Linux System Performance Review

IGEL endpoint ITC00E0C52A6A64

Overall assessment	Critical resource pressure during workload: CPU saturation plus severe memory/swap exhaustion.
Operating system	IGEL OS 12 - version 12.9.0+3 (kernel 6.18.6, x86_64)
Hardware	IGEL M350C; AMD Ryzen Embedded R1505G; 2 cores / 4 threads; 4 GB installed DDR4; 7.3 GB eMMC
Capture window	200 seconds, 2026-07-31 07:33:57 to 07:37:18 MDT
Primary constraint	Memory capacity. Available RAM fell from 1.7 GiB to 100 MiB and swap from 1.0 GiB free to approximately 0 MiB free.
Secondary constraint	CPU was 85.1% busy on average; 110 of 199 samples were at least 90% busy.
Likely workload driver	Microsoft Edge activity, with browser processes reaching approximately 606 MiB RSS per process and up to 150% CPU; memory reclaim (kswapd0) averaged about 30% CPU when present.

Executive conclusion

This system is under-sized for the observed workload. The 4 GB physical-memory configuration (3.2 GiB usable by the OS) is the dominant limitation. During the test, both RAM and the 1.9 GiB zram swap device were consumed, triggering sustained page reclaim, swap traffic, elevated system CPU, blocked tasks, and intermittent I/O wait. The storage device itself delivered low read latency most of the time, so the disk is not the original cause; it is being stressed as a consequence of memory pressure.

Recommended priority: reduce browser/session memory demand or move the workload to hardware with at least 8 GB RAM. For heavier multi-tab browser use, embedded web applications, or concurrent VDI/browser workloads, 16 GB provides substantially more headroom.

1. System Summary

Category	Detail	Assessment
OS NAME	IGEL OS 12	Current captured platform identity
OS VERSION	12.9.0+3	Exact /etc/os-release VERSION
OS base	Ubuntu-compatible; VERSION_ID 20.04, focal	IGEL-managed immutable OS architecture
Kernel	6.18.6 x86_64	4 logical CPUs detected
Hostname	ITC00E0C52A6A64	Capture target
System manufacturer	IGEL Technology GmbH	From dmidecode
Product / family	M350C / Picasso	From dmidecode
System serial	10D3G1005A194302F7	From dmidecode
Baseboard	IGEL M350C; serial 00E0C52A6A64	From dmidecode
BIOS	INSYDE Corp. M350C V:3.2.13-10012019; 2019-10-01	UEFI capable; older firmware
Processor	AMD Ryzen Embedded R1505G with Radeon Vega Gfx	2 cores, 4 threads; 2.0 GHz current / 3.3 GHz max
CPU cache	L1 192 KiB; L2 1 MiB; L3 4 MiB	dmidecode values
Installed memory	4 GB: 2 x 2 GB DDR4-2400 SODIMM	Dual-channel; non-ECC; max array capacity reported as 64 GB
Usable memory	3.2 GiB	Lower than installed due to platform/device reservations
Swap	1.9 GiB zram	Nearly fully consumed by end of capture
Primary storage	7.3 GB eMMC (/dev/mmcblk0)	Small-capacity embedded storage
Network	wlan0 active at 10.0.0.18; enp1s0 disconnected	No link errors or drops observed
Collection duration	200 seconds	199 one-second performance samples
DMI boot status	No errors detected	dmidecode System Boot Information
Onboard devices	Broadcom 5762 Ethernet; Realtek ALC701 sound	Reported enabled by DMI

Note: DMI data is firmware-reported and can contain generic placeholders. The TPM entry, for example, reports specification version 0.0 and should not be treated as proof of a functional TPM.

