

Solaris 10 Overview

Introduction

The purpose of this document is to outline, from an IT Production viewpoint, the main key features of the new Solaris version 10 from Sun Microsystems.

Specifically, the aim is to identify features of Solaris 10 which would be of benefit to IT Production teams, and improvements in scalability, reliability, resilience, and the other “Production-readiness” criteria (¹).

Scope

The scope of this document concerns the supportability and production readiness of Solaris 10 only. Issues involving development suitability are out of scope of this document.

The target audience for this document are IT Production Managers and Support teams, and IT Architects.

Approach

Information on Solaris 10 has been obtained from documentation from Sun Microsystems, and opinions posted on the Internet.

¹ For details of what constitutes “Production-Readiness”, refer to the web site at <http://www.dennisadams.net/productionit.htm>, where these terms are defined and expanded.

Preamble: Solaris 10 from Sun Microsystems

Sun launched version 10 of its 64-bit Solaris operating system in the first quarter of 2005.

One of the significant ways in which this version differs from previous versions of Sun Solaris is that version 10 is an Open Source distribution. It therefore ranks as a strong competitor to Linux on Sparc systems.

Secondly, Solaris 10 is not only available on Sun Sparc systems, but is also available for x86-based platforms. This reflects Sun's on-off-on relationship with x86 technology.

At one time, earlier versions of Solaris were not available on x86. Then with the takeover of Cobalt Systems, Sun began to deploy x86 systems again (typically on Linux). Arguably, it has been pressure from existing customers that has finally persuaded Sun to agreed to continue to support Solaris on the Intel / AMD chips.

The fact that Solaris is now available on x86 architectures is another argument for comparing Solaris as a potential rival to Linux.

Thirdly, the release of Solaris 10 was followed by a decision to Open Source the actual Sun Sparc chip itself. There has been some debate about what Sun is attempting to achieve by this action, but it appears that they have opened the Pandora's box to Open Source of Sun intellectual property.

Main features of Solaris 10

The **Process Rights Management** capability makes it possible for system admins to administer systems without having the unlimited power provided by the Unix “root” account.

Support for **Secure Execution** (digitally signed binaries) has been introduced. In future, it would be possible to lock a system down so that only valid, signed executables from a list of trusted authorities will be allowed to run.

The new **ZettaByte File System** (ZFS) is a combination journalled file system (replacement for UFS) and volume manager. This will be available as an update to Solaris 10. It includes 128-bit addressing, 64-bit checksums, copy-on-write (to minimise recovery times) realtime checksum validation and single management system syntax for both the file management and volume management. ZFS could become a powerful replacement for UFS. It could also potentially replace VxFS (the Veritas/Symantec file system), although VxVM includes database-specific backup facilities, mirroring and clustering.

The **Solaris Container capability** is based on Zones and a Resource Manager. Zones isolate application components from each other, even though they share the same instance of Solaris 10. This means that separate application (e.g. databases, firewalls, mail servers, and web servers) can run at separate IP addresses and as if they were on separate machines.. (See the section on “Sun Virtualisation and Partitioning”, below.).

Dynamic Tracing (DTrace) is a software framework and scripting language (“D”) for tracing system performance in real time, and to aid the tuning of applications.

So-called “**Predictive Self-Healing**” consists of Solaris Fault Manager and Solaris Service Manger. These tools improve the serviceability of the Solaris implementation. Solaris Fault Manager identifies failing components (e.g. CPUs, memory chips, I/O devices) and off-lines them from the running system in a controlled way. Solaris Service Manager is responsible for automatic start/restart of failing software components.

Solaris 10 also includes an **improved TCP/IP** stack with performance enhancements and support for IPv6 etc., and a re-designed Network Cache Architecture (NCA). Sun claim a 50% improvement in TCP/IP throughput. Stream Control Transmission Protocol (SCTP) and Session Initiation Protocol (SIP) are supported for implementing Voice over IP.

Solaris 10 is available for both Sparc chips and AMD Opteron or x86. On the x86 platform, there is a binary compatibility mode with Linux binaries, so that they can run unchanged on Solaris.

Solaris Futures: Sparc or x86?

The decision to support x86 as well as Sparc for Solaris raises the question of what hardware platform to run Solaris 10 on. Since the effort to port to Solaris 10 may be the same, users have the option to chose alternative hardware platforms, or even migrate to Linux instead.

Robert Frances Group (“RFG”) is one of the analyst organisations that has assessed the implications for IT Managers of the Open Sourcing of Solaris. They have concluded that, effectively, Open Source Solaris on x86 is now in the same market at Open Source Linux on x86 ⁽²⁾.

One of the points that they highlight is that Solaris is available under a more restrictive “Open Source” agreement (The Common Development and Distribution License, or “CDDL”), which effectively precludes code re-use outside Solaris itself. There is still some measure of control over the future development of Solaris. And work done on Solaris would remain within the Solaris code base, and could not find its way legally into, say, a Linux kernel. This may force developers to make a decision whether to work on Linux or Solaris, but not both. In view of the more restrictive conditions of CDDL, it is possible that many developers would choose to concentrate their efforts on Linux development.

Secondly, RFG point out that Sun has been unable to gain very high volume production for Sparc chips, and that they are in danger of falling behind AMD and others in performance, power consumption and heat generation ⁽³⁾. Therefore, they may have to concentrate on the AMD market in future, rather than relying on Sparc technology.

