

## Abstract

Identify market applications that have a definitive value proposition running on top of RAMCloud on the NEC Micro Modular Server, RAMCloud-in-a-Box. The project used RAMCloud-in-a-Box infrastructure to demonstrate that it is the ideal foundation for deploying and operating RAMCloud application workloads requiring low latency inter-process communications. Delivery of validation testing on the application(s) that require scalable, low latency storage IO will assist in developing target markets for RAMCloud. We analyzed applications that require low latency inter-process communications, and identified which market application(s) can benefit from RAMCloud's Service Delivery optimization and application acceleration. To validate the performance characteristics that these types of application(s) require, we ported one of them to RAMCloud and measured performance at the application layer.

## ExMatEx Introduction

- Multi-scale materials
- Application driven
- Computer science focused
- ASCR Co-Design Center - \$4M/yr for 5 yrs



- Starting third year
- Work in many areas: molecular dynamics, proxies, programming models, DSLs, multi-scale algorithms, vendor interface, runtimes, software stacks, etc.
- More info at <http://exmatex.org>
- Code: <https://github.com/exmatex>

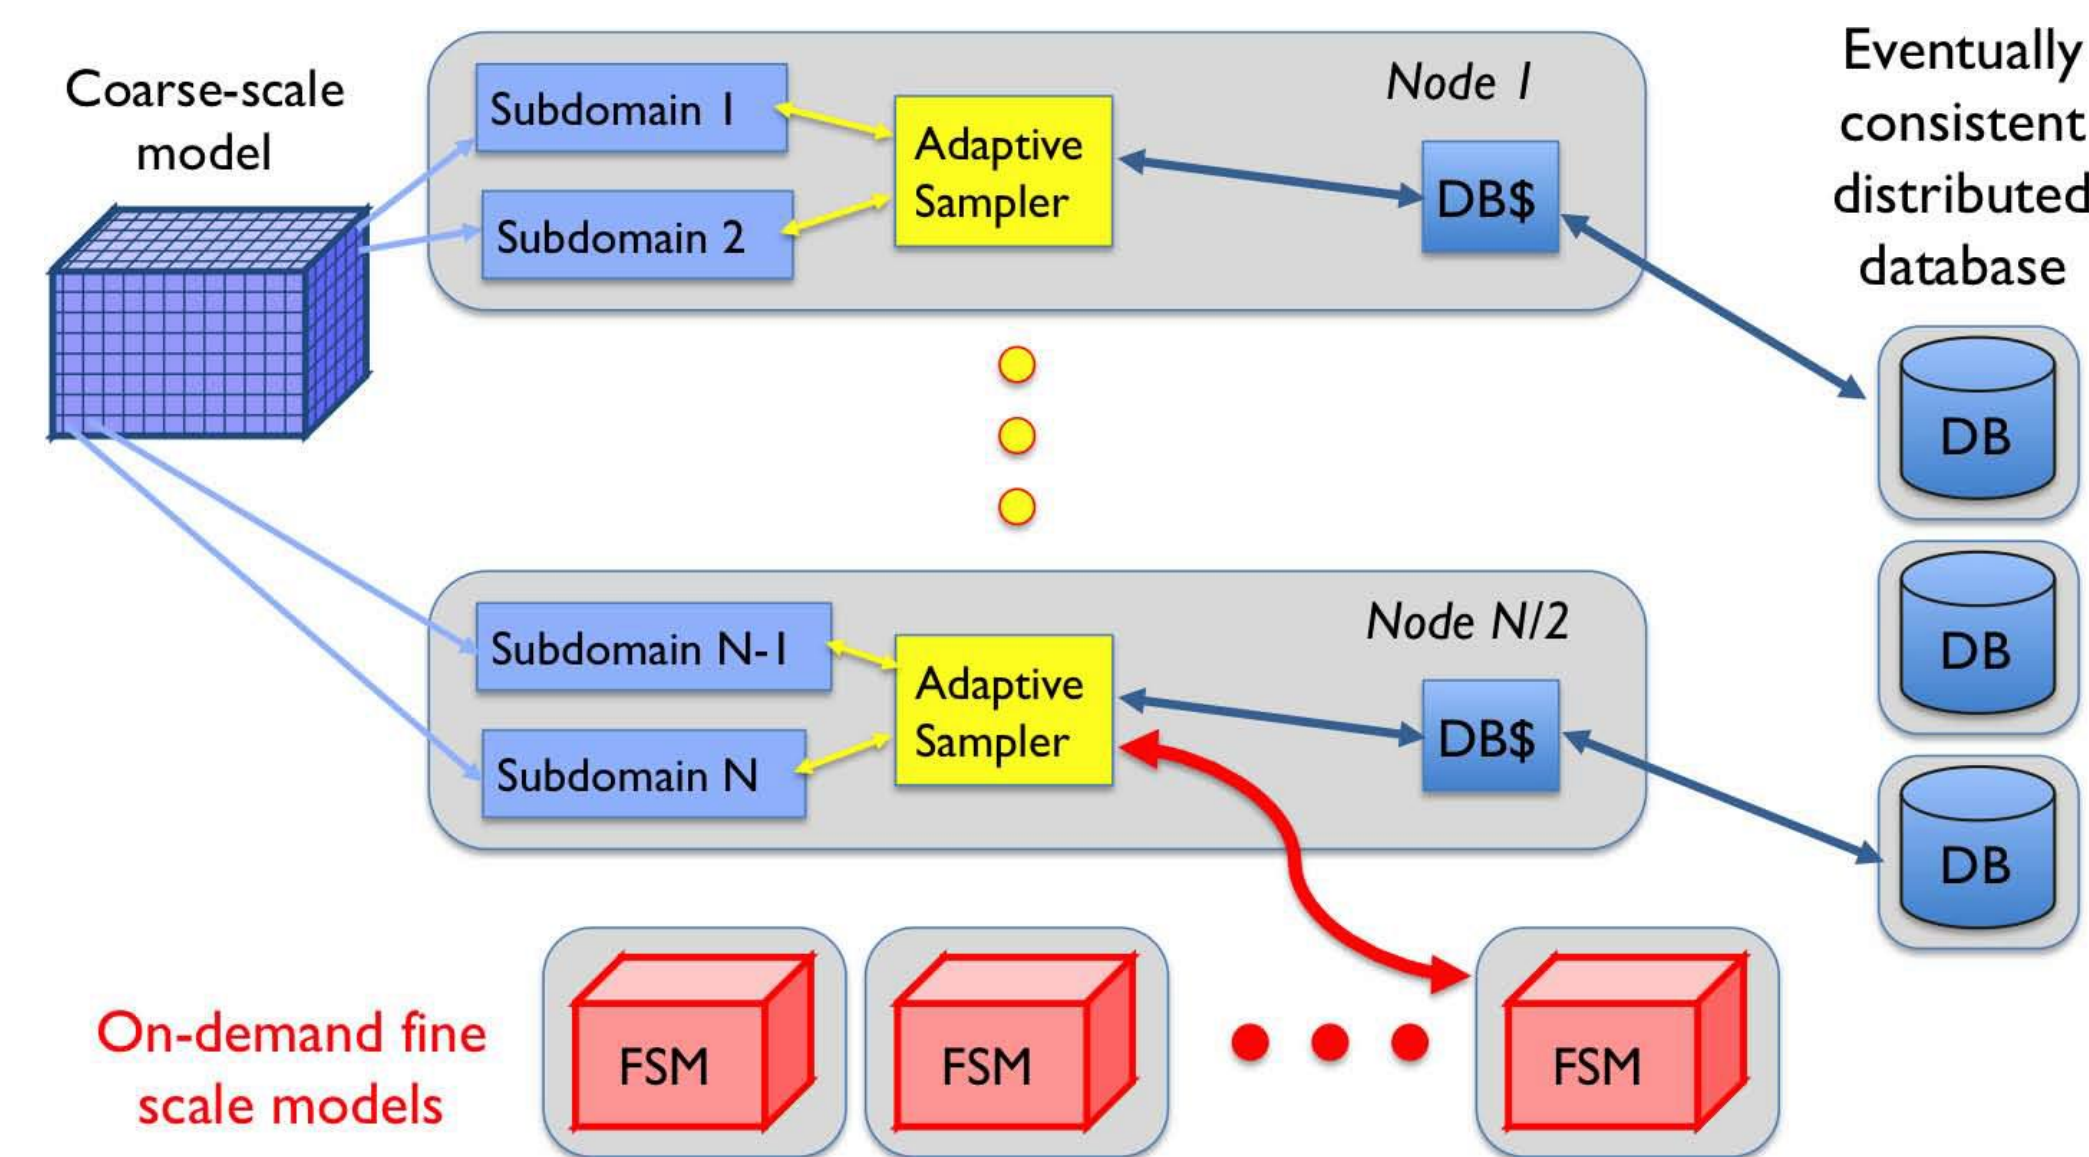


Figure 1. ExMatEx Application Architecture

## ExMatEx Application Orchestration

- 2 components today
  - Tasking
  - Databases
- Dynamic tasking
  - Hierarchical multi-scale
  - Independent computations
  - Uphill dependencies
  - Variable granularity
  - Application still logically synchronous at time-step boundary
- Databases (use largely driven by cloud/web influence)
  - Use to accelerate overall computation (sub-scale calls dominate)
  - Simple key/value stores
  - 60-200 doubles down, dozen doubles up
  - Obviously, queries must be much faster than fine scale calls

