



STORAGE REFERENCE ARCHITECTURE

Graphcore IPU based systems with DDN AI400X2



Table of Contents

1	Introduction	3
2	Graphcore system overview	4
2.1	Innovate at massive scale	4
2.2	Disaggregated to scale with your needs	5
2.3	Scale-out with IPU-Fabric	5
2.4	Datacentre compatibility	5
3	DDN A ³ I solutions overview	6
3.1	DDN AI400X2	6
3.2	Accelerating AI with DDN shared parallel architecture	7
3.3	Limitless scaling and end-to-end AI enablement.....	7
4	Reference architecture	8
4.1	Poplar hosts	8
4.2	Network connections.....	9
4.3	AI400X2 appliance	9
4.4	System configuration.....	9
4.4.1	Poplar host configuration	9
4.4.2	Storage configuration	10
4.4.3	Networking configuration	10
5	Performance results.....	11
5.1	ResNet 50	11
5.1.1	ResNet 50 performance tests	11
5.2	BERT Large	12
5.2.1	BERT Large performance tests	12
5.3	Standardized storage benchmark (AI Workload).....	13
5.4	FIO and application-level sanity-checks.....	15
5.4.1	Each FIO job on each host has exclusive access to its own data file	15
5.4.2	FIO jobs on a host all access the same file	16
5.4.3	The FIO jobs share the same file with the same job on other test hosts	16
5.4.4	The FIO jobs on all hosts all share one single file	17
5.5	Conclusion on benchmarks.....	18
6	Achieve the most from AI with Graphcore and DDN	19
7	Revision history	20



1 Introduction

Graphcore provides machine intelligence compute systems. These systems are built around the Intelligence Processing Unit (IPU), a new processor specifically designed for machine learning. The IPU's unique architecture enables the exploration and deployment of entirely new types of AI workloads – to drive advances in machine intelligence.

DDN AI data platforms are widely recognized and trusted globally. The unique DDN shared parallel architecture ensures high levels of performance, scalability, security, and reliability for AI applications at every scale. DDN is the leading provider of at-scale data storage solutions, powering the most demanding machine learning programs in every industry today.

Organizations everywhere are pursuing AI to transform and evolve their activities. Graphcore and DDN bring intelligent compute and storage platforms together in a converged AI infrastructure solution that enables and accelerates machine learning workloads at every scale. These fully integrated solutions provide optimal performance for a wide range of AI workflows and simple transition from prototype to production.

This document describes a fully validated reference architecture developed jointly by Graphcore and DDN. This solution integrates Graphcore systems built using IPU-M2000™ IPU-Machines™ and DDN AI400X2 appliances to provide a turnkey AI infrastructure that's easy to deploy and delivers optimal performance and efficiency.

All integration and set-up information in this document is also valid for new systems built using Bow-2000™ IPU-Machines, however performance with Bow-2000s will be up to 40% higher.

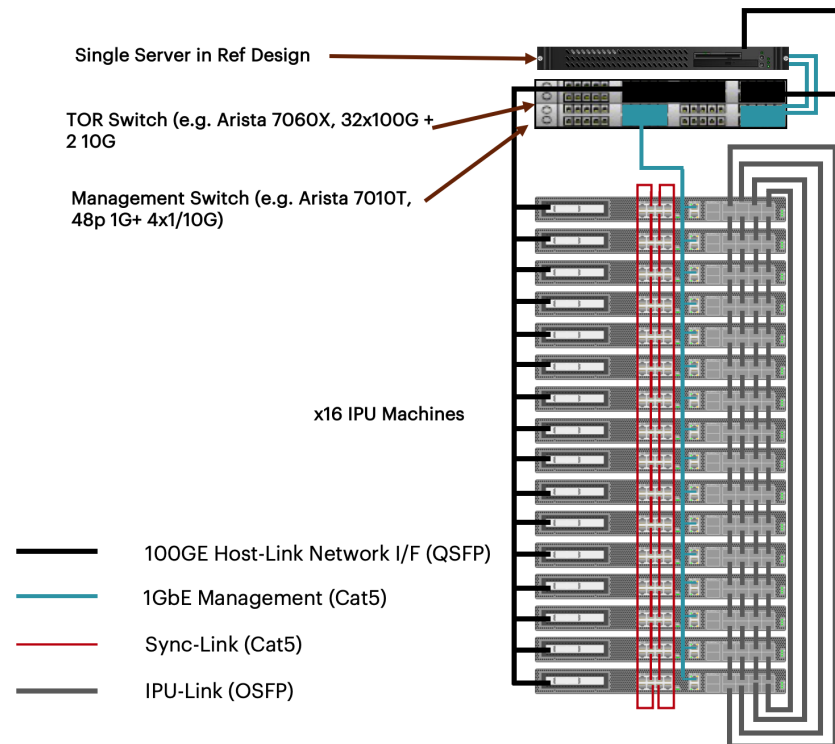
The information in this document applies to Graphcore IPU Pod systems, which covers both classic IPU-POD™ systems (such as the IPU-POD₆₄) and Bow™ Pod systems (such as Bow-Pod₆₄). The term IPU-Machine refers to the blades installed in your system, so IPU-M2000 in IPU-POD systems and Bow-2000 in Bow Pod systems.

