Parallel Sysplex and Data Sharing turn 30! : A Retrospective and Lessons Learned

Peter Enrico

Enterprise Performance Strategies, Inc.

Peter.Enrico@epstrategies.com

performance.questions@epstrategies.com

EPS: We do z/OS performance…

- Pivotor - Reporting and analysis software and services
  - Not just reporting, but analysis-based reporting based on our expertise
- Education and instruction
  - We have taught our z/OS performance workshops all over the world
- Consulting
  - Performance war rooms: concentrated, highly productive group discussions and analysis
- Information
  - We present around the world and participate in online forums

z/OS Performance workshops available

During these workshops you will be analyzing your own data!

- WLM Performance and Re-evaluating Goals
  - February 19-23, 2024
- Parallel Sysplex and z/OS Performance Tuning
  - August 20-21, 2024
- Essential z/OS Performance Tuning
  - September 16-20, 2024
- Also… please make sure you are signed up for our free monthly z/OS educational webinars! (email contact@epstrategies.com)

Pivotor's Comprehensive Report Sets for Immediate Performance Analysis

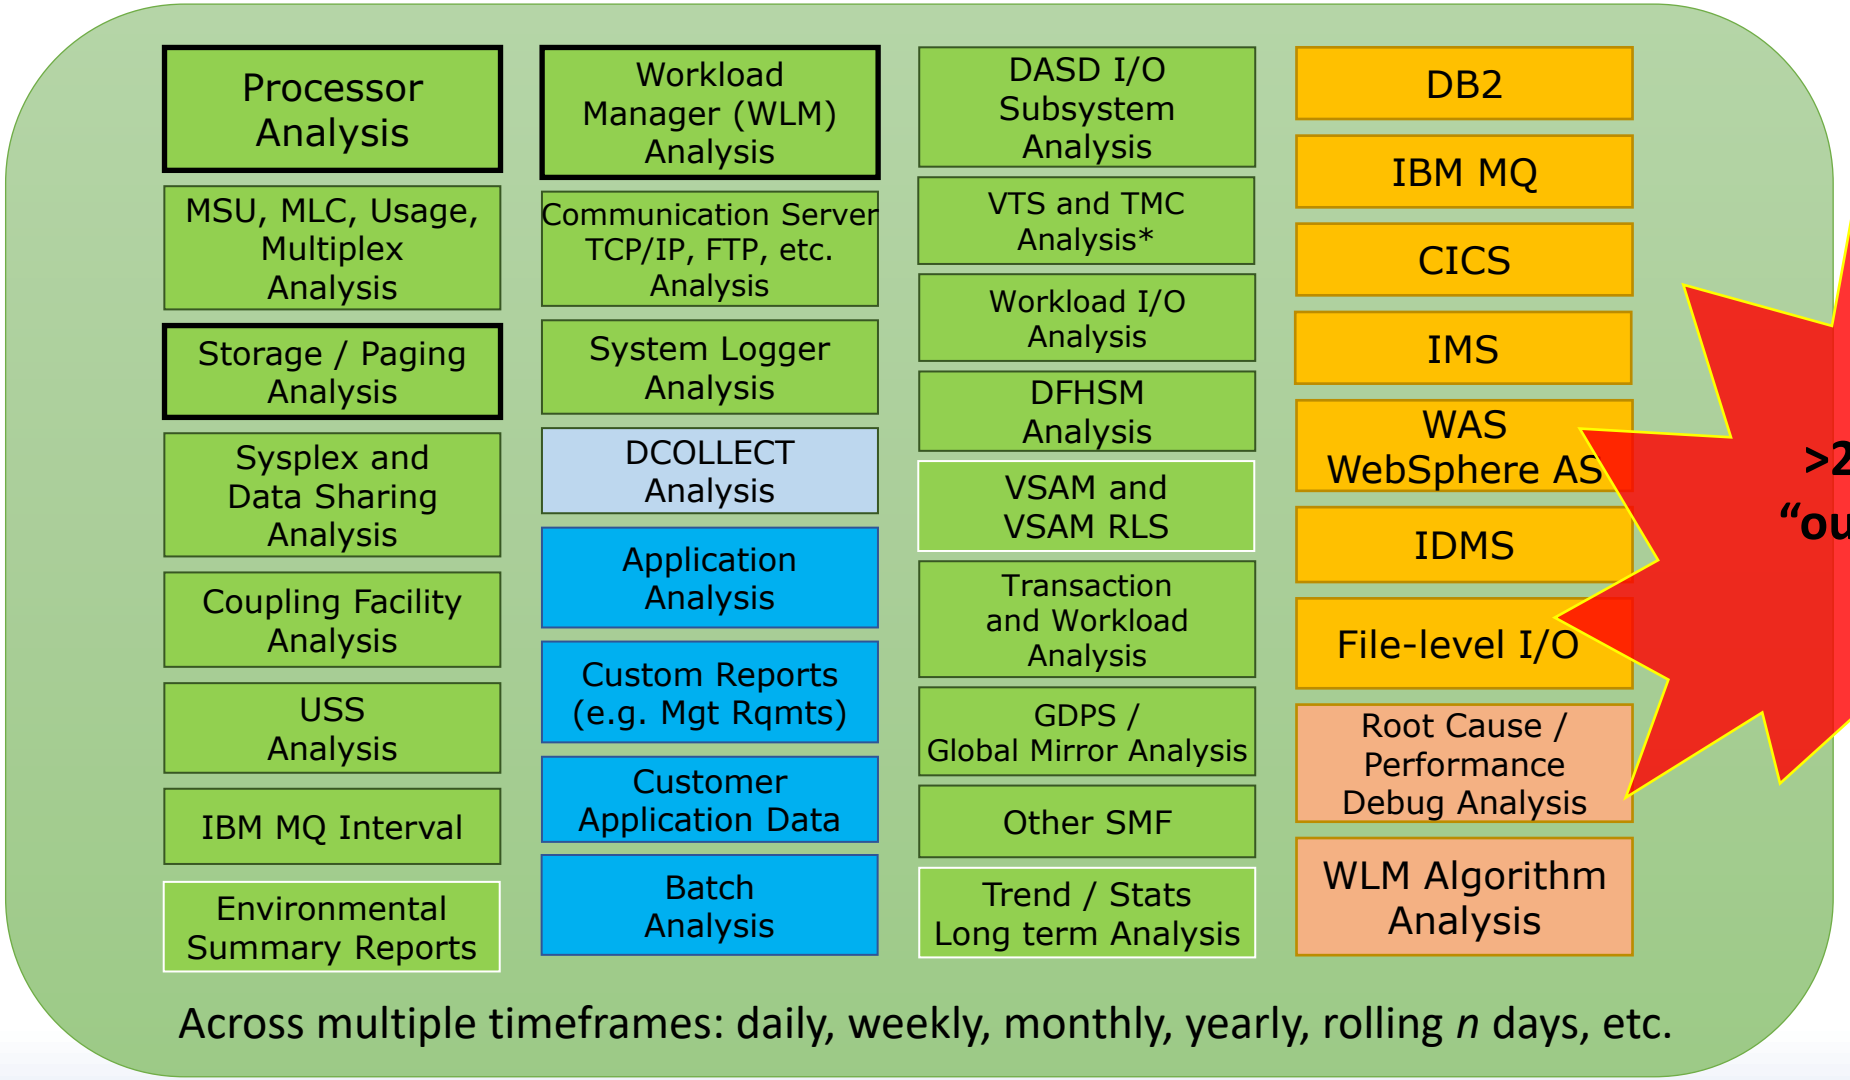

Like what you see?

- If you want a free cursory review of your environment, let us know!
  - We're always happy to process a day's worth of data and show you the results
  - See also: http://pivotor.com/cursoryReview.html
- We also have a ***free*** Pivotor offering available as well
  - 1 System, SMF 70-72 only, 7 Day rolling retention
  - That still encompasses over 100 reports!
- We offer free bi-monthly z/OS performance educational Webinars
  - Sign our guestbook to get invited
  - Or send an email to
    - Peter.Enrico@EPStrategies.com

Like what you see?

- Free z/OS Performance Educational webinars!
  - The titles for our Fall 2023-2024 webinars are as follows:
    - ✓ *LPAR Configurations to Avoid*
    - ✓ *How Different are High, Medium, and Low Pool Processors?*
    - ✓ *CPU Critical: A Modern Revisit of a Classic WLM Option*
    - ✓ *Mainframe Efficiency at High Utilizations (Bob Rogers)*
    - ✓ *I/O, I/O It's Home to Memory We (Should) Go*
    - ✓ *30th Anniversary of WLM : A Retrospective and Lessons Learned*
    - ✓ *Mainframe Efficiency at High Utilizations (presented by Bob Rogers)*
    - ✓ *Understanding and Measuring Warning Track on z/OS*
    - *30th Anniversary of Parallel Sysplex - A Retrospective and Lessons Learned*
    - *Introducing Pivotor Outlier Detection and Analysis*
    - *Batch Initiators – WLM Managed or JES Managed?*
    - *AI on Z: Exploring Common AI Terms on System Z*
    - *Analyzing 'Per CPU' Utilizations*
    - *AI on Z: Exploring new SMF Measurements*
  - Let me know if you want to be on our mailing list for these webinars
- If you want a free cursory review of your environment, let us know!
  - We're always happy to process a day's worth of data and show you the results
  - See also: http://pivotor.com/cursoryReview.html

Note:
This is a fast-paced presentation with lots of quotes

- Sit back, relax, and hear this story
- These slides will be made available, and I strongly encourage you to read through the quotes more closely than I will be doing during this presentation.

What it means to be 30

- We often celebrate landmark anniversaries such as 30

- 30th anniversaries are just moments in time during which we tend to
  - Look back and reflect from where we came
  - How we got here
  - Step back and assess the current state of affairs
  - Look into the future

- MVS 5.1 – First official release of the Parallel Sysplex and Data Sharing
  - Available 1994

Defining the Scope of this Anniversary

- The title of this presentation mentions “Parallel Sysplex and Data Sharing”
- Remember, we have 30 years of history, but this session is only 1 hour
  - Several slides are just FYI for you to look at later.
- Understand there is no ‘one product’ or ‘a single offering’ being discussed here
  - We need to remember that Parallel Sysplex and Data Sharing are not a ‘product’, but rather a collection of functionality across hardware and software that help provide a series of solutions to several issues that, at the time, needed to be addressed (and in fact, some still need to be addressed).
  - Most of these issues have still not be resolved by ‘distributed share-nothing’ systems
- So, what is really be discussed is a particular direction set by IBM, backed by a series of solutions, for the future of the mainframe platform

The Conception of Parallel Sysplex and Data Sharing

Everything has a beginning…

- If this is the 30th anniversary of Parallel Sysplex and Data Sharing, then when was the conception?

- All my research and interviews show that there was no single conceptual event
  - Instead, they were created by the growing awareness of certain hardware and software limitations that, if not addressed, would not allow for the continued growth of mainframe workloads

- First came the creation of Sysplex,
  which then evolved into Parallel Sysplex and Data Sharing

- But why?

