



Linux kernels 2.2, 2.4, and 2.5 performance comparison

Duc Vianney, Ph. D.

IBM Linux Technology Center

LinuxWorld Expo 2002, San Francisco

August 13, 2002

IBM

Linux Technology
Center

Outline

- Motivation
 - ▶ Is 2.4 slower than 2.2?
- Kernel 2.4 features compared to 2.2
- Kernel 2.5 features compared to 2.4
- Benchmarks and measurements
- Test environments: UP, SMP-1, and SMP-4
- Direct performance comparison
- SMP kernel overhead comparison
- Scalability analysis
- Performance ratio analysis



Is 2.4 slower than 2.2?

On Wed, 2 May 2001, Duc Vianney wrote:

>

> Has anyone seen performance degradations between 2.2.19 and 2.4.x

Yes.

The signal handling one is because 2.4.x will save off the full SSE2 state, which means that the signal stack is almost 700 bytes, as compared to <200 before. This is sadly necessary to be able to take advantage of the SSE2 instructions - and on special applications the win can be quite noticeable. This one you won't be able to avoid, although you shouldn't see it on older hardware that do not have SSE2 (you see it because you have a PIII).

.....

2.4.4 will give noticeably better numbers for fork and fork+exec. However, the scheduling optimization that does that actually breaks at least "bash", and it appears that we will just undo it during the stable series. Even if the bug is obviously in user land (and a fix is available), stable kernels shouldn't try to hide the problems.

Kernel 2.4 features compared to 2.2

- Bigger size
- Support more architectures
- Resource management
- Support for enterprise level
- Shared memory
- New system file caching
- Filesystems
- Block devices
- Raw I/O device
- Networking

The IBM logo, consisting of the letters 'IBM' in a bold, sans-serif font, is positioned in the bottom left corner of the slide. The letters are white and set against a dark blue background.

Linux Technology
Center

Kernel 2.5 features compared to 2.4

- Rewrite block I/O, frame buffer layer, IDE subsystem
- New scheduler for improved scalability
- New kernel device structure (kdev_t)
- New driver model & unified device tree
- Preemptible kernel option
- Support for Next Generation Posix Threading (NGPT)
- Support for page tables in highmem
- Support new architectures (AMD and PowerPC 64-bit)
- Fast lightweight userspace semaphores
- NAPI network interrupt mitigation
- Fast walk dcache
- Direct pagecache
- VM reverse mappings

Benchmarks and measurements

■ Benchmarks

- ▶ Imbench (<http://sourceforge.net/projects/Imbench>)
- ▶ PTS (<http://sourceforge.net/projects/ptsbench>)

■ Hardware configuration

- ▶ SMP 4-way, 500 MHz, 2 9.5 GB SCSI drives, 2.5 GB RAM

■ Kernels

- ▶ Stock 2.2.19, 2.4.18, and 2.5.25 from <http://www.kernel.org>

■ Performance metric

- ▶ time in microseconds or throughput bandwidth

■ Performance ratio

- ▶ time based (or bandwidth)
- ▶ greater (or smaller) than 100% means numerator is slower (or faster)

■ Kernel under test

- ▶ UP, SMP-1, and SMP-4

Test environments: UP, SMP-1, and SMP-4

- Linux distribution
 - ▶ SUSE 7.3
- UP
 - ▶ stock kernel built as a uniprocessor kernel
 - ▶ configured and reboot with uniprocessor kernel
- SMP-1
 - ▶ stock kernel built as a multiprocessor kernel
 - ▶ configured as SMP-1 with "maxcpus=1" option in lilo.conf
- SMP4
 - ▶ stock kernel built as a multiprocessor kernel
 - ▶ configured as SMP-4 with "maxcpus=4" option in lilo.conf

The IBM logo, consisting of the letters 'IBM' in a bold, sans-serif font, is positioned in the bottom left corner of the slide. The letters are white and set against a dark blue background.

Linux Technology
Center

Direct performance comparison - UP

- Direct kernel performance comparison
 - ▶ UP vs. UP between 2.2.19, 2.4.18, and 2.5.25
- Performance metrics
 - ▶ Latency of simple functions
 - ▶ Interprocess communication latency
 - ▶ Context switching latency
 - ▶ Memory read rate
 - ▶ Filesystem

The IBM logo, consisting of the letters 'IBM' in a bold, sans-serif font, is positioned in the bottom left corner of the slide. The letters are white with a blue outline.

Linux Technology
Center

Latency of simple functions

- Time to perform a simple system function
- All times are in microseconds

	UP	UP	UP	Ratio	Ratio	Ratio
	2.2.19	2.4.18	2.5.25	2.4.18/2.2.19	2.5.25/2.2.19	2.5.25/2.4.18
Simple syscall	0.6112	0.5849	0.6332	96%	104%	108%
Simple read	0.96	0.9395	1.0888	98%	113%	116%
Simple write	0.8141	0.8231	0.9033	101%	111%	110%
Simple stat	5.7109	5.5176	5.909	97%	103%	107%
Simple fstat	1.1901	1.2959	1.5076	109%	127%	116%
Simple open/close	6.9306	7.4851	7.8082	108%	113%	104%
Signal handler installation	1.719	1.673	1.733	97%	101%	104%
Signal handler overhead	2.307	5.206	5.211	226%	226%	100%
Protection fault	1.245	1.262	1.278	101%	103%	101%
Pipe latency	6.5255	7.8873	8.5057	121%	130%	108%
AF_UNIX sock stream latency	13.646	13.514	14.173	99%	104%	105%
Process fork+exit	274.8	325.06	347.67	118%	127%	107%
Process fork+execve	1323.6	1431.5	1559.3	108%	118%	109%
Process fork+/bin/sh -c	64894	7242	7955	11%	12%	110%
Geometric mean	10.624	9.9677	10.706	94%	101%	107%

Interprocess communication latency

- "hot potato" benchmark - a token is passed back and forth between two processes
- Time interprocess communication latency via UDP/IP, TCP/IP, RPC/tcp, and RPC/udp using socket.
- Time interprocess connection latency via TCP/IP
- All tests using local host
- All times are in microseconds

	UP	UP	UP	Ratio	Ratio	Ratio
	2.2.19	2.4.18	2.5.25	2.4.18/2.2.19	2.5.25/2.2.19	2.5.25/2.4.18
UDP latency	24.40	28.11	29.08	115%	119%	103%
TCP latency	41.56	41.54	47.13	100%	113%	113%
RPC/tcp latency	97.04	106.75	111.18	110%	115%	104%
RPC/udp latency	68.19	77.95	80.06	114%	117%	103%
TCP/IP connection cost	140.64	158.19	160.51	112%	114%	101%
Geometric mean	62.37	68.76	72.17	110%	116%	105%

Context switching latency

- Processes are connected in a ring of Unix pipes. Each process reads a token from its pipe, possibly does some work, and then writes the token to the next process
- All times are in microseconds
- Data are geometric mean of context switching latency time of 2, 4, 8, 16, 32, 64 and 96 processes doing a summing work over a specific array size of KB

	UP	UP	UP	Ratio	Ratio	Ratio
Array size	2.2.19	2.4.18	2.5.25	2.4.18/2.2.19	2.5.25/2.2.19	2.5.25/2.4.18
0 KB	2.4398	2.5444	2.5195	104%	103%	99%
4 KB	5.2189	4.9444	4.9465	95%	95%	100%
8 KB	11.389	7.9481	7.8722	70%	69%	99%
16 KB	17.534	16.907	17.008	96%	97%	101%
32 KB	33.231	34.472	35.004	104%	105%	102%
64 KB	110.94	88.603	87.438	80%	79%	99%
Geometric mean	14.52	13.15	13.12	91%	90%	100%

Memory read rate

- Allocates the specified amount of memory, zeros it, and then times the reading of that memory as a series of integer loads and adds. Each four byte integer is loaded and added to accumulator.
- Data are geometric mean of memory size from 0.5MB to 256MB
- Bandwidth is in megabytes/sec
- Ratios are throughput - higher is better

