Если вы видите что-то необычное, просто сообщите мне.

Анализ производительности Docker контейнеров с помощью встроенных инструментов

Контейнеризация меняет организацю развертывания и использвания програмного обеспечения. Вы можете развертывать почти любое ПО надежно всего с помощью одной команды. И с помощью платформы окрестрации вроде Kubernetes и DC/OS, даже производственное развертывание становится простым.

Возможно вы уже игрались с Docker, и запускали несколько контейнеров. Но одна вещь которую вы не встречали, это то как Docker ведет себя в зависимости от различных нагрузок.

Снаружи Docker контейнеры могут выглядить как черные ящики, но поянтно, что большинство людей хотят получать с них метрики и производить анализ.

В этой статье, мы настроим маленький CrateDB кластер, с Docker и затем пройдемся по полезным командам Docker, которые позволят нам взглянуть на производительность.

Начнем с небольшого вступления для CreateDB.

Настроим Docker кластер для CrateDB

CrateDB это открытая, распределенная SQL база данных которая упрощает хранение и анализ массивного количества данных машин в реальном времени. Она масштабируема горизонтально, высокодоступная и запускается в отказоустойчивых кластерах, которая работает очень хорошо в виртуальных и контейнеризированных окружениях.

У вас уже может быть Docker кластер с CrateDB который вы можете использовать. Или даже, запускать любые контейнеры.

Если вы хотите настроить маленький кластер CrateDB для экспериментов с с метриками производительности, вы можете следовать инструкции от CrateDB.

Docker метрики

Главные параметры производительности контейнеров, который нас интересует в этой статье это загрузка ЦПУ, память, блокировка I/O, и сеть.

Docker предоставляет различные возможности для получения этих метрик:

- Использовать команду `docker stats`
- Использование REST API выставленную Docker демоном
- Чтение cgroups для превдо файлов.

Однако, метрики покрывают эти механизмы, не одинаково.

Для примера, `docker stats` предоставляет верхнеуровневую картинку ресурсов, которую достаточно для большинства пользователей. В то время, как псевдо файлы cgroup предоставляют детальную аналитику которая может полезна для глубокого анализа производительности контейнера.

Мы обсудим эти три возможности.

Начнем с докер команд:

`docker stats`

Эта команда показывает живые денные всех запущенных контейнеров.

Отметим, что если вы указажет остановленный контейнер, команда выполнится успешно, но не будет никакого вывода.

Чтобы ограничить данные для одного или несколких контейнеров, вы можете указать список имен контейнеров или ID разделенные пробелом.

Для трех нодного кластера, вывод может выглядеть следующийм образом:

```
$ docker stats
CONTAINER       CPU %    MEM USAGE / LIMIT        MEM %     NET I/O             BLOCK
I/O           PIDS
2f2697df4b79    0.21%    336.2 MiB / 1.952 GiB    16.82%    21.7 MB / 8.51 MB    2.57 MB / 119
kB    48
9f71cde6529e    0.16%    295.1 MiB / 1.952 GiB    14.76%    42.9 kB / 3.81 kB    0 B / 119
kB        45
75b161da6562    0.21%    351.8 MiB / 1.952 GiB    17.60%    44.5 kB / 3.81 kB    59.4 MB / 119
kB    45
```

Давайте подробнее разеберм эти столбцы:

- CONTAINER - колонка отображает ID контейнера.
- CPU - отображает возможности хоста. Для примера если у вас есть два контейреа, оба будут использовать один и тот же процессор, и каждый будет использовать по максимуму, `docker stats` команда для каждого контейнера будет отображать "50%"

использования. Однако с точки зрения самого контейнера он использует его на всю мощь.

- MEM USAGE / LIMIT и MEM % колонка отражает количество памяти используемое конейнером,
  The MEM USAGE / LIMIT and MEM % columns display the amount of memory used by the container, along with the container memory limit, and the corresponding container utilization percentage.If there is no explicit memory limit set for the container, the memory usage limit will be the memory limit of the host machine.Note that like the CPU % column, these columns report on host utilization. The NET I/O column displays the total bytes received and transmitted over the network by the corresponding container.For example, in the above output, container 2f2697df4b79 received 21.7 MB and sent 8.51 MB of data. The BLOCK I/O section displays the total bytes written and read to the container file system. The PIDS column displays the number of kernel process IDs running inside the corresponding container.

Next, let's take a look at the REST APIs exposed by Docker daemon. REST API

The Docker daemon listens on unix:///var/run/docker.sock, which only allows local connections by the root user. When you launch Docker, however, you can bind it to a different port or socket.

Like docker stats, the REST API continuously reports a live stream of CPU, memory, and I/O data. However, the API provides longer, live-streaming chunks of JSON, with metrics about the container.

