Cloud-Native Day
2020 Korea
LIVE!

Manage | 멀티 클라우드 환경의 통합 관리와 분산된 워크로드 서비스 메시

황주필 (jupilh@vmware.com)

vmware

Cloud-Native Day
2020 Korea
LIVE!

Today's Sessions

1. 마이크로서비스 디자인
2. 마이크로서비스 구축
3. 마이크로서비스 런타임 & 네트워크
4. **마이크로서비스 모니터링 및 관리**

vmware

마이크로서비스 모니터링 및 관리

- 로깅 모니터링은 어떻게 하나?
- 서비스메시로 관리하고 싶은데 멀티 클라우드 환경이라 가능할까?
- 여러개의 Kubernetes 클러스터 버전 및 접근 Policy등은 어떻게 관리해야 할까?

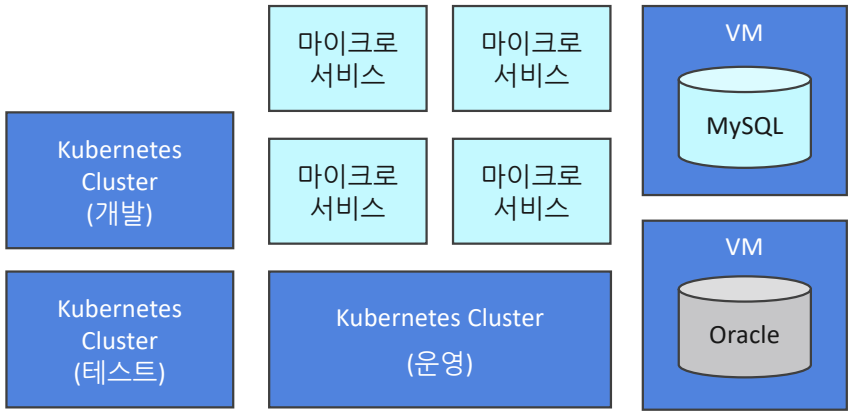

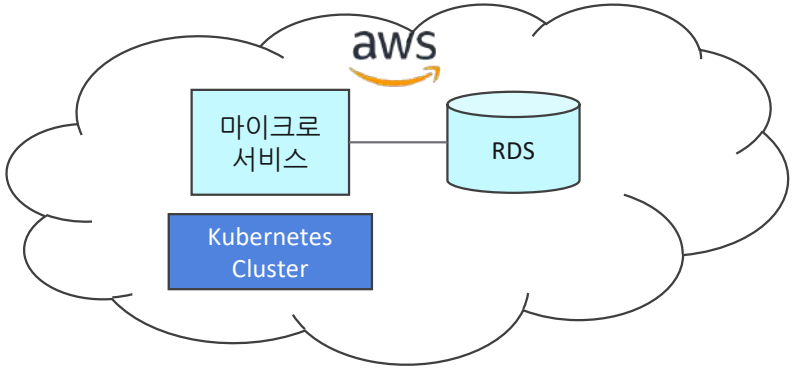

여러 클라우드 Kubernetes 클러스터 환경

기업의 Kubernetes 환경의 기술적 인 엔트로피는 응용 프로그램과 함께 지속적으로 증가

vmware

모니터링 & 관리

무엇을 집중 관리해야 할까?

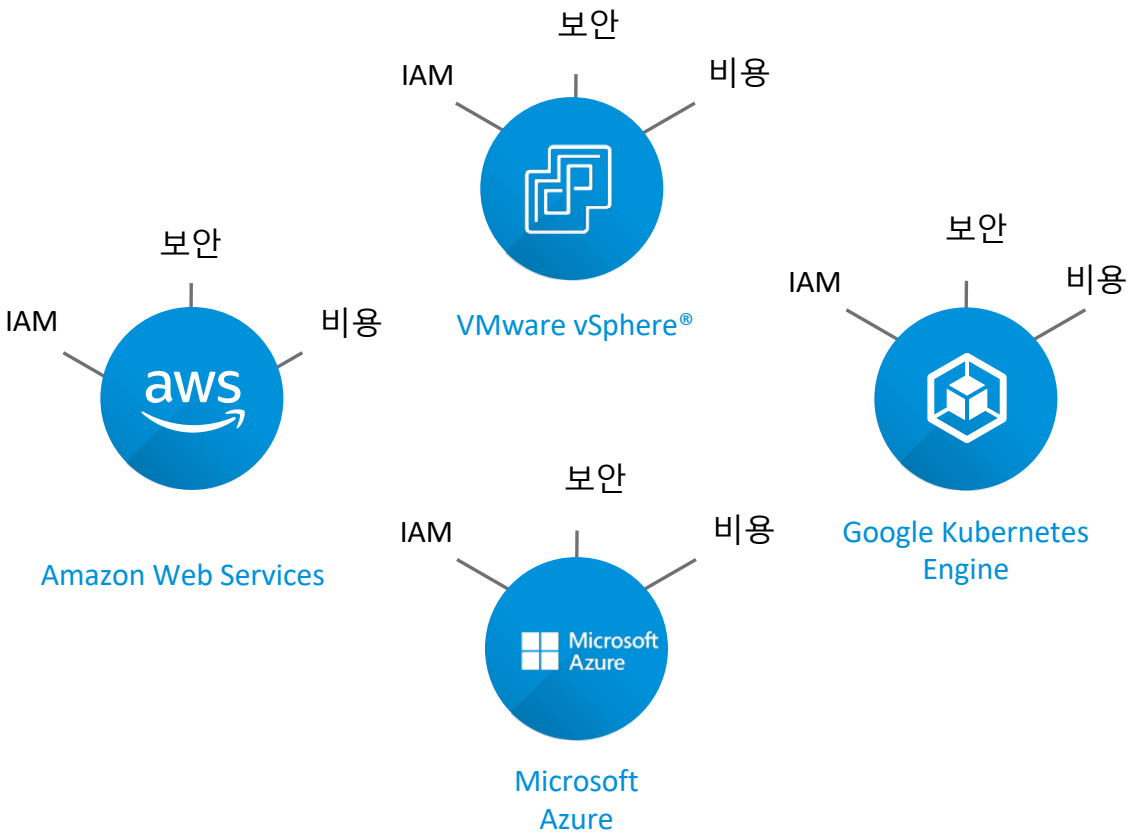

공통 관심사
(Cross-Cutting Concern)

반드시 지원해야 하지만 개별적으로 대응 및 관리하기에 비용이 많이 들고 비효율적이고 스케일하지 않은 것

- 환경마다 다른 기술을 이용한 수동 구성 관리
- Kubernetes 클러스터에 대한 별도의 정책 적용 (엑세스, 네트워크, 보안, etc)
- 클라우드 비용의 시각화 및 최적화
- 기술과 도구의 학습

VMware Tanzu 포트폴리오

Kubernetes 를 중심으로 하는 현대 애플리케이션을 위한 제품 및 서비스

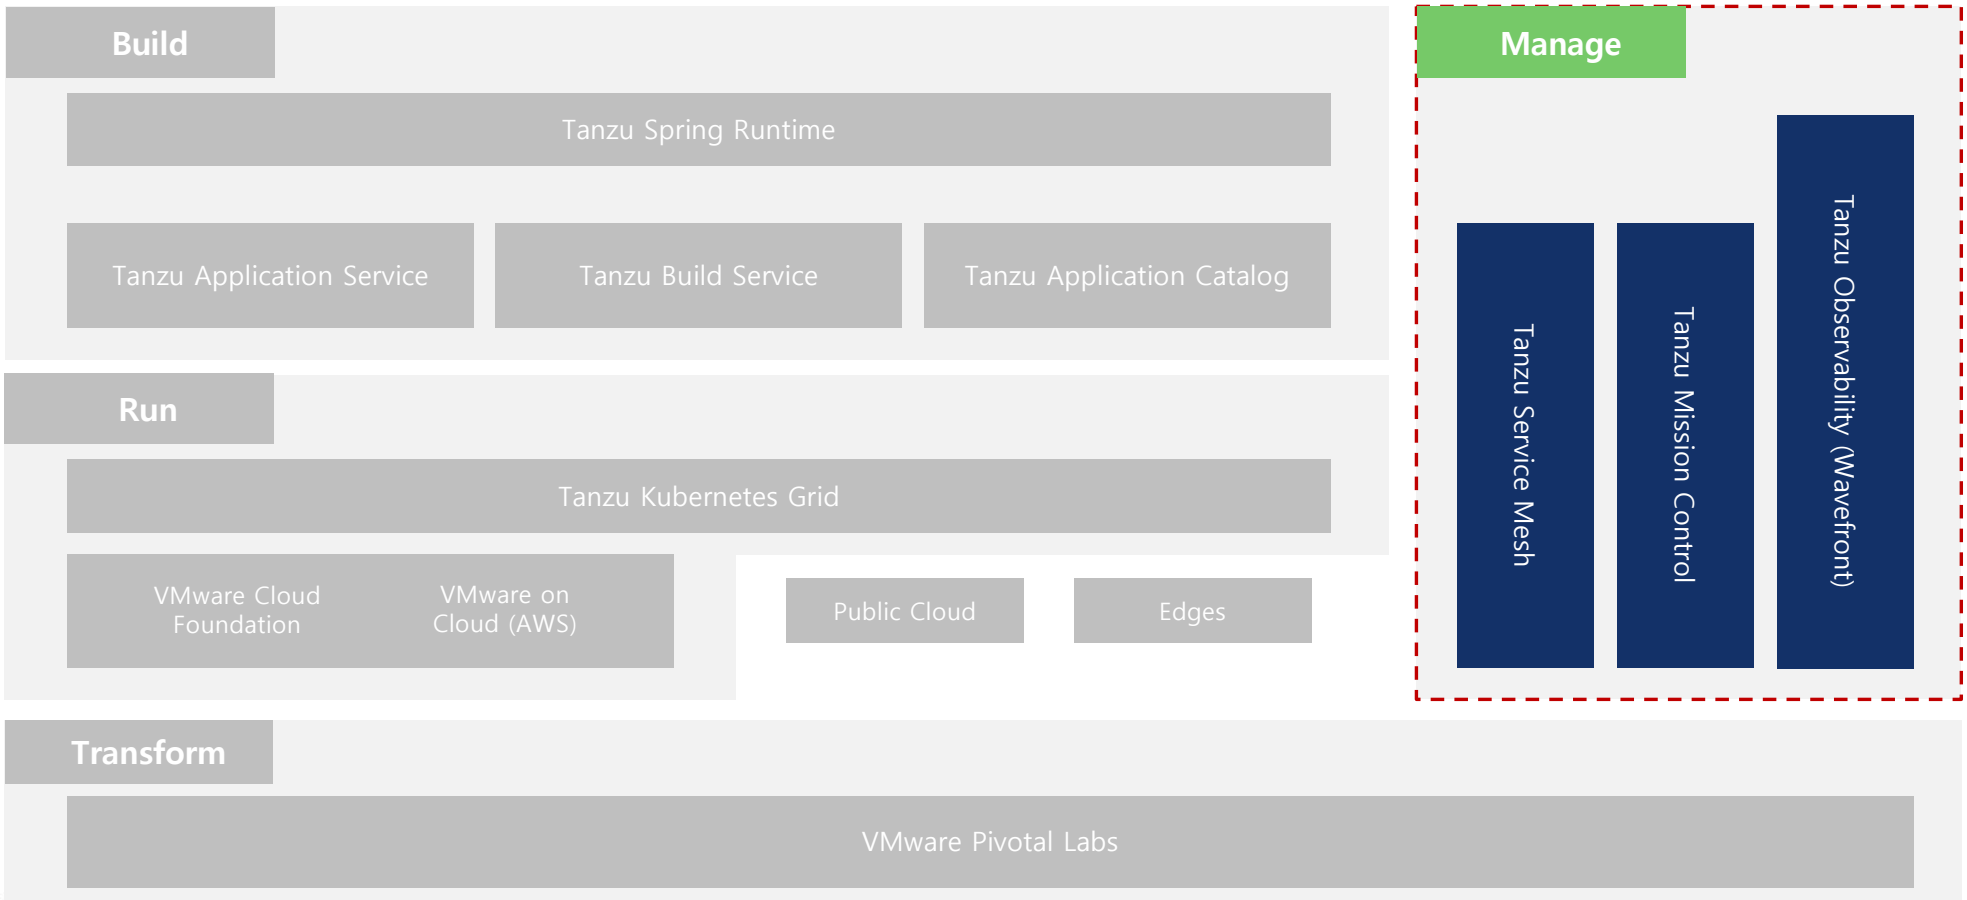

Cloud-Native Day
2020 Korea
LIVE!

멀티 클라우드, 멀티 플랫폼 – 통합 관리

Tanzu Mission Control

vmware

Kubernetes 환경 통합 관리

멀티 클라우드 및 멀티 팀에서 사용되는 Kubernetes 클러스터에 대해 통합된 중앙집중식 프로비저닝, 정책 기반 관리 및 운영

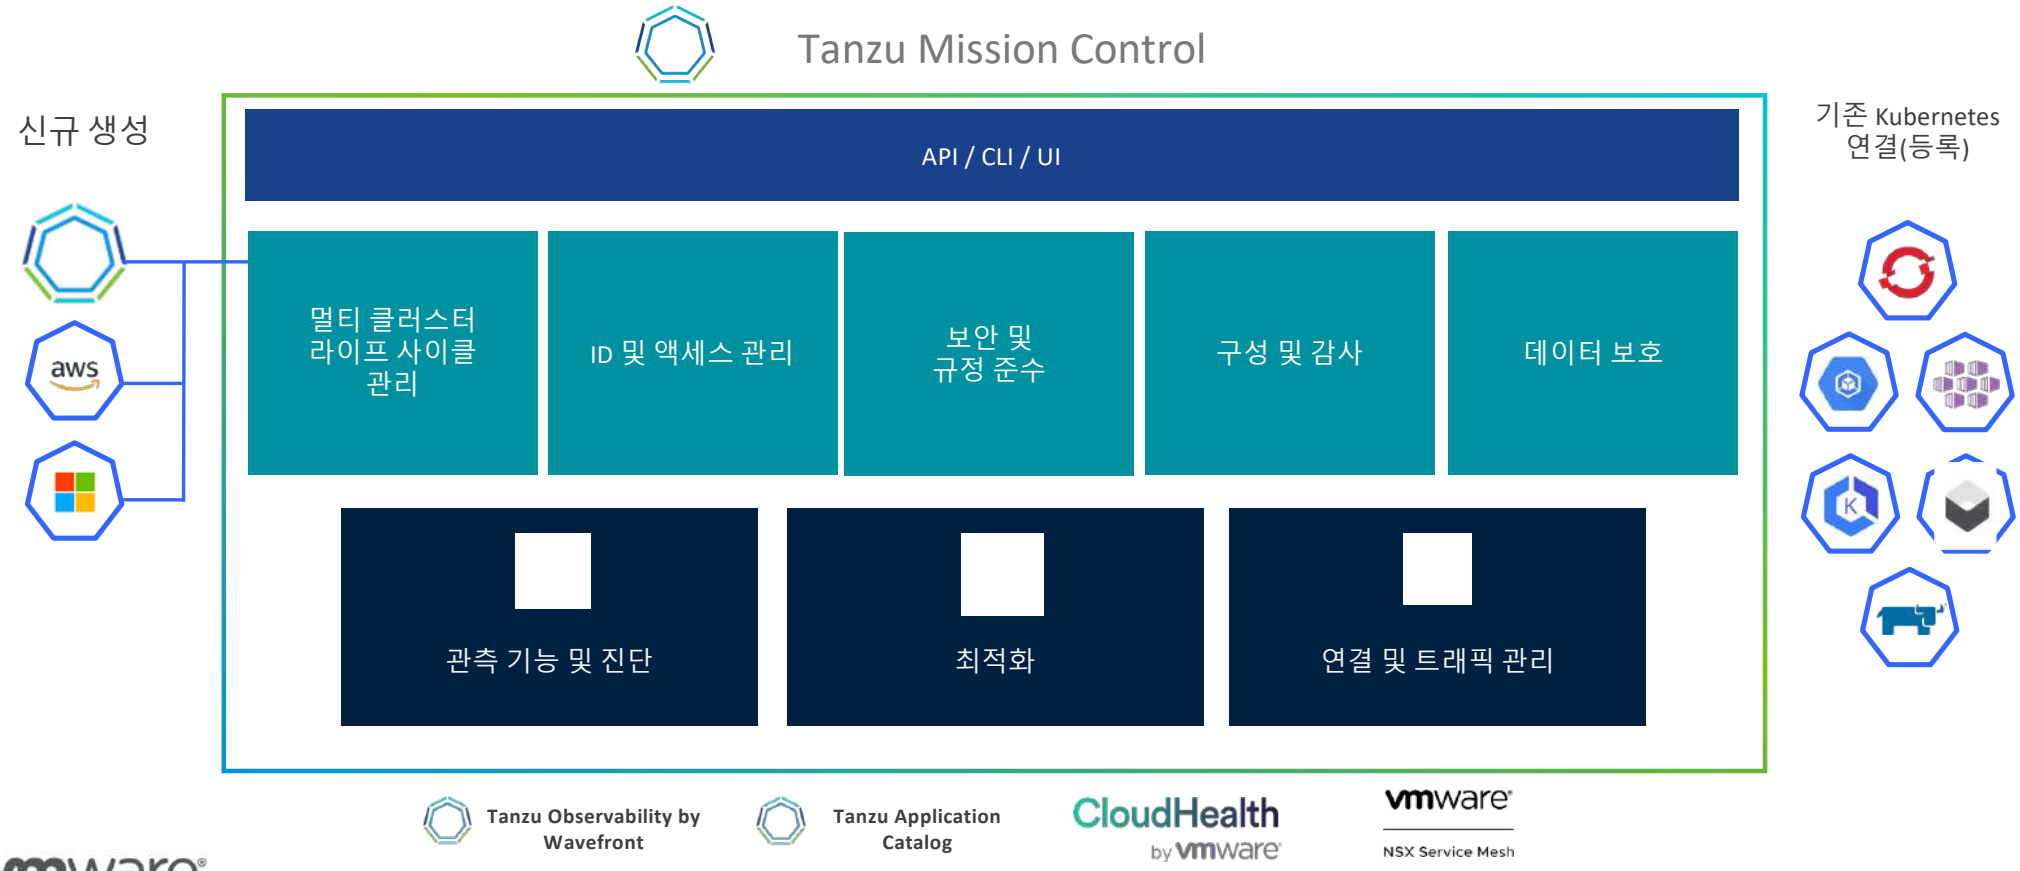

vmware

Tanzu Mission Control 아키텍처

클러스터 그룹과 애플리케이션 워크스페이스

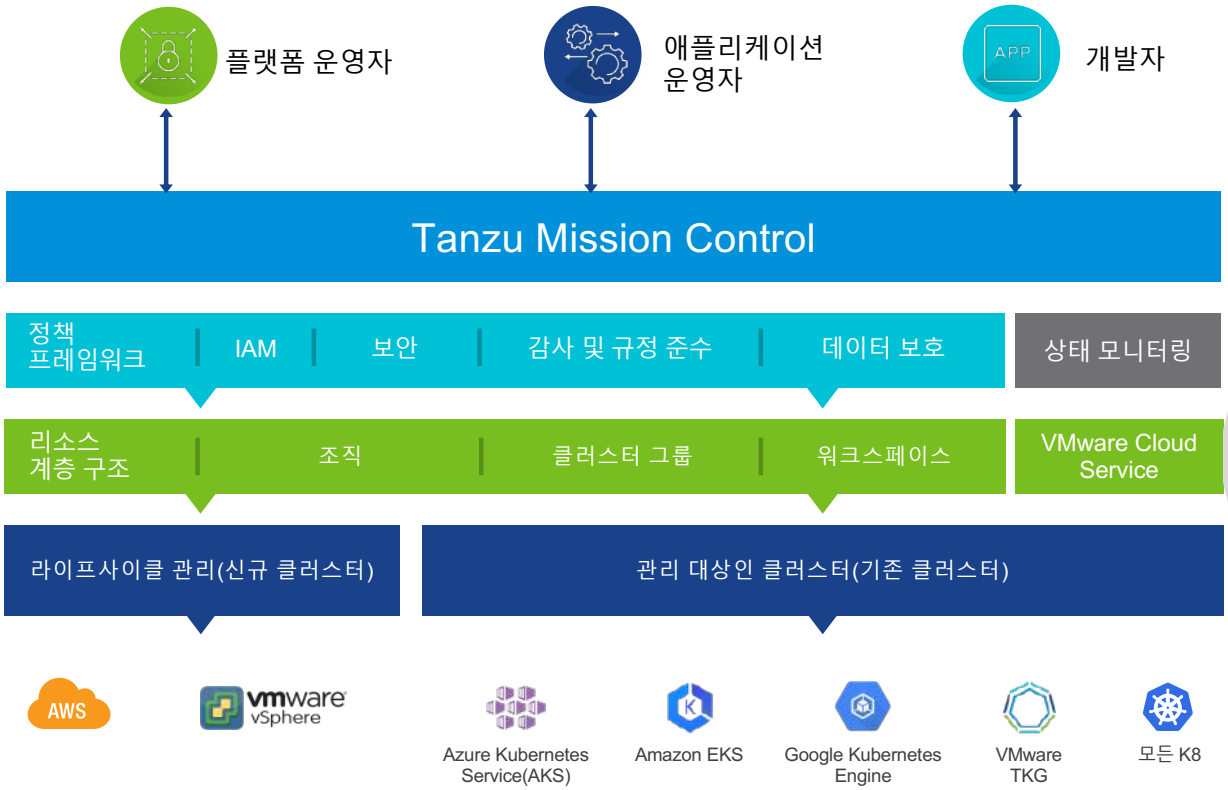

리소스 계층 구조

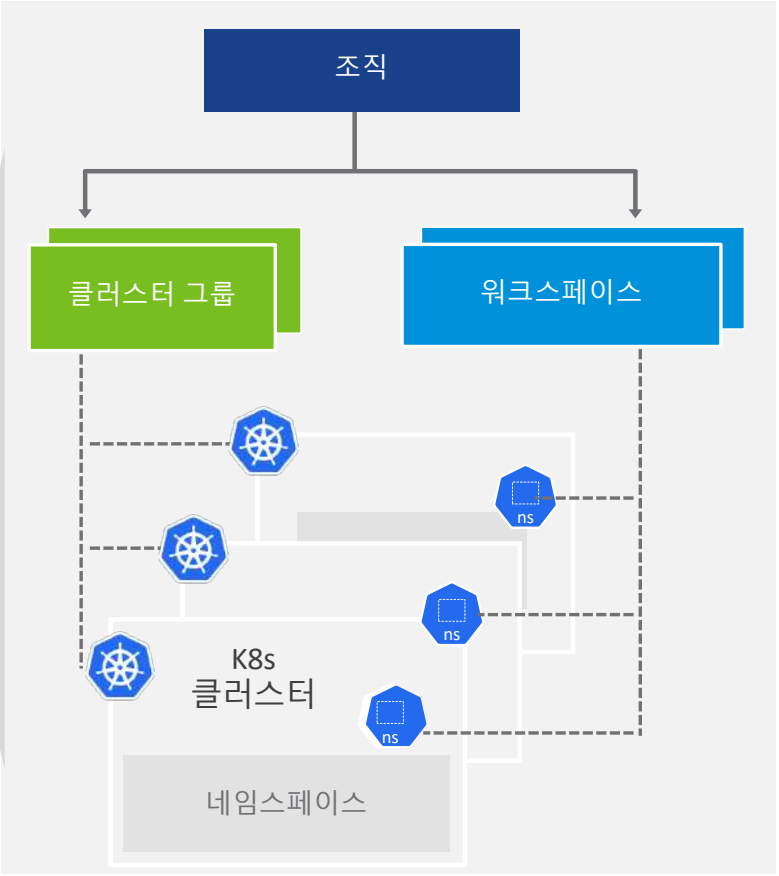

라이프사이클 관리

CNCF 준수하는 모든 K8s 클러스터를 연결하고 관리 가능

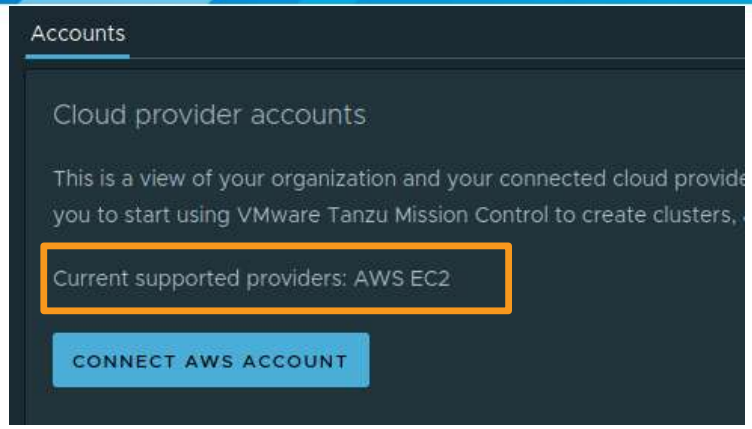

(2020/05 시점: vSphere 를 포함하여 타 클라우드는 향후 추가 예정)

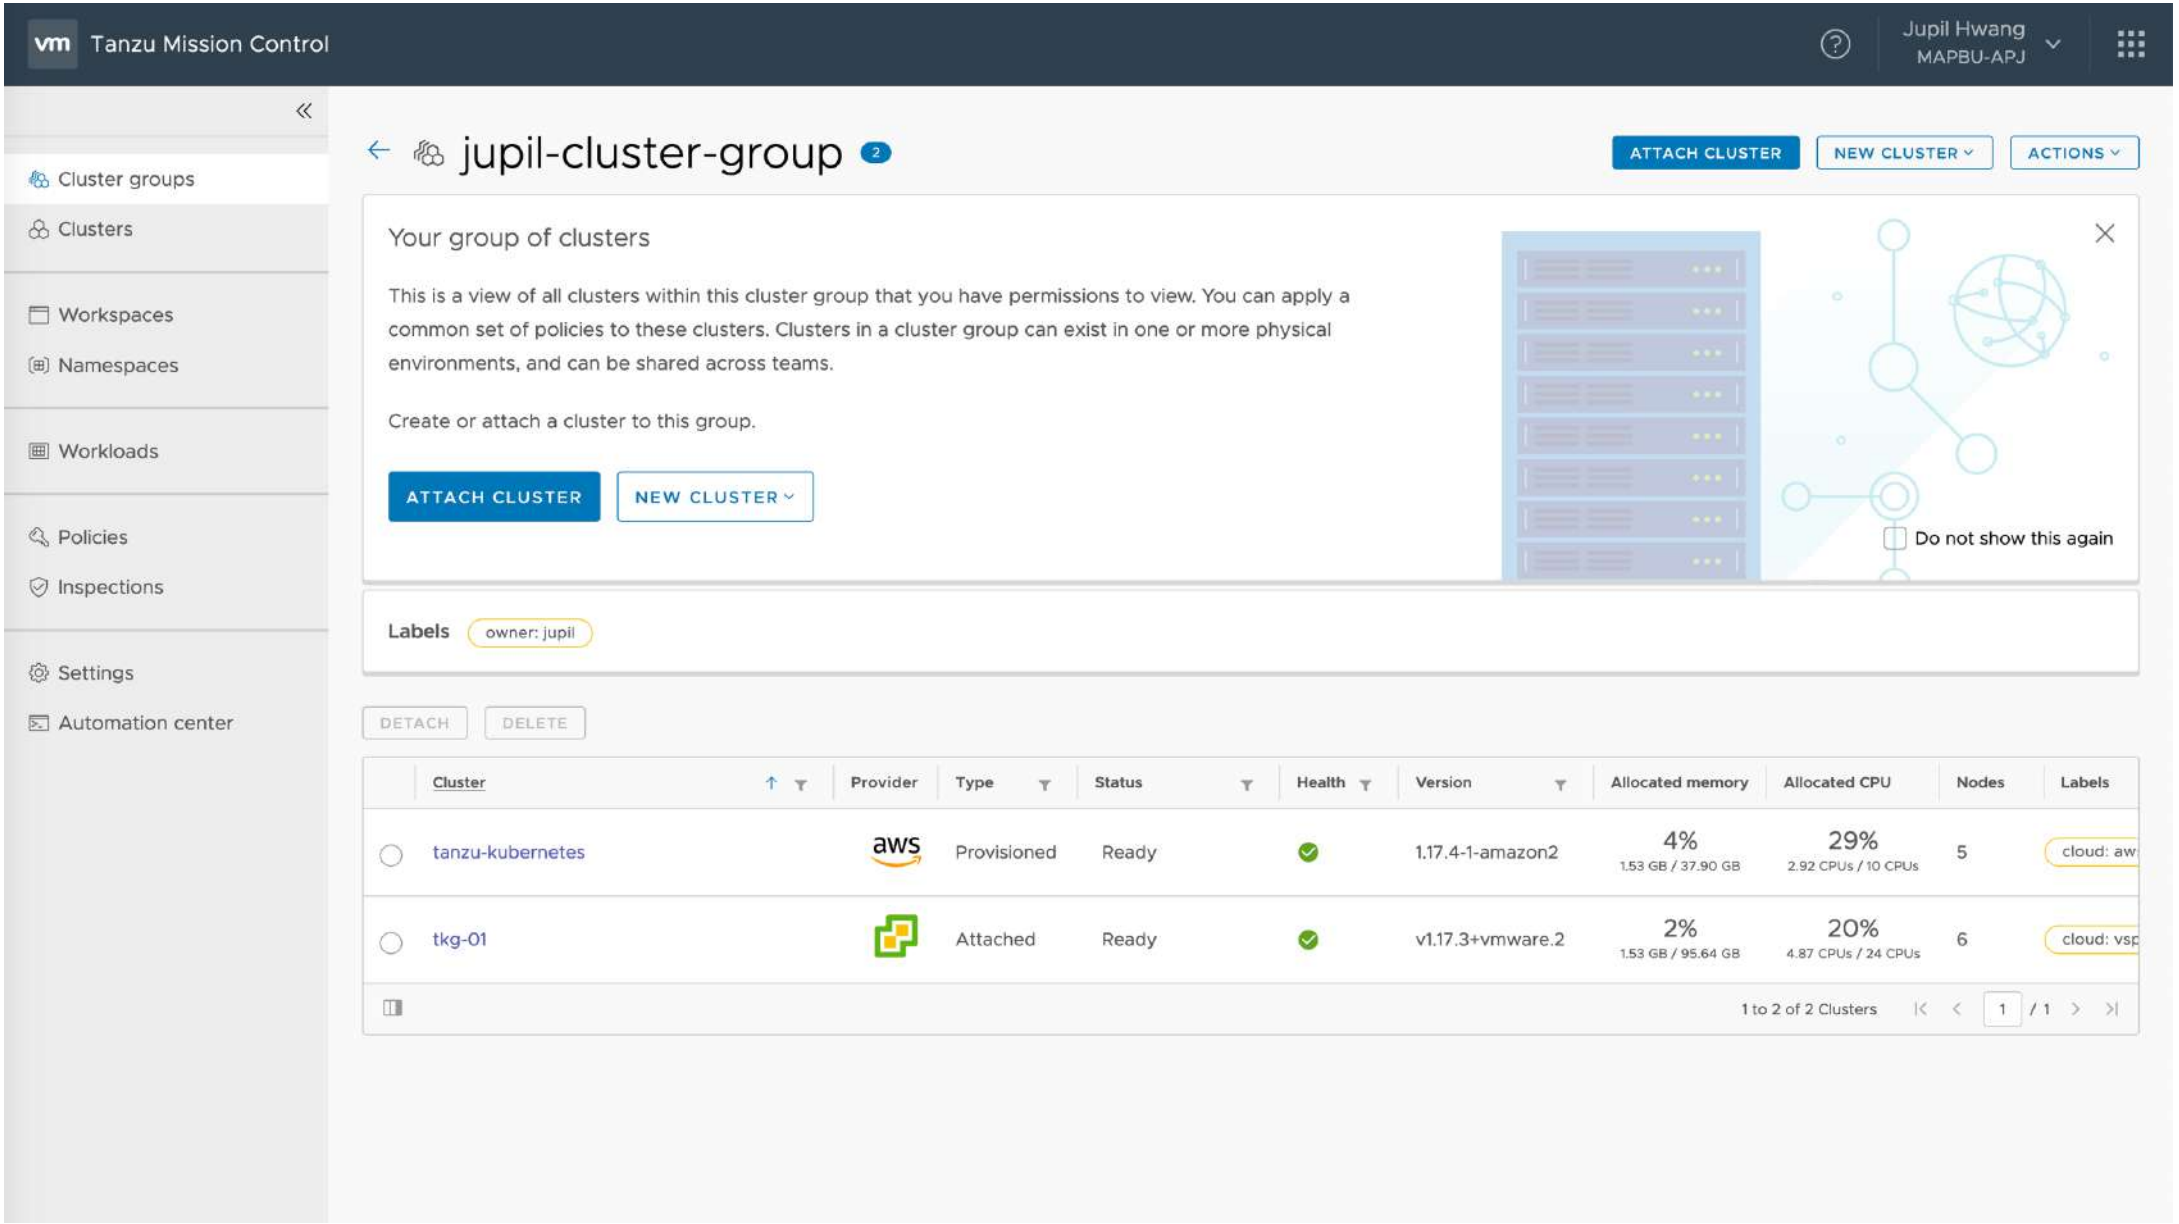
Tanzu Mission Control
Jupil Hwang
MAPBU-APJ
Cluster groups
Clusters
Workspaces
Namespaces
Workloads
Policies
Inspections
Settings
Automation center
jupil-cluster-group
2
ATTACH CLUSTER
NEW CLUSTER
ACTIONS
Your group of clusters
This is a view of all clusters within this cluster group that you have permissions to view. You can apply a common set of policies to these clusters. Clusters in a cluster group can exist in one or more physical environments, and can be shared across teams.
Create or attach a cluster to this group.
ATTACH CLUSTER
NEW CLUSTER
Do not show this again
Labels
owner: jupil
DETACH
DELETE
Cluster
Provider
Type
Status
Health
Version
Allocated memory
Allocated CPU
Nodes
Labels
tanzu-kubernetes
aws
Provisioned
Ready
1.17.4-1-amazon2
4%
1.53 GB / 37.90 GB
29%
2.92 CPUs / 10 CPUs
5
cloud: aw
tkg-01
Attached
Ready
v1.17.3+vmware.2
2%
1.53 GB / 95.64 GB
20%
4.87 CPUs / 24 CPUs
6
cloud: vsp
1 to 2 of 2 Clusters
1
/ 1

