

Statistic Analysis of The Linux Kernel Development

Tsugikazu SHIBATA
NEC
27, September 2010

\$ who am i

- ✓ Tsugikazu Shibata
- ✓ OSS Promotion center, NEC
- ✓ Member of the Board of The Linux Foundation
- ✓ Fostering industry's developer
- ✓ Maintaining a Japanese version of HOWTO at [Documentation/ja_JP/HOWTO](https://www.linux.org/docs/ja_JP/HOWTO)

利
奈
久
州
上
等



Agenda

- Status of latest kernel
- Status of series of kernel development
- Status of distributions

Status of Latest Kernel

- Latest release: 2.6.35
Released: 1, Aug. 2010
- Latest mailline kernel version : 2.6.36-rc5
Released: 20, Sep. 2010
- Latest stable release:
2.6.35.5
2.6.34.7
Long-term stable release
2.6.27.54
2.6.32.22

Linux 2.6.35 status

Released at: August 1st, 2010

Size of tar.bz2: 69,305,709
(=69.3MB)

of Lines: 13,465,772
(=13.5ML)

of Files: 33,315

Comparison with other software

- How large Linux is?
- Firefox is a great Web browser!

Software	Source code size of tar.bz2	# of Lines	# of Files
Linux-2.6.35	69.3MB	13.5Mlines	33,315
Firefox-3.5.12	23.9MB	4.2Mlines	8,346

Comparison with other software

- How large Linux is?
- No GCC? No life !

Software	Source code size of tar.bz2	# of Lines	# of Files
Linux-2.6.35	69.3MB	13.5Mlines	33,315
Firefox-3.5.12	23.9MB	4.2Mlines	8,346
GCC-4.5.1	86.1MB	11.4Mlines	65,841

Linux top directory

```
$ls -F ./linux-2.6.35
```

```
COPYING          REPORTING-BUGS  include/        samples/  
CREDITS          arch/           init/           scripts/  
Documentation/   block/         ipc/            security/  
Kbuild           crypto/        kernel/        sound/  
MAINTAINERS     drivers/       lib/            tools/  
Makefile        firmware/     mm/             usr/  
README          fs/           net/            virt/
```


% of 2nd directories

```
$ls -F ./linux-2.6.35
```

```
COPYING          REPORTING-BUGS  include/        samples/  
CREDITS          arch/           init/           scripts/  
Documentation/   block/          ipc/            security/  
Kbuild          crypto/         kernel/         sound/  
MAINTAINERS      drivers/        lib/            tools/  
Makefile         firmware/       mm/             usr/  
README          fs/             net/            virt/
```

Directory name	# of Lines under dir	%
driver	6,883,548	51.2
arch	2,932,985	21.8
fs	971,067	7.2
net	629,148	4.7
Sound	614,601	4.6

of lines and % in 2nd directories

Linux-2.6.35:	Lines	%
drivers	6883548	51.12%
arch	2932985	21.78%
fs	971067	7.21%
net	629148	4.67%
sound	614601	4.56%
include	428283	3.18%
Documentation	350579	2.60%
kernel	165071	1.23%
firmware	119664	0.89%
mm	75406	0.56%
scripts	64357	0.48%
crypto	48441	0.36%
tools	47680	0.35%
security	46650	0.35%
lib	36863	0.27%
block	18767	0.14%
“.”	12990	0.10%
ipc	7784	0.06%
virt	5323	0.04%
init	4393	0.03%
samples	1207	0.01%
usr	965	0.01%
	-----	-----
Total:	13465772	100%

