



# Rocks based Virtual Cluster Management System: GriVon

## *Project Overview & VMware Roll*

<http://code.google.com/p/grivon/>

Takahiro Hirofuchi

t.hirofuchi at aist.go.jp


*National Institute of Advanced Industrial  
Science and Technology*

OSGC2008

# AIST Virtual Cluster Project

GriVon (Code name)



- Multi-site Virtual Cluster System
  - Easy-to-use distributed computing environment via single-system-image clusters
  - Maximum flexibility and scalability for cluster management



Cluster  
Manager  
Rocks Frontend

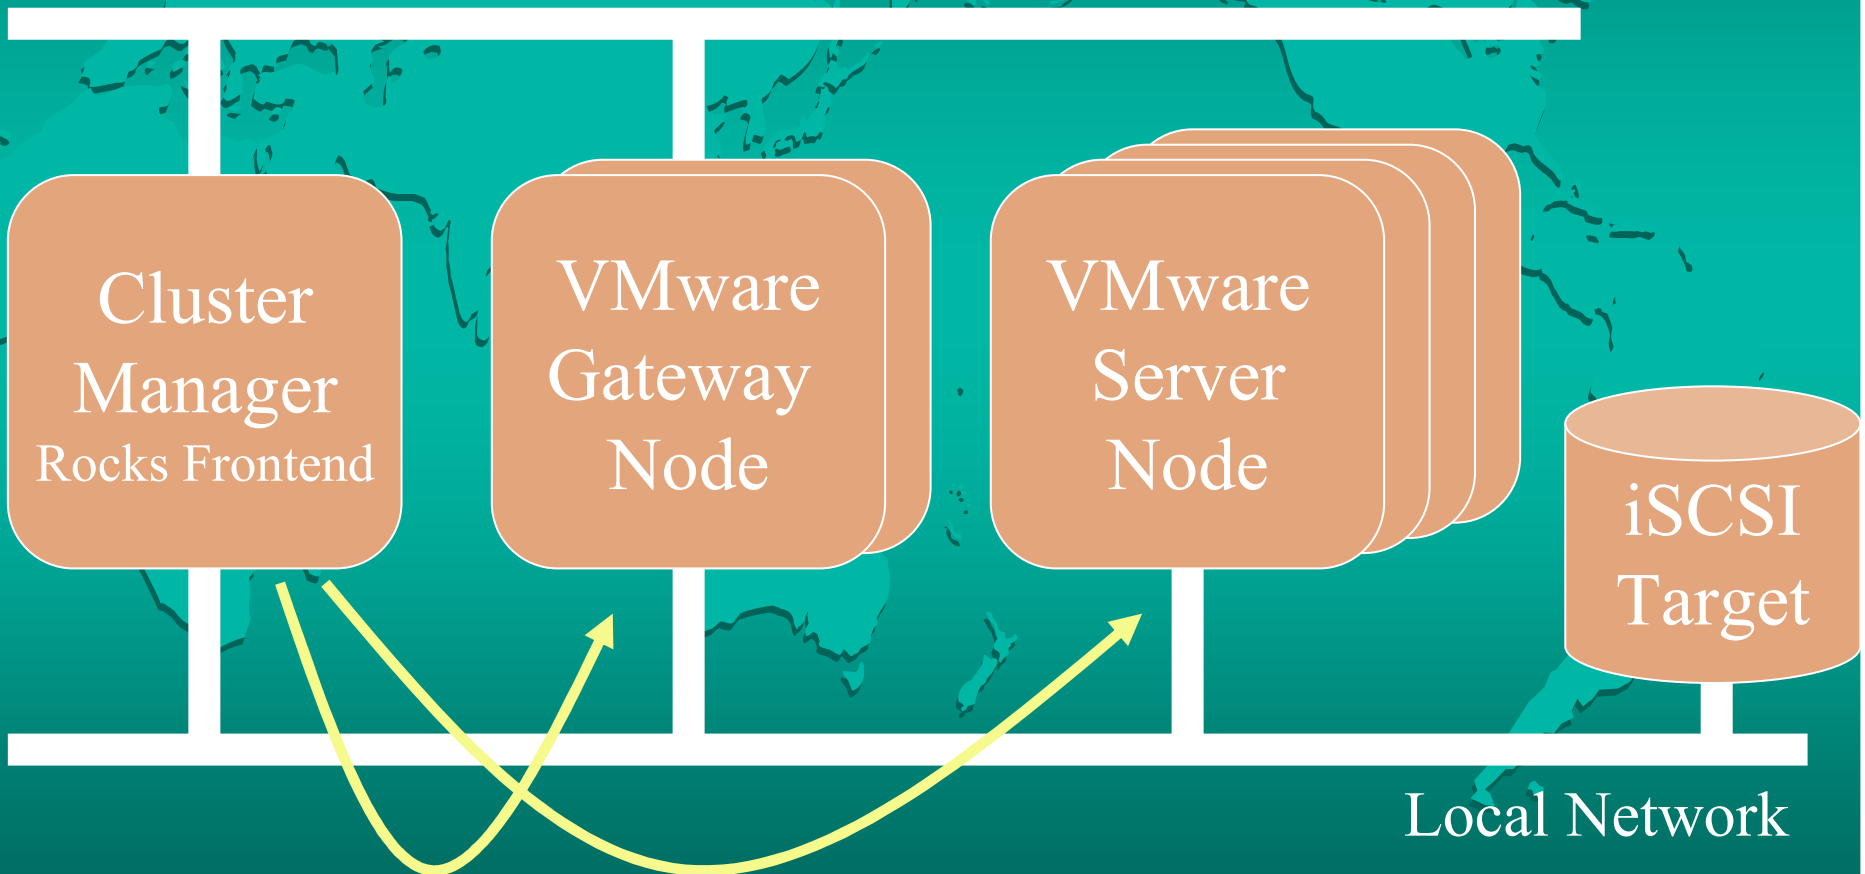
VMware  
Gateway  
Node

VMware  
Server  
Node



# GriVon System Overview

Public Network



Cluster  
Manager  
Rocks Frontend

VMware  
Gateway  
Node

VMware  
Server  
Node

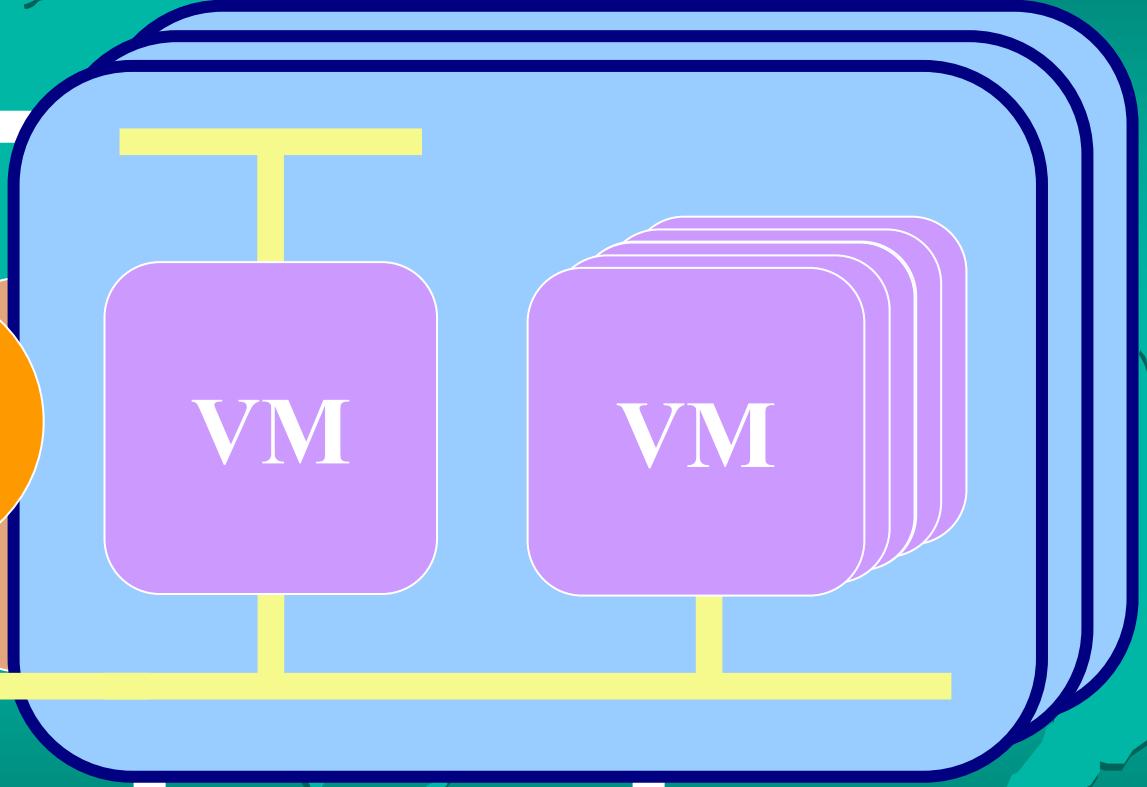
iSCSI  
Target

install

Local Network

Cluster  
Management  
Rocks Frontend

VPN



# Project Status

- A prototype system works for demos.
  - SC2007
  - waku-waku Akihabara Conference Demo 2008
- Todo (a lot, but now)
  - improve implementation quality
  - release code
  - get users



# **VMware Roll**

A node virtualization mechanism  
in  
the AIST virtual cluster system



# Design Criteria

- License Management
  - A VMware Server instance needs a unique serial key
