



Enabling Elastic Cloud Integration with Kubernetes

Tim Whitcomb


NRL Marine Meteorology Division, Monterey, CA

U.S. Naval Research Laboratory
Advancing research further than you can imagine ®

**20th ECMWF Workshop on HPC in Meteorology
Bologna, Italy
October 2023**

Distribution Statement A: Approved for public release. Distribution is unlimited.

How did we get here?



A Successful Public/Private Partnership

- Continuity of operations proposal, project management, and funding for labor from Fleet Numerical Meteorology and Oceanography Center (FNMOC)
- Funds for cloud computing from Navy Commercial Cloud Services (NCCS; PMW-270)
- Cloud brokerage and account support from NAVFAC Information Technology Center
- Azure solution architects and HPC subject matter experts from Microsoft
- Demonstration design, implementation, and execution by Naval Research Laboratory (NRL)

FY20 Demonstration project: Real-Time COAMPS-TC
Ensemble with Hybrid GFS/NAVGEM forcing

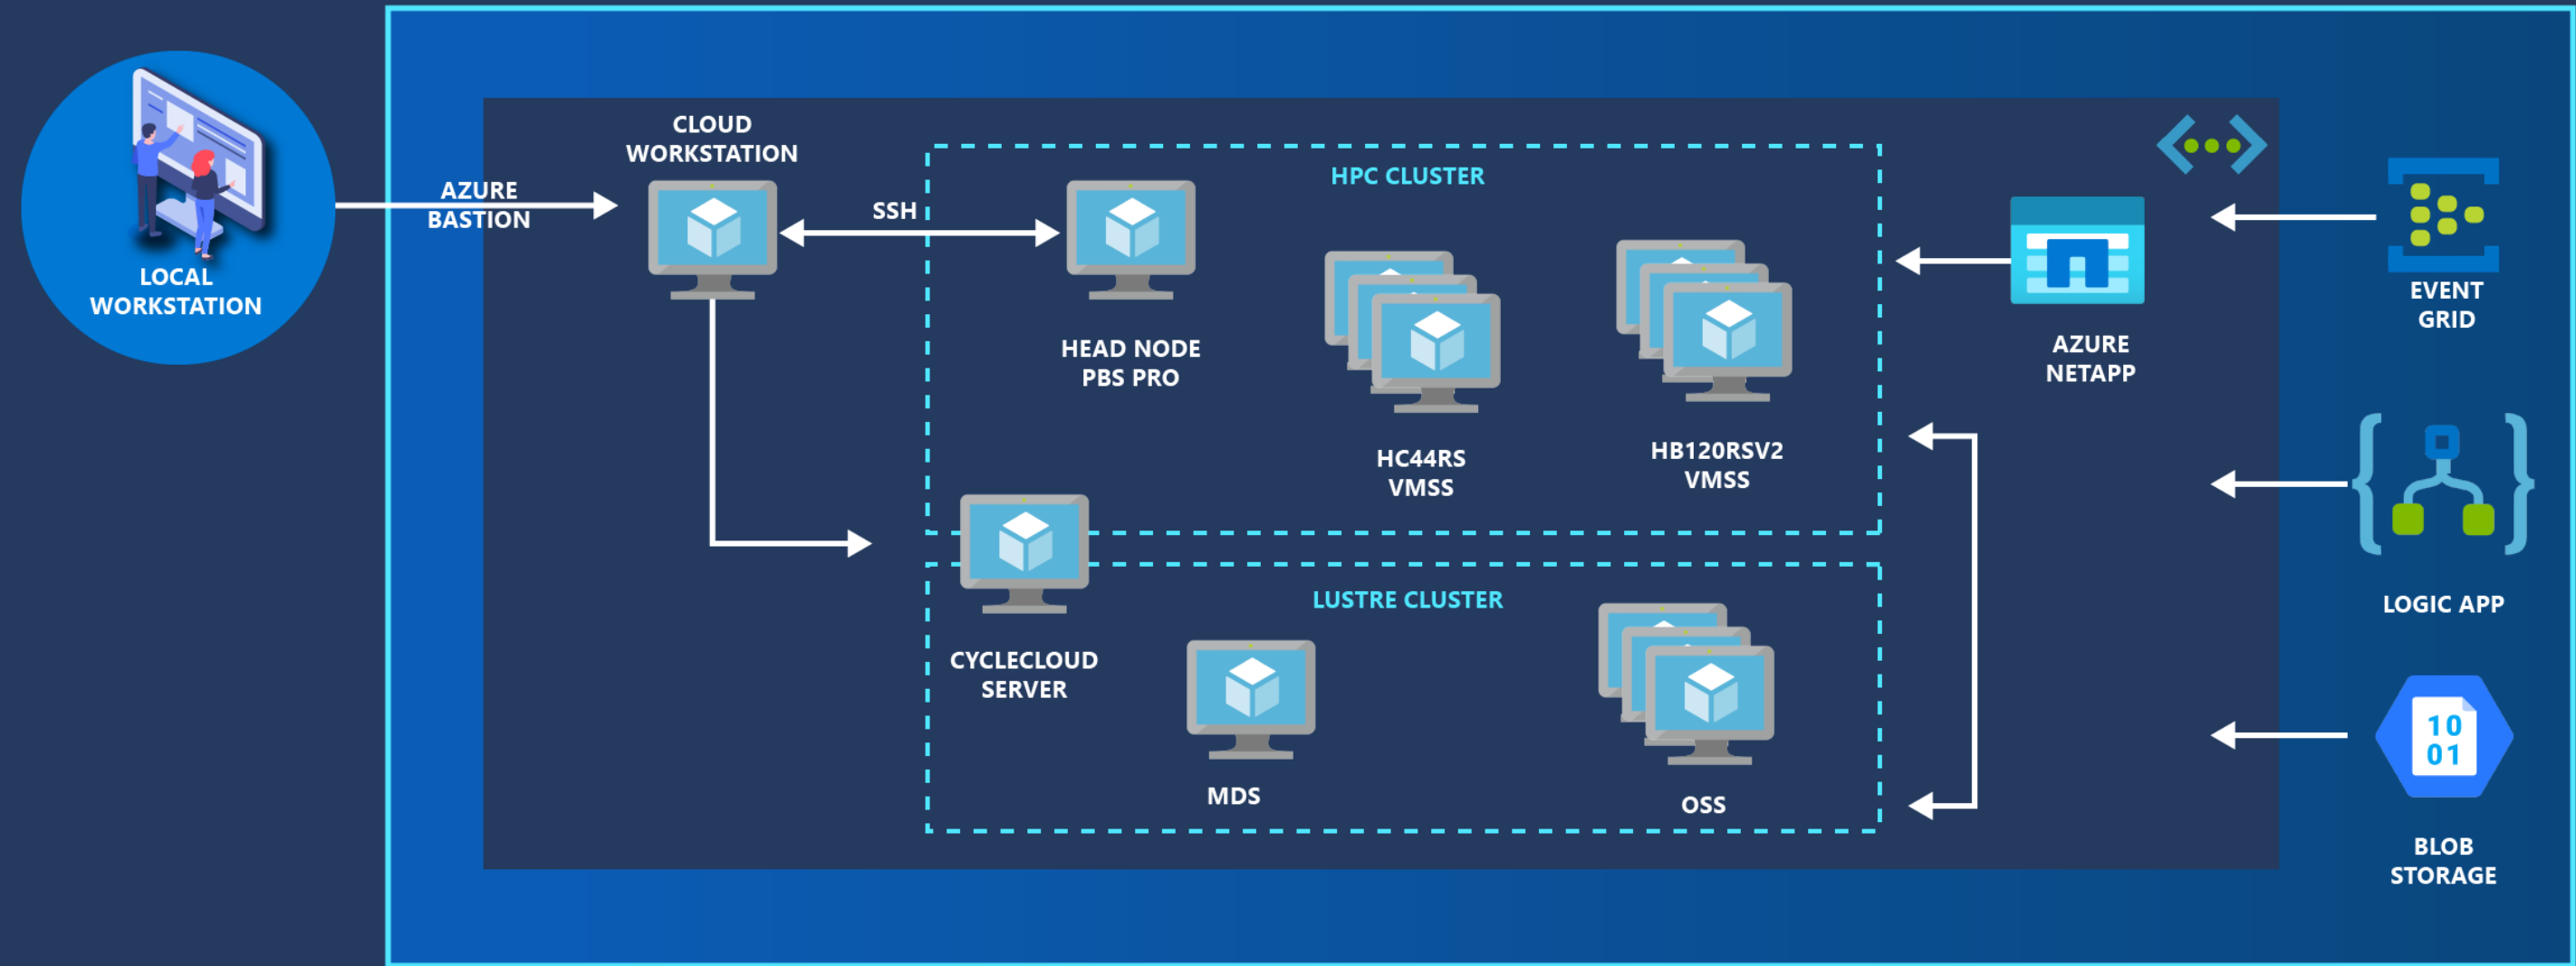


Microsoft
Azure



External Access Control (Azure Bastion)

