

Tools for Linux System Management in the Data Center Environment



Prepared for:
Linux Technology Days 2003
December 4, 2003

Scott Courtney
Sine Nomine Associates



Contents

- Introduction
- Challenges of Linux in the Data Center
- Management Tools and Methods
- Monitoring Tools
- Backup/Disaster Recovery



Introduction

- This talk will cover...
 - Challenges facing Linux administrators in the data center
 - Suggestions of strategic approaches to problems
 - Useful Open Source and commercial software tools
- This talk will *not*...
 - Endorse specific commercial products
 - Address selection of hardware



Challenges of Linux in the Data Center

- Linux is still a distributed operating system
- Linux developers initially focused on single servers, not data centers
- Need for integration with existing management environment and tools



Management Tools: Strategies

- Design to minimize the problem in the first place
 - Shared read-only filesystems for static data
 - Configurations in LDAP or SQL rather than flat files, wherever practical
 - Horizontal vs. vertical scalability decisions
- Look at your existing server management tools
 - Can you extend what you already have to a new platform?
 - If not, is this the time to consider a broader-based tool set for the enterprise as a whole?



Management Tools: Basics

- SuSE Linux Enterprise Server components
 - Auto YaST for software distribution
 - Data Center “Pack” for console access, admin SSO, and distributed command support
- rsync and its ilk for automated file distribution
- CVS can be useful for managing config files, web pages, etc., not just source code



Management Tools: VMWare

- Uses “Virtual 86” mode in Intel CPU for *partial* hardware virtualization
- Virtual hardware environment, not just software emulation of API set
- Overhead is reasonable, but not negligible
- Several versions available: Workstation, GSX, ESX
- ESX version has features like resource capping useful in larger enterprises
- <http://www.vmware.com/>



Management Tools: VMWare (cont'd)

- Transactional filesystem model supports “snapshots” and rollback of changes
- Surprisingly good virtualization and compatibility even with badly-behaved guest o.s. such as W2K
- Seriously RAM-intensive, even in a test environment
- ESX and GSX versions need really beefy hardware
- **Use for consolidating under-utilized servers**
- **Use for creating highly dynamic on-demand test environments**



Management Tools: AFS



- Andrew File System
Open Source, portable (15+ platforms, including Linux, UNIX, Windows)
- Network-based, distributed, from the ground up
- Unlike NFS, works well over WAN or Internet
- Kerberos security and sophisticated ACLs
- Local cache servers, transparent write-back
- Complex to set up, but significant benefits in large-scale environments



Monitoring Tools

- Availability monitoring
 - SNIPS (<http://www.netplex-tech.com/snips/>)
Very flexible, customizable; primitive installation
 - Big Brother (<http://bb4.com/>)
Simple, popular, widely -supported
 - Argus (<http://argus.tcp4me.com/>)
Nice web-based front end, relatively simple (but useful) monitoring of available services
 - Sysmon (<http://www.sysmon.org/>)
 - Heartbeat (<http://www.nolongervalid.com/heartbeat/>)
Availability monitoring, also useful for hot failover
 - Centralized syslogd monitoring, and various utilities to parse and report from those log files



Monitoring Tools

- Availability monitoring
 - Argus




Monitoring Tools

- Performance monitoring
 - SNIPS
(sometimes with custom modules)
 - vmstat, iostat
(probably part of your distro)
 - Loggerithim (<http://www.loggerithim.org/>)
Concise display of load, uptime, etc. -- basics
 - In mainframe environment, *use what VM can do!*



Monitoring Tools

- Performance monitoring
 - Loggerithim



Backup/Disaster Recovery Tools (Strategic Approach)

- First, minimize the scope of the backup task
 - Shared data methods, such as NFS or AFS
 - Careful filesystem planning (static filesystems backed up only when something changes)
 - Storage area network (SAN)
- DR is more than just backup and restore
 - What are you already doing on other platforms?
- Backup and restore from user perspective vs. from recovery perspective



Backup/Disaster Recovery: Products



- Commercial offerings
 - Tivoli Storage Manager (TSM, formerly ADSM)
 - CA BrightStor
 - Veritas
- Open Source offerings
 - Amanda
- SAN drive array built-in or vendor-specific features



Resources on the Web

- Official Linux Home Page
<http://www.linux.org/>
- Sine Nomine Associates
<http://www.sinenomine.net/>



Contact Information

Scott D. Courtney
Senior Engineer
Sine Nomine Associates
43596 Blacksmith Square
Ashburn VA 20147-4606
scourtney@sinenomine.net
<http://sinenomine.net/>



Presentation Copyright © 2003 by Sine Nomine Associates. All Rights Reserved.
Some graphics Copyright Rainhead Media (<http://www.rainheadmedia.com/eng/>), used by permission.
Trademarks are acknowledged as the property of their respective owners and are used for reference only.