Drupal cache vs aiScaler vs Netscaler on the Amazon Cloud

Our Setup:

Drupal 7.23, clean install, anonymous and block cache enabled. We chose Drupal for this report, but similar results can be expected with Wordpress, Joomla, Magento or any other CMS

- m1.small, LAMP, No ADC (ADC=Application Delivery controller = reverse proxy)
- m1.medium, LAMP, No ADC
- t1.micro, aiScaler (unthrottled and throttled)
- m1.large, aiScaler
- m1.large, Netscaler VPX Platinum 10Mbps

ApacheBench shows how many requests per second the server is able to handle with a certain number of simultaneous connections. By combining the measurements, we can see how each solution scales. We found that Drupal does not scale well, even when using the integrated anonymous cache and LAMP optimizations according to the instance size. The results are based on a clean Drupal install, with a great amount of probability that the performance will get even worse by using plug-ins and a custom themes.

There are two major advantages of using aiScaler vs the built-in cache: higher scalability and independence on the back end code, plugins and themes. Even with a free micro-instance and aiScaler, you can get stable application performance under increasing load. The other advantage is that performance won't change with time when you begin to use Drupal more extensively. This allows you to save money by avoiding load balancing or bigger instances on the backend.

Throttling of the AWS micro-instance

We have done a separate test with the micro-instance to simulate a heavy load for 10 minutes and see how throttling influences ApacheBench results. From the graph (yellow line)we can see that t1.micro even under throttling still delivers 5x the performance of Drupal, without an ADC. One could expect much worse results from an instance, which wasn't designed for heavy workloads, but aiScaler shows to be very efficient in CPU usage, maintaining stability and speed.

Sample command used for testing:
ab -c 20 -n 5000 -H 'Accept-Encoding: gzip' http://site.com/
c – number of connections
n – number of total requests
H – we enable gzip to simulate real browser behavior

