

IBM Z Full Stack Solutions Leveraging the IBM z17 Server

Mark Gambino, z/TPF Chief Architect
IBM Distinguished Engineer

2025 TPF Users Group Conference
May 5-7, Austin, TX

IBM Z



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.



An Exchange

Statement

“I run my workload on **x86** servers”



Response

“Why are you using an operating system and database method that are decades old running on ancient HW technology that runs at only **5 MHz** with **16-bit addressing**?”

An Exchange

Statement

“I run my workload on **x86** servers”



Response

“Why are you using an operating system and database method that are decades old running on ancient HW technology that runs at only **5 MHz** with **16-bit addressing**?”



Rebuttal

“Wow, your assumptions are just wrong. Yes, **x86** technology was around in the 1970’s, but why on earth would you think I’m using **8086** cores today?? My modern servers are running at GHz speed using 64-bit addressing. Also, while **Linux** and **SQL** were created decades ago, they have evolved over time.”

A Second Exchange

Statement

“I run my workload
on **IBM Z** servers”



Response

“Why are you using an
operating system and
database method that
are decades old
running on ancient HW
technology that runs at
only **17 MHz** with
24-bit addressing?”

A Second Exchange

Statement

“I run my workload on **IBM Z** servers”



Response

“Why are you using an operating system and database method that are decades old running on ancient HW technology that runs at only **17 MHz** with **24-bit addressing**?”



Rebuttal

“Wow, your assumptions are just wrong. Yes, **IBM Z** technology was around in the 1970's, but why on earth would you think I'm using an **IBM 3033** today?? My modern servers are running at GHz speed using 64-bit addressing. Also, while **TPF** and **TPFDF** were created decades ago, they have evolved over time.”

Two Server Architectures Have Survived, and There Are Key Differences Between Them

	IBM Z	x86
Designed for	High-performance, multiple enterprise-level workloads with a focus on extreme scale and security	General purpose, each server running a single and smaller workload
HW Reliability and Redundancy	High levels of built-in redundancy and fault tolerance, ensuring continuous operation even during hardware failures	Commodity HW
Operating Systems Supported	z/TPF , z/OS, z/VM, Linux	Linux , Windows
Typical Architecture	Applications and database co-located on a small number of physical servers	Separate application and database tiers on 100's to 1000's of physical servers

Integrated Design versus Piecemeal Plug-And-Pray

	z/TPF on IBM Z	Linux on x86
Operating System	IBM (z/TPF)	Vendor 1
Middleware	IBM (included with z/TPF)	Vendors 2, 3, and 4
Database	IBM (included with z/TPF)	Vendors 5 and 6
Languages Environments	IBM (included with z/TPF)	Vendors 7 and 8
Server HW	IBM (IBM Z)	Vendors 9 and 10
Hypervisor	IBM (included with IBM Z)	Vendor 11
HW Accelerators	IBM (included with IBM Z)	Vendors 12 and 13
Storage	IBM or other vendors	Vendor 14
Platform and Application Monitoring	IBM or other vendors	Vendors 15, 16, and 17

Integrated Design versus Piecemeal Plug-And-Pray

	z/TPF on IBM Z	Linux on x86
Operating System	IBM (z/TPF)	Vendor 1
Middleware	IBM (included with z/TPF)	Vendors 2, 3, and 4
Database	IBM (included with z/TPF)	Vendors 5 and 6
Languages Environments	IBM (included with z/TPF)	Vendors 7 and 8
Server HW	IBM (IBM Z)	Vendors 9 and 10
Hypervisor	IBM (included with IBM Z)	Vendor 11
HW Accelerators	IBM (included with IBM Z)	Vendors 12 and 13
Storage	IBM or other vendors	Vendor 14
Platform and Application Monitoring	IBM or other vendors	Vendors 15, 16, and 17

IBM full stack solution designed specifically for your production workloads

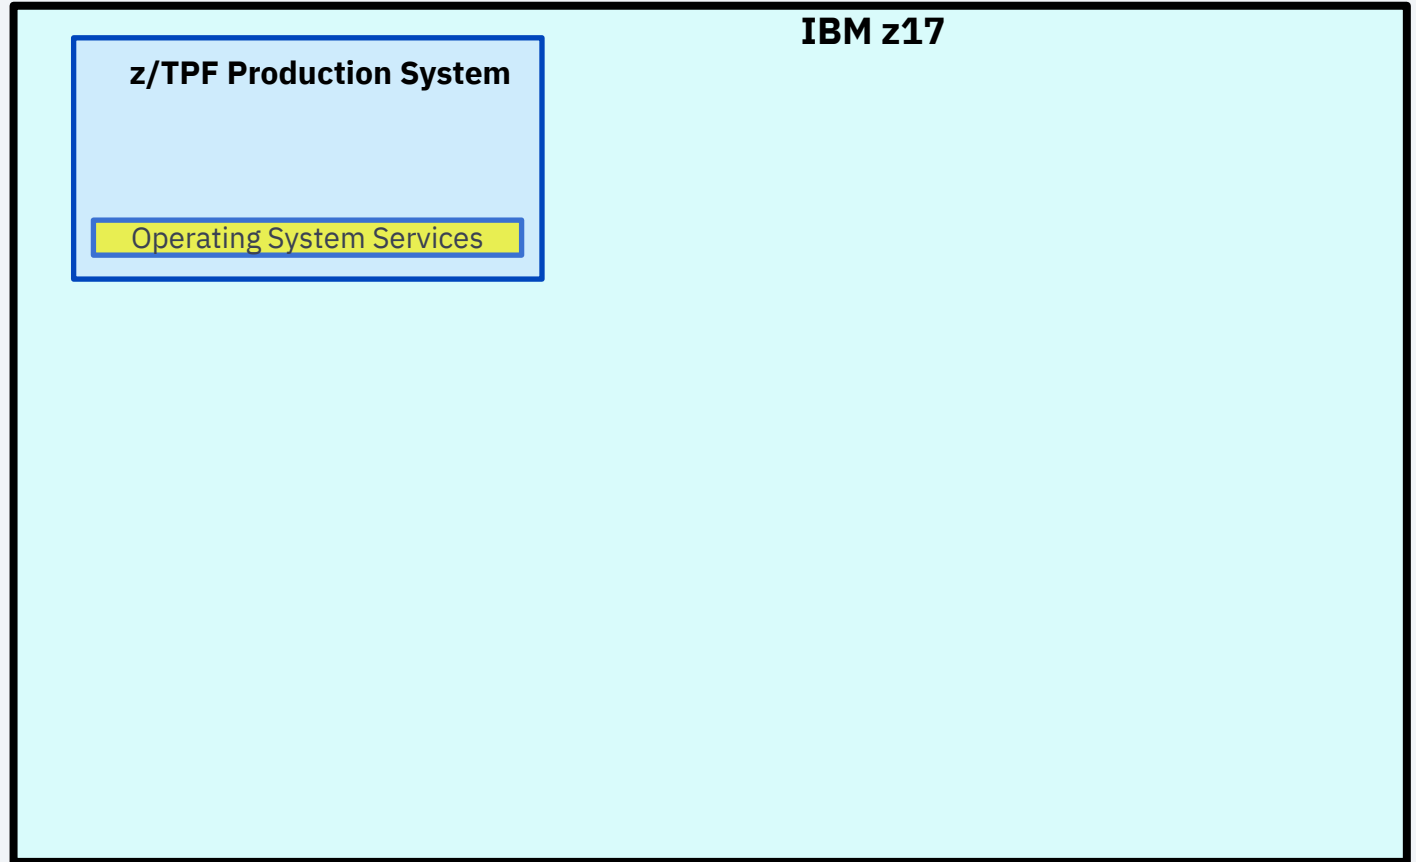
Integrated Design versus Piecemeal Plug-And-Pray

	z/TPF on IBM Z	Linux on x86
Operating System	IBM (z/TPF)	Vendor 1
Middleware	IBM (included with z/TPF)	Vendors 2, 3, and 4
Database	IBM (included with z/TPF)	Vendors 5 and 6
Languages Environments	IBM (included with z/TPF)	Vendors 7 and 8
Server HW	IBM (IBM Z)	Vendors 9 and 10
Hypervisor	IBM (included with IBM Z)	Vendor 11
HW Accelerators	IBM (included with IBM Z)	Vendors 12 and 13
Storage	IBM or other vendors	Vendor 14
Platform and Application Monitoring	IBM or other vendors	Vendors 15, 16, and 17

IBM full stack solution ensures HW and SW compatibility

Which HW and SW versions are compatible across all these vendors?

IBM Z Integrated Full Stack Solution



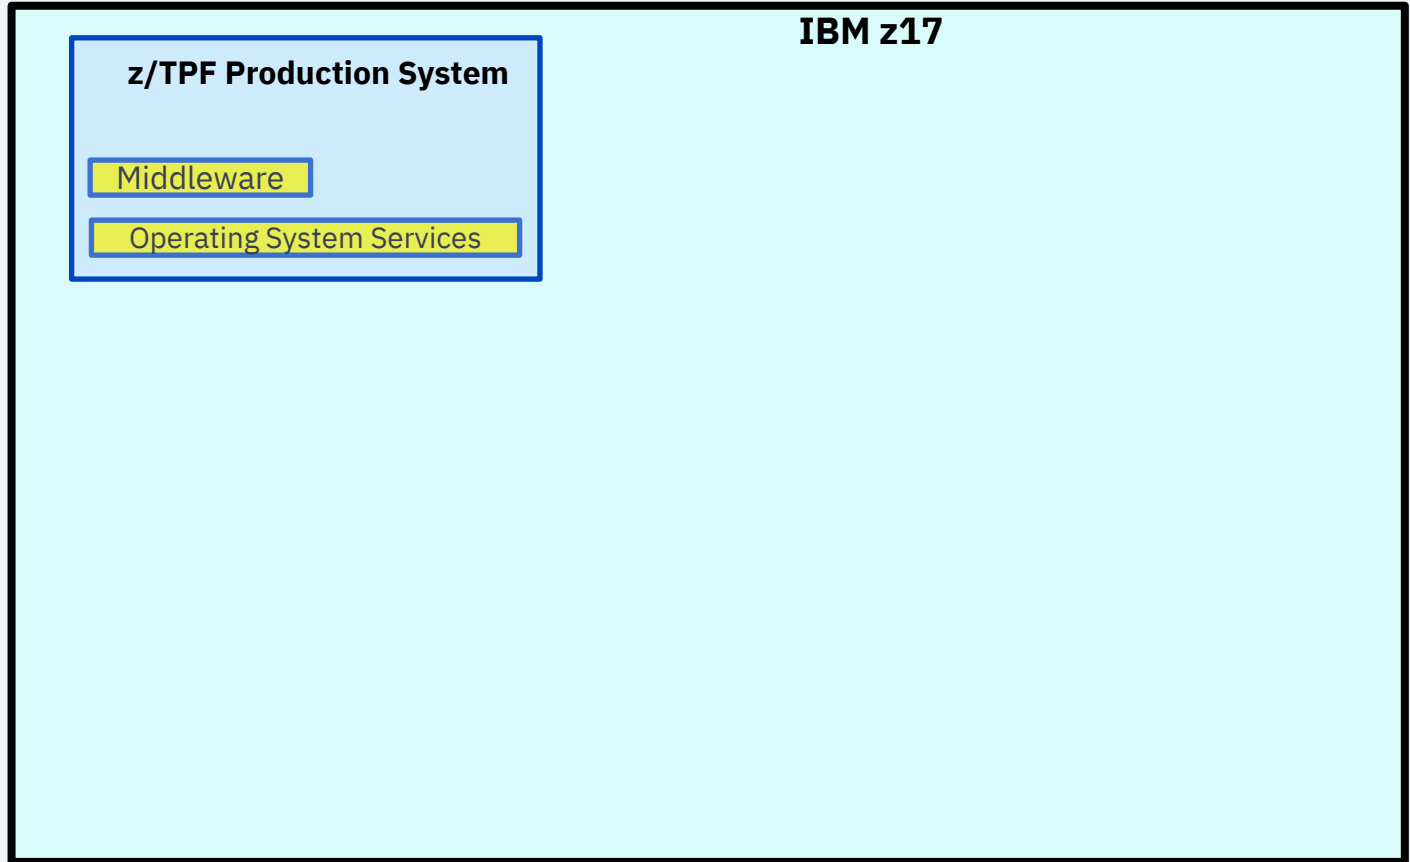
Operating System Services

High performance services using a mix of open standards and value-add z/TPF services

