Ramping Up Web Server Memcached Capabilities with Hybrid-Core Computing

Convey Computer White Paper

CONVEY COMPUTER™

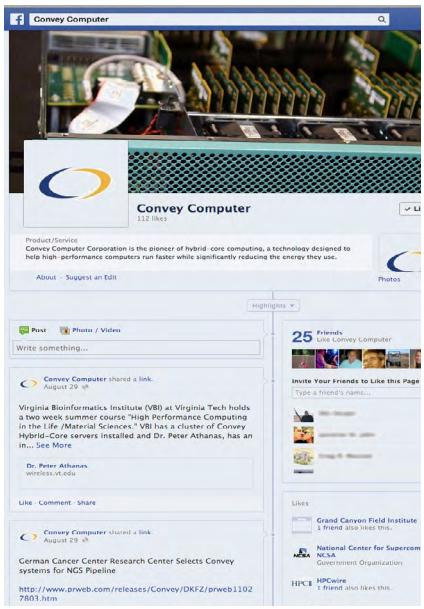

Ramping Up Web Server Memcached Capabilities with Hybrid-Core Computing

Contents

Introduction

Constantly increasing network traffic and rapidly growing datasets present ever increasing scaling challenges for memcached, a high-performance, distributed memory caching system used by most of the top web sites in the world. Convey's hybrid-core technology provides a solution.

The rapid growth of the Internet, and the emergence of hyperscale computing are placing new demands on the flow of web traffic. For example, companies like Amazon, Google, and Facebook have created hyperscale data centers to handle exponentially increasing traffic driven by social networking and on-line commerce. The problem of delivering content to consumers in a real-time manner is compounded by this explosive growth.

Enter Memcached

Memcached is a key technology employed to accelerate the delivery of content that resides/persists in mass storage systems.

Memcached, which stands for Memory Cache Daemon, is a general-purpose distributed memory caching system. It was developed in 2003 because LiveJournal, an early social blogging site, was experiencing rapid growth and its users were complaining about slow response times.

One of the major bottlenecks was the latency associated with reading from the database. To deal with this problem, the journal's founder, Brad Fitzgerald, created memcached. It was a smashing success, eliminating the bottleneck to the delight of LiveJournal users.

Today memcached is implemented in data centers around the world. More than 85% of the top 20 web sites – including such luminaries as YouTube, Zynga, Facebook, Flickr, Slashdot, Twitter and Wikipedia – have installed the cache on their servers. In fact, a full 50% of the top 5,000 web sites use memcached. In addition, many public and private cloud providers have installed the system as a front end to their database engines.

The cache is implemented as a network service. Its mission is simple – speed up dynamic, database driven web sites by caching data in RAM to reduce the number of times an external database or API has to be read. Clients issue "gets" to the memcached layer and if it "hits" in memcached, the content is delivered immediately to the client. If not, the clients ask the designated database to fetch the content, which can then be written into cache.

CONVEY COMPUTER™

Memcached uses a client-server architecture. The servers maintain a key-value associative array that is populated and queried by the clients. Keys can be up to 250 bytes long and values, by default, are limited to one megabyte in size.

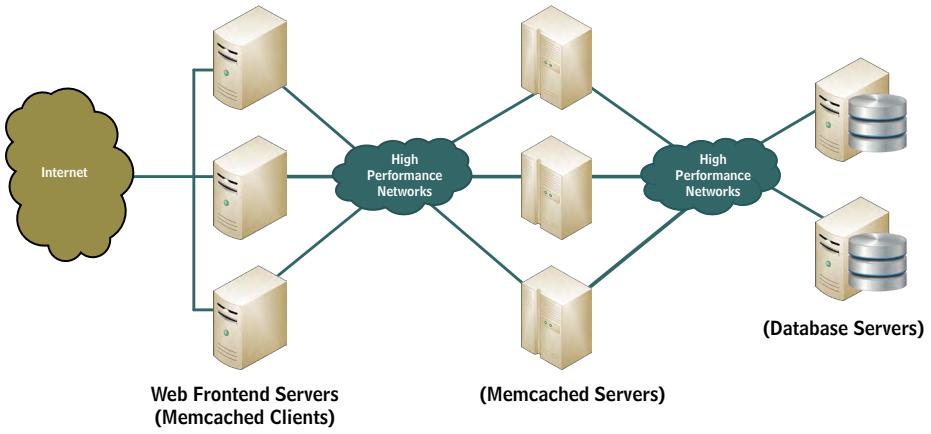

Figure 1. Memcached architecture.

When response times begin to creep into an unacceptable range, there can be dire consequences.

Memcached does not feature persistent memory storage and it has no built in security features. Distributed as open source software, the cache runs on standard x86 platforms with Linux usually serving as the operating system.

Rejuvenation Time

Memcached has been performing dutifully for more than a decade, but the software is beginning to show its age.

