




The Platform for **What's Next**

vmware[®]

by **Broadcom**

Kenneth Hjelmär

Lead Architect – Atea

Allan Kjær

Lead Architect - Atea





Private Cloud udkonkurrerer Public Cloud på 3 kritiske områder



Sikkerhed




Omkostninger




Kontrol

Hvad er VCF?

- VMware VCF – Jeres egen Private Cloud
- Administrere jeres miljø som hyperscalers
- Samlet overblik
- Nedbrydning siloerne i IT



VMware Cloud Foundation 9 – The Unified Platform




Samlet UI (Næsten)

Antallet af forskellige UI's bliver mindre.

Ting konsolideres ind i disse UIs:


- VCF Operations
- vCenter



Niveauer i management af VCF


- VCF Private Cloud
- VCF Fleet
 - Operations
 - Automation
- VCF Instance
- VCF Domains

VCF Private Cloud



Niveaueer i management af VCF

- VCF Private Cloud
- VCF Fleet
 - Oportations
 - Automation
- VCF Instance
- VCF Domains



Nye nøgle features fra vSphere til VCF

- Bedre Lifecycle management
- Authentication på tværs af produkter
- Certificate management
- Password management
- License management

Multi Tenancy

- VCF Automation har en ny Multi Tenancy
- vCenteret har VPC
 - VPC som vi kender det fra Public Cloud
 - Hver "Tenant" kan nemt oprettet ny netværks

Hvordan kommer man fra vSphere til VCF 9

Afhænger af konfigurationen

- Kun vSphere
- vSphere med Aria Operation
- vSphere med Aria Suite Lifecycle og Aria Operation


Udvikling af næste generation VCF



Infrastruktur leveret med
udviklerhastighed



Private AI
as a Service



Cyber
Resilient Data




vSphere Kubernetes Service (VKS)

(The artist formerly known as Tanzu Kubernetes Grid)



Udvikler Selvstændighed




Accelerated Path
to Production



IT Kontrol



vSphere Supervisor



- ESX cluster = **Supervisor**
- ESX hosts = **Worker nodes** afvikler vores workloads
- **Spherelet** er en specialudviklet version af Kubelet der integrerer med ESX
- Supervisor består af et 3 **Control Plane VMs** der udstiller API adgang
- **API adgang** tillader brugerne at manage og deploy ressourcer via Kubernetes CLI (kubectl)

Indbygget Kubernetes Runtime og Multi-Cluster Management



DEVELOPER/
CONSUMER



PLATFORM
ENGINEER /
CONSUMER



CLOUD ADMIN

VMware Cloud Foundation

Modern
Applications



AI/ML
Workloads



Distributed Database



Data
Streaming



Java EE, EAP,
Middleware



Collaboration
Platforms



Multi-cluster Management | Lifecycle | Policies



vSphere Kubernetes Service (VKS)



Pinniped



vSphere CSI Driver



Antrea



CertManager



Prometheus



Harbor



Carvel



Guest Cluster
Auth Service



Calico



Velero



Istio Service Mesh



FluentBit

VKS Core Packages

VKS Standard Packages

VM Service



Network Services



Backup Service



Storage and Data
Services



GitOps
Integration



Identity and
Access Control



Open Extensible for
CNCF Compliant
3rd Party Services

vSphere Supervisor



COMPUTE




STORAGE



NETWORKING

vSphere Superv



- Upstream Conformant
CNCF Certificeret
Kubernetes runtime
- Opgraderbar udenfor
vSphere/vCenter
releases
- N-2 Kubernetes
versioner supporteret

VKS Services

Core and Standard Packages in VKS

VKS Core Packages

Installed by default, and lifecycle-managed within VKr

Authentication & Access



Pinniped



Guest-cluster-auth-service

DevOps Tooling, Secret & Package Management



Kapp-controller



Secret-gen controller



Storage



vSphere CSI Driver

Networking (CNI)



