

Linux Performance Tuning



Jaqui Lynch

Circle4 Consulting

UKCMG 2002 Session 1B6

Personal Perspective

- ◆ Technology
 - ◆ LVM – Logical volume manager
 - ◆ WLM – Workload manager – already on AIX
 - ◆ Linux??? IRD??? Watch this space
 - ◆ LPAR technology – already on AIX
- ◆ Linux is Linux
 - ◆ Suse, Redhat, Debian, IBM
 - ◆ Security and performance
- ◆ Capacity & Scalability Viewpoint
 - ◆ XPZ
 - ◆ Test on X, Develop on P, Production on Z



Circle4 Consulting

2

Agenda

- ◆ Various Commands
 - ◆ CPU
 - ◆ Memory
 - ◆ Disk and I/O
 - ◆ Network



Circle4 Consulting

3

Commands for CPU

- ◆ vmstat
- ◆ sar
- ◆ ulimit –u unlimited
 - ◆ Set max procs per user to unlimited
- ◆ top
- ◆ mpstat
- ◆ /proc/cpuinfo
- ◆ ps



Circle4 Consulting

4

more /proc/cpuinfo

```
processor      : 0
vendor_id     : GenuineIntel
cpu family    : 6
model         : 5
model name    : Pentium II (Deschutes)
stepping      : 2
cpu MHz       : 400.920144
cache size    : 512 KB
fdiv_bug      : no
hlt_bug       : no
sep_bug       : no
f00f_bug      : no
fpu           : yes
fpu_exception : yes
cpuid level   : 2
wp            : yes
flags         : fpu vme de pse tsc msr pae mce cx8 apic sep
               mtrr pge mca cmov pat pse36 mmx osfxsr
bogomips      : 399.77
```



CPU Time

- ◆ Real
 - ◆ Wallclock time
- ◆ User state CPU
 - ◆ Time spent running a users program
 - ◆ Includes library calls
 - ◆ Can be affected by compile optimization and writing efficient code
- ◆ System code CPU
 - ◆ Time spent in the system state on the users behalf
 - ◆ Kernel calls & all I/O routines
 - ◆ Can be affected by using blocking for I/O transfers
- ◆ I/O Time & Network Time
 - ◆ Time spent moving data and servicing I/O requests



uptime

- ◆ 12:23pm up 86 days, 4:42, 4 users, load average: 0.44, 0.61, 0.68
- ◆ 11:53am up 56 days 4:59 1 user load average: 5.45, 5.45, 5.46
- ◆ 09:20am up 23 days 15:23 1 user load average: 0.61, 0.65, 0.56
- ◆ Provides system load average
- ◆ Rough Indicator
- ◆ Counts as runnable all jobs waiting on disk or nfs I/O
- ◆ Last minute, last 5 minutes, last 15 minutes
 - ◆ Single user workstation <2
 - ◆ Light <3 4-7 heavy >7 ??????



top on linux

```
12:24pm up 86 days, 4:43, 4 users, load average: 0.66, 0.63, 0.68
119 processes: 115 sleeping, 3 running, 0 zombie, 1 stopped
CPU states: 25.3% user, 7.6% system, 0.0% nice, 67.0% idle
Mem: 517356K av, 514272K used, 3084K free, 550780K shrd, 33728K buff
Swap: 525288K av, 6916K used, 518372K free 152520K cached
```

PID	USER	PRI	NI	SIZE	RSS	SHARE	STAT	LIB	%CPU	%MEM	TIME	COMMAND
22339	sybase	19	0	16384	15M	16236	R	0	12.8	3.1	3095m	datasv
14590	jaqui	4	0	1268	1268	1016	S	0	7.9	0.2	0:00	ssh
14511	webid	12	0	12448	12M	8828	S	0	3.8	2.4	0:00	httd
25935	webid	7	0	14828	14M	6888	S	0	2.7	2.8	2:43	httd
14586	root	6	0	1060	1060	836	R	0	2.3	0.2	0:00	top
10589	webid	2	0	13256	12M	7348	S	0	0.5	2.5	0:26	httd
14514	webid	1	0	12588	12M	8616	S	0	0.3	2.4	0:00	httd
14585	webid	0	0	11336	11M	11000	S	0	0.3	2.1	0:00	httd
25930	root	1	0	11160	10M	11008	S	0	0.1	2.1	0:11	httd
25938	webid	0	0	14740	14M	7072	S	0	0.1	2.8	2:31	httd
25970	webid	0	0	15548	15M	6868	S	0	0.1	3.0	2:16	httd
13675	root	0	0	1272	1240	1028	R	0	0.1	0.2	0:01	sshd



