

Hardware Monitoring with the new IPMI Plugin v2

Werner Fischer, Technology Specialist Thomas-Krenn.AG

6. OSMC / Nuremberg / Germany
29th November 2011



Thomas-Krenn.AG[®]
The server experts



Introduction

who I am


Werner Fischer



working for a
Server vendor




Linux user
since 2001



who I am not

Kernel or H/W
developer



Introduction

who is **Thomas-Krenn.AG**[®]
The server experts



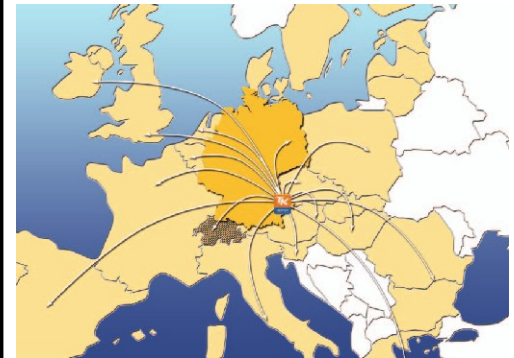
Server & accessories
"Made in Germany"



based in Freyung,
Bavaria



serving all over Europe



Some questions ...

Should I use monitoring?

Some questions ...

Should I use monitoring?

It depends on what you
want to do in your free time...

remember yesterday evening?



Some questions ...

All drives of the RAID 6 O.K.?




Some questions ...

All network connections O.K.?




Some questions ...

All FANs O.K.?






Some questions ...

All power supplies O.K.?






Some questions ...

Can we monitor all these servers?






Some questions ...

Can we monitor all these servers?
Easily?



Some questions ...

Can we monitor all these servers?
Easily?
With one single tool?



Some questions ...

Can we monitor all these servers?

Easily?

With one single tool?

**Yes,
we can!**

Agenda

- 1) IPMI overview**
- 2) Plugin implementation**
- 3) Live demo**
- 4) Common pitfalls**

Intelligent Platform Management Interface

- **IPMI developed by Intel, HP, NEC, Dell**
 - 1998: IPMI v1.0
 - 2001: IPMI v1.5
 - 2004: IPMI v2.0

IPMI main features

1

Monitoring
(temp, fans, ...)

2

Recovery Control
(power on/off/reset)