2. Performance Findings

CPU

Metric	Observed	Interpretation
Average CPU busy	85.1%	High sustained utilization for a 4-thread system
Peak CPU busy	100%	At least one sample had no idle CPU
Average user CPU	46.1%	Application workload was substantial
Average system CPU	32.3%	Very high kernel overhead, consistent with reclaim and I/O activity
Average / peak I/O wait	6.3% / 64.1%	Intermittent stalls, not continuously storage-bound
Samples >=90% busy	110 of 199 (55.3%)	Sustained CPU saturation

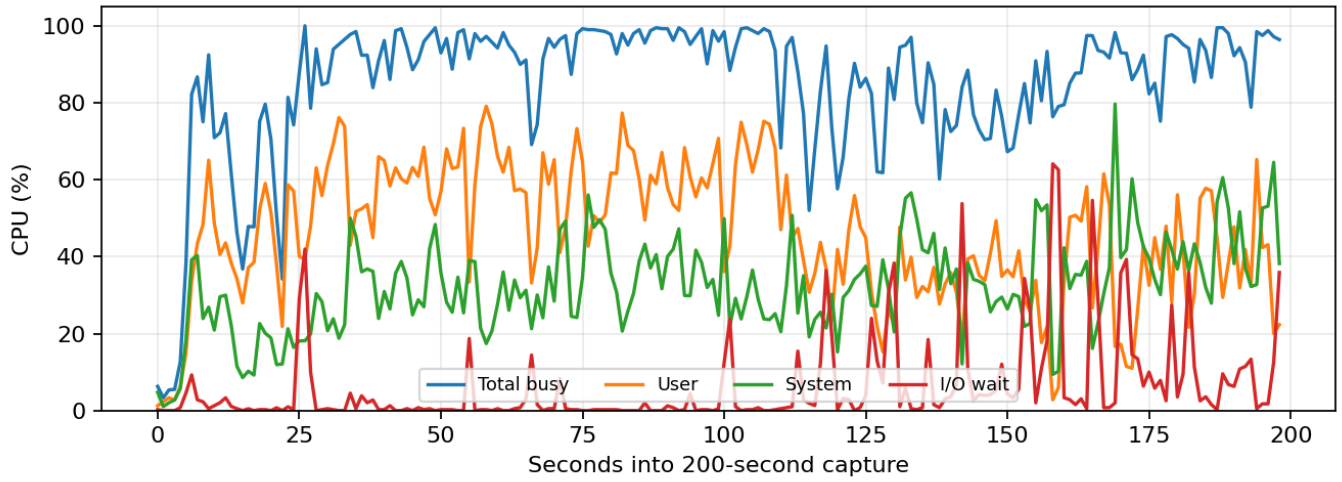

The CPU trace begins mostly idle, then rapidly enters a sustained high-load period. System CPU frequently rises alongside user CPU, indicating that the kernel is doing significant work rather than merely executing applications.

Memory and swap

Point	RAM used / available	Swap used / free
Capture start	1.5 GiB used / 1.7 GiB available	993 MiB used / 1.0 GiB free
Capture end	3.2 GiB used / 100 MiB available	1.9 GiB used / approximately 0 MiB free

Critical finding: the endpoint moved from moderate memory pressure to near-total exhaustion in about 200 seconds. vmstat recorded average swap-in of 6.1 MiB/s and swap-out of 12.6 MiB/s, with peaks of 35.4 MiB/s in and 85.5 MiB/s out. The kswapd0 kernel thread averaged about 30% CPU across its active samples, confirming intensive memory reclaim.

3. Storage and Network

Storage - /dev/mmcblk0

Metric	Average	Peak	Assessment
Read throughput	8.41 MiB/s	59.78 MiB/s	Workload is read-heavy during pressure event
Write throughput	2.70 MiB/s	21.67 MiB/s	Bursty writes, likely swap/cache activity
Read latency	3.53 ms	106.0 ms	Average is acceptable; isolated large spikes
Write latency	10.16 ms	232.6 ms	Higher and bursty under pressure
Queue depth	1.51	33.54	Severe transient queueing
Device utilization	22.0%	96.8%	Near saturation at peak, but not continuously

