

Sine Nomine Associates

Creating Linux Appliances on z/VM

David Boyes
Sine Nomine Associates
MVMUA, January 24, 2005

© 2005 Sine Nomine Associates

Sine Nomine Associates

Note: Work in Progress!

- This presentation is a work in progress for zExpo in April. You are the lucky guinea pigs for the first shot at it.

Please feel free to comment – all improvements welcome!

© 2005 SNA

What is an Appliance?

- Thinking about it, a Linux guest is very similar to a dedicated microcontroller:

For our purpose, an appliance is a dedicated Linux guest that provides or augments a function that the host OS cannot (or is unlikely to) provide.

Appliances vs Embedded Systems

- Embedded System
 - Minimize footprint
 - Minimize writable storage
 - Focus on at most a few services (not general purpose)
 - Usually operate with little or no human intervention
- Appliances
 - Not necessarily constrained for footprint and R/O space, but desirable to keep to minimum
 - Less focused, but still desirable to optimize for single or small number of function
 - Sometimes operated with no human intervention, but not design requirement

Sine Nomine Associates

Major Components

- Container
- Connectivity
- Minimal Services
- Selective Hardening
- Host System Interfaces
- Applications

© 2005 SNA

Sine Nomine Associates

Container

- Container = VM userid
- Provides:
 - Boot loader
 - Most error recovery
 - Basic dispatcher function
 - Definitions for integration points and devices
- Equivalent to I/O bus and device interface backplane for embedded system
- Typically:
 - 16-32M
 - Single 400 cyl minidisk system, or
 - DCSS based /usr

© 2005 SNA

Connectivity

- Major difference in VM guest design: network support is not optional
- Console support is less desirable, but is useful as trivial logging device (in lieu of syslog)
- Unique devices:
 - CTCA
 - IUCV
 - UR devices (rdr/prt/pun)
- Unusual omission: /dev/lp emulation for dev 00E.

Minimal Services

- Design feature: Install ONLY what you absolutely need.
- Difficult to perform with RH and SuSE
 - Packages compiled with many options, which prereqs a lot of large stuff
 - Running usable system usually requires 128M and 1000 cyles or more.
- Easy with Debian and Slackware
 - Packages compiled with minimum options to make package function.
 - Running usable system can be built in 12M and less than 300 3390 cyl
- Both approaches work, but require different attacks for hardening and application support.

Minimal Services

- Strict separation of R/W and R/O data
 - Application and OS data are separate
 - /usr commonly R/O (good use for DCSS)
 - /usr/local, /var, /etc usually R/W

Selective Hardening

- Minimal Install approach requires only basic password and setuid executable monitoring, as only the services that are actually needed are installed.
- This approach assumes that you start with an absolutely bare-bones installation and add only what you need.

Both Debian and Slackware are very amenable to this approach; the "minimum install" for these distributions fits the bill nicely.
- Pare-down approach relies on comprehensive scripts to disable only the unnecessary function.

Use of RH and SuSE as a base dictates this option, as their packages enable some known vulnerabilities, but can be semi-automated using some analysis scripts.

Example: The Bastille Project maintains a set of scripts for common distributions that removes a major set of common vulnerabilities. See www.bastille.org for details.

Host System Interfaces

- IP network sockets
 - Most widely used in concert with CMS Pipelines
 - NFS is also a handy trick.
- UR driver by Malcolm Beattie
 - Hidden in an obscure ITSC redbook, but very handy
- IUCV support
 - Part of low-level code in the IUCV network driver
 - Needs generalization, but is abstract enough to be usable
- CTC support
 - Similar to IUCV driver, but useful for communicating with non-VM environments (idea: NJE/IP gateway for JES?)
- CMSFS
 - Reads CMS minidisks
- E2
 - Pipe stages for ext2fs support for CMS applications

Applications

- The actual meat of the problem, but probably the most open-ended part:
 - What applications fit well here?
 - What are some things that have already been tried?
 - SMTP server replacement
 - SSL Enabler
 - NFS Server for CMS data
 - Tape Server
 - WWW Server
 - NJE/IP Gateway for z/OS (under construction)

Sine Nomine Associates

Q&A

© 2005 SNA

Sine Nomine Associates

Contact Info

Sine Nomine Associates
info@sinenomine.net
www.sinenomine.net

© 2005 SNA

Sine Nomine Associates

© 2005 SNA