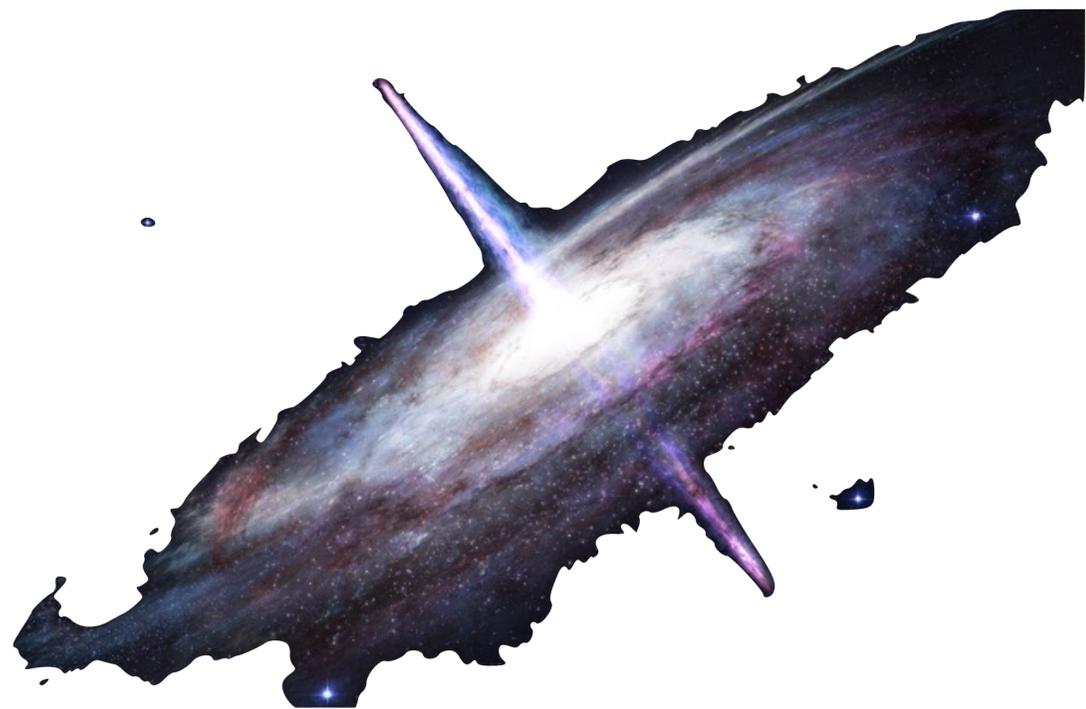
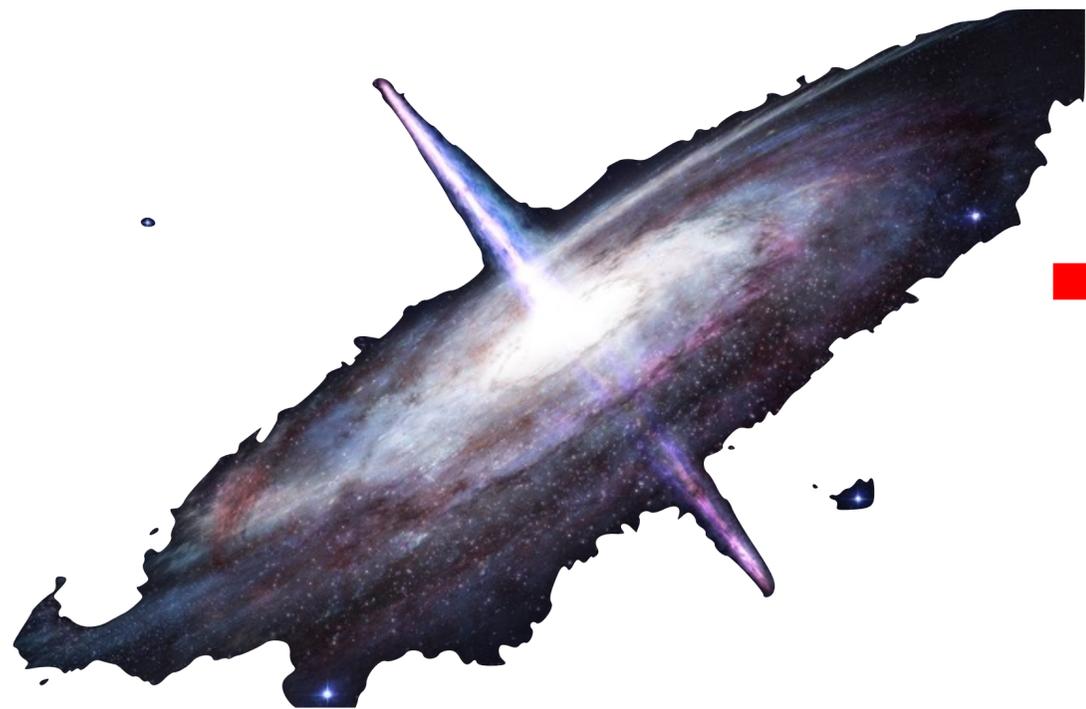
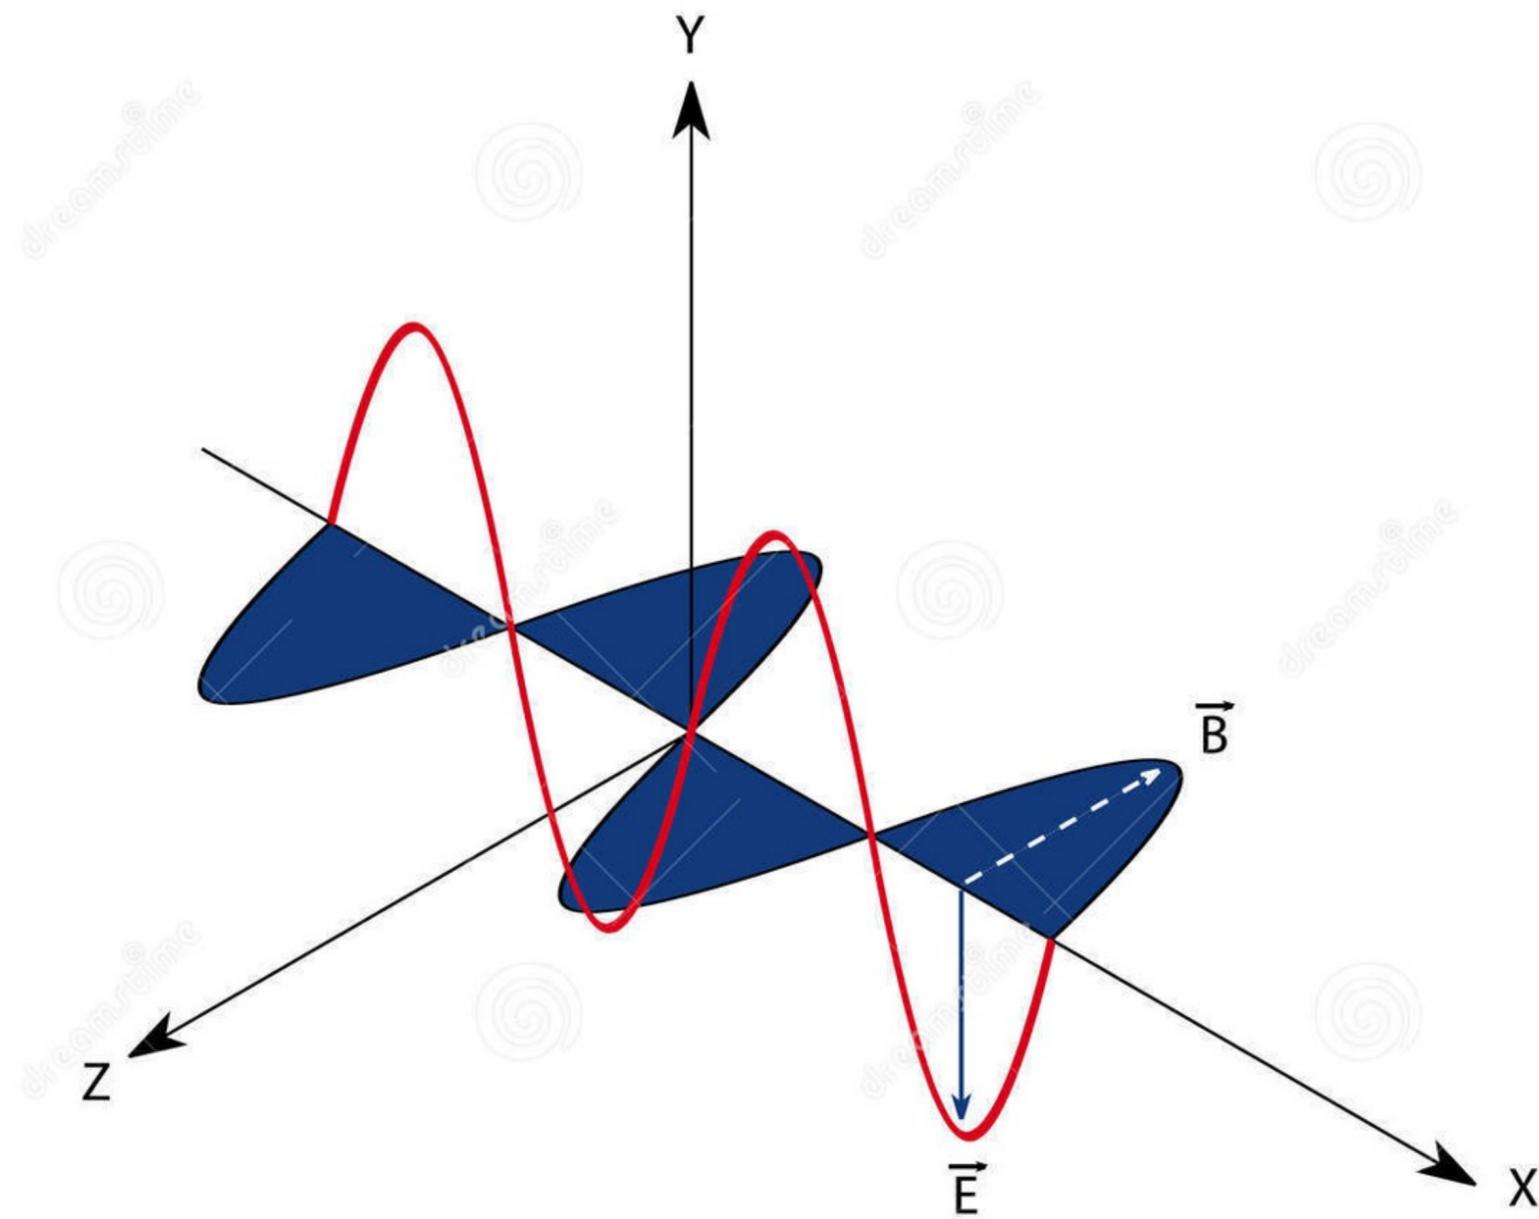
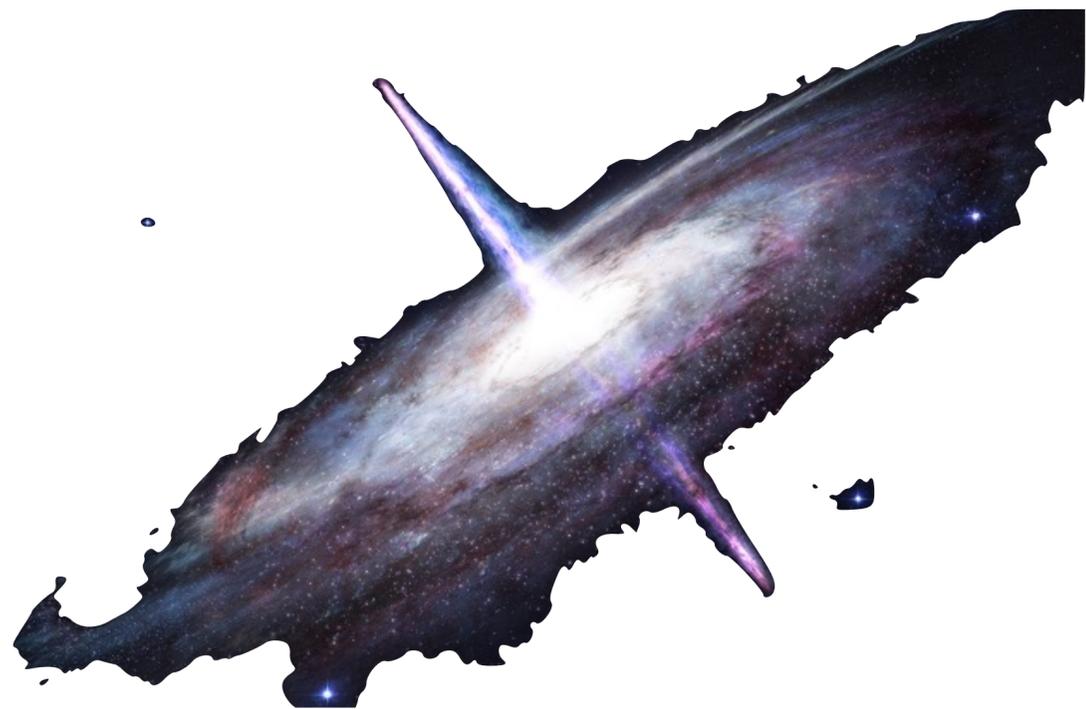
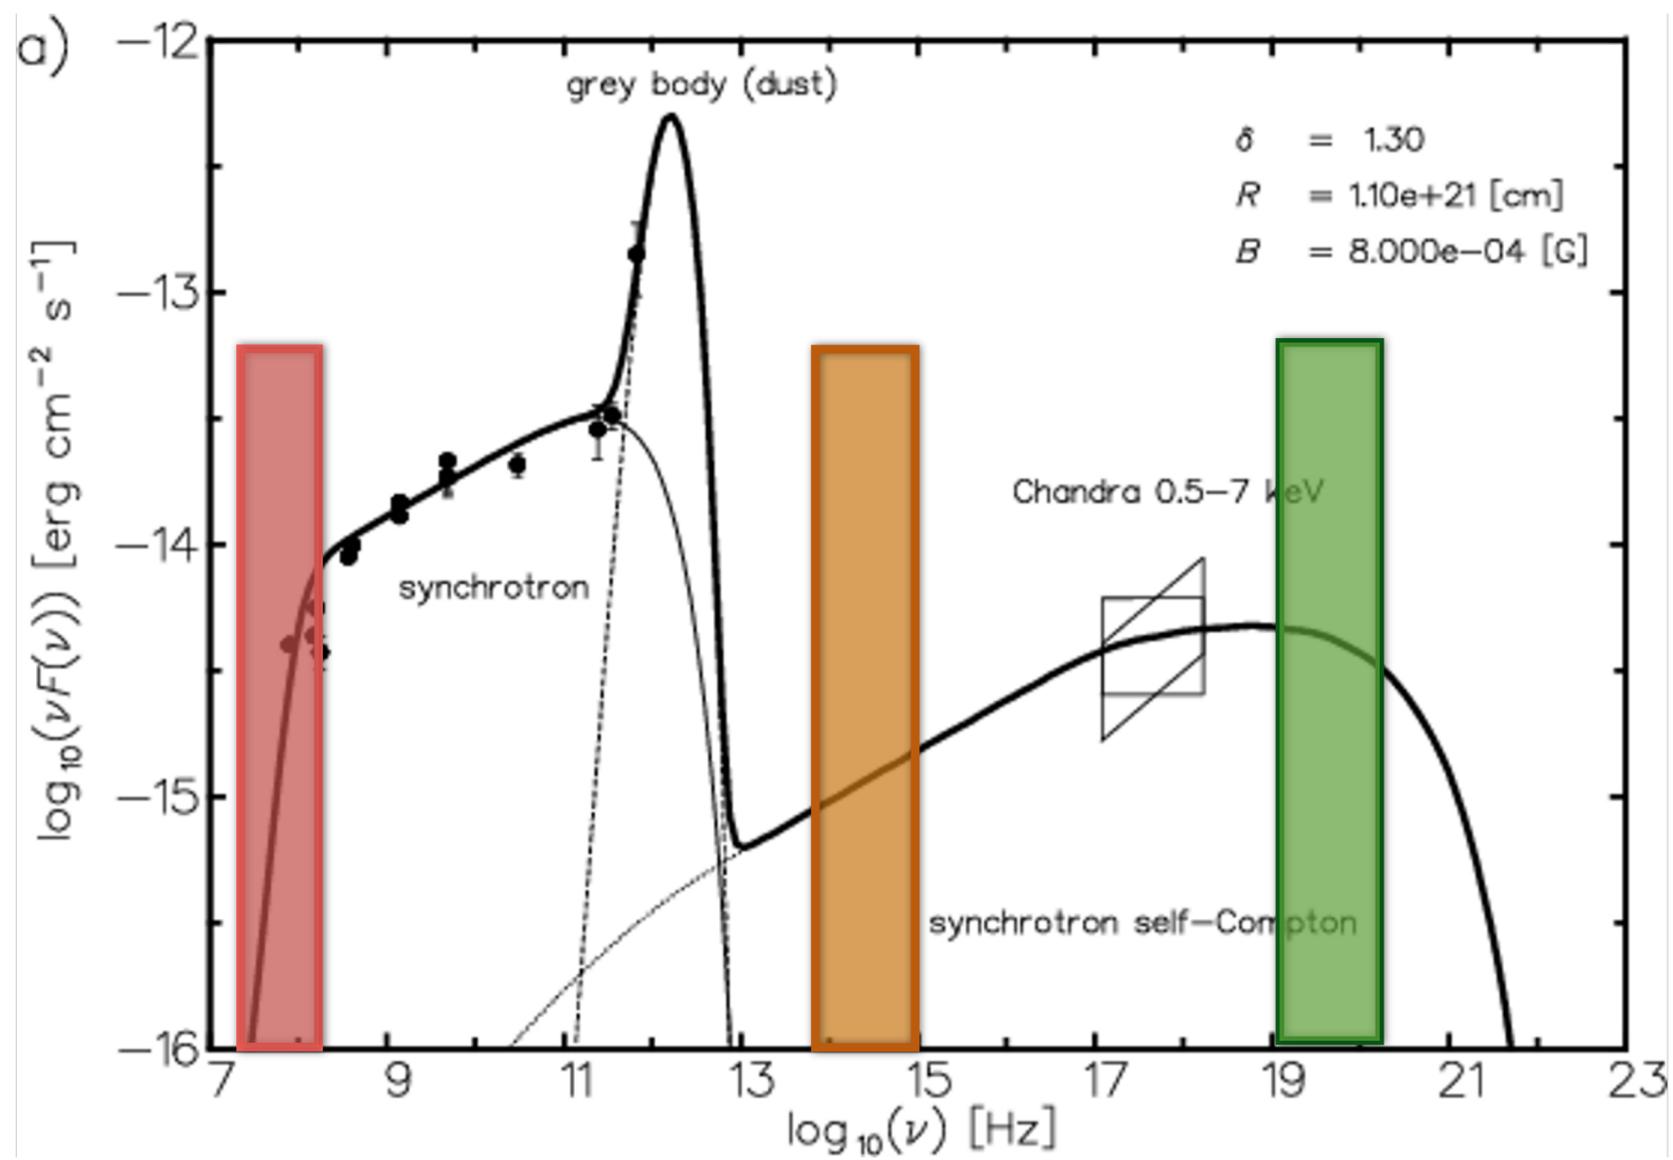
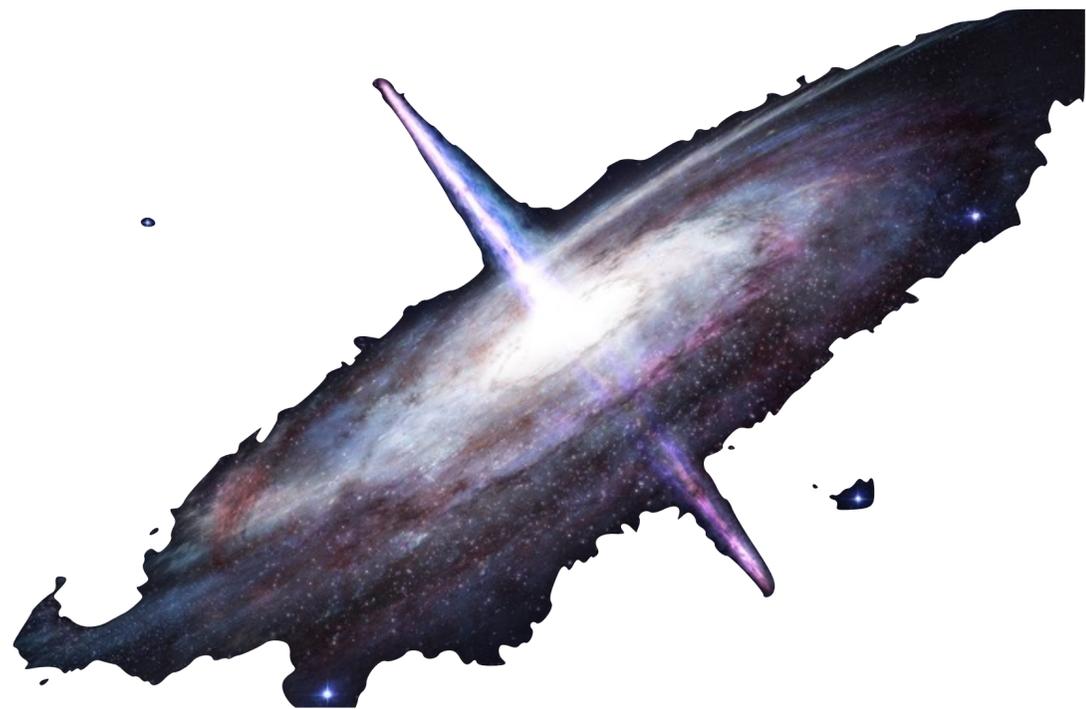


$$\sigma_T = \frac{T_{sys}}{\sqrt{\Delta_\nu \cdot \Delta t}}$$





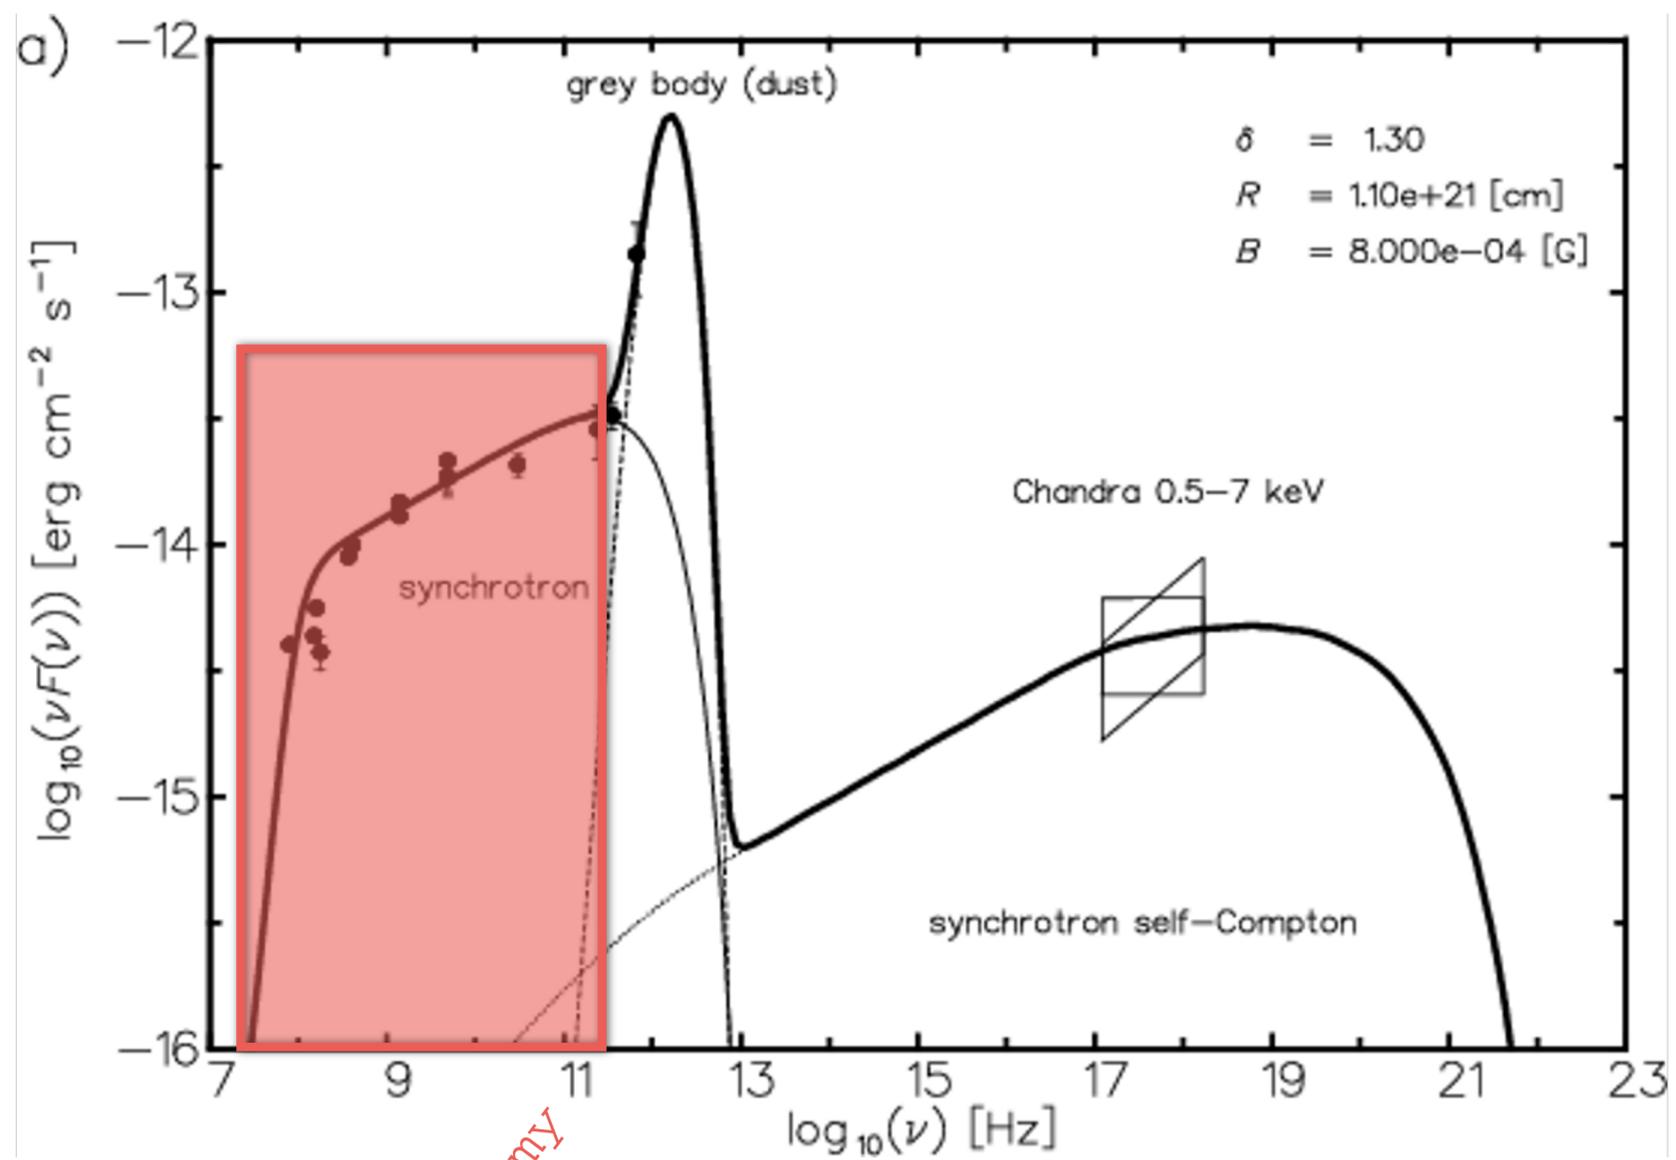
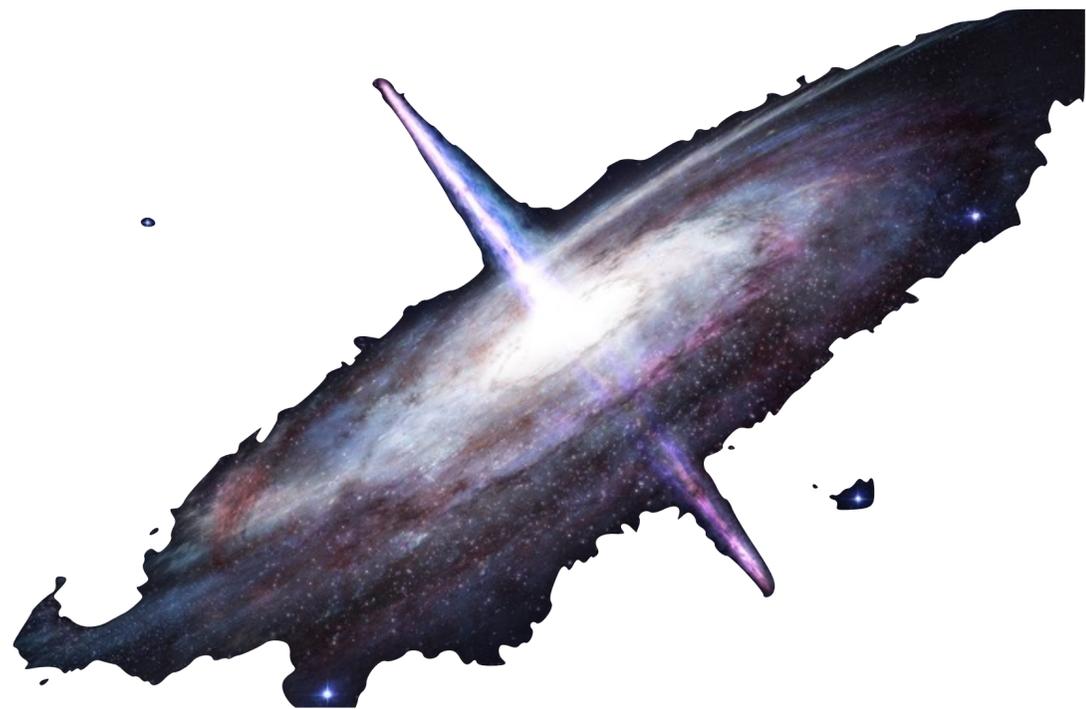




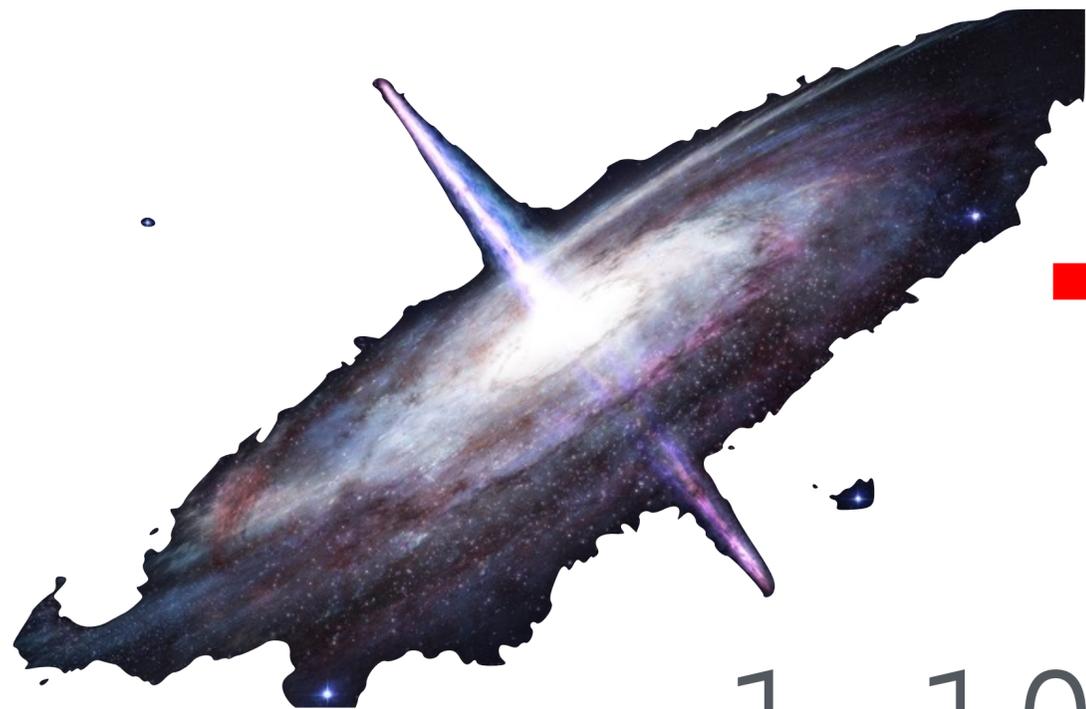
FM radio  
television

optical

gamma

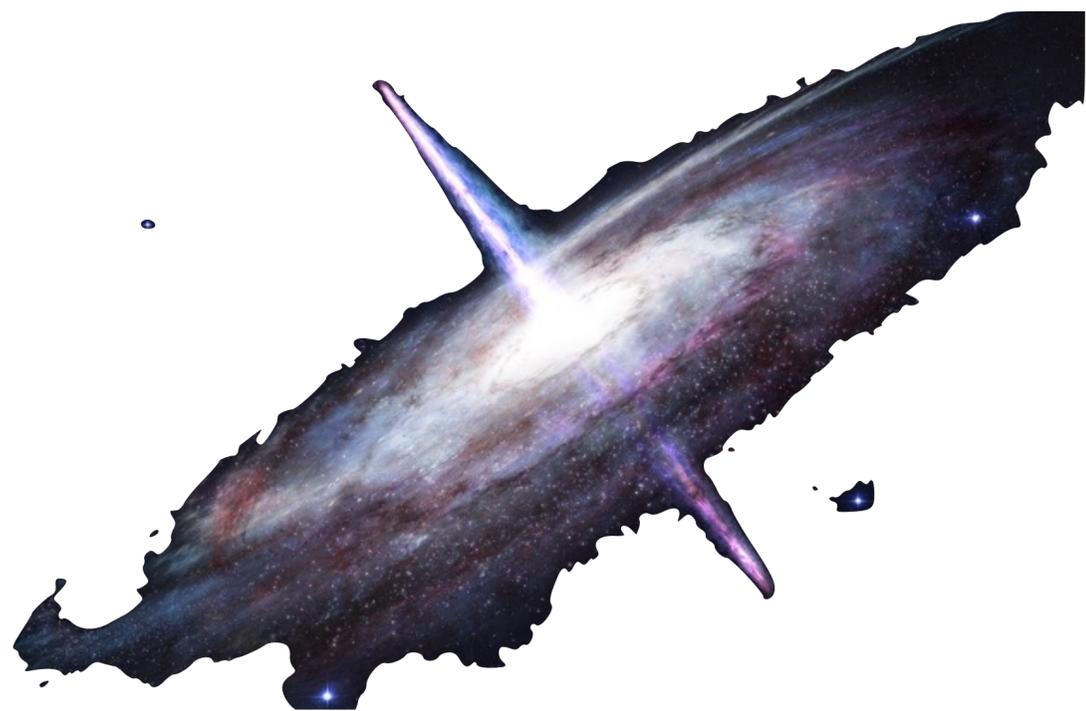


radio astronomy



1-100 GHz RF signals

Undesirable to sample/digitize/process as one chunk!



analog



digital

# VLBI Data:

acquisition, formats, transport and tools

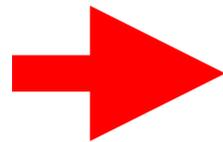
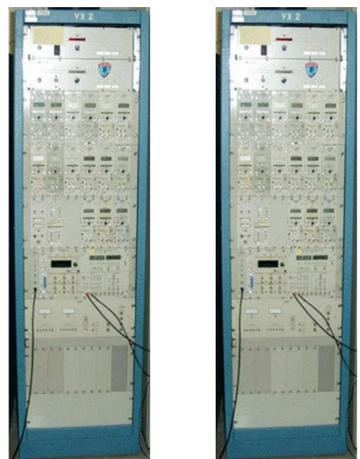
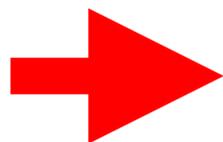
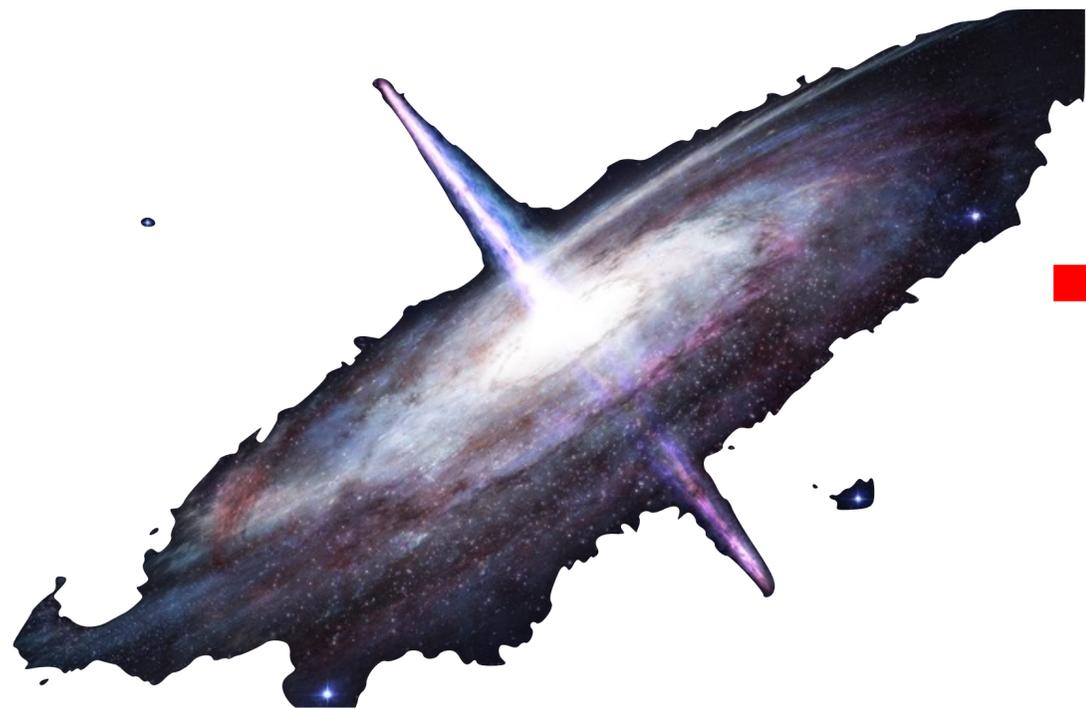


**JIVE**

Joint Institute for VLBI

ERIC

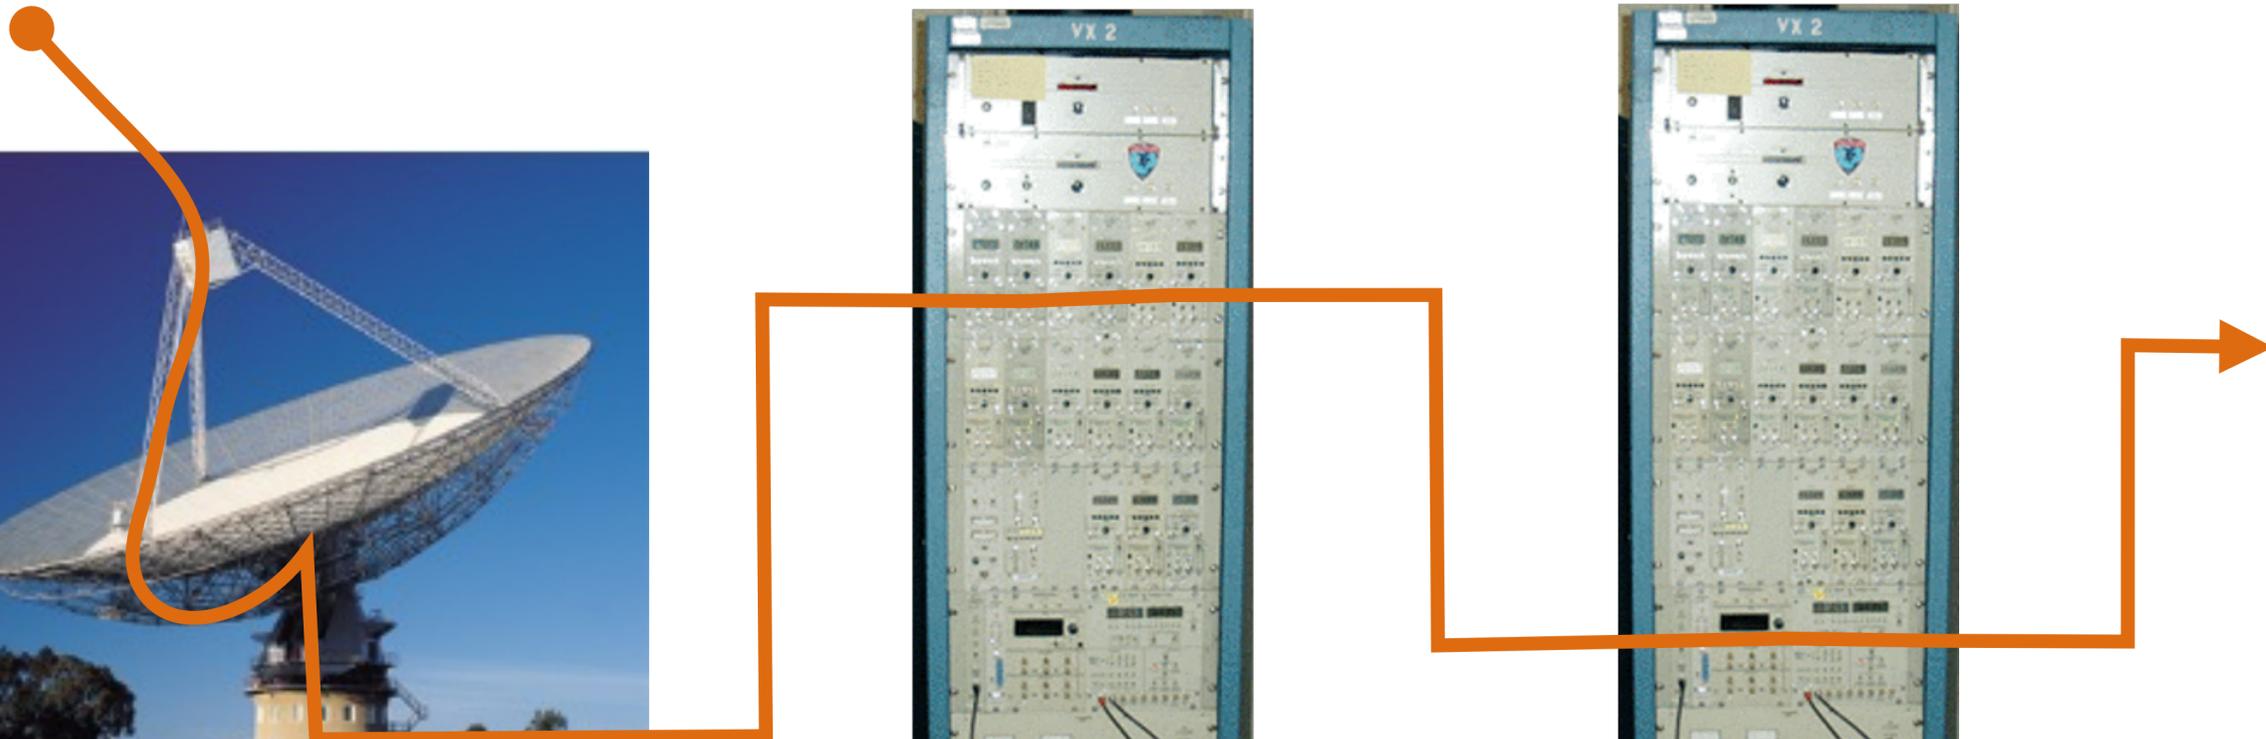
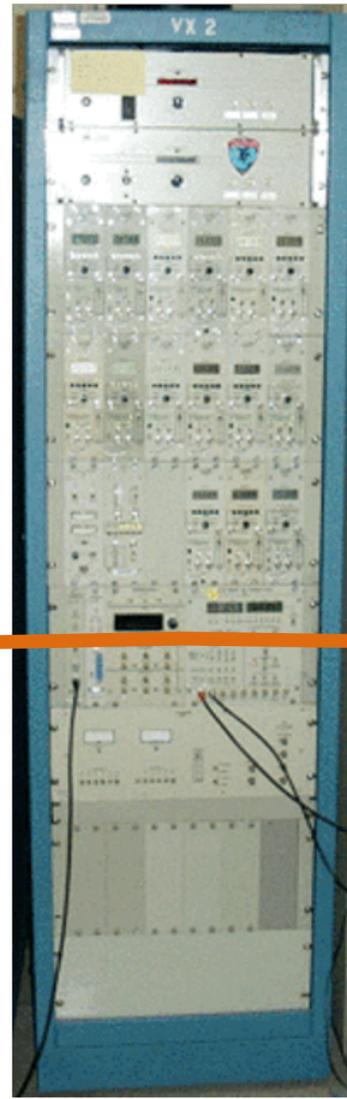
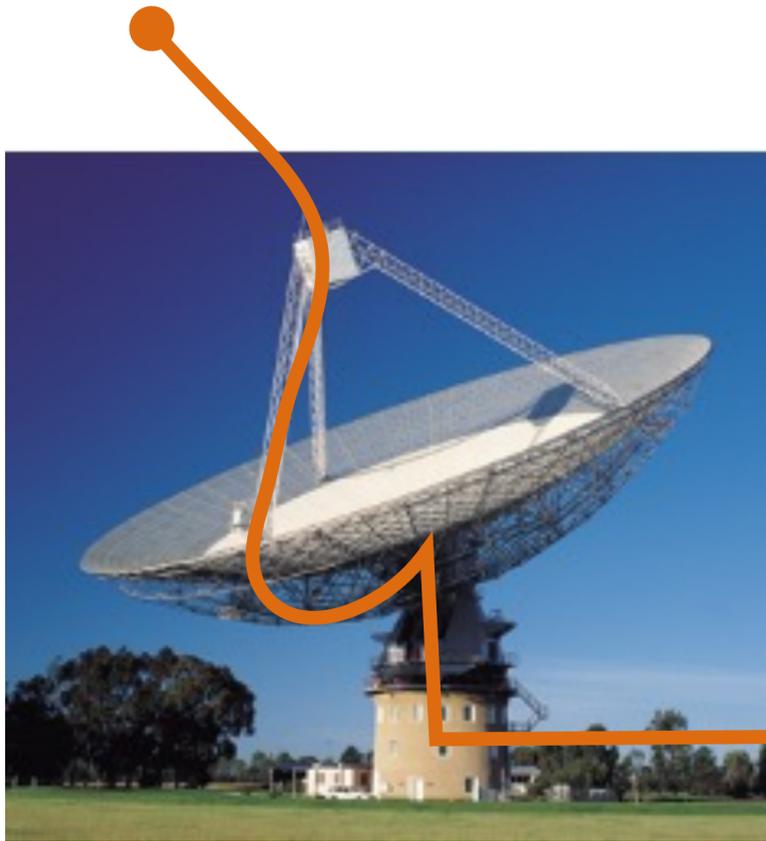
Marjolein Verkouter



analog

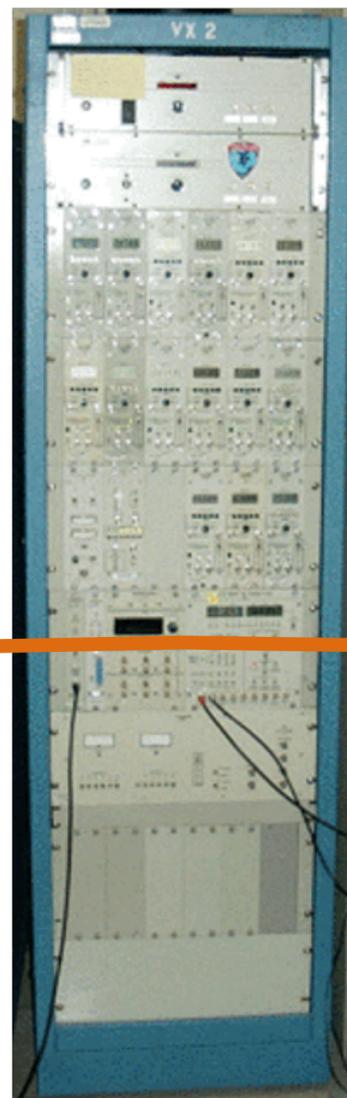
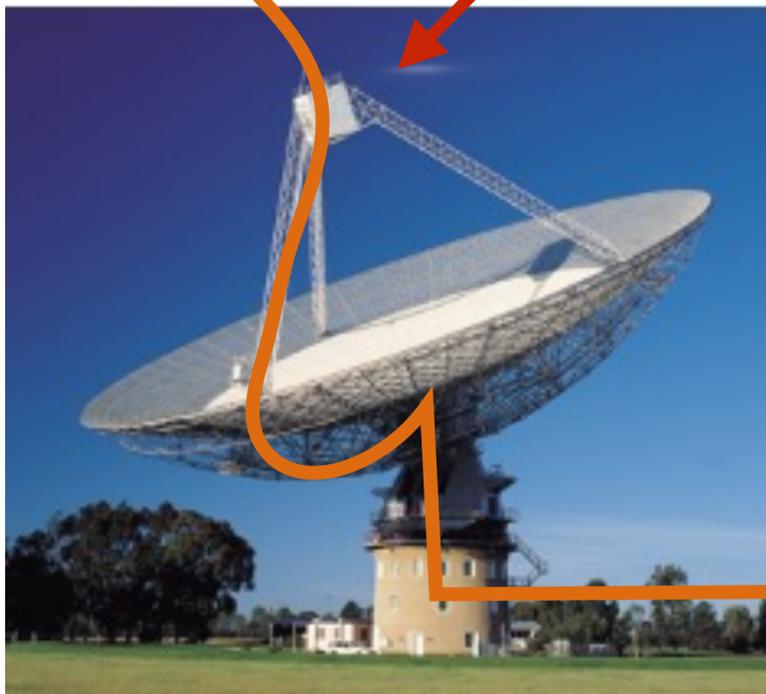
digital





sky freq  
1-100 GHz

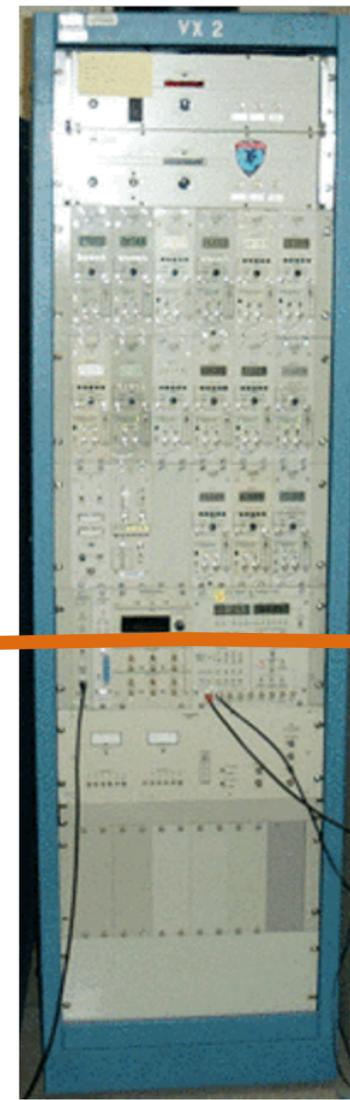
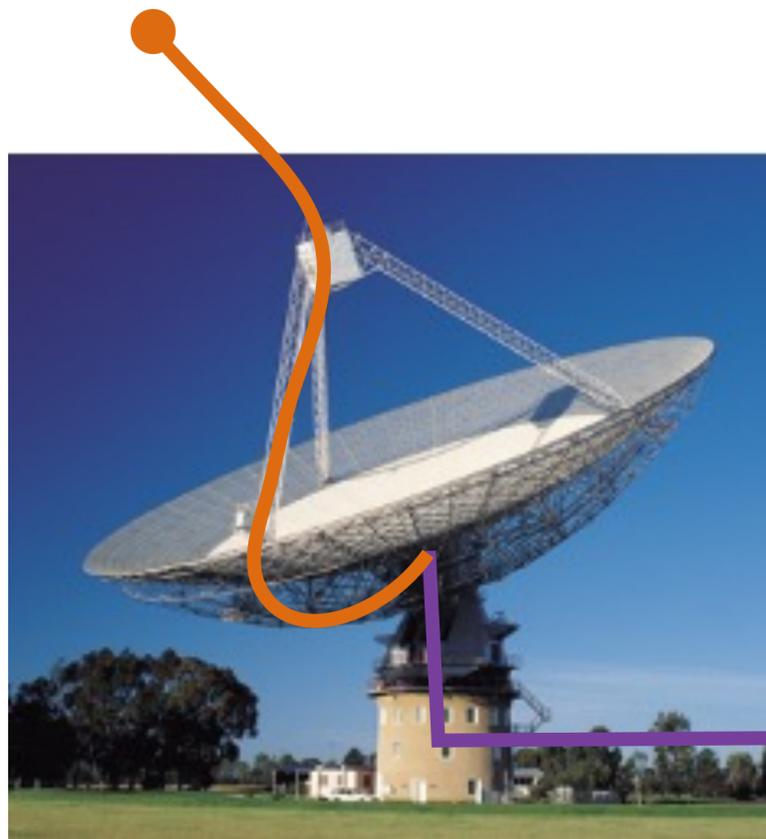
telescope front end



mix sky with LO

sky freq  
1-100 GHz

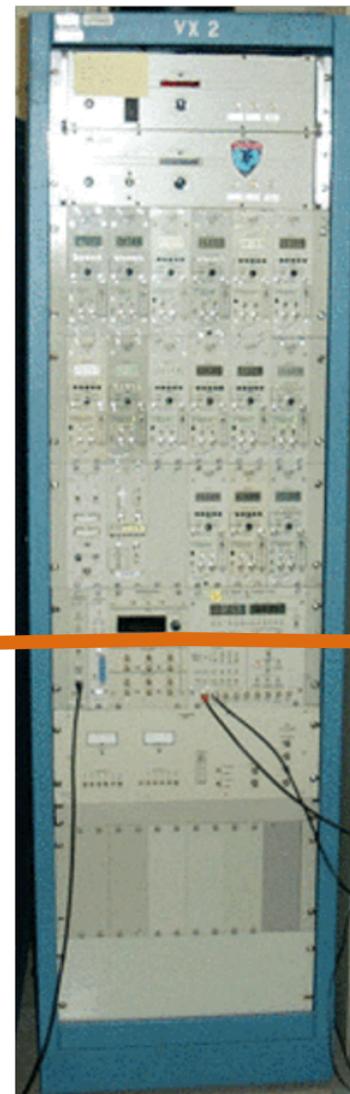
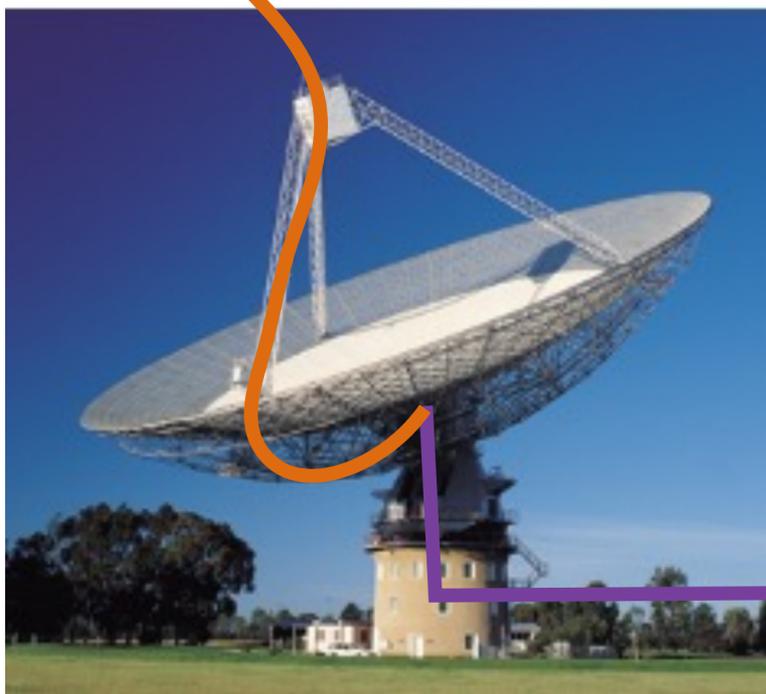
IF  
0-1 GHz



*mix sky with LO*

sky freq  
1-100 GHz

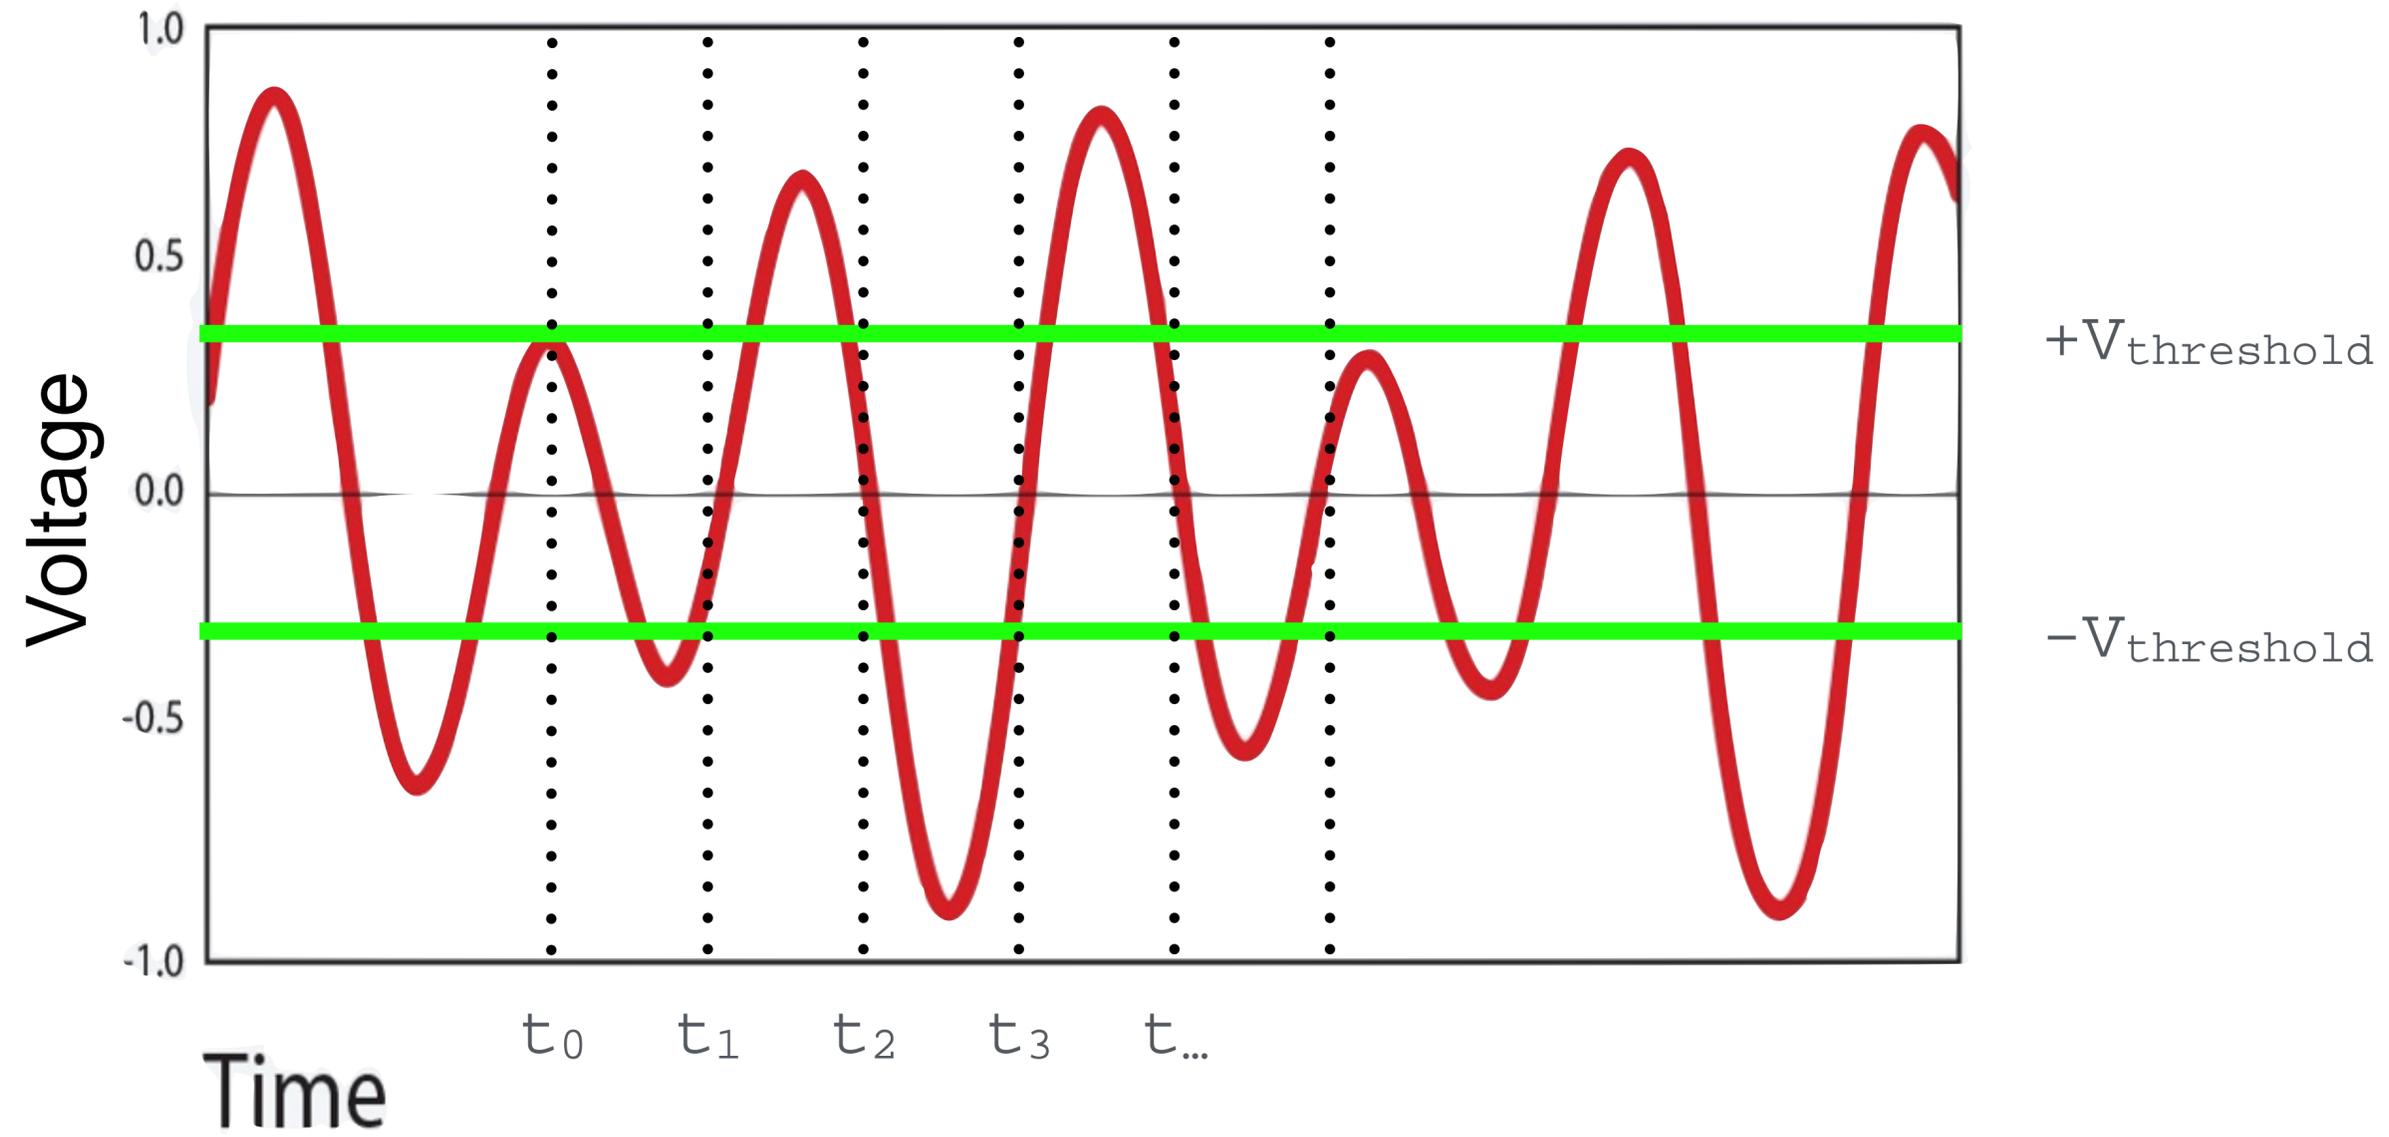
IF  
0-1 GHz

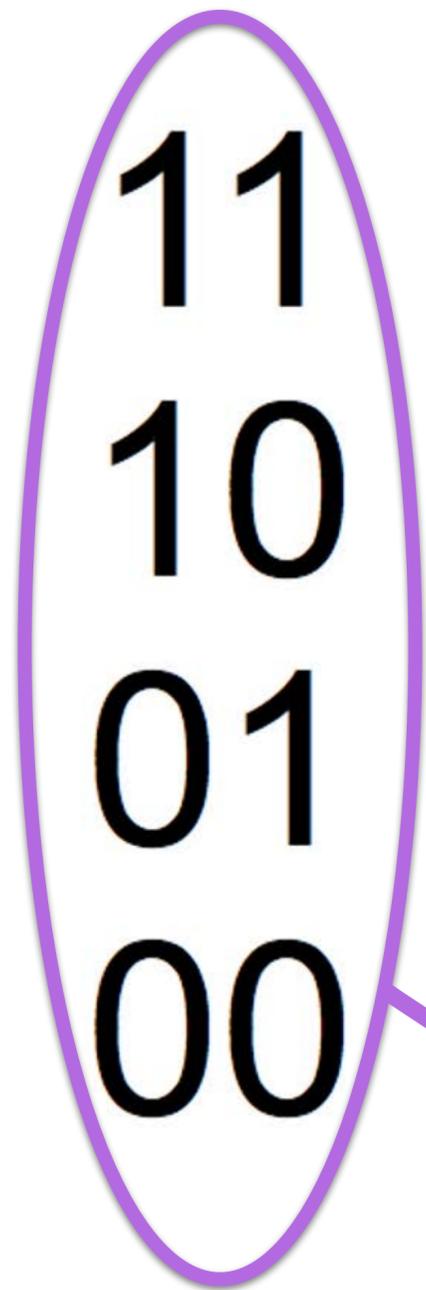
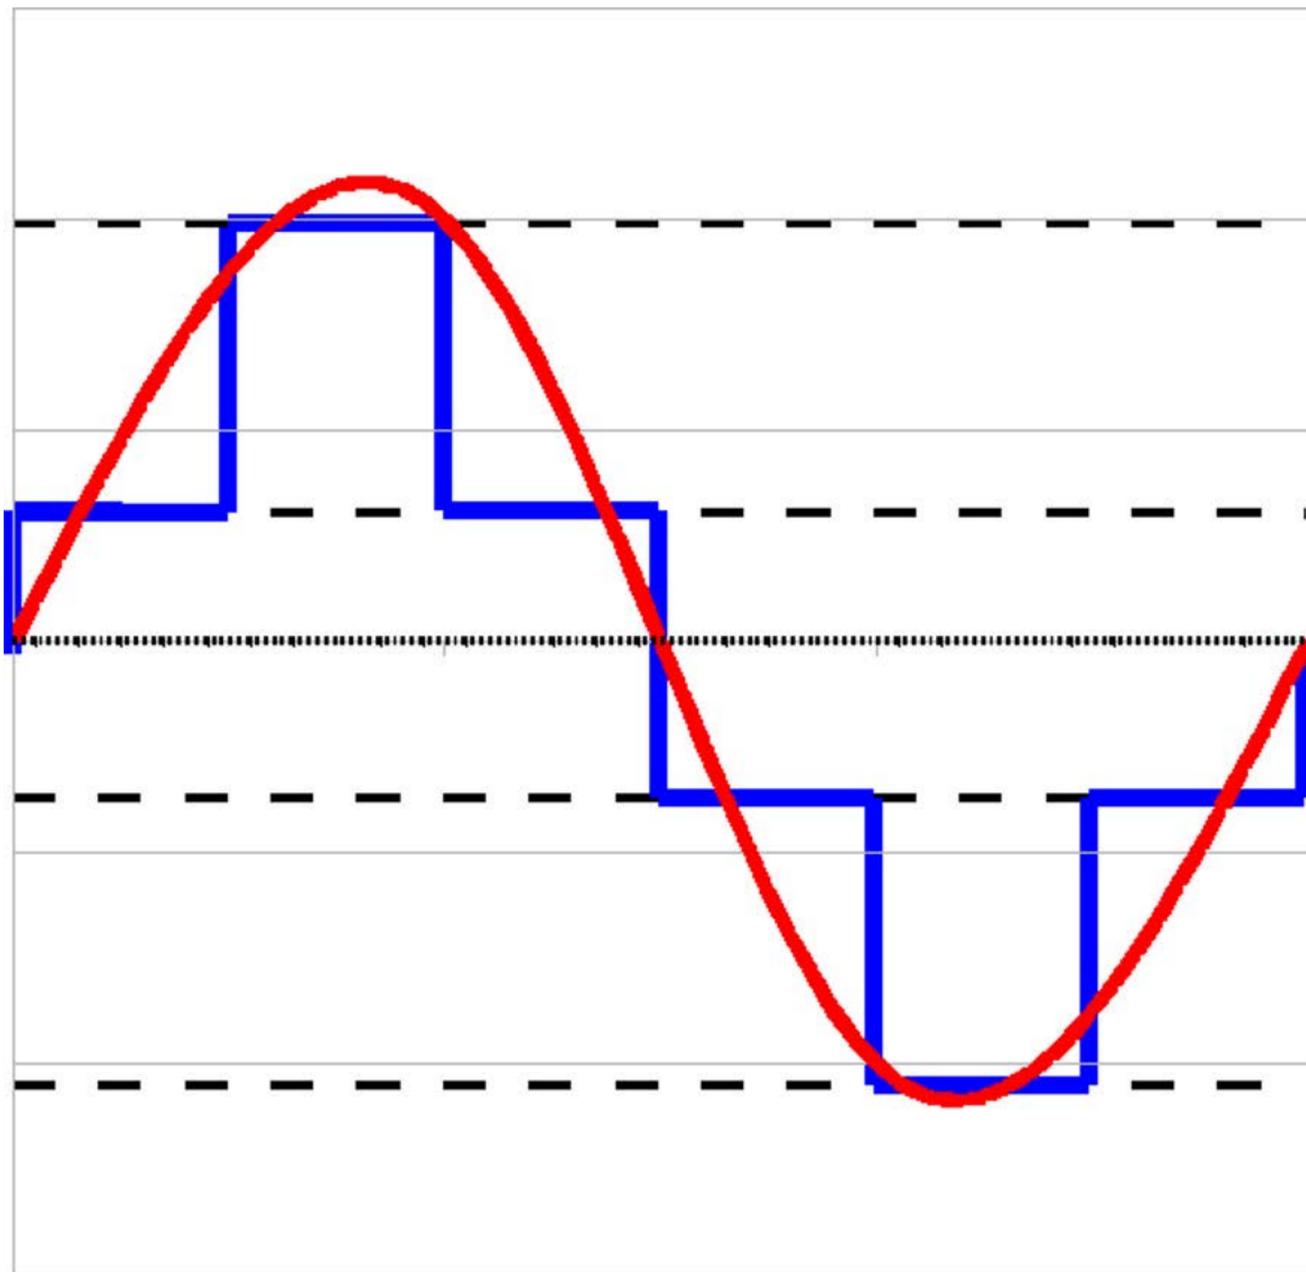


