Google Cloud

“Without networking, there is no cloud.”

Cloud IPv6 Innovation

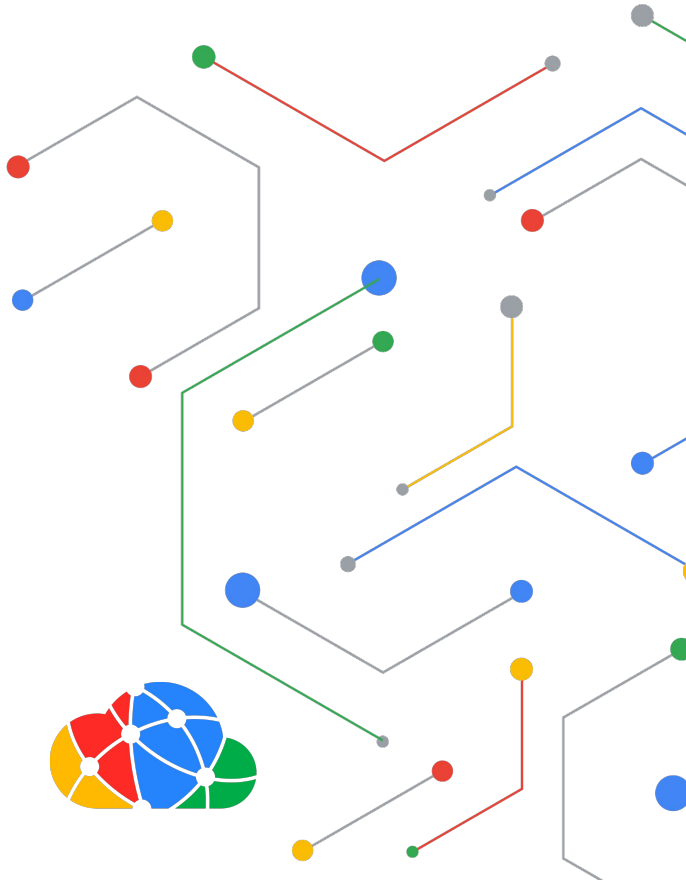

Shaowen Ma, shaowen@google.com
Group Product Manager
July 2022

Agenda

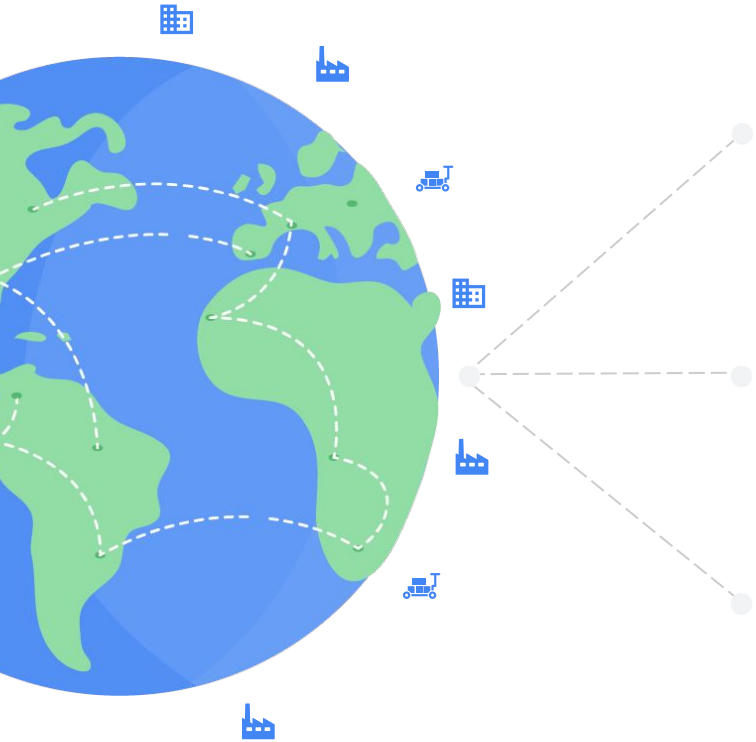

Cloud Migrate to IPv6

Underlay B2/B4/Jupiter/Andromeda

IPv6 ULA and VPC Design

IPv6 Cloud Products

DNS, Interconnect, LB,

IPv6 K8S and Hybrid Cloud

Summary

Many Cloud IPv6 Innovations

Global IPv6 adoption

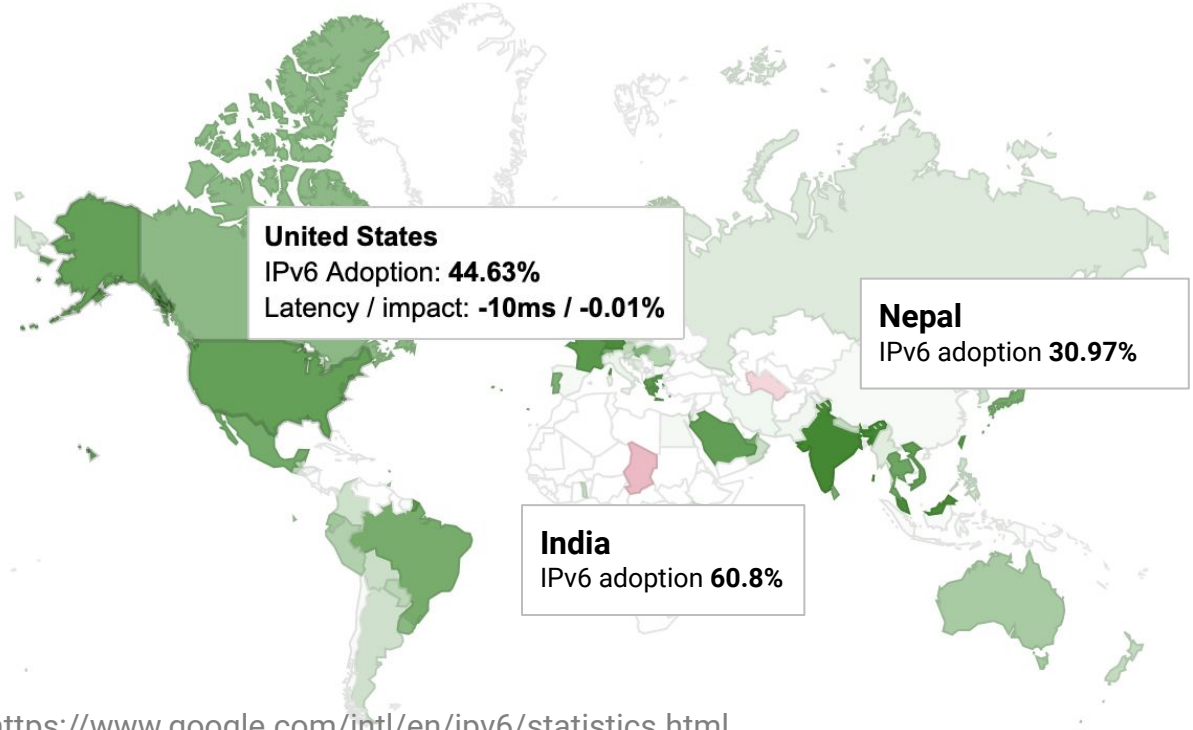

https://www.google.com/intl/en/ipv6/statistics.html

World | Africa | Asia | Europe | Oceania | North America | Central America | Caribbean | South America

