

Research Computing at a Business University

Jason Wells
jwells@bentley.edu
Bentley University
Academic Technology Center
Waltham, MA, USA

J. Eric Coulter
jecoulte@iu.edu
Indiana University
Science Gateways Research Center
Bloomington, IN, USA

ABSTRACT

Research Computing demands are exploding beyond traditional disciplines due to the proliferation of data in all walks of life. At Bentley University ("Bentley"), a business university in the Boston area, this expansion has been most readily seen in our Accounting, Economics, Mathematics, and Natural Sciences departments. The result has been a small effort to build a research computing capability at this small New England university. This poster will serve as an overview of the steps taken to build such an effort at a business university, the revelations we have had along the way, and our plans for the future.

CCS CONCEPTS

• **Human-centered computing - field studies**; • **Social and professional topics** → *Systems planning; Hardware selection; System management*;

KEYWORDS

HPC, Long Tail of Science, Non-traditional disciplines, Single Threaded, Windows

ACM Reference Format:

Jason Wells and J. Eric Coulter. 2019. Research Computing at a Business University. In *Practice and Experience in Advanced Research Computing (PEARC '19)*, July 28-August 1, 2019, Chicago, IL, USA. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3332186.3333161>

1 INTRODUCTION

Bentley is a small business university located 30 minutes from Boston in Waltham, Massachusetts. It has 449 full and part-time faculty, around 5,000 graduate and undergraduate students, and 24 PhD candidates. Most of the university's majors are in business disciplines, or focus on business disciplines from an external perspective (e.g. Health Economics). As a result, we can say that our Research Computing effort targets non-traditional High Performance Computing (HPC) disciplines, and due to a lack of NSF funding, is solidly within the Long Tail of Science[1].

This poster discusses two main questions: whether Bentley researchers have a need for high performance computing, and whether

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

PEARC '19, July 28-August 1, 2019, Chicago, IL, USA

© 2019 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-7227-5/19/07...\$15.00

<https://doi.org/10.1145/3332186.3333161>

or not its possible for a school such as Bentley to provide for that need. We will approach the first question through the introduction of a new metric, and the second from the mindset of providing local hardware resources in a centralized manner.

2 MOTIVATION

Bentley's researchers overwhelmingly conduct their research in three applications (in order): SAS, Stata, and R. The primary issue encountered is that SAS and Stata are limited to a single node ("node bound") and are often running functions or libraries that are limited to a single core ("core bound"). R technically has the same limitations but can overcome them through certain methods.

It is these limitations that perhaps explain the lack of 'technical efficiency' found with Economics (Bentley's main HPC consumers) by Apon et al. [2], and these concerns are not unknown. The Kansas City Federal Reserve Bank's Center for the Advancement of Data and Research in Economics (CADRE) has sought to meet them with nodes equipped with higher speed processors, increased memory amounts, and interactive capabilities [3]. Regardless, it is these limitations that drive the research questions above.

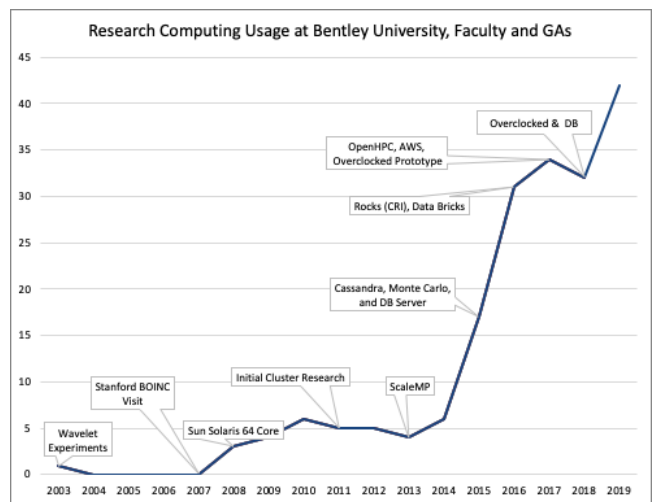


Figure 1: To illustrate our growth, our number of Faculty and GA/RAs over the years is charted, along with some of the major milestones.

