1. Linux Boot Process

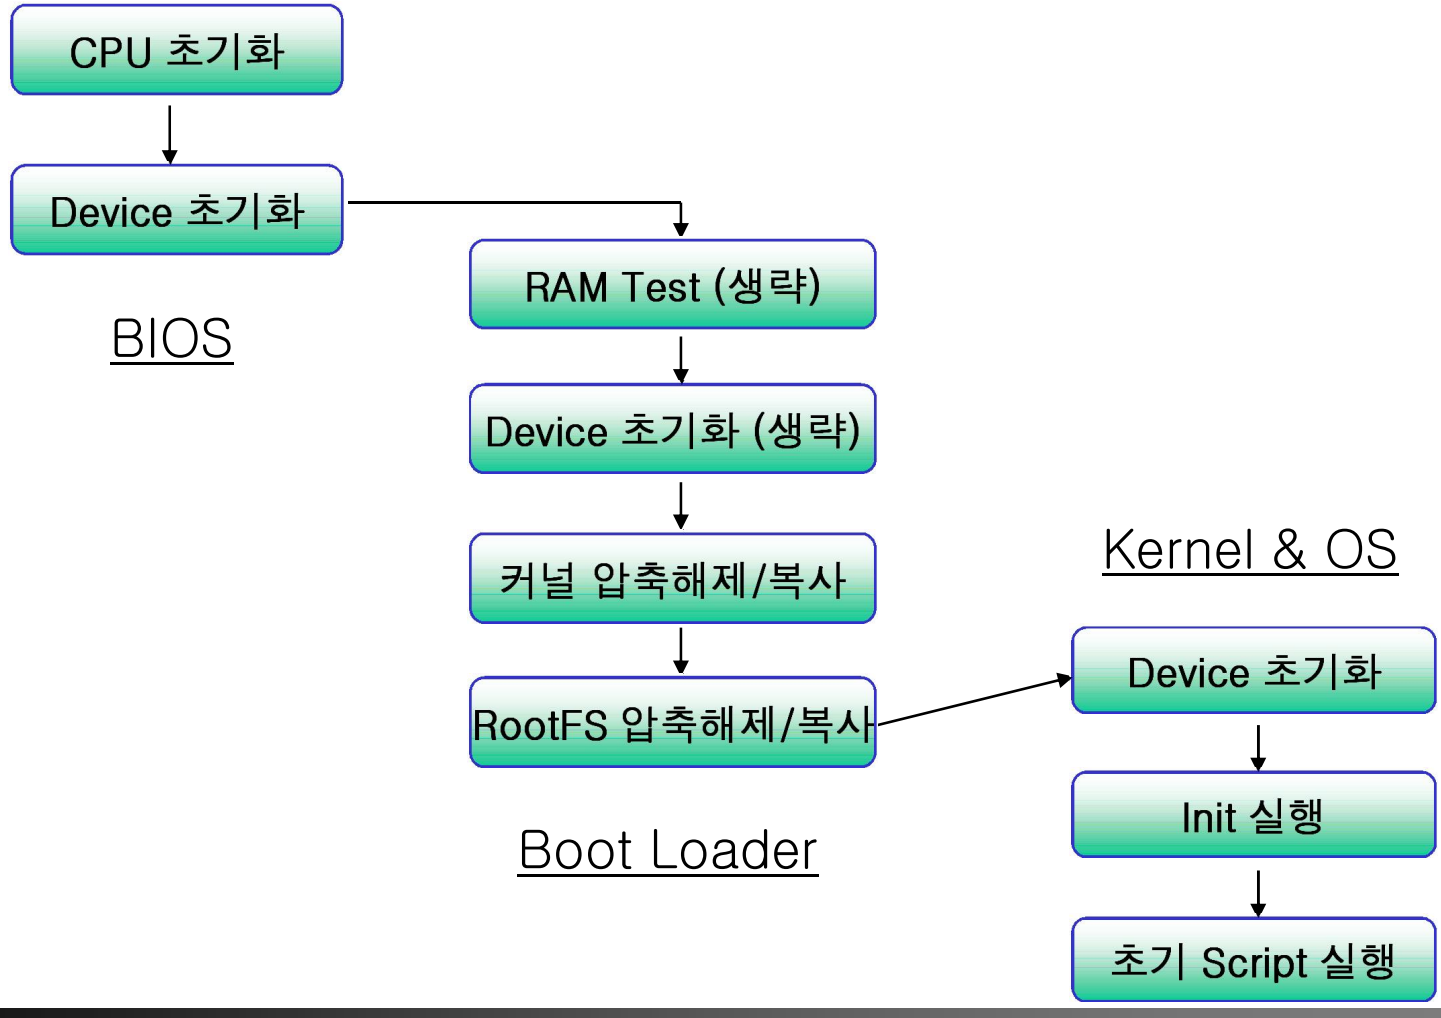

BIOS

- PC의 전원을 입력 후 ROM-BIOS 실행
- Hardware의 물리적 손상 검사 및 초기화
  (CPU, Memory 등의 각종 Device)
- Booting 매체를 검색
  (Hard disk, Floppy disk, CD Rom 등)
- Booting 매체의 MBR(Master Boot Record : 0 sector)에 있는 Boot Program을 Memory에 적재

BIOS(cont)

Boot Program

Boot Loader 또는 Boot Sector라고도 하며, Linux의 Boot Loader를 LILO(Linux Loader)라 함.

(RedHat ver9.0 에서는 GRUB라 함.)

BIOS는 LILO를 Memory에 적재한 후, 모든 제어권을 LILO에게 넘기고 Memory에서 사라짐.

LILO에서 어떤 OS로 Booting을 할지 선택을 하기도 함.


GRUB version 0.93 (638K lower / 128960K upper memory)
Red Hat Linux (2.4.20-8)
Use the ↑ and ↓ keys to select which entry is highlighted.
Press enter to boot the selected OS, 'e' to edit the
commands before booting, 'a' to modify the kernel arguments
before booting, or 'c' for a command-line.
The highlighted entry will be booted automatically in 9 seconds.
redhat.

(2) Boot Loader

