

Zdeněk Salvét

## Virtual machine control tools in METACentrum

CESNET Technical Report X/2007

### 1 Introduction

This technical report describes first generation of commandline and WWW portal tools developed to control virtual machine environment in METACentrum.

### 2 Virtualized environment

We have deployed two different virtualization systems, namely *Xen* [3] and *VServer* [2]. Each of them has its own set of advantages and disadvantages, Xen is used on machines with two or four processors (processor cores) and centralized memory subsystem while on 8-processor machines with NUMA memory architecture (AMD Opteron-based) that are not yet supported by Xen in an optimal way, Vserver is used.

### 3 Control service and command-line tools

We developed set of client-server control tools *xman* that enables easy and secure remote control of virtual machines (viewing state, starting, stoping etc.) for administrators and later, when full authorization architecture is developed, even for ordinary users. Each physical machine (instance of virtual machine monitor in possible multilevel scenario) runs server side – **xmand** daemon which can be managed using **rxm** command line utility.

Currently, following operations are provided:

<code>help</code>	Print help.
<code>list</code>	List information about domains.
<code>create</code>	Create a domain.
<code>destroy</code>	Terminate a domain immediately.
<code>reboot</code>	Reboot a domain.
<code>shutdown</code>	Shutdown a domain.
<code>pause</code>	Pause execution of a domain.
<code>unpause</code>	Unpause a paused domain.
<code>mem-set</code>	Adjust the current memory usage for a domain.
<code>vcpu-set</code>	Set the number of VCPUs for a domain.
<code>describe</code>	Provide domain's XML description.
<code>migrate</code>	Migrate a domain to another machine.
<code>dmesg</code>	Read Xend's message buffer.
<code>info</code>	Get information about the xen host.

### 3.1 Archicture and implementation

In order to support both current virtualization systems and possible new ones that can appear later in an uniform manner, *libvirt* toolkit [1] is used inside `xmand` server to interact with control interfaces of respective virtual machine monitors and their supporting tools. Inside Xen backend, *libvirt* library communicates with `xend` management daemon using its XMLRPC protocol interface over local socket (it has to be enabled with (`xend-unix-server yes`) setting in `xend`'s configuration file (`/etc/xen/xend-config.sxp`).

Every access to `xman` service (`xmand` daemon) is subject to strong user or client service authentication using Kerberos 5 authentication framework which is used as primary authentication protocol across all main services in METACentrum. Both popular implementations of the framework, that is MIT Kerberos 5 and Heimdal, can be used to build `xmand` and `rxm` executables (there is small difference in source codes used for each one).

In current version, access control is simplified and membership of authenticated client in list of physical machine administators (which is kept in standard `.k5login` text file after beeing centrally disseminated by Perun service) is required for all operations. We are going to query more advanced authorization service containing dynamic information on user jobs when available and allow ordinary users full control over virtual machines where they cannot negatively influence other users' work.

Information on all important actions such as new connection establishment and command execution is logged to system log by `xmand`.

## 4 WWW portal interface

Intuitive WWW interface has been implemented on top of command-line client tool. Since an access to virtual machine management is currently limited to system administrators we have deployed it on METACentrum internal administrative portal. We plan deployment on user portal as soon as authorization scheme is extended.

The screenshot shows a Mozilla browser window with the address bar containing `https://mizar.ics.muni.cz:880/admin/virtual/virt_`. The main content area displays a table with the following data:

<a href="#">skirit80-1.ics.muni.cz</a>	running	7.12.2007 16:33:17
<a href="#">skirit80-2.ics.muni.cz</a>	occupied	7.12.2007 16:33:17
<a href="#">skirit81-1.ics.muni.cz</a>	free	7.12.2007 16:33:23
<a href="#">skirit81-2.ics.muni.cz</a>	free	7.12.2007 16:33:23
<a href="#">skirit84-1.ics.muni.cz</a>	running	7.12.2007 16:24:34
<a href="#">skirit84-2.ics.muni.cz</a>	free	7.12.2007 16:24:34
<a href="#">skurut68-1.cesnet.cz</a>		
<a href="#">skurut68-2.cesnet.cz</a>		

Below the table is a legend titled "Vysvětlivky:" with the following entries:

- free**: doména je online, nemá přidělené zdroje, šlo by do ní submitnout a tak ji překlopit do running stavu
- occupied-would-preempt**: doména je online, nemá zdroje (používá je ta druhá, running doména), šlo by do ní submitnout jen privilegovaný job a tím preemptnout druhou doménu
- occupied**: doména je online, nemá zdroje (používá je ta druhá, running doména), nelze do ní submitnout
- running-preemptible**: doména je online, běží job, mohla by přijmout další, mohla by být preemptnuta jinou doménou
- running**: doména je online, běží job, mohla by přijmout další, nemůže být preemptnuta jinou doménou
- preempted**: doména je online, ale zdroje jí sebrala jiná doména (ta je running), může v ní být job, nelze do ní submitnout
- frozen**: doména je offline, je suspendována, může v ní být job, nelze do ní submitnout

At the bottom of the legend area, a grey bar contains the text: "Šedě podsvícený řádek upozorňuje na zastaralé informace."

