
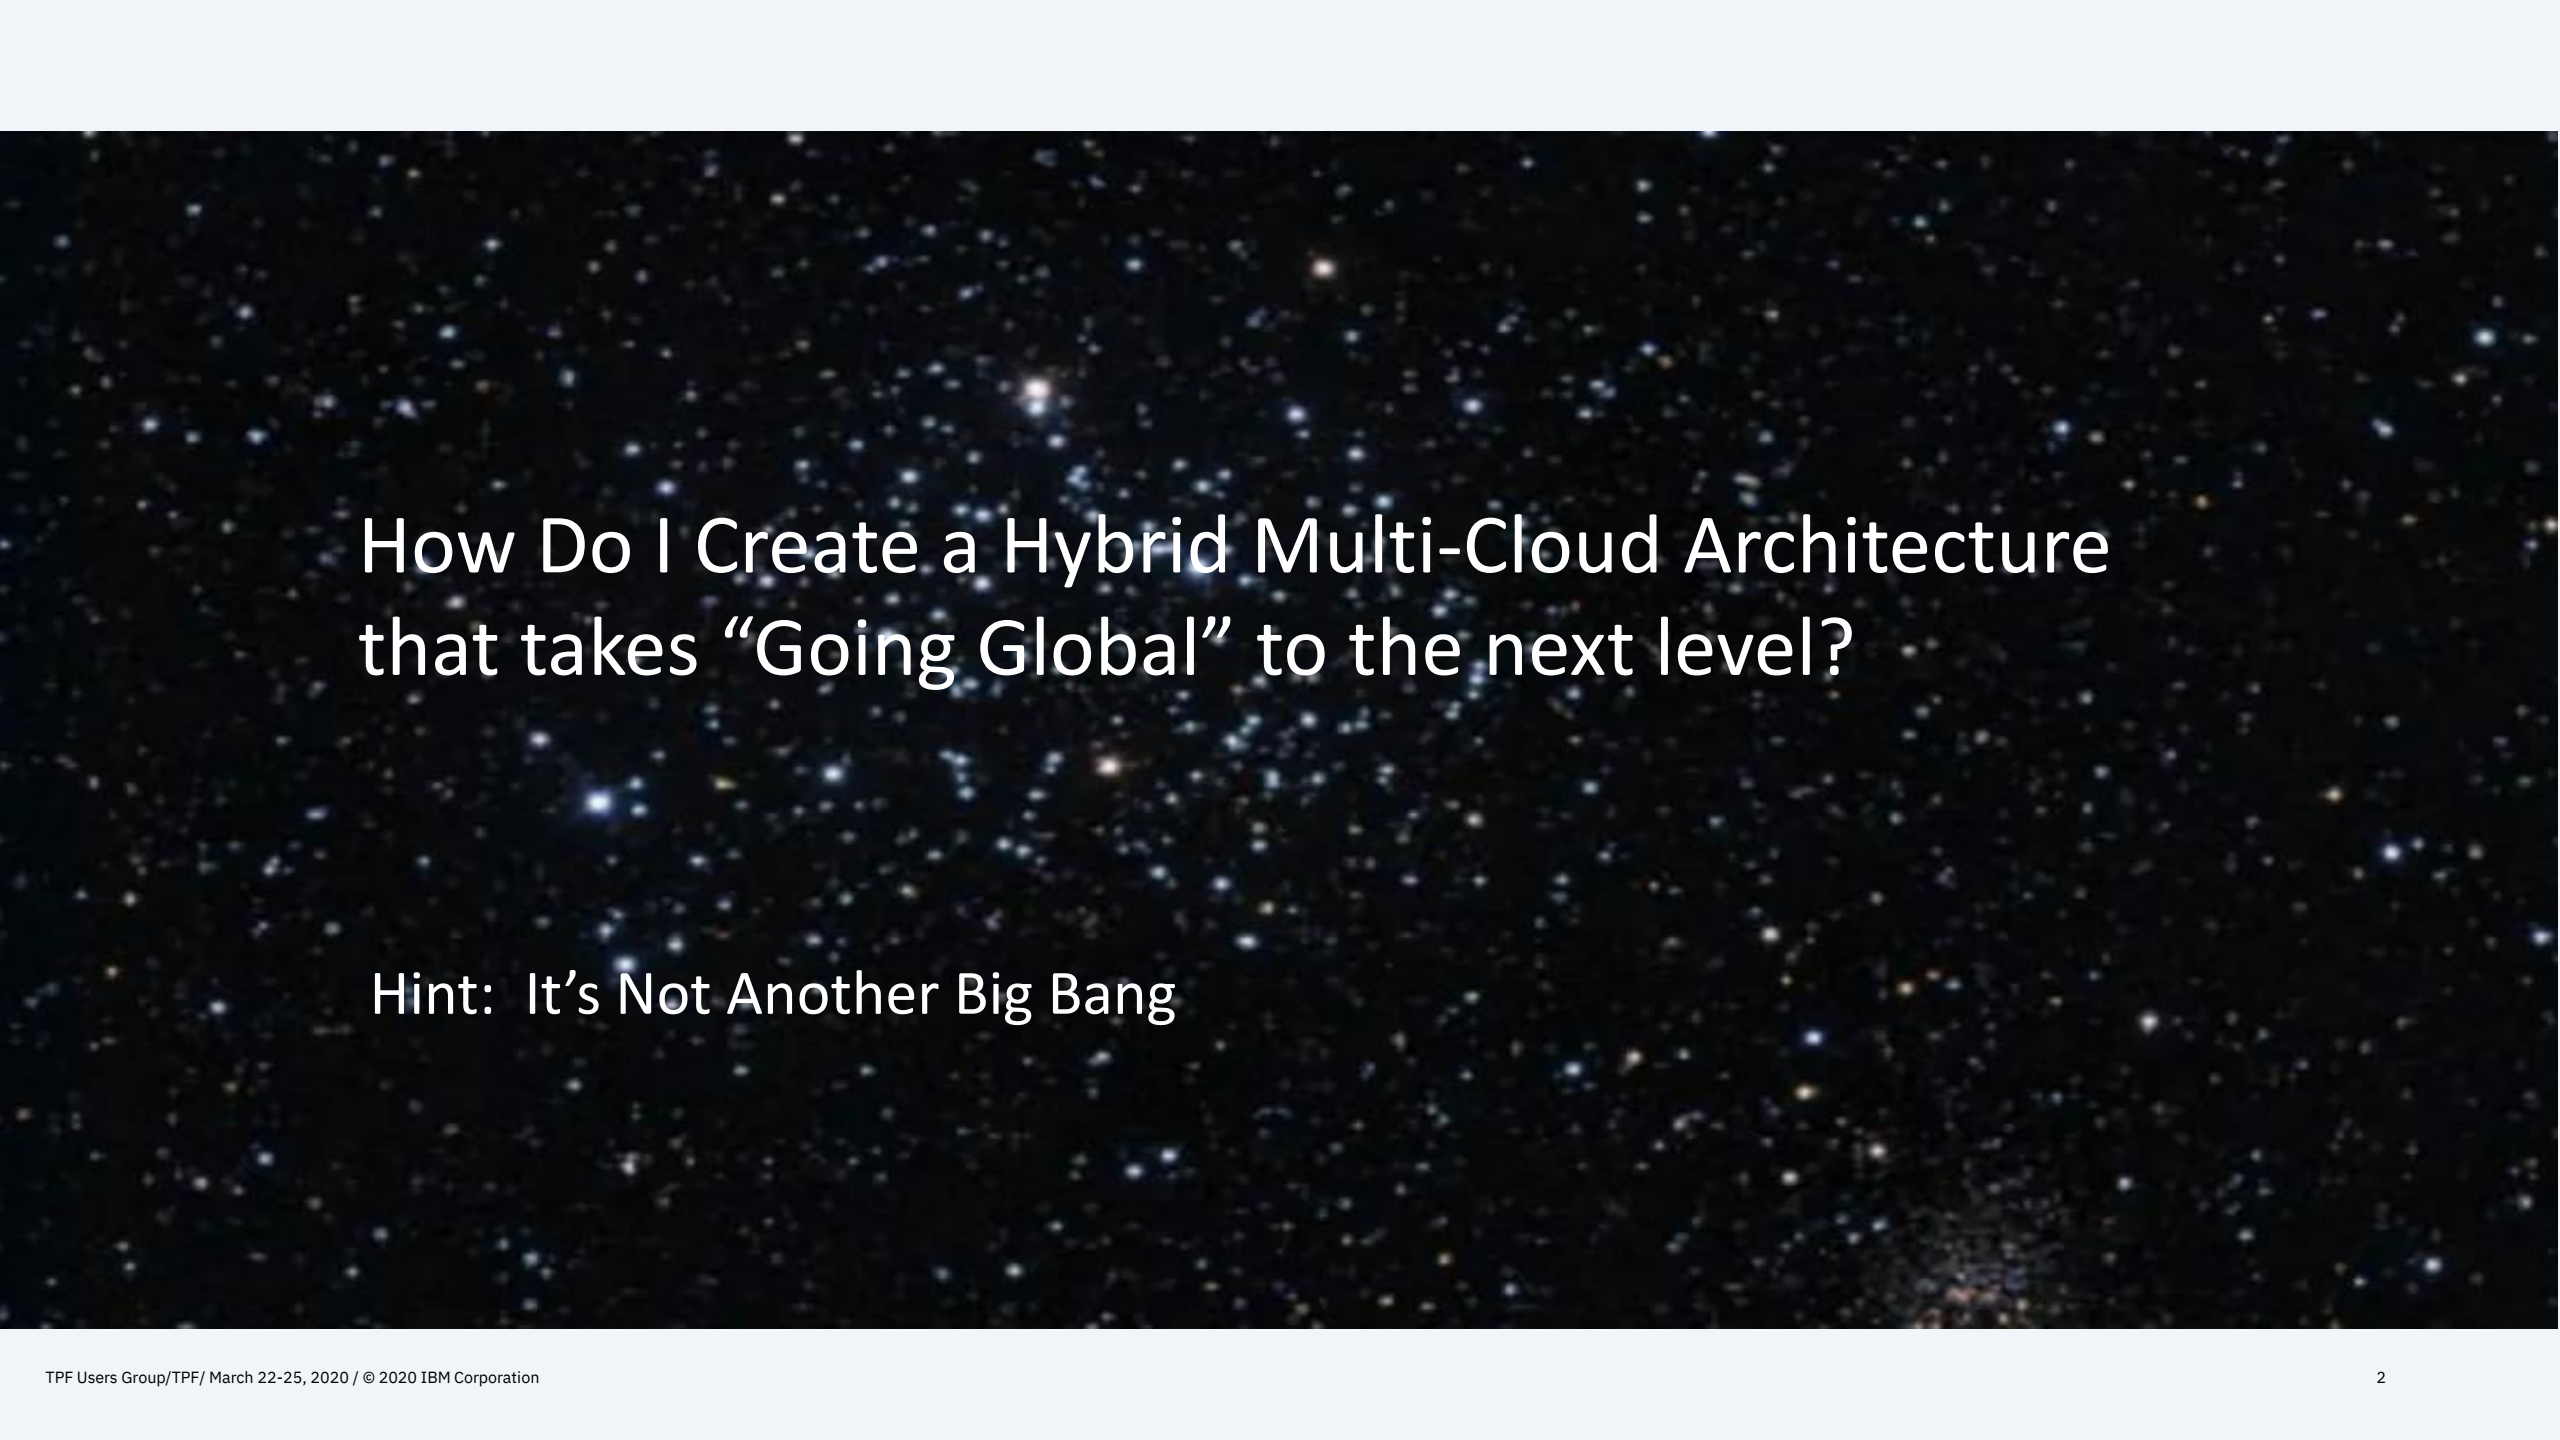


IBM z15 and Hybrid Cloud




Mark Gambino
IBM z/TPF Chief Architect



How Do I Create a Hybrid Multi-Cloud Architecture
that takes “Going Global” to the next level?