RAM test

Memory의 상태를 test

Memory의 크기를 인식

(Memory Scanning이라고도 하며, 이 과정을 생략하고 미리 크기만 setting 되었을 수도 있음.)

(2) Boot Loader

Device 초기화

Device 마다 Register setting, IRQ 할당, Base Address할당 등의 초기화 함.

(경우에 따라 BIOS 대신 Boot Loader가 할 수도 있고, Linux Kernel에서도 초기화를 다시 하므로 Boot Loader가 하지 않을 수도 있음.)

(2) Boot Loader

Kernel 압축해제 / 복사

압축된 Kernel image(vmlinuz)를 Memory의 임의의 번지에 복사 후 압축을 푼다.
다시 Memory의 설정된 Kernel 시작 번지로 Kernel image를 복사한다.

(2) Boot Loader

RootFS 압축해제 / 복사

압축된 RootFS(Root File System) image를 Memory의 임의의 번지에 복사 후 압축을 품.

다시 Memory의 설정된 RootFS 시작 번지로 RootFS image를 복사함.

```
  Booting 'Red Hat Linux (2.4.20-8)'

root (hd0,0)
 Filesystem type is ext2fs, partition type 0x83
kernel /vmlinuz-2.4.20-8 ro root=LABEL=/
   [Linux-bzImage, setup=0x1400, size=0x11098a]
initrd /initrd-2.4.20-8.img
   [Linux-initrd @ 0x7ea1000, 0x3e25c bytes]

Uncompressing Linux... Ok, booting the kernel.
_
```

