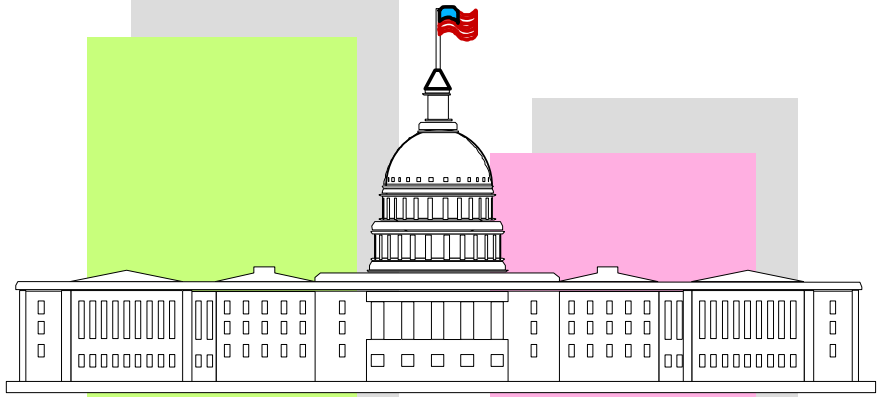

Washington Systems Center
Advanced Technical Support

# LSPR Background

# zSeries Capacity Planning

John L. Fitch
jlfitch @ us.ibm.com

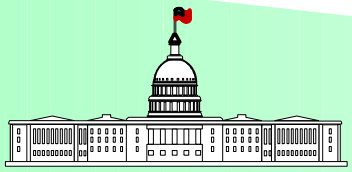

IBM Washington Systems Center

# Large Systems Performance Reference

*Purpose*
*To demonstrate,*
*through measurement,*
*zSeries & S/390 processor*
*capacity running various SCP's*
*and associated workload environments*

LSPR is a trademark of the IBM Corporation

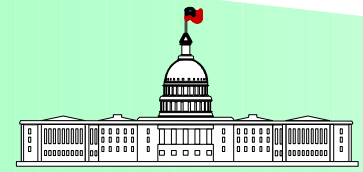



# LSPR Documentation

## Large Systems Performance Reference

Technical Bulletin **SC28-1187** (available via Internet)

- Processor capacity planning background
- Metrics for expressing processor capacity
- zSeries & S/390 workload environments
- Using **LSPR** data
- Validating a new processor capacity expectation
- **LSPR** ITR ratios for zSeries processors

## LSPR document and data from the Internet

http://www.ibm.com/servers/eserver/zseries/lspr/

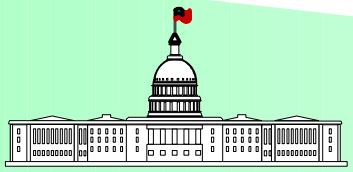



# LSPR Background

1. A set of representative SCP/workload environments
   - SCP's: **z/OS**, **z/VM**, and **Linux**
   - Workloads: **Batch** ⟷ **Online**
2. A methodology that focuses on processor capacity
   - No significant external constraints
   - Equivalent (reasonably high) processor utilization
3. A metric to communicate the results
   - **ITR** (**I**nternal **T**hroughput **R**ate)
   - Transactions or Jobs per processor busy second

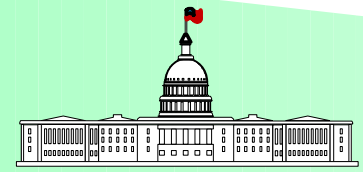



# LSPR Background

Focus is on processor capacity, without regard to

- External resources, including . . .
  - Central storage
  - Expanded storage
  - Channels
  - DASD devices
  - Operator activities
- Special features, such as . . .
  - Vector processing
  - ADMF
  - DB2 Sort Assist
  - Compression
  - Encryption

Capacity is represented by ITR values depicting each processor in its best light. Therefore, ratios between ITR values provide reasonable capacity relationships for capacity planning purposes.

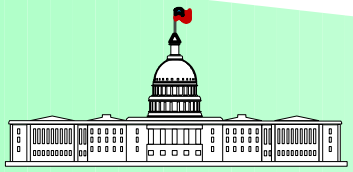



# The LSPR Advantage

## Workloads exploit

- ✓ **SCP** (**z/OS**, **z/VM**, **Linux**, etc.)
  JES2, VTAM, DFSMS, RACF, RMF/SMF
- ✓ **Application Subsystems**
  CICS, DB2, IMS, WebSphere
- ✓ **Access Methods**
  VSAM, BSAM, QSAM, BDAM
- ✓ **Program Products**
  ASSEMBLER, COBOL, FORTRAN, PL/1
  DFSORT, GDDM, ISPF, PL/1, SLR, SCRIPT
- ✓ **Transactions / Jobs**
  Includes a broad mix of activities
  Realistic end-user / initiator count
  Representative inter-arrival times

Representative work run in a representative way

No kernels or synthetic workloads

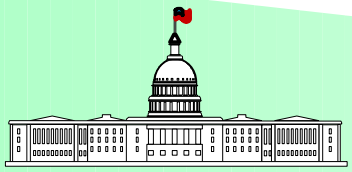



# LSPR Workloads

## z/OS

- **CB-L** Commercial batch - long, CPU-intensive jobs
- **CB-S** Commercial batch - short, I/O-intensive jobs
- **WASDB** WebSphere applications and DB2 data base
- **OLTP-W** Web enabled access to CICS/DB2 database
- **OLTP-T** Traditional IMS Data Systems online workload

## z/VM

- **CMS** CMS interactive user workload
- **WASDB/LVm** WebSphere under Linux as many small guests

## Linux

- **WASDB/L** WebSphere under Linux
- **EAS-AS/L** SAP Application server

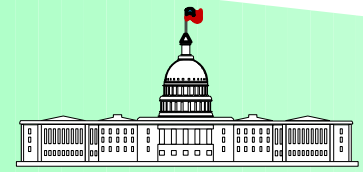



# LSPR Capacity Metrics

## ETR

- **E**xternal **T**hroughput **R**ate
- ✓ Bears direct relationship to work done
- ✓ Computed as work ÷ elapsed time
- ✓ Characterizes capacity of the entire processing system

## ITR

- **I**nternal **T**hroughput **R**ate
- ✓ Bears direct relationship to work done
- ✓ Computed as work done ÷ processor busy time
- ✓ Characterizes capacity of the processor itself

## ITRR

- **I**nternal **T**hroughput **R**ate **R**atio
- ✓ Statement of one processor's capacity relative to another
- ✓ The LSPR bottom-line

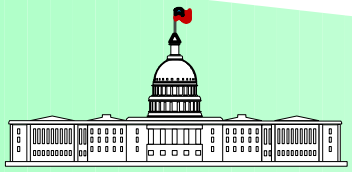



# ETR and ITR

**System #1:** *"the better system"*

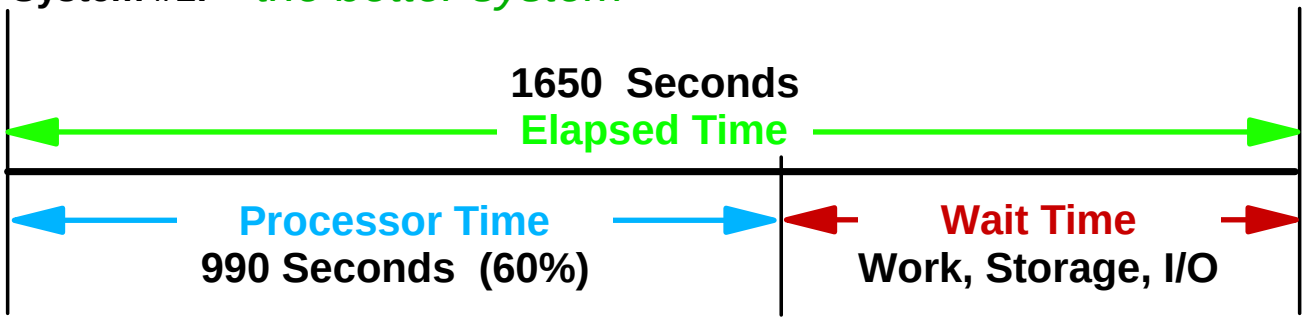


**System #2:** *"the better processor"*

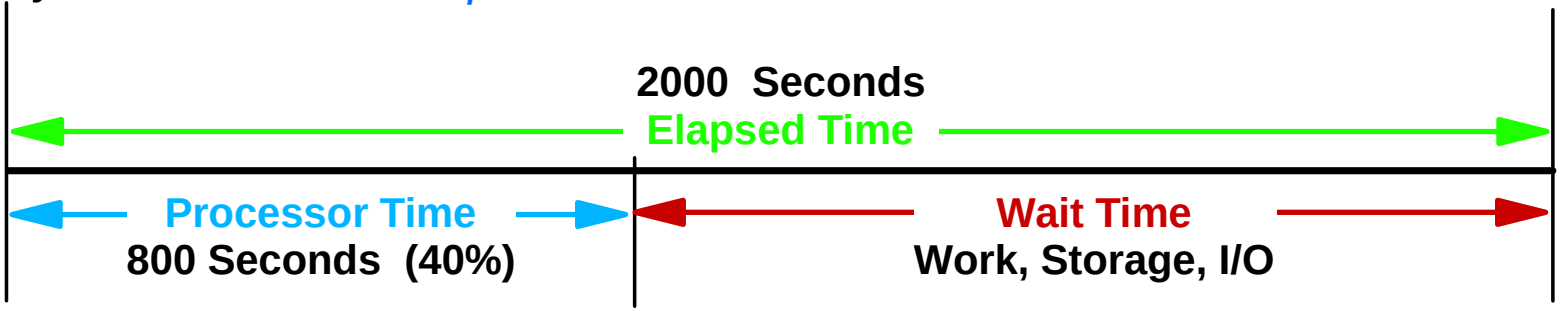

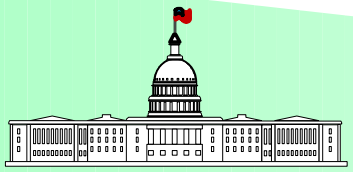



# ETR

External Throughput Rate

**Actual work done ÷ elapsed time**

Provides a methodology for comparing ***System Capacity***

- System is measured at/near a fully-loaded operating level
- Any resource is a potential capacity inhibitor, including
  - Central Processor
  - Memory (C-store and E-store)
  - Channels, CU's and I/O devices
  - TP network
  - Operator actions
- A response time criteria is commonly used to define when the system is saturated

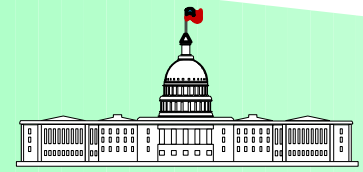



# ITR

Internal Throughput Rate

## Actual Work Done ÷ Processor Busy Time

Provides a methodology for comparing ***Processor Capacity***

- Processor is the only capacity inhibitor
  - Resources such as memory, channels, DASD, etc., are adequate to support the workload being measured
  - Minimal CPU is consumed managing resource constraints
- Processors are measured at equal utilization
- Processors are measured at reasonably high utilization
  - Batch is measured (start to end) at/near 100% utiliization
  - Online is measured (steady-state) at 90% utilization

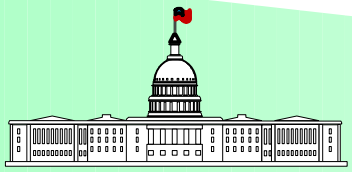



# ITR / ETR Relationship

$$\textbf{Utilization} = \frac{\text{Processor Busy Time}}{\text{Elapsed Time}}$$

$$\textbf{ETR} = \frac{\text{Units-of-Work}}{\text{Elapsed Time}}$$

$$\textbf{ITR} = \frac{\text{Units-of-Work}}{\text{Processor Busy Time}} = \frac{\textbf{ETR}}{\textbf{Utilization}}$$

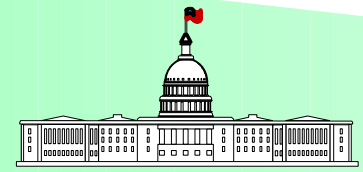



# LSPR Measurement Example

## Measurement Steps

1. Approximate user count to acheive target processor utilization
2. Setup system and restore data to startup condition
3. Logon estimated number of users (staggared over time)
4. Add/drop users to acheive target processor utilization
5. Allow system to acheive steady state
6. Measure for period deemed to be a repeatable sample
7. Analyze monitor data for problems/bottlenecks

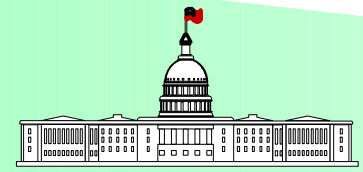



# LSPR Measurement Example

## z/OS with OLTP-W

| Measured data | 2064-2C2 | 2064-2C4 | Ratio |
|---|---|---|---|
| Elapsed seconds | 1,201.30 | 900.78 | |
| Processor seconds | 1,072.16 | 828.63 | |
| Transaction count | 201,374 | 298,609 | |
| **Calculated data** | | | |
| **Utilization** | **89.25 %** | **91.99 %** | |
| **ETR** | **167.630** | **331.500** | **1.98** |
| **ITR** | **187.821** | **360.365** | **1.92** |

