

## Solaris Containers Q&A

### **Q: Are zones supported on Solaris 10 x86?**

A: Yes, Solaris 10 containers are support on both x86 and Sparc platforms. Note that the level of virtualization containers provide is not at a hardware level. There is no dependency on processor architecture.

### **Q: What kinds of applications,if any, cannot be run within a zone?**

A: Any non-privileged (setuid not root) should run in a zone. Most setuid applications will most likely run in a zone. The following paper provides a good overview of the types of applications that may not work in a non-global zone.

[http://developers.sun.com/solaris/articles/application\\_in\\_zone.html](http://developers.sun.com/solaris/articles/application_in_zone.html)

### **Q: From the global zone does a PS snapshot give you a list of all the processes (i.e. does it include processes in non-global zones)?**

A: ps (1) can show you all processes, as well as processes per zone. two new options to ps (-z and -Z) give you zone specific information. Most process related commands (prstat, ptree, pkill and pgrep) have zone specific options.

### **Q: Can I use other vendors filesystems, like Veritas, and ADIC within single zones?**

A: From the Veritas SF4.1 release notes:

VERITAS Volume Manager Software Limitations

No Support for Local Zones

VERITAS Volume Manager does not support local zones.

VERITAS File System Software Limitations

Local Zone Support

The following features are supported with local zones:

- \* Access to a VxFS file system in the global zone from the local zone of a lofs file system
- \* Access to ODM files from local zones
- \* Concurrent I/O with files from local zones
- \* ktrace command in the global zone only
- \* odmstat command in both global and local zones

The following features are unsupported with local zones:

- \* Admin ioctls
- \* Administration commands
- \* VSM
- \* VFS/VxMS
- \* Quick I/O and CQIO

\* Cluster File System (CFS)

**Q: What is Solaris PRM?**

A: Process rights management enables processes to be restricted at the command, user, role, or system level. The Solaris OS implements process rights management through privileges. Privileges decrease the security risk that is associated with one user or one process having full superuser capabilities on a system. A privilege is a discrete right that a process requires to perform an operation. The following paper provides a good overview on process rights management.

[http://developers.sun.com/solaris/articles/application\\_in\\_zone.html](http://developers.sun.com/solaris/articles/application_in_zone.html)

**Q: Is there one kernel per zone?**

A: No. There is one kernel for the entire system, both for the global zone, as well as each of the non-global zones. Each zone runs a zsched and init process.

**Q: Is VxVM? supported in a zone?**

A: Simple answer is no, but a filesystem built on a Veritas Volume is supported in a zone. You cannot access a Veritas device volume directly in a zone.

**Q: Is each non-global zone itself a process on the system? If you looked at PS from the global zone, can you thererore track how much CPU a zone is using by looking at one process in PS?**

A: No. A zone is not a process on the system. It is a virtualized enviroment provided by Solaris kernel. A process running in a zone is a regular process, but with a restricted view of the system. This process only has visibility to other processes in the same zone .

A process in a non-global zone has visibility from the global zone. You can use the prstat(1M)command from the global zone,you can find out how much CPU processes in the non-global zones are consuming.

**Q: Can processes communicate using sockets between zones?**

A: zones can only communicate via the networking api such as sockets.

However there a limitation related to raw sockets.

The paper on 'Bringing your applicaiton into the zone' has more information on this. :

[http://developers.sun.com/solaris/articles/application\\_in\\_zone.html](http://developers.sun.com/solaris/articles/application_in_zone.html)

**Q: Can we still use Kstat under a zone? Does Solaris also have AIX Advanced accounting like feature to get information about all zones through one accounting sub system?**

A: Yes, kstat can be run both in the global as non-global zone. While I am not familiar with AIX advanced accounting feature, this might be similar to the extended accounting subsystem in Solaris. Extended accounting does support zones and projects.

The [Solaris Containers--What They Are and How to Use Them Blueprint](#) goes in great detail about Extended accounting and zones.

**Q: Would it be possible to assign a different tape drive to each zone?**

A: Yes. You can export the device from the global zone to the non-global zone. Beware of device sharing issues that can occur if the device is shared among several zones.

**Q: Can the zone configuration administration starting with creating it and ending with shutting it down, can everything be fully scripted?**

A: Yes. Please refer to the zonecfg(1M) and zoneadm(1M) man pages for further detail.

**Q:Our experiments show there are many unresolved issues with packaging and patching tools. Unexpected behaviour, unable to install files into non-global zones etc. We have some difficulties in packaging software and integrating a Solaris 10 jumpstart server into our environment. What's the timeline for a stabilized and "full" functioning zone/pkg/patch environment?**

A: Most of the issues are probably due to the inherited-pkg-dir (/usr, /lib, /sbin mounted in read only mode as loop-back file system. To work around these issues, use the whole-root method to create the zone.

**Q: Can I hot swap devices between local zones? Example, a USB drive or a CD is part of local zone tree, then removed, and added to local zone?**

A: Zones have only a view of devices that were explicitly configured when the zone is created. Adding devices to the system do therefor not automatically show up to the local zone. You have to export the device from the global zone to the local zone via zonecfg(1M). This will require a zone reboot (not of the entire system). This is much faster than a full system reboot.

