IBM Z

IBM Server Technology Review

IBM Z and LinuxONE III

Monte Bauman

IBM Z and LinuxONE Technical Support Specialist

IBM Columbus

mbauman@us.ibm.com

June 2020

IBM

Let's Get Started…

INTRODUCTION

Topics

Topics

Topics

A little history

THE Z CORE STORY

IBM Z – Processor Roadmap

1.4B
45 nm
z196
9/2010

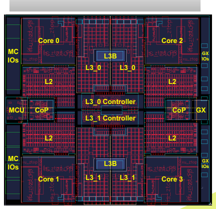

Top Tier Single Thread Performance,System Capacity

Accelerator Integration

Out of Order Execution

Water Cooling

PCIe I/O Fabric

RAIM

Enhanced Energy Management

2.75B
32 nm
zEC12
8/2012

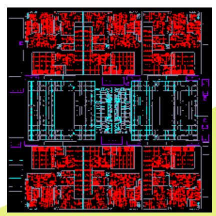

Leadership Single Thread, Enhanced Throughput

Improved out-of-order

Transactional Memory

Dynamic Optimization

2 GB page support

Step Function in System Capacity

4.0B
22 nm
z13
1/2015

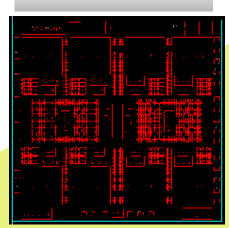

Leadership System Capacity and Performance

Modularity & Scalability

Dynamic **SMT**

Supports two instruction threads

SIMD

PCIe attached accelerators

Business Analytics Optimized

6.1B
14 nm
z14
7/2017

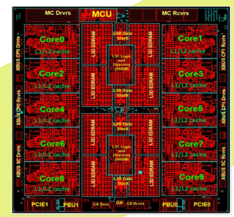

Pervasive encryption

Low latency I/O for acceleration of transaction processing for DB2 on z/OS

Pause-less garbage collection for enterprise scale JAVA applications

New SIMD instructions

Optimized pipeline and enhanced SMT

Virtual Flash Memory

9.1B
14 nm
z15
9/2019

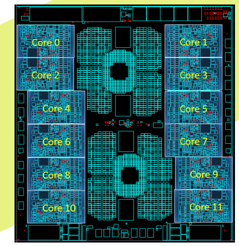

Focus on power efficiency and new on-chip architectures

Improved and enlarged caches

Optimized Out-of-Order architecture

Binary Floating point enhancements

IBM Integrated Accelerator for zEDC (On-chip compression support

Enhanced Cryptographic Coprocessor (CPACF)

Topics

z15 Continues the CMOS Mainframe Heritage

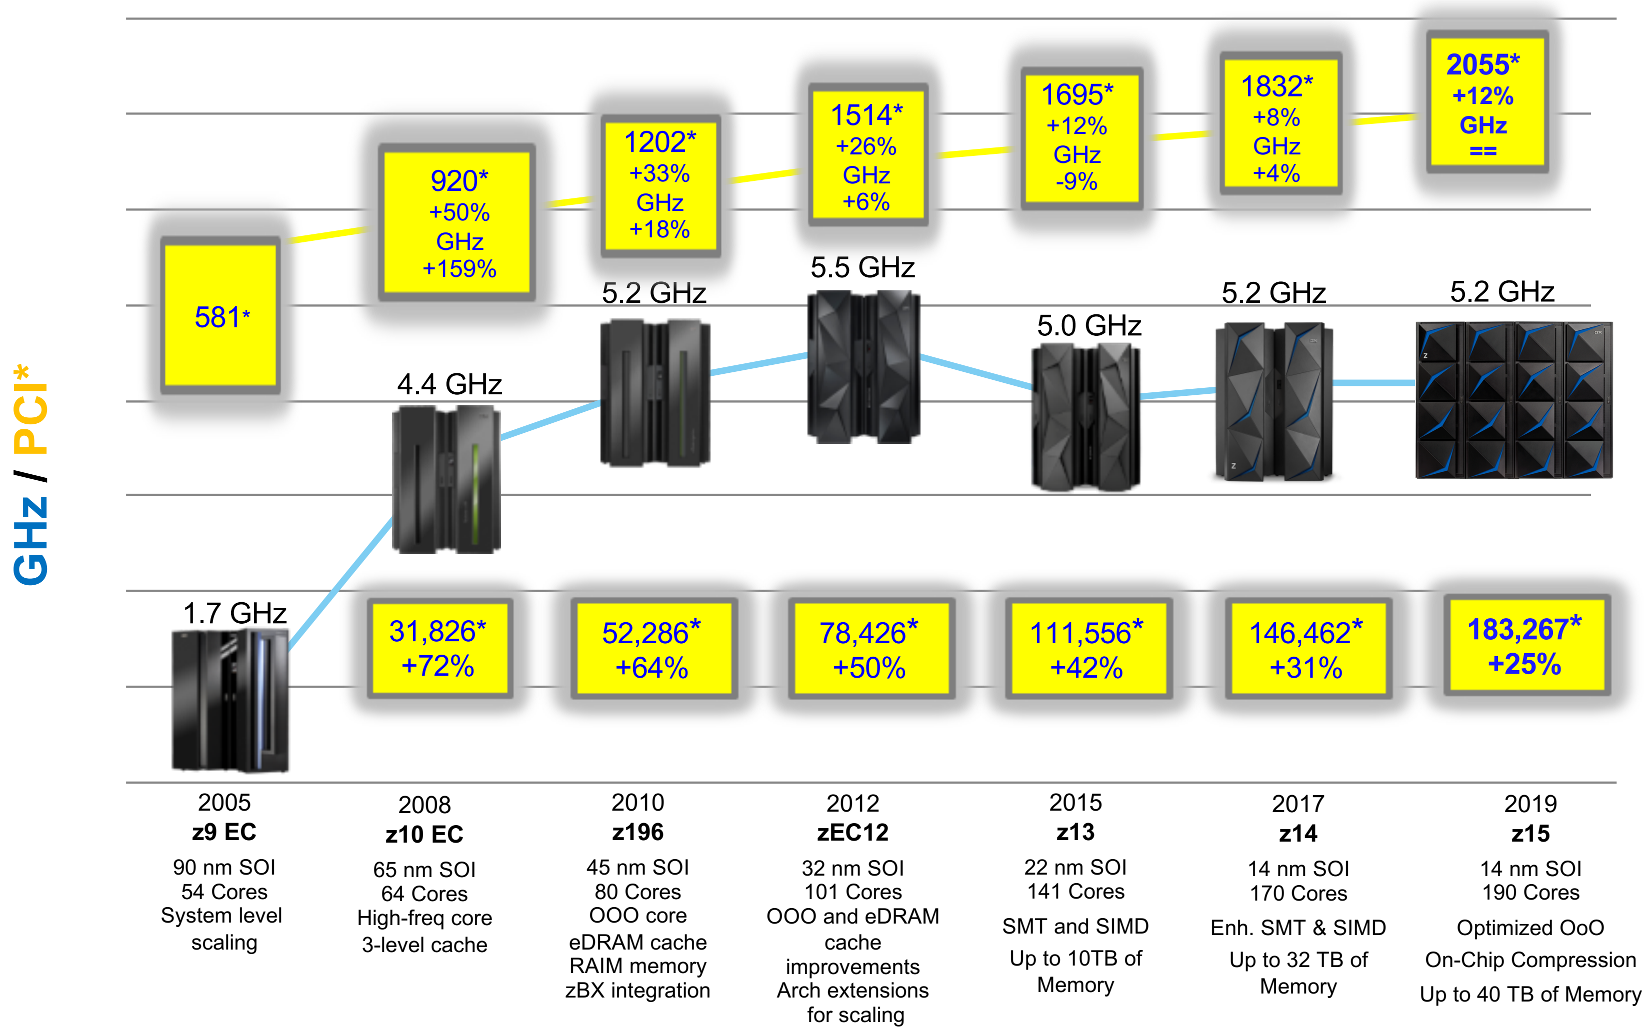

* PCI Tables are NOT adequate for making comparisons of IBM Z processors. Additional capacity planning required

Extending IBM z15 and LinuxONE III

Built on IBM z15 chip technology to address new markets	Cloud native development and deployment	Encryption everywhere to ensure data privacy	Resiliency and availability in simplified package

IBM z15 Model T01

Flexible compute design
- Available in one to four 19” frames based on capacity needs
- Secure Execution for Linux improves tenant density for fortified, virtual machine level granularity
- Two power options – iPDU for electrical efficiency and Bulk Power Assembly for water cooled, internal battery or balance power

Processor Units (PUs)
- Up to 190 client configurable cores
- 14% single-thread performance improvement and 25% maximum system capacity growth over IBM z14™
- Integrated accelerators for crypto and compression
- More investments for Java™ and analytics applications

To the Data
- Encryption of data in flight between z15 and new IBM DS8900F for secure fiber channel endpoint protection
- Significant scalability improvements — up to 20% more I/O channels and 50% more logical Coupling Facility connections
- Faster SSL/TLS handshake performance on IBM z15 with Crypto Express7S compared to z14 with Crypto Express6S
- IBM HyperLink technology for low latency connectivity

IBM z15™
Machine Type: 8561
Model T01

Memory
- Up to 40 TB RAIM Memory design

CPC Drawer	Customer PUs	Max Memory
1	34	8 TB
2	71	16 TB
3	108	24 TB
4	145	32 TB
5	190	40 TB

IBM z15 Model T02

Flexible compute design
- 19” form factor, single or 3 phase power
- Second drawer option for enhanced high availability
- Secure Execution for Linux improves tenant density for fortified, virtual machine level granularity

Processor Units (PUs)
- Feature based sizing 4, 6/13, 6/21, 6/31, 6/65
- 2-8 SAPs and 1-2 spares designated per system
- Up to 40 LPARs
- Integrated accelerators for cryptography and compression

Memory
- RAIM Memory design – Min 64 GB – Max 16 TB
- 160 GB Fixed HSA
- IBM Virtual Flash Memory for improved availability

To the Data
- Significant scaling improvements for coupling – 2x more Coupling Express Long Reach, ICP channel paths and ICF processors and 50% more logical coupling channel paths than z14 ZR1
- IBM HyperLink technology for low latency connectivity

IBM z15™
Machine Type: 8562
Model T02

Feature Based Sizing

CPC Drawer	Customer PUs	Max Memory
Max4	4	2 TB
Max13	13	4 TB
Max21	21	4 TB
Max31	31	8 TB
Max65	65	16 TB

Topics

Let’s Get to the Core of this …

THE Z15 CORE

Topics

z15 Core

Topics

5.2GHz

Pipeline

z15
Core

Topics

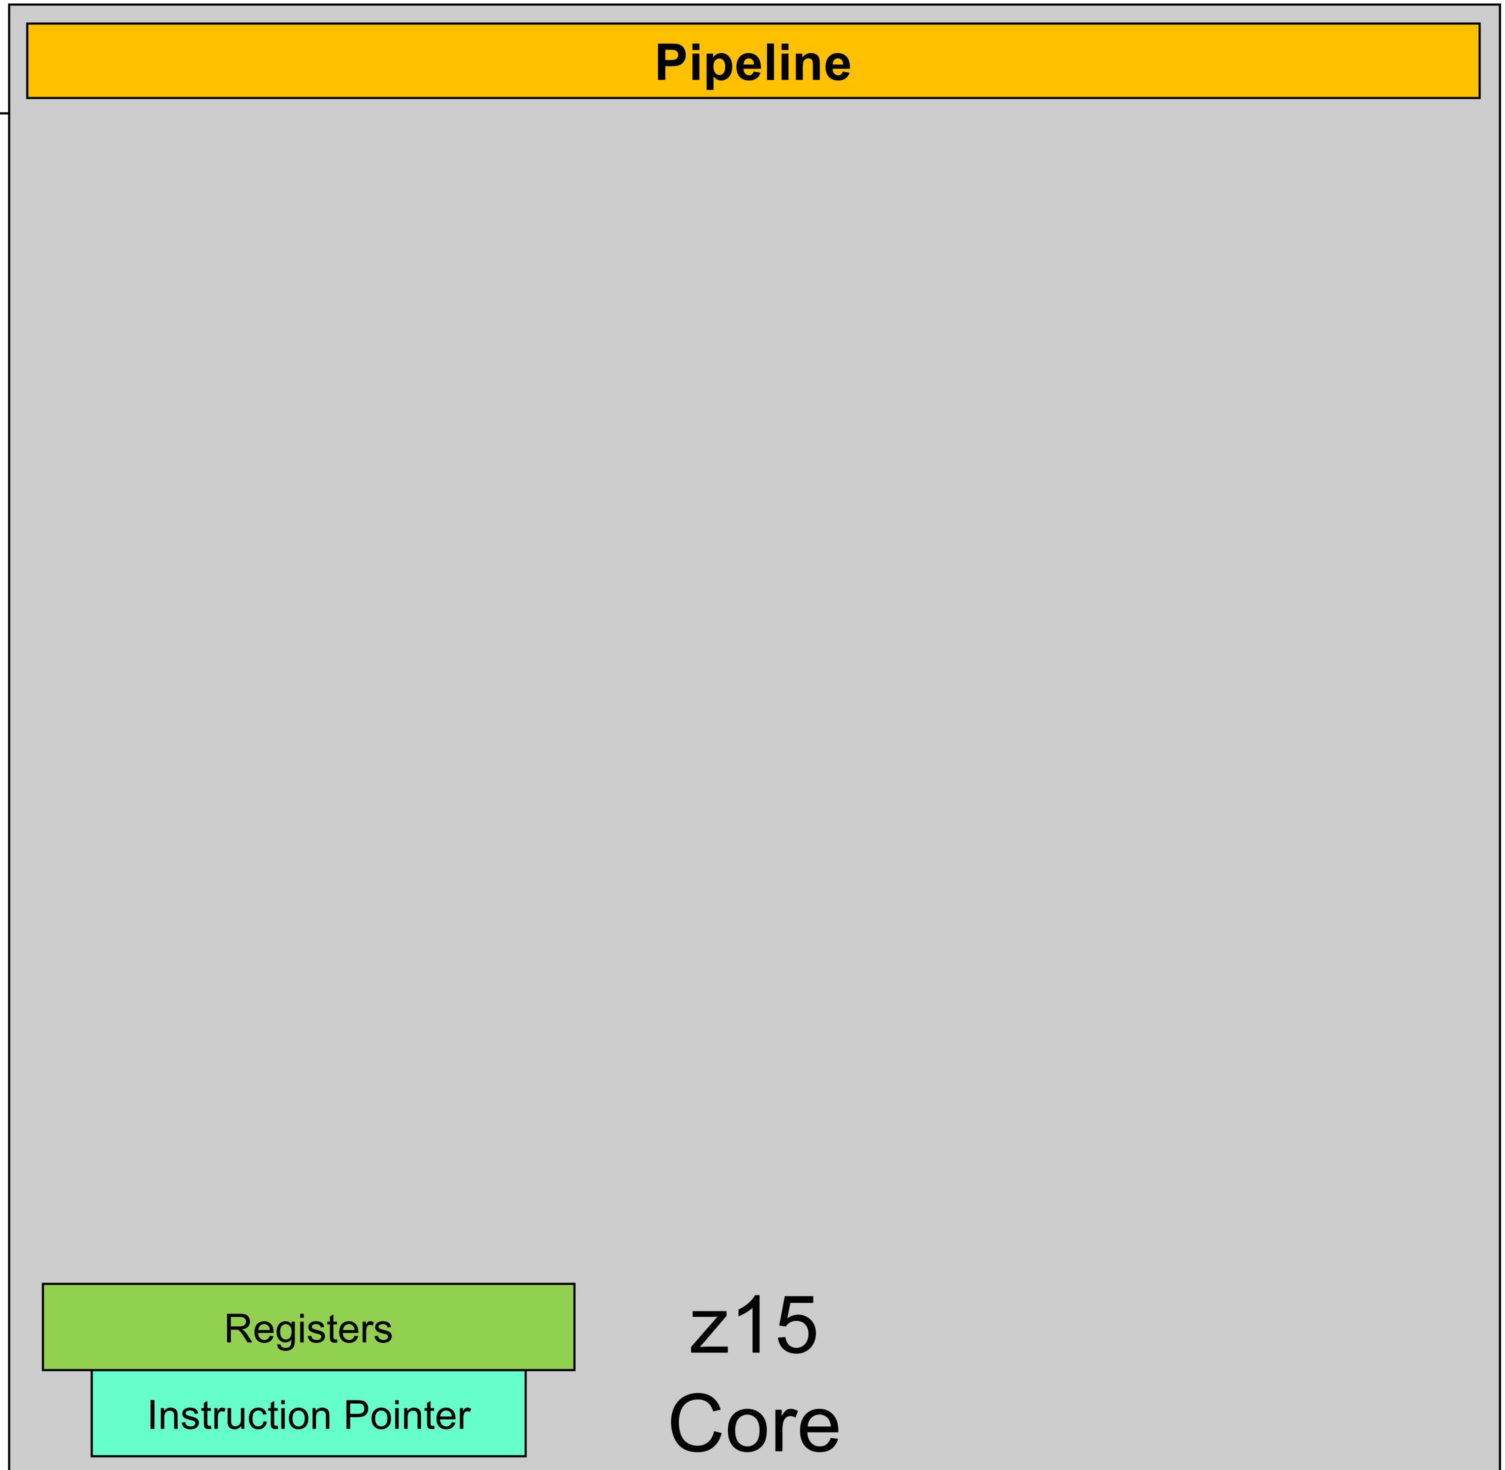
Pipeline
Registers
Instruction Pointer
z15
Core

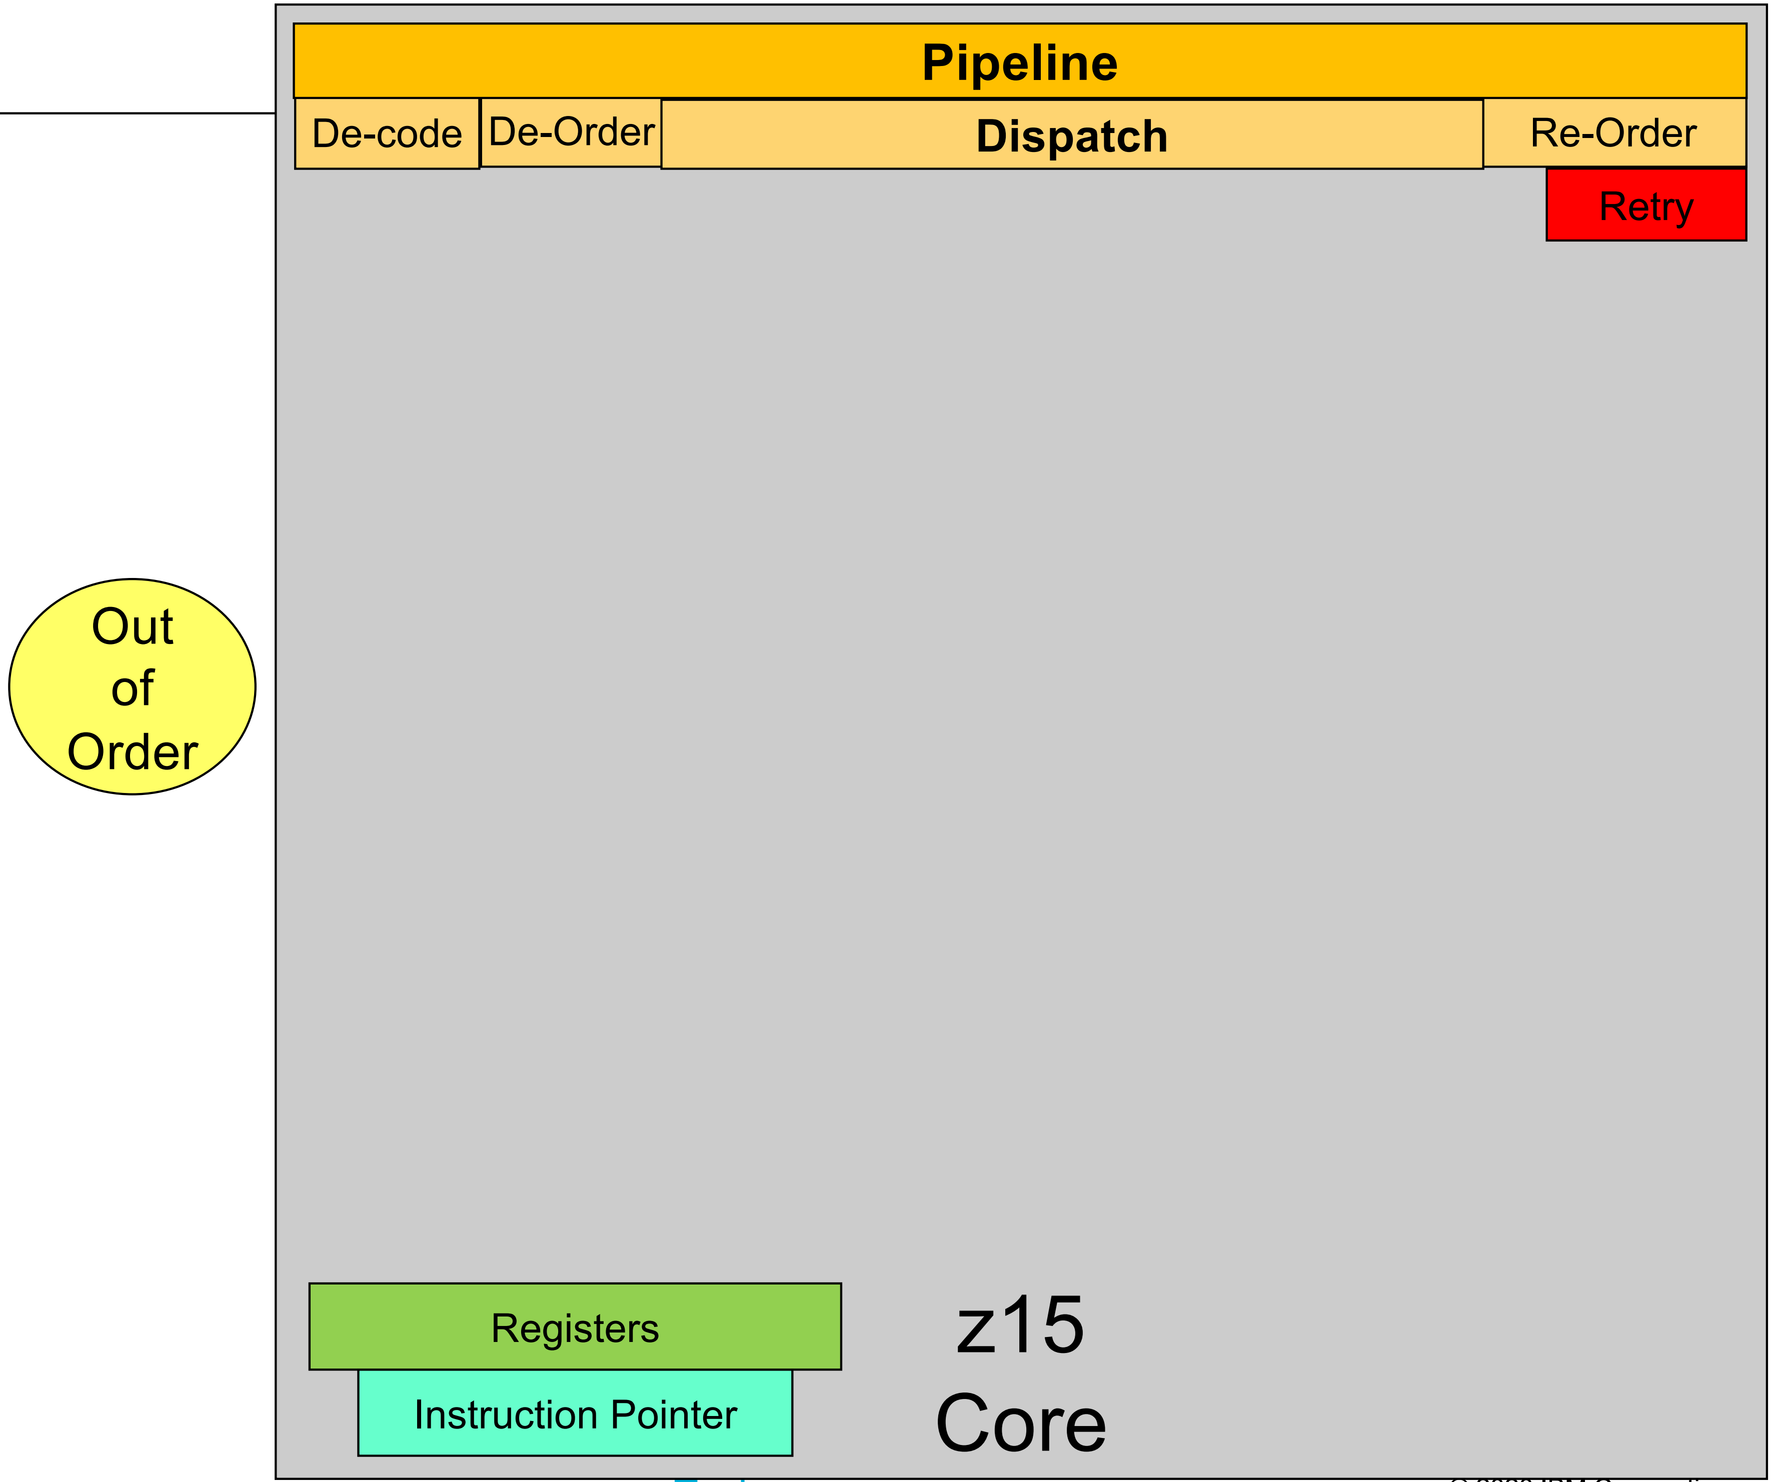
Pipeline
De-code
De-Order
Dispatch
Re-Order
Retry
Out of Order
Registers
Instruction Pointer
z15 Core

Topics

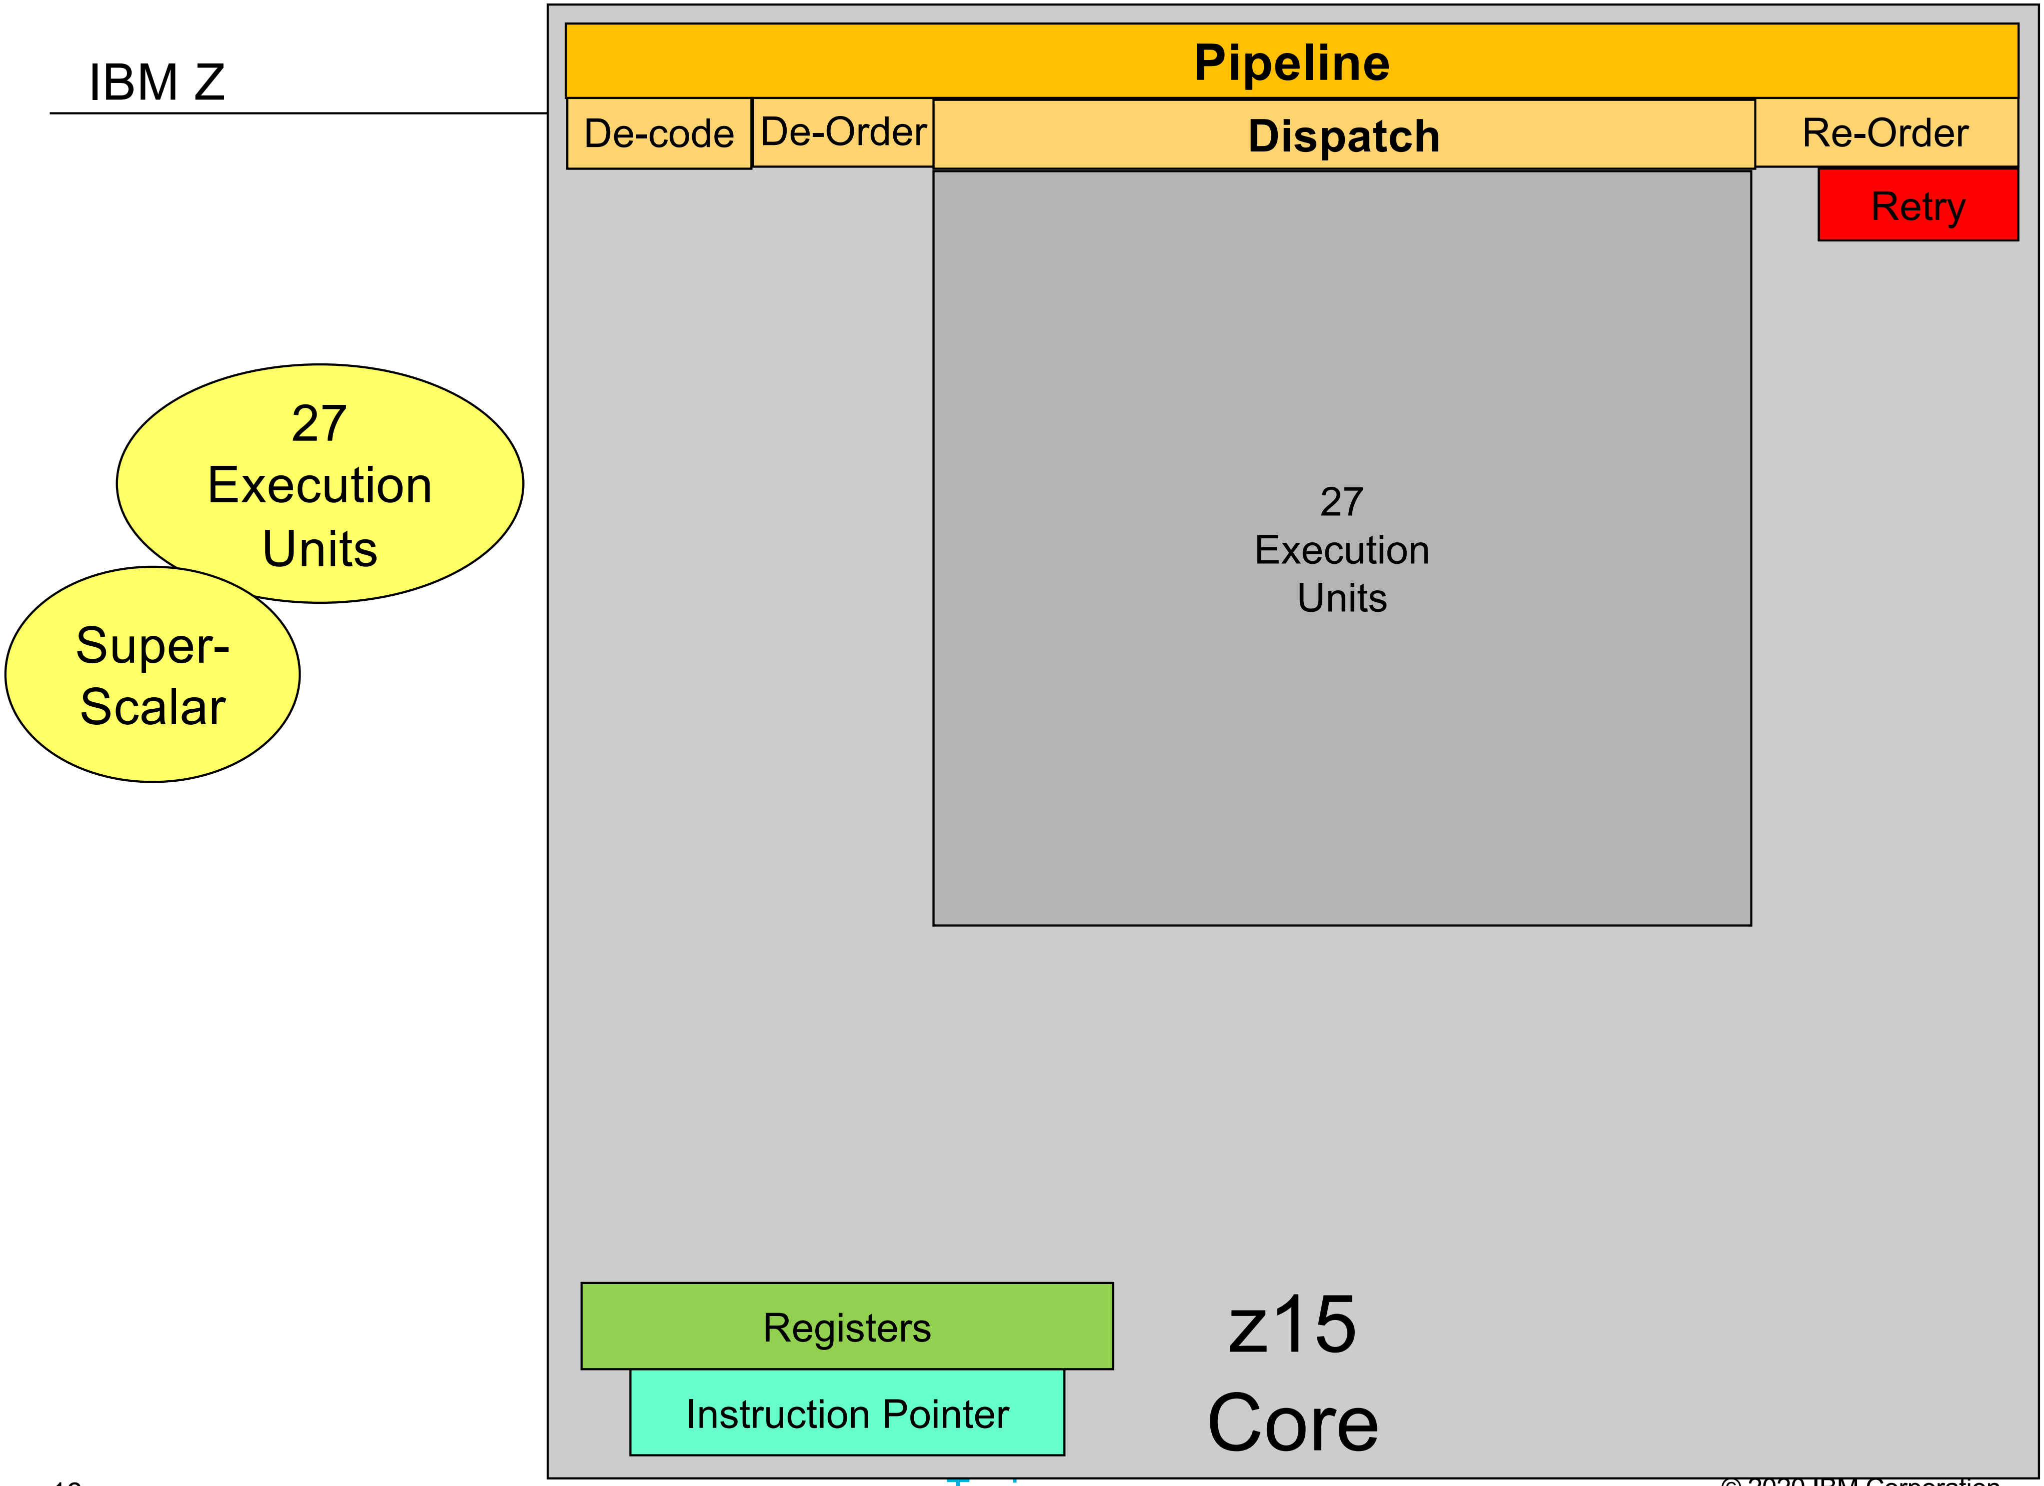
Pipeline
De-code
De-Order
Dispatch
Re-Order
Retry
27
Execution
Units
27
Execution
Units
Super-
Scalar
Registers
Instruction Pointer
z15
Core

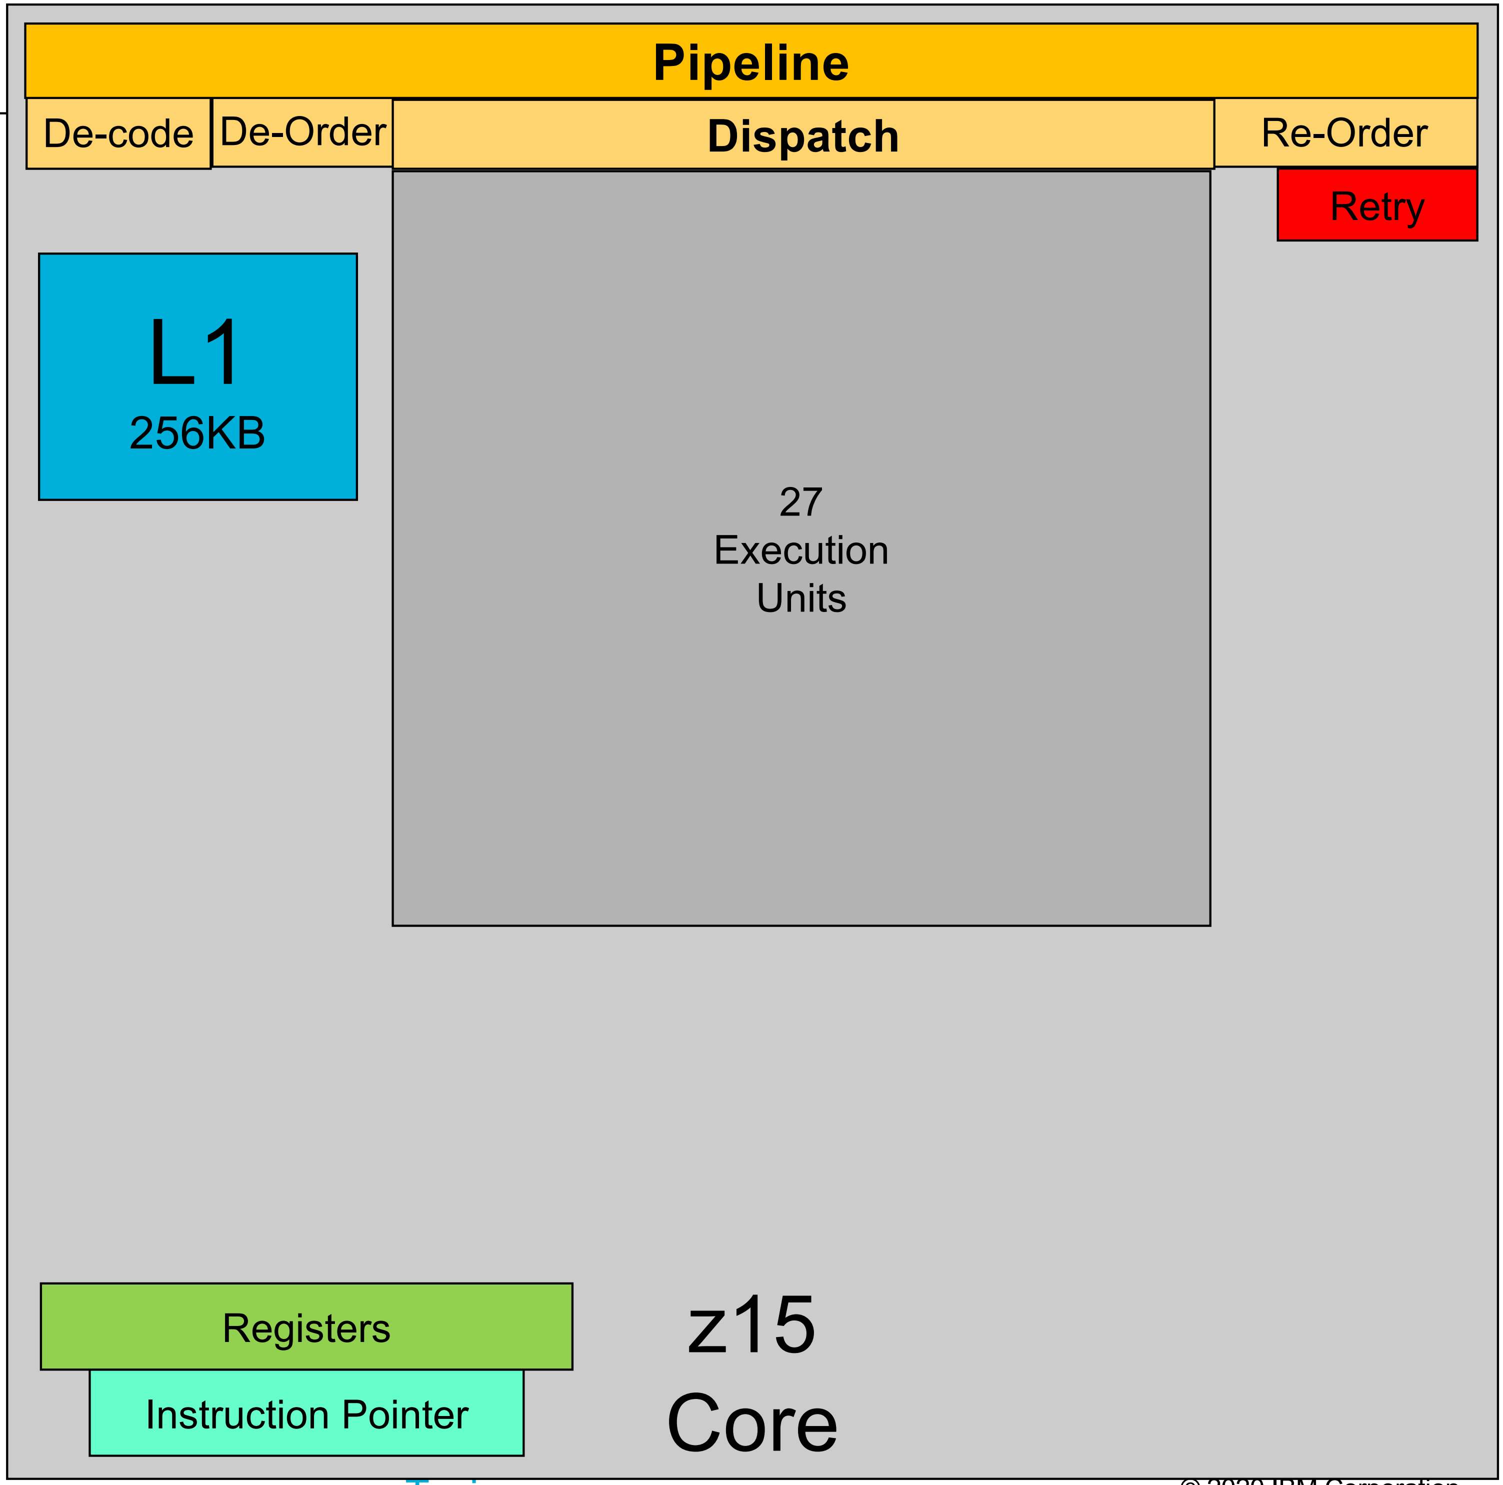
Pipeline
De-code
De-Order
Dispatch
Re-Order
Retry
L1
256KB
27
Execution
Units
Registers
Instruction Pointer
z15
Core

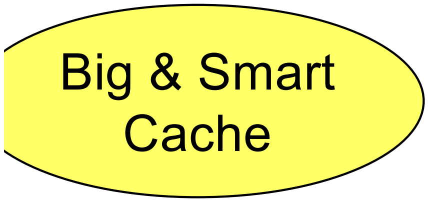
Big & Smart
Cache

Topics

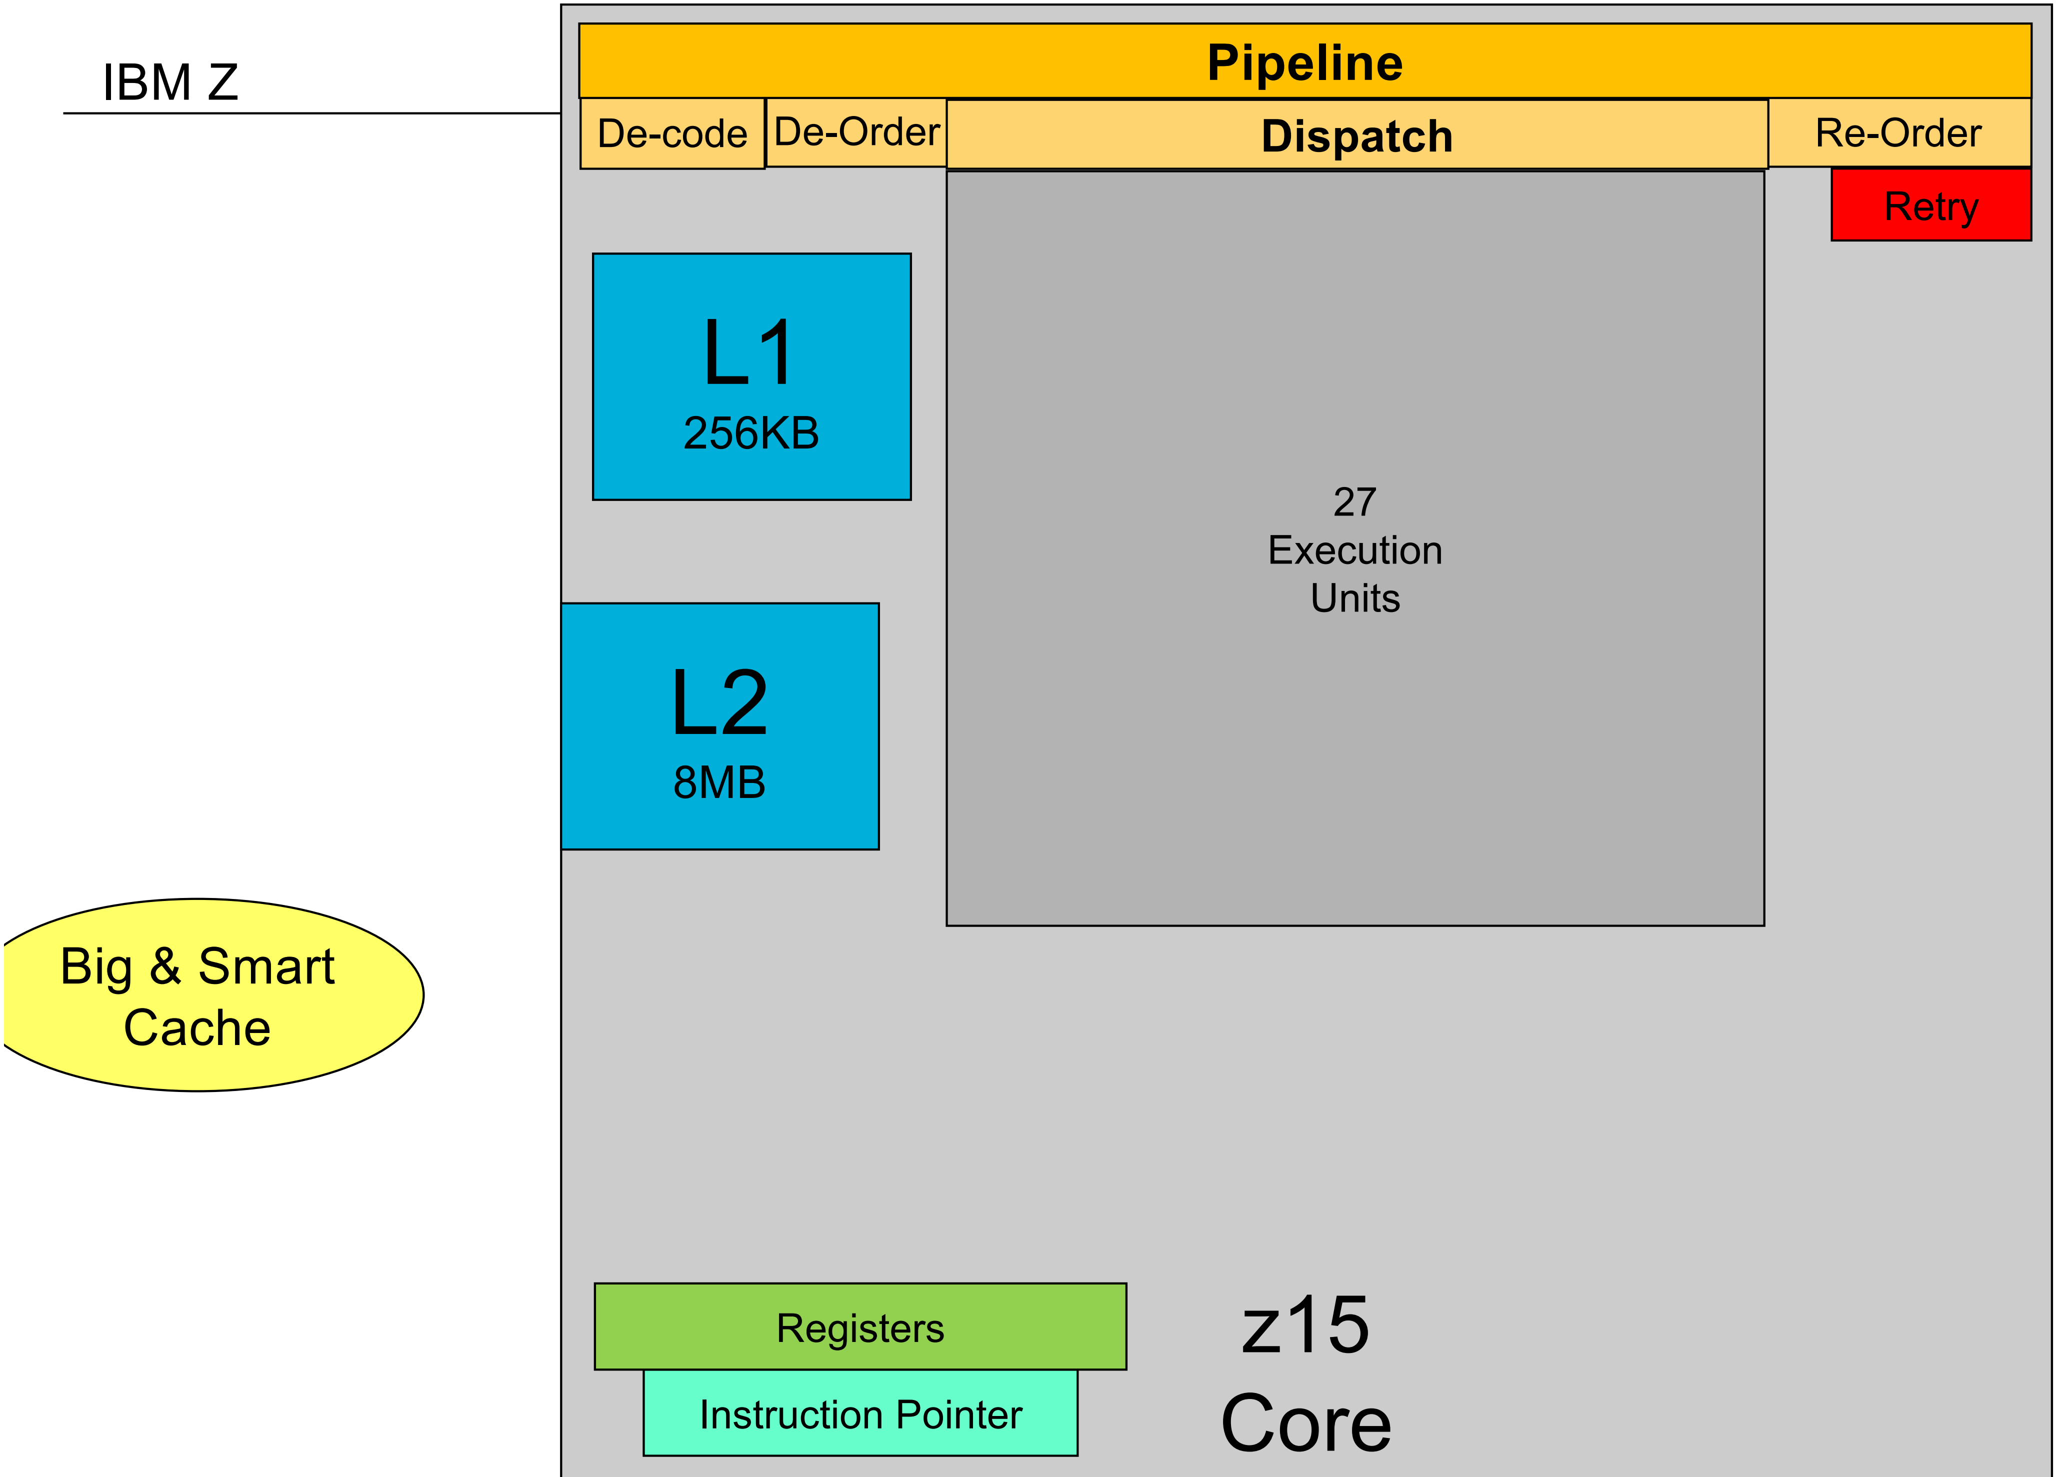
Pipeline
De-code
De-Order
Dispatch
Re-Order
Retry
L1
256KB
27
Execution
Units
L2
8MB
Big & Smart
Cache
Registers
Instruction Pointer
z15
Core

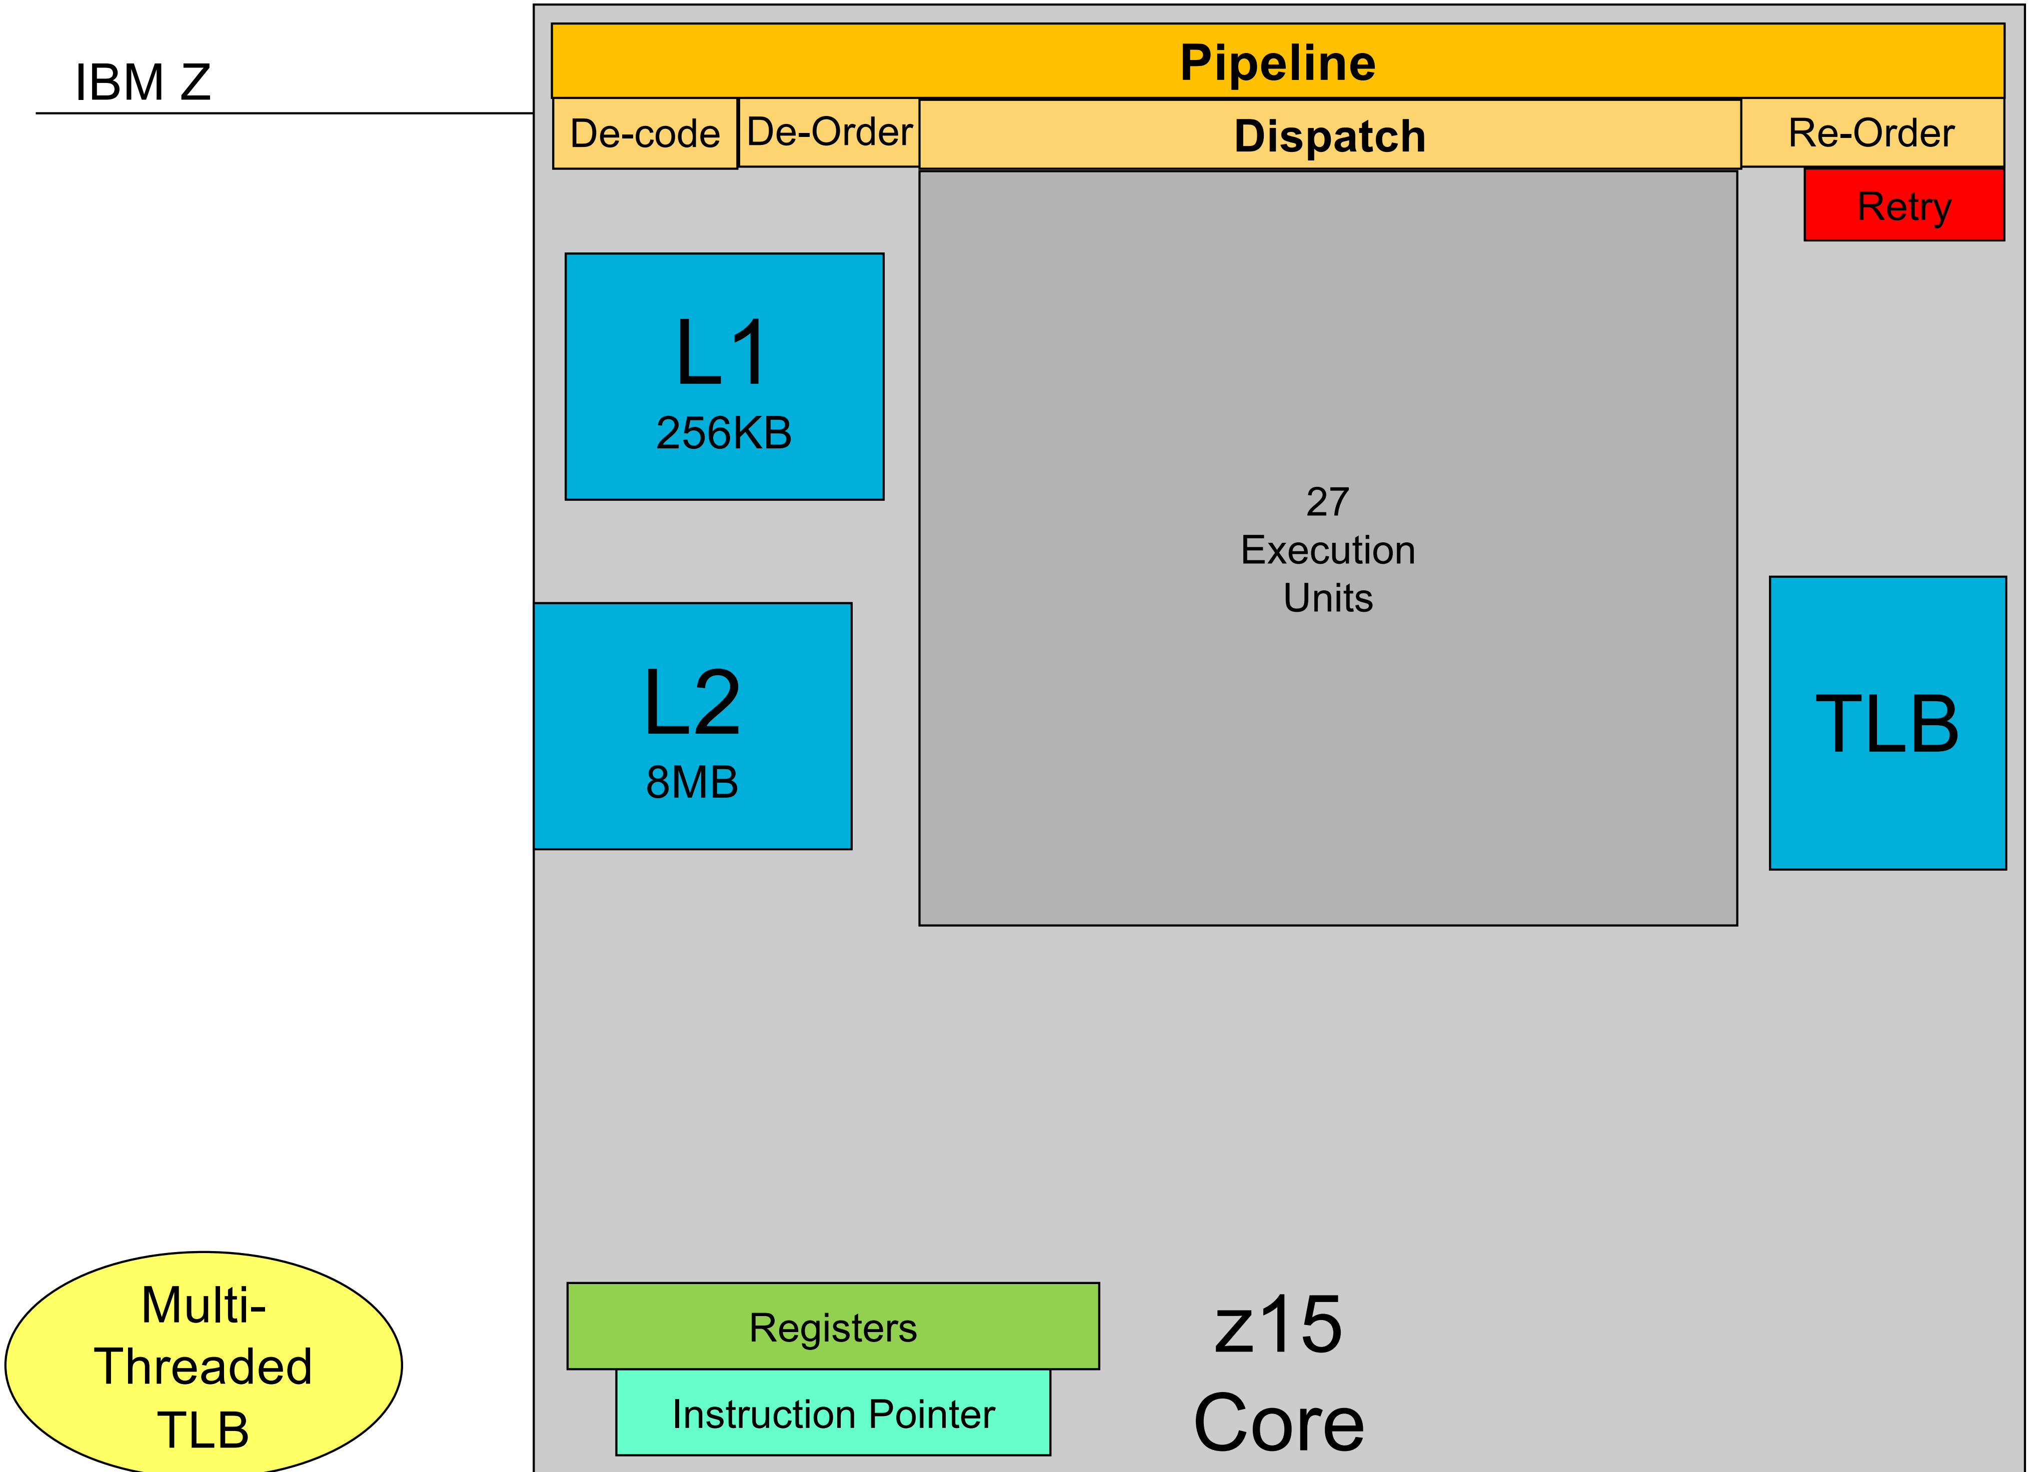
Pipeline
De-code
De-Order
Dispatch
Re-Order
Retry
L1
256KB
27
Execution
Units
L2
8MB
TLB
Registers
Instruction Pointer
z15
Core
Multi-
Threaded
TLB

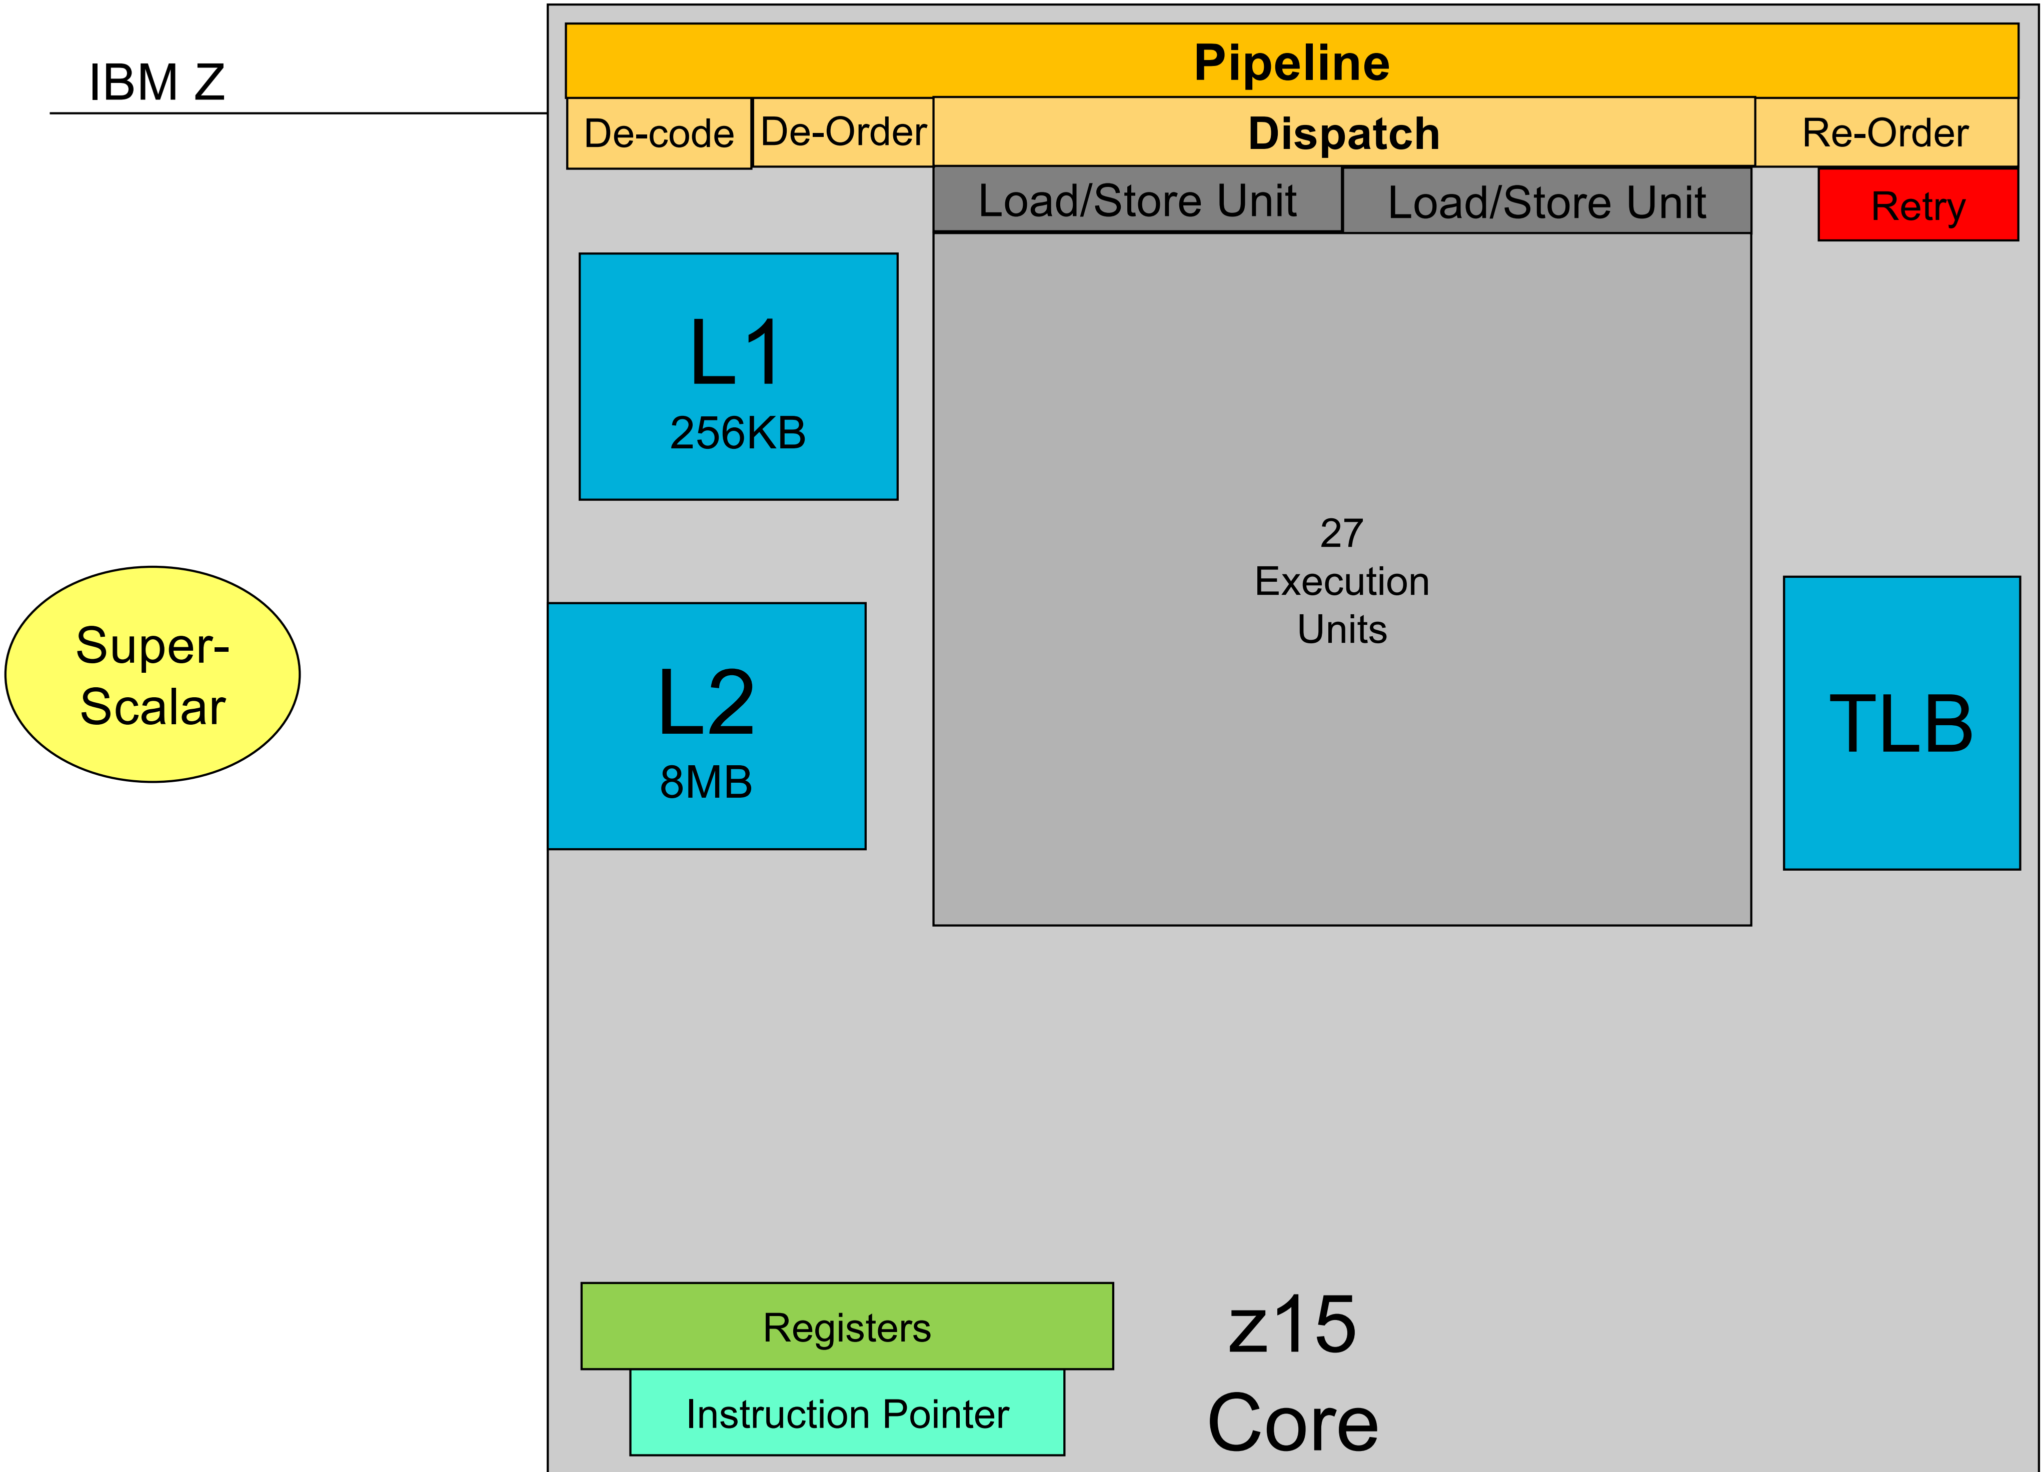
Super-
Scalar
Pipeline
De-code
De-Order
Dispatch
Re-Order
Load/Store Unit
Load/Store Unit
Retry
L1
256KB
27
Execution
Units
L2
8MB
TLB
Registers
Instruction Pointer
z15
Core