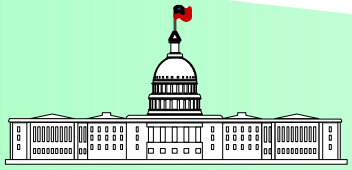

IBM Washington Systems Center

**Absolute ITR Values**
**z/OS LSPR Workloads as of 05/13/2003**

| Processor | Features | CB-L | CB-S | WASDB | OLTP-W | OLTP-T |
|---|---|---|---|---|---|---|
| **zSeries 900 Turbo** | | | | | | |
| 2064-2C1 | 1W | 0.12748 | 3.92286 | 253.601 | 97.610 | 196.468 |
| 2064-2C2 | 2W | 0.25277 | 7.06646 | 489.024 | 187.821 | 377.053 |
| 2064-2C3 | 3W | 0.37557 | 10.07036 | 720.555 | 275.406 | 549.791 |
| 2064-2C4 | 4W | 0.49588 | 12.93456 | 948.196 | 360.365 | 714.681 |
| 2064-2C5 | 5W | 0.61371 | 15.65907 | 1,171.945 | 442.700 | 871.725 |
| 2064-2C6 | 6W | 0.72906 | 18.24387 | 1,391.804 | 522.408 | 1,020.921 |
| 2064-2C7 | 7W | 0.84192 | 20.68898 | 1,607.771 | 599.492 | 1,162.271 |
| 2064-2C8 | 8W | 0.95230 | 22.99439 | 1,819.847 | 673.950 | 1,295.774 |
| 2064-2C9 | 9W | 1.06019 | 25.16010 | 2,028.032 | 745.783 | 1,421.429 |
| 2064-210 | 10W | 1.16559 | 27.18611 | 2,232.326 | 814.990 | 1,539.238 |
| 2064-211 | 11W | 1.26851 | 29.07243 | 2,432.729 | 881.572 | 1,649.200 |
| 2064-212 | 12W | 1.36895 | 30.81905 | 2,629.240 | 945.528 | 1,751.314 |
| 2064-213 | 13W | 1.46690 | 32.42597 | 2,821.861 | 1,006.859 | 1,845.582 |
| 2064-214 | 14W | 1.56236 | 33.89319 | 3,010.590 | 1,065.565 | 1,932.003 |
| 2064-215 | 15W | 1.65535 | 35.22071 | 3,195.428 | 1,121.645 | 2,010.577 |
| 2064-216 | 16W | 1.74584 | 36.40854 | 3,376.375 | 1,175.100 | 2,081.304 |
| **zSeries 990** | | | | | | |
| 2084-301 | 1W | 0.20000 | 6.17591 | 408.637 | 153.050 | 303.443 |
| 2084-302 | 2W | 0.38975 | 11.08934 | 777.099 | 294.670 | 582.525 |
| 2084-303 | 3W | 0.57570 | 15.77285 | 1,138.765 | 431.510 | 849.637 |
| 2084-304 | 4W | 0.75784 | 20.22642 | 1,493.636 | 563.569 | 1,104.779 |
| 2084-305 | 5W | 0.93617 | 24.45007 | 1,841.710 | 690.847 | 1,347.950 |
| 2084-306 | 6W | 1.11069 | 28.44378 | 2,182.989 | 813.345 | 1,579.151 |
| 2084-307 | 7W | 1.28140 | 32.20757 | 2,517.471 | 931.063 | 1,798.381 |
| 2084-308 | 8W | 1.44831 | 35.74142 | 2,845.158 | 1,044.000 | 2,005.640 |
| 2084-309 | 9W | 1.61141 | 39.04535 | 3,166.049 | 1,152.157 | 2,200.929 |

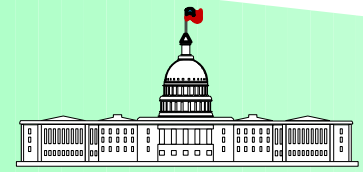



# How LSPR Data Can Be Used

1. Determine relative capacity for a potential new processor, to that of a currently installed processor, based on the SCP/workload environment
2. Scale relative capacity to any single processor's assumed (**MIPS**) rating (generally the currently installed processor would be used)
3. Define a mix of workloads to represent the current production workload, providing a tailored capacity expectation for potential new processors
4. Predict processor life assuming workload growth rates for each of the various production workload components
5. Assess capacity for workload consolidation when planning to merge various workload components running on different processors onto a single processor
6. Analyze workload metrics for the LSPR workloads and assesses the effects of changing them (e.g.; CPU intensity and inter-arrival time) on the number of users that could be supported.

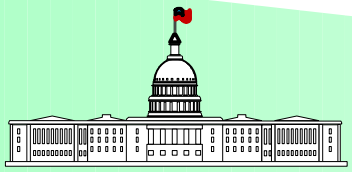



# Developing an LSPR Mixed-Workload Capacity Relationship

**Formula to compute LSPR capacity ratio for a mixed-workload**

**A *Harmonic Mean* calculation**

$$\frac{1}{\left(\frac{\%\ \text{CB-L}}{\text{CB-L ITRR}} + \frac{\%\ \text{CB-S}}{\text{CB-S ITRR}} + \frac{\%\ \text{WASDB}}{\text{WASDB ITRR}} + \frac{\%\ \text{OLTP-W}}{\text{OLTP-W ITRR}} + \frac{\%\ \text{OLTP-T}}{\text{OLTP-T ITRR}}\right)}$$

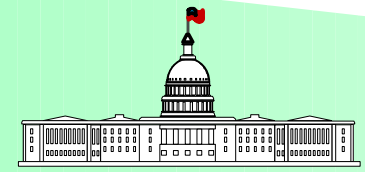



# Mainframe Capacity Growth

## Maximum Capacity for a Single OS Image

MIPS rating for Average Workload (LSPR-Mix)

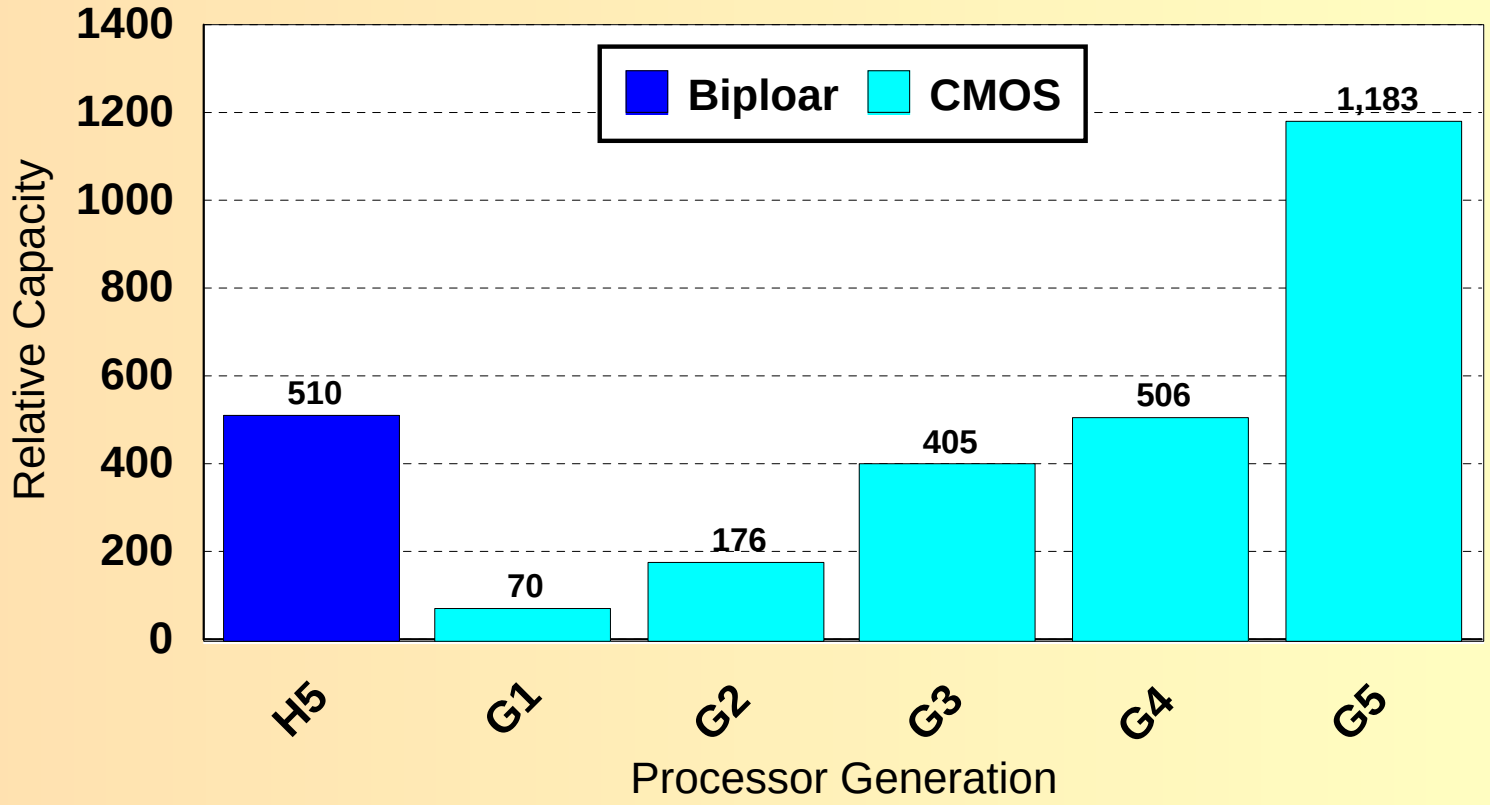

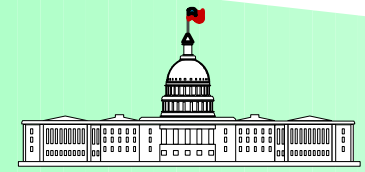

IBM Washington Systems Center

# Mainframe Capacity Growth

## Maximum Capacity for a Single OS Image

MIPS rating for Average Workload (LSPR-Mix)

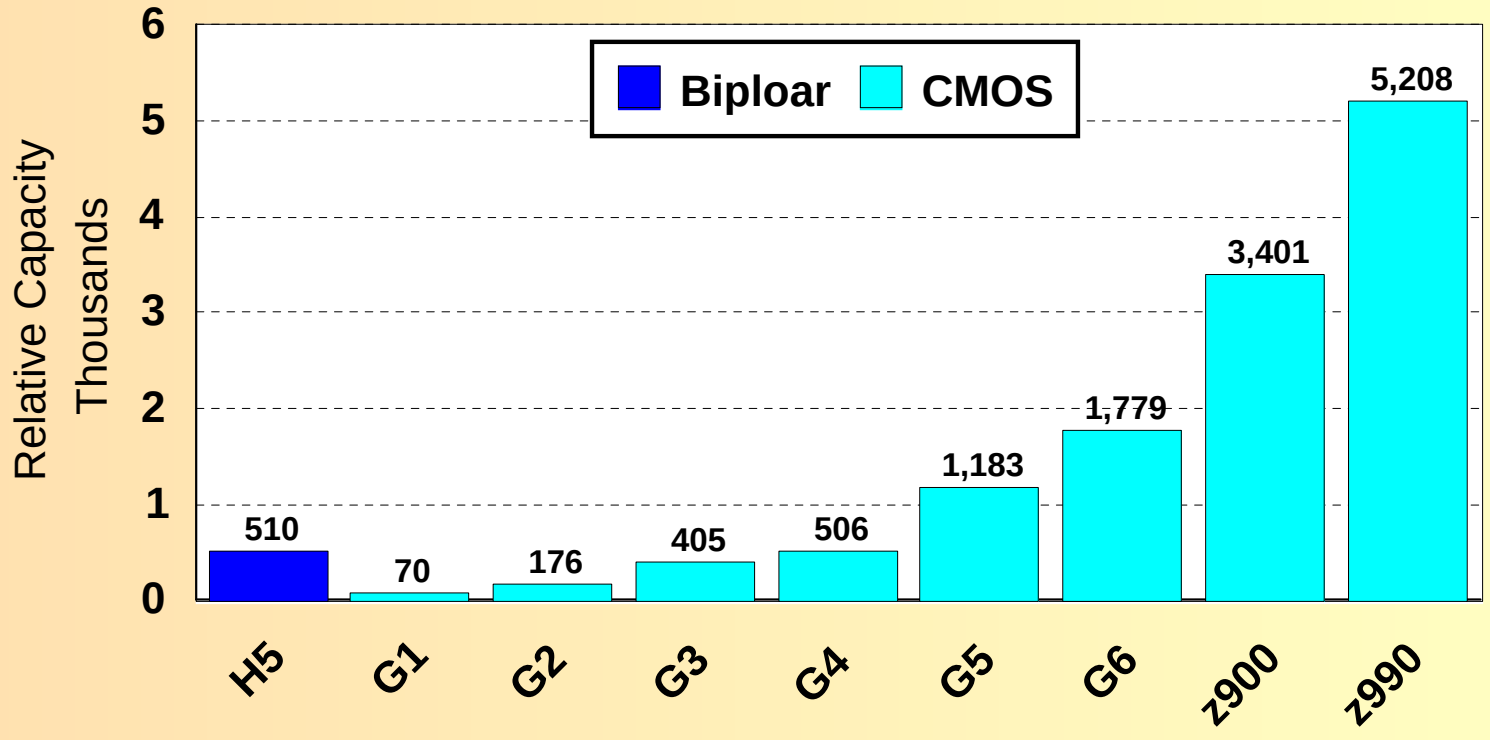

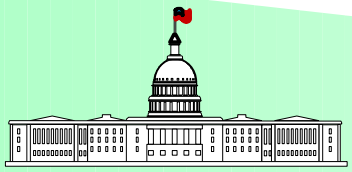

IBM Washington Systems Center

# Changes across IBM Processor Families that Affect Capacity

## Bipolar

### H5

Slow cycle time
Complex hardware logic
Super scalar
Branch history table
Instruction prefetch
Multiple decode
Multile E-units
Out-of-order execution
High speed buffer
- L1 - 2x128KB
- L2 - 2x4MB shared

## CMOS

### G4

Fast cycle time
Hardware / uCode logic
Sequential decode
In-order execution
High speed circuits
High speed buffer
- L1 - 64KB
- L2 - 4x768KB private
- L2.5 - 2MB

### G5

Faster cycle time
More hardware logic
G4 plus ...
Branch history table
More sophisticated HSB
- L1 - 256KB
- L2 - 2x4MB private

### z990

Faster cycle time
More hardware logic
G5 plus ...
Super scalar
Multiple inst per cycle
2nd level TLB
64-bit addressing
Up to 4 books (32 CPs)
More sophisticated HSB
- L1 - per CP
  - 256KB instruction
  - 256KB data
- L2 - per book
  - 32MB shared

LPAR-mode only

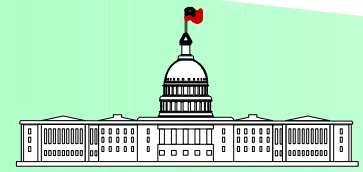



# MIPS as a Processor Metric

- **M**illions of **I**nstructions **P**er **S**econd
- ✓ Meaningless Indicator of Performance
- ✓ A ***ball-park*** (average) indicator of processor capacity
- ✗ No correlation to actual instruction rate
- ✗ Implies some type of ***universal scale***, loosely associated with actual Instruction Execution Rate (**IER)**
- ✗ A ***single-number metric***, which is insensitive to workload