To see this yourself, access the API like so:

$ curl -v --unix-socket /var/run/docker.sock
http://localhost/containers/CONTAINER_ID/stats

Here, replace CONTAINER_ID with the ID of the container you want to inspect.

You should receive a JSON stream.

Here's what it looks like:

Analyzing Docker Container Performance: JSON stream

This is a bit of a mess. So let's run it through a JSON pretty printer and take a closer look.

There are several fields in the JSON data. For this post, we're only going to look at performance specific data.

First, the cpu_stats object:

"cpu_stats": { "cpu_usage": { "total_usage": 20902022446, "percpu_usage": [9406810955, 11495211491], "usage_in_kernelmode": 5040000000, "usage_in_usermode": 14470000000 }, "system_cpu_usage": 139558680000000, "online_cpus": 2, "throttling_data": { "periods":0, "throttled_periods":0, "throttled_time":0 } }

Let's look at its keys, one by one.

The cpu_usage contains an object with the following keys:

total_usage Total CPU usage in nanoseconds. percpu_usage Per core CPU usage in nanoseconds. A sum total of all the usage stats in this object. usage_in_kernelmode System CPU usage in nanoseconds. usage_in_usermode User CPU usage in nanoseconds.

Next up is system_cpu_usage. This value represents the host's cumulative CPU usage in nanoseconds. This includes user, system, and idle.

The online_cpus value represents the number of CPU core on the host machine.

CPU utilization is one of the key factors needed to judge the overall load on the system and as you can see above, the Docker daemon REST API provides comprehensive CPU usage stats, so you can monitor and adjust your deployment as needed.

Now, let's move on to to the memory_stats object:

"memory_stats": { "usage": 310312960, "max_usage": 328871936, "stats":{ "active_anon": 305885184, "active_file": 954368, "cache": 2039808, "dirty": 16384, "hierarchical_memory_limit": 9223372036854771712, "hierarchical_memsw_limit": 9223372036854771712, "inactive_anon": 0, "inactive_file": 1081344, "mapped_file": 139264, "pgfault": 154346, "pgmajfault": 0, "pgpgin": 152351, "pgpgout": 77175, "rss": 305881088, "rss_huge": 0, "swap": 0, "total_active_anon": 305885184, "total_active_file": 954368, "total_cache": 2039808, "total_dirty": 16384, "total_inactive_anon": 0, "total_inactive_file": 1081344, "total_mapped_file": 139264, "total_pgfault": 154346, "total_pgmajfault": 0, "total_pgpgin": 152351, "total_pgpgout": 77175,

"total_rss": 305881088, "total_rss_huge": 0, "total_swap": 0, "total_unevictable": 0, "total_writeback": 0, "unevictable": 0, "writeback": 0 }, "limit": 2096177152 }

There's a lot of data here, and we don't need to know what all of it means.

Here are the most important bits for getting started:

```
The cache value is the memory being used by the container that can be directly mapped to block devices. In simpler terms, this as a measure of file operations (open, read, write, and so on) being performed against the container file system.
The rss value is memory that doesn't correspond to anything mapped to the container file system. That includes stacks, heaps, and anonymous memory maps.
The mapped_file value is the memory mapped by the processes inside the container. Files are sometimes mapped to a segment of virtual memory to improve I/O performance.
The swap value is the amount of swap currently used by processes inside the container. Swap, as you may know, is file system based memory that is used when the physical memory (RAM) has run out.
```

Next up is the blkio_stats object:

"blkio_stats": { "io_service_bytes_recursive": [{ "major": 259, "minor": 0, "op": "Read", "value": 16039936 }, { "major": 259, "minor": 0, "op": "Write", "value": 122880 }, { "major": 259, "minor": 0, "op": "Sync", "value": 16052224 }] }

This object displays block I/O operations performed inside the container.

The io_service_bytes_recursive section contains the number of objects representing the bytes transferred to and from the container file system by the container, grouped by operation type.

Within each object, the first two fields specify the major and minor number of the device, the third field specifies the operation type (read, write, sync, or async), and the fourth field specifies the number of bytes. cgroups Pseudo Files

cgroups pseudo files are the fastest way to read metrics from Docker containers.

cgroups pseudo files do not require root access by default, so you can simply write tools around these files without any extra fuss.

Also, if you are monitoring many containers per host, cgroups pseudo files are usually the best approach because of their lightweight resource footprint.

The location of cgroups pseudo files varies based on the host operating system. On Linux machines, they are generally under /sys/fs/cgroup. In some systems, they may be under /cgroup instead.

To access cgroups pseudo files, you need to include the long ID of your container in the access path.

You can set the $CONTAINER_ID environment variable to the long ID of the container you are monitoring, like so:

export CONTAINER_ID=$(docker inspect --format="{{.Id}}" CONTAINER_NAME)

Here, replace CONTAINER_NAME with the name of your container.

Alternatively, you can set $CONTAINER_ID manually running docker ps --no-trunc and copying the long ID from the command output.

