SC Timing Acquisition Services

Production Deployment Integration Tests and Mitigation

Marcio Donadio / Control Systems Eng, / TID-CDS-ACS

06/08/2023

SLAC NATIONAL ACCELERATOR LABORATORY

Stanford University | U.S. DEPARTMENT OF ENERGY

Outline

Introduction

Acquisition services
Current architecture

Analysis of the implemented fixes

IOCs were crashing
Network stack busy
Latency
Other fixes

Improvements

Implemented since December 2022
Rethinking the strategy

Stress tests performed

Ready to test
Focus of the tests
BSA - The dimension of the load
BLD - General View
BSAS - General View
Python script to handle camonitor

Results of the test in production

Adding one IOC at a time
BLD has high load
Curious side effect used as a tool
Comparison between 2 ATCA vs 1 ATCA
BSAS can use a lot of resources, too
camonitor put a challenge to lcls-srv01

Conclusions and recommendations

Tests were successful!
BSAS
Plan for BLD not feasible with current CPUs
Acknowledgements

1

Introduction

Introduction

Acquisition services

BSA = Beam Synchronous Acquisition. Organizes data in arrays, all aligned with the same timestamp.

BSSS = Beam Synchronous Scalar Service. Samples data from BSA arrays and publishes it as a single number.

BSAS = Beam Synchronous Acquisition Service. Organizes statistical data in a 2D table containing dozens of columns and 1000 rows. It is published as a new EPICS 7 NTTable record type. Mainly used by machine learning algorithms.

BLD = Beam Line Data. Sends data directly to a multicast network, without storing it. There's no PV to see the data contents. Mainly used by the photon side.

Introduction

Acquisition services

Material is being created to explain each acquisition service in a friendly way:

Carolina's slides about BSA and BSSS:
https://slac-my.sharepoint.com/:p:/r/personal/carolina_slac_stanford_edu/_layouts/15/Doc.aspx?sourcedoc=%7Bc564a140-0cee-4ce0-be54-143fa71e8f03%7D&action=default&ct=1686160150001&or=OWA-NT&cid=2d35ed76-3162-0f33-b030-35c76bd315dd&_SRM=2%3AS%3A85&file=HLA_SC_BSA_2023.pptx

TID's Confluence Page (still under construction):
https://confluence.slac.stanford.edu/display/ppareg/BSA%2C+BSSS%2C+BLD%2C+and+BSAS+explained

bsaDriver module README's files.