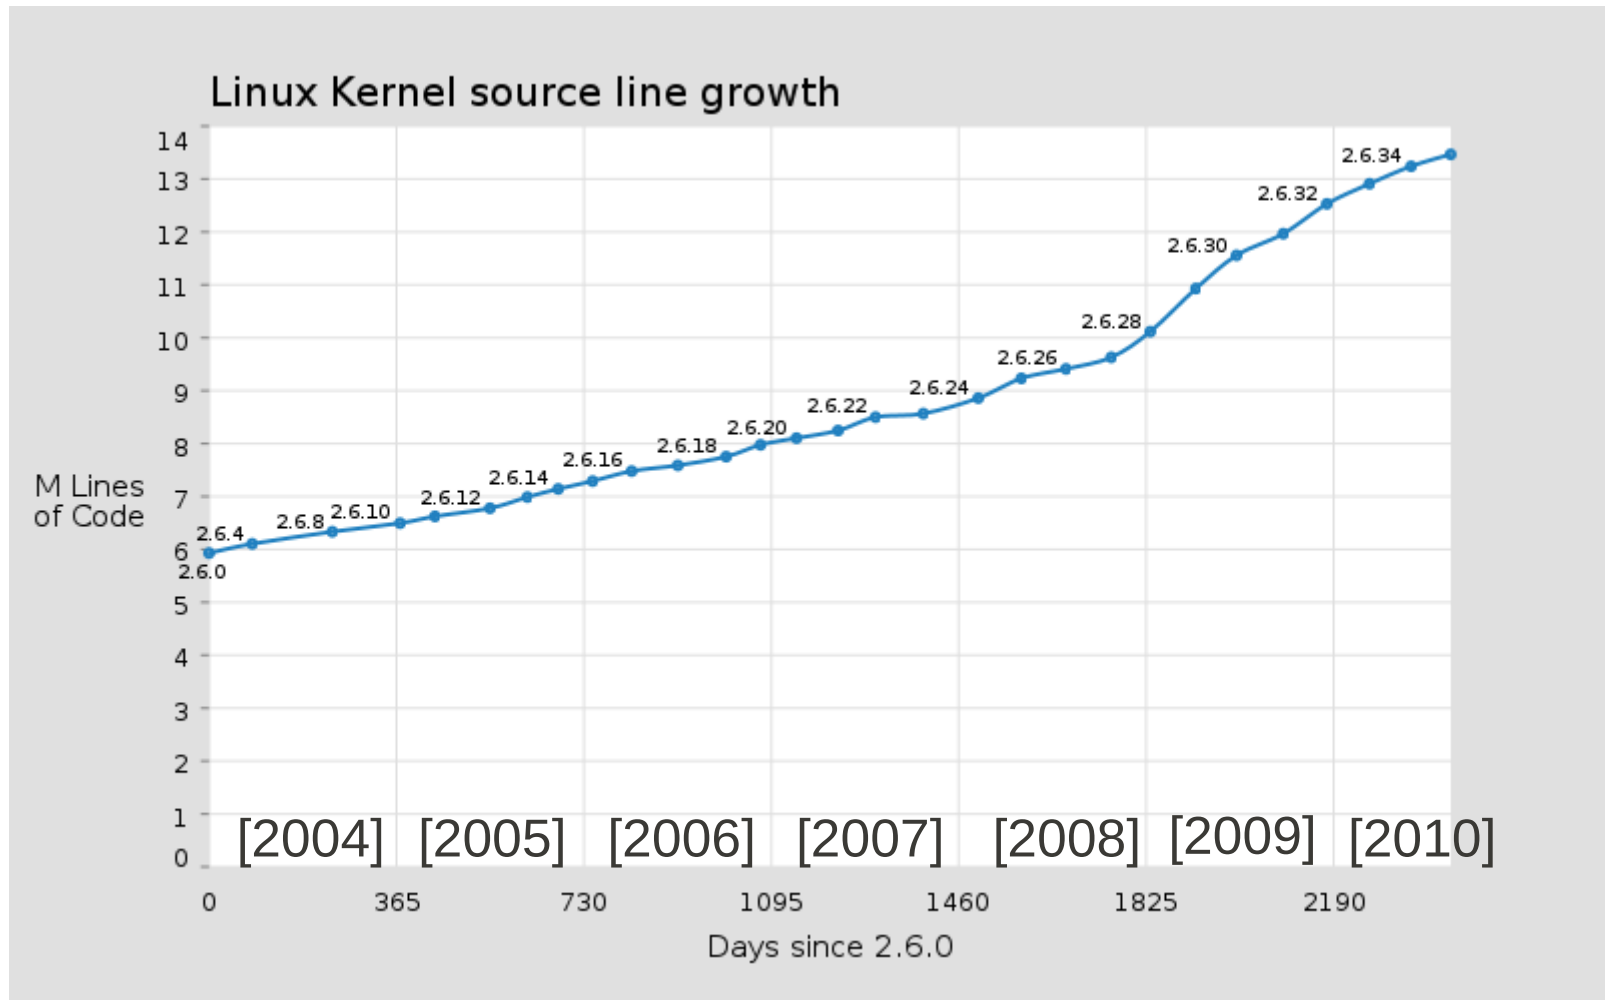
Summary of 2.6.35

- 69.3MBytes of tar.bz2
 - Reasonably large compared with other software
- 13.5MLines of code
 - Huge value of open source development result
- Major parts are
 - Driver 51%
 - Arch 22%
 - Fs 7%
 - Others are under 1%

Status of series of Kernel development

Source code growth

- 2.4ML grew in 2009, about 1ML in this year



of lines in each releases

Version	Lines	added	%

2.6.25	9232541	-	4.2%
2.6.26	9411790	179249	1.9%
2.6.27	9630023	218233	2.3%
2.6.28	10115662	485639	5.0%
2.6.29	10930802	815140	8.1%
2.6.30	11557329	626527	5.7%
2.6.31	11966482	409153	3.5%
2.6.32	12530196	563714	4.7%
2.6.33	12910203	380007	3.0%
2.6.34	13241101	330898	2.6%
2.6.35	13465772	224671	1.7%

Added : dfference with previous release

of Lines/Files developed per year

- Over One million of lines had been added each years

	# of Lines added	# of Files added
2008	1,255,979	2,192
2009	2,414,534	5,230
2010	935,576	2,831

Notes: Each year defined as follows:

2008: 2.6.24-2.6.28

2009: 2.6.28-2.6.32

2010: 2.6.32-2.6.35

of lines developed per day

- More than 4000 lines and 10 files are added per day

Year	# of Lines/Day	# of Files/Day
2008 (335 Days)	3749	6.5
2009 (378 Days)	6388	13.8
2010 (207 Days)	4520	13.7

