



nils.kanning@gwdg.de

Nils Kanning

Intelligent Platform Management Interface (IPMI)

How to Avoid a Bike Ride to the Data Center

Table of contents

- 1 Introduction
- 2 Basic IPMI Commands
- 3 Further Topics
- 4 Summary

Why IPMI?

- Want to find out why compute job on Emmy node gcn2960 failed
- No chance because node crashed and is not reachable?

```
# ssh gcn2960
```

```
ssh: connect to host gcn2960 port 22: No route to host
```

- Next step: Bike ride to data center and press reset button?
- No, reset node via IPMI

```
# ipmitool -I lanplus -H gcn2960-bmc -U admin power reset
```

What is IPMI?

- Intelligent platform management interface
- Manage server even if powered off or unresponsive
 - ▶ Power on, restart and even control server as if on-site
 - ▶ Monitoring, logging, access hardware information
- Network connection to dedicated controller independent of CPU and OS
- IPMI specification led by Intel and supported by over 200 vendors

IPMI Architecture

- Baseboard management controller (BMC): Heart of IPMI
- Network: Dedicated LAN port (preferred) or shared with host
- Satellite controller for chassis connected to BMC via bus
- Sensors on motherboard and for chassis

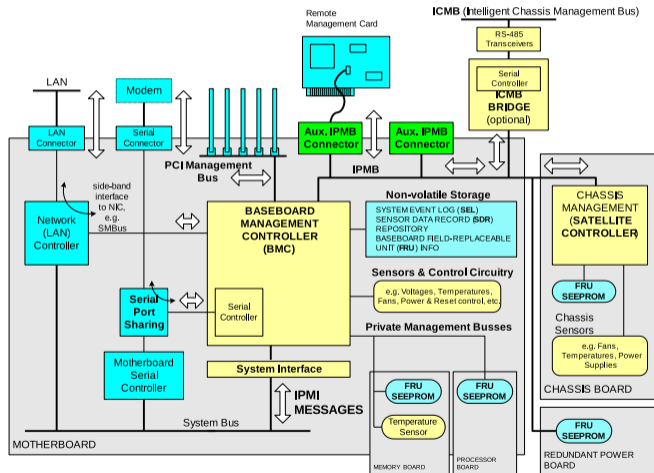


Image source: IPMI Specification v2.0, Revision 1.1, 2013

System Management Software

■ IPMITool

- ▶ Access server remotely via LAN

```
# ipmitool -I lanplus -H <BMC-IP-ADDRESS> -U <USER> <COMMAND>
```

- ▶ Access local server (requires OpenIPMI kernel modules)

```
# ipmitool <COMMAND>
```

- ▶ Command overview

```
# ipmitool [<COMMAND>] help
```

■ Some other software packages

- ▶ OpenIPMI: Kernel modules and higher-level abstraction via libraries
- ▶ FreeIPMI: Alternative to IPMITool, no need for kernel modules

Power and Chassis

■ Overview of power commands

```
# ipmitool -I lanplus -H gcn2960-bmc -U admin power help
```

Password:

chassis power Commands: status, on, off, cycle, reset, diag, soft

■ Check chassis status

```
# ipmitool -I lanplus -H gcn2960-bmc -U admin chassis status
```

```
System Power           : on
Power Overload          : false
Power Interlock         : inactive
Main Power Fault        : false
Power Control Fault     : false
[...]
```

System Event Log (SEL)

- Analyze cause of crash/problem
- List of system events stored in non-volatile memory (IDs in hexadecimal)

```
# ipmitool -I lanplus -H gcn2960-bmc -U admin sel list
339 | 07/20/2022 | 12:35:01 | Temperature #0x21 | Upper Non-critical
445 | 02/07/2023 | 16:10:29 | Power Supply #0x50 | Power Supply AC lost
44d | 02/15/2023 | 08:42:15 | System Event #0x83 | OEM System boot event
45a | 02/16/2023 | 20:33:09 | Power Unit #0x01 | Power off/down
[...]
```

- Details of specific event (ID in decimal)

```
# ipmitool -I lanplus -H gcn2960-bmc -U admin sel get <ID>
```

- Some events hardware-specific: [Intel SEL Troubleshooting Guide](#)

Sensor Data Repository (SDR)

■ Current readings of all sensors

```
# ipmitool [...] sdr list
System Airflow | 52 CFM | ok
Inlet Temp | 33 degrees C | ok
CPU0 VR Temp | 44 degrees C | ok
MBrd Temp | 36 degrees C | ok
System Fan 1A | 7728 RPM | ok
Fan 1 Present | 0x00 | ok
Mem 0 VR Temp | 36 degrees C | ok
PS1 Input Power | 732 Watts | ok
BB +12.0V | 12.32 Volts | ok
[...]
```

