z/TPF Dynamic CPU Capacity Support Preview

Mark Gambino
z/TPF Chief Architect

IBM **z/TPF**
April 3rd, 2017

Disclaimer

Any reference to future plans are for planning purposes only. IBM reserves the right to change those plans at its discretion. Any reliance on such a disclosure is solely at your own risk. IBM makes no commitment to provide additional information in the future.

Dynamic CPU Capacity Hill Statement

Dynamic CPU Capacity

The service provider VP of operations can run more workload on his existing CPU hardware to lower costs and immediately increase CPU capacity to achieve SLAs.

Sub Hill One

The service provider can handle a sustained increase in workload without needing to take an outage.

Sub Hill Two

The service provider can maximize CPU resources to lower its hardware costs.

Sub Hill Three

The service provider can selectively run utilities even during peak volumes without impacting real-time transactions.

Let's talk about…

Timmy the capacity planner projects that because of transaction growth, his z/TPF production system will need 1-2 more engines in the coming year, but does not know exactly when the increased rate of transactions will occur.

As-Is Scenario

An additional 2 engines need to be purchased and added during a scheduled outage, which is months before they will be needed, but that’s the scheduled outage window.

Timmy is also hoping the business forecasts are accurate and that an additional 2 engines will be sufficient because if not, there might not be enough capacity to handle the workload and adding even more capacity will require another outage.

To-Be Scenario

With Dynamic CPU support, additional engines can be added (but not used or paid for until used) during the next scheduled outage.

Just in case the business forecasts are wrong and 2 engines will not be sufficient, Timmy adds 4 additional engines to z/TPF because there are no negative ramifications in doing so.

When workload grows to the point requiring an additional engine, Micky the z/TPF operator can turn on the additional engine(s) immediately.