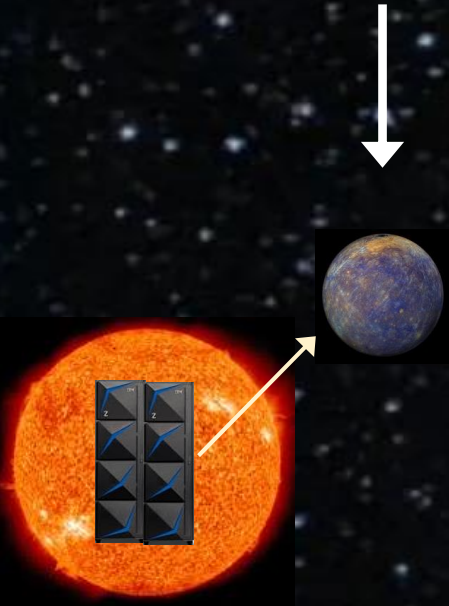
Hint: It’s Not Another Big Bang

System of Record (SOR) Data on IBM Z is Center of Your Business Universe




ILLENIAL

Keep Your Friends Close and Your Analytics Data Closer



ILLENIAL

Data Scientists are from Mars, the Rest of Your Enterprise is from Venus



ILLENIAL

The Majority of Your Users




Group of Users that Do Things a Little Differently




Group of Users that Do Things a Little Differently



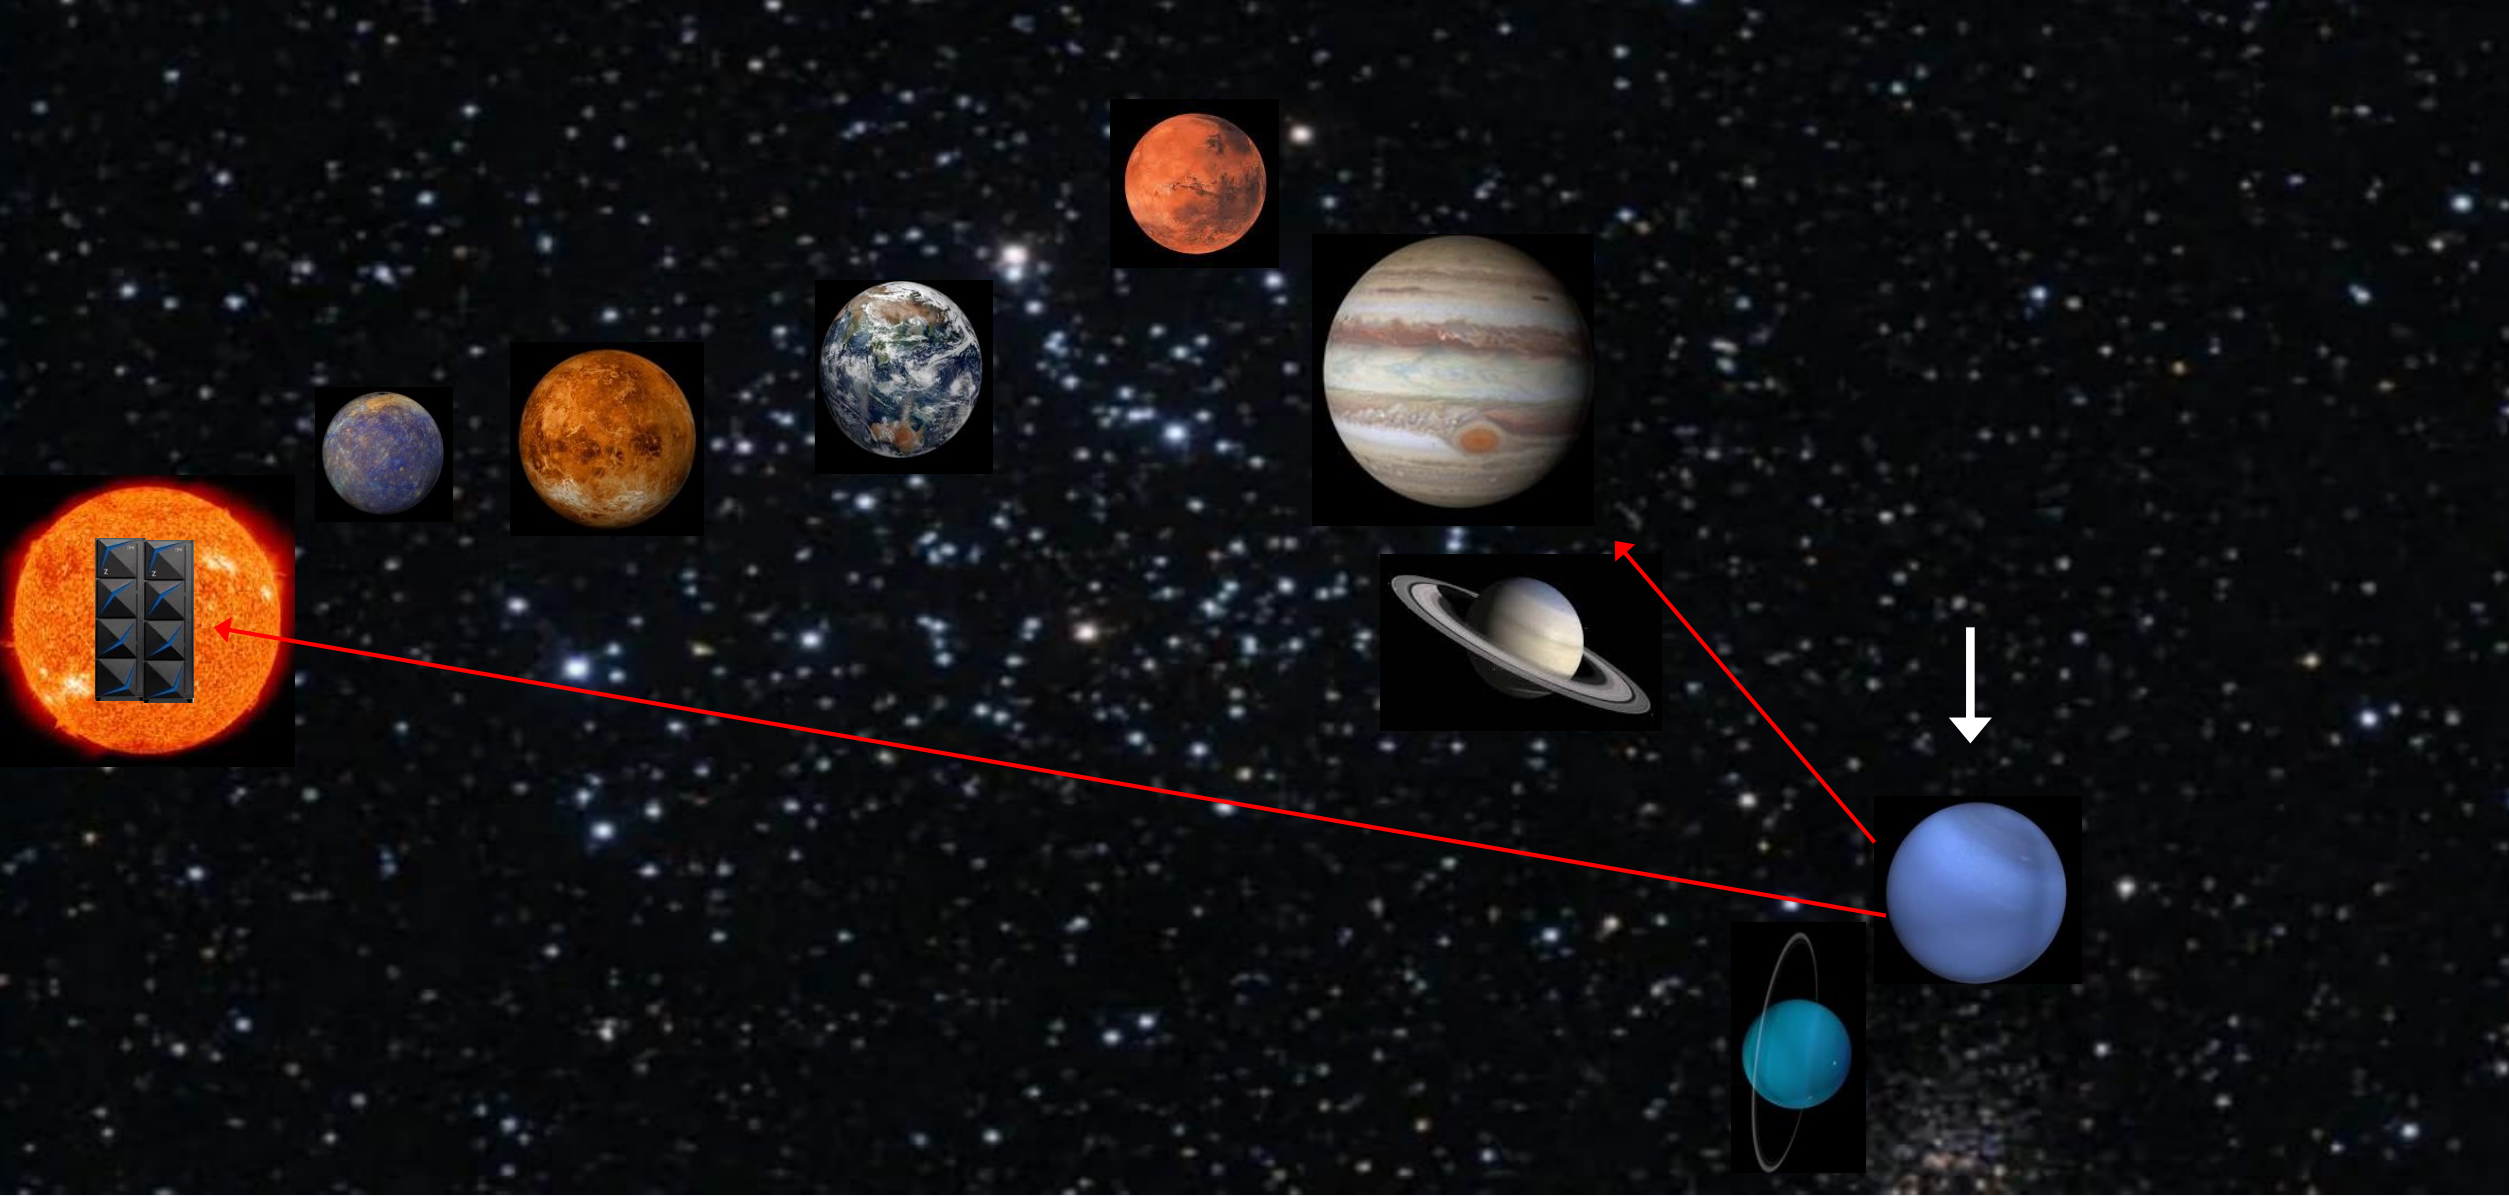
Systems of Engagement (SOE) – Gas Giant (Cloud) Systems




