

IBM Dynamic System Analysis

IBM Systems and Technology Group

Introduction

Today's shrinking administration budgets and staffs mean that not only must x86 servers be highly available but diagnostic tools must make administrators more efficient in preventing server downtime.

IBM® Dynamic System Analysis (DSA) and Dynamic System Analysis Preboot provide business managers with the tools they need to diagnose and solve problems before they occur and are an integral part of the System x[™] systems-management software suite.

DSA provides business managers and IBM support personnel with a single tool to collect configuration and status information and logs and to perform diagnostics on System x servers. It also provides a set of tools and analysis routines that can help pinpoint problems and return a server to a working state more quickly.

The DSA editions

IBM provides the following editions of DSA:

Online DSA runs within the operating system; you do not have to reboot the server to run it. It performs comprehensive data collection of hardware and operating-system information and includes a set of diagnostics that are safe to run while the operating system is running. There are two editions of Online DSA:

- **DSA Portable Edition** is packaged as a self-extracting file that you download from the Web. When you run the file, it self-extracts to a temporary folder and runs Online DSA. After it runs, it automatically deletes the temporary files and folder and leaves the results of the data collection and diagnostics on the server.
- **DSA Installable Edition** is packaged as an InstallShield (Windows®) or RPM (Linux®) file that you download from the Web. When you run the file, it installs DSA on the server. You can then run Online DSA from where it is installed.

Bootable DSA is run from a bootable CD; you must reboot the server to run it. Bootable DSA can perform comprehensive data collection from the hardware and includes the same set of diagnostics as Online DSA. It is provided as an ISO image that you can download from the Web and copy to a CD.

DSA Preboot runs outside of the operating system; you must reboot the server to run it. In addition to all of the capabilities of Online and Bootable DSA, it includes many more diagnostic routines that would be disruptive to run within the operating-system environment (such as resetting devices and causing loss of network connectivity). It also has a graphical user interface (GUI) that you can use to specify which diagnostics to run and to view the diagnostic and data collection results. It is provided as an ISO image or embedded in flash memory in the server.

The benefits of DSA

DSA provides the capability to collect and transmit system data that is relevant to service issues directly to IBM service and support, or to save and transmit this system data as a small compressed file, anywhere that you deem appropriate. This allows for individual and aggregate analysis of your system data, using a post-processing model that includes inventory management, proactive maintenance, environment trend analysis, and on demand service examination.

DSA provides a consistent framework to view and manipulate the system data that is most useful for servicing your server. DSA provides not only hardware-failure indications and associated repair actions but also configuration-error notification and software-version verification.

Online DSA is a tool that you can use to diagnose and solve problems without taking the server offline.

Bootable DSA and DSA Preboot are tools that you can use to diagnose and solve problems with the server offline.

You can run DSA Portable Edition online without installing any software.

DSA is able to diagnose problems with microprocessors, hard disk drives, memory, Ethernet, and tape drives and suggest replacement parts to correct the problems.

DSA performs all these functions with no application prerequisites and no need for a system reboot or reset.

In general, DSA reduces administration costs and supports greater server availability through faster problem resolution.

Improvements in DSA Preboot over IBM Enhanced Diagnostics

Internal IBM test results show a 70% increase in FRU isolation by DSA Preboot over that by IBM Enhanced Diagnostics.

DSA Preboot has all of the comprehensive data collection and data analysis capabilities of Online DSA:

- Baseboard management controller logs
- Critical-details log
- Environmental sensors (fan speeds, temperatures)
- Light path diagnostics status
- Machine-check registers
- Operating-system logs
- PCI slot information
- RAID configuration and logs
- Remote Supervisor Adapter logs
- SMART data
- System code levels
- System configuration information
- Tape drive
- USB
- Video and monitor configuration

In addition, DSA Preboot has diagnostic and data-collection capabilities for the following components and functions:

- Checkpoint panel
- Video memory test

DSA Preboot also has a feature for saving and sending log files to IBM service and support.

