Sticking (with) the Landing

A modern case for Knights Landing in Resource-Constrained Environments

Bryan Johnston
CHPC
CSIR
Cape Town, South Africa
bjohnston@csir.co.za

Charles Crosby
CSIR
Cape Town, South Africa
ccrosby@csir.co.za

Quinn G. Reynolds
Stellenbosch University
Stellenbosch, South Africa
Mintek
Johannesburg, South Africa
quinnr@mintek.co.za

Jennifer M. Schopf
Texas Advanced Computing Center
(TACC), University of Texas at Austin
Austin, USA
jmschopf@tacc.utexas.edu

Christopher J Whalen
Research Data and Communication
Technologies
Garrett Park, USA
cwhalen@researchdata.us

Abstract

The HPC Ecosystems Project has repurposed decommissioned tier-1 HPC systems into entry-level clusters across Africa for over a decade. Stampede2 Knights Landing (KNL) systems are available for global distribution through the Texas Advanced Computing Center's Legacy Computing Program, and to HPC Ecosystems Project partners. To ensure the novel KNL architecture is fit for purpose for sites contemplating adopting the legacy systems, this publication provides a brief performance reference guide for prospective adopters. Benchmark tests were conducted to evaluate Stampede2's KNL processors on modern workloads to help inform prospective adoption decisions. It was concluded that the Stampede2 KNL processors remain particularly suitable for applications that benefit from good memory bandwidth, but that multi-node use is only feasible if high-performance networking is also available.

CCS Concepts

• **Hardware** → **Hardware test**; • **Applied computing** → **Physical sciences and engineering**; • **Social and professional topics**;

Keywords

HPC Ecosystems Project, Legacy Computing, Resource-Constrained Environments, Xeon Phi Knights Landing, Benchmarks

ACM Reference Format:
Bryan Johnston, Charles Crosby, Quinn G. Reynolds, Jennifer M. Schopf, and Christopher J Whalen. 2025. Sticking (with) the Landing: A modern case for Knights Landing in Resource-Constrained Environments. In *Practice and Experience in Advanced Research Computing (PEARC '25), July 20–24, 2025, Columbus, OH, USA*. ACM, New York, NY, USA, 5 pages. https://doi.org/10.1145/3708035.3736075

This work is licensed under a Creative Commons Attribution 4.0 International License.
PEARC '25, Columbus, OH, USA
© 2025 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-1398-9/25/07
https://doi.org/10.1145/3708035.3736075

1 Introduction

Over 35 entry-level HPC sites have been established through the HPC Ecosystems Project's hardware distribution pipeline, leveraging decommissioned tier-1 systems [5, 6], including the Texas Advanced Computing Center (TACC)'s past flagship system, Stampede ("Stampede1")[1]. To enhance support benefits from a common community of practice, the HPC Ecosystems Project partner sites adopt a community OpenHPC software stack [9]. Of the deployed systems in the community, nearly half currently utilise Stampede1 resources.

Launched in 2017 as TACC's flagship system, Stampede2 delivered 18 PFLOPS peak performance and debuted 12th on the TOP500 list[2]. By its 2024 decommissioning, it ranked 110th and featured over 4,200 Intel Xeon Phi 7250 (KNL) nodes and 1,736 Intel Xeon Skylake nodes. Each KNL node included 68 Silvermont-based cores with 4-way SMT, AVX-512, and 16GB of on-package MCDRAM, offering over 3 TFLOPS of double-precision performance—well-suited for memory-bound, parallel workloads[3].

While Stampede2 used Intel Omni-Path Architecture (OPA) interconnect, the OPA interconnect was excluded from benchmarks to ensure a comparative baseline across reference systems that did not have OPA. Stampede2 KNL systems have 96 GiB of DDR4 memory, whereas the KNL reference system used in the benchmarks features 192 GiB. However, for the tested benchmarks, the HBM is expected to be the key factor, with the additional 96GiB having minimal performance impact.

