Architecting the Next Generation DCN for **Flexibility and Scale with Optics and SDN**

Katherine Barabash

IBM Research – Haifa

COSIGN

Cloud Computing:

A style of computing where massively scalable ICT capabilities are delivered as a service over the network

Separation

- *Infrastructure details are not exposed to customers/partners*
- *Workload details are not exposed to the provider*

Isolation

- *Workloads are fully isolated from each other*

Elasticity

- *Resource allocation is automatically adjusted to the dynamically changing workload needs*

Self-service

- *Customers can operate their leased cloud portion through UI/API*

Cloud Computing:

A style of computing where massively scalable ICT capabilities are delivered as a service over the network

Virtualization

- ***Compute –*** *Virtual Machines, Containers*
- ***Storage –*** *Logical Volume, File, Stored Object*
- ***Network***
  - *Traditional – VLANs*
  - *Emerging – virtualization overlays*
    - *Host or infrastructure based*
    - *Software or hardware*

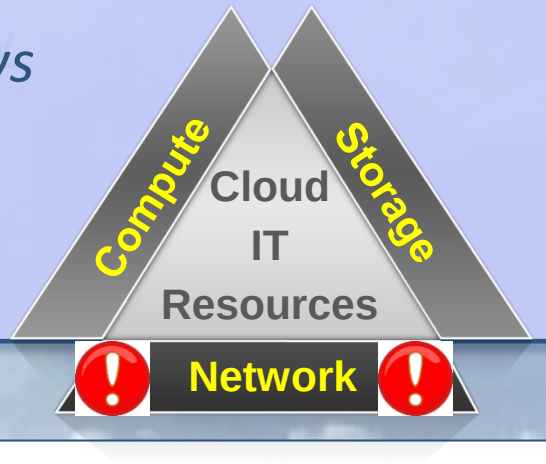

Cloud Computing:

A style of computing where massively scalable ICT capabilities are delivered as a service over the network

Premises:

1. **To deliver Separation, Isolation, Elasticity, and Self-service, ICT resources, namely Compute, Storage, and Network, must be fully and efficiently virtualized**
2. **To cope with the growing demand, Data Centre Network must be redesigned for lower latency, higher capacity, hands-off manageability, and greater flexibility**

COSIGN:

Investigate data plane and control plane technologies for delivering optical advantages to the Data Centre Network

Integrate flexible and controllable COSIGN DCN into the end-to-end cloud orchestration through advanced multi-layer network virtualization