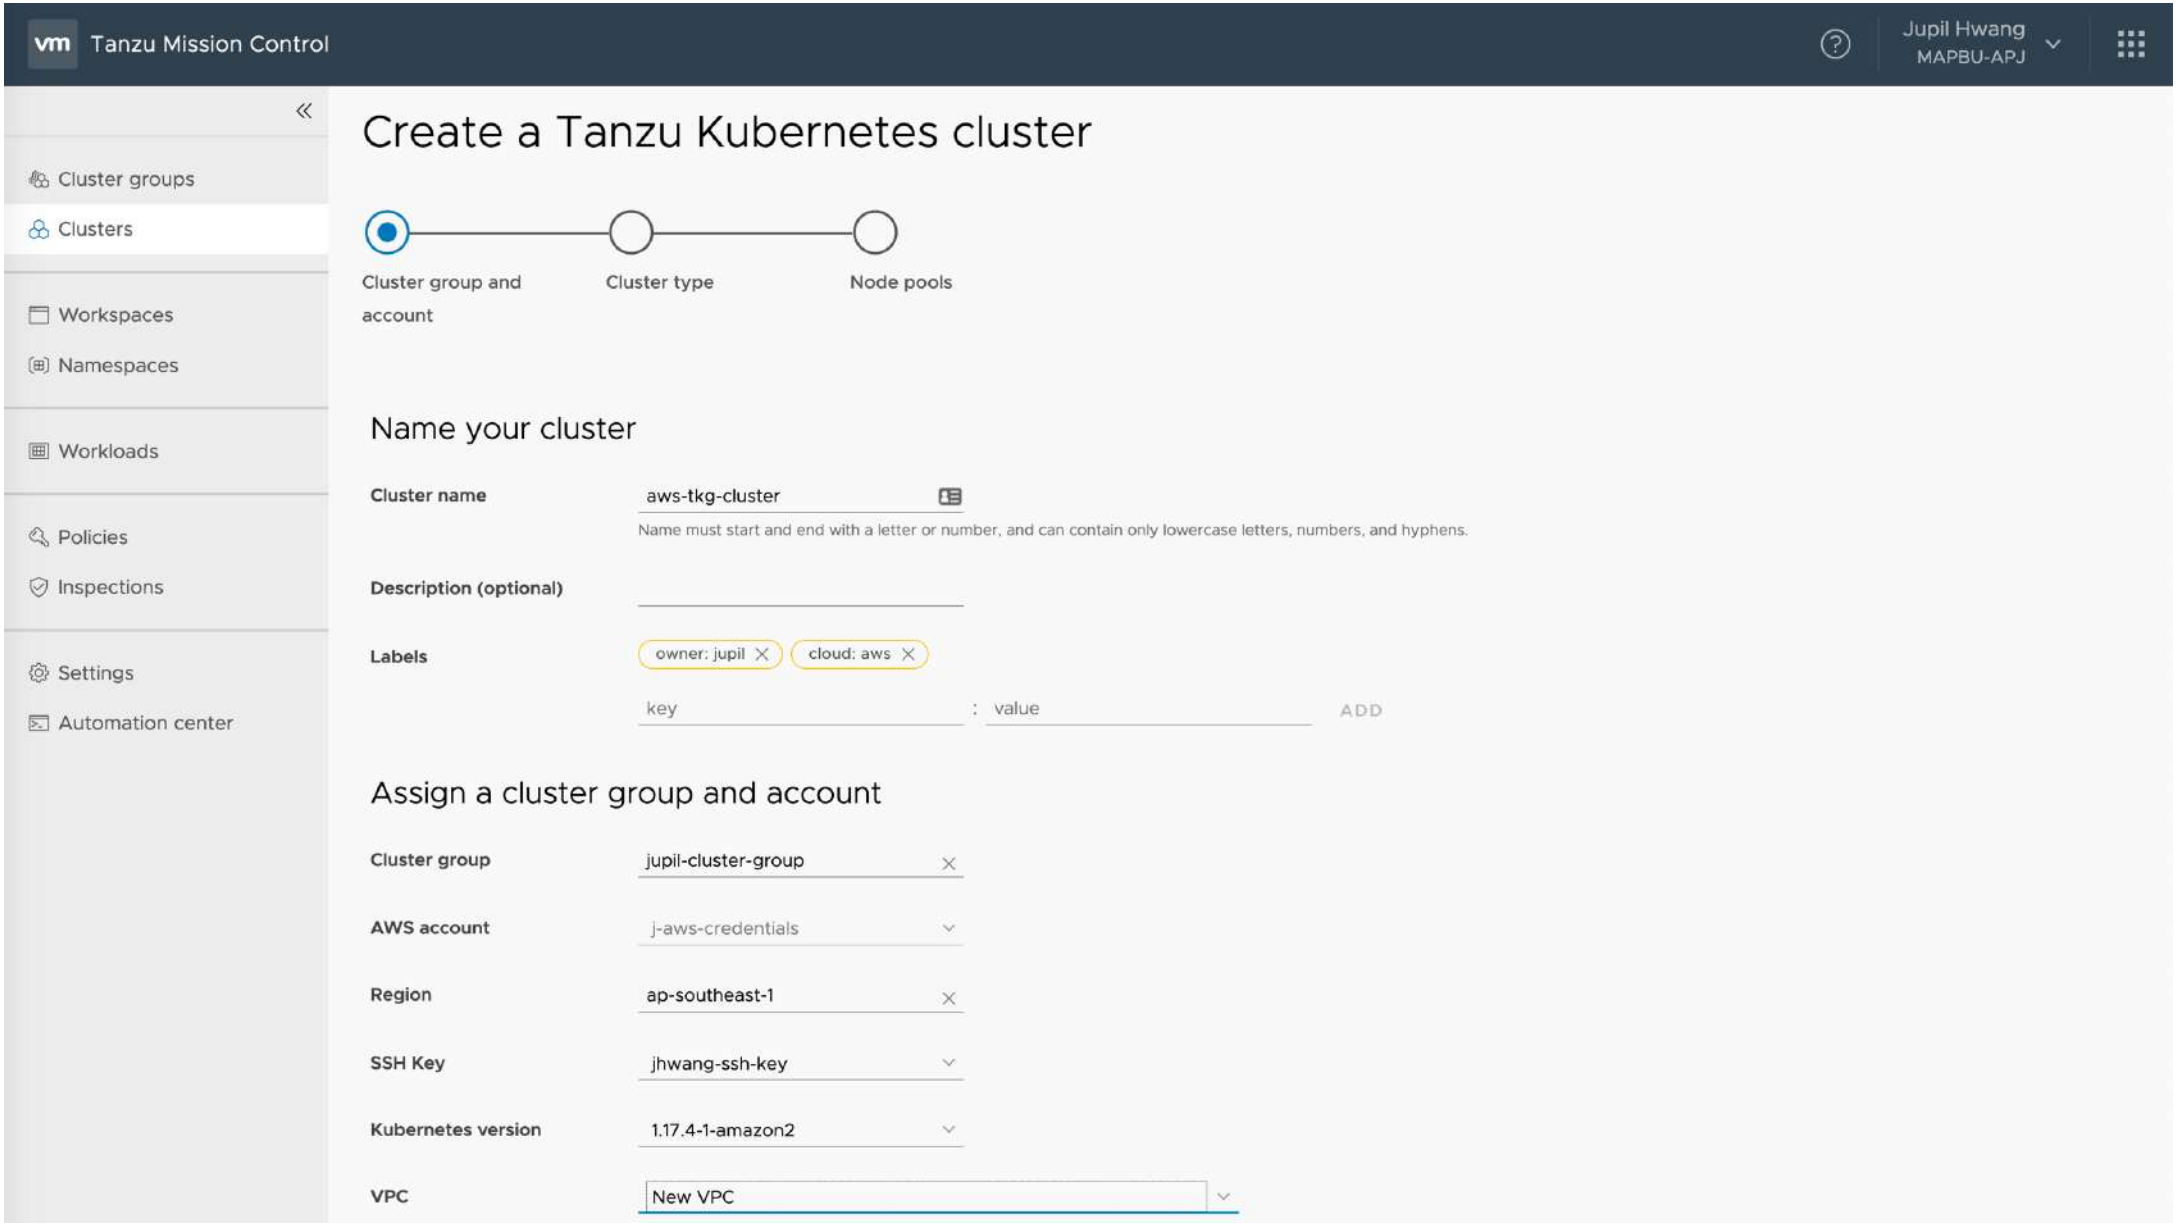
Tanzu Mission Control
Jupil Hwang
MAPBU-APJ
Cluster groups
Clusters
Workspaces
Namespaces
Workloads
Policies
Inspections
Settings
Automation center
Create a Tanzu Kubernetes cluster
Cluster group and account
Cluster type
Node pools
Name your cluster
Cluster name
aws-tkg-cluster
Name must start and end with a letter or number, and can contain only lowercase letters, numbers, and hyphens.
Description (optional)
Labels
owner: jupil
cloud: aws
key
value
ADD
Assign a cluster group and account
Cluster group
jupil-cluster-group
AWS account
j-aws-credentials
Region
ap-southeast-1
SSH Key
jhwang-ssh-key
Kubernetes version
1.17.4-1-amazon2
VPC
New VPC

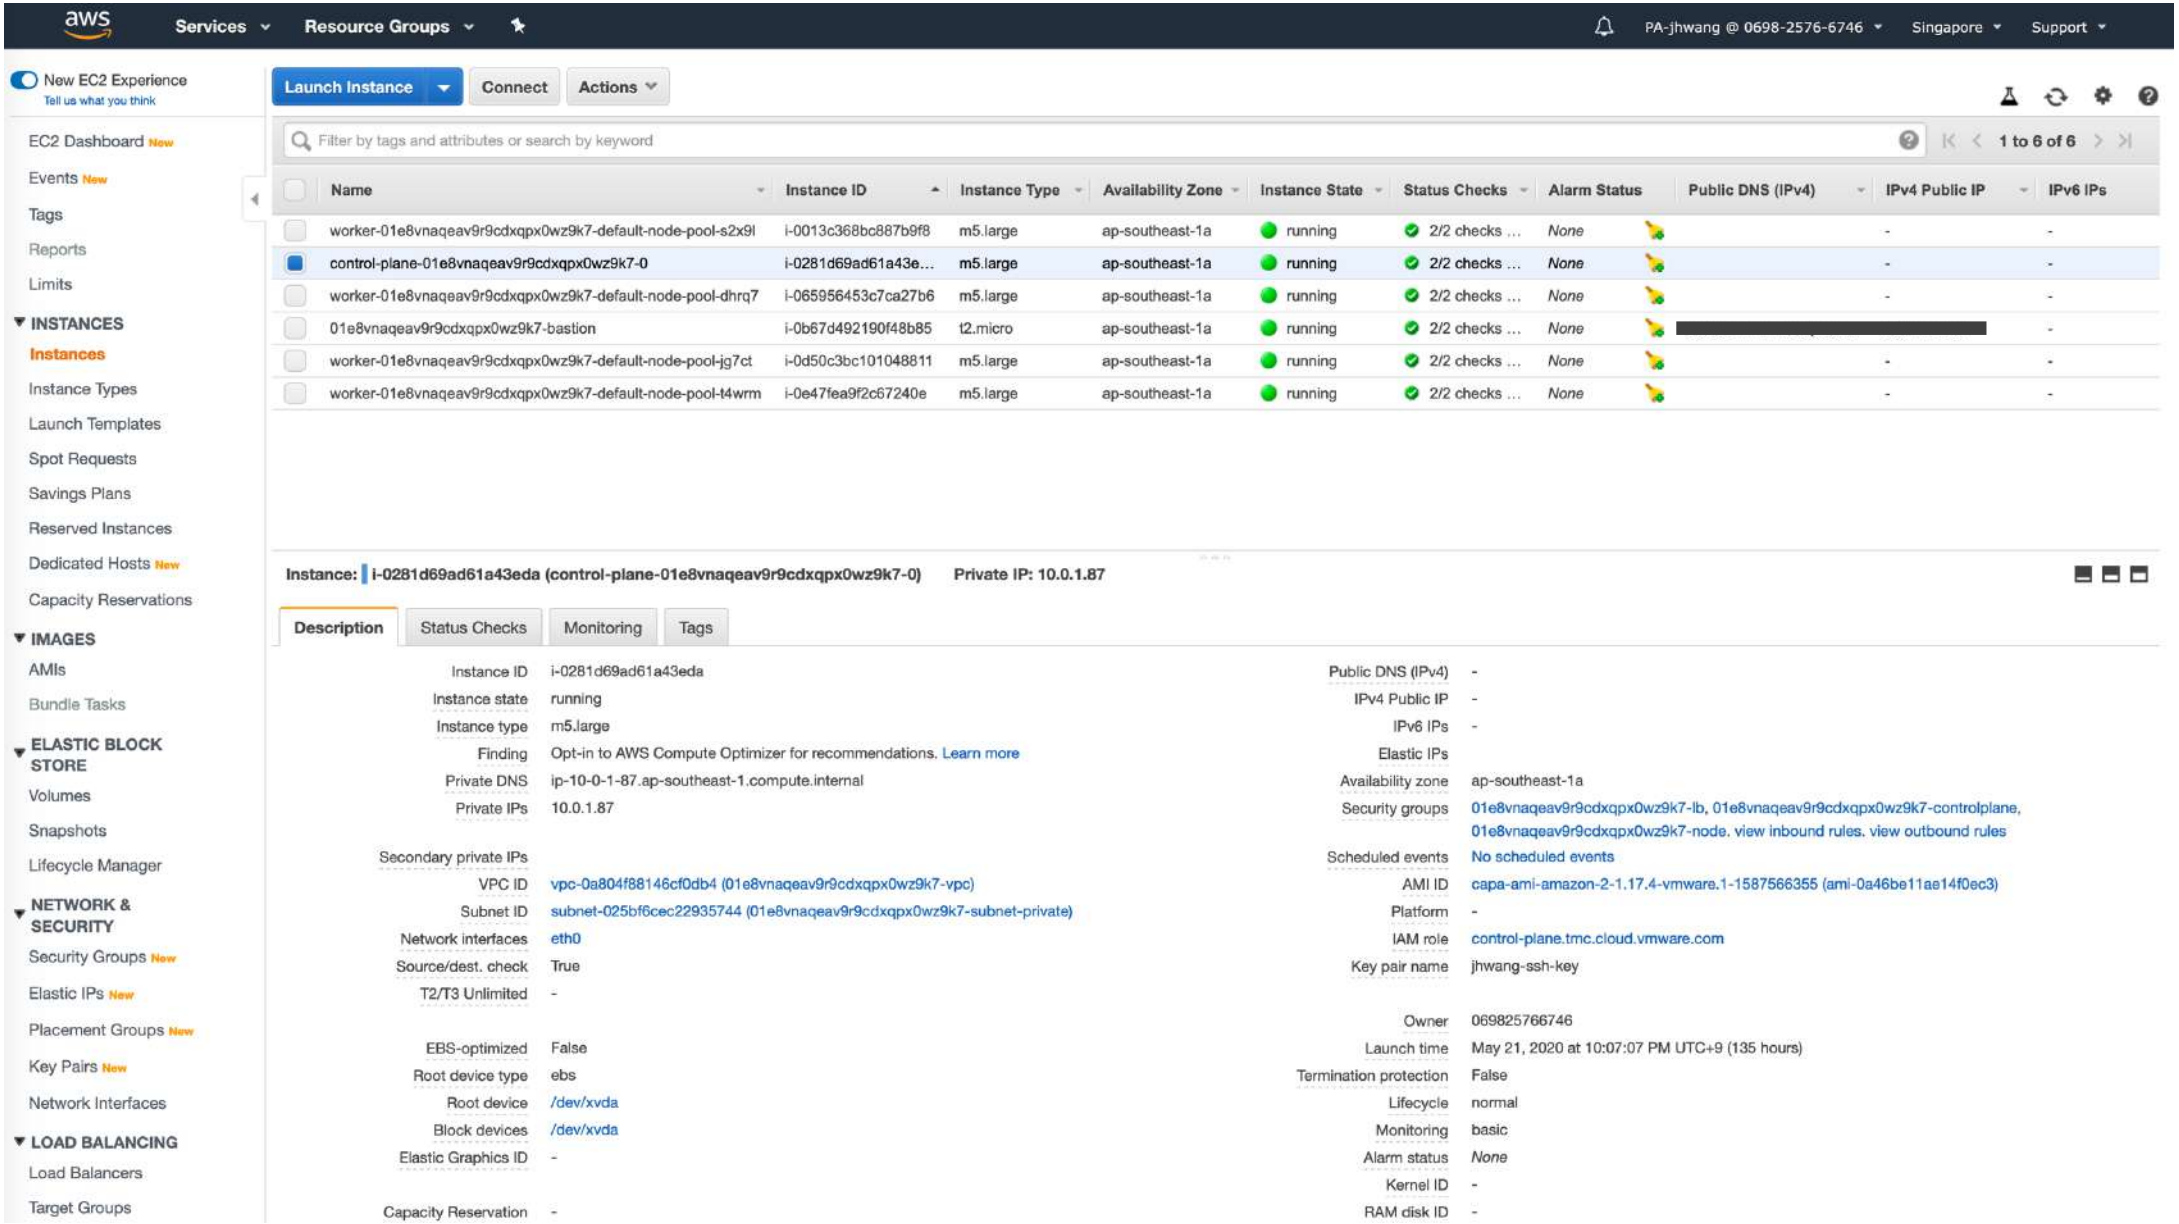
aws
Services
Resource Groups
PA-jhwang @ 0698-2576-6746
Singapore
Support
New EC2 Experience
Tell us what you think
EC2 Dashboard New
Events New
Tags
Reports
Limits
INSTANCES
Instances
Instance Types
Launch Templates
Spot Requests
Savings Plans
Reserved Instances
Dedicated Hosts New
Capacity Reservations
IMAGES
AMIs
Bundle Tasks
ELASTIC BLOCK STORE
Volumes
Snapshots
Lifecycle Manager
NETWORK & SECURITY
Security Groups New
Elastic IPs New
Placement Groups New
Key Pairs New
Network Interfaces
LOAD BALANCING
Load Balancers
Target Groups
Launch Instance
Connect
Actions
Filter by tags and attributes or search by keyword
1 to 6 of 6
Name
Instance ID
Instance Type
Availability Zone
Instance State
Status Checks
Alarm Status
Public DNS (IPv4)
IPv4 Public IP
IPv6 IPs
worker-01e8vnaqeav9r9cdxqpx0wz9k7-default-node-pool-s2x9l i-0013c368bc887b9f8 m5.large ap-southeast-1a running 2/2 checks ... None
control-plane-01e8vnaqeav9r9cdxqpx0wz9k7-0 i-0281d69ad61a43e... m5.large ap-southeast-1a running 2/2 checks ... None
worker-01e8vnaqeav9r9cdxqpx0wz9k7-default-node-pool-dhrq7 i-065956453c7ca27b6 m5.large ap-southeast-1a running 2/2 checks ... None
01e8vnaqeav9r9cdxqpx0wz9k7-bastion i-0b67d492190f48b85 t2.micro ap-southeast-1a running 2/2 checks ... None
worker-01e8vnaqeav9r9cdxqpx0wz9k7-default-node-pool-jg7ct i-0d50c3bc101048811 m5.large ap-southeast-1a running 2/2 checks ... None
worker-01e8vnaqeav9r9cdxqpx0wz9k7-default-node-pool-t4wrm i-0e47fea9f2c67240e m5.large ap-southeast-1a running 2/2 checks ... None
Instance: i-0281d69ad61a43eda (control-plane-01e8vnaqeav9r9cdxqpx0wz9k7-0) Private IP: 10.0.1.87
Description
Status Checks
Monitoring
Tags
Instance ID i-0281d69ad61a43eda
Instance state running
Instance type m5.large
Finding Opt-in to AWS Compute Optimizer for recommendations. Learn more
Private DNS ip-10-0-1-87.ap-southeast-1.compute.internal
Private IPs 10.0.1.87
Secondary private IPs
VPC ID vpc-0a804f88146cf0db4 (01e8vnaqeav9r9cdxqpx0wz9k7-vpc)
Subnet ID subnet-025bf6cec22935744 (01e8vnaqeav9r9cdxqpx0wz9k7-subnet-private)
Network interfaces eth0
Source/dest. check True
T2/T3 Unlimited -
EBS-optimized False
Root device type ebs
Root device /dev/xvda
Block devices /dev/xvda
Elastic Graphics ID -
Capacity Reservation -
Public DNS (IPv4) -
IPv4 Public IP -
IPv6 IPs -
Elastic IPs
Availability zone ap-southeast-1a
Security groups 01e8vnaqeav9r9cdxqpx0wz9k7-lb, 01e8vnaqeav9r9cdxqpx0wz9k7-controlplane, 01e8vnaqeav9r9cdxqpx0wz9k7-node. view inbound rules. view outbound rules
Scheduled events No scheduled events
AMI ID capa-ami-amazon-2-1.17.4-vmware.1-1587566355 (ami-0a46be11ae14f0ec3)
Platform -
IAM role control-plane.tmc.cloud.vmware.com
Key pair name jhwang-ssh-key
Owner 069825766746
Launch time May 21, 2020 at 10:07:07 PM UTC+9 (135 hours)
Termination protection False
Lifecycle normal
Monitoring basic
Alarm status None
Kernel ID -
RAM disk ID -

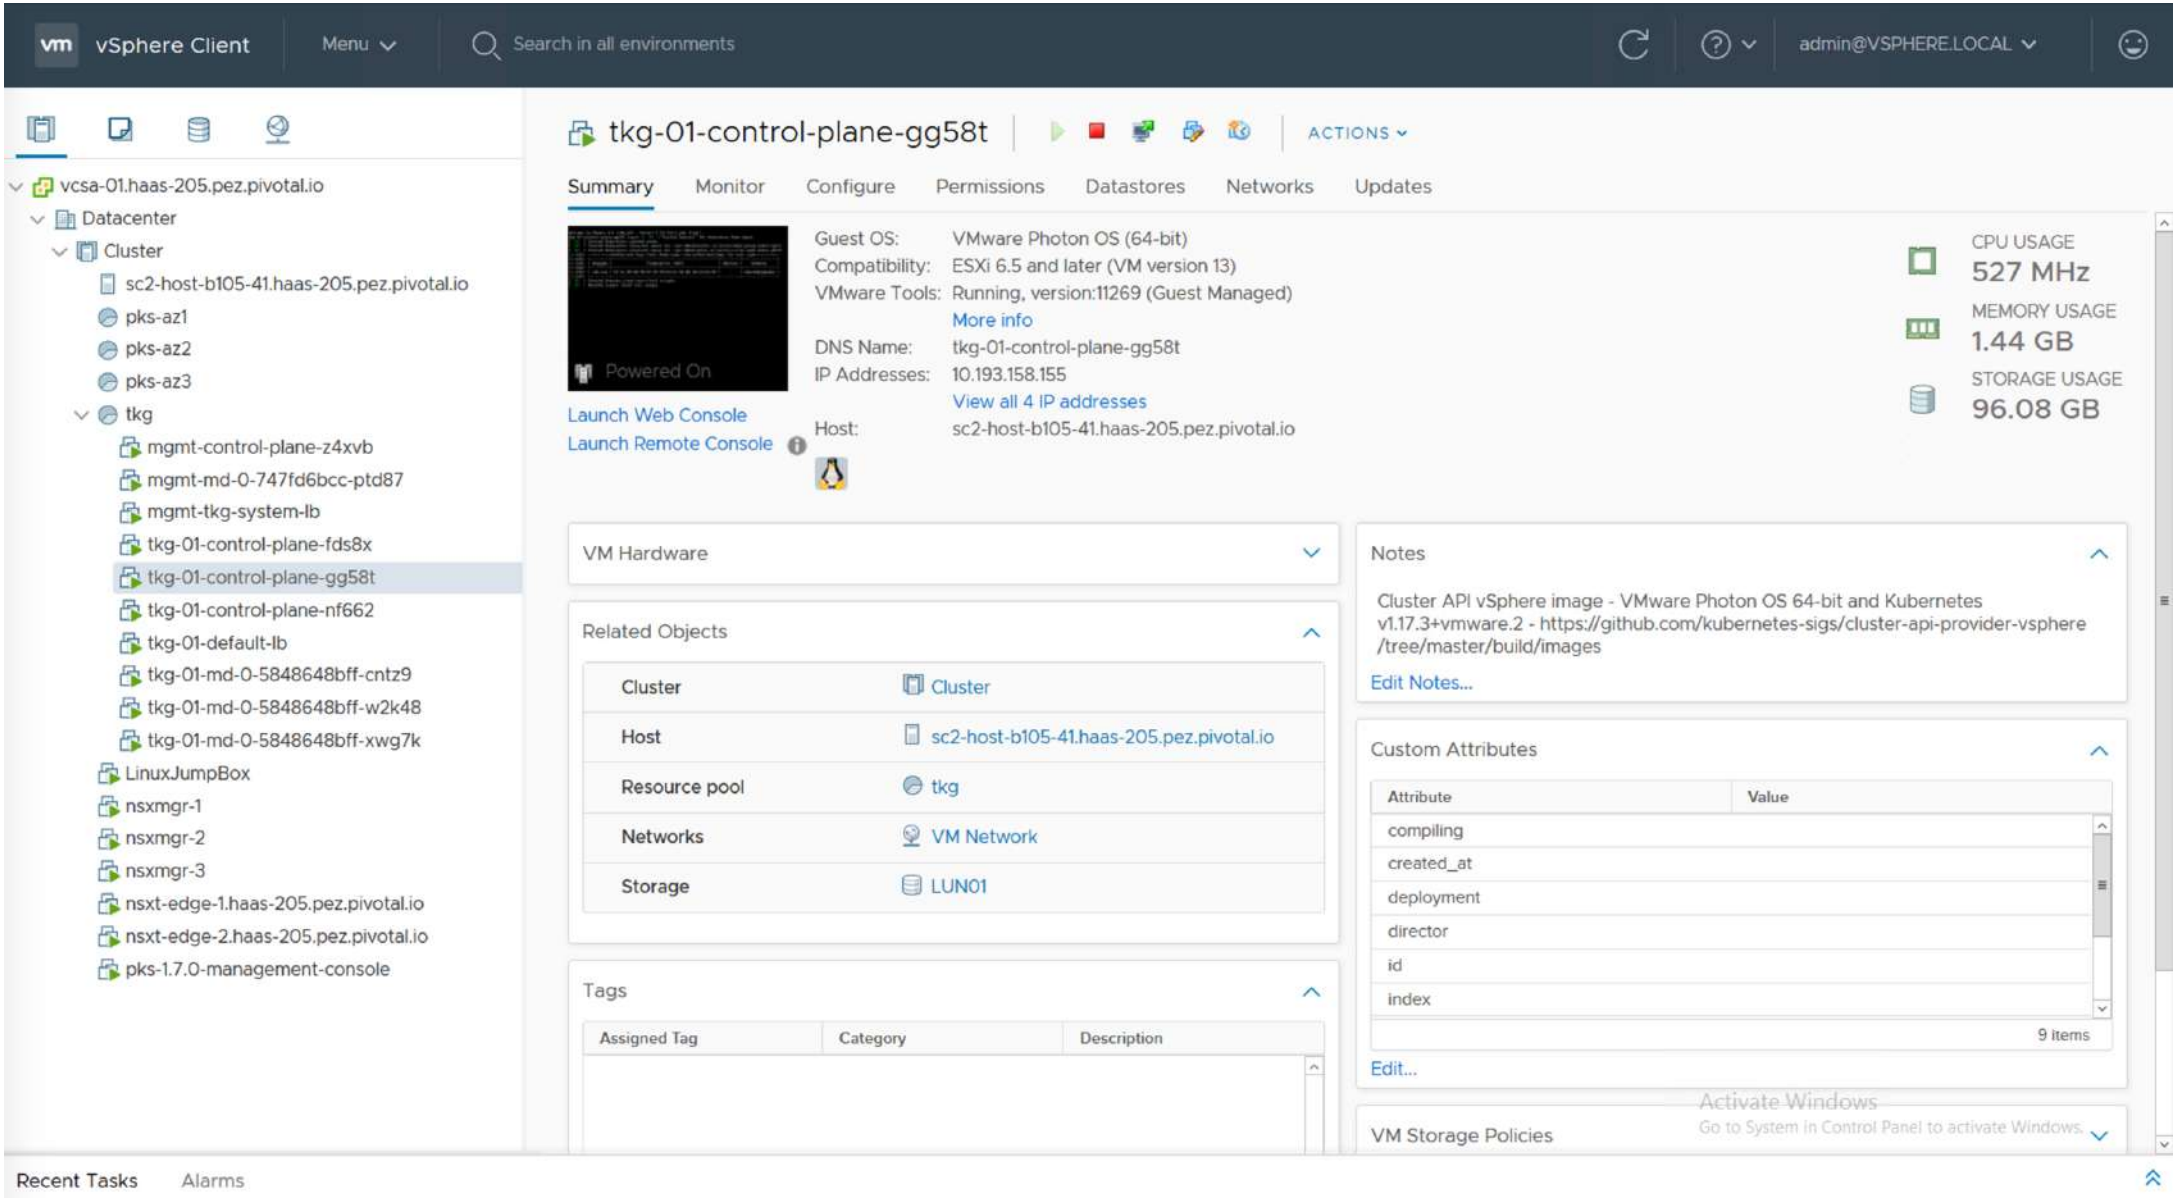
vSphere Client
Menu
Search in all environments
admin@VSPHERE.LOCAL
vcsa-01.haas-205.pez.pivotal.io
Datacenter
Cluster
sc2-host-b105-41.haas-205.pez.pivotal.io
pks-az1
pks-az2
pks-az3
tkg
mgmt-control-plane-z4xvb
mgmt-md-0-747fd6bcc-ptd87
mgmt-tkg-system-lb
tkg-01-control-plane-fds8x
tkg-01-control-plane-gg58t
tkg-01-control-plane-nf662
tkg-01-default-lb
tkg-01-md-0-5848648bff-cntz9
tkg-01-md-0-5848648bff-w2k48
tkg-01-md-0-5848648bff-xwg7k
LinuxJumpBox
nsxmgr-1
nsxmgr-2
nsxmgr-3
nsxt-edge-1.haas-205.pez.pivotal.io
nsxt-edge-2.haas-205.pez.pivotal.io
pks-1.7.0-management-console
tkg-01-control-plane-gg58t
ACTIONS
Summary
Monitor
Configure
Permissions
Datastores
Networks
Updates
Powered On
Launch Web Console
Launch Remote Console
Guest OS: VMware Photon OS (64-bit)
Compatibility: ESXi 6.5 and later (VM version 13)
VMware Tools: Running, version:11269 (Guest Managed)
More info
DNS Name: tkg-01-control-plane-gg58t
IP Addresses: 10.193.158.155
View all 4 IP addresses
Host: sc2-host-b105-41.haas-205.pez.pivotal.io
CPU USAGE
527 MHz
MEMORY USAGE
1.44 GB
STORAGE USAGE
96.08 GB
VM Hardware
Related Objects
Cluster Cluster
Host sc2-host-b105-41.haas-205.pez.pivotal.io
Resource pool tkg
Networks VM Network
Storage LUN01
Tags
Assigned Tag
Category
Description
Notes
Cluster API vSphere image - VMware Photon OS 64-bit and Kubernetes v1.17.3+vmware.2 - https://github.com/kubernetes-sigs/cluster-api-provider-vsphere/tree/master/build/images
Edit Notes...
Custom Attributes
Attribute
Value
compiling
created_at
deployment
director
id
index
9 items
Edit...
VM Storage Policies
Activate Windows
Go to System in Control Panel to activate Windows.
Recent Tasks
Alarms

