



NVIDIA DGX POD 관리서버 구성 with DeepOps

BayNex



AGENDA

GPUPOD 관리시스템 필요성

GPUPOD Management System

GPUPOD 관리 시스템 구성

Orchestration & Job Scheduling

GPUPOD 관리 시스템 구축

Deploying GPUPOD Management System

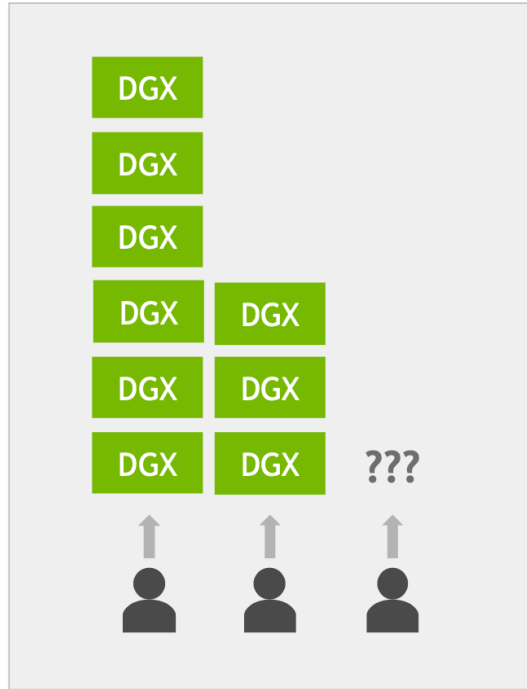
NVIDIA
DGX GPU POD
with Management
The Industry Standard
For AI System

BayNex



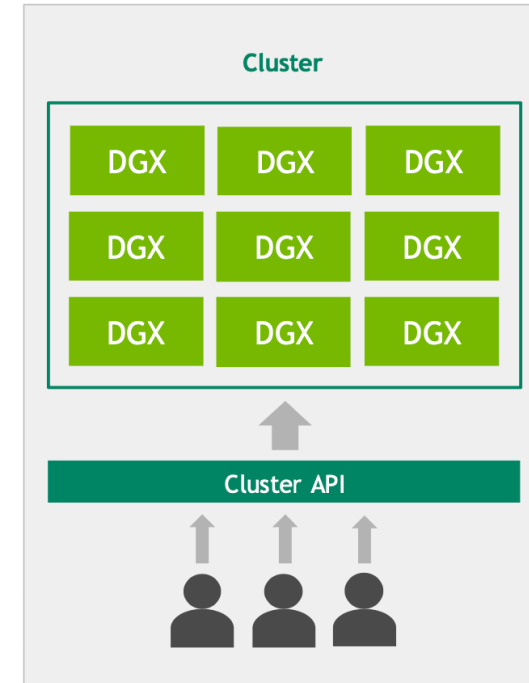
DGX GPU POD Management

Manage resources more effectively



Dedicated Nodes

- User coordination required
- Resources can go unused
- Not scalable
- Difficult to maintain



Cluster with Management

- No user coordination required
- Efficient use of resources
- Highly scalable
- Standardized maintenance

DGX GPU POD

Design Principles



Leverage our ecosystem of proven enterprise-grade software solutions that are fully tested and certified for use on clusters of NVIDIA DGX systems, simplifying the deployment, management, and scaling of AI infrastructure.

