Optimizing Content Delivery Network Design

Andy Banta – Storage Janitor

Powered by MagnitionIO

SNIA COMPUTE, MEMORY, AND STORAGE SUMMIT

Solutions, Architectures, and Community

VIRTUAL EVENT, MAY 21-22, 2024

Andy Banta

Magnition.io (Consultant)
SolidFire (VMware development, acq. by NetApp)
DataGravity (Container exploitation lead)
VMware (iSCSI Tech Lead, IPO)
Sun Microsystems (Initial Fibre Channel development)
Patent, early distributed network projects, data acquisition
@andybanta

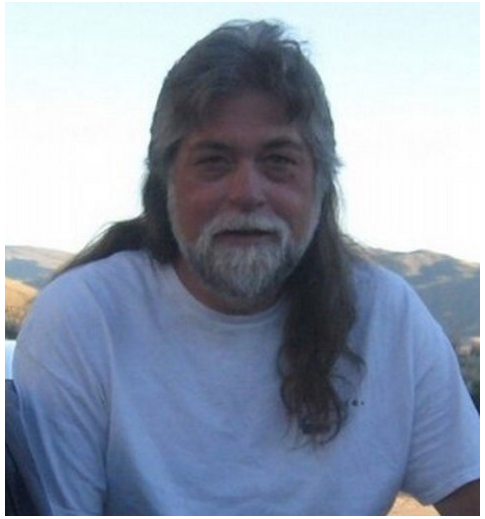

ABOUT MAGNITION

STORAGE PERFORMANCE, REINVENTED

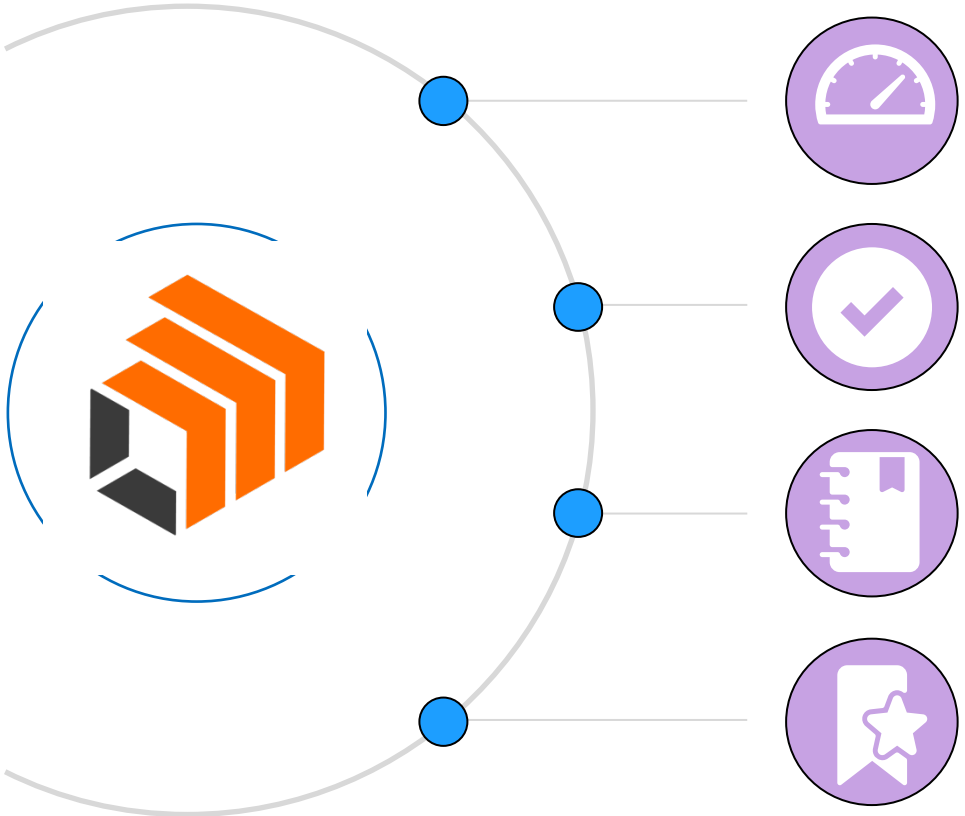

World's First Real-Time Data Placement Optimization

Patented technology is a first for the industry.

Proven At-Scale, with Production Workloads

Use customer traces to fully test diverse workloads in real-time.

Peer-Reviewed and Published in Leading Journals

Multiple industry articles published and reviewed.

Award-Winning, Patented Technology

3-time award winner for innovative technology.

Engineering simulation

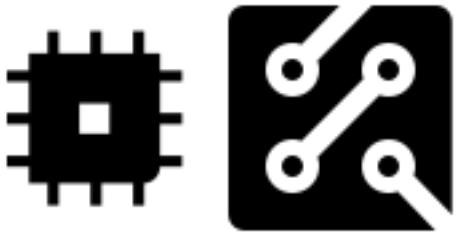

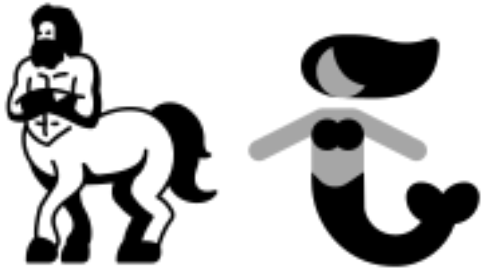

Engineering simulation

- Cheaper, faster, more flexible than system building
- Engineering design uses simulations, why not software?

