

Administering Linux in Production Environments

Æleen Frisch
aefrisch@lorenzian.com
www.aeleen.com

Exponential Consulting, LLC
North Haven, Connecticut, USA

Itinerary

- Introduction
- Production Environment Features
 - Recent Kernel Developments
 - Filesystems: Mundane and Advanced
 - Disk Striping and RAID
 - Parallel Processing and Clustering
 - Enterprise Networking Features
- Deployment Examples
 - File and Print Servers
 - Enterprise User Authentication
 - Beowulf Compute Servers
 - Linux and Databases
 - Linux as an Office PC

Administering
Linux in
Production
Environments

Copyright © 1999-2001,
Exponential Consulting, LLC

2

What is a Production System?

- Real world
- System is a tool
- “Money” is involved

Administering
Linux in
Production
Environments

Copyright © 1999-2001,
Exponential Consulting, LLC

3

Commercial Applications

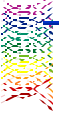
- Major applications are available now:
 - Databases: *Oracle*, *Sybase*, *DB2*, etc.
 - Computational chemistry: *Gaussian 98*
 - CAE: *MSC:Nastran*
 - Others
- Keeping up
 - www.linas.org/linux
 - www.linuxports.com

Administering
Linux in
Production
Environments

Copyright © 1999-2001,
Exponential Consulting, LLC

4

5



Dressing for Success

- Tuxedo vs. Business Suit
- ILM: 11/15/2001

Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC.

5

6




Recent Kernel Developments

Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC.

6

7



2.4 Kernel

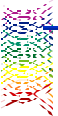
- Feature Lists:
 - lwn.net/2001/0111/a/ww2.4.php3
 - January 2001 Linux Magazine

Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC.

7

8



2.4: Numbers

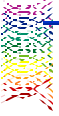
- 64 GB memory
- >2GB files
- 16 Ethernet adapters
- 10 IDE controllers
- SMP support
 - Support for tons of processes
 - Scheduler improvement
- Billions of users/groups

Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC.

8

9

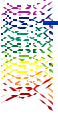


Current Linux Limitations

- Memory size: 64 GB
- File size: 2 TB
- Filesystem size: 2 TB (VFS limitations)
- Filesystem block size=Memory page
 - 4KB (x86)
 - 16KB (IA-64)

Administering Linux in Production Environments
Copyright © 1999-2001, Exponential Consulting, LLC.

10

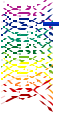


2.4: I/O

- Modes:
 - Separate block device and file I/O
 - Raw devices
- Devices:
 - I2O
 - USB
 - Firewire (IEEE1394)
 - PC Card
 - Infrared
- APCI support
- Graphics: Direct rendering manager
- SCSI2: Tagged command queuing

Administering Linux in Production Environments
Copyright © 1999-2001, Exponential Consulting, LLC.

11

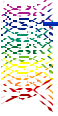


2.4: Networking

- Better multiprotocol support
- Rewritten network layer (firewalls, IP masquerading):
 - Packet filtering
 - Network address translation
- ATM and others
- Can now mount NFS3 shares

Administering Linux in Production Environments
Copyright © 1999-2001, Exponential Consulting, LLC.

12

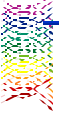


2.4: VFS and Related Facilities

- Single buffer for file caching
 - Eliminates synchronization problems
- Multiple mount points
- LVM in kernel
- RAID rewrite

Administering Linux in Production Environments
Copyright © 1999-2001, Exponential Consulting, LLC.

13



2.4 Bonus Features

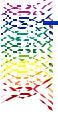
- Memory “Management”
 - New swap space size “recommendations”
 - Fixed around 2.4.10

Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC

12.1

14



2.4.15¹⁶

- ext3
 - “experimental”
- InterMezzo filesystem

Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC

13'

15



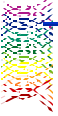
Filesystems: Mundane and Advanced

Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC

14

16



Itinerary

- The VFS
- Local Filesystems
- Journaling Filesystems
- Network Shared Filesystems
- Logical Volume Manager
- Distributed Filesystems

Administering Linux in Production Environments

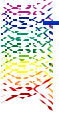
Copyright © 1999-2001, Exponential Consulting, LLC

15

17

Spelling 101

- Filesystem vs. file system



Administering Linux in Production Environments

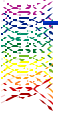
Copyright © 1999-2001, Exponential Consulting, LLC.

16

18

Virtual File System (VFS)

- Kernel subsystem/layer
- Provides a consistent interface for low-level file I/O
- Filesystem need only implement the required functionality using that interface, and it is automatically supported



Administering Linux in Production Environments

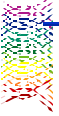
Copyright © 1999-2001, Exponential Consulting, LLC.

17

19

VFS Details

- Structured as an indirection layer
- Specifies low-level entities (objects) ...
 - Inodes, Files, Directories, Superblock,
 - Extended attributes
- ... and required/optional methods for each one



Administering Linux in Production Environments

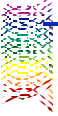
Copyright © 1999-2001, Exponential Consulting, LLC.

18

20

Filesystem Data Structures

- Superblock
 - FS metadata: label, block size, size, # inodes, ...
- Inodes
 - Properties: file type, owners, permissions, times, #links, size, ...
 - Data or Disk addresses or Single/double indirect
- Directory
 - = File that maps file names to inodes

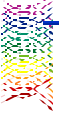


Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC.

19

21



VFS in Action

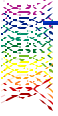
- Command/application invokes system call
- VFS looks up filesystem type in kernel table
- Kernel redirects call to FS-supplied method
- Method runs and accesses disk
 - Device drivers issue needed I/O requests
- Method returns descriptor to desired object
 - Descriptor contains pointers to functions for accessing that object as well as related data (e.g.: mounted filesystem, file, inode, dentry)

Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC

20

22



2.4.15+

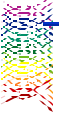
- Journal block device module: designed to add generic journaling capabilities to the VFS

Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC

21

23



Filesystems for Local Disks

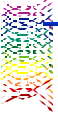
- Many, many supported types:
 - ext2
 - minix
 - CD-ROM: iso9660, MS Joliet extensions,
 - ❖ 2.4: udf for DVD
 - ufs, fat, vfat, umsdos, ntfs, sysv, affs, adfs, hfs, hpfs, qnx4, ...
 - ❖ 2.4: ufs nextstep extensions, efs, ramfs, jffs, cramfs

Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC

22

24



Special filesystems

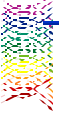
- procfs: /proc
- Pseudo-device: /dev/pts
- devfs:
 - /dev/hda => /dev/ide0/disk0/...
 - **devfsd** to support old device names

Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC

23

25



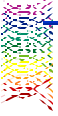
The ext2 Filesystem

- 2GB files and 4TB filesystem
- 255 character filenames
- SetGID directory group ownership inheritance
 - Selectable at mount time
- Variable block sizes
- Performance optimizations:
 - Read-ahead
 - Data block allocation and pre-allocation

Administering Linux in Production Environments
Copyright © 1999-2001, Exponential Consulting, LLC

24

26



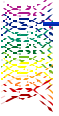
Ext2 FS Tools

- **fsck.ext2**
- **mke2fs**
- **e2label**
- **dumpe2fs**
- **tune2fs**
- **resize2fs**

Administering Linux in Production Environments
Copyright © 1999-2001, Exponential Consulting, LLC

25

27



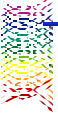
The mke2fs Command

- **-b** *block-size*
 - Default: 1024 bytes
- **-i** *bytes/inode* or **-N** *#inodes*
 - Default: 1 inode per 4096 bytes
- **-m** *reserve%*
 - Default: 5%
- **-L** *label*
- **-c** or **-l** *bad-block-file*
- **-f** *fragment-size*

Administering Linux in Production Environments
Copyright © 1999-2001, Exponential Consulting, LLC

26

28



dump2efs

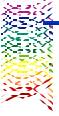
```

dumpe2fs 1.18, 11-Nov-1999 for EXT2 FS 0.5b, 95/08/09
Filesystem volume name: <none>
Last mounted on: <not available>
Filesystem UUID: 03d2f865-390a-4a5a-9162-8f2fc902d3e2
Filesystem magic number: 0xEPF3
Filesystem revision #: 1 (dynamic)
Filesystem features: (none)
Filesystem state: not clean
Errors behavior: Continue
Filesystem OS type: Linux
Inode count: 264928
Block count: 528948
Reserved block count: 26447
Free blocks: 260025
Free inodes: 184468
First block: 0
Block size: 4096
Fragment size: 4096
Last mount time: Sun Dec 3 09:33:59 2000
Last write time: Mon Jan 29 17:41:57 2001
Mount count: 12
Maximum mount count: 20
Last checked: Wed Sep 6 18:54:01 2000
Check interval: 15552000 (6 months)
Next check after: Mon Mar 5 17:54:01 2001
Reserved blocks uid: 0 (user: root)
Reserved blocks gid: 0 (group: root)
    
```

Administering Linux in Production Environments
Copyright © 1999-2001, Exponential Consulting, LLC

27

29



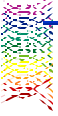
The tune2fs Command

- **-l** to list superblock info
- Modify attributes (mount read-only?)
 - **-C** *max-mounts*
 - **-C** *mount-count*
 - **-i** *check-interval***[dmw]**
 - **-e** *error-behavior*
 - ◊ *continue remount-ro, panic*
 - **-m** *reserved%* **or** **-r** *reserved-blocks*
 - **-g** *gid* **and/or** **-u** *uid*

Administering Linux in Production Environments
Copyright © 1999-2001, Exponential Consulting, LLC.

28

30



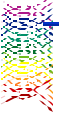
Contributed ext2 Tools

- **ext2fs defrag**
- **ext2fs resize**
- **ext2fsed**
- **ext2undelete**
- See: "Filesystems-HOWTO" (section 6)

Administering Linux in Production Environments
Copyright © 1999-2001, Exponential Consulting, LLC.

29

31



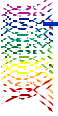
Journaling Filesystems

- A single inherent advantage over "traditional" UNIX filesystems

Administering Linux in Production Environments
Copyright © 1999-2001, Exponential Consulting, LLC.

30

32



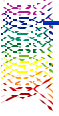
Journaling Linux Filesystems

- ReiserFS: 2.4.1pre7 and later (SuSE 6.4)
 - www.namesys.com
- ext3: coming in standard 2.4.15
 - e2fsprogs.sourceforge.net/ext2.html
 - www.zip.com.au/~akpm/linux/ext3/ext3-usage.html
- SGI's XFS: 1 May 2001
 - 64-bit; streaming video
 - oss.sgi.com/projects/xfs
- IBM JFS: 28 June 2001
 - 64-bit
 - oss.software.ibm.com/developerworks/opensource/jfs

Administering Linux in Production Environments
Copyright © 1999-2001, Exponential Consulting, LLC.

