

# Performance tuning the Linux Kernel on the ARM Architecture

(Playing with a cell phone!)



# Primary Objective: Preventative Maintenance

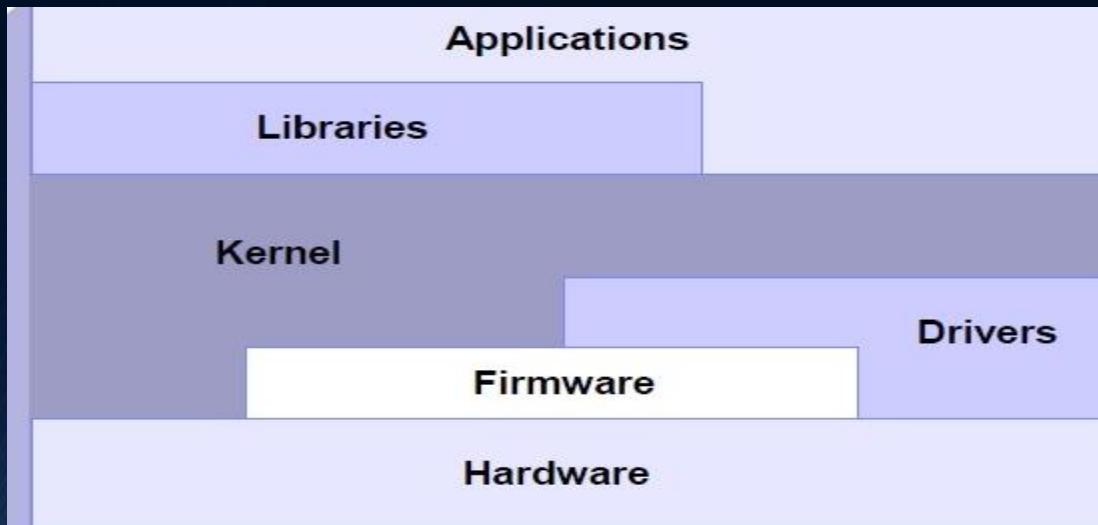


- Attempting to foresee probable bottle-necks and detect areas of poor performance during the very early development stages
- We have the advantage that we are very early in the development of the cluster, which gives us a window of opportunity to change tuning parameters without worrying about denial of service conditions

**“Optimising the linux environment for higher performance usually takes place once a working system is running to specification and unforeseen bottlenecks or bugs occur”**

-Linux Performance and Tuning Guidelines, (RedPaper)

System  
Kernel  
Linux 3.4.61-sun7i+ arm7l



AllWinnerTech SOC A20, ARM® Cortex™-A7 Dual-Core ARM® Mali400 MP2 Complies with OpenGL ES 2.0/1.1  
1GB DDR3 @480M

# What do we tune for? Application specific

- Higher performance configuration will skew system for the intended workload
- Expect poor performance for different workload patterns
- We are not sure at present - IO intensive?
- We want to understand what parameters are relevant for our purposes

