

IPv6&SRv6 Powering Alibaba Cloud AI Computing

2026.03

Network Challenges in AI Computing

Challenge 1

IPv4 Resource Provisioning

Growing model scale and prefill-decode separation are driving AI traffic from intra-cluster to cross-cluster, requiring unique address allocation between clusters instead of reuse.

10 Alibaba Cloud HPN7.0 cluster interconnections require 380 class B IP blocks.

Challenge 2

Virtualization Performance Loss

VXLAN tenant isolation introduces 50-byte packet overhead and encapsulation latency. RDMA-based AI computing demands minimal virtualization loss.

VXLAN introduces 1 μ s latency and 1% throughput loss (1K GPUs training).

Challenge 3

Cross-region bandwidth Contention

ECMP routing causes cross-region AI traffic to compete with Others traffic for backbone queue resources, resulting in packet loss.

DSCP value mapping to hardware queues cannot distinguish between RDMA and TCP traffic.

IPv6/SRv6 Solve Networking Challenges

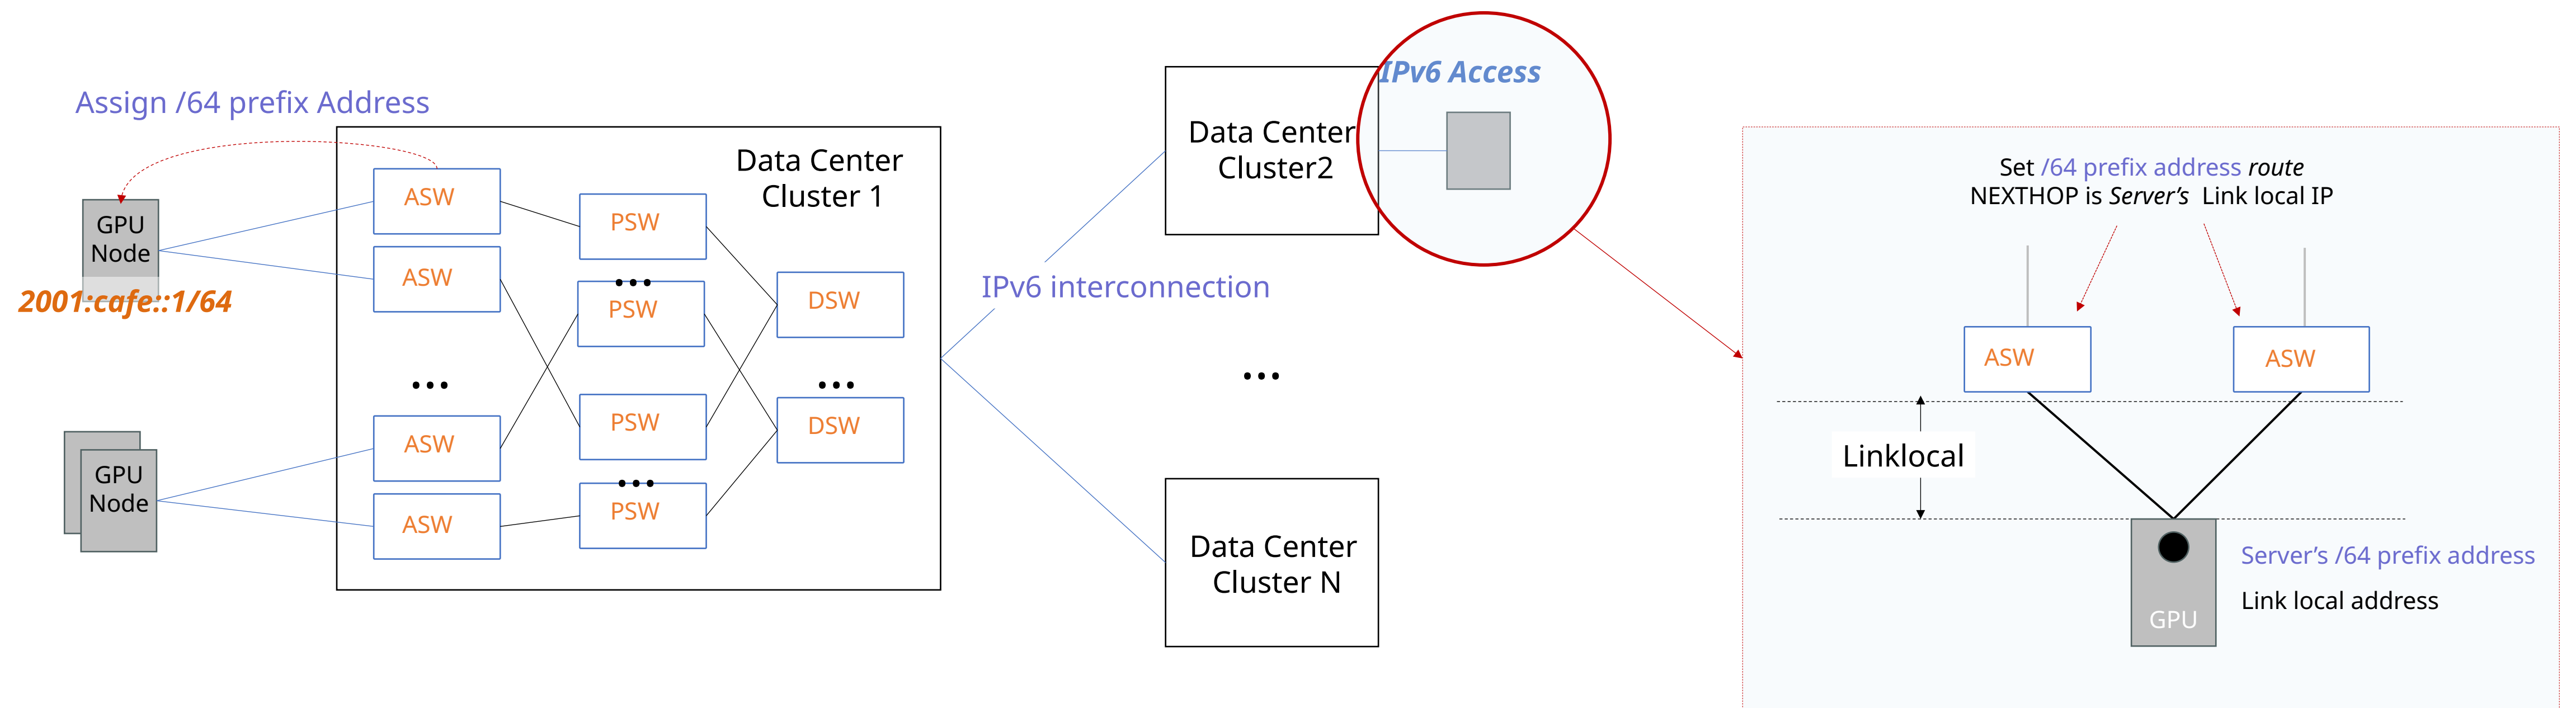
IPv6 Native Data Center

IPv6 Native Cluster

- Provide IPv6-only server access and IPv6 connectivity for servers cross-cluster.
- Assign /64 prefix address to server NICs by default, reserved IPv6 Address for Programming Capability

Layer 3 Mode Access

- Deploy link-local addresses between switches and servers
- Switch route to server's /64 Prefix address through server's link-local IP
- Address the challenge of switches processing ND messages for all IPs in /64 prefix

