

Control system based on a
Highly
Abstracted and
Open
Structure



!CHAOS

Few introductory slides
on Control Systems (CS)

CS evolution

- early in the 60's accelerators and experimental apparatus were remotely controlled by pulling wires from device to control room panels having conventional connectors, knobs, analog gauges of various types
- Pushed by the development of electronic and micro-electronic and, later, by the progress of the information technology (exponential growth of computing power, development of communications networks and an overall reduction of costs and increase of performance) computers entered the game.
- initially used for quasi-on line measurement or feedback
- (expensive) workstations first, (low cost) PC later



CS evolution

- in the early '70s CS included ~ 2500 I/O channels;
- large experimental apparatuses in the half of '90s reached ~ 100,000 I/O channels
- increasing man-power requirements for the developments of CS suggested to avoid home made solutions and independent developments
- tendency is shared development or simply implementation of open source solutions
- in the case of development from scratch limit the number of HW/SW technologies used but offer access points to major languages and SW tools



CS evolution

- nowadays CS are integral part not just of the accelerator, but of the entire infrastructure
- all components of the accelerator are managed by the CS (also the local control during maintenance) which then handles thousands of I / O channels of all types
- CS extend the performance of diagnostic systems and allow implementing new features for handling equipment
- information technology offers solutions for providing operators and scientists with more versatile and simple tools for operation



CS evolution



remote operations

A remote operations center, LHC@FNAL, has been built at Fermilab to make it easier for accelerator scientists and experimentalists working in North America to help commission and participate in operations of the LHC and experiments.





October 15 - 19, 2007
Knoxville Convention Center
Knoxville, Tennessee

General Info ▶
Venue ▶
Registration
Conference Guide
Program
Committees
Program Tracks
Author Info ▶
Vendors ▶
Affiliated Meetings
Participants List

International Conference
on
Accelerator and Large
Experimental Physics
Control Systems

Covering control system topics for
accelerators, telescopes, fusion,
physics detectors, space
exploration, and more

Preliminary Proceedings

Hosted by the
Oak Ridge
National Laboratory
Spallation Neutron
Source
and
Thomas Jefferson
National Accelerator
Facility
Jefferson Lab

Conference Chair: Dave Gurd (ORNL/SNS)
Program Chair: Karen White (JLab)
Conference Coordinator: Lori Lane







Control systems communities



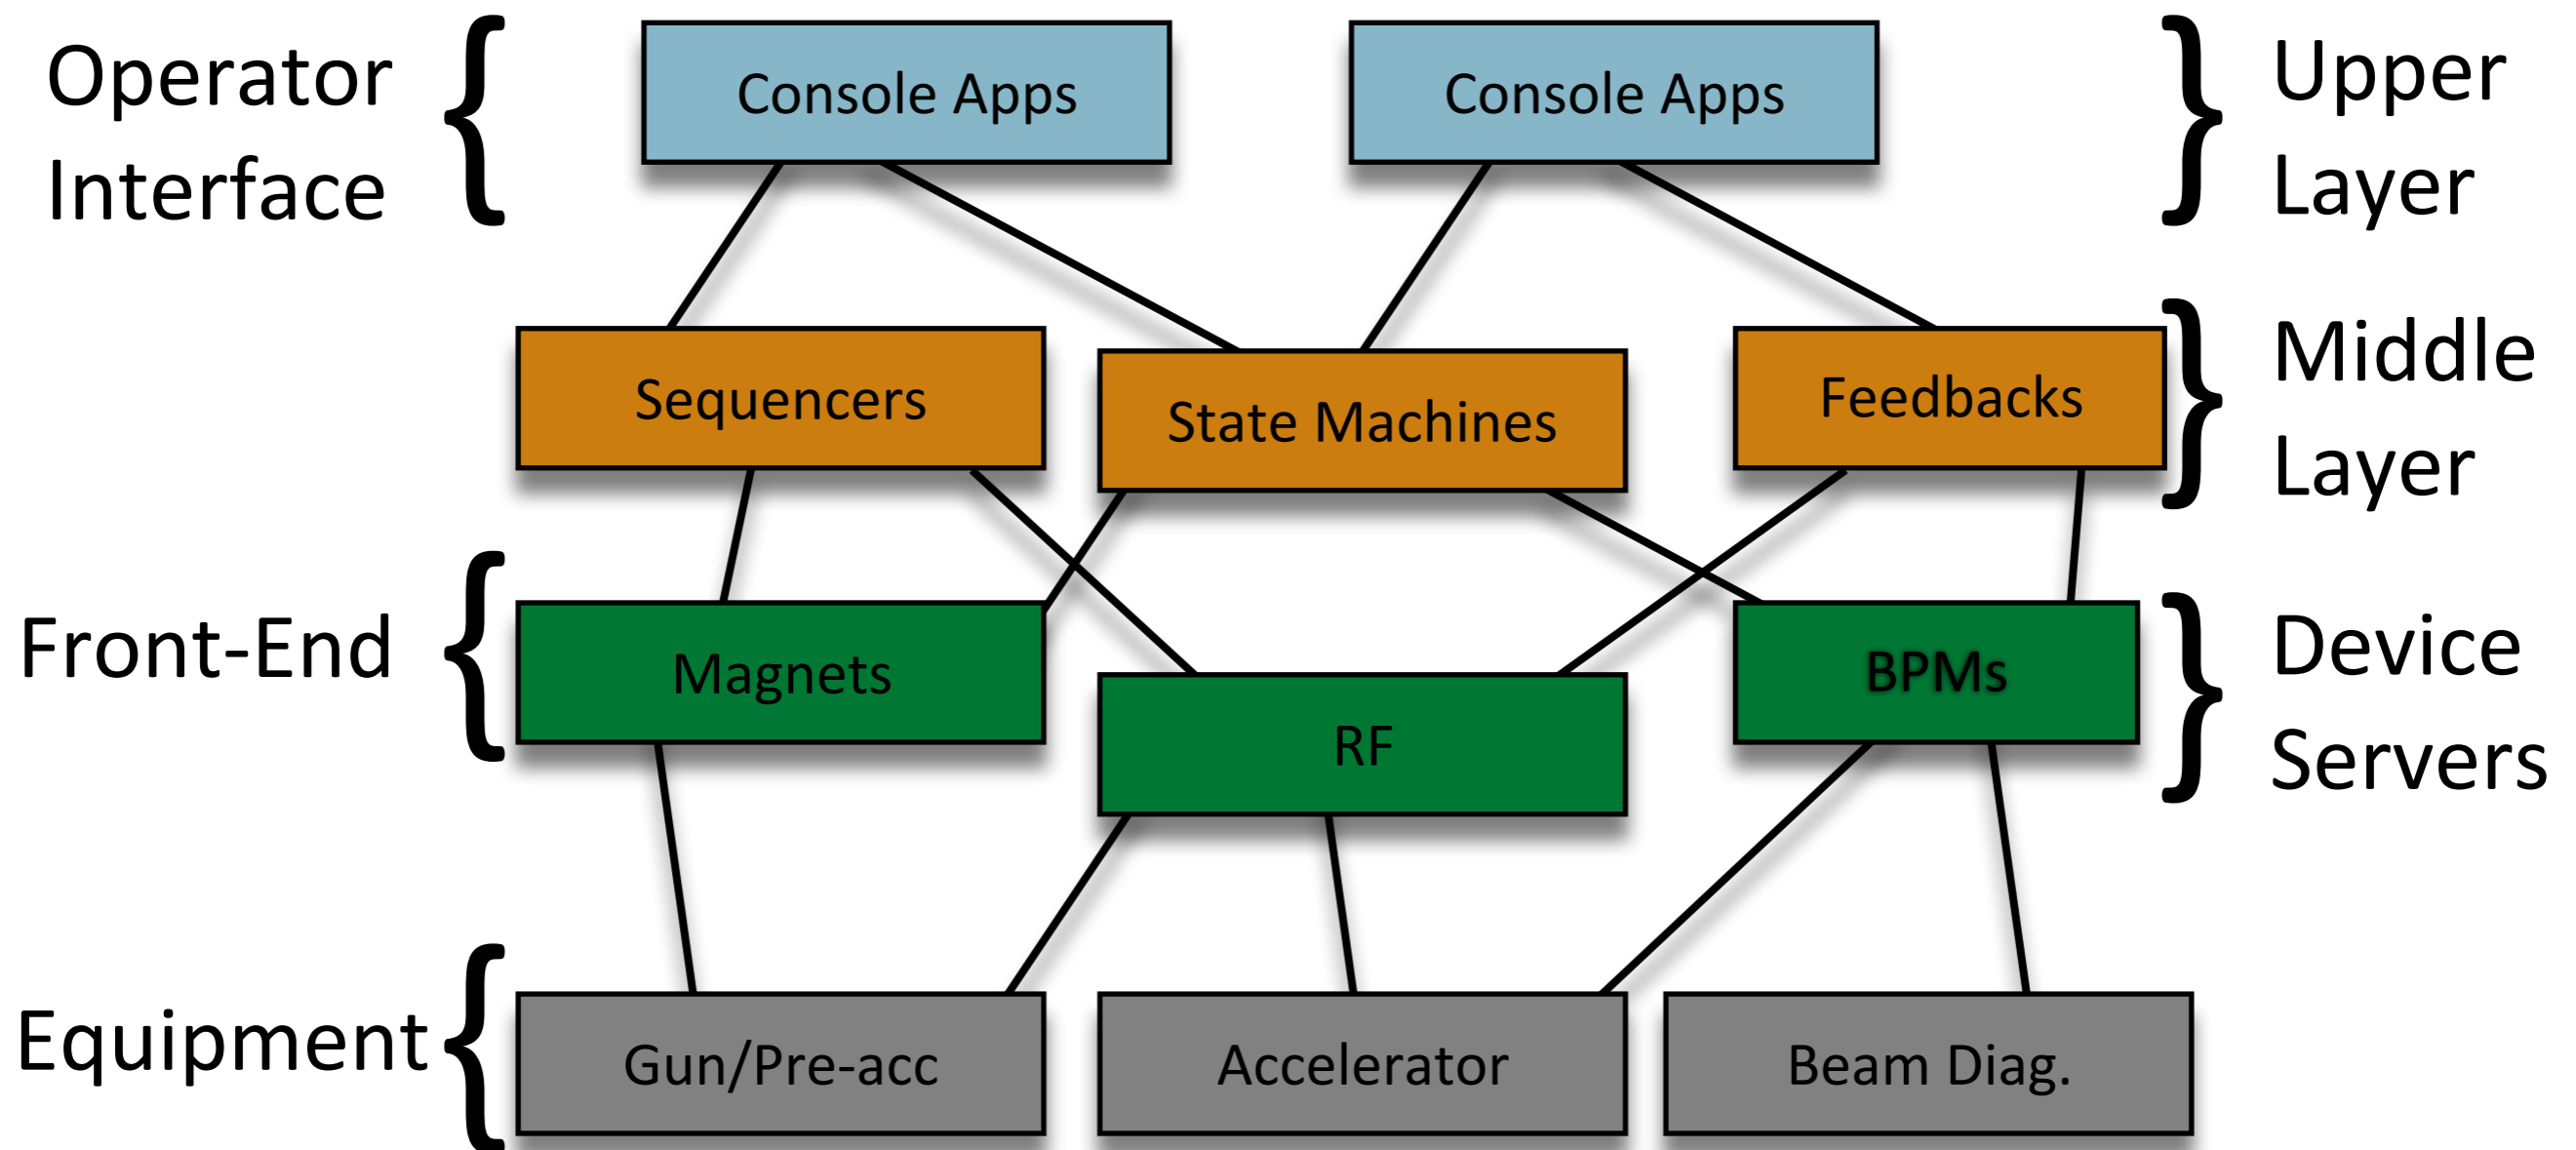
14-17 October 2002 - Frascati, Italy

The 4th International Workshop
on Personal Computers and
Particle Accelerator Controls

PCaPAC2002 is jointly organized
by INFN-LNF and INFN-RomaII
and sponsored by
Istituto Nazionale di Fisica Nucleare




the “standard model”



the “standard model”

First Level

- display data and status of accelerator devices by means of customizable control panels.
- on-line data analysis
- errors and alarms notification
- set-points back up and restore

the “standard model”

Second Level

- communication
- database
- feedback software and “state machine”
- multi-element control procedures
- more...
- access-point for external tools

the “standard model”

Third Level

- interface to accelerator components by means of:
 - I/O channels (ex. an analog voltage output to control the set-point of a magnet power supply)
 - field bus (distributed I/O channels: CANbus, serial etc.)
- interface to stand-alone sub-systems (interlocks, security, components with their own controls)

the “standard model”

distributed controls

- for each level the software development and more appropriate hw can be chosen (with some limitations)
- links between the levels of the system are provided by the communication system
- real time procedures are located near the equipment
- if the communication system is modular, expansions or upgrades of S.diC. can be obtained simply by adding additional third level CPUs

Communication: ethernet networks or...

- communication solutions are now widely based on LAN (Local Area Network)
- early in the '90s alternatives to the network were still considered in medium size CS (VME vertical bus, reflective memory boards)
- though non-deterministic, Ethernet has proven its validity as a physical-layer for communication protocols of CS
- standard networks (1-10 Gb/s) provide sufficient bandwidth, especially if using dedicated networks for CS....
- ..but, if possible, limit data transfer by implementing calculations directly on the front-end

Existing CS: EPICS

What is EPICS?

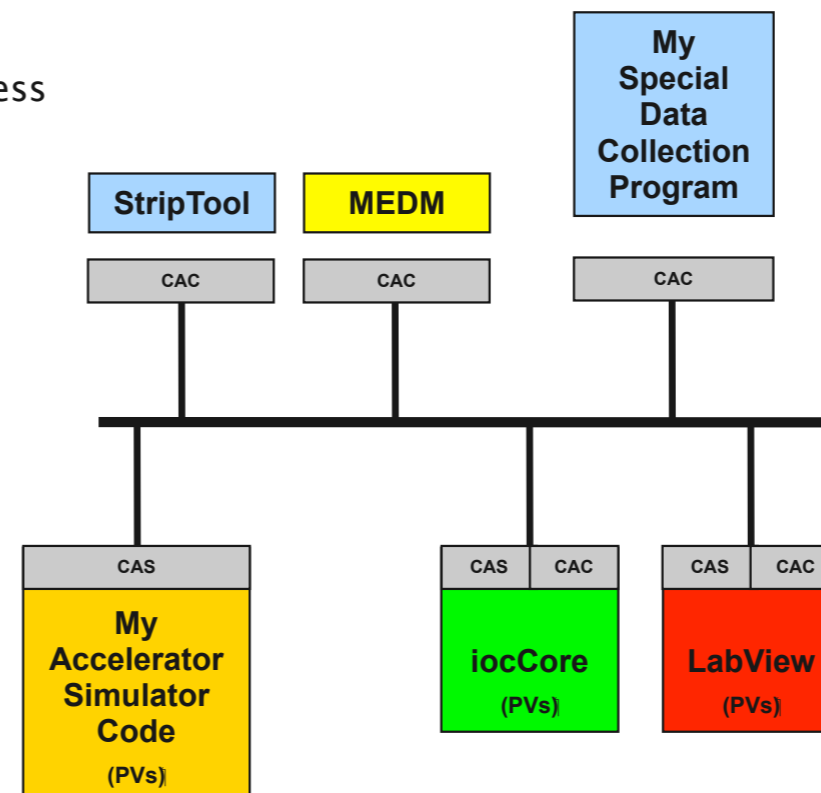
- Early Major Collaborators
 - America
 - ANL, LANL, LBL, ORNL, SLAC, JLAB
 - Europe
 - DESY, BESSY, PSI
 - Asia
 - KEK
- Recent Major Collaborators
 - America
 - TRIUMF, CLS, BNL
 - Europe
 - DIAMOND Light Source
 - Asia/Australia
 - J-PARC, Australian Synchrotron, IHEP, SSRF, TLS, RRCAT
- Many Other Collaborators/Users

Existing CS: EPICS

What is EPICS?

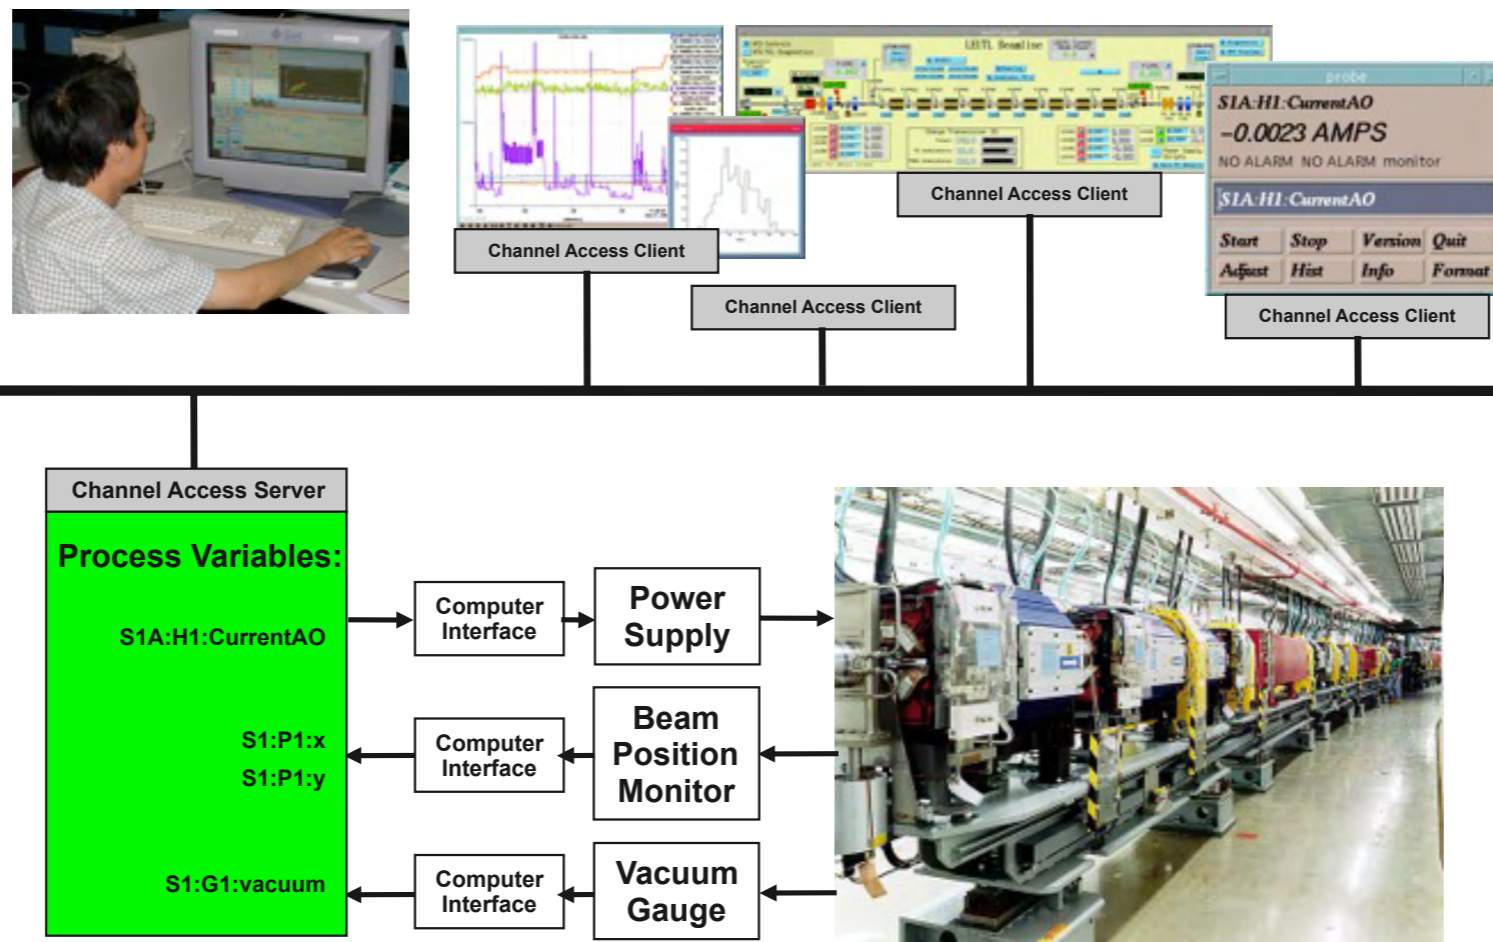
- Any tool/program/application that abides by the Channel Access protocol could be described as “EPICS Compliant”.

EPICS can be viewed as a “toolkit” of EPICS compliant programs. One can select the appropriate tool for their need or develop their own.



Existing CS: EPICS

How does it do it?



January 2009

EPICS Seminar RRCAT

16

Existing CS: TANGO



The TACO Next Generation Objects (TANGO) control system is a **free open source object-oriented** control system for controlling accelerators, experiments and any kind of hardware or software being actively developed by a consortium of (mainly) synchrotron radiation institutes.

TANGO is a distributed control system.

It runs on a single machine as well as hundreds of machines.

TANGO uses the **omniorb** implementation of CORBA as its network protocol.

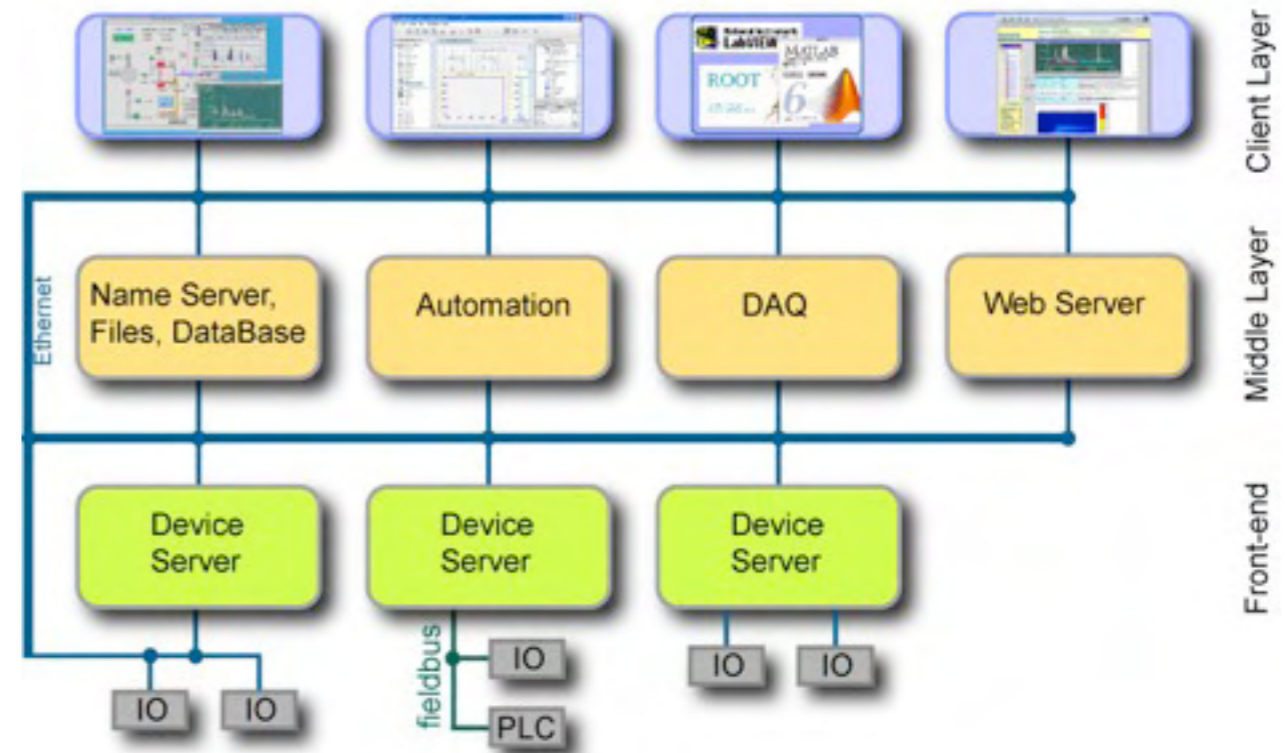
The **client-server** model is the basic communication model.

Communication between clients and servers can be synchronous, asynchronous or event driven.

Existing CS: DOOCS



- The Distributed Object Oriented Control System (DOOCS) it is completely written in C++ language and follows the object oriented programming paradigm [2] since the design idea from the device level, including the device servers, up to the user display is based on objects.




- In general, the object oriented approach implements multiple devices of the same type in a device server process on the lowest tier. A certain device is represented as an instance of a device class of its type. The data of a device, that is accessible as a property of the device from the network, is implemented as data class in a server library.

Existing (industrial) CS: SCADA

- SCADA is an acronym for Supervisory Control and Data Acquisition. SCADA systems are used to monitor and control a plant or equipment in industries such as telecommunications, water and waste control, energy, oil and gas refining and transportation

SCADA systems consist of:

- One or more field data interface devices, usually RTUs, or PLCs, which interface to field sensing devices and local control switchboxes and valve actuators
- A communications system used to transfer data between field data interface devices and control units and the computers in the SCADA central host. The system can be radio, telephone, cable, satellite, etc., or any combination of these.
- A central host computer server or servers (sometimes called a SCADA Center, master station, or Master Terminal Unit (MTU))
- A collection of standard and/or custom software [sometimes called Human Machine Interface (HMI) software or Man Machine Interface (MMI) software] systems used to provide the SCADA central host and operator terminal application, support the communications system, and monitor and control remotely located field data interface devices



back to !CHAOS

!CHAOS is a **INFN** experiment aimed to defining and validating a new paradigm for distributed control systems (DCS) and fast data acquisition (DAQ) for particle accelerators and experimental apparatuses.



Indeed, its flexibility allows a much wider spectrum of applications.

Goals:

- design an innovative control system by implementing the **newest software technologies**, mainly borrowed from the high performance internet services
- introducing the concept of **Control Service** by using a **new services topology** and...
- ...a **new communication solution for live data**, as alternative to client/server
- **high scalability** for each core service;
- **high abstraction** of devices control software;

“the standard model”

control room




accelerator

“the standard model”

control room



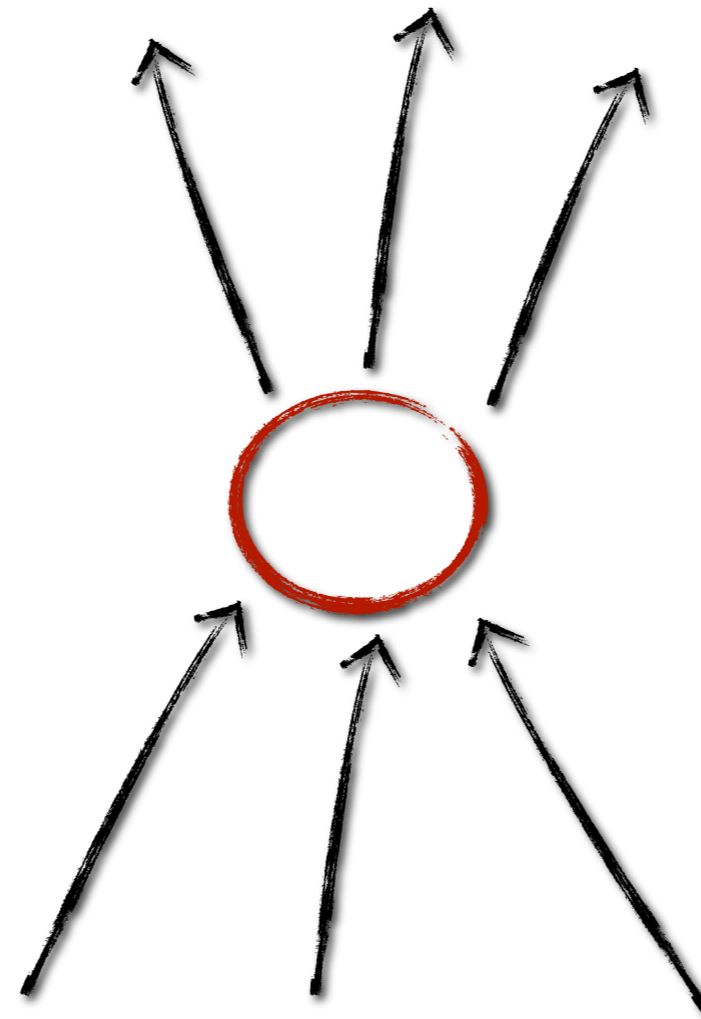
two different topologies



accelerator

the new paradigm


control room



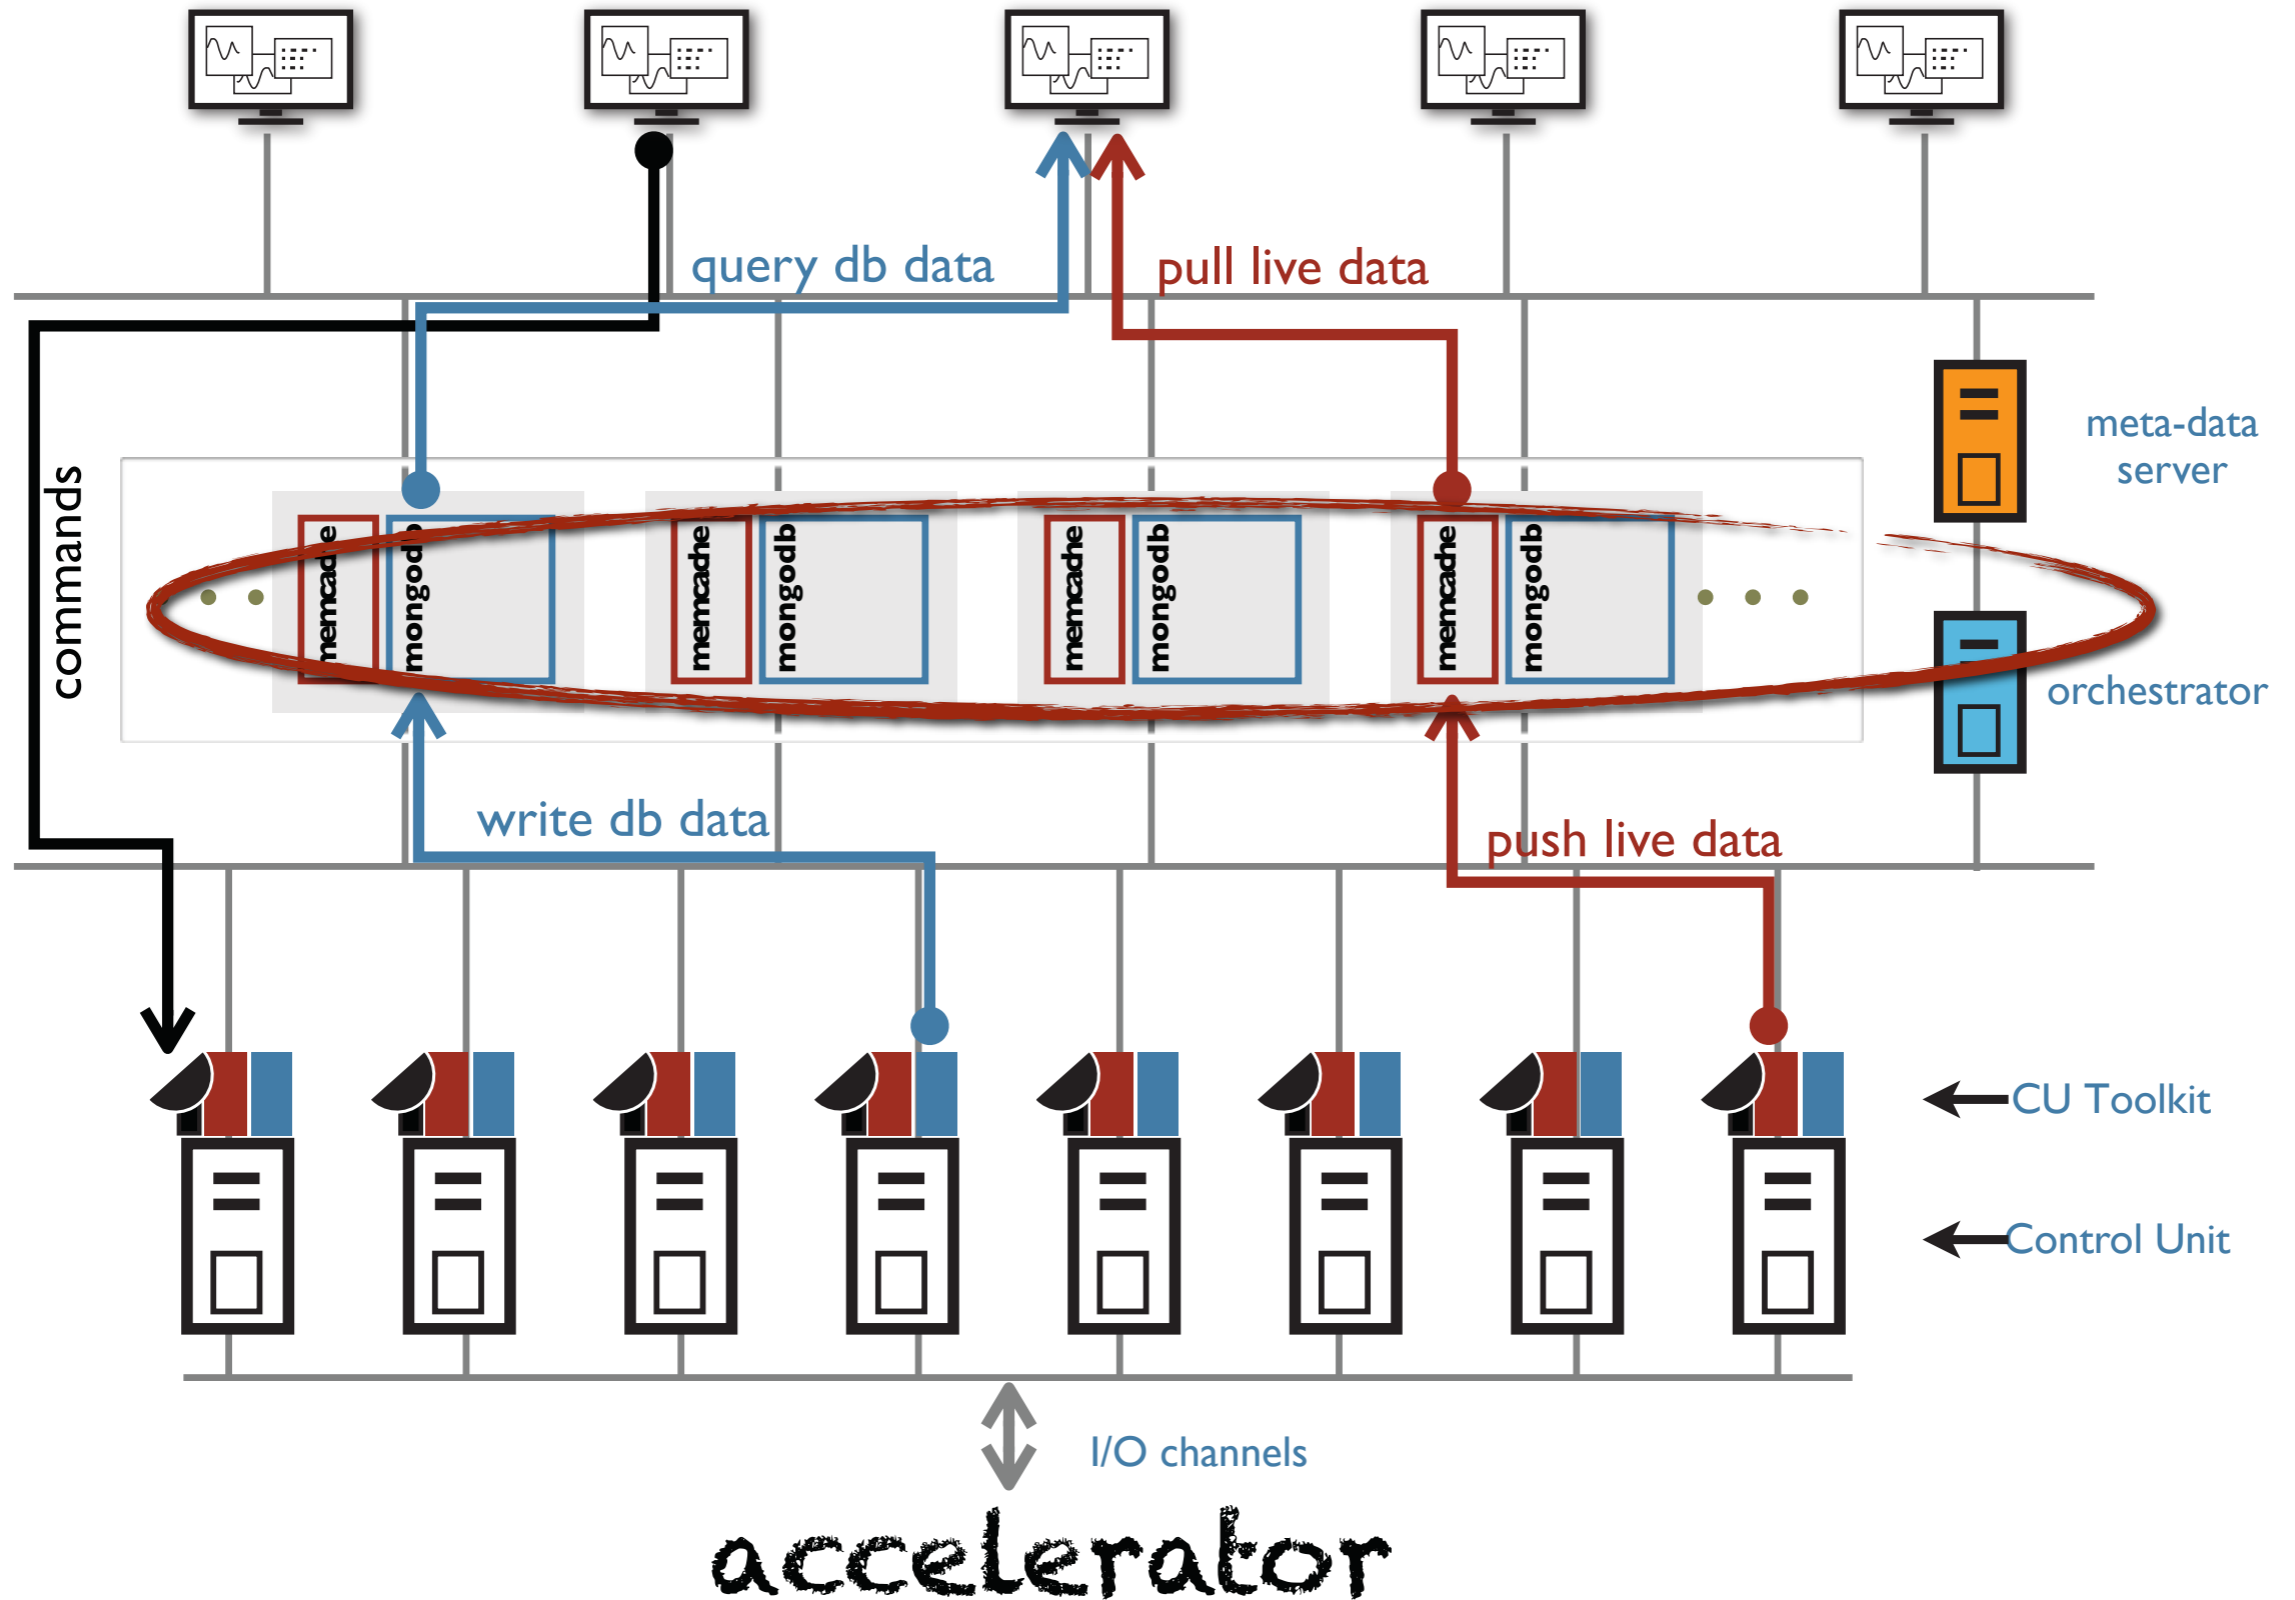
← bottleneck ?