TACC's Legacy Computing Program[4], together with their ongoing partnership with the HPC Ecosystems Project, presents an opportunity for research groups in scientific and high-performance computing to adopt Stampede2 hardware. At the time of writing,

---

[1] Stampede1 (Top500.org) – https://www.top500.org/system/177931/
[2] Stampede2 (Top500.org) – https://www.top500.org/system/179045/
[3] Intel Corporation. Intel® Xeon Phi™ Processor 7250 Specifications – https://www.intel.com/content/www/us/en/ark.html
[4] Texas Advanced Computing Center. TACC Legacy Computing Program – https://tacc.utexas.edu/partnerships/legacy-computing-program/

Table 1: Benchmark System Configurations (derived from `lshw`)

Technical Component	Stampede1 (Ref. #1)	Lengau (Ref. #2)	KNL Benchmark System
Operating System	Rocky Linux 8.10	CentOS 7.3	Rocky Linux 8.10
Intel CPU	Xeon E5-2680 @ 2.70GHz	Xeon E5-2690v3 @ 2.60GHz	Xeon Phi 7250 @ 1.40GHz
TDP	130W	135W	215W
Sockets/Cores/Threads	2s / 16c / 32t	2s / 24c / 24t	1s / 68c / 272t
RAM	32GiB DDR3 1600MHz	128GiB DDR3	192GiB DDR4 1200MHz + 16GiB MCDRAM 7200MHz
Interconnect	1GbE	56Gbps FDR InfiniBand	1GbE

over 100 KNL chassis (4 nodes per chassis) were available for distribution. Considering the variety of specialised scientific computing workloads and their unique requirements, adopting decommissioned hardware with fixed configurations and limited flexibility for customisation does not guarantee performance improvements.

Given KNL's age and architectural nuances, it's essential for prospective adopters to evaluate its relevance to their workloads. This brief reference guide presents preliminary KNL benchmarks on modern software stacks, offering a snapshot of KNL's potential for informed decision-making.

2 Technical Information

The primary audience for KNL distribution in Africa is emerging HPC centres and existing HPC centres refreshing previously adopted hardware (predominantly Stampede1). New HPC centres are likely to prioritise evaluating KNL compatibility with their scientific computing workloads, while existing HPC centres will also care to know the performance delta between their current HPC and new KNL nodes.

Stampede benchmarks were performed on a system configuration that closely aligns with the HPC Ecosystems community configuration (OpenHPC 2.x Slurm configured with 1-Gigabit interconnect [1GbE]).

Benchmarks on Lengau (the region's flagship open science HPC system), while not aligned with HPC Ecosystems community specifications, serve as a baseline to compare the performance of repurposed Stampede and KNL systems to the current flagship HPC accessible to researchers in the region[5]. This demonstrates clearly that high-performance networking such as Infiniband is an integral part of HPC operations, and need to be included with repurposed systems if these are to serve practical purposes.

Despite its bandwidth and latency limitations, 1GbE was selected as a benchmarking baseline due to its ubiquity and cost-effectiveness, particularly among HPC Ecosystems sites and emerging centres, where 1GbE is often the only available interconnect. It provides a realistic, worst-case scenario aligned with the reference guide's scope, making it suitable for normalized, memory- and CPU-focused benchmarks.

Table 1 outlines the system specifications for the reference systems and the KNL benchmark system.

2.1 MCDRAM HBM

MCDRAM can be configured as flat memory, cache, or a hybrid of both; all benchmarks here reflect chosen HBM configuration. Using the high-bandwidth memory instead of the conventional DDR memory yields performance gains exceeding 100% for sparse matrix workloads like CFD and weather modeling.

3 The Known Limitations

3.1 Hardware

Stampede1 and Stampede2 are outdated by today's standards. While still functional, their limitations hinder meaningful comparison to modern systems. Table 2 outlines their key performance constraints and considerations. Only two KNL nodes were available for testing.

Table 2: Key Hardware Limitations Affecting Performance

	Stampede1 (E5-2680)	KNL (Xeon Phi 7250)
CPU	Only 8 cores per socket. Lacks modern AVX extensions.	Lightweight cores. Low clockspeed. No L3 cache.
RAM	Only 32GiB per node.	Only 16GiB of fast MCDRAM 1600MHz.
Notes	12-Year old technology. Support discontinued 2019.	Discontinued in 2018. Memory configuration must be tailored to match the specific needs of each application.

3.2 Methodology

Due to the use of 1GbE networking and a limited number of nodes, the benchmarking exercise could not meaningfully assess weak or strong scaling, both of which require larger node counts and high-performance interconnects. As a result, the focus shifted to evaluating what could be achieved within the constraints of a small-scale deployment. Testing was performed by a single user under controlled conditions, which resulted in minimal data scatter. Error bars were therefore omitted from the graphs to reduce clutter.