## Customer Database Requirements

- Really a "caching" service
  - Cache previous sub-scale calls
  - Typically in-memory
  - May want persistence too (subsequent runs of similar problem)
- Key/value API to start (additional functionality/APIs possible)
- Queries are tasks like any other task (same requirements)
- Transactional not required (i.e. eventual consistency OK)
  - Transactional writes may be required in the future
- Beyond caching sub-scale calls
  - Fault tolerance
  - EOS
  - Material properties
  - Many others

## ExMatEx on RAMCloud-in-a-Box

- Implemented Proxy service for Redis API calls used in ExMatEx
- Proxy service translates Redis API function calls into RAMCloud API calls
  - redisConnect = RAMCloud::RAMCloud::Write
  - redisCommand = RAMCloud::RAMCloud::Read
  - HMSET
  - HMGET
  - HDEL
  - SADD
  - SMEMBERS
  - DEL
  - FLUSHALL
  - FLUSHDB
  - ECHO
  - redisFree
  - redisAsyncConnect
  - freeReplyObject

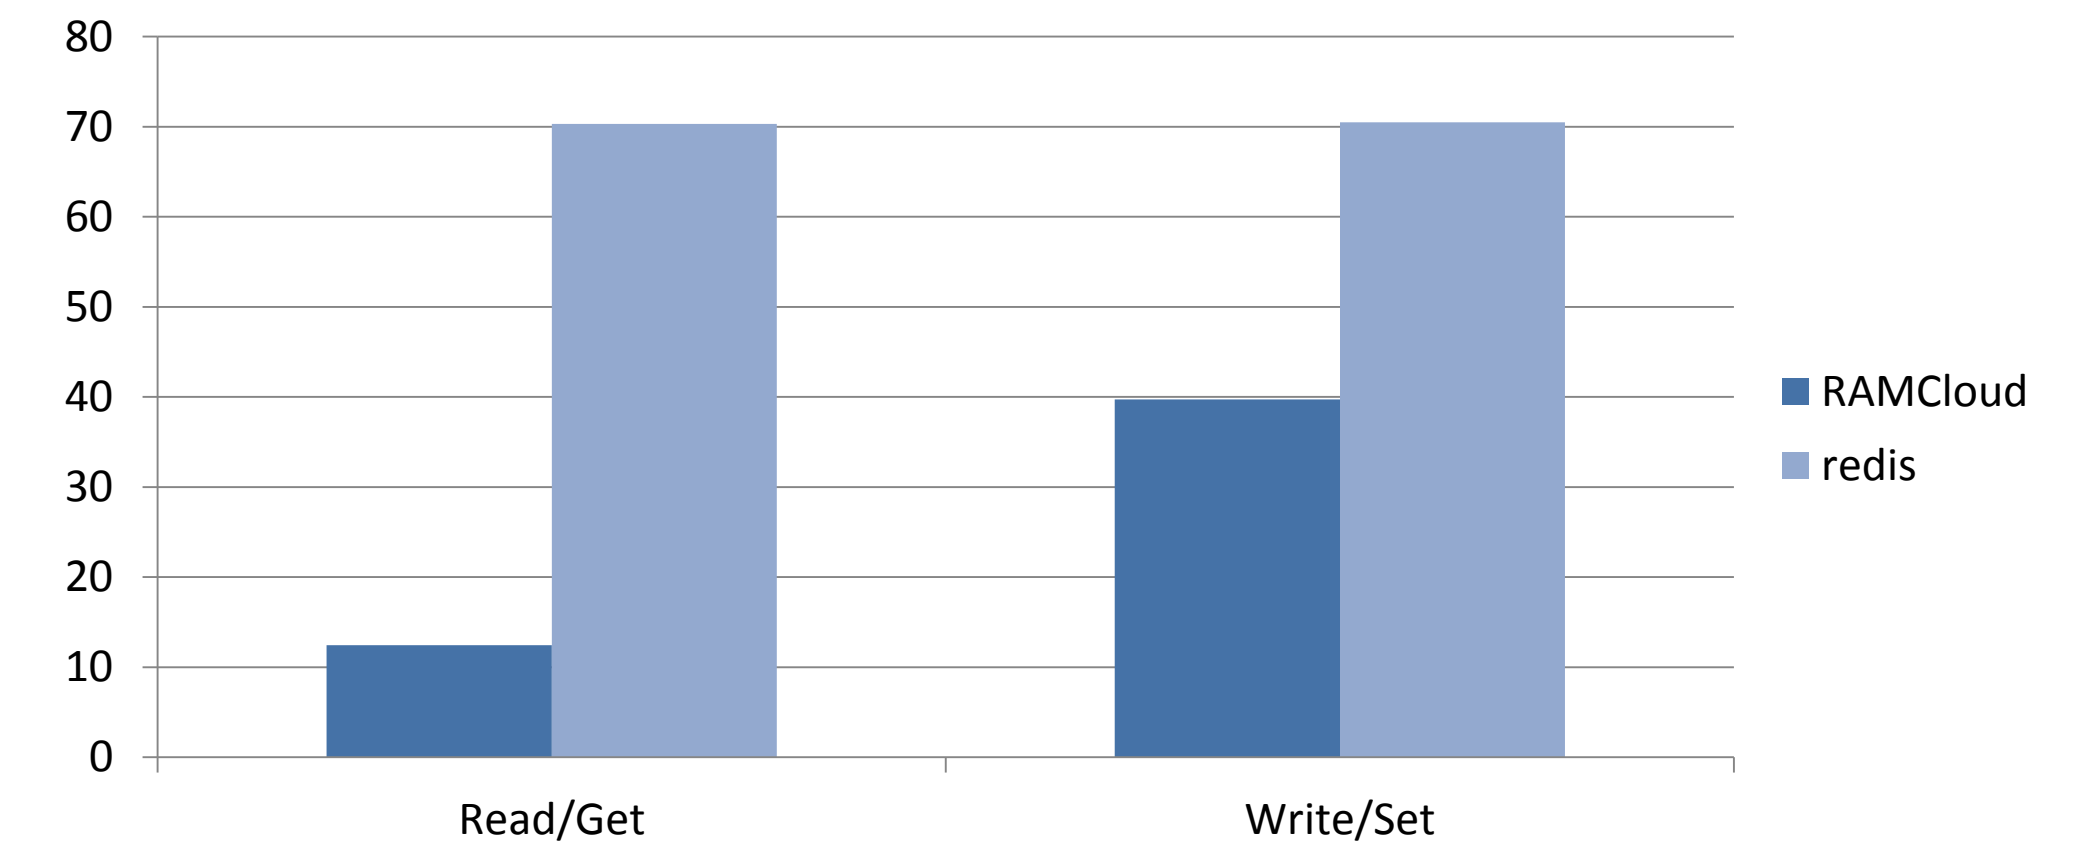


Chart 1. – RAMCloud vs. redis application Read/Write performance comparison

## BioTech and Scientific Market

- The global bioinformatics market reached the mark of around US\$ 3.8 Billion in 2013 and it is anticipated to grow at a CAGR of around 19.3% during 2015-2019<sup>1</sup>
- Both service providers, KIO Networks and BroadCloud, in the project expressed a need in this market based on their current customer bases
- Crosses into the High Performance Computing market with Molecular Modeling applications
  - Flexibility in protocol choices such as Infiniband and TCP/IP make RAMCloud-in-a-Box attractive in this market as NoSQL competition does not support Infiniband
