

RM4UOS11-60LFEXP Series



QUICK SPECS

Stormax with Solaris Administration SAN

www.stormax.co



Executive Overview 2

The Need for Simplified Data Management and Administration 2

Making Data Accessible..... 3

Supporting a Wide Range of Storage Devices 4

Simplifying Storage Management 5

Delivering I/O Performance..... 5

Keeping Systems and Data Secure 8

Understanding How Systems and Applications are Performing 8

Configuring Storage Devices With COMSTAR..... 9

COMSTAR Task Map 11

COMSTAR Configuring 12

Hardware overview..... 13

Post Error Codes..... 12

Executive Overview

Organizations in every industry rely on information stored in the datacenter to maintain a competitive stance and achieve corporate objectives. Whether using enterprise applications to manage people and processes, bring new products to market, or make discoveries faster, companies depend on their computing infrastructure to deliver data on demand.

While many operating systems provide data management features, deriving the greatest value from information requires a platform designed for all of the pieces to work together. Oracle is the only company that delivers an integrated and optimized technology stack from applications and operating environments, to virtualization and management tools, and storage and networking solutions. At the heart of this platform is the Oracle® Solaris operating system, the premier platform for enterprise computing, providing a robust foundation and built in innovation for data management solutions. This technical white paper explains how upgrading to Oracle Solaris can help IT organizations simplify administration and make better use of storage and data management infrastructure.

Organizations—from small businesses and government agencies, to global enterprises, research organizations, and universities—depend on information. As data volumes continue to rise, datacenter managers must cope with expanding storage infrastructure while providing around-the-clock access to data stored on reliable and secure media. In addition, stringent compliance regulations are forcing organizations to retain data for longer time frames. These issues are compounded by service-level demands that require greater operational efficiency, and economic pressures that necessitate cost reductions. The key to success is finding the right platform for storage infrastructure deployment.

Now datacenters can take advantage of powerful features offered only by the Oracle Solaris 11 operating system to create intelligent, feature-rich storage servers that can ease infrastructure complexity and reduce operational costs.

The Need for Simplified Data Management and Administration

Organizations looking to consolidate systems or reduce management complexity have historically had to rely on third-party solutions that deliver file services. Yet exponential growth in the amount and type of data is changing organizational needs for file services—needs that these third-party solutions often fail to address. For example, companies are looking to optimize performance and capacity utilization based on the type, size, or read/write access requirements for files. Intelligent features, such as virtualization, automatic file migration, and policy-based recovery, can change the way cost-effective storage management is done.

Enterprises need more than basic file systems and data copy functions: a wider array of services is needed to meet increasingly complex needs. Because the operating system occupies a critical position in the computing system hierarchy—below applications and above hardware—it can integrate data services directly and provide an ideal platform for deploying network-attached storage or storage servers.

Making Data Accessible

Whether supporting enterprise users, developers, or educational institutions, storage infrastructure must be able to give users access to growing volumes of data. Oracle Solaris is a major evolution in the Oracle Solaris family, incorporating several technologies aimed at simplifying data management and user access.

A New and Easier Way to Manage Data

With greater amounts of information needing to be stored and accessed, robust data management is key. Built into the operating system is Oracle Solaris ZFS, an enterprise-class, general-purpose file system designed to ease data and storage management complexity. Unlike file systems that require a separate volume manager, Oracle Solaris ZFS integrates traditional file system functionality with built-in volume management techniques and data services. Breaking free of the typical one-to-one mapping between file system and associated volumes, Oracle Solaris ZFS shares space dynamically between multiple file systems from a single storage pool to enable more efficient use of storage devices.

By automatically allocating space from pooled storage when needed, Oracle Solaris ZFS simplifies storage management and gives organizations the flexibility to optimize data placement. Many administrative tasks are automated to speed performance and eliminate common errors. For example, creating file systems is fast and easy. There is no need to configure, or reconfigure, underlying storage devices or volumes—these tasks are handled automatically when devices are added to a storage pool.