IBM Z Integrated Full Stack Solution

Middleware Included

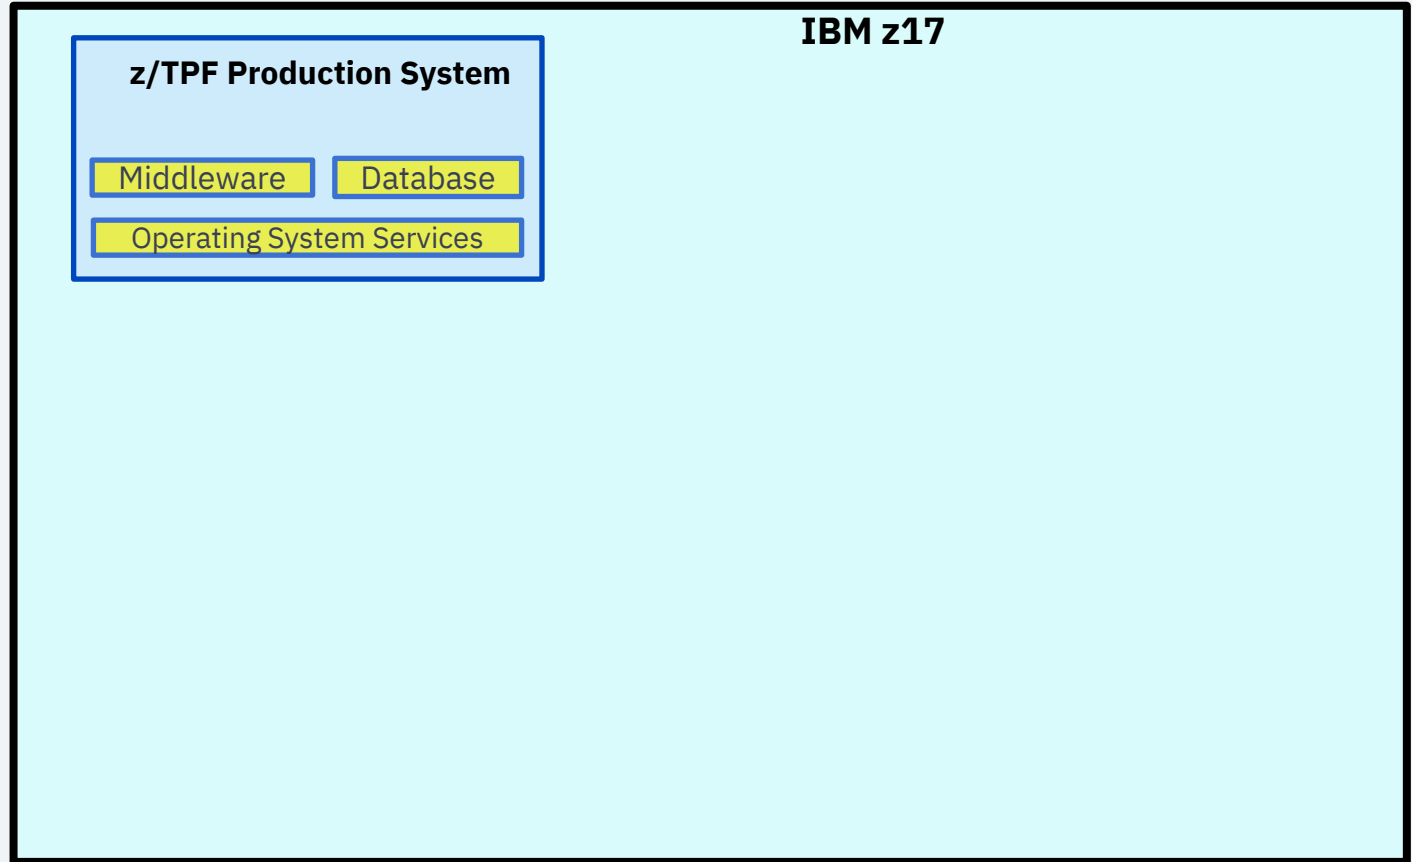
Communicate with the rest of your hybrid cloud and millions of end users using **secure standard** protocols (REST, HTTP, MQ) and data formats (JSON, XML)



IBM Z Integrated Full Stack Solution

Database Access Services Included

Access and update millions of database records per second in a secure and **always consistent** manner



IBM Z Integrated Full Stack Solution

Your Applications

Transactions seamlessly and efficiently flow between programming languages (assembler, C, C++, Java) allowing for **progressive app modernization**

