

Some OCS Experiences

Cosmin Deaconu

UChicago/EFI/KICP

October 29, 2020

docker setup

- Rather than starting processes manually, decided to pursue recommended `docker` setup.
- Immediate problem: `CGroupsV2`
 - ▶ Fedora (all my personal computers) and EL8 (most powerful machine I have root access to in the lab) use `CGroupsV2`, which `docker` has yet to support
 - ▶ “Drop-in replacements” for `docker` like `podman` are nice (no need for root or a server process!), and are the official solution, but doesn’t support `docker-compose`. There is something called `podman-compose`, but I didn’t since I read it doesn’t really work yet.
 - ▶ This is all fixable (e.g. by configuring kernel to `CGroupsV1`, or setting up everything with `podman`, or making it work with `podman`), but a a bit of a pain.
- So I set it up on my workstation in the office, a 5-year-old machine running EL7.
 - ▶ Obviously I could have used an Ubuntu VM, which would probably work, but I’ve been on the Fedora/EL train for the last 15 years and I suspect I’m not the only one.
 - ▶ `docker` works just fine on EL7, but enforcing `SELinux` got in the way, probably because the mounted `docker` directory was in my filesystem. After many tweaks to policy and running into unlogged problems, I just admitted defeat and `setenforce 0`.
- So the experience is not quite out-of-the box on non-Ubuntu systems. Is that an issue for anybody but me?

Some very simple stress testing with FakeDataAgent

- Enabled instances of FakeDataAgent, which just generates random junk at a configurable “sample rate” and “number of channels” which can be made large.
 - ▶ For each instance, set number of channels to 20, sample rate to 1000 (more than most reasonable housekeeping, but not a crazy amount of data... less than an audio stream).
 - ▶ Each is running in its own container, all on the same host.
- Keep in mind this is running on a 5-year-old workstation (8 GB of ram, quad-core haswell i5-4590, with 3.3 GHz Base Clock).
- System usage with 0-16 enabled shown in next slides
 - ▶ Just snapshots from `docker stats`, CPU usage obviously varies
 - ▶ I/O amounts don't correspond to same time period... so don't read too much into those.

“Control” (i.e. the defaults configured in the OCS getting started)

CONTAINER	CPU %	MEM USAGE / LIMIT	MEM %	NET I/O	BLOCK I/O	PIDS
cssiteconfigs_ocs-LSSIM_1	0.39%	35.56 MiB / 7.582 GiB	0.46%	71 kB / 82.9 kB	0 B / 0 B	3
ocssiteconfigs_ocs-client_1	0.00%	456 KiB / 7.582 GiB	0.01%	5.67 kB / 656 B	0 B / 0 B	1
ocssiteconfigs_sisock-crossbar_1	0.24%	103.7 MiB / 7.582 GiB	1.34%	43 kB / 36.7 kB	0 B / 0 B	2
ocssiteconfigs_ls240-sim_1	0.29%	14.03 MiB / 7.582 GiB	0.18%	69.2 kB / 62.3 kB	0 B / 0 B	4
influxdb	0.04%	23.09 MiB / 7.582 GiB	0.30%	41.4 kB / 7.07 kB	336 kB / 266 kB	13
ocssiteconfigs_grafana_1	0.02%	15.04 MiB / 7.582 GiB	0.19%	19.1 kB / 3.49 kB	98.3 kB / 0 B	12
ocssiteconfigs_ocs-web_1	0.00%	3.297 MiB / 7.582 GiB	0.04%	8.52 kB / 656 B	0 B / 0 B	5
ocssiteconfigs_ocs-influx-publisher_1	0.11%	39.85 MiB / 7.582 GiB	0.51%	32.6 kB / 37.6 kB	0 B / 0 B	5