RANK	IPV6%	COUNTRY / REGION
1	60.8%	India
2	53.4%	Belgium
3	51%	Germany
4	48.4%	Malaysia
5	46.9%	Switzerland
6	44.3%	France
7	44.1%	Saudi Arabia
8	43.7%	Luxembourg
9	43.2%	Viet Nam
10	43%	Greece
11	41.9%	Japan
12	39.6%	Chinese Taiwan
13	39.3%	Mexico
14	38.9%	United Arab Emirates
15	38.8%	United States
33	20.6%	China

Why IPv6? IoT+5G

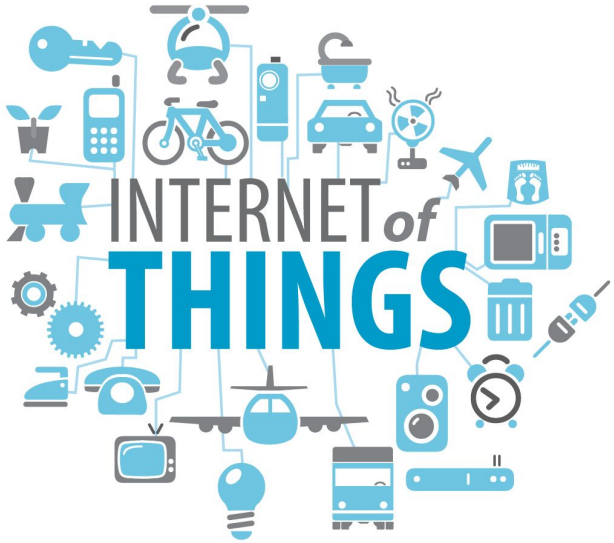

Smart City
Video Surveillance
Smart Lighting
Waste Management
Driver Tracking
ADAS

