IBM z16 and Hybrid Cloud with z/TPF

2023 TPF Users Group Conference
April 24-26, Dallas, TX

Mark Gambino, IBM z/TPF Chief Architect
IBM Distinguished Engineer

Agenda

- z16 benefits for z/TPF clients
- Some z16 data for all you nerds out there
- Halftime Show
- Why z16 is built for Hybrid Cloud
- When, why, how Hybrid Cloud

IBM z16 is built to build

We built a powerful and secure platform for business.
Let’s build the future of yours.

Predict and Automate for Increased Decision Velocity

Secure with a Cyber Resilient System

Build a Sustainable Infrastructure

Modernize with Hybrid Cloud

Where Are You in Your AI Journey

- **43%** accelerated their rollout of AI as a result of the COVID-19 pandemic
- **Data Security** is the biggest area of AI investment in the next 12-24 months
- **49%** say getting insights where and when they are needed is a big challenge
- Over **90%** of companies using AI say their ability to explain how AI arrived at a decision is critical

- Top reasons why companies are investing in automation tooling:
  1. Driving greater efficiencies
  2. Cost savings
  3. Faster problem resolution
  4. Giving valuable time back to employees to work on other tasks

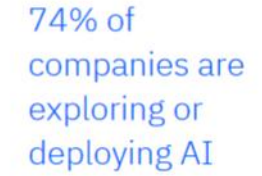

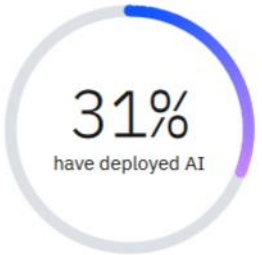

IBM z16 Integrated Accelerator for AI (AIU)

Industry first integrated on-chip AI accelerator designed for high-speed, latency-optimized inferencing

Accelerated AI at scale

Up to **300 billion** deep learning inference requests per day with 1ms response time

Speed to scale with transaction volume

Up to **19x higher throughput** and **20x lower response time** **co-locating** applications and inferencing

Real-time insights when you need them

Infuse AI into **every transaction** while still meeting the most stringent SLAs

Improve business outcomes

Scoring 100% of transactions can **reduce fraud loss by over $100M** per year for a large bank

Seamlessly Leverage AI Accelerator on IBM z16

- Bring machine learning & deep learning models to IBM z16 with ONNX/DLC
- Exploit IBM Integrated Accelerator for AI (AIU) for best inference performance
- Repeatable practice for different vendors to leverage IBM z16 and AIU

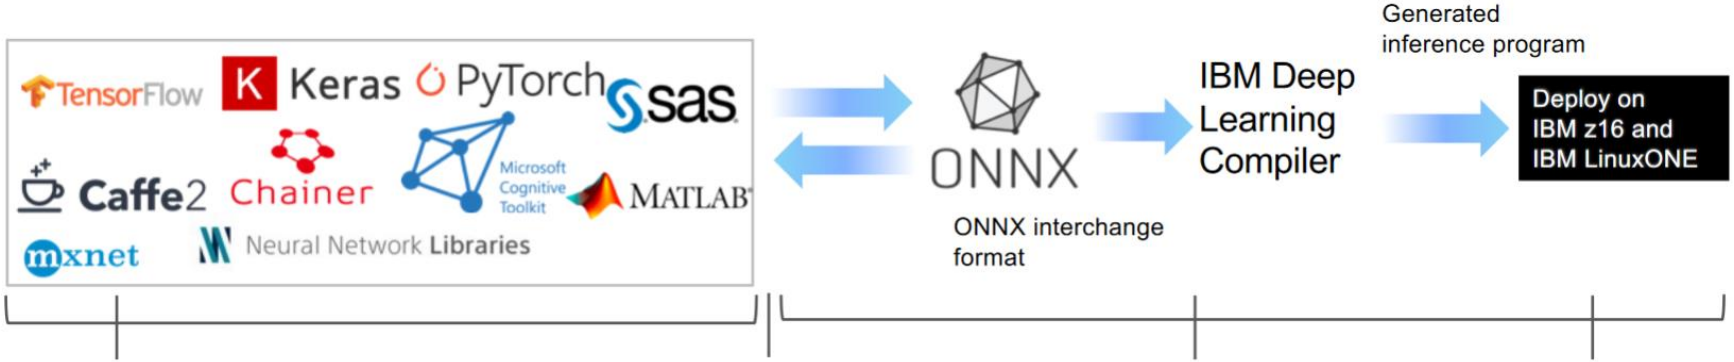

Build and train model in any popular framework on any platform of your choice

Use ONNX, an opensource tool for framework interoperability. Models are converted to the ONNX interchange format

The IBM DLC (Deep Learning Compiler), optimized for performance and new libraries, generates a program from the model for execution on Linux® on zSystems

Deploy on IBM z16 and IBM LinuxONE to infuse model into workload application leveraging AIU

Before: Real-Time Fraud Detection

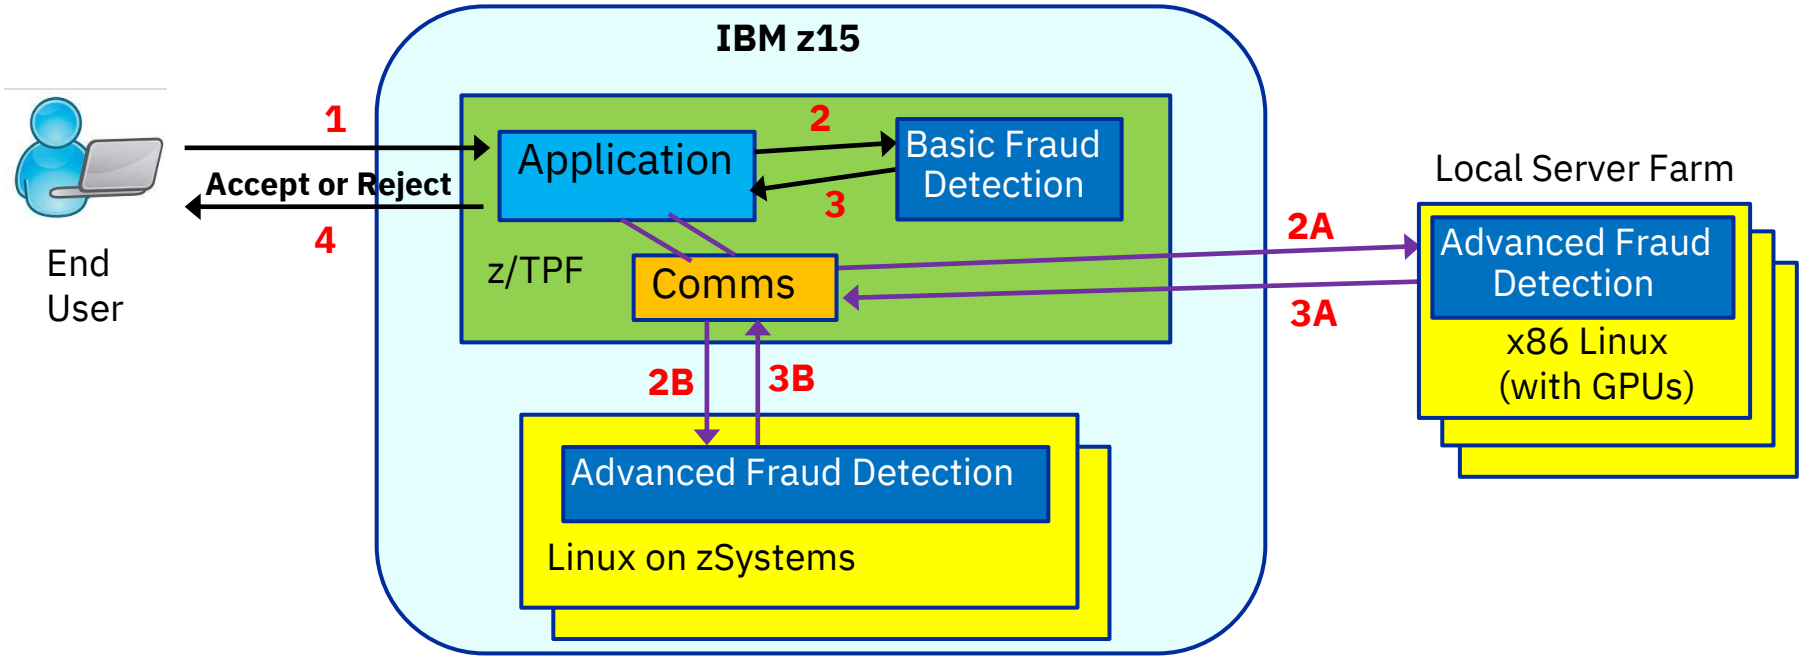

Basic checks done on z/TPF. More advanced Fraud Detection code was created and deployed on Linux servers, but only a **subset** of transactions used that because of cost and response time SLAs.

Now: Real-Time Fraud Detection

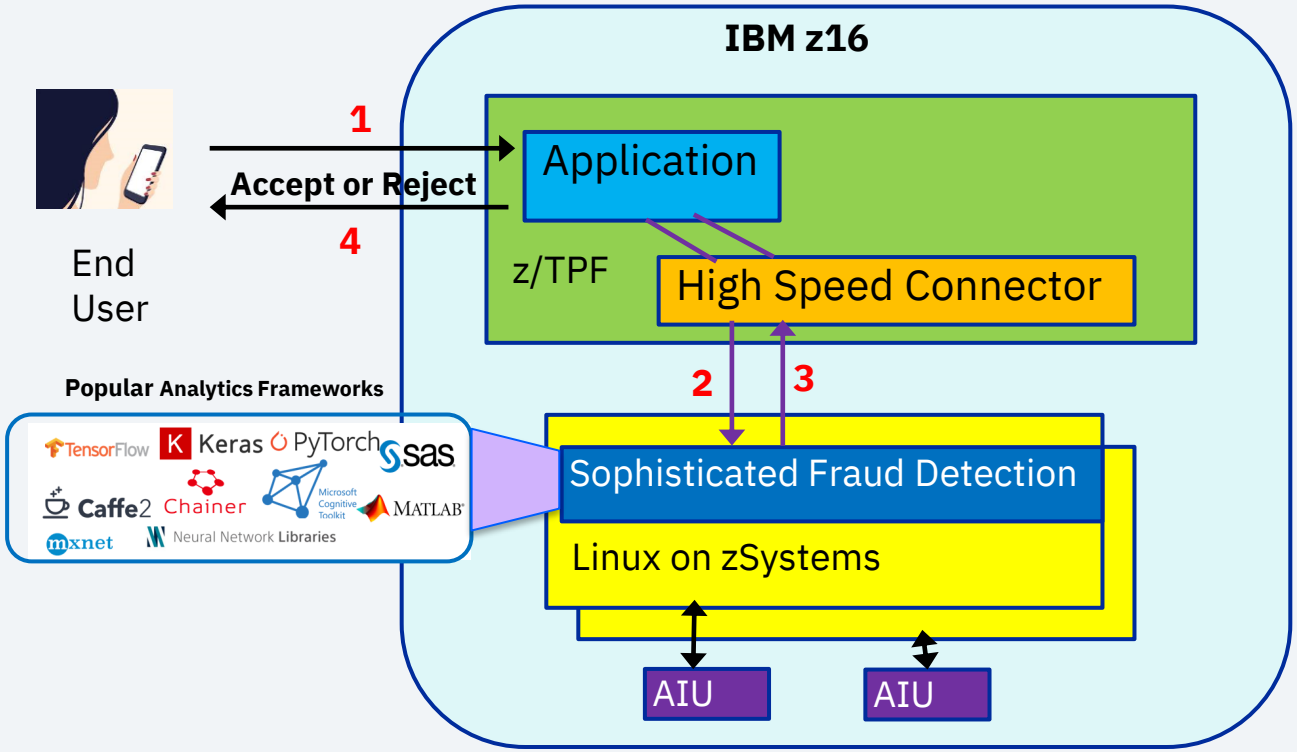

Do advanced Fraud Detection on **every** transaction leveraging popular AI frameworks and on-chip HW AI acceleration (AIU) to scale and meet your response time SLAs.

Also Leverage AIU for Post-Transaction Processing and AIOps

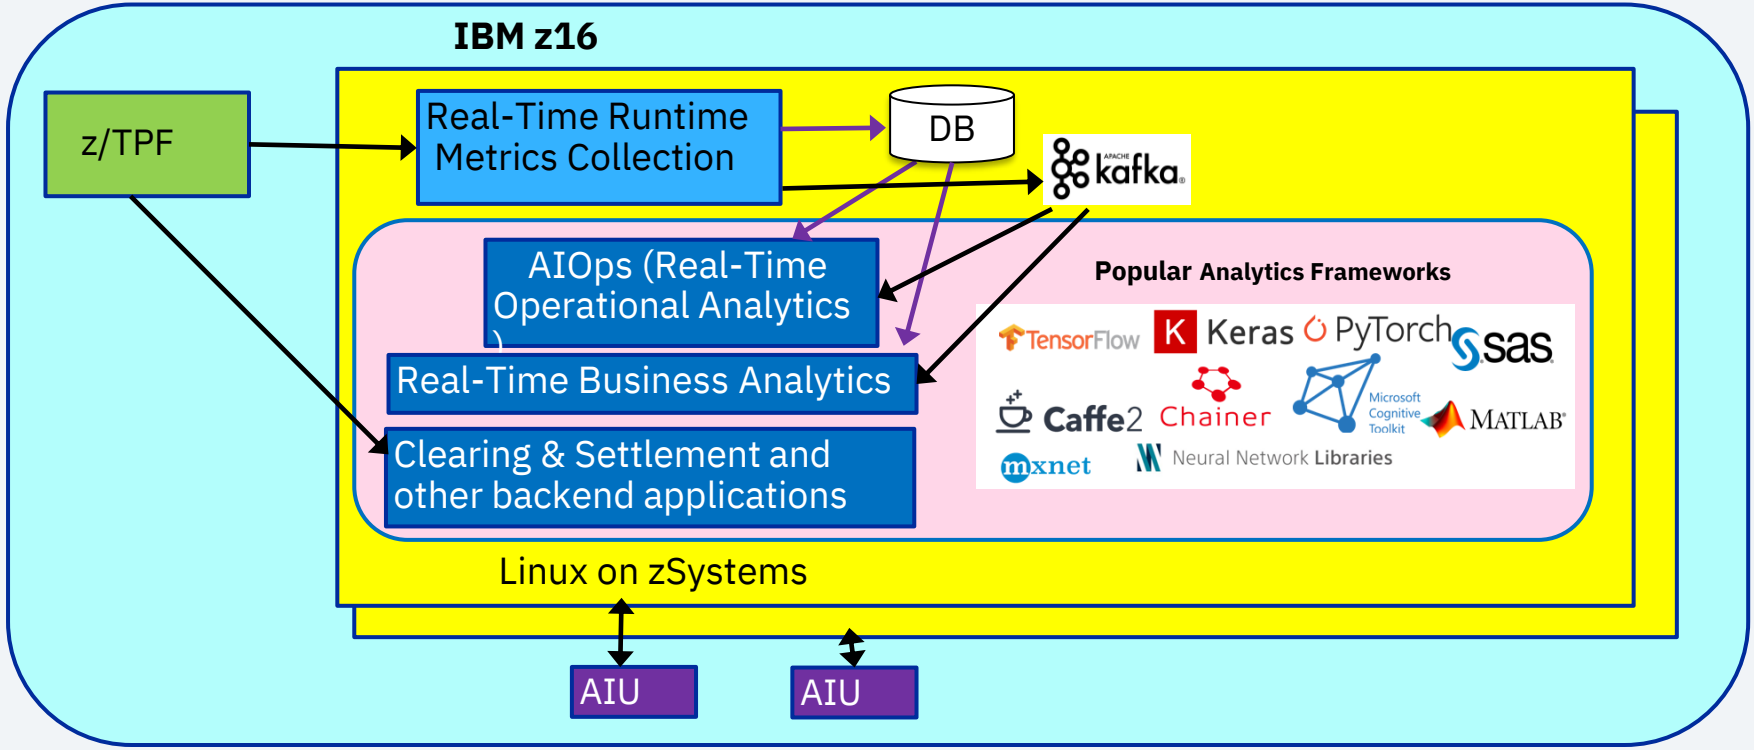

z16 Security Highlights

- **Industry-first cyber resilient system**
  - Multiple layers of firmware secured with dual-signature scheme including a quantum safe algorithm
- Secure memory (transparent memory encryption)
  - 8-channel RAIM (50% less overhead*)
- HW accelerated **quantum safe symmetric key cryptography** (AES-256)
  - Over **700K operations per second per I-stream** on z/TPF with 4K data size (22% increase*)
  - 28% higher throughput for 1000-byte messages over TLS on z/TPF*
- New Crypto-Express8 adapter for HW accelerated RSA operations
  - Requires z/TPF APAR PJ46507 (April 2022)

* compared to z15

Business Continuity is a Key Aspect of Cyber Resiliency

Proactive Outage Avoidance

With extreme weather events becoming more and more frequent, a proactive approach to delivering continuous service is needed.

