

Hardware and Software Engineered to Work Together



An Oracle White Paper
February 2013

Oracle's x86 Systems: The Best x86 Platforms for Oracle Solaris



Executive Summary.....	2
More Advanced Fault Management.....	3
Oracle Solaris FMA with Oracle's x86 Systems	4
Oracle Solaris FMA with Intel Xeon Processors.....	4
Total Problem Visibility	4
Oracle Solaris-Aware Service Processor	5
Component Replacement Made Easy	6
Service Automation Through Co-Engineering.....	7
Oracle Auto Service Request and Oracle Solaris.....	7
Automated Parts Delivery	8
Oracle Solaris FMA and Oracle Auto Service Request on Non-Oracle x86 Systems.....	8
Simplified Installation.....	9
Identification of Disks.....	9
Availability of Drivers	9
Preinstall Option	9
Oracle Enterprise Manager Ops Center.....	10
Oracle's x86 Systems	10
More For Less	11
Conclusion	11

Executive Summary

With over 50,000 businesses and institutions running Oracle Solaris today, whether on the cloud or on-premise in data centers, Oracle Solaris has clearly established itself as the world leader in UNIX-based operating systems. Oracle Solaris includes many unique and innovative technologies that are not commonly available in other OSs, such as Oracle Solaris ZFS, Dynamic Tracing (DTrace), predictive fault detection, built-in virtualization, and advanced security.

As a result of a broad strategic alliance between Oracle and Intel, Oracle Solaris has become widely deployed on Intel-based x86 systems from a variety of different hardware vendors. Intel has embraced Oracle Solaris as a mainstream UNIX operating system for enterprise-class, mission-critical systems based on Intel Xeon processors.

While Intel Xeon processor-based systems from any hardware OEM can be certified to run Oracle Solaris, Oracle has engineered Oracle Solaris to work with Oracle's x86 systems to provide benefits above and beyond those that can be achieved with non-Oracle x86 systems. Many organizations are using Oracle Solaris on non-Oracle x86 systems and are, therefore, not reaping the benefits of the enhanced features and cost savings of running Oracle Solaris on Oracle's x86 systems.

This white paper highlights the significant ways in which running Oracle Solaris on Oracle's x86 systems results in key benefits, including the following:

- Oracle Solaris for x86 servers is developed and tested on Oracle's x86 systems.
- More advanced fault diagnosis is provided through integration with the service processor.
- Auto Service Requests are sent to My Oracle Support for all hardware problems.
- Easier serviceability with accurate component identification is provided.
- Installation is simpler due to deep integration with Oracle Enterprise Manager.
- Integration of the technologies above results in completely automated diagnosis and repair.
- Total cost of ownership is reduced by avoiding per-socket subscription costs.

Oracle's x86 systems are also the building blocks of Oracle's engineered systems and are subjected to highly demanding workloads. As a result, businesses using Oracle's x86 systems will benefit from the higher reliability gained from this hardening, allowing system administrators to focus more on innovation and less on the mundane activities of server lifecycle management.

More Advanced Fault Management

The true advantage of running Oracle Solaris on Oracle's x86 systems comes with the features that Oracle has engineered together. A primary example of this is the Oracle Solaris Fault Management Architecture (FMA). Developed and enhanced over the past decade, Oracle Solaris FMA is composed of a set of diagnosis engines that process raw error events from the hardware and provide an automated and intelligent method for problem diagnosis.

Oracle Solaris FMA also includes a set of agents that respond to those fault events, such as offlining a faulty CPU thread or retiring a memory page on a DIMM. These advanced, self-healing features help reduce unplanned downtime by isolating a problem at runtime and keeping applications running.

Having Oracle Solaris FMA on servers helps organizations reduce operational expenses by finding problems faster, identifying and ordering replacement parts more quickly, and reducing server downtime. By letting the advanced diagnosis engines in Oracle Solaris determine the root cause of a problem, IT personnel spend less time debugging problems. In turn, this enables a server to be restored to service significantly faster, translating directly into business value for the enterprise.