As-Is Environment

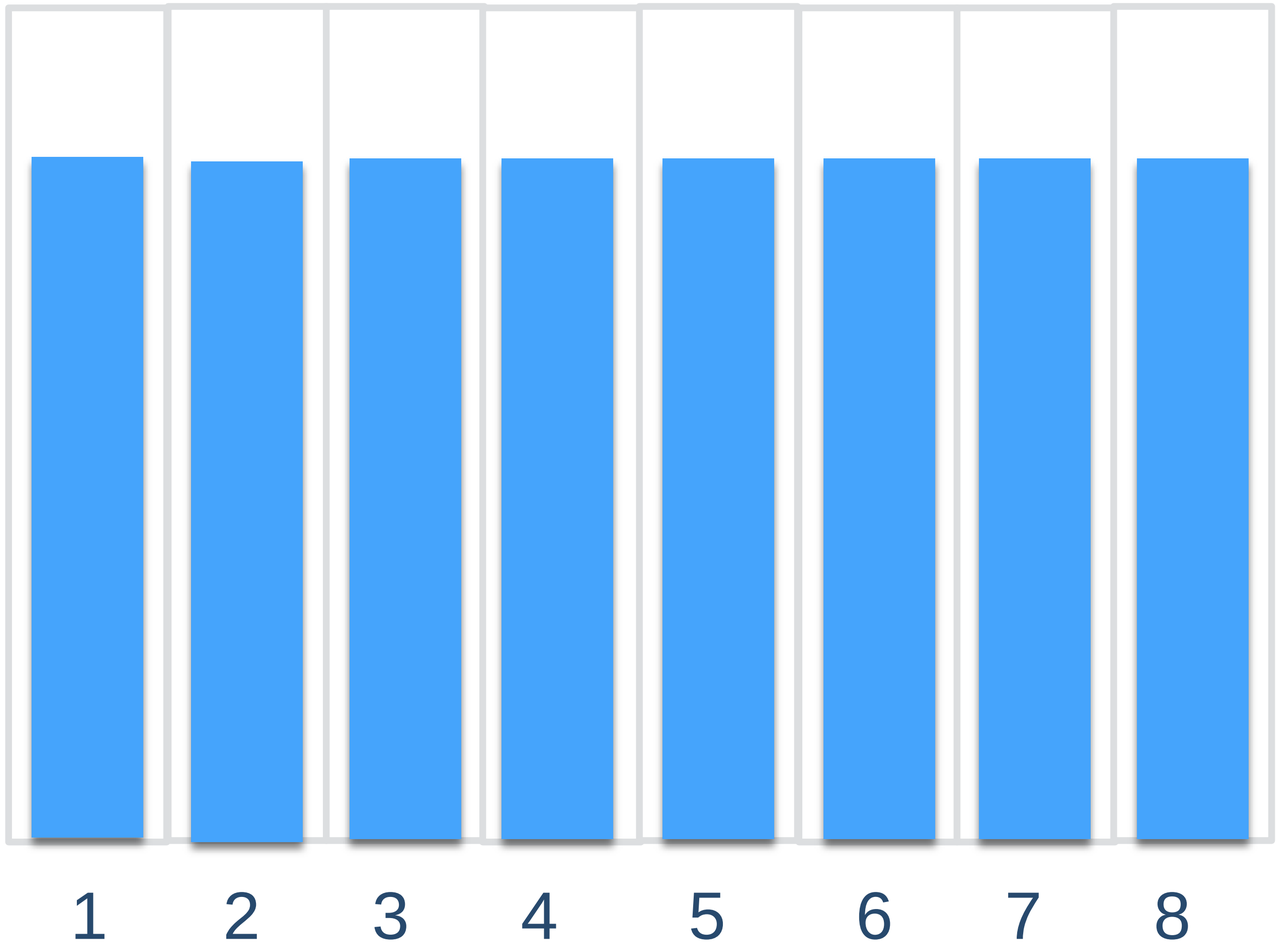

z/TPF LPAR

To-Be Environment – Define the I-stream Cap

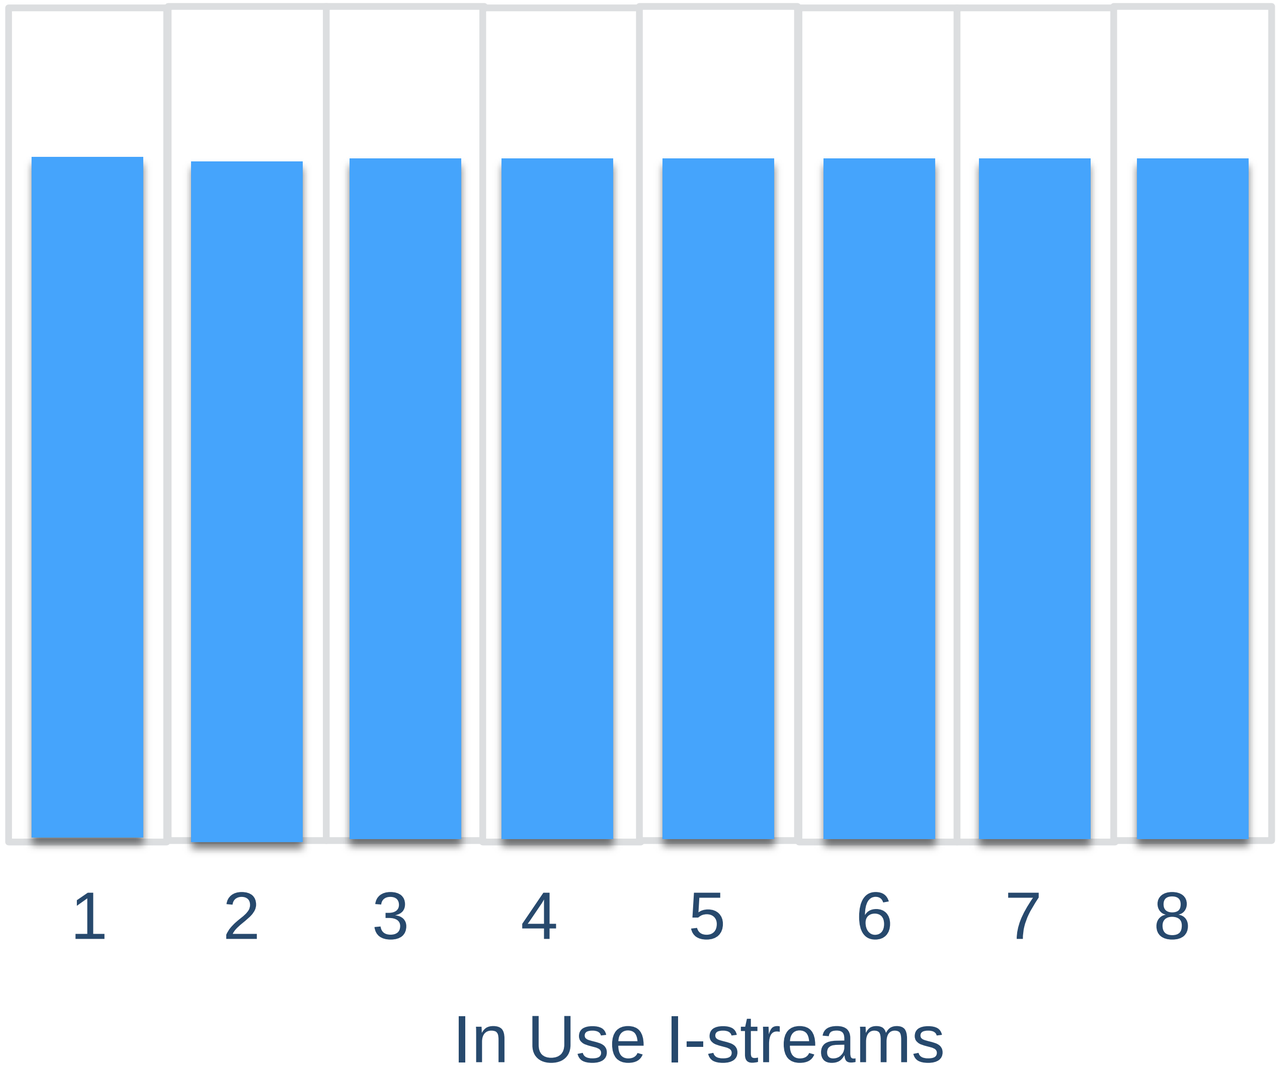

z/TPF LPAR

I-stream cap = 8

To-Be Environment – Reconfigure LPAR to Add Additional Engines and IPL z/TPF

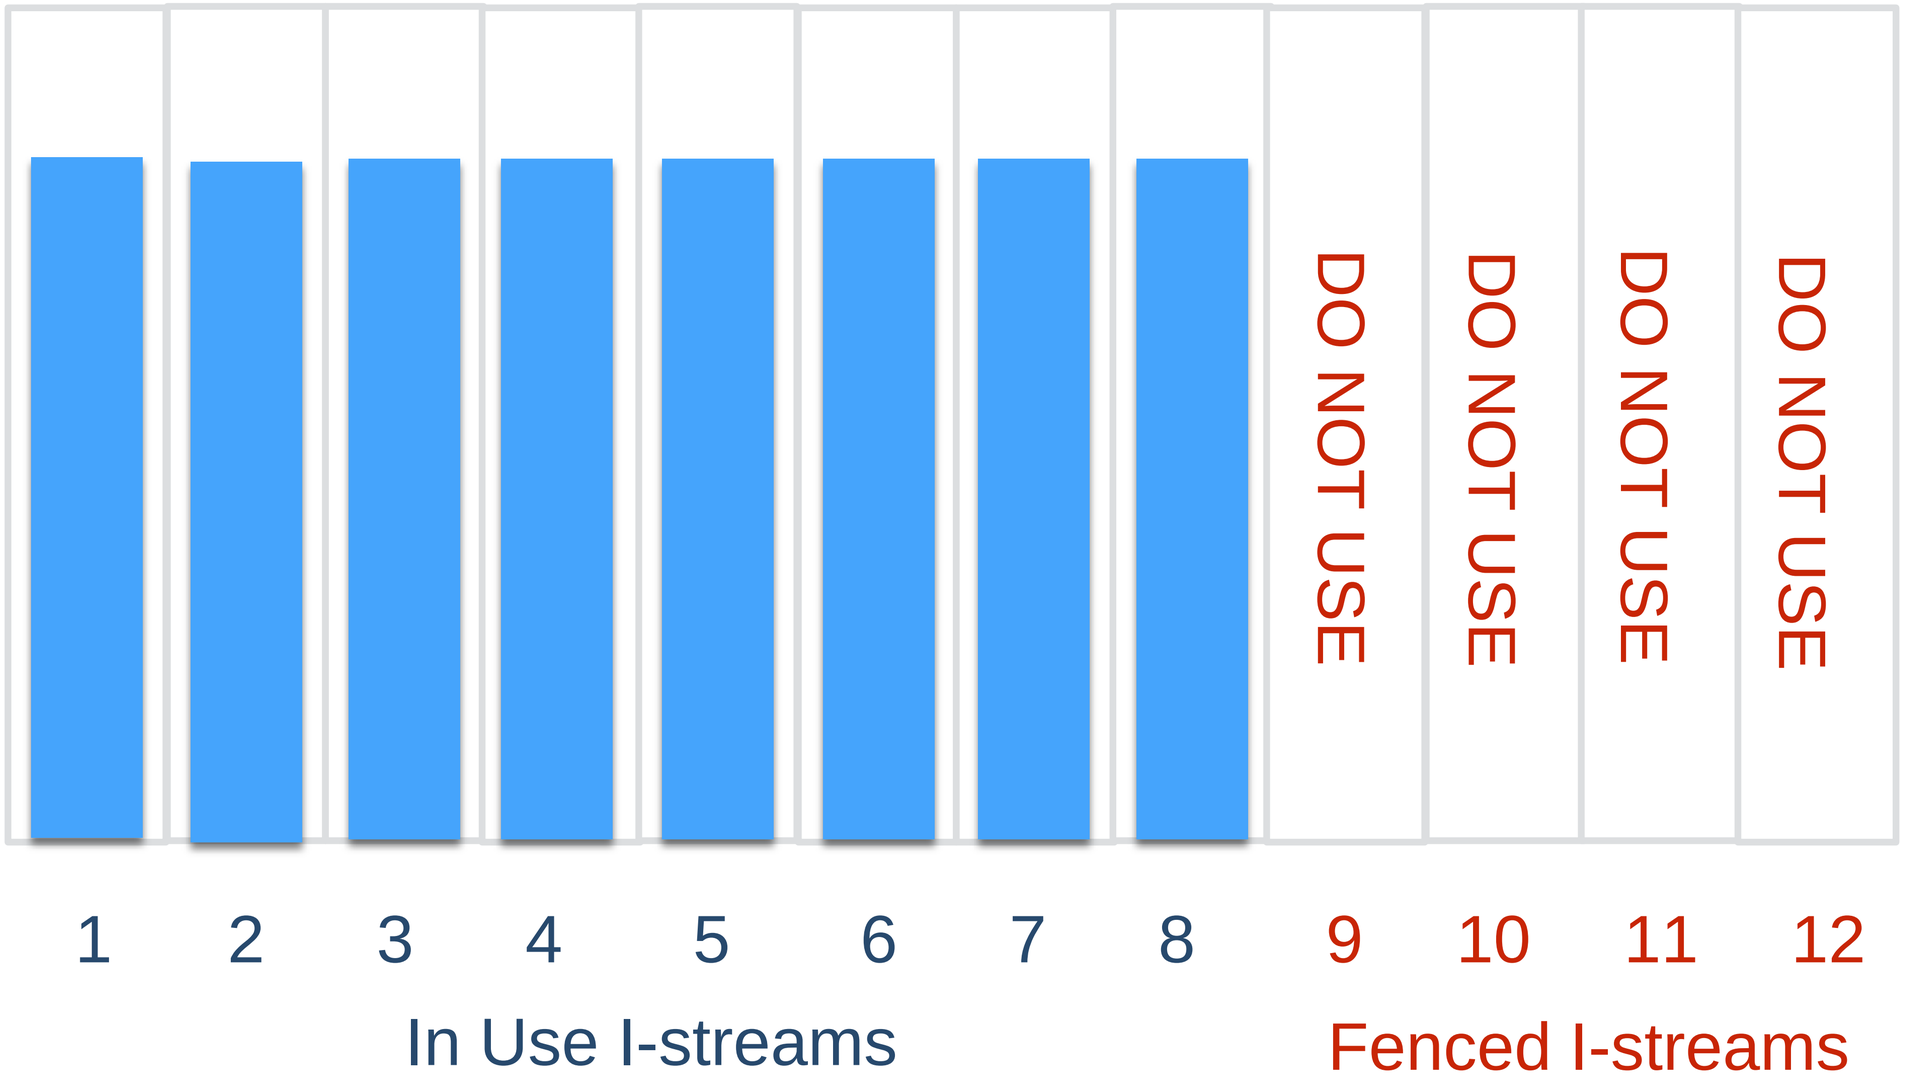

z/TPF LPAR

I-stream cap = 8

To-Be Environment – When Workload Grows, Increase I-stream Cap to Immediately Start Using the 9th Engine

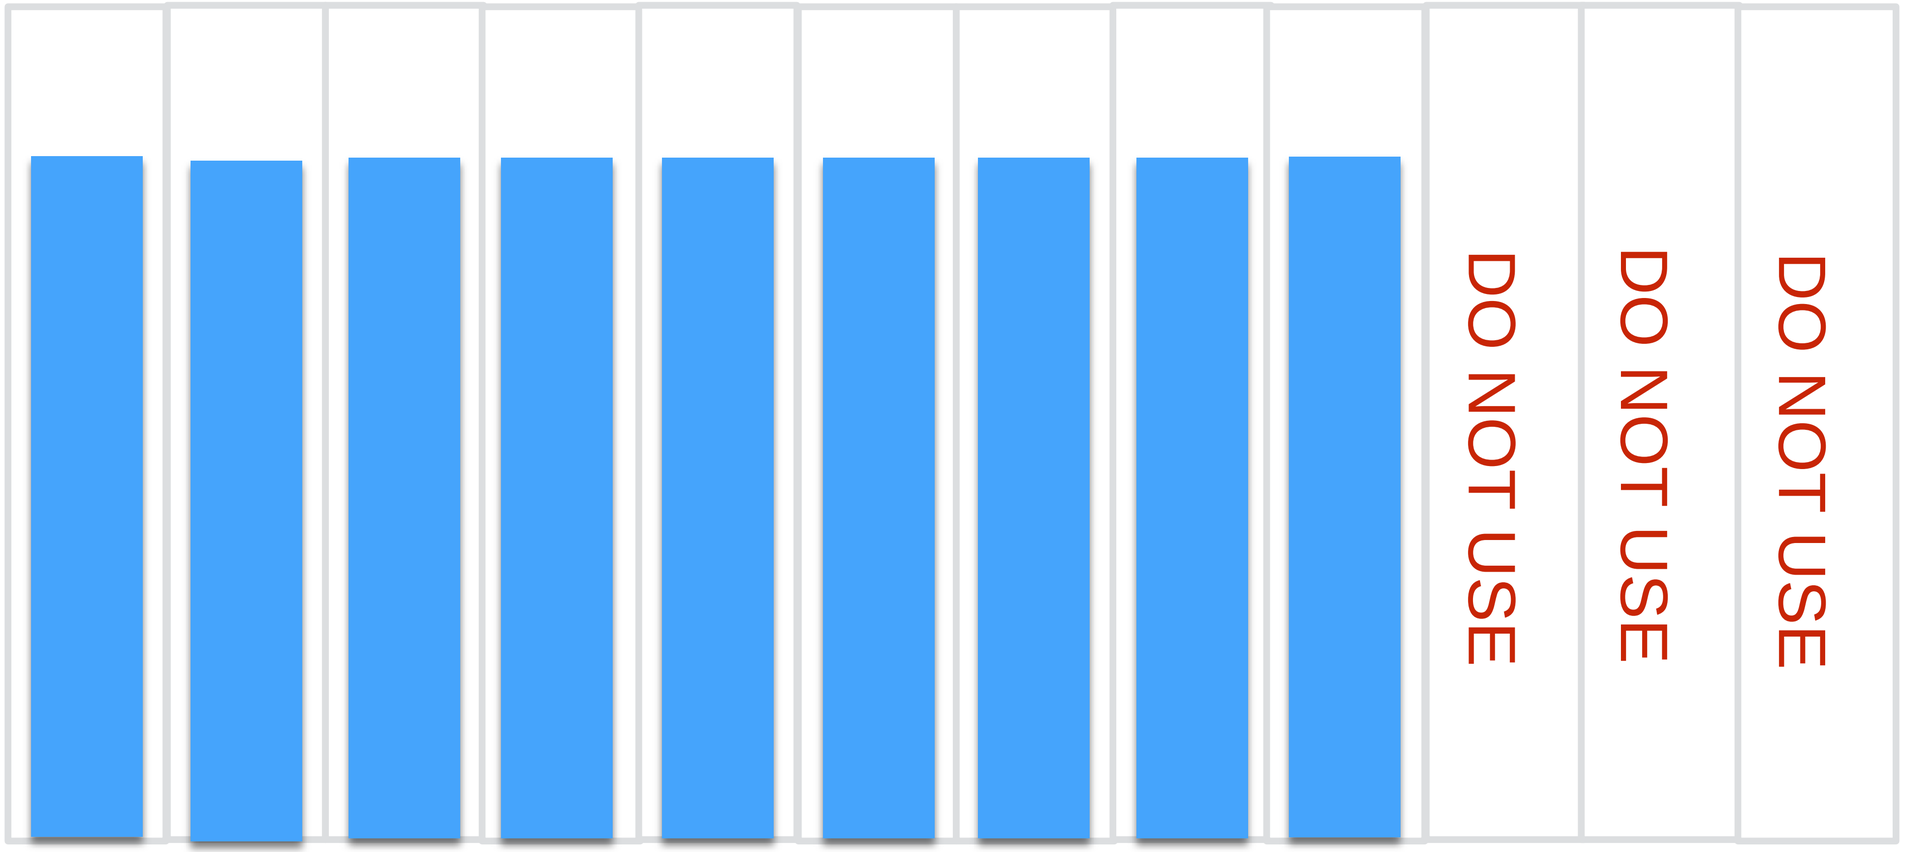

z/TPF LPAR

I-stream cap = 9

1 2 3 4 5 6 7 8 9 10 11 12

In Use I-streams

Fenced I-streams

To-Be Environment – Business Grows Beyond Original Projections so Increase I-stream Cap to Use 10th and 11th Engines

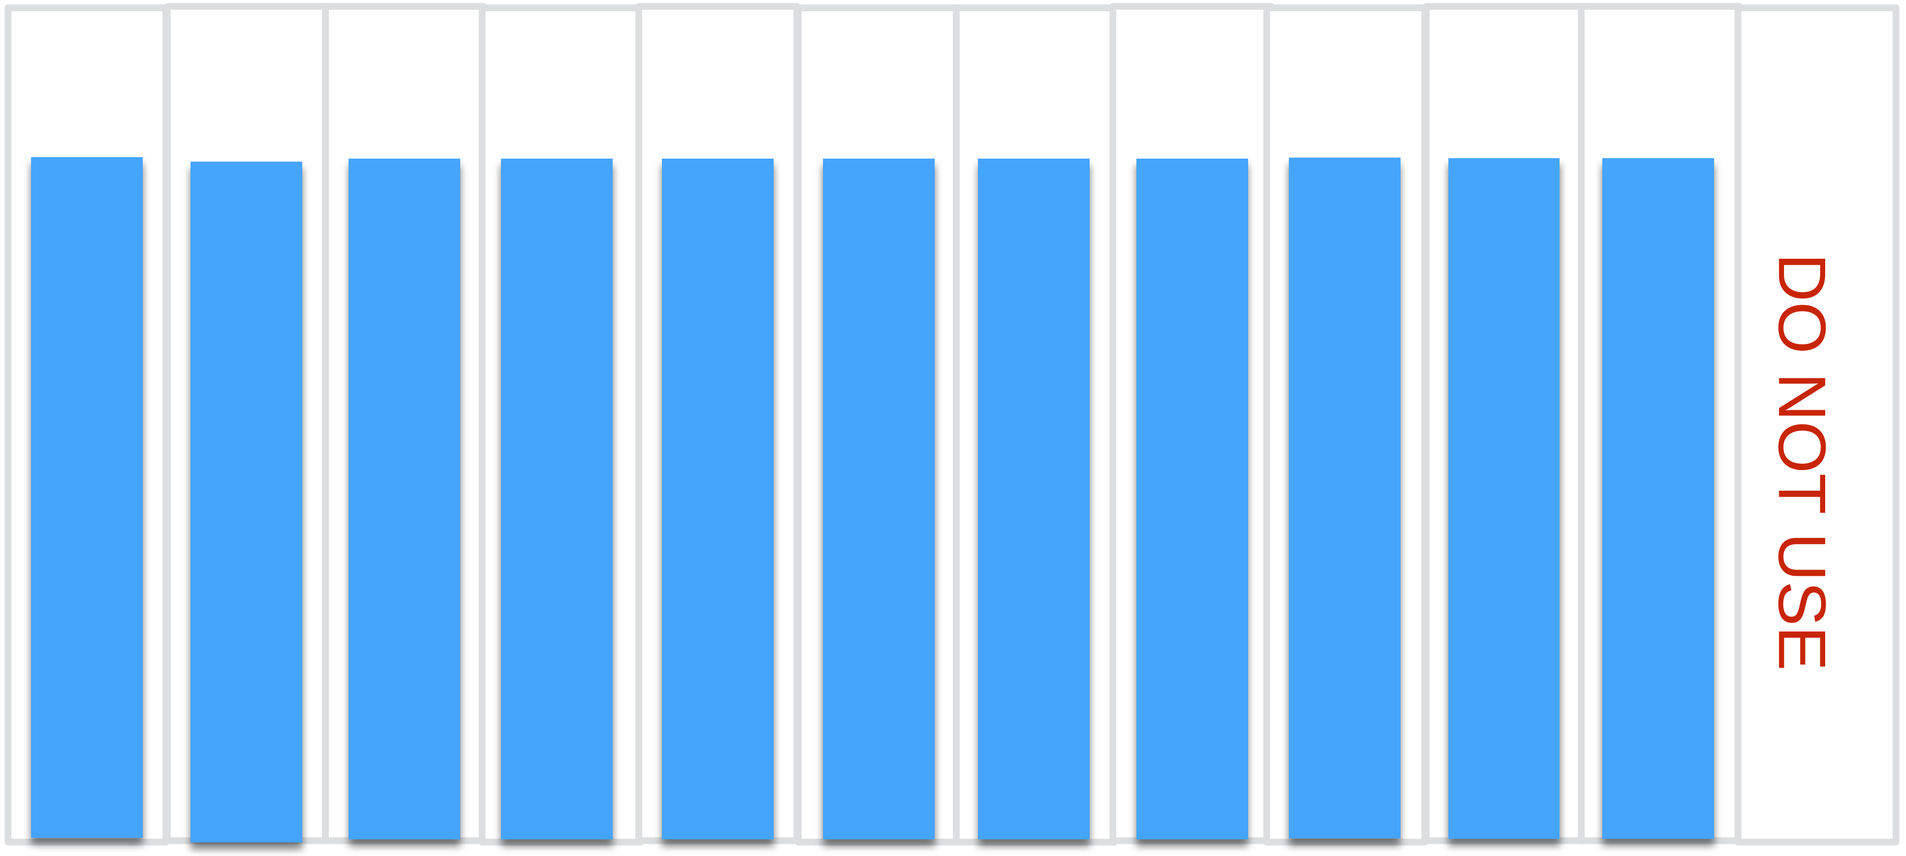

z/TPF LPAR

I-stream cap = 11

1 2 3 4 5 6 7 8 9 10 11 12

In Use I-streams

Fenced I-stream

To-Be Scenario

Timmy’s boss is thrilled because he was able to meet the current and unknown future growing business needs of his company without requiring outages to add the necessary capacity to the z/TPF server.