You need to be able to migrate your critical workloads to an alternate site **before** your business gets impacted.

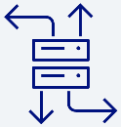

Disaster Recovery and DR Testing

In the event of an unplanned outage, including cyber attacks, the ability to **rapidly restore** operations and service is paramount.

The ability to test that **production workloads can be shifted** and run at full capacity is critical for ensuring continuous availability during unplanned outages is key

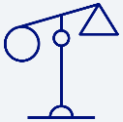

Business Continuity Compliance

Regulation around business continuity and disaster recovery policies are increasing and becoming more stringent.

These regulations mandate that businesses be able to switch over full production loads to a secondary data center and **operate there for extended periods of time.**

Site Facility Maintenance

Site facility and building maintenance is an ongoing activity for businesses. Upgrading for health, environmental, and safety purposes or other improvements **sometimes requires closures.**

The ability to continue to provide 24x7 service to your customers is more important than ever.

IBM z16 Flexible Capacity for Cyber Resiliency

Designed to help organizations proactively reduce the impact of downtime by dynamically shifting their critical workloads to an alternate site for business continuity

Greater Flexibility

Dynamically shift production capacity between z16 systems at different sites.

Can be used for proactive outage avoidance, business continuity compliance, disaster recovery and DR test scenarios.

Be confident that production workloads can be seamlessly shifted to an alternate site and **still meet production SLAs.**

Complete Client Control

Remotely transfer capacity **– no on-site personnel required** after initial set up.

Flexibility over duration of capacity transfer, production **can remain at the alternate site for up to one year.**

Integrates with System Recovery Boost for **faster system and workload startup**

Simplified Compliance

Simplify business continuity **compliance and improve audit readiness** by using the **same procedures for both for DR testing and real unplanned disasters**.

Test recovery procedures for unplanned outages to provide near-continuous availability and disaster recovery.

An Outdoor Hot Tub in the Middle of Winter

Great for the legs after a full day of skiing

Watch that electric meter spin

Sustainability is a Rapidly Growing Area of Focus

- Over **2,500 climate laws and policies** exist worldwide
- 86% of companies have a **sustainability strategy**, but **only 35% have acted** on them
- 51% of CEOs name **sustainability as their greatest challenge** for their organization over the next 2–3 years
- If you don’t know what **ESG** is yet, you will

Which Architecture Would Mother Nature Choose

- OR -

z16 Server running z/TPF, consolidated Linux servers, or both

Thousands of x86 servers

Which Architecture Does Mother Nature Choose

z16 Server running z/TPF, consolidated Linux servers, or both

Reduction in:

- **Electricity consumption**
- **Floor space**
- **Software costs**
- **CO_2e footprint**

And as added bonuses:

- **Better performance**
- **More secure**
- **More reliable**
- **Easier to manage**

Frustrating 2-Word Phrases Learned in the Pandemic

Frustrating 2-Word Phrases Learned in the Pandemic

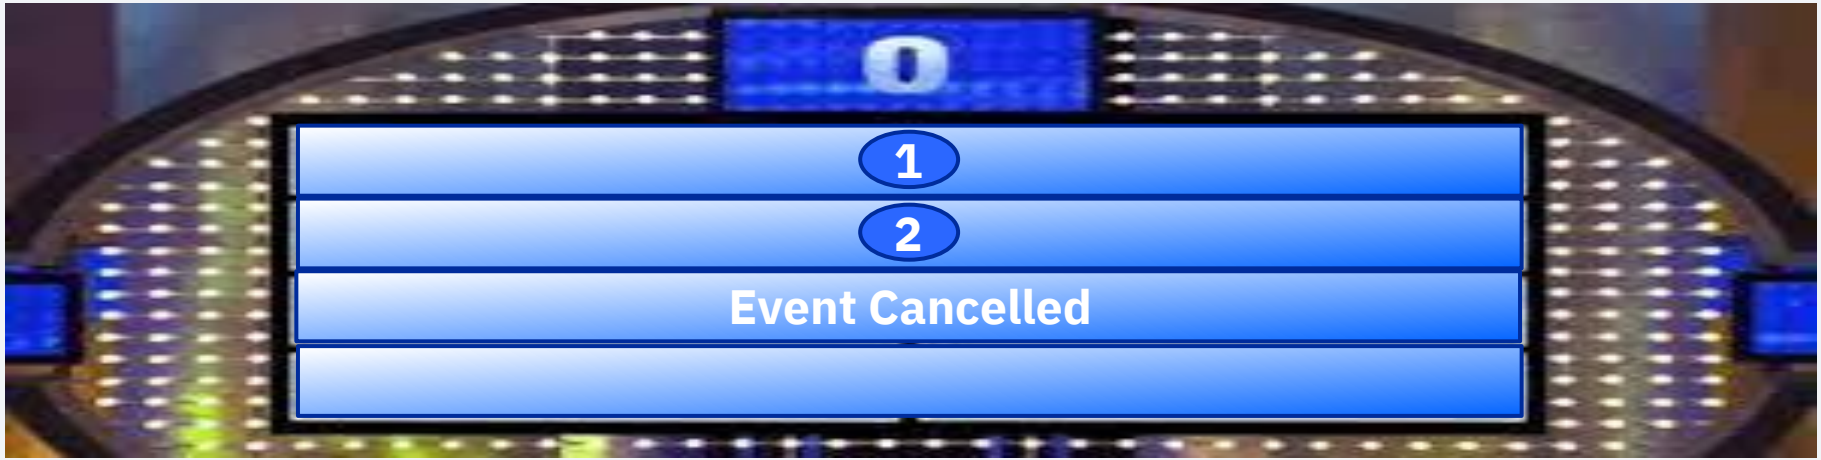

Frustrating 2-Word Phrases Learned in the Pandemic

Frustrating 2-Word Phrases Learned in the Pandemic

Supply chain **disruptions** were numerous, including impacting the ability for many companies to obtain more computer servers to handle increased transaction loads

Capacity on Demand - Most IBM z16 and LinuxONE servers have dark cores that can be enabled to **turn on extra CPU capacity that exists in the boxes you already have!**

z16 CPU and Memory Improvements (Compared to z15)

- 11% more CPU capacity per core*
- **25% more CPU capacity per drawer*** (Large single LPAR growth)
- 40% more CPU capacity in maximum configuration*
- Redesigned memory cache with flatter topology
  - **Faster data access (50% more memory cache capacity per core)**
  - Reduced average access latency
  - Semi-private L2 per core, virtual L3 per chip, virtual L4 per drawer
- Up to 32TB of memory per LPAR (100% more)

* Based on published processor ratings, your results may vary

IBM z16 Telum Processor

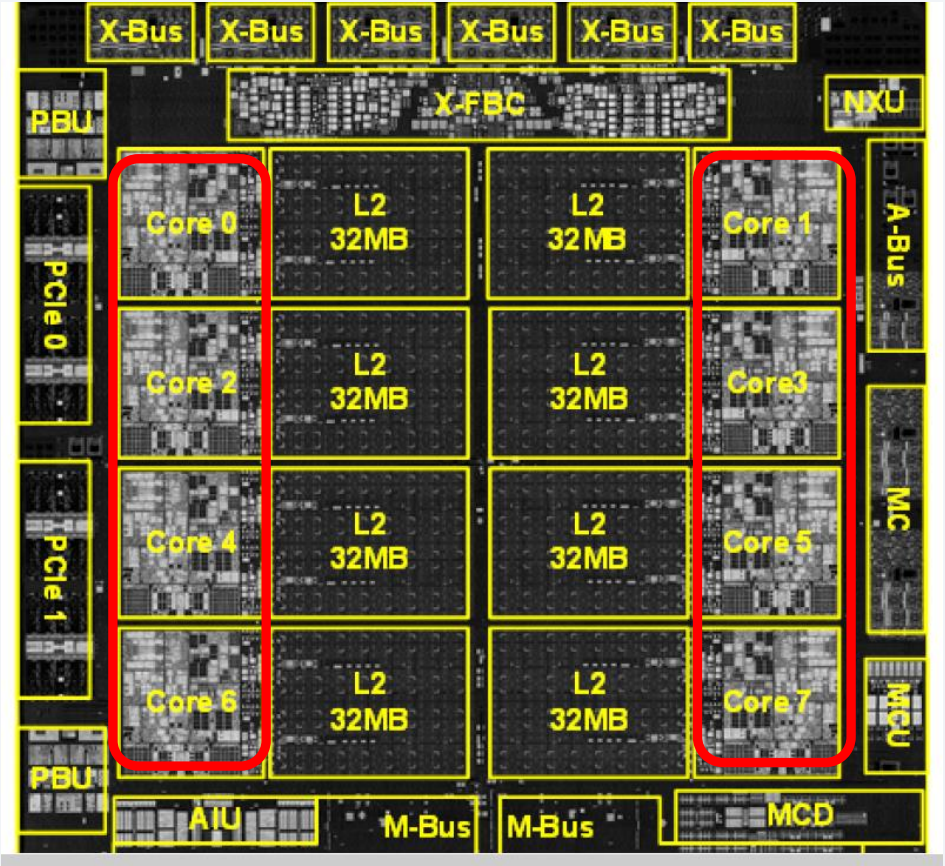

- **8 cores per chip (5.2 GHz)**

IBM z16 Telum Processor

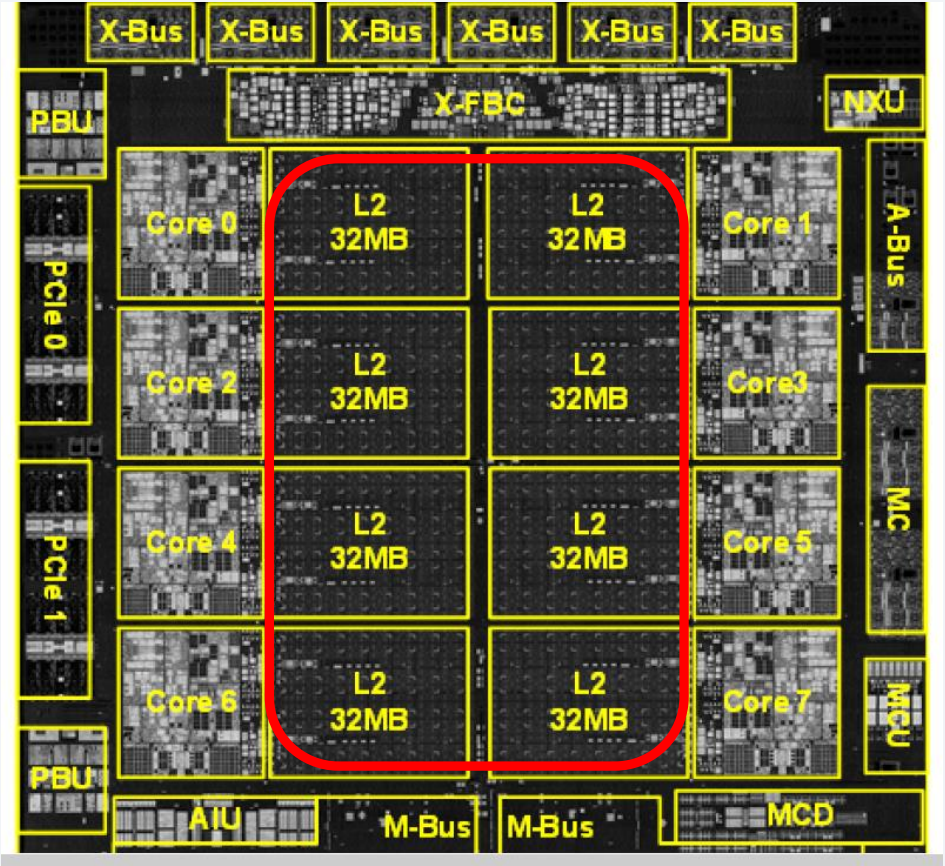

- 8 cores per chip (5.2 GHz)
- **256MB memory cache on chip**

IBM z16 Telum Processor

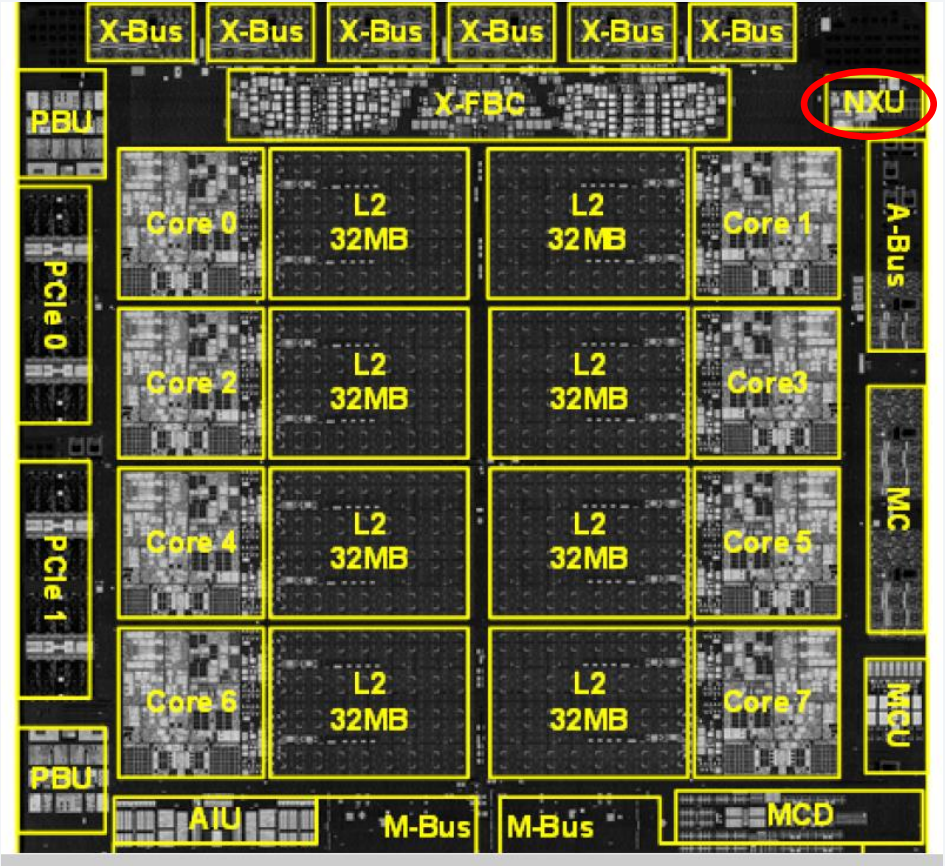

- 8 cores per chip (5.2 GHz)
- 256MB memory cache on chip
- **Compression HW accelerator on chip**
  - **IBM Z Integrated Accelerator for zEnterprise Data Compression (zEDC)**