**Capacity is dependent on workload type and partition configuration**

**Actual capacity can vary +/- 20% or more from relationships determined via MIPS tables**

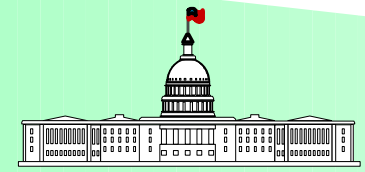



# MIPS Table

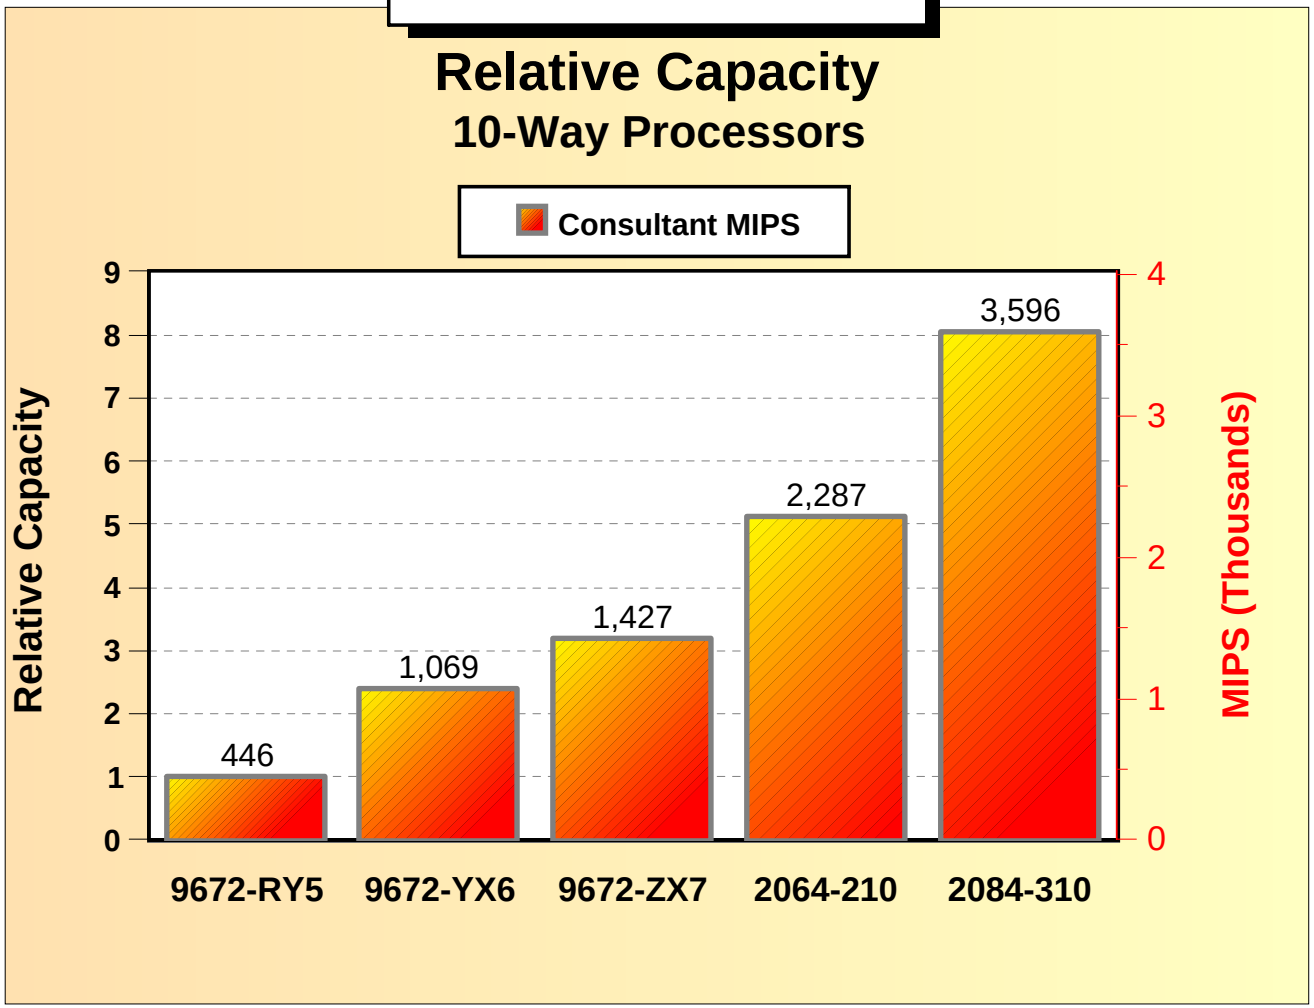
Relative Capacity
10-Way Processors
Consultant MIPS
Relative Capacity
MIPS (Thousands)
0
1
2
3
4
5
6
7
8
9
0
1
2
3
4
446
1,069
1,427
2,287
3,596
9672-RY5
9672-YX6
9672-ZX7
2064-210
2084-310