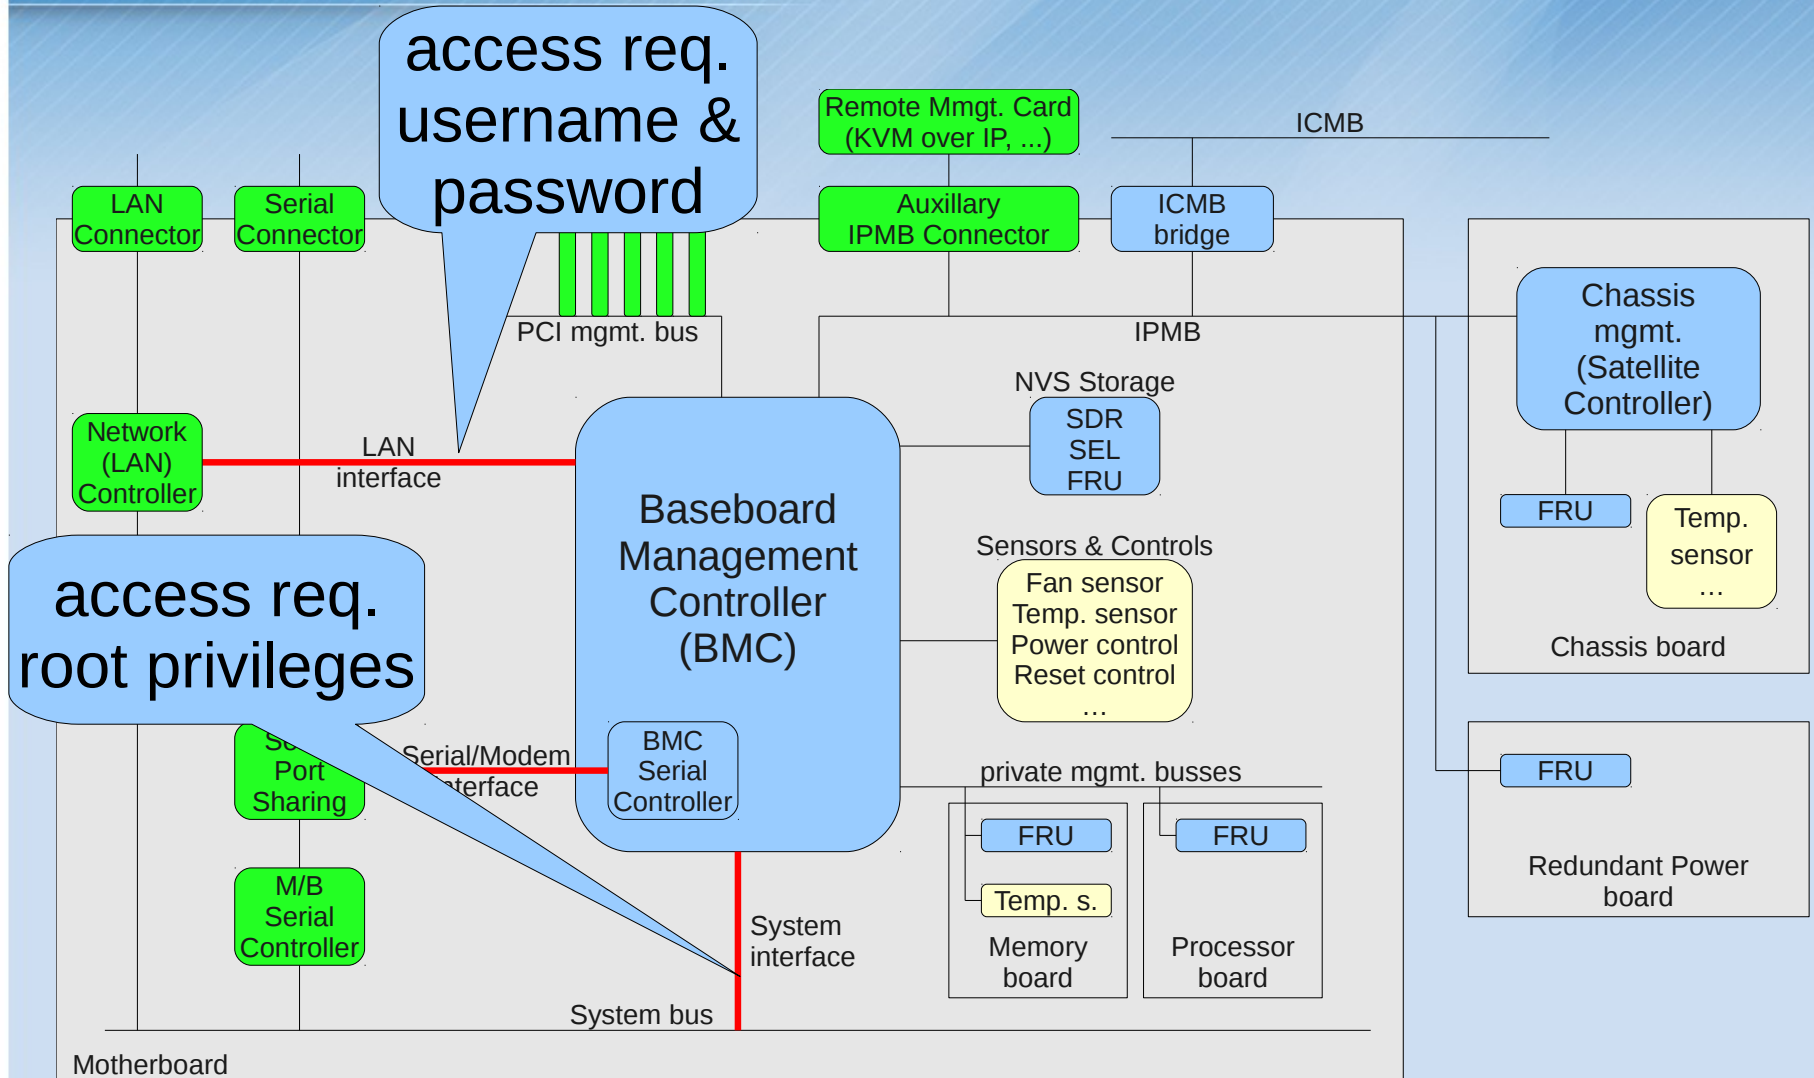
3

Logging
(System Event Log)

4


Inventory
(FRU information)

IPMI overview



IPMI Channel Privilege Levels (LAN access)

Privilege Level	Allows
User	<ul style="list-style-type: none">• query sensors
Operator	<ul style="list-style-type: none">• nearly all IPMI commands• but no changing of out-of-band interfaces
Administrator	<ul style="list-style-type: none">• all IPMI commands



use privilege level 'User'
for monitoring purposes

Example: remote control with ipmitool

```
[user@adminpc ~]$ ipmitool -I lan -H 192.168.1.211 \  
                    -U admin power status
```

Password:

Chassis Power is off

```
[user@adminpc ~]$
```

```
[user@adminpc ~]$ ipmitool -I lan -H 192.168.1.211 \  
                    -U admin power on
```

Password:

Chassis Power Control: Up/On

```
[user@adminpc ~]$
```

```
[user@adminpc ~]$ ipmitool -I lan -H 192.168.1.211 \  
                    -U admin power status
```

Password:

Chassis Power is on

```
[user@adminpc ~]$
```

IPMI Sensor Classes (1/2)

Discrete

multiple states possible:

- up to 15 states
- each state is reflected by a bit
- multiple state bits can active

can provide:

- generic states
- sensor-specific states

other class similar to discrete:

- OEM: discrete sensor where the meaning of the states (offsets) are OEM defined


Threshold

changes event status on:

- analog reading compared to threshold values

provides:

- analog reading of the sensor