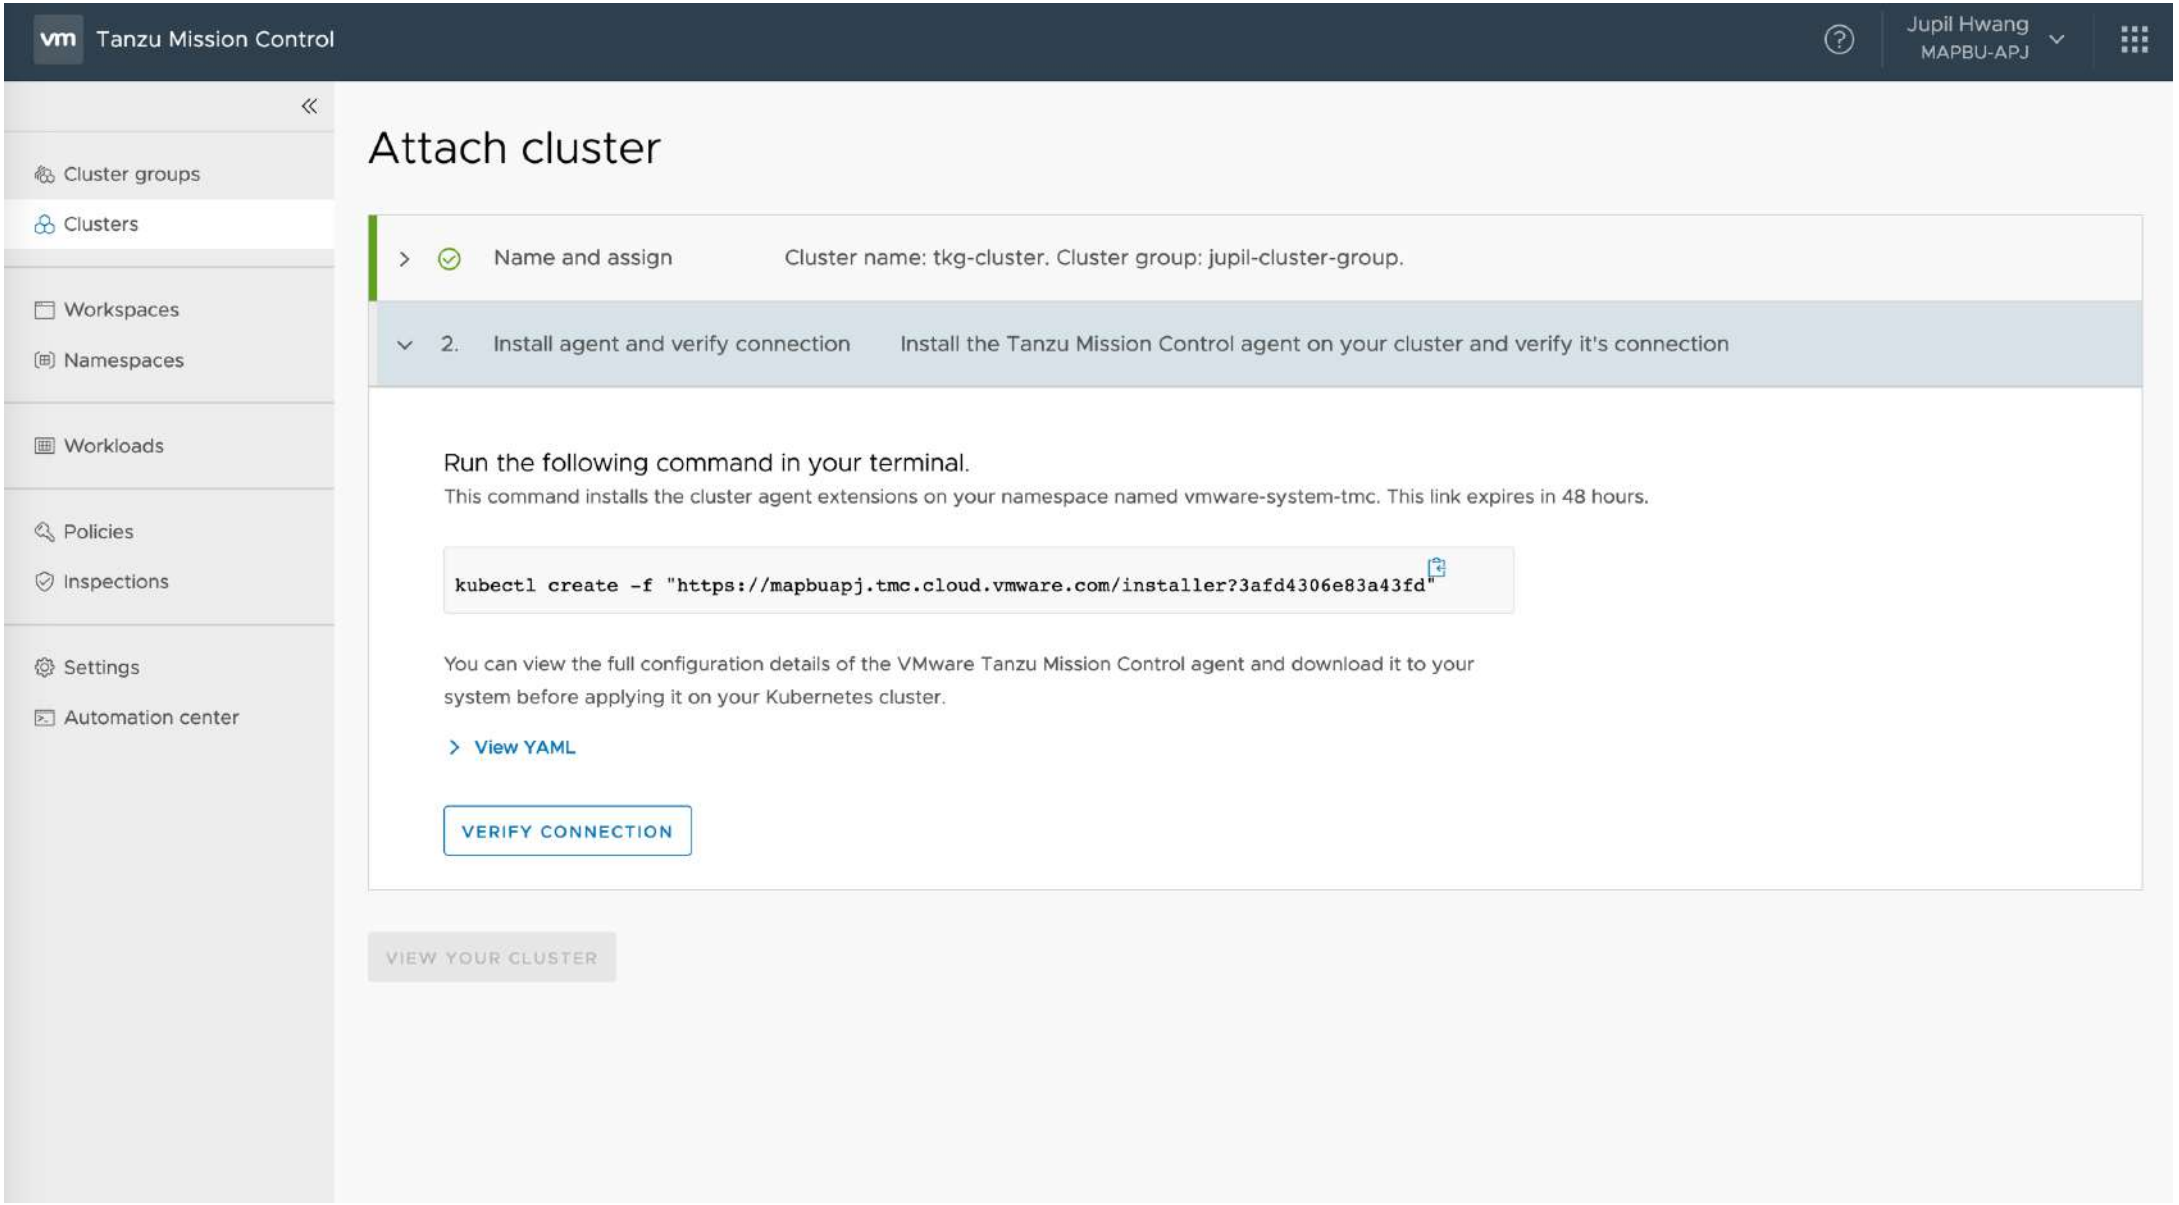
Tanzu Mission Control
Jupil Hwang
MAPBU-APJ
Cluster groups
Clusters
Workspaces
Namespaces
Workloads
Policies
Inspections
Settings
Automation center
Attach cluster
Name and assign
Cluster name: tkg-cluster. Cluster group: jupil-cluster-group.
2. Install agent and verify connection
Install the Tanzu Mission Control agent on your cluster and verify it's connection
Run the following command in your terminal.
This command installs the cluster agent extensions on your namespace named vmware-system-tmc. This link expires in 48 hours.
kubectl create -f "https://mapbuapj.tmc.cloud.vmware.com/installer?3afd4306e83a43fd"
You can view the full configuration details of the VMware Tanzu Mission Control agent and download it to your system before applying it on your Kubernetes cluster.
View YAML
VERIFY CONNECTION
VIEW YOUR CLUSTER

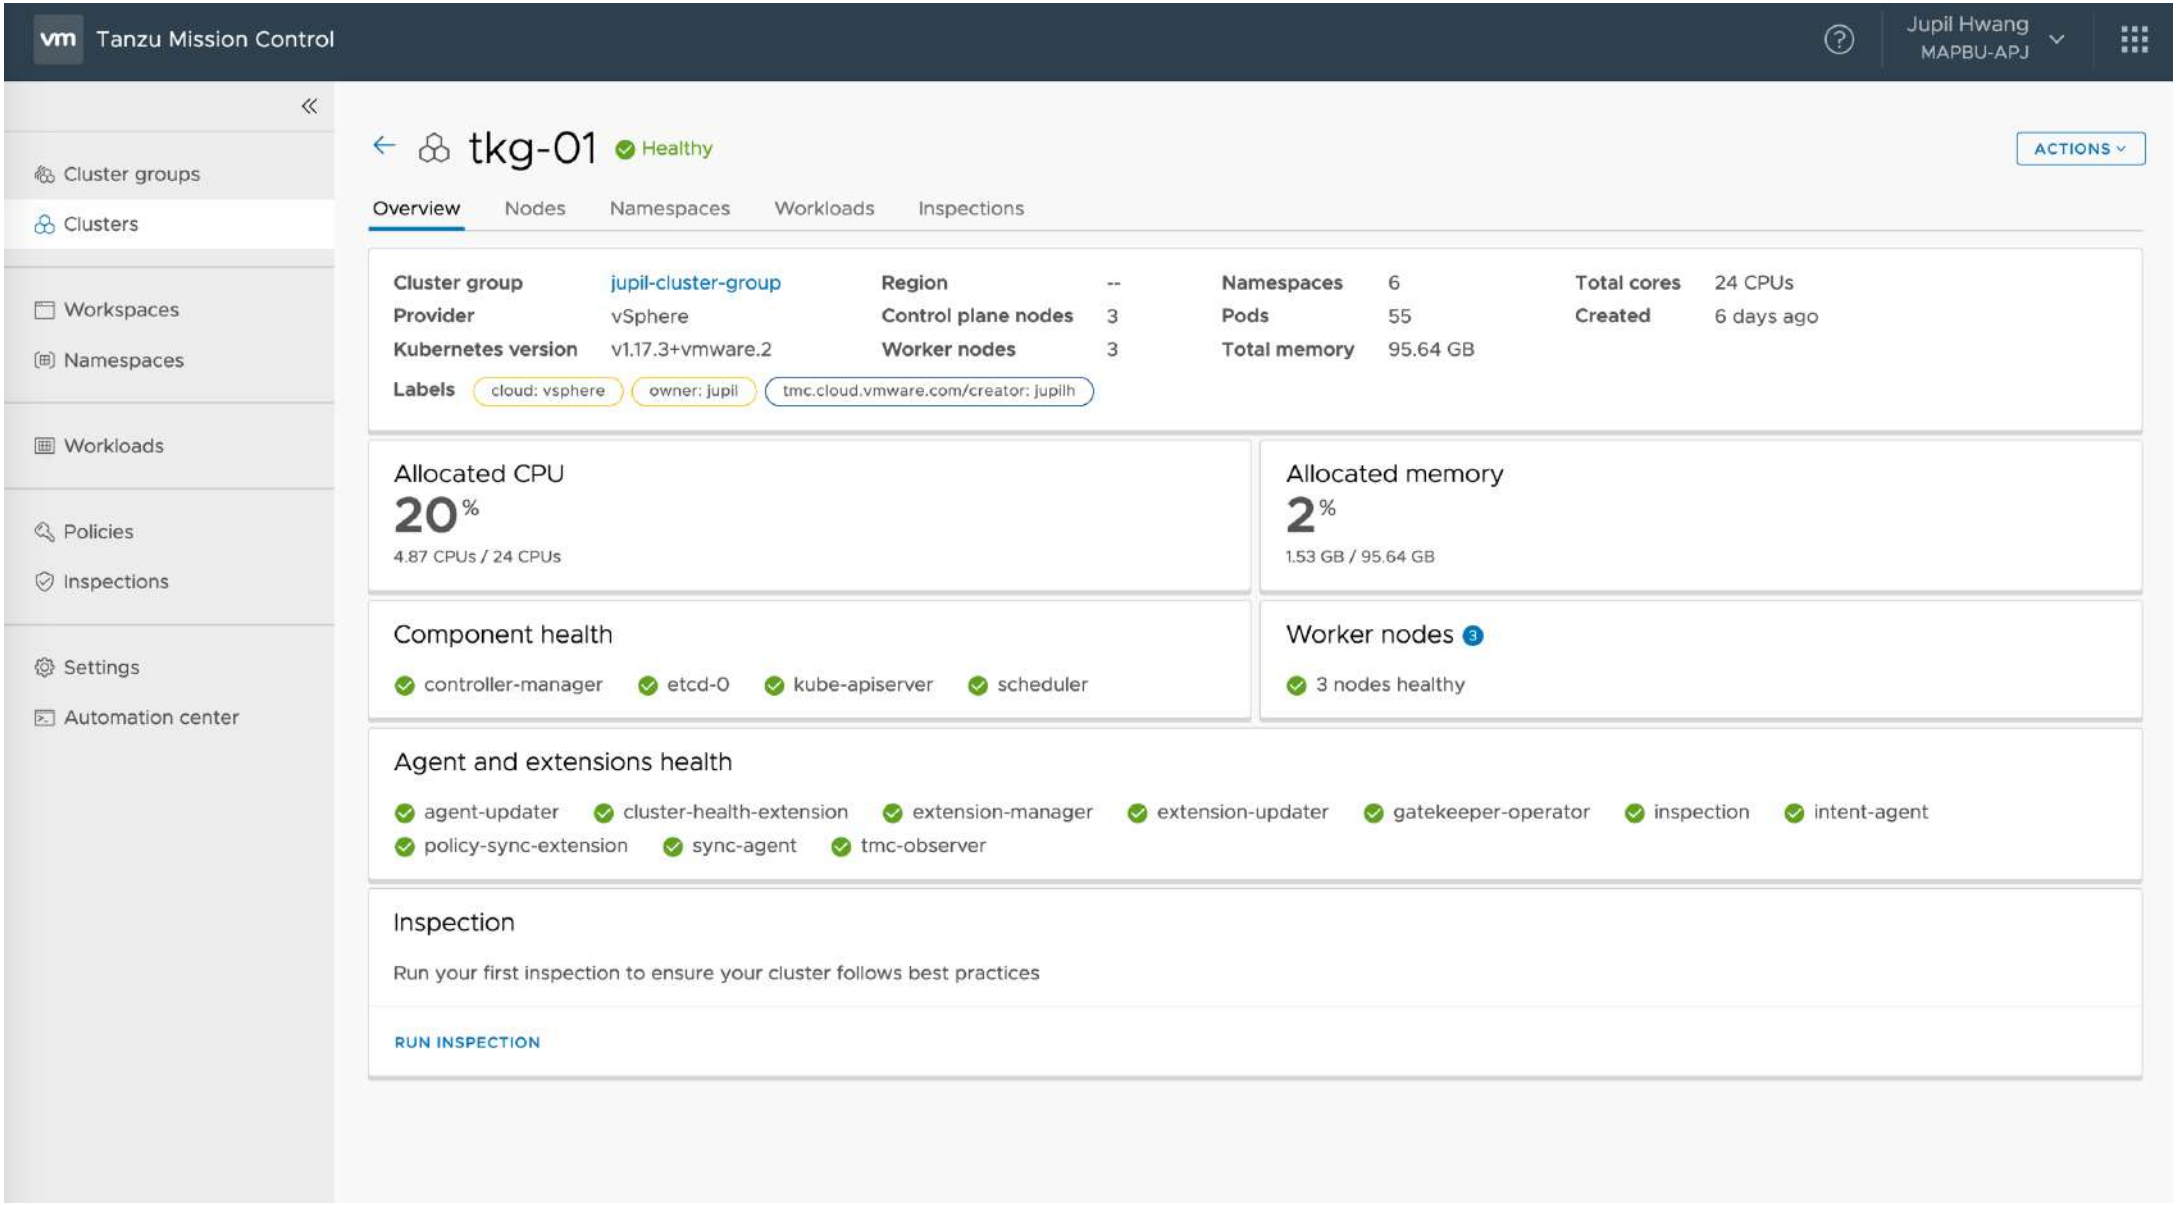
Tanzu Mission Control
Jupil Hwang
MAPBU-APJ
Cluster groups
Clusters
Workspaces
Namespaces
Workloads
Policies
Inspections
Settings
Automation center
tkg-01
Healthy
ACTIONS
Overview
Nodes
Namespaces
Workloads
Inspections
Cluster group
jupil-cluster-group
Region
--
Namespaces
6
Total cores
24 CPUs
Provider
vSphere
Control plane nodes
3
Pods
55
Created
6 days ago
Kubernetes version
v1.17.3+vmware.2
Worker nodes
3
Total memory
95.64 GB
Labels
cloud: vsphere
owner: jupil
tmc.cloud.vmware.com/creator: jupilh
Allocated CPU
20%
4.87 CPUs / 24 CPUs
Allocated memory
2%
1.53 GB / 95.64 GB
Component health
controller-manager
etcd-0
kube-apiserver
scheduler
Worker nodes
3
3 nodes healthy
Agent and extensions health
agent-updater
cluster-health-extension
extension-manager
extension-updater
gatekeeper-operator
inspection
intent-agent
policy-sync-extension
sync-agent
tmc-observer
Inspection
Run your first inspection to ensure your cluster follows best practices
RUN INSPECTION

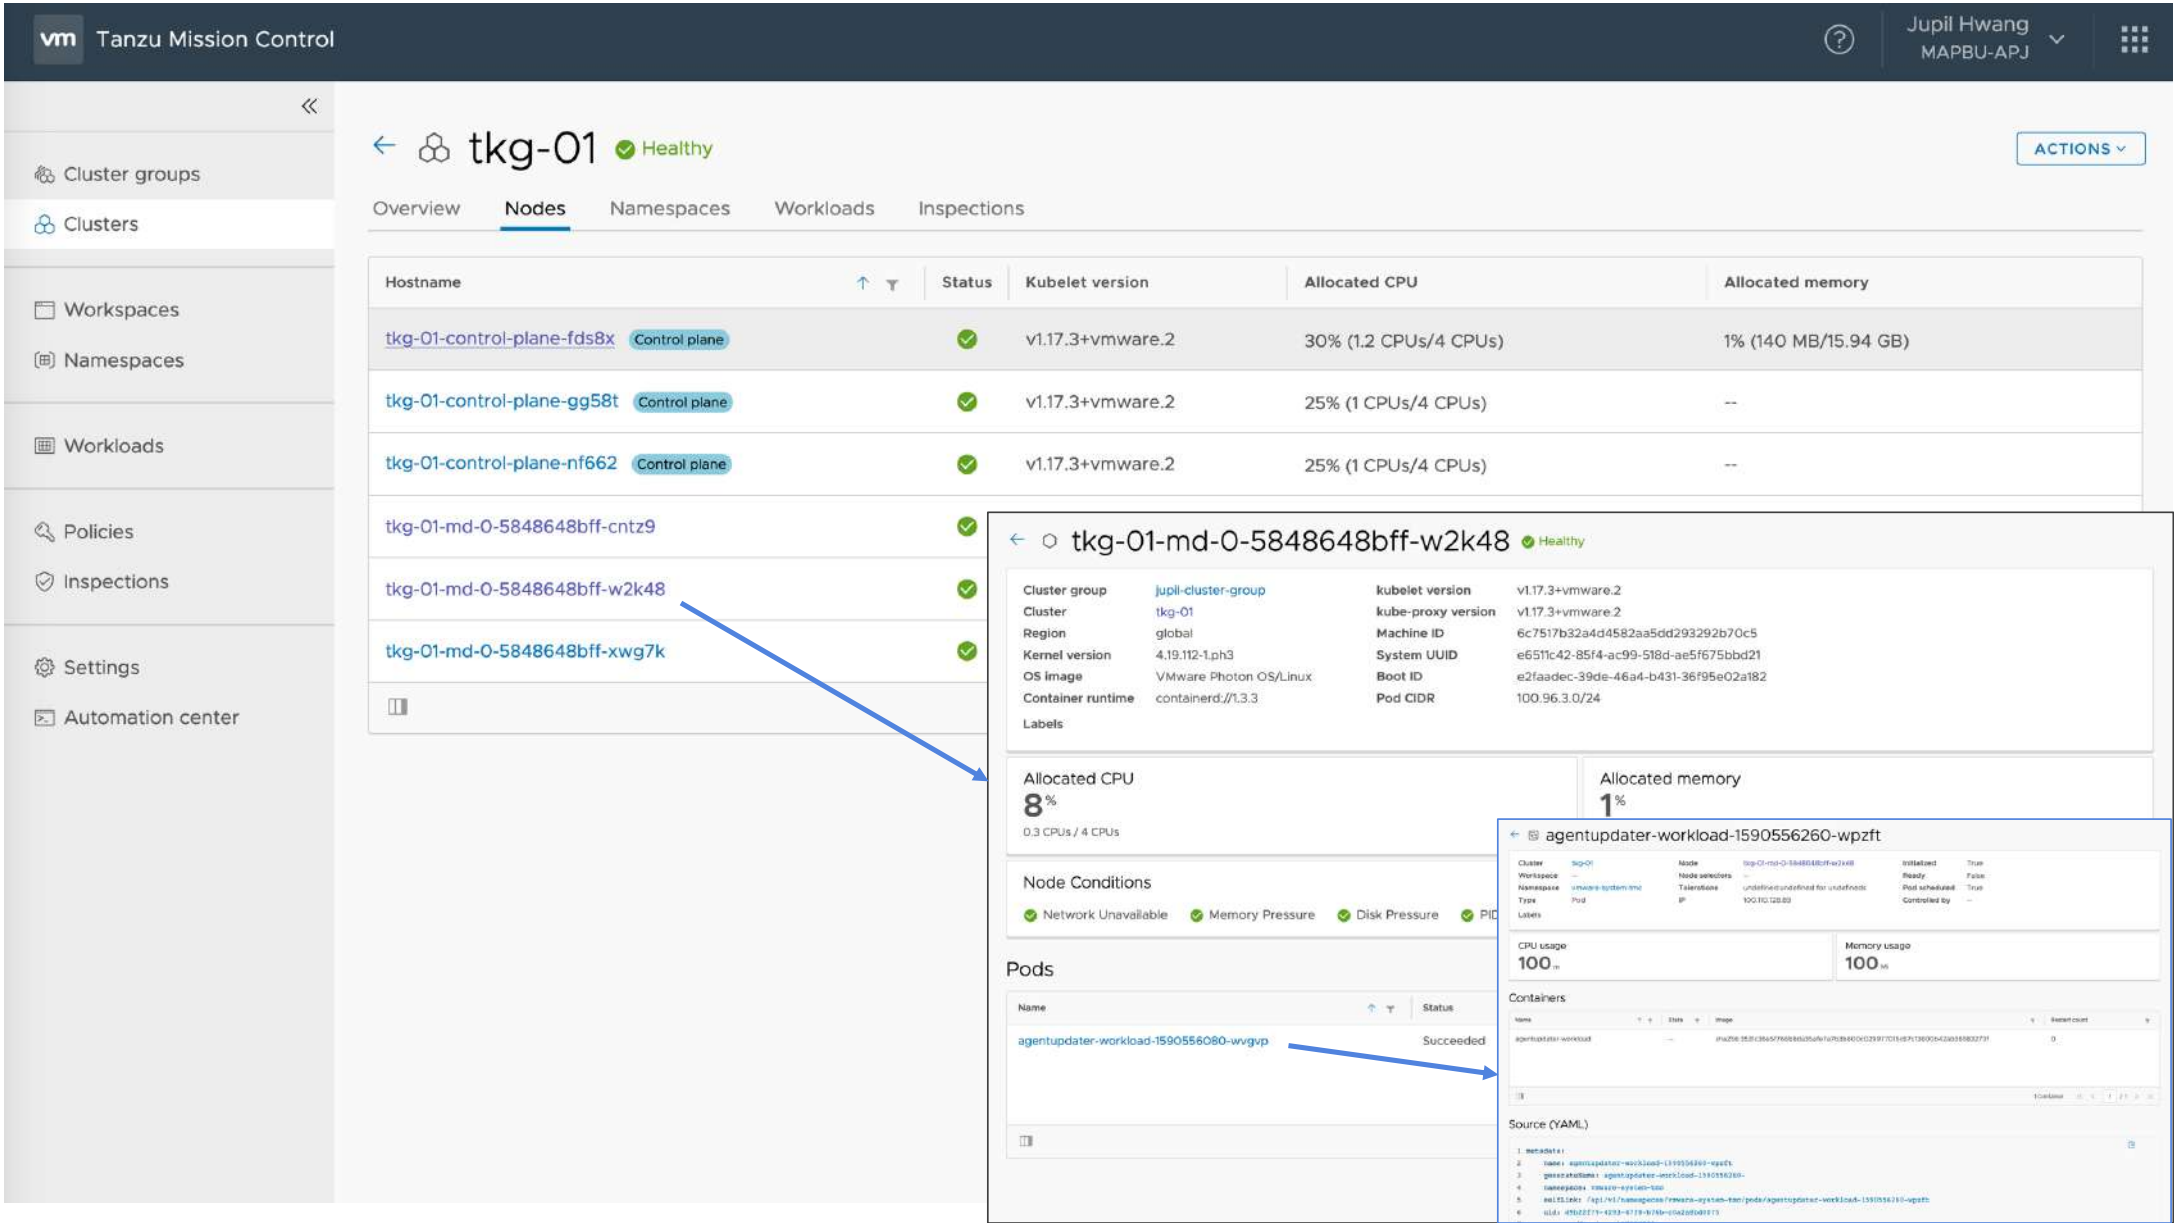
Tanzu Mission Control
Jupil Hwang
MAPBU-APJ
Cluster groups
Clusters
Workspaces
Namespaces
Workloads
Policies
Inspections
Settings
Automation center
tkg-01
Healthy
ACTIONS
Overview
Nodes
Namespaces
Workloads
Inspections
Hostname
Status
Kubelet version
Allocated CPU
Allocated memory
tkg-01-control-plane-fds8x
Control plane
v1.17.3+vmware.2
30% (1.2 CPUs/4 CPUs)
1% (140 MB/15.94 GB)
tkg-01-control-plane-gg58t
Control plane
v1.17.3+vmware.2
25% (1 CPUs/4 CPUs)
--
tkg-01-control-plane-nf662
Control plane
v1.17.3+vmware.2
25% (1 CPUs/4 CPUs)
--
tkg-01-md-0-5848648bff-cntz9
tkg-01-md-0-5848648bff-w2k48
tkg-01-md-0-5848648bff-xwg7k
tkg-01-md-0-5848648bff-w2k48
Healthy
Cluster group
jupil-cluster-group
Cluster
tkg-01
Region
global
Kernel version
4.19.112-1.ph3
OS image
VMware Photon OS/Linux
Container runtime
containerd://1.3.3
Labels
kubelet version
v1.17.3+vmware.2
kube-proxy version
v1.17.3+vmware.2
Machine ID
6c7517b32a4d4582aa5dd293292b70c5
System UUID
e6511c42-85f4-ac99-518d-ae5f675bbd21
Boot ID
e2faadec-39de-46a4-b431-36f95e02a182
Pod CIDR
100.96.3.0/24
Allocated CPU
8%
0.3 CPUs / 4 CPUs
Allocated memory
1%
Node Conditions
Network Unavailable
Memory Pressure
Disk Pressure
Pods
Name
Status
agentupdater-workload-1590556080-wvgvp
Succeeded
agentupdater-workload-1590556260-wpzft
CPU usage
100
Memory usage
100
Containers
Source (YAML)