COSIGN – Next Generation Data Center Network

Combining Optics and SDN In Next Generation Data Center Network

COSIGN – Data Center Network Transformation

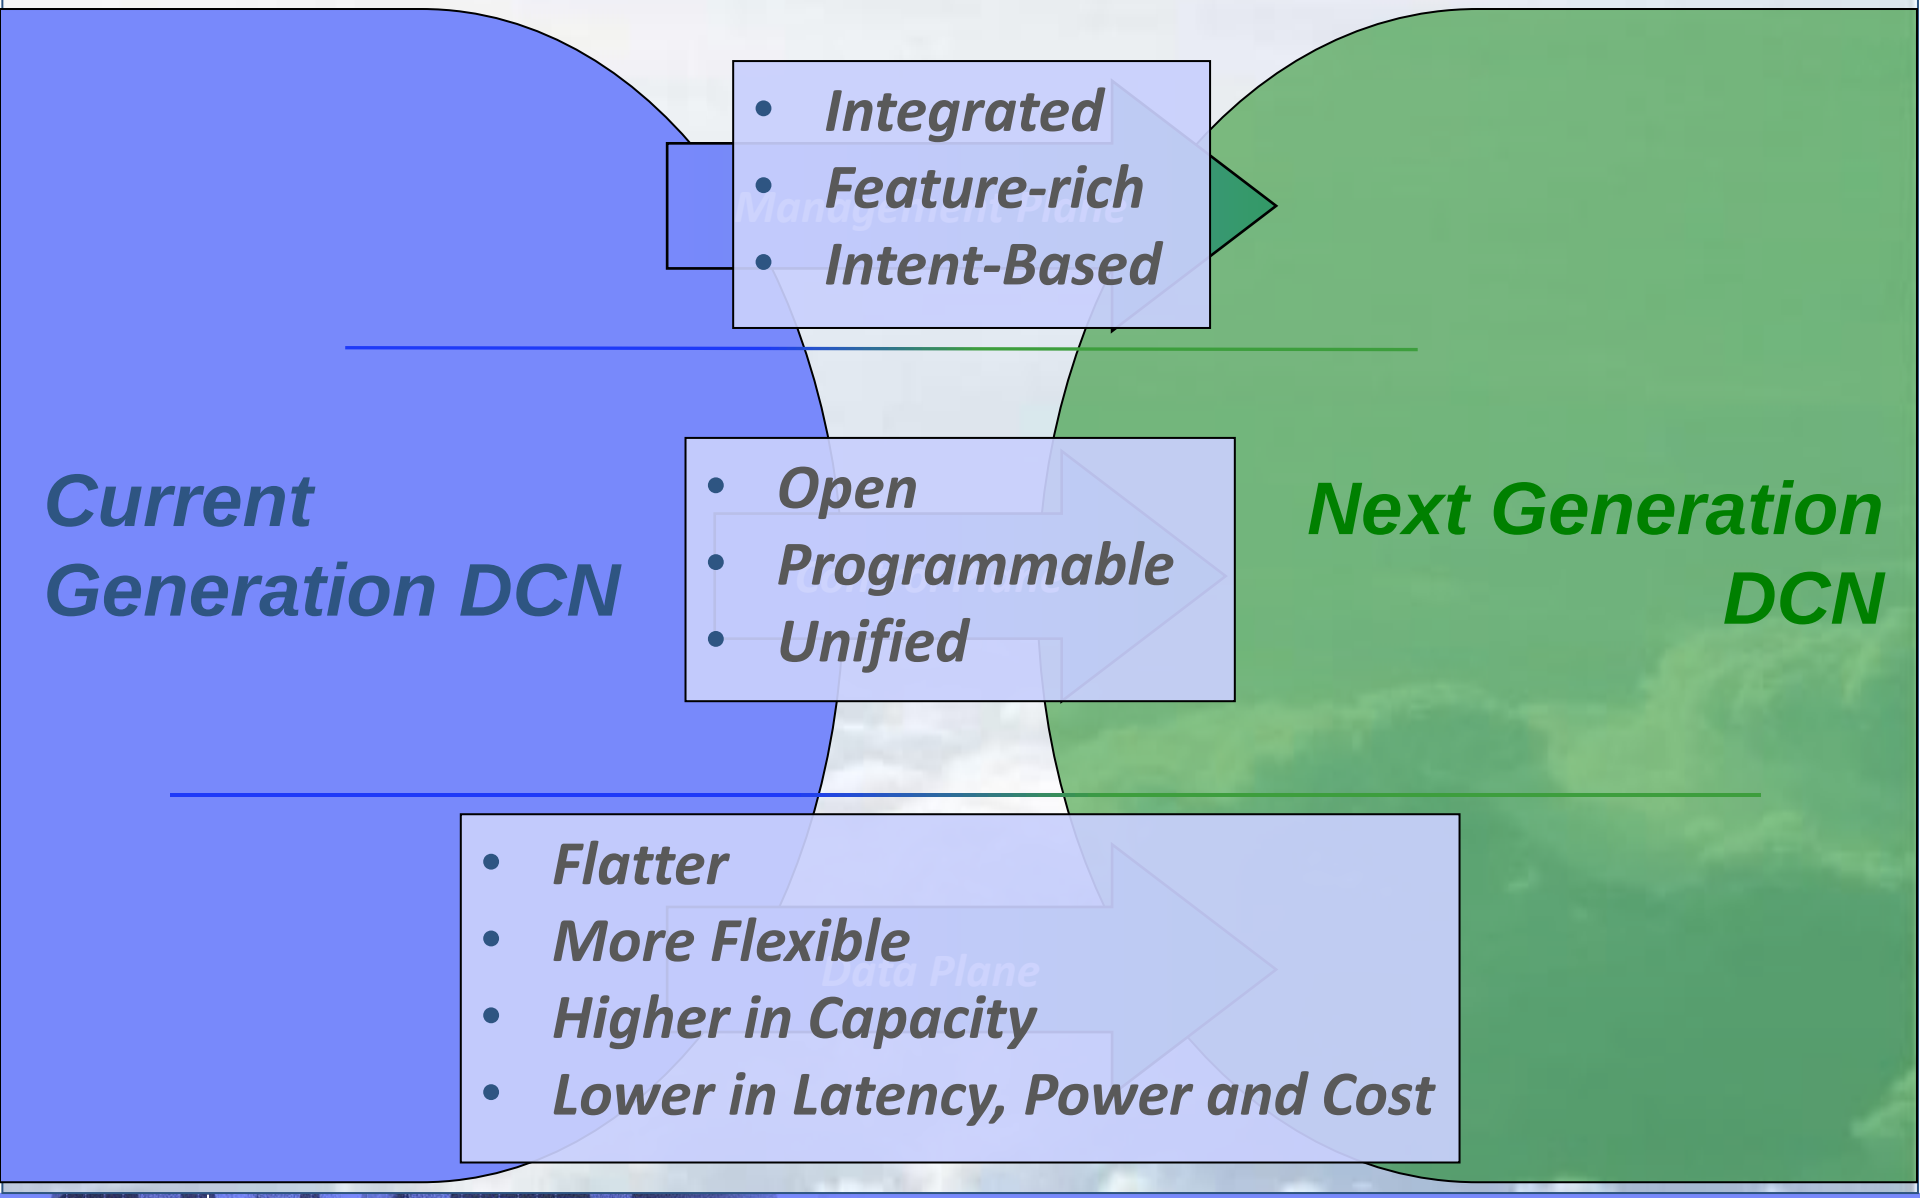

COSIGN – Data Plane Transformation

Current Generation DCN

- Hierarchical
- Rigid
- Over provisioned
- Segregated
- Inefficient for east-west flows