Dynamic CPU Capacity Hill Statement

Dynamic CPU Capacity

The service provider VP of operations can run more workload on his existing CPU hardware to lower costs and immediately increase CPU capacity to achieve SLAs.

Sub Hill One

The service provider can handle a sustained increase in workload without needing to take an outage.

Sub Hill Two

The service provider can maximize CPU resources to lower its hardware costs.

Sub Hill Three

The service provider can selectively run utilities even during peak volumes without impacting real-time transactions.

Sample z Systems Server with 2 Shared LPARs

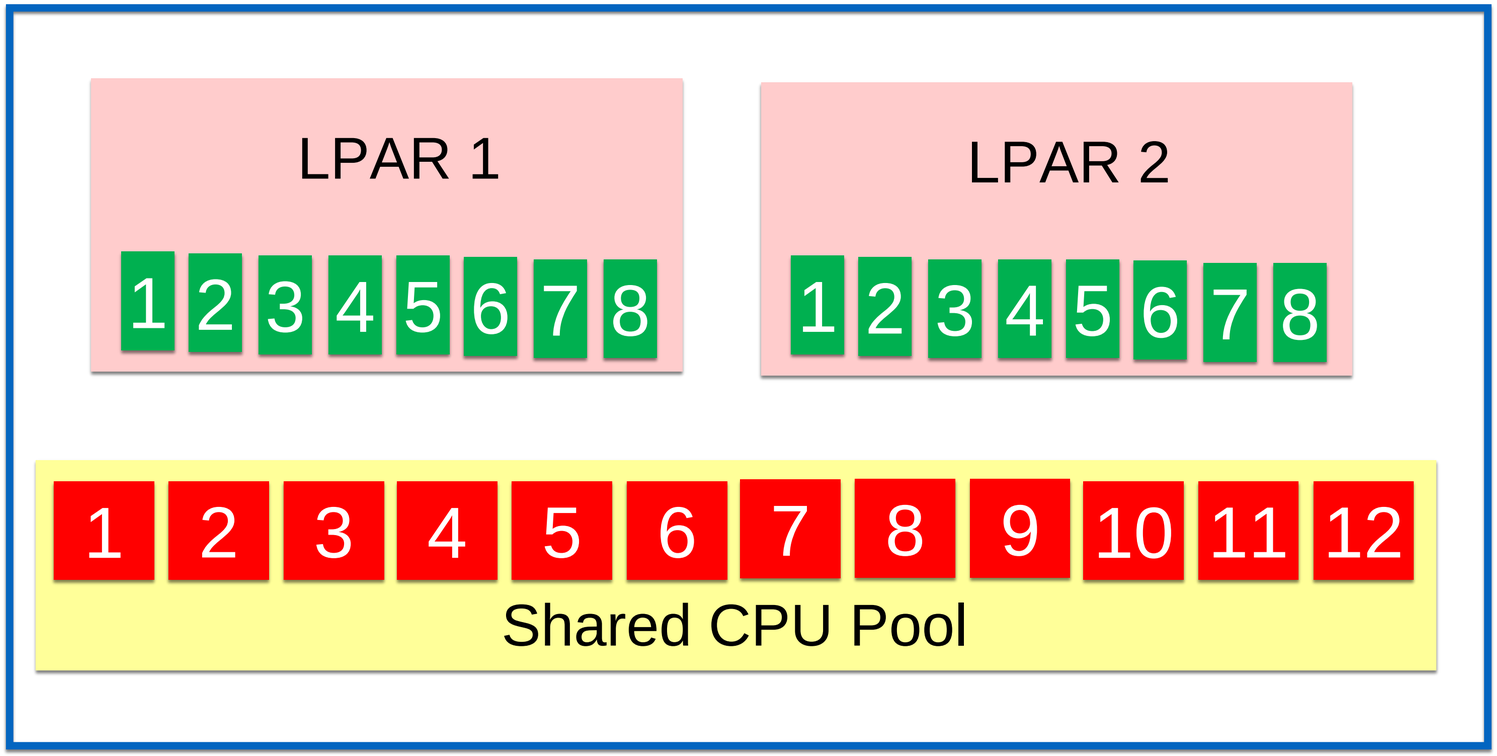

= logical I-stream (CPU)

= physical I-stream (CPU)

- z Systems Server
  - Model 712
  - 12 physical engines (CPUs)
- LPAR 1
  - z/TPF
  - 8 logical engines (I-streams)
  - Shared PR/SM weight = 50
- LPAR 2
  - z/TPF
  - 8 logical engines (I-streams)
  - Shared PR/SM weight = 50

Let's Look Inside the Box
Warning - It's Going To Get Technical Now!

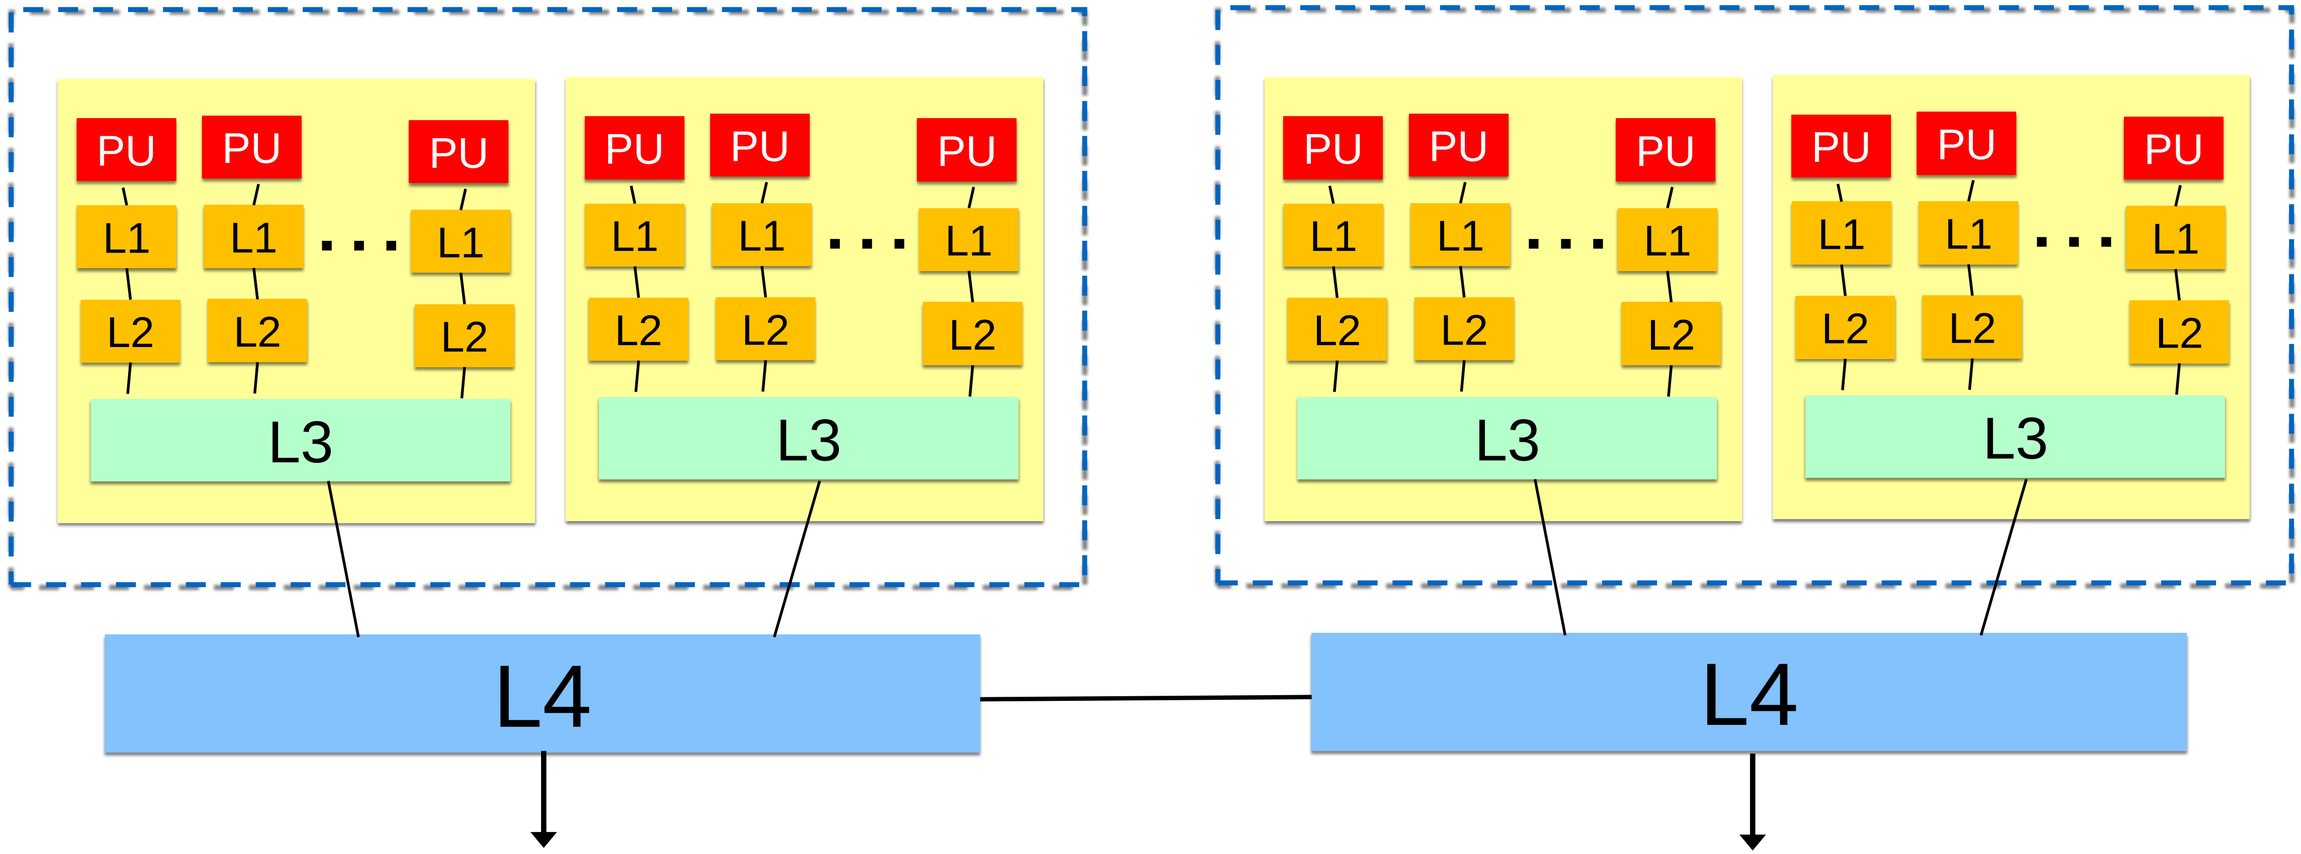

ECB 1 on Logical Engine 1 (I-stream 1) in LPAR 1 is Executing on Physical Engine 1