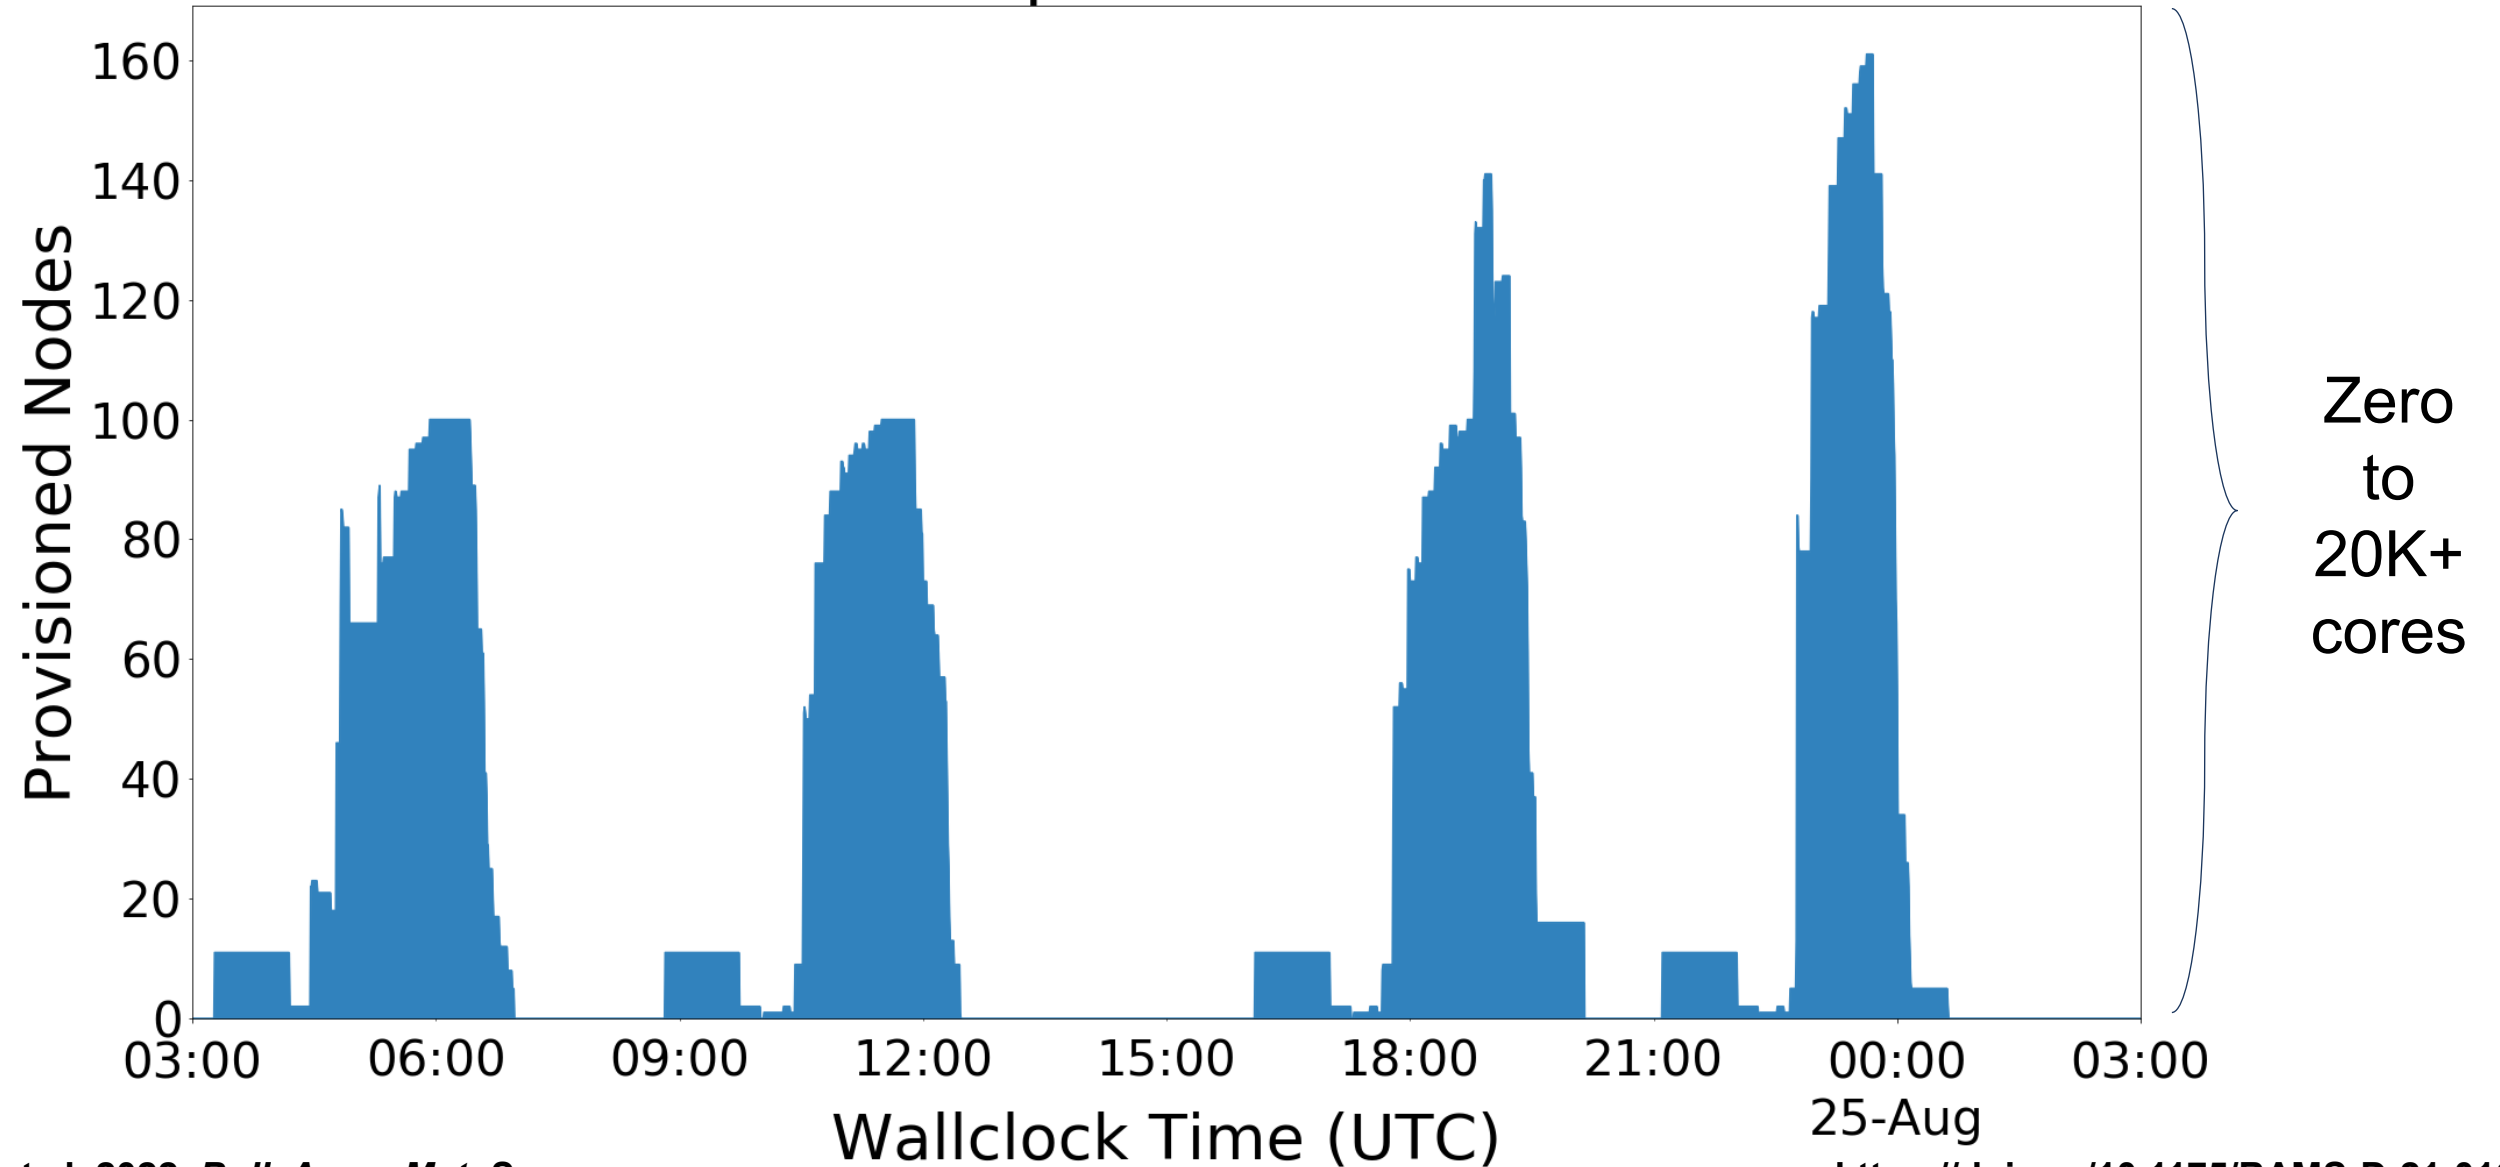
Traditional(ish) HPC Architecture




Leverage Managed Services

Dynamic Provisioning for Dynamic Load

24-Hour Snapshot of Cluster Size



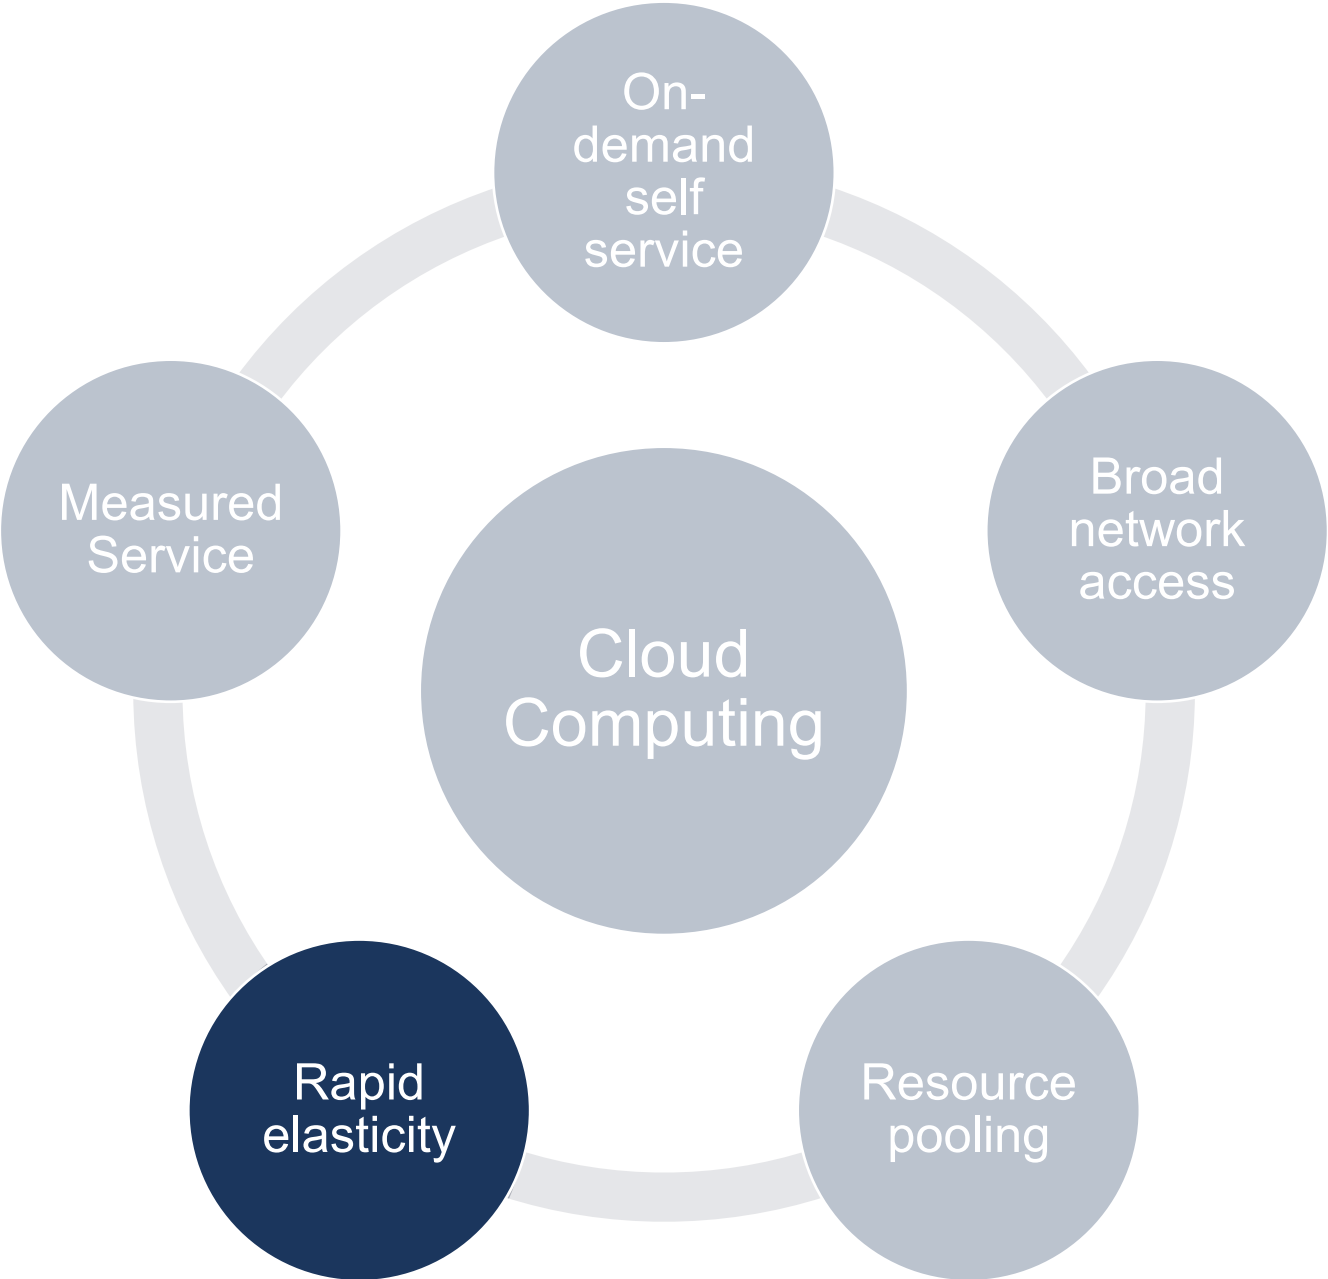
Cloud Computing



“Cloud computing is a model for enabling ubiquitous, convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction.”

- The NIST Definition of Cloud Computing (Mell & Grance, 2011)

Cloud Computing