- discr. threshold comparison status bit



IPMI Sensor Classes (2/2)

Discrete

```
[root@test ~]# ipmitool sdr get "PS2 Status"
Sensor ID          : PS2 Status (0x71)
Entity ID         : 10.2 (Power Supply)
Sensor Type (Discrete) : Power Supply
States Asserted   : Power Supply
                   [Presence detected]
                   [Power Supply AC lost]
Assertion Events  : Power Supply
                   [Presence detected]
                   [Power Supply AC lost]
Assertions Enabled : Power Supply
                   [Presence detected]
                   [Failure detected]
                   [Predictive failure]
                   [Power Supply AC lost]
[...]
```

```
Deassertions Enabled : Power Supply
[...]
```

Threshold

```
[root@test ~]# ipmitool sdr get "Fan 1"
Sensor ID          : Fan 1 (0x50)
Entity ID         : 29.1 (Fan Device)
Sensor Type (Analog) : Fan
Sensor Reading     : 5719 (+/- 0) RPM
Status            : ok
Nominal Reading    : 6708.000
Normal Minimum     : 2451.000
Normal Maximum     : 10965.000
Lower critical     : 1720.000
Lower non-critical : 1978.000
Positive Hysteresis : 86.000
Negative Hysteresis : 86.000
Minimum sensor range : Unspecified
Maximum sensor range : Unspecified
Event Message Control : Per-threshold
Readable Thresholds : lcr lnc
Settable Thresholds : lcr lnc
Threshold Read Mask : lcr lnc
Assertion Events    :
Assertions Enabled  : lnc- lcr-
Deassertions Enabled : lnc- lcr-
```

