Testing summary

MAMAS 0.7.0 performance tests.

MAMAS 0.7.0 follows performance tests of version 0.6.0. The results should be comparable when the new features are not enabled (especially DCL which may impact performance noticeably when a big number of devices is sending data.

There are three different test types in MAMAS testing:

- stability tests ... runs for a number of hours.
- benchmark tests ... to measure different HW / VMs for comparability and clue for sizing deployments.
- peek tests ... to discover MAMAS SW and deployments limits and breaking points.

Next to the test results, HW specs for environments, test description and methodology is given in the document.

The MAMAS 0.7.0 early adopter 2 is designed for up to mid-sized deployments (up to 30000 devices).
As such, not all defined tests were performed. Also, not all environments were ready at the time of MAMAS 0.7.0 testing.
Find out more in test overview.

General MAMAS 0.7.0 performance testing verdict:

| | | |
|---|---|---|
| PASSED | 2023-01-15 | Initial test set results: Stab-2-2:C1; Stab-4-2:C2; Stab-8-2:C3; Stab-16-2:C4; Stab-32-2:H1; Bench-4-noSSL:C1,C2,C3,C4,H1; Bench-8-noSSL:C1,C2,C3,C4,H1; Bench-16-noSSL:C2,C3,C4,H1; Bench-4-SSL:C1,C2,C3,C4,H1; Bench-8-SSL: C1,C2,C3,C4,H1; Bench-16-SSL:C2,C3,C4,H1; |

Document updates:

Testing overview

As mentioned earlier (Testing summary), there are three types of performance tests. All are here because we need to measure:
- number of accepted and dropped requests (and resulting percentage) FOR ALL TESTS
- time for peak and benchmark tests
- memory usage for stability tests
- CPU utilization (%) for stability tests
- mamasd log review to make detail search for warnings and errors (for stability tests)
- number of pushing threads causing requests drops+resends (peak test)

Next to measurements there are a few important calculated metrics:
- ideal maximal throughput (benchmark test)
- typical maximal throughput (20% of ideal) (benchmark test)
- number of pushes per second (benchmark and peak test)
- maximal recommended # of devices per deployment type (HW / VM specs)

| Environment specs | | CPU cores | CPU threads | RAM MiB | stability test | benchmark | peak test |
|---|---|---|---|---|---|---|---|
| Linux container (Proxmox) | C1 | 1 | 1 | 512 | | | |
| | C2 | 2 | 2 | 1024 | | | |
| | C3 | 4 | 4 | 2048 | | | |
| | C4 | 8 | 8 | 4096 | | | |
| AWS instance | A1 (small) | | | | | | |
| | A2 (medium) | | | | | | |
| | A3 (big) | | | | | | |
| Baremetal server Core i7-8C-32G | H1 | 8 | 16 | 32768 | | | |

| Tests specs | ID | SSL? | # of devices | # of threads | # of loops | # of minutes | Human-time | Notes |
|---|---|---|---|---|---|---|---|---|
| Stability | Stab-2-2 | Yes | 1000 | 2 | N/A | 120 | 2 hours | |
| | Stab-4-2 | Yes | 3000 | 4 | N/A | 120 | 2 hours | |
| | Stab-8-2 | Yes | 10000 | 8 | N/A | 120 | 2 hours | |
| | Stab-16-2 | Yes | 30000 | 16 | N/A | 120 | 2 hours | |

| | | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | Stab-4-48 | Yes | 3000 | 4 | N/A | 2880 | 48 hours | |
| | Stab-4-672 | Yes | 3000 | 4 | N/A | 40320 | 4 weeks | |
| Benchmark | Bench-4-noSSL | No | 3000 | 4 | 10 | N/A | | |
| | Bench-8-noSSL | No | 3000 | 4 | 10 | N/A | | New in 0.7.0 |
| | Bench-16-noSSL | No | 10000 | 16 | 10 | N/A | | |
| | Bench-32-noSSL | No | 30000 | 32 | 10 | N/A | | |
| | Bench-4-SSL | Yes | 3000 | 4 | 10 | N/A | | |
| | Bench-8-SSL | Yes | 3000 | 4 | 10 | N/A | | New in 0.7.0 |
| | Bench-16-SSL | Yes | 10000 | 16 | 10 | N/A | | |
| | Bench-32-SSL | Yes | 30000 | 32 | 10 | N/A | | |
| Peak | PeakR-64-max | No | 60000 | 64 | N/A | 10 | | |
| | PeakR-256-max | No | 60000 | 256 | N/A | 10 | | |
| | PeakR-1024-max | No | 60000 | 1024 | N/A | 10 | | |
| | PeakD-64-max | No | 60000 | 64 | N/A | 10 | | |
| | PeakD-256-max | No | 60000 | 256 | N/A | 10 | | |
| | PeakD-1024-max | No | 60000 | 1024 | N/A | 10 | | |
| | PeakF-64-max | No | 60000 | 64 | N/A | 10 | | |
| | PeakF-256-max | No | 60000 | 256 | N/A | 10 | | |
| | PeakF-1024-max | No | 60000 | 1024 | N/A | 10 | | |