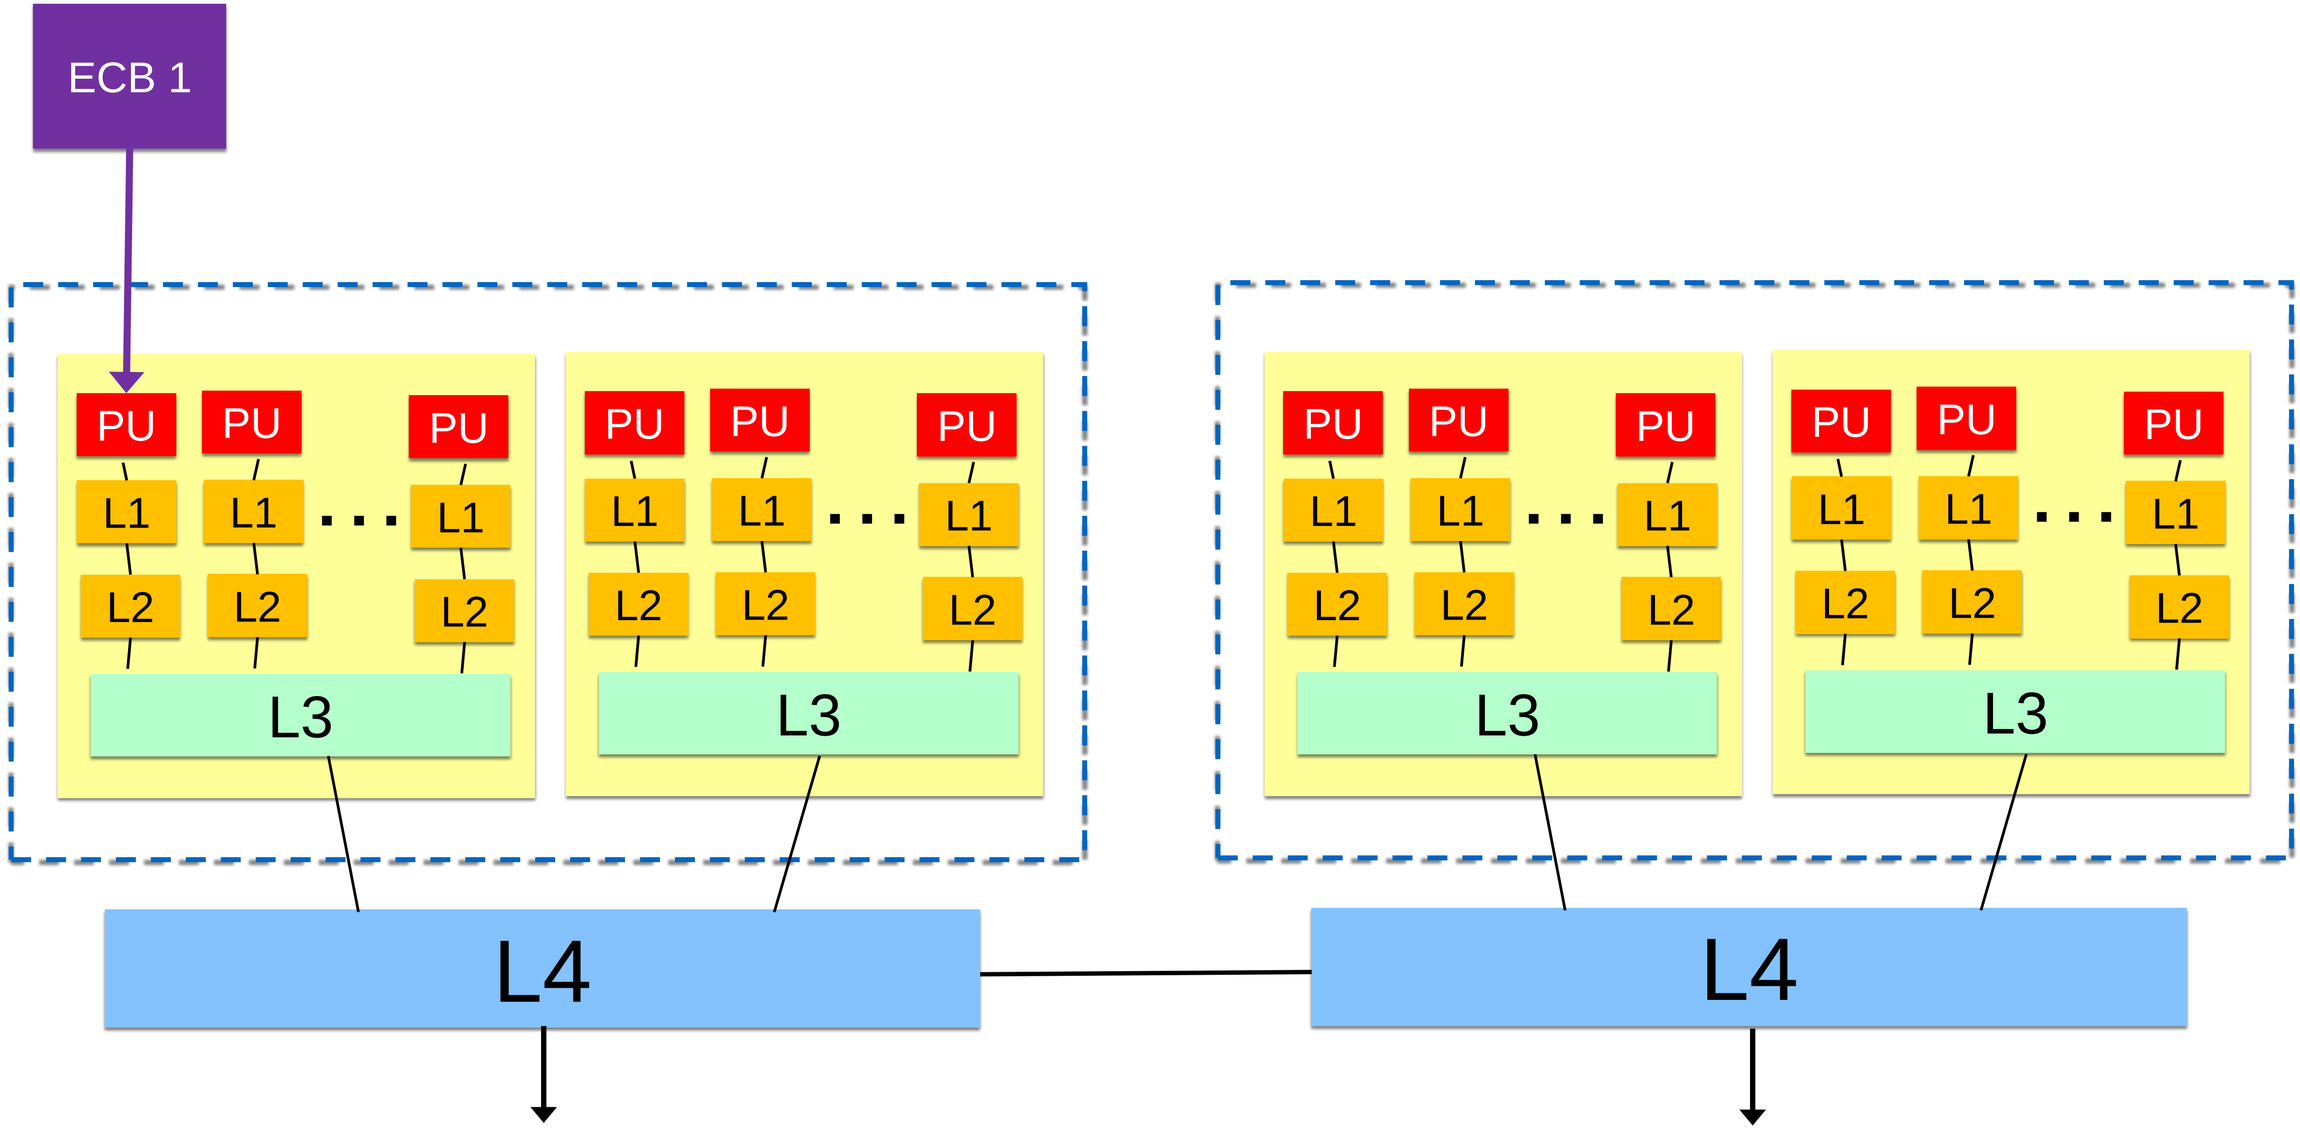

ECB 1 Does I/O and Gives Up Control – When the I/O Completes, it will Continue on Logical Engine 1 in LPAR 1, but on what Physical Engine?

