Inventory Managed Resources/Workloads

Peter Enrico

Email: Peter.Enrico@EPStrategies.com

Enterprise Performance Strategies, Inc.
3457-53rd Avenue North, #145
Bradenton, FL 34210
http://www.epstrategies.com
http://www.pivotor.com

Voice: 813-435-2297
Mobile: 941-685-6789

EPS
z/OS Performance
Education, Software, and
Managed Service Providers

PIVOTOR®

Creators of Pivotor®

Contact, Copyright, and Trademark Notices

Questions?

Send email to Peter at Peter.Enrico@EPStrategies.com, or visit our website at http://www.epstrategies.com or http://www.pivotor.com.

Copyright Notice:

Trademarks:

Enterprise Performance Strategies, Inc. presentation materials contain trademarks and registered trademarks of several companies.

The following are trademarks of Enterprise Performance Strategies, Inc.: **Health Check®, Reductions®, Pivotor®**

The following are trademarks of the International Business Machines Corporation in the United States and/or other countries: IBM®, z/OS®, zSeries® WebSphere®, CICS®, DB2®, S390®, WebSphere Application Server®, and many others.

Other trademarks and registered trademarks may exist in this presentation

Abstract

- **Inventory Managed Resources/Workloads**

- Before doing any sort of performance valuation, it is super important to first understand both your managed resources and your workloads.

- During the webinar, Peter Enrico will go through a list of important details you should learn about both your resources and workloads to help you prepare for any performance analysis assignment.

EPS: We do z/OS performance…

- We love to bring you these complimentary webinars and to teach you how to do what we do!

- But remember, we would love to work with you!
  - *Pivotor - Performance reporting and analysis of your z/OS measurements (i.e. SMF, etc.)*
    - Not just reporting, but coset effective analysis-based reporting based on our expertise

  - *Performance Educational Workshops (while analyzing your own data)*
    - Essential z/OS Performance Tuning
    - Parallel Sysplex and z/OS Performance Tuning
    - WLM Performance and Re-evaluating Goals

  - *Performance War Rooms*
    - Concentrated, highly productive group discussions and analysis

The Performance Balancing Act

- Performance on z/OS is about finding an optimal balance among 3 areas

Optimal Workload Performance (so Customers are happy)

Resource Optimization (so Management is happy)

Optimal MSU Consumption (so Financial people are happy)
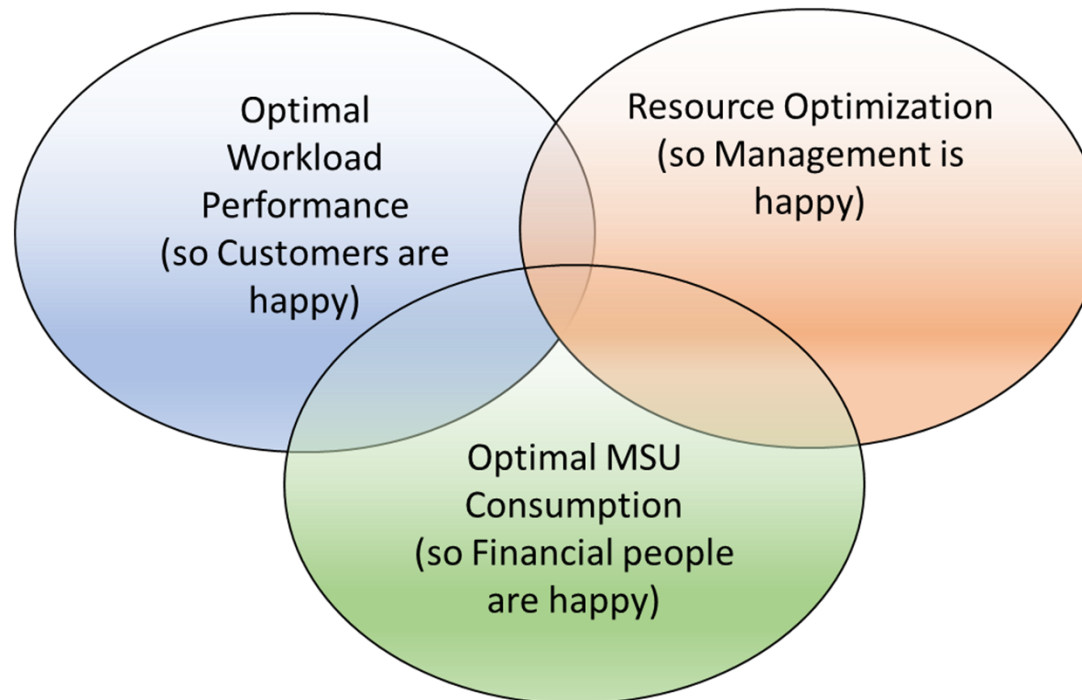
Cookbook Approach to Performing a
z System Performance Health Check

- ## High level steps for revisiting your WLM setup and service definition
    - Step 1:   Inventory Your Managed Resources
    - Step 2:   Inventory System Workloads
    - Step 3:   Understand Current WLM Definitions and System Parameters
    - Step 4:   Learn How to Interpret Measurements
    - Step 5:   Analyze the basic system resources
        - Processor, Storage, DASD I/O, Virtual Storage,
    - Step 6:   Analyze your workload performance
        - Including WLM Service Definition and goals
    - Step 7:   Analyze Sysplex communication and resources
        - XCF, Coupling Facility and Coupling Facility Structures
    - Step 8:   Analyze your Subsystem work managers
        - CICS, IMS, DB2, WebSphere, MQ, etc, etc, etc…
    - Step 9:   Analyze your applications
    - Step 10:   Write a report

Thought Exercise

That is:

- If you are newbie, or someone in a new role, it is always good to become familiar with the resources and workloads.
  - Explore, discover and document!

- If you are an experienced person, it is always good to revisit your understanding of resources and workloads.
  - Make the job for the next person a little easier. Document!