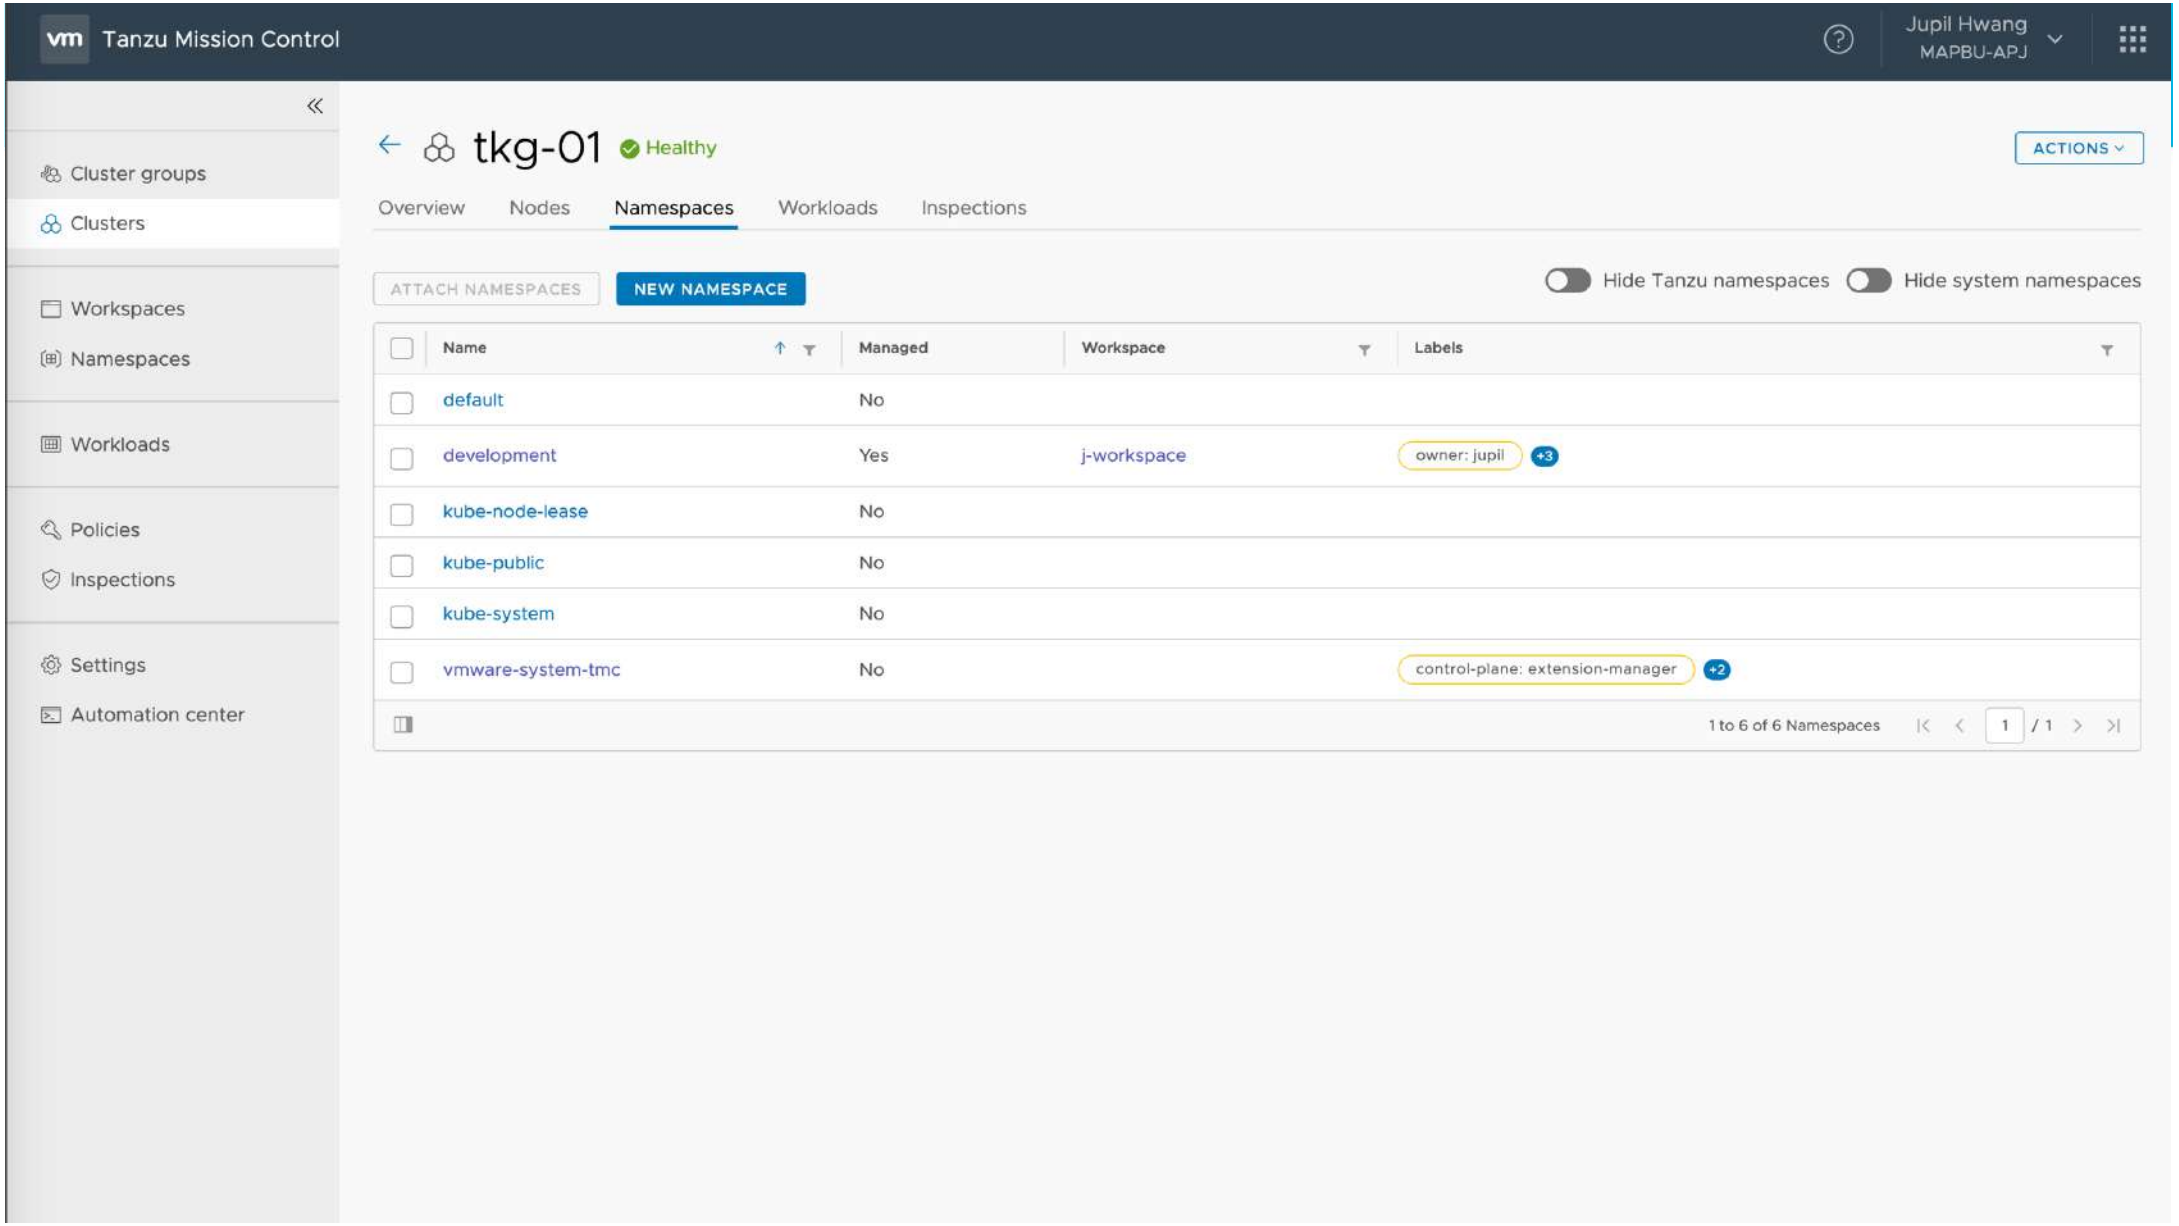
Tanzu Mission Control
Jupil Hwang
MAPBU-APJ
Cluster groups
Clusters
Workspaces
Namespaces
Workloads
Policies
Inspections
Settings
Automation center
tkg-01
Healthy
ACTIONS
Overview
Nodes
Namespaces
Workloads
Inspections
ATTACH NAMESPACES
NEW NAMESPACE
Hide Tanzu namespaces
Hide system namespaces
Name
Managed
Workspace
Labels
default
No
development
Yes
j-workspace
owner: jupil
+3
kube-node-lease
No
kube-public
No
kube-system
No
vmware-system-tmc
No
control-plane: extension-manager
+2
1 to 6 of 6 Namespaces
1
/ 1

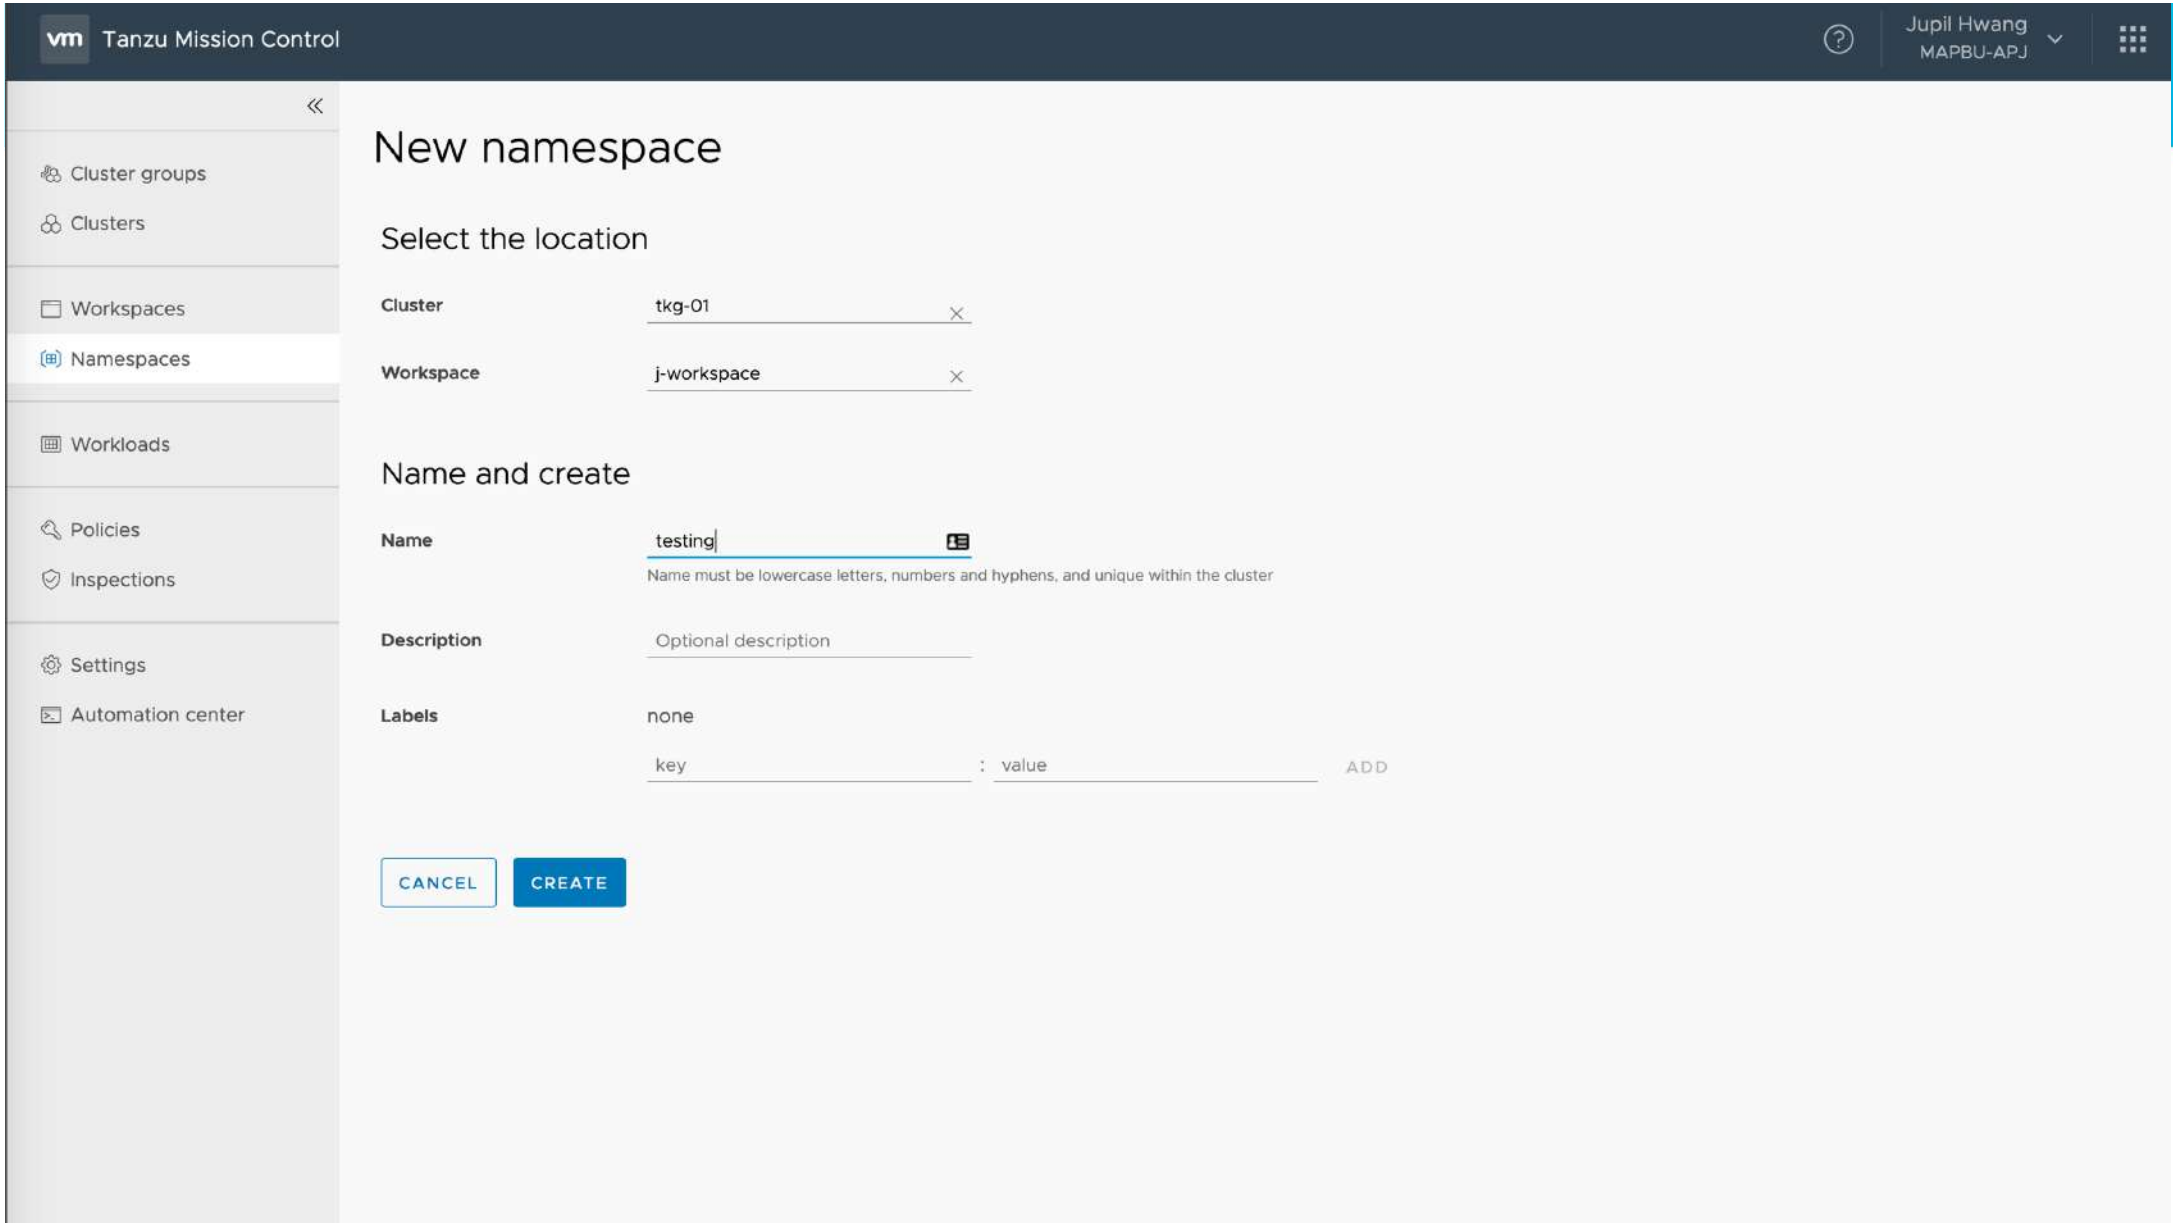
Tanzu Mission Control
Jupil Hwang
MAPBU-APJ
Cluster groups
Clusters
Workspaces
Namespaces
Workloads
Policies
Inspections
Settings
Automation center
New namespace
Select the location
Cluster
tkg-01
Workspace
j-workspace
Name and create
Name
testing
Name must be lowercase letters, numbers and hyphens, and unique within the cluster
Description
Optional description
Labels
none
key
value
ADD
CANCEL
CREATE

보안 및 구성관리

플랫폼
운영자

애플리케이션
운영자

Tanzu Mission Control

보안 /
구성관리

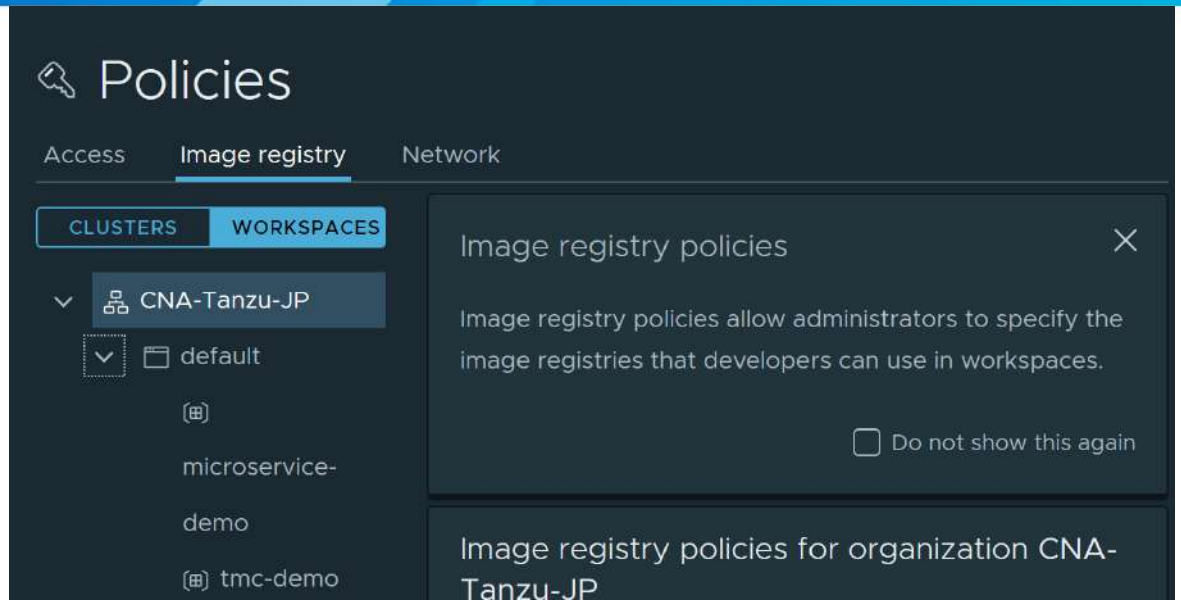

정책유형	
엑세스	Tanzu Mission Control IAM 롤 적용。K8s 롤 바인딩의 자동 설정
이미지 레지스트리	파드 Pull을 할 이미지 레지스트리의 제한。안전하지 않은 이미지의 실행 방지
네트워크	파드와 네임스페이스의 입출력 트래픽의 통신 제어
파드 보안	root 권한 실행 금지, 루트 파일 시스템 마운트 금지 등 실행 제한

정책 프레임워크

적용한 그룹 개체에 대한 정책 일괄 적용

정책이 계층 구조 트리를 따라 하위로 전달됨

전체적으로 복수 개의 클러스터들에 대해 조직, 클러스터 그룹 레벨로 정책 적용

다수의 클러스터에 걸쳐 존재하는 다수의 네임스페이스들에 대해 그룹핑한 워크스페이스 단위 정책 적용

Direct Policy: 계층 구조 상에서 노드에 적용

Inherited Policy: 계층 구조의 상위 노드에서 정책 상속

클러스터 및 관리 대상인 네임스페이스는 언제든지 단일 상위 노드를 가질 수 있음

플랫폼운영자 관점

클러스터 그룹을 통해「클러스터」정책을 제어

애플리케이션 운영자 관점

워크스페이스을 통해「네임스페이스」의 정책을 제어

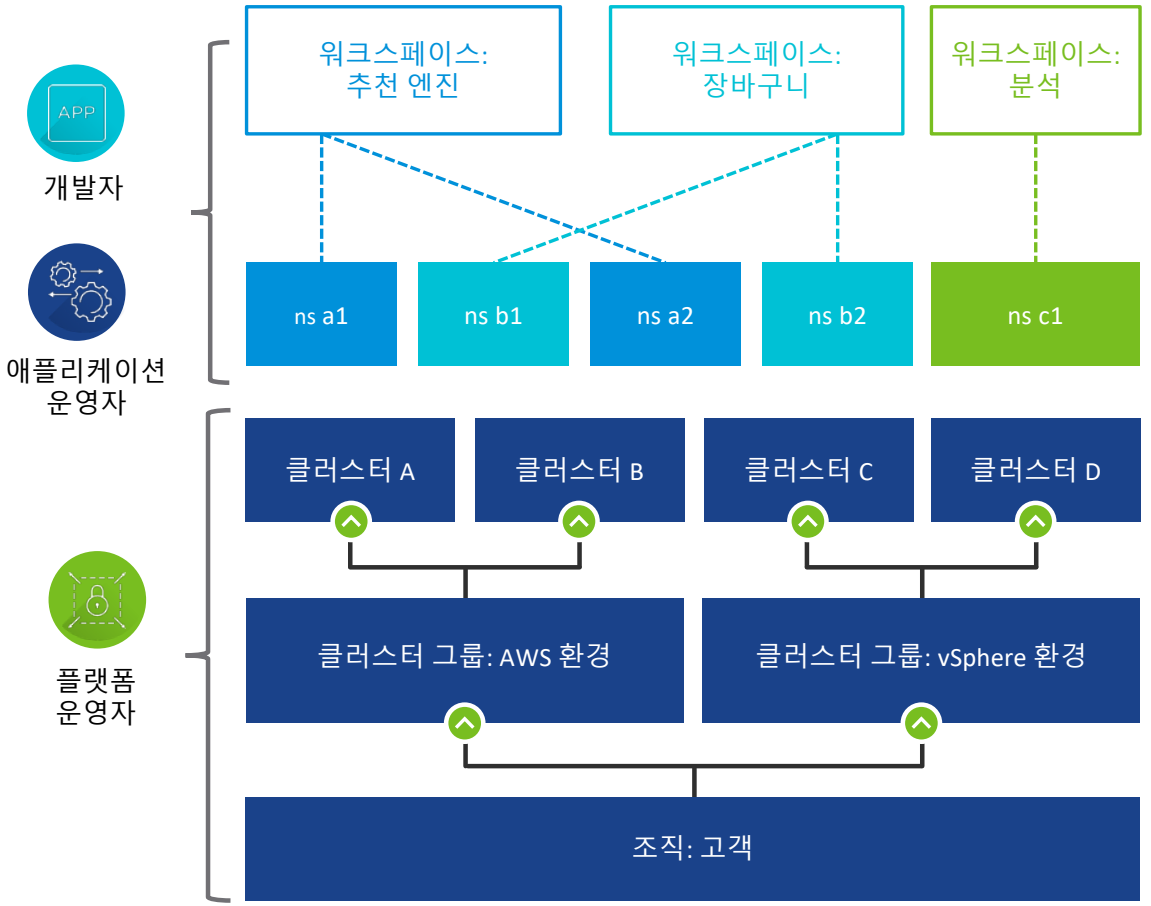

ID / 엑세스관리

전체 K8s 클러스터에서 통합 ID 관리와 엑세스 정책

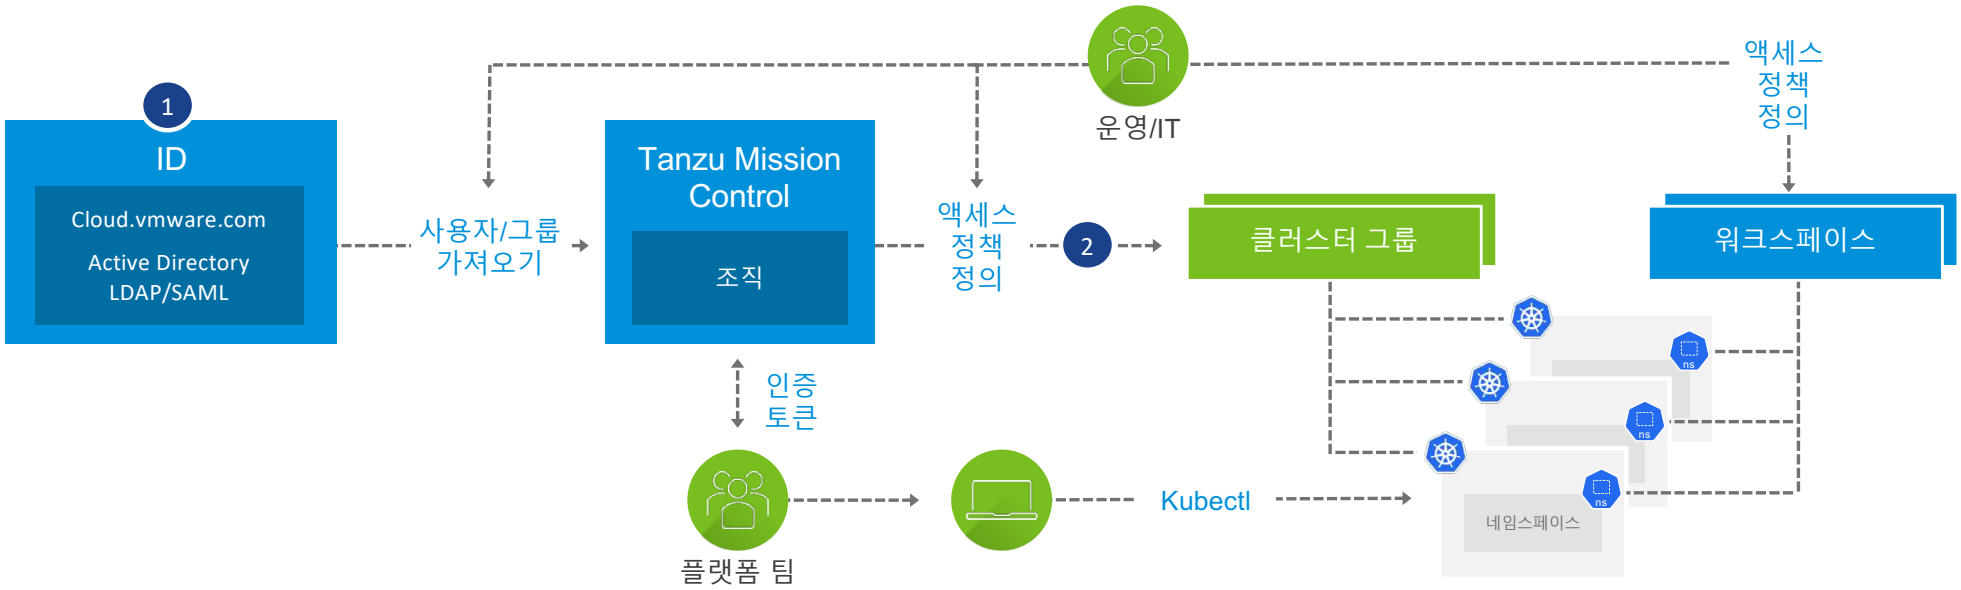

인프라 팀에서 클러스터 그룹 및 역할 매핑을 통해 다수의 클러스터에 대한 사용자 액세스를 한번에 정의 : Org Admin, User

개발자가 클러스터에 대한 셀프 서비스 액세스 확보 : Cluster Admin, Edit, View etc..

