

Oracle Solaris Virtualization

Oracle Solaris 11 Express

Introduction

Virtualization, no longer the new kid on the block, has matured to become a core technology in today's enterprise datacenter. Whether the goal is to reduce costs by consolidating old, costly, underutilized systems or to increase revenue by implementing agile business practices, virtualization is here to stay.

Oracle continues to innovate and integrate right across the virtualization spectrum with technologies engineered to fit all your current needs. Should it be a requirement for hypervisor based technologies, such as Oracle VM for SPARC or Oracle VM for x86, desktop based solutions like Oracle VirtualBox, or a cross platform, extremely lightweight OS virtualization technology such as Oracle Solaris Zones, the right solution is available now.

Oracle Solaris Zones is built-in OS virtualization which continues to build on it's long and distinguished pedigree. One of the most highly adopted, highly used, mature virtualization technologies, Oracle Solaris Zones was first introduced as a core part of Oracle Solaris 10. As of Oracle Solaris 11 Express, Oracle Solaris Zones becomes even more central to both the application and the end user, and continues to introduce enhancements with new improvements and features including:

- Integration into the new packaging system
- Improved observability
- Increased control over administration
- Tight integration with ZFS
- Support for Oracle Solaris 10 Zones
- Integration with the new Oracle Solaris 11 network stack architecture

Zones and Image Packaging

Oracle Solaris Zones are fully integrated with the newly introduced Oracle Image Packing System (IPS). This is particularly useful in the area of Oracle Solaris Zones package management and updating. The process of updating a zone becomes simpler and more efficient, for example a zone can now be updated using only one command. By combining this functionality with ZFS and boot environments, the ability to perform completely safe system upgrades is easily achievable for systems running Oracle Solaris Zones.

In addition, zone instances now have much less software installed (by default) when compared to previous releases. The content of the zone can be customized at install time, allowing additional packages to be added to meet business goals. Starting from a small core of software can also enhance security, a practice known as security minimization.

Finally, IPS allows you to create packages filled with dependency information. This lets you define (and update) the recommended set of packages for your enterprise, so that you can ensure a consistent software base across all of your zones.

In combination, the above features allow the system administrator to reduce the time to deploy or update their zone environments. These features also combine to reduce the length and risk of planned maintenance windows. And flexible packaging content allows different application requirements to be met without pursuing the "install everything" approach to software installation.

Zones Observability

Oracle Solaris 11 Express, includes a new and highly integrated **zonestat** utility. **zonestat** reports highly accurately on the CPU, memory, and resource control

Oracle Solaris Virtualization

Oracle Solaris 11 Express

utilization of running zones. Each zone's utilization is reported as a percentage of both system resources and the zone's configured limits. **zonestat** prints a series of reports at specified intervals. Optionally, the utility can print one or more summary reports.

The **zonestat** utility can be used from within a zone, for a localized view of zone system resources, or from the global zone for a system wide view of zone system resources.

With this improvement in observability it is much easier to identify resource bottlenecks or misbehaving applications. The system administrator is able to take prompt remedial action, accurately addressing the identified problem. Capacity planning is also greatly simplified. In addition, **zonestat** also helps the administrator quickly understand how zones and resource management have been configured. This can be particularly useful when taking over administration duties on an unfamiliar system.

Delegated Administration

With Oracle Solaris 11 Express the administration of Oracle Solaris Zones becomes much more flexible. You have the ability to delegate common zone administration tasks for specific zones to different administrators using Role-Based Access Control (RBAC). With delegated administration, for each zone, a user or set of users may be identified with the permissions to *login*, *manage* or *clone from* that zone. These specific authorizations will be interpreted by the appropriate commands running in the global zone to allow access at the correct authorization level to the correct user.

This is particularly powerful in a shared environment where it is desirable to allow specific users to only manage zones that are relevant to their role. For

example you may be hosting different customers side by side in the same environment.