Step 1:
Inventory Your Managed Resources
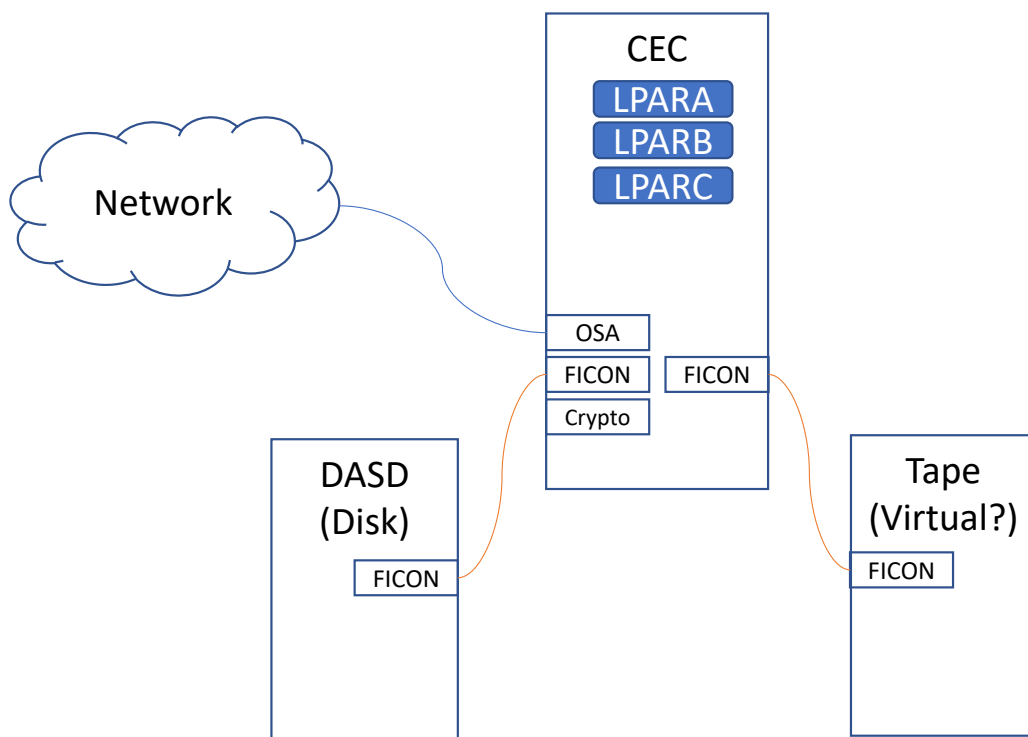
When analyzing any basic system resources

- Incredible performance analysis and problem debug insights can be gained when you have a great understanding of physical resource environment
  - What primary resources are available
  - How is each resource physically configured
  - How is each resource logically configured
  - What are the usage patterns of each of these resources
    - Usage breakdown by physical machine
    - Usage breakdown by logical entity
  - What is the usage patterns of the resource by workloads
    - Example: by WLM service class, and by WLM report class
  - How is the resource being utilized by each address space and transaction

- Basically, we want to draw many pictures

Step 1: Inventory the basic system resources

- Processor configuration, utilizations, contentions, and usage of functions
    - Inventory processor machine resources – CPs, zIIPs, ICFs, IFLs, etc.
    - Physical and logical configuration – number of logical engines to each partition, weights, mins and maxes, when LPARs are active, etc.
    - Inventory processor usage for machine, by LPAR, by workload, etc.
    - Exploitation of functions that control processor usage – Caps, HiperDispatch pooling, SMT, crypto, etc.
- Storage
    - Inventory storage controllers, I/O vendor, machine and model,
    - Physical and logical configurations
    - Inventory storage usage at
        - DASDplex level
        - Storage controller level
        - Storage group level
        - Logical volume
    - Channels and interfaces
    - Peripheral details and usage (such as PAVs)
- Memory
    - Inventory physical memory available on the machine
    - Inventory memory configured to each partition
    - Inventory memory usage by workload, by area, by page size
    - Paging dataset usage and configuration
- Coupling facility
- Virtual storage
- Etc.