While Oracle Solaris FMA is a built-in feature of the operating system, its capabilities can vary dramatically based on the server on which it is running. This is because Oracle Solaris FMA requires hardware-specific rules to be built into the diagnosis engines and, more importantly, it requires a deep understanding of the hardware interfaces of all the components in the system that could potentially fail. This level of detail is lacking on non-Oracle x86 systems, which curtails the efficacy of Oracle Solaris FMA. Since Oracle Solaris engineers perform custom engineering of the diagnosis engines for the hardware interfaces resident on Oracle's x86 systems, only Oracle's x86 systems have complete Oracle Solaris FMA capabilities, as shown in Figure 1.

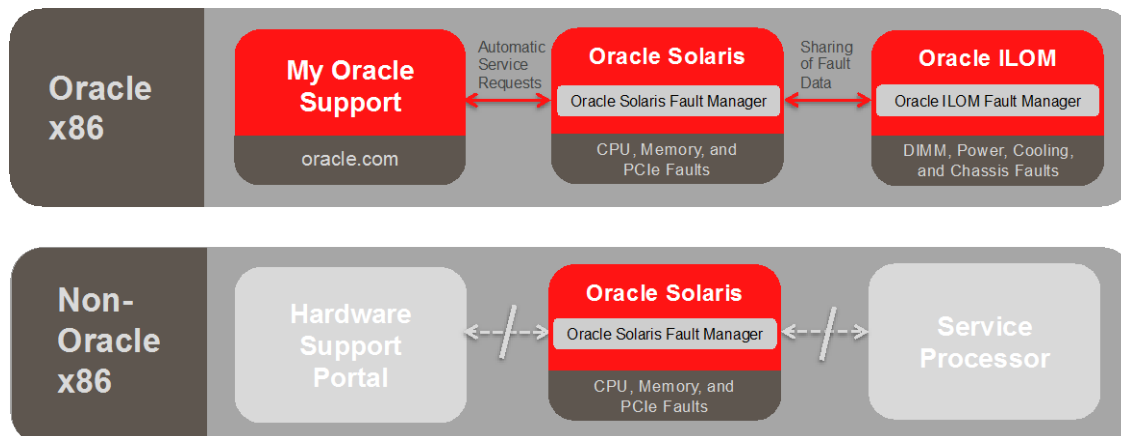


Figure 1. Complete Oracle Solaris FMA capabilities are available only on Oracle's x86 systems.

Oracle Solaris FMA with Oracle's x86 Systems

The advanced fault management capabilities of Oracle Solaris on Oracle's x86 hardware are the result of years of cumulative engineering investments made across multiple product families and generations. Combined innovations in Oracle Solaris FMA and hardware design have resulted in a synergistic progression of new features that benefit end users.

Oracle understands that organizations want a holistic approach to fault management in which all aspects of server health are examined together. This approach prevents system administrators from having to mine and collate error conditions from lots of different sources, such as operating system logs, device-specific tools, and service processors.

Oracle Solaris FMA also filters out noise and unnecessary service calls by determining which error conditions are the result of faulty hardware and must be addressed. In addition, Oracle has engineered Oracle Solaris FMA to diagnose problems in all parts of its servers, including the CPU, DIMMs, PCIe devices, fans, and power supplies.

Oracle Solaris FMA with Intel Xeon Processors

Built out of Oracle's partnership with Intel, Oracle Solaris FMA was instrumented with diagnosis ability for correctable memory errors for DIMMs, correctable errors (CEs) for CPUs, and fault detection for PCIe errors. For CEs, the component can be offlined by Oracle Solaris FMA if the CEs occur too frequently. Also, Oracle Solaris FMA captures error state for CPU uncorrectable errors and reports them upon the next server reboot.

Oracle Solaris has the capacity to offline individual processor strands (no further software threads are scheduled on the affected strand), retire individual pages of memory (4 KB granularity), and cease using problematic I/O devices.

Oracle's x86 systems, all of which use Intel Xeon processors, were used to develop the integration between Oracle Solaris FMA and Intel Xeon processors. In addition, the integration of Oracle Solaris FMA with Intel Xeon processors has been tested and hardened by Oracle on Oracle's x86 systems to ensure maximum visibility into the system.