ECB 1

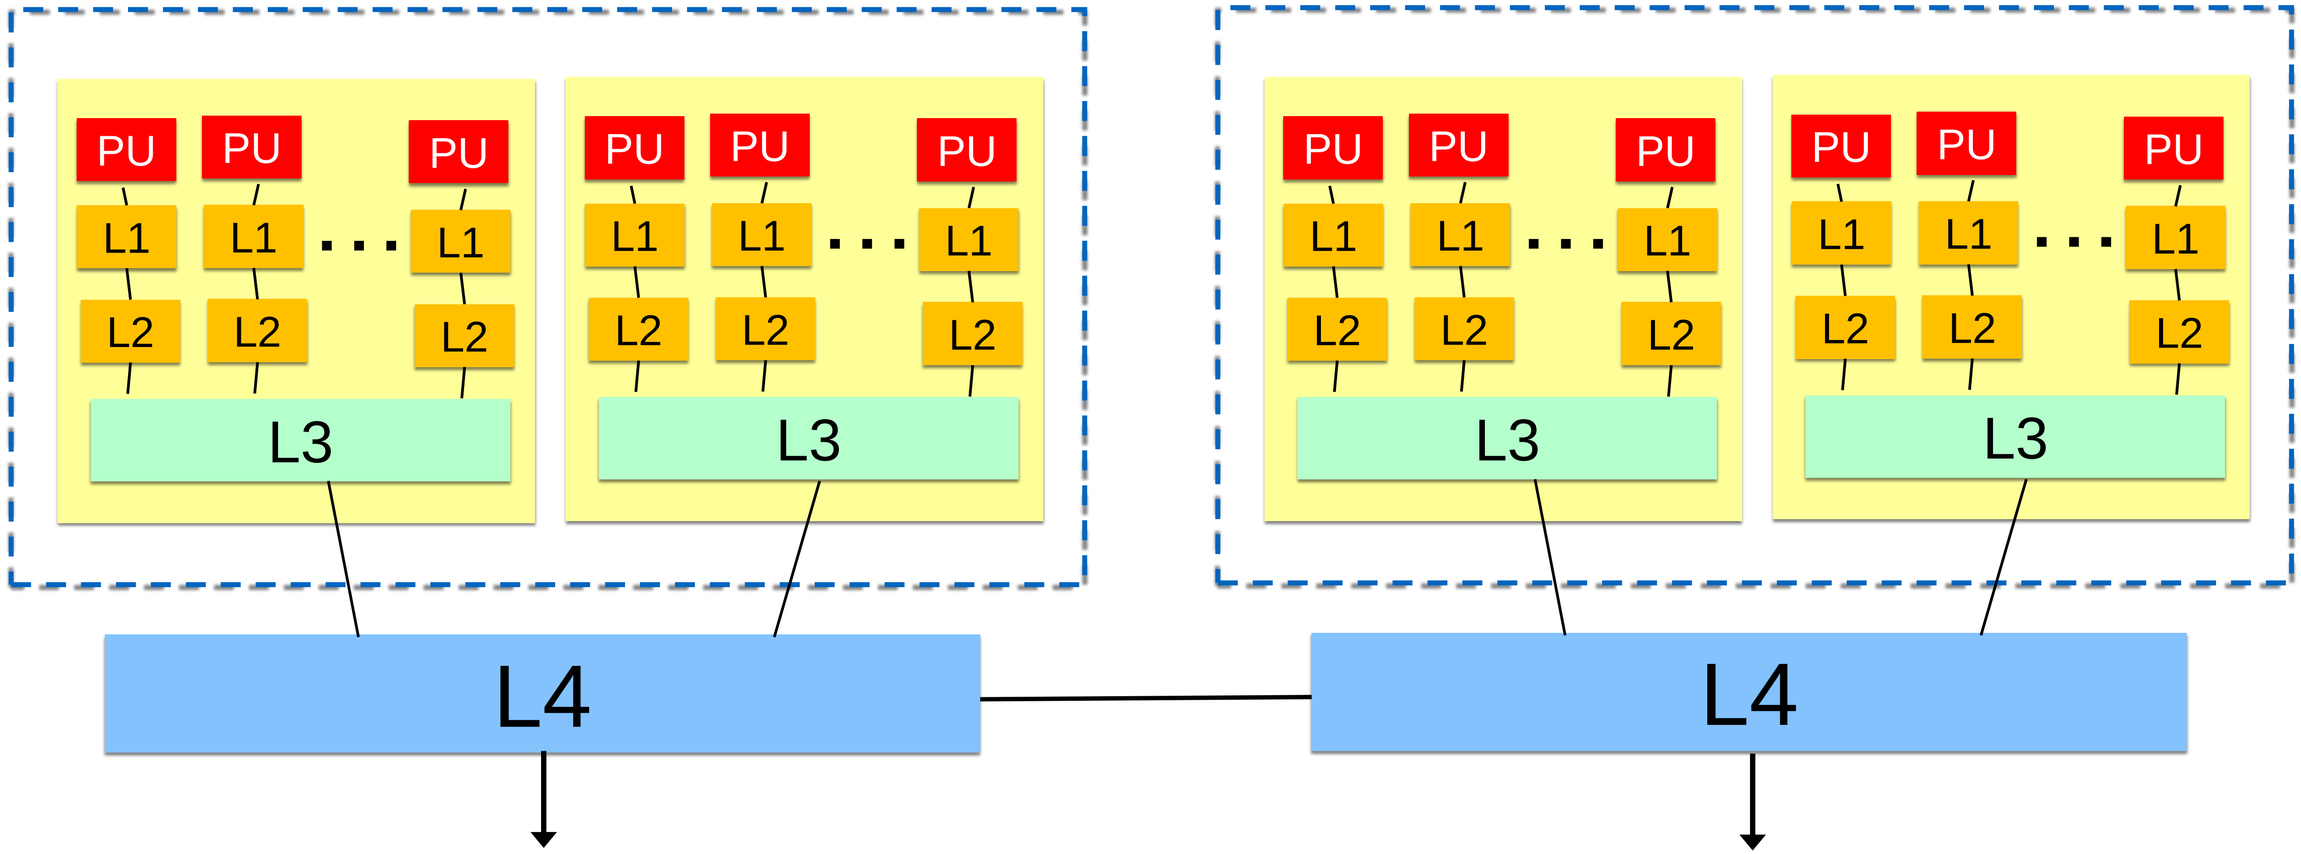

Best Choice – On the Same Physical Engine as Before

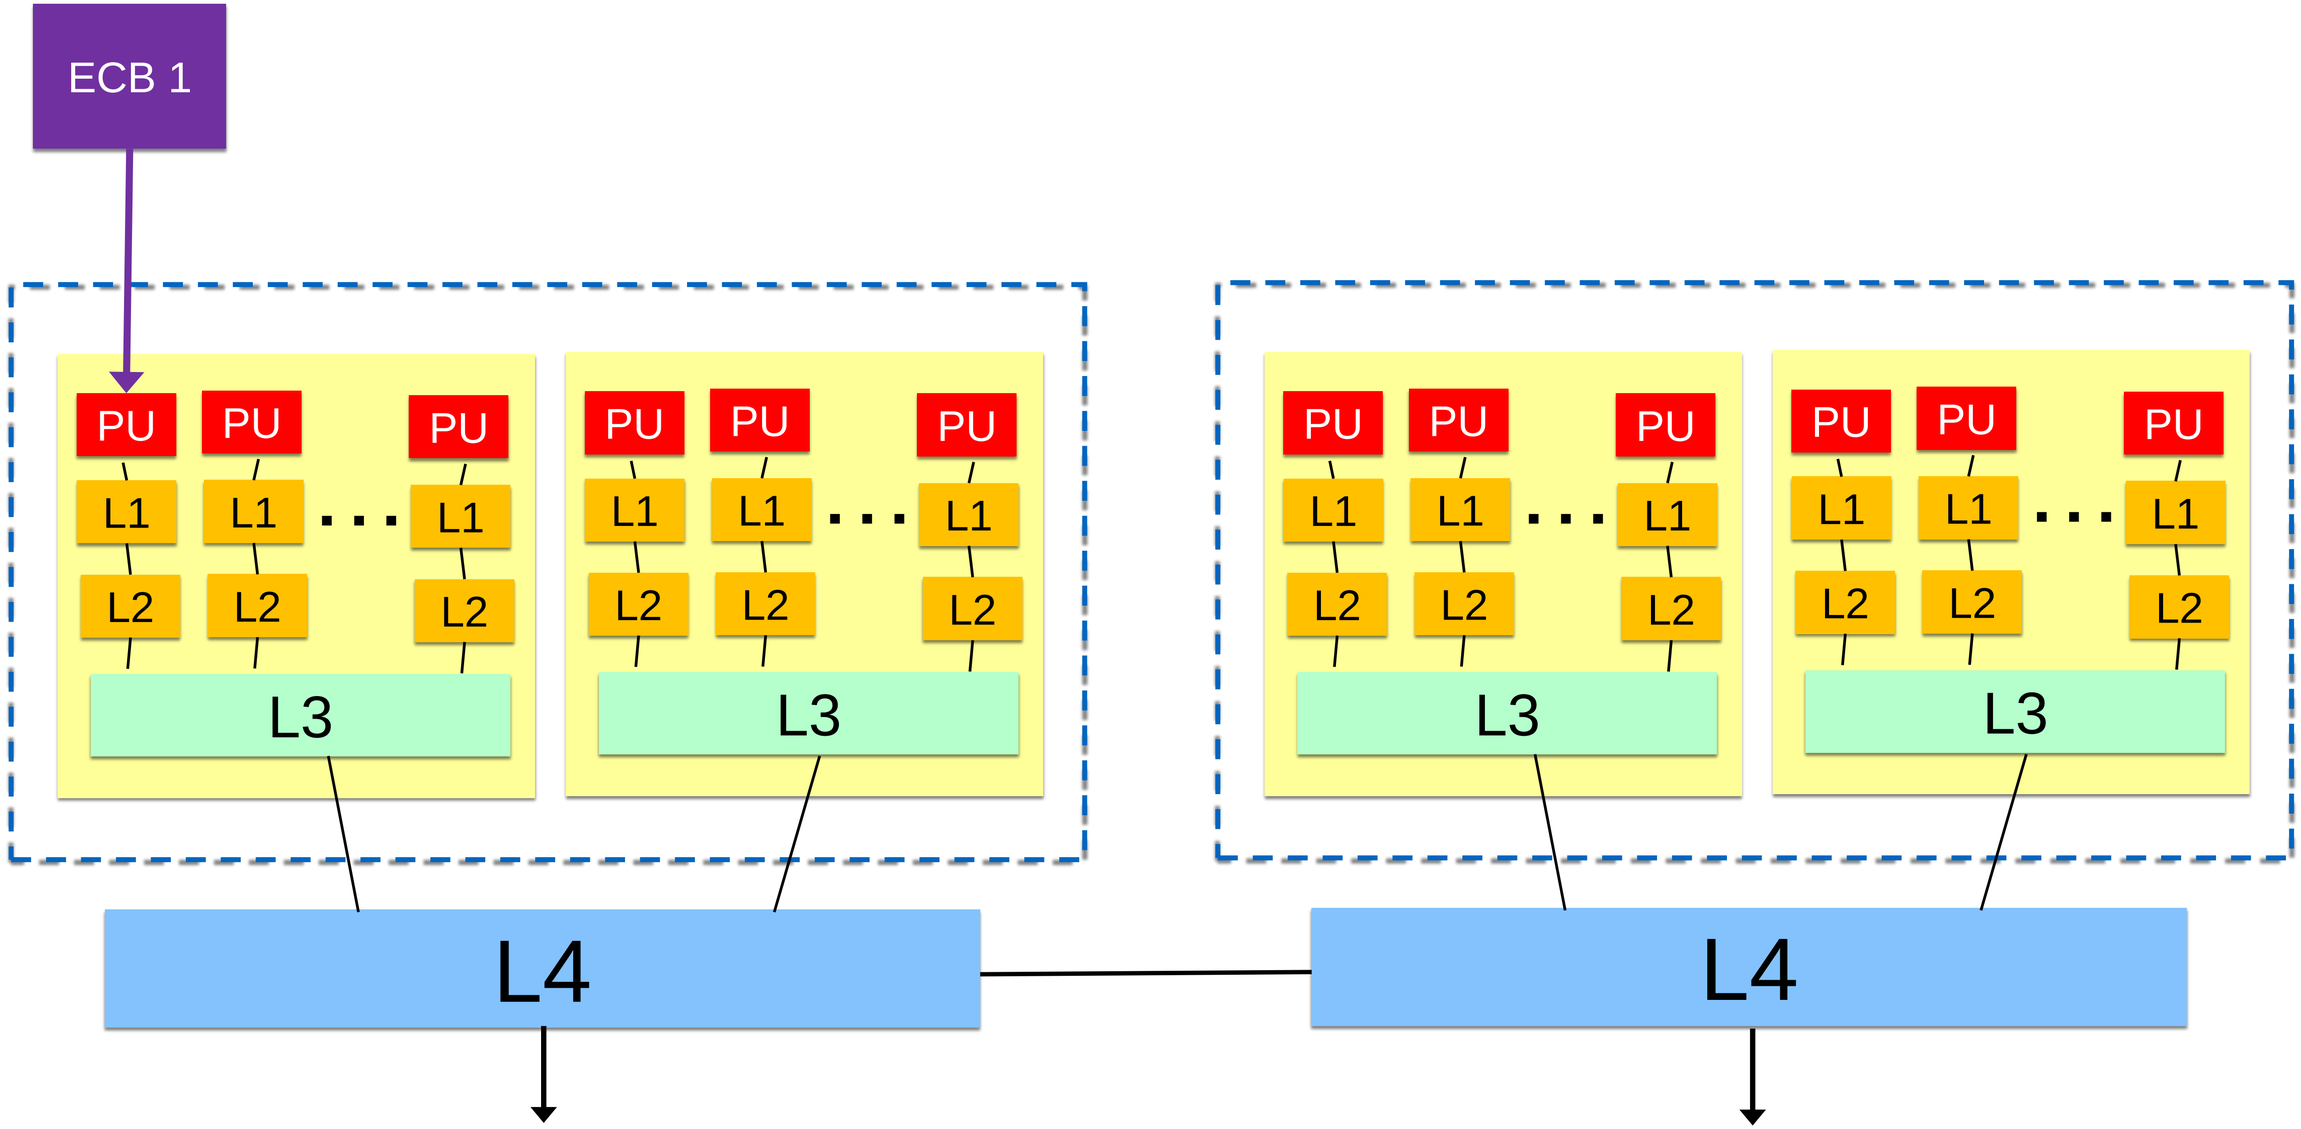

Second Best Choice – On a Physical Engine on the Same Chip as Before

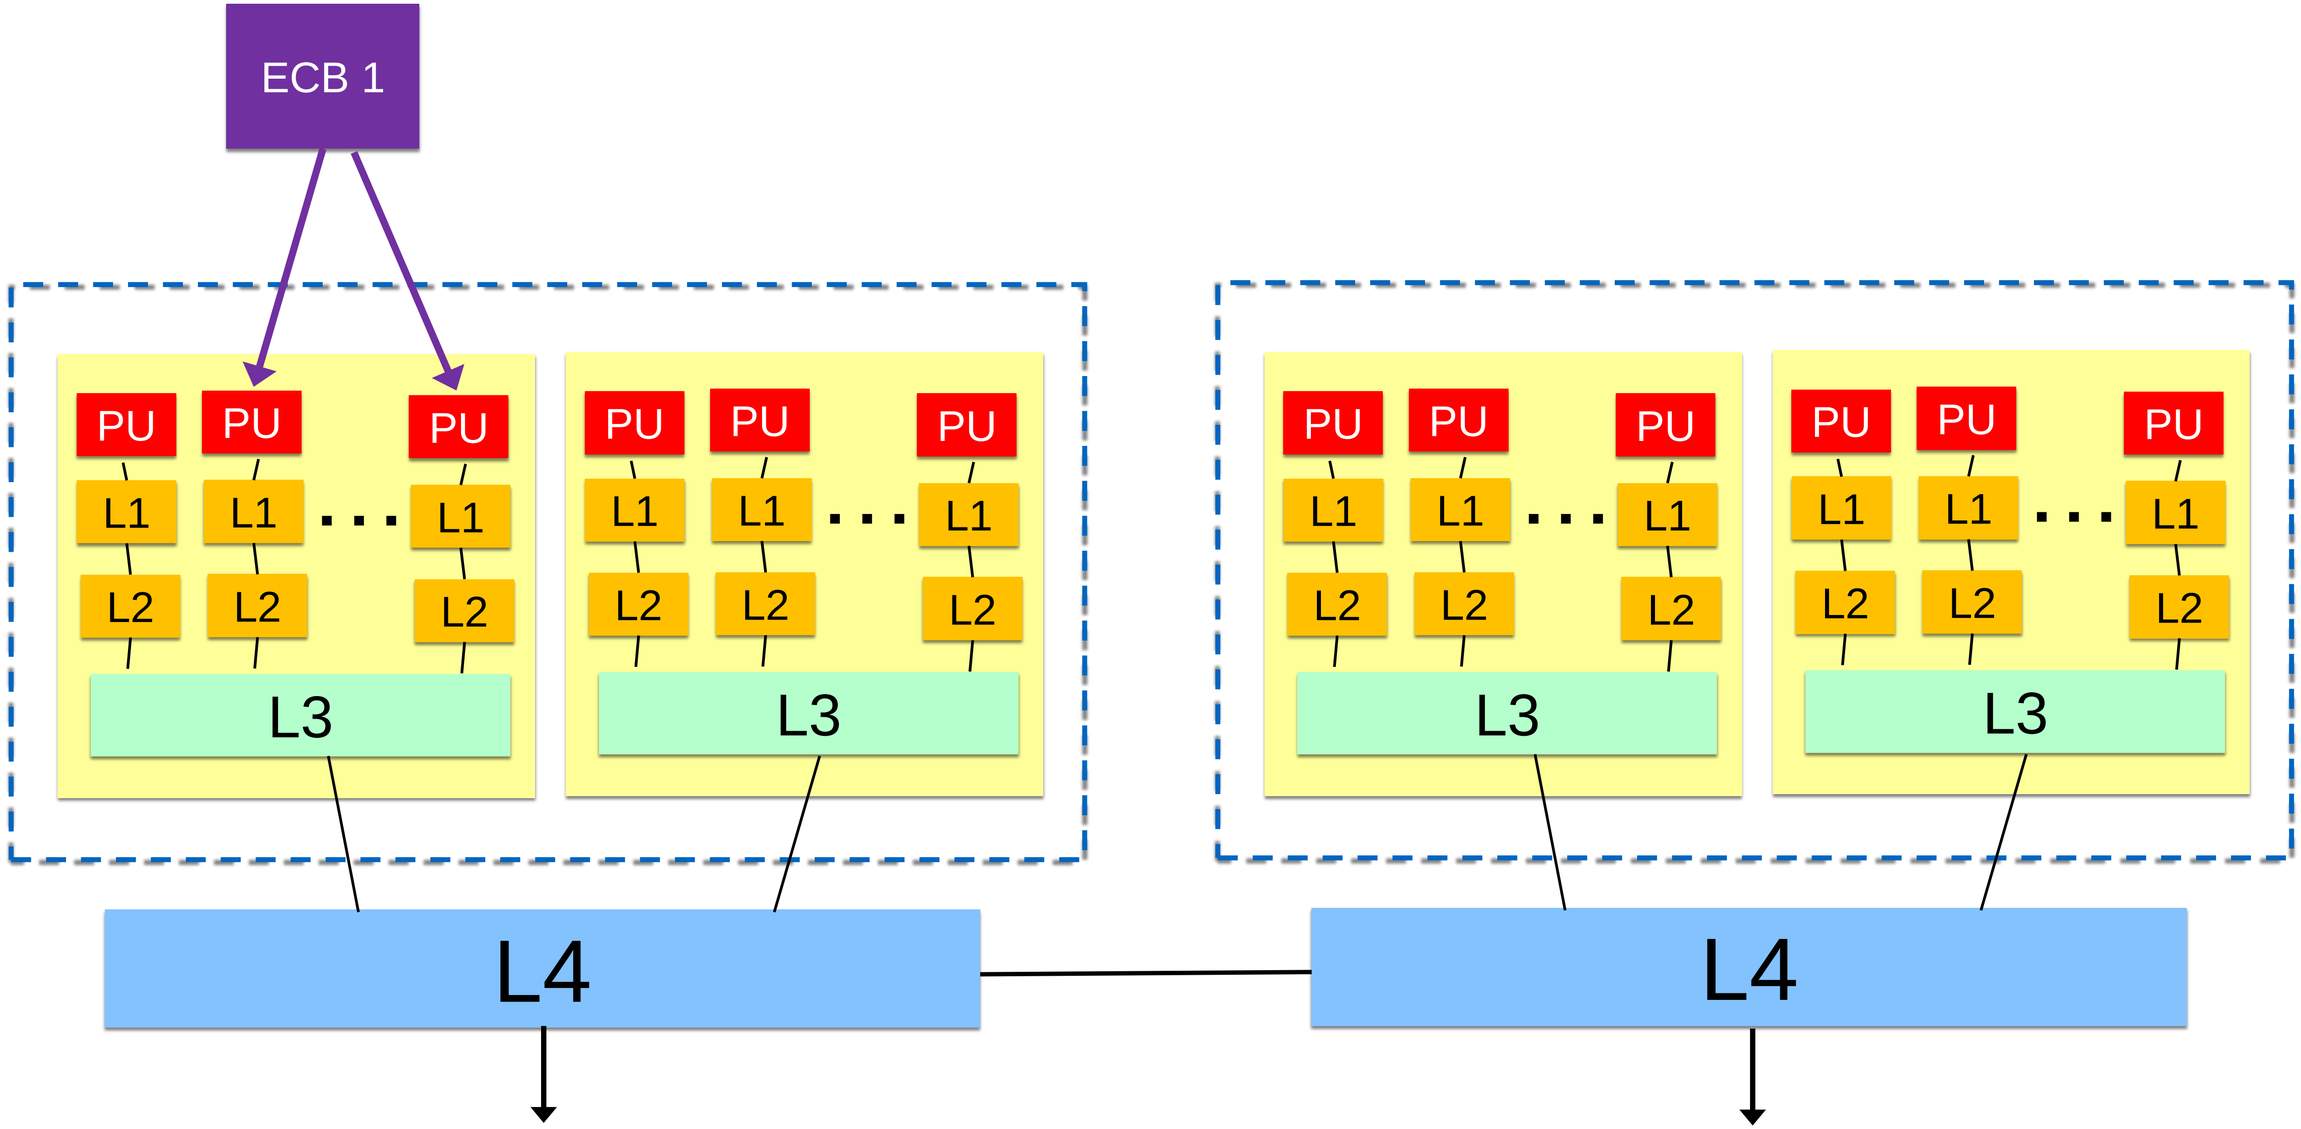

Third Choice – On a Physical Engine on the Same Node as Before

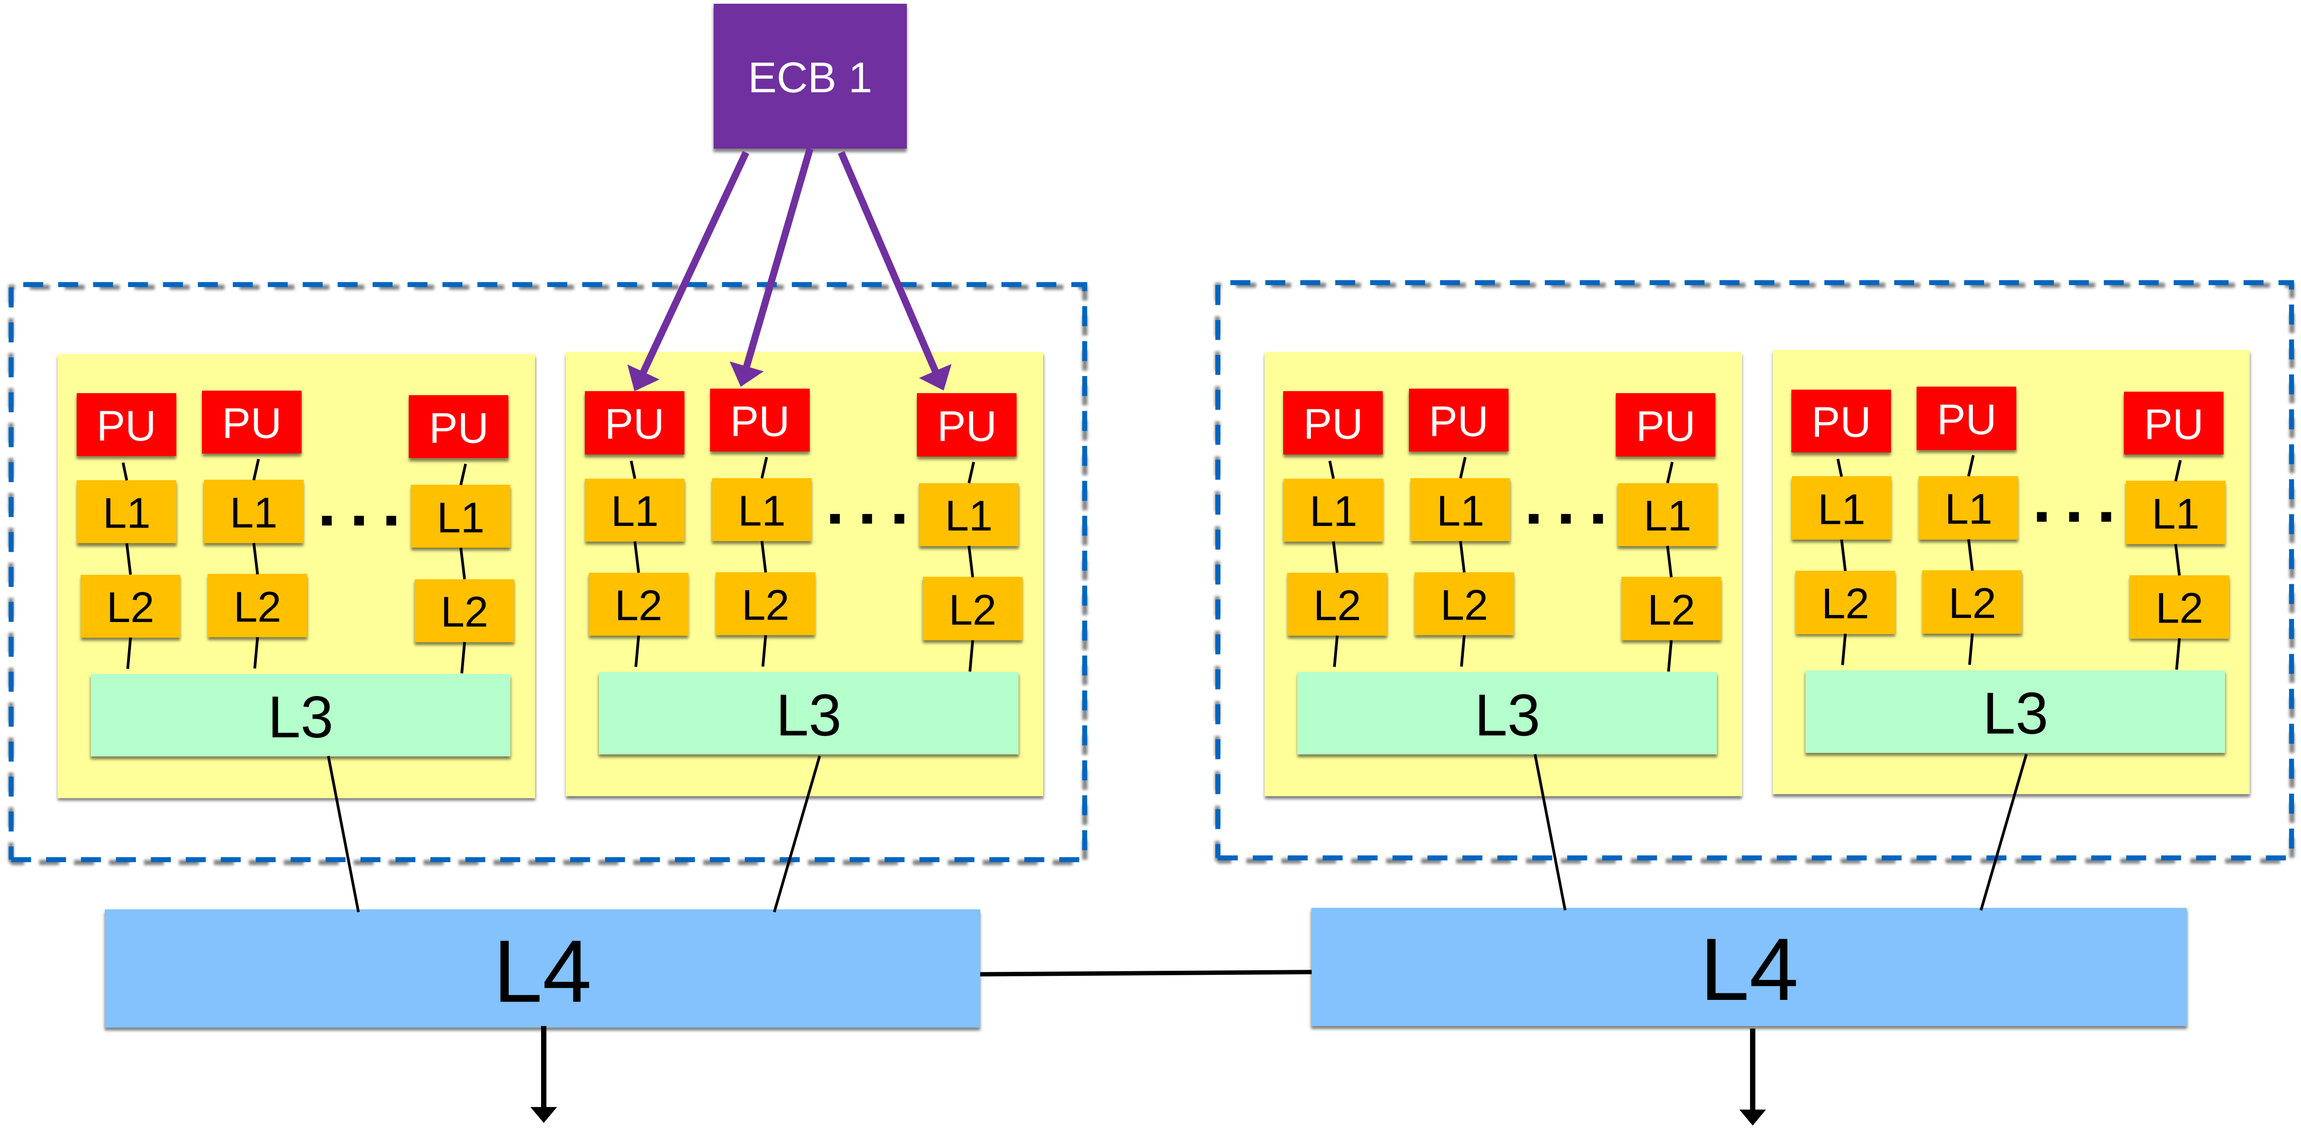

Last Choice – On a Physical Engine a Different Node than Before

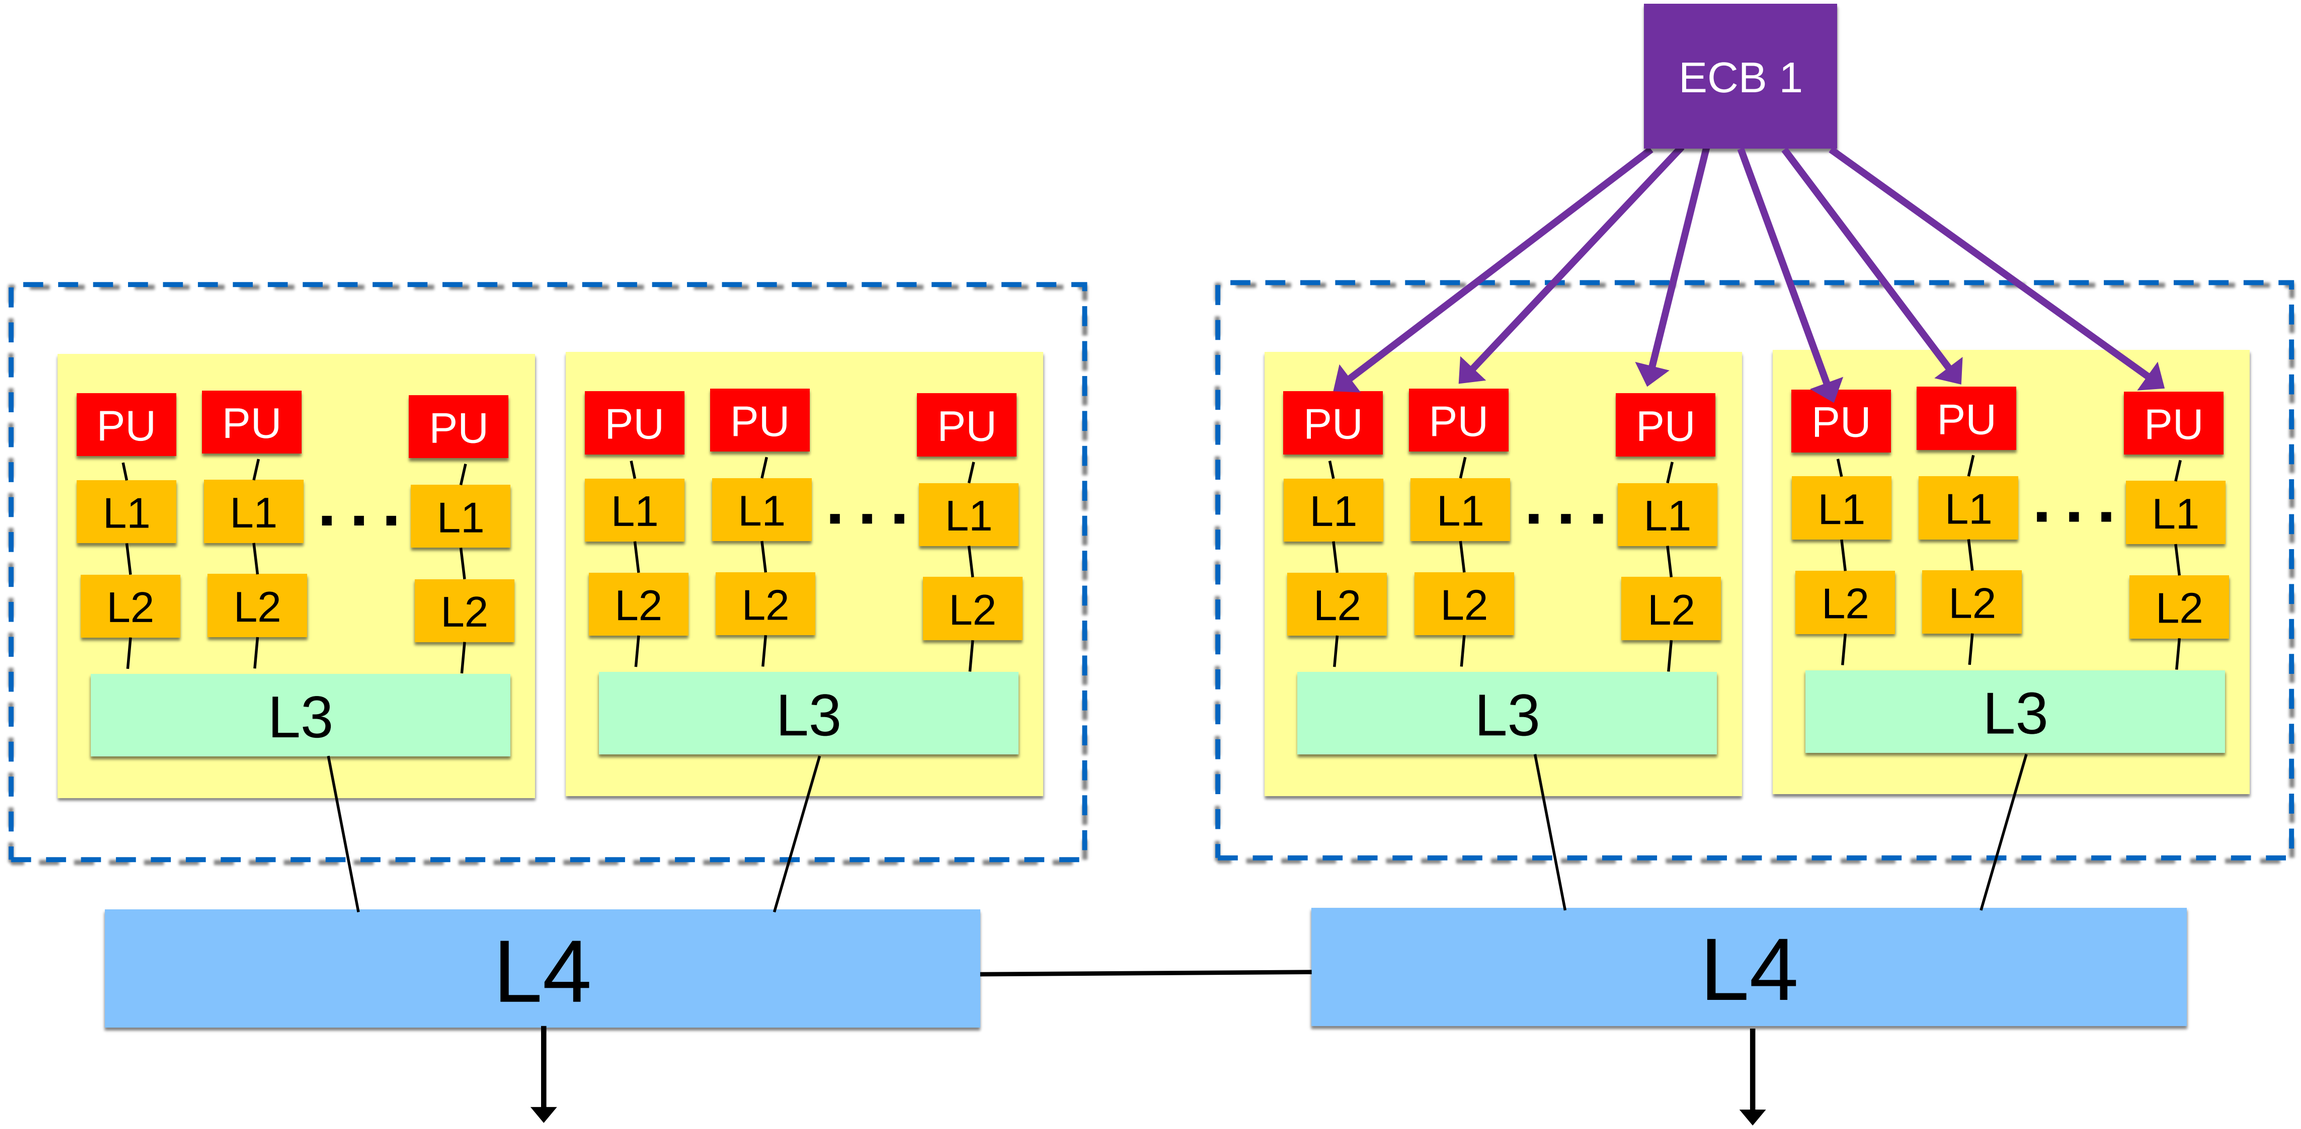

As-Is Environment – Low/Medium Traffic

z/TPF LPAR1
60% CPU busy across 8 engines

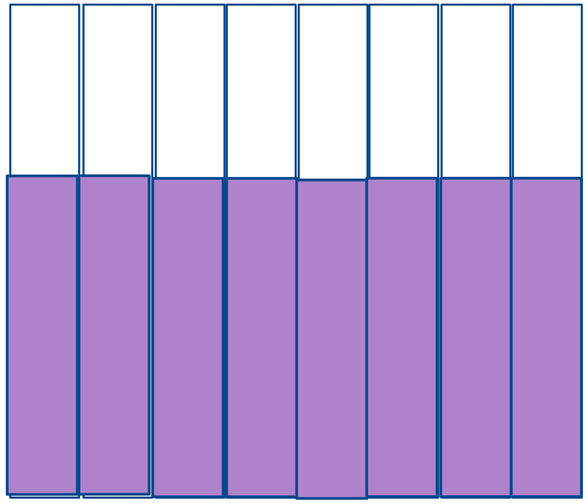

1 2 3 4 5 6 7 8

In-use I-streams

z/TPF LPAR2
20% CPU busy across 8 engines

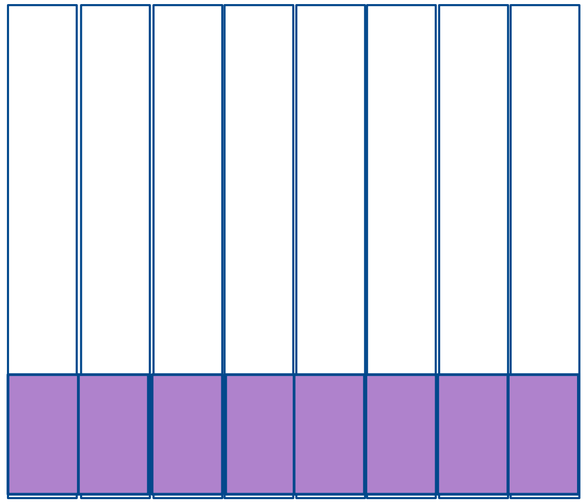

1 2 3 4 5 6 7 8

In-use I-streams

16 Logical In-use

1 2 3 4 5 6 7 8 9 10 11 12

12 Physical

Physical Engines

To-Be Environment – Same Low/Medium Traffic Workload with HiperDispatch Enabled

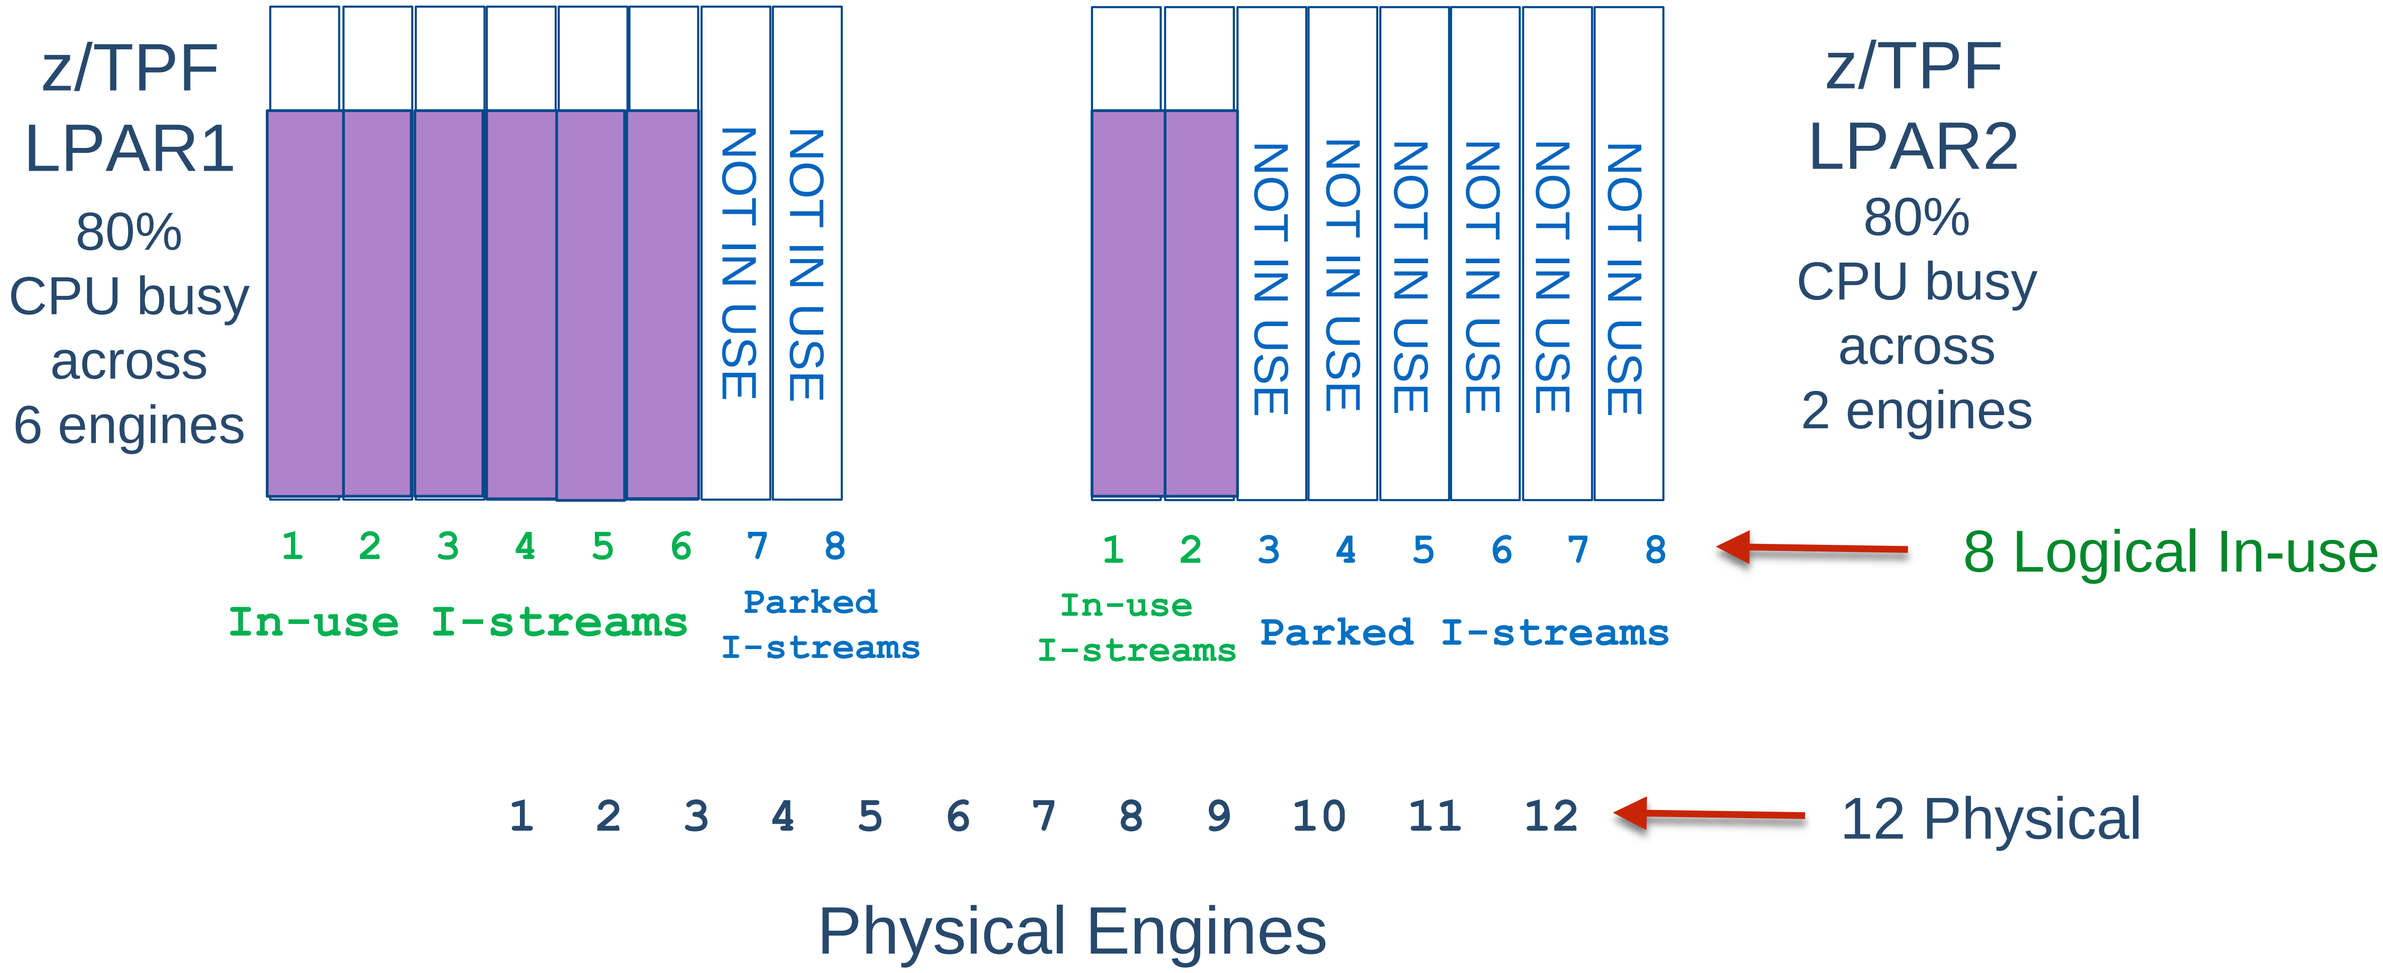

As-Is Environment – High Traffic on Both LPARs

z/TPF LPAR1
72.5% CPU busy across 8 engines

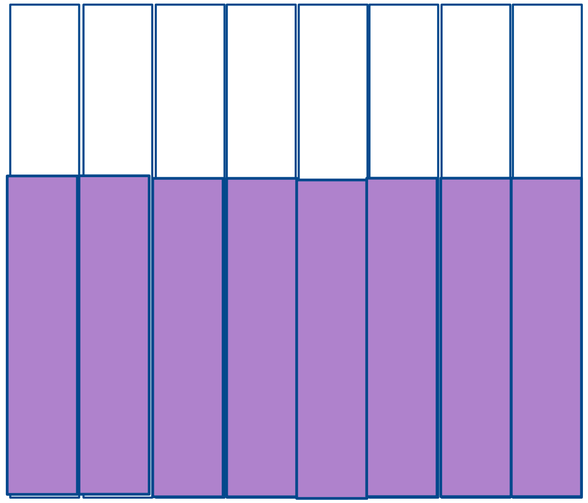

1 2 3 4 5 6 7 8

In-use I-streams

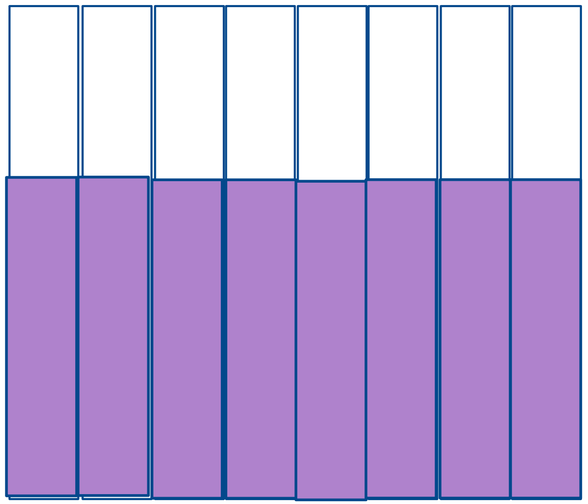

1 2 3 4 5 6 7 8

In-use I-streams

z/TPF LPAR2
