

SHARE in Anaheim  
March 2011



# Manage and Monitor Your z/OS UNIX Services Environment Using WLM

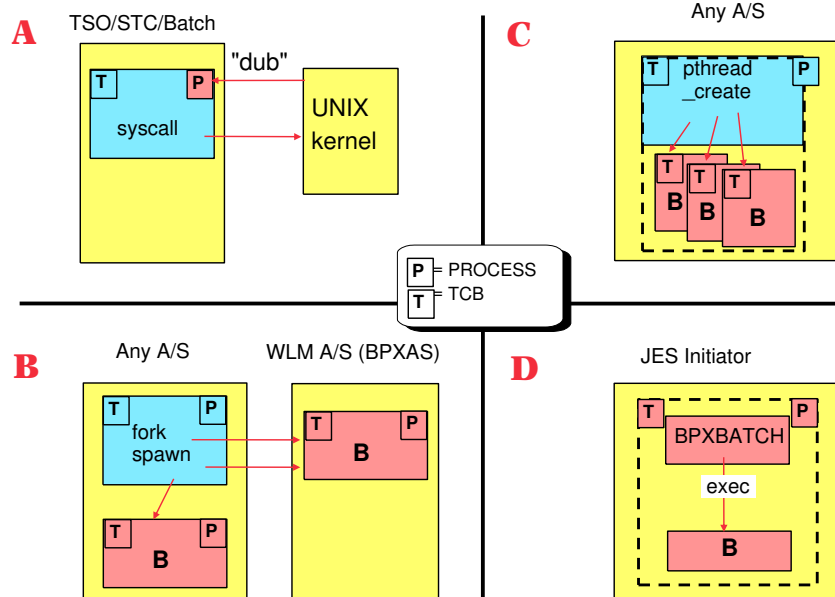
Session 8911

Glenn Anderson, IBM Technical Training



© 2011 IBM Corporation

## Running UNIX Applications in z/OS



## UNIX System Services WLM Usage

- Create and maintain

a pool of address spaces for fork and spawn function calls.

- Classify and manage

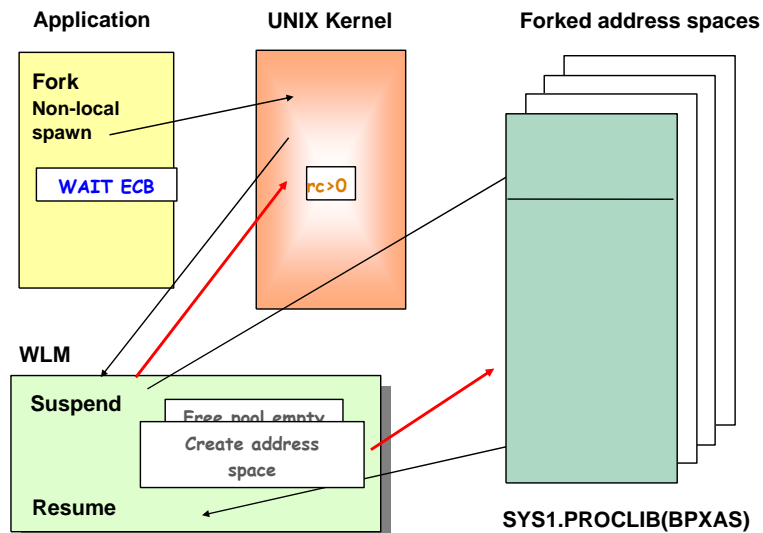
all the UNIX System Services address spaces.

- Propagate an enclave

across a fork, spawn or pthread\_create.

© Copyright IBM Corporation 2011

## UNIX System Services Pool Address Space



© Copyright IBM Corporation 2011

## Display BPXAS Initiators

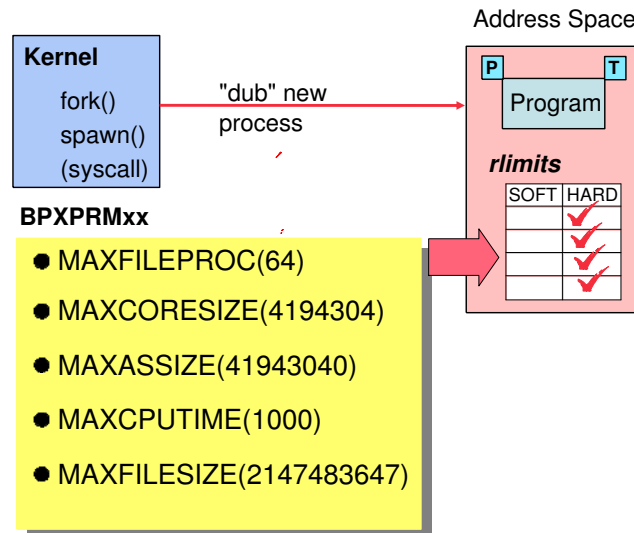
```

BPXAS  BPXAS  IEFPROC  IN  IO  A=005D  PER=NO  SMC=000
      PGN=N/A  DMN=N/A  AFF=NONE
      CT=011.337S  ET=24.19.32
      WUID=STC01224  USERID=OMVSUSR
      WKL=OMVS      SCL=UV40I2  P=1
      RGP=N/A      SRVR=NO  QSC=NO
      ADDR SPACE  ASTE=7F94E740

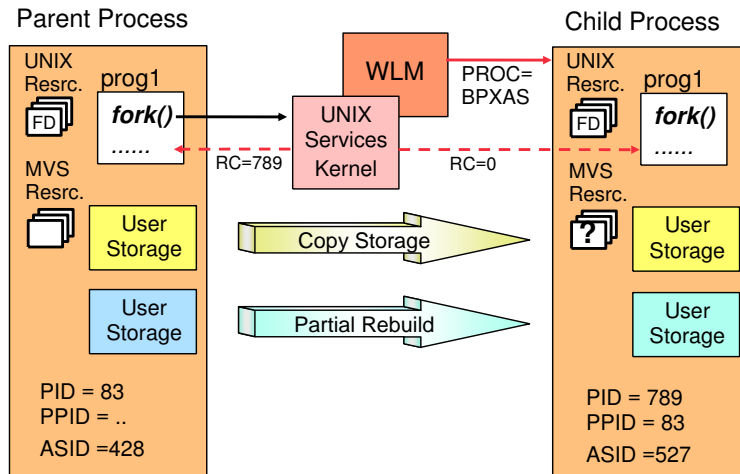
BPXAS  BPXAS  IEFPROC  OWT  IO  A=0060  PER=NO  SMC=000
      PGN=N/A  DMN=N/A  AFF=NONE
      CT=000.002S  ET=NOTAVAIL
      WUID=STC01355
      WKL=SYSTEM  SCL=SYSSTC  P=1
      RGP=N/A      SRVR=NO  QSC=NO
      ADDR SPACE  ASTE=7F94E800
    
```