A Simple Mainframe Environment

CEC

- LPARA
- LPARB
- LPARC

OSA
FICON | FICON
Crypto

Network

DASD (Disk)
FICON

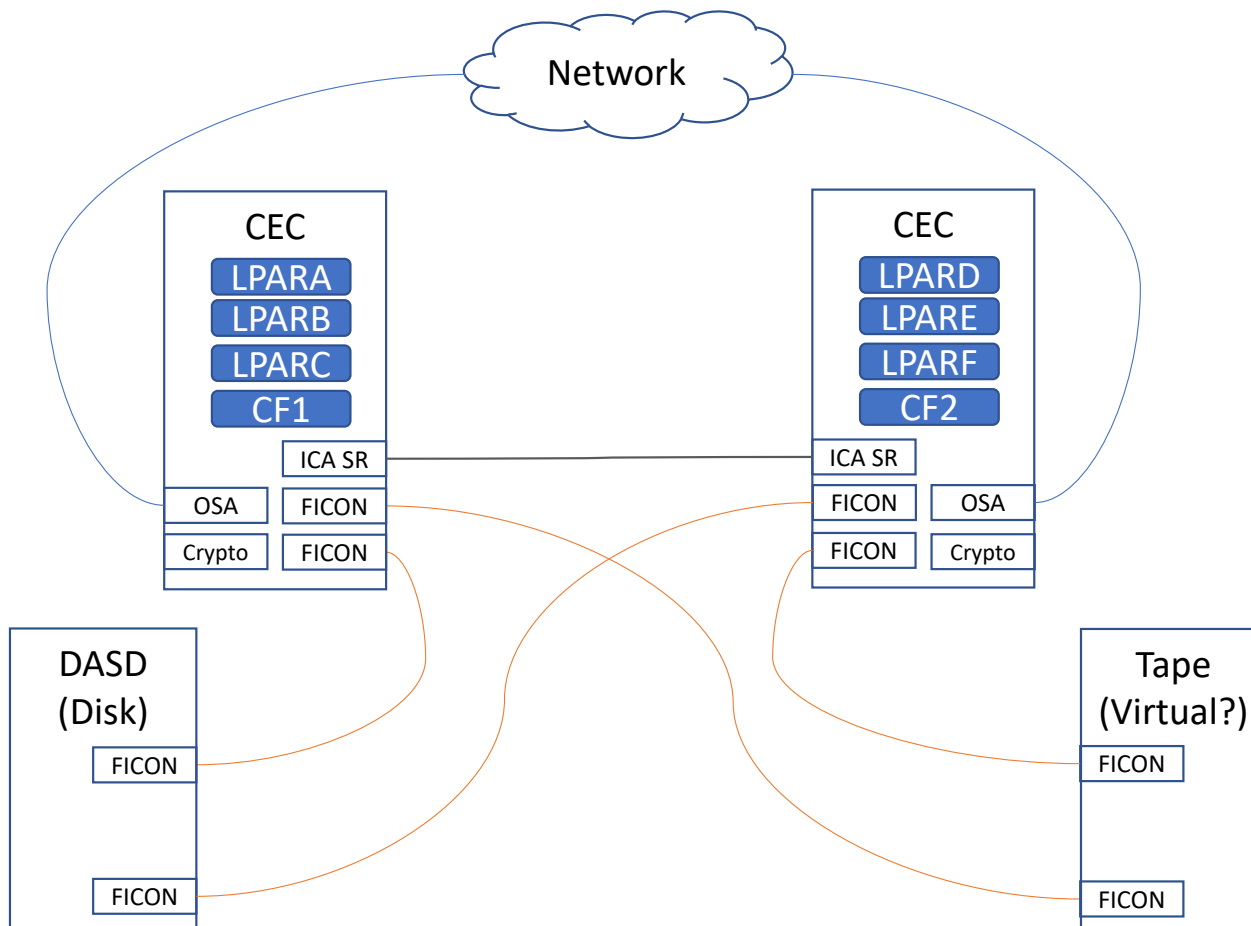Tape (Virtual?)
FICON

- Not to scale
- I/O card positions for drawing convenience only
- Should have multiple of all components (e.g. multiple FICON connections from multiple cards)
- There may be channel-attached printers as well (possibly going through FICON to ESCON convertors)
  - ESCON not natively available on z12 and later machines
- Multiple options exist for each type of interface card.

A Common Sysplex

FICON CTCs not shown, but quite possibly present

Network

CEC
- LPARA
- LPARB
- LPARC
- CF1

ICA SR

OSA | FICON

Crypto | FICON

CEC
- LPARD
- LPARE
- LPARF
- CF2

ICA SR

FICON | OSA

FICON | Crypto

DASD (Disk)
- FICON
- FICON

Tape (Virtual?)
- FICON
- FICON

A Somewhat More Complex Sysplex

Network

CEC
- LPARA
- LPARB
- LPARC

CE LR
zEDC | ICA SR
OSA | FICON
Crypto | FICON

CF1

CF2

CEC
- LPARD
- LPARE
- LPARF

ICA SR
CE LR | zEDC
FICON | OSA
FICON | Crypto

There's a lot of places to look for performance issues!
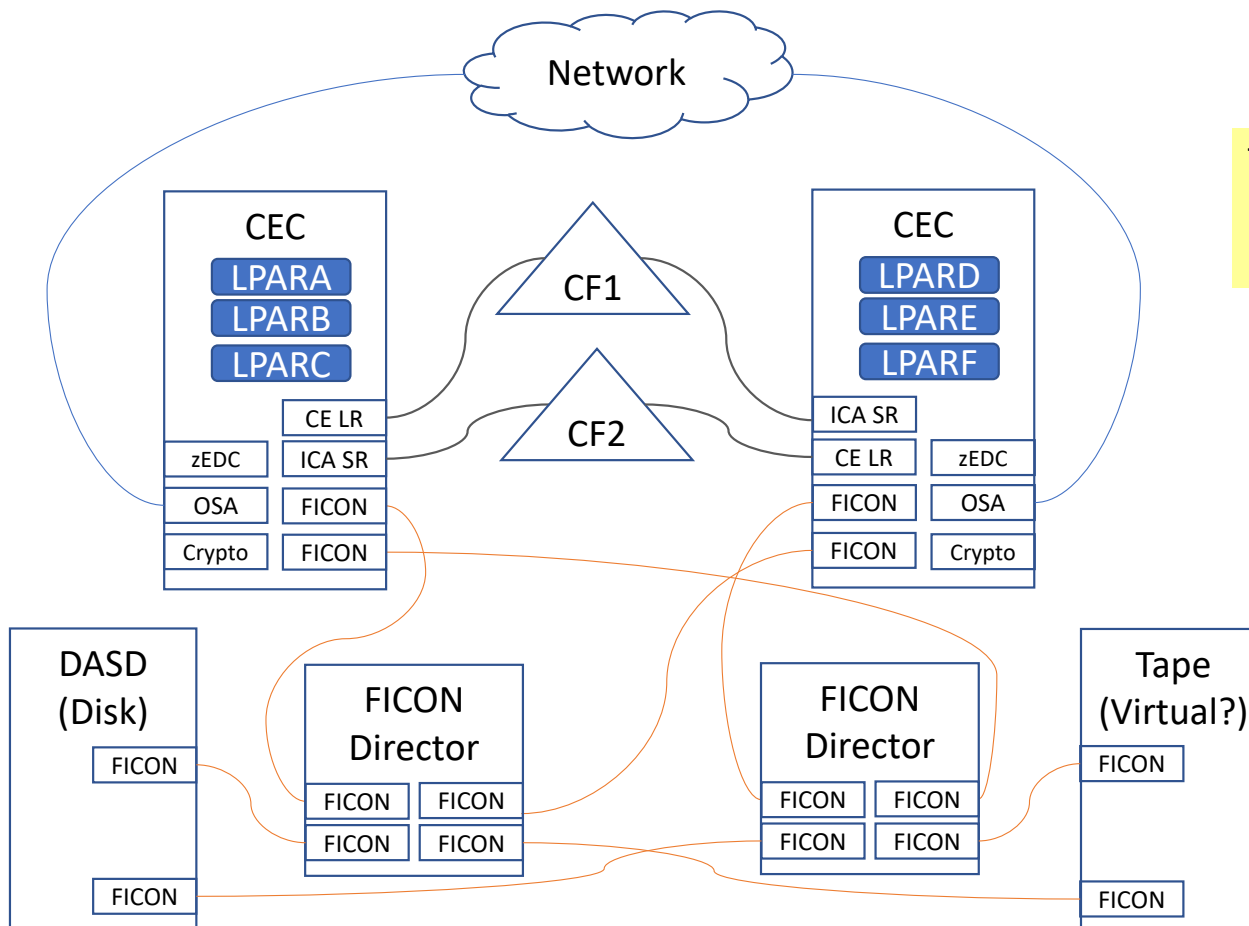
CE LR implies the machines might be in different data centers.