mpstat

mpstat 2 10

Linux 2.2.5-15smp (jaqui.abc.com) 05/15/01

	CPU	%user	%nice	%system	%idle	intr/s
12:28:27						
12:28:29	all	90.00	0.00	4.00	6.00	236.50
12:28:31	all	86.00	0.00	7.50	6.50	199.50
12:28:33	all	86.50	0.00	10.00	3.50	251.00
12:28:35	all	87.00	0.00	5.00	8.00	199.50
12:28:37	all	84.00	0.00	3.50	12.50	197.50
12:28:39	all	87.00	0.00	9.00	4.00	192.50
12:28:41	all	90.50	0.00	7.50	2.00	230.00
12:28:43	all	86.50	0.00	7.50	6.00	248.50
12:28:45	all	82.00	0.00	10.00	8.00	292.50
12:28:47	all	87.50	0.00	3.50	9.00	189.00
Average:	all	86.70	0.00	6.75	6.55	223.65



Circle4 Consulting

9

mpstat -P ALL

mpstat -P ALL 2 10

Linux 2.2.17 (jaqui.abc.com) 05/15/01

	CPU	%user	%nice	%system	%idle	intr/s
11:29:18 AM						
11:29:20 AM	all	1.50	0.00	3.25	95.25	102.00
11:29:20 AM	0	0.00	0.00	0.00	100.00	102.00
11:29:20 AM	1	3.00	0.00	6.50	90.50	102.00

	CPU	%user	%nice	%system	%idle	intr/s
11:29:20 AM						
11:29:22 AM	all	1.50	0.00	1.00	97.50	101.50
11:29:22 AM	0	0.00	0.00	0.00	100.00	101.50
11:29:22 AM	1	3.00	0.00	2.00	95.00	101.50

.....

Average:	CPU	%user	%nice	%system	%idle	intr/s
Average:	all	0.30	0.00	0.42	99.28	103.00
Average:	0	0.00	0.00	0.00	100.00	103.00
Average:	1	0.60	0.00	0.85	98.55	103.00



Circle4 Consulting

10

sar -u

sar -u -f sardata.out

Linux 2.2.5-15smp (jaquil.abc.com) 05/15/01

	CPU	%user	%nice	%system	%idle
12:34:44					
12:34:46	all	53.50	0.00	9.00	37.50
12:34:48	all	79.00	0.00	11.50	9.50
12:34:50	all	80.00	0.00	6.50	13.50
12:34:52	all	59.00	0.00	8.00	33.00
12:34:54	all	75.50	0.00	15.50	9.00
Average:	all	69.40	0.00	10.10	20.50



Circle4 Consulting

11

sar -U

sar -U ALL -f sardata.out

Linux 2.2.17 (jaquil.abc.com) 05/15/01

	CPU	%user	%nice	%system	%idle
10:12:43 AM					
10:12:45 AM	0	2.50	0.00	1.50	96.00
10:12:45 AM	1	93.00	0.00	4.00	3.00
10:12:47 AM	0	96.50	0.00	3.50	0.00
10:12:47 AM	1	0.00	0.00	2.00	98.00
10:12:49 AM	0	34.50	0.00	2.50	63.00
10:12:49 AM	1	62.50	0.00	2.00	35.50
10:12:51 AM	0	48.00	0.00	2.00	50.00
10:12:51 AM	1	49.00	0.00	3.00	48.00
10:12:53 AM	0	29.00	0.00	2.50	68.50
10:12:53 AM	1	67.50	0.00	2.50	30.00
10:12:55 AM	0	0.00	0.00	1.50	98.50
10:12:55 AM	1	97.00	0.00	3.00	0.00
10:12:57 AM	0	74.50	0.00	10.00	15.50
10:12:57 AM	1	85.00	0.00	8.50	6.50
Average:	0	38.20	0.00	5.20	56.60
Average:	1	55.95	0.00	4.85	39.20