One fake data agent

CONTAINER	CPU %	MEM USAGE / LIMIT	MEM %	NET I/O	BLOCK I/O	PIDS
ocssiteconfigs_ocs-LSSIM_1	0.45%	35.57 MiB / 7.582 GiB	0.46%	132 kB / 158 kB	0 B / 49.2 kB	3
ocssiteconfigs_ocs-fake-data-agent_1	3.51%	37.3 MiB / 7.582 GiB	0.48%	39.8 kB / 6.72 MB	0 B / 24.6 kB	5
ocssiteconfigs_ocs-client_1	0.00%	452 KiB / 7.582 GiB	0.01%	6.08 kB / 656 B	0 B / 0 B	1
ocssiteconfigs_ocs-web_1	0.00%	3.297 MiB / 7.582 GiB	0.04%	10.1 kB / 656 B	0 B / 0 B	5
ocssiteconfigs_ocs-influx-publisher_1	34.82%	43.77 MiB / 7.582 GiB	0.56%	6.77 MB / 24 MB	0 B / 20.5 kB	5
ocssiteconfigs_ls240-sim_1	0.27%	14.03 MiB / 7.582 GiB	0.18%	130 kB / 121 kB	0 B / 0 B	4
influxdb	4.55%	66.53 MiB / 7.582 GiB	0.86%	24 MB / 47.9 kB	2.2 MB / 20.5 MB	13
ocssiteconfigs_sisock-crossbar_1	3.85%	110 MiB / 7.582 GiB	1.42%	6.82 MB / 6.79 MB	0 B / 0 B	2
ocssiteconfigs_grafana_1	0.04%	17.21 MiB / 7.582 GiB	0.22%	22.8 kB / 3.43 kB	541 kB / 22 MB	14

2 fake data agents

CONTAINER	CPU %	MEM USAGE / LIMIT	MEM %	NET I/O	BLOCK I/O	PIDS
ocssiteconfigs_ocs-fake-data-agent2_1	3.58%	37.3 MiB / 7.582 GiB	0.48%	75.9 kB / 15.7 MB	0 B / 20.5 kB	5
ocssiteconfigs_ocs-LSSIM_1	0.41%	35.57 MiB / 7.582 GiB	0.46%	299 kB / 366 kB	0 B / 61.4 kB	3
ocssiteconfigs_ocs-fake-data-agent_1	3.48%	37.3 MiB / 7.582 GiB	0.48%	83 kB / 15.7 MB	0 B / 12.3 kB	5
ocssiteconfigs_ocs-client_1	0.00%	452 KiB / 7.582 GiB	0.01%	7.01 kB / 656 B	0 B / 0 B	1
ocssiteconfigs_sisock-crossbar_1	7.08%	114.5 MiB / 7.582 GiB	1.48%	31.7 MB / 31.6 MB	0 B / 0 B	2
ocssiteconfigs_ocs-influx-publisher_1	58.68%	44.84 MiB / 7.582 GiB	0.58%	31.4 MB / 112 MB	0 B / 0 B	5
ocssiteconfigs_ls240-sim_1	0.30%	14.03 MiB / 7.582 GiB	0.18%	291 kB / 282 kB	0 B / 0 B	4
influxdb	12.66%	157.4 MiB / 7.582 GiB	2.03%	112 MB / 129 kB	16.4 kB / 117 MB	14
ocssiteconfigs_grafana_1	0.01%	15.96 MiB / 7.582 GiB	0.21%	24.5 kB / 3.8 kB	0 B / 22.1 MB	14
ocssiteconfigs_ocs-web_1	0.00%	3.297 MiB / 7.582 GiB	0.04%	14.4 kB / 656 B	0 B / 0 B	5

4 fake data agents

CONTAINER	CPU %	MEM USAGE / LIMIT	MEM %	NET I/O	BLOCK I/O	PIDS
ocssiteconfigs_ocs-fake-data-agent2_1	3.66%	37.29 MiB / 7.582 GiB	0.48%	49.7 kB / 7.32 MB	0 B / 12.3 kB	5
ocssiteconfigs_ocs-fake-data-agent_1	3.51%	37.3 MiB / 7.582 GiB	0.48%	41.8 kB / 7.29 MB	0 B / 12.3 kB	5
ocssiteconfigs_ocs-client_1	0.00%	452 KiB / 7.582 GiB	0.01%	6.21 kB / 656 B	0 B / 0 B	1
ocssiteconfigs_ocs-LSSIM_1	0.44%	35.56 MiB / 7.582 GiB	0.46%	153 kB / 180 kB	0 B / 61.4 kB	3
ocssiteconfigs_ocs-fake-data-agent3_1	3.65%	37.3 MiB / 7.582 GiB	0.48%	59.1 kB / 7.7 MB	0 B / 28.7 kB	5
ocssiteconfigs_ocs-fake-data-agent4_1	3.68%	37.29 MiB / 7.582 GiB	0.48%	49 kB / 7.88 MB	0 B / 16.4 kB	5
ocssiteconfigs_ls240-sim_1	0.26%	14.03 MiB / 7.582 GiB	0.18%	149 kB / 138 kB	0 B / 0 B	4
ocssiteconfigs_sisock-crossbar_1	15.09%	113.4 MiB / 7.582 GiB	1.46%	30.4 MB / 47.3 MB	0 B / 0 B	2
ocssiteconfigs_ocs-influx-publisher_1	75.09%	51.8 MiB / 7.582 GiB	0.67%	30.1 MB / 102 MB	0 B / 0 B	5
influxdb	16.07%	128.7 MiB / 7.582 GiB	1.66%	102 MB / 131 kB	369 kB / 107 MB	14
ocssiteconfigs_ocs-web_1	0.00%	3.277 MiB / 7.582 GiB	0.04%	14.7 kB / 656 B	0 B / 0 B	5
ocssiteconfigs_grafana_1	0.02%	15.34 MiB / 7.582 GiB	0.20%	27.6 kB / 3.49 kB	188 kB / 21.8 MB	14

