25TH EDITION
MPLS SD&AI NET WORLD
9/11 APRIL 24
palais des congrès de paris

Bell Canada
Enabling the Telco Edge

Daniel Bernier
Technical Director

Kashif Islam
Principal Architect

Journey towards service programming with SRv6 at Bell

Ongoing journey at Bell towards an **application oriented** network

Prologue - SR-MPLS introduction

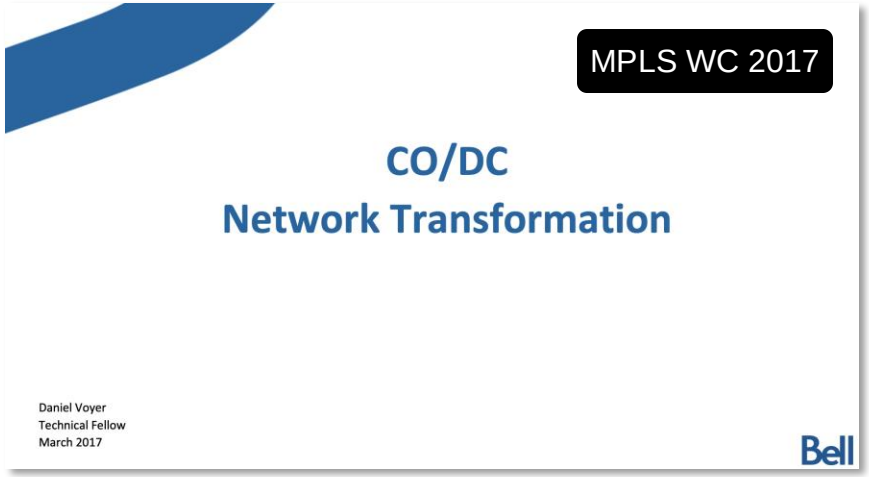

Chapter 1 : Enabling in-network SFC with SRv6 and P4

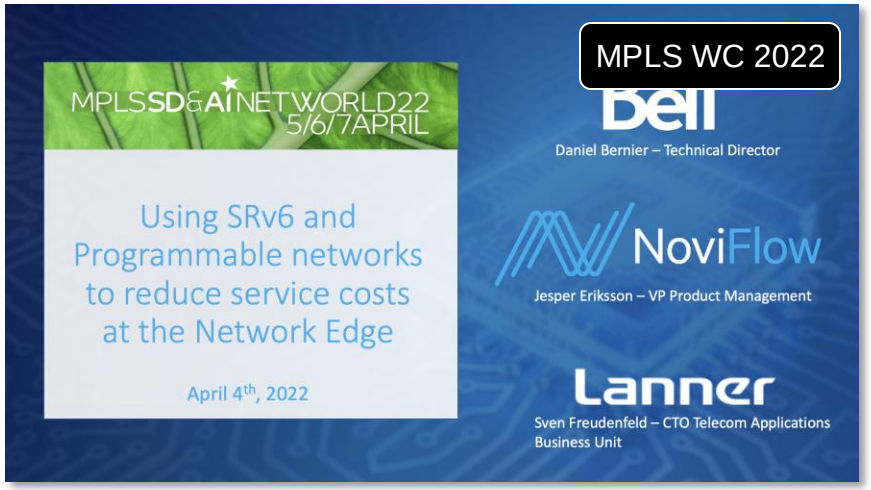

Chapter 2 - Leveraging SRv6 for simplified multi-cloud

Bell Red Hat

Multi-Cloud SRv6 Deployment Update

Some new development since last year helping us moving the needle further

- Most issues regarding IPv6 addressing limitations have been fixed or not required anymore
- Successful interop between Cilium Enterprise with "locator auto-tuning" and IOS-XR using Wide Lib structure.

Upcoming support for Cilium inside ROSA on OCP.

Google Cloud

- ~~Support for IPv4 protocol in next-header after IPv6 encap~~
- Ability to filter IPv4 prefixes out of VPC peering policies.
- ~~Support for customer-defined ULA addressing space (ie not forced fd20::/20)~~
- ~~Support for BYOIPv6 address space including ULA~~

aws

- ~~Support for ULA addressing space (roadmap)~~
- ~~Support for BYOIPv6 ULA address spaces~~
- Support for flexible CIDR ranges for IPv6
- Support of hashing based on IPv6 flow labels.

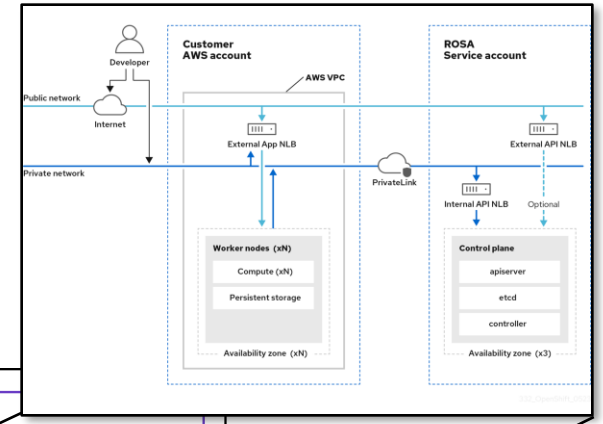

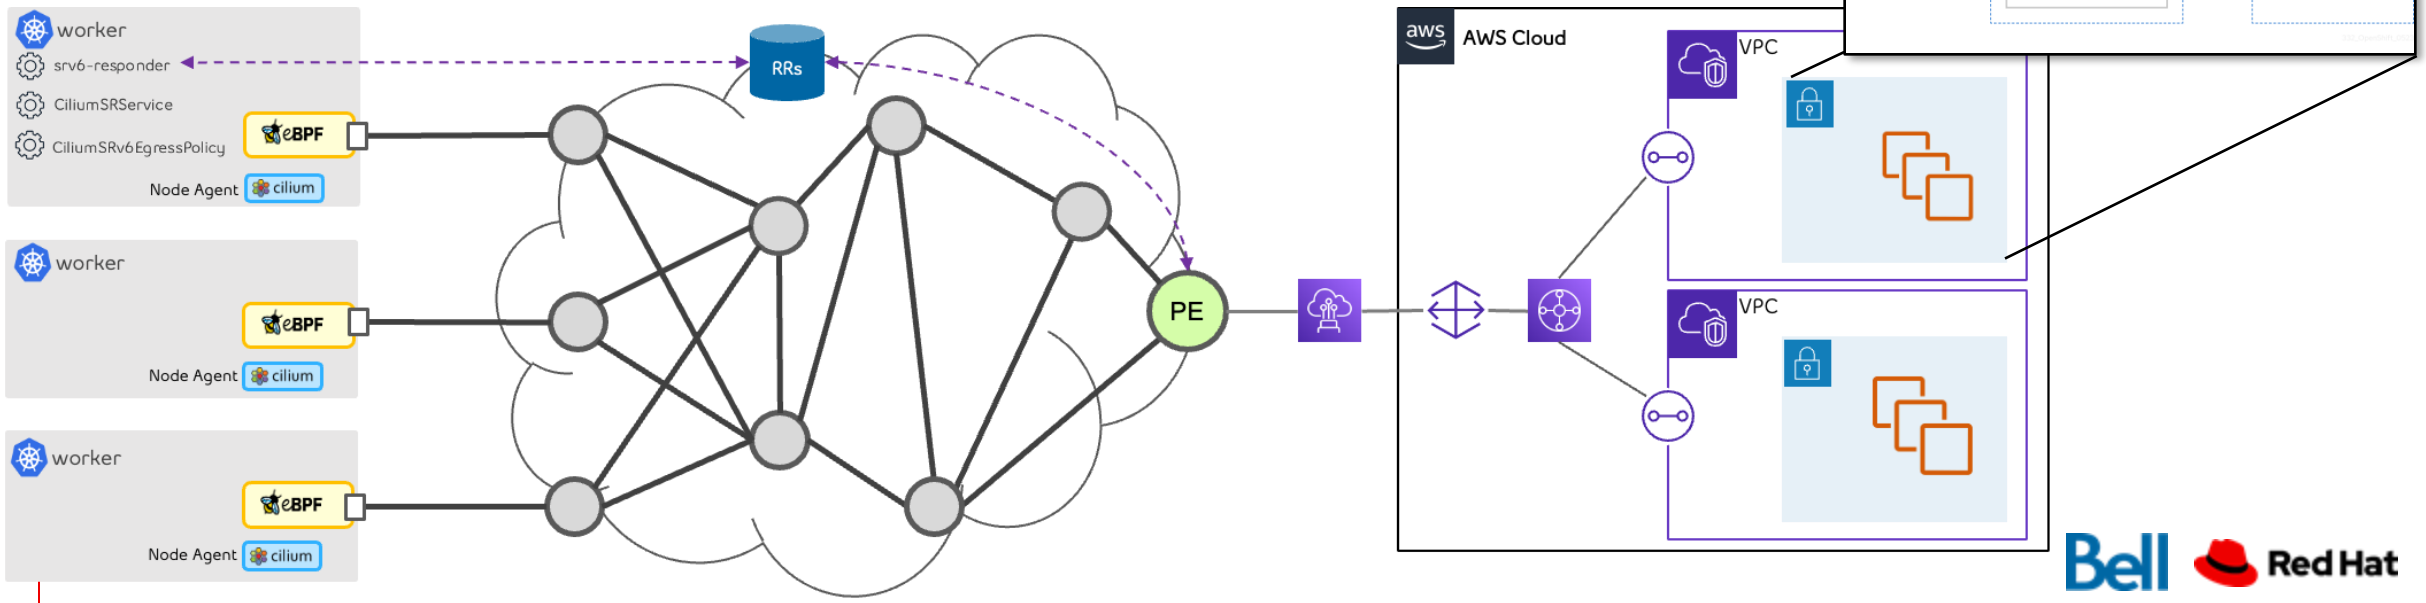

Bell Red Hat

Automating SRv6 the cloud-native way

- SRv6 configuration build using Cilium Enterprise CRDs
- Configurations pushed through GitOps either at cluster build or day 2

sidpool.yaml 361 B

```yaml
apiVersion: isovalent.com/v1alpha1
kind: IsovalentSRv6LocatorPool
metadata:
  name: pool0
  labels:
    # This label will be used in the next step
    export: "true"
spec:
  # behaviorType: Base
  behaviorType: uSID
  prefix: fccc:cc00:0800::/40
  structure:
    locatorBlockLenBits: 32
    locatorNodeLenBits: 16
    functionLenBits: 16
    argumentLenBits: 0

```

bgppeer.yaml 1.09 KiB

```yaml
apiVersion: "cilium.io/v2alpha1"
kind: CiliumBGPPeeringPolicy
metadata:
  name: rr
spec:
  nodeSelector:
    matchLabels:
      # kubernetes.io/hostname: acm1
      bgppeer: rr
  virtualRouters:
  - localASN: 64577
    srv6LocatorPoolSelector:
      matchLabels:
        export: "true"
    exportPodCIDR: true
    mapSRv6VRFs: true
    neighbors:
    - peerAddress: "fccc:cc00:a::1/128"
    # - peerAddress: "fc00:1::254/128"
      # eBGPMultihopTTL: 255
      peerASN: 64577
      families:
      - afi: ipv4
        safi: mpls_vpn
      # - afi: ipv4
      #   safi: unicast
      # - afi: "ipv6"
      #   safi: "unicast"
    - peerAddress: "fccc:cc00:800::253/128"
      # eBGPMultihopTTL: 255
      peerASN: 64577
      families:
```

srv6-covenml.yaml 318 B

```yaml
apiVersion: cilium.io/v2alpha1
kind: CiliumSRv6VRF
metadata:
  name: covenml
spec:
  vrfID: 20
  importRouteTarget: "960:6960"
  exportRouteTarget: "960:6960"
  locatorPoolRef: pool0
  rules:
  - selectors:
    - podSelector:
        matchLabels:
          vrf: covenml
    destinationCIDRs:
    - 192.168.137.128/25

```

pod-covenml.yaml 263 B

```yaml
apiVersion: v1
kind: Pod
metadata:
  labels:
    run: covenml
    vrf: covenml
  name: covenml
  namespace: default
spec:
  containers:
  - command:
    - sleep
    - "10000000"
    image: alpine
    imagePullPolicy: Always
    name: test
  restartPolicy: Always
```


Bell Red Hat

Now how could we make private cloud deployment even simpler ?