Circle4 Consulting

12

ps aux

ps aux

USER	PID	%CPU	%MEM	VSZ	RSS	TTY	STAT	START	TIME	COMMAND
root	1	0.0	0.0	1096	52	?	S	Feb18 0:03		init [3]
root	2	0.0	0.0	0	0	?	SW	Feb18 4:31		[kflushd]
root	3	0.0	0.0	0	0	?	SW	Feb18 0:00		[kpiod]
root	4	0.0	0.0	0	0	?	SW	Feb18 7:54		[kswapd]
root	5	0.0	0.0	0	0	?	SW<	Feb18 0:00		[mdrecoveryd]
bin	254	0.0	0.0	1092	216	?	S	Feb18 0:01		portmap
root	301	0.0	0.0	1280	372	?	S	Feb18 0:24		syslogd -m 0
jaqui	29044	0.0	0.1	1264	548	pts/3	S	10:08 0:00		-ksh
webid	29173	0.5	2.4	14432	12640	?	S	10:14 0:02		/usr/local/apache/bin/httpd -d /usr/local/apache
webid	29174	0.3	2.4	14840	12664	?	S	10:14 0:01		/usr/local/apache/bin/httpd -d /usr/local/apache
webid	29175	0.2	2.4	14728	12580	?	S	10:14 0:01		/usr/local/apache/bin/httpd -d /usr/local/apache
webid	29176	0.3	2.3	13908	12308	?	S	10:14 0:01		/usr/local/apache/bin/httpd -d /usr/local/apache
webid	29177	0.2	2.4	14724	12612	?	S	10:14 0:00		/usr/local/apache/bin/httpd -d /usr/local/apache
webid	29178	0.1	2.3	13784	12148	?	S	10:14 0:00		/usr/local/apache/bin/httpd -d /usr/local/apache
webid	29195	0.1	2.3	13888	12208	?	S	10:15 0:00		/usr/local/apache/bin/httpd -d /usr/local/apache
webid	29206	0.4	2.4	14728	12788	?	S	10:15 0:01		/usr/local/apache/bin/httpd -d /usr/local/apache



Commands for Memory

- ◆ /proc/meminfo
- ◆ sar
- ◆ /proc/sys/vm/*
- ◆ vmstat
- ◆ Lsmmod or look at /proc/modules
- ◆ vmstat
- ◆ swapon
- ◆ free
- ◆ top



vmstat (Linux)

vmstat 2 10

procs				memory				swap		io				system				cpu	
r	b	w	swpd	free	buff	cache	si	so	bi	bo	in	cs	us	sy	id				
1	0	0	6916	13700	27876	156100	0	0	4	4	1	4	3	4	1				
2	0	0	6916	9520	28196	160324	0	0	2121	348	202	138	79	4	17				
4	0	0	6916	4020	28560	163352	0	0	1520	0	219	272	80	15	5				
3	0	0	6916	1732	28532	165408	0	0	1226	250	234	208	73	8	20				
1	2	0	6916	3016	25628	164360	0	0	1250	569	247	395	78	9	13				
2	1	0	6916	2296	21976	164476	0	0	1430	485	245	250	79	11	10				
2	0	0	6916	2444	20208	164480	0	0	660	140	372	396	74	8	18				
2	0	0	6916	2360	19084	165736	0	0	1835	134	211	248	82	9	9				
1	0	0	6916	1192	17908	168180	0	0	2443	161	185	138	88	5	7				
2	0	0	6916	2824	16000	168544	0	0	2122	251	202	160	81	4	14				



swapon and free

swapon -s

Filename	Type	Size	Used	Priority
/dev/sda7	partition	525292	6916	-1

free

	total	used	free	shared	buffers	cached
Mem:	517356	515076	2280	465748	10708	179340
-/+ buffers/cache:	325028	192328				
Swap:	525288	6916	518372			

free -m

	total	used	free	shared	buffers	cached
Mem:	505	477	28	443	36	120
-/+ buffers/cache:	320	184				
Swap:	512	6	506			