31

33



ext3

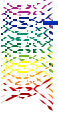
- **mke2fs**'s **-j** option
- Specify separate journal device: **-J device**
- Convert existing with **tune2fs -j**
 - Space considerations
- Interoperability with/as ext2

Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC

32

34



Reiser Filesystem

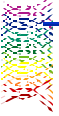
- Designed for performance, especially on directories with lots of files, as well as efficient usage of disk space
- Tools:
 - **mkreiserfs [-b n]** *n must be 4*
 - **reiserfsck**
 - **resize_reiserfs -s [±]size[kM]**

Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC

33

35



Reiser mount Options

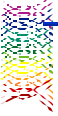
- **mount -o remount,resize=*new-size***
- **mount -o conv**
 - Reiser 3.5 to 3.6
- **mount -o notail**

Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC

34

36



A Testimonial

>From Source Forge

<http://ftp.sourceforge.net/> has 850GB storage, half of which is reiserfs, half is ext2. Both filesystems have been running flawlessly for > 4 months of production (actually longer, but wasn't reiserfs before). That server pushes between 15Mbit and 50Mbit/sec, and pulls/syncs about 2-5Mbit/sec, 24x7.

reiserfs also powers the CVS tree filesystem for cvs-mirror.mozilla.org (also tokyojoe.sourceforge.net), which is the one and only anonymous CVS checkout point for mozilla. That server has run flawlessly under very heavy load since its birth.

I don't get involved in kernel politics, but as a production filesystem, reiserfs is ok in my book.

Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC

35

37

XFS

- `mkfs -t xfs -l size=6000b /dev/hde1`
- `mount -t xfs -o logbufs=8,logbsize=32768 ...`

Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC

36

38

JFS

- `mkfs -t jfs -s 8 /dev/hde1`
- `logredo device`

Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC

37

39

Comparison

	100% Released	Max. size	Block size	Resize	Dyn. Index	Use as Boot FS	RAID	LVM	NFSv3	In std. kernel	Ext. Attr.	ACLs	Sparse files	Interp.
ext3	N	1TB	4K	Y	N	Y	Y	Y	Y	Y	Y	Y	N	n/a
Reiser	Y	1TB	4K	Y	Y	Y	Y	Y	Y	Y	4	N	4	n/a
XFS	Y	1TB	4K	Y	Y	Y	Y	Y	Y	N	Y	Y	Y	Y
JFS	Y	1TB	4K	Y	Y	Y	Y	Y	Y	N	N	N	Y	N

Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC

38

40

Journaling and dump

- ext3: **yes**
 - Long history
- ReiserFS: **no**
- XFS: **yes** (`xfs {dump,restore}`)
- JFS: **no**

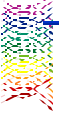
Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC

38'

Thanks Chris Marble!

41



Mounting Remote Filesystems

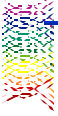
- NFS
- Samba: smbfs
- AFS
- ncpfs

Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC

39

42



Linux NFS

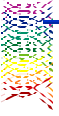
- Kernel-based NFSv2
- Kernel-based NFSv3
- User space NFSv2

Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC

40

43



NFS Components

- Daemons:
 - portmap, rpc.mountd, rpc.nfsd
 - ◊ rpcinfo -p [host]
 - rpc.lockd, rpc.statd
- /etc/exports file
 - *in* host(options) ...


```

/data2 dalton(rw) pascal(ro) henry(rw,all_squash)
*.vader.com(rw,sync)
/data3/new 192.123.12.0/255.255.255.0(ro)
                    
```

Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC

41

44



Other exports Options

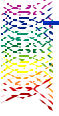
- sync, async
- root_squash, all_squash
- anonuid=*n*₁, anongid=*n*₂

Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC

42

45



Useful mount Options

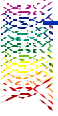
- **bg** **retry=//**
- **intr** *or* **soft**
- **retrans=//**
- **nosuid**
- **rsize=8192, wsize=8192** (max. for v2)

Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC

43

46



UID Mapping

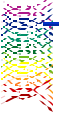
- **rpc.ugidd**
- **map_daemon** /etc/exports option

Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC

44

47



NFS and Security

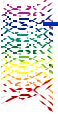
- Limit access to specific machines; avoid blanket exports
- Export read-only whenever practical
- Don't export group-writable files
- Don't export system files (incl. executables)
- Don't use the **insecure** option (allows access from any port)

Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC

45

48



Sharing Filesystems: Samba

- Samba is the solution for sharing filesystems and printers between UNIX and Windows systems:
 - www.samba.org
- Books:
 - Gerald Carter with Richard Sharpe, *Teach Yourself Samba in 24 Hours* (SAMS, Indianapolis, 1999); ISBN: 0-672-31609-9
 - Robert Eckstein, David Collier-Brown and Peter Kelly, *Using Samba* (O'Reilly, 2000); ISBN: 1-56592-449-5

Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC

46

Samba

- Implements ~~CIFS~~ SMB protocol under UNIX
 - Server Message Block is the native protocol for Microsoft networking file/printer sharing:

TCP/IP	OSI	Microsoft LAN
Application	Application	SMB
Transport	Presentation	NetBIOS
Internet	Session	
Network Access	Transport	
	Network	
	Data Link	
	Physical	

Administering Linux in Production Environments
Copyright © 1999-2001, Exponential Consulting, LLC

47

Samba

- Features:
 - Filesystem sharing
 - Printer sharing
 - Master browser
 - Domain security
 - Primary domain controller (alpha)

Administering Linux in Production Environments
Copyright © 1999-2001, Exponential Consulting, LLC

48

Samba Daemons

- **smbd** (TCP)
 - Provides sharing services
- **nmbd** (UDP):
 - Handles NetBIOS name server requests
- Execution options:
 - Standalone: **-D**
 - Via **inetd**

Administering Linux in Production Environments
Copyright © 1999-2001, Exponential Consulting, LLC

49

Samba with inetd

- For ~~inetd~~ control, modify `/etc/inetd.conf`

```
netbios-ssn stream tcp nowait root /usr/sbin/smbd smbd
netbios-ns dgram udp wait root /usr/sbin/nmbd nmbd
```
- and `/etc/services`:


```
netbios-ssn 139/tcp
netbios-ns 137/udp
```

Administering Linux in Production Environments
Copyright © 1999-2001, Exponential Consulting, LLC

50

53

The Samba Configuration File

- /etc/smb.conf
 - [global] section
 - Exported filesystem sections
 - [homes] section: defaults for user home directories
 - ❖ Actual services created on the fly
- **testparm**: verify configuration file structure

51

54

global Section

```
[global]
hosts allow = vala pele
hosts deny = lilith
valid users = dagmar, @chem, @phys, @bio, @geo
invalid users = root, admin, bin, system, daemon
max log size = 2000
username map = /etc/smbusers
```

KB

Map file entries: *linux = translation* (usually Windows)

```
aefrisch = aeleen
sysadmin=Administrator
chem = @chemistry
```

52'

55

Defining Shares

<pre>[chemdir] path = /chem/data/new comment = New Data read only = no case sensitive = yes force group = chemists read list = dagmar, @chem, @phys write list = @chem browseable= no admin users = chavez</pre>	<p><i>Define a directory for export.</i></p> <p><i>Local path to be shared.</i></p> <p><i>Description of filesystem.</i></p> <p><i>Filesystem is not read-only.</i></p> <p><i>Filenames are case sensitive.</i></p> <p><i>Map all users to this group.</i></p> <p><i>Users/groups w/ read access.</i></p> <p><i>Write access list.</i></p> <p><i>Exclude from browse lists.</i></p> <p><i>Administrative users.</i></p>
--	---

53'

Group=NIS netgroup (&) or UNIX group (+)

56

Connecting to Samba Shares

- Run **net use** as usual on the Windows system:
 - **net use s: \\dalton\chemdir**
 - **net use x: \\dalton\homes**
 - **net use x: \\dalton\username**

55

57

User Home Directories

```
[homes]
comment = Home directories
writeable = yes
valid users = %S
```

- Effect of map files:
 - `\\server\Administrator ⇒ \\server\sysadmin`

Administering Linux in Production Environments
Copyright © 1999-2001, Exponential Consulting, LLC

54'

58

Samba Utilities

- **smbstatus**
- **smbrun**
- **smbclient**
- **smbtar**

- GUI admin tool: **swat**
 - Runs within a browser

Administering Linux in Production Environments
Copyright © 1999-2001, Exponential Consulting, LLC

56

59

Mounting Windows Filesystems

- **smbmount**
 - Two versions: smbfs vs. standalone
 - ❖ `smbmount //dalton/chem -c 'mount /mnt ...'`
 - ❖ `smbmount //dalton/chem /mnt`
- `mount -t smbfs \`
 - o `username=aefrisch,password=xxx \`
 - `//dalton/chem2 /chem`

Administering Linux in Production Environments
Copyright © 1999-2001, Exponential Consulting, LLC

57

60

Passwords and Samba

- Add to `[global]` section:

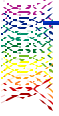

```
encrypt passwords = yes
security = user
```
- Run **mksmbpasswdsh** to create initial Samba password file:
 - `cat /etc/passwd | /path/mksmbpasswdsh`
 - ❖ Owner: root
 - ❖ Mode: 600
 - ❖ Directory mode: 500

Administering Linux in Production Environments
Copyright © 1999-2001, Exponential Consulting, LLC

58

61

More Authentication Options



- Authentication by a different server:


```
security = server
password server = host
encrypt passwords = yes
```
- Local (UNIX) authentication:


```
security = user
encrypt passwords = no
```
- Windows domain participation:


```
security=domain
workgroup = domain
password server = pdc bdc1 bdc2
encrypt passwords = yes
```

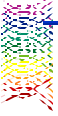
■ Samba server as the PDC

Administering Linux in Production Environments
Copyright © 1999-2001, Exponential Consulting, LLC.

59

62

Advanced Filesystem Features



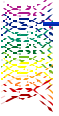
- Larger than 1 disk partition
 - Expandable on the fly
- Distributed across network
 - Load balancing
- Faster I/O
- Fault tolerant

Administering Linux in Production Environments
Copyright © 1999-2001, Exponential Consulting, LLC.

60

63

Logical Volumes



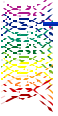
- Dynamically-resizable filesystem consisting of multiple, independent disk partitions (*physical volumes*), upon which a virtual structure is overlaid:
 - *Volume groups/virtual disks*, divisible into
 - *Logical volumes/virtual partitions*, which hold
 - *Filesystems*

Administering Linux in Production Environments
Copyright © 1999-2001, Exponential Consulting, LLC.

61

64

Linux Logical Volumes

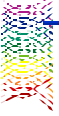


- Logical Volume Manager (lvm)
 - Still developing
- Veritas Volume Manager
 - \$\$\$\$

Administering Linux in Production Environments
Copyright © 1999-2001, Exponential Consulting, LLC.

62

65



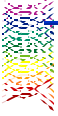
LVM

- www.sistina.com/products_lvm.htm
- Maximum filesystem size: 2TB
 - Limits: 99 VGs, 256 LVs
- Maximum logical volume: 256 GB with standard 4MB physical extent size

Administering Linux in Production Environments
Copyright © 1999-2001, Exponential Consulting, LLC.