Sleeping Address space

## Dubbing A New Process

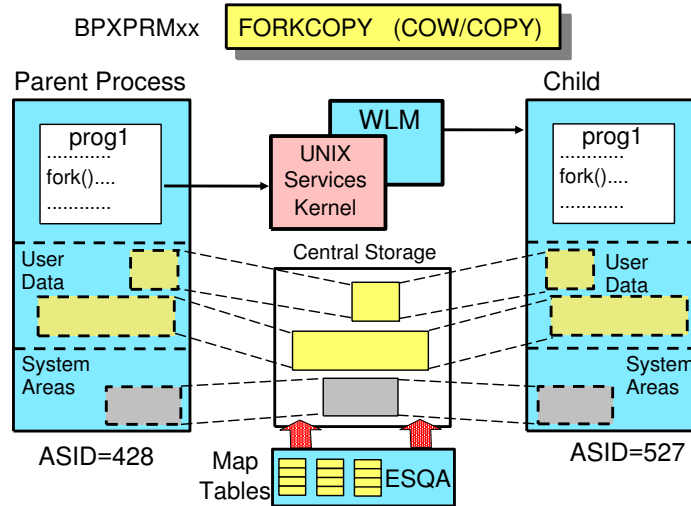


## Create a Process Using Fork



© Copyright IBM Corporation 2011

## Use Shared Pages for Fork Function



© Copyright IBM Corporation 2004

WORKLOAD ACTIVITY																			
z/OS V1R11		SYSPLEX UTCPLXJ8			START 07/22/2009-08.00.00			INTERVAL 000.30.00			MODE = GOAL		PAGE 1						
		RPT VERSION V1R11 RMF			END 07/22/2009-08.30.00														
POLICY ACTIVATION DATE/TIME 07/02/2009 06.41.24																			
----- SERVICE CLASS PERIODS																			
REPORT BY: POLICY=WLMPOLO1 WORKLOAD=BATCH SERVICE CLASS=BATDI SCR RESOURCE GROUP=*NONE PERIOD=1 IMPORTANCE=DISC																			
CRITICAL =NONE																			
-TRANSACTIONS- TRANS-TIME HHH.MM.SS.TTT --DASD I/O-- ---SERVICE--- SERVICE TIME ---APPL %--- --PROMOTED-- ----STORAGE----																			
AVG	3.69	ACTUAL	13.406	SSCHRT	16.0	I OC	418574K	CPU	1710.758	CP	95.14	BLK	0.000	AVG	865.56				
MPL	3.69	EXECUTION	9.736	RESP	1.0	CPU	475211K	SRB	1.596	AAPCP	0.00	ENQ	0.711	TOTAL	3197.37				
ENDED	6	QUEUED	1.048	CONN	0.7	MSO	104385K	RCT	0.085	IIPCP	0.00	CRM	0.000	SHARED	0.68				
END/S	0.00	R/S AFFIN	2.097	DISC	0.1	SRB	443384	IIT	0.180			LCK	11.580						
#SWAPS	44	INELIGIBLE	524	Q-PEND	0.2	TOT	998613K	HST	0.000	AAP	0.00			-PAGE-IN RATES-					
EXCTD	0	CONVERSION	0	I OSQ	0.0	/SEC	554779	AAP	0.000	IIP	0.00			SINGLE	0.0				
AVG ENC	0.00	STD DEV	14.527					IIP	0.000					BLOCK	0.0				
REM ENC	0.00					ABSRPTN	150K							SHARED	0.0				
MS ENC	0.00					TRX SERV	150K							HSP	0.0				
GOAL: DISCRETIONARY																			
RESPONSE TIME EX PERF AVG --EXEC USING%-- ----- EXEC DELAYS % ----- -USING%-- --- DELAY % --- %																			
SYSTEM		VEL% INDX	ADRSP	CPU	AAP	IIP	I/O	TOT	CPU	Q	I/O		CRY	CNT	UNK	IDL	CRY	CNT	QUI
								MPL											
JCO	--N/A--	64.1	3.7	31	0.0	0.0	0.5	18	17	0.8	0.1		30	0.0	50	0.8	0.0	0.0	0.0

-----STORAGE-----

AVG 865.56 → Avg number of storage frames allocated to active ASIDs

TOTAL 3197.37

SHARED 0.68 → Total number of shared storage pages allocated to ASIDs

-PAGE-IN RATES-

SINGLE 0.0

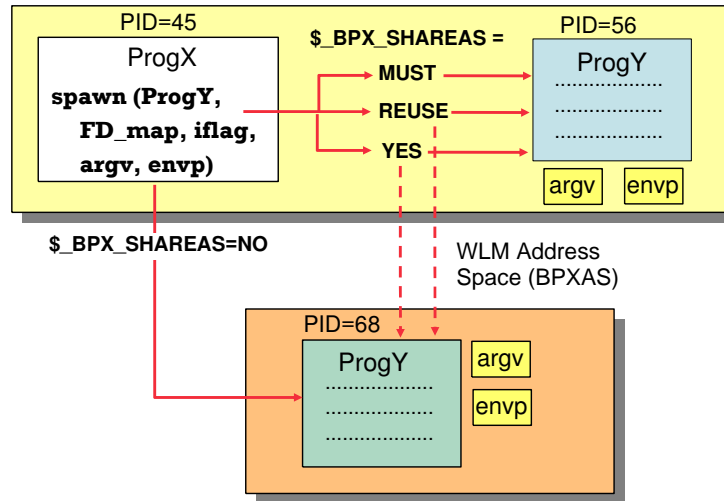
BLOCK 0.0

SHARED 0.0

HSP 0.0

## Create a Process Using Spawn



Parent Process Address Space




© Copyright IBM Corporation 2011

## Starting a Program With Exec

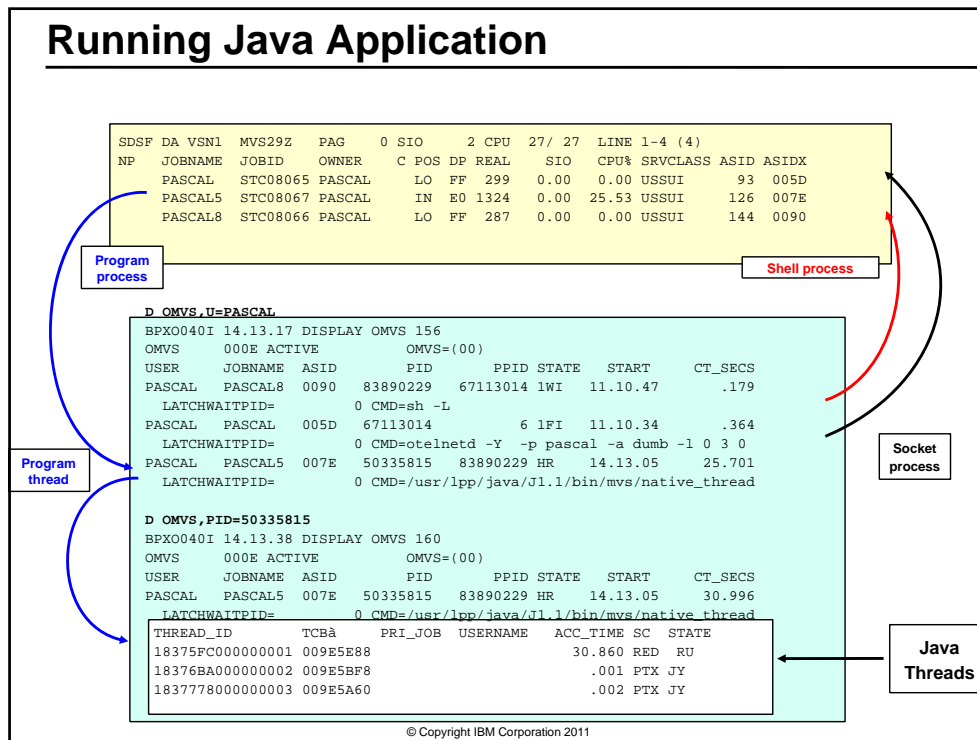
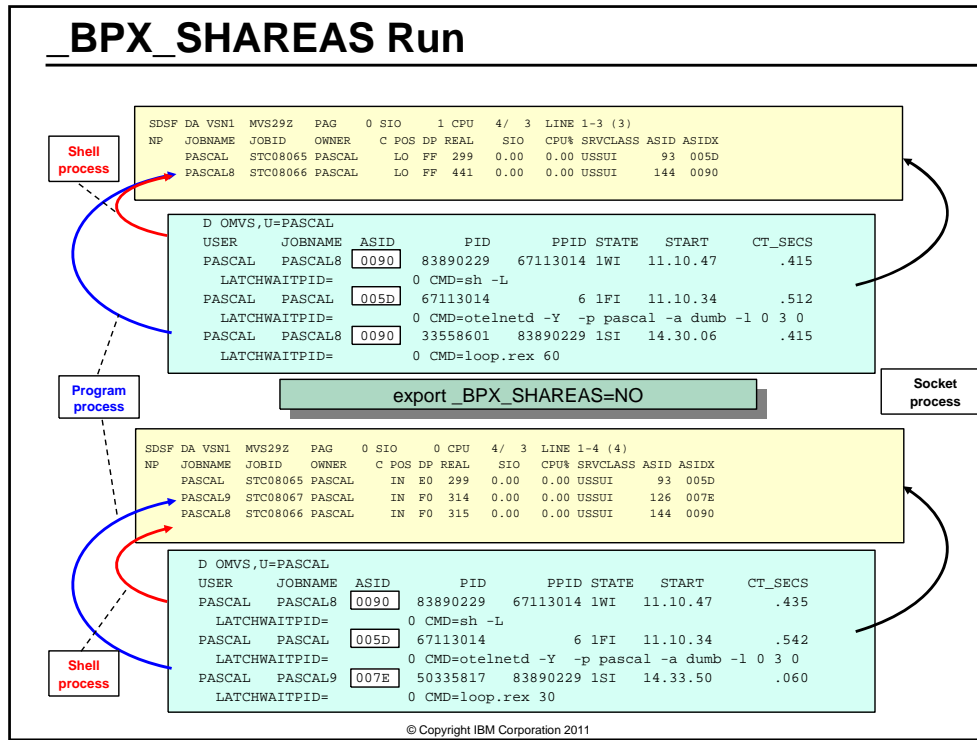
Process before EXEC