Topics

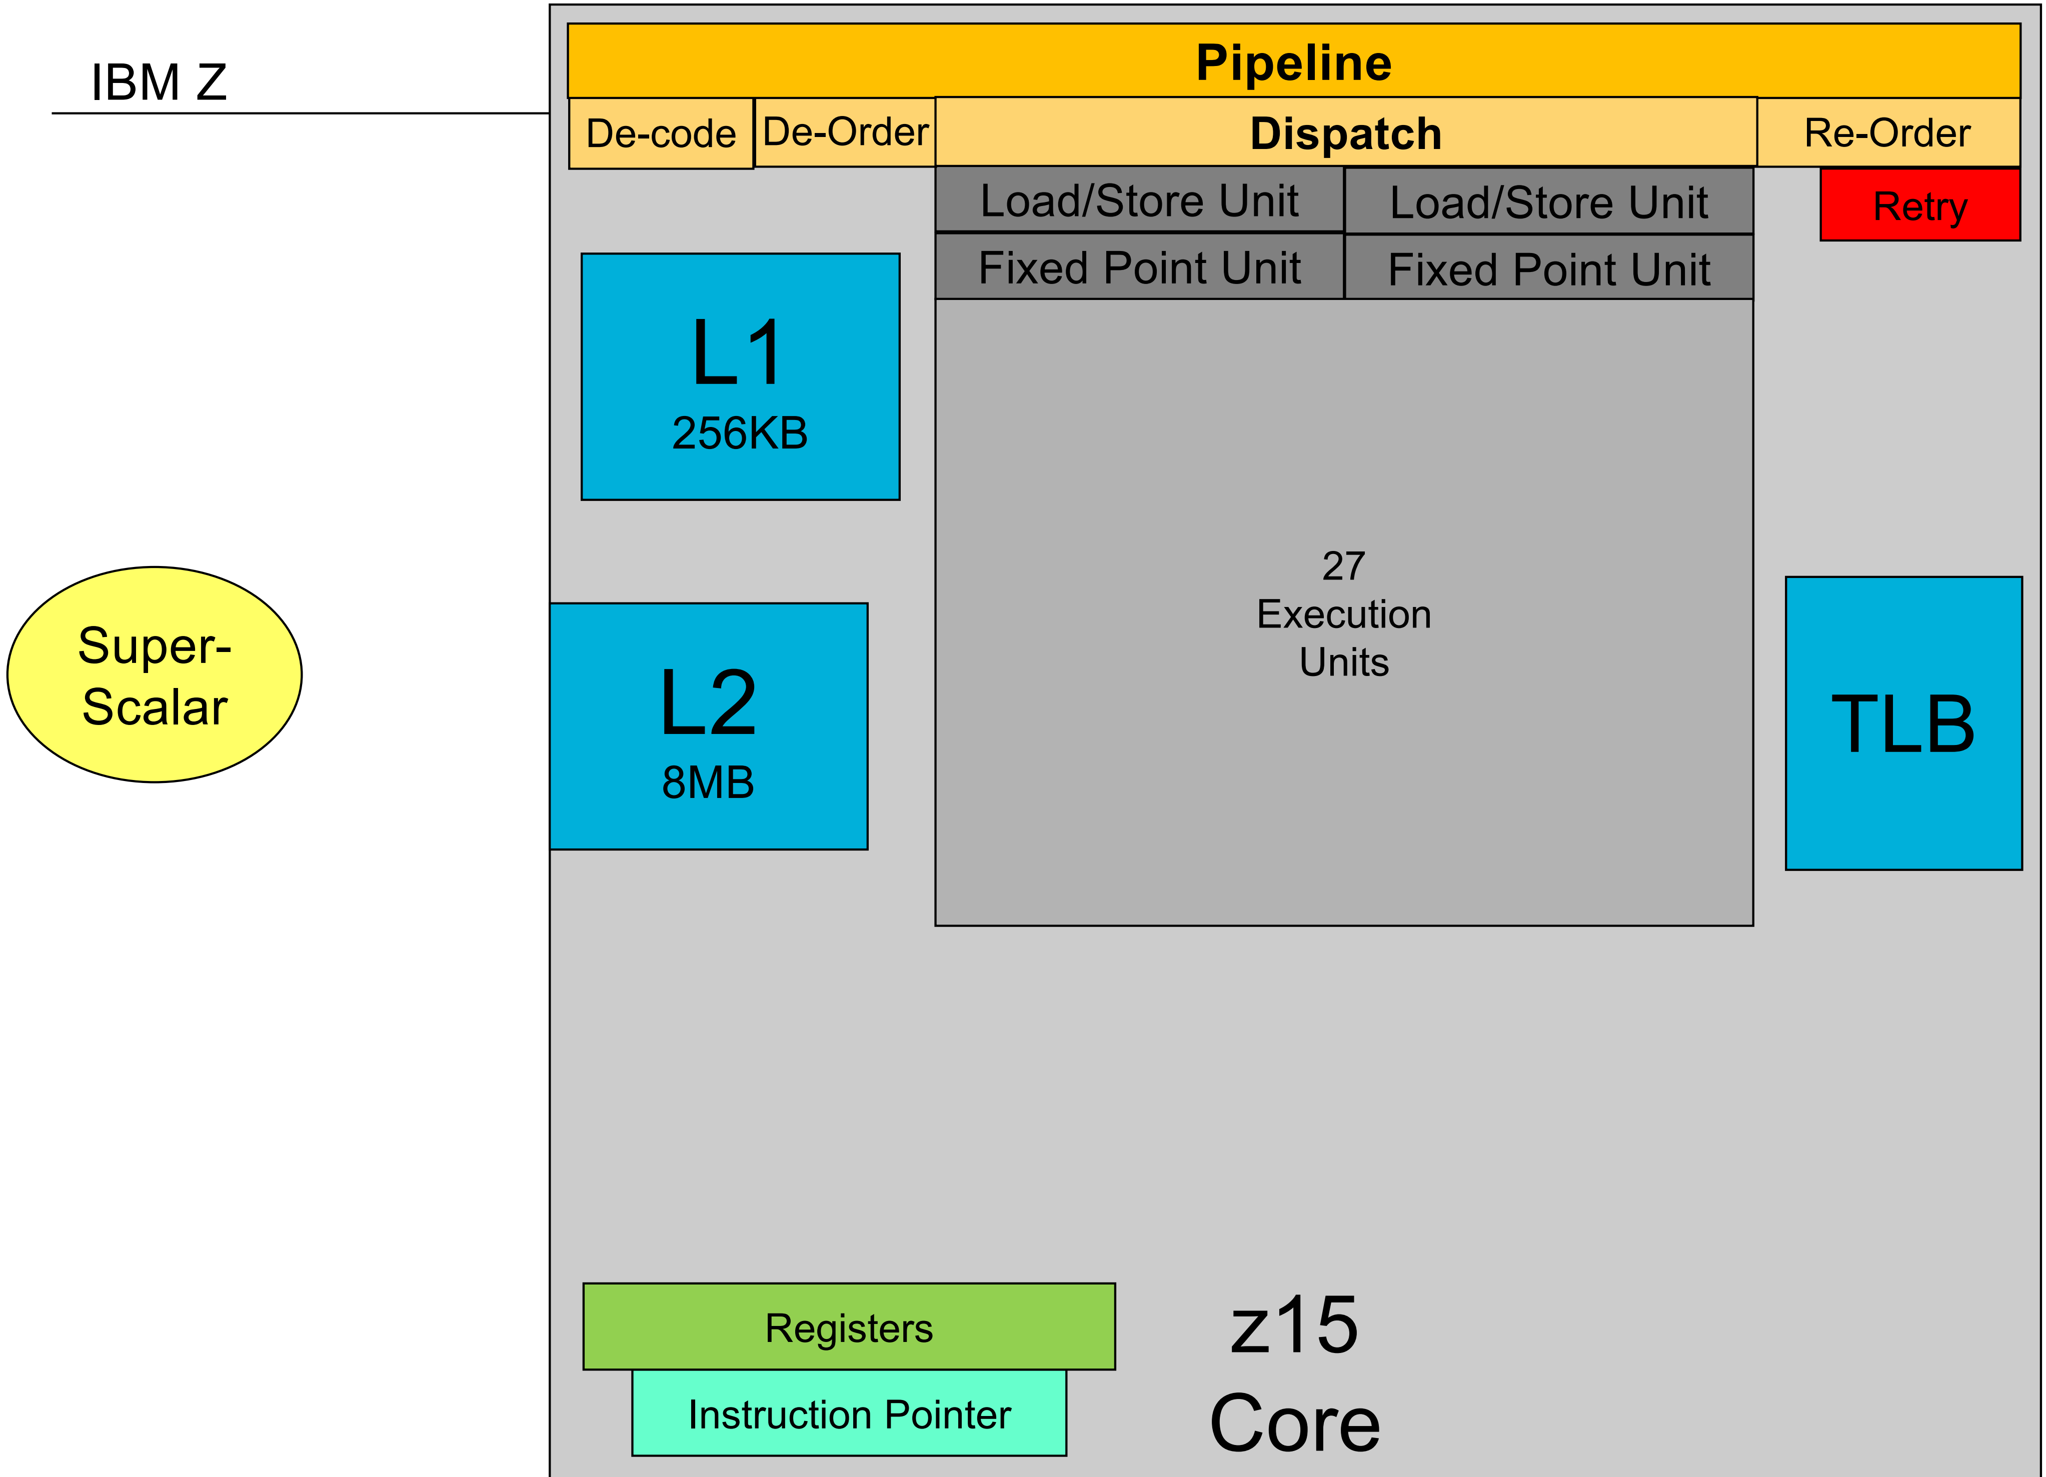
Super-
Scalar
Pipeline
De-code
De-Order
Dispatch
Re-Order
Load/Store Unit
Load/Store Unit
Retry
Fixed Point Unit
Fixed Point Unit
L1
256KB
27
Execution
Units
L2
8MB
TLB
Registers
Instruction Pointer
z15
Core

Topics

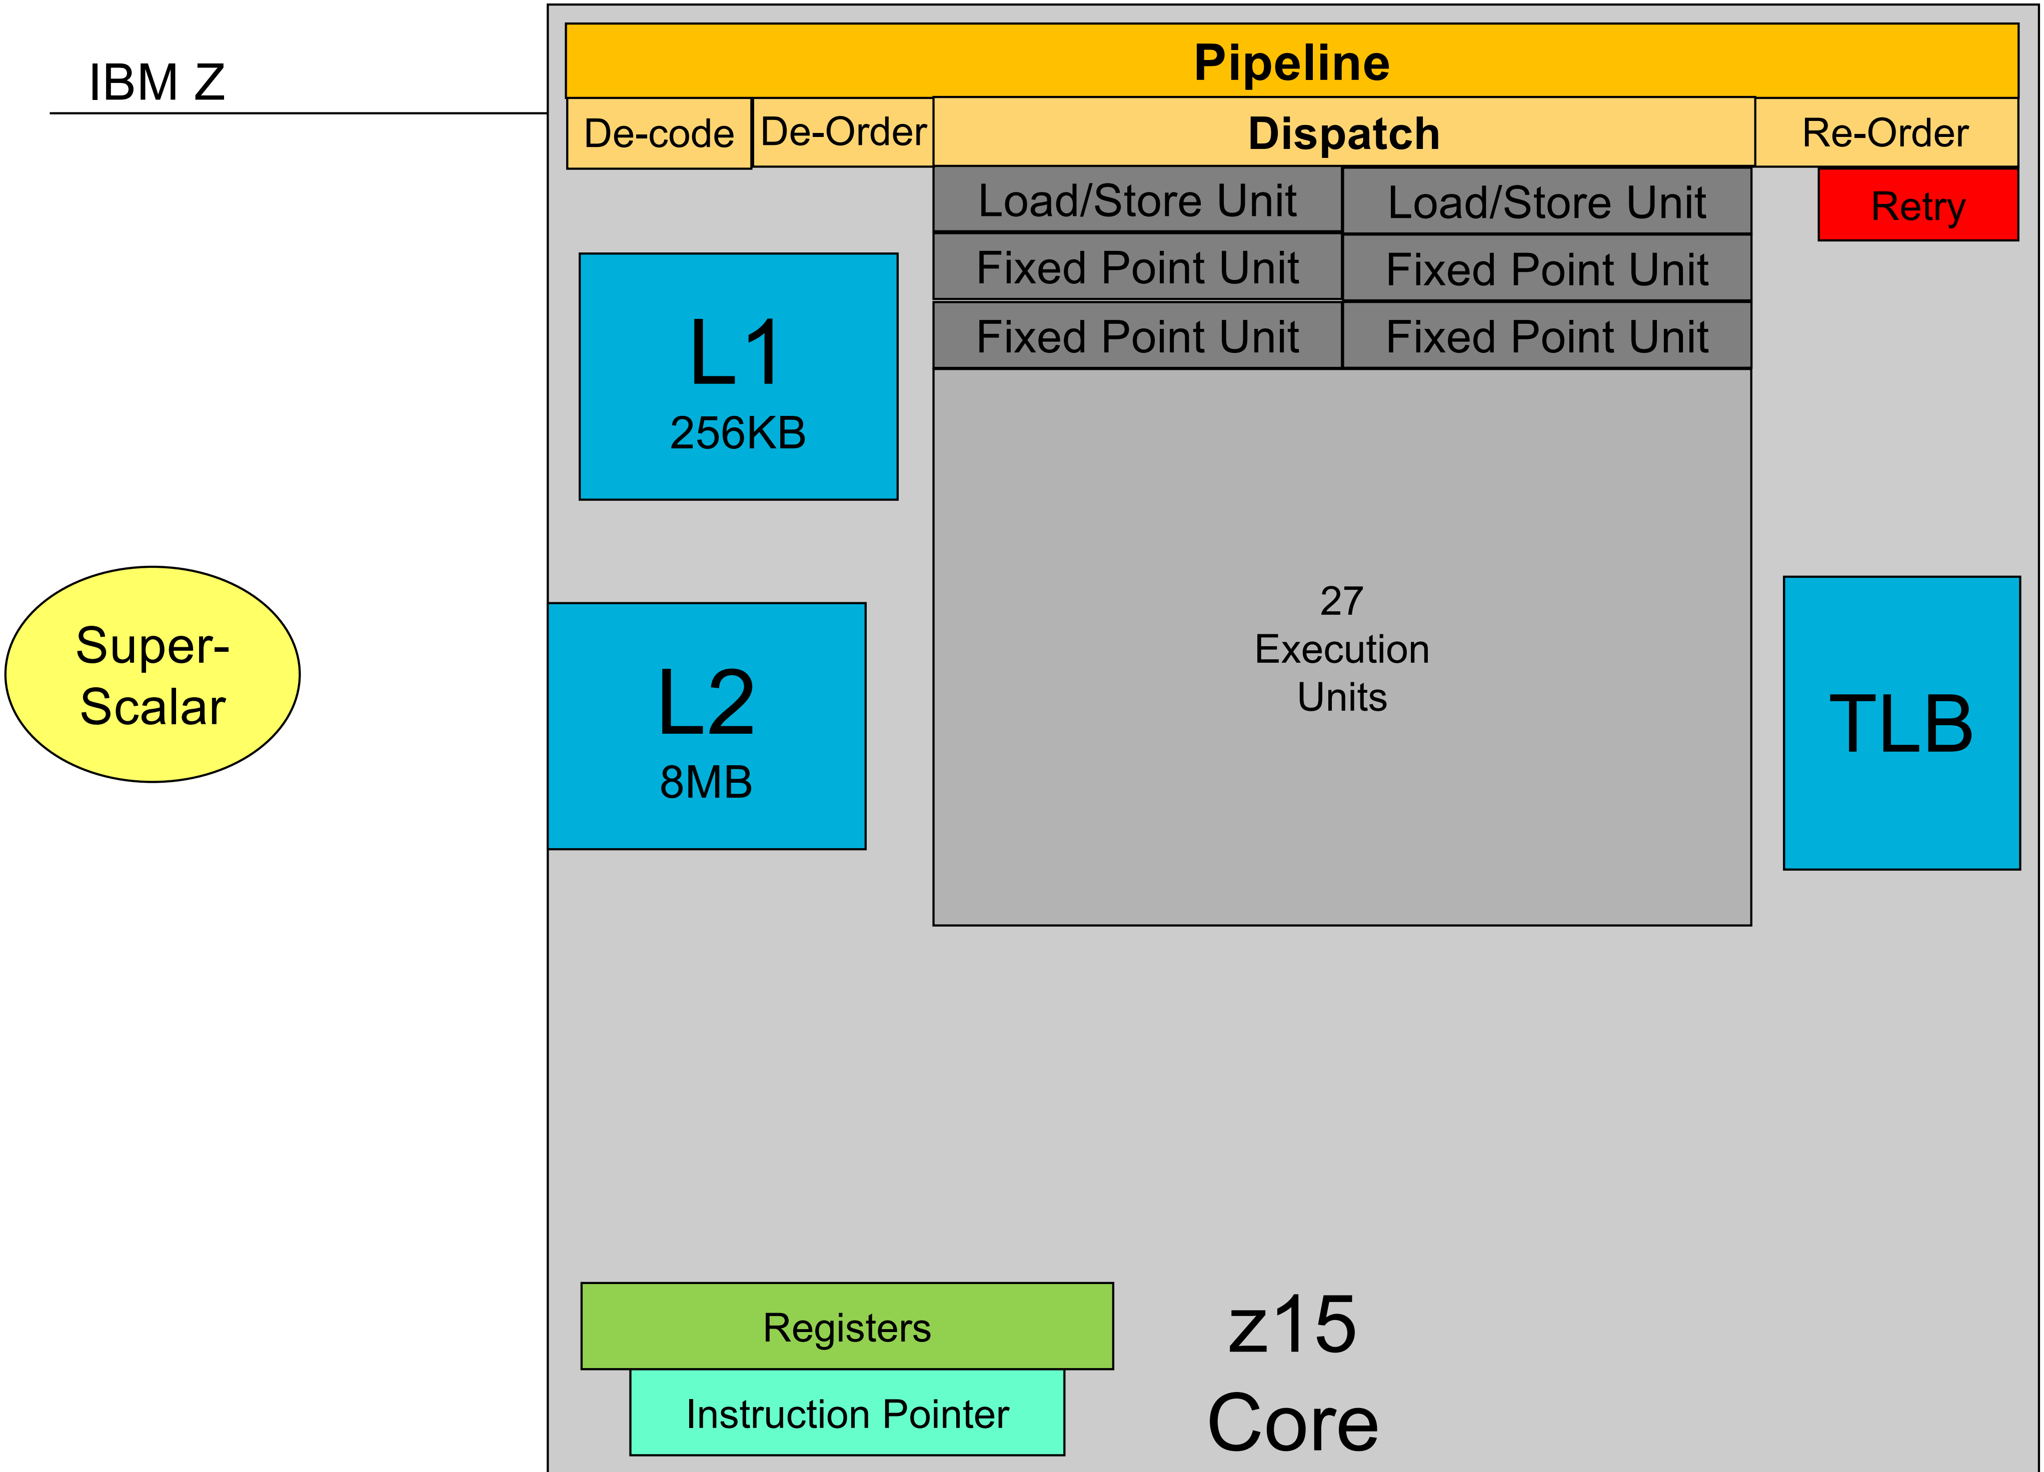
Super-
Scalar
Pipeline
De-code
De-Order
Dispatch
Re-Order
Load/Store Unit
Load/Store Unit
Retry
Fixed Point Unit
Fixed Point Unit
Fixed Point Unit
Fixed Point Unit
L1
256KB
L2
8MB
27
Execution
Units
TLB
Registers
Instruction Pointer
z15
Core

Topics

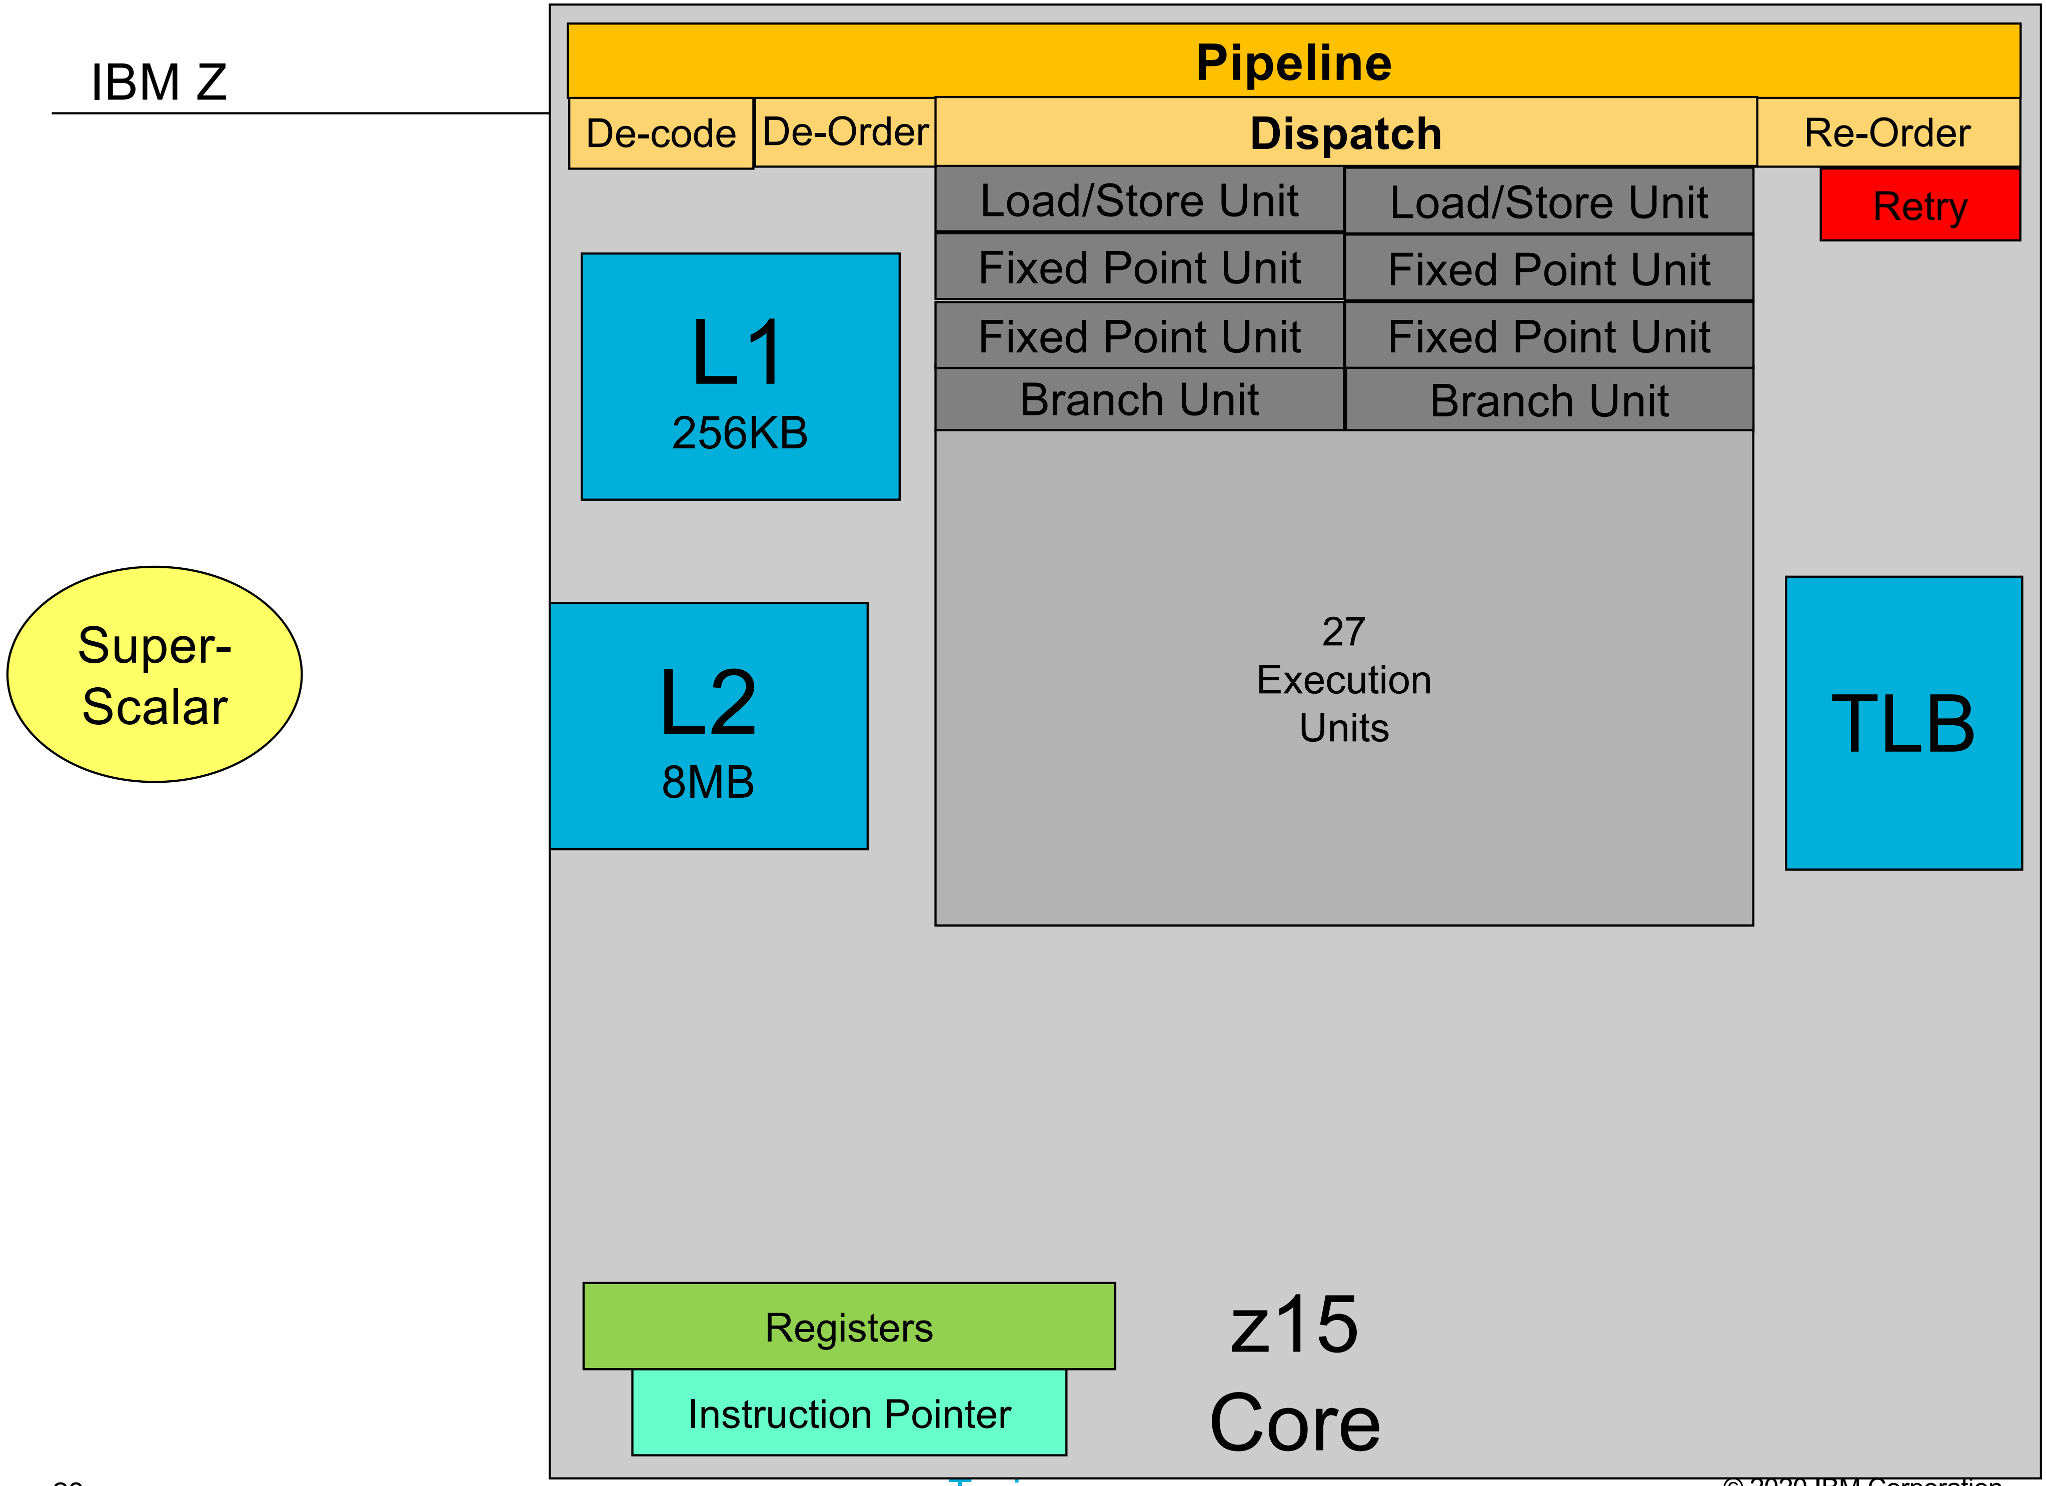
Pipeline
De-code
De-Order
Dispatch
Re-Order
Load/Store Unit
Load/Store Unit
Retry
Fixed Point Unit
Fixed Point Unit
Fixed Point Unit
Fixed Point Unit
Branch Unit
Branch Unit
L1
256KB
L2
8MB
27
Execution
Units
TLB
Super-
Scalar
Registers
Instruction Pointer
z15
Core

Topics

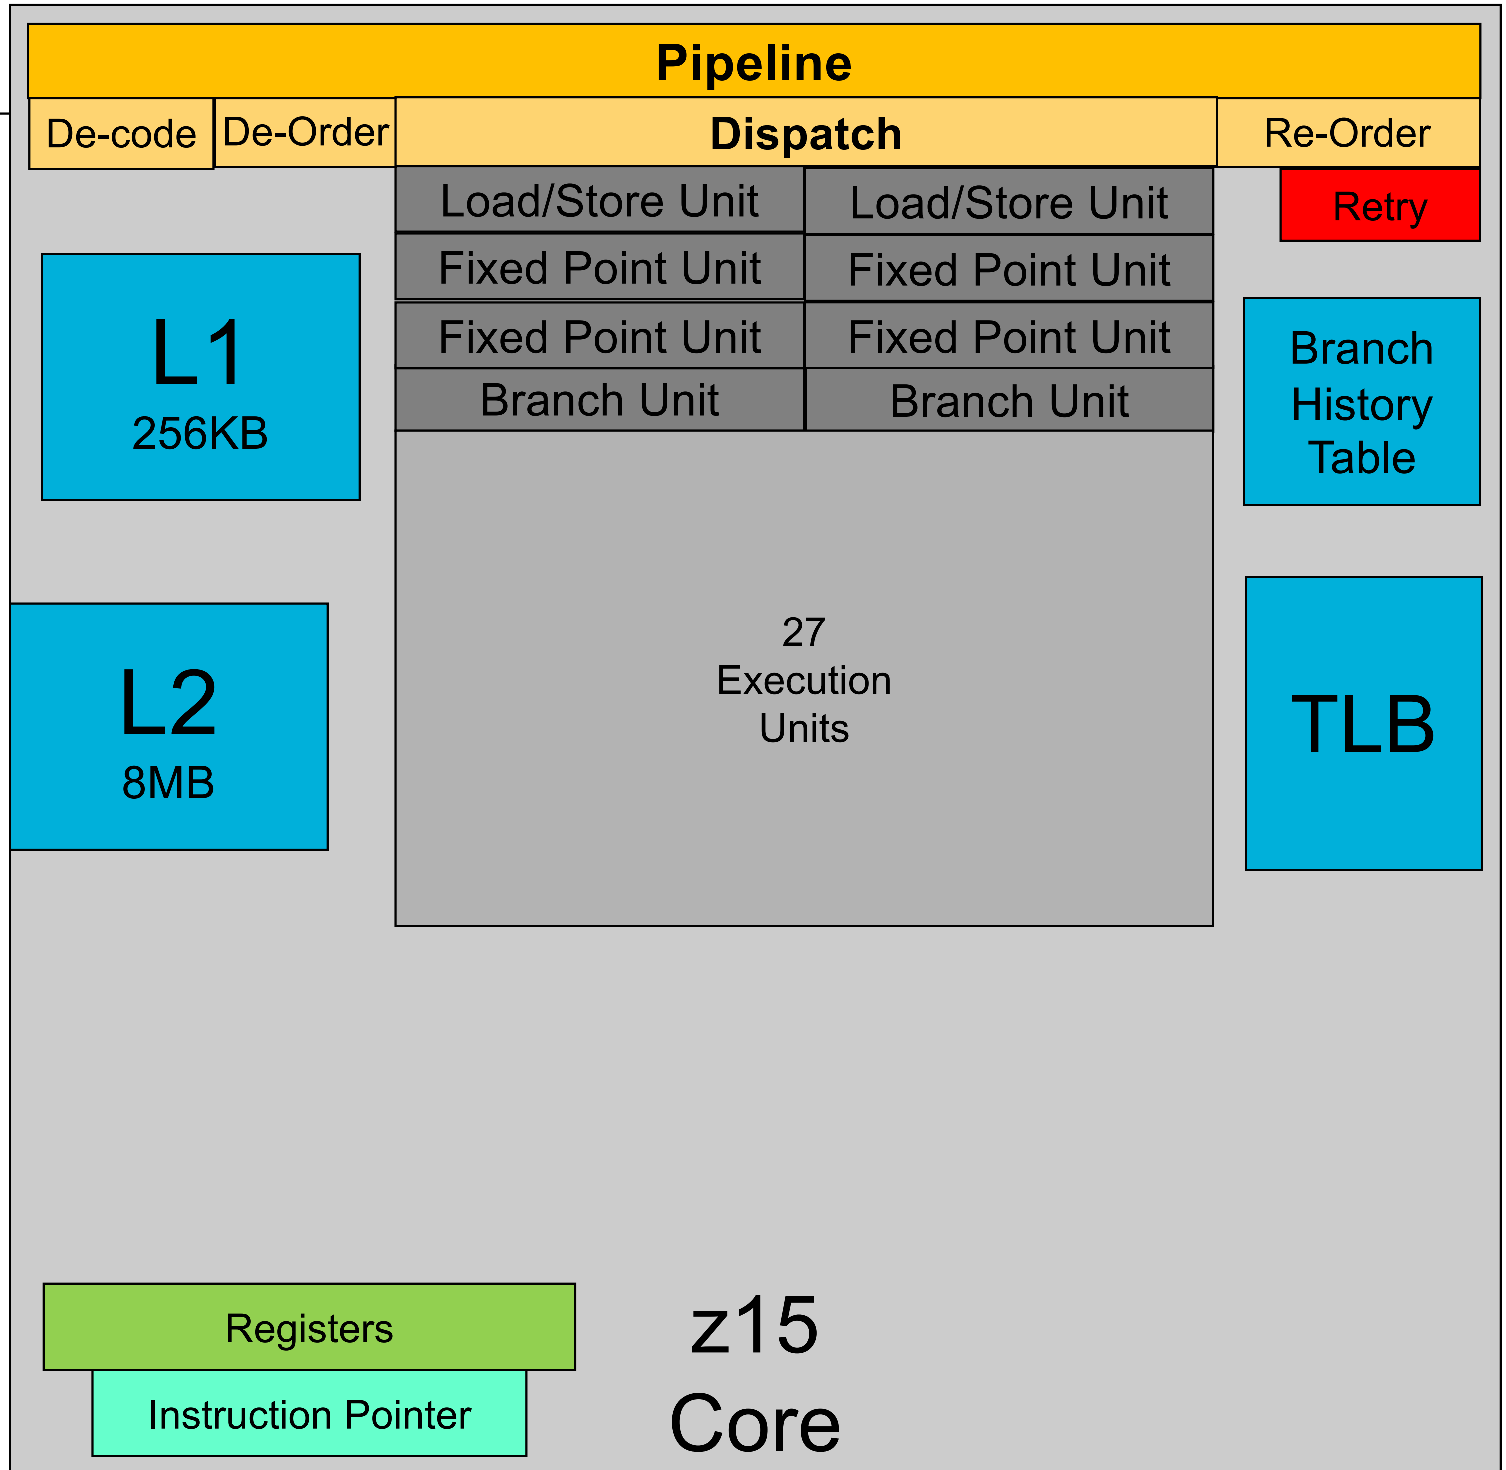
Pipeline
De-code
De-Order
Dispatch
Re-Order
Load/Store Unit
Load/Store Unit
Retry
Fixed Point Unit
Fixed Point Unit
Fixed Point Unit
Fixed Point Unit
Branch Unit
Branch Unit
L1
256KB
Branch
History
Table
27
Execution
Units
L2
8MB
TLB
Registers
Instruction Pointer
z15
Core

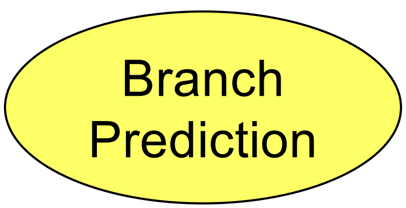
Branch
Prediction

Topics

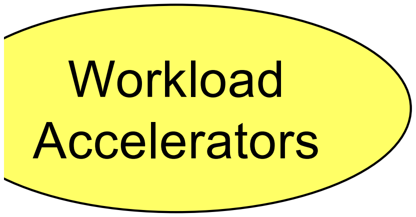
Workload
Accelerators

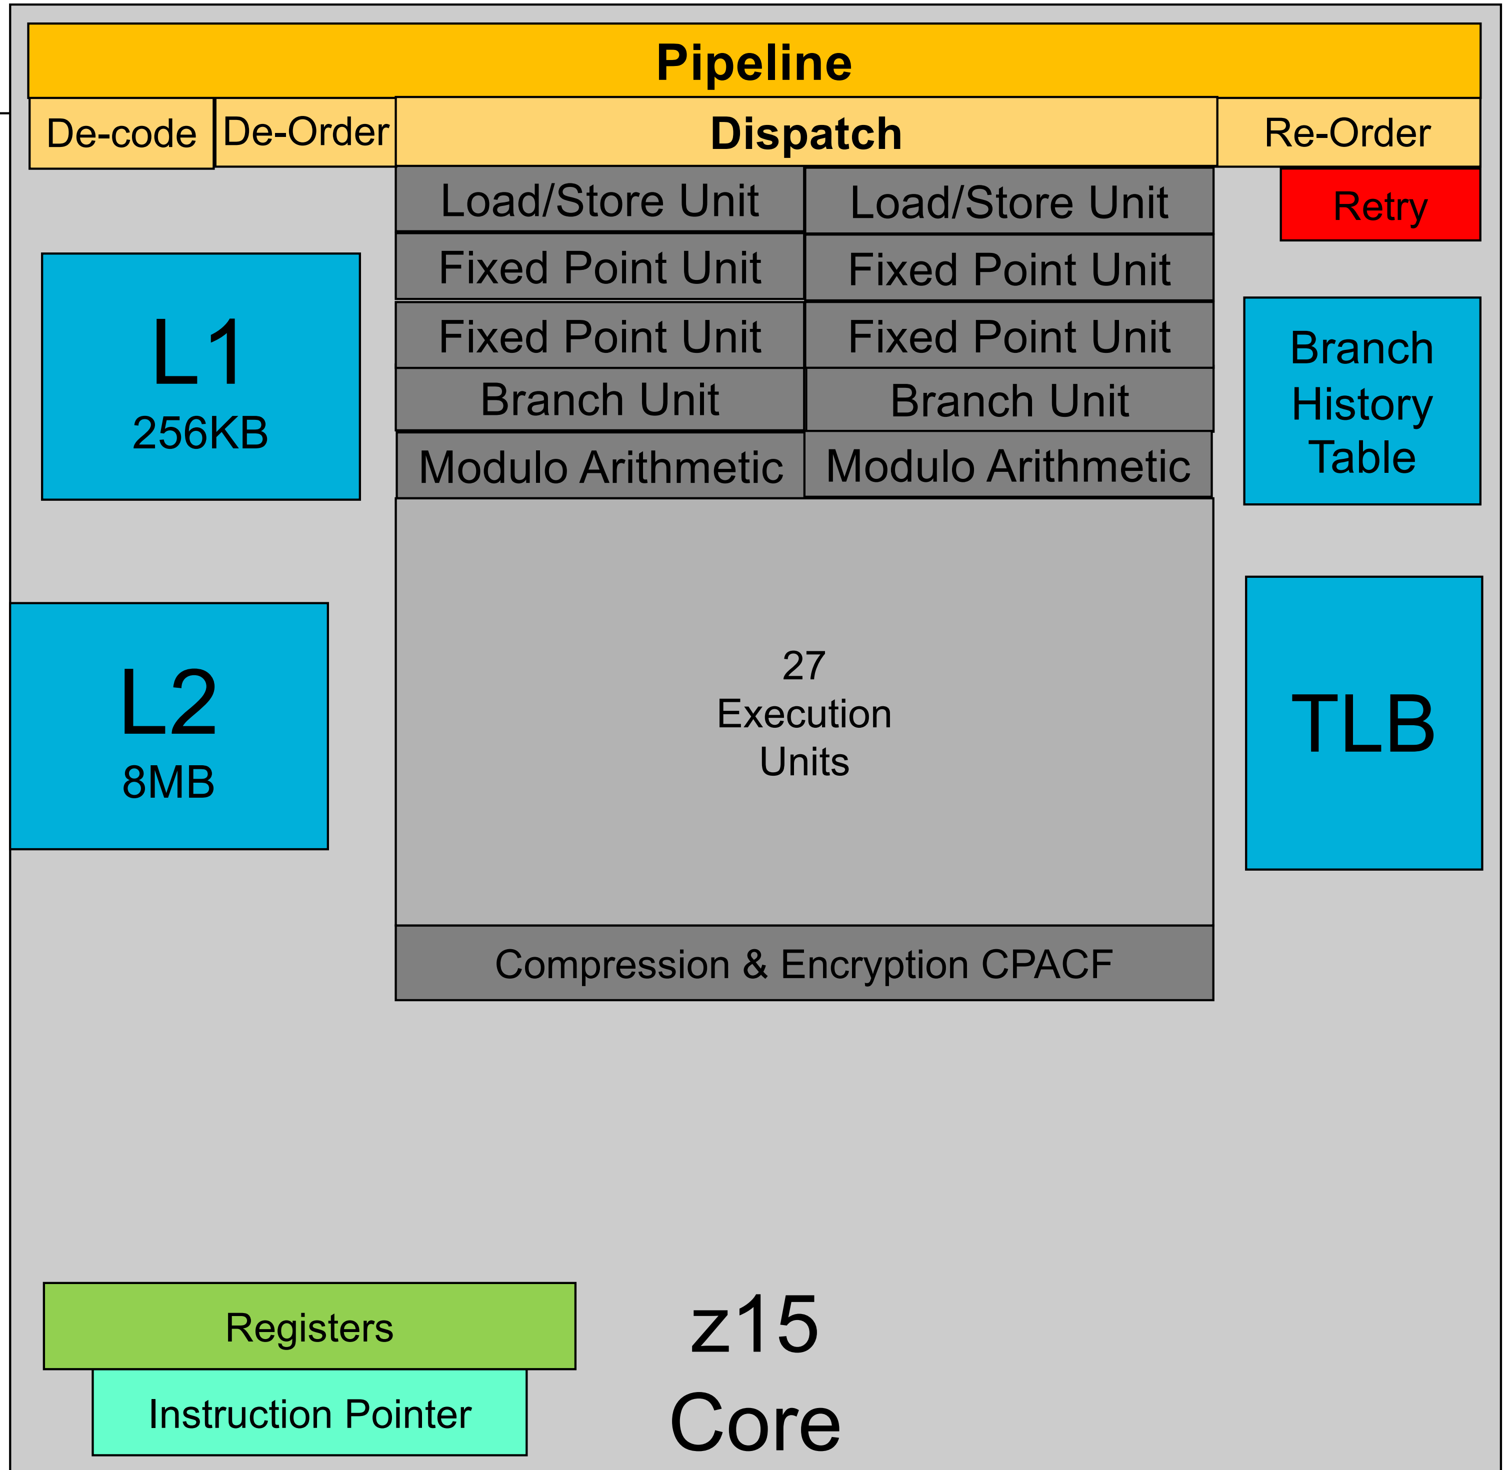
Pipeline
De-code
De-Order
Dispatch
Re-Order
Load/Store Unit
Load/Store Unit
Retry
Fixed Point Unit
Fixed Point Unit
Fixed Point Unit
Fixed Point Unit
Branch Unit
Branch Unit
Modulo Arithmetic
Modulo Arithmetic
L1
256KB
Branch
History
Table
L2
8MB
27
Execution
Units
TLB
Compression & Encryption CPACF
Registers
Instruction Pointer
z15
Core

Topics

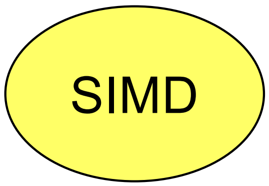
SIMD

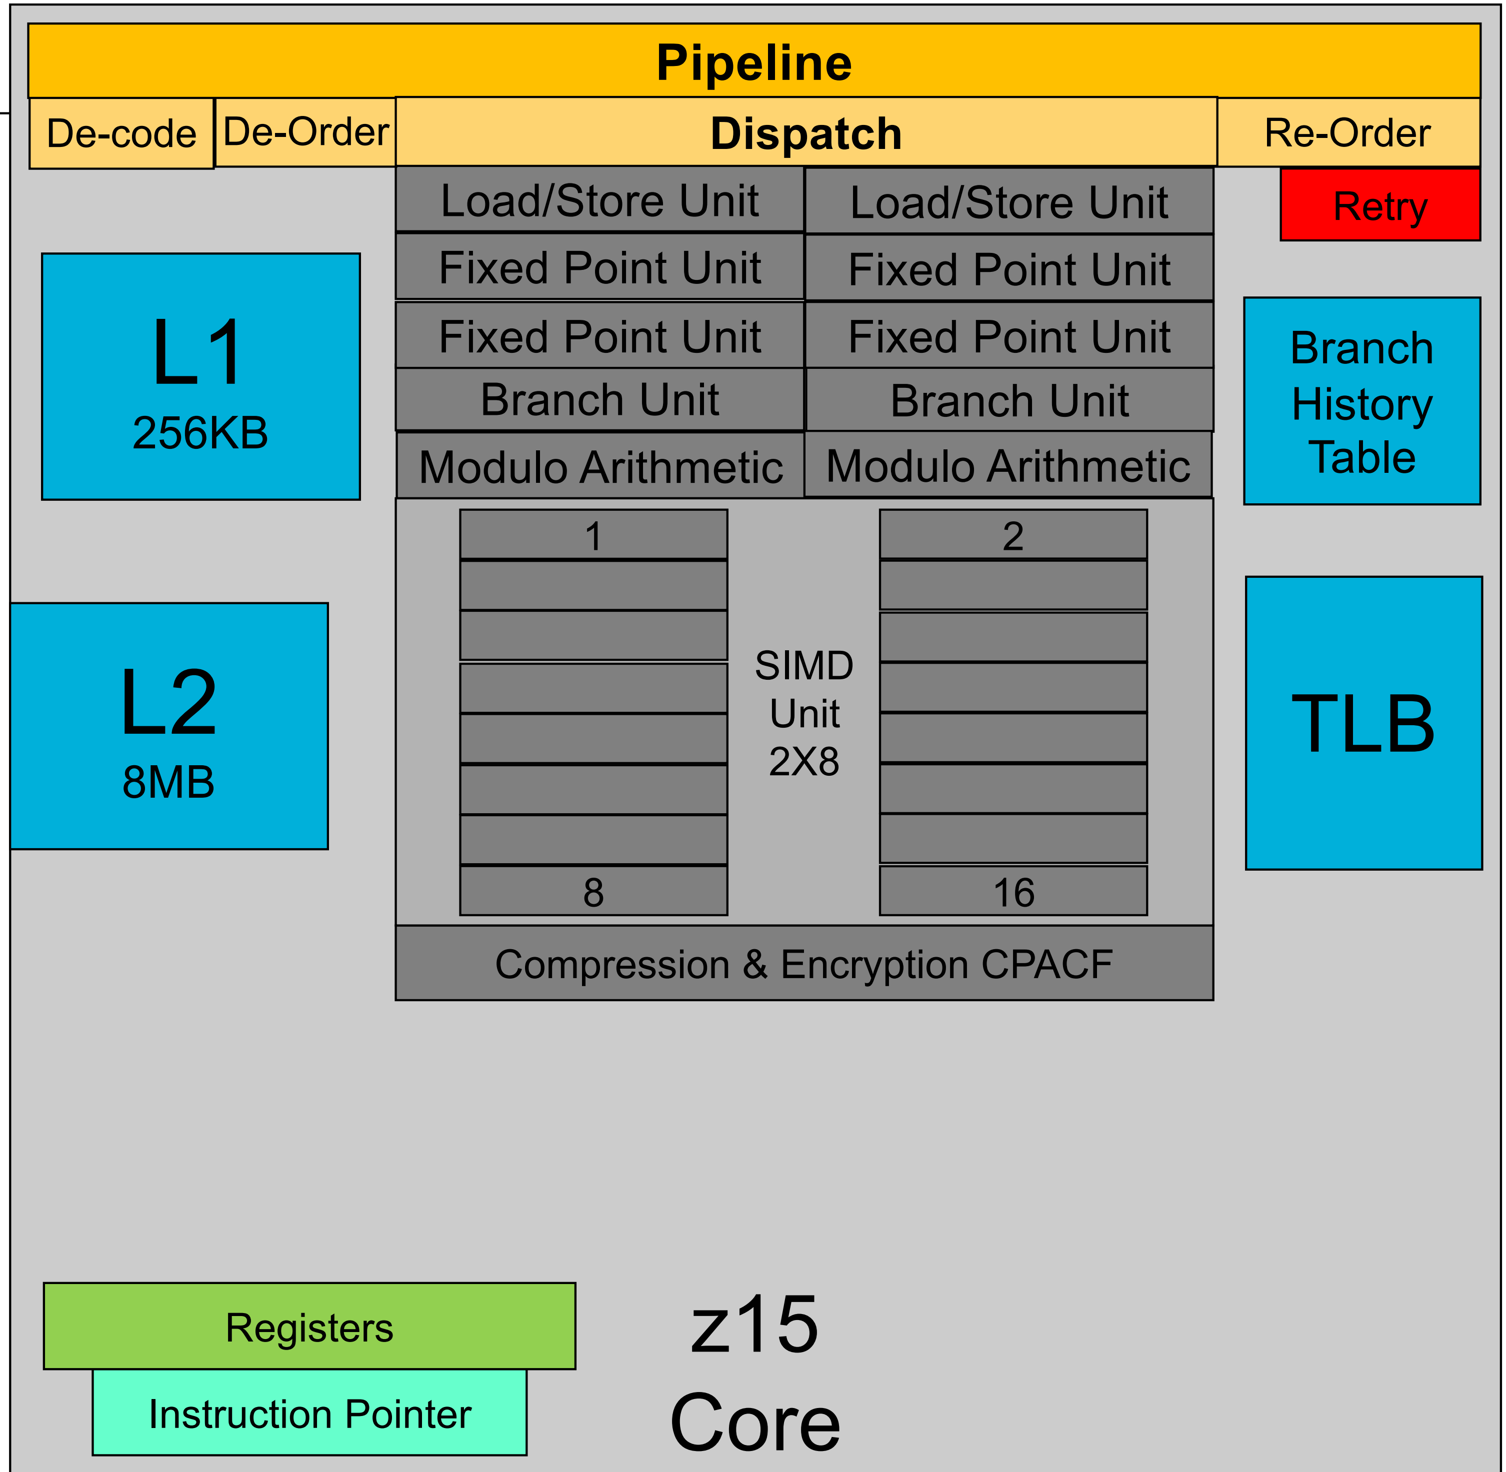
Pipeline
De-code
De-Order
Dispatch
Re-Order
Retry
Load/Store Unit
Load/Store Unit
Fixed Point Unit
Fixed Point Unit
Fixed Point Unit
Fixed Point Unit
Branch Unit
Branch Unit
Modulo Arithmetic
Modulo Arithmetic
L1
256KB
Branch History Table
L2
8MB
1
2
SIMD Unit 2X8
8
16
TLB
Compression & Encryption CPACF
Registers
Instruction Pointer
z15 Core

Topics

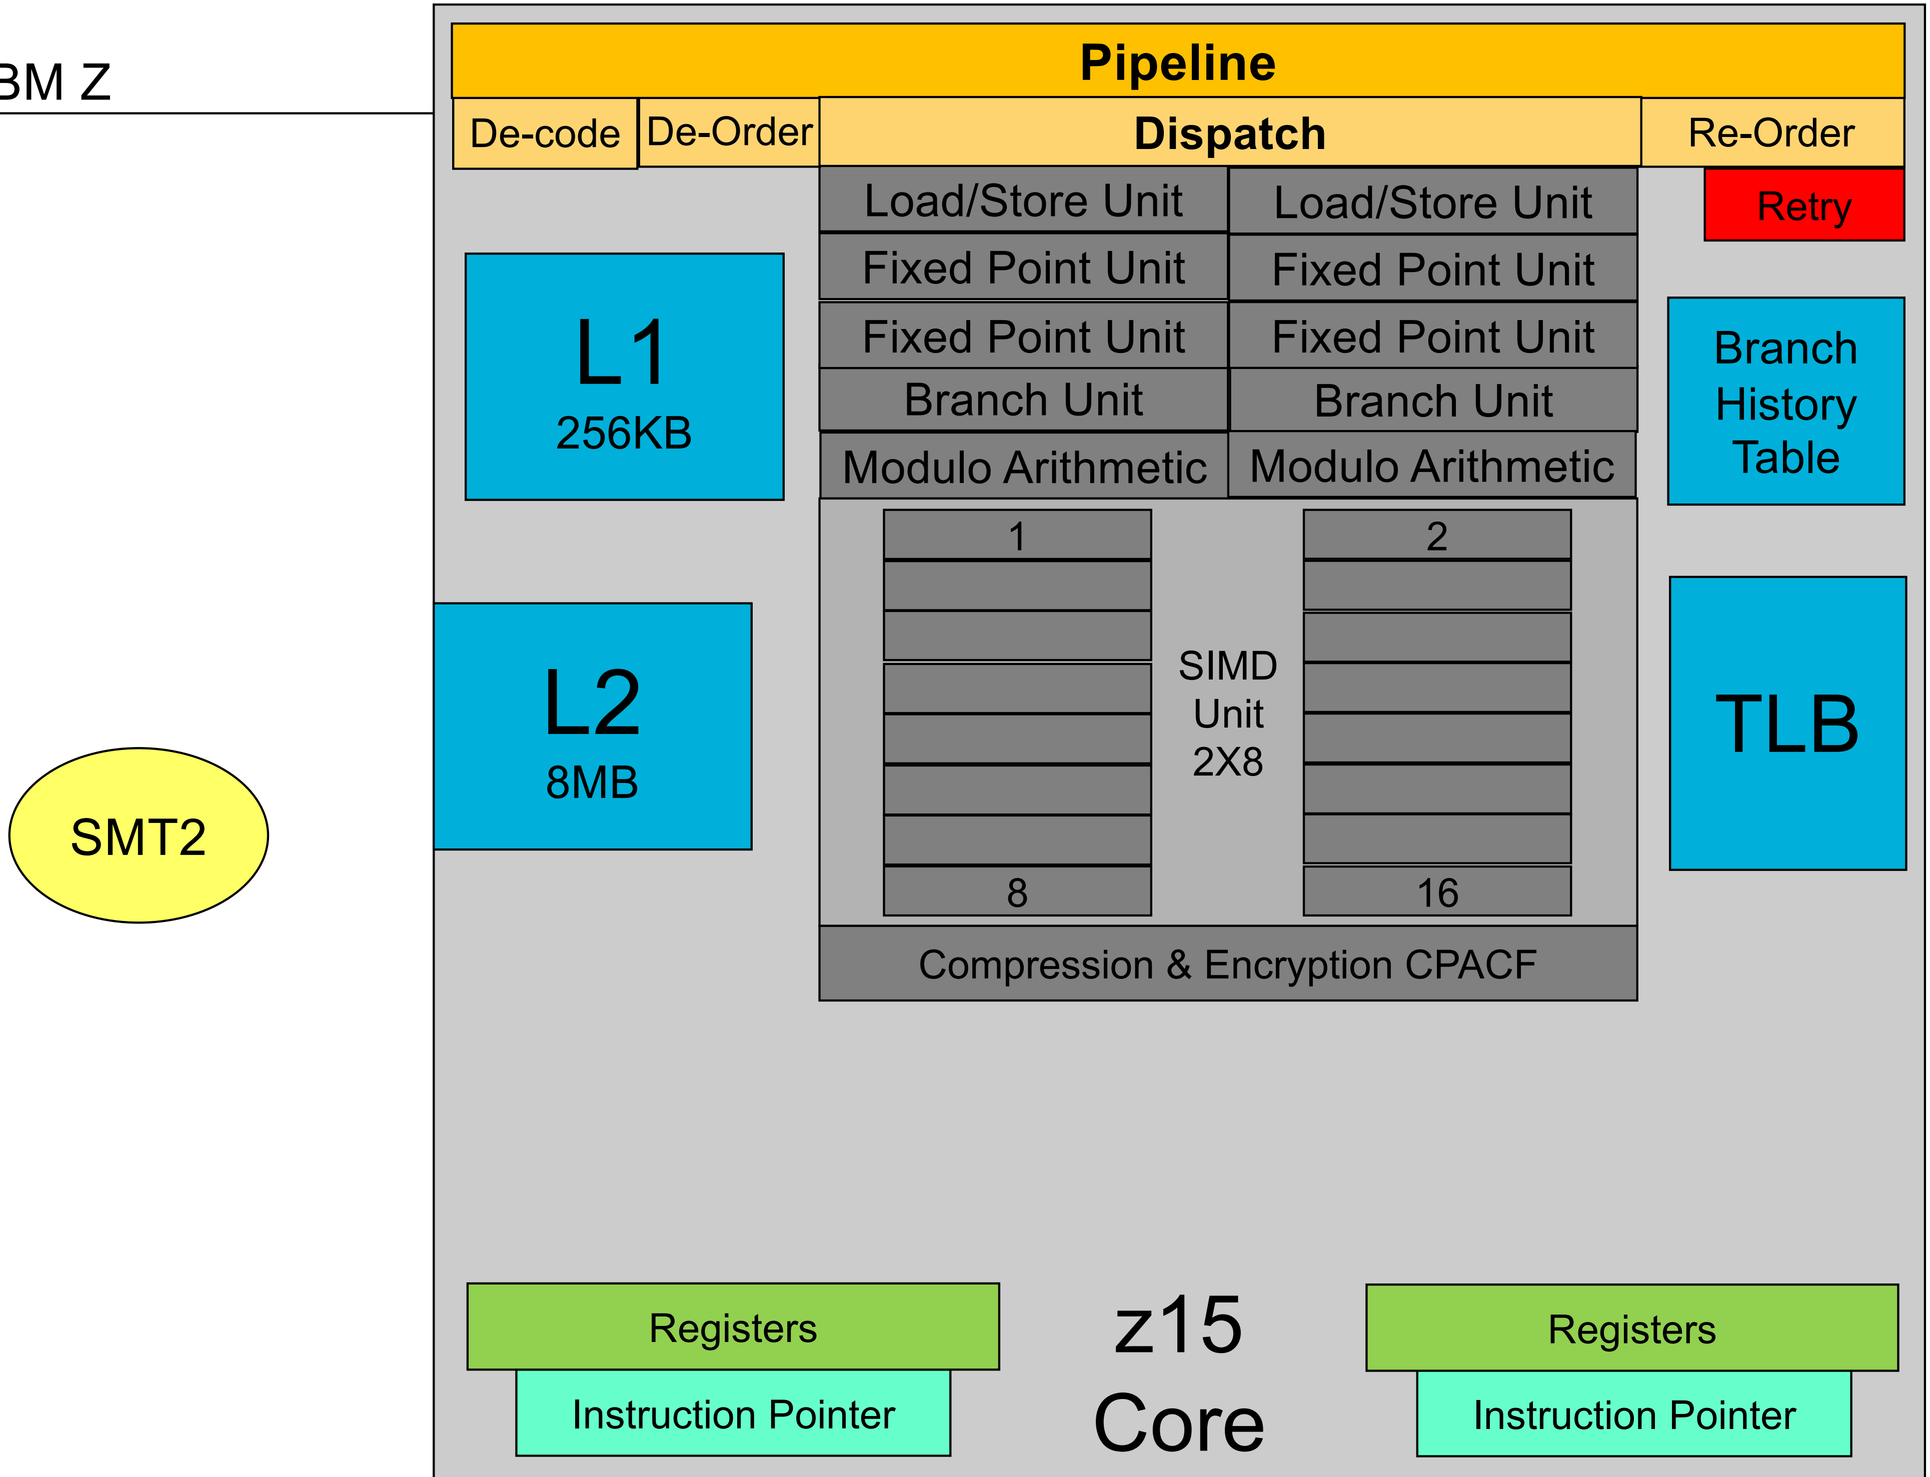
Pipeline
De-code
De-Order
Dispatch
Re-Order
Retry
Load/Store Unit
Load/Store Unit
Fixed Point Unit
Fixed Point Unit
Fixed Point Unit
Fixed Point Unit
Branch Unit
Branch Unit
Modulo Arithmetic
Modulo Arithmetic
L1
256KB
Branch History Table
L2
8MB
1
2
SIMD Unit 2X8
8
16
TLB
SMT2
Compression & Encryption CPACF
Registers
Instruction Pointer
z15 Core
Registers
Instruction Pointer

Topics

New Instructions:

57% reduction in CPU for COBOL 4.2 aps using ABO on z15

58% reduction in CPU moving from COBOL 4.2 to COBOL 6.3

20% more throughput for general Java workloads

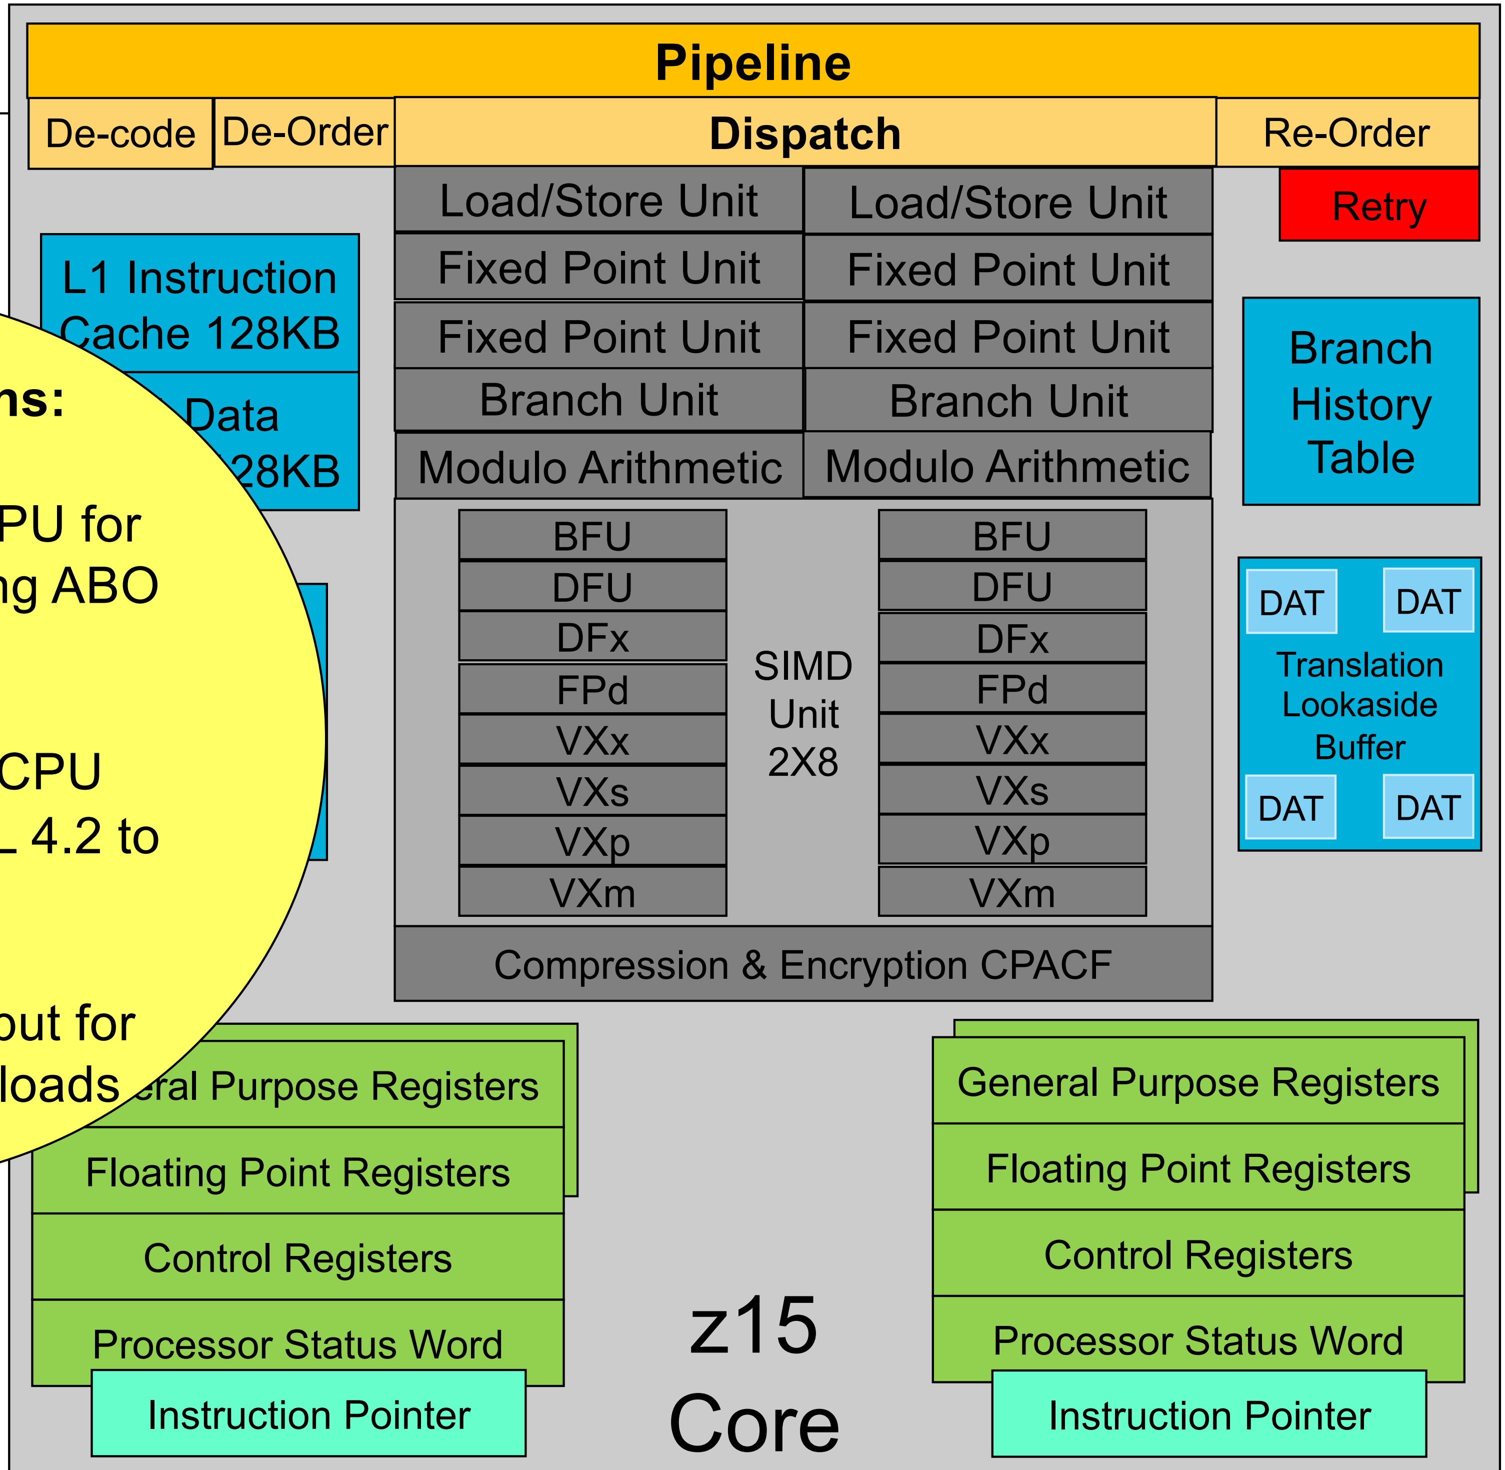

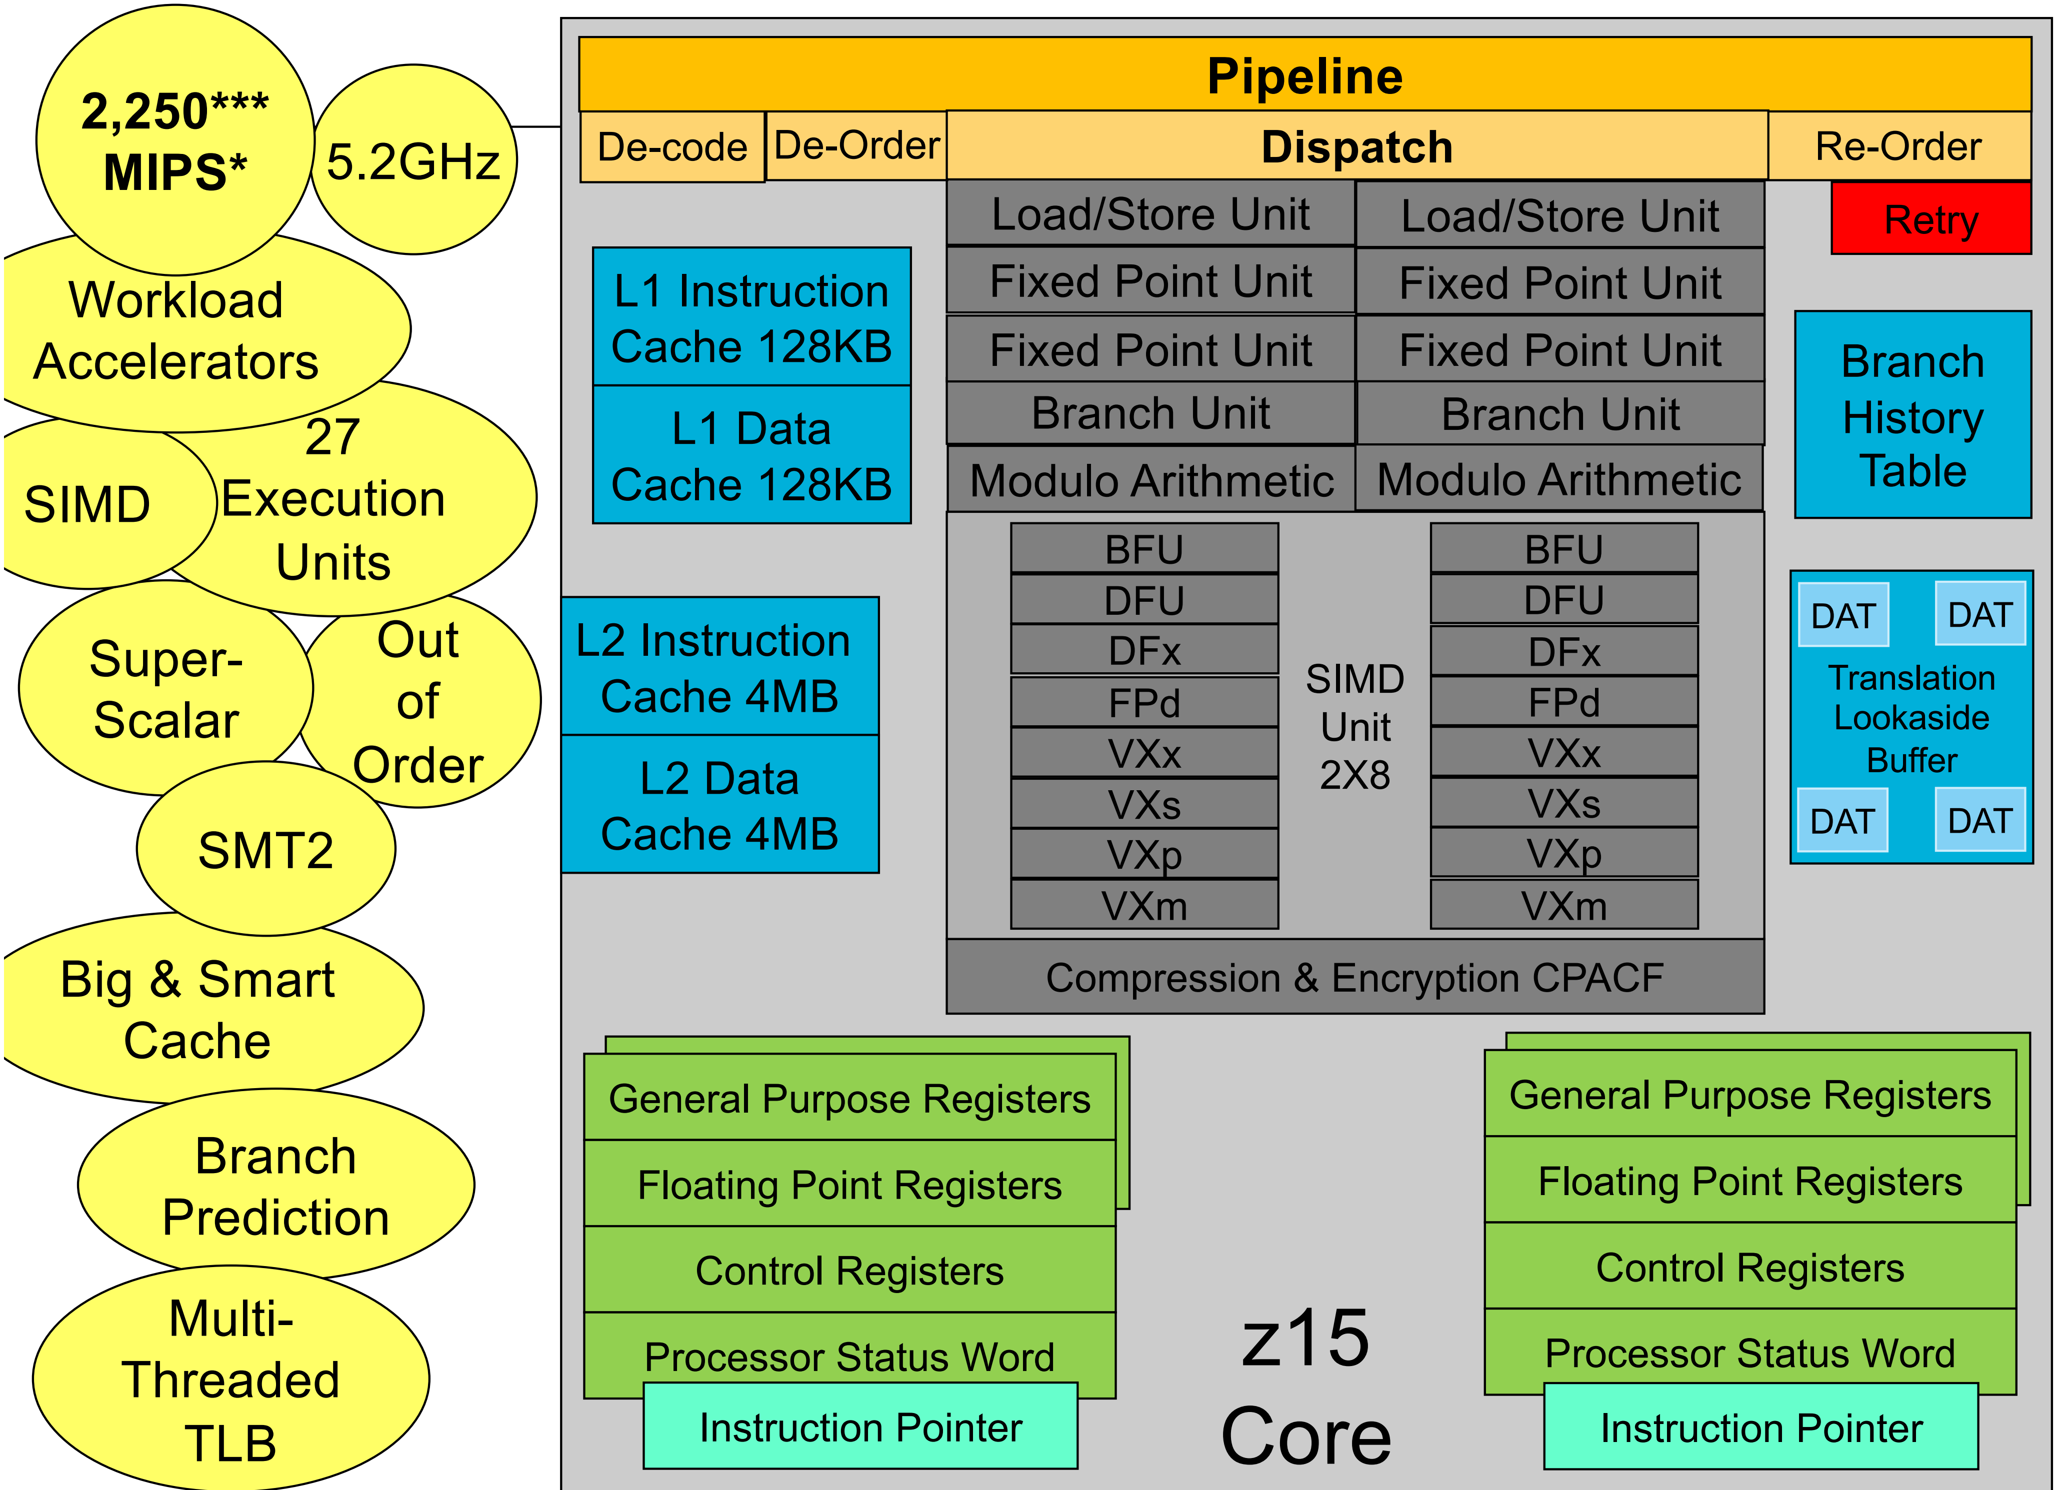
2,250***
MIPS*
5.2GHz
Workload Accelerators
SIMD
27 Execution Units
Super-Scalar
Out of Order
SMT2
Big & Smart Cache
Branch Prediction
Multi-Threaded TLB
Pipeline
De-code
De-Order
Dispatch
Re-Order
Retry
Load/Store Unit
Load/Store Unit
Fixed Point Unit
Fixed Point Unit
Fixed Point Unit
Fixed Point Unit
Branch Unit
Branch Unit
Modulo Arithmetic
Modulo Arithmetic
L1 Instruction Cache 128KB
L1 Data Cache 128KB
L2 Instruction Cache 4MB
L2 Data Cache 4MB
Branch History Table
BFU
DFU
DFx
FPd
VXx
VXs
VXp
VXm
SIMD Unit 2X8
BFU
DFU
DFx
FPd
VXx
VXs
VXp
VXm
DAT
DAT
Translation Lookaside Buffer
DAT
DAT
Compression & Encryption CPACF
General Purpose Registers
Floating Point Registers
Control Registers
Processor Status Word
Instruction Pointer
z15 Core
General Purpose Registers
Floating Point Registers
Control Registers
Processor Status Word
Instruction Pointer

Cores to Chips

9 Billion Things to Like…

THE Z15 CHIP

Topics

9.1B
Transistors

z15
Chip

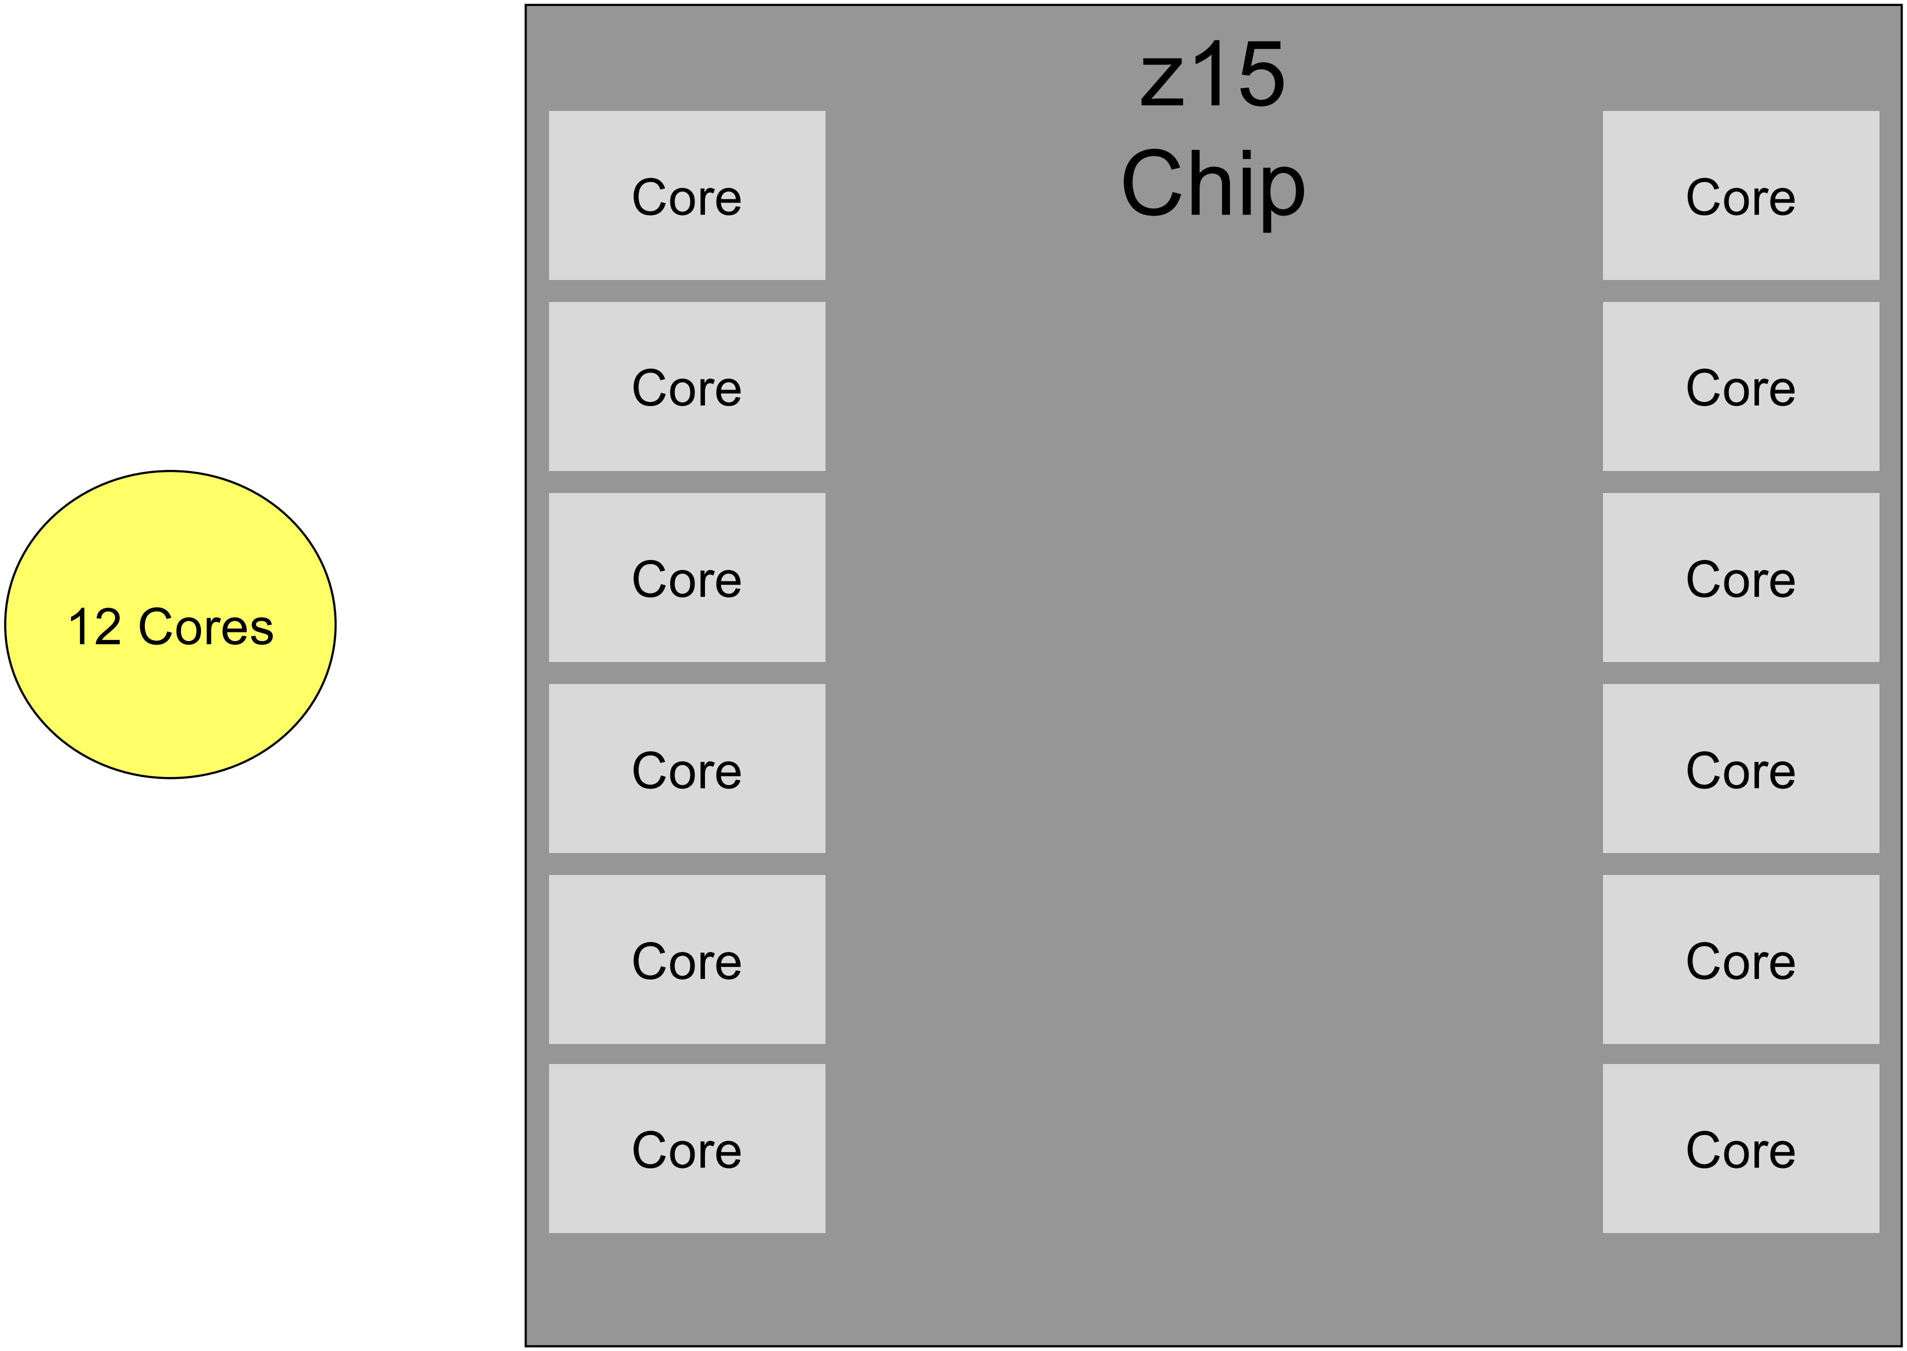
12 Cores
z15
Chip
Core
Core
Core
Core
Core
Core
Core
Core
Core
Core
Core
Core

L1/L2/L3 Total 259MB cache

More (much more) big & smart cache

ECC-protected cache arrays

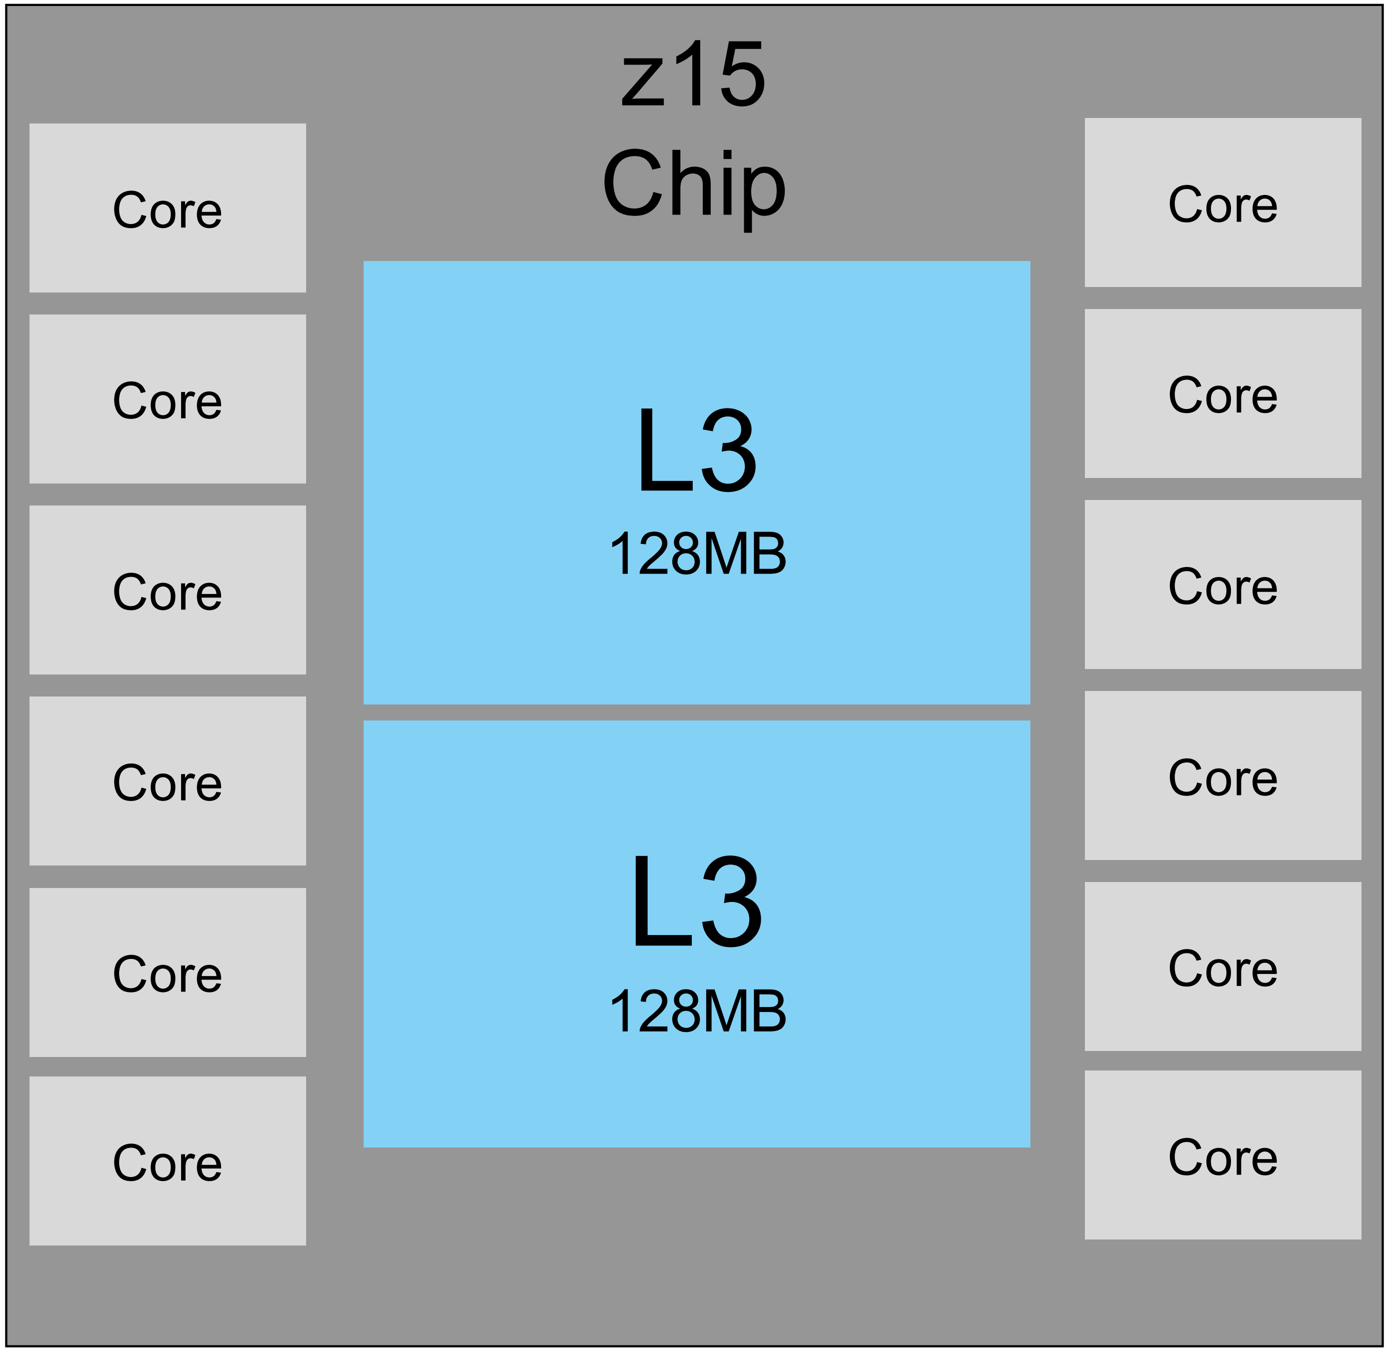

Topics

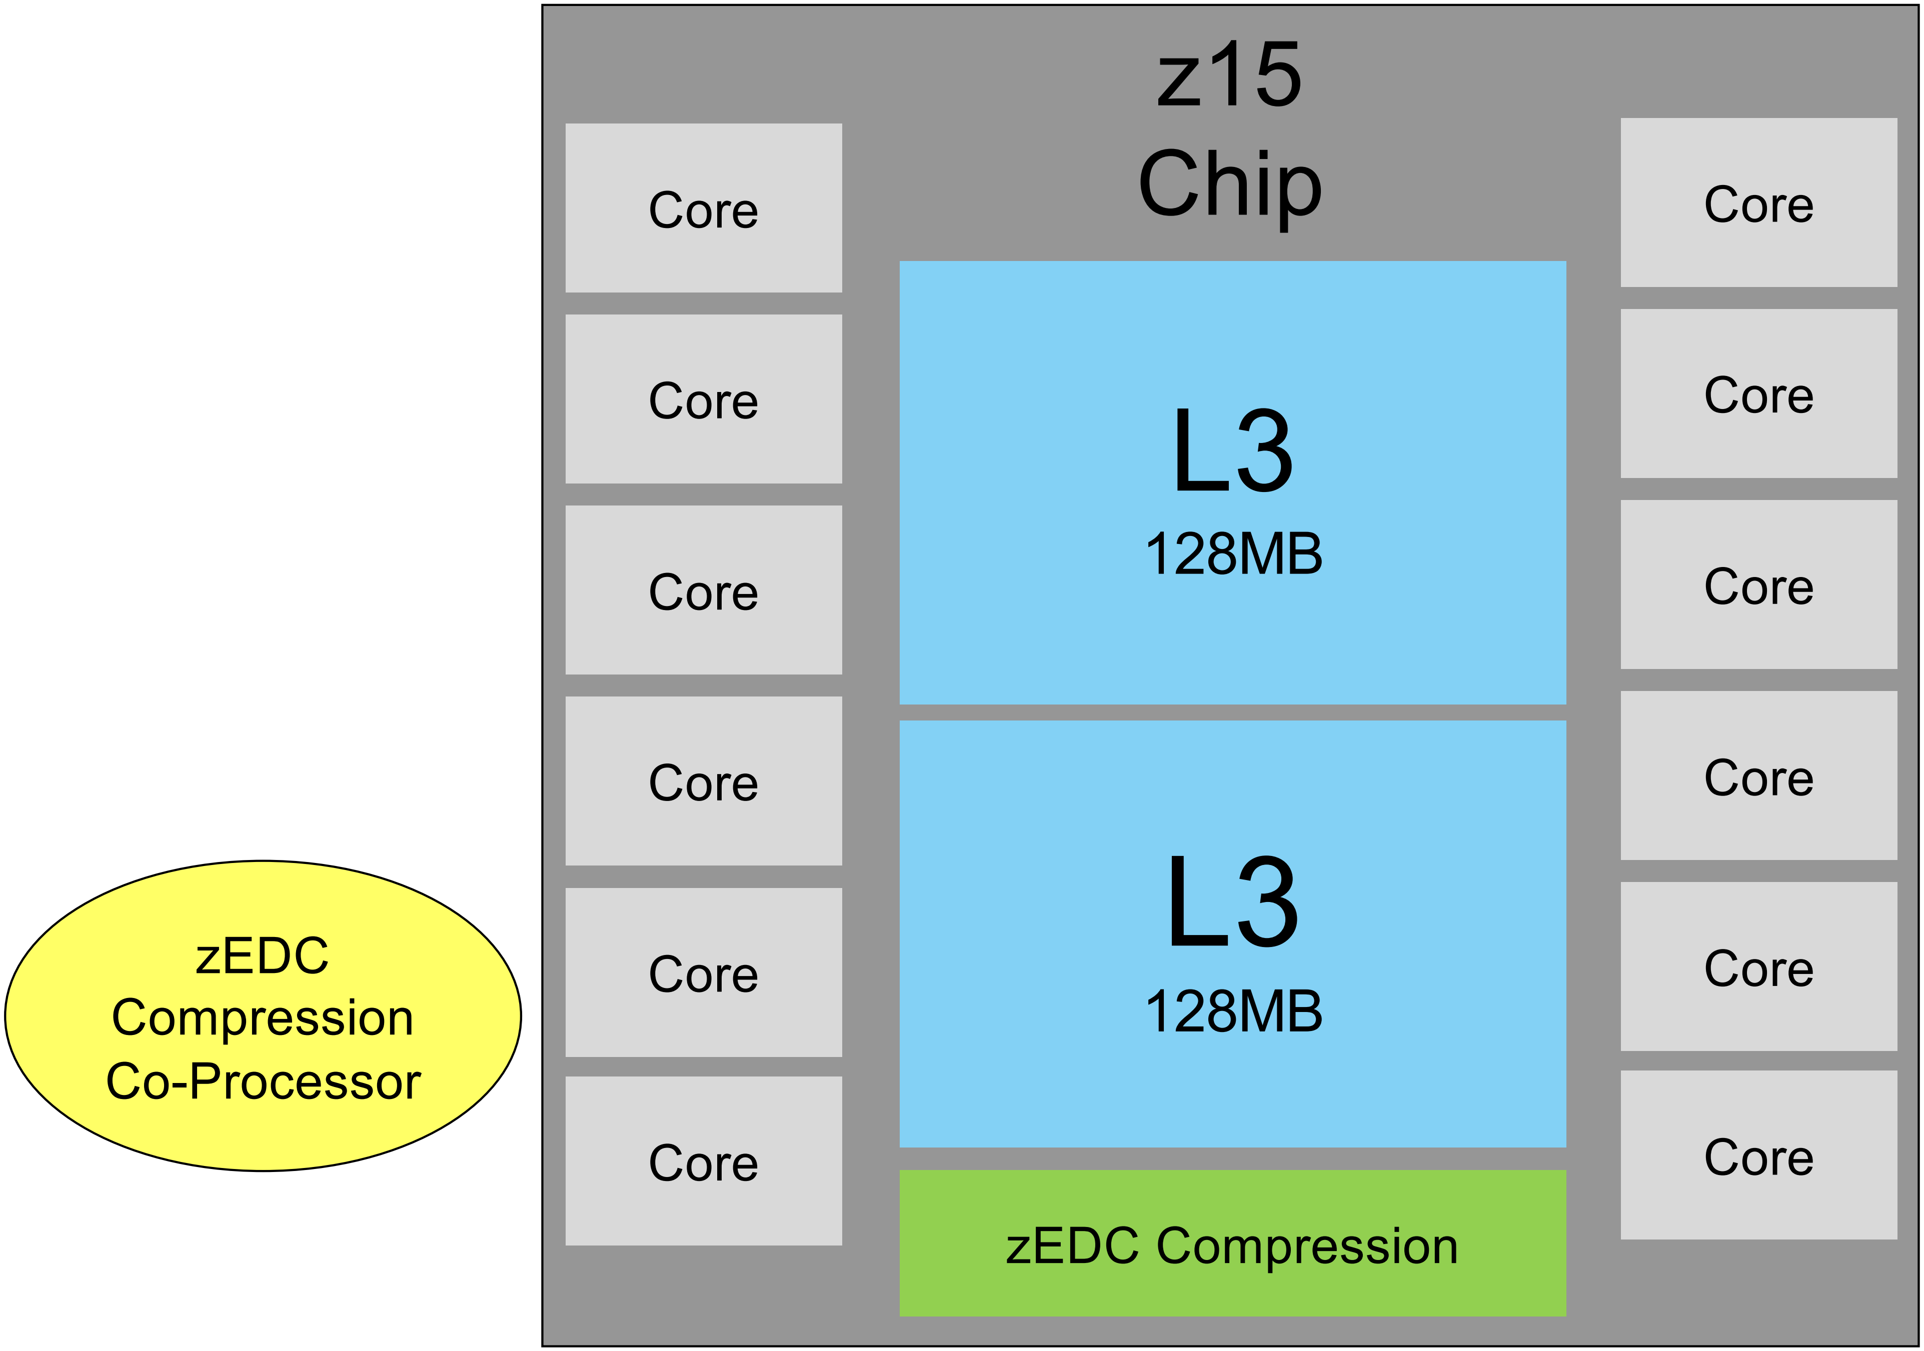
z15
Chip
Core
Core
Core
Core
Core
Core
Core
Core
Core
Core
Core
Core
L3
128MB
L3
128MB
zEDC Compression
zEDC
Compression
Co-Processor

Topics

9.1B Transistors

12 Cores

L1/L2/L3 Total 259MB cache

More (much more) big & smart cache

ECC-protected cache arrays

zEDC Compression Co-Processor

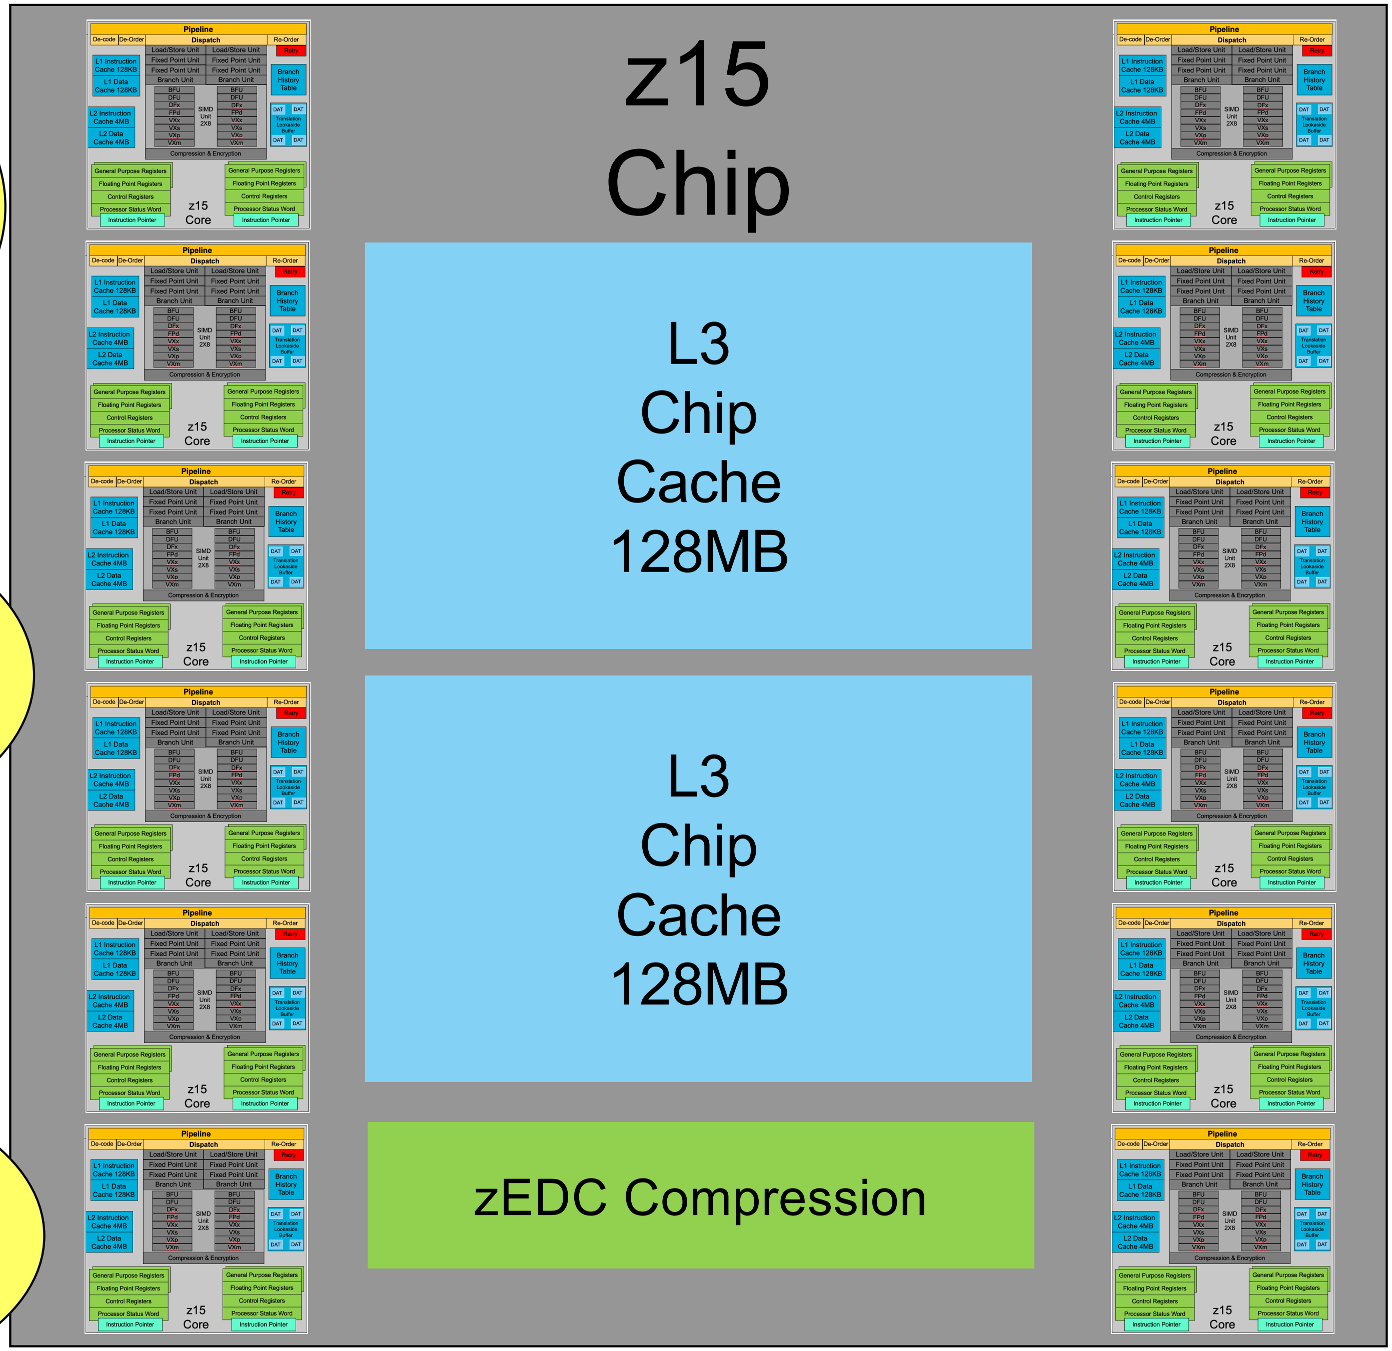

Topics

Chips to Drawers

THE Z15 DRAWER

Topics

Mainframe Motherboard

z15 Drawer

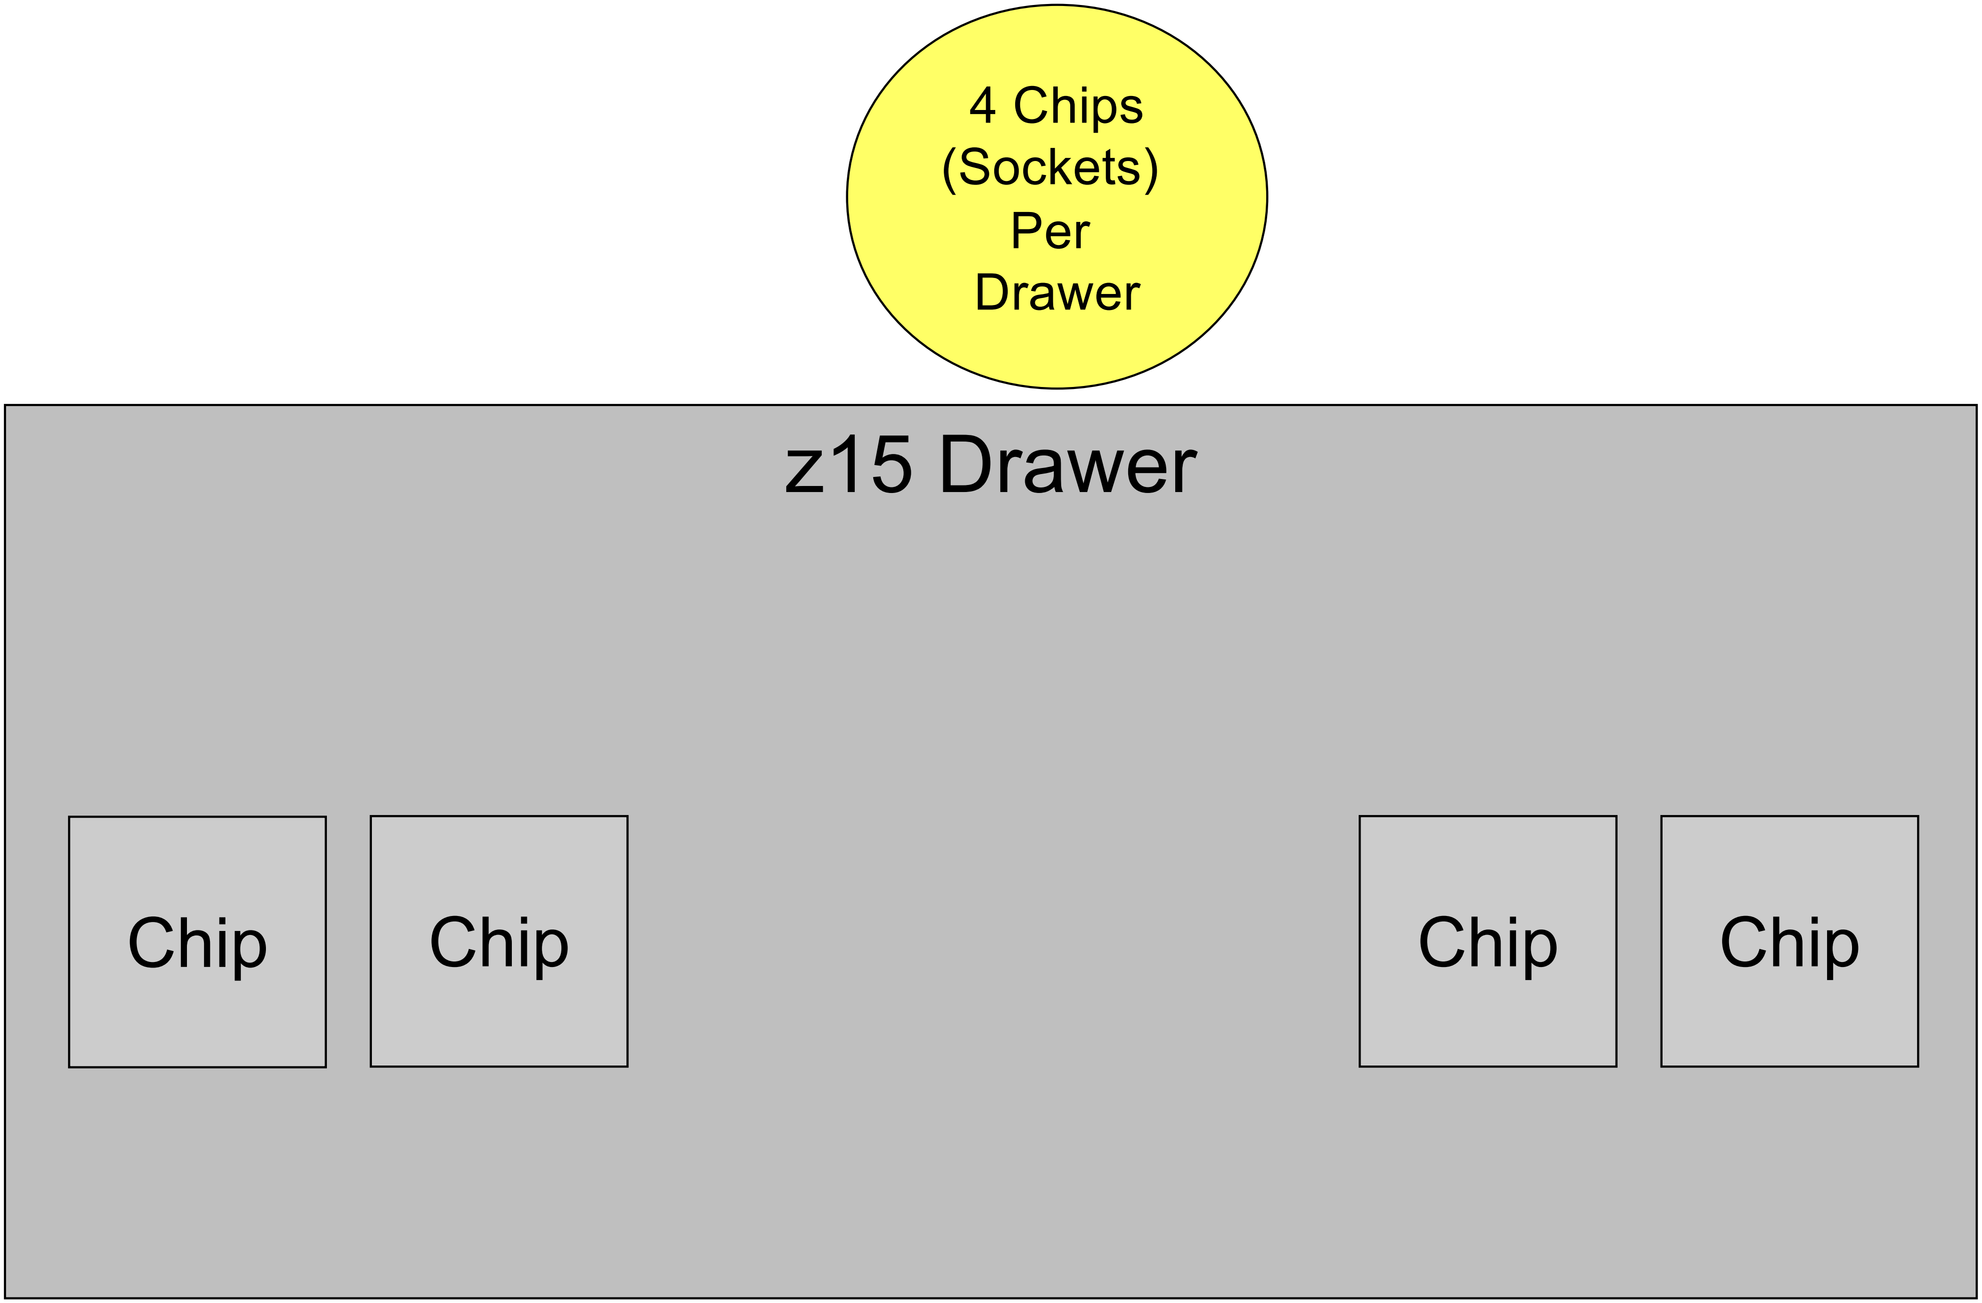
4 Chips
(Sockets)
Per
Drawer
z15 Drawer
Chip
Chip
Chip
Chip

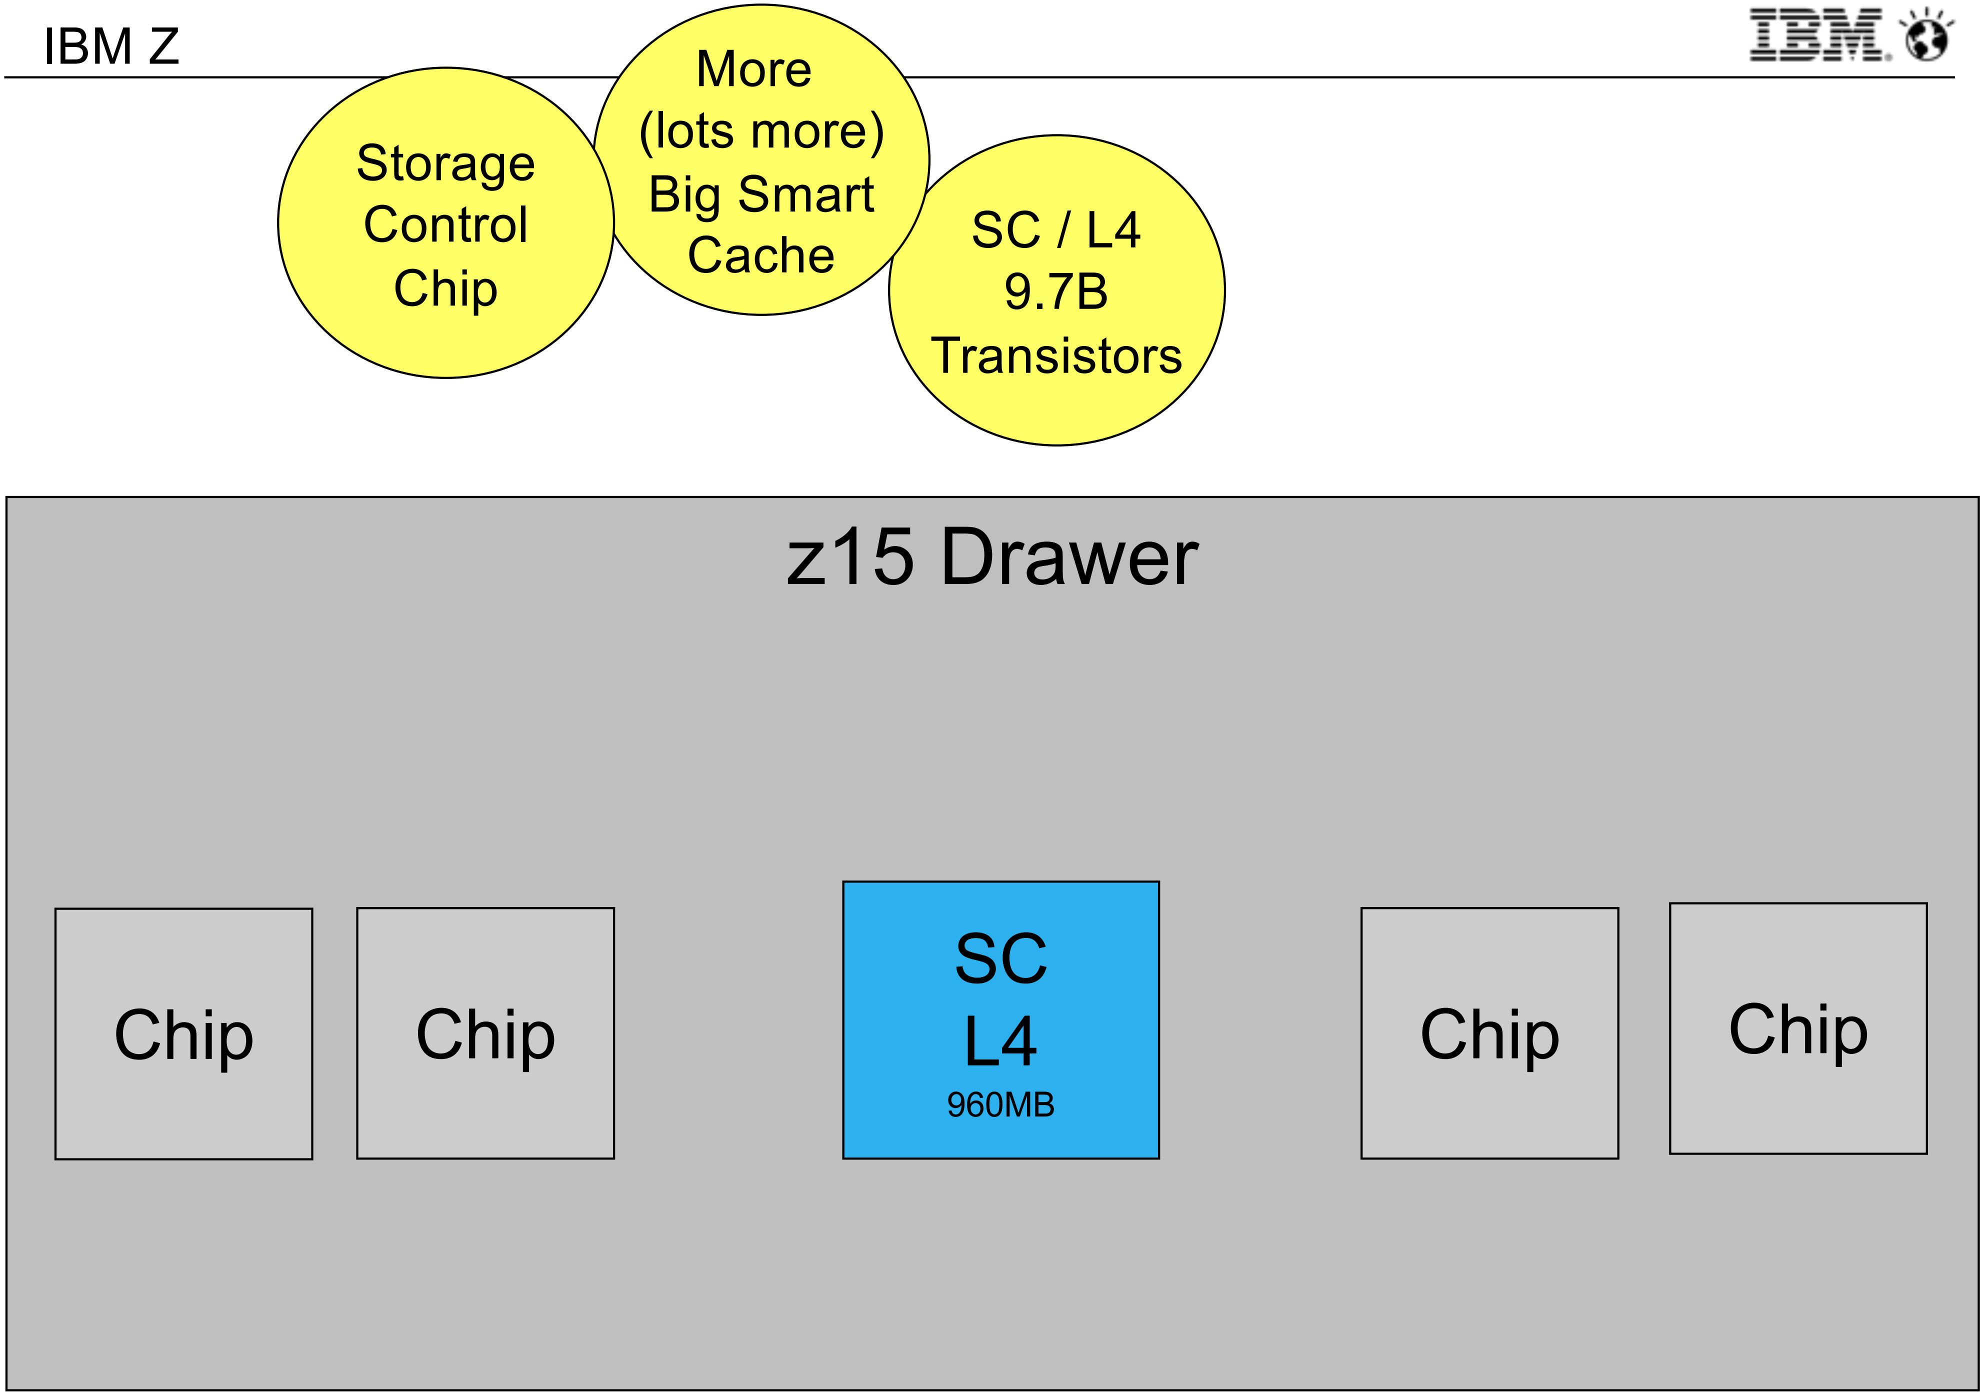
Storage Control Chip
More (lots more) Big Smart Cache
SC / L4 9.7B Transistors
z15 Drawer
Chip
Chip
SC L4 960MB
Chip
Chip

Up to 8TB Useable RAIM RAM

The 5th DIMM is "RAIM" backup for the other 4

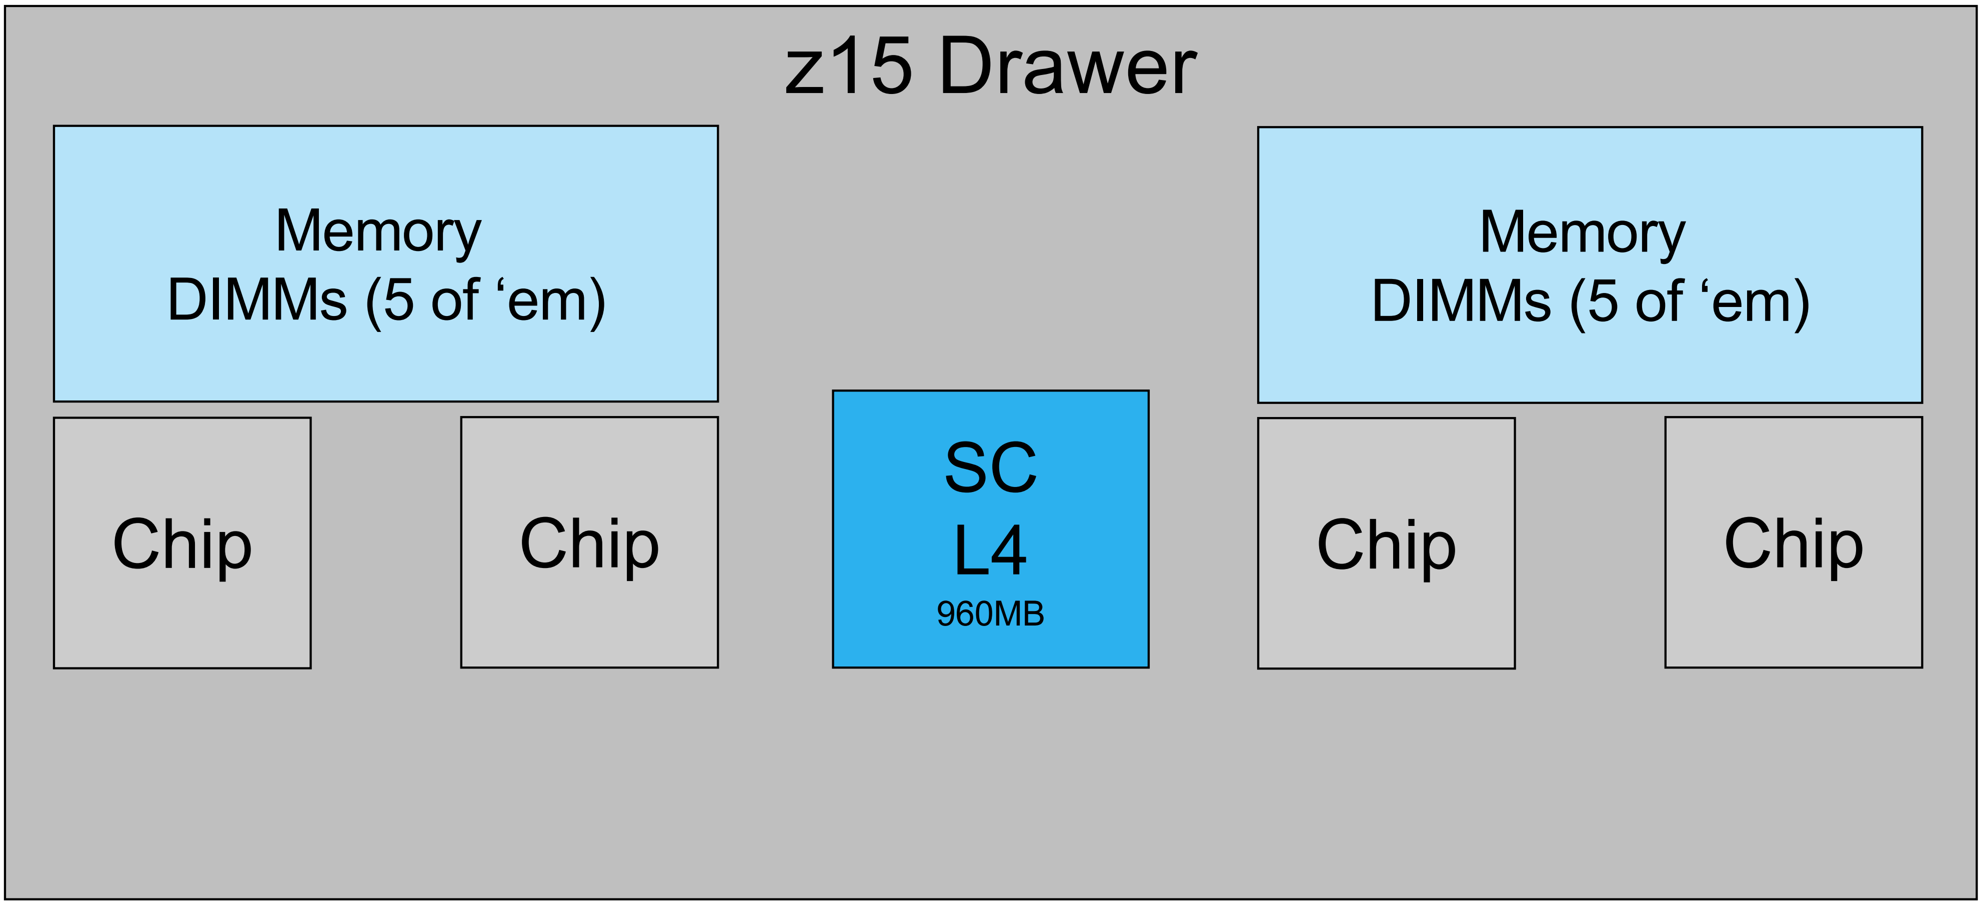

Topics

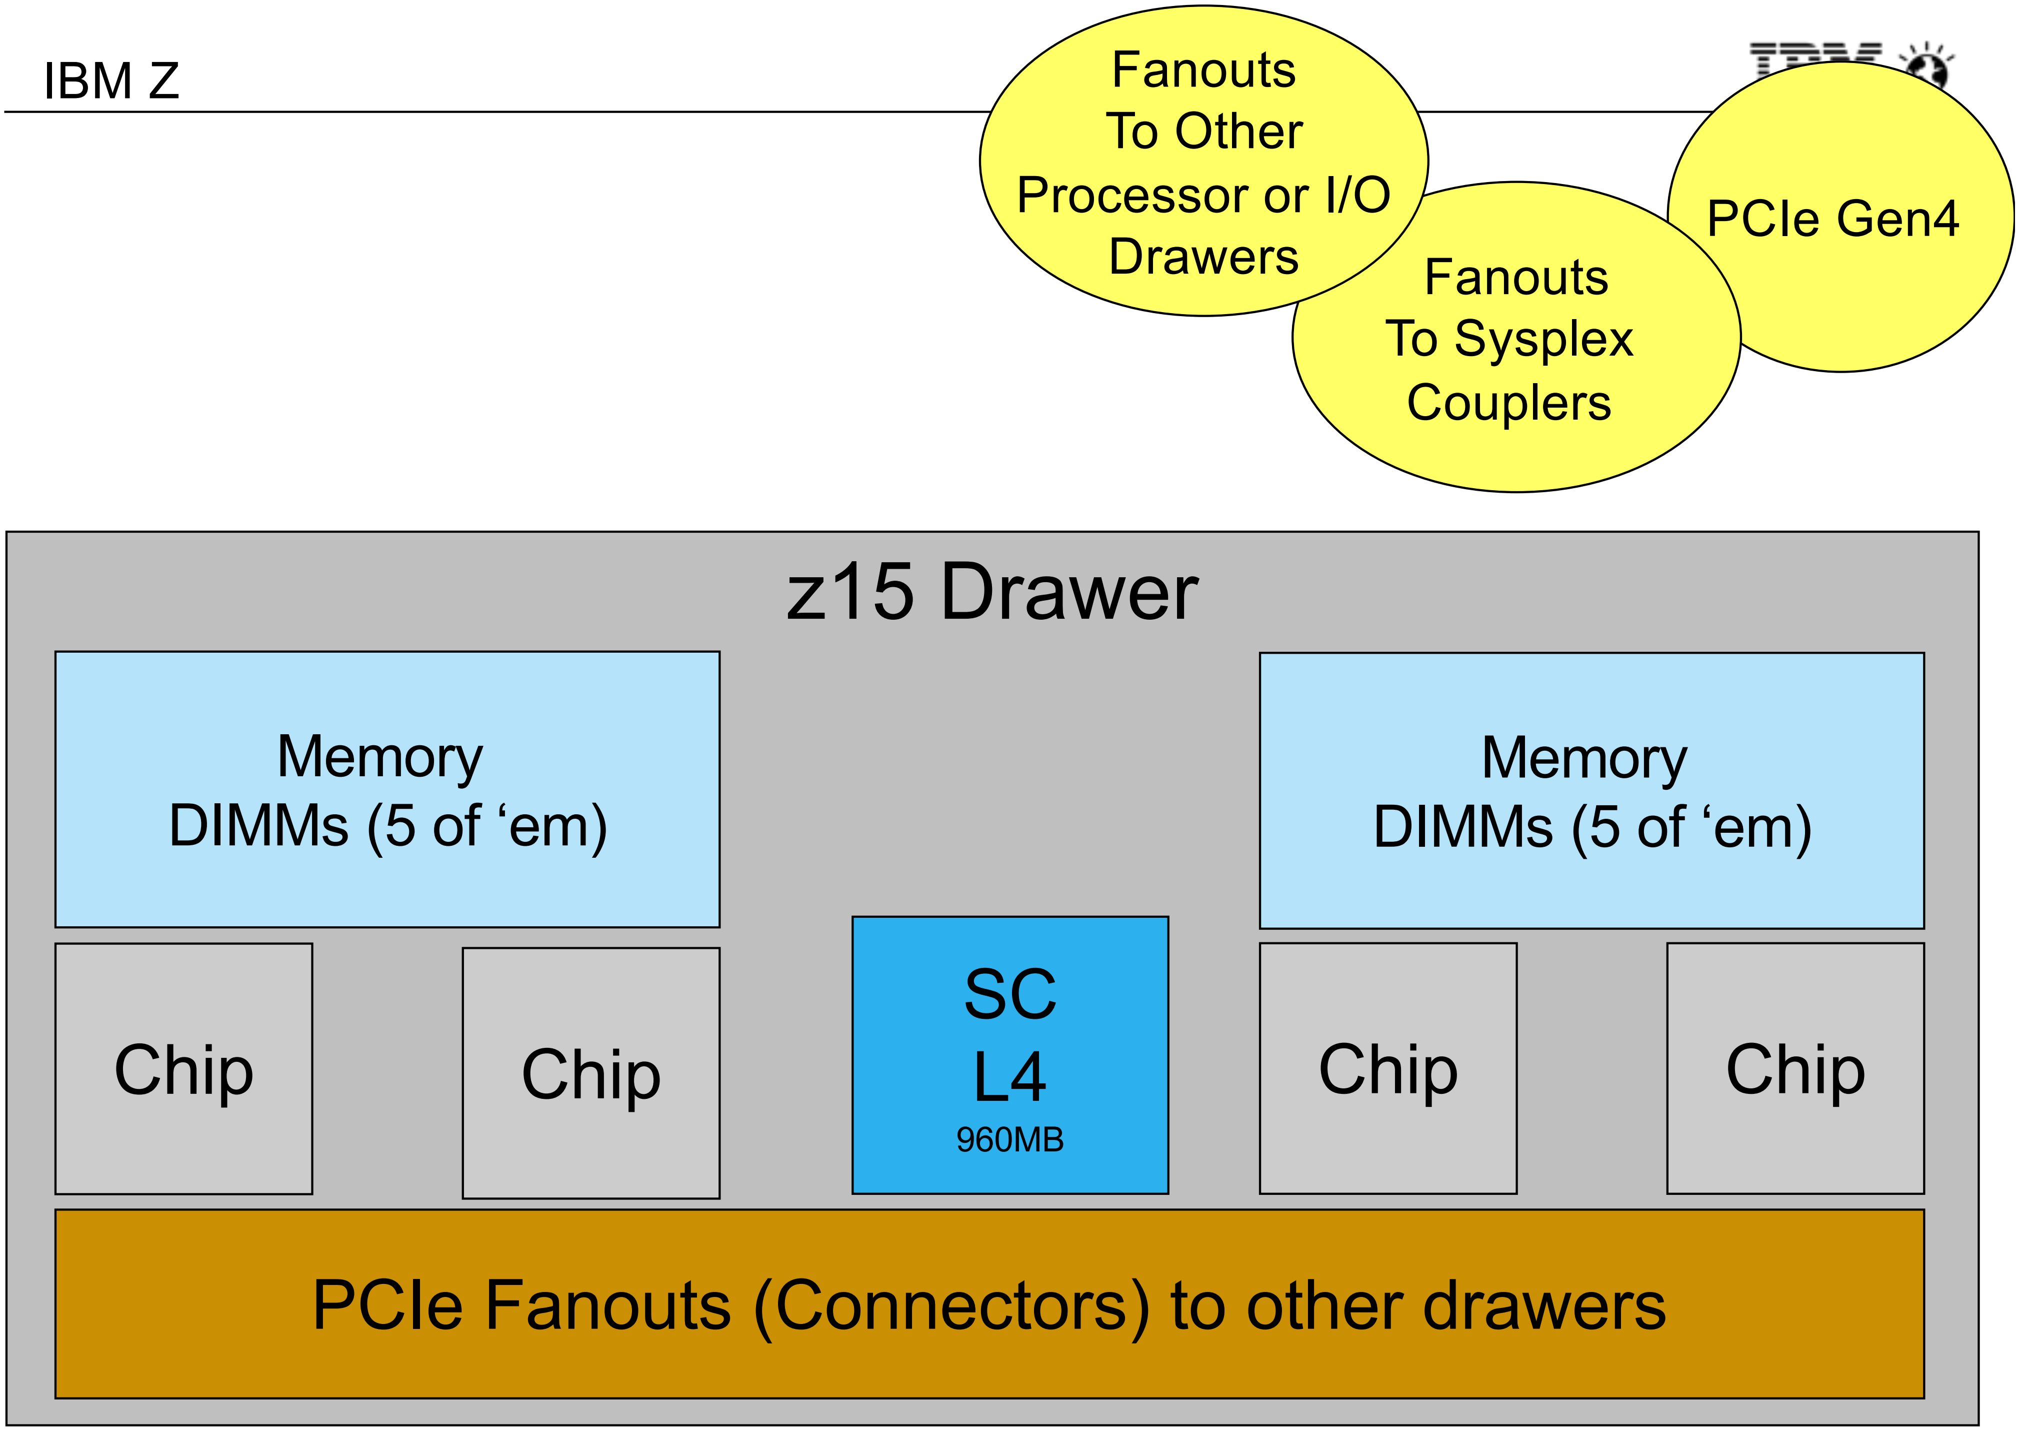
Fanouts
To Other
Processor or I/O
Drawers
Fanouts
To Sysplex
Couplers
PCIe Gen4
z15 Drawer
Memory
DIMMs (5 of ‘em)
Memory
DIMMs (5 of ‘em)
Chip
Chip
SC
L4
960MB
Chip
Chip
PCIe Fanouts (Connectors) to other drawers

z15 Drawer

Mainframe Motherboard

SC / L4 9.7B Transistors

More (lots more) Big Smart Cache

4 Chips Per Drawer

Up to 8TB Useable RAIM RAM

Fanouts To Other Processor and I/O Drawers

Fanouts To Sysplex Couplers

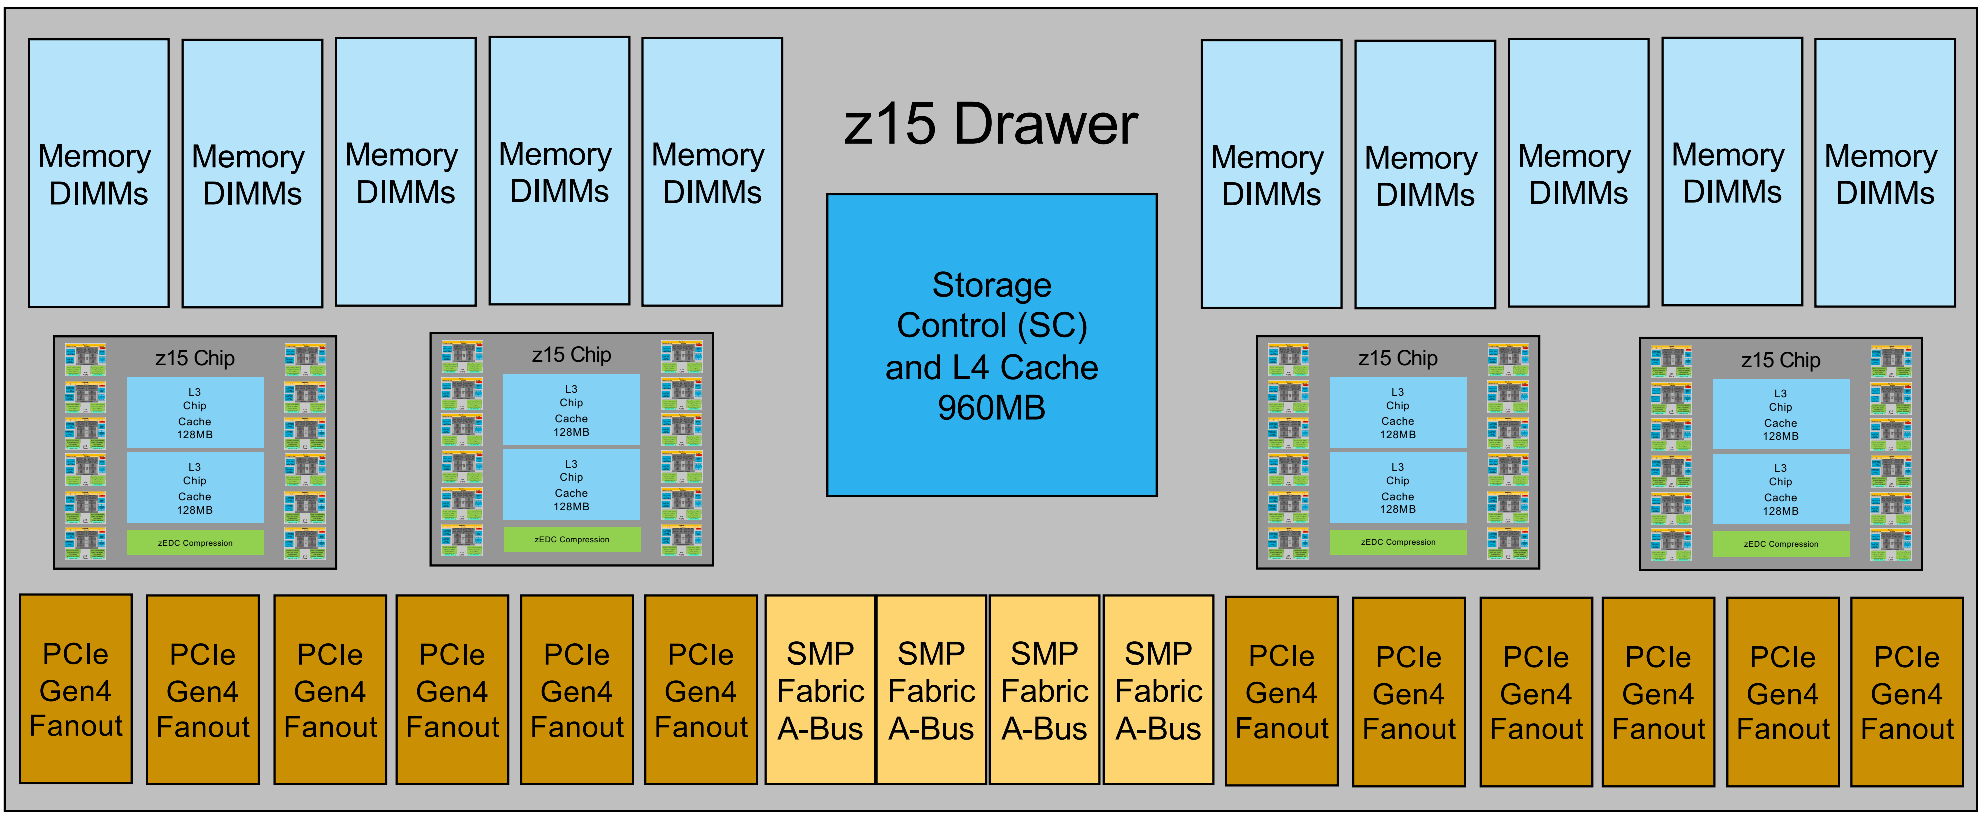

Topics

Drawers to CEC

CEC … not Keg

Central Electronics Complex

Processing Nest

THE Z15 CEC
(THE LINUXONE III LT2 CEC)