accelerator

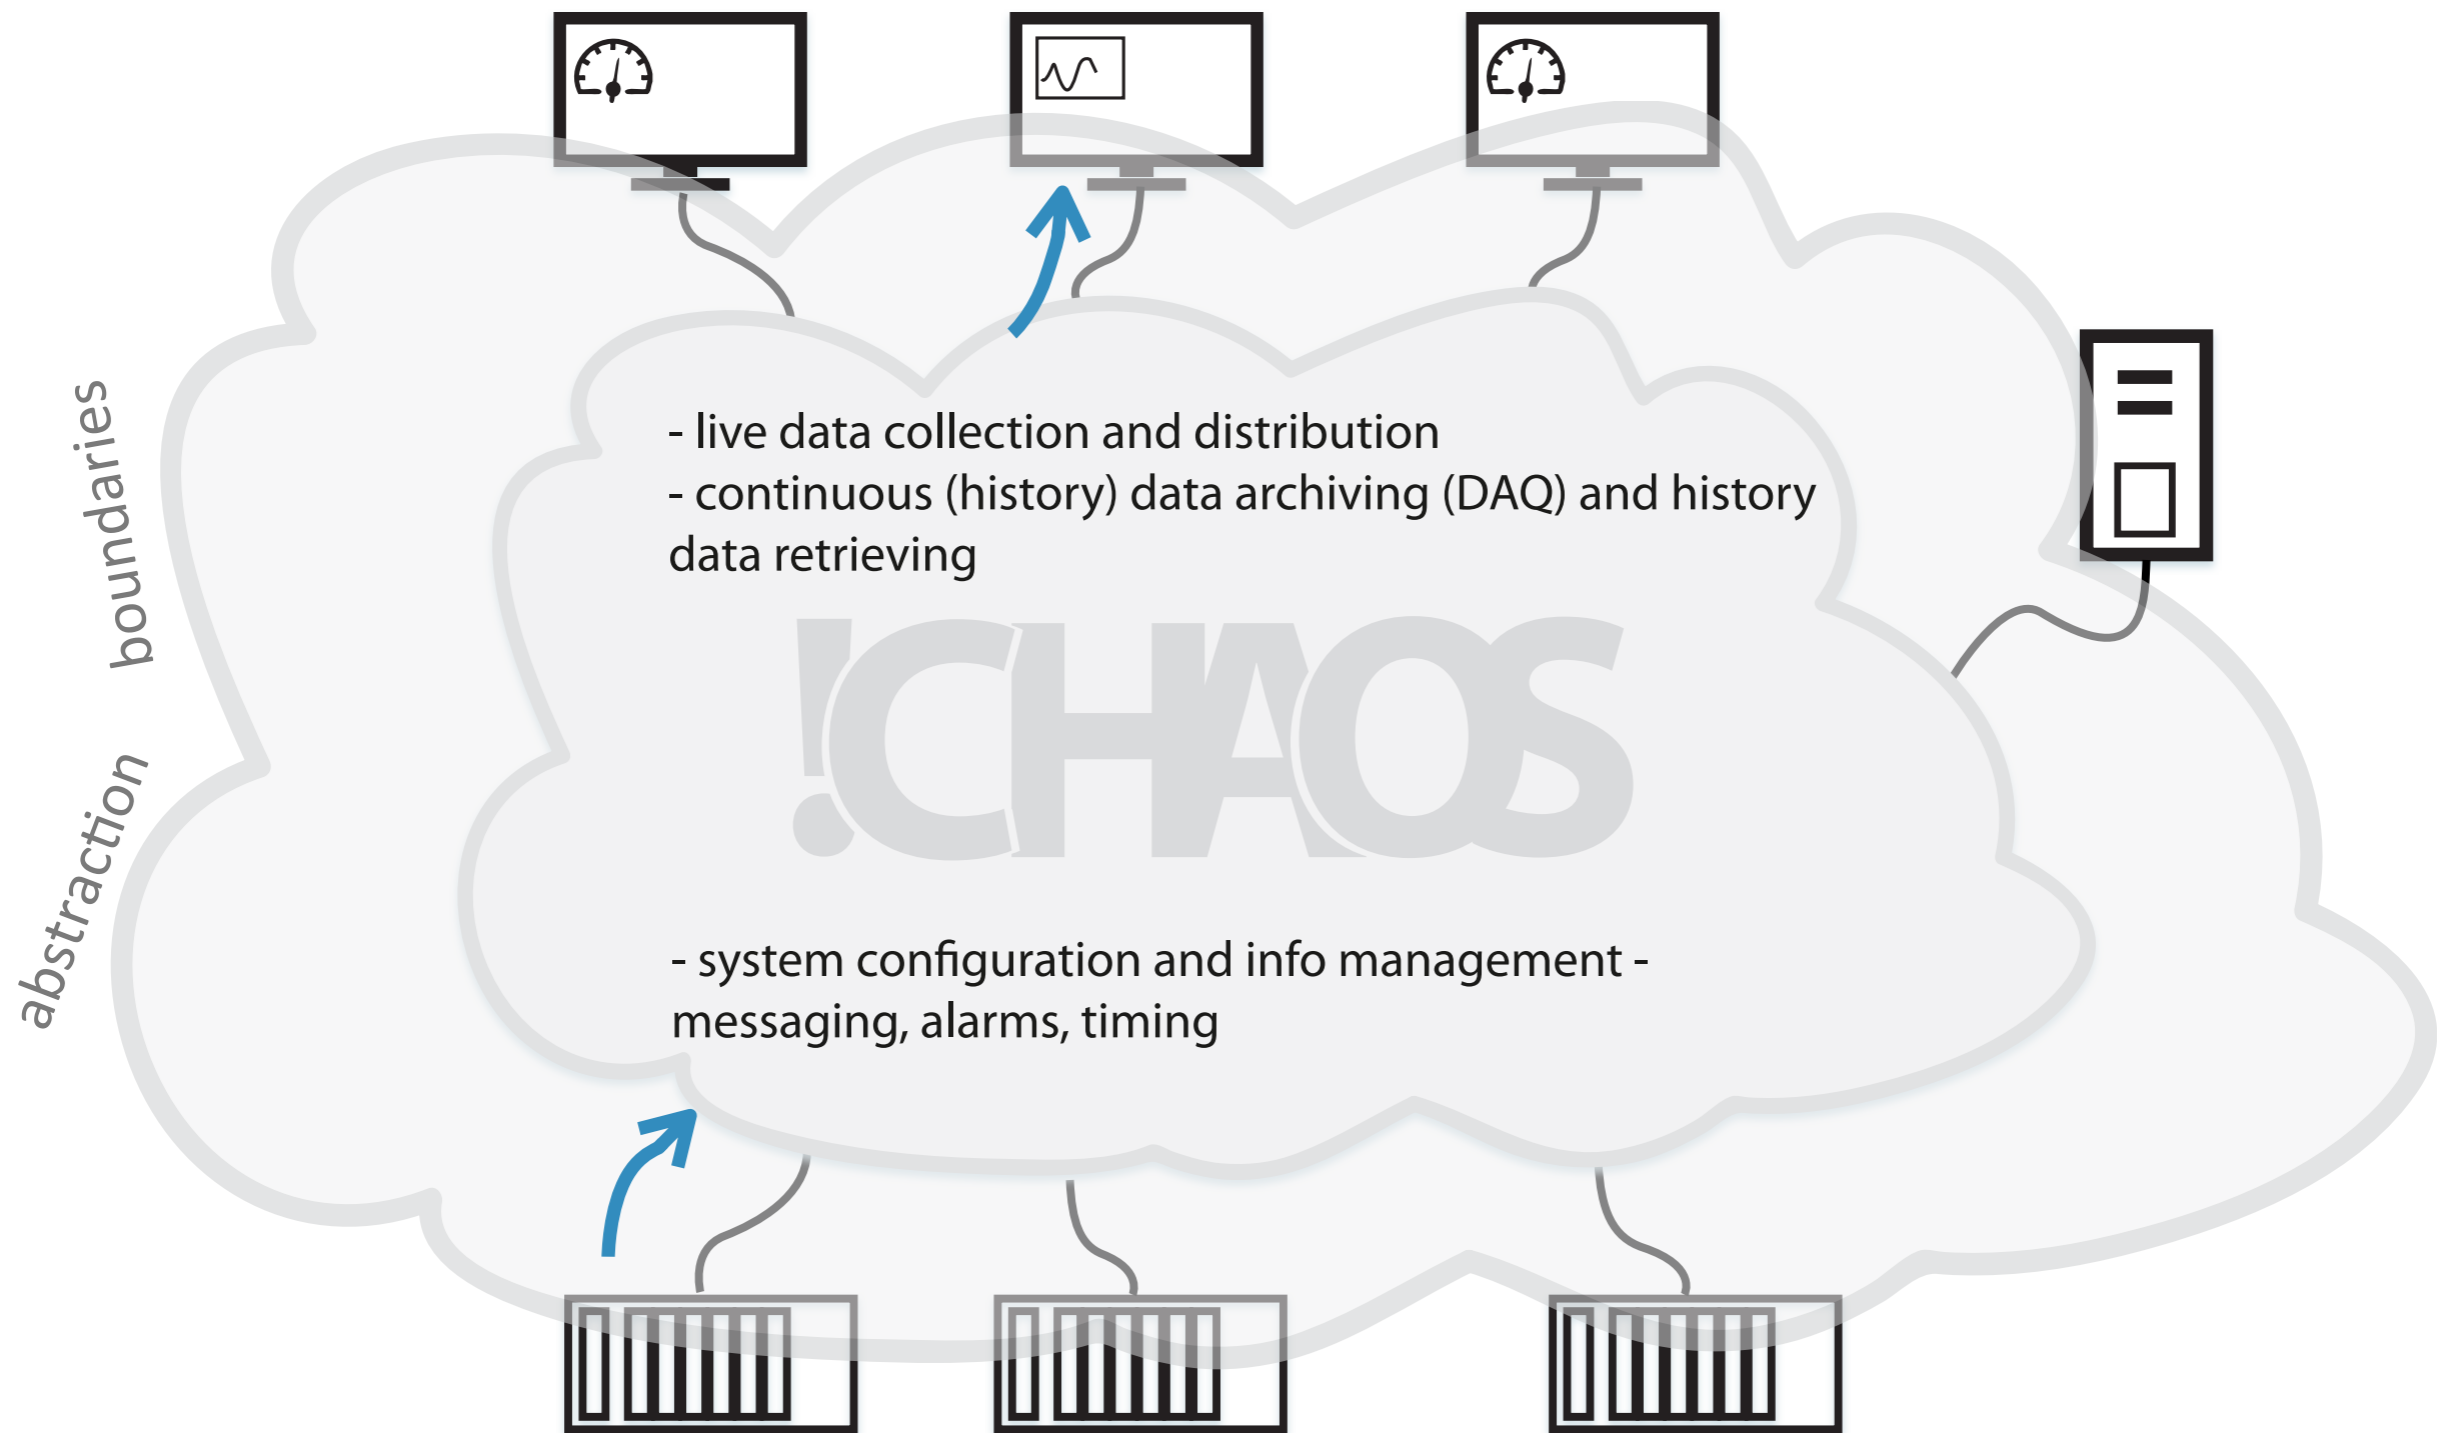
the new paradigm: scalability eliminates bottleneck