	UP	UP	UP	Ratio	Ratio	Ratio
	2.2.19	2.4.18	2.5.25	2.4.18/2.2.19	2.5.25/2.2.19	2.5.25/2.4.18
Memory read	684.49	726.2	714.74	106%	104%	98%
Memory partial read	1491.3	1672.3	1608.2	112%	108%	96%
Memory write	448.08	465.28	469.12	104%	105%	101%
Memory partial write	660.27	679.08	688.99	103%	104%	101%
Memory partial read/write	573.54	589.05	593.64	103%	104%	101%
Memory bzero	577.08	581.71	573.79	101%	99%	99%
Geometric mean	681.24	713.09	708.56	105%	104%	99%

Filesystem

- Creating and deleting small files

- ▶ Number of file creations/sec and removals/sec
- ▶ Ratios are throughput - higher is better

File Size	UP 2.2.19	UP 2.4.18	UP 2.5.25	Ratio 2.4.18/2.2.19	Ratio 2.5.25/2.2.19	Ratio 2.5.25/2.4.18
Creations						
0k	7372	16326	8991	221%	122%	55%
1k	5554	9346	6138	168%	111%	66%
4k	5371	8641	5854	161%	109%	68%
10k	4297	5686	4324	132%	101%	76%
Geometric mean	5544	9305	6114	168%	110%	66%
Removals						
0k	91199	142187	108260	156%	119%	76%
1k	62270	61946	47950	99%	77%	77%
4k	62270	65011	48035	104%	77%	74%
10k	54801	45893	36866	84%	67%	80%
Geometric mean	66349	71598	55063	108%	83%	77%

Direct performance comparison - SMP-1

- Direct kernel performance comparison
 - ▶ SMP-1 vs. SMP-1 between 2.2.19, 2.4.18, and 2.5.25
- Performance metrics
 - ▶ Latency of simple functions
 - ▶ Interprocess communication latency
 - ▶ Context switching latency
 - ▶ Memory read rate
 - ▶ Filesystem
 - ▶ Command line execution time

The IBM logo, consisting of the letters 'IBM' in a bold, sans-serif font, with horizontal lines through the letters.

Linux Technology
Center

Latency of simple functions

- Time to perform a simple system function
- All times are in microseconds

	SMP-1	SMP-1	SMP-1	Ratio	Ratio	Ratio
	2.2.19	2.4.18	2.5.25	2.4.18/2.2.19	2.5.25/2.2.19	2.5.25/2.4.18
Simple syscall	0.641	0.6911	0.7104	108%	111%	103%
Simple read	1.0717	1.28	1.3908	119%	130%	109%
Simple write	1.0294	1.1384	1.2408	111%	121%	109%
Simple stat	5.8784	7.9122	6.9939	135%	119%	88%
Simple fstat	1.324	1.5526	1.6643	117%	126%	107%
Simple open/close	7.6028	10.018	9.1533	132%	120%	91%
Signal handler installation	1.775	1.765	1.811	99%	102%	103%
Signal handler overhead	2.654	5.431	5.761	205%	217%	106%
Protection fault	1.546	1.423	1.469	92%	95%	103%
Pipe latency	9.6865	11.225	11.551	116%	119%	103%
AF_UNIX sock stream latency	17.926	20.198	20.477	113%	114%	101%
Process fork+exit	305.33	370.4	399	121%	131%	108%
Process fork+execve	1401	1515.8	1694.5	108%	121%	112%
Process fork+/bin/sh -c	65820	7543	8269	11%	13%	110%
Geometric mean	12.128	12.203	12.636	101%	104%	104%

Interprocess communication latency

- "hot potato" benchmark - a token is passed back and forth between two processes
- Time interprocess communication latency via UDP/IP, TCP/IP, RPC/tcp, and RPC/udp using socket.
- Time interprocess connection latency via TCP/IP
- All tests using local host
- All times are in microseconds

	SMP-1	SMP-1	SMP-1	Ratio	Ratio	Ratio
	2.2.19	2.4.18	2.5.25	2.4.18/2.2.19	2.5.25/2.2.19	2.5.25/2.4.18
UDP latency	33.01	36.80	36.24	111%	110%	98%
TCP latency	52.46	52.27	54.61	100%	104%	104%
RPC/tcp latency	112.57	116.30	122.86	103%	109%	106%
RPC/udp latency	81.53	87.41	92.07	107%	113%	105%
TCP/IP connection co	170.47	184.62	191.05	108%	112%	103%
Geometric mean	77.02	81.57	84.38	106%	110%	103%

Context switching latency

- Processes are connected in a ring of Unix pipes. Each process reads a token from its pipe, possibly does some work, and then writes the token to the next process
- All times are in microseconds
- Data are geometric mean of context switching latency time of 2, 4, 8, 16, 32, 64 and 96 processes doing a summing work over a specific array size of KB

	SMP-1	SMP-1	SMP-1	Ratio	Ratio	Ratio
Array size	2.2.19	2.4.18	2.5.25	2.4.18/2.2.19	2.5.25/2.2.19	2.5.25/2.4.18
0 KB	3.1619	3.4631	3.4966	110%	111%	101%
4 KB	6.2551	5.8056	5.9542	93%	95%	103%
8 KB	8.5763	8.9936	9.2643	105%	108%	103%
16 KB	16.951	18.412	18.158	109%	107%	99%
32 KB	39.712	35.228	38.295	89%	96%	109%
64 KB	86.65	77.624	84.81	90%	98%	109%
Geometric mean	14.65	14.45	15.00	99%	102%	104%

Memory read rate

- Allocates the specified amount of memory, zeros it, and then times the reading of that memory as a series of integer loads and adds. Each four byte integer is loaded and added to accumulator.
- Data are geometric mean of memory size from 0.5MB to 256MB
- Bandwidth is in megabytes/sec
- Ratios are throughput - higher is better

	SMP-1	SMP-1	SMP-1	Ratio	Ratio	Ratio
	2.2.19	2.4.18	2.5.25	2.4.18/2.2.19	2.5.25/2.2.19	2.5.25/2.4.18
Memory read	688.31	730.04	692.45	106%	101%	95%
Memory partial read	1556.7	1690.7	1596.3	109%	103%	94%
Memory write	448.41	479.65	462.34	107%	103%	96%
Memory partial write	652.18	703.03	681.14	108%	104%	97%
Memory partial read/write	560.67	604.47	581.77	108%	104%	96%
Memory bzero	568.61	583.48	566.57	103%	100%	97%
Geometric mean	681.18	726.3	697.08	107%	102%	96%

Filesystem

- Creating and deleting small files

- ▶ Number of file creations/sec and removals/sec
- ▶ Ratios are throughput - higher is better

File Size	SMP-1 2.2.19	SMP-1 2.4.18	SMP-1 2.5.25	Ratio 2.4.18/2.2.19	Ratio 2.5.25/2.2.19	Ratio 2.5.25/2.4.18
Creations						
0k	7294	15756	8501	216%	117%	54%
1k	5523	8776	5768	159%	104%	66%
4k	5315	8205	5482	154%	103%	67%
10k	4249	5520	4024	130%	95%	73%
Geometric mean	5492	8896	5735	162%	104%	64%
Removals						
0k	87566	102501	88692	117%	101%	87%
1k	59446	44346	41021	75%	69%	93%
4k	59435	43435	40958	73%	69%	94%
10k	51951	30877	29528	59%	57%	96%
Geometric mean	63317	49689	45800	78%	72%	92%