2 Graphcore system overview

IPU Pod systems are created by connecting multiple IPU-Machines, third party CPU server units, and storage appliances, allowing powerful and flexible AI infrastructure designs for machine intelligence training and inference workloads.

2.1 Innovate at massive scale

The core building block of an IPU Pod is the IPU-Machine, contained in a slim 1U blade. This is the fundamental compute engine for machine intelligence from Graphcore, containing four IPU processors, accelerators designed from the ground up for AI. An individual IPU-Machine can deliver up to 1 petaFLOPS (IPU-M2000) or 1.4 petaFLOPS (Bow-2000) of AI compute and has up to 260GB Memory (3.6GB In-Processor-Memory™ and up to 256 GB Streaming Memory™), enabling it to handle the most demanding of machine intelligence workloads.



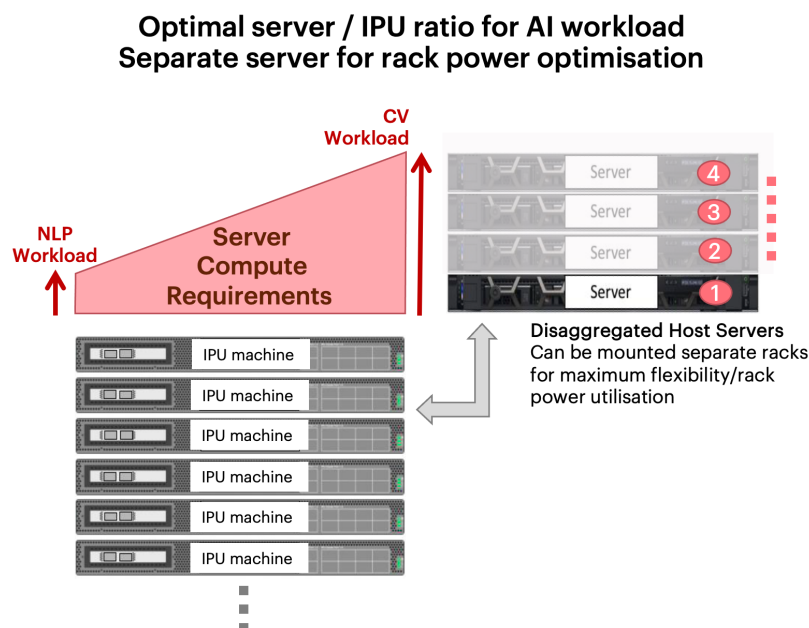
The IPU-Machine has a flexible, modular design, so you can start with one and scale to thousands. IPU-Machines can work as standalone systems or can be interconnected in racks. Individual IPU Pod racks can then be interconnected with other Pods using the 2.8Tbps high-bandwidth, near-zero latency IPU-Fabric™ interconnect architecture, to grow to supercomputing scale.

The classic IPU-POD₆₄ reference design is a rack solution containing 16 IPU-M2000 IPU-Machines, one to four host servers (the default is one host server in the reference configuration), network switches and platform software. It is designed to deliver 16 petaFLOPS of AI compute in an efficient, flexible and pre-qualified configuration. The equivalent Bow based system (Bow Pod₆₄) adopts the same form factor and architecture, with

16 Bow-2000 IPU-Machines, delivering an increase in performance of up-to 40% whilst reducing energy consumption by up-to 16%.

2.2 Disaggregated to scale with your needs

AI workloads have very different compute demands. For production deployment, optimizing the ratio of AI to host compute can maximize performance and efficiency, and improve the total cost of ownership. IPU Pods are disaggregated systems, this means that the ratio between the number of host servers and switches and the number of IPU-Machine units is not fixed. The system can be built so that it is ideally matched to the production workload.