Business Partners that Want Data in Older Format over Older Protocols



Bad Guys in the Distance that Want Your Data and to Wreak Havoc




Is It a Workload? Is It Not a Workload?



The Heart of Your Hybrid Cloud Architecture is z/TPF on IBM Z

Availability

- Services and data
- Always consistent data
- Elasticity
- Pay-as-you-grow
- Deploy application changes on the fly
- Expand databases without downtime




Open & Connected

- REST, HTTP, Kafka, MongoDB, MQ
- JSON, XML, DFDL
- C/C++, Java
- Eclipse-based IDE
- Realtime Analytics leveraging Spark, Jupyter, Grafana, etc.

The Heart of Your Hybrid Cloud Architecture is z/TPF on IBM Z

Performance

- Ability to massively scale up and out
- Consistent low response times
- Cache terabytes of data in memory




Secure

- Data protected on disk **and** in memory
- Secure (TLS) network connections
- Encrypt everything using strong pervasive encryption
- Secure code loads and execution environment

The Heart of Your Hybrid Cloud Architecture is z/TPF on IBM Z

App Development

- Multiple UT options
- Test framework
- Thousands of Java packages exist
- Tooling to identify hot spots and resource usage/changes
- Progressively modernize



Highly Efficient

- Process multiple workloads
- Virtualize resources
- Leverage accelerators for I/O, network, crypto, compression
- Lowest cost per transaction

Question: Will IBM z15 Make My Business Even Better?

Choose your answer:

- A. Yes
- B. Absolutely
- C. Of Course
- D. Definitely

Problem Statement

Unplanned outages are painful enough, but recovering to steady state can also be time consuming and financially challenging

What your operators and coverage staff is thinking as this happens:

- Should I re-activate services in a staged manner to avoid a surge that could lead to yet another outage?
- Should I add extra CPU capacity and is there enough time to make that decision?
- How long will the outage impact my SLAs?
- Why didn't I take a vacation day today?

z15 System Recovery Boost (SRB) reduces the impact of outages

- **Extra CPU capacity** for bringing z/TPF back up faster and to process the transactional backlog once the system is back up
 - You can activate **all services** immediately following the IPL
- **Automatically done** - no operator intervention needed
- **No charge** for using the extra CPU capacity
- **First in industry** type solution

System Recovery Boost Types for z/TPF

- **Speed Boost**

- Applicable only if z/TPF is running on a subcapacity machine (400, 500, or 600 series machine)
- During the boost period, z/TPF LPAR runs at full speed (700 series speed)

- **I-Stream Boost**

- Enables all fenced I-streams to be used during the boost period
- Requires Dynamic CPU support such that there are extra (fenced) engines available for use

- Can use one, both, or neither boost

- Boost period is from when z/TPF needs to IPL until a fixed amount of time after the system is processing work again (back in NORM state)

Conclusion

System Recovery Boost allows you to:

- Minimize the business impacts of unplanned outages
- Reduce the time needed for planned outages

Imagine You Own an Auto Salvage Yard



Ugly old car

250,000 cubic inches



Car Compactor



Recyclable
Materials

50,000 cubic inches

Your Revenue is Directly Related to How Quickly You Can Crush Cars

Now What If I Told You The Following Was Available



Car Crusher 2.0

Can Crush Cars 100x Faster!

Background

Data Compression

- Would like to compress data before transmitting over the network to reduce network bandwidth requirements
- z/TPF has supported the *zlib* standard data compression library since 2008

Problem Statement

Compressing data using SW is expensive in terms of CPU and time

For example, on a z14 using z/TPF zlib SW compression:

- Compressing 64K of data takes around 3.2 milliseconds
- Compressing 256K of data takes around 12.8 milliseconds


Only cost effective for very large messages

z/TPF now supports HW data compression when running on z15

- **On chip** HW compression that is incredibly fast!
 - Integrated Acceleration for z Enterprise Data Compression
- **No application changes required** for code using zlib
 - z/TPF will automatically use HW compression if running on z15
- New tpf_compress API created for even higher performance

Technical Details

Compressing 64K of Data Example



Conclusion

- z15 allows z/TPF to compress data **10 to 498 times faster** than was possible on z14
 - The larger the data size, the higher the savings
- Now cost effective to compress data as low as 4K in size
- Compressing data is super fast, but expanding compressed data is up to 3x faster than that
 - Depends on the complexity of the data
- Opens up many new use cases
 - z/TPF systems code
 - Including, but not limited to, networking
 - Application level

