

Monitoring PowerEdge™ Servers in Solaris™ 10 Environment

Framework for Monitoring
Dell™ PowerEdge Servers
running Solaris 10

By Ahmad Ali
Dell | Enterprise Operating Systems



8/18/2008

THIS WHITE PAPER IS FOR INFORMATIONAL PURPOSES ONLY, AND MAY CONTAIN TYPOGRAPHICAL ERRORS AND TECHNICAL INACCURACIES. THE CONTENT IS PROVIDED AS IS, WITHOUT EXPRESS OR IMPLIED WARRANTIES OF ANY KIND.

Dell, the *Dell* logo, *OpenManage*, and *PowerEdge* are trademarks of Dell Inc; *Solaris* and *Sun* are registered trademarks of Sun Microsystems Inc., in the United States and other countries.

Other trademarks and trade names may be used in this document to refer to either the entities claiming the marks and names or their products. Dell disclaims proprietary interest in the marks and names of others.

©Copyright 2008 Dell Inc. All rights reserved. Reproduction in any manner whatsoever without the express written permission of Dell Inc. is strictly forbidden. For more information, contact Dell.

THE INFORMATION IN THIS DOCUMENT IS SUBJECT TO CHANGE WITHOUT NOTICE AND IS PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND. THE ENTIRE RISK ARISING OUT OF THIS INFORMATION REMAINS WITH THE USER OF THE INFORMATION. IN NO EVENT SHALL DELL BE LIABLE FOR ANY DIRECT, CONSEQUENTIAL, INCIDENTAL, SPECIAL, PUNITIVE OR OTHER DAMAGES, EVEN IF DELL HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

TABLE OF CONTENTS

INTRODUCTION	4
STATUS COLLECTION	5
SYSTEM COMPONENT STATUS	5
INTERNAL RAID STORAGE STATUS	5
SAS 6/iR Status.....	6
PERC 6/i Status.....	6
STATUS NOTIFICATION	7
SENDING SNMP ALERTS	7
EVENT ENTRIES IN SYSLOG.....	8
SENDING EMAIL MESSAGES	8
AN EXAMPLE SCRIPT	8
CONCLUSION	9
REFERENCES	10

INTRODUCTION

Monitoring servers is an unquestionable necessity in any enterprise environment. The Dell OpenManage™ stack¹ provides multiple options for monitoring PowerEdge servers; including both out-of-band and in-band status monitoring options. Dell Remote Access Controllers² (DRACs) and Baseboard Management Controller³ (BMC) provide capabilities for automatically monitoring server status remotely using industry standard SNMP protocol or email messages. These monitoring methods^{4 5} are operating system (OS) agnostic. In-band management and monitoring solutions provided by the OpenManage suite of applications are OS-dependent and currently not supported on Dell PowerEdge servers running the Sun® Solaris Version 10 operating system. This creates a gap in Solaris 10 environments where in-band monitoring methods are implemented; furthermore, DRAC and BMC do not provide monitoring capabilities for internal PERC 6/i, or SAS 6/iR, storage subsystems. This lack of monitoring capability creates a gap for internal storage subsystem monitoring on PowerEdge servers running Solaris 10.

This paper describes a framework for in-band monitoring of both PowerEdge system components and internal storage subsystems in a Solaris 10 environment using tools readily available. This framework, as seen in Figure-1, utilizes native Solaris tools (`ipmitool`, `raidctl`) and the LSI MegaCli tool for collecting system and internal storage subsystems status. Status notification is performed using SNMP alerts, email messages or syslog entries, where syslog is monitored for events. The paper also identifies commands for monitoring specific components, and provides examples of scriptable notifications.