Data analysis

DSA provides the following analysis capabilities in addition to diagnostic tests, stress tests, and data collections capabilities:

- A view that shows for each I/O device its corresponding device drivers and slot information
- Translation of log sense pages for supported tape drives
- Detection of missing Remote Supervisor Adapter and IPMI device drivers (Online DSA only)
- Recommended actions for solving ServeRAID[™] problems
- Recommended actions for solving problems that are reported by light path diagnostics
- Memory Predictive Failure Analysis®: single-bit and double-bit error detection
- Validation of the configuration and cabling for scalability/RXE

Critical details

DSA highlights certain error conditions in the server that the field expertise of the service organization has determined are critical to problem determination. The following items are displayed as critical details:

- A Remote Supervisor Adapter or Remote Supervisor Adapter II is installed, but the correct device drivers are not installed (Online DSA only).
- A BMC service processor is installed, but the correct device drivers are not installed (Online DSA only).
- An ISMP service processor is installed, but the correct device drivers are not installed (Online DSA only).
- Light path diagnostics LEDs are lit.
- Unsupported physical disks are in the RAID array.
- A ServeRAID 8i, 8k, or 7t adapter is installed, but ServeRAID Manager is not installed (Online DSA only).
- ServeRAID Predictive Failure Analysis errors, parity errors, or hard errors are in the log.
- RAID device-driver messages are in the Windows event log (Online DSA only).
- ServeRAID firmware, BIOS code, and device-driver versions for the ServeRAID 8k adapter are mismatched.
- ServeRAID data scrubbing is disabled.
- The ServeRAID configuration has a bad stripe count.

Which DSA edition to use

Follow these general guidelines to determine which DSA edition to use:

- If the server will boot, use Online DSA.
- If the server will not boot, or if comprehensive diagnostics are needed, use DSA Preboot, if the server supports it.
- Otherwise, use Bootable DSA.

The capabilities of each DSA edition are summarized in the following table.

	Online DSA	Bootable DSA	DSA Preboot	
Diagnostic	Has limited diagnostic	Has limited diagnostic	Has thorough diagnostic	
capabilities	capabilities	capabilities	capabilities	
Use mode	Portable: Runs from a temporary folder and then is deleted Installable: Runs from installed files on the server	Is a self-contained bootable image	Is a self-contained bootable image	
Operating system	Runs within the operating system. Is supported in Windows 2000; Windows 2003; Red Hat® Enterprise Linux 3, 4, and 5; and SUSE Linux Enterprise Server 9 and 10.	Can be booted in the absence of a supported, installed, or working operating system	Can be booted in the absence of a supported, installed, or working operating system	
Device drivers	Requires that the correct device drivers be installed in the server for any devices that it tests	Contains all the device drivers that are required by the devices that it tests	Contains all the device drivers that are required by the devices that it tests	
Data collection from the operating system	Can collect information from the operating system	Cannot collect information from the operating system, such as operating-system logs, network settings, or installed applications, system components, hotfixes, device drivers, or services	Cannot collect information from the operating system, such as operating-system logs, network settings, or installed applications, system components, hotfixes, device drivers, or services	
System support	Is supported on all IBM- warranted servers	Is supported on all IBM- warranted servers	Is supported only on IBM System x3950 M2 (Athena) servers (as of December 2007)	

The DSA Preboot startup process

At the beginning of the DSA Preboot startup process, the system BIOS loads the DSA Preboot files from the embedded USB key. This screen is displayed for approximately 4 seconds.

Note: The total startup time for DSA Preboot is currently ~10 minutes as measured on the x3850M2/x3950M2. Enhancements are planned for 1Q08 to decrease the startup time for DSA Preboot.

Next, the device drivers are loaded and installed.