Total Problem Visibility

When Oracle Solaris is running on Oracle's x86 hardware, it automatically detects the presence of Oracle Integrated Lights Out Manager (Oracle ILOM), the embedded service processor that is part of every server from Oracle. After detecting Oracle ILOM, Oracle Solaris initializes a high-speed connection with Oracle ILOM for use by Oracle Solaris FMA to exchange information related to server topology and fault conditions.

On a typical server, the host operating system and the service processor have mutually exclusive (although sometimes partially overlapping) subsystems that they manage. The host operating system has ownership of the CPU, memory, and I/O subsystems while the service processor presides over the fans, power supplies, DIMMs, and other miscellaneous chassis components. For these reasons, data center managers are often forced to monitor the health of the host operating system and the service processor as if they were separate entities.

Having a dedicated interconnect between Oracle Solaris and Oracle ILOM allows Oracle Solaris FMA to have visibility into the entire system, including the power and cooling subsystems for the server as well as other chassis components. When Oracle ILOM detects that an ambient temperature is too hot or a power supply is failing, Oracle Solaris is informed of these problems. Using the Oracle Solaris FMA interfaces, such as the `fmdump` command-line interface or the SUN-FM-MIB for SNMP traps, system administrators can view all open problems within a system—from CPU, memory, and I/O faults to failures in power and cooling subsystems—in one place.

Because Oracle Solaris on Oracle's x86 systems provides a single view of all problems on a system, data center managers and administrators can depend on this operating system and hardware integration for complete system diagnosis, eliminating the need to connect to multiple management entities.

Oracle Solaris–Aware Service Processor

Oracle designed the high-speed communication path between Oracle Solaris and the Oracle ILOM service processor to be fully bidirectional. This allows Oracle ILOM to be aware of and communicate with Oracle Solaris. In turn, this enables Oracle ILOM to receive fault information diagnosed by Oracle Solaris FMA. Once received by Oracle ILOM, notifications of faults diagnosed by Oracle Solaris can be sent as SNMP traps, e-mails, or `syslog` messages, as well as viewed from the Oracle ILOM Web interface or command-line interface.

Oracle's x86 systems are the only Intel-based systems in which the service processor and Oracle Solaris are tightly integrated. As a result, organizations running Oracle's x86 servers have the flexibility to choose how they want to monitor the servers. Some data center managers prefer to monitor the health of a server through the operating system, while others prefer to use the dedicated management port connected to Oracle ILOM. Either way, the managers are provided with insight into the total server health. This simplifies the management infrastructure by not requiring enterprise management tools to communicate with both the host operating system and the service processor for the purpose of server management.

Component Replacement Made Easy

With the complete lifecycle of a server in mind, Oracle Solaris simplifies finding and replacing components when it is running on Oracle's x86 systems. When a hardware problem occurs, data center managers need to be notified as soon as possible, and they need to figure out what steps to take to resolve the issue. Oracle has engineered Oracle Solaris FMA to provide special serviceability features on Oracle's x86 systems, including the ability to identify a failed component using the same label as the one that is silkscreened onto various locations inside and outside of the chassis. In addition, Oracle Solaris FMA also has the ability to light fault LED indicators to help guide a technician to the server and faulted component. This is a result of the deep level of engineering interaction between the server and the operating system development teams.

Oracle Solaris on Oracle's x86 systems reduces operational expenses by requiring less specialized training to locate faulted components. Printed on the chassis of any x86 server from Oracle are short labels that clearly identify each component or slot. CPUs are labeled CPU n , DIMMs are labeled D n , fan modules are labeled FM n , and so on (where n is the component number). Because Oracle Solaris and Oracle's x86 systems are engineered together, Oracle Solaris is aware of these labels. When a component in any of Oracle's x86 servers fails, Oracle Solaris FMA will diagnose this problem and specify the failed component by the very same label as the one that is printed on the chassis. On a non-Oracle server, Oracle Solaris is not aware of the hardware-specific nomenclature; therefore, it might not be obvious how to locate a failed component.