!CHAOS Control System topology



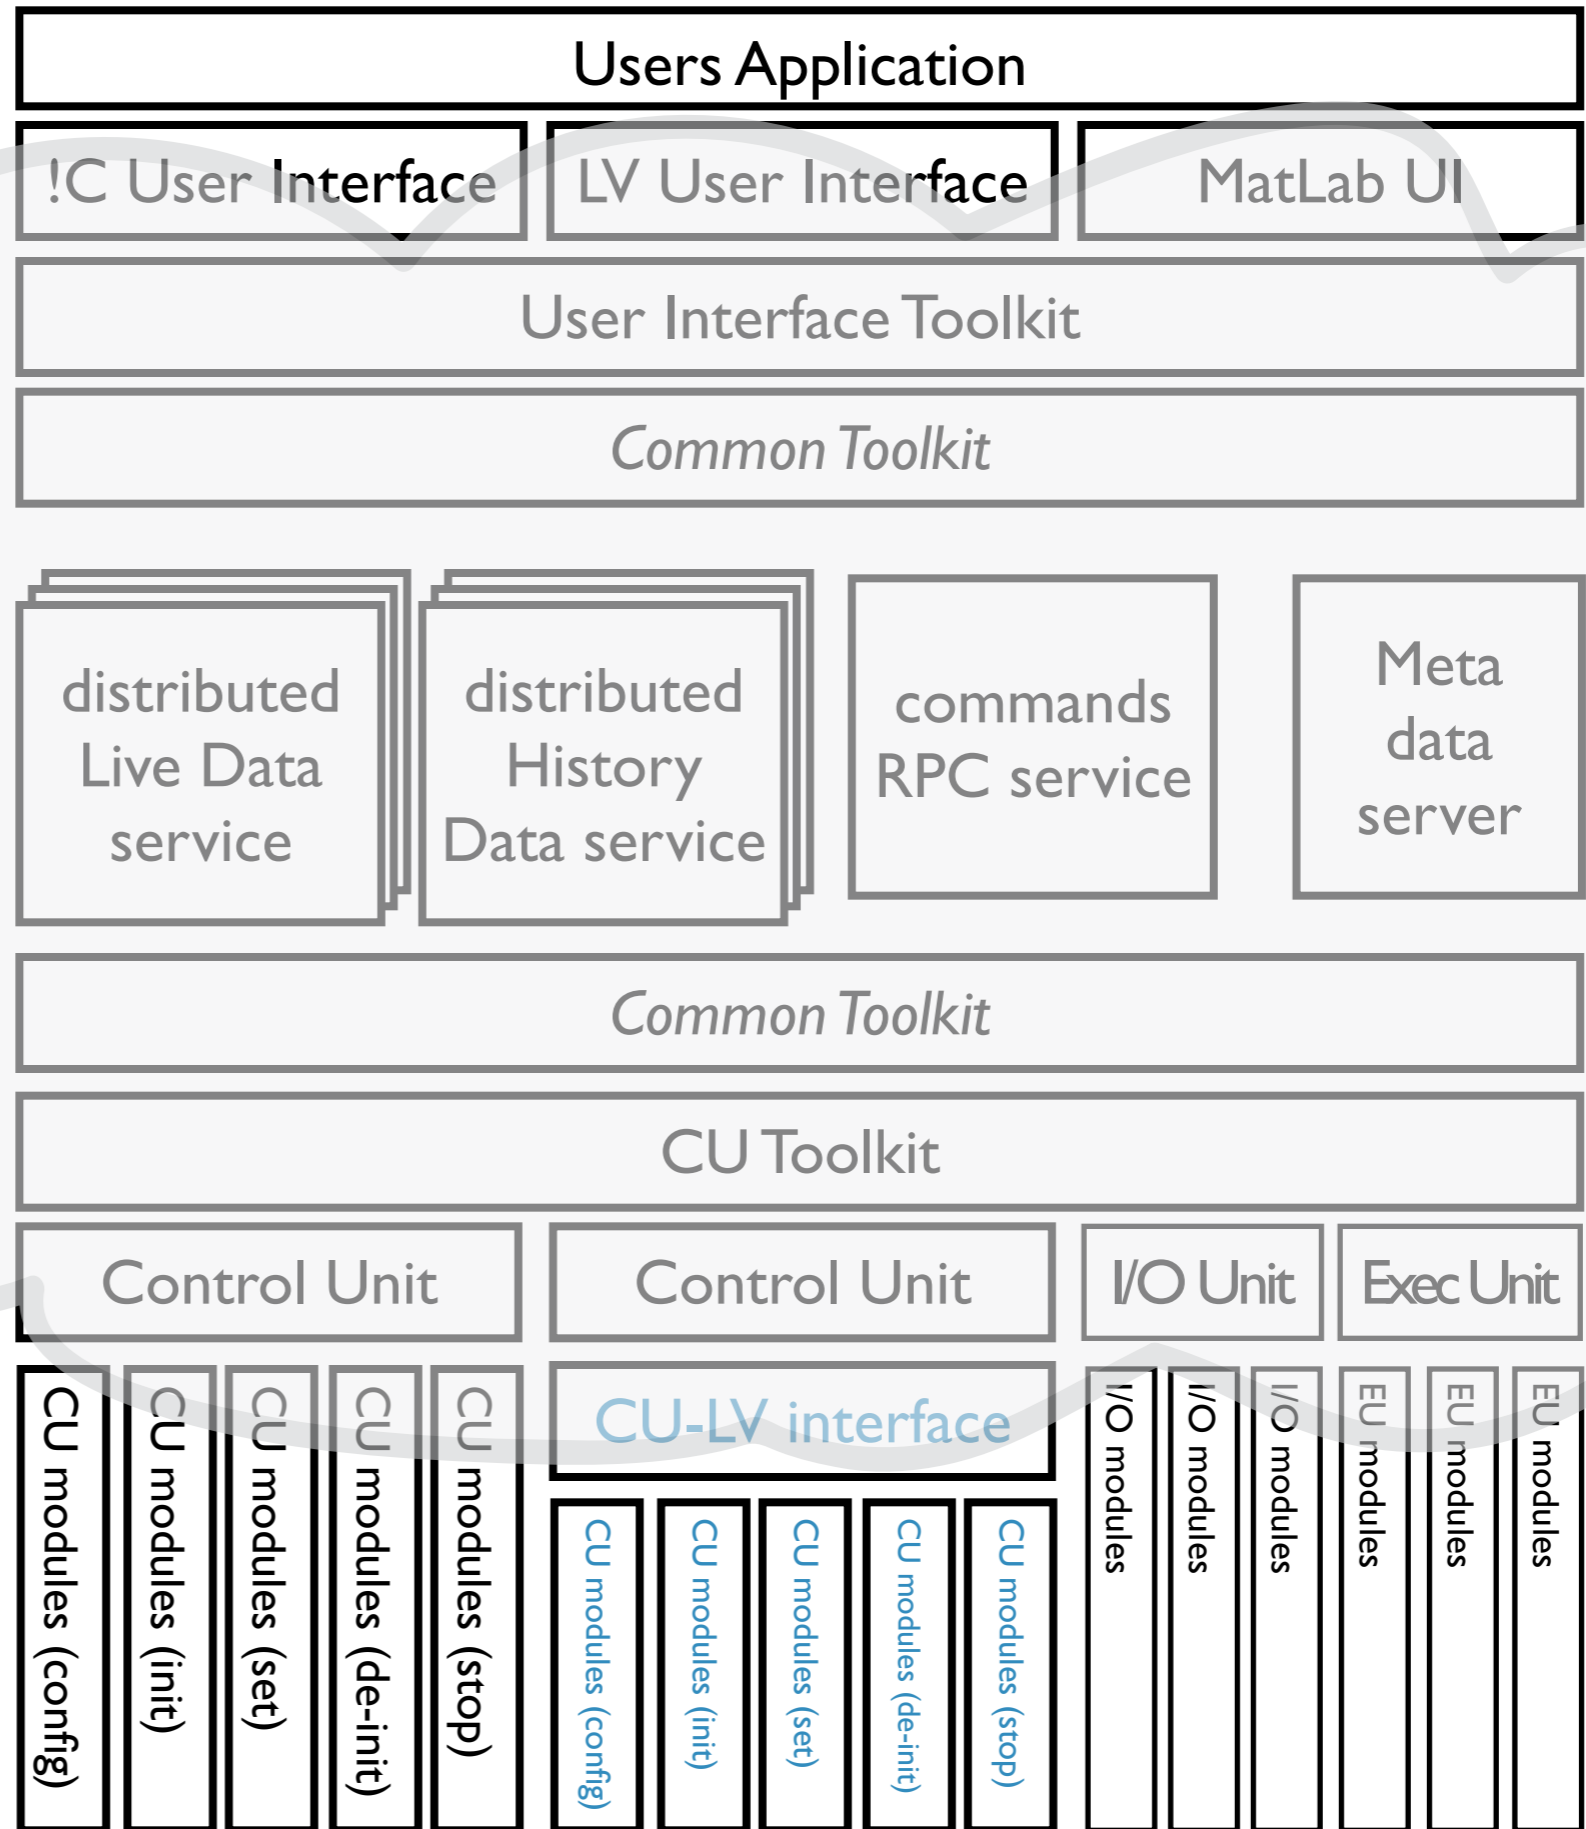
“data consumer” clients



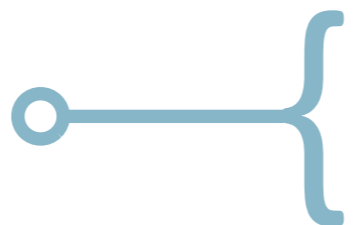

“data producer” clients

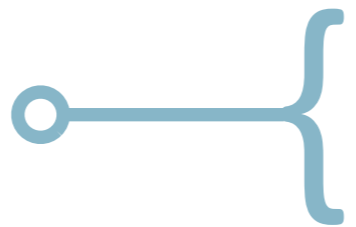



!CHAOS
infrastructure

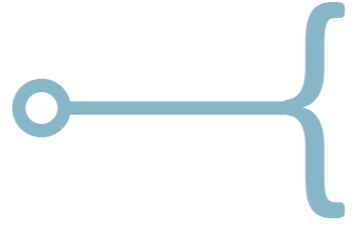



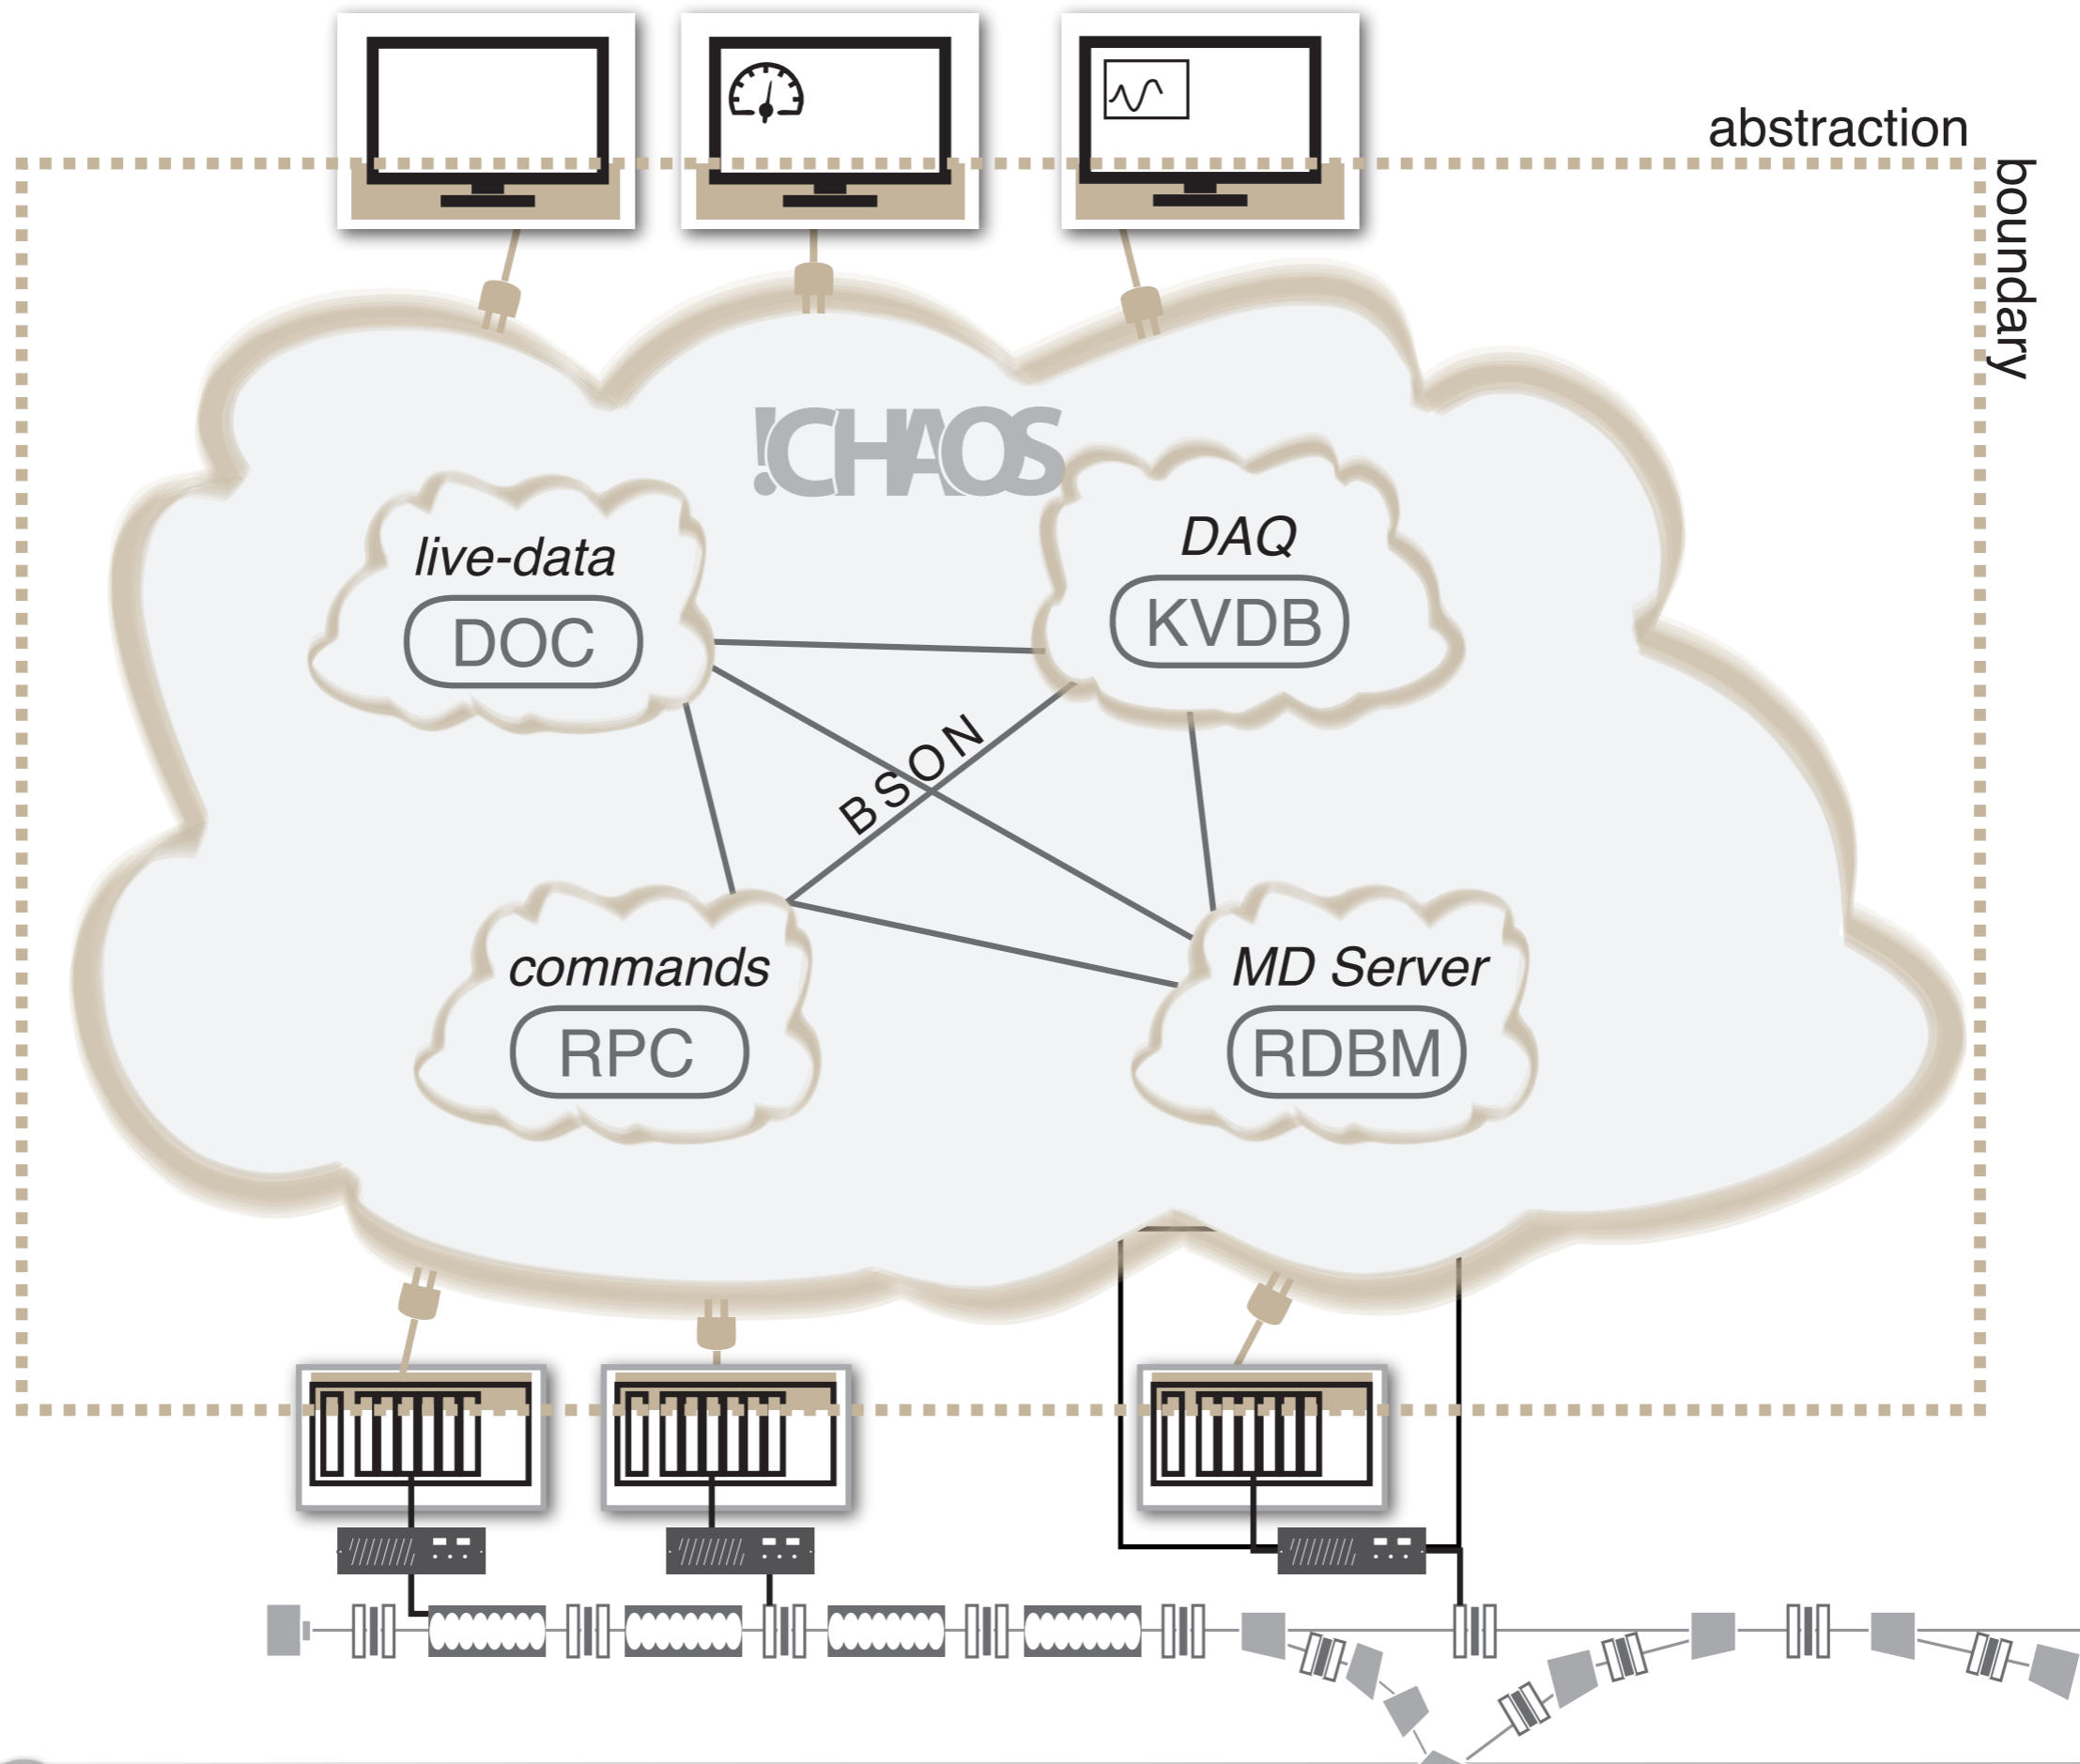
!CHAOS keywords

scalability   of performance
and size

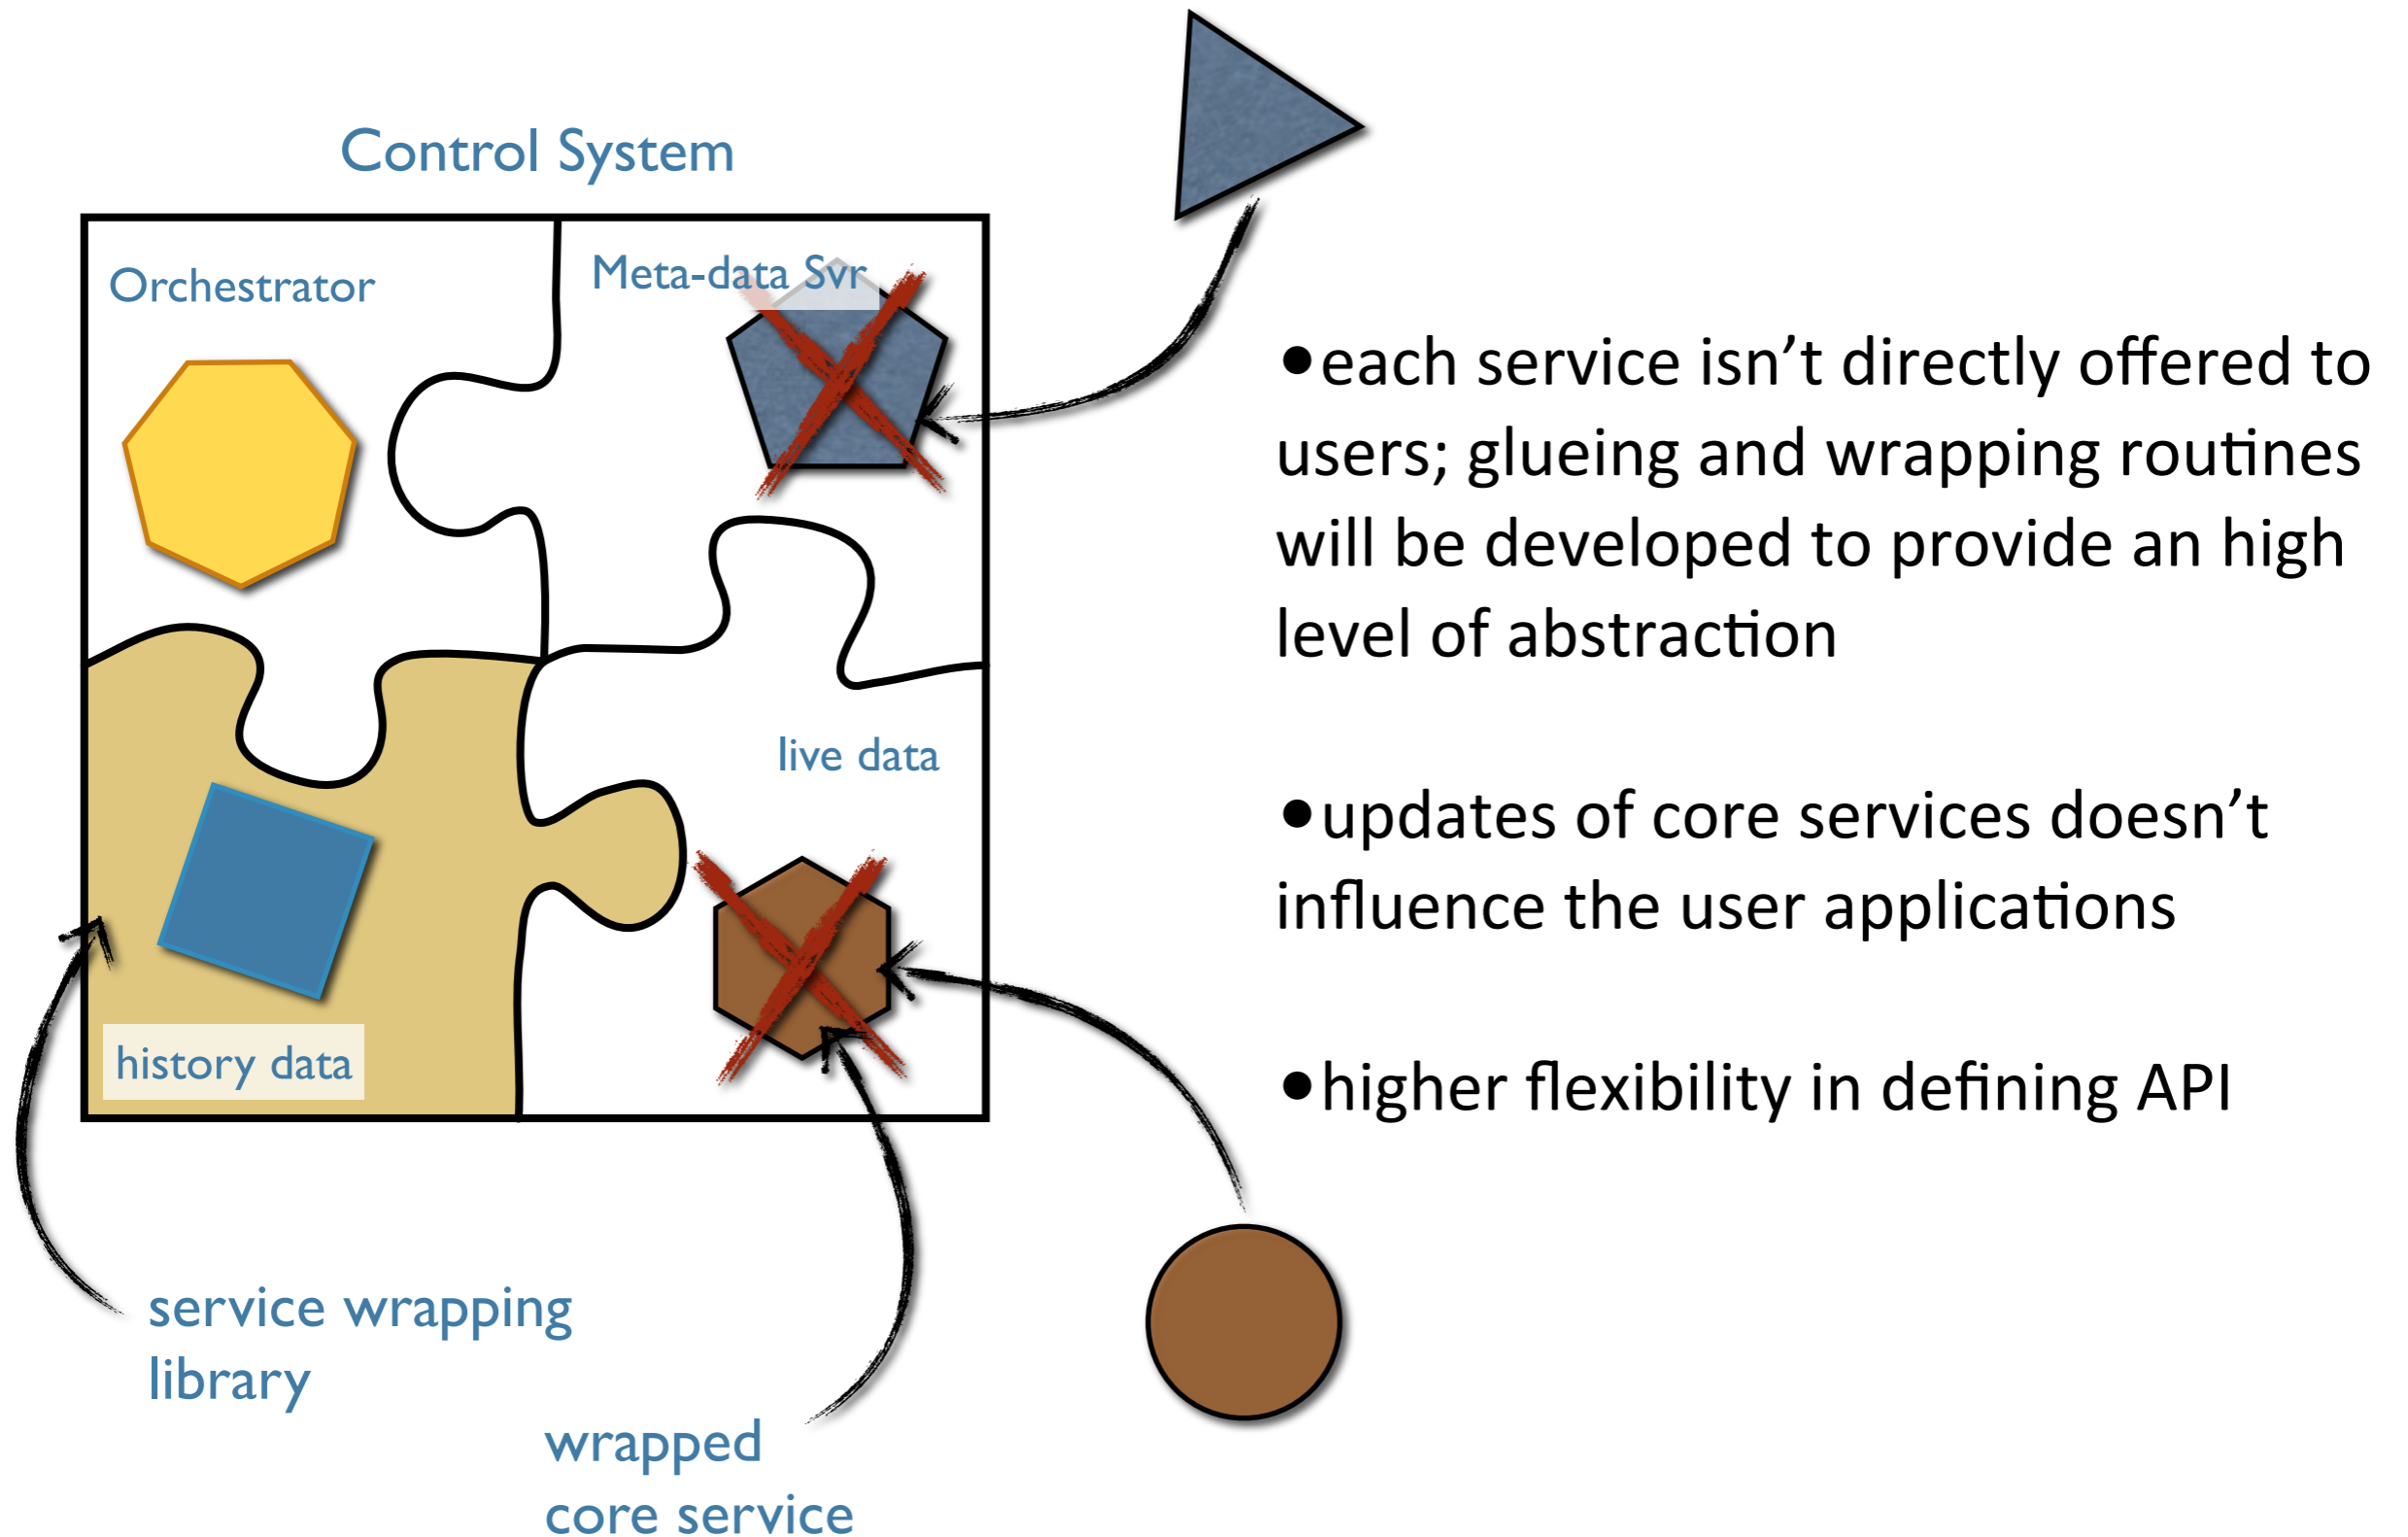
abstraction   of services
and device

integration   all services provided by the SC

**fast, easy and
modular
customization**   devices' drivers
development only



Abstraction of components

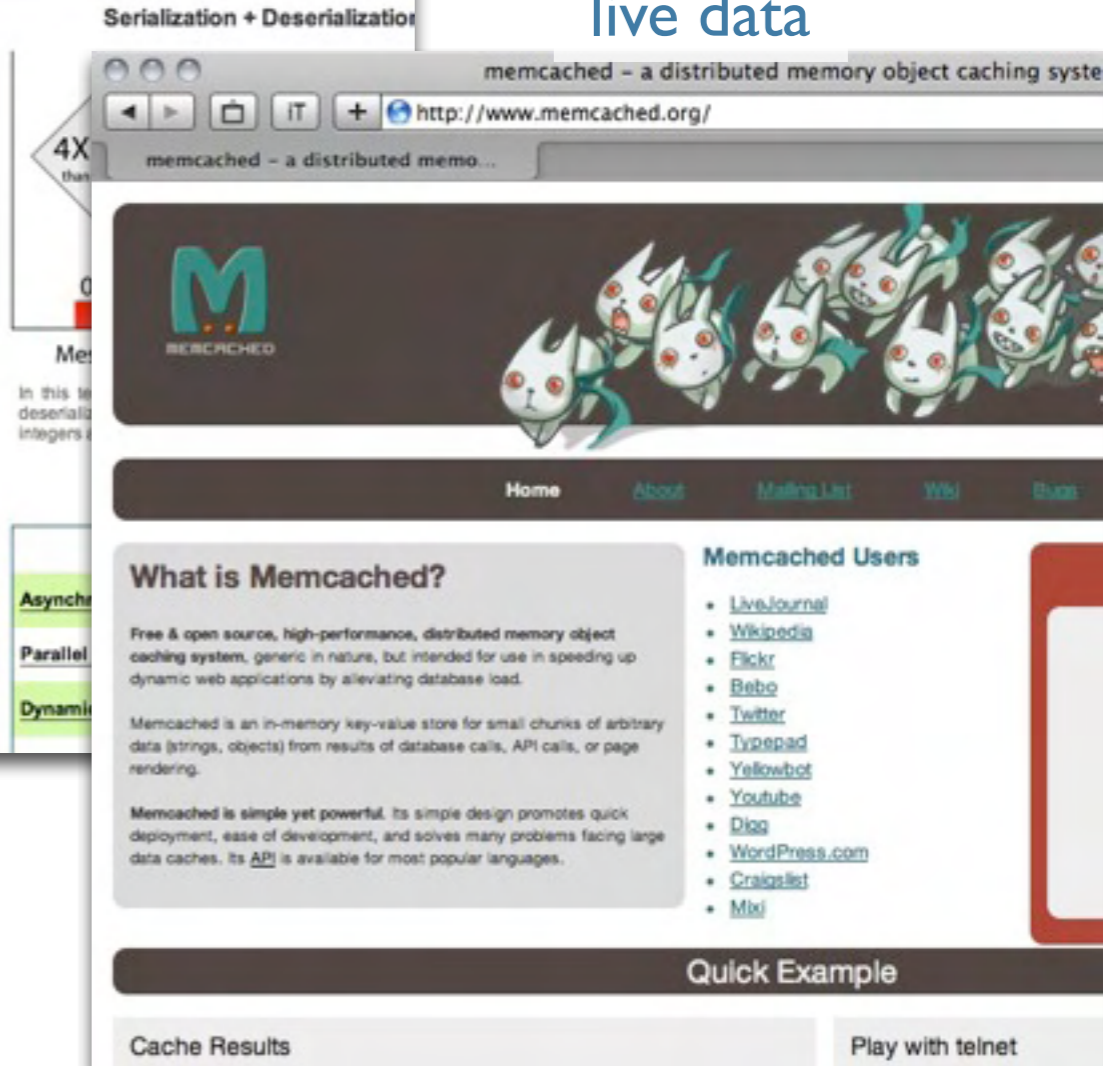
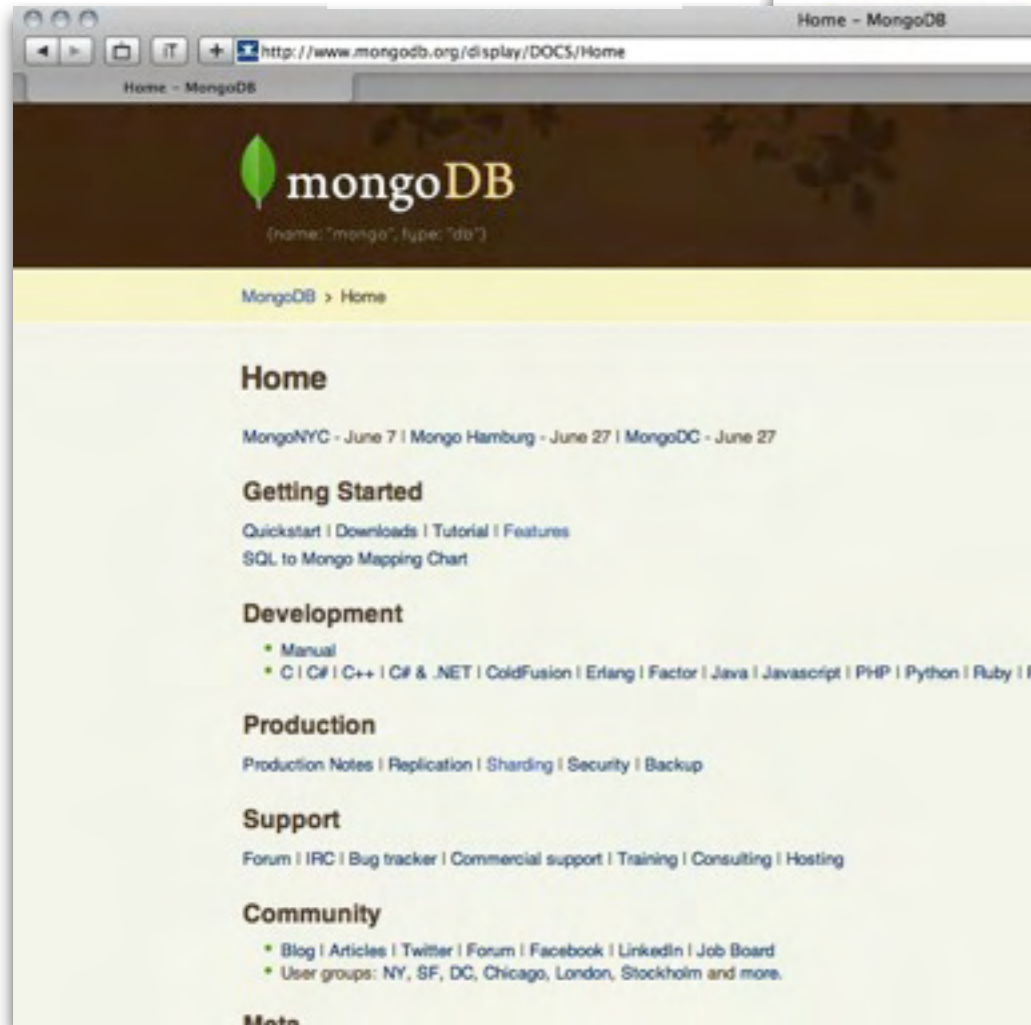
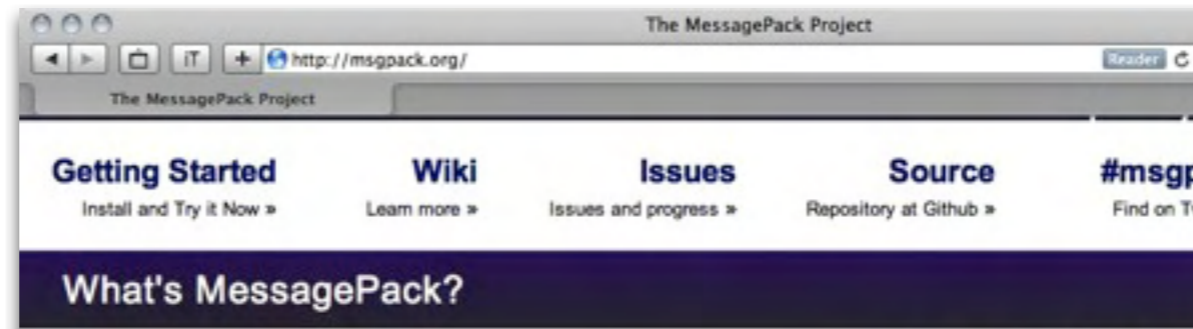