IBM z16 Telum Processor

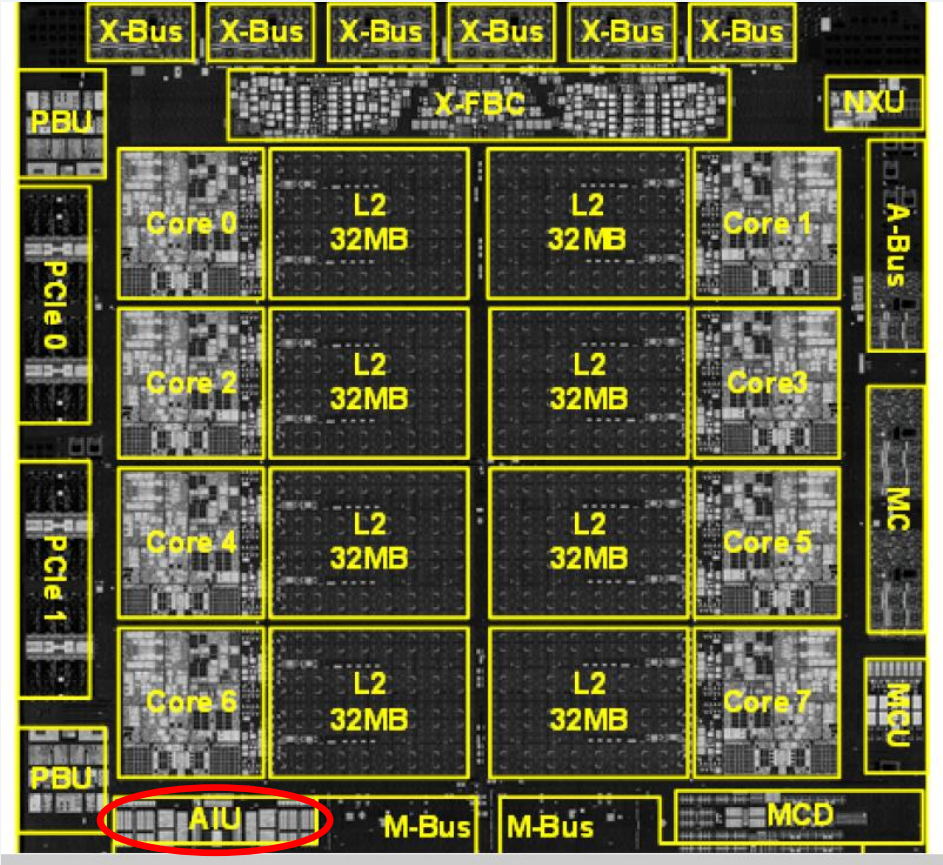

- 8 cores per chip (5.2 GHz)
- 256MB memory cache on chip
- Compression HW accelerator on chip
  - IBM Z Integrated Accelerator for zEnterprise Data Compression (zEDC)
- **AI HW accelerator on chip**
  - **IBM Z Integrated Accelerator for AI (AIU) new for z16**

IBM z16 Telum Processor

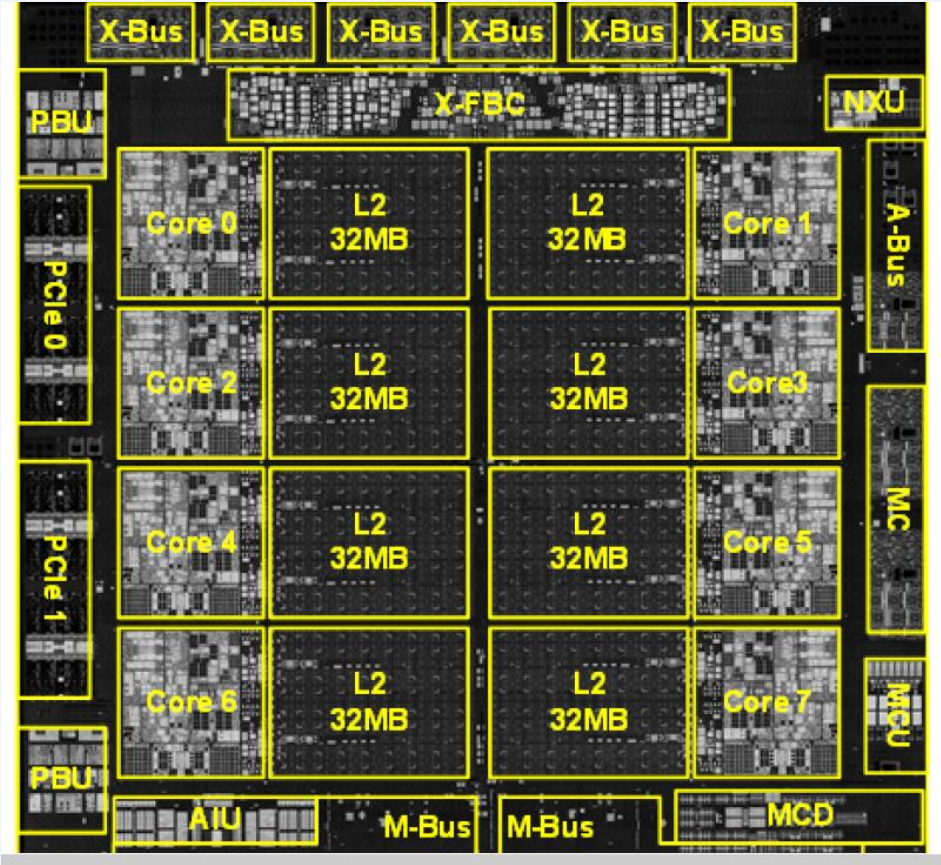

- 8 cores per chip (5.2 GHz)
- 256MB memory cache on chip
- Compression HW accelerator on chip
  - IBM Z Integrated Accelerator for zEnterprise Data Compression (zEDC)
- AI HW accelerator on chip
  - IBM Z Integrated Accelerator for AI (AIU) new for z16
- **22.5 billion transistors on each Dual Chip Module (DCM)**

IBM z16 Core

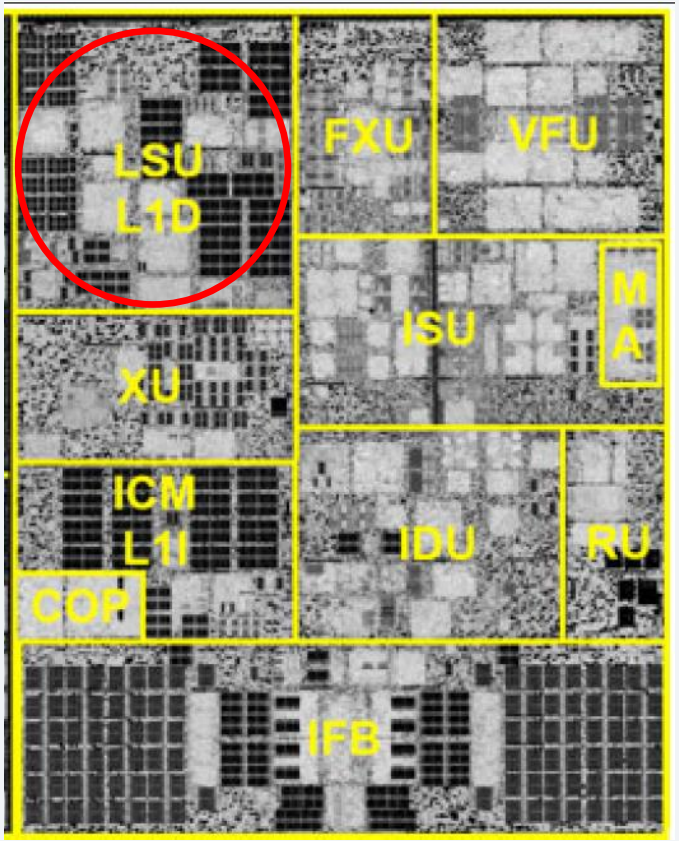

- **L1 data cache (128 KB)**

IBM z16 Core

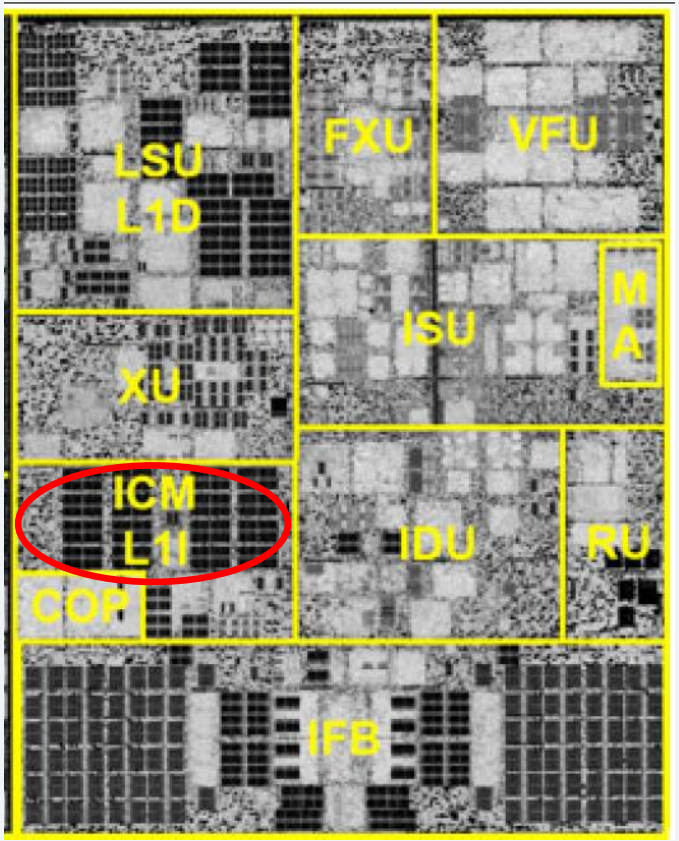

- L1 data cache (128 KB)
- **L1 instruction cache (128 KB)**

IBM z16 Core

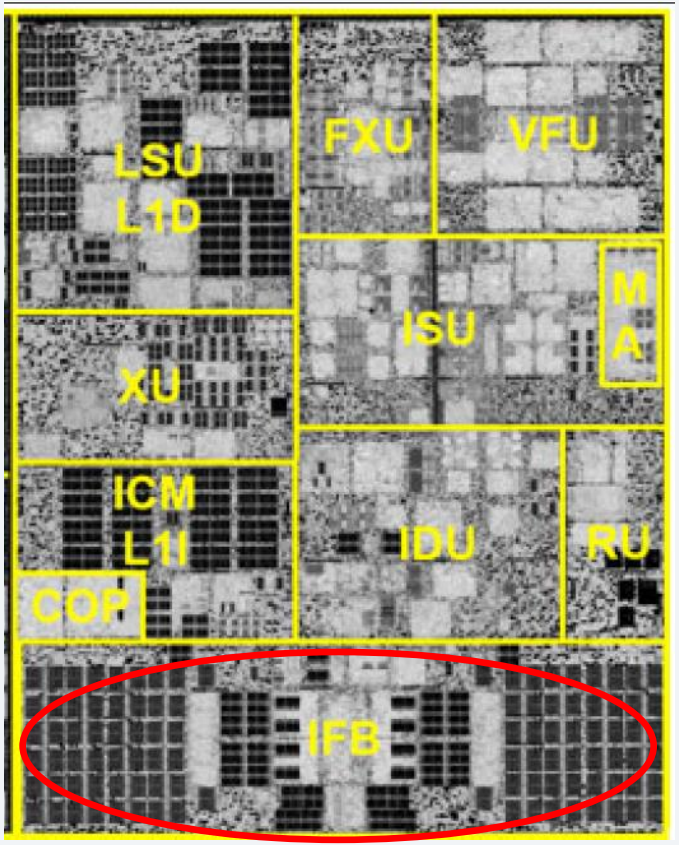

- L1 data cache (128 KB)
- L1 instruction cache (128 KB)
- **Instruction fetch and branch prediction**

IBM z16 Core

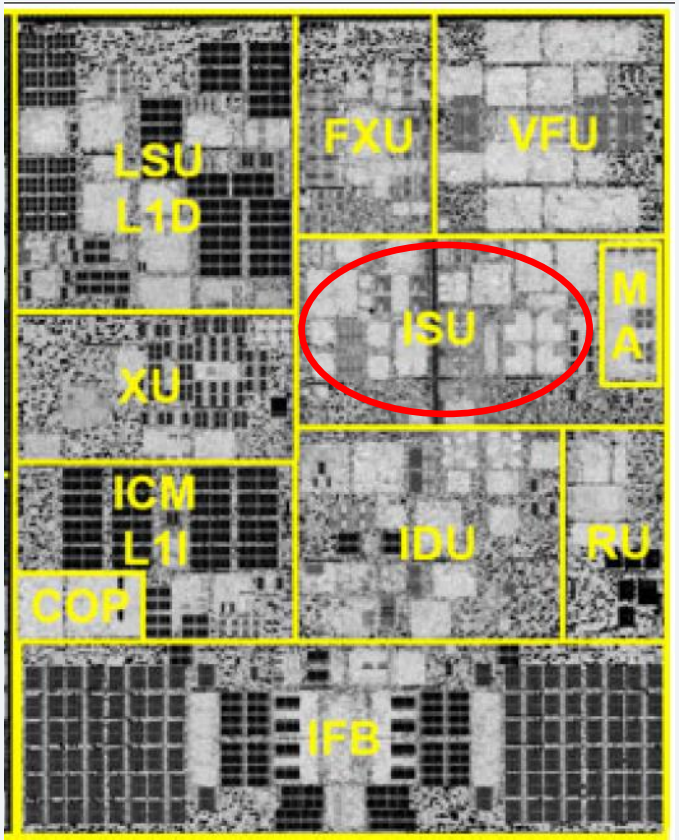

- L1 data cache (128 KB)
- L1 instruction cache (128 KB)
- Instruction fetch and branch prediction
- **Instruction sequence unit handles the out-of-order execution pipeline**