  - Serial Key Pool
  - Distributed as a Roll source
- Multi-Frontend
  - A virtual frontend for each virtual cluster
  - A (virtual) frontend may exist in a remote cluster.
- Maximum Flexibility for Virtual Networking
  - Multiple virtual network interfaces for a virtual node
  - A VLAN is dynamically assigned for a new virtual cluster
  - No reinstallation for new settings

# Rocks Commands

add host vmware  
remove host vmware

create host vmware  
destroy host vmware

add host vlan  
remove host vlan  
list host vlan  
config host vlan

add host vmware interface  
remove host vmware interface  
list host vmware interface

add host vmwarenet  
remove host vmwarenet  
list host vmwarenet  
config host vmwarenet

add vmwarekey  
set host vmwarekey  
remove host vmwarekey  
list vmwarekey  
config host vmwarekey

start host vmware  
stop host vmware  
reboot host vmware  
list host vmware

suspend host vmware  
resume host vmware

plugins for removing host

# Database Tables

- vmware\_nodes
  - Remove Node Colum
    - The node inside a VM may be registered to another frontend.
- vmware\_macs
  - add IfIndex Colum
    - its interface number inside a VM (e.g., eth0, eth1 ...)
- vmware\_disks
- vmware\_vmnets
- vmware\_vmnet\_members
- vmware\_serials