mix sky with LO

BBC, A-to-D  
conversion

# Analog Waveform

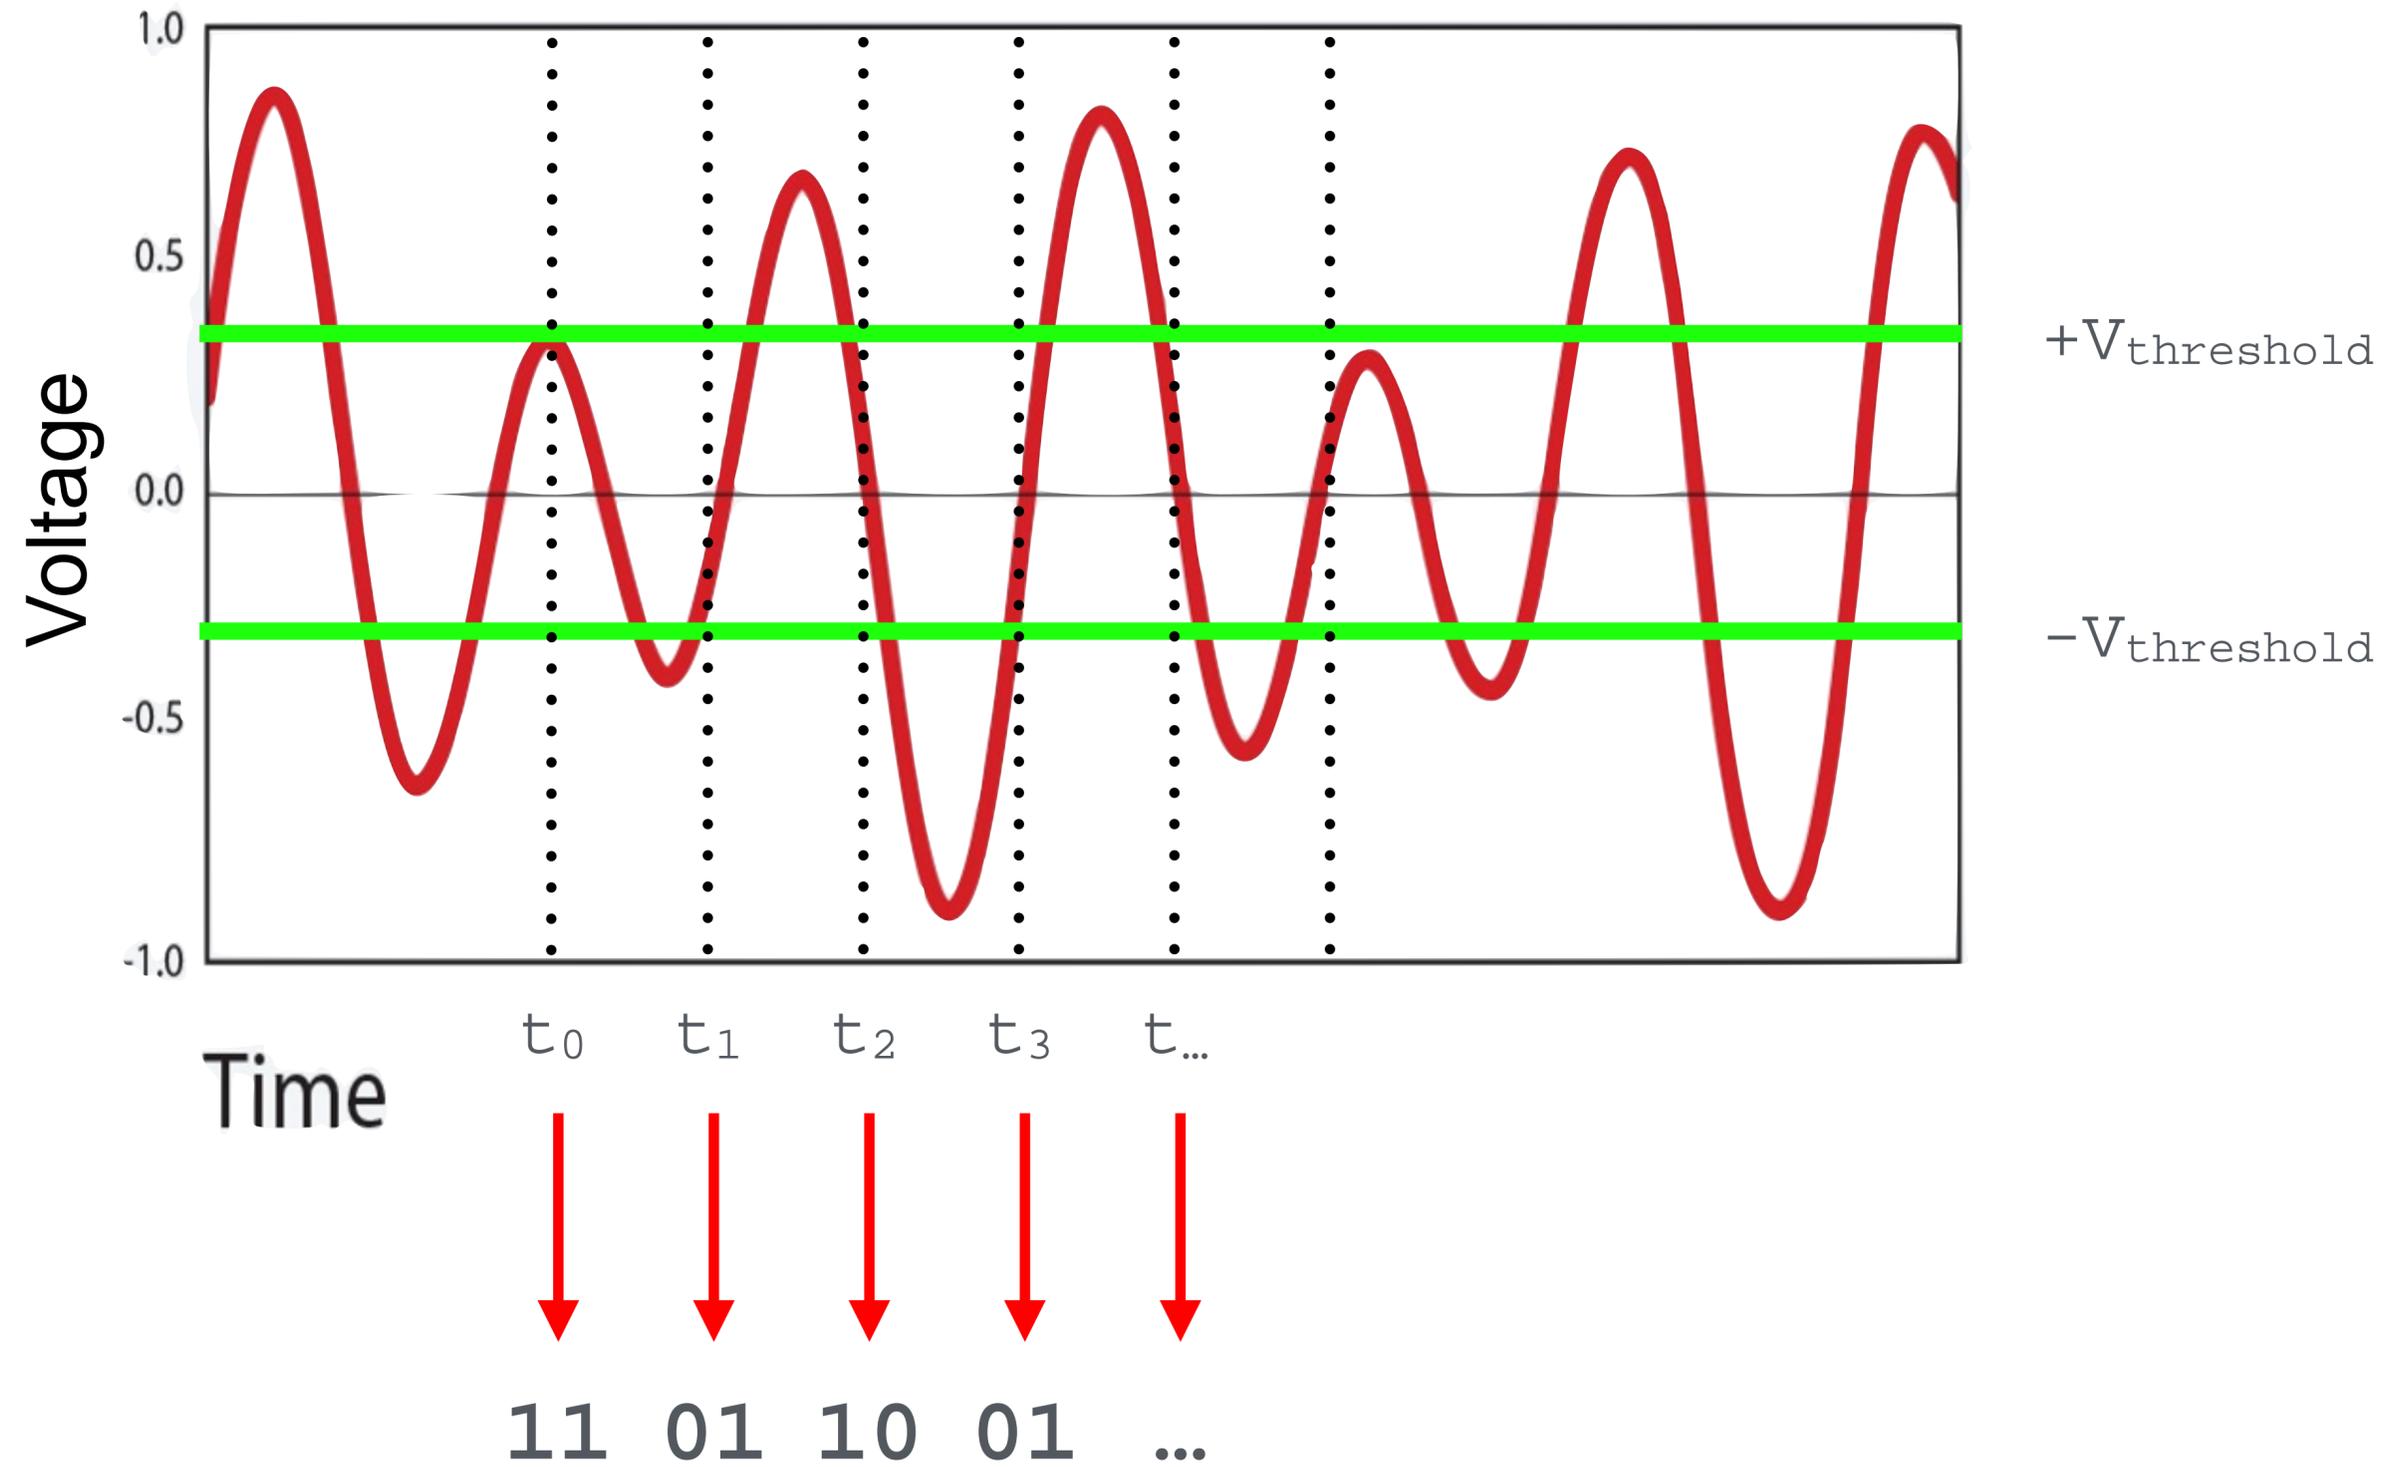




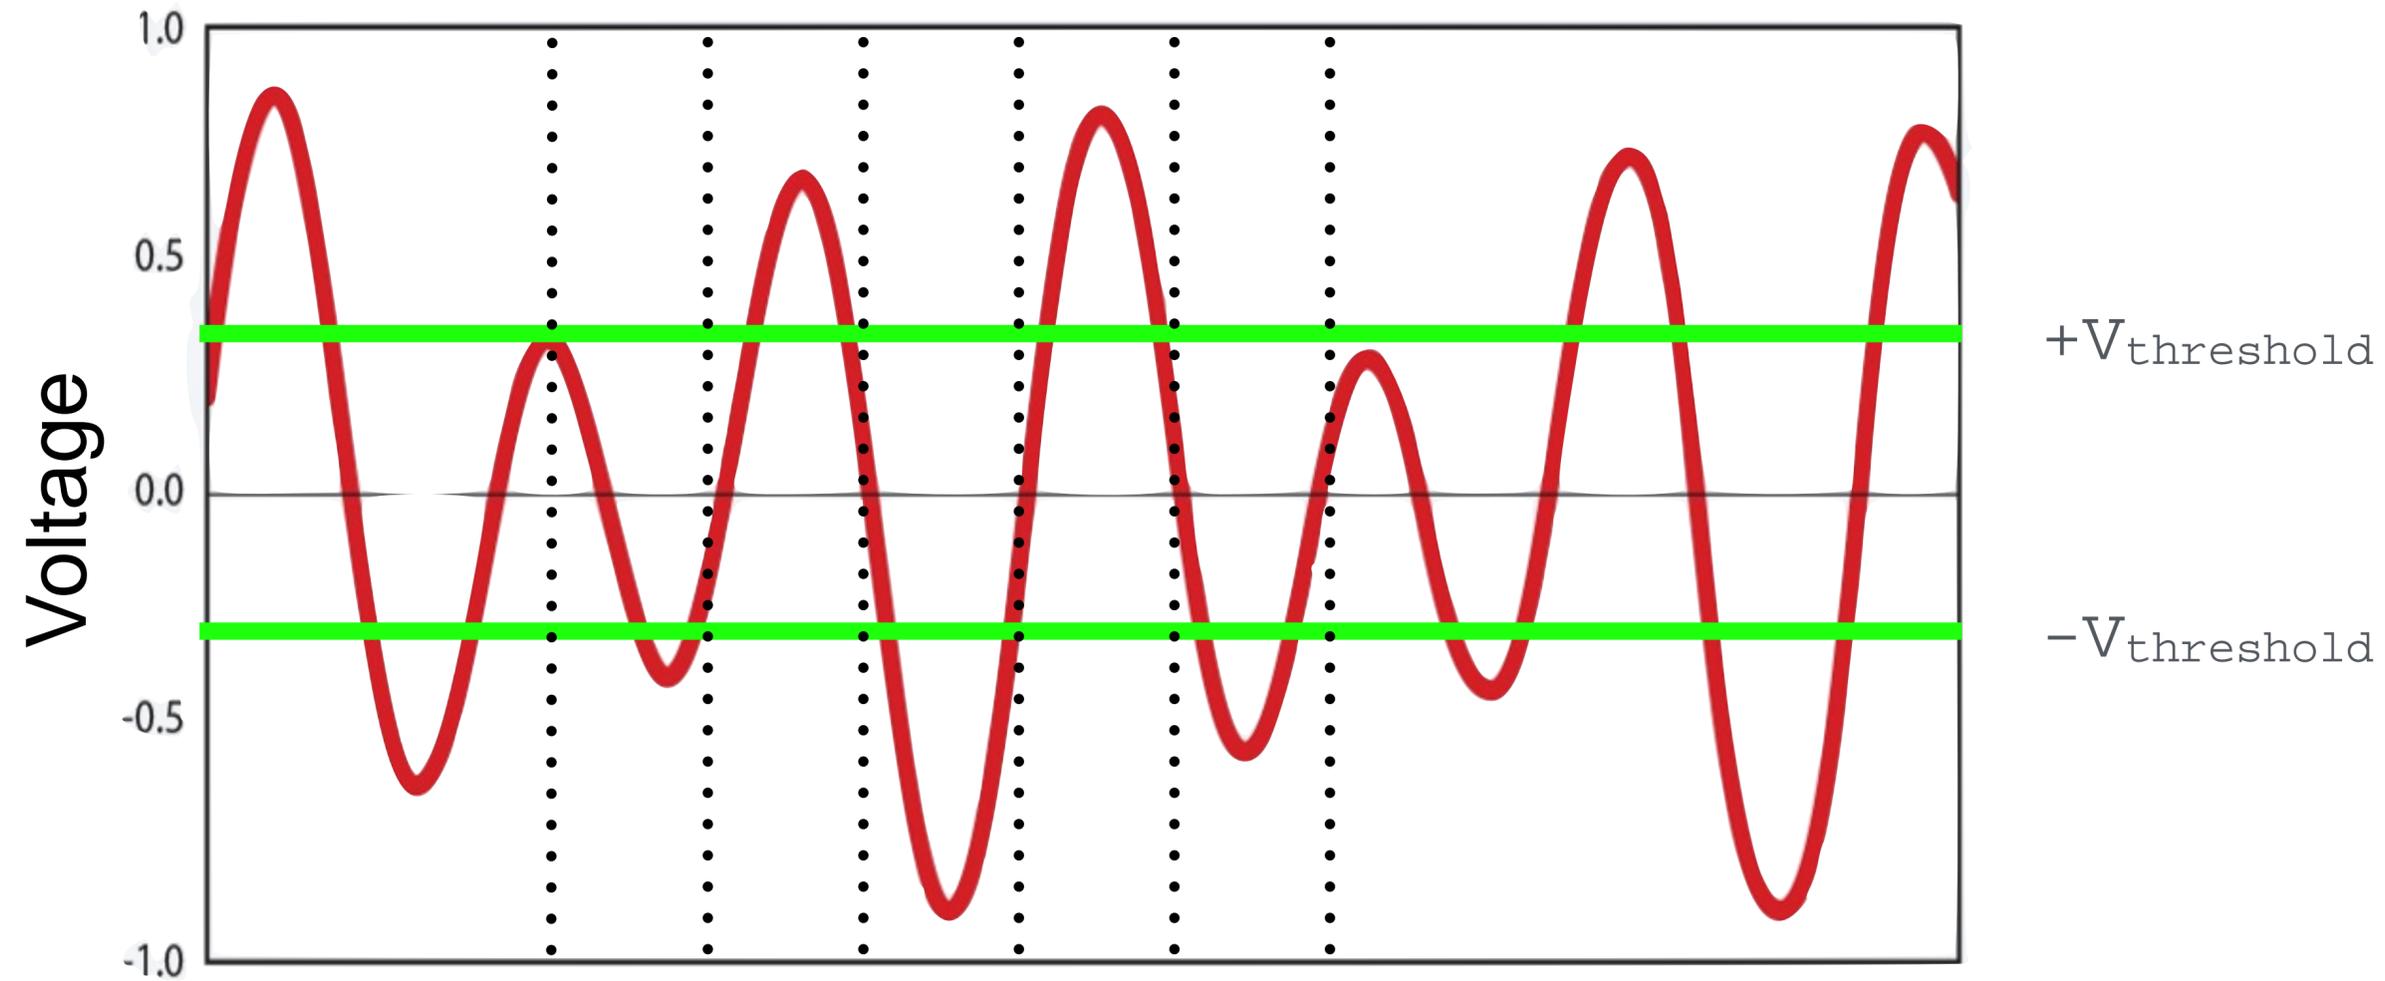
2-bit  
binary  
encoded  
value

- analog signal
- discretized signal

# Analog Waveform

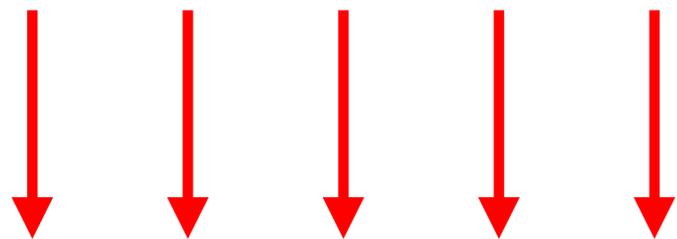


# Analog Waveform



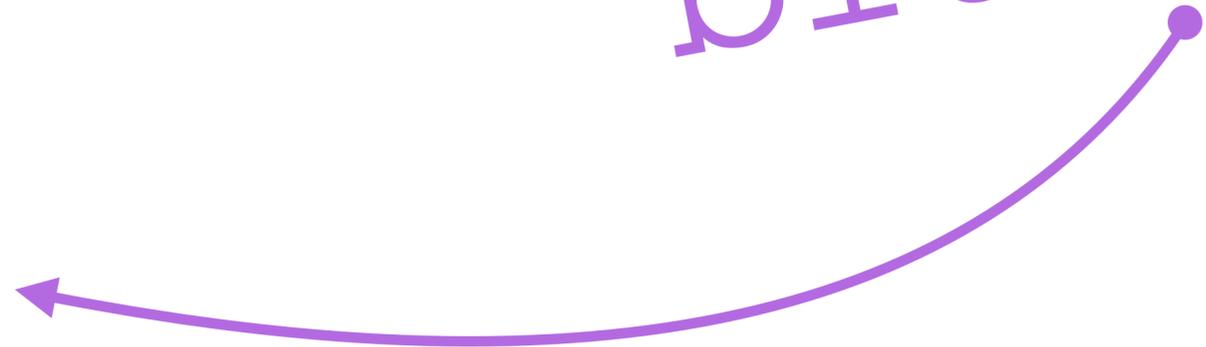
Time

$t_0$   $t_1$   $t_2$   $t_3$   $t_{\dots}$



**11 01 10 01 ...**

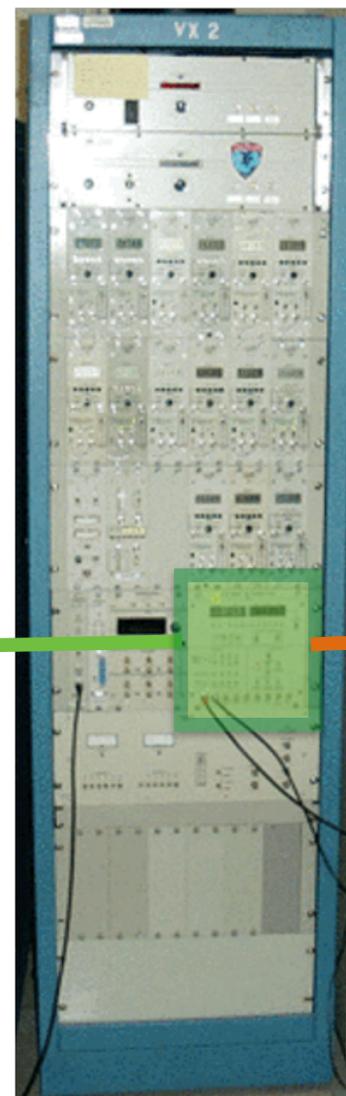
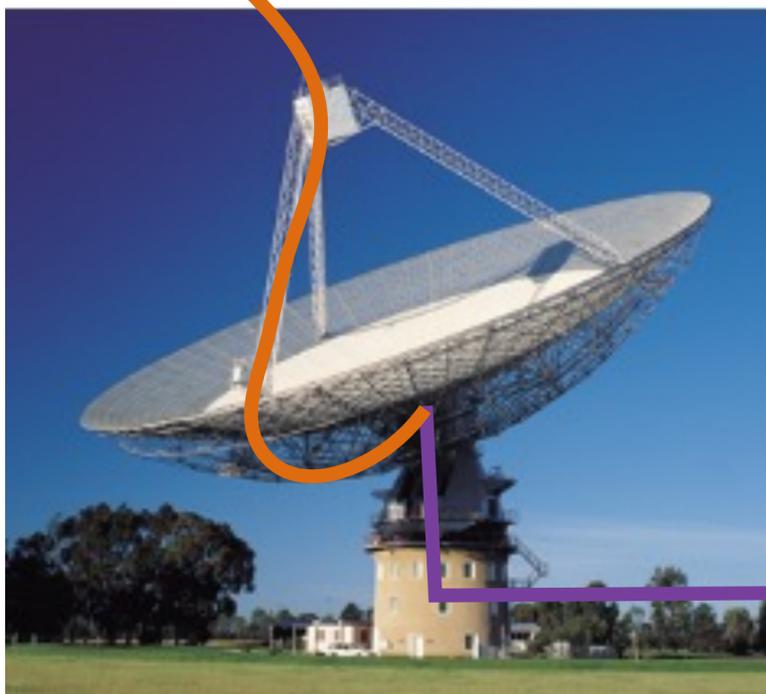
*bitstream*



sky freq  
1-100 GHz

IF  
0-1 GHz

bitstreams  
 $n \cdot 10^6$  Bps



mix sky with LO

BBC, A-to-D  
conversion

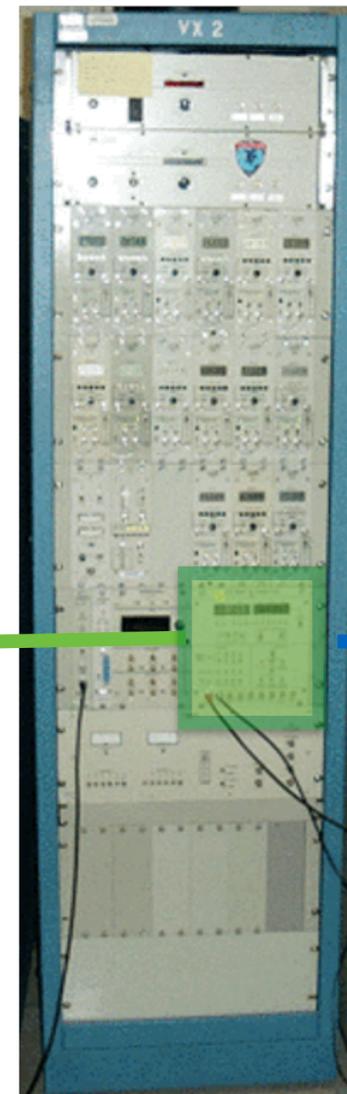
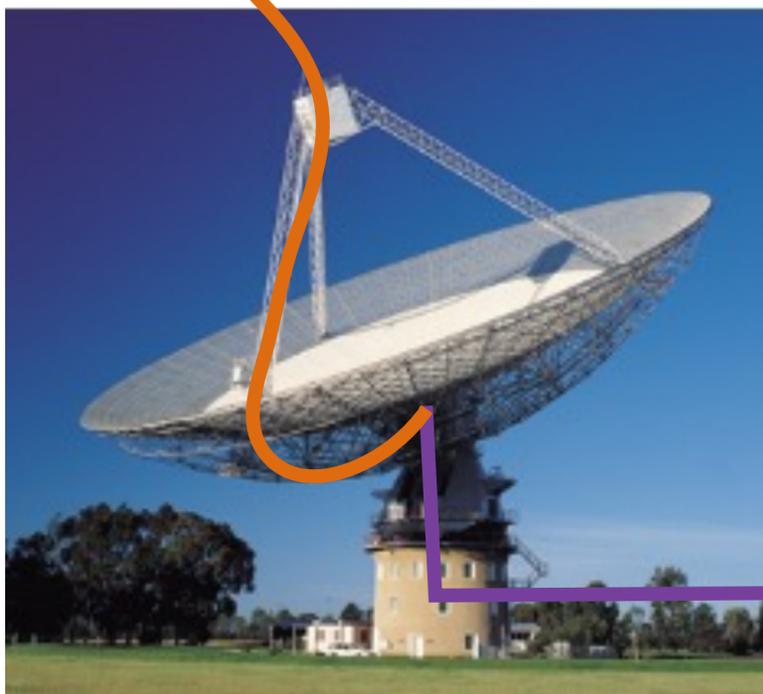
formatting

sky freq  
1-100 GHz

IF  
0-1 GHz

bitstreams  
 $n \cdot 10^6$  Bps

VLBI format  
dataframes



mix sky with LO

BBC, A-to-D  
conversion

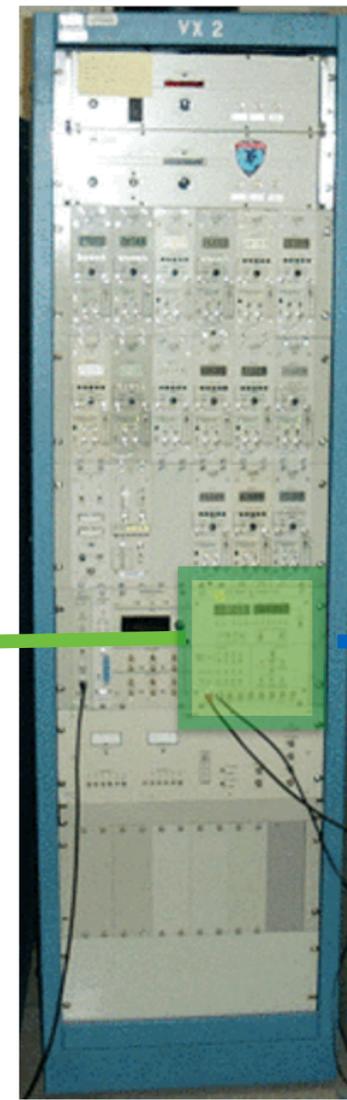
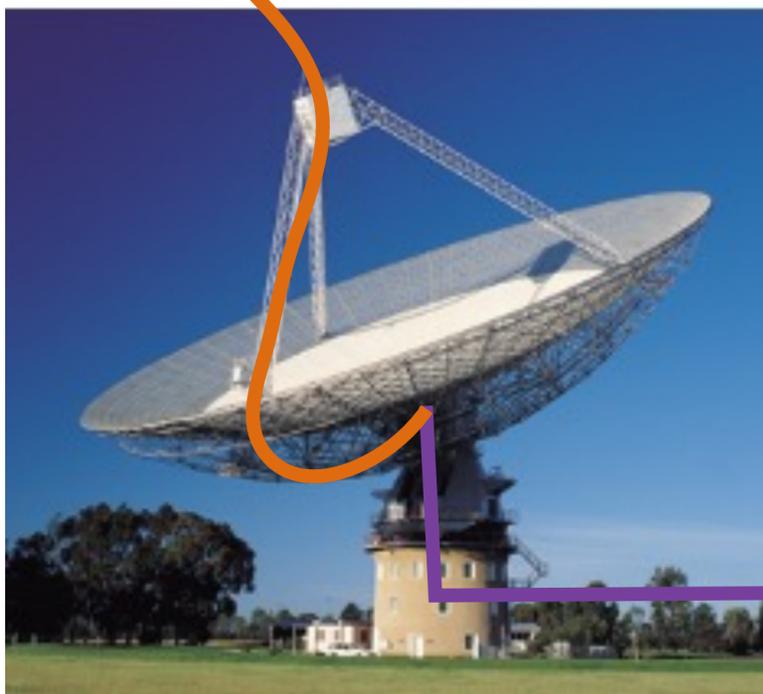
formatting

sky freq  
1-100 GHz

IF  
0-1 GHz

bitstreams  
 $n \cdot 10^6$  Bps

VLBI format  
dataframes



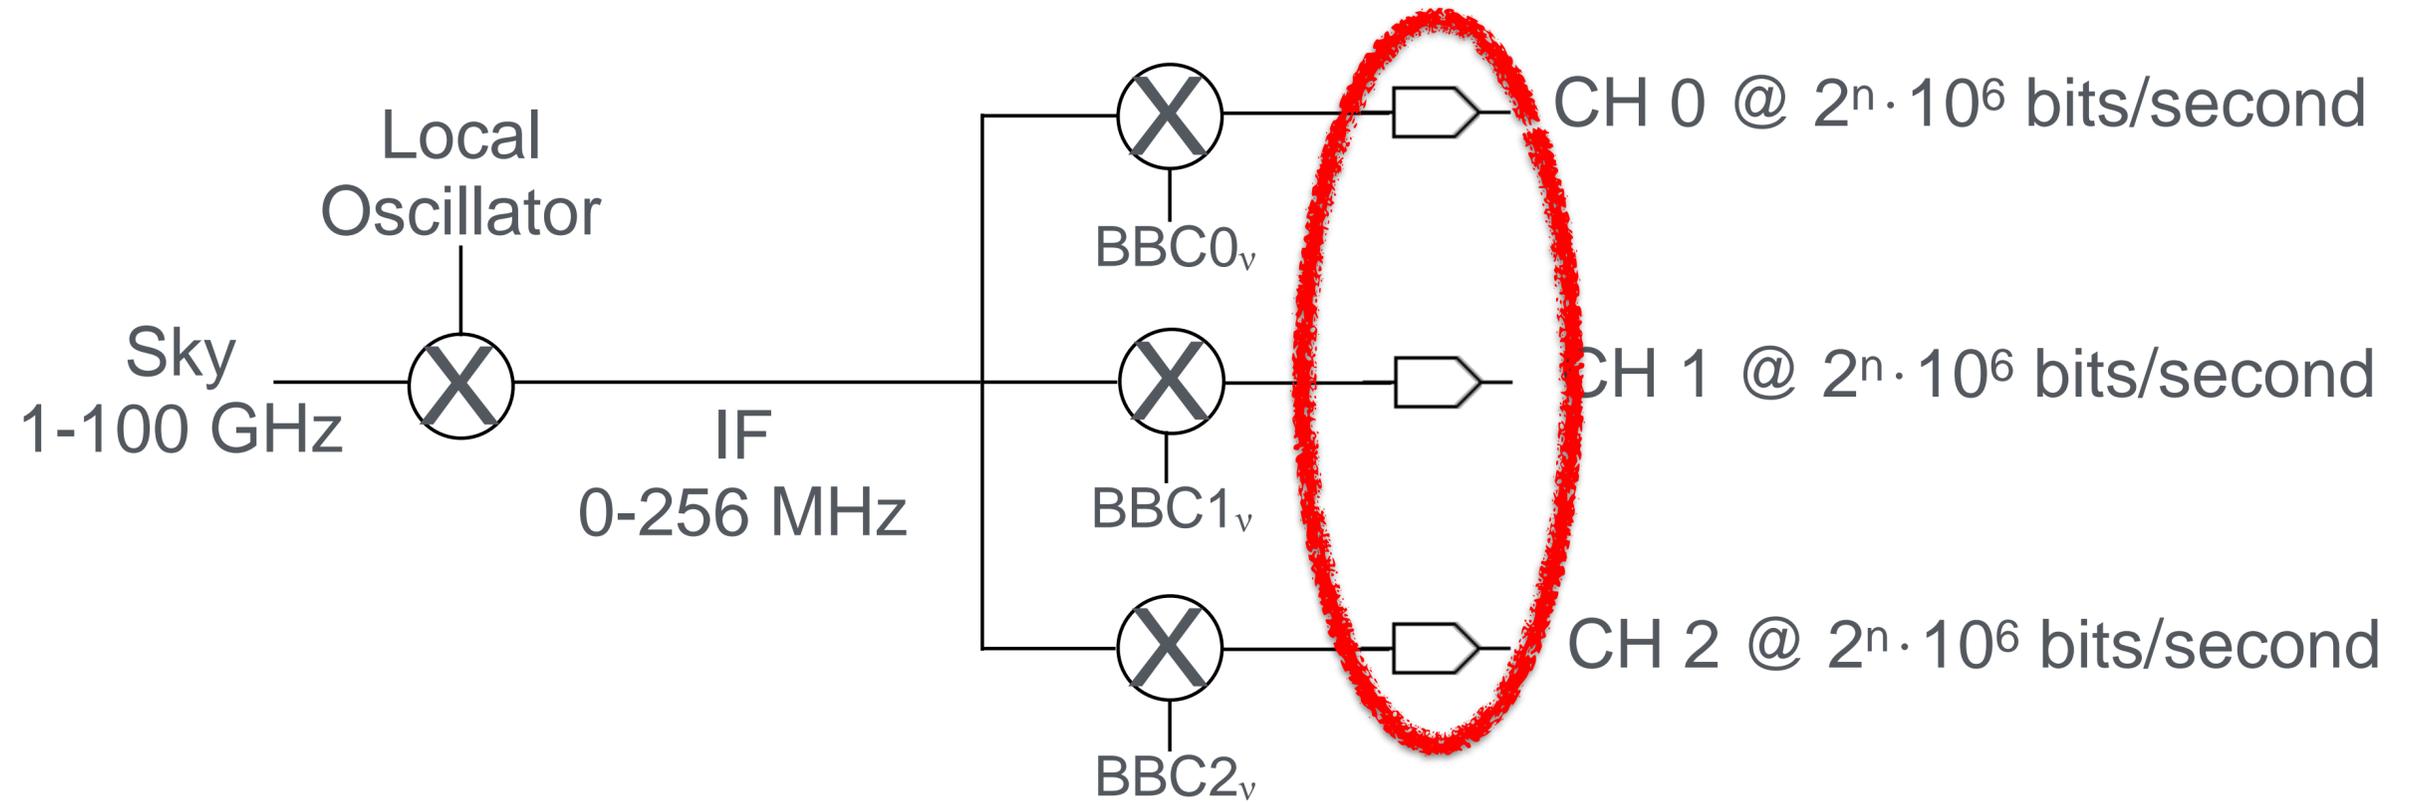
mix sky with LO

BBC, A-to-D  
conversion

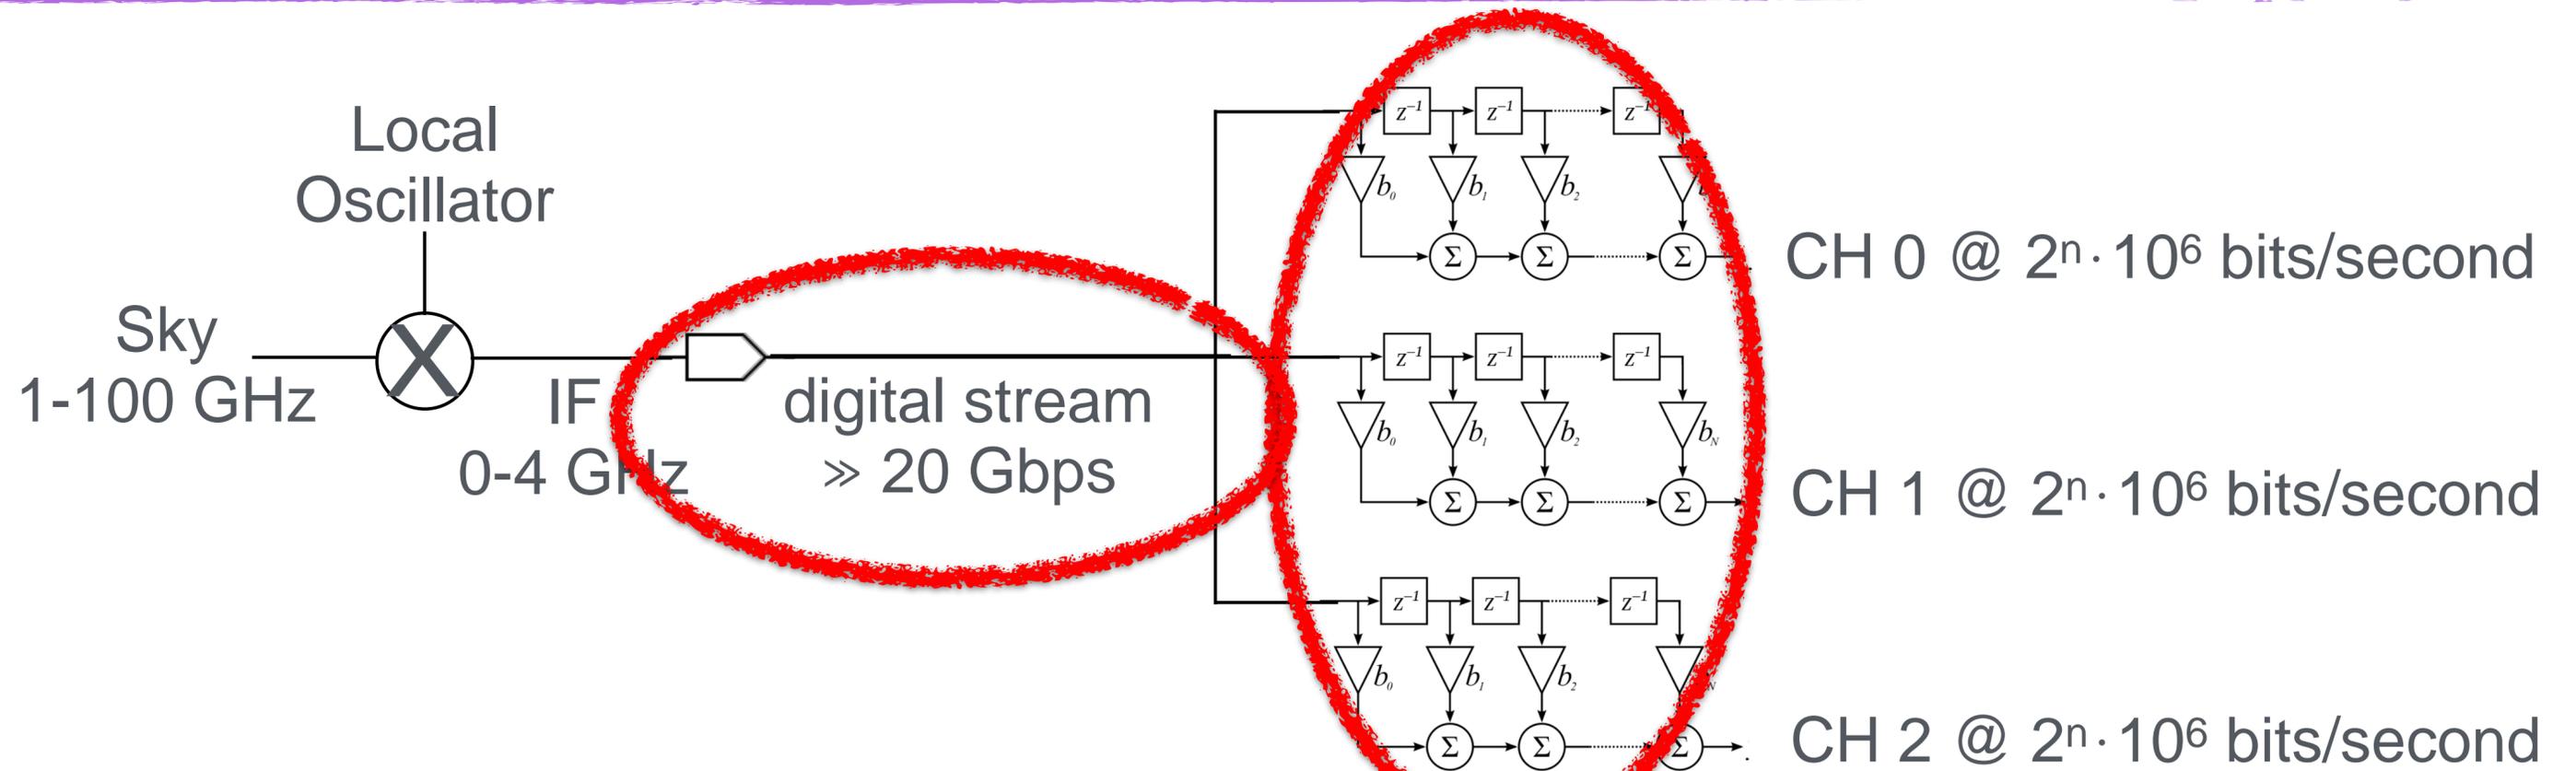
formatting

recording

Old



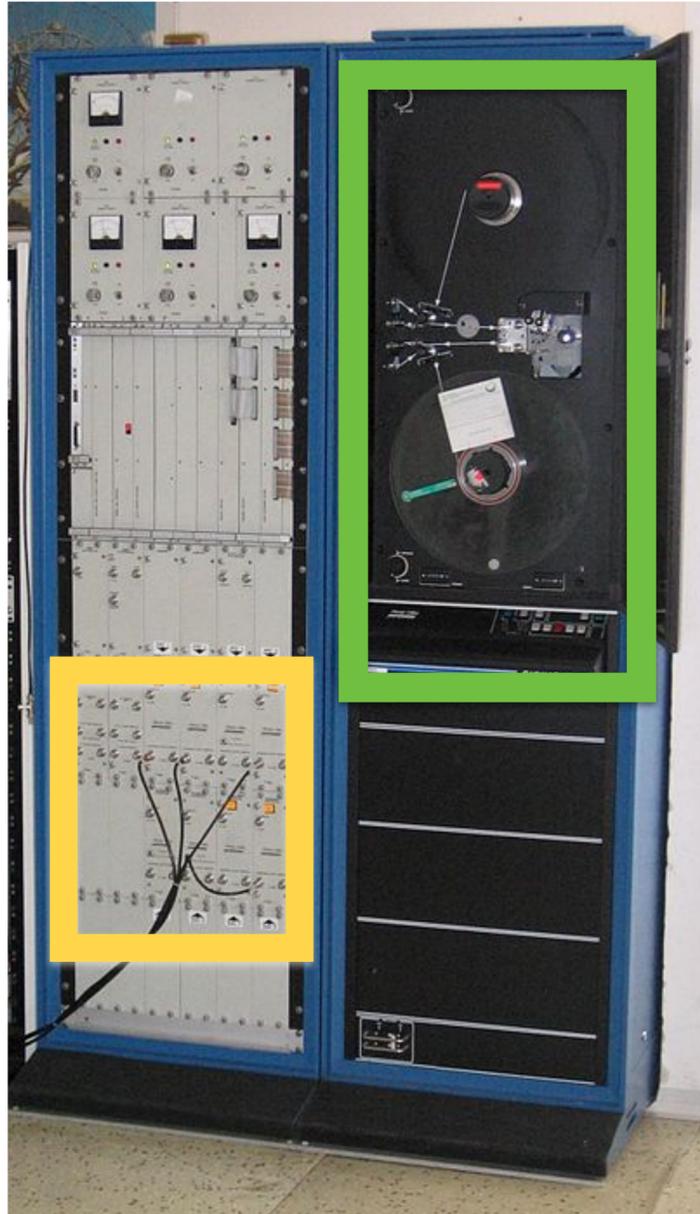
New



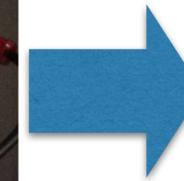
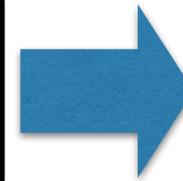
mixer



analog  $\Rightarrow$  digital



Analog Backend

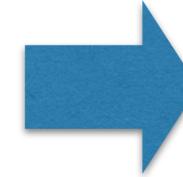
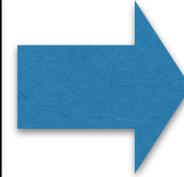


Recorder

64 x RS422  
at (max) 10 Mbps / link

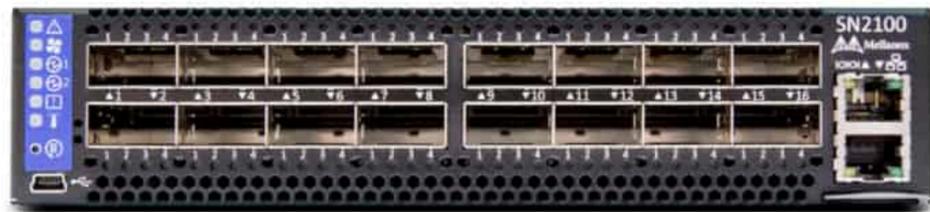


Digital Backends

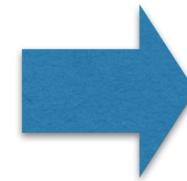
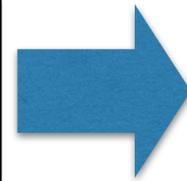


Recorder

$n \times \text{CX4}$   
at 10 Gbps / link



Digital Backend



Recorder

Recorder

$n \times$  SFP+/fibre  
at 10 Gbps / link

Getting as many  
bits on disk as  
possible

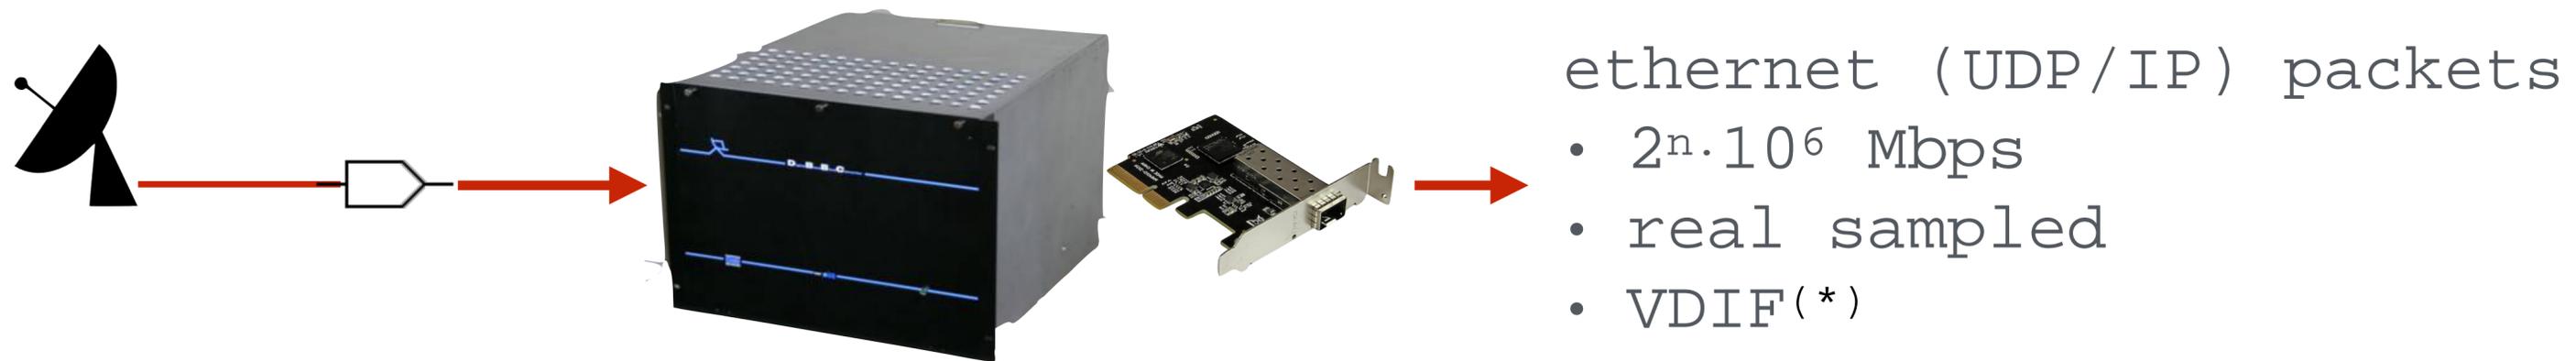
$$\sigma_T = \frac{T_{sys}}{\sqrt{\Delta \nu \cdot \Delta t}}$$

more = better

$$Sensitivity = \frac{1}{\sigma_T}$$

$$f_{\text{sample}} = 2 \cdot \Delta \nu$$

# Contemporary VLBI recorders



(\*) VLBI Data Interchange Format - <https://vlbi.org/vlbi-standards/vdif/>

# Contemporary VLBI recorders



Mark6



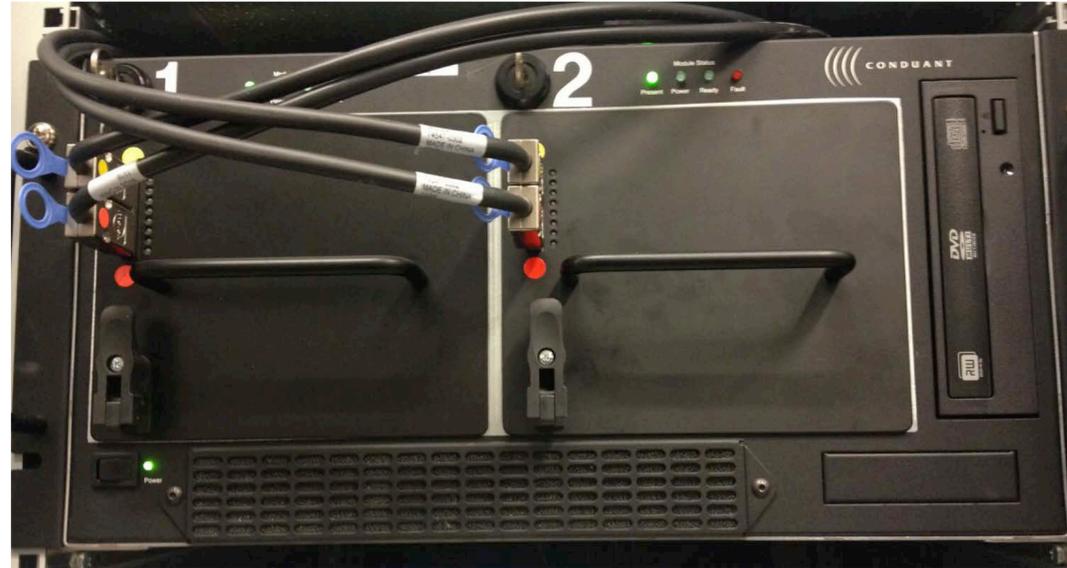
FlexBuff



OCTADISK

...?

# Contemporary VLBI recorders



Mark6



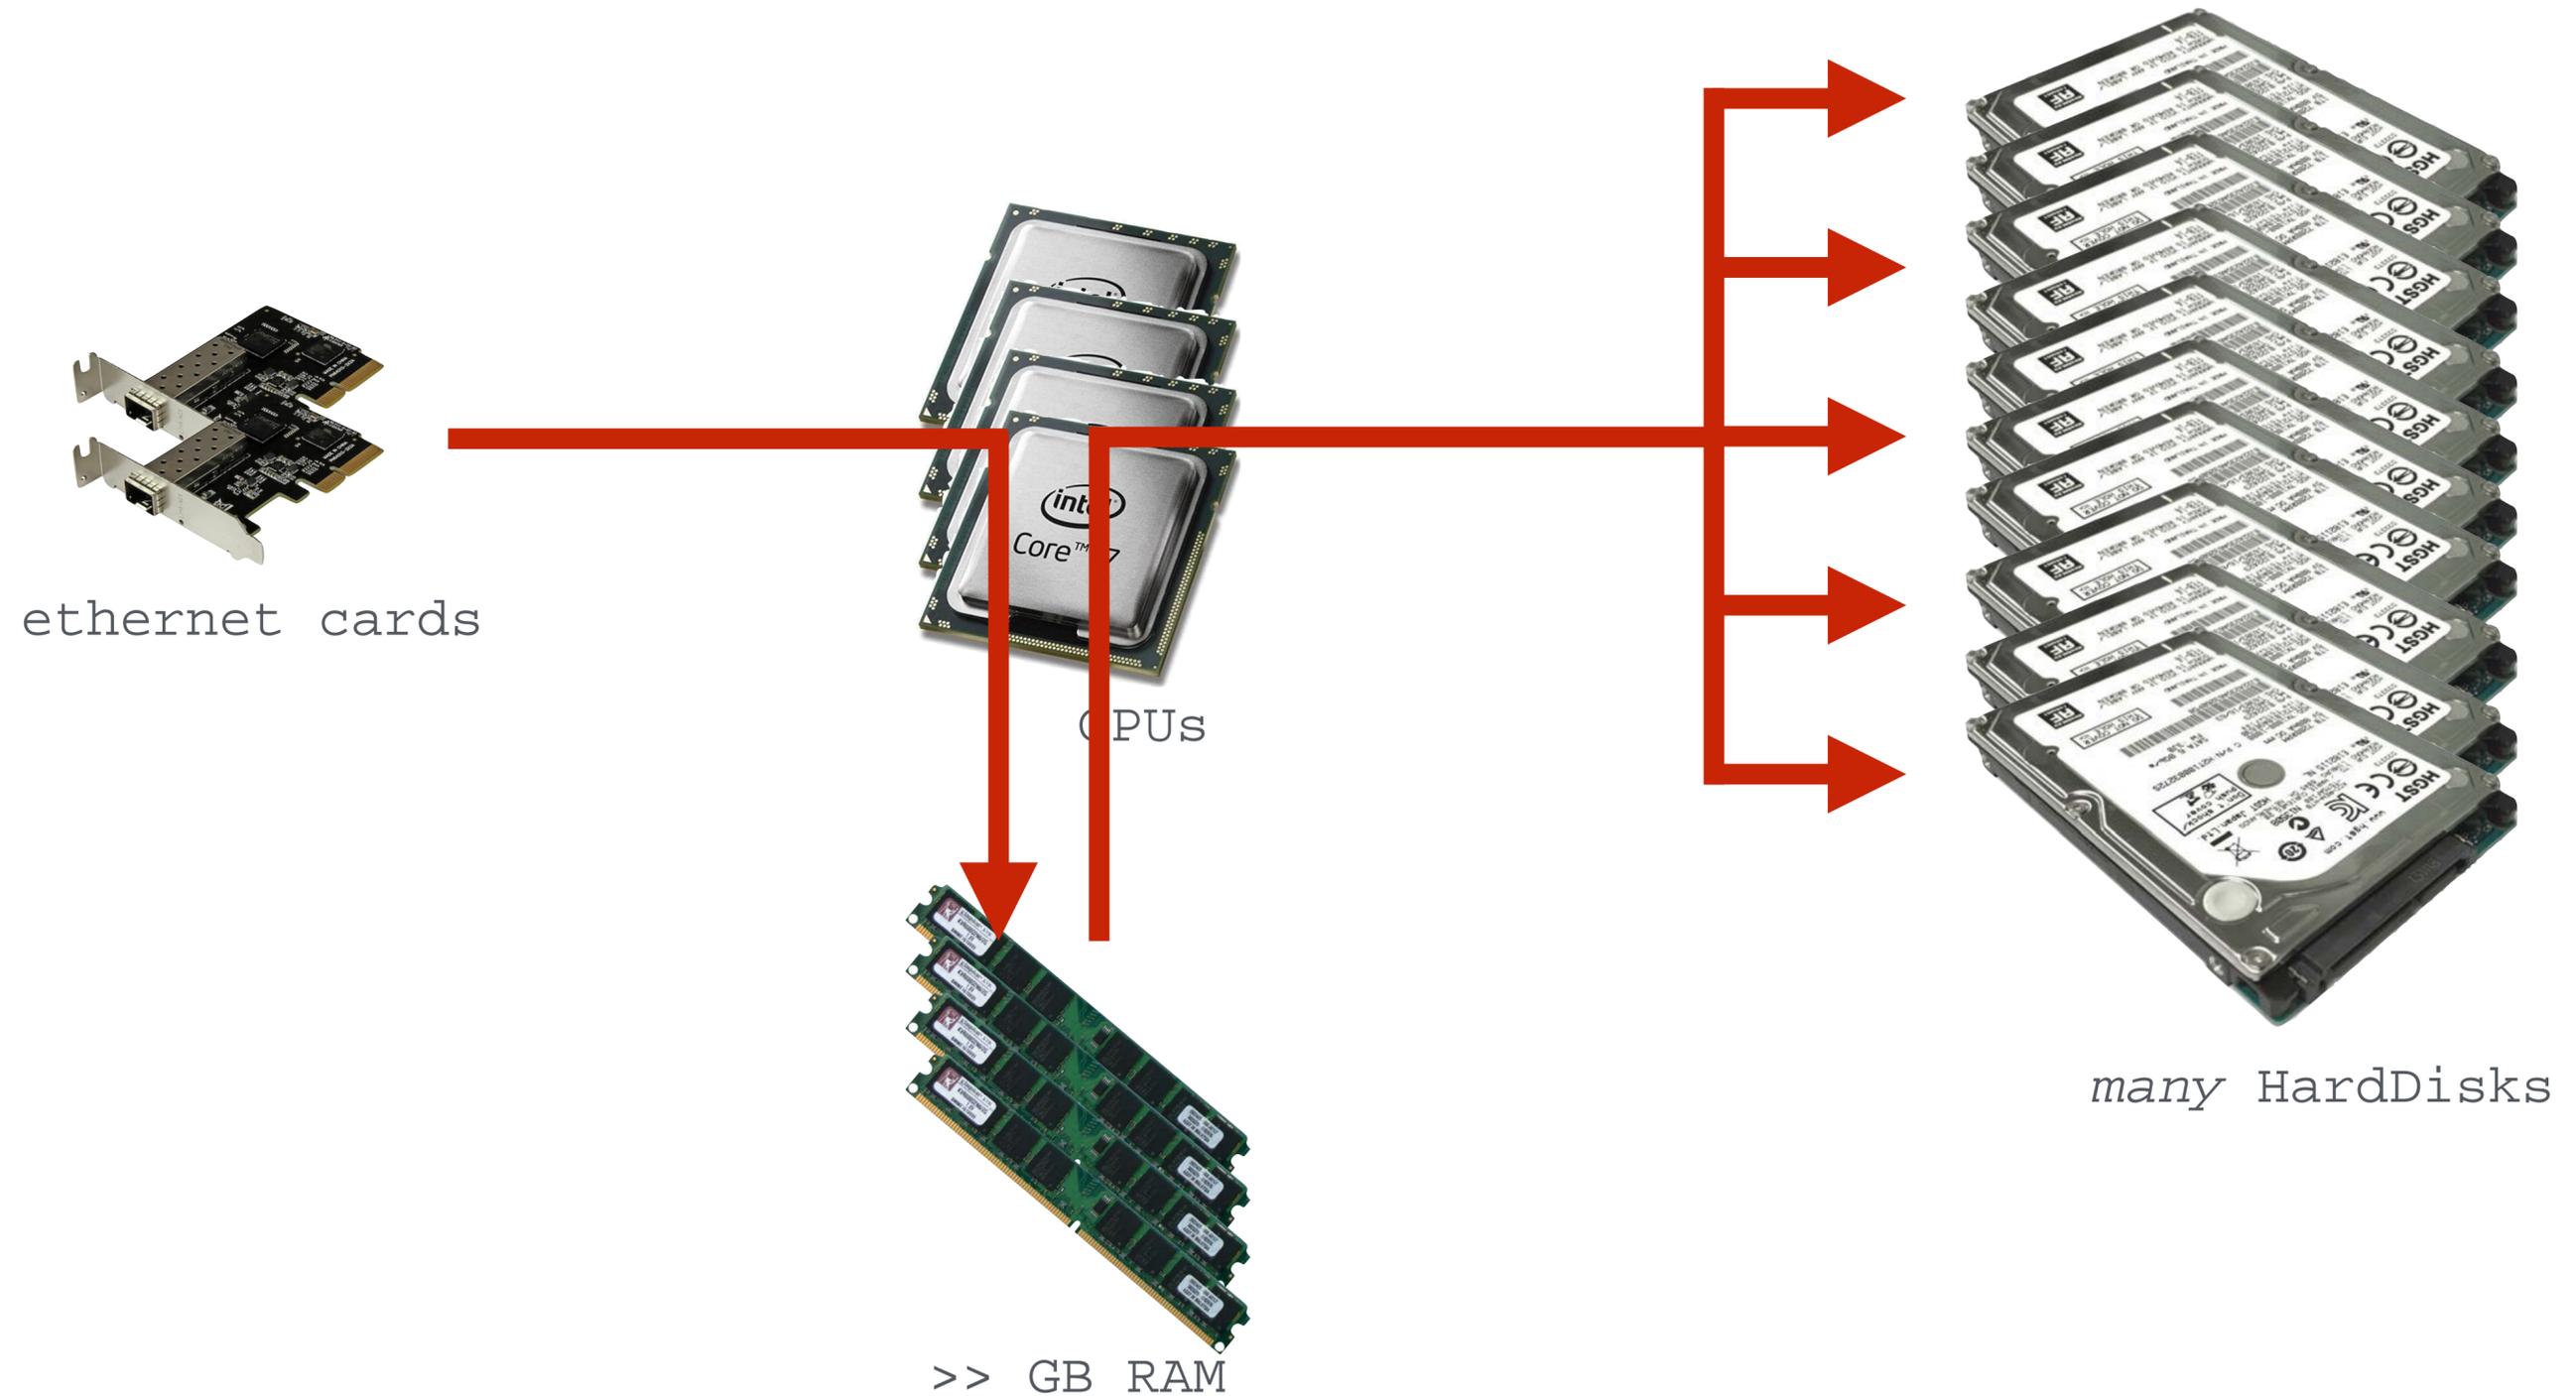
FlexBuff



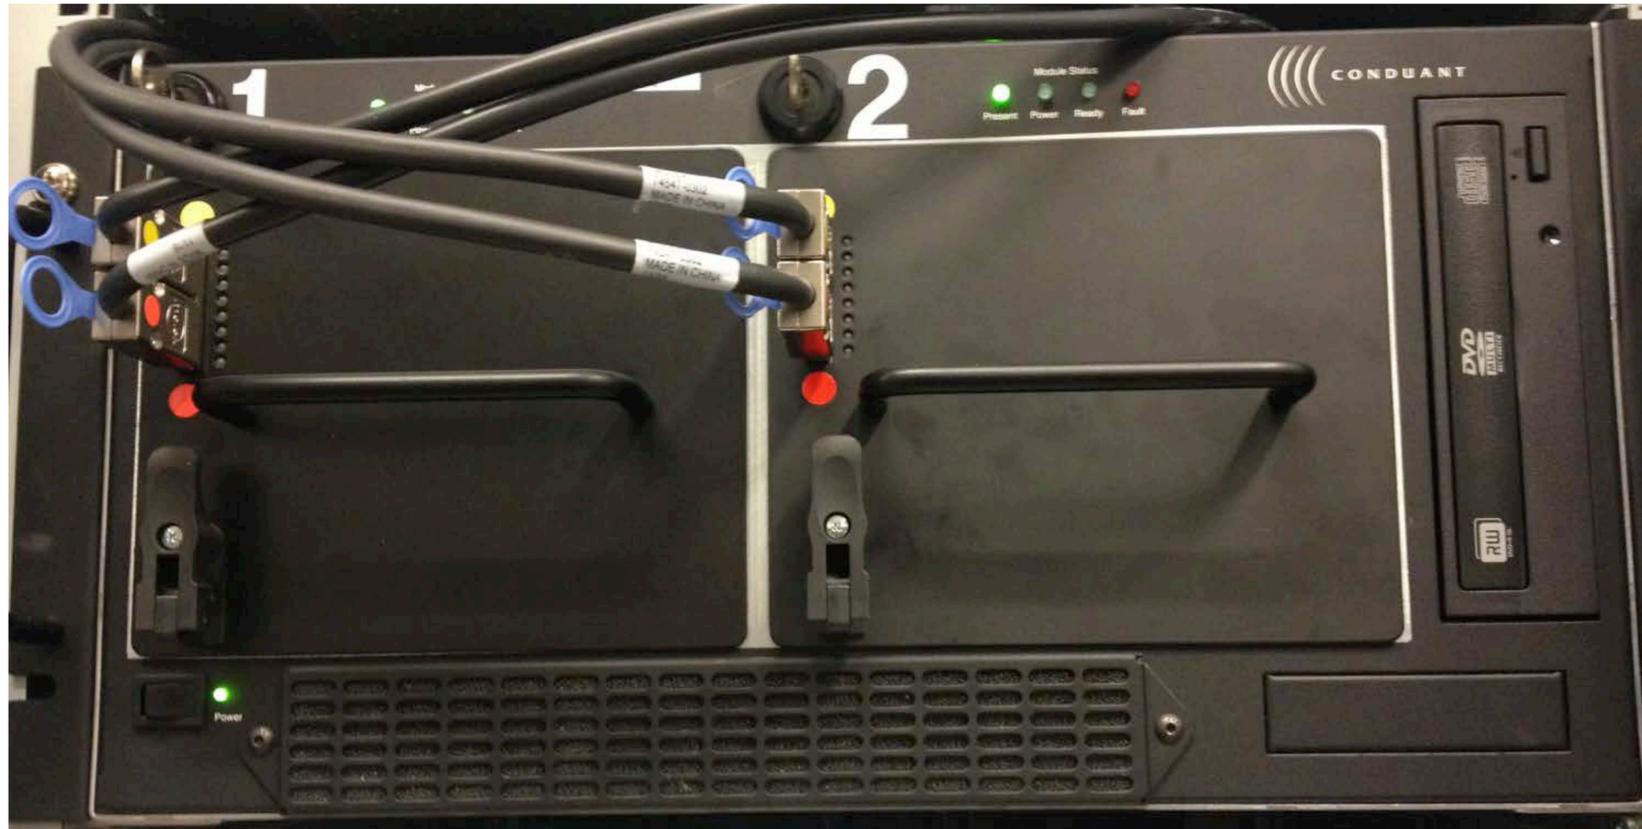
OCTADISK



# Contemporary ethernet recorders

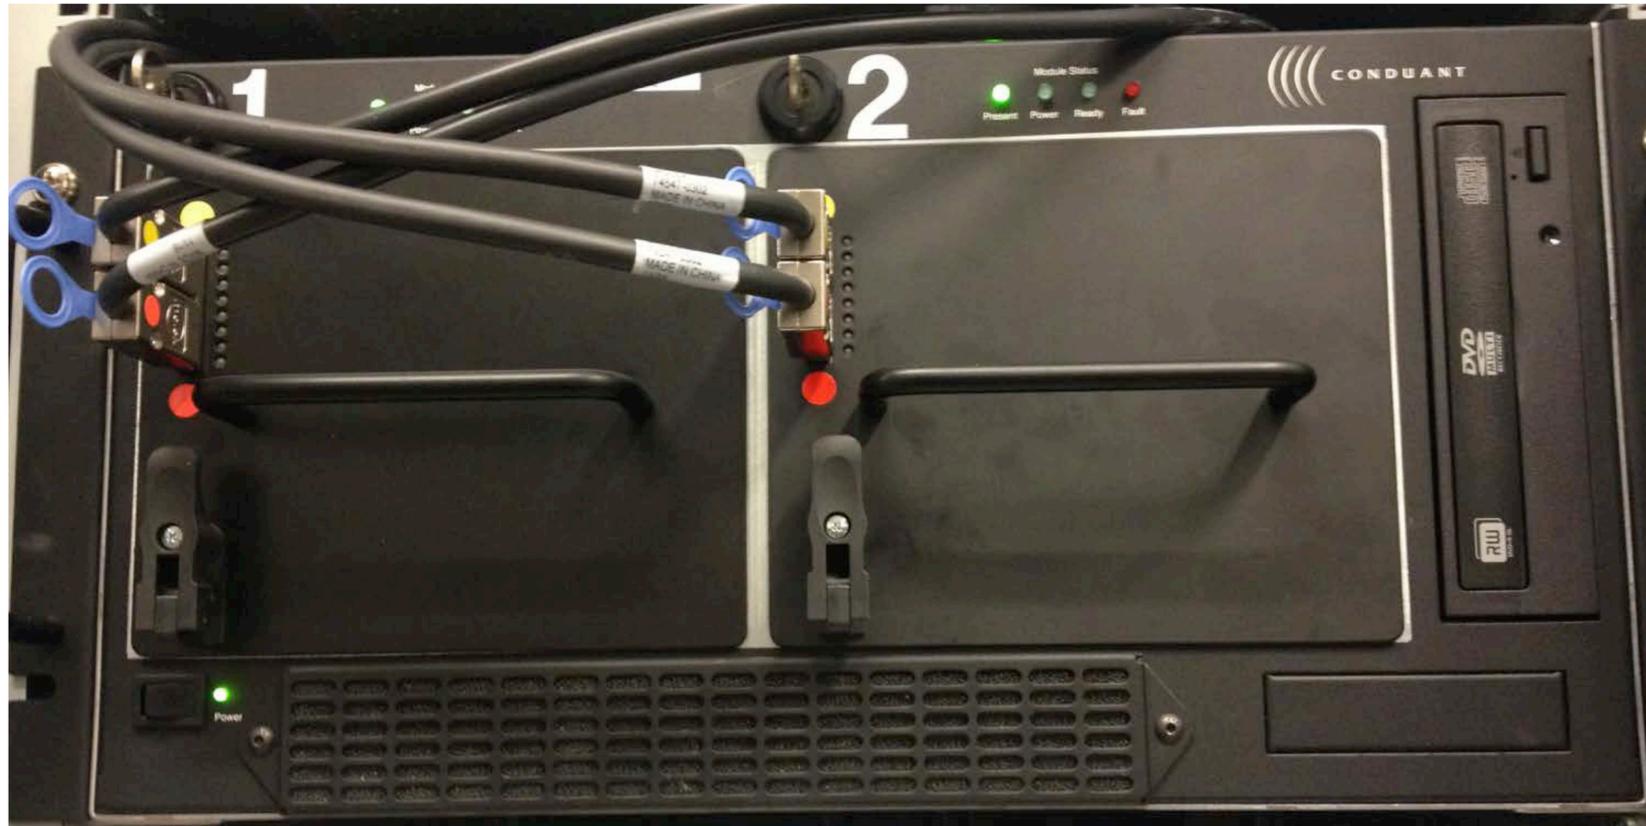


# Contemporary ethernet recorders



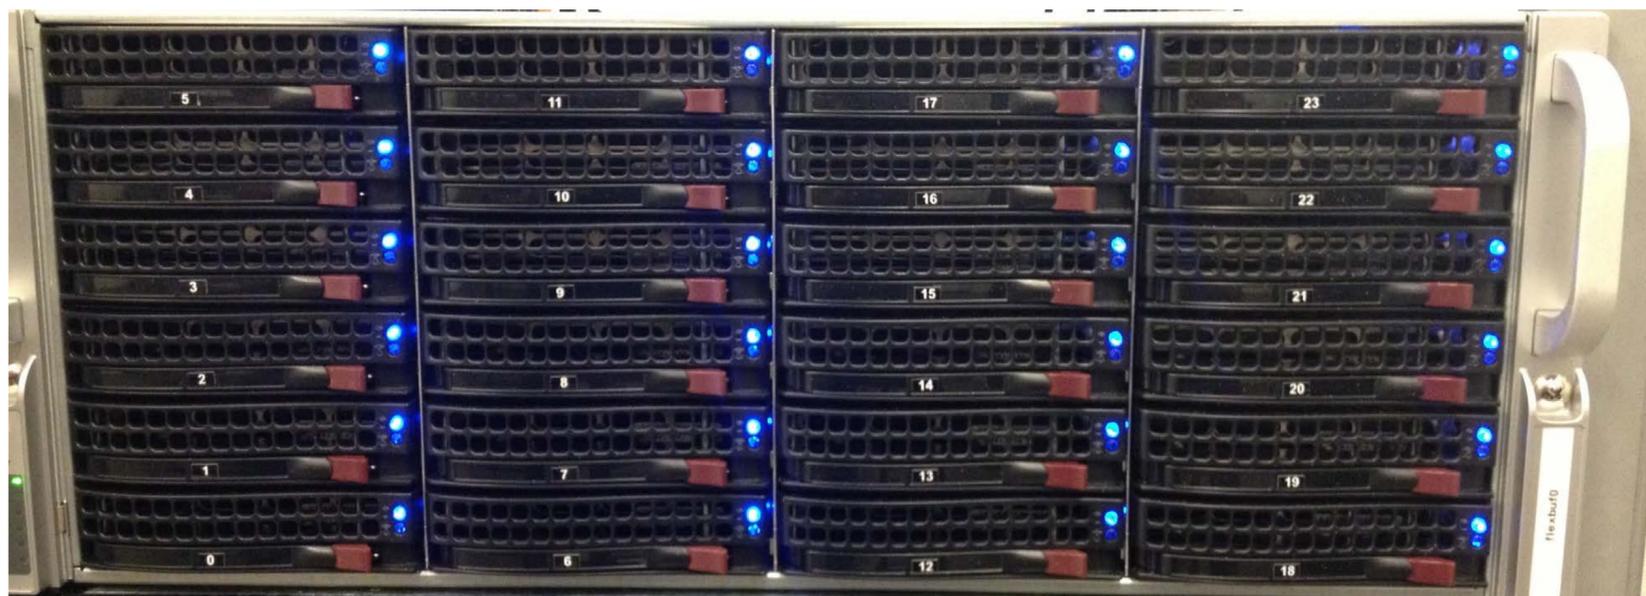
- Mark6 (MIT Haystack/Conduant)
- proprietary hardware
  - only one supplier (Conduant Corp.)
  - $\approx$  8 Gpbs
  - 30 k€ (inc. 32 x 10 TB HDD)

# Contemporary ethernet recorders



Mark6 (MIT Haystack/Conduant)

- proprietary hardware
- only one supplier (Conduant Corp.)
- $\leq 8$  Gpbs ( $\leq 16$  Gpbs with expander)
- 30 k€ (32 x 10 TB HDD + expander)



FlexBuff (Metsähovi / JIVE)

- fully customizable, fully COTS
- $n$  Gpbs
- ~20 k€ (inc. 36 HDD)

*Concept: A. Mujunen, Metsähovi*  
*Productionalized: JIVE*

# Contemporary ethernet recorders

The only tangible difference between the systems. The rest is semantics/software.

