

NanoBSD and the art of timekeeping

FOSDEM 2012



- Rudi van Drunen



- Edu: Electronics Design Engineer
- Senior Consultant & CTO Competa IT
 - Bridging the Gap between Business and Technology
- CTO XlexiT Technology B.V.
 - Wireless / Embedded / Networking
- Tech Guru Wireless Leiden
 - Largest wireless community network in NL

- Rudi van Drunen




- Edu: Electronics Design Engineer
- Senior Consultant & CTO Competa IT
 - Bridging the Gap between Business and Technology
 - YAY: We're hiring (ao. developers)
- CTO XlexiT Technology B.V.
 - Wireless / Embedded / Networking
- Tech Guru Wireless Leiden
 - Largest wireless community network in NL

Agenda

- The story: Motivation
- The case : A NTP appliance
- The hard work : NanoBSD
- The results : pretty graphs
- The Fun: A Demo ??


Time Nuts (1)



Tom van Baak



time-workshop



Cesium cell




Rubidium Clock

Time Nuts (2)



<http://www.leapsecond.com/great2005>



R.van.Drunen@competa.com

Slide 6

I want to be on time...

- So do my systems ...

- logging
- jobs

- syncing clocks

- local : networked
- global : networked

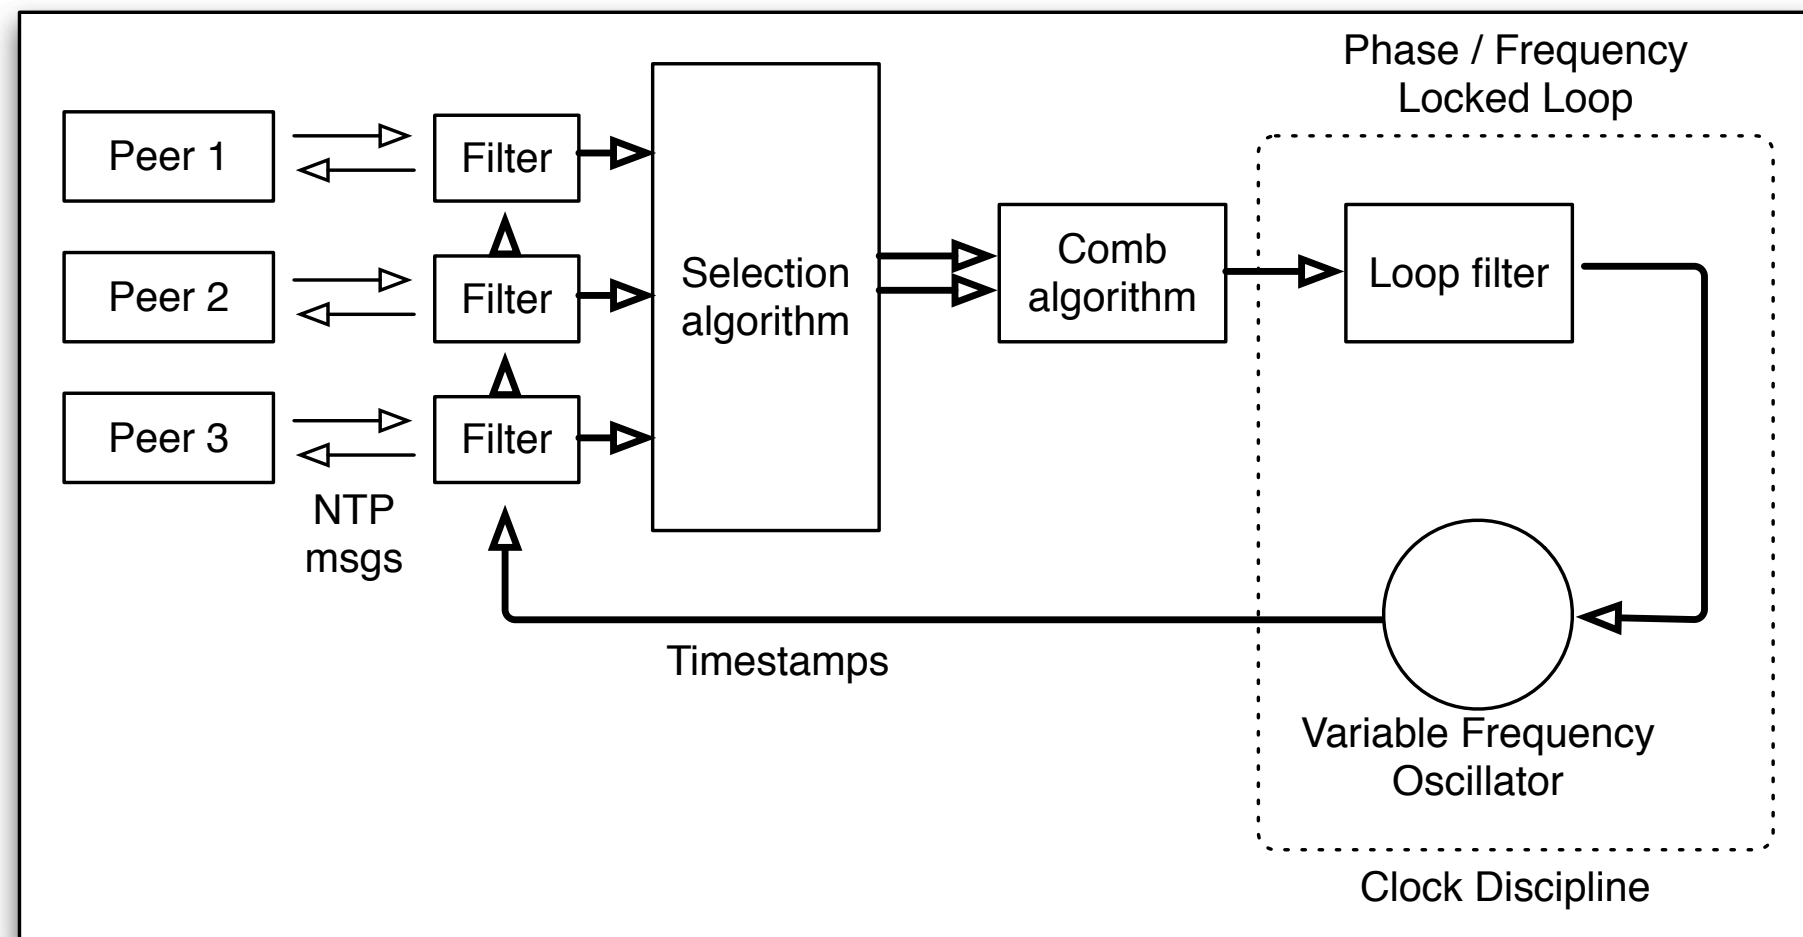


competa™

NTP

- NTP (Network time protocol)
 - accuracy : 10's of ms .. $< 1 \mu\text{s}$