3 THE NEED

The data sets our faculty use to construct models and simulations have steadily been increasing in size, causing longer computation




Figure 2: This chart compares performance on the Single Threaded Passmark benchmark between various Bentley resources and a variety of supercomputers available to researchers through XSEDE, as well as a few Top500 Intel based supercomputers. For researchers in many disciplines at Bentley, the high per-thread performance is more valuable than aggregate performance of thousands of cores.

times that have pushed the number of possible execution times per year down. These execution times can be considered "Research Opportunities Per Year" (ROPY), a useful metric when discussing research-oriented return on investments (ROI) for computational infrastructure. The metric is calculated by the number of units in a year, divided by the programs' execution time in the same units, as so:

$$ROPY = \text{units}_{year} / \text{execution time}_{unit}$$

To show how the metric is calculated assume one faculty member's (Faculty A) program takes four months to run. Because faculty have access to our HPC resources at all hours of the day, regardless of the day, we can calculate their ROPY as 12 months per year, divided by the four months their program takes to run:

$$ROPY_A = 12_{months} / 4_{months\ to\ run\ program}$$

Or

$$ROPY_A = 3$$

Other researchers have much shorter program execution times, such as two hours (Faculty B), giving these researchers ROPYs in the thousands:

$$ROPY_B = 8760_{hours\ in\ a\ year} / 2_{hours} = 4380$$

Our efforts have focused on helping both groups increase their ROPYs. To date, the ROPY 3 researcher is now a ROPY 26, and several researchers with ROPYs in the thousands have reached the tens of thousands.

Given the large increase in the number of researchers at Bentley, as shown in Figure 1, and our success in helping them using the ROPY metric, we have concluded that there is a clear need for HPC resources at Bentley.

Table 1: The different generations of Bentley's HPC, Data Science, and Storage efforts.

Generation	HPC / Batch	Data Science	Storage
First	ScaleMP	Cassandra Docker	HDFS
Second	Rocks	Data Bricks	
Third	OpenHPC Windows		

4 EFFORTS TO DATE

For our first generation of research computing (see Table 1), we began by convincing our Systems group to give us its old virtualization Dell M600 hardware in 2013, which was used for HPC and Data Science purposes. For HPC, we made a three node ScaleMP cluster and for Data Science a 16 node Cassandra/Docker cluster. The ScaleMP cluster was difficult for our faculty to use however, so for our second HPC generation in 2016, we opted for a 16 node Rocks cluster. It was at this time that XSEDE's Campus Resource Integration (CRI)[4, 5] group was brought in to help roll out the Rocks cluster, a feat we likely would not have been able to do alone due to the lack of official support channels. Around this time, we also noticed that the Cassandra/Docker cluster was causing too much focus on teaching students how to use it rather than focusing on the data science they had intended. That resulted in our second Data Science generation, Data Bricks[6], a Software as a Service (SaaS) based supplier of Apache Spark, Python, and R-Studio notebooks.

Security concerns forced the retirement of the old Dell M600 hardware in 2018, so the remaining on-campus HPC efforts moved into a third generation with Dell R430s, overclocked servers from the Quantitative Finance world, and one monster 4U server with 5 NVIDIA Tesla Graphical Processing Unit (GPU) cards scrounged from around campus. CRI came to the rescue again and helped us build an OpenHPC cluster with this hardware. Since then, we