# Add VLAN interfaces


```
vizzy:# rocks add host vlan vmware-server-0-0 viface=eth0.100
Added VLAN with VID == 100 to IF -:eth0:-
```

```
vizzy:# rocks add host vlan vmware-server-0-0 viface=eth0.101
Added VLAN with VID == 101 to IF -:eth0:-
```

```
vizzy:# rocks add host vlan vmware-server-0-0 viface=eth0.102
Added VLAN with VID == 102 to IF -:eth0:-
```

\* The added VLAN interfaces are now active.  
The next reinstallation also sets the interfaces automatically.

```
vizzy:# rocks list host interface vmware-server-0-0
SUBNET  IFACE      MAC                IP                NETMASK  GATEWAY  MODULE  NAME
private eth0       00:30:1b:b3:24:72  10.255.255.254   255.0.0.0  -----  tg3     vmware-server-0-0
----- eth0.100   -----
----- eth0.101   -----
----- eth0.102   -----
```



# Add VMware Network Settings

```
vizzy:# rocks add host vmwarenet vmware-server-0-0  
iface=eth0 type=bridge
```

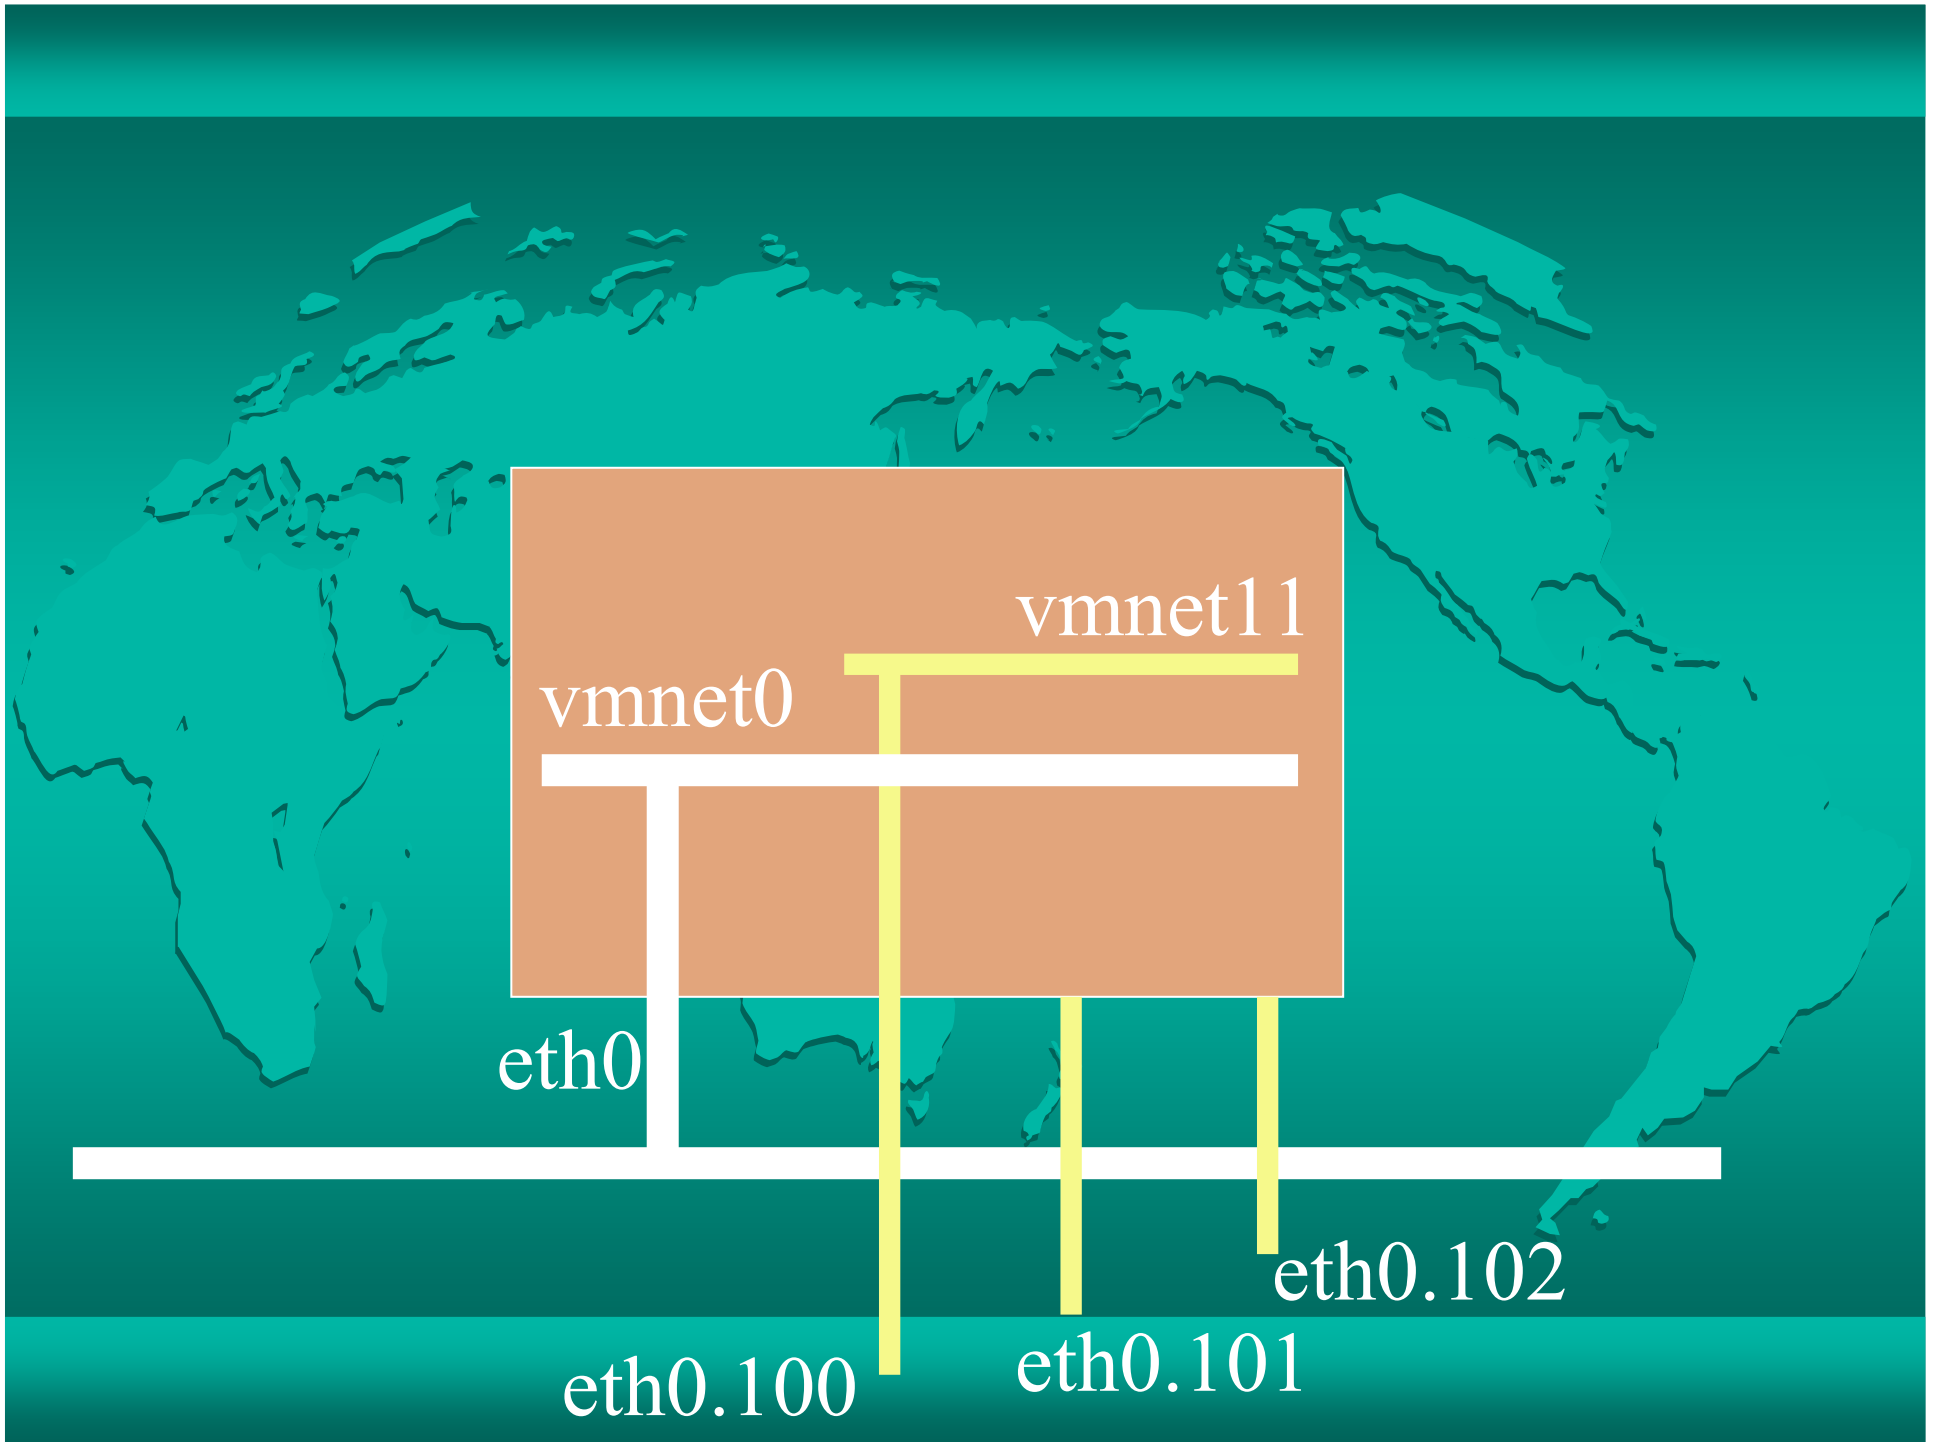
```
vizzy:# rocks add host vmwarenet vmware-server-0-0  
iface=eth0.100 type=bridge
```

```
vizzy:# rocks list host vmwarenet
```

| VMNET   | TYPE   | DEVICE   | HOSTADDR | NETMASK |
|---------|--------|----------|----------|---------|
| vmnet0  | bridge | eth0     | -----    | -----   |
| vmnet11 | bridge | eth0.100 | -----    | -----   |

\* The added VMware network settings are now active.  
The next reinstallation also sets the VMware network settings automatically.

\* It supports "bridge", "NAT", and "host-only" network settings for each physical network interface. "bridge" is already implemented.