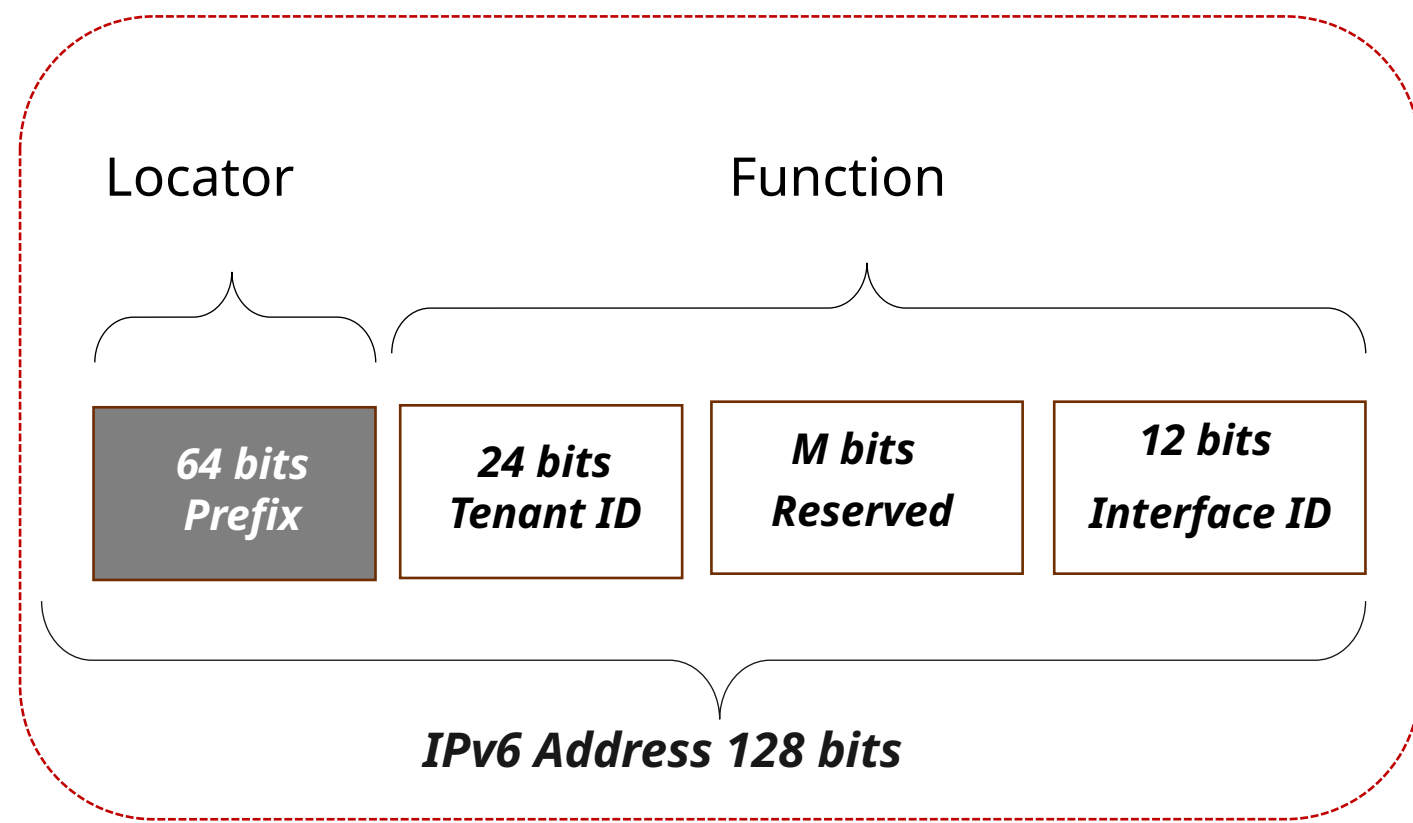


IPv6/SRv6 Solve Networking Challenges

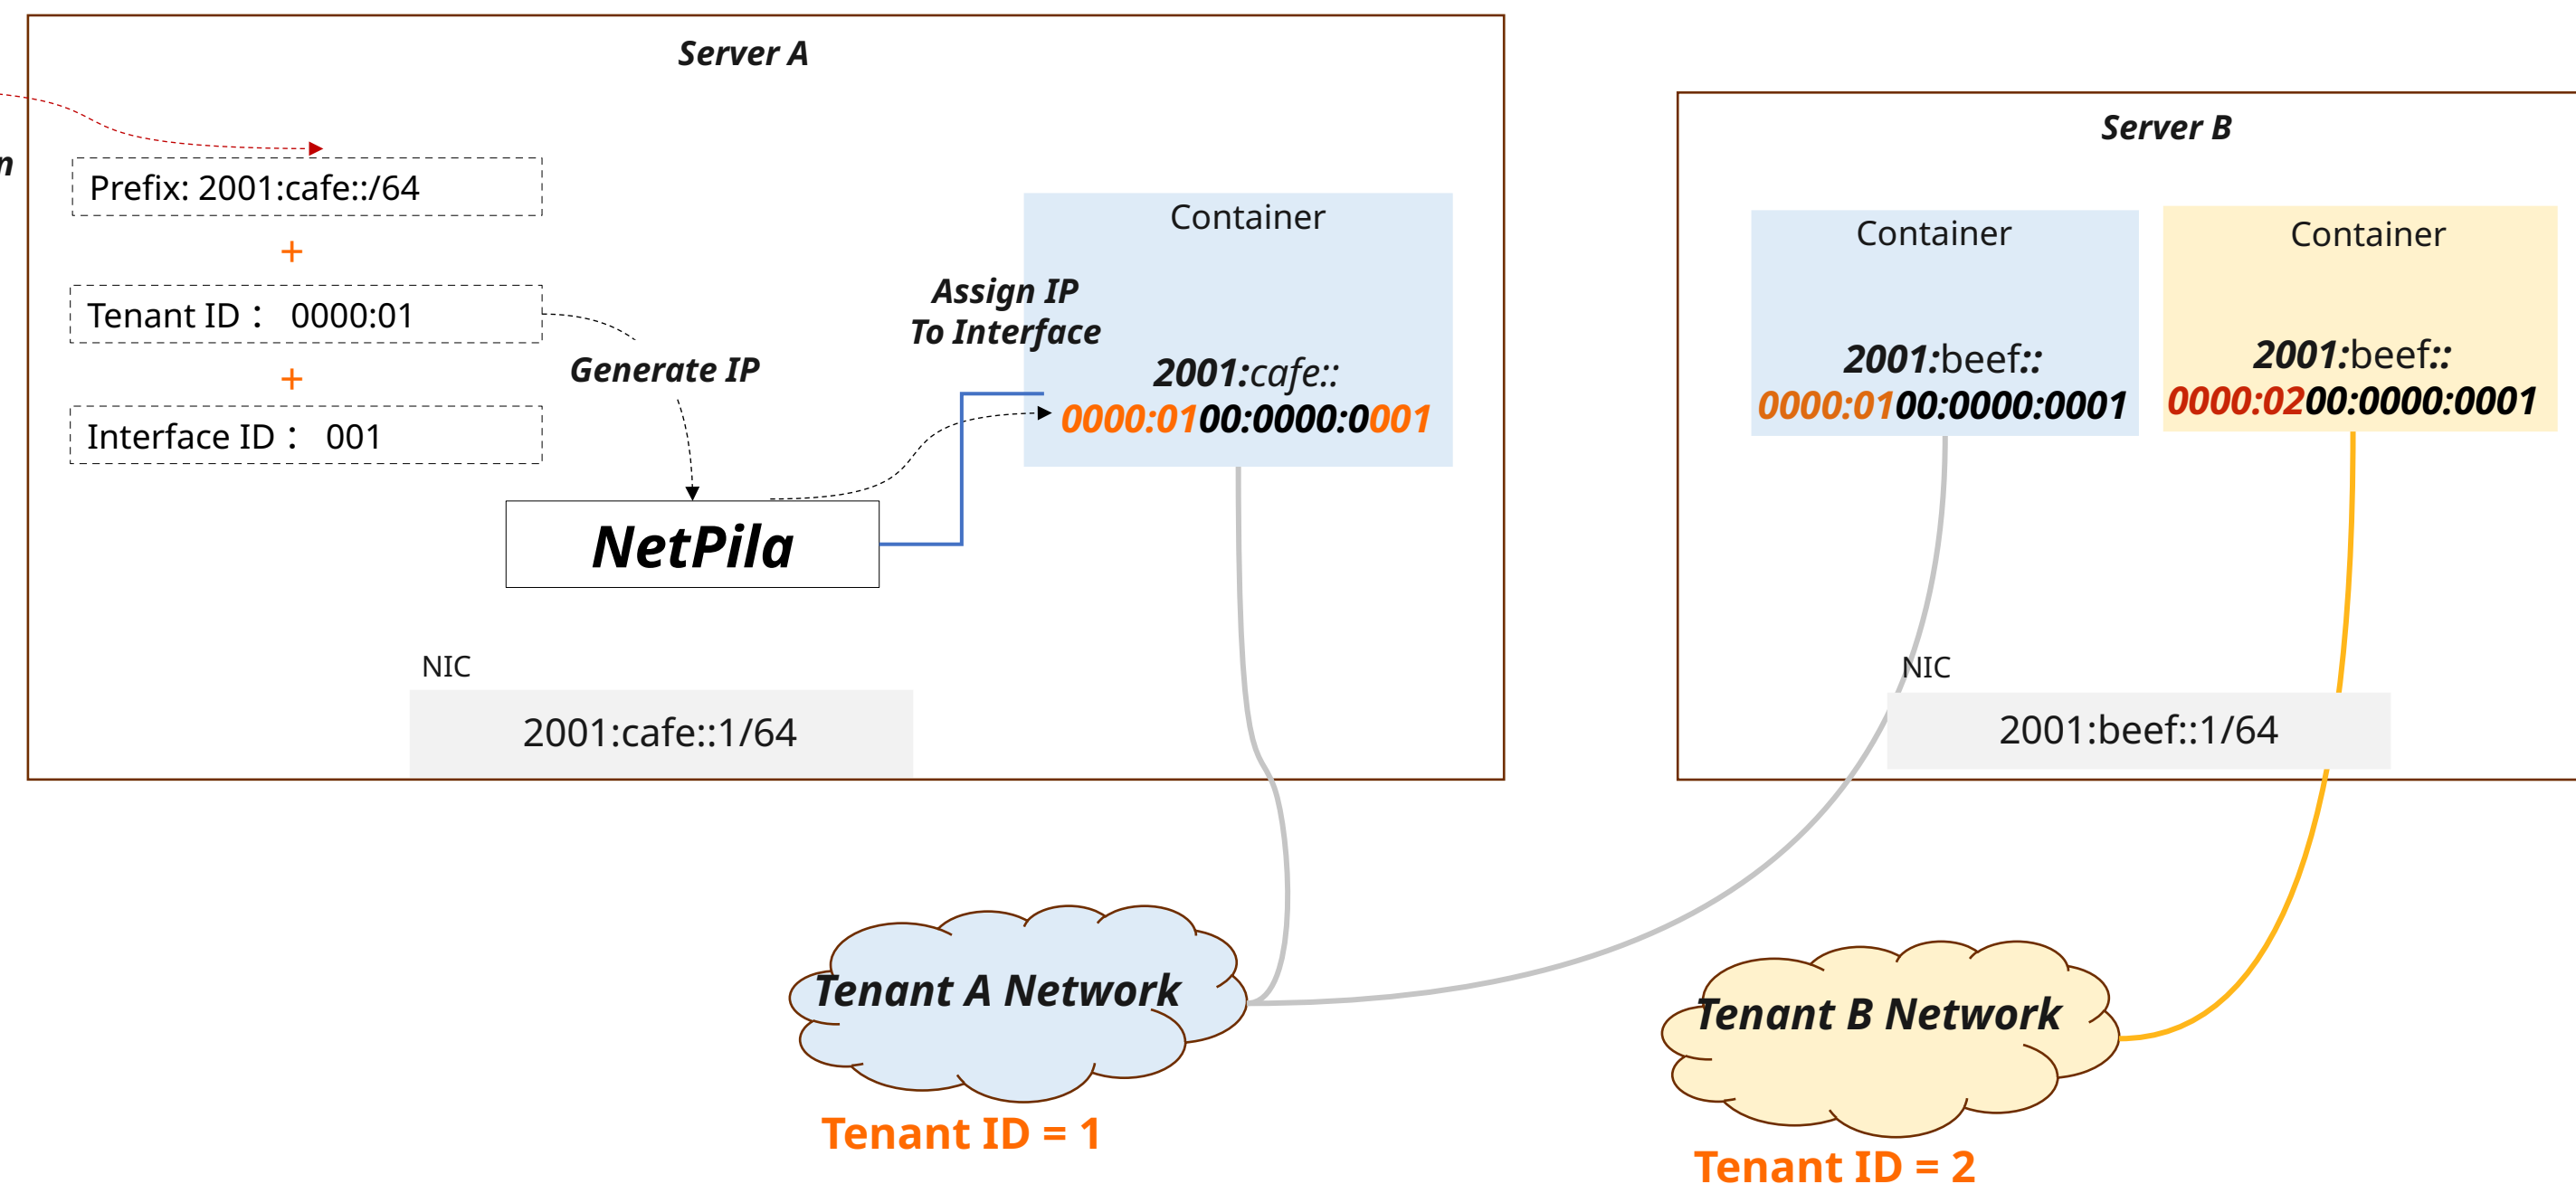
SRv6 Container Network- Netpila

Interface IP Generation Rule Definition