Management System Principles

DGX GPU POD 시스템의 인프라 배포, 관리를 쉽게 하고 확장을 자유롭게 하는 관리 시스템

SUPPORTING AI INFRASTRUCTURE

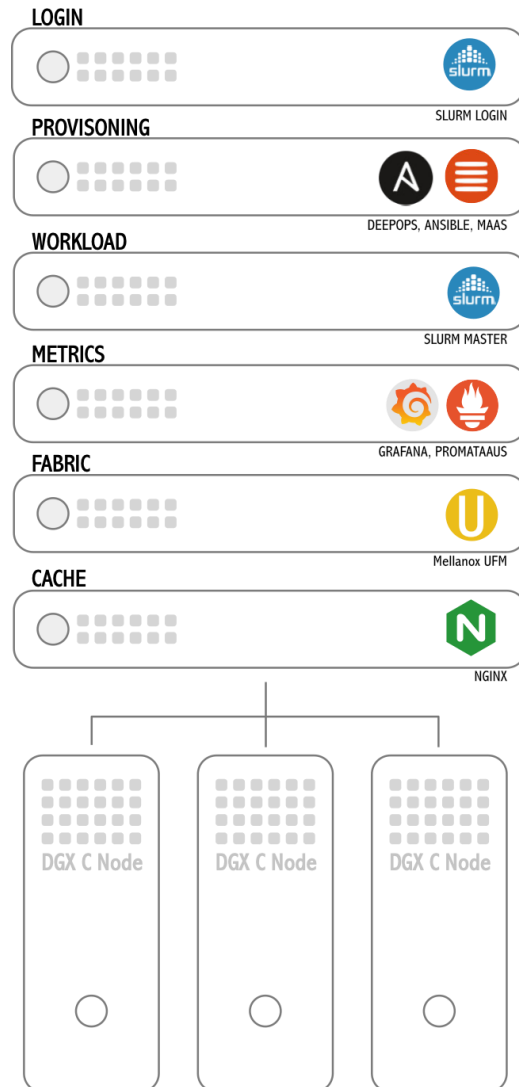
Orchestration &
Job Scheduling
Management System

BayNex



Management Node SOFTWARE

DeepOps with UFM



Fast deployment of Ansible provisioning platform with DeepOps

- Clone DeepOps repository
- Run the initial setup script

Docs on using DeepOps to provision your cluster:

- <https://github.com/NVIDIA/deepops>

Provisioning Management

Ansible



Tool for scripting and automating deployment of software and configuration across the cluster

- 사용이 쉽다
- 에이전트가 필요없다

Ansible Documentation:

- <https://docs.ansible.com>

OS Provisioning Management

MAAS



Automated server provisioning and easy network setup

- Bare metal cloud with on-demand servers
- Super fast install from scratch
- Custom images with pre-installed apps

MAAS Documentation:

- <https://maas.io/docs>

Job Schedule Management Slurm



Open source job scheduler for running large parallel batch jobs.

- fault-tolerant
- highly scalable cluster management
- job scheduling system for large and small Linux clusters.

Slurm Documentation:

- <https://slurm.schedmd.com>
- <https://github.com/NVIDIA/deepops/tree/master/docs/slurm-cluster>

Monitoring Management Promataus with Grapana



Open source monitoring tool for gathering metrics from a large number of nodes

- 강력한 검색
- 훌륭한 시각화
- 간단한 조작

Promataus Documentation

<https://prometheus.io/docs>

Cache Mangement

NGINX



open source software for web serving, reverse proxying, caching, load balancing, media streaming, and more.

- More Scalablity
- Good Perfomance

More info [nginx-cache-proxycy](#):

- <https://github.com/NVIDIA/deepops/blob/master/docs/container/nginx-docker-cache.md>