core services


commands/events

history db

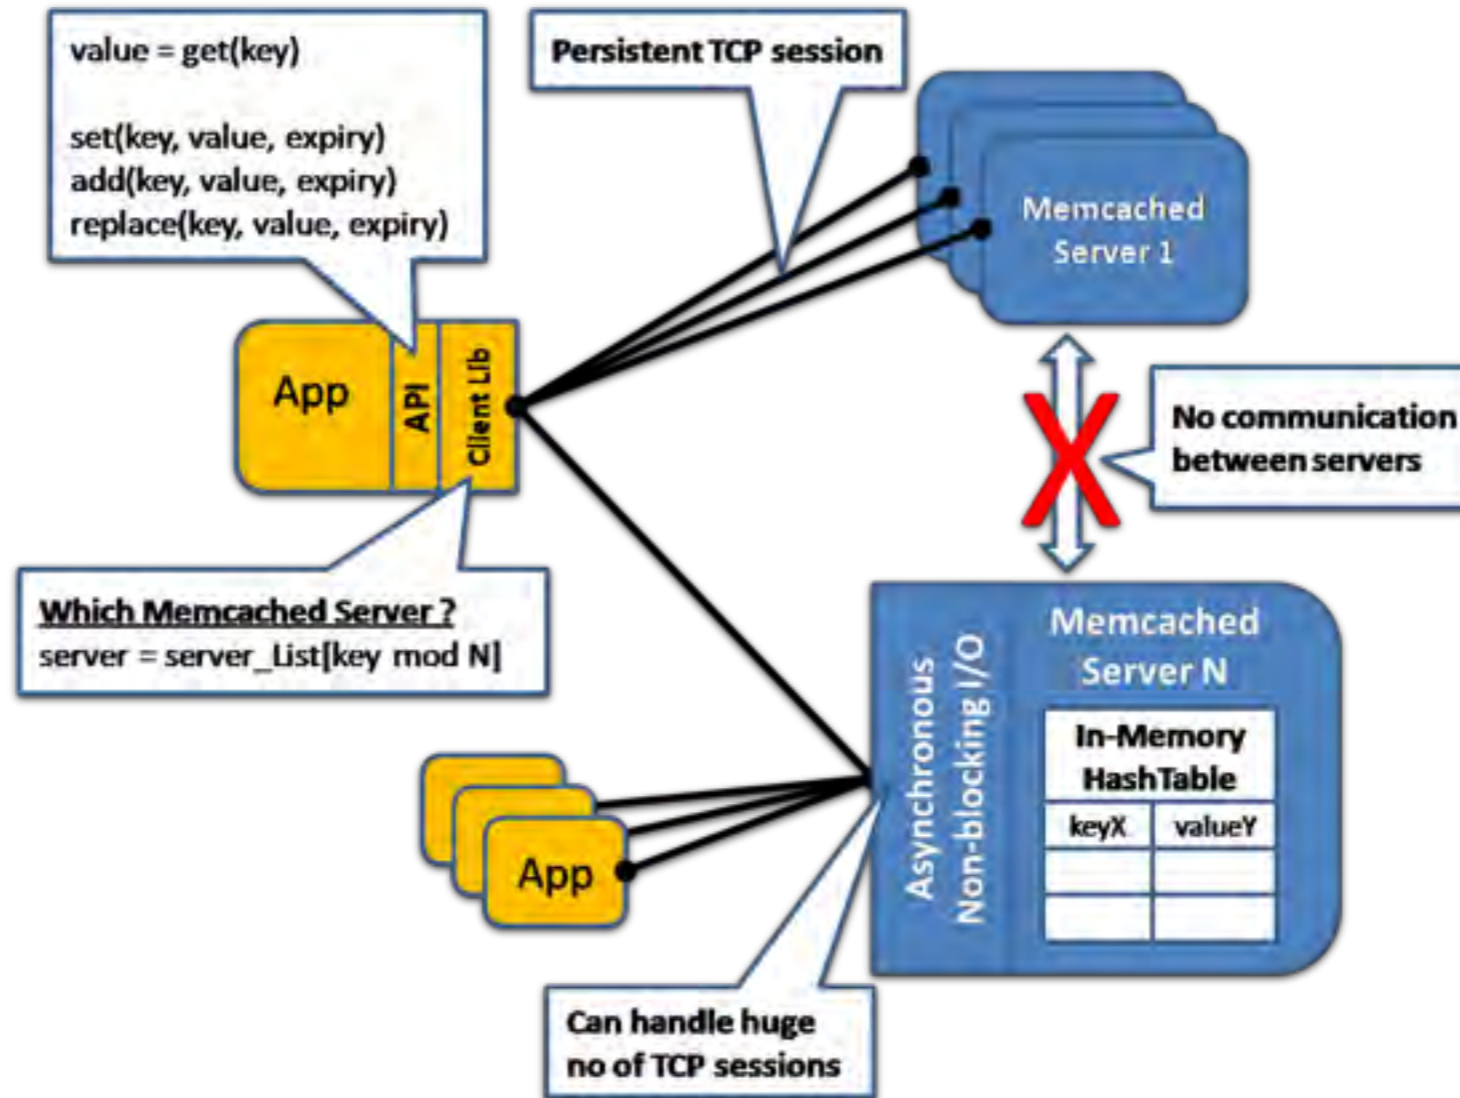
live data



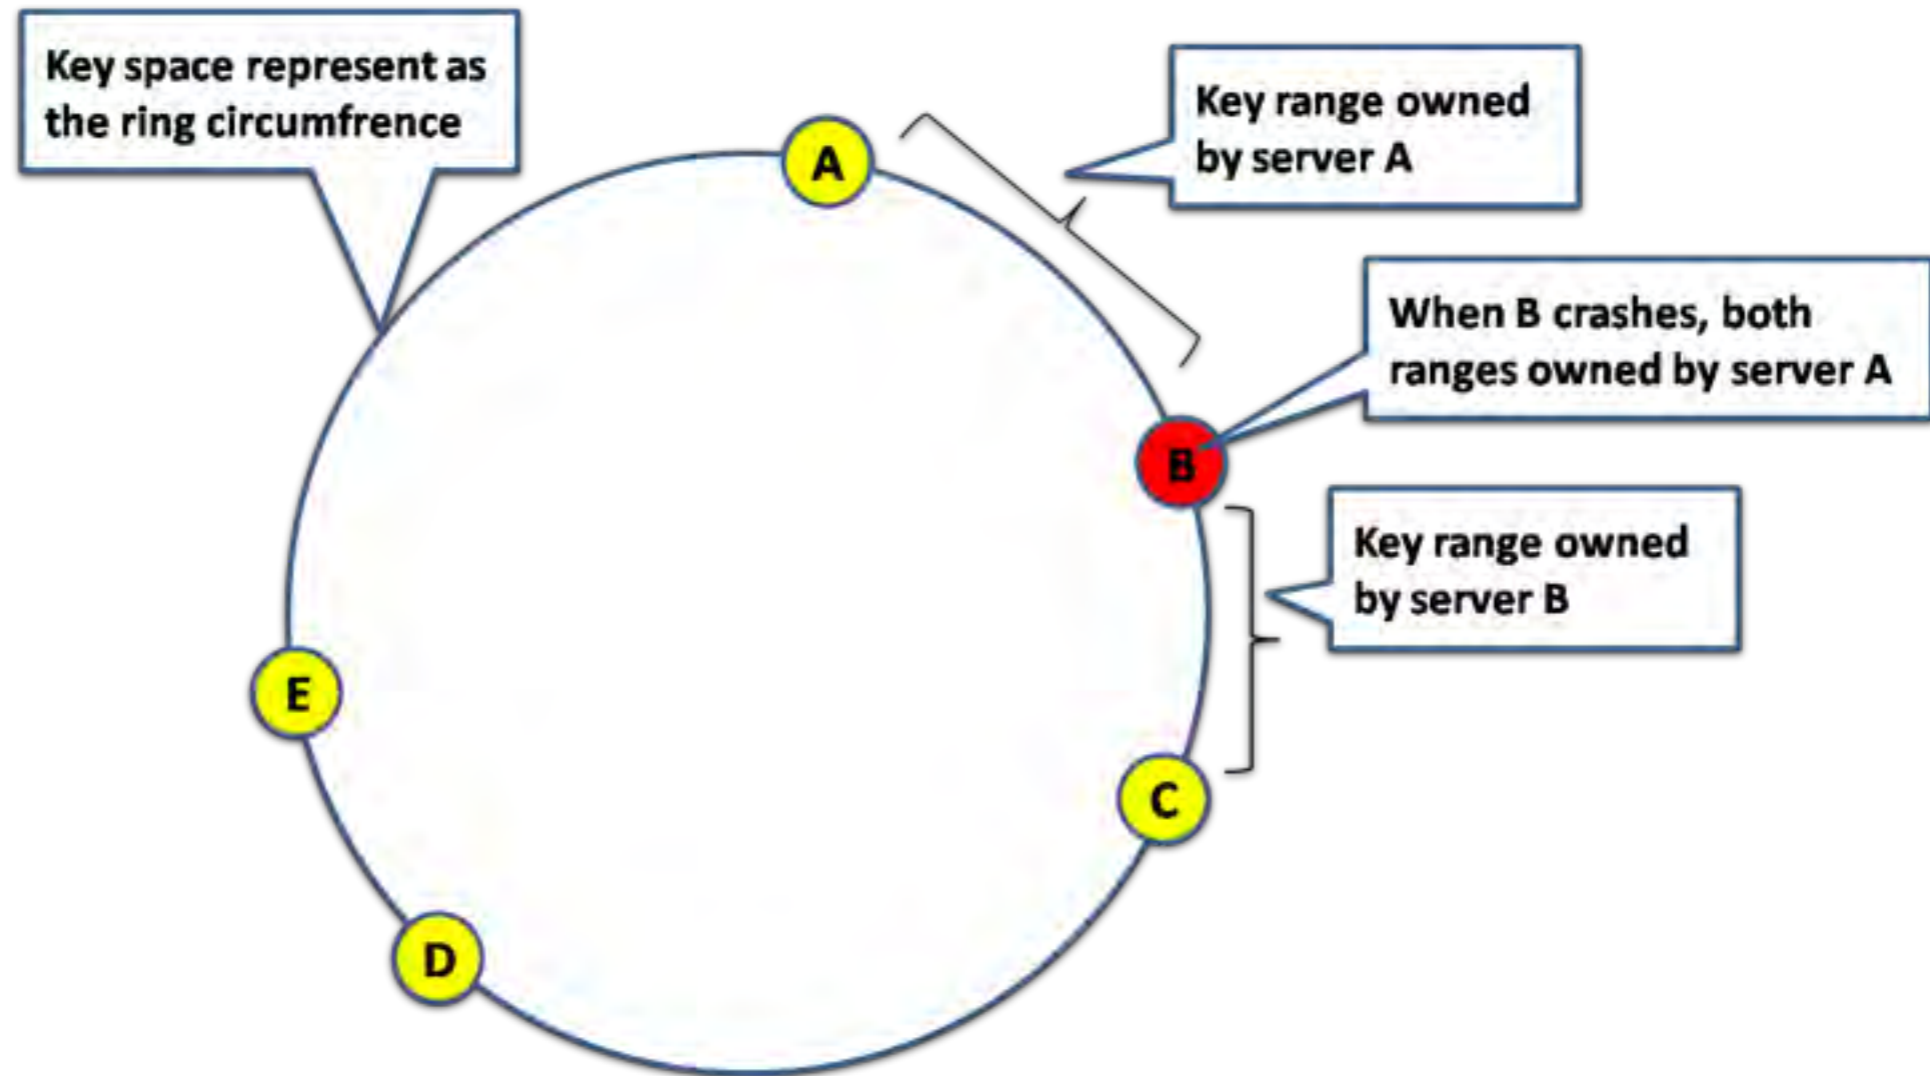
memcached



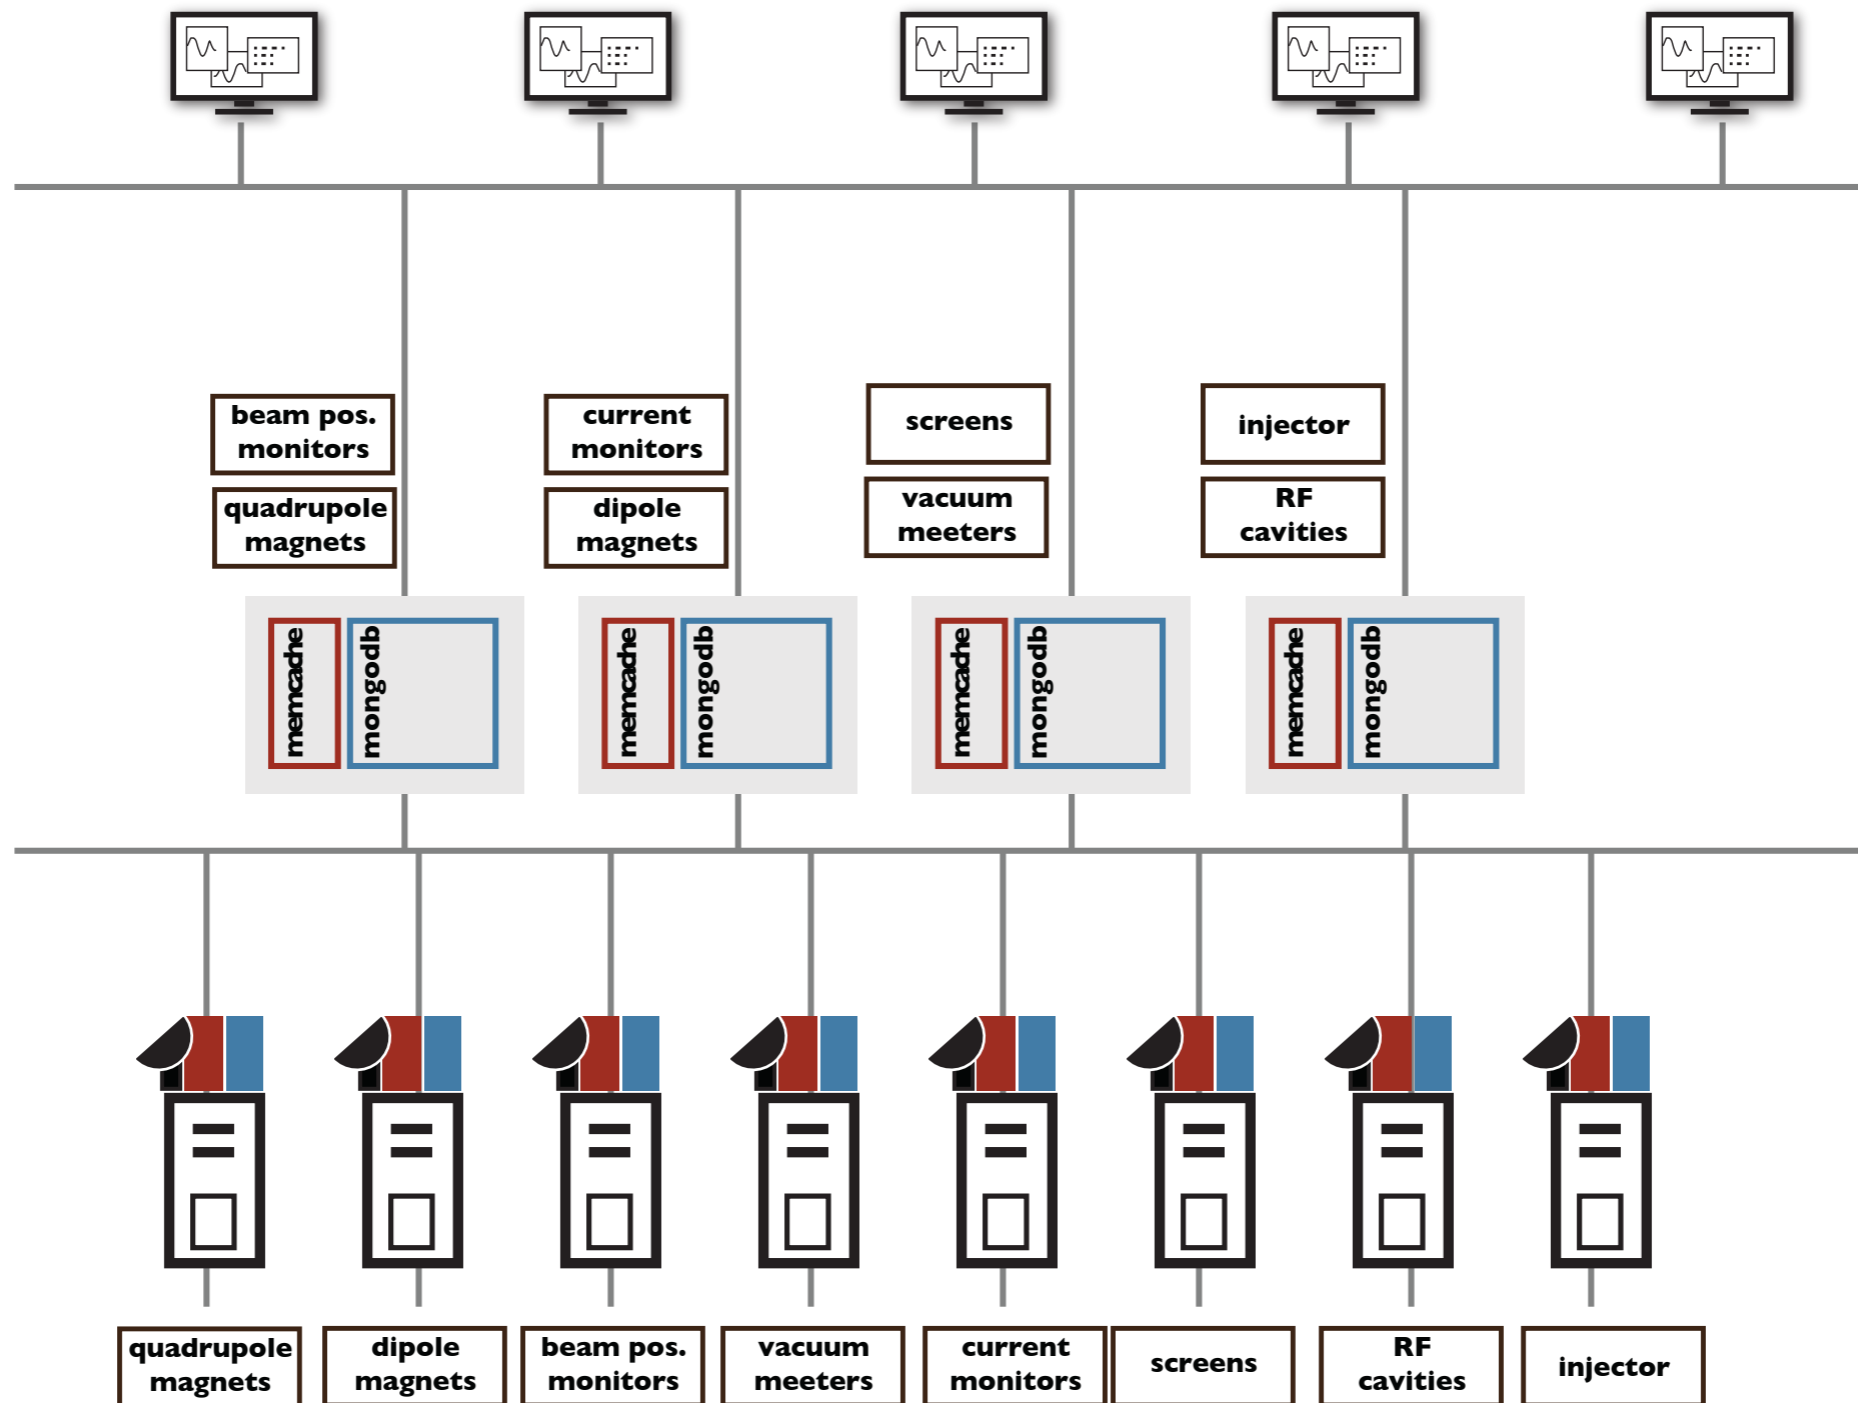
memcached



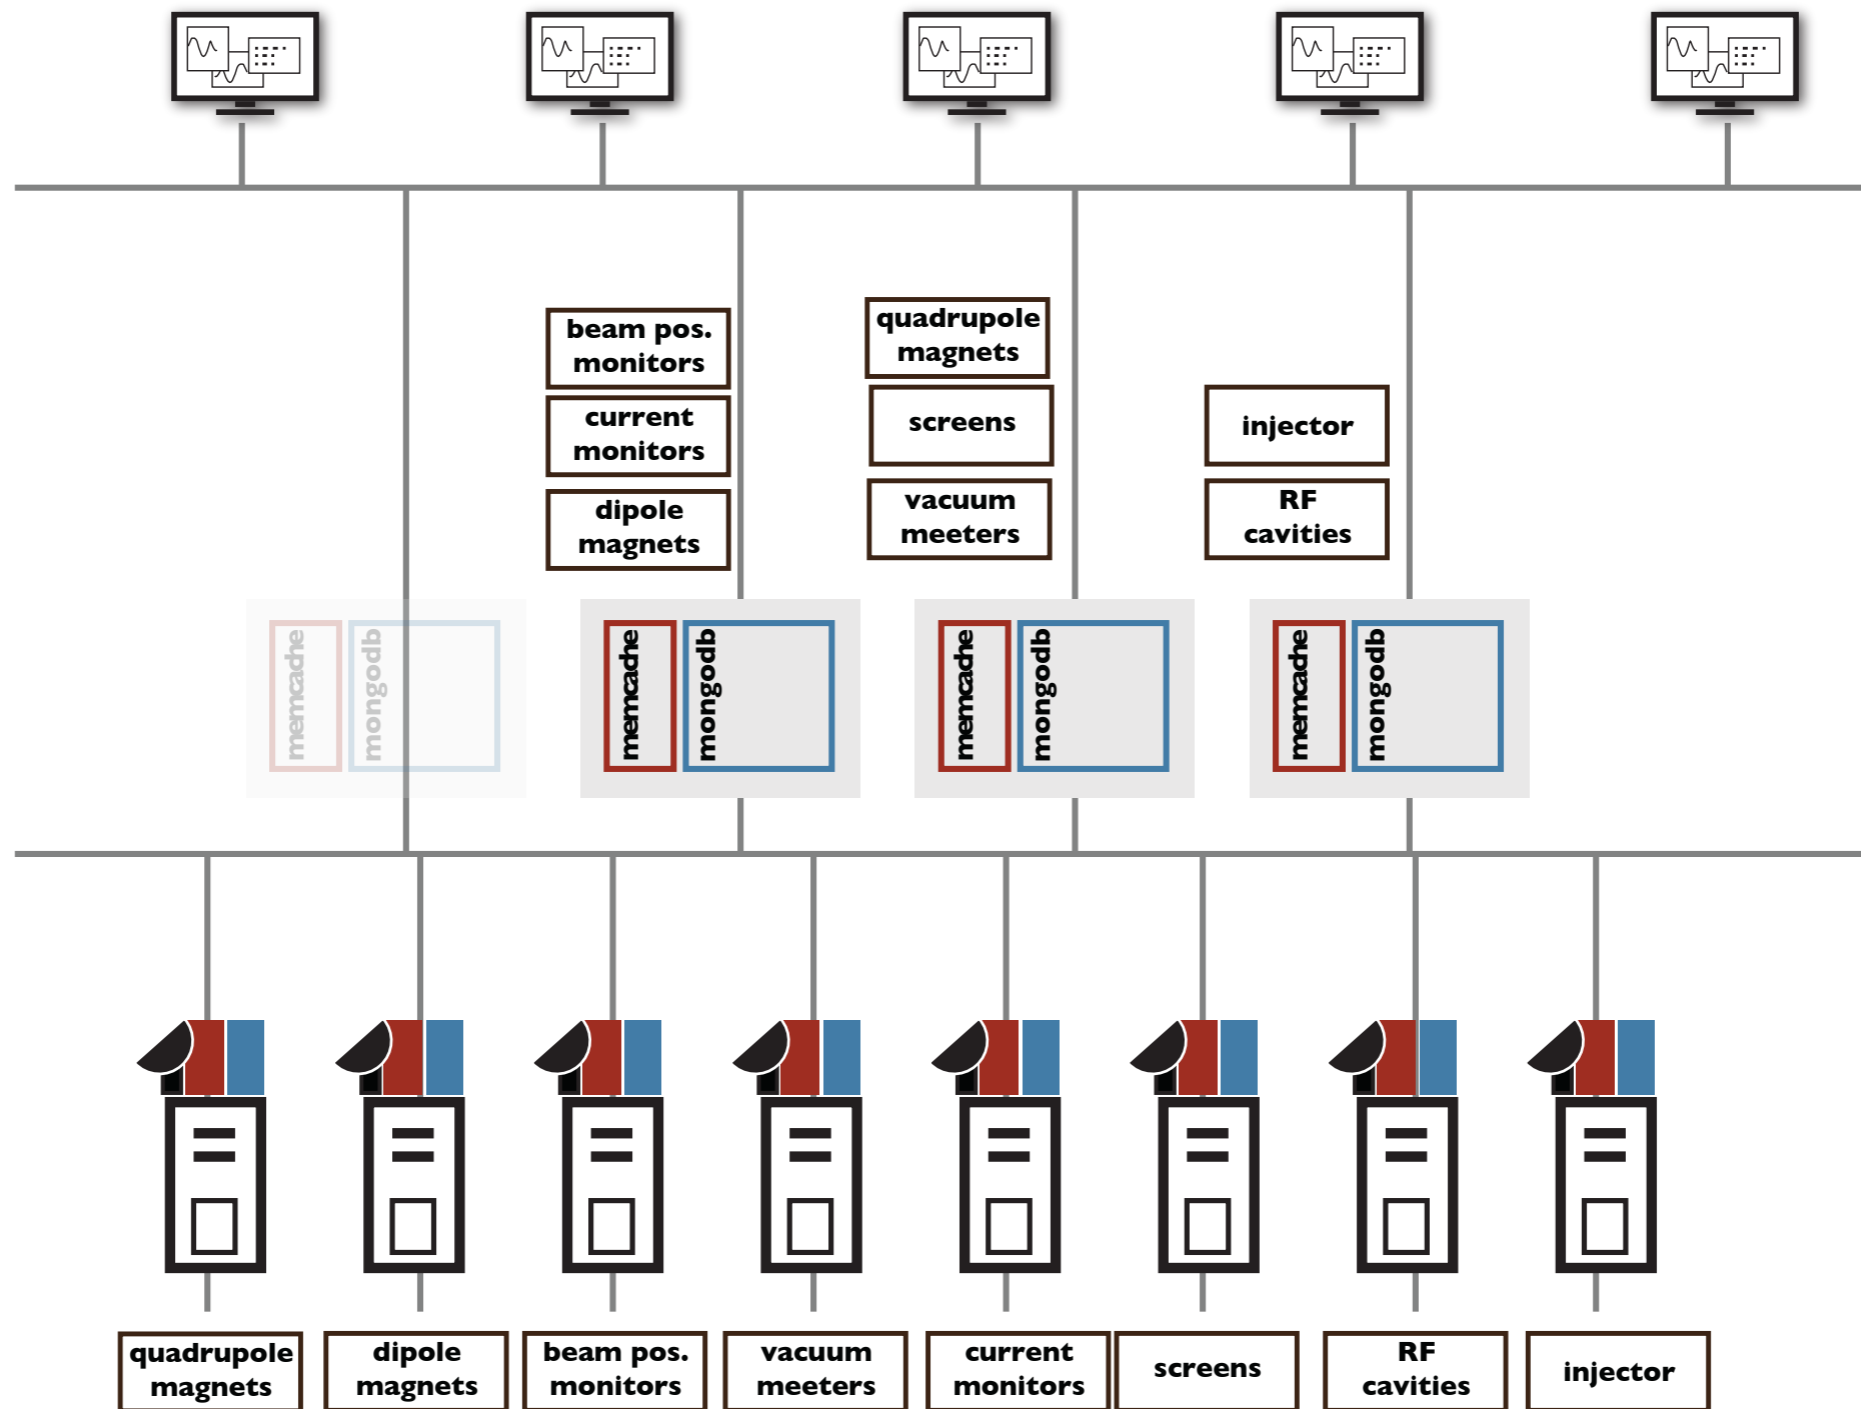
memcached



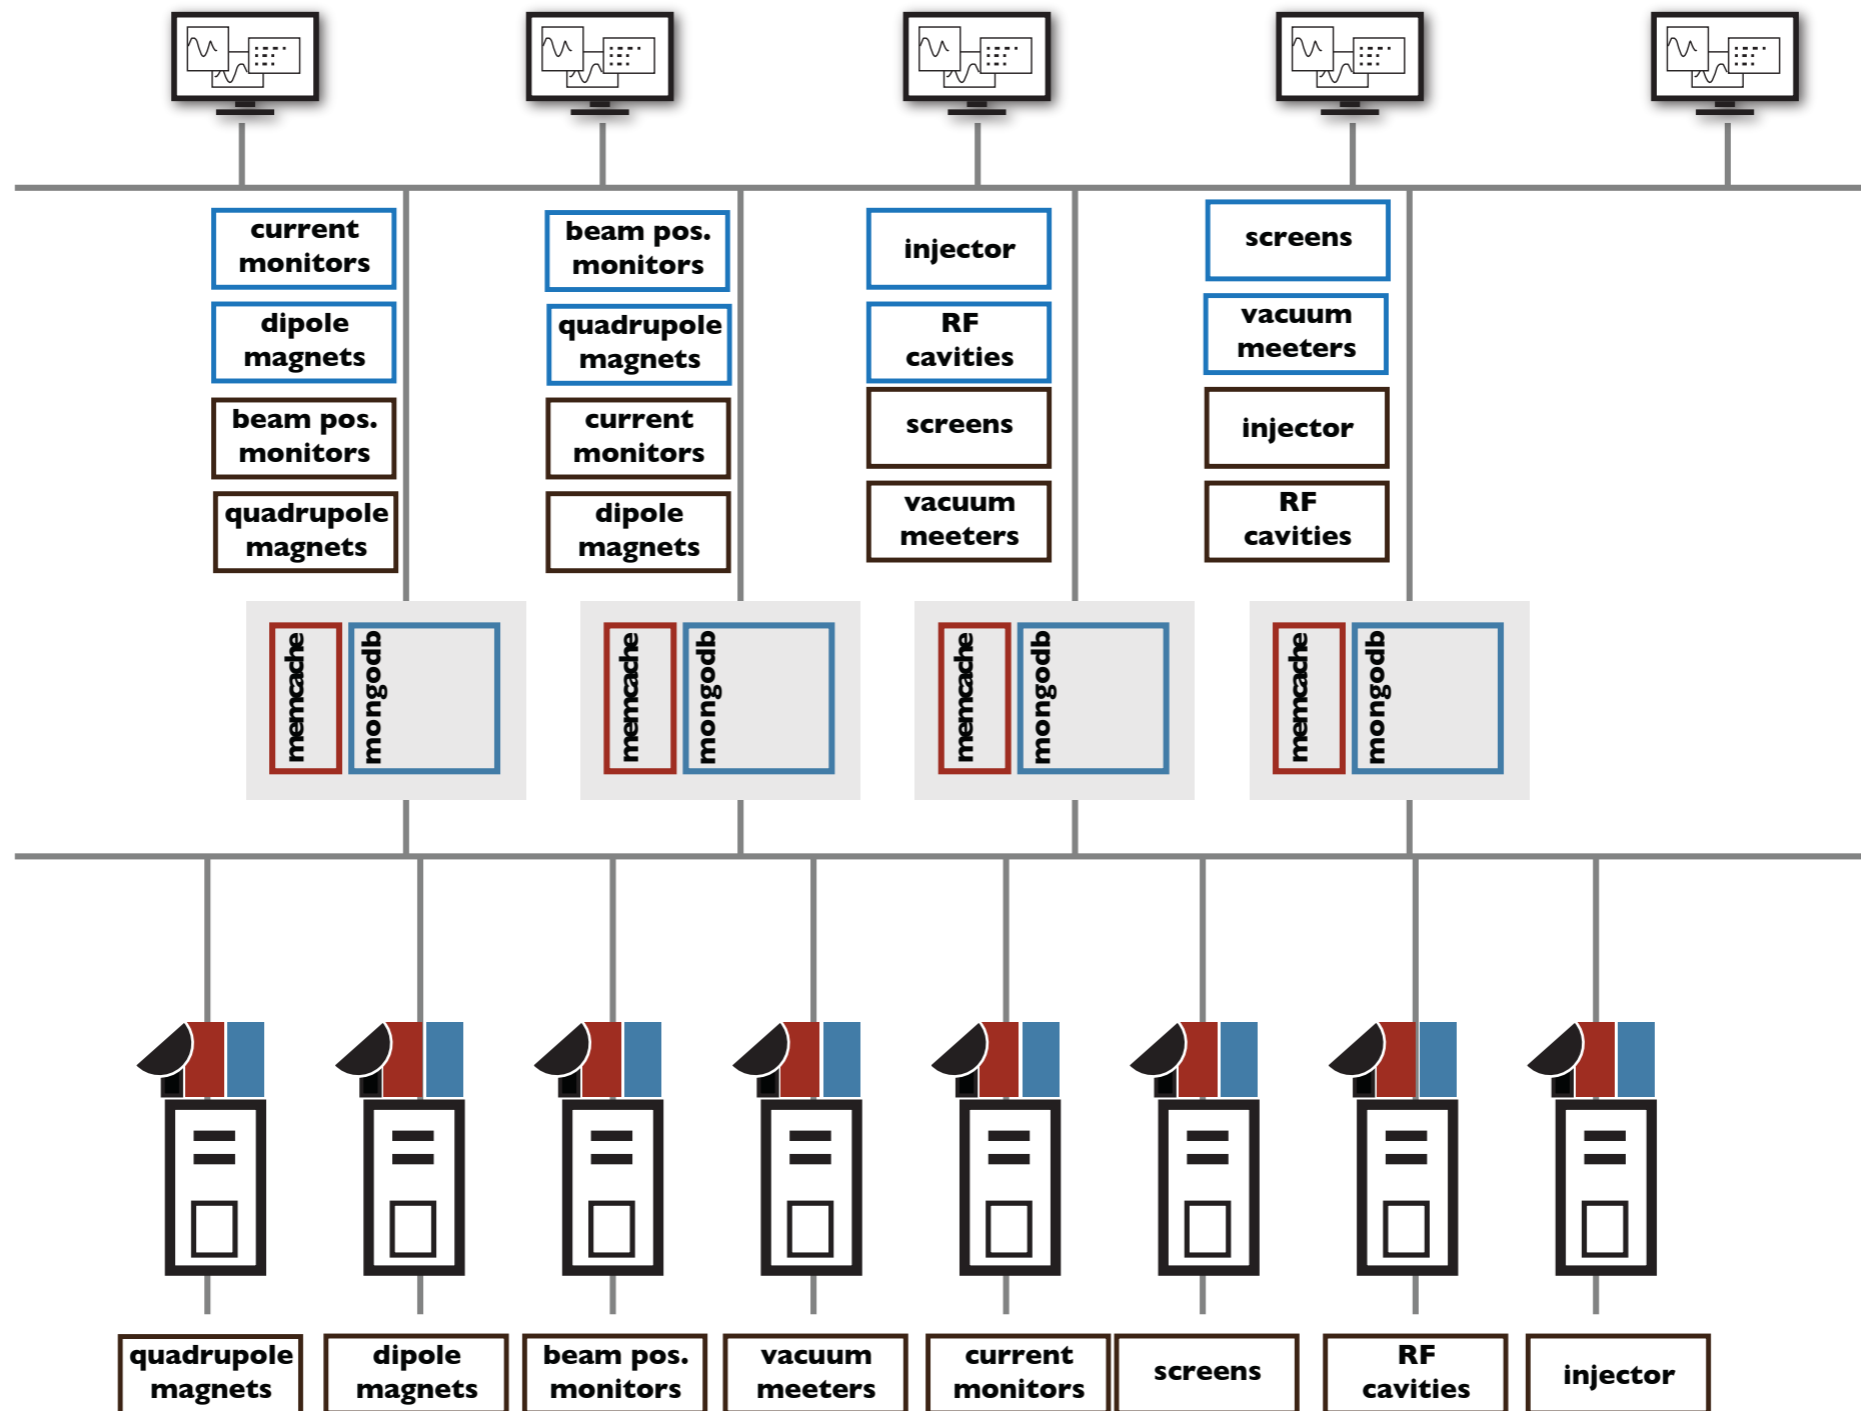
memcached used for live-data



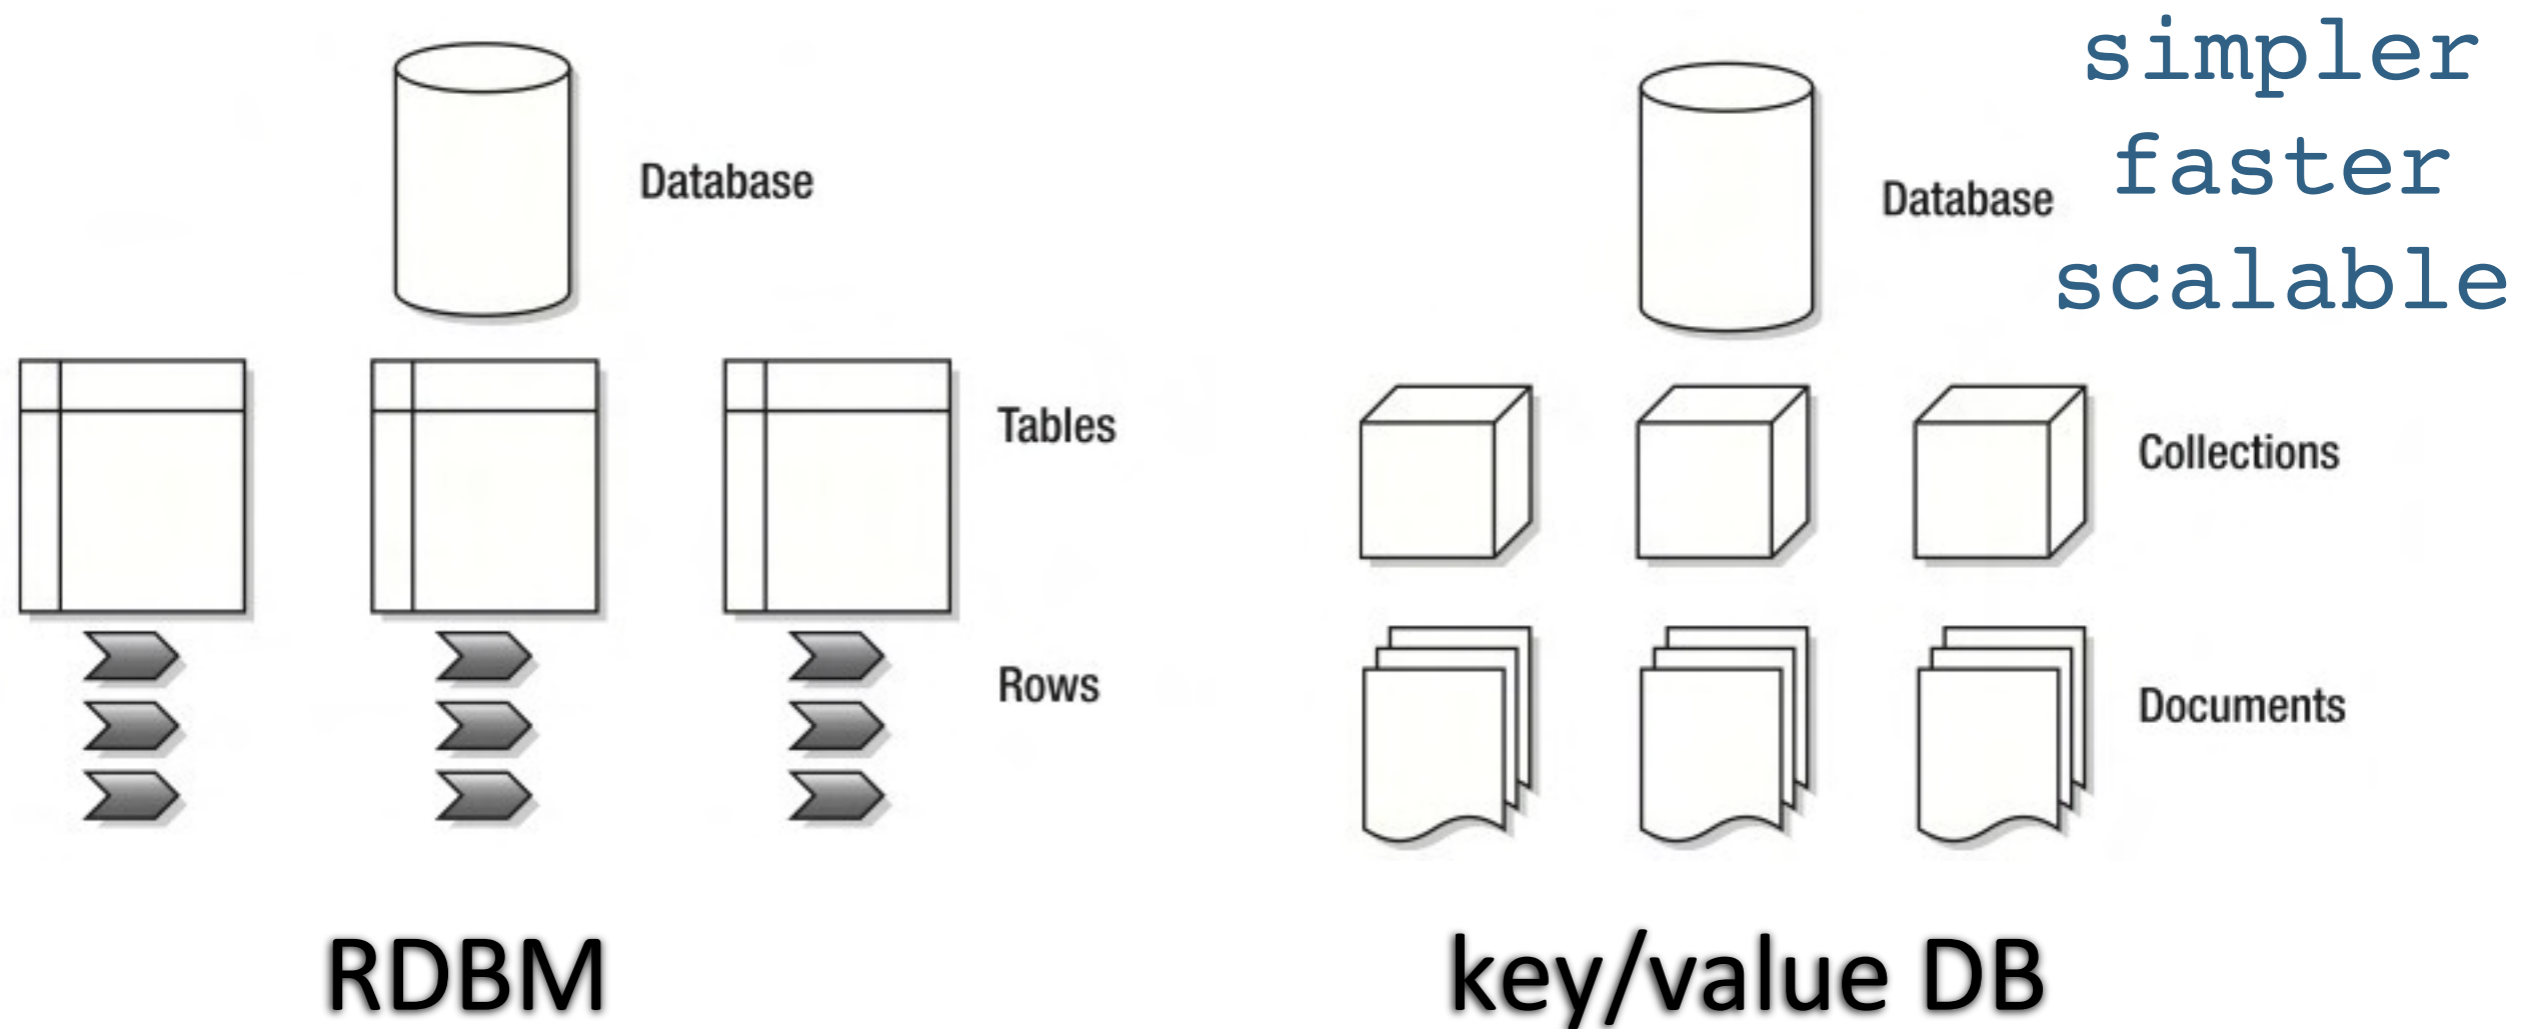
memcached - failover



memcached - failover&mirroring



No-SQL key/value DB (ex. MongoDB)



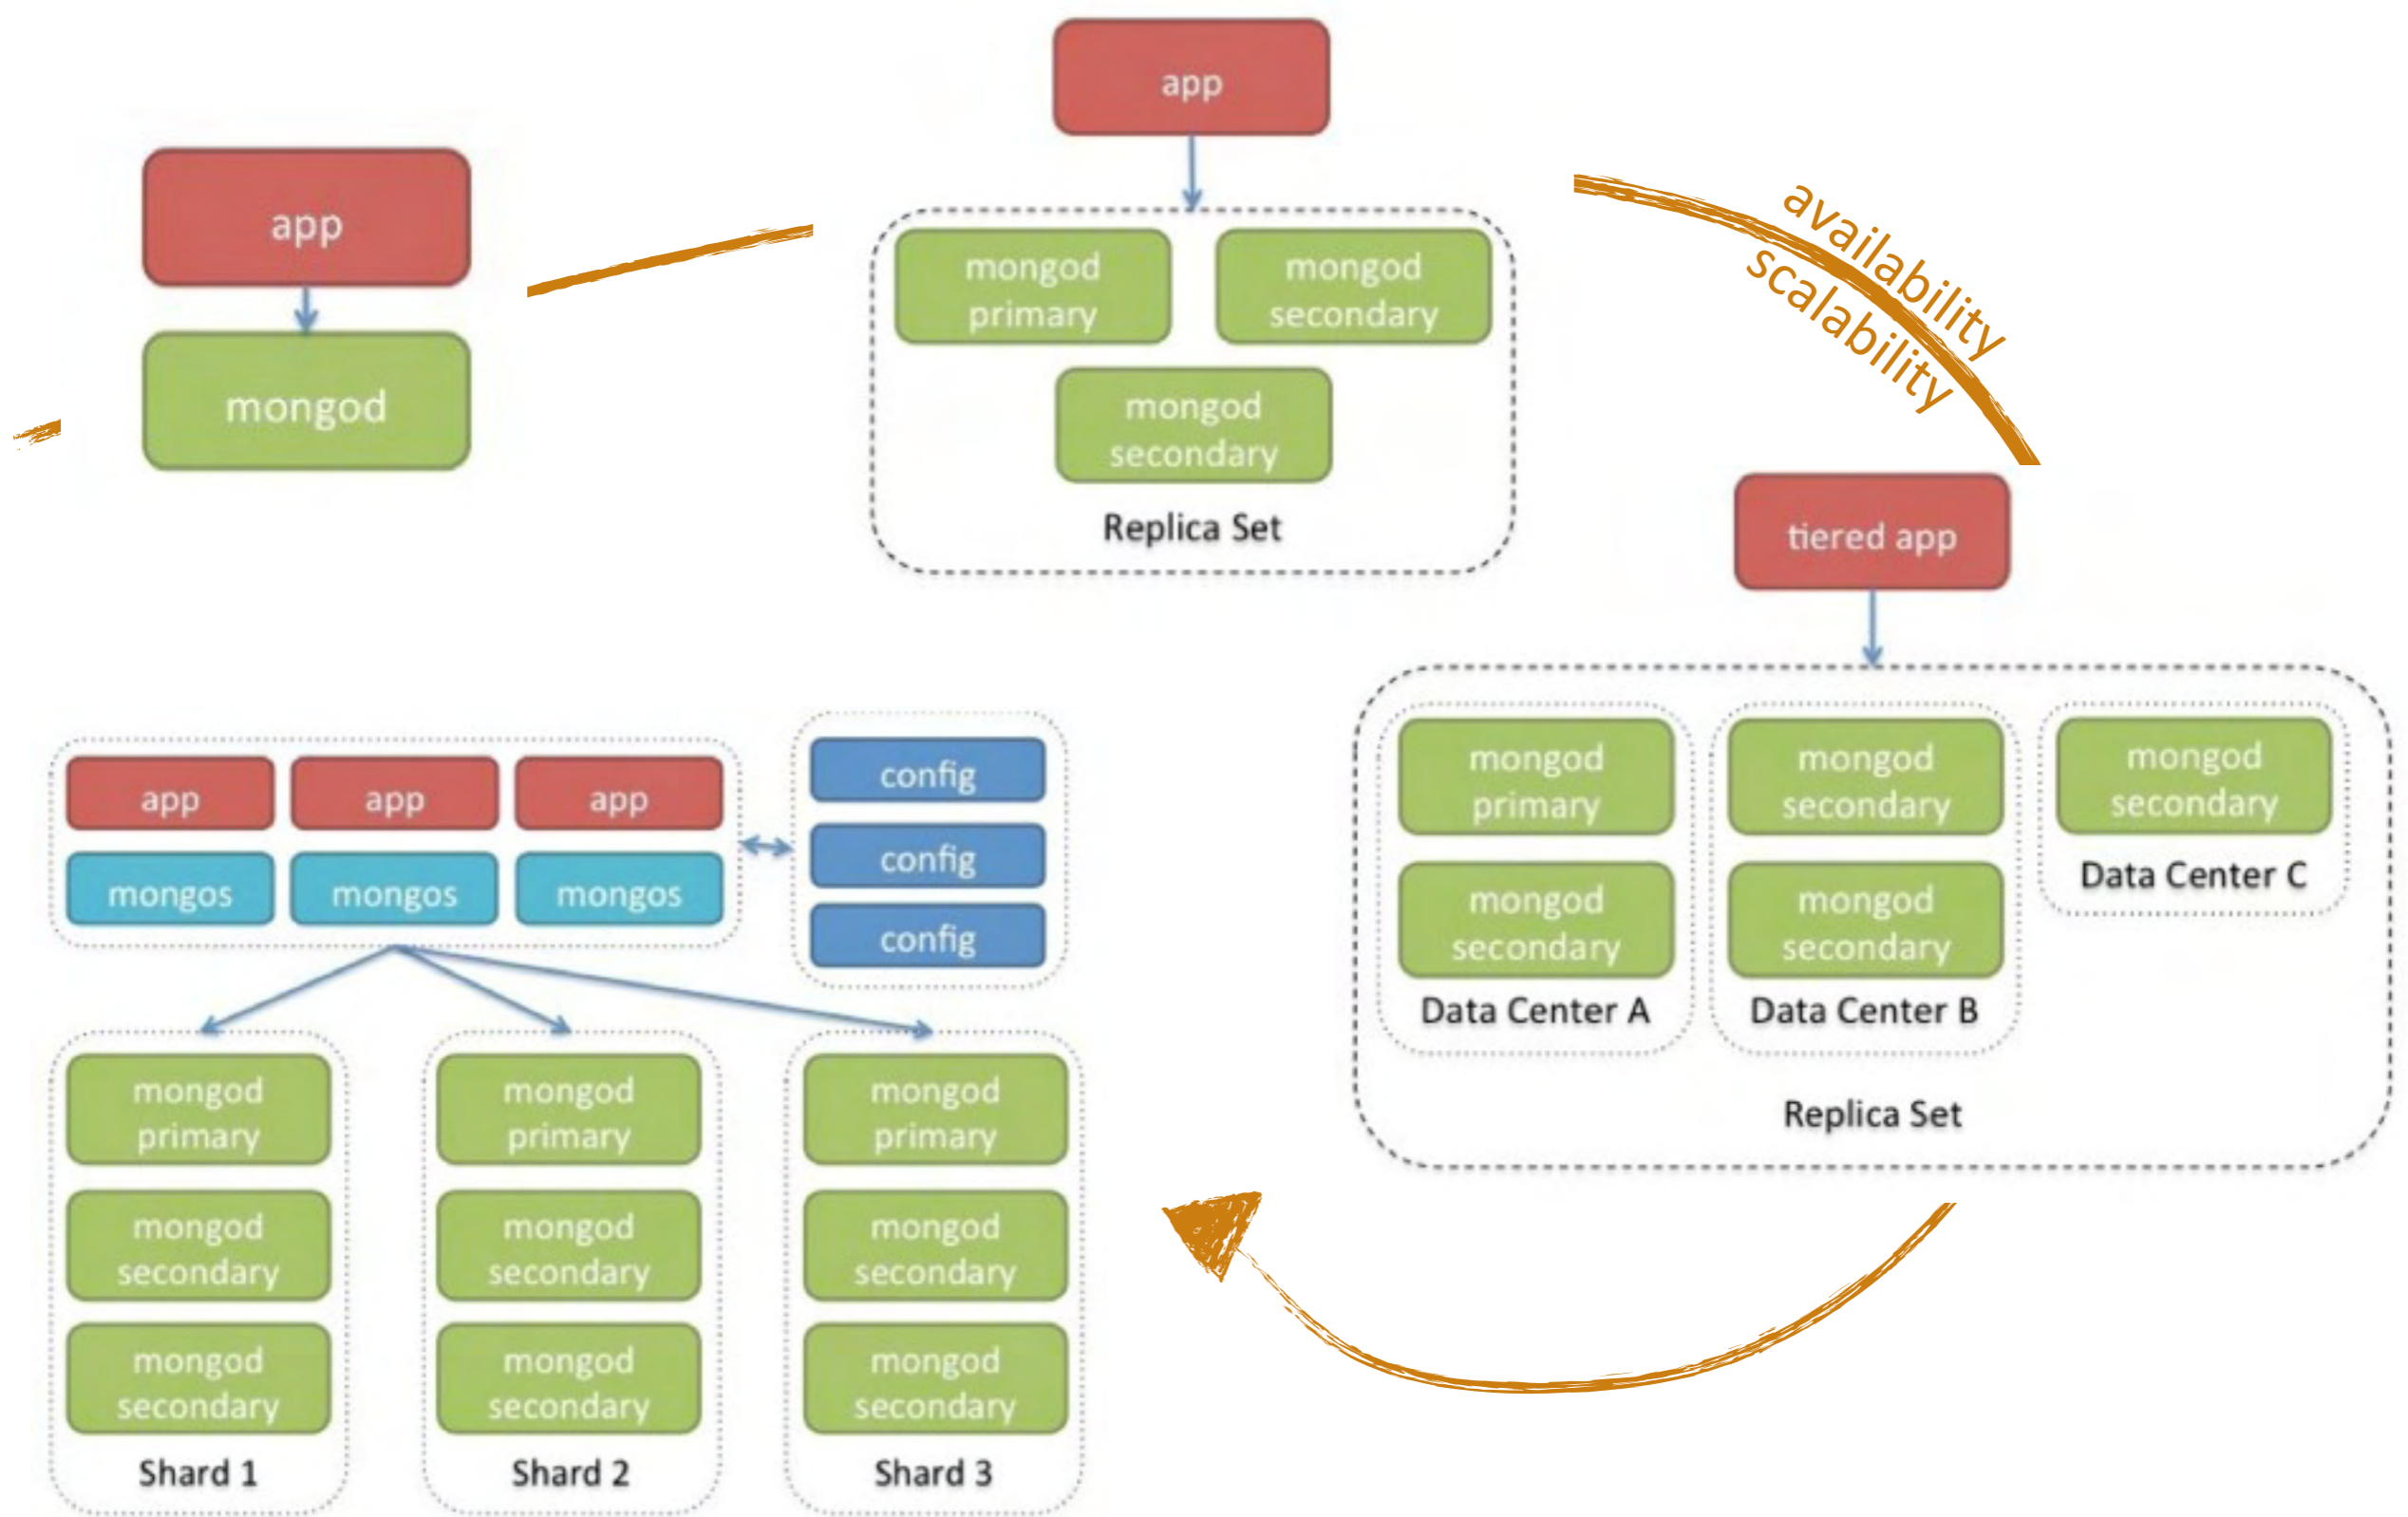
RDBM

key/value DB

```
I_media
I_cds
  I_id, artist, title, genre, releasedate
I_cd_tracklists
  I_cd_id, songtitle, length
```

```
I_media
I_items
  I_<document>
```