Antrea



Calico

Cloud Provider Integration



vSphere-cpi

VKS Standard Packages

Included for customer to opt-in, and lifecycle-managed by VCF

Networking



Contour



External DNS



Istio Service Mesh



CertManager



Windows gMSA

Observability



FluentBit



Telegraf



Prometheus

DevOps Tooling



Harbor

Scaling & Resilience



Velero



Cluster Autoscaler

YAML Manifest Definerer Kubernetes Control Plane / Worker Nodes

```
apiVersion: cluster.x-k8s.io/v1beta1
kind: Cluster
metadata:
```

```
  name: demo-VKS
  namespace: demo-workload
```

```
% kubectl config use-context demo-workload
Switched to context "demo-workload".
```

```
spec:
  clusterNetwork:
    services:
      cidrBlocks: ["198.51.100.0/12"]
    pods:
      cidrBlocks: ["192.0.2.0/16"]
  topology:
    class: tanzukubernetescluster
    version: v1.30.1+vmware.1-fips-tkg.5
```

```
% kubectl get tkr
v1.30.1+vmware.1-fips-tkg.5
```

```
    metadata:
      replicas: 3
  workers:
    replicas: 2
```

```
  variables:
    - name: vmClass
      value: best-effort-2xlarge
    - name: storageClass
      value: vsan-default-storage-policy
```

```
% kubectl get vmclass
best-effort-2xlarge  8      64Gi      22d
```

```
% kubectl get storageclass
vsan-default-storage-policy
```

YAML Manifest Definerer nu også virtuelle maskiner

```
apiVersion: vmoperator.vmware.com/v1alpha1
kind: VirtualMachine
metadata:
  name: ubuntu-vm
  namespace: demo-workload
spec:
  networkInterfaces:
    - networkName: "vm-62-dpg"
      networkType: vsphere-distributed
  className: best-effort-small
  imageName: vmi-b1167122bcb69408d
  powerState: poweredOn
  storageClass: raid-1
  vmMetadata:
    configMapName: ubuntu-vm-cfm
    transport: CloudInit
```

VM Operator now supports Secrets
as an alternative to ConfigMap

```
% kubectl config use-context demo-workload
Switched to context "demo-workload".
```

```
% kubectl get vmclass
```

NAME	CPU	MEMORY	AGE
best-effort-large	4	16Gi	22d
best-effort-small	2	4Gi	22d


```
% kubectl get vmi
```

NAME	PROVIDER-NAME	CONTENT-LIBRARY-NAME
IMAGE-NAME	OS-TYPE	FORMAT
vmi-b1167122bcb69408d	clitem-b1167122bcb69408d	cl-eb2a2047486a2caa8
ubuntu-22.04	ubuntu64Guest	OVF
		26h


```
apiVersion: v1
kind: ConfigMap
metadata:
  name: ubuntu-vm-cfm
  namespace: demo-workload
data:
  user-data: |
    base64 encoded config
  hostname: ubuntu-vm
```

```
#cloud-config
ssh_pwauth: yes
users:
  - default
  - name: ubuntu
    ssh_authorized_keys:
      - ssh-rsa xxx
    sudo: ALL=(ALL) NOPASSWD:ALL
  groups: sudo
```


VMware vSphere Kubernetes Services



Indbygget GitOps i VCF



Indbygget monitorering i VCF




- VCF Operations integration
- Overvåger Supervisor/VKS
- Drill down i Kubernetes objekter
- Detaljerede metrics til troubleshooting
- OOTB dashboards/Metrics/Alerts




Udvilker selvstændighed

- ✓ Definer app requirements i kode
- ✓ Push til Git, trigger deployment
- ✓ Brug eksisterende værktøjer (VS Code, GitHub)
- ✓ Benyt kun godkendte services fra katalog
- ✓ Samlet overblik over app



Accelerated Path to Production



IT Kontrol

- ✓ Kontrol over hvilke services er tilgængelig for udvikleren
- ✓ Oprethold org-level isolation og politikker
- ✓ Overvåg cross-org brug og aktivitet
- ✓ Pre-definer infrastruktur og sikkerheds politikker
- ✓ Udstil kun godkendte services i et katalog

