

# CS615 - Aspects of System Administration

## Performance Tuning and Optimization

---

Department of Computer Science  
Stevens Institute of Technology  
Jan Schaumann

`jschauma@cs.stevens.edu`

`http://www.cs.stevens.edu/~jschauma/615A/`

## Symptoms

---

“The system feels slow.”

## Symptoms

---

“How much time does it take to run my job?”

# Symptoms

---



## Diagnosis

---

Know your applications.

## Diagnosis

---

Know your applications.

Know your users.

## Diagnosis

---

Know your applications.

Know your users.

Know your traffic patterns.

## Diagnosis

---

Know your applications.


Know your users.

Know your traffic patterns.

*Know your systems.*

## Diagnosis

---



## Monitoring System Activity

---

“How much time does it take to run my job?”

```
/usr/bin/time <application>
```


# Diagnosis

---



# Diagnosis

---



## Know your systems.

---

### Useful tools:

- simple:

- cron(8)
- ps(1) / top(1)
- time(1)
- uptime(1) / w(1)

- more advanced:

- dtrace(1M)
- iostat(8)
- sa(8)
- systat(1)
- sar(1)
- vmstat(1)

## Remedy

---

If you don't have enough of something, you have very few options:



## Remedy

---

If you don't have enough of something, you have very few options:

You can get more of it.



## Remedy

---

If you don't have enough of something, you have very few options:

You can use less of it.



## Remedy

---

If you don't have enough of something, you have very few options:

You can ration / distribute the amount you do have.



## Monitoring CPU Workload

---

System Load Average: average number of processes in a kernel's run queue.

```
up 1318 days, 13:46, 1 user, load averages: 993.81, 272.91, 1012.18
up   4 days, 18:28, 0 users, load averages: 70.54, 51.09, 47.18
up   8 days, 23:05, 0 users, load averages: 7.07, 16.89, 20.28
up  314 days, 44 mins, 0 users, load averages: 2.78, 3.97, 3.88
```

## Managing CPU Workload

---

Things you can do to keep processes under control:

- “get more of it”:
  - get faster CPUs
  - get more CPUs (per host)
  - get more hosts

## Managing CPU Workload

---

Things you can do to keep processes under control:

- “use less of it”:
  - optimize programs
  - restrict certain CPUs for specific usage (`cpuset(1)`)
  - change a process’s priorities (`nice(1)` and `renice(1)`)
  - create system-wide CPU limits (`ulimit`)
  - suspend and/or kill processes

## Managing CPU Workload

---

Things you can do to keep processes under control:

- “ration / redistribute the amount you have”:
  - schedule processes at off-hours
  - rewrite program(s) to use multithreading
  - parallelize program(s) across multiple hosts

## Memory Performance

---

Things to consider:

- *paging vs. swapping*
- monitoring memory usage with `vmstat(1)` and `sar(1)`
- use of `ulimit` (or `limit`) to set the `stacksize`, `datasize`, `memorysize`
- adjust buffer cache
- managing swap area

## Disk Performance

---

Mostly covered in previous lecture. To summarize:

- hardware limitations
- filesystem usage
- I/O subsystem configuration
- filesystem housekeeping

## Jan's Words of Wisdom

---

Thursday, April 16th, 2009

## Words of Wisdom

---

Don't try to fix  
what isn't broken.

## So many knobs...

---

```
$ sysctl -a
kern.maxproc = 148
kern.maxfiles = 548
kern.ngroups = 16
kern.iov_max = 1024
kern.mbuf.msize = 256
kern.mbuf.mclbytes = 2048
kern.mbuf.nmbclusters = 1024
kern.mbuf.mblowat = 16
kern.mbuf.mcllowat = 8
vm.nkmempages = 32768
vm.bufcache = 15
vm.bufmem = 11900928
vm.bufmem_lowater = 9876480
vm.bufmem_highater = 79011840
net.inet.tcp.sendspace = 32768
net.inet.tcp.recvspace = 32768
```

## So many knobs...

---

```
$ man 4 options
options MAXUSERS=integer
options MAXUPRC=integer
options NOFILE=integer
options MAXFILES=integer
options TCP_SENDSIZE=value
options TCP_RECVSPACE=value
options TCP_INIT_WIN=value
options SYSVSHM
options SHMMAXPGS=value
options NMBCLUSTERS=value
options NKMEMPAGES=value
options BUFPAGES=value
options MAXTSIZ=bytes
options DFLDSIZ=bytes
options LIMITMEM=value
options NVNODE=integer
```

## Reading

---

### Processes:

- `fork(2)` and `exec(2)`
- <http://www.cse.fau.edu/~roy/cop4604.02s/notes/process.html>

### System Activity Monitoring:

- `systat(1)`, `ps(1)`, `vmstat(1)` et al

### Managing/Monitoring CPU Workload:

- `at(1)` and `batch(1)`
- `cpuset(1)`
- `nice(1)`

### Various kernel tuning:

- `options(4)`
- <http://www.netbsd.org/guide/en/chap-tuning.html>