DASD (Disk)
- FICON
- FICON

FICON Director
- FICON | FICON
- FICON | FICON

FICON Director
- FICON | FICON
- FICON | FICON

Tape (Virtual?)
- FICON
- FICON

Breakdown of General-Purpose Processor

- We always needed to understand the break down of CP CPU consumption

Total Physical Processor (CEC) Time						
Total LPAR Dispatch Time (Partition 1)			etc	Total LPAR Dispatch Time (Partition N)		*PHYSICAL*
Effective LPAR Dispatch Time		Partition LPAR Mgmt Time	...	Effective LPAR Dispatch Time		Partition LPAR Mgmt Time
MVS Captured Time	MVS Un-captured Time			MVS Captured Time	MVS Un-captured Time	LPAR Time to Manage LPAR

Workload CPU

Service Class CPU

Service Class Period CPU

Address Space or Enclave CPU (TCB,SRB,RCT,IIT,HST)

Workload CPU

Service Class CPU

Service Class Period CPU

Address Space or Enclave CPU (TCB,SRB,RCT,IIT,HST)
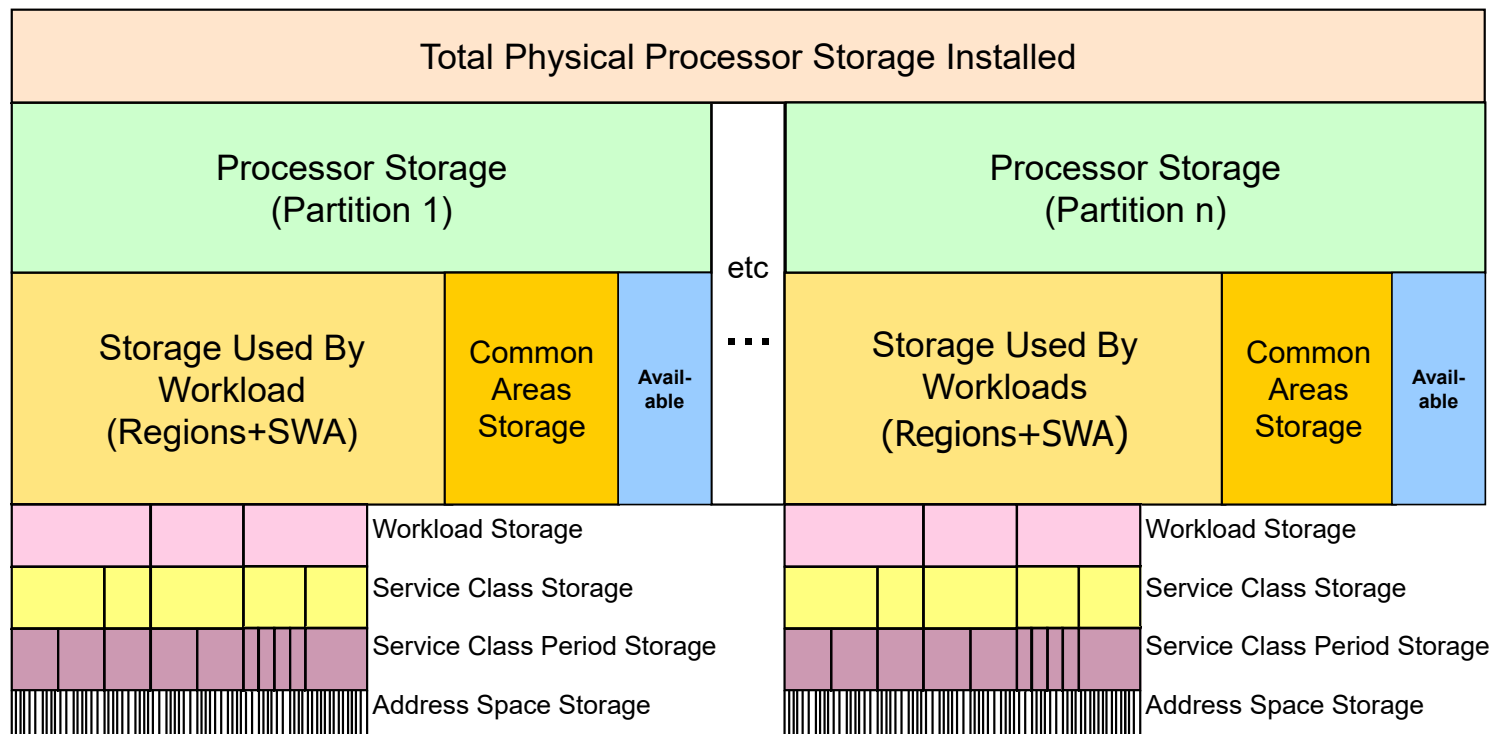
Breakdown of zIIP Engine Time