Notes: Each year means as follows:


2008: 2.6.24-2.6.28

2009: 2.6.28-2.6.32

2010: 2.6.32-2.6.35

Number of developers sent patches

- More than 1000 of developers patch are merged since 2.6.25



Developer sent patch a year

- More than two thousands of developers patch had been merged in each year

	# of developers
2008	2195
2009	2467
2010	2014

Notes: This data comes from gitdm with multiple release as below:

2008: 2.6.24-2.6.28

2009: 2.6.28-2.6.32


2010: 2.6.32-2.6.35

of lines developed per year developer

	# of Lines	# of developers	# of lines/developers
2008	1,255,979	2195	572
2009	2,414,534	2467	978
2010	935,576	2014	464

Development Process

- Merge window and Release candidates
 - 2 weeks of merge window just after kernel released
 - Each RC takes a week or 10 days, continue 8-9 times



Kernel	2.6.26	2.6.27	2.6.28	2.6.29	2.6.30	2.6.31
Max -RC	-rc9	-rc9	-rc9	-rc8	-rc8	-rc9

Days spend for each releases w/ RCs


Version	Release	Days	-RC

2.6.26	2008/07/13	87	- rc9
2.6.27	2008/10/09	88	- rc9
2.6.28	2008/12/24	76	- rc9
2.6.29	2009/03/23	89	- rc8
2.6.30	2009/06/10	79	- rc8
2.6.31	2009/09/09	91	- rc9
2.6.32	2009/12/3	85	- rc8
2.6.33	2010/02/24	83	- rc8
2.6.34	2010/05/16	81	- rc7
2.6.35	2010/08/01	77	- rc6