63

66



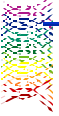
LVM Commands

- **pv** : create, change, display, move, scan
- **vg** : create, change, display, ck, cfgbackup/restore, export, extend, reduce, remove, split, merge, scan
- **lv** : create, change, display, extend, reduce, remove, rename, scan
 - **e2fsadmin** (requires **resize2fs**)
 - ❖ Resize filesystem and its underlying logical volume in a single operation

Administering Linux in Production Environments
Copyright © 1999-2001, Exponential Consulting, LLC.

64

67



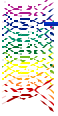
LVM Example

- Set partition type to **0x8E (fdisk)**
- **pvcreate /dev/sdb1 /dev/sdc1**
- **vgcreate vg1 /dev/sdb1 /dev/sdc1**
 - /dev/vg1/group
 - /etc/lvmconf/vg1.conf
- **lvcreate -L 2g -n biolv -r 8 -C y**
 - 8 read-ahead sectors, contiguous
- **mke2fs ... /dev/vg1/biolv**
- **mount /dev/vg1/biolv /somewhere**

Administering Linux in Production Environments
Copyright © 1999-2001, Exponential Consulting, LLC.

65

68



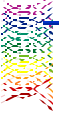
LVM Snapshots

- Designed to ensure consistent backups
 - **lvcreate --size /m --snapshot --name snap1 \ /dev/my_vg/homes**
 - **mount /dev/my_vg/snap1 /somewhere**
 - Back up /somewhere
 - **umount /somewhere**
- Uses standard VMM copy-on-write

Administering Linux in Production Environments
Copyright © 1999-2001, Exponential Consulting, LLC.

66

69



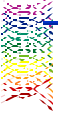
Distributed Filesystems

- Various focuses:
 - Handle client network file access and outages gracefully and seamlessly
 - ❖ Goal: remote files are indistinguishable from local files
 - Distribute network I/O among various servers
 - ❖ High availability
 - ❖ Redundancy
 - Share storage among various clients
- Lots of overlap

Administering Linux in Production Environments
Copyright © 1999-2001, Exponential Consulting, LLC

67

70



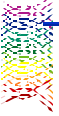
History

- AFS
- Coda
- Aura
- InterMezzo

Administering Linux in Production Environments
Copyright © 1999-2001, Exponential Consulting, LLC

68

71



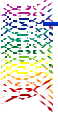
InterMezzo

- www.inter-mezzo.org
- Designed for high availability
- A palimpsest on Coda
 - Designed to be simpler

Administering Linux in Production Environments
Copyright © 1999-2001, Exponential Consulting, LLC

69

72



Concepts

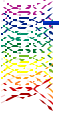
- *Fileset*: directory subtree (“folder collection”)
 - Basic unit served
 - Currently is the entire filesystem
- Change log
 - Filesystem-like journaling to track modifications
- Replication
 - Server to (duplicate) server
 - Client to server after reconnect
- Concurrent conflict handling
 - Current: detect and die

Administering Linux in Production Environments
Copyright © 1999-2001, Exponential Consulting, LLC

70

73

Components



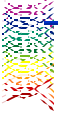
- Lento: cache manager and file server
 - User space daemon
- Presto: kernel module (intermezzo.o)

Administering Linux in Production Environments
Copyright © 1999-2001, Exponential Consulting, LLC.

71

74

Installation



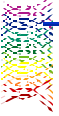
- Kernel support
- Kernel loopback support
- Initial RAM disk for booting
- Group InterMezzo (GID: 4711)
- **mkizofs [-t ext3] -r fsetname j /dev/hda n**
 - ReiserFS and XFS: maybe

Administering Linux in Production Environments
Copyright © 1999-2001, Exponential Consulting, LLC.

72

75

Converting Existing ext2



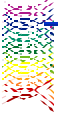
- **mount -t ext2 -o loop /tmp/cache /izo0**
- **mkdir -p /izo0/.intermezzo/nameofdb**
- **chgrp -R InterMezzo /izo0/.intermezzo**
- **chmod 700 /izo0/.intermezzo**
- **touch /izo0/.intermezzo/nameof{kml,lm,last_rcvd}**
- **tune2fs -j /tmp/cache**
- **umount /izo0**

Administering Linux in Production Environments
Copyright © 1999-2001, Exponential Consulting, LLC.

73

76

Configuration: Common



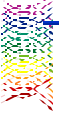
- ```
etc/intermezzo/serverdb
<serverdb> <server name="library" ipaddr="192.168.0.3" /> </serverdb>
```
- ```
etc/intermezzo/setdb
<setdb> <fileset name="test" servername="library">
<replicator>clientA</replicator>
</fileset> </setdb>
```
- ```
etc/fstab
/tmp/fs0 /izo0 intermezzo
loop, fileset="test", prestodev=/dev/intermezzo0,
mtp1=/izo0, cache_type=ext3, noauto 0 0
```
- ```
etc/conf.modules
alias char-major-185 intermezzo
```

Administering Linux in Production Environments
Copyright © 1999-2001, Exponential Consulting, LLC.

74

77

More Initial Setup



- `# dd if=/dev/zero of=/tmp/fs0 bs=1024 count=10k`
`# mkzifs -F /tmp/fs0`
- `# mknod /dev/intermezzo0 c 185 0`
`# chmod 700 /dev/intermezzo0`
- `# mkdir /izo0`
`# mount /izo0`
`# lento`

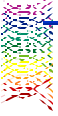
Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC

75

78

Configuration: Server



- `/etc/intermezzo/sysid`
`<sysid name="library" psdev="/dev/intermezzo0" bindaddr="192.168.0.3" />`

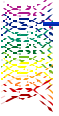
Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC

76

79

Configuration: Client



- `/etc/intermezzo/sysid`
`<sysid name="clientA" psdev="/dev/intermezzo0" bindaddr="192.168.0.20" />`

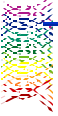
Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC

77

80

Checking the Configuration



- `config_check`

Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC

78

81

Global File System (GFS)

- www.sistina.com/products_gfs.htm
- Network-shared storage
 - Asynchronous journaling
 - Intelligent locking mechanisms
 - ❖ DMEP device/memexpd
 - Large files (64-bit addressing)
- Requires supported host-bus adapter
 - High speed interconnect like Fibre Channel
 - Storage is placed directly on the network

Administering Linux in Production Environments
Copyright © 1999-2001, Exponential Consulting, LLC

79

82

GFS Example

Administering Linux in Production Environments
Copyright © 1999-2001, Exponential Consulting, LLC

80

83

Setting up GFS

- Select hardware approach
 - FC=>Multiported SCSI=>Network Block Device
 - ❖ Future: Gigabit Ethernet TCP/IP storage
 - DMEP hardware=>IP daemon
- Plan configuration:
 - Lock server
 - Topology
 - Component storage
 - Power
 - Clients

Administering Linux in Production Environments
Copyright © 1999-2001, Exponential Consulting, LLC

81

84

More ...

- Kernel configuration: lots of patches
 - GFS & hardware (FC, for example)
- Build & distribute GFS software
- Configure

Administering Linux in Production Environments
Copyright © 1999-2001, Exponential Consulting, LLC

82

85

Configuring GFS

- Pool:
 - Abstraction for shared storage devices
 - Subpools for devices of same type

The diagram shows three 'GFS Client' boxes at the top, each connected to a central 'Storage Area Network' cloud. Below the network is a 'Network Storage Pool' box containing several 'Sub-pool' boxes, each with a 'Storage Device' icon.

Administering Linux in Production Environments
Copyright © 1999-2001, Exponential Consulting, LLC.

83 ■ Cluster information device (cidev): shared partitions which holds metadata

86

Configuring ...

- **gfsconf**
 - Lock sources
 - Callback ports
 - Timeouts
 - Clients
 - STOMITH methods
 - ❖ "Shoot the other machine in the head"
- Failover for **memexpd**
 - Still potential bottleneck

Administering Linux in Production Environments
Copyright © 1999-2001, Exponential Consulting, LLC.

84

87

Architecture

Also 2-layered:

- GLOCKS at filesystem level (lock caching)
- Modules to talk to devices

The diagram shows a vertical stack: 'User Space' (blue) at the top, followed by 'VFS' (yellow), 'GFS' (green), and 'Pool' (yellow). Below 'GFS' is a 'Locking Interface' (green) box. From the 'Locking Interface', arrows point to 'IP Lock' (orange) and 'DMEP' (orange) boxes. 'IP Lock' has an arrow pointing to a cloud labeled 'To Lock Server'. 'DMEP' has an arrow pointing to a cloud labeled 'Callbacks to other clients'. Below the 'Pool' are several yellow disk icons representing storage devices.

Administering Linux in Production Environments
Copyright © 1999-2001, Exponential Consulting, LLC.

85

88

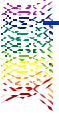
Disk Striping and RAID

Administering Linux in Production Environments
Copyright © 1999-2001, Exponential Consulting, LLC.

86

89

RAID



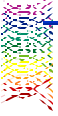
- Redundant Array of Inexpensive Disks
- Choices
 - Software
 - Hardware
 - ❖ Controller
 - ❖ Standalone device

Administering Linux in Production Environments
Copyright © 1999-2001, Exponential Consulting, LLC.

87

90

Software RAID Levels



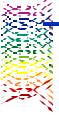
- Linear
 - Best large transfer I/O bandwidth
 - No loss of storage capacity
- 1: Disk mirroring
 - Best data redundancy
 - Good performance on small transfers

Administering Linux in Production Environments
Copyright © 1999-2001, Exponential Consulting, LLC.

88

91

More RAID ...



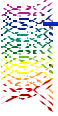
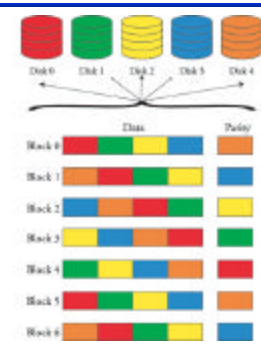
- 4: Disk striping with parity disk
 - Best fault-tolerant sequential file access
 - Not vulnerable to single disk failures
 - Parity disk is a bottleneck for writes
- 5: Disk striping with rotating parity block
 - Optimizes I/O operations/sec
 - Not vulnerable to single disk failures

Administering Linux in Production Environments
Copyright © 1999-2001, Exponential Consulting, LLC.

89

92

RAID 5

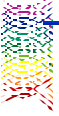



Administering Linux in Production Environments
Copyright © 1999-2001, Exponential Consulting, LLC.

90

93

Hybrid Levels



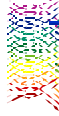
- RAID 0+1
 - Mirroring of a striped disk: two striped disks which are mirrors of one another. Data is striped across each stripe set, and the same data is sent to both striped disks. Thus, this RAID variation provides both I/O performance advantages and fault tolerance.
- RAID 1+0
 - Striping across mirror sets: Similar in intent to RAID 1+0, it provides equivalent performance advantages and slightly better fault tolerance in that it is easier to rebuild the RAID device after a single disk failure (since the data on only one disk is affected).

Administering Linux in Production Environments
Copyright © 1999-2001, Exponential Consulting, LLC.