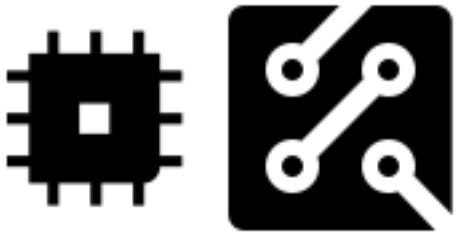

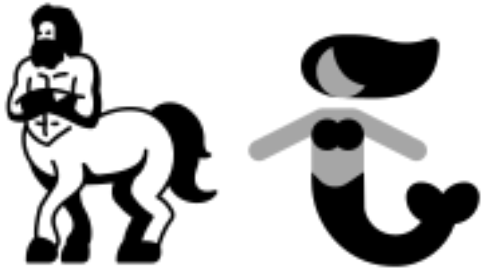

What is Content Delivery

Content Delivery Networks

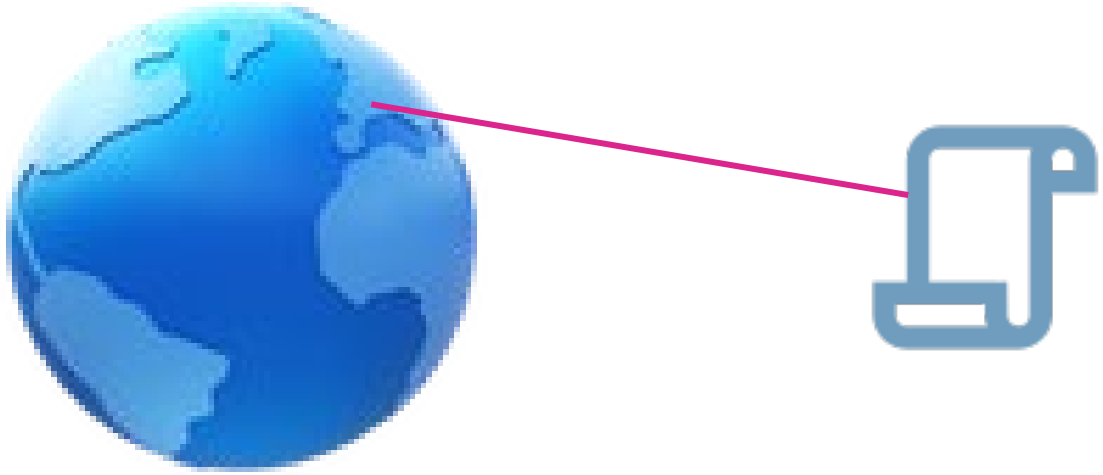

Content Delivery Networks

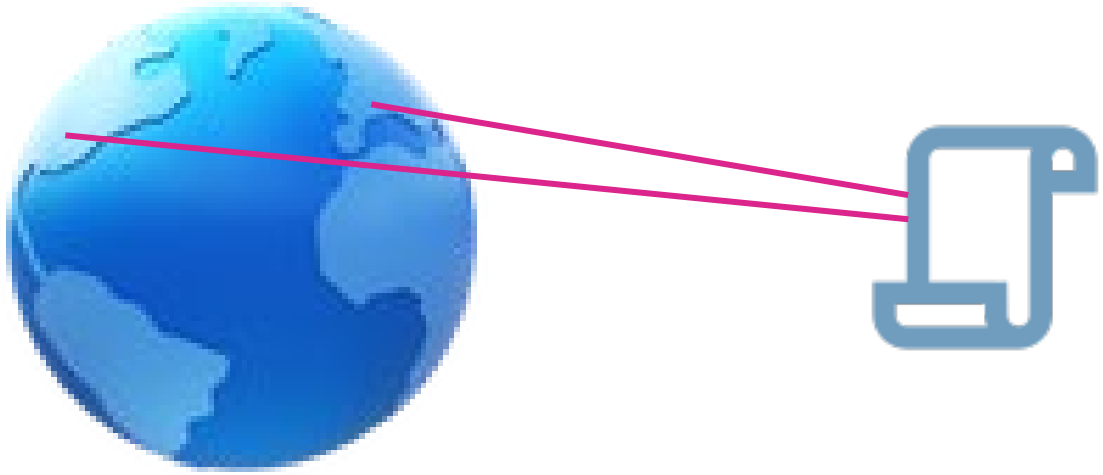