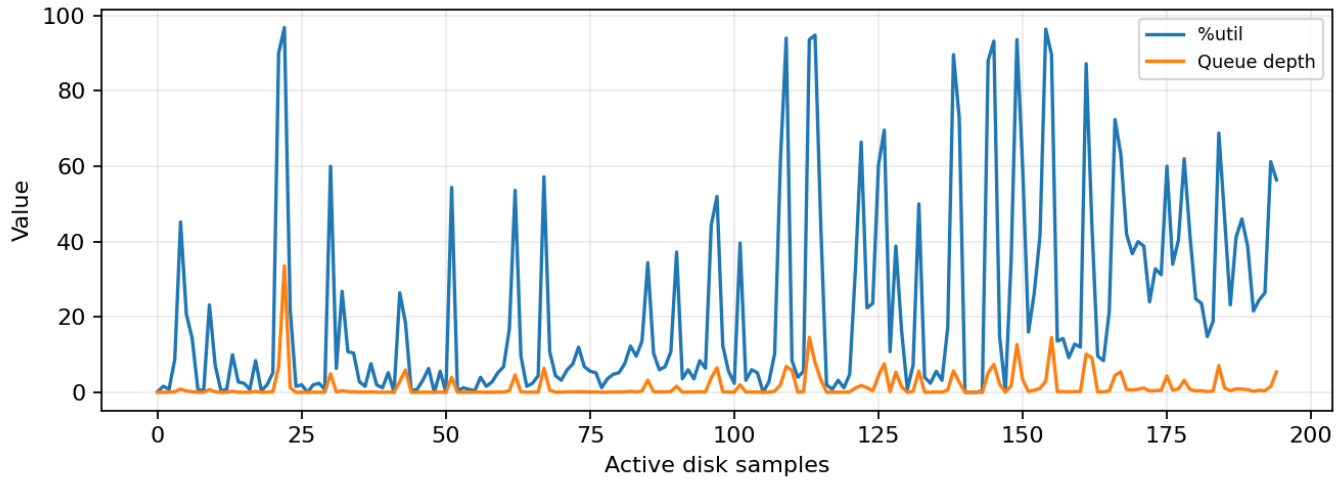

The eMMC device reaches 96.8% utilization and queue depth 33.5 during bursts. However, average utilization is about 22%, average read latency is 3.5 ms, and the initial baseline is nearly idle. This pattern indicates secondary storage pressure caused by swapping and workload startup rather than a continuously failing or inherently saturated disk.

Network

Interface	State / address	Traffic during capture	Errors
wlan0	UP; 10.0.0.18/24; default route via 10.0.0.1	Avg RX 355.3 KiB/s; peak RX 20.9 MiB/s; avg TX 49.6 KiB/s	0 errors, 0 drops
enp1s0	UP administratively, NO-CARRIER / DOWN	No traffic	0 errors, 0 drops

Wireless traffic included a short download burst. Established connections were primarily Microsoft Edge HTTPS sessions plus the IGEL Remote Management Agent on TCP 8443. TCP queues were generally empty, and the collected interface counters show no packet errors or drops.

4. Process-Level Analysis

The process data points to Microsoft Edge as the primary application workload and memory consumer. Because Edge uses many subprocesses, command-level averages combine many samples and should be interpreted directionally rather than as a single-process lifetime average.

Process	Peak CPU	Peak RSS	Peak read/write rate	Interpretation
msedge	150%	620,852 KiB (~606 MiB)	45.5 / 19.6 MiB/s	Primary application pressure; multiple renderer/browser processes
kswapd0	97%	Kernel thread	-	Heavy memory reclaim and swap management
WPEWebProcess	82%	145,488 KiB (~142 MiB)	27.6 / 0 MiB/s	Embedded web-content workload
mutter	25%	76,324 KiB (~74.5 MiB)	13.5 / 0 MiB/s	Display compositor affected by overall load
Xorg	12%	28,280 KiB (~27.6 MiB)	2.4 / 0 MiB/s	Normal display-server activity

Important observations

• Browser process CPU can exceed 100% because pidstat reports usage across multiple logical CPUs. A 150% reading means roughly one and a half logical CPUs were consumed during that interval.