Provision things when you need them and delete them when you don't.


Toward holistic NWP ecosystems

Focus on applications downstream of large models & data assimilation




Provision

- Cloud resources
- Core infrastructure



Deploy

- Applications
- Data connections




Scale

- Facilitate elasticity

Provision

Cloud resources and core infrastructure

Infrastructure as Code (IaC)



HashiCorp

Terraform

Describe infrastructure definition and configuration with text files that are managed like source code. Provide a cloud-agnostic method to deploy cloud resources.

Play 2: Adopt Infrastructure as Code

“[T]he goal is to automate all infrastructure provisioning and configuration in a repeatable, consistent way...”

Infrastructure as Code (IaC)



HashiCorp

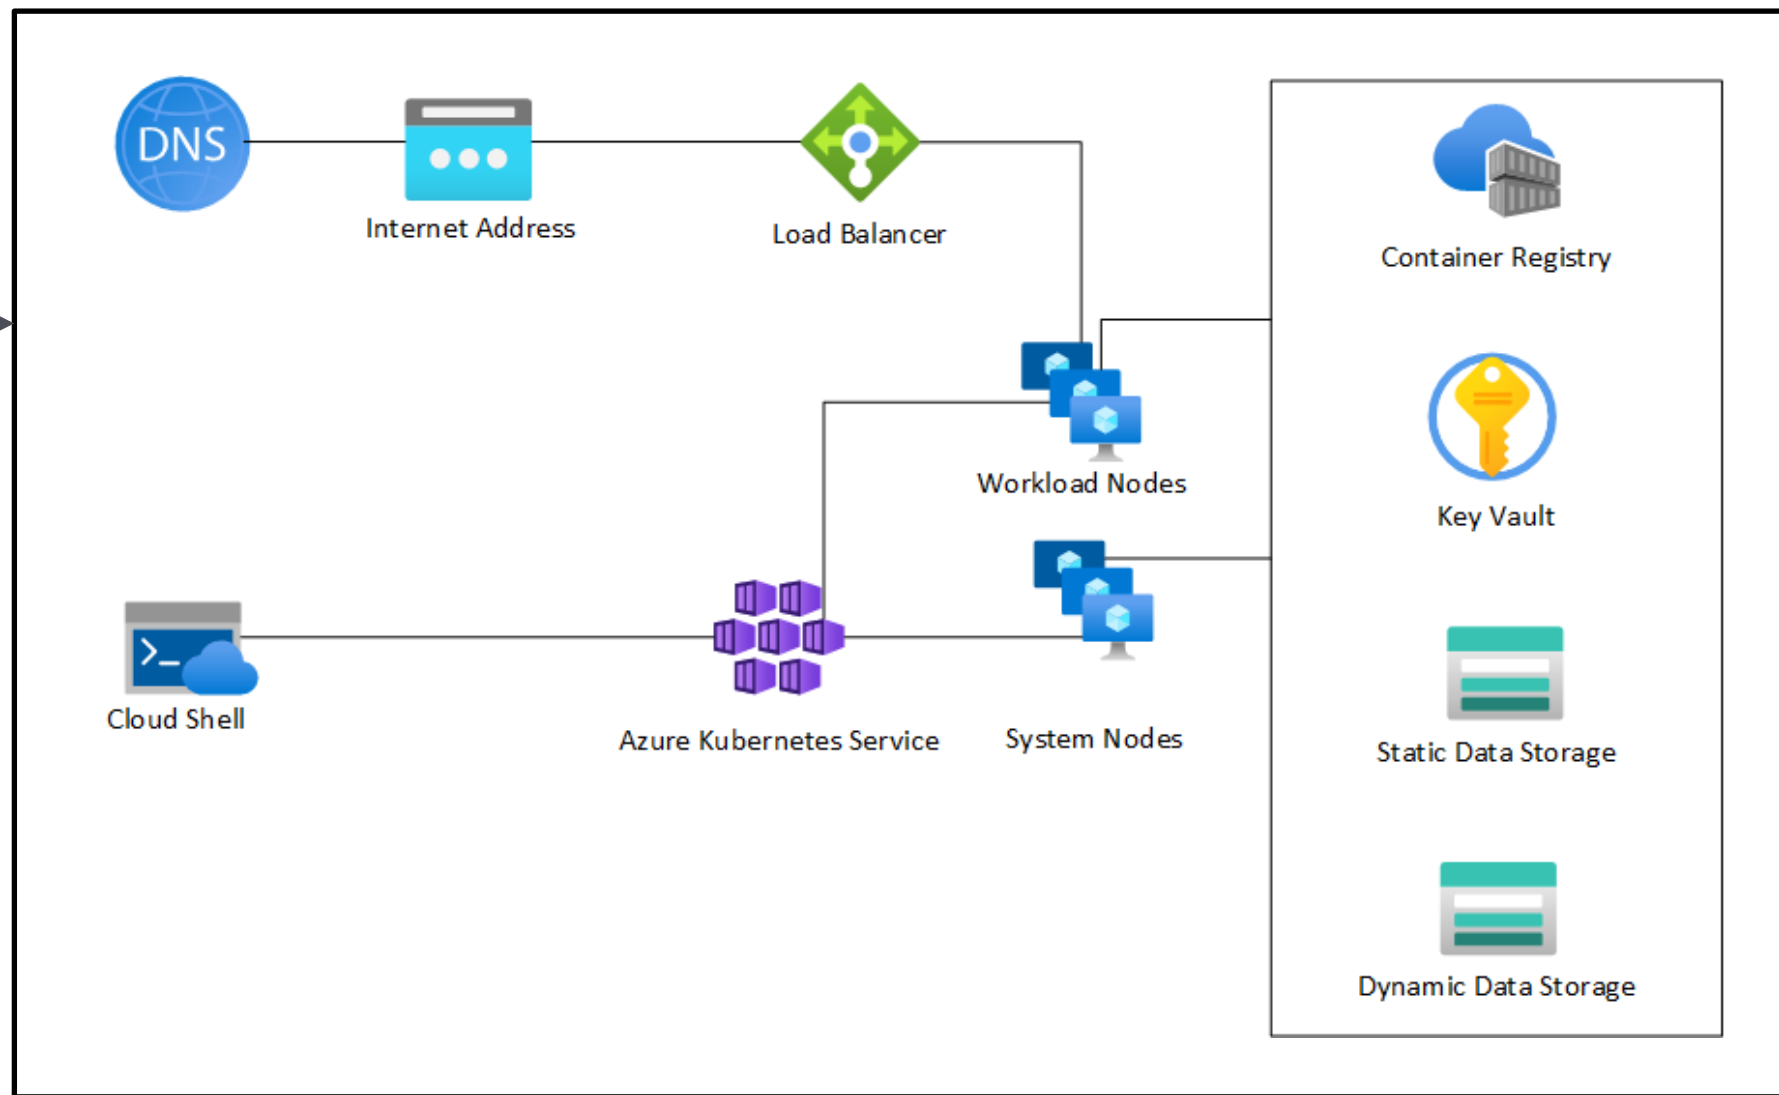
Terraform

Describe infrastructure definition and configuration with text files that are managed like source code. Provide a cloud-agnostic method to deploy cloud resources.