Figure 1: Upper level page of WWW interface

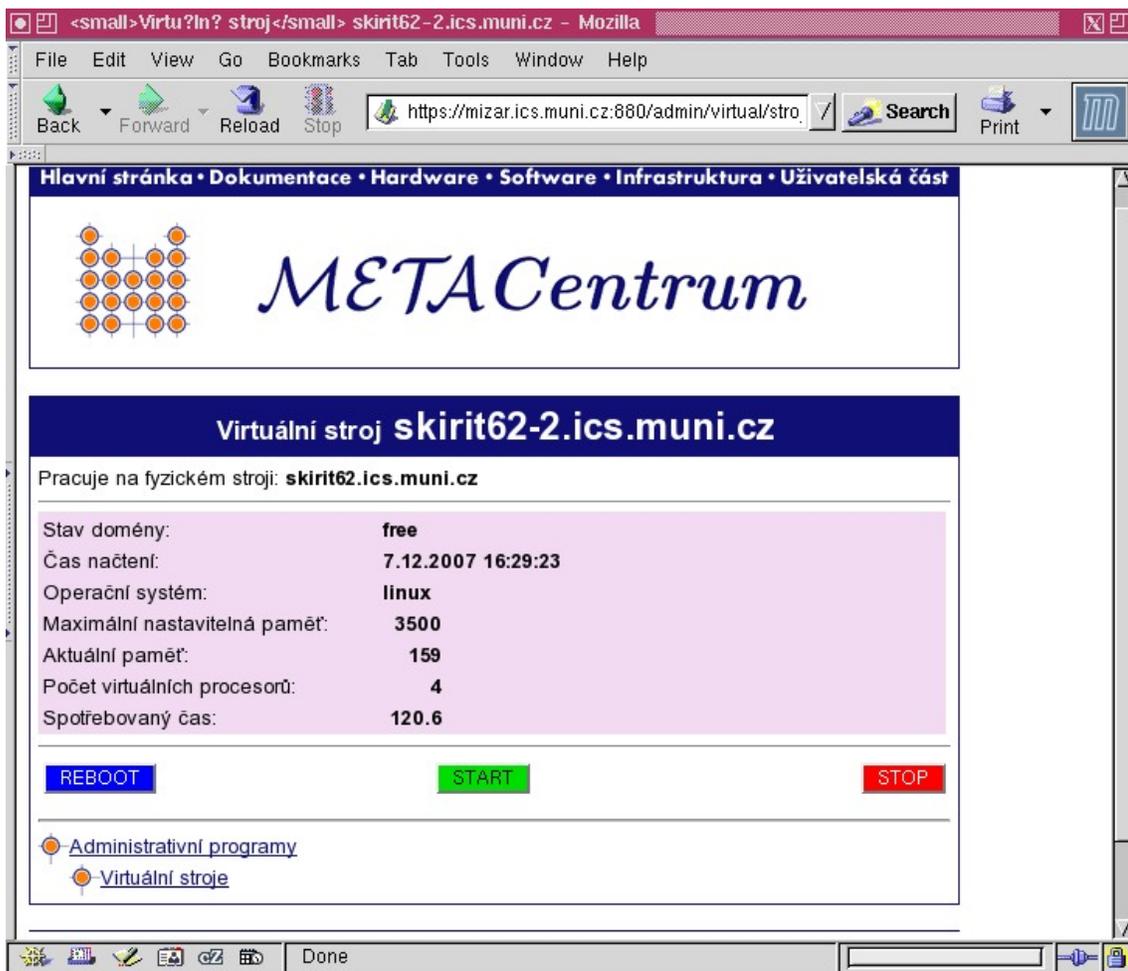


Figure 2: Details and actions available for particular virtual machine

## References

- [1] Libvirt virtualization API. <http://libvirt.org/>.
- [2] Linux vserver. <http://linux-vserver.org/>.
- [3] Xen virtual machine monitor. <http://www.cl.cam.ac.uk/Research/SRG/netos/xen/index.html>.