The basic driving forces

- Capacity
  - The need for greater capacity for transactional workloads than could be provided by a single system
- Scalability
  - Non-linear scalability of the solutions of that time limited options for growth
- Performance
  - Workloads, especially transactional workloads, require optimal responsiveness
- RAS
  - Reliability - Avoidance of single points of failure
  - Availability - Continuous availability required, and not just high availability
  - Serviceability - The maintenance needs of multi-system enterprises

Why Parallel Sysplex and Data Sharing?

What finally got this into the budget for development?

First came Loosely Coupled configurations

- Prior to base Sysplex there 'loosely coupled' configurations
  - There were already products that were providing multi-system functions in loosely coupled environments
    - For example: JES2, JES3, GRS, and IRLM
  - Each would develop their own means of intersystem communication and management
  - This allowed for greater capacity and availability than just tightly coupled multiprocessor environment

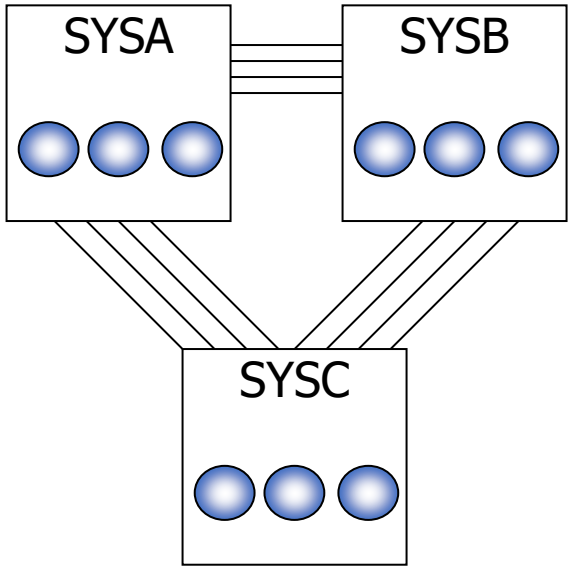

- But many RAS issues still needed to be addressed
  - Biggest was sympathy sickness
    - Systems needed capability to send/receive messages in a time manner
    - Access to I/O was reserve/release on DASD
  - Additional systems management was required sharing of data sets, printers, consoles, etc.

Old GRS Ring – Classic example of sympathy sickness

- One very large multi-system environment RAS issue was the phenomena of sympathy sickness
  - Most notably the sympathy sickness that came from designs using reserve/release on DASD and message passing over CTC where the message passing protocols required both systems to be capable of sending and receiving messages in a timely manner
  - The smallest, slowest, or failing systems resulted in degradations of the other systems
- I/O had a whole set of other issues
  - Imagine the possibility of a non-communicating system still updating a file or database!

GRS Ring Using CTCs

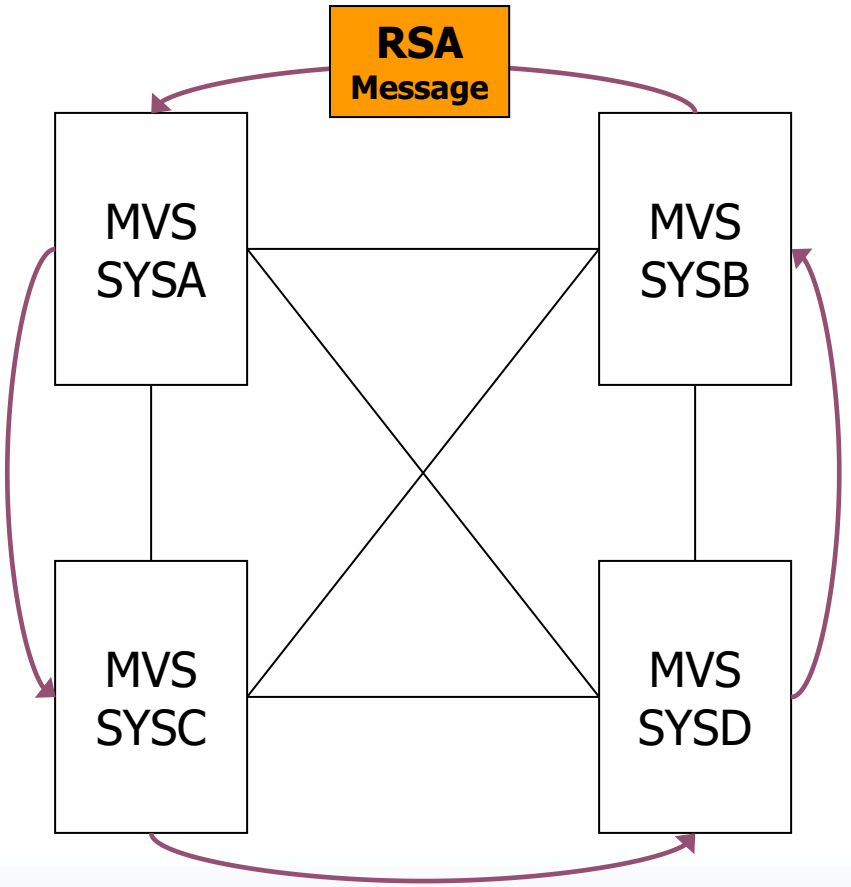

This led to base Sysplex Support
(Circa MVS/ESA 4.1)

- Base Sysplex support provided some major functionality for inter-system communication
  - Group services so products could create and maintain multiple instances of functionality on multiple systems that could be recognized by each other
  - High performance message passing capabilities (i.e. XCF)
  - Services to monitor system and subsystem health
- While this added capacity and RAS benefits over loosely coupled systems, there was still a long way to go

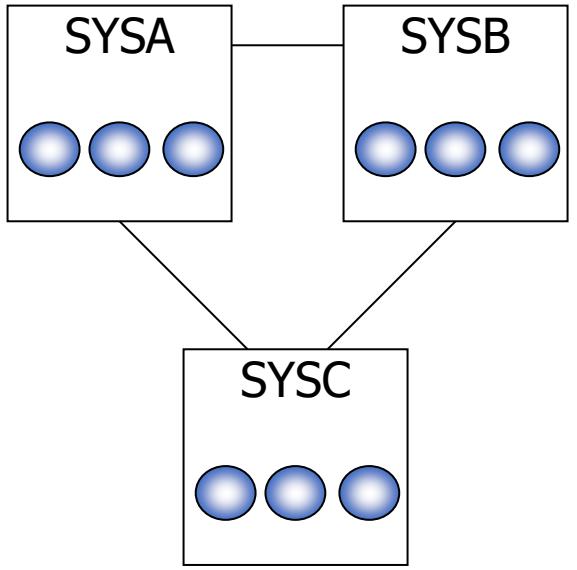

But what is the story of Parallel Sysplex and Data Sharing?

- One of my original goals while researching this subject was to talk to a number of IBMers in the trenches at the time
  - and I did
  - and I will admit I still want to talk to more to evolve this presentation

- However, I was extremely fortunate to track down one of the primary, if not 'the primary' person responsible for the realization of parallel Sysplex and data sharing

- Mike Swanson
  - Retired IBM Fellow
  - I will admit a lot of this presentation is based off the words and interview with of Mike Swanson
  - Thank you!

Why parallel Sysplex?

- At the time, data sharing on MVS already existed
  - IBM had already developed IMS 2-way data sharing in a base Sysplex environment

- Mike Swanson
  - IBM Fellow - Retired
  - *"The most compelling reason was the customer need for greater capacity for their transactional workloads than could be provided by a single system.*

    Evidence was building that having to use IMS two-way data sharing with IRLM message passing as the locking mechanism was not going to scale. Several customers had already experienced the degradation and it was easy to show that going beyond 2-way sharing was not going to work very well."

IMS 2-way data sharing

- At the time, what made this very clear was the limitations of IMS 2-way data sharing with IRLM message passing as the locking mechanism
- Two way data sharing worked, but scaling to more CPs was showing scalability issues, and adding a third system had big degradations
- *Mike Swanson (*IBM Fellow – Retired)
  - *Jimmy Strickland put a lot of effort into surfacing the IRLM locking issue and through his work IBM gained insight that later supported the design and implementation for the CF.*

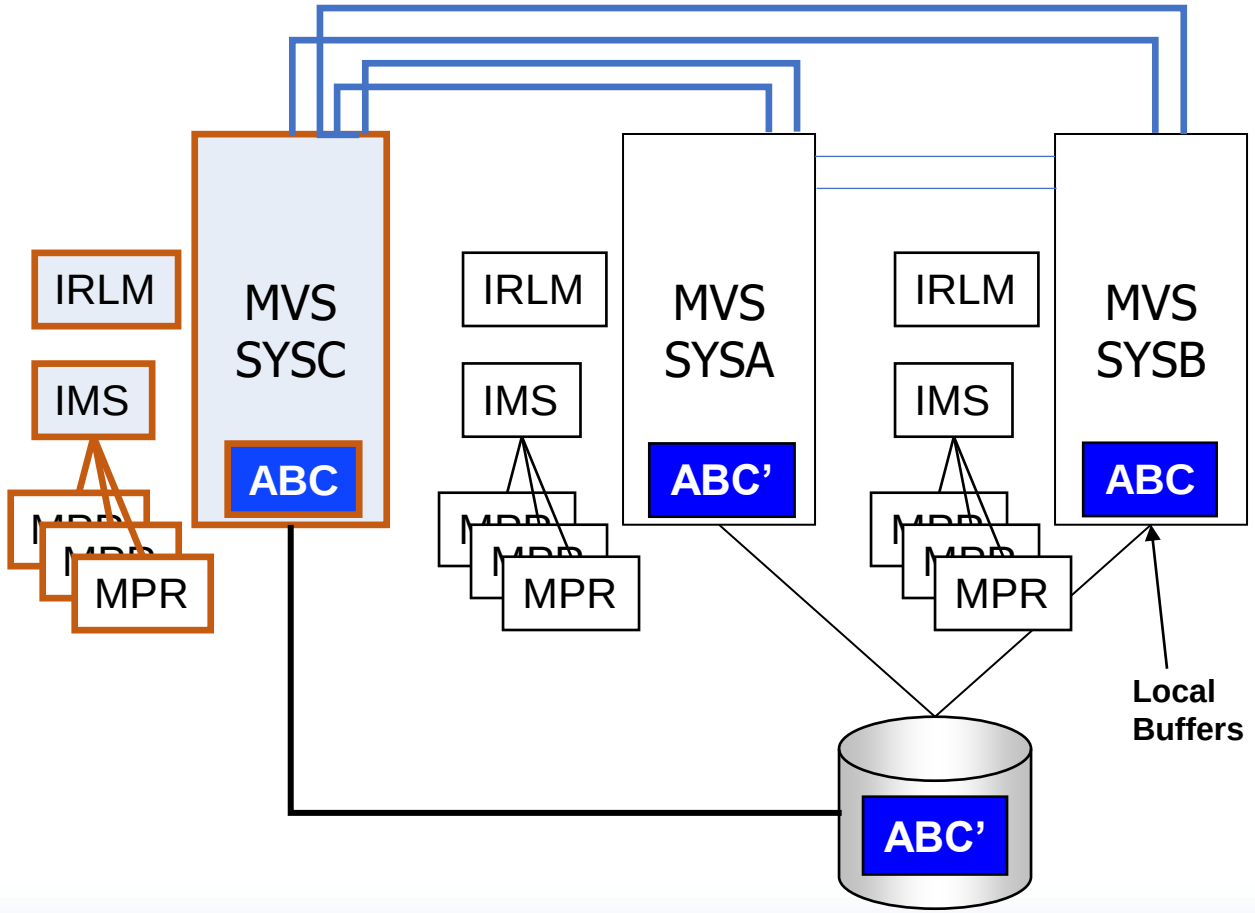

Why parallel Sysplex?