PCI: Bridge: 0000:00:11.0 IO window: 2000-2fff MEM window: e8900000-e89fffff PREFETCH window: 50000000-500fffff apm: BIOS version 1.2 Flags 0x03 (Driver version 1.16ac) highmem bounce pool size: 64 pages Total HugeTLB memory allocated, 0 VFS: Disk quotas dquot_6.5.1 Dquot-cache hash table entries: 1024 (order 0, 4096 bytes) Initializing Cryptographic API io scheduler noop registered (default) io scheduler anticipatory registered io scheduler deadline registered io scheduler cfq registered Limiting direct PCI/PCI transfers. uesafb: framebuffer at 0xf0000000, mapped to 0xf8880000, using 6144k, total 1638 4k vesafb: mode is 1024x768x16, linelength=2048, pages=0 vesafb: protected mode interface info at 00ff:44f0 Jesafb: scrolling: redraw uesalb: scrolling: realaw uesafb: Truecolor: size=0:5:6:5, shift=0:11:5:0 bootsplash: scanning last 2MB of initrd for signature bootsplash 3.1.6-2004/03/31: looking for picture...<6>...no good signature found Console: switching to colour frame buffer device 128x48 fb0: VESA VGA frame buffer device isapnp: Scanning for PnP cards.. isapnp: Scanning for PnP cards... isapnp: No Plug & Play device found Real Time Clock Driver v1.12ac PNP: No PS/2 controller found. Probing ports directly. serio: i8042 AUX port at 0x60,0x64 irg 12 serio: i8042 KBD port at 0x60,0x64 irg 1 Serial: 8250/16550 driver \$Revision: 1.90 \$ 4 ports, IRQ sharing enabled serial8250: ttyS0 at I/0 0x3f8 (irg = 4) is a 16550A serial8250: ttyS1 at I/0 0x2f8 (irg = 3) is a 16550A Floppy drive(s): fd0 is 1.44M FDC 0 is a post-1991 82077 RAMDISK driver initialized: 16 RAM disks of 307200K size 1024 blocksize loop: loaded (max 8 devices) Kernel Build: 2:dsa21:mcp:2007:08:16:215505:5:0:0: Uniform Multi-Platform E-IDE driver Revision: 7.00alpha2 ide: Assuming 33MHz system bus speed for PIO modes; override with idebus=xx PIIX4: IDE controller at PCI slot 0000:00:07.1 PIIX4: chipset revision 1 PIIX4: not 100% native mode: will probe irgs later ide1: BM-DMA at 0x1058-0x105f, BIOS settings: hdc:DMA, hdd:pio

After the device drivers are loaded, DSA Preboot is expanded and run. A menu is displayed from which you can select to run DSA Preboot in text mode from the console (by typing cmd) or to start the GUI (by typing gui).

Starting IBM DSA Preboot v1.00 ...
Extracting...
gui = Enter GUI Environment.
copy = Copy DSA presults to removable media.
exit = Quit program.
 Note: This viil reboot the system.
help = Display this help message.
Please enter a command. (Type 'help' for commands)
>

From the GUI, you can run diagnostics on the server or collect the system configuration, status, and logs. Then, you can transmit the diagnostics results or collected data to IBM service over the Internet or copy the data to a removable medium.

IBM Dynamic System Analysis	
Hardware Diagnostic Utility	
Click on the desired section.	
Diagnostics Run diagnostic tests on system components	
System Information Show detailed information about your system.	
Send System Information to IBM Collect and send the system information to IBM for analysis.	
Copy System Information to Local Media Collect and copy the system information to local media for later analysis.	
Help Learn more about the features of DSA.	
License View the DSA license agreement.	
Exit Quit this utility.	

The following illustration shows an example of the System Information screen after **Hardware** is selected from the options in the left pane.