No-SQL key/value DB (ex. MongoDB)



No-SQL key/value DB (ex. MongoDB)

CSV

```
Surname, Forename, Phone Number  
Membrey, Peter, +852 1234 5678  
Thielen, Wouter, +81 1234 5678
```

key/value data storage:
JSON serialization

JSON

```
{  
  "forename": "Peter",  
  "surname": "Membrey",  
  "phone_numbers": [  
    "+852 1234 5678",  
    "+44 1234 565 555"  
  ]  
}
```

```
{  
  "forename": "Peter",  
  "surname": "Membrey",  
  "numbers": [  
    {  
      "phone": "+852 1234 5678"  
    }, {  
      "fax": "+44 1234 565 555"  
    }  
  ]  
}
```

BSON

- **aim to be binary version of JSON**
- **key - value syntax**
- **every pair is serialized bringing:**
 - **key name**
 - **type and dimension**
 - **value**

BSON


BSON to encode

```
{“key1” : “hello”  
  “key2” : 125}
```



binary representation

```
\x1B\x00\x00\x00  
\x02key 1\x00\x06\x00\x00\x00hello\x00  
\x10key2\x00\x7D\x00\x00\x00  
\x00
```



BSON Total Length

BSON


BSON to encode

```
{“key1” : “hello”  
 “key2” : 125}
```



binary representation

```
\x1B\x00\x00\x00  
\x02key1\x00\x06\x00\x00\x00hello\x00  
\x10key2\x00\x7D\x00\x00\x00  
\x00
```




- type string,
- key “key1”
- dimension and value of the string hello

BSON

BSON to encode

```
{“key1” : “hello”  
 “key2” : 125}
```



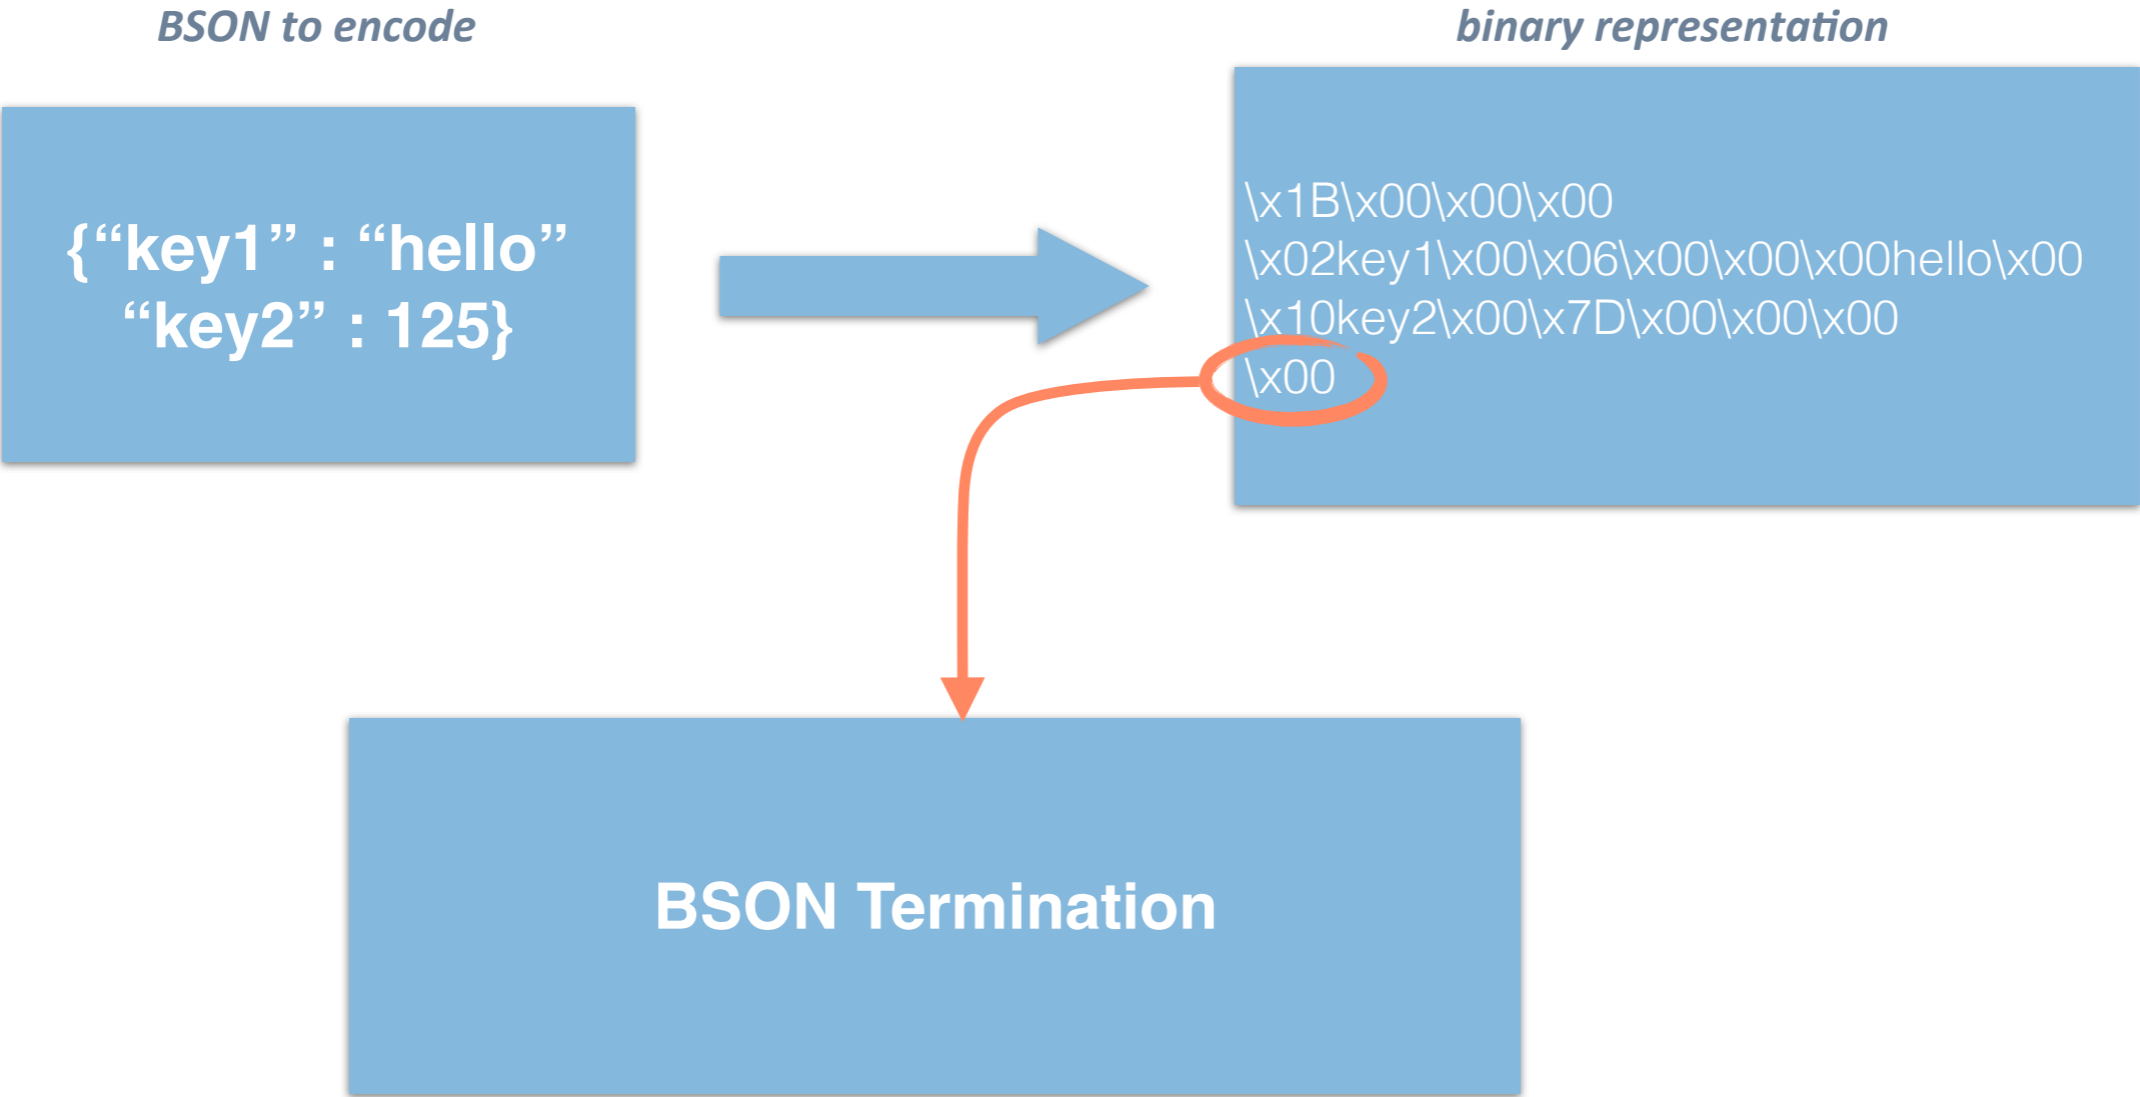
binary representation

```
\x1B\x00\x00\x00  
\x02key1\x00\x06\x00\x00\x00hello\x00  
\x10key2\x00\x7D\x00\x00\x00  
\x00
```



- type int32,
- key “key2”
- int32 value

BSON



BSON optimization

- a possible optimization of BSON would be separating the description (keys, structures and types) from the data (binary representations of the values)
- the serialized data must contain all information for de-serialize and calculate
 - key offset
 - total length of the BSON

BSON optimization

BSON to encode

```
{“key1” : “hello”  
  “key2” : 125}
```



```
{“key1” : cstring : dataoffset  
  “key2” : int32 : dataoffset}
```

```
{“hello”:125}
```

BSON optimization

```
{“key1” : cstring:offset  
“key2” : int32:offset}
```

```
{“hello”:125}
```

```
getString(“key1”) == “hello”  
getInt32(“key2”) == 125
```

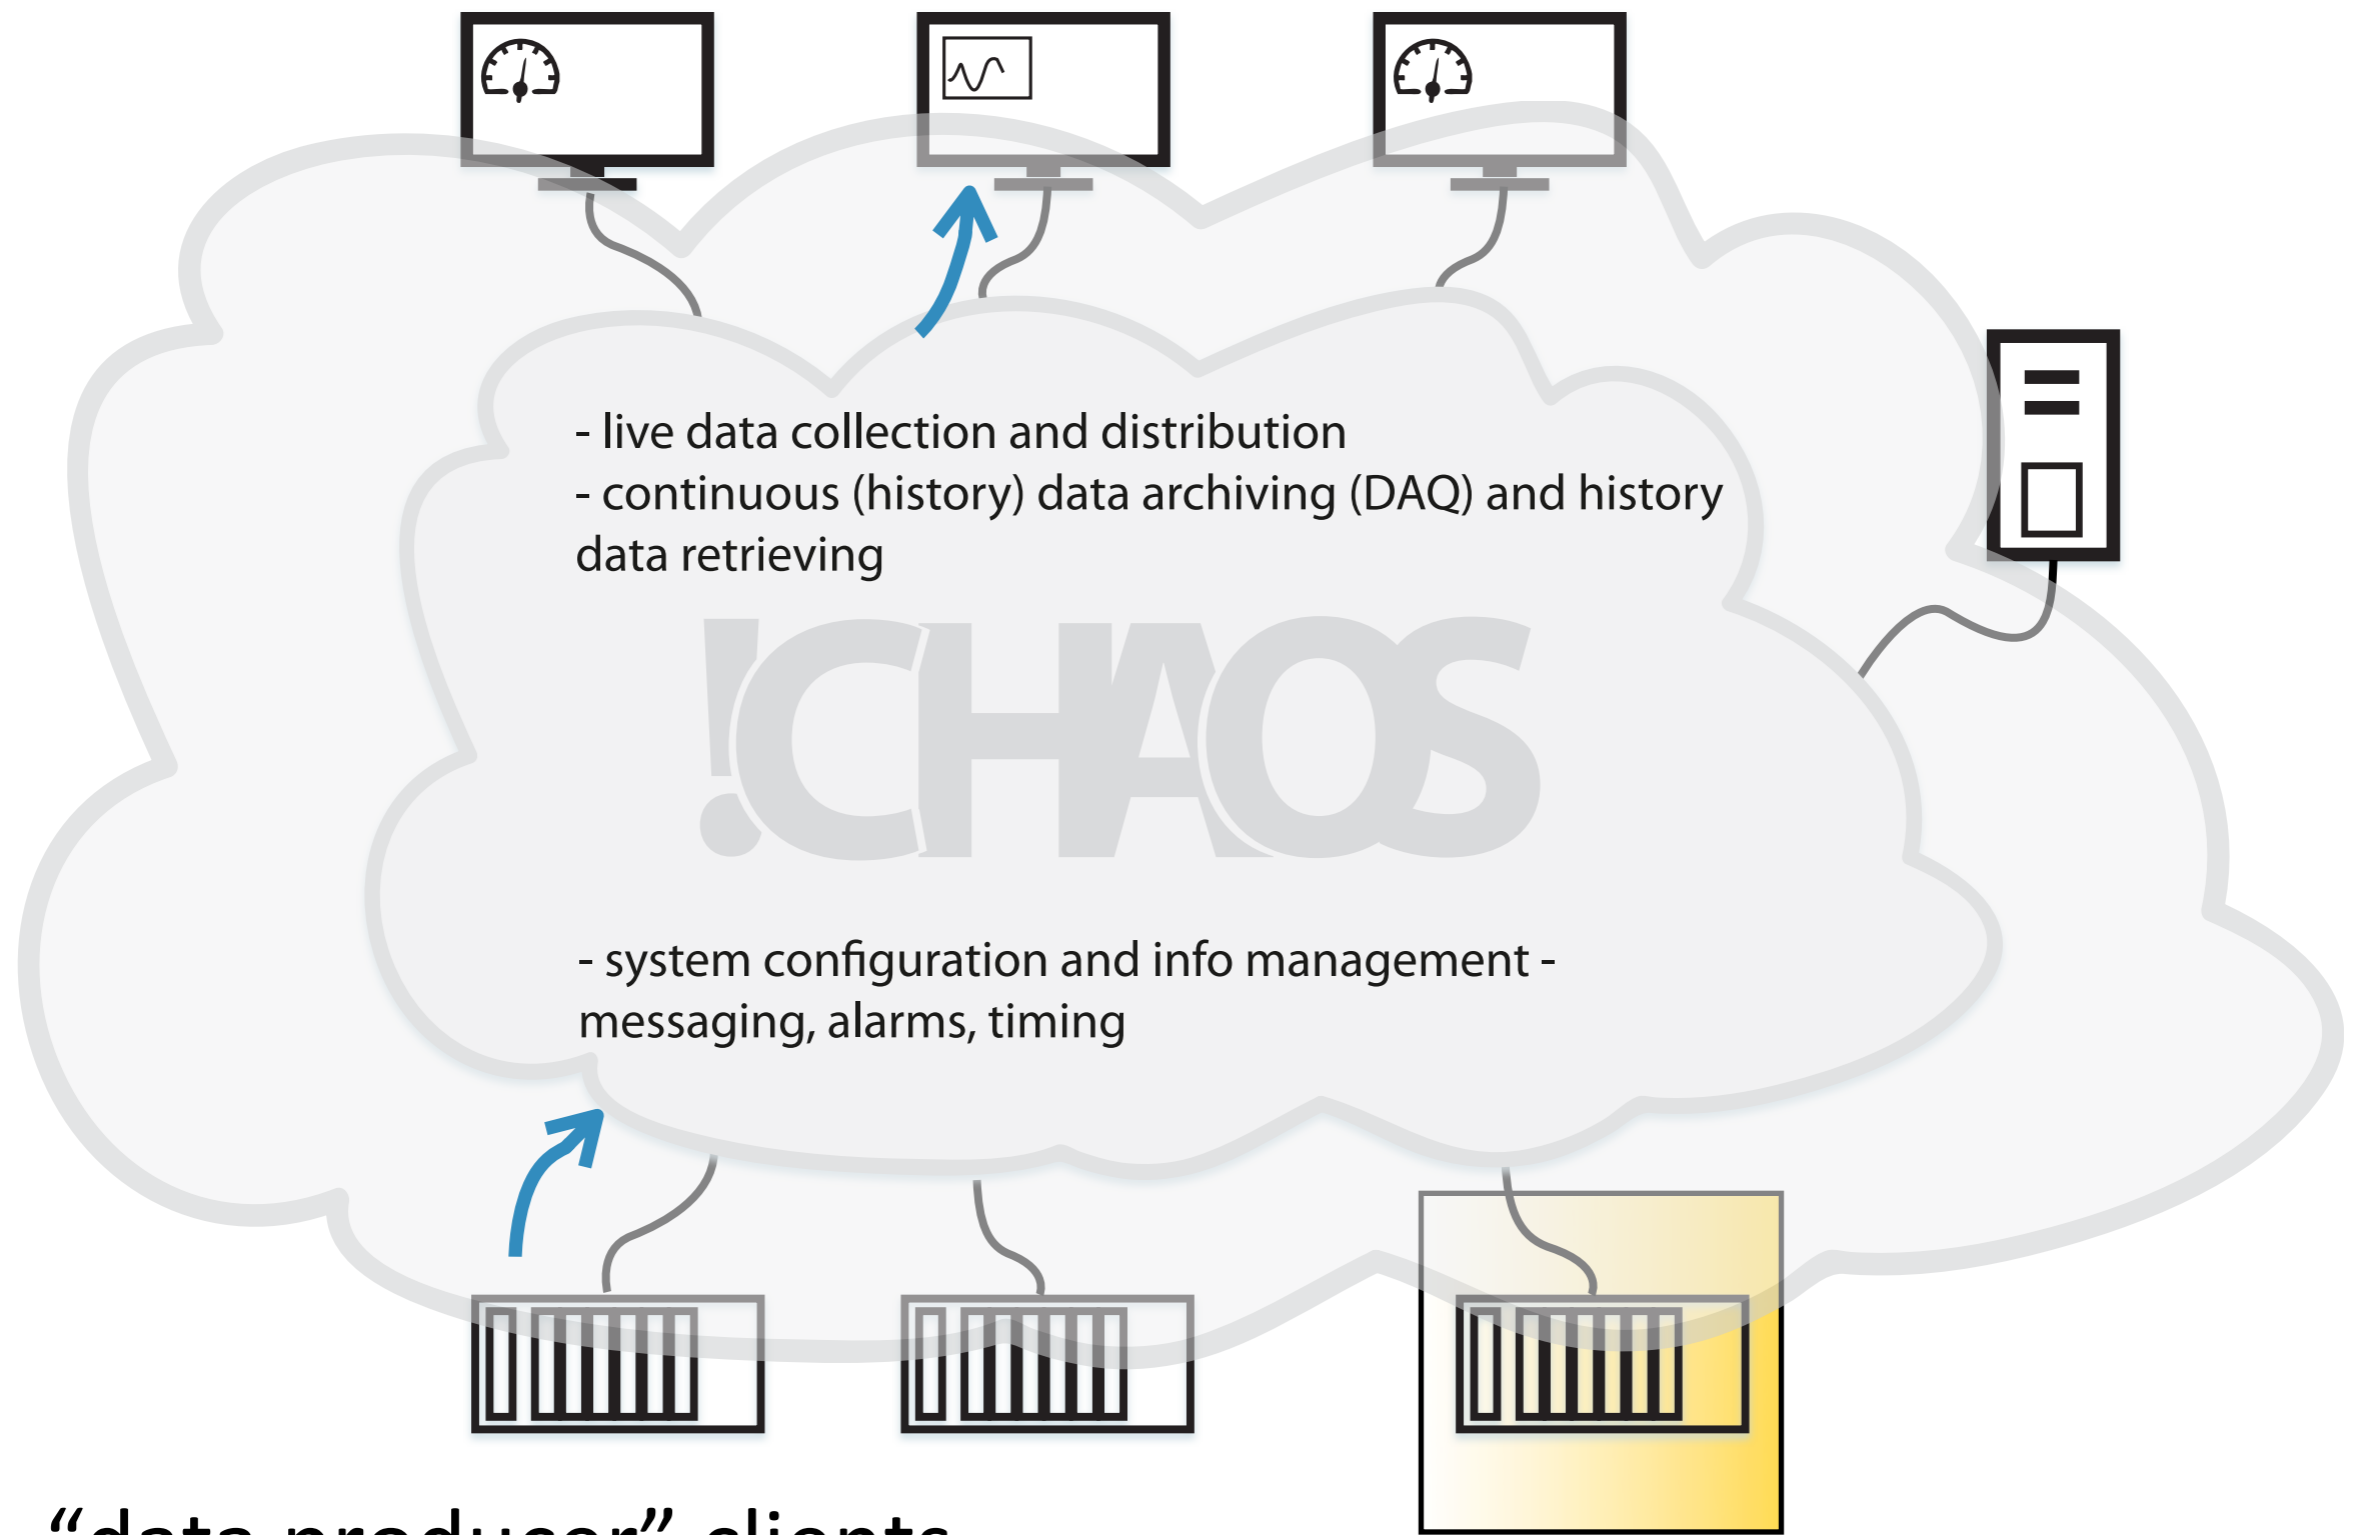
BSON optimization

```
{“key1” : cstring:offset  
“key2” : int32:offset}
```

```
{“world”:228}
```


```
getString(“key1”) == “world”  
getInt32(“key2”) == 128
```