Introduction

Current architecture

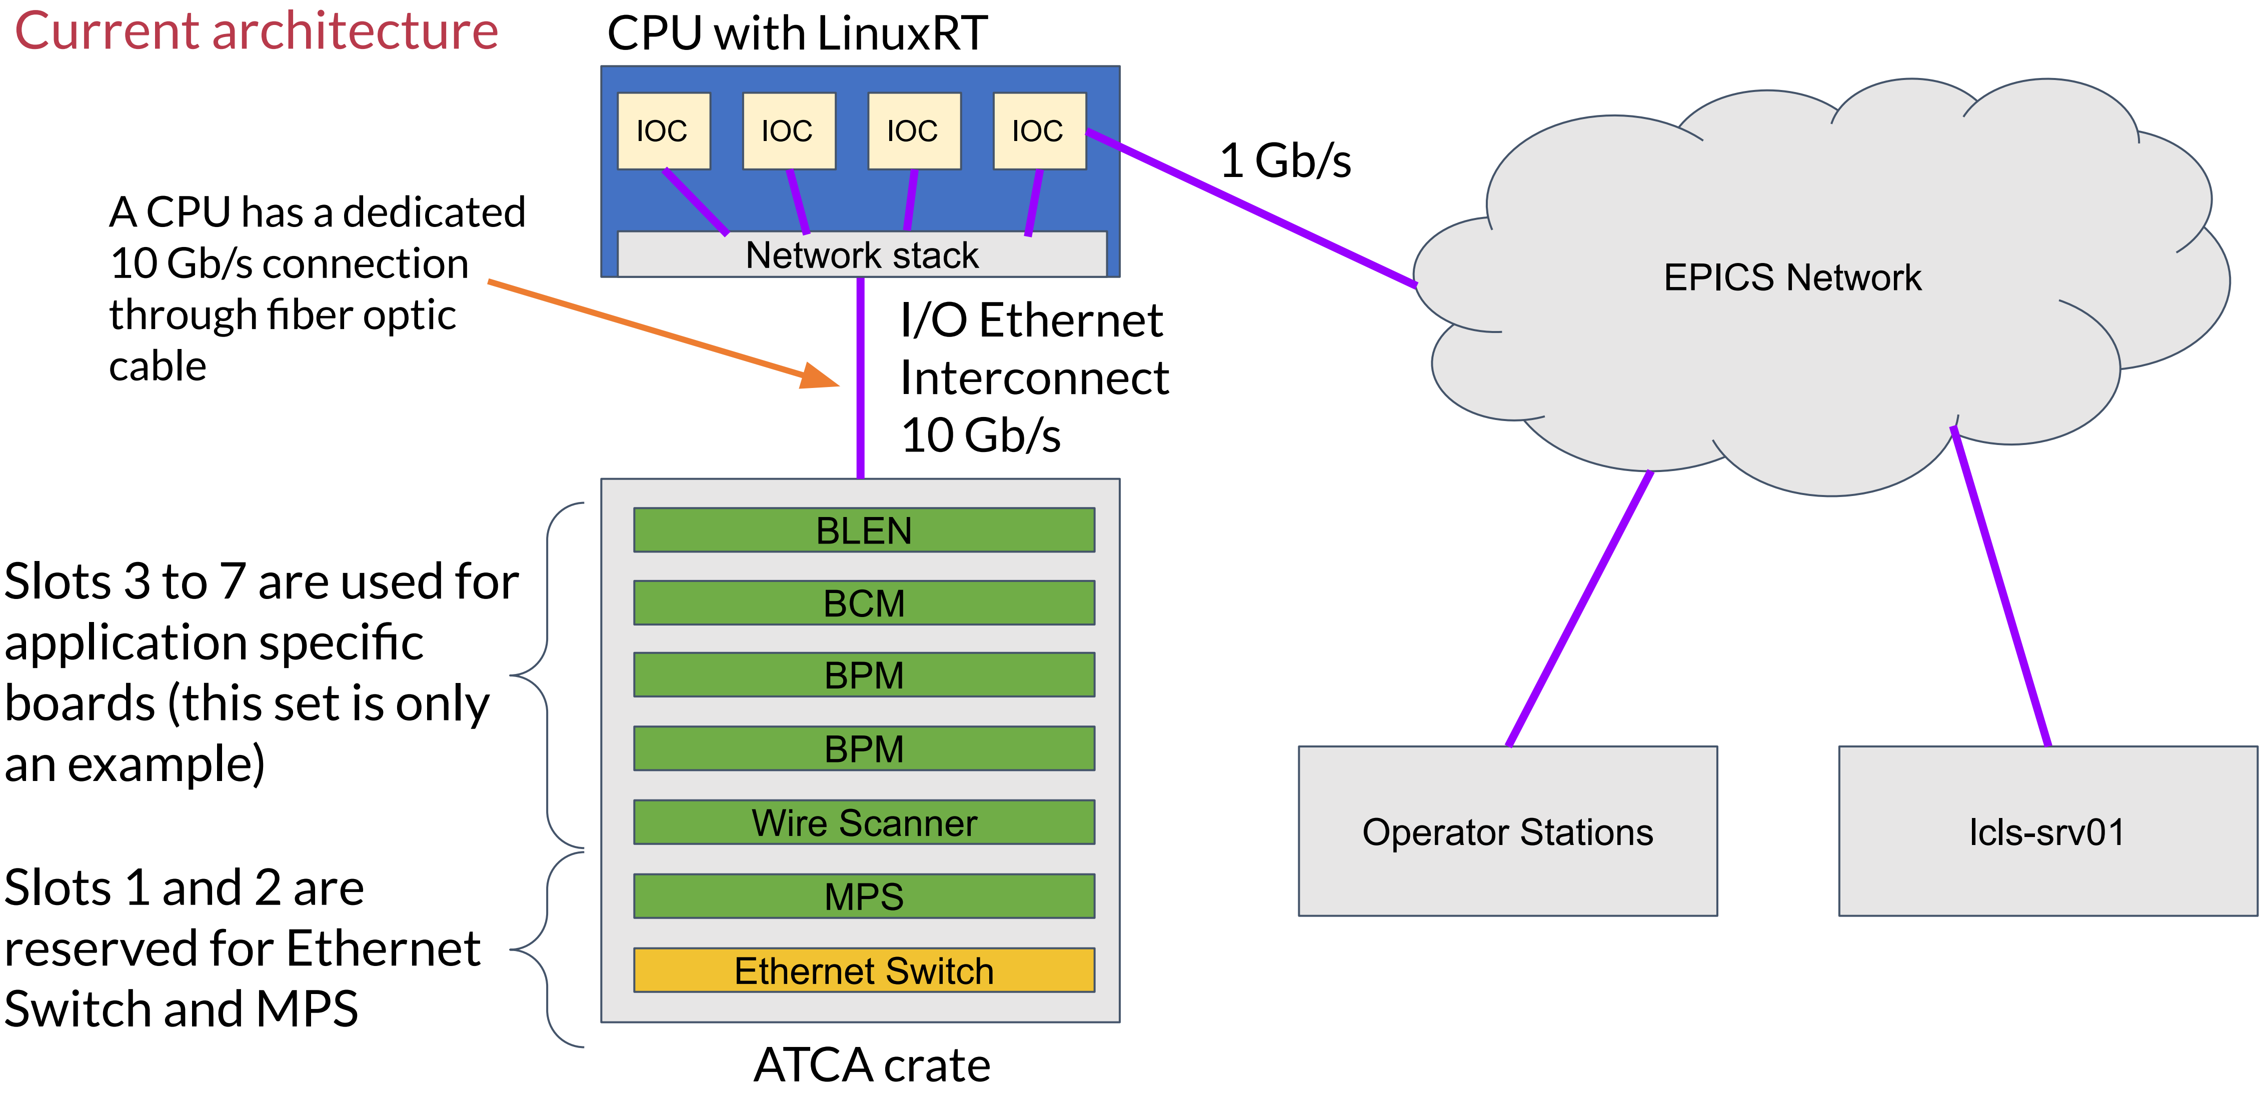

Introduction

Current architecture

Different valid ways of connecting CPUs and ATCAs

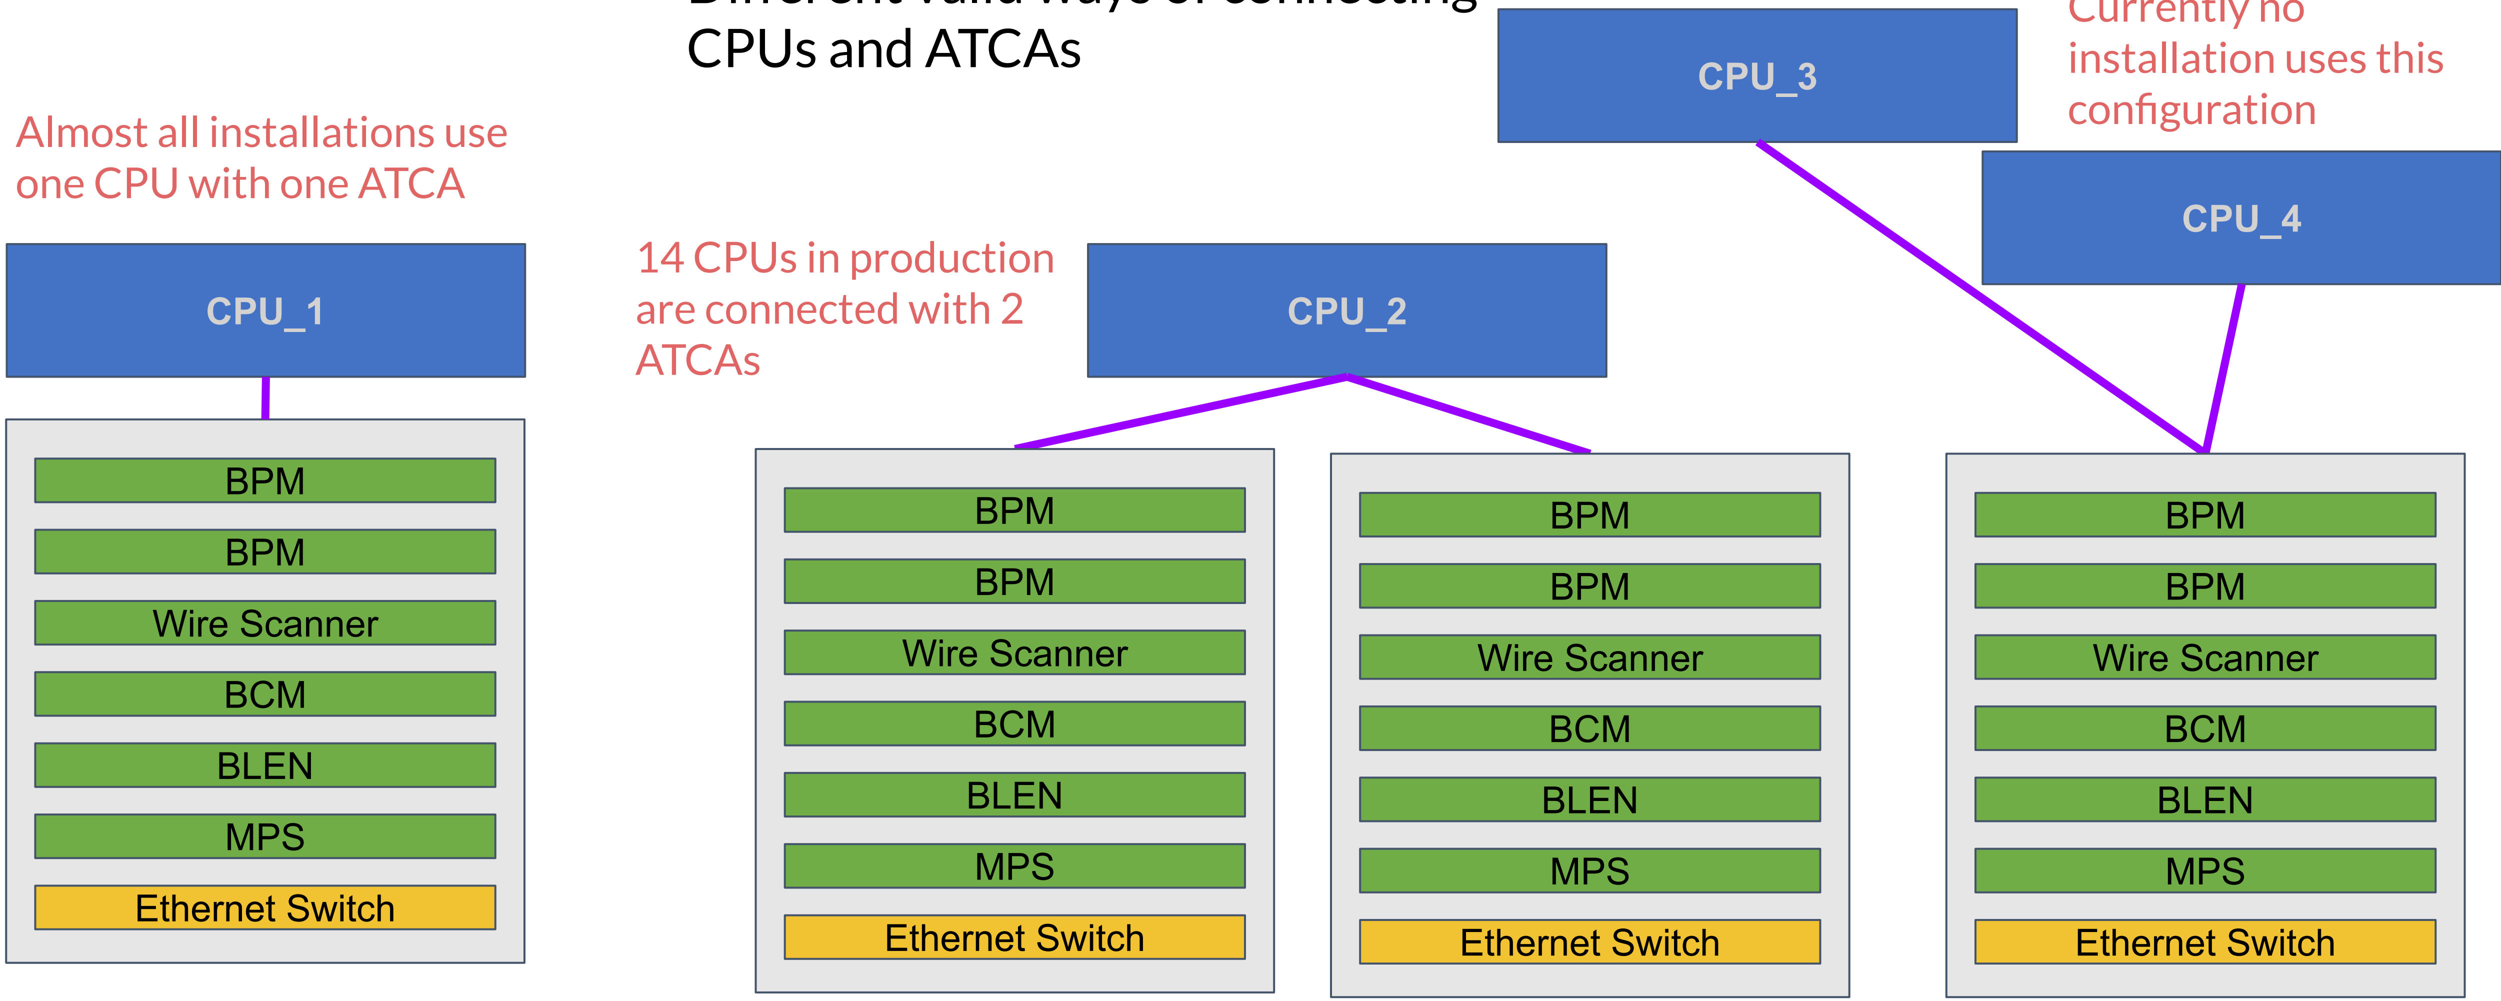

2

Analysis of the implemented fixes

Analysis of the implemented fixes

IOCs were crashing

Even after testing using a test stand with a testing application, when in production:

CATERS:

- 157205 - BPM IOC crash at startup after module update.
- 158662 - cpu-l0b-sp02 MPS and BPM IOC crashes.
- 160141 - ATCA BPM IOCs occasionally crash with error CPSW Error: No response -- timeout.
- 145695 - Common Platform BPM IOCs sometimes stop receiving data from firmware.
- 143545 - ATCA Cavity BPM IOCs affect others in the CPU.

A thorough technical explanation with the investigation is available here: https://confluence.slac.stanford.edu/x/GqiREg

Analysis of the implemented fixes

Network stack busy

During a high transfer of data, like reading from Fault Buffers, all carrier boards will send data to the CPU at the same time.

Backpressure:

- the CPU usage by the Kernel increases above a limit.
- IOCs start to get delayed data.
- IOC “believes” that a connection was broken and quit.

