

DGX-1 Evaluation



Please fill out your answers to the best of your abilities.
Your detailed responses will help Scan to effectively assist you ahead of your Proof of Concept.

Company Details

Company name:

Company address:

Company registration no. :

Company VAT no.

Company website:

Company sector:

Number of employees:

Describe Company:

Contact Details

Contact name:

Position within company:

Do you have purchasing authority: Yes No

Email:

Phone:

Mobile:

Skype:

DGX-1 qualification

How many users will leverage the DGX-1 POC system:

What field of research are you contracting:

Are you currently using Tesla GPU products:

What is your preferred start date of your POC?
Please provide three options

Which frameworks are you currently using from the list below:

| | | | | | | | | | |
|-----------|--------------------------|--------|--------------------------|----------------|--------------------------|----------|--------------------------|---------|--------------------------|
| Baremetal | <input type="checkbox"/> | DIGITS | <input type="checkbox"/> | MatConvNet | <input type="checkbox"/> | OpenDeep | <input type="checkbox"/> | Theano | <input type="checkbox"/> |
| CAFFE | <input type="checkbox"/> | CAFFE2 | <input type="checkbox"/> | Microsoft CNTK | <input type="checkbox"/> | Purine | <input type="checkbox"/> | DL4J | <input type="checkbox"/> |
| Chainer | <input type="checkbox"/> | Julia | <input type="checkbox"/> | Minerva | <input type="checkbox"/> | Pylearn2 | <input type="checkbox"/> | CUDA | <input type="checkbox"/> |
| Keras | <input type="checkbox"/> | MXnet | <input type="checkbox"/> | TensorFlow | <input type="checkbox"/> | Torch | <input type="checkbox"/> | PyTorch | <input type="checkbox"/> |
| Other | <input type="text"/> | | | | | | | | |

In relation to your indicated frameworks of choice please confirm if they have been modified? Yes No

What would you expect to be the average data set size for upload?

0-500GB 500GB-1TB 1TB-50TB 50TB+

Please provide a brief description of you POC application (500 characters max):

How does your POC application fit into your organizations larger key application(s), research or goals?

What are the THREE main criterion for success of your POC?

What software (Deep Learning or otherwise) frameworks are you planning to use for your application? Would you be using NVIDIA's version of the Deep Learning frameworks or would you like to use your own modified versions of the frameworks?

Are you familiar with (NV) Docker usage and running containerized applications?

Containerised Frameworks (NVIDIA optimized frameworks for DGX-1 Pascal and/or Volta are DIGITs, Caffe, CUDA, TensorFlow, Caffe2 Theano, MxNet, CNTK, Torch, PyTorch

What additional software would you require (beyond NVIDIA DL frameworks) for your application, or for (pre-) processing your data?

What data are you planning to use in your application? How large are your datasets? What is the predominant datatype of your data (fixed point, single precision, floating point, double precision floating point, text others? Please provide some description as to the structure and the nature of the data you plan to use.

What are the benchmarks metrics that you would use for your application? What would be your benchmarking strategies to have a like-for-like comparison?

What is the Multi-GPU strategy in your application (in terms of code and data staging/processing across GPU's)? Are you pursuing multicore/multi GPU combined Strategy?

How (if at all) do you plan to utilise memory management across GPU's and GPU-CPU combination?

Are you evaluating speed vs accuracy tests? What are your evaluation strategies for such tests?

Which Machine Learning/Deep Learning (or other) algorithms do you plan to use for your POC and how have they been adopted for GPU's and in particular DGX-1? Which CUDA (or CUDA based) Libraries would best support your algorithms application?

Are you testing scalability? If so, what is the sizing of the production environment?

Onsite DGX-1 POC (Subject to Credit Worthiness and Availability)

| | | |
|--|------------------------------|-----------------------------|
| Do you have 3U identified for the DGX-1? | Yes <input type="checkbox"/> | No <input type="checkbox"/> |
| Is the rack at least 35 Inches (88.9cm) deep? | Yes <input type="checkbox"/> | No <input type="checkbox"/> |
| Are four C14 outlets available? | Yes <input type="checkbox"/> | No <input type="checkbox"/> |
| Are you planning to use Infiniband? | Yes <input type="checkbox"/> | No <input type="checkbox"/> |
| Will the DGX-1 be connected to an NFS Server? | Yes <input type="checkbox"/> | No <input type="checkbox"/> |
| What is the inlet temperature of the rack where DGX-1 will be installed (degrees C)? | <input type="text"/> | |
| Does the cooling (CRAC, etc) currently have 3.2kW of excess capacity? | Yes <input type="checkbox"/> | No <input type="checkbox"/> |

Please Note

DGX-1 management network connectivity requires copper ethernet (RJ45 connectors)

DGX-1 requires HTTP access to developer.download.nvidia.com and international.download.nvidia.com for package retrieval

DGX-1 requires HTTPS access to apt.dockerproject.org for Container retrieval

DGX-1 Purchase

What is your timeframe to purchase/deployment?

- A . Immediate Purchase/deployment
- B . 0-6 Months
- C . 6-12 Months
- D . Not Sure

By completing and returning this form, you are agreeing that the services offered by Scan Computers are for internal research and development purposes only.

The platform, performance and results must NOT be used for external publications and/or reviews and/or public benchmarks.

Please return your completed form to servers@scan.co.uk