Concerns related to cost, reliability and long-term software sustainability, whilst important, are beyond the scope of this preliminary investigation. In practice across HPC Ecosystems deployments, repurposed HPC systems have demonstrated remarkable durability, with hardware failures occurring infrequently. This is consistent

---

[5] Lengau (Top500.org) – https://www.top500.org/system/178793/

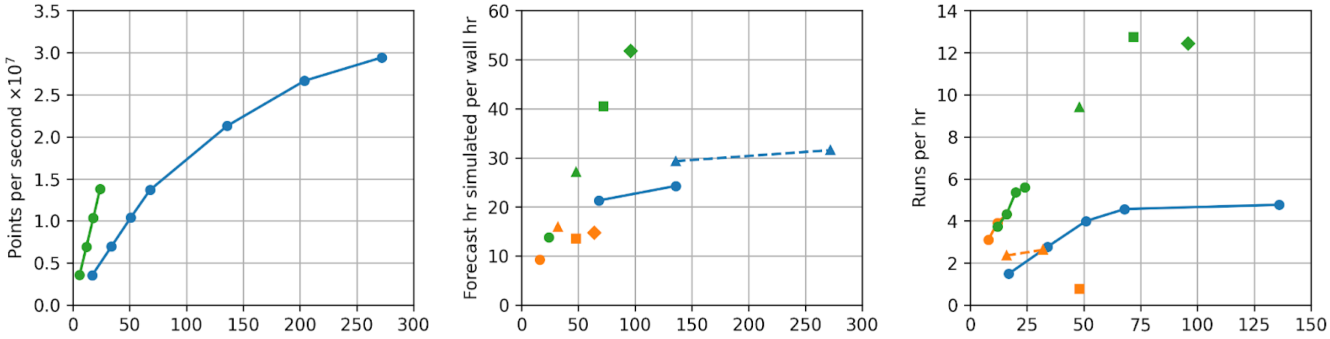

(a) OpenMP

(b) WRF

(c) Palabos

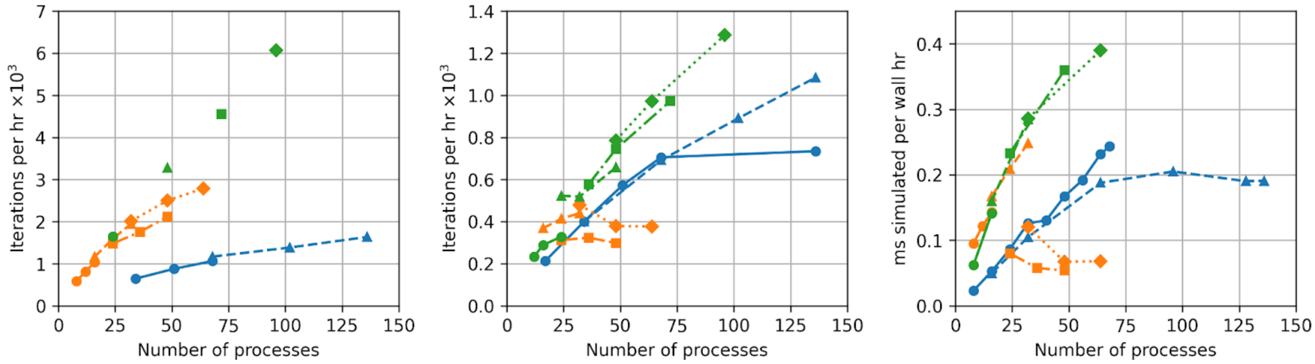

(d) Ansys Fluent

(e) OpenFOAM simpleFoam

(f) OpenFOAM plasmaArc

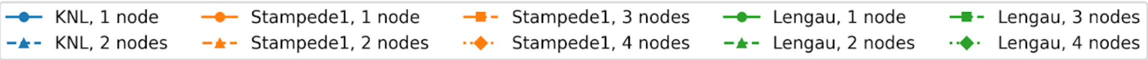

Figure 1: Benchmarking Results of Stampede1, Lengau, and KNL nodes: multiple applications, higher is faster

with our experience at South Africa's CHPC, where financial and other considerations have mandated operating the main HPC system well beyond its expected retirement date.

4 Benchmarking

Figure 1 offers a collection of performance results for each of the Stampede1, Lengau, and KNL reference systems based on a series of application benchmarks whose findings are explored in Section 5.

5 Findings

5.1 Benchmark results

OpenMP (graph 'a'): This Mandelbrot area calculation with no SIMD instructions clearly illustrates relevant characteristics of the KNL processor: Near-perfect scaling over the entire range of 68 physical cores, with further less efficient scaling on multi-threads, up to the full 272 threads. Per-core performance is poor, but the large number of cores and threads allows gives peak single node performance of nearly double that of a Lengau node. This test requires very little memory, but benefits from HBM. The MCDRAM is equally effective when configured as cache or as the preferred part of a combined memory space. The best performance was achieved with Intel's modern OneAPI compiler.