Using IaC might slow initial deployments – the *first* time.
Benefits far outweigh the additional cost up front.

“[T]he goal is to automate all infrastructure provisioning and configuration in a repeatable, consistent way...”

Infrastructure as Code (IaC)



Infrastructure as Code (IaC)

```
resource "azurerm_virtual_network" "vnet" {  
  name = "vnet-${var.project}-${var.environment}-${var.location}"  
  location = var.location  
  resource_group_name = azurerm_resource_group.resource_group.name  
  address_space = ["10.1.0.0/16"]  
}
```

```
resource "azurerm_subnet" "k8s_nodes_subnet" {  
  name = "snet-k8s-${var.project}-${var.environment}-${var.location}"  
  resource_group_name = azurerm_resource_group.resource_group.name  
  virtual_network_name = azurerm_virtual_network.vnet.name  
  address_prefixes = ["10.1.2.0/24"]  
}
```

IaC isn't just for storage or virtual machines – all the way down to the *network architecture* that is used.

Deploy

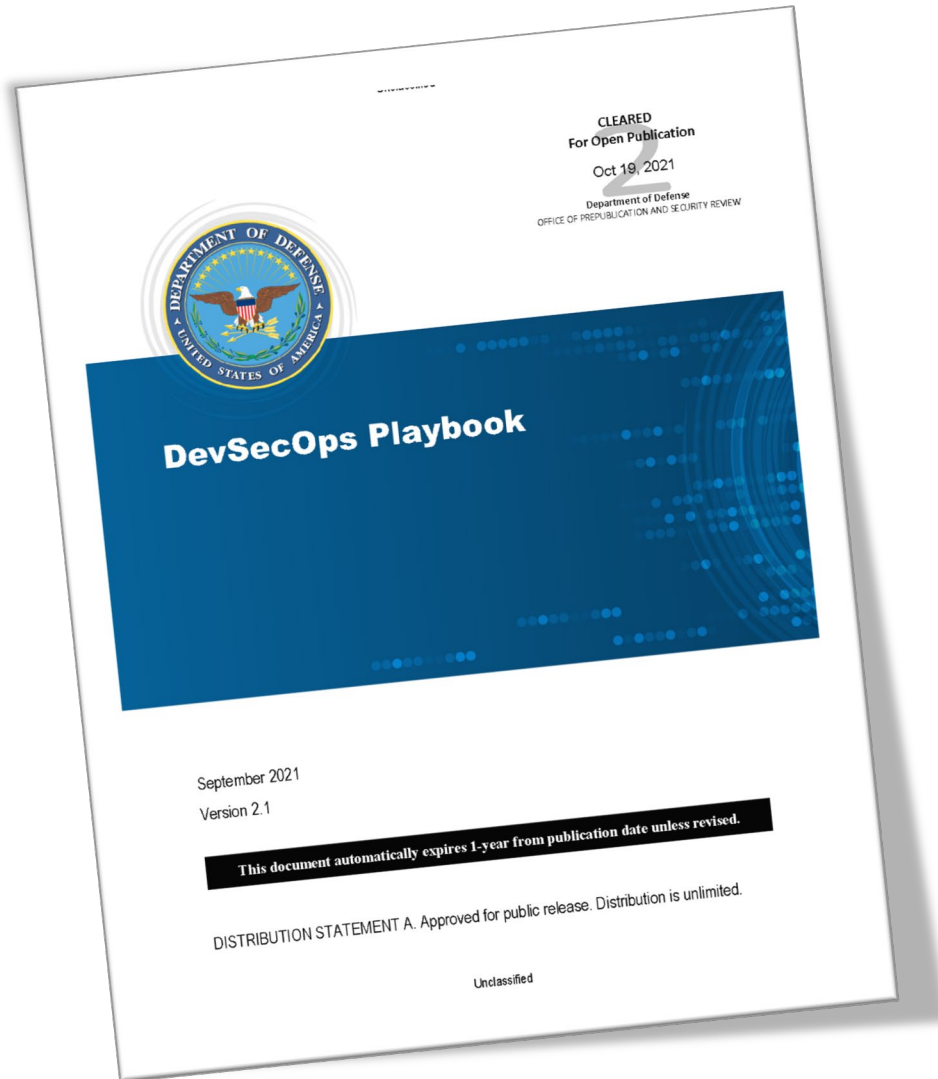
Applications & Data Connections



Kubernetes: Enabling technology for vendor independence



kubernetes



Play 3: Adopt Containerized Microservices

*“DoD DevSecOps Strategy requires ...
Kubernetes ... for container orchestration.”*

Kubernetes (k8s) provides a platform with consistent APIs to support deployments on-prem and across cloud vendors.

Kubernetes: Enabling technology for vendor independence




kubernetes



Play 3: Adopt Containerized Microservices

*“DoD DevSecOps Strategy requires ...
Kubernetes ... for container orchestration.”*