For example, NLP models require very little server CPU interaction and utilization, whereas CNN-based workloads such as CV require a larger proportion of scalar computing and would benefit from more server CPU being involved. The system can be tailored for the workload.

2.3 Scale-out with IPU-Fabric

IPU-Fabric is Graphcore's innovative low-latency, all-to-all IPU interconnect. Eliminating communication bottlenecks with reliable deterministic performance it is highly efficient whatever your scale.

2.4 Datacentre compatibility

IPU Pod systems bring together powerful IPU compute with a choice of best-in-class datacentre technologies and systems from leading technology providers in flexible, pre-qualified configurations, to ensure your datacentre is operating with maximum efficiency and performance, while making your datacentre AI deployments simpler and faster.

In this document, the DDN AI400X2 appliance is evaluated as a storage solution to support AI workloads running on an IPU-POD₆₄.



3 DDN A³I solutions overview

DDN A³I solutions (Accelerated, Any-Scale AI) are architected to achieve the most from at-scale AI training applications running on Graphcore IPU. They provide predictable performance, capacity, and capability through a tight integration between DDN AI400X2 appliances and IPU Pod systems. Every layer of hardware and software engaged in delivering and storing data is optimized for fast, responsive, and reliable access.

The deep integration of DDN AI appliances with IPU Pods ensures a reliable experience. DDN A³I solutions are highly configurable for flexible deployment in a wide range of environments and scale seamlessly in capacity and capability to match evolving workload needs. DDN A³I solutions are deployed globally and at every scale, from a single AI training system all the way to the largest AI infrastructures in operation today.

DDN brings the same advanced technologies used to power the world's largest supercomputers in a fully integrated package that's easy to deploy and manage. DDN A³I solutions are proven to maximum the benefits of using at-scale AI workloads on Graphcore systems.

3.1 DDN AI400X2

The AI400X2 appliance is a fully integrated and optimized shared data platform with predictable capacity, capability, and performance. Every AI400X2 appliance delivers over 90 GB/s throughput and 3M IOPS directly to the IPU-Machines in the IPU Pod system. Shared performance scales linearly as additional AI400X2 appliances are integrated to the AI infrastructure. The all-NVMe configuration provides optimal performance for a wide variety of workload and data types and ensures that IPU Pod operators can achieve the most from at-scale AI applications, while maintaining a single, shared, centralized data platform.

The AI400X2 appliance integrates the DDN A³I shared parallel architecture and includes a wide range of capabilities, including automated data management, digital security, and data protection, as well as extensive monitoring. The AI400X2 appliances enables IPU Pod operators to go beyond basic infrastructure and implement complete data governance pipelines at-scale.

The AI400X2 appliance integrates with IPU Pod systems over an Ethernet network. It is available in multiple all-NVMe capacity configurations. Optional hybrid configurations with integrated HDDs are also available for deployments requiring high-density deep capacity storage. Contact DDN Sales for more information.



3.2 Accelerating AI with DDN shared parallel architecture

The unique DDN A³I shared parallel architecture and client protocol ensures high levels of performance, scalability, security, and reliability. Multiple parallel data paths extend from the drives all the way to containerized applications running on the IPU processors. With DDN's true end-to-end parallelism, data is delivered with high-throughput, low-latency, and massive concurrency in transactions. This ensures applications achieve the most from IPU based systems with all processor cycles put to productive use. Optimized parallel data-delivery directly translates to increased application performance and faster completion times. The DDN A³I shared parallel architecture also contains redundancy and automatic failover capability to ensure high reliability, resiliency, and data availability in case a network connection becomes unavailable.