Days to develop each releases

- Average 81 days in last 4 releases
- Average 84 days in last 10 releases



Linux-next for next merge window

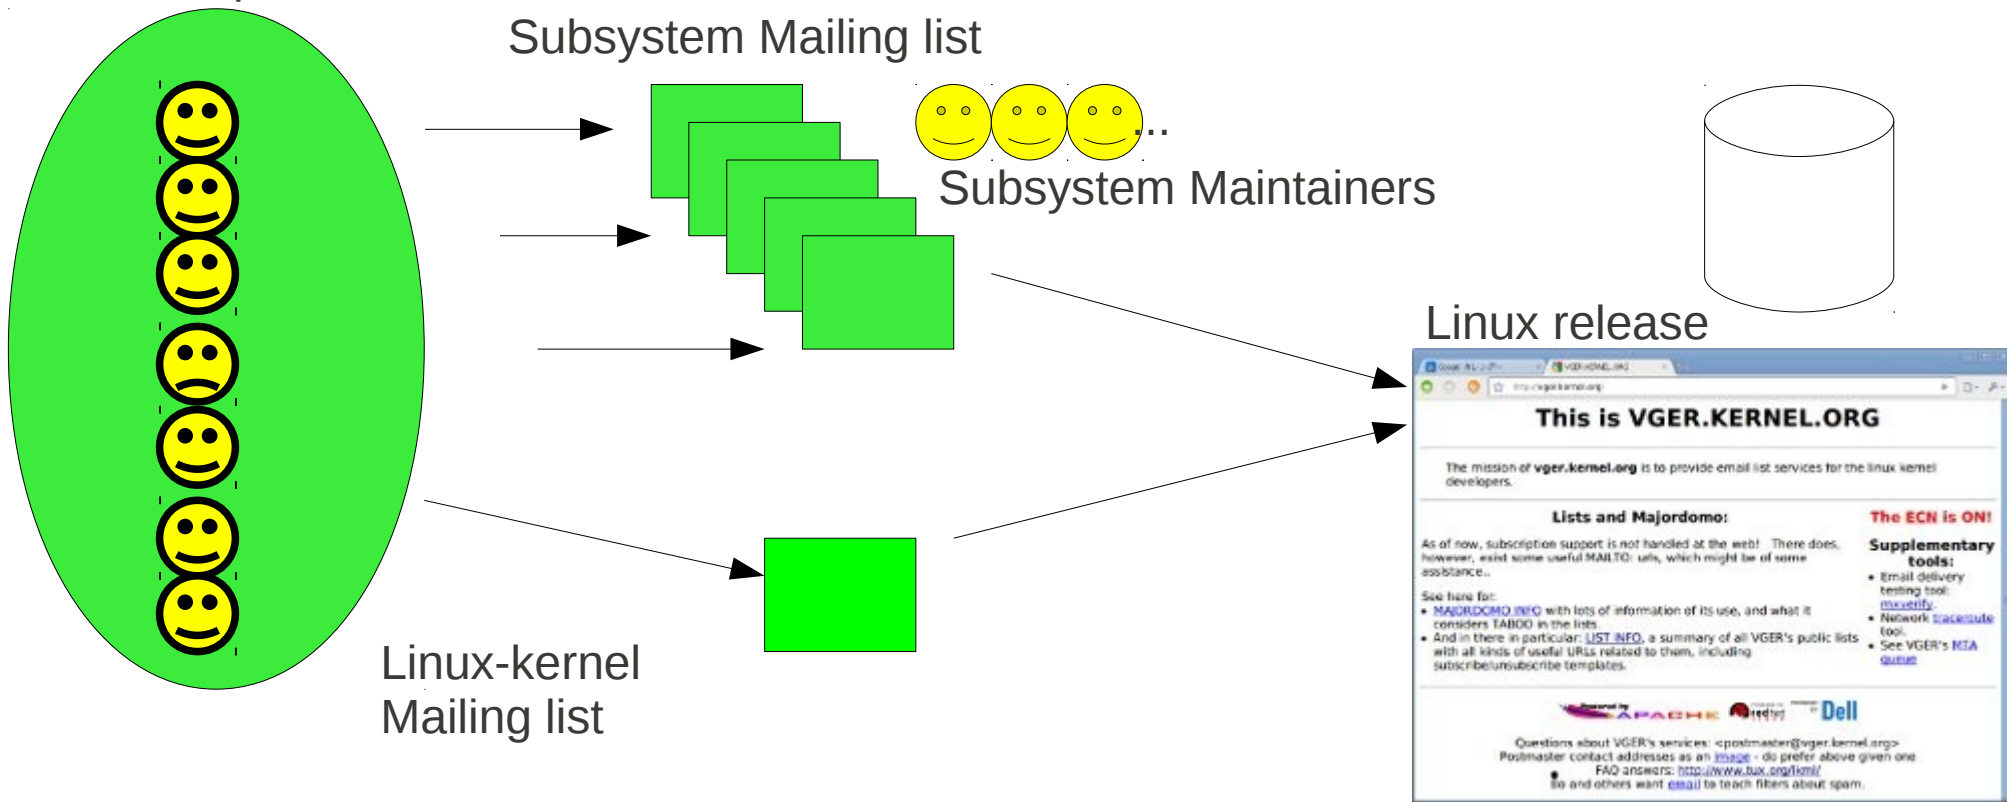
- Do automated build test and maintainers get notification about build problems on almost daily basis
- Merging 171 of subsystem trees to prepare next merge window (It was 141 in 2009 October)