- At the time there were also growing issues regarding the single CP MIPS limitations, as well as MP effects. Amdahl was looking to create faster and faster processors, but IBM wanted more multi-processors (MPs)
  - Should growth be Vertical? Horizonal? Or Both?
- Mike Swanson (IBM Fellow – Retired)
  - *"An equally compelling reason was based on the bipolar machine design having MIPS limitations. A single machine was not projected to be sufficient for an increasing number of customer workloads, and without some mechanism to drastically improve performance multi system solutions were not going to be viable. Work efforts to create a hardware locking assist facility were not accepted by the business."*
  - Additional clarification from Mike:
    - " A call out (Jimmy) Strickland and the shaping of IBM's direction some of which was influenced by early work of Rick Baum."

Large Bipolar CPs versus smaller CMOS CPs

- At the time IBM was developing CMOS processors
  - Current bipolar machines at the time were 10-ways with about 50 MIPS per engine
    - Less per engine if you take into consideration MP effects
  - The first CMOS processors were projected to be about 12-15 MIPS per engine that came in 6-packs
    - Less per engine if you take into consideration MP effects (so 6 CMOS CPs were about 1 bipolar CP)

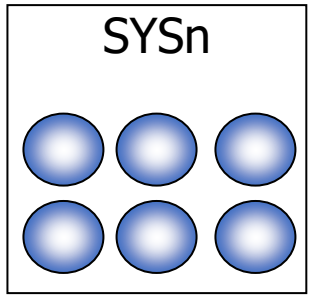

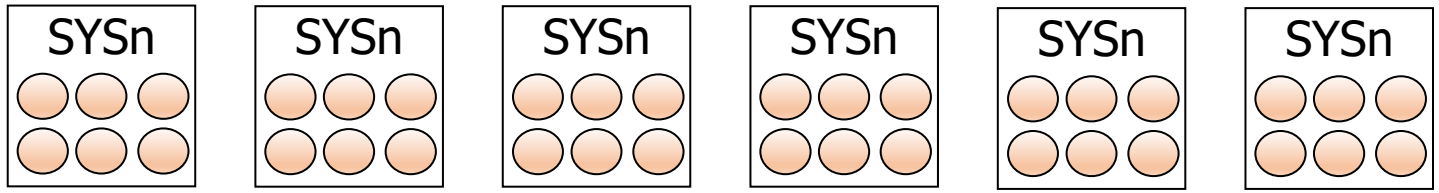

- So the challenge was to create a larger environment from smaller components
  - Could keep attempting to grow larger bipolar CPs, but that was not the future
  - A technology transition was needed until CMOS processors could get bigger
  - To accomplish this, there was a need for workloads to run on a collection of smaller systems, have a shared database and all the middleware to allow for systems management

But why a Coupling Facility?

- It was all about the need for high performance intersystem locking, share queues, and shared caches

- Mike Swanson
  - IBM Fellow - Retired
  - *"Locking in data sharing both IMS and DB2 at the time relied on contention detection with resolution of contention managed in software with different protocols in each product. So, contention detection became a critical part of enabling high performance data sharing.*

    Being able to have a common queue for work request was also a critical requirement so a means for having a shared queue or list was a critical requirement. The ability to pass messages as quickly as possible and to have either point to point or broadcast capabilities was a gate to multisystem performance and availability as well."

Why a Coupling Facility?

- Thus, the concept of a coupling facility and CF links were born

- Mike Swanson
  - IBM Fellow - Retired
  - *"Each of these factors drove the design and definition of the coupling facility and the means for sending/receiving data and notifications to/from the CF. It was also recognized that the overhead in the operating system and hardware for switching between units of work was greater than could be accepted.*

    It became a requirement for accessing the CF that no operating system task switch and related hardware cache disruption should occur. That led to the design of the CF links and the performance requirement for synchronous transfer between systems and the CF."

So then, why Data Sharing?

- There were/are 3 data sharing models common in distributed systems

Single Data Server Model

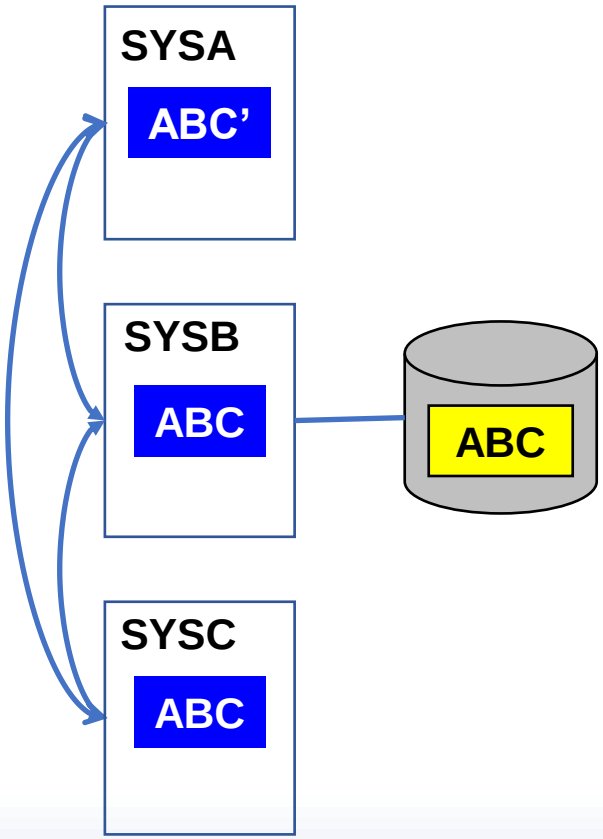

Partitioned Data Model

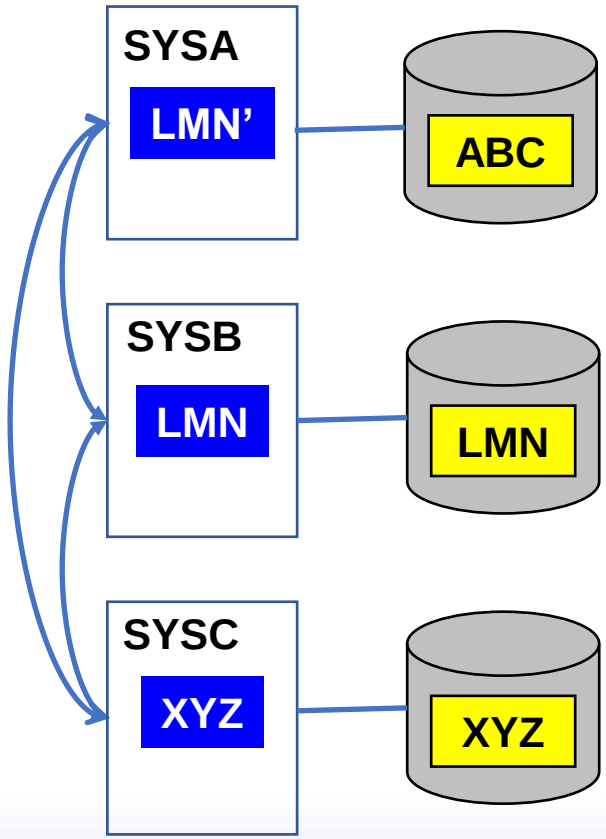

Message Model

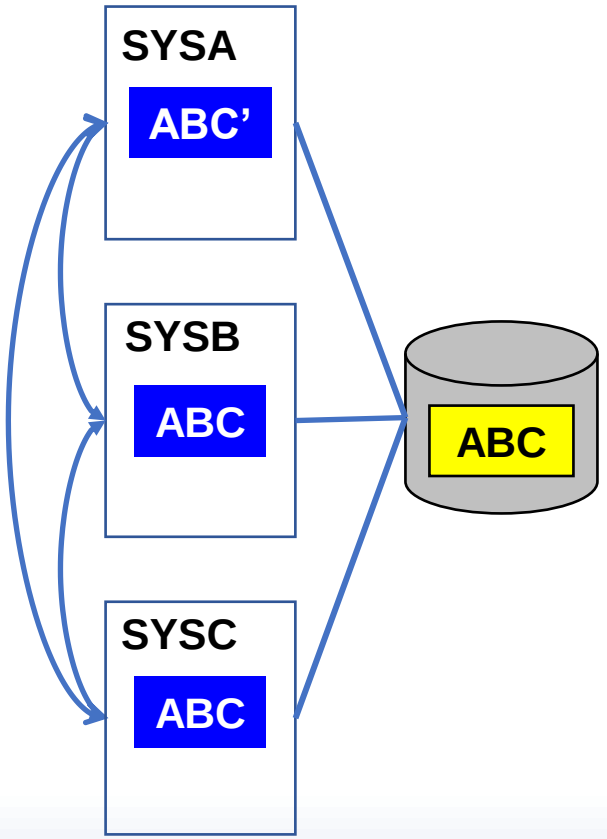

But why data sharing?

- Mike Swanson
  - IBM Fellow - Retired

  - *"A big debate at the time was on the need for data sharing at all. There were large factions within IBM and outside IBM driving for distributed solutions. In the end the idea of keeping the data shared, no requiring partitioning the data and the resulting systems management and performance problems was accepted.*

    Also, with lots of performance work having been done to back up the claim, it was accepted that keeping the data as close to the process that was using it was the best solution for performance, availability and systems management.

    That then led to the design of the CF cache structure and the key requirement for cross system invalidation notification that did not rely on correct operation of the target hardware/software."