Weather Research and Forecasting model (graph 'b'): WRF is a code with great practical relevance to the ecosystems community. On Stampede1, the best performance was obtained by using all 16 physical cores for MPI processes and no additional OpenMP threads. Using 1 GbE networking, scaling was only satisfactory up to two nodes. On KNL, using two threads per physical core (68 cores and 2 OpenMP threads for a total of 136 processes per node), and the HBM configured as cache memory produced the best performance. A KNL node produced approximately 2.5 times the performance of a Stampede1 node, which is in broad agreement with results in [3, 4].

Palabos (graph 'c'): Palabos is an open-source Lattice-Boltzmann application for fluid dynamics. Performance is somewhat affected by memory bandwidth, thus requiring HBM either as cache or as preferred memory space. The application benefits from high-performance networking. With 1GbE, adding more nodes did not improve performance for either Stampede1 or Stampede2 nodes. Lengau's 56 Gbps FDR Infiniband allowed scaling to three nodes for this test case, which is quite small.

Ansys Fluent (graph 'd'): Although the commercially licensed Fluent segregated solver for fluid dynamics is very similar in principle to simpleFoam, the implementation and performance profile differ dramatically. It is evident that the performance is much less sensitive to memory bandwidth and network performance. Satisfactory scaling was obtained up to 4 Stampede1 nodes, notwithstanding the 1GbE network. The KNL performance shows that the high-bandwidth memory, configured either as cache or as preferred memory, cannot mask the weakness of the lightweight processor cores. Because Fluent is commercially licensed per core, running the software on KNL nodes is not advised.

simpleFoam (graph 'e'): simpleFoam is a standard segregated solver for incompressible flow from the OpenFOAM library. The result from Lengau illustrates the good scaling obtained from 56 Gbps FDR Infiniband. By contrast, Stampede1 performance does not scale beyond two nodes on 1 Gbps ethernet, similar to the WRF benchmark. In similar fashion to the WRF results, the KNL per-node performance is reasonable, with a single KNL node marginally outperforming two Lengau nodes, in similar fashion to the results reported in [4, 8]. Using the multi-threaded capability of the KNL processor does not produce improved performance. MCDRAM should be configured either as cache or as preferred memory space. Although the KNL performance continues to scale over both available compute nodes, when compared to the Lengau result it is evident that the 1GbE network is a serious constraint.

OpenFOAM® plasmaArc (graph 'f')[6]: plasmaArc is a proprietary magnetohydrodynamics solver built on the OpenFOAM library. On the Stampede1 system, peak performance was obtained when using two full nodes (32 cores total). In-node scaling was not linear, with only small gains observed when moving from 8 to 16 cores per node. Beyond two nodes the performance degraded rapidly – this is once more likely due to the latency and bandwidth limitations of the 1GbE interface, as shown by the better multi-node scaling obtained on Lengau. On the KNL system peak performance was obtained when using one full node (68 physical cores with no multi-threading), and matched the peak performance reached on two nodes of the Stampede1 system. In-node scaling for a single node was largely linear, provided high-bandwidth memory (HBM) was prioritized and the problem size fitted completely into the available HBM. Multi-node scaling was poor, again most likely due to bottlenecks at the 1GbE interface.

5.2 Insights

The results showed consistently that performance per core on KNL was generally poor, but high core and thread counts combined with HBM consistently offset this in most benchmarks.

The performance of a single KNL node was equivalent to between one and two Stampede1 nodes, with overall peak performance of the KNL cluster matching or exceeding the peak performance of the Stampede1 system.

Scaling of several of the benchmarks remained linear up to a full node core count on KNL, indicating that a more performant interconnect than 1GbE may provide additional benefits; in this case however, the performance deltas observed when extending the benchmarks to multiple nodes were variable and problem-dependent.

Multithreading performance on KNL was seen to offer anything from zero benefit to significant improvements depending on the specifics of the software in use - in particular, codes taking advantage of OpenMP hybrid parallelisation appeared to gain more from multithreading on all platforms.

MCDRAM delivered significant performance gains, especially at higher core counts. Achieving this requires problem sizing to fit the high-speed memory and proper cluster and middleware configuration.

Fixed hardware configurations aren't always practical, so evaluating available resources is key. Testing KNL with GbE highlighted advantages from its core count, MCDRAM, and power efficiency. Despite interconnect limits, performance aligned with early evaluations [1, 2, 7].