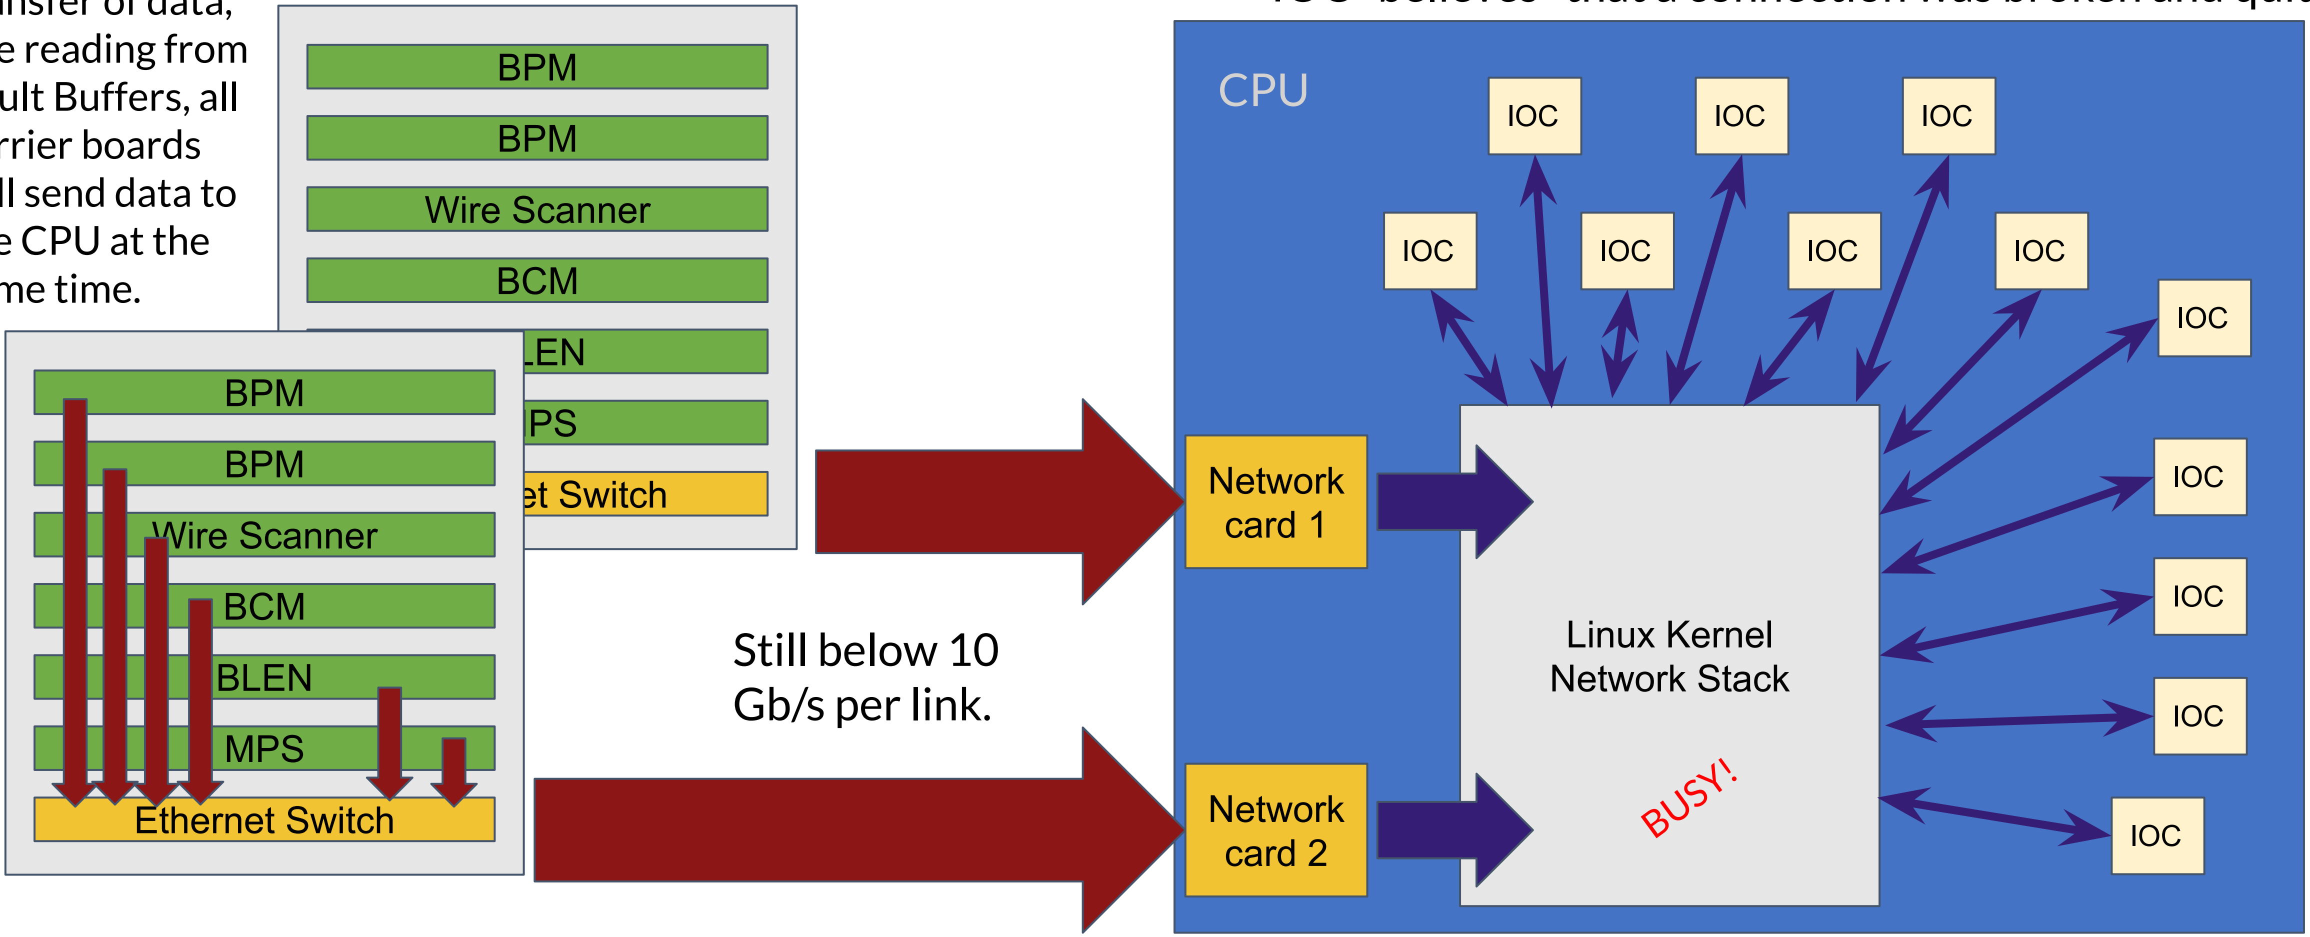

Analysis of the implemented fixes

Network stack busy

One of the fixes was to reduce the amount of data when receiving Fault Buffers. The trade off is that now the complete fault buffer will take several seconds to arrive.