- Similar to End.X SID, defined new fields: Tenant ID, Interface ID, and Reserved field
- Tenant ID: 24-bit default, aligned with VXLAN ID bit length
- Interface IP = Locator (GUA /64 prefix) + Tenant ID + Interface ID



Similar to End.X SID

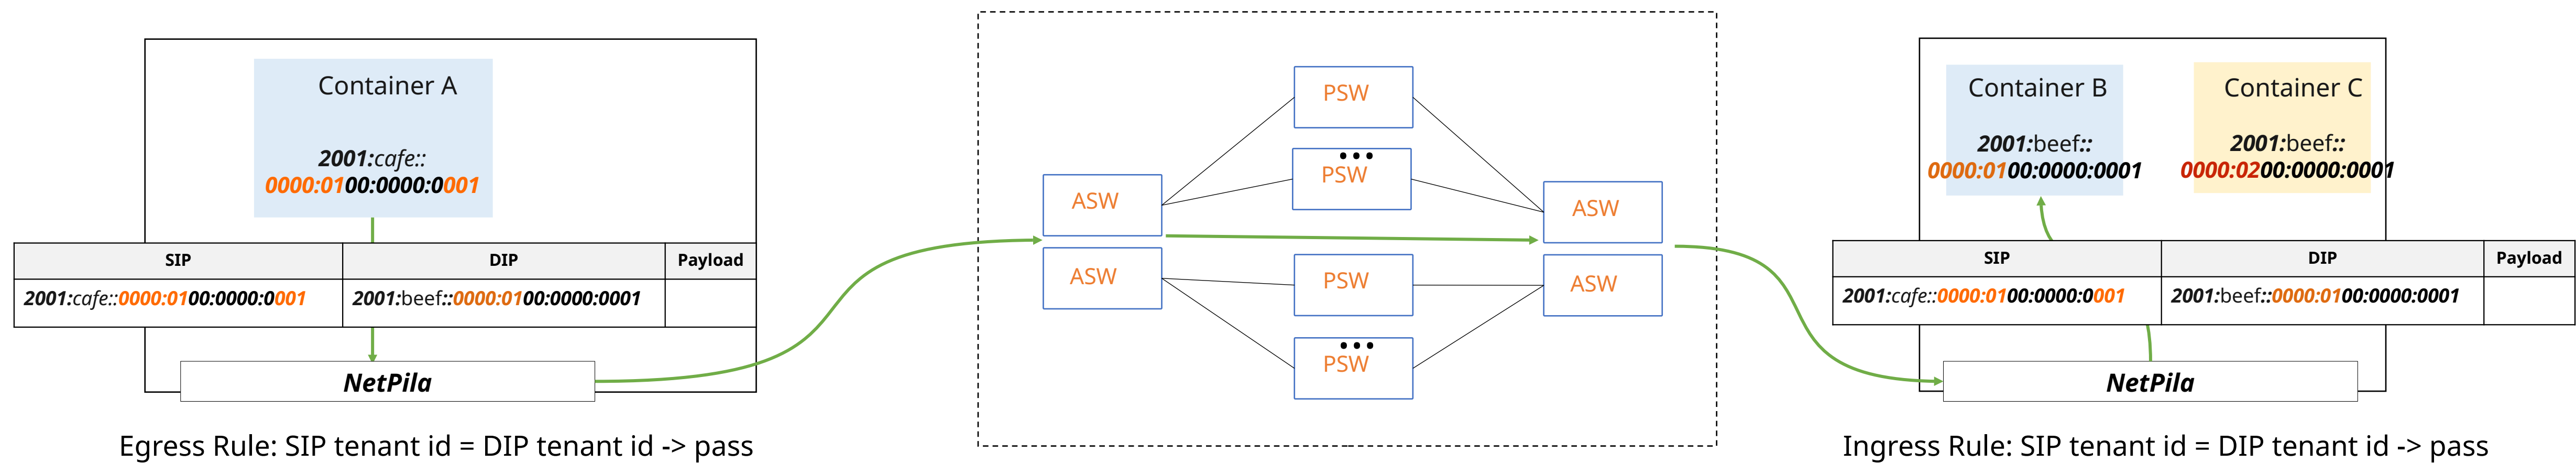


IPv6/SRv6 Solve Networking Challenges

SRv6 Container Network- Netpila

SRBE Forwarding Model & Lightweight Network Isolation

- SRBE model: Multi-path quality is consistent within the data center, requiring no path selection
- Receiving Container interface IP is placed in Destination IP field of Packet
- Matching Tenant ID field of DIP/SIP at sending and receiving node enables tenant isolation