# Add Virtual Machines

```
vizzy:# rocks add host vmware vmware-server-0-0 cpus=1  
mem=768 disksize=20  
added VM on node "vmware-server-0-0" slice "0"
```

```
vizzy:# rocks add host vmware vmware-server-0-0 cpus=1  
mem=768 disksize=20  
added VM on node "vmware-server-0-0" slice "1"
```

\* These commands only add VM entries into the DB.



# Add VMs' Network Interfaces

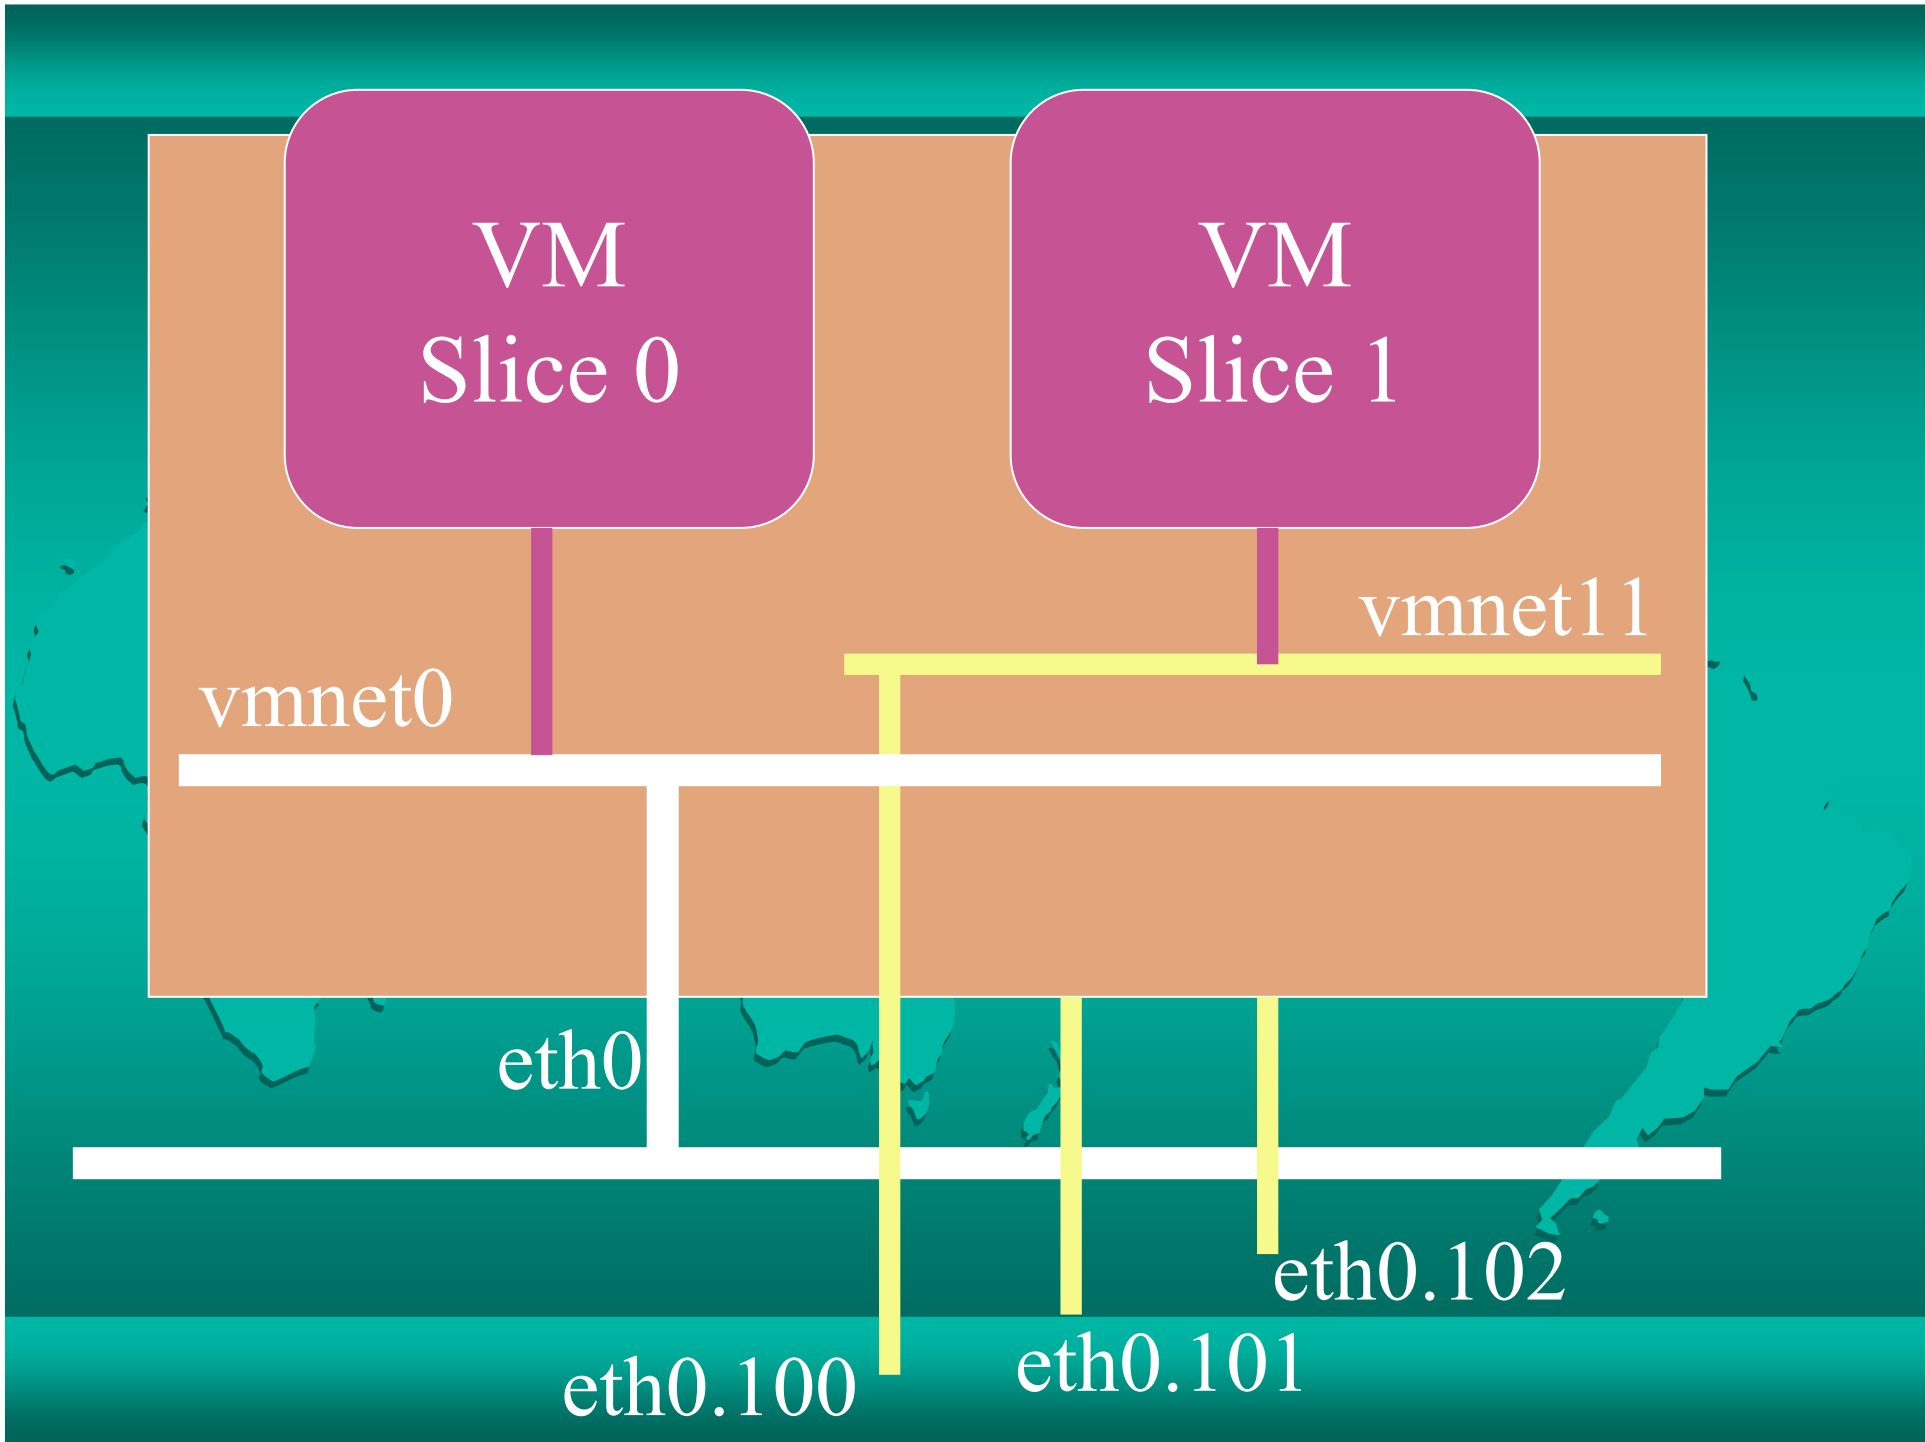
```
vizzy:# rocks add host vmware interface vmware-server-0-0  
slice=0 vmnet=vmnet0
```

```
adding vmware interface to 1 slices  
add 00:50:56:00:00:01 to slice0@vmware-server-0-0  
(vmnodeid 22) as ifindex 0; bound to vmnet0
```

```
vizzy:# rocks add host vmware interface vmware-server-0-0  
slice=1 vmnet=vmnet11
```

```
adding vmware interface to 1 slices  
add 00:50:56:00:00:02 to slice1@vmware-server-0-0  
(vmnodeid 23) as ifindex 0; bound to vmnet11
```