| Test run map 0.7.0 | C1 | C2 | C3 | C4 | A1 | A2 | A3 | H1 |
|---|---|---|---|---|---|---|---|---|
| Stab-2-2 | Yes | | | | | | | |
| Stab-4-2 | | Yes | | | | | | |
| Stab-8-2 | | | Yes | | | | | |
| Stab-16-2 | | | | Yes | | | | Yes |
| Stab-4-48 | | | | | | | | |
| Stab-4-672 | | | | | | | | |
| Bench-4-noSSL | Yes | Yes | Yes | Yes | | | | Yes |
| Bench-8-noSSL | Yes | Yes | Yes | Yes | | | | Yes |
| Bench-16-noSSL | No | Yes | Yes | Yes | | | | Yes |
| Bench-32-noSSL | No | No | No | No | | | | Yes |
| Bench-4-SSL | Yes | Yes | Yes | Yes | | | | Yes |
| Bench-8-SSL | Yes | Yes | Yes | Yes | | | | Yes |

Testing overview

| | | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| Bench-16-SSL | No | Yes | Yes | Yes | | | | Yes |
| Bench-32-SSL | No | No | No | No | | | | No |
| PeakR-64-max | | | | | | | | |
| PeakR-256-max | | | | | | | | |
| PeakR-1024-max | | | | | | | | |
| PeakD-64-max | | | | | | | | |
| PeakD-256-max | | | | | | | | |
| PeakD-1024-max | | | | | | | | |
| PeakF-64-max | | | | | | | | |
| PeakF-256-max | | | | | | | | |
| PeakF-1024-max | | | | | | | | |

| Stability tests | | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| ID / Environment → | | | Stab-2-2 / C1 | Stab-4-2 / C2 | Stab-8-2 C3 | Stab-16-2 C4 | Stab-16-2 H1 | Stab-4-48 H1 |
| Test specs | | Devices | 1000 | 3000 | 10000 | 30000 | 30000 | |
| | | Emulator threads | 2 | 4 | 8 | 16 | 16 | |
| | | Set: Time / minutes | 120 | 120 | 120 | 120 | 120 | |
| | | Set loop time / sec | 200 | 200 | 200 | 200 | 200 | |
| HW specs | | CPU cores | 1 | 2 | 4 | 8 | 8 | |
| | | RAM / MiB | 512 | 1024 | 2048 | 4096 | 32768 | |
| Resources | Pre-test | Available MiB OS | 275.00 | 804.00 | 1600.00 | 2900.00 | 28000.00 | |
| | | Used MiB OS | 236 | 219 | 451 | 1100 | 2000 | |
| | | mamasd RSS (MiB) | 37 | 37 | 37 | 37 | 40 | |
| | Post-test | Available MiB OS | 237 | 633 | 1100 | 1800 | 28000 | |
| | | Used MiB OS | 274 | 390 | 870 | 2200 | 2600 | |
| | | mamasd RSS (MiB) | 57 | 95 | 271 | 665 | 491 | |
| | | machine load/uptime | 1.24 | 1.69 | 2.33 | 2.53 | 1.37 | |
| | | machine load per CPU | 1.24 | 0.85 | 0.58 | 0.32 | 0.17 | |
| Results | Test numbers | Processesd loops | 31 | 30 | 28 | 28 | 29 | |
| | | Total time | | | 7336 | 7311 | 7345 | |
| | | Loop time / sec | 0 | 0 | 262 | 261 | 253 | |
| | | Data-pushes | 31000 | 90000 | 280000 | 840000 | 870000 | |
| | Errors | Data ERR/retries | 0 | 0 | 0 | 0 | 0 | |
| | | Crashes | 0 | 0 | 0 | 0 | 0 | |
| | | Reboots | 0 | 0 | 0 | 0 | 0 | |
| | | Log entries | 0 | 1 | 1 | 0 | 1 | |
| | | Log analysis | | 1 | 1 | 0 | 1 | |
| | Calculations | RSS increase MiB | 20 | 58 | 234 | 628 | 451 | |
| | | RSS increase % | 54.05 | 156.76 | 632.43 | 1697.30 | 1127.50 | |
| | | MiB per device | 0.020 | 0.019 | 0.023 | 0.021 | 0.015 | |
| | | RAM utilization % | 53.52 | 38.09 | 22.02 | 26.86 | 6.10 | |
| | | Max safe # of devs | 1869 | 7877 | 45410 | 111709 | 491520 | |
| | Summary | RAM utilization note | significant | reasonable | ok | ok | perfect (low) | |
| | | CPU utilization note | significant | utilized | utilized | reasonable | perfect (low) | |
| | | Verdict | PASSED | PASSED | PASSED | PASSED | PASSED | TBD |