Topics

T01/LT1 Model	Drawers	Cores	Memory (RAIM)
Max34	1	0 to 34	Up to 8TB

Processor Drawer

Start at 250 MIPS

Sub-Uni's:
400 series
500 series
600 series
(up to 34-way)

Start at 256MB RAIM

Topics

T01/LT1 Model	Drawers	Cores	Memory (RAIM)
Max34	1	0 to 34	Up to 8TB
Max71	2	0 to 71	Up to 16TB

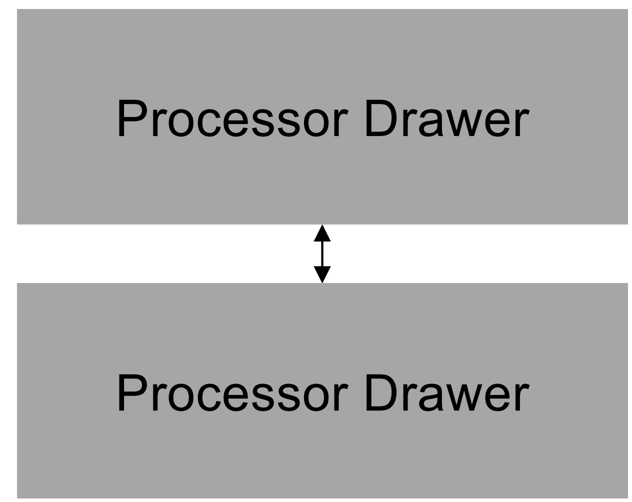

Fully Interconnected SMP Fabric

T01/LT1 Model	Drawers	Cores	Memory (RAIM)
Max34	1	0 to 34	Up to 8TB
Max71	2	0 to 71	Up to 16TB
Max108	3	0 to 108	Up to 24TB

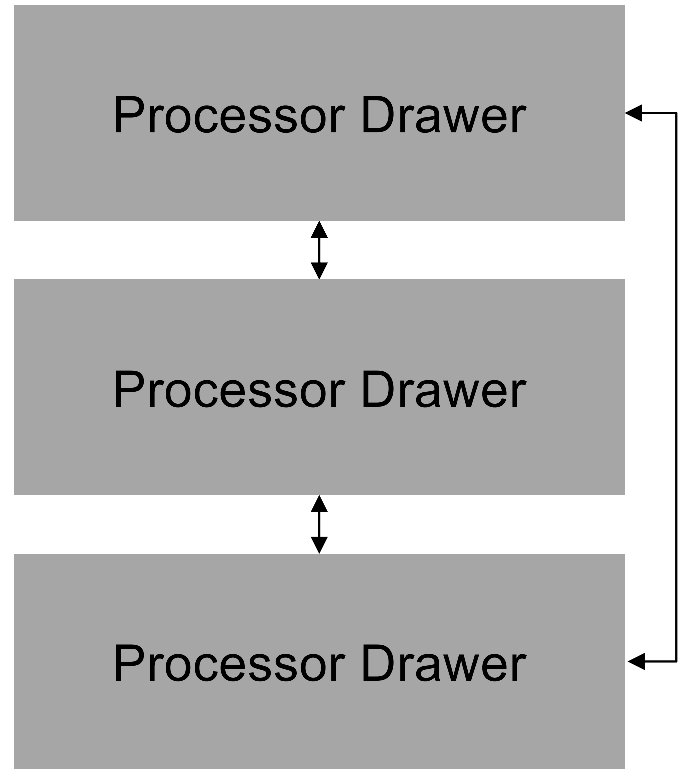

T01/LT1 Model	Drawers	Cores	Memory (RAIM)
Max34	1	0 to 34	Up to 8TB
Max71	2	0 to 71	Up to 16TB
Max108	3	0 to 108	Up to 24TB
Max145	4	0 to 145	Up to 32TB

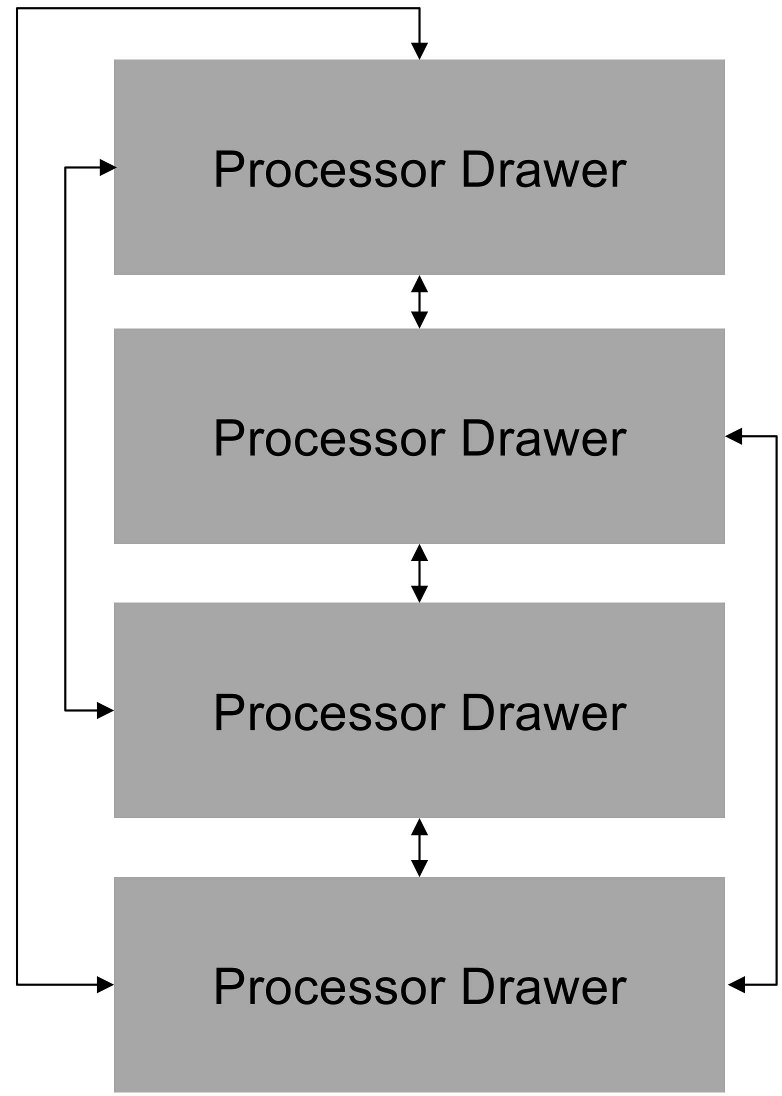

T01/LT1 Model	Drawers	Cores	Memory (RAIM)
Max34	1	0 to 34	Up to 8TB
Max71	2	0 to 71	Up to 16TB
Max108	3	0 to 108	Up to 24TB
Max145	4	0 to 145	Up to 32TB
Max190	5	0 to 190	Up to 40TB

Scale to 200,000 MIPS

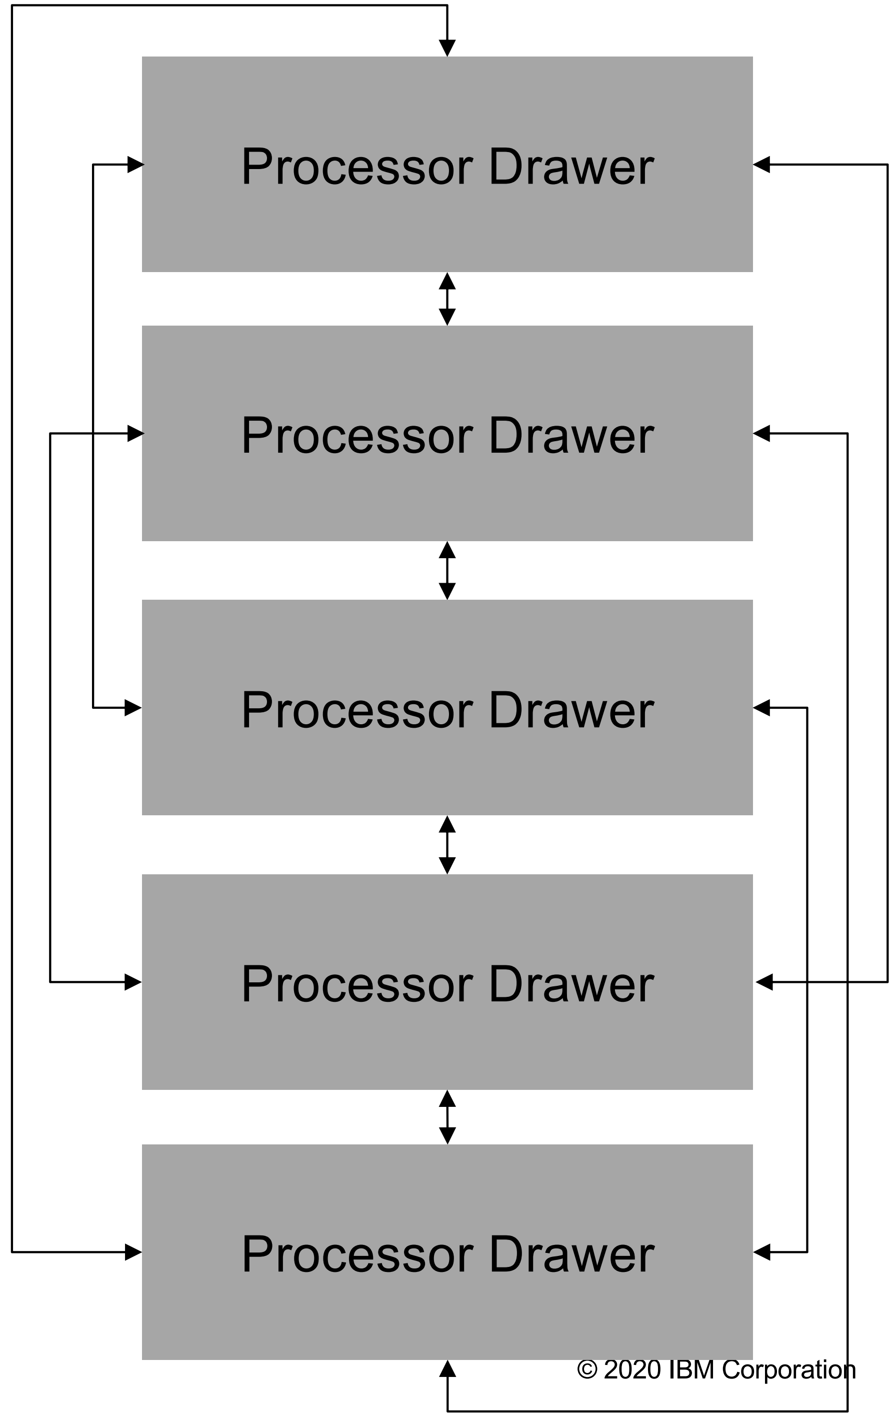

Topics

CEC: z15 T01 --- LinuxONE III LT1

T01/LT1 Model	Drawers	Cores	Memory (RAIM)
Max34	1	0 to 34	Up to 8TB
Max71	2	0 to 71	Up to 16TB
Max108	3	0 to 108	Up to 24TB
Max145	4	0 to 145	Up to 32TB
Max190	5	0 to 190	Up to 40TB

Start at 250 MIPS

0 to 190 Cores

256MB to 40TB RAIM

Scale to 200,000 MIPS

Sub-Uni's:
400 series
500 series
600 series
(up to 34-way)

Fully Interconnected SMP Fabric

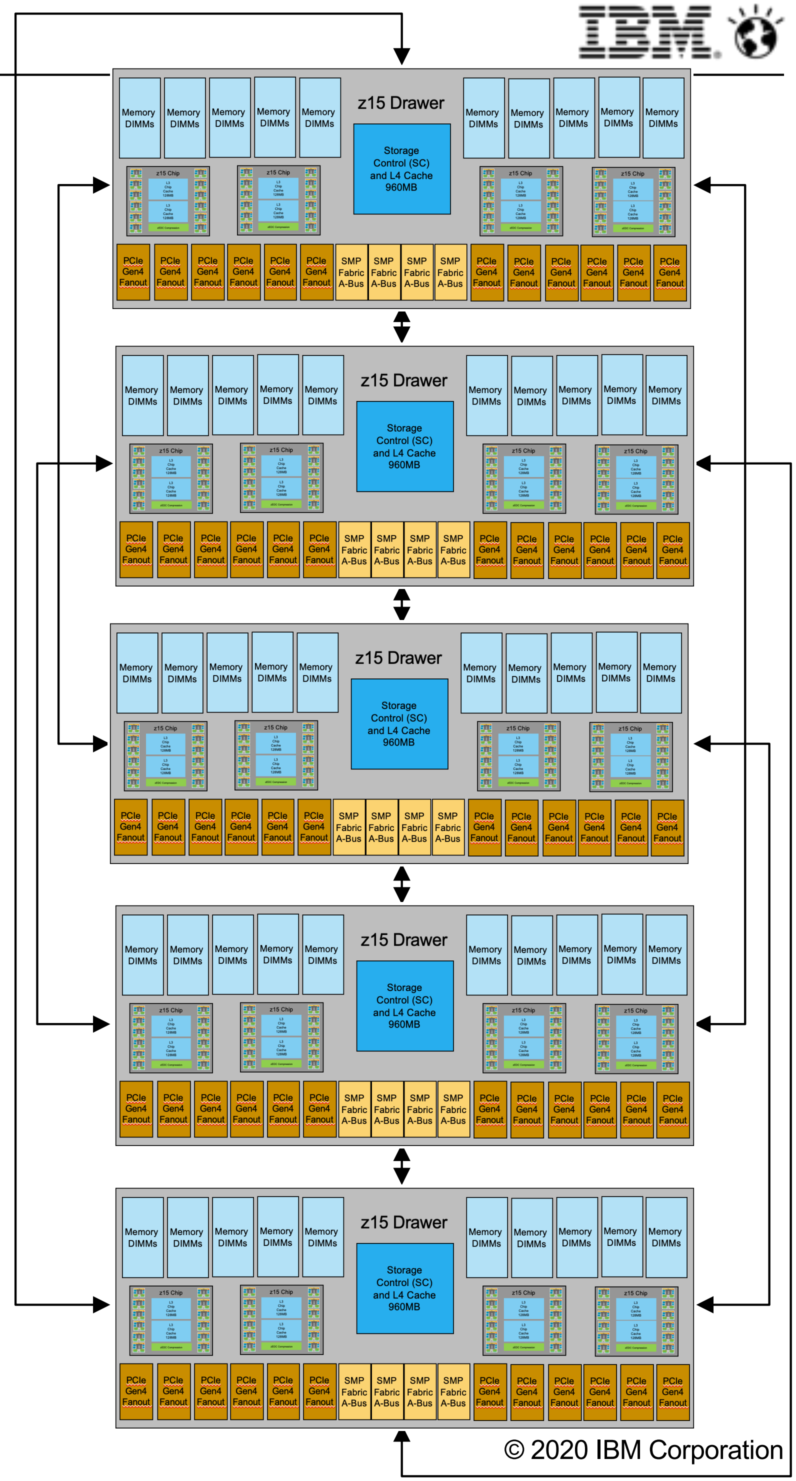

CEC: z15 T02 ------- LinuxONE III LT2

T02/LT2 Model	Drawers	Cores	Memory (RAIM)
Max04	1	0 to 4	Up to 2TB
Max13	1	0 to 13	Up to 4TB
Max21	1	0 to 21	Up to 4TB
Max31	1	0 to 31	Up to 8TB
Max65	2	0 to 65	Up to 16TB

Max04 Max13 Max21 Max31

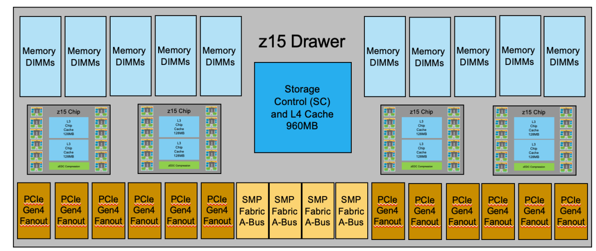

Max65

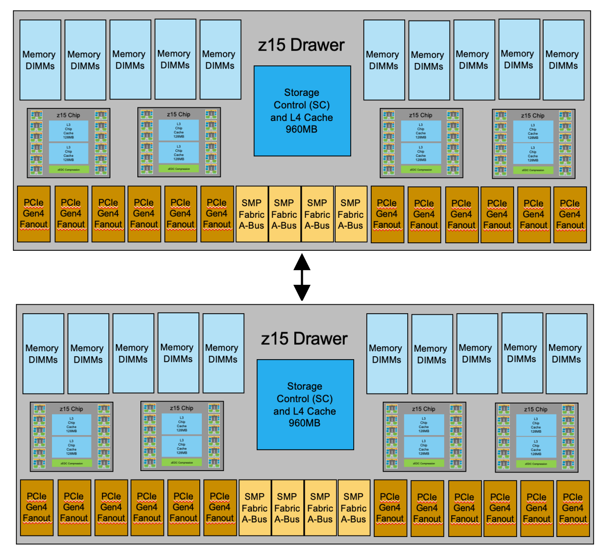

Start at 100 MIPS

0 to 65 Cores

256MB to 16TB RAIM

Sub-Uni's: "A" series through "Z" series (up to z/OS 6-way)

Fully Interconnected SMP Fabric

Rack’em and Stack’em

THE Z15 RACK(S)
THE LINUXONE III RACK(S)

Topics

Following…

- Following are several z15 T01 (or LinuxONE III LT1) server rack-buildout scenarios
  - and several z15 T02 (or LinuxONE III LT2) rack scenarios

- They are examples
- They illustrate key concepts
- There are many more

A z15 T01 single frame system requires 78% less floor space than a z14.
A z15 T01 two-frame system requires 56% less floor space than a z14.
A z15 T01 three-frame system requires 34% less floor space than a z14.
A z15 T01 four-frame system requires 12% less floor space than a z14.

Topics

z15 T02 (LT2) Rack: Scenario 1 – iPDU powered I/O Heavy

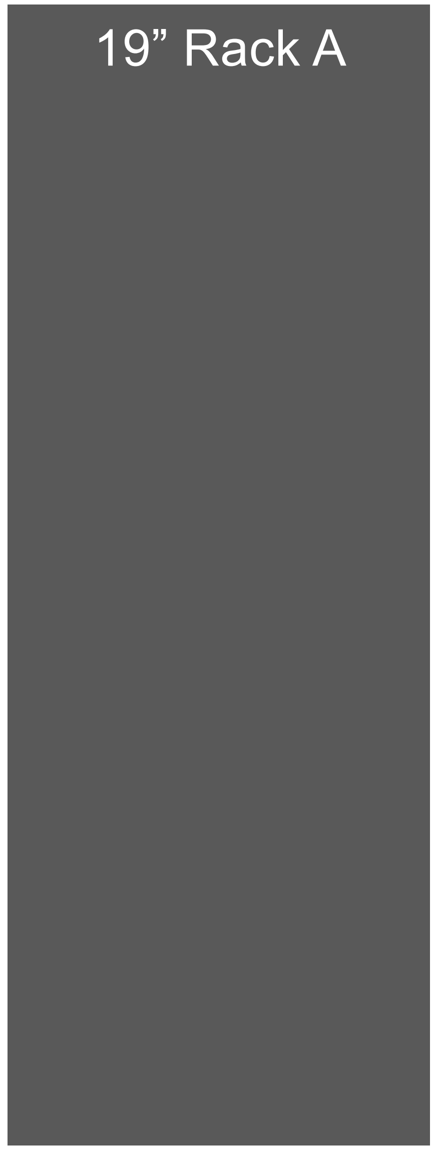

z15 T02 (LT2) Rack: Scenario 1 – iPDU powered I/O Heavy

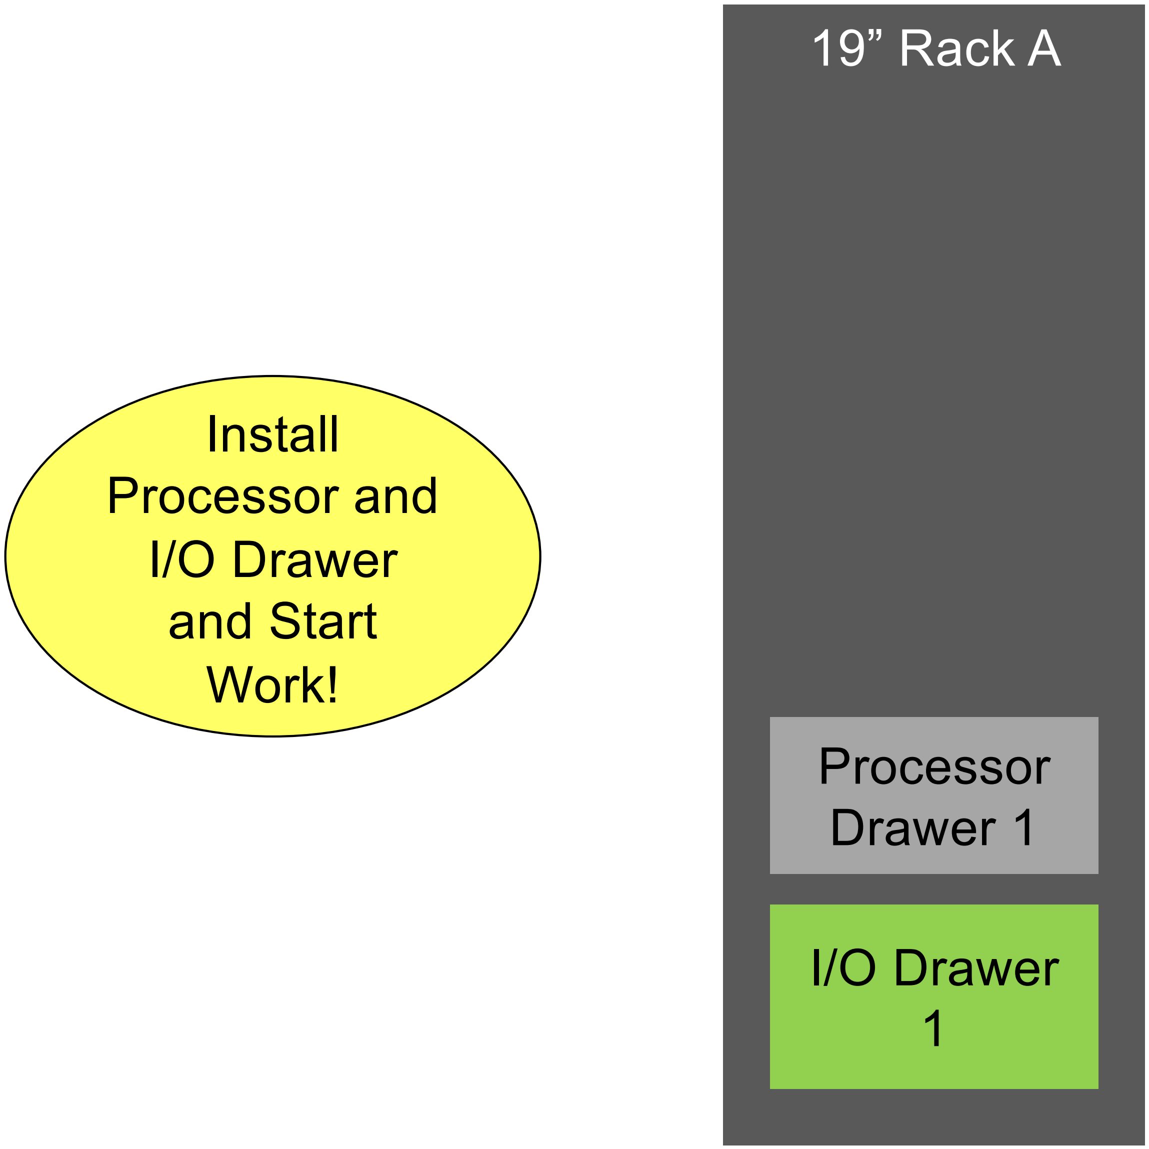

Topics

z15 T02 (LT2) Rack: Scenario 1 – iPDU powered I/O Heavy

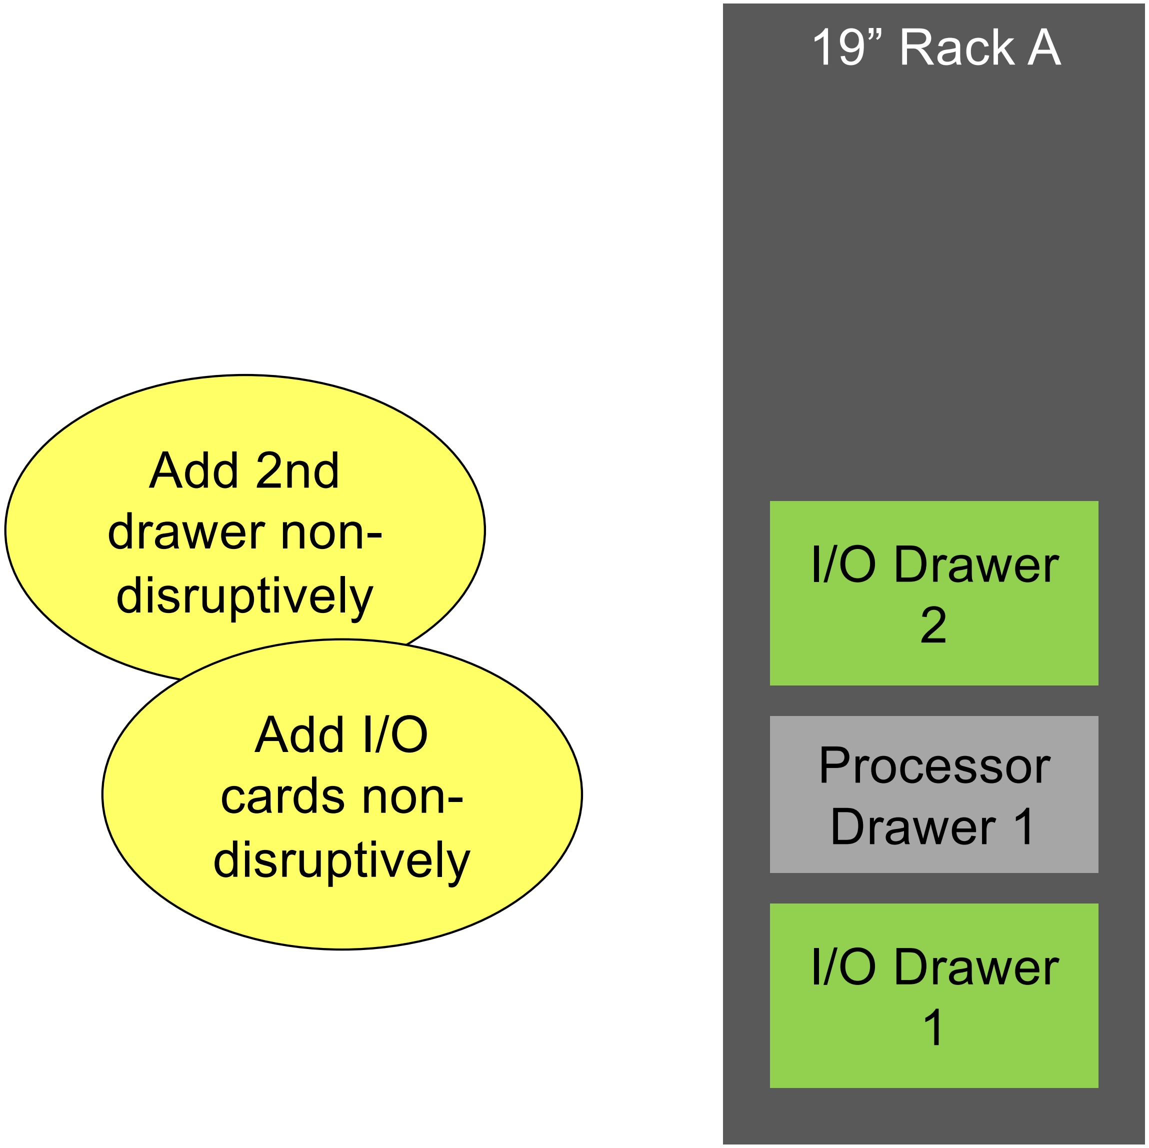

Topics

z15 T02 (LT2) Rack: Scenario 1 – iPDU powered I/O Heavy

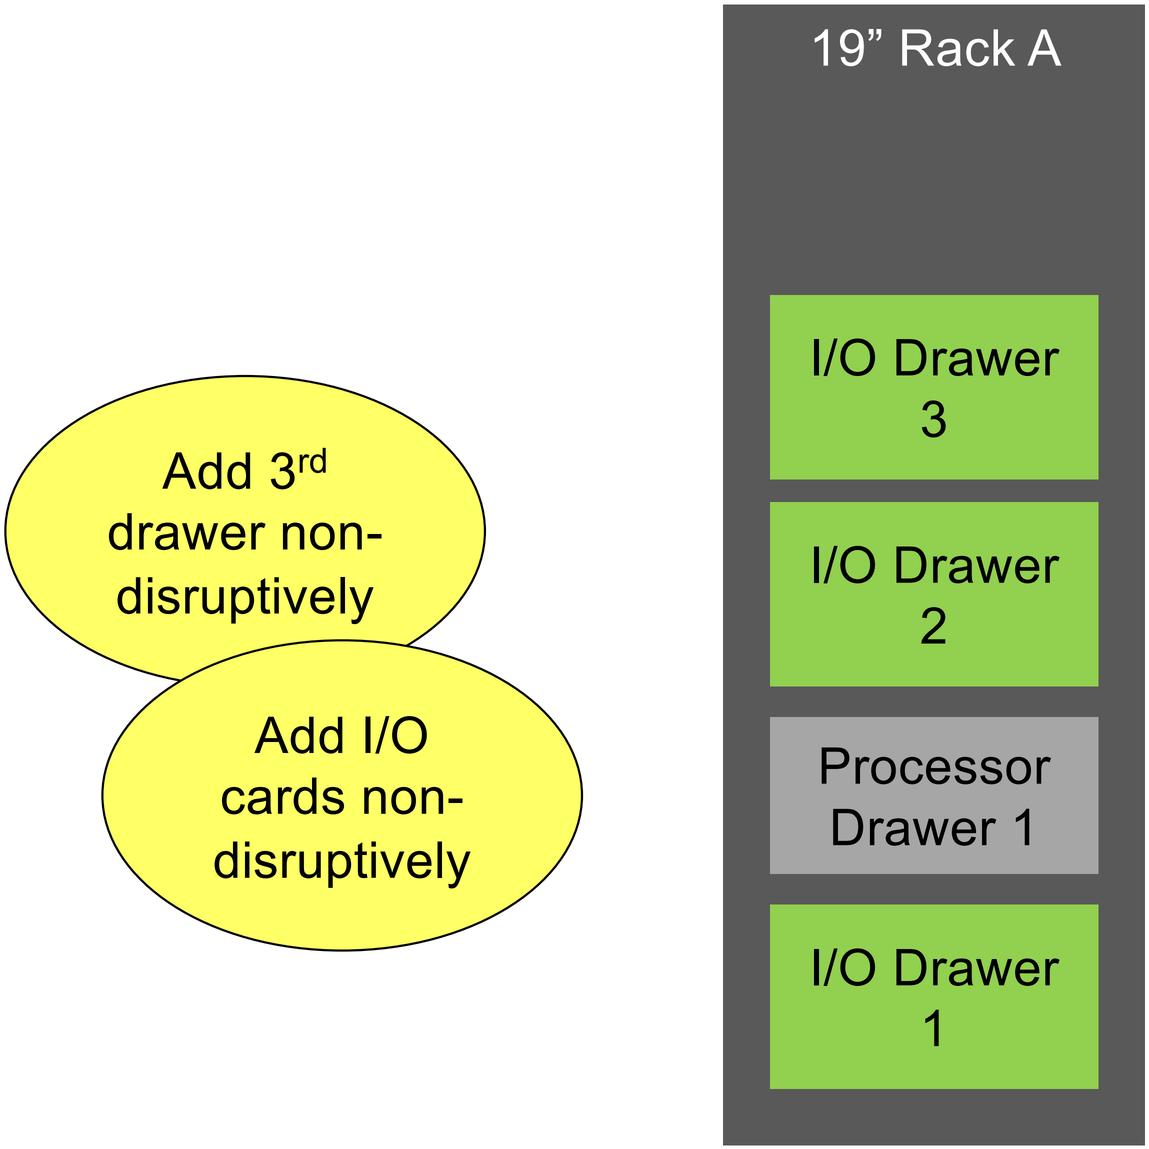

Topics

z15 T02 (LT2) Rack: Scenario 1 – iPDU powered I/O Heavy

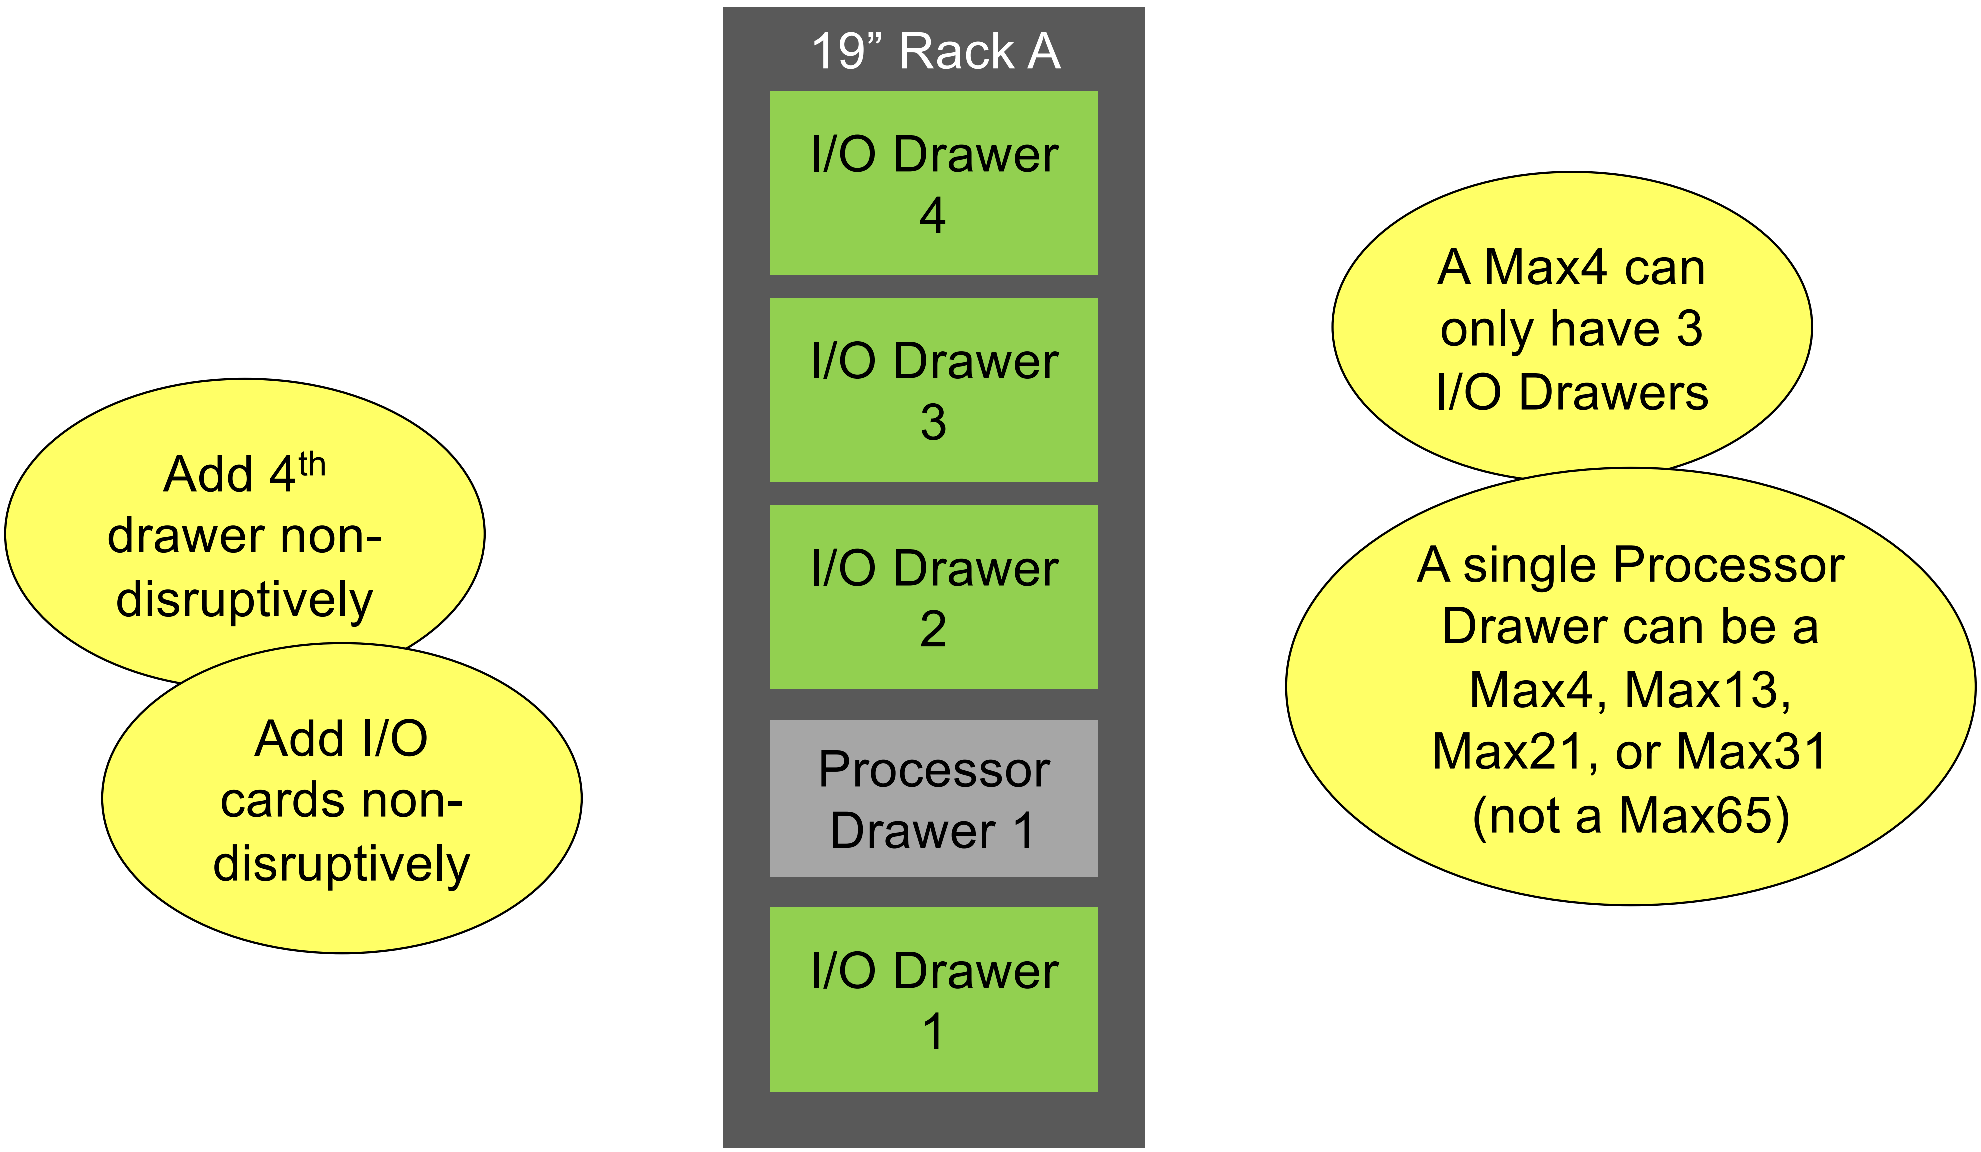

z15 T02 (LT2) Rack: Scenario 1 – iPDU powered CPU Heavy

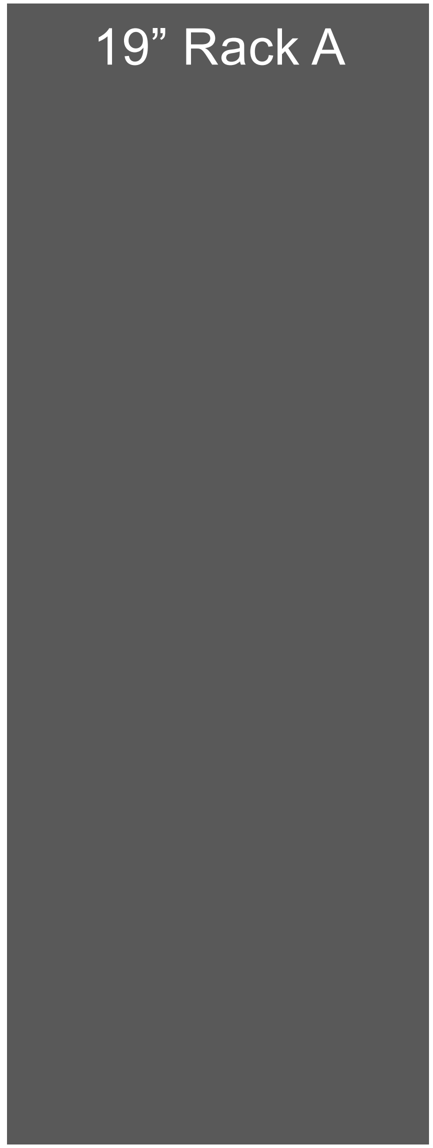

z15 T02 (LT2) Rack: Scenario 1 – iPDU powered CPU Heavy

Install Processor and I/O Drawer and Start Work!

But since anticipating heavy CPU growth, put a reserve in for a future Processor Drawer

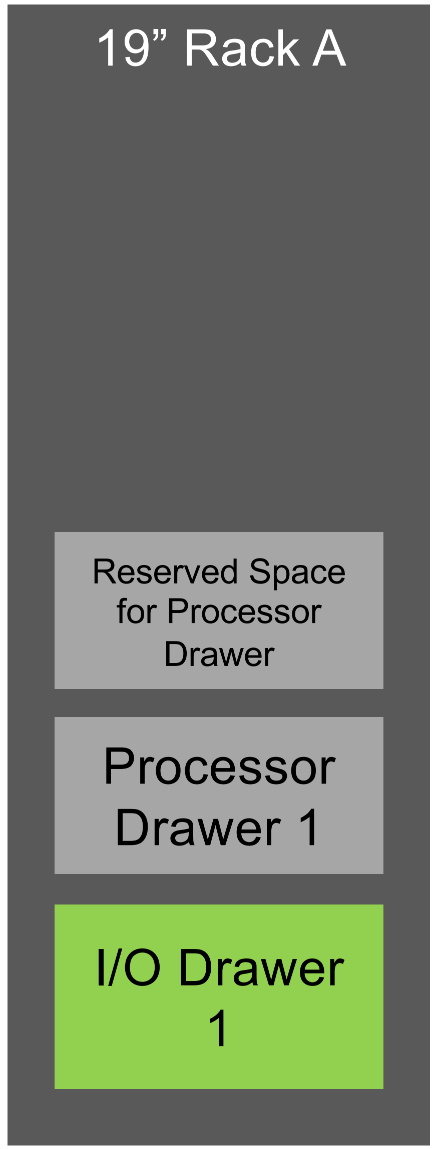

Topics

z15 T02 (LT2) Rack: Scenario 1 – iPDU powered CPU Heavy

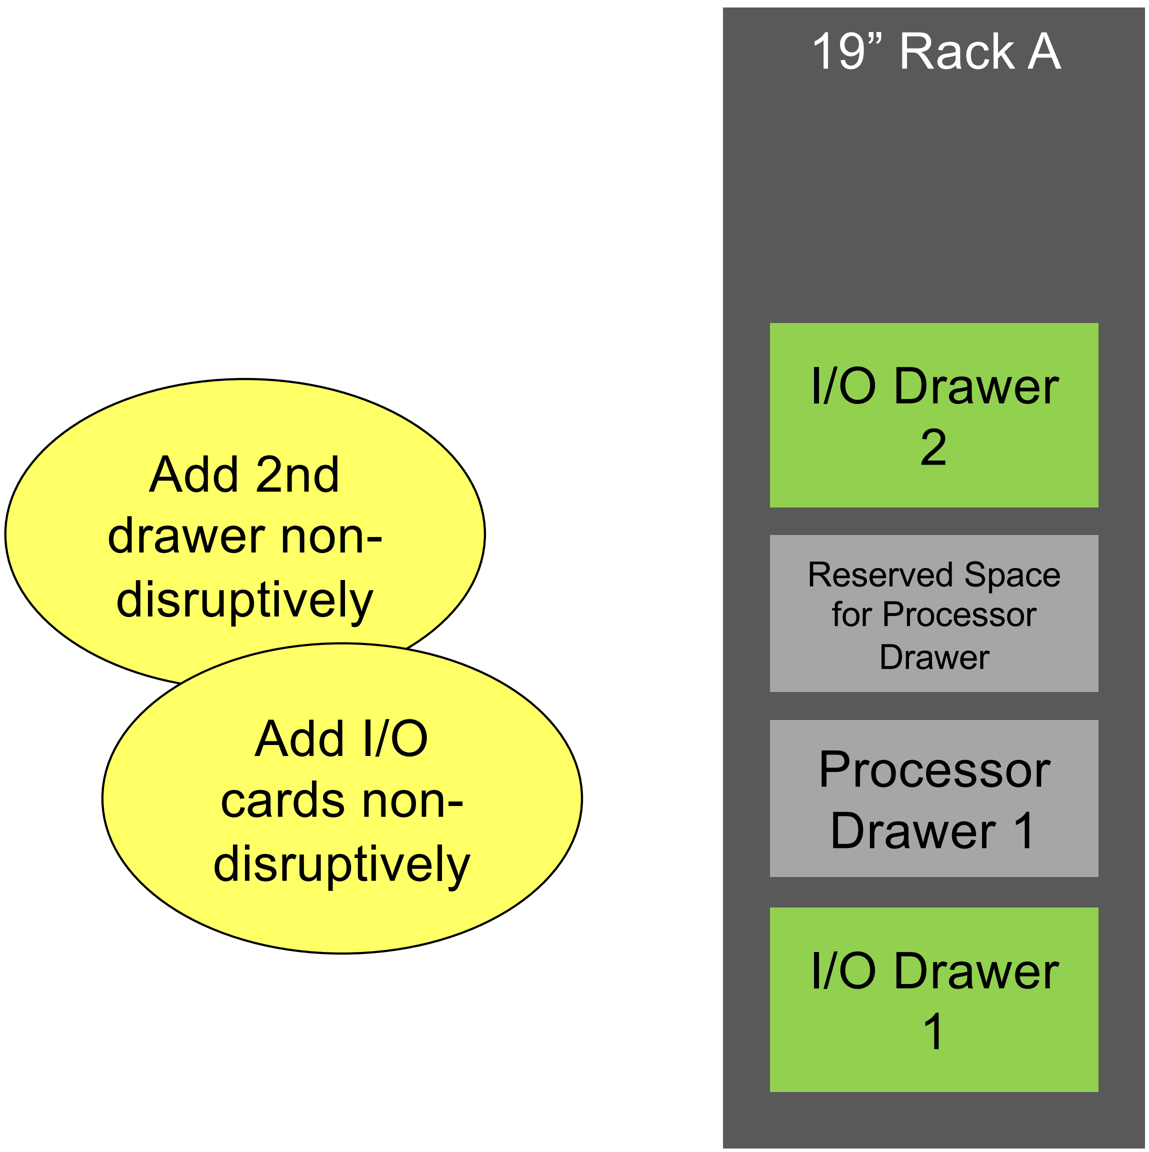

z15 T02 (LT2) Rack: Scenario 1 – iPDU powered CPU Heavy

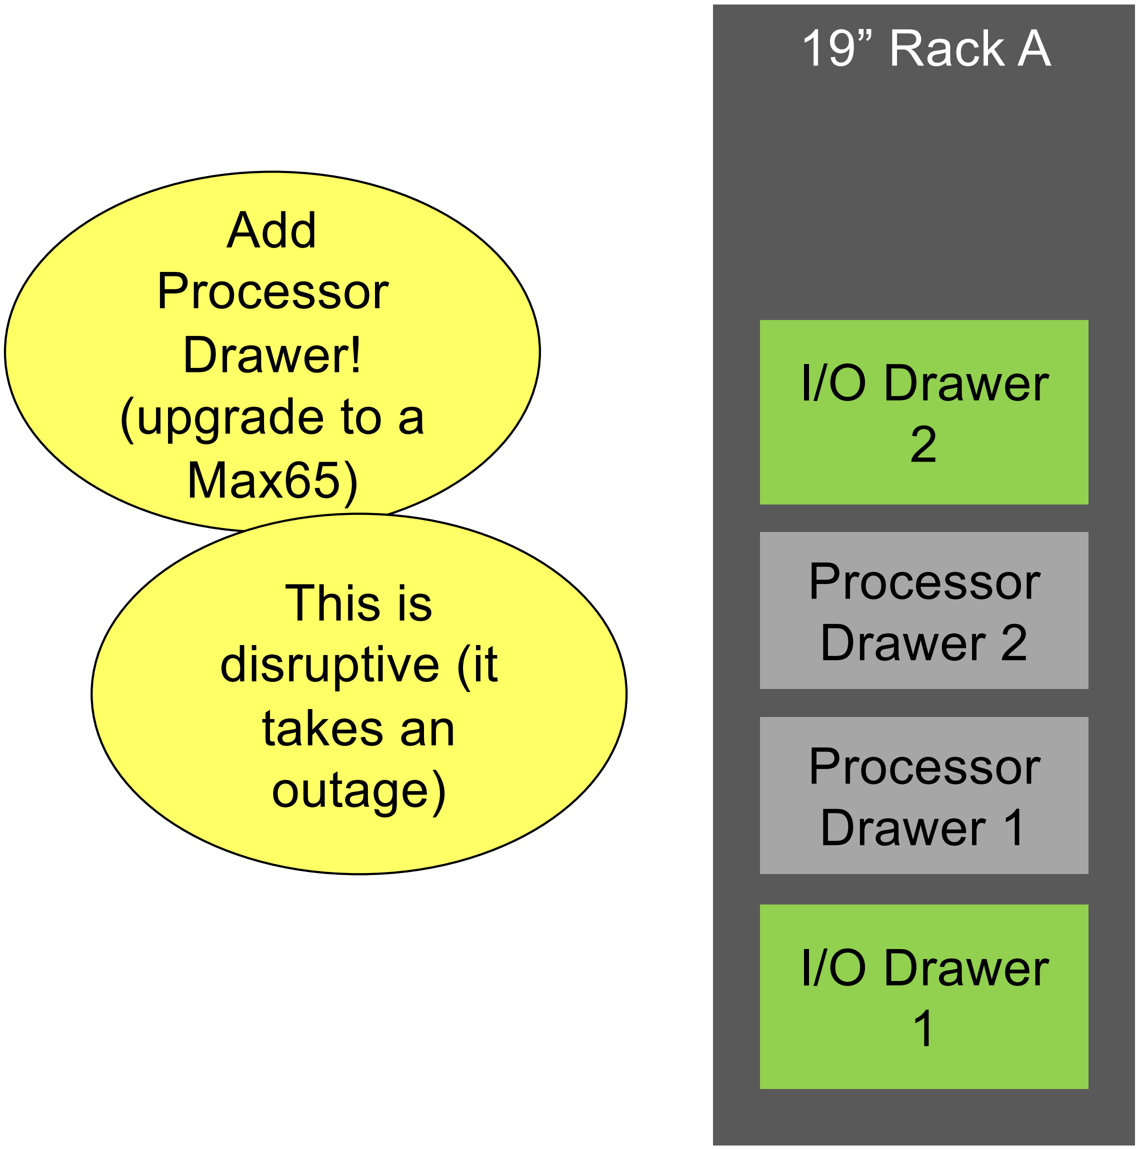

z15 T02 (LT2) Rack: Scenario 1 – iPDU powered CPU Heavy

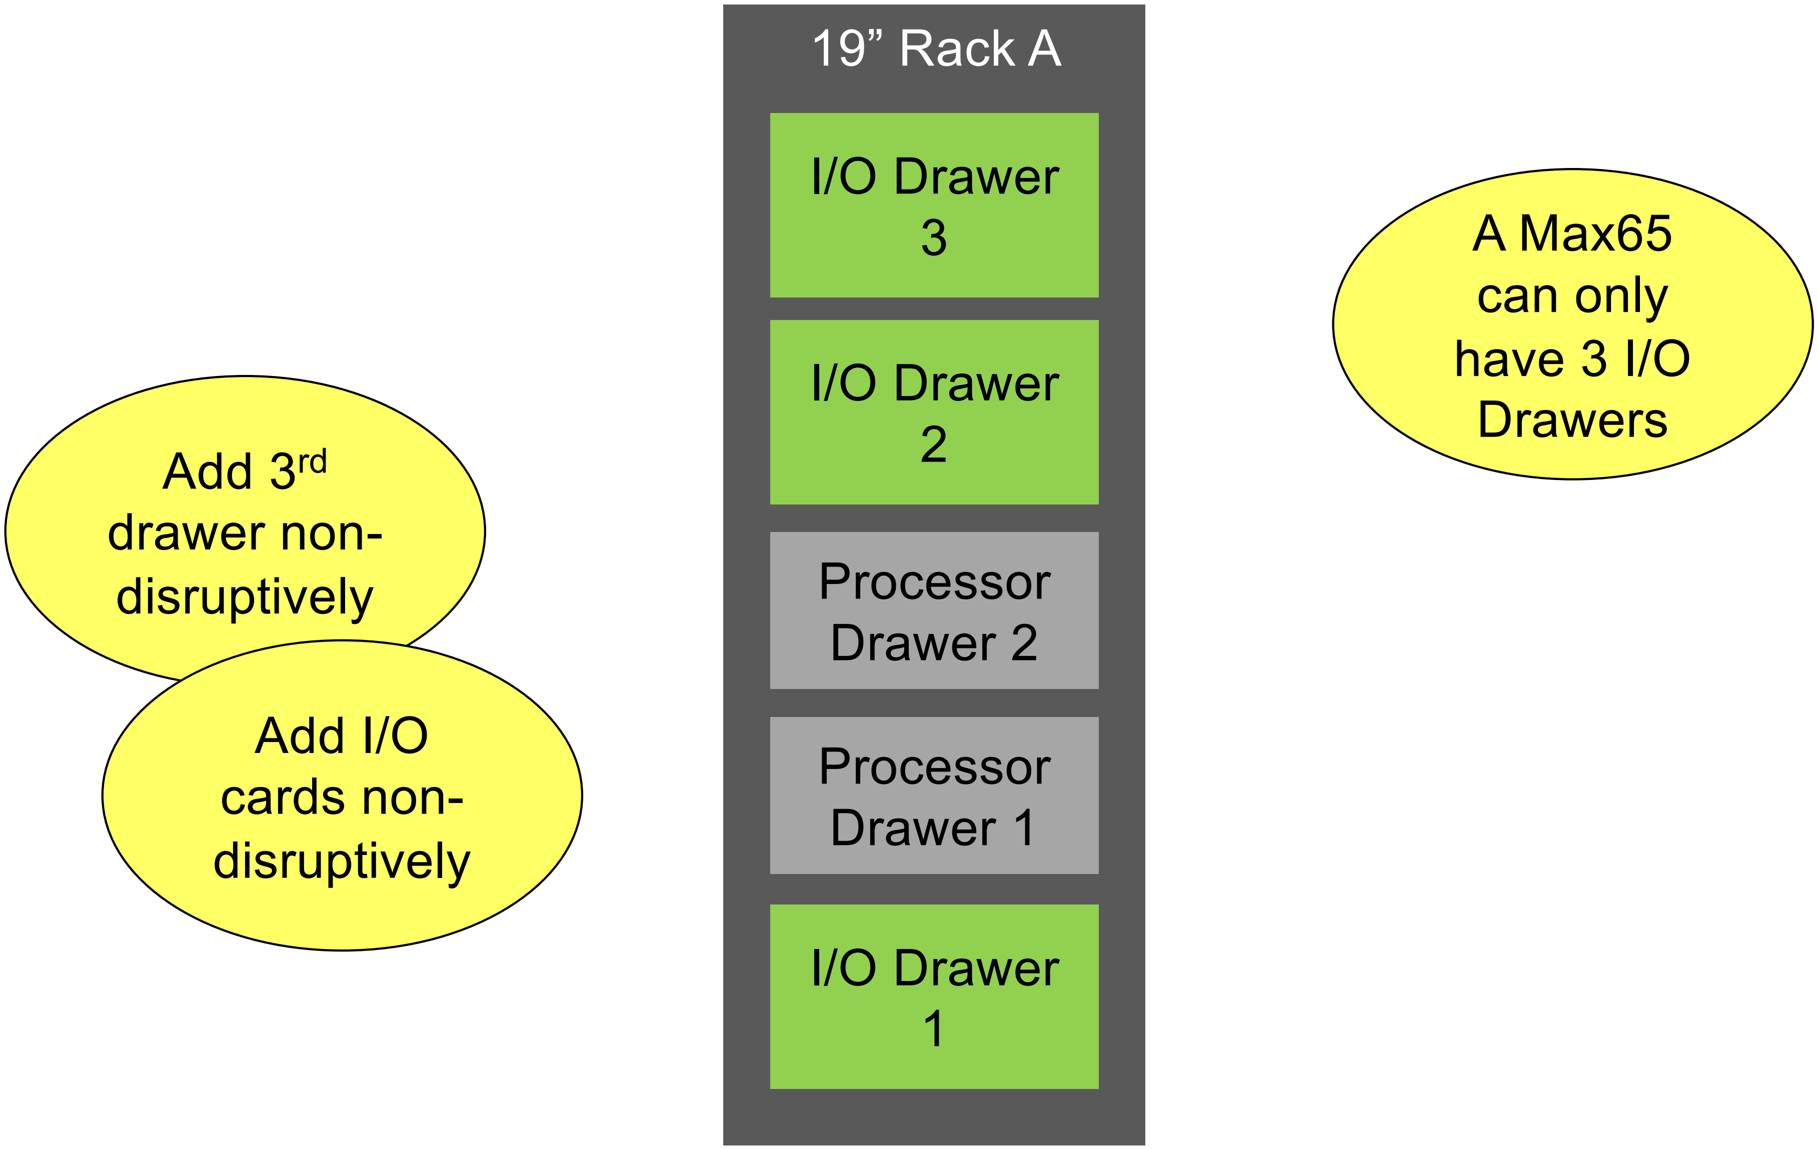

z15 T01 (LT1) Racks: Scenario 1 – iPDU powered I/O Heavy

IBM Supplied 19” Rack (with cool covers!)