The Xeon Phi's 10-year lifespan, ending in 2020, left behind sparse documentation and a largely inactive support community. Fine-grained KNL optimizations for modern stacks are rare, though tools like OpenHPC still offer guidance for OPA. Integrating HBM with current schedulers remains challenging and requires significant effort to uncover.

Although community support is limited, KNL remains relatively easy to configure due to its x86 foundation. Performance hinges on proper HBM setup. Within the HPC Ecosystems Project, hardware uniformity and an active virtual community not only improves deployment support, but also presents a broader community to contribute towards maintaining and updating documentation for KNL on modern software stacks.

6 Conclusions

KNL continues to offer value for select workloads, especially when compared to older Xeon generations. Future work will expand benchmarking across diverse scientific applications, while also addressing operational factors like cost and reliability. The aim is to produce a comprehensive adoption guide for modern KNL deployments.

Acknowledgments

This paper is published by permission of Mintek, and was supported in part by internal funding on Science Vote Project PDR-24002. The authors would like to acknowledge the Texas Advanced Computing Center, University of Texas at Austin and the broad community of volunteers from the HPC Ecosystems Project Community Partners and Affiliates. Some editing of grammar and phrasing was performed using Microsoft CoPilot.

References

[1] Tyler Allen, Christopher S. Daley, Douglas Doerfler, Brian Austin, and Nicholas J. Wright. 2017. Performance and Energy Usage of Workloads on KNL and Haswell Architectures. In *Proceedings of the International Workshop on Performance Modeling, Benchmarking and Simulation of High Performance Computer Systems (PMBS)*, Vol. 10724. Springer, 236–249. doi:10.1007/978-3-319-72971-8_12

[2] G. N. Barakos and M. A. Woodgate. 2017. *KNL Performance Comparison: HLBM.* Technical Report. University of Glasgow, CFD Laboratory. http://archer.ac.uk/community/benchmarks/archer-knl/KNL_perf_HLBM.pdf

[3] Samuel Elliott. 2015. Performance Analysis and Optimization of the Weather Research and Forecasting Model (WRF) on Intel® Multicore and Manycore Architectures. https://sc15.supercomputing.org/sites/all/themes/SC15images/src_poster/poster_files/spost161s2-file2.pdf

[4] Intel Corporation. 2016. Intel® Xeon Phi™ Processor Applications Showcase. White paper. https://www.intel.com/content/dam/develop/external/us/en/documents/nov2016-intel-xeon-phi-processor-showcase.pdf Accessed: 2025-06-26.

---

[6] plasmaArc solver – https://github.com/quinnreynolds/plasmaArc

[5] Bryan Johnston. 2019. HPC Ecosystems Project: Facilitating Advanced Research Computing in Africa. In *Proceedings of the Practice and Experience in Advanced Research Computing (PEARC '19)*. Association for Computing Machinery, New York, NY, USA, Article No. 107, 1–3. doi:10.1145/3332186.3333264

[6] Bryan Johnston, Lara Timm, David Macleod, and John Poole. 2024. Ten Years of the HPC Ecosystems Project: Transforming HPC in Africa for the Past Decade. In *Proceedings of the Practice and Experience in Advanced Research Computing (PEARC '24)*. Association for Computing Machinery, New York, NY, USA, 1–8. doi:10.1145/3626203.3670537

[7] George S. Markomanolis and Saber Feki. 2016. Performance Evaluation of Weather Research and Forecast (WRF) on Intel Knights Landing Processor (KNL). https://www.ixpug.org/docs/sc16bof/08-sc16_knl_v5.pdf

[8] Ravi Ojha, Prasad Pawar, Sonia Gupta, Michael Klemm, and Manoj Nambiar. 2017. Performance Optimization of OpenFOAM* on Clusters of Intel® Xeon Phi (TM) Processors. In *2017 IEEE 24th International Conference on High Performance Computing Workshops (HiPCW)*. 51–59. doi:10.1109/HiPCW.2017.00019

[9] Karl W. Schulz, C. Reese Baird, Yiannis Georgiou, David Brayford, Gregory M. Kurtzer, Thomas Sterling, Derek Simmel, Nirmala Sundararajan, and Eric Van Hensbergen. 2016. Cluster Computing with OpenHPC. In *Proceedings of the First Workshop on HPC Systems Professionals (HPCSYSPROS '16)*. Association for Computing Machinery, New York, NY, USA. doi:10.1145/1235