To further help in identifying and locating faulted components, problems diagnosed by Oracle Solaris result in fault LED indicators being illuminated. This mechanism, which is present only when Oracle Solaris is running on Oracle's x86 systems, takes advantage of the connection with Oracle ILOM. When Oracle Solaris FMA diagnoses a fault, a fault message is sent to Oracle ILOM, which in turn lights the appropriate indicators. This includes lighting an indicator for the specific component that has failed as well as lighting the overall chassis fault indicator.

The ability of Oracle Solaris to light indicators for faulted components on Oracle servers allows for easier and faster service, enabling technicians to quickly restore services that were brought down by the fault. While x86 systems from other hardware vendors also have fault indicators, these indicators cannot indicate problems diagnosed by Oracle Solaris.

On Oracle's x86 systems with Oracle Solaris FMA, there is no guesswork or trial and error to find the right component location. Once a fault message that indicates a faulted component is received, a technician can locate that failed component easily. Any customer-replaceable component that has failed can simply be removed and a new one can be inserted.

Service Automation Through Co-Engineering

Oracle Auto Service Request is a secure, scalable, customer-configurable software solution for Oracle hardware products that resolves problems faster by providing automatic service request generation and priority service request handling for specific faults. If a fault is detected, this software automatically creates a service request with the My Oracle Support portal. This feature helps Oracle's service personnel to dispatch spare parts directly to the data center location, and the customer is notified that a service request has been opened.

Oracle Auto Service Request, which relies on both Oracle Solaris and Oracle ILOM to diagnose server problems, saves organizations money by reducing downtime and improving the efficiency of maintenance operations. This enables IT administrators and solutions developers to concentrate on innovation rather than on the mundane tasks of server maintenance.

Oracle Auto Service Request and Oracle Solaris

Oracle Auto Service Request is the result of collaborative engineering in server hardware, firmware, and Oracle Solaris. Oracle Auto Service Request depends on Oracle Solaris to automatically diagnose problems and send notifications of those problems back to Oracle's support portal. Because of the advanced nature of Oracle Solaris FMA, Oracle Auto Service Request depends on it to make an accurate diagnosis of a problem to determine when a service request should automatically be created.

Another key feature of Oracle Auto Service Request is communicating the part number and serial number of the component that has failed back to Oracle. This feature relies on Oracle Solaris having an understanding of Oracle's x86 server component inventory and access to the part numbers electronically encoded in those components. By automatically placing the part number in the service request to Oracle, Oracle Auto Service Request makes ordering replacement parts very simple.

Special features were engineered into Oracle Solaris to enable this service, such as the `asr-notify` daemon that subscribes to fault events and uses secure Web service protocols to communicate with Oracle's service portal. However, when Oracle Solaris is running on non-Oracle x86 systems, the Oracle Auto Service Request service is not functional.

Organizations using Oracle Auto Service Request benefit from the tremendous amount of engineering by Oracle that went into understanding how servers can fail, how to detect hardware problems, and how to diagnose the problems accurately using Oracle Solaris. While non-Oracle hardware vendors have their own technologies for communicating problems back to their respective support organizations, those technologies are not tightly integrated with Oracle Solaris. Having a set of software and firmware components that have all been engineered and tested together is the only way to ensure accurate diagnosis of problems.

Automated Parts Delivery

Oracle also provides a rapid service request feature to customers by which replacement parts can be sent out automatically when a service request is created. Oracle Solaris FMA, in conjunction with the service processor, automatically determines that a problem has occurred and notifies Oracle. Combining Oracle Auto Service Request with this feature means a problem is found and a replacement part is dispatched to a customer without any human involvement.

By combining Oracle Solaris, Oracle's x86 servers, Oracle Auto Service Request, and rapid service requests, Oracle provides a completely automated diagnosis and repair solution that provides organizations with unparalleled value and savings.

Oracle Solaris FMA and Oracle Auto Service Request on Non-Oracle x86 Systems

When Oracle Solaris is running on non-Oracle x86 systems, the Oracle Solaris FMA feature is still present. It can diagnose a subset of possible problems, including those based off the standard Intel chipset or those that use the same parts as found in Oracle's x86 hardware. However, since complete fault diagnosis depends on having built-in logic around device topology and integration with the service processor, IT staff cannot depend on Oracle Solaris FMA as the only source for identifying field-replaceable component failures. In addition, Oracle Solaris compliance suites do not check for the integrity and correctness of Oracle Solaris FMA on non-Oracle x86 systems.