- Restricts workload placement

Data Plane Topology

Next Generation DCN

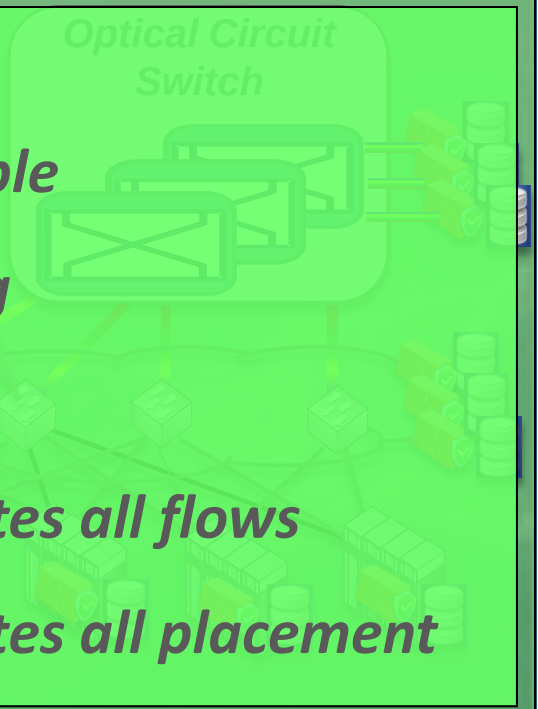

- Flat
- Reconfigurable
- Non-blocking
- Converged
- Accommodates all flows
- Accommodates all placement

COSIGN – Data Plane Transformation

Current Generation DCN

- Many types of interconnects
- Many types of cables
- Occupy lots of space
- Hard to manage