Content Delivery Networks

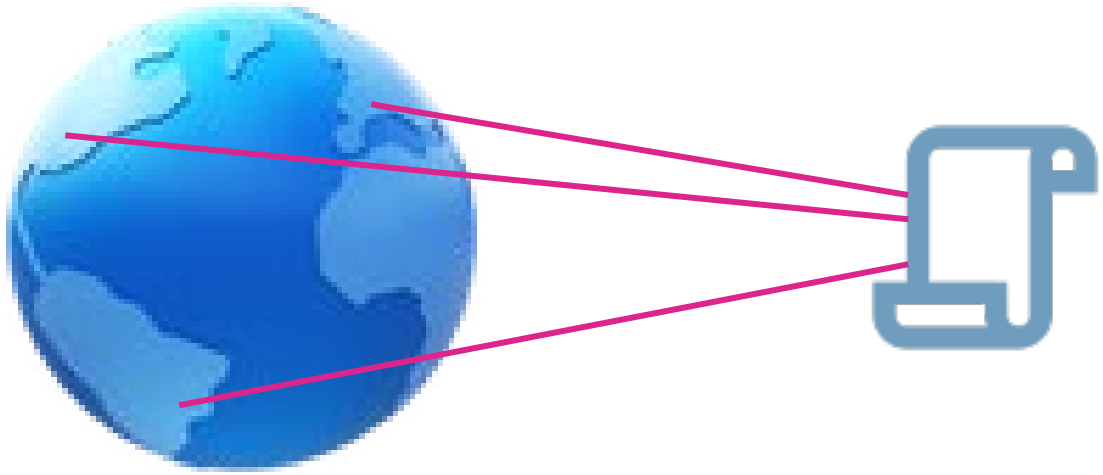

Content Delivery Networks

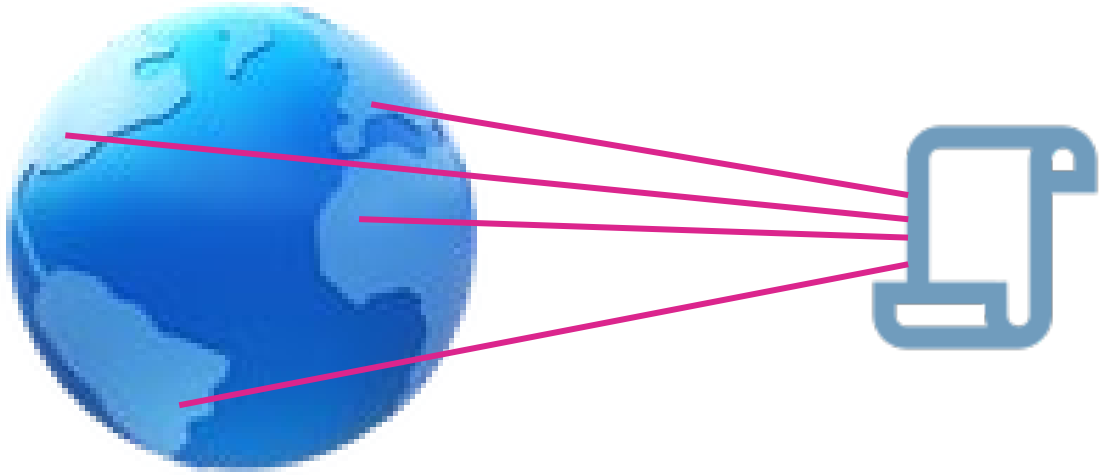

Content Delivery Networks

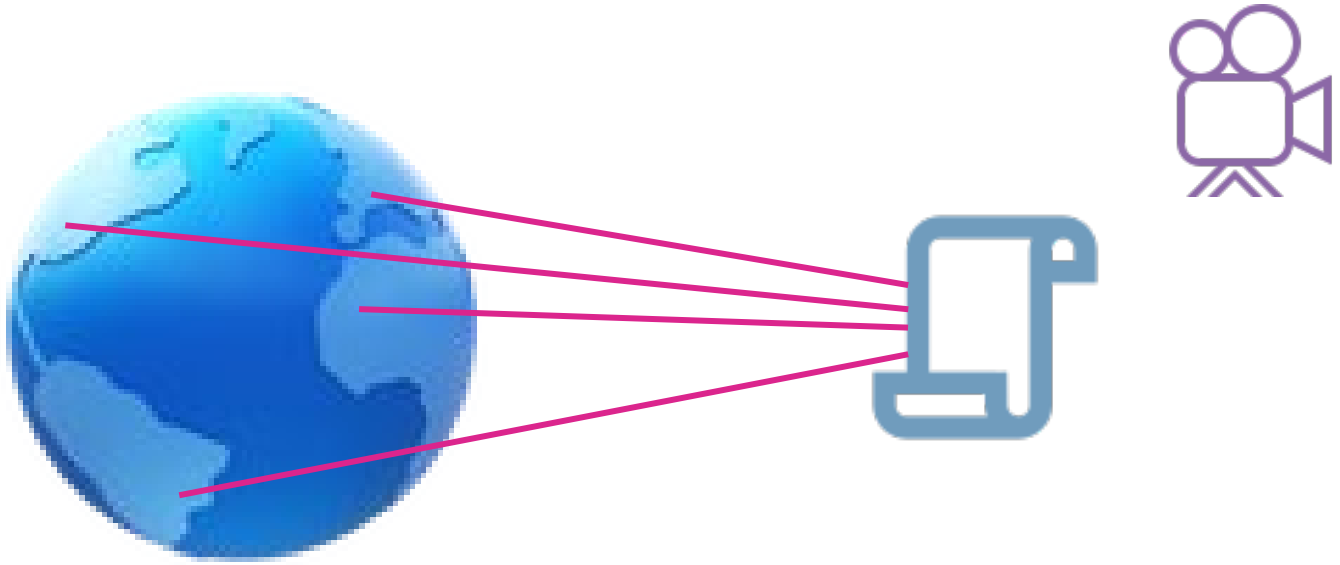

Content Delivery Networks

