



Big Data on Arm in Practice with ThunderX

Jeremy Hallum, Research Computing Manager,
Advanced Research Computing - Technology
Services, University of Michigan



**Linaro
connect**
San Diego 2019

Introduction

- Who are we?
- “The Ivy of the Midwest”
- 1.55 billion USD in annual research expenditures.
- Broad portfolio of academic research, with a university and a full health system.
- 19 schools and colleges united by a football team.



Who is ARC-TS

- ARC-TS provides comprehensive computational and storage tools for U-M researchers.
- HPC, storage, “big data”, enclaves, and networking.
- 14,000 cores of Intel HPC
- 5 PB of high speed storage
- 10 PB of tape archive.
- One goal: Make data easy to move and analyze, at whatever scale.



**ADVANCED
RESEARCH COMPUTING
TECHNOLOGY SERVICES**

UNIVERSITY OF MICHIGAN

Why Arm for Data Science?

- There are a number of reasons the ThunderX platform looks interesting for data science
 - High core count for many independent tasks
 - Rapid data flow between CPU and networking
 - Each node is 40 Gbit native.
 - Gigabyte's ThunderX motherboard/chassis is capable of managing enough high capacity drives per node for HDFS .
- Diversity in architectures drives innovation in research.
- Understand new methodologies for research to stay competitive.

Marvell collaboration

- 2016, there was mutual interest to explore ThunderX for Data Science in an applied setting.
- Acquired 40 data nodes from Gigabyte
 - ~4000 ARMv8 cores.
 - 25 TB of RAM.
 - 3 PB of HDFS
- First goal: Build an analytics environment using Hadoop for research use.
 - Tools: Spark, Hive, pyspark, SparkR, parquet.



In the beginning...

- Start from x86_64.
- “It’s java, so it should be easy, right?”
- We built a working cluster, but we had to rebuild everything along the way
- Complex stack of software dependent on each other to work properly
- Two biggest pain points:
 - Bytecode dependencies on x86 in libhadoop.so.
 - Protobuf
- The first is kind of painful, but can be worked around.
- The second was and still is painful without heavy patching.



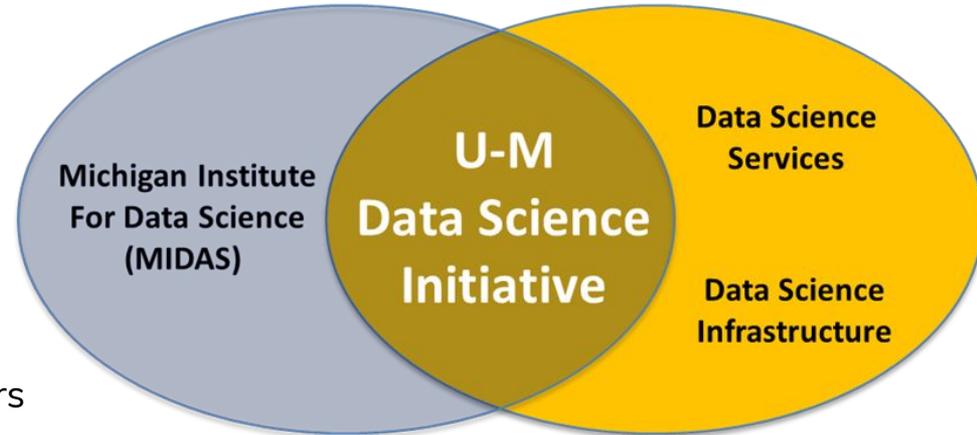
Thanks to Linaro...

- We didn't know about Linaro at first until we heard a reference on the CentOS list.
- We repackaged RPMs to solve many of our initial problems.
- We reached out to Linaro and opened up a dialog about collaborating.
- Switched from our own in-house build process to an in-house/BigTop hybrid.
 - Manage the cluster with Ansible, using the BigTop packages.



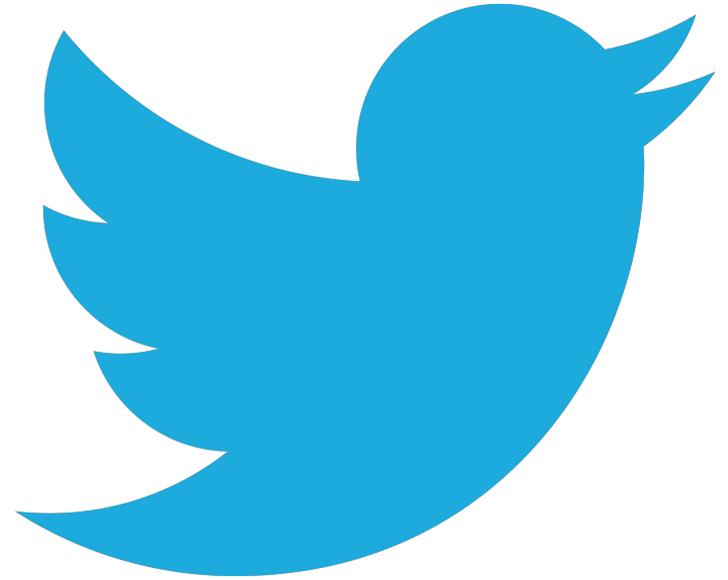
Current Use

- A dozen research groups, and a half dozen courses use the cluster for teaching and analysis.
- Hosts the campus Twitter repository, and the campus Reddit repository.
- Initial use cases: Social Sciences.
 - Contextual post/tweet analysis
 - 'Fake news' detection
 - Understanding the effects of disasters on communities through social networks.