System H Network Info Hardware	ardware Inv							
Network Info		entorv						
Hartware		entory						
	rocessor/Cor	е		CBULO	CRU1	CRU 2	CRU2	
BCUpfo	ame Ianufacturer			GenuineIntel		CFU2	CFUS	
	ersion			Pentium(R) Pro	0	0	0	
Firmware	tepping			<86 Family 6, Model 14,				
Service Processor Config				Stepping 8 (6e8)				
Environmentals	urrent Clock Speed tatus			2133 CPU Enabled	Unnonulated	U	U	
	uper-threading Supp	orted		false	false	false	false	
LSI	yper-threading Statu	\$	i	Not Enabled	Unknown	Unknown	Unknown	
Lightpath	evel 1 Cache Operati	onal Mode		Write Back				
	evel 1 Cache Enabled	(at boot tim	e)	true Italita Baali		_		
SP Built-In Self Test	evel 2 Cache Operation	(at boot time	e)	hrue				
SP Logs	evel 3 Cache Operati	onal Mode	-,					
Merged Devices	evel 3 Cache Enabled	(at boot tim	e)					
Diagnostic Log	lemorv							
Diagnostic Log	ag	RAM slot #0	RAM slot	#1 RAM slot #2	RAM slot #3			
DSA Error Log	ank	RAM slot #0	RAM slot	#1 RAM slot #2	RAM slot #3			
Home	apacity	1,024 Megabytes	0 Bytes	0 Bytes	O Bytes			
M	lemory Type	DRAM	Unpopula	ted Unpopulated	Unpopulated			
S S	peed (Mhz)	Unknown	Unknown	Unknown	Unknown			
	rrorCorrectionType	None	Unknown	Unknown	Unknown			
	lanufacturer							
	artnumber							
	ersion							
				Status	Siz	e	Number Of Partitions	
1d	lev/hdc					12	1.3 MB	0
N	ame	Туре			itatus	Siz	e	Number Of Partitions
/6	lev/sda					80	iB	0
S	vstem Card I	nformati	on					
N	ame	Туре			Manufacturer	Ve	rsion	Serial Number
44	40BX Desktop Referen	ice Moth	er Board		ntel Corporation	n No	ne	None
	ation							

Diagnostics

This section briefly describes each of the tests that are available in DSA Preboot.

CPU/coprocessor stress test

This test stresses the microprocessor by performing arithmetic operations with whole numbers and floating-point numbers and validating the results.

Storage subsystem

The following storage subsystem tests are available.

Disk self-test (DST)

This test performs the industry-standard diagnostics that are provided by most hard disk drive vendors.

Disk stress test

This test verifies successful read operation from various locations on the disk.

Optical drive

Verify media installed

This test identifies the CD-ROM drives, checks the status of each drive, and checks for the presence of media. If media is present, a "Pass" result is recorded in the log. If media is not present, a "No media found" result is recorded in the log, and the test ends.

Drive self-test

This test identifies the CD-ROM drives, checks the status of each drive, and checks for the presence of media. If media is present, the diagnostic runs, and the result is recorded in the log. If media is not present, a "No media found" result is recorded in the log.

Read-error test

This test identifies the CD-ROM drives, checks the status of each drive, and checks for the presence of media. If media is present, it verifies that the drive can read from the disk, and the result is recorded in the log. If media is not present, a "No media found" result is recorded in the log, and the test ends.

Tape drive

The following tape drive tests are available.

Tape presence test

This test performs a load/unload command on the drive. If a tape cannot be loaded or is not present, a "Failed" result is recorded in the log.

Read/write self-test

The function of this test varies depending on the tape drive model and vendor. Typically, it performs an internal verification of the drive microcode and a read/write/verify operation on the tape media.

Tape alert self-test

This test determines whether a Predictive Failure Analysis alert has been detected. During normal operation, most tape drives monitor their own operation, and if they detect impending problems, they set a flag.

Broadcom

The following tests are supplied by Broadcom.

MAC registers test

This test exercises most internal registers in a chip. Each chip contains many internal registers, and each register can contain some read-only bits and read/write bits. Read-only bits that cannot be overwritten and read/write bits that can be written properly are tested. If all registers pass the test, a "Pass" result is recorded in the log. If any register fails the test, a "Fail" result is recorded in the log, and the test ends. This test causes the network adapter to go offline.