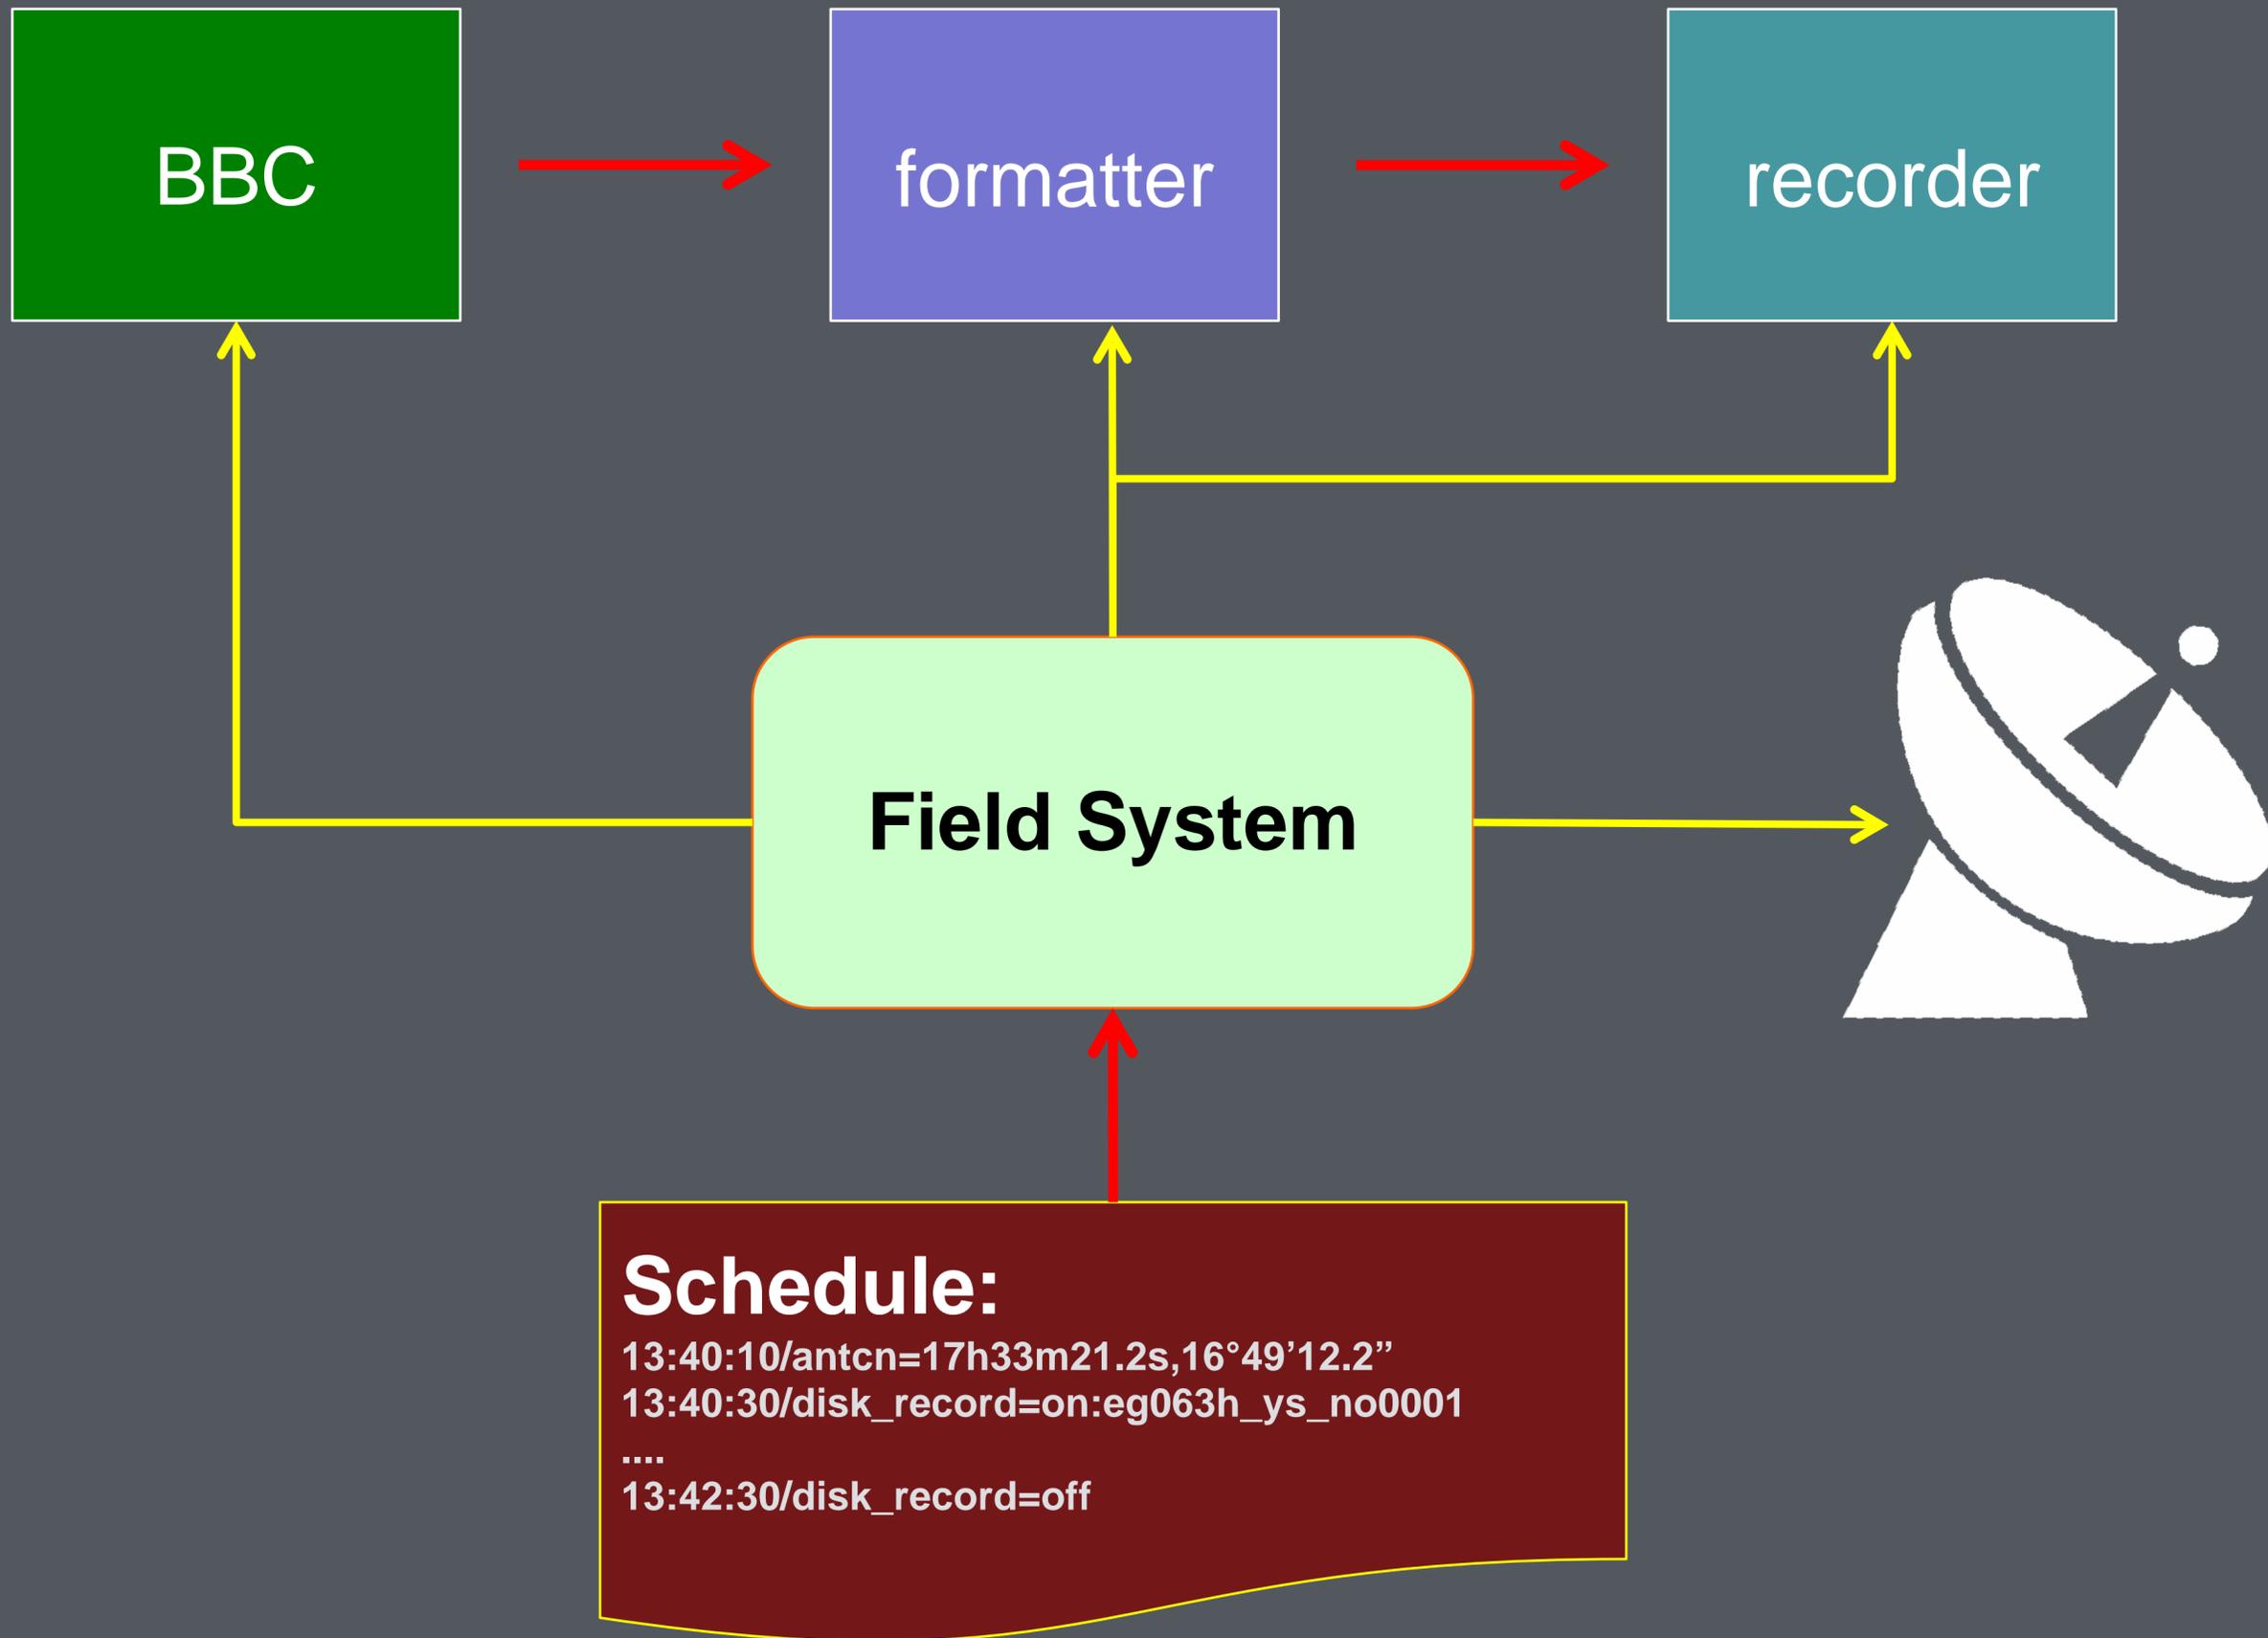
Mark6 removable disk packs



FlexBuff fixed disks



And now for something  
completely different.

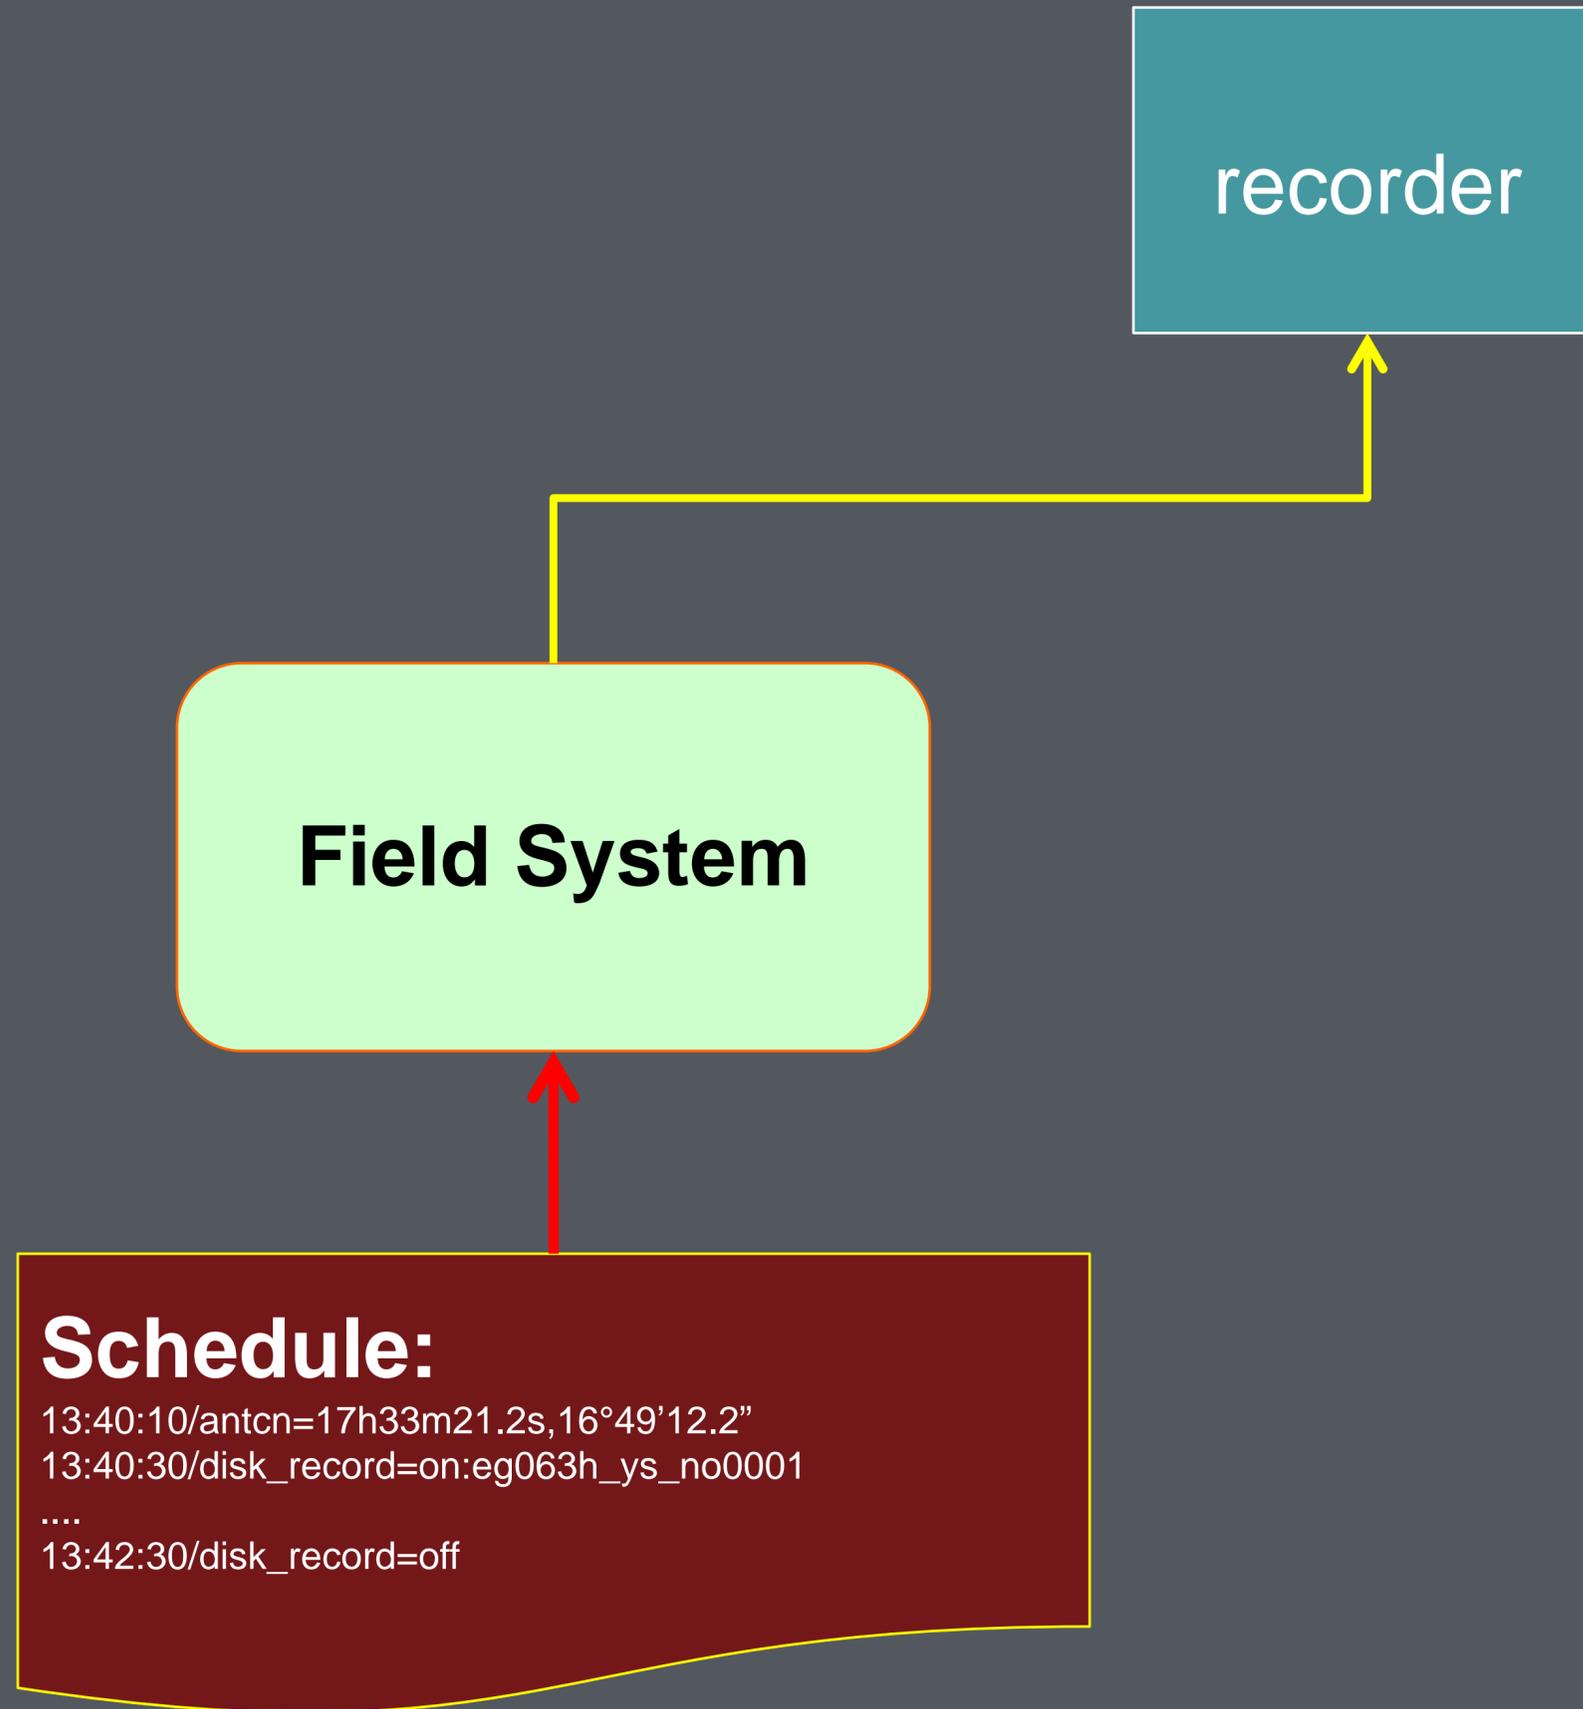


recorder

**Field System**

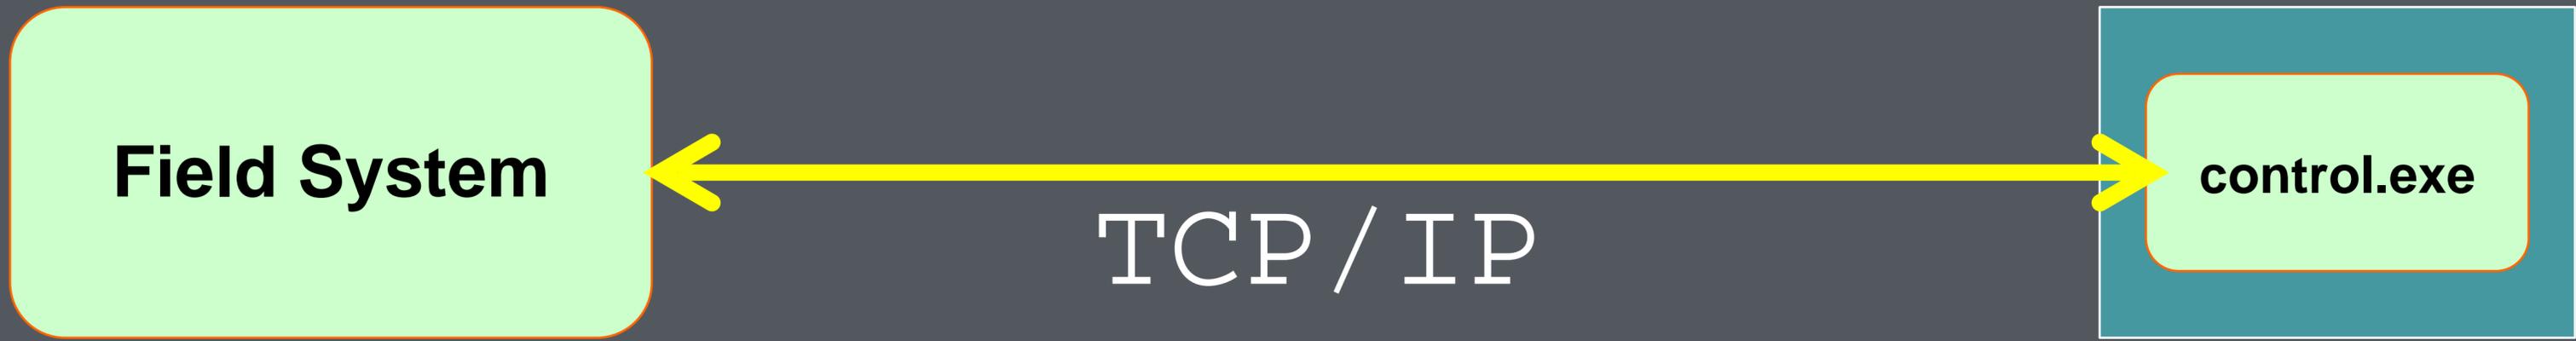
**Schedule:**

13:40:10/antcn=17h33m21.2s,16°49'12.2"  
13:40:30/disk\_record=on:eg063h\_ys\_no0001  
....  
13:42:30/disk\_record=off









**Field System**

TCP / IP

**control.exe**

# **VSI/S - V(LBI) S(tandard) I(nterface) / S(oftware)**

- ASCII based
- every **query** or **command** yields **reply**

# VSI/S - V(LBI) S(tandard) I(nterface) / S(oftware)

- ASCII based
- every **query** or **command** yields **reply**

<command> = <param1> [ : <param2> : ... ];

<query> ? [ <param1> : <param2> : ... ];

```
<command> = <param1> [ : <param2> : ... ];  
!<command> = <result code> [: <value1> : ... ];
```

```
<query> ? [ <param1> : <param2> : ... ];  
!<query> ? <result code> [: <value1> : ... ];
```



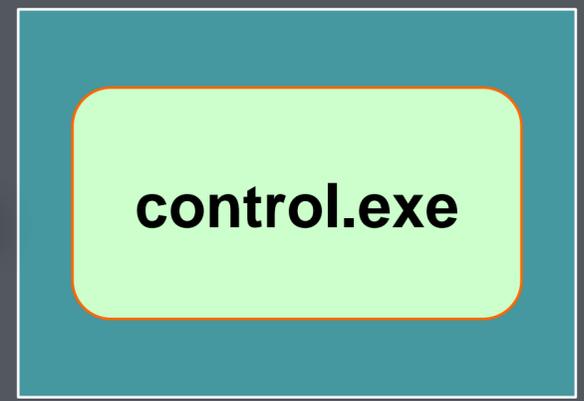
```
mode = mark4 : 64;
```

**Field System**

**control.exe**

```
mode = mark4 : 64;
```

```
<< execute  
command >>
```



```
mode = mark4 : 64;
```

```
<< execute  
command >>
```

```
!mode = 0 ;
```

# VSI/S - V(LBI) S(tandard) I(nterface) / S(oftware)

**<result code>** is an ASCII integer as follows:

- 0 - action successfully completed
- 1 - action initiated or enabled, but not completed
- 2 - command not implemented or not relevant to this DTS
- 3 - syntax error
- 4 - error encountered during attempt to execute command
- 5 - currently unable to service request; try again later
- 6 - inconsistent or conflicting request
- 7 - no such keyword
- 8 - parameter error

SS_rev?	p. 49	Get StreamStor firmware/software revision levels (query only)
OS_rev?	p. 17	Get system information (query only)
protect	p. 24	Get details of operating system (query only)
recover	p. 28	Set/remove erase protection for active module
reset	p. 31	Recover record pointer which was reset abnormally during recording
SS_rev?	p. 32	Reset Mark 5 unit (command only)
SS_rev?	p. 37	Get StreamStor firmware/software revision levels (query only)

### 8.2 Record/Play

mode	p. 31	Set data recording/playback mode
play	p. 38	Play data from current or specified play point
play_rate	p. 40	Set playback data rate; set tvgr rate
record	p. 44	Turn recording on/off; assign scan label
scan_play	p. 52	Play scan specified by current value of scan
scan_set	p. 53	Set scan for scan_check, scan_play, scan_set
skip	p. 55	Skip forward/backward scan

### 8.3 Data Checking

data_check?	p. 15	Check data starting at position
scan_check?	p. 50	Get scan parameters (query only)
track_check?	p. 61	Check data on selected tracks
track_set	p. 63	Select tracks for monitoring

### 8.4 Data Transfer

disk2file	p. 22	Transfer data from Mark 5 to file
disk2net	p. 24	Transfer data from Mark 5 to network
file2disk	p. 28	Transfer data from file to Mark 5
in2net	p. 30	Transfer data directly from network to network
net_protocol	p. 35	Set network data-transfer protocol
net2disk	p. 33	Transfer data from network to Mark 5
net2out	p. 34	Transfer data directly from network to output

### 9.2 System Setup and Monitoring

lpps_source	p. 12	Select source of lpps synchronization tick
clock_set	p. 16	Specify frequency and source of the CLOCK driver
DOT?	p. 27	Get DOT (Data Observe Time) clock information
DOT_inc	p. 28	Increment DOT clock
DOT_set	p. 29	Set DOT clock on next external lpps tick
error?	p. 31	Get error number/message (query only)
mode	p. 35	Set data recording/readback mode
status?	p. 51	Get system status (query only)

### 9.3 Data Checking

disk2net	p. 25	Transfer data between start-scan and stop-scan pointers
file2disk	p. 32	Transfer data from file to Mark 5
in2net	p. 34	Transfer data directly from network to network
net_protocol	p. 37	Set network data-transfer protocol
net2disk	p. 36	Transfer data from network to Mark 5
record	p. 41	Turn recording on/off; assign scan label

### 7.2 System Setup and Monitoring

error?	p. 18	Get error number/message (query only)
mode	p. 22	Set data recording/readback mode
status?		Get system status (query only)

### 7.4 Data Transfer

net_protocol	p. 23	Set network data-transfer protocol
record	p. 29	Turn recording on/off; assign scan label
fill_pattern	p. 20	Set StreamStor 32 bit fill pattern
packet	p. 29	Set packet acceptance criteria

Different hardware  
=  
different command set

MARK 5A COMMAND SET

MARK 5B DIM COMMAND SET

MARK 5C DIM COMMAND SET

# Contemporary ethernet recorders

The only tangible difference between the systems. The rest is semantics/software.

Mark6 removable disk packs



FlexBuff fixed disks



# High speed packet capture

This does not come for free ..

- O/S defaults wrong for this use case
- O/S network buffer sizes too small
- spread (interrupt) over >> cores
- pays no attention to hardware layout

Without tuning get < 4 Gbps lossless ..

# High speed packet capture

Tuning is a topic of its own

# Origin of the differences

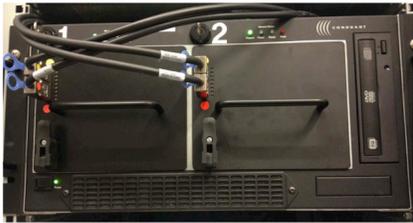
Choice of recording software

# Origin of the differences

Choice of recording software

cplane / dplane

- MIT Haystack

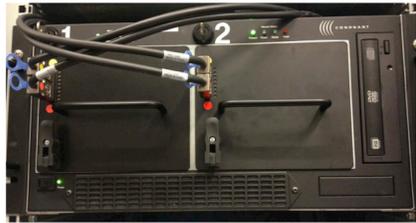


# Origin of the differences

Choice of recording software

cplane / dplane

- MIT Haystack



jive5ab(\*)

- Joint Institute for VLBI in Europe



(\*) <https://github.com/jive-vlbi/jive5ab>

# Consequence(s) of the choice

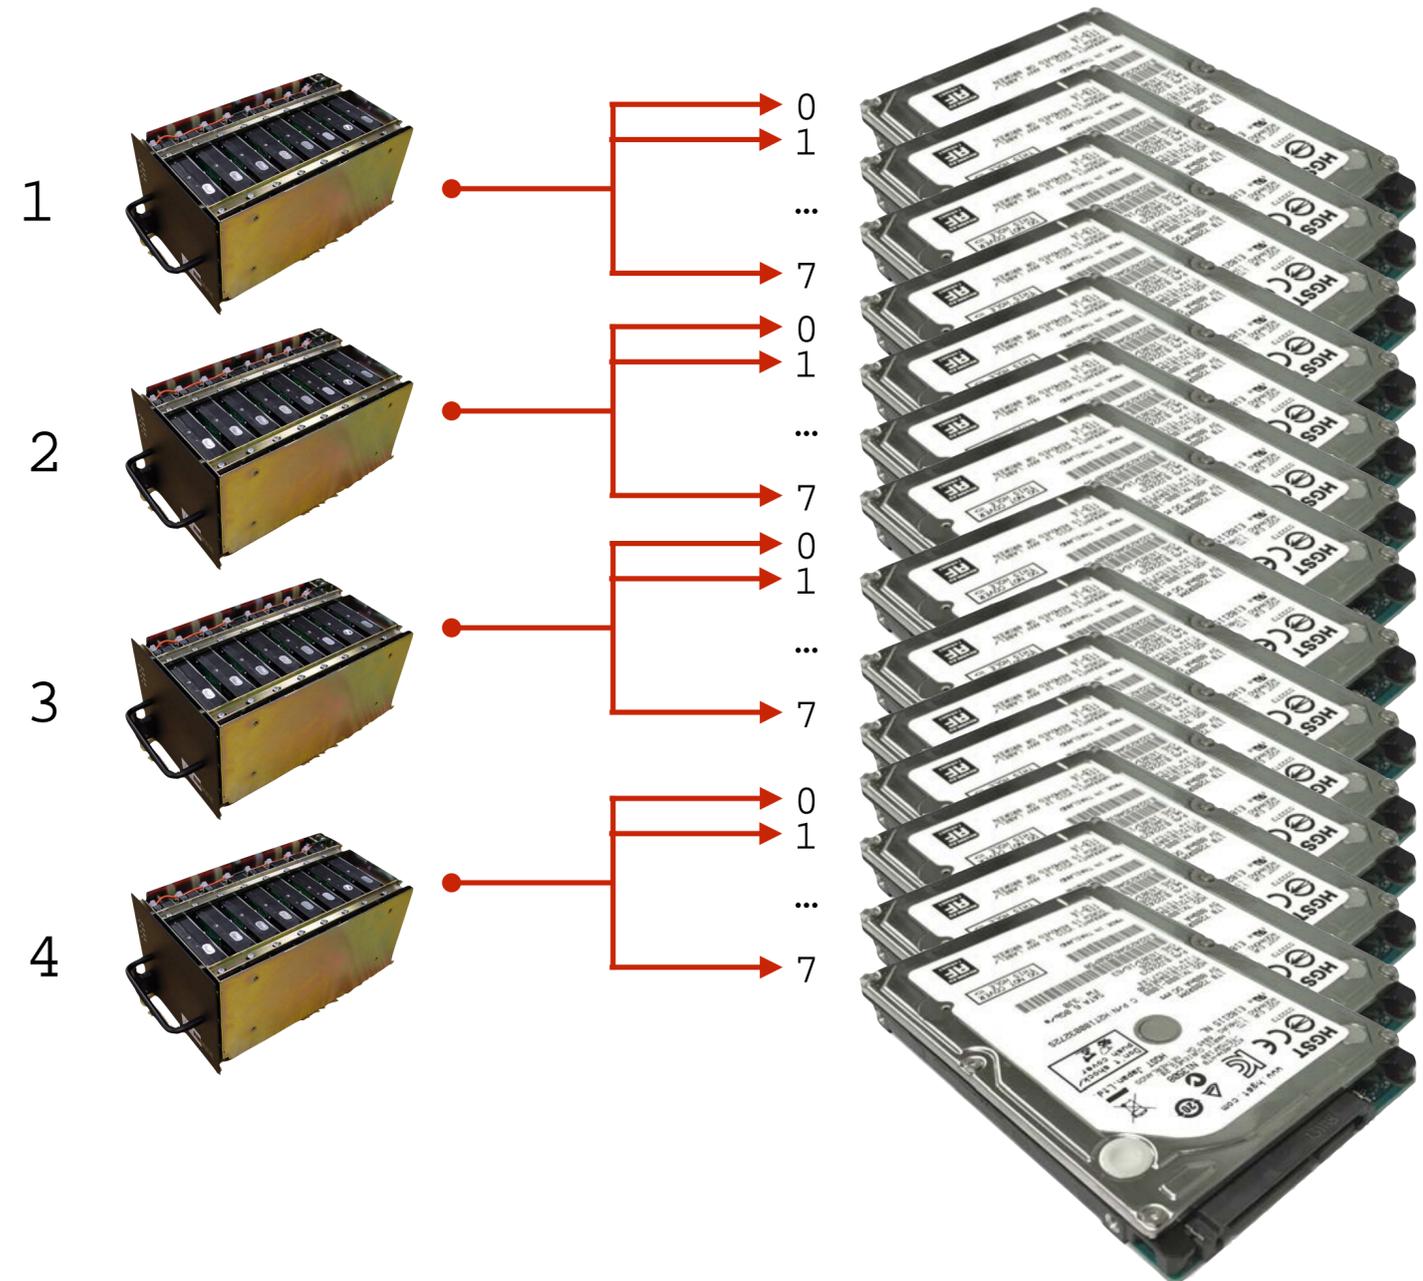
# Consequence(s) of the choice

Mount points for the data disks



# Consequence(s) of the choice

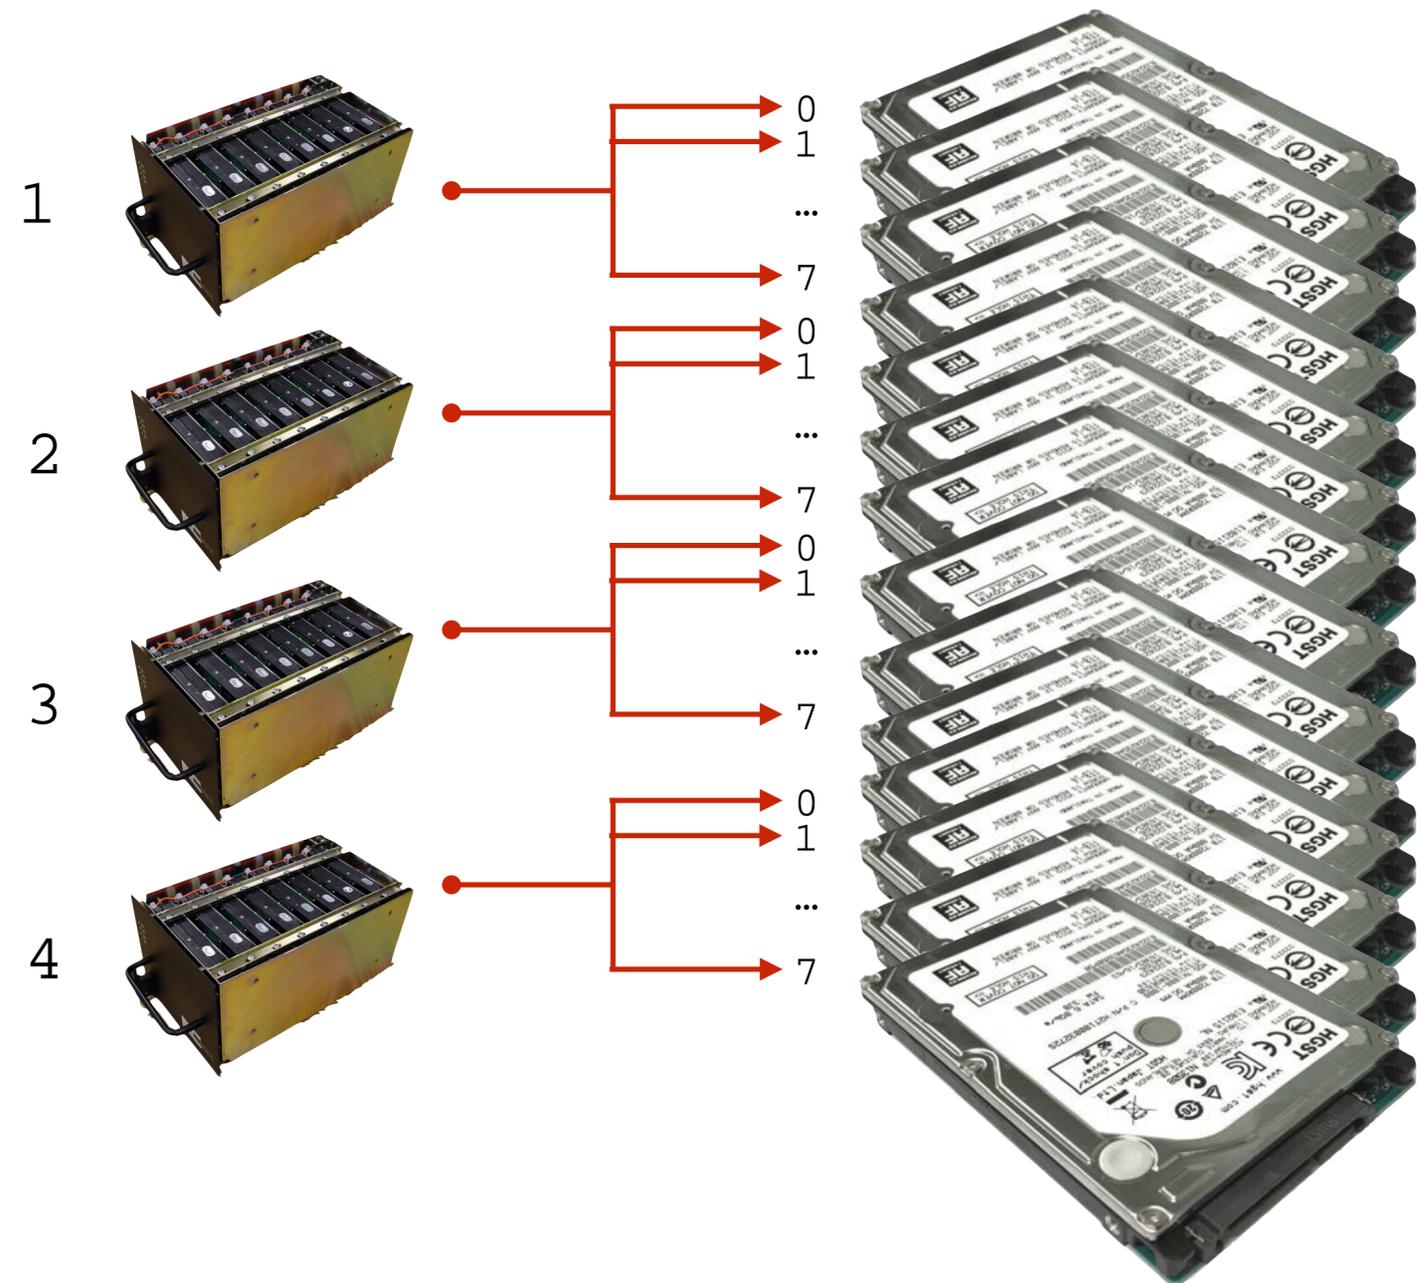
Mount points for the data disks (cplane, Mark6)



# Consequence(s) of the choice

Mount points for the data disks (cplane, Mark6)

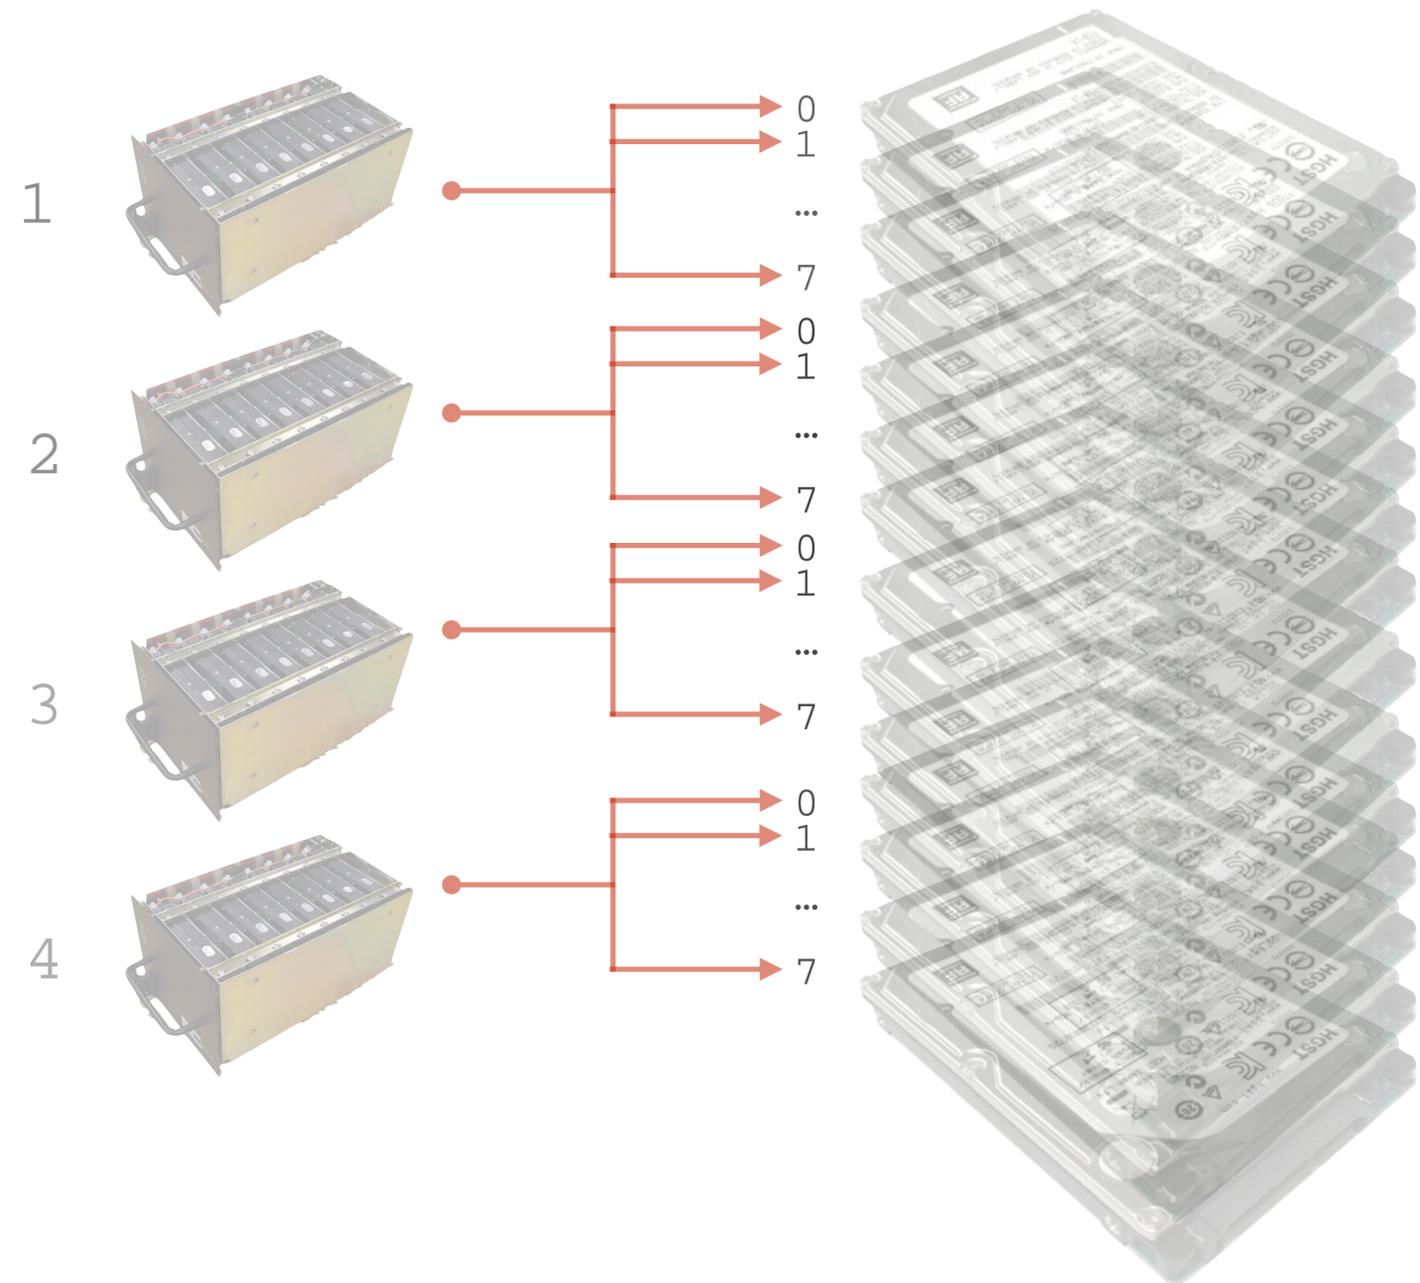
```
/mnt/disk/1/0/data  
  /1/1/data  
  ...  
  /1/7/data  
/mnt/disk/2/0/data  
  ...  
  /2/7/data  
  ...  
/mnt/disk/4/7/data
```



# Consequence(s) of the choice

Mount points for the data disks (cplane, Mark6)

```
/mnt/disk/1/0/data  
  /1/1/data  
  ...  
  /1/7/data  
/mnt/disk/2/0/data  
  ...  
  /2/7/data  
  ...  
/mnt/disk/4/7/data
```



regex: /mnt/disk/[1-4]/[0-7]/data

# Consequence(s) of the choice

Mount points for the data disks (jive5ab, \*)

```
/mnt/... ???  
/data/... ???  
/... ???
```

*jive5ab doesn't care,  
BUT ...*

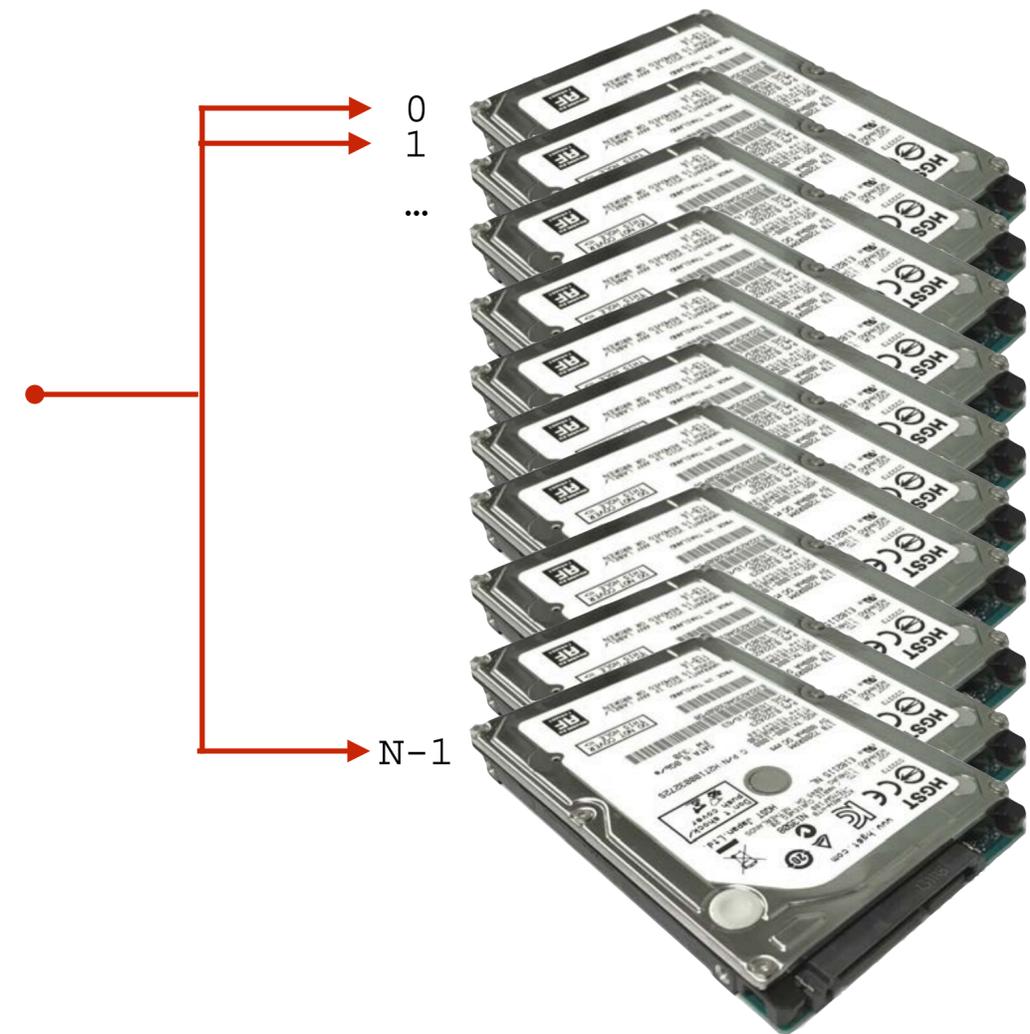


# Consequence(s) of the choice

Mount points for the data disks (jive5ab, DEFAULT startup)

```
/mnt/disk0  
/mnt/disk1  
...  
/mnt/diskN-1
```

*jive5ab looks for those at startup  
does NOT mount them!*



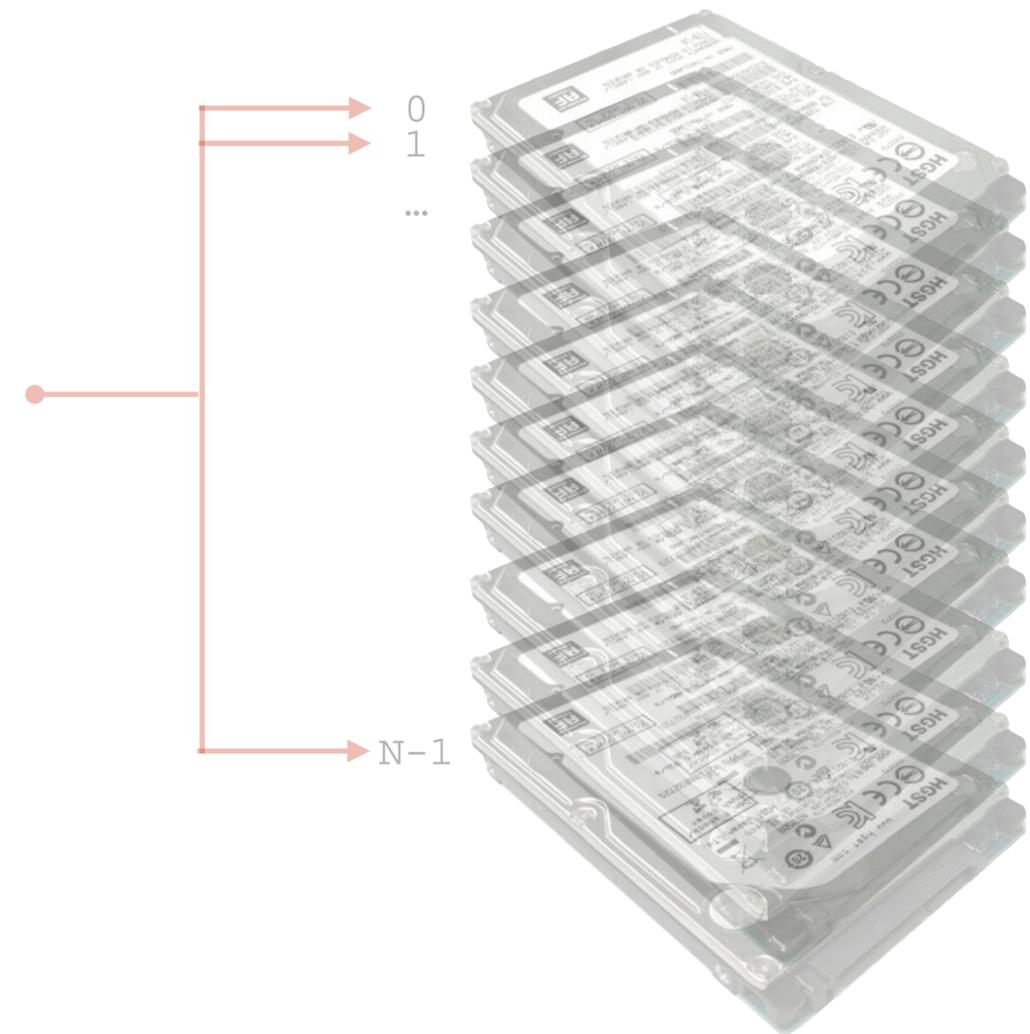
# Consequence(s) of the choice

Mount points for the data disks (jive5ab, DEFAULT startup)

```
/mnt/disk0  
/mnt/disk1  
...  
/mnt/diskN-1
```

*jive5ab looks for those at startup  
does NOT mount them!*

regex: /mnt/disk[0-9]+

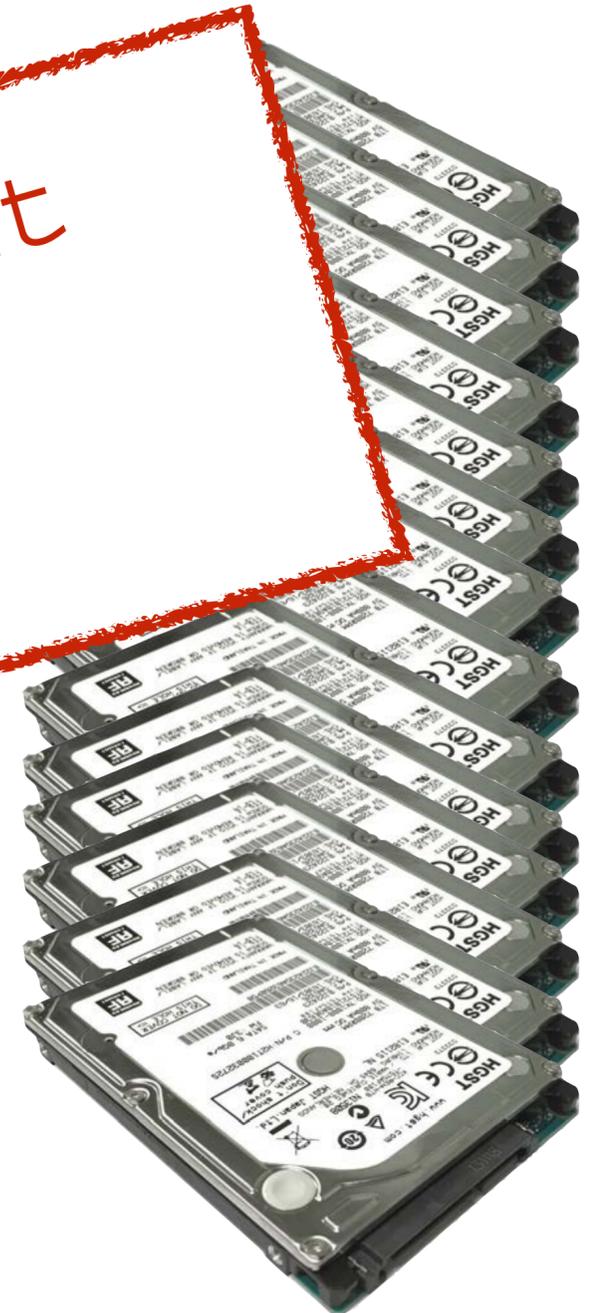
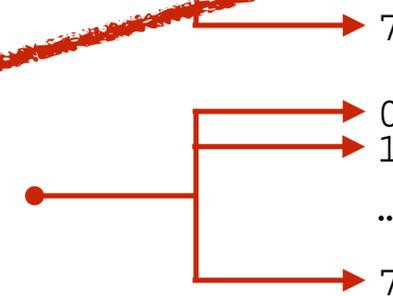


# Consequence(s) of the choice

Mount points for the data disks (jive5ab, -6 command line option)

```
/mnt/disk/1/0/data  
  /1/1/data  
  ...  
  /1/7/d  
/mnt/disk  
/mnt/disk
```

jive5ab looks for those at  
startup  
does NOT mount them!



# Consequence(s) of the choice

Mount points for the data disks (jive5ab, \*)

```
set_disks? ; (*)
```

```
⇒ !set_disks? 0 : /mnt/disk0 : /mnt/disk1 : ... ;
```

```
set_disks = 12 : /path/to/sd* ; (*)
```

```
⇒ !set_disks = 0 : 18 ;
```

```
set_disks? ; (*)
```

```
⇒ !set_disks? 0 : /mnt/disk/1/0/data : .. : /path/to/sdA : ... ;
```

(\*) <https://github.com/jive-vlbi/jive5ab/blob/master/doc/jive5ab-documentation-1.11.pdf>

# Consequence(s) of the choice

Mount points for the data disks (jive5ab, command line script)

```
$> m6sg_mount
```

(\* ) [https://github.com/jive-vlbi/jive5ab/blob/master/scripts/m6sg\\_mount](https://github.com/jive-vlbi/jive5ab/blob/master/scripts/m6sg_mount)

# Consequence(s) of the choice

Mount points for the data disks (jive5ab, command line script)

```
$> m6sg_mount 134
```



(\* ) [https://github.com/jive-vlbi/jive5ab/blob/master/scripts/m6sg\\_mount](https://github.com/jive-vlbi/jive5ab/blob/master/scripts/m6sg_mount)

# Consequence(s) of the choice

Mount points for the data disks (jive5ab, command line script)

```
$> m6sg_mount -u 3
```



(\* ) [https://github.com/jive-vlbi/jive5ab/blob/master/scripts/m6sg\\_mount](https://github.com/jive-vlbi/jive5ab/blob/master/scripts/m6sg_mount)

# Consequence(s) of the choice

How the packets are read from the network

# Consequence(s) of the choice

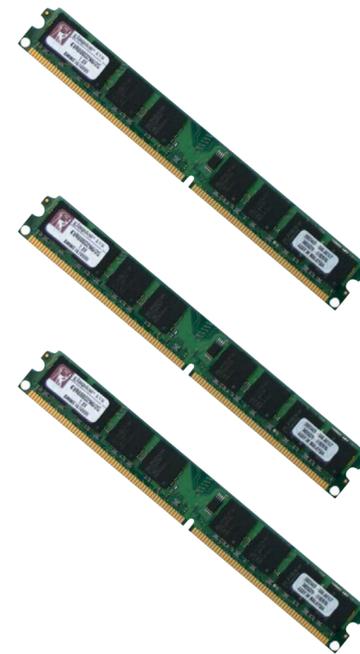
How the packets are read from the network (cplane / dplane)



➔  
\$> tcpdump -c 10M /dev/eth1



➔  
\$> tcpdump -c 10M /dev/ethN



10 MByte  
blocks

# Consequence(s) of the choice

How the packets are read from the network (cplane / dplane)



\$> tcpdump

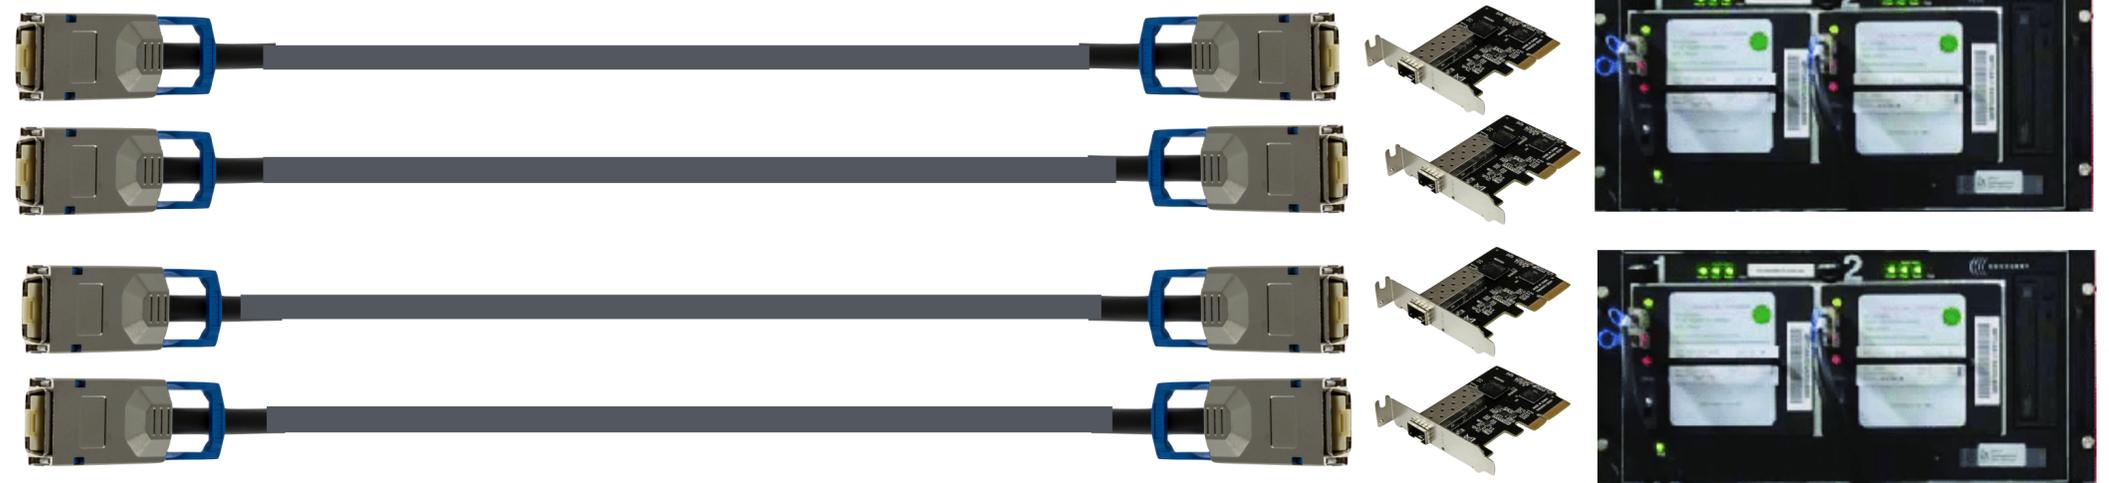
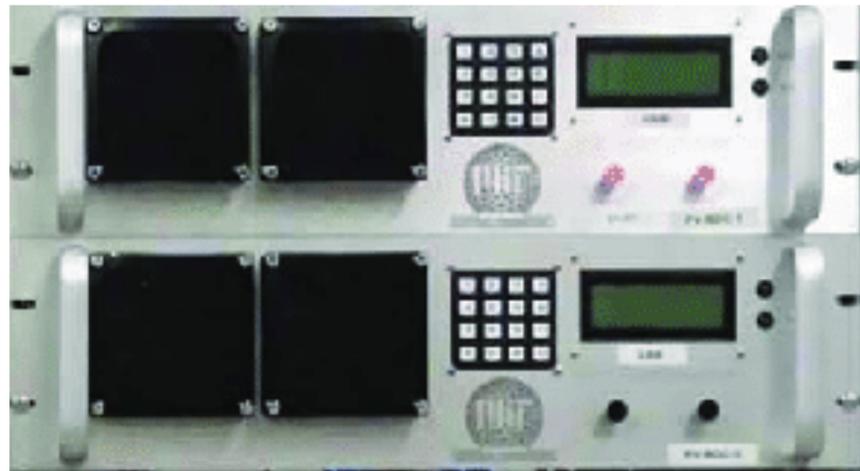
No valid network configuration necessary  
on the network cards or digital back end!