Campus Twitter Data Repository

- <https://bit.ly/2kL2EOC>
- 750+TB of Twitter dechase data. (10% of all data from twitter).
- 20 GB of tweets per day into the system
- Available to all researchers on campus upon request and justification.



Lessons Learned: BigTop

- BigTop is straightforward to use and the benefits are immense. It simplified things for us greatly.
- New Red Hat/CentOS packages make the process simple.
- Fully reproducible Hadoop cluster, rebuildable in 2 days.



Lessons Learned: Software Partnerships



- Conda is a pre-packaged open-source python distribution.
- Consumable at multiple scales, from laptop to HPC cluster.
- Our researchers are used to the Conda python distribution for their work, inclusion of an ARM64 Conda python distribution would be a huge plus for the ecosystem.

Lessons Learned: Distributions

- Early packaging focused on Debian/Ubuntu/Fedora builds.
- We are primarily a CentOS shop.
- We spent a lot of energy early translating Fedora packaging to Red Hat/CentOS.
- Recent transitions to include Red Hat/Centos Packaging have been helpful for support for us.



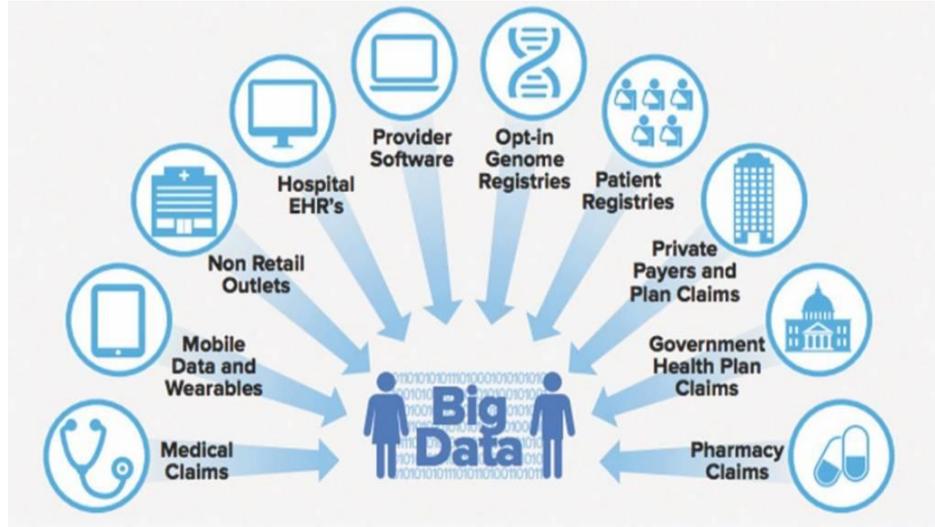
Lessons Learned: Software Availability



- Newer versions of spark/pyspark in the ecosystem.
- There has been talk of Spark/pyspark not developing against arm64 for future versions.
- How do we develop consensus and encourage developers to build for Arm64?

Future Goals

- Social Science data is a fraction of available data on campus.
- The bulk of data is restricted in some fashion
 - Health system data
 - Claims data (Private Insurance/Public Benefits)
- Data Analytics on restricted data sets.
 - 'Precision Health' Data Analytics
 - Opioid crisis root causes
 - Population studies across multiple types of claims data for success/failure rates..
- Additional tools to ensure the safety of data on the system.
- Genomic analysis via Spark.



Future goals

- The data analytics climate is rapidly changing.
- Yarn/Hadoop as a platform is becoming niche.
- Machine Learning(ML) and GPU data analytics are taking the data world by storm.
- Look into alternate methodologies and apply to ARM64 at scale.
- Software and Hardware.



Battling a changing software ecosystem

- Spark is only the tip of an iceberg.
- The ML ecosystem is dynamic and still changing rapidly.
- Worse, because of the dynamism, much of the software can be considered 'beta' and doubly difficult to build/debug and support on ARM64 in production.

Conclusion

- In partnership with Marvell, we have built a Hadoop cluster at scale for researcher use.
- It is currently serving Umich researchers.
- We look for opportunities to grow into the ML analytics space.
- Thanks: to the Data Science Team at UMich: Seth Meyer, Marc Patton, and Matt McLean, who all have worked on this project.

Thank You - Contact

Questions: Jeremy Hallum <jhallum@umich.edu>

ARC-TS: <http://arc-ts.umich.edu/>

Twitter: @ARCTS_UM

Thank you

Join Linaro to accelerate deployment of your
Arm-based solutions through collaboration

contactus@linaro.org



9Boards is a range of specifications with boards and peripherals offering different performance levels and features in a standard footprint.



Linaro
connect
San Diego 2019