Process after EXEC



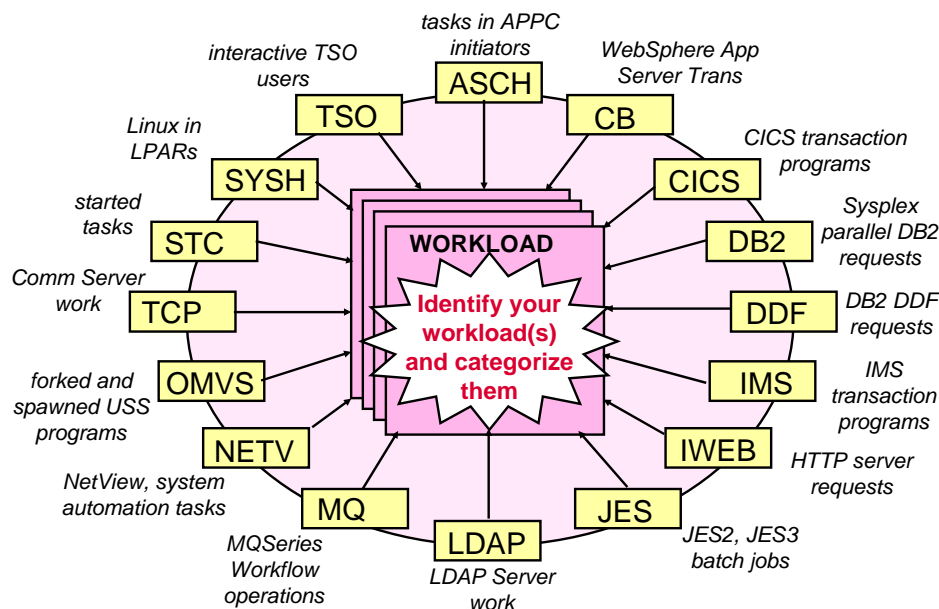
© Copyright IBM Corporation 2011



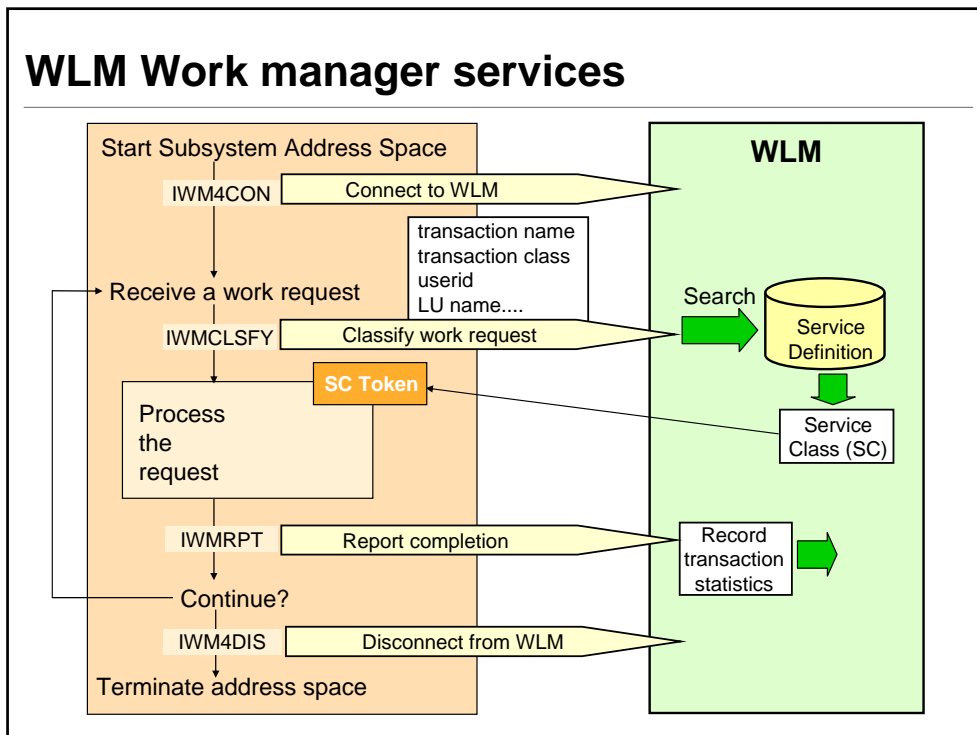
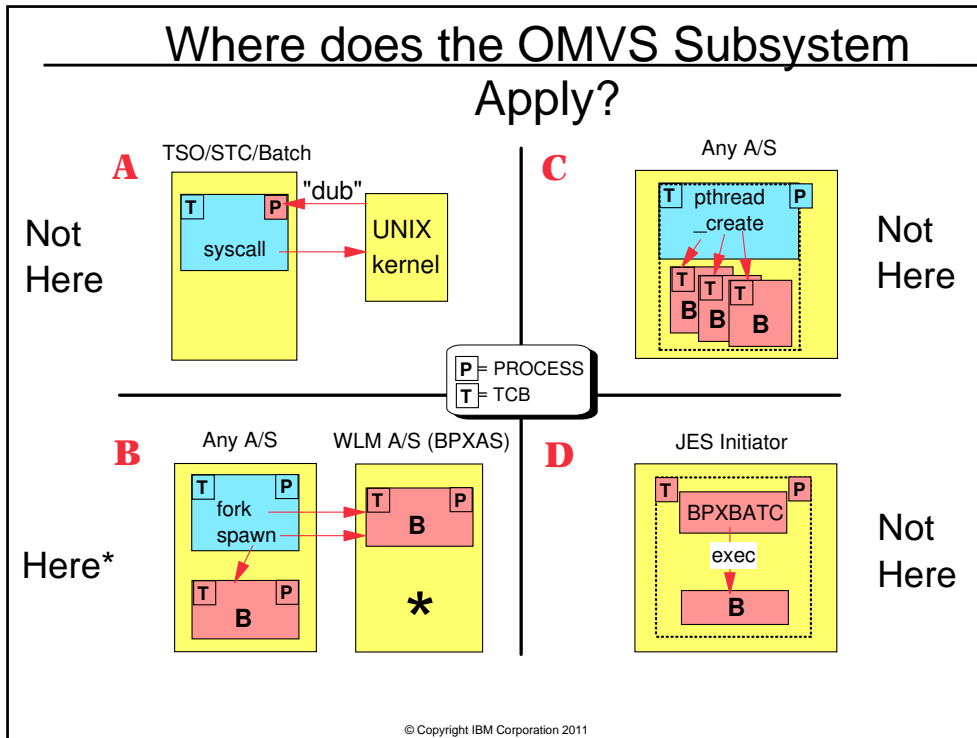
# What is a WLM Transaction?