- Rethinking the need to use BGP (or even IGP) for HBN (host-based networking)
- Collaboration with Cisco on design of Lightweight Host Routing
- Leveraging LLDP-TLVs to a prefix (not limited to SRv6 Locator)
- Fabric can be a pure IPv6 routing device
- Automated through GitOps either via operators or daemonset

Cisco-Bell Innovation
Simplifying locator advertisement

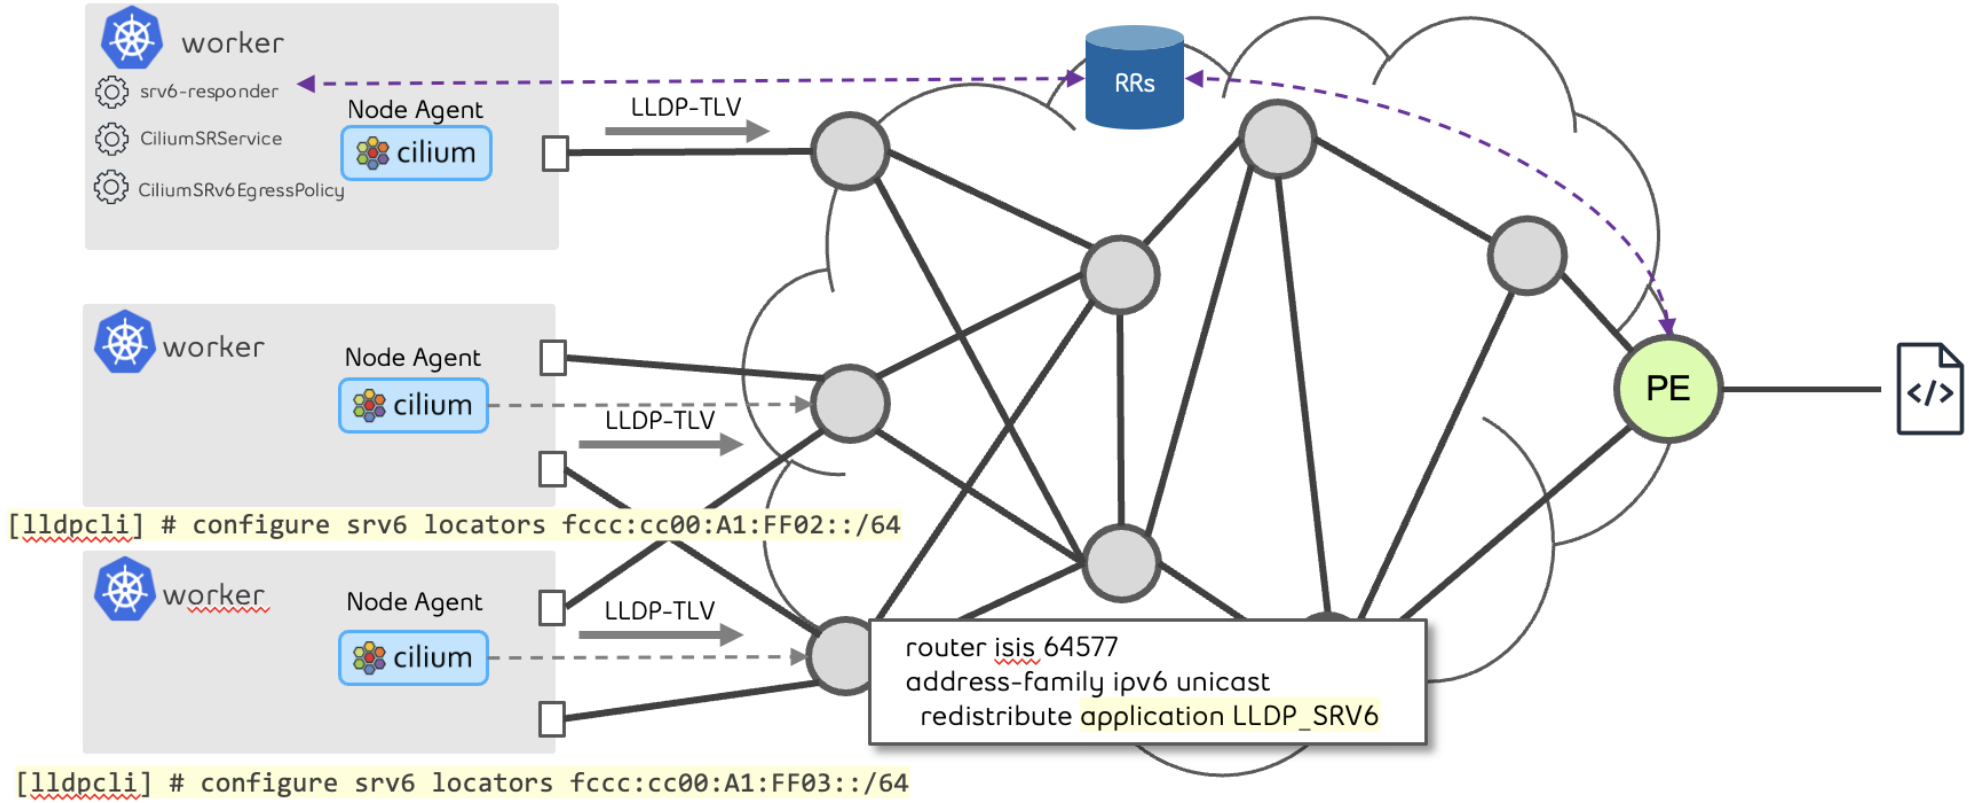

How should we actually auto-configure and deploy across our k8s fleet ?

Bell Red Hat

Edge fleet management … The problem of scaling out

- A large fleet of Kubernetes clusters is necessary for deploying edge clouds
- Central management for large scale deployments is key to success
- The goal of "*human-less*" provisioning required a new approach
- Bonus … could we take care of our k8s clusters on public cloud as well ?

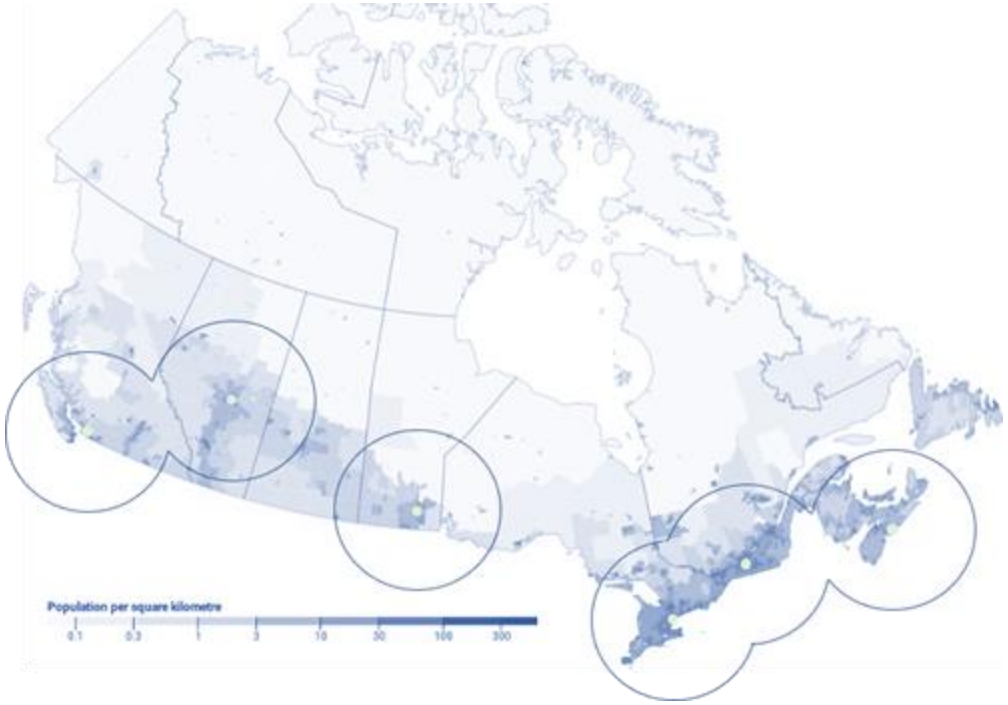

Global DCs covering 90% of population with regular latency

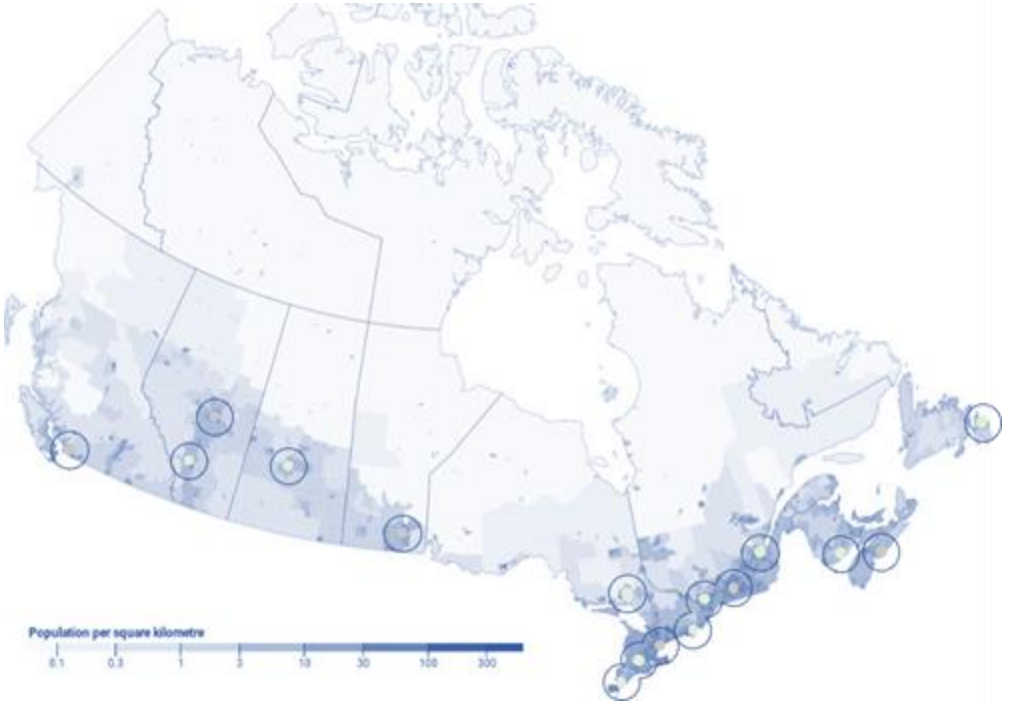

Edge clouds to increase the predictable latency on the other areas

As we expand into farther locations – centralized automated operations is key

Bell Red Hat

The Bell Vision - Central Management for Edge Cloud

- Be Declarative
- Commoditized Deployment
- Accelerated Development
- Allow quick handoff to operations