IBM z17

z/TPF Production System

Applications

Middleware

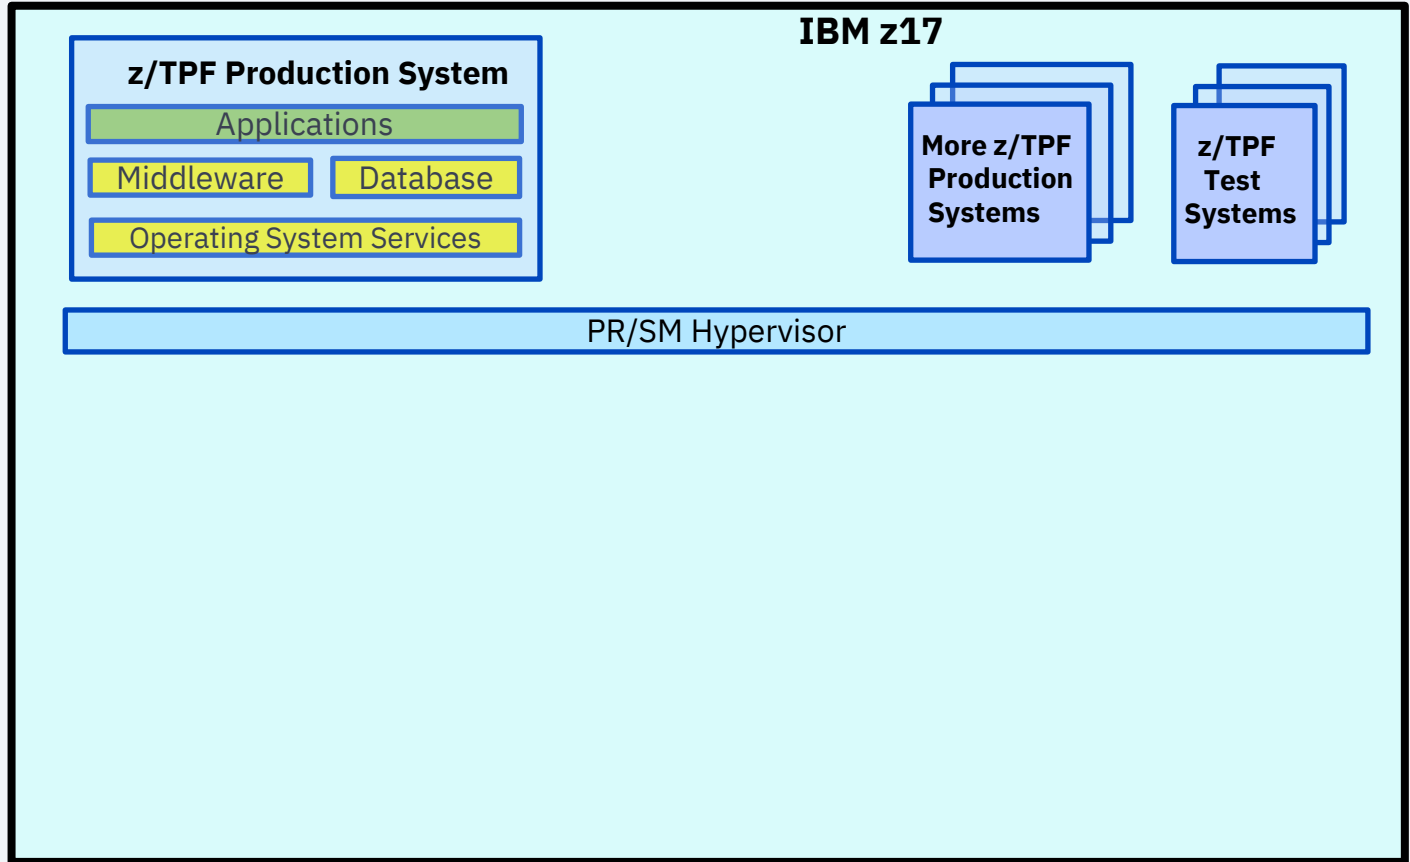
Database

Operating System Services

IBM Z Integrated Full Stack Solution

PR/SM Hypervisor

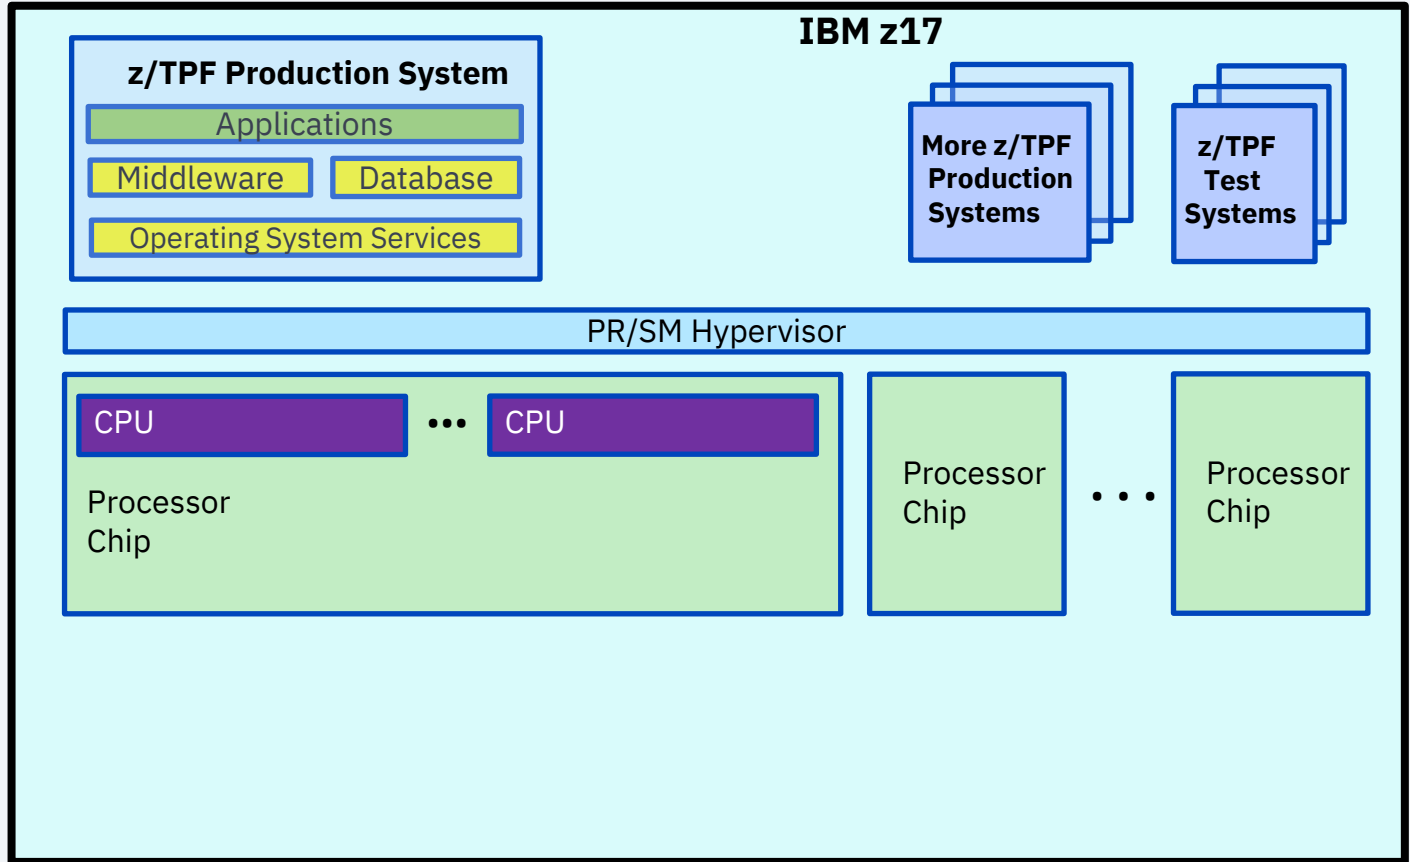
Intelligent sharing of CPU cores by multiple servers (LPARs) to make optimal use of your CPU resources



IBM Z Integrated Full Stack Solution

CPU Cores

Up to 208 high
frequency (5.5 GHz)
CPU cores on z17



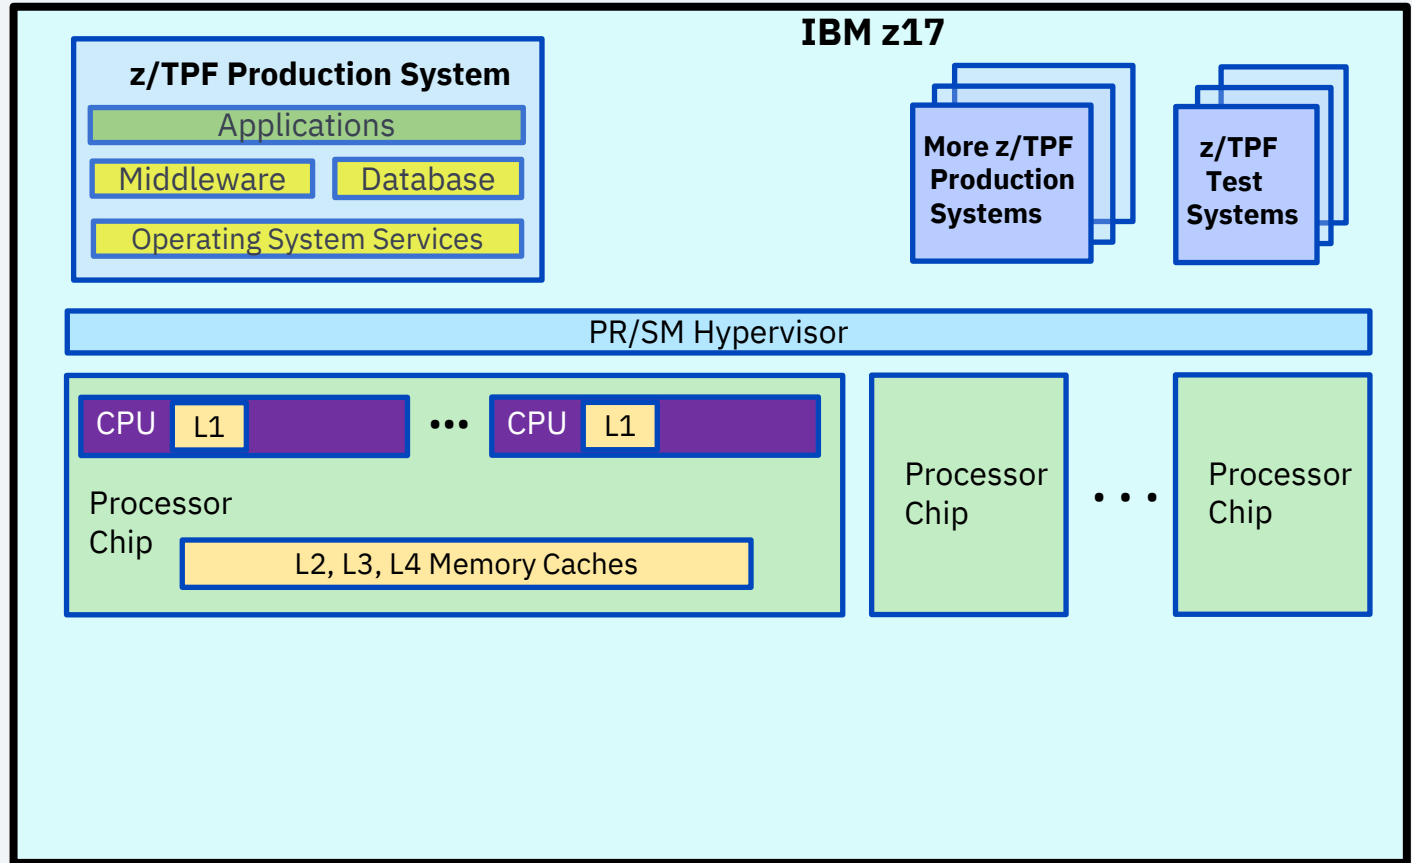
High Frequency CPU Cores Are Nice, But...

- A faster CPU has limited value if it cannot access the data it needs in a timely manner
- Accessing main memory is *slow*; therefore, HW memory caches are essential for real-time data-rich workloads.
- Bigger is better when it comes to memory caches
- Another critical factor is the memory coherency manager firmware design:
 - Determines which data to cast out of cache
 - Guarantees consistency of data across all memory caches and main memory
 - Algorithm for how to handle very frequently updated blocks of memory

IBM Z Integrated Full Stack Solution

Memory Caches

Best in industry multi-layered design to maximize use of cache resources (your own core cache, caches on your processor chip, and caches in other processor chips in the drawer)



What an “AI” Search Engine Said

Q: What is the largest L2 cache?



What an “AI” Search Engine Said

Q: What is the largest L2 cache?

A: L2 cache is a bit slower to access than the L1 cache, but the trade off is that it is much, much larger—on the order of **an entire megabyte** on Zen 4 and a **full two megabytes** on Raptor Cove.

What an “AI” Search Engine Said

Q: What is the largest L2 cache?

A: **L2 cache** is a bit slower to access than the L1 cache, but the trade off is that it is much, much larger—on the order of **an entire megabyte** on Zen 4 and a **full two megabytes** on Raptor Cove.

Reality

	Zen 4	Raptor Cove	IBM z16
L1	I-32 KB D-32 KB	I-32 KB D-48 KB	I-128 KB D-128 KB
L2	1 MB	2 MB	32 MB
L3	32 MB	36 MB	256 MB
L4	None	None	2048 MB (2 GB)
Memory	256 GB	256 GB	40,000 GB (40 TB)

Teaser: These got even bigger with z17!

How Important is an L4 Memory Cache?

Simulated a z/TPF reservations workload running on an IBM z15 server with its L4 memory cache disabled...

How Important is an L4 Memory Cache?

Simulated a z/TPF reservations workload running on an IBM z15 server with its L4 memory cache disabled...