 - 10..20 million servers / clients
 - runs on every OS :-)
 - now in version 4 (ntp4)
 - port 123 (tcp/udp)

NTP




why ?

- Better than using Internet NTP
- Fun and uses no internet :-)
 - logging (timestamps)
 - stock trading
 - radio / tv
 - datalogging / measurement / control
 - "trajectorycontrol"
 - gaming

Building a stratum 1

- Small embedded board
 - SC520 processor
 - High res timer
- GPS as time source
 - NMEA + PPS output
- Other time sources
 - DCF 77 receiver (Easy using a Conrad receiver on serial)
 - LORAN C receiver (Using SDR)

Soekris 4501 (1)



Soekris 4501 (2)



GPS 18 OEM

- Garmin hockey puck GPS
- GPS 18 **lvc**
 - 5V power supply
 - serial output
 - NMEA sentences 4800 bd
 - PPS line (TTL (5V) level)



Connection

- GPS
 - Serial port (Tx / Rx line of UART)
 - 4800 baud NMEA sequences
 - `$GPRMC,113425,A,5210.7572,N,00429.7643,E,000.1,330.3,101209,001.0,W*64`
 - DCD line of UART
 - pulse-per-second (start of UTC)
- Power
 - 5 Volts to GPS from board

Extra PPS accuracy

- Use high-res timer on SC520
- Kernel patch by PHK
 - `OPTIONS CPU_ELAN`
 - `OPTIONS CPU_SOEKRIS`
 - `OPTIONS CPU_ELAN_PPS`
- pps wired to GPIO-0 pin and to the TIMER1IN pin of the SC520

Software (1)