IBM z16 Core

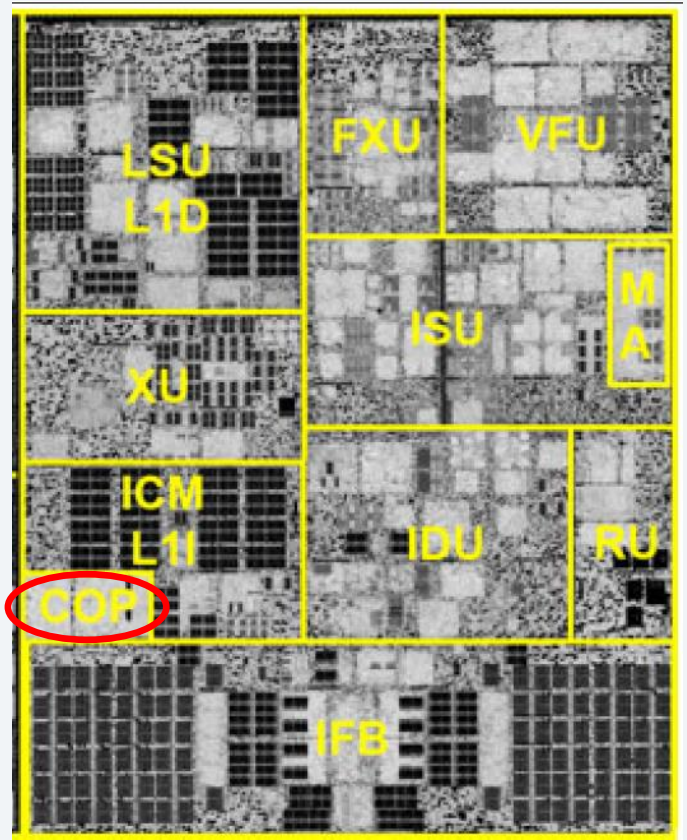

- L1 data cache (128 KB)
- L1 instruction cache (128 KB)
- Instruction fetch and branch prediction
- Instruction sequence unit handles the out-of-order execution pipeline
- **Crypto HW accelerator on chip**
  - **Central Processor Assist for Cryptographic Functions (CPACF)**

z16 Models and Upgrade Options

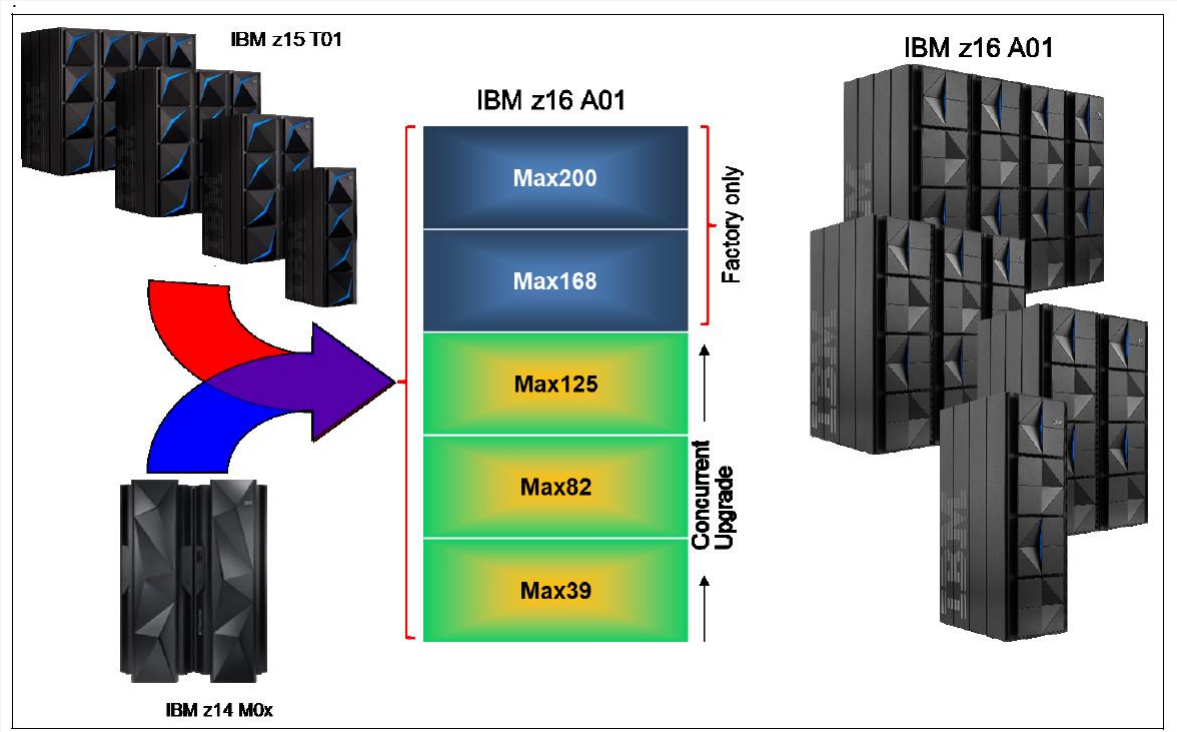

Model	Cores	Drawers	Memory
Max39	Up to 39	1	Up to 10TB
Max82	Up to 82	2	Up to 20TB
Max125	Up to 125	3	Up to 30TB
Max168	Up to 168	4	Up to 40TB
Max200	Up to 200	4	Up to 40TB

Halftime Show

Eventually Your Architecture Looked Like This

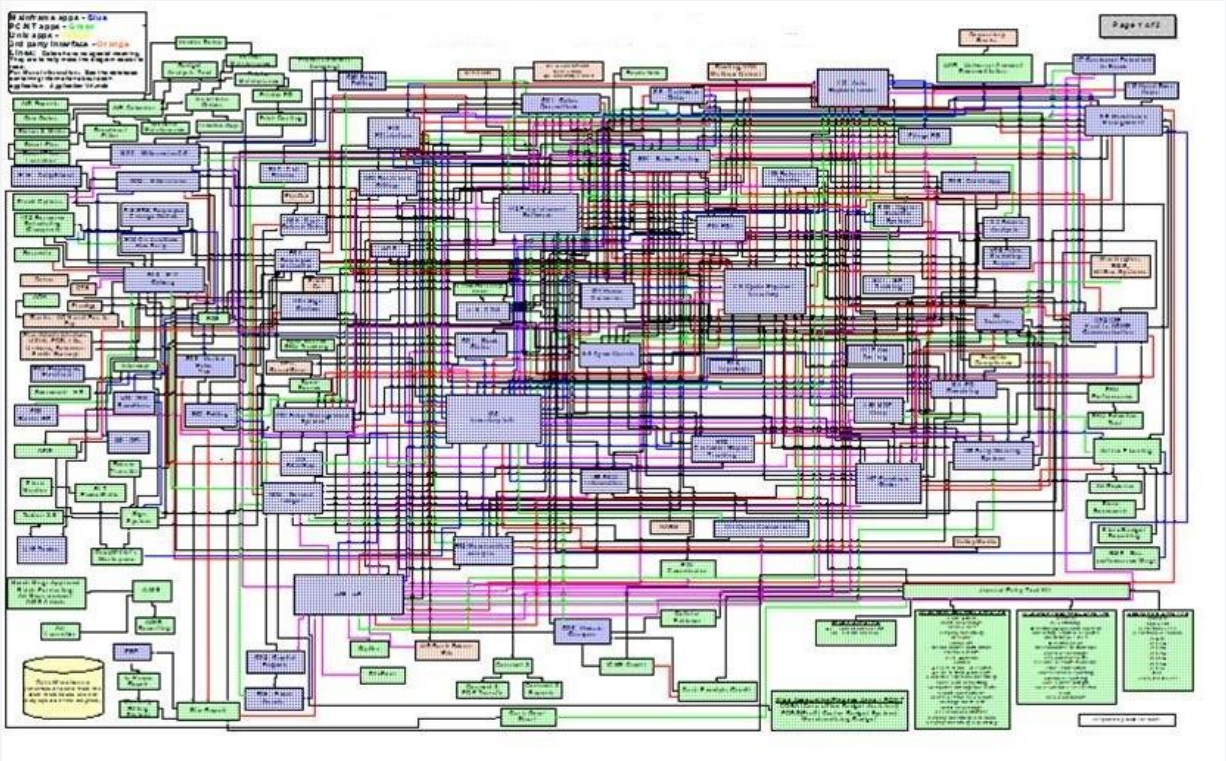

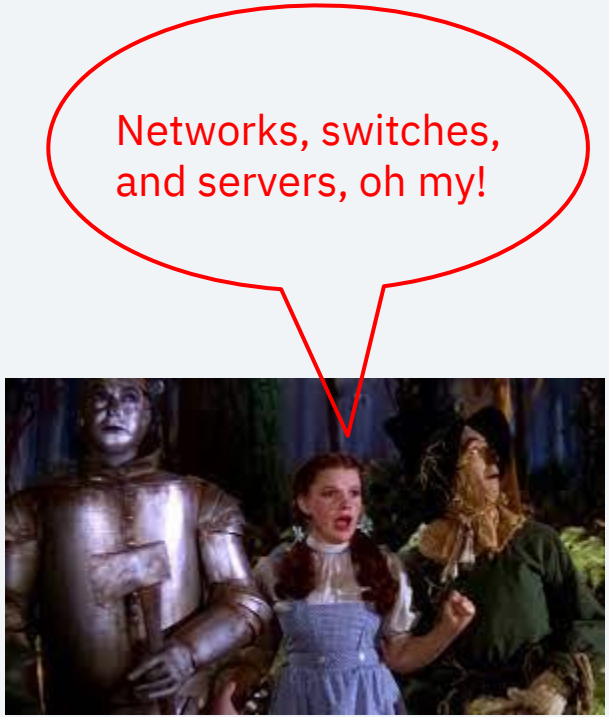

What Do The Following Statements Have in Common

- Texas is not that big
- Everyone attending the TPFUG is in bed by 8PM
- Every workload should run on IBM zSystems
- Every workload should run in the cloud

They are all wrong!

Deliver Digital Innovation without Sacrificing Performance or Resilience

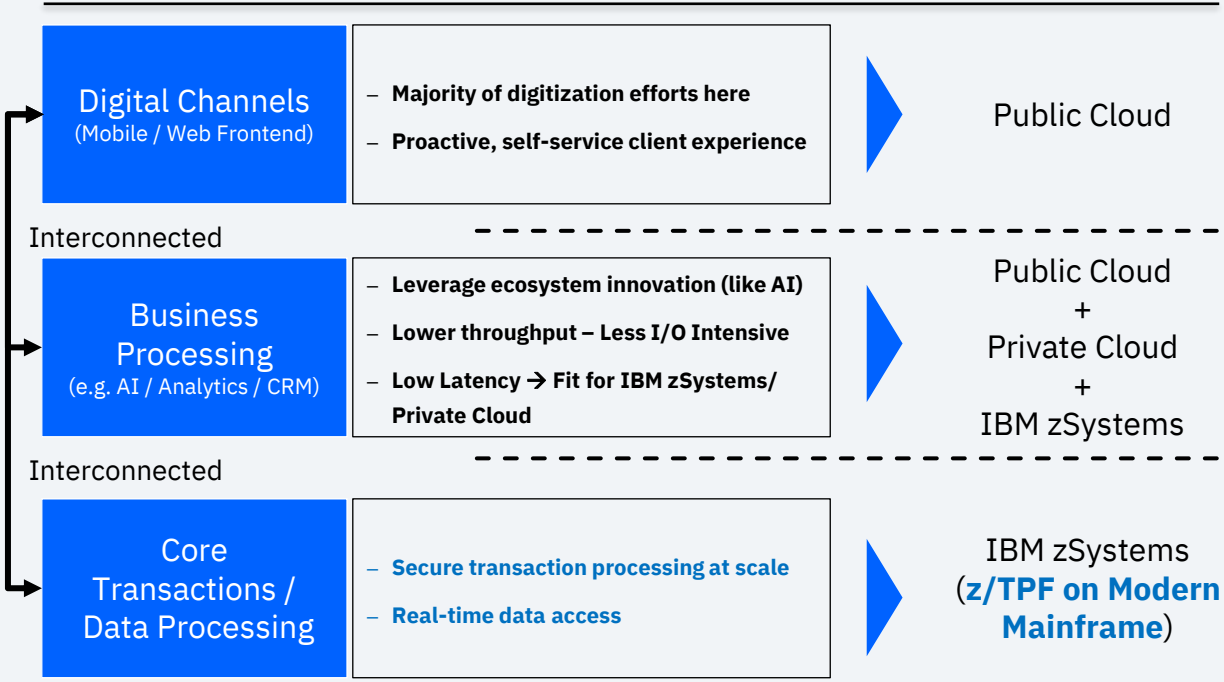

Fit-for-Purpose - Choosing the right platform based on workload requirements is critical first step to modernization

IBM zSystems delivers the scalability, and availability required for core systems … while cloud brings the innovation for digital channels and analytics. Private Cloud can be x86 or Linux on zSystems

Critical to ensure workloads across different platforms are seamlessly interconnected – at the app and data layer

Modernization = Hybrid Cloud = zSystems AND Cloud, not either or

President Biden Visits IBM Poughkeepsie (Oct 6, 2022)

"**IBM has more patents** than any other U.S. company — than any other"

"**It's here at this factory where America's future is literally being built** because of the groundbreaking CHIPS and Science Act "

"As we saw during the pandemic **when factories** that make these chips **shut down** around the world, **the global economy, literally, comes to a screeching halt.**"

If I Had a Nickel for Every Time I've Heard...