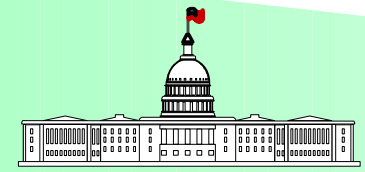

IBM Washington Systems Center

# LSPR

## Relative Capacity
### 10-Way Processors
z/OS V1R4

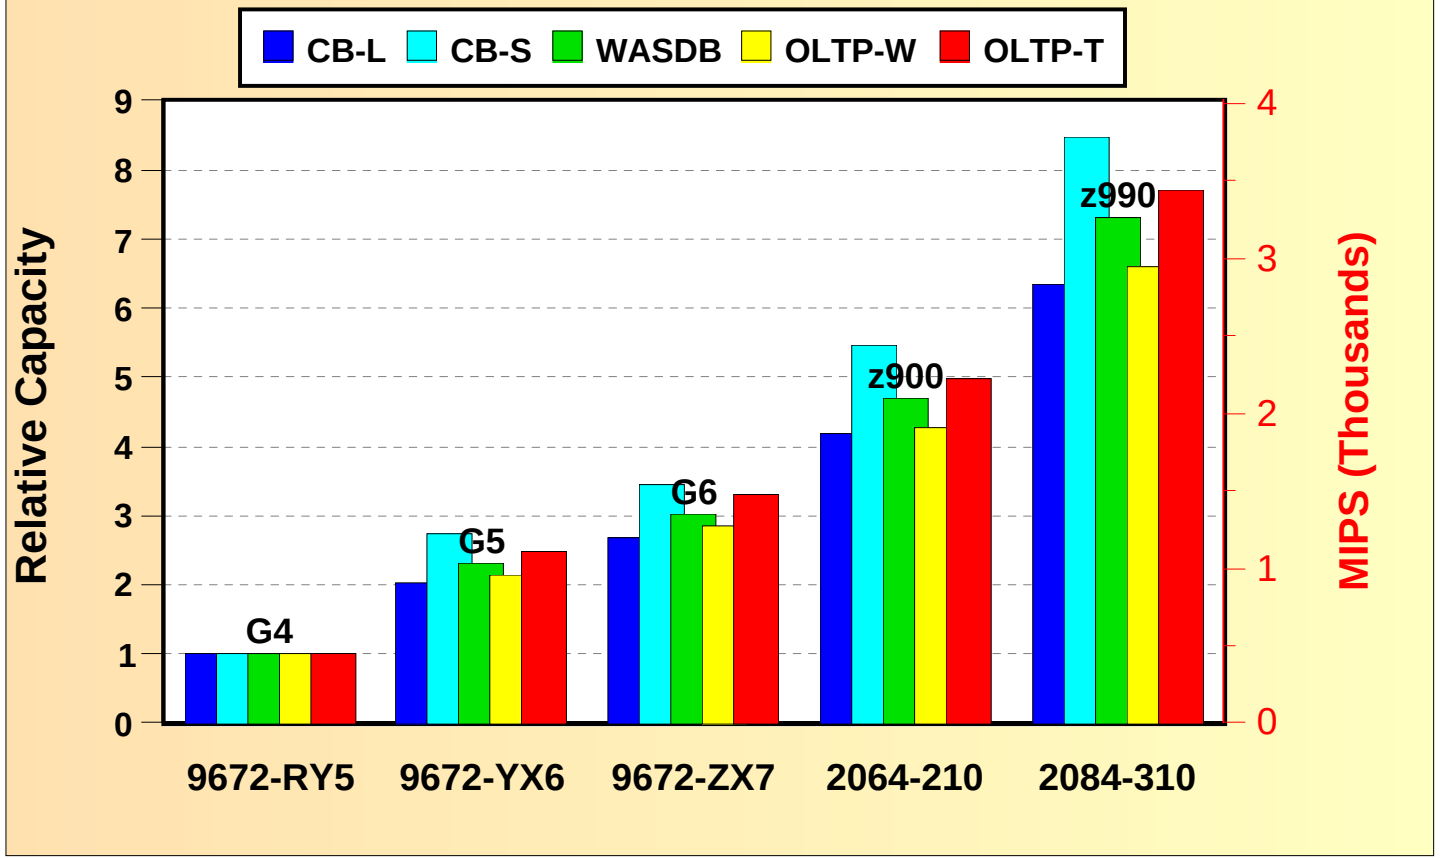

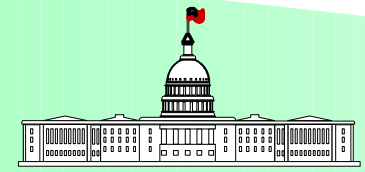

IBM Washington Systems Center

# LSPR

## Relative Capacity
### 1-Way Processors
### z/OS V1R4

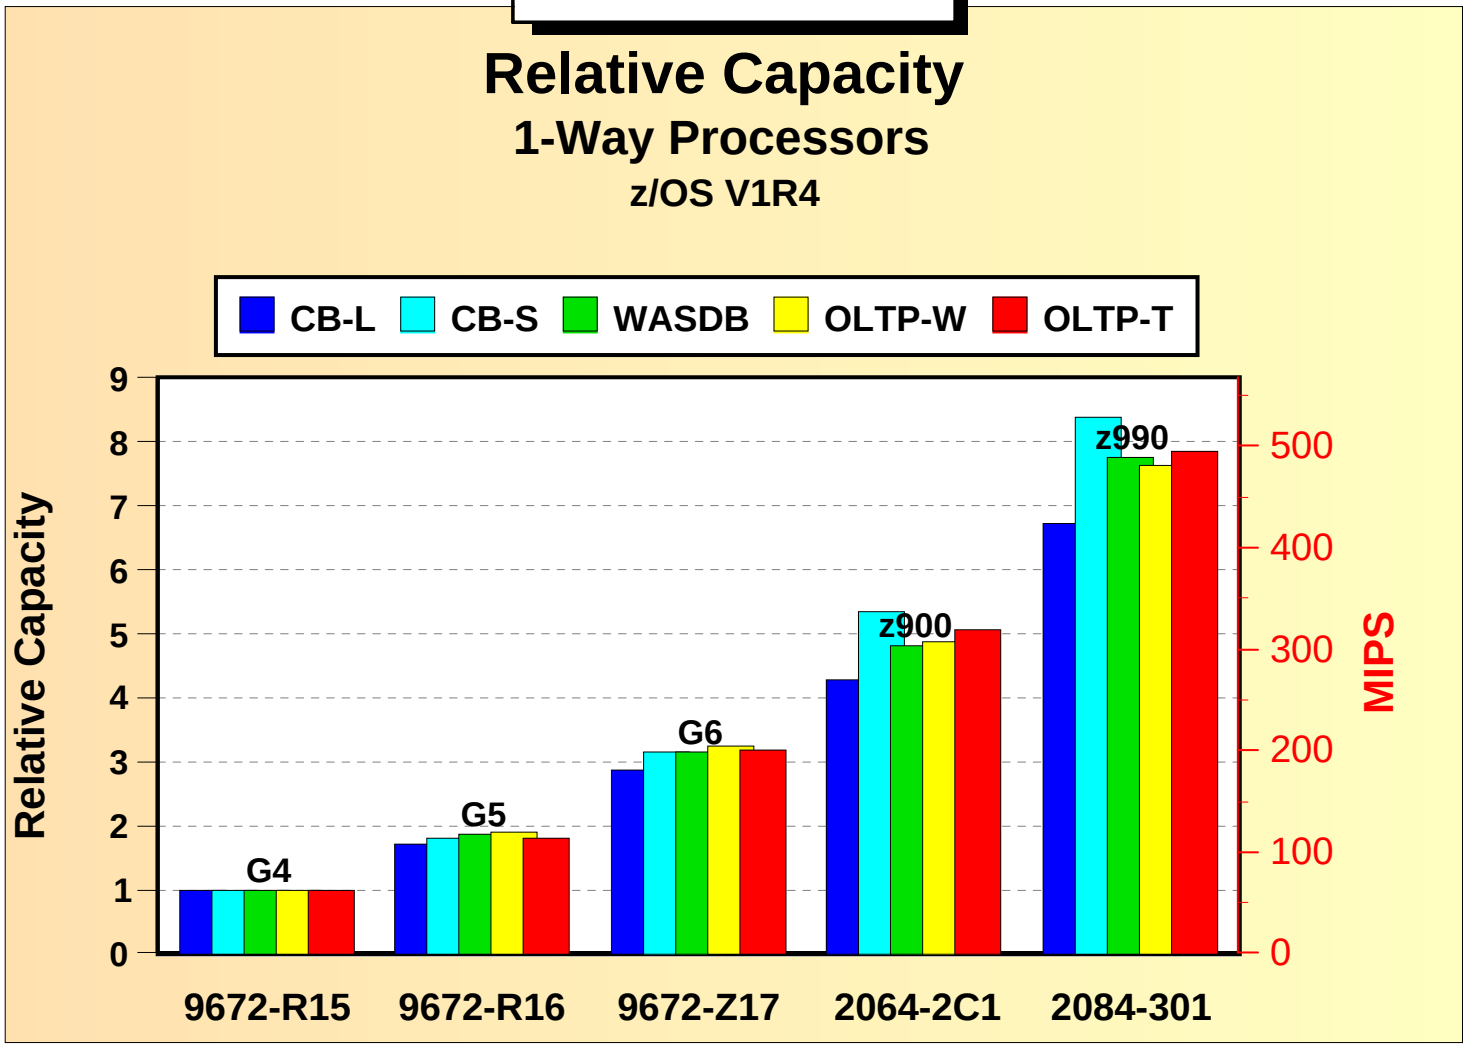

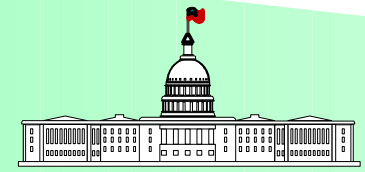

IBM Washington Systems Center

# LSPR

## Relative Capacity
### Slightly Increasing Capacity Progression
z/OS V1R4

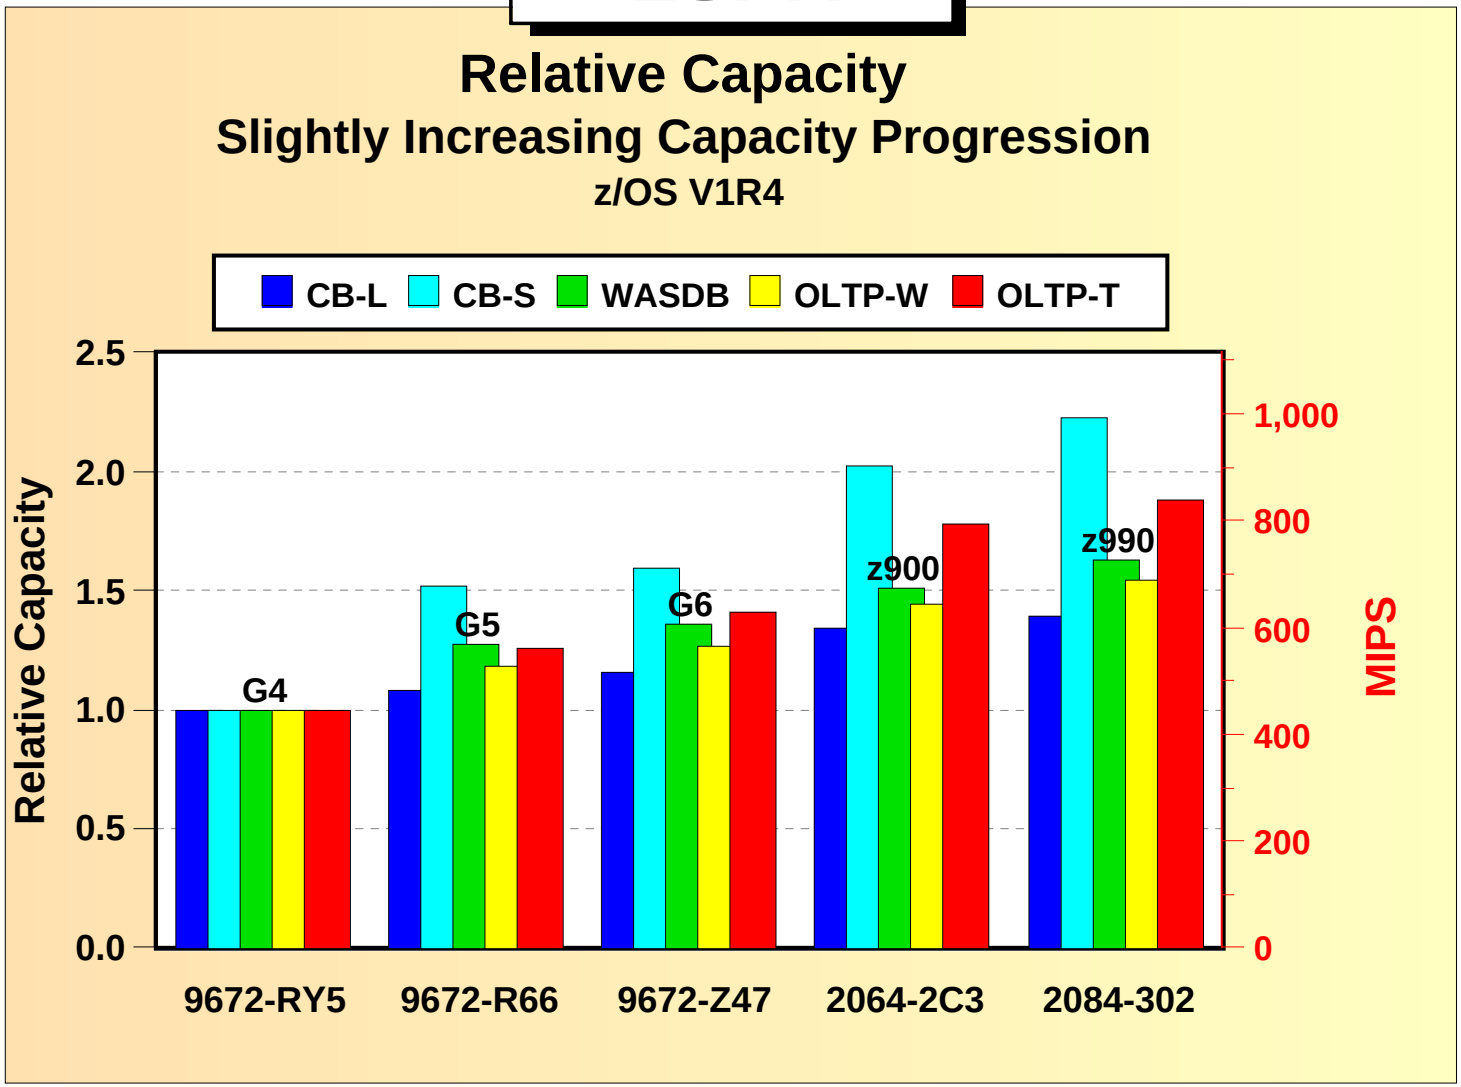

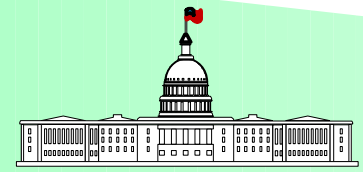



## Factors That Affect Processor Capacity Relationships

# Architecture

***zArchitecture*** *(including* ***S/390****)*

*The only architecture covered by LSPR*

# Hardware Design

*Each* ***zArchitecture*** *(****S/390****) processor family has its own unique internal design*

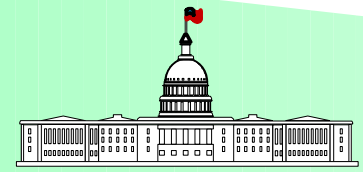



# Factors That Affect Processor Capacity Relationships

## Hardware

- Basic processor cycle time
  - \+ OPCODE efficiency
- High Speed Buffer (HSB) design
- Multi-Processor (N-way) efficiency
- Special hardware / microcode features

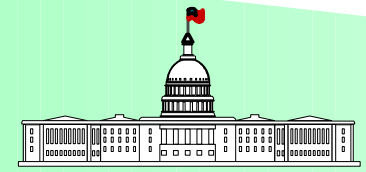



# Factors That Affect Processor Capacity Relationships

an external view

## Software

- CPU time per transaction / job
- % Supervisor or % Problem state
- I/O rate per second or per unit-of-work
- Working set size
- Etc . . .

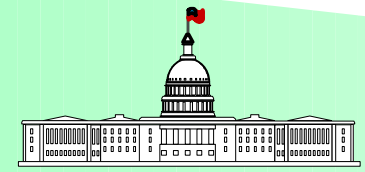



# Factors That Affect Processor Capacity Relationships


## Software

- Instruction set usage
  - OPCODE frequency and sequencing
- Storage reference patterns
- Dispatch rate
- Use of special features

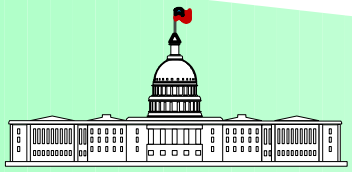



# Factors That Affect Processor Capacity Relationships

| Software | | Hardware |
|---|---|---|
| Instruction set usage | ➡ | OPCODE efficiency |
| Storage reference patterns | ➡ | HSB design |
| Dispatch rate | ➡ | N-way efficiency |
| Use of special features | ➡ | Special hardware / microcode |
| Varies by: SCP/workload environment | | Varies by: Processor family design |