8 fake data agents

CONTAINER	CPU %	MEM USAGE / LIMIT	MEM %	NET I/O	BLOCK I/O	PIDS
cssiteconfigs_ocs-client_1	0.00%	452 KiB / 7.582 GiB	0.01%	6.29 kB / 656 B	0 B / 0 B	1
ocssiteconfigs_ocs-fake-data-agent8_1	4.22%	37.3 MiB / 7.582 GiB	0.48%	101 kB / 16.1 MB	0 B / 12.3 kB	5
ocssiteconfigs_ocs-fake-data-agent_1	3.58%	37.3 MiB / 7.582 GiB	0.48%	75.6 kB / 16.3 MB	0 B / 32.8 kB	5
ocssiteconfigs_ocs-fake-data-agent6_1	3.64%	37.31 MiB / 7.582 GiB	0.48%	81.5 kB / 16.3 MB	0 B / 0 B	5
ocssiteconfigs_ocs-fake-data-agent2_1	3.66%	37.31 MiB / 7.582 GiB	0.48%	87.3 kB / 16.3 MB	0 B / 24.6 kB	5
ocssiteconfigs_ocs-fake-data-agent3_1	3.61%	37.3 MiB / 7.582 GiB	0.48%	83.3 kB / 16.3 MB	0 B / 0 B	5
ocssiteconfigs_ocs-fake-data-agent4_1	3.57%	37.3 MiB / 7.582 GiB	0.48%	81.4 kB / 16.9 MB	0 B / 20.5 kB	5
ocssiteconfigs_ocs-fake-data-agent5_1	3.65%	37.32 MiB / 7.582 GiB	0.48%	75.8 kB / 16.9 MB	0 B / 16.4 kB	5
ocssiteconfigs_ocs-fake-data-agent7_1	3.55%	37.31 MiB / 7.582 GiB	0.48%	95.4 kB / 17.3 MB	0 B / 12.3 kB	5
ocssiteconfigs_ocs-LSSIM_1	0.45%	35.57 MiB / 7.582 GiB	0.46%	338 kB / 406 kB	0 B / 69.6 kB	3
influxdb	11.35%	141.8 MiB / 7.582 GiB	1.83%	240 MB / 258 kB	0 B / 290 MB	14
ocssiteconfigs_ocs-influx-publisher_1	85.89%	268.7 MiB / 7.582 GiB	3.46%	132 MB / 241 MB	0 B / 12.3 kB	5
ocssiteconfigs_ocs-sisock-crossbar_1	27.44%	118.2 MiB / 7.582 GiB	1.52%	133 MB / 169 MB	0 B / 0 B	2
ocssiteconfigs_ls240-sim_1	0.29%	14.03 MiB / 7.582 GiB	0.18%	324 kB / 313 kB	0 B / 0 B	4
ocssiteconfigs_grafana_1	0.03%	16.25 MiB / 7.582 GiB	0.21%	29.1 kB / 3.82 kB	57.3 kB / 22 MB	12
ocssiteconfigs_ocs-web_1	0.00%	3.297 MiB / 7.582 GiB	0.04%	18.9 kB / 656 B	0 B / 0 B	5