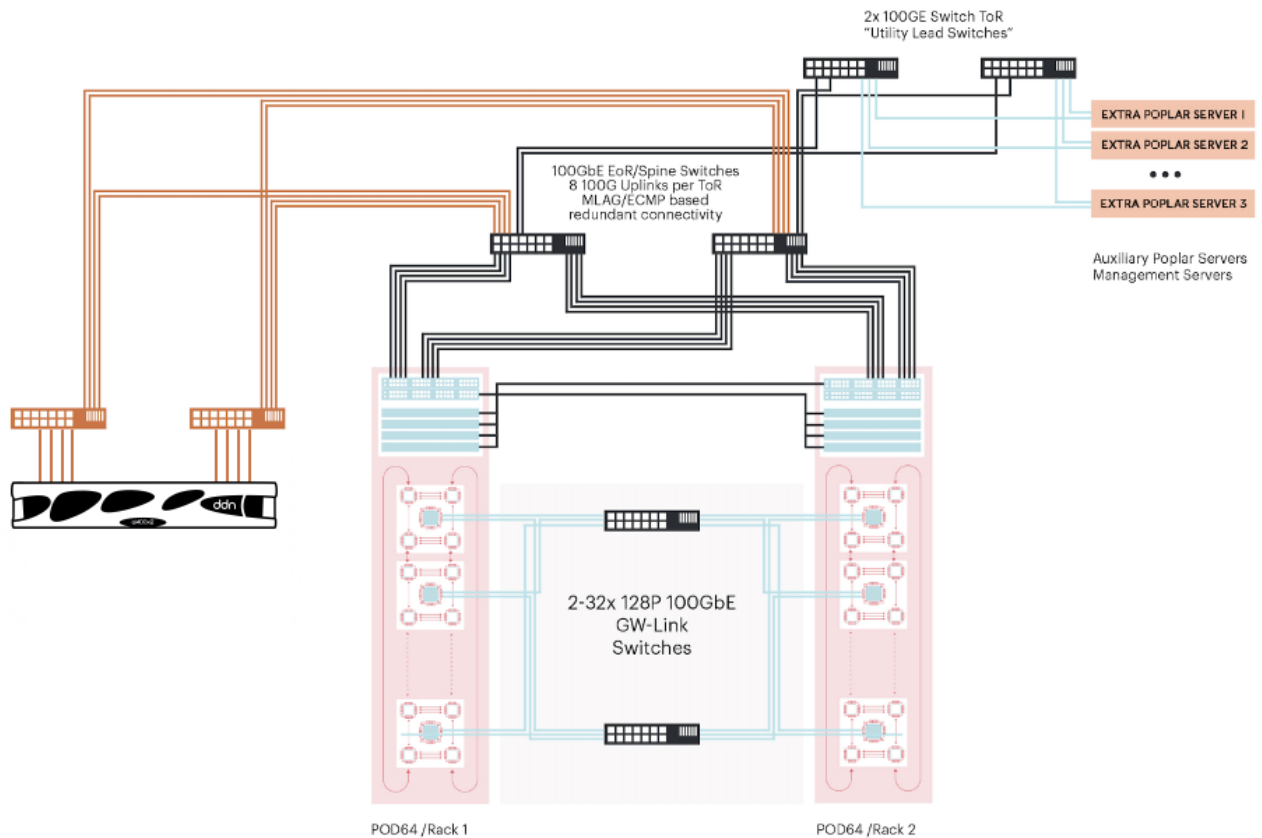
3.3 Limitless scaling and end-to-end AI enablement

DDN A³I solutions enable and accelerate end-to-end data pipelines for AI workflows at every scale. The DDN shared parallel architecture enables concurrent and continuous execution of all phases of AI workflows from a single data repository. This eliminates the management overhead and risks of moving data between storage locations. At the application level, data is accessed through a standard, highly inter-operable file interface, for a familiar and intuitive user experience. Optimized data delivery ensures best application performance possible, from ingest through data preparation, training, inference, and validation. Advanced data management capabilities ensure automated data handling throughout, making the AI400X2 an ideal complement for Graphcore IPU Pod systems at every scale.



4 Reference architecture

This reference architecture integrates IPU-Machines with DDN AI400X2 appliances. It delivers fully optimized end-to-end AI acceleration on IPU processors. This solution greatly simplifies the deployment of IPU Pod configurations, ensures nominal performance and efficiency for maximum IPU processor saturation, and high levels of scalability.



This reference architecture is designed to deliver an optimal balance of technical and economic benefits for a wide range of common use cases for AI. Using the AI400X2 appliance as a building block, solutions can scale linearly, predictably, and reliably in performance, capacity, and capability. For configurations with requirements beyond the base reference architecture, it's simple to scale the data platform with additional AI400X2 appliances.

4.1 Poplar hosts

The Poplar host is the head node for the IPU Pod.

In the Pod system configuration used for the purpose of this document, the Poplar host consists of the following:

- Hardware: 8 x Dell PowerEdge R640 Intel Xeon Platinum 8168 with 96 hyper-threads, 768GB RAM, Mellanox connectX-5.