# 794

net.ipv4.conf.default.arp_accept = 0	0	net.ipv4.neigh.default.gc_interval = 30	net.ipv4.route.gc_interval = 60	ws = 0	net.ipv6.conf.default.proxy_ndp = 0
net.ipv4.conf.default.arp_announce = 0	net.ipv4.conf.lo.proxy_arp = 0	net.ipv4.neigh.default.gc_stale_time = 60	net.ipv4.route.gc_min_interval = 0	= 18096	net.ipv6.conf.default.regen_max_retry = 3
net.ipv4.conf.default.arp_filter = 0	net.ipv4.conf.lo.proxy_arp_pvlan = 0	net.ipv4.neigh.default.gc_thresh1 = 128	net.ipv4.route.gc_thresh = 32768	24130	36192
net.ipv4.conf.default.arp_ignore = 0	net.ipv4.conf.lo.rp_filter = 1	net.ipv4.neigh.default.gc_thresh2 = 512	net.ipv4.route.gc_timeout = 300	net.ipv6.conf.default.router_probe_interv	
net.ipv4.conf.default.arp_notify = 0	net.ipv4.conf.lo.secure_redirects = 1	net.ipv4.neigh.default.gc_thresh3 = 1024	net.ipv4.route.max_size = 524288	al = 60	
net.ipv4.conf.default.bootp_relay = 0	net.ipv4.conf.lo.send_redirects = 1	net.ipv4.neigh.default.locktime = 100	net.ipv4.route.min_adv_mss = 256	net.ipv6.conf.default.router_solicitation_d	
net.ipv4.conf.default.disable_policy = 0	net.ipv4.conf.lo.shared_media = 1	net.ipv4.neigh.default.mcast_solicit = 3	net.ipv4.route.min_pmtu = 552	elay = 1	
net.ipv4.conf.default.disable_xfrm = 0	net.ipv4.conf.lo.src_valid_mark = 0	net.ipv4.neigh.default.proxy_delay = 80	net.ipv4.route.mtu_expires = 600	net.ipv6.conf.default.router_solicitation_i	
net.ipv4.conf.default.force_igmp_version = 0	net.ipv4.conf.lo.tag = 0	net.ipv4.neigh.default.proxy_qlen = 64	= net.ipv4.route.redirect_load = 2	interval = 4	
net.ipv4.conf.default.forwarding = 0	net.ipv4.conf.tun0.accept_local = 0	net.ipv4.neigh.default.retrans_time_ms	net.ipv4.route.redirect_number = 9	net.ipv6.conf.default.router_solicitations = 3	
net.ipv4.conf.default.log_martians = 0	net.ipv4.conf.tun0.accept_redirects = 1	= net.ipv4.neigh.default.ucast_solicit = 3	net.ipv4.route.redirect_silence = 2048	net.ipv6.conf.default.temp_preferred_lft = 86400	
net.ipv4.conf.default.mc_forwarding = 0	net.ipv4.conf.tun0.accept_source_route = 1	net.ipv4.neigh.default.unres_qlen = 35	net.ipv4.rt_cache_rebuild_count = 4	net.ipv6.conf.default.temp_valid_lft = 604800	
net.ipv4.conf.default.medium_id = 0	net.ipv4.conf.tun0.arp_accept = 0	net.ipv4.neigh.default.unres_qlen_bytes	net.ipv4.tcp_abc = 0	net.ipv6.conf.default.use_tempaddr = 2	
net.ipv4.conf.default.promote_secondaries = 0	net.ipv4.conf.tun0.arp_announce = 0	= 65536	net.ipv4.tcp_abort_on_overflow = 0	= net.ipv6.conf.eth0.accept_dad = 1	