91

94

Configuring RAID





- Enable kernel support
 - May need patches
 - ❖ ftp.kernel.org/pub/linux/daemons/raid
 - ❖ Red Hat installs for you
- Special files: `/dev/md*`
- Configuration file: `/etc/raidtab`
 - `mkraid` device
- Persistent superblock
 - Automatic detection of RAID entities
 - `raidstart/raidstop` to control manually

Administering Linux in Production Environments
Copyright © 1999-2001, Exponential Consulting, LLC.

92

95

Kernel

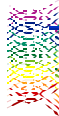



Administering Linux in Production Environments
Copyright © 1999-2001, Exponential Consulting, LLC.

92.1

96

Kernel Support



- `make xconfig=>`Block Devices
 - Multiple devices driver support
 - Autodetect RAID partitions
 - RAID levels
 - ❖ Linear
 - ❖ RAID-0
 - ❖ RAID-1
 - ❖ RAID-4/RAID-5

Administering Linux in Production Environments
Copyright © 1999-2001, Exponential Consulting, LLC.

93

97

Sample /etc/raidtab Entries

```
raiddev /dev/md0
raid-level 0
nr-raid-disks 2
chunk-size 64
persistent-superblock 1

device /dev/sdc1
raid-disk 0

device /dev/sdb1
raid-disk 1
```

```
raiddev /dev/md0
raid-level 1
nr-raid-disks 2
persistent-superblock 1

device /dev/sdc1
raid-disk 0

device /dev/sdb1
raid-disk 1
```

Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC.

94

98

RAID 0+1

```
raiddev /dev/md0
raid-level 0
nr-raid-disks 2
chunk-size 64
persistent-superblock 1
device /dev/sdc1
raid-disk 0
device /dev/sdd1
raid-disk 1

raiddev /dev/md2
raid-level 1
nr-raid-disks 2
persistent-superblock 1
device /dev/md0
raid-disk 0
device /dev/md1
raid-disk 1

raiddev /dev/md1
raid-level 0
nr-raid-disks 2
chunk-size 64
persistent-superblock 1
device /dev/sde1
raid-disk 0
device /dev/sdf1
raid-disk 1
```

Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC.

95

99

RAID 5 with a Hot Spare

```
raiddev /dev/md0
raid-level 0
nr-raid-disks 5
persistent-superblock 1
device /dev/sdc1
raid-disk 0
device /dev/sdd1
raid-disk 1
device /dev/sde1
raid-disk 2
device /dev/sdf1
raid-disk 3
device /dev/sdg1
raid-disk 4
device /dev/sdh1
space-disk 0
```

Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC.

96

100

General RAID Considerations

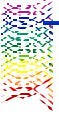
- RAID 0 stripe size matters for performance
 - Best value high dependant on typical I/O transfer size
 - ❖ Defaults are poor for large I/O operations
 - No substitute for testing (trial and error)
- Underlying filesystem block size = 4KB
 - **mke2fs -b 4 ...**
- Don't overload controllers
- Spend the money if you have it
 - RAID 5 overhead ~23% !!

Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC.

97

101



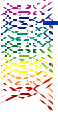
Parallel Processing and Clustering

Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC

98

102



Configuring Compute Servers

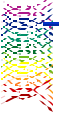
- Parallel program execution
 - SMP: shared memory
 - Distributed parallel: Beowulf and others
 - Simulates shared memory for discrete systems
 - Can be combined
- Clusters
 - High availability

Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC

99

103



Beowulf

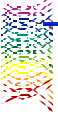
- www.beowulf.org
- www.extreme-linux.org
- An idea, obsession, religion

Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC


100

104



Two Extremes

- Stone Soupercomputer (Oak Ridge): stonesoup.esd.ornl.gov



Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC

101

105

The Other End of the Spectrum



- The Hive (NASA goddard):
newton.gsfc.nasa.gov/thehive/



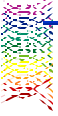
Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC.

101'

106

Creating a Beowulf System



- Hardware/software installation
- Interconnect
 - Ethernet or better
- Kernel changes
 - Generally integrated into kernel source tree
 - Channel bonding
- Parallel computing environment
 - Administrative setup
- Modified (or parallel-ready) applications

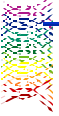
Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC.

102'

107

Appropriate Hardware



- Memory
- I/O bandwidth
 - Disk
 - Network
- CPU
- Disk
 - Partitioning

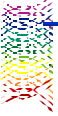
Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC.

102.2

108

System Setup



- System configuration (e.g., DHCP)
 - Static addressing
- Distributed filesystem choice
- Physical labeling:
 - Node name
 - Configuration
 - MAC Address

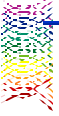
Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC.

102.3

109

Interconnection



- Cluster/farm vs. to outside world
 - Switches
 - Security on “world” node
- Channel bonding
- Topology

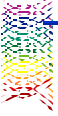
Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC.

102.4

110

Configuration



- Setup vs. ongoing
- Automounting
- Distributing configuration files
- Interprocess communication
 - rsh
 - ssh

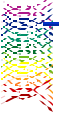
Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC.

102.5

111

Evaluating Pre-Packaged Systems



- Do the math to compute *price-performance*
 - Whole system price vs. sum of parts
 - ❖ Discrete systems from commodity vendors
 - ❖ NICs
 - ❖ Switch or hub
 - ❖ Cabeling
 - ❖ Time?
 - Weight price by processor speed
 - ❖ Pre-packaged systems often use older processors

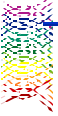
Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC.

104

112

Example



- Box with 4 Compaq EV5s plus proprietary interconnect for \$40K
- **versus**
- Compaq DS20 (2 EV6s with shared memory), listing at around \$35K

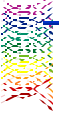
Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC.

105

113

Other Considerations



- Power
- Air conditioning
- Space

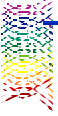
Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC.

105.1

114

SMP



- Enable in kernel:
 - 2.2.x
 - Symmetric multi-processing support: yes
 - Memory type range register (MTRR): yes
 - RTC support: yes
 - Advanced power management: no
- Parallel application
 - OpenMP
 - ❖ Compiler that supports it

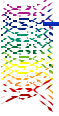
Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC.

106

115

Clustering



- Not designed for single application/job performance
- Purpose is to combine multiple systems to be presented as a single computing resource to users:
 - Linux Virtual Server (LVS)
 - High availability
 - Load balancing
 - Shared storage

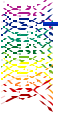
Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC.

109

116

Linux Virtual Server (LVS)



- Multiple servers appear as a single system to users (one IP address)
- Some load balancing
 - Implemented via a designated server
 - Fairly simple algorithms
- www.LinuxVirtualServer.org

Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC.

110

117

LVS

The diagram illustrates the LVS (Linux Virtual Server) architecture. It shows a central cloud labeled 'Internet/Intranet' connected to a 'Load Balancer' (LVS) and a 'Server Cluster'. The Server Cluster consists of multiple servers, each with its own IP address. The LVS is responsible for distributing traffic across the Server Cluster. The IP Cluster is also shown, representing the virtual IP addresses used for the cluster. The diagram is titled 'High Availability of Linux Virtual Server'.

Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC

110.1

High Availability of Linux Virtual Server

118

High Availability Linux

- High-Availability Linux Project: linux-ha.org
 - Compare to fault tolerance: quick recovery vs. never failing
 - Multiple servers
 - Redundant communications channels
 - Shared disks (or distributed filesystem)
 - Resource groups: everything needed for some service/application to work

Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC

111

119

Current Components

- **heartbeat**: detects failed servers
- **mon**: service monitoring daemon
- **fake**: provides IP-address takeover
- Architecture for a more elaborate facility

Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC

112

120

Commercial Products

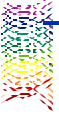
- TurboLinux Cluster Server
 - \$995 for 2 nodes; \$1995 for unlimited
 - www.turbolinux.com
- SGI/SuSE: FailSafe (*in progress*)
- Legato Cluster
- LifeKeeper (SteelEye)
- ...

Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC

113

121



Piranha Project (Red Hat)

- Free!
- ≡LVS

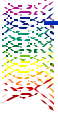
Administering Linux in Production Environments

- Control daemon
- GUI administration tool
- Failover vs. Clustering

Copyright © 1999-2001, Exponential Consulting, LLC

114' ▪ sources.redhat.com/piranha

122



Piranha Options

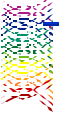
- Routing:
 - Network Address Translation (NAT)
 - Direct
 - IP Tunneling
- Scheduling
 - Round robin
 - Fewest connections
 - With/without weighting (assigned, load averages)

Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC

115

123



Typical Piranha Setup

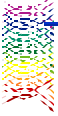
- One routing node
 - Supports one or more virtual servers:
 - ✦ IP address, protocol, port triple
 - Visible from the “external” world
 - Connected to a private network holding the real servers
- Multiple servers
 - Do real work
 - Static data only!

Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC

116

124



Piranha Components

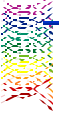
- IPVS and other kernel code
- **lvs** daemon: manages routing table
- **nanny** daemon: monitors a server, updates routing info as appropriate
- **pulse** daemon: handles failovers
- **piranha**: GUI configuration/management tool
- /etc/lvs.cf: configuration file

Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC

117

125



Enterprise Networking Features

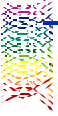
Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC.

118

126

Standard Services Review



- DNS
 - Dynamic
 - ❖ www.technopagan.org/dynamic/ddns-primer.html
 - Load balancing
 - ❖ www.cs.twsu.edu/~hcvillia/acads/project
- DHCP
- Automounting
 - **autofs not amd**
 - ❖ No symbolic links!

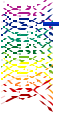
Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC.

119

127

High Speed Networking



- 100 mbps Ethernet: all many sites need
 - All mass-market chipsets supported
 - Issues with dual speed switches
 - ❖ Autonegotiation works best if both ends have it
- Gigabit Ethernet: common chipsets supported
 - www.beowulf.org/linux/drivers/

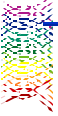
Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC.

120

128

MLPPP



- Multilink PPP: used to combine two or more PPP links into one, higher bandwidth virtual connection
 - Any connections can be used: modems, ISDN (WANs)
 - Linux implementation is basic: www.linux-mp.terz.de
- Configuration
 - Kernel patches
 - Updated **pppd**
 - Start multiple lines:


```
pppd /dev/ttyS0 multilink
wait for ppp0 interface to be up as usual
pppd /dev/ttyS1 multilink mp-join ppp0 mp-nonp noip
wait for ppp0 interface to be up as usual
pppd /dev/ttyS2 multilink mp-join ppp0 mp-nonp noip
```

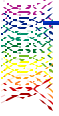
Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC.