Data Plane Cabling

Next Generation DCN

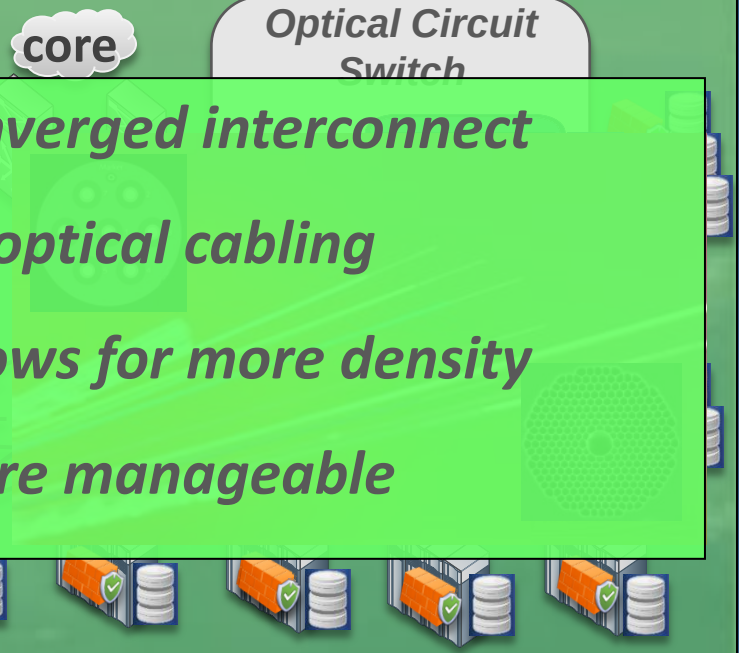

- Converged interconnect
- All optical cabling
- Allows for more density
- More manageable

COSIGN – Control Plane Transformation

Current Generation DCN

- Closed
- Static
- Heterogeneous
- No global visibility
- Slow (HW) innovation cycles

Control Plane

Next Generation DCN

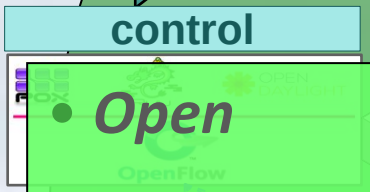

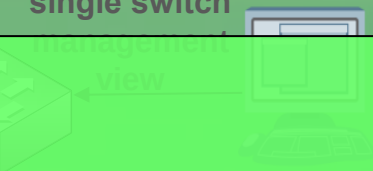

- Open
- Programmable
- Unified
- Logically Centralized
- Fast (SW) innovation cycles

COSIGN – Management Plane Transformation

Current Generation DCN

App level

- Manual configuration steps
- Slow
- Error prone
- Based on overprovisioning

Management Plane

Next Generation DCN

- Automated
- Streamlined
- Optimized
- Allows for right-sizing

DCN for Virtualized Environments

- Same functional requirements but on a far larger scale
  - Large (up to 100K) of physical servers per data center
  - *10 for the Virtual Machines
  - *100 for containers
- More dynamic lifecycle
- Less tolerance for delays and for manual steps

DOVE – Distributed Overlay Virtual Network

Network-as-a-Service abstraction

- *Connectivity under policy*

Centralized control plane (SDN)