Even more fine grained control is possible. For example, a system administrator can be given full control of a specific zone, whereas a developer can be restricted to just login to the zone from the global zone.

Zones and ZFS

Oracle Solaris 11 Express builds on the work already done integrating ZFS with Oracle Solaris Zones. It is already possible to manually place zones on ZFS and delegate datasets to non-global zones. With this release of Oracle Solaris 11 Express a zfs dataset, known as the Zones Boot Environment dataset (ZBE), will be automatically provisioned for the zone upon install.

In addition to reducing the number of administration steps, the above improvement has one immediate advantage: Zones now participate fully in the boot environment (BE) framework. When a system is upgraded using IPS, all of the zones on that system will be cloned using ZFS cloning. Upon reboot of the system, the newly updated clones of the zones will be active. This ensures both a clean upgrade process and the ability to roll-back in the event that something goes wrong.

In a hosted environment, end-users within zones can use ZFS filesystems, including creation, destruction, snapshotting and cloning from within the zone. The global zone system administrator can focus on provisioning top-level datasets to zones and on setting space quotas. They are freed from coordinating with non-global zone users about filesystem layouts. The end result is that global zone administrators can provide flexible environments which are still safe and easy to update.

Oracle Solaris Virtualization

Oracle Solaris 11 Express

Oracle Solaris 10 Zones

Oracle Solaris 10 Zones introduces the ability to run Oracle Solaris 10 environments on top of the Oracle Solaris 11 Express operating system. For users already running their applications either in zones or on bare metal on Oracle Solaris 10 systems, Virtual to Virtual (V2V) and Physical to Virtual (P2V) tools are provided to help relocate workloads into Solaris 10 Zones.

Due to the Oracle Solaris Binary Application Guarantee Program, applications can be moved over to Oracle Solaris 11 Express. Solaris 10 Zones provides an option for those applications whose environment doesn't lend itself to easy duplication. Perhaps the environment is complex, or a quick move without reinstallation is required. Solaris 10 Zones is a proven, tested and fully supported option for the quick adoption of Oracle Solaris 11 Express. With Solaris 10 Zones your applications can be moved over to and benefit immediately from all the new features available in Oracle Solaris 11 Express.

Network integration

Oracle Solaris Zones support for networking technologies is significantly enhanced in Oracle Solaris 11 Express. And networking in general is greatly improved, thanks to a powerful new network stack architecture also known as the Crossbow project. This provides sophisticated virtual network interfaces (VNICs), virtual network switches (called "etherstubs"), and the ability to define, monitor, and resource manage network traffic flows.

Oracle Solaris 11 Express combines these new facilities with an important technology called Exclusive-IP zones. An Exclusive-IP Zone provides a complete, tunable, manageable and independent networking engine to each

zone; such a zone can configure DHCP, IP Routing, IP multipathing, IPsec, and more. However, in Solaris

10, exclusive-IP zones have a fundamental limitation: you need to dedicate a physical network interface (NIC) to the zone.

Oracle Solaris 11 Express resolves this limitation by introducing the ability to create high performance virtual NICs (VNICs) atop physical NICs. The result is that you can create an unlimited number of exclusive-IP zones, even on a system with a single physical NIC.

This capability has also been extended to Oracle Solaris 10 Zones. This is particularly critical when performing P2V conversions. With exclusive IP instances available inside Oracle Solaris 10 Zones, you are able to consolidate more physical systems while maintaining the original systems network identities and ability to control the IP stack.

Finally, you can allocate and limit the bandwidth dedicated to zones using VNICs. This innovative bandwidth control system can in many cases program the underlying physical NICs to implement these controls in hardware. For example, a flow might be defined to represent all HTTP traffic. You could then monitor the traffic on and assign bandwidth controls to such flows.

By way of illustration, after a consolidation exercise one zone may contain a web based sales application for your products. A second zone may contain the mail server for your employees. By assigning bandwidth limits to the VNICs it is possible to ensure that the revenue generating sales application is never starved for network bandwidth by the internal mail server.