

Semaphores -

The number of semaphores and semaphore arrays that your system will use is dependent on the maximum number of agents that will be allowed to be up at any one time. If the number of available semaphores is exceeded DB2 will crash and a Severe error will be logged to the instance's ~/sqllib/db2dump/db2diag.log file. The error will look similar to -

```
2007-12-13-14.23.30.683763-360 I42676881A390   LEVEL: Severe
PID   : 1265           TID  : 4398046711072PROC : db2ipccm 0
INSTANCE: bgpsysdb     NODE  : 000
FUNCTION: DB2 UDB, common communication, sqlccGetIPCs, probe:50
RETCODE : ZRC=0x850F0081=-2062614399=SQLO_SSEM_EXCEED_MAX
        "Requesting too many semaphores"
        DIA8336C Requested too many semaphores.
```

Changing the Parameter -

System wide semaphore levels are set in /proc/sys/kernel/sem . Current limits can also be viewed via ipcs -ls. A typical system might look like -

```
cat /proc/sys/kernel/sem
256 4194304 256 16384
bgpsysdb@bgpfs12:~/sqllib/db2dump> ipcs -ls
```

```
----- Semaphore Limits -----
max number of arrays = 16384
max semaphores per array = 256
max semaphores system wide = 4194304
max ops per semop call = 256
semaphore max value = 32767
```

To set change this setting you can modify /etc/sysctl.conf to have the values you want by adding the line -

```
kernel.sem = 256 4194304 256 16384
```

In order to make the change take effect immediately you can run sysctl -p /etc/sysctl.conf

Recommendations -

The 4 values you'll see in /proc/sys/kernel/sem are (from right to left on the line) -

```
semmsl – max semaphores per set or array
semms – max semaphores system wide
```

semopm – max operations per semop call
semmni – max semaphore sets or arrays system wide

Typically, semmni should be set to a value 2 times the value for db2's MAXAGENTS setting + MAX_CONNECTIONS. So if you have MAXAGENTS at 4000 and MAX_CONNECTIONS at 8000 that would be $2 \times 4000 + 8000 = 16000$. So the above value of 16384 is sufficient. The value of SEMMSL is generally safe at any value from 100-256. SEMOPM is typically set to SEMMSL. SEMMNS should be set to a value equal to SEMMSL*SEMMNI. Above we have it set to 4194304 (256*16384).

Depending on the value of your MAXAGENTS in db2 you can decrease the value of SEMMNS if desired -

MAXAGENTS	SEMMNS
1500 or less	4096
1500-2500	8192
2500-4000	16384

How to Investigate -

You can monitor a systems semaphore usage real time by issuing the following command as root-

```
watch ipcs -u
```

You'll want to run that as applications run, and adjust the /proc/sys/kernel/sem values if any of the maximums get met. Here is typical output from the command -

```
----- Shared Memory Status -----  
segments allocated 46  
pages allocated 1775525  
pages resident 346760  
pages swapped 587591  
Swap performance: 0 attempts      0 successes  
  
----- Semaphore Status -----  
used arrays = 343  
allocated semaphores = 430  
  
----- Messages: Status -----  
allocated queues = 22  
used headers = 0  
used space = 0 bytes
```