Claim 1: My Pilot Project to Move One Workload to Cloud was Successful; therefore, I Can Move Everything to Cloud

Digging into the details, that pilot workload…

- Was written decades ago when zSystems was the only option
- Throughput is measured in transactions per **minute**
- Database profile is predominantly **read-only**
- Is the only application that uses this database
- Is **not OLTP**

Reality - Core zSystems workloads…

- Have transaction rates **1000x** higher with demanding **response time SLAs**
- Are **write-intensive** updating multiple databases in an **always consistent** manner
- Have databases shared by multiple applications, all of which need to be **co-located** to meet SLAs

Claim 2: Moving Code Off of zSystems Will Reduce My Costs

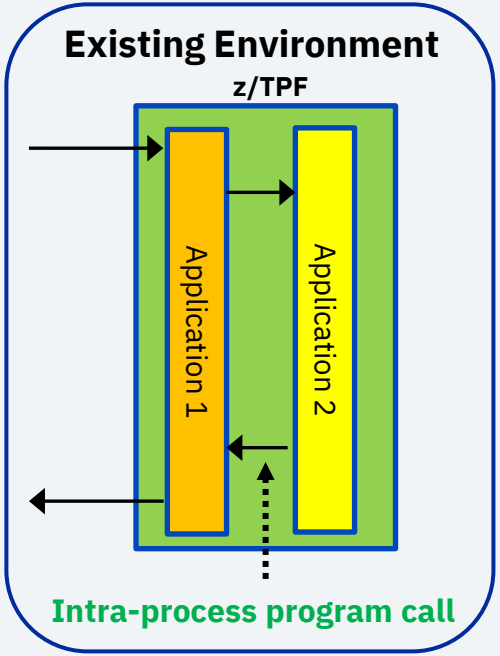

Execution Flow:

1. Application 1 makes a single call to Application 2 inside the same process (ECB)
2. Application 2 executes code and can access data created/retrieved by Application 1 (same ECB)
3. Application 2 returns to Application 1 via program return call
4. Any results and data created/retrieved by Application 2 is available to Application 1 (same ECB)

Simple Example, Single Call

Claim 2: Moving Code Off of zSystems Will Reduce My Costs

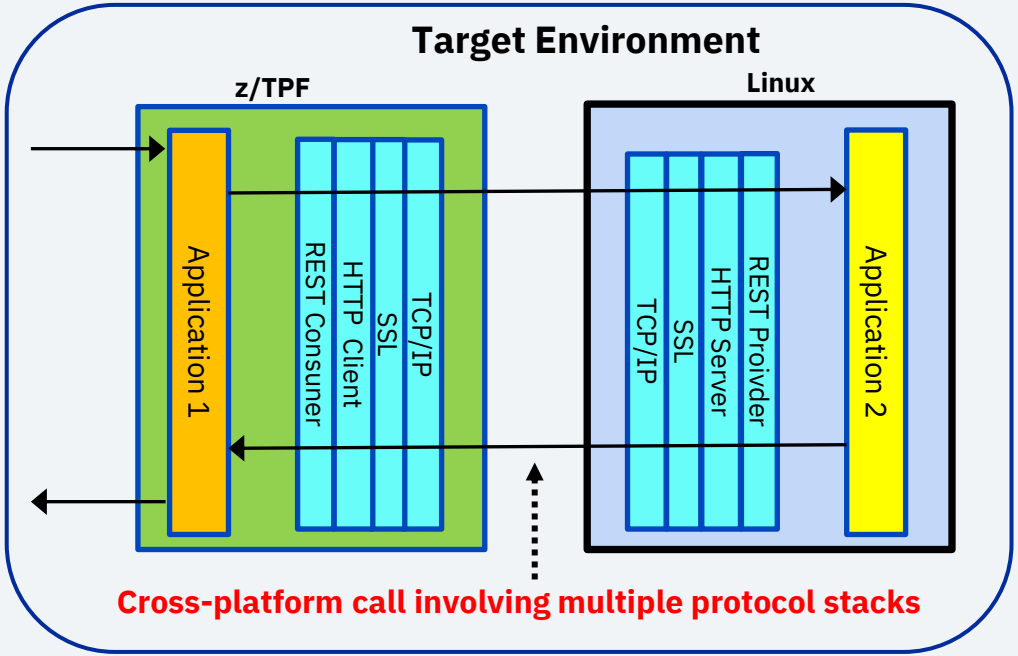

Execution Flow:

1. Application 1 makes a single call to Application 2 using REST (or other middleware) passing all the data needed by Application 2
2. Middleware, security, and network stacks on z/TPF and Linux transform and pass the information to Application 2. Load balancers, firewalls, and API Management nodes may also be involved.
3. Application 2 executes code and sends the response data via REST (or other middleware)
4. Middleware, security, and network stacks on Linux and z/TPF transform and pass the results to Application 1

Claim 2: Moving Code Off of zSystems Will Reduce My Costs

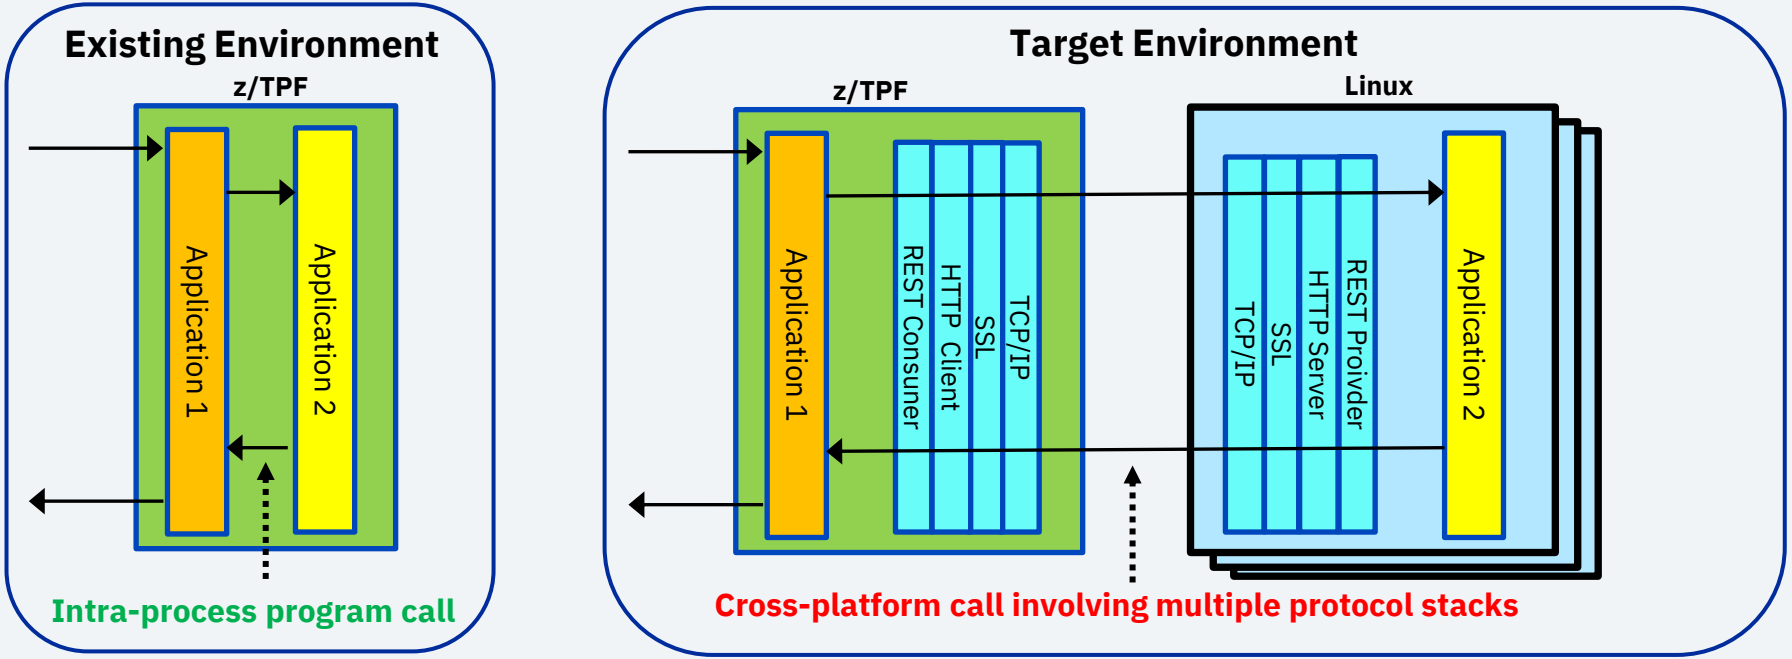

The Reality

- CPU costs on z/TPF and overall **IT costs can go up**
- Have **another platform to manage**, including High Availability (HA), D/R, security, network
- **More complex environment**, more components that could fail, harder to debug

Claim 3: Moving Code Off of zSystems Will Reduce My Costs

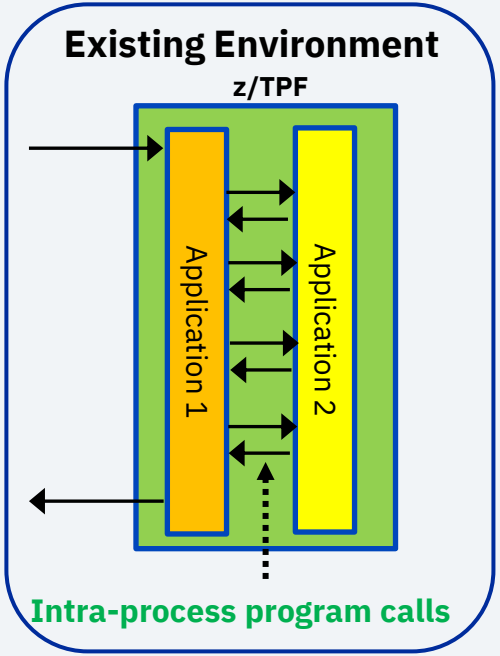

Execution Flow (Do These Steps 4 Times):

1. Application 1 makes a call to Application 2 inside the same process (ECB)
2. Application 2 executes code and can access data created/retrieved by Application 1 (same ECB)
3. Application 2 returns to Application 1 via program return call
4. Any results and data created/retrieved by Application 2 is available to Application 1 (same ECB)

More Complex Example, Chatty Application Makes Multiple Calls

Claim 3: Moving Code Off of zSystems Will Reduce My Costs

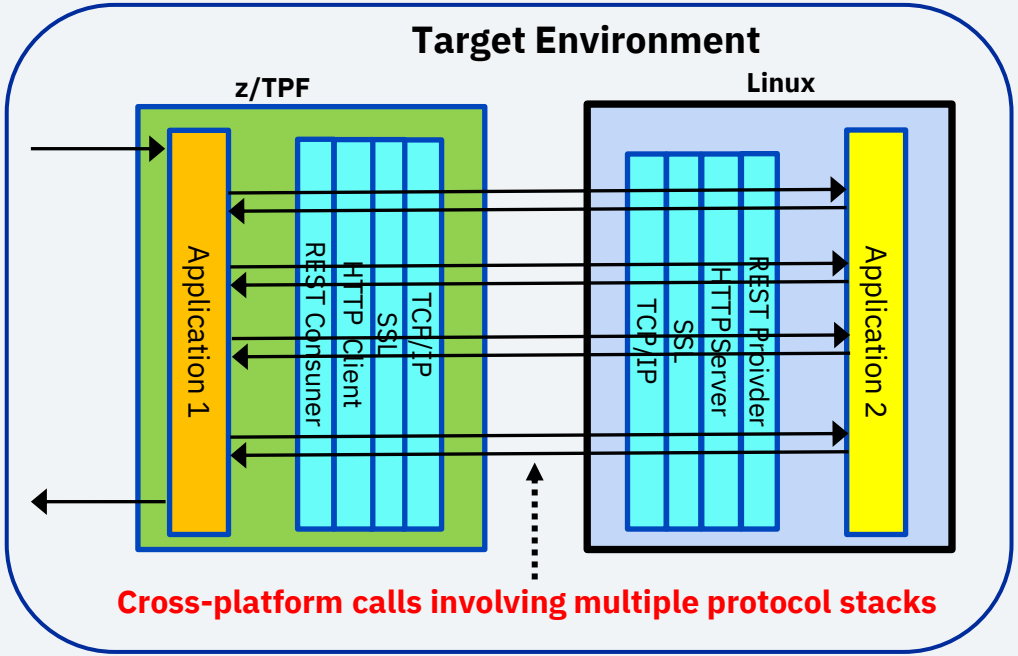

Execution Flow (Do These Steps 4 Times):

1. Application 1 makes a call to Application 2 using REST (or other middleware) passing all the data needed by Application 2
2. Middleware, security, and network stacks on z/TPF and Linux transform and pass the information to Application 2. Load balancers, firewalls, and API Management nodes may also be involved. If stateful application, sending all 4 requests to the same Linux server might be necessary.
3. Application 2 executes code and sends the response data via REST (or other middleware)
4. Middleware, security, and network stacks on Linux and z/TPF transform and pass the results to Application 1

Claim 3: Moving Code Off of zSystems Will Reduce My Costs

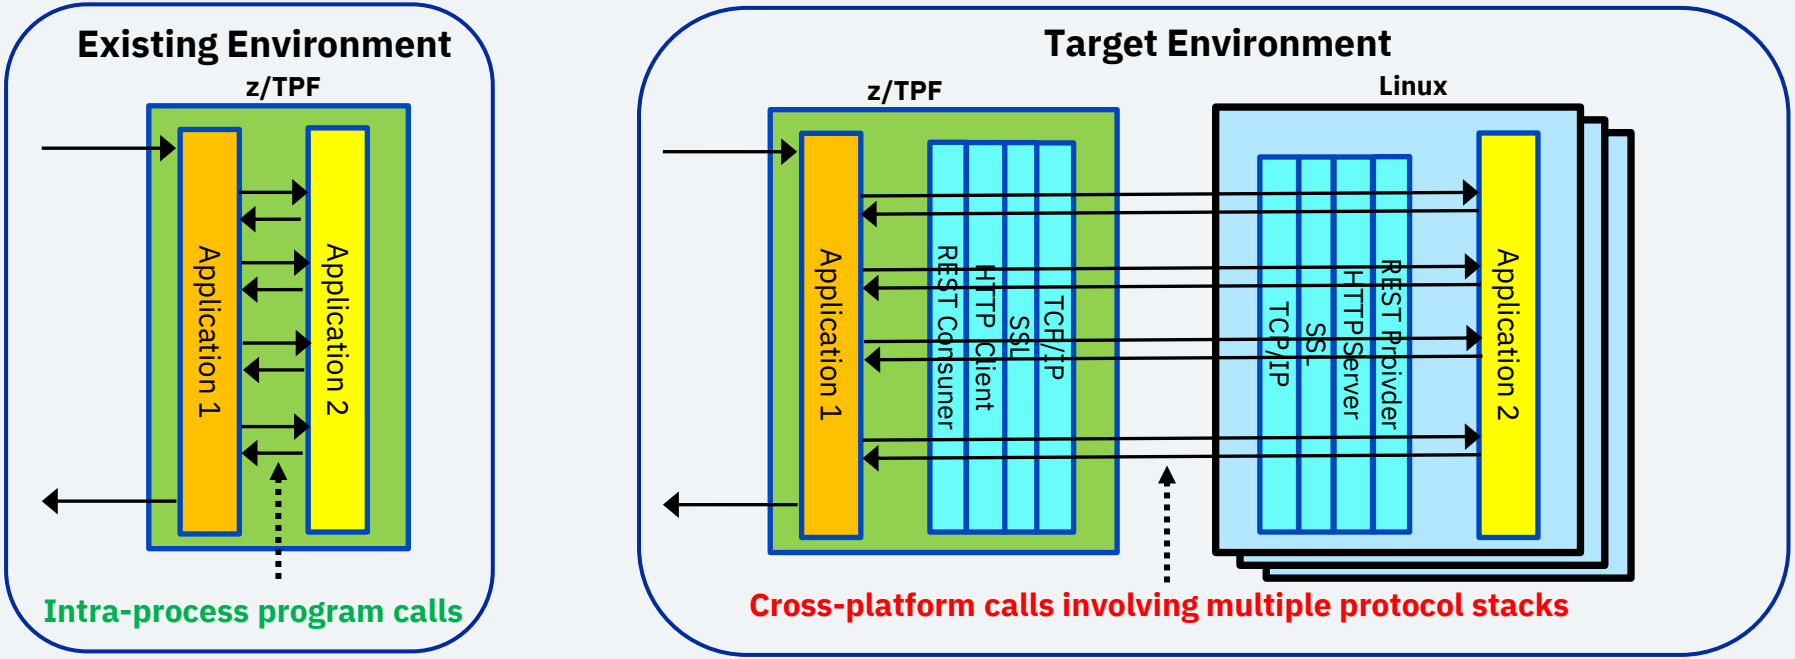

The Reality

- CPU costs on z/TPF and overall **IT costs almost always go up**
- Increased response times can **impact SLAs and ability to scale**
- **Have another platform to manage** , including High Availability (HA), D/R, security, network
- **More complex environment**, more components that could fail, harder to debug