72.5% CPU busy across 8 engines

16 Logical In-use

1 2 3 4 5 6 7 8 9 10 11 12

12 Physical

Physical Engines

To-Be Environment – High Traffic on Both LPARs with HiperDispatch Enabled

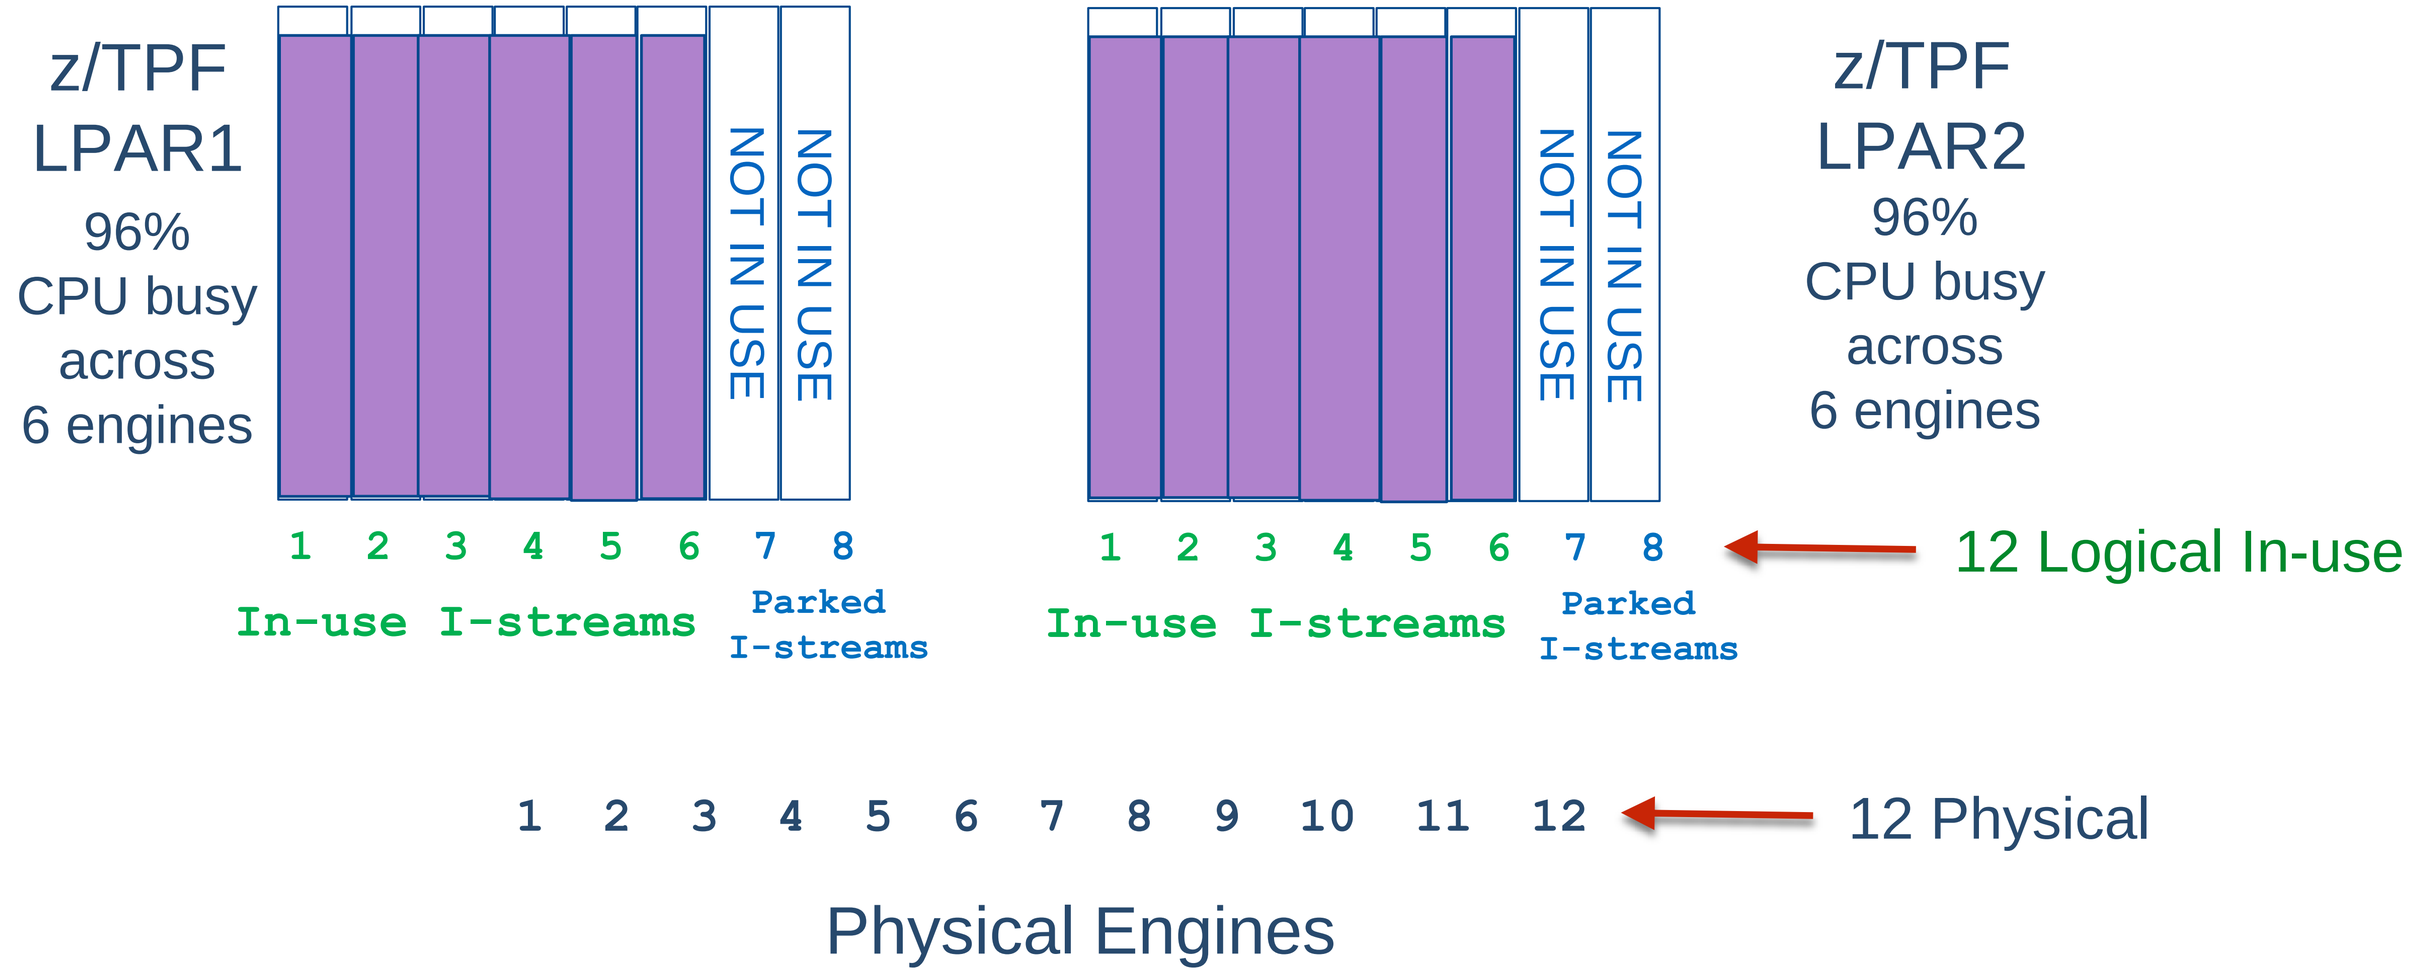

Dynamic CPU Capacity Hill Statement

Dynamic CPU Capacity

The service provider VP of operations can run more workload on his existing CPU hardware to lower costs and immediately increase CPU capacity to achieve SLAs.

Sub Hill One

The service provider can handle a sustained increase in workload without needing to take an outage.

Sub Hill Two

The service provider can maximize CPU resources to lower its hardware costs.

Sub Hill Three

The service provider can selectively run utilities even during peak volumes without impacting real-time transactions.

Let's talk about…

Dave is the VP of operations for service provider HiJynxs Associates Ltd. (HAL). Ever since the original contract in 2001 with customer JupiterBound, HAL was able to run z/TPF utilities at night when traffic volume was low.

As-Is Scenario

JupiterBound expands its business across multiple continents so there are no multi-hour “low traffic periods” anymore; therefore, HAL is now forced to run some utilities during high traffic periods.