121

129

Virtual Private Networks



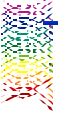
- VPNs use the public Internet as the communications link between sites
 - Data is tunneled in encrypted form
 - Protocols:
 - ❖ PPTP: many security problems
 - ❖ IPSec: emerging standard
 - ❖ ssh: See David Sifry, "Creating VPNs with Linux," *Linux Magazine*, Spring 1999.
 - Virtual Private Server (VPS): Linuxcare
 - www.strongcrypto.com
- VPN is combined with masquerading for use behind firewalls

Administering Linux in Production Environments
Copyright © 1999-2001, Exponential Consulting, LLC

122

130

TCP Wrappers



- Adds (primarily) host-level access control to **inetd**-based network services
- **tcpd** replaces service executable in `/etc/inetd.conf`:

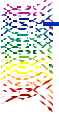

```
#service socket prot wait? user program args
tcpd    stream  tcp    nowait  root    /usr/sbin/tcpd  :inetd
```
- `/etc/hosts.allow` and `.deny` control access
 - allow: `tcpd : LOCAL, mycomp.com, 192.100.43`
 - deny: `ALL : ALL`
- Logs to **syslog** facility

Administering Linux in Production Environments
Copyright © 1999-2001, Exponential Consulting, LLC

123

131

Tripwire



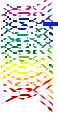
- www.tripwiresecurity.com
 - Originated with COAST/CERIAS
 - Free and commercial versions for Linux
 - "2.0 adds many enhancements"
- What it does:
 - Documents a known system state
 - ❖ Multiple cryptographic signatures for every item
 - Many algorithm choices
 - Compares current state to the stored state

Administering Linux in Production Environments
Copyright © 1999-2001, Exponential Consulting, LLC

124

132

Setting Up and Using Tripwire



- Perform immediate after a clean installation/upgrade
 - Database security (read-only or cryptographic signatures as well)
- Configure system, specifying files to check and ignore
 - Features to handle log files and the like
- Set up and enable automatic periodic monitoring
- Look and act upon the reports

Administering Linux in Production Environments
Copyright © 1999-2001, Exponential Consulting, LLC

125

133

Other Monitoring Tools

- **saint**: www.wwdsi.com/saint
- **nmap**: www.insecure.com/nmap

Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC.

126

134

UPS Configuration

- Connect UPS device to serial port
 - Non-shared IRQ!
 - Cable with 10 Ohm resistor (from manufacturer)
- Run powerd (or other daemon)
 - Master: `powerd /dev/ttyS0 -port n`
 - Client: `powerd -host dalton n`
- Add inittab entries for power-related events:
 - powerfail
 - powerfailnow
 - powerok

Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC.

129

135

File Servers

- ★ _____
- ★ _____
- ★ _____
- ★ _____

Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC.

130

136

Stronger File Permissions: POSIX Access Control Lists

- More fine-detailed control of file access
 - Specifiable on a per-user/per-group basis
 - Default ACLs flow from the directory location
 - acl.bestbits.at
- Limits total number & size
- ACLs and NFS/Samba

Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC.

131

Example

```
u::rwx
g::rwx
o:---
u:chavez:rwx
g:chem:r-x
m:r-x
```

- Entries are applied as follows:
 - User owner uses u::
 - Specified users use u:'s
 - Group members use all applying g:'s (incl. g:: if applicable)
 - ◊ Not accumulated!
 - Everyone else uses o:
 - Mask (m:) sets maximum access level except for o access

Copyright © 1999-2001, Exponential Consulting, LLC.

132

Enabling ACLs

- Patch and build kernel
 - Set development option
- Get/build utilities
- Patch and rebuild ext2fs utilities
- **mount ... -o acl**
- Use **setfacl** and **getfacl** to set permissions
- Do frequent backups

Copyright © 1999-2001, Exponential Consulting, LLC.

133

Adding ACL Capability

*	Linux 2.4.13 or (x86-acpi)	Get the kernel source tree from http://www.kernel.org or from a mirror, and the other Checkpoint patch from http://www.kernel.org or from a mirror. Then get these patches: kernel/patch kernel/patch	0.1.23 (x86) 0.1.23 (x86)
*	Linux 2.4.14 (x86-mpc)	Get the kernel source tree from http://www.kernel.org or from a mirror. Then get these patches: kernel/patch kernel/patch	0.1.24 0.1.24
*	Linux 2.2.28 (x86-mpc)	Get the kernel source tree from http://www.kernel.org or from a mirror. Then get these patches: kernel/patch kernel/patch	0.1.23 0.1.23
*	x86_64 1.5	Get the package (kernel 1.23 or kernel) from http://www.kernel.org Needed for manipulating access control lists	0.1.23 0.1.23
*	FreeBSD 4.1	Get the base package from http://www.freebsd.org or from a mirror, and apply the patch: kernel/patch kernel/patch	0.1.23 0.1.23
*	FreeBSD 4.1.1 (alpha)	Get the base package from http://www.freebsd.org and apply the patch: kernel/patch kernel/patch	0.1.23 0.1.23
*	Extended Attribute utilities	Check for existing extended attributes	0.1.25
*	Other stuff	kernel/patch kernel/patch kernel/patch	

Copyright © 1999-2001, Exponential Consulting, LLC.

134

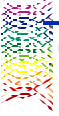
Backing Up ACLs

- **ACLs:**
 - **getfacl -R --skip-base / > backup.acl**
 - **setfacl --restore=backup.acl**
 - **aget -sdR -e base64 / > backup.ea**
 - **aset -B backup.ea**
- **Files with ACLs:**
 - **star H=exustar ... > file.tar**
 - **star -p -x < file.tar**
 - **ftp.fokus.gmd.de/pub/unix/star**

Copyright © 1999-2001, Exponential Consulting, LLC.

134.1

141



Encryption

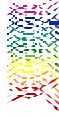
- **Steganography**
 - Filesystem hidden in the low bits of each byte of an audio file
 - ban.joh.cam.ac.uk/~adm36/StegFS/
- **Encrypted filesystems:**
 - ❖ Loop Device mechanism: EncryptionHOWTO.sourceforge.net/
 - ❖ PPDD: linux01.gwdg.de/~alatham/ppdd.html
 - ❖ CFS: <http://drt.ailis.de/crypto/linux-disk.html>
 - ❖ TCFS: tcfs.dia.unisa.it

Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC.

135

142



General Setup Process


- Patch/configure/build kernel
- Get/build utilities
- Patch/rebuild standard FS tools

Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC.

136

143



Print Servers

- ★ _____
- ★ _____
- ★ _____
- ★ _____

Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC.

137

144



LPRng

- Enhanced BSD lpd
- Better networking support
 - Smarter clients
- www.lprng.com

Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC.

137.1

145

lpc

Subcommand	Purpose
<code>hold queue [job]</code>	Place the specified job or all jobs in the queue into a hold state, preventing them from printing.
<code>release queue [job]</code>	Alter the specified hold print (job) to print.
<code>holdall queue</code>	Place all new jobs entering the queue into the hold state. The <code>holdall</code> is designed to ensure the hold-all job will not be explicitly released.
<code>move old-queue job new-queue</code>	Transfer the specified print (job) between queues.
<code>release old-queue new-queue</code>	Release jobs specified to the old queue to the new queue. Specify <code>all</code> for the later to be all jobs.
<code>stop queue [job]</code>	Stop the specified job.
<code>kill queue</code>	Equivalent to <code>stop job</code> then <code>kill</code> the output job, and then reset the queue.
<code>enable printer[@host]</code>	Configure whether the specified speed device is active or not.
<code>reset printer[@host]</code>	Place the specified speed device in reset in configuration file.
<code>clean queue class [job]</code>	Limit printing from the specified queue to jobs in the specified class(es), where class is usually a comma-separated list of one or more class letters (see below). The keyword <code>all</code> removes any current class restrictions in effect.

Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC

137.2

146

Printcap Entries

```

hp:
:lp=/dev/lp0
:cm=HP Laser Jet printer
:lf=/var/log/lpd.log
:af=/var/adm/pacct
:filter=/usr/local/lib/filters/iftph
:tc=.common

laser:
:oh=10.0.0/24
:lp=printers@matisse
:tc=.common

common:
:sd=/var/spool/lpd/%P
:mx=0
    
```

Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC

137.3

147

LPRng Capabilities

- Bounce queues
- Printer pools
- Local and/or remote filtering
- Program-generated printcap file

Access control: `lpd.perms`

```

REJECT SERVICE=X NOT REMOTEHOST=*.ahania.com
REJECT SERVICE=X REMOTEHOST=dalton,hamlet
ACCEPT SERVICE=C SERVER REMOTEGROUP=printop
LPC=top,q,hold,release
ACCEPT SERVICE=R,M,C REMOTEUSER=chavezPRINTER=test
REJECT SERVICE=* PRINTER=test
    
```

Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC

137.4

148

CUPS

- Common UNIX Printing System
- www.cups.org
- Network-based printing
- Separates job processing and device spooling
- Compatible commands

Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC

137.5

149

cupsd.conf

ServerName painters.ahania.com
 ServerAdmin root@ahania.com
 ErrorLog /var/log/cups/error_log
 AccessLog /var/log/cups/access_log
 PageLog /var/log/cups/page_log
 LogLevel info
 MaxLogSize 1048571
 PreserveJobFiles No
 RequestRoot /var/spool/cups
 User lp
 Group sys
 TempDir /var/spool/cups/tmp
 MaxClients 100
 Timeout 300
 Browsing On
 ImplicitClasses On

*Server name.
CUPS admin's email address.
Log file locations.*

*Printer accounting data.
Log detail (debug, warn, error).
Rotate log files when current > this.
Don't keep files after job completes.
Spool directory.
Server user and group owners.*

*CUPS temporary directory.
Max. client connections to server.
Printing timeout period in seconds.
Let clients browse for printers.
Implicit classes are enabled.*

Administering Linux in Production Environments
Copyright © 1999-2001, Exponential Consulting, LLC

137.6

150

Security

- Access control
- Authentication
- Encryption

<Location /classes/checks>
 Encryption Always
 AuthType Digest
 AuthClass Group
 AuthGroupName finance
 Order Deny,Allow
 Deny From All
 Allow From 10.100.67.0/24
 </Location>

*Applies to class named checks.
Always encrypt.
Require valid user account and password.
Restrict to members of the finance group.*

*Deny all access...
Except from this subnet.*

Administering Linux in Production Environments
Copyright © 1999-2001, Exponential Consulting, LLC

137.7

151

Sharing Printers via Samba

[laser4]
 printable = yes
 comment = LaserWriter on dalton
 public = yes
 postscript = yes
 printer name = laz4
 printer driver = *Microsoft Driver name*

[global]
 load printers = yes
 printcap name = /etc/printcap.samba
 printing = bsd | sysv | aix | hpux

[printers]
 writeable = no
 path = /tmp
 auto services = laz4 laz5 monet

- Note that this makes *all* printers in the specified file available!