Claim 4: I Can Move Any Application Anywhere

Existing Environment

- Real-time (latency-sensitive) workload
- Transaction involves services (applications) on 3 different platforms in the data center
  - These platforms can be private (on prem) cloud or zSystems (z/TPF, z/OS, Linux on Z)
  - Each platform can have multiple servers for high availability and ability to scale
- Diagram illustrates 5 API calls between platforms, but in reality there could be dozens of calls
- Co-location of the platforms enables acceptable transaction response time and ability to meet SLAs

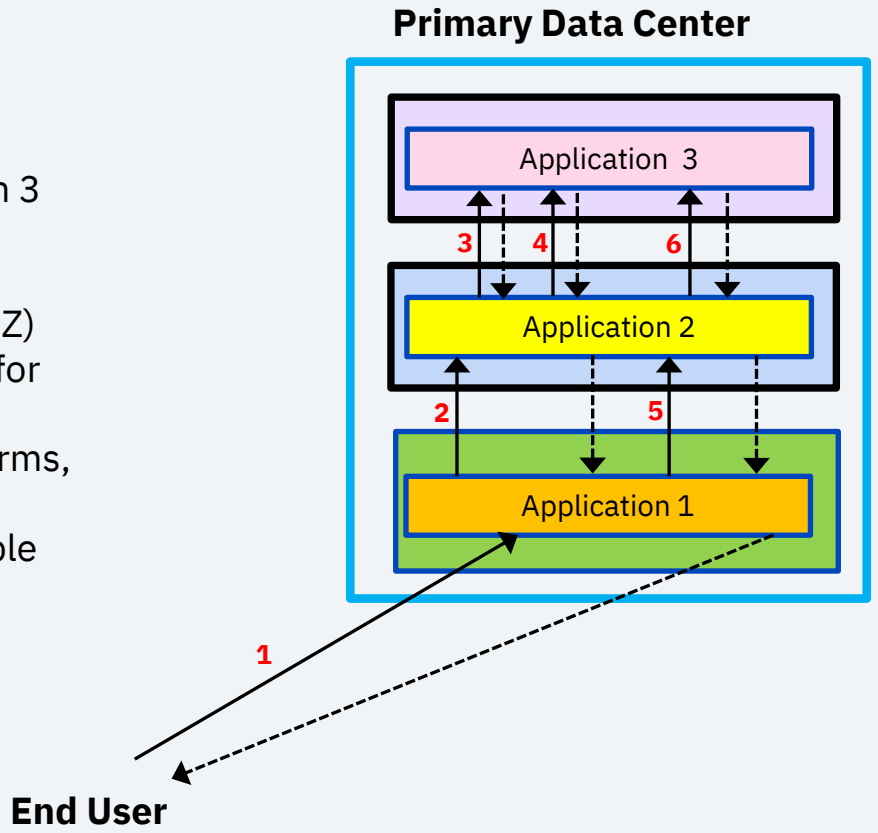

Claim 4: I Can Move Any Application Anywhere

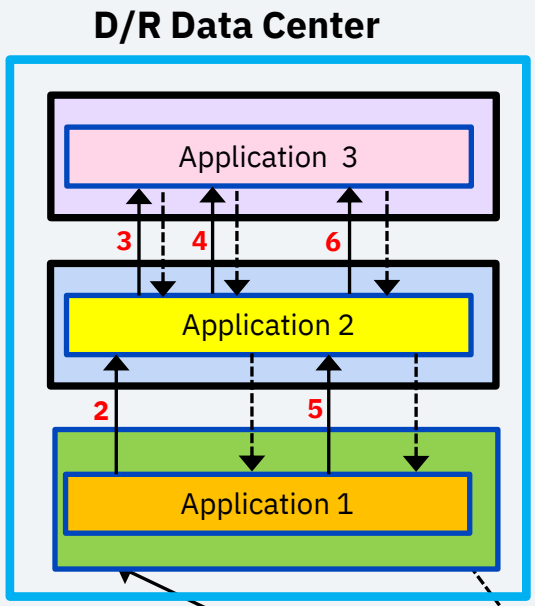

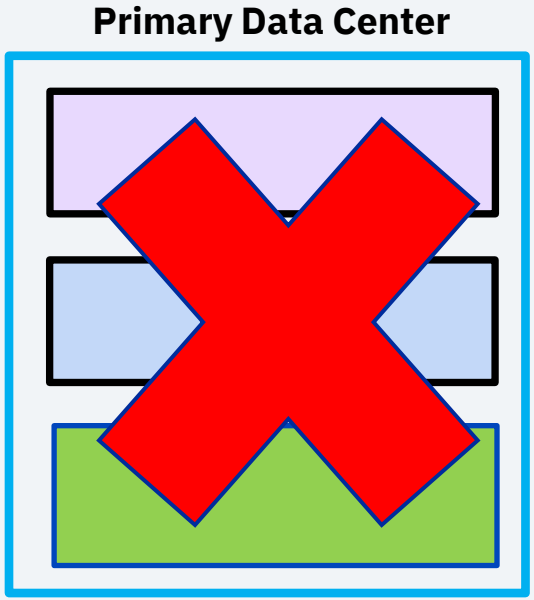

Existing Environment

If primary data center fails, the **entire workload** moves to the D/R site

End User

Claim 4: I Can Move Any Application Anywhere

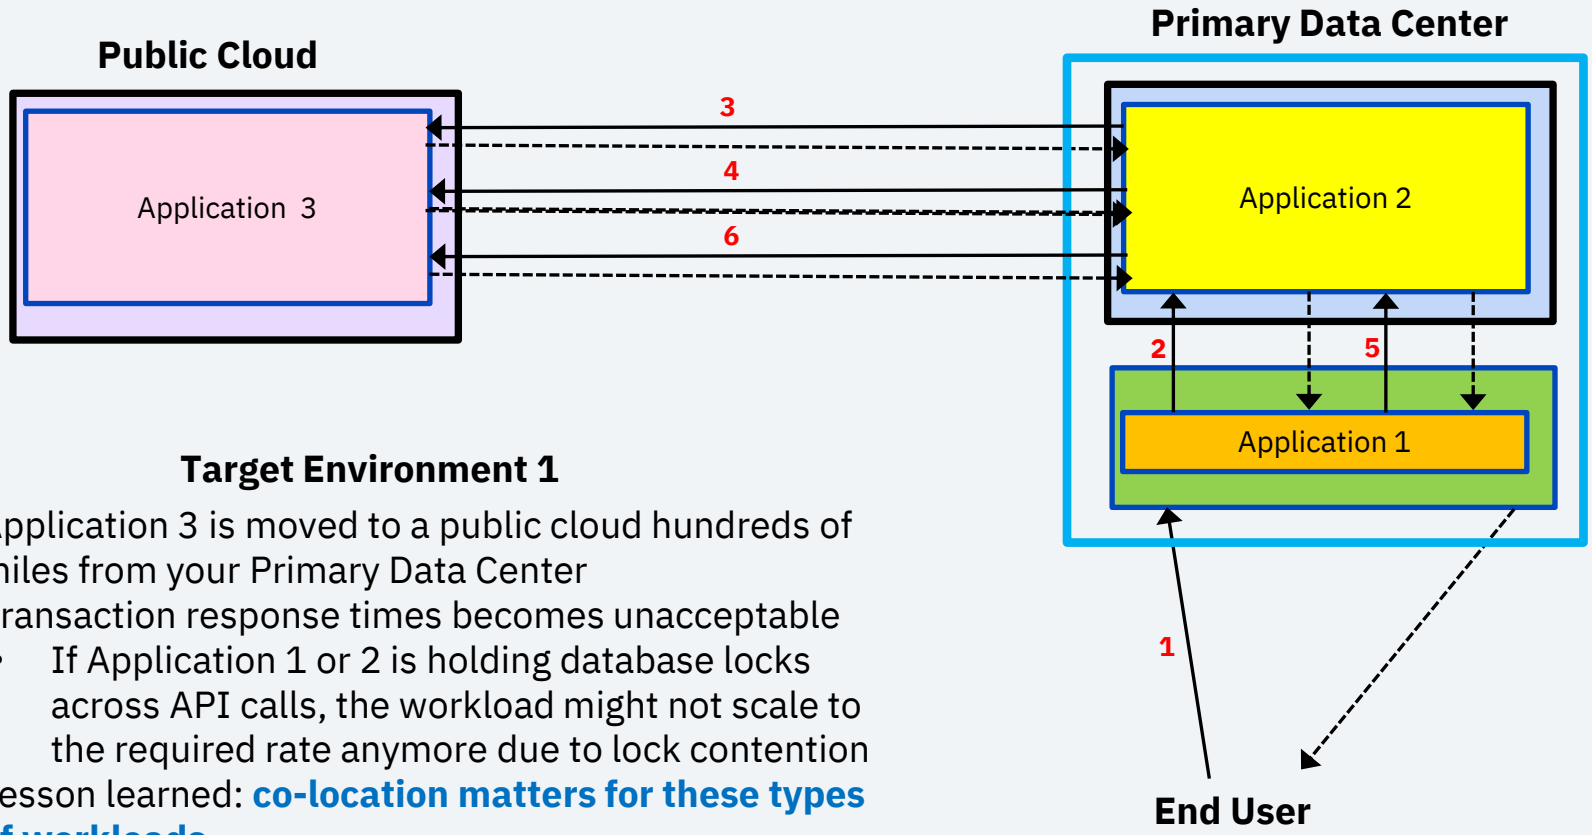

Target Environment 1

- Application 3 is moved to a public cloud hundreds of miles from your Primary Data Center
- Transaction response times becomes unacceptable
  - If Application 1 or 2 is holding database locks across API calls, the workload might not scale to the required rate anymore due to lock contention
- Lesson learned: **co-location matters for these types of workloads**

Claim 4: I Can Move Any Application Anywhere

Target Environment 2

- Applications are moved to 3 different service providers
- Assumes the 3 data centers are close enough in proximity to meet transaction response time and other SLAs for this workload