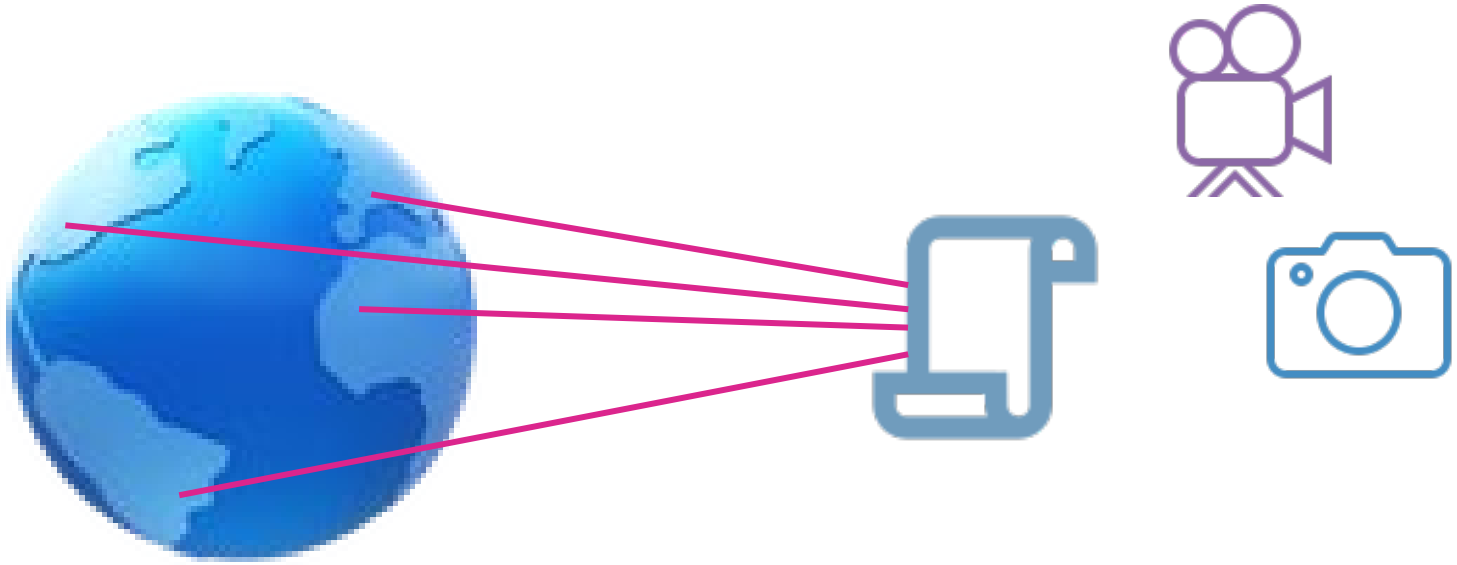

Content Delivery Networks

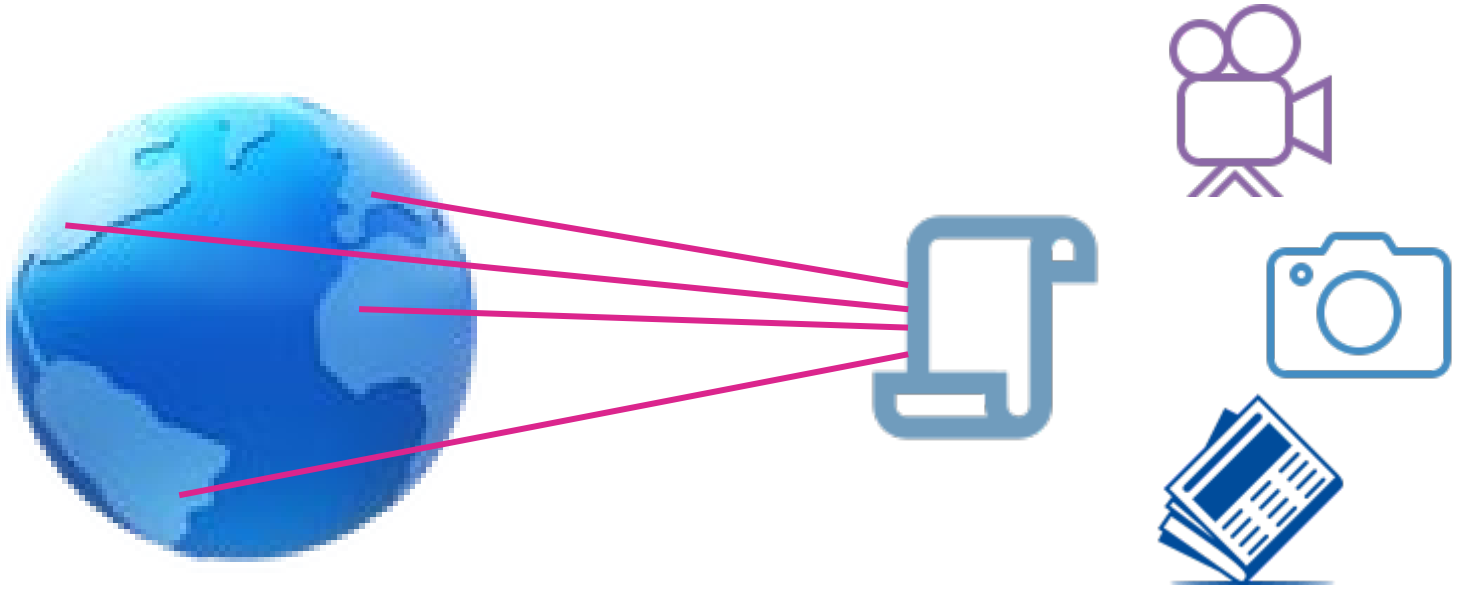

Content Delivery Networks

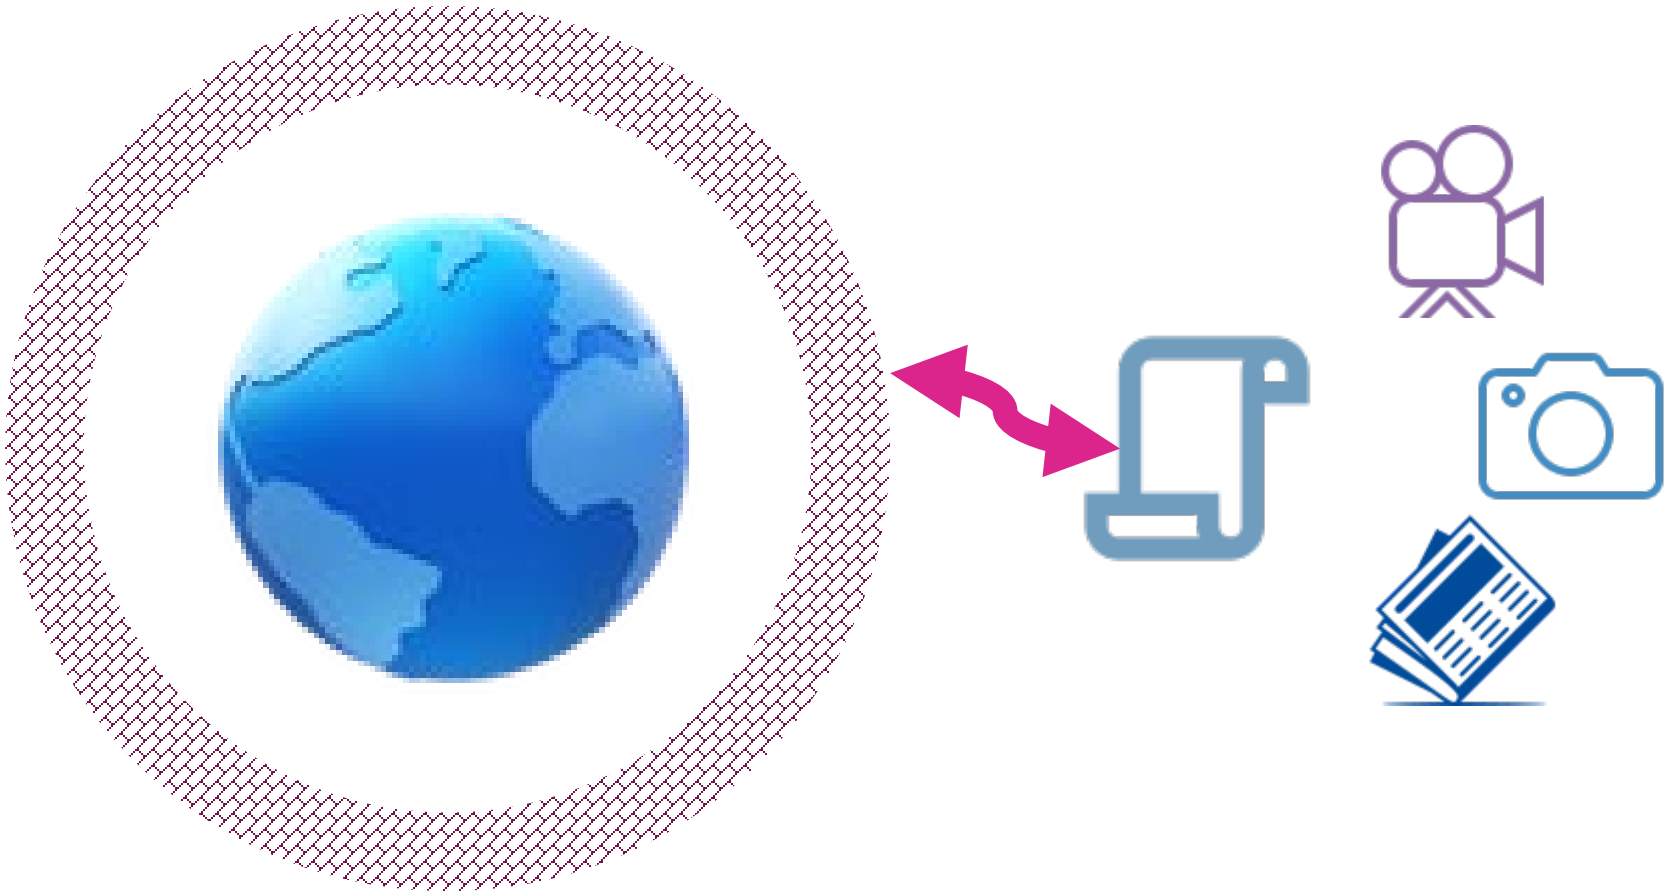

Content Delivery Networks

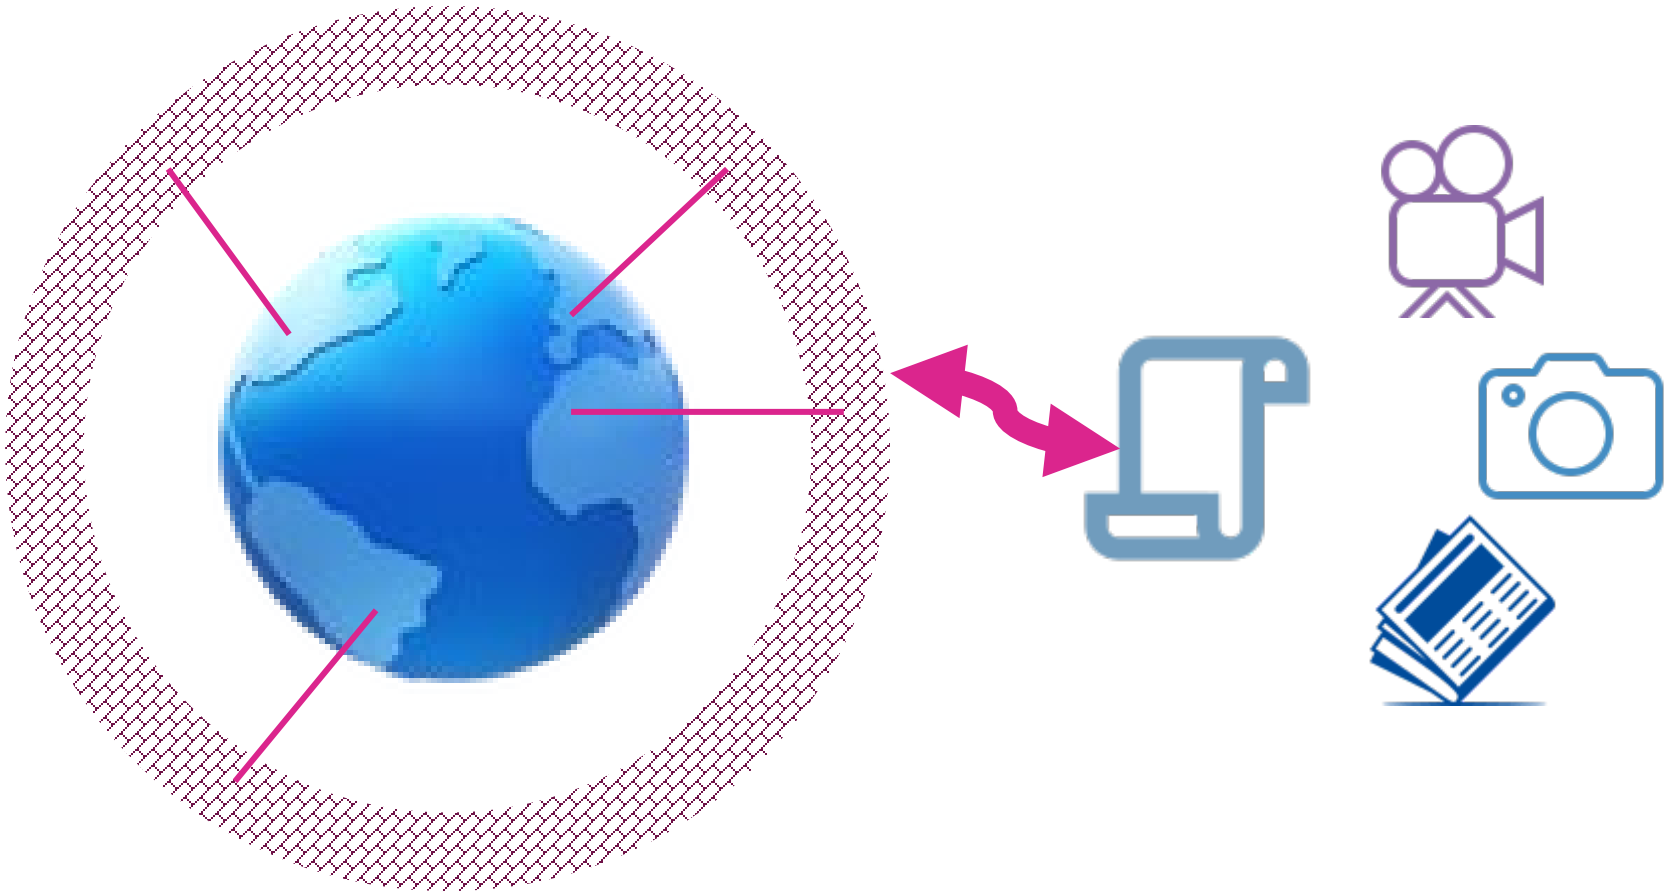