19” Rack A

Topics

z15 T01 (LT1) Racks: Scenario 1 – iPDU powered I/O Heavy

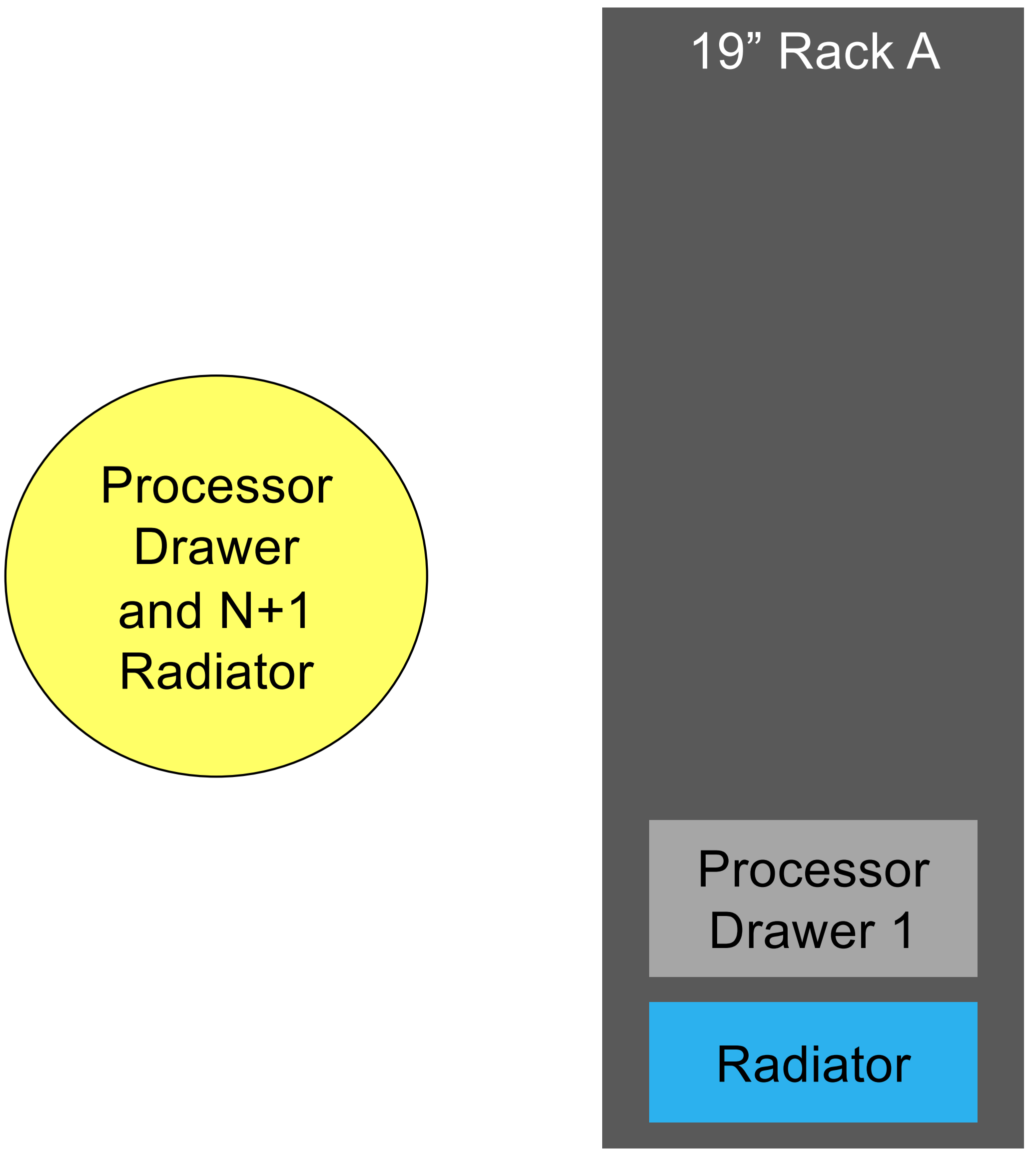

z15 T01 (LT1) Racks: Scenario 1 – iPDU powered I/O Heavy

Install I/O Drawer

Install I/O Cards

Connect Your I/O

…

Start Operations

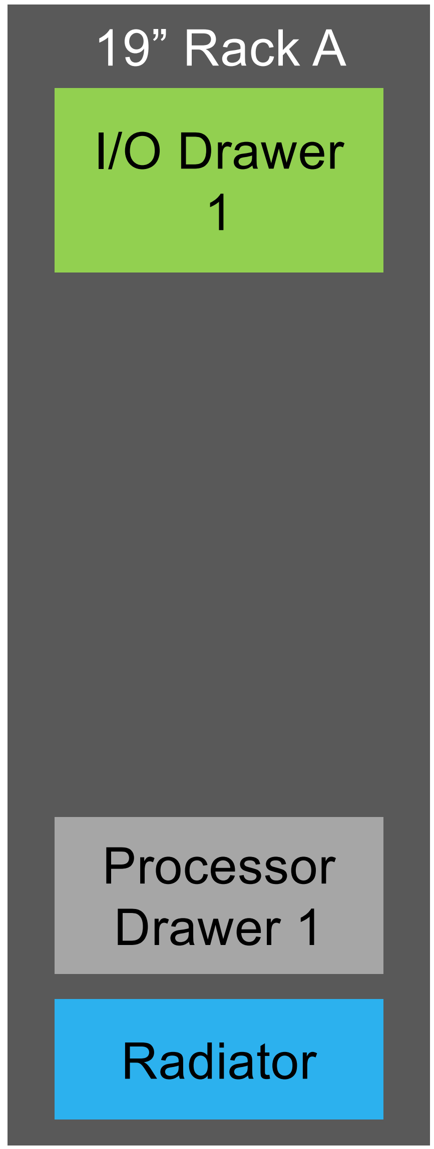

Processor Drawers use PCIe Gen4 Tech and Dual PCIe Gen 3 Fanouts to I/O Drawers

Topics

z15 T01 (LT1) Racks: Scenario 1 – iPDU powered I/O Heavy

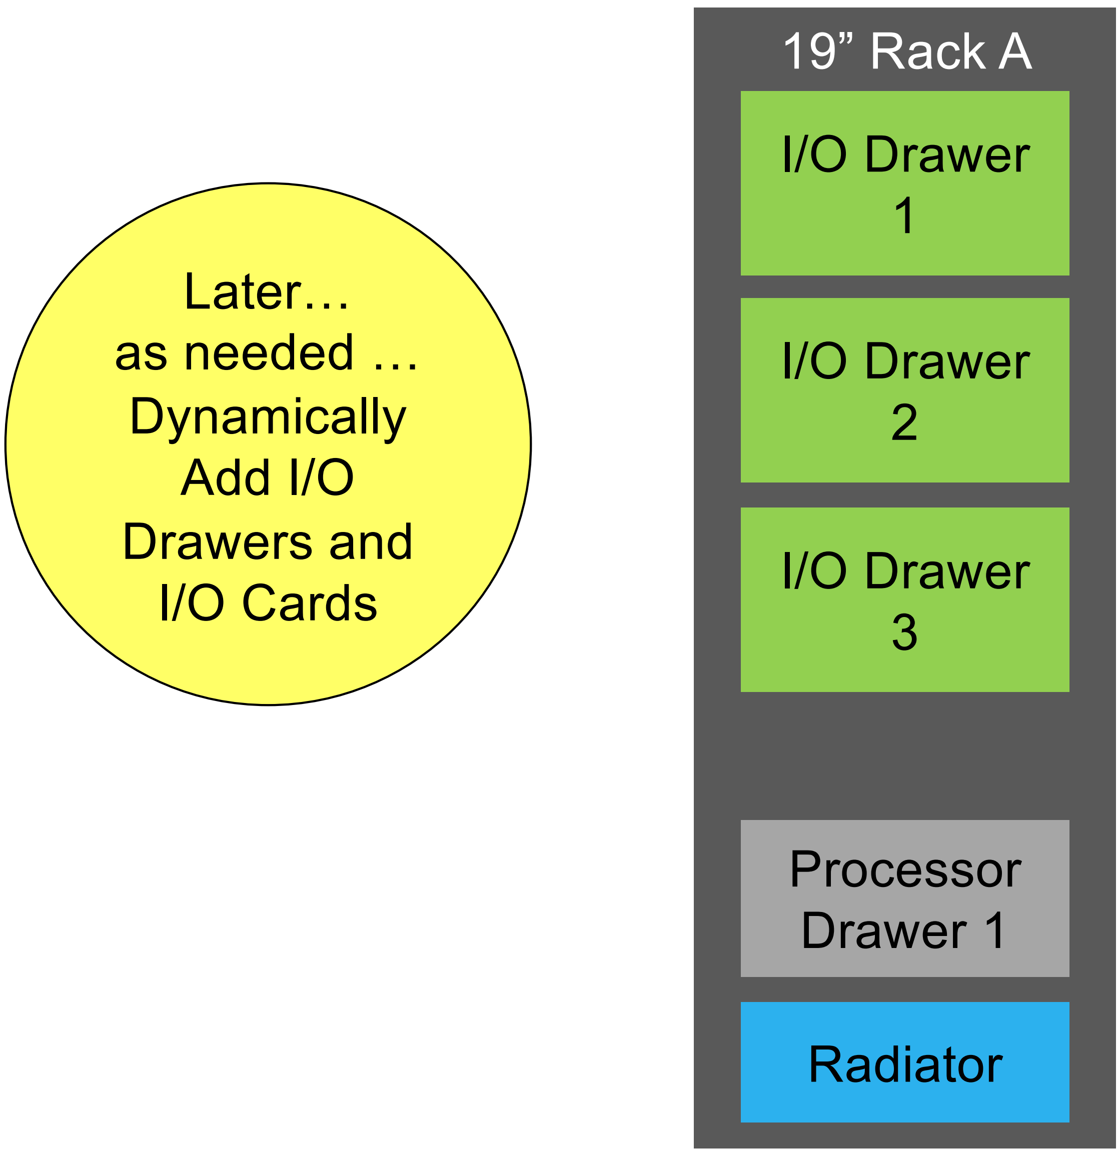

z15 T01 (LT1) Racks: Scenario 1 – iPDU powered I/O Heavy

Grow to a 2nd Rack for More I/O Capacity

Presuming floorspace and power are available, the addition of this rack can be NON-disruptive!