Conclusion

$$1 + 1 = 5$$

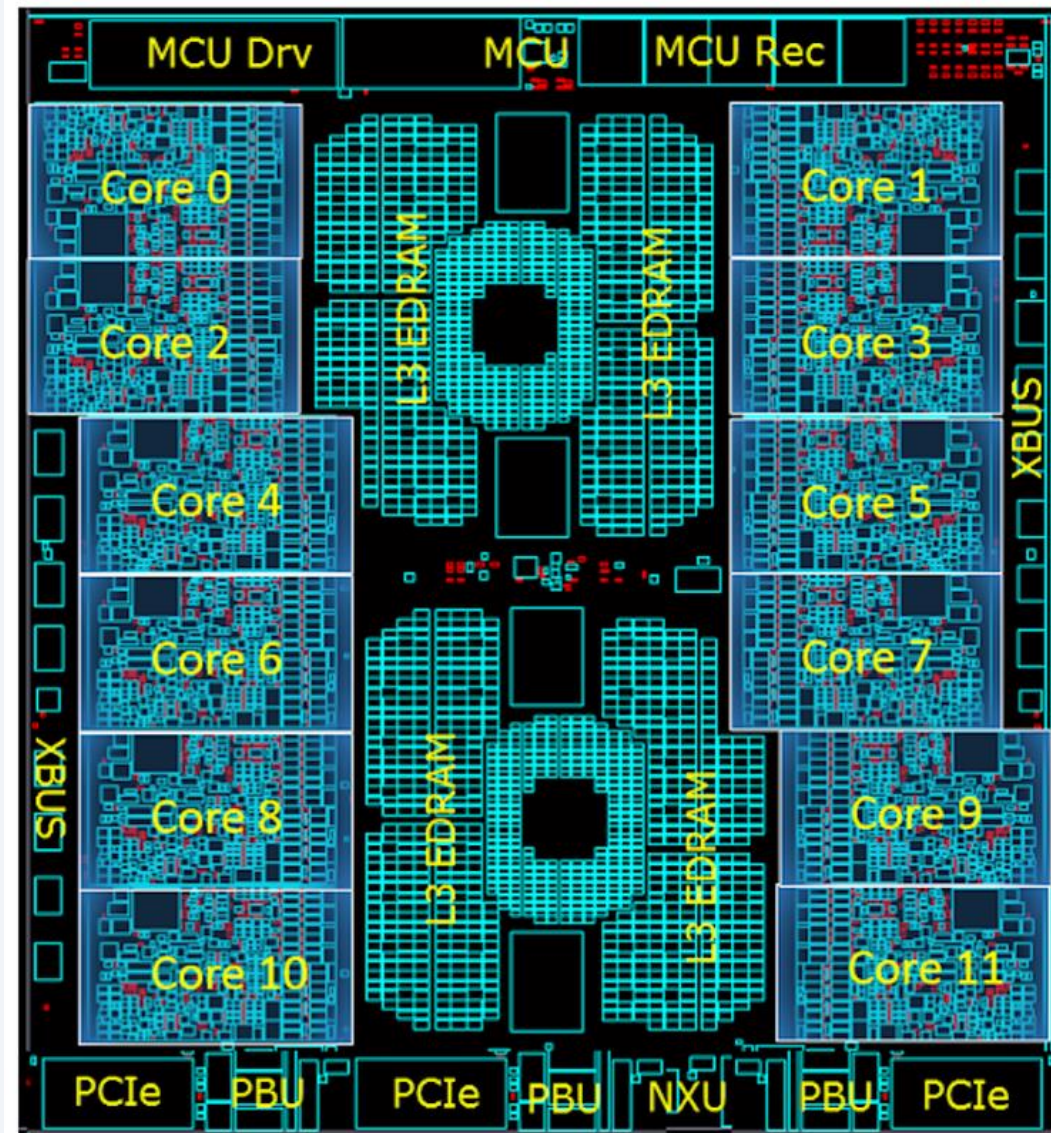
**Pervasive
Compression**

**Pervasive
Encryption**

- Less network bandwidth used
- Reduction in CPU used
- Fewer DASD I/Os
- Less memory used
- Better transaction response time

Lower IT Costs and Better Performance!

IBM z15 Central Processor (CP) Chip



9.2 billion transistors!

Technical Details

	z14	z15	z15 Improvement
Configurable Cores	170	190	12%
Cores per CP Chip	10	12	20%
L1 cache size per core	128 KB (I-cache) 128 KB (D-cache)	128 KB (I-cache) 128 KB (D-cache)	
L2 cache size pre core	2 MB (I-cache) 2 MB (D-cache)	4 MB (I-cache) 4 MB (D-cache)	2x 2x
L3 cache size per chip	128 MB	256 MB	2x
L4 cache size per drawer	672 MB	960 MB	43%
Max memory size	32 TB	40 TB	25%
Execution stream	10 instructions/cycle	12 instructions/cycle	20%
2 GB Page TLB Entries	64	256	4x
Rack Size	24"	19"	Standard Size

z15 Related z/TPF APARs

- **APAR PJ45474 (April 2019)**
 - Needed if you are using new **OSA-Express7S** adapter
- **APAR PJ45850 (September 2019)**
 - Needed if you are using new **Crypto Express7S** adapter
- **APAR PJ45872 (October 2019)**
 - Needed if want to exploit **HW compression** in C/C++ code
 - zlib and z/TPF unique compression APIs
- **APAR PJ45880 (October 2019)**
 - Needed if want to exploit **HW compression** in Java code
- **APAR PJ45739 (November 2019)**
 - Needed if want to exploit **System Recovery Boost (SRB)**