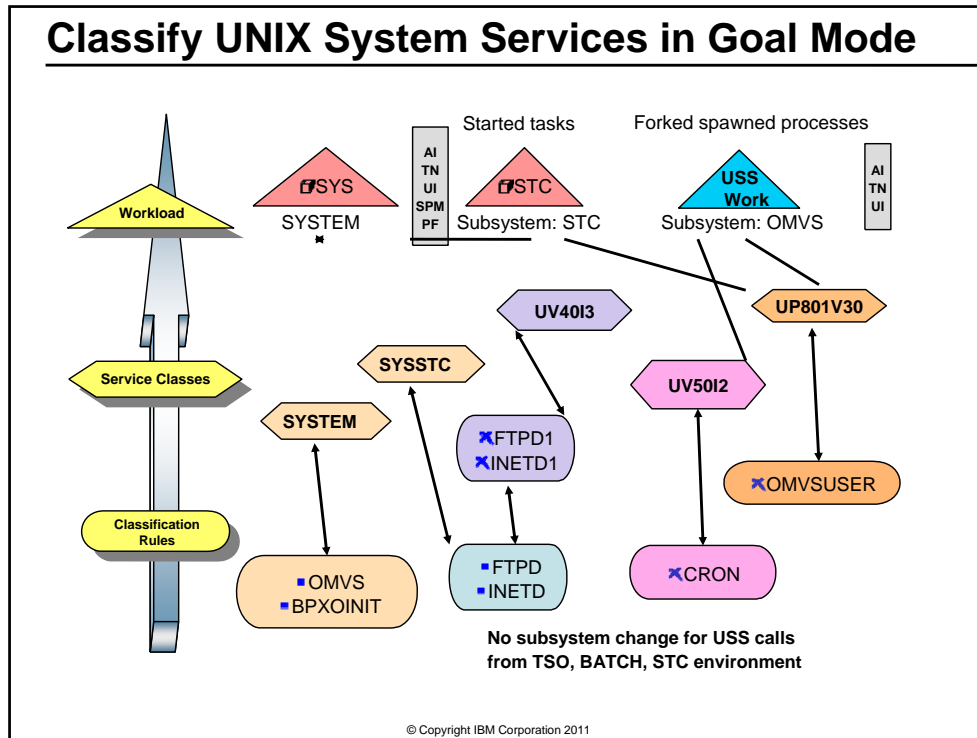
- **A WLM transaction represents a WLM "unit of work"**
  - basic workload entity for which WLM collects a resource usage value
  - foundation for statistics presented in workload activity report
  - represents a single subsystem "work request"
- **Subsystems can implement one of three transaction types**
  - **Address Space:**
    - ▶ WLM transaction measures all resource used by a subsystem request in a **single address space**
    - ▶ Used by JES (a batch job), TSO (a TSO command), OMVS (a process), STC (a started task) and ASCH (single APPC program)
  - **Enclave:**
    - ▶ Enclave created and destroyed by subsystem for each work request
    - ▶ WLM transaction measures resources used by a single subsystem request across **multiple address spaces and systems**
    - ▶ Exploited by "new workload" subsystems - Component Broker(WebSphere), DB2, DDF, IWEB, MQSeries, LDAP, NETV
  - **CICS/IMS Transactions**
    - ▶ Neither address space or enclave oriented - special type
    - ▶ WLM transaction measures resource used by a single CICS/IMS transaction program request

# WLM Workload Types









## OMVS Subsystem Panel

```

Subsystem-Type Xref Notes Options Help
-----
Command ==> Modify Rules for the Subsystem Type Row 1 to 9 of 9
SCROLL ==> PAGE

Subsystem Type . . : OMVS Fold qualifier names? Y (Y or N)
Description . . .

Action codes: A=After C=Copy M=Move I=Insert rule
              B=Before D=Delete row R=Repeat IS=Insert Sub-rule
              More ==>

Action Qualifier Selection Row 1 to 8 of 8
-----
1 Command ==>
1 Select a type with "/"
1 Sel Name Description
1 - AI Accounting Information RCGIAN
1 - PX Sysplex Name RCANNA
1 - SY Sysname ITSOFN
1 - SYG Sysname Group ERWCTG1
1 - TN Transaction Name WAS
1 - TNG Transaction Name Group
*****
1 - UI Userid
1 - UIG Userid Group
*****
***** Bottom of data *****
  
```

© Copyright IBM Corporation 2011

## USS Classification Rules Definitions Sample

```

Subsystem-Type Xref Notes Options Help
-----
                          Modify Rules for the Subsystem Type          Row 1 to 5 of 5
Command ==> _____ SCROLL ==> PAGE

Subsystem Type . : OMVS          Fold qualifier names?  Y  (Y or N)
Description . . . OMVS users

Action codes:  A=After      C=Copy      M=Move      I=Insert rule
                B=Before    D=Delete row  R=Repeat    IS=Insert Sub-rule
                                                More ==>

-----Qualifier-----
Action  Type      Name      Start      Service      Report
-----
_____ 1  AI        UNI*      1          UP801V30    OMVSUSER
_____ 1  TN        PASCALX   _____
_____ 1  UI        PASCAL    _____
_____ 1  UI        OMVSUSR   _____
***** BOTTOM OF DATA *****
    
```

## USS Service Classes Definitions Sample

```

Service-Class View Notes Options Help
-----
                          Service Class Selection List          Row 20 to 37 of
Command ==> _____

Action Codes: 1=Create, 2=Copy, 3=Modify, 4=Browse, 5=Print, 6=Delete,
              /=Menu Bar

Action Class Description Workload
-----
_____ UP801V30 USS Percent 80%<1s vel 30 I2 OMVS
_____ USSAI USS based on Accounting Info OMVS
_____ USSSTN USS based on Jobname OMVS
3 _____ USSUI USS based on userid OMVS
_____ UV30I2 USS velocity importance 2 OMVS
_____ UV30I3 USS velocity importance 3 OMVS
_____ UV40I2 USS velocity importance 2 OMVS
_____ UV50I1 USS velocity importance 1 OMVS
    
```

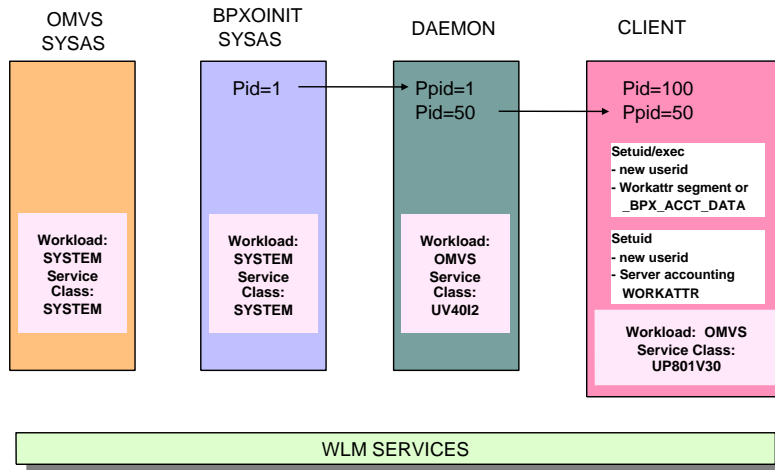
```

Service Class Name . . . . . : USSUI
Description . . . . . : USS based on userid
Workload Name . . . . . : OMVS (name or ?)
Base Resource Group . . . . . : (name or ?)

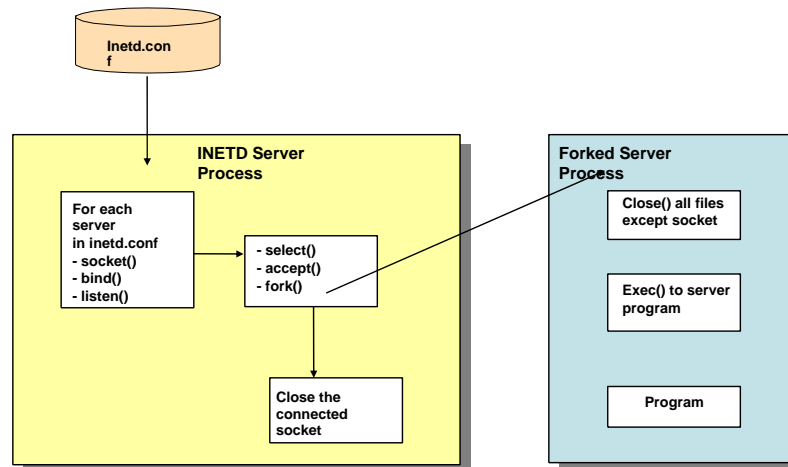
Specify BASE GOAL information. Action Codes: I=Insert new period,
E=Edit period, D=Delete period.

---Period--- -----Goal-----
Action # Duration Imp. Description
-----
_____ 1 1000 2 80% complete within 00:00:01.000
_____ 2 3 Execution velocity of 30
    
```