Command line execution time

	SMP-1 2.2.19	SMP-1 2.4.18	SMP-1 2.5.25	Ratio 2.4.18/2.2.19	Ratio 2.5.25/2.2.19	Ratio 2.5.25/2.4.18
ar -_create_an_archive_file	6613.33	7123.32	7650	108%	116%	107%
ar -_output_contents_of_archive	3846.67	3980	4386.67	103%	114%	110%
ar -_add_files_to_end_of_archive	5493.34	5766.67	6133.34	105%	112%	106%
ar -_replace_a_file_in_archive	5213.34	5480	5870	105%	113%	107%
ar -_output_table_of_contents	4020	4203.33	4570	105%	114%	109%
ar -_Extract_a_file_in_archive	4606.67	4990	5493.33	108%	119%	110%
awk -_read_a_line	19809.5	20175.4	20975.6	102%	106%	104%
awk -_increment_a_variable	28453.3	28594.4	29813.3	100%	105%	104%
awk -_split_a_line_into_fields	58075.1	58403.8	60140.9	101%	104%	103%
basename -_get_basename	1656.67	1776.67	2000	107%	121%	113%
basename -_get_basename_using_suffix	1643.34	1773.33	1993.34	108%	121%	112%
cat -_display_a_large_file	12540	12826.7	13360	102%	107%	104%
cat -_display_file_w/_all_action_options	4726.66	5120	5366.67	108%	114%	105%
chgrp -_modify_group_	2583.34	2773.33	3073.34	107%	119%	111%
chmod -_modify_permissions_using_755	1853.34	1963.33	2170	106%	117%	111%
chmod -_modify_group_of_using_+x	1796.66	1963.33	2160.01	109%	120%	110%
chown -_modify_owner	2696.67	2893.32	3153.33	107%	117%	109%
cmp -_compare_not_=_files,_display_lines	5983.94	6066.66	6440	101%	108%	106%
cmp -_compare_not_=_files,_display_byte	2993.34	3140	3406.67	105%	114%	108%
cmp -_compare_=_files	3111.11	3266.66	3556.34	105%	114%	109%
compress -_compress_large_file,_forced	5441.77	5760	6066.66	106%	111%	105%
compress -_compress_C_file,force_to_std	72676.1	73602.2	76236.6	101%	105%	104%
cp -_copy_a_file_<100_bytes	1813.34	1986.67	2176.67	110%	120%	110%
cp -_copy_a_file_>25000_bytes	1906.66	2053.33	2306.66	108%	121%	112%
cp -_copy_a_file_>10000_bytes	2020	2160	2420	107%	120%	112%
cpio -_un-archive_a_file	2920	3150	3440	108%	118%	109%
cpio -_create_an_archive_file	3786.67	4206.67	4560	111%	120%	108%



Command line execution time (cont'd)

	SMP-1	SMP-1	SMP-1	Ratio	Ratio	Ratio
	2.2.19	2.4.18	2.5.25	2.4.18/2.2.19	2.5.25/2.2.19	2.5.25/2.4.18
csplit_-_split_w/_prefix	14160	15186.7	16040	107%	113%	106%
csplit_-_split_by_pattern	14193.3	15160	16033.3	107%	113%	106%
cut_-_cut_a_file_by_character	5653.34	5760	6093.33	102%	108%	106%
cut_-_cut_a_file_by_fields	5600	5713.34	6053.33	102%	108%	106%
dd_-_bs=2048,_conv=lcase	2250	2460	2720	109%	121%	111%
dd_-_with_bs=4096	2023.34	2236.66	2433.34	111%	120%	109%
dd_-_with_bs=512	2203.33	2470	2723.33	112%	124%	110%
dd_-_with_bs=65536	2060	2223.33	2486.67	108%	121%	112%
dd_-_bs=2048,_conv=swab	2196.67	2393.33	2670	109%	122%	112%
diff_-_differ_w/_embedded_blanks,_-be	4457.83	4653.34	4986.67	104%	112%	107%
diff_-_compare_2_small_files_(<52_lines)	1726.66	1860	2080	108%	120%	112%
diff_-_compare_large_files_=(10,000_l)	15555.6	16727.3	17925.9	108%	115%	107%
diff_-_compare_not_equal_files_(1300_l)	8313.33	8426.66	8906.67	101%	107%	106%
diff_-_compare_equal_files_(1300_lin)	8286.66	8400	8880	101%	107%	106%
dirname_-_get_the_directory_name	1640	1766.66	2003.33	108%	122%	113%
du_-_get_disk_usage_of_each_file	10566.7	10906.7	12226.7	103%	116%	112%
du_-_get_disk_usage_of_total_sum_only	6256.66	6300	6780	101%	108%	108%
du_-_get_disk_usage_of_each_dir/subdir	6607.15	6734.69	7142.85	102%	108%	106%
echo_-_a_message_(100_char)	1726.67	1797.5	2036.67	104%	118%	113%
egrep_-_search_for_extended_pattern_#1	2520	2673.34	2933.34	106%	116%	110%
egrep_-_search_for_extended_pattern_#2	3673.34	3823.33	4140	104%	113%	108%
egrep_-_search_for_simple_pattern	3070	3333.33	3520	109%	115%	106%
ex_-_add_delete_write_quit_a_large_file	86666.7	89166.7	93809.5	103%	108%	105%
fgrep_-_search_for_"confidence"_pattern	2153.33	2280	2553.34	106%	119%	112%
fgrep_-_search_for_"the"_pattern	2513.34	2653.33	2900	106%	115%	109%



Command line execution time (cont'd)

	SMP-1	SMP-1	SMP-1	Ratio	Ratio	Ratio
	2.2.19	2.4.18	2.5.25	2.4.18/2.2.19	2.5.25/2.2.19	2.5.25/2.4.18
find -_find_a_modified_file_(stat'ed)	7100	7003.33	7573.33	99%	107%	108%
find -_find_a_file_(not_stat'ed)	5820	5810	6300	100%	108%	108%
grep -_search_a_large_file_suppress_msg	6473.34	6806.67	7126.67	105%	110%	105%
grep -_search_a_large_file	6473.33	6826.66	7106.67	105%	110%	104%
grep -_search_for_an_extended_pattern	2973.34	3100	3380	104%	114%	109%
grep -_search_for_a_specific_pattern	2353.33	2496.67	2763.34	106%	117%	111%
grep -_search_a_large_file_for_labian	81666.7	82222.2	83333.3	101%	102%	101%
head -_display_first_31+_lines	1793.33	1900	2140	106%	119%	113%
head -_display_first_51+_lines	1796.67	1923.33	2150	107%	120%	112%
head -_display_66+_lines	1790	1913.33	2153.34	107%	120%	113%
join -_join_files_with_formated_output	3180	3280	3573.33	103%	112%	109%
join -_join_files_(regular)	3206.67	3323.34	3540	104%	110%	107%
ld -_link_10_object_files	51233.3	52208.8	54766.7	102%	107%	105%
ls -_format_ls_output_(cpu_time)	63734.9	6010	6306.67	9%	10%	105%
ls -IR	91393.9	33253.3	34016.1	36%	37%	102%
ls -_list_with_defaults	61971.8	4300	4686.66	7%	8%	109%
m4 -_test_predefined_macros	7466.66	7550.2	8066.67	101%	108%	107%
mail -_send_a_message	3654.61	5060.25	5893.34	138%	161%	116%
make -_compile_C_code	11826.7	11760	12680	99%	107%	108%
mkdir -_make_a_directory	1940	2016.67	2280	104%	118%	113%
more -_display_a_file	8995.99	9186.66	9658.64	102%	107%	105%
nl -_number_all_lines	103037	103333	106533	100%	103%	103%
nl -_number_lines_in_a_file_(defaults)	102370	102296	105704	100%	103%	103%
nroff -_text_proccess_a_file	1373333	1314444	1355000	96%	99%	103%
paste -_paste_horizontally	2553.34	2700	2953.33	106%	116%	109%
paste -_paste_vertically	2580	2700	2936.67	105%	114%	109%



Command line execution time (cont'd)

	SMP-1 2.2.19	SMP-1 2.4.18	SMP-1 2.5.25	Ratio 2.4.18/2.2.19	Ratio 2.5.25/2.2.19	Ratio 2.5.25/2.4.18
pr -_paginate_a_file_(columns)	10086.7	10233.3	10746.7	101%	107%	105%
pr -_paginate_with_header	12760	12786.7	13446.7	100%	105%	105%
pr -_paginate_a_file	12600	12773.3	13386.7	101%	106%	105%
ps -_report_process_status	87317.1	32206.6	32610.4	37%	37%	101%
pwd -_print_working_directory	1710	1813.34	2066.67	106%	121%	114%
rm -_remove_a_large_file	1800.01	1753.33	1970	97%	109%	112%
rm -_remove_a_set_of_files	2286.67	2233.34	2533.34	98%	111%	113%
rmdir -_remove_a_directory	1760	1813.34	1926.67	103%	109%	106%
sdiff -_display_different_lines	7693.33	7873.33	8426.67	102%	110%	107%
sdiff -_display_identical_lines	7700	7900	8426.67	103%	109%	107%
sed -_edit_a_file_using_streams	52849.5	52891.6	54377.5	100%	103%	103%
sh -_execute_a_complex_shell_script	788667	273333	300000	35%	38%	110%
sort -_sort_a_large_file_(10,000_lines)	65211.3	65662.7	67365.6	101%	103%	103%
split -_split_a_file_into_50-line_seg	7293.33	8393.33	9226.67	115%	127%	110%
split -_split_a_file_using_defaults	5963.85	6266.66	6500	105%	109%	104%
sum -_display_checksum	5746.67	5826.67	6146.66	101%	107%	105%
tail -_display_file_8000_byte_from_EOF	1860	1980	2206.67	106%	119%	111%
tail -_display_end_of_file	1833.34	1933.33	2170	105%	118%	112%
tar -_create_a_tar_file	47500	53333.3	55625	112%	117%	104%
tar -_display_archive_table_of_contents	45261	46066.7	47933.3	102%	106%	104%
tee -_append_to_output_file	2256.67	2343.33	2374.67	104%	105%	101%
tee -_display_and_write_file	2100	2276.66	2533.33	108%	121%	111%
tr -_replace_all_numbers_with_#	1853.34	1956.66	2186.66	106%	118%	112%
tr -_replace_A-E_with_1-5	10346.7	10640	11080	103%	107%	104%
tr -_translate_lowercase_to_uppercase	1860	1970	2190	106%	118%	111%
tr -_translate_{}_to_()	1846.67	1986.67	2173.33	108%	118%	109%