- *Scalability*
- *Flexibility*
- *Ease of operations*

Overlay data plane

- *Isolation*
- *Scalability*
- *Mobility*
- *Technology independence*

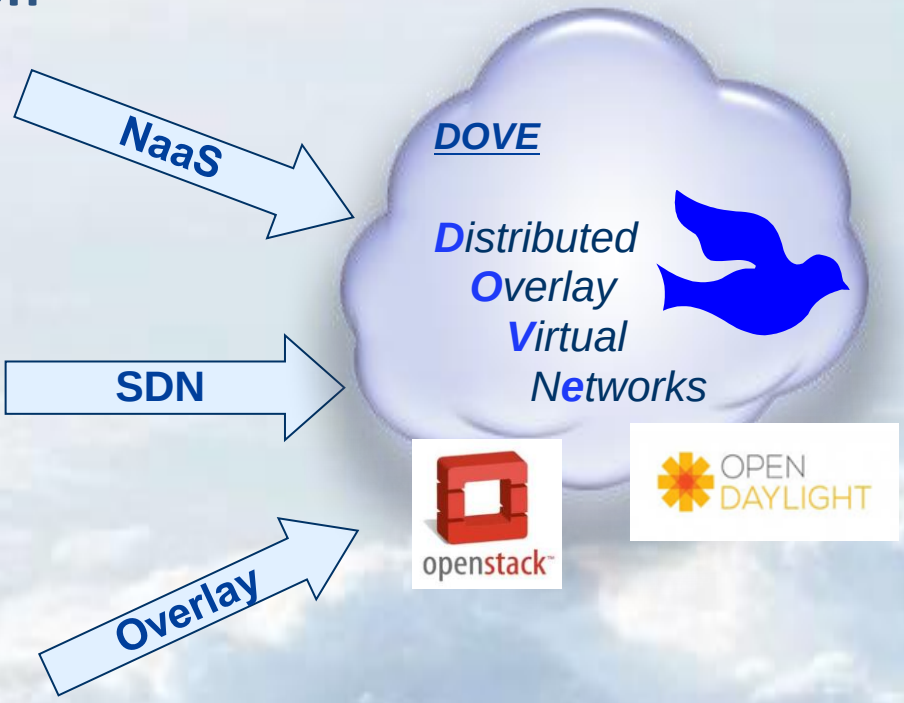

DOVE – Distributed Overlay Virtual Network

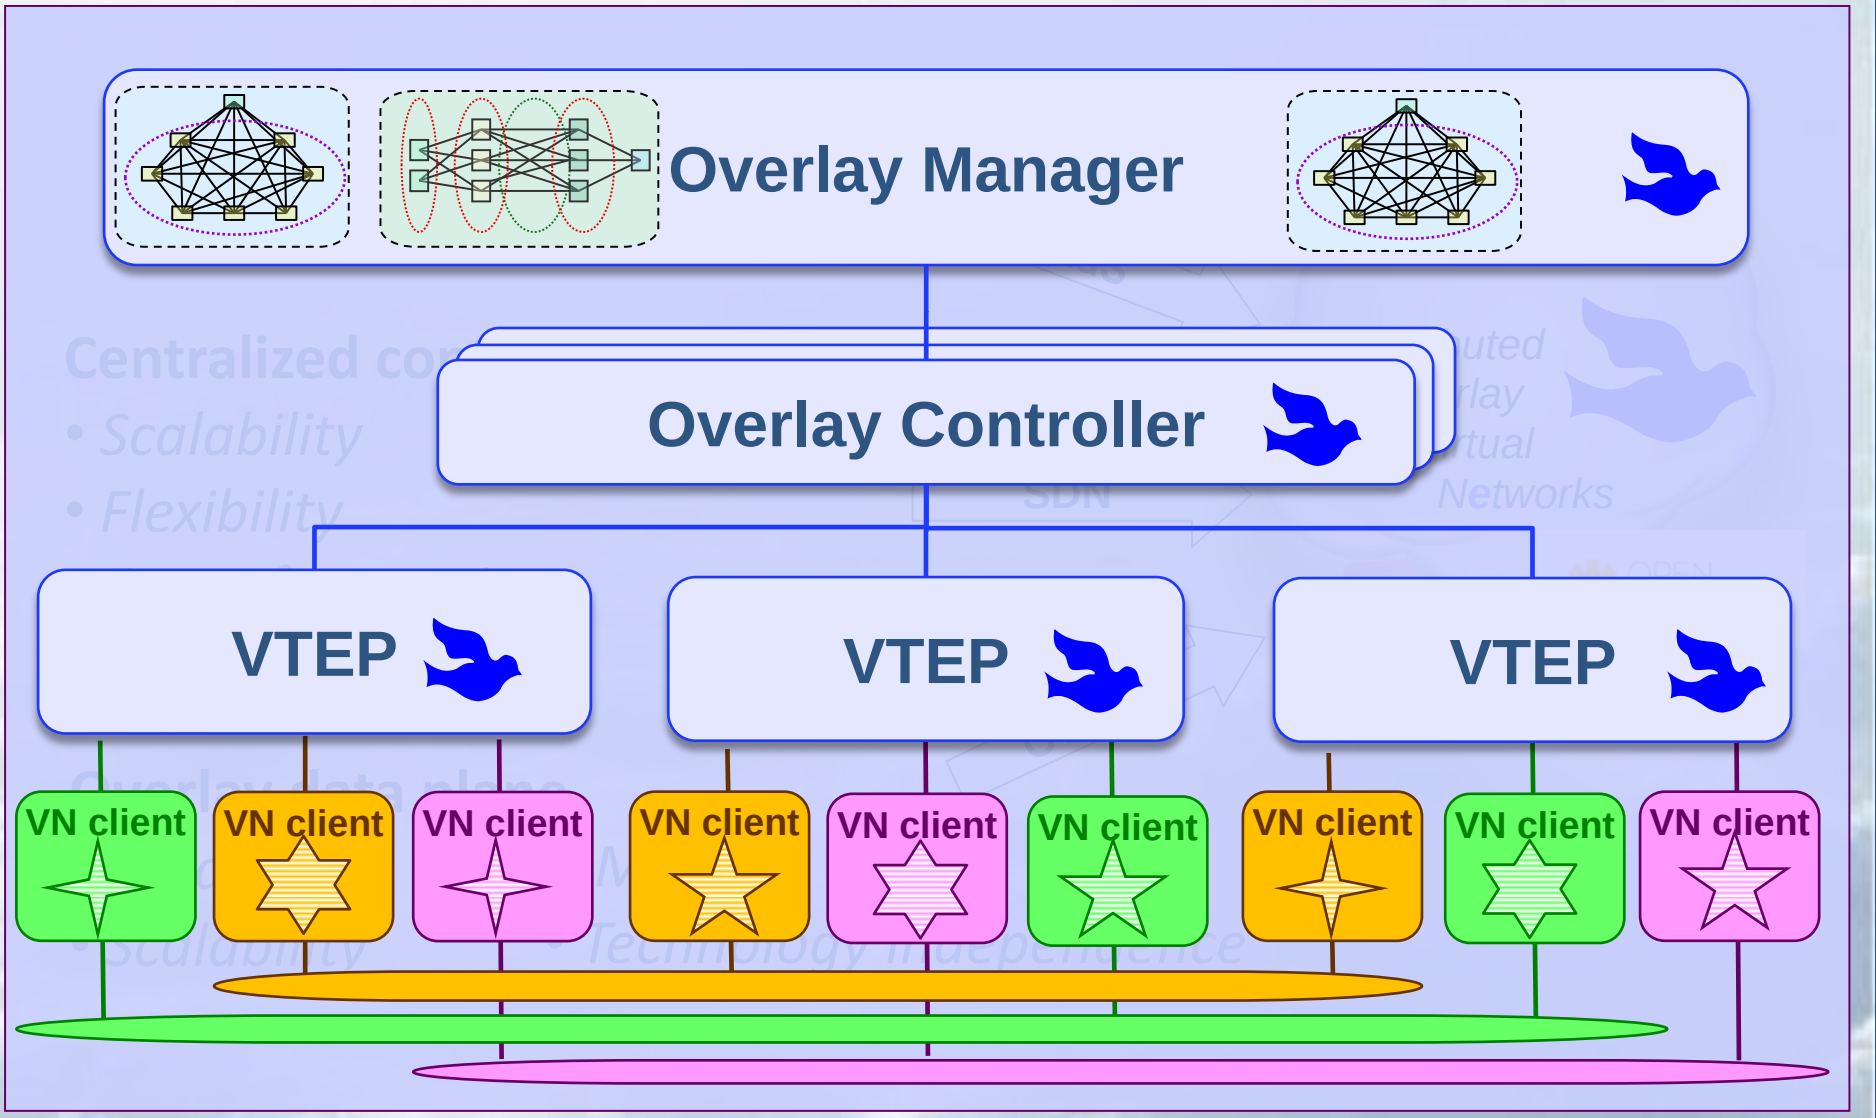

DOVE – Distributed Overlay Virtual Network

Complete isolation from the physical network is a mixed blessing

- **Pros**
  - Can use own addresses
  - VN clients can easily migrate
  - VTEPs enclose the heterogeneity
- **Cons**
  - Can not differentiate between the physical paths
  - Can not interoperate with the physical devices
  - Can not reliably monitor and troubleshoot