## Workload Classification



## Generic Listener Program



## Telnet Login

```
IEE114I 11.59.49 2000.363 ACTIVITY 407
JOBS   M/S   TS  USERS   SYSAS   INITS   ACTIVE/MAX VTAM   OAS
00014  00021 00001   00028   00040   00001/00020   00030
LLA    LLA    LLA    NSW   S   VLF    VLF    VLF    NSW   S
APPC   APPC   APPC   NSW   S   ASCH   ASCH   ASCH   NSW   S
SDSFSRV SDSFSRV SDSF   NSW   S   RRS    RRS    RRS    NSW   S
NET    NET    VTAM   NSW   S   RACF   RACF   RACF   NSW   S
JES2   JES2   IEFPROC NSW   S   TSO    TSO    TCAS   OWT   S
TCPIPOE STEP1  TCPIP   OWT   A   INETD1 STEP1  OMVSUSR OWT   AO
INETD5  *OMVSEX OMVSUSR OWT   AO
```

Forked  
Address Space

```
D J,BPXAS
BPXAS  BPXAS  IEFPROC OWT IO A=00A3 PER=NO SMC=000
PGN=N/A DMN=N/A AFF=NONE
CT=000.033S ET=239.456S
WUID=STC01054 USERID=OMVSUSR
WKL=OMVS SCL=UV40I2 P=1
RGP=N/A SRVR=NO QSC=NO
ADDR SPACE ASTE=7F94D8C0
```

pid/ppid

```
D OMVS,A=ALL
OMVS  000E ACTIVE OMVS=(00)
OMVSUSR INETD1 0067 16842797 1 1FI 11.32.57 .036
LATCHWAITPID= 0 CMD=/usr/sbin/inetd /etc/inetd.conf
OMVSUSR INETD5 00A3 83951729 16842797 1FI 11.59.33 .026
LATCHWAITPID= 0 CMD=otelnetsd -l
```

## Telnet Login (2)

```
SDSF DA X9 (ALL) PAG 0 SIO 2596 CPU 28/ 27 LINE 1-2 (2)
COMMAND INPUT ==> SCROLL ==> CSR
PREFIX-PAS* DEST=(ALL) OWNER=* SORT=CPU%#D
NP JOBNAME C Pos DP Real Paging SIO CPU% EXCP-Cnt SrvClass ASID ASIDX
PASCAL6 LO FF 335 0.00 0.00 0.00 358 UP801V30 123 007B
PASCAL LO FF 317 0.00 0.00 0.00 177 UP801V30 163 00A3
```

password  
validation  
setuid()  
setgid()

```
BPXAS  BPXAS  IEFPROC OWT IO A=00A3 PER=NO SMC=000
PGN=N/A DMN=N/A AFF=NONE
CT=000.017S ET=007.998S
WUID=STC01054 USERID=PASCAL
WKL=OMVS SCL=UP801V3 P=1
RGP=N/A SRVR=NO QSC=NO
ADDR SPACE ASTE=7F94D8C0
```

```
BFX0040I 12.19.23 DISPLAY OMVS 426
OMVS  000E ACTIVE OMVS=(00)
USER  JOBNAME ASID PID PPID STATE START CT_SECS
PASCAL PASCAL6 007E 83951681 83951729 1CI 12.04.39 .096
LATCHWAITPID= 0 CMD=sh -L
PASCAL PASCAL 00A3 83951729 16842797 1FI 11.59.33 .051
LATCHWAITPID= 0 CMD=otelnetsd -Y lig32-239-184-70.emea.lig-di
```

```
D TCPIP,,N,ALLC
EZZ2500I NETSTAT CS V2R10 TCPIPOE 430
USER ID CONN LOCAL SOCKET FOREIGN SOCKET STATE
BPXOINIT 0000000B 0.0.0.0.10007 0.0.0.0.0 LISTEN
INETD1 000B79FE 9.100.193.138.23 32.239.184.70.2029 ESTBLSH
D TCPIP,,N,BYTE
INETD1 0000001434 0000000083 00023 32.239.184.70.2029 ESTBLSH
```

## Telnet Accounting

```

SDSF DA VSN1 MVS29Z PAG 0 SIO 2 CPU 3/ 3 LINE 1-2 (2)
NP JOBNAME JOBID OWNER C POS DP REAL SIO CPU% SRVCLASS ASID ASIDX
TILLARD STC08074 TILLARD LO FF 306 0.00 0.00 UP801V30 104 0068
  
```

```

alu pascal WORKATTR(WAACNT(UNIXACCT))
lu pascal noracf workattr
USER=PASCAL

WORKATTR INFORMATION
-----
WAACNT= UNIXACCT
***
  
```

WLM classification rules  
associate USSAI service  
class for the qualifier type  
AI=UNIXACCT

```

SDSF DA VSN1 MVS29Z PAG 0 SIO 3 CPU 5/ 4 LINE 1-3 (3)
NP JOBNAME JOBID OWNER C POS DP REAL SIO CPU% SRVCLASS ASID ASIDX
PASCAL STC08092 PASCAL LO FF 301 0.00 0.00 USSAI 93 005D
PASCAL3 STC08093 PASCAL LO FF 279 0.00 0.00 USSAI 144 0090
  
```

© Copyright IBM Corporation 2011

## Display z/OS UNIX (OMVS) Information

```

BPXO040I 14.20.47 DISPLAY OMVS 492
OMVS 000E ACTIVE OMVS=(00)
USER JOBNAME ASID PID PPID STATE START CT_SECS
OMVSUSR BPXOINIT 0022 1 0 MRI 15.56.17 33.807
LATCHWAITPID= 0 CMD=BPXPINPR
SERVER=Init Process AF= 0 MF=00000 TYPE=FILE
DFS DFSCM 001D 65539 1 1R 15.56.17 15.317
LATCHWAITPID= 0 CMD=IOECMINI
OMVSUSR JGATE 0049 16842757 1 1WI 15.56.20 144.219
LATCHWAITPID= 0 CMD=-sh -c cd /usr/lpp/ctg/bin;ctgstart -noi
OMVSUSR CSQ9CHIN 0042 65542 1 1R 15.56.18 140.124
LATCHWAITPID= 0 CMD=CSQXJST
OMVSUSR CSQ9CHIN 0042 16842759 1 1R 15.56.22 140.124
LATCHWAITPID= 0 CMD=CSQXDISP
OMVSUSR CSQ9CHIN 0042 16842760 1 1R 15.56.22 140.124
LATCHWAITPID= 0 CMD=CSQXDISP
TCPIPOE TCPIPOE 0051 83951627 1 1R 14.58.10 1736.926
LATCHWAITPID= 0 CMD=EZBTMST
OMVSUSR SYSLOGD8 004D 33620006 1 1FI 15.57.01 159.967
LATCHWAITPID= 0 CMD=/usr/sbin/syslogd -f /etc/syslog.conf
OMVSUSR JGATE 0049 16842793 16842757 HKI 15.57.01 144.219
LATCHWAITPID= 0 CMD=/usr/lpp/java/J1.1/bin/./bin/mvs/native
OMVSUSR INETD1 0067 16842797 1 1FI 11.32.57 .141
LATCHWAITPID= 0 CMD=/usr/sbin/inetd /etc/inetd.conf
  