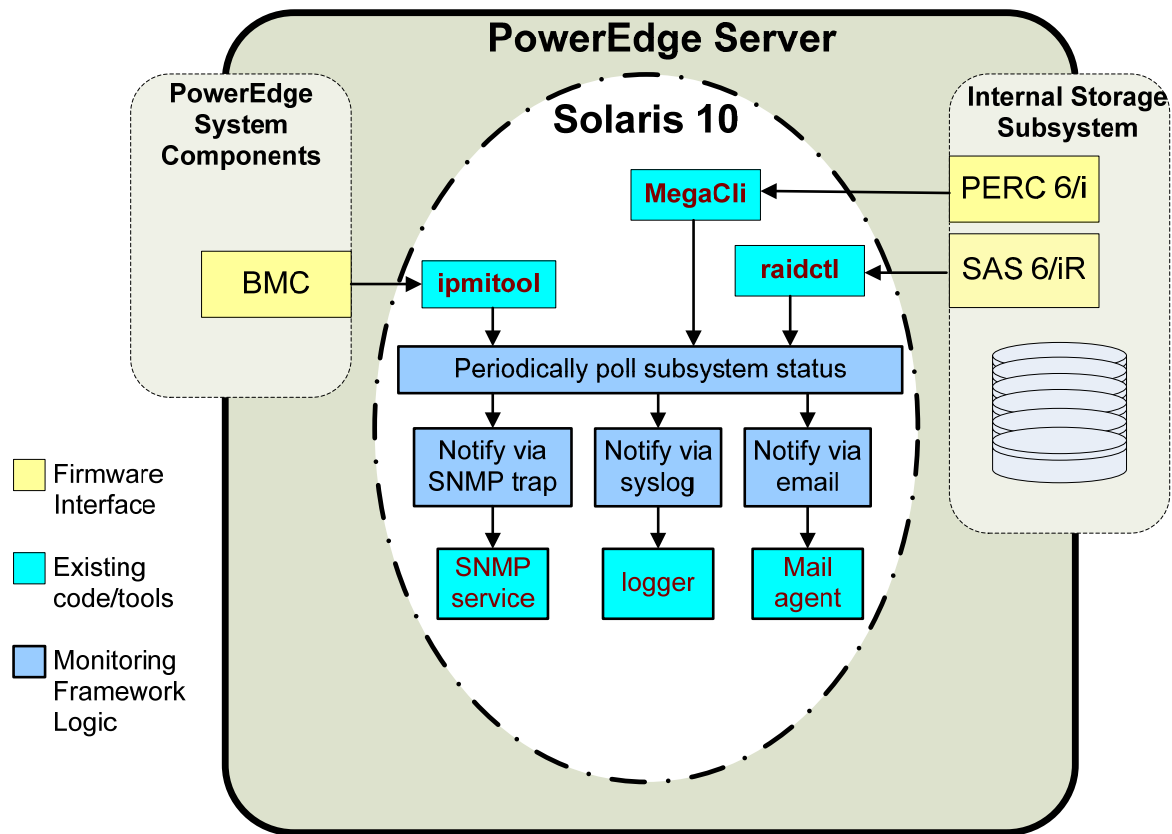


Figure 1: Proposed framework for in-band monitoring in a Solaris 10 environment

The audience for this paper should be familiar with the following areas since they will implement this solution: system management and monitoring functions, scripting, SNMP and MAIL tools.

STATUS COLLECTION

All PowerEdge servers supporting Solaris 10 have a Baseboard Management Controller (BMC) that exposes an industry-standard IPMI interface to the operating system. Solaris10 includes a device driver that works with this interface. The `ipmitool`, packaged with the Solaris 10 operating system, can interact with the BMC and collect the status information for system components like fan RPM, temperature probe reading and memory ECC errors. Similarly, `raidctl` is included natively with the operating system and can be used to collect the status of internal RAID storage attached to the Dell SAS 6/iR controller. In addition, `MegaCli` from LSI must be installed in order to collect status information for internal storage attached via the Dell PERC 6/i controller.

SYSTEM COMPONENT STATUS

This section provides examples for using the `ipmitool` to poll the status of system components via the BMC.