10 MByte  
blocks

# Consequence(s) of the choice

How the packets are read from the network (cplane / dplane)



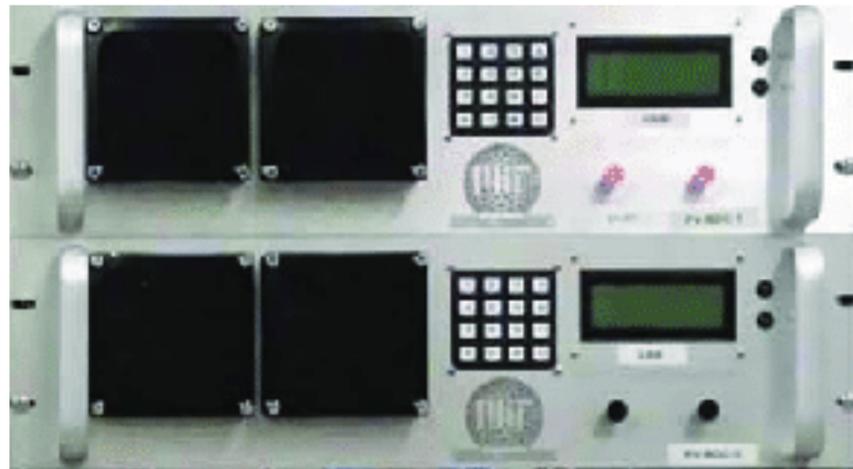
Configuration: 😊

Flexibility: 😞

# Consequence(s) of the choice

How the packets are read from the network (cplane / dplane)

BAND A, R+L



R+R



BAND B, R+L

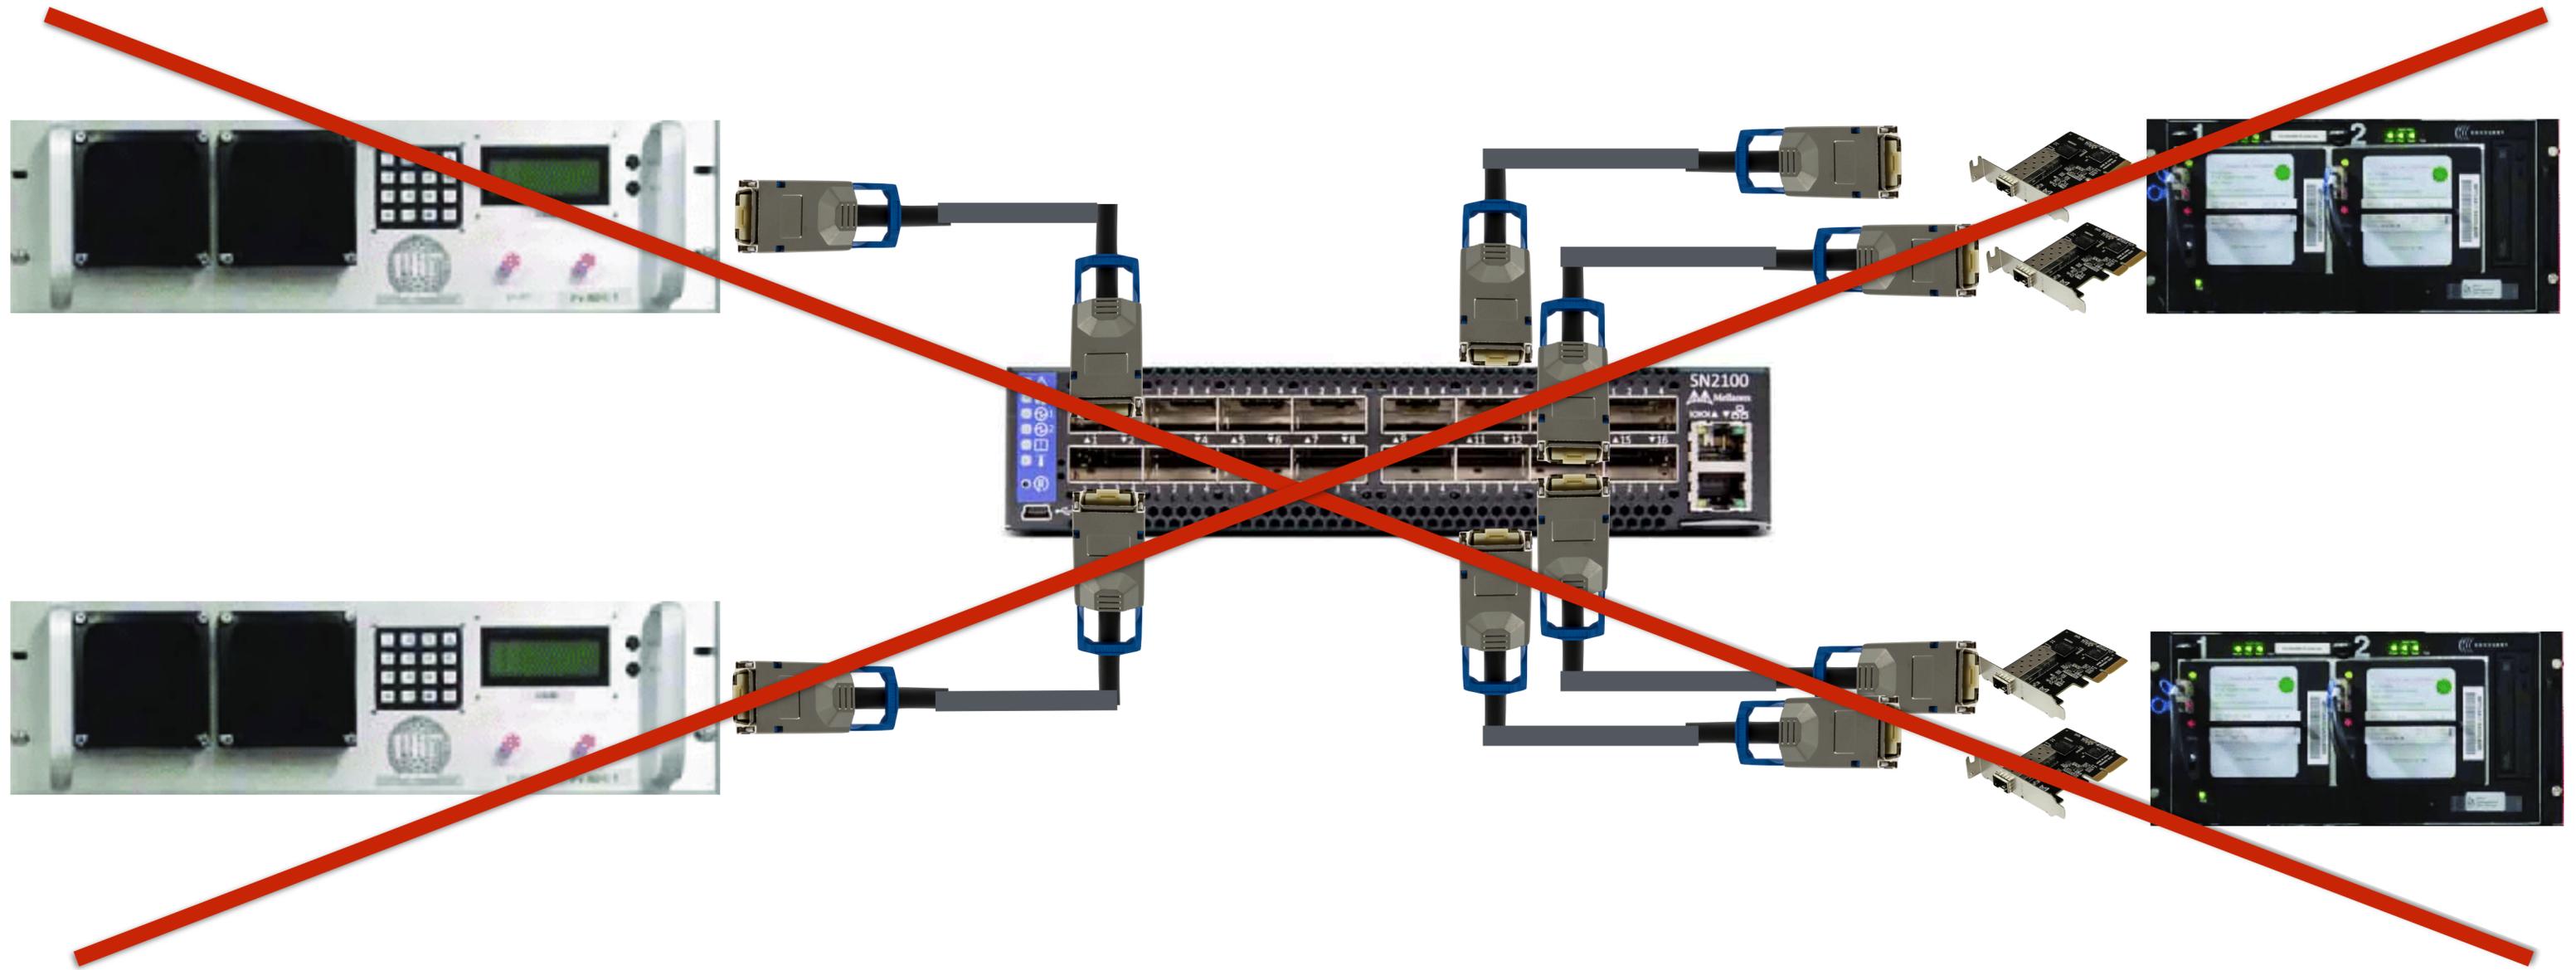
L+L

Installation: 😊

Flexibility: 😞

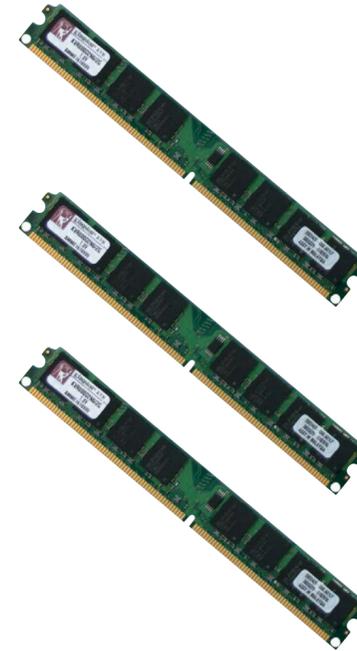
# Consequence(s) of the choice

How the packets are read from the network (cplane / dplane)



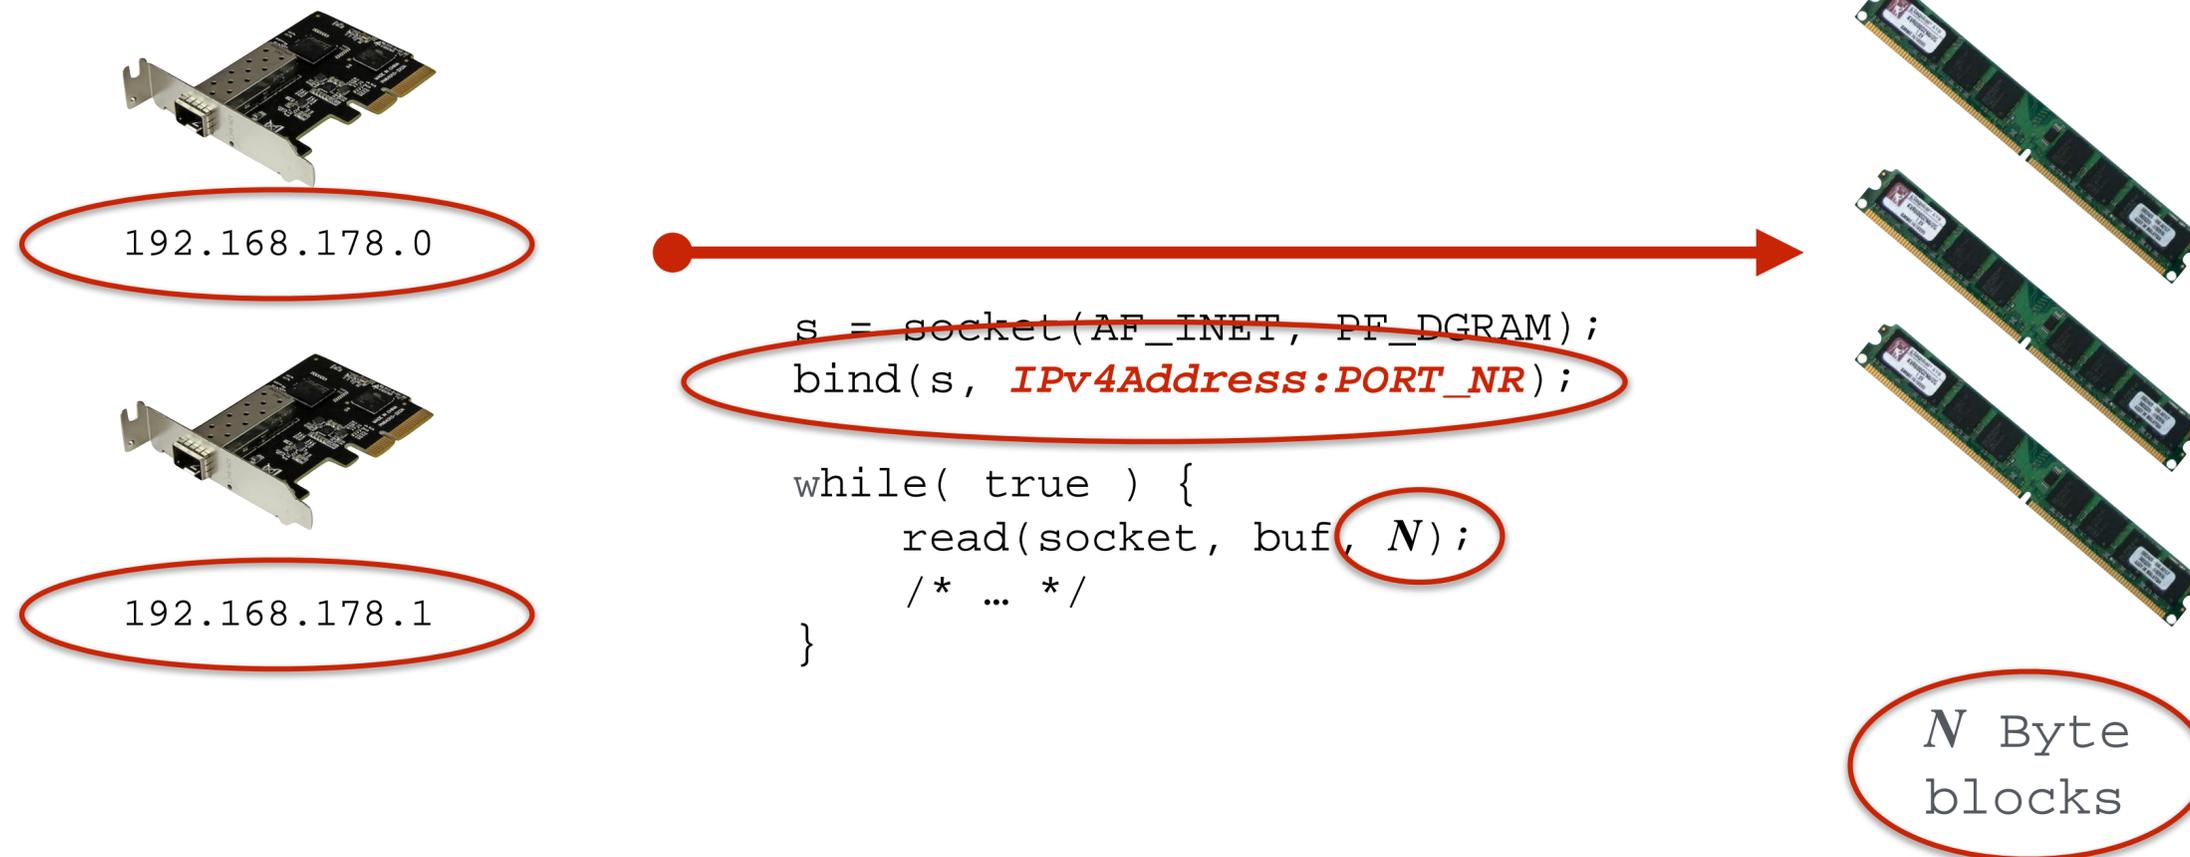
# Consequence(s) of the choice

How the packets are read from the network (jive5ab)



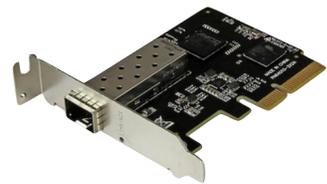
# Consequence(s) of the choice

How the packets are read from the network (jive5ab)



# Consequence(s) of the choice

How the packets are read from the network (jive5ab)



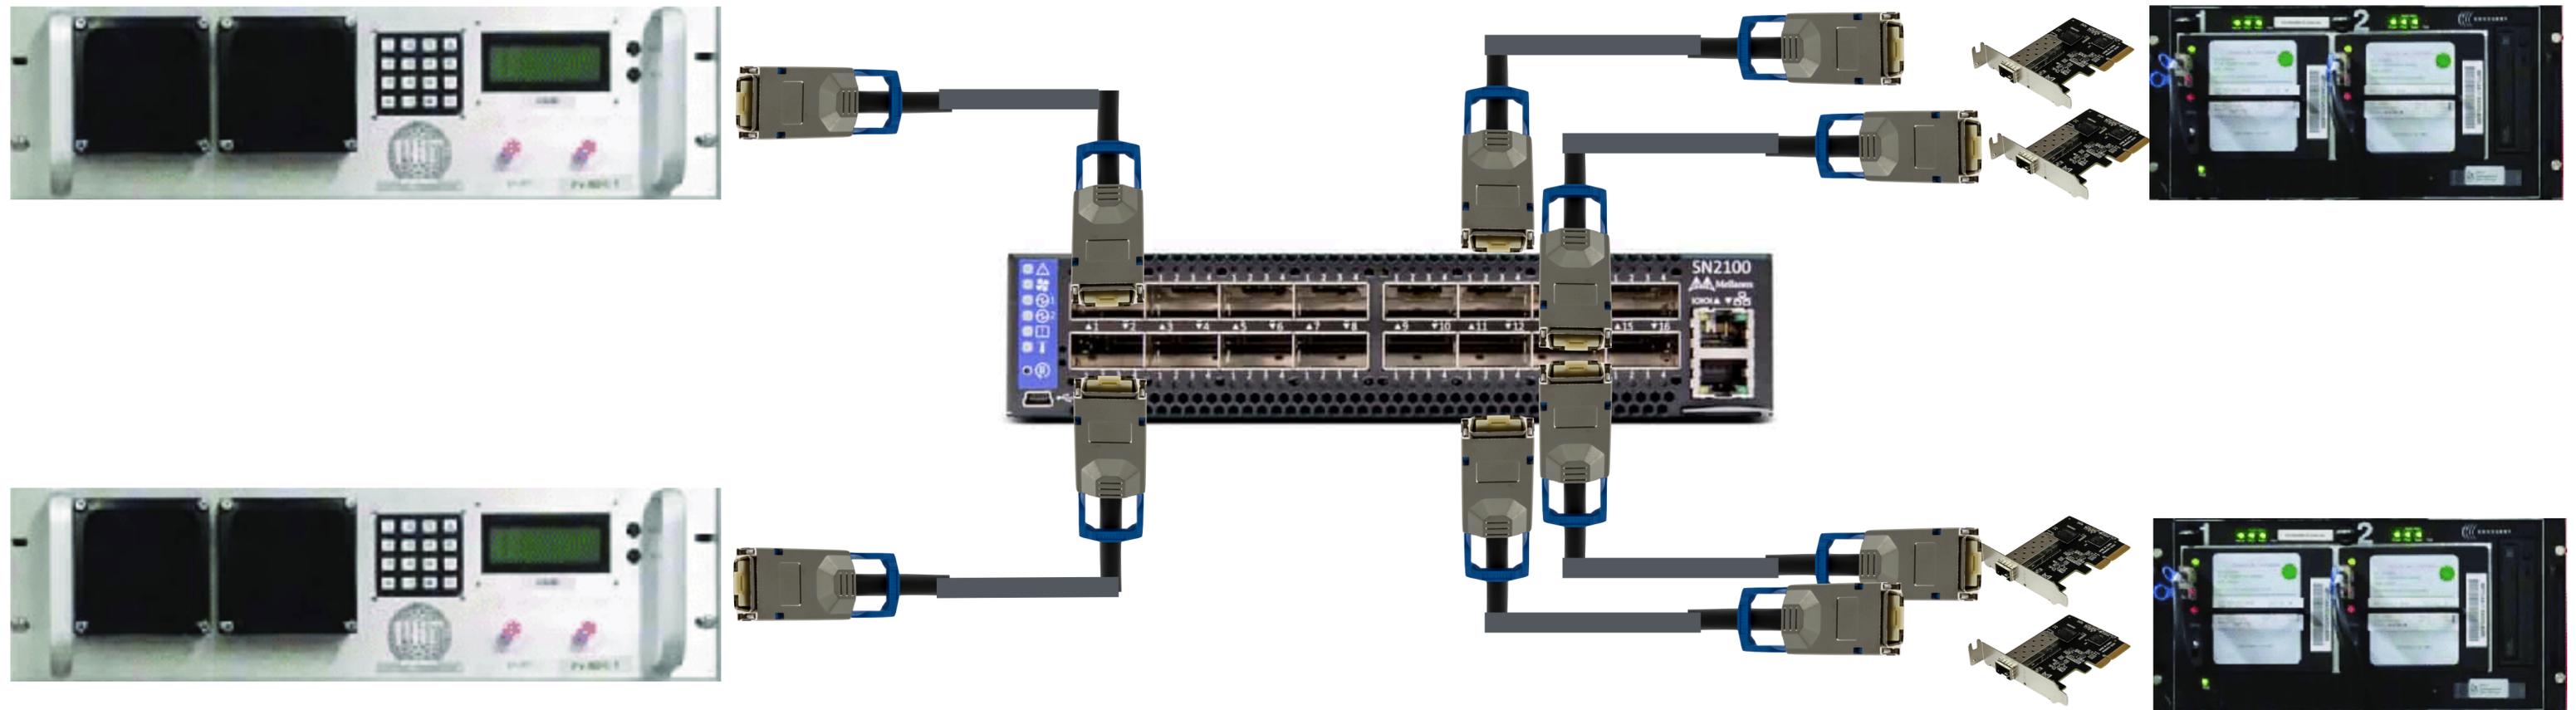
192.168.178.0

**REQUIRE** valid network configuration on  
the network cards and digital back ends!

*N* Byte  
blocks

# Consequence(s) of the choice

How the packets are read from the network (jive5ab)



Configuration: 😞

Flexibility: 😊

# Consequence(s) of the choice

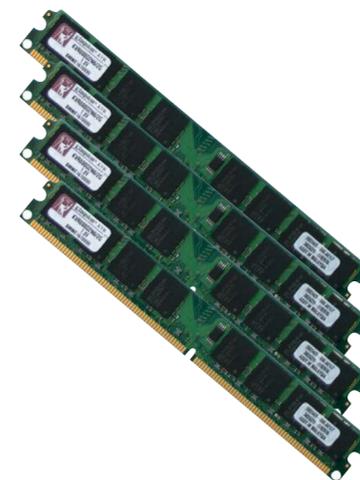
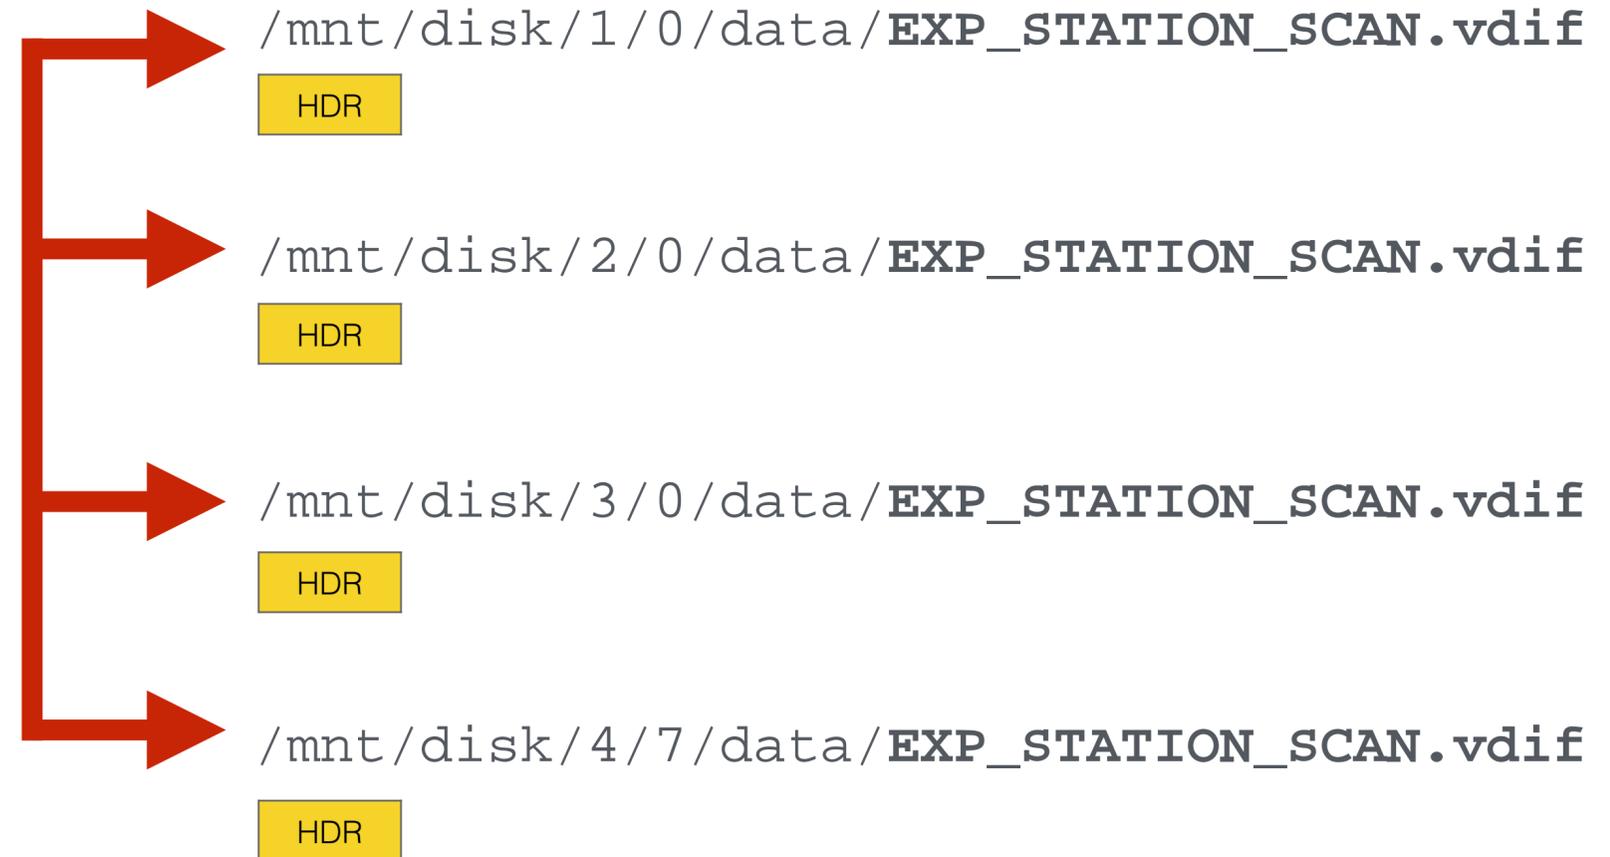
What gets written to disk

# Consequence(s) of the choice

What gets written to disk (cplane / dplane Mark6 format)



ethernet cards

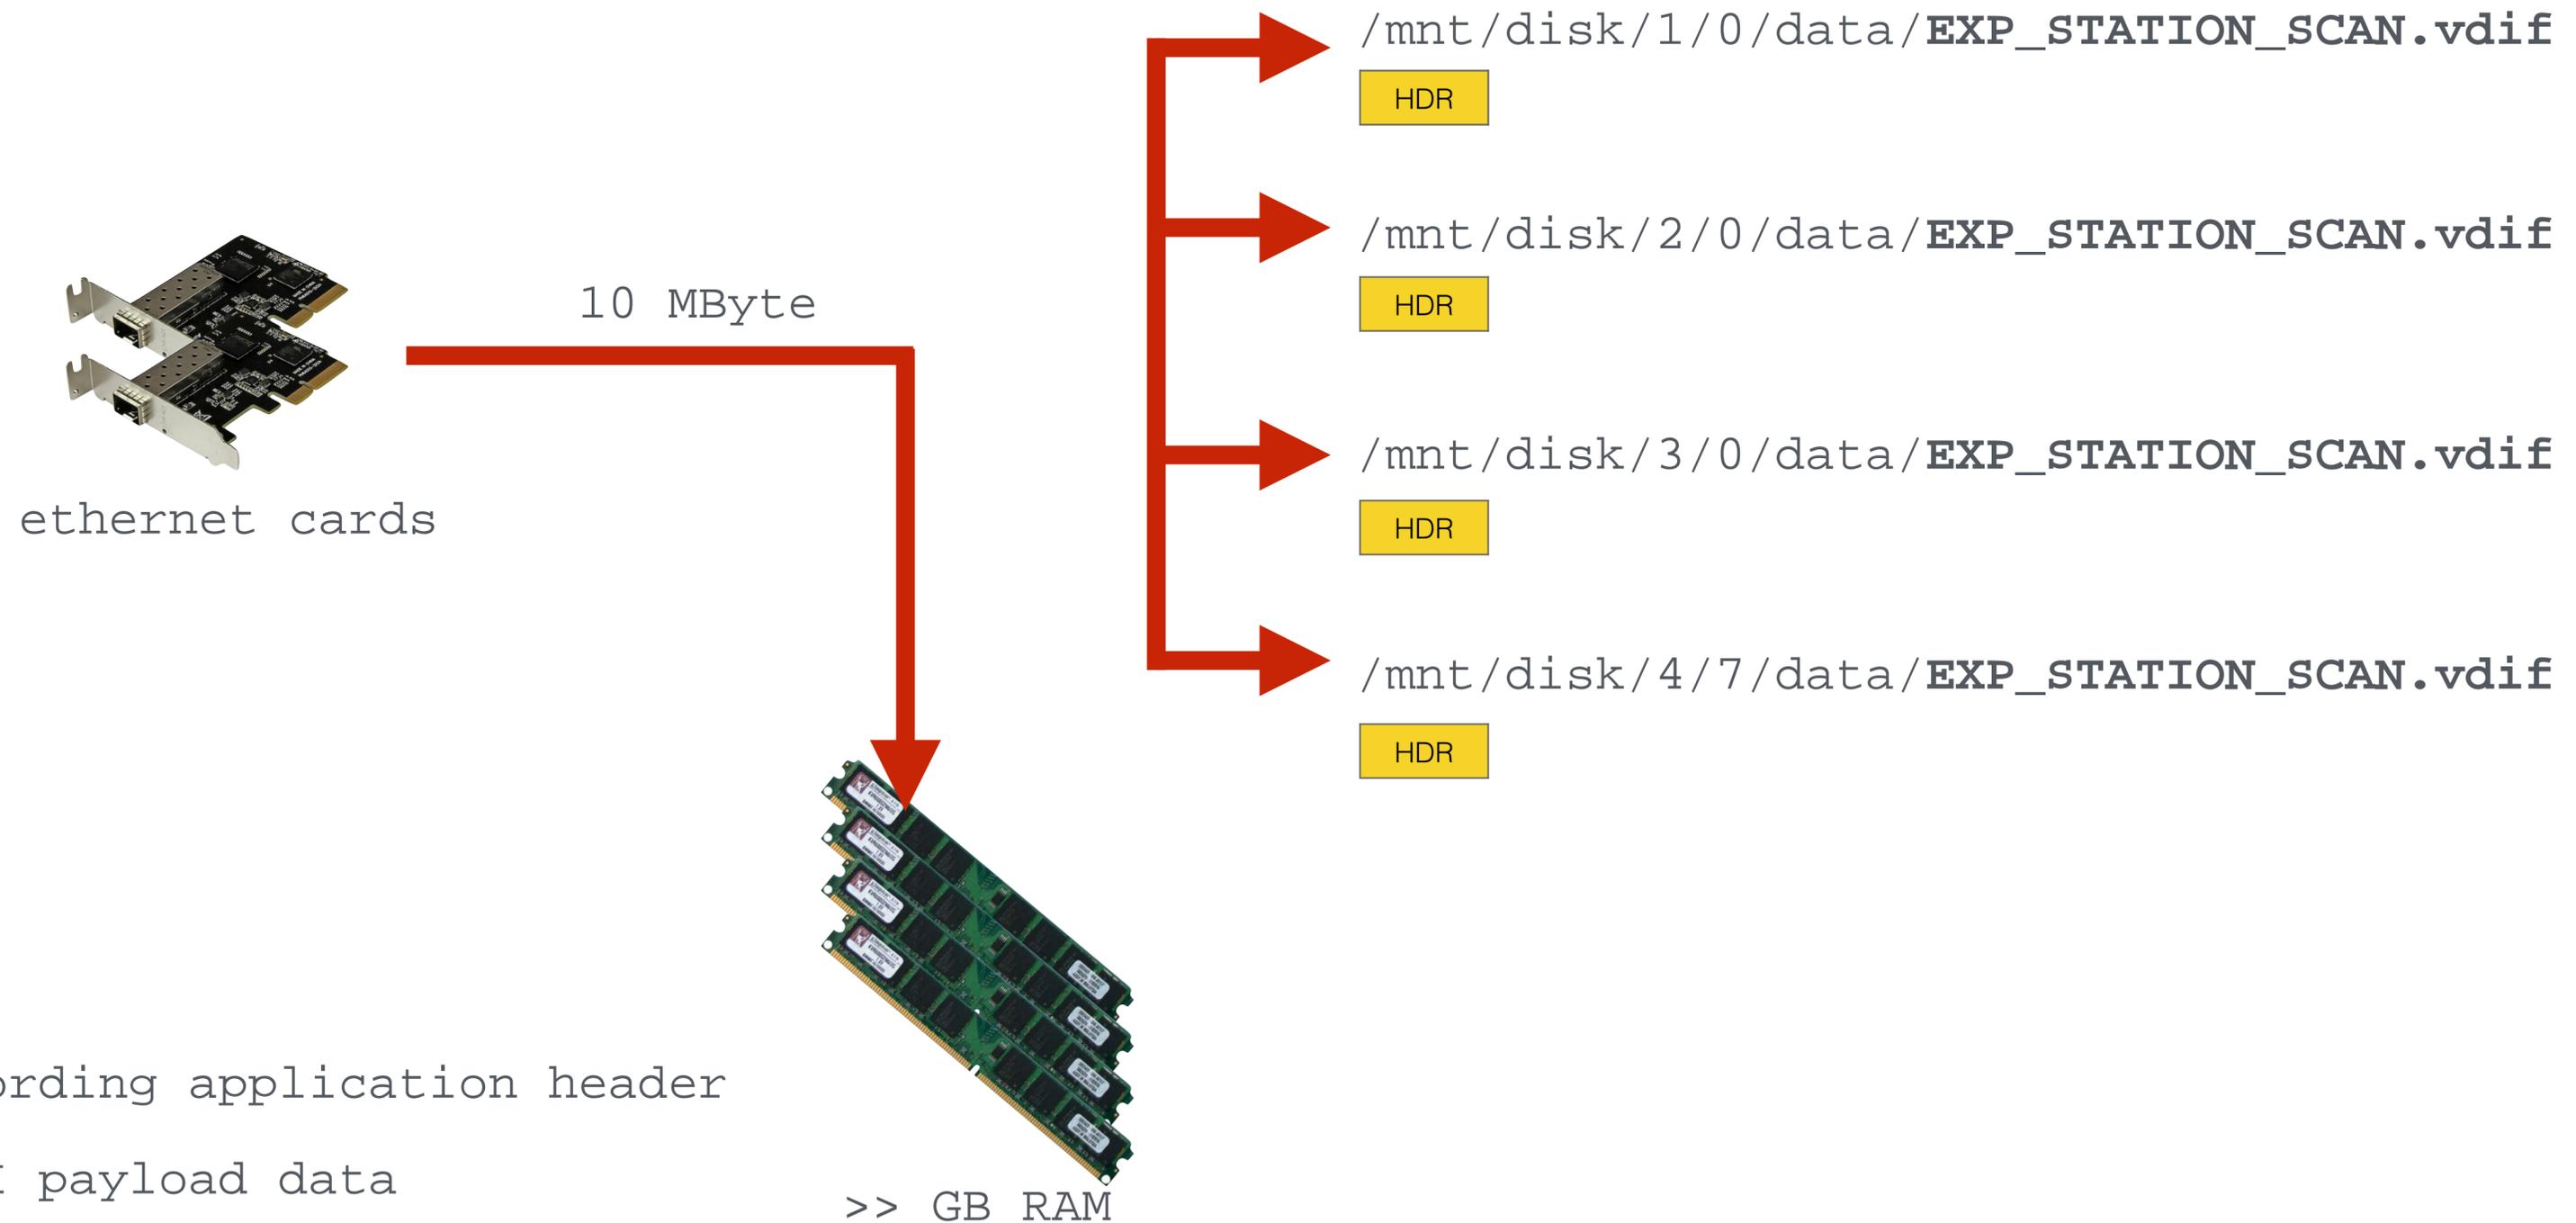


>> GB RAM

-  Recording application header
-  VLBI payload data

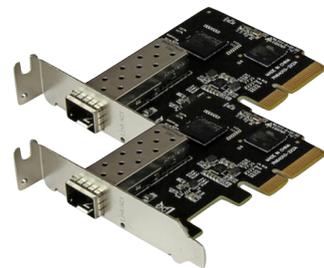
# Consequence(s) of the choice

What gets written to disk (cplane / dplane Mark6 format)

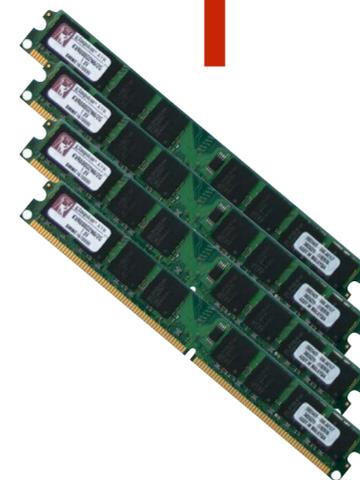


# Consequence(s) of the choice

What gets written to disk (cplane / dplane Mark6 format)



ethernet cards



>> GB RAM

/mnt/disk/1/0/data/EXP\_STATION\_SCAN.vdif

HDR 0

/mnt/disk/2/0/data/EXP\_STATION\_SCAN.vdif

HDR

/mnt/disk/3/0/data/EXP\_STATION\_SCAN.vdif

HDR

/mnt/disk/4/7/data/EXP\_STATION\_SCAN.vdif

HDR



Recording application header



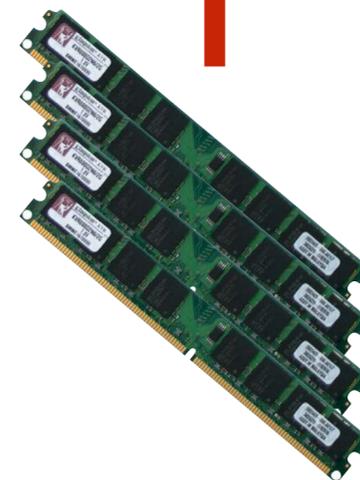
VLBI payload data

# Consequence(s) of the choice

What gets written to disk (cplane / dplane Mark6 format)



ethernet cards



>> GB RAM

/mnt/disk/1/0/data/EXP\_STATION\_SCAN.vdif



/mnt/disk/2/0/data/EXP\_STATION\_SCAN.vdif



/mnt/disk/3/0/data/EXP\_STATION\_SCAN.vdif



/mnt/disk/4/7/data/EXP\_STATION\_SCAN.vdif



Recording application header



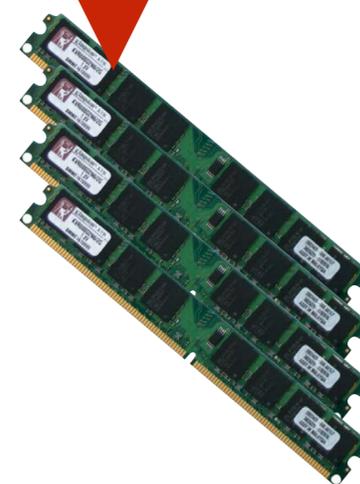
VLBI payload data

# Consequence(s) of the choice

What gets written to disk (cplane / dplane Mark6 format)



10 MByte



/mnt/disk/1/0/data/EXP\_STATION\_SCAN.vdif



/mnt/disk/2/0/data/EXP\_STATION\_SCAN.vdif



/mnt/disk/3/0/data/EXP\_STATION\_SCAN.vdif



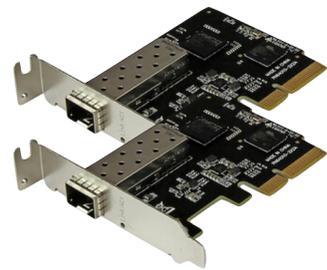
/mnt/disk/4/7/data/EXP\_STATION\_SCAN.vdif



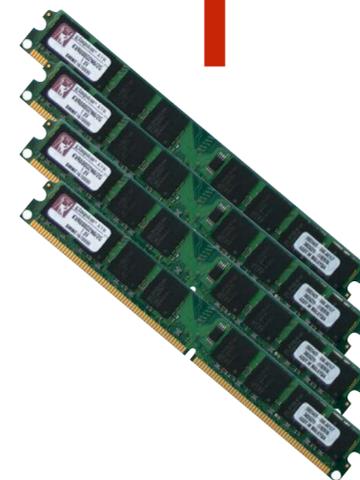
-  Recording application header
-  VLBI payload data

# Consequence(s) of the choice

What gets written to disk (cplane / dplane Mark6 format)



ethernet cards



>> GB RAM

/mnt/disk/1/0/data/EXP\_STATION\_SCAN.vdif



/mnt/disk/2/0/data/EXP\_STATION\_SCAN.vdif



/mnt/disk/3/0/data/EXP\_STATION\_SCAN.vdif



/mnt/disk/4/7/data/EXP\_STATION\_SCAN.vdif



Recording application header



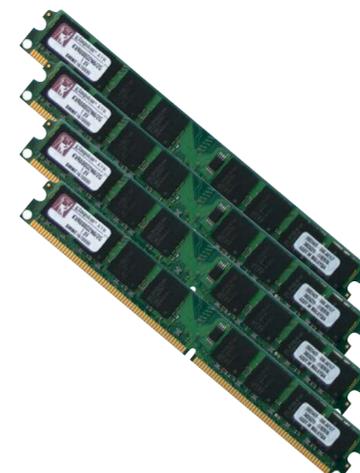
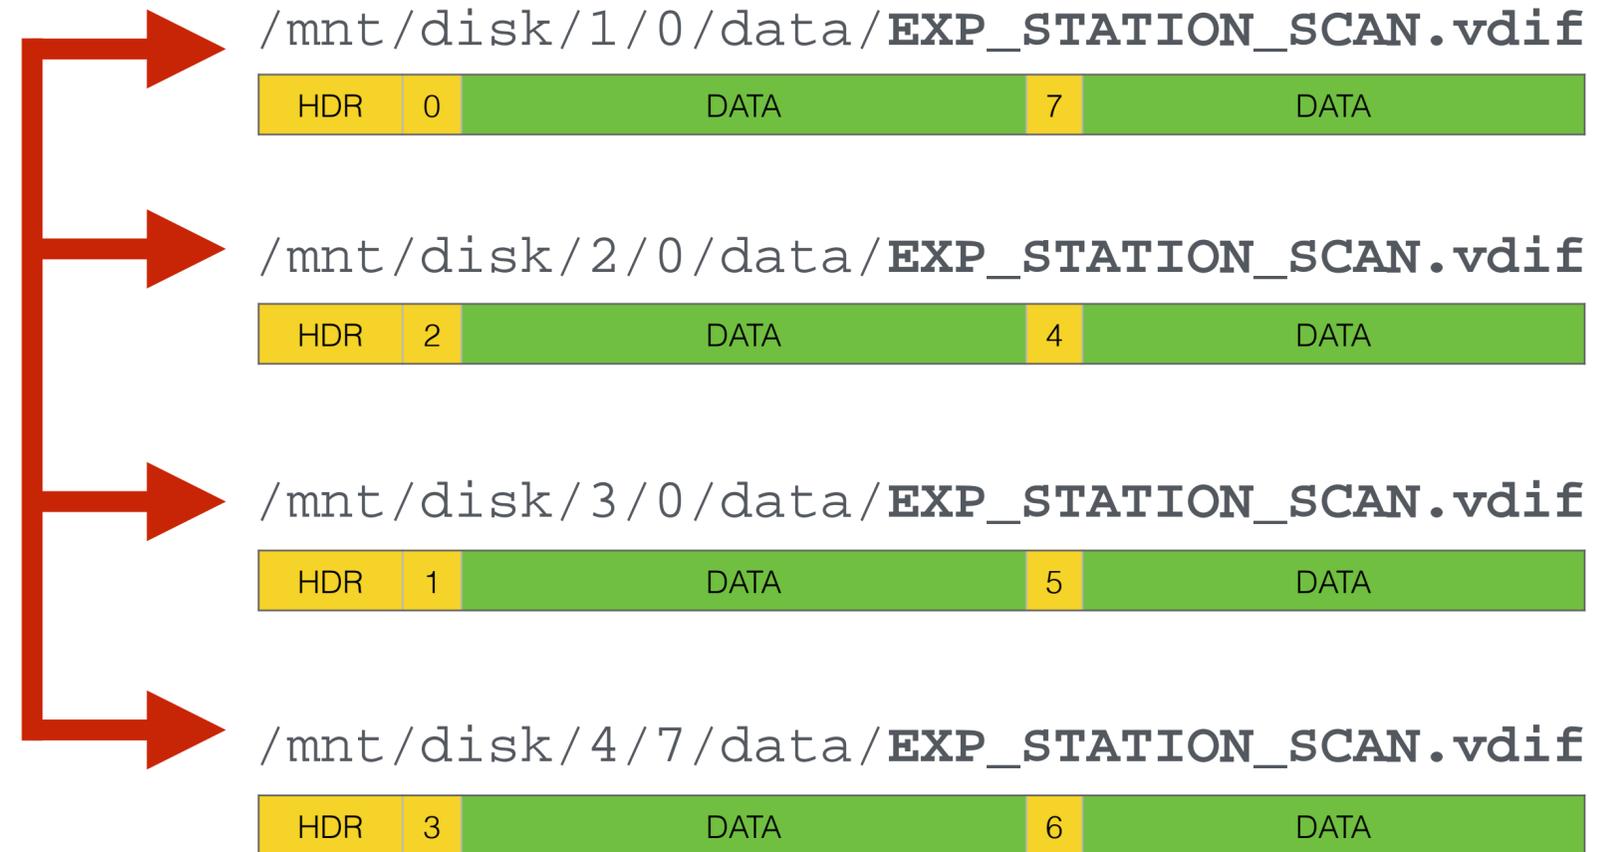
VLBI payload data

# Consequence(s) of the choice

What gets written to disk (cplane / dplane Mark6 format)



ethernet cards



>> GB RAM

 Recording application header  
 VLBI payload data

# Consequence(s) of the choice

What gets written to disk (cplane / dplane Mark6 format)

```
/mnt/disk/1/0/data/EXP_STATION_SCAN.vdif
```

```
/mnt/disk/2/0/data/EXP_STATION_SCAN.vdif
```

```
/mnt/disk/3/0/data/EXP_STATION_SCAN.vdif
```

```
/mnt/disk/4/7/data/EXP_STATION_SCAN.vdif
```

## Inconvenient for e-shipping

 Recording application header

 VLBI payload data

# Consequence(s) of the choice

What gets written to disk (cplane / dplane Mark6 format)

/mnt/disk/1/0/data/EXP\_STATION\_SCAN.vdif



/mnt/disk/2/0/data/EXP\_STATION\_SCAN.vdif



/mnt/disk/3/0/data/EXP\_STATION\_SCAN.vdif



/mnt/disk/4/7/data/EXP\_STATION\_SCAN.vdif



VDIF content cannot be used directly

 Recording application header

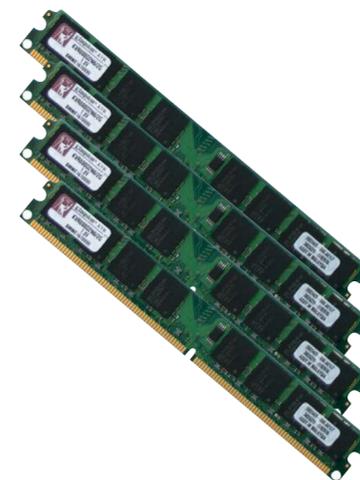
 VLBI payload data

# Consequence(s) of the choice

What gets written to disk (jive5ab, FlexBuff/vbs format (DEFAULT))



ethernet cards

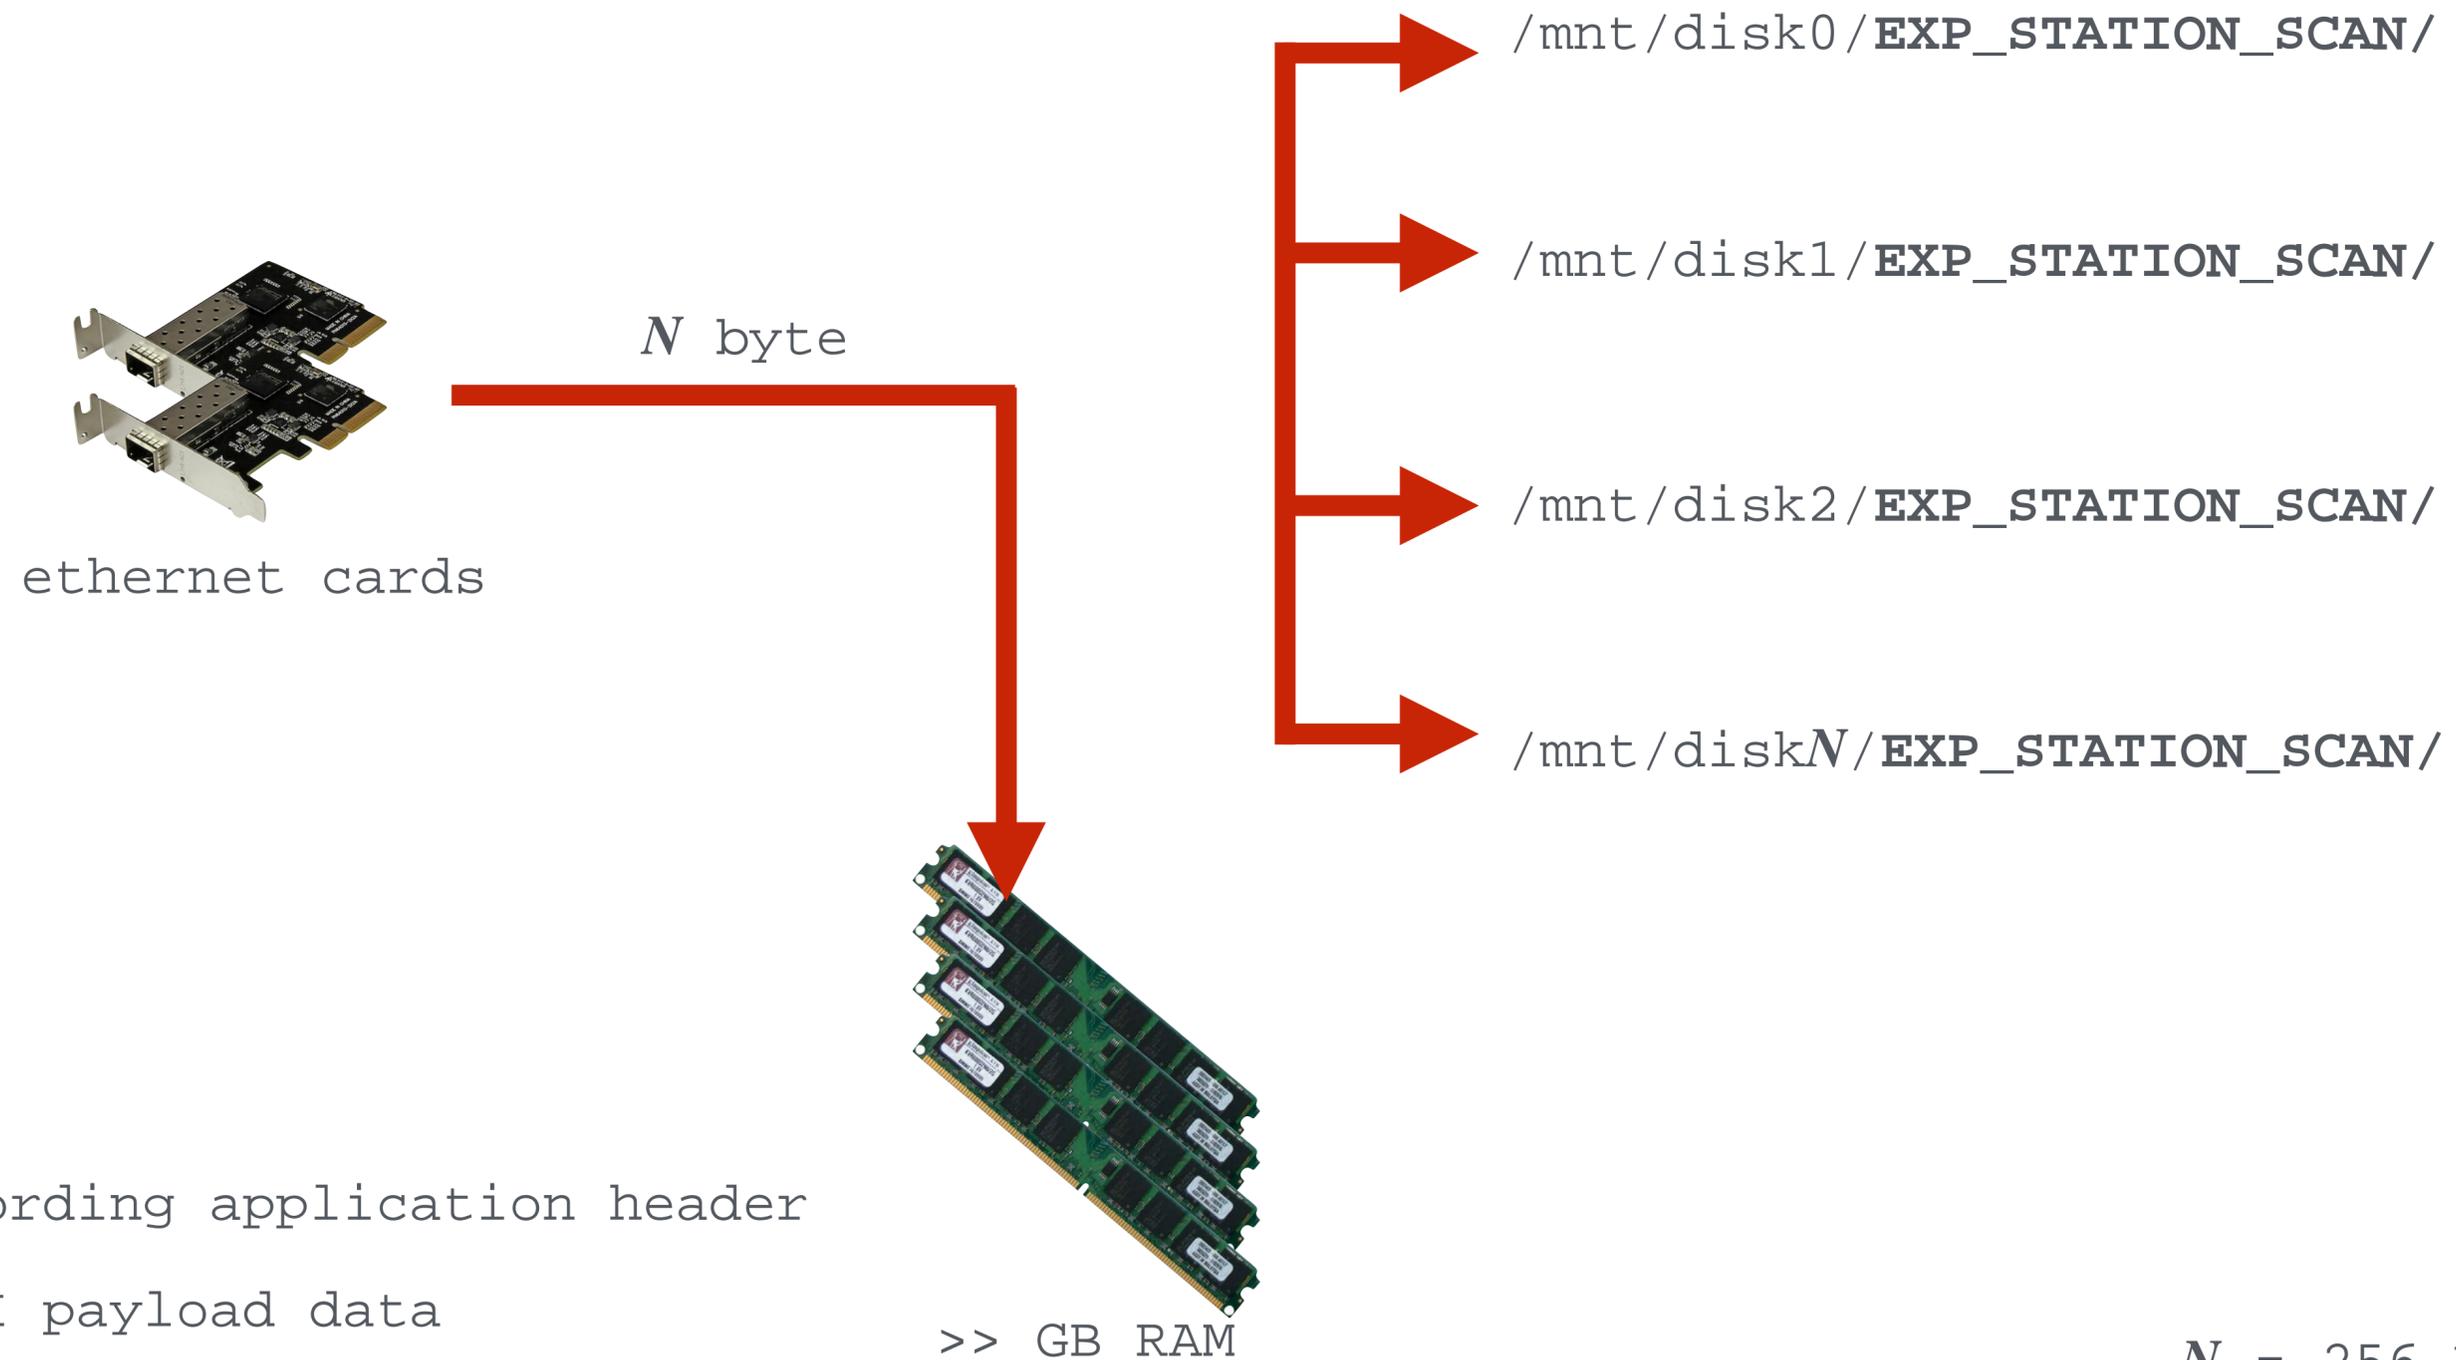


>> GB RAM

-  Recording application header
-  VLBI payload data

# Consequence(s) of the choice

What gets written to disk (jive5ab, FlexBuff/vbs format (DEFAULT))



-  Recording application header
-  VLBI payload data

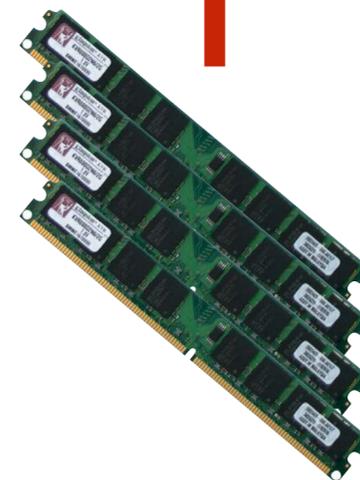
$N = 256$  MB typically

# Consequence(s) of the choice

What gets written to disk (jive5ab, FlexBuff/vbs format (DEFAULT))



ethernet cards



>> GB RAM

/mnt/disk0/EXP\_STATION\_SCAN/  
EXP\_STATION\_SCAN.0000000

/mnt/disk1/EXP\_STATION\_SCAN/

/mnt/disk2/EXP\_STATION\_SCAN/

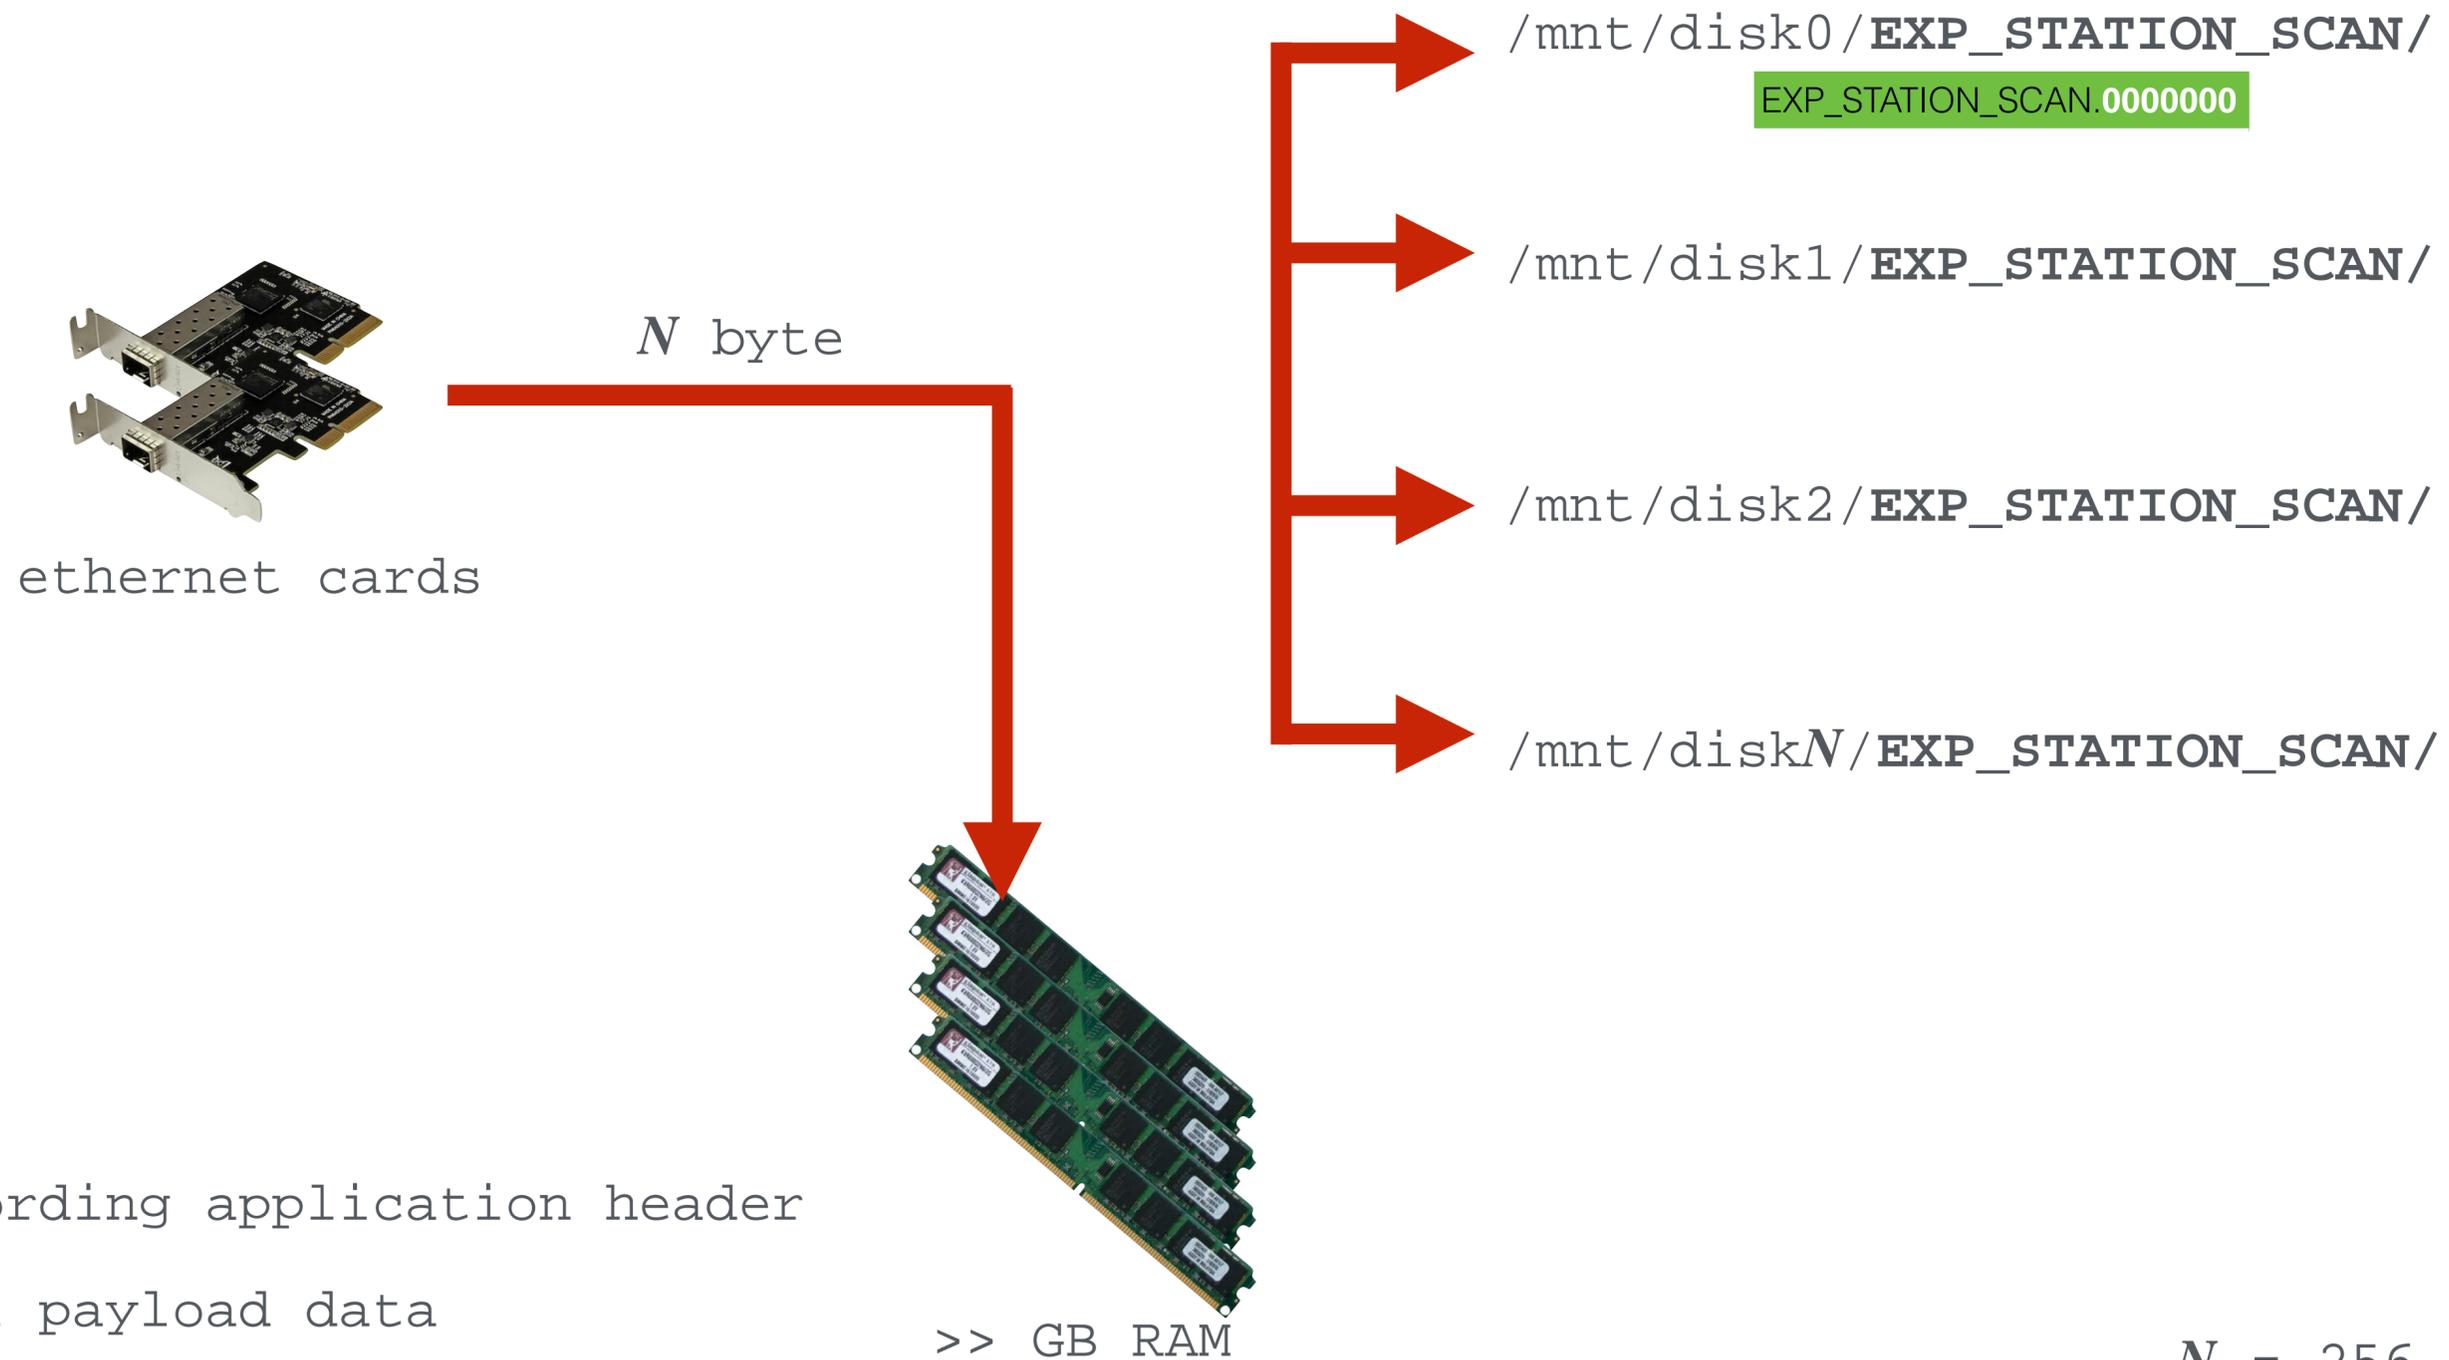
/mnt/diskN/EXP\_STATION\_SCAN/

-  Recording application header
-  VLBI payload data

$N = 256$  MB typically

# Consequence(s) of the choice

What gets written to disk (jive5ab, FlexBuff/vbs format (DEFAULT))



-  Recording application header
-  VLBI payload data

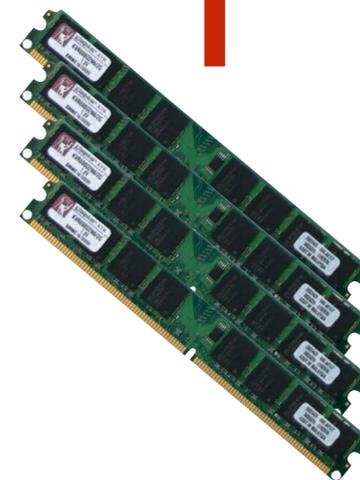
$N = 256$  MB typically

# Consequence(s) of the choice

What gets written to disk (jive5ab, FlexBuff/vbs format (DEFAULT))



ethernet cards



>> GB RAM

/mnt/disk0/EXP\_STATION\_SCAN/  
EXP\_STATION\_SCAN.0000000

/mnt/disk1/EXP\_STATION\_SCAN/

/mnt/disk2/EXP\_STATION\_SCAN/  
EXP\_STATION\_SCAN.0000001

/mnt/diskN/EXP\_STATION\_SCAN/

-  Recording application header
-  VLBI payload data

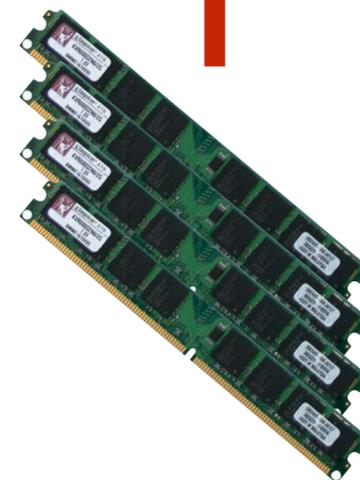
$N = 256$  MB typically

# Consequence(s) of the choice

What gets written to disk (jive5ab, FlexBuff/vbs format (DEFAULT))



ethernet cards



>> GB RAM

/mnt/disk0/**EXP\_STATION\_SCAN/**

EXP\_STATION\_SCAN.0000000

EXP\_STATION\_SCAN.0000007

/mnt/disk1/**EXP\_STATION\_SCAN/**

EXP\_STATION\_SCAN.0000002

EXP\_STATION\_SCAN.0000004

/mnt/disk2/**EXP\_STATION\_SCAN/**

EXP\_STATION\_SCAN.0000001

EXP\_STATION\_SCAN.0000006

/mnt/diskN/**EXP\_STATION\_SCAN/**

EXP\_STATION\_SCAN.0000003

EXP\_STATION\_SCAN.0000005



Recording application header



VLBI payload data

$N = 256$  MB typically

# Consequence(s) of the choice

What gets written to disk (jive5ab, FlexBuff/vbs format (DEFAULT))

```
/mnt/disk0/EXP_STATION_SCAN/
```

```
EXP_STATION_SCAN.0000000
```

```
EXP_STATION_SCAN.0000007
```

```
/mnt/disk1/EXP_STATION_SCAN/
```

```
EXP_STATION_SCAN.0000002
```

```
EXP_STATION_SCAN.0000004
```

```
/mnt/disk2/EXP_STATION_SCAN/
```

```
EXP_STATION_SCAN.0000001
```

```
EXP_STATION_SCAN.0000006
```

```
/mnt/diskN/EXP_STATION_SCAN/
```

```
EXP_STATION_SCAN.0000003
```

```
EXP_STATION_SCAN.0000005
```

Unique file names: e-transfer ✓



Recording application header



VLBI payload data

# Consequence(s) of the choice

What gets written to disk (jive5ab, FlexBuff/vbs format (DEFAULT))

```
/mnt/disk0/EXP_STATION_SCAN/
```

```
EXP_STATION_SCAN.0000000
```

```
EXP_STATION_SCAN.0000007
```

```
/mnt/disk1/EXP_STATION_SCAN/
```

```
EXP_STATION_SCAN.0000002
```

```
EXP_STATION_SCAN.0000004
```

```
/mnt/disk2/EXP_STATION_SCAN/
```

```
EXP_STATION_SCAN.0000001
```

```
EXP_STATION_SCAN.0000006
```

```
/mnt/diskN/EXP_STATION_SCAN/
```

```
EXP_STATION_SCAN.0000003
```

```
EXP_STATION_SCAN.0000005
```

Only VLBI data: can use any VDIF tool



Recording application header



VLBI payload data

# Consequence(s) of the choice

What gets written to disk (jive5ab - you have a choice)

compiled in default:

```
vbs (FlexBuff) format
```

command line: set default format

```
$> jive5ab --format mk6|flexbuff
```

runtime: set format (VSI/S)

```
record = mk6 : 0|1 ;
```

# FlexBuff how-to

1. buy/repurpose machine
2. install+tune operating system (any POSIX)
3. connect (many) disks
4. can mount as /mnt/diskNNN?  
yes: done  
no: remember path/regex  
in FS `jive5ab.ctl` add `set_disks= path/regex;`
5. configure network card(s)
6. get + build jive5ab  

```
$> git clone https://github.com/jive-vlbi/jive5ab.git
$> mkdir build && cd build && cmake -DSSAPI_ROOT=nossapi ..
$> make [-j <n_cpu>] [VERBOSE=1]
```
7. profit! `$> jive5ab -m3 [options]`

# Tuning

the key to high speed

packet capture

# CPU/core limitations

Achieving *very* high capture rates

# NO HYPERTHREADING

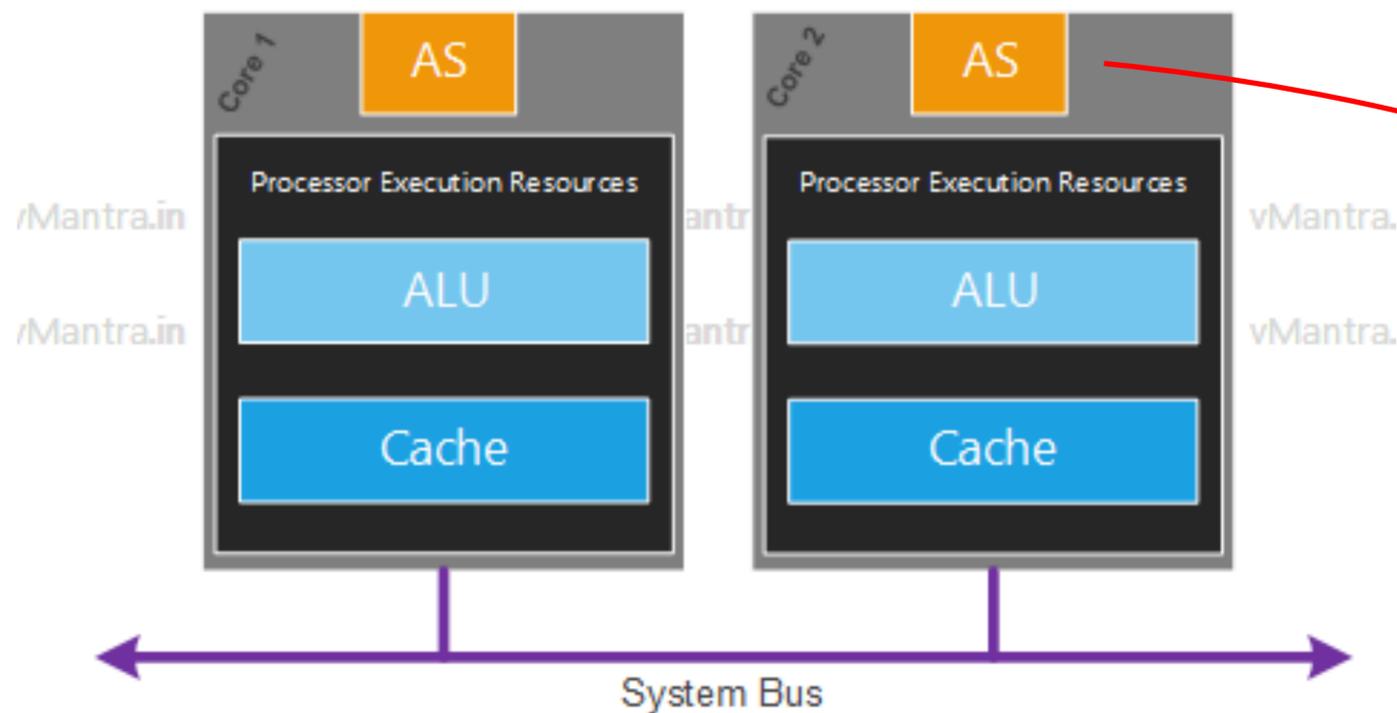
# CPU/core limitations

Achieving *very* high capture rates

# CPU/core limitations

Achieving *very* high capture rates

Dual Core Processor without Hyper-Threading

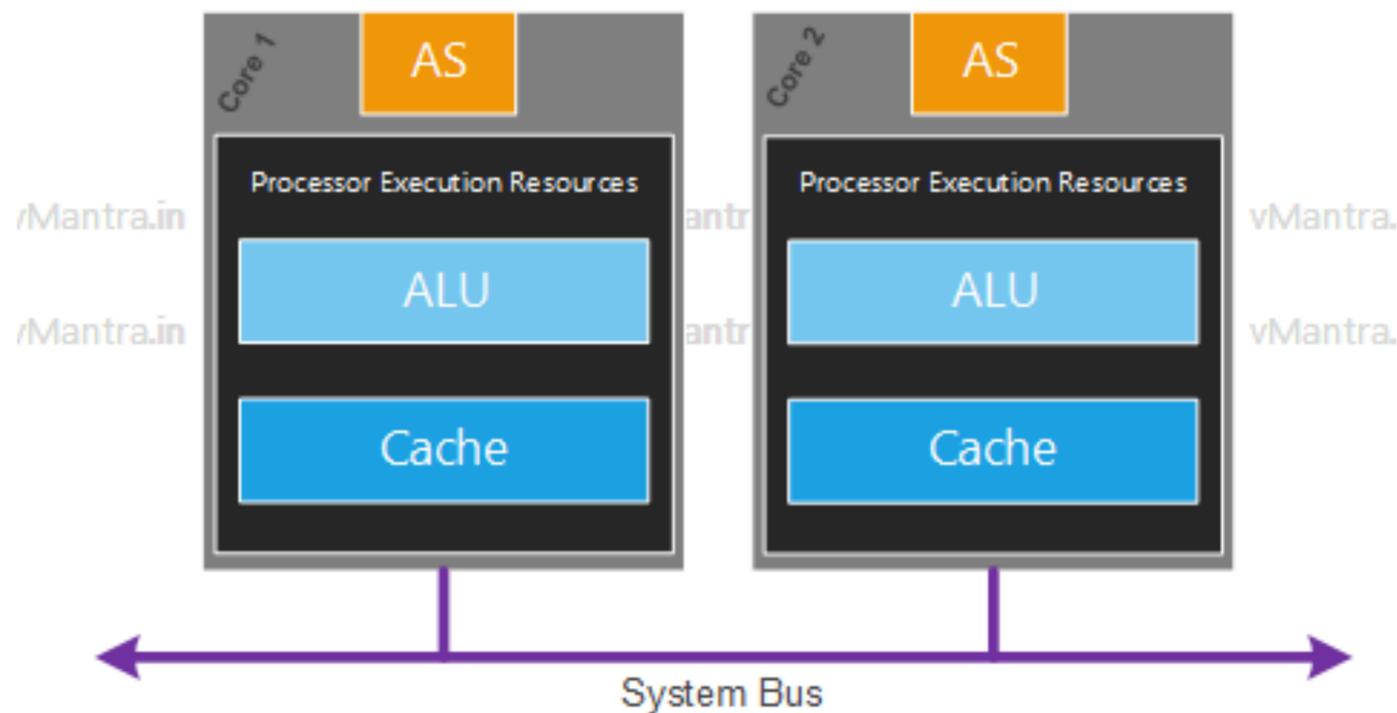


AS = Architecture State  
which means  
"CPU registers and flags &cet."