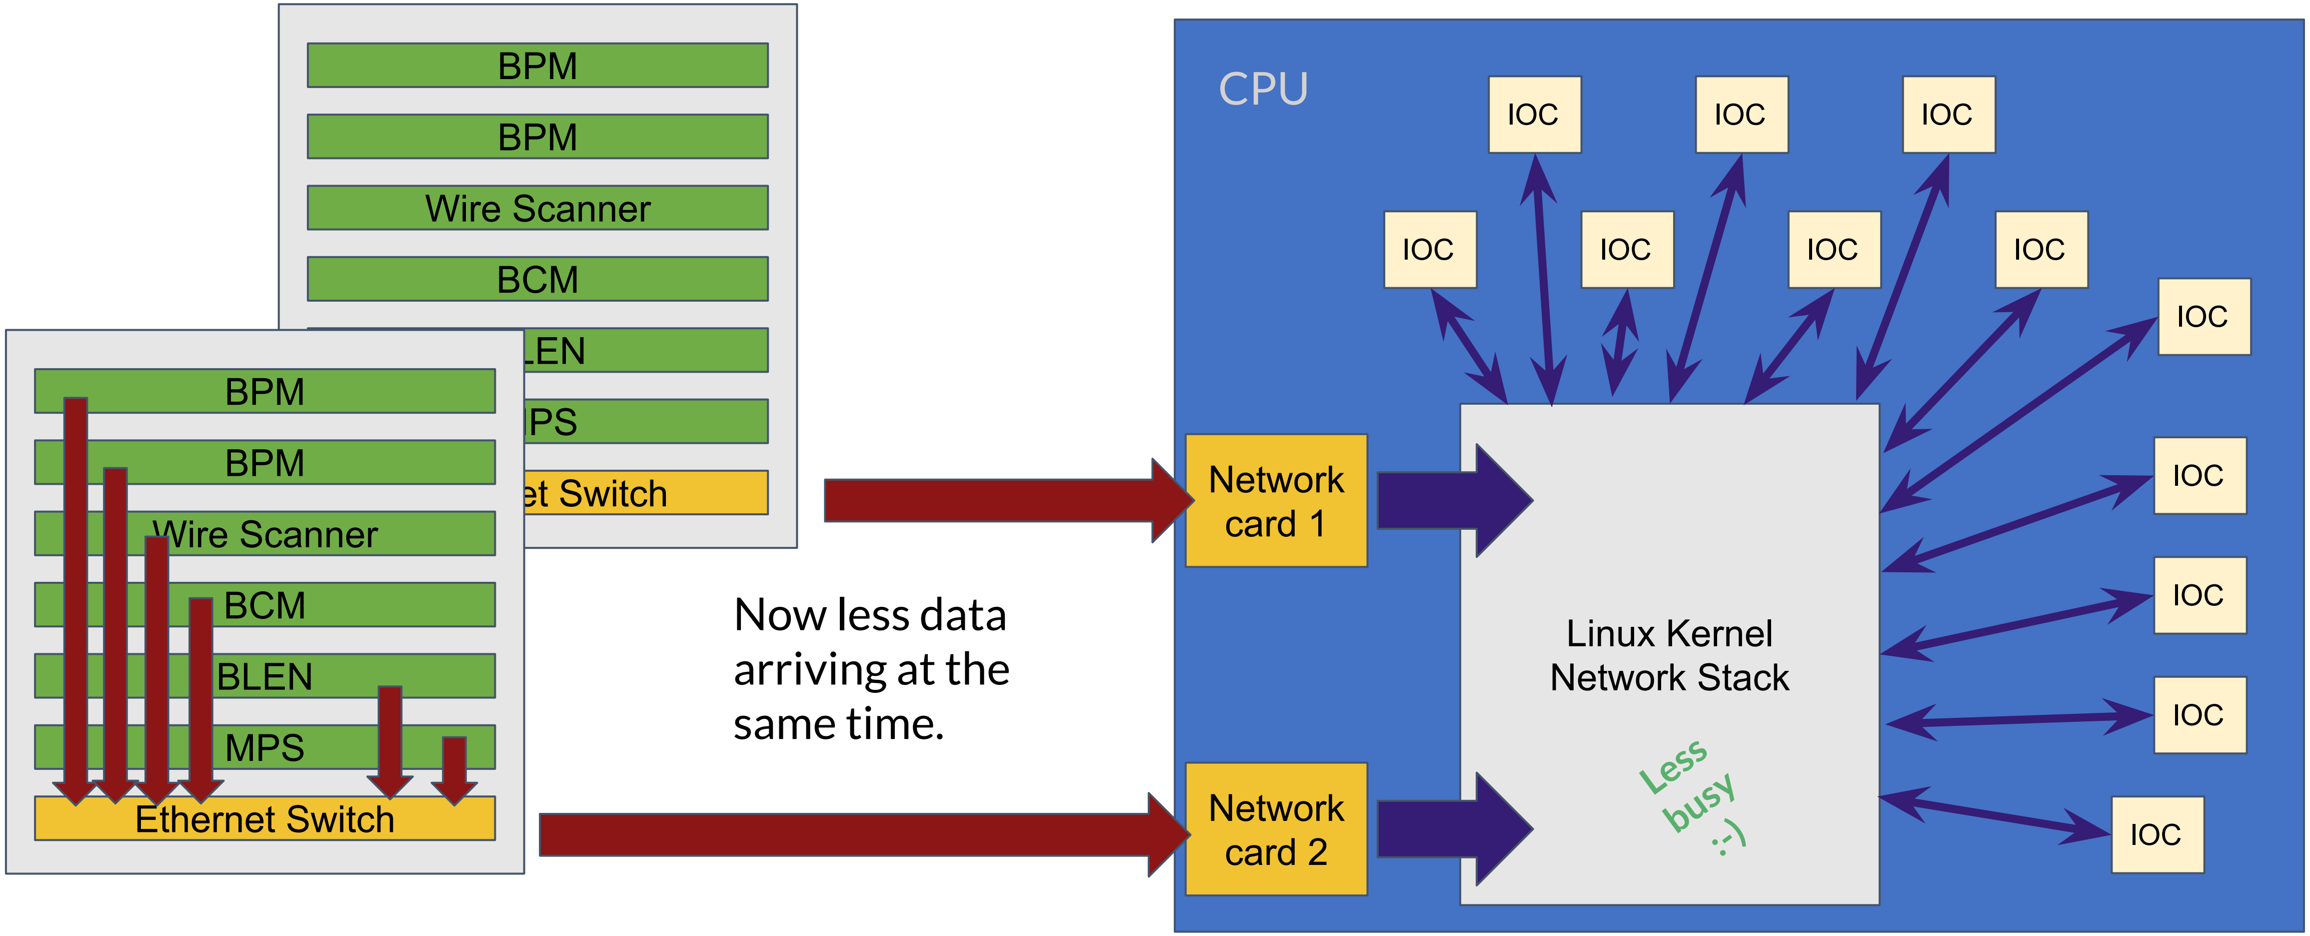

Analysis of the implemented fixes

Latency

Latency plays a big role when asking for data piece by piece. The time between every ask and answer sums up.

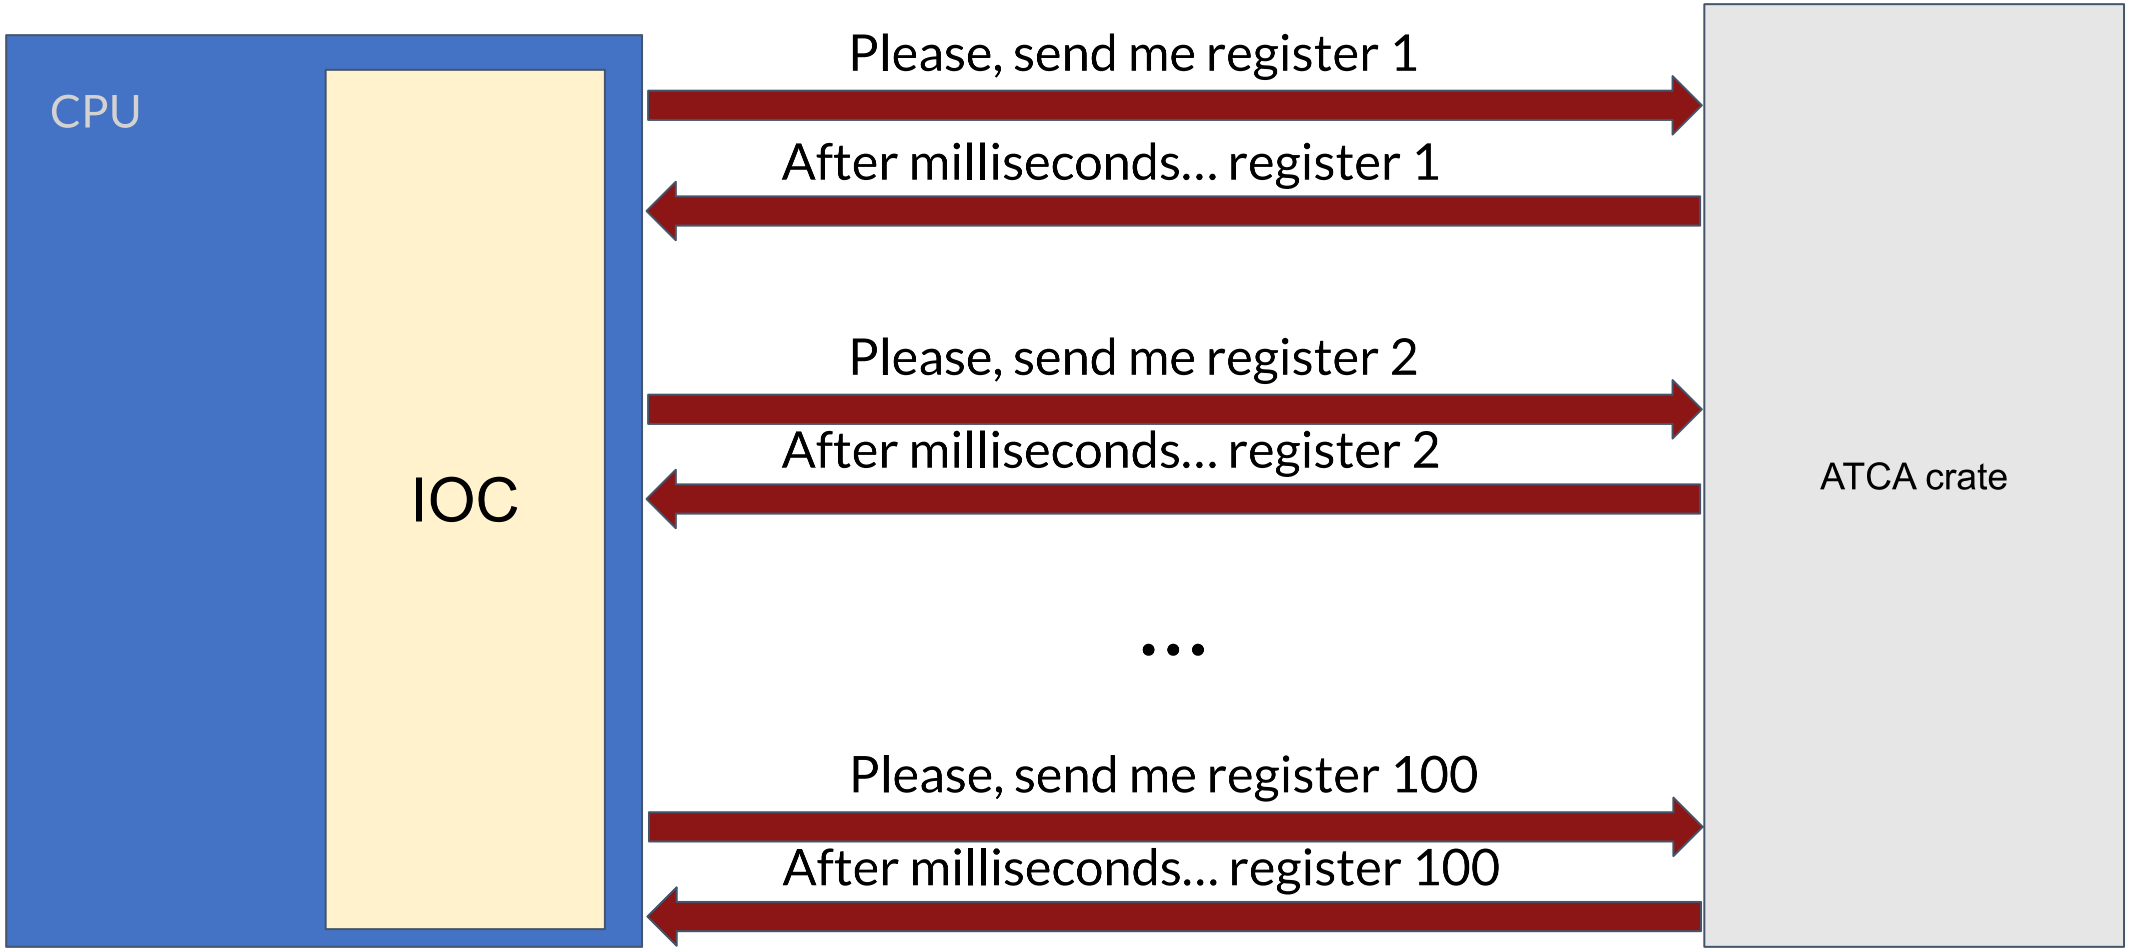

Analysis of the implemented fixes

Latency

A fix for TPG was to ask for multiple registers at the same time. One single answer has only one unit of latency.