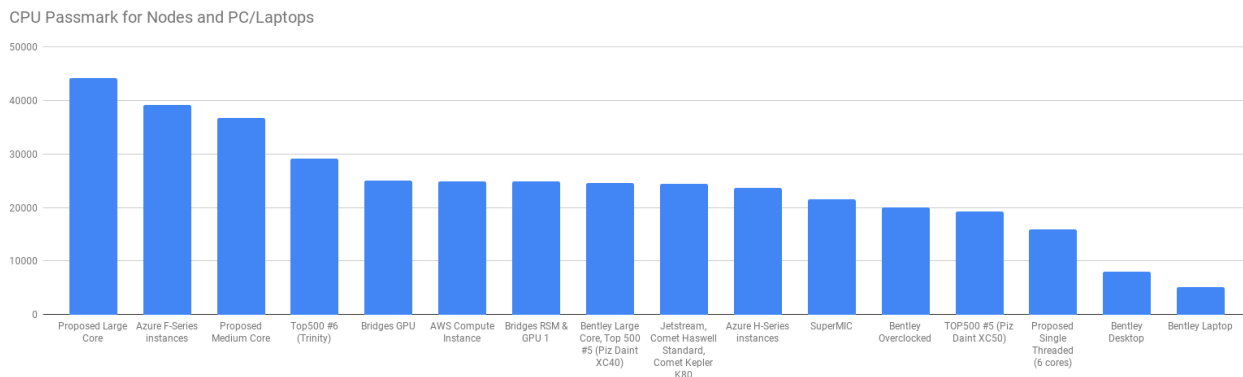


Figure 3: This chart compares performance on the CPU Passmark benchmark between various Bentley resources and a variety of supercomputers available to researchers through XSEDE, as well as a few Top500 Intel based supercomputers. For researchers in many disciplines at Bentley, the high per-node performance is more valuable than aggregate performance of thousands of cores.

have added three of the overclocked servers running Windows and delivered to faculty by Ericom, a Citrix competitor. Additionally, we are adding 36-core count, and database servers to the OpenHPC and Windows environments.

First Generation storage efforts started in 2018 with a Hadoop File System (HDFS) instance, but died due to poor performance. This issue is discussed in more detail later in this abstract.

5 ADDRESSING THE DEMAND

Past efforts have largely focused on learning about the advantages and disadvantages of HPC. We have concluded that while clusters make management easy, a greater emphasis is needed on per node and per core performance, as well as on how user friendly our resources are. We address these requirements by focusing on the different elements of a HPC system: Compute, Storage, Networking, and Operating System. A modern concern is whether cloud computing can handle each need, so we briefly address it in each section as well.

5.1 Compute

5.1.1 Methodology. We utilize a large database of benchmarking figures using the popular Passmark software to compare processors for two benchmarks: 1) the Single Threaded Mark which represents the core bound functions and libraries mentioned above, and 2) a CPU Mark consisting of many CPU tests beyond the Single Threaded Mark, which represents the node bound applications.

5.1.2 Results. Bentley’s overclocked nodes are aimed at the per core concern, and utilize a single 10-core Intel i7-6950x processor running at 4.394GHz with a single threaded Passmark score of 2,148 which was amongst the top 10 single threaded Passmark scores in 2017 when purchased. Figure 2 shows how this score compares to other supercomputers and even Bentley’s standard laptops and desktops. With Bentley desktops being faster on this benchmark than all of the supercomputers it made sense to embrace the overclocked nodes. Recently, however, new CPUs have been released

which could negate the need for overclocking. For instance, Intel Xeon E-2186G based nodes (denoted on Figure 2 as “Proposed Single Threaded”) provide a dramatic increase over even the overclocked nodes.

Another factor to consider is the node based performance. CPU Mark provides the best basis for this comparison. Figure 3 shows how our two 18 core Intel E5-2695 v4 node (“Bentley Large Core” on Figures 2 and 3) compare with the other supercomputer nodes. Our researchers are more likely to favor our servers than apply for an XSEDE allocation, so while our Bentley Large Core node is in the middle of the pack, it suffices for our node bound needs. Here too we see an opportunity to increase our speeds by offering a few dual Intel Xeon Gold 6148 nodes (“Proposed Medium Core” on Figures 2 and 3), and at least one node with 28 core Intel Xeon Platinum 8180 in quad or eight way configuration (“Proposed Large Core” on Figures 2 and 3). This server would increase the ROPY of our ROPY 26 client the most, but it will come at a price of USD90K.

We specifically note Azure F-Series, H-Series, and Amazon Web Services (AWS) Compute Instances CPUs on these charts to show how these services aim their HPC solutions. Bentley’s single threaded focus is nearly absent on Figure 2 (our desktops are faster), and while not covered in this poster, the cost for many cores in the highest of CPU Mark cloud instances is very high. Unfortunately those nodes would actually be used for our longest running codes.