Allows evolving from pure GitOps towards config as data and Nephio KRM based approach

https://docs.nephio.org/docs/guides/install-guides/install-on-openshift/

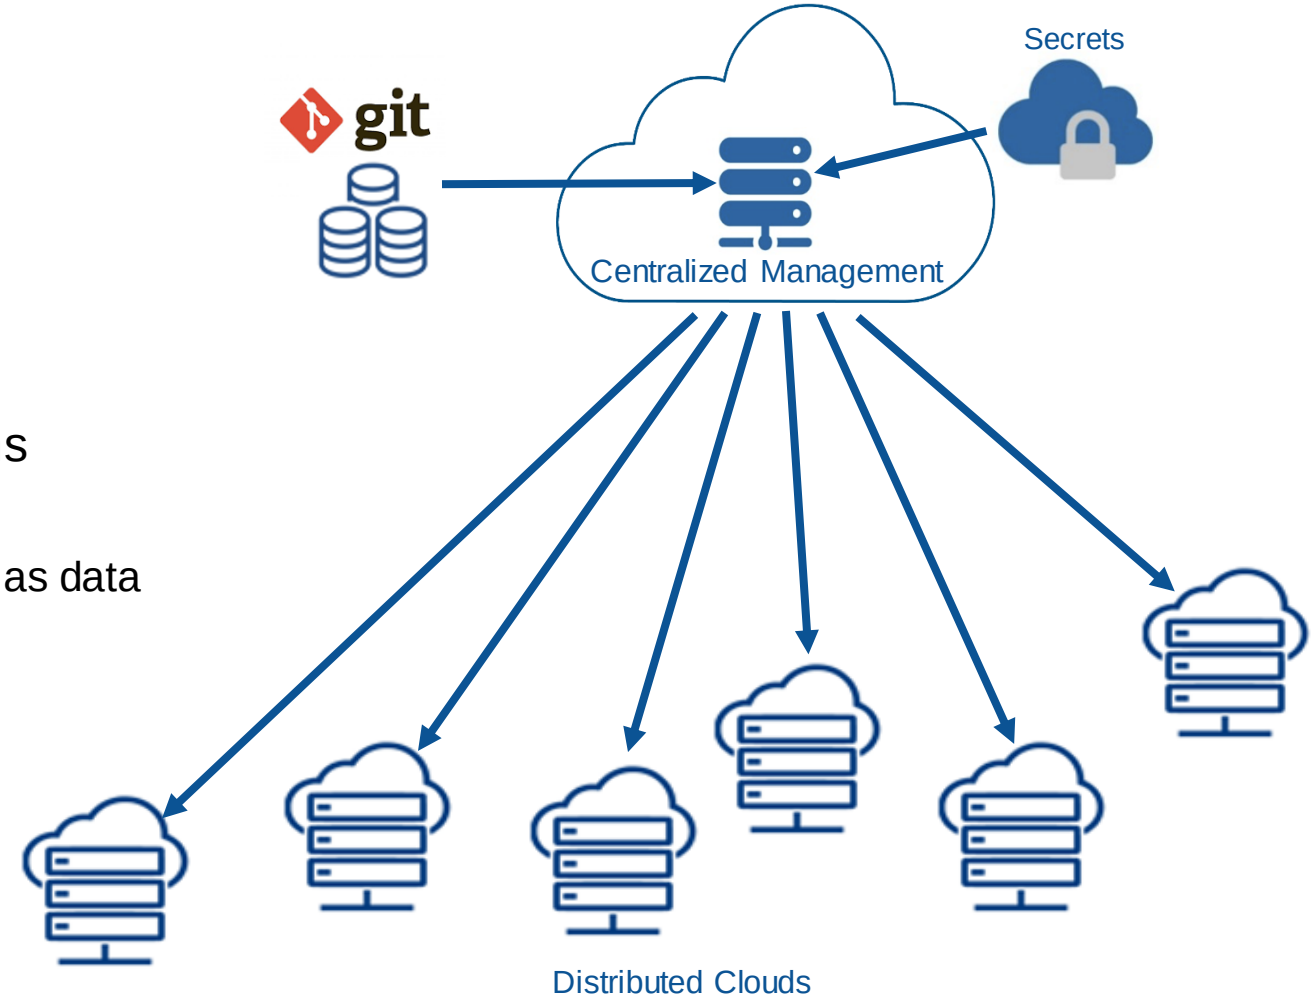

Requirement to support "centrally" managed various fleets of clusters

Bell Red Hat

Build vs Buy Analysis

- Extensive study done internally for a "Build vs Buy" approach

Buy: *turnkey but limited choices*

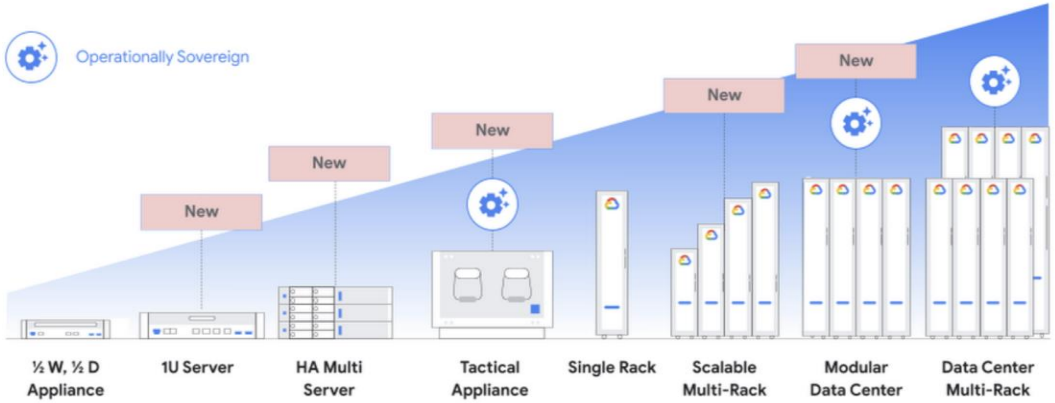

Build*: flexible but complex*

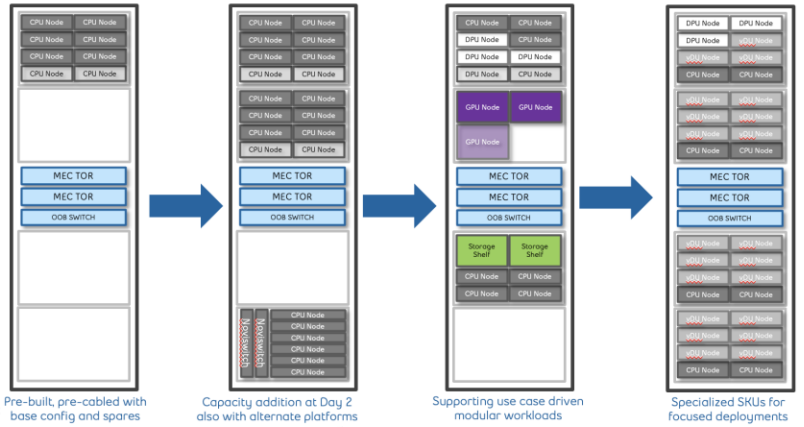

If building, must fill the complexity gaps by working with a partner

Bell Red Hat

Executing Bell's Telco Cloud Vision

MPLS WC 2023: Edge Cloud – The Next Frontier

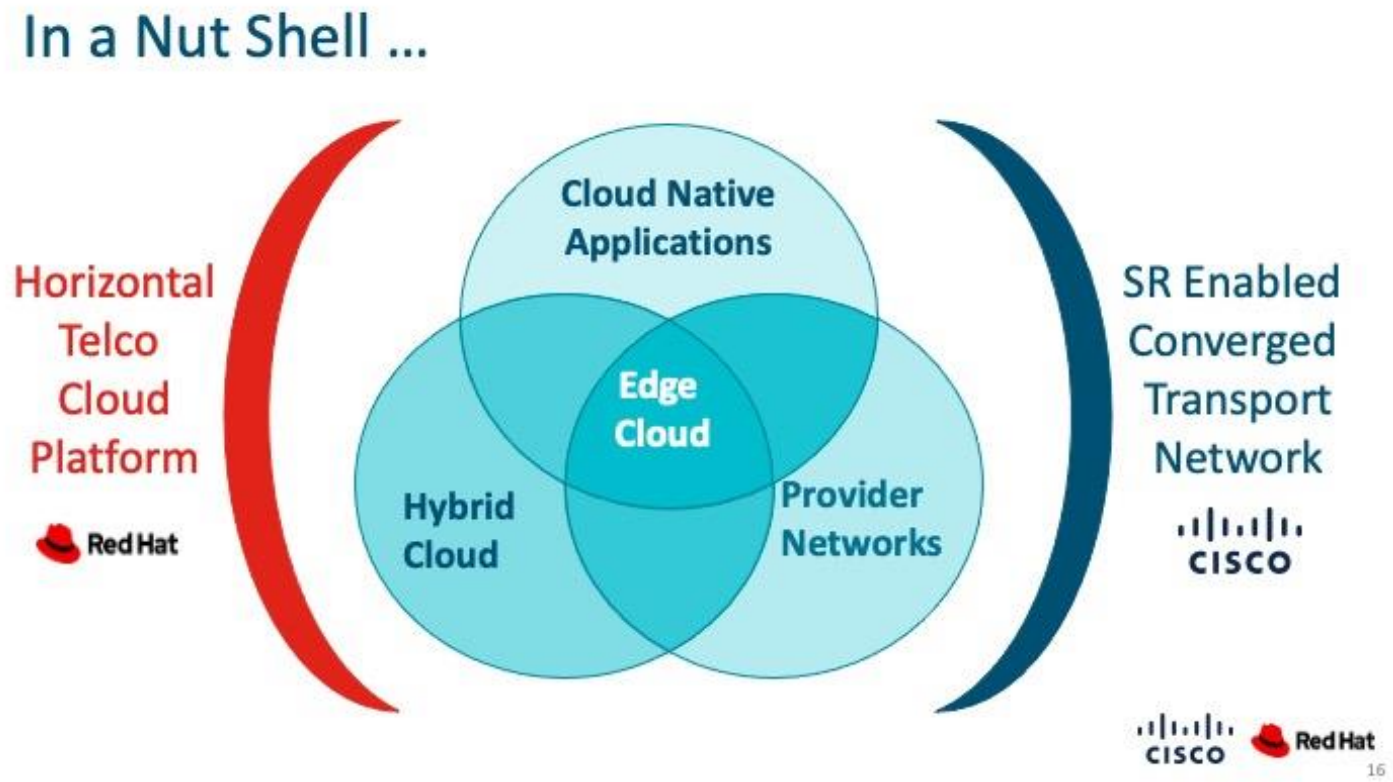

Red Hat: Bringing Hybrid Cloud, Cloud-Native Apps and Provider Networks together!!