Connectivity
Stadium / Campus
In-Building
Wind / Solar farms

Healthcare
Remote Monitoring
Virtual Care
Ambulance Health
COVID Response Site

AR/VR
Gaming / Sporting
Search / Ads
Remote training
Connected workforce

IPv4	IPv6
2^32=4.2B(10^9)	2^128=3.4x10^38

4.5 billion

5G | WiFi | Fixed

By 2025, over 4.5 billion IoT connections, of which majority will be 5G

Sources: Omdia, Gartner, Ericsson

Google Cloud

A snapshot for IPv6 Innovation

Google innovations in networking

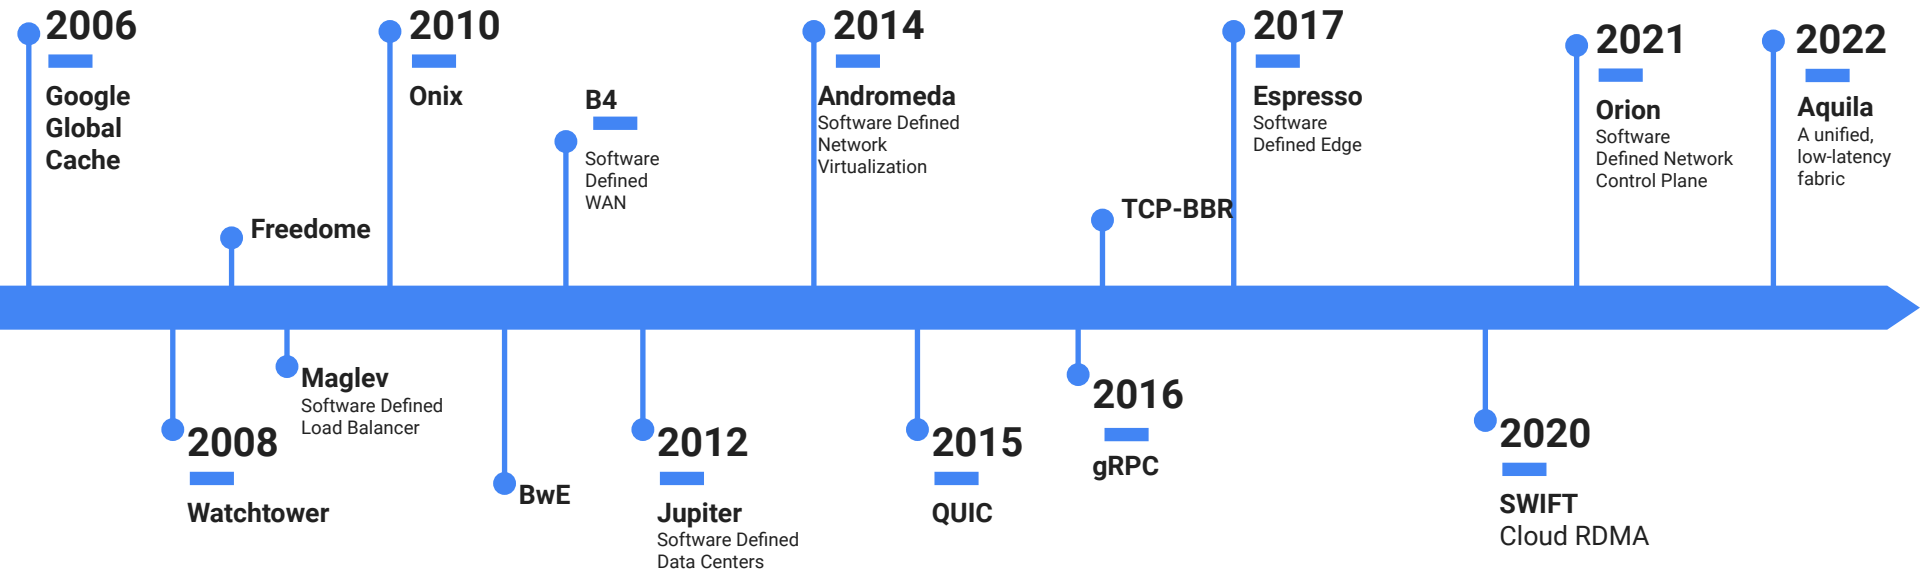

Cloud Global Network IPv6 Transformation

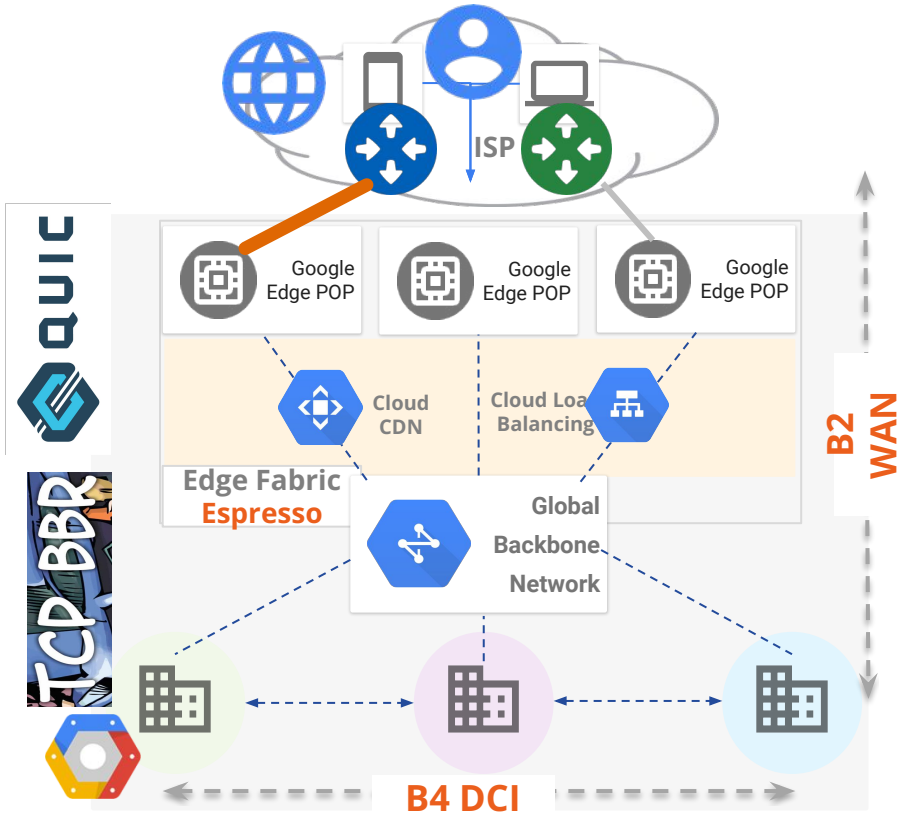

SD-WAN/Interconnect	IPSec BGP IPv6
Edge Pop/LB/CDN/DNS	MPLSoGRE IPv6
Global Backbone, B2/B4	MPLS/SR/ 6to4/Dual Stack
Data Center Cluster	IPv6-only
K8S/Storage/GPU Overlay Services	IPv6 Only, 6to4 IPinIP, RoCEv2
Applications, Search/Gmail/Youtube etc	IPv4/IPv6

RoCE nvm EXPRESS GPU Direct Cloud RDMA TensorFlow

Key IPv6 Benefit: No NAT, no Performance downgrade

NAT44
Common
VPC design

NAT64/46
IPv6 to IPv4 translation

NAT66
No NAT for IPv6 VPC

- No more NAT (Network Address Translation)
- Auto-configuration
- No more private address collisions
- Better multicast routing
- Simpler header format
- Simplified, more efficient routing
- True quality of service (QoS), also called "flow labeling"
- Built-in authentication and privacy support
- Flexible options and extensions

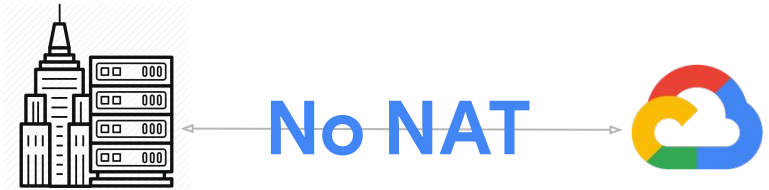

Cloud IPv6 Network Products

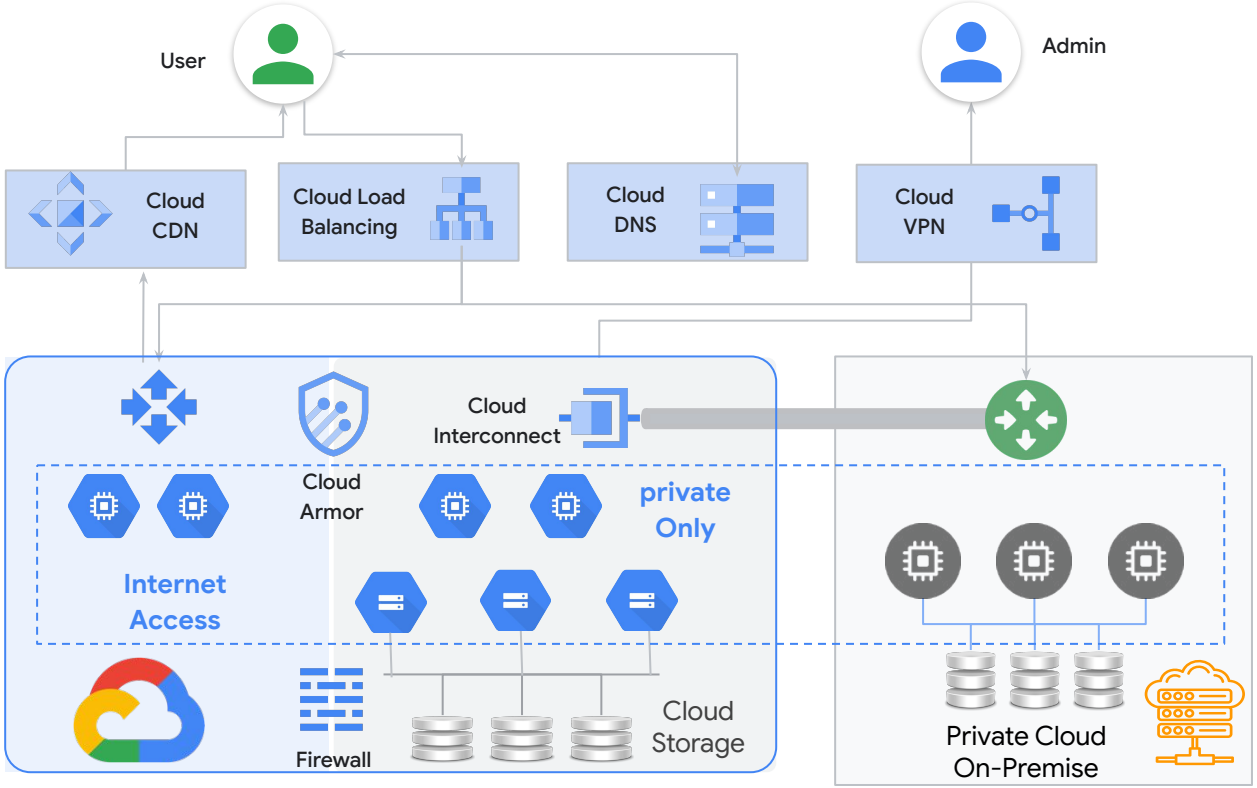

IPv6 VPC Still need Private Address(ULA)!