\* These commands only add VM entries into the DB.



# Create & Start VM

```
vizzy:# rocks create host vmware vmware-server-0-0 slice=0
```

```
vizzy:# rocks start host vmware vmware-server-0-0 slice=0
```

```
vizzy:# rocks list host vmware
```

| VMNODE                           | #CPU     | MEM        | #NIC     | #DISKS   | STATE       |
|----------------------------------|----------|------------|----------|----------|-------------|
| <i>slice0@vmware-server-0-0:</i> | <i>1</i> | <i>768</i> | <i>1</i> | <i>1</i> | <i>on</i>   |
| <i>slice1@vmware-server-0-0:</i> | <i>1</i> | <i>768</i> | <i>1</i> | <i>1</i> | <i>none</i> |

```
vizzy:# vmware-server-console -h vmware-server-0-0 &
```

```
vizzy:# vncviewer vmware-server-0-0:5900 &
```

\* VNC port number is "5900 + slice".



# Feature Summary

- License management
- Minimum (but sufficient?) VLAN support
- Flexible virtual networking mechanism
  - Multiple interfaces
  - No reinstallation
- Everything is done in the Rocks framework!

# Future Work

- Xen
  - “rocks commands” and database tables are similar, but not completely same.
- Auto frontend installation
- Multi-site support
- iSCSI support
- Virtual cluster reservation GUI
  - The prototypes of the above 4 features are already implemented in AIST

# Conclusion

- GriVon
  - Advanced virtual cluster system for large-scale computer centers and datacenters
  - Multi-site virtual cluster support
- VMware Roll
  - Very stable full virtualization
  - Flexible virtual network settings for GriVon
  - Available at <http://code.google.com/p/griyon/>

# Appendix





# VMware Server EULA

## 9.1(b)

use the Software solely for your own internal information processing services and computing needs in connection with permitted uses of the Software, including the hosting of computer application-based services from a Virtual Machine and provision of such services via an internal or external network, provided such services **may not consist of services to a third party that provide primarily computing or processing power (such as *utility computing or grid computing*) or any computer application-based service that is traded, rented, leased or sold on a Virtual Machine basis;**

**1. Go <http://www.vmware.com/download/server/>**

[Solutions](#) | [Products](#) | [Technology](#) | [Services](#) | [Resources](#) | [Customers](#) | [Partners](#) | [About Us](#)

Home > Downloads > Server Virtualization Products > VMware Server

VMware  
Server

Download VM

**2. Get Free Serial Numbers!**

To use the versions below, you will need to register for your free serial number(s).

Looking for the Server 2.0 (Beta)? [Download Here](#)

**VMware Server 1.0.5**  
Latest Version: 1.0.5 | 3/14/07

[Download Now](#)

**3. Download VMware Server & Console binary of Linux RPM**

[Version History](#)

[Drivers & Tools](#)

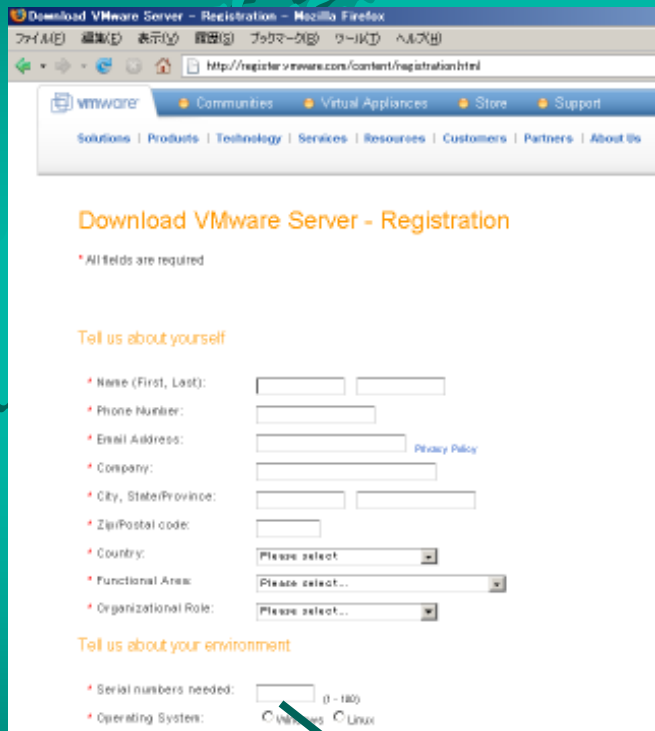
[Open Source](#)

**VMware Server 1.0.4**

Version: 1.0.4 | 9/18/07 | Build: 56528

[Download](#) | [Release Notes](#)