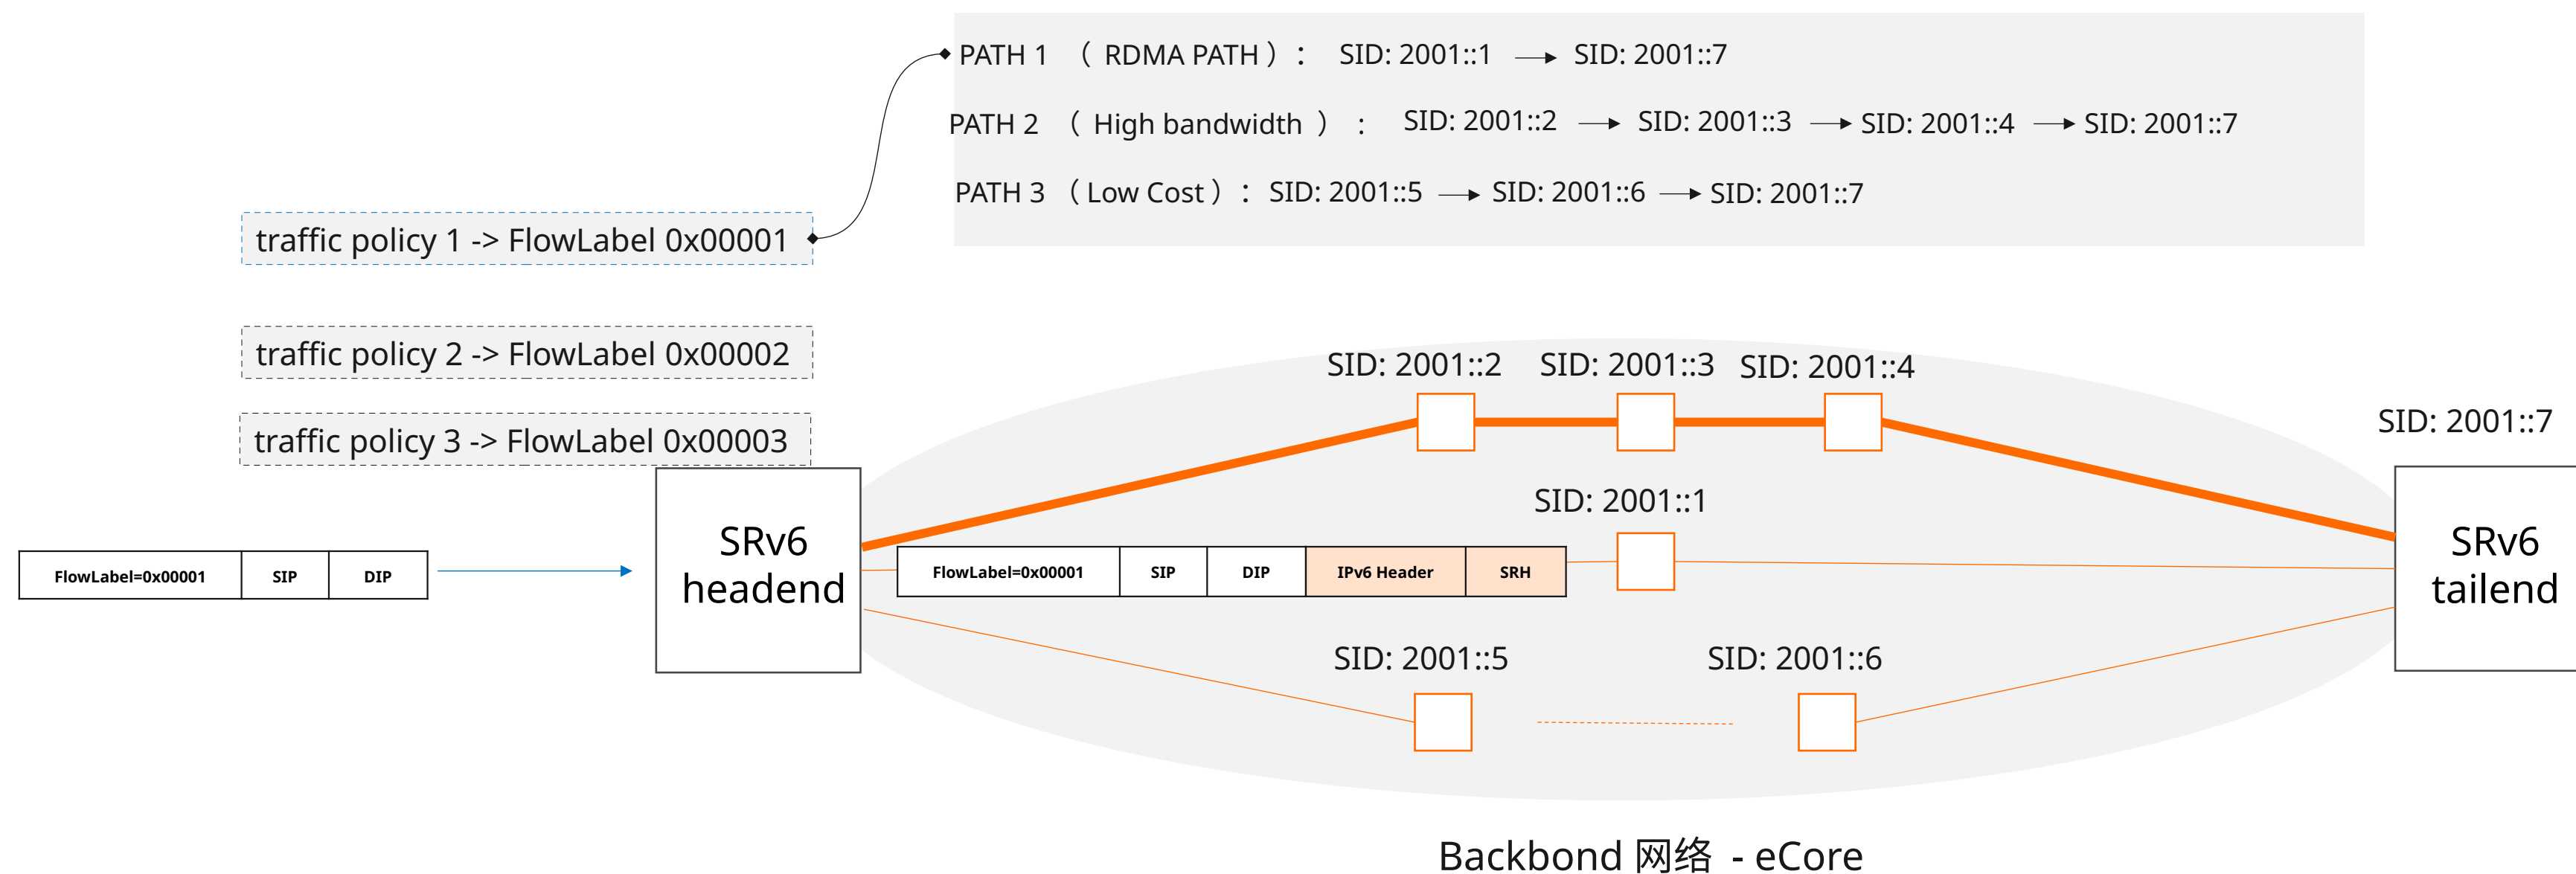


IPv6/SRv6 Solve Networking Challenges

SRv6-Based Backbone Network - eCore

SRv6-Based Cross-Region Path Selection

- At the backbone SRv6 headend, configure flow label and SID-List Path mapping.
- Based on Traffic Policy, match the flow label value of packets and encapsulate them with outer IPv6 header and SRH;