Administering Linux in Production Environments
Copyright © 1999-2001, Exponential Consulting, LLC

138

152

Printing to a Windows Printer

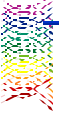
- Standard remote system printcap entry:
picasso:lp=:m=:vala:rp=picasso:
- Using smbprint:
gaughin:sd=dir:lp=/dev/null:if=/usr/sbin/smbprint:af=file:
 - Place a .config file in the specified spool directory containing password:
 server=zoas
 service=matisse
 password=pwd

Administering Linux in Production Environments
Copyright © 1999-2001, Exponential Consulting, LLC

139

153

Font Management Summary



- XFree86Config FontPath
- fonts.dir and fonts.scale files
 - **mkfontdir**
 - **type1inst**
- Font server: **xfst**
- Ghostscript Fontmap file

```
OctavianMT-Italic (oci____.pfb) ;
OctavianMT (ocr____.pfb) ;
OctavianMT-Roman /OctavianMT ;
```

Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC

139.1

154

Enterprise User Authentication



- ★ _____
- ★ _____
- ★ _____
- ★ _____


Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC

140

155

User Authentication



- UNIX standard mechanisms
 - Data files: passwd and group
 - Shadow password file
 - ✦ MD5
- PAM

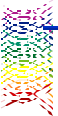
Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC

141

156

PAM



- Pluggable Authentication Modules
 - Linux and Solaris 7
- Per-application (service) configuration files in /etc/pam.d
 - Actual authentication performed by modules: shared libraries (.so files)

Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC

142

157

PAM Configuration Files

- Entry format: *type severity path args*
 - *type*: purpose
 - ❖ *auth* – user authentication
 - ❖ *account* – account attributes/controls
 - ❖ *session* – pre/post service activities (logging to syslog)
 - ❖ *password* – causes password change if applicable
 - *severity*: how results affect outcome
 - ❖ *required*: failure => access denied
 - ❖ *requisite*: immediate required
 - ❖ *sufficient*: success => access granted immediately
 - ❖ *optional*: result used only if nothing else is deterministic
 - *path args*: path to module and arguments to it

Administering Linux in Production Environments
Copyright © 1999-2001, Exponential Consulting, LLC

143

158

Example: rlogin

```
##PAM-1.0
auth      required /lib/security/pam_securetty.so
auth      sufficient /lib/security/pam_rhosts_auth.so
auth      required /lib/security/pam_pwdb.so shadow nullok
auth      required /lib/security/pam_nologin.so
account   required /lib/security/pam_pwdb.so
account   required /lib/security/pam_time.so
password  required /lib/security/pam_cracklib.so
password  required /lib/security/pam_pwdb.so
           nullok use_authok md5 shadow
session   required /lib/security/pam_pwdb.so
```

Will this work??

Administering Linux in Production Environments
Copyright © 1999-2001, Exponential Consulting, LLC

144

159

Example: su

```
##PAM-1.0
auth      required /lib/security/pam_rootok.so
auth      required /lib/security/pam_wheel.so groupadmins
auth      required /lib/security/pam_pwdb.so shadow md5 nullok
account   required /lib/security/pam_pwdb.so
password  required /lib/security/pam_pwdb.so
session   required /lib/security/pam_pwdb.so
```

Administering Linux in Production Environments
Copyright © 1999-2001, Exponential Consulting, LLC

145

160

PAM Modules

- *pam_deny* (**account, auth, passwd, session**)
- *pam_permit* (**account, auth, passwd, session**)
 - Deny/allow all access by always returning failure/success (respectively). These modules do not log, so stack them with *pam_warn* to log the events.
- *pam_warn* (**account, auth, passwd, session**)
 - Log information about the calling user and host to syslog.
- *pam_access* (**account**)
 - Specify system access based on user account and originating host/domain as in the widely-used logdaemon facility. Its configuration file is */etc/security/access.conf*.

Administering Linux in Production Environments
Copyright © 1999-2001, Exponential Consulting, LLC

146

161

More...

- ***pam_pwdb* (account, auth, passwd, session)**
- ***pam_unix* (account, auth, passwd, session)**
 - Two modules for verifying and changing user passwords. When used in the **auth** stack, the modules check the entered user password. When used as an **account** module, they determine whether a password change is required or not (based on password aging settings in the shadow password file); if so, they delay access to the system until the password has been changed.
 - When used as a **password** component, the modules update the user password. In this context, the **shadow** (use the shadow password file) and **use_authok** options are useful; the latter forces the modules to set the new password to one provided by a previous module in the stack and should accordingly be set when a password checking module is used.

Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC.

147

162

And More...

- ***pam_cracklib* (passwd)**
 - Password triviality checking. Needs to be stacked with *pam_pwdb* or *pam_unix*. See the separate discussion below.
- ***pam_pwcheck* (passwd)**
 - Another password checking module, checking that the proposed password conforms to the settings specified in */etc/login.defs* (discussed previously in this chapter).
- ***pam_env* (auth)**
 - Set/unset environment variables with a PAM stack. It uses the configuration file */etc/security/pam_env.conf*.

Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC.

148

163

Still More...

- ***pam_issue* (auth)**
- ***pam_motd* (session)**
 - Display an issue or message of the day file at login. The issue file is displayed before the username prompt, and the message of the day file is displayed at the end of a successful login process.
- ***pam_krb4* (auth, passwd, session)**
 - Interface to Kerberos user authentication.
- ***pam_lastlog* (auth)**
 - Adds an entry to the */var/log/lastlog* file which contains data about each user login session.

Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC.

149

164

More Again...

- ***pam_limits* (session)**
 - Sets user process resource limits (*root* is not affected), as specified in its configuration file, */etc/security/limits.conf*. This file contains entries of the form:


```
name hard/soft resource limit-value
```

where *name* is a user or group name or an asterisk (indicating the default entry). The second field indicates whether it is a soft limit, which the user can increase if desired, or a hard limit (the upper bound which the user cannot exceed). The final two fields specify the resource in question and the limit assigned to it.

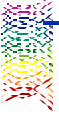
Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC.

150

165

More and More...



- **pam_listfile (auth)**
 - Deny/allow access based on a list of usernames in an external file:

```
auth required pam_listfile.so onerr=fail sense=allow \
file=/etc/ftpusers item=user
```

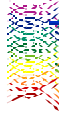
- **pam_mail (auth, session)**
 - Displays a message indicating whether the user has mail. The default mail file location (*/var/spool/mail*) can be changed with its **dir** argument.
- **pam_mkhomedir (session)**
 - Creates the user's home directory if it does not already exist, copying files from */etc/skel* to the new directory.

Administering Linux in Production Environments
Copyright © 1999-2001, Exponential Consulting, LLC.

151

166

Some More...




- **pam_nologin (auth)**
 - Prevents non-*root* logins if the file */etc/nologin* exists, the contents of which are displayed to the user.
- **pam_rhosts (auth)**
 - Performs traditional */etc/rhosts* and *~/.rhosts* password-free authentication for **rsh** and **rlogin** sessions between networked hosts (see chapter 8).
- **pam_rootok (auth)**
 - Allow *root* access without a password.
- **pam_securetty (auth)**
 - Prevents *root* access unless the current terminal line is listed in the file */etc/securetty*.

Administering Linux in Production Environments
Copyright © 1999-2001, Exponential Consulting, LLC.

152

167

...The Final Two



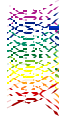
- **pam_time (account)**
 - Restrict access by time of day by user, group, tty and/or shell.
- **pam_wheel (auth)**
 - Designed for the **su** facility, this module prevents root access to any user who is not a member of a specified group (**group=name** option), which defaults to GID 0. You can reverse the logic of the test to deny *root* access to members of a specific group by using the **deny** option along with **group**.

Administering Linux in Production Environments
Copyright © 1999-2001, Exponential Consulting, LLC.

153

168

More on PAM



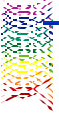
- The other service:


```
auth required /lib/security/pam_warn.so
auth required /lib/security/pam_deny.so
...
```
- Module configuration files stored in */etc/security*
 - Example: *time.conf* specifies hours when users may access defined PAM services

Administering Linux in Production Environments
Copyright © 1999-2001, Exponential Consulting, LLC.

154

169



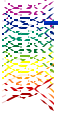
PAM Futures

- More modules
 - Debugged modules
- New syntax for *severity*:
 - return-val=action* [, *return-val=action* [...]]
 - Example: `success=ok,open-err=ignore,authtok_a`

Administering Linux in Production Environments
Copyright © 1999-2001, Exponential Consulting, LLC.

155

170



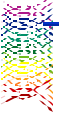
SmartCards

- Hardware token
 - CryptoCard: www.cryptocard.com
- Card plus reader
 - MUSCLE: Mvmt. for use of smart cards in Linux env.
 - ❖ www.linuxnet.com/smartcard/tutorial.html
 - Schlumberger cards
 - Some with biometric authentication integrated into the reader
 - Radius standard: PAM module

Administering Linux in Production Environments
Copyright © 1999-2001, Exponential Consulting, LLC.

156

171



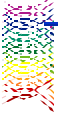
Networking and Authentication

- Beyond the single system
 - rdist, rsynch
 - NIS
 - LDAP

Administering Linux in Production Environments
Copyright © 1999-2001, Exponential Consulting, LLC.

157

172



LDAP

- Lightweight Directory Access Protocol
- Protocol for accessing general directory services (namespace)
 - DAP, X.500
 - Global structure is defined but little used
- Software: www.openldap.org
- Migration tools: www.padl.com/tools.html
 - ❖ Also PAM modules

Administering Linux in Production Environments
Copyright © 1999-2001, Exponential Consulting, LLC.

158

173

Basics

- Directory not database
 - Entries (objects) and attributes
- Distinguished name:
 - cn=Aeleen Frisch, ou=People, dc= ahania, dc=com
 - uid=chavez, o= SomeCo, c=US
- RDN

Administering Linux in Production Environments
Copyright © 1999-2001, Exponential Consulting, LLC

158.1

174

Example User Data

```
dn: uid=chavez,o= SomeCo ,c=US
uid: chavez
cn: Rachel Chavez
objectClass: posixAccount
objectClass: account
gecos: Rachel Chavez
userpassword (crypt) xxxxxxxxxxxx
loginShell: /bin/tcsh
uidNumber: 278
gidNumber: 250
homeDirectory /home/chavez
```

- More object classes
- More attributes
 - First name
 - Last name
 - Email address
 - Kerberos data
 - Phone number
 - Picture
 - Whatever

Administering Linux in Production Environments
Copyright © 1999-2001, Exponential Consulting, LLC

159

LDIF format

175

Deployment Process

- Design
- LDAP server:
 - Open LDAP **slapd** (LDAP 3)
 - Netscape Directory Server (LDAP 3)
 - ♦ Replication servers
- Configure
 - /etc/openldap/{slapd,ldap}.conf
- Migrate existing data into LDAP databases

Administering Linux in Production Environments
Copyright © 1999-2001, Exponential Consulting, LLC

160

176

Linux Software

- DB manager: GNU **gdbm** (www.fsf.org) or BerkeleyDB (www.sleepycat.com).
- Transport Layer Security (TLS/SSL) libraries (www.openssl.org).
- The Cyrus SASL libraries (asg.web.cmu.edu/sasl).

```
# rpm -qa | egrep -i '(db|sasl|ssl)'
db-3.1.17-13
gdbm-1.8.0-225
sdb-2001.1.18-0
sdb_en-2001.1.18-0
cyrus-sasl-1.5.24-5
openssl-0.9.6-21
openssl-devel-0.9.6-21
```

Administering Linux in Production Environments
Copyright © 1999-2001, Exponential Consulting, LLC