Send patch and mailing list

- 129 mailing lists in VGER.KERNEL.ORG.
- Subscribers in Linux-kernel mailing list are 5591 (last year: 5491).
- Total subscribers of 129 mailing lists are 34464 (last year 31594).

Developers



Linux-kernel Mailing list


Subsystem Mailing list

Subsystem Maintainers

Linux release

Number of MAINTAINERS

- 847 maintainers in 2.6.35



Choosing feature of Linux

- Make manuconfig (or xconfig)
- There are 693 of Kconfig files in 2.6.35
- 10484 of configuration items


```
ファイル(E) 編集(E) 表示(V) 端末(T) ヘルプ(H)
.config - Linux Kernel v2.6.25.1 Configuration

Linux Kernel Configuration
Arrow keys navigate the menu. <Enter> selects submenus --->. Highlighted
letters are hotkeys. Pressing <Y> includes, <N> excludes, <M> modularizes
features. Press <Esc><Esc> to exit, <?> for Help, </> for Search. Legend: [*]
built-in [ ] excluded <M> module < > module capable

  General setup --->
  [*] Enable loadable module support --->
  -* Enable the block layer --->
  Processor type and features --->
  Power management options --->
  Bus options (PCI etc.) --->
  Executable file formats / Emulations --->
  Networking --->
  ||| Device Drivers --->
  Firmware Drivers --->
  File systems --->
  Kernel hacking --->
  Security options --->
  -* Cryptographic API --->
  [*] Virtualization --->
  Library routines --->
  ---
  Load an Alternate Configuration File
  v(+)
```