5.2 Storage

Our first generation storage efforts began with a HDFS instance. It performed poorly chiefly because this parallel virtual file systems speed up access by storing the same file on three drives. We only saw a 3x speedup of the base drive speed, which was paltry for traditional hard drives. We tried to fix this using three RAID sets but ran into another problem. While RAID sets can be used to increase read speeds, the SAS channels they make use of have maximum speeds. As a result, even with 14 SSDs spread across two SAS Raid channels, the fastest read speeds we ever saw was 2.8GBps (1.4GBps per SAS channel), even though the SSDs are each

capable of 450MBps individually. With 14 SSDs we should have seen a 13x speed up (5.850GBps). The only traditional solution, as a result, would be many hard drives, in Raid sets, in many servers, a solution a small school like Bentley would not be able to employ.

But there is a solution. While Solid State Drives are faster than Hard Drives, Bentley desktops and laptops have Non-Volatile Memory Express (NVMe) drives, which are an order of magnitude faster than even SSDs (Samsung PM961 on Figure 4). This leads to a strange situation where the read speed of the NVMe drives (2800 MBps/2.8GBps/22.4Gbps) in our laptops and desktops is faster than our supercomputer’s network fabric (10Gbps). To truly offer a faster storage experience than our client’s desktops and laptops, we are considering rolling out converged NVMe drives in a parallel virtual file system (probably BeeGFS), at the same HDFS 3x replication, to yield about 8.4GBps. We can offer this solution using just one server with eight PCIe interfaces, and eight M.2 to PCIe adapters. Please see next section for network implications.

As far as cloud services, we have not yet heard of high speed storage efforts to the degree that we are considering with this service, and past experience with storage costs (10TB is about USD1K/month with AWS EBS and compute for us) give us pause when considering moving to the cloud.

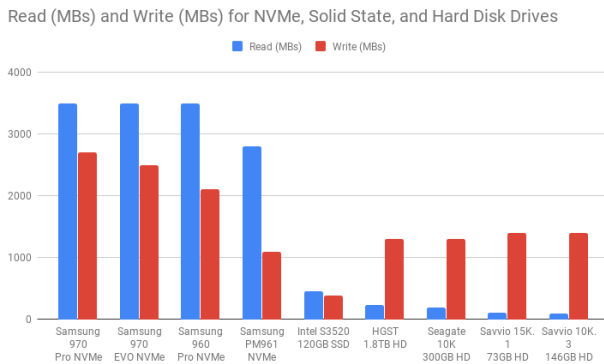


Figure 4: This chart compares read/write performance across a NVMe, SSD, and standard HDD drives. Most large scale resources focus more on sheer size of storage rather than speed, though there are a handful of exceptions, such as PSC-Bridges AI-GPU nodes. For clients with moderate data needs and high I/O requirements, NVMe drives offer a huge performance boost.

5.3 Networking

As a result of the storage speed issue, our 10Gbps network simply will not suffice for future needs. To successfully carry the traffic for NVMe drives (at 3x replication) and metadata, we would need to be running 100Gbps networking speeds to each compute node, which is available in Ethernet or newer Infiniband protocols. Higher replication rates will be possible but the network will limit the usefulness of such rates. Cloud efforts occasionally discuss 25-40Gbps options for Read Direct Memory Applications (RDMA) and networking, but we have not see such speeds for storage connections.

The decision on Ethernet vs. Infiniband is a difficult one. We do not yet have researchers using Message Passing Interface (MPI) libraries or RDMA, so latency has not really been a concern for us. In the end we will likely choose the least expensive total cost for ownership option that gives us the speeds required while also considering future requirements.

5.4 Operating Systems

Bentley researchers mirror their corporate counterparts and utilized Windows nearly entirely until a decade ago, when Apple computers began to creep onto campus. As can be see in Figure 5, 84.3% of our full-time faculty use Windows now, with the remainder on Macs. Linux nearly does not exist. This limited the number of faculty willing to work with our non-Windows HPC solutions. For this reason we began offering Windows nodes in 2018 and have seen a dramatic uptake since then (see Figure 1). Most importantly, clients began finding our resources through the Ericom software on their own and simply started using it spontaneously.