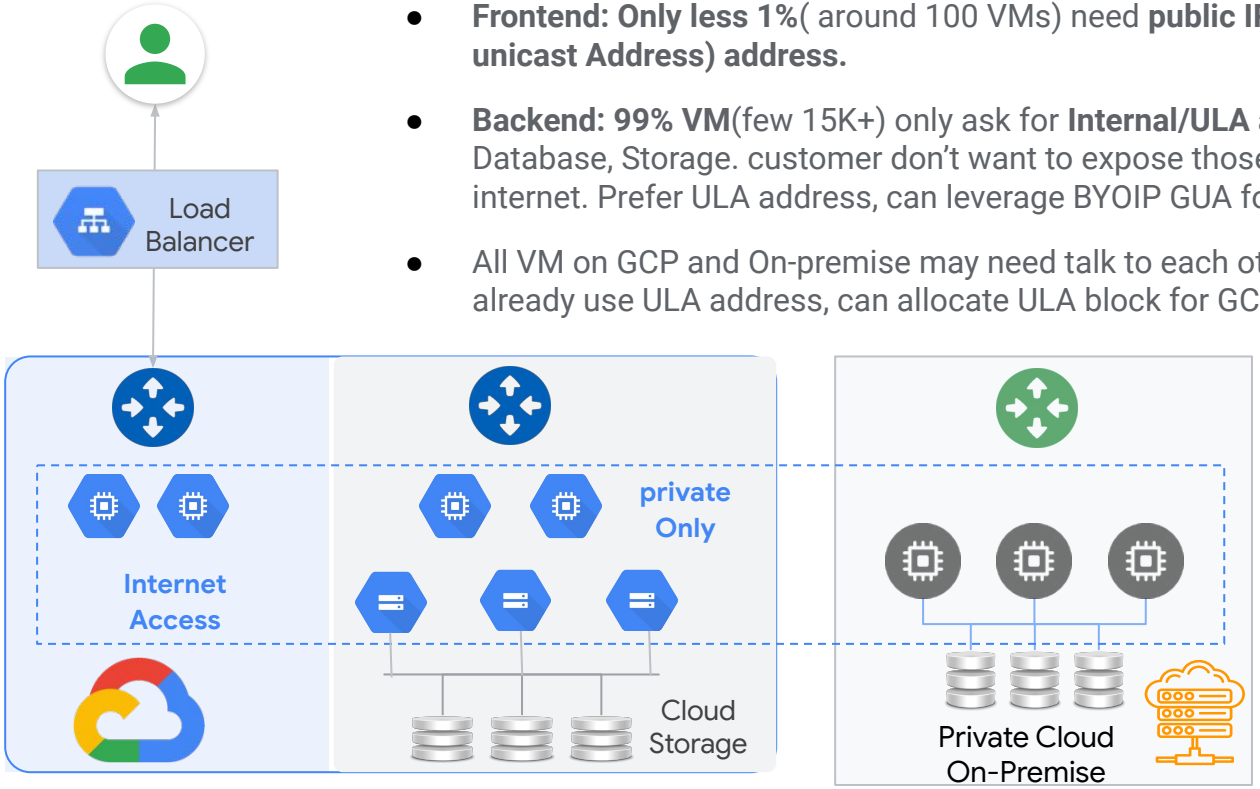

- **Frontend: Only less 1%**(around 100 VMs) need **public IPv6 GUA(global unicast Address) address.**
- **Backend: 99% VM**(few 15K+) only ask for **Internal/ULA address**, include Database, Storage. customer don't want to expose those application to internet. Prefer ULA address, can leverage BYOIP GUA for now.
- All VM on GCP and On-premise may need talk to each others. Private cloud already use ULA address, can allocate ULA block for GCP

IPv6 Addressing, GUA/ULA
Global unicast Address/Unique Local Addressing

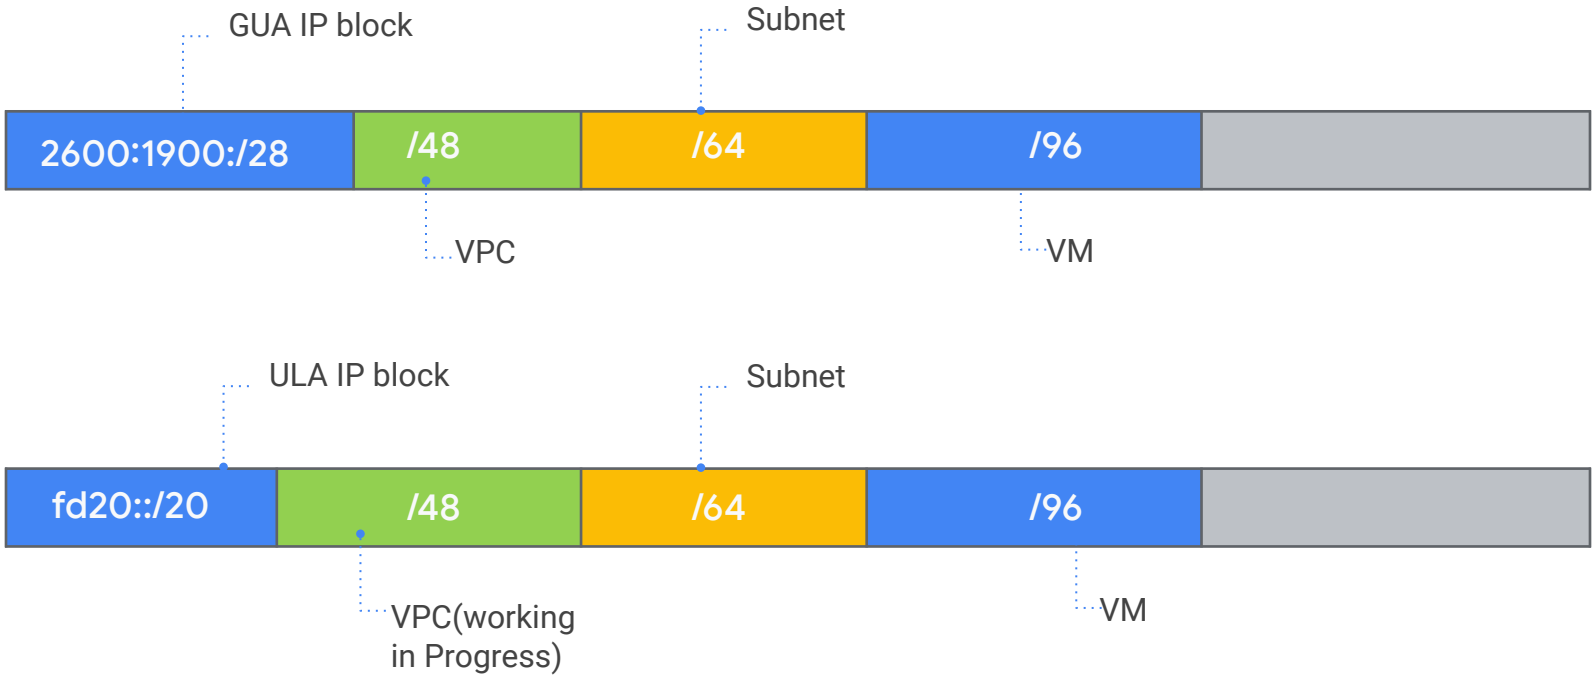

IPv6 NAT66? (Other Cloud Provider Solution)