- We need to understand how PR/SM allocates the zIIP processor resource
  - In all measurements zIIPs
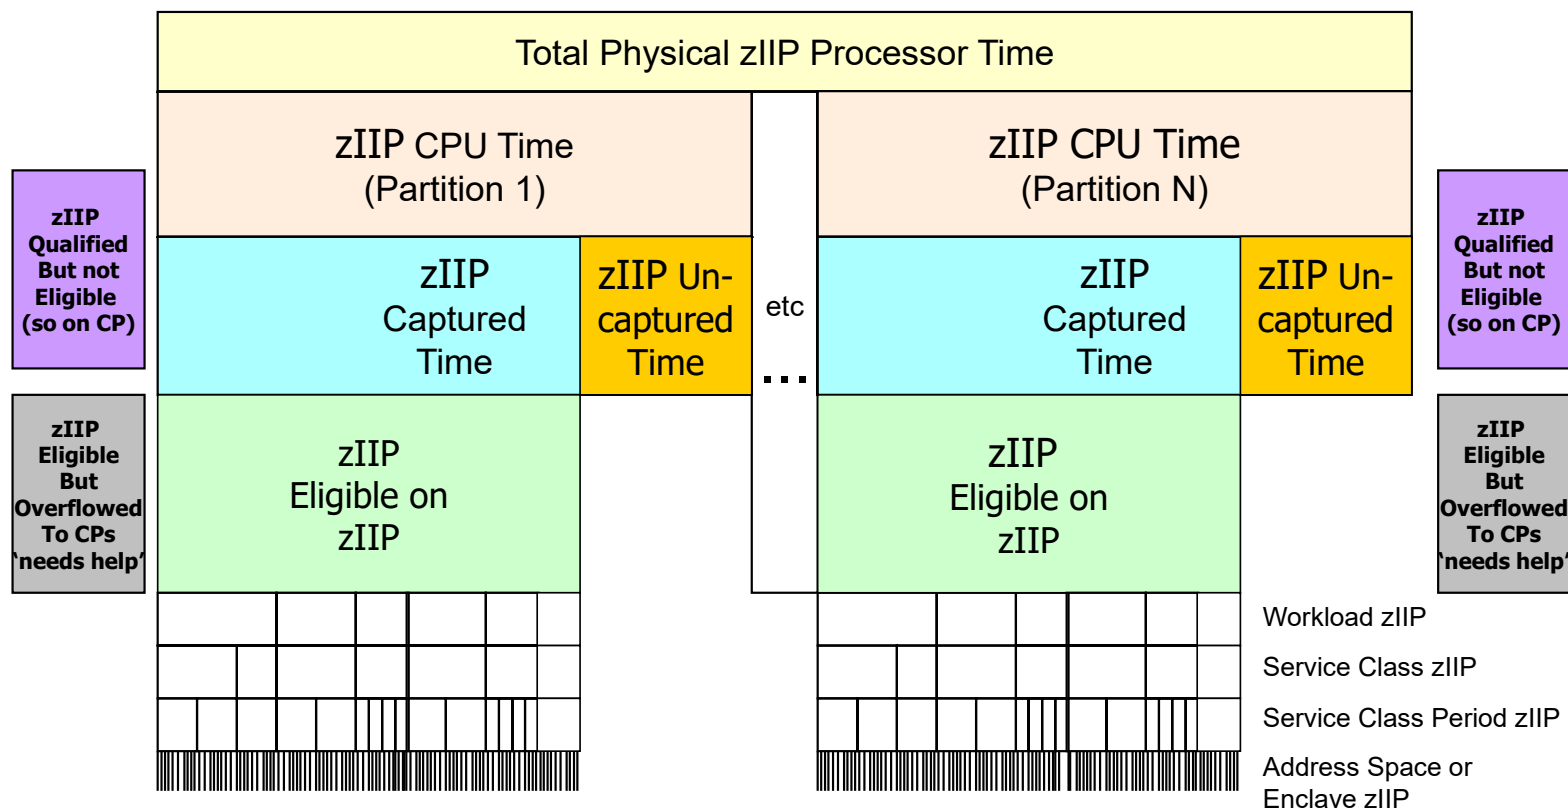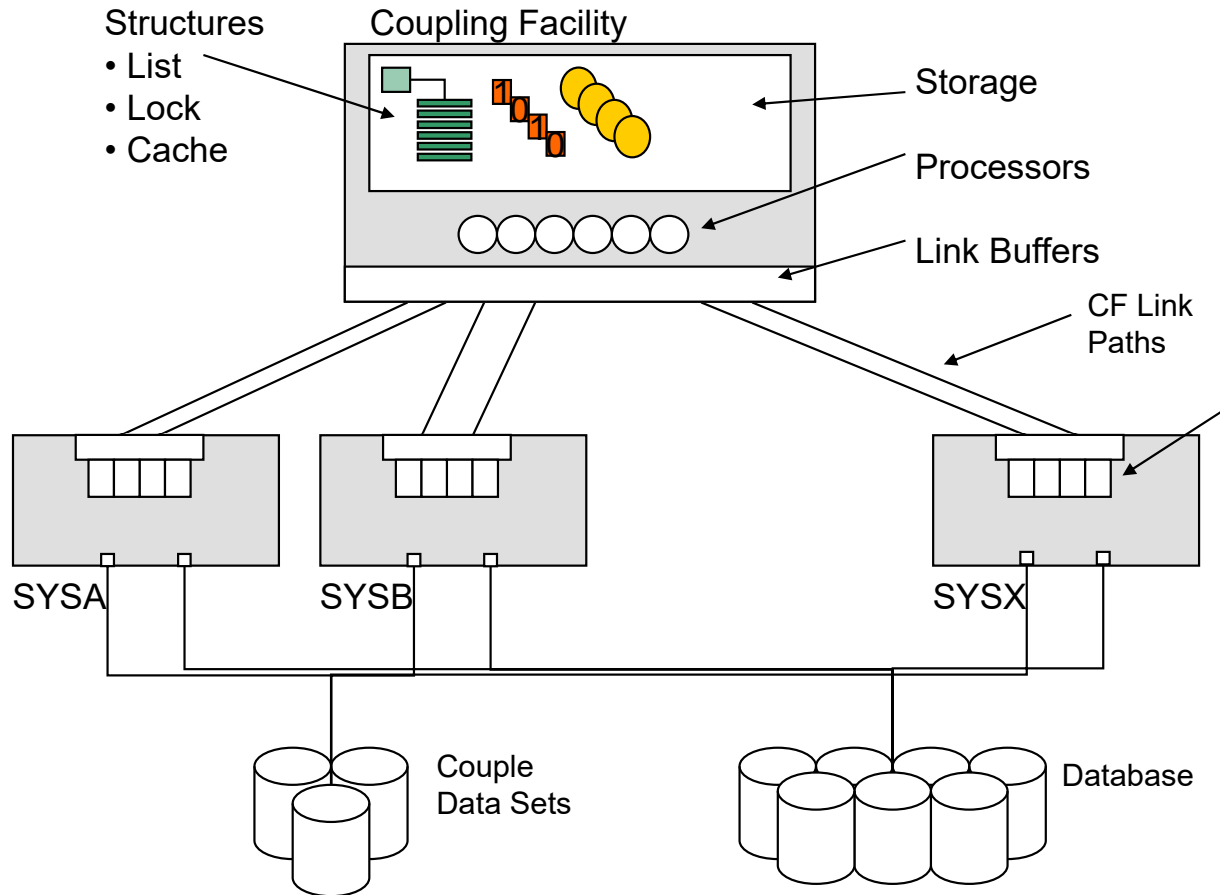
Breakdown of Storage

- We always needed to understand the break down of Storage usage

Understand Your Coupling Environment

Enterprise Performance Strategies, Inc. ©

Lots of other things to inventory

- When it comes to resources, there is so much more to inventory
  - Tape resources and configurations
  - Network resources and configurations
  - Hardware and software code levels
  - Disaster recovery configurations
  - Etc…

- The important point is, gain an understanding of your resources
  - Any knowledge about the resources will help you during any performance analysis, problem debug, and any exercise to maintain the environment.