| Log analysis explained | |
|---|---|
| 1 | WARN: Mid-air collision when creating default profile, skipping… : when multiple new devices with same default profile pushes data for the first time at the same time, only one default profile is created. Shows especially under artificial conditions (all devices with same pair of device type X FW version). Not an error. |

| Benchmark tests | | bench-4-SSL | bench-8-SSL | bench-16-SSL | bench-32-SSL | bench-4-noSSL | bench-8-noSSL | bench-16-noSSL |
|---|---|---|---|---|---|---|---|---|
| Devices | | 3000 | 3000 | 10000 | 30000 | 3000 | 3000 | 10000 |
| Loops | | 10 | 10 | 10 | 10 | 10 | 10 | 10 |
| Emulator Threads | | 4 | 8 | 16 | 32 | 4 | 8 | 16 |
| Regular loop time | | 200 | 200 | 200 | 200 | 200 | 200 | 200 |
| Container 1 | CPUs / MiBs | 1 / 512 | | | | | | |
| | Avail RAM pre | 246 | 246 | | | 246 | 246 | |
| | Avail RAM post | 176 | 117 | | | 69 | N/A | |
| | Load post | 3.21 | 3.36 | | | 3.1 | 11.82 | |
| | Load post / CPU | 3.21 | 3.36 | | | 3.1 | 11.82 | 0 |
| | Total pushes | 30000 | 30000 | | | 30000 | 30000 | 100000 |
| | Failed pushes | 0 | 0 | | | 0 | 23618 | |
| | Failed % | 0.00 | 0.00 | #DIV/0! | #DIV/0! | 0.00 | 78.73 | 0.00 |
| | mamasd restarts | 0 | 0 | | | 0 | yes as noDB | |
| | DB errors | 0 | 0 | | | 0 | Multiple + OOM | |
| | Elapsed time | 210.43 | 178.20 | | | 87.08 | 473.68 | |
| | Pushes / sec | 142.57 | 168.35 | #DIV/0! | #DIV/0! | 344.51 | 13.47 | #DIV/0! |
| | Ideal max # devs | 28513 | 33670 | #DIV/0! | #DIV/0! | 68902 | 2695 | #DIV/0! |
| | Rec. max # devs | 5703 | 6734 | #DIV/0! | #DIV/0! | 13780 | 539 | #DIV/0! |
| | Notes | | | | | | DB OOM | |
| Container 2 | CPUs / MiBs | 2 / 1024 | | | | | | |
| | Avail RAM pre | 647.00 | 567.00 | 374.00 | | 468.00 | 460.00 | 421.00 |
| | Avail RAM post | 571.00 | 545.00 | 406.00 | | 450.00 | 440.00 | 345.00 |
| | Load post | 1.40 | 1.80 | 2.65 | | 2.64 | 2.57 | 2.44 |
| | Load post / CPU | 0.70 | 0.90 | 1.33 | 0.00 | 1.32 | 1.29 | 1.22 |
| | Total pushes | 30000.00 | 30000.00 | 100000.00 | 300000.00 | 30000.00 | 30000.00 | 100000.00 |
| | Failed pushes | 0.00 | 0.00 | 0.00 | | 0.00 | 0.00 | 0.00 |
| | Failed % | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 |
| | mamasd restarts | 0.00 | 0.00 | 0.00 | | 0.00 | 0.00 | 0.00 |
| | DB errors | 0.00 | 0.00 | 0.00 | | Yes, 69 | 0.00 | 0.00 |
| | Elapsed time | 171.71 | 114.66 | 342.60 | | 88.50 | 62.22 | 174.85 |
| | Pushes / sec | 174.71 | 261.64 | 291.89 | #DIV/0! | 338.98 | 482.16 | 571.92 |
| | Ideal max # devs | 34943 | 52329 | 58377 | #DIV/0! | 67797 | 96432 | 114384 |
| | Rec. max # devs | 6989 | 10466 | 11675 | #DIV/0! | 13559 | 19286 | 22877 |

| | | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | Notes | | | TG border | | | | |
| Container 3 | CPUs / MiBs | 4 / 2048 | | | | | | |
| | Avail RAM pre | 1400 | 1400 | 1300 | 1200 | 1200 | 1200 | 1100 |
| | Avail RAM post | 1400 | 1400 | 1300 | 831 | 1200 | 1200 | 1100 |
| | Load post | 1.34 | 1.45 | 1.75 | 2.4 | 2.11 | 1.9 | 1.73 |
| | Load post / CPU | 0.335 | 0.3625 | 0.4375 | 0.6 | 0.5275 | 0.475 | 0.4325 |
| | Total pushes | 30000 | 30000 | 100000 | 300000 | 30000 | 30000 | 100000 |
| | Failed pushes | 0 | 0 | 0 | 3494 | 0 | 0 | 0 |
| | Failed % | 0.00 | 0.00 | 0.00 | 1.16 | 0.00 | 0.00 | 0.00 |
| | mamasd restarts | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| | DB errors | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| | Elapsed time | 186.21 | 104.81 | 338.34 | 1019.07 | 82.55 | 59.17 | 182.12 |
| | Pushes / sec | 161.11 | 286.23 | 295.56 | 290.96 | 363.42 | 507.01 | 549.09 |
| | Ideal max # devs | 32222 | 57246 | 59112 | 58191 | 72683 | 101403 | 109818 |
| | Rec. max # devs | 6444 | 11449 | 11822 | 11638 | 14537 | 20281 | 21964 |
| | Notes | | | TG border | TG overload | | | |
| Container 4 | CPUs / MiBs | 8 / 4069 | | | | | | |
| | Avail RAM pre | 2500 | 2500 | 2400 | 2200 | 2600 | 2600 | 2500 |
| | Avail RAM post | 2400 | 2400 | 2400 | 2100 | 2500 | 2500 | 2500 |
| | Load post | 1.52 | 1.83 | 2.31 | 2.9 | 2.71 | 2.54 | 2.36 |
| | Load post / CPU | 0.19 | 0.22875 | 0.28875 | 0.3625 | 0.33875 | 0.3175 | 0.295 |
| | Total pushes | 30000 | 30000 | 100000 | 300000 | 30000 | 30000 | 100000 |
| | Failed pushes | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| | Failed % | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 |
| | mamasd restarts | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| | DB errors | 0 | 0 | 0 | 0 | 0 | 77 | 0 |
| | Elapsed time | 158.14 | 103.93 | 336.37 | 1020.23 | 87.79 | 51.48 | 181.01 |
| | Pushes / sec | 189.71 | 288.66 | 297.30 | 294.05 | 341.72 | 582.75 | 552.46 |
| | Ideal max # devs | 37941 | 57731 | 59459 | 58810 | 68345 | 116550 | 110491 |
| | Rec. max # devs | 7588 | 11546 | 11892 | 11762 | 13669 | 23310 | 22098 |
| | Notes | | | TG border | TG overload | | | |
| | CPUs / MiBs | 8 / 32768 | | | | | | |
| | Avail RAM pre | 29000 | 29000 | 29000 | 29000 | 28000 | 28000 | 28000 |
| | Avail RAM post | 29000 | 29000 | 28000 | 28000 | 28000 | 28000 | 28000 |