16 fake data agents

CONTAINER	CPU %	MEM USAGE / LIMIT	MEM %	NET I/O	BLOCK I/O	PIDS
ocssiteconfigs_ocs-fake-data-agent3_1	3.68%	37.3 MiB / 7.582 GiB	0.48%	89.2 kB / 18.9 MB	0 B / 28.7 kB	5
ocssiteconfigs_ocs-fake-data-agent5_1	5.46%	37.3 MiB / 7.582 GiB	0.48%	118 kB / 18.9 MB	0 B / 20.5 kB	5
ocssiteconfigs_ocs-fake-data-agent_1	3.68%	37.3 MiB / 7.582 GiB	0.48%	124 kB / 19.1 MB	0 B / 20.5 kB	5
ocssiteconfigs_ocs-fake-data-agent16_1	3.71%	37.31 MiB / 7.582 GiB	0.48%	154 kB / 19.1 MB	0 B / 0 B	5
ocssiteconfigs_ocs-fake-data-agent2_1	3.65%	37.3 MiB / 7.582 GiB	0.48%	151 kB / 18.8 MB	0 B / 24.6 kB	5
ocssiteconfigs_ocs-fake-data-agent4_1	3.38%	37.31 MiB / 7.582 GiB	0.48%	92.6 kB / 18.6 MB	0 B / 0 B	5
ocssiteconfigs_ocs-fake-data-agent10_1	3.62%	37.31 MiB / 7.582 GiB	0.48%	104 kB / 18.6 MB	0 B / 28.7 kB	5
ocssiteconfigs_ocs-fake-data-agent12_1	3.69%	37.31 MiB / 7.582 GiB	0.48%	136 kB / 19 MB	0 B / 28.7 kB	5
ocssiteconfigs_ocs-client_1	0.00%	448 KiB / 7.582 GiB	0.01%	15.7 kB / 656 B	131 kB / 0 B	1
ocssiteconfigs_ocs-fake-data-agent15_1	5.58%	37.31 MiB / 7.582 GiB	0.48%	155 kB / 19.1 MB	0 B / 0 B	5
ocssiteconfigs_ocs-fake-data-agent7_1	3.62%	37.31 MiB / 7.582 GiB	0.48%	132 kB / 19 MB	0 B / 0 B	5
ocssiteconfigs_ocs-fake-data-agent6_1	3.65%	37.3 MiB / 7.582 GiB	0.48%	138 kB / 20.2 MB	0 B / 41 kB	5
ocssiteconfigs_ocs-fake-data-agent8_1	3.56%	37.31 MiB / 7.582 GiB	0.48%	101 kB / 20.2 MB	0 B / 32.8 kB	5
ocssiteconfigs_ocs-fake-data-agent9_1	3.67%	37.31 MiB / 7.582 GiB	0.48%	135 kB / 19.7 MB	0 B / 20.5 kB	5
ocssiteconfigs_ocs-fake-data-agent14_1	3.67%	37.3 MiB / 7.582 GiB	0.48%	106 kB / 19.5 MB	0 B / 32.8 kB	5
ocssiteconfigs_ocs-fake-data-agent13_1	3.71%	37.3 MiB / 7.582 GiB	0.48%	157 kB / 21.3 MB	0 B / 16.4 kB	5
ocssiteconfigs_ocs-fake-data-agent11_1	3.62%	37.31 MiB / 7.582 GiB	0.48%	112 kB / 21.1 MB	0 B / 0 B	5
ocssiteconfigs_ocs-LSSIM_1	0.31%	35.57 MiB / 7.582 GiB	0.46%	419 kB / 500 kB	0 B / 53.2 kB	3
ocssiteconfigs_sisock-crossbar_1	67.38%	119.7 MiB / 7.582 GiB	1.54%	312 MB / 351 MB	0 B / 0 B	2
ocssiteconfigs_ocs-web_1	0.00%	3.277 MiB / 7.582 GiB	0.04%	23.1 kB / 656 B	213 kB / 0 B	5
ocssiteconfigs_grafana_1	0.01%	16.77 MiB / 7.582 GiB	0.22%	35.2 kB / 3.88 kB	90.1 kB / 21.4 MB	14
ocssiteconfigs_ls240-sim_1	0.24%	14 MiB / 7.582 GiB	0.18%	406 kB / 386 kB	0 B / 0 B	4
ocssiteconfigs_ocs-influx-publisher_1	81.50%	850.6 MiB / 7.582 GiB	10.96%	309 MB / 275 MB	49.2 kB / 0 B	5
influxdb	13.62%	136.6 MiB / 7.582 GiB	1.76%	274 MB / 818 kB	655 kB / 314 MB	14

16 fake data agents, after a while

I guess some of the processes ran out of memory and probably the computer is spending all of its time swapping. Yuck.