Step 2:
Inventory System Workloads

When analyzing any basic system workloads

- Incredible performance analysis and problem debug insights can be gained when you have a great understanding of your customers and workloads

- It is simple to say:
  - "We have TSO, batch, CICS, DB2, DDF, some WAS, and a bunch of started tasks"

- Instead, you need to dig in, and really understand your workloads
  - Understand your workloads from a business point-of-view
  - Understand your workloads from software point-of-view
  - Understand your workloads from a transactional point-of-view
  - Map the business point-of-view to the software point-of-view

Understand your workloads from a business point-of-view

- **What is your business?**
  - It is easy enough to say you are a bank, or insurance company, service provider, etc.

- **But how does this translate to the business and the customer interactions?**

- **Example for a bank**
  - Day time banking and tellers
  - Online banking
  - Investment services
  - Legal requirements that may put requirements on batch
  - Backoffice services such as payroll
  - Etc…

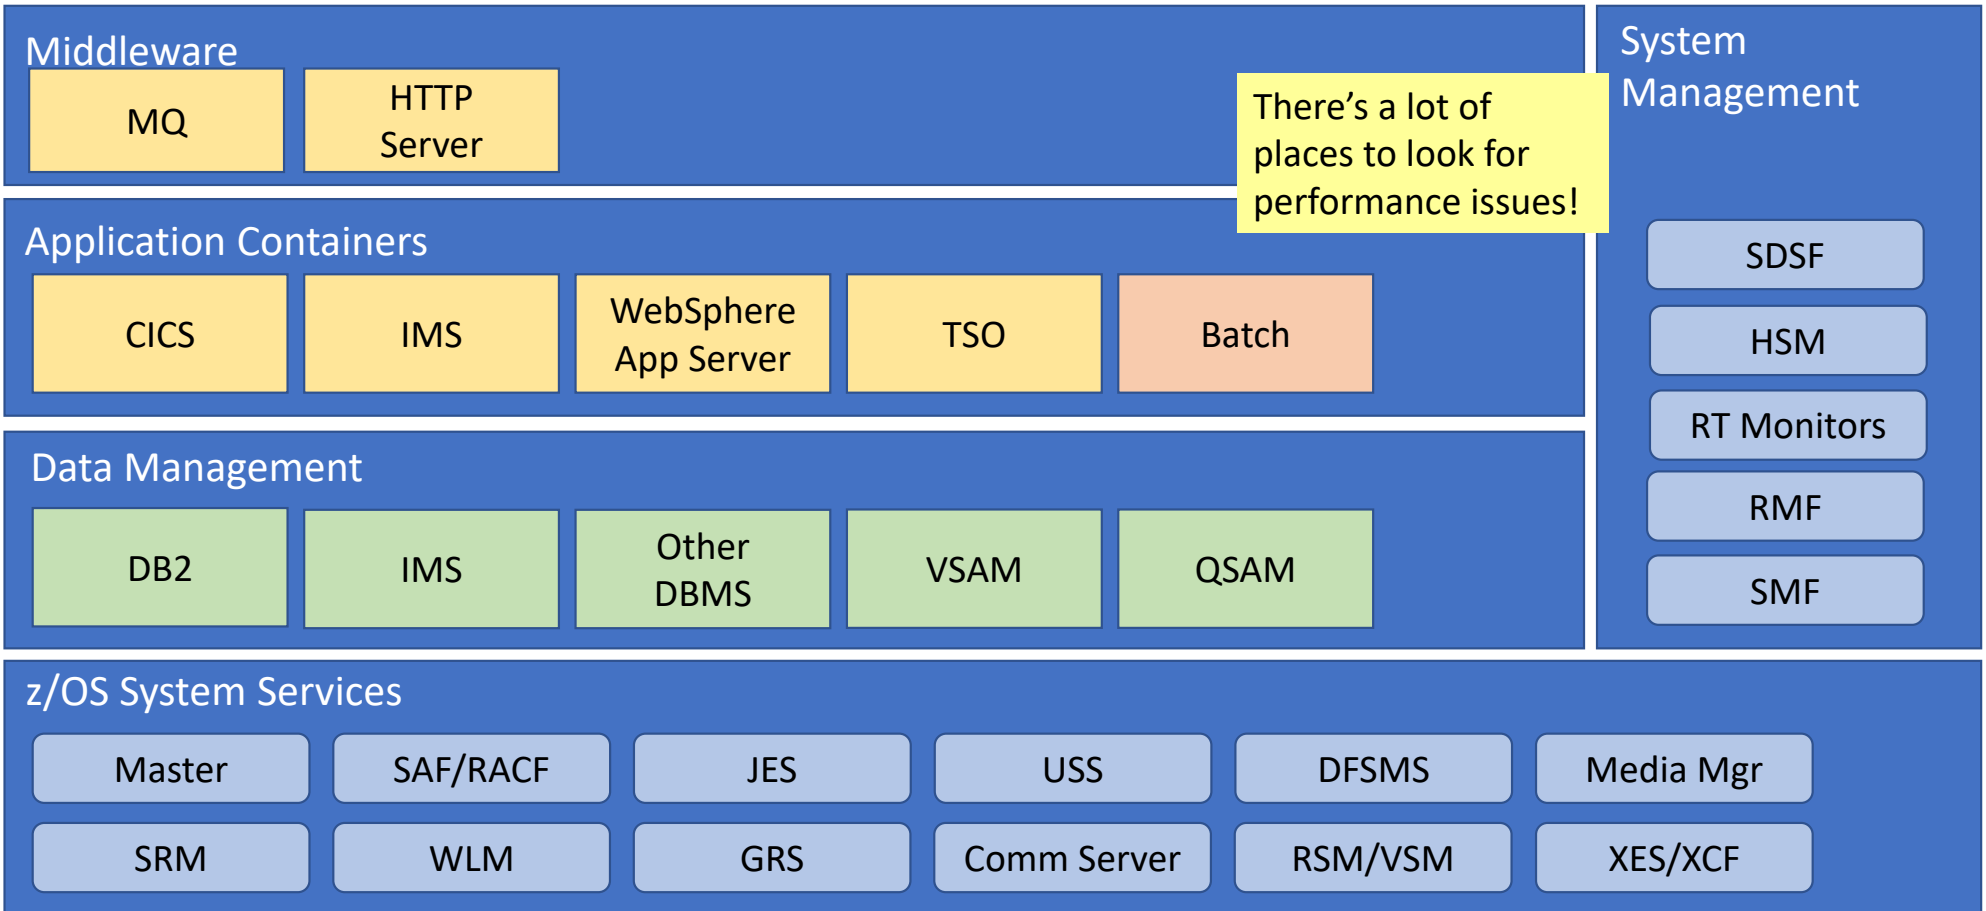- **Always great to have business metrics to map to performance data**