net.ipv4.conf.default.proxy_arp = 0	net.ipv4.conf.tun0.arp_filter = 0	net.ipv4.neigh.eth0.anycast_delay = 100	net.ipv4.tcp_adv_win_scale = 1	net.ipv6.conf.eth0.accept_ra = 1	
net.ipv4.conf.default.proxy_arp_pvlan = 0	net.ipv4.conf.tun0.arp_ignore = 0	net.ipv4.neigh.eth0.app_solicit = 0	net.ipv4.tcp_allowed_congestion_control	net.ipv6.conf.eth0.accept_ra_defrtr = 1	
net.ipv4.conf.default.rp_filter = 1	net.ipv4.conf.tun0.notify = 0	net.ipv4.neigh.eth0.base_reachable_time_ms = 30000	net.ipv4.conf.all.accept_source_route = 0	net.ipv6.conf.eth0.accept_ra_info_max = 1	
net.ipv4.conf.default.secure_redirects = 1	net.ipv4.conf.tun0.bootp_relay = 0	net.ipv4.neigh.eth0.delay_first_probe_time_ms = 5	net.ipv4.conf.all.autoconf = 1	net.ipv6.conf.eth0.accept_ra_rtr_pref = 1	
net.ipv4.conf.default.send_redirects = 1	net.ipv4.conf.tun0.disable_policy = 0	net.ipv4.neigh.eth0.proxy_delay = 80	net.ipv4.conf.all.dad_transmits = 1	net.ipv6.conf.eth0.accept_ra_defrtr = 1	
net.ipv4.conf.default.shared_media = 1	net.ipv4.conf.tun0.medium_id = 0	net.ipv4.neigh.eth0.proxy_qlen = 64	net.ipv4.conf.all.disable_ipv6 = 0	net.ipv6.conf.eth0.accept_ra_info_max = 1	
net.ipv4.conf.default.src_valid_mark = 0	net.ipv4.conf.tun0.promote_secondaries = 0	net.ipv4.neigh.eth0.retrans_time_ms = 60	hybla htcp vegas veno scalable lp yeahnet	net.ipv6.conf.eth0.accept_ra_rtr_pref = 0	
net.ipv4.conf.default.tag = 0	net.ipv4.conf.tun0.proxy_arp = 0	net.ipv4.neigh.eth0.locktime = 100	net.ipv6.conf.all.force_mld_version = 0		
net.ipv4.conf.eth0.accept_local = 0	net.ipv4.conf.tun0.proxy_arp_pvlan = 0	net.ipv4.neigh.eth0.mcast_solicit = 3	net.ipv6.conf.all.force_tllao = 0		
net.ipv4.conf.eth0.accept_redirects = 1	net.ipv4.conf.tun0.rp_filter = 1	net.ipv4.neigh.eth0.proxy_delay = 80	net.ipv6.conf.all.forwarding = 0		
net.ipv4.conf.eth0.accept_source_route = 1	net.ipv4.conf.tun0.secure_redirects = 1	net.ipv4.neigh.eth0.proxy_qlen = 64	net.ipv6.conf.all.hop_limit = 64		
net.ipv4.conf.eth0.arp_accept = 0	net.ipv4.conf.tun0.send_redirects = 1	net.ipv4.neigh.eth0.retrans_time_ms = 30000	net.ipv6.conf.all.max_addresses = 16		
net.ipv4.conf.eth0.arp_announce = 0	net.ipv4.conf.tun0.shared_media = 1	net.ipv4.neigh.eth0.delay_first_probe_time_ms = 5	net.ipv6.conf.all.max_desync_factor = 600		
net.ipv4.conf.eth0.arp_filter = 0	net.ipv4.conf.tun0.src_valid_mark = 0	net.ipv4.neigh.eth0.tag = 0	net.ipv6.conf.all.autoconf = 1		
net.ipv4.conf.eth0.arp_ignore = 0	net.ipv4.conf.tun0.tag = 0	net.ipv4.icmp_echo_ignore_all = 0	net.ipv6.conf.eth0.dad_transmits = 1		
net.ipv4.conf.eth0.arp_notify = 0	net.ipv4.icmp_echo_ignore_broadcasts = 1	net.ipv4.icmp_echo_ignore_all = 0	net.ipv6.conf.eth0.force_mld_version = 0		