CONTAINER	CPU %	MEM USAGE / LIMIT	MEM %	NET I/O	BLOCK I/O	PIDS
ocssiteconfigs_ocs-fake-data-agent16_1	0.30%	103.3 MiB / 7.582 GiB	1.33%	1.42 MB / 188 MB	14.4 MB / 0 B	5
ocssiteconfigs_ocs-fake-data-agent12_1	0.15%	138 MiB / 7.582 GiB	1.78%	1.38 MB / 188 MB	1.01 MB / 28.7 kB	5
ocssiteconfigs_ocs-client_1	0.00%	428 KiB / 7.582 GiB	0.01%	24.5 kB / 656 B	131 kB / 0 B	1
ocssiteconfigs_ocs-fake-data-agent15_1	0.15%	121.6 MiB / 7.582 GiB	1.57%	1.38 MB / 188 MB	2.43 MB / 537 kB	5
ocssiteconfigs_ocs-fake-data-agent7_1	0.15%	99.83 MiB / 7.582 GiB	1.29%	1.37 MB / 187 MB	10.5 MB / 0 B	5
ocssiteconfigs_ocs-fake-data-agent6_1	0.15%	135.2 MiB / 7.582 GiB	1.74%	1.11 MB / 189 MB	2.14 MB / 41 kB	5
ocssiteconfigs_ocs-fake-data-agent8_1	0.15%	132.5 MiB / 7.582 GiB	1.71%	815 kB / 189 MB	2.33 MB / 1.49 MB	5
ocssiteconfigs_ocs-fake-data-agent14_1	0.48%	99.76 MiB / 7.582 GiB	1.28%	859 kB / 188 MB	4.38 MB / 32.8 kB	5
ocssiteconfigs_ocs-LSSIM_1	0.23%	18.68 MiB / 7.582 GiB	0.24%	3.99 MB / 4.81 MB	14.3 MB / 2.9 MB	2
ocssiteconfigs_sisock-crossbar_1	0.01%	618.8 MiB / 7.582 GiB	7.97%	3.02 GB / 2.71 GB	105 MB / 0 B	2
ocssiteconfigs_ocs-web_1	0.00%	3.145 MiB / 7.582 GiB	0.04%	31.9 kB / 656 B	213 kB / 0 B	5
ocssiteconfigs_grafana_1	0.02%	7.465 MiB / 7.582 GiB	0.10%	56.3 kB / 8.29 kB	245 MB / 21.5 MB	14
ocssiteconfigs_ls240-sim_1	0.27%	5.562 MiB / 7.582 GiB	0.07%	3.83 MB / 3.84 MB	1.06 MB / 0 B	4
ocssiteconfigs_ocs-influx-publisher_1	1.24%	4.747 GiB / 7.582 GiB	62.61%	2.28 GB / 1.86 GB	1.68 GB / 0 B	5
influxdb	0.06%	126.8 MiB / 7.582 GiB	1.63%	1.86 GB / 4.54 MB	1.57 GB / 2.72 GB	15

Some preliminary observations

- ocs-influx-publisher CPU usage gets high, but plateaus.
 - ▶ Is it dropping packets? The fact that NET I/O ratio gets below one makes me think it likely is?
 - ▶ Memory usage explodes. Maybe because it can't keep up with incoming firehose?
- Crossbar CPU scales with number of FakeDataAgents.
 - ▶ Didn't go past 16 agents, but the trend is clear.
 - ▶ A quick look at crossbar.io suggests that the router should be using PyPy rather than CPython for best performance (assuming set up with CPython but did not check)
- Didn't measure this for all configurations, but load average on machine started around 6 with 16 fake data agents (not much else going on this machine), and after a while, ran out of ram, load average exploded, no work was being done.
- This test might be too pessimistic, and although I've read some documentation and browsed through a bit of a code, I'm still a novice at this so could be doing something deeply wrong.