PHY register test

This test exercises most MII registers (in copper media PHY) in a chip. Each chip contains many MII registers, and each register can contain some read-only bits and read/write bits. Read-only bits that cannot be overwritten and read/write bits that can be written properly are tested. If all registers pass the test, a "Pass" result is recorded in the log. If any register fails the test, a "Fail" result is recorded in the log, and the test ends. This test causes the network adapter to go offline.

NVRAM test

This test verifies the firmware checksum, including the bootstrap record and configuration block in the EEPROM chip. If both checksums are valid, a "Pass" result is recorded in the log. If any checksum fails the test, a "Fail" result is recorded in the log, and the test ends.

Internal memory test

This test exercises most of the internal memory in a chip. Each chip contains memory to support on-chip microprocessor operation. The memory is tested a double-word at a time. The test writes various patterns and reads them back to verify that they were written correctly. If all memory passes the test, a "Pass" result is recorded in the log. If any memory fails the test, a "Fail" result is recorded in the log, and the test ends. This test causes the network adapter to go offline.

Interrupt test

This test verifies the interrupt capability of a chip. It clears all interrupt statuses, triggers the network adapter to generate an interrupt, and waits for an interrupt to come into the driver interrupt handler. If the interrupt occurs before a timeout, a "Pass" result is recorded in the log. If the interrupt does not occur before a timeout, a "Fail" result is recorded in the log, and the test ends. This test causes the network adapter to go offline.

MAC loopback test

This test performs a MAC layer loopback test. It sends an Ethernet packet to the network adapter and waits for a packet to be returned. If the packet is returned before a timeout, a "Pass" result is recorded in the log. If the packet is not returned before a timeout, a "Fail" result is recorded in the log, and the test ends.

PHY loopback test

This test performs a PHY layer loopback test. It sends an Ethernet packet to the network adapter and waits for a packet to be returned. If the packet is returned before a timeout, a "Pass" result is recorded in the log. If the packet is not returned before a timeout, a "Fail" result is recorded in the log, and the test ends.

LED test

This test flashes the LEDs on the network adapter. If the LEDs flash correctly, a "Pass" result is recorded in the log. If the LEDs do not flash correctly, a "Fail" result is recorded in the log.

Memory

To test as much available memory as possible, the memory test runs in a stand-alone environment that is separate from the main DSA Preboot application. There is no memory test in the main DSA Preboot application. After you press F2 during system startup, the memory test is loaded, and a text-based menu is displayed from which you can select a memory-test option, view results, or exit to the DSA Preboot main menu. If you take no action, after a short delay, a quick memory test runs, and then the DSA Preboot main menu is displayed. After the main DSA Preboot application starts, you must restart the server to re-enter the stand-alone memory-test environment.

Memory test

This tests the entire installed system memory. It writes complex data patterns in most of the available memory and then reads and verifies the data. The memory area that is reserved to hold the diagnostic code is not written; however, to flush out ECC errors in reserved memory, the diagnostic reads these reserved memory locations.

Memory stress test

This test focuses on fixed-pattern, linear-pattern, random-pattern, walking-bus-pattern, and toggle-bus-pattern tests.

Systems-management hardware

The following systems-management hardware tests are available.

Remote Supervisor Adapter restart test

This test restarts the service processor, at which time the service processor runs a built-in selftest (BIST). If the device passes the test, a "Pass" result is recorded in the log. If the device fails the test, a "Fail" result is recorded in the log.

BMC I2C test

This test returns the statuses of the various I2C buses in the server. If a bus passes the test, a "Pass" result is recorded in the log. If a bus fails the test, a "Fail" result is recorded in the log, with a list of devices that might be causing the problem.

Security subsystem

This test verifies that the hardware can be initialized correctly and that the Trusted Platform Module is enabled and configured to protect the system.