- Software: Ubuntu 18.04.06 LTS 5.4.0-66 kernel

4.2 Network connections

The network configuration recommended for IPU-POD₆₄ classic and Bow-Pod₆₄ systems is as follows:

- 1G management network: Arista DCS-7010T-48-F
- Top-of-Rack for 100G connectivity: Arista DCX-7060CX-32S-ES-F
- Connected via 8way 100G LAG to: Arista 7060PX4-32-F2 as SPINE 400G.

4.3 AI400X2 appliance

DDN recommends an all-NVMe AI400X2 appliance for Graphcore IPU-POD₆₄ classic and Bow-Pod₆₄ systems. Each AI400X2 appliance arrives fully configured, ready to connect to the IPU Pod via eight 200GbE network links for high-performance data transfers, and provides dedicated 1GbE network ports for management and monitoring.

The DDN AI400X2 in the reference architecture validation is configured with 24 NVMe devices and is running DDN EXAScaler software version 6.0.0.

4.4 System configuration

This section describes the system configuration tested by Graphcore and DDN for the IPU-POD₆₄ classic reference platform. The information below is also applicable to Bow Pod₆₄ systems.

This section includes configuration settings for:

- Poplar host
- Storage
- Networking

4.4.1 Poplar host configuration

The Poplar framework can be used to define graph operations and control the execution and profiling of code on the IPU, as well as to configure the host. A single 100G Ethernet interface for the storage connection is configured with the Mellanox OFED device driver.

The Poplar host has a lossless Ethernet (PFC) configuration applied, but the storage servers and network core switches do not. As a result, PFC is only partly effective and packet loss and retransmission occurs during these tests. The impact on performance has not been quantified.

On each client, the storage driver is configured through the standard DDN procedure using the exa-client package and its deployment script, as follows.

Install the EXA driver:

```
./exa_client_deploy.py --install
```

Configure and tune the EXA driver:



```
./exa_client_deploy.py --configure --lnets "tcp(interface)" --yes
```

Mount the filesystem:

```
$(sshpass -p DDNSolutions4U ssh ${EXAVM} mount_lustre_client --dry)
```

A `fstab` entry can be added using the output of the mount command.

4.4.2 Storage configuration

The DDN AI400X2 is configured using the default manufacturing settings.

The Poplar host accesses a single namespace shared filesystem called `ai400x2`.

The specific tuning shown below is added on the AI400X2 to increase TCP performance. This will be set by default in the next version of EXAScaler (from version 6.0.1) and does not need manual change.

```
/etc/modprobe.d/lustre.conf # On each AI400X2 VM

(...)

options ksocklnd conns_per_peer=15
```

4.4.3 Networking configuration

The following information refers to the system networking configuration:

- The MTU for the network interfaces on the storage network is 9000 (Poplar hosts and AI400X2 VMs)
- The LNET to access the AI400X2 filesystem uses TCP (no RDMA through RoCE)
- Specific tunings are transparently set on the Poplar host through the DDN deployment script during the configuration, use `--dry-run` option to display the list
 - Relax ordering is set on Mellanox adapters
 - Some `ethtool` tunings are applied
 - IP routes and rules are updated to scatter the work across all interfaces



5 Performance results

All benchmarking results in this section are based on the IPU-POD₆₄ classic system with IPU-M2000 IPU-Machines.

Performance tests were undertaken using:

- ResNet 50
- BERT Large
- Standardized Storage Benchmark for AI Workloads
- FIO

For the ResNet and BERT tests, it is important to keep in mind that the focus is on the IPU performance. The storage solution needs to be fast enough to keep the processors fed with all the data they can consume.

5.1 ResNet 50

ResNet 50 tests were performed using 8 concurrent hosts as AI400X2 clients. The results (averaged over multiple runs) from the first steps of the training phase are shown here compared with local NVMe drives. Only the first steps result in reads from the storage because data is held in filesystem cache subsequently.

5.1.1 ResNet 50 performance tests

The following tables summarise the benchmark results in images per second.

Step	NVMe	AI400X2
1	17,805	16,492
2	17,738	16,537
3	17,773	16,223
4	17,854	16,987
5	18,315	18,330
6	18,285	18,272
7	18,253	18,250
8	18,211	18,251
9	18,210	18,227

The following table demonstrates the impact of the number of concurrent jobs on the first training step. This shows that up to 8 concurrent clients have negligible impact on the throughput.