160.1

177

slapd.conf

```
# /etc/openldap/slapd.conf
include      /etc/openldap/schema/core.schema
pidfile     /var/run/slapd.pid
argsfile    /var/run/slapd.args

database    ldbm
suffix      "dc=ahania,dc=com"
rootdn      "cn=Manager,dc=ahania,dc=com"
# encode with slapdpasswd -h '{MD5}' -s secret -v -u
rootpw      {MD5}Xr41lOzQ4PCOq3aQ0qbuaQ==
directory   /var/lib/ldap
```

Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC

160.2

178

Starting the Server

- **/etc/init.d/ldap start**
- **Enabling configuration:**
 - **E.g.: SuSE /etc/rc.config:**

```
TART_LDAP="yes"
```
- **Populate database/migrate data**
 - **PADL Perl scripts**

Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC

160.3

179

Clients

- **gq, web2ldap, kldap, ...**
- **Netscape**

Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC

161

180

Client Configuration File

- **/etc/ldap.conf or /etc/openldap/ldap.conf**

```
# /etc/openldap/ldap.conf
URI ldap://bella.ahania.com
BASE dc=ahania,dc=com
```

Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC

161.1

181

System Level Clients

- LDAP-aware applications:
 - **sendmail** and **Postfix**
 - **Apache**
- Using for authentication:
 - **pam_ldap**
 - **nss_ldap**

Administering Linux in Production Environments
Copyright © 1999-2001, Exponential Consulting, LLC

161.2

182

Schema for Authentication

- /etc/openldap/schema/*.schema

objectClass	Schema file
top	cppe
person	cppe
organizationalPerson	cppe
inetOrgPerson	cppe
account	cpbine
posixAccount	ns
shadowAccount	ns
inetLocalMailRecipient	inet*mailrecipient

Administering Linux in Production Environments
Copyright © 1999-2001, Exponential Consulting, LLC

161.3

183

nss_ldap

- `/etc/nsswitch.conf`

```
nss_base_passwd    ou=People,dc=ahania,dc=com
nss_base_shadow    ou=People,dc=ahania,dc=com
nss_base_group     ou=Group,dc=ahania,dc=com
```
- `/etc/hsswitch.conf`

```
passwd: files ldap
shadow: files ldap
```

Administering Linux in Production Environments
Copyright © 1999-2001, Exponential Consulting, LLC

161.4

184

pam_ldap

- `/etc/pam.d/rlogin`

```
##PAM-1.0
auth    required    /lib/security/pam_securetty.so
auth    required    /lib/security/pam_nologin.so
auth    sufficient  /lib/security/pam_rhosts_auth.so
auth    sufficient  /lib/security/pam_ldap.so
auth    required    /lib/security/pam_unix.so
auth    required    /lib/security/pam_mail.so
account sufficient  /lib/security/pam_ldap.so
account required   /lib/security/pam_unix.so
password sufficient /lib/security/pam_ldap.so
password required  /lib/security/pam_unix.so strict=false
session required   /lib/security/pam_unix.so debug
```

Administering Linux in Production Environments
Copyright © 1999-2001, Exponential Consulting, LLC

161.5

185

Limiting Access by Host and/or Group

- ldap.conf**

```
# Allow user access to hosts      # Only this group can access this host
pam_check_host_attr yes          pam_groupdn cn=dalton.ahania.com,...
                                pam_member_attribute uniqueMember
```
- Directory entries:**

<pre>List of allowed hosts dn: uid=aefriach,ou=People,... objectClass: account objectClass: posixAccount .. host: milton.ahania.com host: shelley.ahania.com host: yeats.ahania.com</pre>	<pre># List of allowed users on the local host dn: cn=dalton.ahania.com,... objectClass: ipHost objectClass: device objectClass: groupOfUniqueNames cn: dalton cn: dalton.ahania.com uniqueMember: uid=chavez,ou=People,dc=... uniqueMember: uid=carter,ou=People,dc=...</pre>
---	--

Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC

161.6

186

Security: Access Control

- slapd.conf**

```
# simple access control: read-only except passwords
access to dn="*.*,dc=ahania,dc=com" attr=userPassword
  by self write
  by dn=root,ou=People,dc=ahania,dc=com write
  by * auth
access to dn="*.*,dc=ahania,dc=com"
  by self write
  by * read
```

Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC

161.7

187

Security: TLS₃₈₉ and SSL₆₃₆

- Key generation:**

```
cd /usr/ssl/cert
openssl req -newkey rsa:1024 -keyout slapd_key.pem \
-x509 -days 365 -out slapd_cert.pem
openssl rsa -in slapd_key.pem -out slapd_key.pem
chown ldap:ldap *.pem (if appropriate)
chmod 600 sl*.pem
```
- Startup:**

```
slapd -h "ldap:// ldap://"
```
- slapd.conf**

```
TLS_CIPHER_SUITE HIGH:MEDIUM:+SSLv2
TLS_CERTIFICATE_FILE /usr/ssl/certs/slapd_cert.pem
TLS_CERTIFICATE_KEY_FILE /usr/ssl/certs/slapd_key.pem
```

Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC

161.8

188

Security: Other

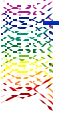
- Kerberos**
- SASL**

Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC

161.9

189



Compute Servers

Administering Linux in Production Environments

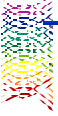
Copyright © 1999-2001, Exponential Consulting, LLC.

- ★ _____
- ★ _____
- ★ _____
- ★ _____

162

190

Helpful Installation and Configuration Tools



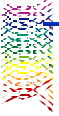
- **KickStart** (RedHat)
 - Automated installation/reinstallation
 - Ghost, Drive Image
- **cfengine** (www.iu.hioslo.no/cfengine)
 - System monitors and corrects itself by comparing to a stored “healthy” state

Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC.

171'

191



(Some) cfengine Capabilities

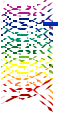
Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC.

- Check and configure the network interface.
- Edit text files for the system and for all users.
- Make and maintain symbolic links, including multiple links from a single command.
- Check and set the permissions and ownership of files.
- Delete junk files which clutter the system.
- Systematic, automated mounting of NFS filesystems .
- Checking for the presence of important files and filesystems .
- Controlled execution of user scripts and shell commands.
- Process management.

171.1

192



Pieces (1.6)

Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC.

- **cfengine**
- **cfid**
- **cfrun**
- **cfengine.conf**
- **cfid.conf**

171.2

193

A Simple cfengine.conf

```
control:
actionsequence = ( tidy links )
access = ( chavez root )

links:
/bin -> /usr/bin

tidy:
/tmp pattern=* age=7 recurse=inf
```

Administering Linux in Production Environments
Copyright © 1999-2001, Exponential Consulting, LLC.

171.3

194

Actions

- mountables, mountinfo, mountall, umount, addmounts
- directories, files, links, required
- copy, tidy, editfiles, disable
- shellcommands, processes
- netconfig, broadcast, resolve, defaultroute
- checktimezone
- mailcheck
- module

Administering Linux in Production Environments
Copyright © 1999-2001, Exponential Consulting, LLC.

171.4

195

Examples 1

```
copy:
$(masteretc)/hosts.denydest=/etc/hosts.deny o=root mode=0644
$(masteretc)/ntp.drift dest=/etc/ntp.drift mode=644
$(masteretc)/shells dest=/etc/shells mode=644

linux:
$(masteretc)/rc.config dest=/etc/rc.config o=root mode=644

dbdvest.Hr03::
/dev dest=/backup/dev server=dalton r=1 backup=false
```

Administering Linux in Production Environments
Copyright © 1999-2001, Exponential Consulting, LLC.

171.5

196

Examples 2

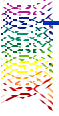
```
files:
/etc/security mode=600 owner=root group=sys recurse=inf
action=fixall
/usr/bin checksum=md5 action=warnall
/home recurse=inf include=*.txt action=compress
/private acl=secure1 action=fixall

acl:
{ secure1
method:overwrite
fstype:nt
user:chavez:rx:allowed
user:mark:all:allowed
user:toreo:read:allowed
group:dummy:all:denied
}
```

Administering Linux in Production Environments
Copyright © 1999-2001, Exponential Consulting, LLC.

171.6

197



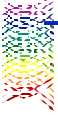
Classes

- OS
 - ultrix, sun4, sun3, hpux, hpux10, aix, solaris, osf, irix4, irix, irix64
 - freebsd, netbsd, openbsd, bsd4_3, newsos, solarisx86, aos,
 - nextstep, bsdos, linux, debian, cray, unix_sv, GnU, NT
- host
- host group name
- day of the week
- hour (Hr02)
- minute (Min33)
- 5 minute interval (Min00_05)
- day (Day1)
- month
- year (Yr2001)
- name

Administering Linux in Production Environments
Copyright © 1999-2001, Exponential Consulting, LLC.

171.7

198



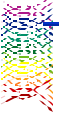
Class Examples

- myclass.solaris.Monday.Hr01::
- sun4|ultrix|osf::
- myhosts.aix!vader::
- December.Day31.Friday::
- any::

Administering Linux in Production Environments
Copyright © 1999-2001, Exponential Consulting, LLC.

171.8

199



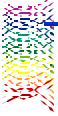

Distributing the Resources

- Batch processing
 - Multiple queues across many servers
 - Priorities
 - Resource limits
 - Access control
 - Authentication
 - Logging and accounting
- Load balancing

Administering Linux in Production Environments
Copyright © 1999-2001, Exponential Consulting, LLC.

162.1

200

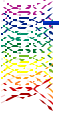
PBS

- Portable Batch Scheduler
 - Designed for maximal resource usage
 - www.pbspro.com
- Intelligence and control reside in the scheduler, not in the queue attributes
- Site-definable scheduling via configuration and/or hooks for routines
- Separate pre-staging vs. running vs. cleanup

Administering Linux in Production Environments
Copyright © 1999-2001, Exponential Consulting, LLC.

162.2

201



PBS Implementation

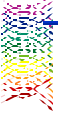
- Master server: **pbs_server**
- Per system: **pbs_mom**

Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC

162.3

202



Others

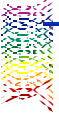
- Maui Scheduler
 - www.mhpcc.edu/maui
- IBM Load Leveler

Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC

162.4

203



Parallel Environments

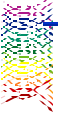
- PVM, MPI
- BSP
- DSM schemes
 - Linda

Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC

163

204



Parallelization Strategies

- Master/worker paradigm is dominant
 - Same code for both
- Deciding where to parallelize is the most important question
 - Discrete task -based strategies
 - ❖ When free, worker gets the next bit of work to do
 - Domain decomposition
 - ❖ Worker is preassigned a subset of the entire operation

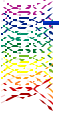
Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC

103

205

Computing π Numerically



```

get n as input
h = 1.0/double(n);
sum=0.0;
for (i=1; i<=n; i++) {
    sum += 4.0 / (1.0 + (h*(i - 0.5))**2); }
mypi = h*sum;
print the result
  
```

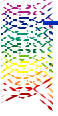
Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC.

164

206

Message Passing Interface (MPI)



- Processes explicitly send data to one another
 - MPI_Send and MPI_Receive
- www.erc.msstate.edu/labs/hpcl/projects/mpi

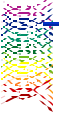
Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC.