**Q: Can zones be bound to a CPU? And how does Fair Share Scheduler work on a multi-way server?**

A: Yes, a zone can be bound to a resource pool or processor set. In addition, there are two zone-wide resource controls (zone.cpu-shares and zone.max-lwps) control the cpu-shares within a resource pool or processor set. Within a zone, shares can be subdivided even more using FSS and projects.

The [Solaris Containers--What They Are and How to Use Them Blueprint](#) walks you through a sample case on how to do this.

**Q: Has any work been done regarding testing [NetBackup?](#) with this? Esp. with a media server role?**

A: NBU Server (Master and Media) - will "not" be supported in Local Zones. NBU Client plans are to test and support it in Sept 2005 timeframe.

**Q:Are individual containers able to be set to receive an ip address from a DHCP server?**

A: No DHCP support yet. but it is something that will be there in Solaris 10 update 1.

**Q: Can a process signal that a local zone should be terminated? I'm curious how one builds an automated zone scheduler for, say, isolated unit tests of a system application.**

A: The only official way to stop a zone is from the global zone via zoneadm(1M).

**Q: How do you control what zones processes started via /etc/rc2.d start in. Is there a separate rc2.d for each zone?**

A: Yes. Also each zone has their own smf(5) database.

**Q: How does it handle the same url from 2 webservers in different zones?**

A: I suspect you might be inquiring on whether zones provides some sort of loadbalancing. It does not do this differently as if the two webservers would be running on two different machines. So you would need to set up loadbalancing seperately.

**Q: Can each zone can have its own /opt?**

A: Each zone can have its own /opt, or it can be shared from the global zone (in the sparse root scenario) using lofs. By default, each zone has its own /opt.

**Q: Can we assign one network interface to more than one zone?**

A: The physical network interface is shared among multiple zones. Zones will create logical interfaces on top the same physical interface. However, zones will not be able to see the traffic of the other zones

**Q: What happens when we install softwares on /usr/local of global zone?**

A: It depends if /usr/local and /usr is shared with the global zone. In the whole root model, you can create and install /usr/local in your local zone. In the sparse root model where /usr is a read only loopback from the global zone you can make the /usr/local rw available by adding the fs in zonecfg or even by using symlinks. (the Solaris Forums detail the different ways of doing this)

[Solaris zones forum](#)

**Q: What does the (1M) notation stand for?**

A: (1M) behind a command refers to the section in the main pages.

**Q: Will the backup application see all zones, ie all files?**

A: Yes, a backup utility running in the global zone will see all the files (i.e. all the files in the global zone and non-global zones).

**Q: If a process in one zone goes into a spin, could this crash the other zones?**

A: It will consume lots of CPU resources. To prevent this from happening, bind the zone to a resource pool which can be used to limit the amount of CPU resources a process in a zone can consume.

**Q: How are you going to configure HA VCS in the zoning configuration?**

A: Veritas Cluster Server (CS) engine, agents do not run in a local zone, only agent entry points. A zone is treated as a resource in a service group, so only failover to a different system is supported. Only script based agent entry points are supported at this time. Appendix B of the VCS 4.1 User's Guide has more details:

[http://www.veritas.com/home/pubs/vcs/4.1/solaris/to\\_mfg/en/vcs\\_users.pdf](http://www.veritas.com/home/pubs/vcs/4.1/solaris/to_mfg/en/vcs_users.pdf)

**Q: Is MVFS (ClearCase?) Supported?**

A: The current schedule has ClearCase support for Solaris 10 Sparc by the end of the year and a beta release of Solaris 10 x86 will be out around Q2CY2006.

**Q: Are there any pre-defined zone cfgs like profile which I can apply on a zone?**

A: We do not provide specific zone cfgs, although various scripts are provided in the documentation. For example the admin guide has specific info on this.

**Q: What is the minimum hardware required to implement Containers?**

A: zones require extra hard drive space to put zone specific files. The amount of disk space depends on whether you use the whole root or sparse root model. In the whole root model, you'll need about 2GB extra disk space. Zones by itself does not require hardware assistance to work (as domains do). We have tested hundreds of zones on a simple Ultra10. Size the server based on the applications you will be running. Zones have only little overhead outlined above.

**Q: From the file system point of view, how many zone models are?**

A: Not sure about this question, but there are two methods used to create the zone's root file system when a zone is created.

The two methods are referred to as the sparse-root and whole-root model. The following article has a good explanation of the two methods

[http://www.sun.com/bigadmin/features/articles/zones\\_partition.html](http://www.sun.com/bigadmin/features/articles/zones_partition.html). This same article demonstrates the various methods of making filesystems available to a zone.

**Q: How much info is available regarding Solaris cluster and Veritas cluster and how they work with this (if VxVM? isn't fully supported, what about Solstice disksuite etc.). I can see some customer's trying to do some very interesting things with clusters and containers.**

A: You cannot export a Solaris Volume Manager metadevice or a VxVM virtual disk to a zone. Please see the above answer to the VCS question. We are working on support for Sun Cluster Resource Groups within local zones. The RG would be able to failover to another local zone, either on the same system or on a different system. This will be available in an update to Solaris 10

**Q: Do containers support weblogic?**

A: Yes, BEA has tested and support WebLogic running in both global and non-global zones.