# SMOA Devices - distributed power management utility for Green IT infrastructure


Krzysztof Kurowski
Poznań Supercomputing and Networking Center
Applications Department
ul. Noskowskiego 10
61-704 Poznań, Poland
krzysztof.kurowski@man.poznan.pl

Ariel Oleksiak
Poznań Supercomputing and Networking Center
Applications Department
ul. Noskowskiego 10
61-704 Poznań, Poland
ariel@man.poznan.pl

Michał Witkowski
Poznań Supercomputing and Networking Center
Applications Department
ul. Noskowskiego 10
61-704 Poznań, Poland
michal.w.witkowski@gmail.com


## 1. INTRODUCTION

Rising energy prices and the increasing environmental awareness of the general public are playing an ever-increasing role in the world of IT and computational science. After years of constructing computer systems solely for their sheer performance, power consumption became one of the main criteria for evaluating a computer system and metrics, such as performance-per-watt, are receiving increasing attention.

The push for increased energy-efficiency is also present in the field of desktop computing and even home appliances. All major modern operating systems feature sophisticated power management capabilities based on open industry standards such as ACPI. While these features can reduce the power consumption of a running machine, they do not address the problem of computers which are idling. A typical desktop PC, according to our measurements, can consume between 100 and 140 watts while idling without any gain in productivity. The issue can be mitigated by suspending or shutting down idling machines, although such solution entails the problem of waking the machines up afterwards. Moreover, computers and other electronic devices can draw a significant amount of power even when they are turned off, but still connected to utilities. Our experiences show that power supplies of some servers can draw up to 40W when the machine is turned off. Thus management of power at outlet level can also bring energy savings.

Even the smallest power consumption reduction, achieved using the techniques already available in modern computer hardware and software, can cause enormous net-savings on company office or cluster scales. However, introducing a power saving policy on a larger scale requires a sophisticated distributed system for power management, control and detailed metering. Such a system should be easy to deploy, regardless of network security policies, and be able to operate in a multi-vendor environment.

In this poster, we present an easily extensible, flexible and standards-based distributed monitoring and control system for energy consumption management. The main purpose of its creation was to combine various power management related tools, solutions and monitoring equipment using a standardized, easy-to-deploy network protocol.

Our system provides information about available computing resources for each managed machine and exposes power management capabilities available in major operating systems. It also offers means of bringing up machines after they are powered off, hibernated or suspended in order to reduce idling times. Moreover, control and measurement of power provided to the machines at outlet level is also possible. All these capabilities are exposed using a unified interface, making our solution suitable either for administrative use or as a power management tool for other mechanisms, e.g. batch schedulers.

An important feature of the SMOA Devices platform is that it is designed to cope with network restrictions. It is suitable for deployment even in heavily-firewalled environments, thanks to a novel use of the XMPP protocol.

## 2. ARCHITECTURE

The SMOA Devices architecture integrates various tools and solutions across different levels ranging from integrated measuring devices to OS power management policies. Each of them is represented in the SMOA Devices platform as a separate *Node* having its own JID (Jabber ID) in the XMPP network. Every *Node* plays one of three roles: *Device* (machine) node, *Waker* node or *Meter* node. The role played by the *Node* is determined by the *handlers*, software plugins described shortly in section 3, that are enabled.

The *Device* node is responsible for interacting with the machine's operating system. It provides our platform with the power management capabilities already available on the machines. The *Device* node allows for CPU frequency scaling, power management profiles configuration, system temperature monitoring, process resource utilization and providing means for safe shutdown/suspend of the machine's operating system.

The *Waker* node is responsible for bringing up machines after their shutdown by the means of Wake-on-LAN, IPMI or other vendor-supplied means. The network location of the *Waker* node is crucial. In case of Wake-on-LAN (WoL) the *Waker* needs to be placed in the same network segment as the target machines, since the WoL magic packet is delivered by network broadcast. When IPMI is used, the *Waker* would need to be placed on a machine bordering normal and high-security management networks.

*Meter* nodes are interfaces between the SMOA Devices platform and the power measurement/control hardware. The type of hardware is irrelevant, it could be some embedded measuring devices, manageable power distribution units (PDU) or intelligent server cases. The responsibility of the *Meter* is to provide a unified mapping of node JIDs to measurement information and outlet control functions.

The *SMOA Devices Service* is the central point of the architecture integrating capabilities provided by the *Device, Waker* and *Meter* nodes, so that the management functionality and information about a physical machine are provided in a coherent way, regardless of their source. Of course multiple *Services* can be used to manage a pool of *Nodes* in order to increase redundancy and scalability.

## 3. SOFTWARE FLEXIBILITY AND INTEROPERABILITY

The software components of the SMOA Devices architecture need to be highly portable and easy-to-extend, thus Python was chosen as the language of implementation. The *Node* code is available for all major platforms (Windows, Linux and Mac OS X), providing the same core functionality regardless of the OS involved. All components are highly modular with each featuring the concept of a *handler*. Each *handler* provides some dedicated functionality, e.g. supplying information about CPU load or sending Wake-on-LAN requests. *Handlers* can be easily added or swapped-out to adapt the platform to specific scenarios. For example, in server environment one might want to write a *handler* to provide *Waker* and *Meter* capabilities using vendor-supplied solutions, such as IBM BladeCenter Advanced Management Module.

The platform's software also features a simple notification mechanism similar to a push-pull model, allowing software components or users to subscribe to various information sources. The information in these sources can be provided either periodically (e.g. power consumption statistics) or in an event-based manner (e.g. alarms). In both cases, XMPP provides the perfect protocol for such uses.

For end-user interaction SMOA Devices packages contain a dedicated plugin for a widely popular IM client Pidgin. The aim of the plugin is to blur the gap between management tools and end-user applications, thus all managed machines are visible as “buddies” on the contact list. Since each machine has a dedicated JID, its “on-line” state represents its power state. The contact's status message contains the hostname for easier identification. The plugin also allows the user to change the powerstate of the machines via right-click menu items. The functionality of the plugin brings the management functionality closer to the desktop.

## 4. APPLICATION SCENARIOS

While there is a significant effort to reduce the power consumption of individual components of computer systems, the most savings can be achieved by powering off computers altogether. As far as economy is concerned, an idling computer consumes quite a lot of power without producing much value. Even computers (or other appliances) that are turned off, but still connected to utilities, draw a certain amount of power that may be significant in deployments of a larger scale. SMOA Devices can be used in different environments and scenarios to counter these problems. Possible application scenarios, both in HPC and office environments, are listed in the poster.

## 5. CONCLUSION

SMOA Devices platform faces the challenge of integrating methods and tools available on different abstraction levels ranging from power outlet control, through operating system calls, to networked remote wake-up. This kind of vertical integration, combined with easily-deployable communication protocol, asynchronous notifications and wide-ranging monitoring and control capabilities, make SMOA Devices a well-suited solution for many application scenarios in the upcoming years of increased environmental awareness and the green computing revolution.