



With Redis Labs, CAPI goes Mainstream, Big Time!

When IBM announced CAPI (Coherent Accelerator Processor Interface) exclusively on POWER8 in 2014, its practical use and the applications it benefited didn't immediately leap to mind. So, when Redis Labs announced that applying CAPI technology to in-memory database applications provided significant cost advantages over similar Intel-based systems; it caught our attention.

CAPI technology reduces latency by allowing a processor to directly access¹ data on a flash device. It operates at speeds far faster than typical I/O devices. CAPI supports as much as 40TB of flash memory. Previously, large Redis NoSQL databases on x86 servers would be at a severe cost disadvantage because the database had to reside in main memory which required a large number of x86 servers to handle a terabyte or more of data.

Redis Labs²

Redis Labs and IBM collaborated to implement a NoSQL solution using POWER8 and CAPI. It took several weeks of programming by Redis Labs using IBM provided APIs. Redis needs just 6 POWER8 systems to support a 40TB database. Table 1 (below) compares a possible Redis Labs POWER8 configured solution with a Dell configuration and an Amazon EC2 solution³. A Dell system with 512MB of main storage per server would need 80 servers to handle a 40TB database⁴. The Redis Labs solution needs only 6 P8 servers vs. the Dell-Intel's 80 servers. This is a 13 to 1 (80/6) server advantage ratio⁵ over the Dell!

ROI – a 40TB Redis deployment

	EC2	Dell	IBM P8 + FlashSystem (8247-22L)
# of servers	164	80	6
RAM	40TB	41TB	5TB
Flash	-	-	40TB
Persistence storage	164TB	1200TB	86TB
3yr costs	\$4,823,568	\$4,689,980	\$1,292,632
IBM savings %	73.20%	72.44%	
IBM yearly savings	\$1,607,856	\$1,132,449	

redislabs

Table 1 Source: Redis Labs - Prices are based on list prices that were current as of January 2015

¹ IBM provides excellent whitepapers about CAPI and NOSQL at the bottom of this article: <http://www-304.ibm.com/webapp/set2/sas/f/capi/home.html>

² See <https://redislabs.com/>, more detail is available here: <https://en.wikipedia.org/wiki/Redis>

³ These are configuration comparisons; all costs are based on calculations using system list prices.

⁴ Note these comparisons are not benchmarks as no actual measurements were made on either configuration.

⁵ Note all ratios are based on counts of servers not server racks.



IBM

IBM developed and delivered CAPI to the OpenPOWER Foundation anticipating members would find additional applications for it. As evidenced in presentations at the recent OpenPOWER Foundation summit, IBM's hopes are bearing fruit. As an example, Nallatech's⁶ uses CAPI to attach FPGA (Field programmable gate arrays)⁷ to a POWER8 system. FPGA allow programming of microprocessors in the field to perform specialized functions for networks, app acceleration, etc.

IBM has also done calculation exercises to highlight potential advantages for Redis database using CAPI on POWER8. One calculation shows that managing a 12TB database requires slightly more than 24 x86 servers versus only a single POWER8 server. Thus, they can claim a 24 to 1 (server consolidation ratio) advantage for POWER8. While, technically true, we believe most customers would want a second server (as backup protection) in a production environment⁸.

The Redis Lab's database deployment comparison (discussed above) required only 6 POWER8 servers, a 13 to 1 reduction ratio, even though the database is larger (40TB) and includes processor redundancy. It is closer to a real-life production data center example. Even if we add a server for backup to IBM's example (cited above) the 24 to 1 ratio becomes 24 to 2 or 12 to 1. This is very close to the Redis Lab result.

To our knowledge, neither IBM nor Redis Labs has publicly discussed any comprehensive benchmark data comparing real configurations of x86 servers versus POWER8 CAPI servers. IBM has case studies in the works that focus on POWER8 Linux. We'll discuss some of these when they are available.

Redis Labs has a benchmark on their web site comparing their service with several of their competitors⁹. The results are very favorable to Redis Labs.