- **30% loss of performance!**

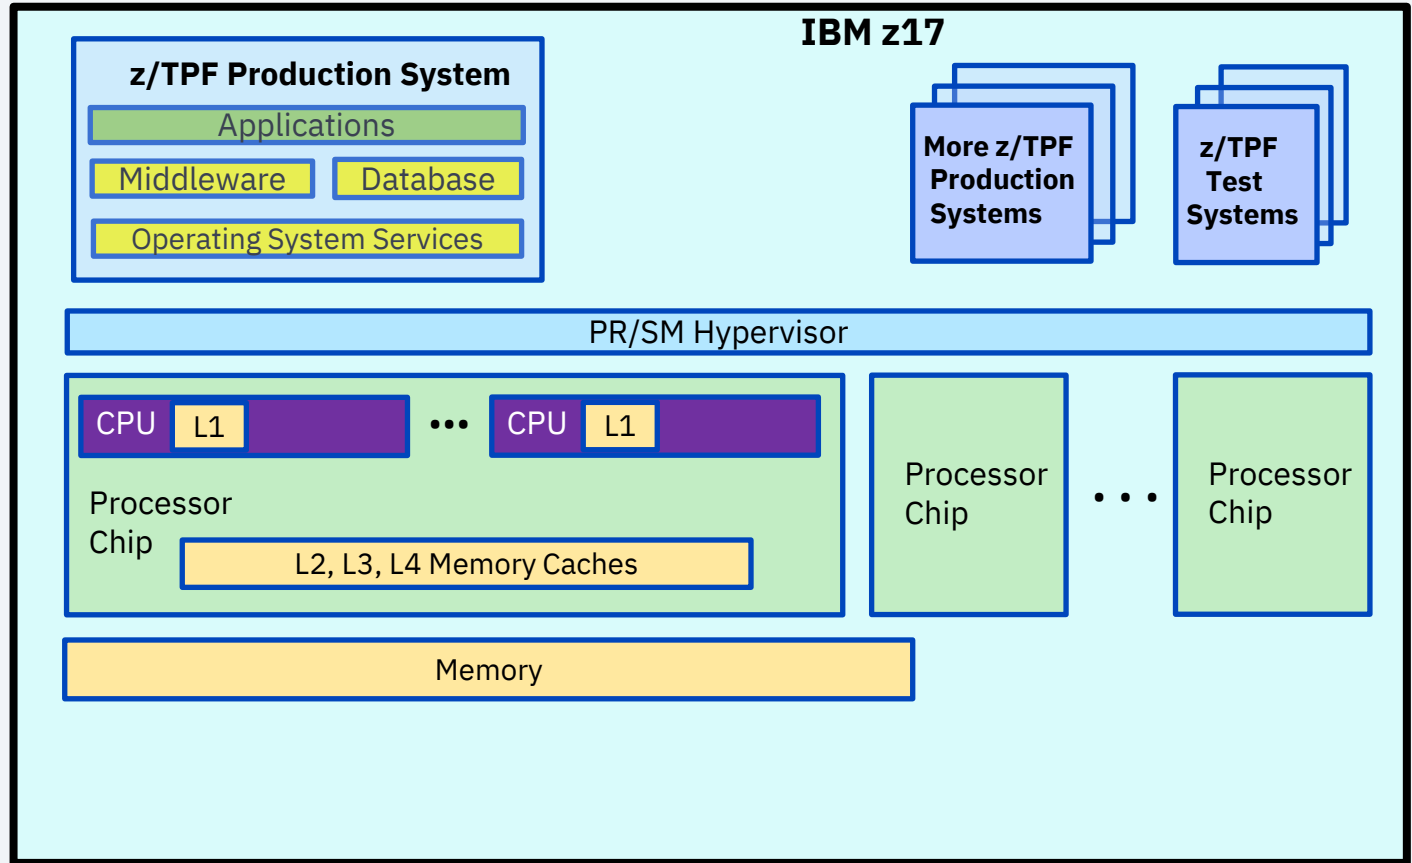
Having to access more data from main memory on x86 because no L4 cache is bad enough, but because Linux does memory paging, **accessing memory might require an I/O operation and be even slower!**

- There is no memory paging on z/TPF (**no I/O ever to access memory**)

IBM Z Integrated Full Stack Solution

Memory

Industry best up to **64 terabytes (TB)** on z17 enabling you to keep an enormous volume of data in memory that is available to your applications with **zero I/O latency**

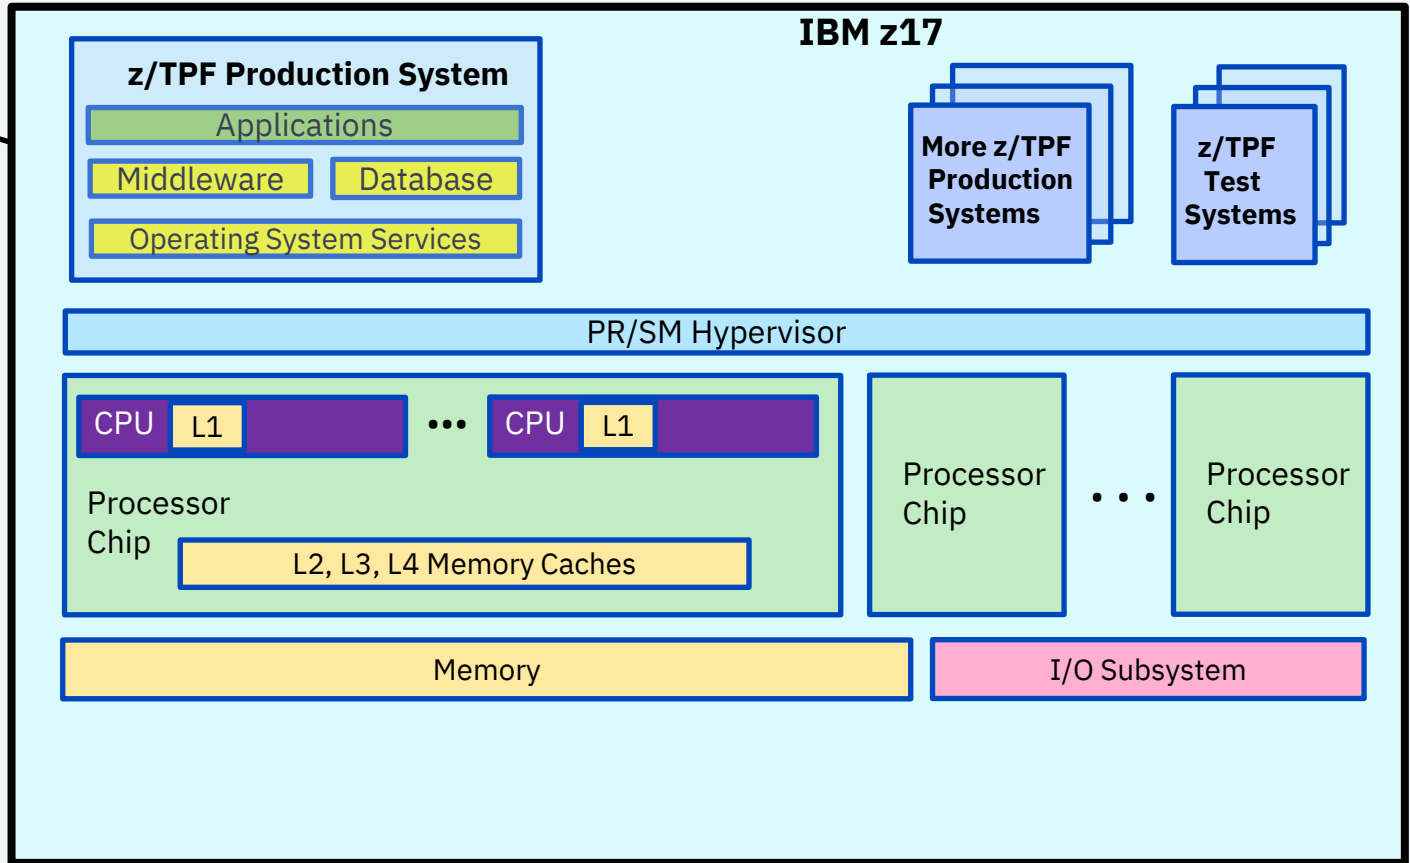


IBM Z Integrated Full Stack Solution

IBM DS8000 G10 Storage
with Safeguarded Copy

I/O Subsystem

Consistent and low response times, even at millions of I/O operations per second (IOPS), accessing and updating secure data in an always consistent manner



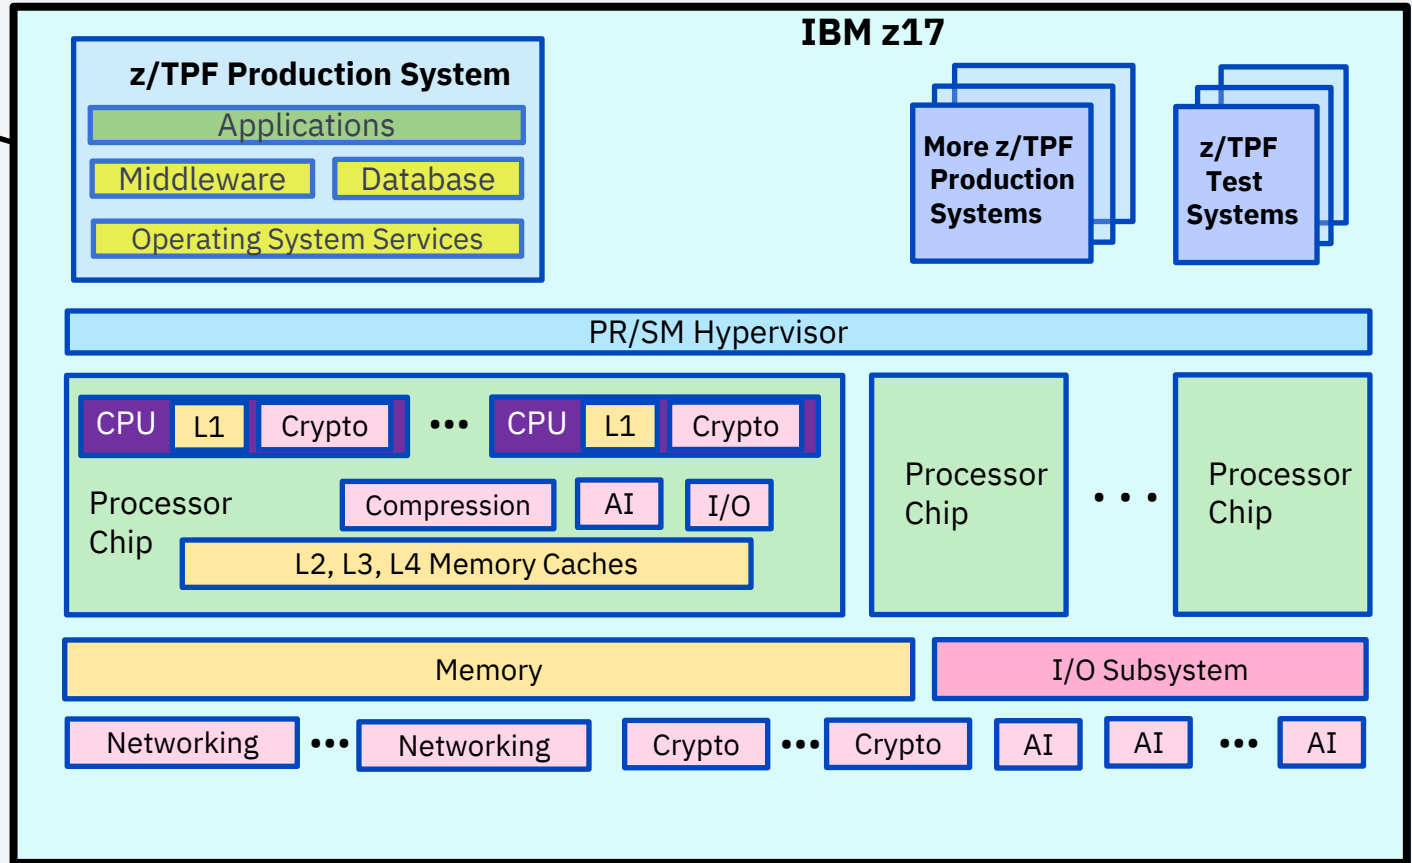
IBM Z Integrated Full Stack Solution

IBM DS8000 G10 Storage
with Safeguarded Copy

HW Accelerators

Specialized HW in the core, processor chip, and server to do I/O, networking, crypto, and AI at scale.

A new on-chip I/O accelerator, two types of AI HW accelerators and a new network adapter on z17.