- Embedded distribution
 - nanobsd (see handbook, std. distro)
 - config kernel straightforward
 - PHK patches
 - config userland to include ntp
 - NTP needs drivers for PPS_API and NMEA (use ports collection)
 - define device links for /dev/gps0 /dev/pps
 - sysctl.conf:
 - `machdep.elan_gpio_config=-----P...-----`

Software (2)

- ntp.conf
 - Select right clock mode
 - Local NMEA clock (mode 20)
 - Local PPS discipline (mode 22)
 - Select fudge factor
 - delay between UTC second and rising edge pps (datasheet gps)
 - Set logging /statistics file locations
 - Set security and access from other hosts
 - Define (internet ?) peers

configs

```
timelord# more /etc/ntp.conf
#
server 127.127.20.0 prefer minpoll 4 maxpoll 4
fudge 127.127.20.0 time1 0.000015
server 127.127.22.0 minpoll 4 maxpoll 4
#
driftfile /etc/ntp/ntp.drift
statsdir /etc/ntp/
#
statistics clockstats
statistics rawstats
statistics loopstats
```

```
timelord# more /etc/devfs.conf
# Let NTP know where to find its clock
link      cuad1          gps0
link      elan-mmcr      pps0
```

BSD

- One (1) distribution
 - Clean version control
- Simple build
 - Kernel
 - Userland
- Experience !