Using the CF

- Needed:
  - Highspeed locking mechanism
  - Buffer validation
  - Data caching
  - List processing
    - Serialized or non-serialized
- Required:
  - Alleviate intersystem communication, and instead rely on the system communicating with the common storage area that would be known as the coupling facility

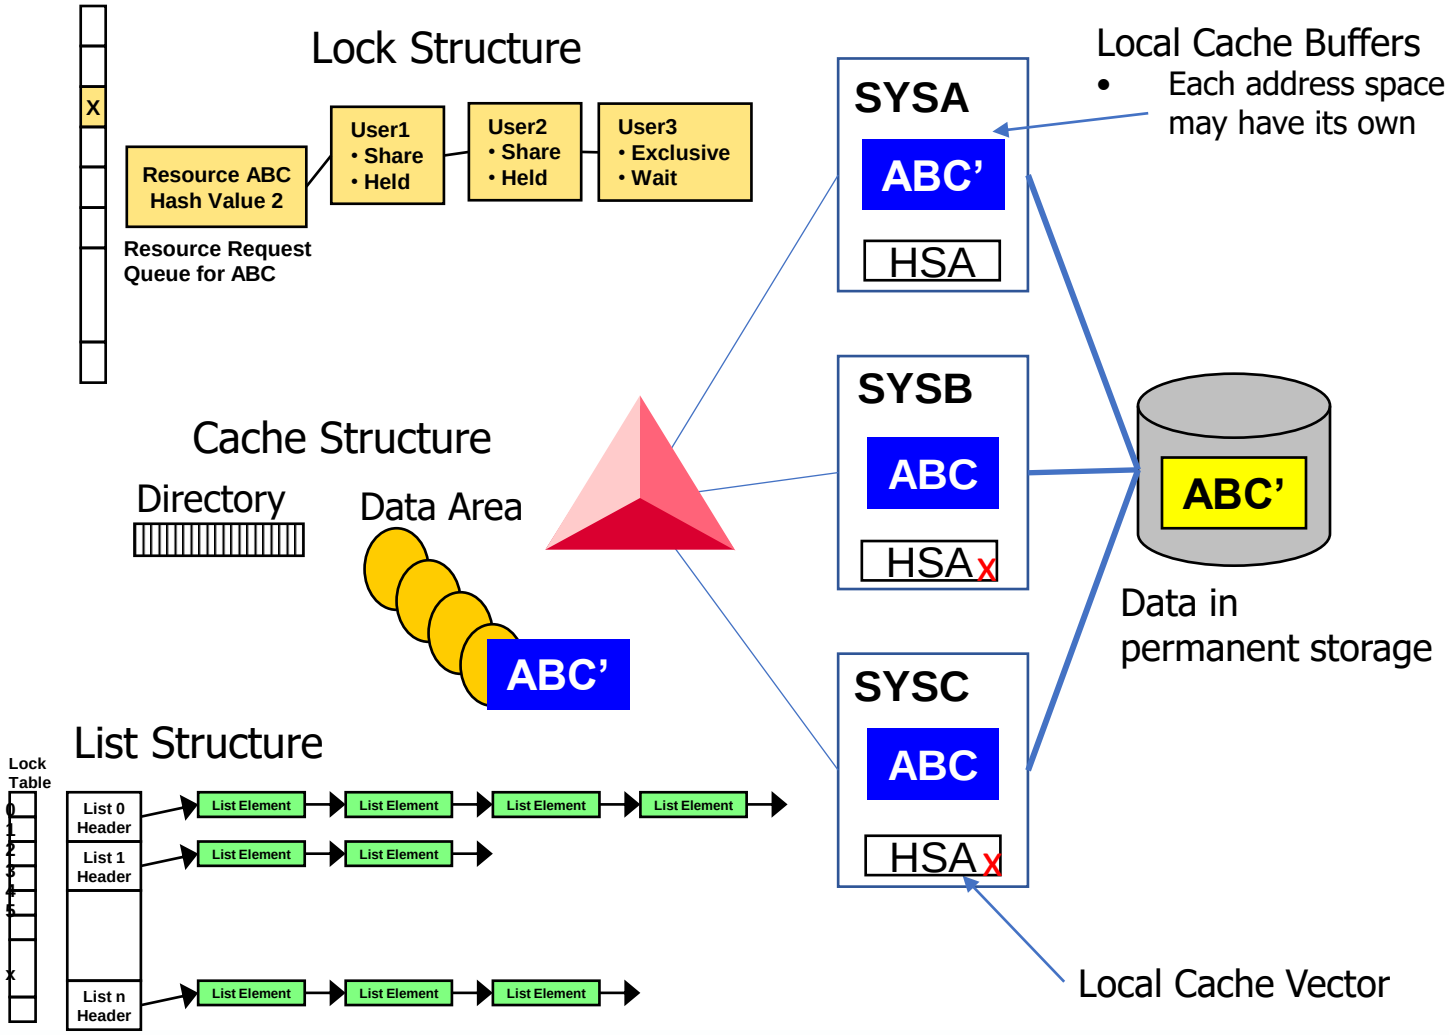

Lots of other facilities were needed

There were lots of other facilities that were needed, but two of the biggest considerations also included:

- Sysplex timer
  - Needed to keep all the system clocks in sync to enable the sequencing of events
- Fencing
  - Needed to enable every system in the Sysplex to isolate all other systems images
  - This would ensure things like I/O resources would be released by a failing system to allow for continued use by the surviving system images
  - After all… if a system were to fall out of ‘contact’ it would be terrible if it were still updating data. Data transfer had to be halted

How did the design evolve?

Evolution of the design

- Mike Swanson
  - IBM Fellow - Retired

  - *"The inability of IBM to support the lock assist facility and the introduction of Sysplex with group services, signaling and status monitoring were foundational pieces.*

    Exploitation of parallel Sysplex occurred over multiple years. The idea was to start with one data type and one transaction manager and grow the functions over time. IMS data with IRLM and the lock manager and IMS as the transaction manager were the starting point.

    During design and development of parallel Sysplex hardware, microcode, operating system and subsystem exploitation there were on going meetings with each of the labs in which technical exchanges formulated proposals for how to design and build a multisystem, data sharing, transaction balancing, highly if not continuously available and manageable platform."

Evolution of the design – Design Councils

- A work effort for something as massive as Parallel Sysplex could not be developed in a vacuum
  - IBM knew to project like Parallel Sysplex done it was to take the brain power, hard work, and sweat of a wide array of people and skills
- Multiple design councils were formed across IBM and with IBM customers
  - Software technical design council
  - Hardware technical design council
  - Customer design council

Evolution of the design – Design Councils

- Customer design council

- Mike Swanson
  - IBM Fellow - Retired
  - *"A customer design council was formed to steer the design and prioritization of functions being designed and committed in each lab. Gary Ferdinand, newly moved back to Poughkeepsie from managing DB2 at STL, was given the mission of leading the customer council. Without him and his constant communication to senior management there would have been no parallel Sysplex.*

    Jack Isenberg and I did the early leg work to go to each lab, meet with senior design skills, convey the idea of a parallel Sysplex and work to find a meaningful product meeting the IBM and customer goals."
  - Additional comment from Mike:
    - *"It is also very significant that there was one customer that worked with IBM on building the very first parallel Sysplex that supported IMS data sharing. Senior IBM management committed to support that effort and keeping that commitment gave drive, urgency and reality to the technical work."*

Evolution of the design – Design Councils

- Hardware design council

- Mike Swanson
  - IBM Fellow - Retired
    - *“Run by Jim Rymarczyk that hammered out the hardware including processor instructions, CF link attributes and CF implementation.*
    - *Some of the key players were Joe Williams and Ambrose Verdibello. Architecture was required much of which was done by Dave Elko and Audrey Helffrich.*
    - *There were many others.”*

Evolution of the design – Design Councils

- Software design council

- Mike Swanson
  - IBM Fellow - Retired
    - *"On the software side the Design council had key design skills from all the areas in which change was needed for data sharing, workload distribution.....*
    - *Some of the areas and names were: Vern Watts & Dick Hannan for IBM, Inderpal Narang and Chandrasekara Mohan for DB2, Tim Holloway and Ken Davies for CICS, Jim Hall and Jim Arnold for VTAM, George Wang did performance work, and a slew of Poughkeepsie people including Jeff Nick, Jeff Frey, Jeff Aman, Carl Clark, Alan Little, Dave Surman.*
    - *There were a lot of players."*

What was developed and delivered in the first years?

Any Good Stories?

- Mike Swanson
  - IBM Fellow - Retired
  - *"I'm not so sure about this question Peter. There are lots of stories but I'm not sure most of them should be told!*
  - *There were way too many late-night beer drinking sessions with paper napkins and a pencil where Jeff (Frey), Jeff (Nick) and I tried to address concerns of various exploiting products with proposals to changes in the CF, microcode and OS.*
  - *There was the break between Christmas and New Years where we came to the realization the existing design for locking in the OS support code for the CF was not going to have the parallelism needed. Jeff (Nick) and I redesigned the locking hierarchy to a much more granular level, found all the code where locks were obtained/released and rewrote the design specs with the new hierarchy."*

Parallel Sysplex Announced in 1994

- A lot of stuff that needs to be addressed to do it right

- Some of the key announcement letters included the following:
  - IBM MVS/ESA SP Version 5 Release 1 and OpenEdition Enhancements
    - April 6, 1994
    - Announcement Number 294-152
  - S/390 Parallel Sysplex Overview
    - April 6, 1994
    - Announcement Number 194-080
  - IBM S/390 Coupling Facility 9674 Model C01
    - April 6, 1994
    - Announcement Number 194-082

Parallel Sysplex Structure

- A parallel Sysplex is made up of a variety of key software, hardware, microcode components
- Functionality:
  - Data sharing
  - Systems management
  - Distributed workloads