사용자/그룹을 Tanzu Mission Control 롤(조직 관리자, 사용자)에 매핑하면, 쿠버네티스 ClusterRoleBinding 으로 반영

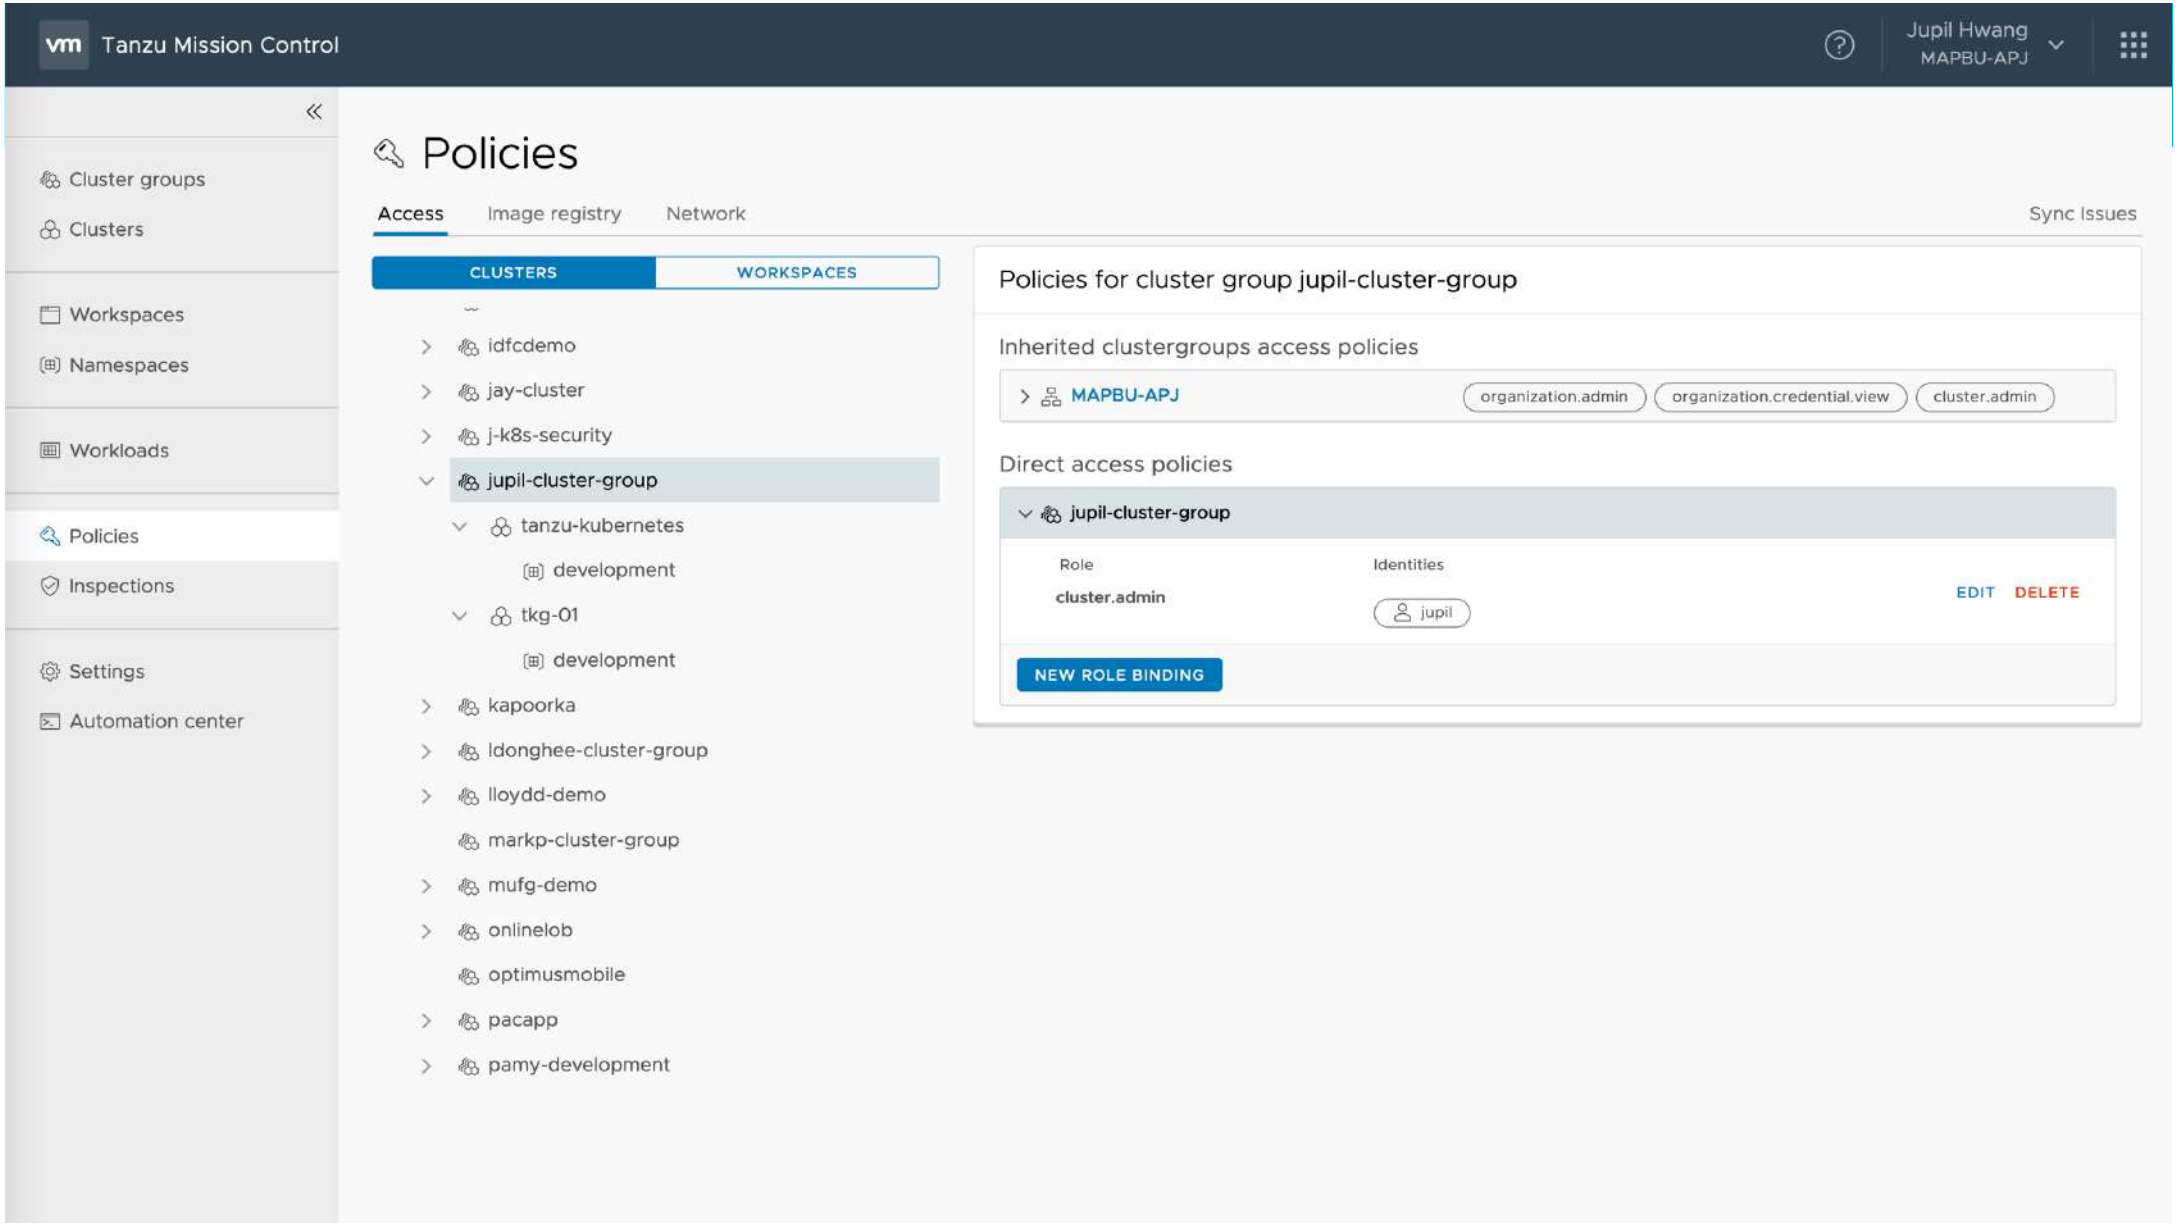
Tanzu Mission Control
Jupil Hwang
MAPBU-APJ
Cluster groups
Clusters
Workspaces
Namespaces
Workloads
Policies
Inspections
Settings
Automation center
Policies
Access
Image registry
Network
Sync Issues
CLUSTERS
WORKSPACES
idfcdemo
jay-cluster
j-k8s-security
jupil-cluster-group
tanzu-kubernetes
development
tkg-01
development
kapoorka
ldonghee-cluster-group
lloydd-demo
markp-cluster-group
mufg-demo
onlinelob
optimusmobile
pacapp
pamy-development
Policies for cluster group jupil-cluster-group
Inherited clustergroups access policies
MAPBU-APJ
organization.admin
organization.credential.view
cluster.admin
Direct access policies
jupil-cluster-group
Role
Identities
cluster.admin
jupil
EDIT
DELETE
NEW ROLE BINDING

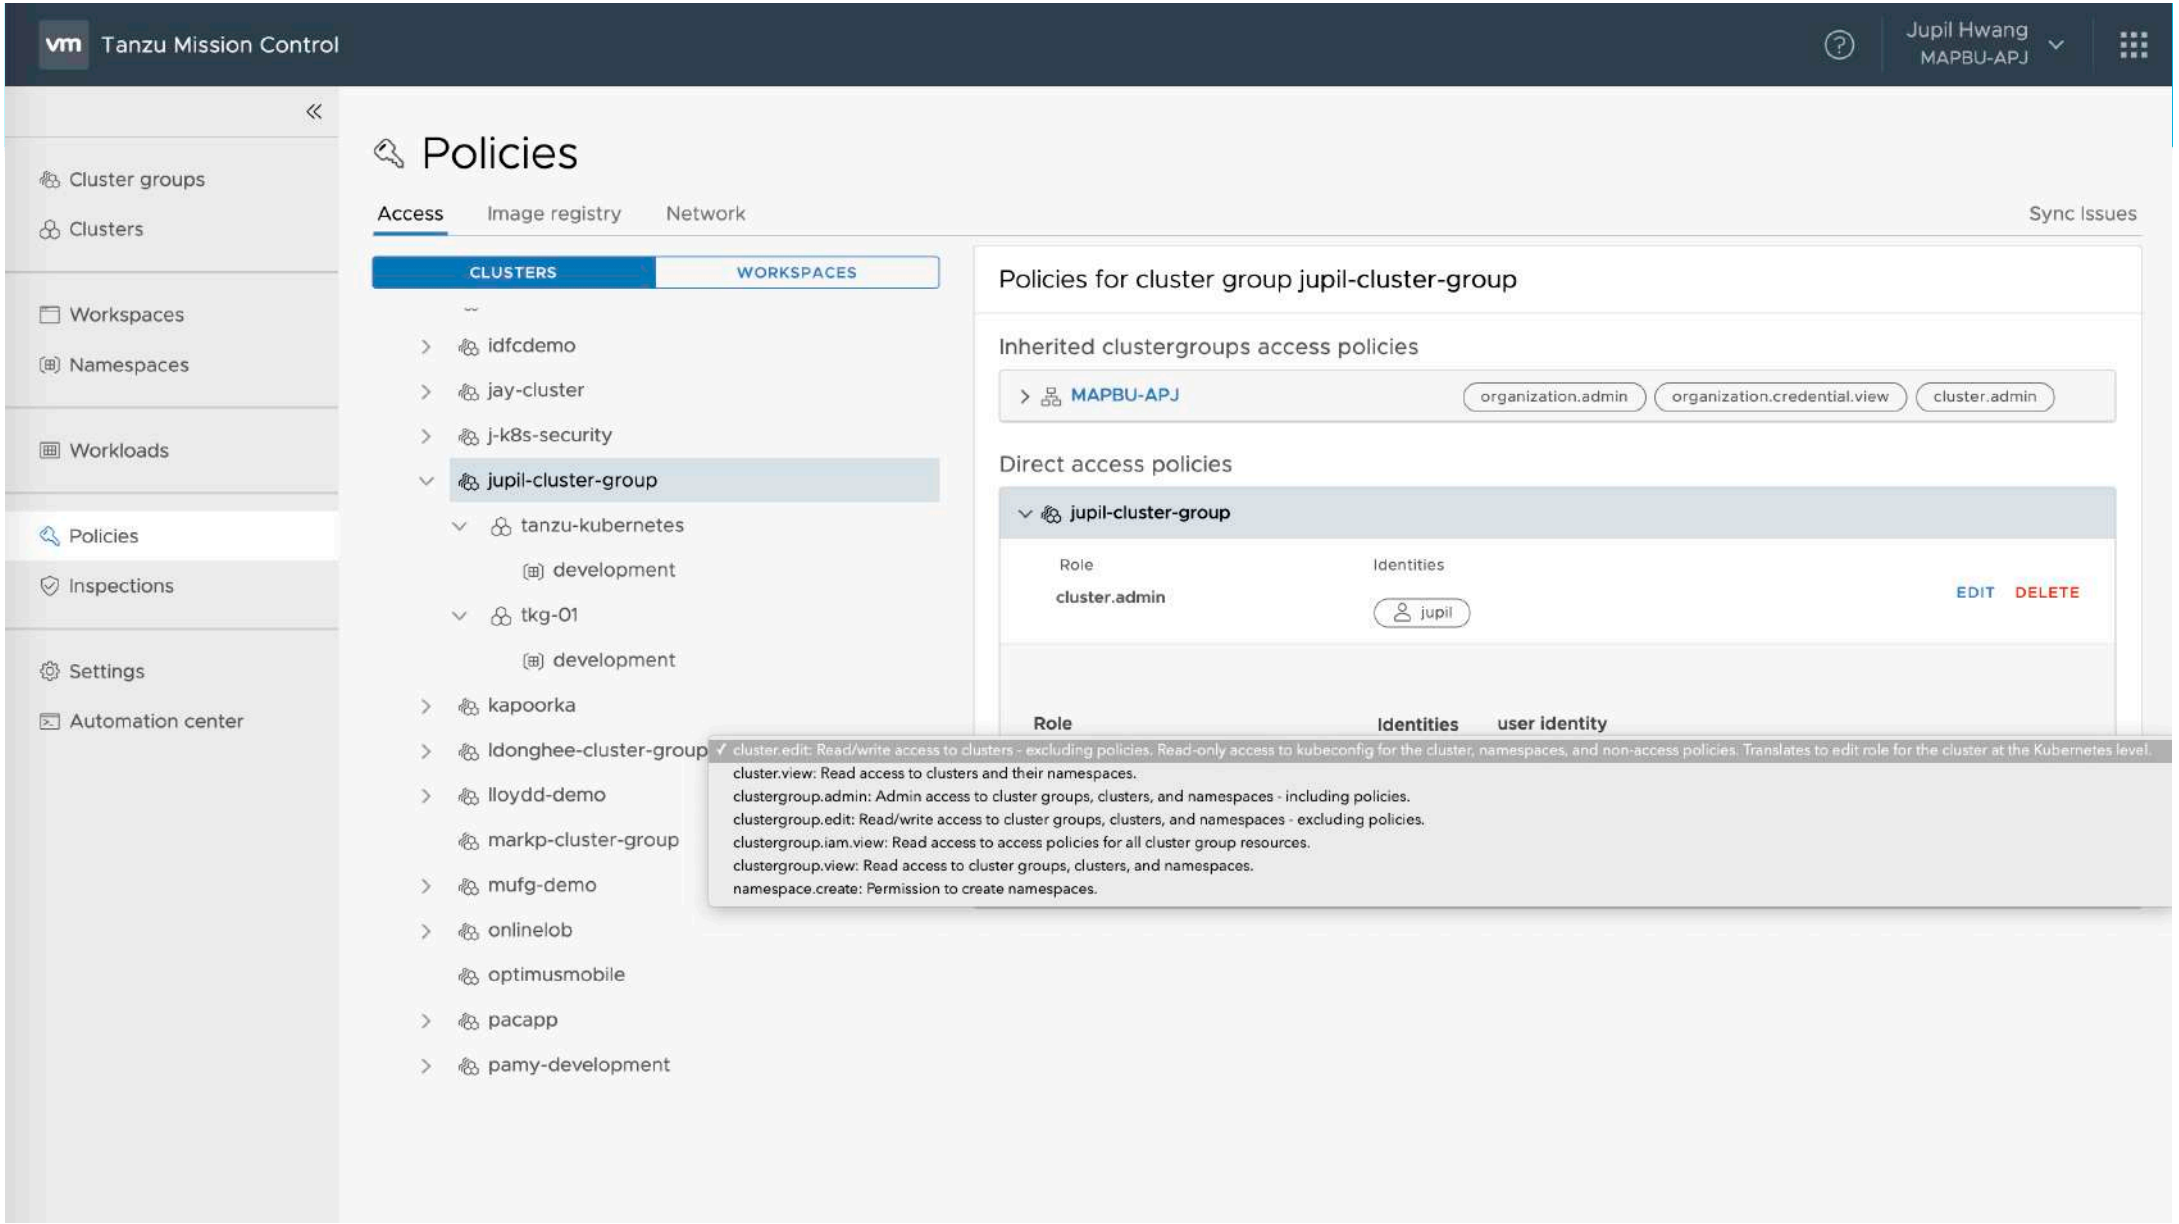
Tanzu Mission Control
Jupil Hwang
MAPBU-APJ
Cluster groups
Clusters
Workspaces
Namespaces
Workloads
Policies
Inspections
Settings
Automation center
Policies
Access
Image registry
Network
Sync Issues
CLUSTERS
WORKSPACES
idfcdemo
jay-cluster
j-k8s-security
jupil-cluster-group
tanzu-kubernetes
development
tkg-01
development
kapoorka
ldonghee-cluster-group
lloydd-demo
markp-cluster-group
mufg-demo
onlinelob
optimusmobile
pacapp
pamy-development
Policies for cluster group jupil-cluster-group
Inherited clustergroups access policies
MAPBU-APJ
organization.admin
organization.credential.view
cluster.admin
Direct access policies
jupil-cluster-group
Role
Identities
cluster.admin
jupil
EDIT
DELETE
Role
Identities
user identity
cluster.edit: Read/write access to clusters - excluding policies. Read-only access to kubeconfig for the cluster, namespaces, and non-access policies. Translates to edit role for the cluster at the Kubernetes level.
cluster.view: Read access to clusters and their namespaces.
clustergroup.admin: Admin access to cluster groups, clusters, and namespaces - including policies.
clustergroup.edit: Read/write access to cluster groups, clusters, and namespaces - excluding policies.
clustergroup.iam.view: Read access to access policies for all cluster group resources.
clustergroup.view: Read access to cluster groups, clusters, and namespaces.
namespace.create: Permission to create namespaces.

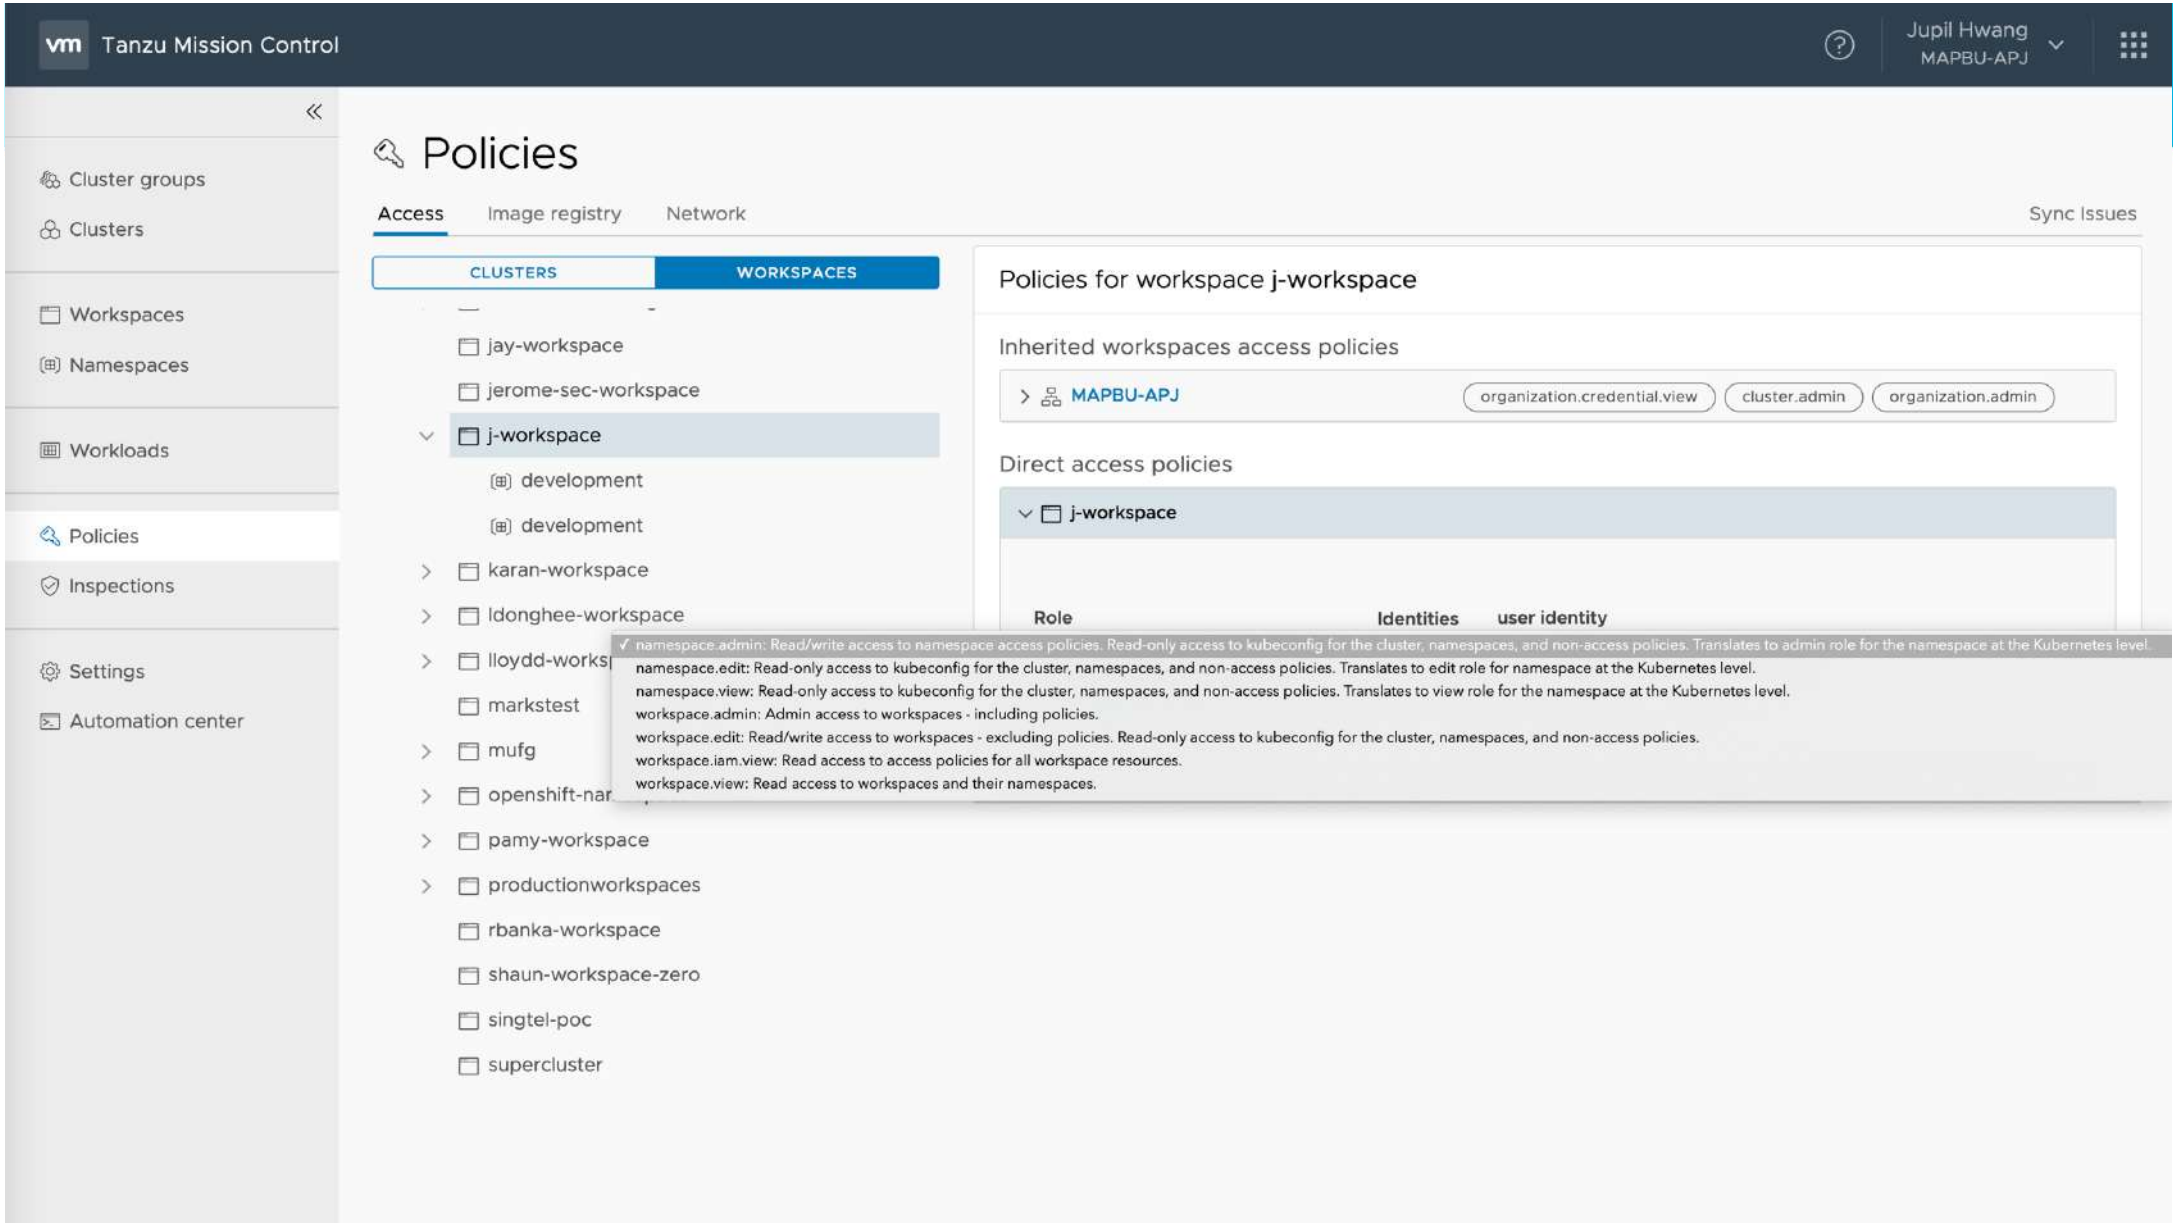
Tanzu Mission Control
Jupil Hwang
MAPBU-APJ
Cluster groups
Clusters
Workspaces
Namespaces
Workloads
Policies
Inspections
Settings
Automation center
Policies
Access
Image registry
Network
Sync Issues
CLUSTERS
WORKSPACES
jay-workspace
jerome-sec-workspace
j-workspace
development
development
karan-workspace
ldonghee-workspace
lloydd-works
markstest
mufg
openshift-nar
pamy-workspace
productionworkspaces
rbanka-workspace
shaun-workspace-zero
singtel-poc
supercluster
Policies for workspace j-workspace
Inherited workspaces access policies
MAPBU-APJ
organization.credential.view
cluster.admin
organization.admin
Direct access policies
j-workspace
Role
Identities
user identity
namespace.admin: Read/write access to namespace access policies. Read-only access to kubeconfig for the cluster, namespaces, and non-access policies. Translates to admin role for the namespace at the Kubernetes level.
namespace.edit: Read-only access to kubeconfig for the cluster, namespaces, and non-access policies. Translates to edit role for namespace at the Kubernetes level.
namespace.view: Read-only access to kubeconfig for the cluster, namespaces, and non-access policies. Translates to view role for the namespace at the Kubernetes level.
workspace.admin: Admin access to workspaces - including policies.
workspace.edit: Read/write access to workspaces - excluding policies. Read-only access to kubeconfig for the cluster, namespaces, and non-access policies.
workspace.iam.view: Read access to access policies for all workspace resources.
workspace.view: Read access to workspaces and their namespaces.

Pod Security 정책

Pod 스펙 상에서 보안에 민감한 측면 제어(예: privileged containers)

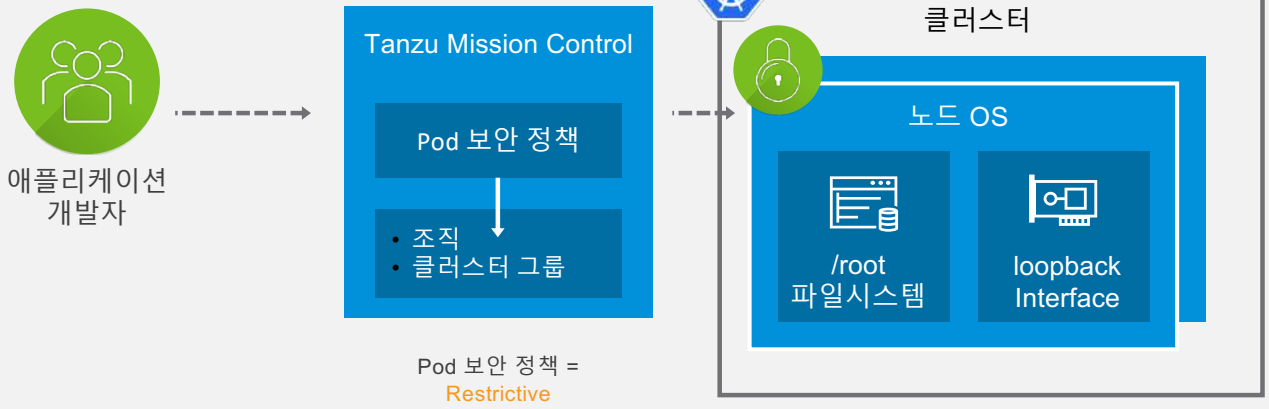

Pod가 수행할 수 없는 기능

- 루트 파일 시스템 마운팅
- 호스트 네트워크에 바인딩
- 호스트 볼륨 마운팅
- 등등

방식

Pod 보안 정책이 적용되면 Admission Controller가 클러스터에서 활성화됩니다.