competa™

NanoBSD

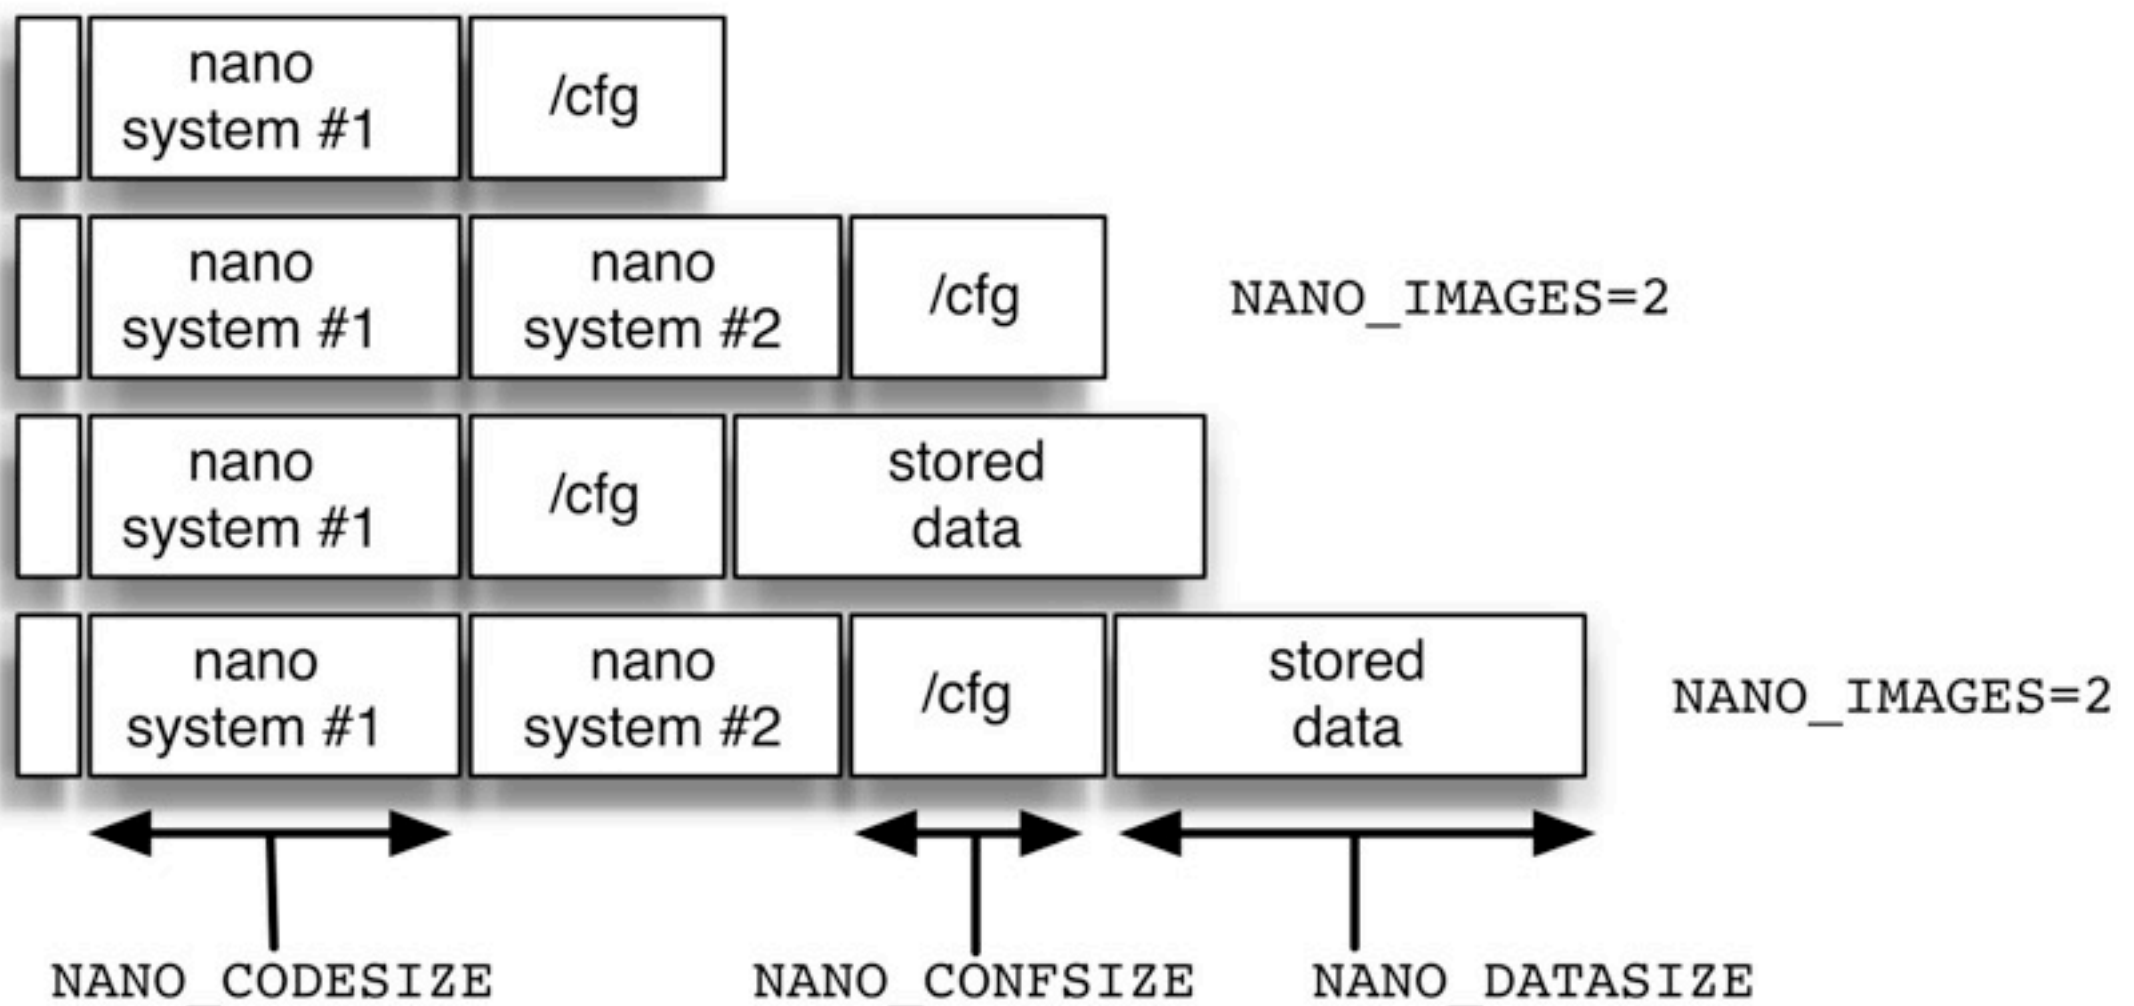
- A Script in the BSD (source) tree
 - `/usr/src/tools/tools/nanobsd/`
 - Builds a complete system
 - result: Flash image
 - Takes a config file
 - Takes packages to be included