net.ipv4.conf.eth0.bootp_relay = 0	net.ipv4.icmp_errors_use_inbound_ifaddrs = 1	net.ipv4.icmp_echo_ignore_broadcasts = 1	net.ipv6.conf.eth0.force_tllao = 0		
net.ipv4.conf.eth0.disable_policy = 0	net.ipv4.icmp_errors_use_ipoptions = 1	net.ipv4.icmp_errors_use_inbound_ifaddrs = 1	net.ipv6.conf.eth0.proxy_ndp = 0		
net.ipv4.conf.eth0.disable_xfrm = 0	net.ipv4.icmp_ignore_bogus_error_responses = 1	net.ipv4.icmp_errors_use_ipoptions = 1	net.ipv6.conf.eth0.forwarding = 0		
net.ipv4.conf.eth0.force_igmp_version = 0	net.ipv4.icmp_max_memberships = 20	net.ipv4.icmp_ignore_bogus_error_responses = 1	net.ipv6.conf.eth0.hop_limit = 64		
net.ipv4.conf.eth0.forwarding = 0	net.ipv4.icmp_max_msf = 10	net.ipv4.icmp_max_memberships = 20	net.ipv6.conf.eth0.max_addresses = 16		
net.ipv4.conf.eth0.log_martians = 0	net.ipv4.icmp_max_ttl = 600	net.ipv4.icmp_max_msf = 10	net.ipv6.conf.eth0.max_desync_factor = 60		
net.ipv4.conf.eth0.mc_forwarding = 0	net.ipv4.icmp_peer_minttl = 120	net.ipv4.icmp_max_ttl = 600	net.ipv6.conf.eth0.autoconf = 1		
net.ipv4.conf.eth0.medium_id = 0	net.ipv4.icmp_peer_threshold = 65664	net.ipv4.icmp_max_ttl = 1000	net.ipv6.conf.eth0.eth0.disable_ipv6 = 0		
net.ipv4.conf.eth0.promote_secondaries = 0	net.ipv4.icmp_ratelimit = 1000	net.ipv4.icmp_max_ttl = 1000	net.ipv6.conf.eth0.force_mld_version = 0		
net.ipv4.conf.eth0.proxy_arp = 0	net.ipv4.icmp_ratemask = 6168	net.ipv4.icmp_max_ttl = 3	net.ipv6.conf.eth0.force_tllao = 0		
net.ipv4.conf.eth0.proxy_pvlan = 0	net.ipv4.icmp_max_memberships = 20	net.ipv4.icmp_max_ttl = 35	net.ipv6.conf.eth0.proxy_ndp = 0		
net.ipv4.conf.eth0.rp_filter = 1	net.ipv4.icmp_max_msf = 10	net.ipv4.icmp_max_ttl = 35	net.ipv6.conf.eth0.forwarding = 0		
net.ipv4.conf.eth0.secure_redirects = 1	net.ipv4.icmp_max_ttl = 600	net.ipv4.icmp_max_ttl = 75	net.ipv6.conf.eth0.hop_limit = 64		
net.ipv4.conf.eth0.send_redirects = 1	net.ipv4.icmp_max_ttl = 120	net.ipv4.icmp_max_ttl = 90	net.ipv6.conf.eth0.max_addresses = 16		
net.ipv4.conf.eth0.shared_media = 1	net.ipv4.icmp_max_ttl = 120	net.ipv4.icmp_max_ttl = 7200	net.ipv6.conf.eth0.max_desync_factor = 60		
net.ipv4.conf.eth0.src_valid_mark = 0	net.ipv4.icmp_max_ttl = 120	net.ipv4.icmp_max_ttl = 0	net.ipv6.conf.eth0.router_solicitation_delay = 600		
net.ipv4.conf.eth0.tag = 0	net.ipv4.icmp_max_ttl = 120	net.ipv4.icmp_max_ttl = 0	net.ipv6.conf.eth0.mc_forwarding = 0		
net.ipv4.conf.eth0.local_port_range = 32768	me = 5	net.ipv4.icmp_max_ttl = 60	net.ipv6.conf.eth0.mtu = 1500		