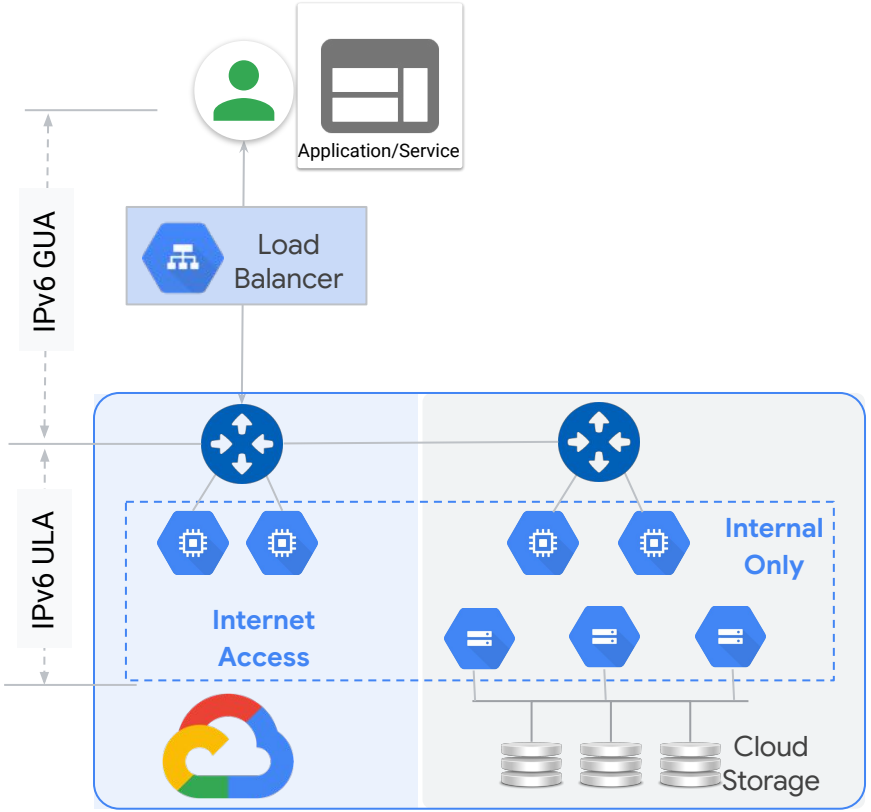

- All VM only have single internal IP address(possible ULA)
- Ingress: User to Cloud traffic use Anycast LB, already **address translation(similar with NAT)**
- Egress: VM use ULA as Src IP, and Andromeda replace Src address with GUA, 1:1 NAT, **GUA not visible to VM.**
- Easier to connect to On-Premise Private Cloud ULA via Private BGP peering.
- IPv4 public cloud providers first use Public IPv4, then they facing security and shortage of IPv4, later all cloud provider changed to only allocate RFC1918 to VM, and NAT on vRouter.
- Only provide ULA on VM, **NAT on virtual Router**