- An end user customer was evaluating NoSQL key value stores for genetic sequencing application
  - Expressed ease of deployment and an appliance approach key for them
  - Scalability and ease of expansion are crucial
  - Require 1-2 second recovery, within RAMCloud-in-a-Box parameters
  - Extremely fast data access with low latency
  - C API interface required
  - Open source with support must be in place

## Conclusions

- RAMCloud-in-a-Box is a simple, fast NoSQL storage appliance that eases setup and integration
  - 100 Byte Reads in 12.453µs, 6x faster than redis
  - 100 Byte Writes in 39.735µs, 2x faster than redis
  - 1K Byte Reads in 19.2µs, and Writes in 60.450µs
- Tightly integrated and optimized system saves time and money in setup and deployment
- Extremely fast data access performance with low latency
- Fast resilient data writes are 2x faster than redis that is non-resilient
- Perpetual data availability for systems as a result
- RAMCloud-in-a-Box is a massively scalable solution vs. redis as the current NoSQL key value market leader
- RAMCloud-in-a-Box uses 75% less rack space, and 25% of the energy requirements of other systems
  - The scale of 680 RAMCloud-in-a-Box systems would be required to meet the customer need for faster key value storage

## Contact

Robert Brown  
Software CPA, LLC  
Email: [bob@software-cpa.com](mailto:bob@software-cpa.com)  
Website: [www.software-cpa.com](http://www.software-cpa.com)  
Phone: 760-207-3762

## References

1. RNCOS - Global Bioinformatics Market Outlook 2019.