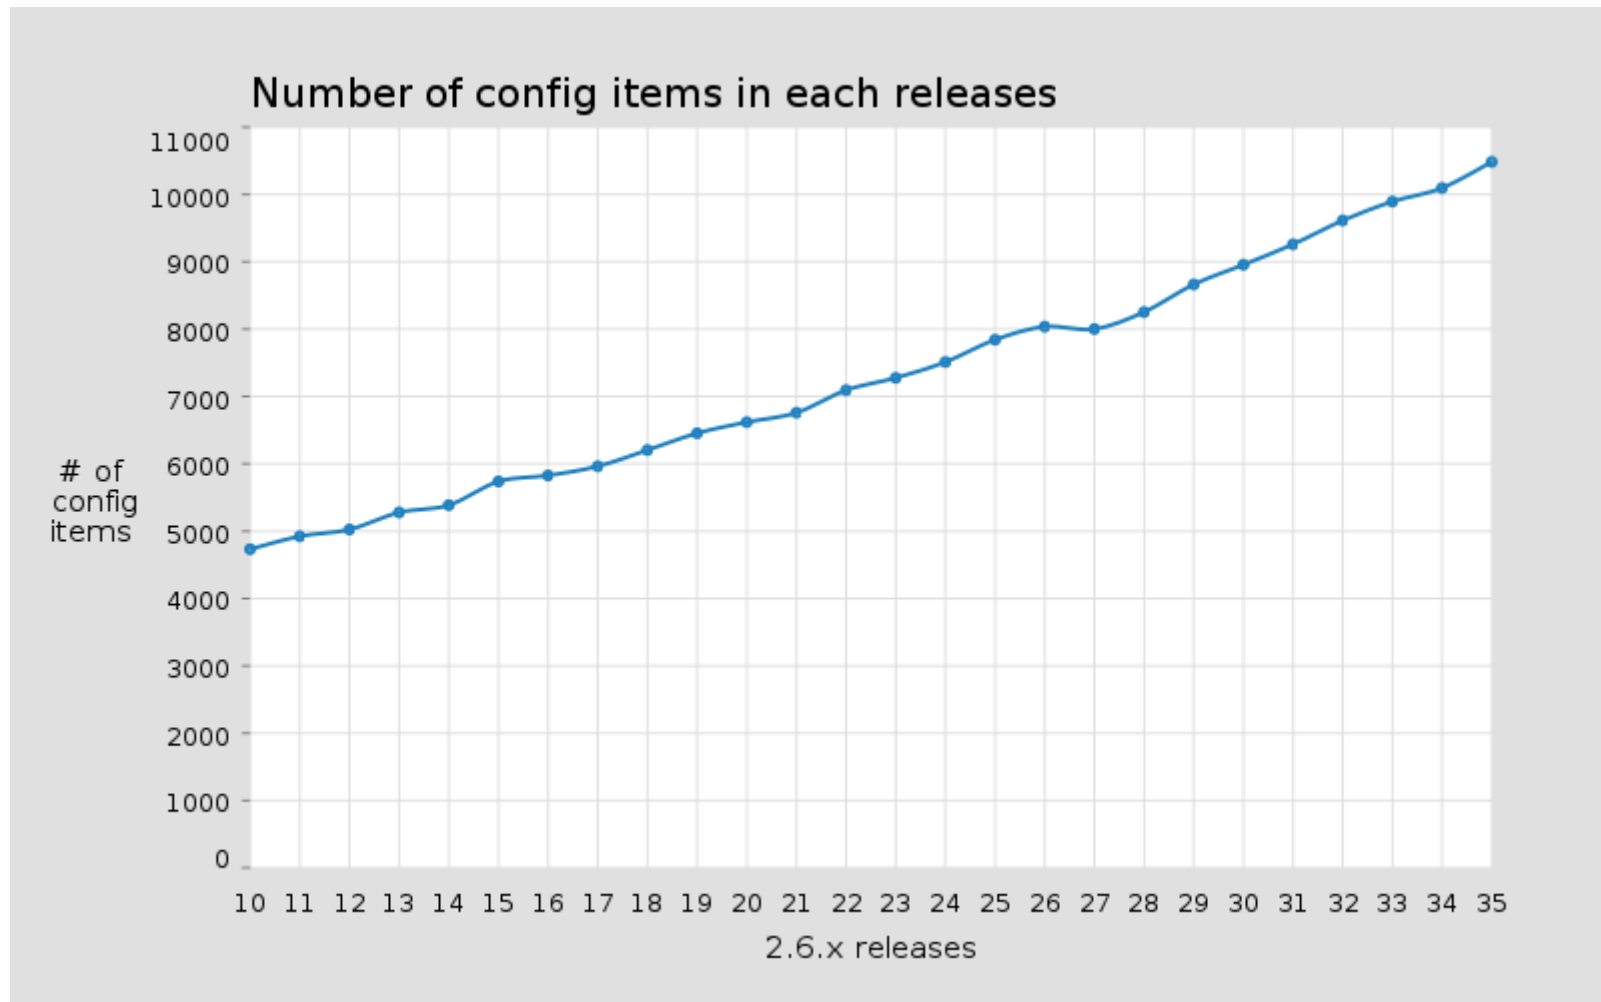
Growth of configuration files

- 693 of configuration files in 2.6.35



of configuration items

- 10484 of configuration items in 2.6.35



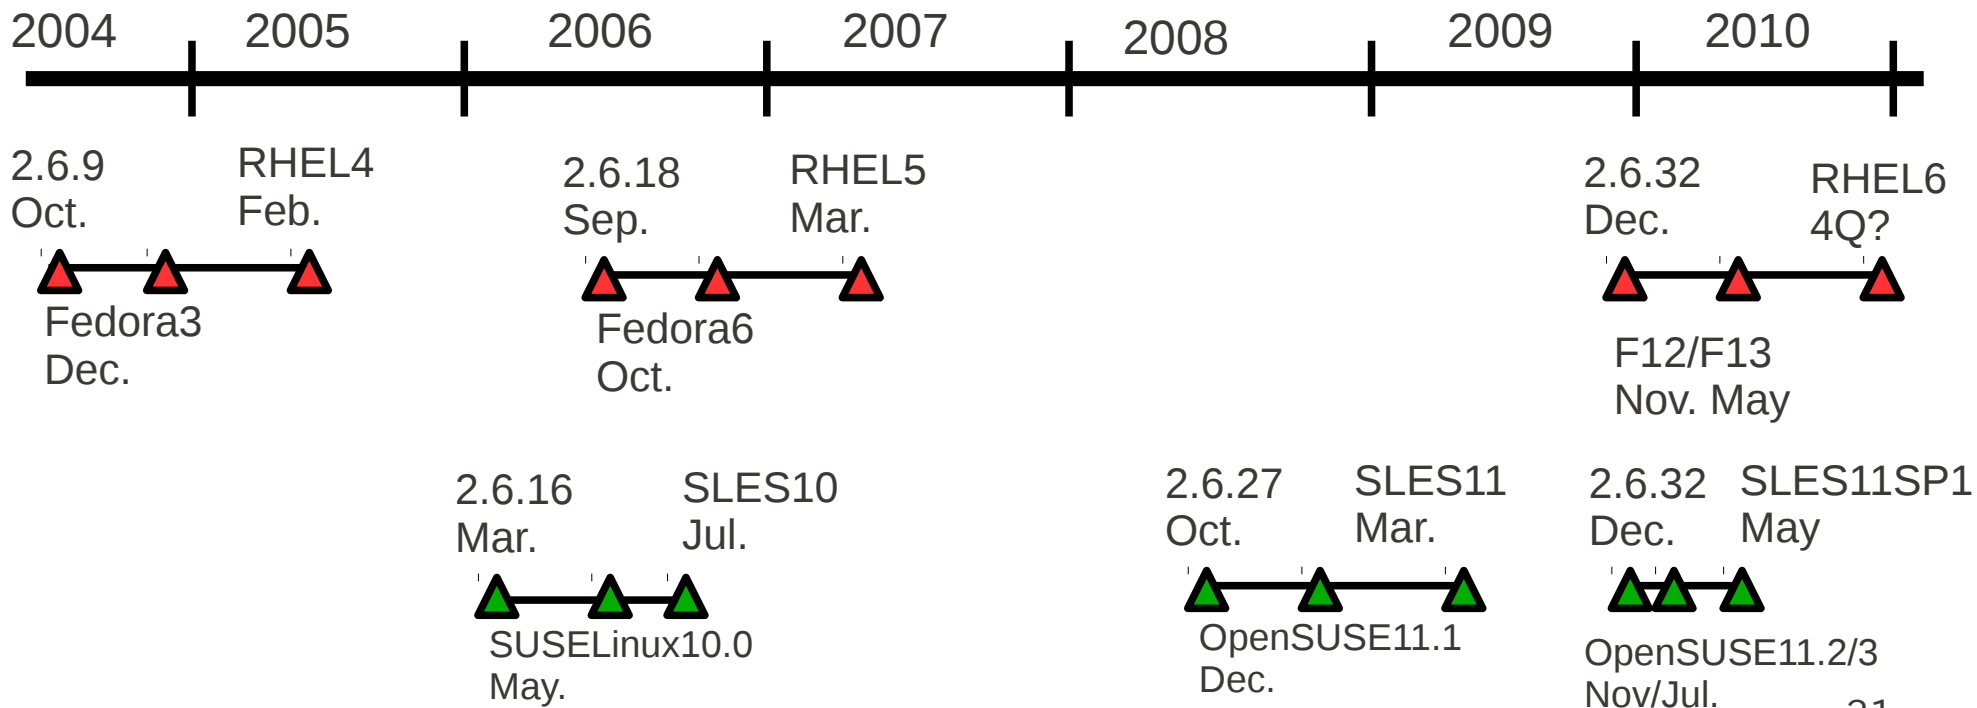
Summary

- 7 years of stable 2.6 kernel development process since 2004
- 1200 of developers send patches in each releases within about 80 days
- Yearly 1.2-2.4MLines, 2000-5000 Files added
- Yearly 2000-2500 developers send patches
- Daily 4000Lines, 10 Files added
- Well coordinated development rules and process with great number of Maintainers

Status of Distributions

Linux upstream and distros

- Distros take 9-12 months from upstream Kernel release to provide commercial products for it's integration and testing.
- New versions will be released every 3-4 years.