net.ipv4.conf.lo.accept_local = 0	61000	net.ipv4.neigh.tun0.gc_stale_time = 60	net.ipv6.conf.eth0.optimistic_dad = 0		
net.ipv4.conf.lo.accept_redirects = 1	1 net.ipv4.ip_local_reserved_ports = 1	net.ipv4.neigh.tun0.locktime = 100	net.ipv6.conf.eth0.temp_valid_lft = 604800		
net.ipv4.conf.lo.accept_source_route = 1	net.ipv4.ip_no_pmtu_disc = 0	net.ipv4.neigh.tun0.mcast_solicit = 3	net.ipv6.conf.eth0.use_tempaddr = 2		
net.ipv4.conf.lo.accept = 0	net.ipv4.ip_no_local_bind = 0	net.ipv4.neigh.tun0.proxy_delay = 80	net.ipv6.conf.eth0.default.accept_dad = 1		
net.ipv4.conf.lo.arp_announce = 0	net.ipv4.ip_frag_high_thresh = 262144	net.ipv4.neigh.tun0.proxy_qlen = 64	net.ipv6.conf.eth0.default.accept_ra_defrtr = 1		
net.ipv4.conf.lo.arp_filter = 0	net.ipv4.ip_frag_low_thresh = 196608	net.ipv4.neigh.tun0.retrans_time_ms = 1000	net.ipv6.conf.eth0.default.accept_ra_pinfo = 1		
net.ipv4.conf.lo.arp_ignore = 0	net.ipv4.ip_frag_max_dist = 64	net.ipv4.neigh.tun0.unicast_solicit = 3	net.ipv6.conf.eth0.default.accept_ra_rt_info_		
net.ipv4.conf.lo.arp_notify = 0	net.ipv4.ip_frag_secret_interval = 600	net.ipv4.neigh.tun0.unres_qlen = 35	max_plen = 0		
net.ipv4.conf.lo.bootp_relay = 0	net.ipv4.ip_frag_time = 30	net.ipv4.neigh.tun0.unres_qlen_bytes	net.ipv6.conf.eth0.default.accept_ra_rtr_pref = 1		
net.ipv4.conf.lo.disable_policy = 1	net.ipv4.neigh.default.anycast_delay = 100	= 65536	net.ipv6.conf.eth0.default.accept_redirects = 1		
net.ipv4.conf.lo.disable_xfrm = 1	net.ipv4.neigh.default.app_solicit = 0	net.ipv4.ping_group_range = 0	net.ipv6.conf.eth0.default.accept_source_route = 0		
net.ipv4.conf.lo.force_igmp_version = 0	net.ipv4.neigh.default.base_reachable_time_ms = 30000	me = 5	net.ipv6.conf.eth0.default.autoconf = 1		
net.ipv4.conf.lo.forwarding = 0	net.ipv4.neigh.default.gc_stale_time = 60	net.ipv4.route.error_burst = 500	net.ipv6.conf.eth0.default.dad_transmits = 1		
net.ipv4.conf.lo.log_martians = 0	net.ipv4.neigh.default.gc_stale_time = 60	net.ipv4.route.error_cost = 100	net.ipv6.conf.eth0.default.disable_ipv6 = 0		
net.ipv4.conf.lo.mc_forwarding = 0	net.ipv4.neigh.default.gc_stale_time = 60	net.ipv4.route.gc_elasticity = 8	net.ipv6.conf.eth0.default.force_mld_version = 0		
net.ipv4.conf.lo.medium_id = 0	net.ipv4.neigh.default.gc_stale_time = 60		net.ipv6.conf.eth0.default.force_tllao = 0		
net.ipv4.conf.lo.promote_secondaries = 0	net.ipv4.neigh.default.gc_stale_time = 60		net.ipv6.conf.eth0.default.forwarding = 0		
	me = 5		net.ipv6.conf.eth0.default.hop_limit = 64		
			net.ipv6.conf.eth0.default.max_addresses = 16		
			net.ipv6.conf.eth0.default.optimistic_dad = 0		