Likewise, the Oracle Auto Service Request feature is applicable only with Oracle's hardware. When Oracle Solaris is run on non-Oracle x86 systems, the features in Oracle Solaris that connect and notify Oracle about hardware problems are not present.

Figure 2 summarizes the differences in fault and serviceability features between Oracle's x86 systems and non-Oracle x86 systems.

	Oracle x86	Non-Oracle x86
Diagnosis of correctable and uncorrectable CPU and memory errors on Intel Xeon processor-based servers	✓	✓
Single view of all hardware problems on the server	✓	✗
Identification of faulty components using the same name that is printed on the chassis or motherboard	✓	✗
Fault indicator (LED) turned on for component and server that has a problem	✓	✗
Automatically generated service request for diagnosed problems	✓	✗
Validated and quality tested for each new hardware model	✓	✗

Figure 2. Differences in fault and serviceability features between Oracle's x86 systems and non-Oracle x86 systems.

Simplified Installation

Another major area where organizations can benefit from running Oracle Solaris on Oracle's x86 systems, compared to non-Oracle x86 systems, is the area of Oracle Solaris installation. Oracle has engineered Oracle Solaris to be easier to install on Oracle's x86 servers. Simpler installation reduces operational expenses by reducing the time required to bring a server into production and improving the efficiency of IT staff.

Identification of Disks

When installing Oracle Solaris on a local disk, the user must select a target disk drive. Oracle has simplified this task on Oracle's x86 systems by allowing drives to be identified by the drive bay name that is printed on the outside of the server chassis. For non-Oracle x86 systems, disk drives are identified only by their worldwide ID (WWID).

It is often difficult to determine the physical location of a drive by the WWID, especially since servers are being designed with a large number of disk drives. This feature alleviates the problem of having to figure out the WWID for a given disk prior to installation, which is especially beneficial in the case of a new installation where no operating system has been installed yet to interrogate the disks for the identifier.

Availability of Drivers

Oracle ensures that all Oracle Solaris drivers for its x86 server line are available and tested at the time of release of a new server. Oracle develops Oracle Solaris in conjunction with its servers, writes drivers specifically for devices in its platforms, and ensures that the bundle of drivers and devices are tested together. When running Oracle Solaris on Oracle's x86 systems, system administrators need not hunt for drivers nor be concerned about using an untested combination of software and hardware.

Preinstall Option

As an additional option for organizations, Oracle will preinstall Oracle Solaris on its x86 systems at the time of manufacturing the server. The preinstalled image is always up to date with the latest Support Repository Updates (SRUs) and patches. This removes the burden of having to perform the installation for each server as well as ensuring that the image has all the latest updates. Although not all data center managers elect to use this option, those that do benefit from speedier out-of-box setup times and guaranteed compatibility of the installed software with the given server and its installed options.

Oracle Enterprise Manager Ops Center

Oracle Enterprise Manager Ops Center 12c is an enterprise management tool that allows IT staff to manage all aspects of their servers. In addition to providing detailed hardware monitoring and reporting for hardware problems, Oracle Enterprise Manager Ops Center can provision a bare metal system with an operating system and also configure virtualization. Oracle Enterprise Manager Ops Center can install and configure Oracle Solaris as well as set up Oracle Solaris Zones for the server, all from a single management console.

Because Oracle Enterprise Manager Ops Center 12c was engineered to work best with Oracle's servers, it provides a unique single-pane management view of Oracle's x86 systems. Although Oracle Enterprise Manager Ops Center can install Oracle Solaris onto non-Oracle x86 systems, it provides limited discovery, inventory, and hardware monitoring of those systems. As a result, some operations that might otherwise be automated can require manual intervention. Using Oracle Enterprise Manager Ops Center with Oracle's x86 systems and Oracle Solaris provides more features and better reliability, because it builds on special Oracle hardware interfaces that have all been tested to work together.