• kswapd0 is a symptom, not the root application. Its high CPU usage reflects the kernel trying to recover memory under pressure.

• earlyoom is configured on the endpoint. With only about 100 MiB available and swap exhausted at the end, the system was approaching the point where earlyoom could terminate a preferred process to preserve system responsiveness.

• Blocked tasks peaked at 14 and the run queue peaked at 22, both high for four logical CPUs. This aligns with the periods of I/O wait and reclaim activity.

5. Findings and Severity

Severity	Finding	Evidence	Impact
Critical	Physical memory and swap exhaustion	End: 100 MiB available RAM; zram swap essentially full	Thrashing, latency, process termination risk
High	Sustained CPU saturation	85.1% average busy; 55.3% of samples >=90% busy	Slow UI and application response
High	Kernel reclaim overhead	kswapd0 ~30% average CPU in active samples; heavy swap I/O	CPU diverted from user workload
Medium	Transient eMMC saturation	96.8% peak utilization; queue depth 33.5	Short I/O stalls during swap/write bursts
Low	Network health	No link errors or drops; empty TCP queues	No evidence network is the bottleneck
Advisory	BIOS age	Firmware dated 2019-10-01	Review against current supported firmware baseline

6. Recommendations

1 - Increase memory capacity

Use a configuration with at least 8 GB RAM for this workload. Prefer 16 GB where multiple Edge tabs, web applications, VDI clients, or embedded WPE/WebKit content run concurrently. Confirm the exact supported upgrade options for the M350C platform before changing hardware.

2 - Reduce browser working set

Limit concurrent tabs/windows, disable unnecessary extensions, review background apps, and validate whether Edge Sleeping Tabs / memory-saving controls are permitted in the managed policy set. Investigate the specific Edge page or workload that started near 07:34:03.

3 - Reproduce with process grouping

Repeat the capture while grouping Edge processes by PID, command line, and process type (browser, renderer, GPU, utility). Add periodic `ps` RSS/PSS totals and, where available, `smem -tk` to avoid double-counting shared pages.

4 - Add memory-pressure telemetry

Collect `/proc/pressure/memory`, `/proc/pressure/io`, `sar -r ALL 1`, `sar -W 1`, and `journalctl` messages from earlyoom/OOM during the same interval. PSI will quantify time stalled by memory and I/O pressure.

5 - Correlate workload start

The resource event begins several seconds into the trace. Capture application launch timestamps, Edge URLs/process types, and user actions to identify whether a page load, download, update, video session, or VDI workflow initiated the spike.

6 - Review firmware and storage health

Check the supported IGEL firmware/BIOS baseline for M350C. Also collect eMMC health/lifetime information where exposed by the device and kernel. No direct storage failure is proven by this capture.

7. Data Quality and Limitations

This report is based on a 200-second snapshot. It is sufficient to identify the acute pressure event but not to establish long-term frequency. The collection commands were terminated by timeout with status 124 as designed; this does not indicate collector failure. The final `free -h` snapshot is authoritative for end-state memory, while vmstat showed a constant free-memory field in this capture and was therefore used primarily for swap, run-queue, blocked-task, I/O, and CPU trends. DMI values are firmware-supplied and may include placeholders.

No temperature, filesystem-capacity (`df`), SMART/eMMC health, journal/OOM log, GPU load, or application-specific telemetry was present in the archive. Those items would improve root-cause precision.

Bottom line

The observed performance problem is primarily a memory-capacity problem that cascades into swap activity, high kernel CPU, CPU saturation, and transient eMMC queueing. Network health appears normal. Microsoft Edge is the largest visible application contributor, but the platform has too little memory headroom for the captured workload.