Storage Pools

Unlike most operating systems, which leave volume management to add-on tools, Oracle Solaris integrates volume management functions. Breaking free of the typical one-to-one mapping between the file system and its associated volumes, Oracle Solaris ZFS introduces the storage pool model.

Oracle Solaris ZFS decouples the file system from physical storage in the same way that virtual memory abstracts the address space from physical memory, allowing for more efficient use of storage devices. Space is shared dynamically between multiple file systems from a single storage pool, and is parceled out of the pool as file systems request it. Physical storage can be added to storage pools dynamically, without interrupting services, providing new levels of flexibility, availability, scalability, and performance. When capacity is no longer required by a file system in the pool, it is made available to other file systems.




Figure 1. Virtual storage pools let multiple file systems share storage space.

Supporting a Wide Range of Storage Devices

Organizations of every kind are taking advantage of the wide range of storage devices now available to meet capacity and performance demands—from smaller, low-cost systems to high-performance, high capacity devices, and everything in between. Incorporating so many device types into a storage infrastructure requires an intelligent platform with the facilities to simplify device integration. Oracle Solaris supports a variety of protocols and interface technologies and provides key host bus adapter drivers with the software distribution.

- **Internet SCSI (iSCSI) target support.** Many storage deployments rely on the iSCSI protocol to move data across intranets and manage devices remotely. With iSCSI target support, the operating system can make SCSI storage devices available to clients over the network.

- **Fibre Channel.** Oracle Solaris 10 includes a number of Fibre Channel packages, including libraries based on the T11 FC-HBA specification, Emulex and QLogic device drivers, debugging aids, a FCIP IP/ARP over Fibre Channel device driver, a Fibre Channel transport layer, and much more.

- **Host bus adapter drivers.** A wide range of drivers are included in the operating system, including the Adaptec AdvanceRaid Controller SCSI HBA, Adaptec Ultra320, Advanced Host Controller Interface SATA, LSI MegaRAID SCSI HBA, LSI Hardware RAID HBA, LSI MegaSAS RAID Controller HBA, and Marvell 88SX SATA device drivers.

- **Multiple protocol support.** Several protocols are supported by the operating system, including the Serial Management Protocol for Serial Attached SCSI (SAS) that provides a way to communicate with SAS expanders through the Serial Management Protocol (SMP). A SMP target driver lets users issue SMP requests and receive SMP responses through an interface, giving privileged users the ability to configure and manage SAS domains.

Simplifying Storage Management

With so many devices—and types of devices—at work, finding ways to locate, manage, and make available storage solutions and the data they contain is paramount. Storage software includes tools and technologies that make it easier to find, integrate, manage, and work with storage solutions.

Automatic Device Discovery and Configuration Management

Locating devices on the network and integrating them into the storage infrastructure can be a time-consuming task. Oracle Solaris 10 supports the Internet Storage Name Service (iSNS) protocol, an industry standard for automated discovery, management, and configuration of iSCSI and Fibre Channel devices on a TCP/IP network. The iSNS protocol helps administrators identify, connect to, and manage devices by supporting name registration and discovery, discovery domains and logon control, state change notification, and entity status inquiries. As a result, administrators can obtain information dynamically about accessible devices, manage devices with familiar network management tools, enhance security, and adapt to topology changes.

Delivering I/O Performance

As demand rises and services grow in complexity, datacenter infrastructure must provide massive capacity and fast access applications and data in order to keep pace with business priorities. Companies can add storage capacity easily by augmenting infrastructure with additional hard disk drives and arrays.

Unfortunately, the devices that provide the highest capacity are expensive, and fail to provide the I/O performance needed to keep systems supplied with the data for processing. Now Flash technology is emerging as a strong storage alternative that can help rebalance system and storage I/O performance.

Robust data integrity, reliability, and availability, combined with breakthrough performance and power characteristics, make it possible to create a new class of storage platform that combines the strengths of hard disk drive technology with enterprise solid-state disks (SSDs) based on Flash technology. For example, enterprise SSDs can be placed in a new storage tier to assist hard disk drives by holding frequently accessed data to minimize the impact of disk latencies and improve application performance. By using enterprise SSDs to handle certain types of I/O, and hard disk drives to store massive data sets, a Hybrid Storage Pool delivers significant performance gains without sacrificing capacity.