# Get Free VMware Serials



Download VMware Server - Registration

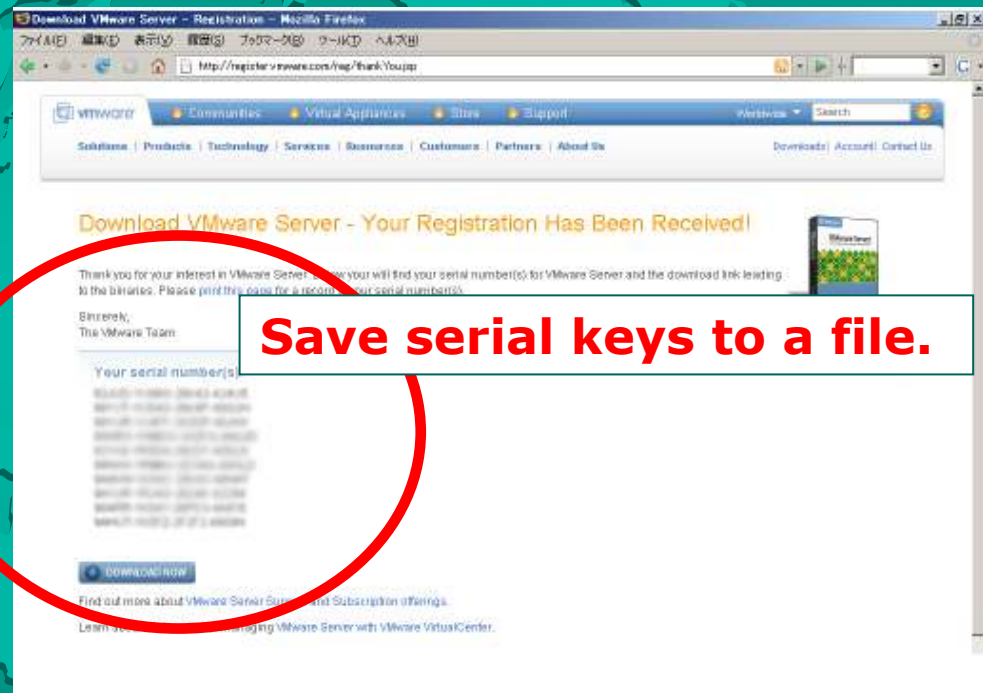
\* All fields are required

Tell us about yourself

- \* Name (First, Last):
- \* Phone Number:
- \* Email Address:  [Privacy Policy](#)
- \* Company:
- \* City, State/Province:
- \* Zip/Postal code:
- \* Country:
- \* Functional Area:
- \* Organizational Role:

Tell us about your environment

- \* Serial numbers needed:  0 - 100
- \* Operating System:  Windows  Linux



Download VMware Server - Your Registration Has Been Received!

Thank you for your interest in VMware Server. You will find your serial number(s) for VMware Server and the download link leading to the binaries. Please print this page for a record of your serial number(s).

Sincerely,  
The VMware Team

Your serial number(s)

[DOWNLOAD NOW](#)

Find out more about VMware Server Solutions and Subscription offerings.  
Learn how to integrate VMware Server with VMware VirtualCenter.

Save serial keys to a file.

Tell us about your environment

- \* Serial numbers needed:  (1 - 100)
- \* Operating System:  Windows  Linux

## Download VMware Server - Your Registration Has Been Received!



Thank you for your interest in VMware Server. Below you will find your serial number(s) for VMware Server and the download link leading to the binaries. Please print this page for a record of your serial number(s).

Sincerely,  
The VMware Team

Your serial number(s):

Serial Number(s) for VMware Server  
Serial Number(s) for VMware Server  
Serial Number(s) for VMware Server  
Serial Number(s) for VMware Server  
Serial Number(s) for VMware Server  
Serial Number(s) for VMware Server  
Serial Number(s) for VMware Server  
Serial Number(s) for VMware Server  
Serial Number(s) for VMware Server  
Serial Number(s) for VMware Server

**Save serial keys to a file.**

DOWNLOAD NOW

Find out more about VMware Server Support and Subscription offerings.  
Learn about the benefits of managing VMware Server with VMware VirtualCenter.

# Installation

- Download a VMware Roll source.
  - Extract it into somewhere
- Download VMware binary RPMs.
  - `VMware-server- $\{ver\}$ .rpm`
  - `VMware-server-console- $\{ver\}$ .rpm`
  - Copy them into  `$\{vmware-roll-root\}$ /RPMS/i386/`
- Do “make roll”
- Install the built roll
  - `“rocks add roll VMware-5.0-0.i386.disk1.iso”`
  - `“rocks enable roll vmware”`
  - `“cd /home/install; rocks-dist dist”`





# Assign Serial Keys to VMware Server Nodes

```
vizzy:# rocks set host vmwarekey vmware-server-0-0
```

```
vizzy:# rocks list vmwarekey
```

| SERIAL                   | HOST              |
|--------------------------|-------------------|
| XXXXX-XXXXX-XXXXX-XXXXX: | vmware-server-0-0 |
| XXXXX-XXXXX-XXXXX-XXXXX: | -----             |
| XXXXX-XXXXX-XXXXX-XXXXX: | -----             |
| XXXXX-XXXXX-XXXXX-XXXXX: | -----             |
| XXXXX-XXXXX-XXXXX-XXXXX: | -----             |

\* The assigned serial key is now active.  
The next reinstallation also sets the serial key to the host automatically.

\* If arg. is "vmware-server", assign all nodes at once.



# Create VMs in Nodes

```
vizzy:# rocks create host vmware vmware-server-0-0 slice=0
# creating slice0@vmware-server-0-0
ssh -x vmware-server-0-0 vmware-vdiskmanager -c -a
lsilogic -s 20Gb -t 0
/state/partition1/vmware/disks/0.scsi0:0.vmdk
Using log file /tmp/vmware-root/vdiskmanager.log
Creating a monolithic growable disk
'/state/partition1/vmware/disks/0.scsi0:0.vmdk'
Virtual disk creation successful.
ssh -x vmware-server-0-0 vmware-cmd -s register
/state/partition1/vmware/0/0.vmx
register(/state/partition1/vmware/0/0.vmx) = 1
```

# Start VMs in Nodes

```
vizzy:# rocks start host vmware vmware-server-0-0 slice=0  
ssh -x vmware-server-0-0 vmware-cmd  
/state/partition1/vmware/0/0.vmx start  
start() = 1
```

```
vizzy:# rocks list host vmware
```

| VMNODE                           | #CPU     | MEM        | #NIC     | #DISKS   | STATE       |
|----------------------------------|----------|------------|----------|----------|-------------|
| <i>slice0@vmware-server-0-0:</i> | <i>1</i> | <i>768</i> | <i>1</i> | <i>1</i> | <i>on</i>   |
| <i>slice1@vmware-server-0-0:</i> | <i>1</i> | <i>768</i> | <i>1</i> | <i>1</i> | <i>none</i> |

```
vizzy:# vmware-server-console -h vmware-server-0-0 &
```

```
vizzy:# vncviewer vmware-server-0-0:5900 &
```

```
* VNC port number is "5900 + slice".
```