# Automating Resource Management for Linux on z/VM: Lessons Learned

*SHARE 2009 Austin – Session 9294*

*Edmund R. MacKenty*
*Rocket Software, Inc.*

*Your primary source for enterprise software*


MAINSTAR

# Why Automation?

- Makes your life easier
- Less opportunity for mistakes
- Encapsulates knowledge of procedure details
- Computers are better at the repetitive tasks
- You can work on design, not implementation details

# Automation of Linux Guest Creation

- Ability to create and destroy Linux guests on demand
- Various z/VM resources must be made available to the automation tool
- Resources are assigned to “pools” by administrators
- Resources are dynamically allocated from these pools, and freed back into them, by the automation tool

# What Resources Are Needed?

- Directory Entries
  - Operations: Add, Delete, Insert item, Remove item
- DASD Devices
  - Operations: Allocate extent, Free extent
- Network Devices
  - Operations: Assign to guest, Remove from guest
- IP Addresses
  - Operations: Assign, Remove

# Mainstar Provisioning Expert

- What it is:
  - A tool for making Linux guests under z/VM
- What it does:
  - Configures the (virtual) hardware and software
- How it does it:
  - Defines guests in the z/VM directory
  - Assigns DASD and network devices to guests
  - Creates and populates Linux filesystems
  - Configures Linux system and application files
- All operations are completely automated

# Cloning vs. Configuring

- Cloning:
  - Creates an exact copy of an existing guest
  - Only the guest and host names and IP address change
  - Is a manual process, or done by scripts
  - Resource management is ad hoc
- Configuring:
  - Each instance can have different virtual hardware
  - Arbitrary arrangements of filesystems
  - Customize instance to specific needs of the workload
  - Process is automated, with well-defined resource management
  - Opens up lots of new possibilities

# Political Issues vs. Technical Issues

- Some Political Issues:
  - Who controls which resources?
  - Can different groups control different resource pools?
  - Who has access to which resource pools?
  - When are resources made available to the tools?
- Some Technical Issues:
  - How are resource pools implemented?
  - How are resource pools used?
  - Which resources does z/VM control?
  - Which resources does Linux control?

# Who Manages Resources?

- Several different groups have roles:
  - z/VM Admins
  - Network Admins
  - Security Admins
  - Linux Admins
- Each group has different interests/needs
- Communication between groups can be a problem
- Terminology differences can create confusion
- Automation tools can do parts of each groups roles

# Directory Entries

- Must have one entry for each guest created
- Parts of the entry vary for each guest:
  - Guest ID
  - Storage (memory) size
  - MDISK assignments
  - Network devices
  - LINKs to shared devices
- Common parts of the entry can be in a directory profile

# Directory Entries

- Political Issues:
  - Who controls the z/VM Directory?
  - How are changes made to it?
  - How is contention avoided?
  - How are security issues resolved?
  - z/VM admins are not used to dynamic entries
- Technical Issues:
  - Is anyone using XEDIT on the directory?
  - Use DIRMAINT directly or the Systems Management API?
  - Do any directory changes require RACF rule changes?
  - What goes into the directory profile?

# DASD Management

- Types of DASD
  - ECKD
  - FBA
  - FCP-attached SCSI
  - VDISK
- Uses of DASD
  - Physical devices
  - Logical devices
  - Filesystems
  - Swap
- Groups of DASD

# DASD Management

- Political Issues:
  - Who can allocate DASD?
  - Who owns the DASD?
  - Which DASD should be used for which purposes?
  - Should DASD be erased when it is freed?
- Technical Issues:
  - Access Permissions
  - Multiple DASD pools
  - How much DASD is needed?
  - Logical Volume Management
  - DATAMOVE

# FCP-Attached SCSI Devices

- Different from traditional DASD, not virtualized by z/VM
- z/VM knows it as a Real Device
- Linux must provide the WWPN and LUN values
- Cannot allocate portions of a device to different guests

# FCP-Attached SCSI Devices

- Political Issues:
  - Who assigns FCP devices?
  - Who knows the WWPNs and LUNs?
  - How do they get handed over to the Linux group?
- Technical Issues:
  - Should there be a managed pool of WWPNs and LUNs?
  - How to tell which devices are in use?
  - Linux has no access to SAN management tools

# DASD Sharing

- Shared DASD must be read-only, because of Linux block caching
- Shared DASD cannot be changed after creation
- Reduces amount of DASD needed for many instances
- Does not reduce storage requirements of instances

# DASD Sharing

- Political Issues:
  - Who can share the DASD?
  - Who owns shared DASD?
- Technical Issues:
  - Linking to the shared DASD
  - Preventing deletion of shared DASD while references exist
  - FCP-attached SCSI devices cannot be shared

# Network Device Management

- Several types of network devices:
  - VSWITCH (Layer 2 and Layer 3)
  - Guest LAN (Hipersocket, OSA)
  - Channel-To-Channel
- Device type depends on subnet definition
- Need network address, netmask, gateway, etc.
- Must support multiple subnets
- Allow for both internal and external connectivity

# Network Device Management

- Political Issues:
  - Who defines the network subnets?
  - Who controls the virtual network devices?
- Technical Issues:
  - How to represent available subnets?
  - Should the user select which subnet an instance is on?
  - Allow multiple network interfaces for an instance?
  - Should the tool create VSWITCHs dynamically?

# IP Address Management

- Static vs. dynamic (DHCP) IP addresses
- IP addresses are defined within some subnet
- Subnets can be shared by several systems
- Available IP's on a subnet are defined by the Network Gods
- Automated tool cannot control all aspects of network
- Would like to assign static IP's from a pool of IP addresses

# IP Address Management

- Political Issues:
  - Who assigns IP address to guests?
  - When are IP addresses assigned?
  - Should IP addresses be reserved in a resource pool?
  - Should IP addresses by assigned dynamically? (not DHCP)
- Technical Issues:
  - Should the user select the specific IP address for an instance?
  - Can any available subnet and IP address be used?
  - Should IP address pool be managed by z/VM or Linux?

# Access Control

- Unfettered automated resource managers can be dangerous
- We must prevent some users from doing some operations
- Users should have access to the minimum they need to do their work
- Need to limit access to:
  - Each type of resource
  - Each instance of each type of resource
  - Each operation on a type of resource

# Access Control

- Political Issues:
  - Who is allowed to allocate resources?
  - Who is allowed to use resources?
  - Who is allowed to free resources?
  - How are production systems protected?
- Technical Issues:
  - How to authenticate users?
  - How to define permissions for each user, object and operation?
  - How to restrict user views based on access permissions?

# Security Management

- Resource automation tools must work with established security policies
- Certain z/VM resources are often protected by RACF
- Creating directory entries, linking devices, etc. may require RACF rules
- Need a way to dynamically create RACF rules when dynamically allocating resources
  - Exit called on z/VM operations
- Security levels and methods vary from site to site

# Security Management

- Political Issues:
  - Who sets up RACF rules?
  - Who is allowed to change the rules?
  - Can RACF rules be allowed to be automated?
- Technical Issues:
  - RACF rules are very site-specific
  - Which operations need RACF operations?
  - How to undo changes to RACF rules?

# Summary

- Automating the management of each type of resource has its own set of problems
- Different groups must be involved for each type of resource
- Automation tools may change when a group assigns a resource, but not which group does the assignment
- What each group will allow to be automated varies from site to site

# Questions?

## Automating Resource Management for Linux on z/VM: Lessons Learned

*SHARE 2009 Austin – Session 9294*

*Contact Information:*
*Presenter:* Ed.MacKenty@RocketSoftware.com
*Company:* http://www.rocketsoftware.com