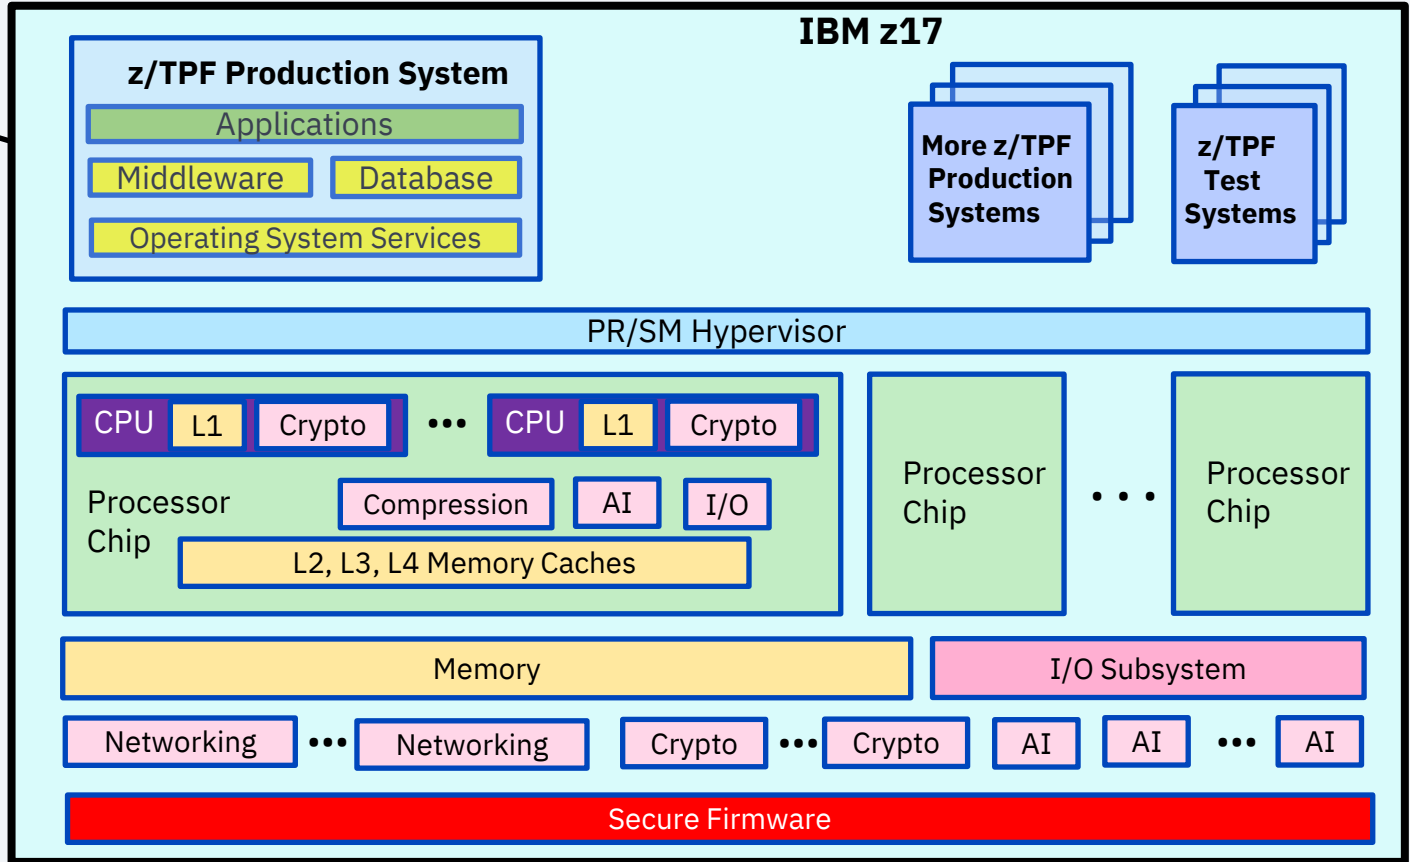


IBM Z Integrated Full Stack Solution

IBM DS8000 G10 Storage
with Safeguarded Copy

Secure Firmware

First introduced on z16 where FW is dual-signed, including via a quantum-safe algorithm, so you are assured the processor code you are running is authentic IBM code.

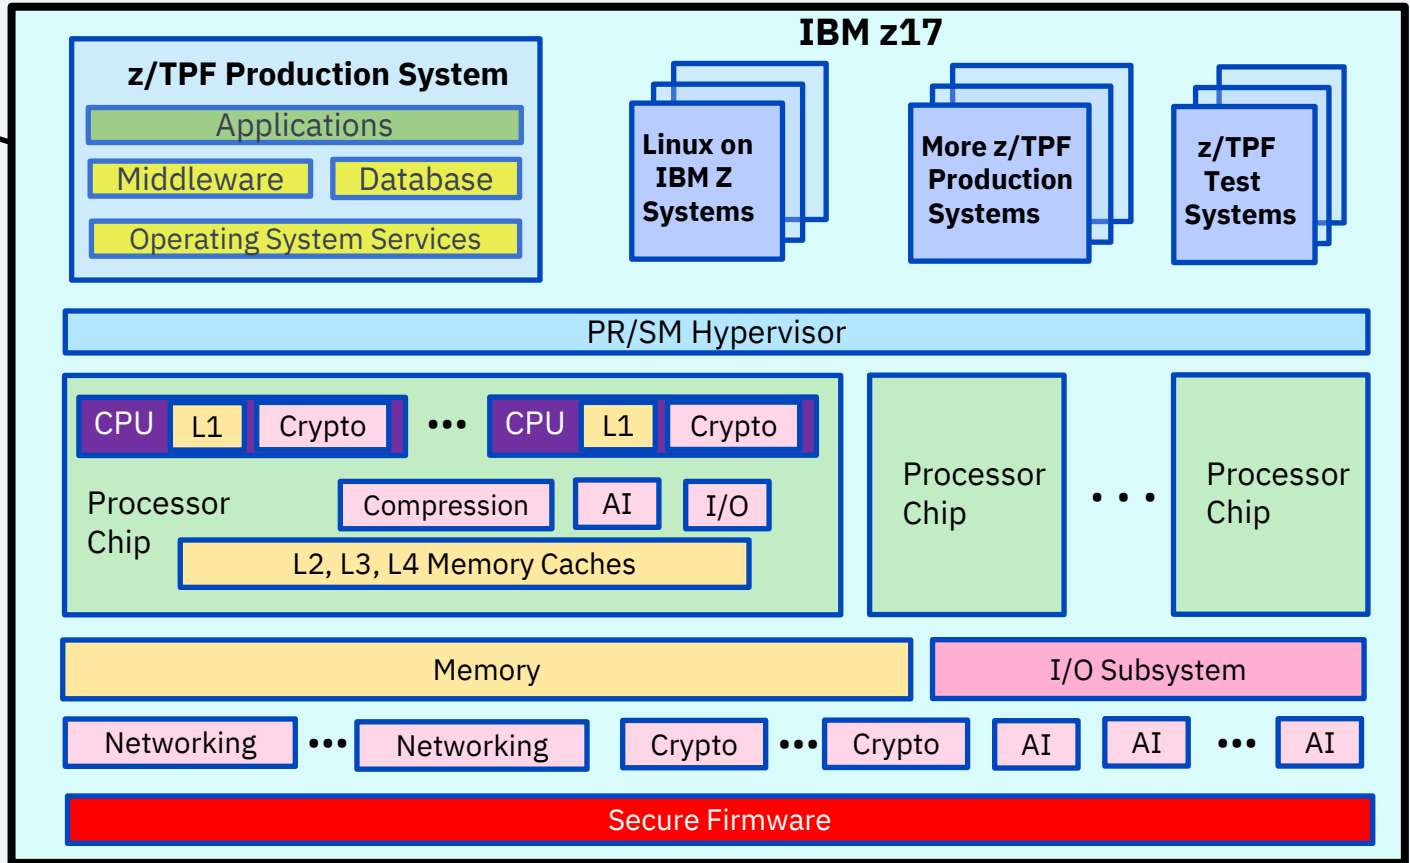


IBM Z Integrated Full Stack Solution

IBM DS8000 G10 Storage
with Safeguarded Copy

Linux on IBM Z

Co-located cooperative processing like fraud detection in real-time using popular AI frameworks such as TensorFlow and PyTorch, leveraging AI HW accelerators.



IBM Z Integrated Full Stack Solution

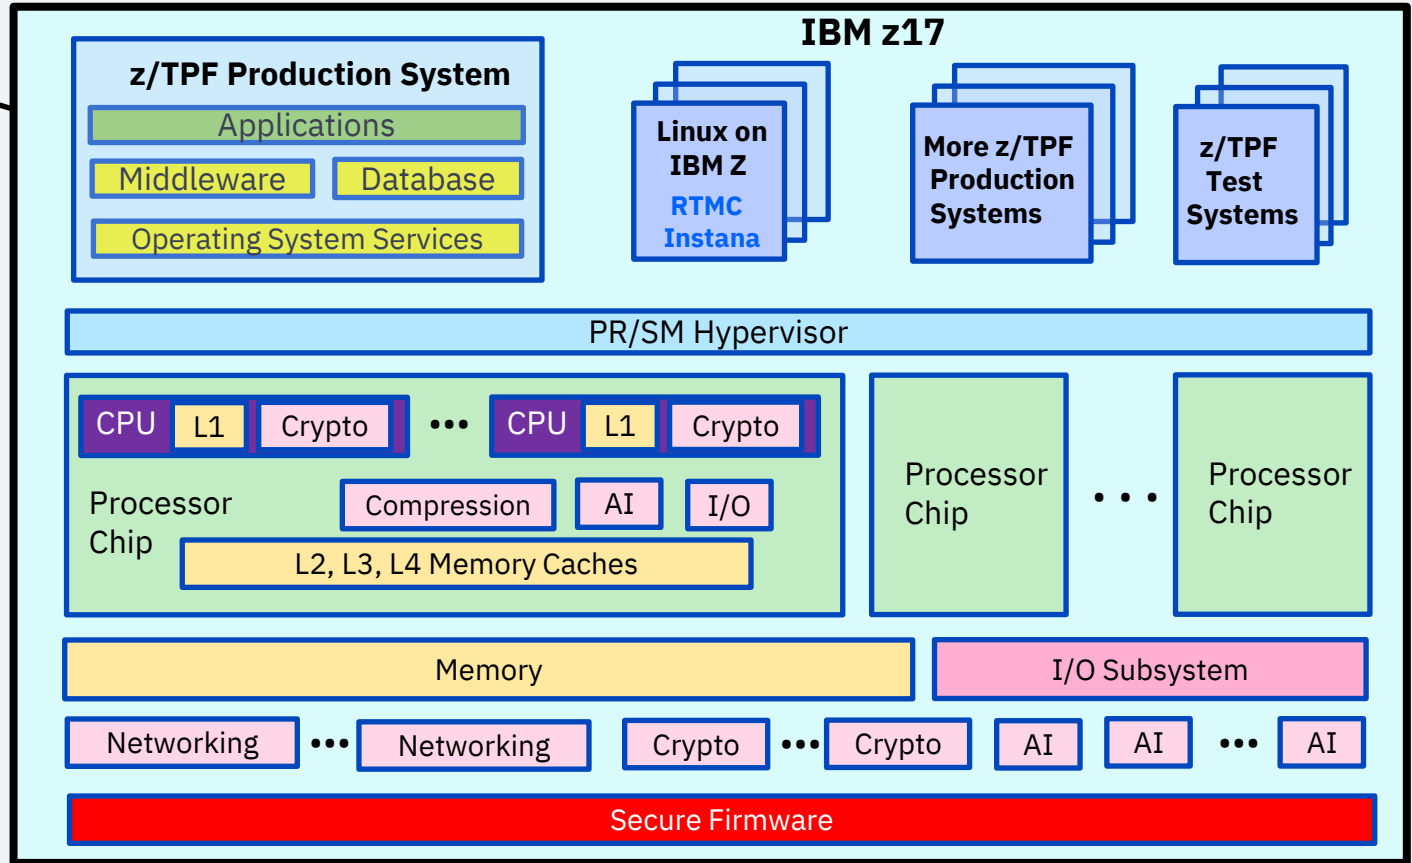
IBM DS8000 G10 Storage
with Safeguarded Copy

Linux on IBM Z


Real-time operational and business analytics using **RTMC**.