```

© Copyright IBM Corporation 2011

## Display Process ID

```

D OMVS,PID=50397330
BFXO040I 00.14.18 DISPLAY OMVS 882
OMVS      000E ACTIVE          OMVS=(00)
USER      JOBNAME ASID        PID          PPID STATE   START
CT_SECS
WEBSRV    WAS      0105     50397330          1 HK      00.13.07
5.820
LATCHWAITPID=          0 CMD=IMWHTTDP
THREAD_ID  TCB@      PRI_JOB  USERNAME  ACC_TIME SC STATE
OF6824000000000000 008FF020 OMVS          5.196 KIN  KU
OF6830000000000001 008DE190          .001 SPM  JY V
OF6848000000000002 008E1BF8          .005 SPM  JY V
OF6854000000000003 008E1D90 OMVS          .001 ANR  JF V
.....
OF6E00000000000064 0082E190          .001 SPM  JY V
OF6E0C000000000065 008BA068 OMVS          .279 SLP  JS
OF6E18000000000066 0083A0F0          .001 PTX  JY
OF6E24000000000067 008E1708          .001 PTX  JY
OF6E3C000000000068 008E1570          .005 STE  JY
OF6E48000000000069 008E13D8          .013 STE  JY
OF6E5400000000006A 008BC0F0 OMVS          .009 SLP  JS V
OF6E6000000000006B 008BC288 OMVS          .039 CTW  JK V
OF6E6C00000000006C 008C2E88 OMVS          .004 CTW  JK V
OF6E7800000000006D 008E1148          .001 STE  JY V
    
```

## Process and Thread status

```

D OMVS,PID=50397330
BFXO040I 00.14.18 DISPLAY OMVS 882
OMVS      000E ACTIVE          OMVS=(00)
USER      JOBNAME ASID        PID          PPID STATE   START   CT_SECS
WEBSRV    WAS      0105     50397330          1 HK      00.13.07   5.820
LATCHWAITPID=          0 CMD=IMWHTTDP
THREAD_ID  TCB@      PRI_JOB  USERNAME  ACC_TIME SC STATE
OF6824000000000000 008FF020 OMVS          5.196 KIN  KU
OF6830000000000001 008DE190          .001 SPM  JY V
OF6848000000000002 008E1BF8          .005 SPM  JY V
OF6854000000000003 008E1D90 OMVS          .001 ANR  JF V
.....
OF6E00000000000064 0082E190          .001 SPM  JY V
OF6E0C000000000065 008BA068 OMVS          .279 SLP  JS
OF6E18000000000066 0083A0F0          .001 PTX  JY
OF6E24000000000067 008E1708          .001 PTX  JY
OF6E3C000000000068 008E1570          .005 STE  JY
OF6E48000000000069 008E13D8          .013 STE  JY
OF6E5400000000006A 008BC0F0 OMVS          .009 ELP  JS V
OF6E6000000000006B 008BC288 OMVS          .039 CTW  JK V
OF6E6C00000000006C 008C2E88 OMVS          .004 CTW  JK V
OF6E7800000000006D 008E1148          .001 STE  JY V
    
```

Current Thread state:

- A Message queue receive wait
- B Message queue send wait
- C Communication system kernel wait
- D Semaphore operation wait
- E Quiesce frozen
- F File system kernel wait
- G MVS Pause wait
- K Other kernel wait (for example, pause or sigsuspend)
- J The thread was pthread created rather than dubbed
- N The thread is medium weight
- O The thread is asynchronous and medium weight
- P Prtrace kernel wait
- Q Quiesce termination wait
- R Running (not kernel wait)
- S Sleeping
- U Initial process thread (heavy weight and synchronous)
- V Thread is detached
- W Waiting for child (wait or waitpid callable service)
- X Creating new process (fork callable service is running)
- Y Thread is in an MVS wait

Last syscall request:

Refer: UNIX System Services  
 Programming: Assembler Callable Services Reference  
 Document Number SC28-1899  
 Example:  
 ANR: Accept and receive  
 SPM: Sigprocmask  
 SLP: Sleep  
 CTW: Cond\_timed\_wait  
 PTX: Pthread exit and get  
 STE: Set Timer Event

## Display z/OS UNIX Options

### D OMVS,OPTIONS

```

BPX0043I 14.17.13 DISPLAY OMVS 485
OMVS      000E ACTIVE          OMVS=(00)
OS/390 UNIX CURRENT CONFIGURATION SETTINGS:
MAXPROCSYS      =      300      MAXPROCUSER      =      500
MAXFILEPROC     =     65535     MAXFILESIZE   = NOLIMIT
MAXCPUPTIME     = 2147483647   MAXUIDS       =      100
MAXPTY          =      256
MAXMMAPAREA     =      4096     MAXASSIZE     = 2147483647
MAXTHREADS      =     10000     MAXTHREADTASKS = 10000
MAXCORESIZE     = 4194304     MAXSHAREPAGES = 32768000
IPCMSGQBYTES    = 262144      IPCMSGQMNUM   = 10000
IPCMSGNIDS      =      500     IPCSEMNIIDS   = 20000
IPCSEMNOPI      =     32767    IPCSEMNSEMS   =      250
IPCshmPAGES     = 25600       IPCshmNIDS    =      500
IPCshmNSEGS     =      1000    IPCshmSPAGES  = 2621440
SUPERUSER       = ROOT        FORKCOPY      = COPY
STEPLIBLIST     =
USERIDALIASTABLE=
PRIORITYPG VALUES: NONE
PRIORITYGOAL VALUES: NONE
MAXQUEUEDSIGS   = 100000      SHRLIBRGNsize = 67108864
SHRLIBMAXPAGES =      4096     VERSION       = MOXR10
SYSCALL COUNTS = YES         TTYGROUP      = TTY
SYSPLEX         = YES         BRML SERVER    = X9
LIMMSG          = NONE

```

© Copyright IBM Corporation 2011

## Display and Set Process ID Limits

```

D OMVS,L,PID=50397330
BPX0051I 00.14.29 DISPLAY OMVS 884
OMVS      000E ACTIVE          OMVS=(00)
USER      JOBNAME ASID        PID      PPID STATE  START
CT_SECS
WEBSRV   WAS      0105    50397330      1 HK      00.13.07
5.830
LATCHWAITPID=      0 CMD=IMWHTTDP
PROCESS LIMITS:    LIMMSG=NONE

```

	CURRENT	HIGHWATER	PROCESS	LIMIT
MAXFILEPROC	25	26	65535	65535 *
MAXFILESIZE	---	---	NOLIMIT	NOLIMIT
MAXPROCUSER	52	90	NOLIMIT	100000 *
MAXQUEUEDSIGS	0	0	100000	50 *
MAXTHREADS	109	109	10000	10000 *
MAXTHREADTASKS	109	109	10000	1000 *
IPCshmNSEGS	0	0	1000	4194304 *
MAXCORESIZE	---	---	4194304	

```
SETOMVS PID=50397330,MAXTHREADS=50
```

```

BPX0015I THE SETOMVS COMMAND WAS SUCCESSFUL.
BPXI040I PROCESS LIMIT MAXTHREADS HAS REACHED 205% OF ITS CURRENT 970
CAPACITY OF 50 FOR PID=50397330 IN JOB WAS RUNNING IN ADDRESS SPACE 0105