In addition, user feedback has suggested that porting applications from Solaris on Sparc to Solaris on AMD is just as difficult as porting them to Linux. Therefore, for some development teams, Linux would a preferred solution, rather than Solaris on x86.

² “Linux 2.6 and Solaris 10. An Analysis of Two Strategies for Enterprise Operating Systems” Robert Frances Group, February 2005.

³ Sun is planning to release a Quad-Core Sparc chip “Niagara” in 2006, which may break this trend.

Sun Virtualisation and Partitioning

Prior to version 10, the hardware system “Domain” configuration facility for Enterprise Servers (E3800 and above) was available to isolate physical components within a chassis for fault independence.

Sun is also participating in the open source Xen project to develop new virtualization technology to allow multiple O/Ss to run simultaneously on a single hardware system.

With Solaris 10, the Solaris Container capability is based on Zones and a Resource Manager. Zones isolate application components from each other, even though they share the same instance of Solaris 10. This means that separate application (e.g. databases, firewalls, mail servers, and web servers) can run at separate IP addresses and as if they were on separate machines..

Therefore, Sun now have four different approaches to Virtualisation, as follows:

Approach	Description	Possible uses
Dynamic System Domains	Multiple O/S copies on the same hardware. Total physical isolation.	Physical partitioning. Server consolidation of multiple O/S onto separate physical partitions in a large physical chassis.
Xen	Multiple O/S copies on the same hardware. Shared Physical Resources similar to VMWare ESX. Note: unlike VMWare, Xen does not perfectly emulate hardware, so the O/S installed on it must have Xen support in the kernel.	Virtual partitioning. Server consolidation of multiple O/S onto a shared chassis. Implements the O/S isolation of Domains, with the ability to grow resources as required.
Solaris Containers	One copy of the O/S with multiple applications and/or users isolated from each other. The Containers or Zones are in the same O/S, but isolated from each other.	Enables surrogate different IP addresses to link to applications such as J2EE containers with isolation in case of failure.
Solaris Resource Manager	One copy of the O/S with multiple applications granted resource utilisation in preference to each other.	Enables individual applications on the same Operating System to have some prioritisation.

Care should be exercised in choosing the appropriate virtualisation technique.

Licensing implications should also be borne in mind.

Issues and Guidelines for Solaris 10 deployment

This section reviews some of the key considerations that have to be planned for when deploying Solaris 10.

- Support for applications on Solaris 10 is a key determinant of its deployment usefulness and likely timescales. Most applications appear to require re-compiling for implementation with Solaris 10.
- The new Solaris Service Manager Framework (“SMF”) alters the startup/shutdown and restart mechanisms in Solaris. New Startup scripts read XML manifest files that document service dependencies. The utility “svcadm” is used to run the svcs command interface to SMF. The boot process “init” is very different. This will all involve re-training of System Admins.
- A key decision for Solaris 10 deployment will be whether to implement ZFS, and/or what to do with Clustering, Veritas File System etc. ZFS will require a new approach to storage management, since it is both a storage manager and a volume manager. ZFS also includes a volume emulator (Zvol) to emulate block-based devices for some applications such as Sybase etc. A management policy decision would need to be taken on whether and how to implement this.
- There are currently no clustering technologies for Solaris 10, although Sun Cluster with ZFS support (HA-ZFS) is expected at some time. Typically, Sun users have implemented clustering using Veritas or Legato in the past. A management policy decision would need to be taken on how to implement clustering support.
- A management decision would have to be made about what forms of Virtualisation to implement (see the section “Sun Virtualisation and Partitioning”). This will impact licensing. For example, if an application is installed in a Container, then Oracle will allow licensing on the basis of the size of the capped container (which it calls a “hard partition” in the case of a Solaris Container), not the size of the underlying hardware system. The number of CPUs or cores in this pool determines the license cost. This could save significant license costs in the case of some applications.
- Databases on Solaris include Oracle and PostgreSQL. PostgreSQL is an open source database, which is integrated into Solaris 10. Sun are working with the PostgreSQL development community to take advantage of advanced features such as Self-Healing and DTrace. Oracle 9i R2 and 10g R1 databases (RAC and non-RAC) have been certified to run in a global zone. Oracle has also stated that they will use Solaris 10 64-bit as their development platform.
- It is recommended that full dependency listings are compiled of all in-house and COTS (Commercial Off-The-Shelf) software and certification on Solaris 10 is validated before proceeding.

- As well as software validation, it is important to put together hardware and interfaces compatibility lists. Sun supply a Hardware Compatibility Lists (HCL) for Solaris 10, which should form the basis of such a list. However, particular attention should be paid to Network and SAN / NAS adaptors etc.
- If it is planned to implement ZFS, it is recommended that migration planning and testing be put in place. In particular, the full functionality of the various NAS / SAN appliances must be tested, including backup / recovery, snap mirroring etc.
- It is essential that all backup /recovery software (e.g. Veritas, Legato, Tivoli Storage Manager etc.) is officially certified by the vendors if ZFS is being implemented.
- Organisations that have multiple copies of Solaris are strongly advised to create an in-house “build” which contains the tailored elements of Solaris that are required for their organisation. This should be provided either as a set of scripts build notes or an in-house build CD media.

About the Author: Dennis Adams has over 20 years experience working with Managers of IT Production Systems and now provides a specialised Consultancy Service to help IT Production Managers take a more Strategic view of IT Production Management. He can be contacted via his company website at <http://www.dennisadams.net>