Raw computing power, fueled by Moore's Law, continues its exponential growth, spawning increasingly large datasets and ramping up traffic on the web. Everyone's favorite buzz word, big data, has become a reality and server clusters everywhere have to contend with its seemingly endless growth. One oft quoted phrased attributed to IBM holds that 90% of the information in the world was created in the last two years. And we're just getting started.

This unprecedented growth in web traffic is having an impact. Latency is rising, and consistently low latency cannot be assured. And consistency is essential – your response times may meet your criteria for worst-case average latency, but that average may be made up of widely fluctuating numbers as the short-term load increases. This makes for a less than optimum user experience.

One solution is to deploy more memcached servers, but this has its own set of problems... Adding more servers also means additional space, power, cooling, network infrastructure and capital expenditures.

When response times begin to creep into an unacceptable range, there can be dire consequences. For example, optimal memcached performance is critical for customer-facing applications. Companies that rely on providing their customers with a consistently good online experience to bring them back to their web site time and again are particularly vulnerable. One Convey client said that a delay of more than a few seconds to deliver content in response to customer requests can be disastrous. The latency is severe enough to risk driving customers away – and they may not return.

More is Less

One solution is to deploy more memcached servers, but this has its own set of problems. For example, suppose your data center's infrastructure includes a single server loaded with memcached middleware to accelerate your workloads across the cluster. But job traffic is inexorably increasing and response times are slowing.

To deal with the resultant latency and throughput problems, you decide to add a second server. Now you are splitting the job in two and dealing with multiple instances of memcached. This change in architecture requires corresponding modifications in how you manage your infrastructure. Adding more servers also means additional space, power, cooling, network infrastructure and capital expenditures. In the meantime, traffic continues to build and costly data center expansion results.

There must be a better way. And there is – it's called hybrid-core computing.

The Hybrid-Core Computing Solution

The solution is simplicity itself – ramp up the processing power of the memcached server to such an extent that it can easily handle an order of magnitude more transactions. With this additional operational headroom, memcached can provide the low latency, consistency, and throughput required for the most demanding web applications.

The solution is simple: The addition of hybrid-core capabilities provides an order of magnitude improvement in memcached throughput while maintaining sub-millisecond response.

The key is Convey's hybrid-core technology. Rather than adding more processors to your existing systems, the Convey approach focuses on the lowest common denominator – the individual processor – and makes it run faster. A lot faster.

Hybrid-core combines standard multi-core x86 processors with standard Field Programmable Gate Arrays. The FPGAs function as reconfigurable coprocessors optimized for handling data intensive applications. Computationally intensive portions of an application are implemented in hardware, delivering much higher throughput than a conventional processor core. A globally shared memory presents the same view of threads on both the x86 cores and the coprocessor.

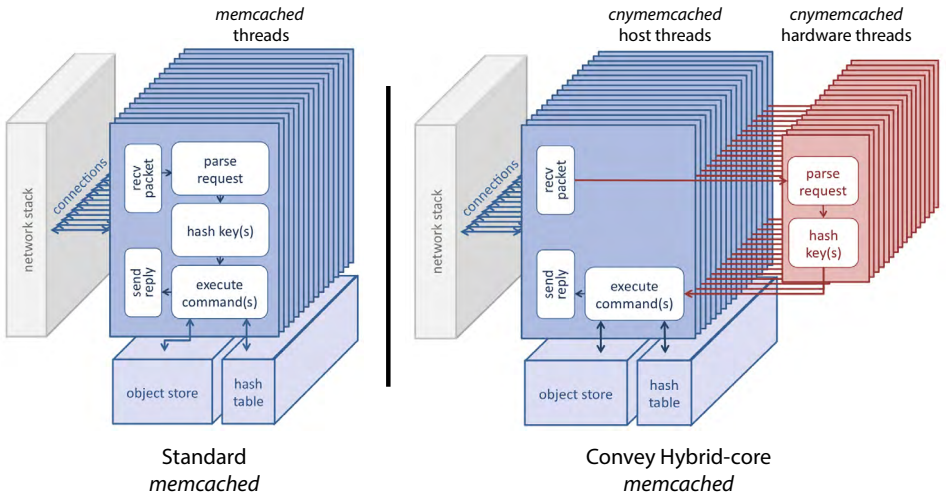

Figure 2. Comparison of "standard" memcached vs Convey hybrid-core implementation.

Using FPGAs to accelerate memcached is a natural fit. The string processing and hash functions required to process the keys in a memcached request can be implemented efficiently on FPGAs, using hardware pipelining and replication to process many keys concurrently. The x86 processors in the system can then be dedicated to network processing, ensuring maximum throughput. In the hybrid-core architecture of the Convey system, the x86 cores and the hardware pipelines on the coprocessor share memory, operating as one unified daemon incorporating both multithreaded execution on the host and custom hardware. This allows memcached to achieve orders of magnitude speedup in handling system requests.