# CPU/core limitations

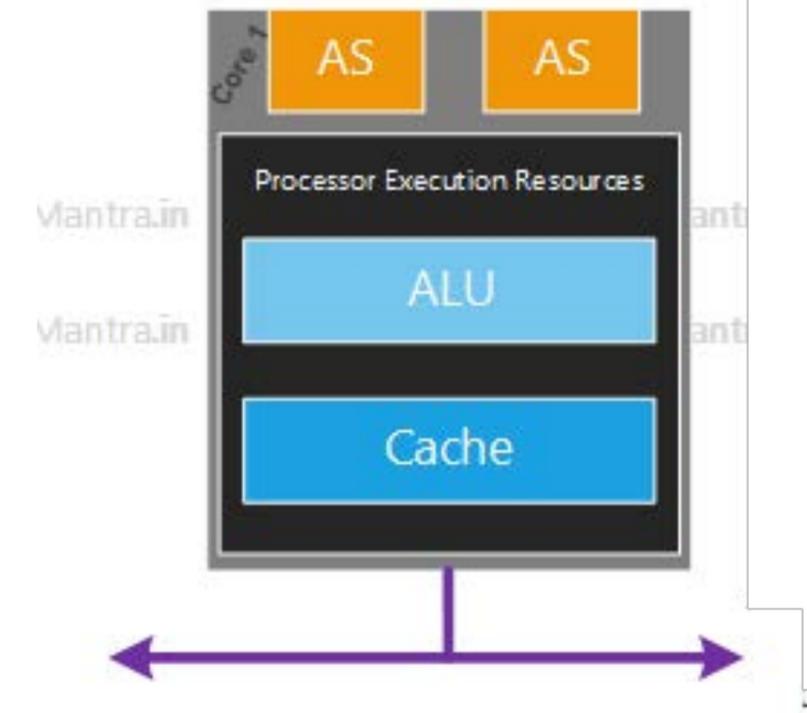
Achieving *very* high capture rates

Dual Core Processor without Hyper-Threading



VS .

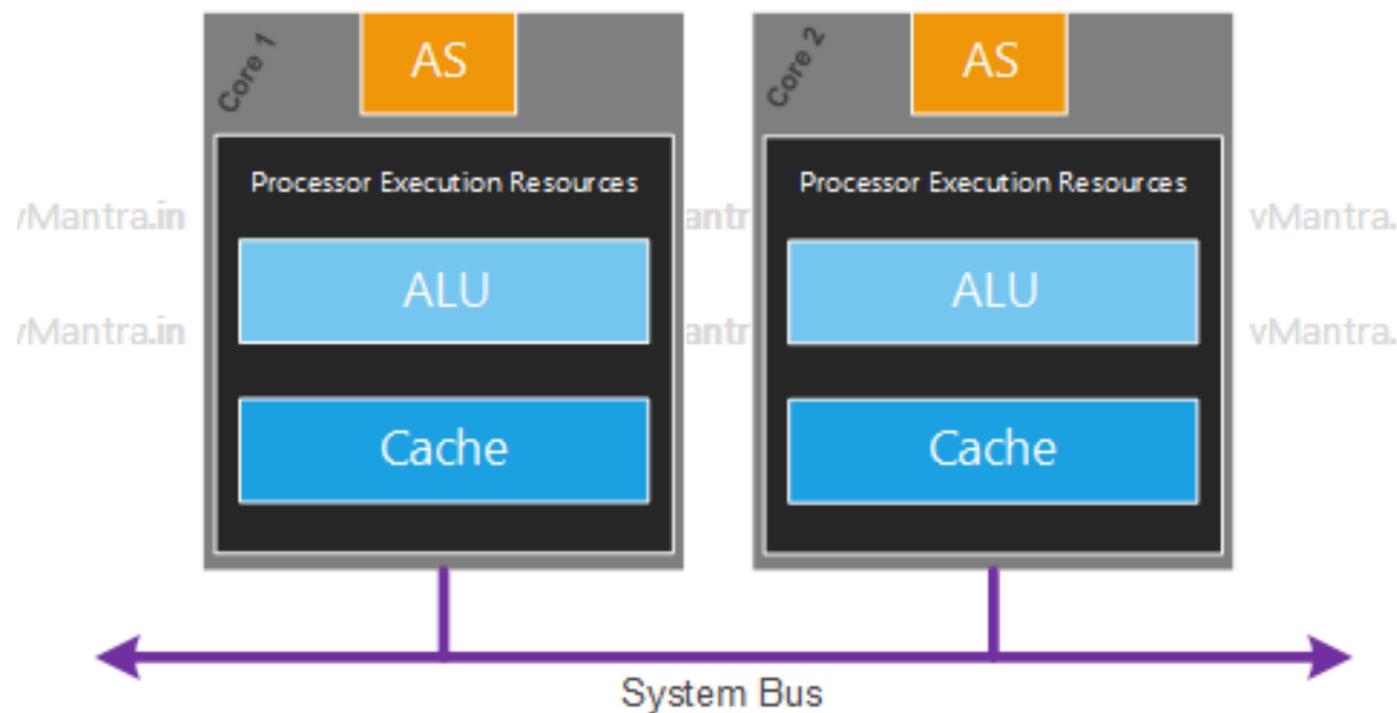
Single Core Processor with Hyper-Threading enabled



# CPU/core limitations

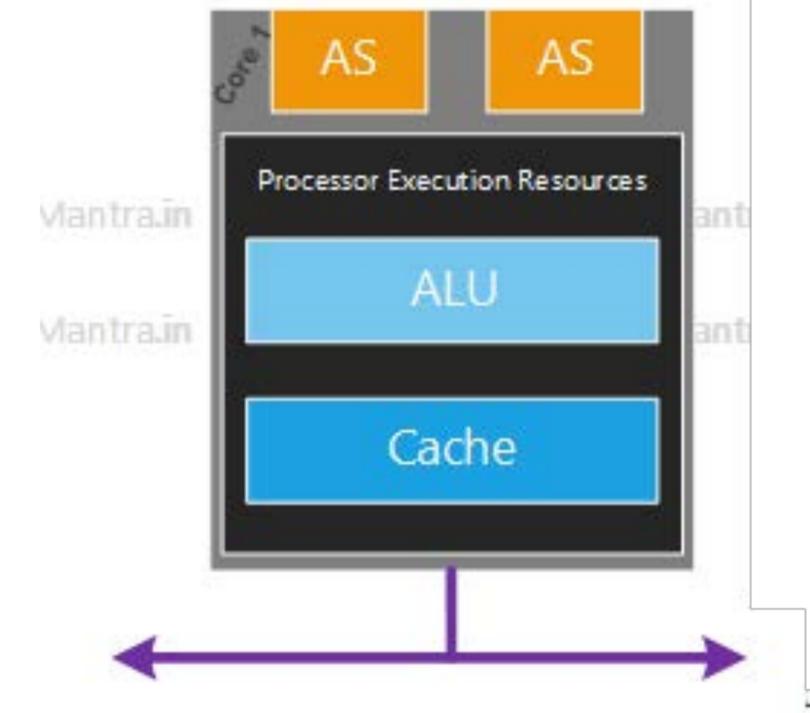
Achieving *very* high capture rates

Dual Core Processor without Hyper-Threading



VS.

Single Core Processor with Hyper-Threading enabled



```
$> grep -E '^processor' /proc/cpuinfo | wc -l  
2
```

```
$> grep -E '^processor' /proc/cpuinfo | wc -l  
2
```

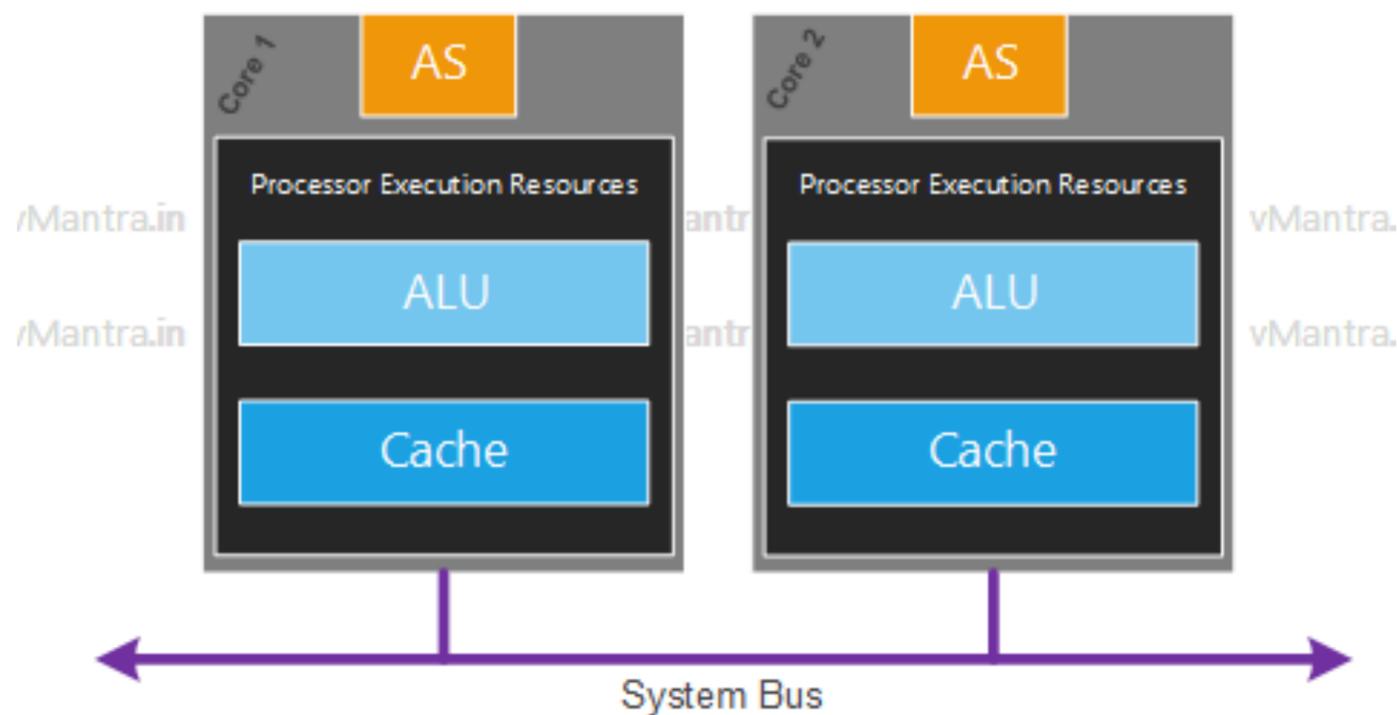
```
top - 16:53:39 up 1:07, 2 users, load average: 0.05, 0.03, 0.00
Tasks: 92 total, 1 running, 91 sleeping, 0 stopped, 0 zombie
%Cpu(s): 0.0 us, 0.2 sy, 0.0 ni, 99.7 id, 0.0 wa, 0.0 hi, 0.0 si, 0.2 st
KiB Mem : 4050652 total, 3573436 free, 120792 used, 356424 buff/cache
KiB Swap: 1048572 total, 1048572 free, 0 used. 3707452 avail Mem
```

PID	USER	PR	NI	VIRT	RES	SHR	S	%CPU	%MEM	TIME+	COMMAND
461	root	20	0	35800	1904	1664	S	0.3	0.0	0:00.12	irqbalance
1115	root	20	0	0	0	0	S	0.3	0.0	0:00.07	kworker/1:1
1	root	20	0	139004	6832	5272	S	0.0	0.2	0:00.87	systemd
2	root	20	0	0	0	0	S	0.0	0.0	0:00.00	kthreadd
4	root	0	-20	0	0	0	S	0.0	0.0	0:00.00	kworker/0:0H
6	root	0	-20	0	0	0	S	0.0	0.0	0:00.00	kworker/0:1H
7	root	20	0	0	0	0	S	0.0	0.0	0:00.00	ksortirqd/0
8	root	20	0	0	0	0	S	0.0	0.0	0:00.06	rcu_sched
9	root	20	0	0	0	0	S	0.0	0.0	0:00.00	rcu_bh
10	root	rt	0	0	0	0	S	0.0	0.0	0:00.00	migration/0
11	root	0	-20	0	0	0	S	0.0	0.0	0:00.00	lru-add-drain
12	root	rt	0	0	0	0	S	0.0	0.0	0:00.01	watchdog/0
13	root	20	0	0	0	0	S	0.0	0.0	0:00.00	cpuhp/0
14	root	20	0	0	0	0	S	0.0	0.0	0:00.00	cpuhp/1
15	root	rt	0	0	0	0	S	0.0	0.0	0:00.00	watchdog/1
16	root	rt	0	0	0	0	S	0.0	0.0	0:00.06	migration/1
17	root	20	0	0	0	0	S	0.0	0.0	0:00.00	ksoftirqd/1

# CPU/core limitations

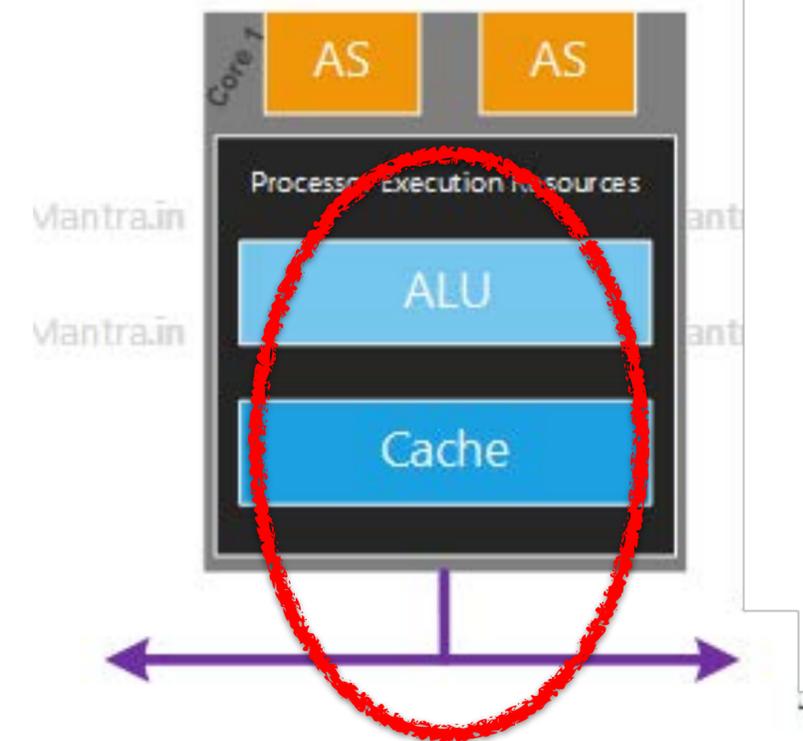
Achieving *very* high capture rates

Dual Core Processor without Hyper-Threading



VS .

Single Core Processor with Hyper-Threading enabled

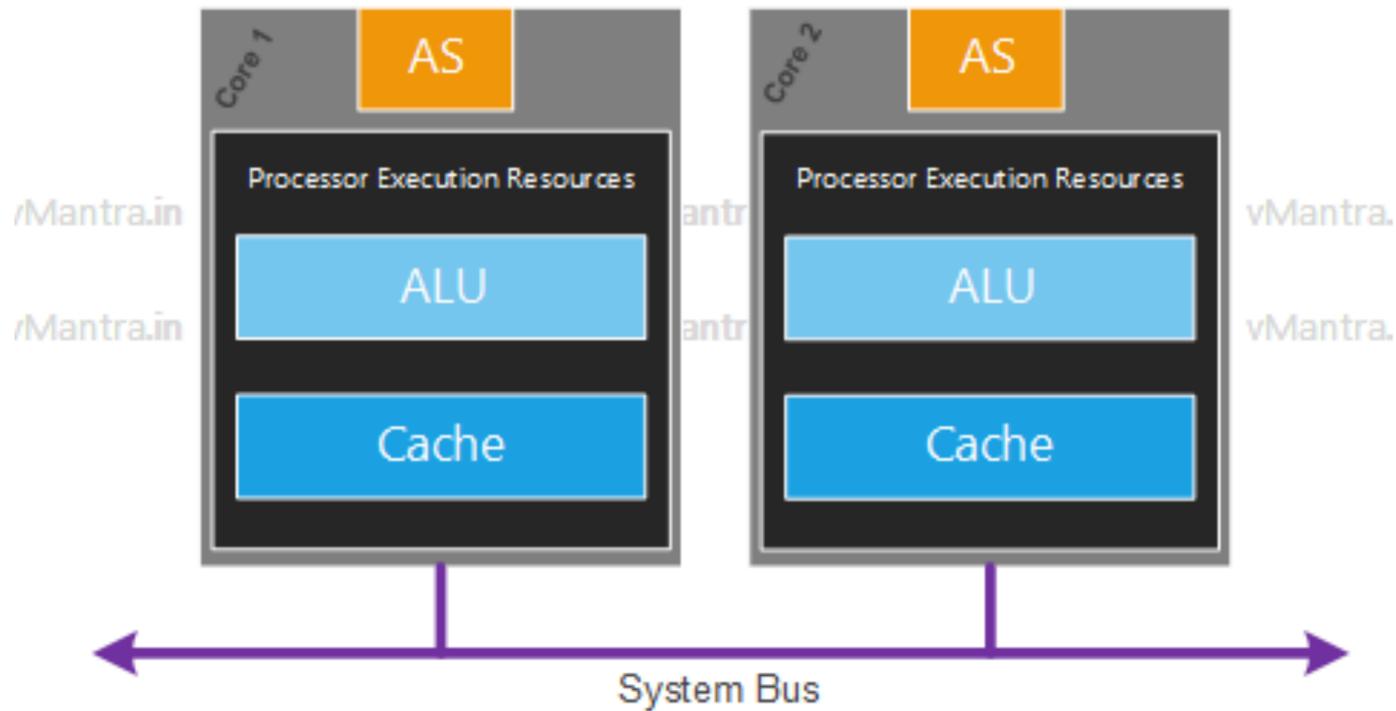


# CPU/core limitations

Achieving *very* high capture rates

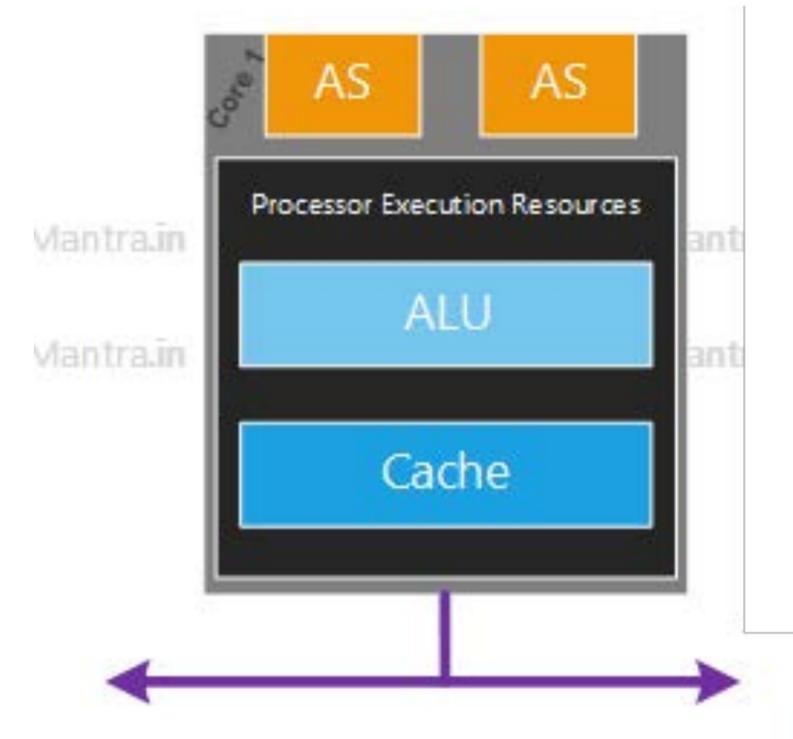
```
thread0 {  
  read(net, ...);  
}
```

```
thread1 {  
  write(disk, ...);  
}
```



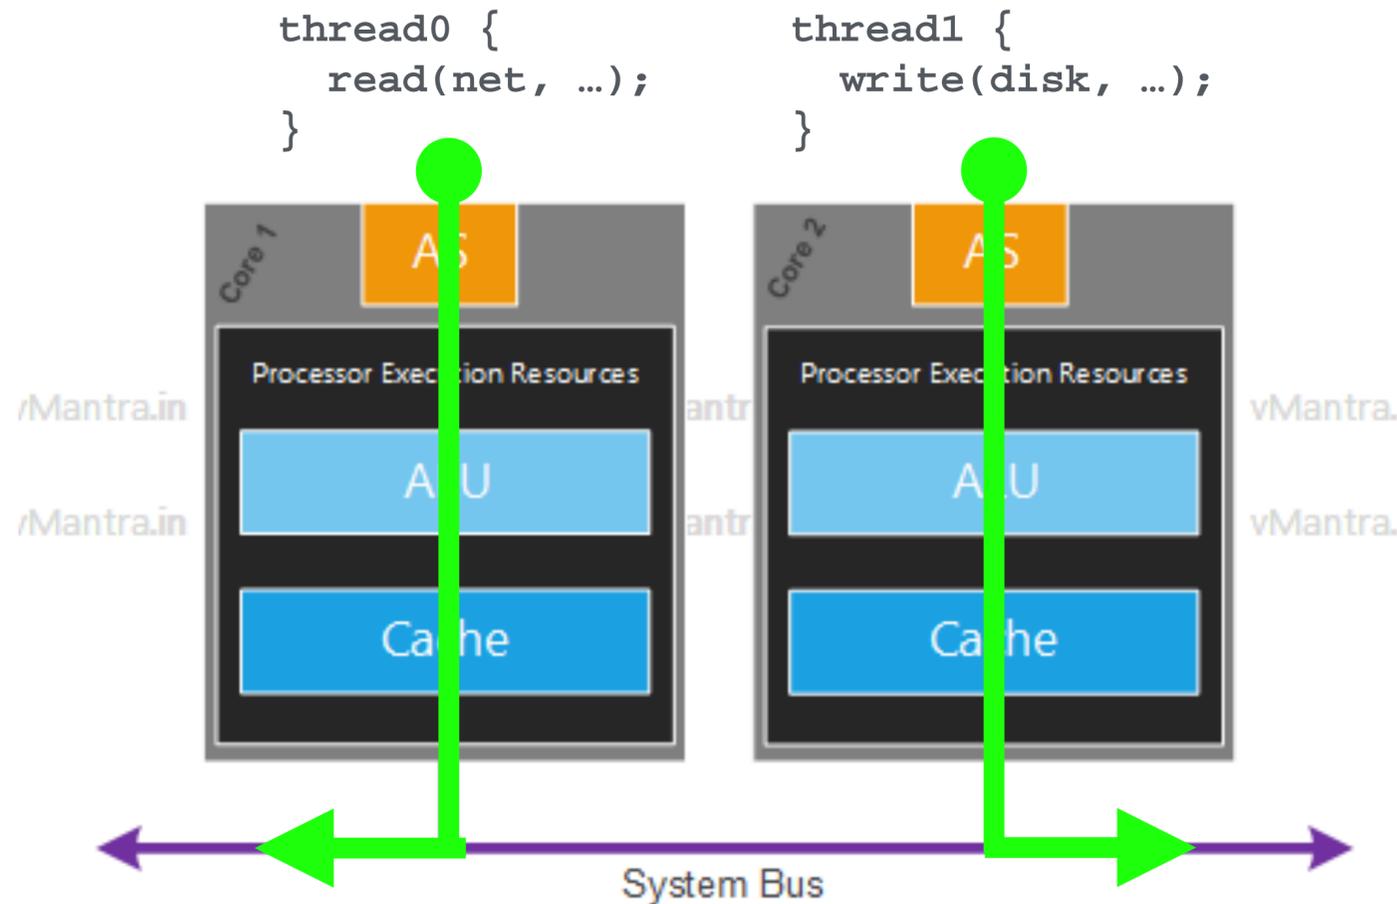
VS .

```
thread0 {      thread1 {  
  read(net, ...);  write(disk, ...);  
}
```

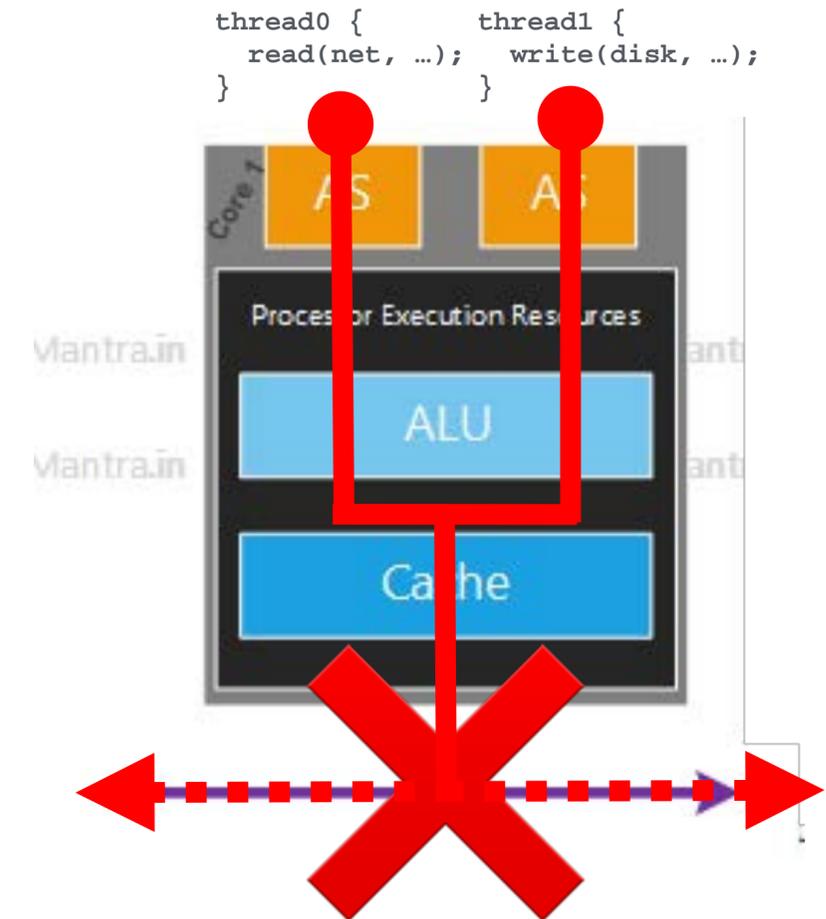


# CPU/core limitations

Achieving *very* high capture rates (#1 NO hyperthreading)



VS .



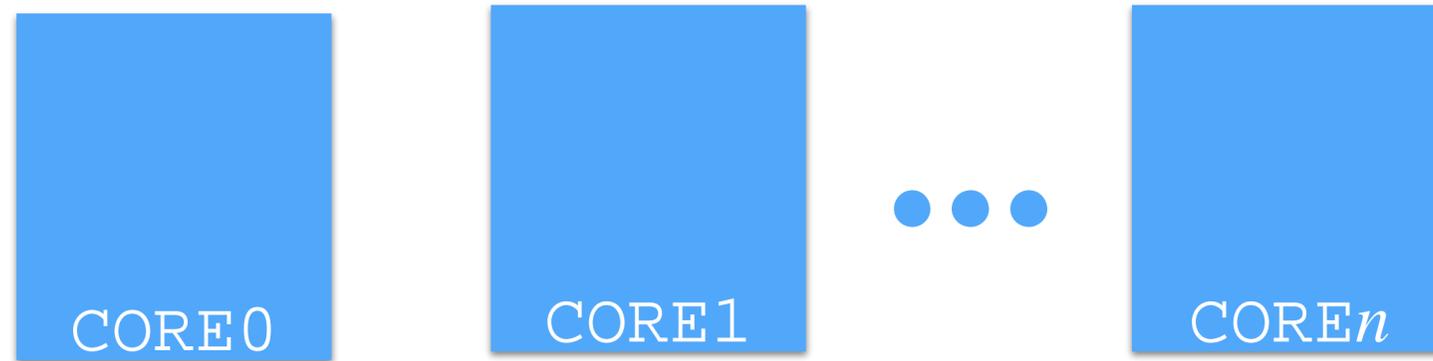
# CPU/core limitations

Achieving *very* high capture rates

# Interrupt pinning

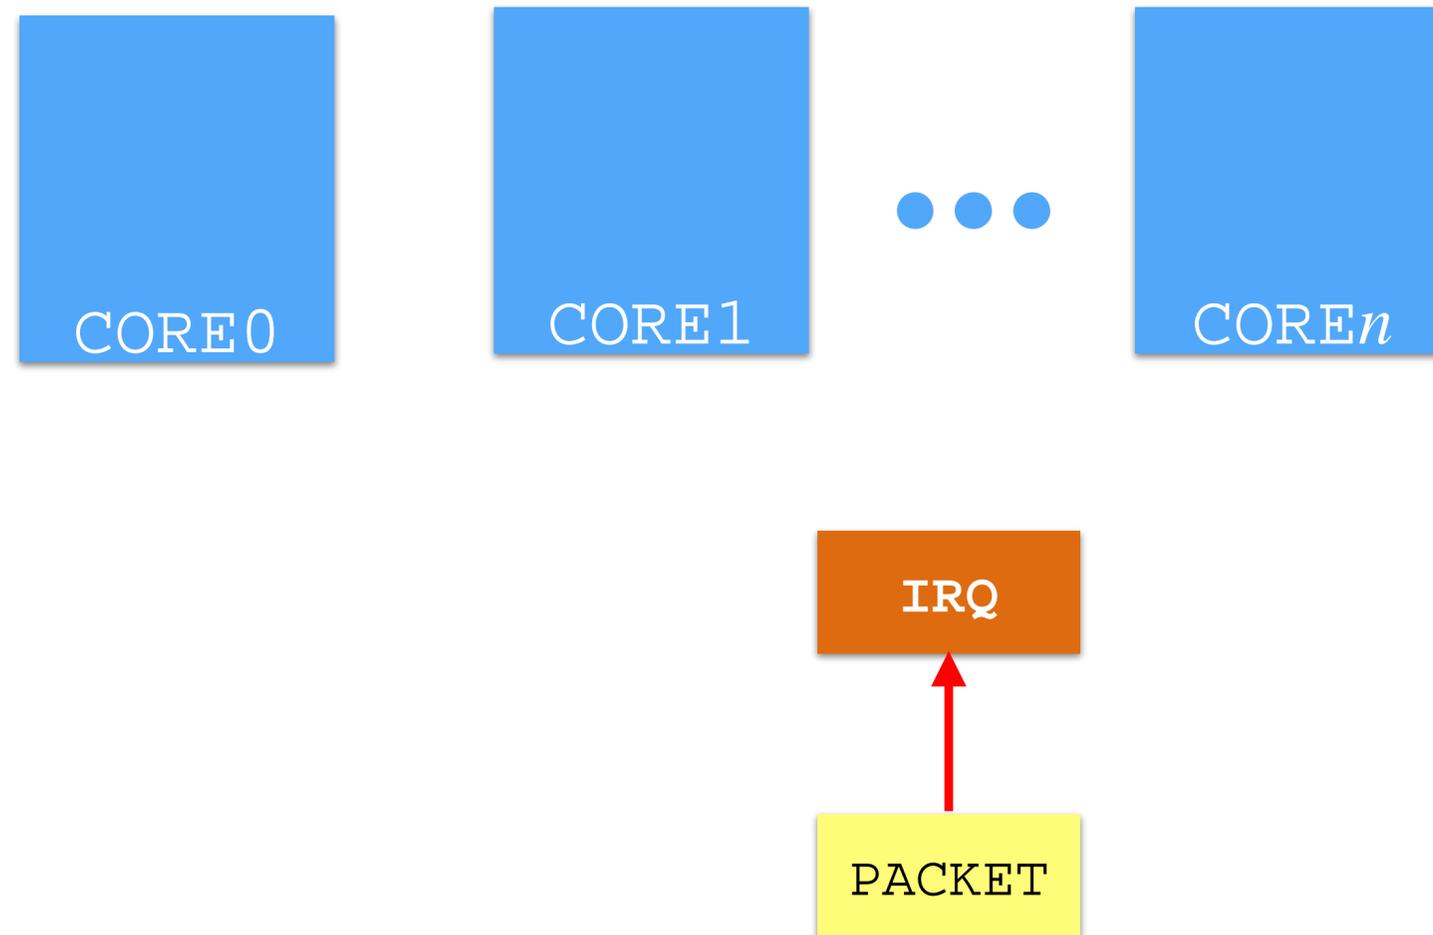
# CPU/core limitations

Achieving *very* high capture rates



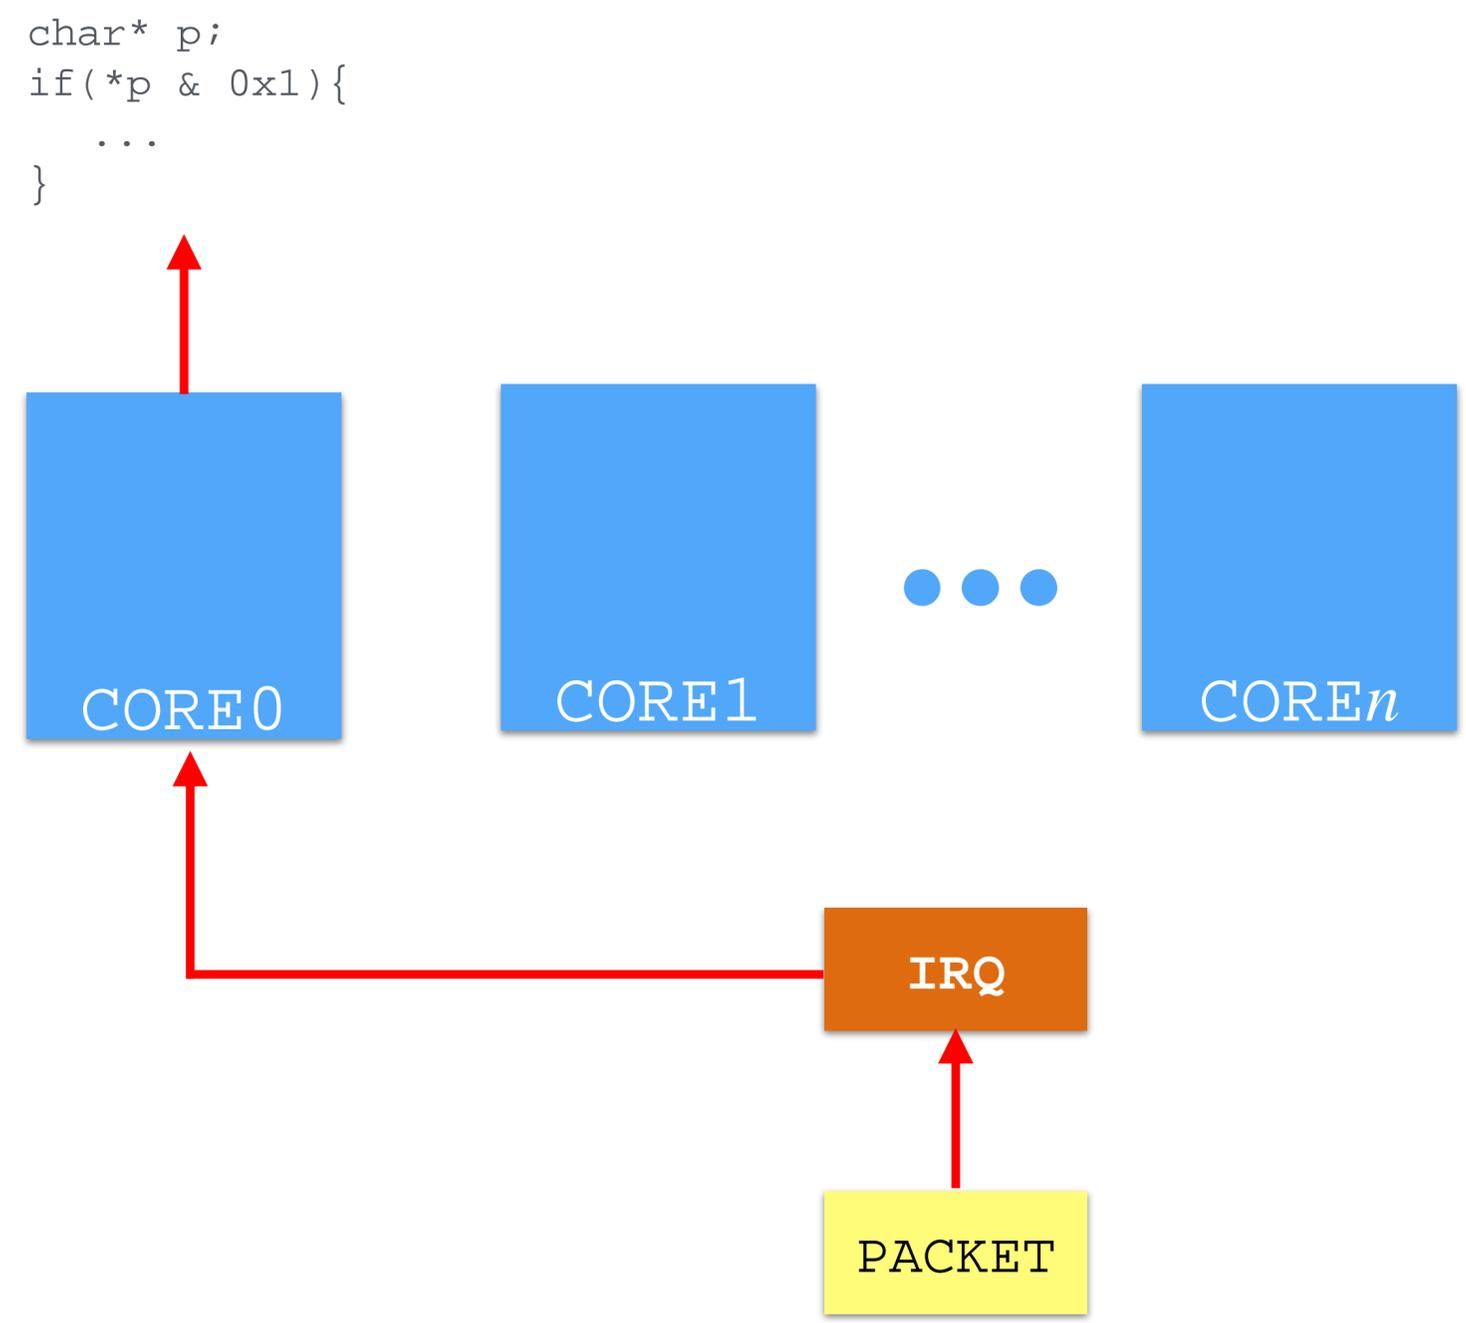
# CPU/core limitations

Achieving *very* high capture rates



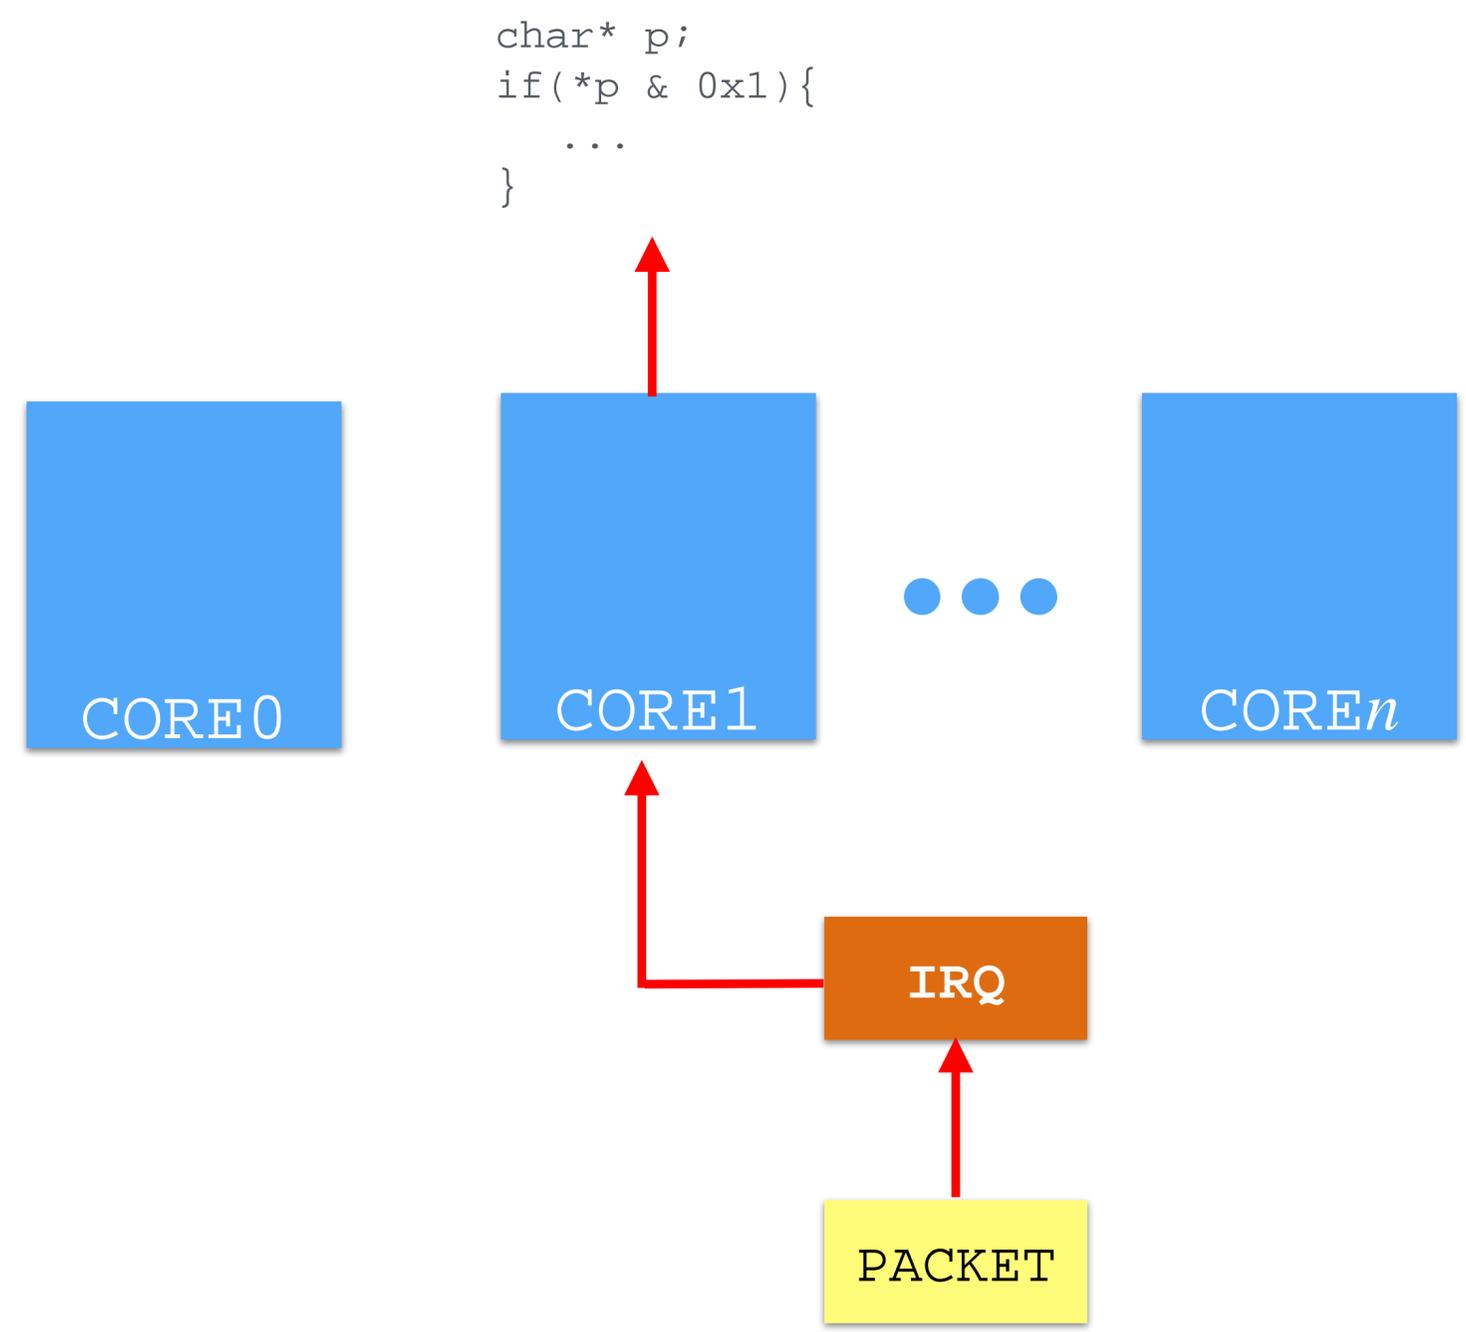
# CPU/core limitations

Achieving *very* high capture rates



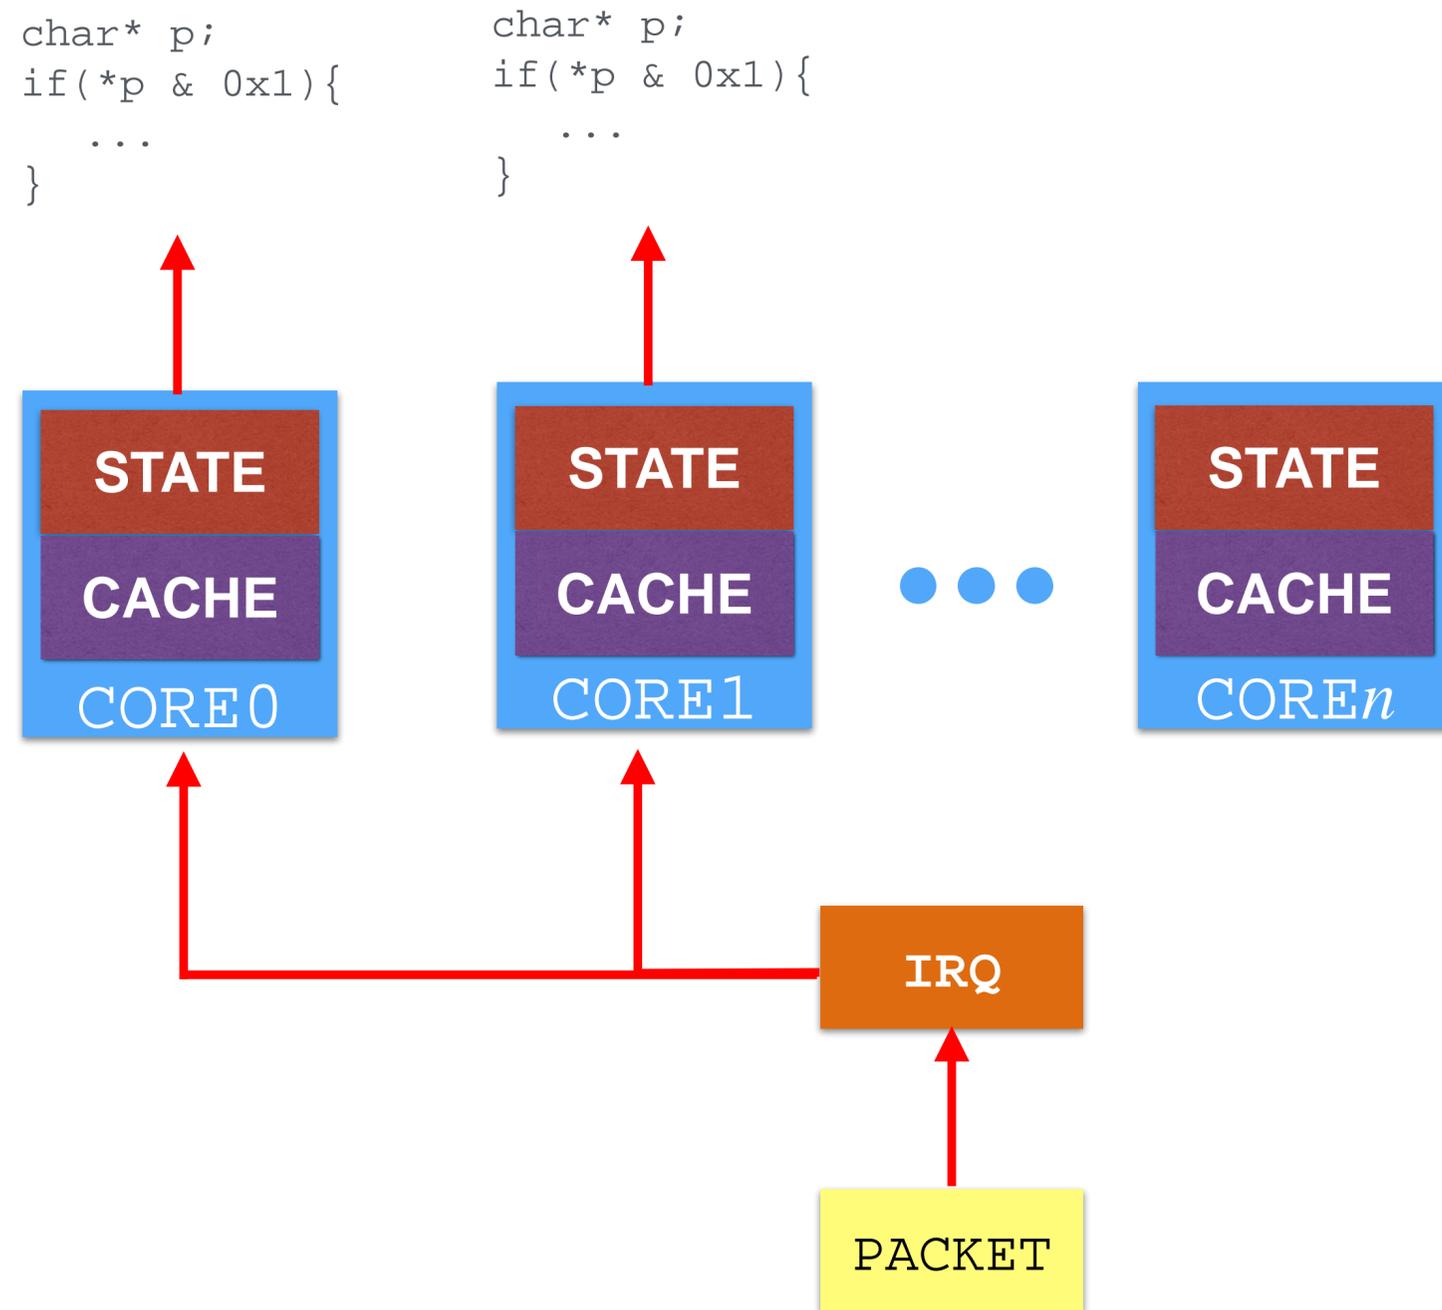
# CPU/core limitations

Achieving *very* high capture rates



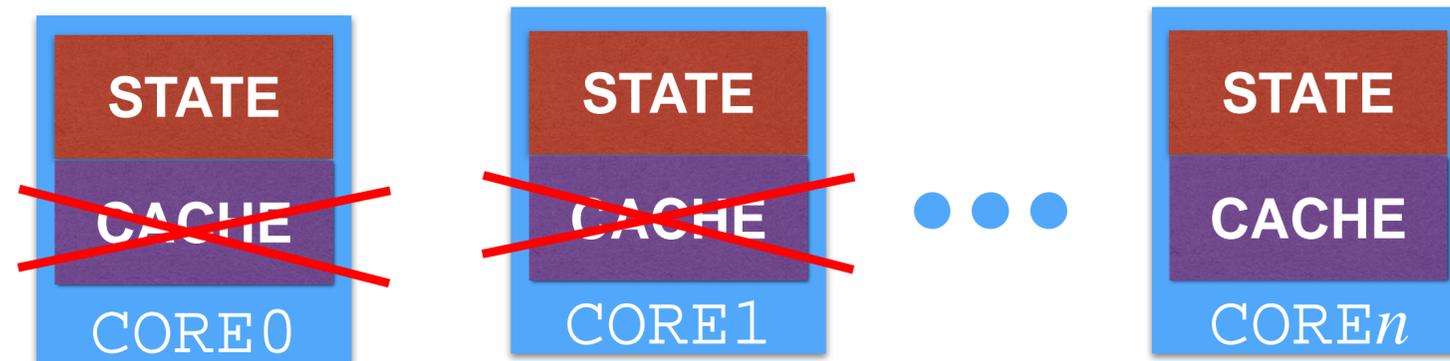
# CPU/core limitations

Achieving *very* high capture rates



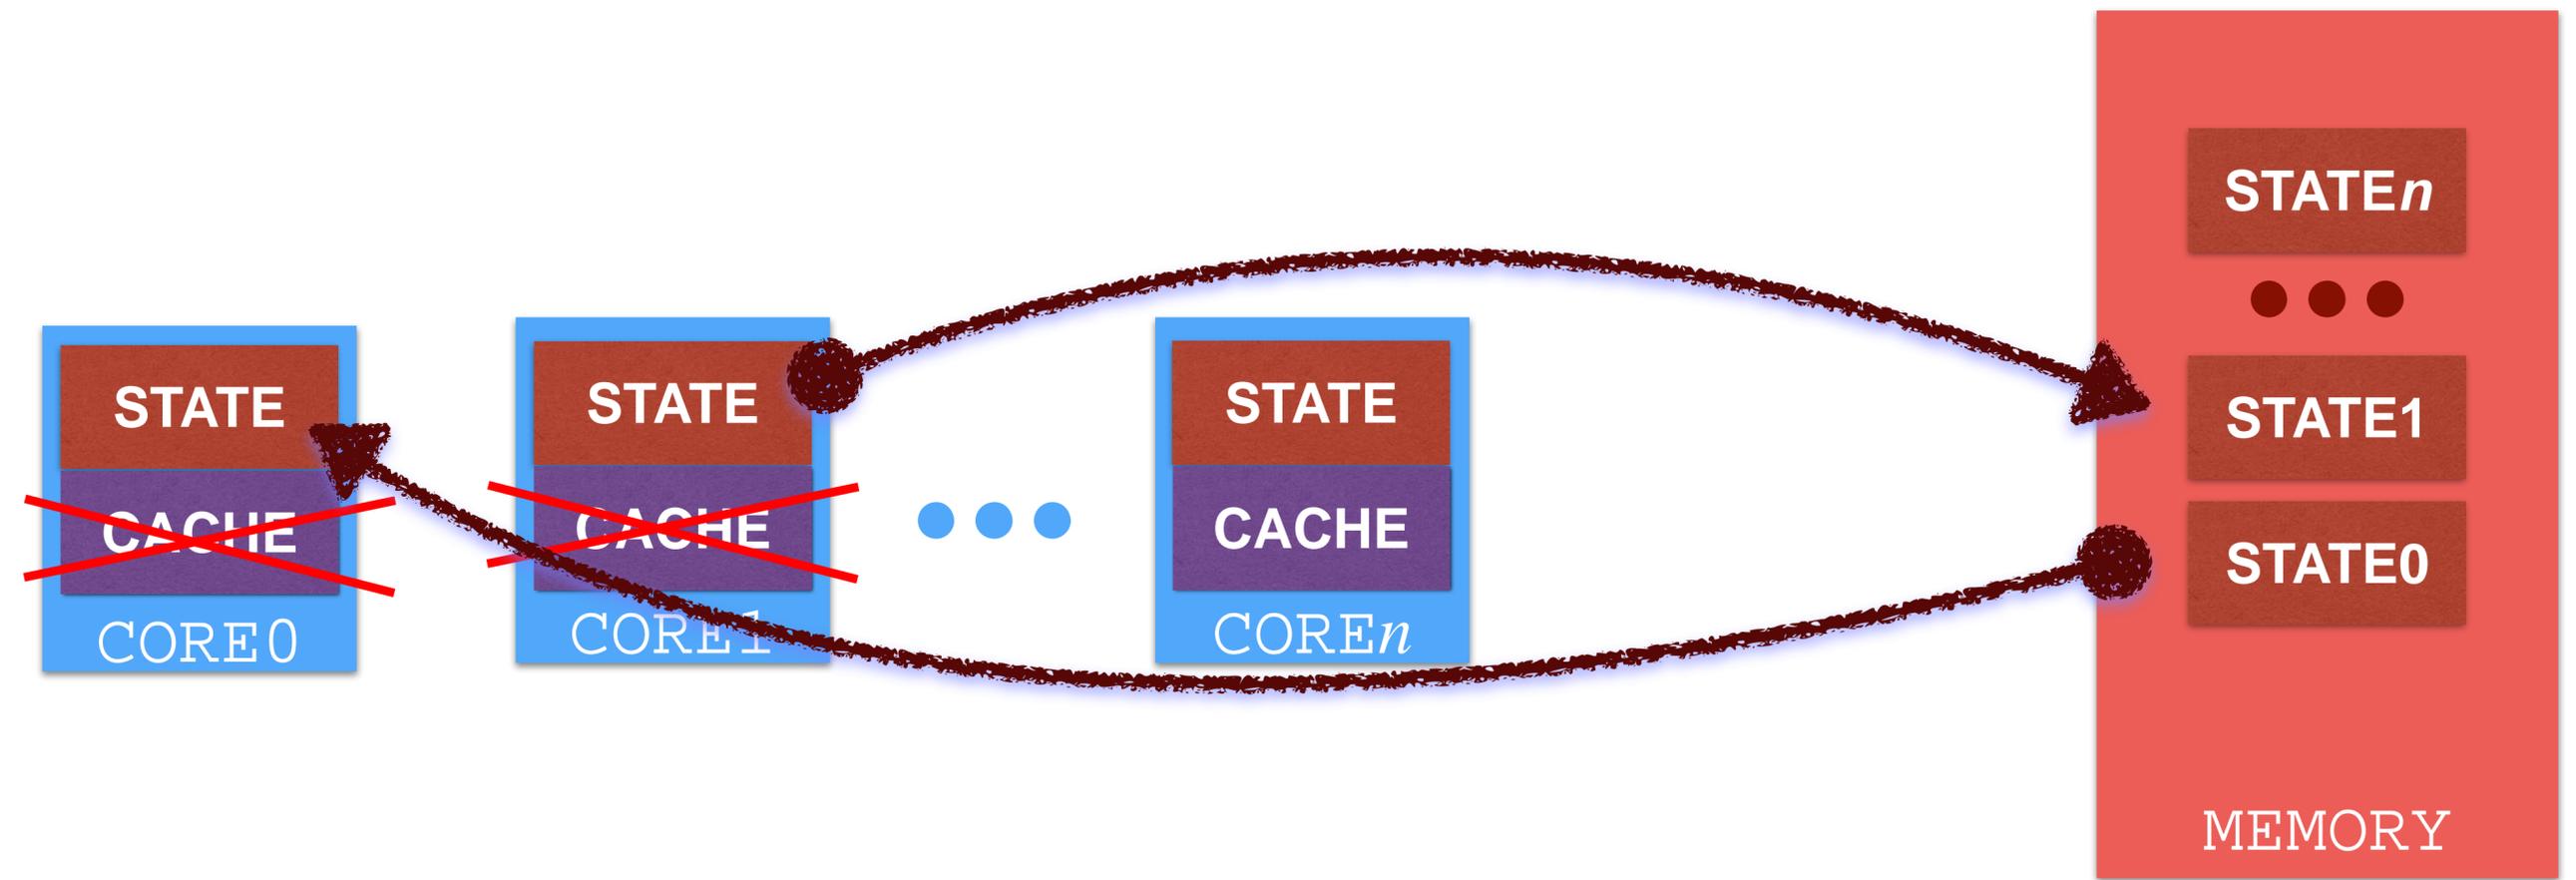
# CPU/core limitations

Achieving *very* high capture rates



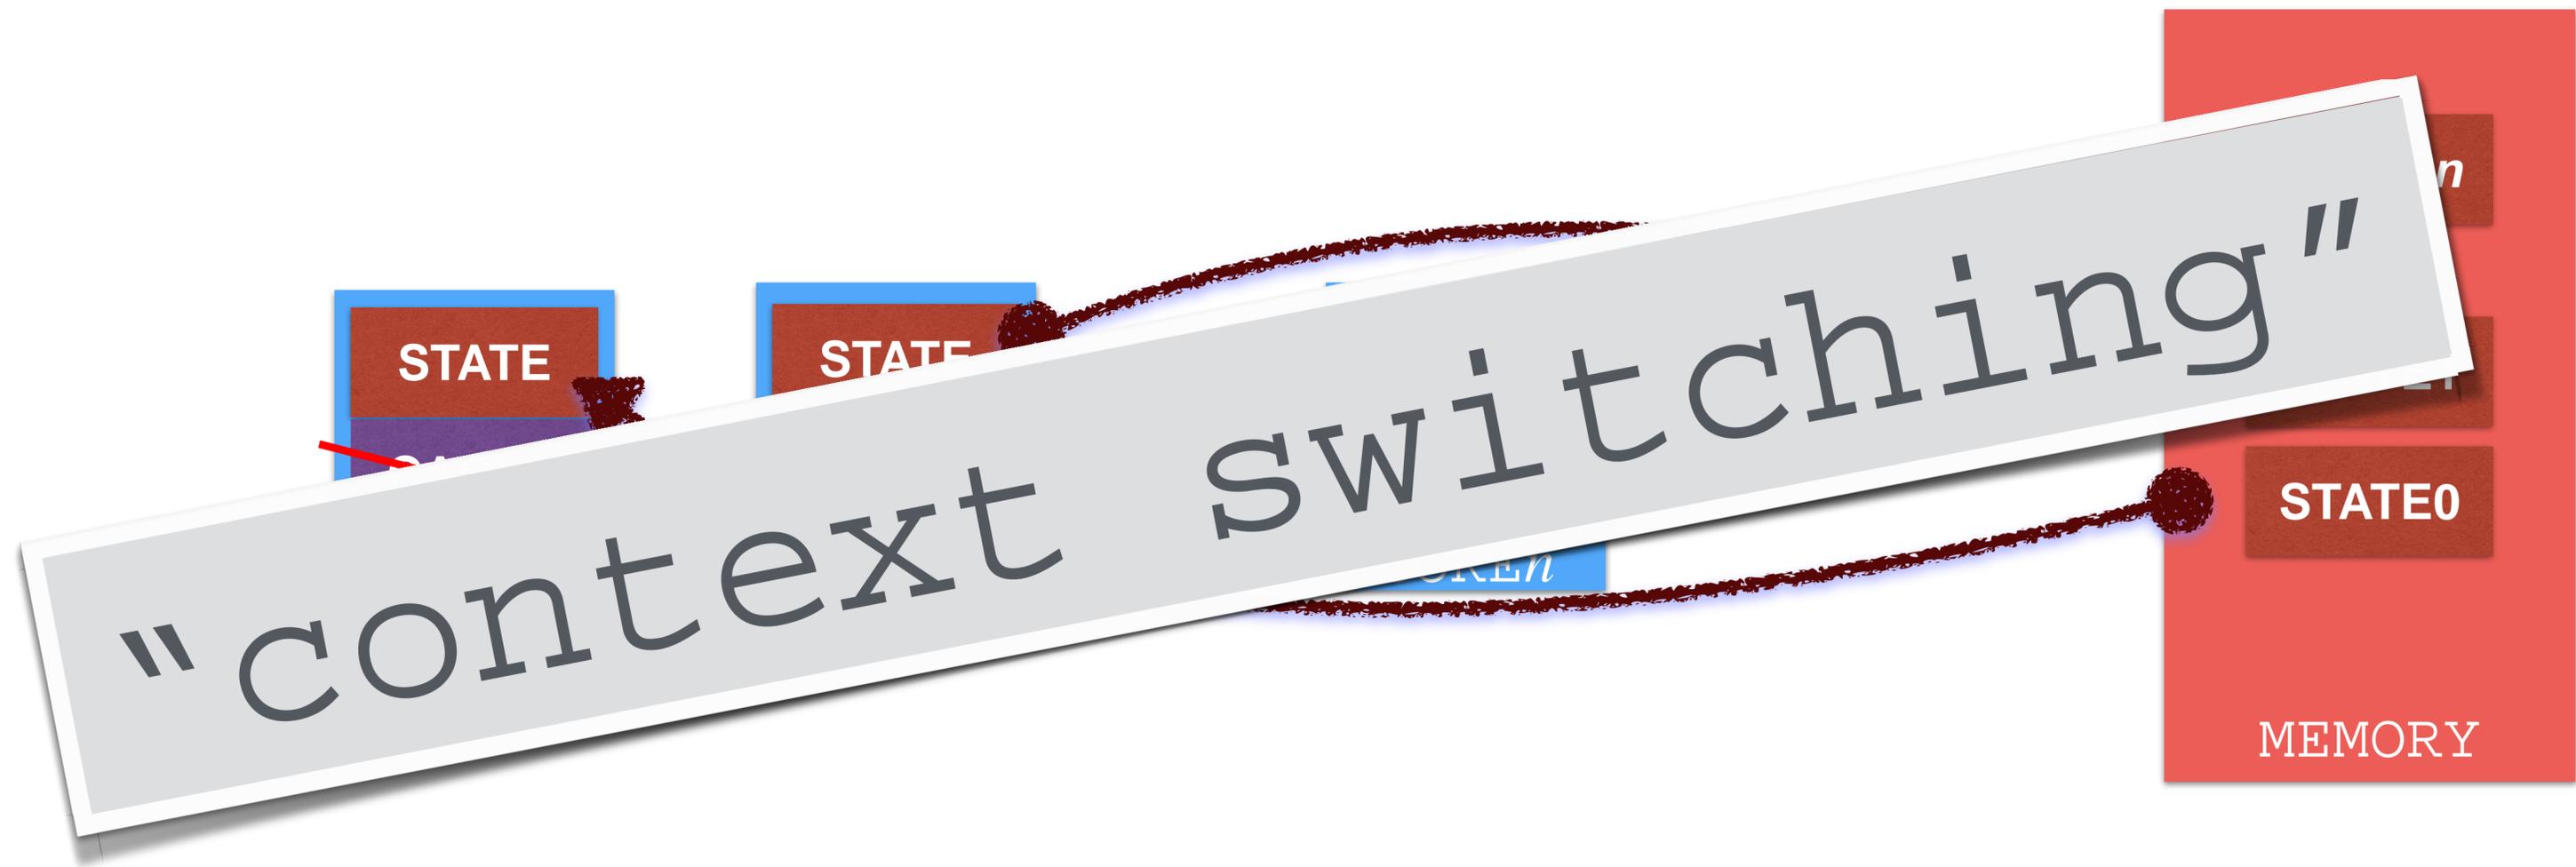
# CPU/core limitations

Achieving *very* high capture rates



# CPU/core limitations

Achieving *very* high capture rates



# CPU/core limitations

Achieving *very* high capture rates (#2 IRQ scheduling)