IPMI Sensor Types

```
root@test:~# ipmi-sensors -L
Temperature
Voltage
Current
Fan
Physical_Security
Platform_Security_Violation_Attempt
Processor
Power_Supply
Power_Unit
Cooling_Device
[...]
```

Example: query sensors with FreeIPMI

```
[root@testserver ~]# ipmimonitoring
```

```
Record_ID | Sensor Name | Sensor Group | Monitoring Status |  
Sensor Units | Sensor Reading  
[...]  
17 | Fan 5 | Fan | Nominal | RPM | 9052.000000  
18 | Fan 6 | Fan | Nominal | RPM | 8060.000000  
19 | PS1 AC Current | Current | Nominal | A | 0.124000  
20 | PS2 AC Current | Current | Nominal | A | 0.992000  
[...]  
36 | Physical Scrty | Physical Security | Critical | N/A |  
'General Chassis Intrusion'
```

Example: interpret discrete sensors(FreeIPMI)

```
root@test:~# cat /etc/freeipmi/freeipmi_interpret_sensor.conf
[...]
```

## IPMI_Physical_Security	
#	
# IPMI_Physical_Security_No_Event	Nominal
# IPMI_Physical_Security_General_Chassis_Intrusion	Critical
# IPMI_Physical_Security_Drive_Bay_Intrusion	Critical
[...]	
# IPMI_Power_Supply_No_Event	Nominal
# IPMI_Power_Supply_Presence_Detected	Nominal
# IPMI_Power_Supply_Power_Supply_Failure_Detected	Critical
# IPMI_Power_Supply_Predictive_Failure	Critical
# IPMI_Power_Supply_Power_Supply_Input_Lost_AC_DC	Critical
[...]	
## IPMI_Memory	
#	
# IPMI_Memory_No_Event	Nominal
# IPMI_Memory_Correctable_Memory_Error	Warning
# IPMI_Memory_Uncorrectable_Memory_Error	Critical

IPMI System Event Log (SEL)

- **stored in non-volatile storage**

```
[root@testserver ~]# ipmitool sel elist
 40 | 06/21/2010 | 14:29:29 | Power Supply PS1 Status | Power Supply AC lost | Asserted
 54 | 06/21/2010 | 14:29:29 | Power Unit Power Redundancy | Fully Redundant
 68 | 06/21/2010 | 14:29:29 | Power Unit Power Redundancy | Redundancy Lost
 7c | 06/21/2010 | 14:29:29 | Power Unit Power Redundancy | Non-Redundant: Sufficient from Redundant
[...]
2fc | 06/21/2010 | 15:20:32 | Physical Security Physical Scrtty | General Chassis intrusion | Asserted
```

```
[root@testserver ~]# ipmitool sel elist
Power Supply PS1 Status | Power Supply AC lost | Asserted
Power Unit Power Redundancy | Fully Redundant
Power Unit Power Redundancy | Redundancy Lost
Power Unit Power Redundancy | Non-Redundant: Sufficient from Redundant
[...]
Physical Security Physical Scrtty | General Chassis intrusion | Asserted
```

Agenda

- 1) IPMI overview
- 2) Plugin implementation
- 3) Live demo
- 4) Common pitfalls

Plugin implementation

- **Bash script**
- **uses FreeIPMI, gawk**

```
# ./check_ipmi_sensor -H 10.10.10.114 -f /etc/ipmi-config/ipmi.cfg
IPMI Status: OK | 'System Temp'=29.000000 'FAN 1'=4185.000000 'FAN
2'=4320.000000 'FAN 3'=4590.000000 'FAN 4'=4320.000000 'FAN
A'=4590.000000 'Vcore'=0.712000 '3.3VCC'=3.392000 '12V'=12.190000
'VDIMM'=1.528000 '5VCC'=5.088000 '-12V'=-11.681000 'VBAT'=3.024000
'VSB'=3.344000 'AVCC'=3.408000
```