Disk

- Flash
 - Reliability, NO MOVING PARTS
 - Power / heat
 - Wear
 - Flash: `-noatime -ro`
 - r/w : use memory filesystem
 - Enable write only when needed

nanoBSD (flash)Disk

boot0



nanobsd

- Usage: `$0 [-bikqv] [-c config_file]`
- `-b` suppress builds (both kernel and world)
- `-i` suppress disk image build
- `-k` suppress buildkernel
- `-n` add `-DNO_CLEAN` to `buildworld`, `buildkernel`, etc
- `-q` make output more quite
- `-v` make output more verbose
- `-w` suppress buildworld
- `-c` specify config file
- `-h` Display usage information.

config file

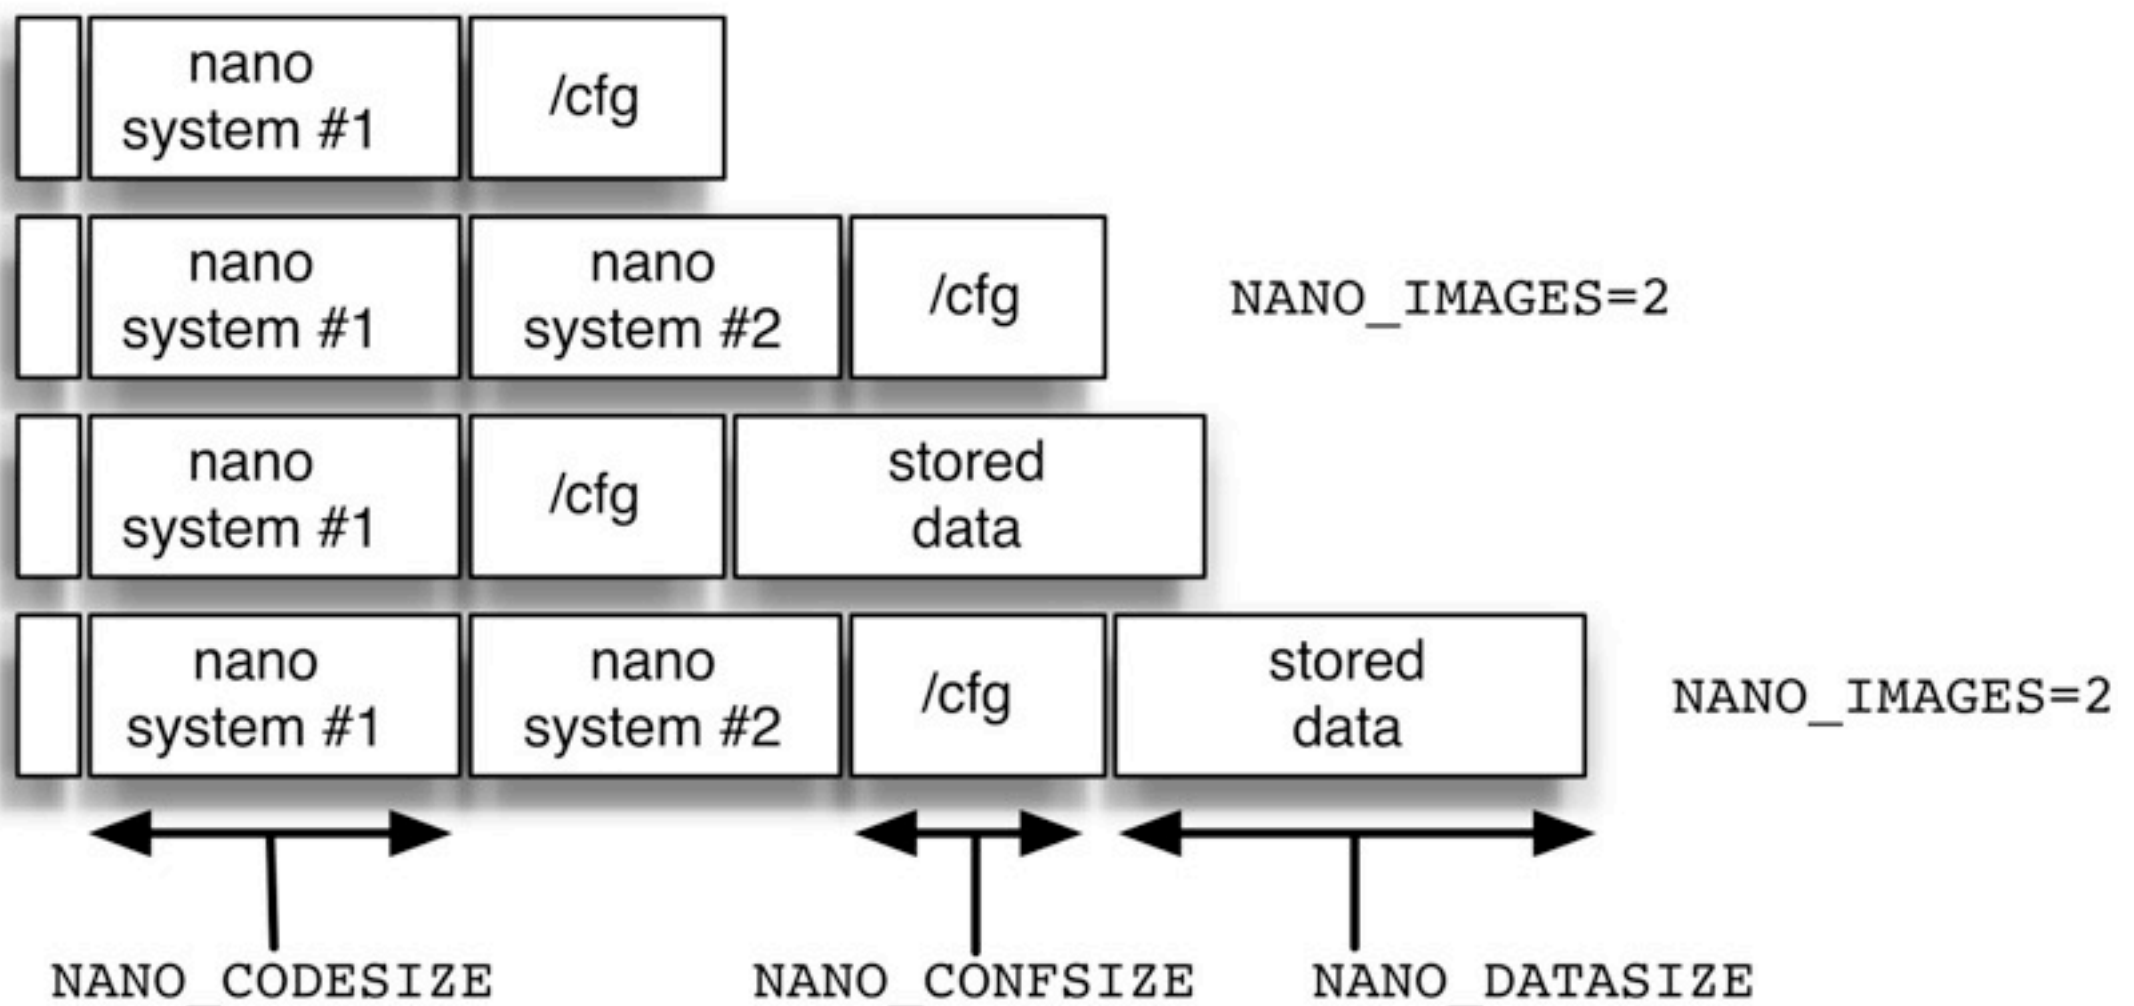
```
#  
NANO_NAME=myconf  
# options passed to buildworld  
# CONF_BUILD=  
# options passed to installworld  
# CONF_INSTALL=  
# options passed to buildworld / installworld  
CONF_WORLD='NO_MUPPET=YES'  
#  
NANO_KERNEL=MY_NANO_KERNEL  
FlashDevice Sandisk 512M  
#
```