Command line execution time (cont'd)

	SMP-1	SMP-1	SMP-1	Ratio	Ratio	Ratio
	2.2.19	2.4.18	2.5.25	2.4.18/2.2.19	2.5.25/2.2.19	2.5.25/2.4.18
uncompress -_force_uncompression	5880	6053.33	6340	103%	108%	105%
vi_-_perform_page_forwards/backwards	328333	326667	336667	99%	103%	103%
vi_-_quit,_no_save,_a_large_file	58750	70000	65000	119%	111%	93%
vi_-_edit_a_small_file	143750	136667	131667	95%	92%	96%
vi_-_write,_quit_a_large_file	48333.3	47500	49166.7	98%	102%	104%
wc_-_count_number_of_char_in_a_file	1796.67	1933.33	2153.34	108%	120%	111%
wc_-_count_number_of_lines_in_a_file	2086.67	2176.67	2466.66	104%	118%	113%
wc_-_count_number_of_words_in_a_file	2826.66	2946.66	3226.66	104%	114%	110%
wc_-_count_lines,_words,_and_char	2826.67	2960	3213.33	105%	114%	109%
yacc_-_generate_a_parsing_program	5160.64	5421.68	5813.33	105%	113%	107%
Geometric Mean				97%	105%	108%



Direct performance comparison - SMP-4

- Direct kernel performance comparison
 - ▶ SMP-4 vs. SMP-4 between 2.2.19, 2.4.18, and 2.5.25
- Performance metrics
 - ▶ Latency of simple functions
 - ▶ Interprocess communication latency
 - ▶ Context switching latency
 - ▶ Memory read rate
 - ▶ Filesystem
 - ▶ Command line execution time

The IBM logo, consisting of the letters 'IBM' in a bold, sans-serif font, is positioned in the bottom left corner of the slide. The letters are white with a blue outline.

Linux Technology
Center

Latency of simple functions

- Time to perform a simple system function
- All times are in microseconds

	SMP-4 2.2.19	SMP-4 2.4.18	SMP-4 2.5.25	Ratio 2.4.18/2.2.19	Ratio 2.5.25/2.2.19	Ratio 2.5.25/2.4.18
Simple syscall	0.6764	0.6914	0.6957	102%	103%	101%
Simple read	1.1347	1.3048	1.4241	115%	126%	109%
Simple write	1.0298	1.146	1.2331	111%	120%	108%
Simple stat	6.1147	7.9797	7.2249	131%	118%	91%
Simple fstat	1.3636	1.5506	1.7197	114%	126%	111%
Simple open/close	7.8668	9.9602	9.6373	127%	123%	97%
Signal handler installation	1.773	1.727	1.77	97%	100%	102%
Signal handler overhead	2.663	5.42	5.661	204%	213%	104%
Protection fault	1.548	1.451	1.462	94%	94%	101%
Pipe latency	9.6916	13.4378	11.998	139%	124%	89%
AF_UNIX sock stream latency	18.6465	33.5746	20.767	180%	111%	62%
Process fork+exit	419.615	520.455	394.92	124%	94%	76%
Process fork+execve	1549.5	1726	1668.3	111%	108%	97%
Process fork+/bin/sh -c	66320	8126	8184	12%	12%	101%
Geometric mean	12.7355	13.348	12.712	105%	100%	95%

Interprocess communication latency

- "hot potato" benchmark - a token is passed back and forth between two processes
- Time interprocess communication latency via UDP/IP, TCP/IP, RPC/tcp, and RPC/udp using socket.
- Time interprocess connection latency via TCP/IP
- All tests using local host
- All times are in microseconds

	SMP-4	SMP-4	SMP-4	Ratio	Ratio	Ratio
	2.2.19	2.4.18	2.5.25	2.4.18/2.2.19	2.5.25/2.2.19	2.5.25/2.4.18
UDP latency	51.92	52.92	35.25	102%	68%	67%
TCP latency	63.53	63.59	54.76	100%	86%	86%
RPC/tcp latency	114.37	114.21	120.28	100%	105%	105%
RPC/udp latency	88.44	90.26	86.23	102%	97%	96%
TCP/IP connection c	123.78	123.88	128.91	100%	104%	104%
Geometric mean	83.79	84.46	76.27	101%	91%	90%



Context switching latency

- Processes are connected in a ring of Unix pipes. Each process reads a token from its pipe, possibly does some work, and then writes the token to the next process
- All times are in microseconds
- Data are geometric mean of context switching latency time of 2, 4, 8, 16, 32, 64 and 96 processes doing a summing work over a specific array size of KB

Array size	SMP-4 2.2.19	SMP-4 2.4.18	SMP-4 2.5.25	Ratio 2.4.18/2.2.19	Ratio 2.5.25/2.2.19	Ratio 2.5.25/2.4.18
0 KB	3.12755	4.55502	3.519	146%	113%	77%
4 KB	5.43843	7.13491	5.7931	131%	107%	81%
8 KB	8.60602	9.96508	11.269	116%	131%	113%
16 KB	18.5262	18.0686	17.322	98%	94%	96%
32 KB	41.0297	29.9279	30.038	73%	73%	100%
64 KB	80.6251	62.0577	68.362	77%	85%	110%
Geometric mean	14.41	14.88	14.19	103%	98%	95%



Memory read rate

- Allocates the specified amount of memory, zeros it, and then times the reading of that memory as a series of integer loads and adds. Each four byte integer is loaded and added to accumulator.
- Data are geometric mean of memory size from 0.5MB to 256MB
- Bandwidth is in megabytes/sec
- Ratios are throughput - higher is better

	SMP-4	SMP-4	SMP-4	Ratio	Ratio	Ratio
	2.2.19	2.4.18	2.5.25	2.4.18/2.2.19	2.5.25/2.2.19	2.5.25/2.4.18
Memory read	702.263	729.161	696.23	104%	99%	95%
Memory partial read	1580.65	1683.76	1650.9	107%	104%	98%
Memory write	460.165	479.051	471.06	104%	102%	98%
Memory partial write	679.362	700.944	691.59	103%	102%	99%
Memory partial read/write	582.988	603.615	589.57	104%	101%	98%
Memory bzero	584.166	583.487	578.88	100%	99%	99%
Geometric mean	700.526	724.981	709.73	103%	101%	98%

Filesystem

- Creating and deleting small files

- ▶ Number of file creations/sec and removals/sec
- ▶ Ratios are throughput - higher is better

File Size	SMP-4 2.2.19	SMP-4 2.4.18	SMP-4 2.5.25	Ratio 2.4.18/2.2.19	Ratio 2.5.25/2.2.19	Ratio 2.5.25/2.4.18
Creations						
0k	15507	15803	8685	102%	56%	55%
1k	8890	8794	5885	99%	66%	67%
4k	8240	8197	5602	99%	68%	68%
10k	5422	5542	4093	102%	75%	74%
Geometric mean	8859	8914	5851	101%	66%	66%
Removals						
0k	114824	102881	90163	90%	79%	88%
1k	48061	43191	41785	90%	87%	97%
4k	47813	43094	41801	90%	87%	97%
10k	33022	30589	30286	93%	92%	99%
Geometric mean	54331	49196	46733	91%	86%	95%