Extend RTMC with AI and ML using popular AI frameworks leveraging AI HW accelerators.

IBM **Instana** Application Performance Monitor (APM) end-to-end monitoring has visibility to z/TPF metrics and trace data via OpenTelemetry.



If I Had A



For every time I've heard ...

“If I convert my z/TPF application to a higher level language, then I can just move it to the cloud.”

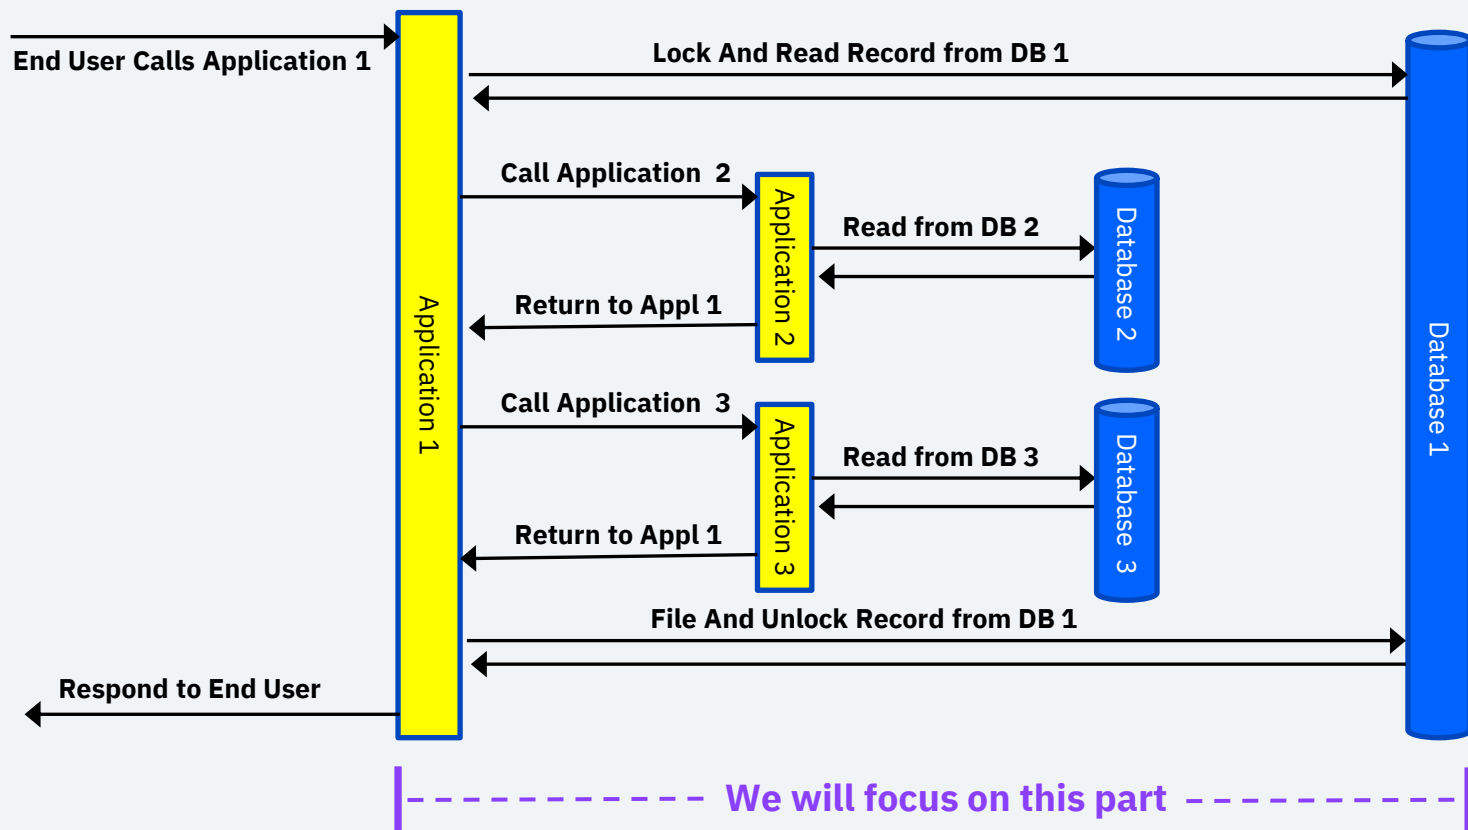
I could buy a ...




Because ...

**Replatforming = rearchitecture and rewrite of the application,
and after all that will it even perform or scale?**

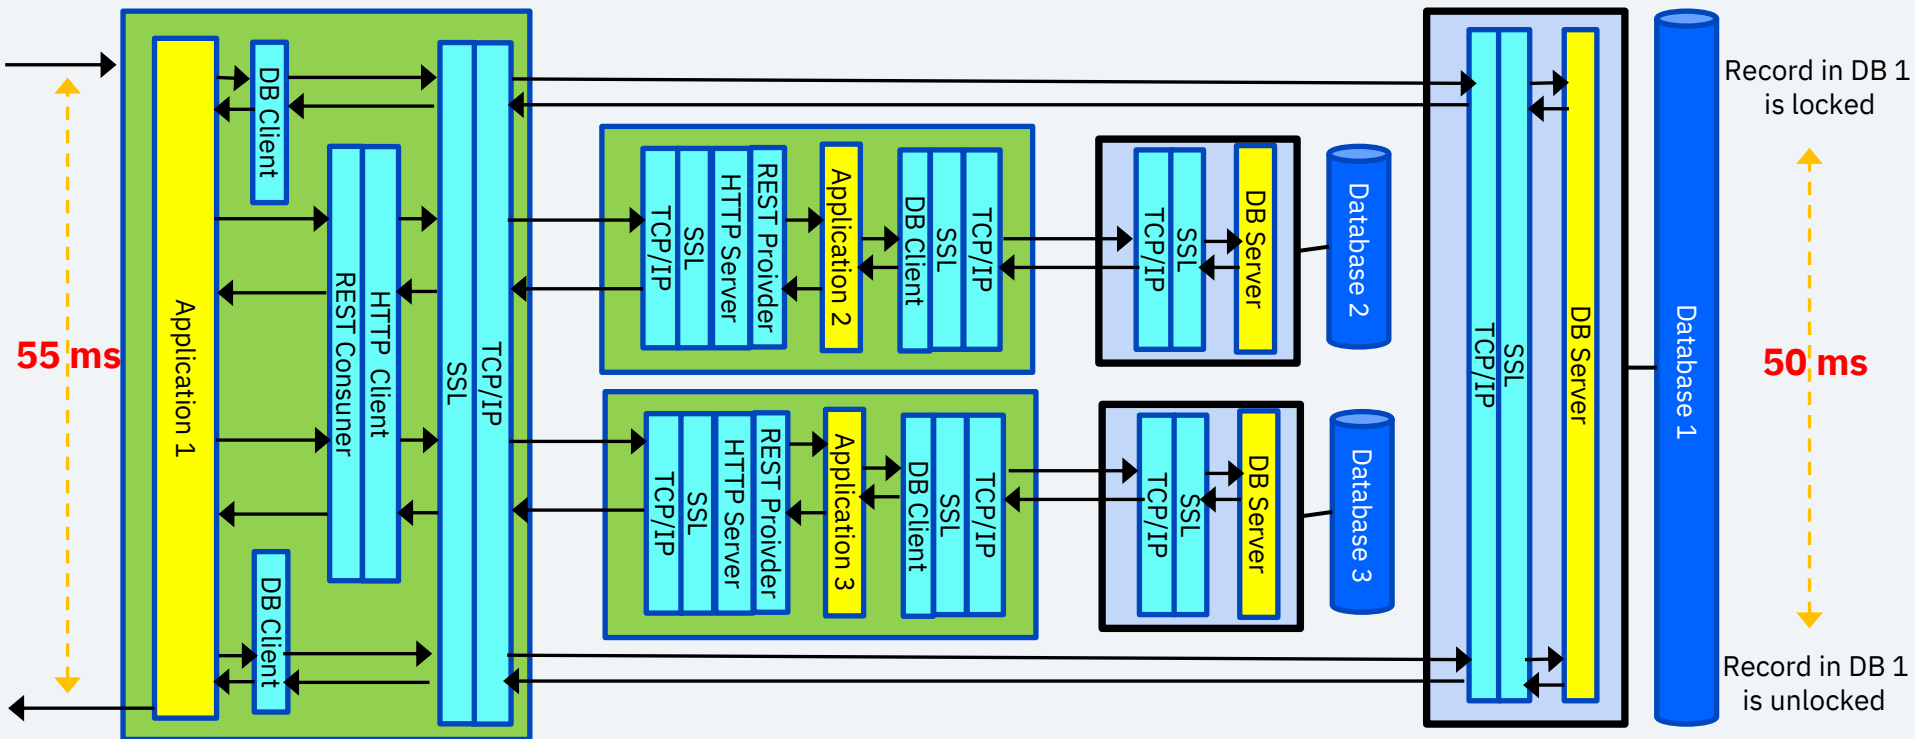
Let's Look at a High-Level Flow Diagram for a Sample Transaction




Typical x86 Implementation – Separate Server Clusters for Each Application and Each Database



Typical x86 Implementation – Latency, CPU Overhead, and DB Lock Hold Times Limit Ability to Scale




z/TPF Implementation – MUCH More Efficient!



Applications invoke each other via
intra-program-to-program calls

z/TPF Implementation – MUCH More Efficient!




Applications invoke each other via
intra-process program-to-program calls

	x86	z/TPF
Transaction Response Time	55 ms	2.2 ms
External Network Flows*	12	0
Protocol Stacks Traversed*	72	0
Database Lock Hold Time	50 ms	2 ms
Maximum Messages per Second For That Database Record	20	500

*** Both of these result in higher latency, higher CPU consumption, and inability to meet response time SLAs**

z/TPF Implementation – MUCH More Efficient!