Oracle's x86 Systems

Oracle's x86 systems are the best x86 platforms for Oracle software. Only Oracle provides customers with an optimized hardware and software stack that comes complete with choice of OS, virtualization software, and cloud management tools—all at no extra charge. Oracle's optimized hardware and software stack has enabled a 10x performance gain in its engineered systems and has delivered world-record benchmark results. Oracle's comprehensive, open standards-based x86 systems provide the best platform to run Oracle software with enhanced reliability for data center environments.

Oracle has a wide range of Intel Xeon processor-based servers to suit the needs of enterprise-class applications that are required to conduct day-to-day operations. These systems range from two- to eight-socket servers. Figure 3 lists Oracle's x86 systems.

Product Name	Features	Primary Application
Sun Blade X3-2B	Two-socket blade module based on Intel Xeon processor E5-2600 product family	Compute node for clustered computing and virtualization needs in a blade form factor
Sun Server X3-2	Two-socket 1RU server based on Intel Xeon processor E5-2600 product family	Compute node for clustered computing and virtualization needs
Sun Server X3-2L	Two-socket 2RU server based on Intel Xeon processor E5-2600 product family	Storage node for large internal storage capacity needs
Sun Server X2-4	Four-socket 3RU server based on Intel Xeon processor E7-4800 product family	Consolidation node for efficient use of resources with high-performance processors and large memory availability
Sun Server X2-8	Eight-socket 5RU server based on Intel Xeon processor E7-8800 product family	SMP node for scale-up computing

Figure 3. Oracle's x86 systems.

More information about the Oracle x86 systems can be found at:

<http://www.oracle.com/us/products/servers-storage/servers/x86>

More For Less

Over and above the aforementioned advantages of running Oracle Solaris on Oracle's x86 systems, a further benefit is the cost savings for licensing and support. In addition, using Oracle's x86 systems will automatically ensure compliance with Oracle Solaris licensing terms, freeing up IT resources from having to track which non-Oracle systems have active subscriptions and which do not. Organizations will pay more to run Oracle Solaris on non-Oracle x86 systems.

When running Oracle Solaris on non-Oracle x86 systems, the cost of each subscription is US\$1000 annually per socket on a system with up to four sockets and US\$2000 annually per socket on larger systems. Additionally, organizations would need to purchase supplementary hardware support from the server vendor. In contrast, Oracle Premier Support provides for comprehensive coverage and a single point of contact for Oracle's x86 systems and Oracle Solaris. There are no additional subscription charges for running Oracle Solaris on Oracle's x86 systems.

When calculating the total cost of ownership of a server, the licensing costs must be factored in, which could contribute significant costs for an Oracle Solaris environment on a non-Oracle server. Using Oracle's x86 systems as the platform on which to run Oracle Solaris-based applications, organizations can avoid the annual Oracle Solaris subscription fees, which can be up to US\$16,000 per year for an eight-socket server. Considering that most enterprise-class data centers use multiple servers, sometimes scaling to thousands of nodes, these subscription costs multiply very quickly, forming a large part of an IT budget. In essence, why pay more and get less?

Conclusion

Oracle Solaris runs better and smarter on Oracle's x86 hardware. By engineering Oracle Solaris to work with Oracle's x86 hardware, Oracle has engineered unique features that bring value to organizations that are looking to run their mission-critical applications. The advanced fault management and serviceability features, as well as improvements to installation and system management, make Oracle's x86 systems the best hardware choice for running Oracle Solaris.

There are also many enhancements to Oracle Solaris currently in development that will add even more advantages to running Oracle Solaris on Oracle's x86 systems in the future. Oracle is committed to engineering both Oracle Solaris and its x86 systems to provide unmatched performance, reliability, and manageability through innovative engineering.



**Oracle's x86 Systems: The Best x86
Platforms for Oracle Solaris**

February 2013

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

Copyright © 2013, Oracle and/or its affiliates. All rights reserved. This document is provided for information purposes only and the contents hereof are subject to change without notice. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group. 1012

Hardware and Software, Engineered to Work Together