/proc/meminfo

more /proc/meminfo

```
Mem: 529772544 436813824 92958720 552947712 32686080 62316544
Swap: 537894912 7081984 530812928
MemTotal: 517356 kB
MemFree: 90780 kB
MemShared: 539988 kB
Buffers: 31920 kB
Cached: 60856 kB
SwapTotal: 525288 kB
SwapFree: 518372 kB
```



sar -B

sar -B -f sardata.out

Linux 2.2.5-15smp (jaquil.abc.com) 05/15/01

	pgpgin/s	pgpgout/s	activepg	inadtypg	inaclnpg	inatarpg
12:34:44						
12:34:46	376.00	1.50	0	0	0	0
12:34:48	441.50	1.50	0	0	0	0
12:34:50	568.00	71.50	0	0	0	0
12:34:52	437.00	1.50	0	0	0	0
12:34:54	450.50	22.50	0	0	0	0
Average:	454.60	19.70	0	0	0	0



sar -r

sar -r -f sardata.out

Linux 2.2.5-15smp (jaquil.abc.com) 05/15/01

	12:34:44	12:34:46	12:34:48	12:34:50	12:34:52	12:34:54	Average:
kbmemfree	62196	455160	59564	57064	59368	62288	60096
kbmemused	455160	62196	457792	460292	457988	455068	457260
%memused	87.98	87.98	88.49	88.97	88.52	87.96	88.38
kbmemshrd	487916	487916	486792	474980	465808	457332	474566
kbuffers	13856	13856	13952	13968	13972	14100	13970
kbcached	93604	93604	94632	95860	96844	97748	95738
kbwpfree	518376	518376	518376	518376	518376	518376	518376
kbwpused	6912	6912	6912	6912	6912	6912	6912
%wpused	1.32	1.32	1.32	1.32	1.32	1.32	1.32



sar -R

sar -R -f sardata.out

Linux 2.2.5-15smp (jaqui.abc.com) 05/15/01

	12:34:44	12:34:46	12:34:48	12:34:50	12:34:52	12:34:54	Average:
frmpg/s							
shmpg/s							
bufpg/s							
campg/s							
12:34:46	-328.00	-131.50	2.00	96.00			
12:34:48	-329.00	-140.50	12.00	128.50			
12:34:50	-312.50	-1476.50	2.00	153.50			
12:34:52	288.00	-1146.50	0.50	123.00			
12:34:54	365.00	-1059.50	16.00	113.00			
Average:	-63.30	-790.90	6.50	122.80			



sar -w and -W

```
sar -w -f sardata.out
```

```
Linux 2.2.5-15smp (jaquil.abc.com) 05/15/01
```

```
12:34:44 cswch/s
12:34:46 630.50
12:34:48 760.50
12:34:50 833.50
12:34:52 723.00
12:34:54 651.50
Average: 719.80
```

```
sar -W -f sardata.out
```

```
Linux 2.2.5-15smp (jaquil.abc.com) 05/15/01
```

```
12:34:44 pswpin/s pswpout/s
12:34:46 0.00 0.00
12:34:48 0.00 0.00
12:34:50 0.00 0.00
12:34:52 0.00 0.00
12:34:54 0.00 0.00
Average: 0.00 0.00
```



Commands for I/O

- ◆ iostat
- ◆ sar
- ◆ /proc/sys/fs/*
- ◆ hdparm
- ◆ /proc/stat
- ◆ df



df

df

Filesystem	1k-blocks	Used	Available	Use%	Mounted on
/dev/sda1	127880	52265	69011	43%	/
/dev/sda8	12762679	9957815	2141620	82%	/home
/dev/sda5	2971111	2097239	720222	74%	/usr
/dev/sda6	762122	277769	444981	38%	/var

df

Filesystem	Size	Used	Avail	Use%	Mounted on
/dev/sda5	539M	155M	383M	29%	/
/dev/sdb6	21G	14G	7.5G	65%	/home
/dev/sda6	2.9G	1007M	1.9G	34%	/usr
/dev/sda9	1.9G	358M	1.6G	18%	/usr/local
/dev/sda8	992M	96M	896M	10%	/var
/dev/sdb7	19G	231M	19G	1%	/web



iostat

iostat 2 10

```
Linux 2.2.5-15smp (jaqui.abc.com) 05/15/01
```

avg-cpu:	%user	%nice	%sys	%idle		
	14.23	0.00	3.53	82.24		
Device:	tps	Blk_read/s	Blk_wrtn/s	Blk_read	Blk_wrtn	
hdisk0	6.71	6.54	6.92	48742592	51551898	
hdisk1	0.00	0.00	0.00	0	0	
hdisk2	0.00	0.00	0.00	0	0	
hdisk3	0.00	0.00	0.00	0	0	

avg-cpu:	%user	%nice	%sys	%idle		
	25.00	0.00	4.50	70.50		
Device:	tps	Blk_read/s	Blk_wrtn/s	Blk_read	Blk_wrtn	
hdisk0	5.50	11.00	0.00	22	0	
hdisk1	0.00	0.00	0.00	0	0	
hdisk2	0.00	0.00	0.00	0	0	
hdisk3	0.00	0.00	0.00	0	0	