The following command reads the BMC Sensor Data Repository (`sdr`),

```
# ipmitool sdr list
```

The command output displays a '|' separated table for all sensors checked - one row per sensor. Each row lists sensor name, value and status. The table displays 'Not Readable' as value for sensors that are not implemented, and displays 'ok' when a sensor value is within set threshold/limits.

The following command is a simple way to get list of sensors that need attention.

```
# ipmitool sdr list |egrep -v "Not Readable|//| ok"
```

Following is example output from a system with a fan failure,

```
FAN 3 RPM          | 0 RPM          | cr
```

The command output shows that FAN3 is not running and its status is critical. This output can simply be used for notification, or more analysis can be done by capturing it. Multiple `sdr` commands can be executed quickly by first caching the repository, and then looking for different subsystems, as in following example,

```
# ipmitool sdr dump <sdr-cache>
# ipmitool -S <sdr-cache> sdr list full
# ipmitool -S <sdr-cache> sdr list fru
```

Consult the `ipmitool` documentation for additional commands. `IPMITool` raw commands can be used to take full advantage of the BMC capabilities⁶.

INTERNAL RAID STORAGE STATUS

Dell PowerEdge servers supporting Solaris 10 have two internal RAID storage options - SAS 6/iR and PERC 6/i. The status of RAID volumes attached to SAS 6/iR can be checked using the `raidctl`

command and the LSI `MegaCli` application respectively. This section provides examples for using these commands to check internal RAID storage status.

SAS 6/iR Status

The following command provides the status of a Dell SAS 6/iR RAID storage subsystem:

```
# raidctl -S
```

The command provides the status of both the RAID volumes and the disks attached to the controller. RAID volumes are identified as `cxydz`, while physical disks are identified as `x.y.z` where `x`, `y` and `z` represent controller number, target ID, and LUN number. For RAID volume, the status can be **OPTIMAL** (operating optimally), **DEGRADED** (operating with reduced functionality), **FAILED** (non-functional), or **SYNC** (disks are syncing). For a physical disk, the status can have a value of **GOOD**, **FAILED**, **MISSING**, or **HSP** (hot spare).

Example output is from this command is below:

```
0 "LSI_1068E"  
c0t1d0 2 0.9.0 0.2.0 1 OPTIMAL  
c0t3d0 2 0.10.0 0.4.0 1 OPTIMAL  
  
0.0.0 GOOD  
0.2.0 GOOD  
0.4.0 GOOD  
0.9.0 GOOD  
0.10.0 GOOD
```

The following command is a simple way of identifying all RAID volumes attached to a SAS controller that are not in an **OPTIMAL** state:

```
# raidctl -S | egrep c[0-9]+t[0-9]+d[0-9] | egrep -v OPTIMAL
```

The following is example output that shows a RAID volume on a SAS 6/iR in a non-OPTIMAL state,

```
c0t6d0 2 0.10.0 0.7.0 1 SYNC
```

The output can be used for notification; check `raidctl` documentation for more command options.

PERC 6/i Status

LSI `MegaCli` provides multiple commands to check the status of RAID volumes attached to a PERC 6/i controller. The number of RAID volumes configured on a PERC 6/i, properties and status for each volume, and the information and status for constituent member disks is provided by entering the following command:

```
# MegaCli -ldpdinfo -a0
```

The following command can be used to check the status of each RAID volume and its constituent HDDs:

```
# MegaCli -ldpdinfo -a0 | egrep "Virtual|^State|Slot|Firmware"
```

The following is example output from the above command; a PERC 6/i sub-system where the third volume (Target ID: 2) is in a **Degraded** state and the disk in slot number 4 is in an **Offline** state.

```
Number of Virtual Disks: 3
Virtual Disk: 0 (Target Id: 0)
State: Optimal
Slot Number: 0
Firmware state: Online
Virtual Disk: 1 (Target Id: 1)
State: Optimal
Slot Number: 1
Firmware state: Online
Slot Number: 2
Firmware state: Online
Virtual Disk: 2 (Target Id: 2)
State: Degraded
Slot Number: 3
Firmware state: Online
Slot Number: 4
Firmware state: Offline
```

