
zonecfg:mdmzone1:fs> set special=/mdm/mdmzone1
zonecfg:mdmzone1:fs> set type=lofs
zonecfg:mdmzone1:fs> end
zonecfg:mdmzone1> remove inherit-pkg-dir dir=/lib
```

```
zonecfg:mdmzone1> remove inherit-pkg-dir dir=/platform
zonecfg:mdmzone1> remove inherit-pkg-dir dir=/sbin
zonecfg:mdmzone1> remove inherit-pkg-dir dir=/usr
zonecfg:mdmzone1> add net
zonecfg:mdmzone1:net> set address=yourIPzone1
zonecfg:mdmzone1:net> set physical=nxge0
zonecfg:mdmzone1:net> set defrouter=yourIPdefaultrouter
zonecfg:mdmzone1:net> end
zonecfg:mdmzone1> verify
zonecfg:mdmzone1> commit
zonecfg:mdmzone1> exit
```

Boot and install the zone

```
zoneadm -z mdmzone1 verify
# zoneadm -z mdmzone1 install
# zoneadm -z mdmzone1 boot
# zlogin -C mdmzone1
```

After the local zone is installed and booted, the installation of SAP NetWeaver Master Data Management is as described in SAP NetWeaver Master Data Management documentation.

Summary

By using Oracle Solaris Containers technology, the customers of SAP NetWeaver Master Data Management can consolidate the SAP NetWeaver Master Data Management server and database server, or multiple SAP NetWeaver Master Data Management servers on a same physical server irrespective on how to allocated the cpus and memory resources. Oracle Solaris is able to allocate them automatically depending on each zone consuming. The SAP NetWeaver Master Data Management installation process on a local zone is the same as on a physical server making the zones technologies transparent to the user. Thus, the customer is able to enjoy more flexibly, higher efficiency and lower cost of ownership.

Deploying SAP NetWeaver Master Data Management on Oracle Solaris Containers
November 2010
Author: Adina Kalin

Oracle Corporation
World Headquarters
500 Oracle Parkway
Redwood Shores, CA 94065
U.S.A.

Worldwide Inquiries:
Phone: +1.650.506.7000
Fax: +1.650.506.7200
oracle.com

Oracle is committed to developing practices and products that help protect the environment

SOFTWARE. HARDWARE. COMPLETE.