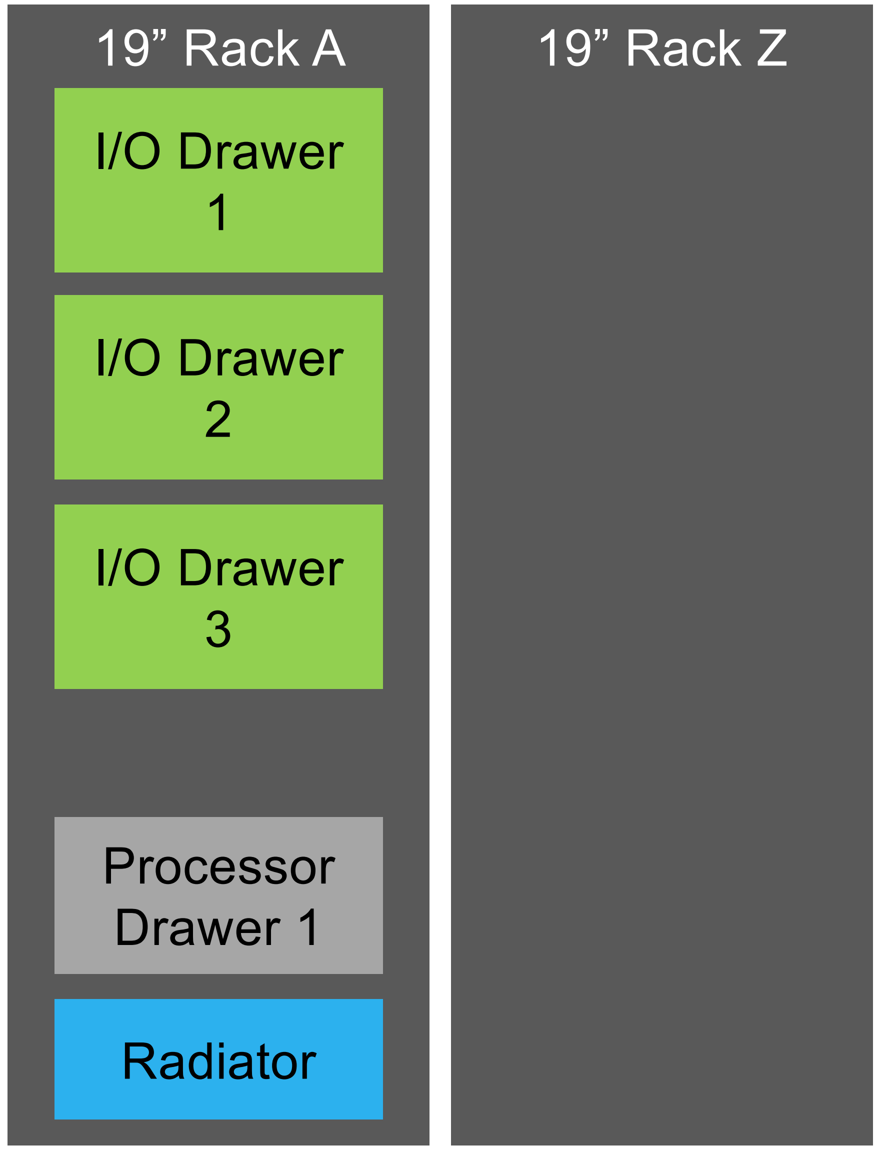

z15 T01 (LT1) Racks: Scenario 1 – iPDU powered I/O Heavy

Once the Rack is Installed, Dynamically Add I/O Drawers and Cards

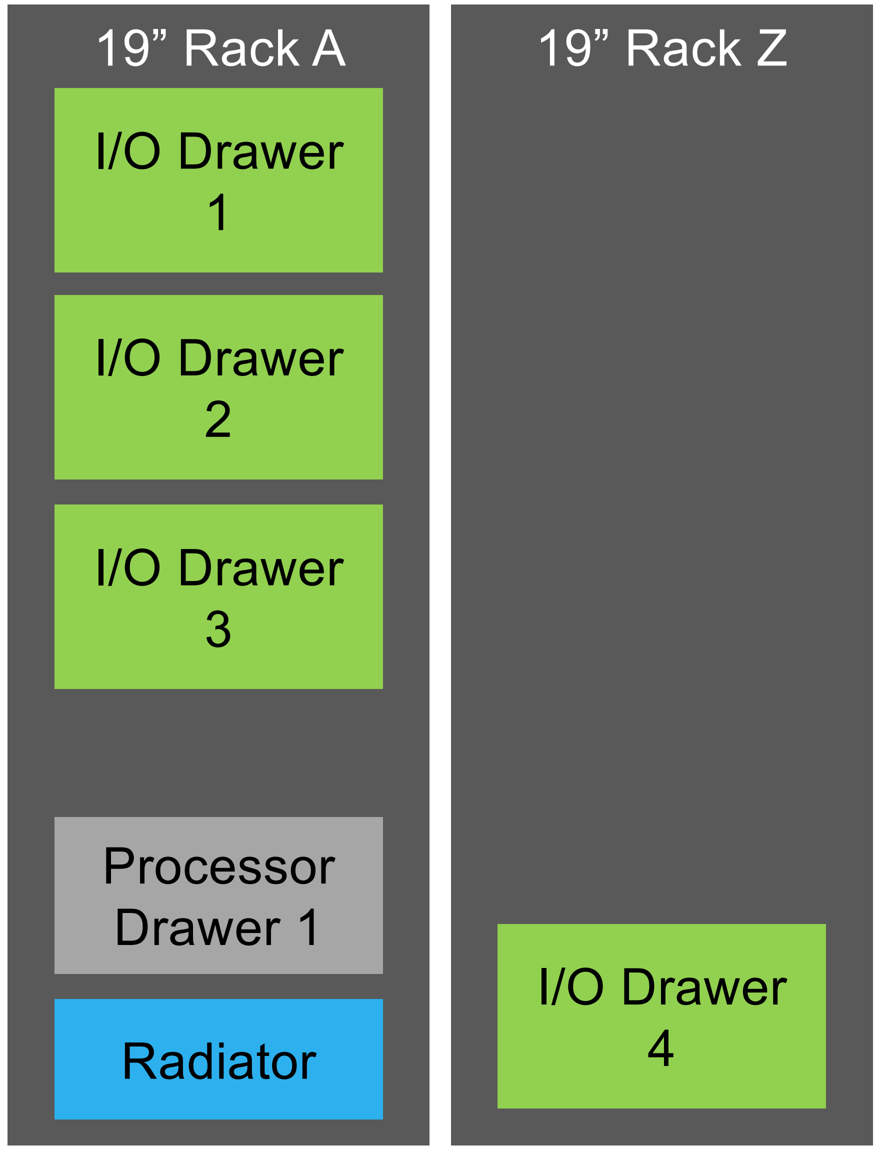

Topics

z15 T01 (LT1) Racks: Scenario 1 – iPDU powered I/O Heavy

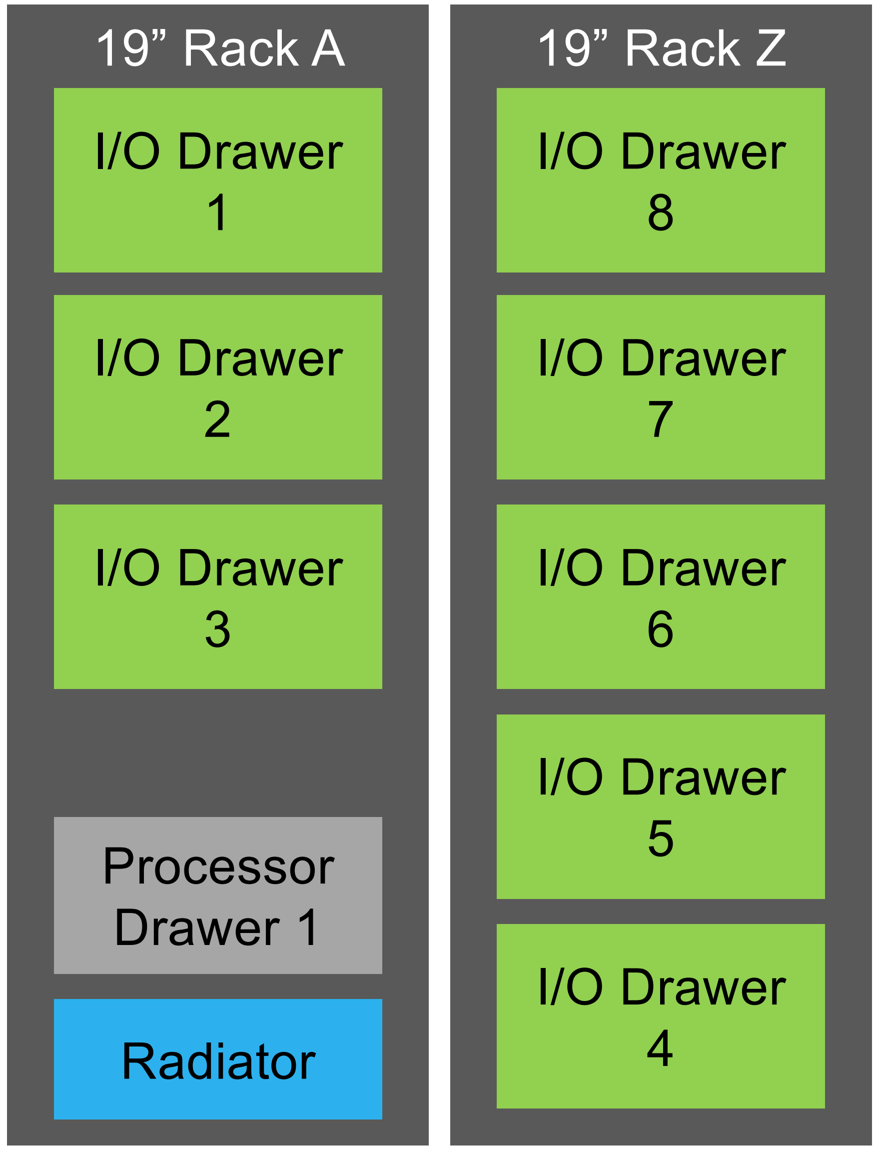

z15 T01 (LT1) Racks: Scenario 1 – iPDU powered I/O Heavy

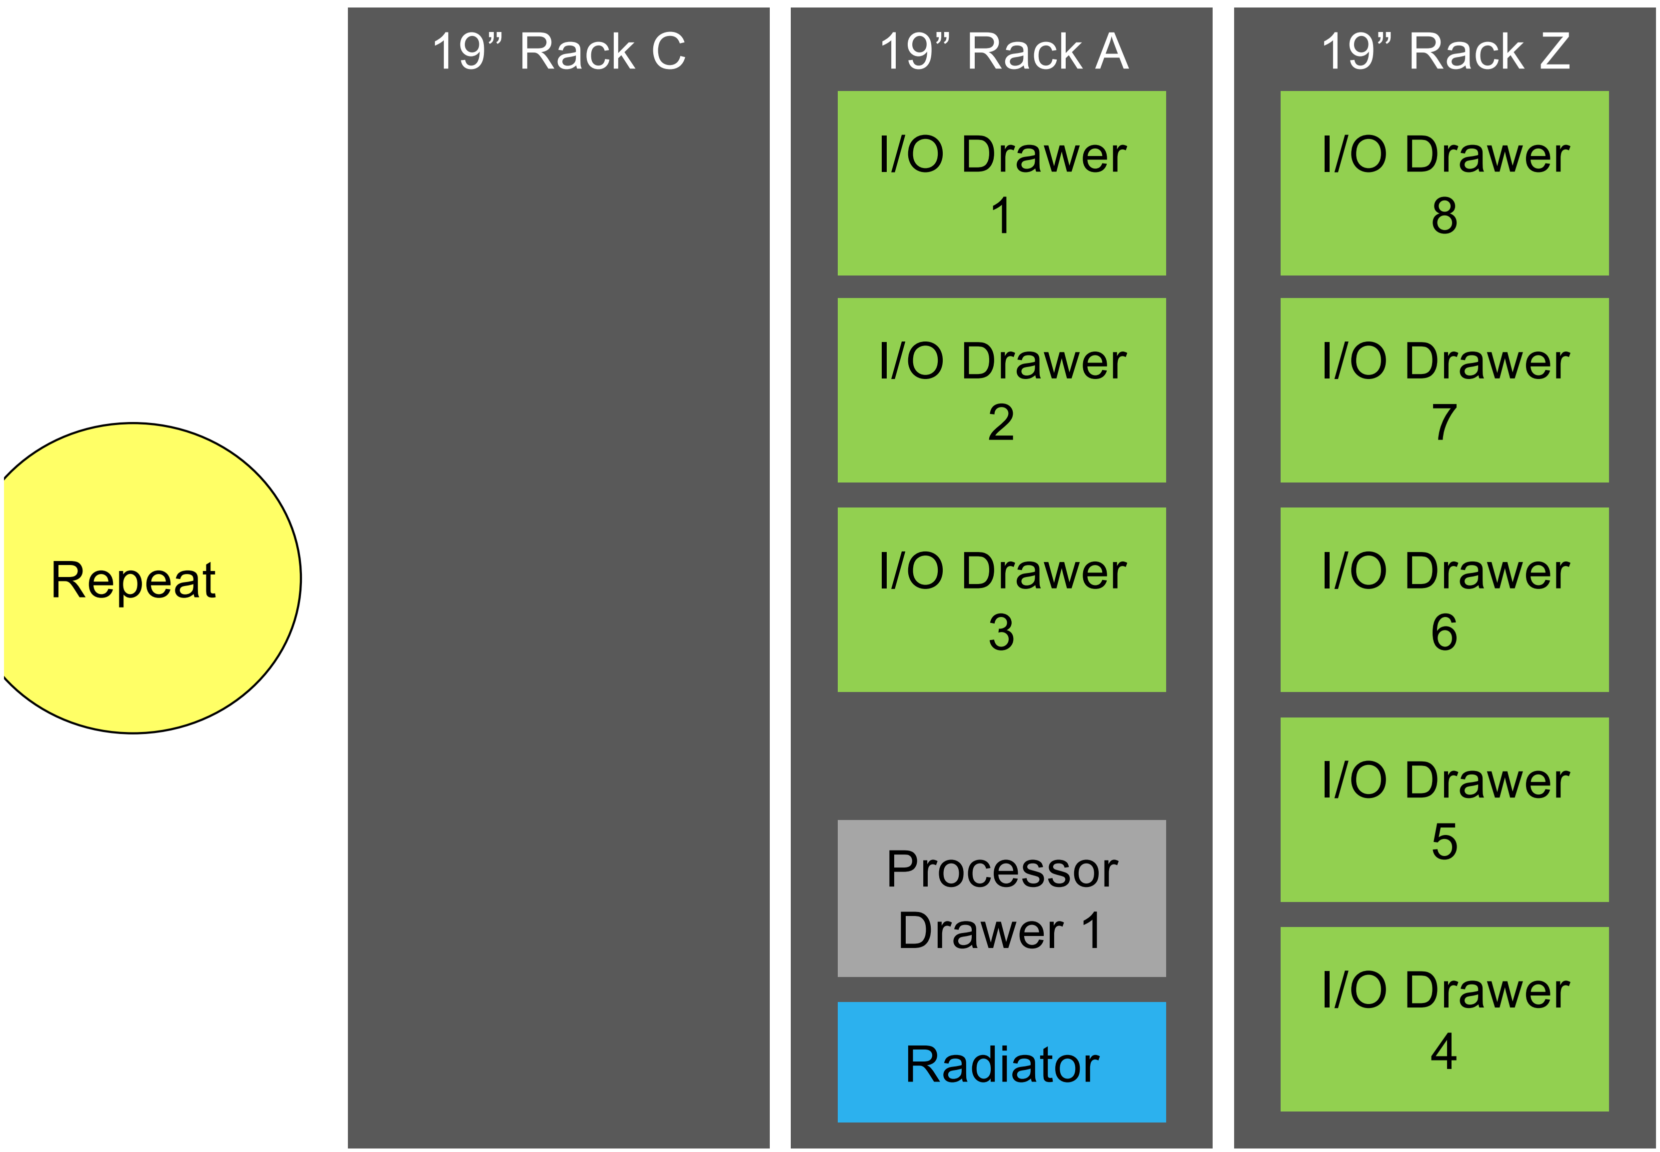

Topics

z15 T01 (LT1) Racks: Scenario 1 – iPDU powered I/O Heavy

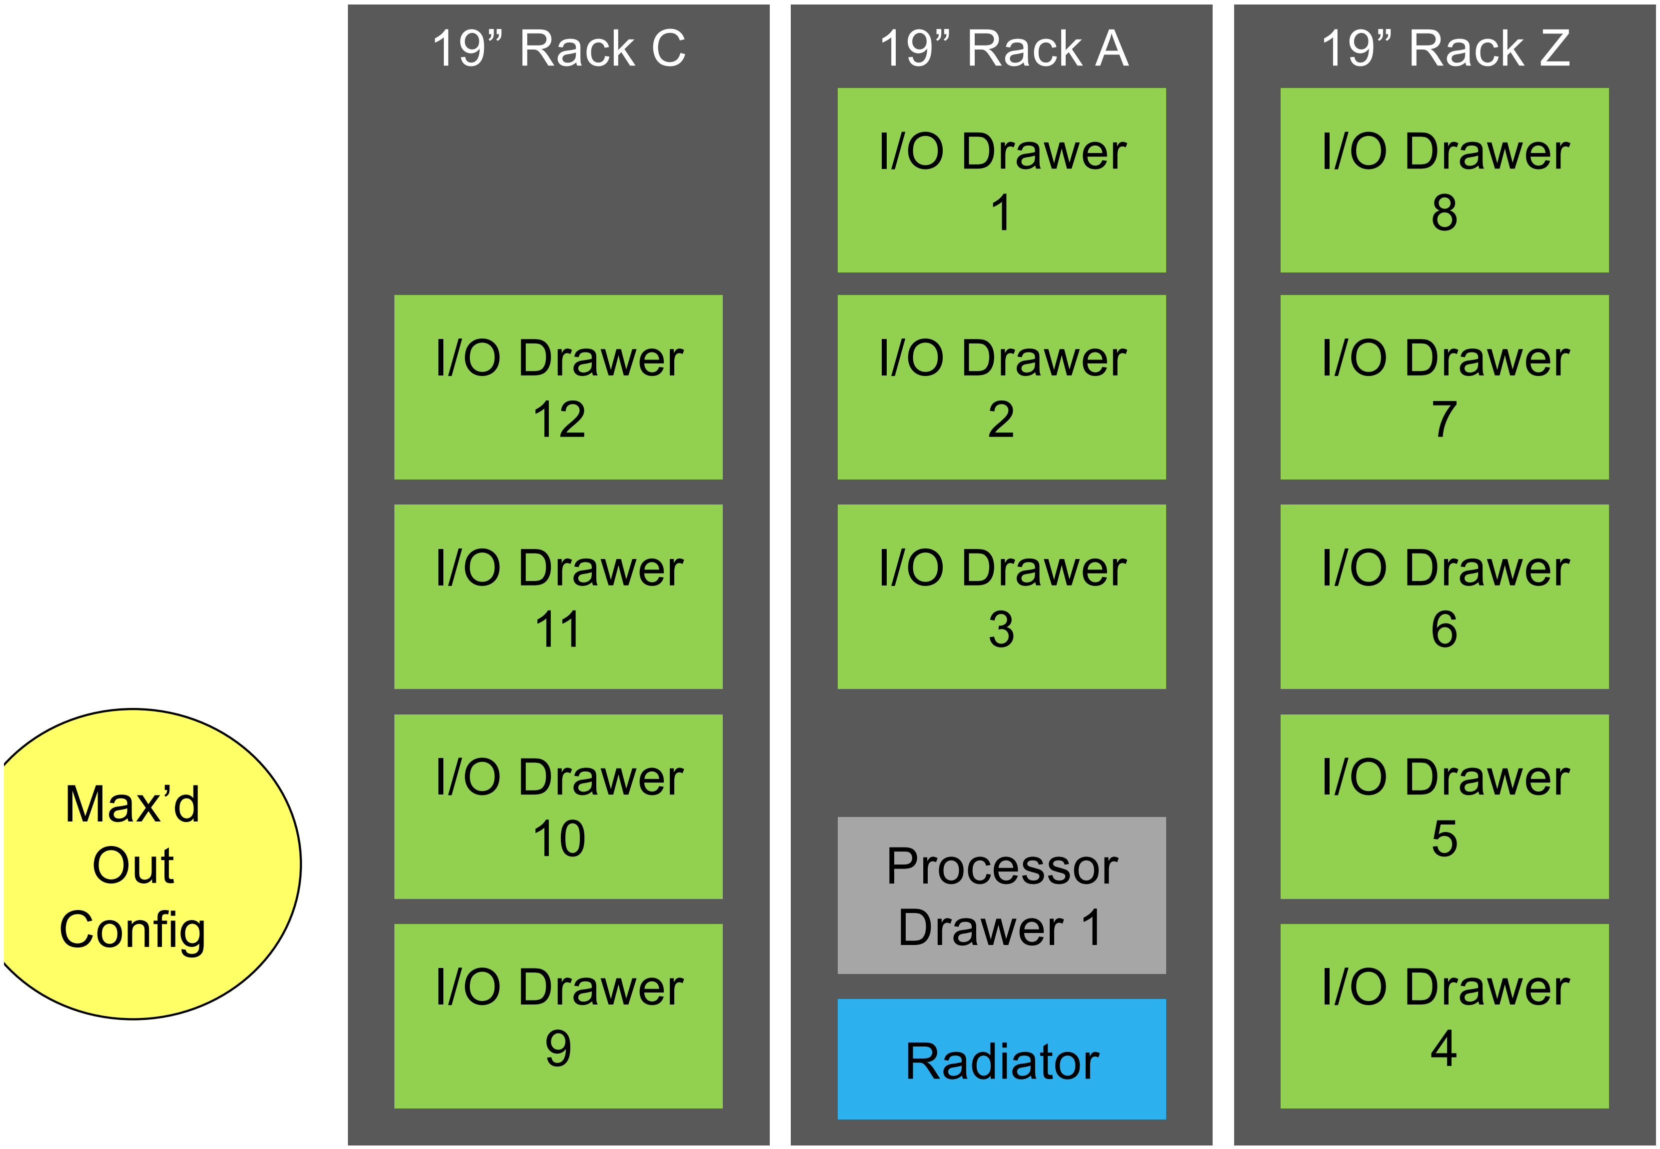

Topics

z15 T01 (LT1) Racks: Scenario 2 – iPDU Powered CPU Heavy

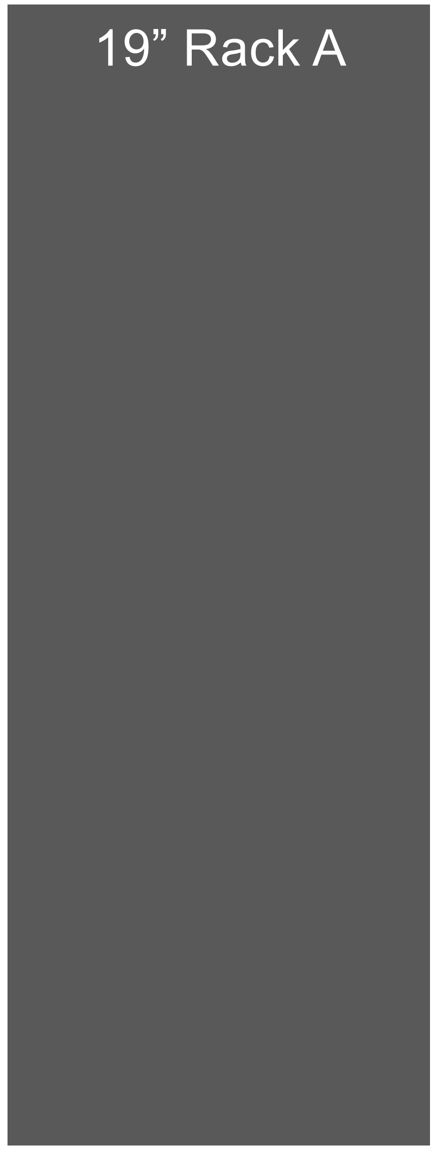

z15 T01 (LT1) Racks: Scenario 2 – iPDU Powered CPU Heavy

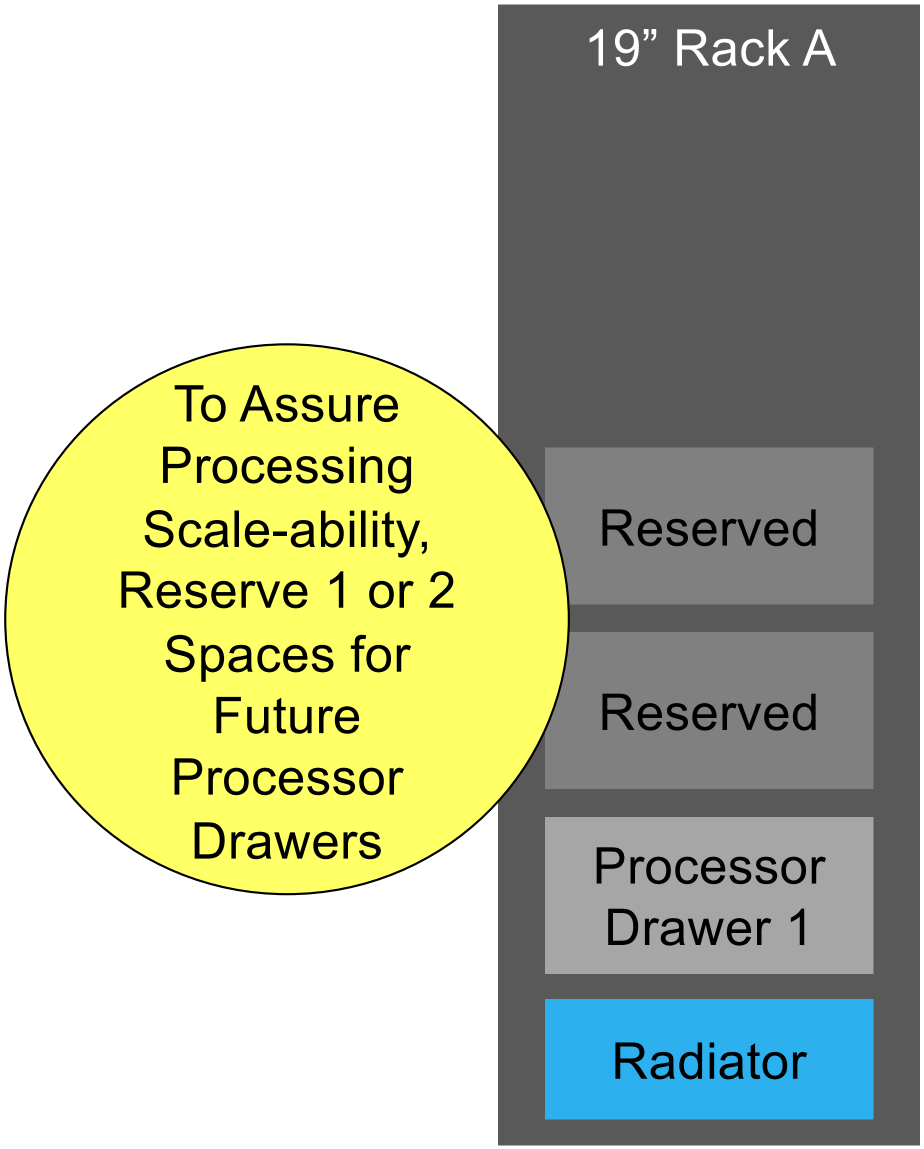

Topics

z15 T01 (LT1) Racks: Scenario 2 – iPDU Powered CPU Heavy

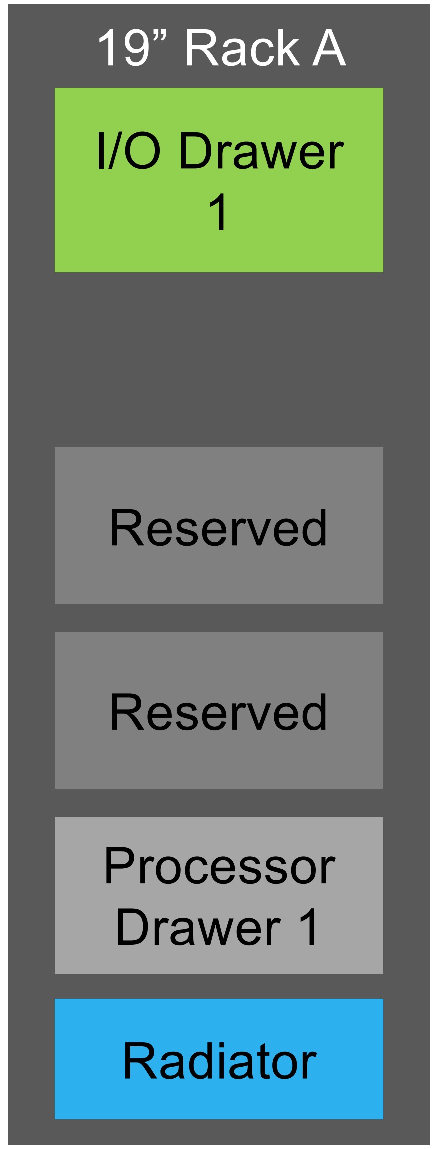

Everything that happens next CAN BE non-disruptive

Topics

z15 T01 (LT1) Racks: Scenario 2 – iPDU Powered CPU Heavy

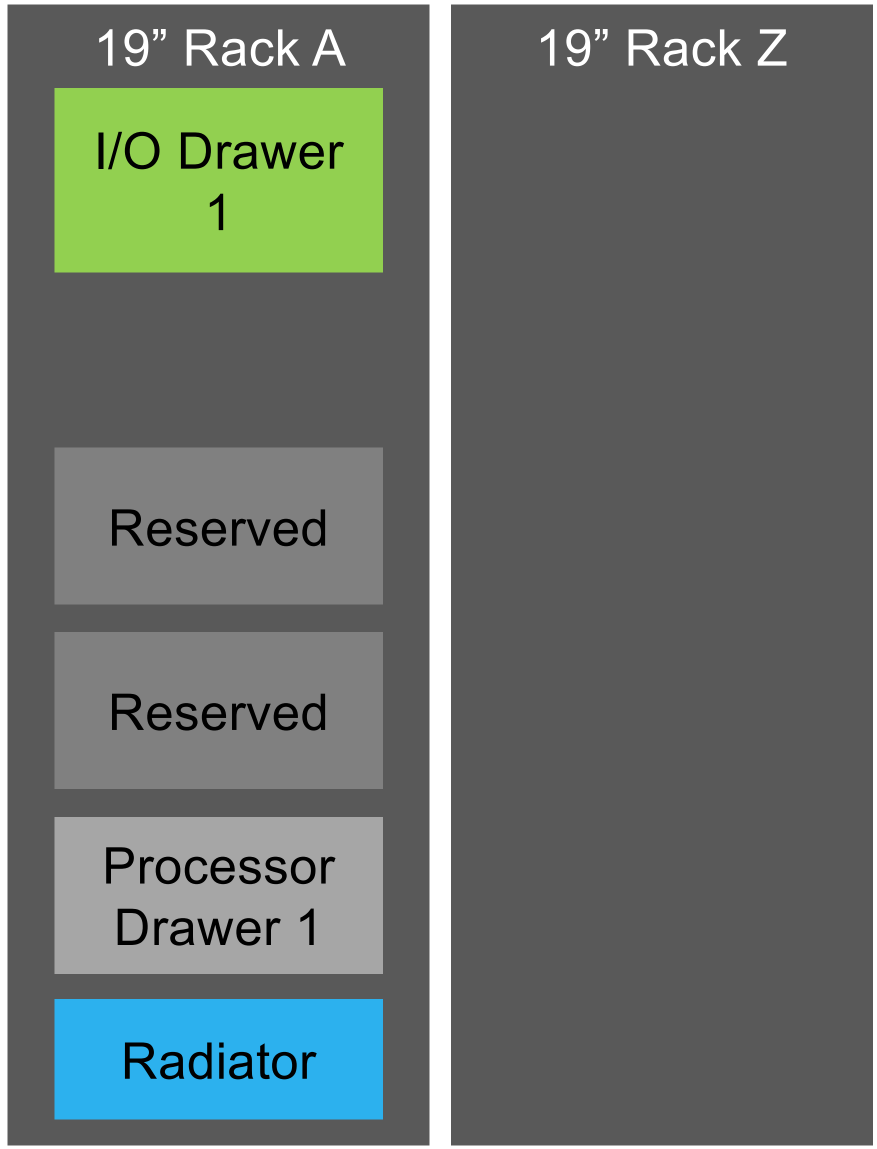

z15 T01 (LT1) Racks: Scenario 2 – iPDU Powered CPU Heavy

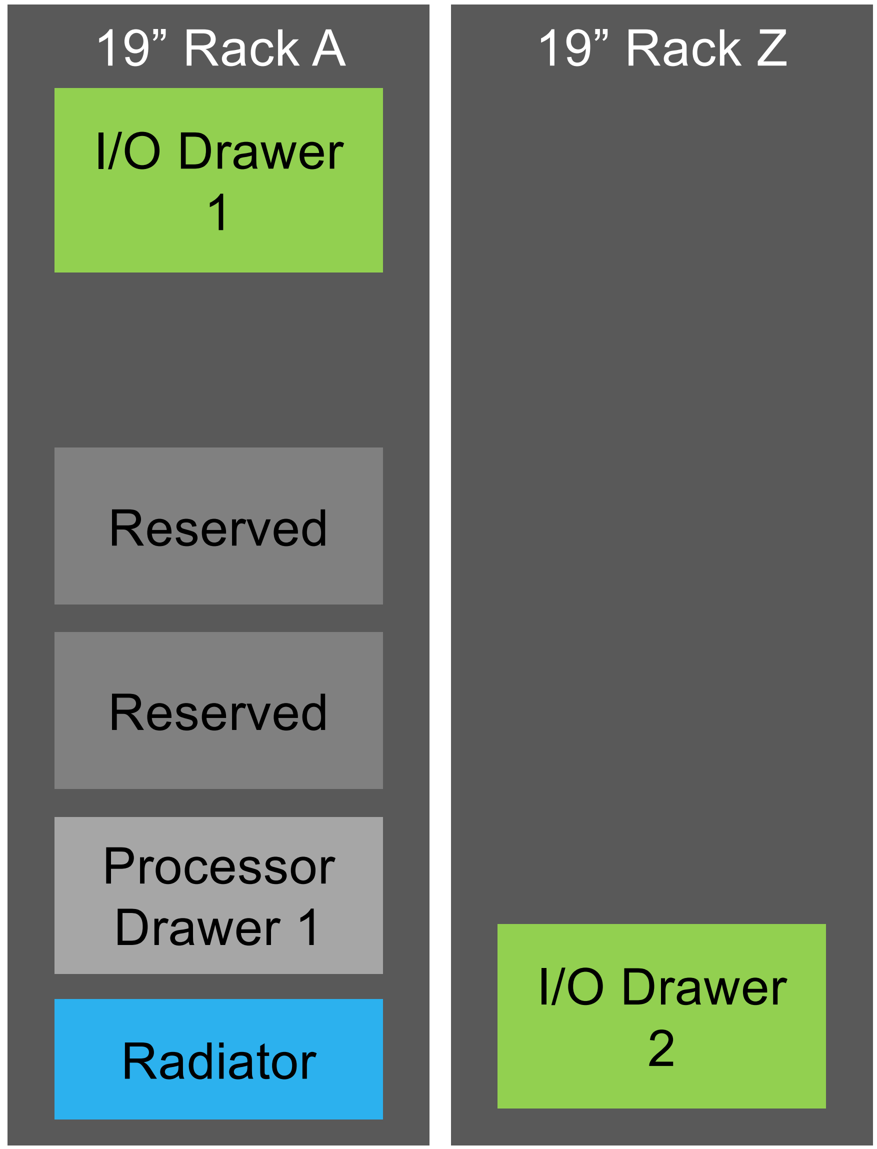

z15 T01 (LT1) Racks: Scenario 2 – iPDU Powered CPU Heavy

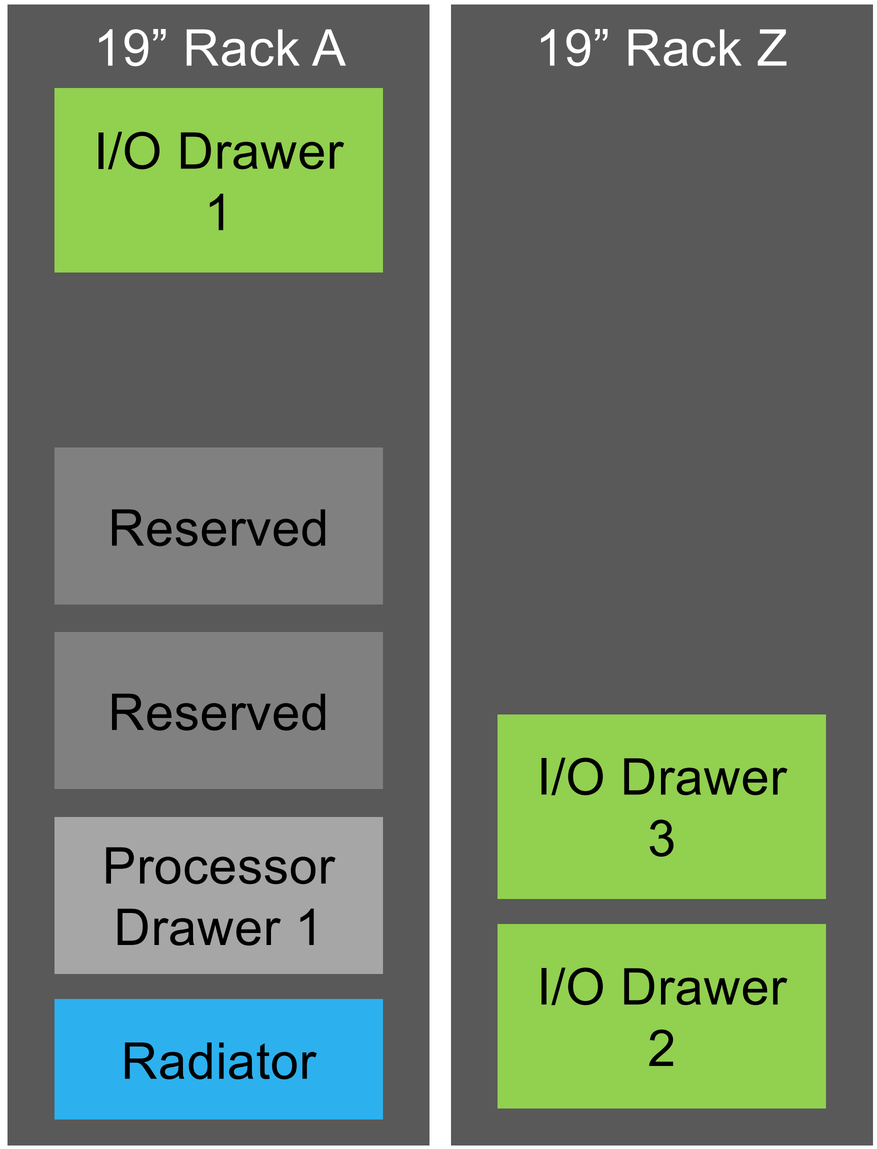

Topics

z15 T01 (LT1) Racks: Scenario 2 – iPDU Powered CPU Heavy

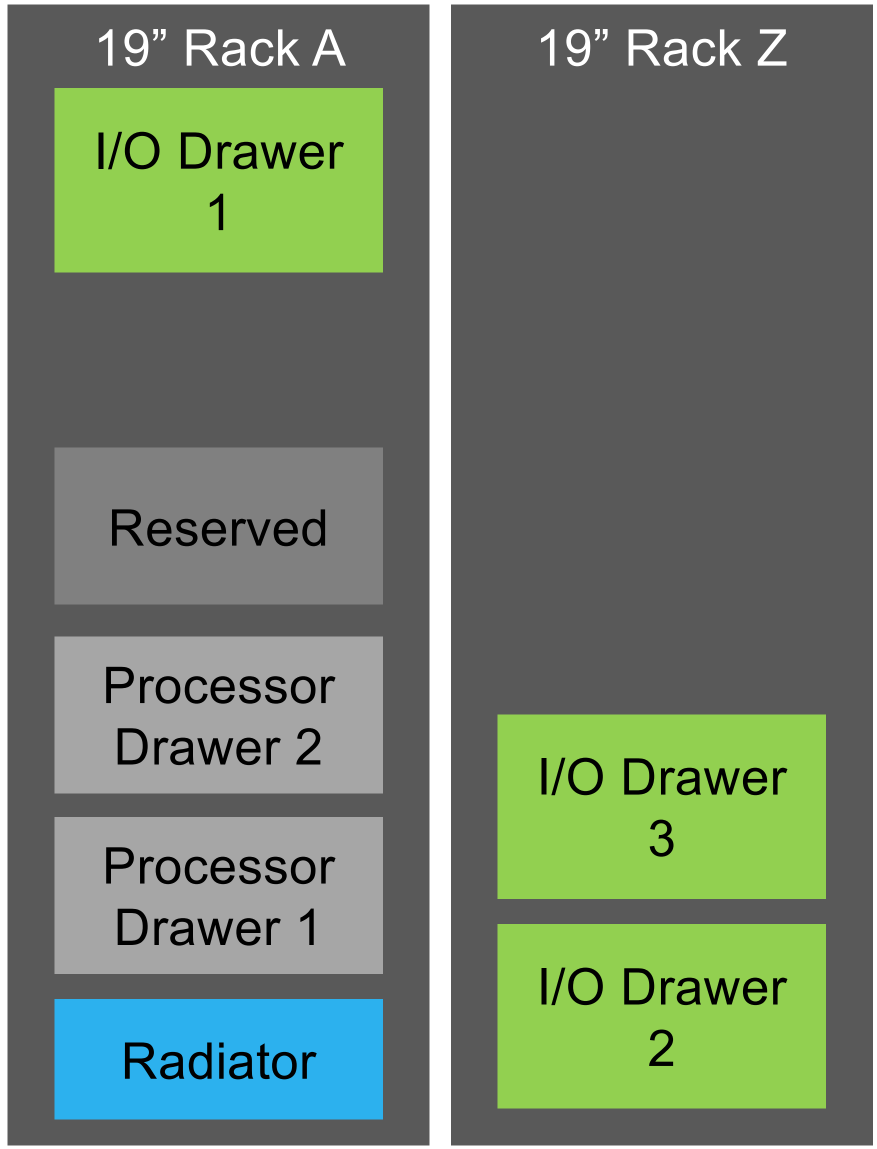

Topics

z15 T01 (LT1) Racks: Scenario 2 – iPDU Powered CPU Heavy

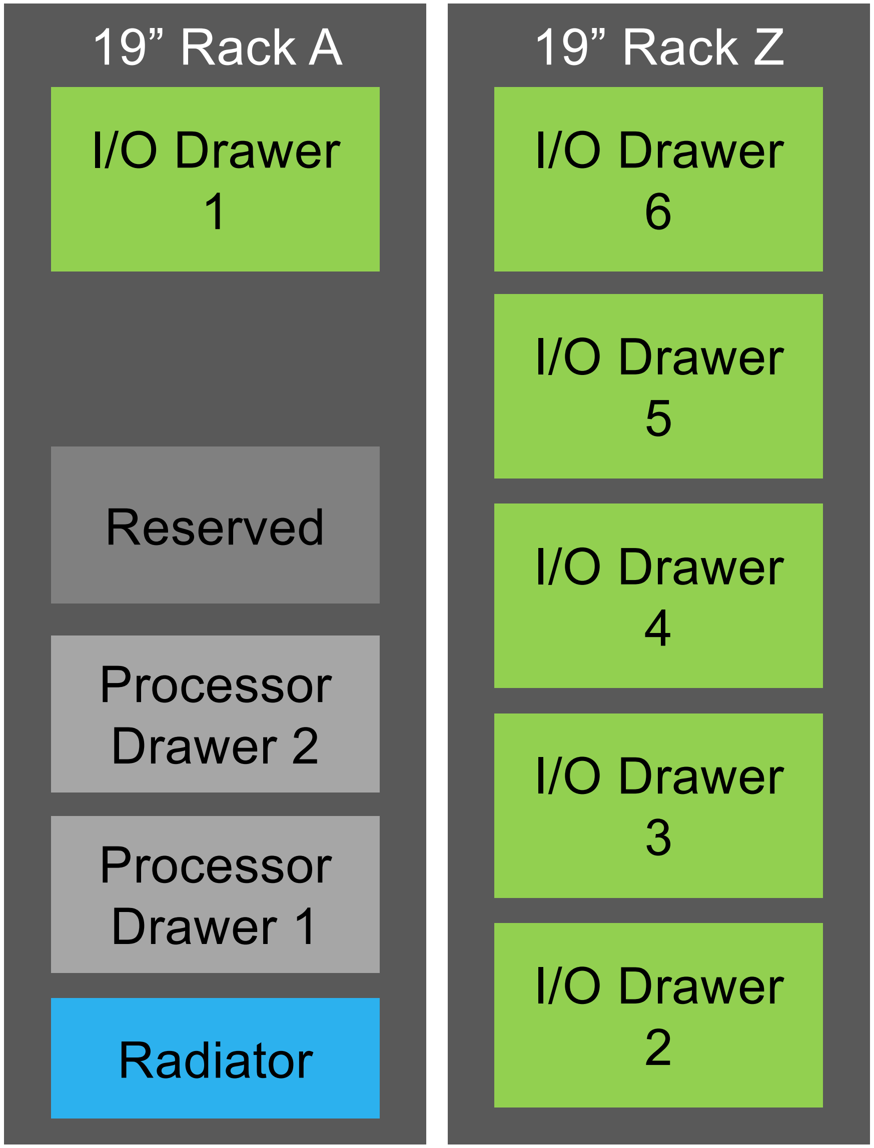

Topics

z15 T01 (LT1) Racks: Scenario 2 – iPDU Powered CPU Heavy

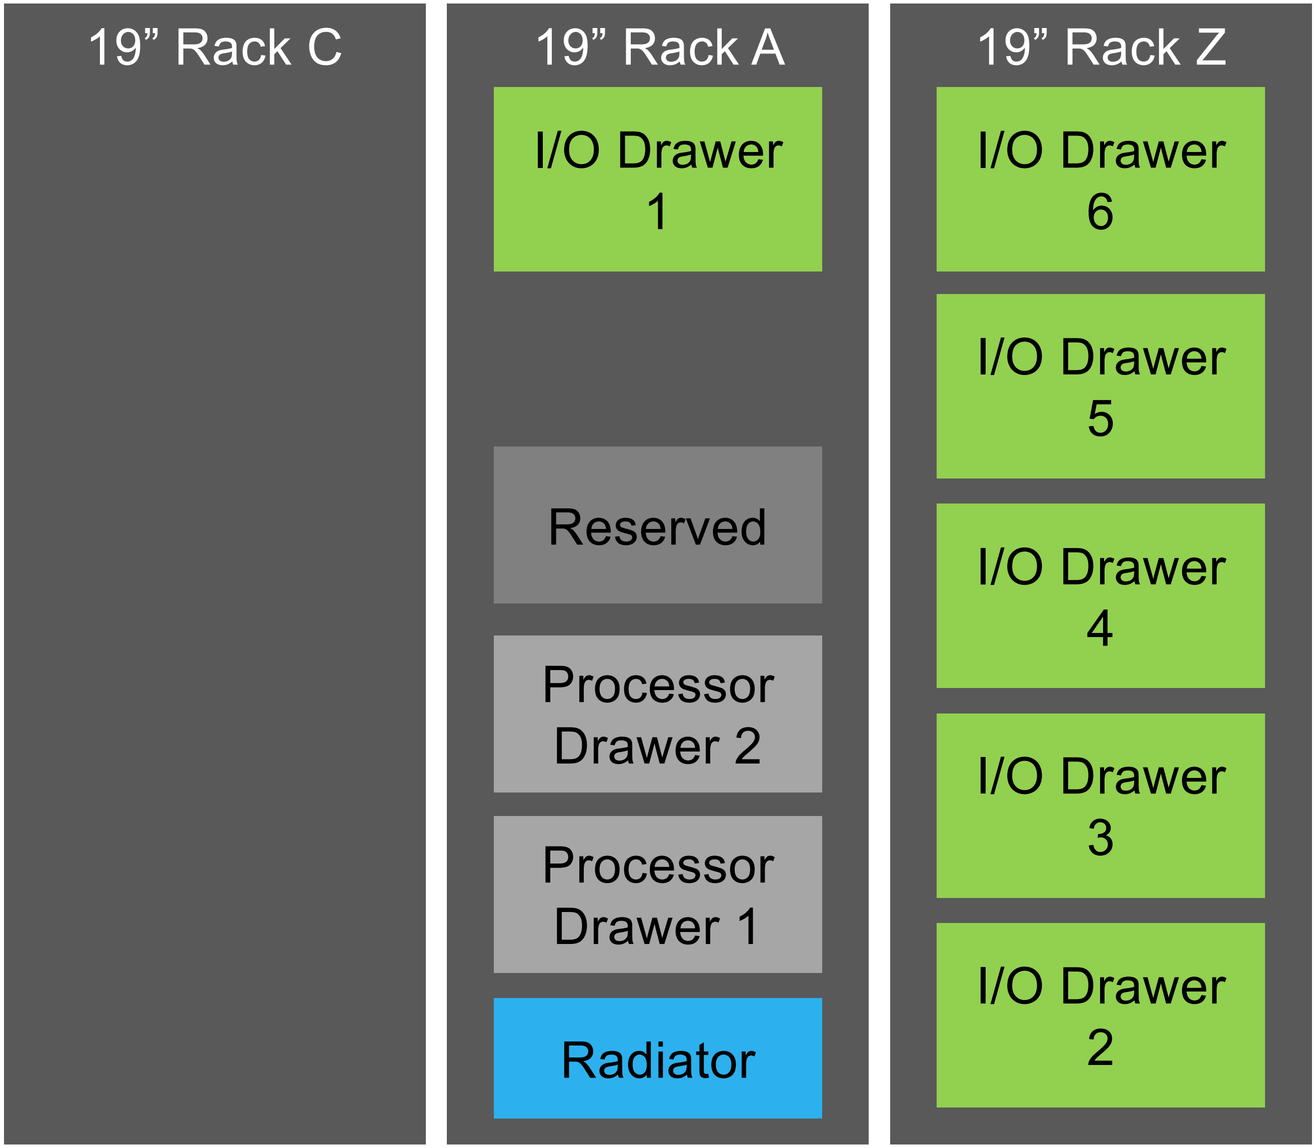

Topics

z15 T01 (LT1) Racks: Scenario 2 – iPDU Powered CPU Heavy

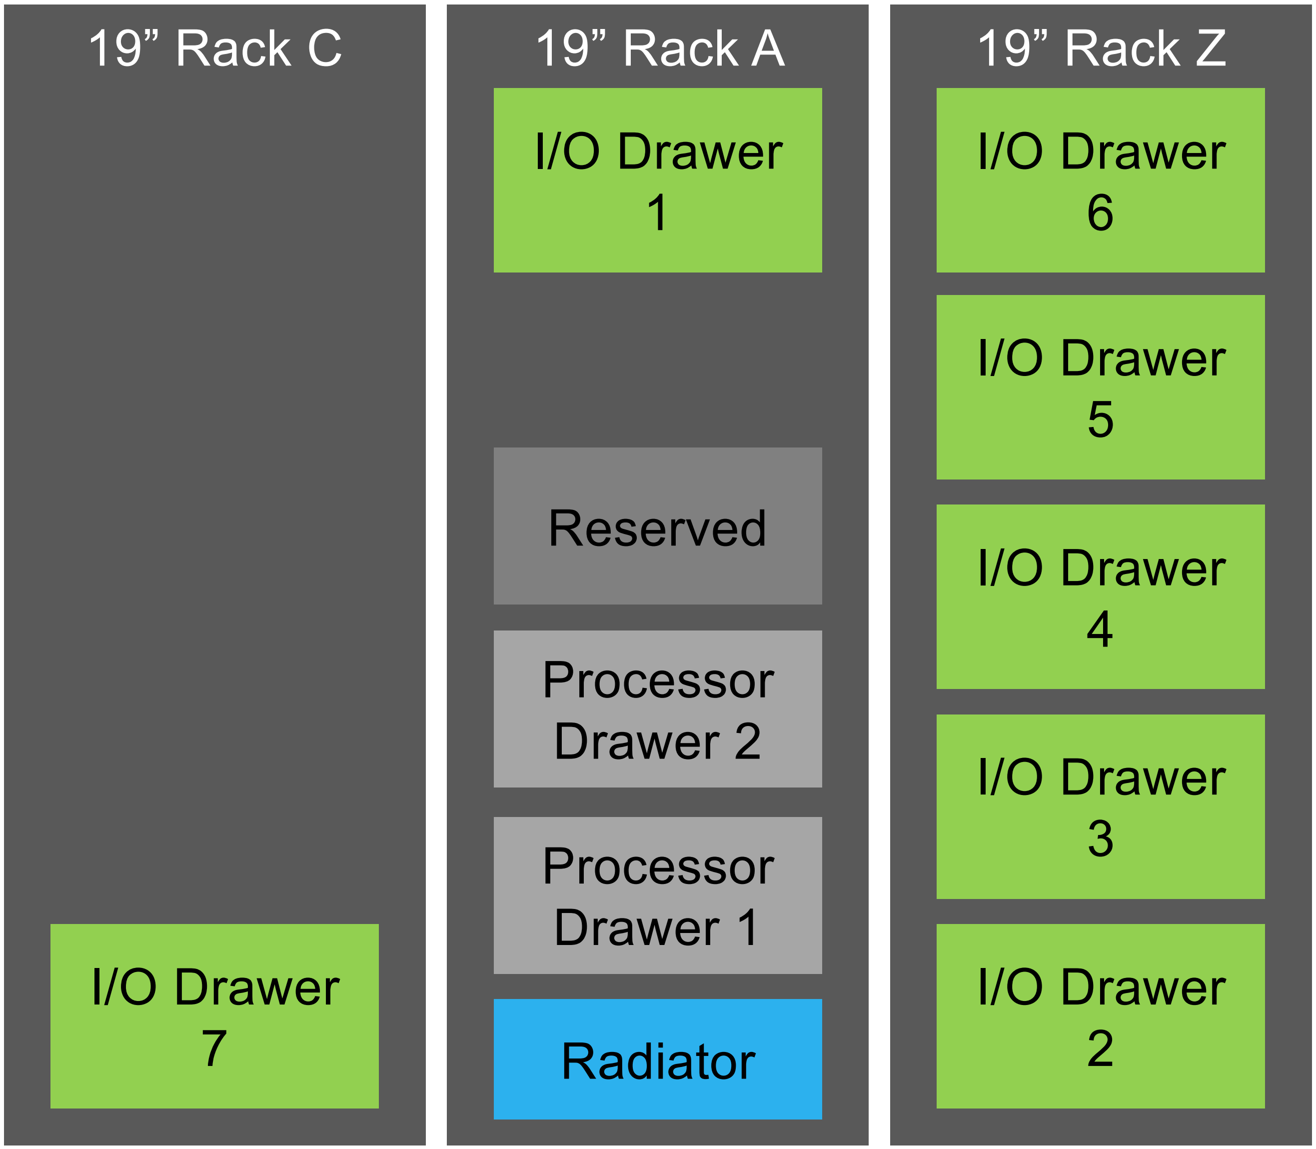

Topics

z15 T01 (LT1) Racks: Scenario 2 – iPDU Powered CPU Heavy

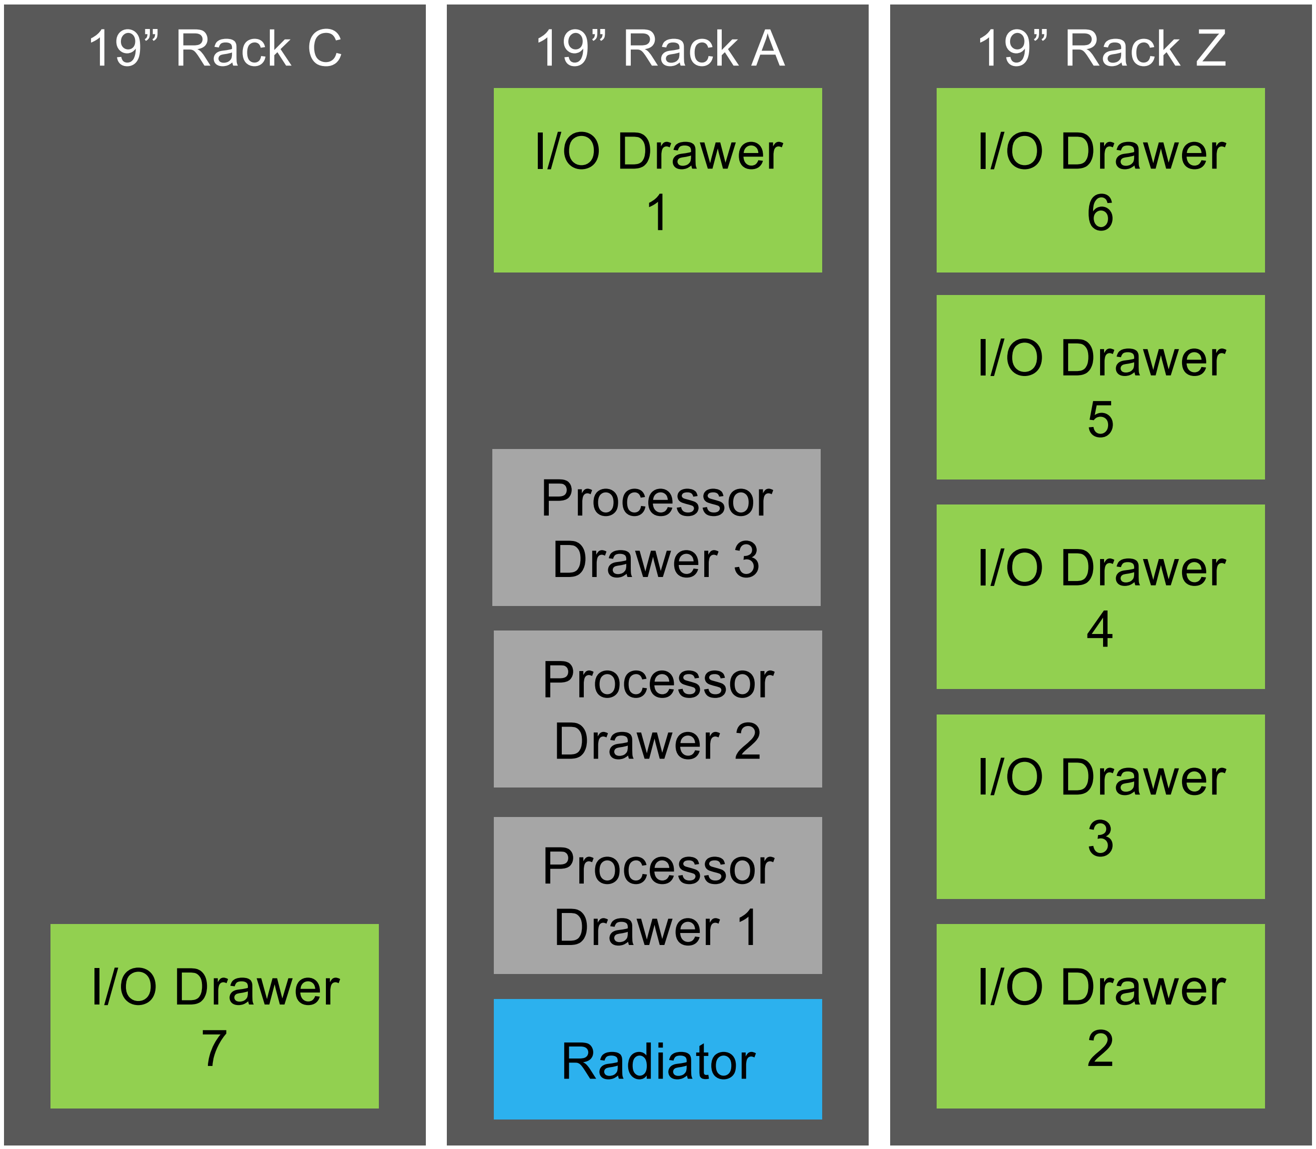

Topics

z15 T01 (LT1) Racks: Scenario 2 – iPDU Powered CPU Heavy

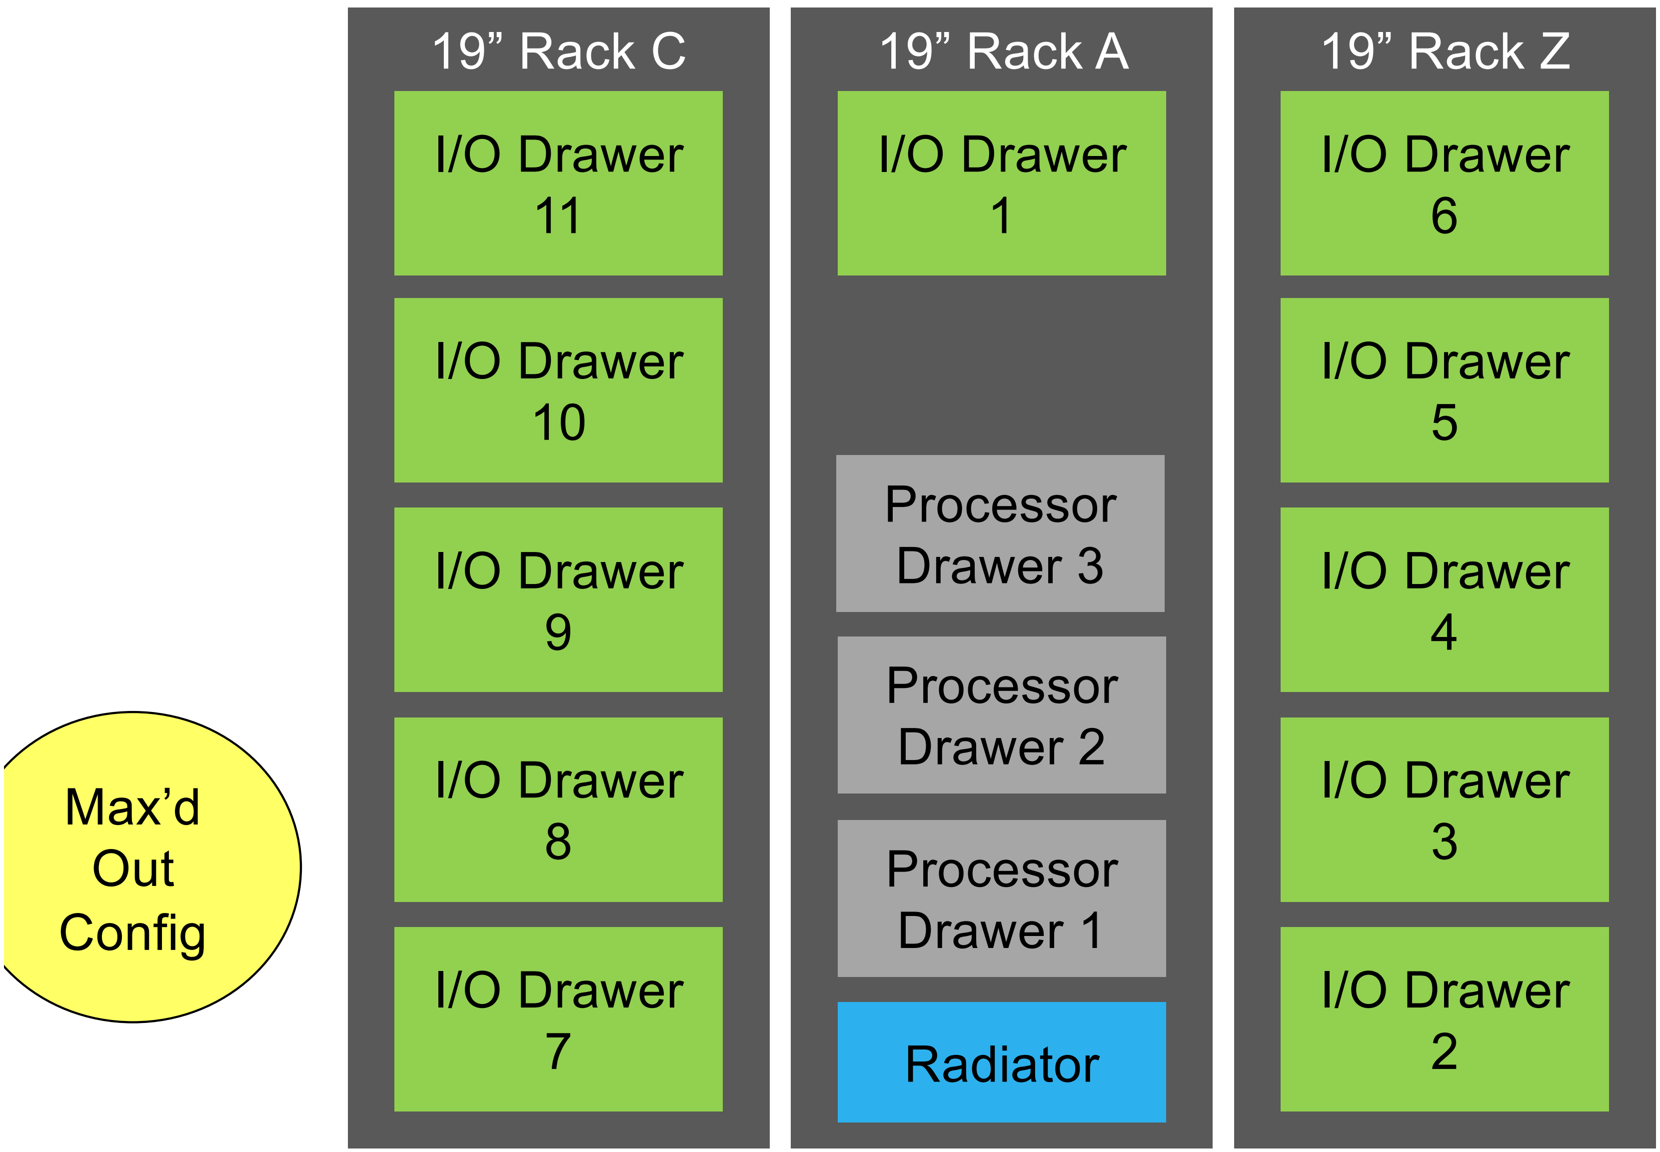

z15 T01 (LT1) Racks: Scenario 3 – iPDU Powered Heavy Heavy

Systems with 4 or 5 Drawers Must be Factory Built

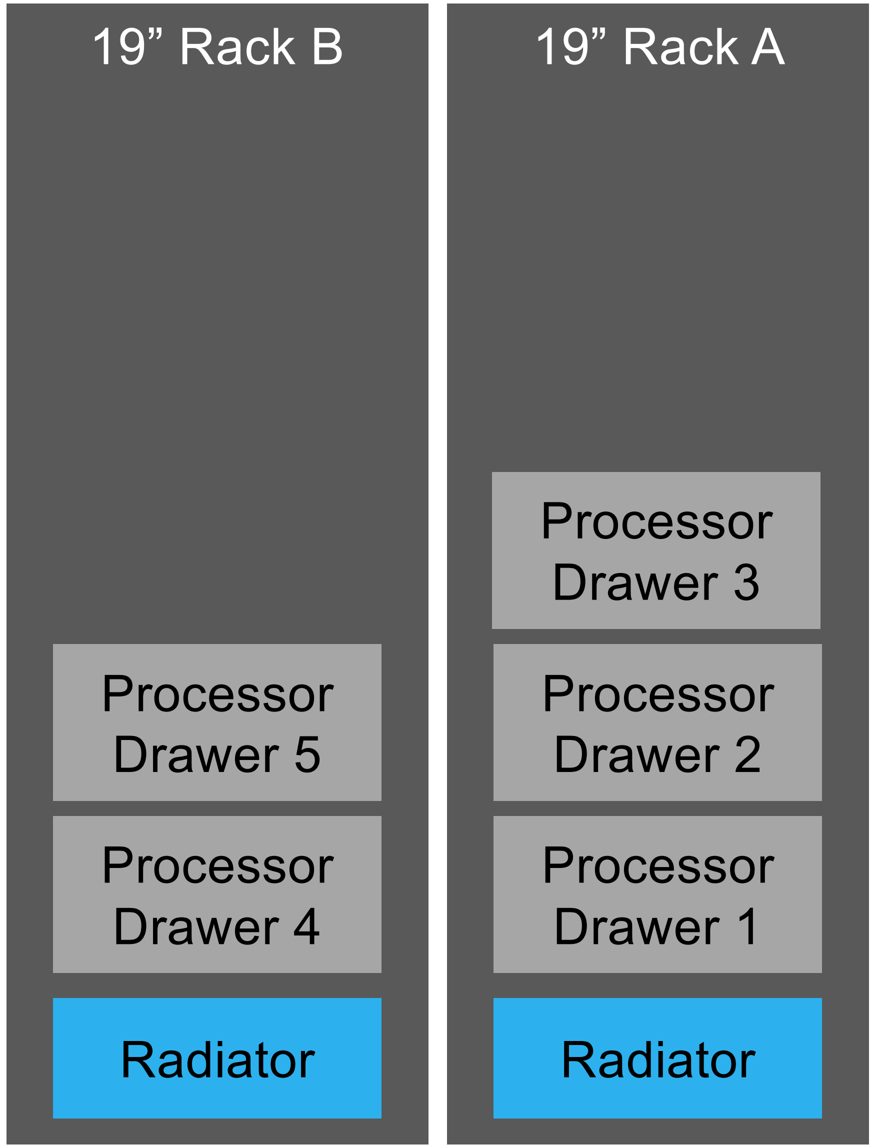

z15 T01 (LT1) Racks: Scenario 3 – iPDU Powered Heavy Heavy

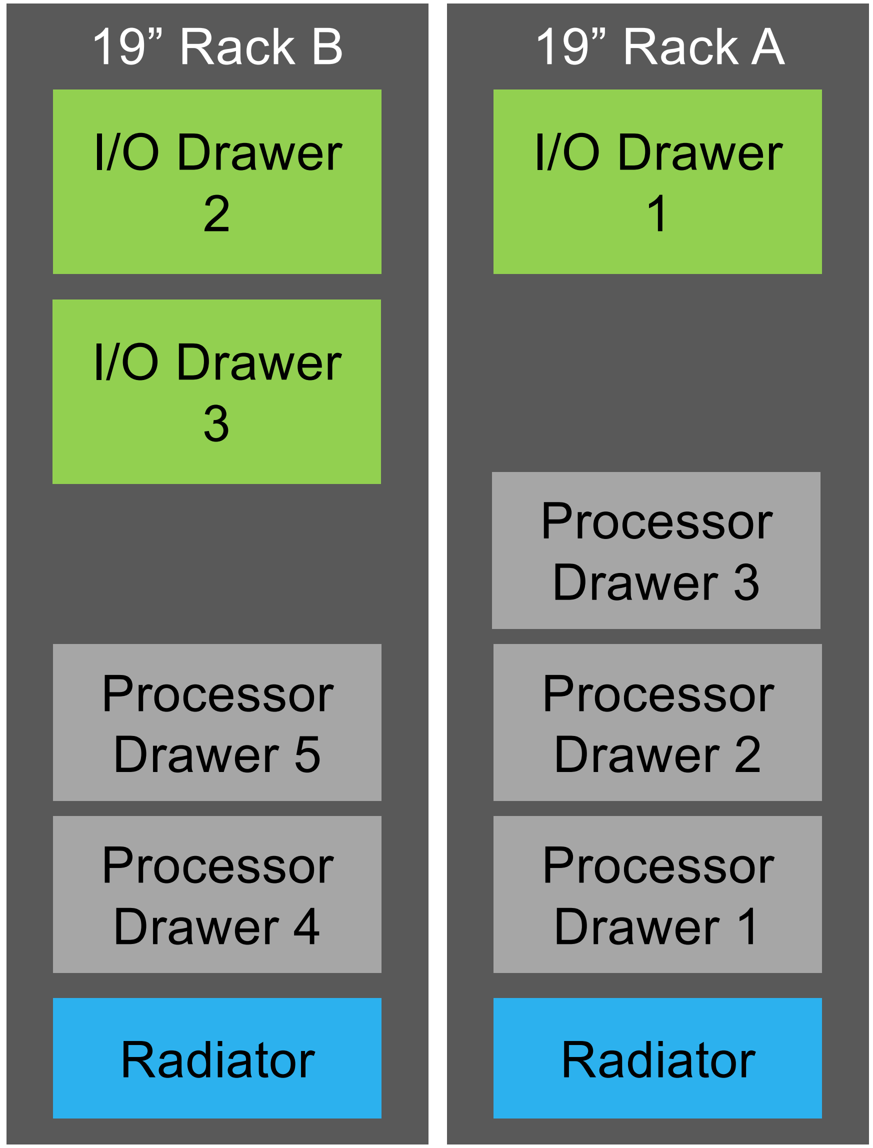

Everything that happens next CAN BE non-disruptive

Topics

z15 T01 (LT1) Racks: Scenario 3 – iPDU Powered Heavy Heavy

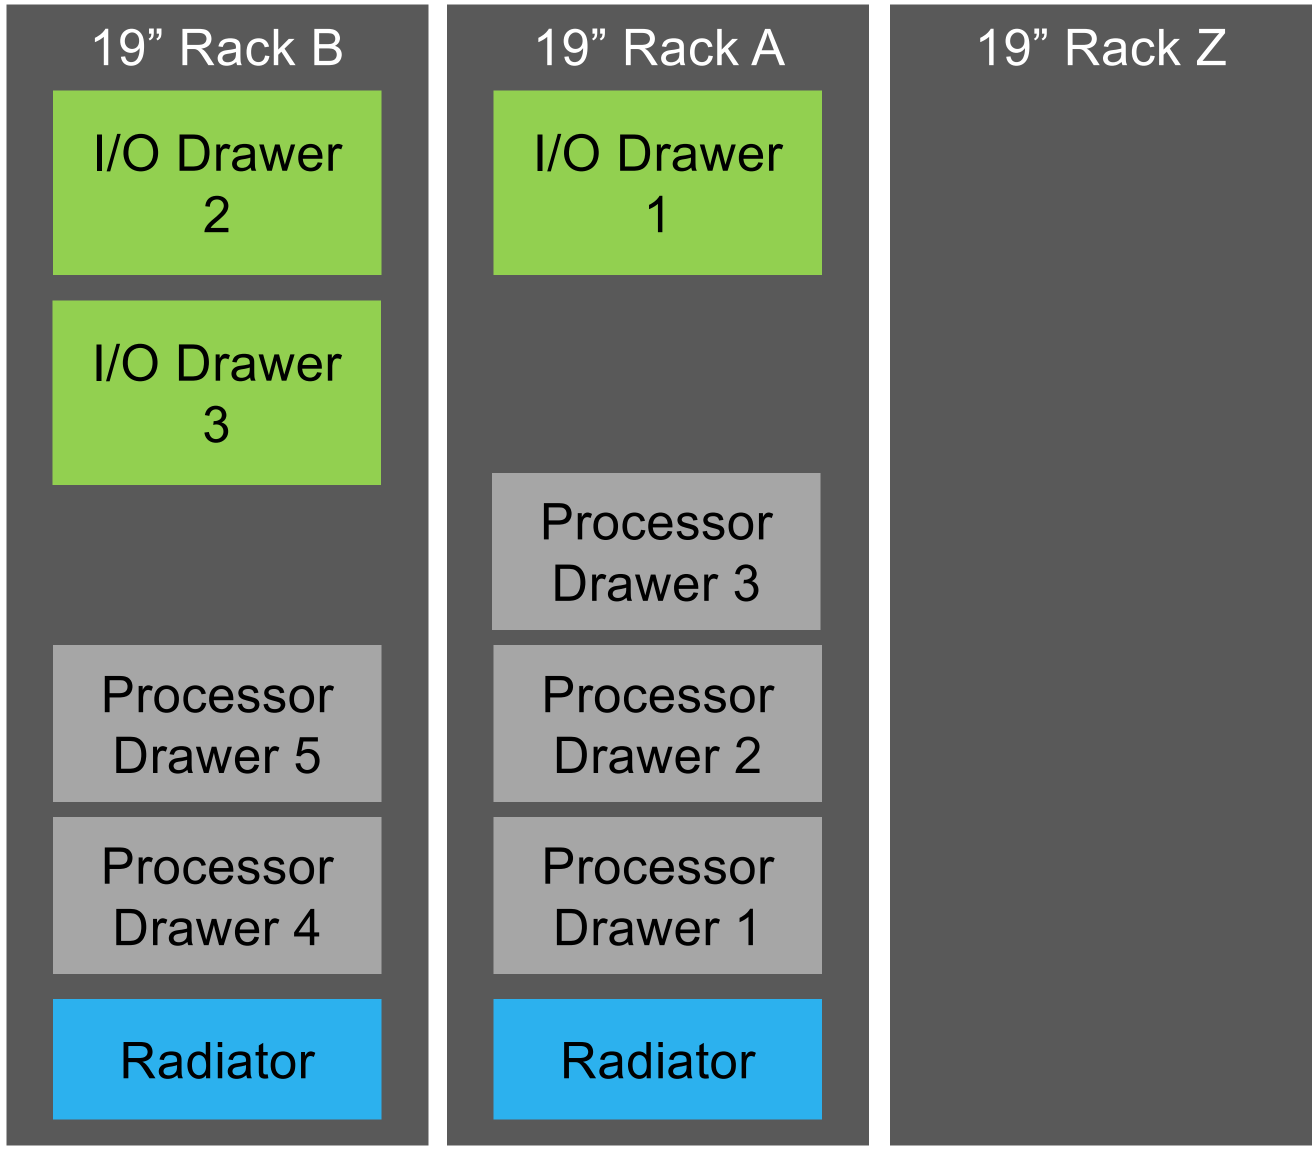

z15 T01 (LT1) Racks: Scenario 3 – iPDU Powered Heavy Heavy

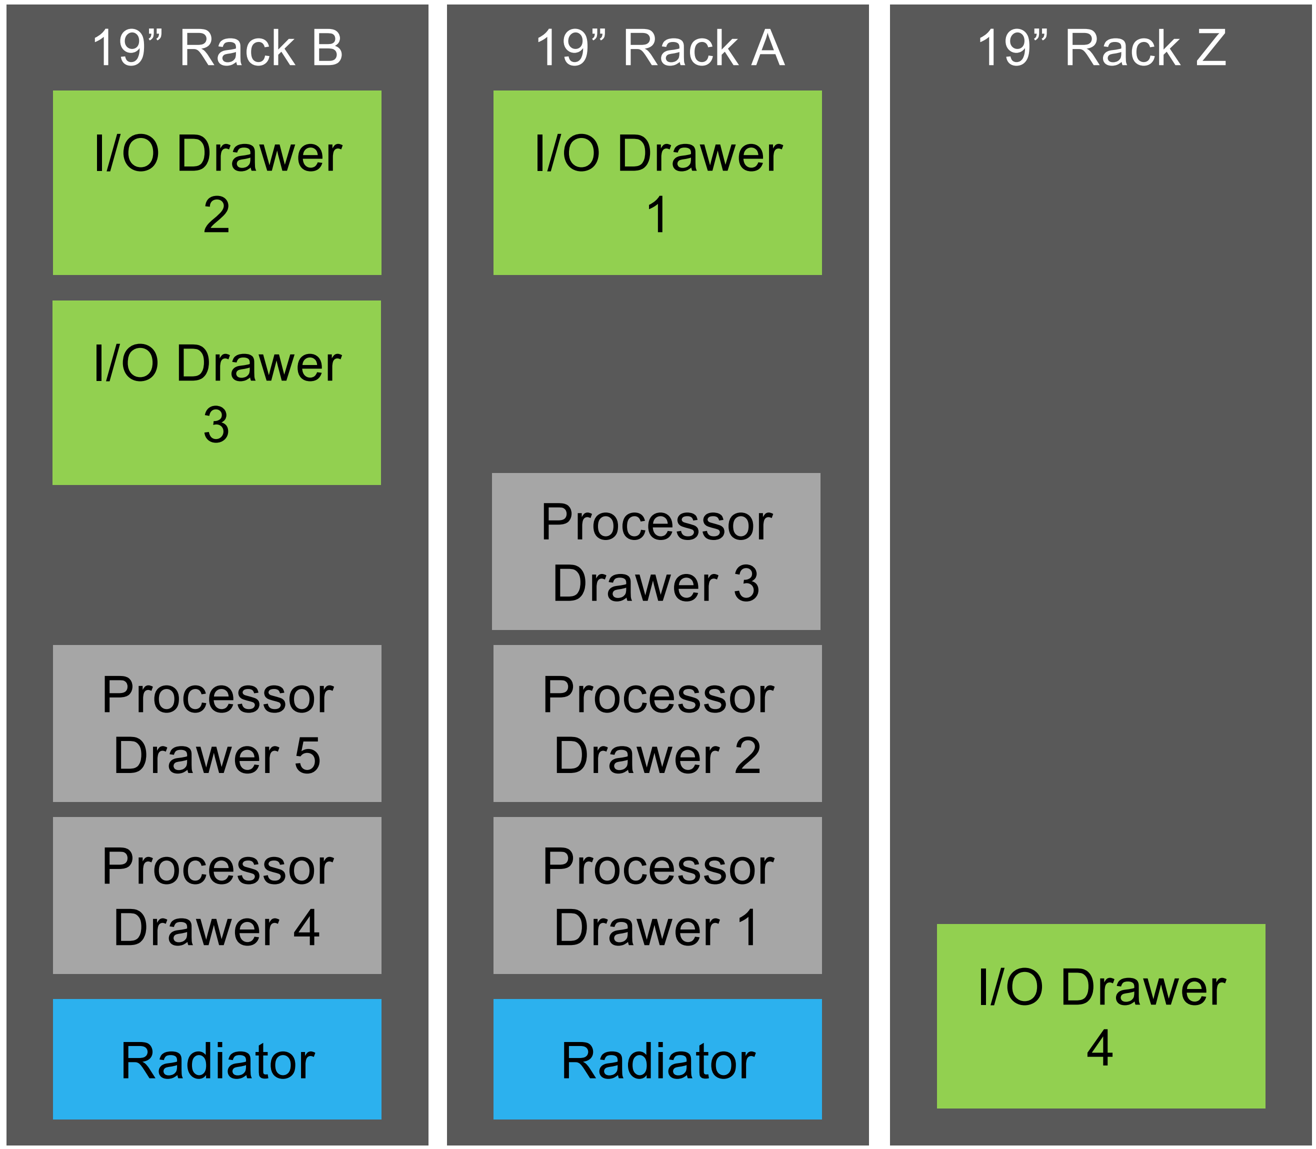

Topics

z15 T01 (LT1) Racks: Scenario 3 – iPDU Powered Heavy Heavy

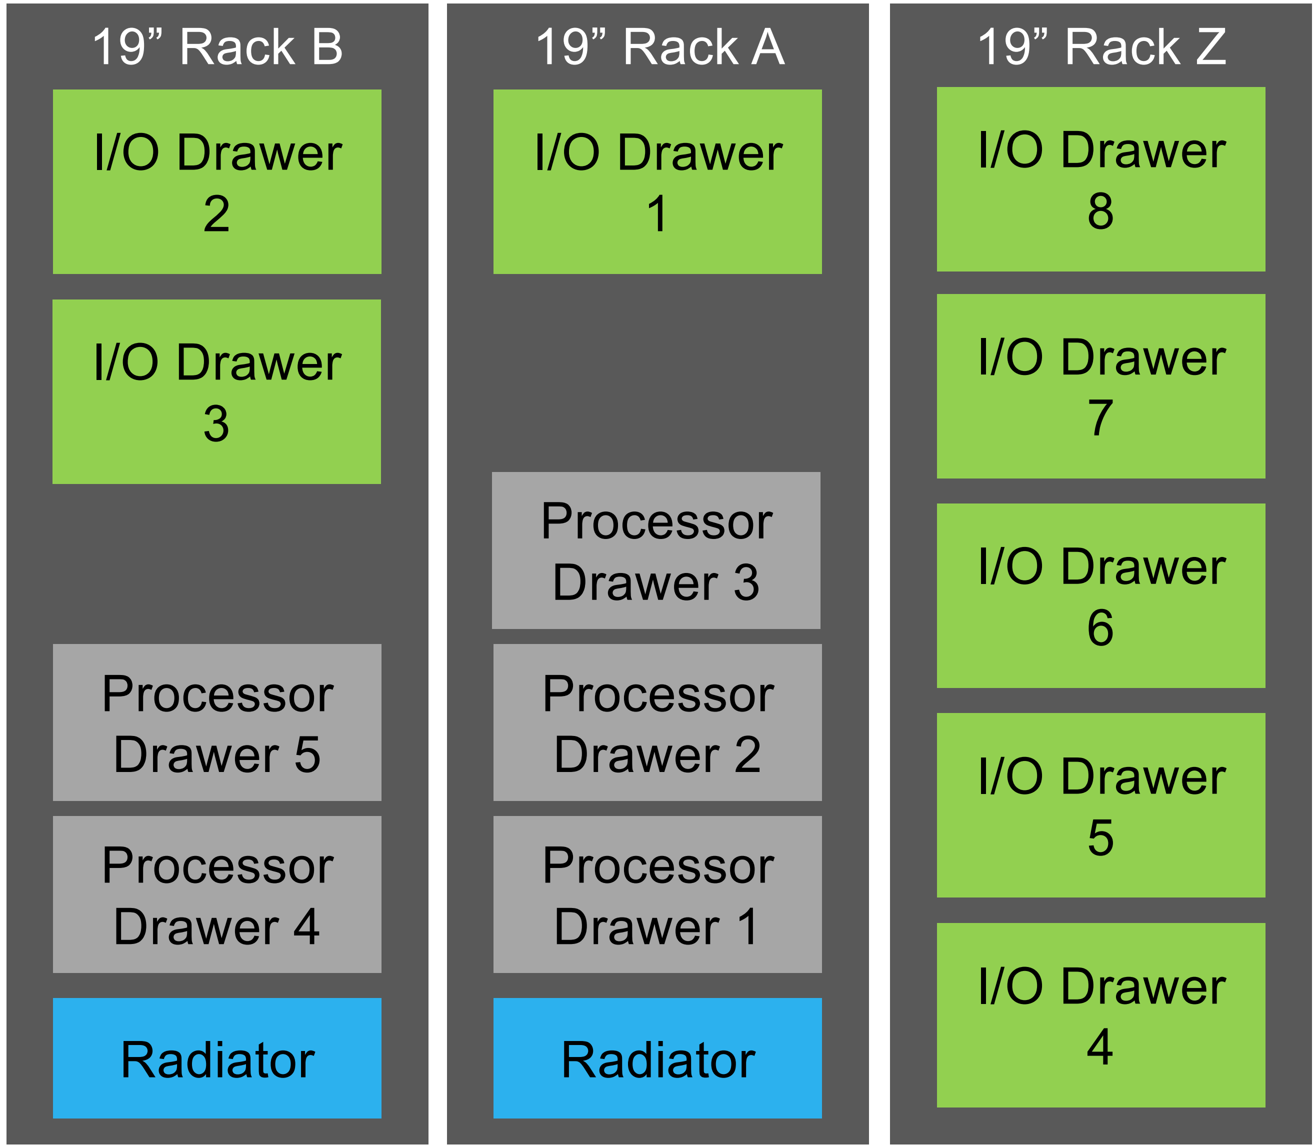

Topics

z15 T01 (LT1) Racks: Scenario 3 – iPDU Powered Heavy Heavy

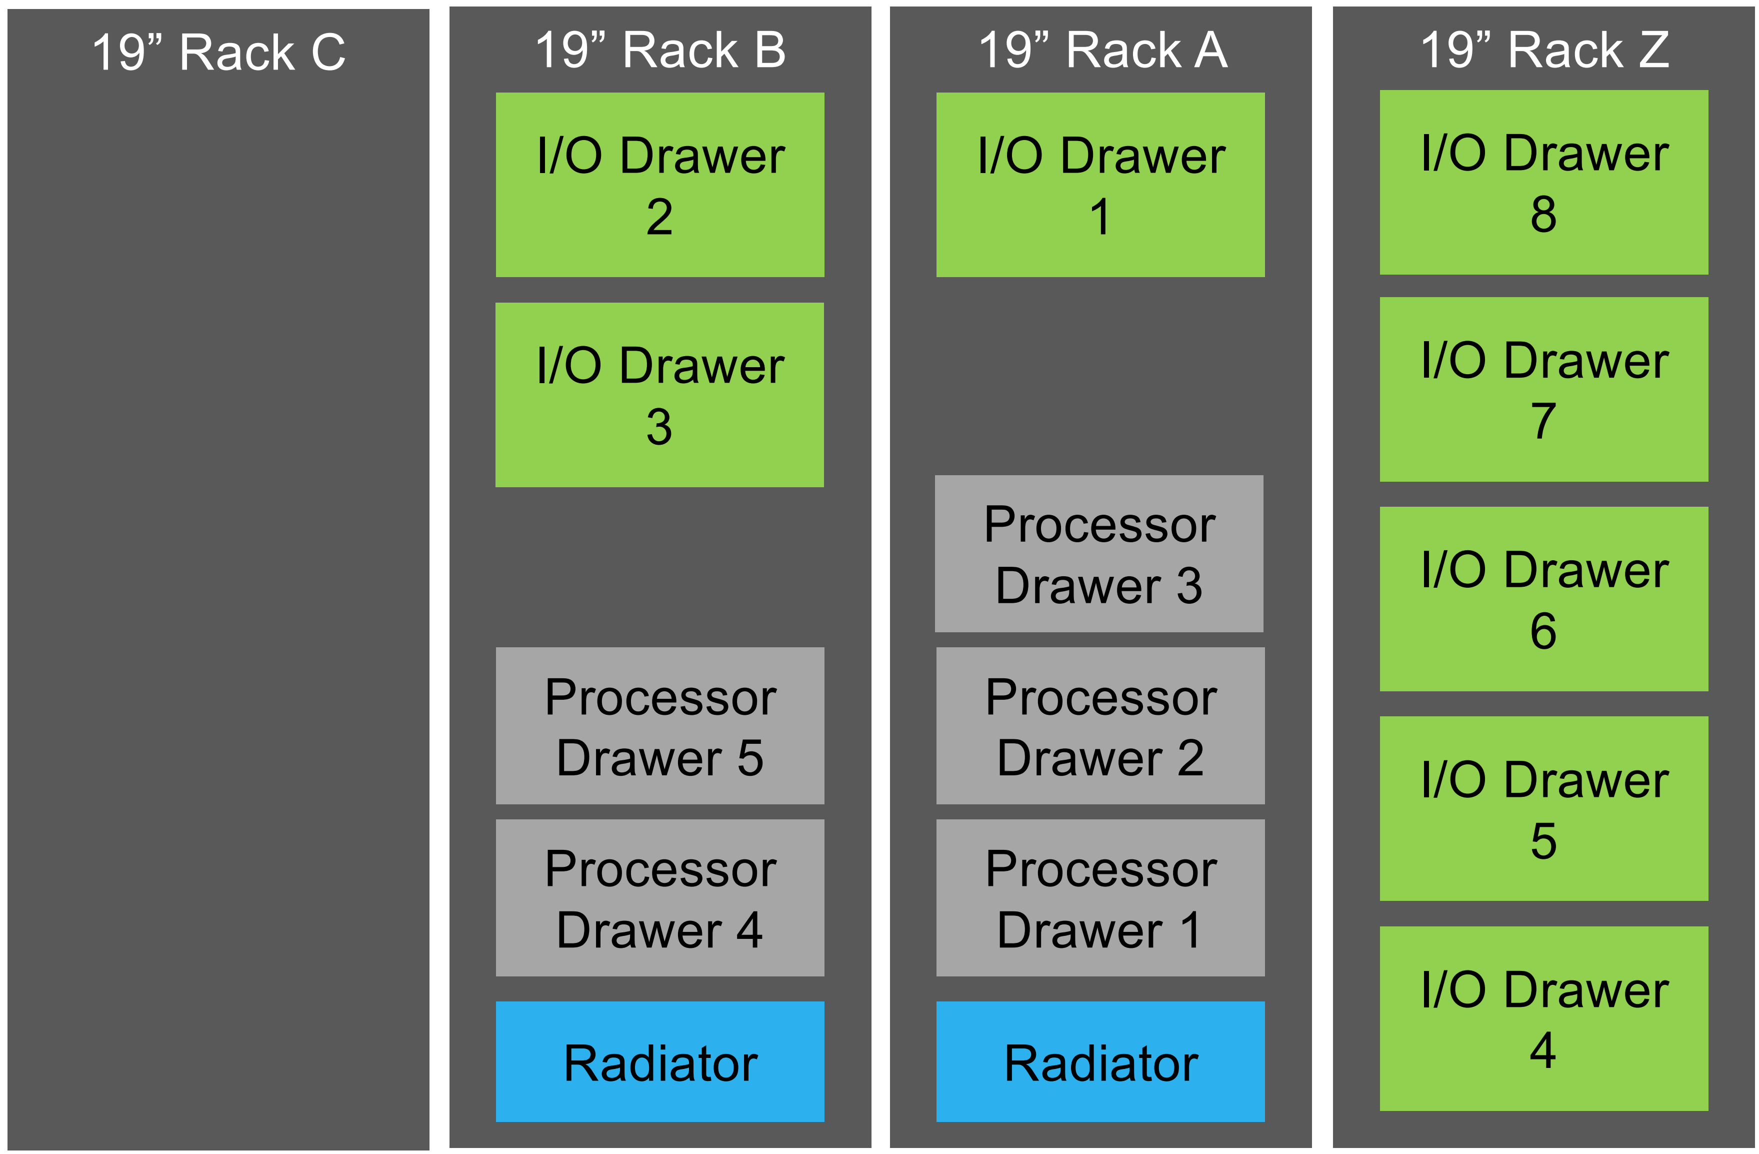

z15 T01 (LT1) Racks: Scenario 3 – iPDU Powered Heavy Heavy

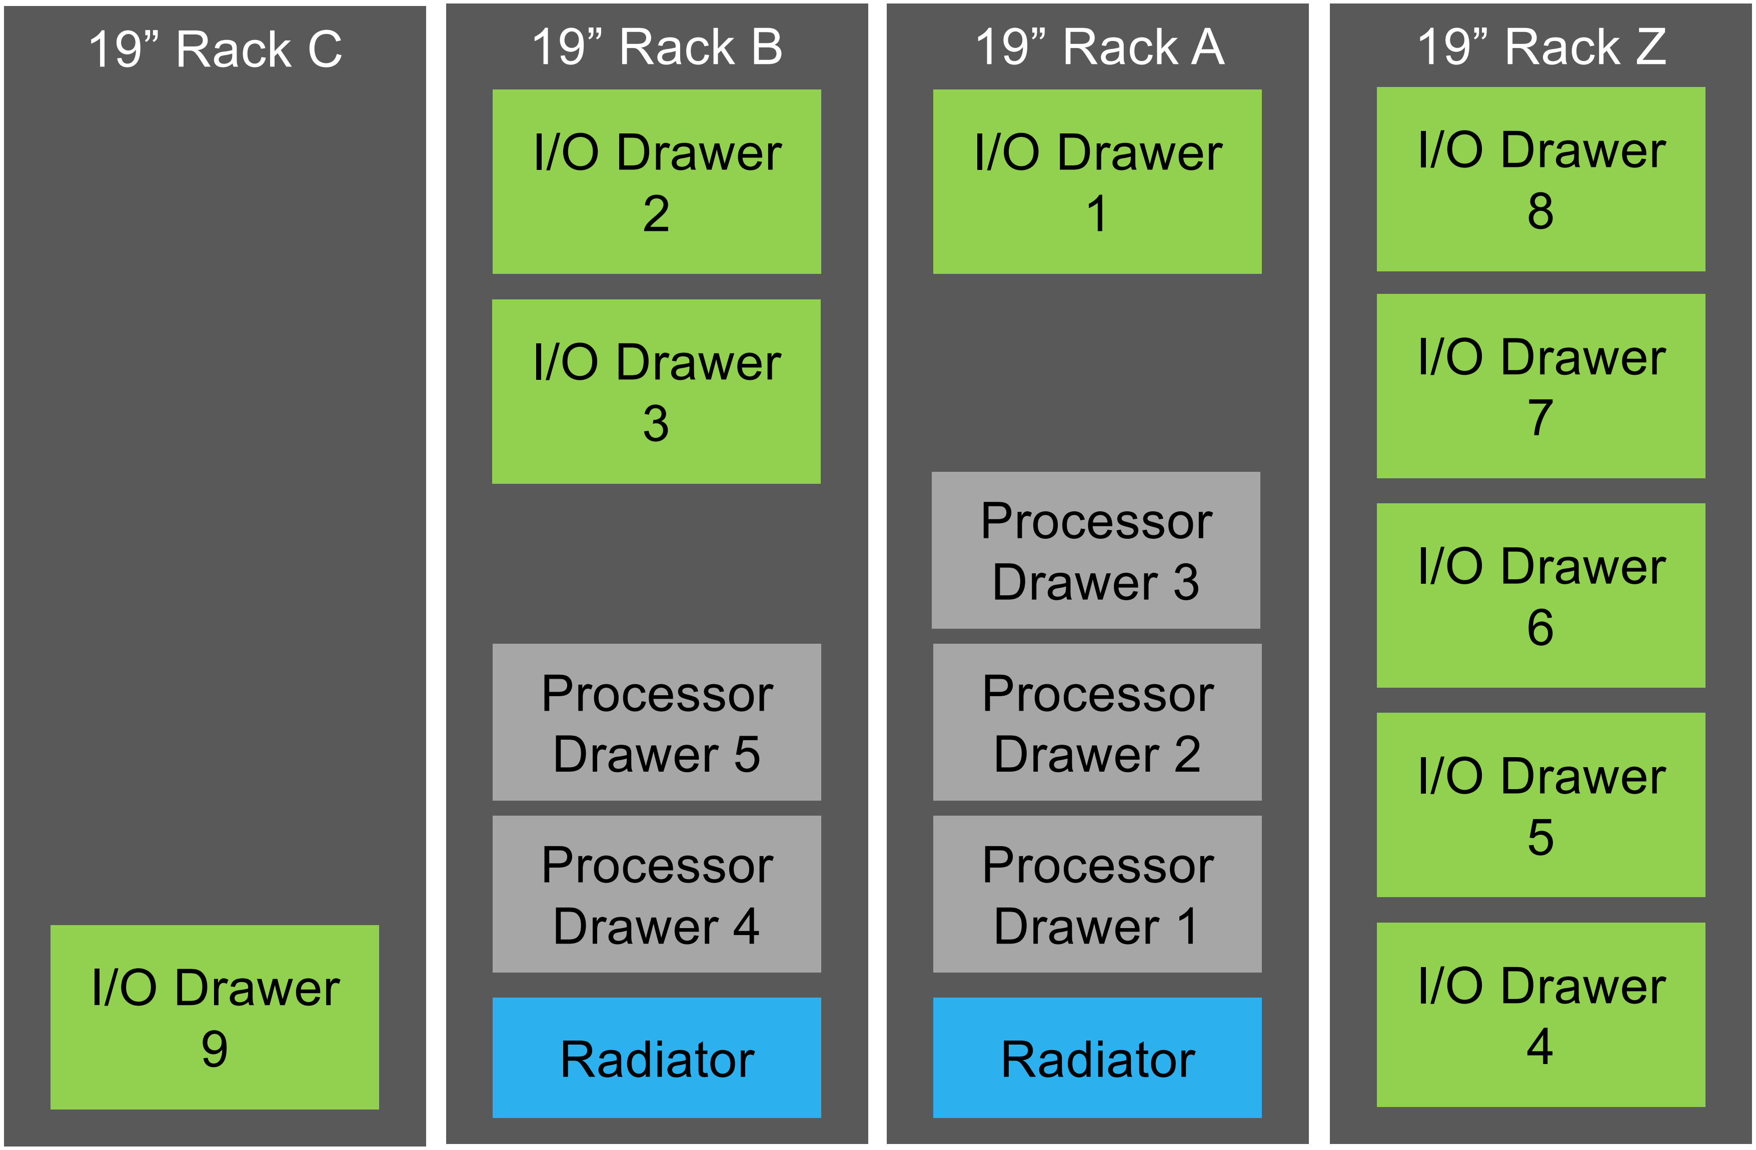

Topics

z15 T01 (LT1) Racks: Scenario 3 – iPDU Powered Heavy Heavy

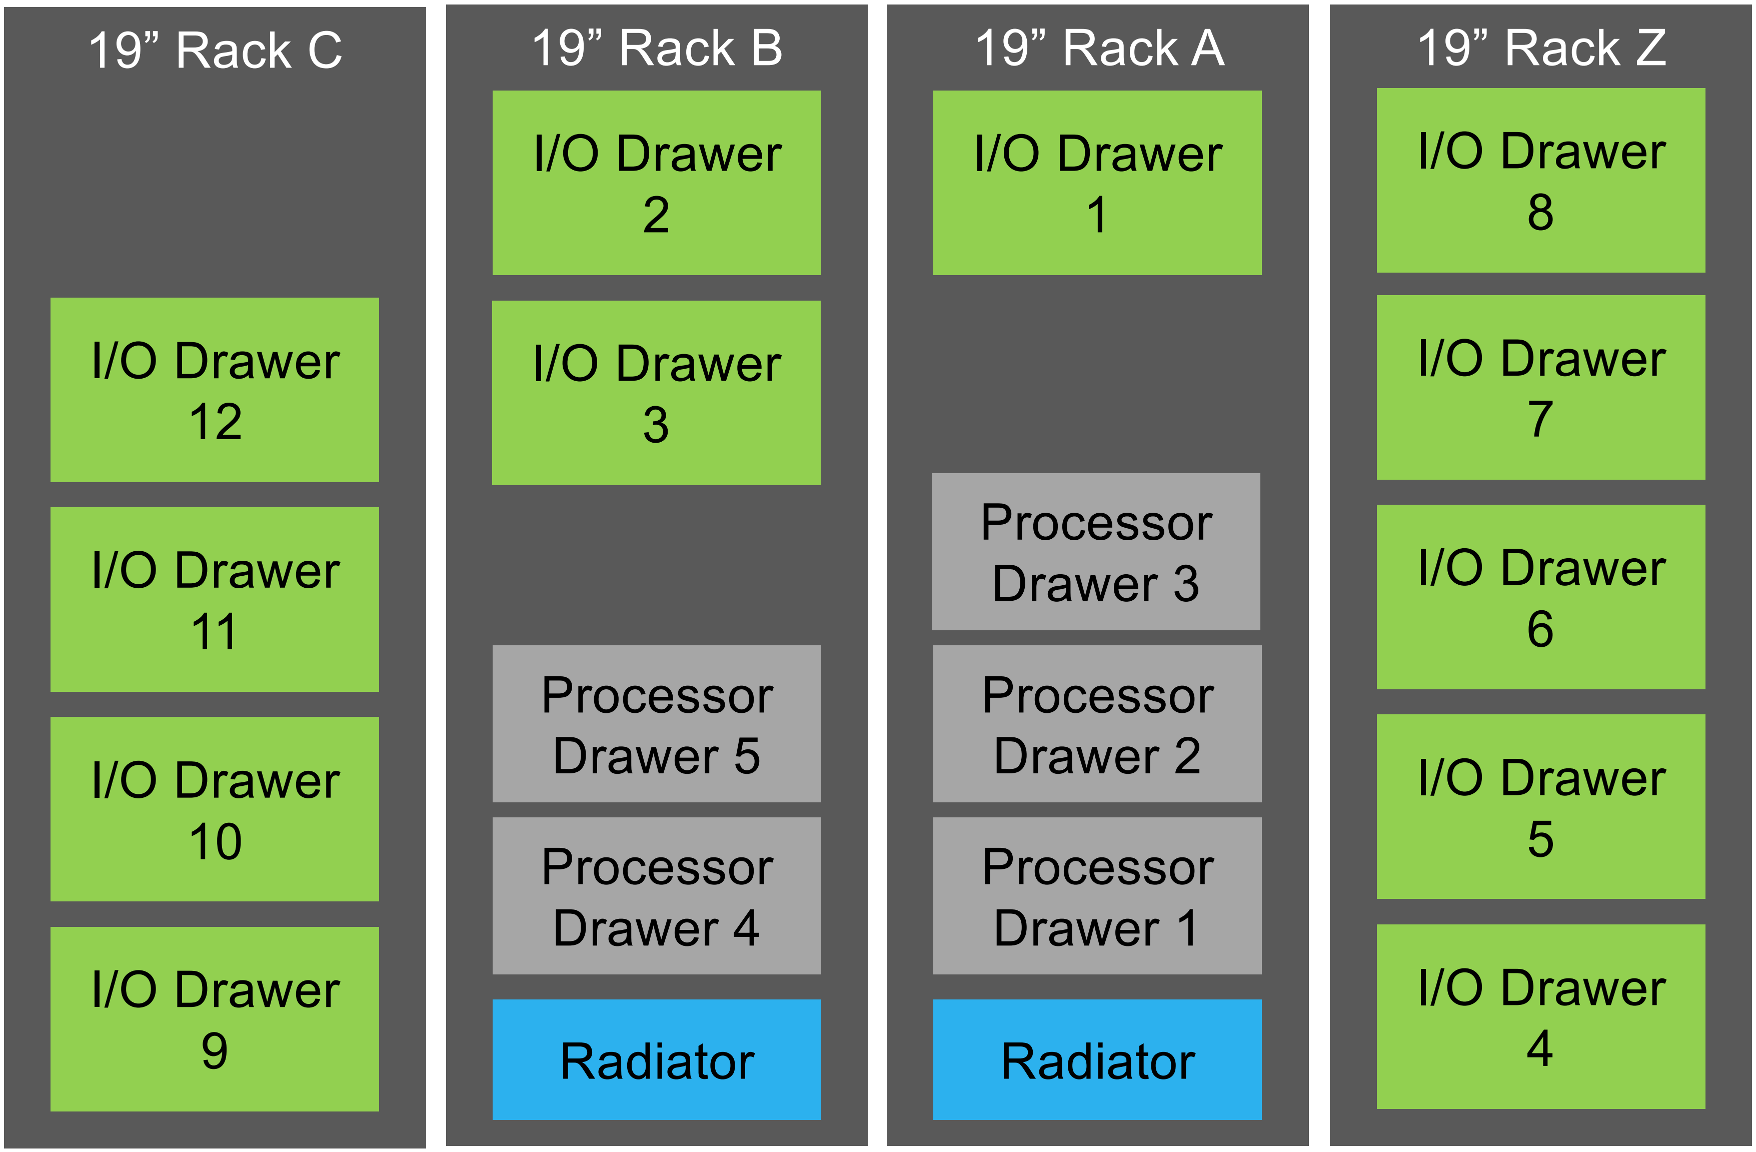

Topics

z15 T01 (LT1) Racks: Scenario 4 – Bulk Power Assembly Powered Heavy Heavy

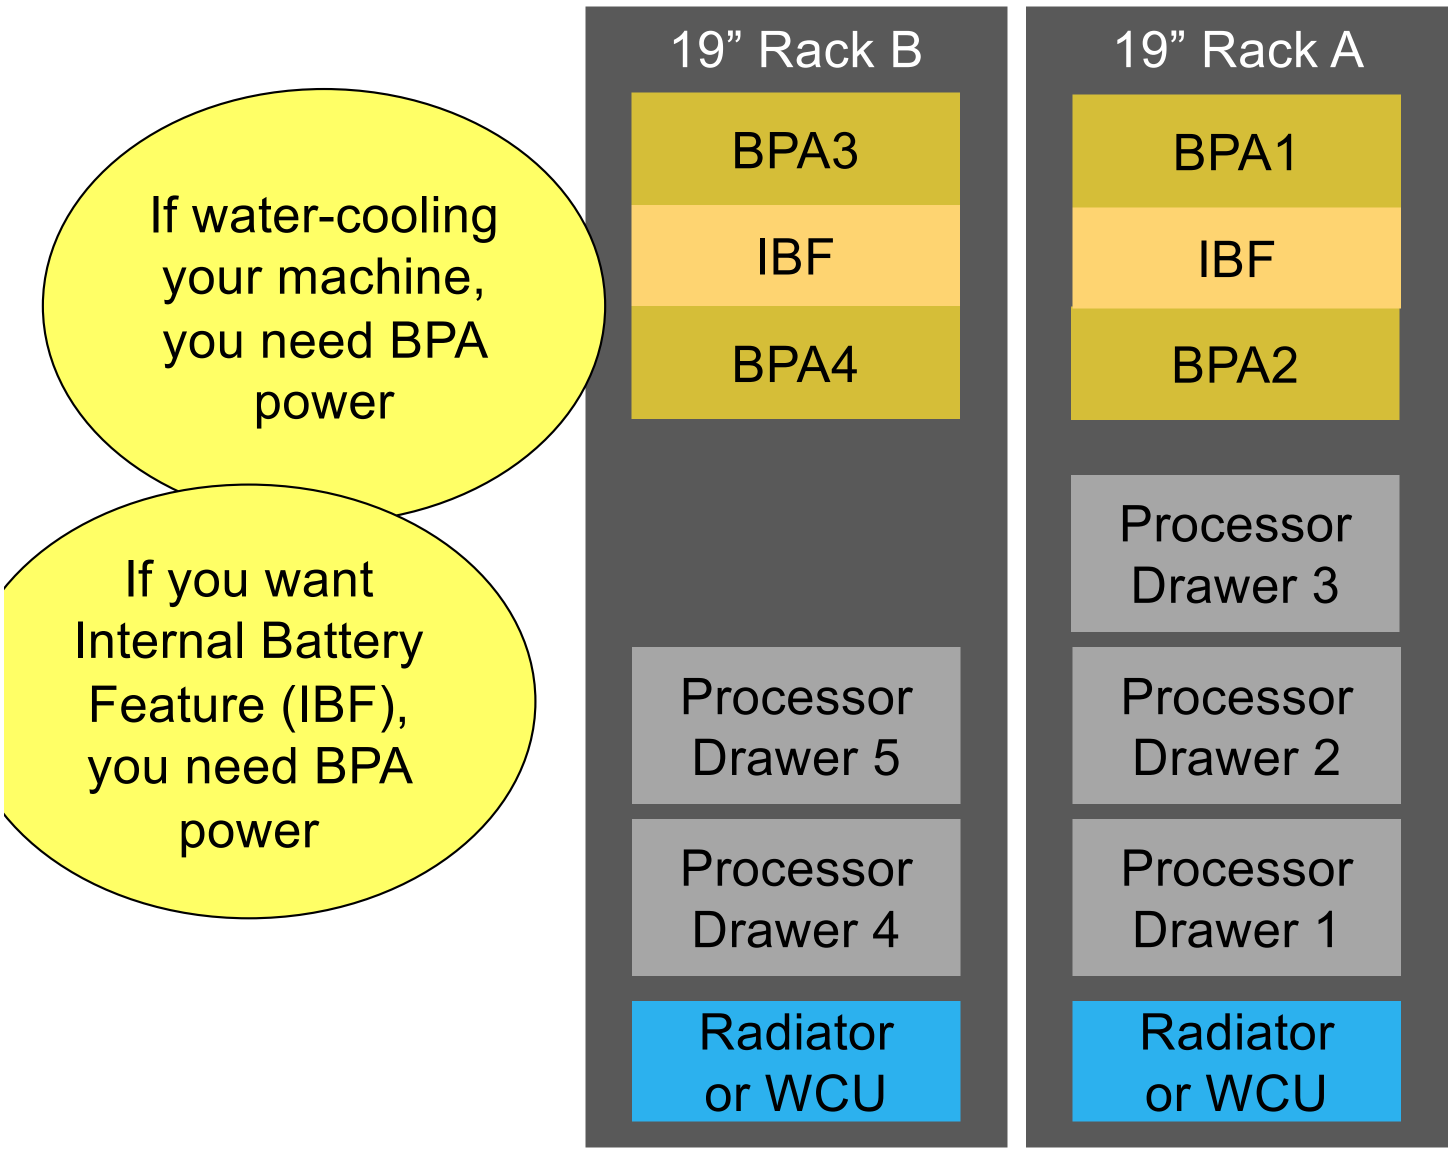

Topics

z15 T01 (LT1) Racks: Scenario 4 – Bulk Power Assembly Powered Heavy Heavy

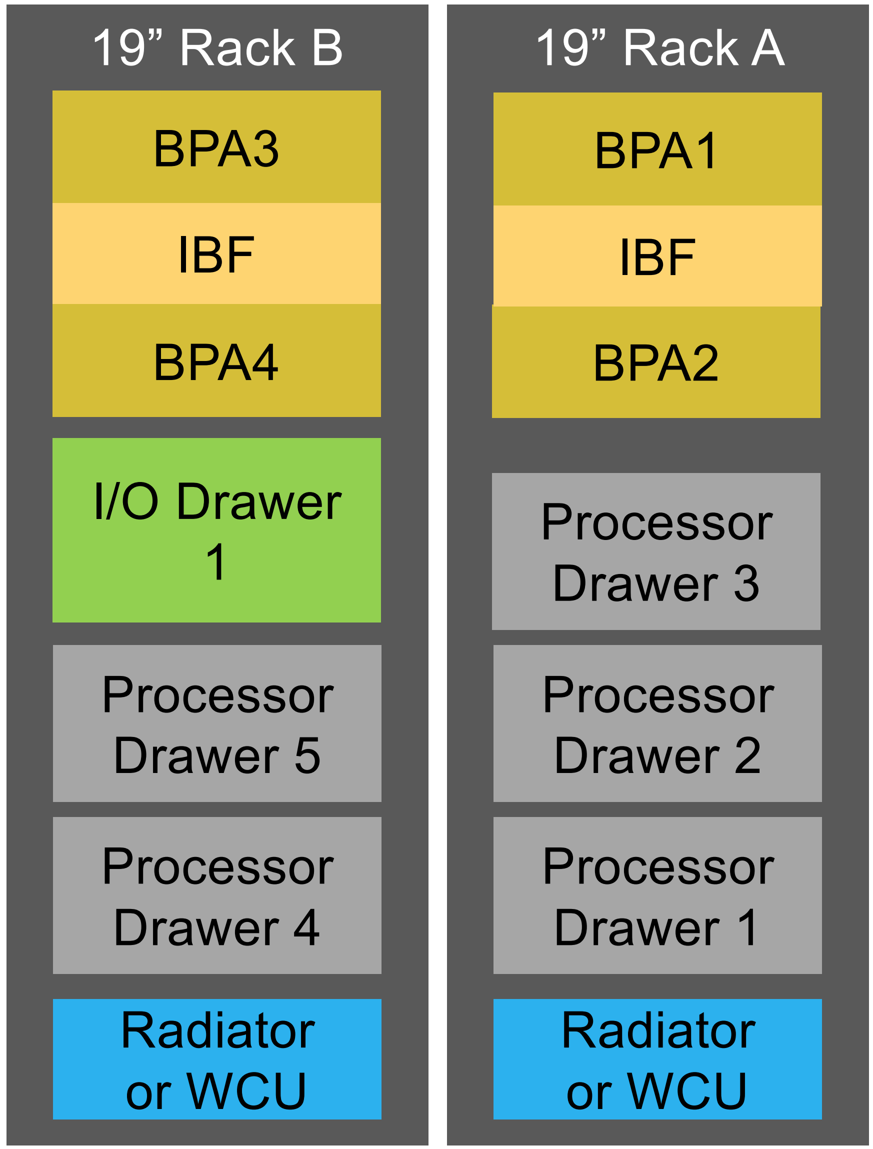

Topics

z15 T01 (LT1) Racks: Scenario 4 – Bulk Power Assembly Powered Heavy Heavy

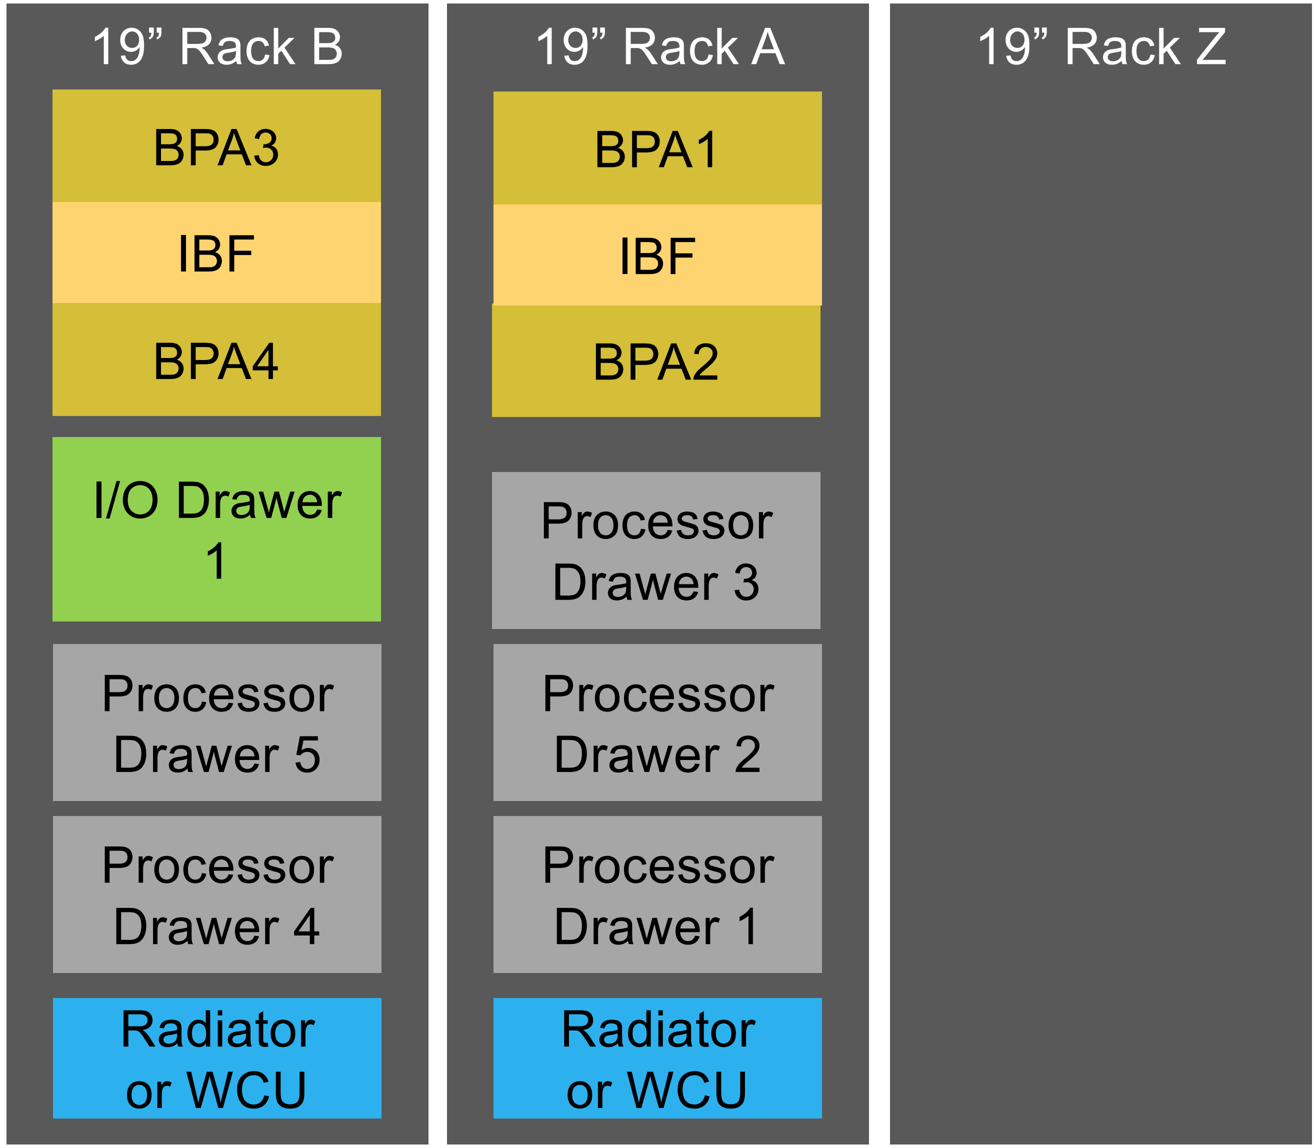

Topics

z15 T01 (LT1) Racks: Scenario 4 – Bulk Power Assembly Powered Heavy Heavy

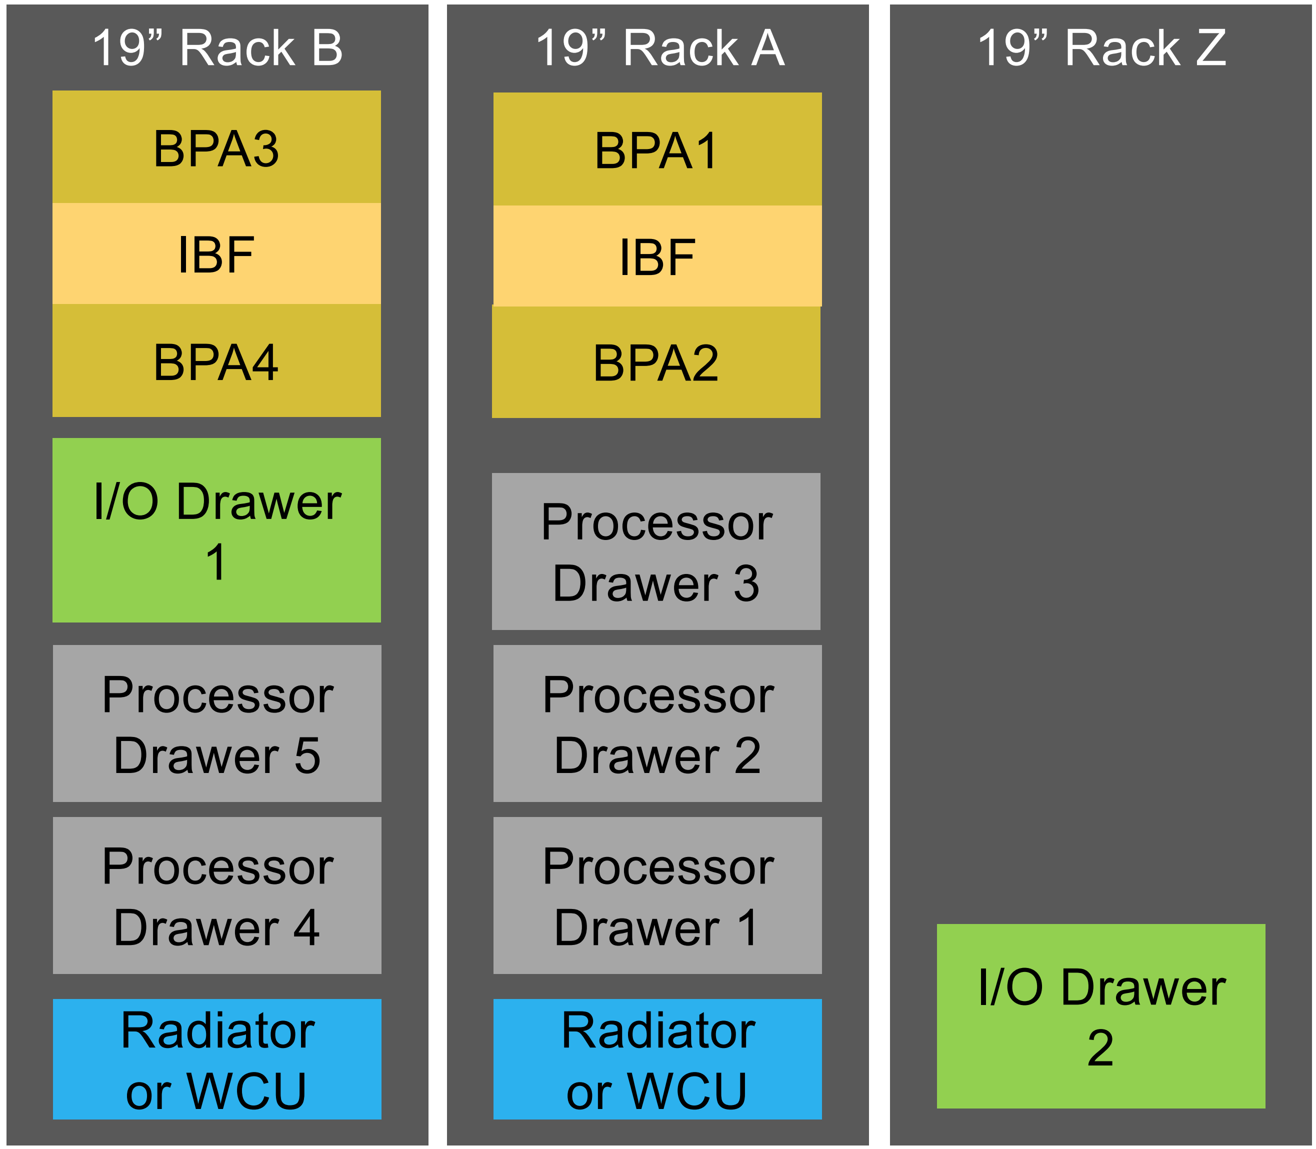

z15 T01 (LT1) Racks: Scenario 4 – Bulk Power Assembly Powered Heavy Heavy

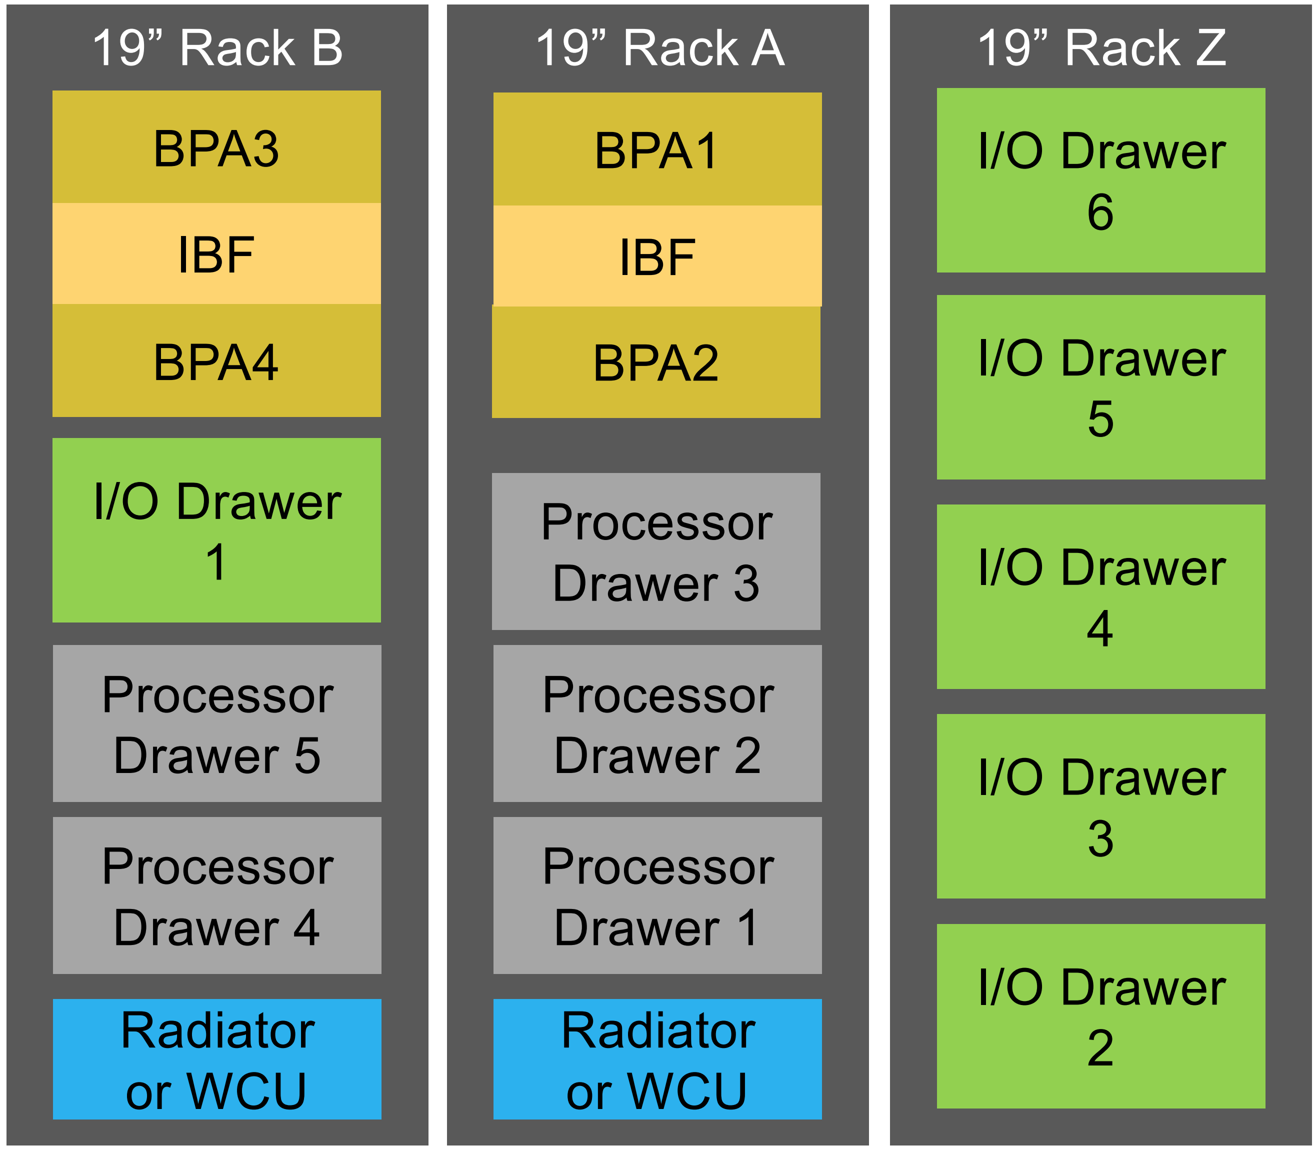

Topics

z15 T01 (LT1) Racks: Scenario 4 – Bulk Power Assembly Powered Heavy Heavy

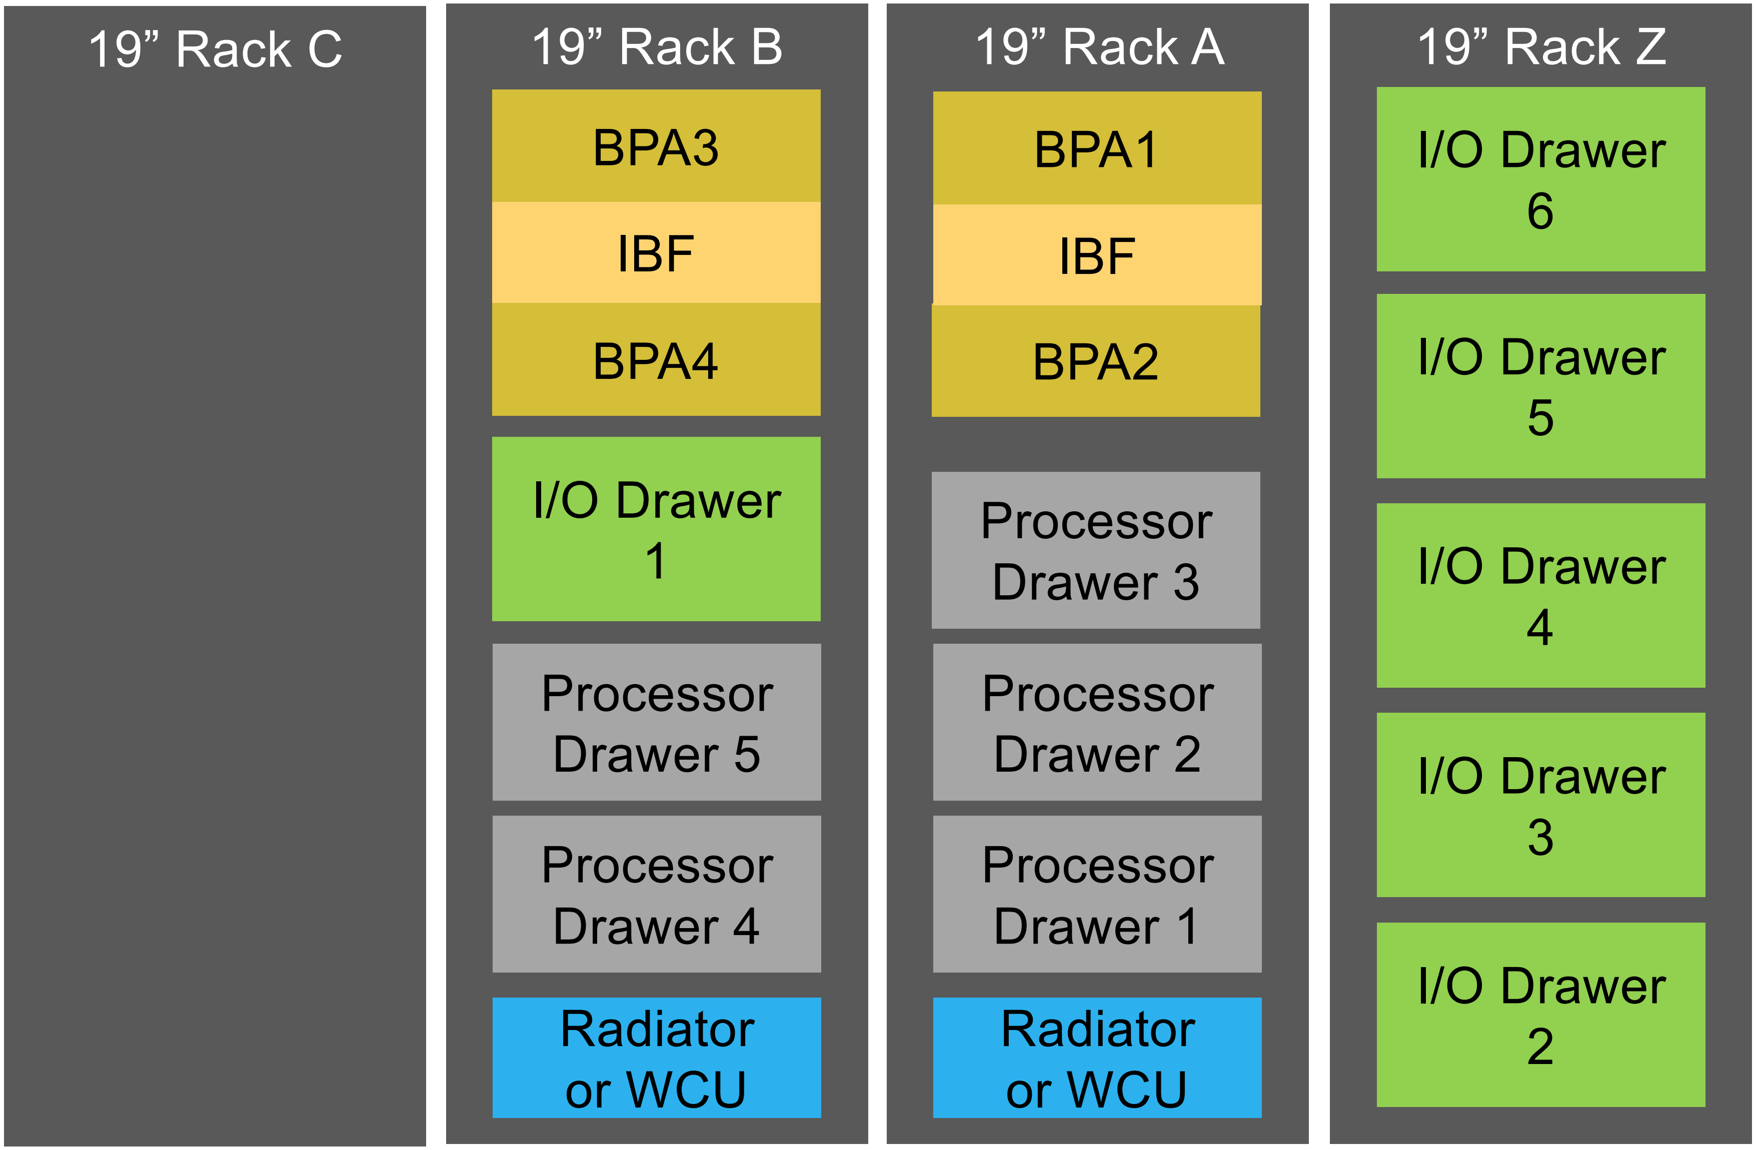

Topics

z15 T01 (LT1) Racks: Scenario 4 – Bulk Power Assembly Powered Heavy Heavy

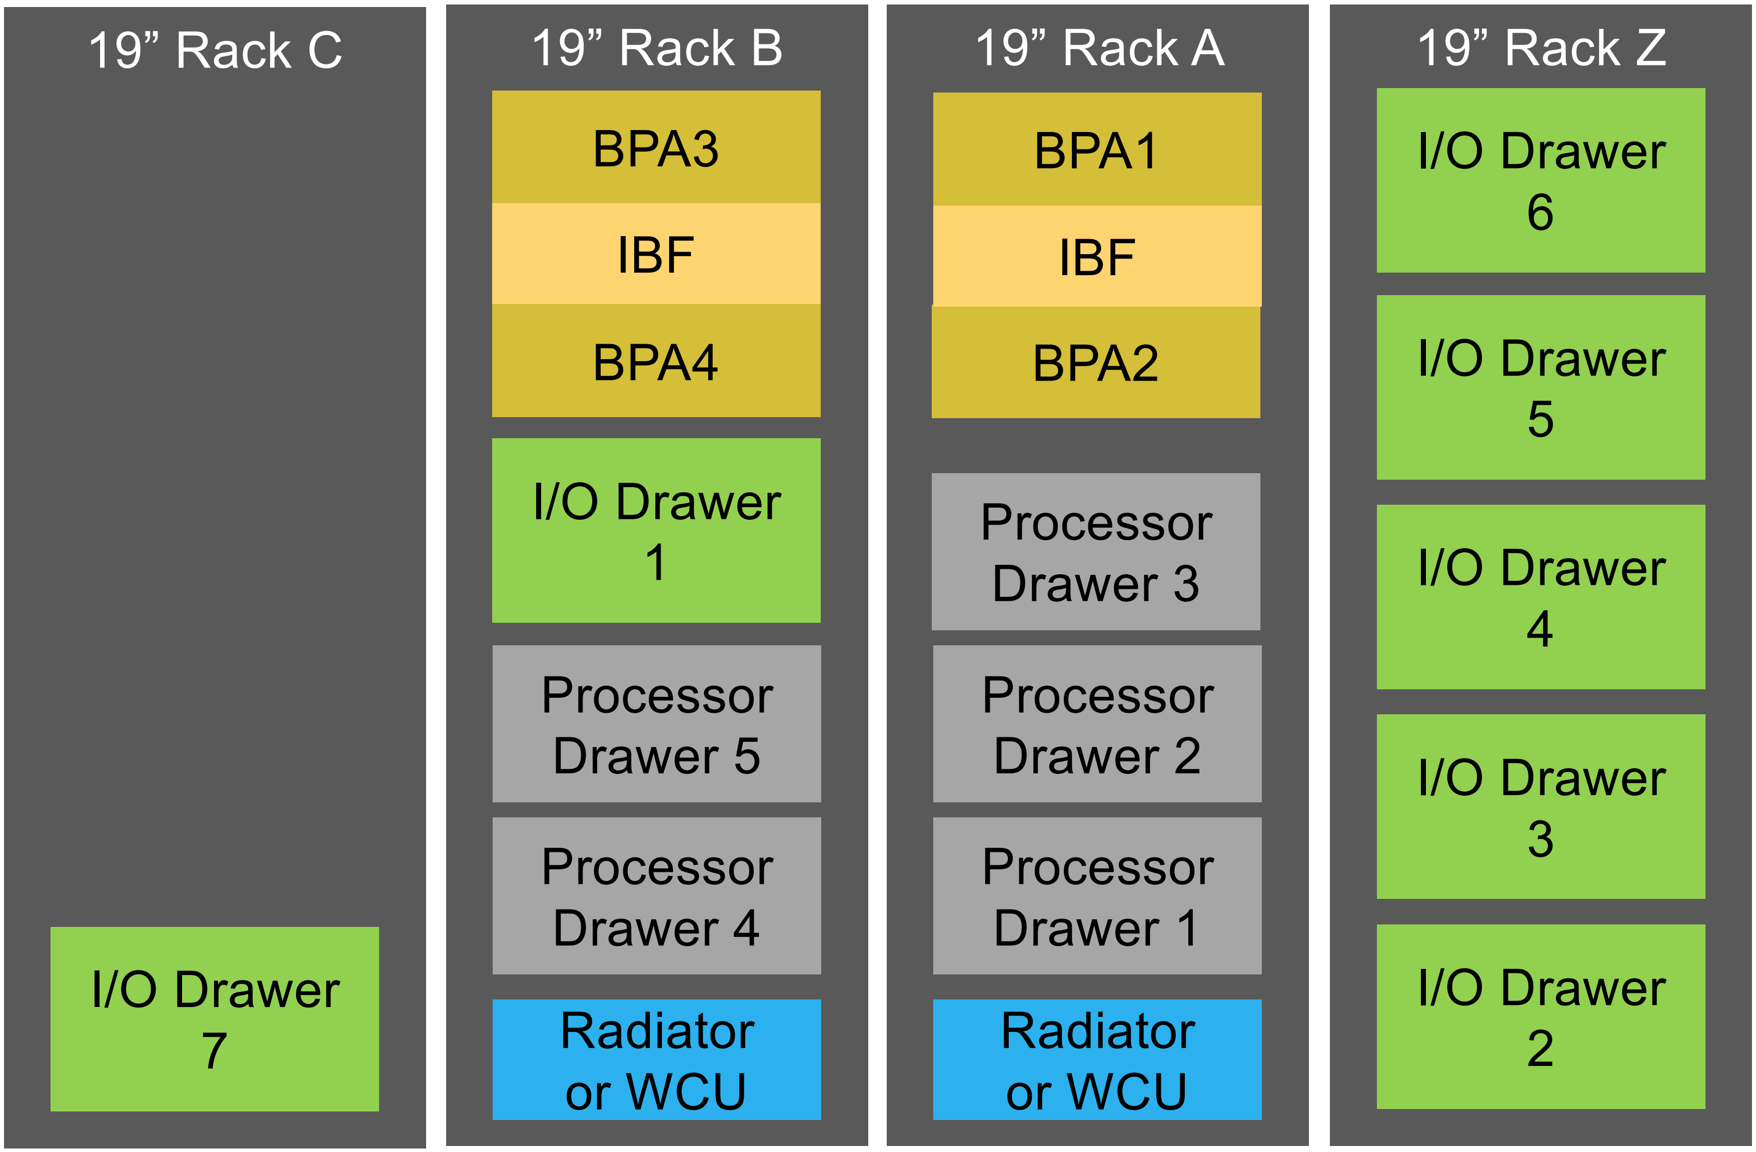

Topics

z15 T01 (LT1) Racks: Scenario 4 – Bulk Power Assembly Powered Heavy Heavy

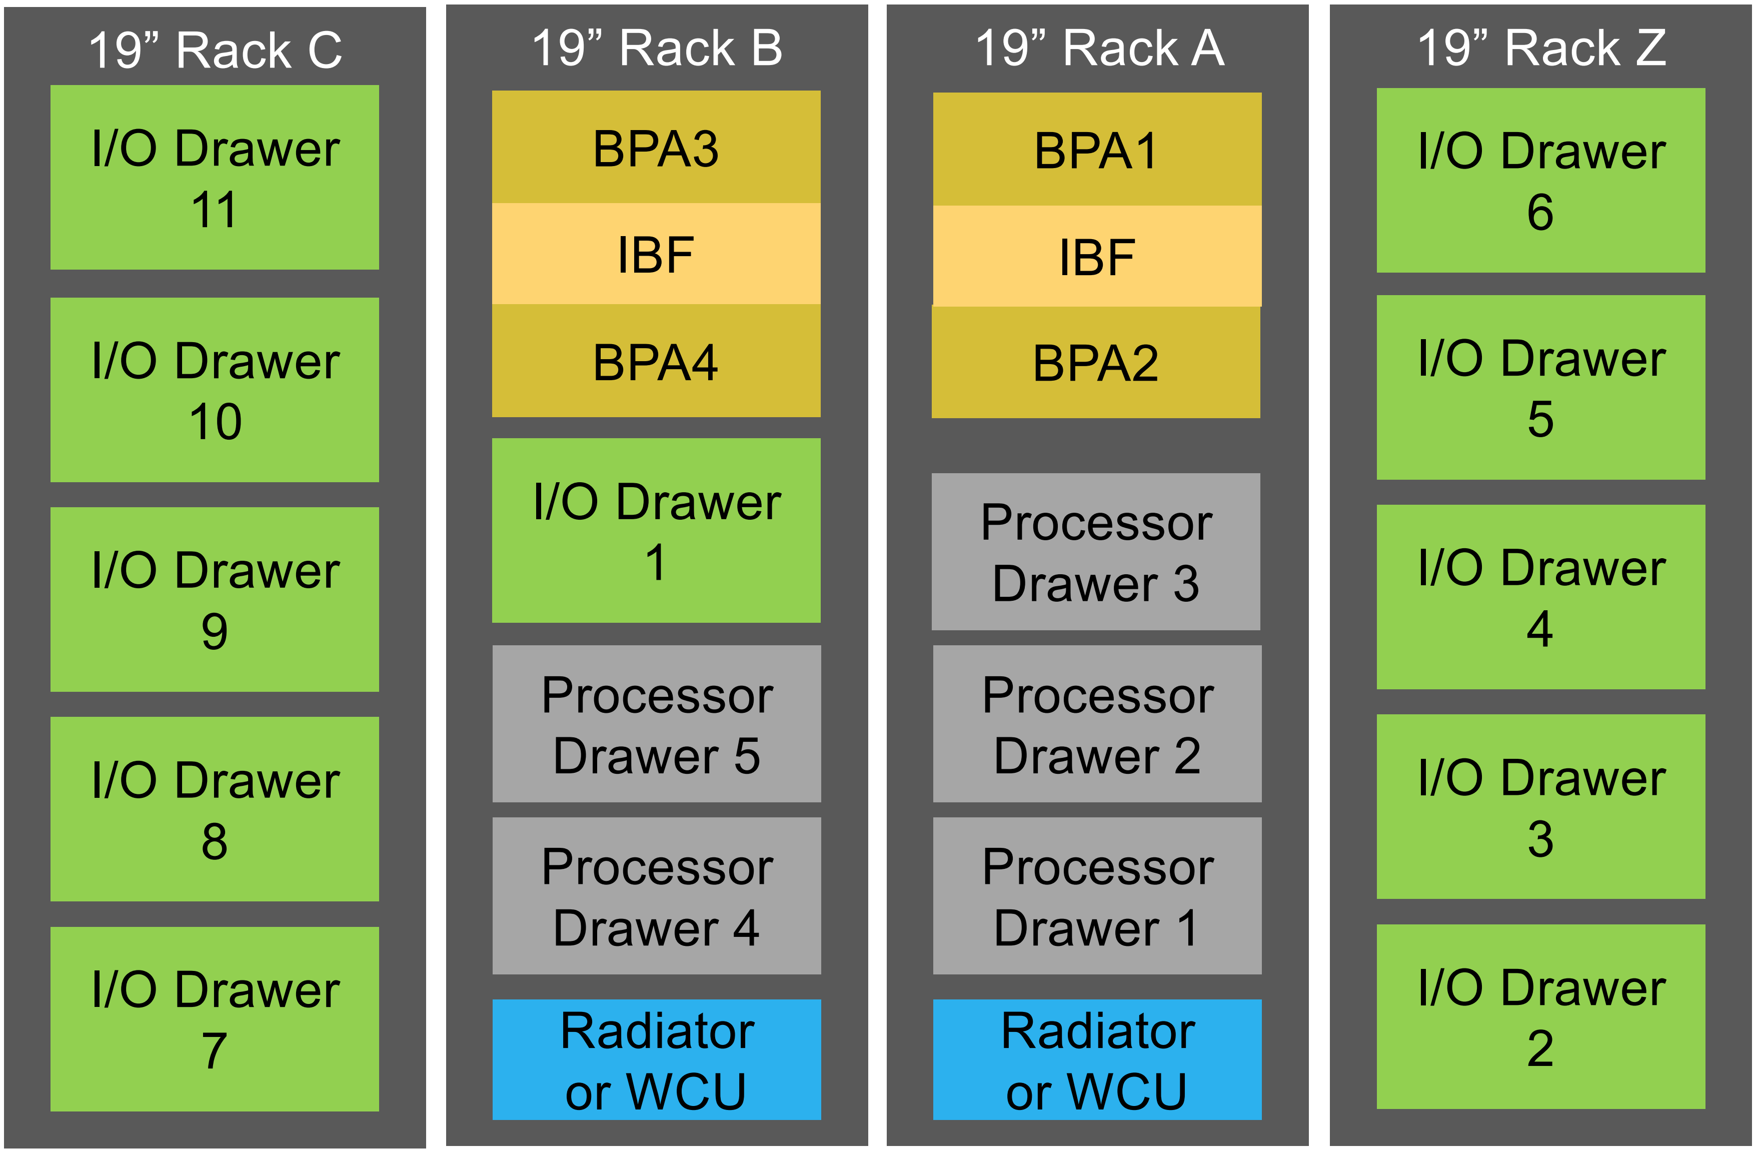

Topics

z15 T01 (LT1) Racks: Scenario 4 - Under the Covers … a Real Rear View

Virtualization Technology

EXTENDING FLEXIBILITY
EXTENDING SECURITY

Topics

LinuxONE Virtualization

The virtualization capabilities of LinuxONE are outstanding and start with a built-in always-on hardware-based hypervisor called “PR/SM” (processor resource/system manager).

PR/SM’s job is created “partitions” … sometimes call logical partitions, or LPARs. A partition is simply a place for an operating system to run. At least 1, but up to 85 partitions can be created by and managed by PR/SM.

LinuxONE runs Linux, so in our partition we could run RHEL or SLES or Ubuntu or an open source version of Linux like ClefOS or others.

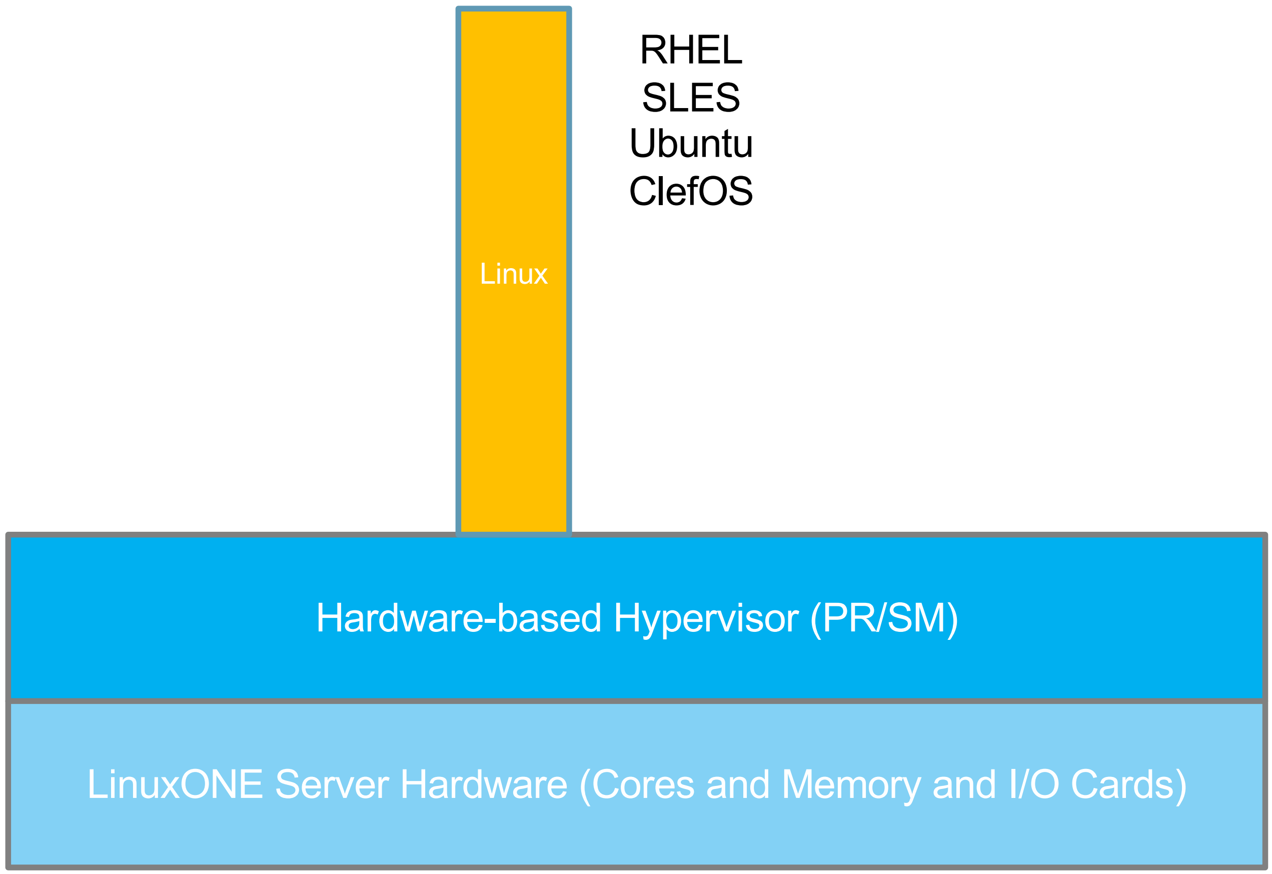

Topics

When multiple partitions are defined and activated simultaneously it is PR/SM's job to allocate processing and memory and I/O resources to the partitions. The sophisticated configuration and management controls provided by PR/SM assure that the partitions that need priority treatment, GET priority treatment!

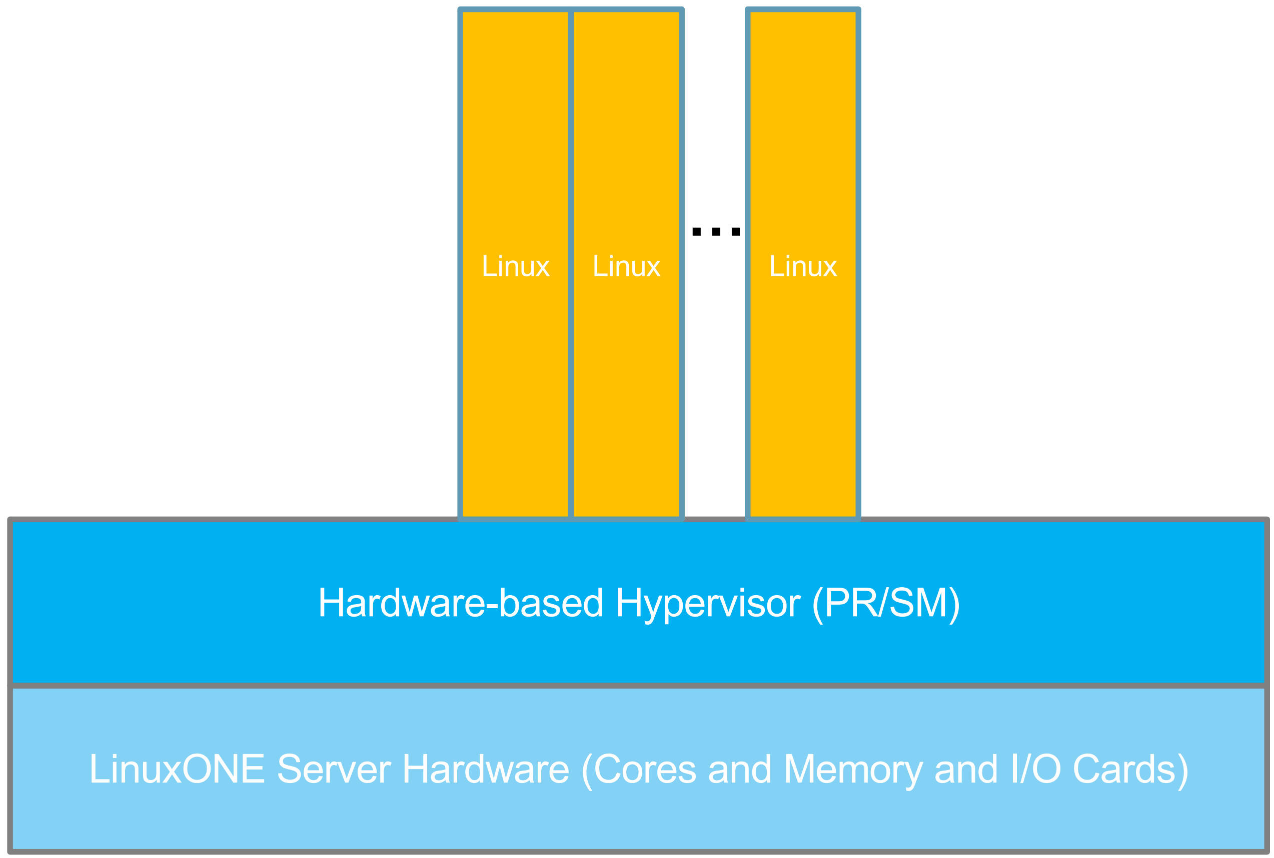

LinuxONE Virtualization

PR/SM offers good flexibility and excellent efficiency, but is limited to 85 partitions. That might not be enough!