Jobs	NVMe	AI400X2
1	17,532	16,262
2	17,898	16,543
3	17,766	16,503
4	17,700	16,428
5	17,804	16,487
6	17,954	16,550
7	17,799	16,538
8	17,795	16,458

5.2 BERT Large

BERT Large pretraining using 8 hosts concurrently. Figures shown are average per host.

5.2.1 BERT Large performance tests

The following tables summarise the throughput benchmark results in items/second.

Filesystem	config=128	config=384
AI400X2	3,708	1,054
NVME	3,703	1,052

These results demonstrate that the storage components do not limit the performance of the system.



5.3 Standardized storage benchmark (AI Workload)

The following results refer to the SPECstorage Solution 2020 benchmark. The benchmark simulates an AI application and runs an increasing number of concurrent workloads against the storage appliance.

AI Jobs	Storage (Ops/sec)	Latency (msec)	Storage (MB/Sec)
10	4350	0.965	956
20	8700	1.012	1911
30	13050	0.992	2864
40	17400	1.016	3821
50	21750	1.008	4778
60	26100	1.027	5732
70	30450	1.041	6687
80	34800	1.035	7644
90	39150	1.068	8597
100	43500	1.099	9552
110	47850	1.087	10506
120	52200	1.115	11462
130	56550	1.154	12417
140	60900	1.166	13372
150	65250	1.210	14326
160	69600	1.254	15281
170	73950	1.257	16238
180	78300	1.265	17190
190	82650	1.390	18146
200	87000	1.337	19104
210	91350	1.305	20059
220	95700	1.825	21014
230	100050	2.000	21971
240	104400	1.662	22925



The SPECStorage results in the table above illustrate that the performance of the system scales as expected, and that the storage components do not restrict the performance of the compute units. This validates the proposed configuration.



5.4 FIO and application-level sanity-checks

We used Flexible IO Tester (FIO), a storage benchmarking and workload simulation tool.

For file-based tests, the results were:

- Aggregate peak throughput of 65GB/s using 512KB block random reads, 8 clients
- Peak single client performance of single test (as above) is ~9GB/s

The following 4 charts show testing across 4 different use cases which represent possible ML application scenarios. This enables us to test the storage appliance using a variety of cases, from multiple independent ML jobs across many servers, to one large scale out training utilizing many servers.

5.4.1 Each FIO job on each host has exclusive access to its own data file

The following diagram shows the Read Throughput (MB/s) recorded by increasing the number of clients for the single AI400X2 appliance connected to the IPU Pod.

The benchmark simulates a workload where multiple agents operate concurrently on separate data.

A total of 524,288 random read operations are executed.

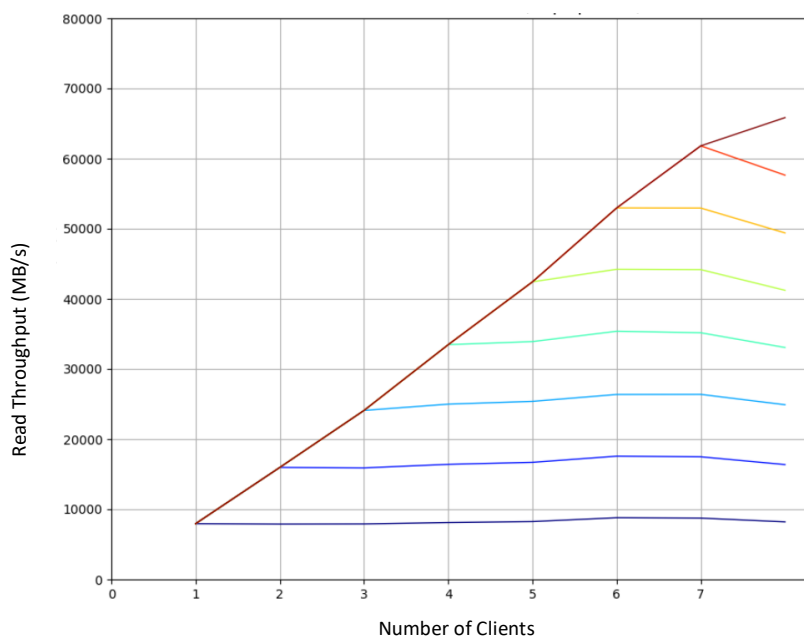


Figure 1 – Each FIO job on each host has exclusive access to its own data file

The results show nearly linear performance, which in the light of the size of the system it was serving is compelling. The reference point is a run load test with one host, achieving approximately 9Gb/s read throughout. The benchmark incrementally adds hosts, one at the time. The results show that the performance scales almost linearly.