Bloomberg

IBM Launches New Mainframe Generation to Combat Cybercrime

Forbes

IBM Galvanizes Its Place In Secure Workloads With New z15 Platform

TechCrunch

The mainframe business is alive and well, as IBM announces new z15

The Motley Fool

IBM Unveils z15 Mainframe, With a Focus on Data Privacy

6. IBM has introduced its next generation of mainframes that it contends will combine the speed and power of the Big Iron with the cost-savings and efficiency of the **hybrid cloud**, **data privacy** and **security controls** for modern workloads. The z15 mainframe systems have 25 percent more memory, and will cut customer costs in half, compared to the old x86-based servers. But it's the cloud, data privacy and security features that distinguish the z15, which can encrypt every file, restrict who can access the keys, and run 1 trillion encrypted web transactions per day.

[INSIDE CLOUD]



DOW JONES

WSJ PRO

CYBERSECURITY

IBM Debuts Tech Meant to Reduce Third-Party Breaches

computing

Encrypt data everywhere ... z15 supports **cloud-native application development** ... *Instant Recovery*, a new feature enabling the system to recover more quickly from both planned and unplanned downtime.



IBM z15: Multicloud makes the case for why mainframes still matter



FUTURUM.

Instant Recovery capability is good news for IBM clients and prospects but bad news for key rivals such as Oracle, Dell, HPE, Microsoft, and Dell

NETWORKWORLD

IBM z15 mainframe, amps-up cloud, security features

YAHOO! FINANCE


The three big ideas that came out of this were instant recovery, cloud native application development, and data privacy. The IBM z15 mainframe has unprecedented encryption, unprecedented security and privacy.



IBM z15 - The Enterprise Platform for Mission-Critical Hybrid Multi-Cloud

Workloads:

- Traditional Applications
- Cloud Native Applications
- Realtime Analytics




Industry Leading:

- Data Privacy and Security
- Business Continuity and Resilience
- Hybrid Cloud Readiness
- Performance
- Platform Management

Thank You

Questions? Comments?



Virtual TPFUG Q&A (Part 1 of 2)

Summary of Q&A from the virtual TPFUG event:

Question	Answer
Q: I was not aware that z/TPF supported pervasive encryption. Is that just for z/TPFDF or all file systems?	A: Currently, automatic DB encryption is just for z/TPFDF databases. However, there are general use APIs that your applications can use to encrypt data before filing that data in other (non-z/TPFDF) databases.
Q: Are there any thoughts to expand automatic DB encryption to include traditional (FIND/FILE) databases too?	A: Yes. Originally, we only did automatic DB encryption for z/TPFDF databases because we could do so in a way that did not increase DASD IO's. When encrypting a DB record, additional information needs to be saved along with the encrypted data such that the system will know how to decrypt the data later on. With FIND/FILE databases, there is no room in existing DB record to save that additional information; therefore, the additional information would need to be saved in a different record on disk; therefore, accessing those DBs would require double the amount of IO. The highly efficient z15 HW compression capability gives us a new idea. We can compress the data in a FIND/FILE DB record to free up space in that record to save the information needs for decryption. The end result is that all of that would fit into that one 4K block; therefore, automatic encryption of FIND/FILE DBs would be possible with no increase in IO's. If this is something your company is interested in, please contact us.
Q: Will the new tpf_compress function work with software compression on a machine older than z15?	A: No, the new tpf_compress function will result in an error if issued on a machine prior to z15. You must use zlib for data compression on older machines. zlib can decompress data compressed with the new tpf_compress API.

Virtual TPFUG Q&A (Part 2 of 2)

Summary of Q&A from the virtual TPFUG event:

Question	Answer
Q: Network Traces with Compression and Encryption... Is there a way to decrypt and decompress data in those traces?	A: No. By definition, data that is encrypted using TLS (SSL) cannot be decrypted after the fact because for security purposes the private keys and transient session keys are not available to any tools. IBM uses network traces to debug problems at the comms (TCP/IP) level, and is not interested in the user data. If you are looking to log application messages, network traces are not the right method to do that anymore, and have not been for some time once secure (TLS) connection started to be mandated, and compression is now another reason. Standard middleware like HTTP (which includes REST) and MQ have user exits where message requests and responses are in the clear that you can use to log using whatever logging mechanism you would like. Later on in this conference you will see an example of how to use logging package.
Q : Can the new z/TPF compress/decompress APIs be used for sending messages between different z/TPF systems?	A: In the conference you will hear about some standard middleware (HTTP and MQ) that we plan to update to use compression so if you use those protocols you can send compressed data to another z/TPF system or other platform. In addition, you can use the compression APIs to send compressed data over your own proprietary protocols too.



Trademarks

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

Notes

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

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

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

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

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

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

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