컨테이너 이미지 정책

허용된 레지스트리를 명시하여 안전하지 않은 위치에서 이미지를 가져오지 않도록 보장

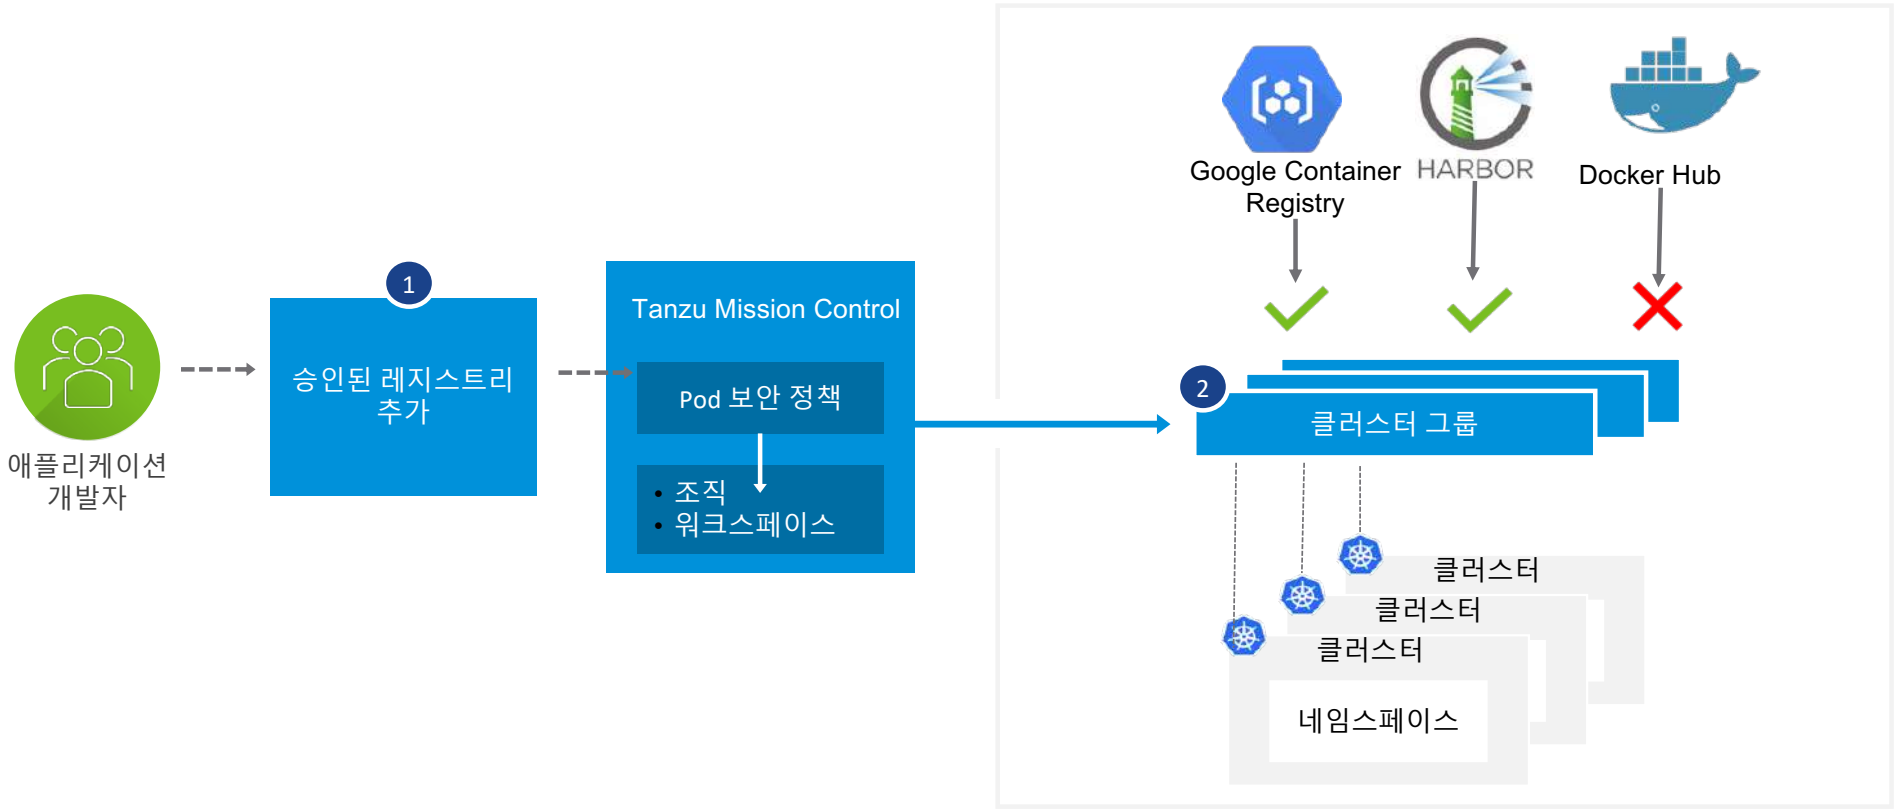

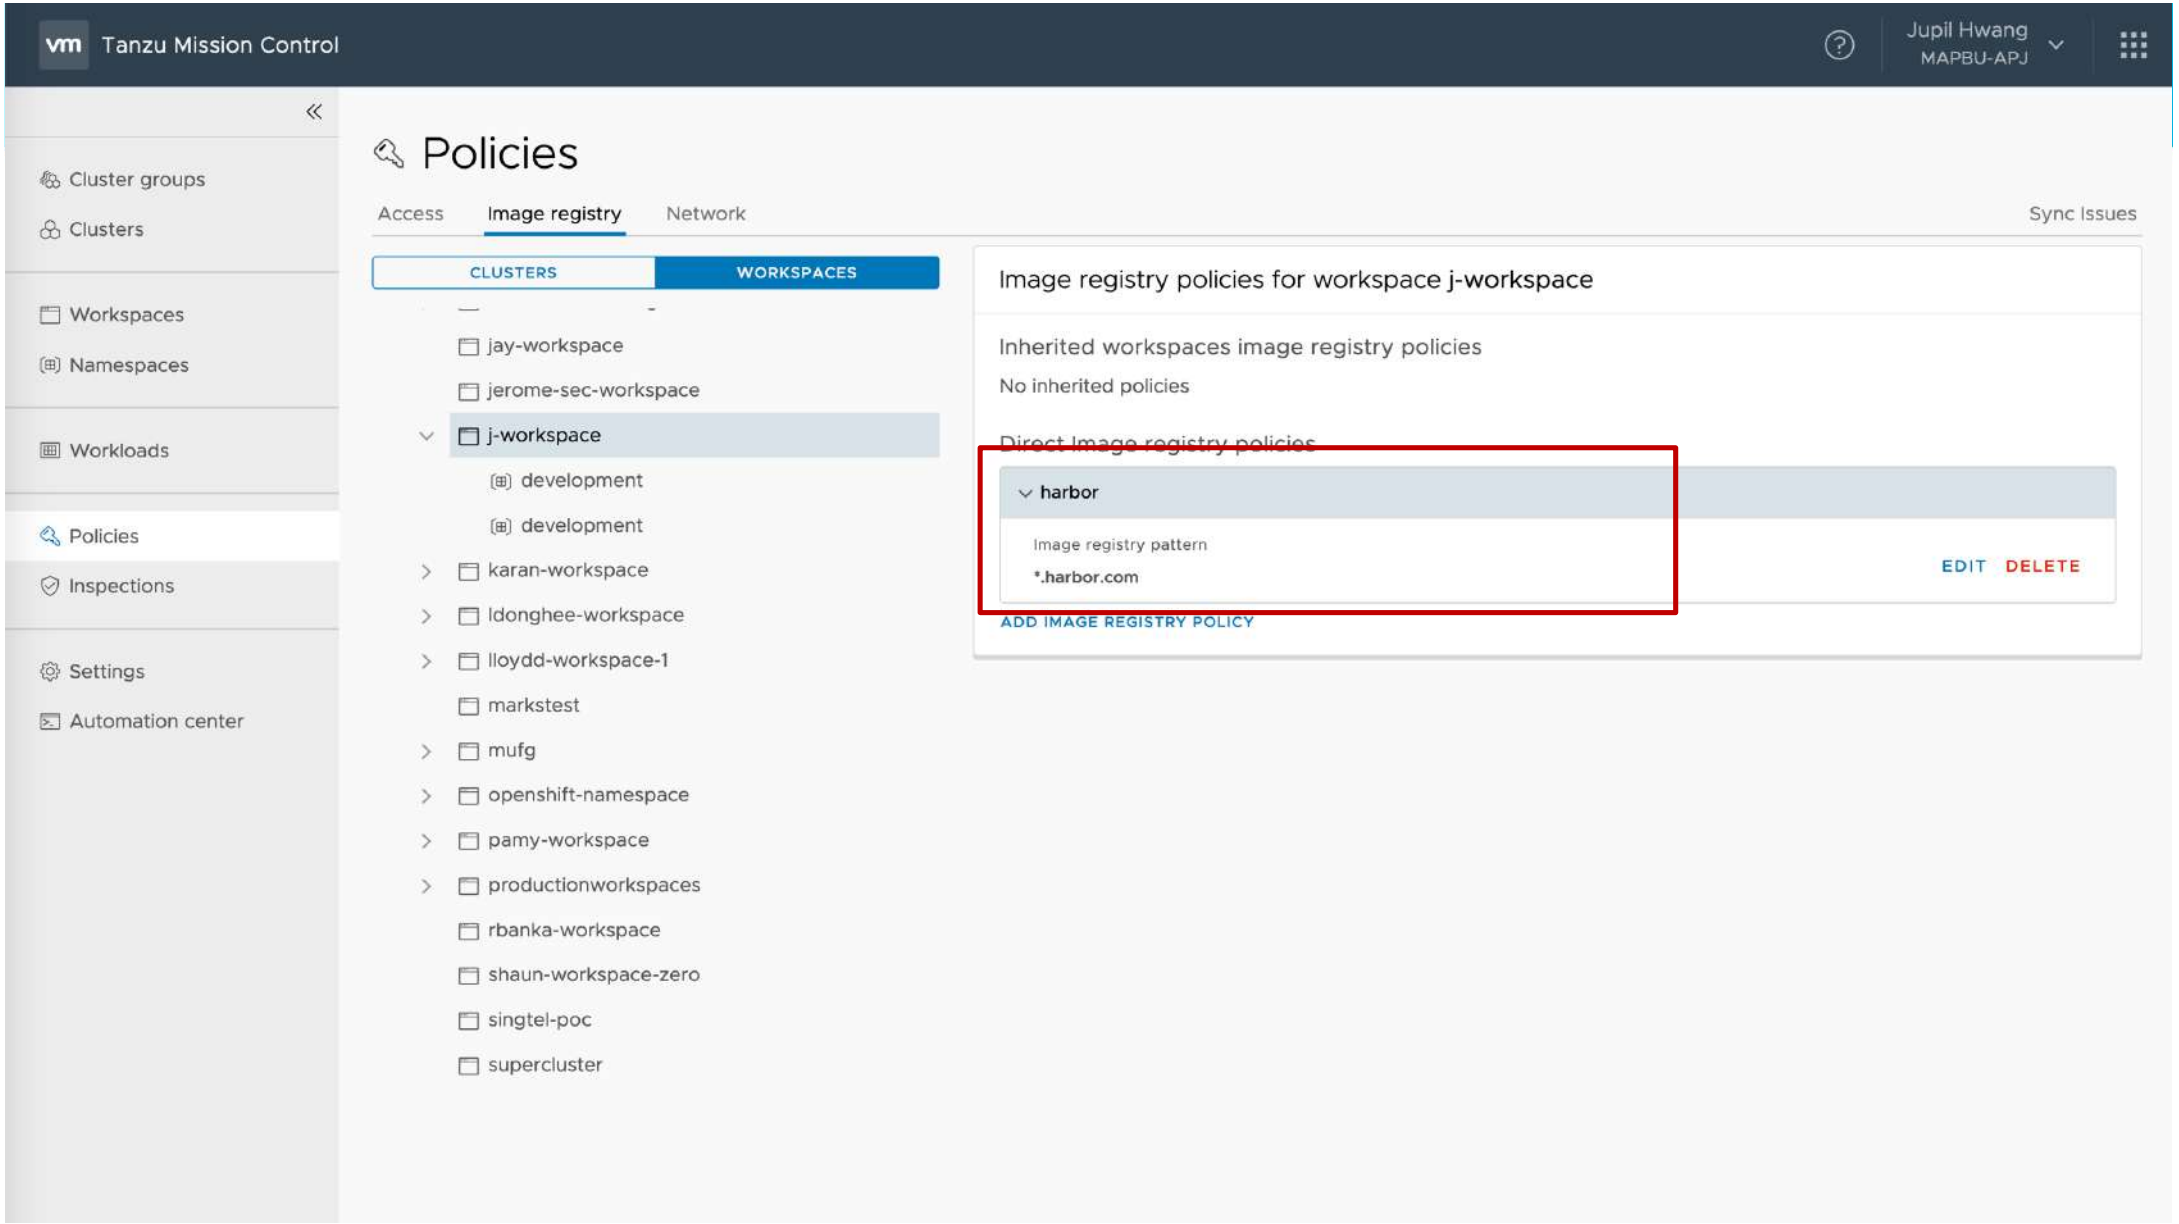
Tanzu Mission Control
Jupil Hwang
MAPBU-APJ
Cluster groups
Clusters
Workspaces
Namespaces
Workloads
Policies
Inspections
Settings
Automation center
Policies
Access
Image registry
Network
Sync Issues
CLUSTERS
WORKSPACES
jay-workspace
jerome-sec-workspace
j-workspace
development
development
karan-workspace
ldonghee-workspace
lloydd-workspace-1
markstest
mufg
openshift-namespace
pamy-workspace
productionworkspaces
rbanka-workspace
shaun-workspace-zero
singtel-poc
supercluster
Image registry policies for workspace j-workspace
Inherited workspaces image registry policies
No inherited policies
Direct Image registry policies
harbor
Image registry pattern
*.harbor.com
EDIT
DELETE
ADD IMAGE REGISTRY POLICY

네트워크 정책

Pod 그룹이 상호 및 다른 네트워크 엔드포인트와 통신할 수 있는 방법 지정

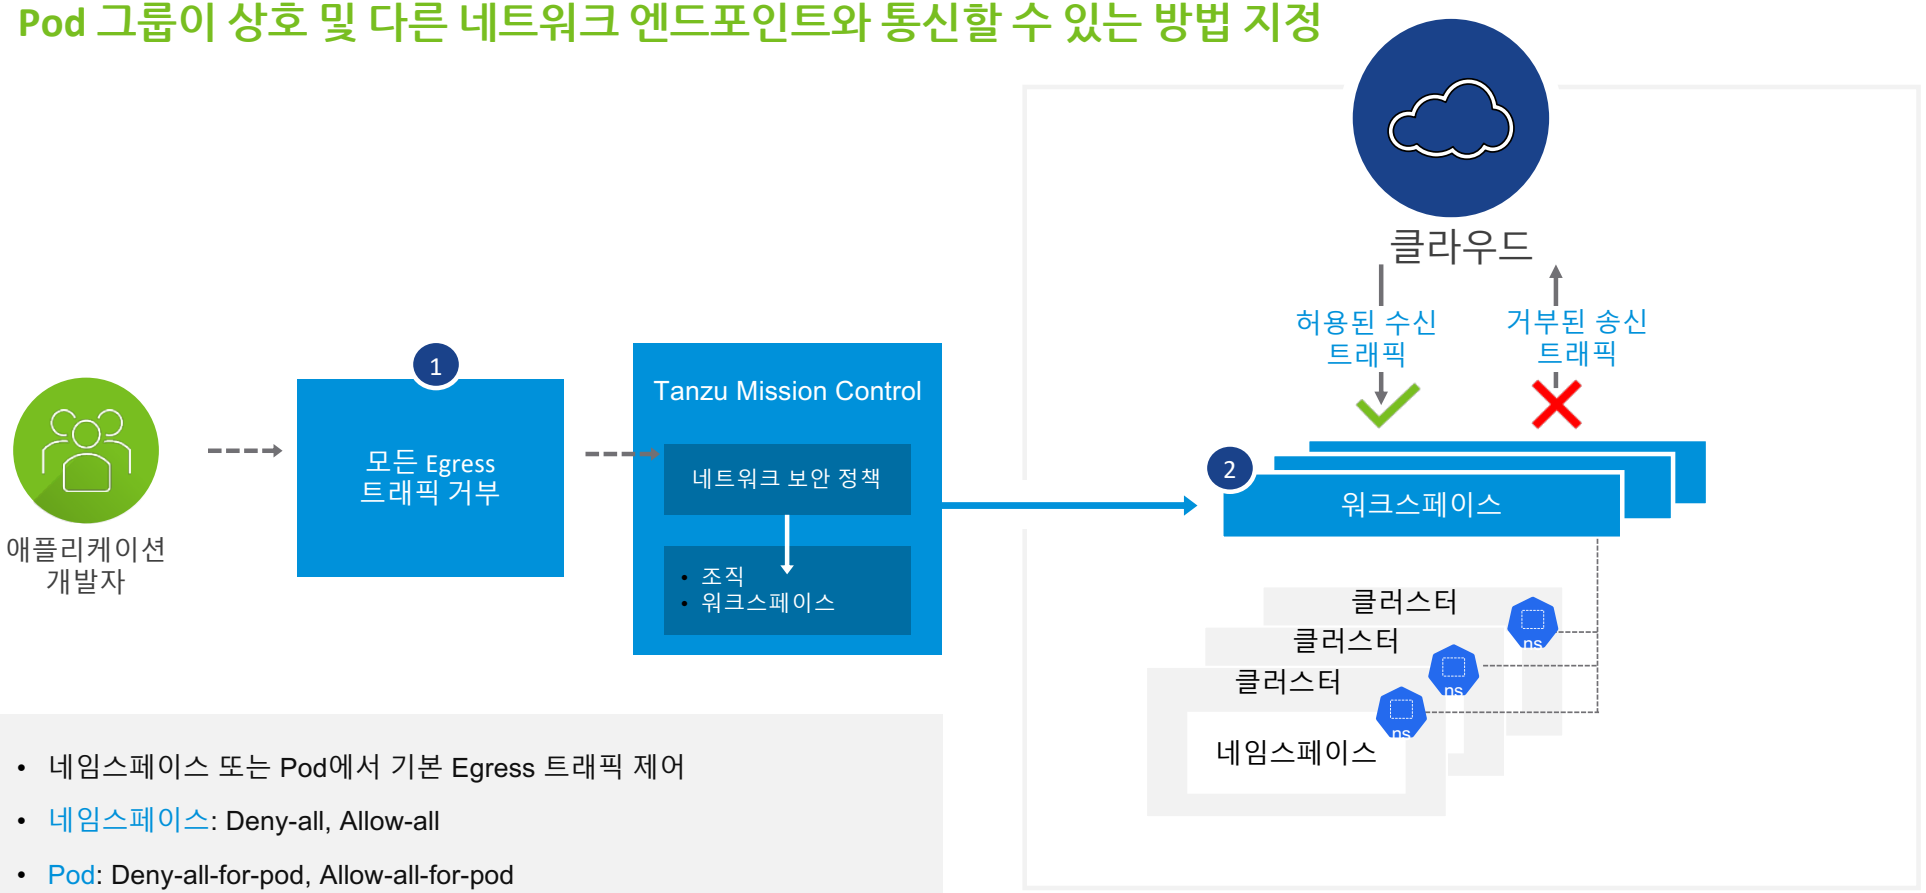

- 네임스페이스 또는 Pod에서 기본 Egress 트래픽 제어
- 네임스페이스: Deny-all, Allow-all
- Pod: Deny-all-for-pod, Allow-all-for-pod

감사 / 컴플라이언스 준수 검사

미리 설정된 조건에 따라 특정 시점의 클러스터 상태를 보고

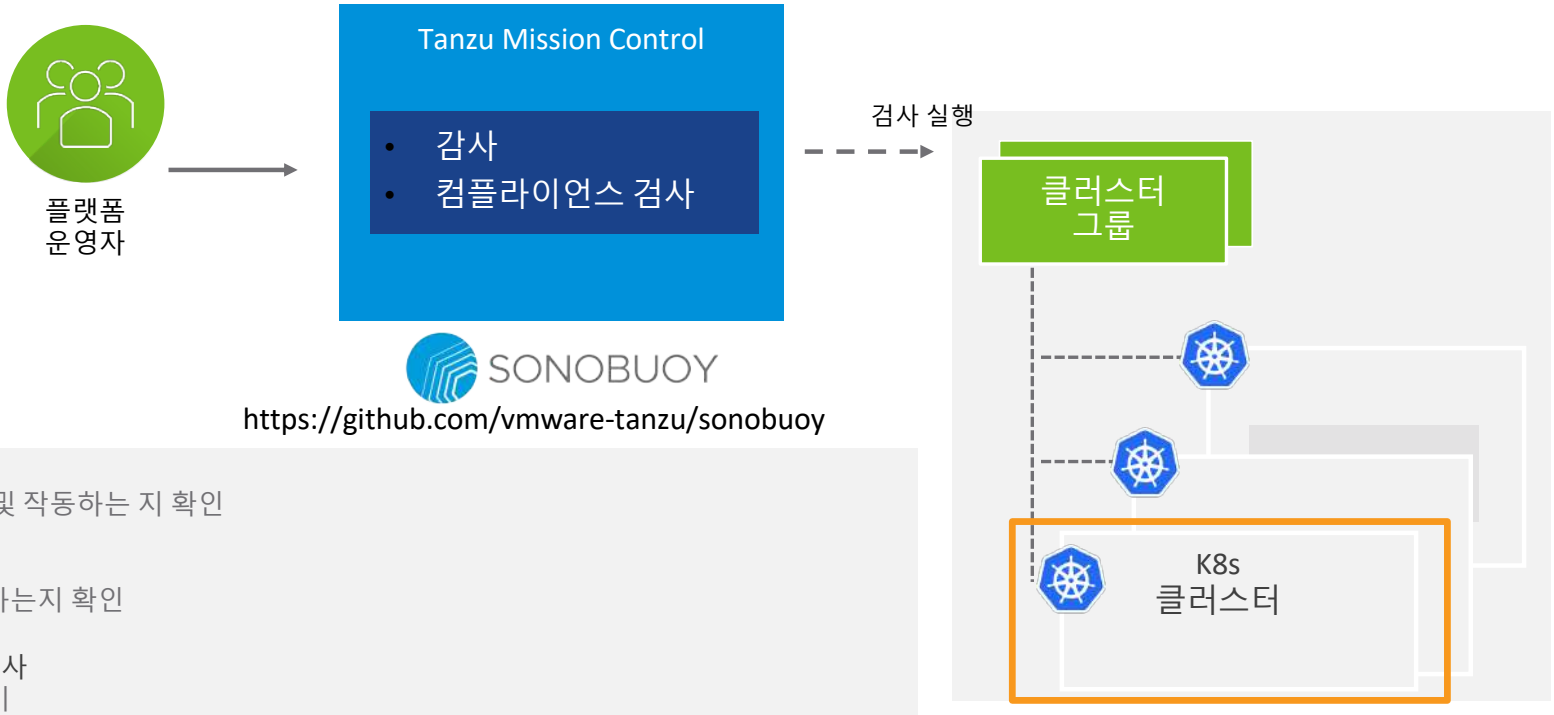

검사 유형

Conformance
- 대상 K8s 클러스터가 제대로 설치 구성 및 작동하는 지 확인

Lite
- 노드의 설정이 K8s 의 요구사항을 충족하는지 확인

오픈소스 Sonobuoy를 확장하여 클러스터 검사
- 전체 시간 경과에 대한 결과 저장 및 보기
- 검사 정책/스케줄
- 시간의 경과에 따른 변화

데이터 보호 정책

전체 조직(Org), 클러스터 그룹, 워크스페이스 단위의 백업 정책 적용

- Kubernetes 개체 & 볼륨 백업
- 볼륨, 네임스페이스 복원
- 클라우드 간 백업

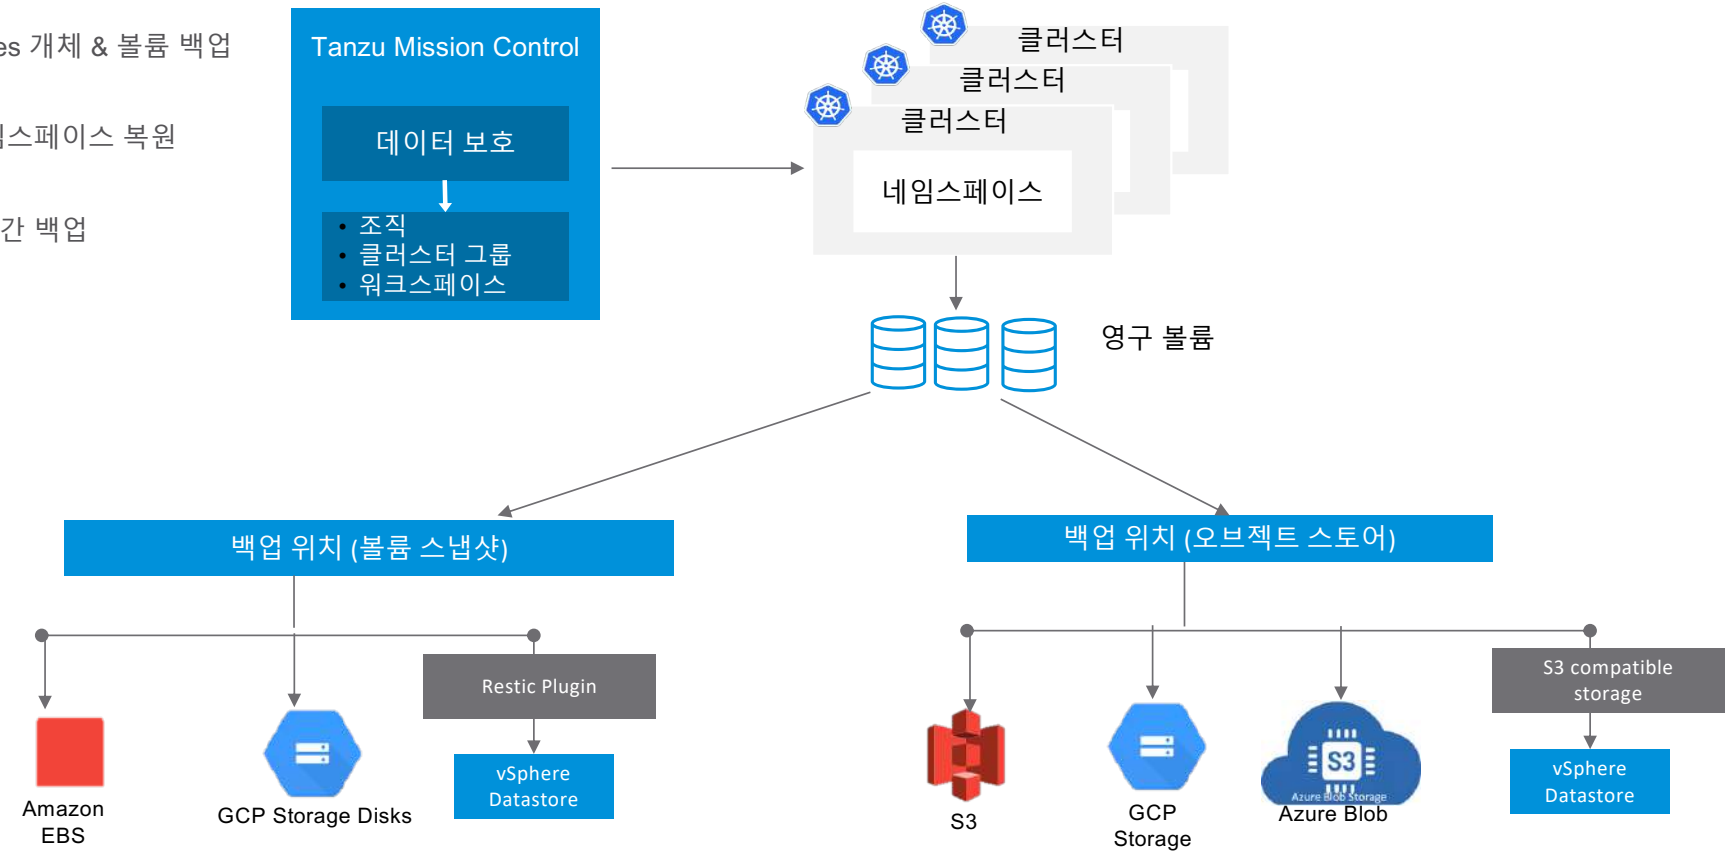

Cloud-Native Day
2020 Korea
LIVE!