Dual NIC solution with VPN and Internet Access

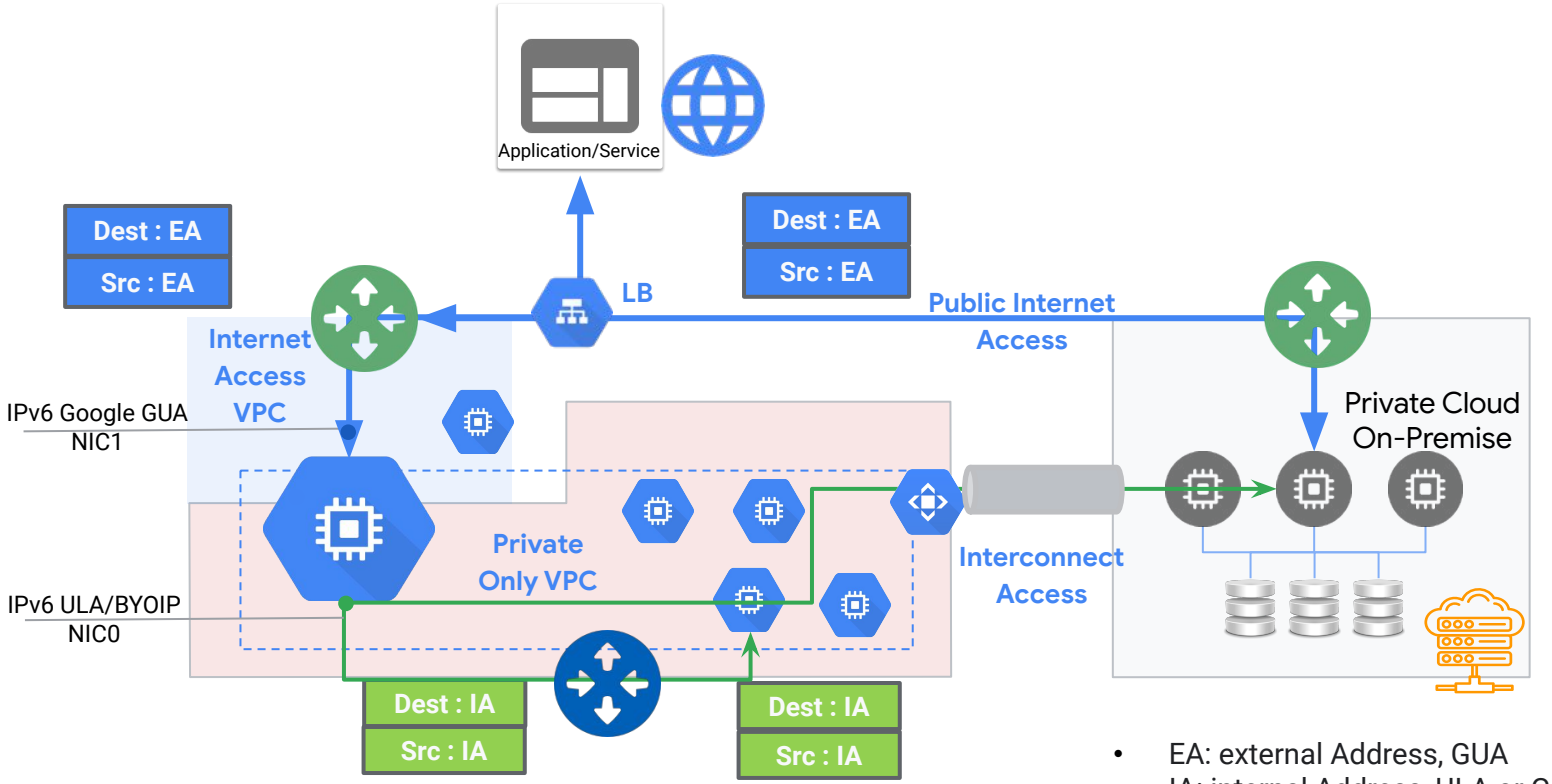

- EA: external Address, GUA
- IA: internal Address, ULA or GUA

Global LB for IPv6 and IPv4 clients

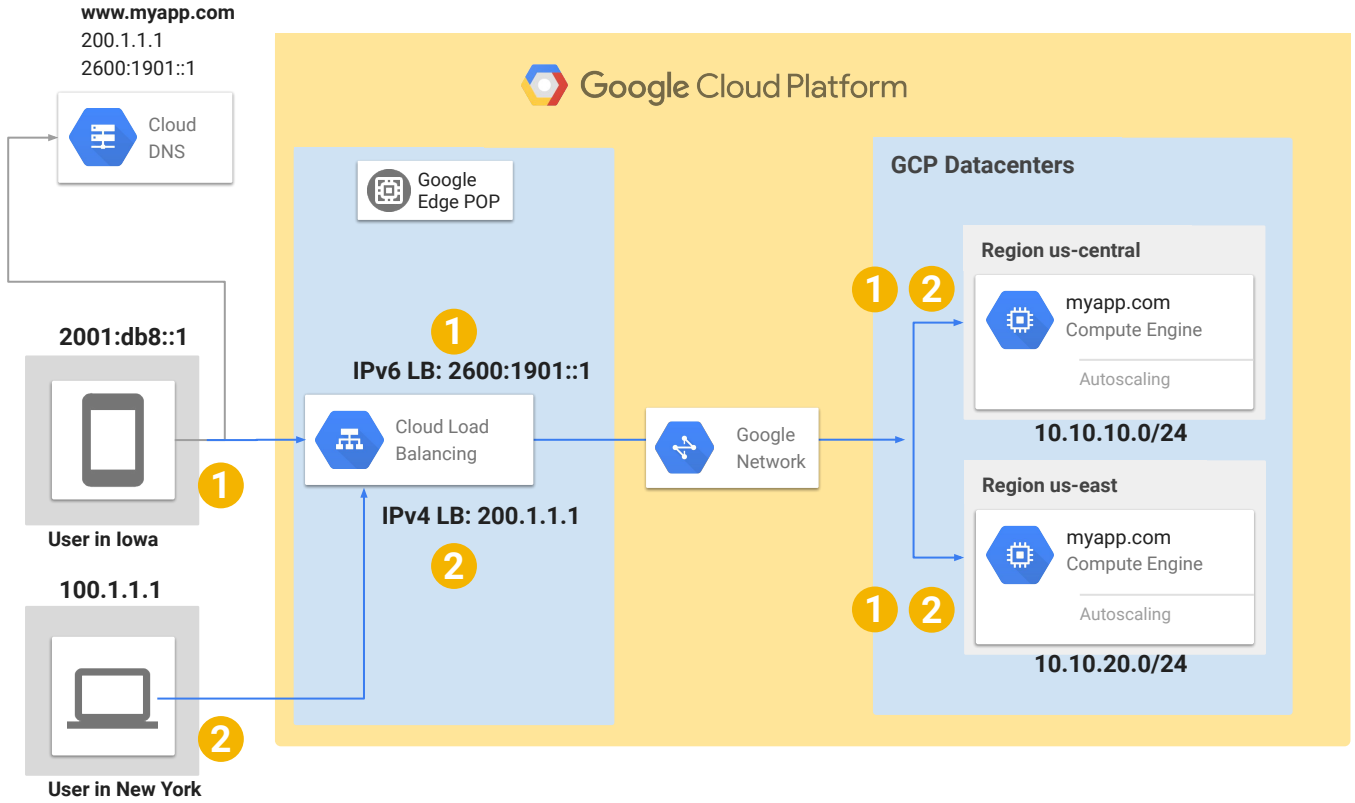

IPv6 Forwarding rules(FRs) i.e. VIPs are free, customers only pay for IPv4 FRs. Rest of the pricing is same.

1 IPv6 clients path

2 IPv4 clients path

Cloud GW IPv6 Support

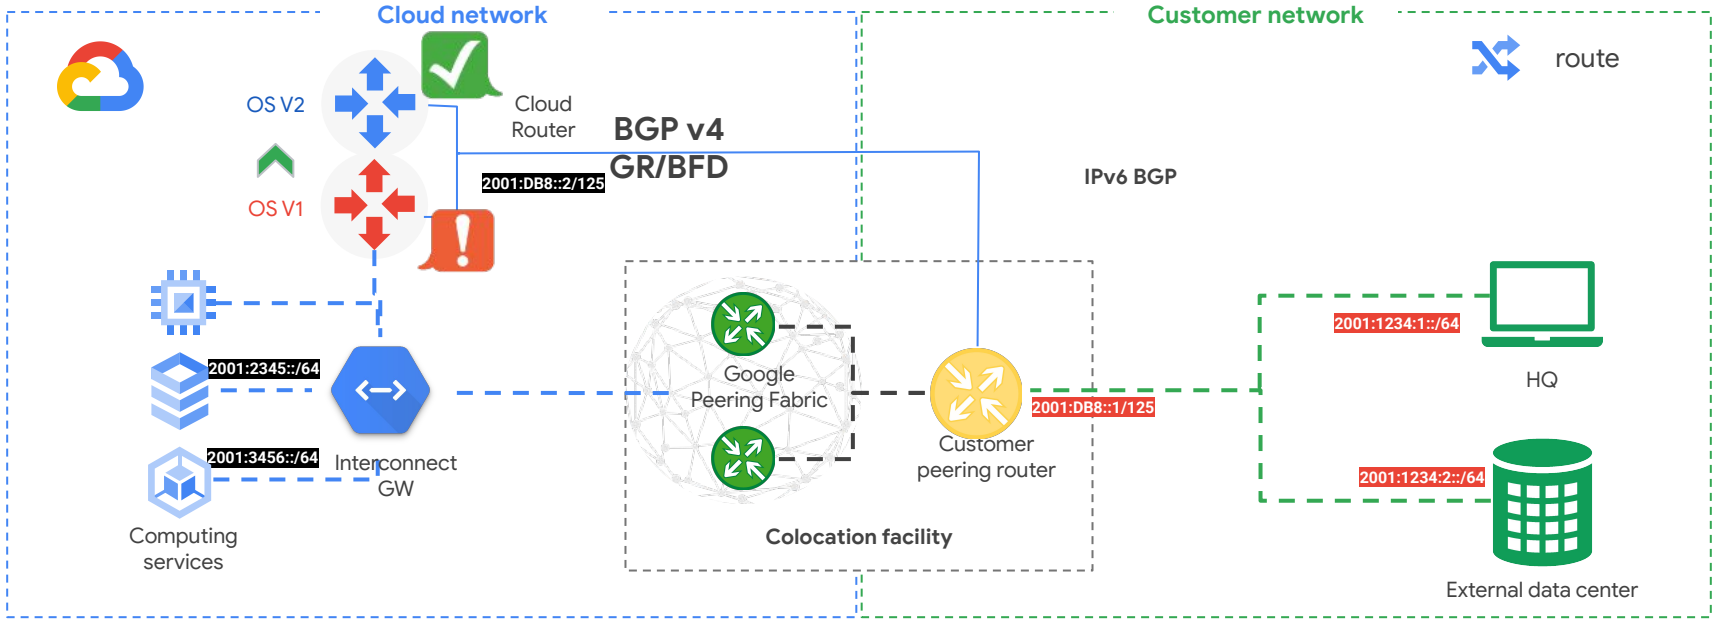

Now support IPv4 BGP session with IPv6 AF.
Later can support IPv6 BGP session with IPv6 AF

Scalable & Dynamic

Scalable IP addressing

IPv6 support and flexible IP address management through multi-Pod CIDR allows you to run more apps with less IPs.