“data consumer” clients

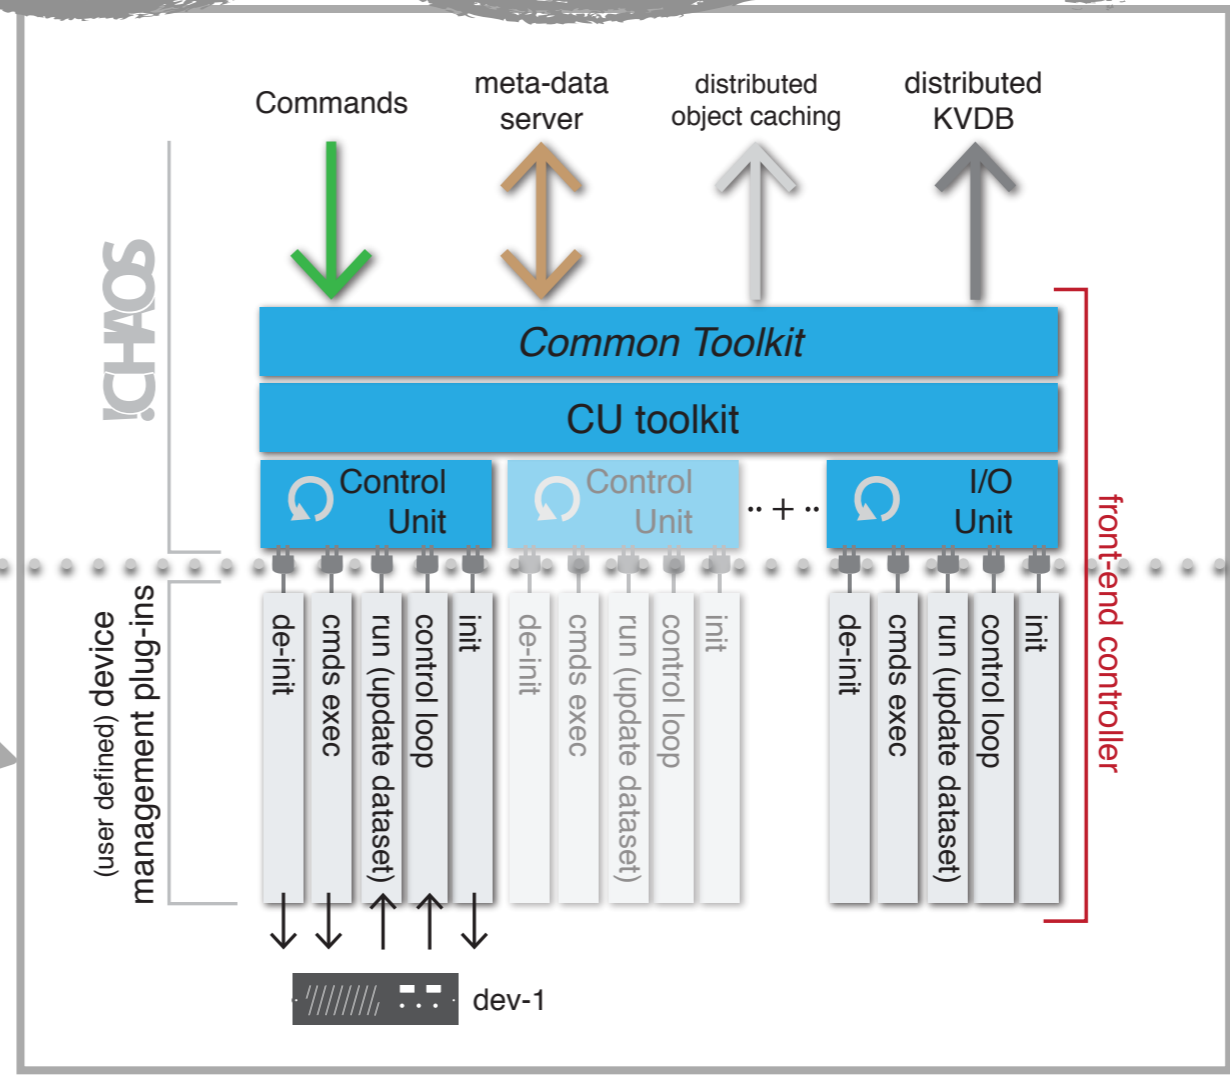


“data producer” clients

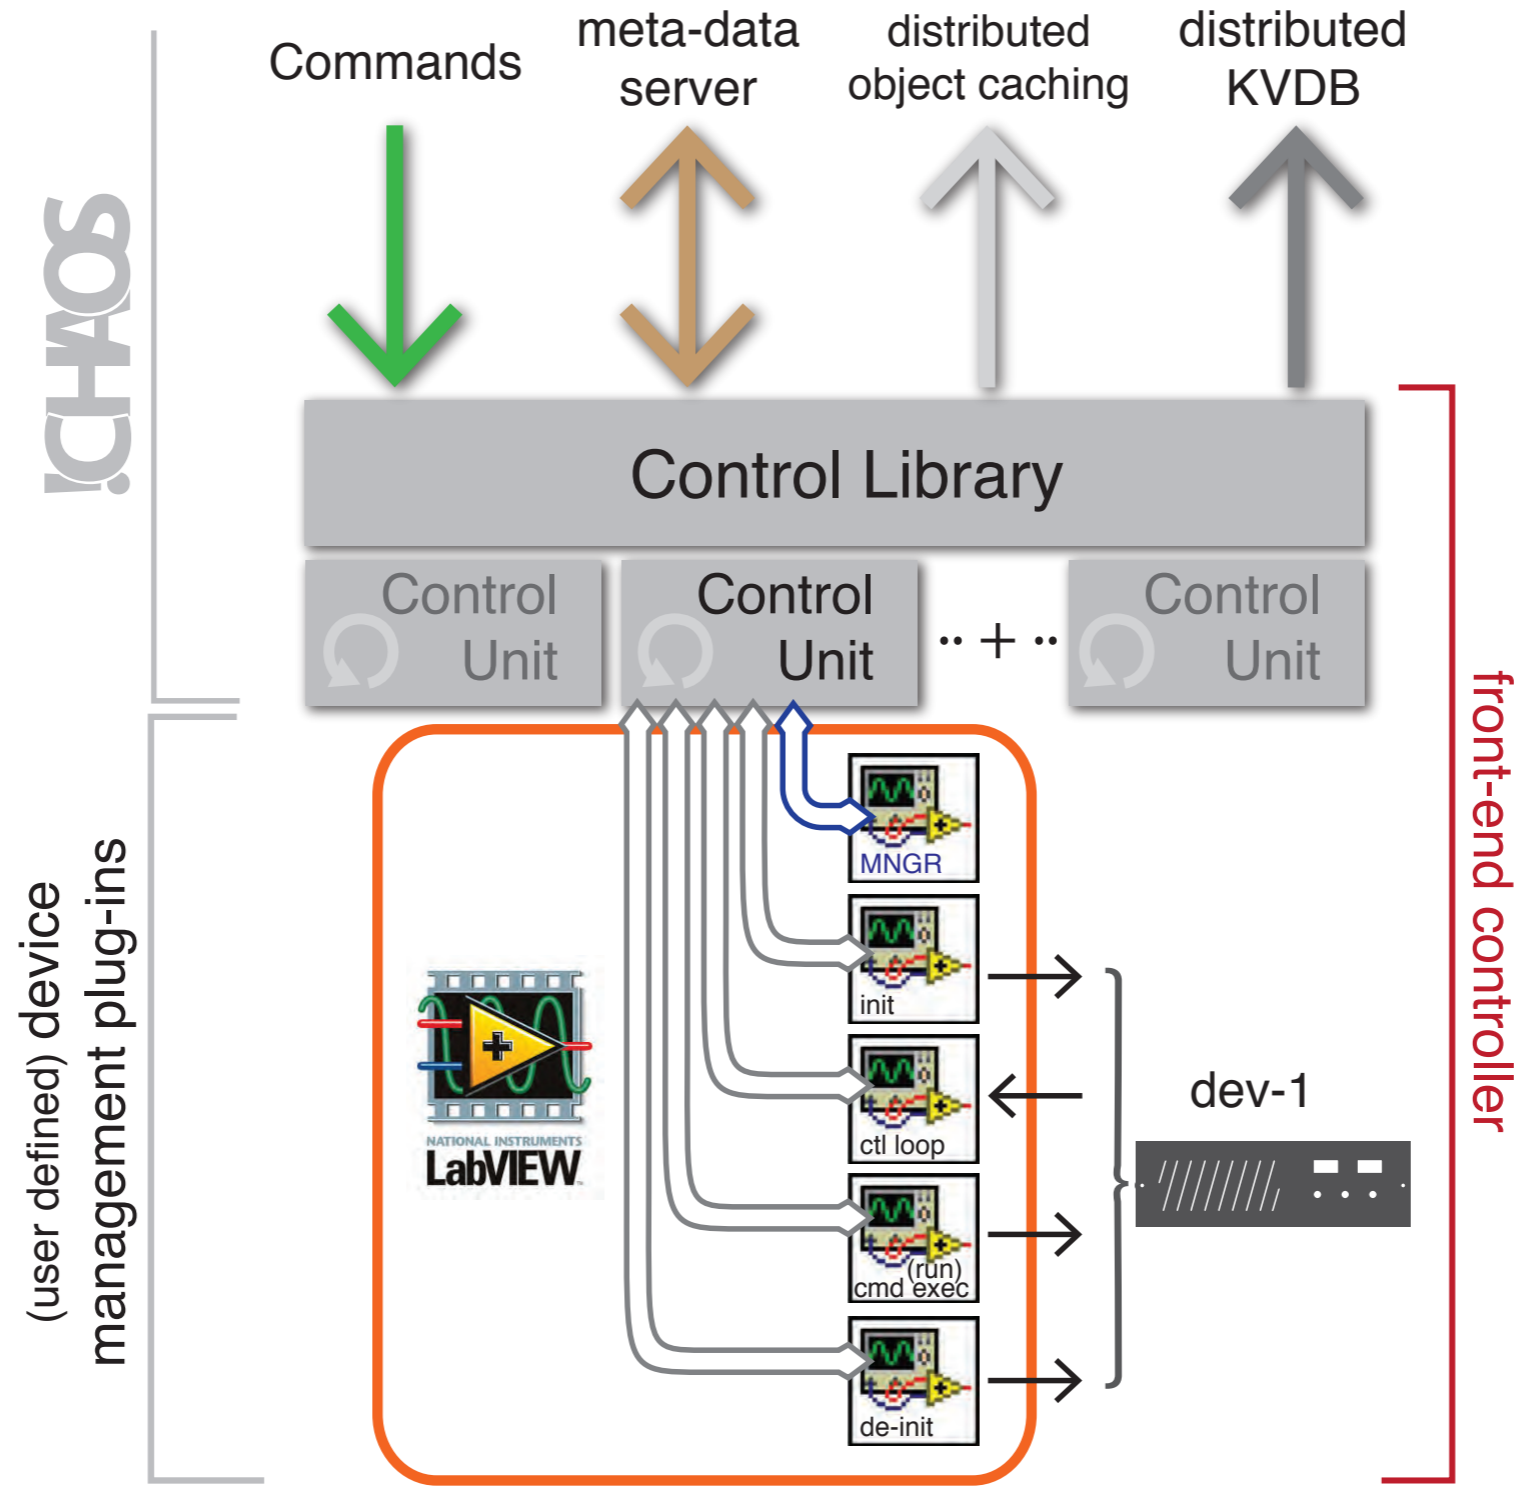
AOS



abstraction boundaries

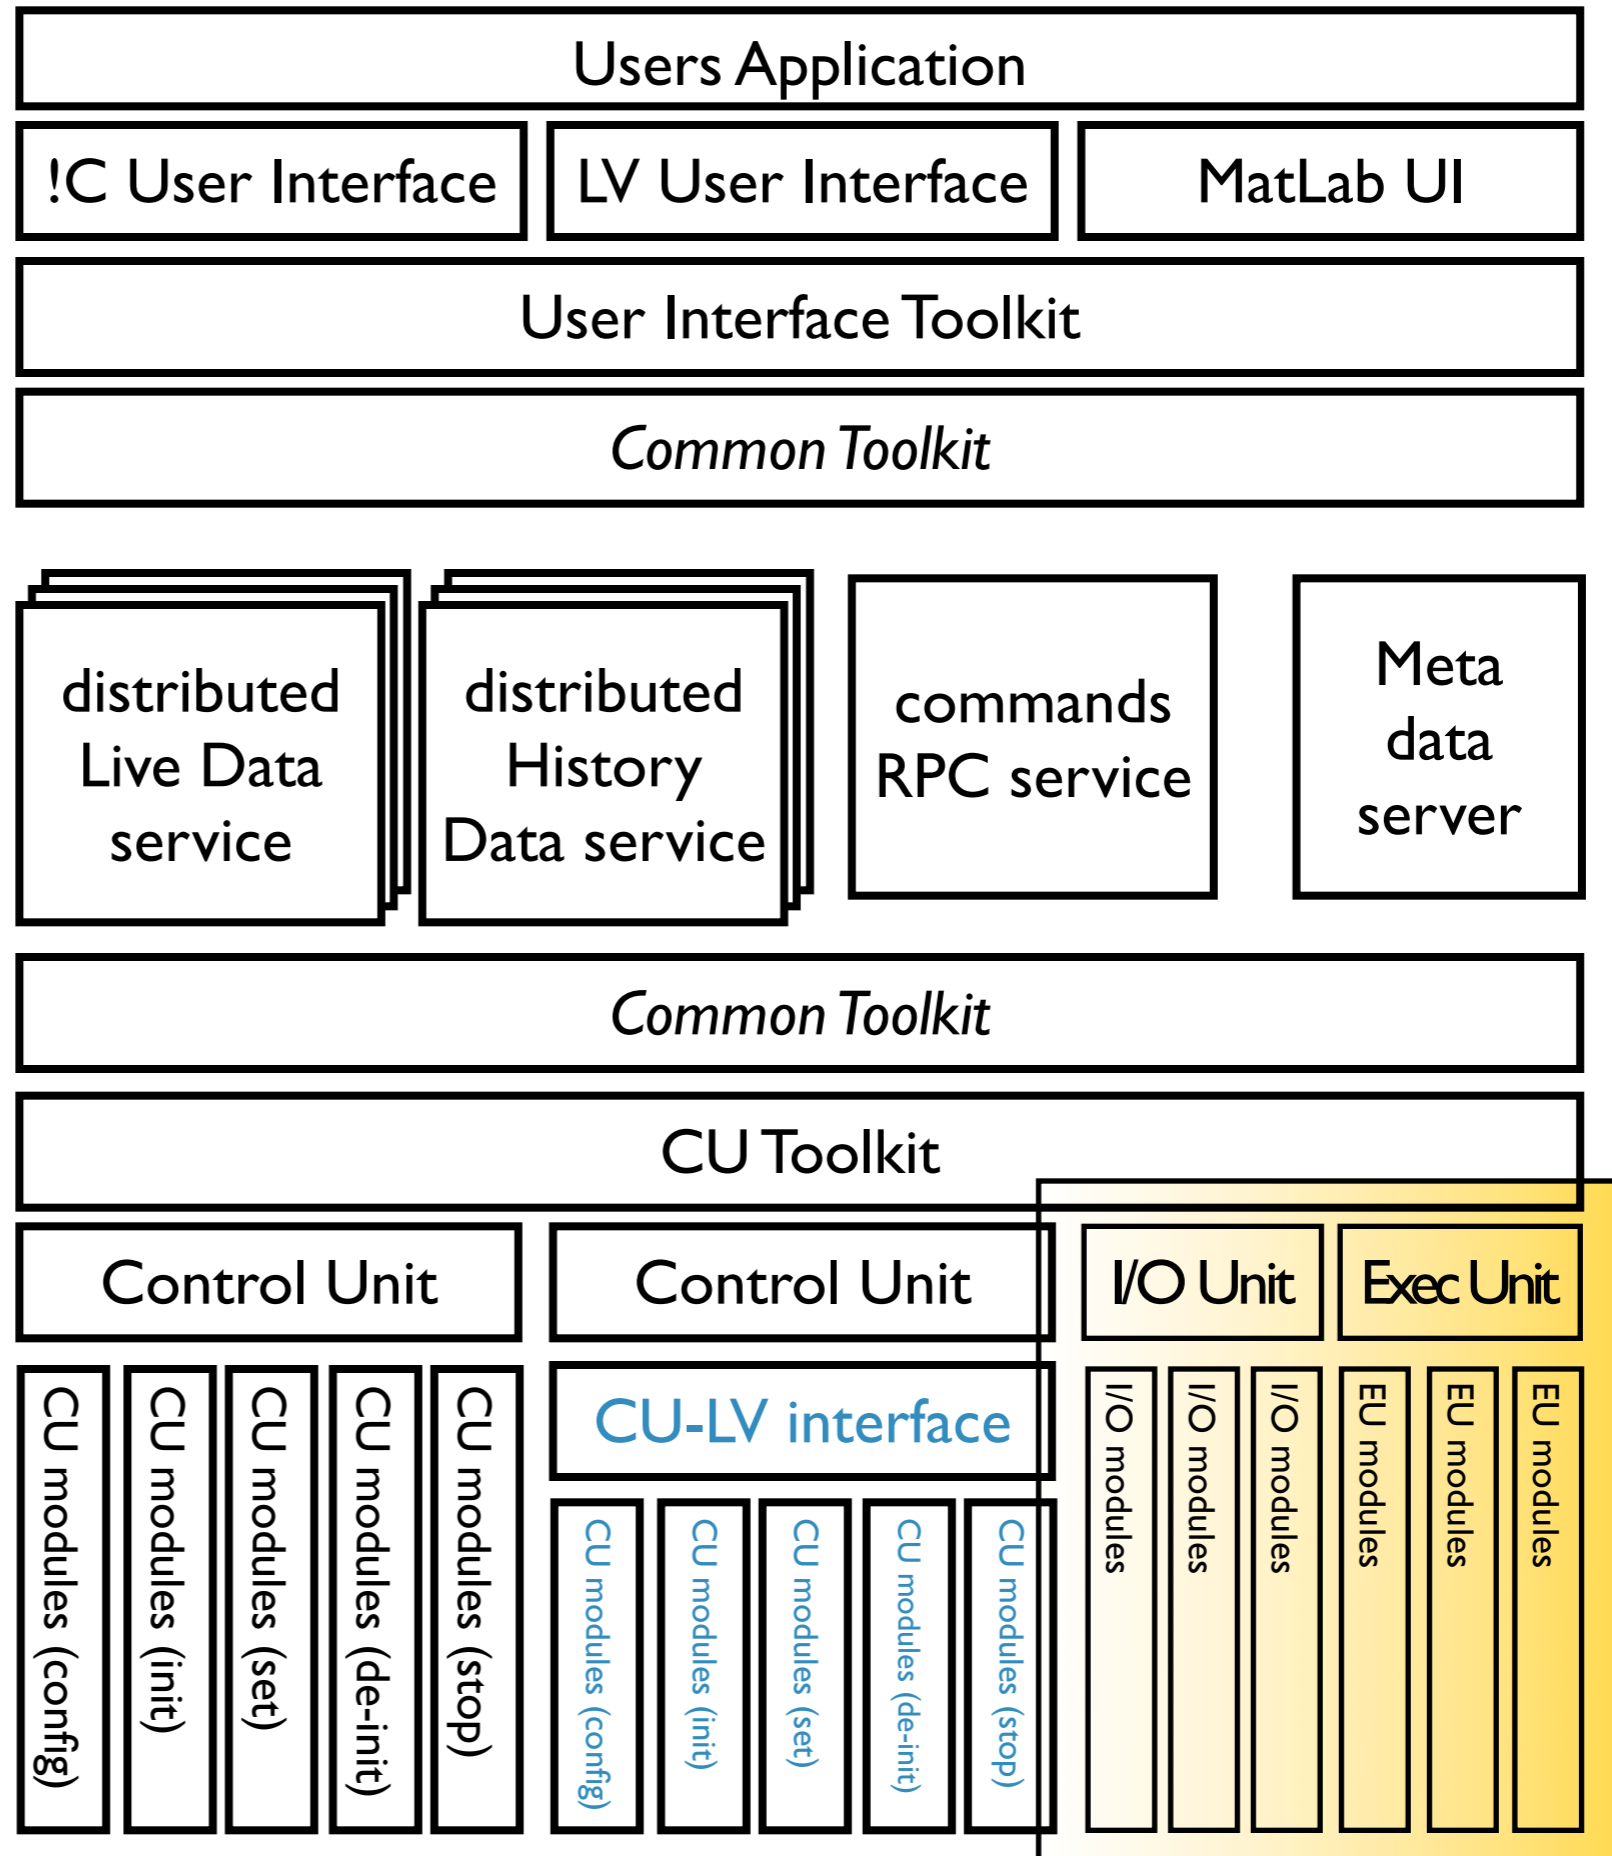


LabVIEW-CHAOS integration (a very basic solution)





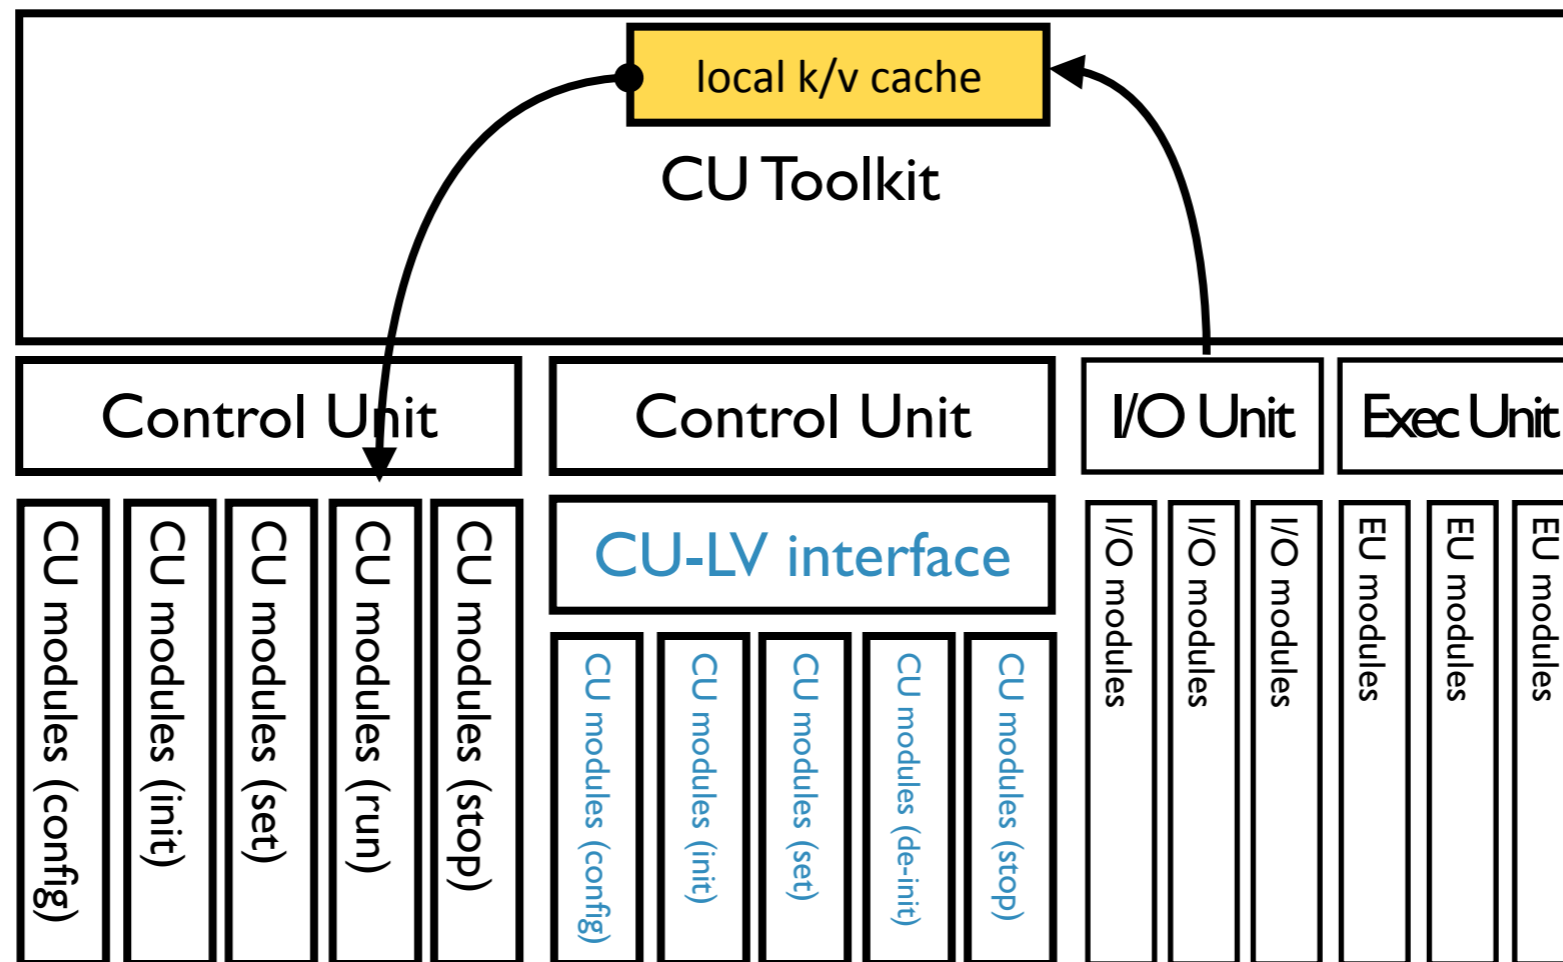
!CHAOS
infrastructure



I/O Unit allows decoupling device logic from signal acquisition

I/O Units, programmed by hw experts, will read Input channels and store data in a local k/v cache (the same technology as for live-data)

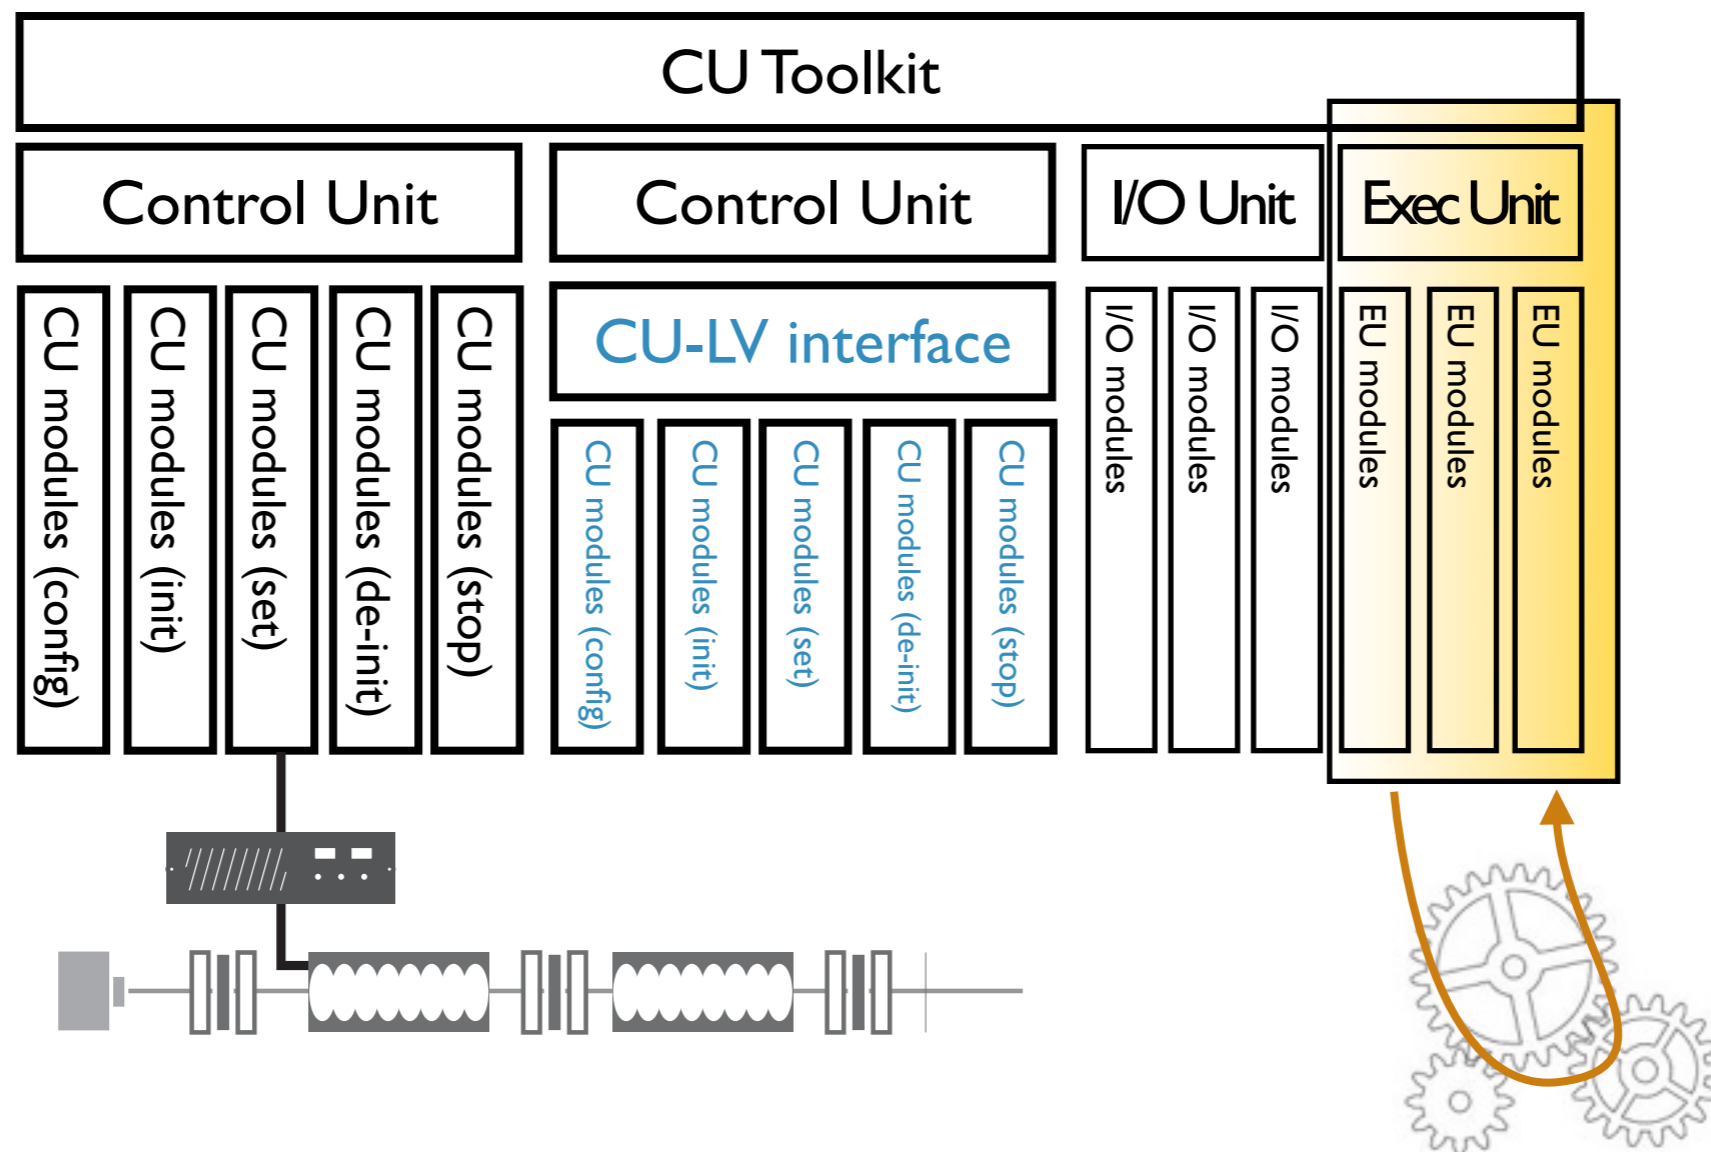
CU programmers will get data from local k/v cache instead of I/O modules directly



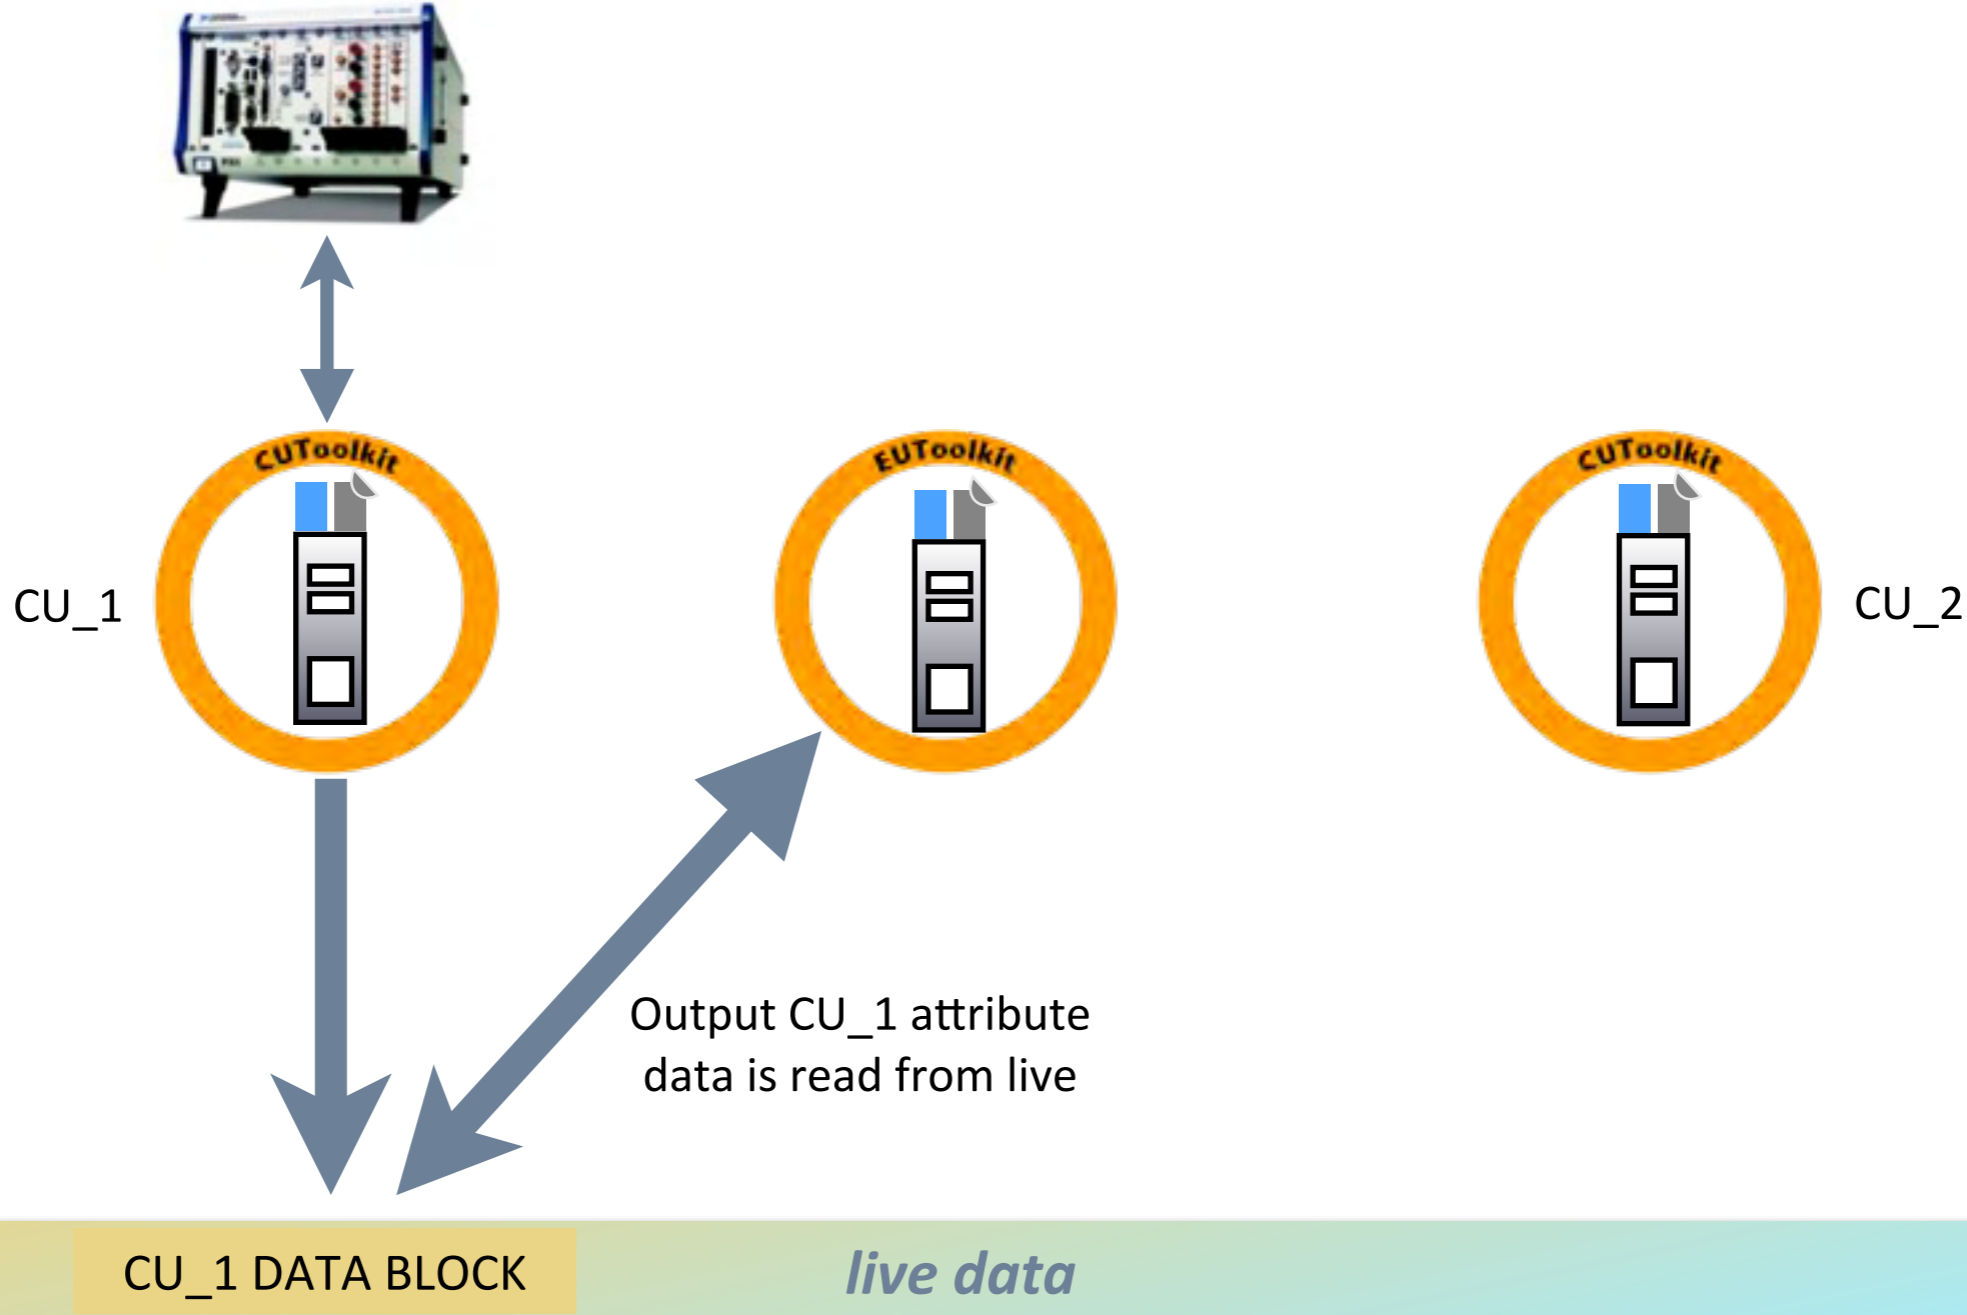
Execution Unit extend the concept of Control Unit to computational tasks

EU are not connected to any device, instead they produce data as result of a logic or calculations. They produce an output data class as consequence of the value defined by the input data class.

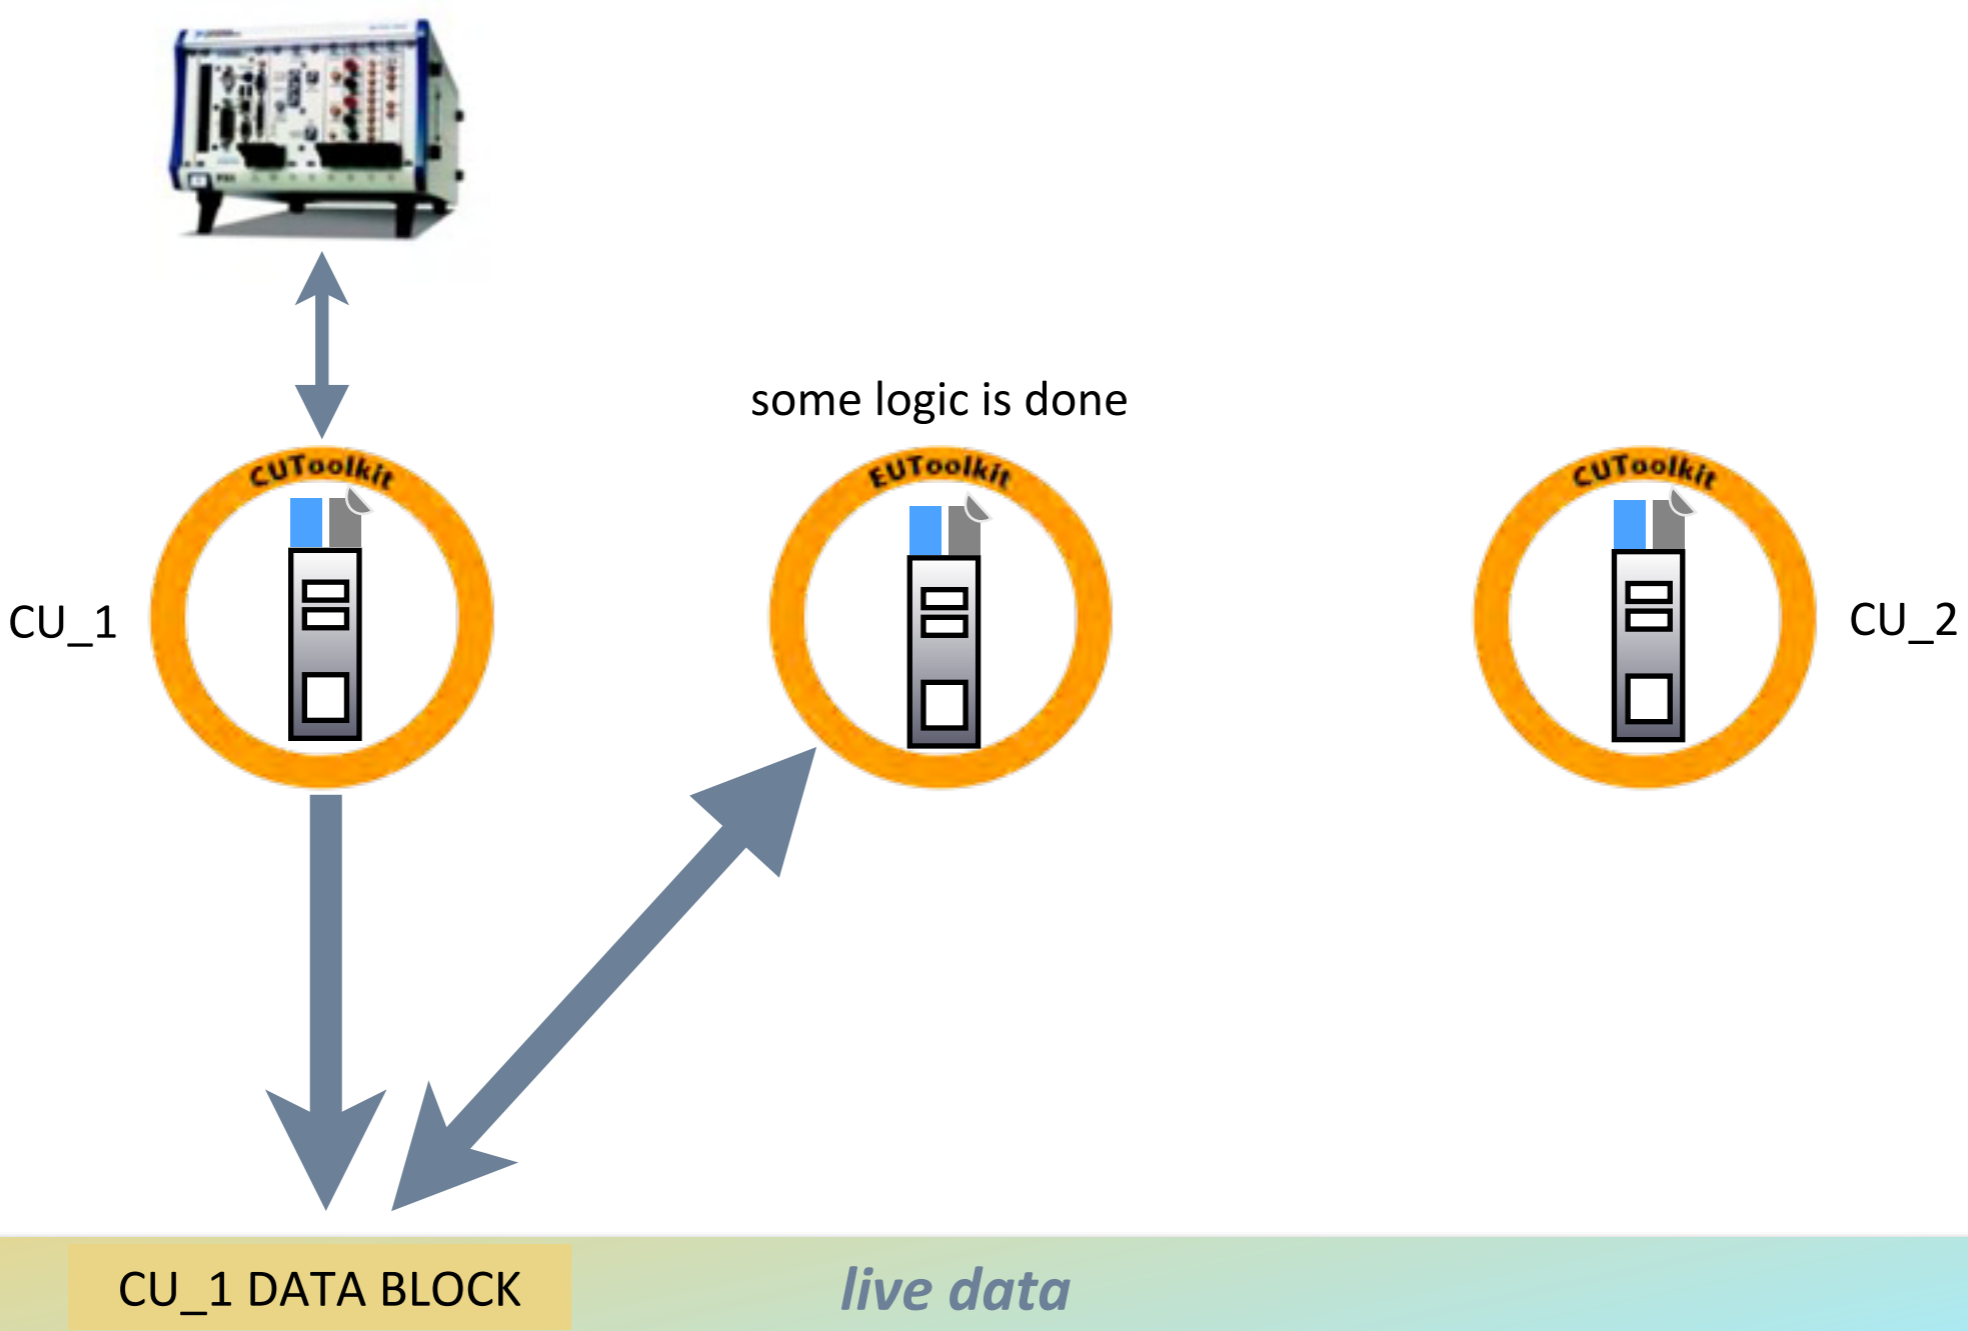
EU wil be used to implement analysis, measurement procedures, feedbacks