For each workload understand the basics

- Who is the end-user / customer?
  - Real customer paying money, or internal customer?
- What is that customer's expectation or SLA?
- What subsystems, database managers, and transaction managers does the workload require and utilize?
  - What products and applications make up this workload?
- What resources does the workload require
  - And what is the typical pattern of resource usage?
- When does the work want to run?
  - Daily daytime activity, nighttime, month end, quarter end, seasonal, ad-hoc?
- From where do the transactions originate?
- Where will this workload run?

Understand your workloads from S/W point-of-view

Middleware

| MQ | HTTP Server |

Application Containers

| CICS | IMS | WebSphere App Server | TSO | Batch |

Data Management

| DB2 | IMS | Other DBMS | VSAM | QSAM |

z/OS System Services

| Master | SAF/RACF | JES | USS | DFSMS | Media Mgr |
| SRM | WLM | GRS | Comm Server | RSM/VSM | XES/XCF |

System Management

| SDSF |
| HSM |
| RT Monitors |
| RMF |
| SMF |

There's a lot of places to look for performance issues!

Understand your workloads from a transactional point-of-view

- **The basic building blocks of workloads**
  - Address spaces and Enclaves
    - Given a workload, what address spaces and enclaves on the system make up that workload
  - Transactions
    - Given a workload, what is a delineated unit of work?
    - What are the key transactions that represent the health of the workload?
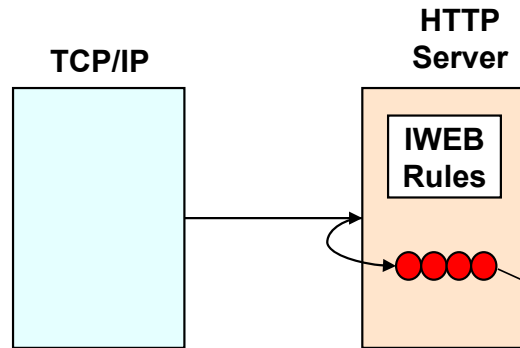 - Dispatchable units – represented by TCBs and SRBs

- **Summarization:**
  - Workloads are composed of address spaces and/or enclaves. These address spaces and enclaves are composed of dispatchable units represented by TCBs and SRBs. Transactions are the delineated units that represent work.

  - This understanding is key to gain an understanding of types of workloads

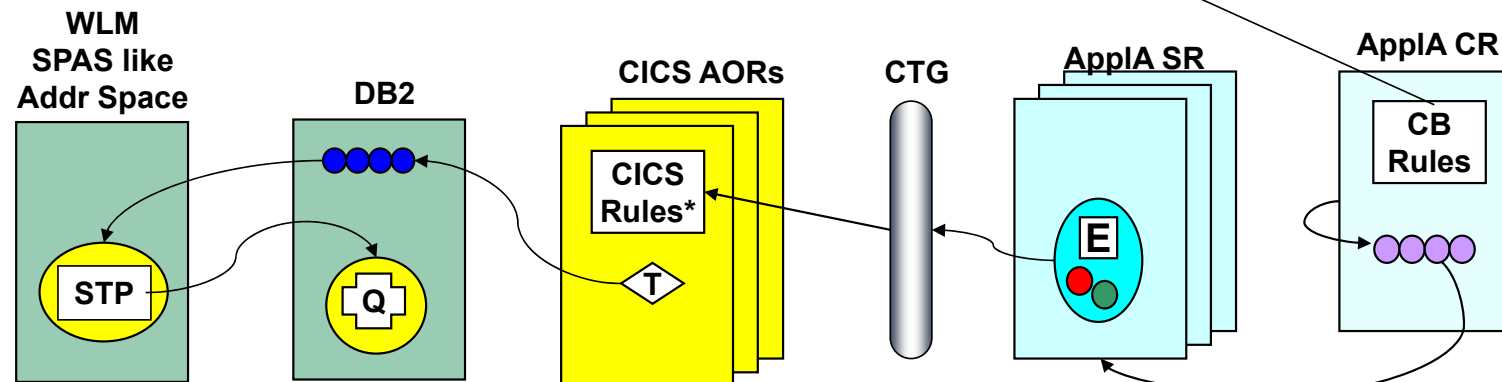# Understand your workloads from a transactional point-of-view

- Map the business unit to the S/W product to the address space and enclave level

- What is a transaction?

TCP/IP

HTTP Server

IWEB Rules

In this scenario, transaction may be classified up to three separate times

WLM SPAS like Addr Space

STP

DB2

Q

CICS AORs

CICS Rules*

T

CTG

ApplA SR

E

ApplA CR

CB Rules

*CICS Rules if WLM Transaction Management enabled

What are the workloads requirements?