With the system fully loaded, with 8 clients, we observe almost 65Gb/s across, through 2 switches. The maximum theoretical performance from this connectivity is 8x 100Gb/s, hence the benchmark result is highly competitive. Similar considerations can be made for all of the FIO benchmarks.

5.4.2 FIO jobs on a host all access the same file

The following diagram shows the Read Throughput (MB/s) recorded by increasing the number of Clients for the single AI400X2 appliance connected to the IPU Pod.

The benchmark simulates multiple IPU training runs using a single server.

A total of 524,288 random read operations are executed.

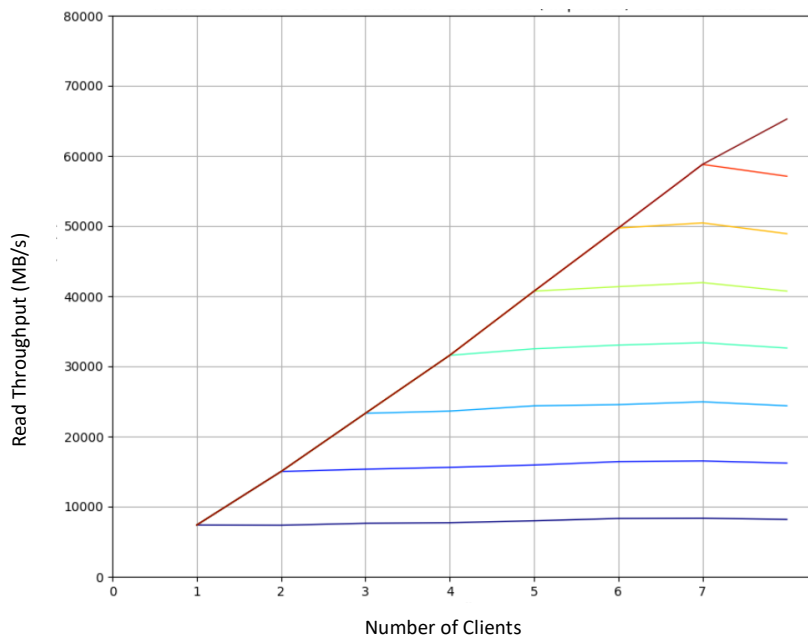


Figure 2 – All FIO jobs on a host access the same file

5.4.3 The FIO jobs share the same file with the same job on other test hosts

The following diagram shows the Read Throughput (MB/s) recorded by increasing the number of clients for the single AI400X2 appliance connected to the IPU Pod.

The benchmark simulates multi IPU training on a scale out system using multiple IPU hosts, where multiple FIO jobs share the same file with the same job on other test hosts.

A total of 524,288 random read operations are executed.

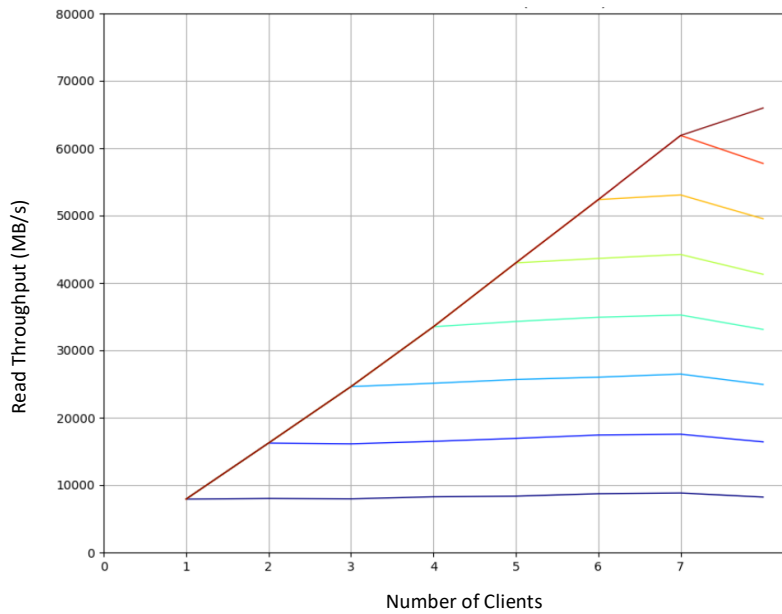


Figure 3 - The FIO jobs share the same file with the same job on other test hosts

5.4.4 The FIO jobs on all hosts all share one single file

The following diagram shows the Read Throughput (MB/s) recorded by increasing the number of clients for the single AI400X2 appliance connected to the IPU Pod.