COSIGN
DCN Architecture

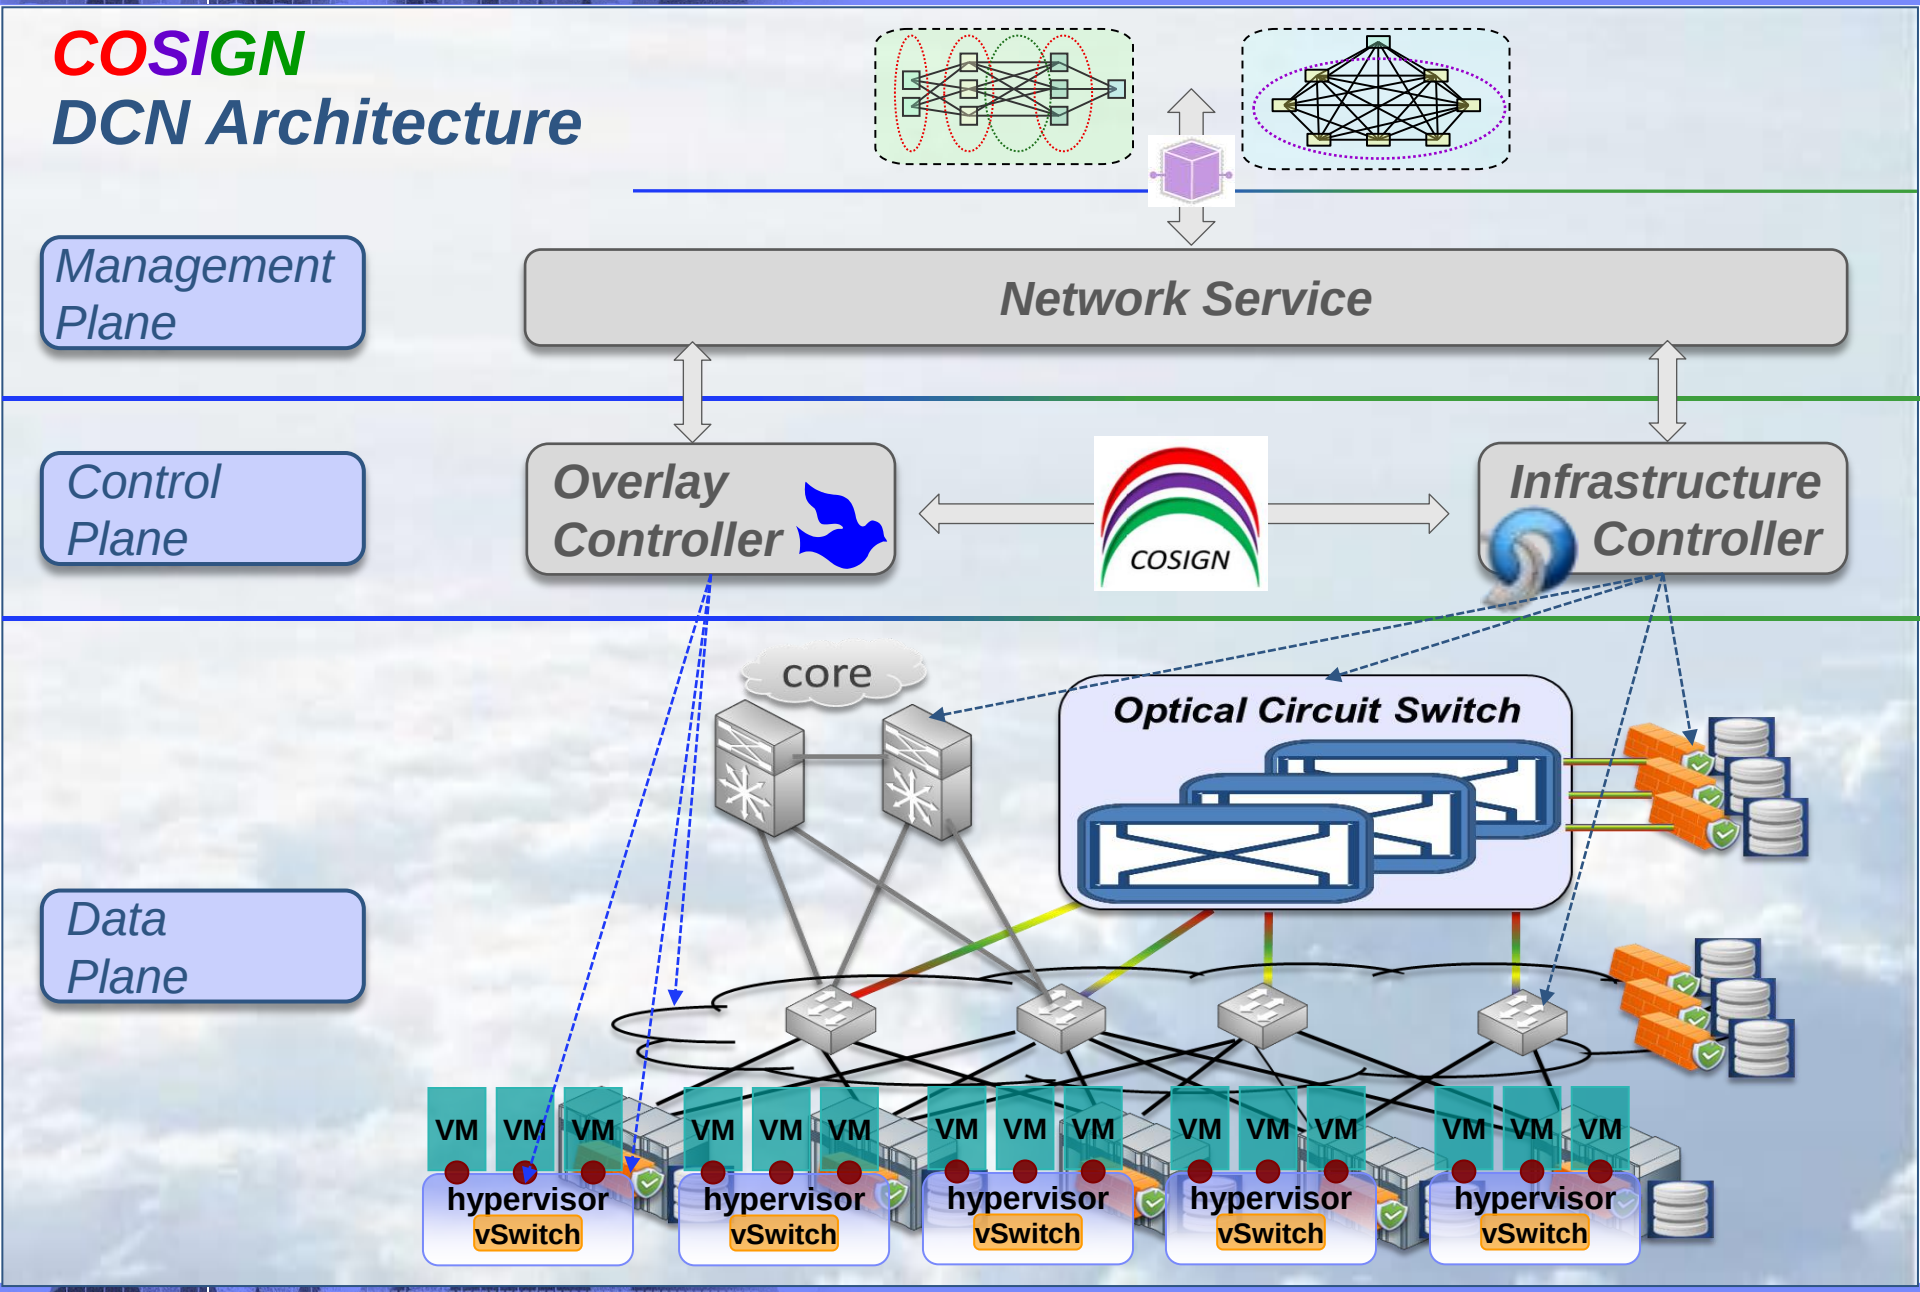

SEVENTH FRAMEWORK PROGRAMME

COSIGN

http://www.fp7-cosign.eu/

DTU

UPC UNIVERSITAT POLITÈCNICA DE CATALUNYA BARCELONATECH

PhotonX Networks

UNIVERSITY OF Southampton Optoelectronics Research Centre

polatis

Venture Photonics Ltd (Venture)

i2cat

NEXTWORKS ENGINEERING FORWARD

TU/e

interoute from the ground to the cloud

University of BRISTOL

ofs A Furukawa Company

IBM

Katherine Barabash at IBM Research