Configuring Media

- Define Flashdevice
 - `FlashDevice <vendor> <ident>`
 - in file `FlashDevice.sub`
- Total sectorcount (`diskinfo(8)`)
 - `NANO_MEDIASIZE`
- Some bioses need explicit
 - `NANO_HEADS`
 - `NANO_SECTORS`

nanoBSD (flash)Disk

boot0



config partition

- Files for /etc
- r/w mounted during boot
- `vi /etc/ntp.conf`
- `mount /cfg; cp /etc/ntp.conf /cfg;`
`umount /cfg`

ramdisk

- /etc
 - default 5 Mbytes
 - `NANO_RAM_ETCSIZE=40960`
- /var
 - `NANO_RAM_TMPVARSIZE=40960`

customizing nano

```
#
cust_motd () (
    echo "My new MOTD" > ${NANO_WORLDDIR}/etc/motd
)
customize_cmd cust_motd
#
# no VGA
customize_cmd cust_comconsole
# ssh in as root
customize_cmd cust_ssh_root
# install from ../nanobsd/Files
customize_cmd cust_install_files
```

steps in building

```
clean_build  
make_conf_build  
build_world  
build_kernel  
clean_world  
make_conf_install  
install_world  
install_etc  
setup_nanobsd_etc  
install_kernel  
run_customize  
setup_nanobsd  
prune_usr  
run_late_customize  
create_${NANO_ARCH}_diskimage  
last_orders
```



competa™

Some admin scripts

- `change_passwd`
- `save_sshkeys`
- `Update`
 - `updatep1`
 - `updatep2`

competa™

New software Please !