Hybrid Storage Pool technology is designed to exceed the performance of Fibre Channel technologies without the additional management and administration complexity of a SAN. Several Oracle Solaris ZFS components are key to Hybrid Storage Pool operation and help accelerate performance.

- The Oracle Solaris ZFS Adaptive Replacement Cache (ARC) is the main file system memory cache and is stored in DRAM.
- The (Level Two Adaptive Replacement Cache (L2ARC) extends the ARC into read-optimized SSDs to provide a large read cache to accelerate reads. The Oracle Solaris ZFS Intent Log (ZIL) is transactional and uses write-based SSDs to provide a large cache to accelerate writes.
- The disk storage pool consists of conventional disk drives. Note that high-performance, more costly disk drives are no longer strictly required to achieve high performance levels given the interposition of SSDs in a Hybrid Storage Pool. Sophisticated file system algorithms in Oracle Solaris ZFS use the ARC in memory and the L2ARC on SSD to determine pre-fetch or data placement during sustained read operations. SSDs accelerate write throughput for Oracle Solaris ZFS synchronous write I/O operations, helping to boost write performance.




Figure 2: Hybrid Storage Pools optimize data placement to improve I/O performance

Other Oracle Solaris Performance Features

Oracle Solaris includes a wide range of features aimed at improving infrastructure performance.

- Intelligent algorithms.

Oracle Solaris ZFS simplifies the code paths from the application to the hardware, delivering sustained throughput. Block allocation algorithms accelerate write operations, and consolidate many small random writes into a single, more efficient sequential operation. An I/O scheduler bundles disk I/O to optimize arm movement and sector allocation to speed throughput.

In addition, an intelligent prefetch performs read ahead for sequential data streaming, and can adapt its read behavior on the fly for more complex access patterns. Furthermore, data is striped automatically across all available storage devices to balance I/O and maximize throughput. Oracle Solaris ZFS immediately begins to allocate blocks from devices as soon as they are added to the storage pool, increasing effective bandwidth as each device is added to the system.

- Better storage space utilization and increased data throughput.

Today, many companies take advantage of Redundant Arrays of Inexpensive Disks, or RAID systems, to achieve varied levels of performance and availability. With Oracle Solaris 10, the benefits traditionally associated with expensive hardware-based RAID implementations can be achieved with built-in software RAID features. The intelligent RAID-Z implementation automatically spans data across disks to deliver higher transfer rates, and uses parity, striping, and atomic operations to ensure the reconstruction of corrupted data. The result: higher transfer rates, greater on-line storage capacity, increased data availability and system reliability, easier management of large amounts of data, and reduced maintenance and downtime.

- Built-in compression techniques.

Oracle Solaris 11 supports multiple compression algorithms, including lzjb, gzip, and gzip-N, to reduce the amount of disk space used to hold files and increase data throughput. Compression is performed on a per block basis to help improve storage device utilization.

Keeping Systems and Data Secure

In today's hyper-connected economy, the ability to safeguard information and protect against security breaches is essential. Oracle Solaris 11 helps keep systems and data secure with built-in consistency checks, encryption capabilities, access controls, and labeling.

Data Integrity Measures

Oracle Solaris ZFS uses several techniques, such as copy-on-write and end-to-end check summing, to keep on-disk data self-consistent and eliminate silent data corruption. Data is written to a new block on the media before changing the pointers to the data and committing the write. Because the file system is always consistent, time-consuming recovery procedures like `fsck` are not required if the system is shut down in an unclean manner. In addition, data is read and checked constantly to help ensure correctness, and any errors detected in a mirrored pool are automatically repaired to protect against costly and time-consuming data loss and (previously undetectable) silent data corruption. Corrections are made possible by a RAID-Z implementation that uses parity, striping, and atomic operations to help reconstruct corrupted data.