Content Delivery Networks

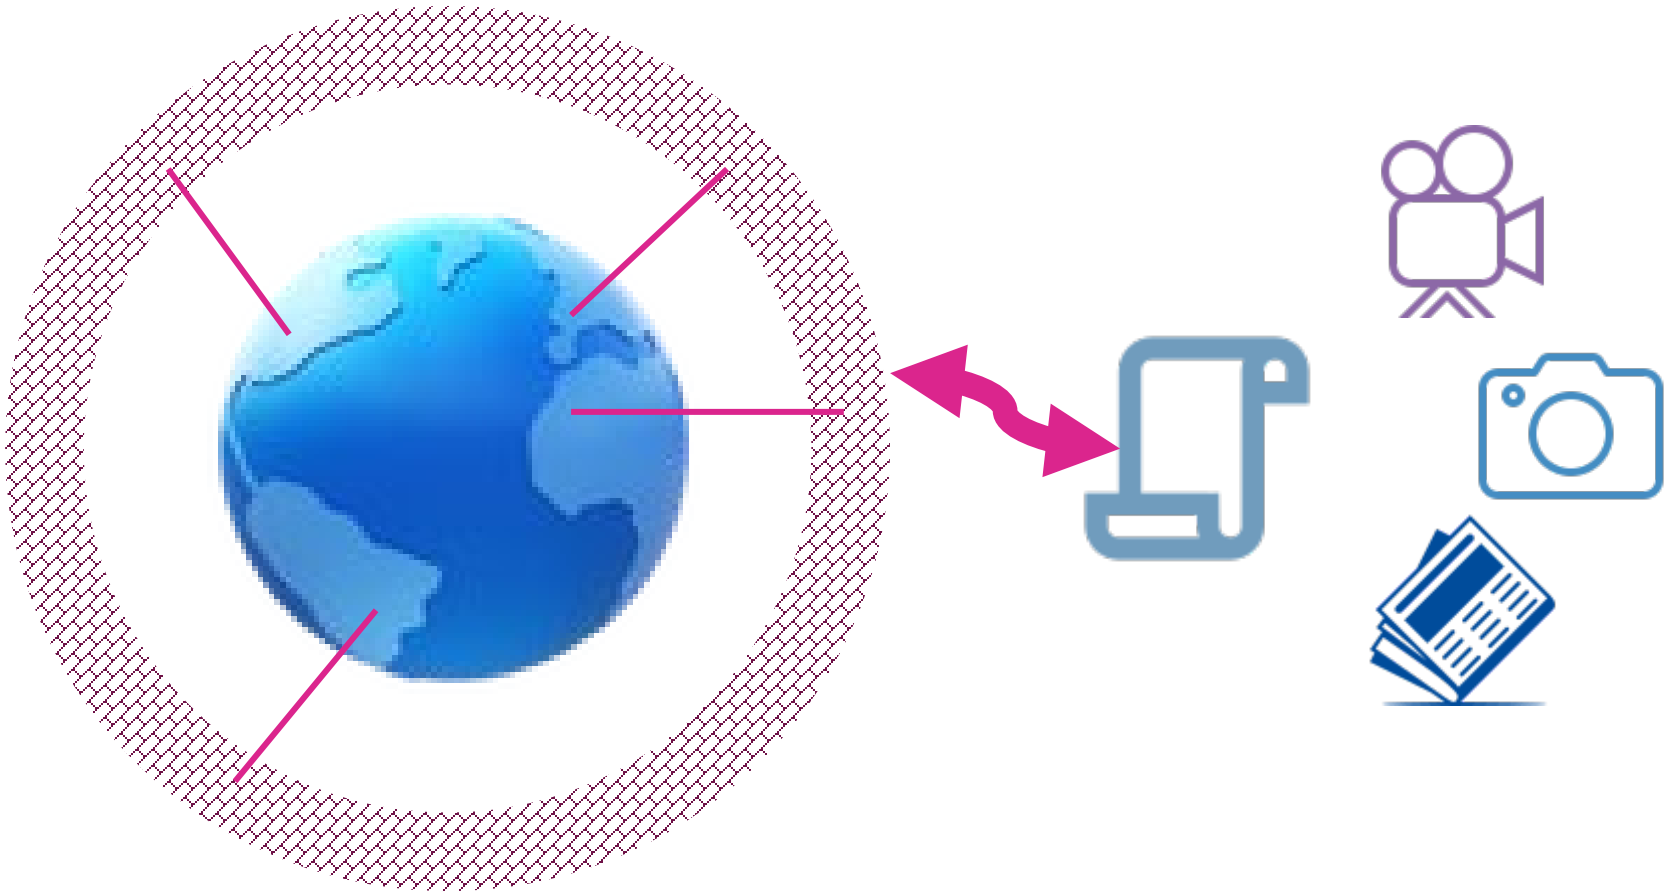

Content Delivery Networks

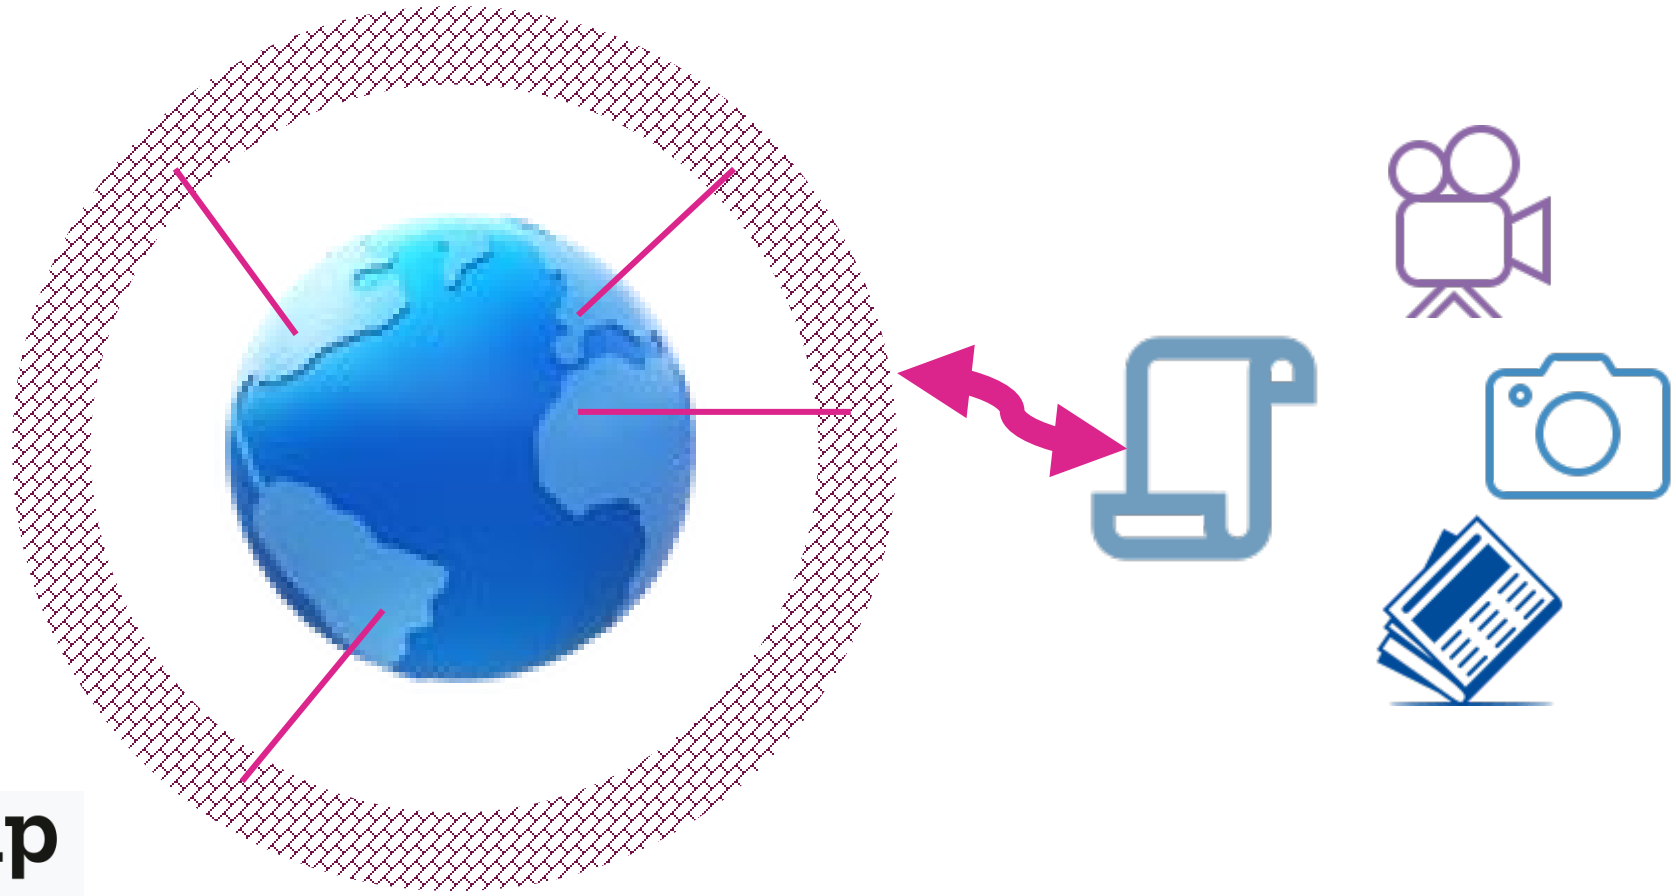

Content Delivery Networks

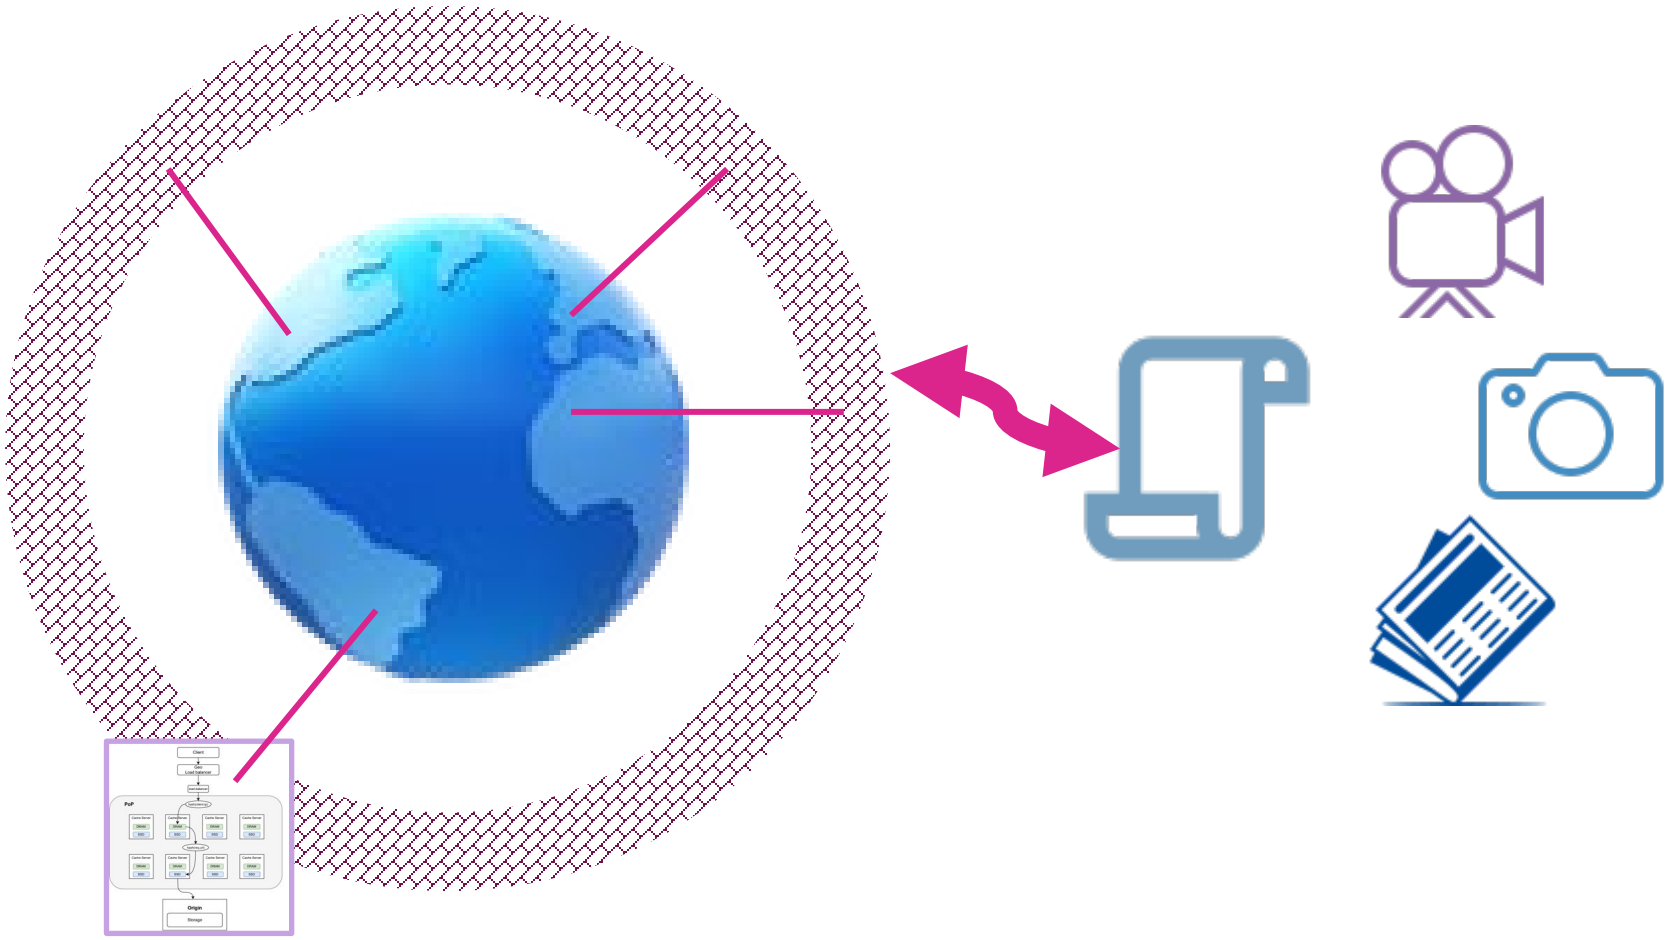

Content Delivery Networks