```

© Copyright IBM Corporation 2011



## Display USS Limits

```

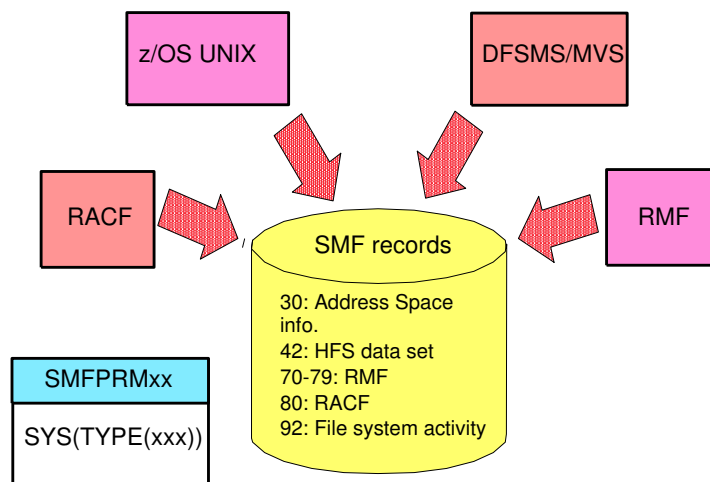
BPXO051I 14.19.05 DISPLAY OMVS 487
OMVS      000E ACTIVE          OMVS=(00)
SYSTEM WIDE LIMITS:          LIMMSG=NONE

```

	CURRENT USAGE	HIGHWATER USAGE	SYSTEM LIMIT
MAXPROCSYS	52	191	300
MAXUIDS	1	2	100
MAXPTYS	0	2	256
MAXMMAPAREA	0	25	4096
MAXSHAREPAGES	49	790	32768000
IPCMSGNIDS	14	48	500
IPCSEMNIDS	0	2	20000
IPCSHMNIDS	0	7	500
IPCSHMSPAGES	0	16	2621440
IPCMSGQBYTES	---	1068	262144
IPCMSGQMNUM	---	267	10000
IPCSHMMPAGES	---	10240	25600
SHRLIBRGNSIZE	0	0	67108864
SHRLIBMAXPAGES	374	710	4096

© Copyright IBM Corporation 2011

## Collect z/OS UNIX Activity Data



© Copyright IBM Corporation 2011

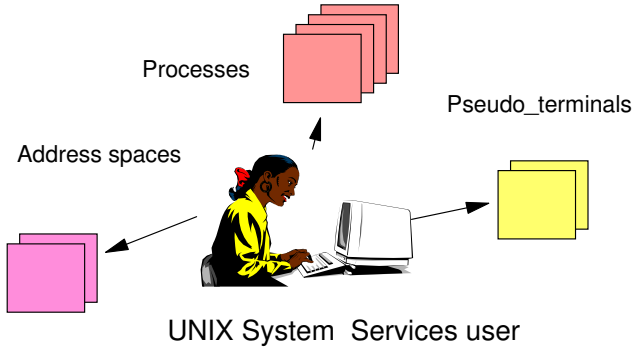
# Tune z/OS Unix Process Resources

## RMF Monitor I Kernel Activity Report

OMVS PROCESS ACTIVITY

MAXIMUM (TOT)	PROCESSES			USERS			PROCESSES PER USERS		
	MINIMUM	AVERAGE	MAXIMUM	MINIMUM	AVERAGE	MAXIMUM	MINIMUM	AVERAGE	MAXIMUM
	300			100			50		
CURRENT (TOT)	53	55.50	56	1	1.000	1			
OVERRUNS (N/S)	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000

- BPXPRMxx**
- MAXUIDS ?
  - MAXPTYS ?
  - MAXRTYS ?
  - MAXPROCUSER ?
  - MAXPROCSYS ?



© Copyright IBM Corporation 2011


# Control IPC Resources

## RMF Monitor I Kernel Activity Report

OMVS INTER-PROCESS COMMUNICATION

MAXIMUM (TOT)	MESSAGE QUEUE IDS			SEMAPHORE IDS			SHARED MEMORY IDS			SHARED MEMORY PAGES		
	MINIMUM	AVERAGE	MAXIMUM	MINIMUM	AVERAGE	MAXIMUM	MINIMUM	AVERAGE	MAXIMUM	MINIMUM	AVERAGE	MAXIMUM
	500			20000			500			2621K		
CURRENT (TOT)	30.00	29.80	30.00	1.000	0.993	1.000	3.000	2.980	3.000	2.000	1.987	2.000
OVERRUNS (N/S)	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000



- BPXPRMxx**
- IPCMSGNIDS
  - IPCSEMNIDS
  - IPCSTMNIDS
  - IPCSTMSPAGES
  - MAXQLIEDSIGS



# Control Memory Usage

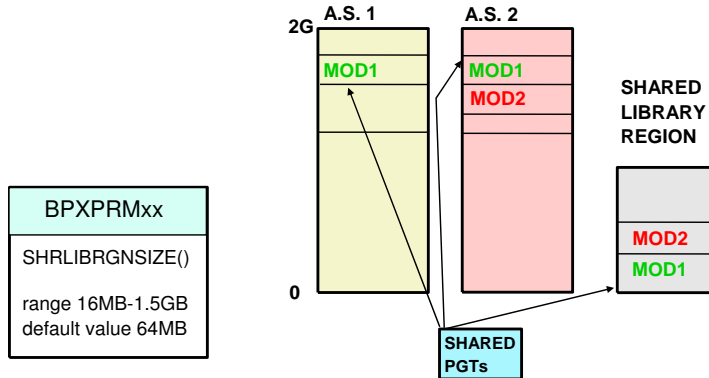
## RMF Monitor I Kernel Activity Report

	OMVS MEMORY MAP						SHARED LIB REGION			QUEUED SIGNALS		
	MEMORY MAP STORAGE PAGES			SHARED STORAGE PAGES			MAX SHARED LIBRARY REGION			MAXIMUM QUEUED SIGNALS		
MAXIMUM (TOT)	MINIMUM	AVERAGE	MAXIMUM	MINIMUM	AVERAGE	MAXIMUM	MINIMUM	AVERAGE	MAXIMUM	MINIMUM	AVERAGE	MAXIMUM
	4096			32.8M			64			100K		
CURRENT (TOT)	0.000	0.000	0.000	341.0	339.9	341.0	0.000	0.000	0.000	0.000	0.000	0.000
OVERRUNS (N/S)	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000



# System Shared Libraries

■ Optimizes Sharing of HFS programs across the system



- ★ Intended for Large HFS Executables
- ★ Modules Shared on a megabyte boundary
- ★ Maps into high end private storage
- ★ Extended attribute I (minusL) set on

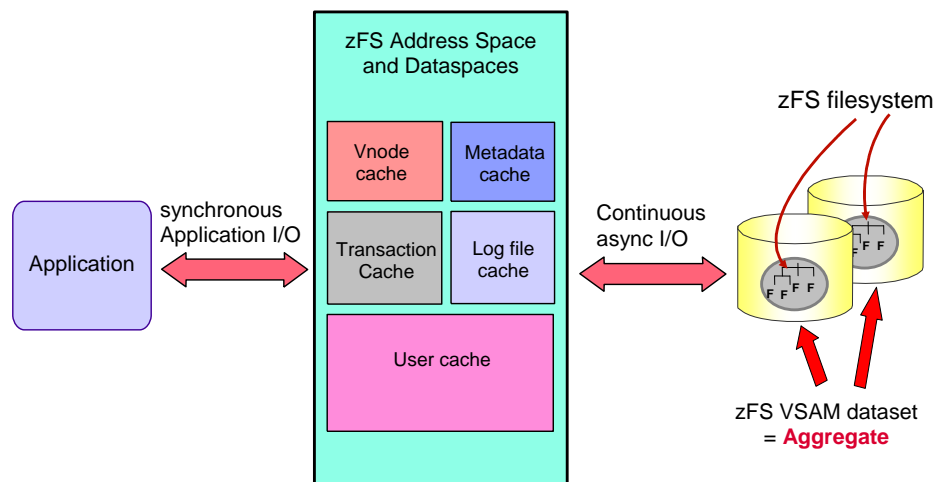


## zFS Reporting with Monitor III

- **RMF Monitor III support for the new UNIX file system - zFS**
  - New Monitor III gathering option ZFS | NOZFS
  - New Monitor III zFS reports provides data on
    - ▶ zFS response time / wait times
    - ▶ zFS cache activity
    - ▶ zFS activity / capacity by aggregate
    - ▶ zFS activity / capacity by filesystem
  - Data helps to control the zFS environment according to
    - ▶ Cache sizes
    - ▶ I/O balancing
    - ▶ Capacity control for zFS aggregates