**Linux default:**

*round-robin IRQ scheduling*

**Solution:**

*pin **ethernet** IRQs to fixed core(s)  
to keep data local to core + cache*

# CPU/core limitations

Achieving *very* high capture rates

## Buffer sizes

# CPU/core limitations

Achieving *very* high capture rates

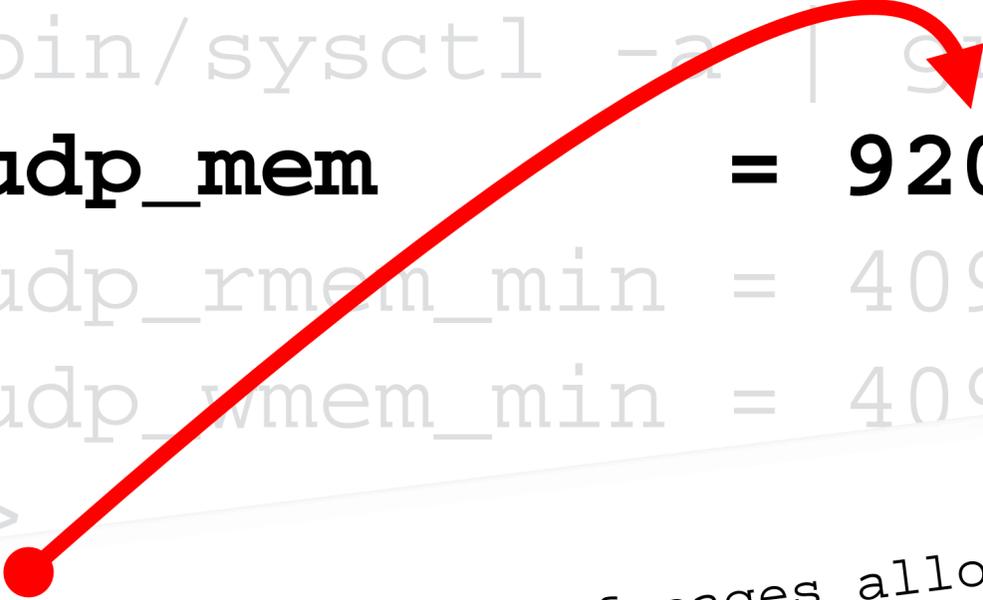
```
$ sudo /sbin/sysctl -a | grep udp
net.ipv4.udp_mem          = 92031  122711  184062
net.ipv4.udp_rmem_min    = 4096
net.ipv4.udp_wmem_min    = 4096
net.core.rmem_max        = 1073741824
<< snip >>
```

<https://www.kernel.org/doc/Documentation/networking/ip-sysctl.txt>

# CPU/core limitations

Achieving *very* high capture rates

```
$ sudo /sbin/sysctl -a | grep udp  
net.ipv4.udp_mem = 92031 122711 184062  
net.ipv4.udp_rmem_min = 4096  
net.ipv4.udp_wmem_min = 4096  
<< snip >>
```



**~370 MB**

Number of pages allowed for queueing by all UDP sockets

<https://www.kernel.org/doc/Documentation/networking/ip-sysctl.txt>

# CPU/core limitations

Achieving *very* high capture rates

```
$ sudo /sbin/sysctl -a | grep -E 'net.core.[rw]mem'  
net.core.rmem_default = 212992  
net.core.rmem_max = 212992  
net.core.wmem_default = 212992  
net.core.wmem_max = 212992
```

# CPU/core limitations

Achieving *very* high capture rates

```
$ sudo /sbin/sysctl -a | grep -E 'net.core.[rw]mem'  
net.core.rmem_default = 212992  
net.core.rmem_max = 212992  
net.core.wmem_default = 212992  
net.core.wmem_max = 212992
```

# CPU/core limitations

Achieving *very* high capture rates

```
$ sudo /sbin/sysctl -a | grep backlog  
net.core.netdev_max_backlog = 1000  
net.ipv4.tcp_max_syn_backlog = 128
```

# CPU/core limitations

Achieving *very* high capture rates

```
$ sudo /sbin/sysctl -a | grep backlog  
net.core.netdev_max_backlog = 1000  
net.ipv4.tcp_max_syn_backlog = 128
```

# CPU/core limitations

Achieving very high capture rates (#3 Tune buffers)

```
$ cat tune.txt (*)
```

```
net.core.rmem_max=201326592
```

```
net.ipv4.udp_mem="536870912 805306368 1073741824"
```

```
net.core.netdev_max_backlog=1048576
```

```
<< snip: there is more! >>
```

```
$ sudo /sbin/sysctl -ptune.txt
```

(\*) See <https://github.com/jive-vlbi/jive5ab/doc/flexbuf.recording.txt> (in 3.1.0-branch and later)  
A documented 'script' to serve as tuning guide: the what, why, and how

# CPU/core limitations

Achieving *very* high capture rates

Single core  
performance

# CPU/core limitations

Achieving *very* high capture rates

Single core: 16-21\* Gbps reliably

- per packet interrupt handling
- memory speed limit

(\* ) *broadly scales with CPU clock frequency*

# CPU/core limitations

Achieving *very* high capture rates (#4 multiple separate streams)

Create separate  $\leq 16$  Gbps streams:

- different IP destination address
- different UDP destination port
- or both

# CPU/core limitations

Achieving *very* high capture rates (#4 multiple separate streams)

Create separate  $\leq 16$

- diff

Setting up + starting  
not really supported



CPU / co

ons

Achieving very high capt

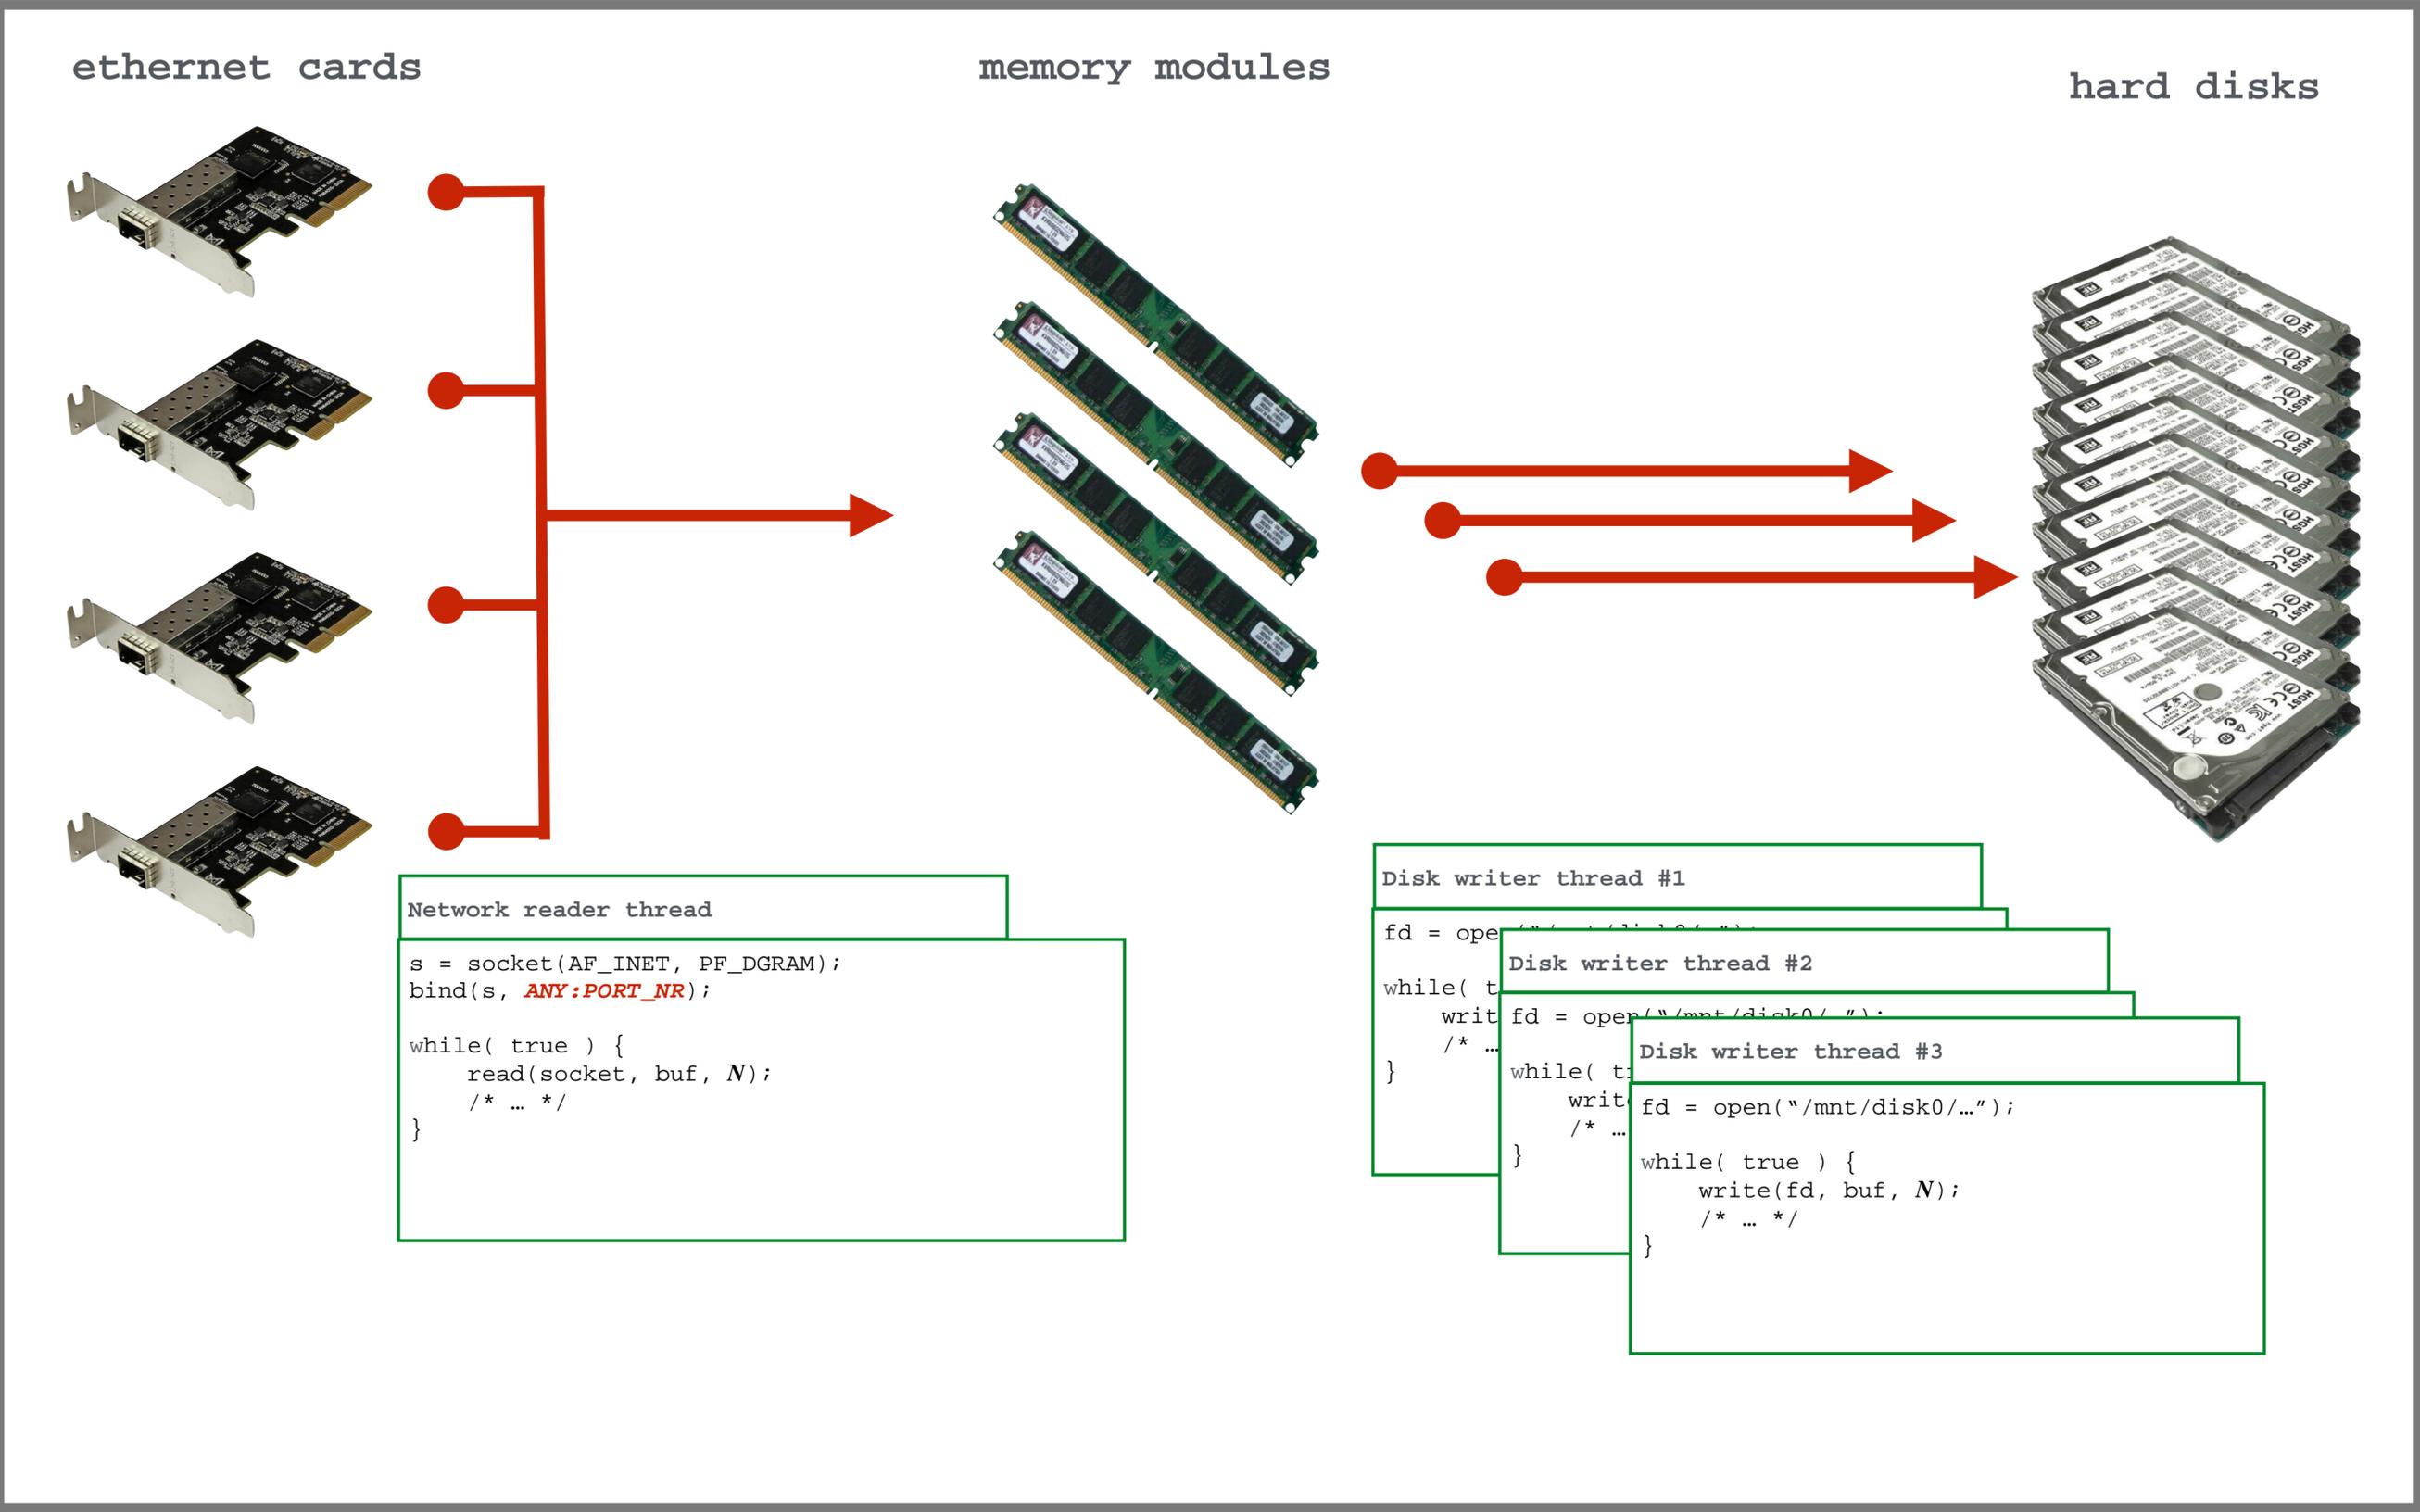
separate streams)



```
net_port = <port> ;  
record = on : <scan> ;
```



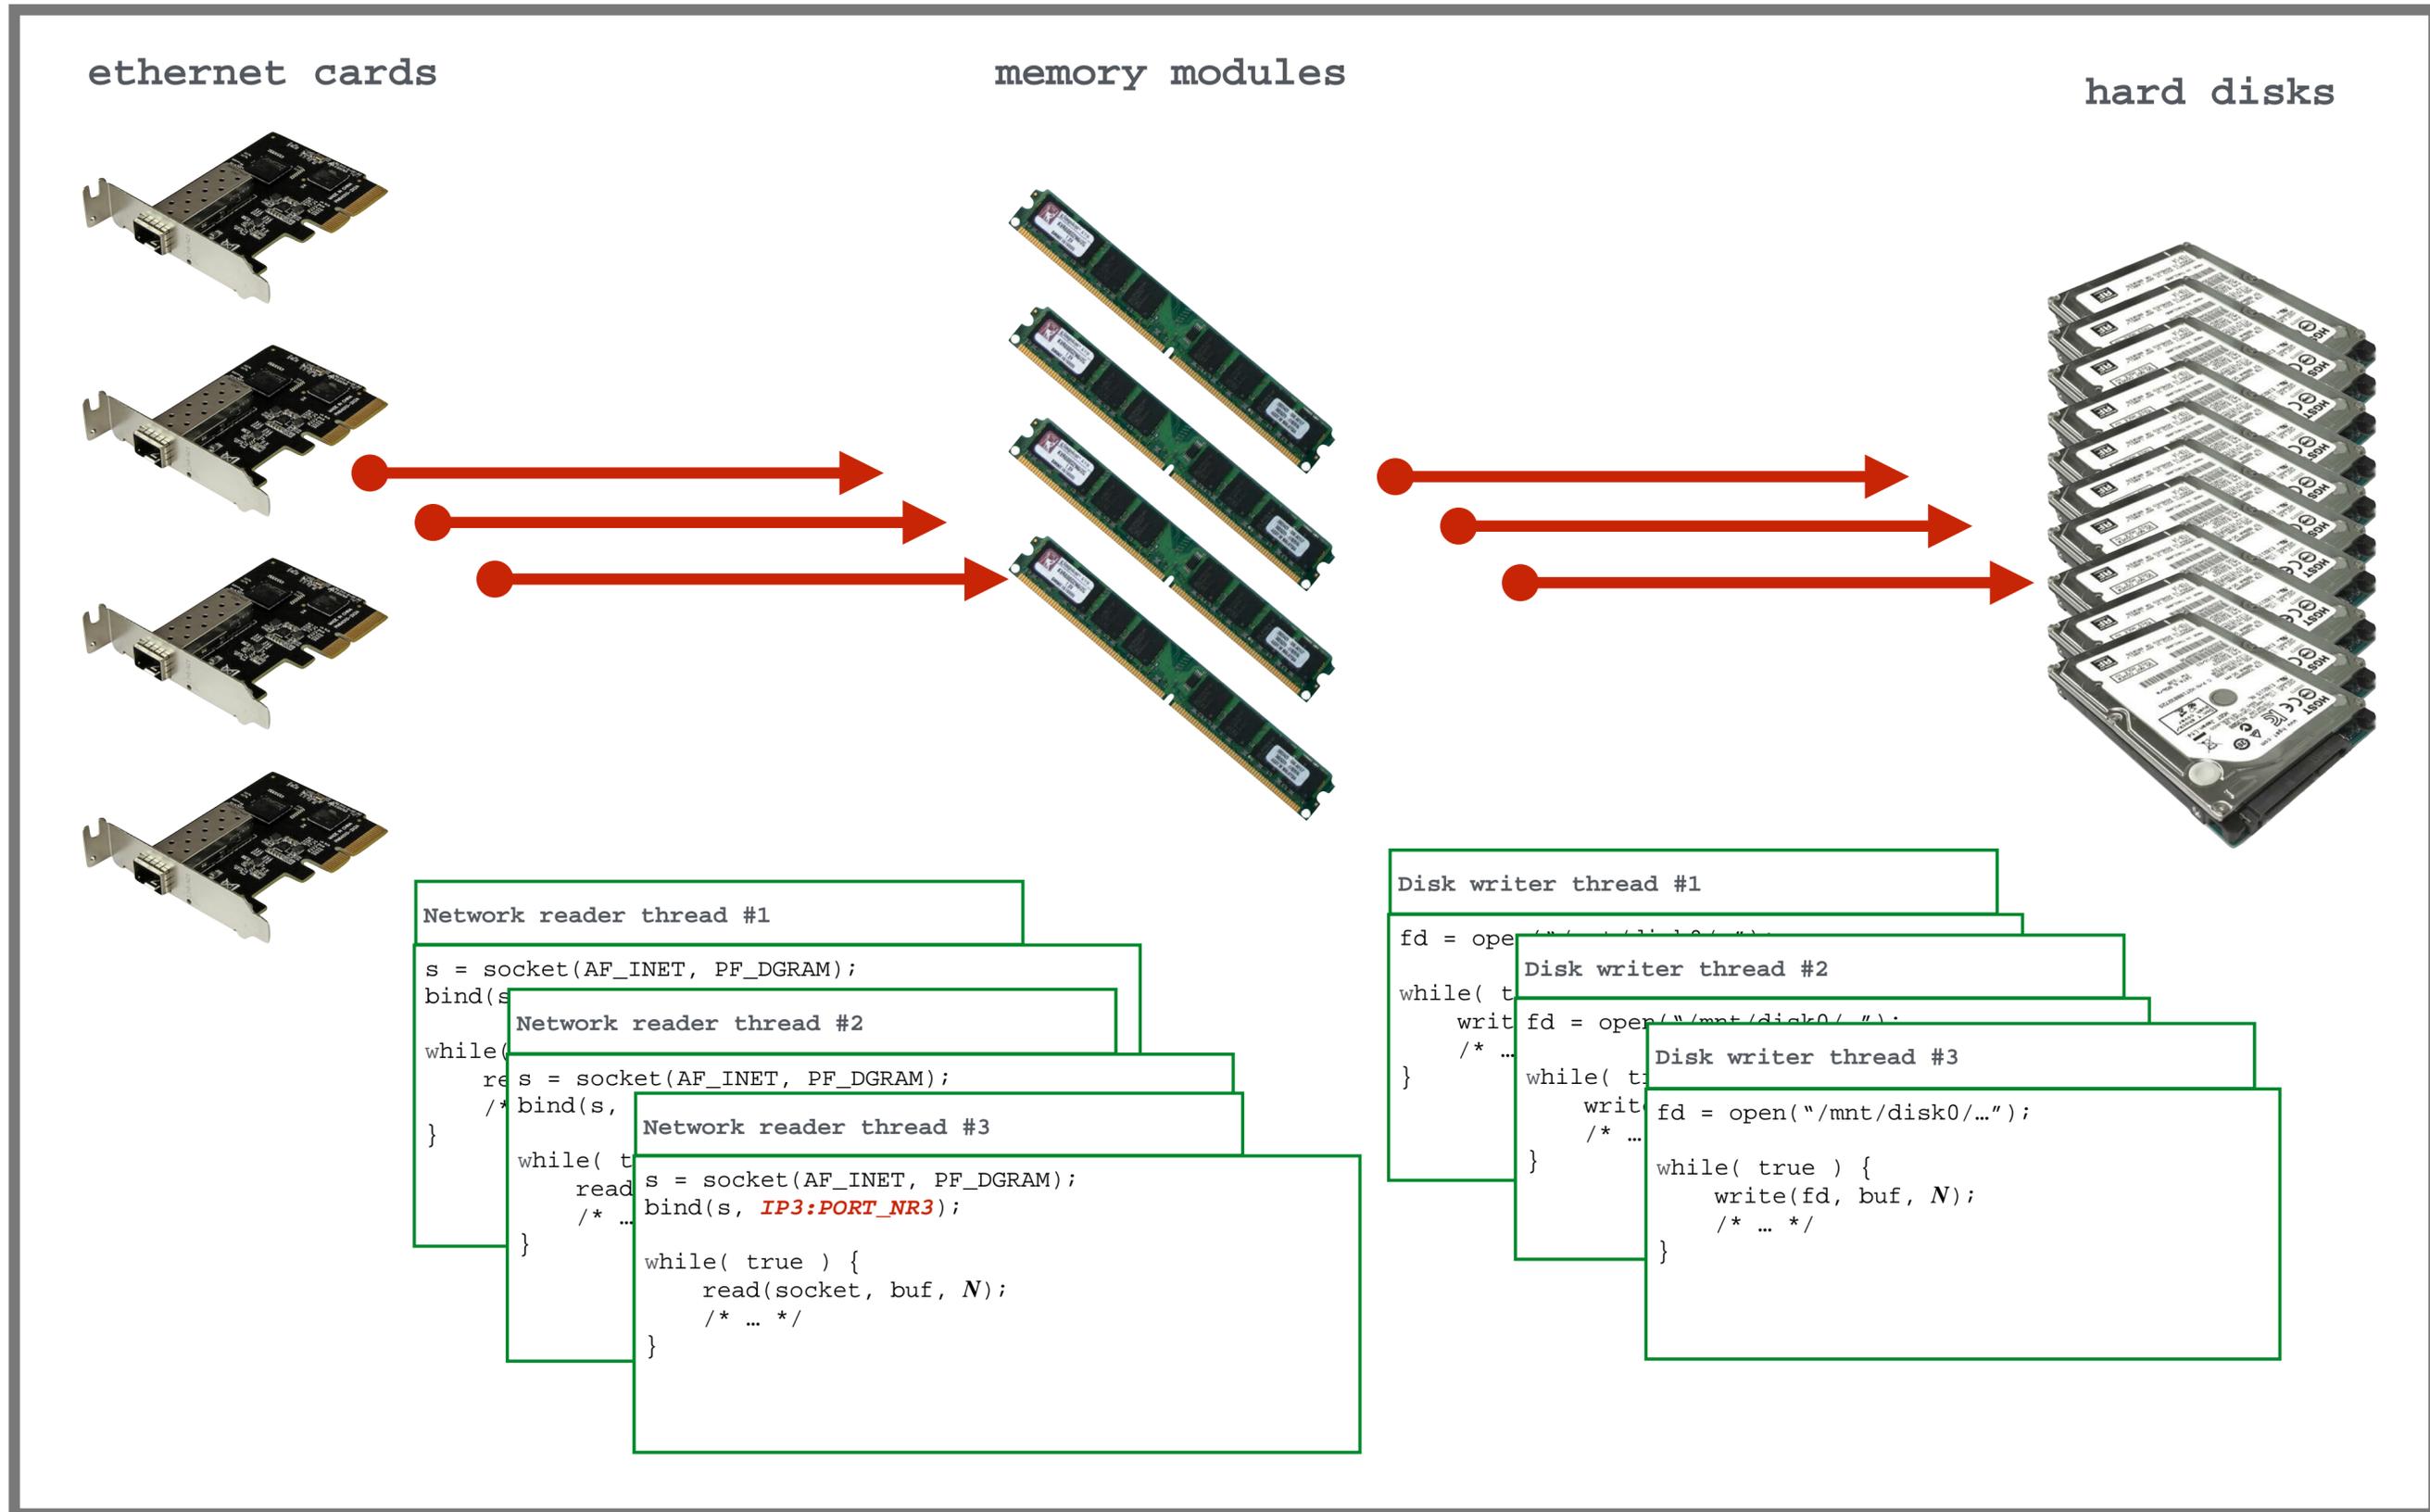
creates **one thread**, reads from network address  
[ip:]port



```
net_port = [ip1@]port1 [:[ip2@]port2]* ;  
record = on : <scan> ;
```



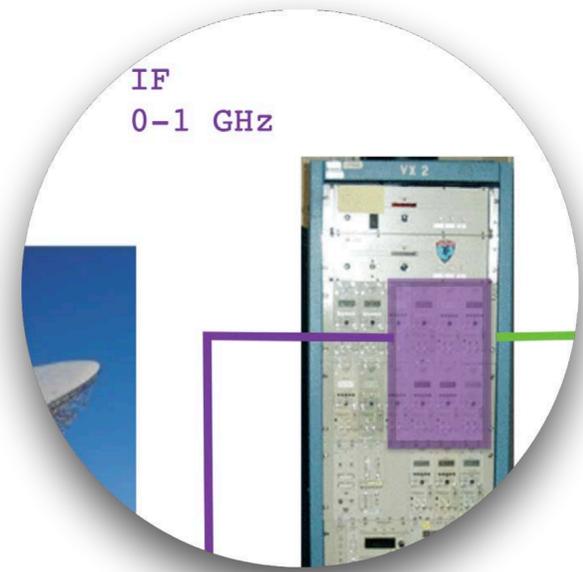
creates **n<sub>port</sub>** threads  
reading from network addresses  
[ip1:]port1,  
[ip2:]port2,  
etc.



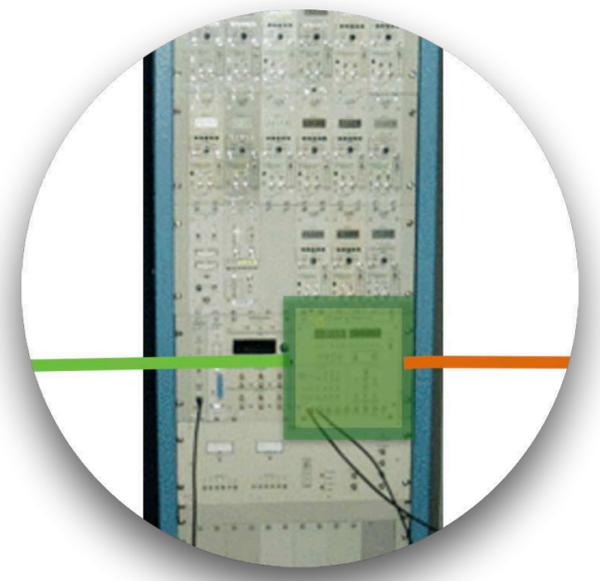
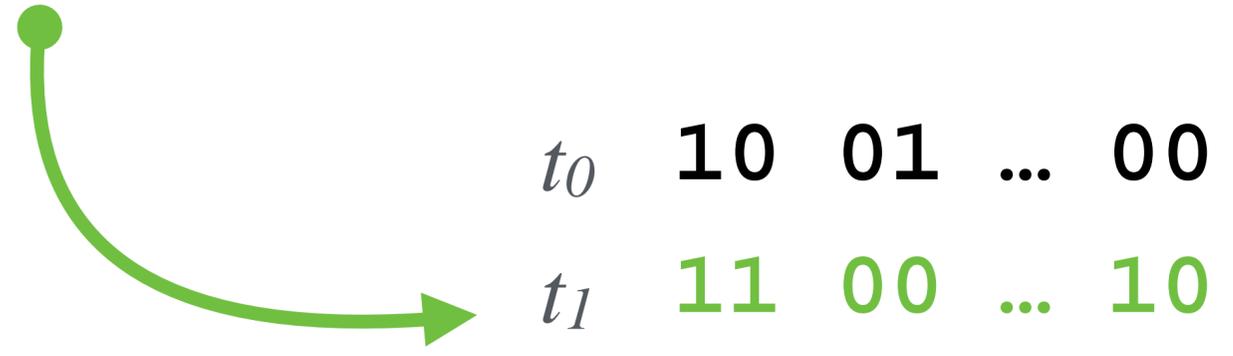
# Data format (s)

what did I just record?

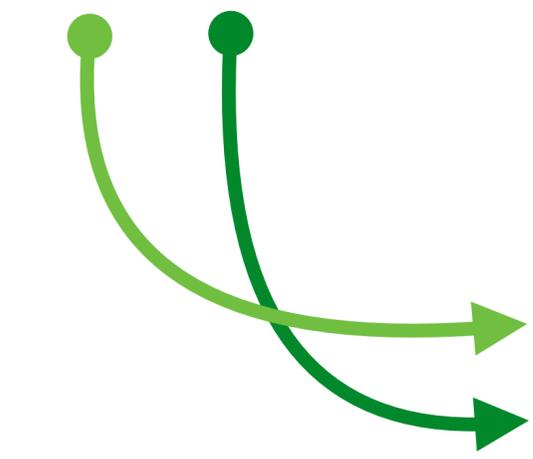
	$t_0$	$t_1$	$t_2$	$t_3$	$t_4$	...									
BBC0 <sub>v</sub>	10	11	01	11	01	10	00	10	10	00	10	01	11	01	10
BBC1 <sub>v</sub>	01	00	10	10	00	11	10	00	10	01	11	01	10	00	01
...															
BBCN <sub>v</sub>	00	10	01	11	01	10	00	01	01	11	01	10	00	11	10



	$t_0$	$t_1$	$t_2$	$t_3$	$t_4$	...
BBC0 <sub>v</sub>	10	11	01	11	01	10 00 10 10 00 10 01 11 01 10
BBC1 <sub>v</sub>	01	00	10	10	00	11 10 00 10 01 11 01 10 00 01
...						
BBCN <sub>v</sub>	00	10	01	11	01	10 00 01 01 11 01 10 00 11 10

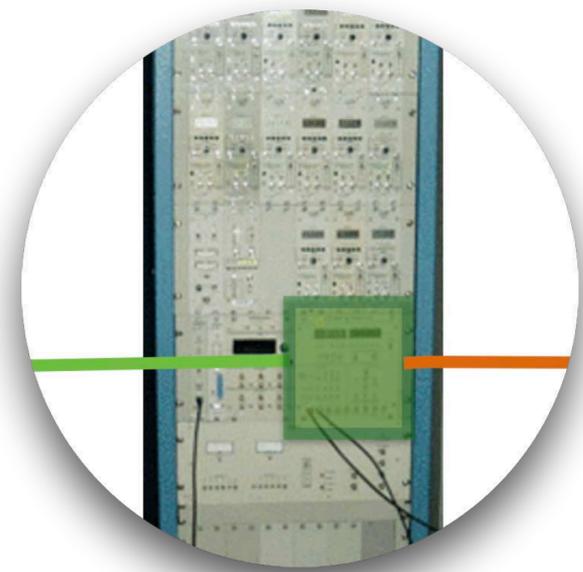


	$t_0$	$t_1$	$t_2$	$t_3$	$t_4$	...
BBC0 <sub>v</sub>	10	11	01	11	01	10 00 10 10 00 10 01 11 01 10
BBC1 <sub>v</sub>	01	00	10	10	00	11 10 00 10 01 11 01 10 00 01
...						
BBCN <sub>v</sub>	00	10	01	11	01	10 00 01 01 11 01 10 00 11 10



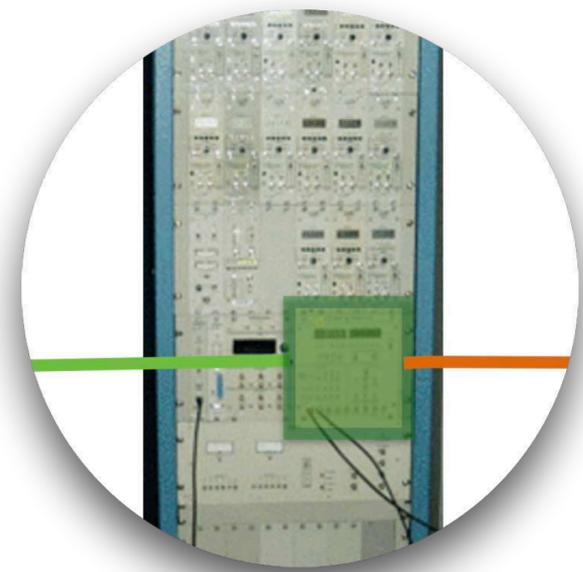
$t_0$	10	01	...	00
$t_1$	11	00	...	10
$t_2$	01	10	...	01
...				

$$\overbrace{\hspace{10em}}^{n_{channel} \times n_{bits \text{ per sample}} \div 8}$$



binary representation  
of format specific header

$t_0$	10	01	...	00
$t_1$	11	00	...	10
	11	11	...	11
	01	01	...	10
$t_2$	01	10	...	01
...				



# Mark5B

adb459ed	390dd99b	45f31dd4	9798e953
5d0a9d3e	81dd1a87	98976aae	70a862a2
e76bae65	2e6a79f4	41d589e3	9d6e991e
31349721	aaddd6bc	df83926a	7a78a5eb
abaddeed	bead04b1	65467925	18762097
0595f62c	c6752d53	2c61a994	fa85d6d9
97c4ac8c	db15ebaf	95a519b5	219b6379
a69a1469	fc7859ab	43c77c54	4a158ea6
9d75f2f7	58ea8499	11f9076b	7d8cf754
552619b8	a79b9ba7	d5e7fe86	44b849d5

*note: this is in 32-bit hexadecimal representation now*

adb459ed	390dd99b	45f31dd4	9798e953
5d0a9d3e	81dd1a87	98976aae	70a862a2
e76bae65	2e6a79f4	41d589e3	9d6e991e
31349721	aadd6bc	df83926a	7a78a5eb
<b>abaddeed</b>	<b>bead04b1</b>	<b>65467925</b>	<b>18762097</b>
<b>0595f62c</b>	<b>c6752d53</b>	<b>2c61a994</b>	<b>fa85d6d9</b>
97c4ac8c	db15ebaf	95a519b5	219b6379
a69a1469	fc7859ab	43c77c54	4a158ea6
9d75f2f7	58ea8499	11f9076b	7d8cf754
552619b8	a79b9ba7	d5e7fe86	44b849d5

*note: this is in 32-bit hexadecimal representation now*

6b5669a1	82c65526	a9b65a1f	63aa656c
54e695f5	9a8d9857	58975895	a7498aa2
2dd189e6	8999ac5b	d964f565	2269c9f6
9ae5296b	a8875542	4b87557d	
0a37956f			
			0000000
	662db168	2979f057	56959798
56256a76	e4ad879b	6bc1d467	c6d23918
94a4ad83	69d2665e	f625c279	6a12b1ab
46cb9389	ada02966	25694a99	a5a3975b
e6e46da8	61ba145a	72669a17	95931b29
dcd9e5a2	1d396981	b492a659	b194a6c5

no magic pattern!

6b5669a1	82c65526	a9b65a1f	63aa656c
54e695f5	9a8d9857	58975895	a7498aa2
2dd189e6	8999ac5b	d964f565	2269c9f6
9ae5296b	a8875542	4b87557d	5197668a
<b>0a37956f</b>	<b>2200066a</b>	<b>000003ec</b>	<b>4005069</b>
<b>00000000</b>	<b>00000000</b>	<b>00000000</b>	<b>00000000</b>
62ec9736	662db168	2979f057	56959798
56256a76	e4ad879b	6bc1d467	c6d23918
94a4ad83	69d2665e	f625c279	6a12b1ab
46cb9389	ada02966	25694a99	a5a3975b
e6e46da8	61ba145a	72669a17	95931b29
dcd9e5a2	1d396981	b492a659	b194a6c5

The standard 32-byte VDIF Data Frame Header is shown in Figure 3.

	Byte 3		Byte 2	Byte 1	Byte 0
	Bit 31 (MSB)				Bit 0 (LSB)
Word 0	$I_1$	$L_1$	Seconds from reference epoch <sub>30</sub>		
Word 1	Un-assigned <sub>2</sub>		Ref Epoch <sub>6</sub>	Data Frame # within second <sub>24</sub>	
Word 2	$V_3$		$\log_2(\#chns)_5$	Data Frame length (units of 8 bytes) <sub>24</sub>	
Word 3	$C_1$	bits/sample-1 <sub>5</sub>	Thread ID <sub>10</sub>	Station ID <sub>16</sub>	
Word 4	EDV <sub>8</sub>			Extended User Data <sub>24</sub>	
Word 5	Extended User Data <sub>32</sub>				
Word 6	Extended User Data <sub>32</sub>				
Word 7	Extended User Data <sub>32</sub>				

Figure 3: VDIF Data Frame Header format; subscripts are field lengths in bits; byte #s indicate relative byte address within 32-bit word (little endian format)

From VLBI Data Interchange Format (2009) - <https://vlbi.org/vlbi-standards/vdif/>

The standard 32-byte VDIF Data Frame Header is shown in Figure 3.

	Byte 3		Byte 2	Byte 1	Byte 0
	Bit 31 (MSB)				Bit 0 (LSB)
Word 0	$I_1$	$L_1$	Seconds from reference epoch <sub>30</sub>		
Word 1	Un-assigned <sub>2</sub>		Ref Epoch <sub>6</sub>	Data Frame # within second <sub>24</sub>	
Word 2	$V_3$		$\log_2(\#chans)_5$	Data Frame length (units of 8 bytes) <sub>24</sub>	
Word 3	$C_1$	bits/sample-1 <sub>5</sub>	Thread ID <sub>10</sub>	Station ID <sub>16</sub>	
Word 4	EDV <sub>8</sub>				
Word 5					
Word 6			Extended User Data <sub>32</sub>		
Word 7			Extended User Data <sub>32</sub>		

can describe > 256 Gbps  
for ~33 years continuously

Figure 3: VDIF Data Frame Header format; subscripts are field lengths in bits; byte #s indicate relative byte address within 32-bit word (little endian format)

From VLBI Data Interchange Format (2009) - <https://vlbi.org/vlbi-standards/vdif/>

The standard 32-byte VDIF Data Frame Header is shown in Figure 3.

	Byte 3		Byte 2	Byte 1	Byte 0
	Bit 31 (MSB)				Bit 0 (LSB)
Word 0	$I_1$	$L_1$	Seconds from reference epoch <sub>30</sub>		
Word 1	Un-assigned <sub>2</sub>		Ref Epoch <sub>6</sub>	Data Frame # within second <sub>24</sub>	
Word 2	$V_3$		$\log_2(\#chns)_5$	Data Frame length (units of 8 bytes) <sub>24</sub>	
Word 3	$C_1$	bits/sample-1 <sub>5</sub>	Thread ID <sub>10</sub>	Station ID <sub>16</sub>	
Word 4	EDV <sub>8</sub>			Extended User Data <sub>24</sub>	
Word 5	Extended User Data <sub>32</sub>				
Word 6	Extended User Data <sub>32</sub>				
Word 7	Extended User Data <sub>32</sub>				

Figure 3: VDIF Data Frame Header format; subscripts are field lengths in bits; byte #s indicate relative byte address within 32-bit word (little endian format)

From VLBI Data Interchange Format (2009) - <https://vlbi.org/vlbi-standards/vdif/>

The standard 32-byte VDIF Data Frame Header is shown in Figure 3.

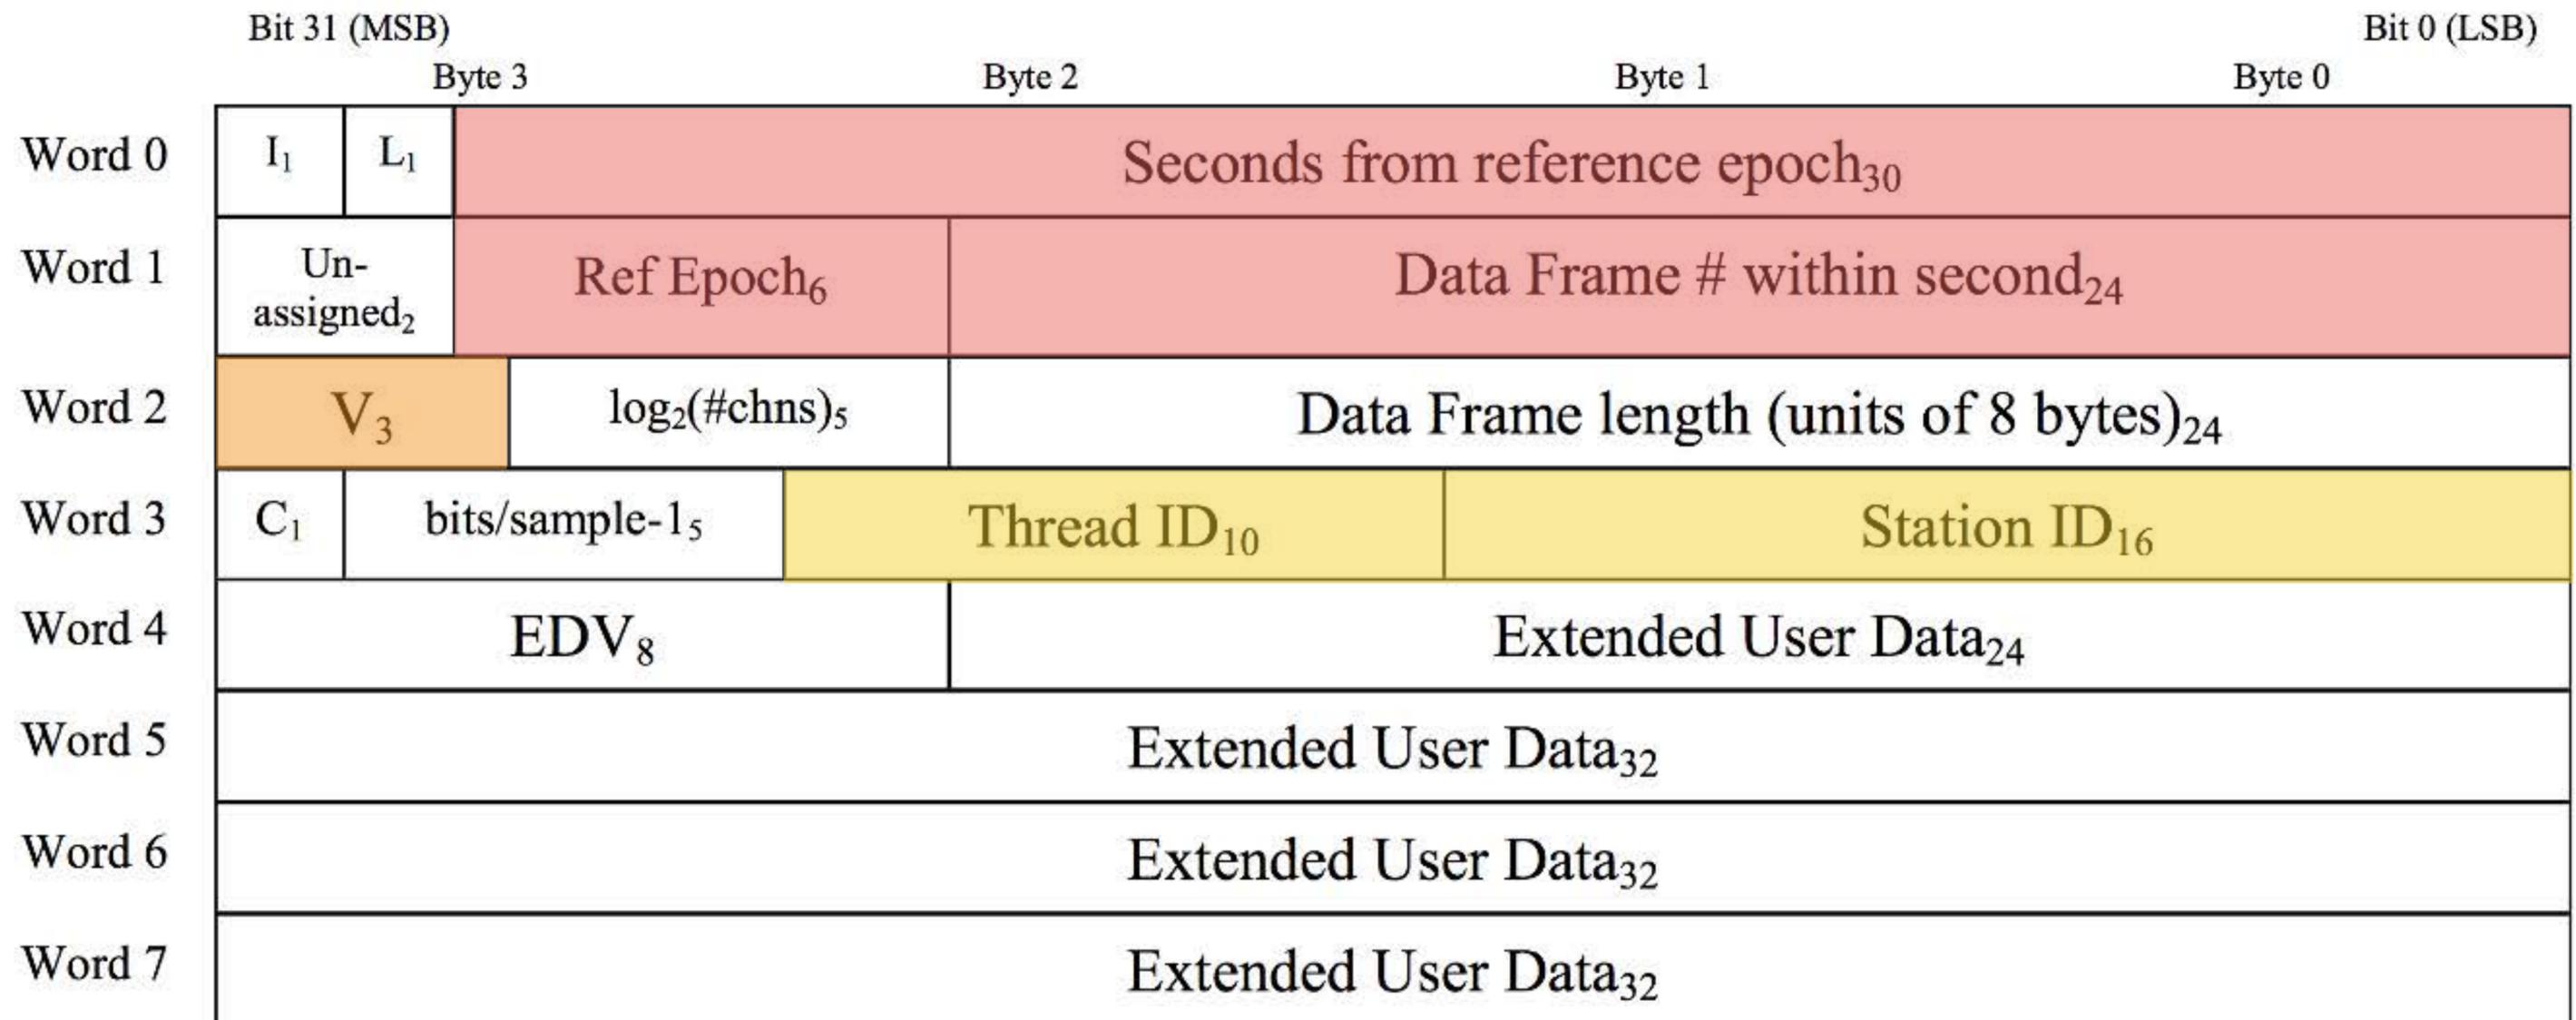


Figure 3: VDIF Data Frame Header format; subscripts are field lengths in bits; byte #s indicate relative byte address within 32-bit word (little endian format)

From VLBI Data Interchange Format (2009) - <https://vlbi.org/vlbi-standards/vdif/>

The standard 32-byte VDIF Data Frame Header is shown in Figure 3.

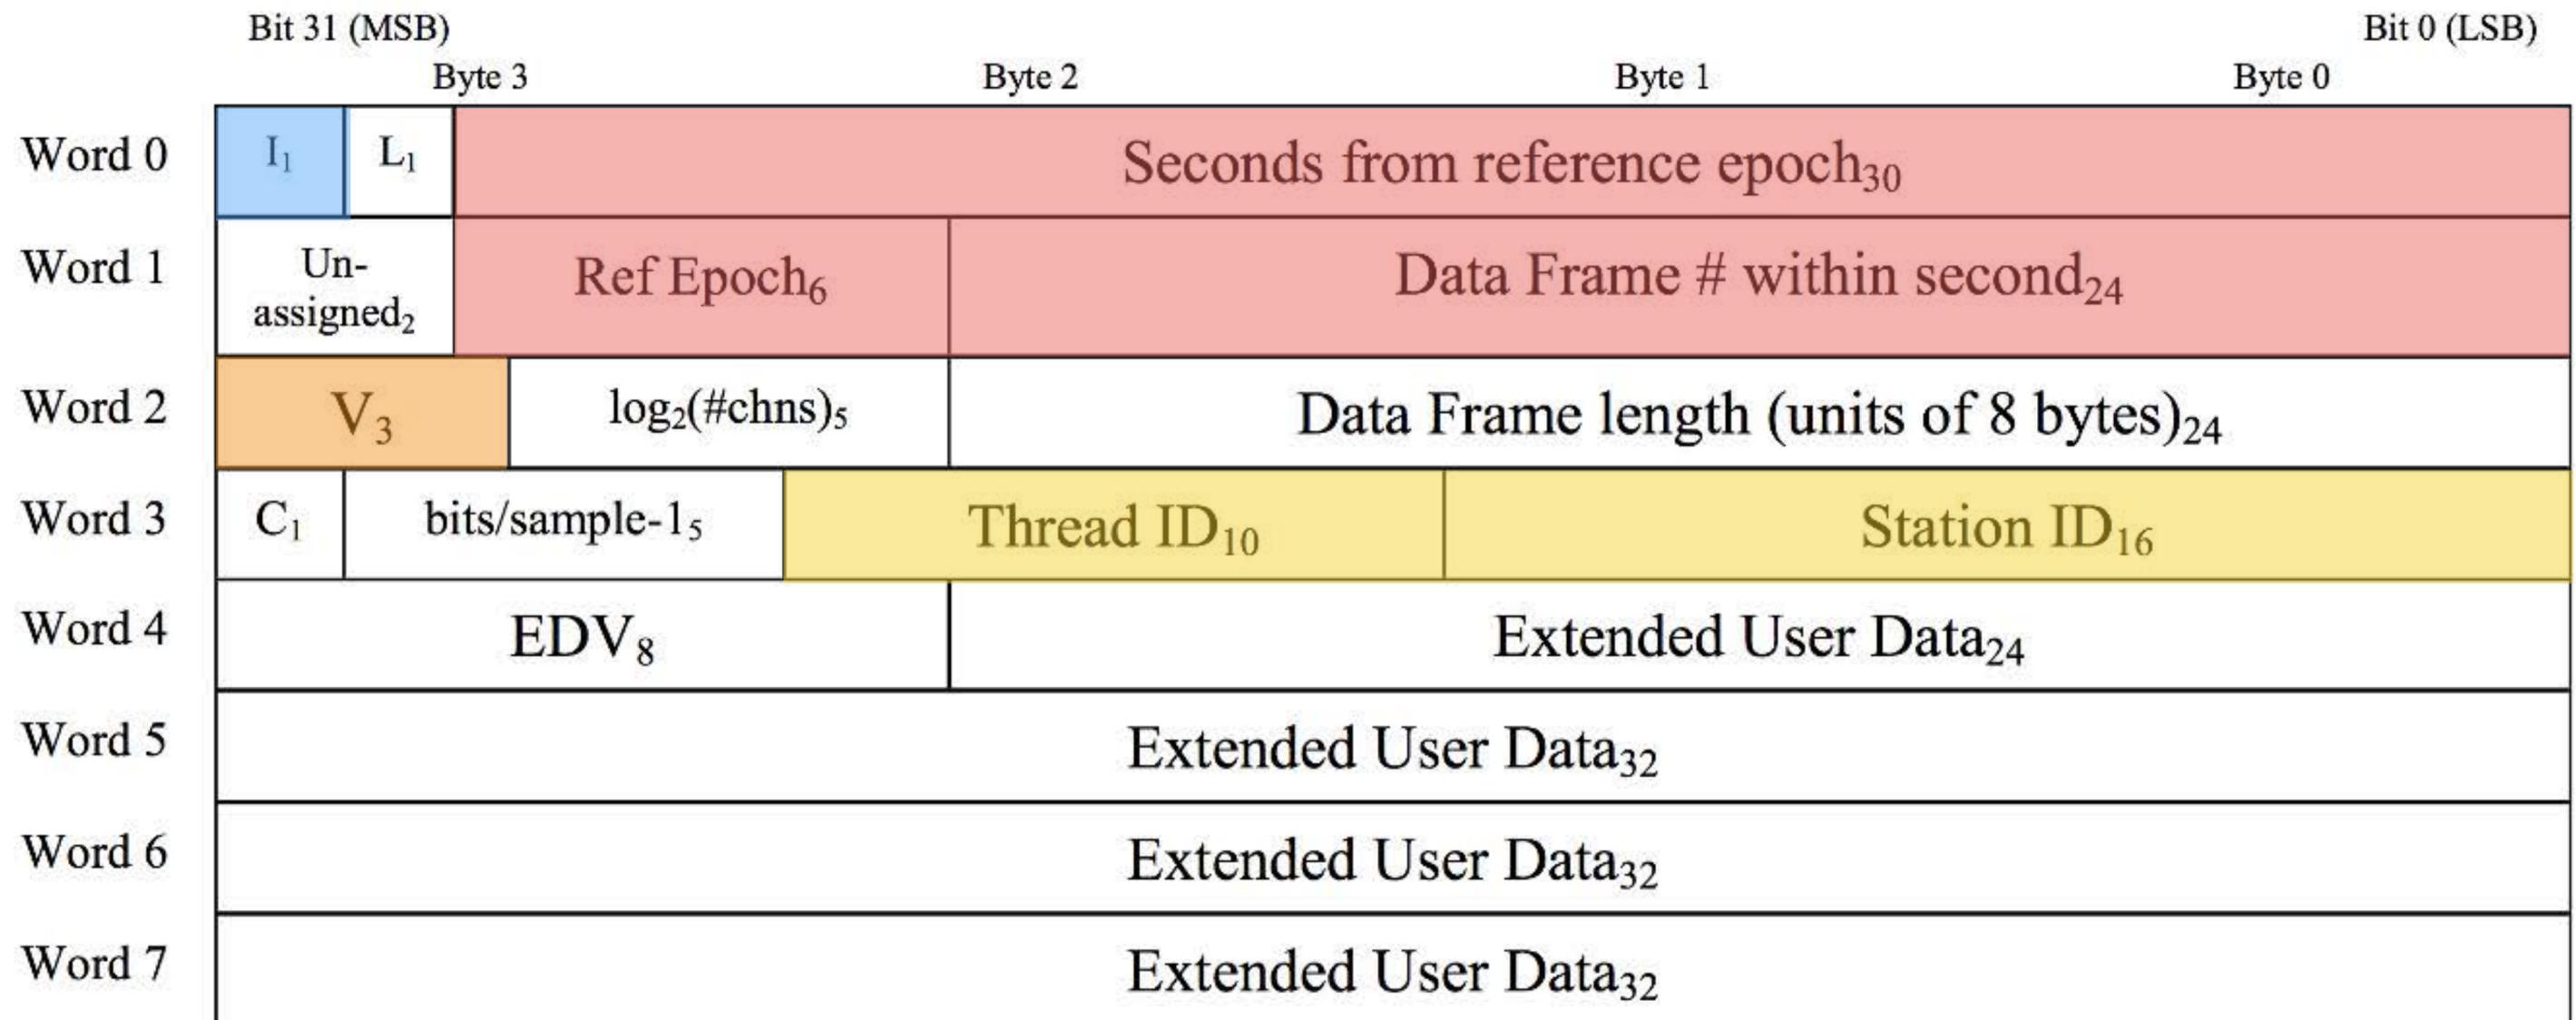


Figure 3: VDIF Data Frame Header format; subscripts are field lengths in bits; byte #s indicate relative byte address within 32-bit word (little endian format)

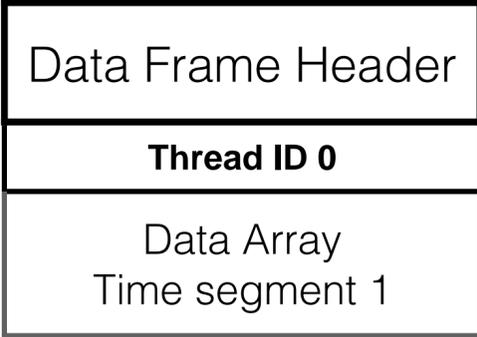
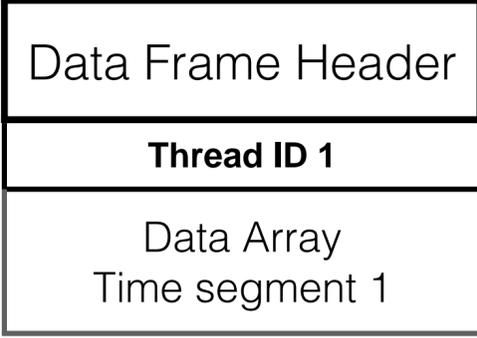
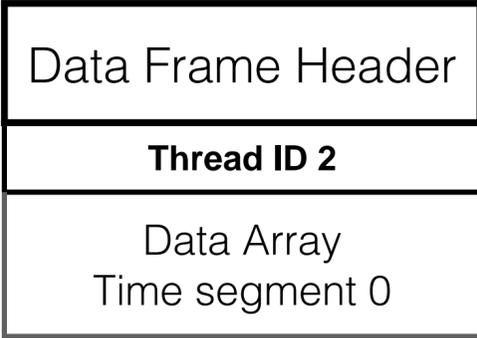
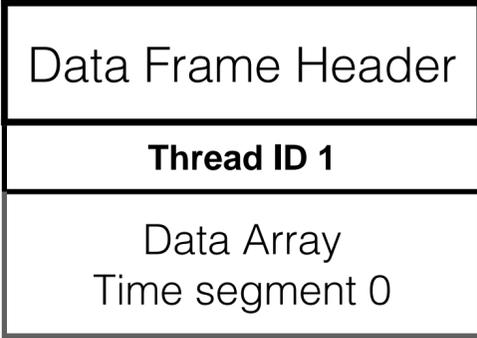
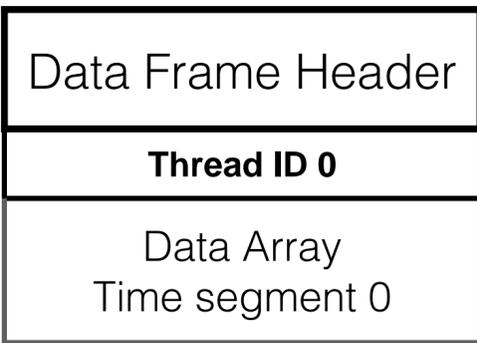
From VLBI Data Interchange Format (2009) - <https://vlbi.org/vlbi-standards/vdif/>

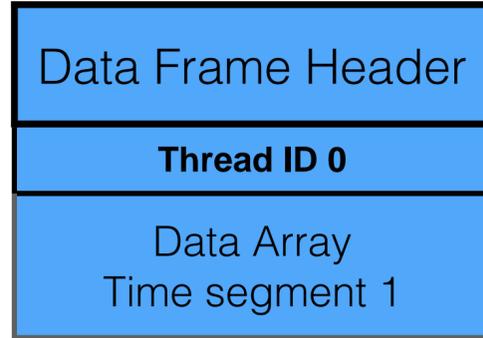
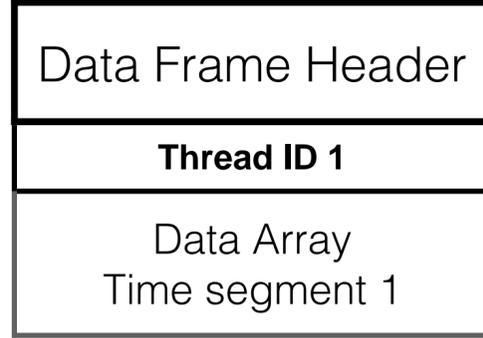
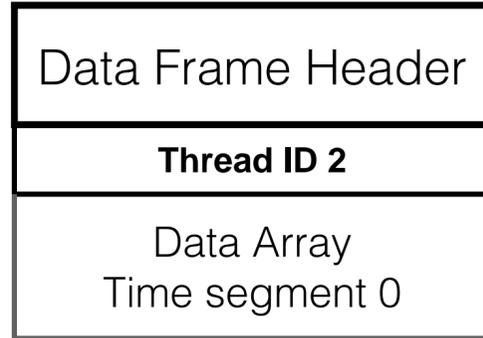
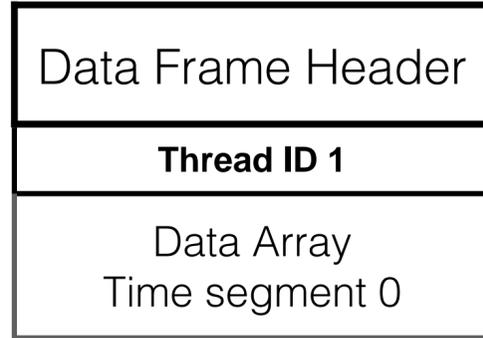
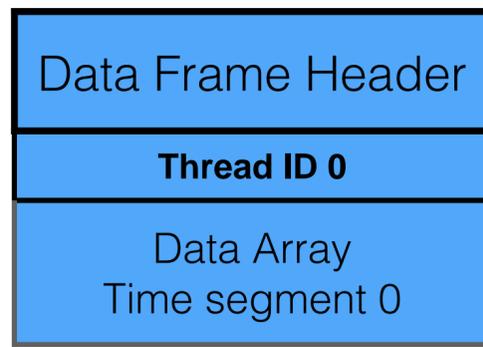
The standard 32-byte VDIF Data Frame Header is shown in Figure 3.