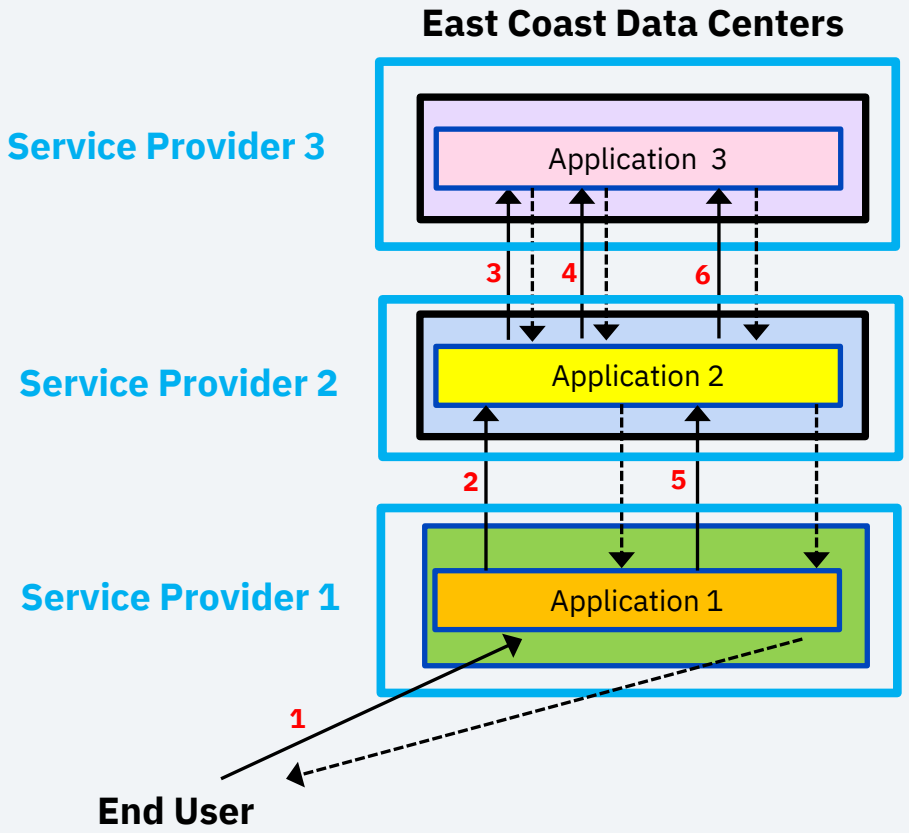

Claim 4: I Can Move Any Application Anywhere

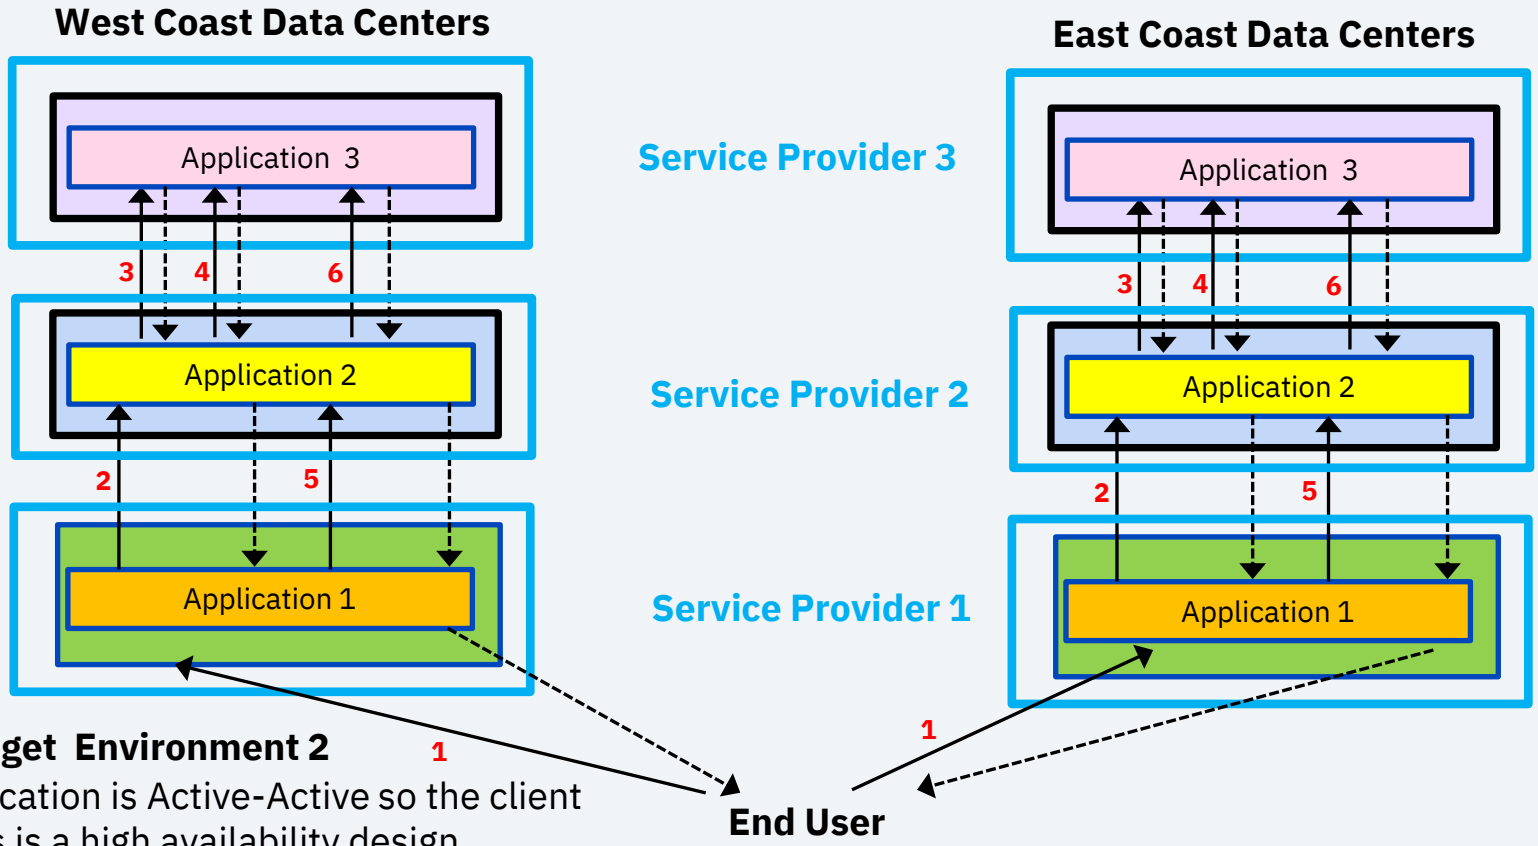

Target Environment 2

Each application is Active-Active so the client **thinks** this is a high availability design...

Claim 4: I Can Move Any Application Anywhere

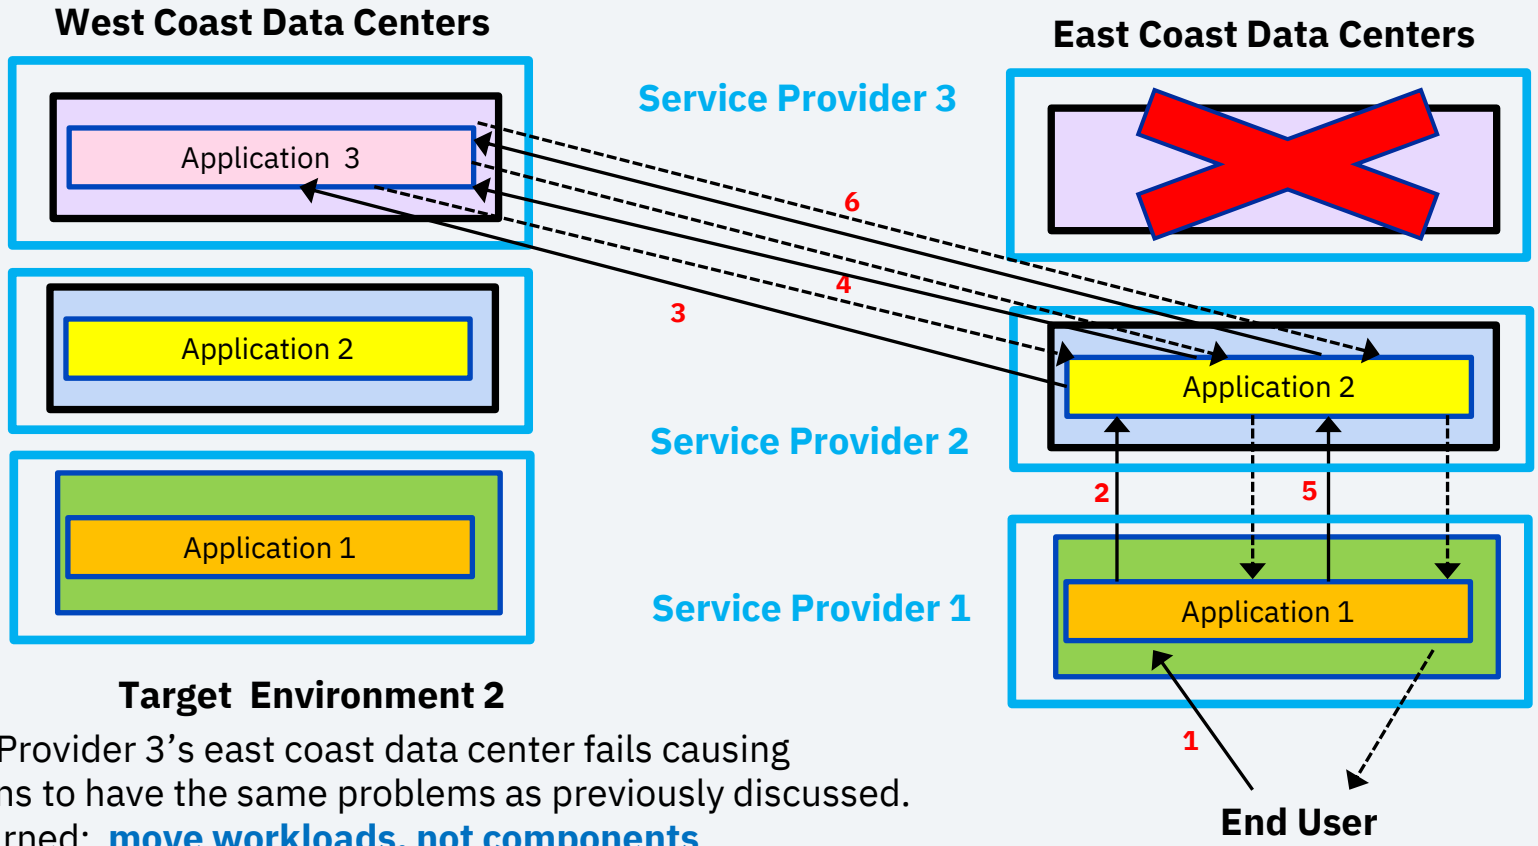

... Service Provider 3's east coast data center fails causing transactions to have the same problems as previously discussed.
Lesson learned: **move workloads, not components**

The Number One Rule

Real Estate

A similar house in a different location can **cost significantly more** and have different qualities of life

High-Volume OLTP Workloads

The same workload in a different architecture can **cost significantly more** and have **lower qualities of service**

Important to **co-locate** the **latency sensitive** parts of the workload

Not Recommended Hybrid Cloud Architectures

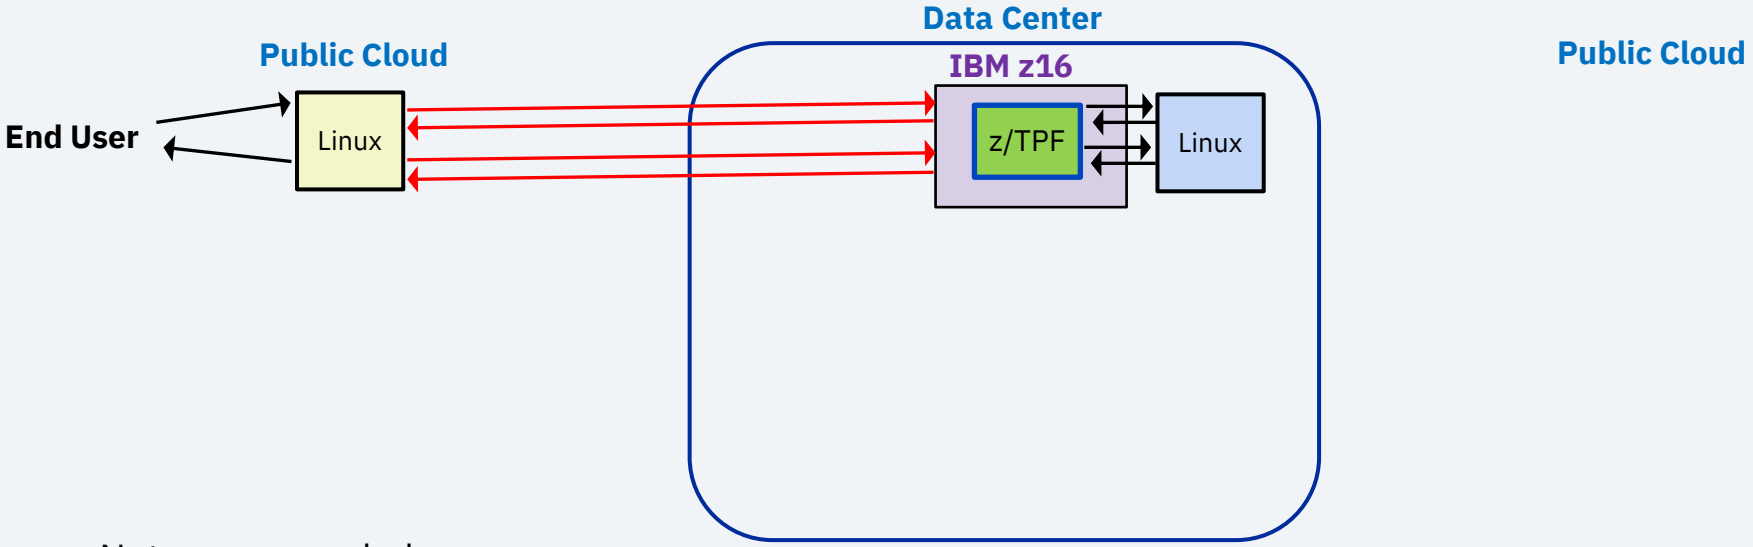

Not recommended:

1. System of Engagement (SOE) in public cloud that makes multiple calls to z/TPF

Not Recommended Hybrid Cloud Architectures

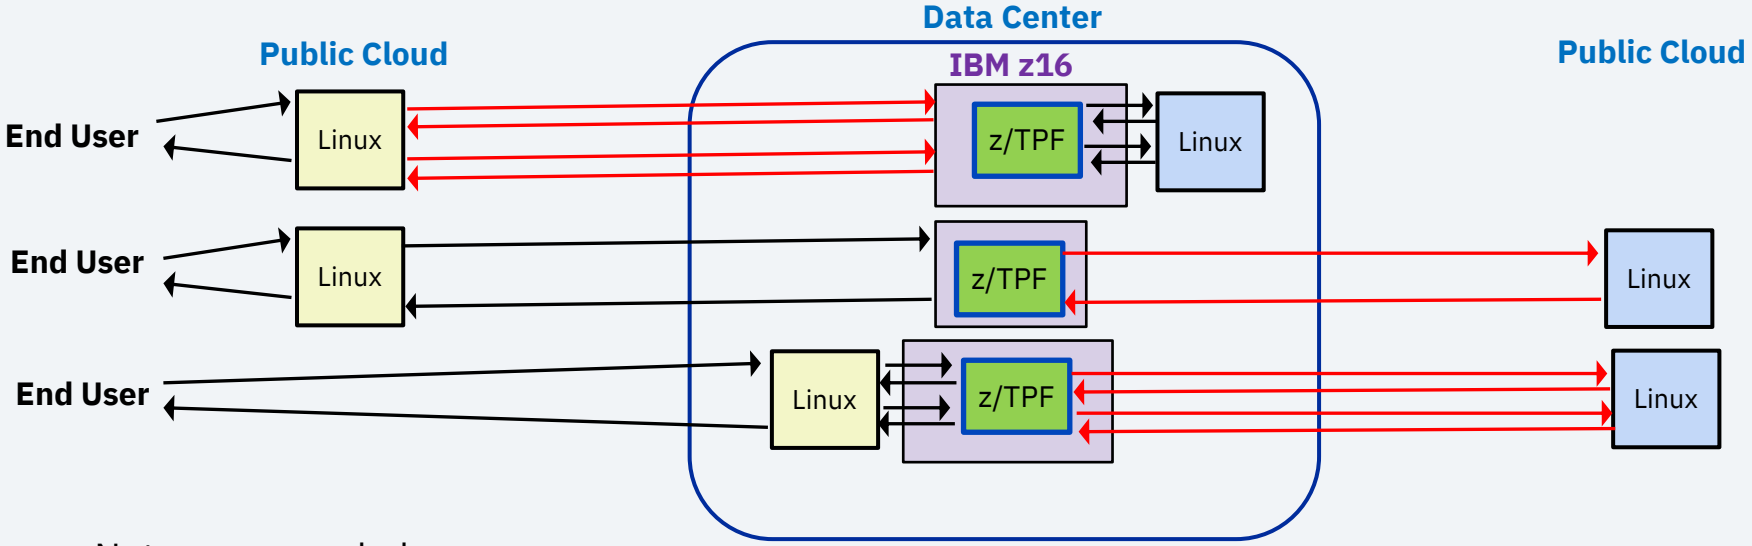

Not recommended:

1. System of Engagement (SOE) in public cloud that makes multiple calls to z/TPF
2. Services consumed by z/TPF **transactions** are in public clouds.
   - Post-transaction data feed from z/TPF (via Business Events, Kafka Producer, and so on) to data and event consumers **can be** in public cloud or on prem