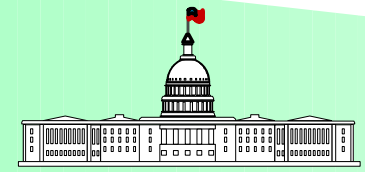



# LSPR

## Relative Capacity
### 1-Way Processors
### OS/390 V2R10

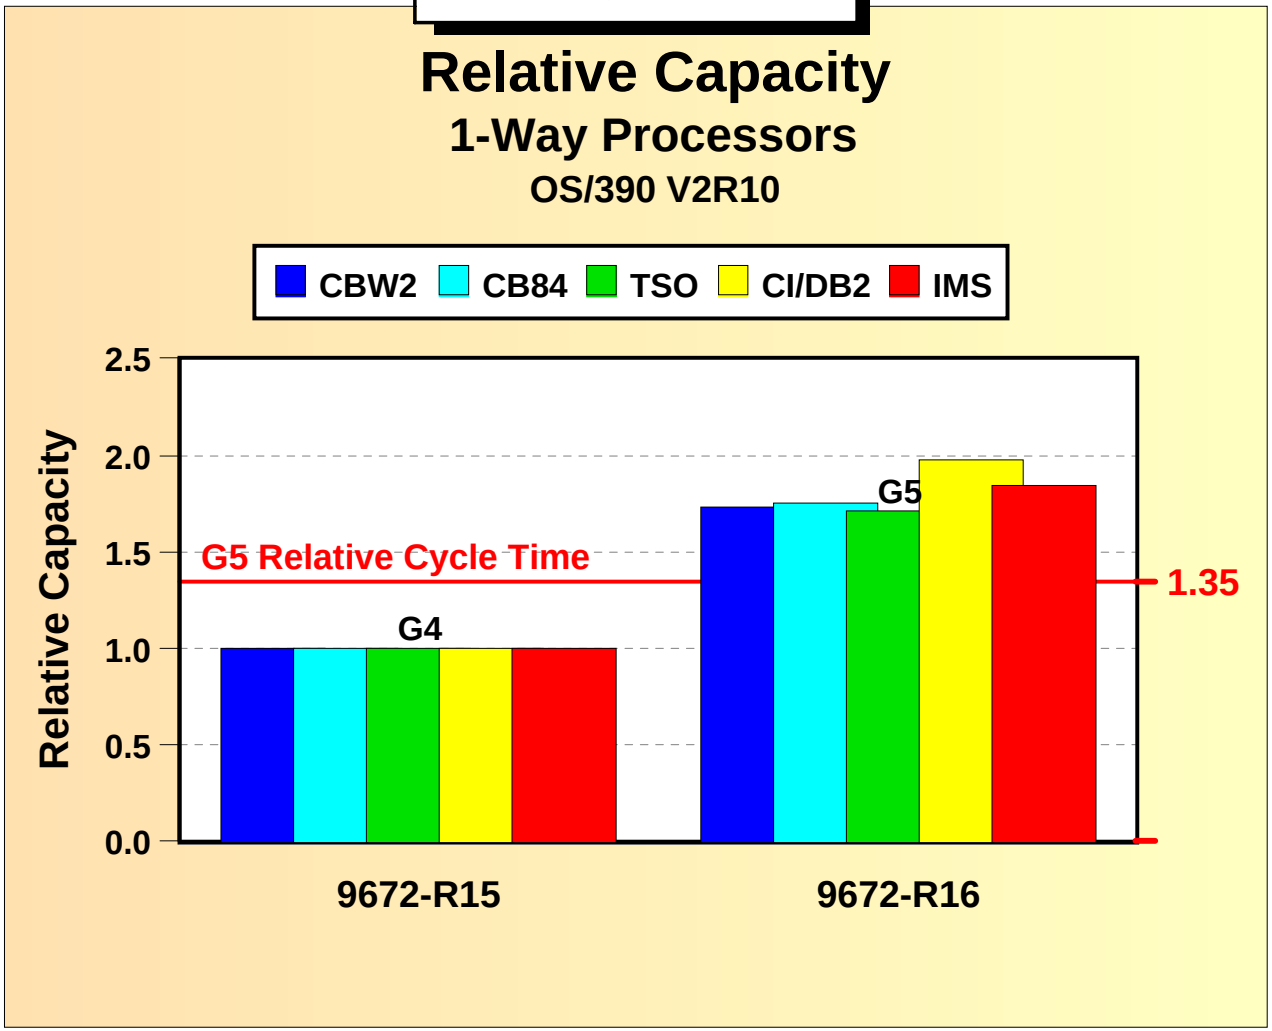

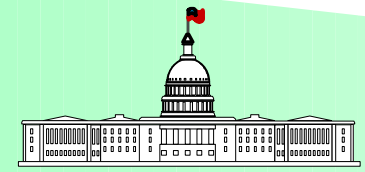

IBM Washington Systems Center

# LSPR

## Relative Capacity
### 10-Way Processors
### OS/390 V2R10

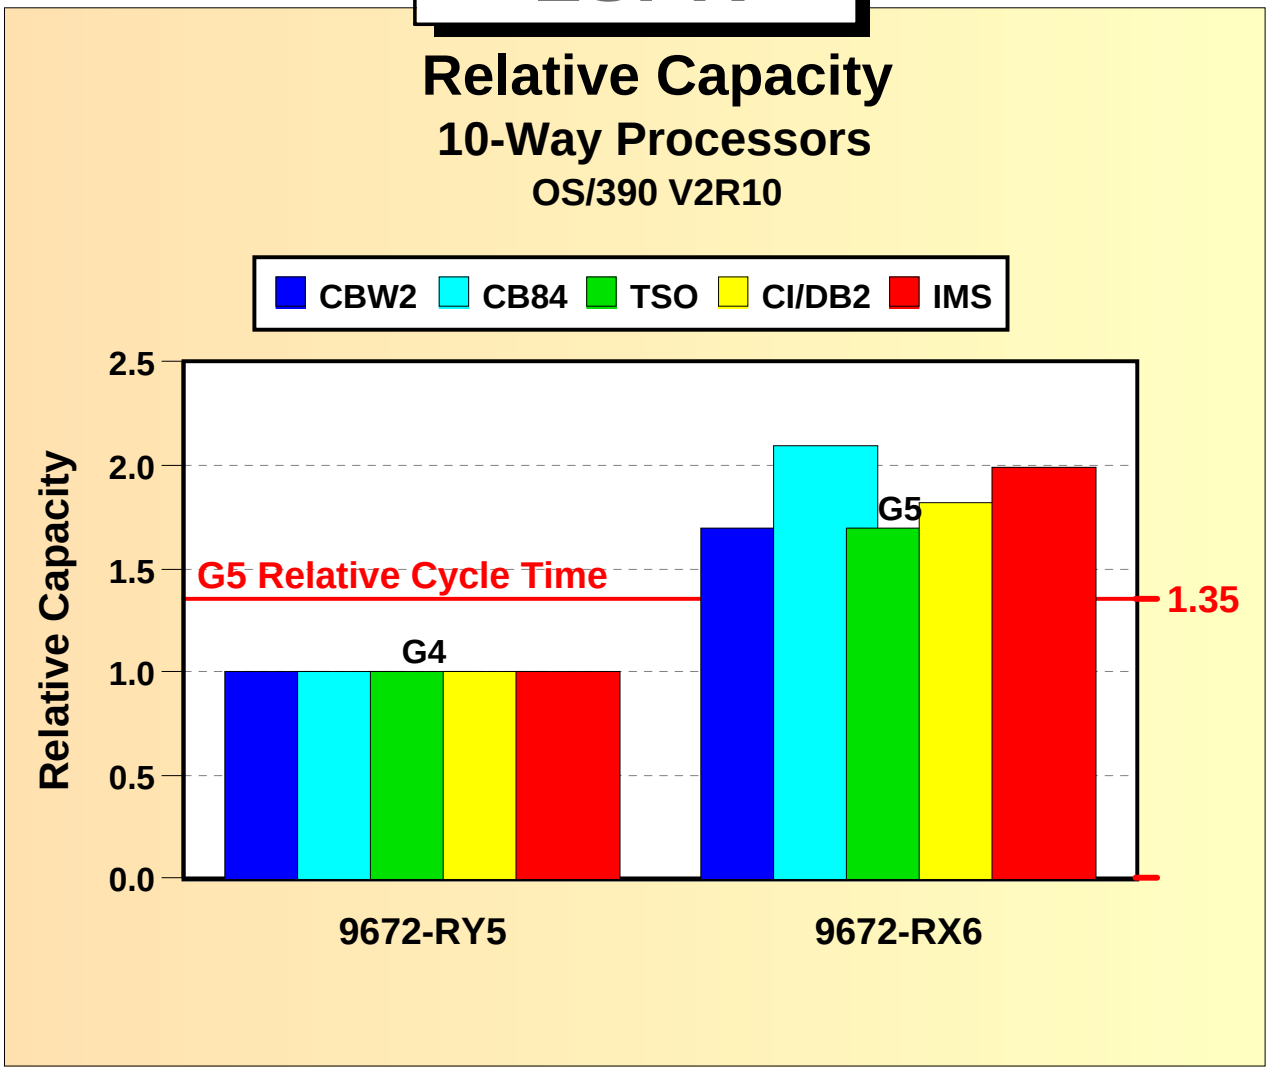

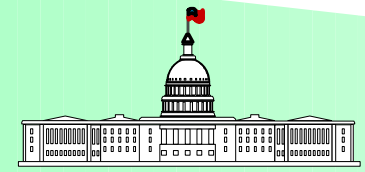



LSPR

# High Speed Buffer Effect

## 3090-180J vs 3090-180S

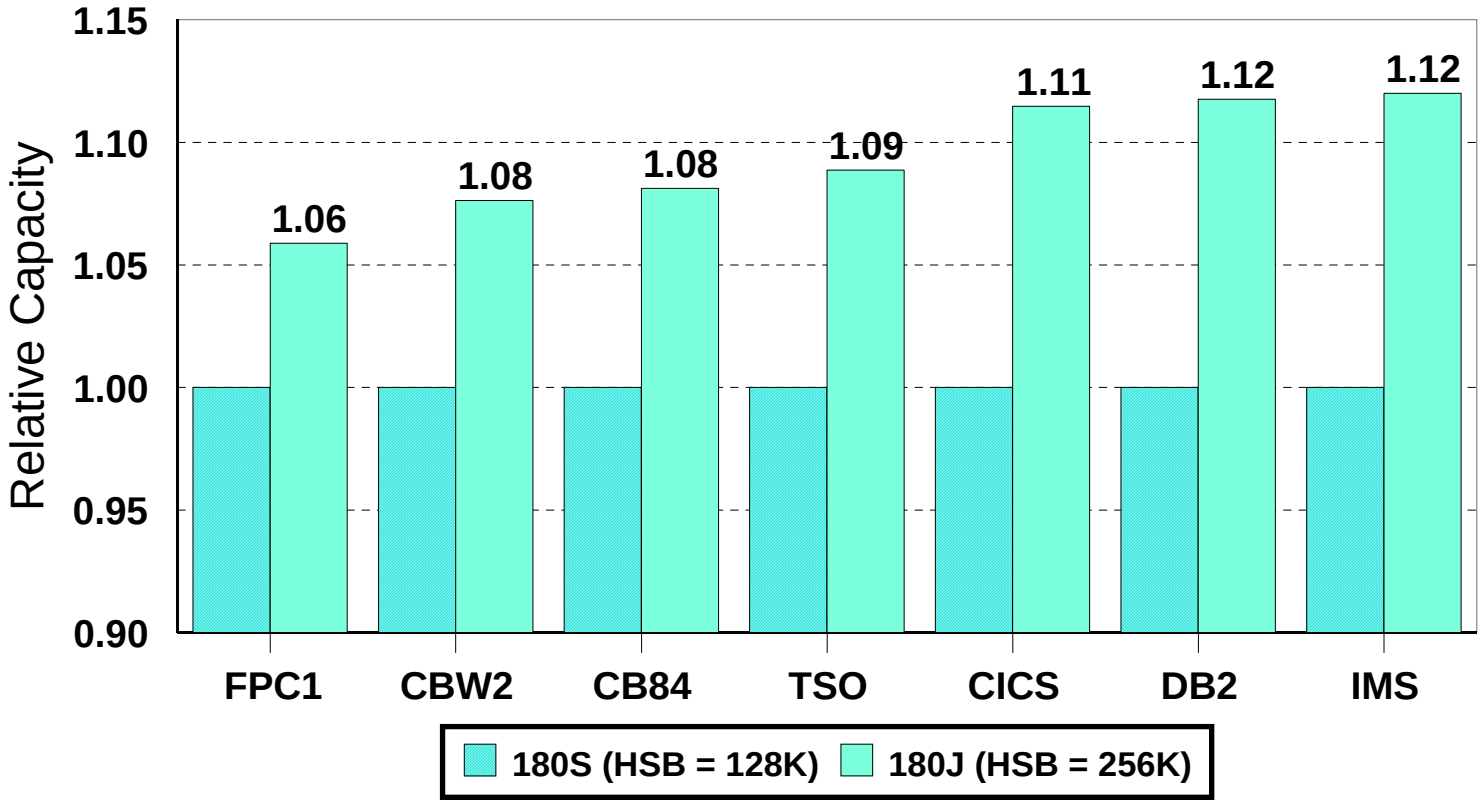

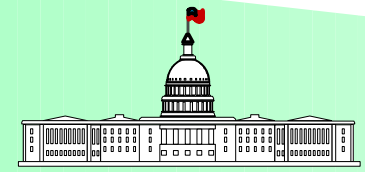



LSPR

# High Speed Buffer Effect

## 3090-600J vs 3090-600S

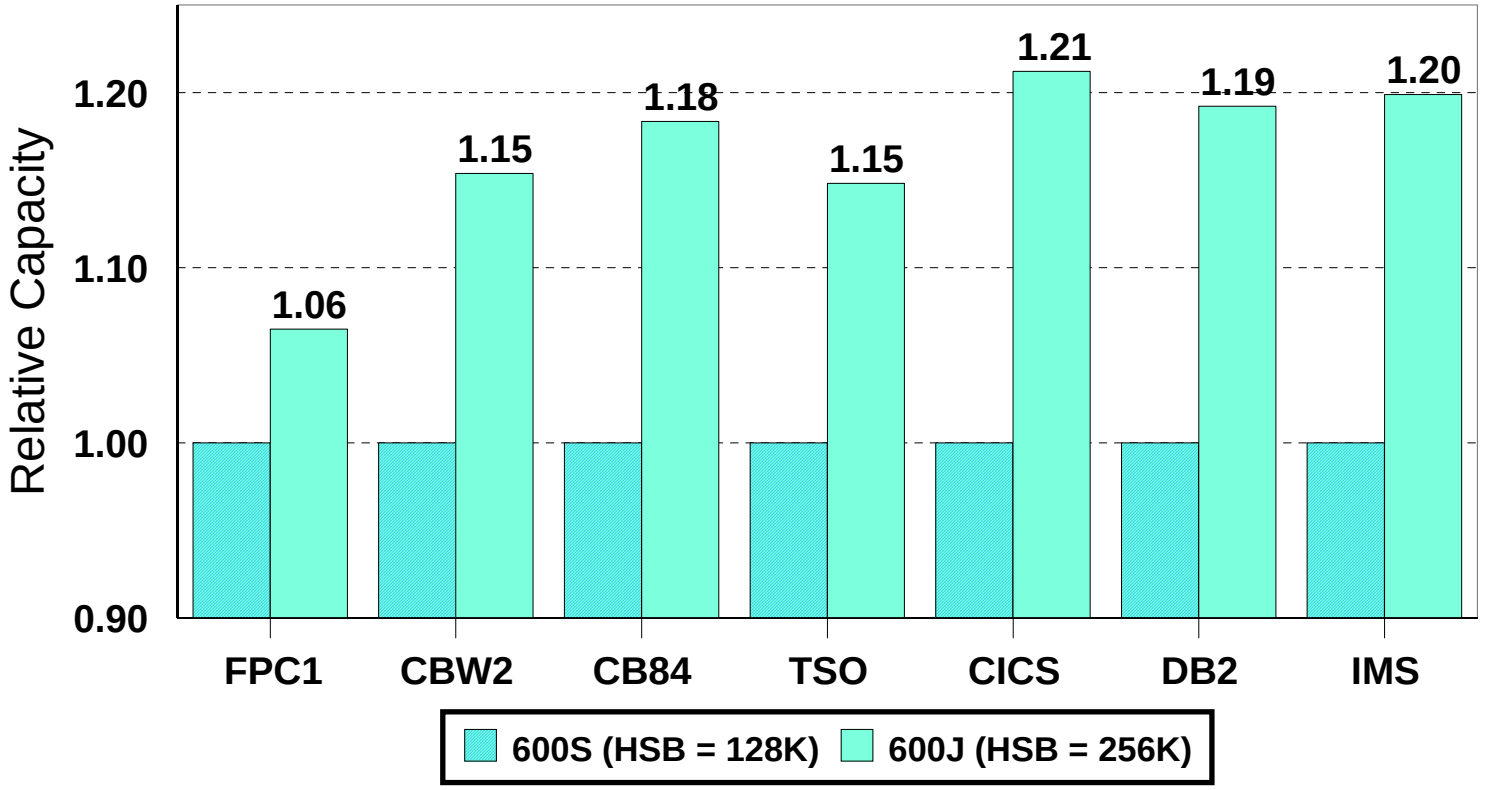

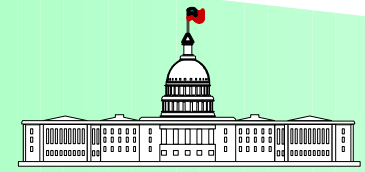



# LSPR

## N-Way Efficiency

9021 711-based (Bipolar)