다양한 환경에 분산된 워크로드 서비스 메시

Tanzu Service Mesh

vmware

멀티 클러스터 – 마이크로 서비스 메시

멀티 플랫폼
멀티 클라우드

중앙 집중식 가시성 과 진단

사용자, 서비스 및 데이터에
대한 글로벌 정책

중앙 집중식 보안, 감사 및
규정 준수

애플리케이션 변경 없음

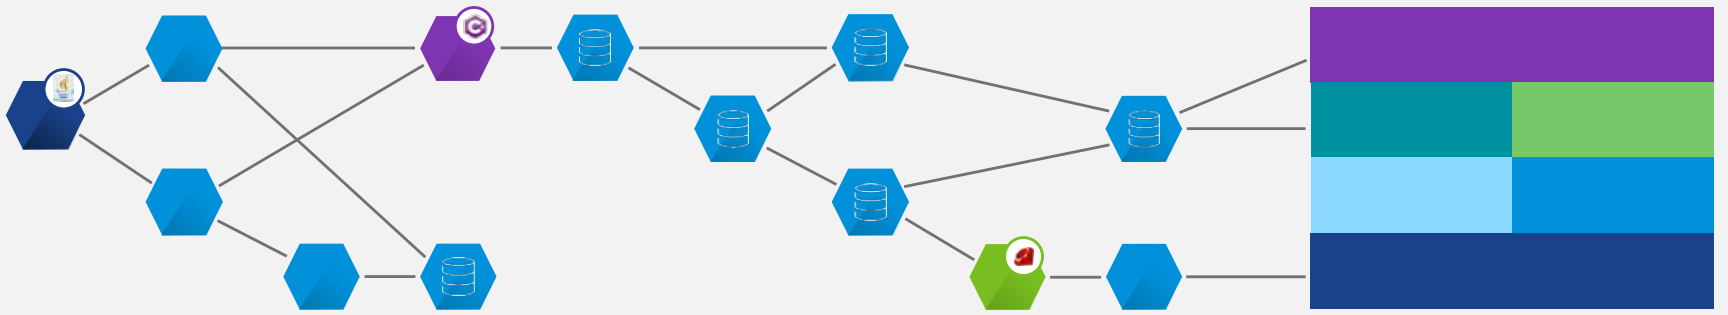

Kubernetes

Public Clouds

VMs

SaaS

Serverless

서비스 메시 - Istio


Services
Istio
Kubernetes

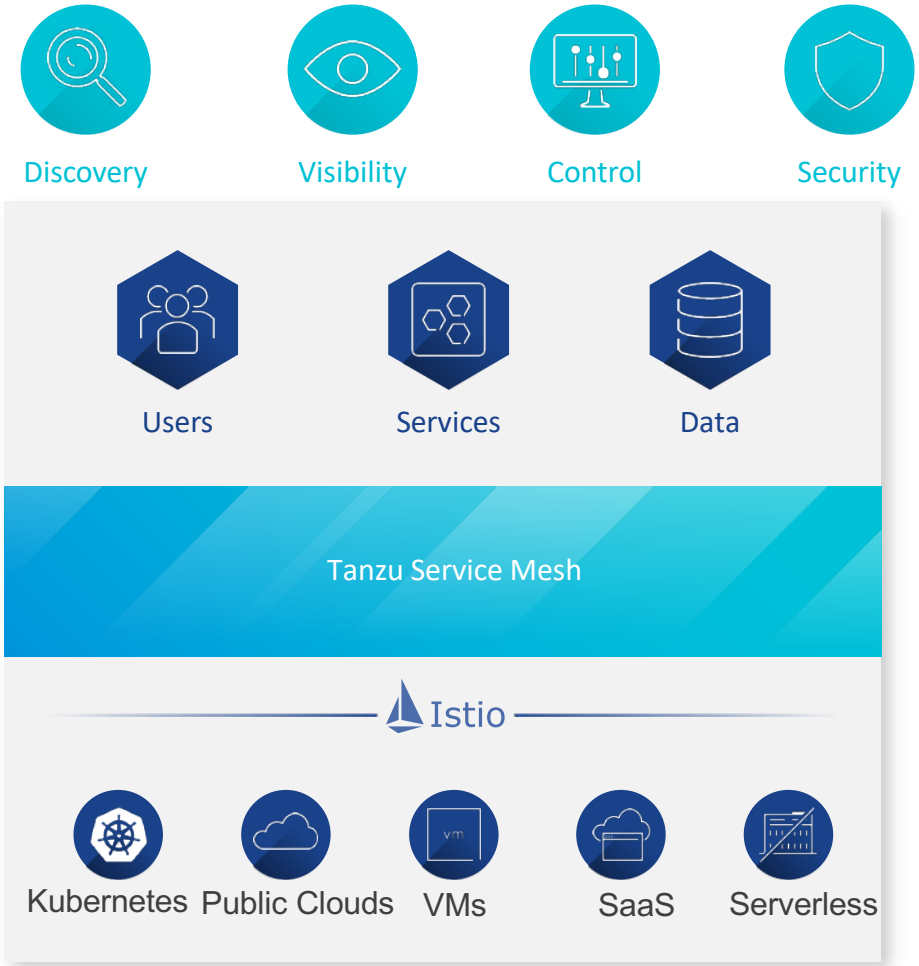
Discovery
Visibility
Control
Security
Users
Services
Data
Tanzu Service Mesh
Istio
Kubernetes
Public Clouds
VMs
SaaS
Serverless

Discover, Observe, Connect and Secure

Apps wherever they are deployed across Kubernetes clusters

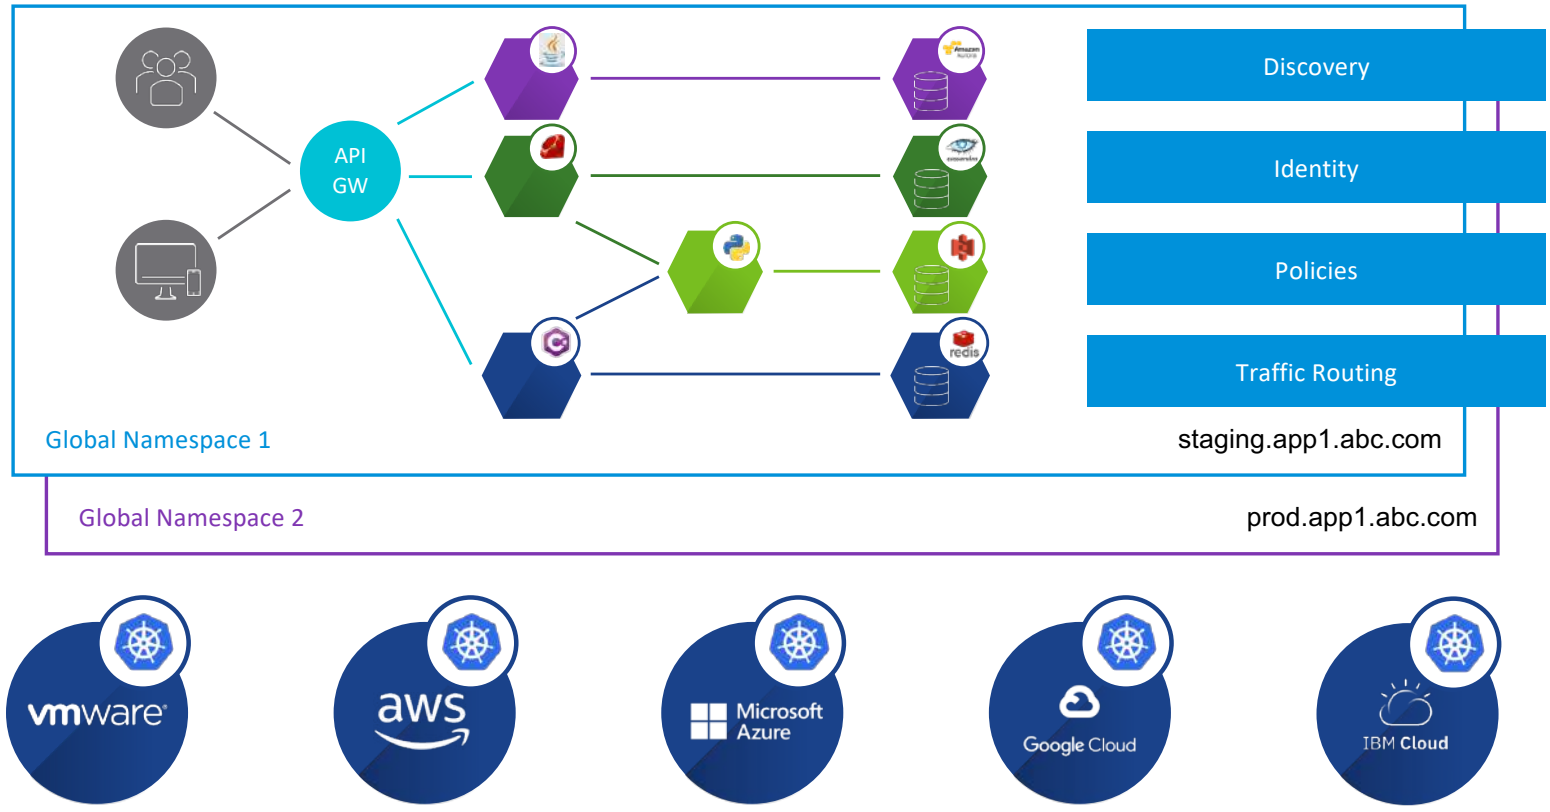

Global Namespace

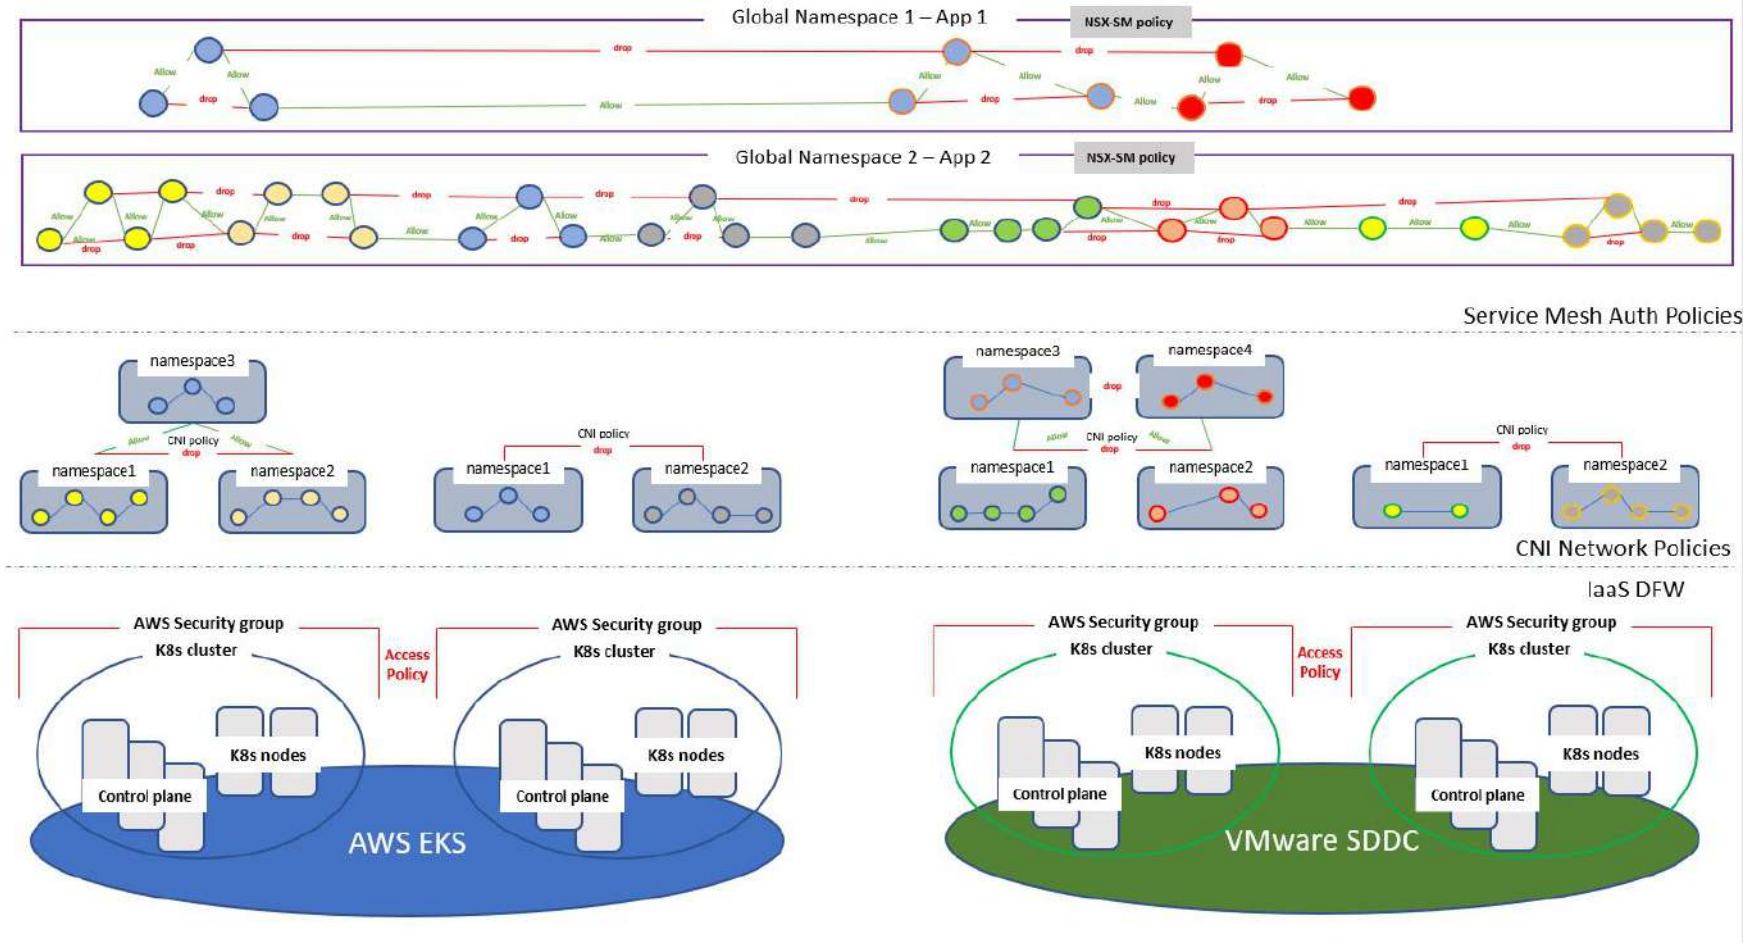

Apply Tanzu Service Mesh to Cluster

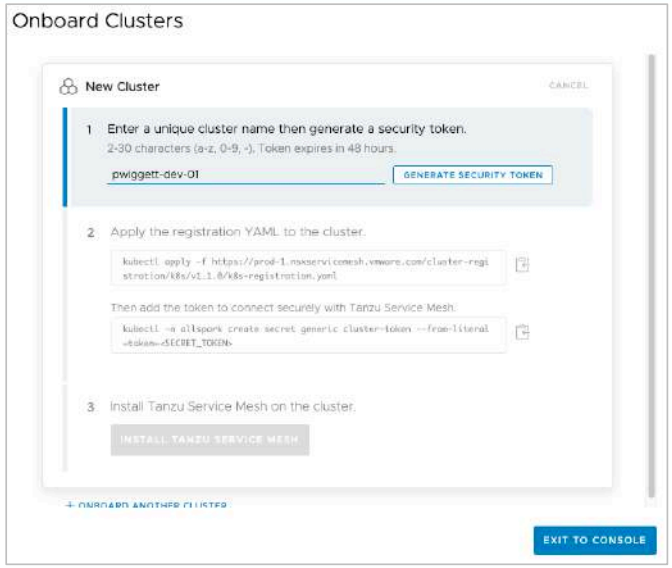

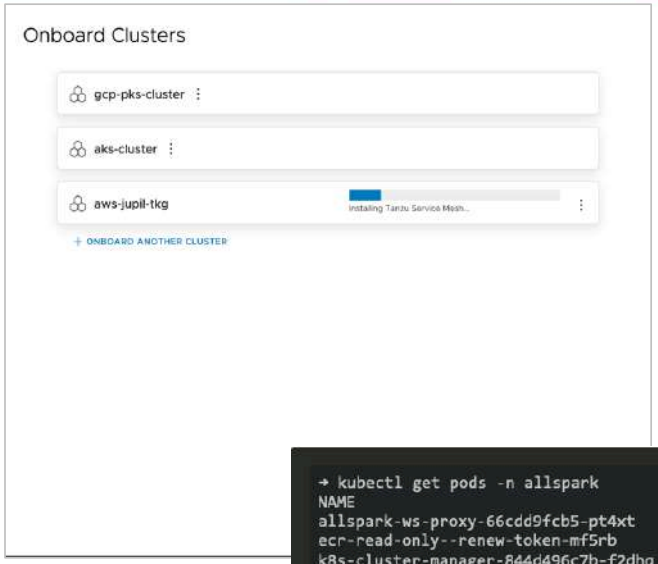


```
→ kubectl get pods -n allspark
NAME                                  READY   STATUS      RESTARTS   AGE
allspark-ws-proxy-66cdd9fcb5-pt4xt    1/1     Running     0          13m
ecr-read-only--renew-token-mf5rb      0/1     Completed   0          13m
k8s-cluster-manager-844d496c7b-f2dhq  1/1     Running     0          13m
telegraf-istio-79b6c4dfd4-7hp5l       1/1     Running     0          7m44s

→ kubectl get pods --n istio-system
NAME                                                      READY   STATUS      RESTARTS   AGE
allspark-telegraf-node-675vd                              1/1     Running     0          6m18s
allspark-telegraf-node-8qbcn                              1/1     Running     0          6m18s
allspark-telegraf-node-qsxn9                              1/1     Running     0          6m18s
istio-citadel-cb8fcd9d7-m9z44                             1/1     Running     0          5m55s
istio-egressgateway-688898bd66-f6979                      1/1     Running     0          5m59s
istio-galley-d7485ccfc-kqm7x                              1/1     Running     0          6m
istio-ingressgateway-db5ff6469-2vcg8                      1/1     Running     0          5m59s
istio-init-crd-10-release-1.3-latest-daily-646vm          0/1     Completed   0          6m22s
istio-init-crd-11-release-1.3-latest-daily-mpgcm          0/1     Completed   0          6m22s
istio-init-crd-12-release-1.3-latest-daily-tkbrj          0/1     Completed   0          6m22s
istio-pilot-69f45cf765-lgh58                              2/2     Running     0          5m56s
istio-policy-59776b88[redacted]-5lj8t                     2/2     Running     2          5m58s
istio-security-post-install-release-1.3-2019-11-27-14-59-cgj8q2   0/1     Completed   0          6m11s
istio-sidecar-injector-58598d54d5-rkzmf                   1/1     Running     0          5m54s
istio-telemetry-54ddcdcdc6-w5br6                          2/2     Running     2          5m57s
istiocoredns-576dd6459c-l559k                             2/2     Running     0          5m59s

→
```

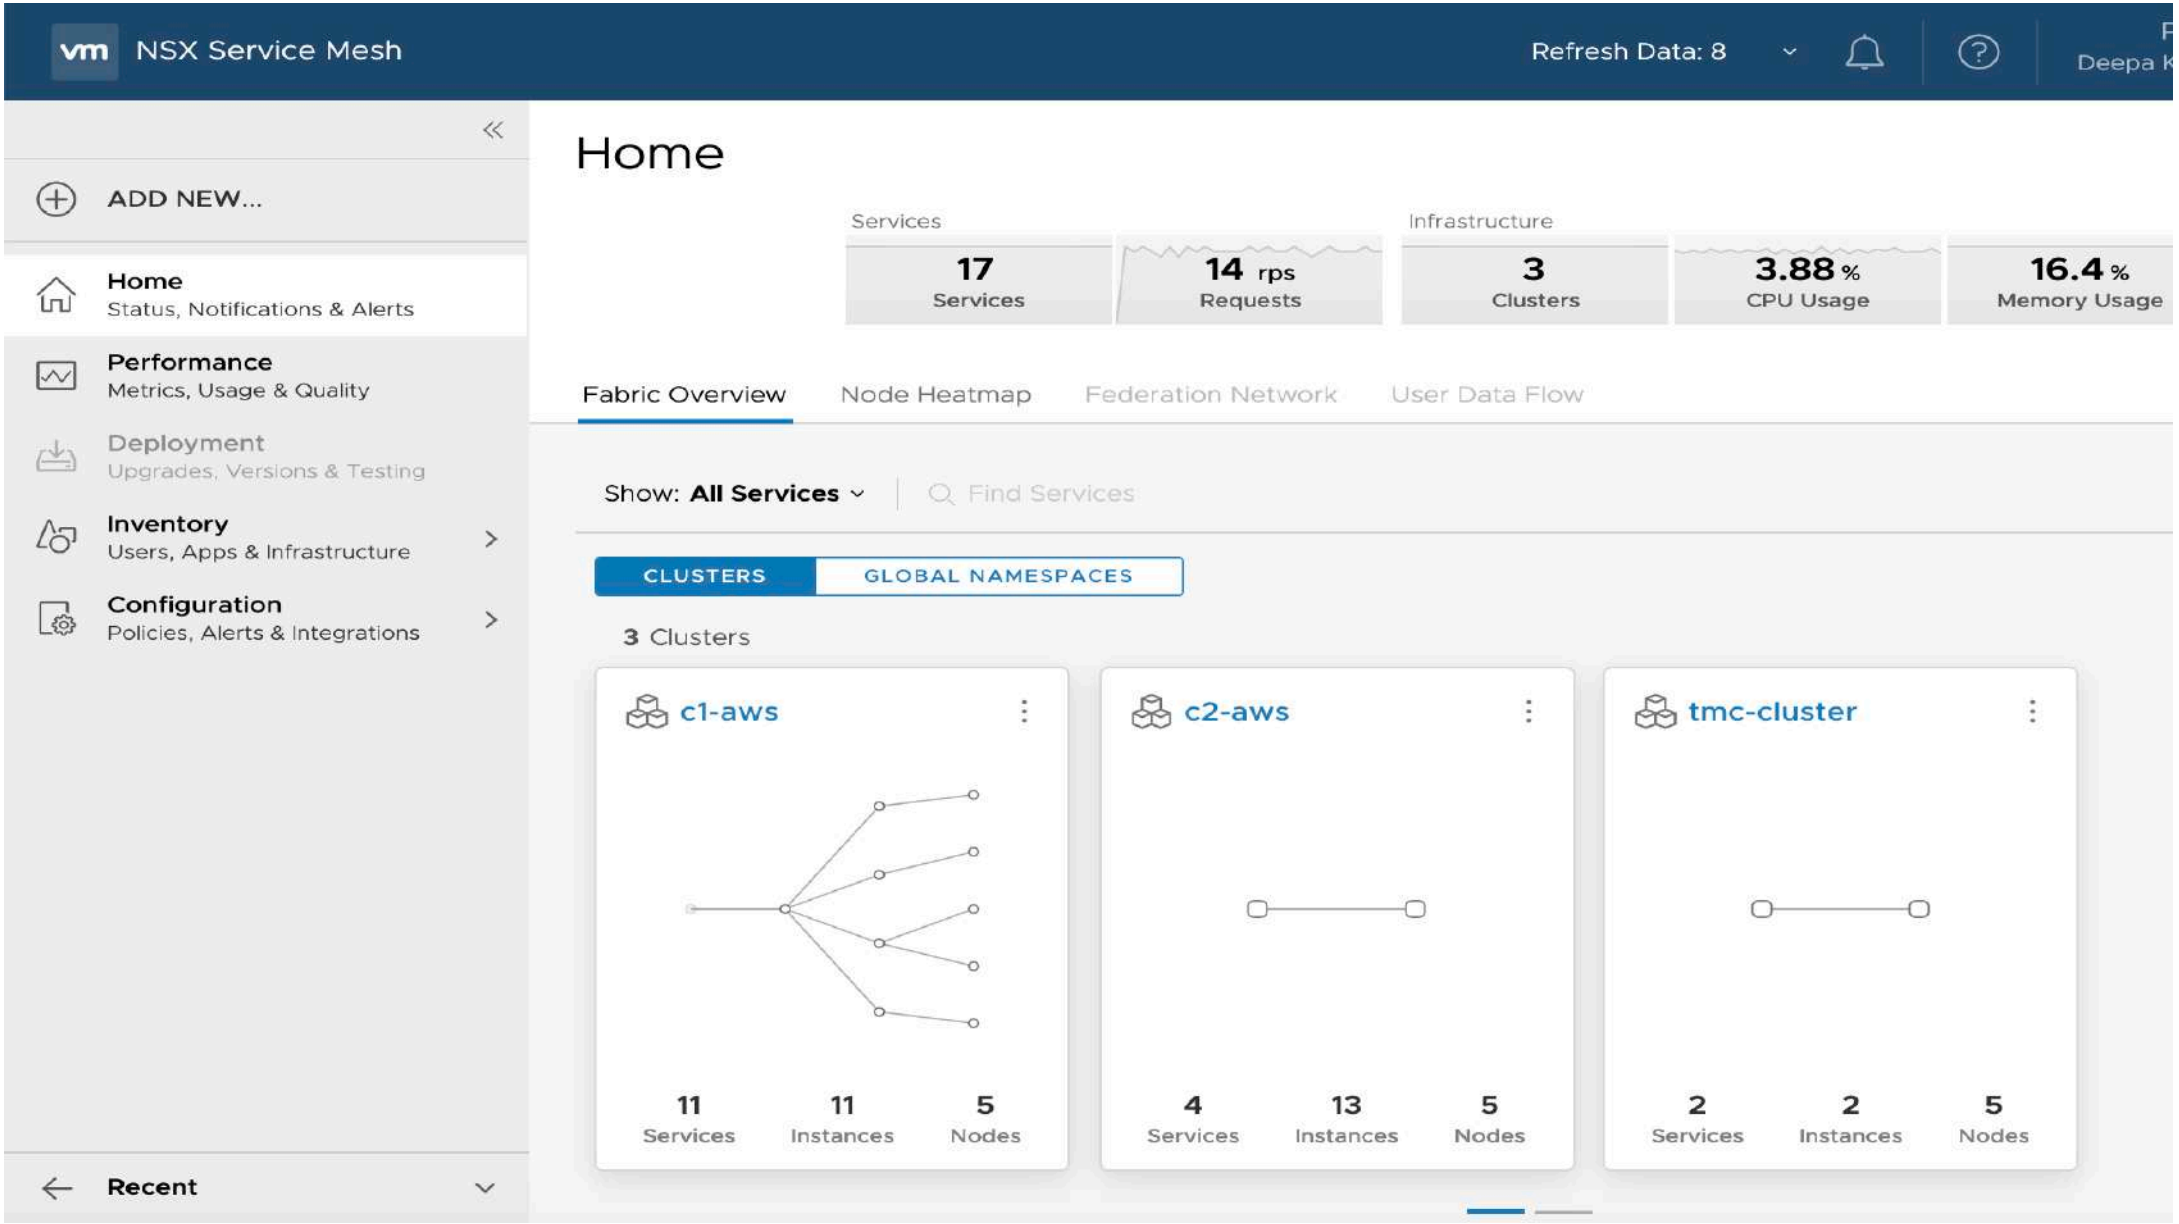
NSX Service Mesh
Refresh Data: 8
Deepa K
ADD NEW...
Home
Status, Notifications & Alerts
Performance
Metrics, Usage & Quality
Deployment
Upgrades, Versions & Testing
Inventory
Users, Apps & Infrastructure
Configuration
Policies, Alerts & Integrations
Recent
Home
Services
17
Services
14 rps
Requests
Infrastructure
3
Clusters
3.88 %
CPU Usage
16.4 %
Memory Usage
Fabric Overview
Node Heatmap
Federation Network
User Data Flow
Show: All Services
Find Services
CLUSTERS
GLOBAL NAMESPACES
3 Clusters
c1-aws
11
Services
11
Instances
5
Nodes
c2-aws
4
Services
13
Instances
5
Nodes
tmc-cluster
2
Services
2
Instances
5
Nodes

pwiggett-dev-01 Requests: **0.11 rps** p99 Latency: **80 ms** Errors: **0.75 %**

REMOVE CLUSTER

Cluster pwiggett-dev-01	**Platform** Kubernetes	**Platform Version** v1.17.3+vmware.1
CPU Capacity (Cores) 6	**Memory Capacity** 22.5 GB	**Disk Capacity** 150 GB
Services 3	**Service Instances** 6	**Nodes** 3

Service Topology | Performance | Services | Service Instances | Nodes | System Details | Users | Data | Configuration

Show: **All Services** | Find Services | SETTINGS | Data: **Last 1 hour**

pwiggett-dev-01 **3** Connected Services

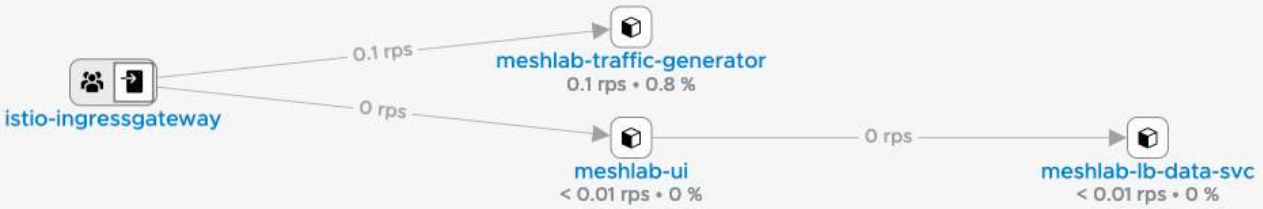

pwiggett-dev-01 Requests: **0.45 rps** p99 Latency: **93 ms** Errors: **0 %**

REMOVE CLUSTER

Cluster
pwiggett-dev-01

CPU Capacity (Cores)
6

Services
3

Platform
Kubernetes

Memory Capacity
22.5 GB

Service Instances
6

Platform Version
v1.17.3+vmware.1

Disk Capacity
150 GB

Nodes
3

Service Topology | Performance | Services | Service Instances | Nodes | System Details | Users | Data | Configuration

CHART SETTINGS | Data: **Last 5 minutes**

REQUESTS

REQUEST COUNT

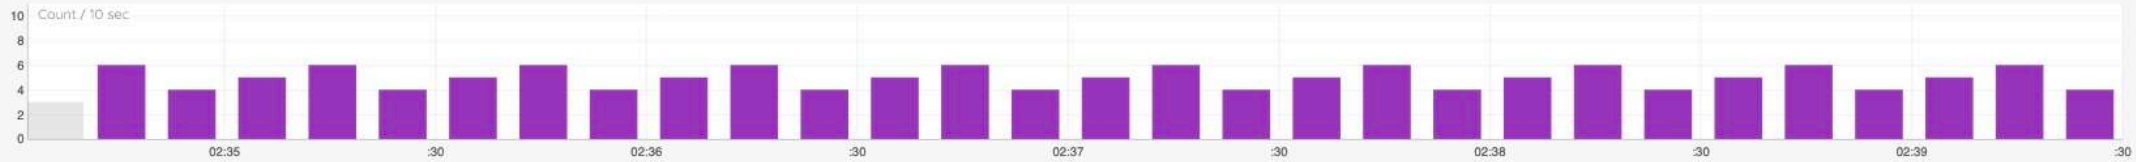

P50 LATENCY


```
kubo@jumper:~$ k get serviceentries
NAME                          HOSTS                              LOCATION        RESOLUTION   AGE
nsxsm.gns.acme.cart           [cart.acme.local]                  MESH_INTERNAL   DNS          8d
nsxsm.gns.acme.cart-redis     [cart-redis.acme.local]            MESH_INTERNAL   DNS          8d
nsxsm.gns.acme.catalog-b      [catalog-b.acme.local]             MESH_INTERNAL   DNS          8d
nsxsm.gns.acme.catalog-mongo  [catalog-mongo.acme.local]         MESH_INTERNAL   DNS          8d
nsxsm.gns.acme.kubernetes     [kubernetes.acme.local]            MESH_INTERNAL   DNS          8d
nsxsm.gns.acme.order          [order.acme.local]                 MESH_INTERNAL   DNS          8d
nsxsm.gns.acme.order-mongo    [order-mongo.acme.local]           MESH_INTERNAL   DNS          8d
nsxsm.gns.acme.payment        [payment.acme.local]               MESH_INTERNAL   DNS          8d
nsxsm.gns.acme.shopping       [shopping.acme.local]              MESH_INTERNAL   DNS          8d
nsxsm.gns.acme.users          [users.acme.local]                 MESH_INTERNAL   DNS          8d
nsxsm.gns.acme.users-mongo    [users-mongo.acme.local]           MESH_INTERNAL   DNS          8d
kubo@jumper:~$ k edit serviceentry nsxsm.gns.acme.catalog-b
Edit cancelled, no changes made.
```

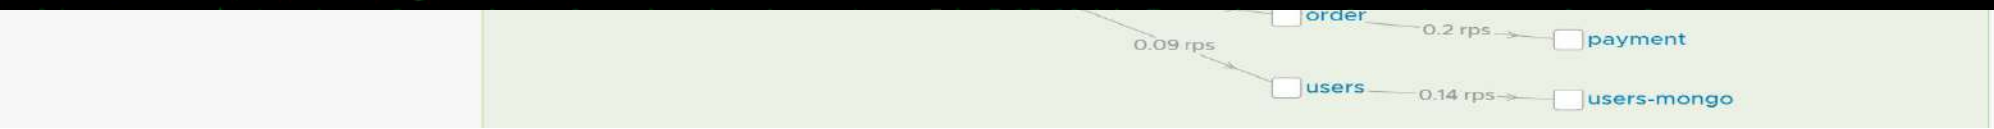

Cloud-Native Day
2020 Korea
LIVE!