© Copyright IBM Corporation 2011

## Overview of zFS Environment



© Copyright IBM Corporation 2011

## zFS Summary Report

```

RMF V1R7      zFS Summary Report      Line 1 of 3
Command ==>>      Scroll ==>> CSR

Samples: 120      System: TRX2 Date: 09/07/04 Time: 16.30.00 Range: 120 Sec

----- Response Time ----- Cache Activity -----
----- Wait% ----- User ----- Vnode -- Metadata - -Trx -
Avg I/O Lock Sleep Rate Hit% Read% Dly% Rate Hit% Rate Hit% Rate
0.27 0.4 80.7 0.0 405.5 100 0.0 0.0 0.0 0.0 10.9 99.9 0.2

-----Aggregate Name----- Size Use% --Mode- FS Read Write
(B/Sec)
OMVS.TRX2.LOCAL.COMPAT.A.ZFS 528K 29.2 R/W CP 1 0 137
OMVS.TRX2.LOCAL.MULTI.A.ZFS 2160K 8.0 R/W MS 3 0 137
RMF.TEST.ZFS.HFSMP1 7200K 16.1 R/W MS 3 137 1058
    
```

**Response time section:**

- ▶ Average response time for zFS request
- ▶ Wait percentages

**Aggregate section:**

- ▶ Capacity data
- ▶ Mount mode
- ▶ # filesystems in the aggregate
- ▶ Read / Write rates (Bytes per second)

**Cache activity:**

- ▶ Request rates
- ▶ Hit ratios
- ▶ % read requests
- ▶ % requests delayed

## Monitor III - zFS I/O Details Popup

```

RMF V1R7      zFS Summary Report      Line 1 of 3
Command ==>>      Scroll ==>> CSR

Samples: 120      System: TRX2 Date: 09/07/04 Time: 16.30.00 Range: 120 Sec

----- Response Time ----- Cache Activity -----
----- Wait% ----- User ----- Vnode -- Metadata - -Trx -
Avg I/O Lock Sleep Rate Hit% Read% Dly% Rate Hit% Rate Hit% Rate
0.27 0.4 80.7 0.0 405.5 100 0.0 0.0 0.0 0.0 10.9 99.9 0.2

-----Aggregate Name----- Size Use% --Mode- FS Read Write
(B/Sec)
OMVS.TRX2.LOCAL.COMPAT.A.ZFS 528K 29.2 R/W CP 1 0 137
OMVS.TRX2.LOCAL.MULTI.A.ZFS 2160K 8.0 R/W MS 3 0 137
RMF.TEST.ZFS.HFSMP1 7200K 16.1 R/W MS 3 137 1058
    
```

zFS Summary - I/O Details by Type

Count	Waits	Cancl	Merge	Type
815	326	0	0	FILE SYSTEM METADATA
346	23	0	29	LOG FILE
1447	175	0	2	USER FILE DATA

Press Enter to return to the Report panel.

■ I/O details per request type

## Monitor III: Vnode Cache and User Cache Details

- Vnode cache details:
  - Request rate, hit ratio
  - Vnode statistics

```

zFS Summary - Vnode Cache Details

Request Rate : 14.2      vnodes      : 65536
Hit%         : 99.9     vnode size  : 168
              ext. vnodes : 65536
              ext. vnode size : 668
              open vnodes  : 12
              held vnodes  : 44

Press Enter to return to the Report panel.
    
```

```

Command ==>> RMP V1R7 zFS
Samples: 120 System: TRX2 Date: 0
Sec

----- Response Time -----
Wait% ----- Cache Activity -----
----- User ----- Vnode -- Metadata -- Trx --
Avg I/O Lock Sleep Rate Hit% Ready Dly% Rate Hit% Rate Hit% Rate
0.27 0.4 80.7 0.0 405.5 100 0.0 0.0 0.0 0.0 10.9 99.9 0.2
    
```

```

zFS Summary - User Cache Details

Read Rate      : 27.4      Size       : 256M
Write Rate     : 16.9     Total Pages : 65536
Read Hit (%)   : 59.4     Free Pages  : 12703
Write Hit (%)  : 100.0    Segments   : 8192
Read Delay (%) : 1.3
Write Delay (%) : 0.0     User Cache readahead: ON
Async Read Rate : 10.1    Storage fixed : NO
Scheduled Write Rate : 83.8
Page Reclaim Writes : 0
Fsyncs        : 0

Press Enter to return to the Report panel.
    
```

- User cache details:
  - Request rates, hit ratios, delays
  - Storage statistics

## Monitor III: Metadata and Transition Cache Details

- Transaction cache details:
  - Transaction rate details
  - Transaction state breakdown

```

zFS Summary - Transaction Cache Details

Transaction Rate : 236.9
EC Merge Rate   : 16.4

Transactions:
Active          : 355      Allocated : 2153
Pending        : 154
Complete       : 82
Free           : 12

Press Enter to return to the Report panel.
    
```

```

Command ==>> RMP V1R7 zFS S
Samples: 120 System: TRX2 Date: 09/
Sec

----- Response Time -----
Wait% ----- Cache Activity -----
----- User ----- Vnode -- Metadata -- Trx --
Avg I/O Lock Sleep Rate Hit% Ready Dly% Rate Hit% Rate Hit% Rate
0.27 0.4 80.7 0.0 405.5 100 0.0 0.0 0.0 0.0 10.9 99.9 0.2
    
```

```

zFS Summary - Metadata Cache Details

Request Rate : 1827.5      Size       : 32M
Hit%         : 99.7       Buffers    : 4096
Storage fixed : NO

----- Metadata Backing Cache Details -----

Request Rate : 227.5      Size       : 32M
Hit%         : 99.7       Buffers    : 4096
Storage fixed : NO

Press Enter to return to the Report panel.
    
```

- Metadata cache details:
  - Request rate, hit ratio
  - Storage statistics
  - Also for Metadata Backing

## Monitor III: zFS File System Activity Report

```

RMF V1R7 zFS Activity Report                               Line 1 of 3
Command ==>                                               Scroll ==> CSR

Samples: 120 System: TRX2 Date: 09/07/04 Time: 10.20.00 Range: 120 Sec

Aggregate Name : ALL

----- File System Name/Mount Point ----- Mode  Limit  Usg%   Operation
OMVS.TRX1.LOCAL.COMPAT.A.ZFS                R/W   383K   2.3    0.0
/SYSTEM/local/ZFS/COMPAT/a
ZFSM1                                        R/W   200K   4.5    0.0
/u/bpmu/zfs1
ZFSM2                                        R/O   200K   4.5    0.0
/u/bpmu/zfs2
ZFSM3                                        N/M   200K   4.5    0.0

```

### ■ zFS file system activity:

- File system name
- Mount information
- Capacity data
- Activity rate (Operations per second)

© Copyright IBM Corporation 2011

## Summary

- Use WLM OMVS Subsystem Rules to manage BPXAS address spaces
- Monitor USS resource usage with RMF
- Use RMF III to monitor USS processes and zFS
- Understand info available in z/OS console displays