OS/390 V1R1

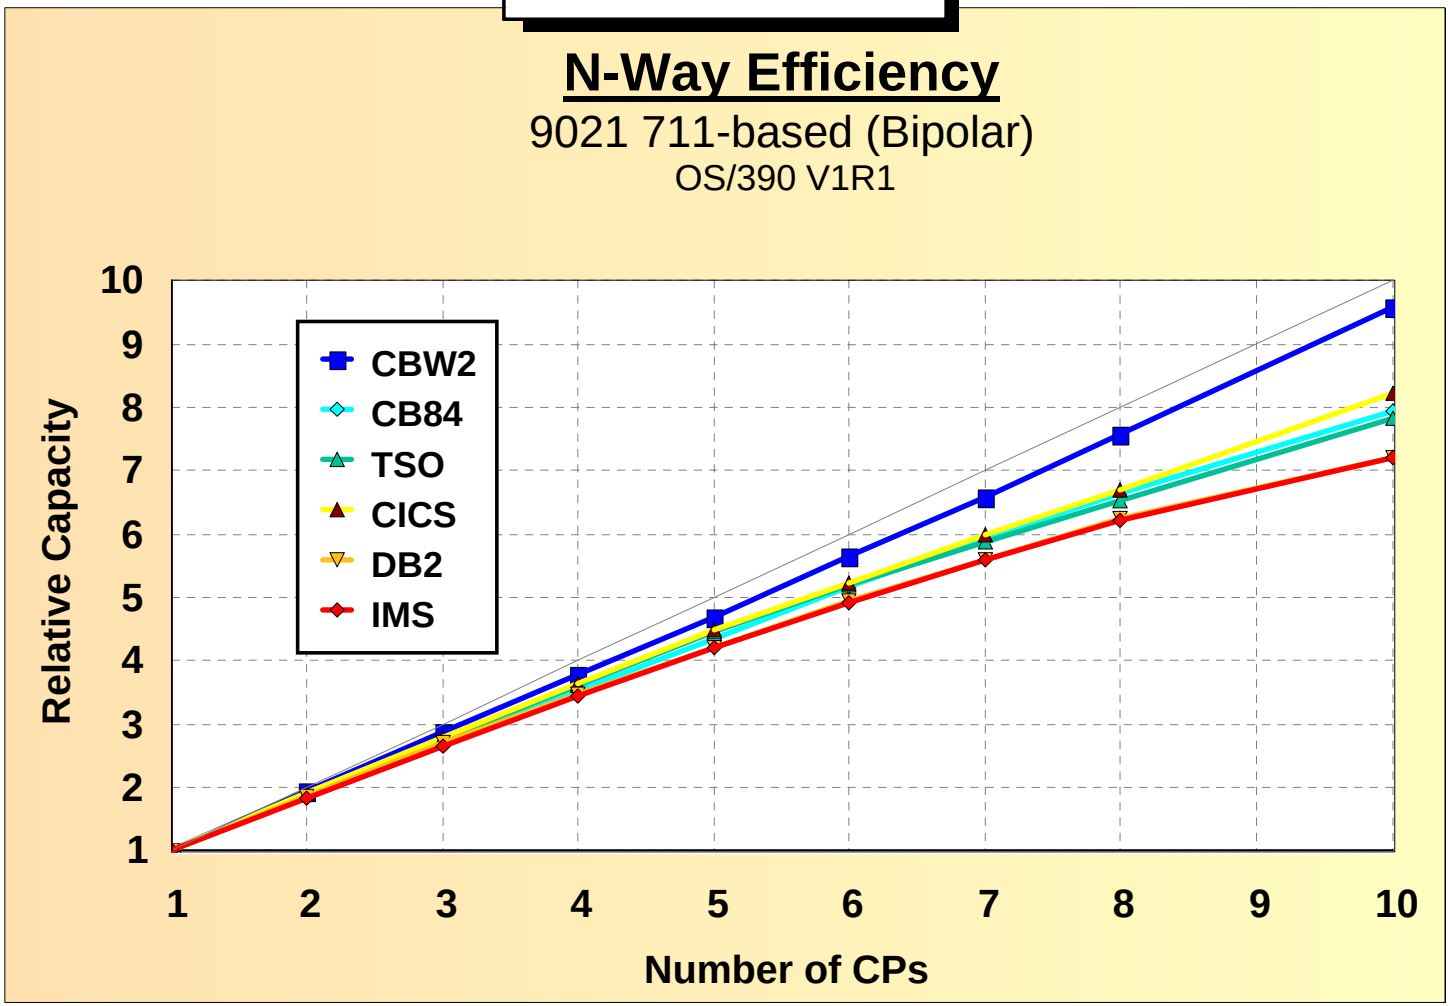

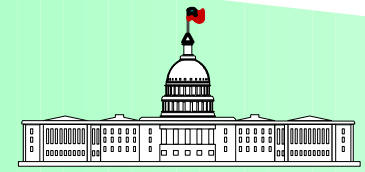



# LSPR

## N-Way Efficiency

G4 (CMOS)

OS/390 V1R1

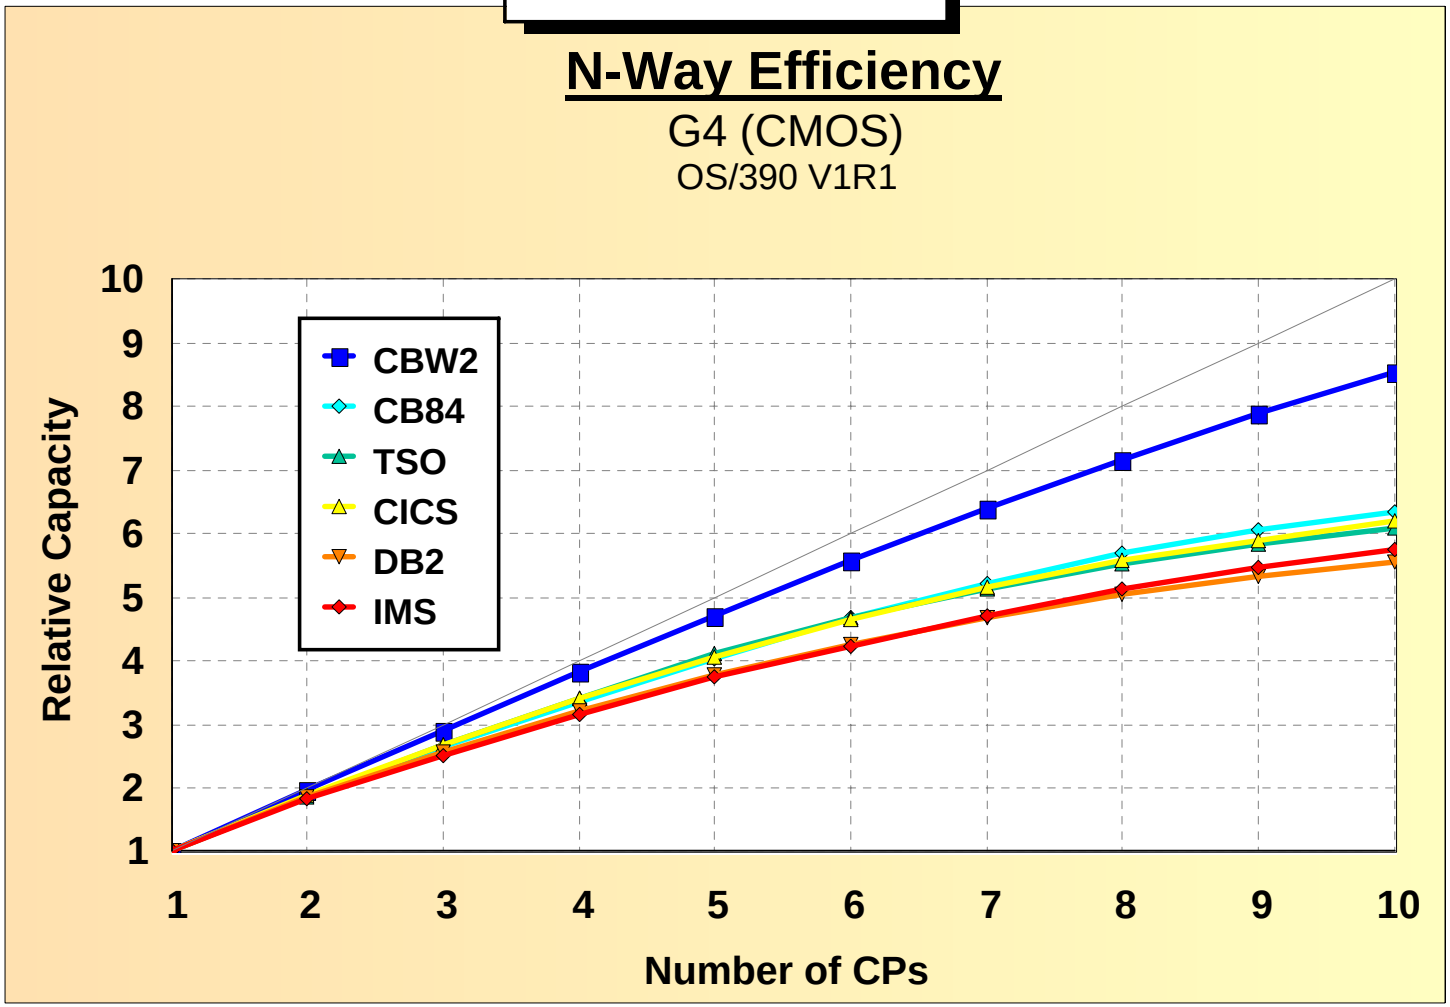

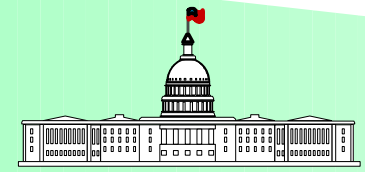



# LSPR

## N-Way Efficiency

G6 (CMOS)