Advanced Traffic Management

Scale from 0 to N with traffic-based autoscaling. Distribute global traffic based on Service capacity and health, allowing you to serve traffic

15,000 Node Clusters

GKE supports the highest scale clusters through networking innovations such as EndpointSlices, container-native load balancing, and scalable IP addressing.

Global Service Discovery

GKE supports the highest scale clusters through networking innovations such as EndpointSlices, container-native load balancing, and scalable IP addressing.

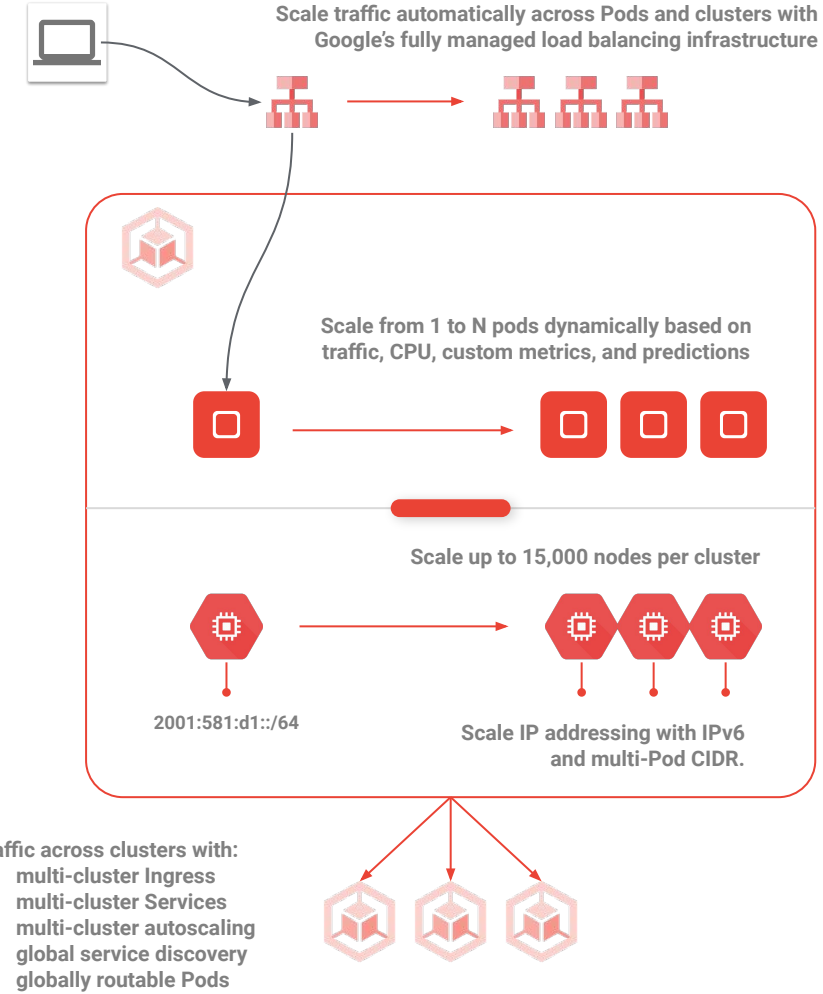

GKE Private Addresses

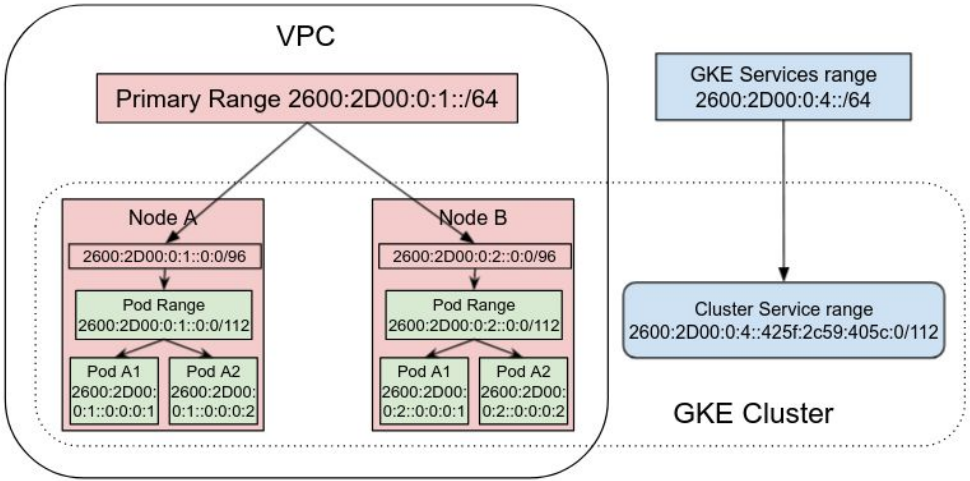

IPv6 address allocation for GKE clusters

```
gcloud compute networks subnets update SUBNET_NAME \
  --stack-type=IPv4_IPv6 \
  --ipv6-access-type=EXTERNAL
```

```
gcloud compute networks subnets create SUBNET_NAME \
  --network=NETWORK \
  --range=IPV4_RANGE \
  --stack-type=IPV4_IPV6 \
  --ipv6-access-type=INTERNAL \
  --region=REGION
```

Dual Stack (IPv4/IPv6)


```
# Valid IP family policies:
# - SingleStack
# - PreferDualStack
# - RequireDualStack
ipFamilyPolicy: PreferDualStack
ipFamilies:
- IPv6
- IPv4
clusterIP: 198.51.100.2
clusterIPs:
- 192.51.100.2
- 2001:D88:1::2
```


```
podIP: 192.0.2.2
podIPs:
- ip: 192.0.2.2
- ip: 2001:D88:2::2
```


```
podCIDR: 192.0.2.0/24
podCIDRs:
- ip: 192.0.2.2/24
- ip: 2001:D88:2::2/36
```

Google Cloud

Dual Stack IPv4/IPv6 Hybrid Cloud

1. Flat Network - Service Mesh , Multi-Cluster Connectivity without NAT or Proxy
2. Migration from on-premises IPv6 deployment to GCP
3. IPv6 standardization - Telcos and SPs