Checkpoint panel test

This test verifies that the diagnostic checkpoint panel display segments light correctly.

EXA port ping test

This test sends test packets between ports in multi-node configurations and verifies the transmissions.



For More Information

IBM System x and xSeries Servers	ibm.com/systems/x			
IBM BladeCenter Server and options	ibm.com/systems/bladecenter			
IBM System x and BladeCenter Power Configurator	ibm.com/systems/bladecenter/powerconfig			
IBM Standalone Solutions Configuration Tool (SSCT)	ibm.com/servers/eserver/xseries/library/configtools.html			
IBM Electronic Service Agent	ibm.com/support/electronic			
IBM ServerProven Program	ibm.com/servers/eserver/serverproven/compat/us			
IBM Technical Support	ibm.com/server/support			

Legal Information

© IBM Corporation 2007 IBM Systems and Technology Group Dept. 4P3A 3039 Cornwallis Road Research Triangle Park, NC 27709

Produced in the USA December 2007 All rights reserved.

For a copy of applicable product warranties, write to: Warranty Information, P.O. Box 12195, RTP, NC 27709, Attn: Dept. JDJA/B203. IBM makes no representation or warranty regarding third-party products or services including those designated as ServerProven or ClusterProven. Telephone support may be subject to additional charges. For onsite labor, IBM will attempt to diagnose and resolve the problem remotely before sending a technician.

IBM, the IBM logo, the e-business logo, Active Memory, iSeries, PowerExecutive, Predictive Failure Analysis, ServeRAID, System i, System Storage, System x, TotalStorage, xSeries and Xtended Design Architecture are trademarks of IBM Corporation in the United States and/or other countries. For a list of additional IBM trademarks, please see http://ibm.com/legal/copytrade.shtml.

InfiniBand is a trademark of InfiniBand Trade Association.

Intel, and Intel Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Linux is a registered trademark of Linus Torvalds.

Red Hat, the Red Hat "Shadow Man" logo, and all Red Hatbased trademarks and logos are trademarks or registered trademarks of Red Hat, Inc., in the United States and other countries.

Lotus, Domino, Notes and SmartSuite are trademarks or registered trademarks of Lotus Development Corporation and/or IBM Corporation.

Microsoft, Windows, and SQL Server are trademarks or registered trademarks of Microsoft Corporation.

TPC Benchmark is a trademark of the Transaction Processing Performance Council.

UNIX is a registered trademark in the U.S. and/or other countries licensed exclusively through The Open Group.

Other company, product and service names may be trademarks or service marks of others.

IBM reserves the right to change specifications or other product information without notice. References in this

publication to IBM products or services do not imply that IBM intends to make them available in all countries in which IBM operates. IBM PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS

OR IMPLIED, INCLUDING THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions; therefore, this statement may not apply to you.

This publication may contain links to third party sites that are not under the control of or maintained by IBM. Access to any such third party site is at the user's own risk and IBM is not responsible for the accuracy or reliability of any information, data, opinions, advice or statements made on these sites. IBM provides these links merely as a convenience and the inclusion of such links does not imply an endorsement.

Information in this presentation concerning non-IBM products was obtained from the suppliers of these products, published announcement material or other publicly available sources. IBM has not tested these products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

MB, GB and TB = 1,000,000, 1,000,000,000 and 1,000,000,000 bytes, respectively, when referring to storage capacity. Accessible capacity is less; up to 3GB is used in service partition. Actual storage capacity will vary based upon many factors and may be less than stated.

Performance is in Internal Throughput Rate (ITR) ratio based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput that any user will experience will depend on considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput improvements equivalent to the performance ratios stated here.

Maximum internal hard disk and memory capacities may require the replacement of any standard hard drives and/or memory and the population of all hard disk bays and memory slots with the largest currently supported drives available. When referring to variable speed CD-ROMs, CD-Rs, CD-RWs and DVDs, actual playback speed will vary and is often less than the maximum possible.