OS/390 V2R4

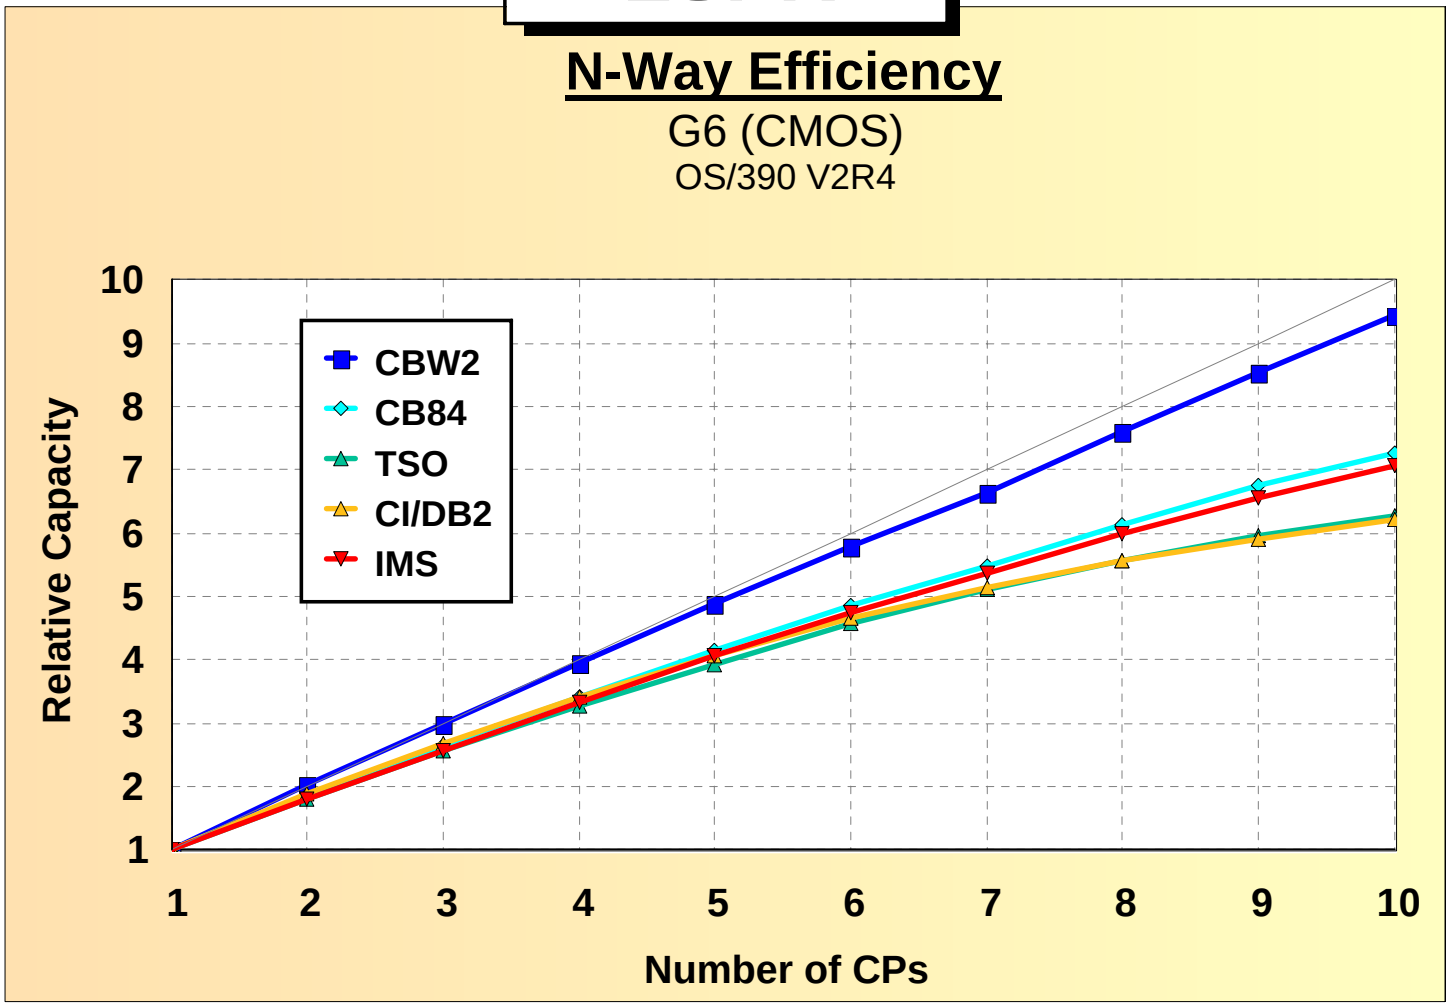

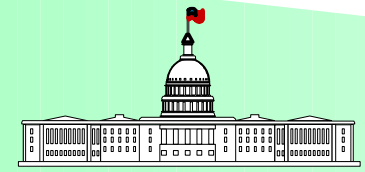



# LSPR

## N-Way Efficiency

z900 (CMOS)
OS/390 V2R10

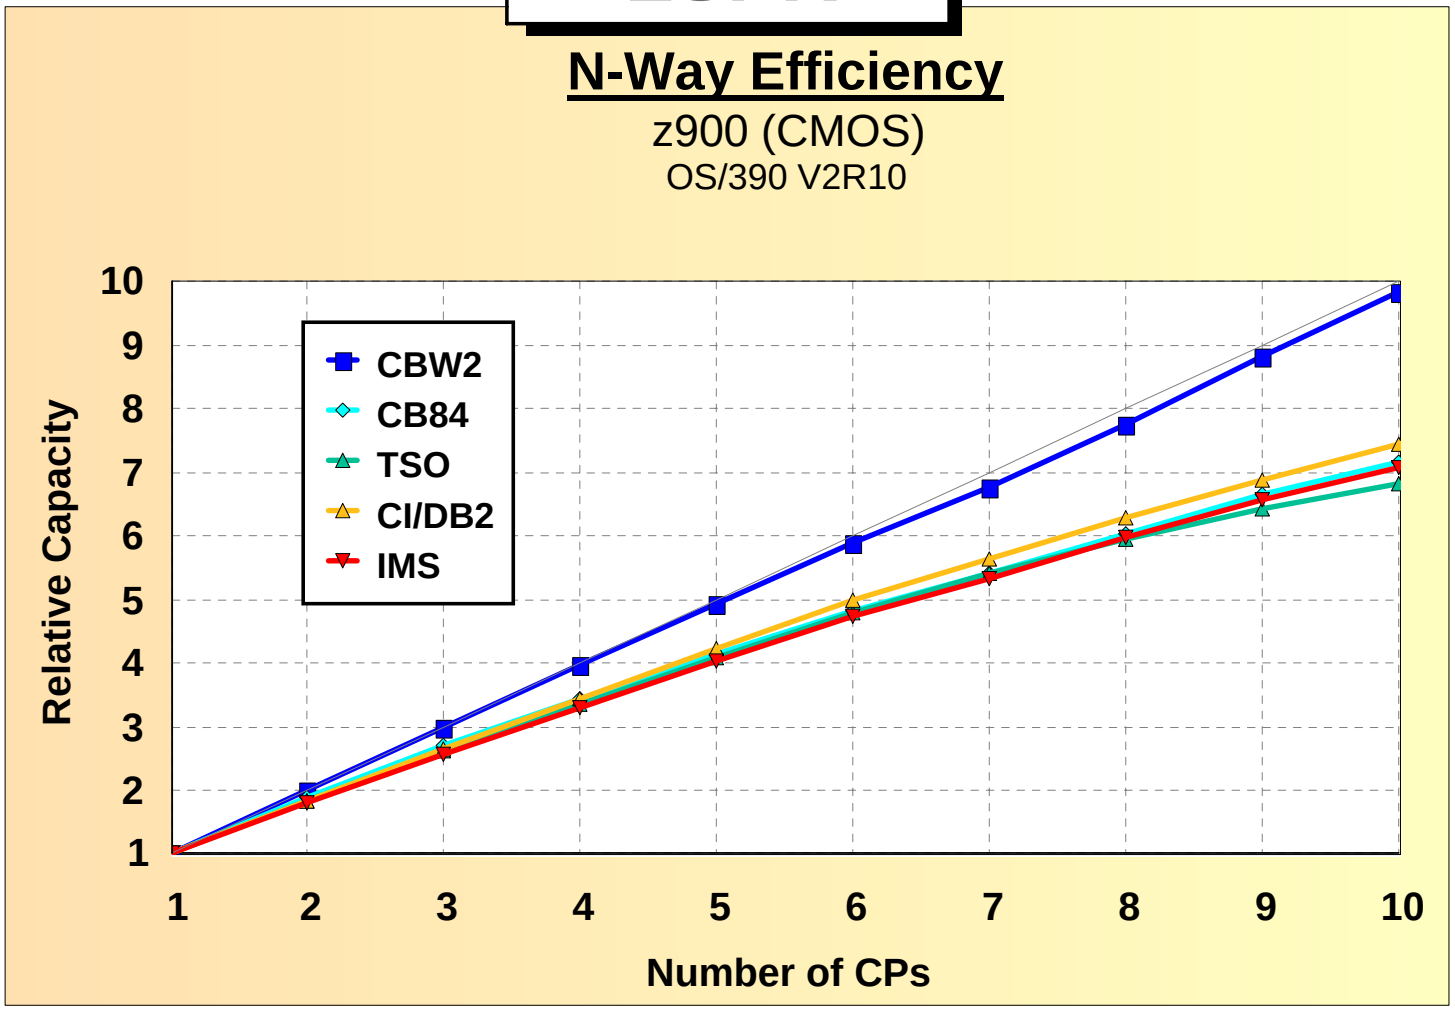

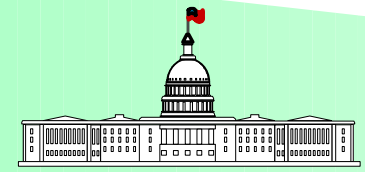



# LSPR

## N-Way Efficiency

z990 (CMOS)
z/OS V1R4

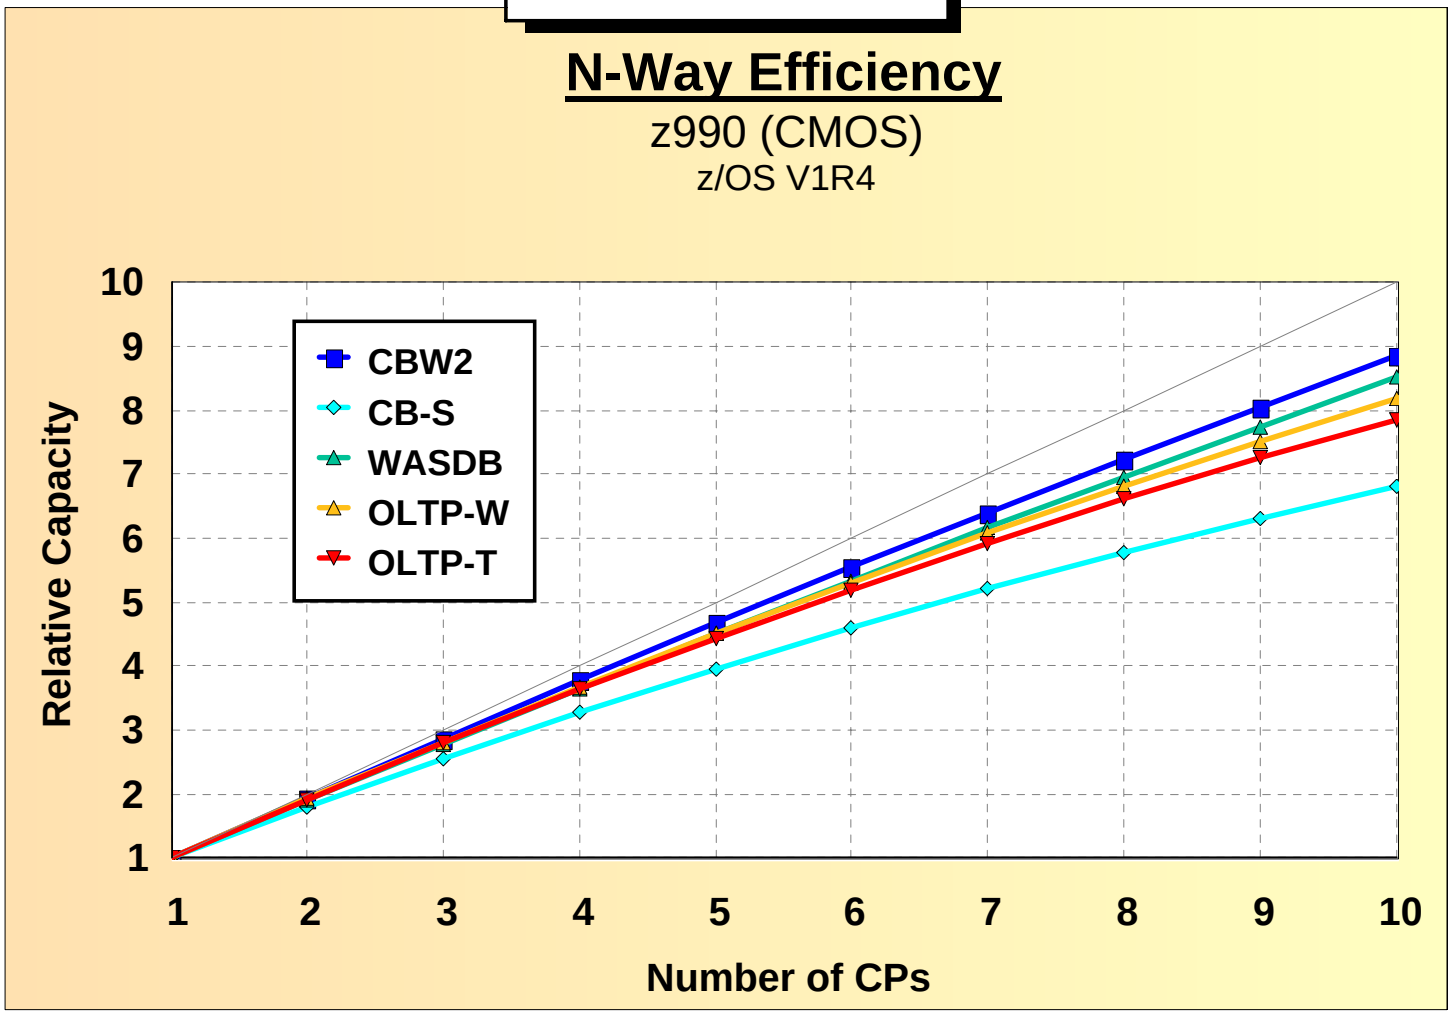

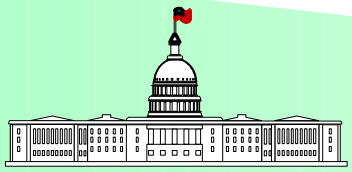



# Factors Affecting Processor Capacity Relationships

***Any given zArchitecture (or S/390) processor's capacity and performance potential is dependent on . . .***

- ***that processor's underlying design***
- ***how the SCP/workload environment interacts with that design***

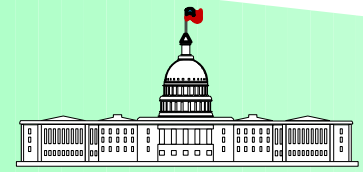



# Scope of DP Workloads

| BATCH | | ONLINE |
|---|---|---|
| JOB or JOB-STEP | Unit of Work | TRANSACTION |
| VERY LARGE | CPU Time per Dispatch | VERY SMALL |
| LOWER | I/O Intensity | HIGHER |
| INFREQUENT | Supervisor Services Requests | FREQUENT |
| VERY LOW | SUPERVISOR State Time | RELATIVELY HIGH |
| LOW | Dispatch Rate | HIGH |
| VERY EFFICIENT | N-Way Exploitation | LESS EFFICIENT |
| NOT DOMINANT | PRIVOPs and Branching | DOMINANT |
| PREDICTABLE | Storage Reference Patterns | RANDOM |
| VERY HIGH | High Speed Buffer Hit Ratio | LOWER |

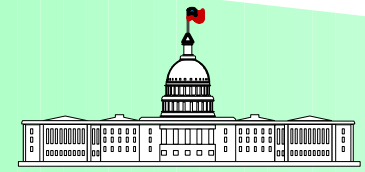



# Scope of LSPR Workloads

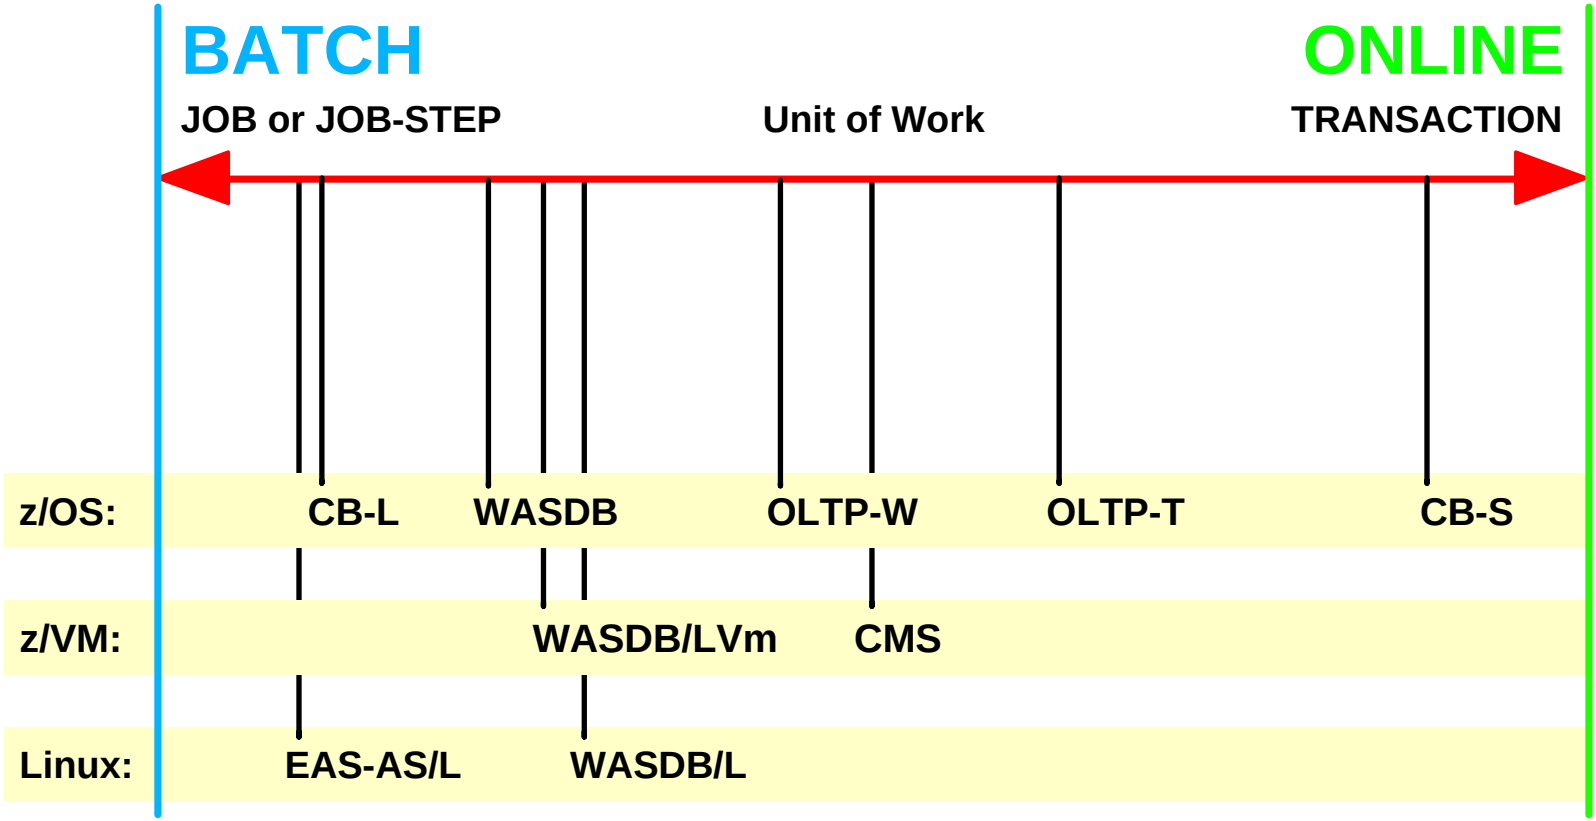

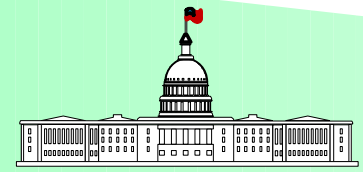



# How Do I Relate My Workload to LSPR ?

## Some general considerations

- Never focus solely on the LSPR workload that provides the best-case capacity relationship
- Always be aware of the LSPR workload with the worst-case capacity relationship
- For batch, don't use **CB-L** or **CB-S** independently without a specific reason; start with a mix of ¾ **CB-L** and ¼ **CB-S**
- For online workloads, factor in some **CB-L** when transactions are more than trivial (normally the case for today's environments)
- Capacity ratios are particularly sensitive to workload when **G1** - **G4** is involved as the "from" processor. Also, when **z990** is "to" processor.
- Request help from IBM to determine which LSPR workload primitives should be considered, and how to best interpret resulting capacity relationships
- Use the predefined workload mixes available in **zPCR** and **PCRW** tools