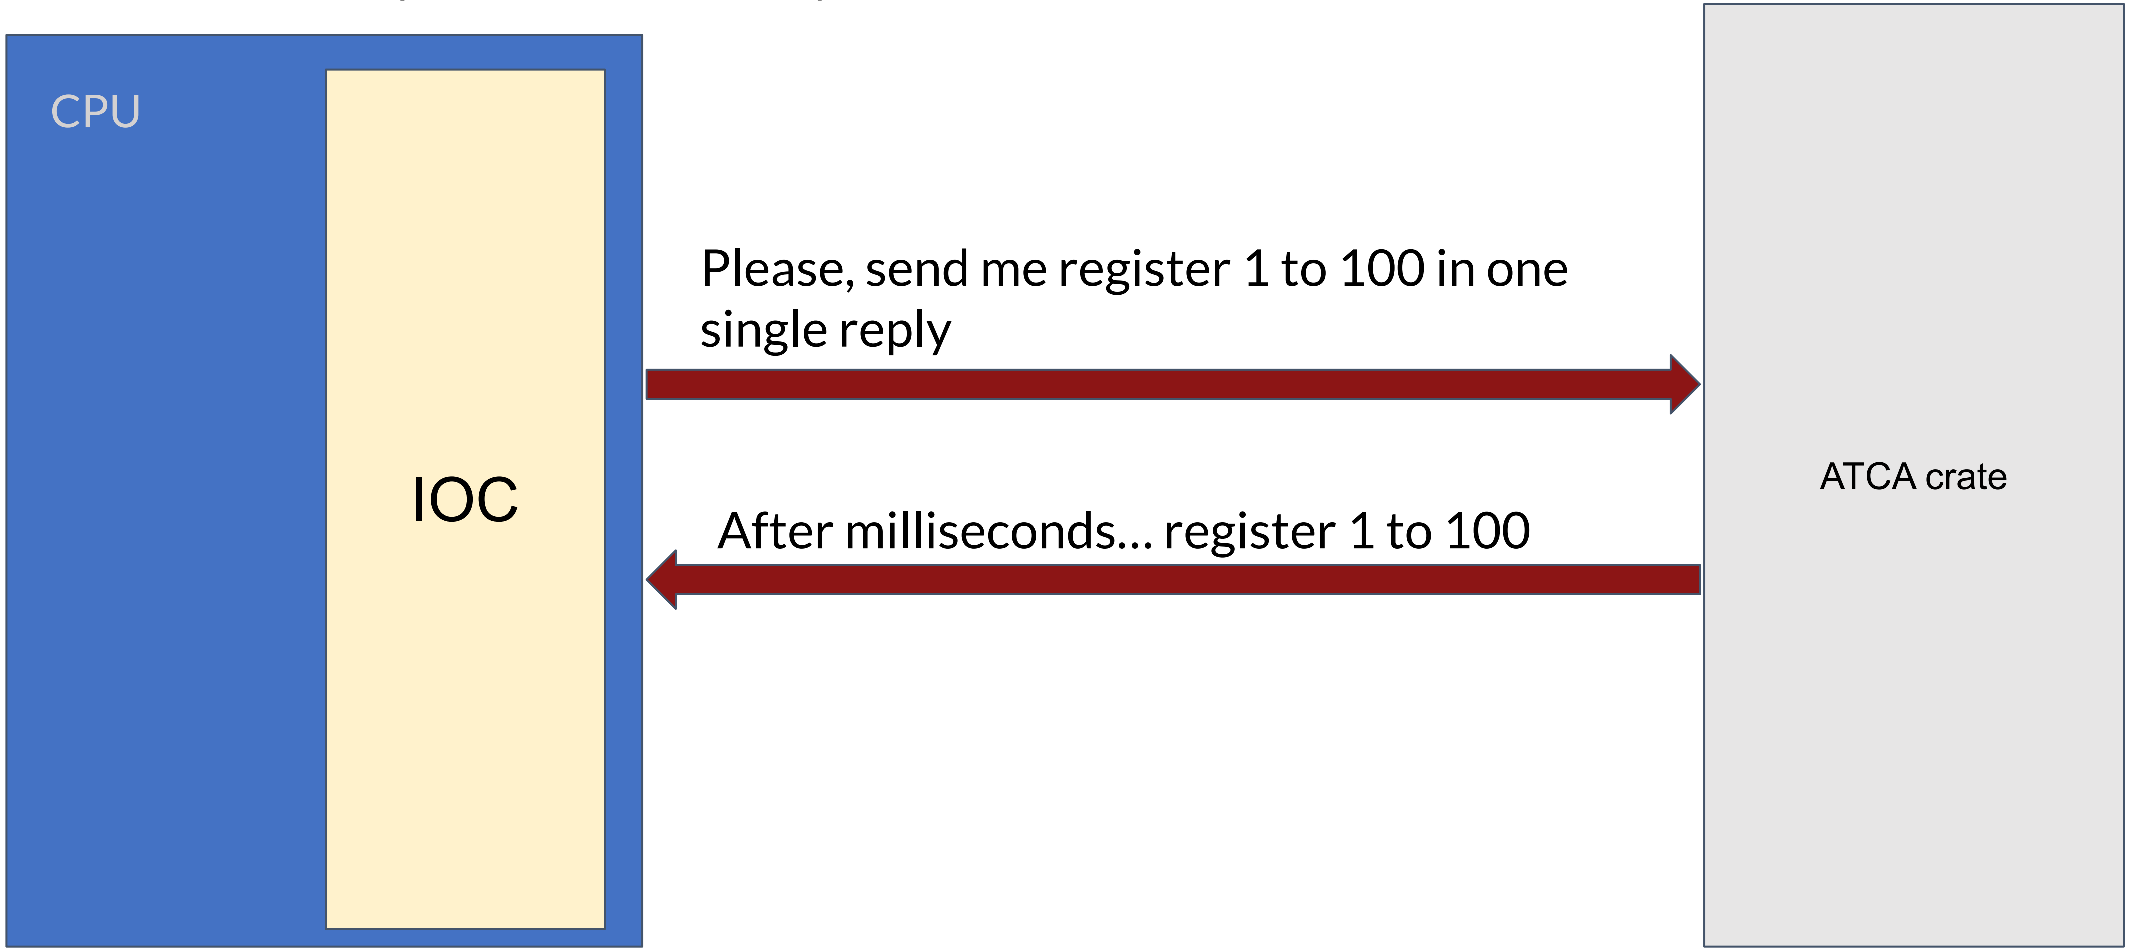

Analysis of the implemented fixes

Other fixes

Changes on the software structure brought both CPU and memory usage down. Less CPU usage means that the network stack has more time to deal with packets coming from the network.

Bugs were found in software and firmware and fixed. One of the bugs was responsible for the so called “BSA buffer freezing”.

3

Improvements

Improvements

Implemented since December 2022

BSA buffer allocation changes: more system buffers, less user buffers

- 29 user buffers.
- 15 system buffers.

Merging BSSS timing filtering control with BSA

- Previously BSSS supported fewer timing filters than BSA. Now there's a perfect match between BSSS and BSA.

NORD now can be used to check up the progress of BSA and Fault buffer.

SC_BSA_ENABLE PV

- It can be used to turn off SC BSA mode when operating with NC beam.

Improvements

Implemented since December 2022

README was completely refactored

- Concepts of what is each acquisition service are described.
- Each acquisition service has its own README now.
- Example of an st.cmd to instantiate all acquisition services.
- Please, read them!

Massive change in PV names after a long work with the naming Czars.

Improvements

Rethinking the strategy

Additionally to the serious problems that we were observing in production, which needed months to fix, we had a list of improvements to implement.

The implementation of the improvements was made at the same time as the problems were being diagnosed and fixed.

A better approach would have been to bring the system to stability and, only later, to work on the new features:

- New features can distract the focus on the diagnostic because it is the same person who implements both.
- It may be difficult to tell if a test with negative results was due to a previous problem or a new one created with the improvement implementation.

4

Stress tests performed

Stress tests performed

Ready to test

All fixes and improvements were tested in test stands, using TPG or a test application.

To make sure all problems were solved, we needed a more complete scenario.

Sonya, Jeremy, and Namrata agreed to lent us one CPU and 2 ATCA crates in production so a better test scenario could be used. Thanks!

In the future we need to think on a test stand similar to what we used in production so we don't depend exclusively on installed equipment.

Stress tests performed

Focus of the tests

11 IOCs running in cpu-l0b-sp02 (production area):

- 2 MPS
- 8 BPMs
- 1 Wire Scanner

4 types of stress test running at the same time:

- Fault buffers
- BLD
- BSAS
- camonitor

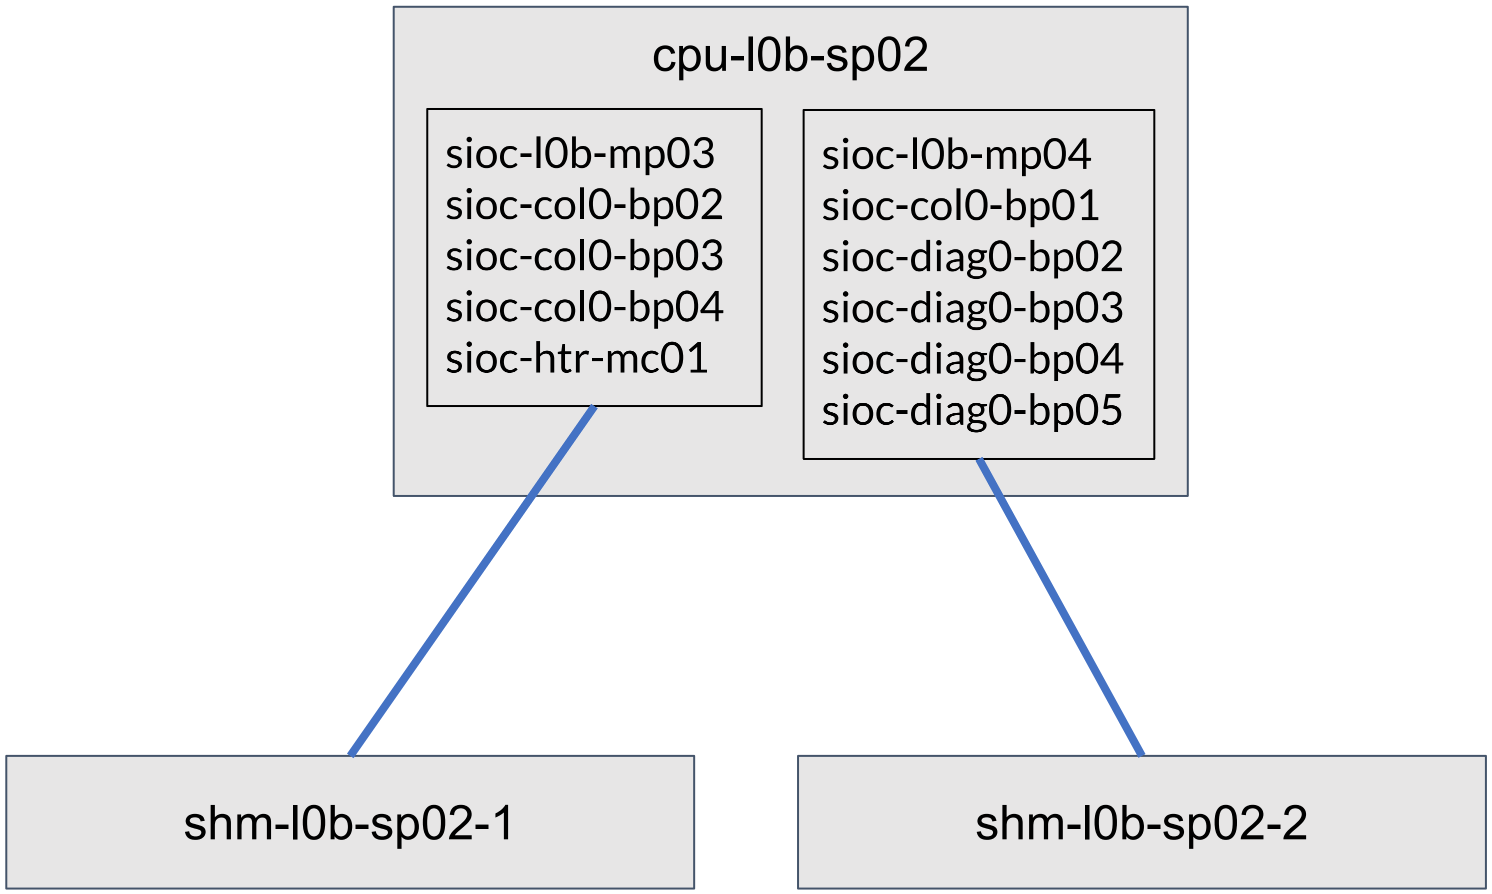

Stress test with fault buffers

The dimension of the load

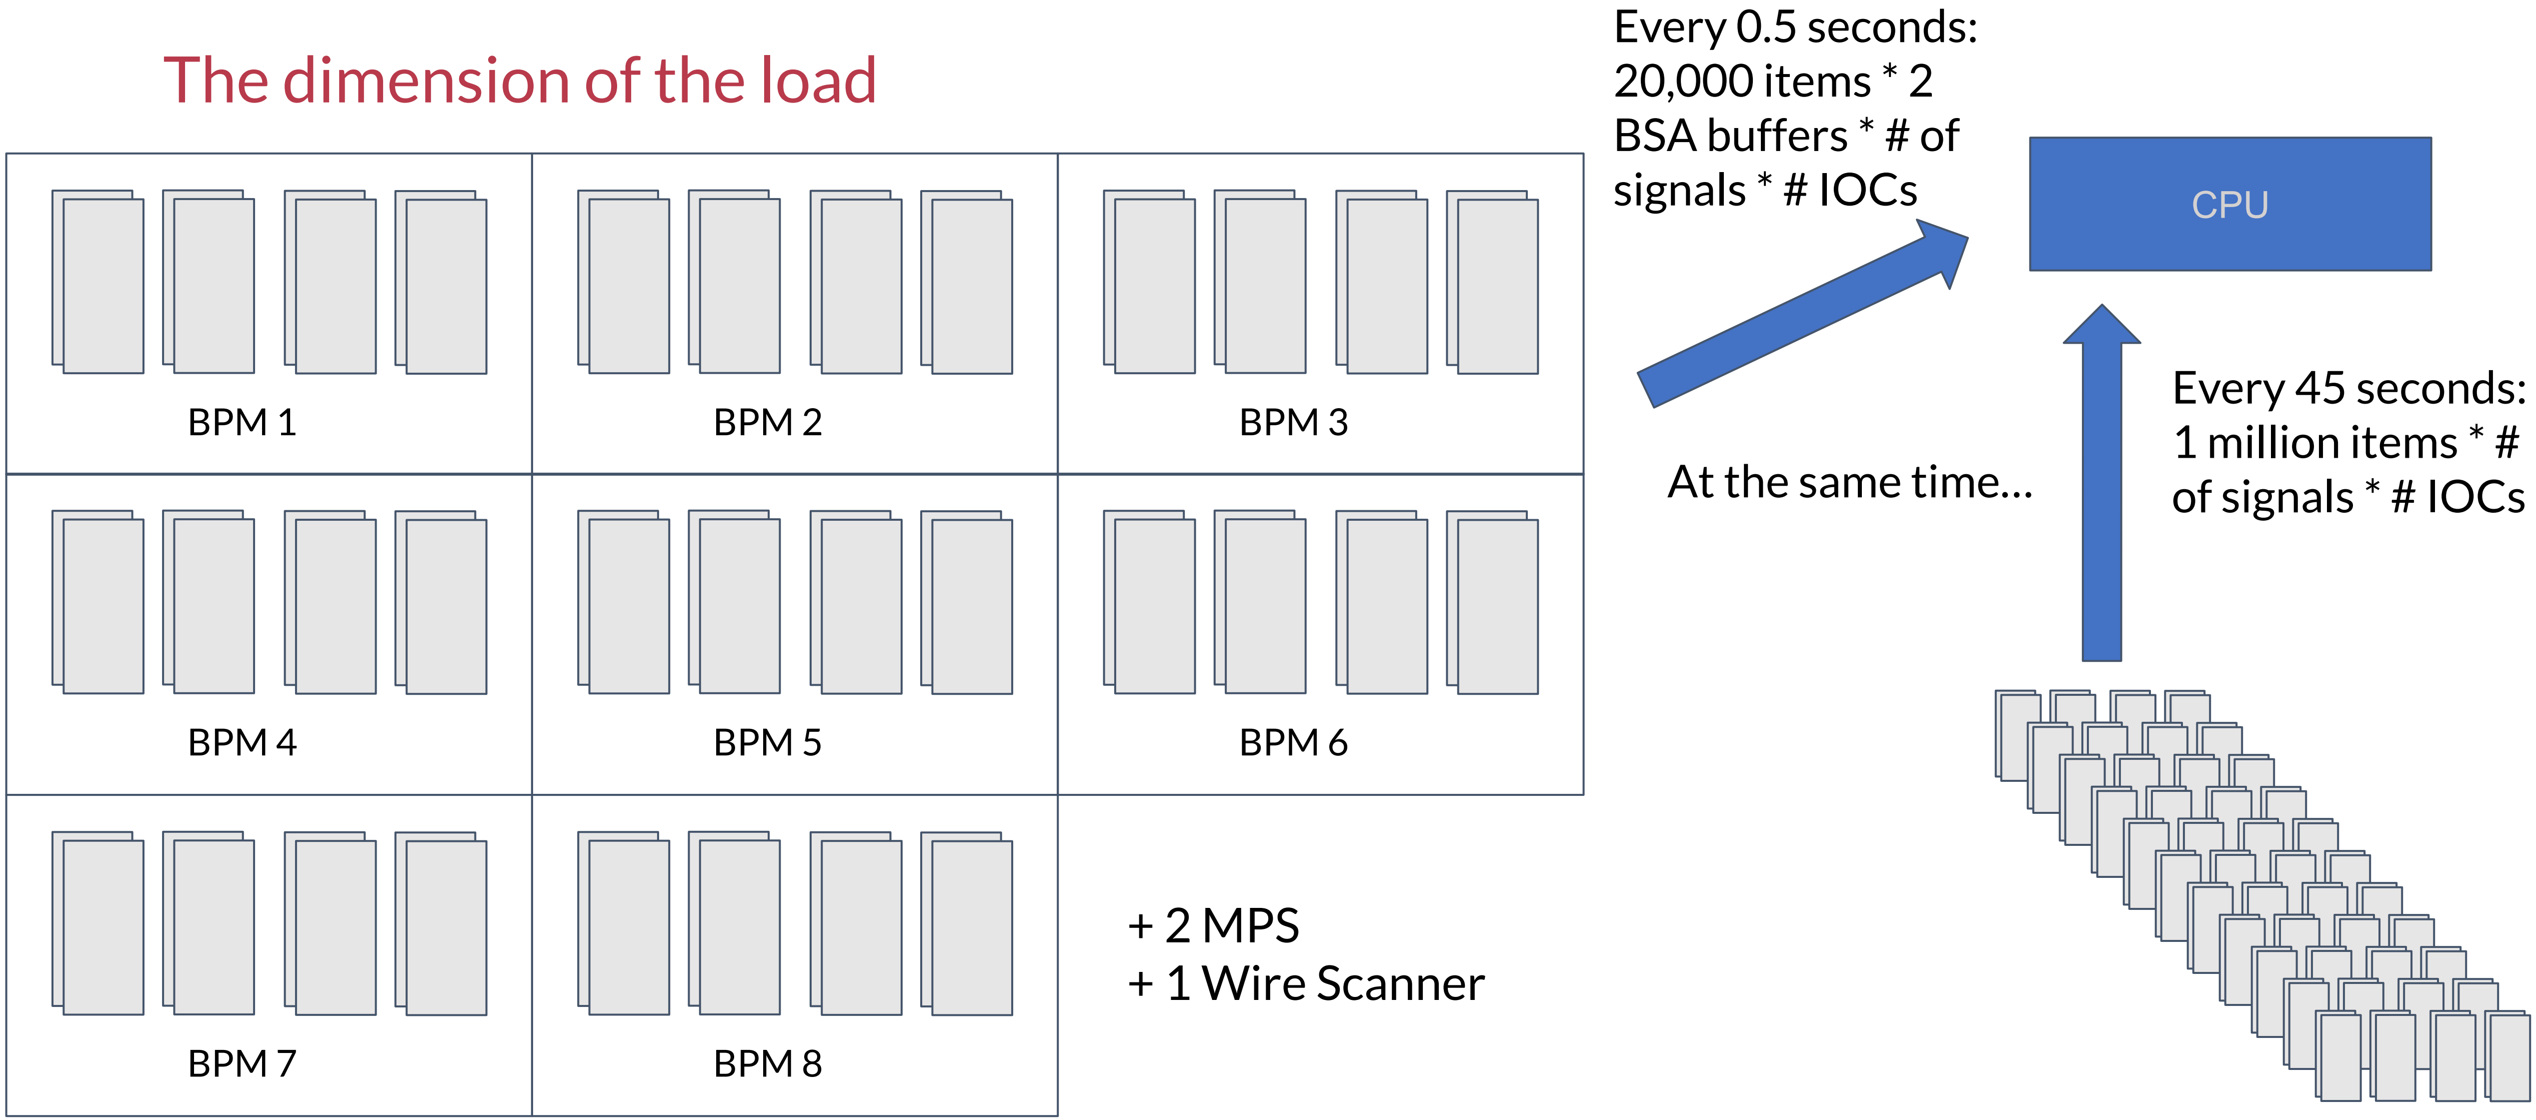

Stress test with BLD

General View

BLD doesn't show data through PVs.