Bell Red Hat

Accelerating Edge Cloud Deployment with GitOps

Declarative

Declare the desired infrastructure as code

Single source of truth in Git

Versioned and Immutable

Auditable, version controlled in Git

enforces immutability

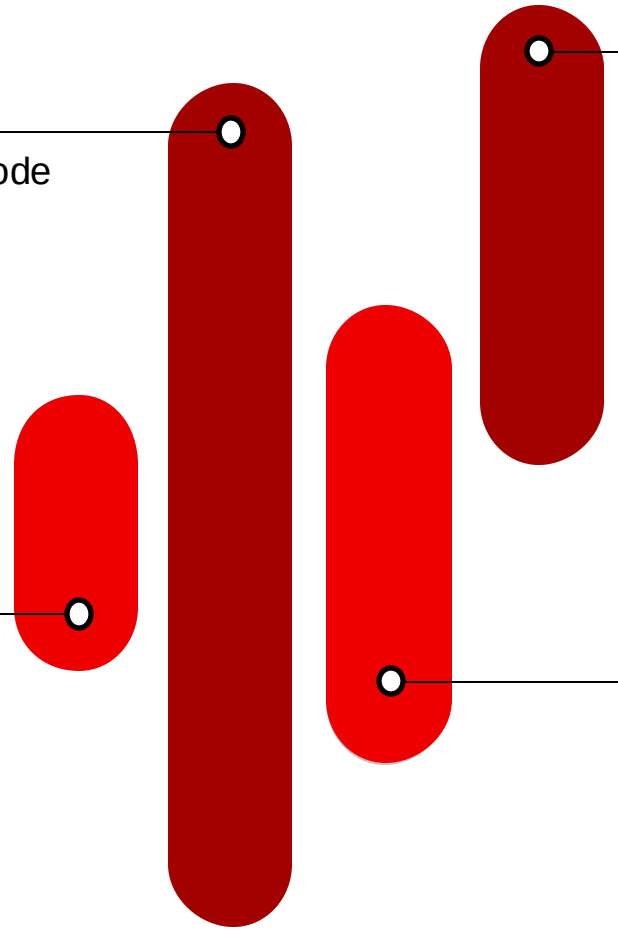

Automatic Sync

GitOps agents provides automatic sync without manual intervention.