Benchmark tests

| | | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| Baremetal H1 | Load post | 0.45 | 0.68 | 1.2 | 1.64 | 1.4 | 1.25 | 1.24 |
| | Load post / CPU | 0.05625 | 0.085 | 0.15 | 0.205 | 0.175 | 0.15625 | 0.155 |
| | Total pushes | 30000 | 30000 | 100000 | 300000 | 30000 | 30000 | 100000 |
| | Failed pushes | 0 | 0 | 569 | 7306 | 0 | 0 | 701 |
| | Failed % | 0.00 | 0.00 | 0.57 | 2.44 | 0.00 | 0.00 | 0.70 |
| | mamasd restarts | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| | DB errors | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| | Elapsed time | 192.76 | 106.74 | 338.67 | 1028.72 | 98.50 | 56.05 | 183.04 |
| | Pushes / sec | 155.63 | 281.06 | 293.59 | 284.52 | 304.57 | 535.24 | 542.50 |
| | Ideal max # devs | 31127 | 56211 | 58719 | 56905 | 60914 | 107047 | 108500 |
| | Rec. max # devs | 6225 | 11242 | 11744 | 11381 | 12183 | 21409 | 21700 |
| | Notes | | | TG border | TG overload | | | |

| Notes explained | |
|---|---|
| TG border | Traffic Generator was almost fully utilized – its limit is around 310 pushes/second and it has 8CPU / 16 cores only. |
| TG overload | Traffic Generator was overloaded, results are not valid. For next benchmark test, we need to distribute TG over multiple machines. |
| DB OOM | DB was killed by OS due to OOM. Caused by the need of saving too much data to DB at one time (HTTP is faster than HTTP encryption). It is of no consequence to real situation though as the HTTP-only benchmark in 4 threads prooved that the container equivalent of 68000 devices whereas the container is recommended up to 3000 devices only. |

| bench-32-noSSL |
| --- |
| 30000 |
| 10 |
| 32 |
| 200 |
| |
| |
| |
| |
| |
| |
| |
| |
| #DIV/0! |
| |
| |
| |
| |
| #DIV/0! |
| #DIV/0! |
| #DIV/0! |
| |
| |
| |
| |
| |
| |
| 0.00 |
| 300000.00 |
| |
| 0.00 |
| |
| |
| |
| #DIV/0! |
| #DIV/0! |
| #DIV/0! |

| |
|---|
| |
| 1000 |
| 753 |
| 1.73 |
| 0.4325 |
| 300000 |
| 3981 |
| 1.33 |
| 0 |
| 0 |
| 531.75 |
| 556.69 |
| 111338 |
| 22268 |
| TG overload |
| |
| 2400 |
| 1900 |
| 2.68 |
| 0.335 |
| 300000 |
| 4519 |
| 1.51 |
| 0 |
| 0 |
| 534.83 |
| 552.48 |
| 110495 |
| 22099 |
| TG overload |
| |
| 28000 |
| 27000 |

| |
|---|
| 1.33 |
| 0.16625 |
| 300000 |
| 15391 |
| 5.13 |
| 0 |
| 0 |
| 544.35 |
| 522.84 |
| 104568 |
| 20914 |
| TG overload |

| |
|---|
| |
| |
| PS as there is no<br>r can handle load |

Peak tests are for 60000 devices which is outside of scope for MAMAS 0.7.0

Comparison

| Comparison with version 0.6.0 results | | | | | | | Notes |
|---|---|---|---|---|---|---|---|
| | | C1 | C2 | C3 | C4 | H1 Stab-4-48 for memory | |
| 0.6.0 | mamasd RSS (post) | 49 | 54 | 112 | | 61 | |
| | KiB per device (est.) | 15 | 7 | 8 | | 8 | |
| | OS used MiB (post, Stab-2-2) | 169 | 216 | 394 | | | |
| | Pushes / sec (bench-4-SSL) | 142 | 182 | 189 | 189 | 193 | |
| 0.7.0 (with DCL) | mamasd RSS (post) | 57 | 95 | 271 | 665 | | |
| | KiB per device (est.) | 20 | 19 | 23 | 21 | | |
| | OS used MiB (post, Stab-2-2) | 274 | 390 | 870 | 2200 | | |
| | Pushes / sec (bench-4-SSL) | 143 | 175 | 161 | 190 | | |
| Comparison 0.6.0→0.7.0 | RSS % (bigger=worse) | 116.33 | 175.93 | 241.96 | | 0.00 | 0.7.0 uses more memory for more features |
| | KiB per device % (bigger=worse) | 133.33 | 271.43 | 287.50 | | 0.00 | |
| | OS used MiB % (bigger=worse) | 162.13 | 180.56 | 220.81 | | #DIV/0! | |
| | Pushes / sec % (bigger = better) | 100.70 | 96.15 | 85.19 | 100.53 | 0.00 | Performace is more/less the same – added DCL and more data but compensated with some optimizations. |