To prevent those utilities from impacting transactional workloads, Dave must install more CPU capacity, and those additional costs either reduce HAL’s profits or increases JupiterBound’s bill reducing customer satisfaction… a lose-lose situation.

To-Be Scenario

JupiterBound expands its business across multiple continents so there are no multi-hour “low traffic periods” anymore; therefore, HAL is now forced to run some utilities during high traffic periods.

HAL can now mark certain utilities as “low priority”, meaning if CPU resources become constrained, these utilities will run at a slower rate as to not impact transactional workloads.

HAL is now able to run utilities at any time of day or night to meet SLA’s and do so on existing HW so both HAL and JupiterBound are happy… a win-win situation.

THANK YOU

Questions or comments?

Mark Gambino
z/TPF Chief Architect

IBM **z/TPF**
April 3rd, 2017

trademarks

IBM, the IBM logo, ibm.com and Rational are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at “Copyright and trademark information” at www.ibm.com/legal/copytrade.shtml.

Notes

Performance is in Internal Throughput Rate (ITR) ratio based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput improvements equivalent to the performance ratios stated here.

All customer examples cited or described in this presentation are presented as illustrations of the manner in which some customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics will vary depending on individual customer configurations and conditions.

This publication was produced in the United States. IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the product or services available in your area.

All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

Information about non-IBM products is obtained from the manufacturers of those products or their published announcements. IBM has not tested those products and cannot confirm the performance, compatibility, or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Prices subject to change without notice. Contact your IBM representative or Business Partner for the most current pricing in your geography.

This presentation and the claims outlined in it were reviewed for compliance with US law. Adaptations of these claims for use in other geographies must be reviewed by the local country counsel for compliance with local laws.