sar I/O rates

```
sar -b -f sardata.out
```

```
Linux 2.2.5-15smp (jaqui.abc.com) 05/15/01
```

	tps	rtps	wtps	bread/s	bwrtn/s
12:34:44					
12:34:46	9.50	8.50	1.00	17.00	2.00
12:34:48	21.00	20.00	1.00	40.00	2.00
12:34:50	43.00	23.00	20.00	46.00	40.00
12:34:52	20.50	19.50	1.00	39.00	2.00
12:34:54	30.00	13.00	17.00	26.00	34.00
Average:	24.80	16.80	8.00	33.60	16.00



Circle4 Consulting

25

iostat rates

- ◆ Add kb/s for all disks by adapter
- ◆ Divide by 1024 to get mb/sec
- ◆ For <70% thrupt you should see
 - ◆ <3.5 scsi-1
 - ◆ <7 scsi-2
 - ◆ <14 scsi-2 FW
 - ◆ <21 scsi-2 FWD



Circle4 Consulting

26

/proc/sys/fs

- ◆ Max values can be set in /proc/sys/fs/*-max
- ◆ file_max
 - ◆ Set to 16384 – max files
- ◆ inode_max
 - ◆ Set to 65536 – max inodes
- ◆ hdparm
 - ◆ Use man hdparm
 - ◆ Sets filesystem read ahead count, drive write cache, drive lookahead prefetch count, drive lookahead



Circle4 Consulting

27

sar -v to show kernel values

```
sar -v -f sardata.out
```

```
Linux 2.2.5-15smp (jaqui.abc.com) 05/15/01
```

	dentunusd	file-sz	%file-sz	inode-sz	super-sz	%super-sz
12:34:44						
12:34:46	14545	266	6.49	12012	6	2.34
12:34:48	14545	269	6.57	12008	6	2.34
12:34:50	14545	262	6.40	12009	6	2.34
12:34:52	14545	276	6.74	12006	6	2.34
12:34:54	14545	285	6.96	12008	6	2.34
Average:	14545	272	6.62	12009	6	2.34

	dquot-sz	%dquot-sz	rtsig-sz	%rtsig-sz
12:34:46	0	0.00	0	0.00
12:34:48	0	0.00	0	0.00
12:34:50	0	0.00	0	0.00
12:34:52	0	0.00	0	0.00
12:34:54	0	0.00	0	0.00
Average:	0	0.00	0	0.00

Reset max values in /proc/sys/fs/*-max



Circle4 Consulting

28

lsof

- ◆ lsof is a 3rd party tool that you can use to look at open files and network connections
- ◆ `ftp://vic.cc.purdue.edu/pub/tools/unix/lsof/README`
- ◆ `cd /usr/local`
- ◆ `tar -xvf soft/tarfiles/lsof_4.51.tar`
- ◆ `cd lsof_4.51`
- ◆ `./Configure aix`
- ◆ ensure you set HASSECURITY to yes
- ◆ `make`
- ◆ `cp lsof /usr/local/bin`
- ◆ `cp lsof.8 /usr/local/man/man8`
- ◆ Try `lsof | grep TCP` or `lsof | grep UDP` or `lsof | more`



Networking

- ◆ `netstat`
- ◆ `nfsstat`
- ◆ `lsof | grep TCP`
- ◆ `lsof | grep UDP`



/proc/sys/net/core

- ◆ Critical on high speed networks such as z/390 hypersockets (24gb/sec) or gigabit ethernet
- ◆ TCP window sizes (2.1.90 and >)
 - ◆ `rmem_max` & `wmem_max`
 - ◆ Default is 64kb – can increase it
- ◆ Default Buffer Settings
 - ◆ Max socket buffer size 64kb
 - ◆ Default TCP socket buffer size 32kb
 - ◆ Default UDP socket buffer size 32kb?



netstat -v

Active Internet connections (w/o servers)

Proto	Recv-Q	Send-Q	Local Address	Foreign Address	State
tcp	0	40	jaquill11.abc.com:ssh	acscube9.abc.com:4215	ESTABLISHED
tcp	0	0	jaquill11.abc.com:15000	jaquill11.abc.com:4377	ESTABLISHED
tcp	0	23	jaquill11.abc.com:4377	jaquill11.abc.com:15000	ESTABLISHED
udp	0	112	jaquill11.abc.com:1031	chaos.abc.com domain	