Continuous Reconciliation

GitOps Agents learns current state and reconcile with declared state

Bell Red Hat

Declare and Forget – The GitOps Approach

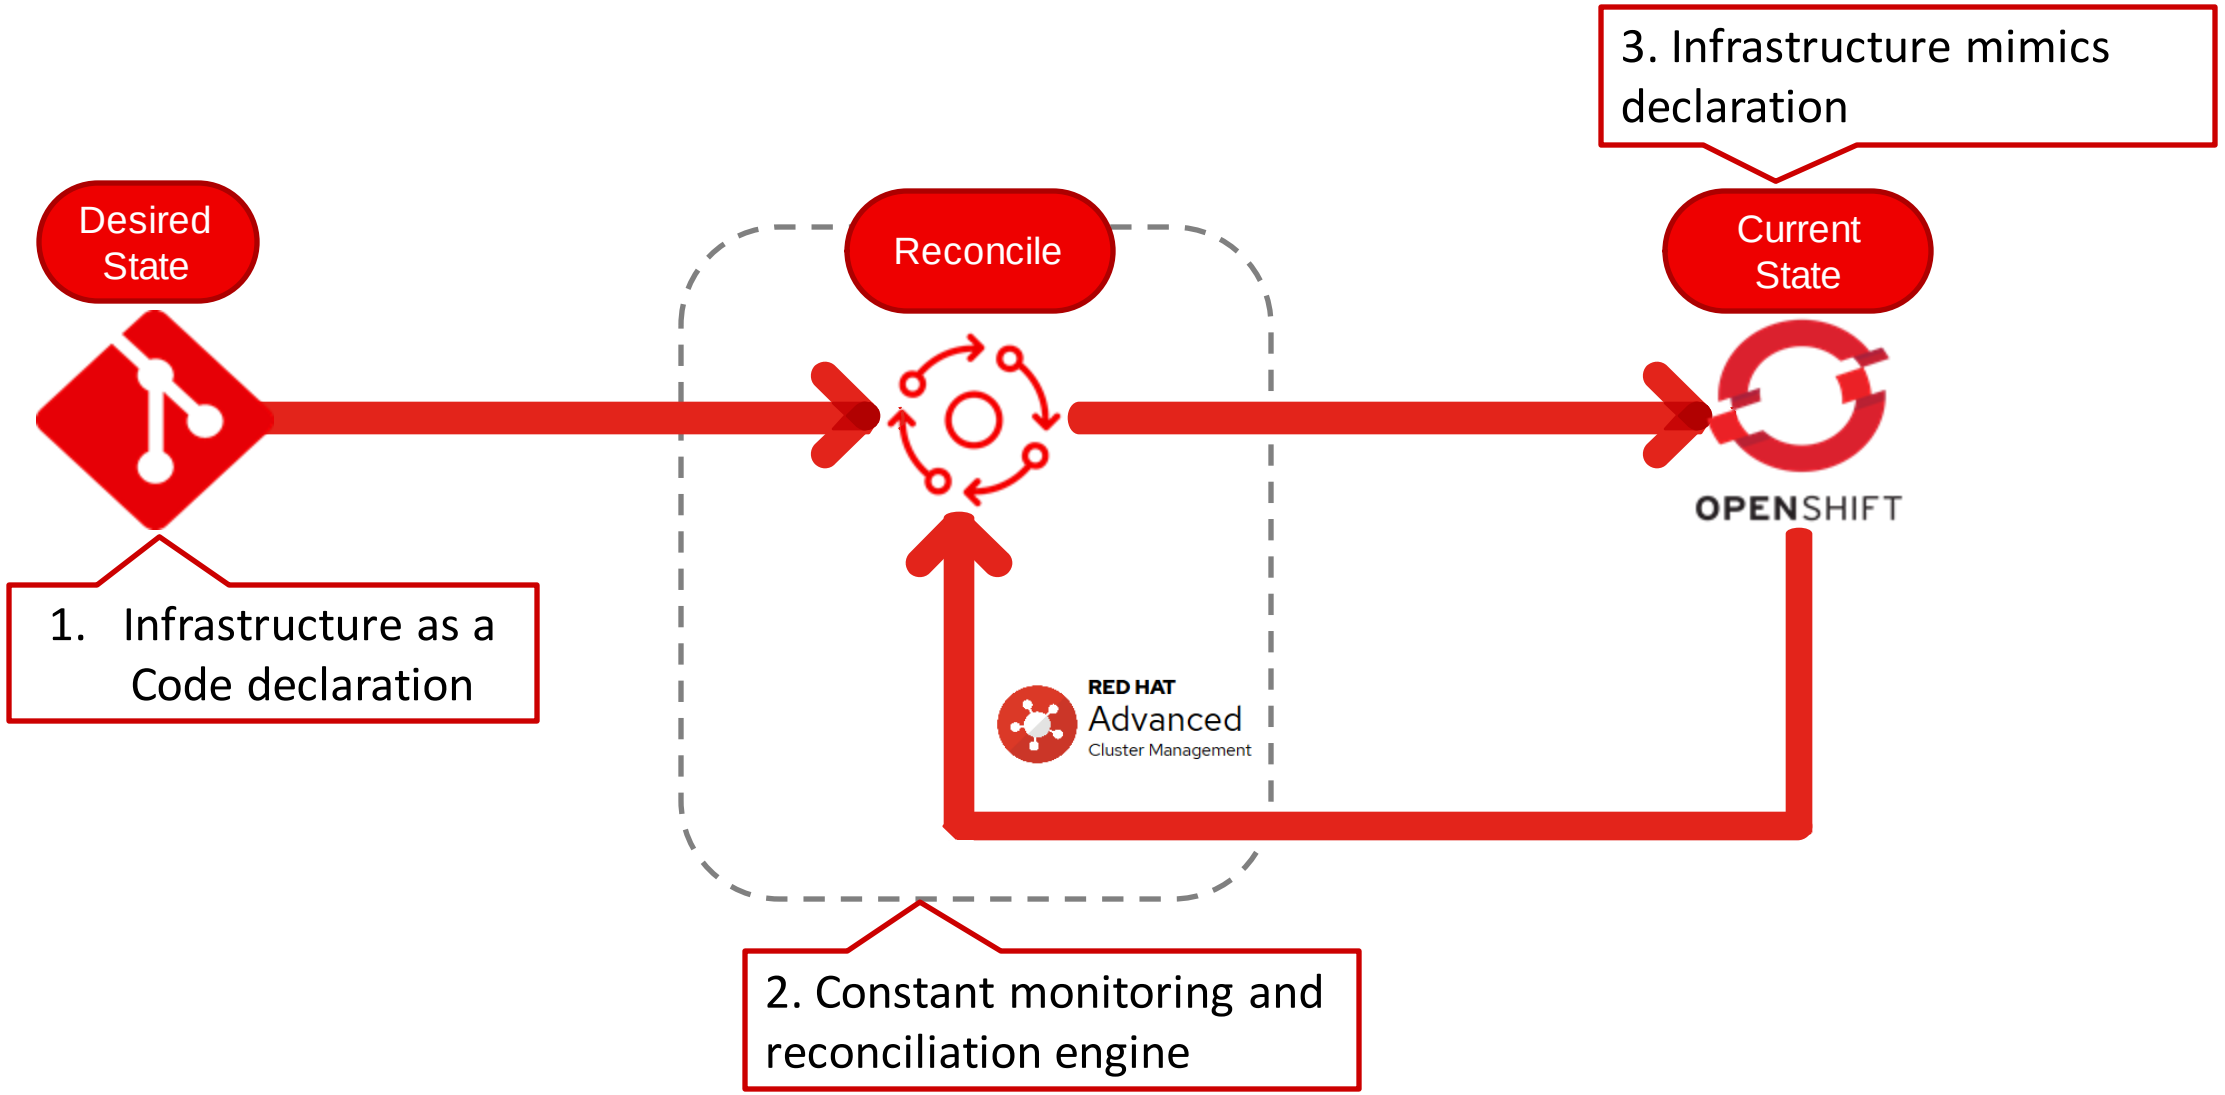

Bell Red Hat

GitOps ZTP for Cloud Platform