Plugin implementation

- **Bash script**
- **uses FreeIPMI, gawk**

```
# ./check_ipmi_sensor -H 10.10.10.114 -f /etc/ipmi-config/ipmi.cfg -v 2
IPMI Status: OK | 'System Temp'=29.000000 'FAN 1'=4320.000000 'FAN
[...]
System Temp = 29.000000 (Status: Nominal)
CPU Temp = 'Low' (Status: Nominal)
FAN 1 = 4320.000000 (Status: Nominal)
FAN 2 = 4320.000000 (Status: Nominal)
FAN 3 = 4590.000000 (Status: Nominal)
[...]
AVCC = 3.408000 (Status: Nominal)
Chassis Intru = 'OK' (Status: Nominal)
PS Status = 'Presence detected' (Status: Nominal)
```

Plugin implementation

- clear illustration in webinterfaces

The screenshot displays a monitoring web interface with a table of services and a detailed view of a critical IPMI service.

Service	Status	Last check	Duration	Info	Attempt	Output
Host: debian6 (2 Items)						
<input checked="" type="checkbox"/> IPMI	CRITICAL	2011-11-29 14:52:19			1 / 3	IPMI Status: C
<input type="checkbox"/> PING	OK	2				OK - Pa

Detailed serviceinfo

Display name	IPMI
Current state	CRITICAL
Output	IPMI Status: Critical [Chassis Intru = Critical]
SERVICE_LONG_OUTPUT	
Performance data	'System Temp'=29.000000 'FAN 1'=4050.000000 'FAN 2'=4320.000000 'FAN 3'=4590.000000 'FAN 4'=4185.000000 'FAN A'=4725.000000 'Vcore'=0.712000 '3.3VCC'=3.392000 '12V'=12.190000 'VDIMM'=1.528000 '5VCC'=5.088000 '-12V'=-11.681000 'VBAT'=3.024000 'VSB'=3.344000 'AVCC'=3.408000
Current check attempt	1
Max check attempts	3
Last check	2011-11-29 14:54:19
Check type	ACTIVE

Agenda

- 1) IPMI overview**
- 2) Plugin implementation**
- 3) Live demo**
- 4) Common pitfalls**

Agenda

- 1) IPMI overview
- 2) Plugin implementation
- 3) Live demo
- 4) **Common pitfalls**

Common pitfalls

- **sensors with state N/A**

```
[...]  
12 | CPU1 Temp | OEM Reserved | N/A | N/A | N/A | 'OEM Event = 0000h'  
13 | CPU2 Temp | OEM Reserved | N/A | N/A | N/A | 'OEM Event = 0000h'  
[...]
```

- **solution shortest-term: exclude (-x opt.)**
- **solution short-term: FreeIPMI update**
tkwiki.cc/FreeIPMI-NA-Sensor

Common pitfalls


- **unrecognized events**

```
[...]  
40 | Status | Cable/Interconnect | Nominal | N/A |  
'Cable/Interconnect is connected'  
41 | RAC Status | Module/Board | N/A | N/A | 'Unrecognized Event =  
0001h' 'Unrecognized Event = 0002h' 'Unrecognized Event = 0004h'  
42 | OS Watchdog | Watchdog 2 | Nominal | N/A | 'OK'  
[...]
```

- **solution shortest-term: ignore unrec. e.**
tkwiki.cc/FreeIPMI-Unrec-Event

Agenda

- 1) IPMI overview
- 2) Plugin implementation
- 3) Live demo
- 4) Common pitfalls



some conclusions ...

Conclusions (1/2)

- **Download:**
www.thomas-krenn.com/en/oss
- **Mailing List:**
lists.thomas-krenn.com
- **Thanks for your contribution:**
Nikolaus Filus, Timme Katz, Lars Meuser, Sebastian Mörchen, Gustav Olsson, Holger Paschke, Andy Spiegl, Ulrich Zehl

Conclusions (2/2)

1 Monitor hardware
with Icinga & IPMI

2 Problems?
They will tell you!

3 It'll save you
time & money



Get German article
on the plugin for free at
tkwiki.cc/ipmi-plugin



Thanks for your time!