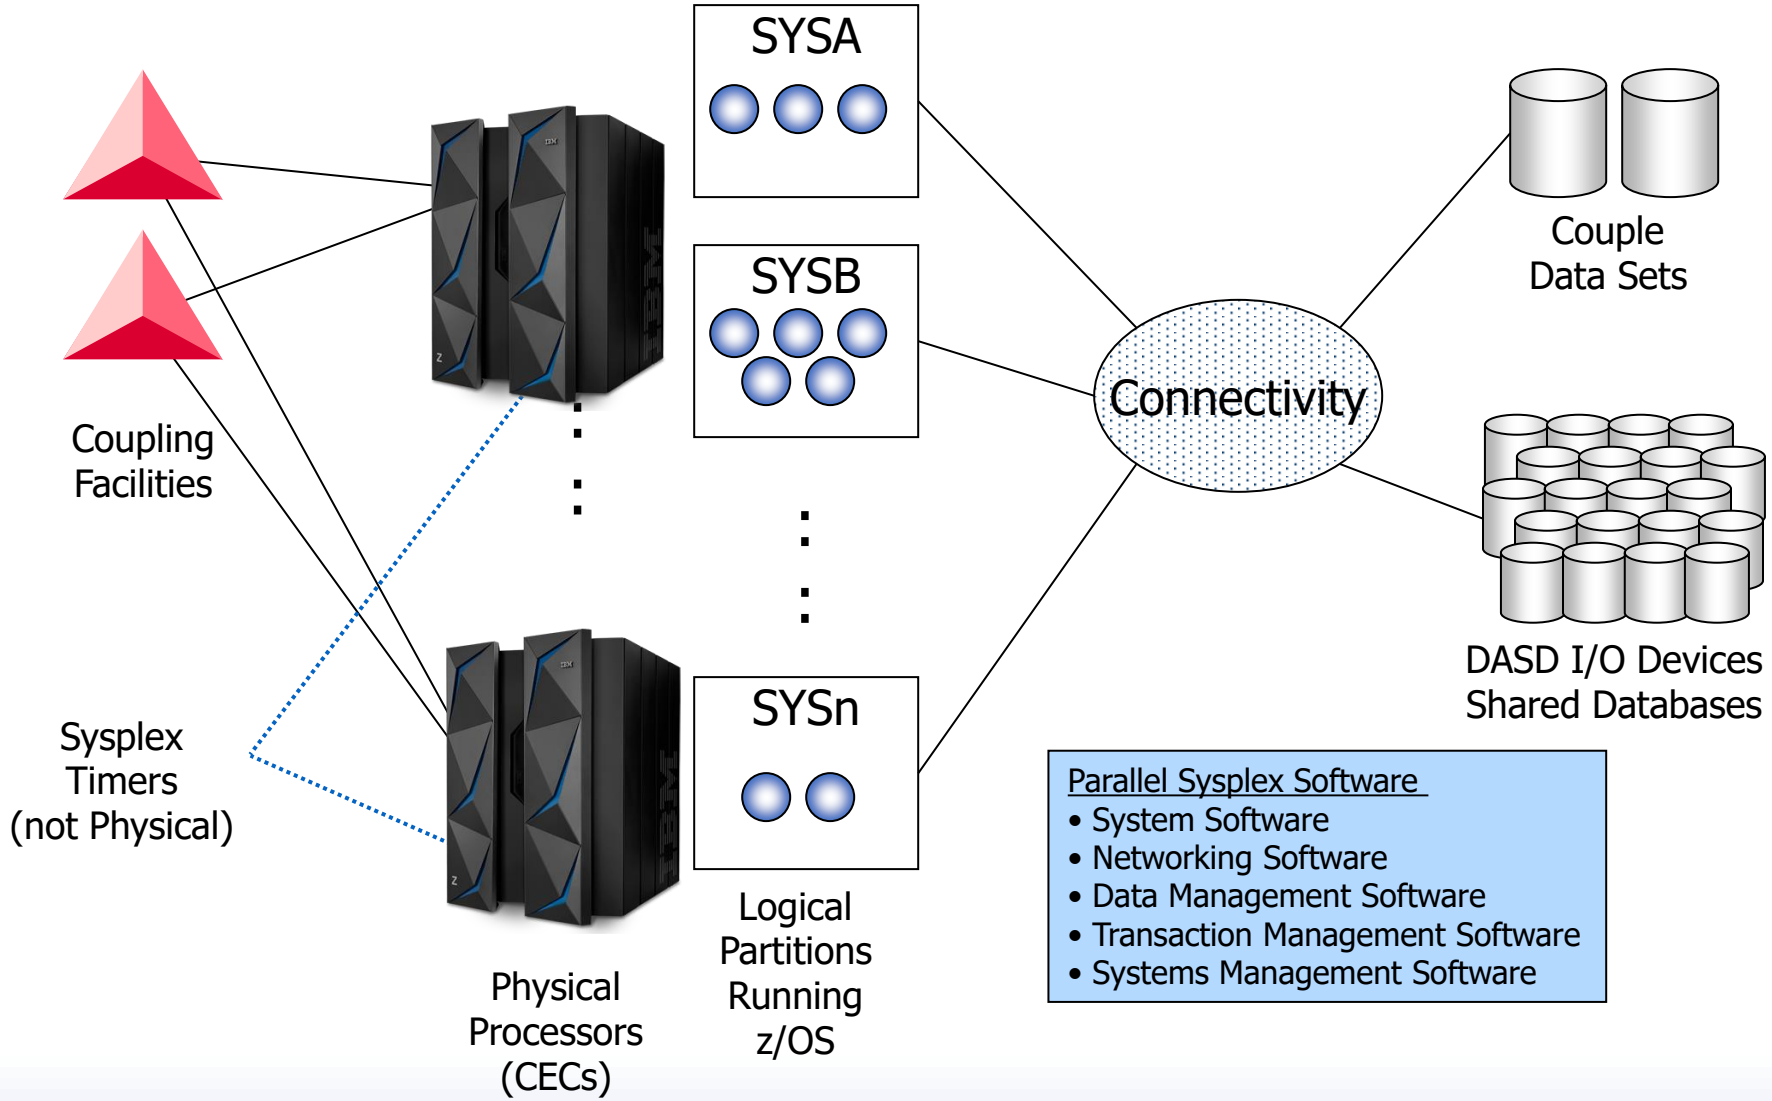

A qualitatively different solution

- Parallel Sysplex was a different way of doing things
- For example, CICS / DB2 in a Sysplex environment

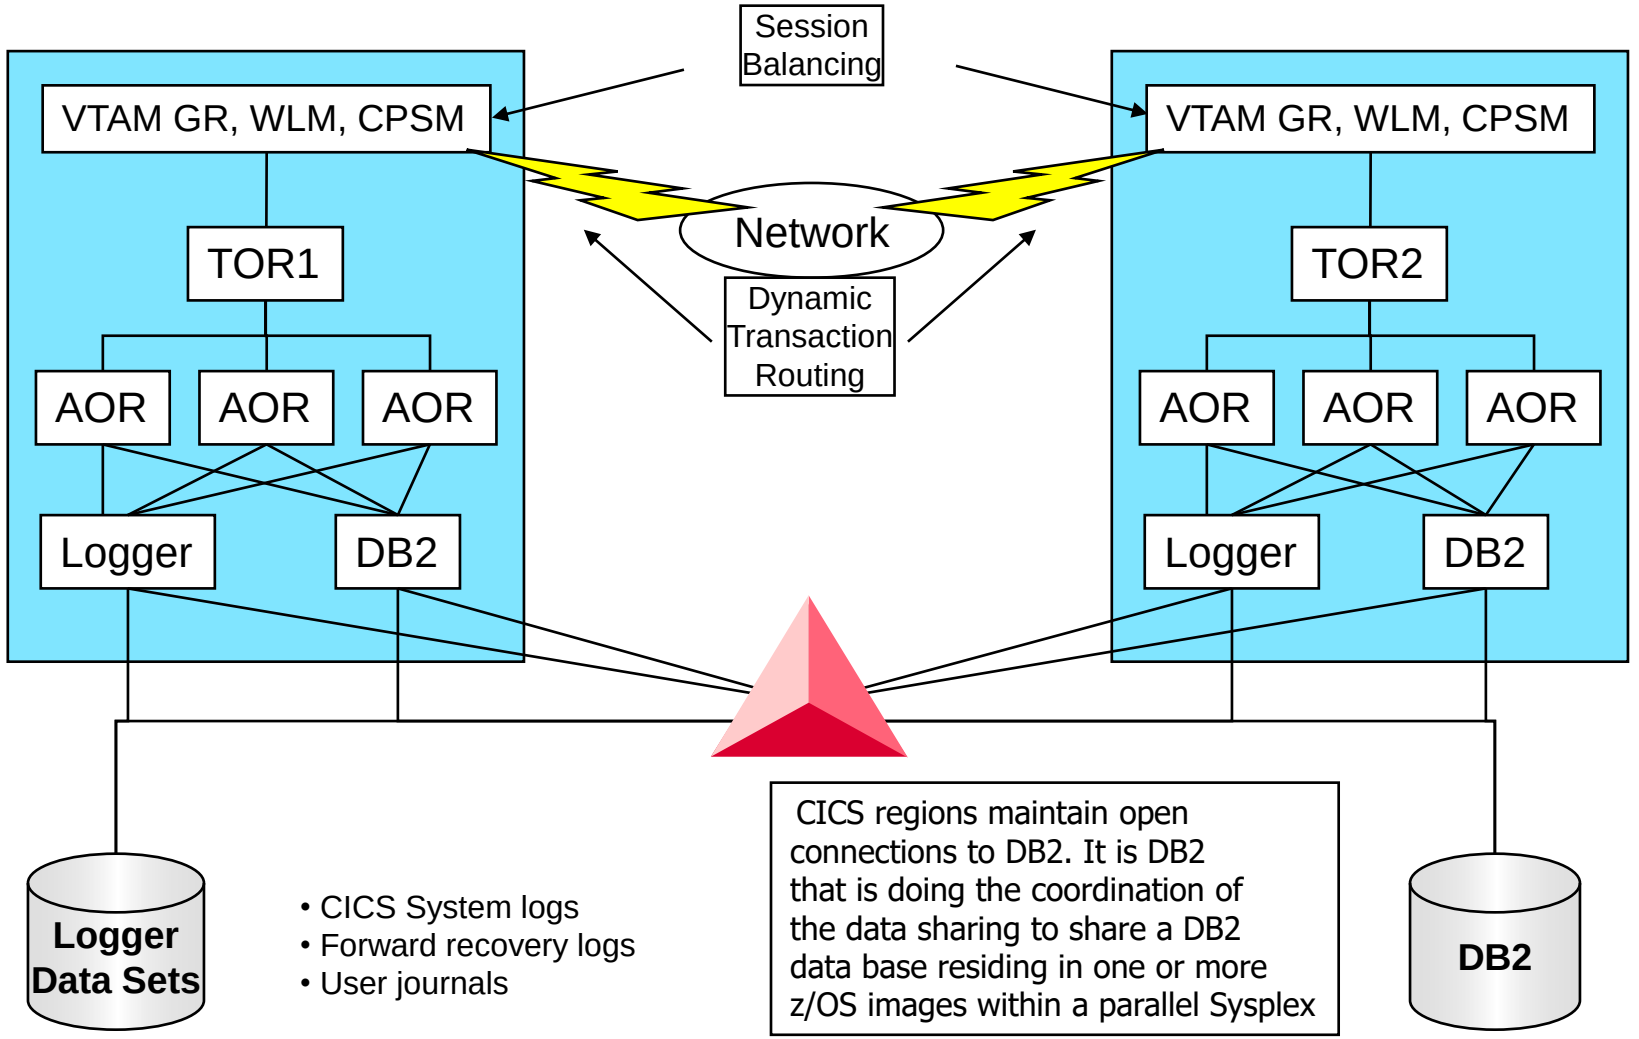

Sysplex Hardware – over a period of years

- Processors
  - zArchitecture processors, and the eventual advent of ICF processors
- Coupling Facility
  - The coupling facility enables high performance multisystem data sharing.
  - Initially a stand-alone coupling facility (9674)
  - Important point: on zArchitecture technology, so provided for huge flexibility which led to ICF engines and ICF LPARs
- Coupling Facility Links
  - CF links provide high speed connectivity between the coupling facility and the exploiting systems
  - Initially sender and receiver pairs, then eventually bidirectional peer mode
- Server Time Protocol for a Sysplex Timer
  - The ability to synchronize the time-of-day (TOD) clocks in multiple CPCs in a Sysplex
  - Initially a hardware timer
  - Utilizes the Server Time Protocol (STP) for synchronization
- Control Units, I/O Devices, Channels, Directors, etc.
  - Storage controllers in a Sysplex provide the increased connectivity necessary among a greater number of systems.

Sysplex Software – over a period of years

- System Software
  - Base system software that is enhanced to support a Sysplex includes the z/OS operating system
- Networking Software
  - Such as VTAM and TCP/IP that supports the attachment of a Sysplex to a network
- Data Management Software
  - The data managers that support data sharing in a Sysplex include IMS DB, DB2, and VSAM
- Transaction Management Software
  - The transaction managers that support a Sysplex include CICS, IMS, WAS, WMQ, and more
- Systems Management Software
  - A number of software products are enhanced to run in a Sysplex and exploit its capabilities
  - The products manage accounting, workload, operations, performance, security, and configuration, and they make a Sysplex easier to manage by providing a single point of control

What has been some of the great successes?