- How update a once running nano ?
 - `ssh nanobox cat _disk.image.gz |
zcat | sh updatep1`

competa™

OK, Let's go !

- Define target Flash size & conf
- Build image
 - `sh nanobsd -C <nanoconfigfile>`
- Coffee !
- Write image to flashcard
 - `dd if=/usr/obj/nanobsd.full/_disk.full of=/dev/da0 bs=64k`

OK, back to reality

- Does it work ?!

competa™

Test

```
timelord# tail -f /var/log/messages

Dec 10 11:16:35 timelord kernel: Timecounter "ELAN" frequency 8333333 Hz quality 1000
.
Dec 10 11:16:35 timelord kernel: Timecounters tick every 6.666 msec
Dec 10 11:16:35 timelord kernel: Elan-mmcr driver: MMCR at 0xc59e4000. PPS support.
Dec 10 11:16:35 timelord kernel: Elan-mmcr Soekris net45xx comBIOS ver. 1.33 20080103 Copyright (C) 2000-2007
Dec 10 11:33:47 timelord ntpd[536]: time reset +957.271867 s
.
Dec 10 11:48:53 timelord ntpd[536]: time reset -1.005068 s
Dec 10 11:48:53 timelord ntpd[536]: kernel time sync status change 2001

timelord# cat /dev/gps0
$GPRMC,131244,A,5210.7630,N,00429.7604,E,000.2,166.5,101209,001.0,W*67
$GPRMC,131245,A,5210.7630,N,00429.7603,E,000.3,166.5,101209,001.0,W*60
$GPRMC,131246,A,5210.7629,N,00429.7602,E,000.4,166.5,101209,001.0,W*6D

timelord# ntpq -p
      remote           refid      st t when poll reach   delay    offset    jitter
=====
+GPS_NMEA(0)       .GPS.           0 1   1   16   377    0.000    0.003    0.015
oPPS(0)            .PPS.           0 1   4   16   377    0.000    0.002    0.015



timelord# ntptime
ntp_gettime() returns code 0 (OK)
  time cecb7092.2d69f480 Thu, Dec 10 2009 13:05:22.177, (.177398329),
  maximum error 1018 us, estimated error 15 us, TAI offset 0
ntp_adjtime() returns code 0 (OK)
  modes 0x0 (),
  offset 1.623 us, frequency -6.086 ppm, interval 1 s,
  maximum error 1018 us, estimated error 15 us,
  status 0x2001 (PLL,NANO),
  time constant 4, precision 0.001 us, tolerance 496 ppm,
```

Results

- Stratum1 NTP server
 - EUR 300
 - max error < 5 μ s
 - a weekend of fun !

competa™

Accuracy / Offset




loopstats file:

```
55175 73274.572 0.000037798 -6.205 0.000015259 0.000611 4
<date><time (s)><clock offset(s)><freq offset (ppm)><jitter (s)><wanderer (ppm)> <clock discipline>
```

R.van.Drunen@competa.com

Slide 38

pll correction



Compare this to the ambient temperature graph :-)

Additions

- Webinterface
- LCD + buttons
- Better crystal stability (lowering jitter)
 - Crystal Oven (10 MHz)
 - Rubidium clock from E-bay (ask me !)
 - TAPR clock module synthesizer
 - 10 MHz -> 33.333 MHz

Thanks to

- Competa IT
 - <http://www.competa.com>
- Paul Henning-kamp
- The FreeBSD Crew
- USENIX

competa™

Any Questions ...



resources

- time-nuts
 - <http://www.leapsecond.com>
- ;login: article
 - http://competa.com/downloads/Articles/USENIX_2009-08_a_home_built_NTP_appliance.pdf
- config files
 - <http://usenix.org/publications/login/2009-08/index.html>