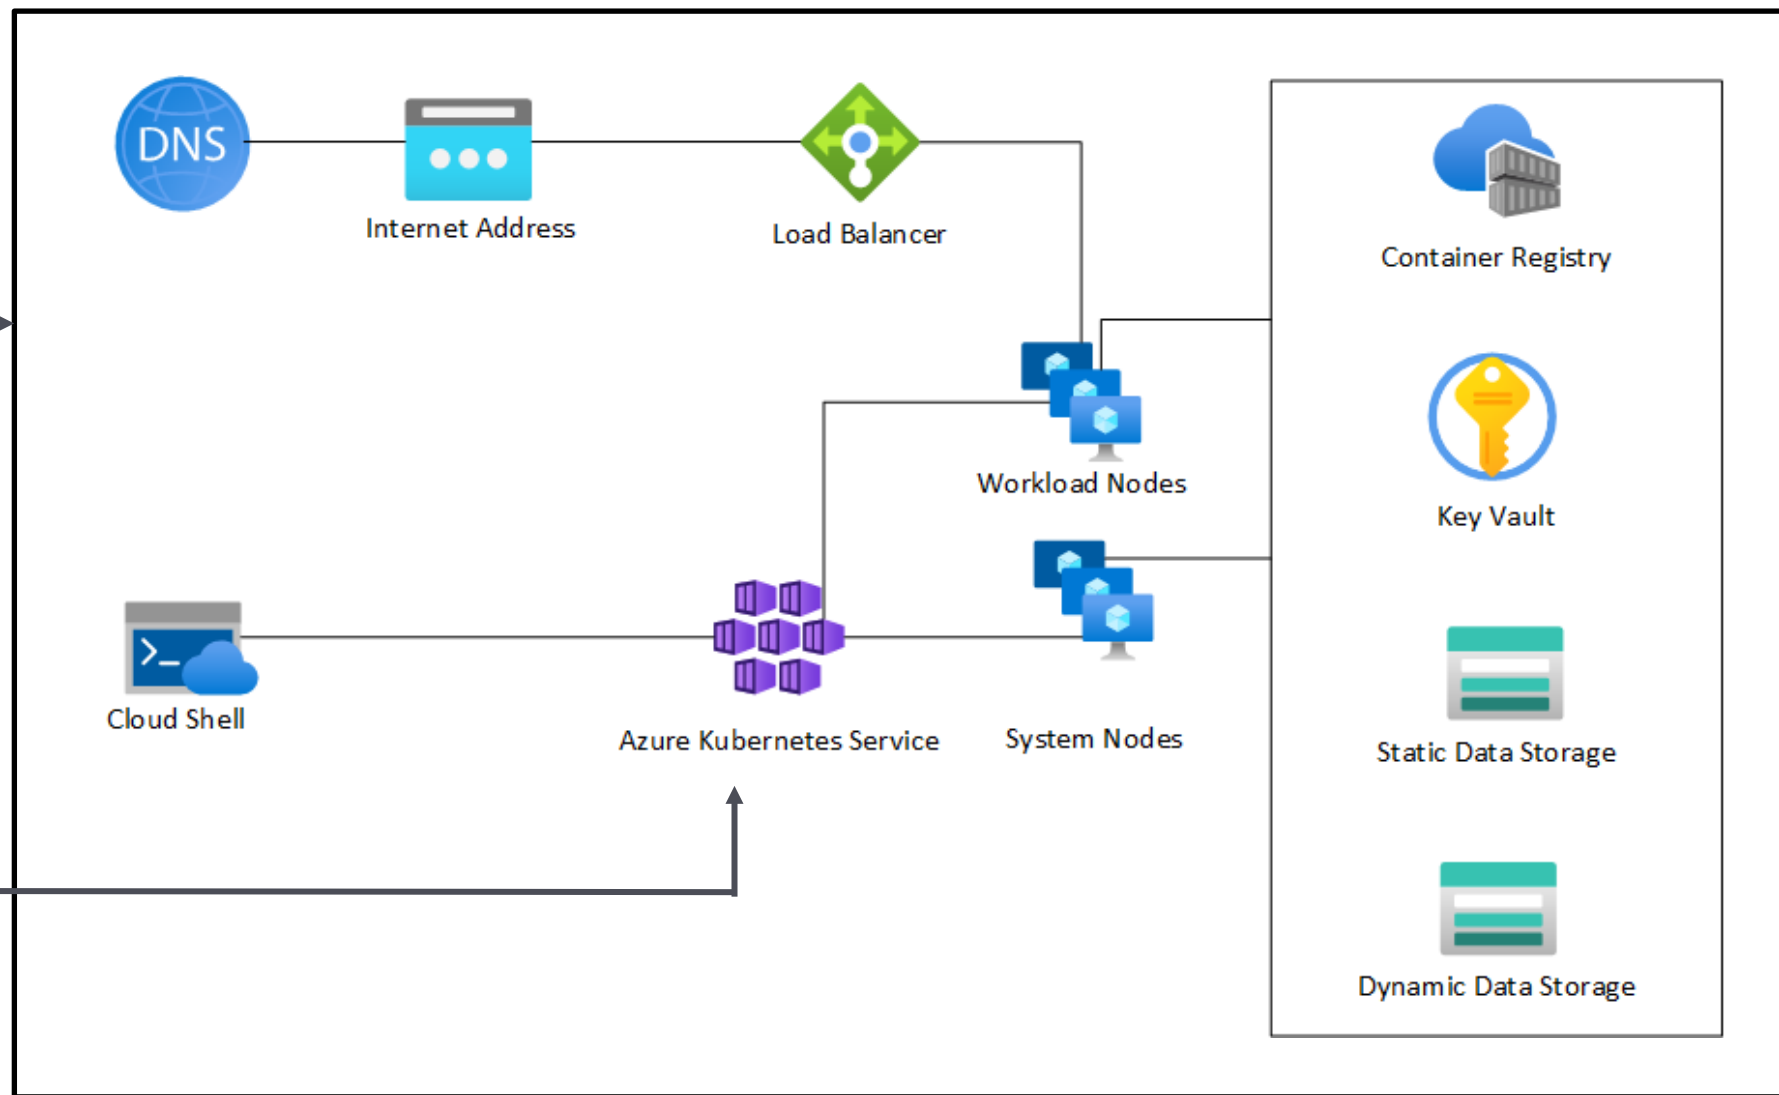
Kubernetes provides a platform for resilient orchestration of containerized workloads and provides abstractions for storage, networking, and other primitives with YAML configuration. Kubernetes enables *clustering* in a loosely-coupled, HPC-adjacent context.