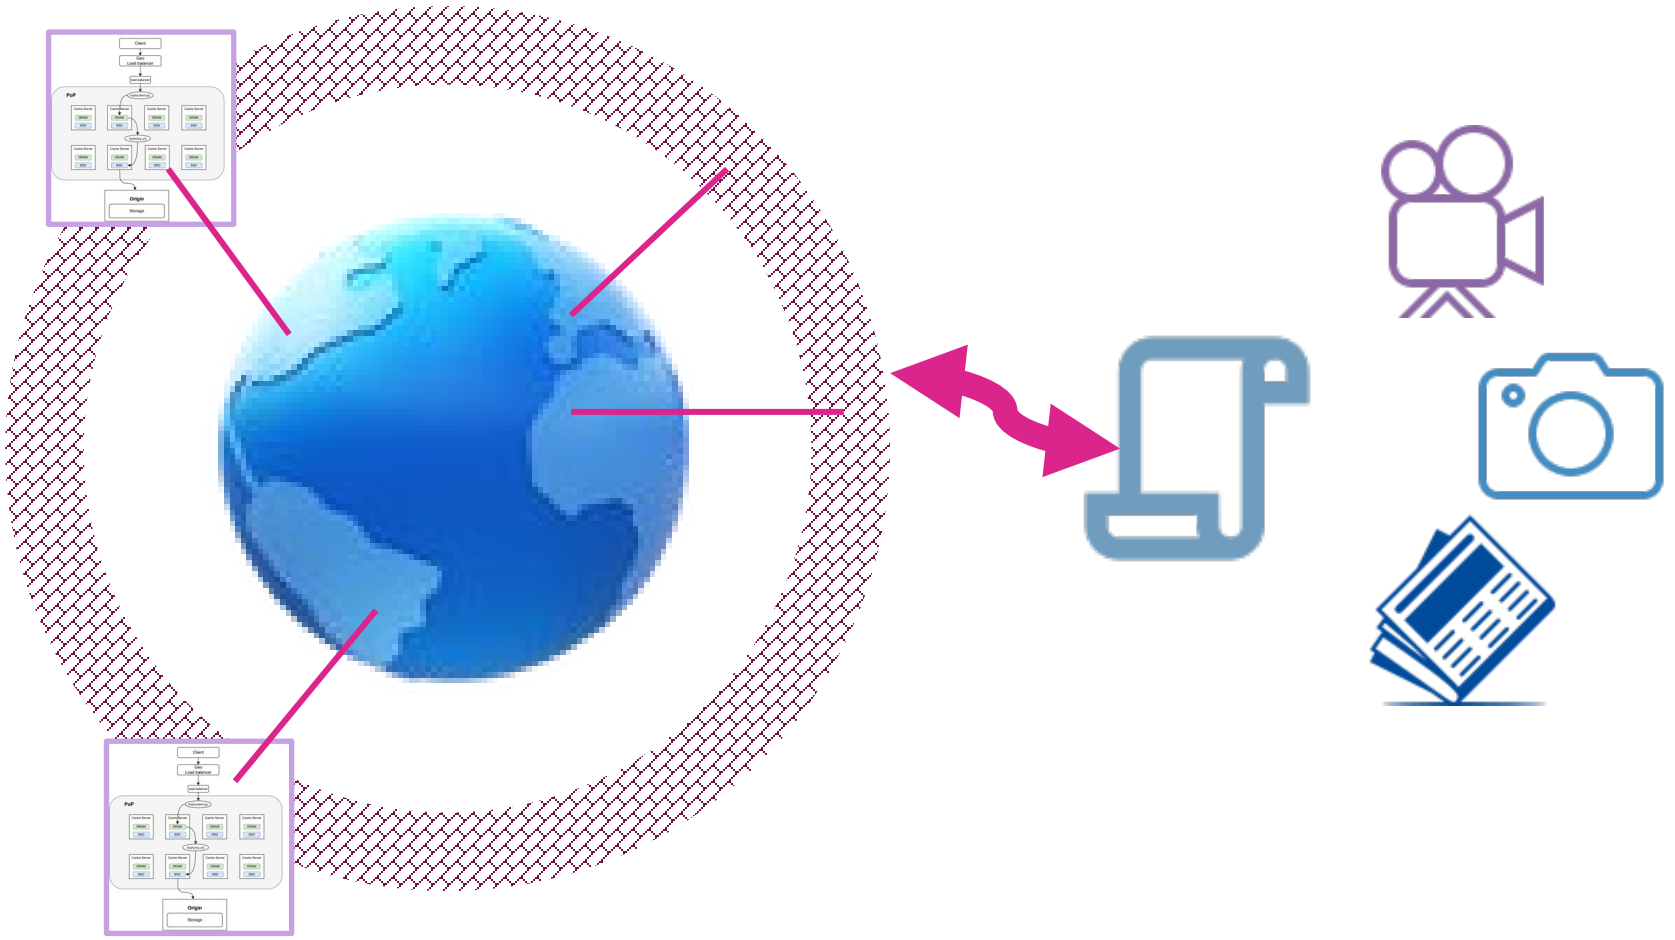

Content Delivery Networks

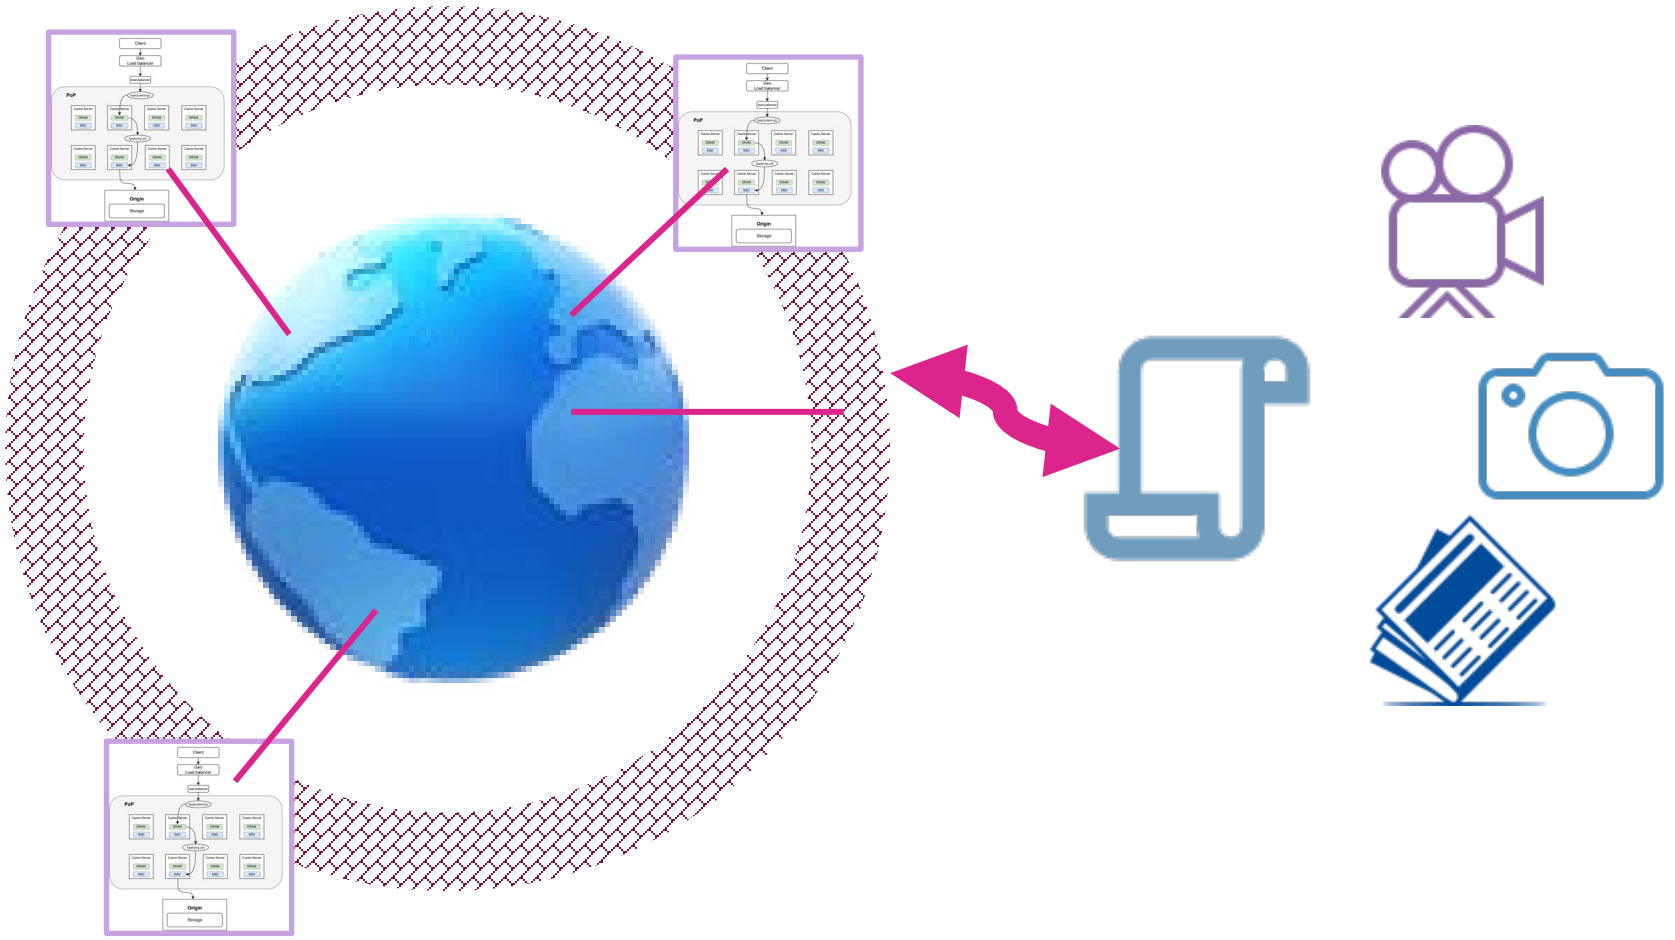

Content Delivery Networks

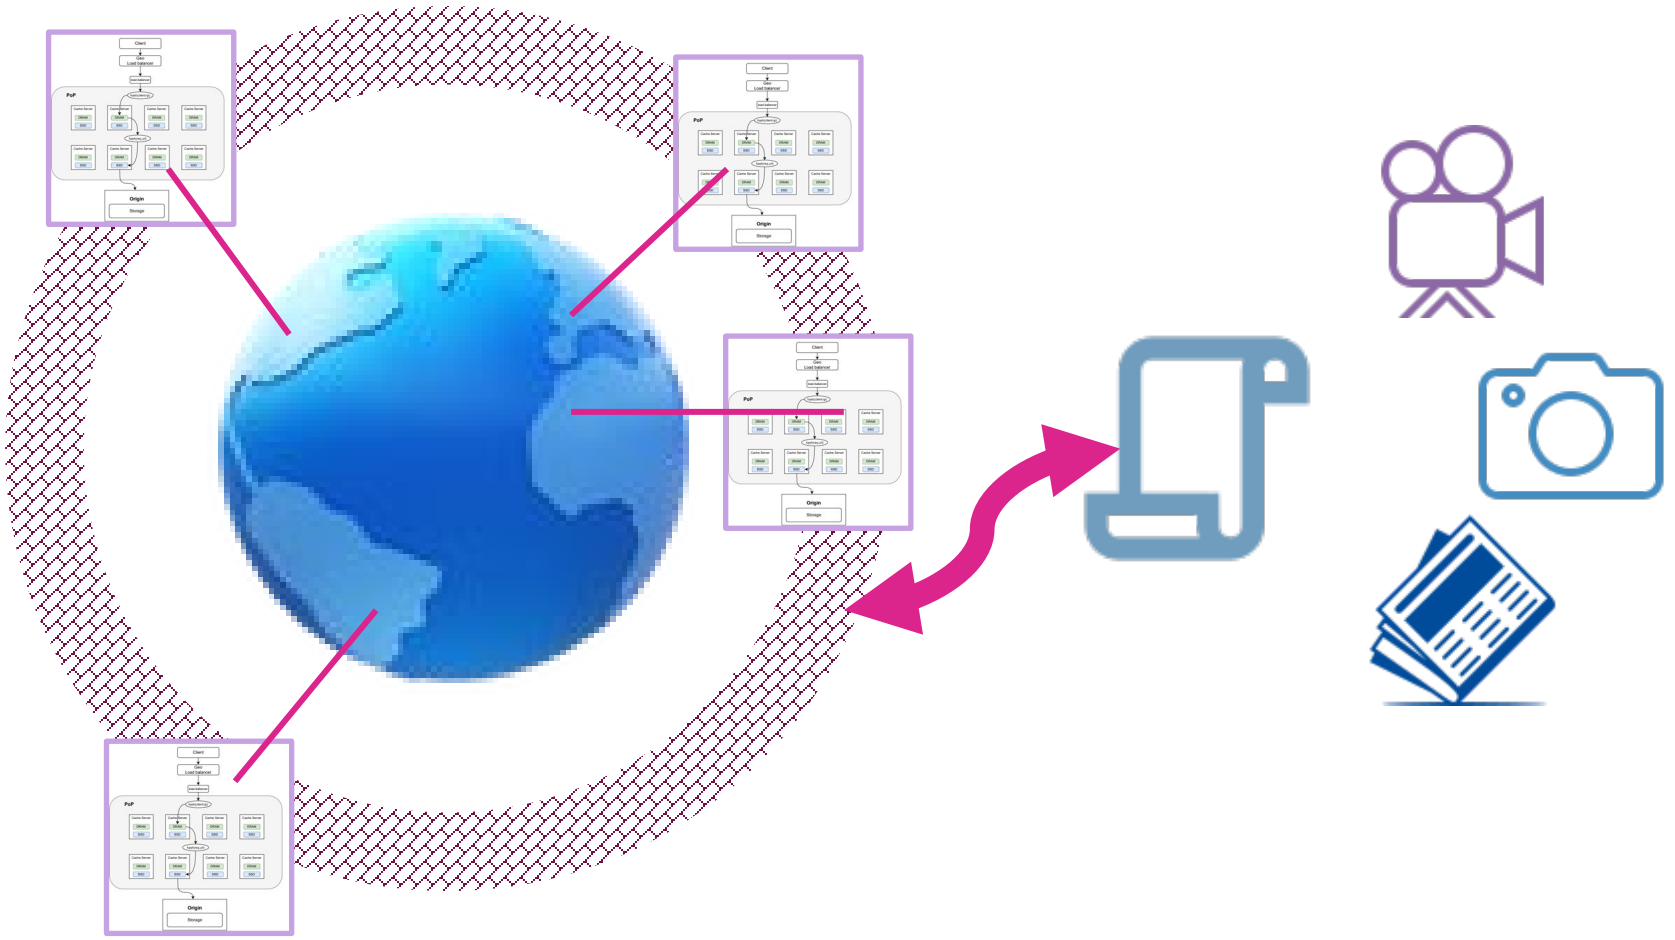

CDN PoP

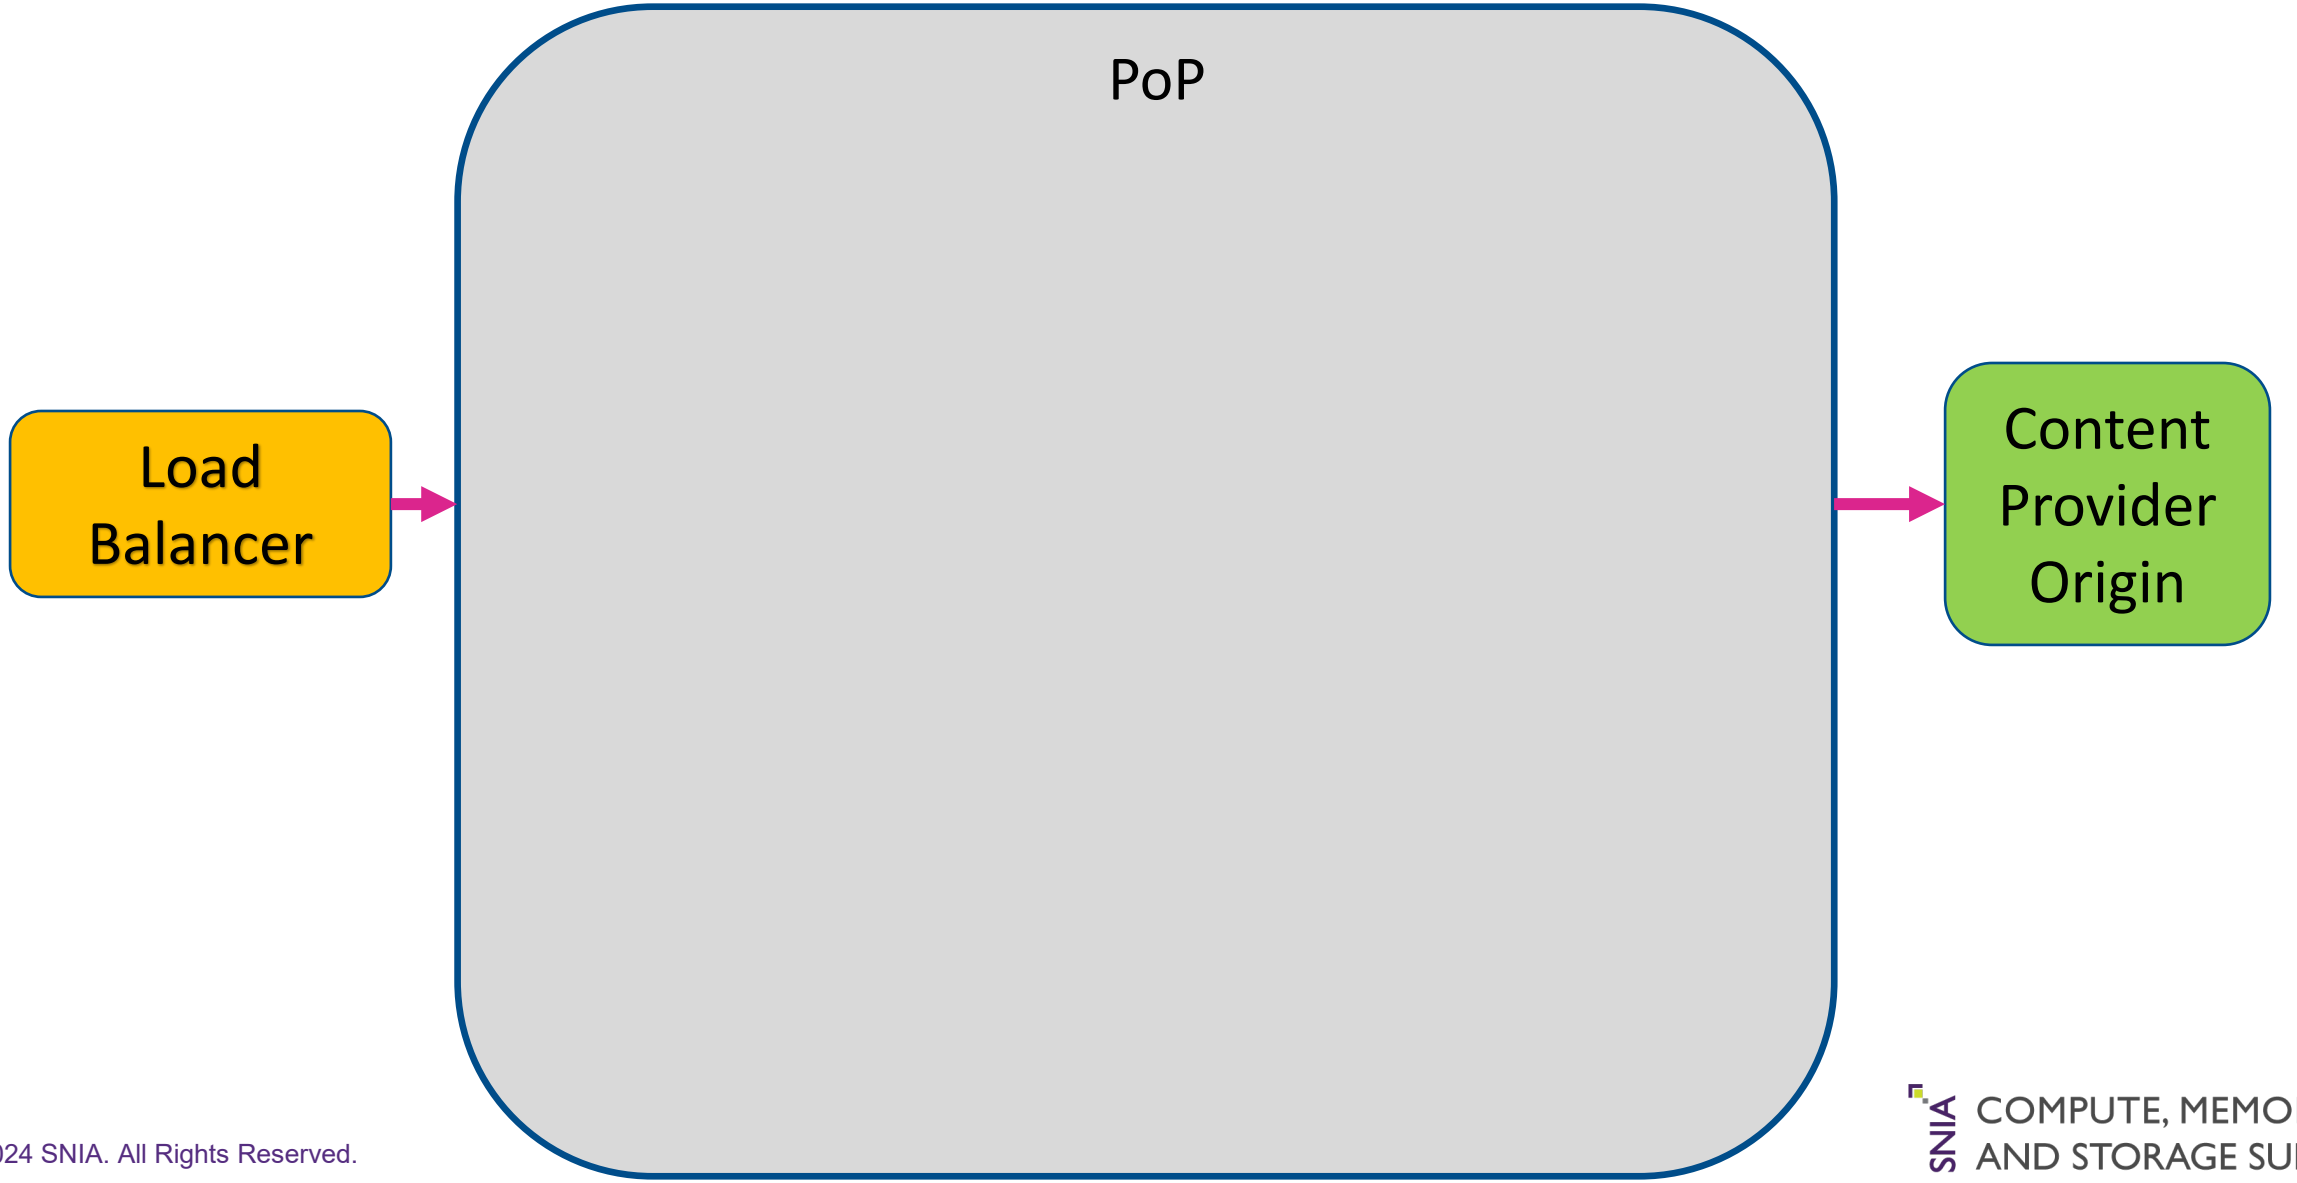

CDN PoP

