

HiFi Caching

Hierarchical Filtering Cache Logic



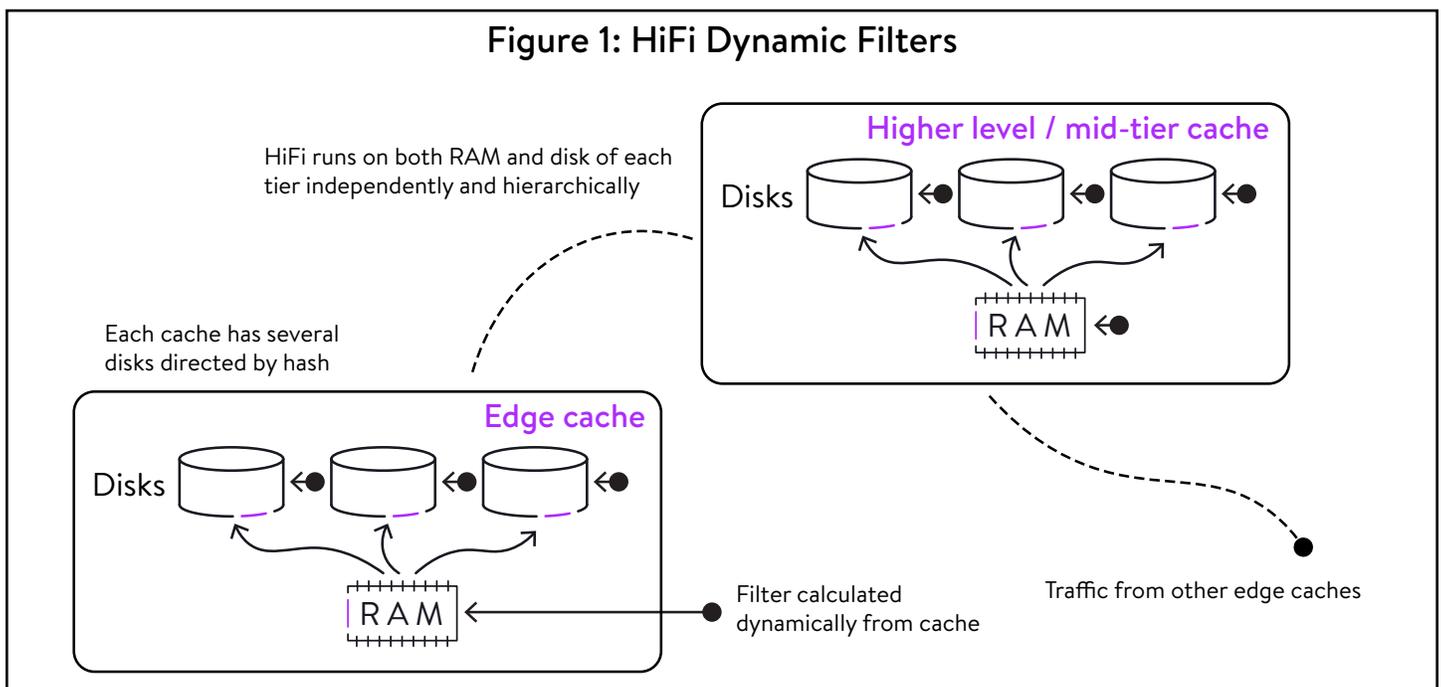
Velocix CDN solutions employ a proprietary caching algorithm designed to improve cache hit rates and deliver better overall system performance. The algorithm is based on hierarchical filtering cache logic, otherwise known as HiFi Caching, to drive higher stream density, reduce loading on local storage systems, and decrease demands placed on upstream caches, origin servers, and network resources.

Why HiFi Caching?

Despite 30 years of research in caching policies, most CDN caches still employ least recently used (LRU) as the asset eviction policy. Unfortunately, LRU policy produces suboptimal cache hit-rates in a variety of circumstances related to video streaming. Some open source caches such as Apache Traffic Server or Squid offer the option to use greedy dual size frequency (GDSF) or least frequency used with dynamic aging (LFU-DA) eviction policies instead of LRU. While these policies produce somewhat better hit-rates than LRU, they do not produce optimal hit rates either. They also do not reduce replacement rates, the rate at which the cache is re-written, as compared to the LRU policy. Reducing replacement rates is key to reducing disk I/O requirements and increasing the life of both spinning and solid state disks.

HiFi Caching Explained

Velocix developed a new caching algorithm called hierarchical filtering (HiFi), which produces superior cache hit-rates and significantly improves cache throughput. HiFi does this by introducing a dynamic filter at each cache in the CDN hierarchy that decides whether to admit a given asset into the cache. This filter is used at all caches, such as disk or RAM caches, at any of the caching hierarchy levels as shown in Figure 1.