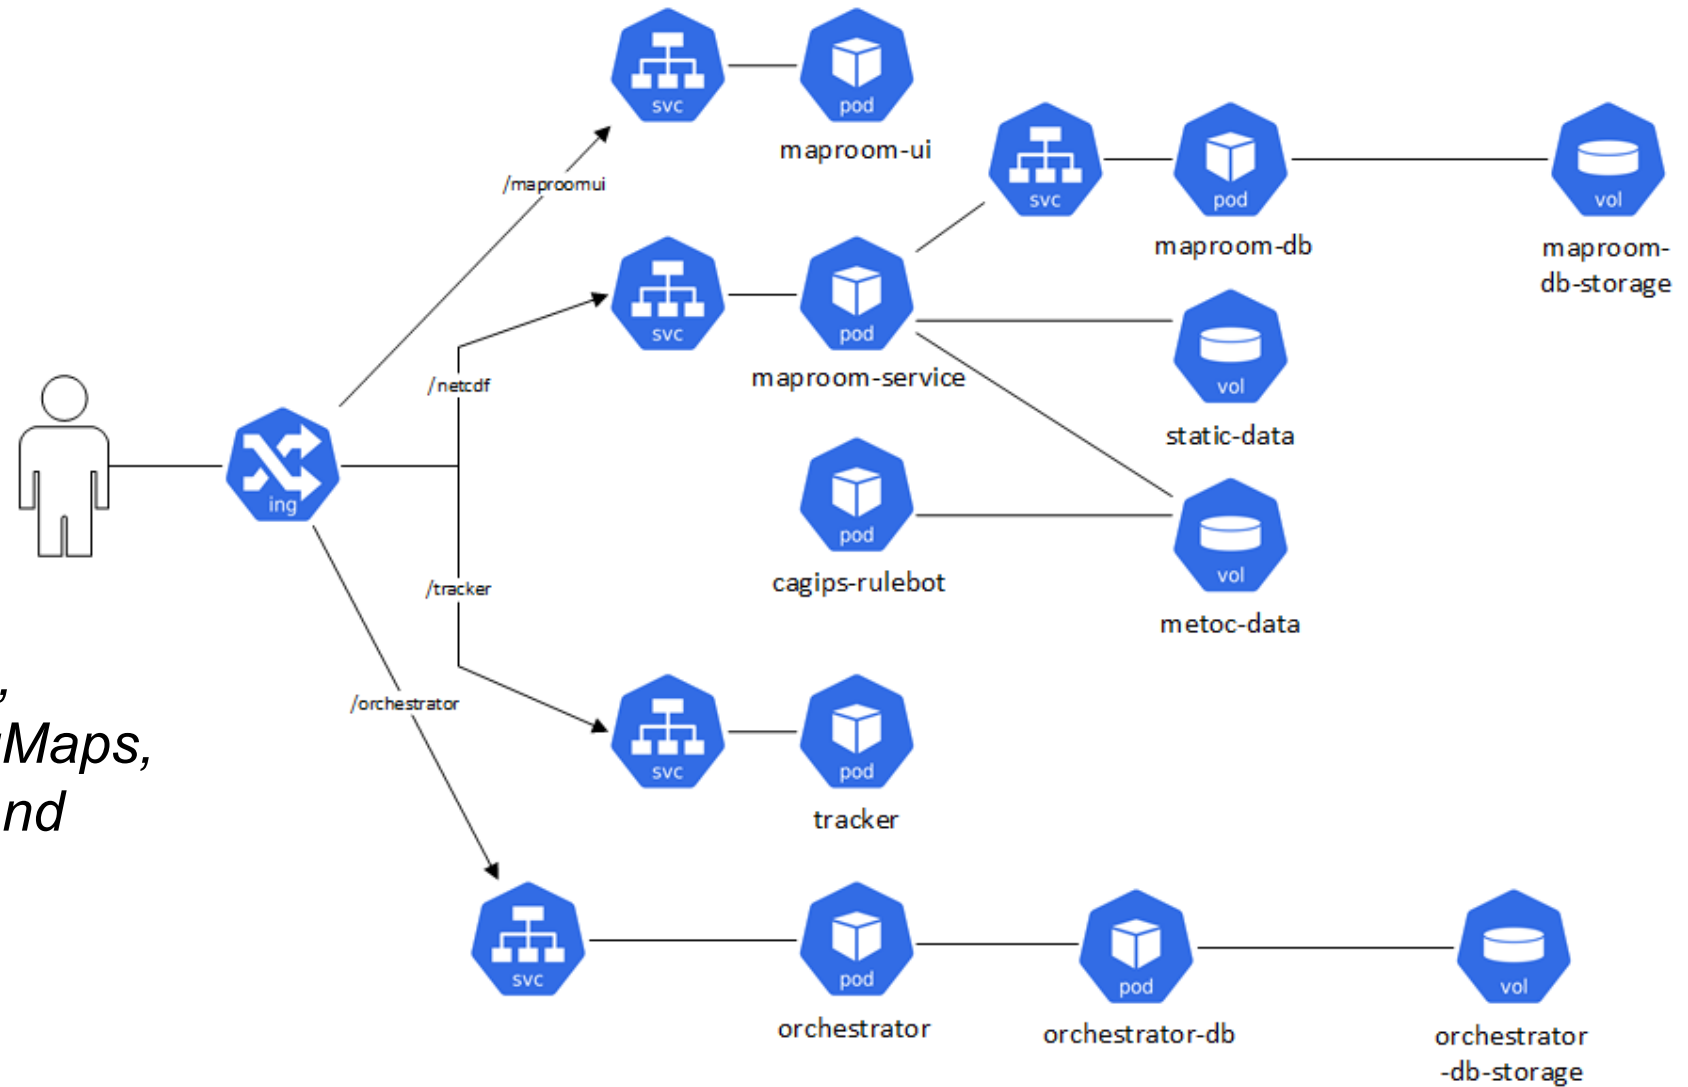
Store the description of applications running on the Kubernetes cluster in a Git repository - rely on GitOps tools like ArgoCD to keep the running applications synchronized with the authoritative repository.

Play 3: Adopt Infrastructure as Code

“GitOps is a paradigm where systems are described ... declaratively, using code to specify the desired state.”



GitOps: A sample set of applications



Resources like Deployments, ReplicaSets, Secrets, ConfigMaps, Endpoints, EndpointSlices, and CRDs omitted for clarity