- Different workload are managed differently
  - Different types of work require different types of goals
    - Long running work verses short running work
  - WLM could manage work differently depending on the type of work
    - Example: WebSphere may both appear to be 'transaction managers' but WLM manages them very differently
  - Workloads in a multi-system Sysplex
    - Work running on multiple systems
    - Work running just a few systems
  - Business objectives and service level agreements
    - Influence importance levels
  - End user objectives
    - Some work is short running and other work is long running
  - Interactions of the workloads
    - For example: Goals and classification for WebSphere transactions will probably be based on exploitation of CICS, DB2, web server, etc.
  - Types of workloads
    - Interactive work will have different types of goals than long running work
    - Some work is enclave based, and other work is address space based, and other work is something in-between

Understanding Workloads and Their Requirements

- z/OS is used to execute work to complete some set of functions
  - Groups of work on z/OS are known as 'Workloads'

- Not all workloads are alike
  - Interactive workloads versus background workloads
  - System workloads versus customer workloads
  - High importance workloads versus low importance workloads
  - High regular volume workloads versus low sporadic volume workloads
  - Short running work versus long running work
  - Distributed workloads versus single system workloads
  - e-business workload versus legacy workloads
  - And much more

- The workloads represent the customer and income

- Performance Fundamental
  - The ability to identify and understand the needs of the workloads

WLM is also a reason to understand workloads

- **In the world of WLM Keep work in each service class relatively homogeneous**
  - WLM takes different actions for different types of work to meet goals
  - Reports may not always reflect reality
  - Example: Don't mix CICS transactions in same service class as TSO

- **Separate unlike work**
  - Don't mix enclave work with non-enclave work
  - Don't mix interactive work with non-interactive work
  - Don't mix participants with non-participants
  - Don't mix server with non-servers
  - Don't mix regions managed towards region goal with regions managed towards trans goals
  - Don't mix Batch in WLM inits with batch in JES inits
  - Don't assign goals to spaces that should truly be in SYSTEM and SYSSTC
  - Don't put stuff into SYSTEM and SYSSTC that should not be there

Understanding Workloads:
Batch Example

- Most installations separate batch into Production versus Test, or High Importance versus low importance

- But understand who your customers are and why the batch workload exists:
  - Normal Production Jobs submitted by a Job Scheduler
  - Critical Path Jobs submitted through a Job Scheduler
  - Ad-hoc Jobs (possibly submitted by a Job Scheduler)
  - Development Jobs
  - Normal System Support Jobs
  - High-Priority System Support Jobs
  - Logs, Archival, Backup, and D/R Jobs
  - Quick Utility Jobs
  - Emergency or Hot Jobs

Understanding Workloads:
# Batch Example

- Also understand additional setup and resources need by the batch
  - Batch jobs that require input from media that is not normally online to a system
  - Batch jobs that require offline media to generate output have setup requirements that involve operator intervention

- May Include:
  - Manual tape mounts (reels or cartridges)
  - ATL/Silo Mounts- normally very quick unless media is not in the ATL
  - Virtual Tape Servers (VTS)- normally very quick unless data is not staged in the disc Cache
  - Direct SYSOUT Writers (not used much)

Understanding Workloads:
Batch Example

- Also understand where the batch originates
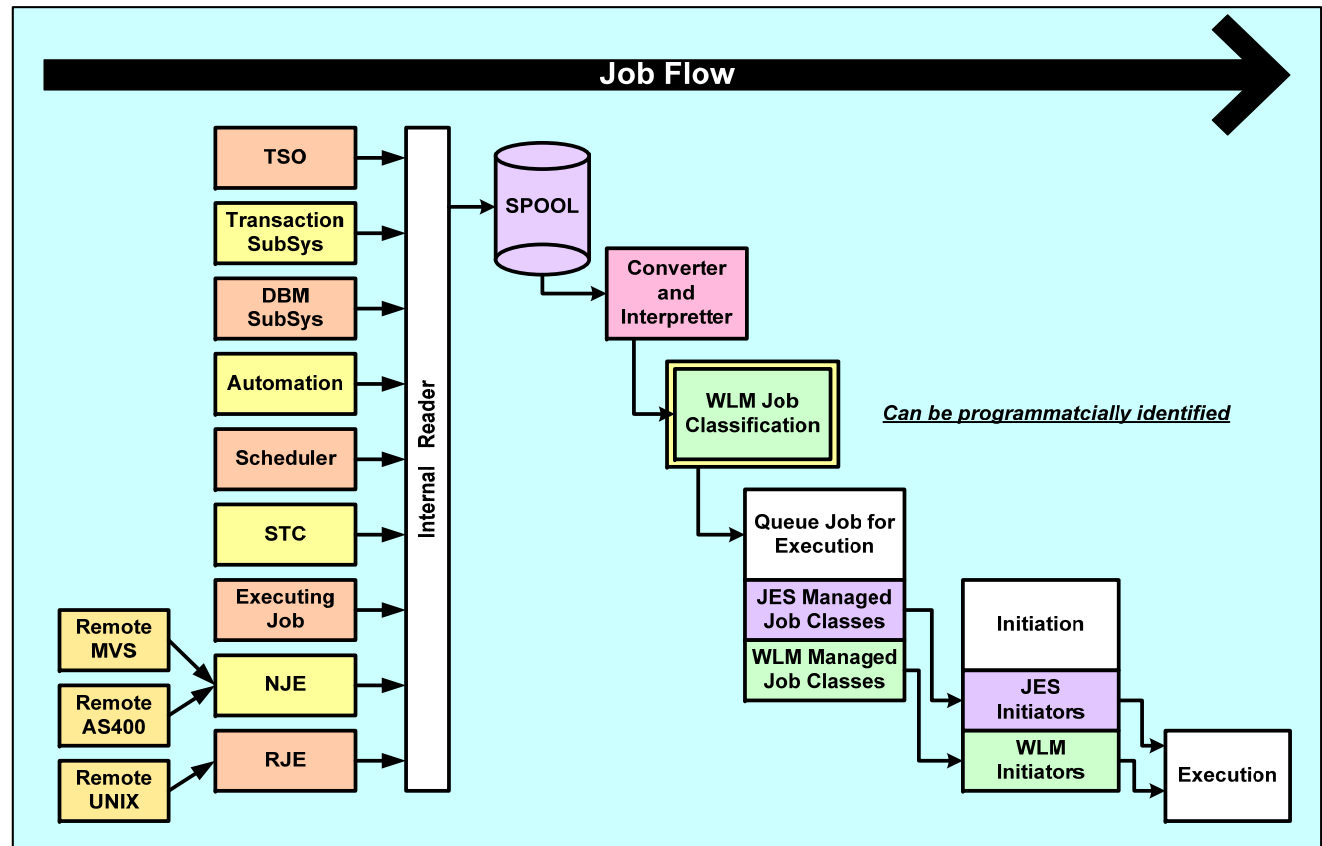
Now do this for all your workloads!

Questions?