Command line execution time

	SMP-4 2.2.19	SMP-4 2.4.18	SMP-4 2.5.25	Ratio 2.4.18/2.2.19	Ratio 2.5.25/2.2.19	Ratio 2.5.25/2.4.18
ar -_create_an_archive_file	7680	7550	7553.34	98%	98%	100%
ar -_output_contents_of_archive	4220	4460	4300	106%	102%	96%
ar -_add_files_to_end_of_archive	6060	6101.3	6066.66	101%	100%	99%
ar -_replace_a_file_in_archive	5893	5186.7	5816.66	88%	99%	112%
ar -_output_table_of_contents	4383	4663.3	4543.33	106%	104%	97%
ar -_Extract_a_file_in_archive	5133	5500	5440	107%	106%	99%
awk -_read_a_line	20000	20488	20325.2	102%	102%	99%
awk -_increment_a_variable	28695	28896	28960	101%	101%	100%
awk -_split_a_line_into_fields	58400	58545	58732.4	100%	101%	100%
basename -_get_basename	1867	2100	1956.67	113%	105%	93%
basename -_get_basename_using_suffix	1887	2020	1940	107%	103%	96%
cat -_display_a_large_file	13433	13587	13193.3	101%	98%	97%
cat -_display_file_w/_all_action_options	5440	5660	5406.67	104%	99%	96%
chgrp -_modify_group_	3060	2940	3026.68	96%	99%	103%
chmod -_modify_permissions_using_755	2020	2320	2033.34	115%	101%	88%
chmod -_modify_group_of_using_+x	2137	2345	2260	110%	106%	96%
chown -_modify_owner	2885	3260	3066.67	113%	106%	94%
cmp -_compare_not_=_files,_display_lines	6213	6425.7	6360	103%	102%	99%
cmp -_compare_not_=_files,_display_byte	3277	3475	3340	106%	102%	96%
compress -_compress_large_file,_forced	6260	6566.3	6000	105%	96%	91%
compress -_compress_C_file,force_to_stdout	73944	75430	74032.3	102%	100%	98%
cp -_copy_a_file_<100_bytes	2053	2330	2123.34	114%	103%	91%
cp -_copy_a_file_>25000_bytes	2187	3376	2253.33	154%	103%	67%
cp -_copy_a_file_>10000_bytes	2273	2417.5	2386.66	106%	105%	99%
cpio -_un-archive_a_file	3217	3546.7	3386.67	110%	105%	95%
cpio -_create_an_archive_file	4247	4703.3	4433.33	111%	104%	94%



Command line execution time (cont'd)

	SMP-4 2.2.19	SMP-4 2.4.18	SMP-4 2.5.25	Ratio 2.4.18/2.2.19	Ratio 2.5.25/2.2.19	Ratio 2.5.25/2.4.18
csplit -_split_w/_prefix	15033	15920	15613.3	106%	104%	98%
csplit -_split_by_pattern	14907	15653	15560	105%	104%	99%
cut -_cut_a_file_by_character	5870	6133.3	5953.33	104%	101%	97%
cut -_cut_a_file_by_fields	5860	6013.3	5920	103%	101%	98%
dd -_bs=2048,_conv=lcase	2553	2675	2673.34	105%	105%	100%
dd -_with_bs=4096	2373	2387.5	2386.67	101%	101%	100%
dd -_with_bs=512	2470	2890	2660	117%	108%	92%
dd -_with_bs=65536	2428	2536.7	2450	104%	101%	97%
dd -_bs=2048,_conv=swab	2460	2586.7	2620	105%	107%	101%
diff -_differ_w/_embedded_blanks,_-be	5033	4880	4880	97%	97%	100%
diff -_compare_2_small_files_(<52 _lines)	2020	2233.3	2046.67	111%	101%	92%
diff -_compare_large_files_=(10,000_l)	16893	16800	18048.8	99%	107%	107%
diff -_compare_not_equal_files_(1300_l)	8833	9280	8740	105%	99%	94%
diff -_compare_equal_files_(1300_lin)	8800	9133.3	8706.66	104%	99%	95%
dirname -_get_the_directory_name	1893	2035	1970	107%	104%	97%
du -_get_disk_usage_of_each_file	10960	11433	11893.3	104%	109%	104%
du -_get_disk_usage_of_total_sum_only	6807	6943.3	6646.67	102%	98%	96%
echo -_a_message_(100_char)	1950	2053.3	2003.33	105%	103%	98%
egrep -_search_for_extended_pattern_#1	2830	2993.3	2883.33	106%	102%	96%
egrep -_search_for_extended_pattern_#2	4033	4220	4053.34	105%	100%	96%
egrep -_search_for_simple_pattern	3393	3610	3440	106%	101%	95%
ex -_add_delete_write_quit_a_large_file	90526	90000	91428.6	99%	101%	102%
fgrep -_search_for_"confidence"_pattern	2387	2503.3	2483.33	105%	104%	99%
fgrep -_search_for_"the"_pattern	2837	2966.7	2843.33	105%	100%	96%



Command line execution time (cont'd)

	SMP-4 2.2.19	SMP-4 2.4.18	SMP-4 2.5.25	Ratio 2.4.18/2.2.19	Ratio 2.5.25/2.2.19	Ratio 2.5.25/2.4.18
find -_find_a_modified_file_(stat'ed)	7600	7513.3	7440	99%	98%	99%
find -_find_a_file_(not_stat'ed)	6200	6266.7	6160	101%	99%	98%
grep -_search_a_large_file_suppress_msgs	6940	7173.3	6936.66	103%	100%	97%
grep -_search_a_large_file	6933	7240	6933.34	104%	100%	96%
grep -_search_for_an_extended_pattern	3277	3440	3313.33	105%	101%	96%
grep -_search_for_a_specific_pattern	2647	2810	2710	106%	102%	96%
grep -_search_a_large_file_for_laban	80538	83333	83333.3	103%	103%	100%
head -_display_first_31+_lines	2007	2170	2100	108%	105%	97%
head -_display_first_51+_lines	2033	2230	2120	110%	104%	95%
head -_display_66+_lines	2060	2270	2123.33	110%	103%	94%
join -_join_files_with_formatted_output	3407	3666.7	3513.34	108%	103%	96%
ld -_link_10_object_files	53033	53567	53800	101%	101%	100%
ls -_format_ls_output_(cpu_time)	64272	6593.3	6180	10%	10%	94%
ls -IR	91697	34036	33232.9	37%	36%	98%
ls -_list_with_defaults	62048	4690	4600	8%	7%	98%
sdiff -_display_different_lines	8513	8253.3	8240	97%	97%	100%
sdiff -_display_identical_lines	8207	8266.7	8240	101%	100%	100%
sed -_edit_a_file_using_streams	53000	52958	53145.6	100%	100%	100%
sh -_execute_a_complex_shell_script	834444	325833	294167	39%	35%	90%
sort -_sort_a_large_file_(10,000_lines)	66854	66667	66182.8	100%	99%	99%
split -_split_a_file_into_50-line_seg	7687	9053.3	9080	118%	118%	100%
split -_split_a_file_using_defaults	6313	6675	6465.86	106%	102%	97%
sum -_display_checksum	5953	6125	6013.33	103%	101%	98%
tail -_display_file_8000_byte_from_EOF	2100	2430	2156.66	116%	103%	89%
tail -_display_end_of_file	2052	2233.3	2120	109%	103%	95%
tar -_create_a_tar_file	48500	52000	54375	107%	112%	105%
tar -_display_archive_table_of_contents	45422	46900	46933.3	103%	103%	100%



Command line execution time (cont'd)