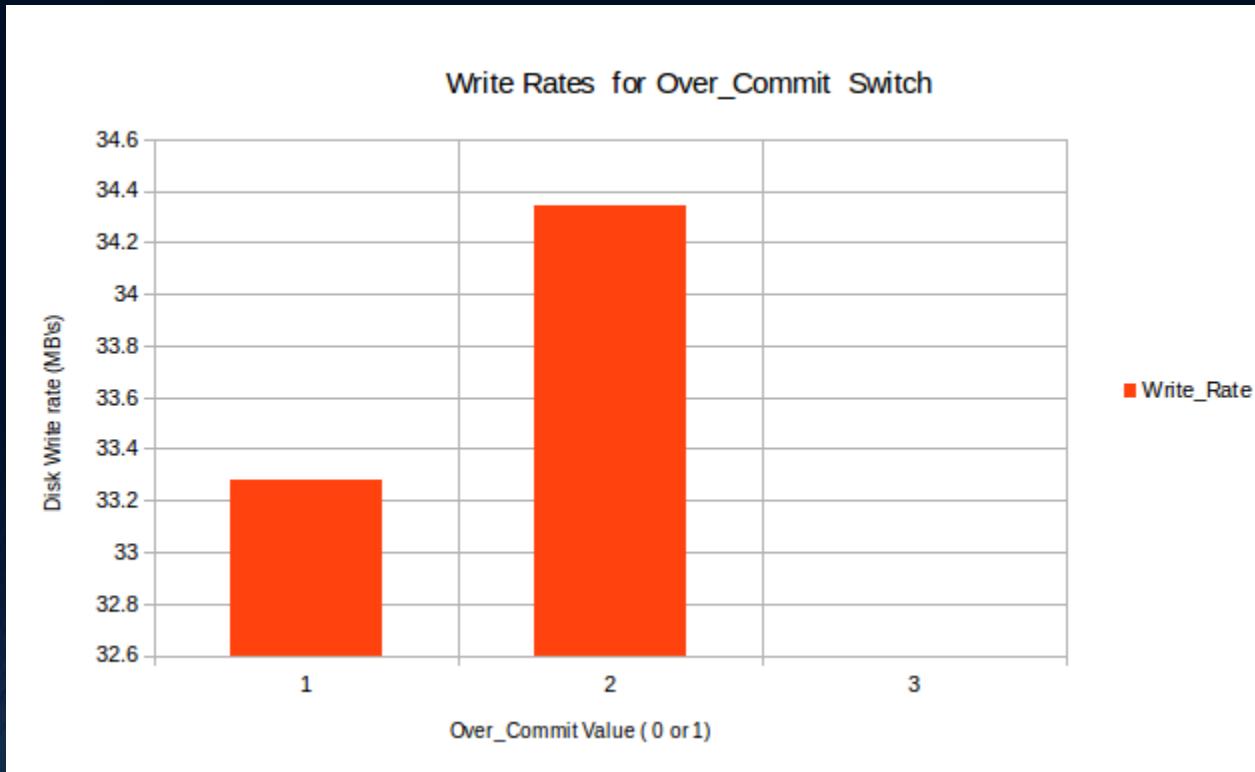
# How do we tune?