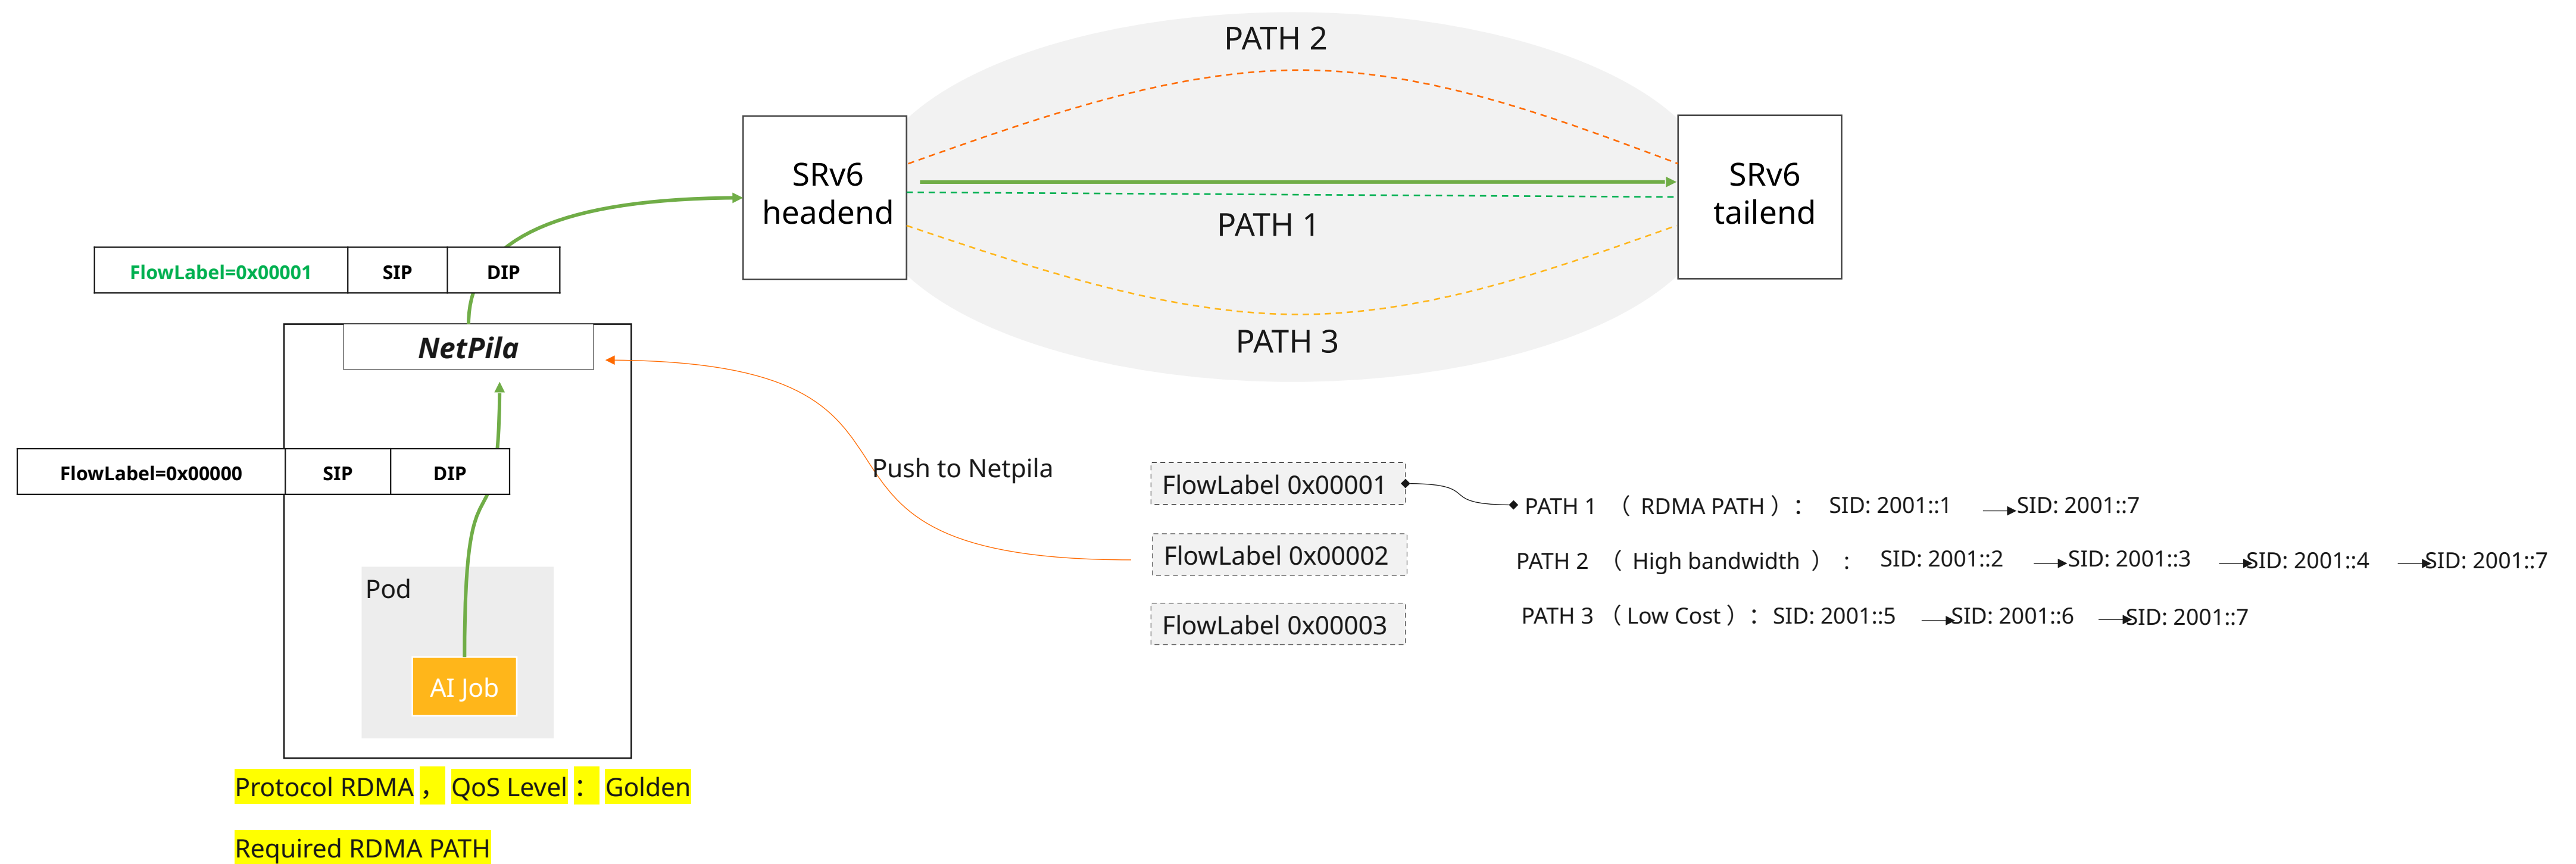


IPv6/SRv6 Solve Networking Challenges

SRv6-Based Backbone Network - eCore

Server-Network Collaboration for Cross-Region Path Selection

- Netpila on the server side receives Flow Label Mapping rules via the control plane.
- Flow Label values in packets are modified according to application required service and protocol type



Conclusion

Deployment Status

Data Center/ Netpila

Deployed **14+** clusters, covering **90,000** servers.

Support Products in Alibaba Cloud

- AI Computing
- Cloud-native Container Services
- Data Base

Business Impact

AI Computing in Alibaba Cloud is now fully based on IPv6/SRv6 solutions.

2% ↑
Throughput

10% ↓
Latency

3x ↑
Job Starup Speed

 Alibaba Cloud | MORE THAN JUST CLOUD