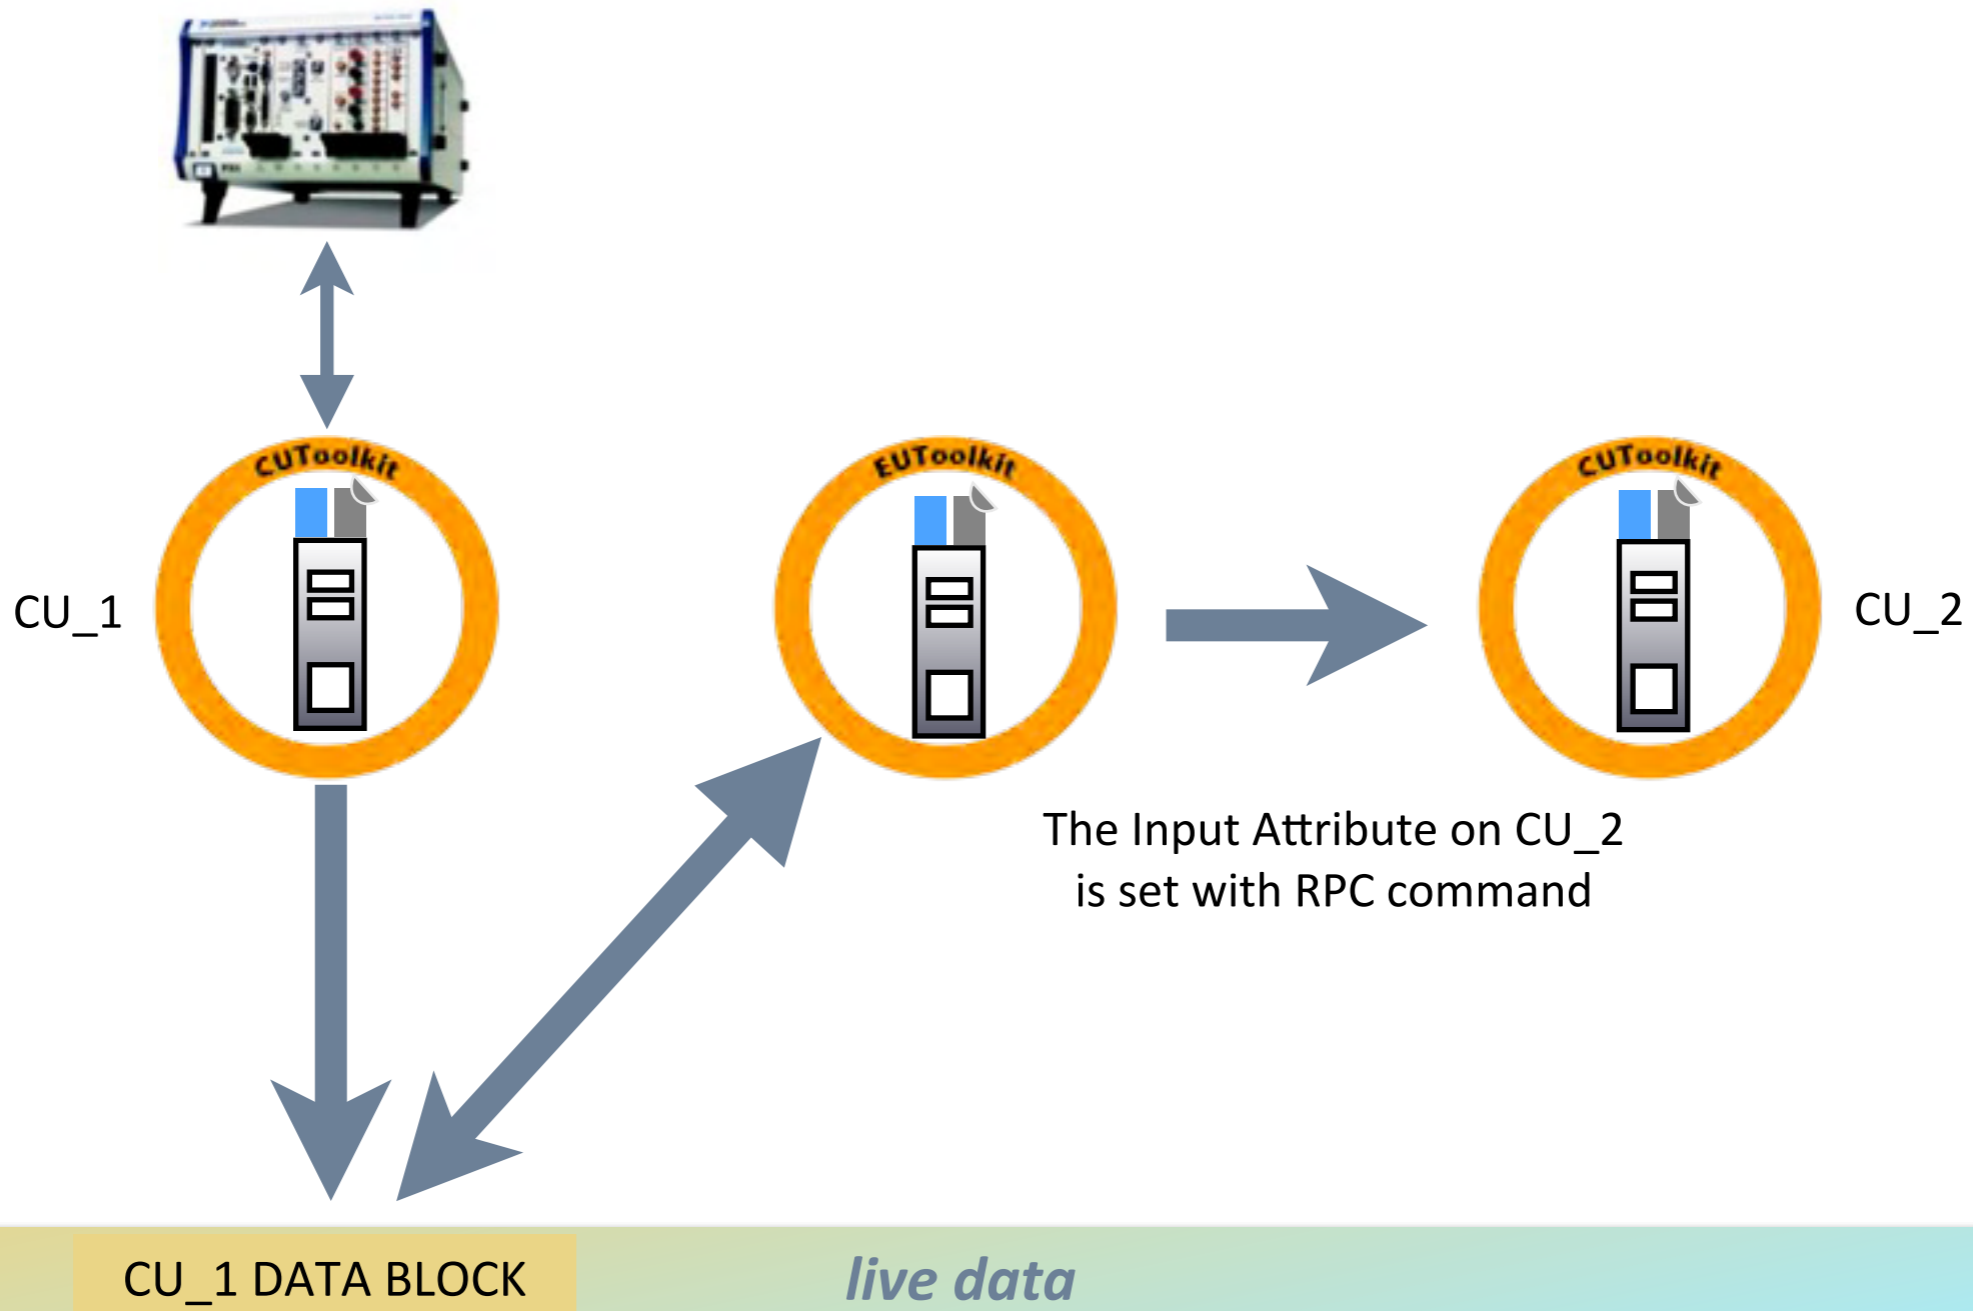
Execution Unit Example 1



Execution Unit Example 1



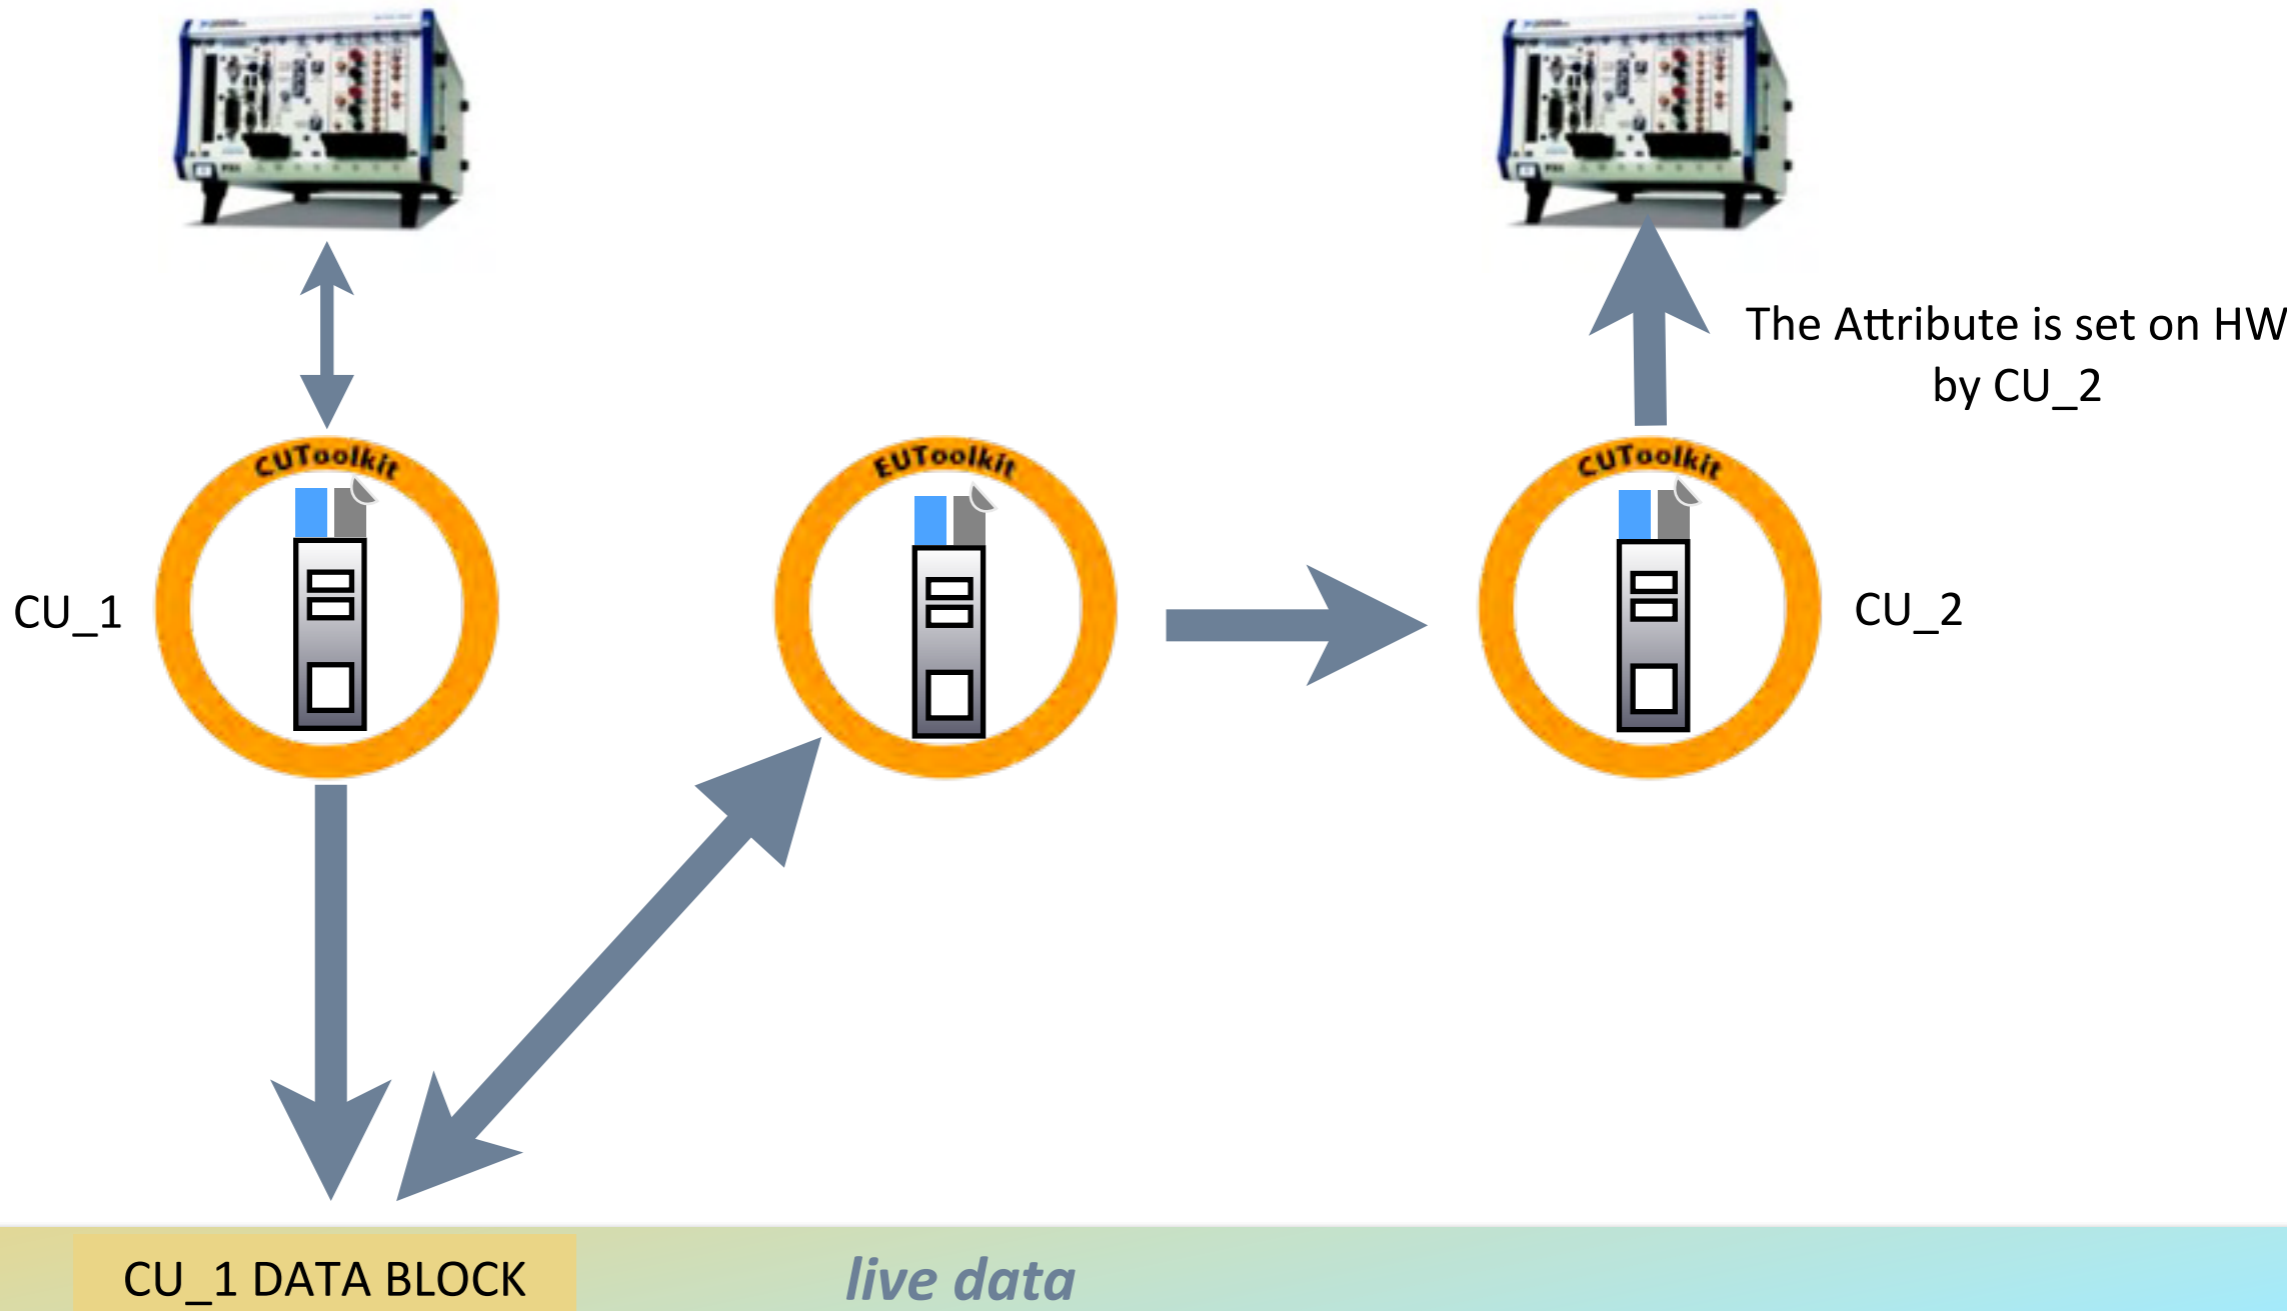
Execution Unit Example 1



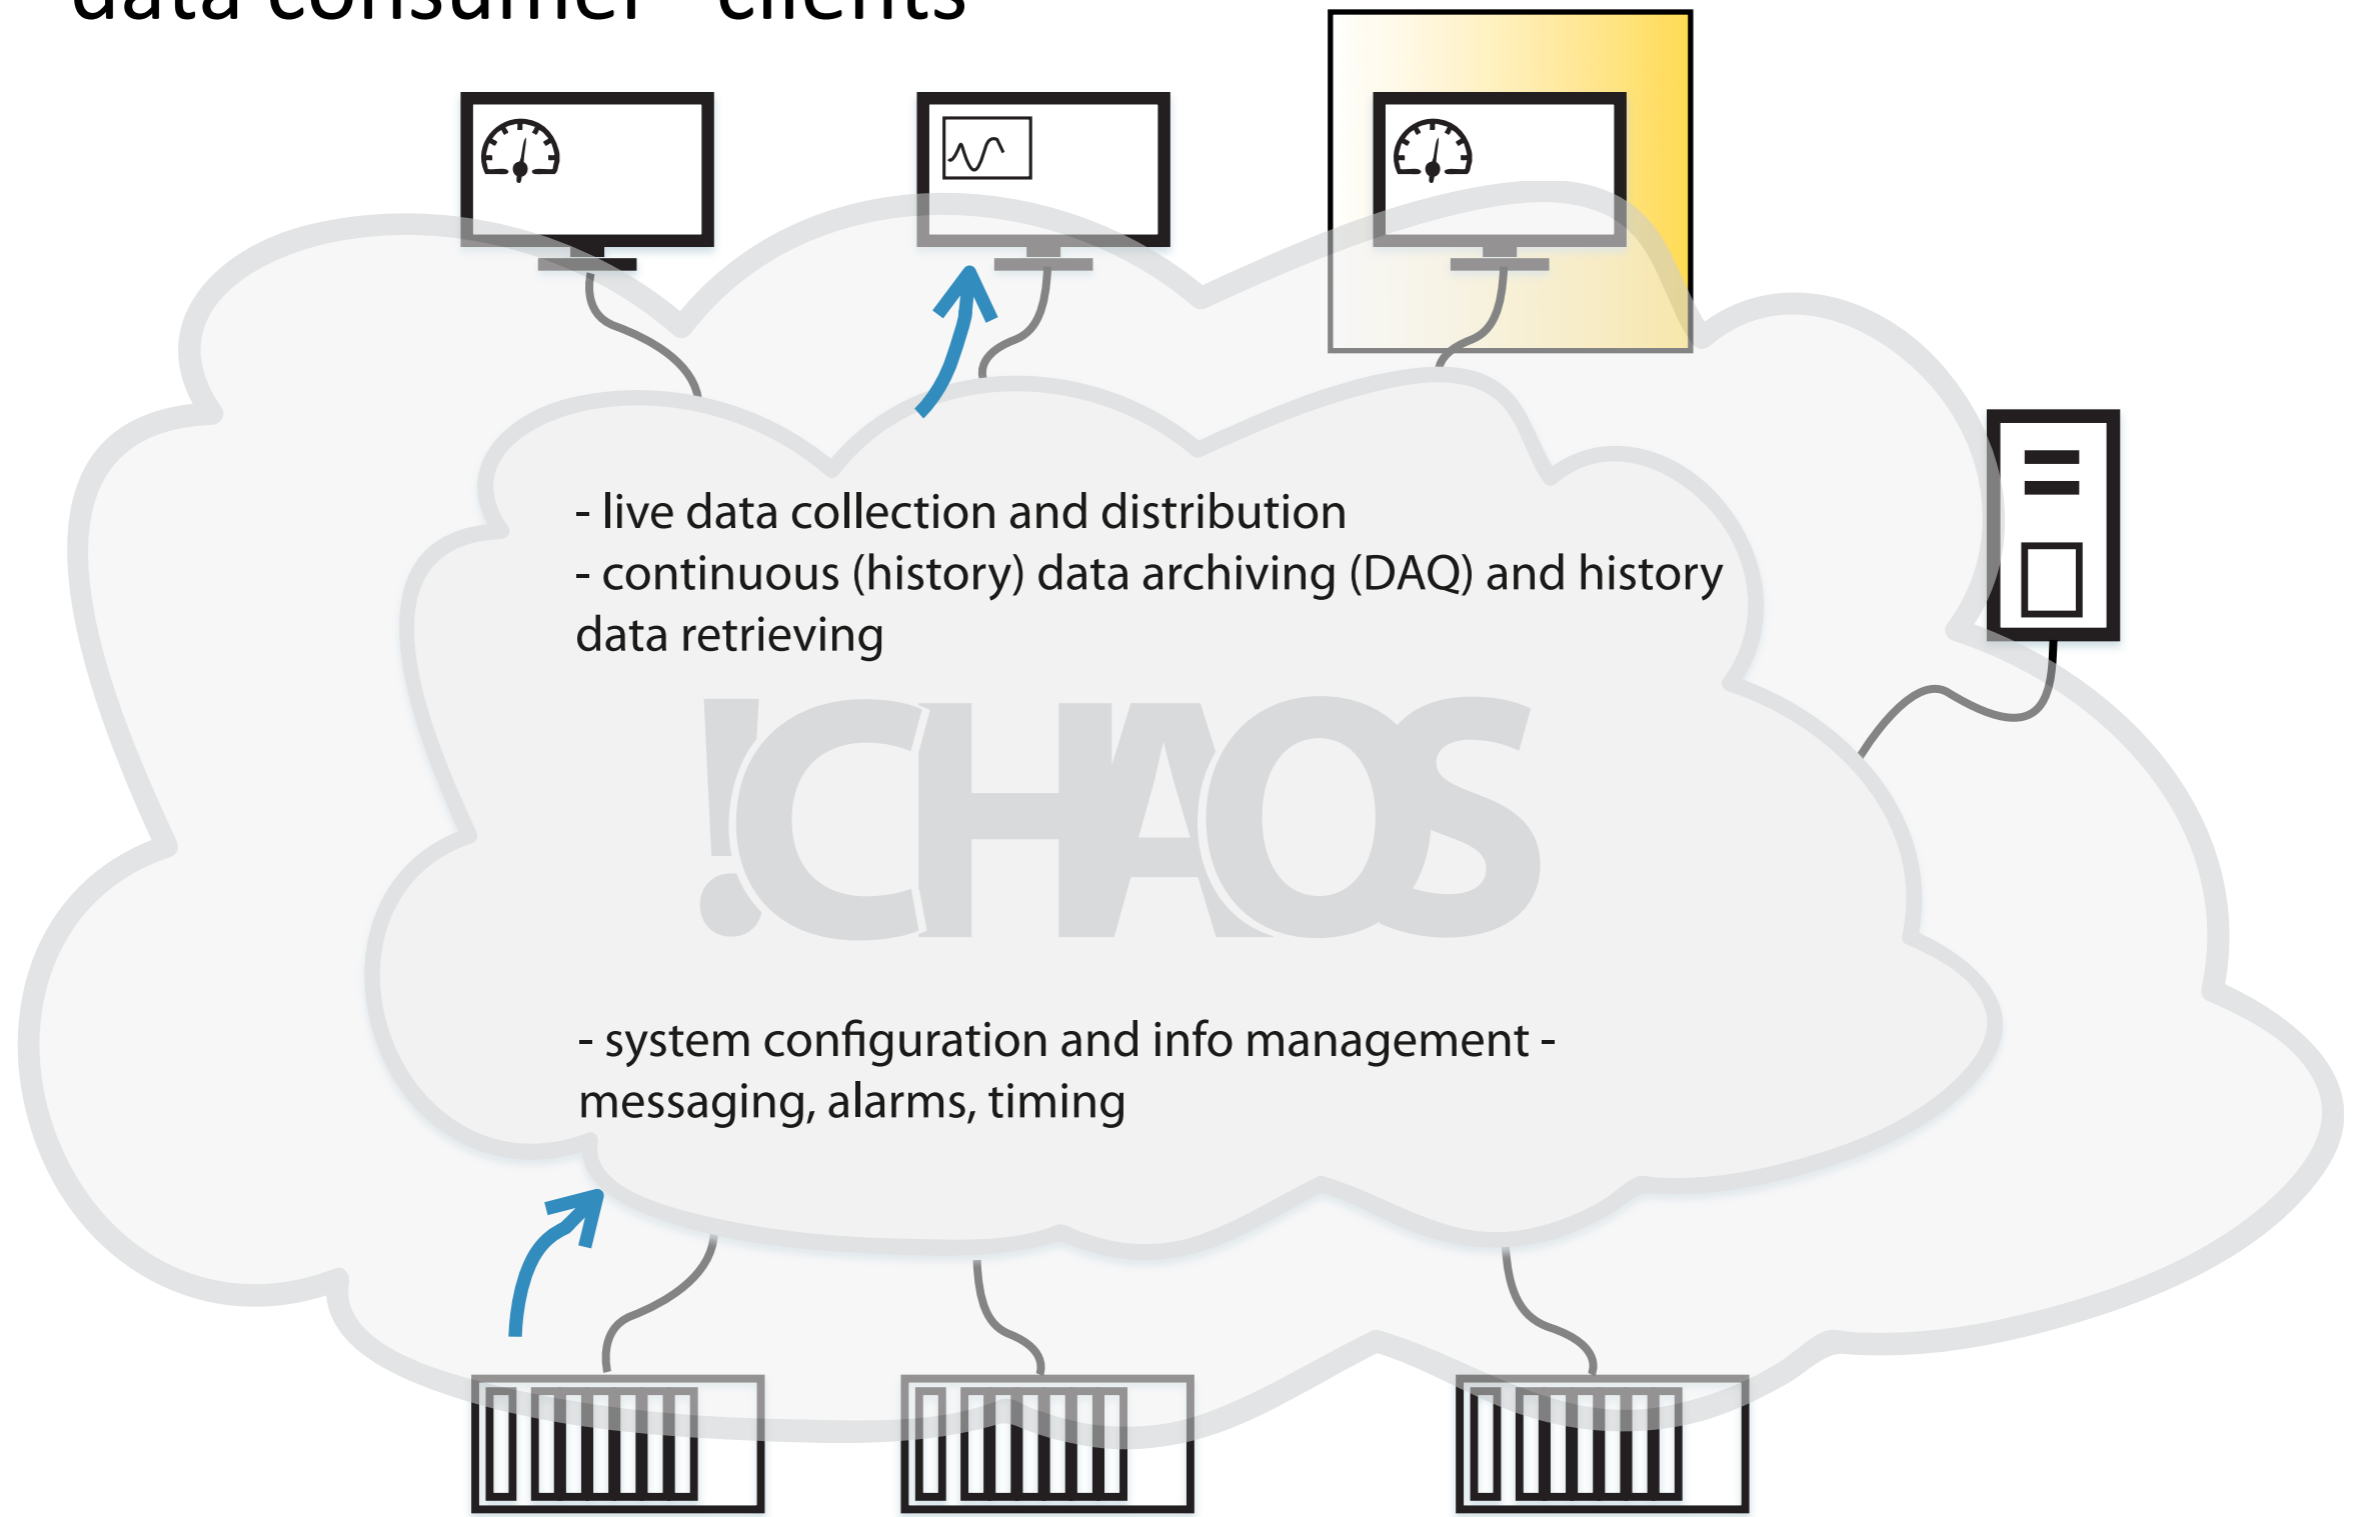
CU_1 DATA BLOCK

live data

Execution Unit Example 1

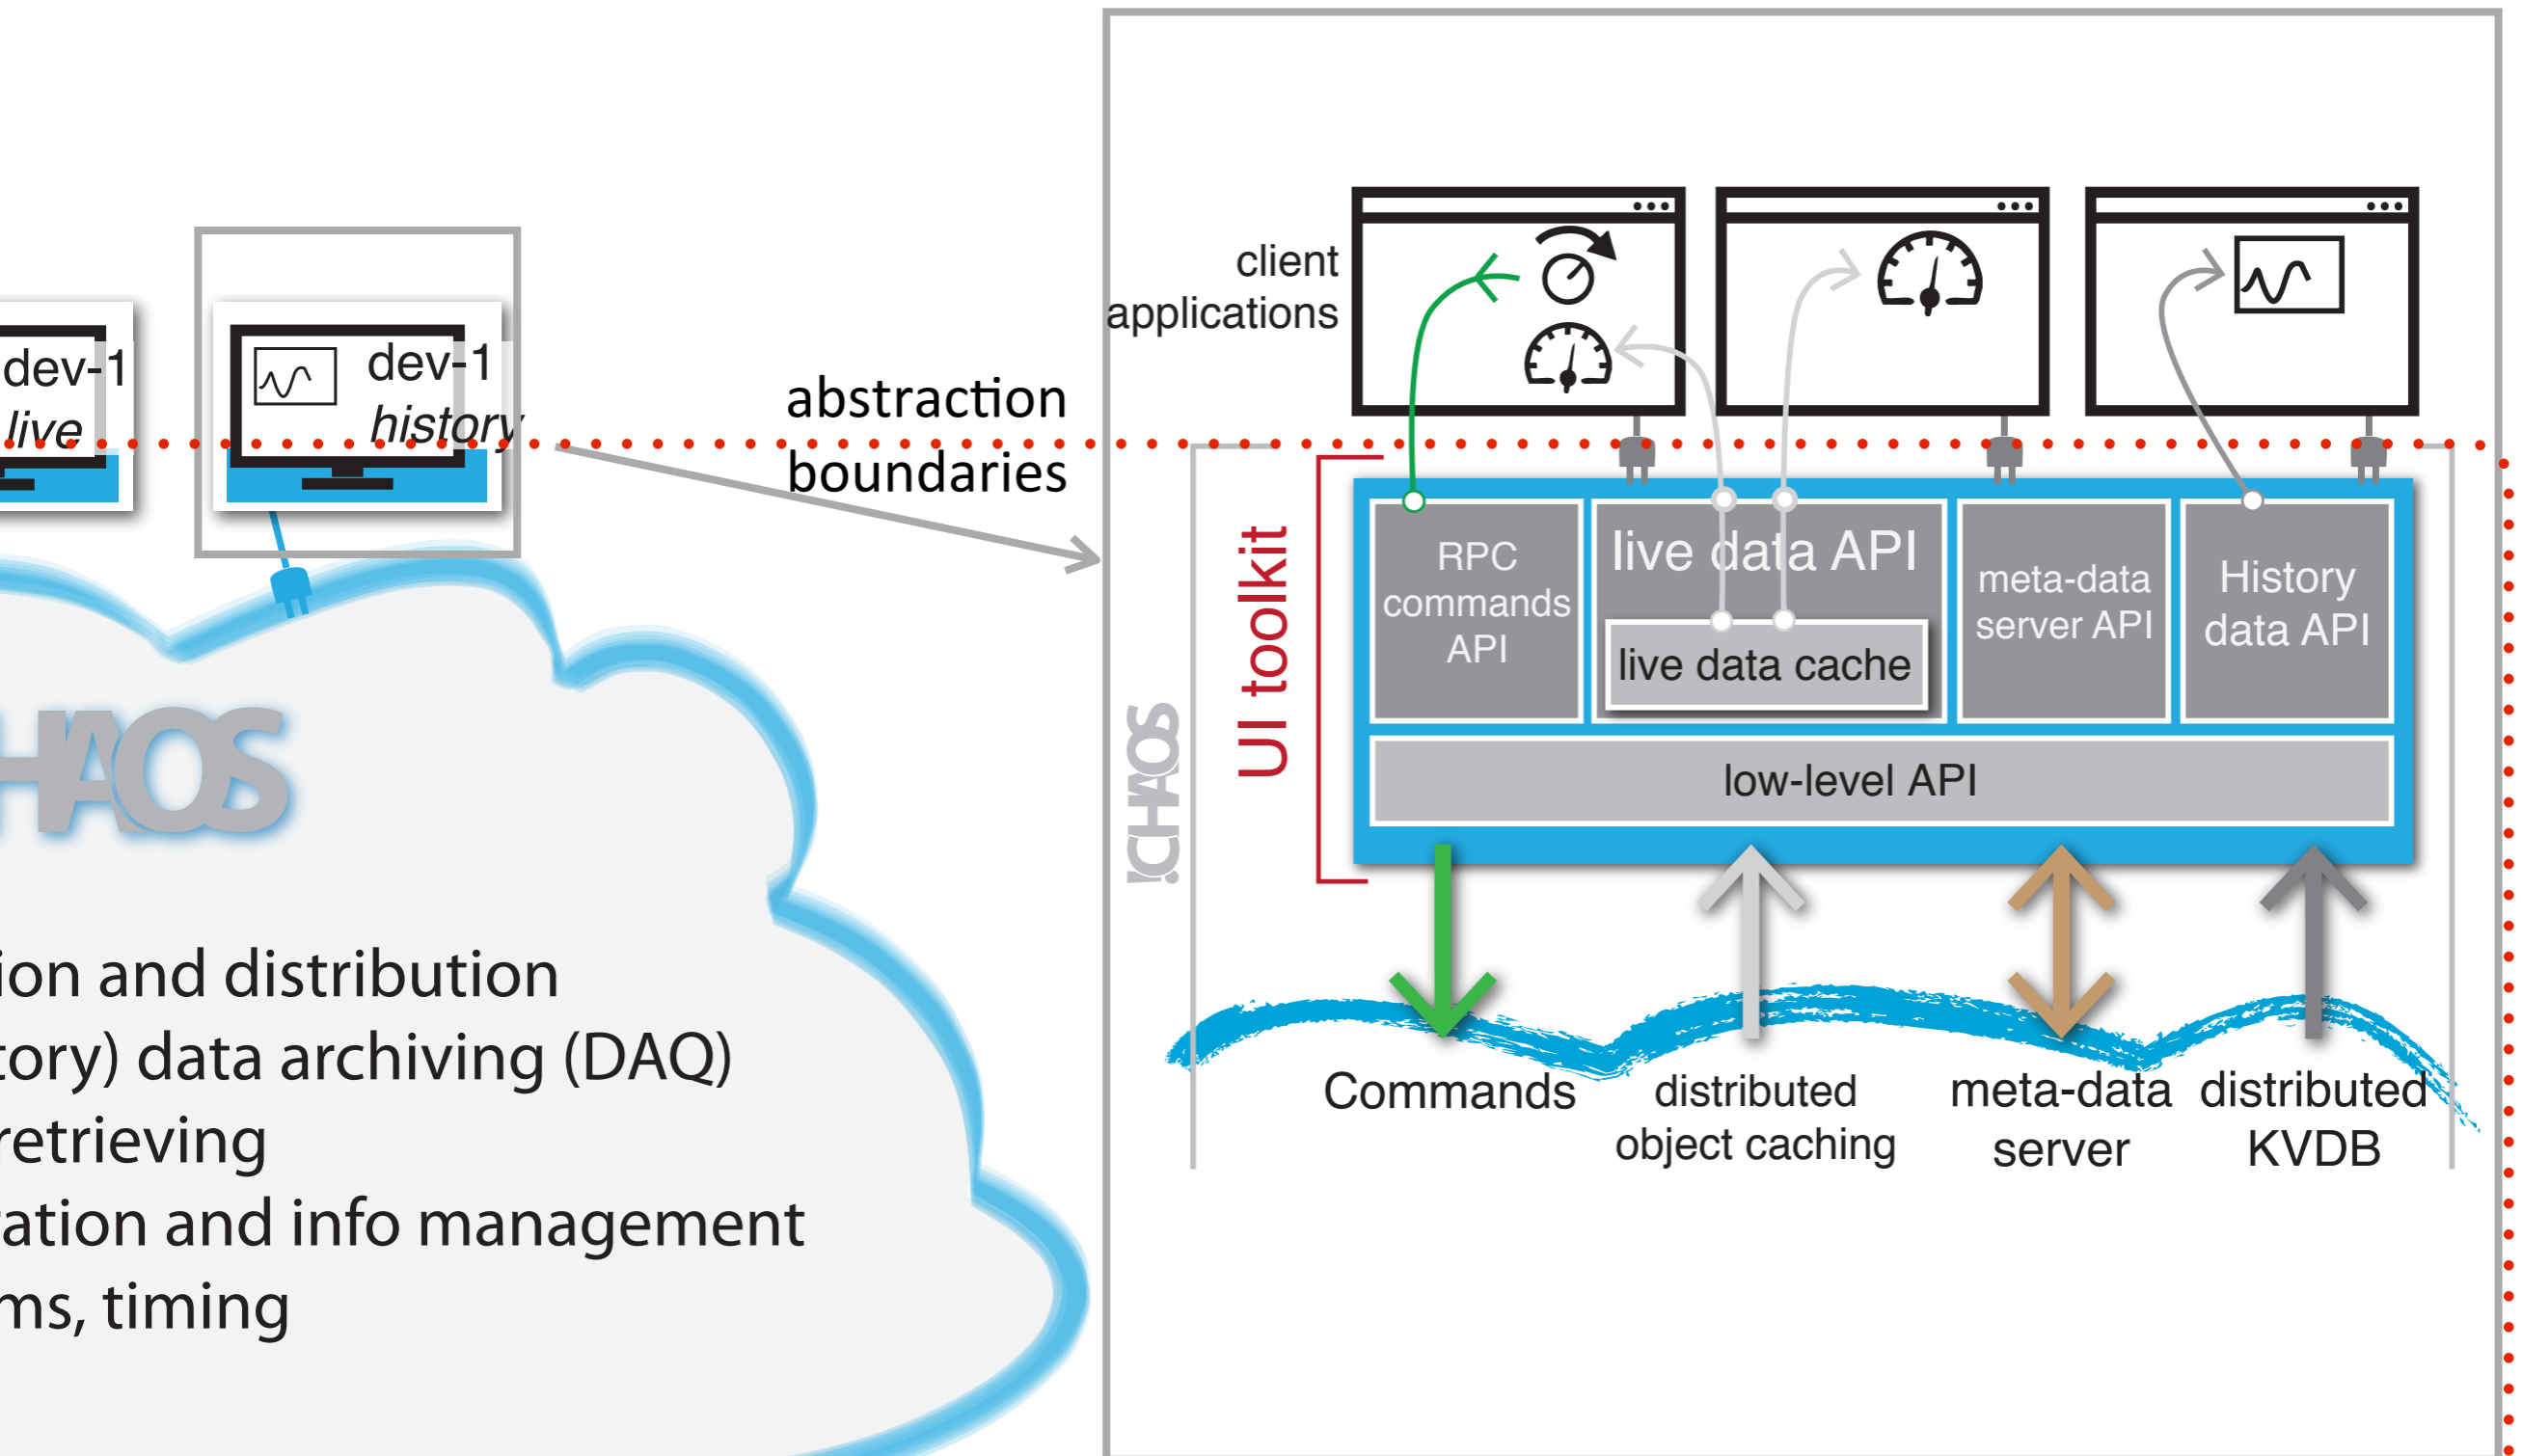


“data consumer” clients

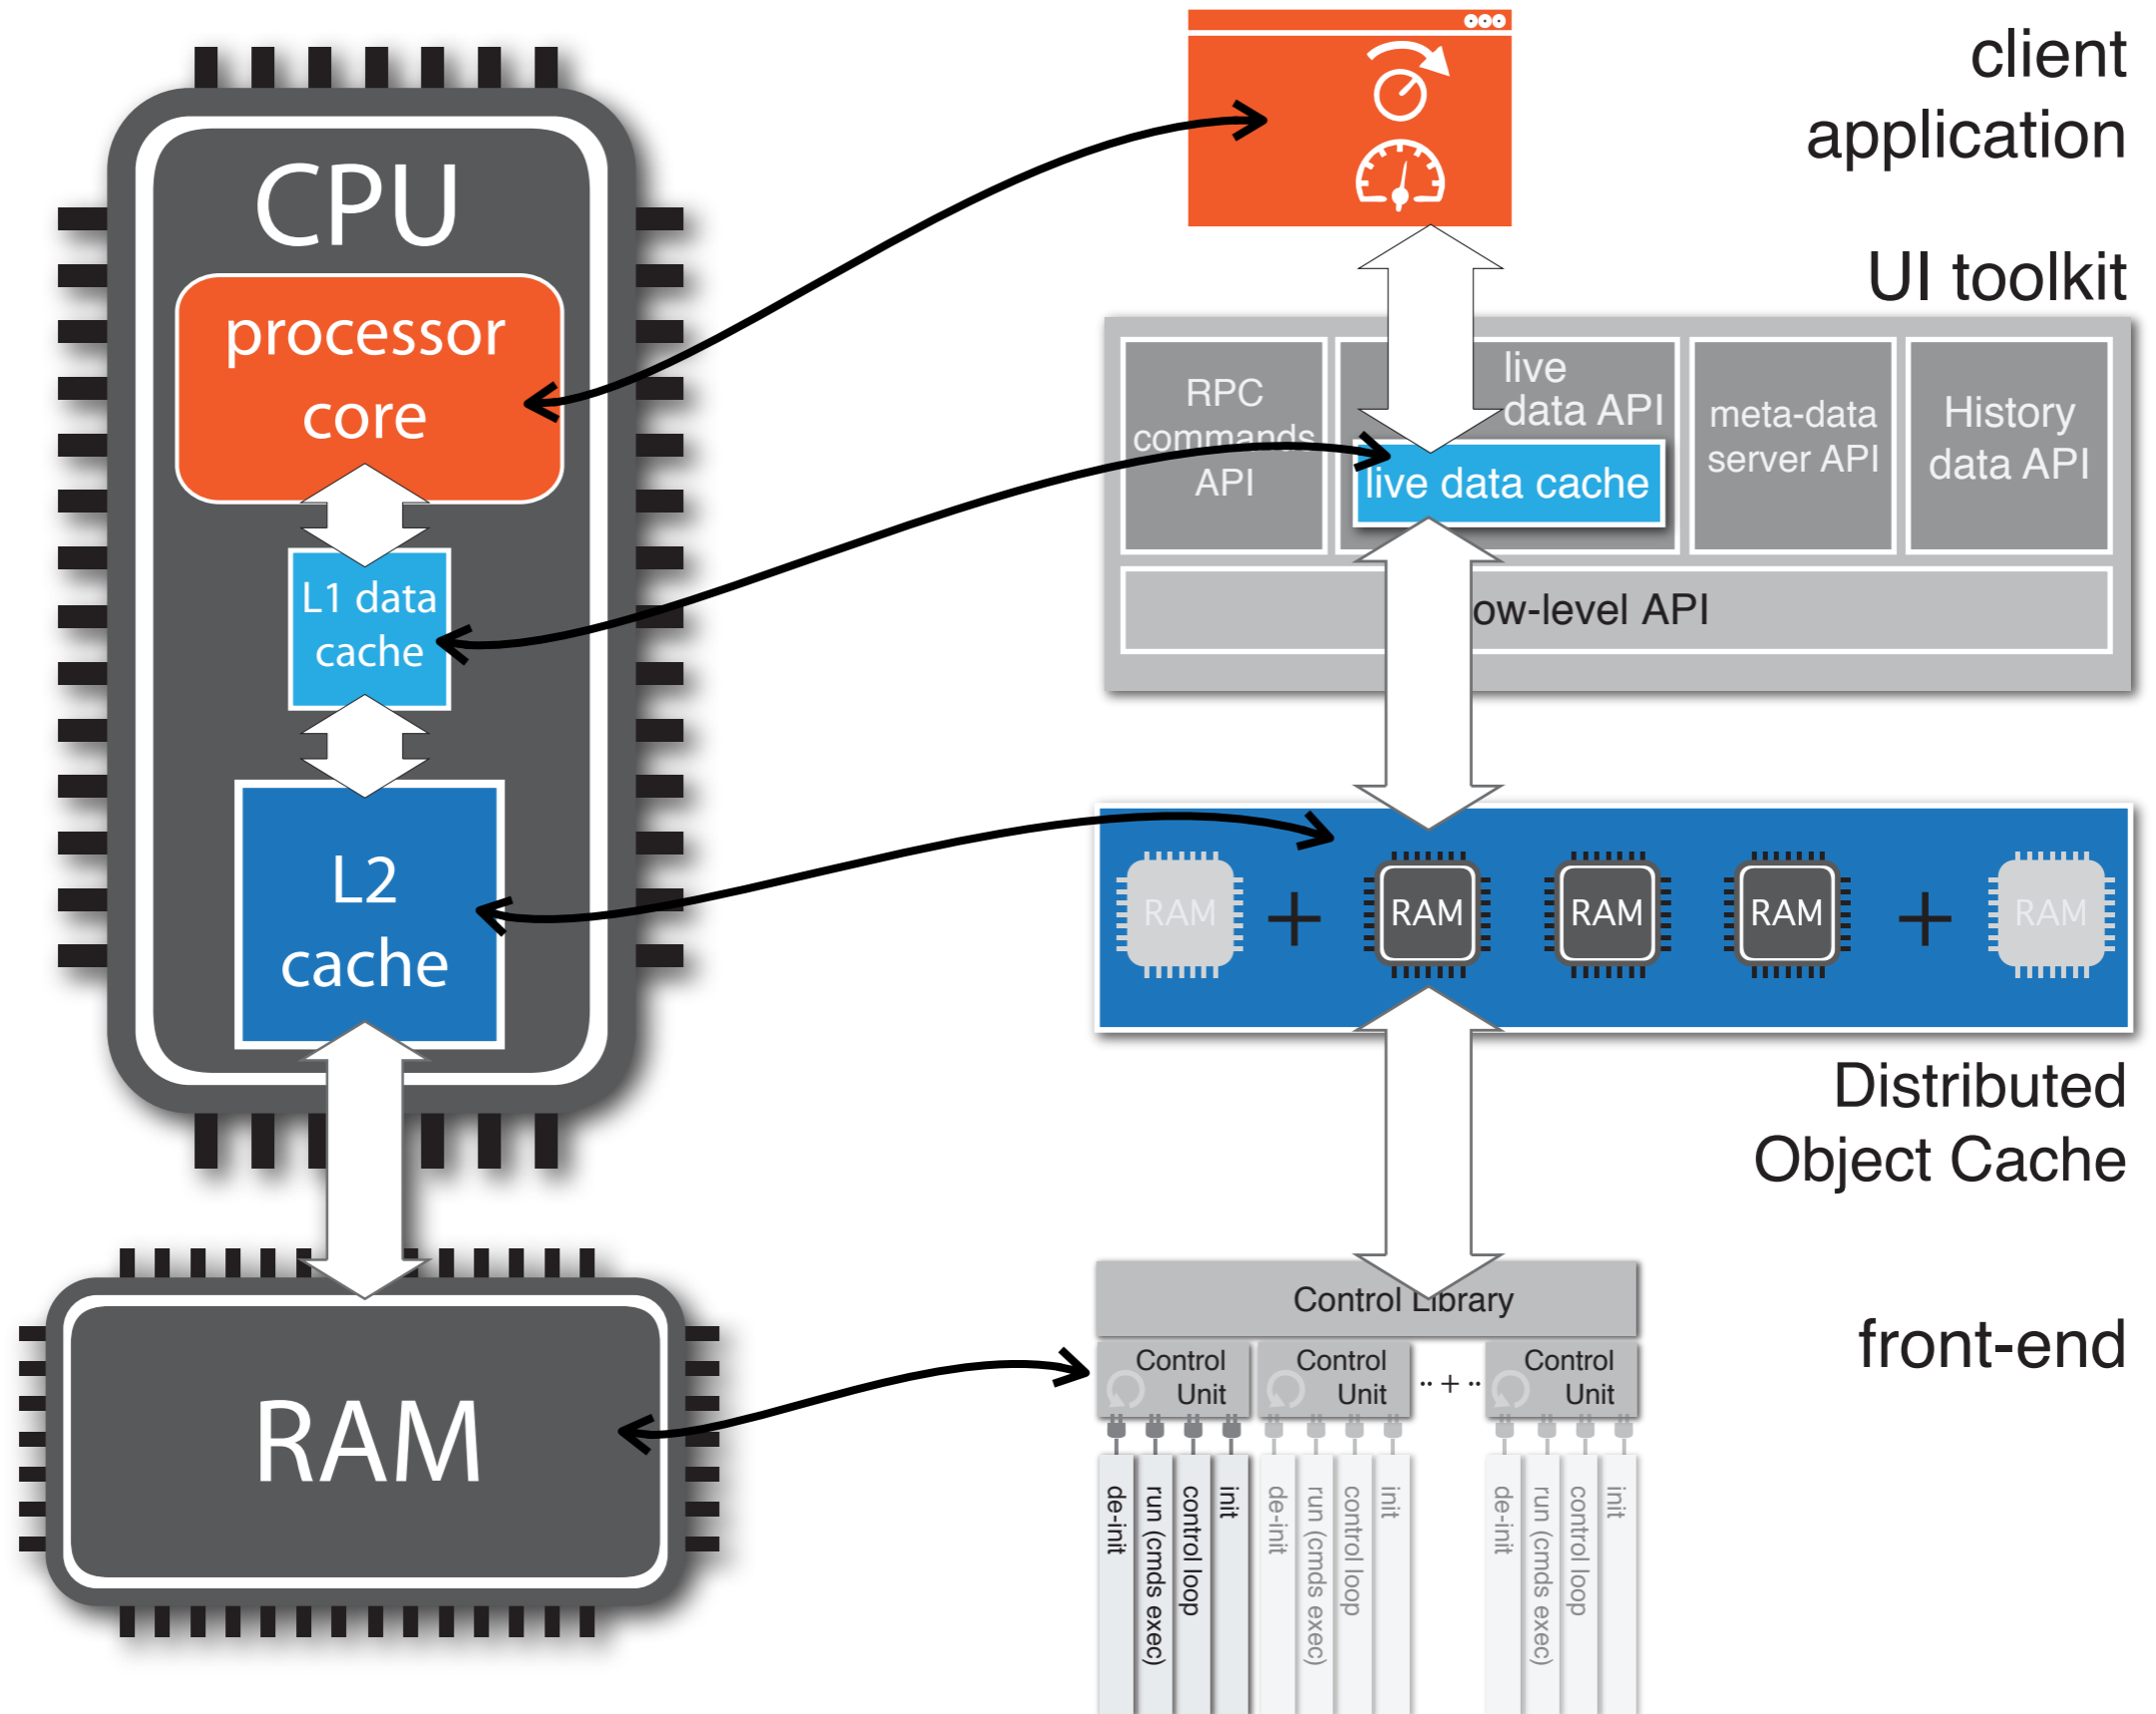


“data producer” clients

client abstraction

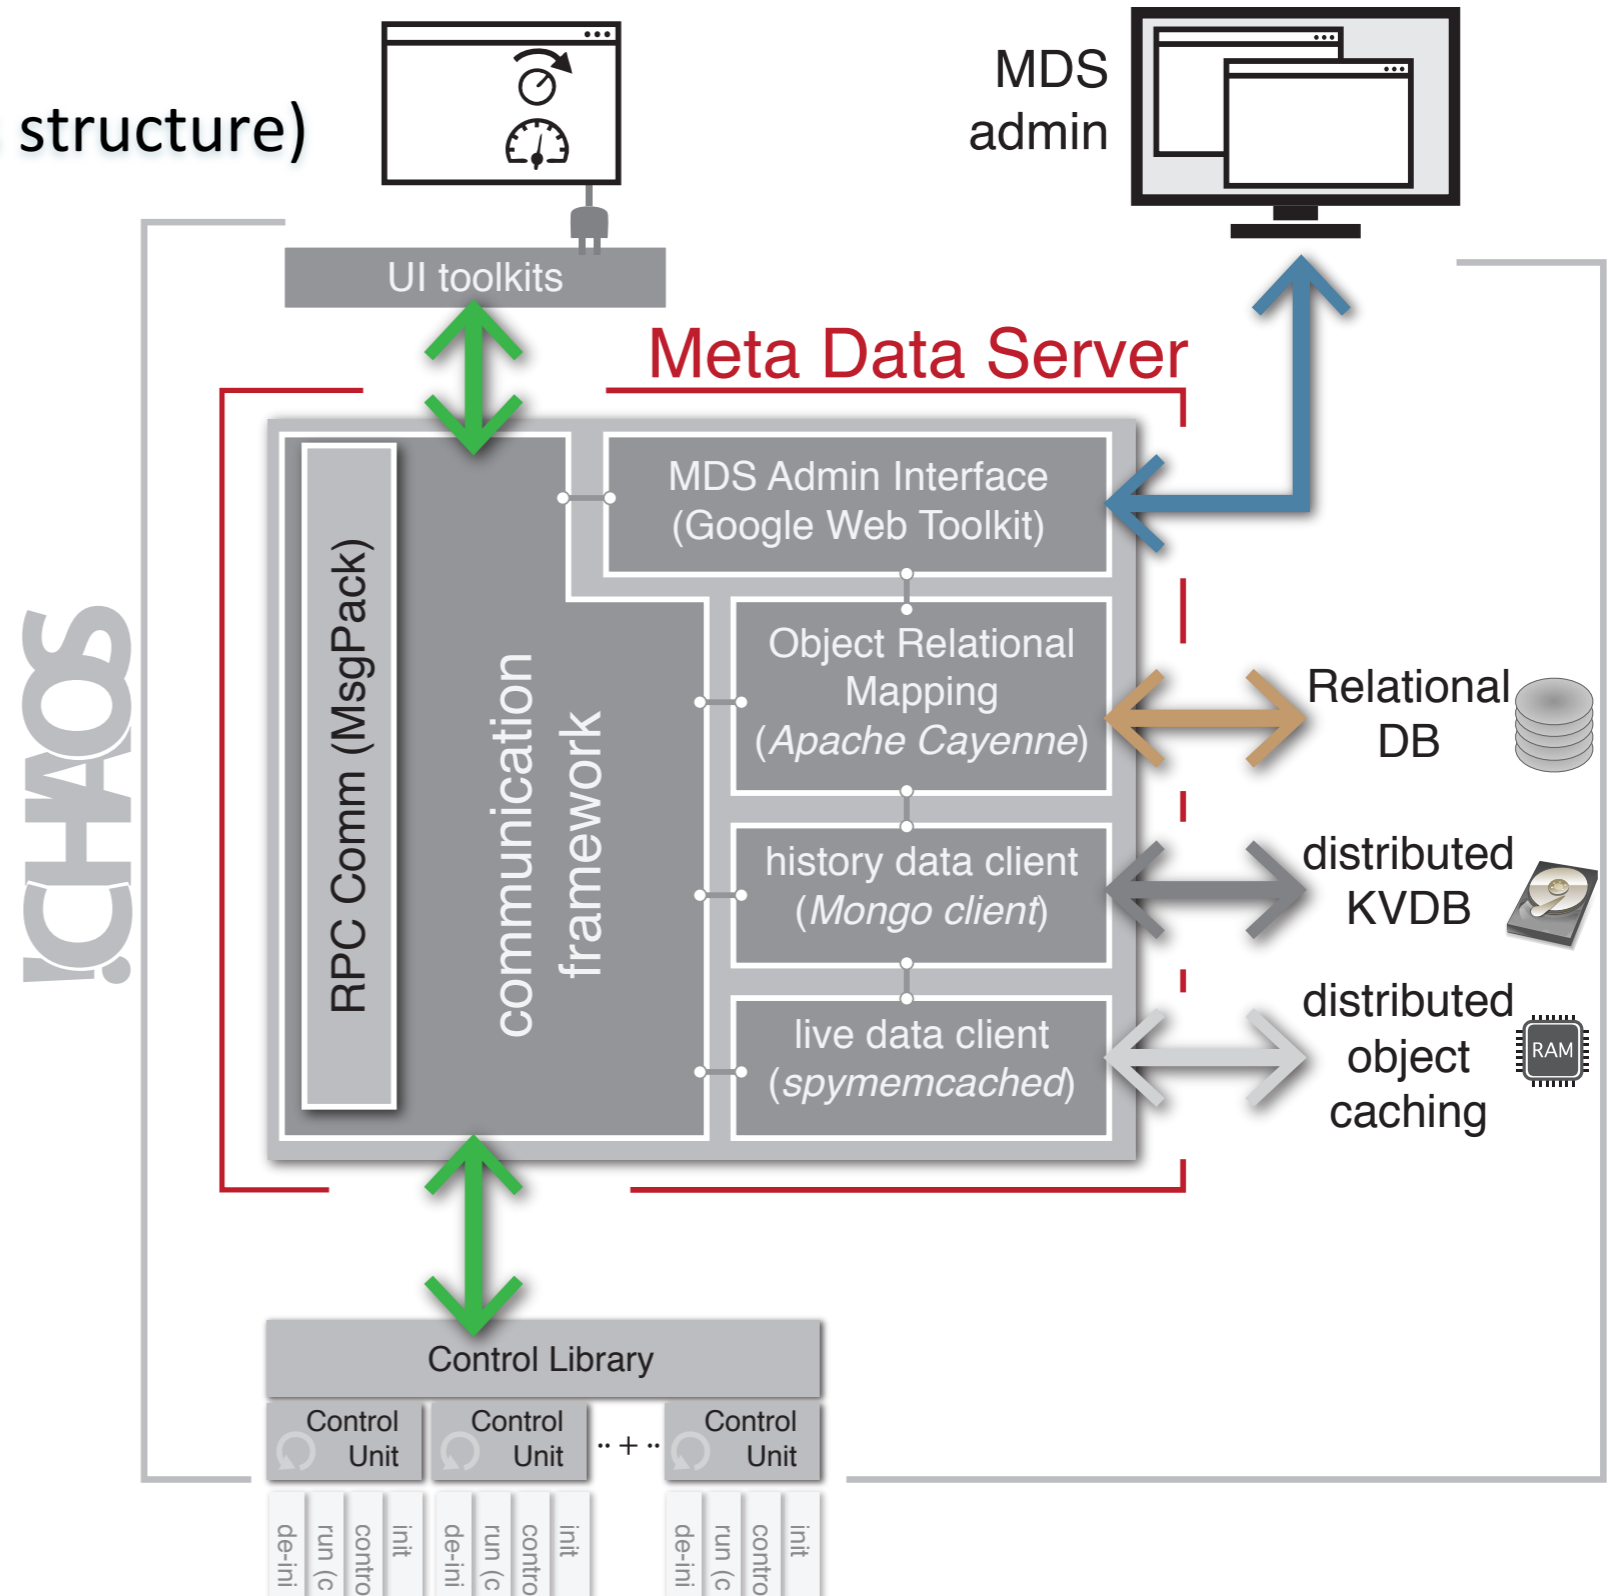


are we simply caching everywhere ?




Meta-data Server

- CU configuration manager
(e.g. managing of pushing data rate)
- Semantic of data (e.g. db records structure)
- Command's list and semantic
- Naming service




test #3.1




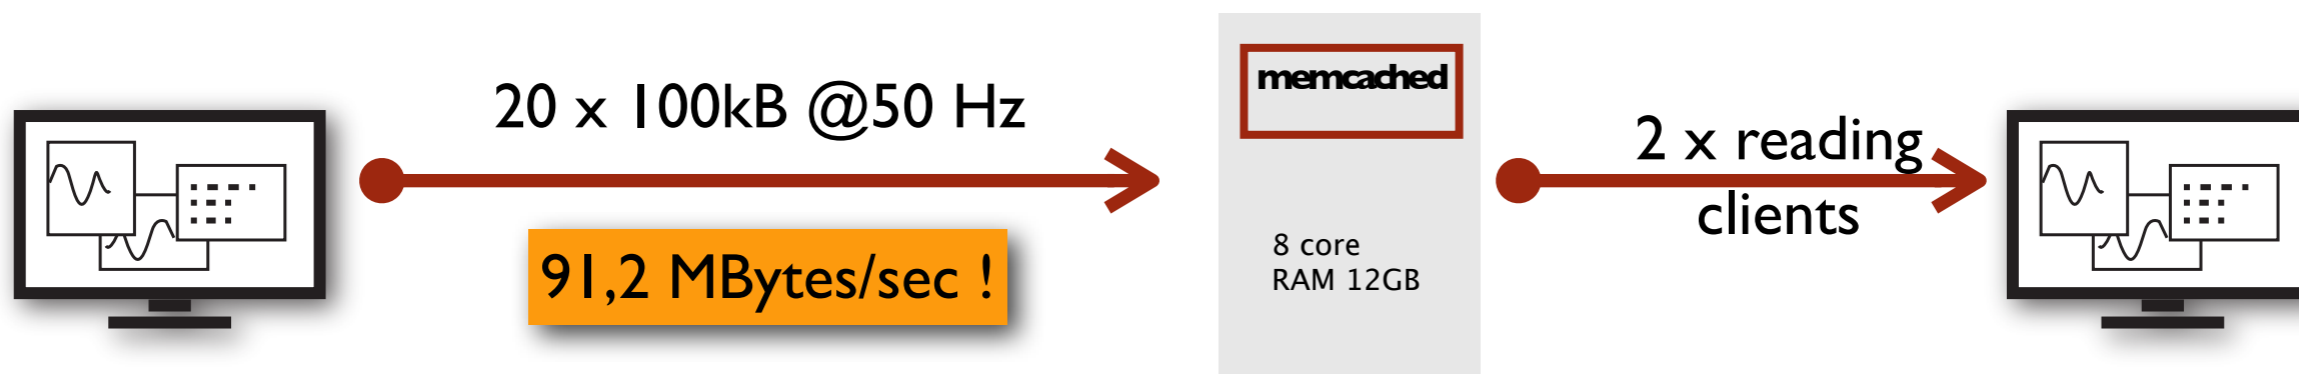
writing every... (msec)	#CU (Write)	#clients (Read)	#servers	#processes/ server	CPU load (%)
20	60	20	1	1	3-5
20	80	20	1	1	4-6
20	80	20	2	1	2-3
50	60	20	1	1	1-3
50	80	20	2	1	0-2
100	60	20	1	1	?
100	80	20	2	1	?

test #3.2




writing every... (msec)	#CU (Write)	#clients (Read)	#servers	#processes/ server	CPU load (%)
20	80	20	1	4 (1 per core)	2-3
20	80	40	1	4 (1 per core)	2-3
		40	1	4 (1 per core)	0

test #4




PID	USER	PR	NI	VIRT	RES	SHR	%CPU	MEM	TIME+	COMMAND
28059	dbuser	15	0	72236	10m	616	11.0	0.1	3:50.82	memcached
28066	dbuser	15	0	129m	5688	628	11.0	0.0	3:09.89	memcached
28052	dbuser	15	0	69812	8024	612	7.0	0.0	2:13.86	memcached
28074	dbuser	15	0	67568	5816	616	4.0	0.0	1:29.09	memcached


s4_hardware1_w20_m20_buff100000_rd10.log



s4_hardware1_w20_m20_buff100000_rd12.log

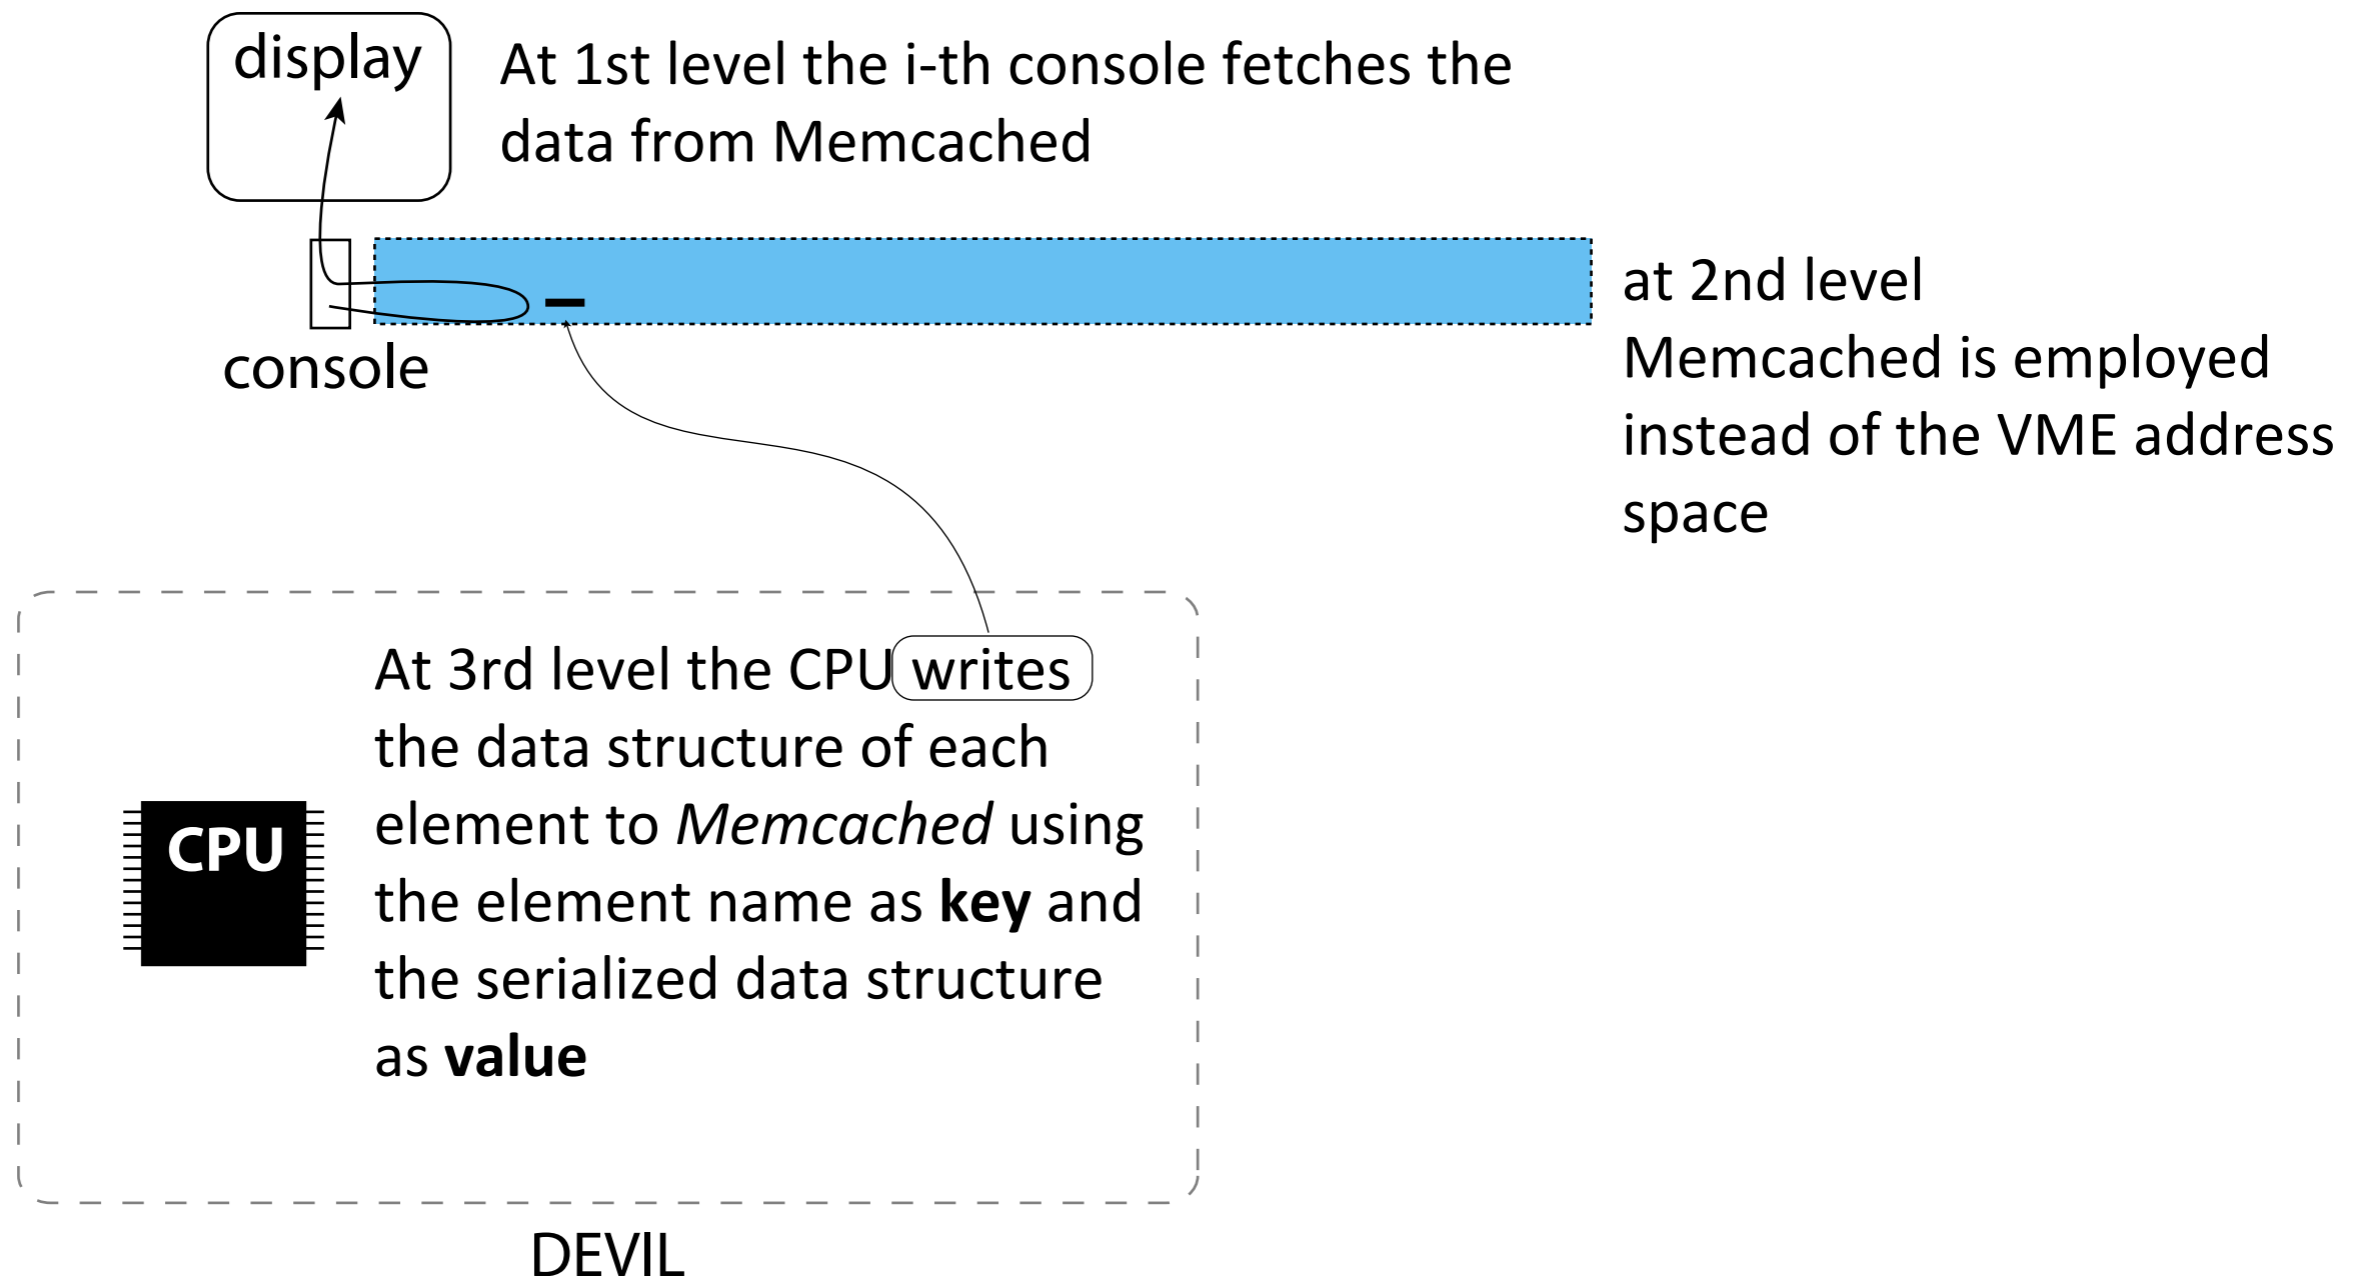


Measured transfer time between front-end CU and a client application via DOC for a different number of concurrent clients reading the same key/value being continuously updated by the CU.



tests at DAFNE

The idea: replace the 2nd level VME address space with the *Memcached* associative memory



Preliminary measurements

data size: 64 bytes for packet read

elemName	859504
status setting	0
status readout	2