The IOC sends the data directly to the network and a client at the other side can subscribe to receive the data.

BLD must run at full rate as a requirement. In other words, bunch by bunch data needs to go to the network.

This requirement is already challenging enough for a stress test, so no additional scenario was tested beyond the default.

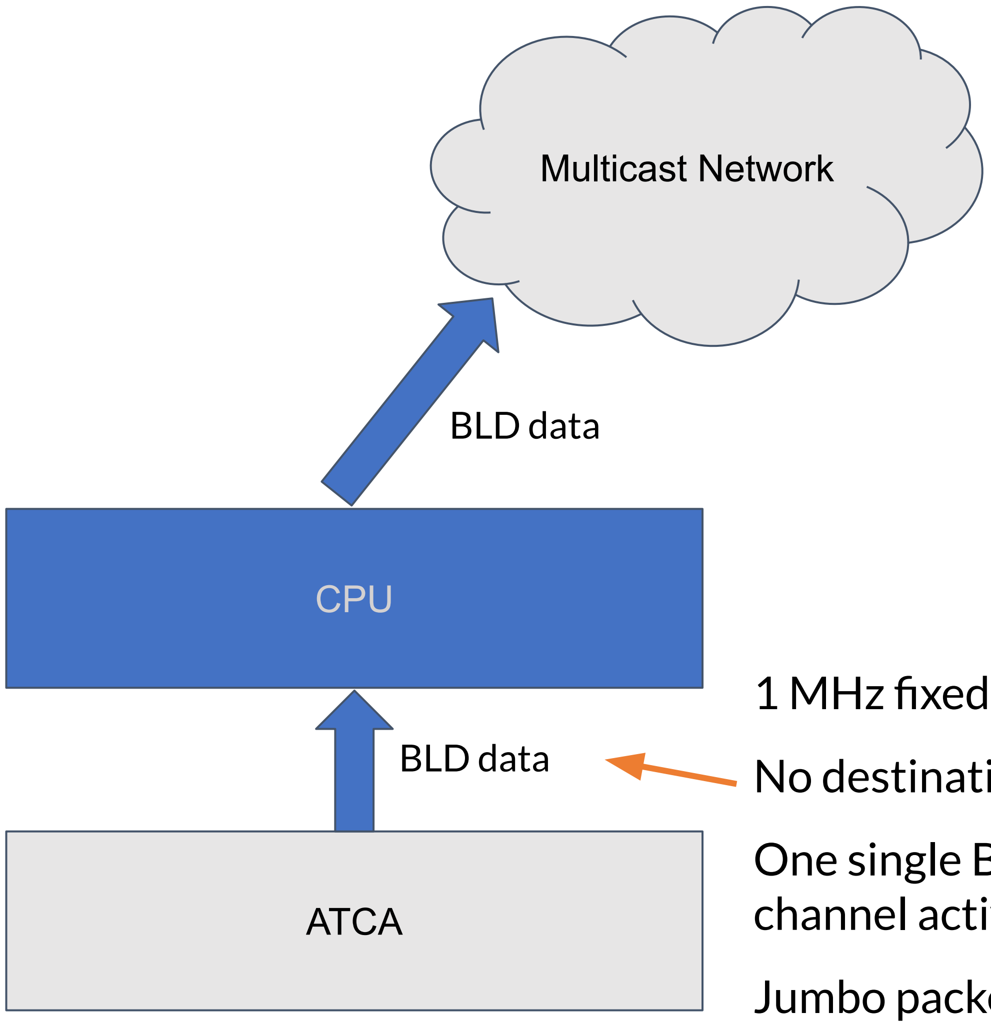

1 MHz fixed rate

No destination

One single BLD channel activated

Jumbo packet sizes

Stress test with BSAS

General View

BSAS provides 4 tables per carrier board. Each table is related to one destination:

- DIAG0
- BSYD
- HXR
- SXR

The number of columns of each table is proportional to the number of signals exposed.

The number of rows can be configured.

secondsPastEpoch	nanoseconds	pulseId	X.CNT	X.VAL	X.AVG	X.RMS	X.MIN	X.MAX	Y.CNT	Y.VAL	etc.

External network load with camonitor

Python script to handle camonitor

A Python script was created to open thousands of camonitors

- Approximately 2400 waveform PVs with 20,000 elements each.
- Approximately 600 waveform PVs with 1 million elements each.

5

Results of the test in production

Results of the test in production

Adding one IOC at a time

11 IOCs use approximately 93% of the CPU (including all cores)

Kernel space tasks takes twice more CPU than user space ones.

This is due to the high activity in the network stack.

Still no BLD.