	SMP-4 2.2.19	SMP-4 2.4.18	SMP-4 2.5.25	Ratio 2.4.18/2.2.19	Ratio 2.5.25/2.2.19	Ratio 2.5.25/2.4.18
tee -_append_to_output_file	2517	2330	2560	93%	102%	110%
tee -_display_and_write_file	2406	2565	2493.33	107%	104%	97%
tr -_replace_all_numbers_with_#	2070	2333.3	2140	113%	103%	92%
tr -_replace_A-E_with_1-5	11060	11040	10826.7	100%	98%	98%
tr -_translate_lowercase_to_uppercase	2090	2363.3	2143.33	113%	103%	91%
tr -_translate_{}_to_()	2067	2380	2133.33	115%	103%	90%
uncompress -_force_uncompression	6273	6473.3	6220	103%	99%	96%
vi -_perform_page_forwards/backwards	321667	325000	330000	101%	103%	102%
vi -_quit,_no_save,_a_large_file	58750	62500	65000	106%	111%	104%
vi -_edit_a_small_file	125000	133000	128333	106%	103%	96%
vi -_write,_quit_a_large_file	48750	50000	48750	103%	100%	98%
wc -_count_number_of_char_in_a_file	2060	2286.7	2126.66	111%	103%	93%
wc -_count_number_of_lines_in_a_file	2435	2472.5	2433.34	102%	100%	98%
wc -_count_number_of_words_in_a_file	3090	3316.7	3160	107%	102%	95%
wc -_count_lines,_words,_and_char	3057	3186.7	3180	104%	104%	100%
Geometric Mean				98%	95%	97%



SMP overhead comparison

- SMP kernel support overhead comparison
 - ▶ UP vs. SMP-1
 - ▶ 2.2.19, 2.4.18, and 2.5.25
- Performance metrics
 - ▶ Latency of simple functions
 - ▶ Interprocess communication latency
 - ▶ Context switching latency
 - ▶ Filesystem

The IBM logo, consisting of the letters 'IBM' in a bold, sans-serif font, is positioned in the bottom left corner of the slide. The background of the slide features a vertical strip on the left side with a blue-to-white gradient, containing faint images of a computer mouse, a globe, and a keyboard.

Linux Technology
Center

Latency of simple functions

- Time to perform a simple system function
- All times are in microseconds

	2.2.19 Ratio SMP-1/UP	2.4.18 Ratio SMP-1/UP	2.5.25 Ratio SMP-1/UP
Simple syscall	105%	118%	112%
Simple read	112%	136%	128%
Simple write	126%	138%	137%
Simple stat	103%	143%	118%
Simple fstat	111%	120%	110%
Simple open/close	110%	134%	117%
Signal handler installation	103%	105%	105%
Signal handler overhead	115%	104%	111%
Protection fault	124%	113%	115%
Pipe latency	148%	142%	136%
AF_UNIX sock stream latency	131%	149%	144%
Process fork+exit	111%	114%	115%
Process fork+execve	106%	106%	109%
Process fork+/bin/sh -c	101%	104%	104%
Geometric mean	114%	122%	118%

Interprocess communication latency

- "hot potato" benchmark - a token is passed back and forth between two processes
- Time interprocess communication latency via UDP/IP, TCP/IP, RPC/tcp, and RPC/udp using socket.
- Time interprocess connection latency via TCP/IP
- All tests using local host
- All times are in microseconds

	2.2.19 Ratio SMP-1/UP	2.4.18 Ratio SMP-1/UP	2.5.25 Ratio SMP-1/UP
UDP latency using localhost	135%	131%	125%
TCP latency using localhost	126%	126%	116%
RPC/tcp latency using localhost	116%	109%	110%
RPC/udp latency using localhost	120%	112%	115%
TCP/IP connection cost to localhost	121%	117%	119%
Geometric mean	123%	119%	117%

Context switching latency

- Processes are connected in a ring of Unix pipes. Each process reads a token from its pipe, possibly does some work, and then writes the token to the next process
- All times are in microseconds
- Data are geometric mean of context switching latency time of 2, 4, 8, 16, 32, 64 and 96 processes doing a summing work over a specific array size of KB

	2.2.19 Ratio SMP-1/UP	2.4.18 Ratio SMP-1/UP	2.5.25 Ratio SMP-1/UP
Array size=0 KB			
2	143%	135%	137%
4	161%	152%	142%
8	129%	155%	136%
16	131%	144%	137%
24	133%	138%	139%
32	121%	135%	141%
64	116%	120%	140%
96	110%	115%	137%
Geometric mean	130%	136%	139%

Filesystem

- Creating and deleting small files
 - ▶ Number of file creations/sec and removals/sec
 - ▶ Ratios are throughput - higher is better

	2.2.19 Ratio SMP-1/UP	2.4.18 Ratio SMP-1/UP	2.5.25 Ratio SMP-1/UP
Creations			
0k	99%	97%	95%
1k	99%	94%	94%
4k	99%	95%	94%
10k	99%	97%	93%
Geometric mean	99%	96%	94%
Removals			
0k	96%	72%	82%
1k	95%	72%	86%
4k	95%	67%	85%
10k	95%	67%	80%
Geometric mean	95%	69%	83%

Scalability analysis

- Which kernel provides better performance scalability
- Performance ratio between SMP-4 over SMP-1 on
 - ▶ 2.2.19
 - ▶ 2.4.18
 - ▶ 2.5.25
- Performance metrics
 - ▶ Latency of simple functions
 - ▶ Interprocess communication latency
 - ▶ Context switching latency
 - ▶ Filesystem
 - ▶ Command line execution time

The IBM logo, consisting of the letters 'IBM' in a bold, sans-serif font, is positioned in the bottom left corner of the slide. The background of the slide features a vertical strip on the left side with a blue-to-white gradient, containing faint images of a computer mouse, a globe, and a document with a waveform.

Linux Technology
Center

Latency of simple functions

- Time to perform a simple system function
- All times are in microseconds

	2.2.19	2.4.18	2.5.25
	SMP-4/SMP-1	SMP-4/SMP-1	SMP-4/SMP-1
Simple syscall	106%	100%	98%
Simple read	106%	102%	102%
Simple write	100%	101%	99%
Simple stat	104%	101%	103%
Simple fstat	103%	100%	103%
Simple open/close	103%	99%	105%
Signal handler installation	100%	98%	98%
Signal handler overhead	100%	100%	98%
Protection fault	100%	102%	100%
Pipe latency	100%	120%	104%
AF_UNIX sock stream latency	104%	166%	101%
Process fork+exit	137%	141%	99%
Process fork+execve	111%	114%	98%
Process fork+/bin/sh -c	101%	108%	99%
Geometric mean	105%	109%	101%

Interprocess communication latency

- "hot potato" benchmark - a token is passed back and forth between two processes
- Time interprocess communication latency via UDP/IP, TCP/IP, RPC/tcp, and RPC/udp using socket.
- Time interprocess connection latency via TCP/IP
- All tests using local host
- All times are in microseconds

	2.2.19	2.4.18	2.5.25
	SMP-4/SMP-1	SMP-4/SMP-1	SMP-4/SMP-1
UDP latency	157%	144%	97%
TCP latency	121%	122%	100%
RPC/tcp latency	102%	98%	98%
RPC/udp latency	108%	103%	94%
TCP/IP connection cost	73%	67%	67%
Geometric mean	109%	104%	90%



Context switching latency

- Processes are connected in a ring of Unix pipes. Each process reads a token from its pipe, possibly does some work, and then writes the token to the next process
- All times are in microseconds
- Data are geometric mean of context switching latency time of 2, 4, 8, 16, 32, 64 and 96 processes doing a summing work over a specific array size of KB

	2.2.19 SMP-4/SMP-1	2.4.18 SMP-4/SMP-1	2.5.25 SMP-4/SMP-1
Array size			
0 KB	99%	132%	101%
4 KB	87%	123%	97%
8 KB	100%	111%	122%
16 KB	109%	98%	95%
32 KB	103%	85%	78%
64 KB	93%	80%	81%
Geometric mean	98%	103%	95%

Filesystem

- Creating and deleting small files
- Number of file creations/sec and removals/sec
- Ratios are throughput - higher is better

	2.2.19 SMP-4/SMP-1	2.4.18 SMP-4/SMP-1	2.5.25 SMP-4/SMP-1
Creations			
0k	213%	100%	102%
1k	161%	100%	102%
4k	155%	100%	102%
10k	128%	100%	102%
Geometric mean	161%	100%	102%
Removals			
0k	131%	100%	102%
1k	81%	97%	102%
4k	80%	99%	102%
10k	64%	99%	103%
Geometric mean	86%	99%	102%