Check it worked, like so:

$ echo $CONTAINER_ID
3d4569e14470937cfeaeb8b32fd3f4e6fa47bbd83e397b3c44ba860854752692

Now we have that environment variable set, let's explore few of the pseudo files to see what's there.

We can start with memory metrics.

You can find the memory metrics in the memory.stat file, located in a directory that looks like this:

/sys/fs/cgroup/memory/docker/$CONTAINER_ID

So, for example:

$ cd /sys/fs/cgroup/memory/docker/$CONTAINER_ID $ cat memory.stat cache 13742080 rss 581595136 rss_huge 109051904 mapped_file 2072576 dirty 16384 writeback 0 pgpgin 386351 pgpgout 267577 pgfault 374820 pgmajfault 93 inactive_anon 40546304 active_anon 541061120

inactive_file 6836224 active_file 6893568 unevictable 0 hierarchical_memory_limit 9223372036854771712 total_cache 13742080 total_rss 581595136 total_rss_huge 109051904 total_mapped_file 2072576 total_dirty 16384 total_writeback 0 total_pgpgin 386351 total_pgpgout 267577 total_pgfault 374820 total_pgmajfault 93 total_inactive_anon 40546304 total_active_anon 541061120 total_inactive_file 6836224 total_active_file 6893568 total_unevictable 0

The first half of the file has the statistics for processes in the container.

The second half of the file (entries starting with total_) has stats for all processes running in the container, including sub-cgroups within the container.

Entries in this file fall into two broad categories: gauges and counters.

Entries like cache are gauges, meaning they can increase or decrease indicating the current value. Other entries, like pgfault, are counters and can only increase.

Next, let's take a look at CPU metrics.

CPU metrics are contained in the cpuacct.stat file, which is located in the following directory:

/sys/fs/cgroup/cpuacct/docker/$CONTAINER_ID

Let's take a look:

$ cd /sys/fs/cgroup/cpuacct/docker/$CONTAINER_ID $ cat cpuacct.stat user 6696 system 850

This file shows us CPU usage, accumulated by the processes of the container. This is broken down into user and system time.

User time (user) corresponds to time during which the processes were in direct control of the CPU, i.e. executing process code. Whereas system time (system) corresponds to the time during which the CPU was executing system calls on behalf of those processes.

Next, let's explore the I/O stats in cgroup files.

I/O metrics are present in blkio.throttle.io_service_bytes and blkio.throttle.io_serviced files present in this directory:

/sys/fs/cgroup/blkio/docker/$CONTAINER_ID

Change into the directory:

$ cd /sys/fs/cgroup/blkio/docker/$CONTAINER_ID

Then:

$ cat blkio.throttle.io_serviced 259:0 Read 3 259:0 Write 1 259:0 Sync 4 259:0 Async 0 259:0 Total 4 Total 4

This shows total count of I/O operations performed by the container under analysis.

$ cat blkio.throttle.io_service_bytes 259:0 Read 32768 259:0 Write 4096 259:0 Sync 36864 259:0 Async 0 259:0 Total 36864 Total 36864

This shows the total bytes transferred during all the I/O operations performed by the container.

Finally, let's look at how to extract network metrics from pseudo files. This is important as network metrics are not directly exposed by control groups. Instead, Docker provides per-interface metrics.

Since each container has a virtual ethernet interface, Docker lets you directly check the TX (transmit) and RX (receive) counters for this interface from inside the container.

Lets see how to do that.

First, fetch the PID of the process running inside the container:

$ export CONTAINER_PID=`docker inspect -f '{{ .State.Pid }}' $CONTAINER_ID`

Then read the file:

$ cat /proc/$CONTAINER_PID/net/dev

This should give you something like:

Inter-| Receive | Transmit face |bytes packets errs drop fifo frame compressed multicast|bytes packets errs drop fifo colls carrier compressed eth0: 44602 271 0 0 0 0 0 5742 83 0 0 0 0 0 0

This shows you the data transfer details for the container's virtual interface eth0. Wrap Up

In this post we took a look at the docker stats command, the Docker REST API, and cgroups pseudo files.

We learnt that there are multiple ways to get statistics from a Docker container. Which method you use will depend on your setup.

```
The docker stats command is good for small scale use, with a few containers running on a single host.
The Docker REST API is good when you have multiple containers running on multiple hosts, and you'd like to retrieve the stats remotely.
The cgroups pseudo files are the fastest and most efficient way to get stats, and are suitable for for large setups where performance is important.
```

While all these options are useful if you're planning to build your own tooling around Docker monitoring, there are several pre-built solutions, including Prometheus, cAdvisor, Scout, DataDog. We'll take a closer look at Docker health monitoring tools in the future.

Revision #4
Created 2021-09-10 07:29:26 UTC by gasick
Updated 2023-04-16 19:26:41 UTC by gasick