Applications invoke each other via
intra-process program-to-program calls

	x86	z/TPF
Transaction Response Time	55 ms	2.2 ms
External Network Flows*	12	0
Protocol Stacks Traversed*	72	0
Database Lock Hold Time*	50 ms	2 ms
Maximum Messages per Second For That Database Record	20	500

*** All of these can prevent a workload
from being able to scale**

IBM Z Full Stack Differentiators Beyond Performance

- **Security** – pervasive encryption, quantum-safe firmware
- **Resiliency** - spare physical cores in IBM Z machines take over if an in-use core fails, redundant array of independent memory (RAIM), drawer failover
- **Capacity on Demand** - dark cores exist on most IBM Z machines that can be enabled to add CPU capacity on your existing IBM Z servers
- **Electrical Costs** - x86 cores run hot and GPUs much hotter. IBM Z cores and AI accelerators consume considerably less electricity
- **Floor Space** – consolidate workloads from thousands of x86 cores into a single IBM Z server
- **Simplification** – far fewer servers means less components to configure, monitor, manage, and fewer things that could fail or “hiccup”
- **Scalable Architecture** – scale up for the most demanding workloads and deploy right-sized solutions for smaller geo's with lower volume workloads

And Now Introducing the IBM z17 Server



IBM Z HW Acceleration “Recent” History


- **z14 – Pervasive Encryption**
 - Enough CPACF HW crypto capability to encrypt all data at rest and in flight
- **z15 – Pervasive Compression**
 - On chip HW compression accelerator capable of 12 gigaBYTES per second
- **z15 – Latest TLS Standard at Scale**
 - Ephemeral Elliptic Curve Cryptography (ECC) operations in CPACF HW
- **z16 – Real-Time Inferencing at Scale**
 - First in industry on chip AI accelerator (AIU)

IBM z17 HW Acceleration Innovations

- **New:** On-Chip I/O Accelerator
- **Enhanced:** On-Chip AI Accelerator (AIU) for Inferencing
- **New:** AI Accelerator Card for Gen AI and Complex Models

Data Processing Unit (DPU) – I/O Accelerator

- Functionality from I/O Adapters' application specific integrated circuits (ASICs) has moved to the DPU on processor chip
- Uses higher I/O density 4-port FICON® cards
 - New FICON-Express32-4P adapters
 - z/TPF APAR PJ48194 (May 2025) is required for data collection to recognize the new adapters
- Improves I/O performance and reduces channel latency
 - Ideal for DASD and tape I/O-intensive workloads
- Improves RAS via DPU clusters
- Reduces power consumption
- If you carry forward older FICON adapters to z17, they will **not** use or get the benefits of the DPU:
 - FICON Express32S adapters
 - FICON Express16SA adapters



Data Processing Unit (DPU) – I/O Accelerator

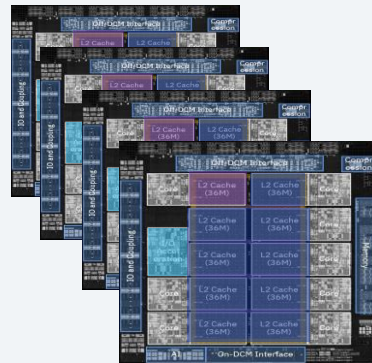
- Up to and including z16, TCP/IP connectivity was via **OSA-Express adapters**
 - Uses the efficient **Queued Direct I/O (QDIO)** protocol
 - Requires IBM specialized logic within the adapter card to enable features such as virtualization and LPAR to LPAR packet routing
- New on z17 are **Network Express adapters**. **What protocol do they use?**
 - **QDIO 2.0**
 - **QDIO+**
 - **QDIO Express**
 - **Enhanced QDIO**
 - **Next Gen QDIO**
 - **Rapido QDIO**
 - **Really Quick QDIO**
 - **Superior Networking Outstanding Bandwidth (SNOB)**

Data Processing Unit (DPU) – I/O Accelerator

- Up to and including z16, TCP/IP connectivity was via **OSA-Express adapters**
 - Uses the efficient **Queued Direct I/O (QDIO)** protocol
 - Requires IBM specialized logic within the adapter card to enable features such as virtualization and LPAR to LPAR packet routing
- New on z17 are **Network Express adapters**
 - Uses the even more efficient **Enhanced QDIO (EQDIO)** protocol
 - Moves IBM specialized logic from the network adapter card to the DPU
 - Similar DPU benefits for networking as discussed for DASD and tape (FICON)
 - EQDIO is architected to handle higher bandwidth networks (25 GbE and beyond)
 - z/TPF APAR PJ46989 (2Q 2025) is required to use Network Express adapters
 - Much more detailed presentation on Network Express and z17 networking tomorrow at the *Communications Subcommittee*
- Some OSA-Express adapters can be carried forward to z17 but will **not** use or get the benefits of the DPU (or the benefits of EQDIO)

Next Generation On-Chip AI Accelerator (AIU)

- IBM z16 introduced the industry first on-chip AI accelerator (AIU)
 - Real-time inferencing at scale using popular AI frameworks running on Linux on Z LPAR next to your z/TPF LPAR
- IBM z17 has the improved next generation AIU
 - 24 trillion of operations per second (TOPS) – 4x improvement
 - New data types now supported, including INT8
 - More efficient processing for existing models
 - Enables new models to also be used
- Large language model (LLM) enhancements to allow broader range of AI models for a comprehensive analysis of both structured and unstructured (textual) data



Next Generation On-Chip AI Accelerator (AIU)

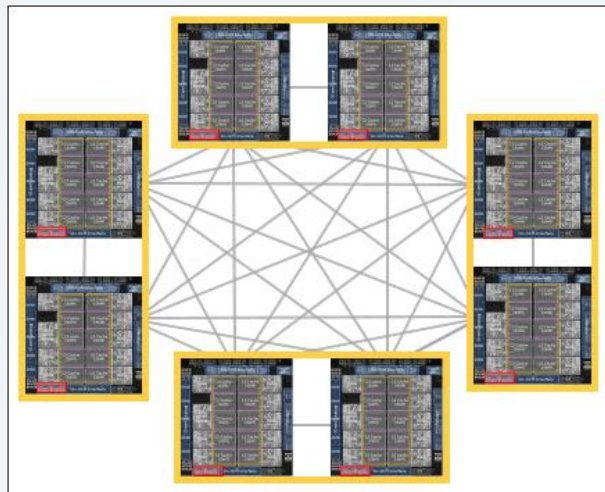
z17 In-Drawer Intelligent AIU Routing

IBM z16:

- A core (CPU) could only use the AIU on its own processor chip

IBM z17:

- A core (CPU) can use **all** AIUs within that drawer
- **8x** more compute capacity available to an AI workload



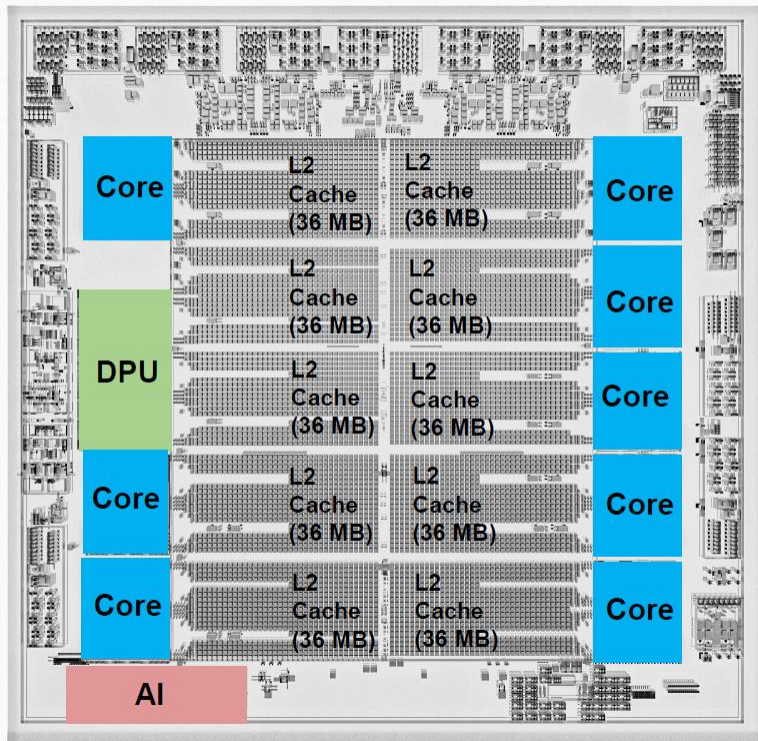
New Spyre AI Accelerator for Gen AI*

- **Purpose-built:** to scale foundation model inferencing (not training) for a curated set of AI models tailored to the use cases that are most relevant to enterprise clients.
 - AI Assistants such as the IBM watsonx Assistants
 - Document processing
 - Classification
- **Scalable AI:** can cluster together multiple Spyre adapters together to handle more complex Gen AI models
- **Sustainable AI:** only **75W** of power per adapter
- **Secure AI:** using a Confidential Computing environment
- **Multi-Model AI:** use the AIU to produce high confidence results for 80-90% of transactions, then send the lower confidence outcomes to Spyre for deeper analysis



* Target availability 4Q 2025

View of IBM z17 Telum II Processor Chip



Each Telum II chip includes:

- Up to 8 CPU cores
- 1 Data Processing Unit (DPU) I/O accelerator
- 10 L2 memory caches
- 1 next gen AI accelerator (AIU)

There are 8 chips in each drawer

z17 Improvements – Size and Speed Matter

	z16	z17
CPU clock speed	5.2 GhZ	5.5 GhZ (+5.8%)
L2 memory cache size	32 MB	36 MB (+12.5%)
L3 memory cache size	8 * 32MB = 256 MB	10 * 36MB = 360 MB (+40%)
L4 memory cache size	8 * 8 * 32MB = 2 GB	8 * 10 * 36MB = 2.88 GB (+40%)
Memory size	4 * 10TB = 40 TB	4 * 16TB = 64 TB (+60%)

- Redesigned translation lookaside buffer level 2 (TLB2) for faster address translation
- Branch prediction improvements
- Out-of-order execution pipeline enhancements

IBM z17 Summary – Part 1 of 2

- **Bigger, Better, Faster Processing with More Capacity for Traditional Core Workloads**
 - CPU improvements (speed, pipeline, prediction accuracy)
 - Larger memory caches and memory coherency manager improvements
 - I/O improvements, including the new on-chip Data Processing Unit (DPU)
 - Network performance improvements with new Network Express adapters
- **More Efficient and Environmentally Friendly Processing**
 - **Reduced physical footprint** with redesigned chip, drawer, and I/O infrastructure
 - **Reduced power consumption** with redesigned drawers, I/O infrastructure, Voltage Control Loop (VCL), and of course the new z17 AI accelerators
 - **Reduced carbon footprint** with redesigned carry forward I/O expansion frames, reduced packaging, and coolant change to propylene glycol
 - Easier install and removal with coolant change to propylene glycol