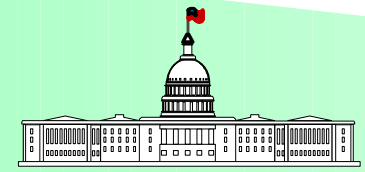



# Scope of LSPR Workloads

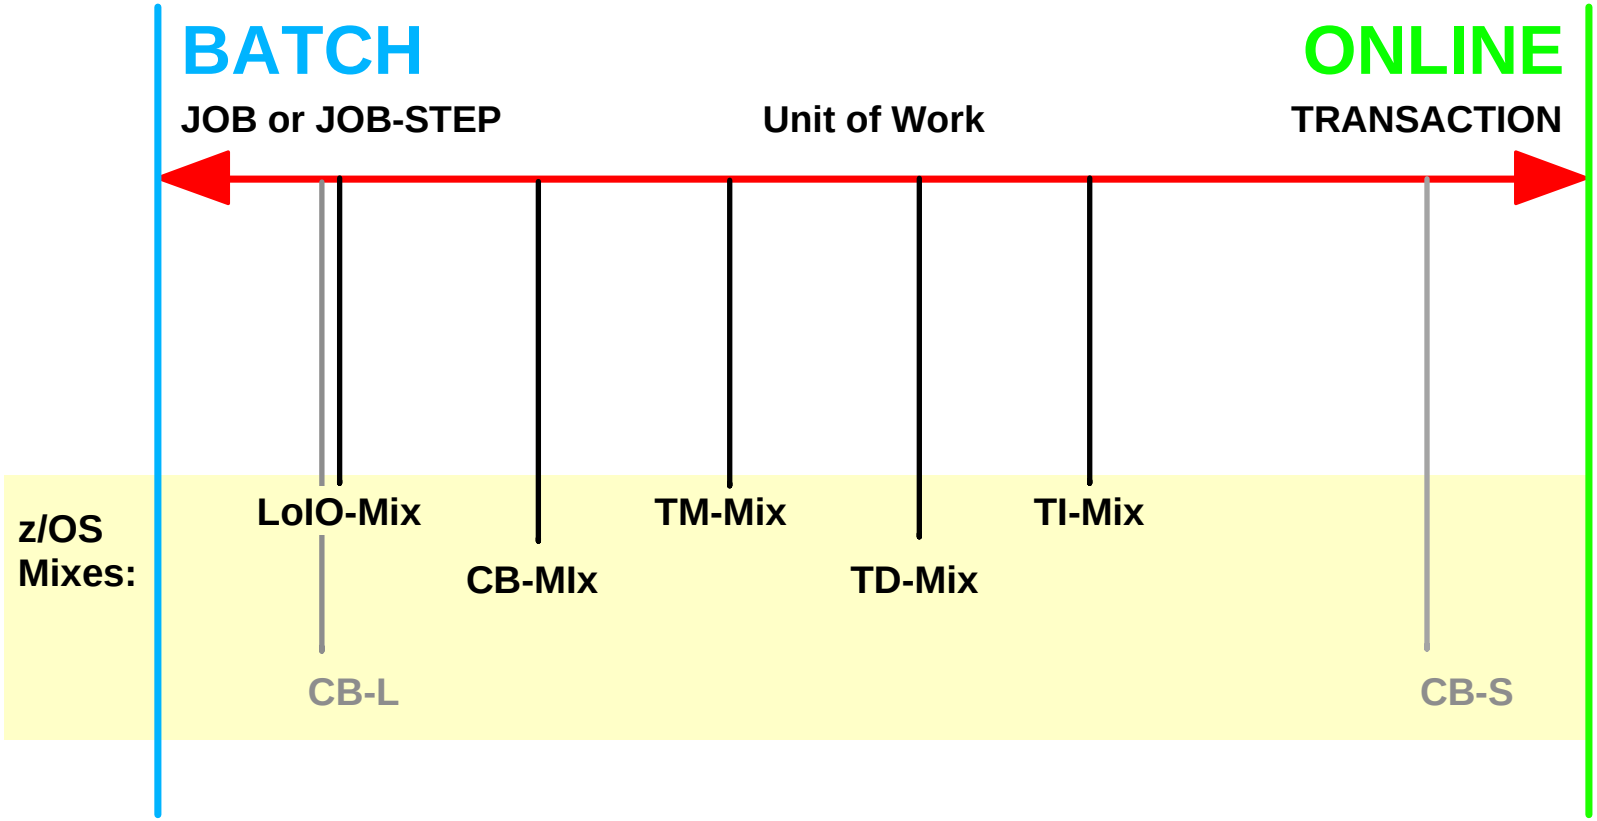

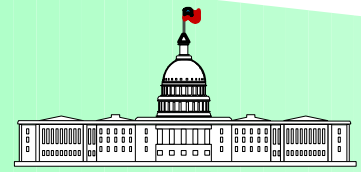



# General Capacity Observations

## Workload Component vs Overall Workload

### (An Example: "B" = 2.0 x "A")

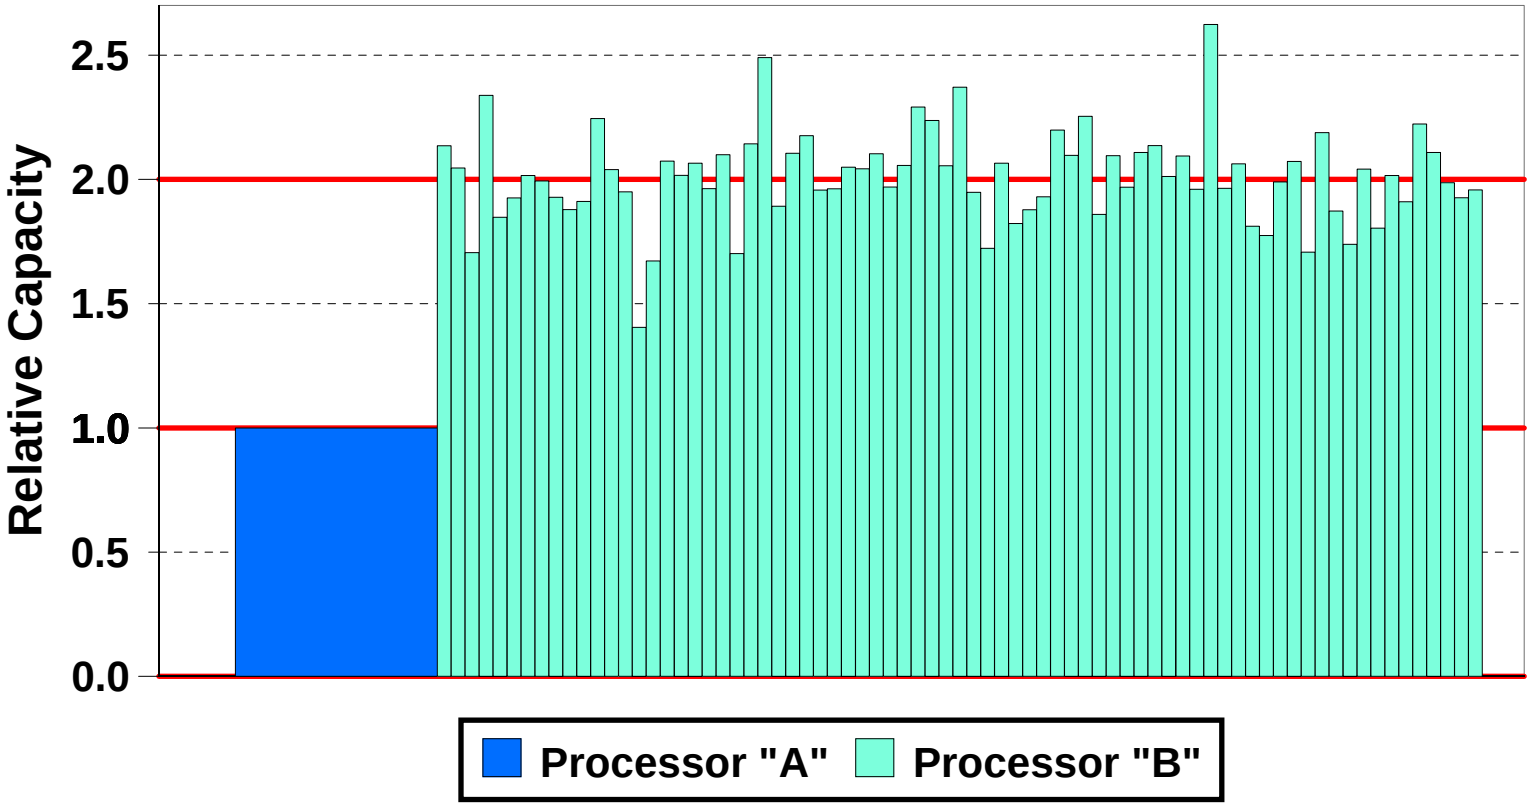

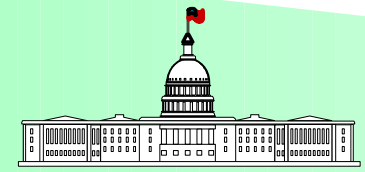



# General Capacity Observations

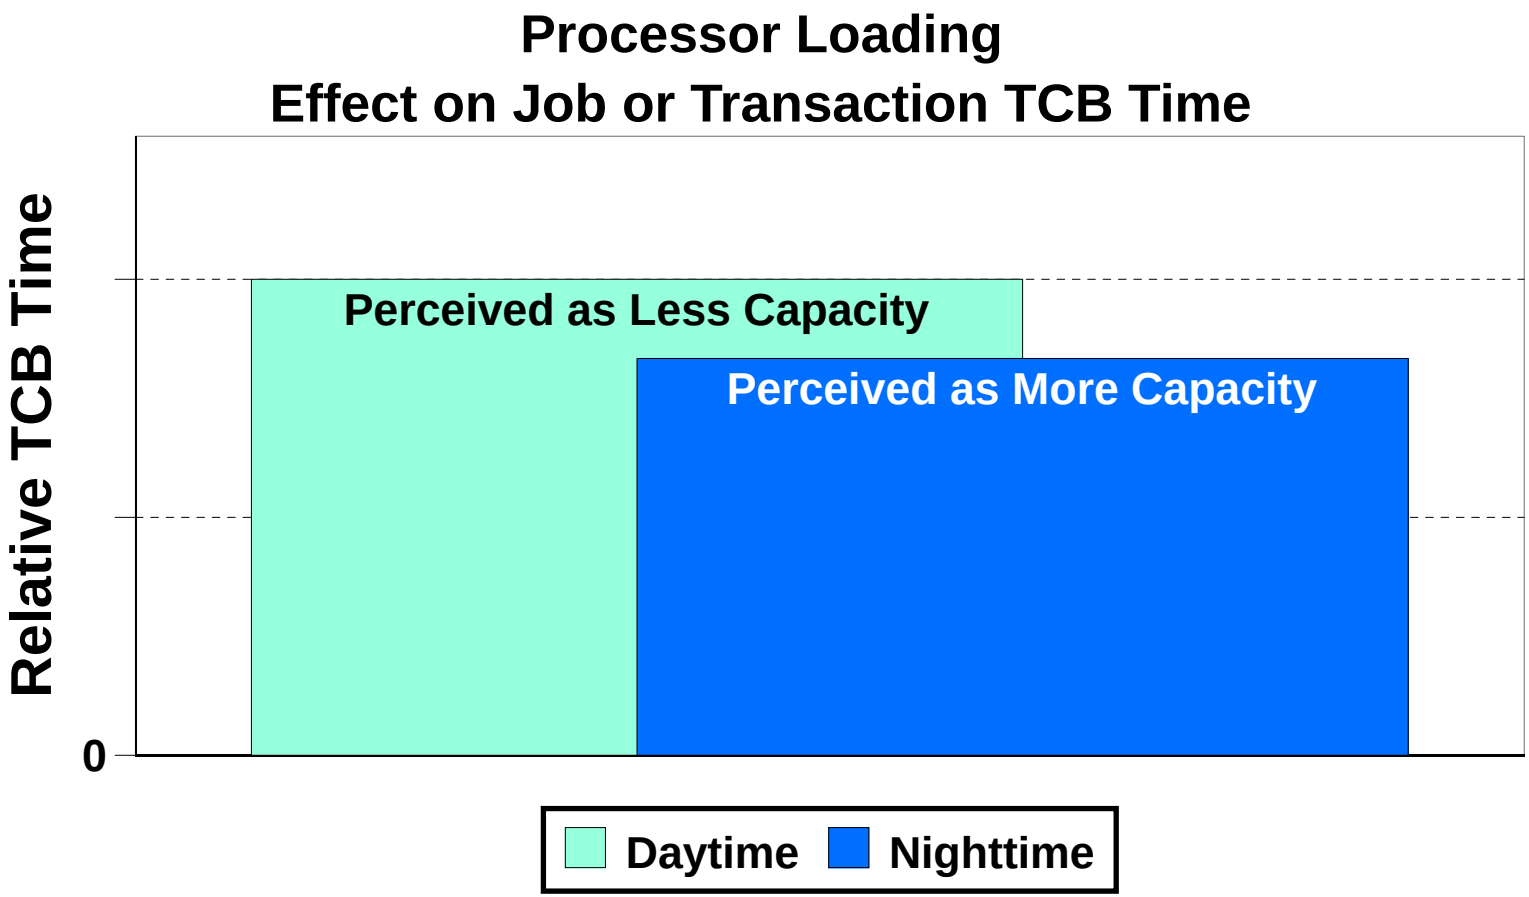

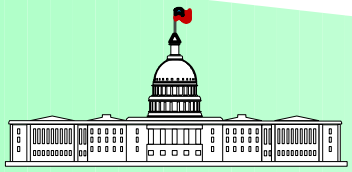

IBM Washington Systems Center

# Assessing Processor Capacity
## Accuracy *vs* Cost