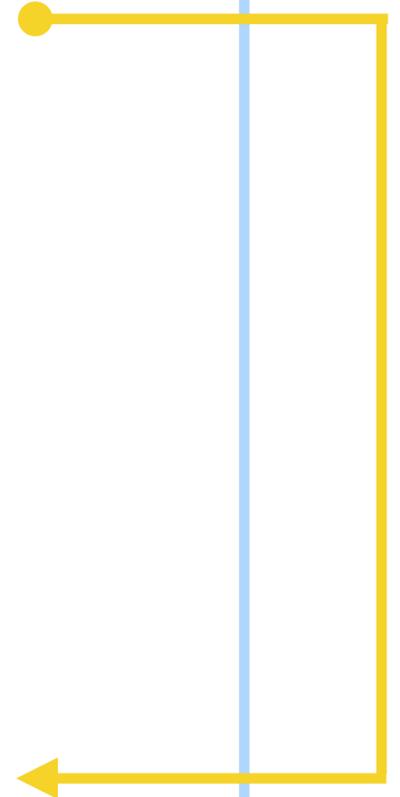
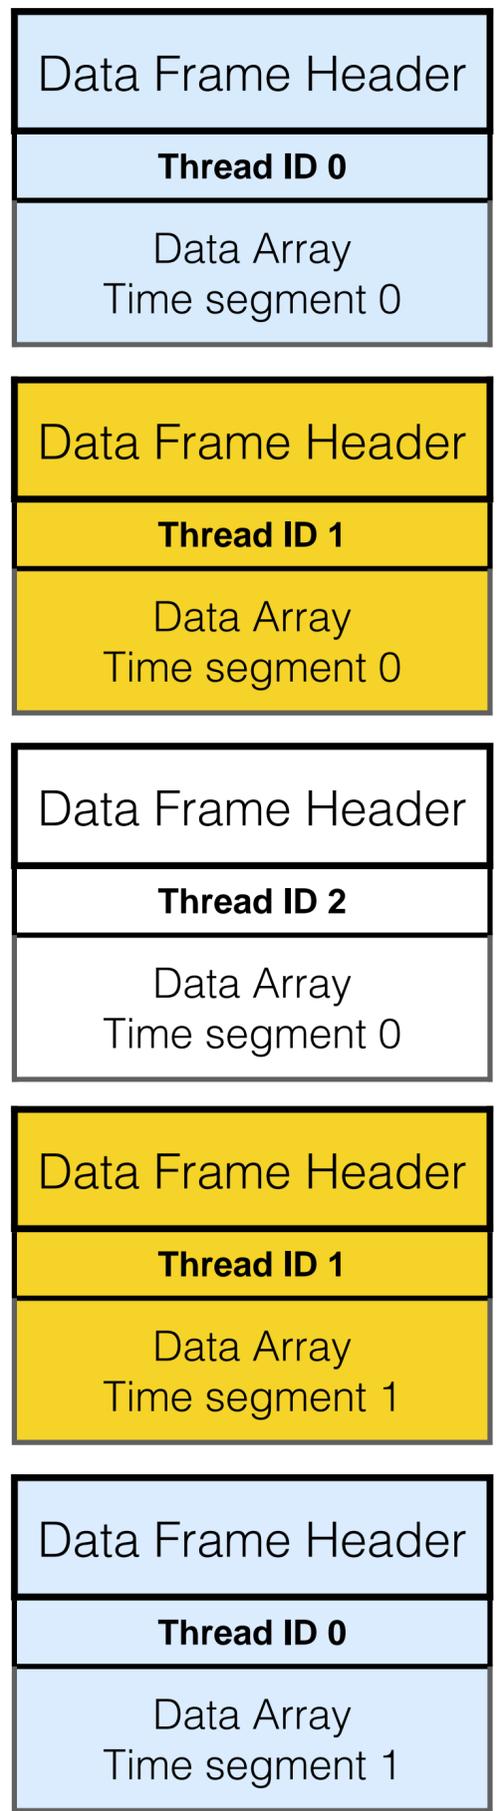
	Byte 3		Byte 2	Byte 1	Byte 0
	Bit 31 (MSB)				Bit 0 (LSB)
Word 0	$I_1$	$L_1$	Seconds from reference epoch <sub>30</sub>		
Word 1	Un-assigned <sub>2</sub>		Ref Epoch <sub>6</sub>	Data Frame # within second <sub>24</sub>	
Word 2	$V_3$		$\log_2(\#chans)_5$	Data Frame length (units of 8 bytes) <sub>24</sub>	
Word 3	$C_1$	bits/sample-1 <sub>5</sub>	Thread ID <sub>10</sub>		Station ID <sub>16</sub>
Word 4	EDV <sub>8</sub>		Extended User Data <sub>24</sub>		
Word 5	Extended User Data <sub>32</sub>				
Word 6	Extended User Data <sub>32</sub>				
Word 7	Extended User Data <sub>32</sub>				

Figure 3: VDIF Data Frame Header format; subscripts are field lengths in bits; byte #s indicate relative byte address within 32-bit word (little endian format)

From VLBI Data Interchange Format (2009) - <https://vlbi.org/vlbi-standards/vdif/>

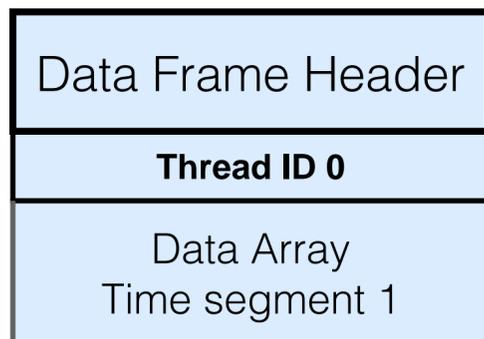
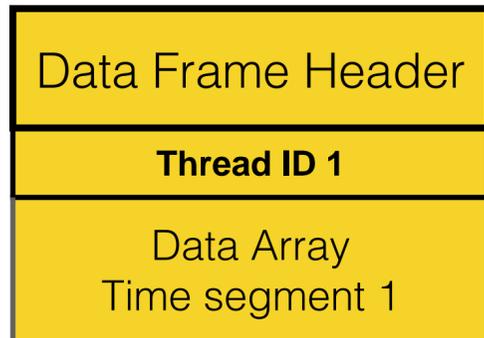
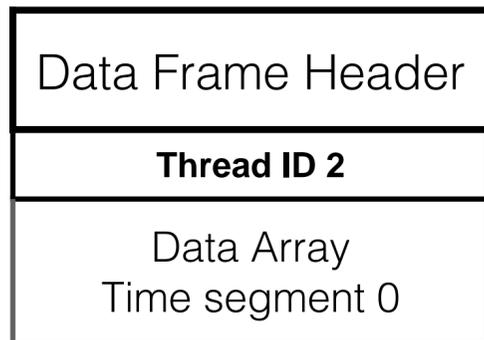
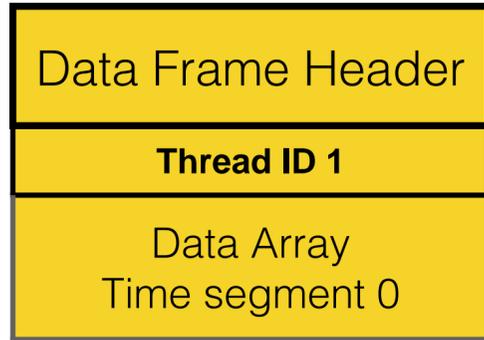
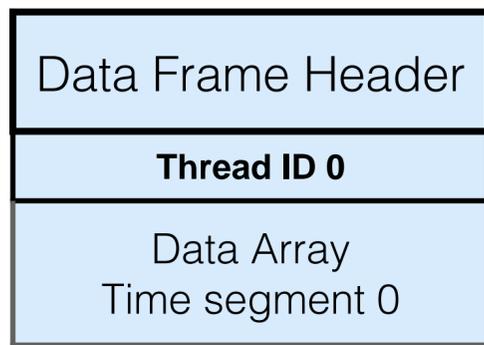






← thread 1, time 0

← thread 1, time 1



*“The VDI specification does not mandate strict Data Frame ordering within a Data Thread, but a best effort should be made to do so.”*

Data is recorded

...

now what?

# Was the recording successful?

What the **CENSORED** did I just record?

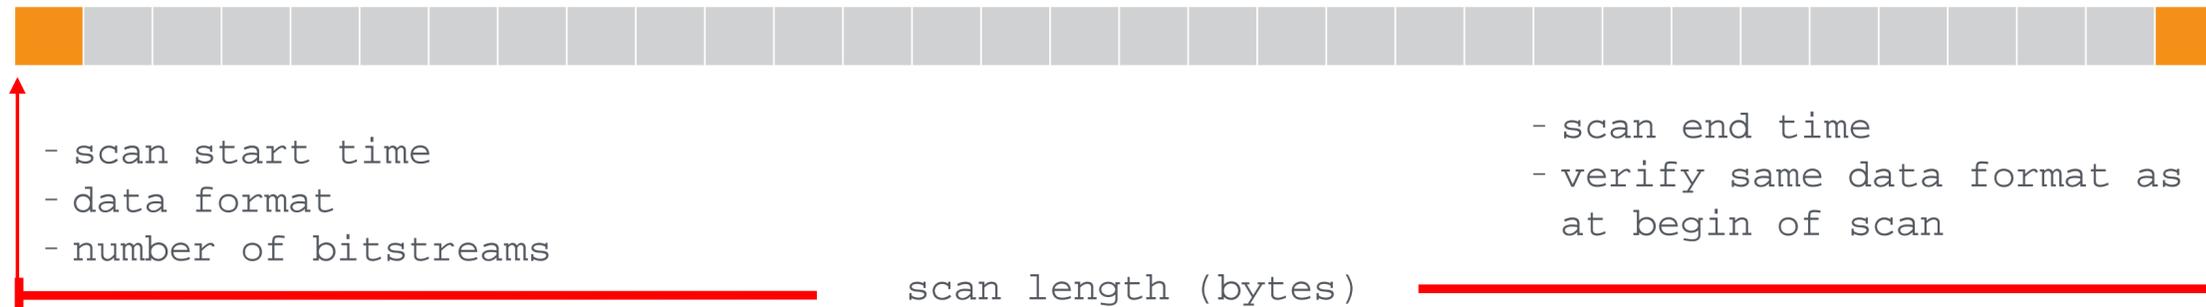
Enter: the **scan\_check?** query

# Was the recording successful?

What the **CENSORED** did I just record?

Enter: the **scan\_check?** query

`EXP_STA_SCAN(vbs/mk6)`



# Was the recording successful?

What the  did I just record?

Enter: the **scan\_check?** query

- detect frame format (MarkIV, VLBA, Mark5B, VDIF)
  - mostly by looking for the magic bit pattern (syncword)
  - VDIF done heuristically so may #FAIL
- detect recorded data rate
- detect if data is missing
  - decode time stamp at begin and end of scan
  - *then*  $((t_{end} - t_{begin}) \times data\_rate) - recording\_length$

# Was the recording successful?

What the  did I just record?

```
# Tell system to find a specific recording
# without argument (default) is 'the last recording you made'
scan_set = n15x2_ef_no0014 ;
!scan_set = 0;
```

```
# What is in that recording?
```

```
scan_check?
```

```
!scan_check? 0 : ? : n15x2_ef_no0014 : VDIF : 2015y297d11h02m01.1389s : 780.786s : 16Mbps : 0 : 8032;
```

- scan name
- detected data type
- start time
- length
- data rate
- format specific information

# Was the recording successful?

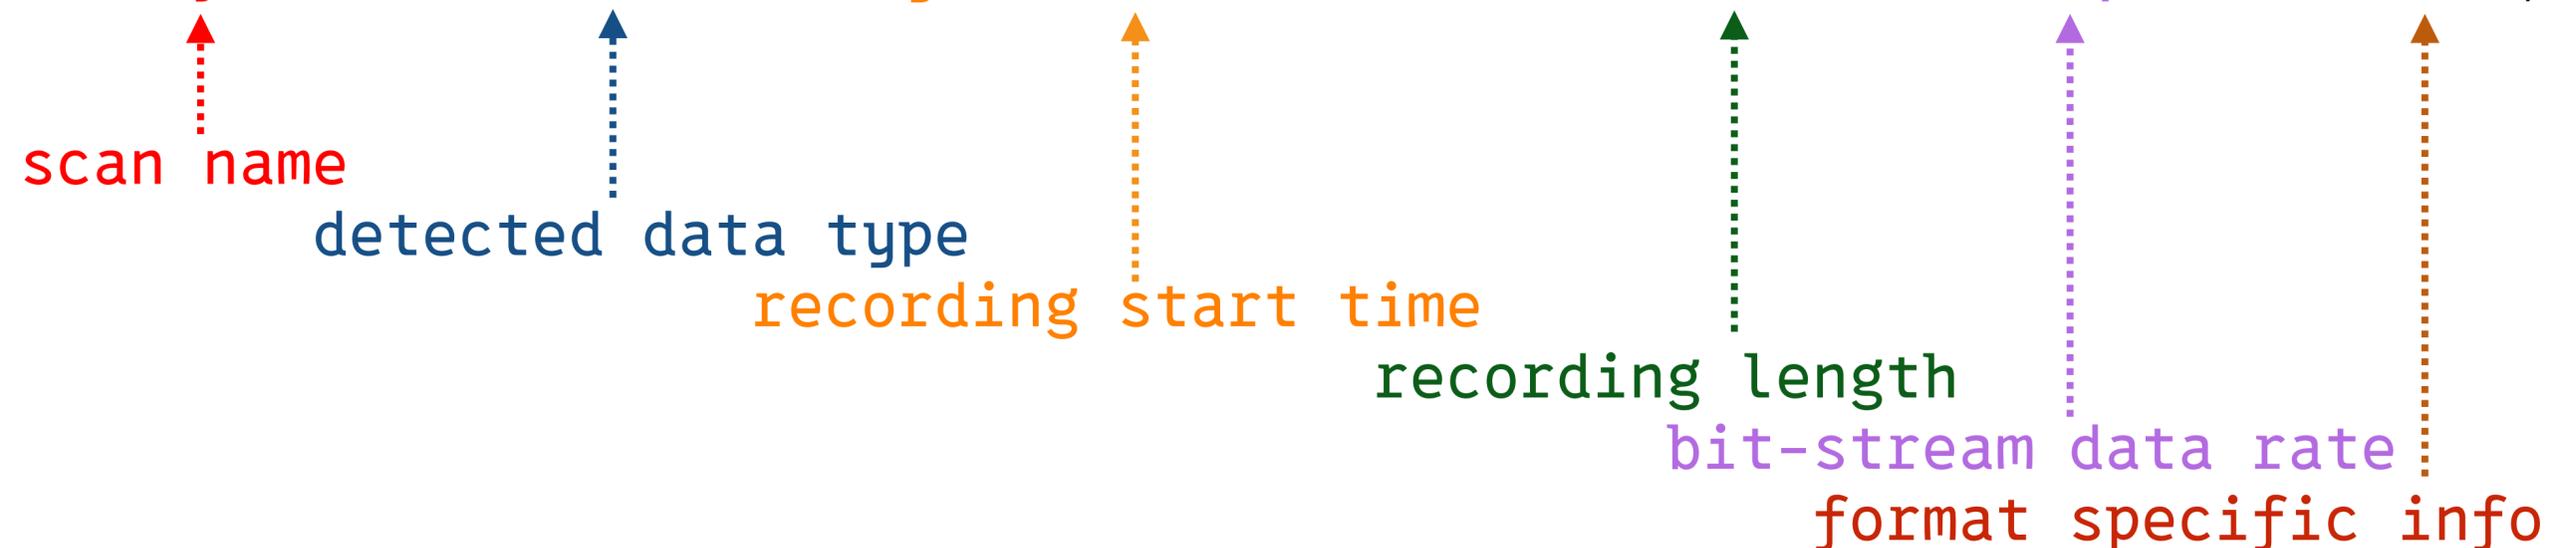
What the **CENSORED** did I just record?

```
# Tell system to find a specific recording
# without argument (default) is 'the last recording you made'
scan_set = n15x2_ef_no0014 ;
!scan_set = 0;
```

**# What is in that recording?**

```
scan_check?
```

```
!scan_check? 0 : ? : n15x2_ef_no0014 : VDIF : 2015y297d11h02m01.1389s : 780.786s : 16Mbps : 0 : 8032;
```



# Adaptable scan\_check algorithm

The simplistic algorithm of "read at start and end" #FAIL

the **scan\_check?** algorithm:

- re-invented + re-implemented

- handle >1 VDIF threads in one "recording"

- handle (very) high data rates

# Adaptable scan\_check algorithm

The simplistic algorithm of "read at start and end" #FAIL

EXP\_STA\_SCAN(vbs/mk6)



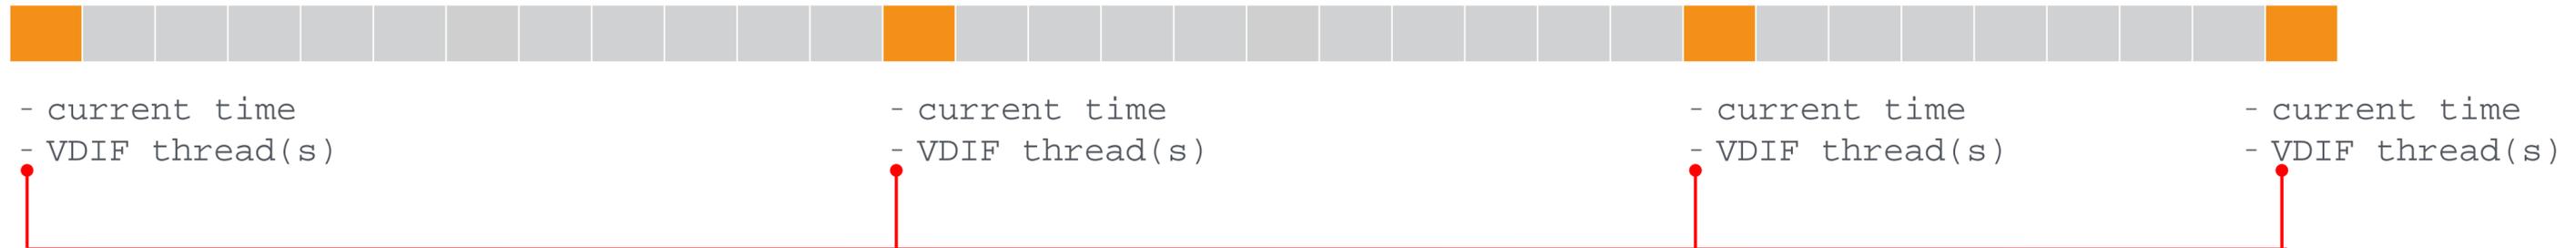
- current time  
- VDIF thread(s)

Combined result:  
• VDIF threads  
• data rate per thread

# Adaptable scan\_check algorithm

The simplistic algorithm of "read at start and end" #FAIL

`EXP_STA_SCAN(vbs/mk6)`



Allows tuning\*

- how much data read in total
  - *more data read = better heuristics, but:*
  - decoding ~ 8MB takes ~ 1-2 s execution time!
- how many sample points/data per sampling point
- can take hint from current `net_protocol` setting
- will be in `jive5ab` 3.2 and up

(\*) See e.g. <https://github.com/jive-vlbi/jive5ab/commit/2304d5fb8e89463a8ea65e151a3eeec7180dbd23>

# Get data to correlator

What options do exist?

# Get data to correlator

What options do exist? (Mark6, ship disk modules)



# Get data to correlator

What options do exist?

```
$> scp /mnt/disk/* io13.mpifr-bonn.de:/path/..
```

???

# Get data to correlator

What options do exist?

```
$> scp /mnt/disk/* io13.mpifr-bonn.de:/path/..
```



# Get data to correlator

What options do exist?

```
$> ftp /mnt/disk/* io13.mpifr-bonn.de:/path/..
```

# Get data to correlator

What options do exist?

```
$> ftp|scp /mnt/disk/* io13.mpifr-bonn.de:/path/.
```

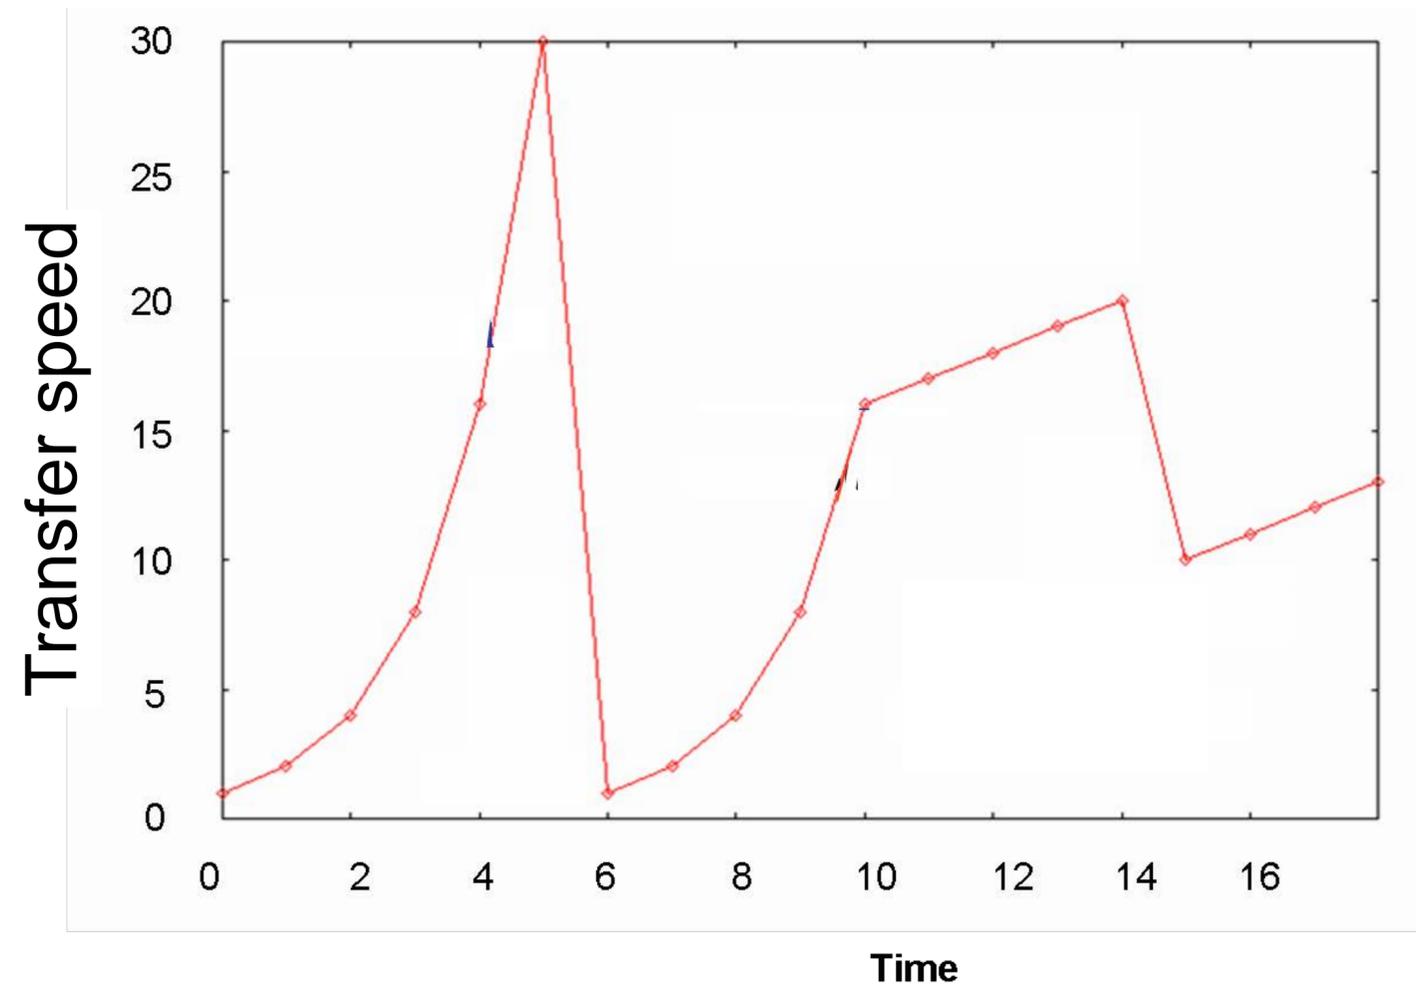


Transmission Control Protocol

[https://en.wikipedia.org/wiki/Transmission\\_Control\\_Protocol](https://en.wikipedia.org/wiki/Transmission_Control_Protocol)

# Get data to correlator

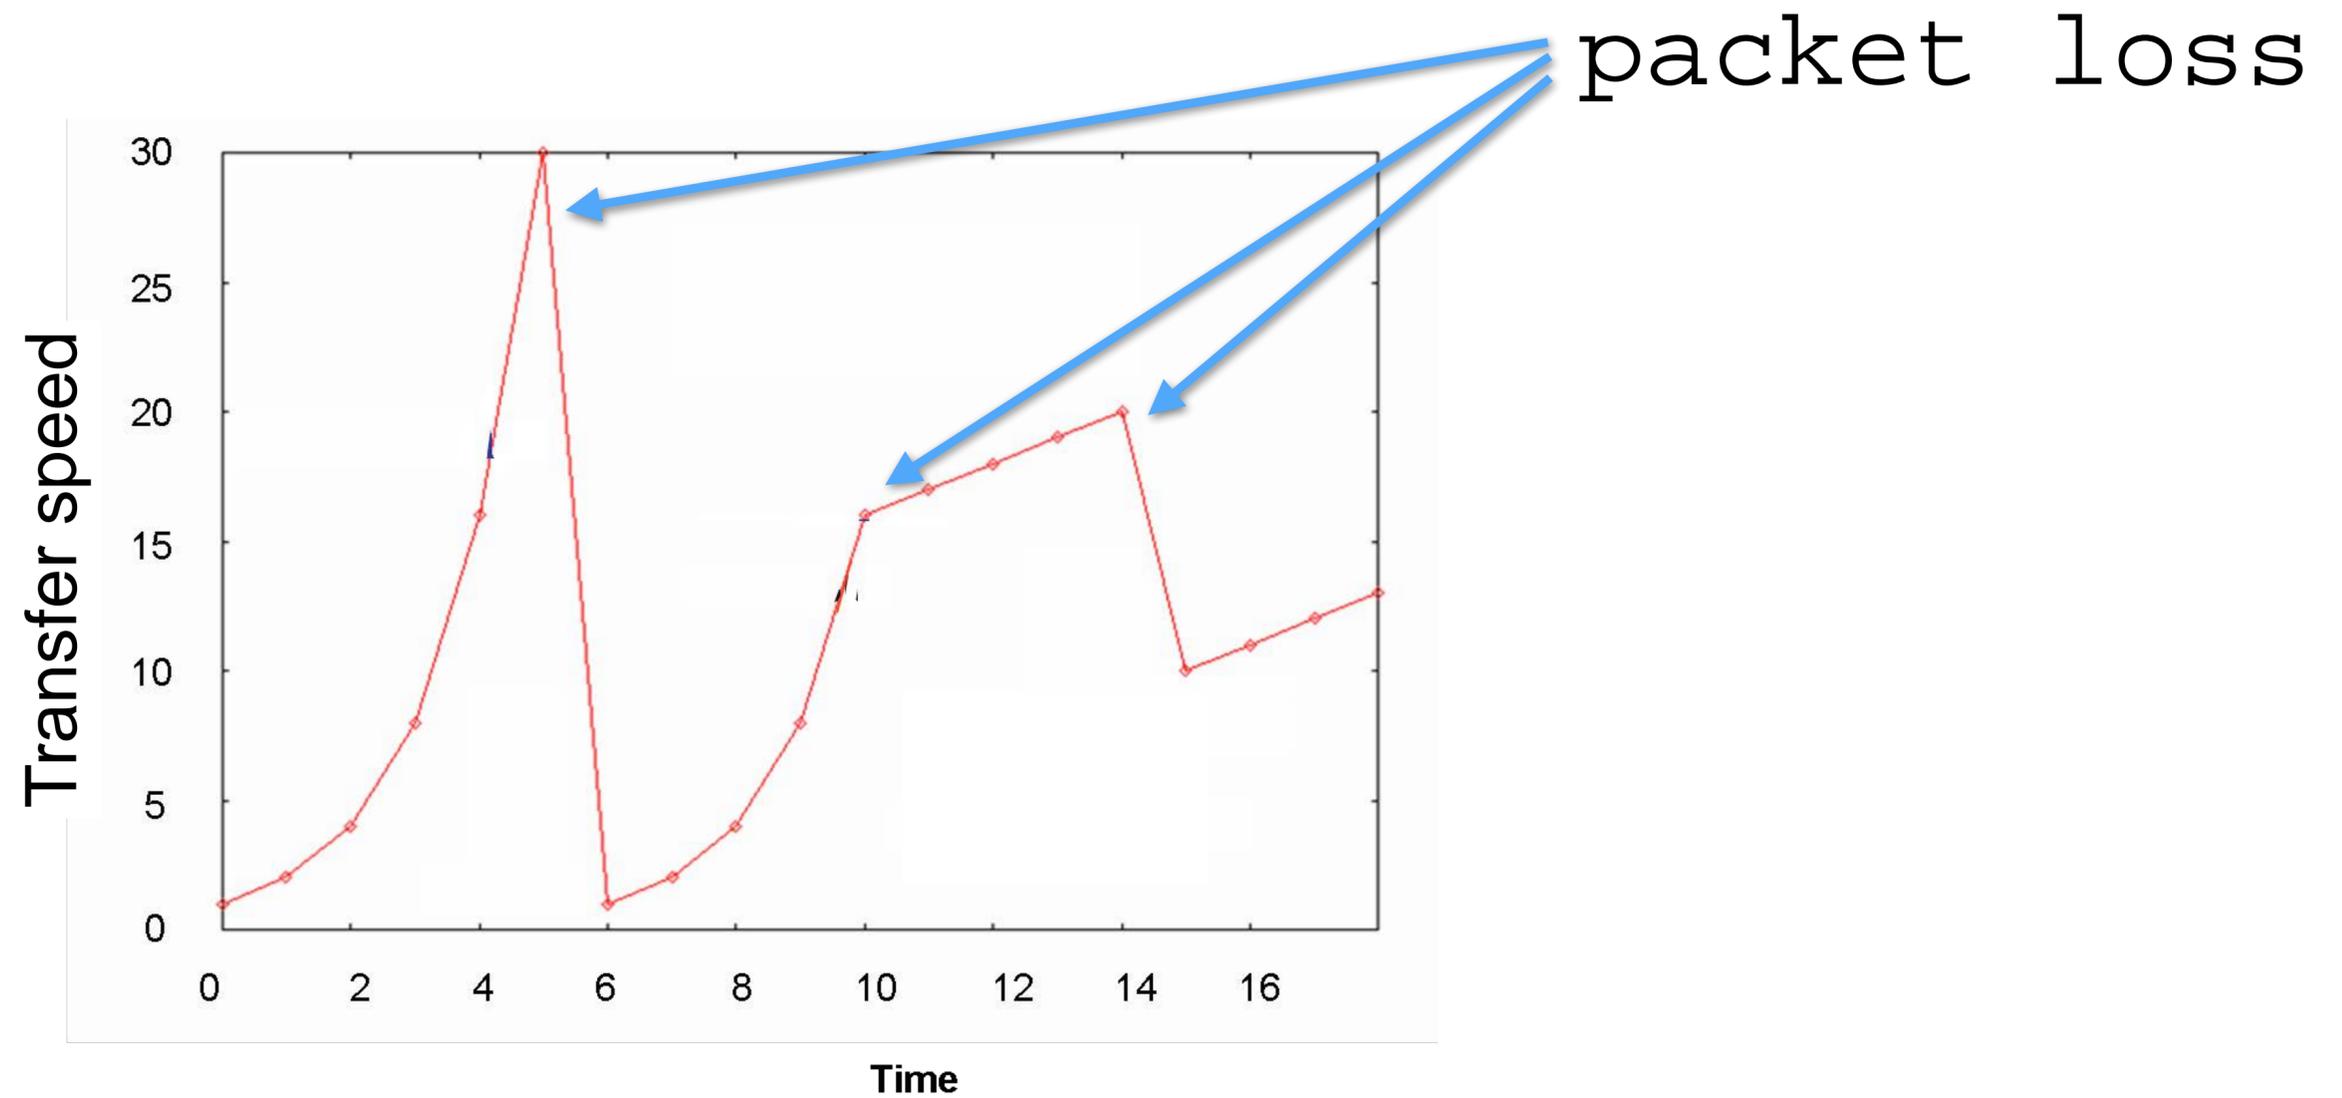
What options do exist?



[https://en.wikipedia.org/wiki/Transmission\\_Control\\_Protocol](https://en.wikipedia.org/wiki/Transmission_Control_Protocol)

# Get data to correlator

What options do exist?



[https://en.wikipedia.org/wiki/Transmission\\_Control\\_Protocol](https://en.wikipedia.org/wiki/Transmission_Control_Protocol)

# Get data to correlator

What options do exist? Use the **UDP Data Transfer** protocol

## UDT

- software library in user space (not in O/S kernel!)
  - require application to actually use the library
- based on connectionless unreliable UDP protocol
- implement TCP-like features:
  - connection oriented
  - reliable

*[https://en.wikipedia.org/wiki/UDP-based\\_Data\\_Transfer\\_Protocol](https://en.wikipedia.org/wiki/UDP-based_Data_Transfer_Protocol)*

# Get data to correlator

What options do exist? Use the **UDP Data Transfer** protocol

## UDT

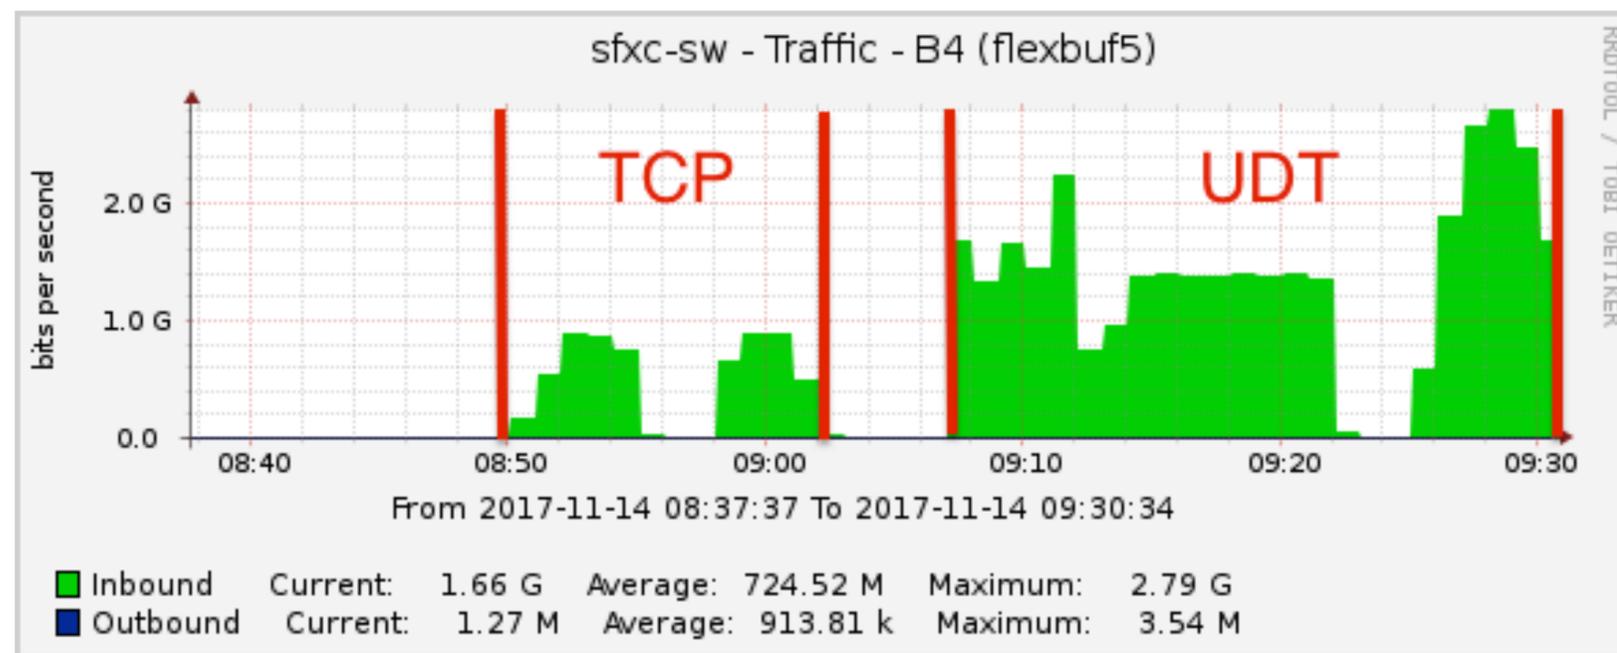
- software library in user space (not in O/S kernel!)
  - require application to actually use the library
- based on connectionless unreliable UDP protocol
- implement TCP-like features:
  - connection oriented
  - reliable
- A LOT faster on long fat links!

*[https://en.wikipedia.org/wiki/UDP-based\\_Data\\_Transfer\\_Protocol](https://en.wikipedia.org/wiki/UDP-based_Data_Transfer_Protocol)*

# Get data to correlator

What options do exist? Use the **UDP Data Transfer** protocol

## UDT



[https://en.wikipedia.org/wiki/UDP-based\\_Data\\_Transfer\\_Protocol](https://en.wikipedia.org/wiki/UDP-based_Data_Transfer_Protocol)

# Get data to correlator

What options do exist? `jive5ab` has `UDT` support built in

`jive5ab`

```
net_protocol = udt : ... ; (*)
```

(\*) <https://github.com/jive-vlbi/jive5ab/blob/master/doc/jive5ab-documentation-1.11.pdf>

# Get data to correlator

What options do exist? `jive5ab` e-shipping



<https://github.com/jive-vlbi/jive5ab/blob/master/scripts/m5copy>

# Get data to correlator

What options do exist? jive5ab e-shipping

```
$> m5copy [options] SRC DST
```

# Get data to correlator

What options do exist? jive5ab e-shipping

```
$> m5copy [options] mk5:///1-10 file:///path/to/
```

# Get data to correlator

What options do exist? `jive5ab e-shipping`

```
$> m5copy [options] mk5:///1-10 file:///path/to/
```

# Get data to correlator

What options do exist? jive5ab e-shipping

```
$> m5copy [options] SRC DST
```

```
mk5://[host][:port]/[module/]scans
```

```
file://[host][:port]/path/to/{dir/|file}
```

```
mk6://[host][:port]/[module/]recording(s)
```

```
vbs://[host][:port]/recording(s)
```

host: host name or IPv4 address (default **localhost**)

port: jive5ab command port (defaults **2620**)

# Get data to correlator

What options do exist? `jive5ab e-shipping`

```
$> m5copy [options] SRC DST
```

## [options]

<code>-p &lt;PORT&gt;</code>	PORT number to use for data connection (default <b>2630</b> )
<code>-m &lt;MTU&gt;</code>	MTU to use for the data transfer (default <b>1500</b> )
<code>-udt</code>	Use UDT as data transfer protocol (default <b>TCP</b> )
<code>-r &lt;RATE&gt;</code>	Maximum transfer rate to use ( <i>only when using UDT</i> )
<code>--resume,</code> <code>--allow_overwrite,</code> <code>--ignore_existing</code>	How to handle file(s)/recording(s) that already exist on the other side

# Get data to correlator

What options do exist? `jive5ab e-shipping`

```
verkouter@Mac> m5copy -udt -mtu 9000 -p 2631 mk6://130.141.242.16:2621/... vbs://flexbuf12.jive.nl/...
```

# Get data to correlator

What options do exist? `jive5ab e-shipping`

```
verkout@Mac> m5copy -udt -mtu 9000 -p 2631 mk6://130.141.242.16:2621/... vbs://flexbuf12.jive.nl/...
```

# Get data to correlator

What options do exist? jive5ab e-shipping

```
verkout@Mac> m5copy -udt -mtu 9000 -p 2631 mk6://130.141.242.16:2621/... vbs://flexbuf12.jive.nl/...
```

# Get data to correlator

What options do exist? `jive5ab e-shipping`

```
verkouter@Mac> m5copy -udt -mtu 9000 -p 2631 mk6://130.141.242.16:2621/... vbs://flexbuf12.jive.nl/...
```

# Get data to correlator

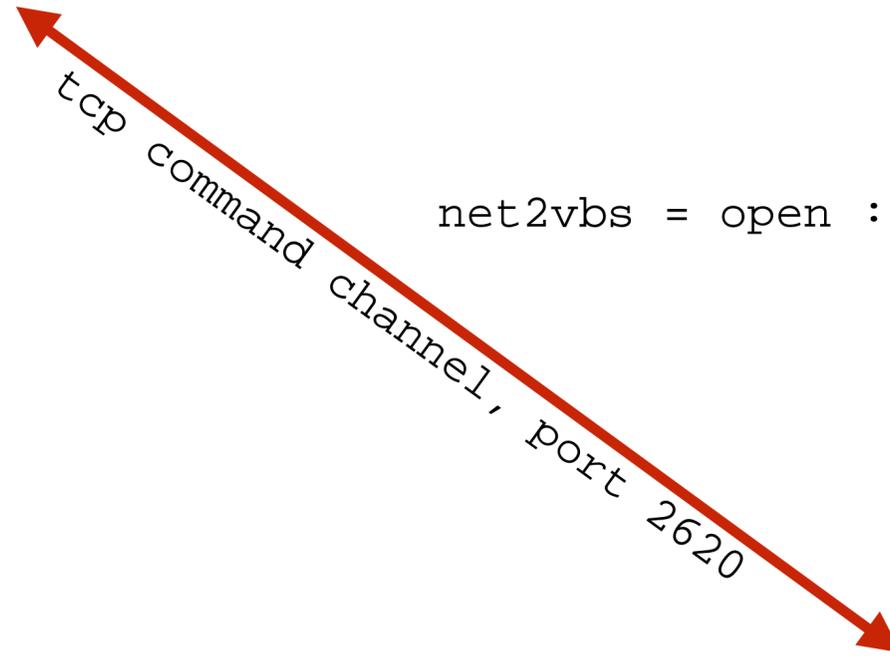
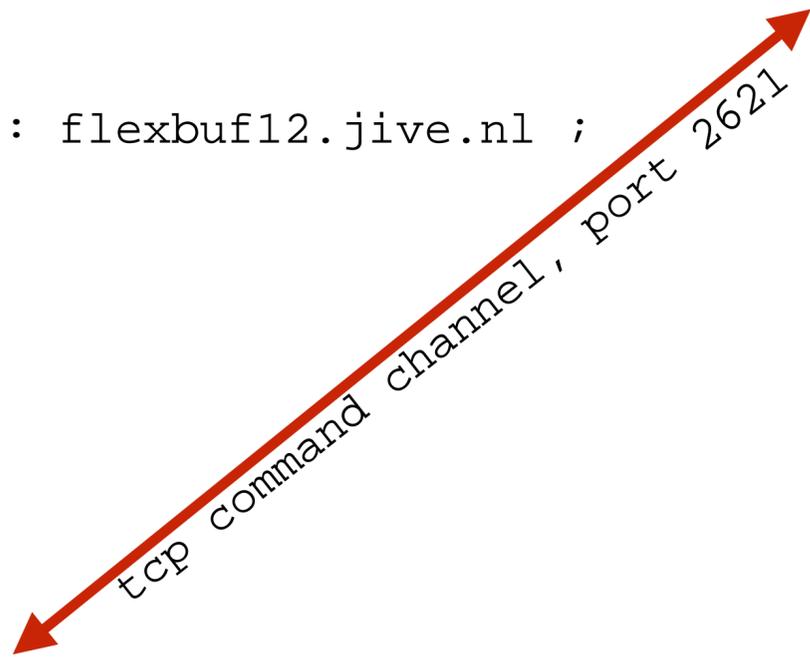
What options do exist? jive5ab e-shipping



```
verkouter@Mac> m5copy -udt -mtu 9000 -p 2631 mk6://130.141.242.16:2621/... vbs://flexbuf12.jive.nl/...
```

```
disk2net = connect : flexbuf12.jive.nl ;
```

```
net2vbs = open : ... ;
```



```
udt data channel port 2631, mtu 9000
```



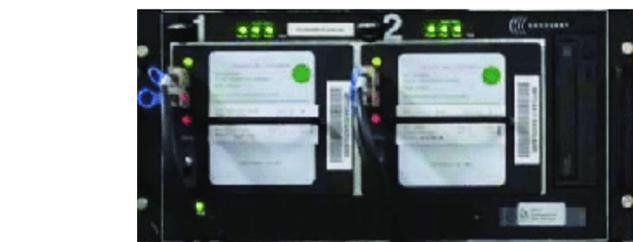
```
$> jive5ab -p 2621 -m 3 -6
```

```
$> jive5ab -m 3
```

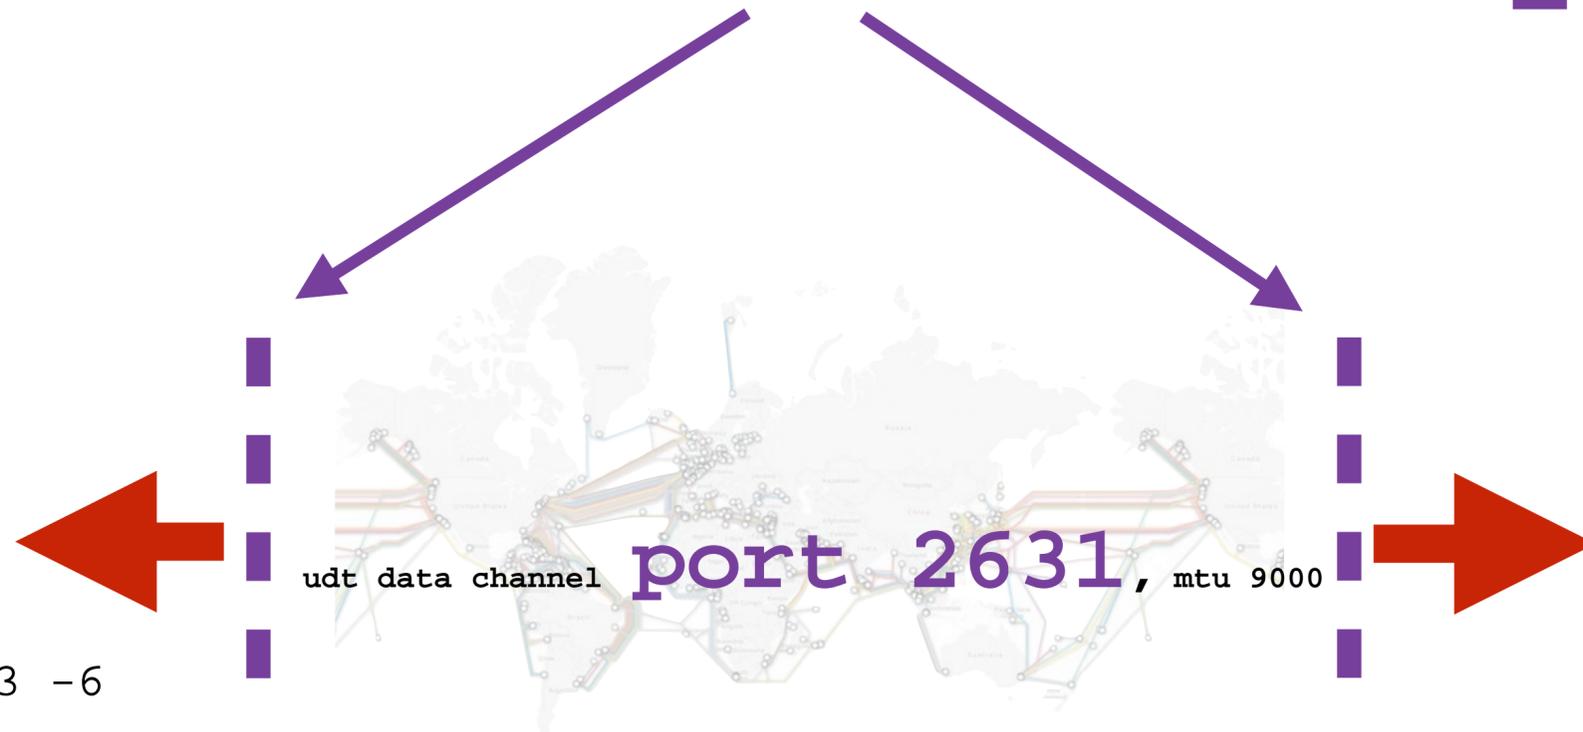
# Get data to correlator

What options do exist? jive5ab e-shipping UDT firewall settings

firewalls must allow  
**bi-directional UDP**  
traffic on data port



```
$> jive5ab -p 2621 -m 3 -6
```



```
$> jive5ab -m 3
```

# DST

# SRC

		Mark5		File		FlexBuff		Mark6	
		lcl	rem	lcl	rem	lcl	rem	lcl	rem
Mark5	lcl								
	rem								
File	lcl								
	rem								
FlexBuff	lcl								
	rem								
Mark6	lcl								
	rem								

DST

SRC

		Mark5		File		FlexBuff		Mark6	
		lcl	rem	lcl	rem	lcl	rem	lcl	rem
Mark5	lcl								
	rem								
File	lcl								
	rem								
FlexBuff	lcl								
	rem								
Mark6	lcl								
	rem								

disk2file

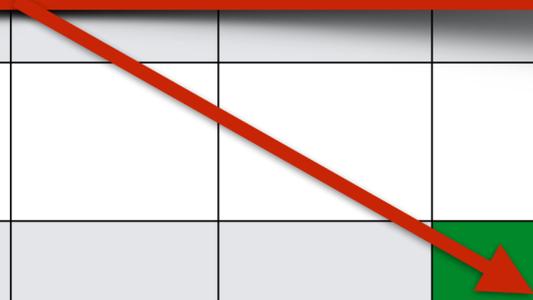


# DST

# SRC

		Mark5		File		FlexBuff		Mark6	
		lcl	rem	lcl	rem	lcl	rem	lcl	rem
Mark5	lcl								
	rem								
File	lcl								
	rem								
FlexBuff	lcl								
	rem								
Mark6	lcl								
	rem								

disk2net + net2vbs



# Get data to correlator

What options do exist? `jive5ab` + `m5copy` GUI frontend

The screenshot shows the 'jive5ab Copy Manager (on ccsbeta)' window. It is divided into two main sections, one for host '0 (10.88.0.50)' and one for host 'aribox (10.88.0.24)'. Each section has a 'Selection' area with a 'Type' dropdown, a 'Reload' button, and a 'Host' dropdown. Below these are tables showing data for each host.

**Host: 0 (10.88.0.50)**

Bank	VSN	#scan	Recording	Size
⊕ A	SHVL-199/16000/1024	474		14.35T
⊖ B				

**Host: aribox (10.88.0.24)**

Disk usage: Total: 82.38T, Used: 60.47T, Free: 17.73T

Experiment	Station	Scan	Size
⊕ EG087A			4.232T
⊕ EG089B			1.911T
⊕ EL053H			13.43G
⊕ EL053I			53.67G
⊕ EL053J			53.67G
⊕ EP087I			2.987T
⊕ F14C2			187.4G
⊕ FBTEST			18.17G
⊕ FT009A			796.4G
⊕ GA035B			4.187T
⊕ GA037A			5.530T
⊕ GA037B			4.096T
⊕ GR035			7.591T
⊕ GR35F			975.9G
⊕ R1680			4.531T
⊕ RAGS14C			1.398T
⊕ RG14C			992.2G
⊕ RG14D			2.277T
⊕ RG14E			1.621T
⊕ RG14H			4.219T

Selection size: no selection

Show file chunks

<https://github.com/jive-vlbi/jive5ab-copy-manager>

# Get data to correlator

What options do exist? `jive5ab` + `m5copy` GUI frontend

The screenshot shows the 'jive5ab Copy Manager (on ccsbeta)' interface. It is divided into two main sections.

**Left Panel:** Shows a file list for 'Mark5' on host '0 (10.88.0.50)'. The table has columns: Bank, VSN, #scan, Recording, and Size. Rows 2 through 6 are selected, with a total selection size of 135.8G (5 items).

Bank	VSN	#scan	Recording	Size
A	SHVL-199/16000/1024	474		14.35T
1			gkftp_wb_no0002	35.70G
2			n16l2_wb_no0001	17.85G
3			n16l2_wb_no0002	17.85G
4			n16l2_wb_no0003	35.76G
5			n16l2_wb_no0004	32.18G
6			n16l2_wb_no0005	32.18G
7			n16l2_wb_no0006	50.09G
8			n16l2_wb_no0008	17.85G
9			n16l2_wb_no0009	35.76G
10			n16l2_wb_no0010	32.18G
11			n16l2_wb_no0011	32.18G
12			n16l2_wb_no0012	50.09G
13			n16l2_wb_no0014	17.85G
14			n16l2_wb_no0015	35.76G
15			n16l2_wb_no0016	32.18G
16			n16l2_wb_no0017	50.09G
17			n16l2_wb_no0019	17.85G
18			n16l2_wb_no0020	17.79G
19			n16l2_wb_no0021	17.85G
20			gs037a_wb_no0037	35.70G
21			gs037a_wb_no0039	35.70G
22			gs037a_wb_no0041	35.70G
23			gs037a_wb_no0043	35.70G
24			gs037a_wb_no0045	35.70G
25			gs037a_wb_no0047	35.70G
26			gs037a_wb_no0049	35.70G
27			gs037a_wb_no0051	35.70G
28			gs037a_wb_no0053	35.70G
29			gs037a_wb_no0055	35.70G
30			gs037a_wb_no0057	10.63G
31			gs037a_wb_no0058	34.51G

**Right Panel:** Shows disk usage for 'FlexBuff' on host 'aribox (10.88.0.24)'. It displays a tree view of experiments and scans. The 'TEST1' experiment is expanded to show a scan 'no0001' with 11 chunks of varying sizes (256.0M to 19.71M). The total disk usage is 60.47T, with 17.72T free.

Disk usage: Total: 82.38T, Used: 60.47T, Free: 17.72T

Experiment	Station	Scan	Chunk	Size
FBTEST				18.17G
FT009A				796.4G
GA035B				4.187T
GA037A				5.530T
GA037B				4.096T
GR035				7.591T
GR35F				975.9G
R1680				4.531T
RAGS14C				1.398T
RG14C				992.2G
RG14D				2.277T
RG14E				1.621T
RG14H				4.219T
TEST1				2.519G
	Wb	no0001		2.519G
			00000000 on disk11	256.0M
			00000001 on disk31	256.0M
			00000002 on disk21	256.0M
			00000003 on disk32	256.0M
			00000004 on disk8	256.0M
			00000005 on disk13	256.0M
			00000006 on disk16	256.0M
			00000007 on disk4	256.0M
			00000008 on disk14	256.0M
			00000009 on disk5	256.0M
			00000010 on disk6	19.71M

<https://github.com/jive-vlbi/jive5ab-copy-manager>

# Get data to correlator

What options do exist? jive5ab + m5copy GUI frontend

The screenshot shows the 'jive5ab Copy Manager (on ccsbeta)' window. It is divided into two main sections. The left section is titled 'Selection' and shows a table of files with columns: Bank, VSN, #scan, Recording, and Size. A context menu is open over the table, showing options: 'Scan check', 'Copy to ...', and 'Clear selection Esc'. The right section is also titled 'Selection' and shows a 'Disk usage' summary: Total: 82.38T, Used: 60.47T, Free: 17.73T. Below this is a table of experiments with columns: Experiment, Station, Scan, Chunk, and Size. A 'Show file chunks' checkbox is checked at the bottom right.

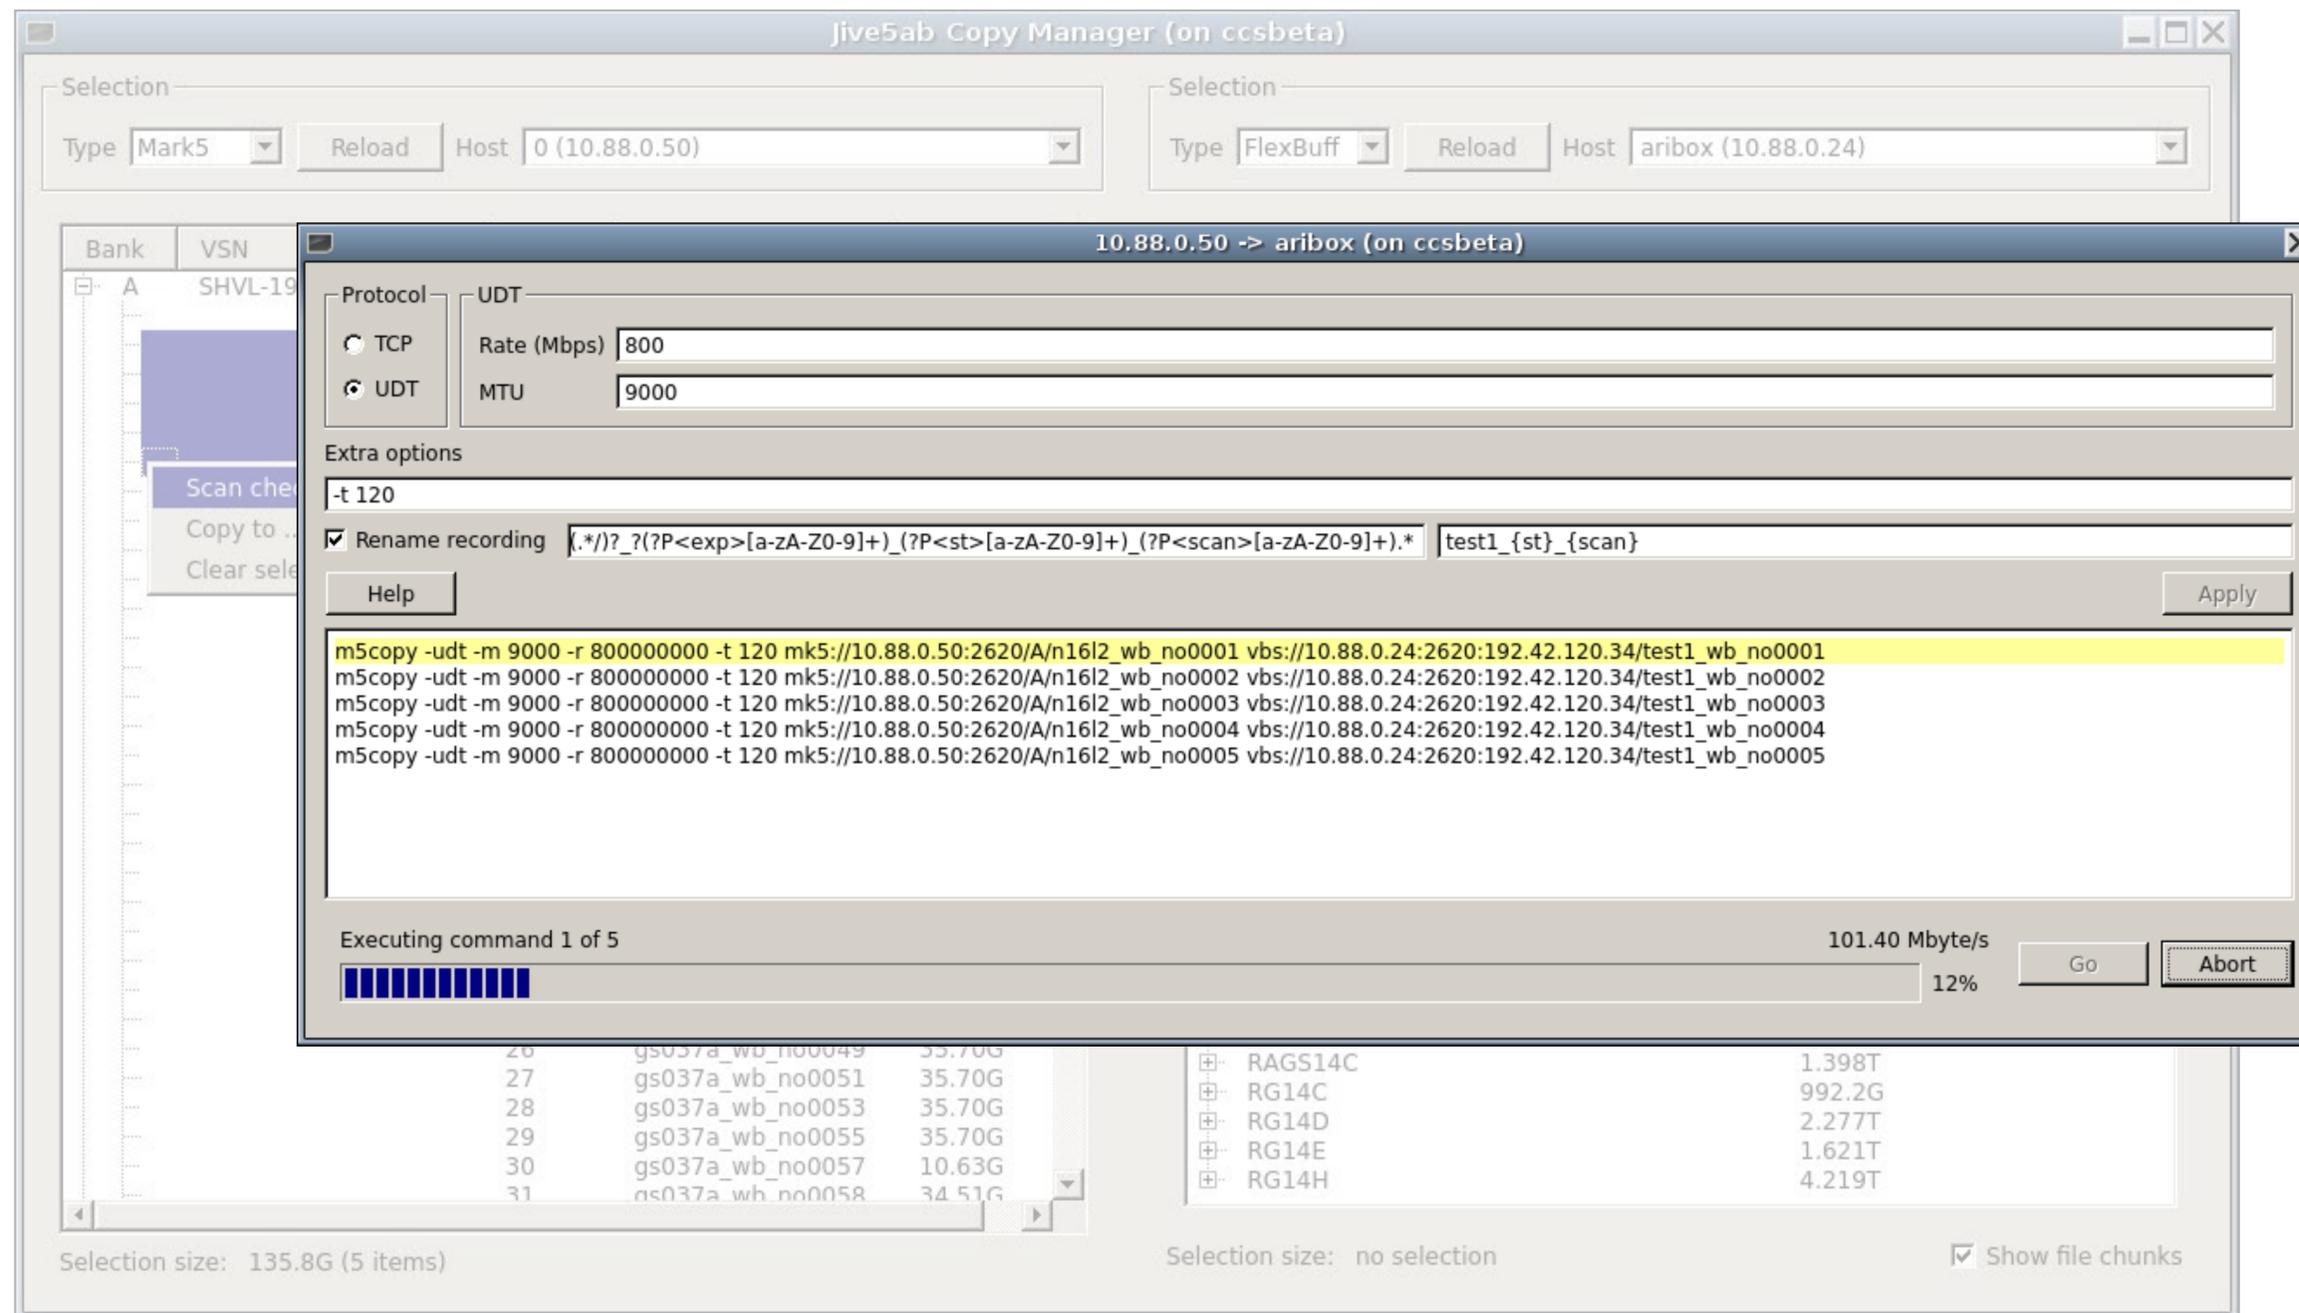
Bank	VSN	#scan	Recording	Size
A	SHVL-199/16000/1024	474		14.35T
		1	gkftp_wb_no0002	35.70G
		2	n16l2_wb_no0001	17.85G
		3	n16l2_wb_no0002	17.85G
		4	n16l2_wb_no0003	35.76G
		5	n16l2_wb_no0004	32.18G
		6	n16l2_wb_no0005	32.18G
		7	n16l2_wb_no0006	50.09G
		8	n16l2_wb_no0008	17.85G
		9	n16l2_wb_no0009	35.76G
		10	n16l2_wb_no0010	32.18G
		11	n16l2_wb_no0011	32.18G
		12	n16l2_wb_no0012	50.09G
		13	n16l2_wb_no0014	17.85G
		14	n16l2_wb_no0015	35.76G
		15	n16l2_wb_no0016	32.18G
		16	n16l2_wb_no0017	50.09G
		17	n16l2_wb_no0019	17.85G
		18	n16l2_wb_no0020	17.79G
		19	n16l2_wb_no0021	17.85G
		20	gs037a_wb_no0037	35.70G
		21	gs037a_wb_no0039	35.70G
		22	gs037a_wb_no0041	35.70G
		23	gs037a_wb_no0043	35.70G
		24	gs037a_wb_no0045	35.70G
		25	gs037a_wb_no0047	35.70G
		26	gs037a_wb_no0049	35.70G
		27	gs037a_wb_no0051	35.70G
		28	gs037a_wb_no0053	35.70G
		29	gs037a_wb_no0055	35.70G
		30	gs037a_wb_no0057	10.63G
		31	gs037a_wb_no0058	34.51G

Experiment	Station	Scan	Chunk	Size
EG087A				4.232T
EG089B				1.911T
EL053H				13.43G
EL053I				53.67G
EL053J				53.67G
EP087I				2.987T
F14C2				187.4G
FBTEST				18.17G
FT009A				796.4G
GA035B				4.187T
	Ef			264.8G
	Jb			693.7G
	Mc			754.0G
	Nt			307.0G
	Sh			773.9G
	Tr			749.9G
	Ys			744.5G
GA037A				5.530T
GA037B				4.096T
GR035				7.591T
GR35F				975.9G
R1680				4.531T
RAGS14C				1.398T
RG14C				992.2G
RG14D				2.277T
RG14E				1.621T
RG14H				4.219T