(3) Kernel & OS

Device 초기화

앞서 Boot Loader에서 언급한 것과 같이 Device 마다 Register setting, IRQ 할당, Base Address할당 등의 초기화 함.

```
Linux version 2.4.20-8 (bhcompile@porky.devel.redhat.com) (gcc version 3.2.2 200
30222 (Red Hat Linux 3.2.2-5)) #1 Thu Mar 13 17:54:28 EST 2003
BIOS-provided physical RAM map:
 BIOS-e820: 0000000000000000 - 000000000009f800 (usable)
 BIOS-e820: 000000000009f800 - 00000000000a0000 (reserved)
 BIOS-e820: 00000000000ca000 - 00000000000cc000 (reserved)
 BIOS-e820: 00000000000dc000 - 0000000000100000 (reserved)
 BIOS-e820: 0000000000100000 - 0000000007ef0000 (usable)
 BIOS-e820: 0000000007ef0000 - 0000000007efc000 (ACPI data)
 BIOS-e820: 0000000007efc000 - 0000000007f00000 (ACPI NVS)
 BIOS-e820: 0000000007f00000 - 0000000008000000 (usable)
 BIOS-e820: 00000000fec00000 - 00000000fec10000 (reserved)
 BIOS-e820: 00000000fee00000 - 00000000fee01000 (reserved)
 BIOS-e820: 00000000fffe0000 - 0000000100000000 (reserved)
0MB HIGHMEM available.
128MB LOWMEM available.
On node 0 totalpages: 32768
zone(0): 4096 pages.
zone(1): 28672 pages.
zone(2): 0 pages.
Kernel command line: ro root=LABEL=/
Initializing CPU#0
Detected 997.536 MHz processor.
Console: colour VGA+ 80x25
Calibrating delay loop... _
```

```
Calibrating delay loop... 1933.31 BogoMIPS
Memory: 124584k/131072k available (1347k kernel code, 5012k reserved, 999k data,
 132k init, 0k highmem)
Dentry cache hash table entries: 16384 (order: 5, 131072 bytes)
Inode cache hash table entries: 8192 (order: 4, 65536 bytes)
Mount cache hash table entries: 512 (order: 0, 4096 bytes)
Buffer-cache hash table entries: 8192 (order: 3, 32768 bytes)
Page-cache hash table entries: 32768 (order: 5, 131072 bytes)
CPU: L1 I cache: 16K, L1 D cache: 16K
CPU: L2 cache: 256K
Intel machine check architecture supported.
Intel machine check reporting enabled on CPU#0.
CPU: Intel Pentium III (Coppermine) stepping 08
Enabling fast FPU save and restore... done.
Enabling unmasked SIMD FPU exception support... done.
Checking 'hlt' instruction... OK.
POSIX conformance testing by UNIFIX
mtrr: v1.40 (20010327) Richard Gooch (rgooch@atnf.csiro.au)
mtrr: detected mtrr type: Intel
PCI: PCI BIOS revision 2.10 entry at 0xfd9a0, last bus=1
PCI: Using configuration type 1
PCI: Probing PCI hardware
PCI: Using IRQ router PIIX [8086/7110] at 00:07.0
PCI: Cannot allocate resource region 4 of device 00:07.1
Limiting direct PCI/PCI transfers.
isapnp: Scanning for PnP cards...
```

```
isapnp: Scanning for PnP cards...
isapnp: No Plug & Play device found
Linux NET4.0 for Linux 2.4
Based upon Swansea University Computer Society NET3.039
Initializing RT netlink socket
apm: BIOS version 1.2 Flags 0x03 (Driver version 1.16)
Starting kswapd
VFS: Disk quotas vdquot_6.5.1
pty: 2048 Unix98 ptys configured
Serial driver version 5.05c (2001-07-08) with MANY_PORTS MULTIPORT SHARE_IRQ SER
IAL_PCI ISAPNP enabled
ttyS0 at 0x03f8 (irq = 4) is a 16550A
ttyS1 at 0x02f8 (irq = 3) is a 16550A
Real Time Clock Driver v1.10e
Floppy drive(s): fd0 is 1.44M
FDC 0 is a post-1991 82077
NET4: Frame Diverter 0.46
RAMDISK driver initialized: 16 RAM disks of 4096K size 1024 blocksize
Uniform Multi-Platform E-IDE driver Revision: 7.00beta-2.4
ide: Assuming 33MHz system bus speed for PIO modes; override with idebus=xx
PIIX4: IDE controller at PCI slot 00:07.1
PIIX4: chipset revision 1
PIIX4: not 100% native mode: will probe irqs later
    ide0: BM-DMA at 0x1470-0x1477, BIOS settings: hda:DMA, hdb:pio