Command line execution time

	2.2.19	2.4.18	2.5.25
	SMP-4/SMP-1	SMP-4/SMP-1	SMP-4/SMP-1
ar -_create_an_archive_file	116%	106%	99%
ar -_output_contents_of_archive	110%	112%	98%
ar -_add_files_to_end_of_archive	110%	106%	99%
ar -_replace_a_file_in_archive	113%	95%	99%
ar -_output_table_of_contents	109%	111%	99%
ar -_Extract_a_file_in_archive	111%	110%	99%
awk -_read_a_line	101%	102%	97%
awk -_increment_a_variable	101%	101%	97%
awk -_split_a_line_into_fields	101%	100%	98%
basename -_get_basename	113%	118%	98%
basename -_get_basename_using_suffix	115%	114%	97%
cat -_display_a_large_file	107%	106%	99%
cat -_display_file_w/_all_action_options	115%	111%	101%
chgrp -_modify_group_	118%	106%	98%
chmod -_modify_permissions_using_755	109%	118%	94%
chmod -_modify_group_of_using_+x	119%	119%	105%
chown -_modify_owner	107%	113%	97%
cmp -_compare_not_=_files,_display_lines	104%	106%	99%
cmp -_compare_not_=_files,_display_byte	109%	111%	98%
compress -_compress_large_file,_forced	115%	114%	99%
compress -_compress_C_file,force_to_stdout	102%	102%	97%
cp -_copy_a_file_<100_bytes	113%	117%	98%
cp -_copy_a_file_>25000_bytes	115%	164%	98%
cp -_copy_a_file_>10000_bytes	113%	112%	99%
cpio -_un-archive_a_file	110%	113%	98%
cpio -_create_an_archive_file	112%	112%	97%

Command line execution time (cont'd)

	2.2.19	2.4.18	2.5.25
	SMP-4/SMP-1	SMP-4/SMP-1	SMP-4/SMP-1
csplit -_split_w/_prefix	106%	105%	97%
csplit -_split_by_pattern	105%	103%	97%
cut -_cut_a_file_by_character	104%	106%	98%
cut -_cut_a_file_by_fields	105%	105%	98%
dd -_bs=2048,_conv=lcase	113%	109%	98%
dd -_with_bs=4096	117%	107%	98%
dd -_with_bs=512	112%	117%	98%
dd -_with_bs=65536	118%	114%	99%
dd -_bs=2048,_conv=swab	112%	108%	98%
diff -_differ_w/_embedded_blanks,_-be	113%	105%	98%
diff -_compare_2_small_files_(<52_lines)	117%	120%	98%
diff -_compare_large_files_=(10,000_l)	109%	100%	101%
diff -_compare_not_equal_files_(1300_l)	106%	110%	98%
diff -_compare_equal_files_(1300_lin)	106%	109%	98%
dirname -_get_the_directory_name	115%	115%	98%
du -_get_disk_usage_of_each_file	104%	105%	97%
du -_get_disk_usage_of_total_sum_only	109%	110%	98%
echo -_a_message_(100_char)	113%	114%	98%
egrep -_search_for_extended_pattern_#1	112%	112%	98%
egrep -_search_for_extended_pattern_#2	110%	110%	98%
egrep -_search_for_simple_pattern	111%	108%	98%
ex -_add_delete_write_quit_a_large_file	104%	101%	97%
fgrep -_search_for_"confidence"_pattern	111%	110%	97%
fgrep -_search_for_"the"_pattern	113%	112%	98%



Command line execution time (cont'd)

	2.2.19	2.4.18	2.5.25
	SMP-4/SMP-1	SMP-4/SMP-1	SMP-4/SMP-1
find_-_find_a_modified_file_(stat'ed)	107%	107%	98%
find_-_find_a_file_(not_stat'ed)	107%	108%	98%
grep_-_search_a_large_file_suppress_msgs	107%	105%	97%
grep_-_search_a_large_file	107%	106%	98%
grep_-_search_for_an_extended_pattern	110%	111%	98%
grep_-_search_for_a_specific_pattern	112%	113%	98%
grep_-_search_a_large_file_for_laban	99%	101%	100%
head_-_display_first_31+_lines	112%	114%	98%
head_-_display_first_51+_lines	113%	116%	99%
head_-_display_66+_lines	115%	119%	99%
join_-_join_files_with_formated_output	107%	112%	98%
ld_-_link_10_object_files	104%	103%	98%
ls_-_format_ls_output_(cpu_time)	101%	110%	98%
ls -IR	100%	102%	98%
ls_-_list_with_defaults	100%	109%	98%
sdiff_-_display_different_lines	111%	105%	98%
sdiff_-_display_identical_lines	107%	105%	98%
sed_-_edit_a_file_using_streams	100%	100%	98%
sh_-_execute_a_complex_shell_script	106%	119%	98%
sort_-_sort_a_large_file_(10,000_lines)	103%	102%	98%
split_-_split_a_file_into_50-line_seg	105%	108%	98%
split_-_split_a_file_using_defaults	106%	107%	99%
sum_-_display_checksum	104%	105%	98%
tail_-_display_file_8000_byte_from_EOF	113%	123%	98%
tail_-_display_end_of_file	112%	116%	98%
tar_-_create_a_tar_file	102%	97%	98%
tar_-_display_archive_table_of_contents	100%	102%	98%



Command line execution time (cont'd)

	2.2.19	2.4.18	2.5.25
	SMP-4/SMP-1	SMP-4/SMP-1	SMP-4/SMP-1
tee -_append_to_output_file	112%	99%	108%
tee -_display_and_write_file	115%	113%	98%
tr -_replace_all_numbers_with_#	112%	119%	98%
tr -_replace_A-E_with_1-5	107%	104%	98%
tr -_translate_lowercase_to_uppercase	112%	120%	98%
tr -_translate_{}_to_()	112%	120%	98%
uncompress -_force_uncompression	107%	107%	98%
vi -_perform_page_forwards/backwards	98%	99%	98%
vi -_quit,_no_save,_a_large_file	100%	89%	100%
vi -_edit_a_small_file	87%	97%	97%
vi -_write,_quit_a_large_file	101%	105%	99%
wc -_count_number_of_char_in_a_file	115%	118%	99%
wc -_count_number_of_lines_in_a_file	117%	114%	99%
wc -_count_number_of_words_in_a_file	109%	113%	98%
wc -_count_lines,_words,_and_char	108%	108%	99%

Scalability ratio analysis

- To study which kernel provides better performance scalability
- Performance ratio comparison between 2.4.18 and 2.5.25
 - ▶ UP ratio
 - ▶ SMP-1 ratio
 - ▶ SMP-4 ratio
- Performance metrics
 - ▶ Latency of simple functions
 - ▶ Interprocess communication latency
 - ▶ Context switching latency
 - ▶ Command line execution time

The IBM logo, consisting of the letters 'IBM' in a bold, sans-serif font, with horizontal lines through the letters.

Linux Technology
Center

Latency of simple functions

- Time to perform a simple system function
- All times are in microseconds

	UP Ratio 2.5.25/2.4.18	SMP-1 Ratio 2.5.25/2.4.18	SMP-4 Ratio 2.5.25/2.4.18
Simple syscall	108%	103%	101%
Simple read	116%	109%	109%
Simple write	110%	109%	108%
Simple stat	107%	88%	91%
Simple fstat	116%	107%	111%
Simple open/close	104%	91%	97%
Signal handler installation	104%	103%	102%
Signal handler overhead	100%	106%	104%
Protection fault	101%	103%	101%
Pipe latency	108%	103%	89%
AF_UNIX sock stream latency	105%	101%	62%
Process fork+exit	107%	108%	76%
Process fork+execve	109%	112%	97%
Process fork+/bin/sh -c	110%	110%	101%
Geometric Mean	107%	104%	95%



Interprocess communication latency

- "hot potato" benchmark - a token is passed back and forth between two processes
- Time interprocess communication latency via UDP/IP, TCP/IP, RPC/tcp, and RPC/udp using socket.
- Time interprocess connection latency via TCP/IP
- All tests using local host
- All times are in microseconds

	UP	SMP-1	SMP-4
	Ratio	Ratio	Ratio
	2.5.25/2.4.18	2.5.25/2.4.18	2.5.25/2.4.18
UDP latency using localhost	103%	98%	67%
TCP latency using localhost	113%	104%	86%
RPC/tcp latency using localhost	104%	106%	105%
RPC/udp latency using localhost	103%	105%	96%
TCP/IP connection cost to localh	101%	103%	104%
Geometric Mean	105%	103%	90%