## /proc/sys

/proc/sys/fs/*	Used to increase the number of open files the OS allows and to handle quota.
/proc/sys/kernel/*	For tuning purposes, you can enable hotplug, manipulate shared memory, and specify the maximum number of PID files and level of debug in syslog.
/proc/sys/net/*	Tuning of network in general, IPV4 and IPV6.
/proc/sys/vm/*	Management of cache memory and buffer.

```
/$ sysctl -a  
/$ sysctl vm > /home/jonathan/vm.txt  
/$ sysctl vm.overcommit_memory  
/$ 0  
/$ sysctl -w vm.overcommit_memory=1  
/$ 1
```

## Diving In: dd command

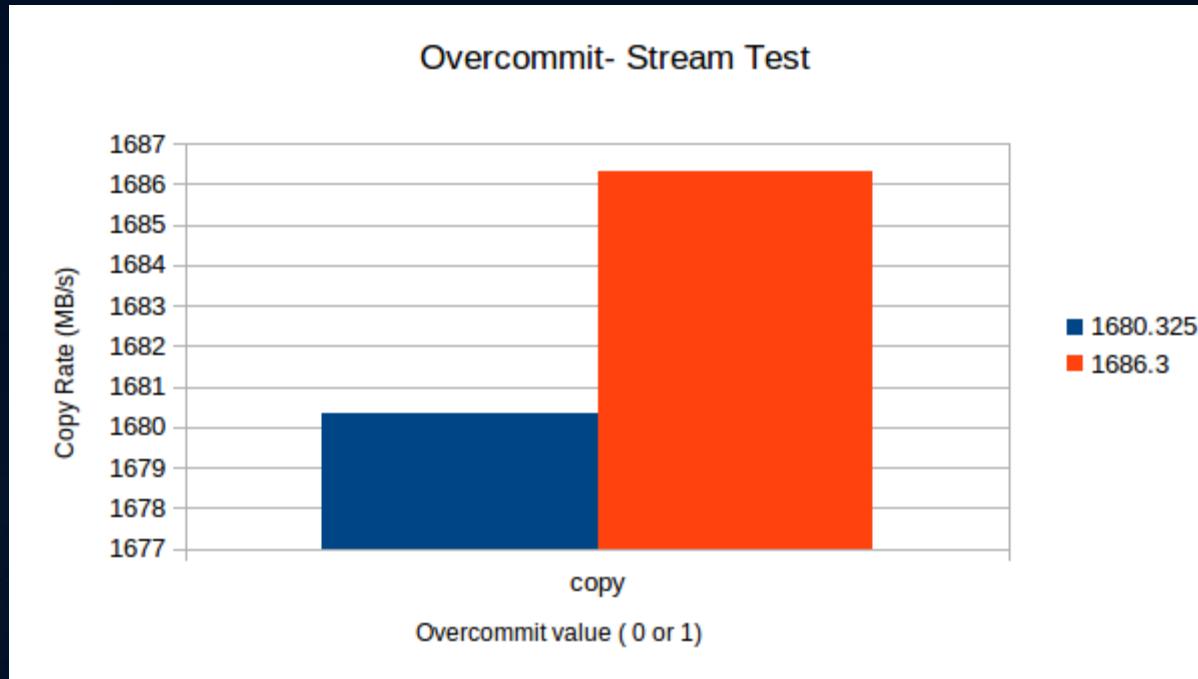


3% improvement in disk write rate

.overcommit\_memory=1

The kernel will grant every memory allocation without checking whether the allocation can be satisfied

# Stream results for overcommit=1



# Stream – popular benchmark for memory

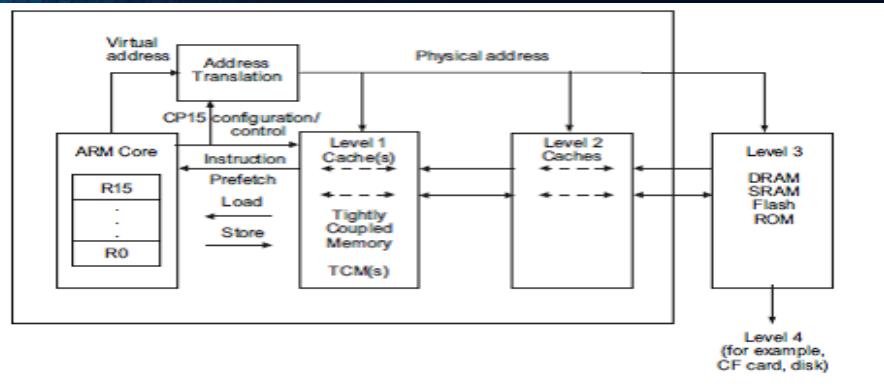
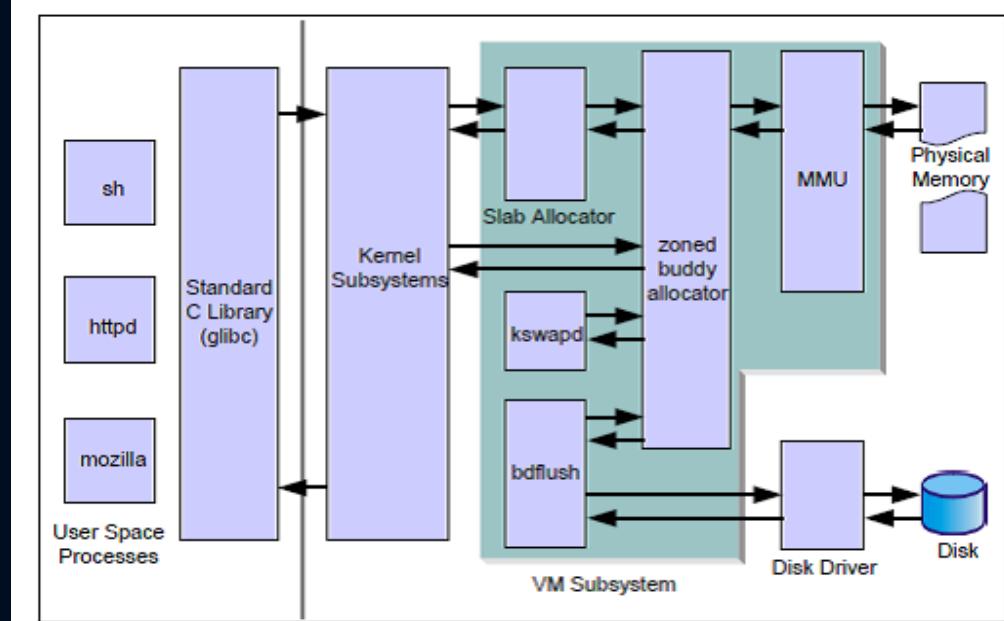
## Four functions