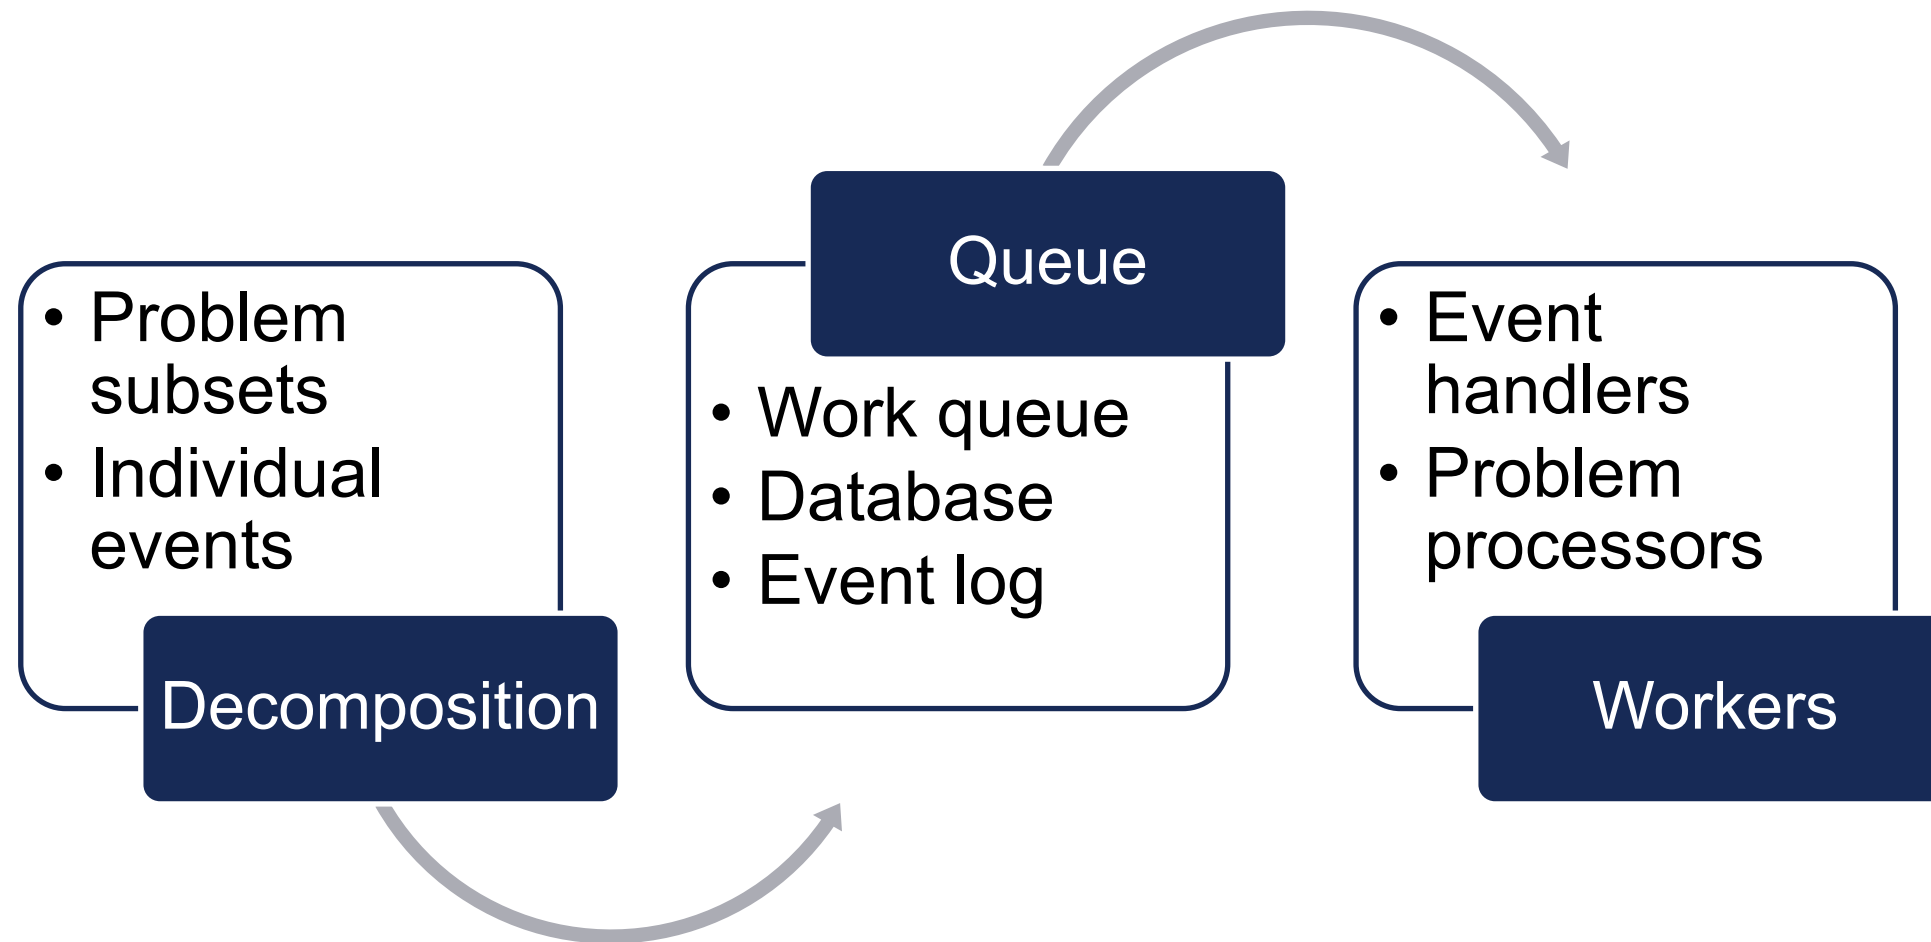
Scale

Architect for elasticity

U.S. NAVAL
RESEARCH
LABORATORY

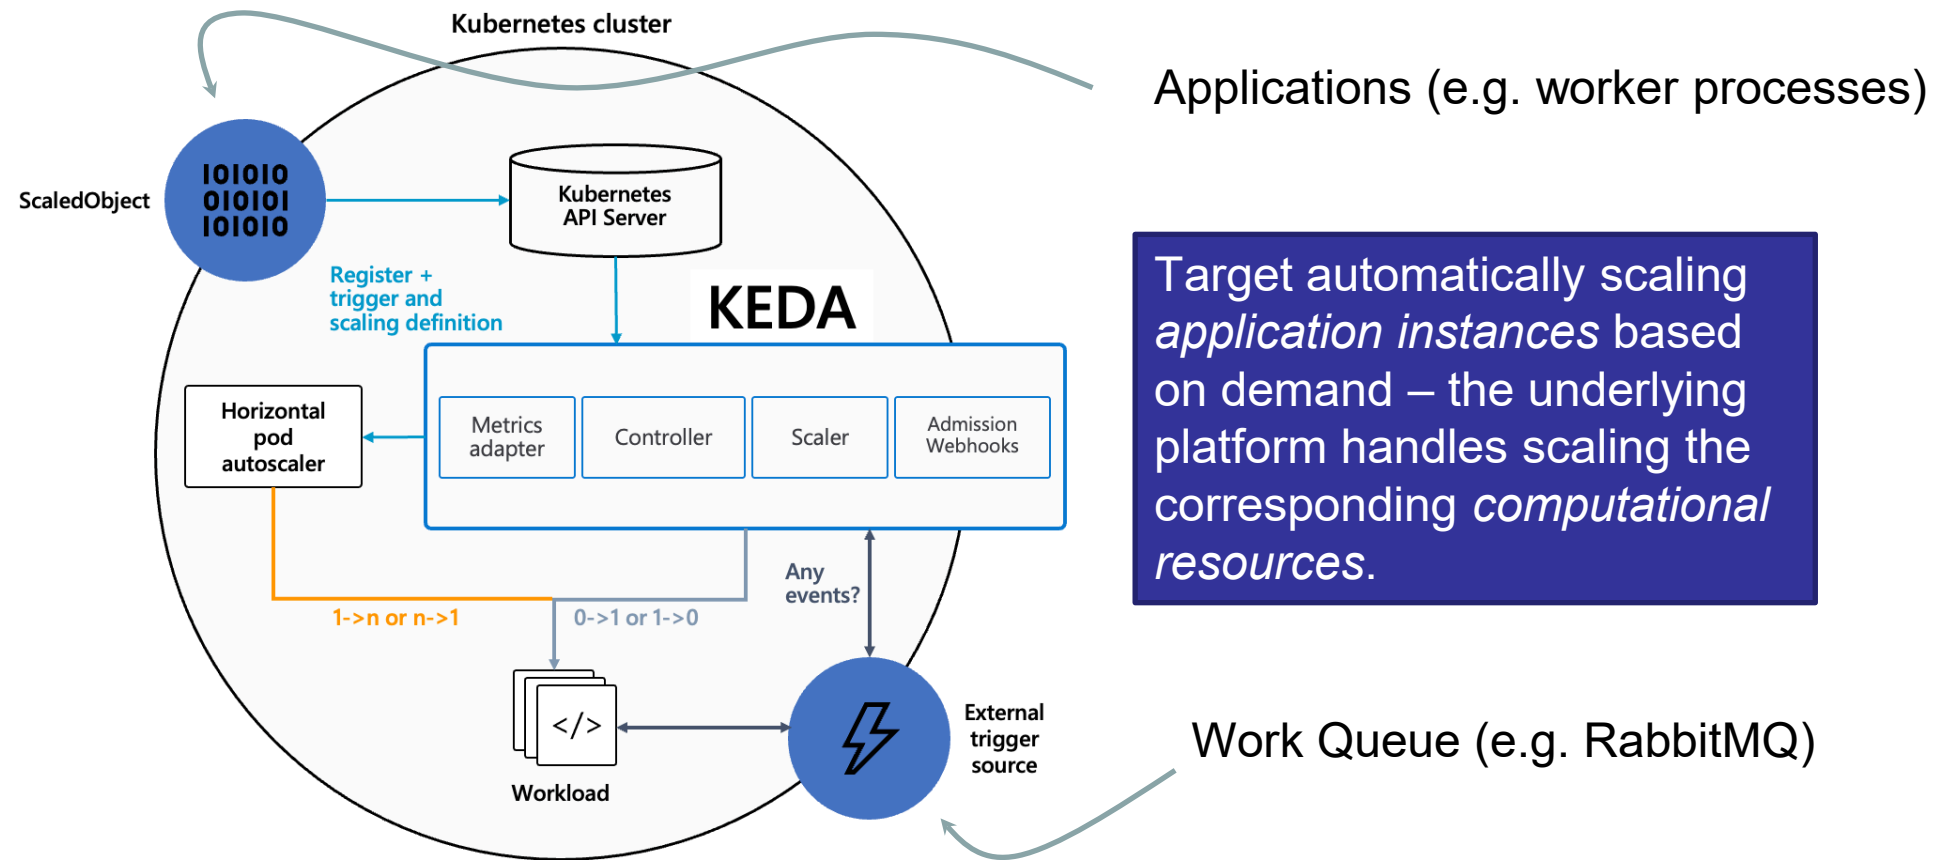


Recipe to drive elastic scaling



This is a *very familiar process* to those of us used to solving problems with HPC.

Scaling with KEDA*



*Kubernetes Event-Driven Autoscaler

Scaling with KEDA

Scaling Up



Scaling Down



Final Thoughts

Elastic Cloud Integration with Kubernetes

Identifying touch points facilitates vendor neutrality

Enables cooperative processing with traditional HPC systems

Widespread tools lead to broader apertures for hiring skilled engineers

Allows HPC to focus on HPC problems