The benchmark simulates multi IPU training on a scale out system using multiple IPU hosts, where the FIO jobs on all hosts all share one single file.

A total of 524,288 random read operations are executed.

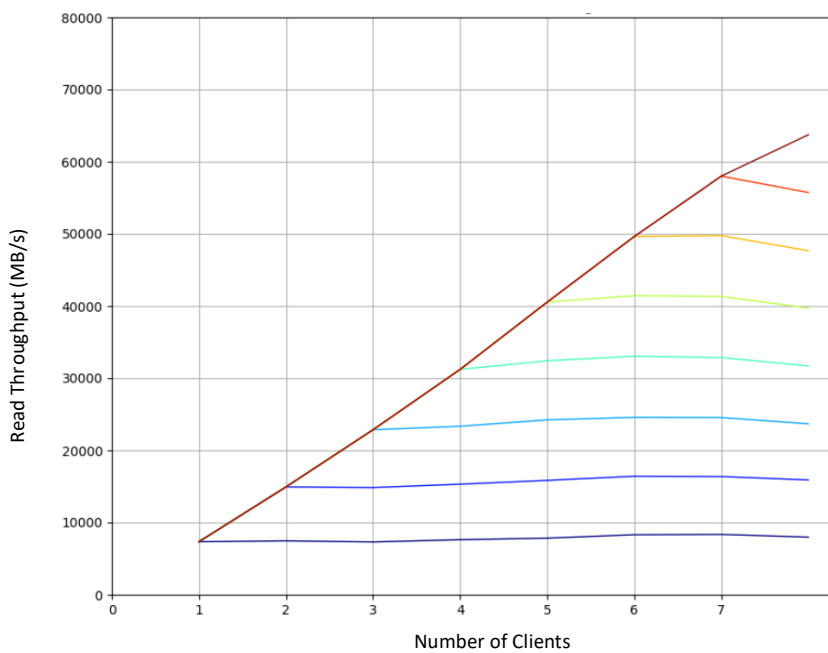


Figure 4 - The FIO jobs on all hosts all share one single file



5.5 Conclusion on benchmarks

The testing and benchmarking have confirmed our expectations, with the storage appliance delivering the required performance and scaling characteristics needed by the IPU Pod system to compute the workloads at full performance.

To summarise, the benefit that using the AI400X2 appliance brings to users of IPU Pod systems (based on IPU-M2000 or Bow-2000) are as follows:

- The DDN appliance can supply the data required to feed very large Pod systems
- The results indicate that the configuration scales as expected
- Performance of the AI400X2 appliance has been reliable across a broad range of use cases
- All tests and benchmarks were carried out using a Graphcore standard infrastructure, with no customizations, hence we expect customer deploying IPU Pod systems with DDN storage to experience these benefits.



6 Achieve the most from AI with Graphcore and DDN

Graphcore and DDN propose a fully integrated and optimized AI Infrastructure solution that delivers optimal AI application performance for environments at every scale. The fully validated solution ensures rapid deployment and ease of management, maximizing the benefits of the combined compute and storage technologies. The solution is proven to deliver peak performance at every scale, up to the largest Pod system configurations.

Contact Graphcore and DDN sales for more information.



7 Revision history

Version	Date	Notes
1.0	14 th of March 2022	Initial release

Trademarks and copyright

Graphcore®, Graphcloud® and Poplar® are registered trademarks of Graphcore Ltd.

Bow™, Bow-2000™, Colossus™, Graphcloud™, In-Processor-Memory™, IPU-Core™, IPU-Exchange™, IPU-Fabric™, IPU-Link™, IPU-M2000™, IPU-Machine™, IPU-POD™, IPU-Tile™, PopART™, PopDist™, PopLibs™, PopRun™, PopVision™, PopTorch™, Streaming Memory™ and Virtual-IPU™ are trademarks of Graphcore Ltd.

All other trademarks are the property of their respective owners.

Copyright © 2016-2022 Graphcore Ltd. All rights reserved.