hda: VMware Virtual IDE CDROM Drive, ATAPI CD/DVD-ROM drive
ide0 at 0x1f0-0x1f7,0x3f6 on irq 14
ide-floppy driver 0.99.newide
ide-floppy driver 0.99.newide
md: md driver 0.90.0 MAX_MD_DEVS=256, MD_SB_DISKS=27
md: Autodetecting RAID arrays.
md: autorun ...
md: ... autorun DONE.
NET4: Linux TCP/IP 1.0 for NET4.0
IP Protocols: ICMP, UDP, TCP, IGMP
IP: routing cache hash table of 1024 buckets, 8Kbytes
TCP: Hash tables configured (established 8192 bind 16384)
Linux IP multicast router 0.06 plus PIM-SM
```

```
Linux IP multicast router 0.06 plus PIM-SM
NET4: Unix domain sockets 1.0/SMP for Linux NET4.0.
RAMDISK: Compressed image found at block 0
Freeing initrd memory: 248k freed
VFS: Mounted root (ext2 filesystem).
Red Hat nash version 3.4.42 starting
Loading scsi_mod.o module
SCSI subsystem driver Revision: 1.00
Loading sd_mod.o module
Loading BusLogic.o module
PCI: Found IRQ 11 for device 00:10.0
scsi: ***** BusLogic SCSI Driver Version 2.1.15 of 17 August 1998 *****
scsi: Copyright 1995-1998 by Leonard N. Zubkoff <lnz@dandelion.com>
scsi0: Configuring BusLogic Model BT-958 PCI Wide Ultra SCSI Host Adapter
scsi0:   Firmware Version: 5.07B, I/O Address: 0x1440, IRQ Channel: 11/Level
scsi0:   PCI Bus: 0, Device: 16, Address: 0xEC800000, Host Adapter SCSI ID: 7
scsi0:   Parity Checking: Enabled, Extended Translation: Enabled
scsi0:   Synchronous Negotiation: Ultra, Wide Negotiation: Enabled
scsi0:   Disconnect/Reconnect: Enabled, Tagged Queuing: Enabled
scsi0:   Scatter/Gather Limit: 128 of 8192 segments, Mailboxes: 211
scsi0:   Driver Queue Depth: 211, Host Adapter Queue Depth: 192
scsi0:   Tagged Queue Depth: Automatic, Untagged Queue Depth: 3
scsi0:   Error Recovery Strategy: Default, SCSI Bus Reset: Enabled
scsi0: *** BusLogic BT-958 Initialized Successfully ***
_
```

```
scsi0: *** BusLogic BT-958 Initialized Successfully ***
scsi0 : BusLogic BT-958
  Vendor: VMware,    Model: VMware Virtual S  Rev: 1.0
  Type:   Direct-Access                       ANSI SCSI revision: 02
scsi0: Target 0: Queue Depth 28, Asynchronous
scsi0: Target 1: Queue Depth 3, Asynchronous
scsi0: Target 2: Queue Depth 3, Asynchronous
scsi0: Target 3: Queue Depth 3, Asynchronous
scsi0: Target 4: Queue Depth 3, Asynchronous
scsi0: Target 5: Queue Depth 3, Asynchronous
scsi0: Target 6: Queue Depth 3, Asynchronous
scsi0: Target 7: Queue Depth 3, Asynchronous
scsi0: Target 8: Queue Depth 3, Asynchronous
scsi0: Target 9: Queue Depth 3, Asynchronous
scsi0: Target 10: Queue Depth 3, Asynchronous
scsi0: Target 11: Queue Depth 3, Asynchronous
scsi0: Target 12: Queue Depth 3, Asynchronous
scsi0: Target 13: Queue Depth 3, Asynchronous
scsi0: Target 14: Queue Depth 3, Asynchronous
scsi0: Target 15: Queue Depth 3, Asynchronous
Attached scsi disk sda at scsi0, channel 0, id 0, lun 0
SCSI device sda: 10485760 512-byte hdwr sectors (5369 MB)
Partition check:
 sda: sda1 sda2 sda3
Loading jbd.o module
Journalled Block Device driver loaded
Loading ext3.o module
Mounting /proc filesystem
Creating block devices
Creating root device
Mounting root filesystem
EXT3-fs: INFO: recovery required on readonly filesystem.
EXT3-fs: write access will be enabled during recovery.
scsi0: Tagged Queuing now active for Target 0
kjournald starting.  Commit interval 5 seconds
EXT3-fs: recovery complete.
EXT3-fs: mounted filesystem with ordered data mode.
Freeing unused kernel memory: 132k freed
INIT: version 2.84 booting
```