The Results

Velocix has tested the performance of the HiFi algorithm in actual field deployments and verified that it delivers significantly improved performance. To showcase the benefits of the algorithm, we also conducted a simulation so we could accurately measure the benefits under controlled conditions. The simulation variables were based on a realistic set of requirements for a CDN supporting Cloud DVR and time shifted TV services. The simulation parameters were as follows:

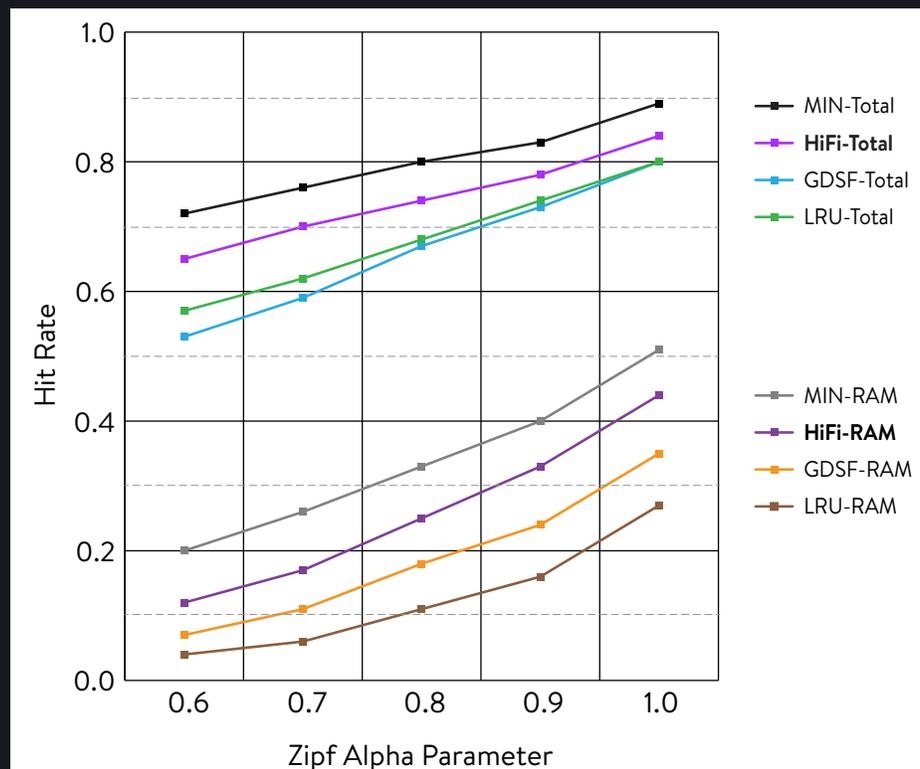
- 10,000 total CDN assets
- CDN edge cache with capacity for 2,000 assets
- 10% of assets replaced in the library every day
- 75% TV assets / 25% movie assets
- TV asset popularity decays by 40% each day
- Movie asset popularity decays by 10% each day

A dynamic environment such as this is difficult for any caching algorithm to achieve optimal hit rates, as asset popularity decays very rapidly. We compared the following caching policies under the same challenging conditions:

- LRU: The policy used most commonly used in CDN caches.
- GDSF: A policy option available in select caches.
- HiFi: An optimized caching algorithm unique to Velocix products.
- MIN: A theoretical boundary on the best hit rate achievable which uses a well known algorithm - Belady's MIN, which takes all future request rates into account to arrive at the absolute optimal caching decision. This algorithm is not practical, but serves as a useful bound on the optimality of algorithms.

Figure 2: Performance of Caching Algorithm

Figure 2 shows the average cache hit rates achieved by each policy at the RAM cache and at the total cache, which includes the disks.



Typically, the popularity of video assets have profiles similar to a Zipf function with alpha values between 0.7 and 0.8. Figure 2 shows that the Velocix HiFi algorithm outperforms LRU in this range by about 150% at the RAM level and about 10% at the total cache level. At caches that support a smaller percentage of the library, HiFi can improve hit rate even more. Using HiFi, replacement rates also decreased by 10 times in most scenarios. The increased cache hit rate and reduced replacement rates offer the following benefits for CDN operators:

- Reduced server hardware costs due to throughput improvements of 50% to 150% per cache.
- Lower power consumption as a result of fewer servers required to support the same traffic volume.
- Fewer maintenance issues because load on disk drives and SSDs has been reduced.
- Decreased costs for network equipment and origin server capacity since HiFi caches reduce demand on upstream CDN resources.
- Better consumer experience as a greater number of session requests can be fulfilled using local server resources.

The benefits of HiFi Caching only improve as content libraries grow. As library capacities grow, cache capacities often don't keep pace, creating more churn on the edge caches. Traditional caching policies like LRU can struggle and deliver much lower hit rates in these scenarios. HiFi Caching's performance characteristics make it an ideal choice for operators that intend to scale their content libraries.

In Summary

Velocix HiFi Caching is an advanced caching algorithm that offers significant performance and efficiency benefits over traditional techniques. In addition to the demonstrated advantages in lab tests, the solution has been proven in the field and is currently supporting high performance streaming video applications at many of the world's largest CDN operators and video service providers. HiFi Caching is unique to Velocix products and is integrated into the company's CDN products.

Please contact a Velocix salesperson to learn more about HiFi Caching and the Velocix CDN.