The greatest successes were satisfying the driving forces

- Capacity
  - Growth could be both horizontal, vertical, or both
- Scalability
  - Near linear scalability
- Performance
  - Greater responsiveness
- RAS
  - Reliability - Avoidance of single points of failure
  - Availability - Continuous availability required, and not just high availability
  - Serviceability - The maintenance needs of multi-system enterprises
- Also:
  - System fencing
  - Greater workload management, distributed transactions, workload balancing

The Great Successes

- I asked a wide assortment of people what they thought were the biggest successes of parallel Sysplex and data sharing
  - Typical answers were
    - Shared everything
      - Mike Swanson: *"Datasharing and distributed workload management were the big win"*
      - Bob Rogers: *"Success… efficient share everything. That is not easy to do."*

- Today, Parallel Sysplex still stands unrivaled in terms of the key goals it was targeted to achieve
  - Mike Swanson: "Keeping the parts of the parallel Sysplex going as processor technology, link technology, … keep changing and still meeting the performance and RAS characteristics is no small feat."

The Great Successes – Scalability!

- Once get over initial performance cost of going into parallel Sysplex, growth becomes more efficient of growing the number of system images rather than making the images larger by adding more engines

- Remember, CMOS engines are not getting much faster so even today largest customers can still grow horizontally

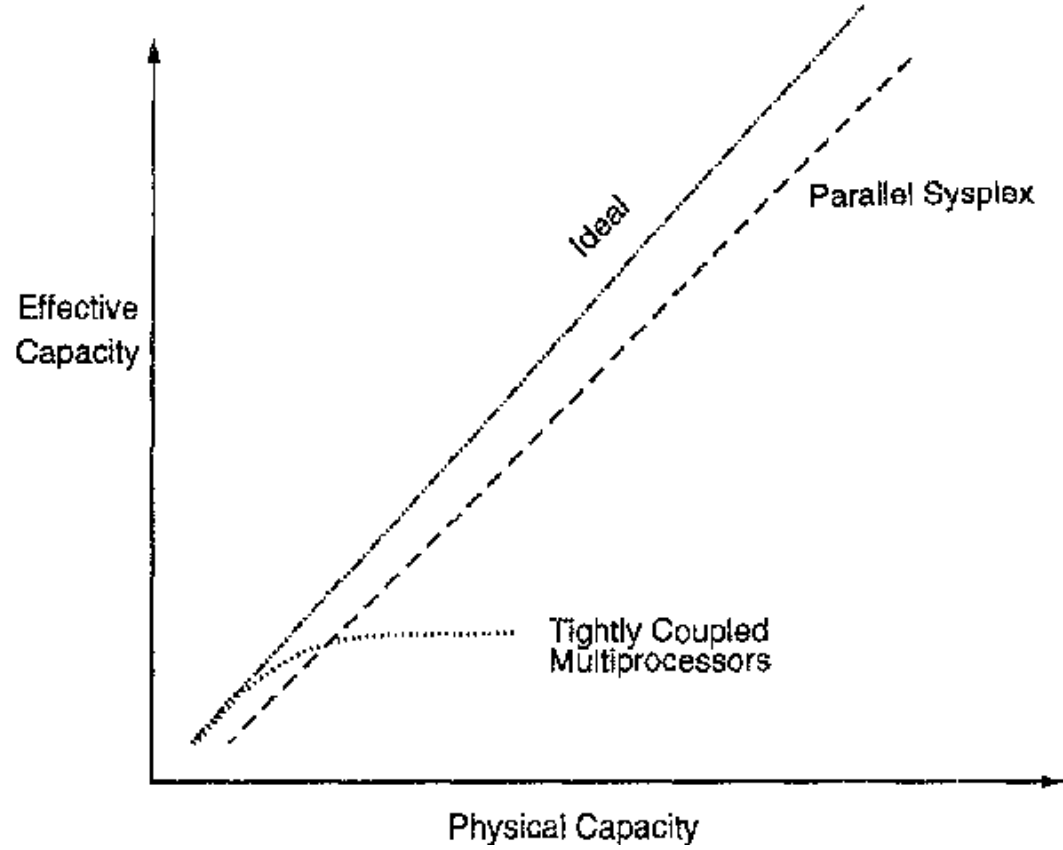

Parallel Sysplex and True Consistency

- Mythili Venkatakrishnan
  - Distinguished Engineer, IBM

  - *"It's also clear that the distributed horizontal scale solutions that are prevalent today struggle with the fact that what can be achieved is usually "eventual consistency"* -- *and this creates a lot of challenges for mission critical applications and systems of record.*

  - *Parallel Sysplex and the data sharing capabilities built on top enable "true consistency"* *and even as enterprise clients evaluate hybrid cloud strategies, what they are fast realizing is that many systems of record require the shared data models that are consistent. "*

Some Genius Moves by IBM

- Not a forced migration, but an environmental option
  - It was never the intention to force migration over to a CF or parallel Sysplex
  - IBM never wanted to require, and so even today a CF or parallel Sysplex are not required

- Coupling Facility was based on s/390 architecture
  - When first introduced the CF was its own machine…
    - Originally it was thought they needed a special 'appliance' for locking, another for caching, etc.
    - It was wisely decided against this approach
  - Because it was based on s/390 architecture and not a special machine, it led the way to CF LPARs, ICF engines, ICF LPARs, etc.

What were some of the disappointments?

Parallel Sysplex Disappointments

- I asked a wide assortment of people what they thought were the biggest disappointments of parallel Sysplex and data sharing

- Not many disappointment were listed
  - The general though was that the benefits far outweighed anything else

- However, typical answers were
  - Overhead
  - Complexity
  - Frustration that pricing forced *unnatural acts*

Disappointment? MVS CPU Cost (MSUs)

- According to IBM, typical observed performance cost for Parallel Sysplex is:
  - 3% - Cost of multisystem management and resource sharing
  - <10% - Cost of data sharing
  - 0.5% - Incremental cost of adding a new system image to the Sysplex

- Reality is much better
  - According to me (Peter Enrico), today the typical performance cost for probably more than 80% to 90% of z/OS shops is between 1% and 6%
    - The biggest costs is due to the costs of lock structures and cache structures for data sharing
    - Most customers do not have high degrees high degrees of data sharing
    - Except for features like logger, most systems management structures have low cost (and most logger activity is off the CF anyway)

Performance Cost to Parallel Sysplex and Data Sharing

- While performance was initially very good, there was still the estimated an initial 10% overhead with an additional .05% each additional system
- But the overhead is directly related to the degree of data sharing, and most customers are not high data sharing environments
- The chart on this slide shows a recent estimated host effect table from IBM
  - Note that IBM still estimates the lowest overheads as being 10% to 11%
  - However, most of today's environment have much less than 9 CF ops/MIPS, and most customers do not have intense data sharing
- Typical customers have host effects of between 1% to 6%

Host effect with primary application in data sharing.
Chart is based on 9 CF ops/Mi – may be scaled linearly for other rates

Host / CF	z13s	z13	z14 ZR1	z14	z15 T02	Z15 T01	z16
z13s CL5	18	18	18	18	18		
z13s 1x IFB	18	18		18			
z13s 12x IFB	17	17		17			
z13s 12x IFB3	12	12		12			
z13s CS5	11	11	11	11	12		
z13 CL5	18	18	18	18	18		
z13 1x IFB	18	18		18			
z13 12x IFB	16	17		17			
z13 12x IFB3	12	12		12			
z13 CS5	11	11	11	11	12		
z14 ZR1 CL5	18	18	18	18	18		18
z14 ZR1 CS5	11	11	11	12	12		12
z14 1x IFB	18	18		21			
z14 12x IFB	16	16		17			
z14 12x IFB3	11	11		12			
z14 CL5	18	18	18	18	18		18
z14 CS5	10	11	11	12	12		12
z15 T02 CL5	18	18	18	18	18		18
z15 T02 CS5	10	11	11	11	11		12
z15 T01 CL5	18	18	18	18	18	18	18
z15 T01 CS5	10	10	11	11	11	11	12
z16 CL5			18	18	18	18	18
z16 CS5			10	10	11	11	10

Source: Coupling Technology Overview and Planning, Gary King, IBM

Disappointment? Complexity

- Acceptance to parallel Sysplex and data sharing was somewhat slow because it was very complex.
  - Designed by rocket scientists and had huge complexity. But… making a parallel Sysplex (i.e. 2) was the most difficult. Going to 3, or more, was much easier
- Was perceived to have great complexity
  - Acceptance was based on what you already knew.
  - Larger shops that had large staffs and expertise found it much easier to exploit parallel Sysplex
  - Many workloads were hard to parallelize
    - CICS transactions that had affinities, databases that could not do multisystem locking, etc.
- For a long while, and even today, some customers find it easier to grow vertically rather than horizontally
  - Depends on availability RAS by customer

Disappointment? PSLC and *Unnatural Acts*

- Migrating to parallel Sysplex and CMOS meant more systems, more overhead, and more complexity
- In 1994 IBM also introduced pricing incentive PSLC to help customers to migrate to parallel Sysplex
- The pricing led to people to do things that were insane.
  - Led to migration of parallel Sysplex, not for efficiency, but rather for price efficiencies.
  - Not for any benefit other than pricing.
  - Not that this is not a good thing, but foolish to incent people to do the wrong thing.