Understanding How Systems and Applications are performing

In order to achieve the highest possible performance levels, organizations need to understand how well applications, systems, and storage are performing. Oracle Solaris 10 includes several tools for assessing operational performance and pinpointing opportunities for improvement.

Oracle Solaris Dynamic Tracing Facility

The Oracle Solaris Dynamic Tracing (DTrace) facility is a dynamic tracing framework that provides top-to-bottom system observability for troubleshooting systemic problems in real time. Designed to quickly identify the root cause of system performance problems, DTrace combines over 100,000 trace points with a powerful scripting language and a simple, interactive command-line interface. It works by safely and dynamically instrumenting the running operating system kernel and applications with trace points (known as *probes*) that are completely passive until enabled. Probes can be enabled quickly for data collection, and then disabled again to minimize performance impacts on the system being examined. Developers and administrators can use this information to quickly identify performance bottlenecks, optimize resource utilization and performance, and quantify resource requirements.

Configuring Storage Devices With COMSTAR

This chapter describes how to configure Common Multiprotocol SCSI TARget, or COMSTAR, a software framework that enables you to convert any Oracle Solaris 11 host into a SCSI target device that can be accessed over a storage network by initiator hosts. This means you can make storage devices on a system available to Linux, Mac OS, or Windows client systems as if they were local storage devices. Supported storage protocols are iSCSI, FC, iSER, and SRP.

COMSTAR and iSCSI Technology (Overview)

iSCSI is an acronym for Internet SCSI (Small Computer System Interface), an Internet Protocol (IP)-based storage networking standard for linking data storage subsystems. By carrying SCSI commands over IP networks, the iSCSI protocol enables you to access block devices from across the network as if they were connected to the local system. COMSTAR provides an easier way to manage these iSCSI *target devices*. COMSTAR utilizes a SCSI TargetMode Framework (STMF) to manage target storage devices with the following components:

- Port providers (or plug-ins) – Implement protocols, such as Fiber Channel (FC) and iSCSI.
- Logical unit providers – Emulate various SCSI devices, such as disk and tape devices.
- The libstmf management library – Provides the COMSTAR management interface. The modules that implement the iSCSI functionality do not interact directly with the underlying transport. In a similar way, the modules that implement the transport protocol are unaware of the SCSI-level functionality that is inherent in the packets they are transporting. Some transport examples are Fibre Channel and iSCSI. The framework separates the execution and cleanup of SCSI commands and the associated resources. This separation simplifies the task of writing SCSI or transport modules.

Use the following to administer these features:

- The `itadm` command manages Internet SCSI (iSCSI) nodes within the SCSI target mode framework.
- The `stmfadm` command configures logical units within the SCSI target mode framework.
- The `srptadm` command manages SCSI RDMAProtocol (SRP) target ports within the SCSI target mode framework.

The following solutions are available to use storage devices in your existing TCP/IP network:

- iSCSI block devices or tape – Translates SCSI commands and data from the block level into IP packets. Using iSCSI in your network is advantageous when you need to have block-level access between one system and the target device, such as a tape device or a database. Access to a block-level device is not locked so that you can have multiple users or systems accessing a block-level device such as an iSCSI target device.

Here are the benefits of using iSCSI targets and initiators in Oracle Solaris:

- The iSCSI protocol runs across existing Ethernet networks.
- You can use any supported network interface card (NIC), Ethernet hub, or Ethernet switch.
- One IP port can handle multiple iSCSI target devices.
- You can use existing infrastructure and management tools for IP networks.
- You might have existing Fibre-Channel devices that can be connected to clients without the cost of Fibre Channel HBAs. In addition, systems with dedicated arrays can now export replicated storage with Oracle Solaris ZFS or UFS file systems.
- The protocol can be used to connect to Fibre Channel or iSCSI Storage AreaNetwork (SAN) environments

with the appropriate hardware. Here are the current limitations or restrictions of using the iSCSI initiator software in Oracle Solaris:

- Support for iSCSI devices that use SLP is not currently available.
- iSCSI targets cannot be configured as dump devices.
- Transferring large amounts of data over your existing network can have an impact on performance.