When the utmost in flexibility is called for, a software hypervisor is called for, and IBM offers z/VM!   z/VM can be installed in a partition, or in multiple partitions. A common design pattern is a z/VM partition for production, another z/VM partition for QA, and a third z/VM partition for Dev and Test.

Mix and Match? Sure. “Native” Linux-running partitions can be configured alongside z/VM partitions. No problem.

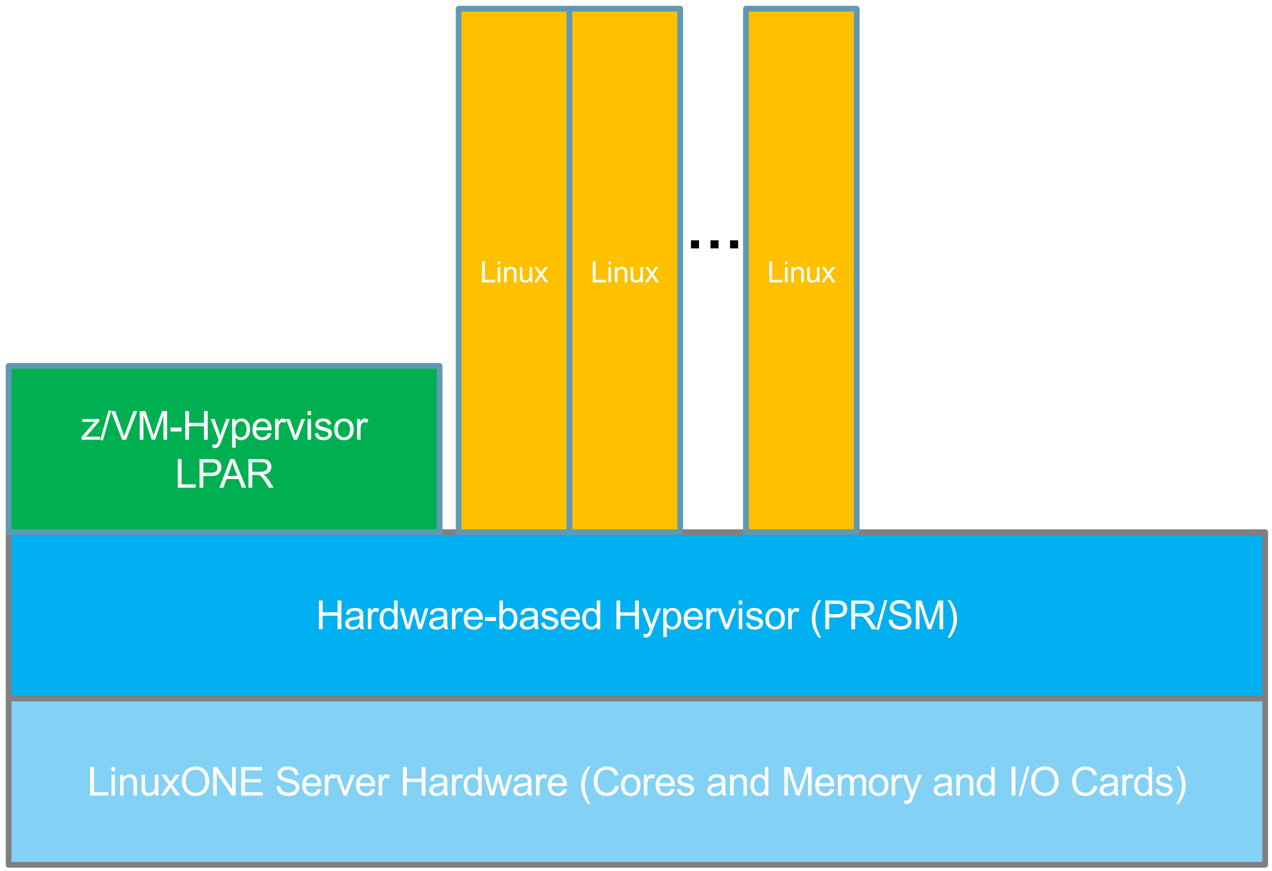

Topics

LinuxONE Virtualization

z/VM is a software hypervisor, and its job is creating one or many (many) “virtual machines” (or sometimes called a “virtual servers”, or even a “virtual guests”). z/VM can create as many virtual servers as the hardware assets allocated to the partition can support. We have customers running z/VM supporting 1000’s of virtual servers with a single copy of z/VM (and let’s not forget about the other copies of z/VM on the same server for the the other environments).

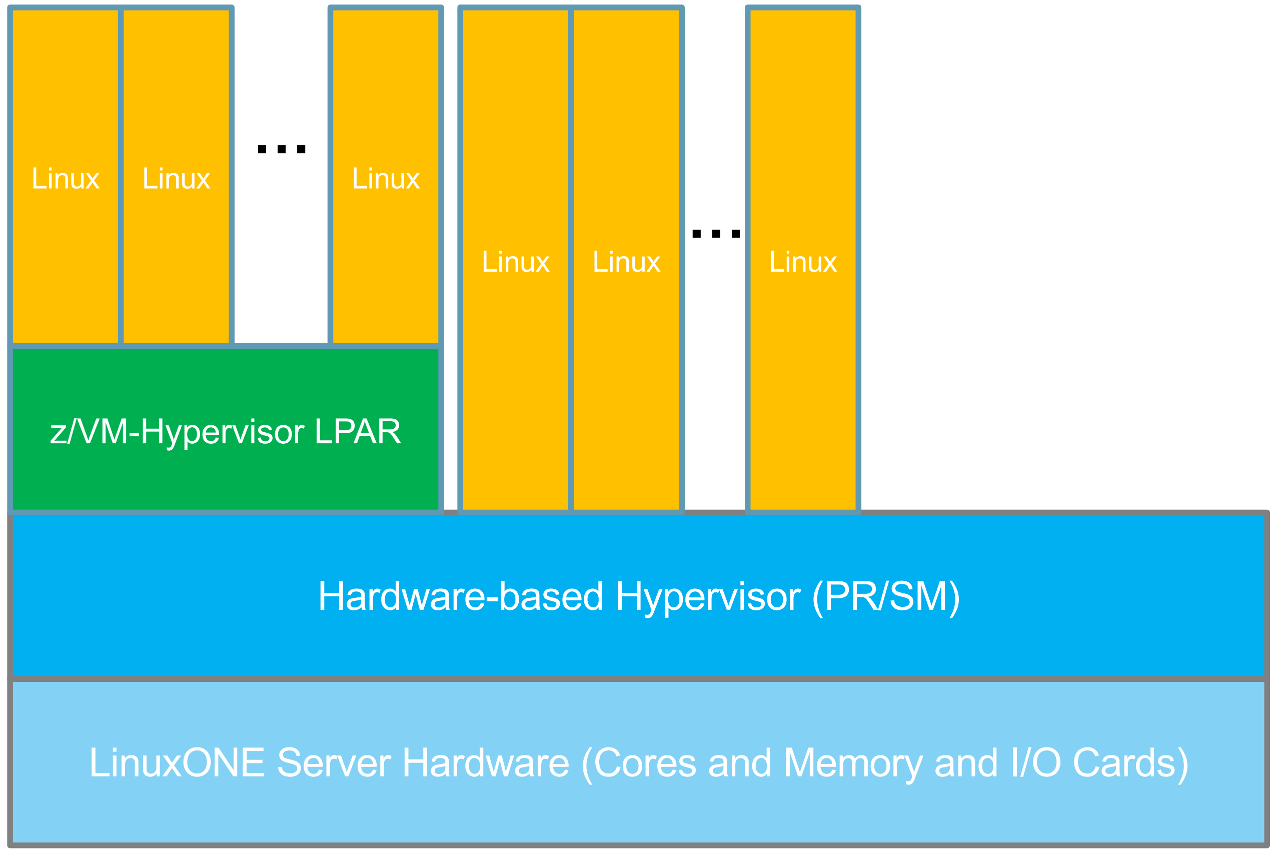

Topics

LinuxONE Virtualization

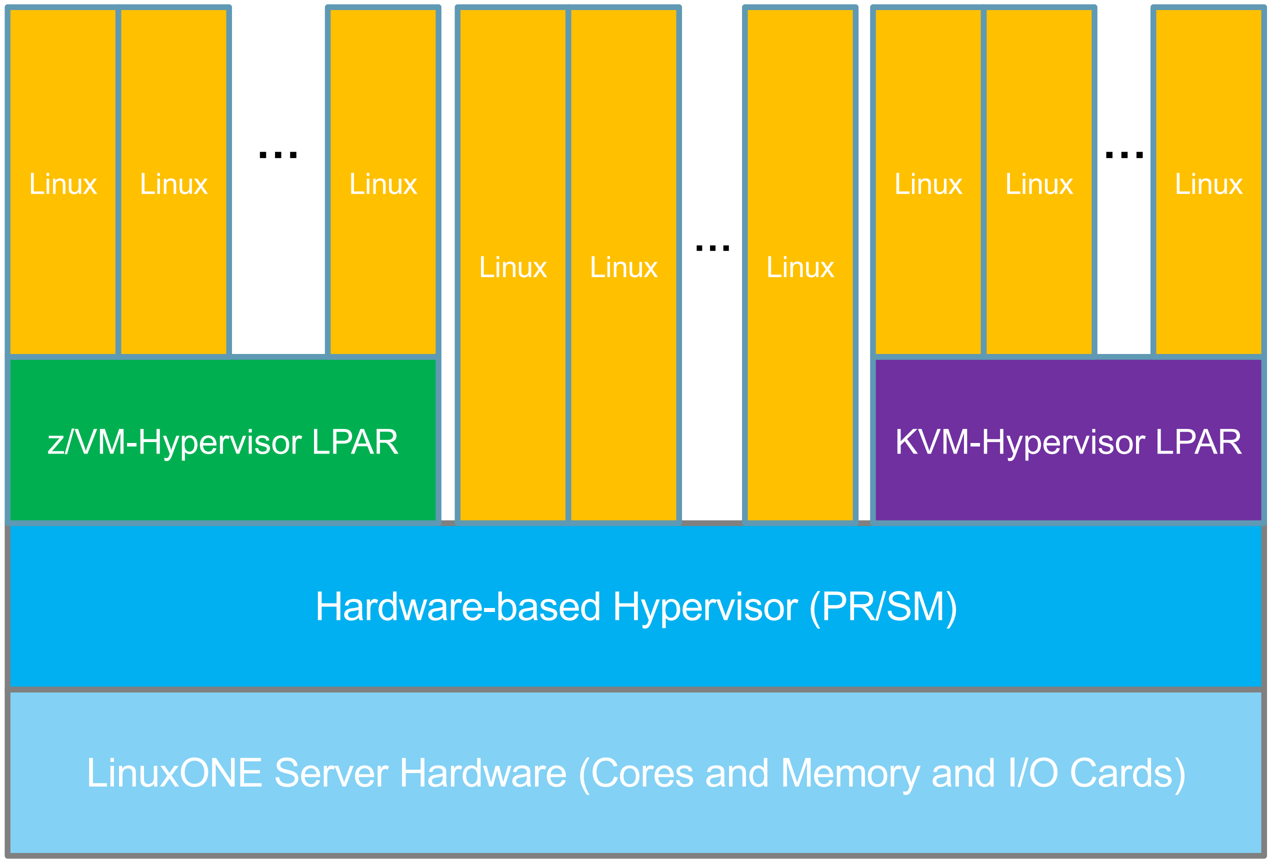

z/VM is great, it has been around a long time, and has a great following. But as stated earlier, IBM and LinuxONE are vested heavily in the open source community, and the open source community has a software hypervisor that it likes very much, called "KVM". LinuxONE supports KVM and similar to z/VM we have clients running many many virtual servers under the control of KVM.

This drawing is accurate, for we can indeed run a z/VM partition, several native-Linux partitions, and a KVM partition simultaneously on the same LinuxONE server!

Topics

LinuxONE Virtualization

To summarize a bit, we saw how PR/SM is built-in and simple and efficient, but is rather limited in the number of partitions it can build (85 on a LinuxONE III). There is a fit for that, for example, we have customers running small numbers of VERY large Oracle databases and doing so in PR/SM mode makes great sense.

But MOST customers require greater flexibility, and as such bring to bare a software hypervisor like z/VM or KVM.

Interesting note. Oracle fully and wholy supports ONE and only ONE non-Oracle hypervisor, and that is z/VM!

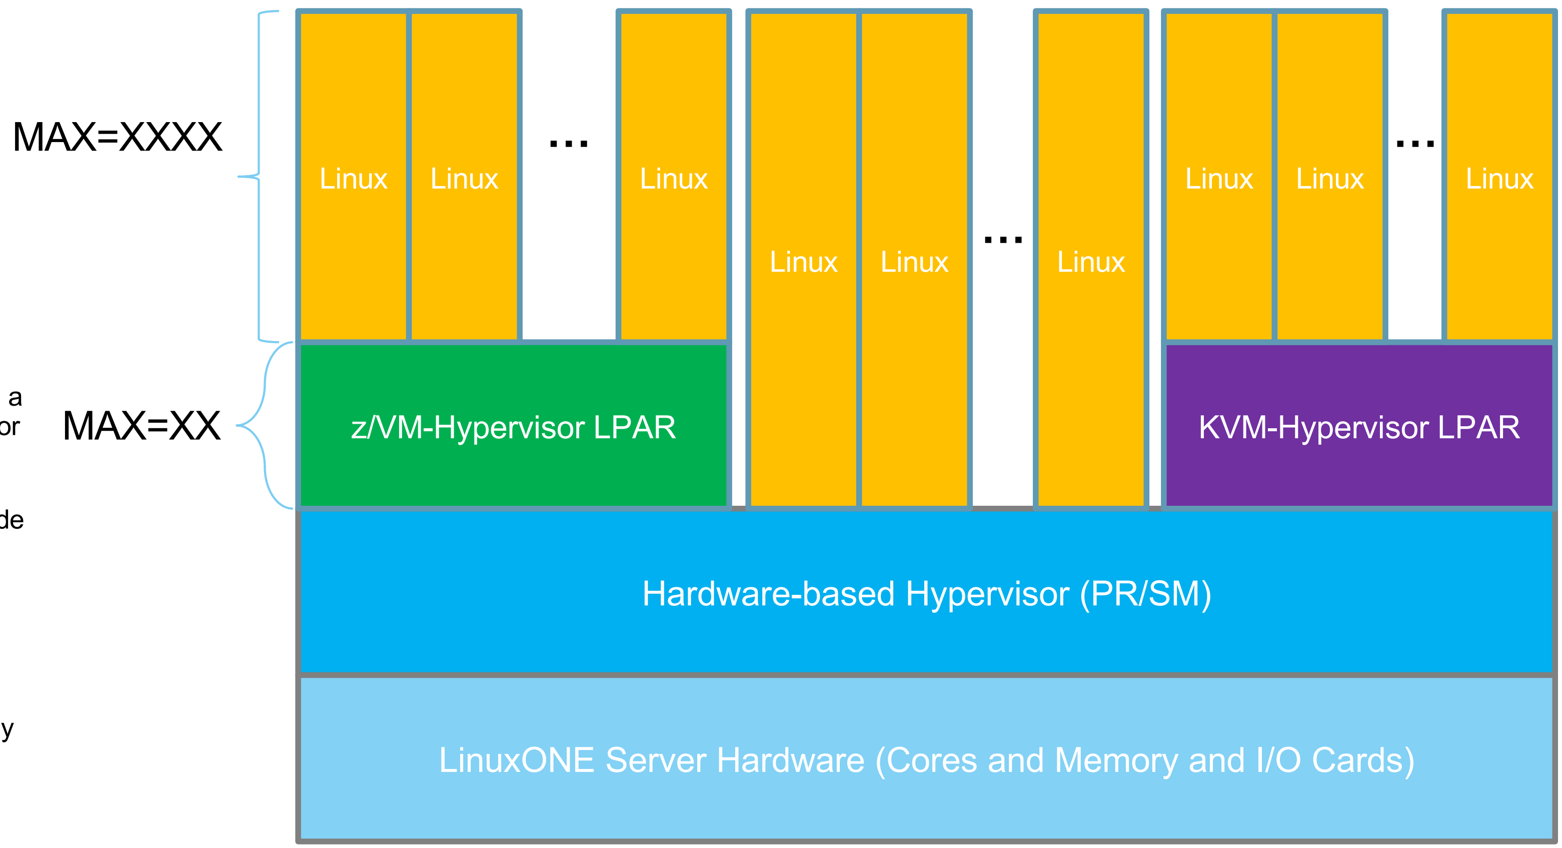

Topics

LinuxONE Networking

It is great to host a ton of virtual servers on a single hardware platform, garnering great efficiency and resource utilization … but what happens when many of those virtual servers need to talk to each other? How many wires do we have to connect to this hardware to get these servers talking to each other?

Luckily, LinuxONE offers terrific "virtual" networking technology at every layer of our virtualization technology stack. PR/SM affords ways for partitions to connect "inside" the box. And both z/VM and KVM offer sophisticated virtual networking capabilities. High performance, low latency, highly secure, and low cost networking is a good thing.

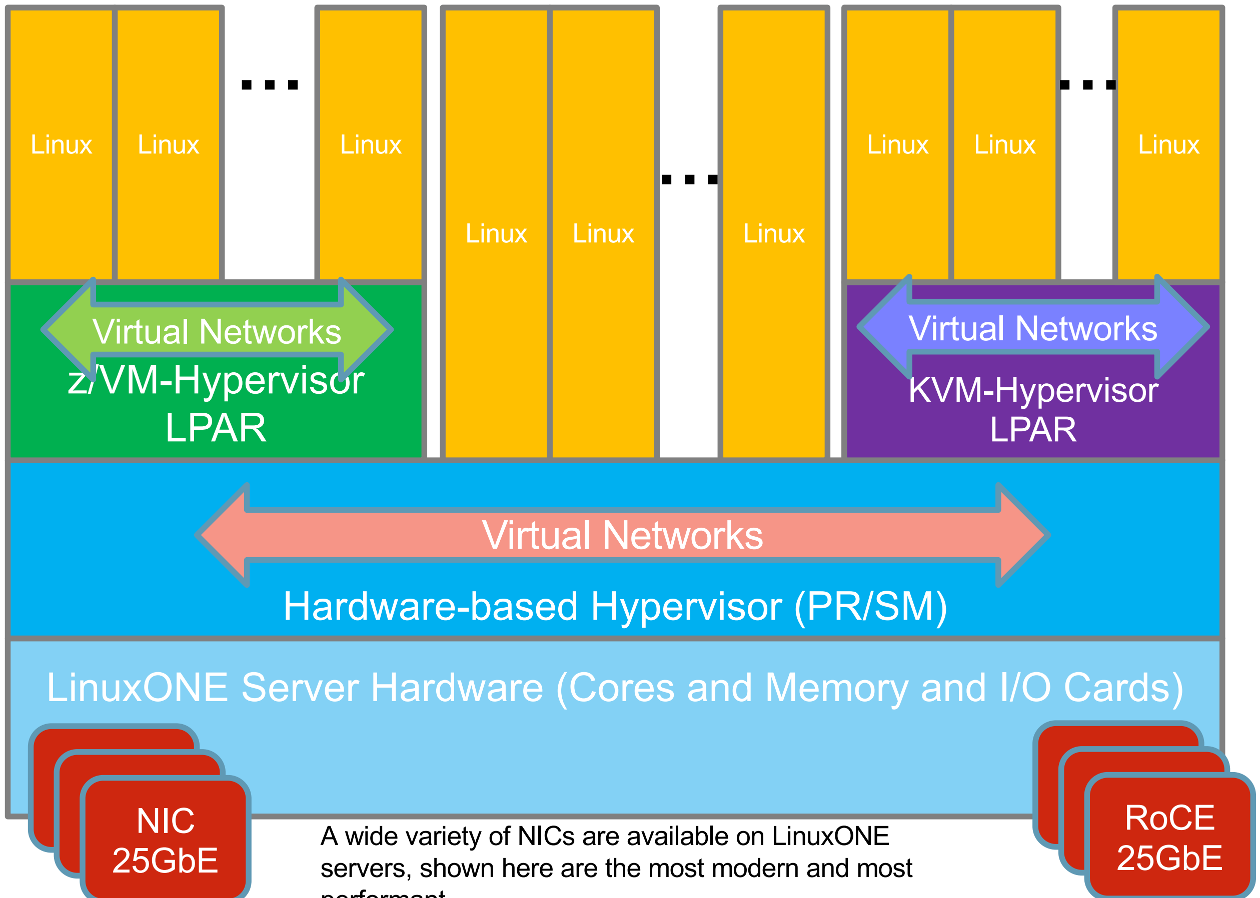

A wide variety of NICs are available on LinuxONE servers, shown here are the most modern and most performant.

Topics

LinuxONE servers need to read and write data. Whether that data resides on SAN-attached disk, or on mainframe-style disks, LinuxONE can attach to it!

LinuxONE storage attachments are fast, and (should you like) numerous.

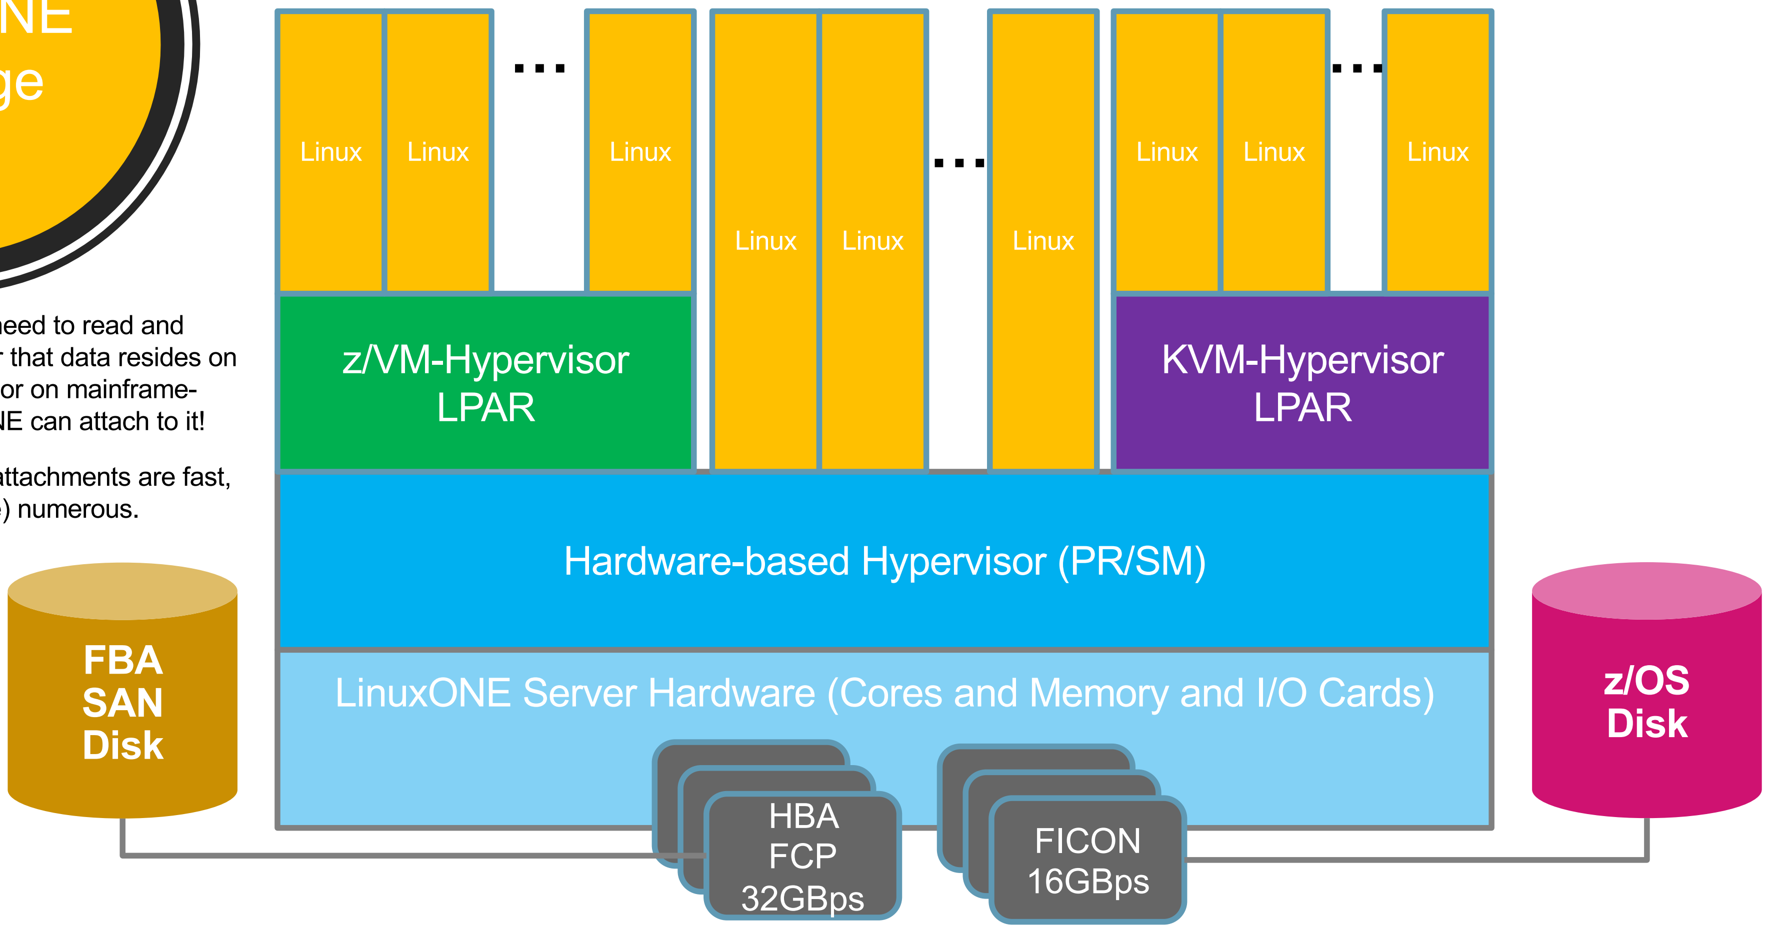

Topics

LinuxONE servers are extremely flexible, commonly running multiple workloads and multiple workload types for multiple workload environments simultaneously.

This drawing aspires to illustrate that flexibility, showing data serving workloads, application serving workloads, messaging and integration workloads, hosted by virtual servers under a variety of hypervisors and/or by native partitions.

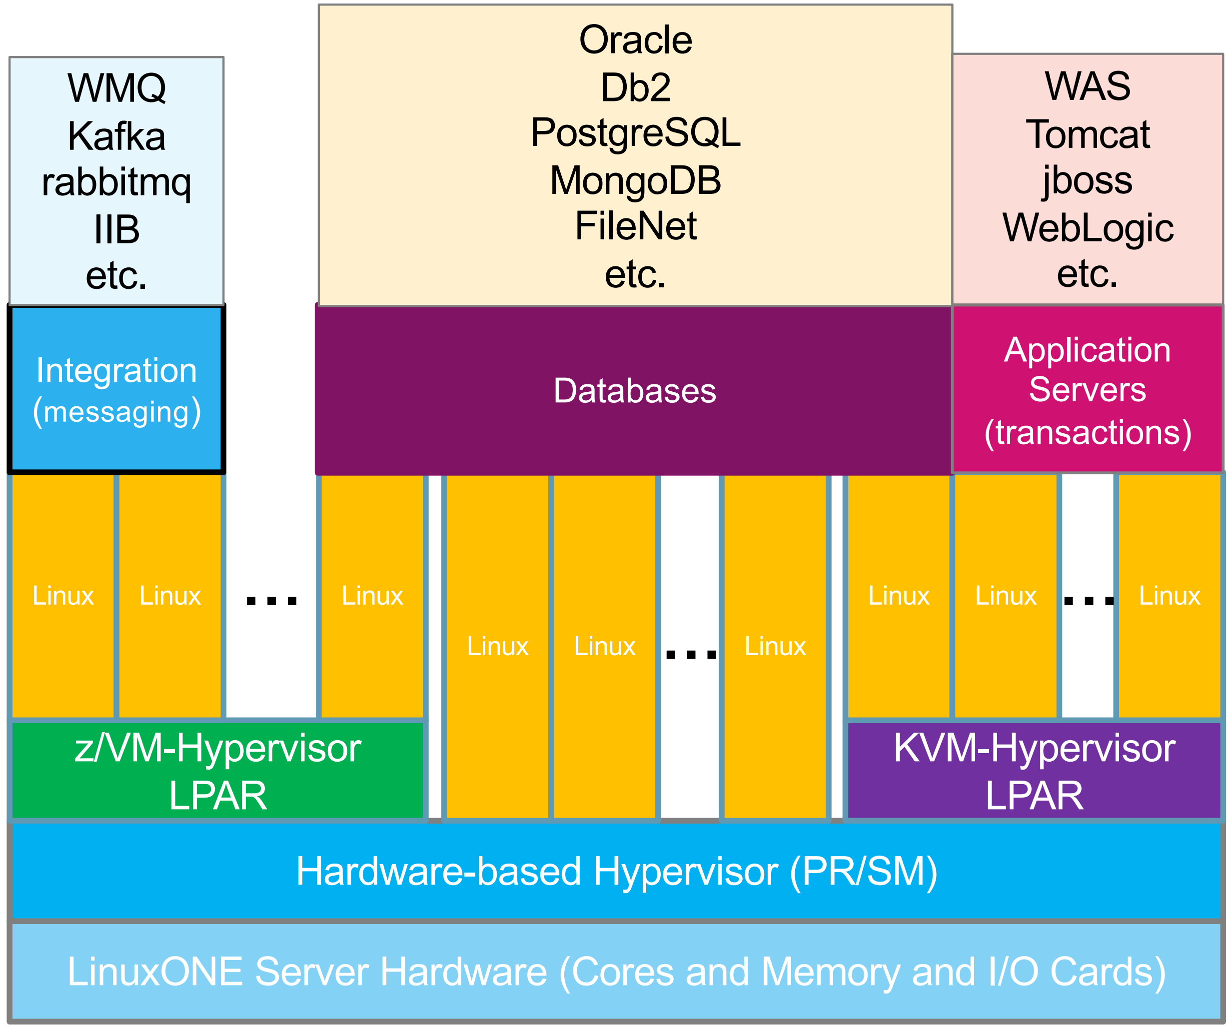

Last fall (September 2019), IBM acquired Redhat, and the rampant adoption of Redhat's Openshift-based technology across IBM's hardware and software portfolio began. The LinuxONE server is no exception, as OpenShift-based products and technologies have been announced and are beginning to "ship" on LinuxONE. Additionally, native Redhat RHEL support and pricing on LinuxONE have seen nice improvements.

This chart illustrates the eventual full participation of LinuxONE across the Redhat and IBM Cloud Pak product set. For customers building cloud-native apps using cloud-native tools and techniques, the LinuxONE platform is ready (or at least it will be in a quarter's time: SOD is for 1H2020).

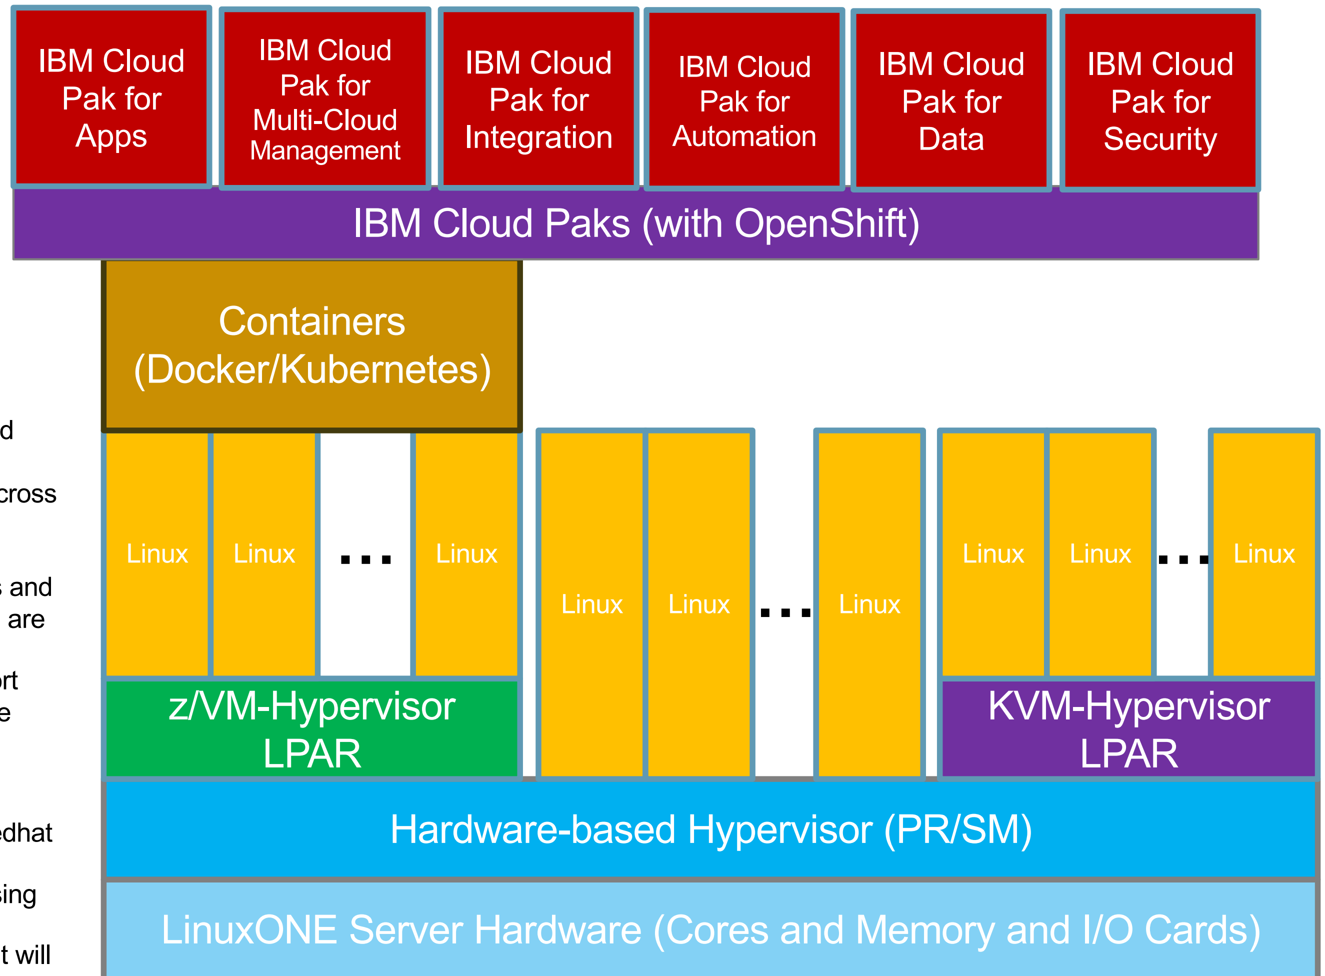

Topics

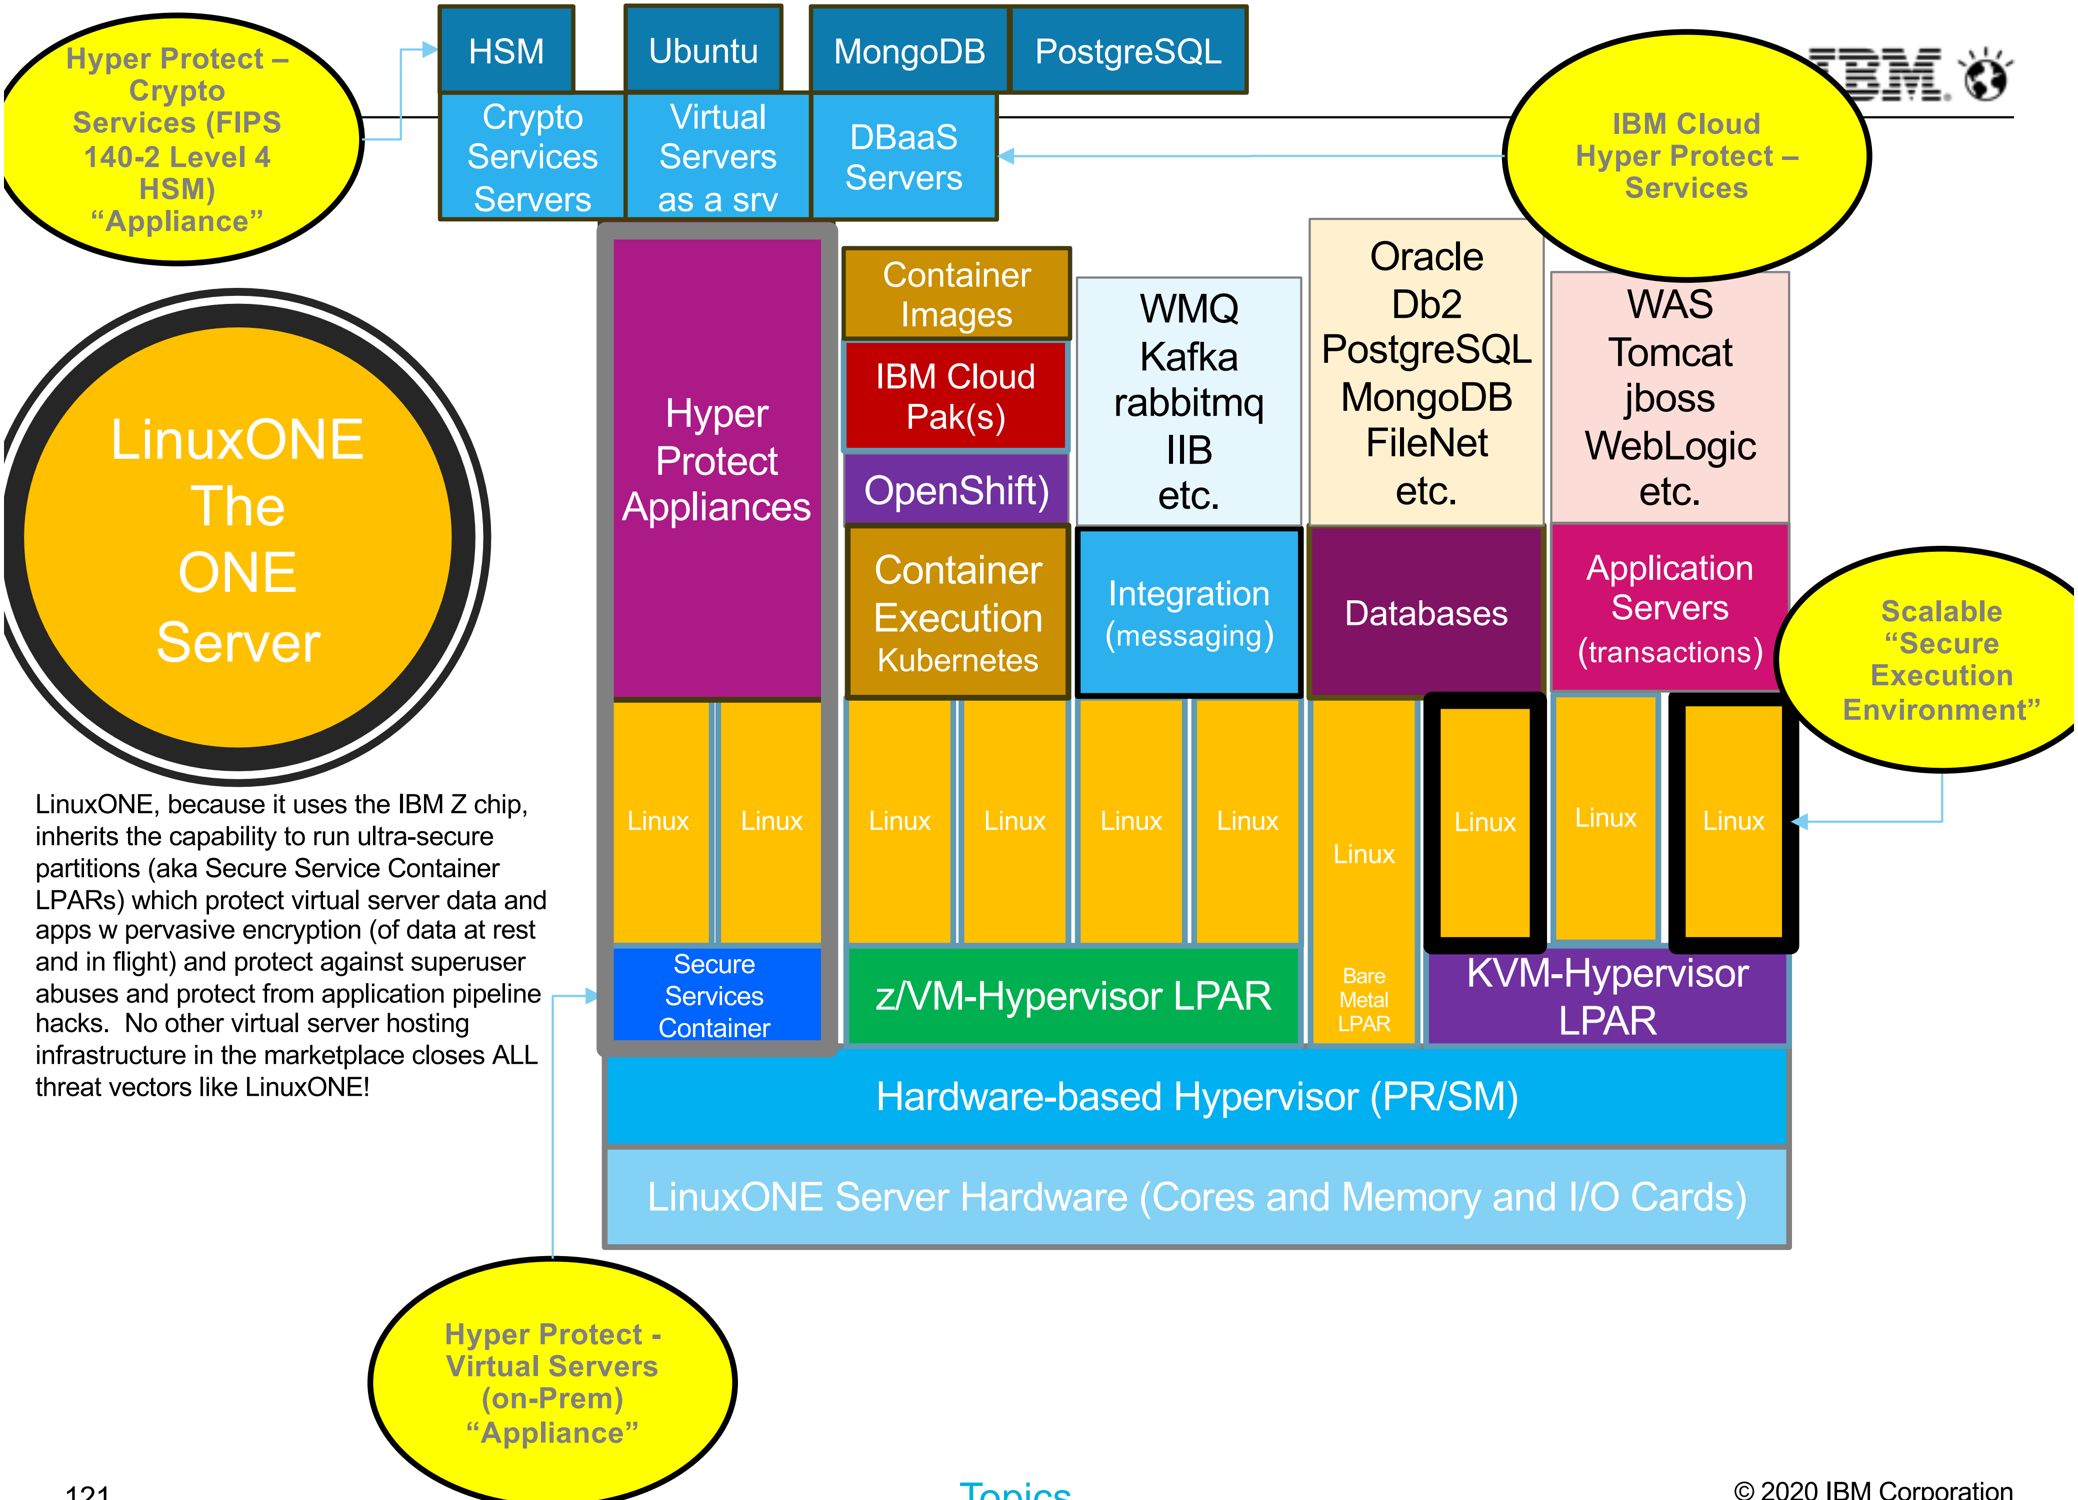

LinuxONE, because it uses the IBM Z chip, inherits the capability to run ultra-secure partitions (aka Secure Service Container LPARs) which protect virtual server data and apps w pervasive encryption (of data at rest and in flight) and protect against superuser abuses and protect from application pipeline hacks. No other virtual server hosting infrastructure in the marketplace closes ALL threat vectors like LinuxONE!