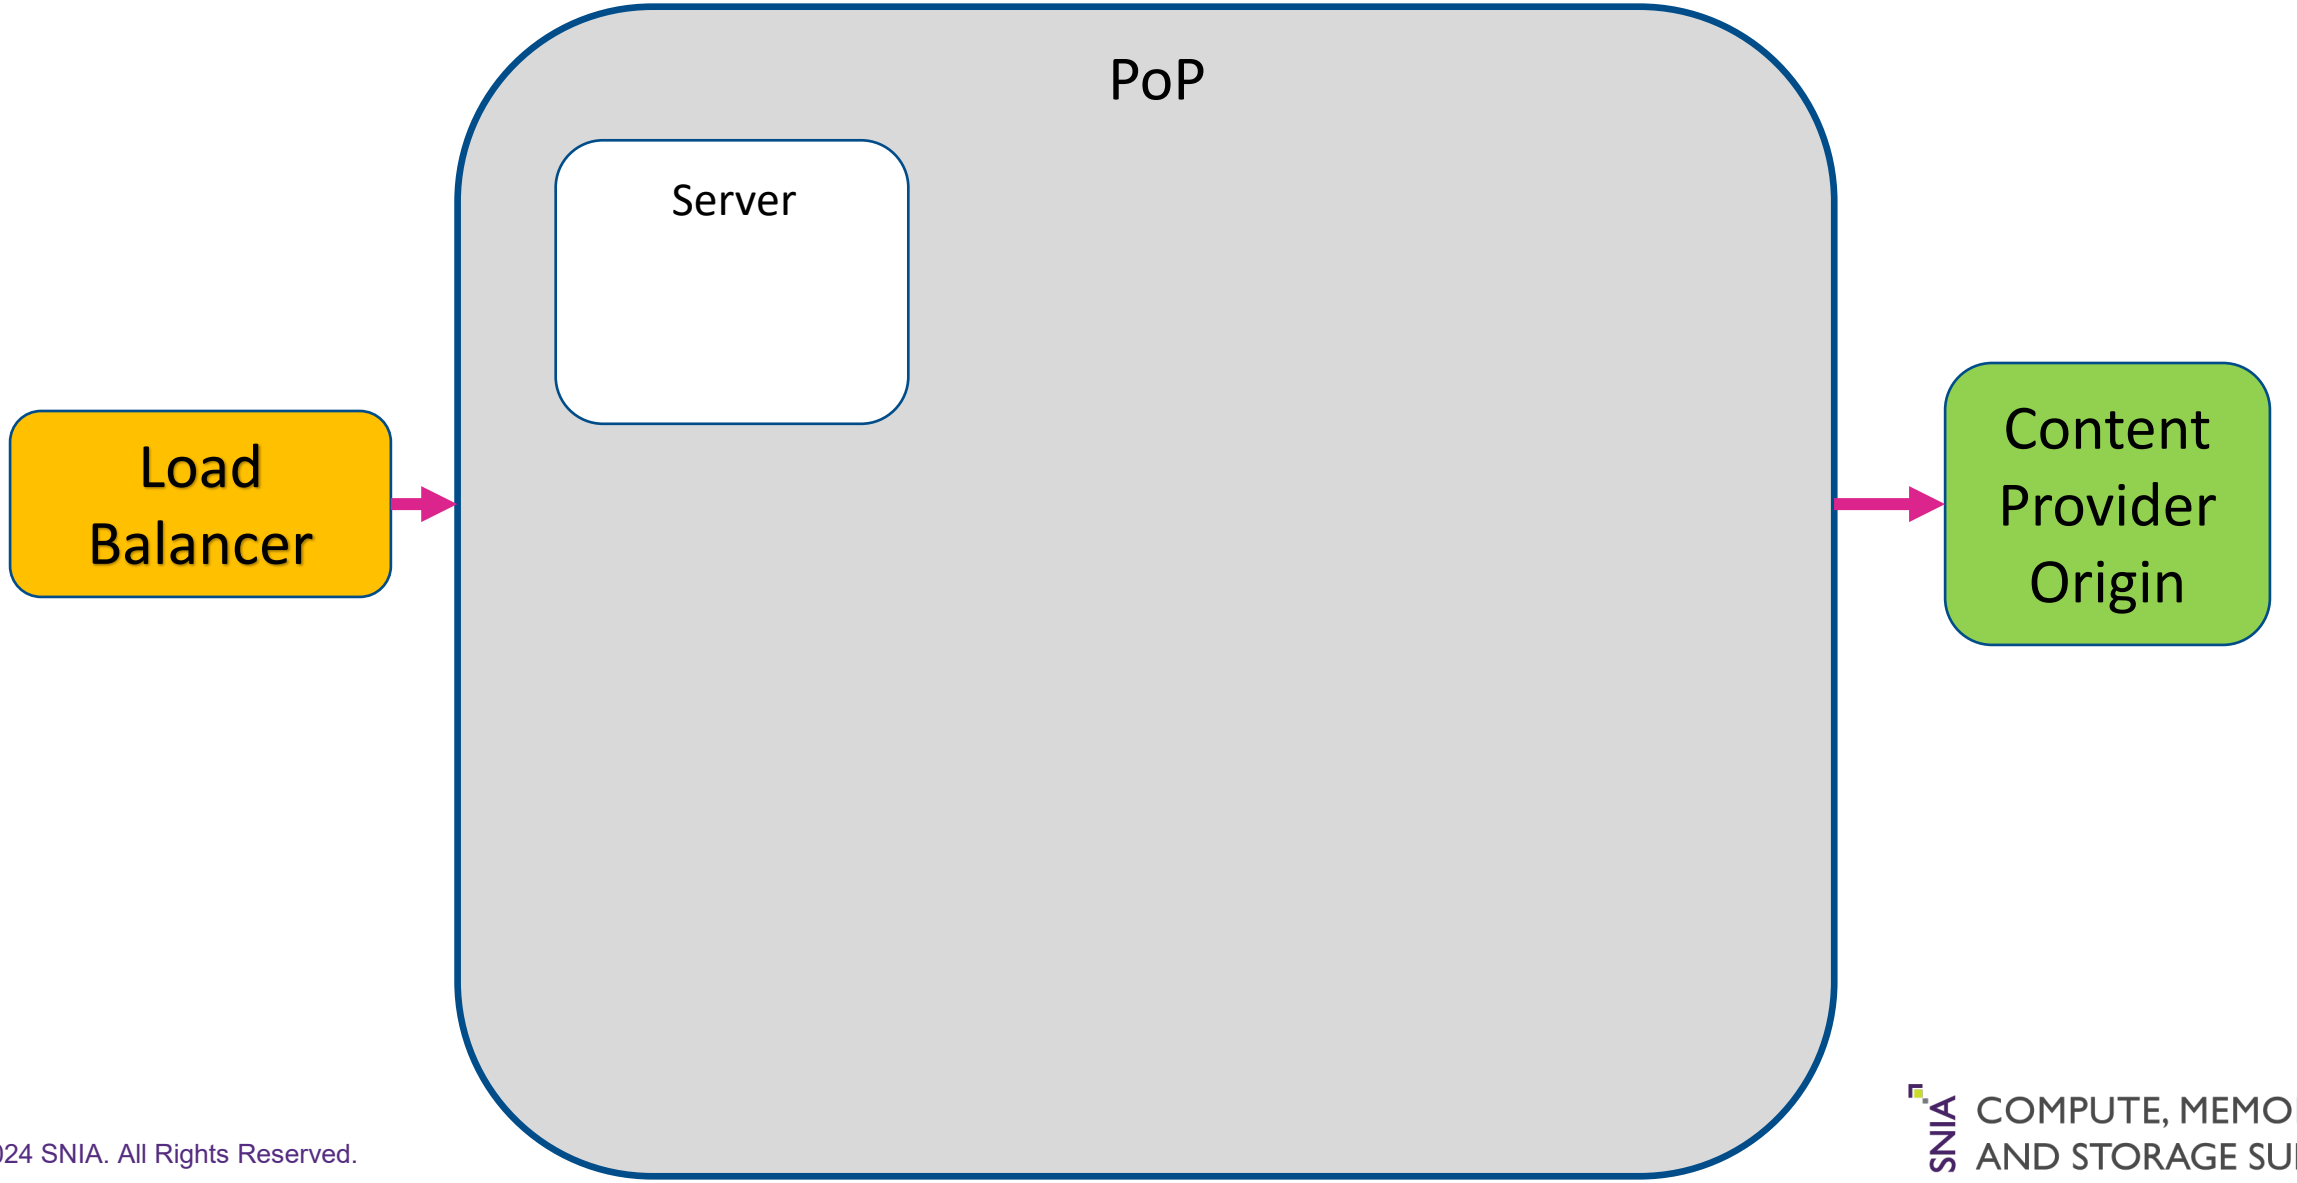

CDN PoP -- Compute

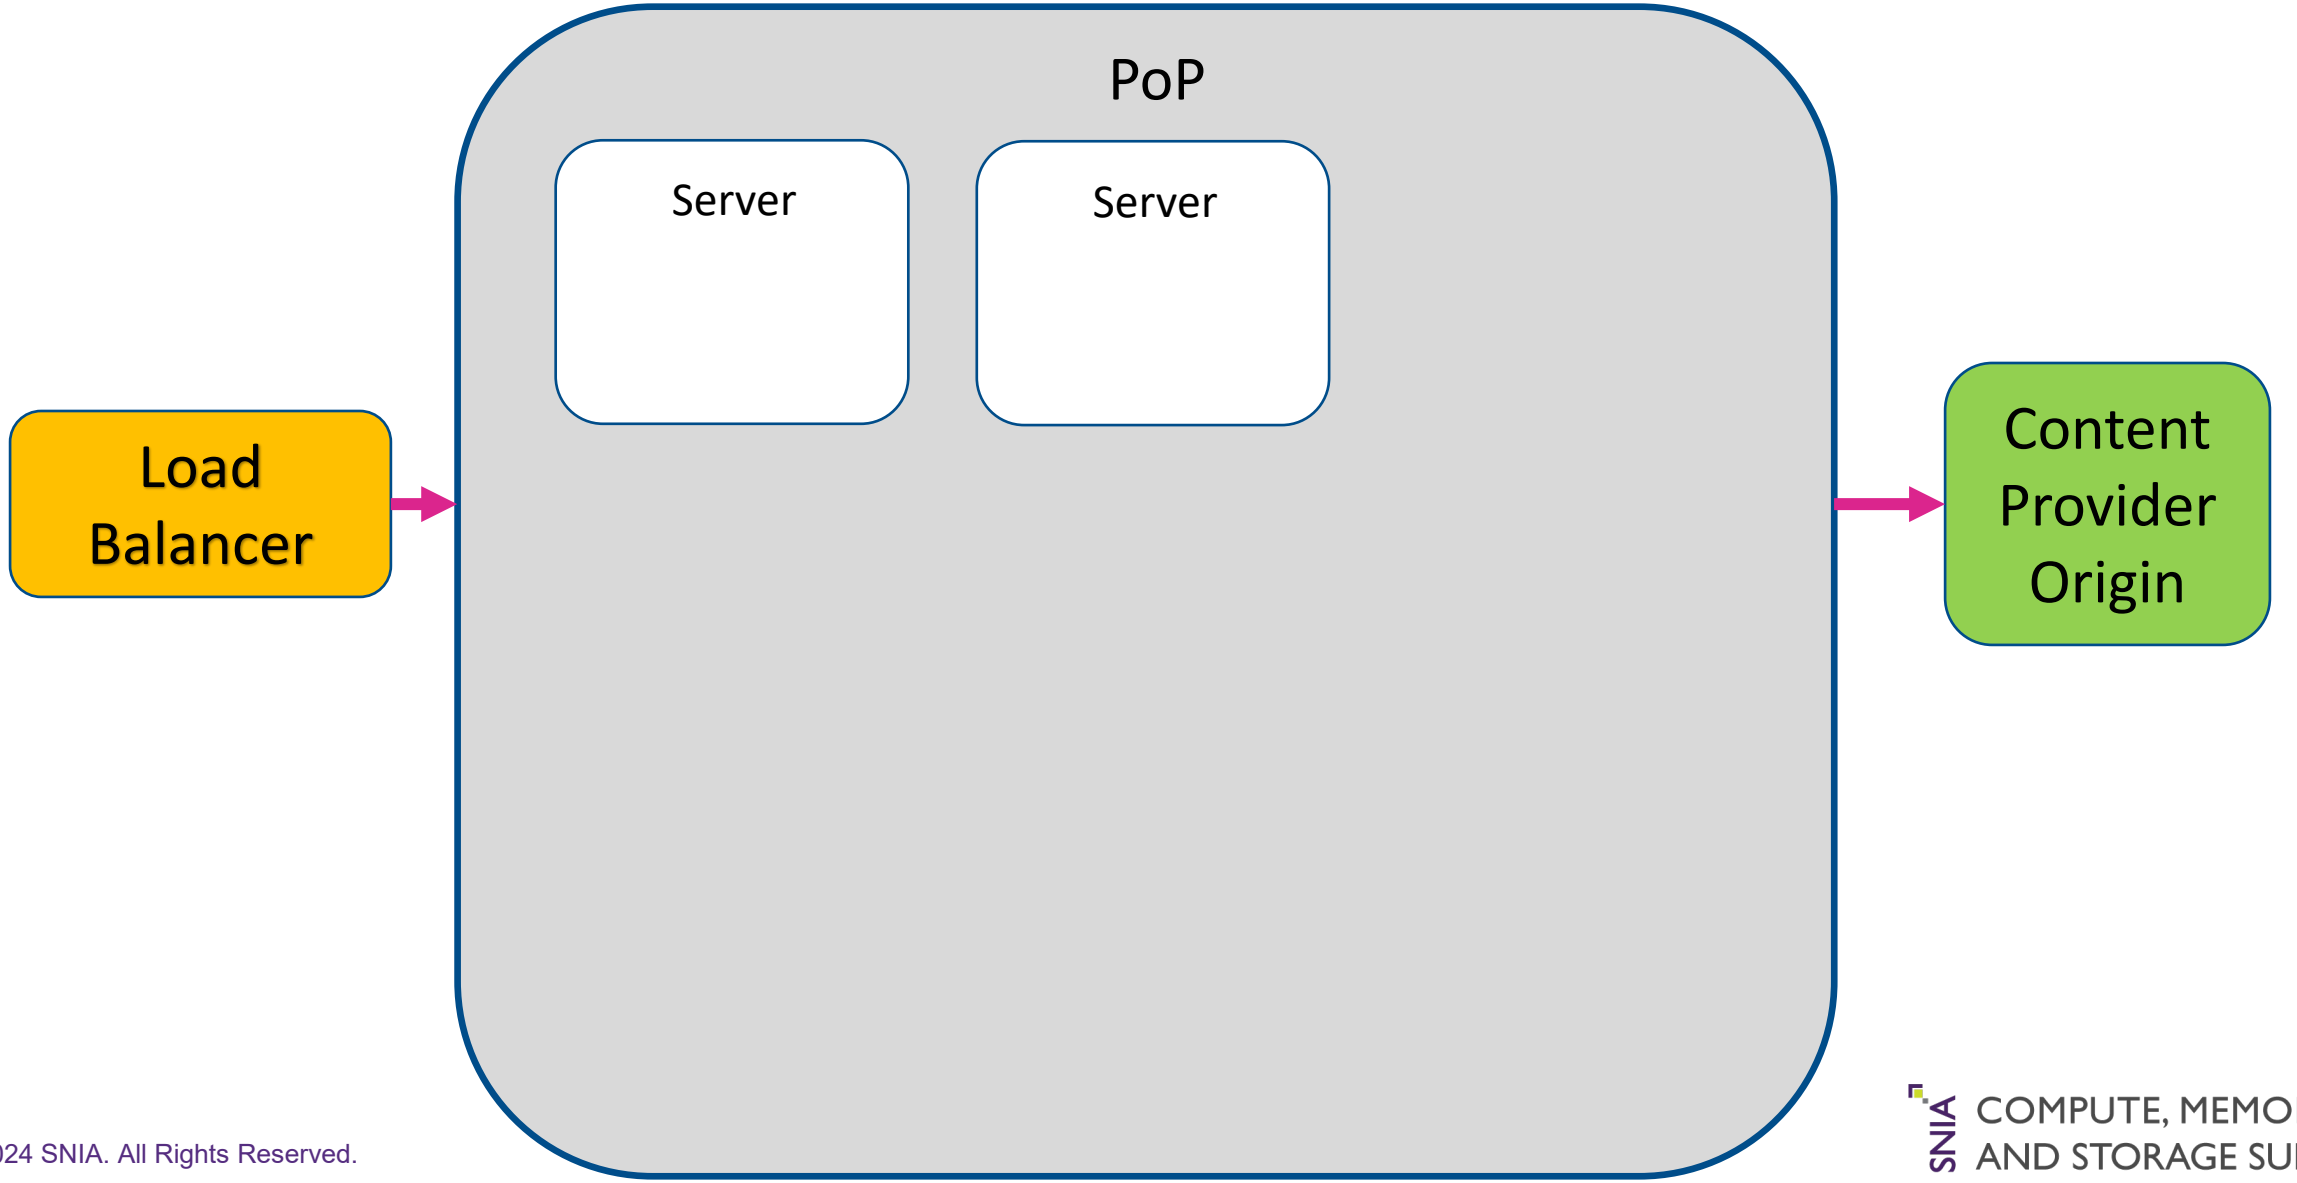

CDN PoP -- Compute

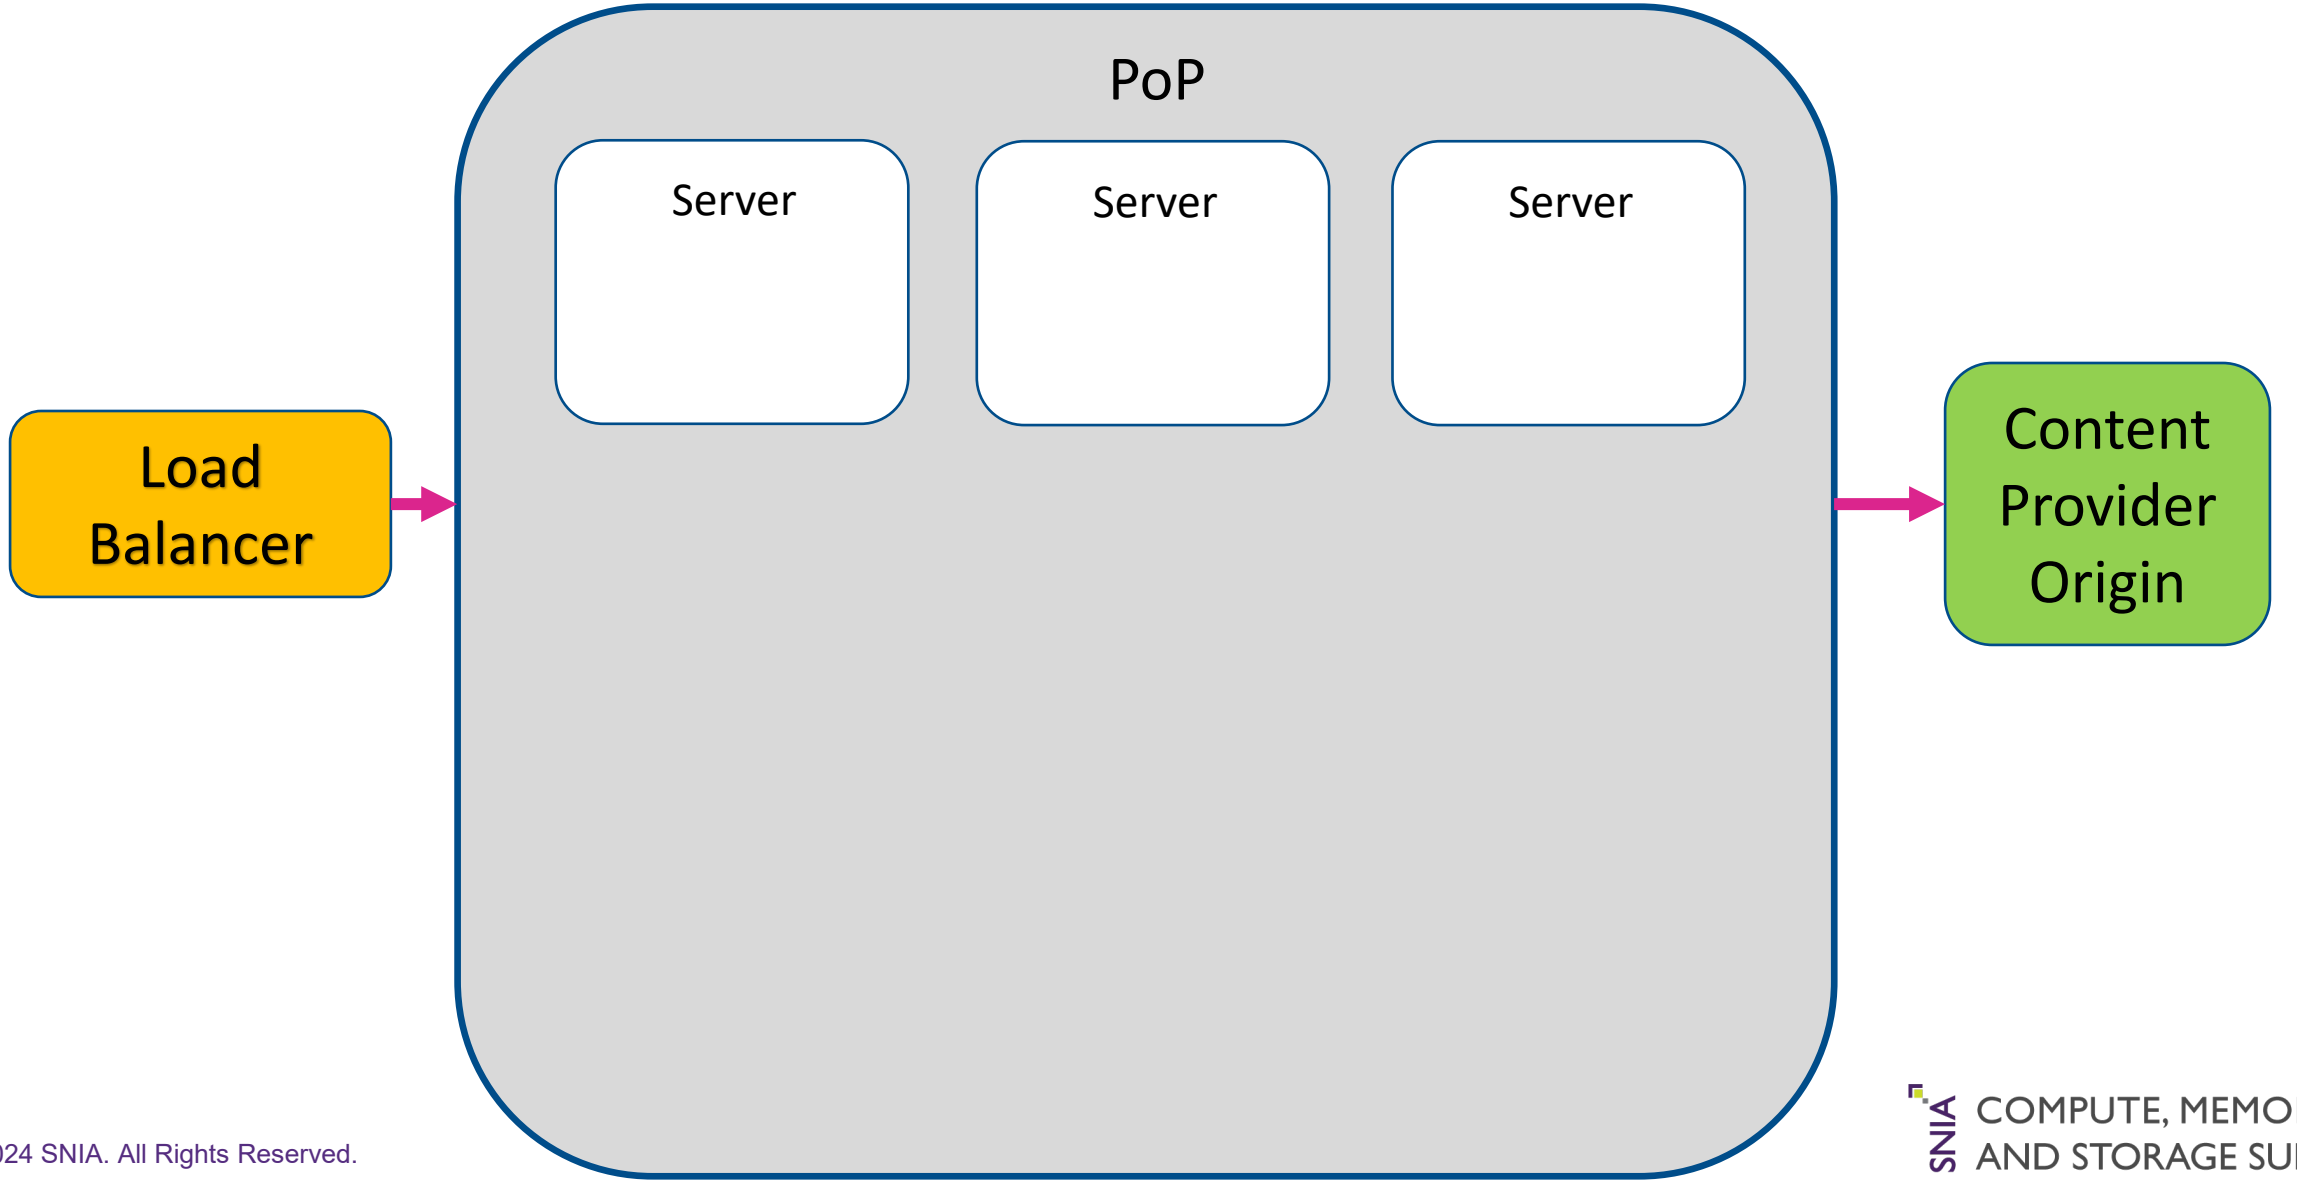

CDN PoP -- Compute

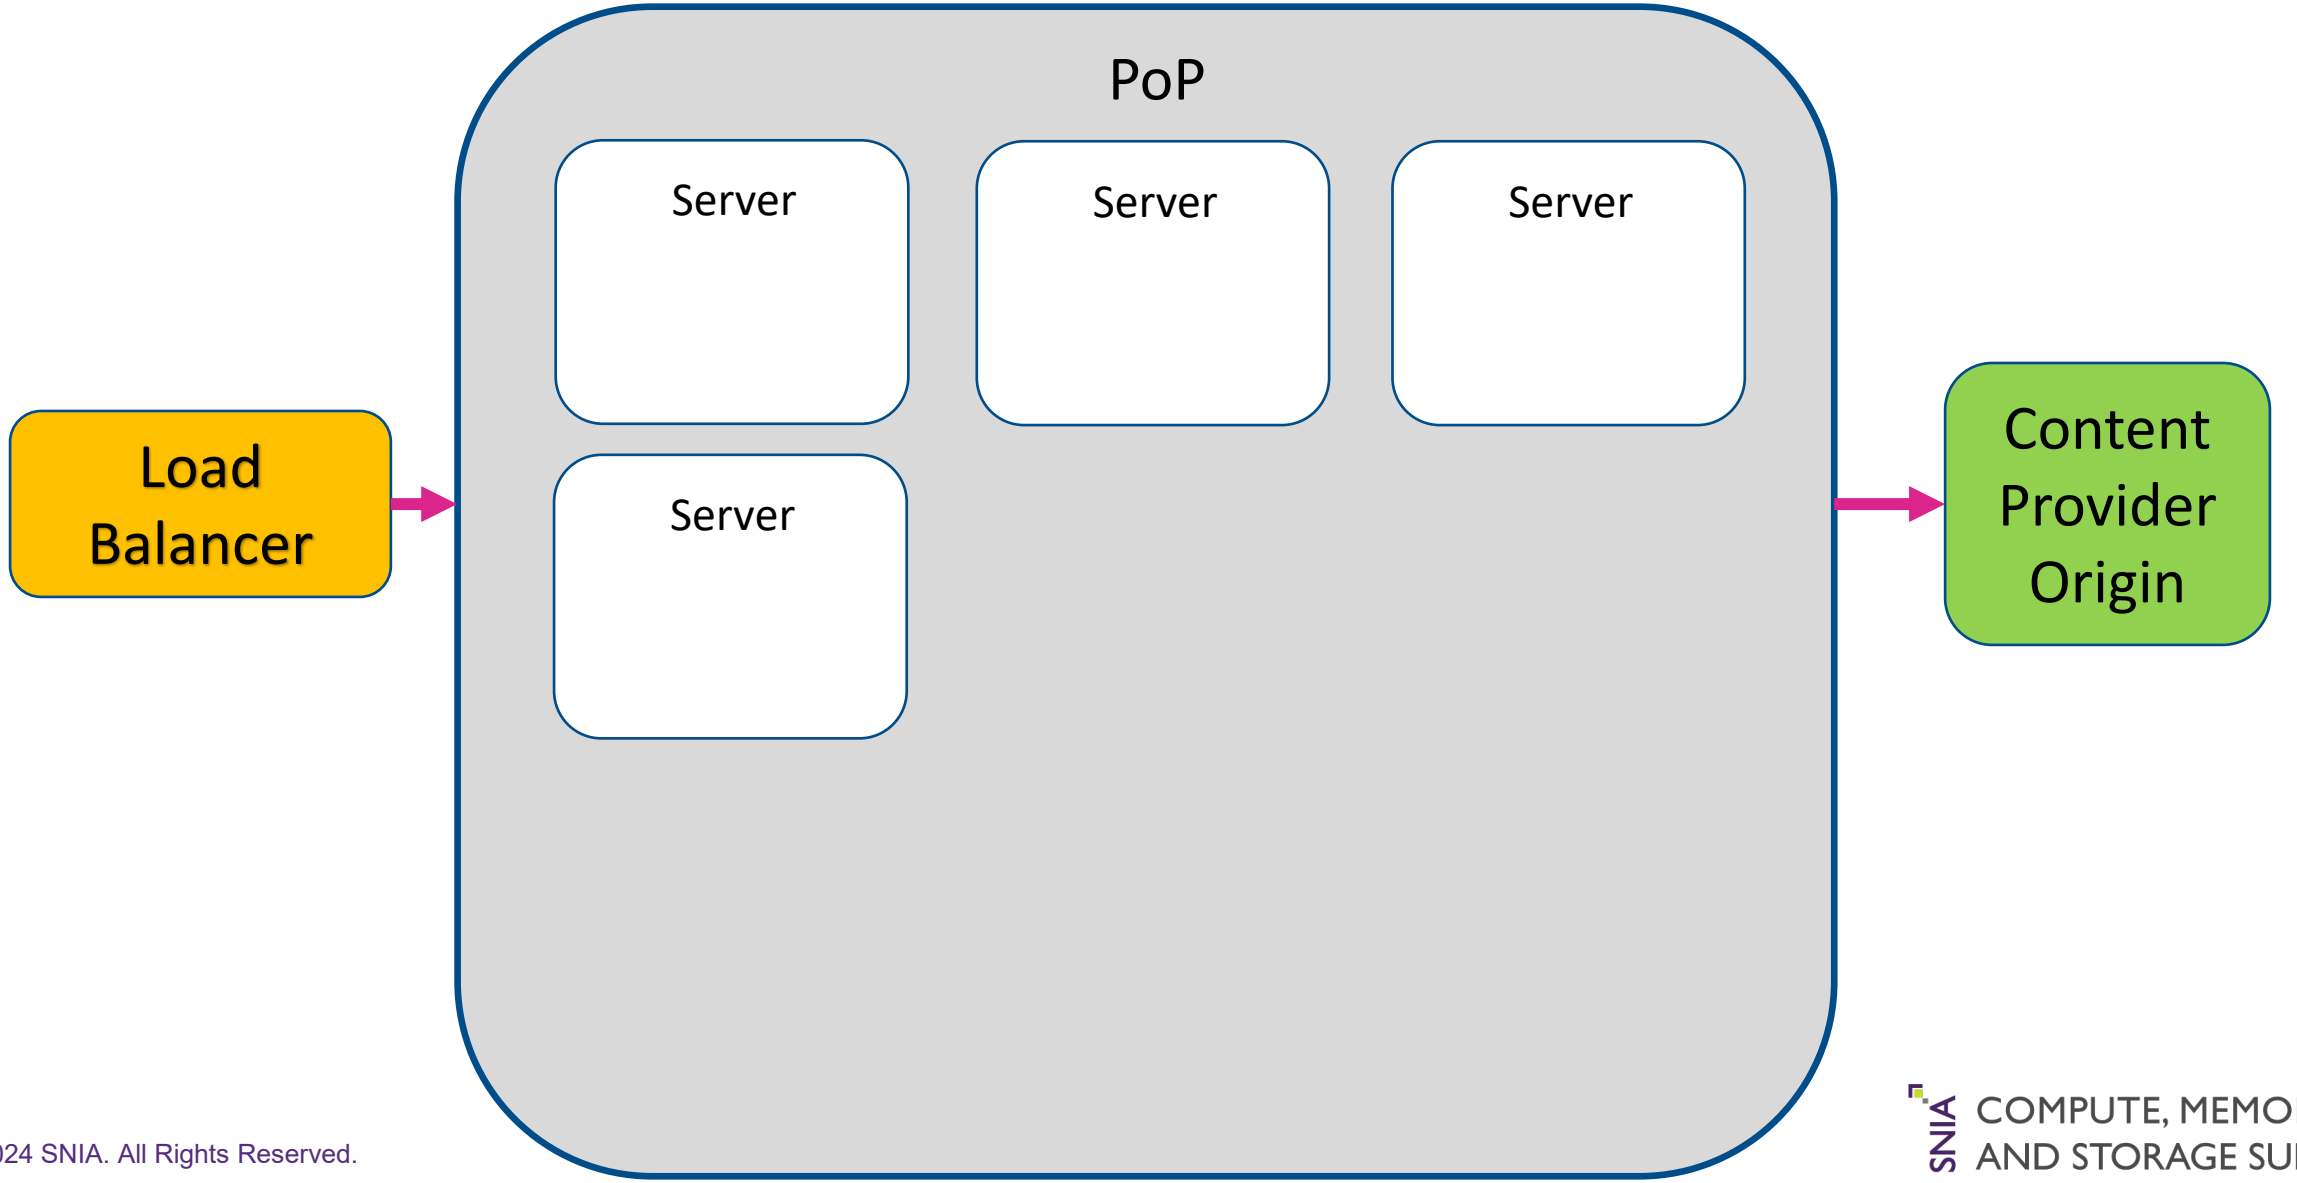

CDN PoP -- Compute