The PERC 6/i battery status for an **Over Temperature** or **Low** capacity condition can be checked using following command:

```
# MegaCli -AdpBbuCmd GetBbuStatus -a0 |egrep "Capacity Alarm|Over Temp"
```

Command options for MegaCli can be listed by running the following command:

```
# "MegaCli -?"
```

You can also consult the additional documentation listed at the end of this paper^{7 8}.

STATUS NOTIFICATION

Status notification can be performed using SNMP alerts, or email messages, or syslog entries where syslog is monitored for events. The status notifications can be classified as *critical*, *warning* and *notice*. Notifications via email messages, or syslog entries, do not have stringent format requirements for the message. For consistency, you want to build messages that are somewhat similar to the corresponding MIB OID that SNMP alerts would use.

SENDING SNMP ALERTS

Solaris 10 delivers a set of SNMP tools that can be used to send alerts to a management node. Both the `snmptrap` from the Net-SNMP package⁹, and `snmp_sendtrap` can be scripted to notify management agents using the Dell OpenManage MIBs. The Server Administrator Instrumentation MIB (filename: 10892.mib), and the Server Administrator Storage Management MIB (filename: dcstorag.mib) can be extracted from an OpenManage suite installation and used for reporting events about system board components and internal storage. These MIBs are documented in current *Dell OpenManage Server Administrator* guides¹⁰. The "Alert Descriptions and Corrective Actions" section¹¹ in the *OpenManage Server Administrator Storage Management User's Guide* provides a mapping of alert descriptions and SNMP trap numbers in `dcstorag.mib`.

EVENT ENTRIES IN SYSLOG

In environments where the syslog is monitored, generating a syslog entry for any error condition is the easiest notification method of server status. The tool for generating syslog entries from scripts is `logger(1)`. It can be used as shown below:

```
# logger -p daemon.error "An error message needing attention"
```

By default, events would be logged in `/var/adm/messages`. Perl users can use `Sys::Syslog` instead of `logger`. Please check documentation on the `logger(1)` tool for more details.

SENDING EMAIL MESSAGES

Sending email messages alerts is very simple in a Solaris environment. The alert message text (`$alert_msg`) contains needed information so that a receiving administrator can prepare to take remedial action. The subject of the email (`$alert_subject`) can be built by simply concatenating the alert level, subsystem name, and the host name, while the alert message can be the body of the email message. The following example command will send the alert message to the root user:

```
echo "$alert_msg" | mailx -s "$alert_subject" root
```

Similarly, the following command can be used to send the alert log file as the body of the email:

```
cat alert_log_file | mailx -s "$alert_subject" root
```