Secure Execution for isolated and secure VM guests

Ensure that only the people with a need-to-know within the organization have access to data in the clear, while still allowing those who don’t to do their jobs efficiently and effectively.

58%

Of security attacks on financial institutions in 2016 were **insider** attacks.

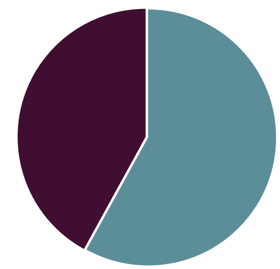

TODAY

- Complex systems for user based access control that are still vulnerable to malicious insiders and inadvertent lapses
- Even encrypted data is at risk while running
- Admins are held liable for data breaches that are due to inadequate security policies and controls

- Guest workloads are completely isolated; easily restrict access on a need-to-know basis and remove liability for system admins
- Data on a running system cannot be seen outside guests
- Scale guest workloads beyond LPAR limitations (85 instances)

I’ll Have Fries with That!

THE SPECIAL SAUCE

Topics

Computer Science Core Performance

$$\left[\frac{\text{Programming Instructions}}{\text{Workload}} + \frac{\text{Memory Instructions}}{\text{Workload}} \right] * \frac{\text{Cycles}}{\text{Instructions}} * \frac{\text{Seconds}}{\text{Cycle}} = \frac{\text{CPU Time (Seconds)}}{\text{Workload}}$$

Fastest Thread

Topics

Computer Science Workload Performance

$$\left[\frac{\text{Programming Instructions}}{\text{Workload}} + \frac{\text{Memory Instructions}}{\text{Workload}}\right] * \frac{\text{Cycles}}{\text{Instructions}} * \frac{\text{Seconds}}{\text{Cycle}} = \frac{\text{CPU Time (Seconds)}}{\text{Workload}}$$

Fastest Thread

$$\frac{\text{CPU Time in Seconds}}{\text{Workload}} + \frac{\text{Synchronous I/O Wait in Seconds}}{\text{Workload}} + \frac{\text{Dispatch-able Wait Time in Seconds}}{\text{Workload}} = \frac{\text{Total Time (Seconds)}}{\text{Workload}}$$

Fastest Server

Topics

Computer Science Core Performance

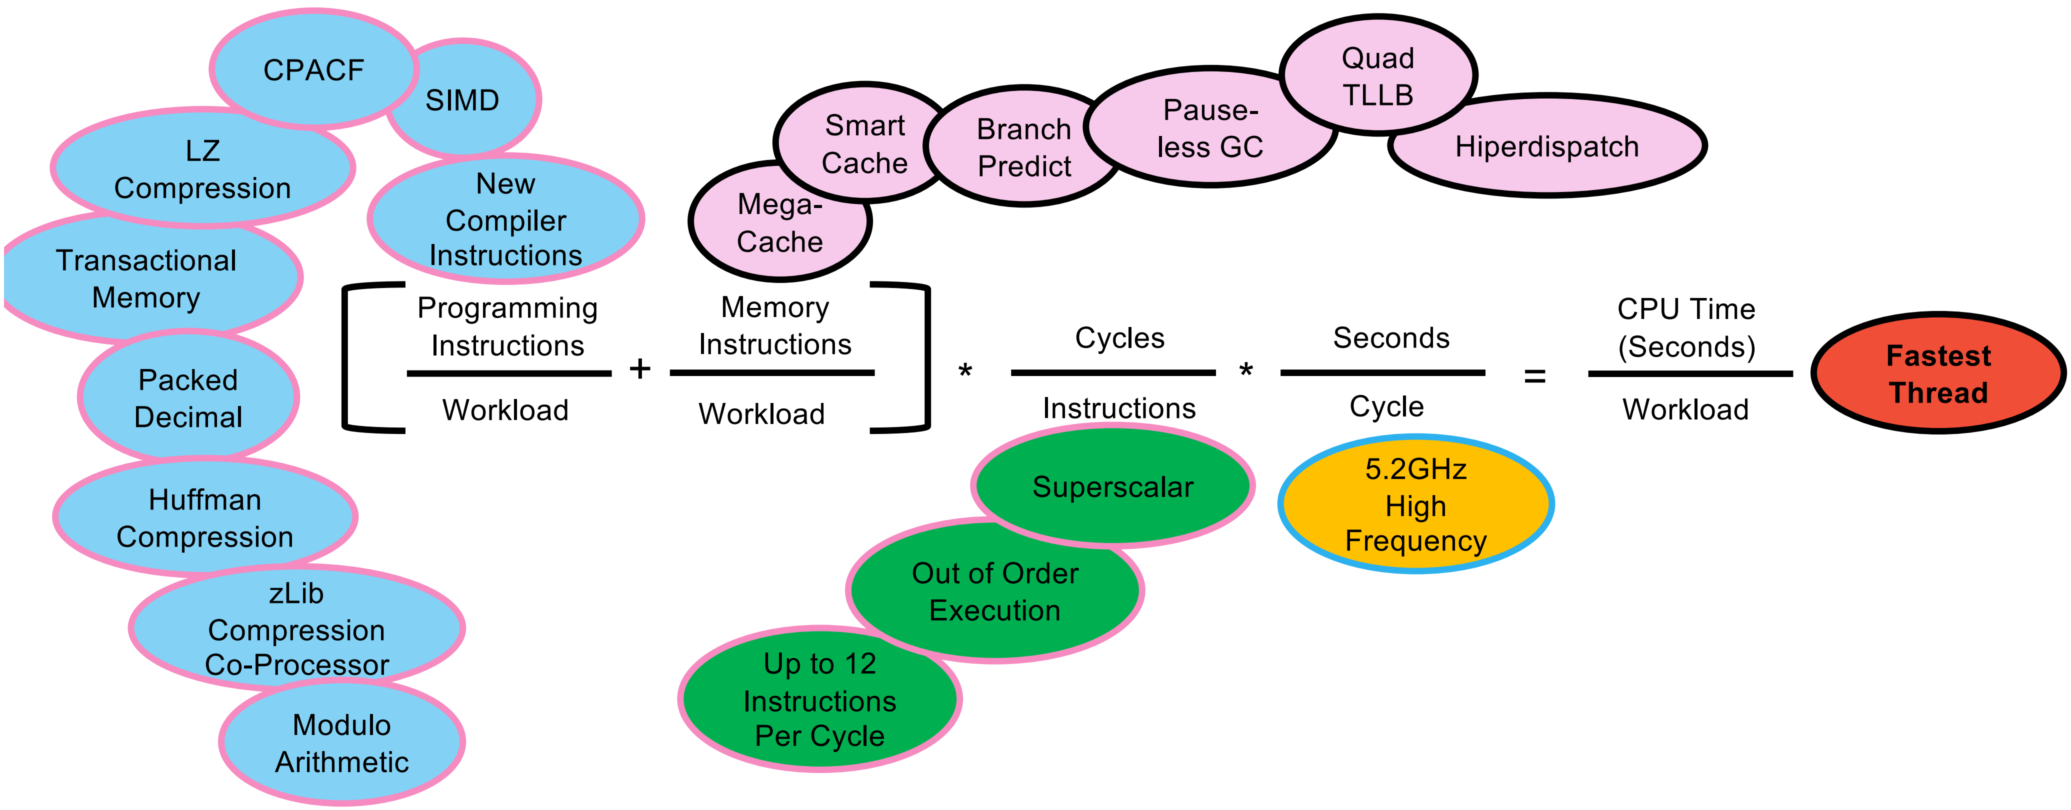

Topics

Computer Science Workload Performance

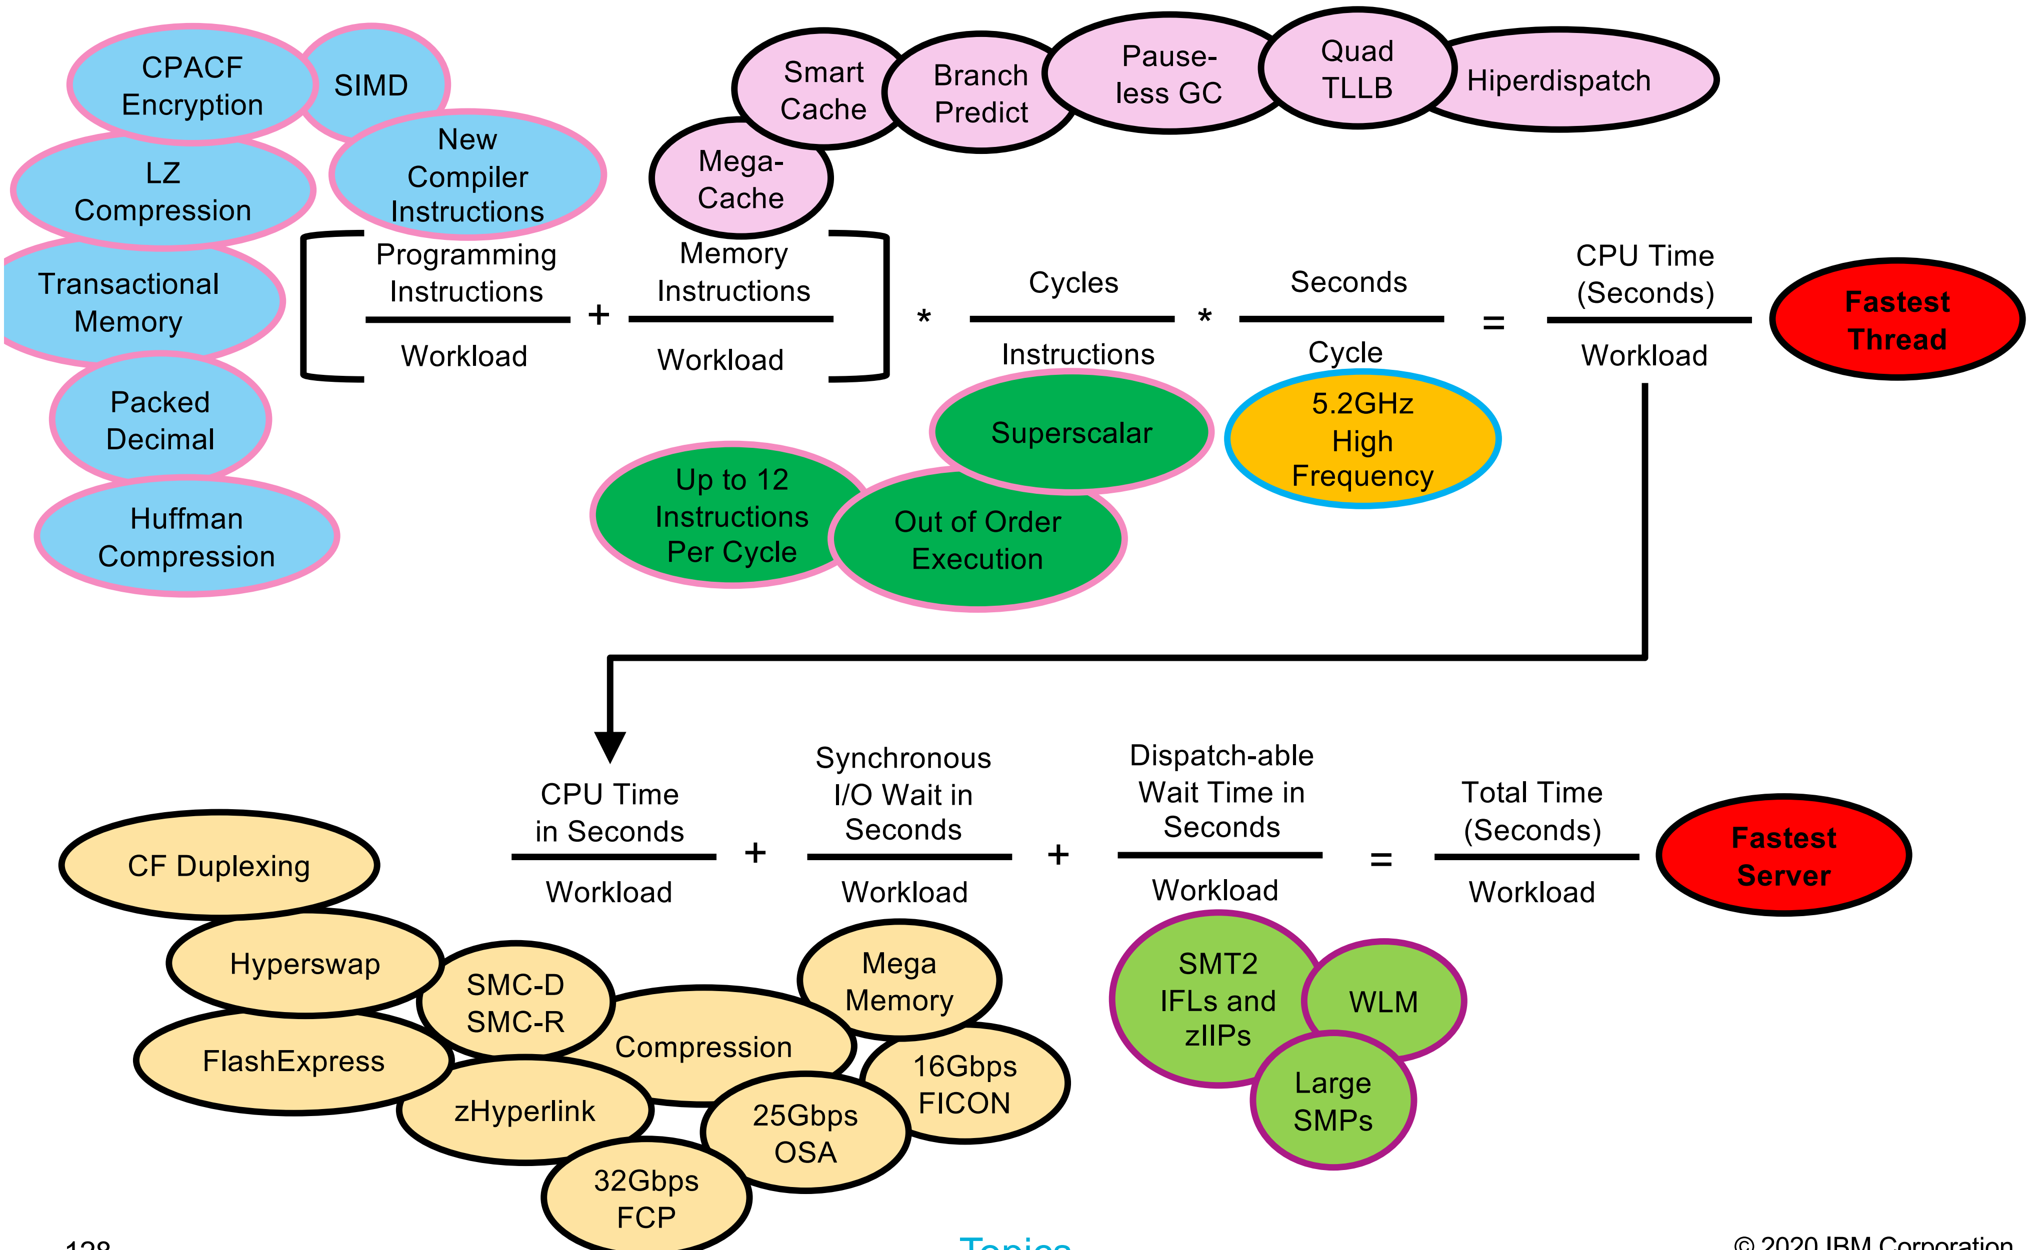

IBM Z Capacity Metrics

		2097 z10	2817 z196	2827 zEC12	2964 z13	3906 z14	8561 z15
Release Date		1Q 2008	3Q 2010	3Q 2012	Q1 2015	3Q 2017	3Q 2019
SMT-1		n/a	n/a	n/a	**1**	**1**	**1**
Max-MIPS/Core	**(MIPS)**	**990**	**1,280**	**1,650**	**1,906**	**2,020**	**2,232**
Max-MIPS/Chip	(MIPS)	3,701	4,833	9,097	13,627	17,783	23,414
Max-MIPS/Drawer	(MIPS)	10,177	16,289	26,604	43,269	50,521	58,197
Max-MIPS/CEC	**(MIPS)**	**43,426**	**68,410**	**103,699**	**154,904**	**195,496**	**240,718**
SMT-2		n/a	n/a	n/a	2	2	2
SMT-2 Factor		n/a	n/a	n/a	20%	25%	27.50%
Max-MIPS/Core	**(MIPS)**	**990**	**1,280**	**1,650**	**2,287**	**2,525**	**2,846**
Max-MIPS/Chip	(MIPS)	3,701	4,833	9,097	16,352	22,229	29,853
Max-MIPS/Drawer	(MIPS)	10,177	16,289	26,604	51,923	63,151	74,201
Max-MIPS/CEC	**(MIPS)**	**43,426**	**68,410**	**103,699**	**185,885**	**244,370**	**306,915**
Max-MIPs/Core/GHz SMT1	(MIPS)	225	246	300	381	388	429
Max-MIPs/Core/GHz SMT2	(MIPS)	225	246	300	457	486	547

Look how much more work we can do!

Look how much more work we do per cycle!!!

Topics

IBM Z Technology Metrics

		2097	z10	2817	z196	2827	zEC12	2964	z13	3906	z14	8561	z15
Release Date			1Q 2008		3Q 2010		3Q 2012		Q1 2015		3Q 2017		3Q 2019
Max-GHz	(GHz)		4.4		5.2		5.5		5.0		5.2		5.2
Chip Transistors	(Billion)		1.00		1.40		2.75		3.99		6.10		9.10
Max Cores per Chip			4		4		6		8		10		12
Max Chips per MCM or Drawer			5		6		6		6		6		4
Max MCMs or Drawers per CEC			4		4		4		4		4		5
Max Chips per CEC			20		24		24		24		24		20
Max Core per CEC			80		96		144		192		240		240
Net Cache/Core	(M)		5.6		15.7		20.8		41.6		30.3		49.6
PU Chip Transistors	(B)		1.00		1.40		2.75		3.99		6.10		9.10
SC Chip Transistors	(B)		1.6		1.5		2.75		7.1		9.7		9.7
MAX KWs			27.5		30.1		27.6		29.8		29.8		28.1
Max-MIPS/CEC/KW	(MIPS)		1579		2273		3757		6238		8200		10922

Look how much
more energy
efficient we are!

Topics

Mainframe CMOS Server MIPS vs. Moore's Law

Year	Single Image MIPS	Moore's Law on Single Image MIPS
1994	61	61
1995	186	61
1996	357	122
1997	446	122
1998	1069	244
1999	1606	244
2000	3061	488
2001	3061	488
2002	3804	976
2003	11391	976
2004	11391	1952
2005	11391	1952
2006	23716	3904
2007	23716	3904
2008	43426	7808
2009	43426	7808
2010	68410	15616
2011	68410	15616
2012	103699	31232
2013	103699	31232
2014	103699	62464
2015	154904	62464
2016	154904	124928
2017	195496	124928
2018	195496	249856
2019	240718	249856

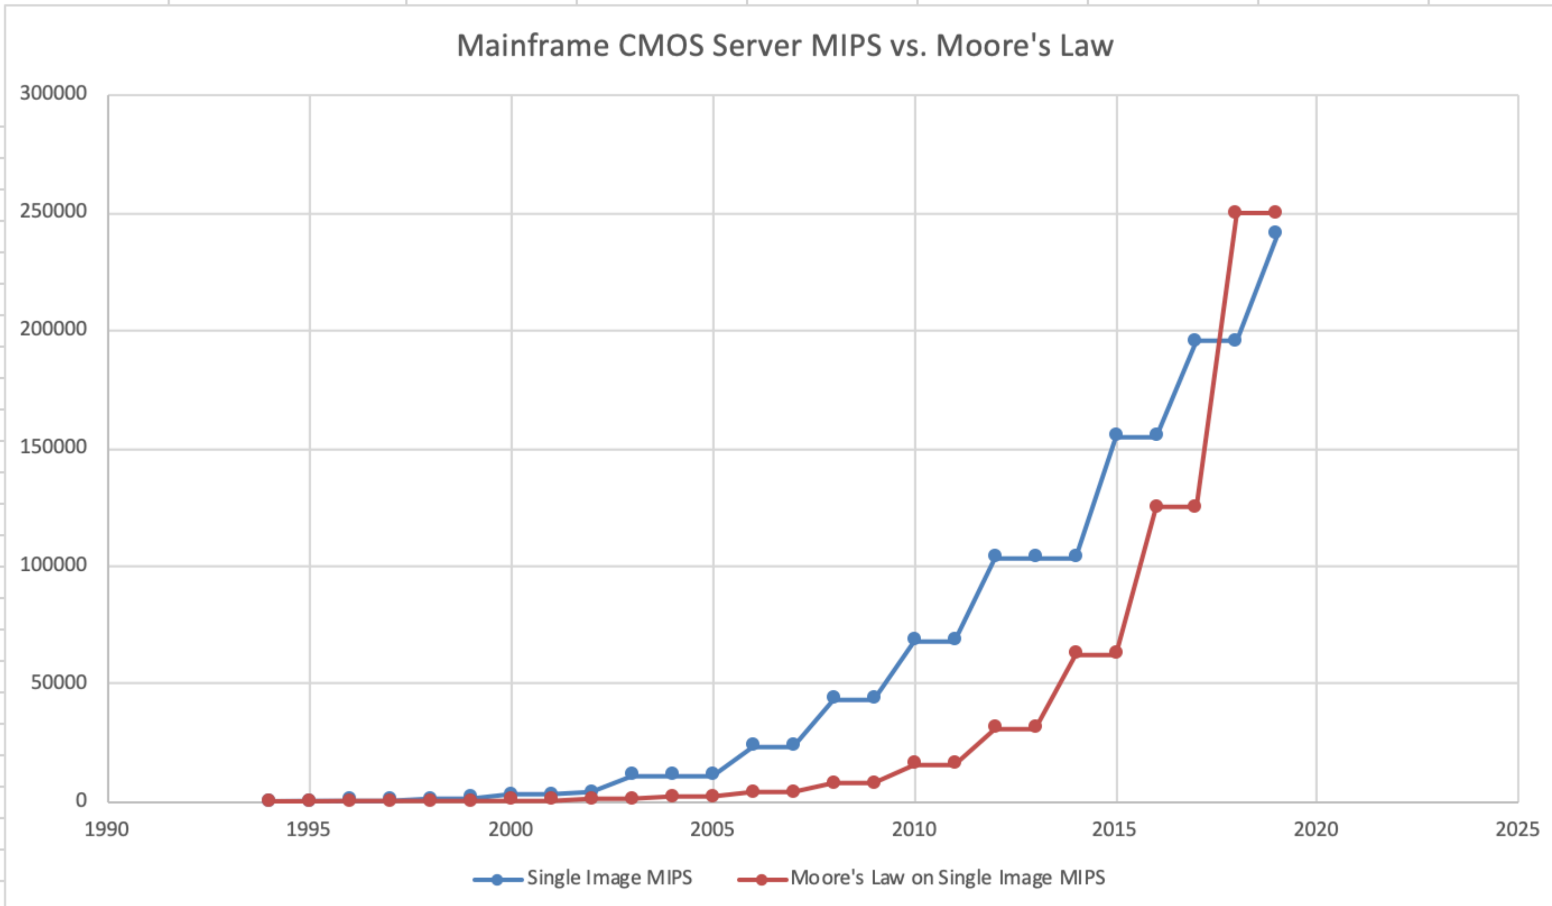

Comparing Cores
An exercise in "example-atory" mathematics

- (A) z Cores are faster (by about 2X) ... 5.2GHz vs. 2.6 (or so)GHz
- (B) z Cores have a ton more cache and z cache is a lot smarter than x86 cache ... so z Cores WAIT (for data and instruction cache loads) a lot less ... something like 5X+ more cache and 2X less waiting
- (C) z Systems have I/O offload to SAPs and to smart I/O cards (OSA cards and FCP cards) ... so z Cores run fewer I/O-managing instructions then x86 (by maybe 25% or some such) ... (a 1.25X factor)
- (D) z Systems, because they are bigger, can run a wider variety of workloads (more VMs and/or more containers) of varying types and varying priorities and as such can consistently fill up processor utilization to higher levels ... something like 80% or 90% for Linux on z compared to something like 40% to 45% on x86 (a 2X factor)
- Doing the math, we get z Core's can do the work of x86 cores at this consolidation rate:
  - (A) 2X * (B) 2X * (C) 1.25X * (D) 2X = 10X
    - For customers who report lower utilization on x86... factor D goes up
    - For customers with higher I/O loads (databases or messaging workloads) ... factor C goes up
    - For customers with older x86 stuff ... factor A goes up and factor B goes up
    - For customers with fewer VMs on VMware ... factor D goes up
    - For customers who segregate VMs and workloads on various VMware instances ... factor D goes up
  - You can play with the factors (and the math) and pretty easily see the 10X go to 50X for fringe cases ... explaining why the CPO team so often sees and says 20 to 50.

Topics

IBM Z Servers Continue to Scale with z15

Each new range continues to deliver:

- New functions
- Pervasive encryption – data centric
- Unprecedented capacity to meet consolidation needs
- Innovative I/O for transactional environments
- Improved efficiency to further reduce energy consumption
- Continues to delivering flexible and simplified on demand capacity

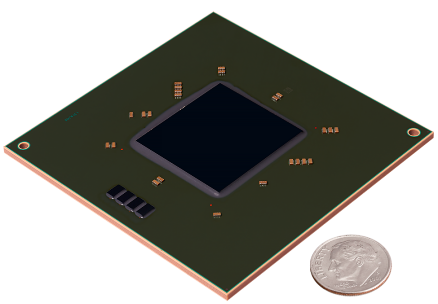

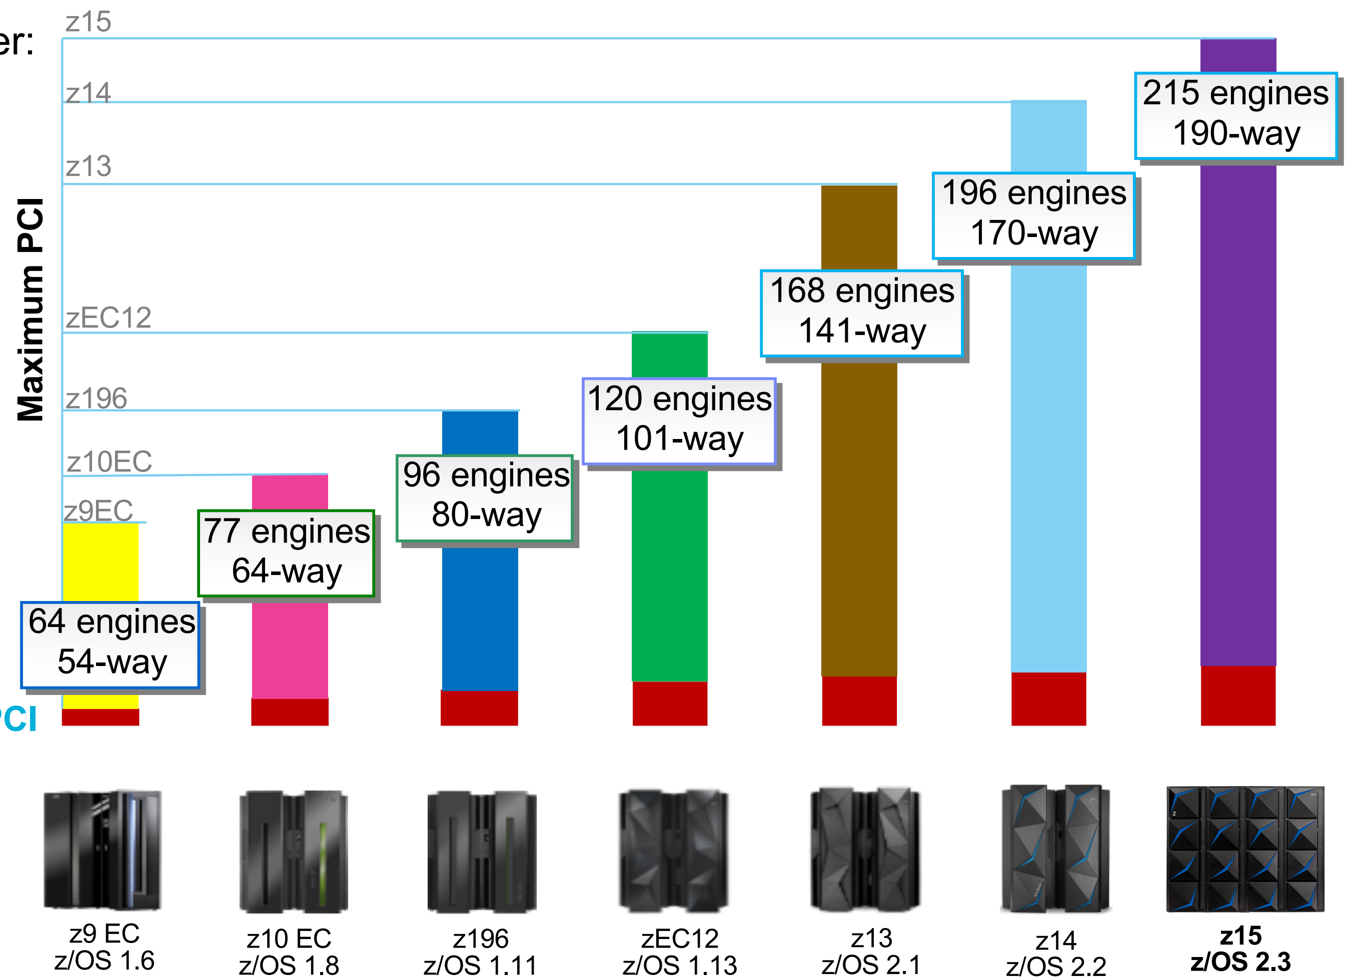

PCI - Processor Capacity Index

Note: OS supports varies for the number of ‘engines’ supported

Balanced system design*

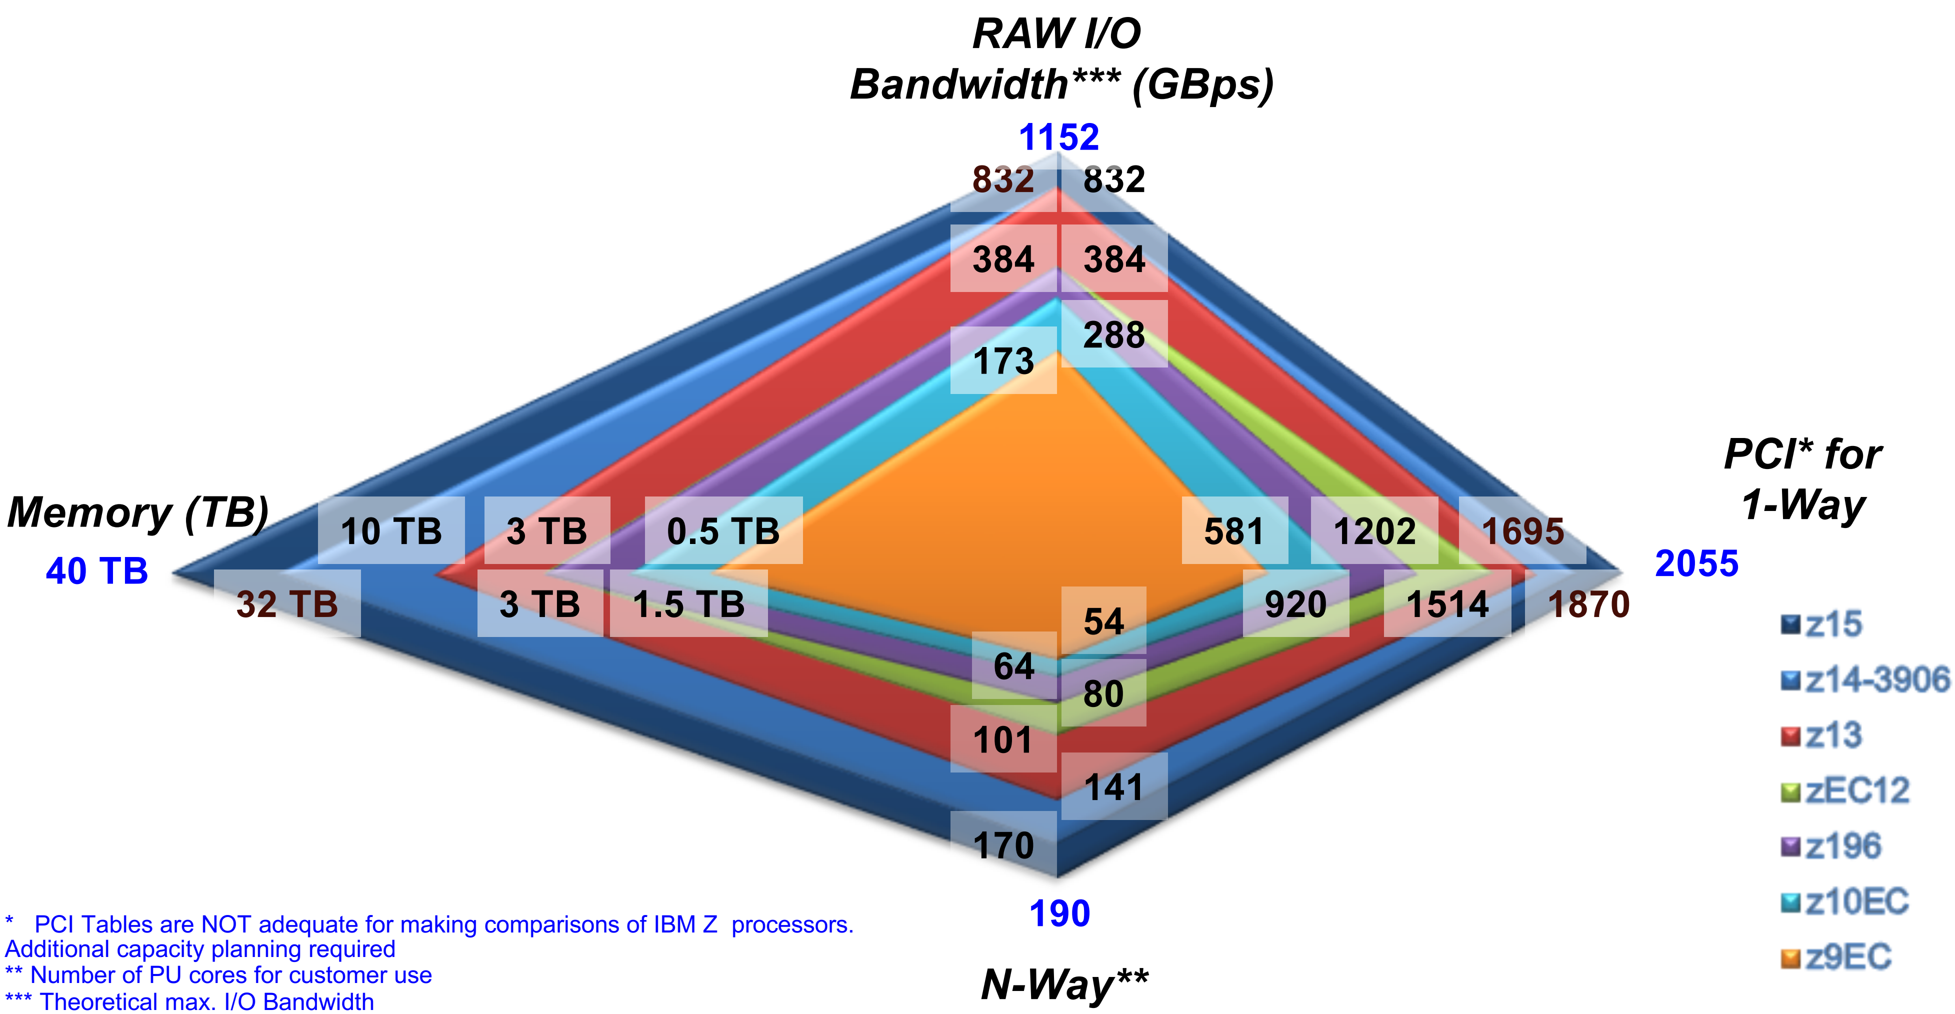

* PCI Tables are NOT adequate for making comparisons of IBM Z processors. Additional capacity planning required
** Number of PU cores for customer use
*** Theoretical max. I/O Bandwidth

Topics

Protection of data that must be shared

New ***z/OS Data Privacy for Diagnostics*** is a z/OS capability **exclusive to z15** with the ability to control access to data shared with business partners and eco-systems

Challenges

Protection from accidentally sharing sensitive data when sending diagnostic information to vendors

Concern for organizations who must comply with GDPR laws and/or other data privacy laws or company mandates

Client Value

Sensitive data tagging APIs combined with machine learning (ML) to detect, tag and redact all tagged data from diagnostic dumps

MVP is working with 1st set of exploiters (Db2, IMS and some DFSMS™ components) to provide the infrastructure to tag sensitive data in z/OS

Tagging does not impact dump times

Supported on IBM z15 running z/OS 2.3 or 2.4

Cryptographic acceleration with z15 hardware

Cryptographic acceleration with Crypto Express7S:

- Improved SSL/TLS handshake performance on z15 with Crypto Express7S compared to z14 with Crypto Express6S
- Updates to Common Cryptographic Architecture (CCA) for security modules that enhance remote ATM key loading, offer new protections for banking payments, and extended compliance support to stay up to date on industry standards

Cryptographic coprocessor on every core with CP Assist for Cryptographic Function (CPACF):

- Enhanced with elliptic curve cryptographic (ECC) algorithms that can help reduce CPU consumption for applications like Blockchain
- Enable an EP11 secure key to be converted to a protected key that can be used by CPACF

Designed for EAL5+ and FIPS 140-2 Level 4

IBM Z investments in Quantum Safe Cryptography

Initial z15 Capability

- Initially delivered via Quantum safe digital signatures for z/OS SMF records
- Agility in algorithms to update as standards evolve
- Acceleration coming in HSM for essential primitives for Quantum Safe cryptography

Timeline of Quantum Advantage vs. Data Lifecycle

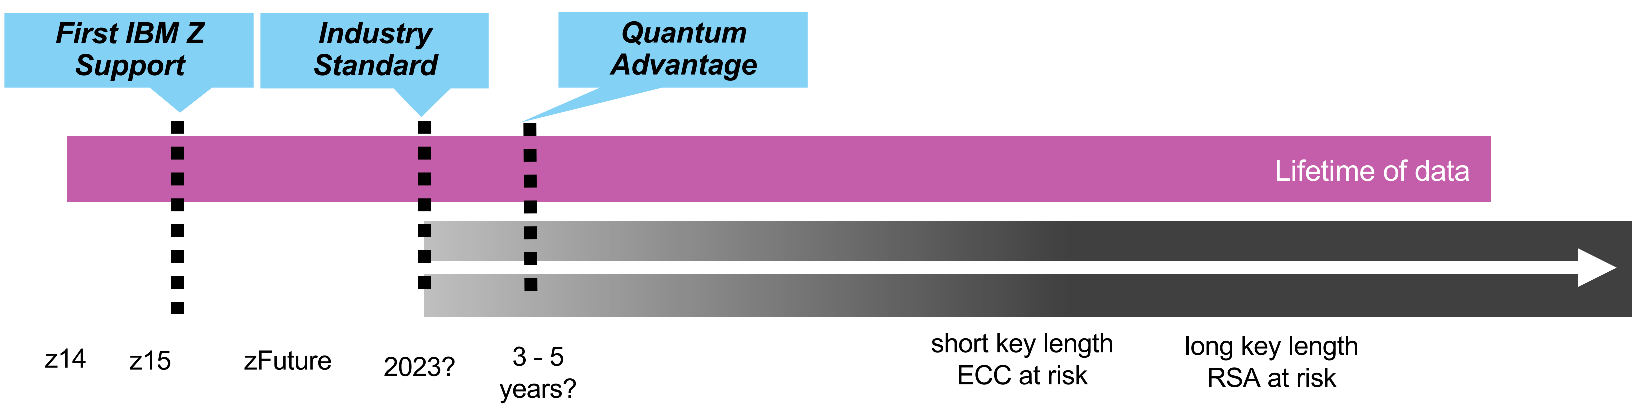

Expanding the z/OS Software Ecosystem

New ***z/OS Container Extensions*** delivers unmodified Linux on Z Docker images running inside z/OS

Challenges

Porting desired software to z/OS requires effort and presents time-to-value and currency concerns

Requirement to deploy dependent software hosted on separate Linux servers leads to complicated z/OS operational procedures and hinders the ability to take full advantage of z/OS Qualities of Service.

Client Value

Modernize z/OS workloads by providing flexibility for development and operations on Z

Integration with other DevOps tools and Linux applications all in z/OS

Maintain operational control and extend z/OS Qualities of Service to Linux software

Make use of existing IT investments by employing Linux within the Z platform

BONUS – workloads run on zCX are zIIP eligible

Topics

zOperational Data Generation and Analytics:

*New offerings for Middleware interdependency data generation, and automated z/OS cross stack analytics

New - z/OS Workload Interaction Correlator & Navigator

Enriched Data

- Generates high frequency (5 seconds) standardized, summarized, and synchronized activity for workload components
- Provides response time interdependencies for CICS®, IMS™, and Db2®
- Embodies First-Failure Data Capture (FFDC) for workload / performance diagnosis

+

Intuitive Analytics

- Visually intuitive, interactive analytics identifies cause and victim workload component peers
- Make more timely, informed decisions from code and test, to system configuration and maintenance, to problem analysis
- User experience delivered as part of Zowe open source framework

Capture all essential Key Performance Indicators

- The combination captures essential Key Performance Indicators (KPIs)
- Gain insight into current and potential future issues by diagnosing the problem faster

* Statement of direction in z15 announcement

Topics

zEDC - BSAM/QSAM Dataset Compression on z15

Integrated Accelerator for zEDC benefits on z15

- Using Integrated Accelerator for zEDC compression with BSAM/QSAM files on z15 can reduce **file size by up to 83%,** while improving **CPU costs by up to 13%** and elapsed times by up to 74% compared to using no compression.
- Integrated Accelerator for zEDC compression on z15 can reduce BSAM/QSAM **file size by up to 6X** compared to not using compression.
- Combining Integrated Accelerator for zEDC compression with BSAM/QSAM file encryption on z15 can improve **elapsed time by up to 72%** while reducing **CPU by up to 7%** compared to not using compression and encryption.

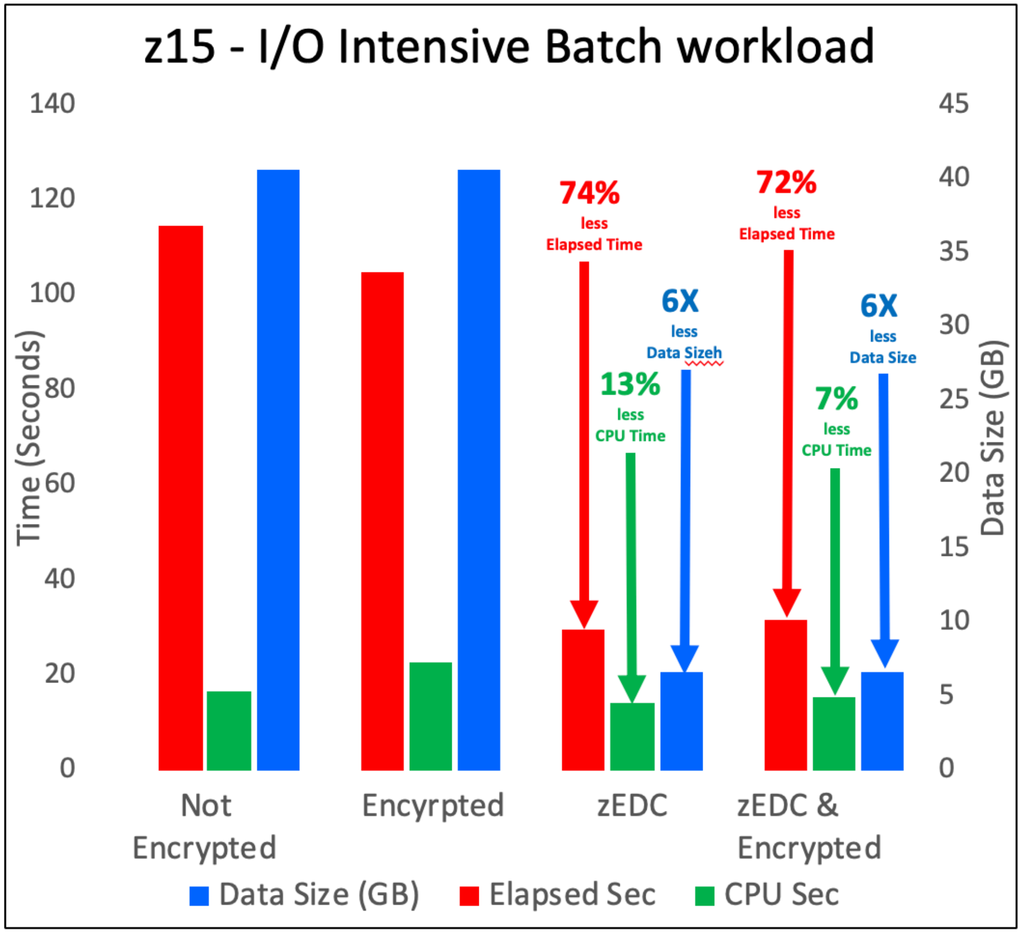

DISCLAIMER: Measurements completed in a controlled environment using a z/OS 2.3 batch workload accessing BSAM and QSAM sequential files . Results may vary by customer based on individual workload, configuration and software levels.

System Recovery Boost

Instant Recovery

Topics

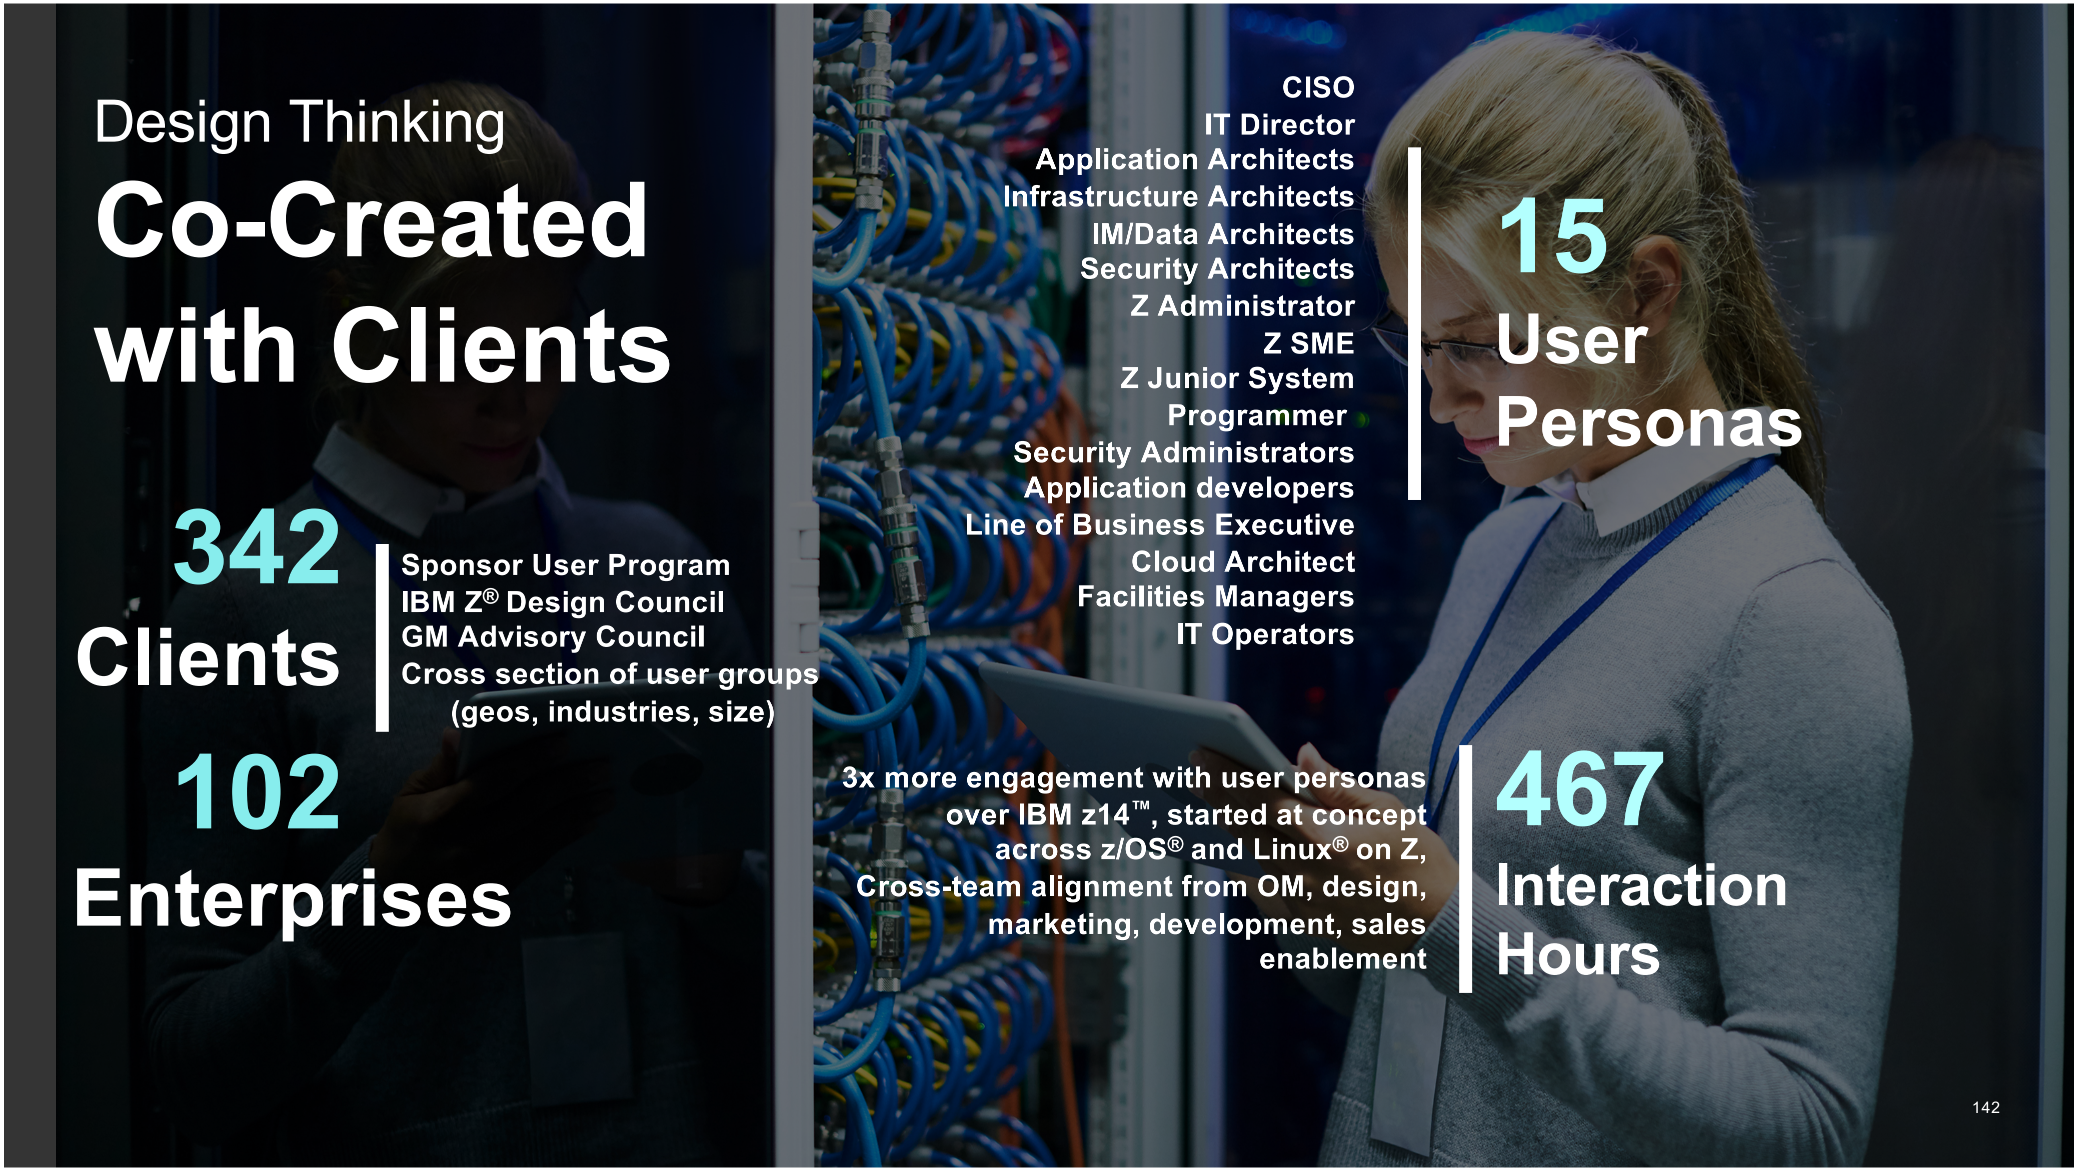
Design Thinking
Co-Created with Clients
342 Clients
Sponsor User Program
IBM Z® Design Council
GM Advisory Council
Cross section of user groups (geos, industries, size)
102 Enterprises
CISO
IT Director
Application Architects
Infrastructure Architects
IM/Data Architects
Security Architects
Z Administrator
Z SME
Z Junior System Programmer
Security Administrators
Application developers
Line of Business Executive
Cloud Architect
Facilities Managers
IT Operators
15 User Personas
3x more engagement with user personas over IBM z14™, started at concept across z/OS® and Linux® on Z,
Cross-team alignment from OM, design, marketing, development, sales enablement
467 Interaction Hours
142

IBM Z – Transformational Resilience

IBM Z is at the forefront to surpass industry availability requirements, maximizing uptime, and empowering your IT system to rapidly and autonomically recover from any disruption.

- Better throughput
- Higher overall server hardware reliability
- Faster recovery when failures occur
- Reduce I/O latency to storage
- Concurrent maintenance improvements
- React faster to workload fluctuations
- Improved workload scaling
- Parallel Sysplex® coupling technology for data sharing and workload balancing

> “ ***IBM Z is in a class of its own: 83% of respondents said their firms achieved five and six nines –99.999% and 99.9999% – or greater uptime.***
>
> ***(ITIC)***

Topics

Anatomy of a planned outage
Visualizing benefits with System Recovery Boost

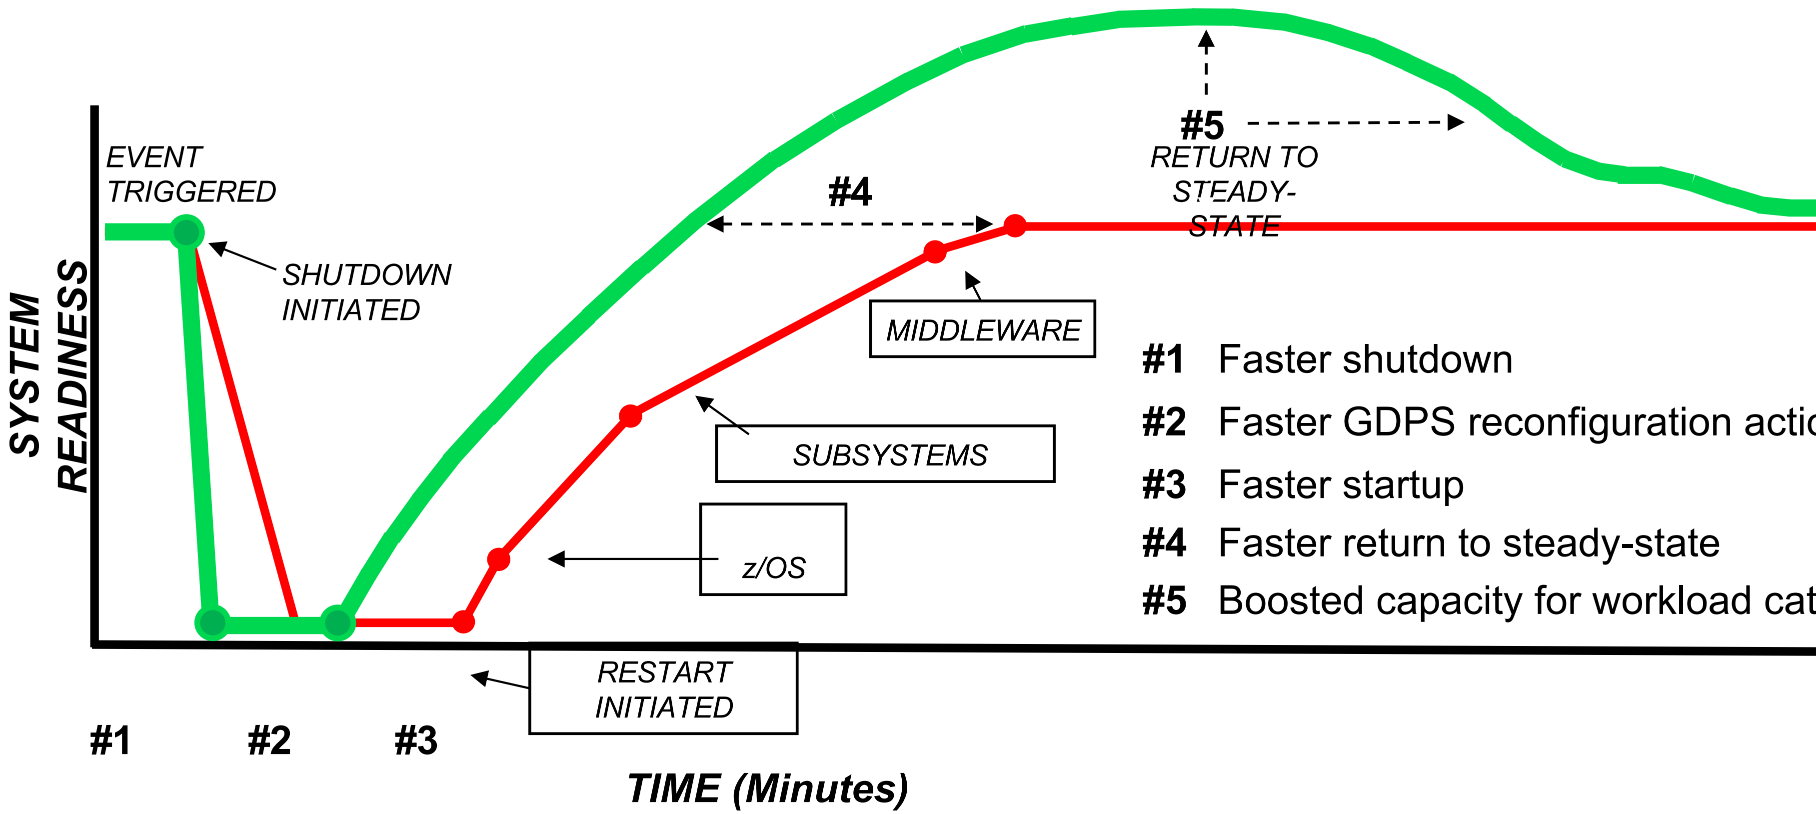

#1 Faster shutdown

#2 Faster GDPS reconfiguration actions

#3 Faster startup

#4 Faster return to steady-state

#5 Boosted capacity for workload catch-up

(Animated Slide – use right arrow key to continue animation)

Topics

IBM System Recovery Boost
Unleash your capacity to maximize your availability

Diminish the impact of any event, planned or unplanned, so you can achieve service level excellence with ***zero increase in IBM software licensing costs***.

Recover workloads substantially faster than on prior Z machines by unleashing additional processing capacity during a fixed-period performance increase on an LPAR-by-LPAR basis.

With System Recovery Boost, you can achieve up to:

2.0x Faster	2.0x Faster	2.5x Faster	2.5x Faster
Return to pre-shutdown service levels	*Processing of transactional backlog*	*Processing of batch backlog*	*GDPS automated startup, shutdown, site switches, etc.*

Performance for the System Recovery Boost period

1 **Processor Capacity Boost using zIIPs**

Provides parallelism and a boost in processor capacity for processing any kind of work during the Boost.

2 **Speed Boost**

Sub-capacity machines gain a boost in processor speed by running the central processors at full-capacity speed during the Boost.

3 **GDPS Reconfiguration**

Increases the speed at which GDPS® drives hardware actions, along with the speed of the underlying hardware services

System Recovery Boost Turbo - Unlock additional "dark cores" for extra zIIP processor capacity

IBM z15 functional comparison to IBM z14

Performance and Scale	Uniprocessor Performance	New up to **14**% performance improvement over IBM z14 (z14)[1]
	System Capacity	New up to **25**% system total z/OS capacity performance improvement over z14[1]
	SMT	New 2nd generation SMT delivers up[to **25%** performance improvements for IFLs and zIIP workloads vs non-SMT on z14
	Cache	New z15 has **86%**more on-chip cache per core versus z14
	Models	Five feature based sizes with up to five CPC drawers (z14 has five models and four drawers)
	Processing cores	New up to 190 cores to configure, up to 170 on z14
	Granular Capacity	New up to 292 capacity settings versus 269 on the z14
	Memory	New up to 40 TB RAIM memory versus 32 TB RAIM memory on z14
	Compression	CMPSC compression and new Integrated Acceleration for ZEDC versus CMPSC and zEDC Express on z14
Virtualization	LPAR virtualization	85 partitions – same as z14
	RoCE adapter	2X the maximum number of RoCE features (up to 16) allowing RoCE to be extended to more workloads versus 8 on z14
	Simplified LPAR management	Enhanced IBM Dynamic Partition Manager allows for config and management of system resources on both
Infrastructure Efficiency	Networking	New OSA-Express7S with improvements over z14 using OSA-Express6S
	HiperSockets® and SMC-D	Up to 32 HiperSockets and memory-to-memory communications with SMC-D offers within-the-box communications for z/OS – same as z14
	FICON	FICON Express16S+ same as IBM z14
	zHPF	zHPF extended distance II offers faster remote site recovery with improved I/O service time improvement when writing data remotely (GDPS® HyperSwap®) same as z14
	IBM zHyperLink	IBM zHyperLink1.1 - short distance z15 channel that can be installed on IBM DS8880 System Storage® for lower latency same on z14
	Forward Error Correction	Industry standard FEC for optical connections for substantially reduced I/O link errors same as z14
	FICON dynamic routing	Dynamic Routing allows for sharing of switches between FICON and FCP without creating separate virtual switches same on z14
	LCSS/Subchannel sets	Up to six LCSS and 4 Subchannel sets – same as z14
	WWPN	I/O serial number migration allows keeping same serial number on replacement server same as z14
	HMC	Next generation HMC with simplified panels, new mobile capabilities, security enhancements (including multi-factor authentication), easier help panels – not on z13. (No Classic Style User Interface on z14)
	Pause-less garbage collection	New enterprise scale Java applications to run without periodic pause for garbage collection on larger & larger heaps same on z14
	IBM Virtual Flash Express	New memory replacement for Flash Express helping improve availability – same as z14
Resiliency and Availability	System Recovery Boost	Enables faster recovery and restoration of service from any planned or unplanned operating system outages not on z14
	Coupling – ICA SR	Short distance coupling with PCIe-based links (ICA SR) – same as z14
	Coupling Express LR	Coupling Express LR – Coupling Express LR will be available on z14
	Coupling scale	50% increase in logical coupling CHPIDs per CPC over z14
	STP	New Simplified STP management with HMC enhancements same as z14
	Sparing	Enhanced integrated sparing on z15 and z13 reducing the number of on site service and maintenance events
	Rack Mounted Accessories	Rack-mounted HMC and TKE options to save space in the data center
	Environmentals	New 19" frame packaging and optional non raised floor, overhead cabling, water cooling, ASHRAE 3 rating - 24" frame packing on z14
	Coupling – HCA-3	No InfiniBand Coupling with HCA-3 InfiniBand Coupling Links on z15 – available on z14
Security	Cryptographic Coprocessor	CPACF for improved performance and true Random Number Generator available on z14
	Crypto Express	New Crypto Express7S with performance increase in accelerator mode plus new algorithms for elliptic curve, SHA, VISA FPE versus z14 Crypto Express6S
	IBM Secure Service Container	Secure deployment of software virtual appliances – available on z14
	Secure Console Access	Protection of sensitive data by using Transport Layer Security (TLS) support in the Open Systems Adapter-Integrated Console Controller (OSA-ICC)

[1] Disclaimer: Based on preliminary internal measurements and projections and compared to the z14. Official performance data will be available upon announce. Results may vary by customer based on individual workload, configuration and software levels. Visit LSPR website for more details at: https://www-304.ibm.com/servers/resourcelink/lib03060.nsf/pages/lsprindex .

Topics

IBM z15 functional comparison to IBM z13

Performance and Scale	Uniprocessor Performance	New up to **25%** performance improvement over IBM z13 (z13)[1]
	System Capacity	New up to **63%** system total z/OS capacity performance improvement over z13[1]
	SMT	New 2nd generation SMT delivers up to **56%** performance improvement for IFLs and zIIP workloads vs non-SMT on z13
	Cache	New z15 has **263%** more on-chip cache per core versus z13
	Models	Five feature based sizes with up to five CPC drawers (z13 has five models and four drawers)
	Processing cores	New up to 190 cores to configure, up to 141 on z13
	Granular Capacity	New up to 292 capacity settings versus 231 on the z13
	Memory	New up to 40 TB RAIM memory versus 10 TB RAIM memory on z13
	Compression	CMPSC compression and new Integrated Acceleration for ZEDC versus CMPSC compression and zEDC Express on z13
Virtualization	LPAR virtualization	85 partitions – same as z13
	RoCE adapter	2X the maximum number of RoCE features (up to 16) allowing RoCE to be extended to more workloads vs 8 on z13 (also no 25 GbE Express2.1 on z13)
	Simplified LPAR management	Enhanced IBM Dynamic Partition Manager allows for config and management of system resources on both
Infrastructure Efficiency	Networking	New OSA-Express7S with improvements over z13 using OSA-Express5S
	HiperSockets and SMC-D	Up to 32 HiperSockets and memory-to-memory communications with SMC-D offers within-the-box communications for z/OS – same as z13
	FICON	FICON Express16S+ versus FICON Express16S on z13
	zHPF	zHPF extended distance II offers faster remote site recovery with improved I/O service time improvement when writing data remotely (GDPS® HyperSwap®) same as z13
	IBM zHyperLink	IBM zHyperLink1.1 - short distance z15 channel that can be installed on IBM DS8880 System Storage for lower latency not on z13
	Forward Error Correction	Industry standard FEC for optical connections for substantially reduced I/O link errors same as z13
	FICON dynamic routing	Dynamic Routing allows for sharing of switches between FICON and FCP without creating separate virtual switches same on z13
	LCSS/Subchannel sets	Up to six LCSS and 4 Subchannel sets – same as z13
	WWPN	I/O serial number migration allows keeping same serial number on replacement server same as z13
	HMC	Next generation HMC with simplified panels, new mobile capabilities, security enhancements (including multi-factor authentication), easier help panels – not on z13. (No Classic Style User Interface on z13)
	Pause-less garbage collection	New enterprise scale Java applications to run without periodic pause for garbage collection on larger & larger heaps not on z13
	IBM Virtual Flash Express	New memory replacement for Flash Express helping improve availability – not available on z13
Resiliency and Availability	System Recovery Boost	Enables faster recovery and restoration of service from any planned or unplanned operating system outages not on z13
	Coupling – ICA SR	Short distance coupling with PCIe-based links (ICA SR) – same as z13
	Coupling Express LR	Coupling Express LR – Coupling Express LR will be available on z13
	Coupling scale	50% increase in logical coupling CHPIDs per CPC over z13
	STP	New Simplified STP management with HMC enhancements not available on z13
	Sparing	Enhanced integrated sparing on z15 and z14 reducing the number of on site service and maintenance events
	Rack Mounted Accessories	Rack-mounted HMC and TKE options to save space in the data center
	Environmentals	New 19" frame packaging and optional non raised floor, overhead cabling, water cooling ASHRAE 3 rating - 24" frame packing and no ASHRAE 3 on z13
	Coupling – HCA-3	No InfiniBand Coupling with HCA-3 InfiniBand Coupling Links on z15 – available on z13
Security	Cryptographic Coprocessor	CPACF for improved performance and true Random Number Generator versus z13
	Crypto Express	New Crypto Express7S with **a** performance increase in accelerator mode plus new algorithms for elliptic curve, SHA, VISA FPE versus z13 Crypto Express5S
	IBM Secure Service Container	Secure deployment of software virtual appliances – available on z13
	Secure Console Access	Protection of sensitive data by using Transport Layer Security (TLS) support in the Open Systems Adapter-Integrated Console Controller (OSA-ICC)

[1] Disclaimer: Based on preliminary internal measurements and projections and compared to the z14. Official performance data will be available upon announce. Results may vary by customer based on individual workload, configuration and software levels. Visit LSPR website for more details at: https://www-304.ibm.com/servers/resourcelink/lib03060.nsf/pages/lsprindex.

The End is Near

THE END IS INEVITABLE

The Last Page

Questions?
Comments?
Requests?

THANK YOU for your Time

Topics