Android and Linux kernel

- Android is using relatively new kernel
- Took 5-7 month after kernel released



Android version	Code name	Released date	Kernel version	Kernel release date
1.1		9, Feb. 2009		
1.5	Cupcake	30, Apr. 2009	2.6.27	9.Oct.2008
1.6	Donut	15, Sep. 2009	2.6.29	23, Mar 2009
2.0, 2.1	Eclair	26, Oct. 2009	2.6.29	23, Mar 2009
2.2	Froyo	20, May 2010	2.6.32	6, Jan 2010
?	Ginger Bread	4Q, 2010	2.6.33 or 34	Feb, 2010 May, 2010

MeeGo and Linux kernel

- MeeGo is also taking latest kernel
- 2-3 month after kernel released

Version	Released date	Kernel version	Kernel release date
1.0	26, May 2010	2.6.33	24, Feb. 2010
1.01	July 2010	2.6.33.5	
1.1	Oct. 2010	2.6.35	1, Aug 2010



Summary

- Enterprise Distros, community and users are well connected as the eco system
 - Long term stable release and long life supported products makes stable infrastructure of IT system
- Linux is becoming the fundamental part for phones and any of embedded space with distributions such as Android and MeeGo
 - Latest kernel will help further innovation

Conclusion

- Now, Linux is indispensable to create innovative products, social infrastructure and global internet services
- Linux is enough large and rapidly growing. So, single entity cannot be able to compete or own the Linux community
- Some leading companies are trying to create the strong strategic model by using real openness
- What we will be able to do is ...

JOIN THE COMMUNITY,
JOINT THE PARTY OF SUPPORTING LINUX
AND
DO YOUR ACTIVITES IN THE COMMUNITY

END