Active UNIX domain sockets (w/o servers)

Proto	RefCnt	Flags	Type	State	I-Node	Path
unix	4	[]	DGRAM		566	/dev/log
unix	0	[]	DGRAM		728	
unix	0	[]	DGRAM		641	
unix	0	[]	DGRAM		626	
unix	1	[]	STREAM	CONNECTED	586	
unix	1	[]	STREAM	CONNECTED	585	
unix	0	[]	DGRAM		576	



netstat -s

IP & ICMP:

Tcp:

3821432 active connections openings
0 passive connection openings
73725 failed connection attempts
0 connection resets received
64 connections established
377615674 segments received
387277155 segments send out
1430638 segments retransmitted
17161 bad segments received.
1024316 resets sent

Udp:

55053 packets received
10482 packets to unknown port received.
0 packet receive errors
53753 packets sent

TcpExt:

11421 resets received for embryonic SYN_RECV sockets
2879471 packets pruned from receive queue because of socket buffer overrun
5324489 packets dropped from out-of-order queue because of socket buffer overrun
126 ICMP packets dropped because they were out-of-window



Circle4 Consulting

33

netstat -i

Kernel Interface table

Iface	MTU	Met	RX-OK	RX-ERR	RX-DRP	RX-OVR	TX-OK	TX-ERR	TX-DRP	TX-OVR	Flg
eth0	1500	0	94408545	0	0	0	99872636	0	0	0	BRU
lo	3924	0	287611603	0	0	0	287611603	0	0	0	LRU



Circle4 Consulting

34

nfsstat

- ◆ Client and Server NFS Info
- ◆ `nfsstat -cn, -r, -s`
 - ◆ retransmissions due to errors
 - ◆ badcalls
 - ◆ timeouts
 - ◆ waits
 - ◆ reads



Circle4 Consulting

35

Summary

- ◆ Linux has useful free performance tools
- ◆ Download iptraf and sysstat toolsets
- ◆ Find more information at:
 - ◆ <http://www.tunelinux.com>
 - ◆ <http://linuxperf.nl/linux.org>
 - ◆ http://www.psc.edu/networking/perf_tune.html
 - ◆ Enabling high performance data transfers on hosts
- ◆ See following slides for additional notes on issues such as raid, etc
- ◆ jaqui@circle4.com
- ◆ <http://www.circle4.com/jaqui/>



Circle4 Consulting

36

Disk Technologies

- ◆ Arbitrated
 - ◆ SCSI 20 or 40 mb/sec
 - ◆ FC-AL 100mb/sec
 - ◆ Devices arbitrate for exclusive control
 - ◆ SCSI priority based on address
- ◆ Non-Arbitrated
 - ◆ SSA 80 or 160mb/sec
 - ◆ Devices on loop all treated equally
 - ◆ Devices drop packets of data on loop



RAID Levels

- ◆ Raid-0
 - ◆ Disks combined into single volume stripeset
 - ◆ Data striped across the disks
- ◆ Raid-1
 - ◆ Every disk mirrored to another
 - ◆ Full redundancy of data but needs extra disks
 - ◆ At least 2 I/Os per random write
- ◆ Raid-0+1
 - ◆ Striped mirroring
 - ◆ Combines redundancy and performance



RAID Levels

- ◆ RAID-5
 - ◆ Data striped across a set of disks
 - ◆ 1 more disk used for parity bits
 - ◆ Parity may be striped across the disks also
 - ◆ At least 4 I/Os per random write (read/write to data and read/write to parity)
 - ◆ Uses hot spare technology



Intra and Inter Policies

- ◆ Intra Policy
 - ◆ How data is laid out on the disk
 - ◆ Outer edge, outer middle, center, inner middle, inner edge
- ◆ Inter Policy
 - ◆ How data is laid out between/across disks



Other I/O performance notes

- ◆ Mirroring of disks
- ◆ Mirror Write Consistency
- ◆ Write verify
- ◆ Logical volume scheduling policy
 - ◆ Parallel or sequential
- ◆ SCSI (1 disk at a time) vs scsi-2 fw diff
- ◆ RAID and SSA
- ◆ Buffering and two way searches
- ◆ I/O Pacing and Async I/O
- ◆ Striping