InfiniBand Fabric Management UFM

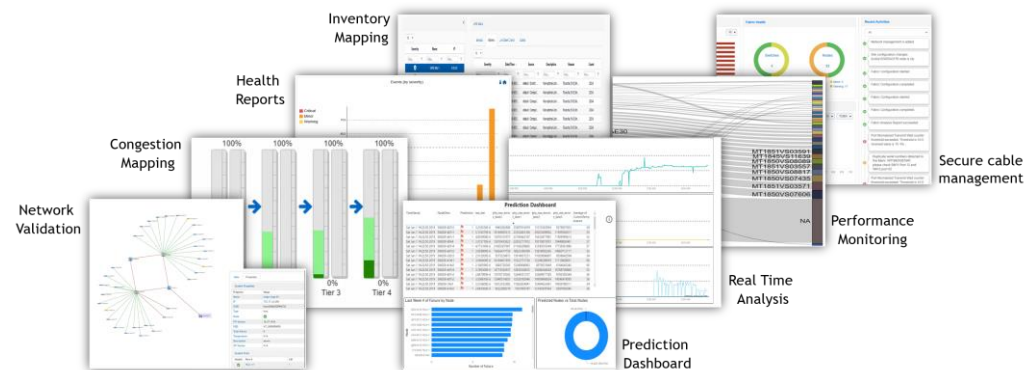


MELLANOX UNIFIED FABRIC MANAGER

- Efficiently monitor
- Manage and preventatively troubleshoot

UFM Documentation

- <https://docs.mellanox.com>



Management Node

LOGIN



SLURM LOGIN

PROVISIONING



DEEPOPS, ANSIBLE, MAAS

WORKLOAD



SLURM MASTER

METRICS



GRAFANA, PROMATAAUS

FABRIC



Mellanox UFM

CACHE



NGINX

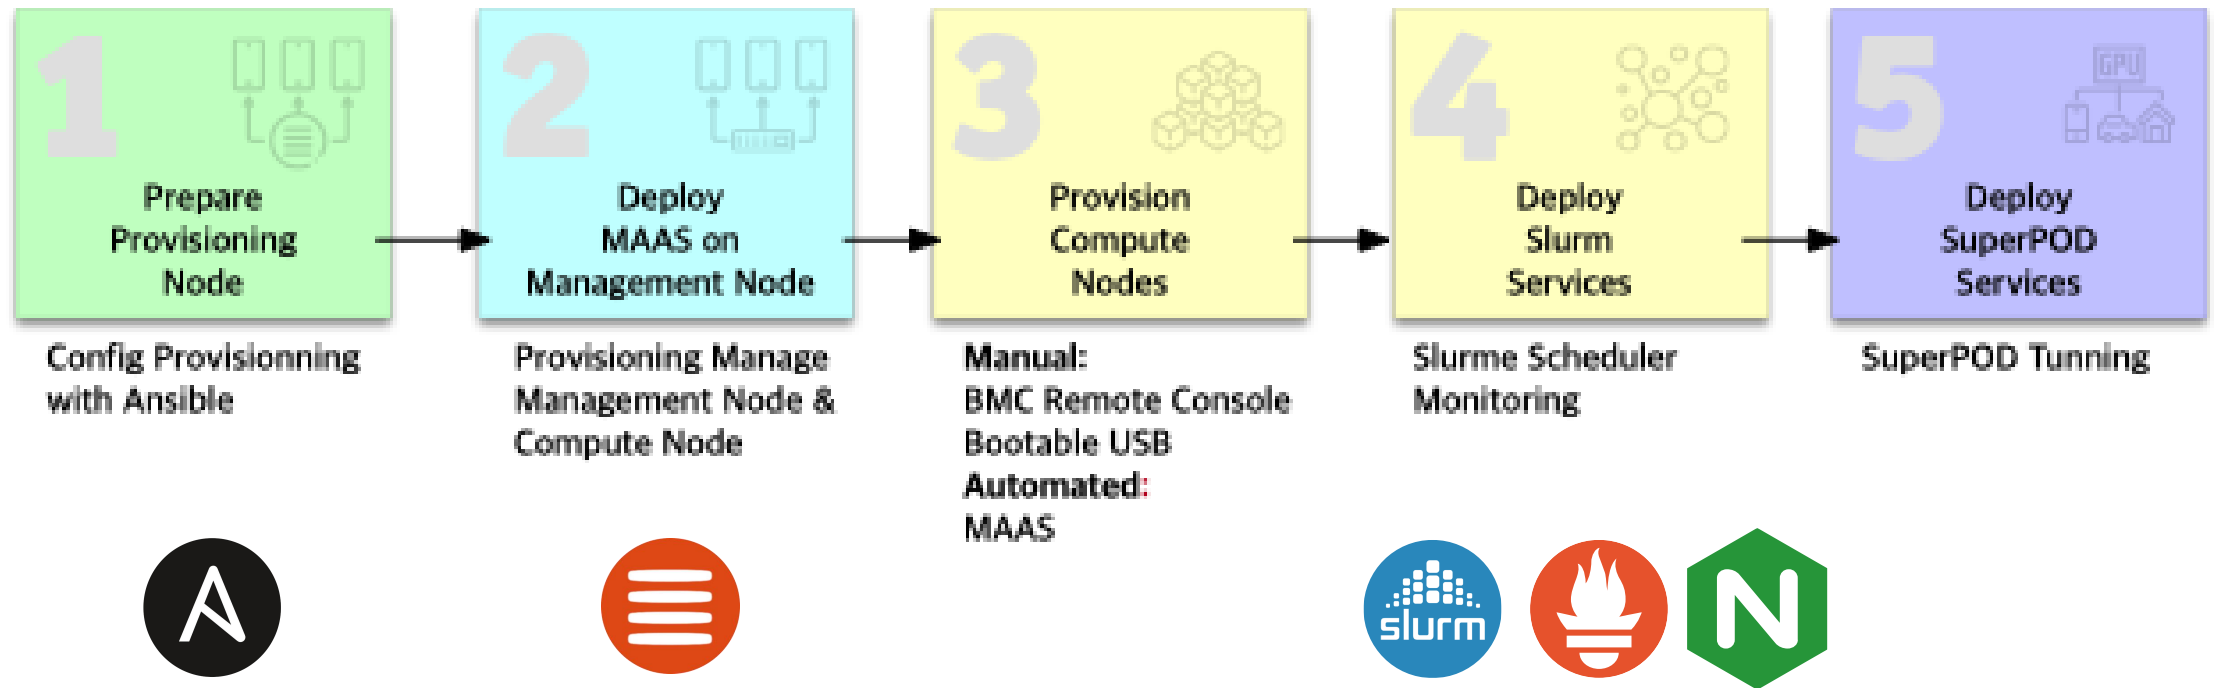
DEEPOPS DEPLOYMENT

OS DEPLYMENT

DEPLOYING SLURM



Deploying Deploying Process



Fast deployment with DeepOps

DEPLOYMENT START

Deploying DeepOps

```
# Clone the DeepOps repository
$ git clone -b 20.12 \
  https://github.com/NVIDIA/deepops

# Run the initial setup script
$ cd deepops && ./scripts/setup.sh
```

Fast deployment of Ansible provisioning platform with DeepOps

- Clone DeepOps repository
- Run the initial setup script

Docs on using DeepOps to provision your cluster:

- <https://github.com/NVIDIA/deepops>

OS DEPLOYMENT

Deploying MAAS

```
# Move Deepops Directory
$ cd ~/deepops
```

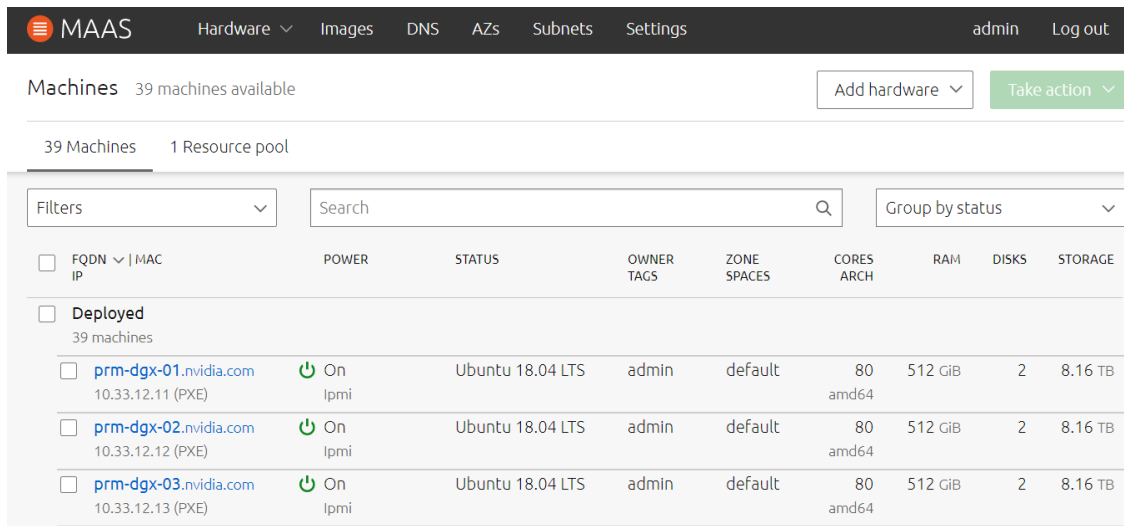
```
# Run the MAAS playbook to install on localhost
$ ansible-playbook \
playbooks/provisioning/maas.yml
```

Fast deployment of MAAS provisioning platform with DeepOps

- Run the MAAS playbook

Docs on using MAAS to provision your cluster:

- <https://github.com/NVIDIA/deepops/blob/master/docs/pxe/maas.md>
- <https://maas.io/docs>



The screenshot shows the MAAS web interface. At the top, there is a navigation bar with the MAAS logo, menu items (Hardware, Images, DNS, AZs, Subnets, Settings), and user information (admin, Log out). Below the navigation bar, there is a section for 'Machines' with '39 machines available'. There are buttons for 'Add hardware' and 'Take action'. Below this, there is a filter section with 'Filters', a search bar, and 'Group by status'. The main content is a table of machines, with a 'Deployed' filter selected, showing 39 machines. The table has columns for FQDN | MAC IP, POWER, STATUS, OWNER TAGS, ZONE SPACES, CORES ARCH, RAM, DISKS, and STORAGE. Three machines are visible in the table:

FQDN MAC IP	POWER	STATUS	OWNER TAGS	ZONE SPACES	CORES ARCH	RAM	DISKS	STORAGE
<input type="checkbox"/> prn-dgx-01.nvidia.com 10.33.12.11 (PXE)	<input type="checkbox"/> On Ipmi	Ubuntu 18.04 LTS	admin	default	80 amd64	512 GIB	2	8.16 TB
<input type="checkbox"/> prn-dgx-02.nvidia.com 10.33.12.12 (PXE)	<input type="checkbox"/> On Ipmi	Ubuntu 18.04 LTS	admin	default	80 amd64	512 GIB	2	8.16 TB
<input type="checkbox"/> prn-dgx-03.nvidia.com 10.33.12.13 (PXE)	<input type="checkbox"/> On Ipmi	Ubuntu 18.04 LTS	admin	default	80 amd64	512 GIB	2	8.16 TB

Job Scheduler DEPLOYMENT

Deploying Slurm

```
# Move Deepops Directory
$ cd ~/deepops

# Edit the Ansible inventory file to list nodes in the cluster
$ nano config/inventory

# (Optional) Change configuration of the Slurm cluster
$ nano config/group_vars/slurm-cluster.yml

# Deploy Slurm cluster
$ ansible-playbook -l slurm-cluster playbooks/slurm-cluster.yml
```

Slurm deployment guide:

<https://github.com/NVIDIA/deepops/tree/master/docs/slurm-cluster>

RESOURCES

Learn more and try

NVIDIA DeepOps: <https://github.com/nvidia/deepops>

DGX POD reference architectures: <https://www.nvidia.com/en-us/data-center/dgx-pod/>

Slurm documentation: <https://slurm.schedmd.com>



감사합니다.



BayNex