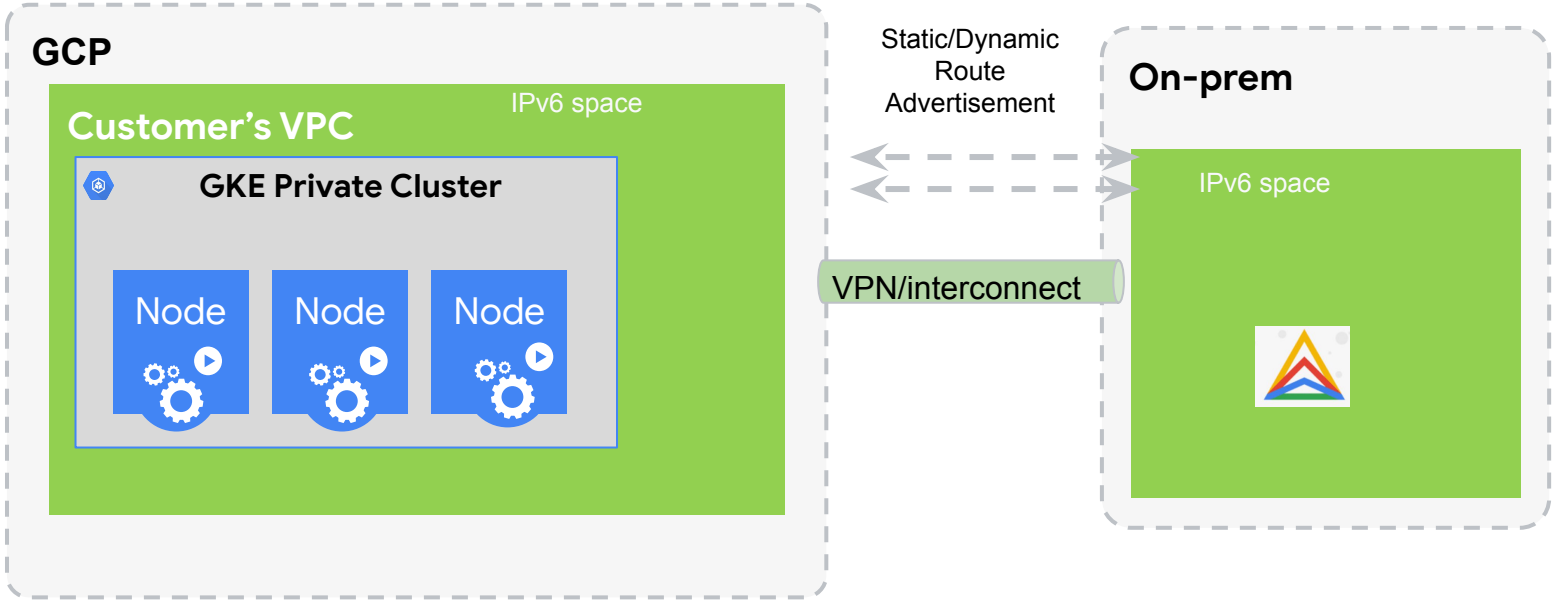

IPv6 Enable GDC for Applications and Network Functions

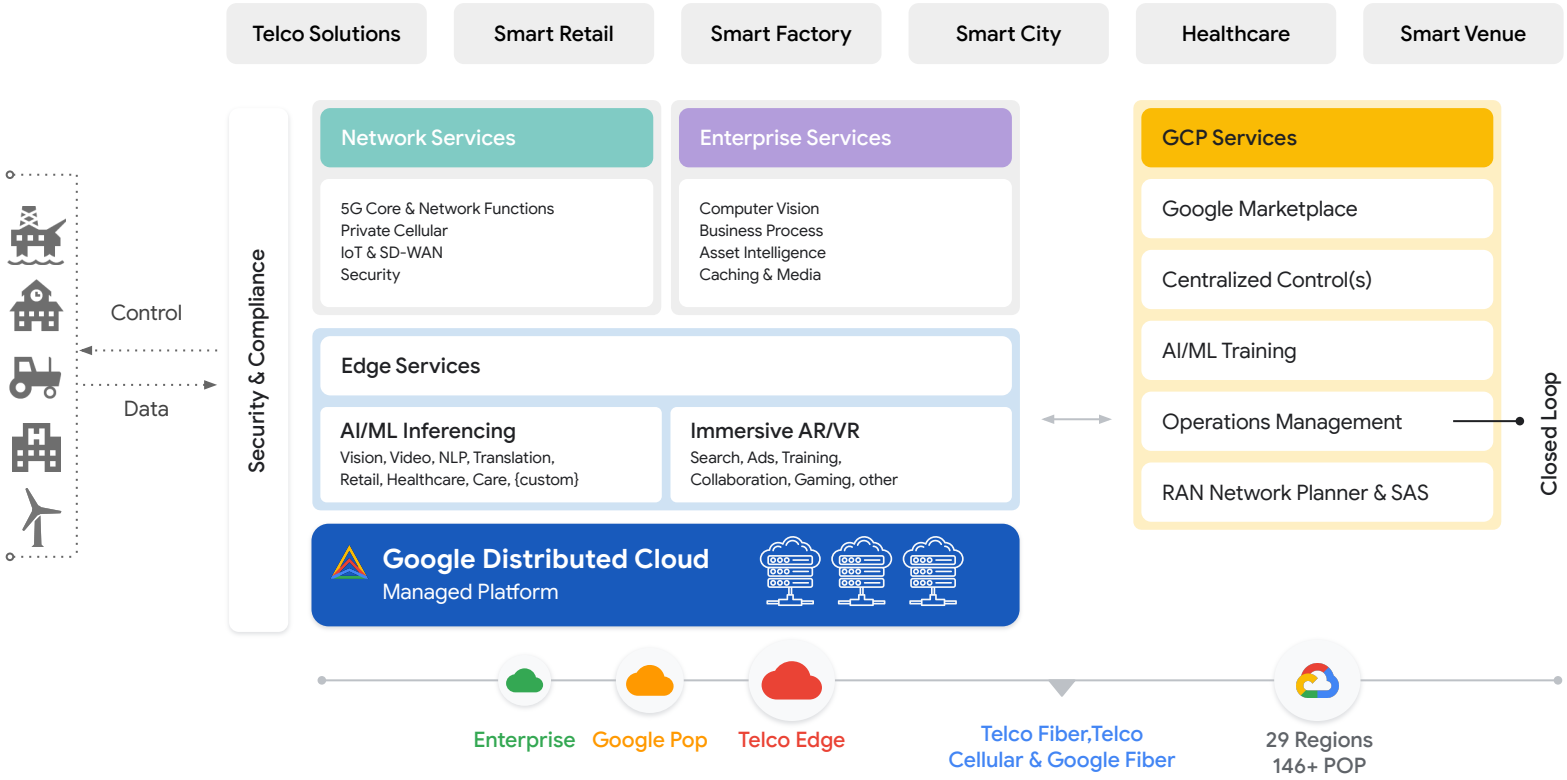

Google

Thank you