■ Details of specific sensor

```
# ipmitool [...] sdr get "MBrd Temp"
Sensor ID : MBrd Temp (0x25)
Entity ID : 7.1 (System Board)
Sensor Type : Temperature (0x01)
Sensor Reading : 34 degrees C
Status : ok
Nominal Reading : 60.000
Normal Minimum : 10.000
Normal Maximum : 105.000
Upper critical : 115.000
[...]
```

■ Sensors crossing thresholds recorded in SEL, thresholds adjustable

Serial Over LAN (SOL)

- Redirect data for serial port of motherboard through IPMI session
- Access to BIOS/UEFI and Linux console possible
- Enable SOL and start session (exit with ~~.)

```
# ipmitool -I lanplus -H gcn2960-bmc -U admin sol set enabled true
# ipmitool -I lanplus -H gcn2960-bmc -U admin sol activate
[SOL Session operational. Use ~? for help]
CentOS Linux 7 (Core)
Kernel 3.10.0-1160.76.1.el7.x86_64 on an x86_64
gcn2960 login:
```

- Enter BIOS/UEFI by changing boot device before starting session

```
# ipmitool -I lanplus -H gcn2960-bmc -U admin chassis bootdev bios
```

BMC Web Console

- Vendor-specific, not part of IPMI specification
 - Graphical interface for IPMI features and more
 - ▶ Keyboard video mouse (KVM):
Remote Control/iKVM over HTML5
 - ▶ Using remote image files: Virtual Media
 - Access via laptop
 - ▶ Forward port of web console
- ```
laptop$ ssh -L 4443:gcx2960-bmc:443 gadm1
```
- ▶ Open in browser: <https://localhost:4443>

The screenshot shows the Intel Integrated BMC Web Console interface. At the top is the Intel logo and the title "Integrated BMC Web Console". Below this is a navigation bar with tabs: System, Server Health, Configuration, Remote Control, Virtual Media, and Summary. The "Summary" tab is selected. On the left side, there is a sidebar menu with links: System Information, FRU Information, CPU Information, DIMM Information, NVMe Information, NIC Information, Storage Information, and Current Users. The main content area displays the "Summary" section, which includes a warning: "KCS Policy Control Mode is Allow All. This". Below this, the "Summary" section lists various system status items: Host Power Status (Host is current), Remote Management Module key (Installed), Device (BMC) Available (Yes), BMC Firmware Build Time, BIOS ID, BMC FW Rev, Backup BMC FW Rev, Build ID, SDR Package Version, Mgmt Engine (ME) FW Rev, Baseboard Serial Number, and Overall System Health (indicated by two green circles). At the bottom right, there is a "Web Session Timeout" section with a dropdown menu set to "30 Min(s)".

Image source: Screenshot of Intel Integrated BMC Web Console

# Telegraf Plugin: Sensor Data as Time Series

- Only current sensor readings available via IPMI
- **Telegraf IPMI plugin** imports readings into InfluxDB
- Telegraf and InfluxDB part of TIG/**TICK** monitoring
- IPMITool executed at regular intervals for all nodes
- No performance impact because via BMC not CPU

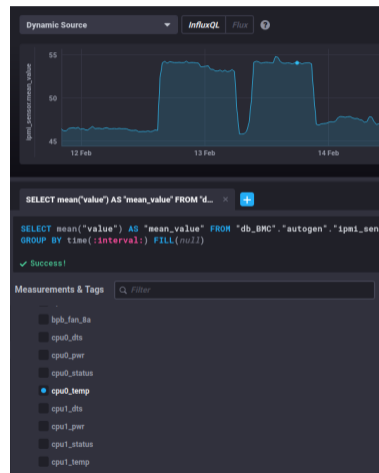


Image source: Screenshot of Chronograf showing IPMI data

# Redfish: IPMI Successor

- Aim: Industry standard replacing IPMI over LAN
- [Redfish specification](#) v1.0 in 2015, actively developed
- Secure and easy to use: RESTful API via HTTPS and JSON data format
- Client applications, browser interfaces and libraries
- Manage IT infrastructure beyond servers: Storage, network, power/cooling



Image source: [DMTF Redfish Forum](#)

# Key Takeaways

- IPMI indispensable tool for admins in particular of large clusters
- Power control, event logs (SEL) and sensor data (SDR) all independent of OS
- Enabled by controller on motherboard (BMC) with separate LAN interface
- IPMITool utility to manage servers: Only essential commands covered in talk
- Exercise: What does this command do?

```
ipmitool raw 0x00 0x02 0x03
```

# IPMI References

- Thomas-Krenn - IPMI Basics

See also other articles in IPMI section

- Thomas-Krenn - IPMITool Cheat Sheet

- IPMI Specification v2.0, Revision 1.1, 2013

Over 600 pages including 25-page non-technical introduction

- ADMIN 03/2010 - IPMI

Old but informative article

- Dan Farmer - IPMI

Critical discussion of IPMI and BMC security