The Convey memcached implementation is functionally identical to the standard memcached, supporting both ASCII and binary protocols, and can be dropped in as a direct replacement for existing memcached servers. Just point the clients at a Convey memcached appliance, and see the benefits of much higher throughput with sub-millisecond latencies.

...as the number of servers increases, the number of data items requested from each server is reduced... This problem is often referred to as the "multiget hole."

Removing the Scalability Barrier

As mentioned earlier, memcached finds itself struggling when confronted with today's data intensive applications. Generation of a typical web page might require hundreds of memcached accesses to data spread across many memcached servers. Most client libraries will optimize the access to those servers by "batching" all the requests to a particular server into one network request. But as the number of servers increases, the number of data items requested from each server is reduced. This can actually increase the total number of network requests required to populate the page. This problem is often referred to as the "multiget hole."

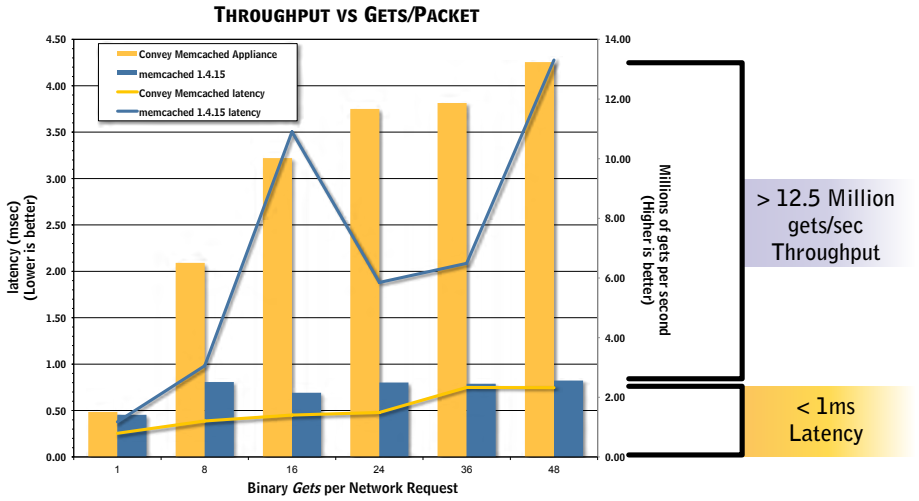

Figure 3. Bandwidth and performance comparison for various multiget packet sizes

The solution is simple: the addition of hybrid-core capabilities provides an order of magnitude improvement in memcached throughput while maintaining sub-millisecond response (Figure 3). The same throughput in gets per second can be sustained with a smaller number of servers, thus avoiding the multiget hole. The Convey solution provides the headroom needed to run the most demanding applications without having to add more memcached servers to the mix.

Summary of Benefits

The combination of memcached and Convey's hybrid-core technology yields a number of benefits.

This unique solution:

- Provides low latency – response times are not only reduced, but their peaks and valleys are minimized as well, making for a better customer experience
- Enables an order of magnitude faster throughput
- Permits significant reductions in power, space, and cooling requirements

Conclusions

Since it was first developed in 2003, memcached has been a singular success. Fully 85% of the top 20 web sites have implemented this open source, high performance, distributed memory caching system to speed up their dynamic web applications.

Now, on its 10 anniversary, memcached could use a tune up. Simply adding more memcached servers has proven to be a less than optimal solution. The trick is to take your existing memcached implementation and boost its capabilities by an order of magnitude or more, providing all the horsepower you need to handle real time caching during peak traffic flows.

The trick is to take your existing memcached implementation and boost its capabilities by an order of magnitude or more, providing all the horsepower you need to handle real time caching during peak traffic flows.

To do this, memcached needs some help, and Convey's hybrid-core platforms provide it. Accelerating memcached performance allows the system to quickly handle database calls, API calls or page rendering with consistently low latency and advanced throughput. Essentially, the Convey hybrid-core platform transforms memcached into a network appliance – all that's required is a change of IP address to begin routing the cached data through the high-speed memcached appliance.

Convey will continue to evolve the memcached/hybrid-core combination. Next generation configurations will see optimized network connections for even lower latency. Multi-get scalability limitations are being addressed. Eliminating the need for additional memcached servers helps keep capital and operating expenses under control. And Convey will continue to work with its customers to develop new solutions leveraging the power of FPGA acceleration and tailored personalities.

High speed memory caching was introduced 35 years ago and memcached is one of its direct descendants. With the addition of hybrid-core technology from Convey, memcached can look forward to a long and robust future.

Convey Computer Corporation
1302 E. Collins Boulevard
Richardson, Texas 75081
Phone: 214.666.6024 Fax: 214.576.9848
Toll Free: 866.338.1768
www.conveycomputer.com