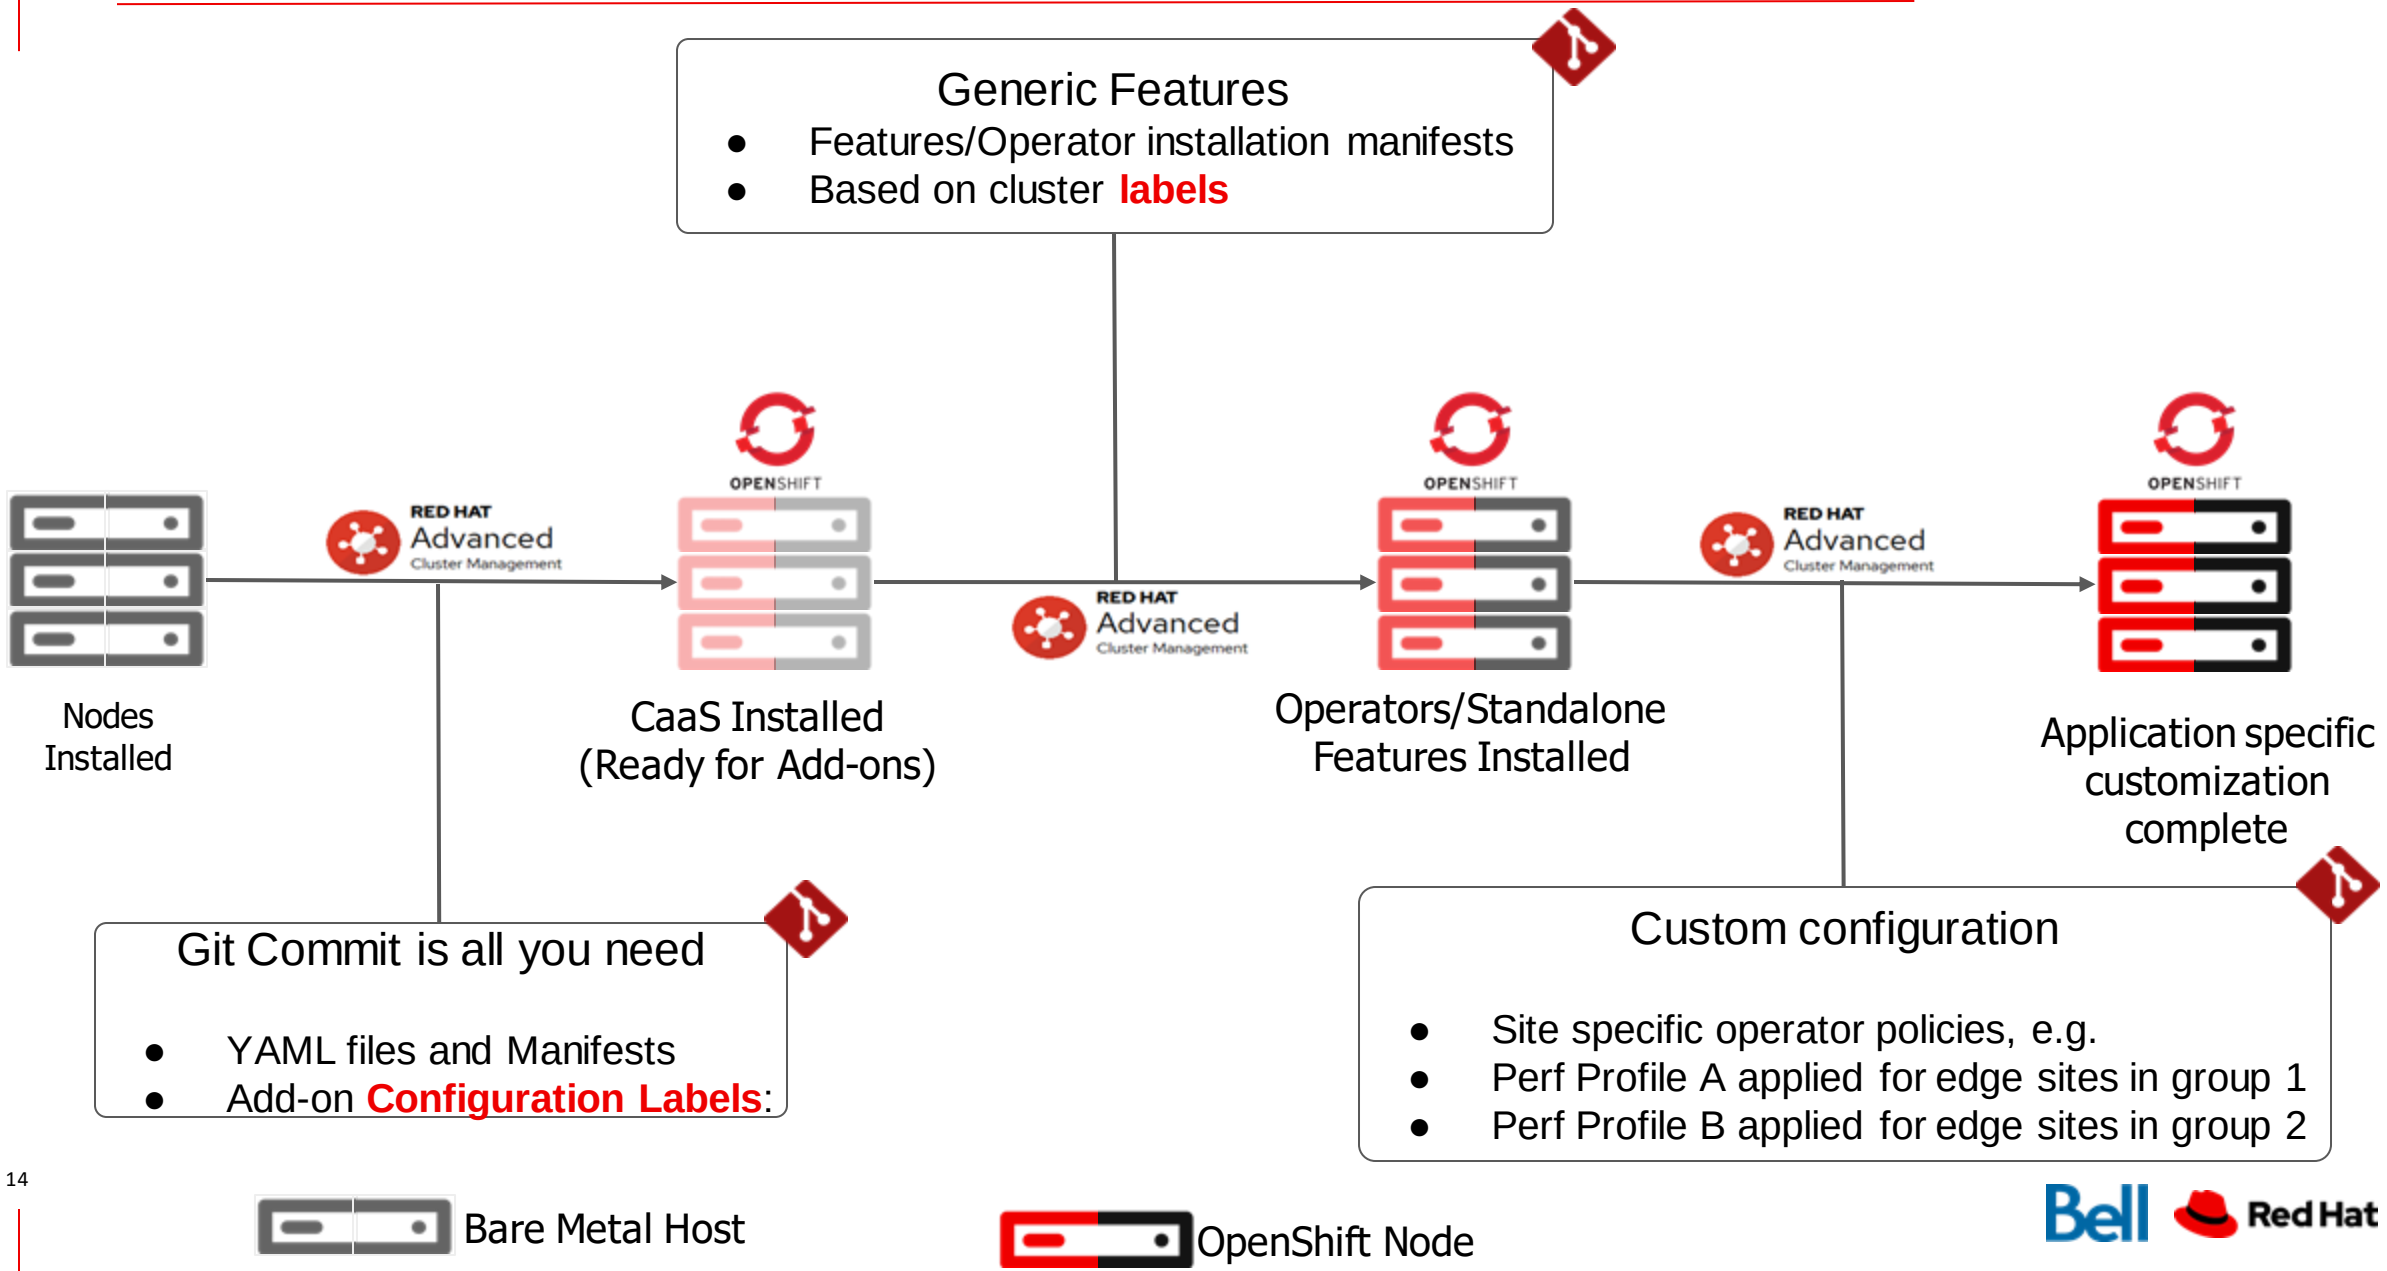

Bell's Vision: Realized !!

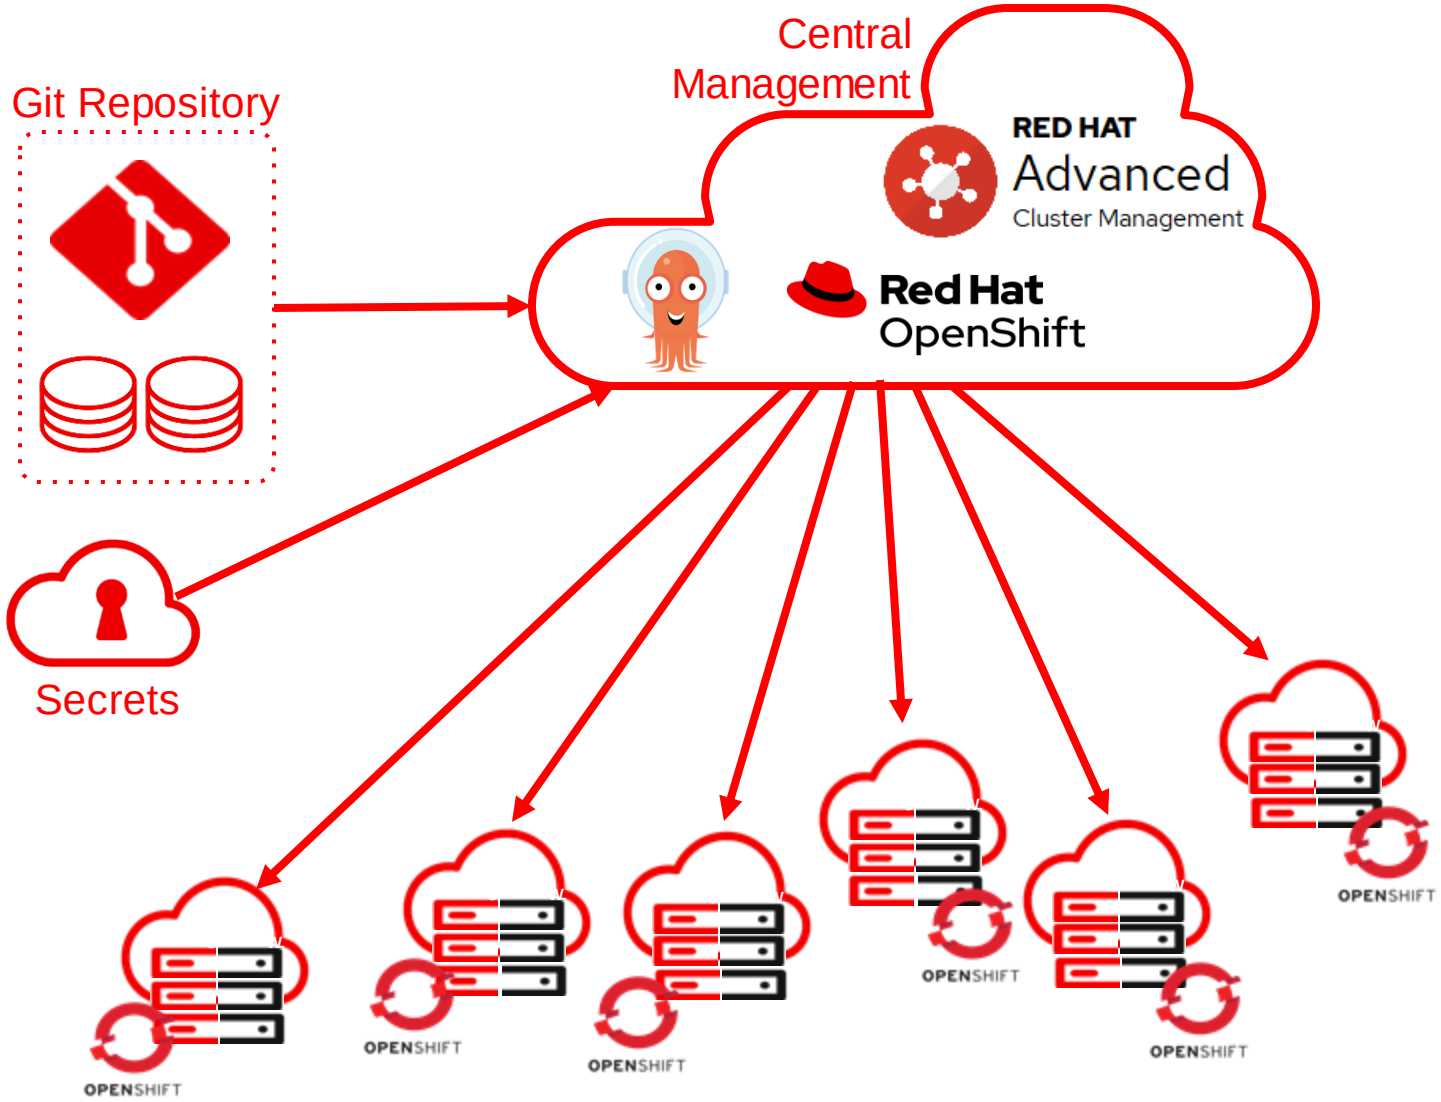

Distributed Cloud with
Red Hat OpenShift

- Be Declarative
- Commoditized Deployment
- Accelerated Development
- Quick handoff to Ops