AN EXAMPLE SCRIPT

The following is a simple script that checks the status of system components and internal RAID systems, and notifies users by logging error conditions in syslog.

```
#!/bin/bash
#
# Check status using ipmitool, raidctl and MegaCli
# Assuming these tools are already in the PATH
#####
# Last updated 8/12/2008
#
timestamp=`date '+%m%d%y-%H%M%S'`

ipmitool sdr list | egrep -v "Not Readable|/| ok" > BMC_${timestamp}.log
while read logline
do
    logger -p DAEMON.ERROR "$timestamp: System needs attention as per BMC"
    logger -p daemon.ERROR "BMC $timestamp: $logline"
done < BMC_${timestamp}.log

raidctl -S | egrep c[0-9]+t[0-9]+d[0-9] | egrep -v OPTIMAL > SAS_${timestamp}.log
while read logline
do
    logger -p daemon.ERROR "$timestamp: SAS RAID needs attention"
    logger -p daemon.ERROR "SAS $timestamp: $logline"
done < SAS_${timestamp}.log
```



```

MegaCli -LdPdInfo -a0|egrep "^Virtual|^Stat|Slot|Firm" > PERC_$(timestamp).log
OIFS=$IFS
IFS=:;
while read -r fld vall
do
  case "$fld" in
    "Virtual Disk")   VD_=$vall      ;;
    "State")          VD_state=$vall ;;
    "Slot Number")   Slot=$vall     ;;
    "Firmware state")
      if [ "$vall" != " Online" ]
      then
        logline="Virtual Disk:$VD_is $VD_state. Disk in Slot$Slot in$vall
state."
        logger -p daemon.ERROR "PERC $(timestamp): $logline"
      fi
      ;;
    esac
done < PERC_$(timestamp).log
IFS=$OIFS;

rm BMC_$(timestamp).log
rm SAS_$(timestamp).log
rm PERC_$(timestamp).log
exit 0

```

CONCLUSION

The framework for in-band monitoring of both PowerEdge system components and internal storage subsystems in a Solaris 10 environment has been outlined in this paper. Examples scripts, as well as command line syntax, have been provided so that readers of this paper can implement this solution in their own environments. In addition to the tools and example scripts and commands listed here, tools such as the `smartmontools`¹² package are available to monitor system components and internal storage subsystems in a Solaris 10 environment; `smartctl` and `smartd` can assist in monitoring disks that are connected to onboard SATA or non-RAID configurations on Dell SAS6/iR, and is supported with Solaris.

Visit Dell support and solution websites^{13 14} for other whitepapers, and useful information, to help support the launch of Solaris 10 on Dell PowerEdge Servers.

REFERENCES

- ¹ The Dell OpenManage suite of applications builds solutions based on industry standards based protocols and practices. See www.dell.com/openmanage
- ² Dell Remote Access Controllers – Manual
<http://support.dell.com/support/edocs/software/smdrac3/>
- ³ Dell OpenManage Baseboard Management Controller – Manual
<http://support.dell.com/support/edocs/software/smbmcmu/>
- ⁴ Monitoring and Managing Agentless Servers - Using Dell OpenManage IT Assistant 8.0 with IPMI
<http://www.dell.com/downloads/global/power/ps4q06-20070158-John.pdf>
- ⁵ Multiple Ways to Efficiently Monitor and Manage Dell PowerEdge Servers
<http://www.dell.com/downloads/global/power/ps4q06-20060320-Bose-OE.pdf>
- ⁶ Using IPMItool raw commands for remote management of Dell PowerEdge Servers
<http://www.dell.com/downloads/global/power/ps4q07-20070387-Babu.pdf>
- ⁷ Managing PERC 6 with MegaCli under Solaris 10
http://linux.dell.com/files/whitepapers/solaris/Managing_PERC6_0714.pdf
- ⁸ MegaRAID SAS Software User's Guide: Chapter-3
http://www.lsi.com/files/docs/techdocs/storage_stand_prod/sas/mr_sas_sw_ug.pdf
- ⁹ Net-SNMP Tutorial -- Commands
<http://www.net-snmp.org/tutorial/tutorial-5/commands/index.html>
- ¹⁰ Dell OpenManage Server Administrator Version 5.4 SNMP Reference Guide
http://support.dell.com/support/edocs/software/svradmin/5.4/en/snmp/pdf/om_54_snmp_ref_gd.pdf
- ¹¹ OMSS UG: Alert Descriptions and Corrective Actions
http://support.dell.com/support/edocs/software/svradmin/5.4/en/omss_ug/html/evntmnr.html#1877763
- ¹² <http://smartmontools.sourceforge.net/index.html>
- ¹³ Dell Solaris Solutions
<http://www.dell.com/solaris>
- ¹⁴ Dell TechCenter Wiki: Solaris
<http://www.delltechcenter.com/page/Solaris>