Test 1: ApacheBench

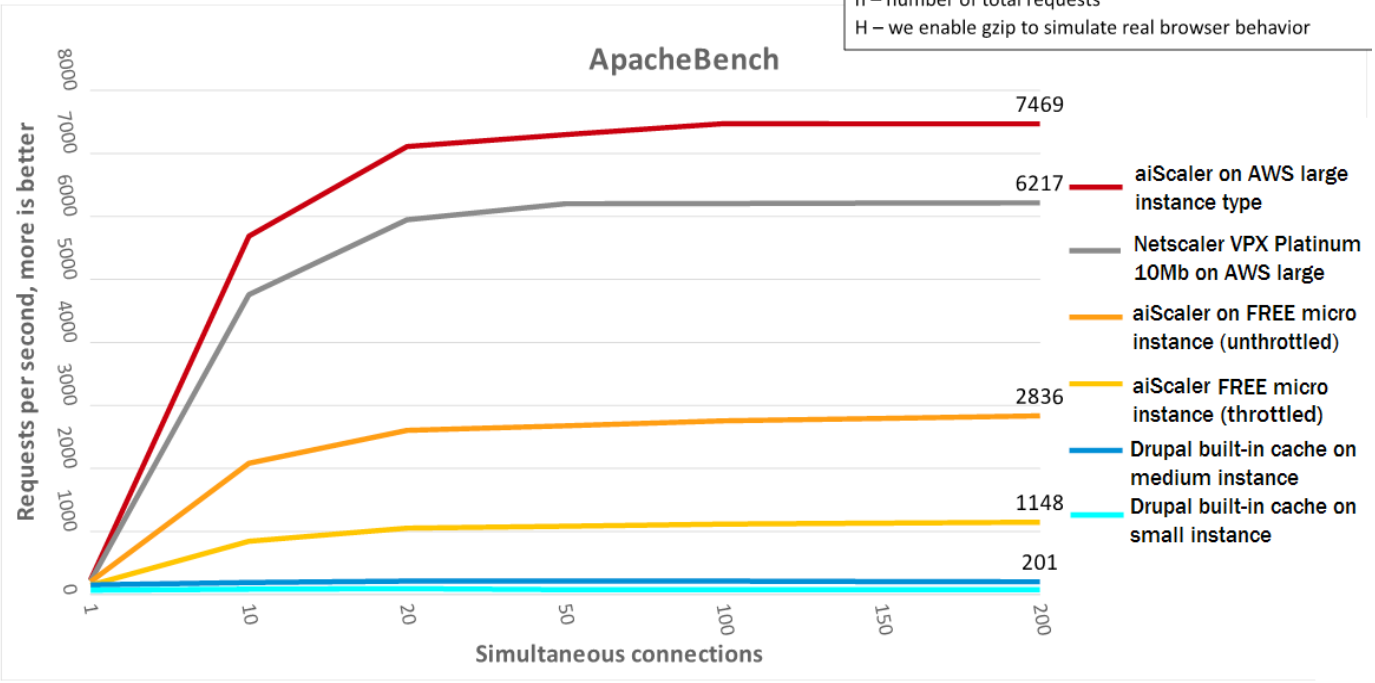

Test 2: Pingdom

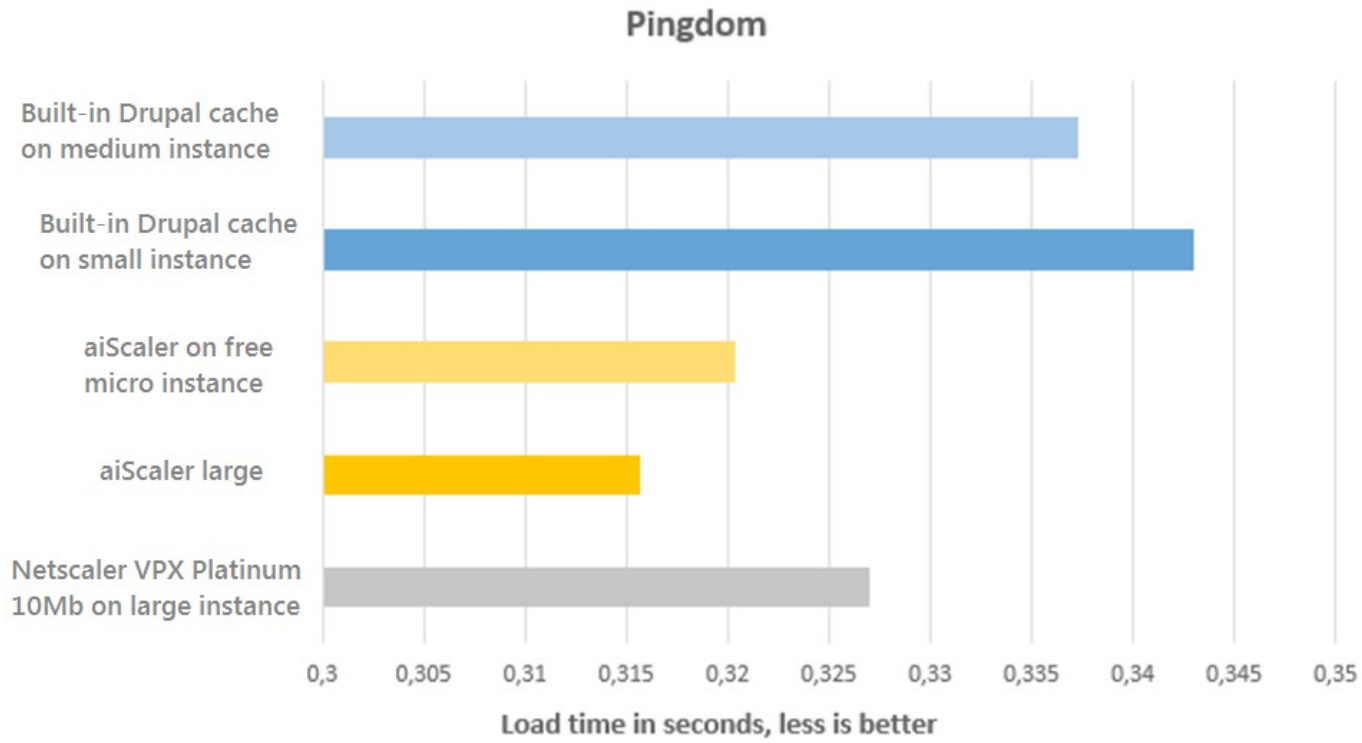

Pingdom shows how much time is needed to load the page. Normally a LAMP setup has to use backend services for each request, like Apache and MySQL, which decrease page loading times. aiScaler serves the contents directly from memory, which results in faster page loading times and reduced use of backend resources. The heavier page you use, the bigger your advantage.