Recommended Hybrid Cloud Architectural Options with z/TPF

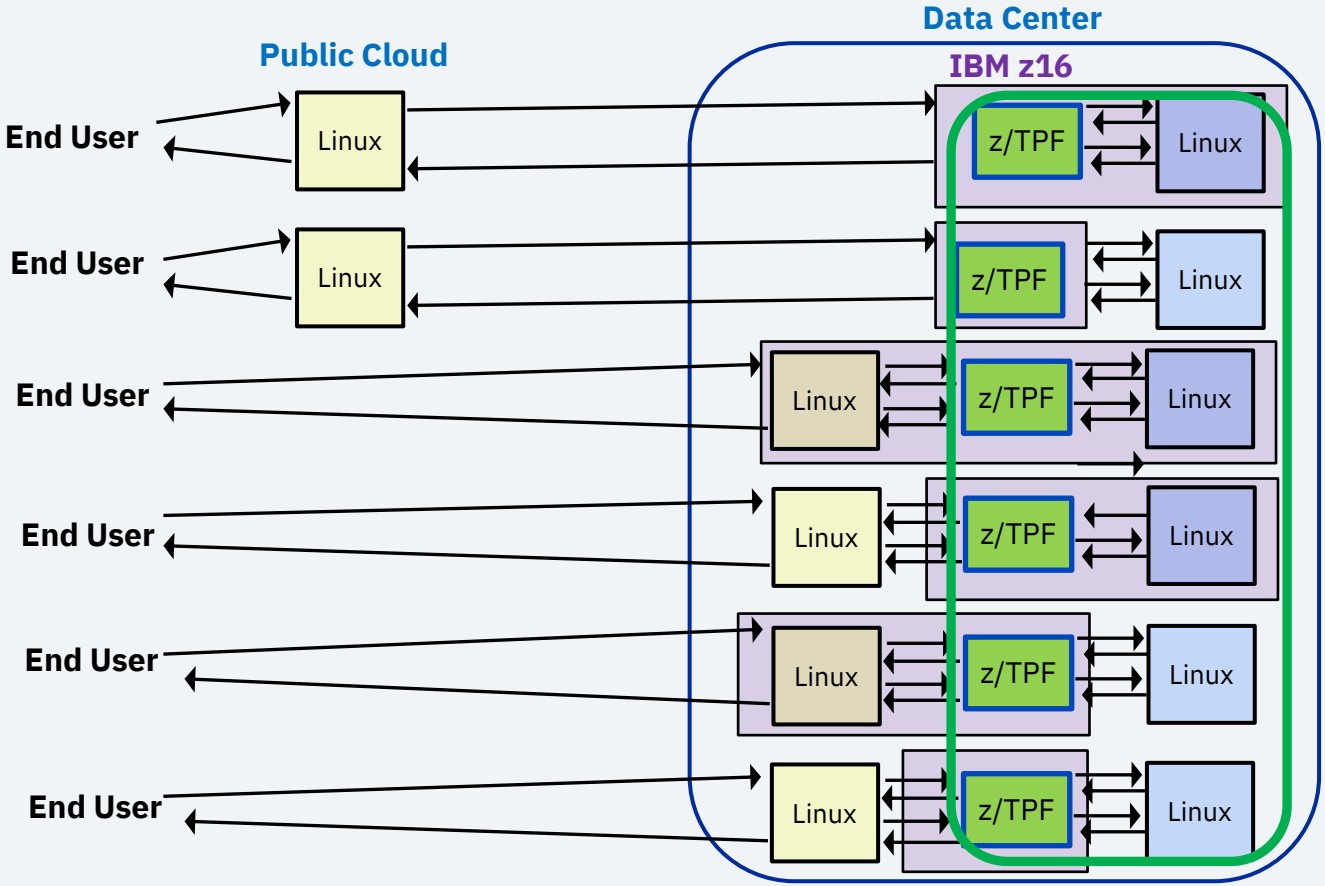

Considerations:

1. Co-locate z/TPF and services consumed by z/TPF transactions, either on Linux on Z or on prem cloud

Recommended Hybrid Cloud Architectural Options with z/TPF

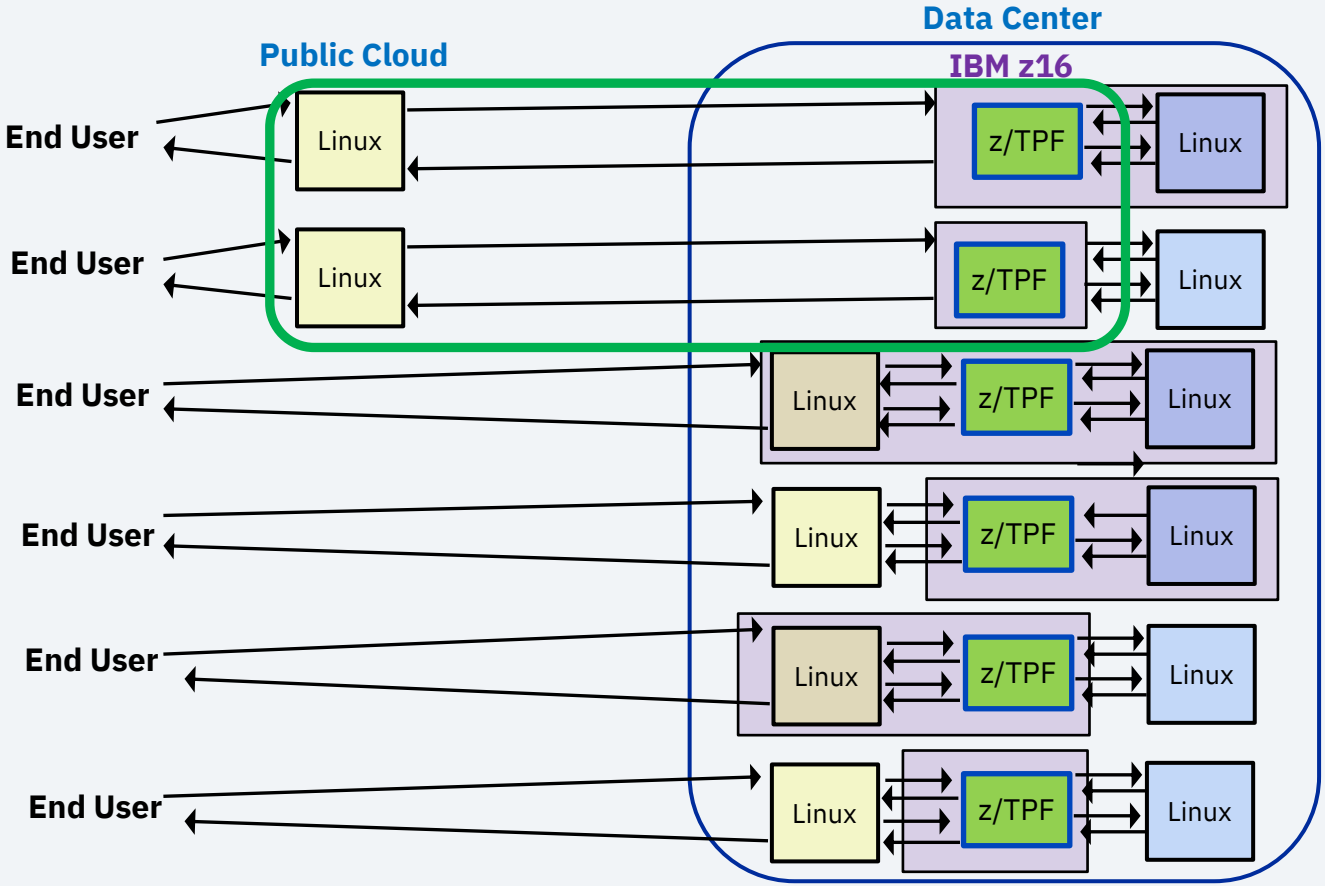

Considerations:

1. Co-locate z/TPF and services consumed by z/TPF transactions, either on Linux on Z or on prem cloud
2. If SOE makes single call to z/TPF, SOE could go in public cloud or co-locate SOE with z/TPF, either on Linux on Z or on prem cloud

Recommended Hybrid Cloud Architectural Options with z/TPF

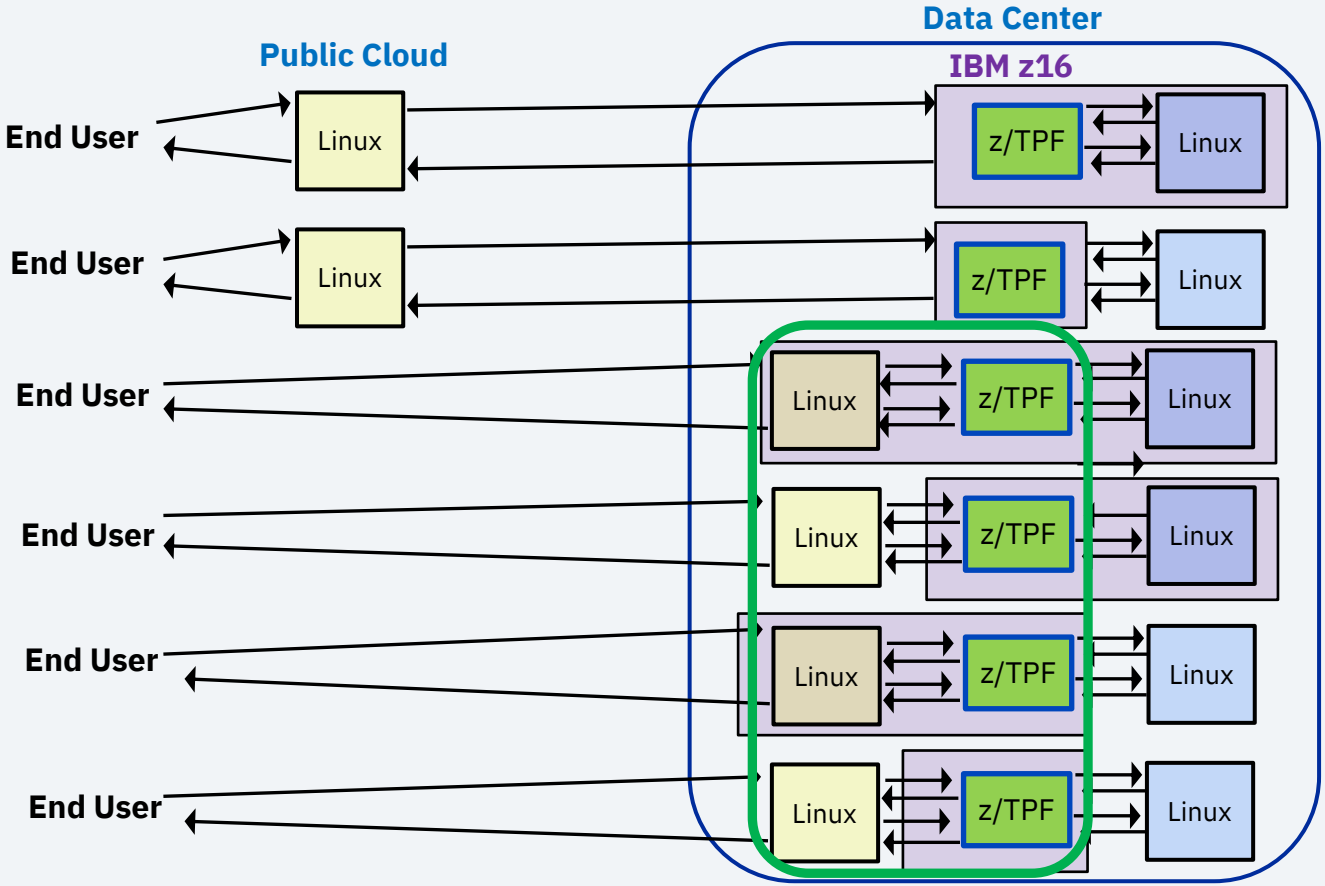

Considerations:

1. Co-locate z/TPF and services consumed by z/TPF transactions, either on Linux on Z or on prem cloud
2. If SOE makes single call to z/TPF, SOE could go in public cloud or co-locate SOE with z/TPF, either on Linux on Z or on prem cloud
3. If SOE makes multiple calls to z/TPF, co-locate SOE with z/TPF, either on Linux on Z or on prem cloud

Further Optimization Leveraging Java Support on z/TPF

- Expose a **macro**service externally that drives many **micro**services on z/TPF
- If that orchestration layer is written in Java, move that into z/TPF to reduce latency and costs

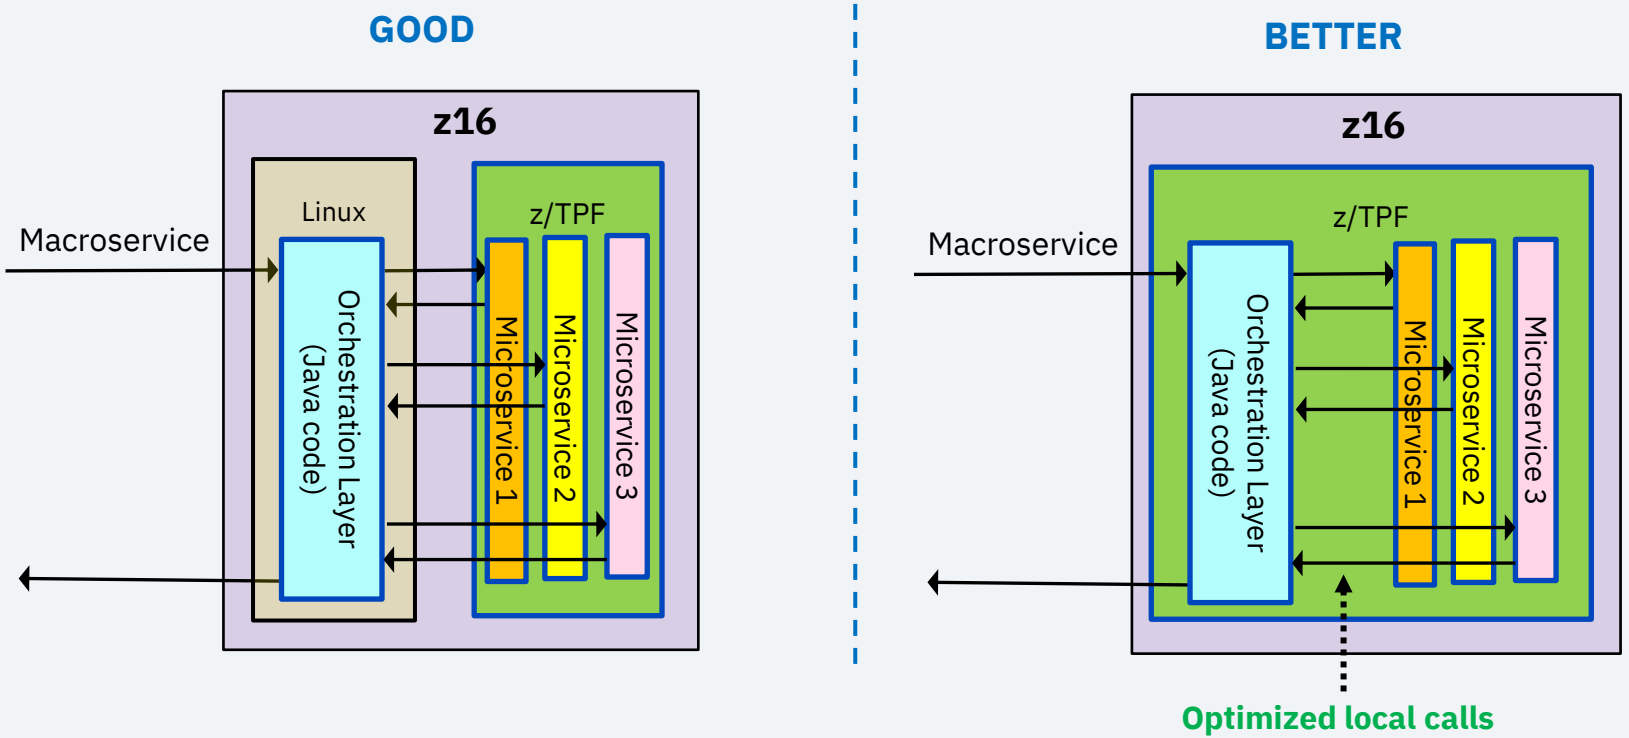

Summary

IBM z16

Predict and Automate for Increased Decision Velocity (AI)

Secure with a Cyber Resilient System

Build a Sustainable Infrastructure

Modernize with Hybrid Cloud

Architect a Hybrid Cloud solution…

That leverages the strengths of zSystems and cloud

Capable of scale, resiliency, failover, DR

Optimized for performance, latency, availability, security, cost, manageability

Thank you