IBM z17 Summary – Part 2 of 2

- **Full Stack Enterprise AI Solutions for Workload Lifecycle Use Cases**
 - Next generation on-chip AI acceleration (AIU) for real-time inferencing at scale
 - New and powerful Spyre AI accelerators for Gen AI and more complex models
 - **Development:** Run watsonx Assistants to improve developer productivity on z17
 - **Production:** Better decision making done “in transaction” for 100% of your workload
 - **Manage:** Extend RTMC with AI and ML for sophisticated real-time operational analytics
 - **Business Optimization:** Extend RTMC with AI and ML for sophisticated real-time business analytics
- **z/TPF APARs for z17**
 - APAR PJ46989 is required to use the new Network Express adapters
 - APAR PJ48194 is required to use the new FICON-Express32-4P adapters

More z/TPF Education and Faster Skills Building

Link to Education Materials on z/TPF Landing Page (<https://www.ibm.com/docs/en/ztpf/latest>)

The screenshot shows the IBM z/TPF Documentation landing page. The browser address bar displays `ibm.com/docs/en/ztpf/2025`. The page header includes the IBM logo, the word "Documentation", and a search bar. The left sidebar contains a "z/Transaction Processing Facility..." section with a "Change version" dropdown set to "2025", a "Show full table of contents" checkbox, and a "Filter on titles" input field. Below this is a "Welcome" section with links to "z/TPF 2025", "z/TPFDF 2025", "TPF Operations Server", and "TPF Toolkit". At the bottom of the sidebar is a link for "Announcements & sales manuals". The main content area features a "Resources" section with a grid of six resource cards. The first card is "IBM TPF Support Portal", the second is "TPF Family of Products Downloads", and the third is "IBM Fix Central". The fourth card, "z/TPF Education and Other Resources", is highlighted with a blue border and a green arrow pointing to it. The fifth card is "z/TPF Community and Blog" and the sixth is "TPF Users Group". Each card includes a brief description and an external link icon.

Resources		
IBM TPF Support Portal Get assistance for the z/TPF product family. Link	TPF Family of Products Downloads Access product software, maintenance, tools, and drivers for the z/TPF product family. Link	IBM Fix Central Provides fixes and updates for TPF Toolkit. Link
z/TPF Education and Other Resources Browse educational content and other resources for the z/TPF product family. Link	z/TPF Community and Blog A web page dedicated to sharing news, knowledge, and insights about the z/TPF product family. Link	TPF Users Group A non-profit organization dedicated to the advancement of the z/TPF operating system. Link

Free Online z/TPF Education

IBM Training

Online, self-paced courses

z/TPF Programming Models

Level: Beginner

Audience: System programmers, application programmers

Learn basic concepts, process models, program environments, and supported programming languages on the z/TPF system.

z/TPF Task Management

Level: Intermediate

Audience: System programmers, application programmers

Learn how transactional work flows through the z/TPF system, including giving up and regaining control under various conditions.

z/TPF Task Management for Utilities

Level: Intermediate

Audience: System programmers, application programmers

Learn what is considered utility or batch work, define what APIs and settings are used to manage utility or batch work, and understand how utility and batch work flows through the system.

← New for 2025

Education on z/TPF TCP/IP Support currently being created

IBM Redbooks



IBM MediaCenter

⇒ Z/TPF INTERNSHIP

z/TPF Detailed Summary
z/TPF Middleware - DFDL
—
Bradd Kadlecik
z/TPF Development

01:04:31

z/TPF Internship:
Introduction to DFDL

z/TPF Detailed Summary
z/TPF Automated Test Framework
—
Jennifer Chariot
z/TPF Development

38:45

z/TPF Internship: z/TPF
Automated Test Framework

z/TPF High Speed
Connector and Enhanced
HTTP Client
—
Jamie Farmer
z/TPF Development

42:02

z/TPF Internship: High
Speed Connector

z/TPF Detailed Summary
z/TPF Middleware – IBM MQ
—
John Muller
z/TPF Development

43:09

z/TPF Internship: MQ for
z/TPF

⇒ TPF TOOLKIT

Customizing Perspectives
TPF Toolkit 4.6

06:57

TPF Toolkit: Customizing
Perspectives

Building TPF Projects
TPF Toolkit 4.6

08:59

TPF Toolkit: Building TPF
Projects

Creating TPF projects
TPF Toolkit 4.6

07:41

TPF Toolkit: Creating TPF
Projects

Resource Markers
Visualizing errors and warnings in your code


16:56

TPF Toolkit: Resource
Markers

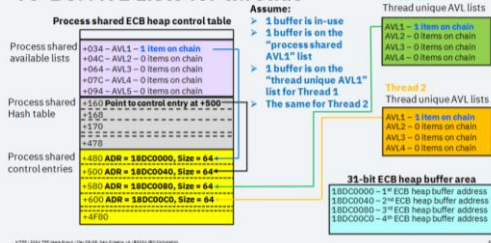
[Load More](#)

IBM Presentations from TPF Users Group Conferences


64-Bit IBM MQ Memory vs 31-Bit IBM MQ Memory (SWBs)



To-Be: AVL Lists for threads




Multiple Unordered Dispatch Queues




Over **200** in just the past decade alone

Also Leverage AIU for Post-Transaction Processing and AIOPs




Technical Details

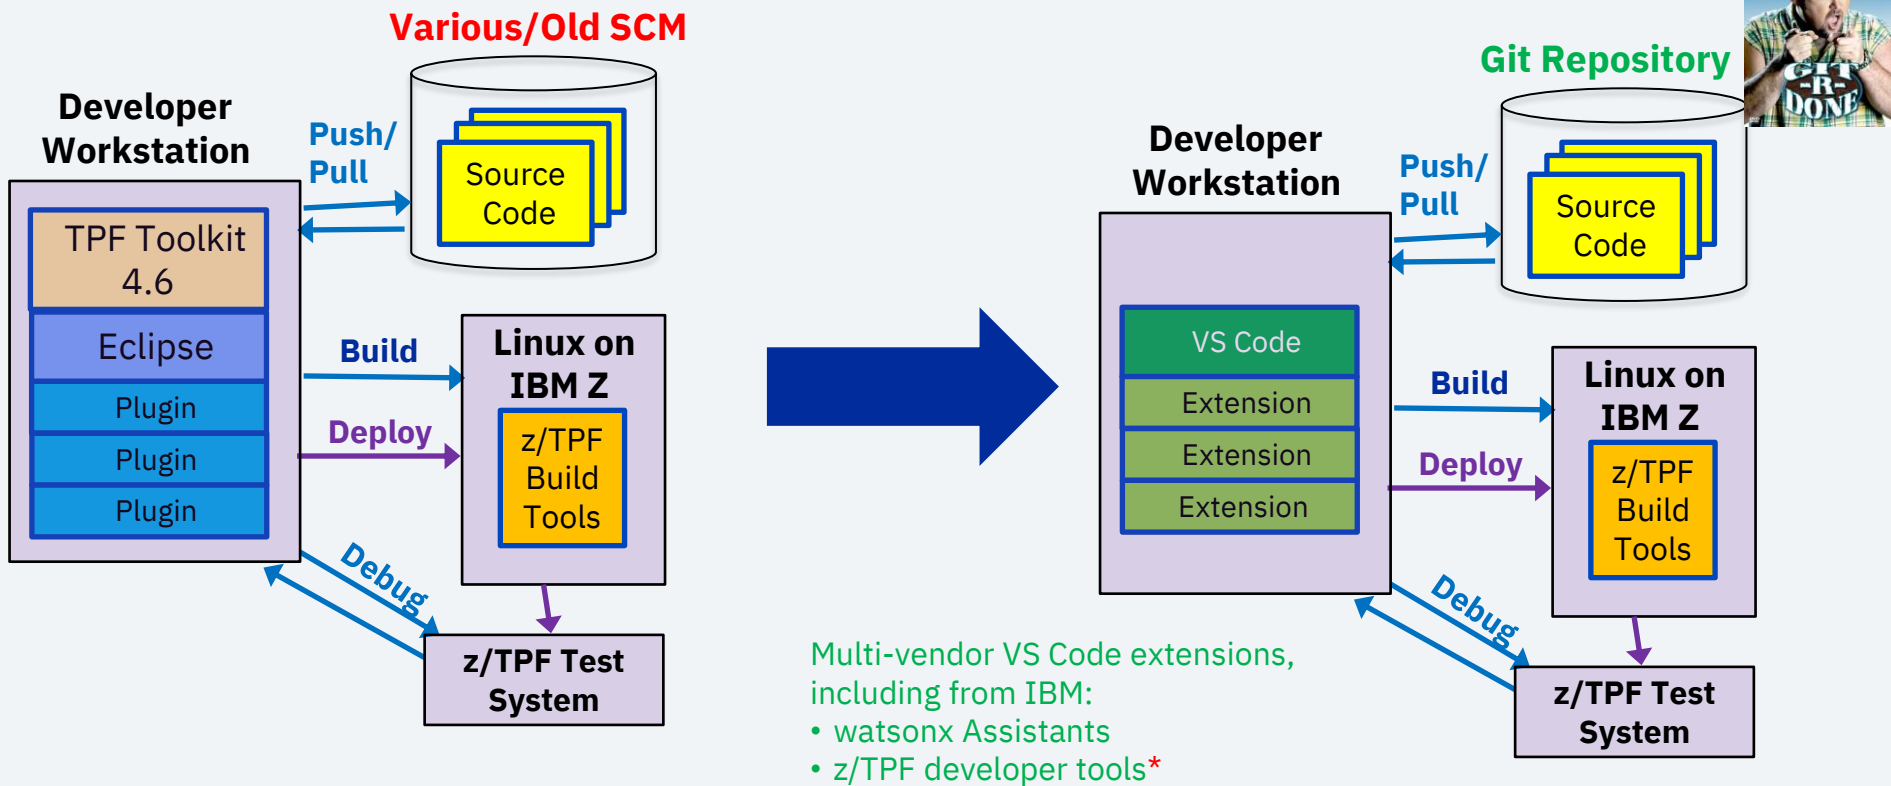


APAR PJ46608 (Jan 2022) – Correlation Analysis Dashboard

Simplified correlation dashboards – select metrics and application dimension.



z/TPF Development Environment Transition Leveraging AI Assistants



Utilize All These Resources to Accelerate Your [z/TPF University](#) Experience



Thank you

© Copyright IBM Corporation 2025. All rights reserved. The information contained in these materials is provided for informational purposes only, and is provided AS IS without warranty of any kind, express or implied. Any statement of direction represents IBM's current intent, is subject to change or withdrawal, and represent only goals and objectives. IBM, the IBM logo, and ibm.com are trademarks of IBM 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 at [Copyright and trademark information](#).