(3) Kernel & OS

Init 실행

Kernel 이 실행되면서 Swapper(pid 0) Process 가 실행 됨. (Memory 관리, Disk 관리, Process 관리 등 수행)

pid 0인 Process는 pid 1인 init Process를 실행 함.

init가 처음 수행하는 파일은 /etc/inittab임.

(3) Kernel & OS

Swapper Process의 기능

Sleep 되어 있는 Process를 Memory에서 Disk 공간에 잠시 Swap시켰다가 그 Process를 사용할 때, 다시 Memory에 적재시키는 일을 함.

(3) Kernel & OS

pid 1인 init Process (/sbin/init)

- File System의 구조를 검사
- File System mount
- Sever Demon을 띄움
- 사용자 login
- login 후 Shell을 띄움

(3) Kernel & OS

/etc/inittab

System의 상태에 따라 해당하는 runlevel에서 init Process가 처리할 일들을 지정하고 있는 파일임.

- runlevel 지정
- /etc/rc.d/rc.sysinit 실행
- /etc/tc.d/rc 실행
- rungetty에 의해 자동 login

(3) Kernel & OS

Runlevel

0 – System을 종료할 때 사용 (Do NOT set init default to this)
1 – Single user mode에서 사용
2 – Multi user, without NFS (The same as 3, if you do not havenetworking)
3 – Full multi user mode에서 사용
4 – 사용자가 정의해서 사용
5 – X를 실행시키기 위해 사용
6 – reboot (Do NOT set init default to this)

(3) Kernel & OS

초기 Script 실행

/etc/rc.d/rc.sysinit
(System 초기화시 맨 먼저 딱 한 번 실행됨.)

/etc/rc.d/rc
(각종 Demon들을 실행시키고 종료시킴.)

(3) Kernel & OS

/etc/rc.d/rc.sysinit

- 기본적인 path 설정
- /etc/sysconfig/network가 있으면 그 script 실행
- Keymap의 loading
- System font loading
- Swap 영역 활성화
- Disk 검사(fsck)
- /proc File System의 mount

(3) Kernel & OS

- Root File System을 rw mode로 다시 mount
- /etc/HOSTNAME file의 설정
- /etc/mtab file에 root와 /proc File System의 entry 추가하기
- Kernel module loading
- System 시간 설정
- 등등…

(3) Kernel & OS

/etc/rc.d/rc

ls /etc/rc.d [Enter]
rc0.d
rc1.d
rc2.d
rc3.d
rc4.d
rc5.d
rc6.d

(3) Kernel & OS

위 각 level 마다 있는 directory는 S와 K로 시작하는 file로 구성되며 거의 대부분 /etc/rc.d/init.d 아래의 file들에 link되어 있음.

S : 해당 service를 실행시키는 file

K : 해당 service를 종료시키는 file

```
INIT: version 2.84 booting
Setting default font (lat0-sun16):                         [  OK  ]

               Welcome to Red Hat Linux
               Press 'I' to enter interactive startup.
Mounting proc filesystem:                                  [  OK  ]
Unmounting initrd:                                         [  OK  ]
Configuring kernel parameters:                             [  OK  ]
Setting clock  (localtime): Mon Apr 25 14:21:27 KST 2005   [  OK  ]
Loading default keymap (us):                               [  OK  ]
Setting hostname localhost.localdomain:                    [  OK  ]
Initializing USB controller (usb-uhci):                    [  OK  ]
Mounting USB filesystem:                                   [  OK  ]
Initializing USB HID interface:                            [  OK  ]
Initializing USB keyboard:                                 [  OK  ]
Initializing USB mouse:                                    [  OK  ]
Your system appears to have shut down uncleanly
Press Y within 1 seconds to force file system integrity check...
Checking root filesystem
/: clean, 145539/565600 files, 603807/1130574 blocks
                                                           [  OK  ]
Remounting root filesystem in read-write mode:             [  OK  ]
Activating swap partitions:                                [  OK  ]
Finding module dependencies:                               [  OK  ]
Checking filesystems
/boot: recovering journal
/boot: clean, 41/26104 files, 12728/104391 blocks
                                                           [  OK  ]
Mounting local filesystems:                                [  OK  ]
Enabling local filesystem quotas:                          [  OK  ]
Enabling swap space:                                       [  OK  ]
INIT: Entering runlevel: 5
```

```
INIT: Entering runlevel: 5
Entering non-interactive startup
Checking for new hardware                                  [  OK  ]
Updating /etc/fstab                                        [  OK  ]
Flushing all current rules and user defined chains:        [  OK  ]
Clearing all current rules and user defined chains:        [  OK  ]
Applying iptables firewall rules:                          [  OK  ]
Setting network parameters:                                [  OK  ]
Bringing up loopback interface:                            [  OK  ]
Bringing up interface eth0:  Error, some other host already uses address 211.57.
110.110.
                                                           [FAILED]
Starting system logger:                                    [  OK  ]
Starting kernel logger:                                    [  OK  ]
Starting portmapper:                                       [  OK  ]
Starting NFS statd:                                        [  OK  ]
Starting keytable:                                         [  OK  ]
Initializing random number generator:                      [  OK  ]
Starting pcmcia:                                           [  OK  ]
Mounting other filesystems:                                [  OK  ]
Starting up APM daemon:                                    [  OK  ]
Starting automount: No Mountpoints Defined                 [  OK  ]
Starting sshd:                                             [  OK  ]
Starting xinetd:                                           [  OK  ]
Starting sendmail:                                         [  OK  ]
Starting sm-client:                                        [  OK  ]
Starting console mouse services:                           [  OK  ]
Starting Canna server:                                     [  OK  ]
Starting crond:                                            [  OK  ]
Starting cups:                                             [  OK  ]
Starting xfs:                                              [  OK  ]
Starting anacron:                                          [  OK  ]
Starting atd:                                              [  OK  ]
_
```