보이지 않는 것은 알 수 없고 알 지 못하면 고칠 수 없다.

Tanzu Observability by Wavefront

vmware

멀티 클라우드 애플리케이션은 심각한 Day 2 Ops 문제를 발생

완전히 새로운 형태의 모니터링 및 분석 기능이 필요

Effect

미리 숨겨진 문제 발견 → 신뢰성, 가용성, 고객 경험 → 지속 불가능한 수동 작업 → 통제되지 않는 비용

Cause

추적할 오브젝트, 서비스, 컨테이너 및 애플리케이션 요소의 폭발적인 증가

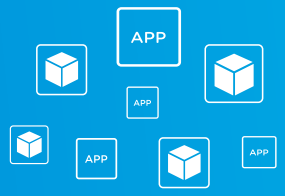

사일로화 된 도구들, 수많은 경고, 제한된 가시성, 지속적 제공 과 자율적인 팀

New World

하이브리드, 네이티브 퍼블릭 클라우드 환경에서 실행되는 현대화된 애플리케이션

vmware Cloud Foundation | vmware TKG | aws | Azure | Google Cloud

Monitoring vs. Observability

Tanzu Observability는 Monitoring이 아닌 복잡한 환경의 모던 애플리케이션을 위한 적극적이며 빠른 Observability 솔루션

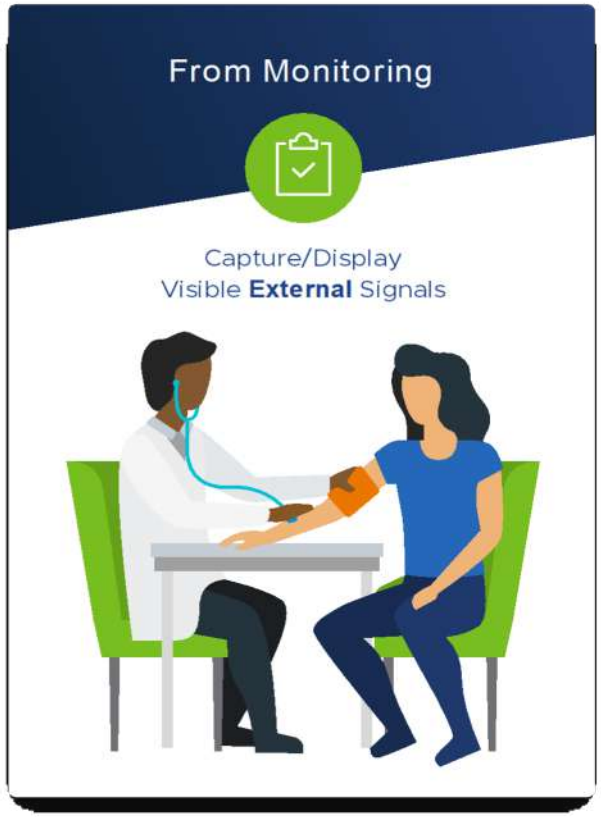

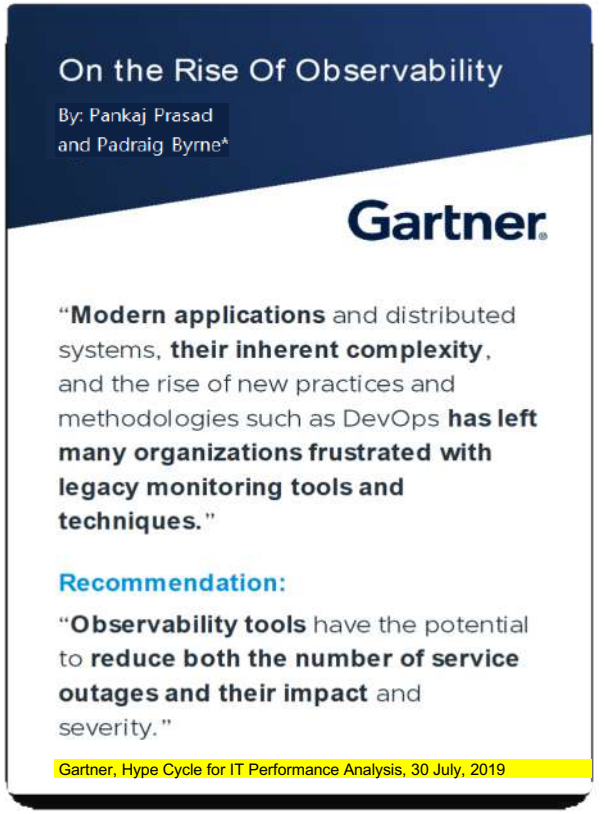

※ Observability : 관찰성, 시스템의 출력 변수(Output variable)를 사용하여 상태 변수(State variable)에 대한 정보를 알아낼 수 있는지를 나타내는 용어

Tanzu Observability by Wavefront 아키텍처

AI 및 분석을 기반으로 한 기능

Ecosystems

200개 이상의 Out of the box 통합으로 확장된 기회와 가치에 대한 새로운 데이터 소스 추가

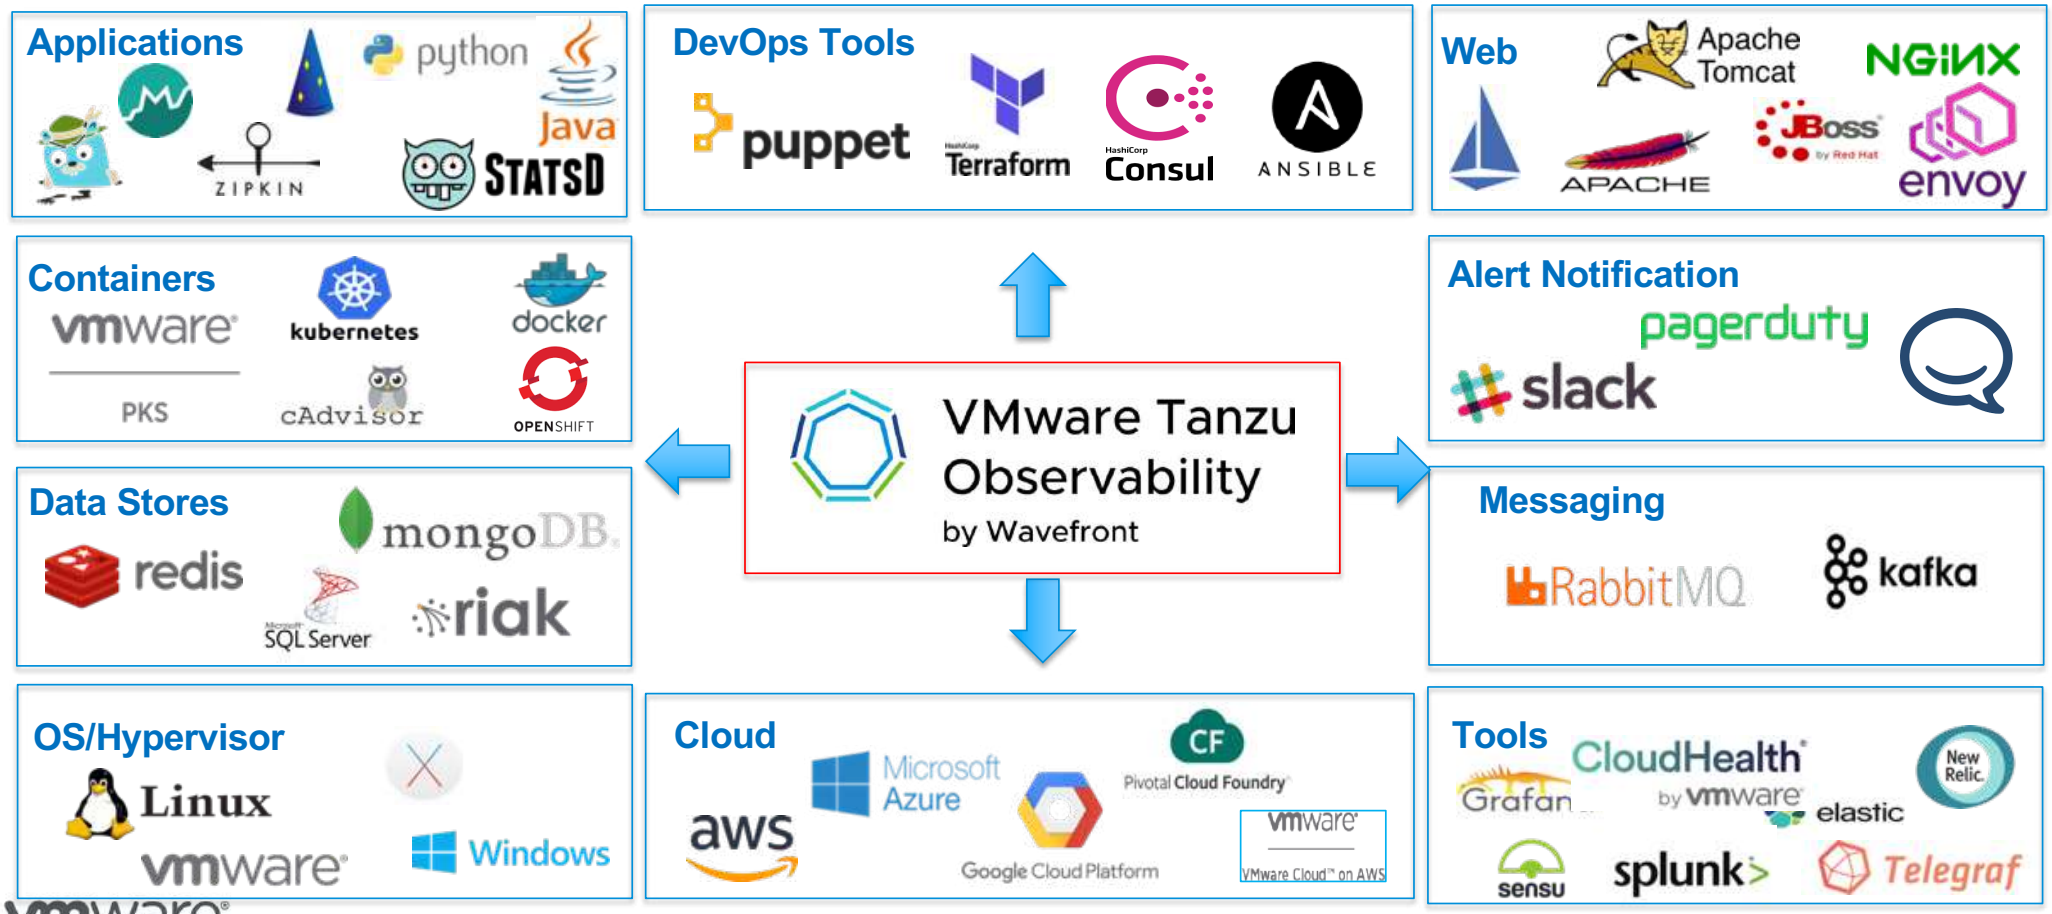

www.wavefront.com/integrations

대시보드

Dashboard 예제

쿠버네티스 대시보드

컨테이너화된 애플리케이션 가시성

Kubernetes Health 모니터링

자원 소비

프로그래밍 방식 알림

AWS EC2 대시보드

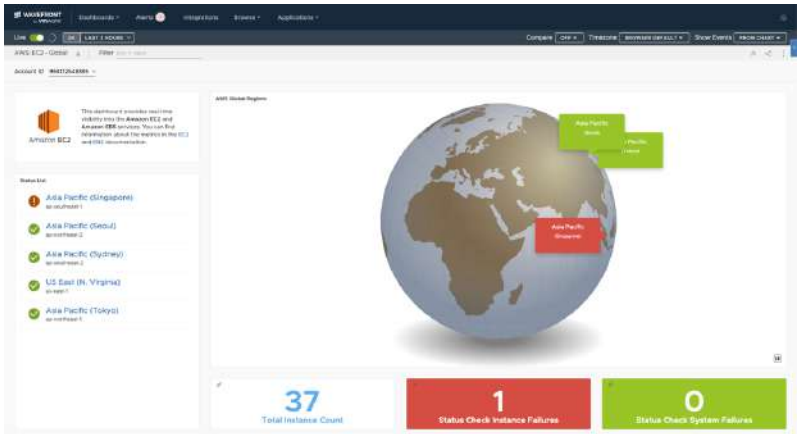

AWS 빌링 대시보드

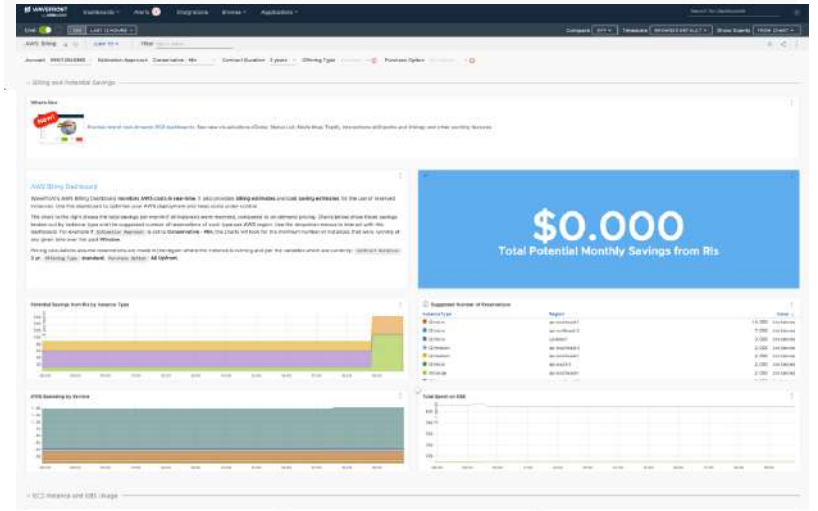

스프링 부트를 위한 Wavefront

Wavefront Spring Boot Starter

pom.xml

```
<dependency>
  <groupId>com.wavefront</groupId>
  <artifactId>wavefront-spring-boot-starter</artifactId>
  <version>2.0.0-RC1</version>
</dependency>

<dependency>
  <groupId>org.springframework.cloud</groupId>
  <artifactId>spring-cloud-starter-sleuth</artifactId>
  <version>2.2.2.RELEASE</version>
</dependency>
```

./mvnw spring-boot:run

```
A Wavefront account has been provisioned successfully and the API token has been saved to disk.

To share this account, make sure the following is added to your configuration:

        management.metrics.export.wavefront.api-token=9c355722-f863-4a9f-9058-e612756ce9d3
        management.metrics.export.wavefront.uri=https://wavefront.surf

Connect to your Wavefront dashboard using this one-time use link:
https://wavefront.surf/us/sG3R8kWxtj
```

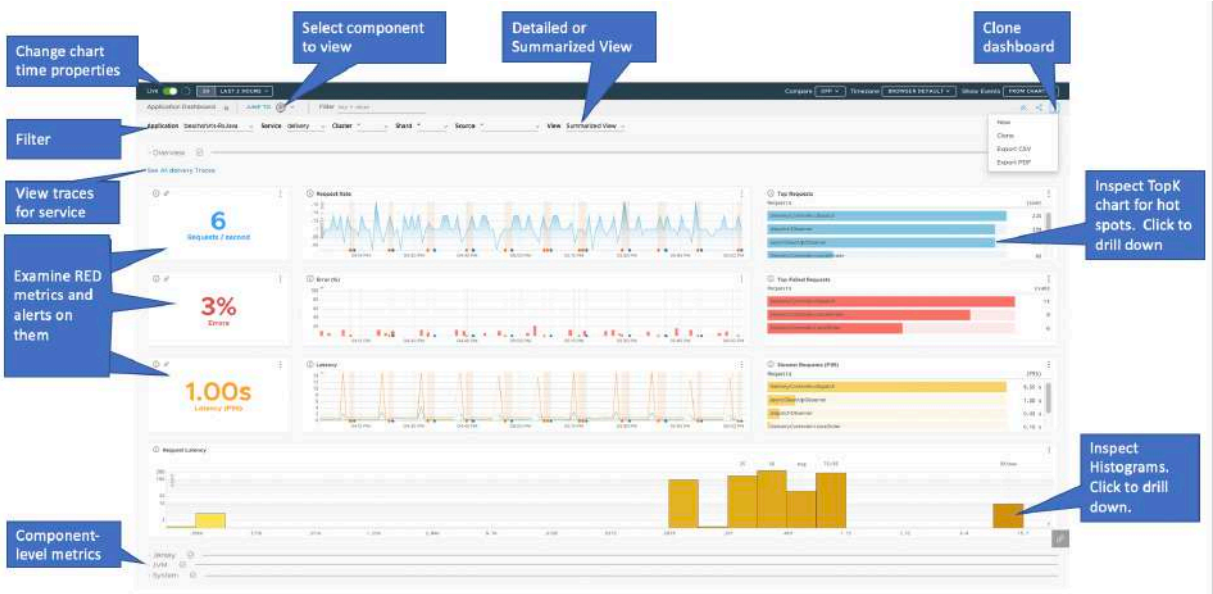

분산 추적

Distributed Tracing

AI Genie

AI Genie 기반 자동 이상 징후 감지 및 예측

Tanzu Observability AI Genie

Chart Editor

Anomaly Detection

Forecasting

손쉬운 사용 및 자동화된 기능

빠른 문제 해결

더욱 간단한 성능 최적화

간단한 지능형 알림

더 간단한 알림 생성	더욱 빠른 문제 해결	성능 최적화를 위한 예측
• 즉시 사용 가능한 이상 징후 및 예측 가시성 • 자동으로 생성된 쿼리를 통한 지능적 알림 • 실제 이상 징후 및 잠재적인 향후 문제 파악	• 모든 서비스 계층의 문제를 신속하게 파악 • 실제 이상 징후 효율적 감지 및 MTTR 감소 • 실시간 운영 중단 과 영향력 높은 인시던트 방지	• 드러난 사용률 패턴에 대한 이해 • 비용 최적화, 애플리케이션과 인프라 효율성 극대화

VMware Tanzu 포트폴리오

Kubernetes 를 중심으로 하는 현대 애플리케이션을 위한 제품 및 서비스

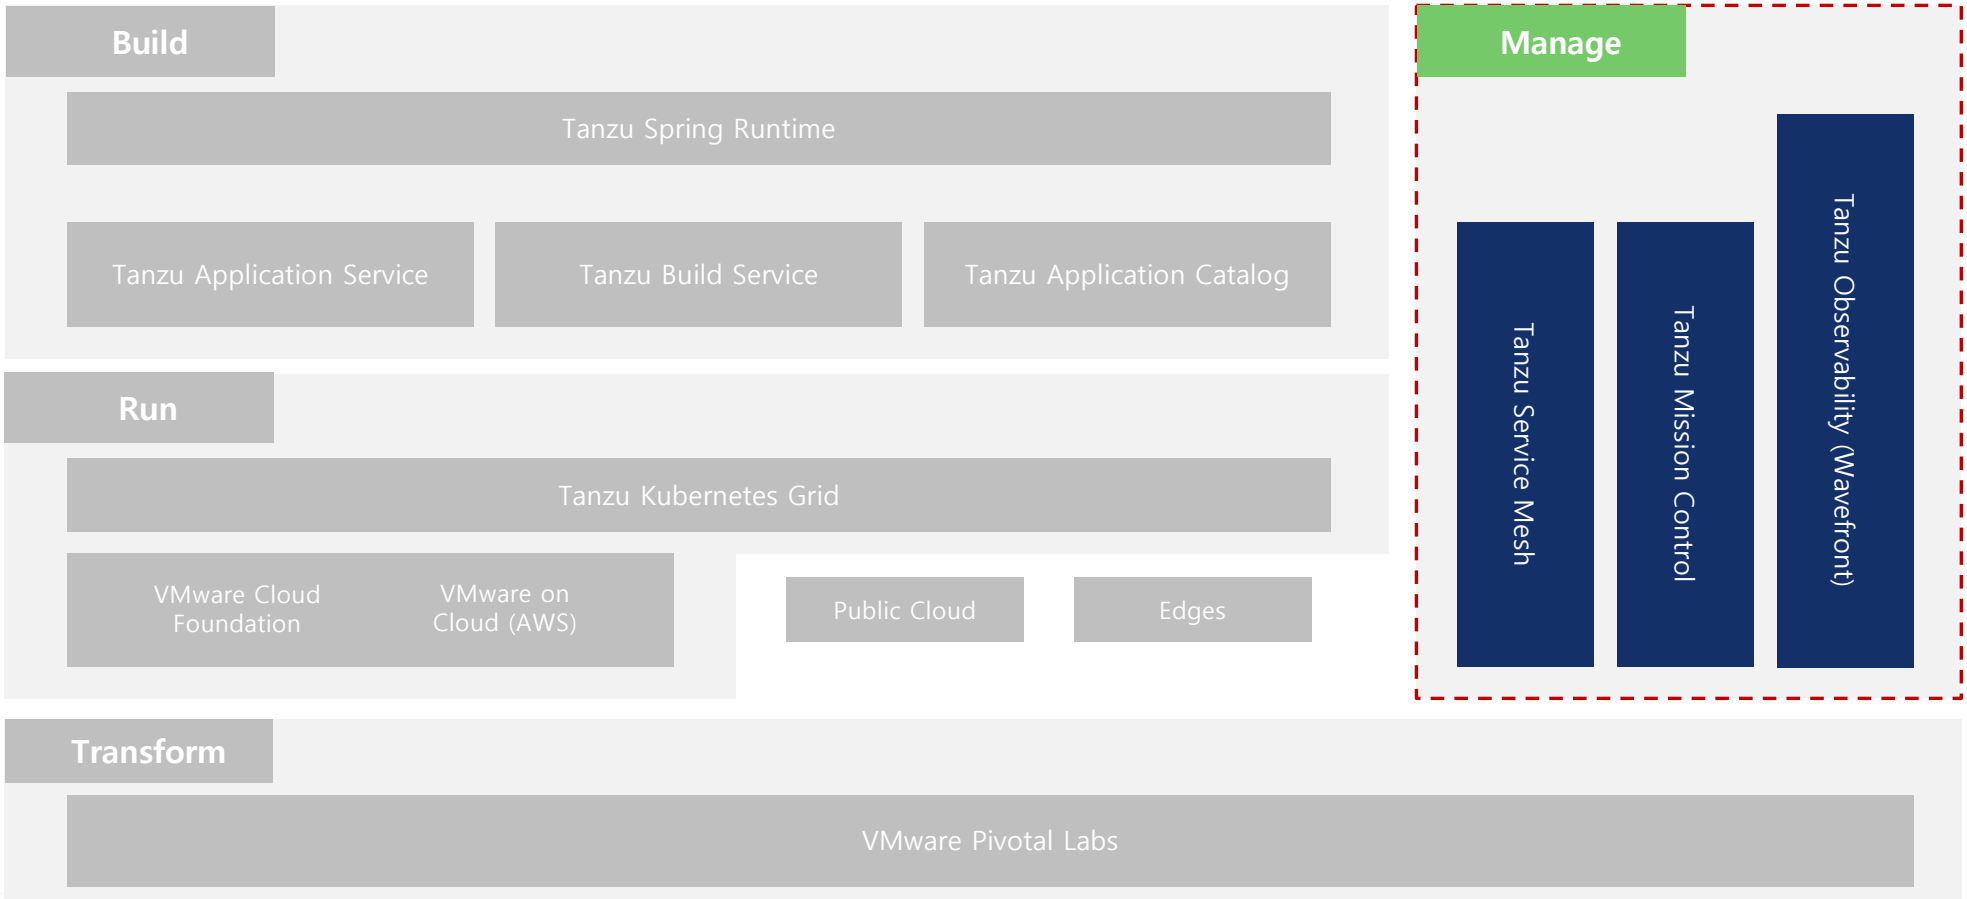

Cloud-Native Day
2020 Korea
LIVE!

고맙습니다.

vmware

Cloud-Native Day
2020 Korea
LIVE!

뉴노멀 비즈니스를 위한 마이크로서비스 전략
Change FASTER

vmware

Today's Sessions

Cloud-Native Day
2020 Korea
LIVE!

1. 마이크로서비스 디자인
2. 마이크로서비스 구축
3. 마이크로서비스 런타임 & 네트워크
4. 마이크로서비스 모니터링 및 관리

vmware

뉴노멀 비즈니스를 위한 마이크로서비스 전략

Modern App & Modern Infrastructure

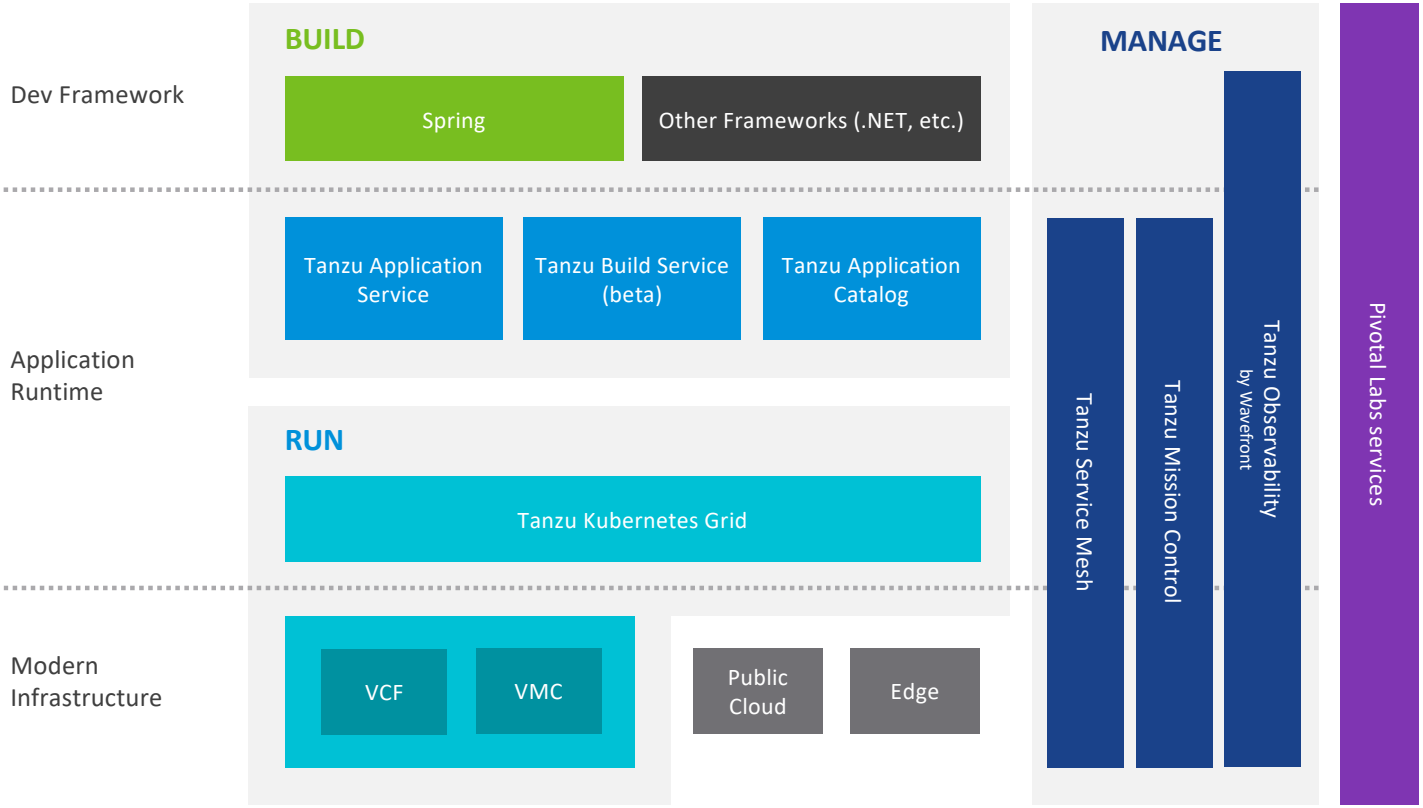

Modern App – 릴리즈 시간 단축

- Spring, .NET, Node.js, 등을 위한 소프트웨어의 출시(Release) 자동화 하는 안전하고 확장 가능한 플랫폼 제공(Modern SW Supply Chain)
- 멀티 클라우드 기업 전반에서 인프라 및 애플리케이션 상태를 관찰하고 분석하기 위한 단일 플랫폼

Modern Infrastructure - 인프라 안정화

- 인프라 및 워크로드 자동화 (Day 1 & 2 Operations)
- SDDC 스택에 K8S Runtime 통합을 통해 VM 및 Container 에 대한 통합 관리 및 가시성 제공
- 멀티 클라우드 환경에서의 일관된 K8S 클러스터 운영 모델 제공

VMware Tanzu Education New Courses

Course
VMware vSphere: Install, Configure, Manage [V7]
The vSphere with Kubernetes: Deploy and Manage [V7]
VMware Cloud Foundation: Management and Operations [V3.9.1]
VMware Tanzu Mission Control: Management and Operations 2020
VMware Tanzu Kubernetes Grid: Install, Configure, Manage v1.0

*모든 교육 문의는 (일정, 금액, 과정설명서) 하기 연락처로 부탁드립니다.

임푸르뫼과장
pim@vmware.com
010-6208-6049

고맙습니다.