Udvikler autonomi og kontrolleret adgang til services

NEW Developer-Ready Services



Built-In Data Services

vSAN native S3
object storage

Secure,
enterprise-grade
Postgres, MySQL



Integrated App Delivery

GitOps with
ArgoCD in VCF(Available)

Istio service
mesh



Guardrails by Default

Policy-as-code
for VPC security

Hardened
containers


Udvikling af næste generation VCF



Infrastructure at the
Speed of the Developer



Private AI
as a Service



Cyber
Resilient Data

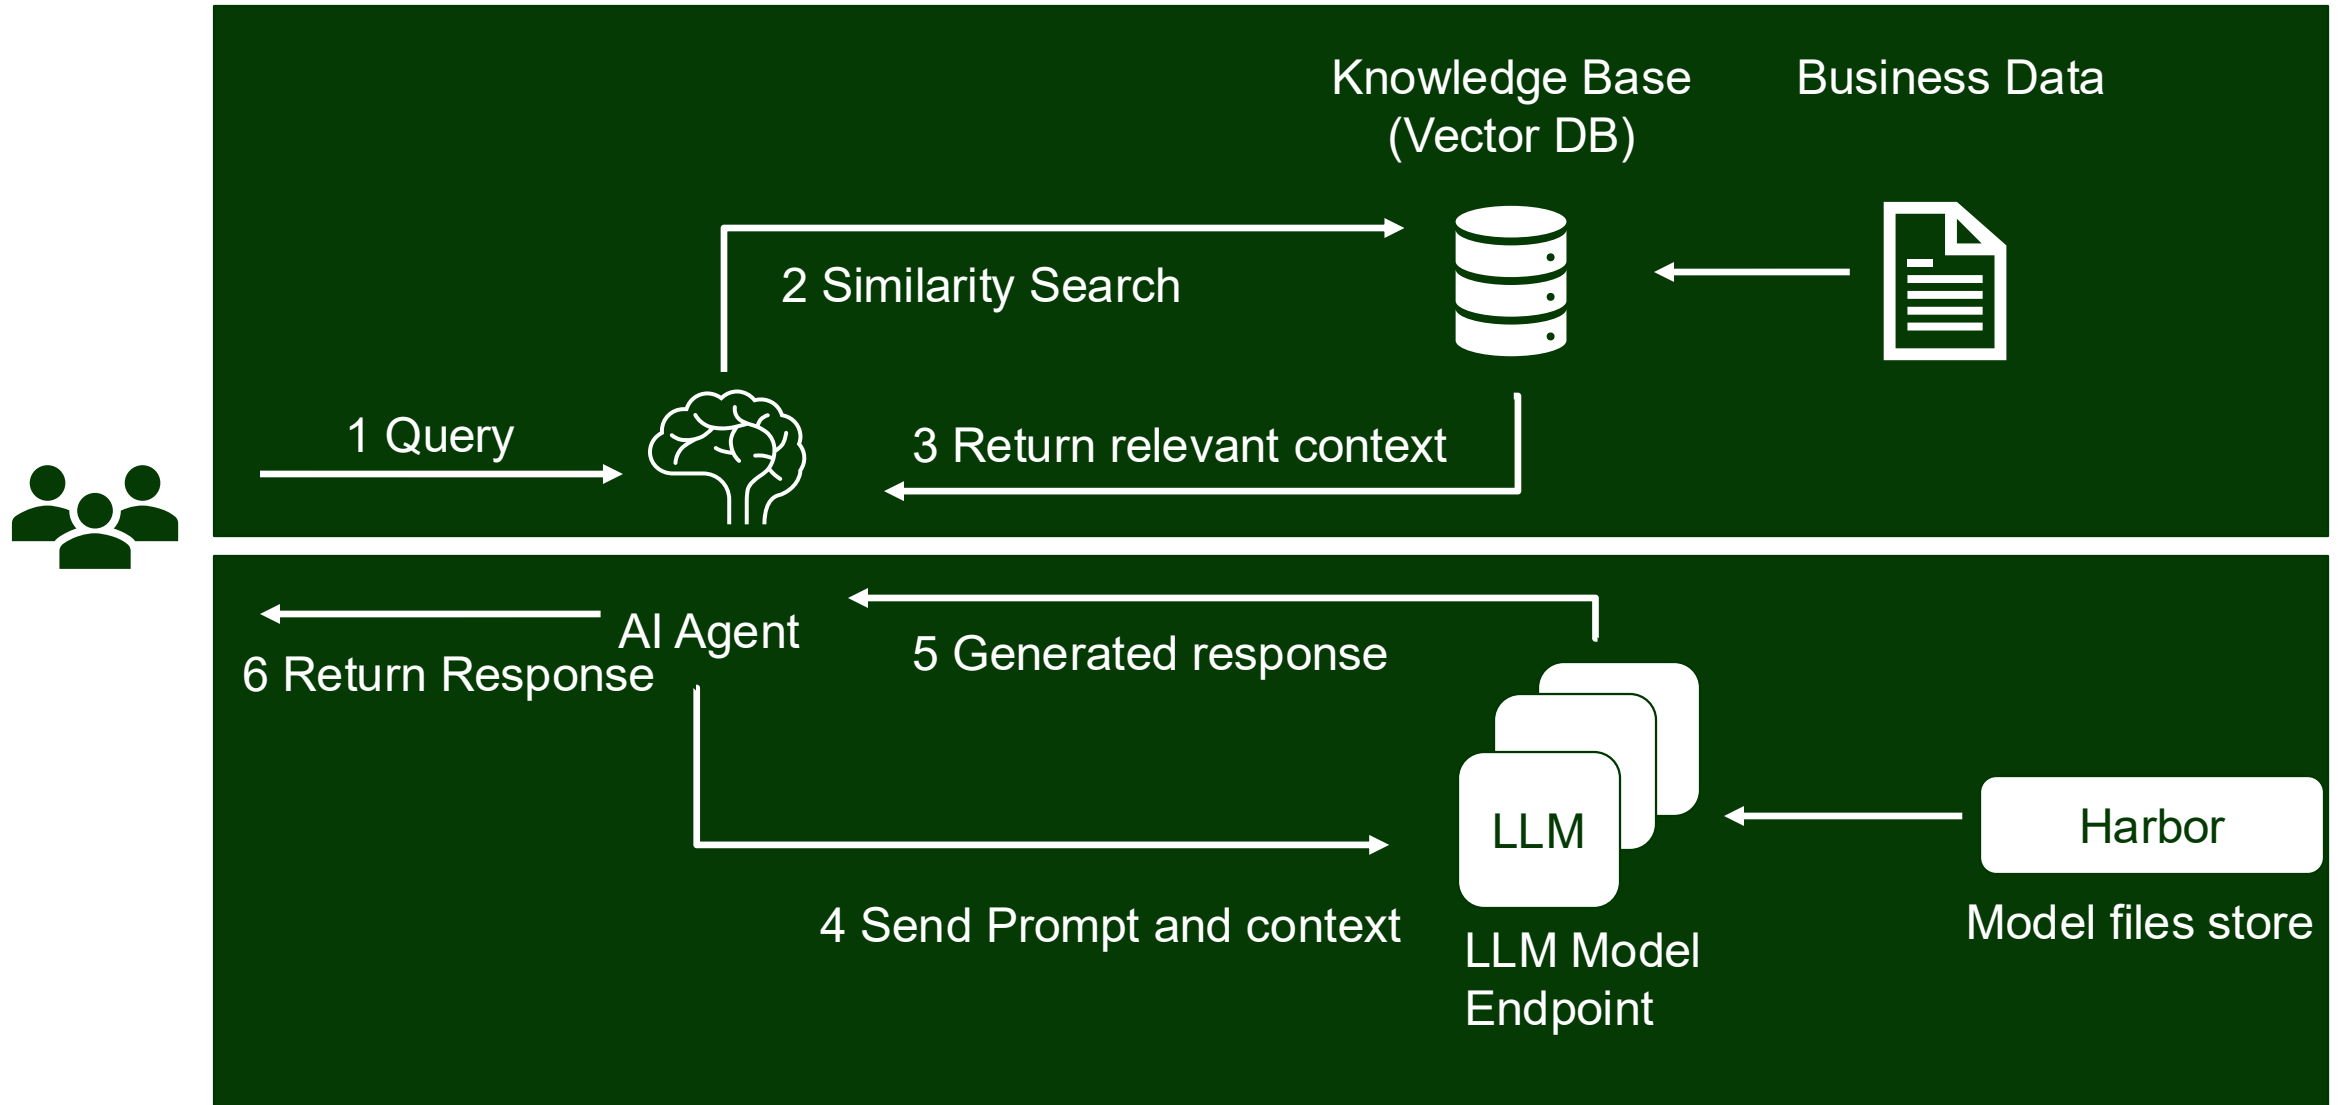
”Add-on”: Private AI

Kører jeres egne AI modeller i den Private Cloud

Den gode nyhed, nu inkl. i VCF9 licensen

Combine your internal, private data with a Large Language Model so users can interact with the data through natural language, securely and at scale

A RAG Agent



Data Sources

- Sharepoint
- Google Drive
- Confluence
- S3-Compatible APIs

The screenshot displays the ATEA Tech Tour interface, specifically the 'Data Sources' configuration page for a 'Customer Support Data' knowledge base. The interface is divided into several sections:

- Left Sidebar:** Contains navigation links for 'Data Indexing and Retrieval', 'Knowledge bases', 'Data sources', 'Agent Builder', and 'Agents'.
- Top Header:** Shows the breadcrumb 'Data Indexing and Retrieval > Knowledge Bases > Customer Support Data' and a 'DELETE KNOWLEDGE BASE' button.
- Knowledge Base Details:** A section containing information about the 'Customer Support Data' knowledge base, including its type, description, index refresh policy, last index, next index, and a list of agents (Customer Support Agent 1 and Customer Support Agent 2).
- Data Sources:** A section listing the data sources used for indexing, including 'Customer Support' and 'Accounting Sheet', both marked as 'Indexed'.
- Index Settings:** A section showing the index status (Completed), embedding model (nvidia/veva-22), chunk size (xxx), token overlap (0), and text splitter (Sentence).
- Document List:** A table listing documents (doc 1 to doc 7) and their status (Indexed).
- Add Data Source:** A section for adding new data sources, showing a progress bar with steps: 1. General Information, 2. Data Source Details, and 3. Review and Complete. The 'Data Source Type' dropdown is open, showing options like 'Google Drive Folder', 'Confluence Space', 'Amazon S3 Compatible Store', and 'Sharepoint Enterprise Site'.

Index Knowledge Base



Select an embedding model, configure index settings, and index the knowledge base.

Embedding model endpoint * ⓘ	Select ▼
Text splitter * ⓘ	Sentence
Chunk size * ⓘ	Enter maximum number of tokens
Token overlap * ⓘ	0

Knowledge Bases


Contains metadata, text chunks and vector embeddings create for each processed document accessible from data source

**** Multiple Data Sources can concurrently contribute to a single Knowlagde base.**

Model as a Service

Flere forskellige RAG Agents kan benytte same LLM

- Sikkert
- Bedre udnyttelse af ressourcer
- Mindre vedligehold



Nyt fra VMware Explore

Private cloud. Private AI. Public success.

Uncover new IT strategies for better results.





Kontakt:

Kenneth Hjelmer

Kenneth.hjelmar@Atea.dk

Kontakt:

Allan Kjær

allan.kjar@rAtea.dk

Kontakt:

Flemming Westervang

Flemming.westervang@Atea.dk

Næste indlæg:
Hybrid by Design:
En strategisk vej til AI og agil drift