- 1.Copy:  $A(i) = b(i)$
- 2.Add:  $a(i) = b(i) + c(i)$
- 3.Scale:  $a(i) = q(bi)$
- 4.Triad:  $a(i) = b(i) + q(ci)$

# /proc/sys/vm

```
vm.block_dump = 0           vm.min_free_kbytes = 16384
vm.dirty_background_bytes = vm.min_free_order_shift = 1
0                           vm.mmap_min_addr = 32768
vm.dirty_background_ratio = vm.nr_pdflush_threads = 0
10                          vm.oom_dump_tasks = 1
vm.dirty_bytes = 0          vm.oom_kill Allocating_task =
vm.dirty_expire_centisecs = 0
3000                         vm.overcommit_memory = 1
vm.dirty_ratio = 20          vm.overcommit_ratio = 50
vm.dirty_writeback_centisecs= vm.page-cluster = 3
= 500                        vm.panic_on_oom = 0
vm.drop_caches = 0           vm.percpu_pagelist_fraction =
vm.extfrag_threshold = 500   0
vm.extra_free_kbytes = 0     vm.scan_unevictable_pages =
vm.hightmem_is_dirtyable = 0 0
vm.laptop_mode = 0            vm.stat_interval = 1
vm.legacy_va_layout = 0      vm.swappiness = 60
vm.lowmem_reserve_ratio =    vm.vfs_cache_pressure = 100
256      32      32
vm.max_map_count = 65530
```

# Where to from here?



1	vm.block_dump	=	0
2	vm.dirty_background_d_bytes	=	0
3	vm.dirty_background_d_ratio	=	5
4	vm.dirty_bytes	=	0
5	vm.dirty_expire_ns	=	3000
6	vm.dirty_ratio	=	10
7	vm.dirty_writeback_centisecs	=	500
8	vm.drop_caches	=	0
9	vm.extfrag_threshold	=	500
10	vm.hugepages淘汰_as_movable	=	0
11	vm.hugetlb_shm_group	=	0
12	vm.laptop_mode	=	0
13	vm.legacy_va_layout	=	0
14	vm.lowmem_reserve_ratio	=	256
15	vm.max_map_count	=	65530
16	vm.memory_failure_early_kill	=	0
17	vm.memory_failure_recovery	=	1
18	vm.min_free_kbytes	=	67584
19	vm.zone_reclaim_mode	=	5
20	vm.min_unmapped_ratio	=	1
21	vm.mmap_min_addr	=	65536
22	vm.nr_hugepages	=	0
24	vm.nr_overcommit_hugepages	=	0
25	vm.pdflush_thresholds	=	0
26	vm.numa_zonelist_order	=	default
27	vm.oom_dump_tasks	=	1
28	vm.oom_kill_allocating_task	=	0
29	vm.overcommit_memory	=	0
30	vm.overcommit_ratio	=	50
31	vm.page-cluster_size	=	3
32	vm.panic_on_oom	=	0
33	vm.percpu_pagelist_fraction	=	0
34	vm.scan_unevictable_pages	=	0
35	vm.stat_interval	=	1
36	vm.swappiness	=	60
37	vm.vfs_cache_pressure	=	100

# Flop/w

# Thank You