Bell Red Hat

Telco Cloud Management: Any Infrastructure, Any Cloud

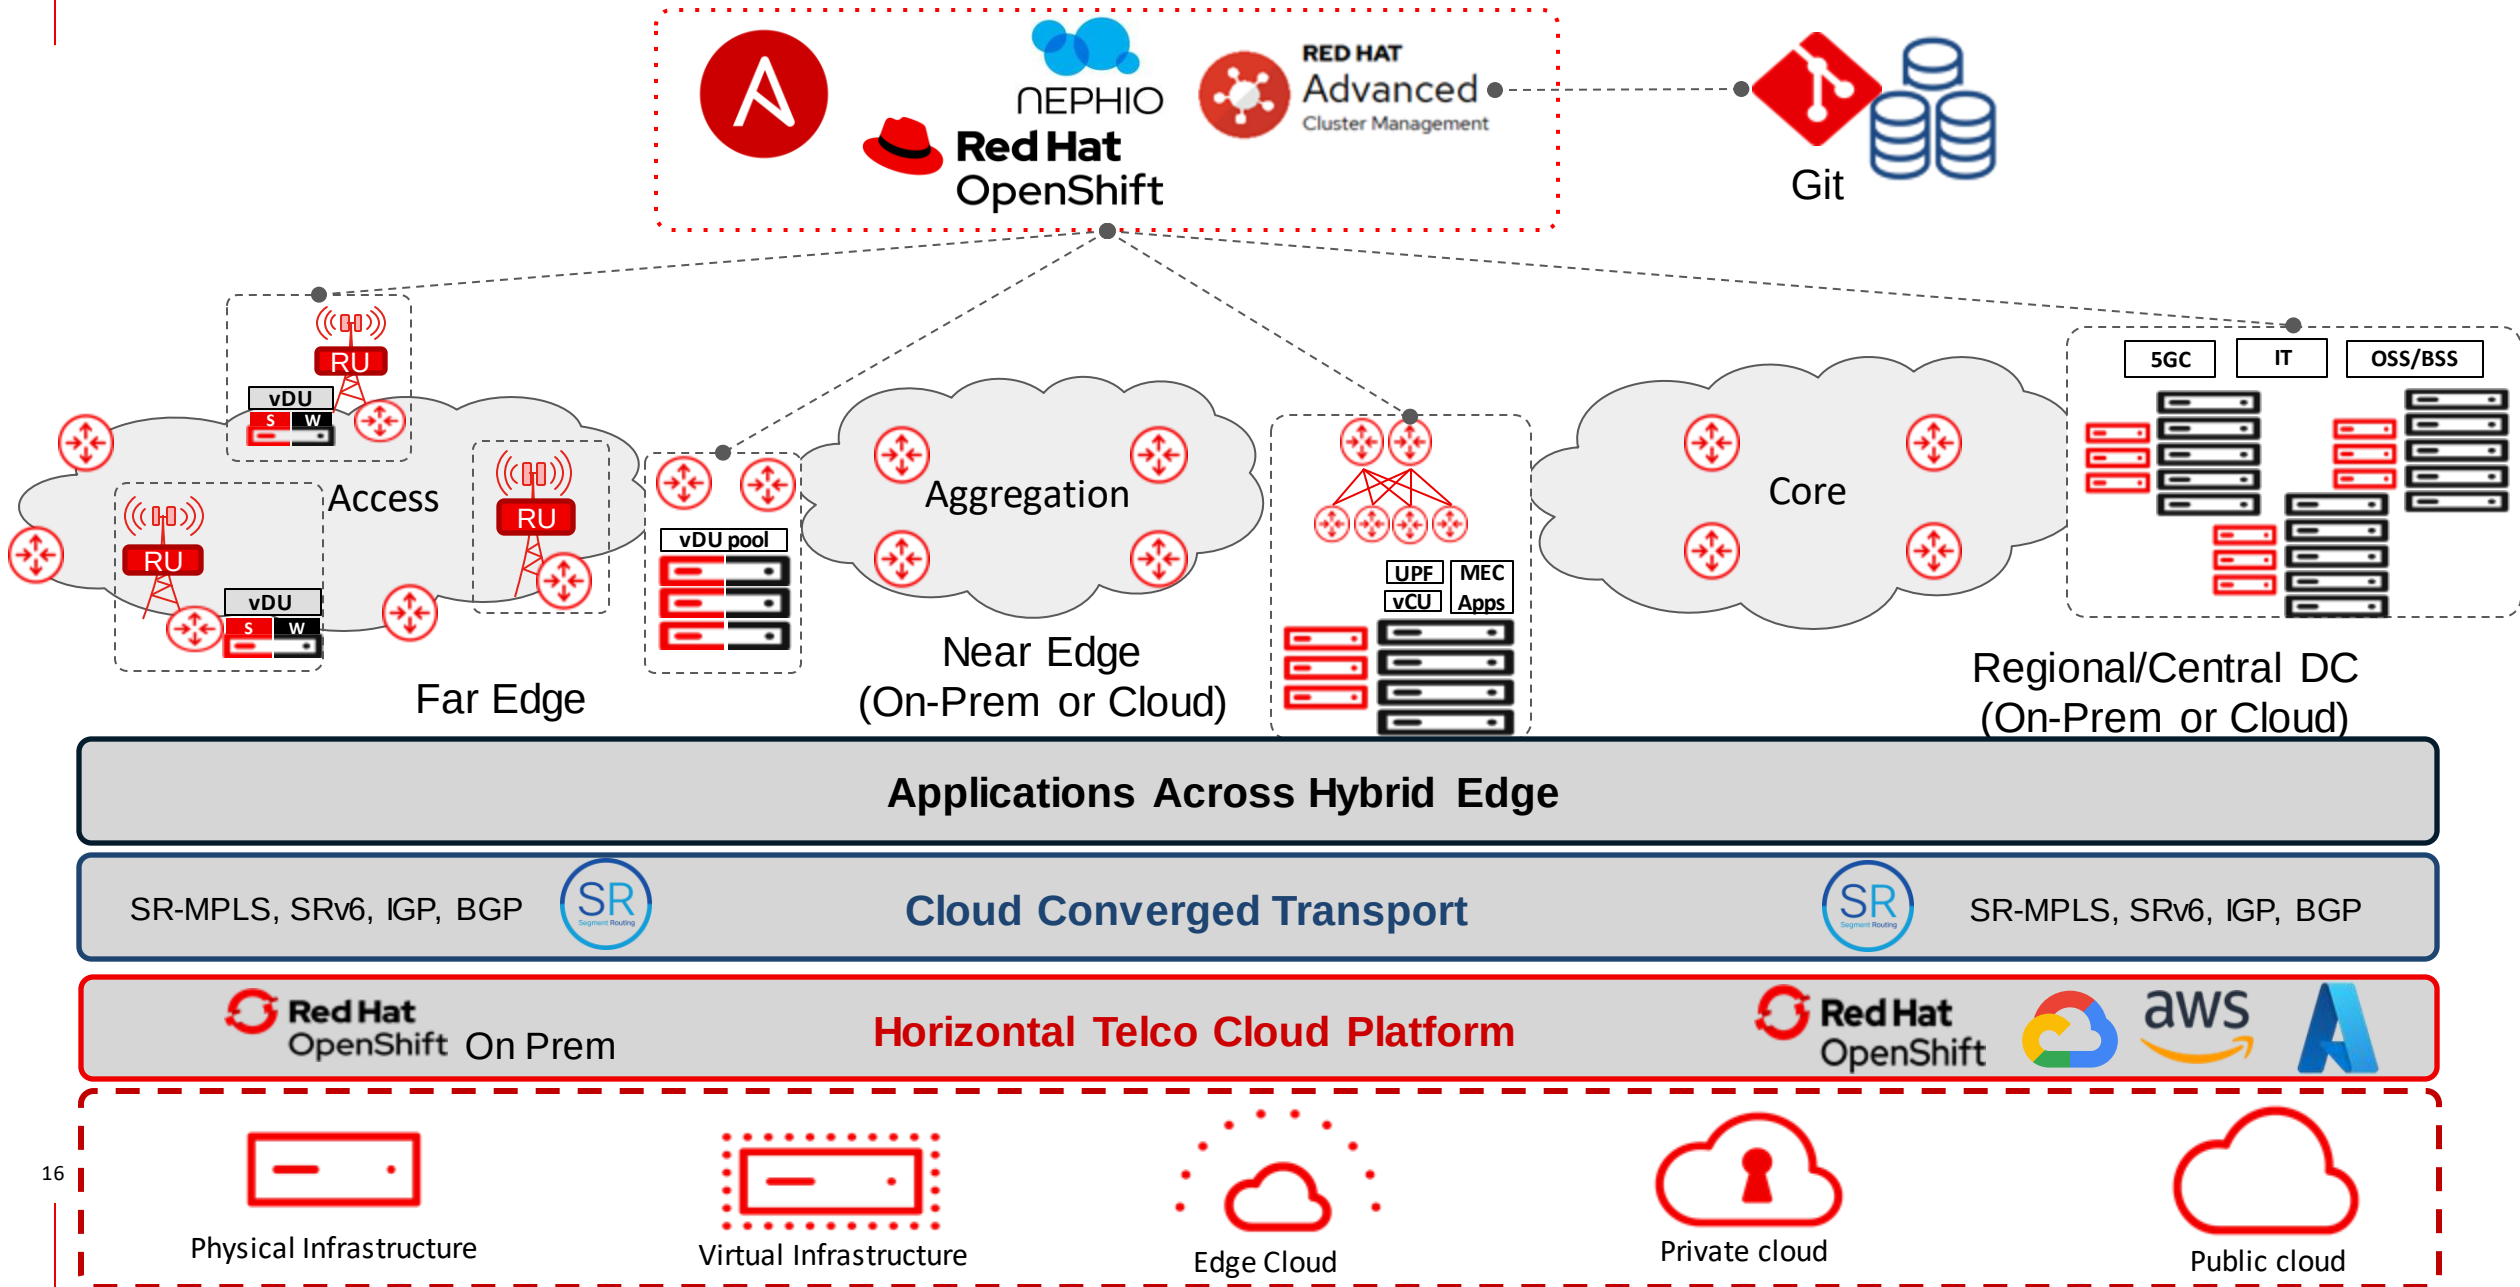

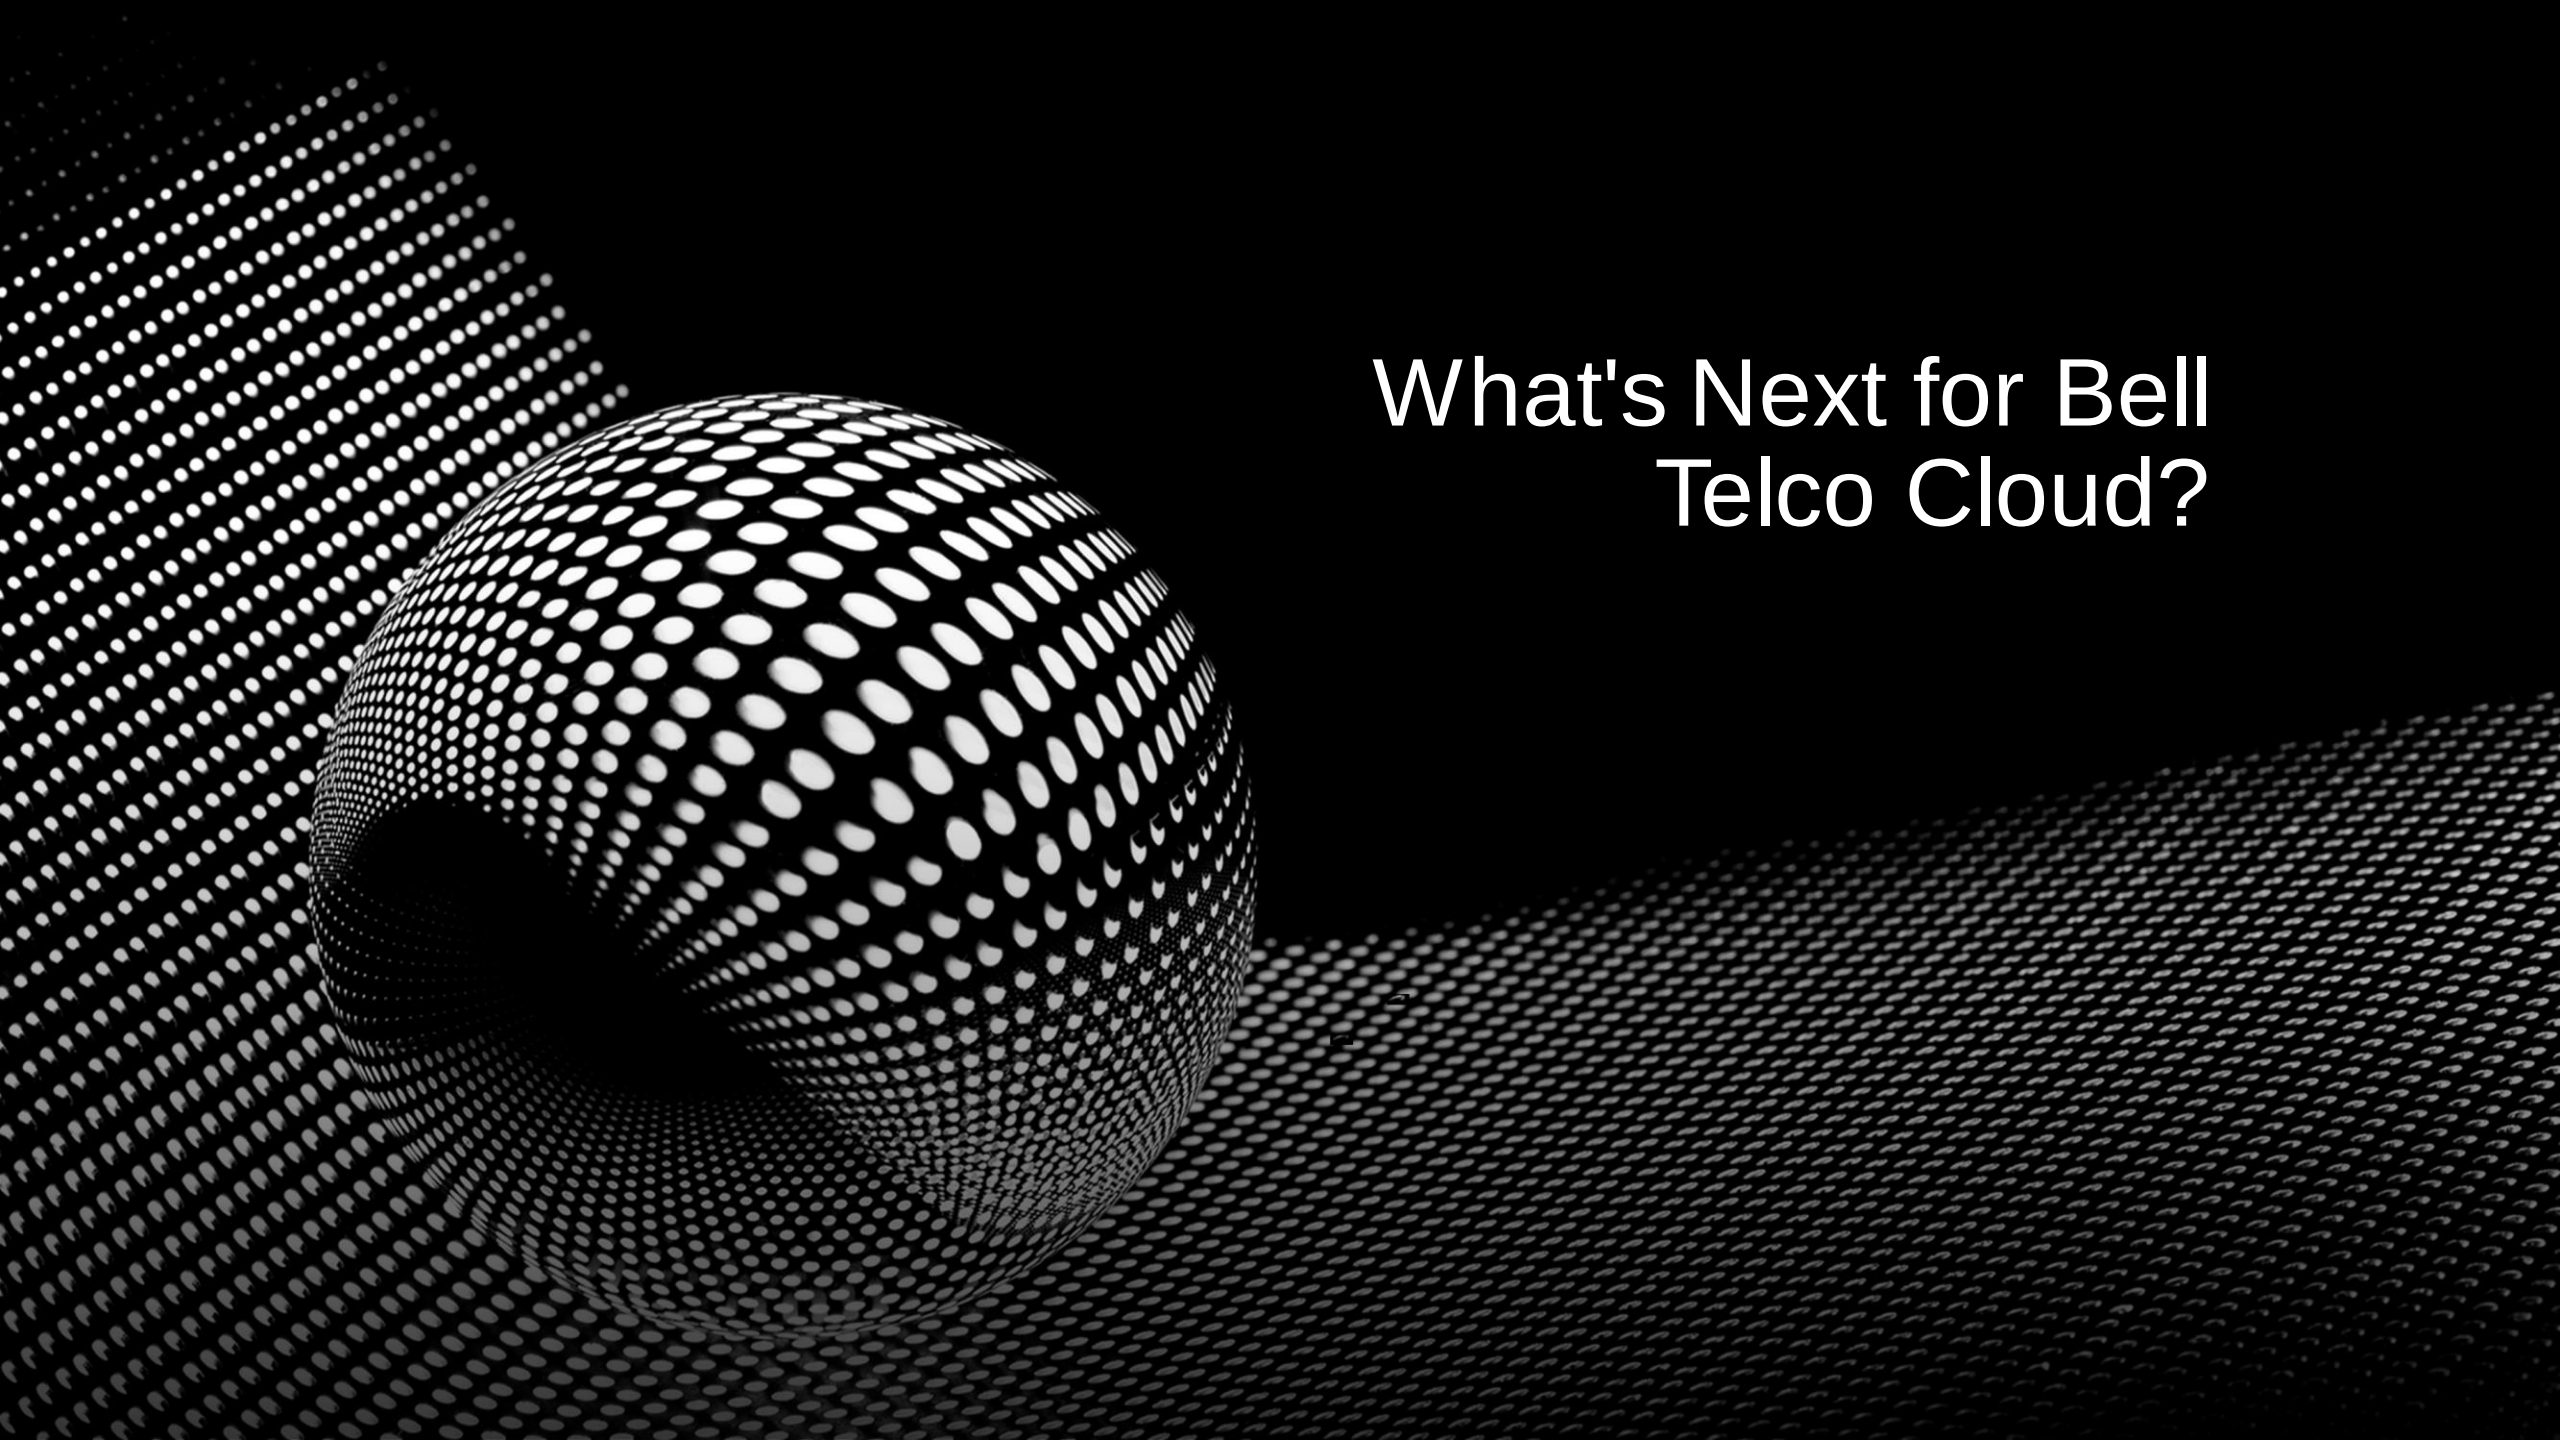
What's Next for Bell
Telco Cloud?