Context switching latency

- Processes are connected in a ring of Unix pipes. Each process reads a token from its pipe, possibly does some work, and then writes the token to the next process
- All times are in microseconds
- Data are geometric mean of context switching latency time of 2, 4, 8, 16, 32, 64 and 96 processes doing a summing work over a specific array size of KB

	UP Ratio 2.5.25/2.4.18	SMP-1 Ratio 2.5.25/2.4.18	SMP-4 Ratio 2.5.25/2.4.18
Array size=0 KB			
2	110%	112%	67%
4	107%	100%	67%
8	115%	101%	79%
16	99%	94%	82%
24	91%	92%	81%
32	93%	97%	79%
64	91%	105%	83%
96	91%	109%	83%
Geometric Mean	99%	101%	77%

Command line execution time

	SMP-1 Ratio 2.5.25/2.4.18	SMP-4 Ratio 2.5.25/2.4.18
ar_-_create_an_archive_file	107%	100%
ar_-_output_contents_of_archive	110%	96%
ar_-_add_files_to_end_of_archive	106%	99%
ar_-_replace_a_file_in_archive	107%	112%
ar_-_output_table_of_contents	109%	97%
ar_-_Extract_a_file_in_archive	110%	99%
awk_-_read_a_line	104%	99%
awk_-_increment_a_variable	104%	100%
awk_-_split_a_line_into_fields	103%	100%
basename_-_get_basename	113%	93%
basename_-_get_basename_using_suffix	112%	96%
cat_-_display_a_large_file	104%	97%
cat_-_display_file_w/_all_action_options	105%	96%
chgrp_-_modify_group_	111%	103%
chmod_-_modify_permissions_using_755	111%	88%
chmod_-_modify_group_of_using_+x	110%	96%
chown_-_modify_owner	109%	94%
cmp_-_compare_not_=_files,_display_lines	106%	99%
cmp_-_compare_not_=_files,_display_byte	108%	96%
compress_-_compress_large_file,_forced	105%	91%
compress_-_compress_C_file,force_to_stdout	104%	98%
cp_-_copy_a_file_<100_bytes	110%	91%
cp_-_copy_a_file_>25000_bytes	112%	67%
cp_-_copy_a_file_>10000_bytes	112%	99%
cpio_-_un-archive_a_file	109%	95%
cpio_-_create_an_archive_file	108%	94%



Command line execution time (cont'd)

	SMP-1 Ratio 2.5.25/2.4.18	SMP-4 Ratio 2.5.25/2.4.18
csplit -_split_w/_prefix	106%	98%
csplit -_split_by_pattern	106%	99%
cut -_cut_a_file_by_character	106%	97%
cut -_cut_a_file_by_fields	106%	98%
dd -_bs=2048,_conv=lcase	111%	100%
dd -_with_bs=4096	109%	100%
dd -_with_bs=512	110%	92%
dd -_with_bs=65536	112%	97%
dd -_bs=2048,_conv=swab	112%	101%
diff -_differ_w/_embedded_blanks,_-be	107%	100%
diff -_compare_2_small_files_(<52_lines)	112%	92%
diff -_compare_large_files_=(10,000_l)	107%	107%
diff -_compare_not_equal_files_(1300_l)	106%	94%
diff -_compare_equal_files_(1300_lin)	106%	95%
dirname -_get_the_directory_name	113%	97%
du -_get_disk_usage_of_each_file	112%	104%
du -_get_disk_usage_of_total_sum_only	108%	96%
echo -_a_message_(100_char)	113%	98%
egrep -_search_for_extended_pattern_#1	110%	96%
egrep -_search_for_extended_pattern_#2	108%	96%
egrep -_search_for_simple_pattern	106%	95%
ex -_add_delete_write_quit_a_large_file	105%	102%
fgrep -_search_for_"confidence"_pattern	112%	99%
fgrep -_search_for_"the"_pattern	109%	96%

Command line execution time (cont'd)

	SMP-1 Ratio 2.5.25/2.4.18	SMP-4 Ratio 2.5.25/2.4.18
find -_find_a_modified_file_(stat'ed)	108%	99%
find -_find_a_file_(not_stat'ed)	108%	98%
grep -_search_a_large_file_suppress_msgs	105%	97%
grep -_search_a_large_file	104%	96%
grep -_search_for_an_extended_pattern	109%	96%
grep -_search_for_a_specific_pattern	111%	96%
grep -_search_a_large_file_for_laban	101%	100%
head -_display_first_31+_lines	113%	97%
head -_display_first_51+_lines	112%	95%
head -_display_66+_lines	113%	94%
join -_join_files_with_formatted_output	109%	96%
ld -_link_10_object_files	105%	100%
ls -_format_ls_output_(cpu_time)	105%	94%
ls -IR	102%	98%
ls -_list_with_defaults	109%	98%
m4 -_test_predefined_macros	107%	95%
make -_compile_C_code	108%	103%
mkdir -_make_a_directory	113%	84%
more -_display_a_file	105%	99%
nl -_number_all_lines	103%	100%
nl -_number_lines_in_a_file_(defaults)	103%	101%
nroff -_text_preprocess_a_file	103%	102%
paste -_paste_horizontally	109%	90%
paste -_paste_vertically	109%	97%



Command line execution time (cont'd)

	SMP-1 Ratio 2.5.25/2.4.18	SMP-4 Ratio 2.5.25/2.4.18
pr -_paginate_a_file_(columns)	105%	99%
pr -_paginate_with_header	105%	100%
pr -_paginate_a_file	105%	99%
ps-_report_process_status	101%	103%
pwd -_print_working_directory	114%	96%
rm -_remove_a_large_file	112%	60%
rmdir -_remove_a_directory	106%	92%
sdiff -_display_different_lines	107%	100%
sdiff -_display_identical_lines	107%	100%
sed -_edit_a_file_using_streams	103%	100%
sh -_execute_a_complex_shell_script	110%	90%
sort -_sort_a_large_file_(10,000_lines)	103%	99%
split -_split_a_file_into_50-line_seg	110%	100%
split -_split_a_file_using_defaults	104%	97%
sum -_display_checksum	105%	98%
tail -_display_file_8000_byte_from_EOF	111%	89%
tail -_display_end_of_file	112%	95%
tar -_create_a_tar_file	104%	105%
tar -_display_archive_table_of_contents	104%	100%
tee -_append_to_output_file	101%	110%
tee -_display_and_write_file	111%	97%
tr -_replace_all_numbers_with_#	112%	92%
tr -_replace_A-E_with_1-5	104%	98%
tr -_translate_lowercase_to_uppercase	111%	91%
tr -_translate_{}_to_()	109%	90%



Command line execution time (cont'd)

	SMP-1	SMP-4
	Ratio	Ratio
	2.5.25/2.4.18	2.5.25/2.4.18
uncompress -_force_uncompression	105%	96%
vi -_perform_page_forwards/backwards	103%	102%
vi -_quit,_no_save,_a_large_file	93%	104%
vi -_edit_a_small_file	96%	96%
vi -_write,_quit_a_large_file	104%	98%
wc -_count_number_of_char_in_a_file	111%	93%
wc -_count_number_of_lines_in_a_file	113%	98%
wc -_count_number_of_words_in_a_file	110%	95%
wc -_count_lines,_words,_and_char	109%	100%
Geometric mean	107%	97%



Legal Statement

- This work represents the views of the author and does not necessarily reflect the views of IBM Corporation.
- The following terms are trademarks or registered trademarks of International Business Machines Corporation in the United States and/or other countries: IBM (logo), e-business (logo).
- A full list of U.S. trademarks owned by IBM may be found at <http://www.ibm.com/legal/copytrade.shtml>.
- Linux is a registered trademark of Linus Torvalds.
- Intel is a trademark or registered trademark of Intel Corporation or its subsidiaries in the United States and other countries.
- Other company, product, and service names may be trademarks or service marks of others.

The IBM logo, consisting of the letters 'IBM' in a bold, sans-serif font, is positioned in the bottom left corner of the slide. The letters are white with a blue outline, set against a dark blue background.

Linux Technology
Center