<https://github.com/jive-vlbi/jive5ab-copy-manager>

# Get data to correlator

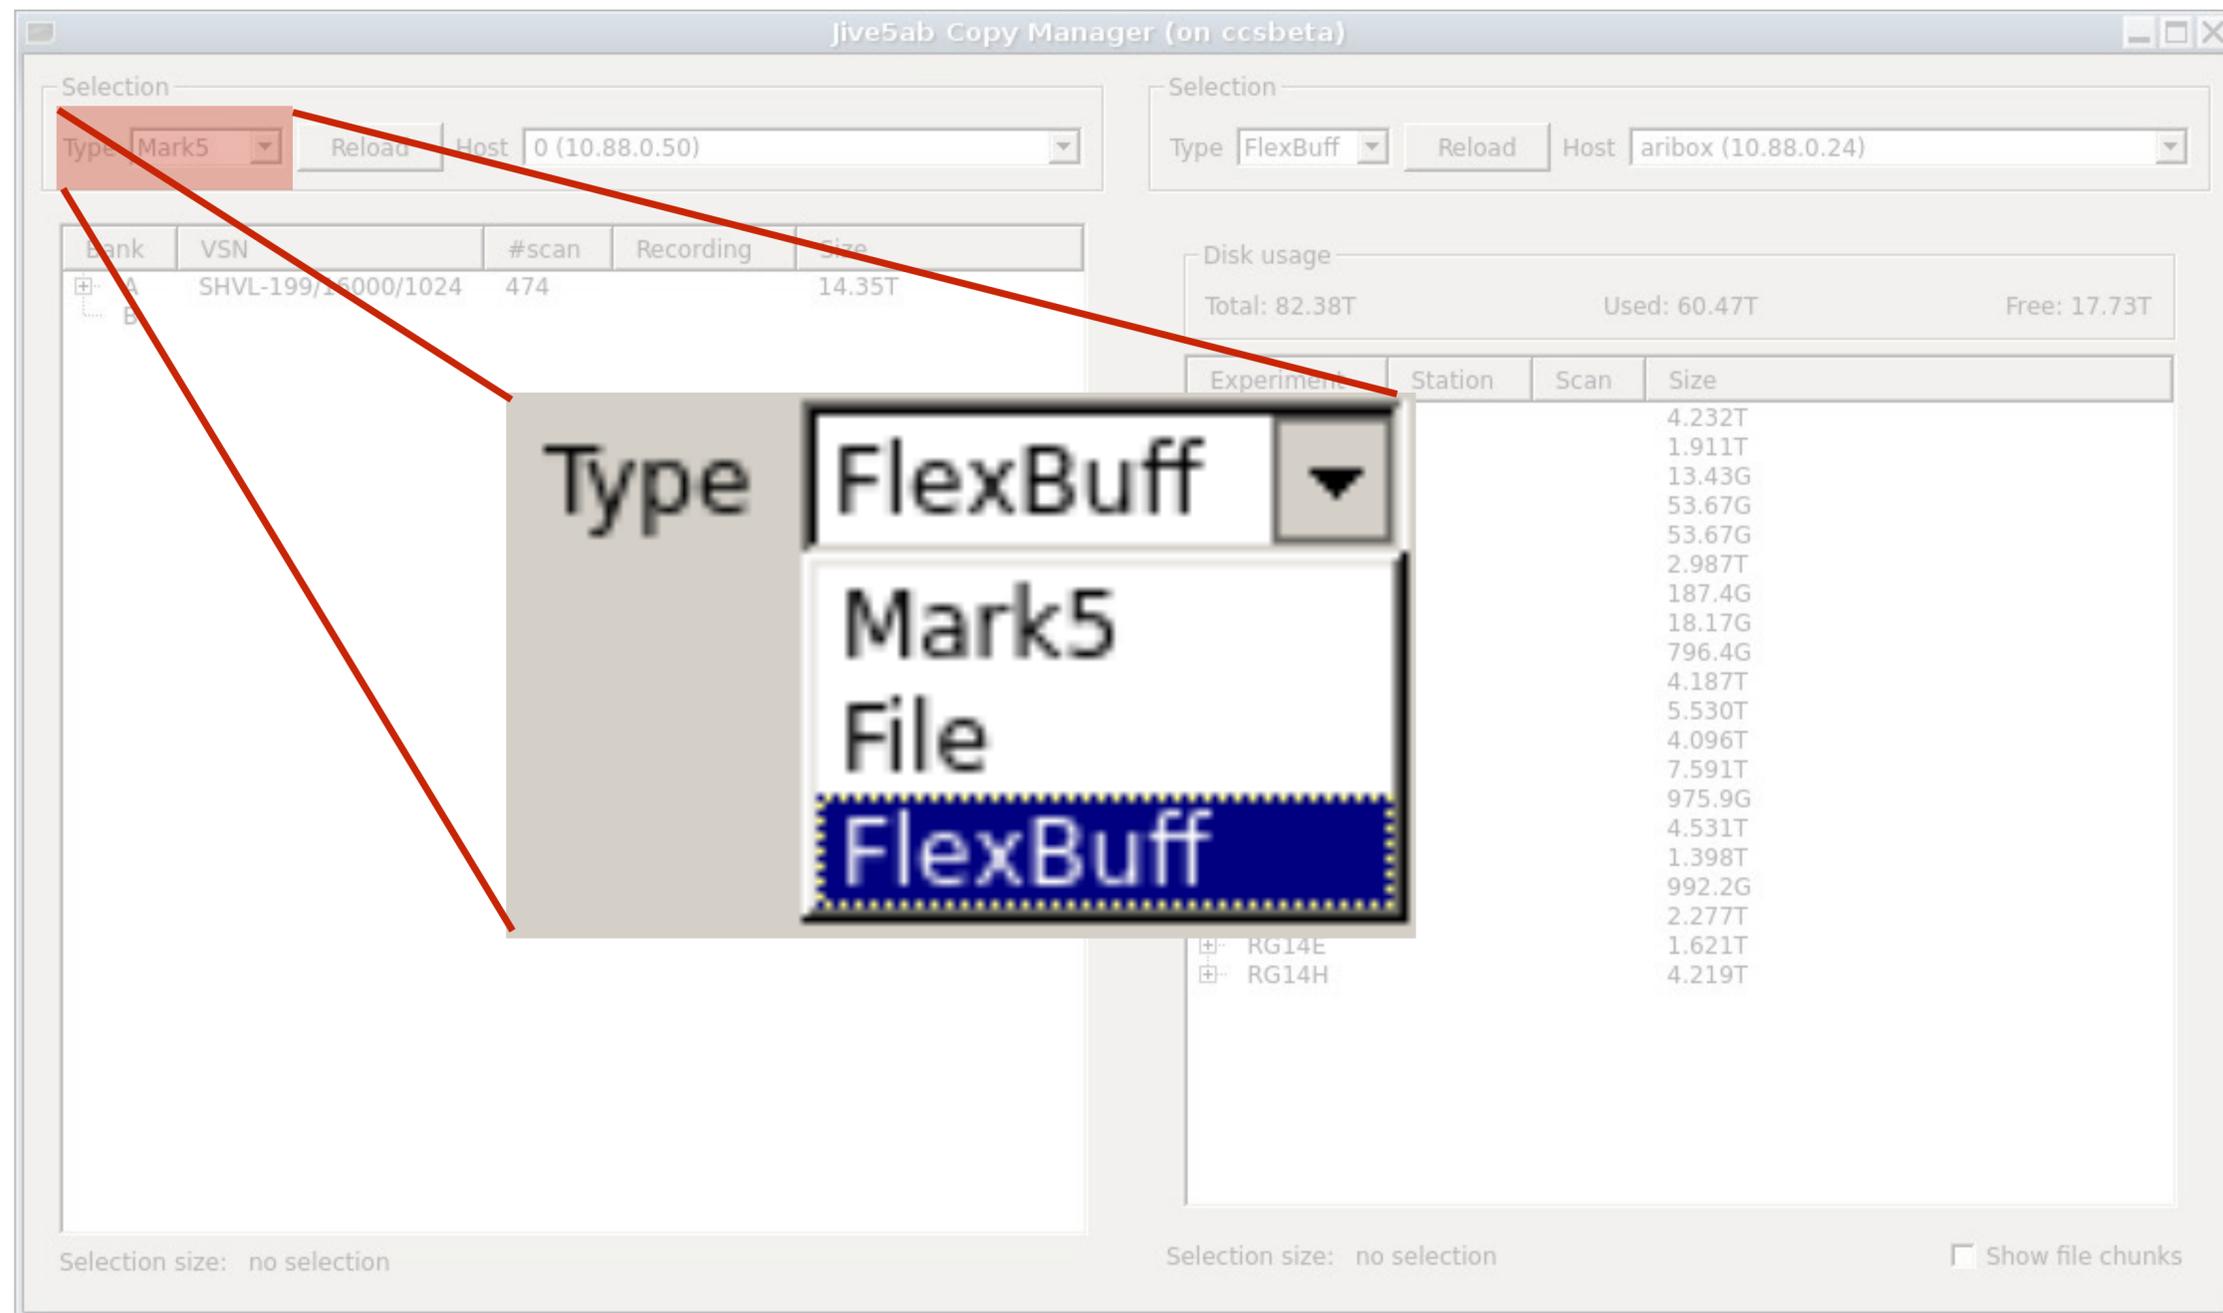
What options do exist? jive5ab + m5copy GUI frontend



<https://github.com/jive-vlbi/jive5ab-copy-manager>

# Get data to correlator

What options do exist? `jive5ab` + `m5copy` **GUI frontend**



<https://github.com/jive-vlbi/jive5ab-copy-manager>

# Get data to correlator

What options do exist? jive5ab + m5copy GUI frontend

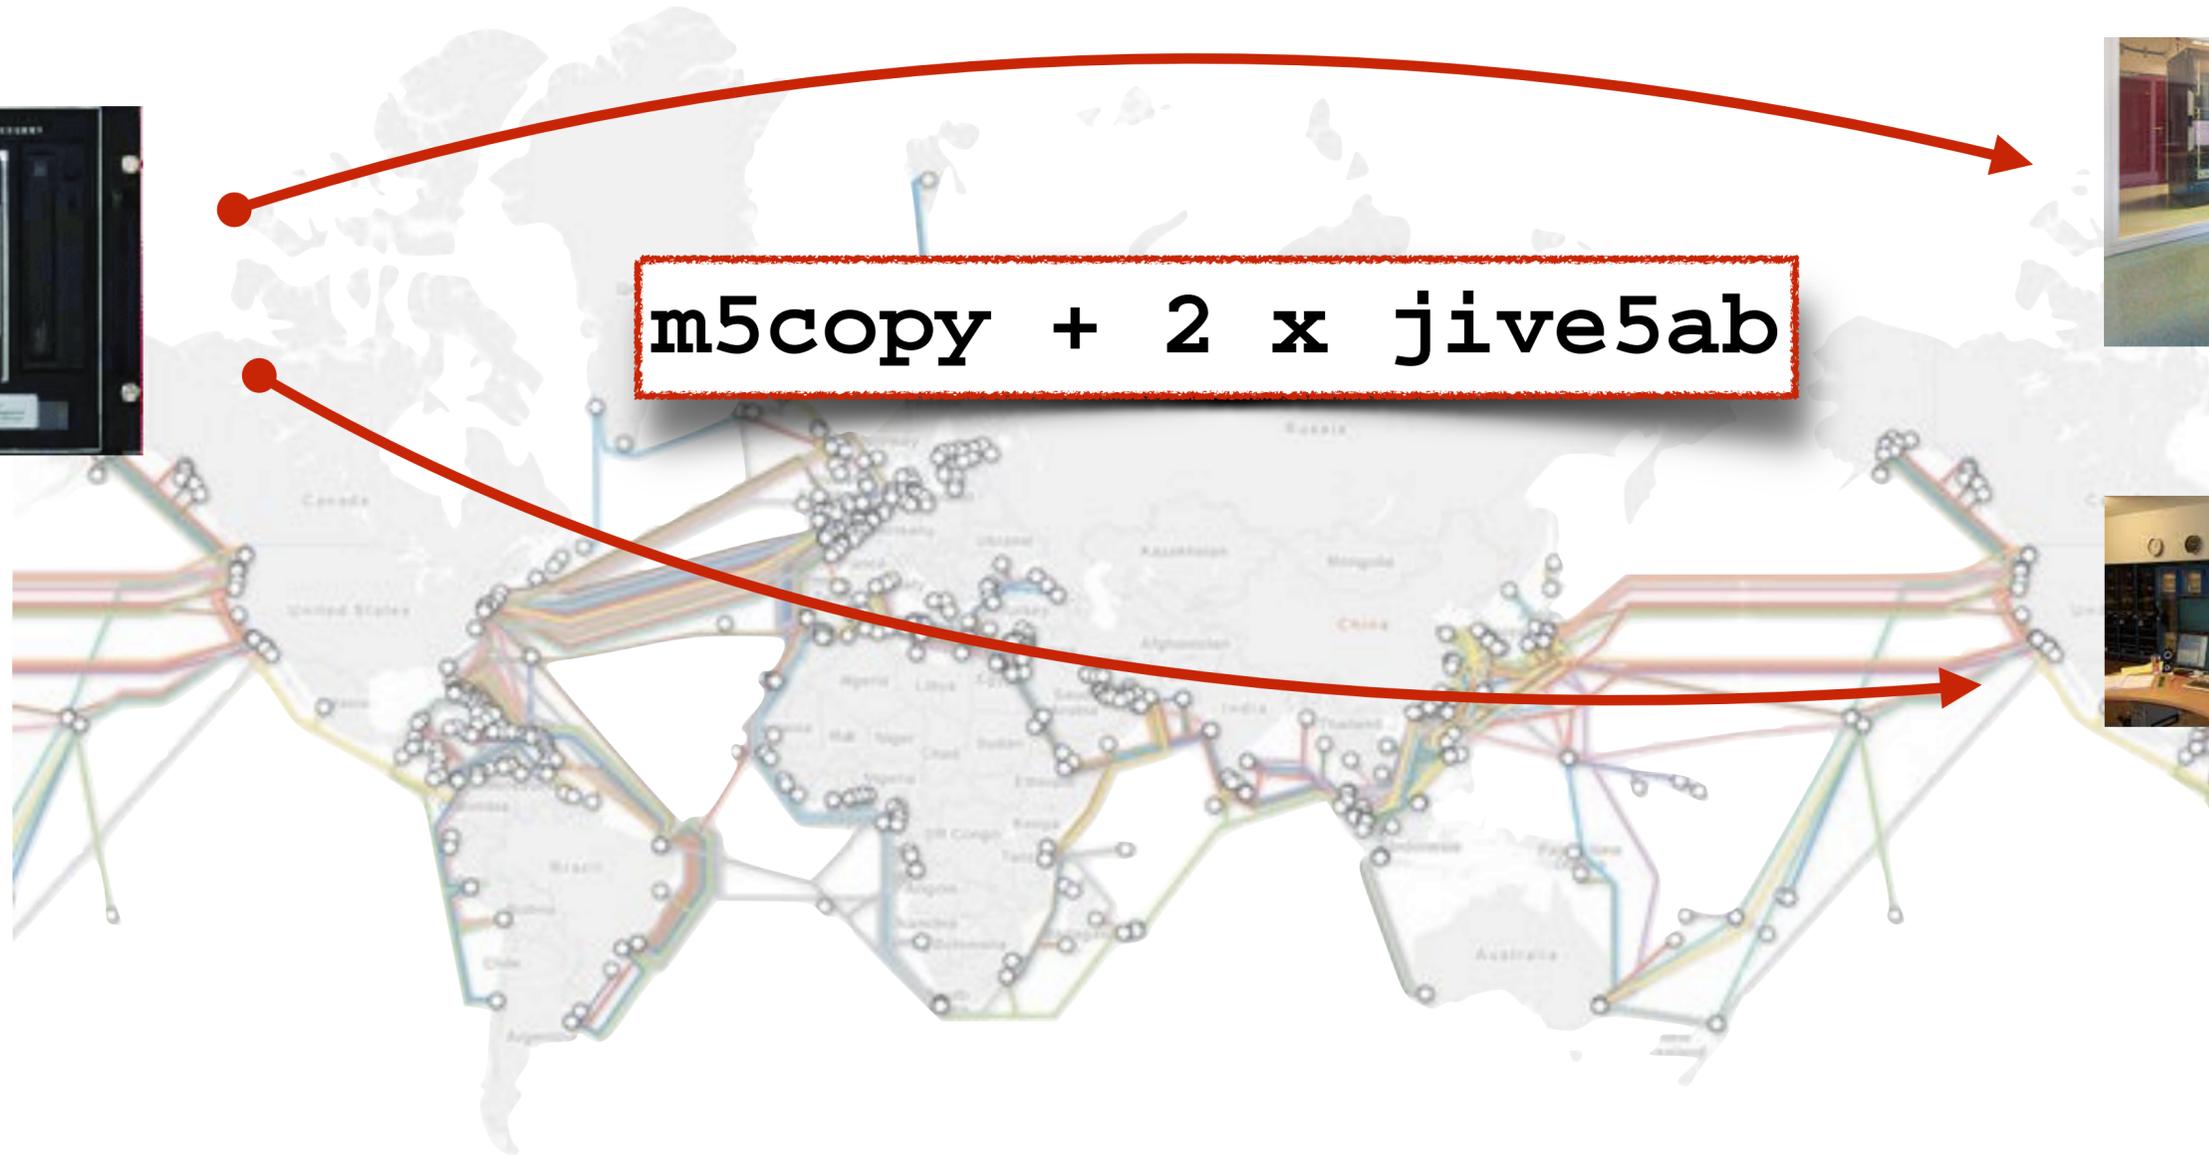
The screenshot shows the FlexBuff manager GUI with two main panels. The left panel shows a list of experiments with columns for Experiment, Station, Scan, and Size. The right panel shows a similar list with a context menu open over a selected item. The Selection panel at the top right has a 'Host' field set to 'mark6-1 (10.88.0.81)'. The bottom status bar shows 'Selection size: 7.328T (99 items)' and a checked 'Mark6 format' checkbox.

Experiment	Station	Scan	Size
GA043			9.51T
GA043A			21.63T
Hn			7.328T
			19G
			.4G
			.4G
			19G
			.4G
			.4G
			43.19G
			172.8G
			43.19G
			43.19G
			69.58G
			70.54G
			70.54G
			64.78G
			69.58G
			70.54G
			115.2G
			43.19G
			43.19G
			69.58G
			70.54G

<https://github.com/jive-vlbi/jive5ab-copy-manager>

# Get data to correlator

What options do exist? (jive5ab + m5copy e-shipping)



**m5copy + 2 x jive5ab**



# Get data to correlator

The advantage of flexbuff-to-flexbuff e-shipping

## Station

## Correlator

/mnt/disk0/**EXP\_STATION\_SCAN/**

EXP\_STATION\_SCAN.0000000

EXP\_STATION\_SCAN.0000007

/mnt/disk1/**EXP\_STATION\_SCAN/**

EXP\_STATION\_SCAN.0000002

EXP\_STATION\_SCAN.0000004

/mnt/disk2/**EXP\_STATION\_SCAN/**

EXP\_STATION\_SCAN.0000001

EXP\_STATION\_SCAN.0000006

/mnt/disk3/**EXP\_STATION\_SCAN/**

EXP\_STATION\_SCAN.0000003

EXP\_STATION\_SCAN.0000005



/mnt/disk0/**EXP\_STATION\_SCAN/**

/mnt/disk1/**EXP\_STATION\_SCAN/**

# Get data to correlator

The advantage of flexbuff-to-flexbuff e-shipping

## Station

## Correlator

/mnt/disk0/**EXP\_STATION\_SCAN/**

EXP\_STATION\_SCAN.0000000

EXP\_STATION\_SCAN.0000007

/mnt/disk1/**EXP\_STATION\_SCAN/**

EXP\_STATION\_SCAN.0000002

EXP\_STATION\_SCAN.0000004

/mnt/disk2/**EXP\_STATION\_SCAN/**

EXP\_STATION\_SCAN.0000001

EXP\_STATION\_SCAN.0000006

/mnt/disk3/**EXP\_STATION\_SCAN/**

EXP\_STATION\_SCAN.0000003

EXP\_STATION\_SCAN.0000005



/mnt/disk0/**EXP\_STATION\_SCAN/**

EXP\_STATION\_SCAN.0000000

EXP\_STATION\_SCAN.0000002

/mnt/disk1/**EXP\_STATION\_SCAN/**

EXP\_STATION\_SCAN.0000001

# Get data to correlator

The advantage of flexbuff-to-flexbuff e-shipping

## Station

## Correlator

/mnt/disk0/**EXP\_STATION\_SCAN/**

EXP\_STATION\_SCAN.0000000

EXP\_STATION\_SCAN.0000007

/mnt/disk1/**EXP\_STATION\_SCAN/**

EXP\_STATION\_SCAN.0000002

EXP\_STATION\_SCAN.0000004

/mnt/disk2/**EXP\_STATION\_SCAN/**

EXP\_STATION\_SCAN.0000001

EXP\_STATION\_SCAN.0000006

/mnt/disk3/**EXP\_STATION\_SCAN/**

EXP\_STATION\_SCAN.0000003

EXP\_STATION\_SCAN.0000005



/mnt/disk0/**EXP\_STATION\_SCAN/**

EXP\_STATION\_SCAN.0000000

EXP\_STATION\_SCAN.0000002

/mnt/disk1/**EXP\_STATION\_SCAN/**

EXP\_STATION\_SCAN.0000001

# Get data to correlator

The advantage of flexbuff-to-flexbuff e-shipping

## Station

## Correlator

/mnt/disk0/**EXP\_STATION\_SCAN/**

EXP\_STATION\_SCAN.0000000

EXP\_STATION\_SCAN.0000007

/mnt/disk1/**EXP\_STATION\_SCAN/**

EXP\_STATION\_SCAN.0000002

EXP\_STATION\_SCAN.0000004

/mnt/disk2/**EXP\_STATION\_SCAN/**

EXP\_STATION\_SCAN.0000001

EXP\_STATION\_SCAN.0000006

/mnt/disk3/**EXP\_STATION\_SCAN/**

EXP\_STATION\_SCAN.0000003

EXP\_STATION\_SCAN.0000005



/mnt/disk0/**EXP\_STATION\_SCAN/**

EXP\_STATION\_SCAN.0000000

EXP\_STATION\_SCAN.0000002

EXP\_STATION\_SCAN.0000007

EXP\_STATION\_SCAN.0000004

/mnt/disk1/**EXP\_STATION\_SCAN/**

EXP\_STATION\_SCAN.0000001

EXP\_STATION\_SCAN.0000003

EXP\_STATION\_SCAN.0000006

EXP\_STATION\_SCAN.0000005

Dealing with  
scattered data  
is a

**CENSORED**

# Dealing with scattered data

# Dealing with scattered data

How do I know what / how much is on those disks?!

# Dealing with scattered data

How do I know what / how much is on those disks?!

## VBS\_FS tools

*[https://code.jive.eu/verkout/vbs\\_fs](https://code.jive.eu/verkout/vbs_fs)*

`vbs_ls, vbs_rm, vbs_fs`

# Dealing with scattered data

How do I know what / how much is on those disks?!

VBS\_FS tools

[https://code.jive.eu/verkout/vbs\\_fs](https://code.jive.eu/verkout/vbs_fs)

**vbs\_ls**, vbs\_rm, vbs\_fs

# Dealing with scattered data

How do I know what / how much is on those disks?!

VBS\_FS tools

[https://code.jive.eu/verkout/vbs\\_fs](https://code.jive.eu/verkout/vbs_fs)

vbs\_ls, **vbs\_rm**, vbs\_fs

# Dealing with scattered data

How do I know what / how much is on those disks?!

VBS\_FS tools

[https://code.jive.eu/verkout/vbs\\_fs](https://code.jive.eu/verkout/vbs_fs)

vbs\_ls, vbs\_rm, **vbs\_fs**

# Dealing with scattered data

How do I know what / how much is on those disks?!

## VBS\_FS tools

[https://code.jive.eu/verkout/vbs\\_fs](https://code.jive.eu/verkout/vbs_fs)

**vbs\_ls**, **vbs\_rm**, vbs\_fs



Python scripts, list/remove vbs+mk6  
modelled after **ls(1)** and **rm(1)**  
many familiar options

# Dealing with scattered data

How do I know what / how much is on those disks?!

## VBS\_FS tools

[https://code.jive.eu/verkout/vbs\\_fs](https://code.jive.eu/verkout/vbs_fs)

vbs\_ls, vbs\_rm, **vbs\_fs**

FUSE virtual file system driver

[https://en.wikipedia.org/wiki/Filesystem\\_in\\_Userspace](https://en.wikipedia.org/wiki/Filesystem_in_Userspace)



# Dealing with scattered data

How do I know what / how much is on those disks?!

\$>

# Dealing with scattered data

How do I know what / how much is on those disks?!

```
$> vbs_ls [options][pattern..]
```

# Dealing with scattered data

How do I know what / how much is on those disks?!

```
$> vbs_ls -lrth haavee*
```

# Dealing with scattered data

How do I know what / how much is on those disks?!

```
$> vbs_ls -lrth haavee*
```

```
Found 4 recordings in 90 chunks,    57.68G
drw-r--r--  jops      flexbuf      12.75G Jun 22 10:27 haavee_vbs_no0001
drw-r--r--  jops      flexbuf       9.75G Jun 22 10:27 haavee_vbs_no0001a
-rw-r--r--  jops      flexbuf     15.64G Jun 22 10:27 haavee_m6_no0001
-rw-r--r--  jops      flexbuf     19.54G Jun 22 10:28 haavee_m6_no0002
```

# Dealing with scattered data

How do I know what / how much is on those disks?!

```
$> vbs_ls -lrth haavee*
```

```
Found 4 recordings in 90 chunks, 57.68G
```

```
drw-r--r-- jops flexbuf 12.75G Jun 22 10:27 haavee_vbs_no0001  
drw-r--r-- jops flexbuf 9.75G Jun 22 10:27 haavee_vbs_no0001a  
-rw-r--r-- jops flexbuf 15.64G Jun 22 10:27 haavee_m6_no0001  
-rw-r--r-- jops flexbuf 19.54G Jun 22 10:28 haavee_m6_no0002
```

# Dealing with scattered data

How do I know what / how much is on those disks?!

```
$> vbs_ls -lrth haavee*
```

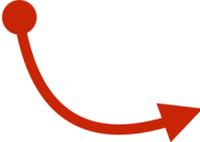
```
Found 4 recordings in 90 chunks, 57.68G
```

```
drw-r--r-- jops flexbuf 12.75G Jun 22 10:27 haavee_vbs_no0001  
drw-r--r-- jops flexbuf 9.75G Jun 22 10:27 haavee_vbs_no0001a  
-rw-r--r-- jops flexbuf 15.64G Jun 22 10:27 haavee_m6_no0001  
-rw-r--r-- jops flexbuf 19.54G Jun 22 10:28 haavee_m6_no0002
```

# Dealing with scattered data

How do I know what / how much is on those disks?!

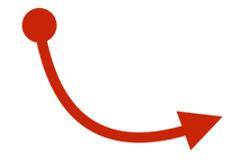
```
$> vbs_ls ... -T ... <pattern> [<pattern> ...]
```

 “accumulate by <pattern>s”

# Dealing with scattered data

How do I know what / how much is on those disks?!

```
$> vbs_ls ... -lTh em117e* em117f* eg088_ys*
```

 “accumulate by <pattern>s”

```
Found 3 recordings in 13426 chunks,      3.11T
drw-r--r-- jops      flexbuf      2.36T Aug 05 13:24 eg088_ys*
drw-r--r-- jops      flexbuf      394.59G Jun 06 20:34 em117e*
drw-r--r-- jops      flexbuf      372.59G Jun 06 18:09 em117f*
```

# Dealing with scattered data

Easy access to recordings using a virtual file system

```
verkouter@flexbuf1:~$ find . -type f -name haavee\* -exec ls -lh {} \;
```

# Dealing with scattered data

Easy access to recordings using a virtual file system

```
verkouter@flexbuf1:~$ find . -type f -name haavee\* -exec ls -lh {} \;
```

```
-rw-r--r-- 1 jops flexbuf 3.3G Jun 22 12:28 ./haavee_m6_no0002
-rw-r--r-- 1 jops flexbuf 2.7G Jun 22 12:27 ./haavee_m6_no0001

-rw-r--r-- 1 jops flexbuf 256M Jun 22 12:26 ./haavee_vbs_no0001/haavee_vbs_no0001.00000007
-rw-r--r-- 1 jops flexbuf 256M Jun 22 12:26 ./haavee_vbs_no0001/haavee_vbs_no0001.00000019
-rw-r--r-- 1 jops flexbuf 256M Jun 22 12:26 ./haavee_vbs_no0001/haavee_vbs_no0001.00000013
-rw-r--r-- 1 jops flexbuf 256M Jun 22 12:26 ./haavee_vbs_no0001/haavee_vbs_no0001.00000025
-rw-r--r-- 1 jops flexbuf 256M Jun 22 12:27 ./haavee_vbs_no0001/haavee_vbs_no0001.00000031
-rw-r--r-- 1 jops flexbuf 256M Jun 22 12:26 ./haavee_vbs_no0001/haavee_vbs_no0001.00000001
-rw-r--r-- 1 jops flexbuf 256M Jun 22 12:27 ./haavee_vbs_no0001/haavee_vbs_no0001.00000049
-rw-r--r-- 1 jops flexbuf 256M Jun 22 12:27 ./haavee_vbs_no0001/haavee_vbs_no0001.00000043
-rw-r--r-- 1 jops flexbuf 256M Jun 22 12:27 ./haavee_vbs_no0001/haavee_vbs_no0001.00000037

-rw-r--r-- 1 jops flexbuf 256M Jun 22 12:27 ./haavee_vbs_no0001a/haavee_vbs_no0001a.00000037
-rw-r--r-- 1 jops flexbuf 256M Jun 22 12:27 ./haavee_vbs_no0001a/haavee_vbs_no0001a.00000001
-rw-r--r-- 1 jops flexbuf 256M Jun 22 12:27 ./haavee_vbs_no0001a/haavee_vbs_no0001a.00000031
-rw-r--r-- 1 jops flexbuf 256M Jun 22 12:27 ./haavee_vbs_no0001a/haavee_vbs_no0001a.00000025
-rw-r--r-- 1 jops flexbuf 256M Jun 22 12:27 ./haavee_vbs_no0001a/haavee_vbs_no0001a.00000019
```

# Dealing with scattered data

Easy access to recordings using a virtual file system

```
$> vbs_fs [options] /path/to/mountpoint
```

# Dealing with scattered data

Easy access to recordings using a virtual file system

/mnt/disk/1/0/data/EXP\_STATION\_SCAN.vdif

HDR	0	DATA	7	DATA
-----	---	------	---	------

/mnt/disk/2/0/data/EXP\_STATION\_SCAN.vdif

HDR	2	DATA	4	DATA
-----	---	------	---	------

/mnt/disk/3/0/data/EXP\_STATION\_SCAN.vdif

HDR	1	DATA	5	DATA
-----	---	------	---	------

/mnt/disk/4/7/data/EXP\_STATION\_SCAN.vdif

HDR	3	DATA	6	DATA
-----	---	------	---	------

/mnt/disk0/EXP\_STATION\_SCAN/

EXP_STATION_SCAN.0000000	EXP_STATION_SCAN.0000007
--------------------------	--------------------------

/mnt/disk1/EXP\_STATION\_SCAN/

EXP_STATION_SCAN.0000002	EXP_STATION_SCAN.0000004
--------------------------	--------------------------

/mnt/disk2/EXP\_STATION\_SCAN/

EXP_STATION_SCAN.0000001	EXP_STATION_SCAN.0000006
--------------------------	--------------------------

/mnt/disk3/EXP\_STATION\_SCAN/

EXP_STATION_SCAN.0000003	EXP_STATION_SCAN.0000005
--------------------------	--------------------------

/path/to/mountpoint/EXP\_STATION\_SCAN.vdif

DATA[0]	DATA[1]	DATA[2]	DATA[3]	DATA[4]	DATA[5]	DATA[6]	...
---------	---------	---------	---------	---------	---------	---------	-----

/path/to/mountpoint/EXP\_STATION\_SCAN

EXP_STATION_SCAN.0000000	EXP_STATION_SCAN.0000001	EXP_STATION_SCAN.0000002	EXP_STATION_SCAN.0000003	...
--------------------------	--------------------------	--------------------------	--------------------------	-----

# Dealing with scattered data

Easy access to recordings using a virtual file system

```
$> vbs_fs [options] /path/to/mountpoint
```

A multi-threaded FUSE virtual file system in C++

- Reconstructs scattered recordings as single files, transparently for:
  - `cplane/dplane` MIT Haystack Mark6 format
  - `jive5ab` FlexBuff/vbs format
- User, group, permissions, modification time reflect actual status of on-disk files
- I/O scheduling done in `vbs_fs`
  - configurable read-ahead
  - serving multiple files and/or multiple users guaranteed optimal *if reading from the same mountpoint*
  - `vbs_fs` can be run multiple times but may degrade I/O performance depending on usage pattern

# Dealing with scattered data

Easy access to recordings using a virtual file system

```
$> mkdir /path/to/mountpoint
```

```
$> vbs_fs [options] /path/to/mountpoint
```

```
$> ls -alh /path/to/mountpoint
```

```
-rw-r--r-- 0 jops flexbuf 16G Jun 22 12:27 /tmp/foo/haavee_m6_no0001  
-rw-r--r-- 0 jops flexbuf 20G Jun 22 12:28 /tmp/foo/haavee_m6_no0002  
-rw-r--r-- 0 jops flexbuf 13G Jun 22 12:27 /tmp/foo/haavee_vbs_no0001  
-rw-r--r-- 0 jops flexbuf 9.8G Jun 22 12:27 /tmp/foo/haavee_vbs_no0001a
```

# Dealing with scattered data

Easy access to recordings using a virtual file system

```
$> mkdir /path/to/mountpoint
```

```
$> vdiif -o /path/to/mountpoint
```

```
$> ls -alh /path/to/mountpoint
```

```
-rw-r--r-- 0 jops flexbuf 16G Jun 22 12:27 /tmp/foo/haavee_m6_no0001
-rw-r--r-- 0 jops flexbuf 20G Jun 22 12:28 /tmp/foo/haavee_m6_no0002
-rw-r--r-- 0 jops flexbuf 13G Jun 22 12:27 /tmp/foo/haavee_vbs_no0001
-rw-r--r-- 0 jops flexbuf 9.8G Jun 22 12:27 /tmp/foo/haavee_vbs_no0001a
```

100% VDIIF useful payload only!  
vbs\_fs strips all cplane headers



# Dealing with scattered data

Easy access to recordings using a virtual file system

```
$> mkdir /path/to/mountpoint
```

```
$> vbs_fs [options] /path/to/mountpoint
```

## [options]

- 6 Scan Mark6 mountpoints for *shrapnel* (default **FlexBuff mountpoints**)
- R <PATH> Add <PATH> to list of directories to scan for *shrapnel*
- I <PATTERN> Only index/scan recordings matching <PATTERN> (default **anything recognizable as recording**)
- n <NUM> Enable read-ahead of <NUM> blocks to speed up data access (default **no read-ahead**)

# Get data to correlator

What options do exist? e-transfer daemon/client

# Get data to correlator

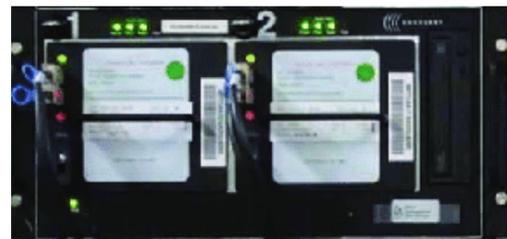
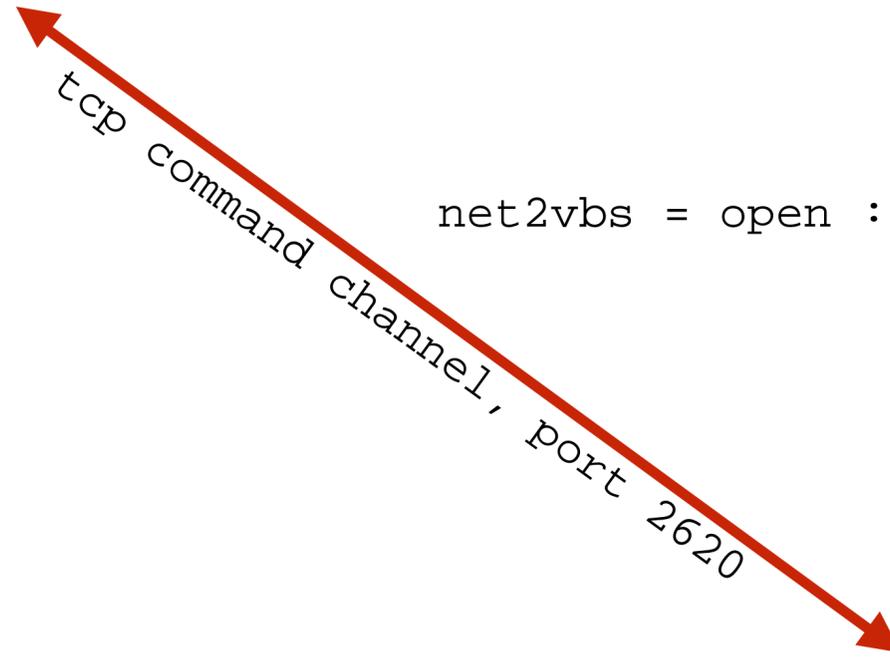
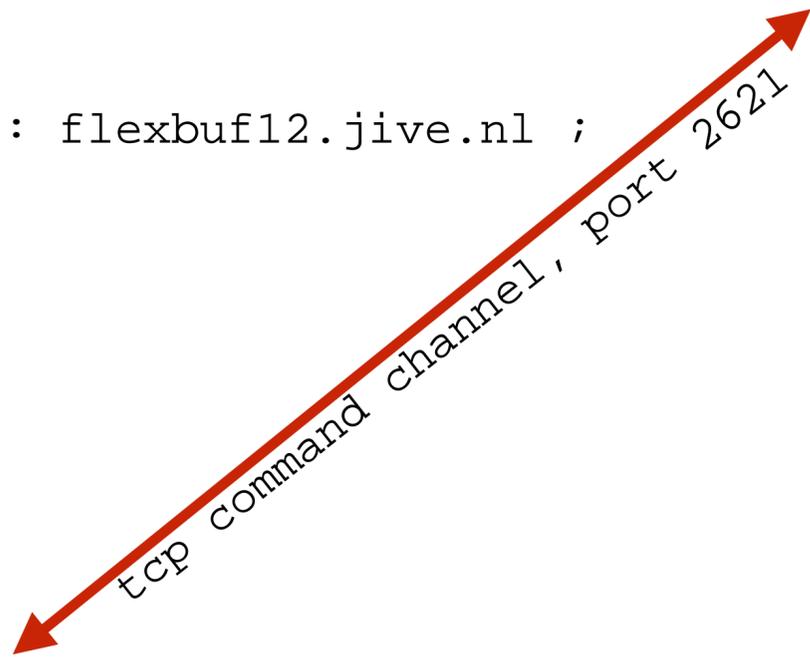
What options do exist? e-transfer daemon/client



```
verkouter@Mac> m5copy -udt -mtu 9000 -p 2631 mk6://130.141.242.16:2621/... vbs://flexbuf12.jive.nl/...
```

```
disk2net = connect : flexbuf12.jive.nl ;
```

```
net2vbs = open : ... ;
```



```
$> jive5ab -p 2621 -m 3 -6
```

```
$> jive5ab -m 3
```

# Get data to correlator

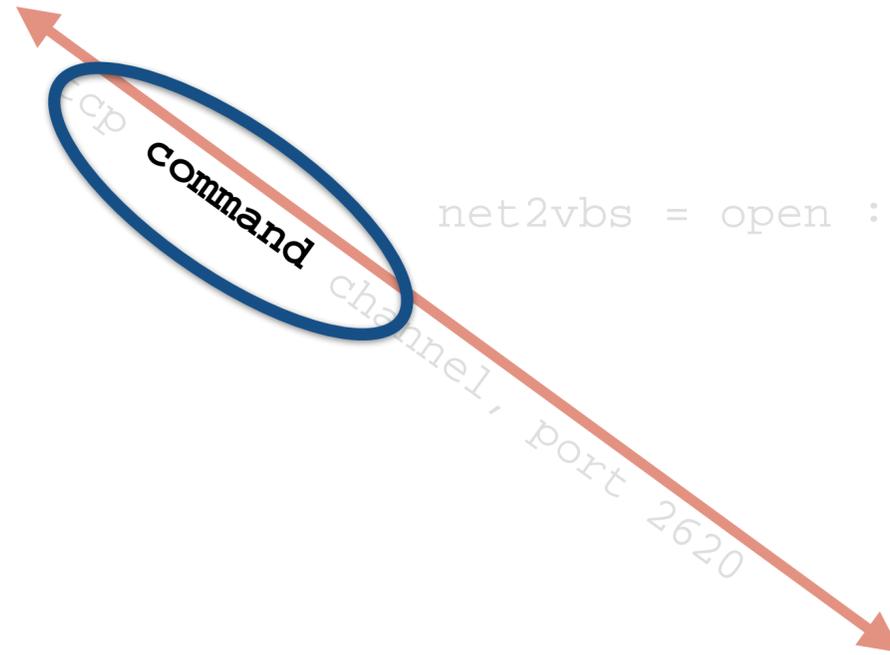
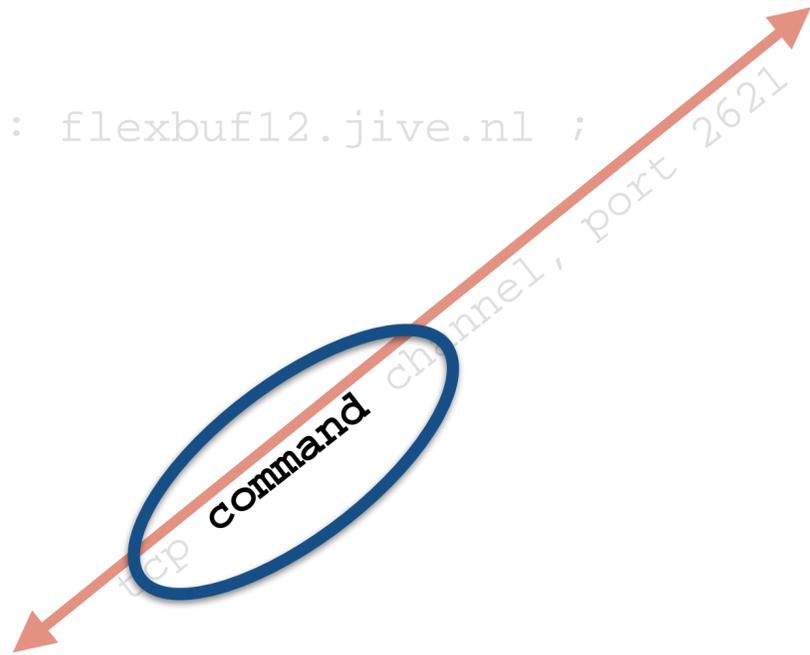
What options do exist? e-transfer daemon/client



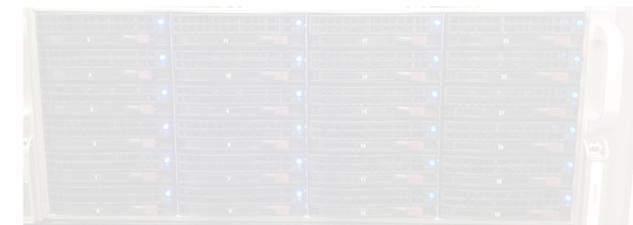
```
verkouter@Mac> m5copy -udt -mtu 9000 -p 2631 mk6://130.141.242.16:2621/... vbs://flexbuf12.jive.nl/...
```

```
disk2net = connect : flexbuf12.jive.nl ;
```

```
net2vbs = open : ... ;
```



```
udt data channel port 2631, mtu 9000
```



```
$> jive5ab -p 2621 -m 3 -6
```

```
$> jive5ab -m 3
```

# Get data to correlator

What options do exist? e-transfer daemon/client

## The e-transfer tools

<https://github.com/jive-vlbi/etransfer>

etd, etc



"e-transfer daemon"  
"e-transfer client"

# Get data to correlator

What options do exist? e-transfer daemon/client

## The e-transfer tools support

- tcp, udt
  - IPv4 + IPv6
- proper daemon
  - multiple clients over one port simultaneously
  - server communicates data channel to client
- also remote-to-remote transfers
- file to file only

# Get data to correlator

What options do exist? e-transfer daemon/client

The e-transfer tools support

- tcp, udt

- IPv4 + IPv6

- rsh

in m5copy the user sets the data channel, not the server admin

- port simultaneously data channel to client

- i-co-remote transfers

- file to file only

# Get data to correlator

What options do exist? e-transfer daemon/client

```
$> ../etd [options] --command tcp:// --data udt://
```

# Get data to correlator

What options do exist? e-transfer daemon/client

```
$> ../etd [options] --command tcp:// --data udt://
```

**--command <protocol>: // [ <host> ] [ : <port> ]**

<protocol>	tcp, tcp6, udt, udt6 (default <b>no default</b> )
<host>	IPv4 or IPv6 address or hostname (default <b>all interfaces</b> i.e. <b>IPv4 0.0.0.0</b> )
<port>	Port number to listen on for incoming command connections (default <b>4004</b> )

# Get data to correlator

What options do exist? e-transfer daemon/client

```
$> ../etd [options] --command tcp:// --data udt://
```

```
--data <protocol>:// [<host>] [:<port>]
```

<protocol>	tcp, tcp6, udt, udt6 (default <b>no default</b> )
<host>	IPv4 or IPv6 address or hostname (default <b>all interfaces</b> i.e. <b>IPv4 0.0.0.0</b> )
<port>	Port number to listen on for incoming data connections (default <b>8008</b> )

# Get data to correlator

What options do exist? e-transfer daemon/client

```
$> ../etd [options] --command tcp://  
      --data tcp://10.88.0.33:8009  
      --data udt://192.42.120.39:8008
```

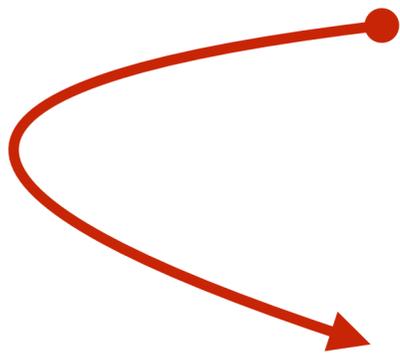
# Get data to correlator

What options do exist? e-transfer daemon/client

```
$> ../etd [options] --command tcp://  
    --data tcp://10.88.0.33:8009  
    --data udt://192.42.120.39:8008
```

e-transfer client will:

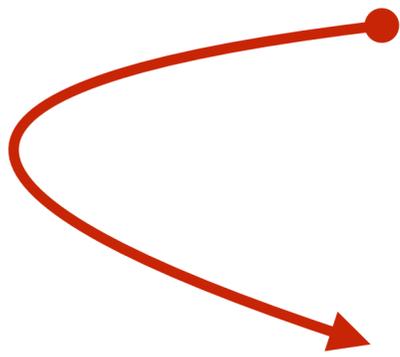
- try to connect to data channels in this order



# Get data to correlator

What options do exist? e-transfer daemon/client

```
$> ../etd [options] --command tcp://  
      --data tcp://10.88.0.33:8009  
      --data udt://192.42.120.39:8008
```



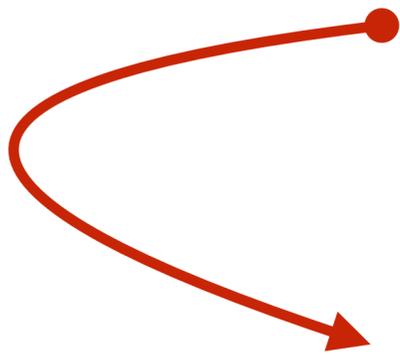
e-transfer client will:

- try to connect to data channels in this order
- internal client uses fast tcp

# Get data to correlator

What options do exist? e-transfer daemon/client

```
$> ../etd [options] --command tcp://  
--data tcp://10.88.0.33:8009  
--data udt://192.42.120.39:8008
```



e-transfer client will:

- try to connect to data channels in this order
- internal client uses fast tcp
- external client uses fast udt

# Get data to correlator

What options do exist? e-transfer daemon/client

```
$> ../etc [options] SRC DST
```

# Get data to correlator

What options do exist? e-transfer daemon/client

```
$> .../etc [options] SRC DST
```

```
/path/to/file*.*
```

# Get data to correlator

What options do exist? e-transfer daemon/client

```
$> .../etc [options] SRC DST
```

```
/path/to/file*.*
```

```
flexbuf4.jive.nl:/path/to/file*.*
```

```
flexbuf4.jive.nl#4005:/path/to/file*.*
```

```
tcp6:flexbuf6.jive.nl:/path/to/file*.*
```

# Get data to correlator

What options do exist? e-transfer daemon/client

```
$> .../etc [options] SRC DST
```

non-standard control port

```
/path/to/file*.*
```

```
flexbuf4.jive.nl:/path/to/file*.*
```

```
flexbuf4.jive.nl #4005:/path/to/file*.*
```

```
tcp6: flexbuf6.jive.nl:/path/to/file*.*
```

non-standard protocol

# Get data to correlator

What options do exist? e-transfer daemon/client

```
$> ../etc [options] SRC DST
```

```
/path/to/{dir/|file}
```

```
130.141.242.16:/path/to/{dir/|file}
```

```
udt:fb7.jive.nl#42267:/path/to/{dir/|file}
```

# Get data to correlator

What options do exist? e-transfer daemon/client

```
$> vbs_fs [options] /mnt/data
```

# Get data to correlator

What options do exist? e-transfer daemon/client

```
$> vbs_fs [options] /mnt/data
```

```
$> ../etc [options] /mnt/data/* fb7.jive.nl:/data/
```

# Get data to correlator

What options do exist? e-transfer daemon/client

```
$> ../etc [options] SRC DST
```

## [options]

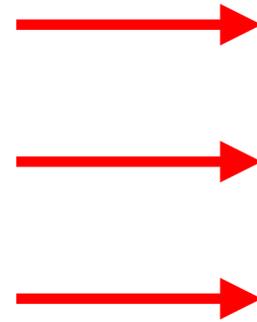
-h, --help	short (long) explanation of the command line
-m <number>	Message level - higher number = more output (default <b>0</b> )
-v	Enable verbose output on each file transferred
--resume, --overwrite, --skipexisting	How to handle file(s) that already exist on the other side (default <b>New</b> i.e. <i>file may not exist!</i> )

... and there's more ...

Thanks  
for  
attention!

(Automatically) collect frames  
in  
multiple recordings?

digital backend(s)



VLBI data recorder(s)

end ( s )

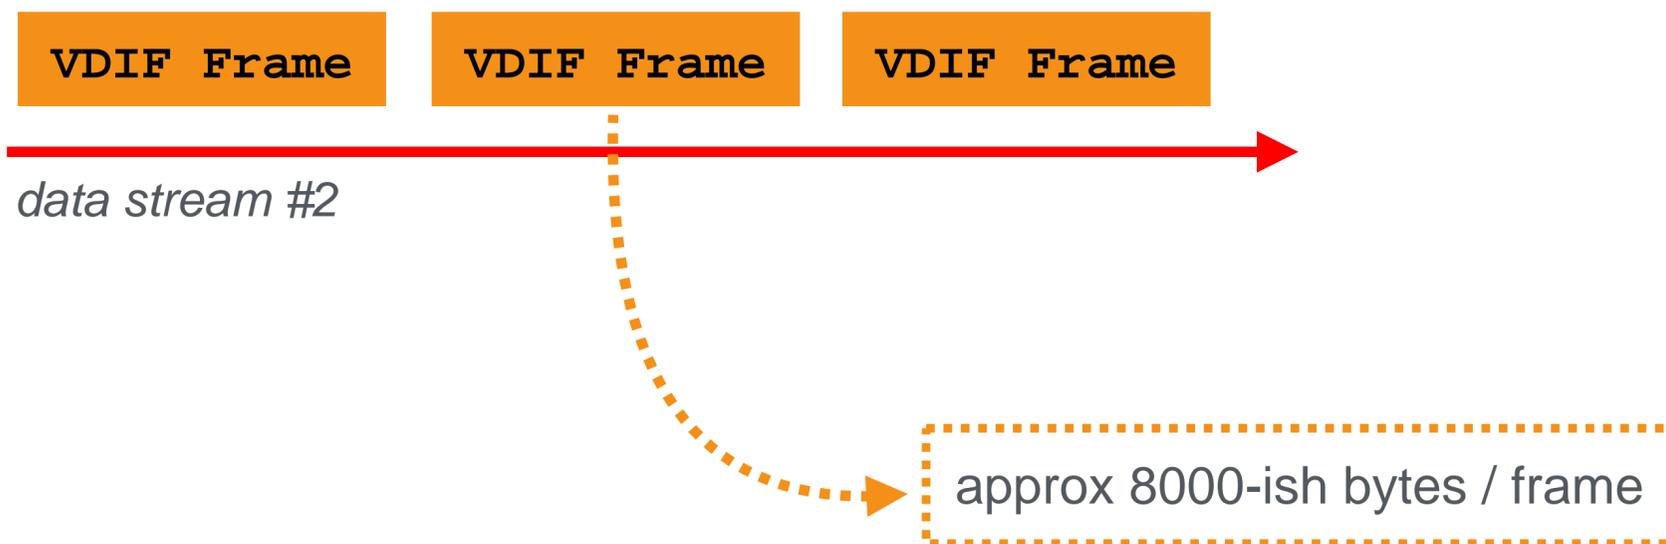
  
*data stream #0*

  
*data stream #1*

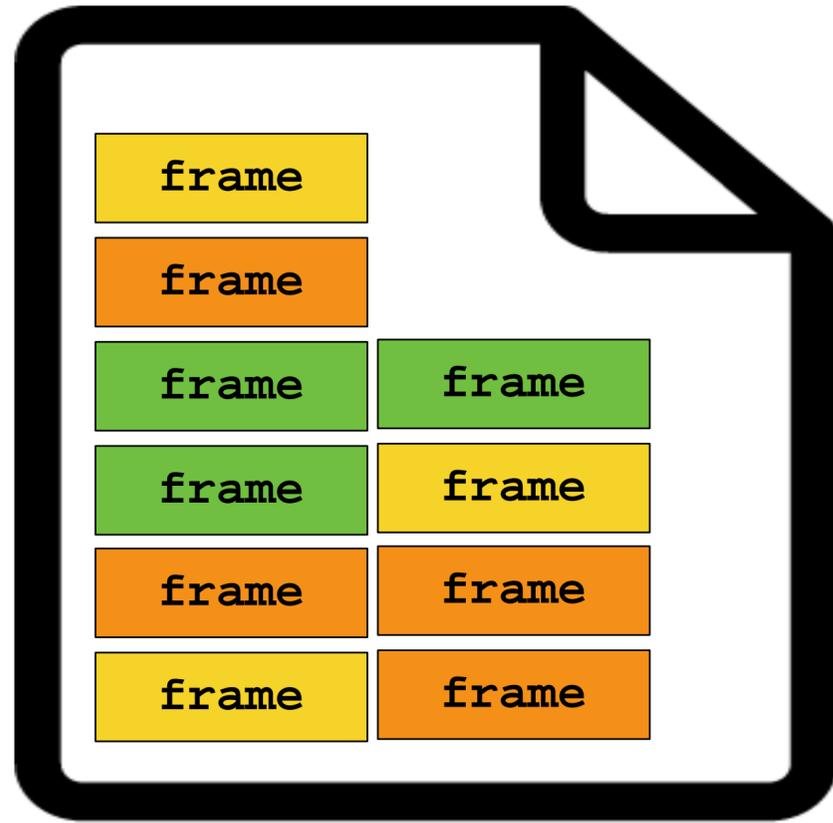
  
*data stream #2*

VLBI

end ( s )



VLBI



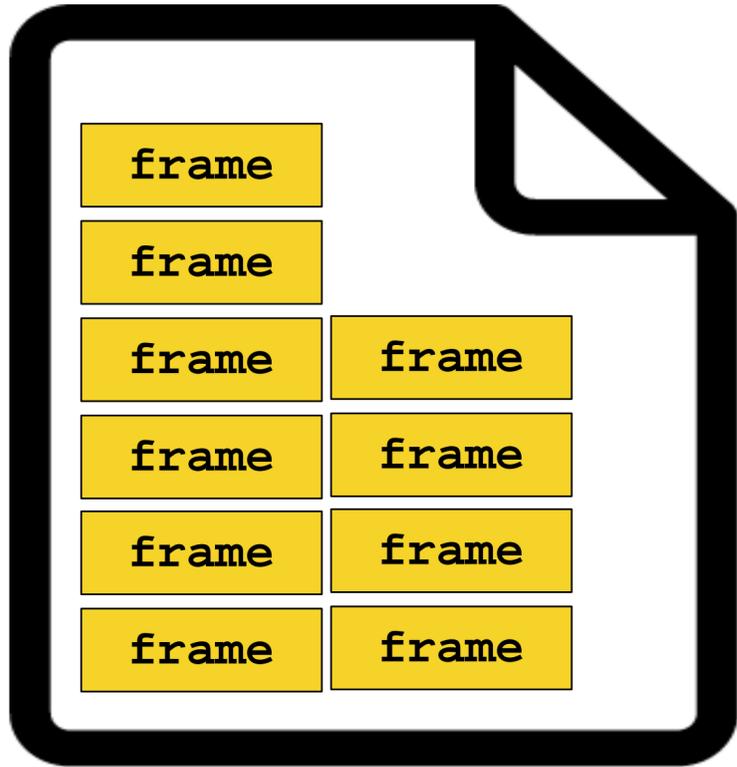
`EXP_STA_SCAN(vbs/mk6)`

```
net_port = [ip1@]port1[=suffix1]:  
           [ip2@]port2[=suffix2];
```

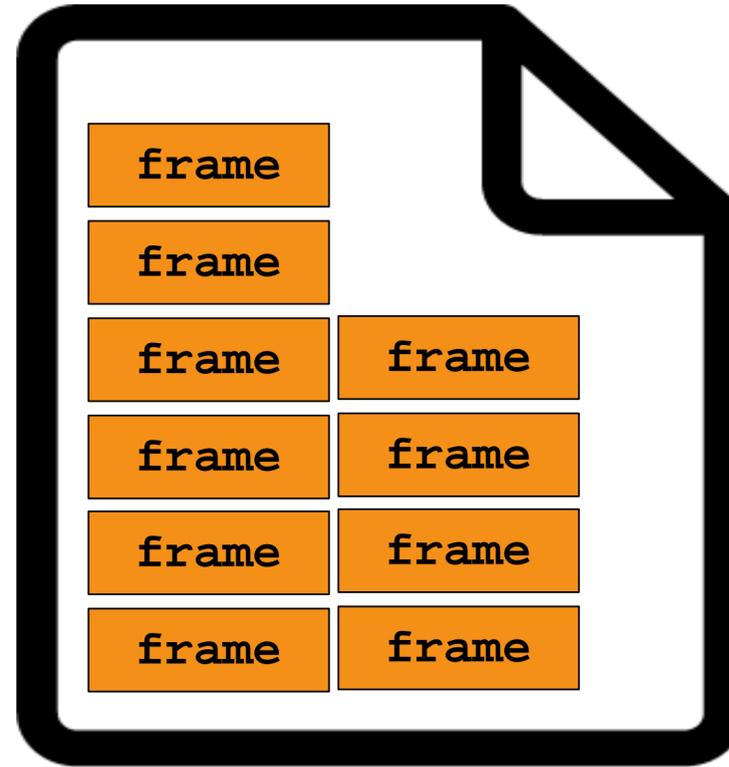
```
net_port = [ip1@]port1[=suffix1]:  
           [ip2@]port2[=suffix2];  
  
record = on : <scan> ;
```



records **all** frames arriving on  
[ip1:]port1 into <scan>\_ds<suffix1>  
[ip2:]port2 into <scan>\_ds<suffix2>  
etc.



EXP\_STA\_SCAN\_ds<suffix1>



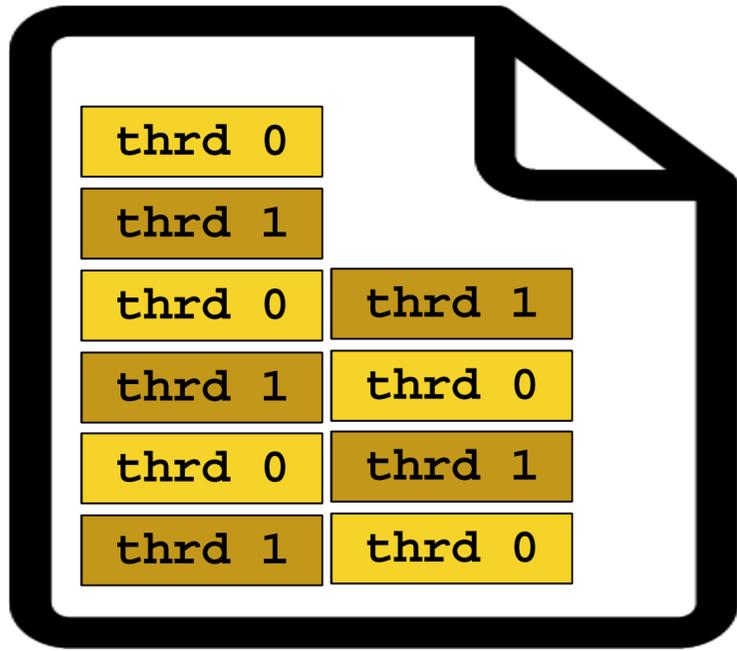
EXP\_STA\_SCAN\_ds<suffix2>

...

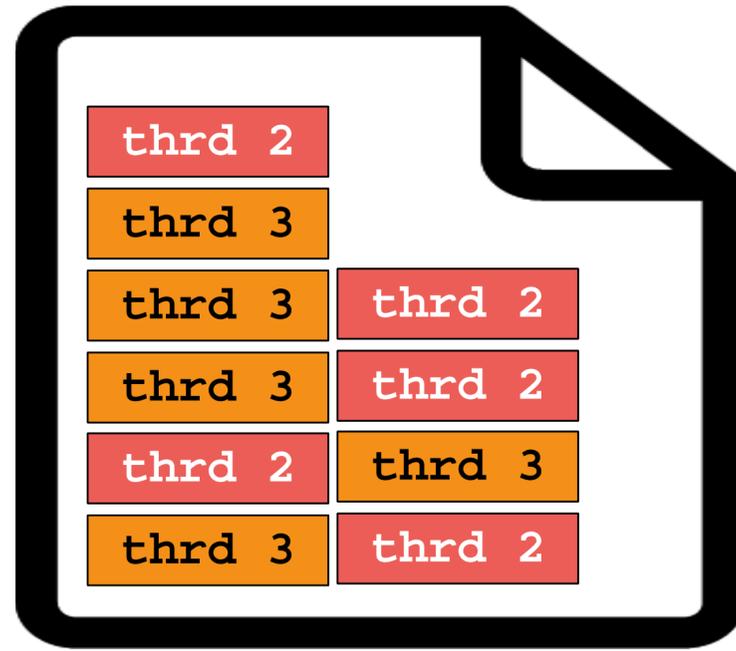
end ( s )



VLBI



EXP\_STA\_SCAN\_ds<suffix1>



EXP\_STA\_SCAN\_ds<suffix2>

...

# The **datastream=** command

See this commit log for brief usage/help/explanation:

<https://github.com/jive-vlbi/jive5ab/commit/74c3f1b35bd258f344def8a50cae805381420549>

The standard 32-byte VDIF Data Frame Header is shown in Figure 3.

	Byte 3		Byte 2	Byte 1	Byte 0
	Bit 31 (MSB)				Bit 0 (LSB)
Word 0	$I_1$	$L_1$	Seconds from reference epoch <sub>30</sub>		
Word 1	Un-assigned <sub>2</sub>		Ref Epoch <sub>6</sub>	Data Frame # within second <sub>24</sub>	
Word 2	$V_3$		$\log_2(\#chns)_5$	Data Frame length (units of 8 bytes) <sub>24</sub>	
Word 3	$C_1$	bits/sample-1 <sub>5</sub>	Thread ID <sub>10</sub>	Station ID <sub>16</sub>	
Word 4	EDV <sub>8</sub>			Extended User Data <sub>24</sub>	
Word 5	Extended User Data <sub>32</sub>				
Word 6	Extended User Data <sub>32</sub>				
Word 7	Extended User Data <sub>32</sub>				

Figure 3: VDIF Data Frame Header format; subscripts are field lengths in bits; byte #s indicate relative byte address within 32-bit word (little endian format)

From VLBI Data Interchange Format (2009) - <https://vlbi.org/vlbi-standards/vdif/>

```
VDIF {threadId=0, stationId="Wf"}  
sent by 192.168.178.1:5000
```

```
VDIF {threadId=1, stationId="Wf"}  
sent by 192.168.178.2:5000
```

```
$> jive5ab -m 3  
s = socket(AF_INET, PF_DGRAM);  
bind(s, 192.168.178.800, 2630);
```

```
while( true ) {  
    ipv4_addr fromAddr;  
    read_from(socket, buf, &fromAddr);  
    /* ... */  
    vdifh = (struct vdif_header*)buf;  
    if( vdifh.threadId == 0 ) {  
        /* ... */  
    }  
}
```

**data stream #0**

```
VDIF {threadId=0, stationId="Wf"}  
sent by 192.168.178.1:5000
```

```
VDIF {threadId=1, stationId="Wf"}  
sent by 192.168.178.2:5000
```

```
$> jive5ab -m 3  
s = socket(AF_INET, PF_DGRAM);  
bind(s, 192.168.178.800, 2630);
```

```
while( true ) {  
    ipv4_addr fromAddr;  
    read_from(socket, buf, &fromAddr);  
    /* ... */  
    vdifh = (struct vdif_header*)buf;  
    if( vdifh.threadId == 0 ) {  
        /* ... */  
    }  
}
```

*data stream #0*

```
VDIF {threadId=0, stationId="Wf"}  
sent by 192.168.178.1:5000
```

```
VDIF {threadId=1, stationId="Wf"}  
sent by 192.168.178.2:5000
```

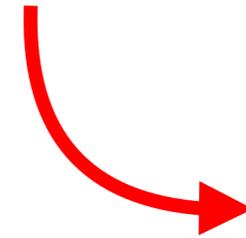
```
$> jive5ab -m 3  
s = socket(AF_INET, PF_DGRAM);  
bind(s, 192.168.178.800, 2630);
```

```
while( true ) {  
    ipv4_addr fromAddr;  
    read_from(socket, buf, &fromAddr);  
    /* ... */  
    vdifh = (struct vdif_header*)buf;  
    if( vdifh.threadId == 0 ) {  
        /* ... */  
    }  
}
```

*data stream #0*

```
datastream = add : <suffix> : <match specification> [: <more matches> ];
```

```
datastream = add : <suffix> : <match specification> [: <more matches> ];
```



Match on any/all of:

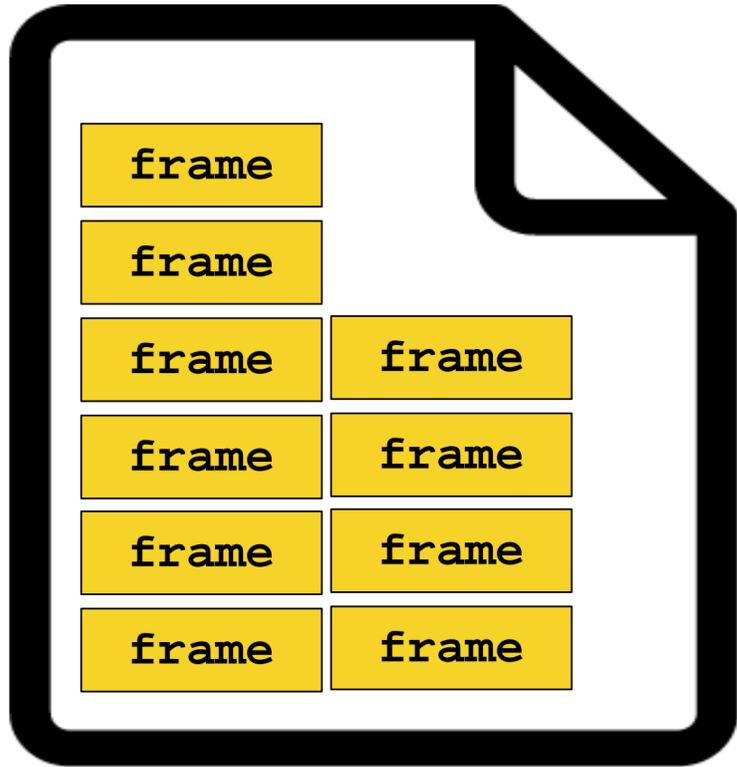
- source IPv4
- source port
- VDIF threadId
- VDIF stationId

See this commit log for brief usage/help/explanation:

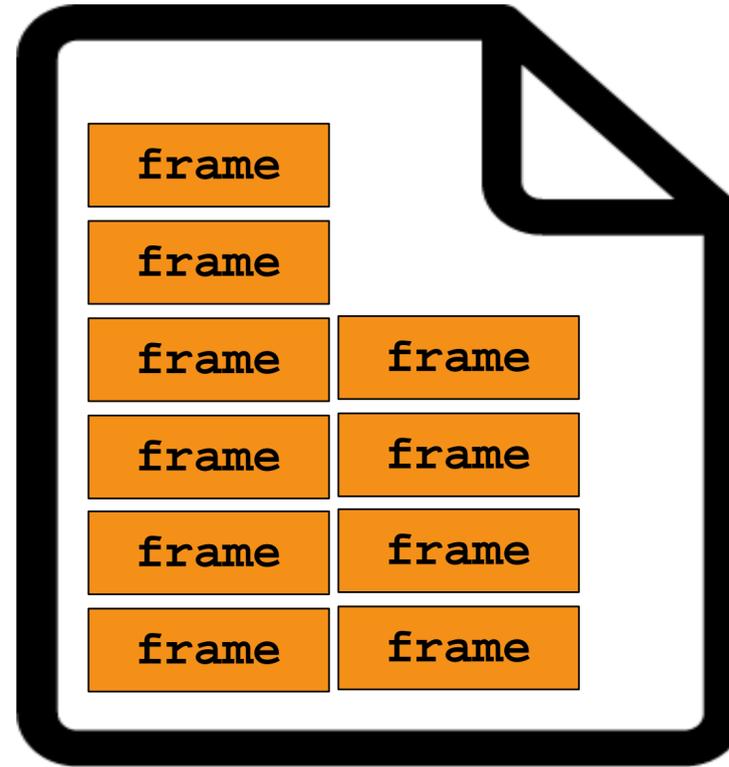
<https://github.com/jive-vlbi/jive5ab/commit/74c3f1b35bd258f344def8a50cae805381420549>

```
datastream = add : <suffix> : <match specification> [: <more matches> ];
```

See this commit log for brief usage/help/explanation:  
<https://github.com/jive-vlbi/jive5ab/commit/74c3f1b35bd258f344def8a50cae805381420549>



EXP\_STA\_SCAN\_ds<suffix1>



EXP\_STA\_SCAN\_ds<suffix2>

...

Thanks  
for  
attention!



```
00:53:20;mk5=record? nthread
```

```
00:53:20/mk5/!record? : 4 : 32 ;
```

```
00:53:31; mk5=net_protocol?
```

```
00:53:31/mk5/net protocol? 0 : udpsnor: 768000000 : 128000000 : 16 ;
```

```
00:53:39; mk5=net_port?
```

```
00:53:39/mk5/!net_port? 0 : 2630 : 2631 : 2632 : 2634 ;
```

```
00:53:20;mk5=record? nthread
```

```
00:53:20/mk5/!record? : 4 : 32 ;
```

```
00:53:31; mk5=net_protocol?
```

```
00:53:31/mk5/net protocol? 0 : udpsnor: 768000000 : 128000000 : 16 ;
```

```
00:53:39; mk5=net_port?
```

```
00:53:39/mk5/!net_port? 0 : 2630 : 2631 : 2632 : 2634 ;
```

00:53:20;mk5=record? nthread

00:53:20/mk5/!record? : 4 : 32 ;

00:53:31; mk5=net\_protocol?

00:53:31/mk5/net protocol? 0 : udpsnor: 768000000 : 128000000 : 16 ;

00:53:39; **mk5=net\_port?**

00:53:39/mk5/!**net\_port?** 0 : **2630 : 2631 : 2632 : 2634 ;**

```
2024.009.15:46:35.17/mk5=version?2024.009.15:46:35.17/mk5/!version? 0 : jive5ab :  
3.1.0-dev : 64bit : Debug :  
koll : 08-Jan-2024 :  
12h26m00s : nossapi :  
/usr/local/src/etransfer ;
```

```
2024.009.15:48:37.54;mk5=scan_check?
```

```
2024.009.15:48:37.59/mk5/scan_check? 0 :?:  
xxx_xx_no0001 : VDIF : 128 :  
2024y009d15h43m13.3117s :  
5.9812s : 256Mbps :  
1060224 : 8032 ;
```

```
2024.009.15:48:42.56;mk5=evlbi?
```

```
2024.009.15:48:42.56/mk5/evlbi? 0 : total: 3062244:  
loss : 0 (0.00%) :  
out-of-order: 0 ( 0.00%) :  
extent: 0 seqnr/pkt ;
```

```
2024.009.15:46:35.17/mk5=version?2024.009.15:46:35.17/mk5/!version? 0 : jive5ab :  
3.1.0-dev : 64bit : Debug :  
koll : 08-Jan-2024 :  
12h26m00s : nossapi :  
/usr/local/src/etransfer ;
```

```
2024.009.15:48:37.54;mk5=scan_check?
```

```
2024.009.15:48:37.59/mk5/scan_check? 0 :?:  
xxx_xx_no0001 : VDIF : 128 :  
2024y009d15h43m13.3117s :  
5.9812s : 256Mbps :  
1060224 : 8032 ;
```

```
2024.009.15:48:42.56;mk5=evlbi?
```

```
2024.009.15:48:42.56/mk5/evlbi? 0 : total: 3062244:  
loss : 0 (0.00%) :  
out-of-order: 0 ( 0.00%) :  
extent: 0 seqnr/pkt ;
```

```
2024.009.15:46:35.17/mk5=version?2024.009.15:46:35.17/mk5/!version? 0 : jive5ab :  
3.1.0-dev : 64bit : Debug :  
koll : 08-Jan-2024 :  
12h26m00s : nossapi :  
/usr/local/src/etransfer ;
```

```
2024.009.15:48:37.54;mk5=scan_check?
```

```
2024.009.15:48:37.59/mk5/scan_check? 0 :?:  
xxx_xx_no0001 : VDIF : 128 :  
2024y009d15h43m13.3117s :  
5.9812s : 256Mbps :  
1060224 : 8032 ;
```

```
2024.009.15:48:42.56;mk5=evlbi?
```

```
2024.009.15:48:42.56/mk5/evlbi? 0 : total: 3062244:  
loss : 0 (0.00%) :  
out-of-order: 0 ( 0.00%) :  
extent: 0 seqnr/pkt ;
```

128 x 256 Mbps = 32 Gbps  
on a standard Mark6  
with  
32 x 10 TB HDD (or larger)