Expanding the cluster management boundaries to ORAN

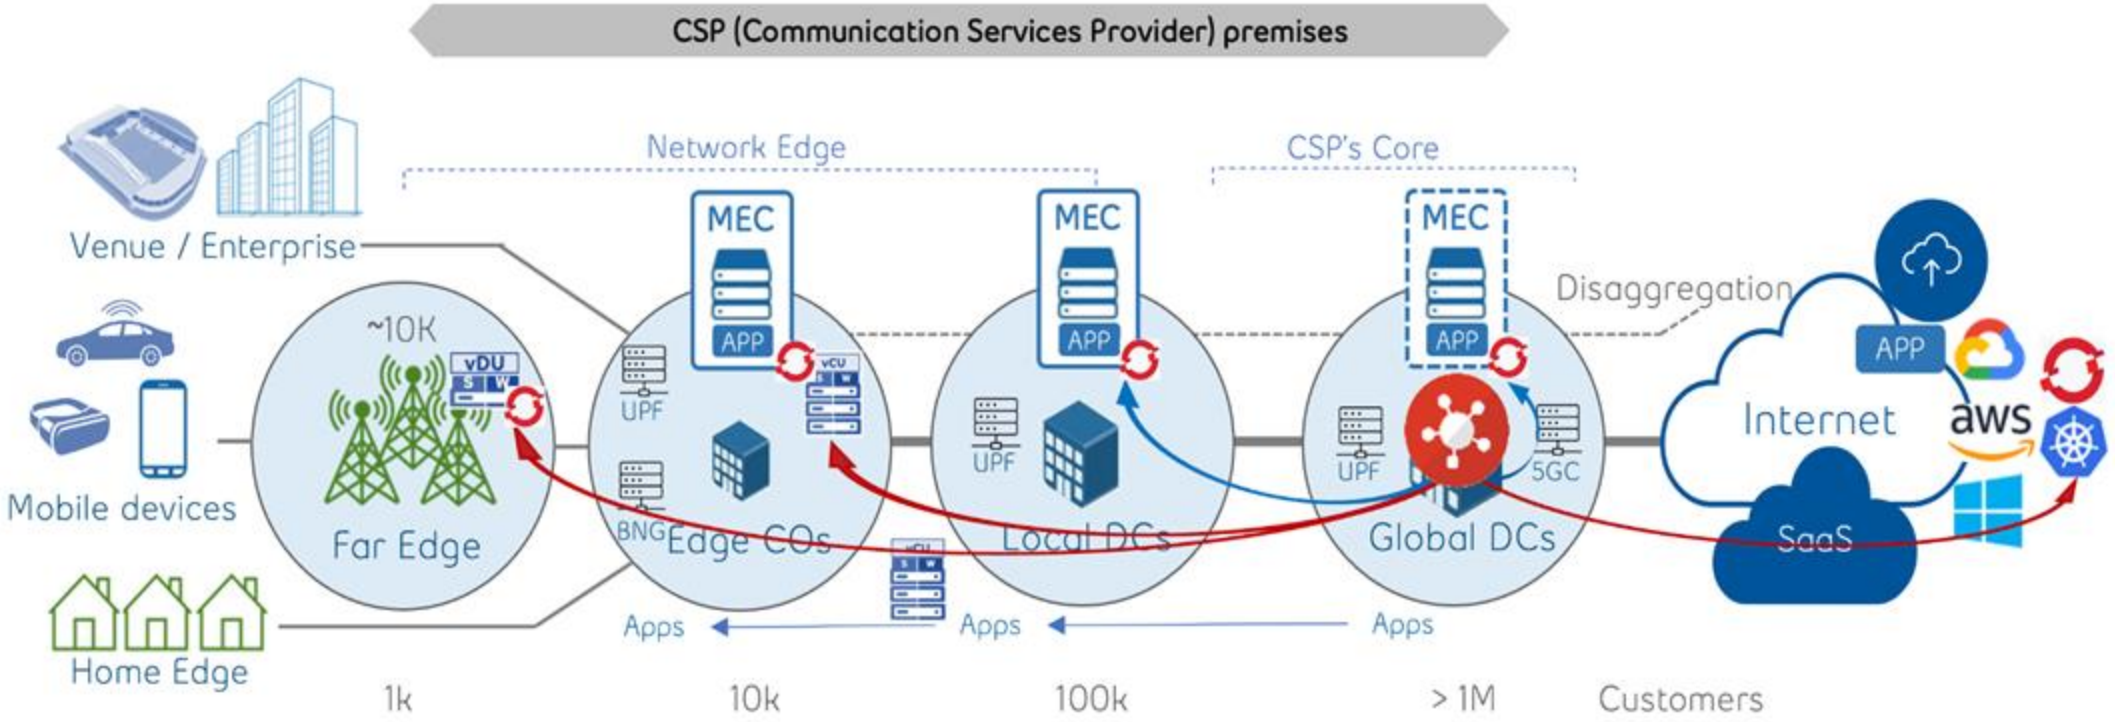

Hypothesis: Red Hat ACM can be leveraged by Bell not only to manage vRAN at the Far Edge, but also Cloud k8s engines

Bell Red Hat

Evolution towards Ubiquitous Computing

- Application hosting form factors are and will continue evolving ?
- We must be able to expand the GitOps approach to evolving frameworks ?
- Can the Kubernetes central management approach be reused for these new applications ?
- At what inflexion do frameworks like BBF USP (TR-369) evolve towards distributed application frameworks ?

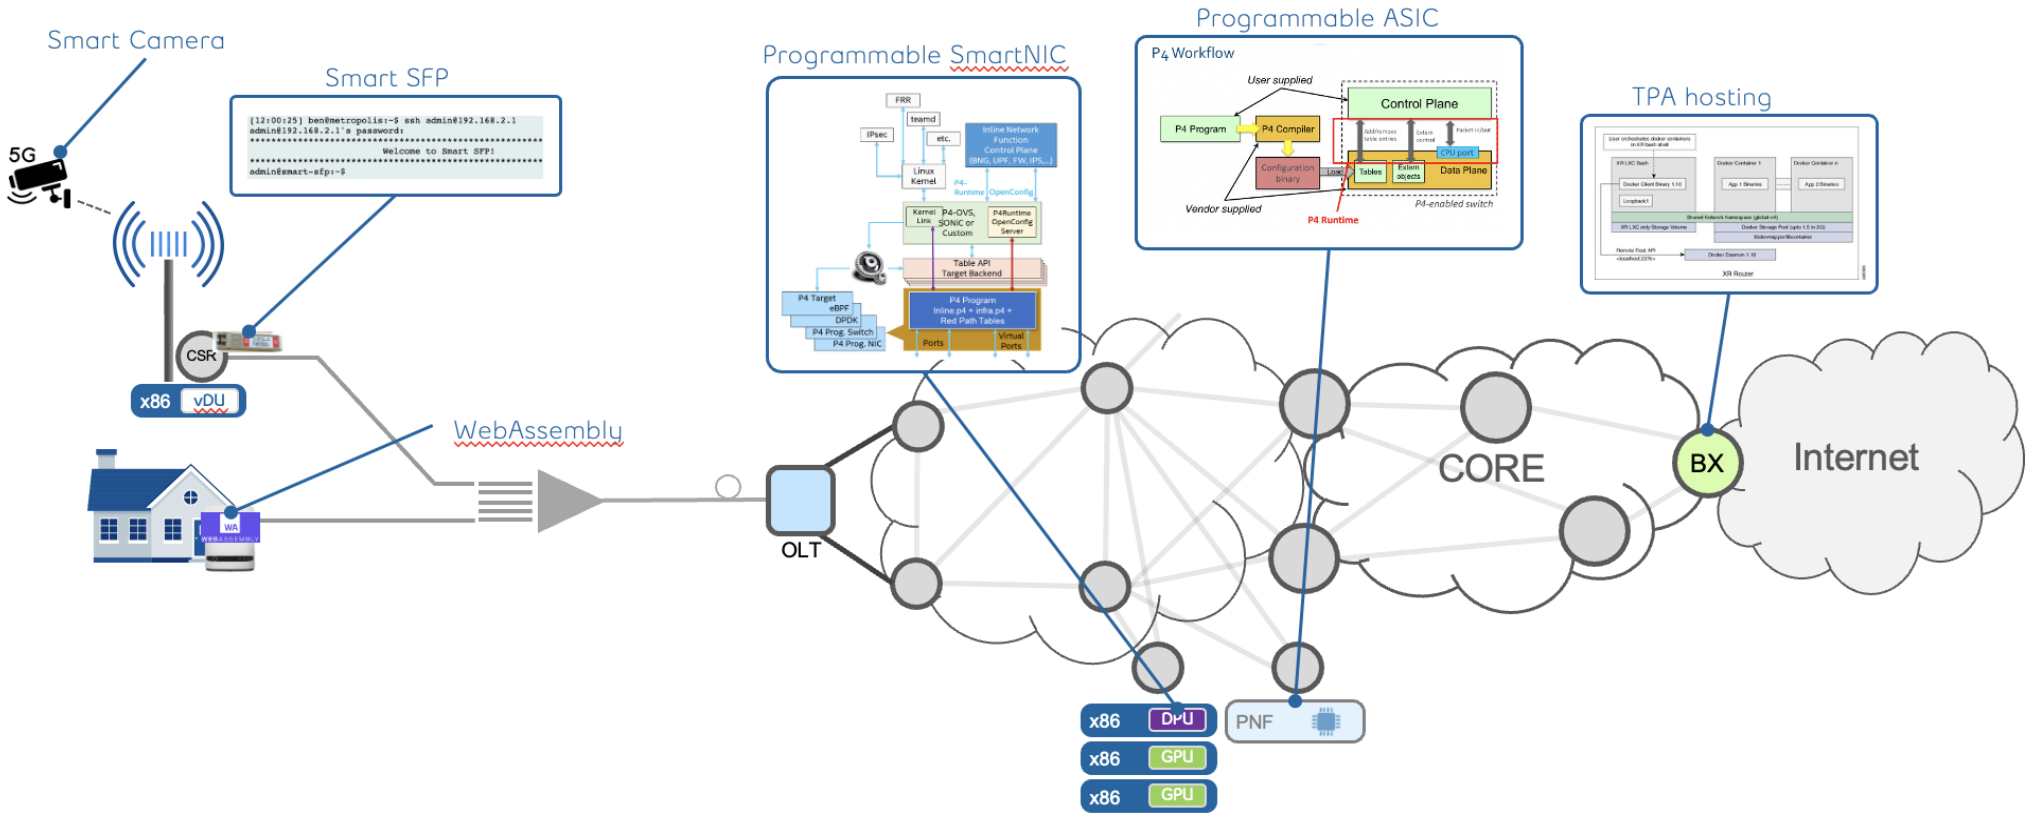

Beyond Edge ... Expanding GitOps and Red Hat ACM

Beyond Openshift management

- ACM managing clusters in Cloud k8s engines (GKE, EKS, AKS)
- Lightweight deployments: ACM + Hypershift

Proactive Management

- Scaling ACM: growing beyond single Hub (i.e. the Multi-cluster Global Hub concept)
- Telemetry collection with ACM

Sustainability

- Improve sustainability of the OpenShift Ecosystem
- ACM to hibernate OCP on Public Clouds (GCP, AWS, Azure)

Thank you!
upperside conferences
paris
© 1994-2024 UppersideConferences. contact@uppersideconferences.com