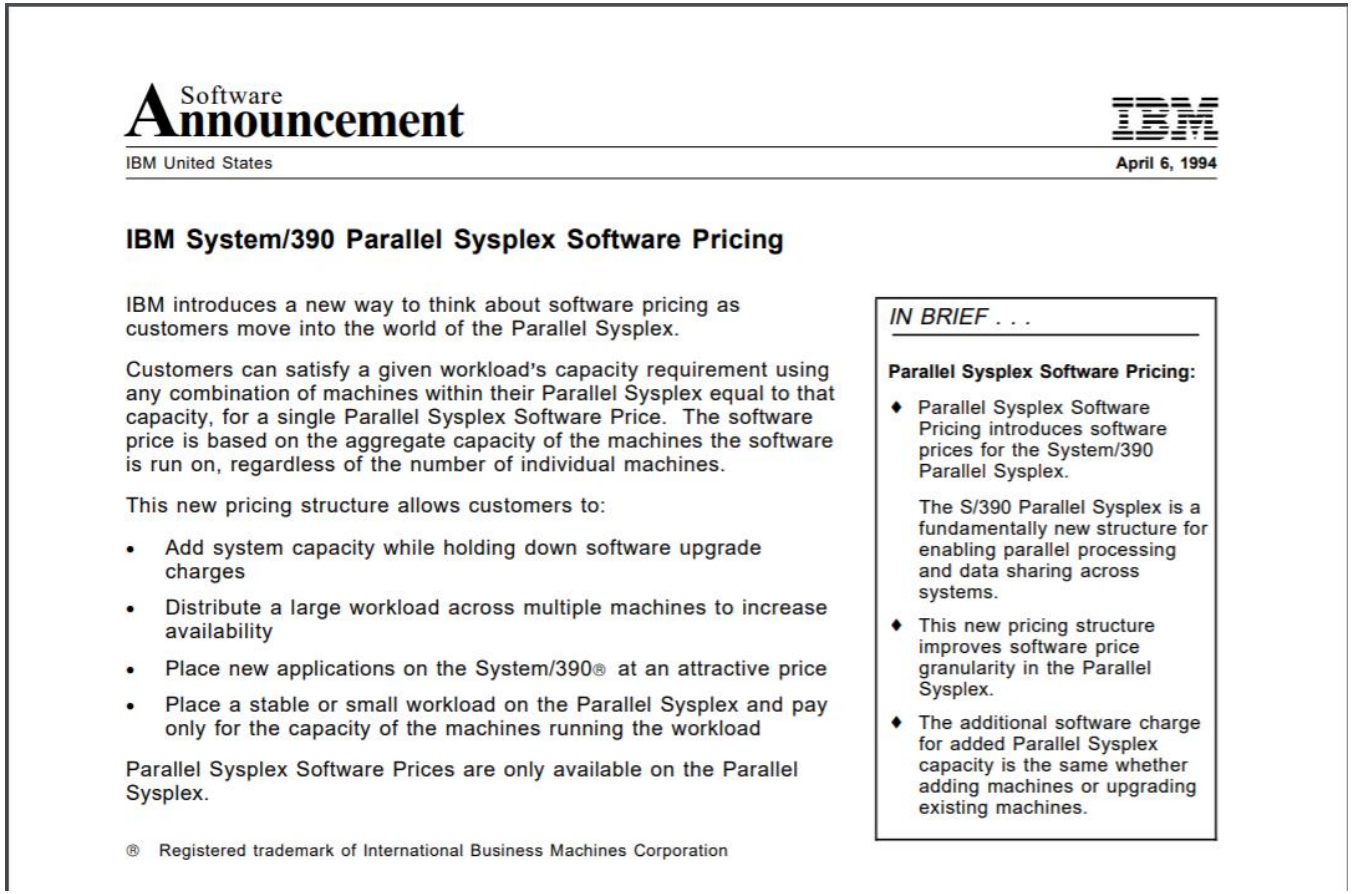

Software Announcement

IBM United States

April 6, 1994

IBM System/390 Parallel Sysplex Software Pricing

IBM introduces a new way to think about software pricing as customers move into the world of the Parallel Sysplex.

Customers can satisfy a given workload's capacity requirement using any combination of machines within their Parallel Sysplex equal to that capacity, for a single Parallel Sysplex Software Price. The software price is based on the aggregate capacity of the machines the software is run on, regardless of the number of individual machines.

This new pricing structure allows customers to:

- Add system capacity while holding down software upgrade charges
- Distribute a large workload across multiple machines to increase availability
- Place new applications on the System/390® at an attractive price
- Place a stable or small workload on the Parallel Sysplex and pay only for the capacity of the machines running the workload

Parallel Sysplex Software Prices are only available on the Parallel Sysplex.

® Registered trademark of International Business Machines Corporation

IN BRIEF . . .

Parallel Sysplex Software Pricing:

- Parallel Sysplex Software Pricing introduces software prices for the System/390 Parallel Sysplex.

  The S/390 Parallel Sysplex is a fundamentally new structure for enabling parallel processing and data sharing across systems.
- This new pricing structure improves software price granularity in the Parallel Sysplex.
- The additional software charge for added Parallel Sysplex capacity is the same whether adding machines or upgrading existing machines.

https://www-01.ibm.com/common/ssi/rep_ca/4/897/ENUS294-154/ENUS294-154.PDF

Example of PSLC Pricing Curve

- The initial buy-in to parallel Sysplex is the most expensive, but as more capacity is added the costs fall off
- But this was on a Sysplex basis
- Some customers would add test and development systems to their production environments just to avoid the initial costs

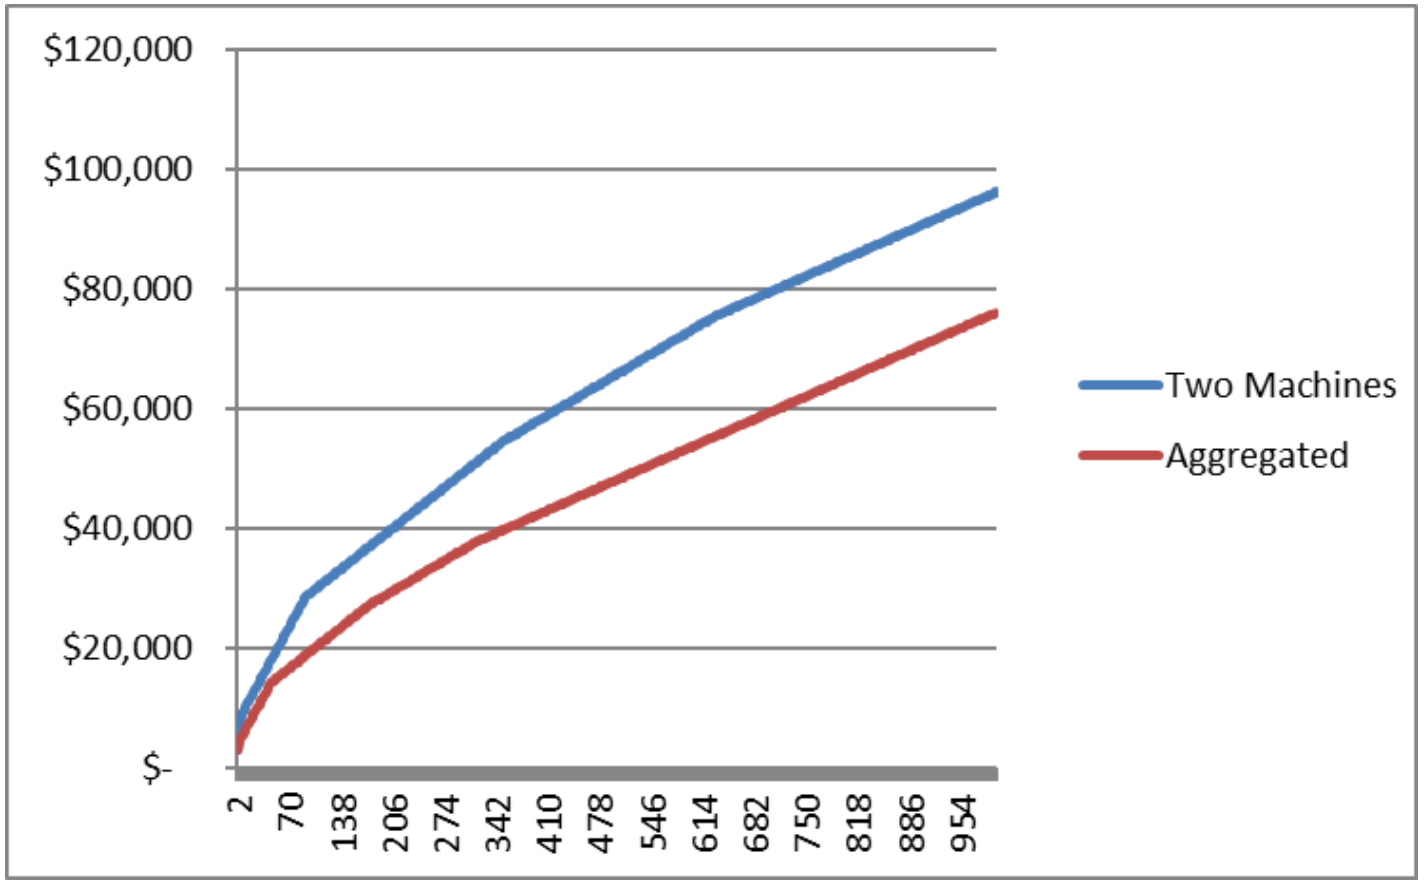

Disappointment? PSLC/E and *Unnatural Acts*

- In late 1994 IBM also introduced pricing incentive PSLC/E
- Required customer's machines be operating in an "actively coupled" environment to qualify
- This led to the creation of Sham-plex to get some pricing benefits
  - Example: Just for JES2 checkpoint, or combine unlike systems/workloads (i.e. production and test) in same Sysplex
  - Not that this is not a good thing, but foolish to incent people to do the wrong thing.

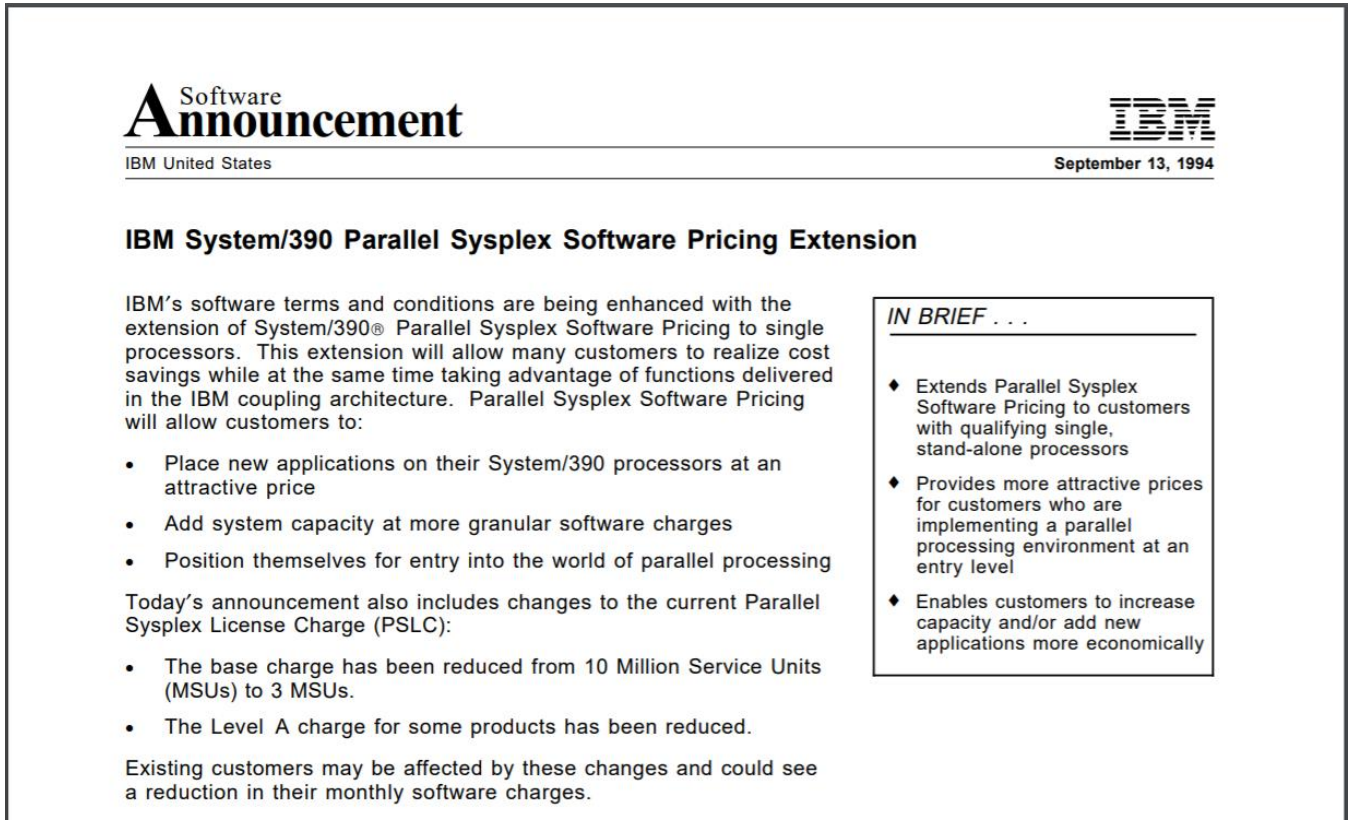

Software Announcement

IBM United States — September 13, 1994

IBM System/390 Parallel Sysplex Software Pricing Extension

IBM's software terms and conditions are being enhanced with the extension of System/390® Parallel Sysplex Software Pricing to single processors. This extension will allow many customers to realize cost savings while at the same time taking advantage of functions delivered in the IBM coupling architecture. Parallel Sysplex Software Pricing will allow customers to:

- Place new applications on their System/390 processors at an attractive price
- Add system capacity at more granular software charges
- Position themselves for entry into the world of parallel processing

Today's announcement also includes changes to the current Parallel Sysplex License Charge (PSLC):

- The base charge has been reduced from 10 Million Service Units (MSUs) to 3 MSUs.
- The Level A charge for some products has been reduced.

Existing customers may be affected by these changes and could see a reduction in their monthly software charges.

IN BRIEF . . .

- Extends Parallel Sysplex Software Pricing to customers with qualifying single, stand-alone processors
- Provides more attractive prices for customers who are implementing a parallel processing environment at an entry level
- Enables customers to increase capacity and/or add new applications more economically

https://www-01.ibm.com/common/ssi/rep_ca/0/897/ENUS294-540/ENUS294-540.PDF

The invention of the Sham-plex

- Customers would have bare minimum Sysplexes to take advantage of pricing

- Forced IBM to make clarifications
  - *"The configuration and operating modes described in this exhibit must be the normal mode of operations for this environment. The OS/390 and MVS Images participating in the above Sysplex functions must account for at least 50% of the total OS/390 and MVS workload on each machine.*

  - *A processor can only be in one Parallel Sysplex for pricing purposes. If the processor is partitioned, and the partitions are in different qualifying Parallel Sysplexes, the customer may select which Parallel Sysplex the processor will be included in for billing."*

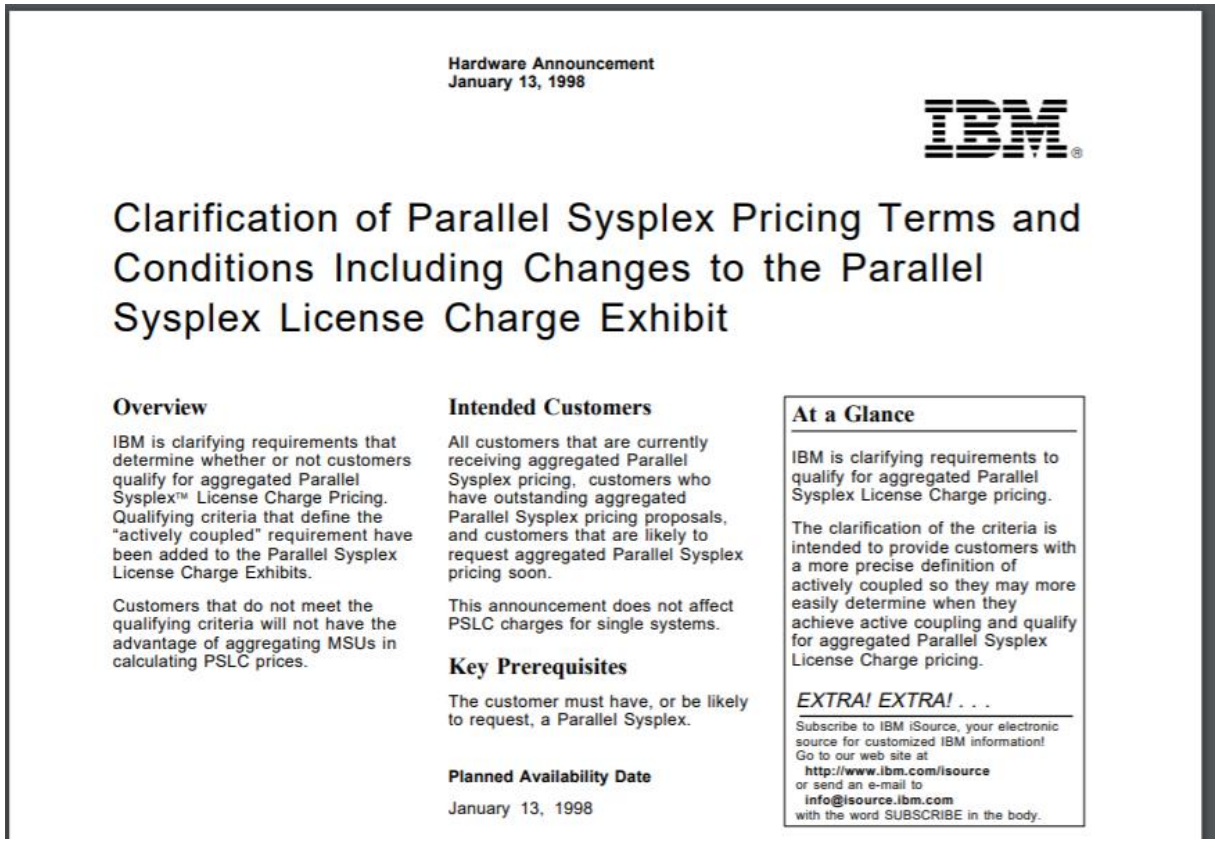

Hardware Announcement
January 13, 1998

IBM

Clarification of Parallel Sysplex Pricing Terms and Conditions Including Changes to the Parallel Sysplex License Charge Exhibit

Overview

IBM is clarifying requirements that determine whether or not customers qualify for aggregated Parallel Sysplex™ License Charge Pricing. Qualifying criteria that define the "actively coupled" requirement have been added to the Parallel Sysplex License Charge Exhibits.

Customers that do not meet the qualifying criteria will not have the advantage of aggregating MSUs in calculating PSLC prices.

Intended Customers

All customers that are currently receiving aggregated Parallel Sysplex pricing, customers who have outstanding aggregated Parallel Sysplex pricing proposals, and customers that are likely to request aggregated Parallel Sysplex pricing soon.

This announcement does not affect PSLC charges for single systems.

Key Prerequisites

The customer must have, or be likely to request, a Parallel Sysplex.

Planned Availability Date

January 13, 1998

At a Glance

IBM is clarifying requirements to qualify for aggregated Parallel Sysplex License Charge pricing.

The clarification of the criteria is intended to provide customers with a more precise definition of actively coupled so they may more easily determine when they achieve active coupling and qualify for aggregated Parallel Sysplex License Charge pricing.

EXTRA! EXTRA! . . .

Subscribe to IBM iSource, your electronic source for customized IBM information! Go to our web site at **http://www.ibm.com/isource** or send an e-mail to **info@isource.ibm.com** with the word SUBSCRIBE in the body.

https://www-01.ibm.com/common/ssi/rep_ca/1/897/ENUS198-001/ENUS198-001.PDF

What was the influence of Parallel Sysplex on the z Platform?

A far reaching and all encompassing

- As mentioned… When parallel Sysplex was designed and developed, it touched nearly every major hardware and software component of the mainframe platform over a period of years!
  - Hardware
    - CMOS, Coupling Facility, Coupling Links, Sysplex Time, etc.
  - System software
    - Such as MVS, JES2, JE3, DFSMS, and so much more
  - Transaction managers
    - Such as CICS, IMS, and (today) WAS
  - Database managers
    - Such as DB2, VSAM, and IMS DB
  - System workload managers
    - Such as z/OS Workload Manager (WLM) and CICSPlex SM, transaction routers and distributors
  - Networking software for balancing
    - Such as VTAM, TCP/IP
  - Operations
    - Such as consoles, systems automation, RACF
  - Many vendor products
  - More…

The influence on the mainframe platform

- A common theme from my interviews was that if not for parallel Sysplex and data sharing, migration off the mainframe would have happened much quicker and much sooner

- Many larger customers could not have grown as needed
  - Once CMOS processors matured, parallel Sysplex gave the option of both horizontal and vertical growth
  - There are some pretty large parallel Sysplex and data sharing environments

- Initial objective of capacity, and is still true
  - Today's (year 2019) CMOS will not get much faster any more
  - Thus, today adding more images will scale much better than high MPs

What is the future of Parallel Sysplex and Data Sharing?

Future of Parallel Sysplex and Data Sharing

- From a technical enhancement point-of-view
  - I have no idea
- Assumptions is that shops will continue to use as they are today
  - Some installations are still growing and needing more capacity
- Remember, today's processors are not going to get much faster, but many customer workloads are continuing to grow
- But interesting questions are
  - As more companies out-source their environments, how will outsources influence/force customer Sysplex decisions
  - Mainframe as a Service (MaaS) cloud will also be interesting
  - As more customers leave the mainframe, the residual workloads left behind may not need parallel Sysplex or data sharing (i.e. Sysplex collapse)

Future of Parallel Sysplex and Data Sharing

- Mike Swanson
  - IBM Fellow - Retired

  - As you know, I've been out of this game for awhile. However, a couple of things you could consider.
  - *… Finally, as a framework for thinking on the future of parallel Sysplex, consider what it is in place to address -- multisystem sharing of data and distribution/balancing of workload.*
  - *With that in mind, as new data, data access, data usage and data relevance to critical customer environments evolves there may be opportunities to extend parallel Sysplex facilities across hardware, microcode, z/OS and subsystems to address data sharing.*
  - *As workloads evolve, both existing and new forms of program execution, there may be ways to extend those facilities to meet the goals of scale, linear growth, high performance and availability.*

Thank you!

- Although I was not personally involved in design of development of any of the parallel Sysplex or data sharing hardware or software development, I know there is still so much more to discover that historic time period of IBM
  - I wish we had more time!
  - There are so many fun stories to tell

- During this presentation I talked to a wide array of people, but I specifically want to thank the Mike Swanson, retired IBM Fellow for his unique and historical insights

- A special thank you to Mythili Venkatakrishnan Distinguished Engineer, IBM

- Also, thank you to Bob Rogers (IBM retired, Trident, friend), Scott Chapman of EPS, and Al Sherkow of I/S Management Strategies for their insights and input