In summary, the server ratios discussed here are configuration comparisons demonstrating potential benefits of the POWER8 CAPI Flash solution. They are not based on actual benchmark or production data. They should not be generalized to other cases outside of this specialized database example.

Redis Labs has POC (proof-of-concept) studies underway with multiple clients. Their product should be fully deployed at a number of production sites later this year. We expect more detailed performance data from production environments will emerge soon after.

⁶ For Nallatech see <http://www.nallatech.com/solutions/openpower-capi-developer-kit-for-power-8/>

⁷ For more on FPGA's see https://en.wikipedia.org/wiki/Field-programmable_gate_array

⁸ IBM points out that for true high availability one would need multiple shards for each data element. Neither the x86 case nor the POWER8 case is configured for high availability, so they are comparable in this respect.

⁹ See <http://lp.redislabs.com/2015-Q1-Redis-Cloud-Service-Benchmark-Lander01.html>



Competitive Responses

We expect Intel competitors to respond to the Redis and IBM comparisons¹⁰. First, they will configure NoSQL systems with more RAM which **will** serve to significantly reduce the required number of servers. For example, it takes only 20 Intel servers to handle a 40 TB database, **if** it is configured with 2TB of RAM. However, the cost of 40TB of RAM will be far higher than 40TB of flash, leaving the Redis/IBM solution with a significant pricing advantage.

In the slightly longer term, Intel can reengineer their chip to provide CAPI-like capabilities. Depending on how this is achieved, it could potentially result in other advantages for Redis Labs. For example, if Intel creates a solution with a significant amount of flash under the covers of a server, a failure would remove that portion of the database from the online configuration. This would give Redis Labs/IBM a significant availability advantage since their database resides on an external device, and not likely to suffer similar isolation.

Benchmark results will eventually be available. From a strictly performance point of view, we believe an in-memory solution may have some advantage over a CAPI-flash solution. However, we believe the price advantage would weigh heavily in favor of Flash based systems.

Summary

There is no question that CAPI technology has demonstrated great potential. Its application to the Redis NoSQL database shows real promise. In addition, we believe that other applications are already being discovered and developed. CAPI is a very real product not a theoretical technology. Redis Labs has multiple POC studies underway. We predict IBM and Redis Labs will have a significant advantage when they report performance in production environments.

In the meantime, some customers will probably wait to see more production results before making significant investments. We don't think that will take long. The other wild card is the speed of the development of applications using terabyte-sized databases that require very high speed performance. The faster such apps and databases proliferate the more pressing will be the demand for systems that can manipulate and manage them. And, that is what makes the CAPI flash solution so attractive.

¹⁰ We assume the size of the market for large in-memory databases is sufficient to justify response. IBM estimates the market might be as large as \$12B over the next several years. We believe that the growth of IOT is a major market driver.



Publication Date: June 4, 2015

This document is subject to copyright. No part of this publication may be reproduced by any method whatsoever without the prior written consent of Ptak Associates LLC.

To obtain reprint rights contact associates@ptakassociates.com

All trademarks are the property of their respective owners.

While every care has been taken during the preparation of this document to ensure accurate information, the publishers cannot accept responsibility for any errors or omissions. Hyperlinks included in this paper were available at publication time.

About Ptak Associates LLC

Our analysts cover a breadth of areas to bring you a "complete picture" on technology trends across the industry. Whether it's Cloud, Mobile, Analytics, Big Data, DevOps, IT Operational Analytics, Workload Optimized systems or some other emerging trend, Ptak Associates analysts cover these trends with a unique perspective that is both deep and broad.

Our clients include both industry leaders and dynamic newcomers. We help IT organizations understand and prioritize their needs within the context of present and near-future IT trends, enabling them to use IT technology effectively in solving business problems. We help technology vendors refine their strategies, and provide them with both market insight and deliverables that communicate the business values of their products and services. We provide all clients with an understanding of how their competitors are playing in their market space, and deliver actionable recommendations.

www.ptakassociates.com