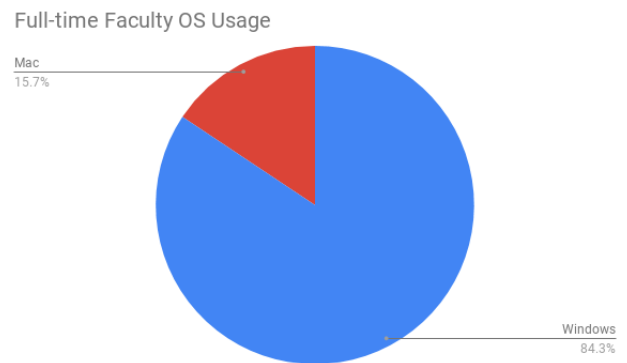


Figure 5: This pie chart documents the number of full-time faculty who currently using Windows, Macs, and Linux as their primary computer Operating System. Please note that Linux is absent from the graph because the number is 0.

Aside from making introductory training interesting, we also have to deploy hardware solutions in both Linux and Windows as a result. The overclocked servers, Large Core servers, and even GPU resources are offered in both operating systems.

6 CONCLUSION

While many of the needs encountered at Bentley are common to non-traditional HPC disciplines, they are not widely articulated by large research computing centers. Researchers in these areas are often left to compute on desktops, pointed towards large systems that are tailored to other needs, or pushed out to the cloud, which, as we have shown, can be superseded in speed by well-considered local solutions. Additionally, moving research onto cloud providers comes with many headaches in the form of billing, security, compliance, systems administration, and software/data migration concerns. <https://www.overleaf.com/4412273942nqvbnjwsrqhg> The biggest gains have been in the realm of compute power. However, upgrading the storage and networking available to researchers will be

invaluable in providing solutions that are truly more powerful than a researcher's laptop. It is also clear from our experience that many researchers in non-traditional disciplines are best served by a variety of hardware and operating system options. While progress in clock speeds has plateaued in recent years, providing access to high Passmark processors of both types is still a viable model for researcher support in many areas which do not require or strongly benefit from many tightly coupled processors. Additionally, this model has proven simpler to implement and support than a "traditional" cluster system, and resulted in happier researchers, faster adoption by researchers, and greatly increased ROPs. Other institutions in the area have also expressed interest in this model, as they too are discovering that not all faculty are well supported by cluster computing. In organizations where the primary goal is supporting research, such systems may be worth investigating in order to broaden the reach of research IT, and to truly bring the computing revolution to all participants in the research game.

ACKNOWLEDGMENTS

This work was partially funded by the Extreme Science and Engineering Discovery Environment (XSEDE), which is supported by National Science Foundation grant number ACI-1548562

REFERENCES

- [1] Patrick Heidorn. Shedding light on the dark data in the long tail of science. *Library Trends*, 57:280–299, 09 2008.
- [2] Amy W. Apon, Linh B. Ngo, Michael E. Payne, and Paul W. Wilson. Assessing the effect of high performance computing capabilities on academic research output. *Empirical Economics*, 48(1):283–312, Feb 2015.
- [3] BJ Lougee, Tim Morley, and Mark Watson. The road to cyberinfrastructure at the federal reserve bank of kansas city. *CADRE Technical Briefings*, Apr 2018.
- [4] Craig A. Stewart, Richard Knepper, James Ferguson, Felix Bachmann, Ian Foster, Andrew Grimshaw, Victor Hazlewood, and David Lifka. What is campus bridging and what is xsede doing about it? In *Proceedings of the 1st Conference of the Extreme Science and Engineering Discovery Environment: Bridging from the eXtreme to the Campus and Beyond*, XSEDE '12, pages 47:1–47:8, New York, NY, USA, 2012. ACM.
- [5] Eric Coulter, Jeremy Fischer, Barbara Hallock, Richard Knepper, and Craig Stewart. Implementation of simple xsede-like clusters: Science enabled and lessons learned. In *Proceedings of the XSEDE16 Conference on Diversity, Big Data, and Science at Scale*, XSEDE16, pages 10:1–10:8, New York, NY, USA, 2016. ACM.
- [6] Databricks - unified analytics. <https://databricks.com/>, 2019. [Online; accessed 21-May-2019; Databricks].