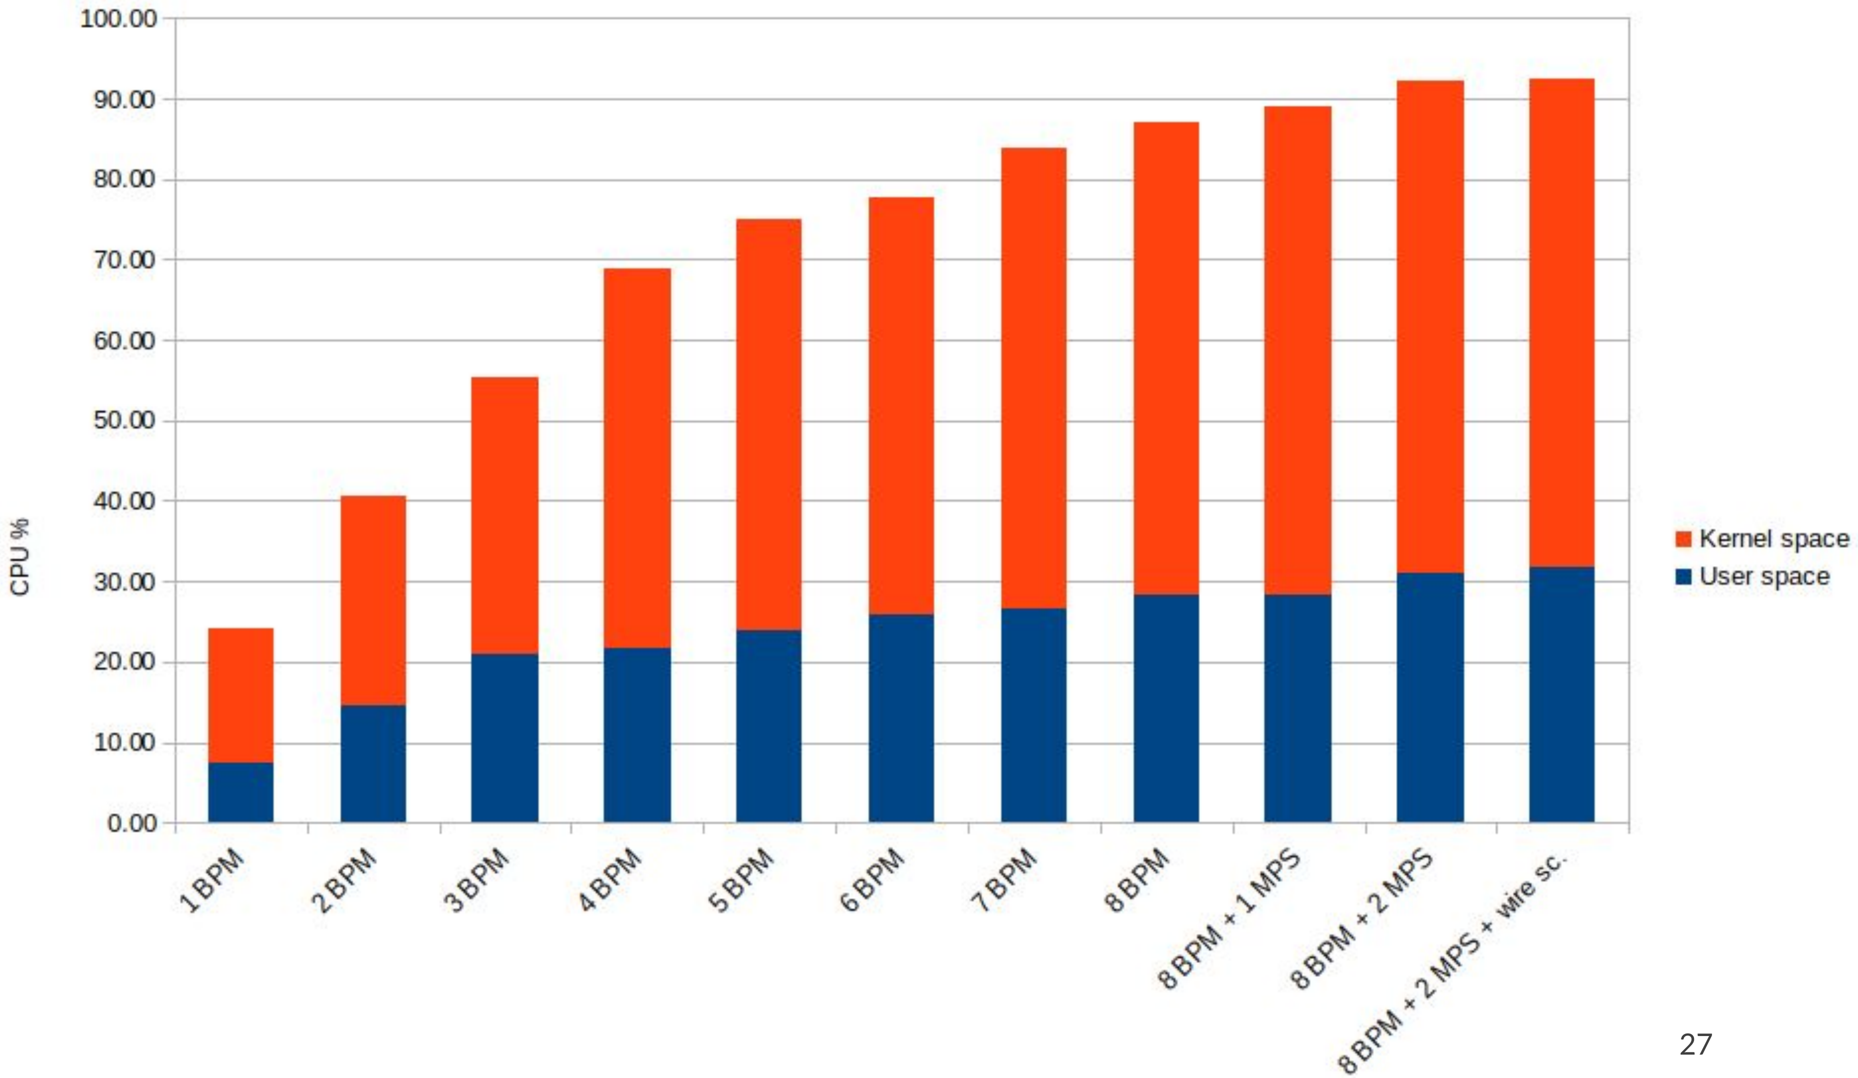

Results of the test in production

BLD has high load

4 IOCs with BLD:

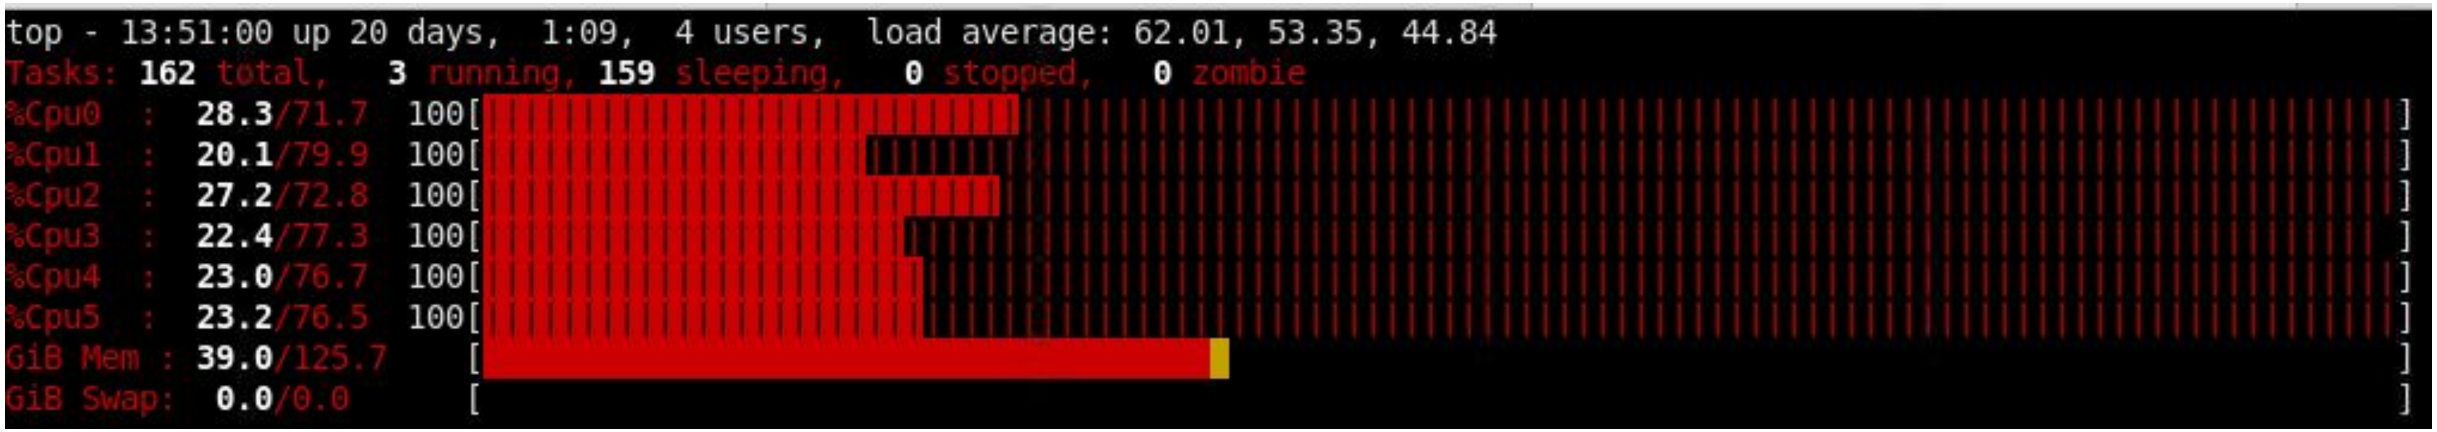

At this scenario:

- "top" takes a lot of time to start.
- ifconfig is slow.
- Autocomplete with Tab is slow.
- Accessing the IOC console is extremely slow (almost 1 minute after issuing iocConsole).
- iocManager showing heartbeat status changing: PRESENT - INTERMITTENT - ABSENT, meaning that the IOC was unable to post monitors for seconds.

Results of the test in production

Curious side effect used as a tool

We discovered that the trigger channel calculated rate is affected when the system is too busy:

```
$ caget TPR:{COL0,DIAG0}:BP0{1,2,3,4,5}:0:CH0{1,2}_RATE | grep -v onnected
TPR:COL0:BP01:0:CH01_RATE   203.5
TPR:COL0:BP02:0:CH01_RATE   111.5
TPR:COL0:BP02:0:CH01_RATE   111.5
TPR:COL0:BP03:0:CH01_RATE   173.5
TPR:COL0:BP03:0:CH01_RATE   173.5
TPR:COL0:BP01:0:CH01_RATE   205.0
TPR:DIAG0:BP02:0:CH01_RATE   226.0
TPR:DIAG0:BP02:0:CH01_RATE   226.0
TPR:DIAG0:BP04:0:CH01_RATE   279.0
TPR:DIAG0:BP04:0:CH01_RATE   279.0
TPR:DIAG0:BP05:0:CH01_RATE   153.0
TPR:DIAG0:BP05:0:CH01_RATE   153.0
```

Triggers were configured to 100 Hz, so the calculated triggers should be close to it.

Results of the test in production

Comparison between 2 ATCA versus 1 ATCA crate

1 ATCA	2 ATCAs
Maximum 2 BLD for having viable IOCs	Maximum 1 BLD for having viable IOCs
Maximum 4 BLDs before system instability	Maximum 3 BLDs before system instability
5 or more BLDs cause IOCs to crash	4 or more BLDs cause IOCs to crash

BLD takes so many resources that there's not much difference between 2 full crates and 1 full crate.

Results of the test in production

BSAS can use a lot of resources, too

BSAS running with 1000 rows didn't add a significant load compared with what was already running in the IOCs.

When configuring 10K or more rows, most of the IOCs crashed. Jeremy mentioned that he saw multiple MPS IOCs crashing throughout the accelerator, not only on cpu-l0b-sp02.

Results of the test in production

camonitor script put a challenge to lcls-srv01

Script was run in lcls-srv01.

When the script opened 3,000 camonitors for waveform records, it used 7 GB of RAM memory and 6.5% of one CPU core of the server.

Eventually lcls-srv01 crashed. We are not sure if the script was responsible for the crash, but it is a prime suspect.

When opening camonitors using PyEpics, keep in mind that it uses around 10 bytes per element of the waveform. Rule of thumb:

- BSA buffer with 20,000 elements ~= 200 kB per buffer
- Fault buffer with 1 million elements ~= 10 MB per buffer

6

Conclusions and recommendations

Conclusions and recommendations

Tests were successful!

After some back and forth with the system in production we got a stable situation.

Under extraordinary heavy load, the IOCs stayed alive for days.

The amount of heavy load that we used during the tests are expected to happen only a few times per day in real life.

Now we understand better the system boundaries.

Still a few issues to address with instances of MPS and Wire Scanner, though.

Conclusions and recommendations

BSAS

BSAS need to have a strong limitation on the number of rows as accidentally configuring it can crash dozens of IOCs in production.

Implementing the limitation on firmware is easier than in software.

Conclusions and recommendations

Plan for BLD is not feasible with current CPUs

From the current plan, these are the CPUs that will work well with BLD:

- cpu-ltus-sp01 (2 IOCs, 1 ATCA):
  - sioc-ltus-bp03
  - sioc-ltus-bp05
- cpu-ltus-sp02:
  - sioc-ltus-bp12
- cpu-ltus-sp03:
  - sioc-ltus-bp14
- cpu-gunb-sp01:
  - sioc-gunb-bp01
- cpu-fees-sp01 (2 IOCs, 1 ATCA):
  - sioc-fees-gd01
  - sioc-fees-gd02

Conclusions and recommendations

Plan for BLD is not feasible with current CPUs

These are the CPUs that can't hold the planned BLDs:

- cpu-unds-sp03 (4 IOCs, 2 ATCA). With 2 ATCAs only one IOC with BLD is possible.
    - sioc-unds-bp02
    - sioc-unds-bp13
    - sioc-unds-bp14
    - sioc-unds-bp15
- cpu-bc1b-sp01 (2 IOCs, 2 ATCA). With 2 ATCAs only one IOC with BLD is possible.
    - sioc-bc1b-bl01
    - sioc-bc1b-bp02
- cpu-l2b-sp03 (2 IOCs, 2 ATCA). With 2 ATCAs only one IOC with BLD is possible.
    - sioc-bc2b-bp01
    - sioc-bc2b-bp02

Conclusions and recommendations

Plan for BLD is not feasible with current CPUs

To address the problem, 2 approaches can be used:

- Add more CPUs and distribute the IOCs between them.
- Use more powerful CPUs.

(more data on equipment cost will be shown later)

The problem happens when we have full rate.

We still have time to address the problem as 1 MHz beam is far away in the future.

No tests were made to measure the impact of BLD with intermediary rates, like 1 kHz or 10 kHz.

Conclusions and recommendations

Plan for BLD is not feasible with current CPUs

If upgrading the system, we have current quotes with 2 options:

- Option 1:
  - SKY-8101 ($4,440) Gen2 Silver 4210R 2.4GHz Max 3.2GHz 10 Cores, 2x 64GB DDR4
  - $4,440 x (14+1 Backup) = $66,600
- Option 2:
  - SKY-8101 ($6,580) Gen2 Gold 6230R ,2.1GHz Max 4.0GHz 26 Cores, 2x 64GB DDR4
  - $6,580 x (14+1 Backup) = $98,700

Conclusions and recommendations

Acknowledgements

Kukhee Kim
Matt Weaver
Marcio Donadio
Larry Ruckman
Carolina Mattison
Mike Skoufis
Sonya Hoobler
Jeremy Mock
Namrata Balakrishnan
Kyle Leleux