165

207

MPI Example



- main routine:

```

MPI_Init(&argc, &argv);
MPI_Comm_size(MPI_COMM_WORLD, &numprocs);
MPI_Comm_rank(MPI_COMM_WORLD, &myid);
if (myid==0) {
    input parameter n
    MPI_Bcast(&n, 1, MPI_INT, 0, ...);
    h = 1.0/double(n); sum=0.0;
    for (i=myid+1; i<=n; i+=numprocs) {
        sum += 4.0 / (1.0 + (h*(i - 0.5))**2); }
    mypi = h*sum;
    MPI_Reduce(&mypi, &pi,
              MPI_DOUBLE, MPI_SUM, 0, ...);
    if (myid==0) { print the result }
    MPI_Finalize();
  }
  
```

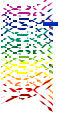
Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC.

166

208

BSP



- Bulk Synchronous Parallel Model
- www.bsp-worldwide.org
- Consists of a small number of operations designed to distribute data during a parallel calculation
 - Structured for asynchronous communication
 - Read/write data from remote process' memory without its participation
 - Send data to a remote process' queue

Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC.

167

209

Simulating Shared Memory

- Linda
 - www.lindaspaces.com
- Provides a virtual shared memory (VSM)
 - High level operations that can be ignorant of the specifics of communication

Administering Linux in Production Environments
Copyright © 1999-2001, Exponential Consulting, LLC

168

210

Master

- master function:


```
get n
out("pi-int", n);
out("pi-task", n);
sum=0.0
worker();
for (j=0; j<=n; j++)
  in("pi-sum", j, ?sum_j);
  sum += sum_j;
}
print sum
return;
```
- For faster performance, change `j` to `?k`
 - ❖ Why??

Administering Linux in Production Environments
Copyright © 1999-2001, Exponential Consulting, LLC

169

211

Worker

- worker function:


```
while (1) {
rd("pi-int", ?n);
in("pi-task", ?i);
if (i==0) { break; }
out("pi-task", i-1);
h = 1.0 / double(n);
sum = 4.0 / (1.0 + (h*(i-0.5))**2);
mypi = h * sum;
out = ("pi-sum", n, mypi);
}
return;
```

Administering Linux in Production Environments
Copyright © 1999-2001, Exponential Consulting, LLC

170

212

OpenMP

- www.openmp.org
- Standard set of compiler directives for shared-memory parallelism
 - Serial code is unmodified

Administering Linux in Production Environments
Copyright © 1999-2001, Exponential Consulting, LLC

107

213

OpenMP Example

```

!$omp parallel do
  do 50 k=1, ntop
    arr(k) = a(k) + b(k-1) * c(k) / (k+1)
  50 continue
!$omp end parallel do

  sum=0.0
!$omp parallel default(shared)
!$omp & reduction(+: sum)
  lots of computations
!$omp end parallel
    
```

Administering Linux in Production Environments
Copyright © 1999-2001, Exponential Consulting, LLC

108

214

Databases

★ _____

★ _____

★ _____

★ _____

Administering Linux in Production Environments
Copyright © 1999-2001, Exponential Consulting, LLC

172

215

Databases

- mySQL
- PostgreSQL
- Commercial servers: Oracle, Sybase, DB2, ...

Administering Linux in Production Environments
Copyright © 1999-2001, Exponential Consulting, LLC

173

216

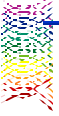
mySQL

- Good for **medium size applications**
 - If you need to save money ...
- Limitations include:
 - Limited SQL
 - ❖ No subqueries
 - ❖ `select sum(total) from pmts where trx_id in (select trx_id from trx_hdr.org where country>1 and paid=0 and ship>'06/01/1999' and join)`
 - ❖ Applies automatic default values on inserts
 - No transactions
- www.mysql.com

Administering Linux in Production Environments
Copyright © 1999-2001, Exponential Consulting, LLC

174

217



mySQL Tools

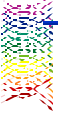
- **xmysql:** web.wt.net/~dblhack
- **xmysqladmin**
 - www.tcx.se/Contrib/
- **kmysql:** www.xnot.com/kmysql

Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC.

175

218



Database Programming Interfaces

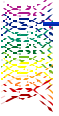
- Perl
 - MySQL
 - DBI and DBD drivers
 - ♦ DBD::ODBC
 - Win32:ODBC
- More in Tck/Tk, Python, ...

Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC.

176

219



Office PC

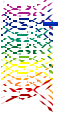
- ★ _____
- ★ _____
- ★ _____
- ★ _____

Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC.

177

220



Office Applications

- Applixware
- Star Office
 - Microsoft response
- Others
 - WordPerfect
 - Newcomers

Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC.

178

Word Processing



- Convert 1 page test file: **Ap4.41 Ap4.42 305.1**
- Import 10 page sample file: **Ap4.41 Ap4.42 305.1**
- Import 187 page manual chapter: **Ap4.41 Ap4.42 305.1**
- Paragraph styles: **Ap4.41 Ap4.42 305.1**
- Character styles: **Ap4.41 Ap4.42 305.1**
- Superscripts: **Ap4.41 Ap4.42 305.1**
- Footnotes: **Ap4.41 Ap4.42 305.1**
- Tables: **Ap4.41 Ap4.42 305.1**
- Sectioning: **Ap4.41 Ap4.42 305.1**
- Dynamic headers/footers: **Ap4.41 Ap4.42 305.1**
- TOC/Index: **Ap4.41 Ap4.42 305.1**
- Equation editor equations: **Ap4.41 Ap4.42 305.1**
- Manually-constructed equations: **Ap4.41 Ap4.42 305.1**
- Pictures: **Ap4.41 Ap4.42 305.1**

Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC

Word Processing



Word Processing Item	Apple i.Mac	Apple i.Mac	Star Office 5.1	Apple 5	Star Office 5.2
Convert 1 page test file	YES	YES	YES	YES	YES
Import 10 page sample file	YES	YES	YES	YES	YES
Import 187 page manual chapter	NO	YES	YES	YES	YES
Paragraph styles	YES	YES	YES	YES	YES
Character styles	NO	NO	YES	NO	YES
Superscripts	NO	NO	YES	NO	YES
Footnotes	YES	YES	YES	YES	YES
Tables	YES	YES	YES	YES	YES
Sectioning	NO	YES	YES	YES	YES
Dynamic headers/footers	NO	NO	NO	NO	NO
TOC/Index	NO	NO	NO	NO	NO
Equation editor equations	YES	YES	YES	YES	YES
Manually-constructed equations	NO	NO	YES	NO	YES
Pictures	NO	YES	YES	YES	YES
Acceptable font substitution	NO	YES	NO	YES	NO
True fonts (embedded in files)	NO	NO	NO	NO	YES
True fonts (externally updated)	NO	NO	NO	NO	YES
Text color	NO	NO	YES	NO	YES
Multiple columns	YES	YES	YES	YES	YES
Color and line draw	YES	YES	YES	YES	YES

Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC

Spreadsheets



- Convert spreadsheet files: **Ap4.41 Ap4.42 305.1**
- Convert multisheet workbooks: **Ap4.41 Ap4.42 305.1**
- Formulas: **Ap4.41 Ap4.42 305.1**
- Text color: **Ap4.41 Ap4.42 305.1**
- Text formatting: **Ap4.41 Ap4.42 305.1**
- Graphs (visible): **Ap4.41 Ap4.42 305.1**
- Graphs (correct): **Ap4.41 Ap4.42 305.1**

Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC

Spreadsheets



Spreadsheet Item	Apple i.Mac	Apple i.Mac	Star Office 5.1	Apple 5	Star Office 5.2
Convert spreadsheet files	YES	YES	YES	YES	YES
Convert multisheet workbooks	YES	YES	YES	NO	YES
Formulas	YES	YES	YES	YES	YES
Text color	YES	YES	YES	YES	YES
Text formatting	NO	YES	YES	YES	YES
Complex and grid formatting	NO	NO	YES	NO	YES
Print, display and sort range	NO	NO	NO	NO	YES
Print to print file	NO	YES	YES	YES	YES
Print to screen	NO	YES	NO	YES	YES

Administering Linux in Production Environments

Copyright © 1999-2001, Exponential Consulting, LLC

225

Presentations

- Convert current PPT files: **Ap4.41 Ap4.42 806.1**
- Convert prev. rev PPT files: **Ap4.41 Ap4.42 806.1**
- MS template: **Ap4.41 Ap4.42 806.1**
- Custom template: **Ap4.41 Ap4.42 806.1**
- Notes pages: **Ap4.41 Ap4.42 806.1**
- Drawn graphics: **Ap4.41 Ap4.42 806.1**
- Imported graphics: **Ap4.41 Ap4.42 806.1**

Administering Linux in Production Environments
Copyright © 1999-2001, Exponential Consulting, LLC

181

226

Presentations

Presentation Item	Ap4.4.1	Ap4.4.2	Star Office 3.1	Ap4.6.5	Star Office 3.2
Convert current PPT files	NO	YES	NO	NO	YES
Convert prev. rev PPT files	NO	YES	YES	NO	NO
MS template	NO	YES	YES	NO	YES
Custom template	NO	YES	NO	NO	YES
Notes pages	NO	YES	YES	NO	YES
Drawn graphics	NO	YES	YES	NO	YES
Imported graphics	NO	YES	YES	NO	YES
Obj charts	NO	YES	YES	NO	YES
Manually created slides	NO	NO	NO	NO	YES
Program development	NO	NO	NO	NO	NO

Administering Linux in Production Environments
Copyright © 1999-2001, Exponential Consulting, LLC

181.1

227

Required Applications: Everyday

- File manager
- Outlook
 - Email
 - Schedule
 - Contacts
- Web browser
- Winzip
- Music (CD, MP3)
- Video
- Burn CD
- PDF viewer
- PDA interface
- Games
- Expenses
- Fax

Administering Linux in Production Environments
Copyright © 1999-2001, Exponential Consulting, LLC

182'

228

Required Applications: Specialty

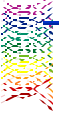
- Personal/Baby database
- Image editing
- Vector-based illustration
- HTML editor
- 3D drawing
- Program development

Administering Linux in Production Environments
Copyright © 1999-2001, Exponential Consulting, LLC

182.1

229

Required Applications: Administrative



- Database with development environment
- Business finances
- Web server
- FTP server
- File sharing
- Printer sharing
- Remote installation facility
- Scripting

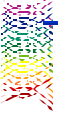
Administering
Linux in
Production
Environments

Copyright © 1999-2001,
Exponential Consulting, LLC.

182.2

230

Select the Right Distribution



- Lots of ported, pre-compiled packages
 - SuSE
 - Recent RedHat
 - ?

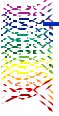
Administering
Linux in
Production
Environments

Copyright © 1999-2001,
Exponential Consulting, LLC.

182.3

231

Into the Future



- “These are Linux’s early days ...”
 - The beginning of the end?
 - The end of the beginning?

Administering
Linux in
Production
Environments

Copyright © 1999-2001,
Exponential Consulting, LLC.

183