(3) Kernel & OS

rungetty에 의한 자동 login

- swapper(pid 0) process에 의해 생성된 init process가 fork()를 함.
- fork()의해 생성된 자식 init가 exec()를 이용하여 rungetty process를 실행.
- rungetty process는 다시 exec()을 하여 login precess를 수행
- 자동 login option(-autologin root)을 통해 login 후 shell을 띄움.

```
Red Hat Linux release 9 (Shrike)
Kernel 2.4.20-8 on an i686

localhost login: _
```

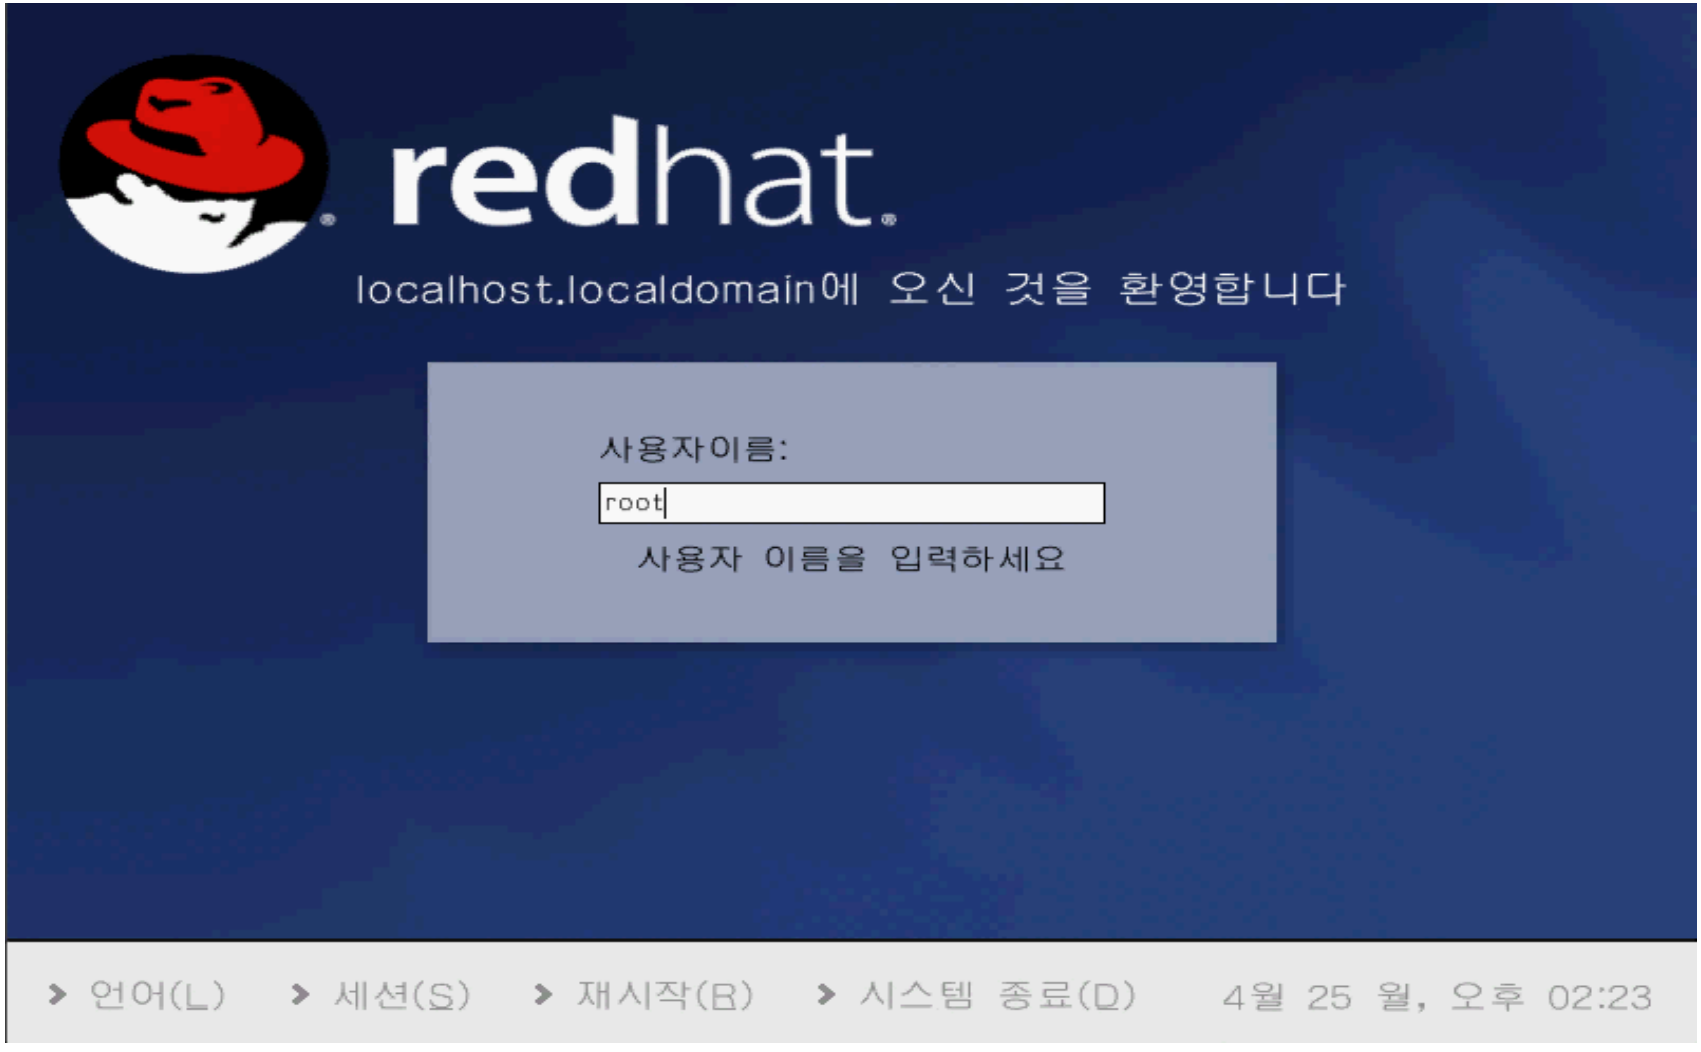
redhat
localhost.localdomain에 오신 것을 환영합니다
사용자이름:
root
사용자 이름을 입력하세요
언어(L)
세션(S)
재시작(R)
시스템 종료(D)
4월 25 월, 오후 02:23

(2) init 과정

런레벨 지정

/etc/r.cd/rc.sysinit 실행 → 패스설정(PATH=/bin:/sbin:/usr/bin)
/ 파일시스템 mount
HOSTNAME 설정
/proc 파일시스템 mount

l3:3:wait:/etc/rc.d/rc 3 실행 → S22network 스크립트 파일 실행
- Network 설정
S50inet 스크립트 파일 실행
- inet daemon 실행

rungetty를 통해 자동 login

↓

bash 실행

(3) /etc/inittab code

id:runlevels:action:process

id : 해당 state를 구분하기 위한 레이블.
runlevels : 그 줄의 내용을 적용하기 위한 runlevel의 목록.
action : entry에 의해 실행되는 process를 어떻게 할 것 인가에 대한 설명.
process : process의 실행 file의 경로와 process 에 넘겨줄 인수.