consoleName	0
errorMask	4
sysFlags	
onLine	<input checked="" type="checkbox"/>
byPass	<input type="checkbox"/>
remote	<input checked="" type="checkbox"/>
busy	<input type="checkbox"/>
HV setting	0.00
HV readout	201.48
faults	0 b0

fetch frequency ~ 100 Hz

with no dependency on the number of fetching consoles
(up to 7 in our test)


Memcached server load (measured with the *top* command)

CPU: 0.3% - 0.7% memory: $\sim 0.1\%$

tests at SPARC

Memcached has been used for storing the beam image from a digital camera of beam diagnostic.


- network: Ethernet @1 Gbps
- image size: 640x480
 - @8 bit = **300 kB**
 - @16 bit = **600 kB**



measured fetch frequency: **8bit ~ 25 Hz, 16bit ~ 13 Hz**; same transfer rate for all consoles fetching images from memcached (up to 4 in our test)

!CHAOS

- aims at the development of a control system for INFN future accelerators and large apparatuses
- introduces a breaking through paradigm based on high performance internet technologies
- profits from INFN long tradition of control systems for particle accelerators, expertise in computing and information technologies
- can be a complement for future INFN interdisciplinary projects
- generates opportunities for collaborations with industries and technology transfer



!CHAOS

<http://chaos.infn.it>

- !CHAOS launch workshop at LNF (Dec. 2011)
- evaluation release (!CHAOS_beta_0.1) ready to download from website
- !CHAOS established as **Open Source project** of INFN
- candidate Control System for SuperB accelerator and for slow-controls of experimental apparatus
- collaboration with National Instruments for integration of LabVIEW in !CHAOS
 - NI-!CHAOS meeting at LNF (March 2012)
 - invited at Big Physics Symposium (Zurich - March 2012)
- collaborations with italian industries (agreements for joint developments in preparation)
- academic collaborations: Univ. Roma TV (Fac. Informatica e Ingegneria), Politecnico di Bari, Università di Cagliari (Fac. Ingegneria)