Configuring COMSTAR (Task Map)

Task	Description
Identify the COMSTAR software and hardware requirements.	Identify the software and hardware requirements for setting up an iSCSI storage network with COMSTAR.
Determine the iSCSI target discovery method.	Determine the iSCSI target discovery method best suited for your environment.
Enable the STMF service.	Enable the STMF service, which provides persistent target information.
Create SCSI logical units and make them available.	Create SCSI logical units (LUNs) and make them available to all hosts or specific hosts for iSCSI or iSER configurations.
Configure the iSCSI target.	Configure the iSCSI target for the iSCSI storage component.
Configure the iSCSI initiator.	Configure the system or systems that initiate SCSI requests to the iSCSI target.
Access the iSCSI disks.	You can access your iSCSI disks with the format utility. You can also enable the iSCSI disks to be available automatically after the system is rebooted.
Restrict LUN access to selected Systems.	You might want to restrict LUN access to specific systems in the network.
Configure Fibre Channel devices.	Configure FC devices with COMSTAR if you have a FC storage array in your environment.
Configure FCoE devices.	Configure Fibre Channel over Ethernet (FCoE) devices with COMSTAR. FCoE functionality is provided through Ethernet interfaces. FCoE ports are logical entities associated with Ethernet interfaces.
Configure SRP devices.	Configure SRP devices with COMSTAR. The SRP (SCSI RDMA Protocol) accelerates the SCSI protocol by mapping the SCSI data transfer phases to Infiniband (IB) Remote Direct Memory Access (RDMA) operations.
(Optional) Set up authentication in your Oracle Solaris iSCSI configuration.	Decide whether you want to use authentication in your Oracle Solaris iSCSI configuration: Consider using unidirectional CHAP or bidirectional CHAP. Consider using a third-party RADIUS server to simplify CHAP management.

Configuring COMSTAR

Configuring your iSCSI targets and initiators with COMSTAR involves the following tasks:


- Identifying hardware and software requirements
- Configuring your IP network
- Connecting and setting up the iSCSI target device
- Configuring the initiators
- Configuring the iSCSI target discovery method
- Creating file systems on your iSCSI disks
- (Optional) Configuring iSCSI authentication between the iSCSI initiator and the iSCSI target
- Monitoring your iSCSI configuration The iSCSI configuration information is stored in the `/etc/iscsi` directory, but it requires no manual administration.

COMSTAR Terminology


Review the following terminology before configuring iSCSI targets and initiators.

Term	Description
discovery	The process that presents the initiator with a list of available targets.
discovery method	The way in which the iSCSI targets can be found. Three methods are currently available: <ul style="list-style-type: none"> ■ Internet StorageName Service (iSNS) – Potential targets are discovered by interacting with one or more iSNS servers. ■ SendTargets – Potential targets are discovered by using a <i>discovery-address</i>. ■ Static – Static target addressing is configured.
initiator	The driver that initiates SCSI requests to the iSCSI target.
initiator group	A set of initiators. When an initiator group is associated with a LUN, only initiators from that group may access the LUN.
iqn or eui address format	An iqn (iSCSI qualified name) address is the unique identifier for a device in an iSCSI network using the form <i>iqn.date.authority:uniqueid</i> . An iSCSI initiator or target is assigned an IQNname automatically when the iSCSI initiator or target is initialized. An eui (extended unique identifier) address consists of 16 hexadecimal digits, and identifies a class of GUIDs that is used in both the SCSI and InfiniBand standards. SRP devices use the eui address format.
logical unit	A uniquely numbered component in a storage system. Also referred to as a LUN. When a LUN is associated with one or more SCSI targets, the target can be accessed by one or more SCSI initiators.
target device	The iSCSI storage component.
target group	A set of targets. A LUN can be made available to all targets in one target group.
target portal group	A list of IP addresses that determines which interfaces a specific iSCSI target will listen to. A TPG contains IP addresses and TCP port numbers

Hardware overview



JBOD Wiring Diagram



Automatic Port Configurations of External SAS Expander Controller