Test 3: WebPagetest

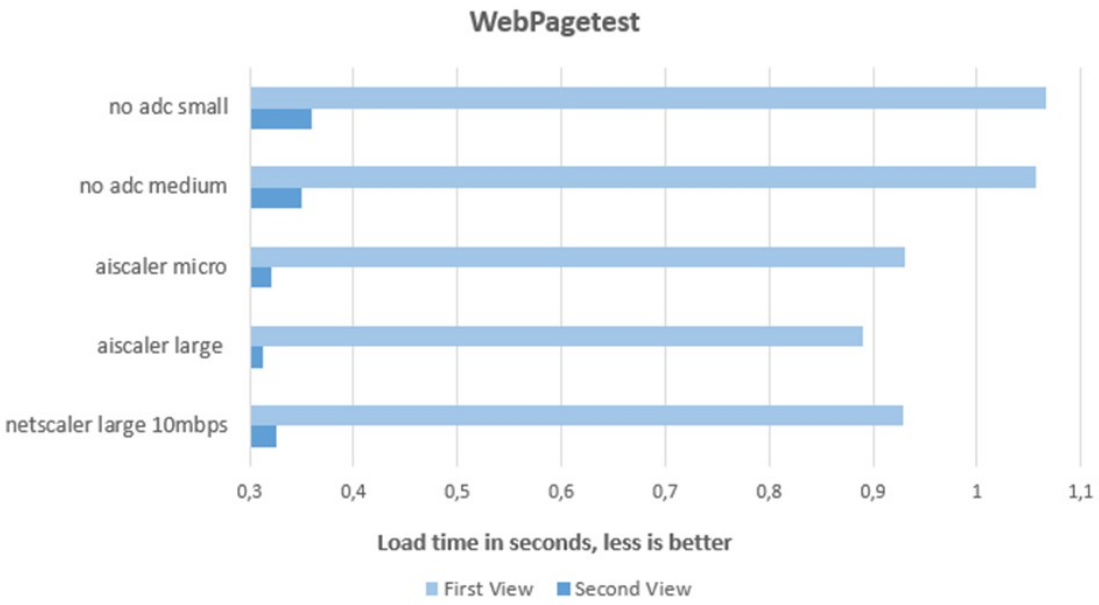

WebPagetest shows similar results to Pingdom, but it adds a first view time, for pages which are not cached yet by the ADC's. It's remarkable to see that the ADC's still deliver slightly lower loading times on uncached pages, which is probably due to better session management.

Conclusion:

aiScaler running on a free micro-instance, can serve 5 to 10 times as many RPS as the built-in Drupal cache. The exact number depends a lot on the throttling applied by Amazon, but even in the worst case scenario, it increased RPS by more than 5 times. We chose Drupal for this report, but similar results can be expected with Wordpress, Joomla, Magento or any other CMS.

We still recommend using a larger instance than the free micro-instance for the following reasons:

- Stability and performance. The micro-instance can be throttled by Amazon, so it should not be used by a professional website. Then the bigger instances have more CPU power. This doesn't linearly translate into performance, but it does mean better stability and performance under heavy load.
- Available memory: aiScaler runs completely in memory. A micro-instance only has 500mb of available RAM. That is not going to be enough for most average websites. Even a small instance will have 1.7GiB and a large instance 7.5GiB. In this report, we used a basic Drupal installation, with a low memory footprint. In real-life situations, the difference in performance will be bigger.
- Professional support. Users of the free instance can only use the online documentation. Paying users can apply for 2 hours of free personal support, with the option of buying extra hours to optimize their configuration.