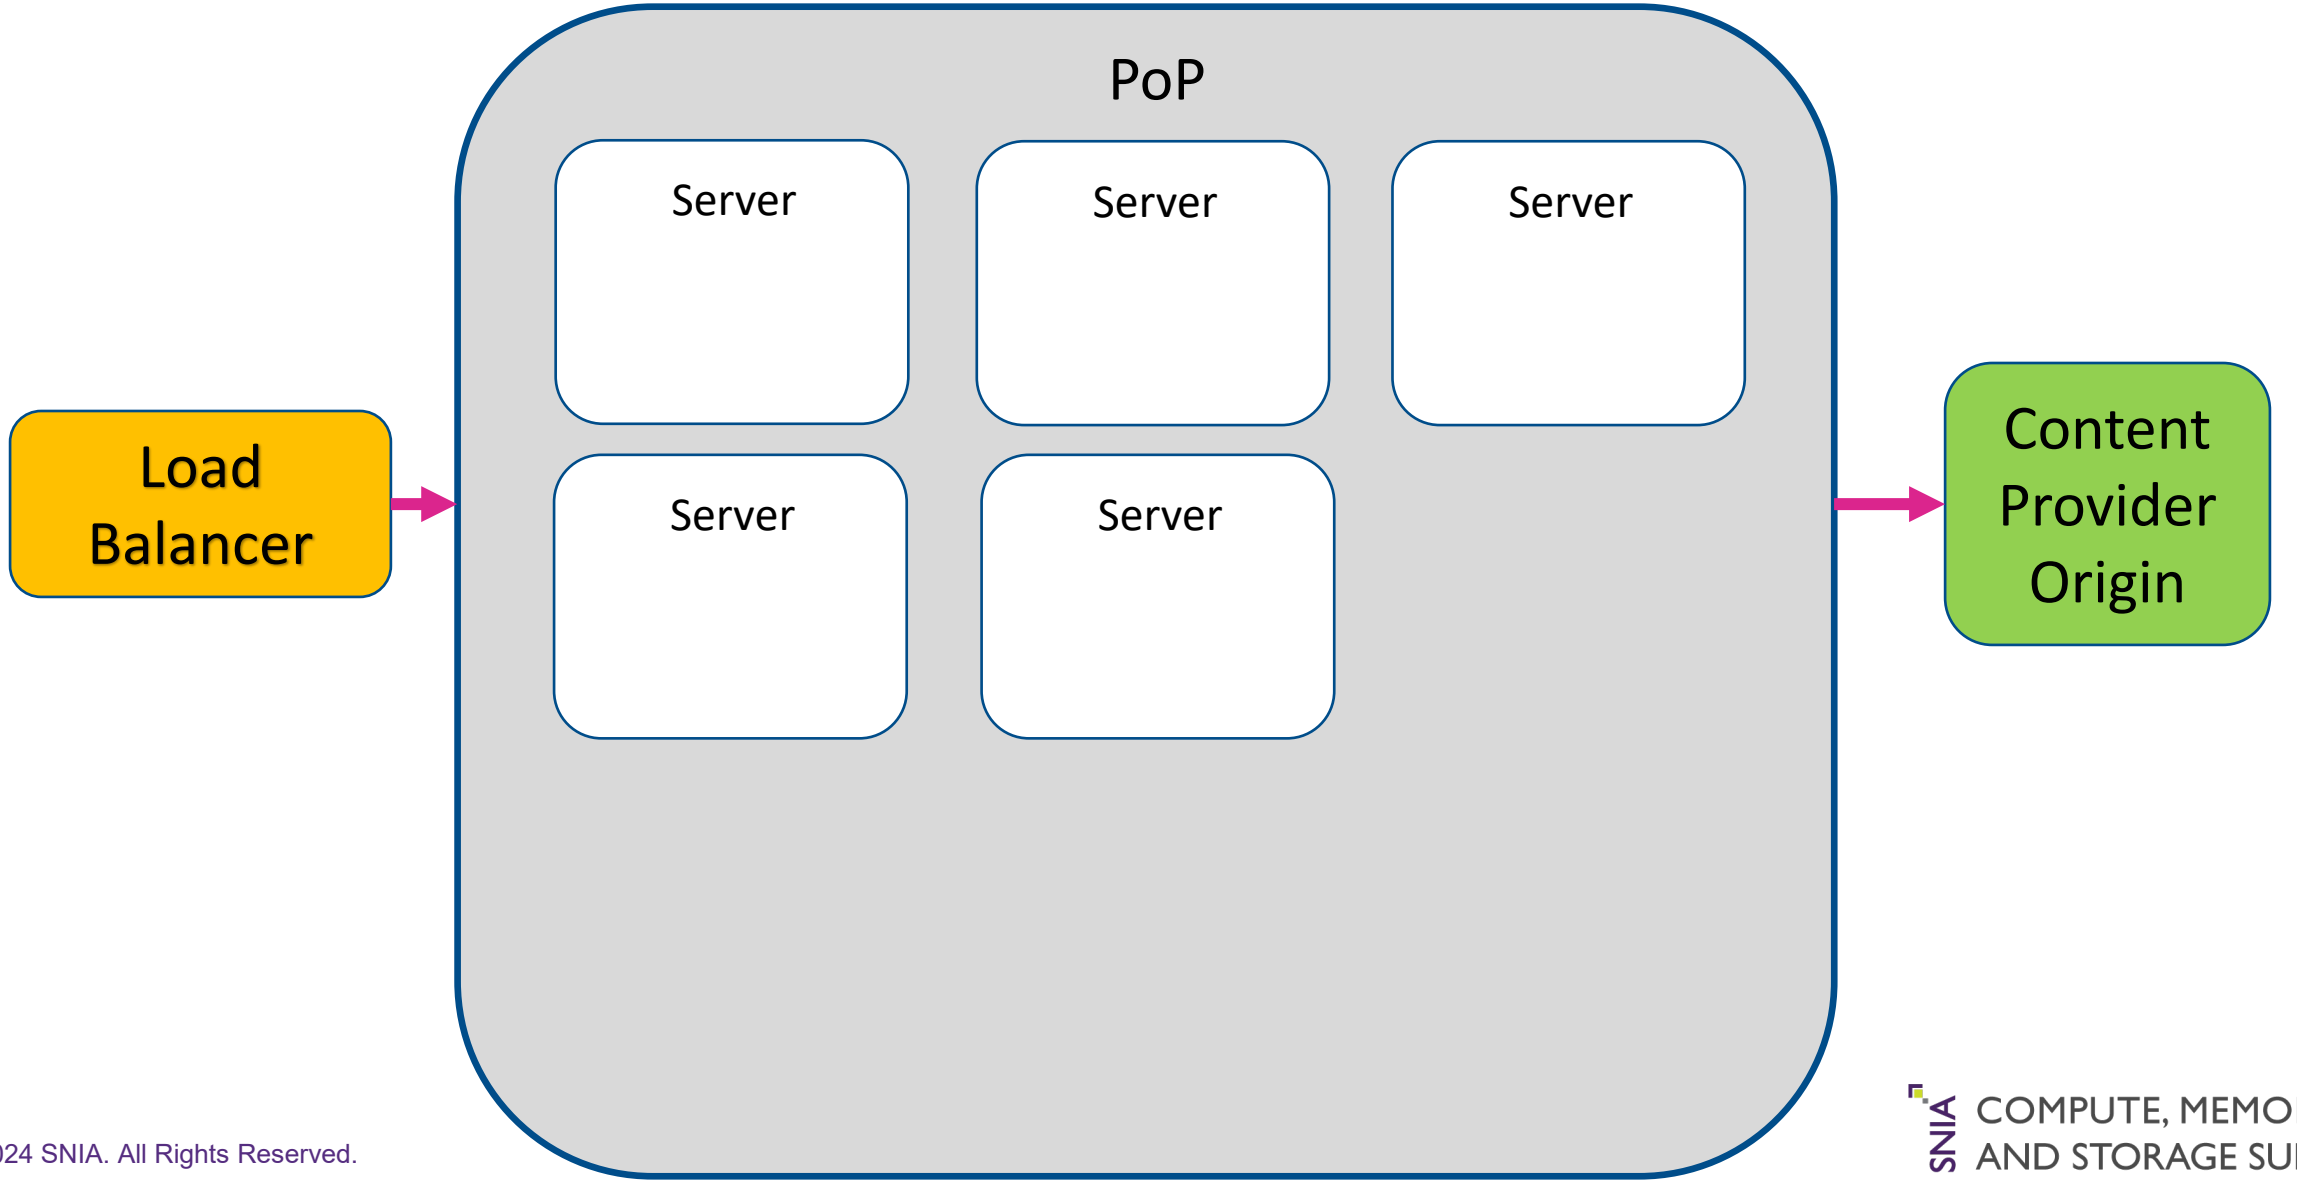

CDN PoP -- Compute

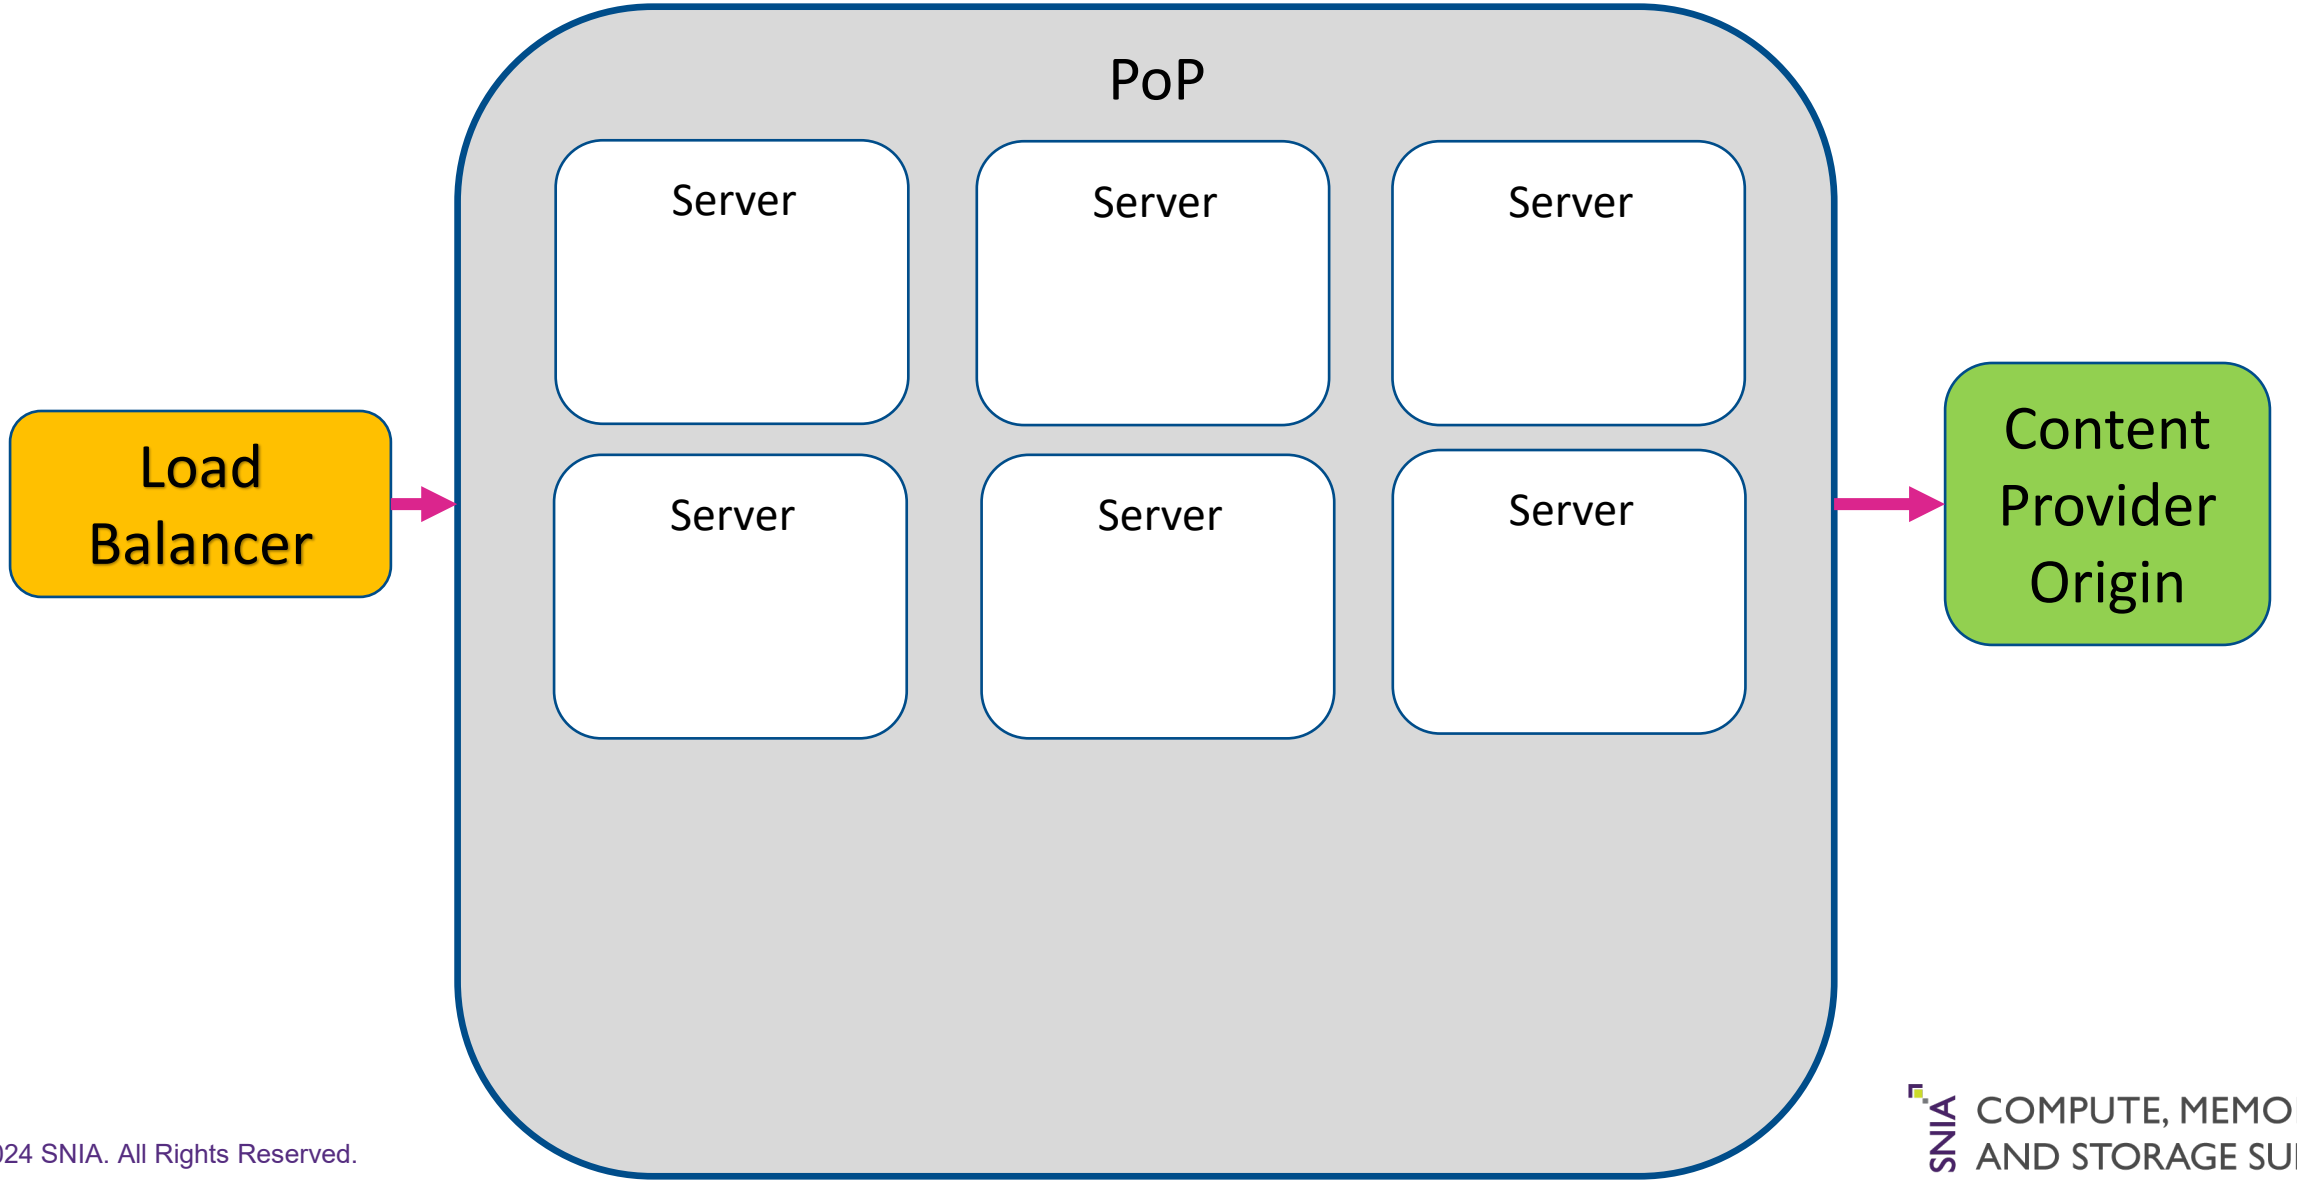

CDN PoP -- Compute

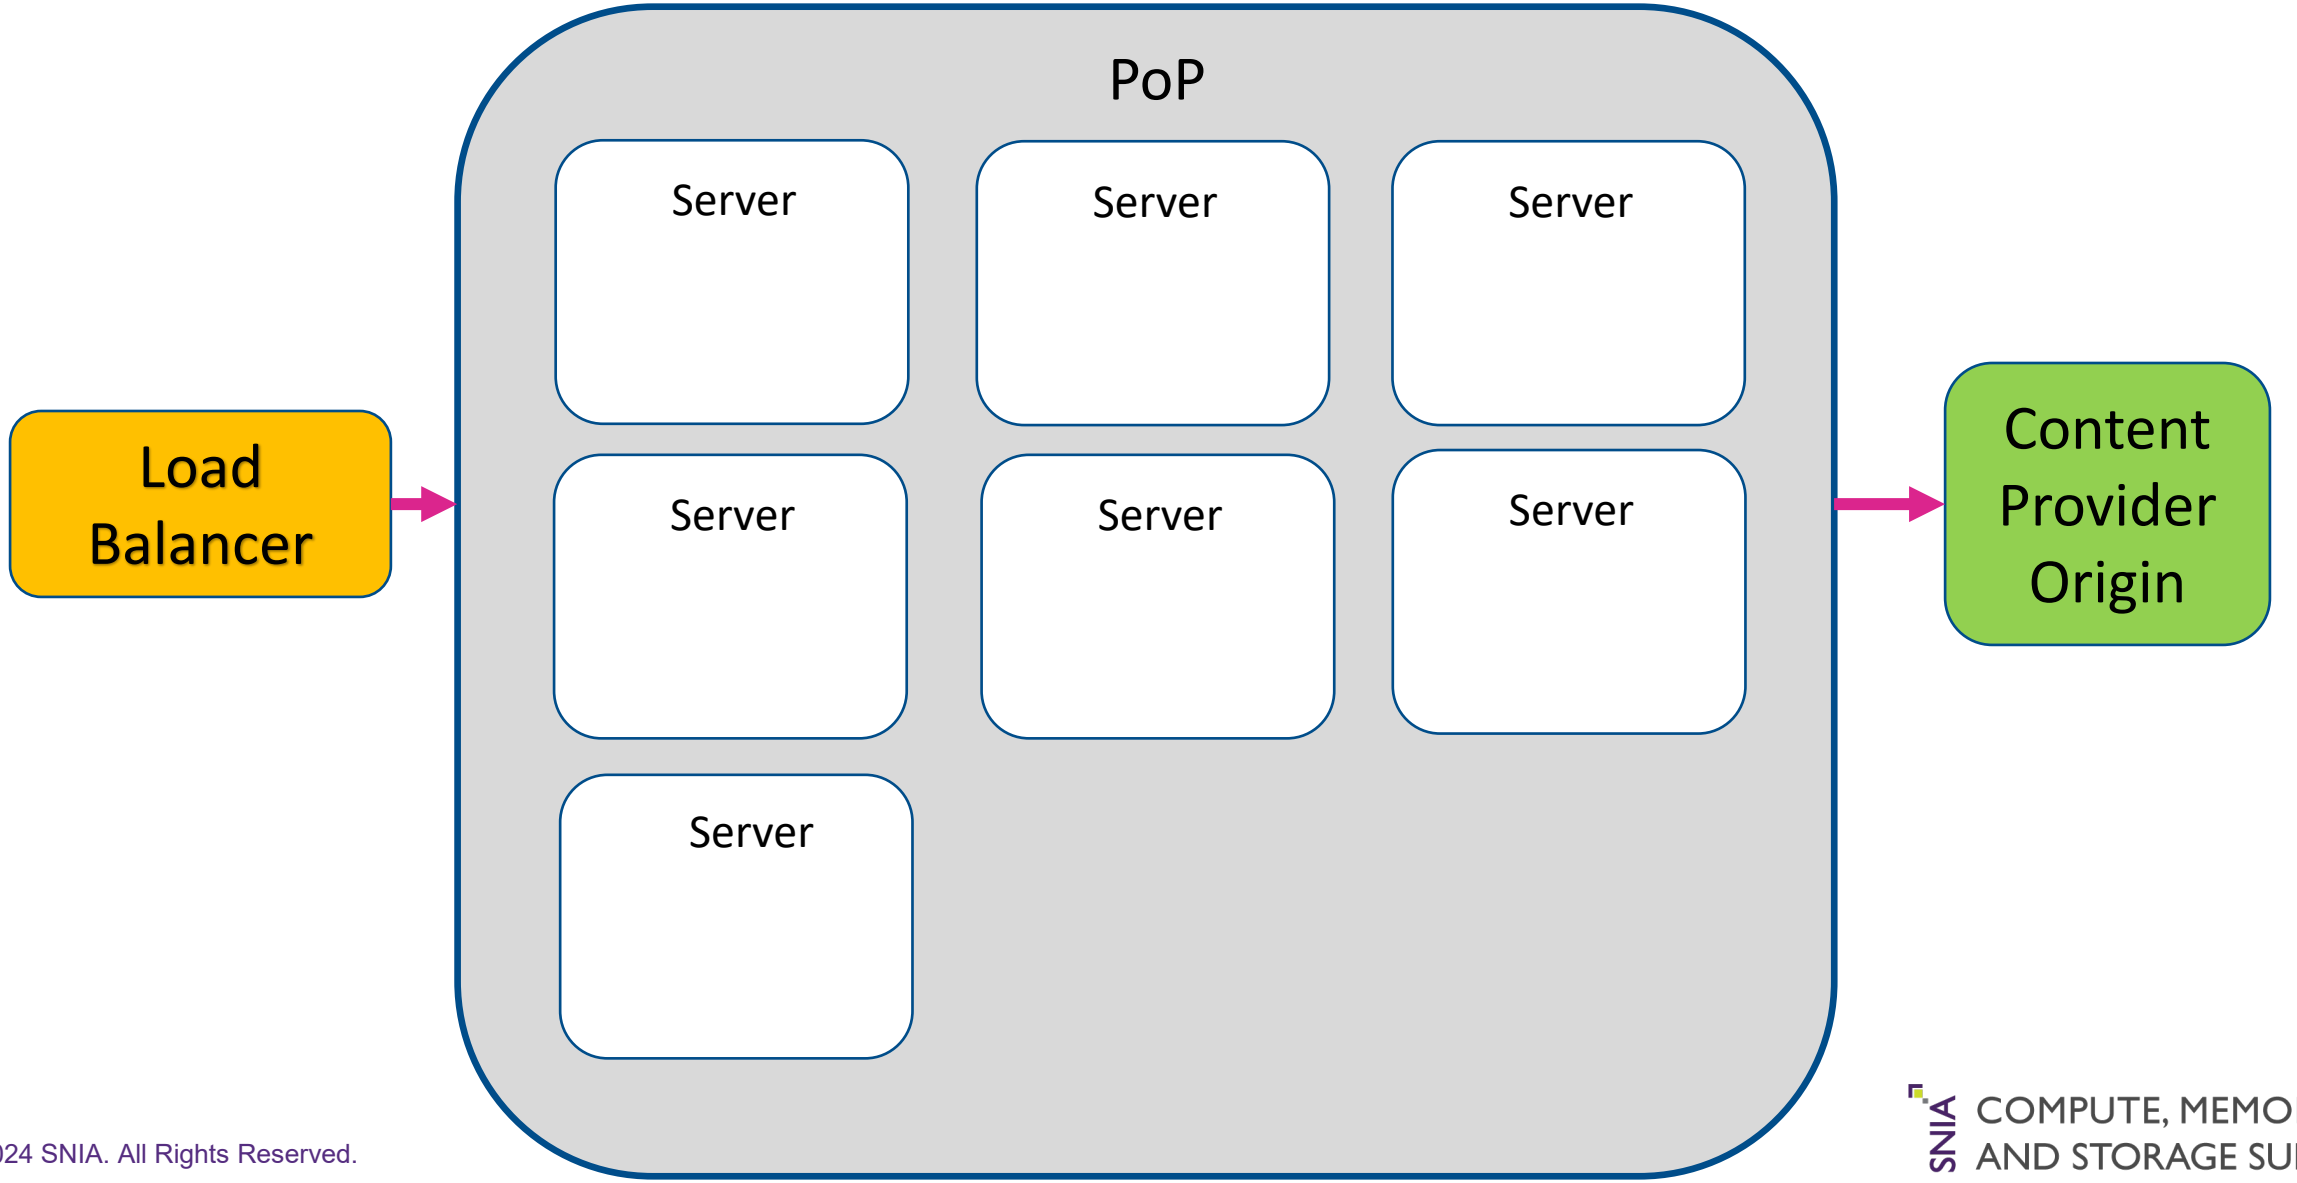

CDN PoP -- Compute

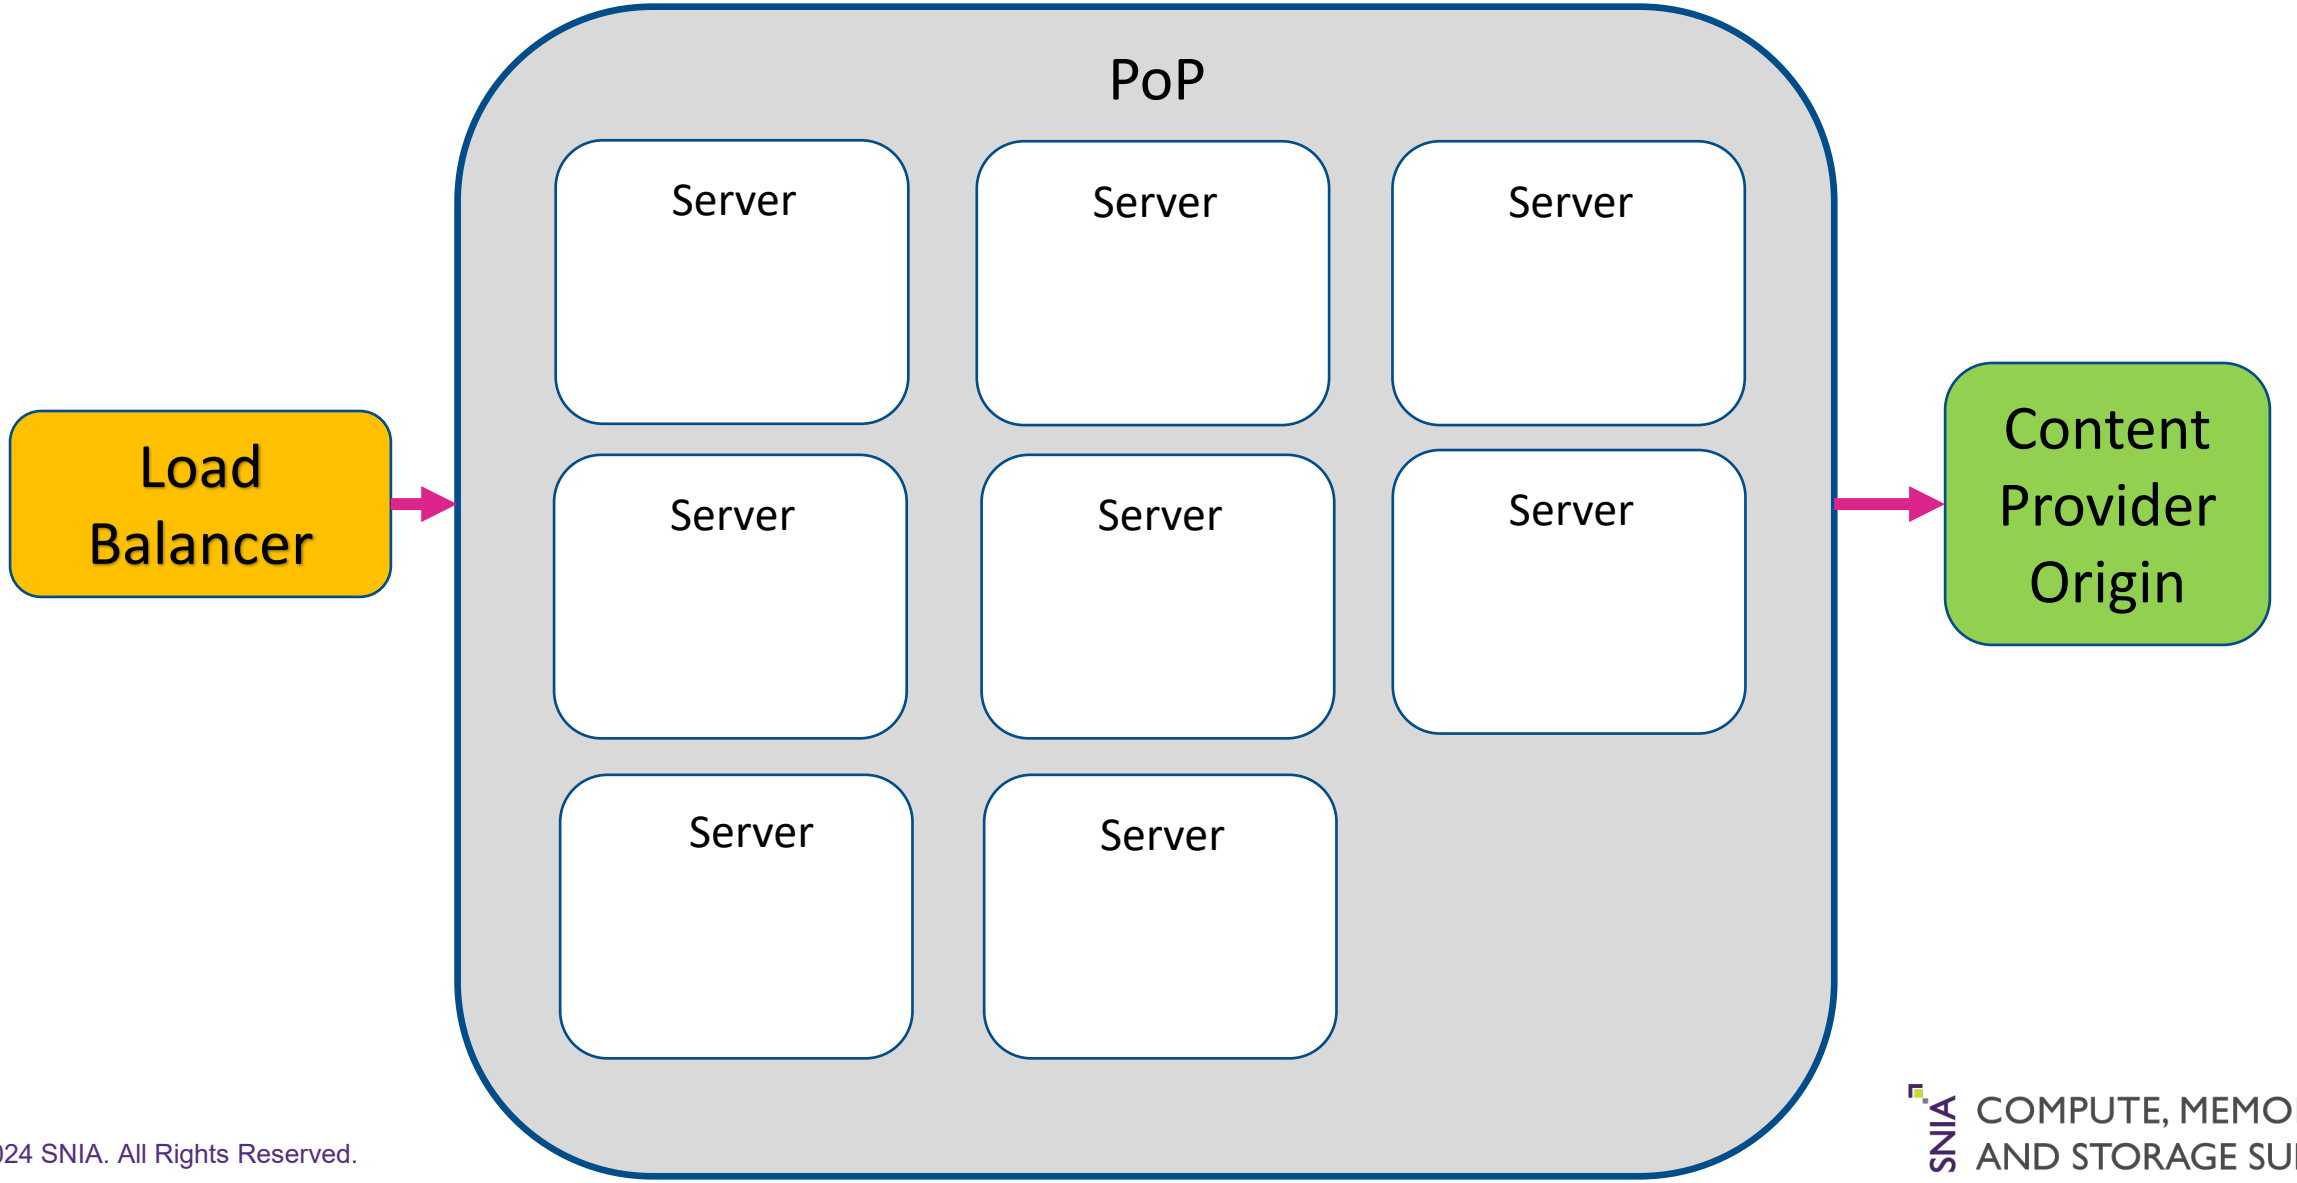

CDN PoP -- Compute

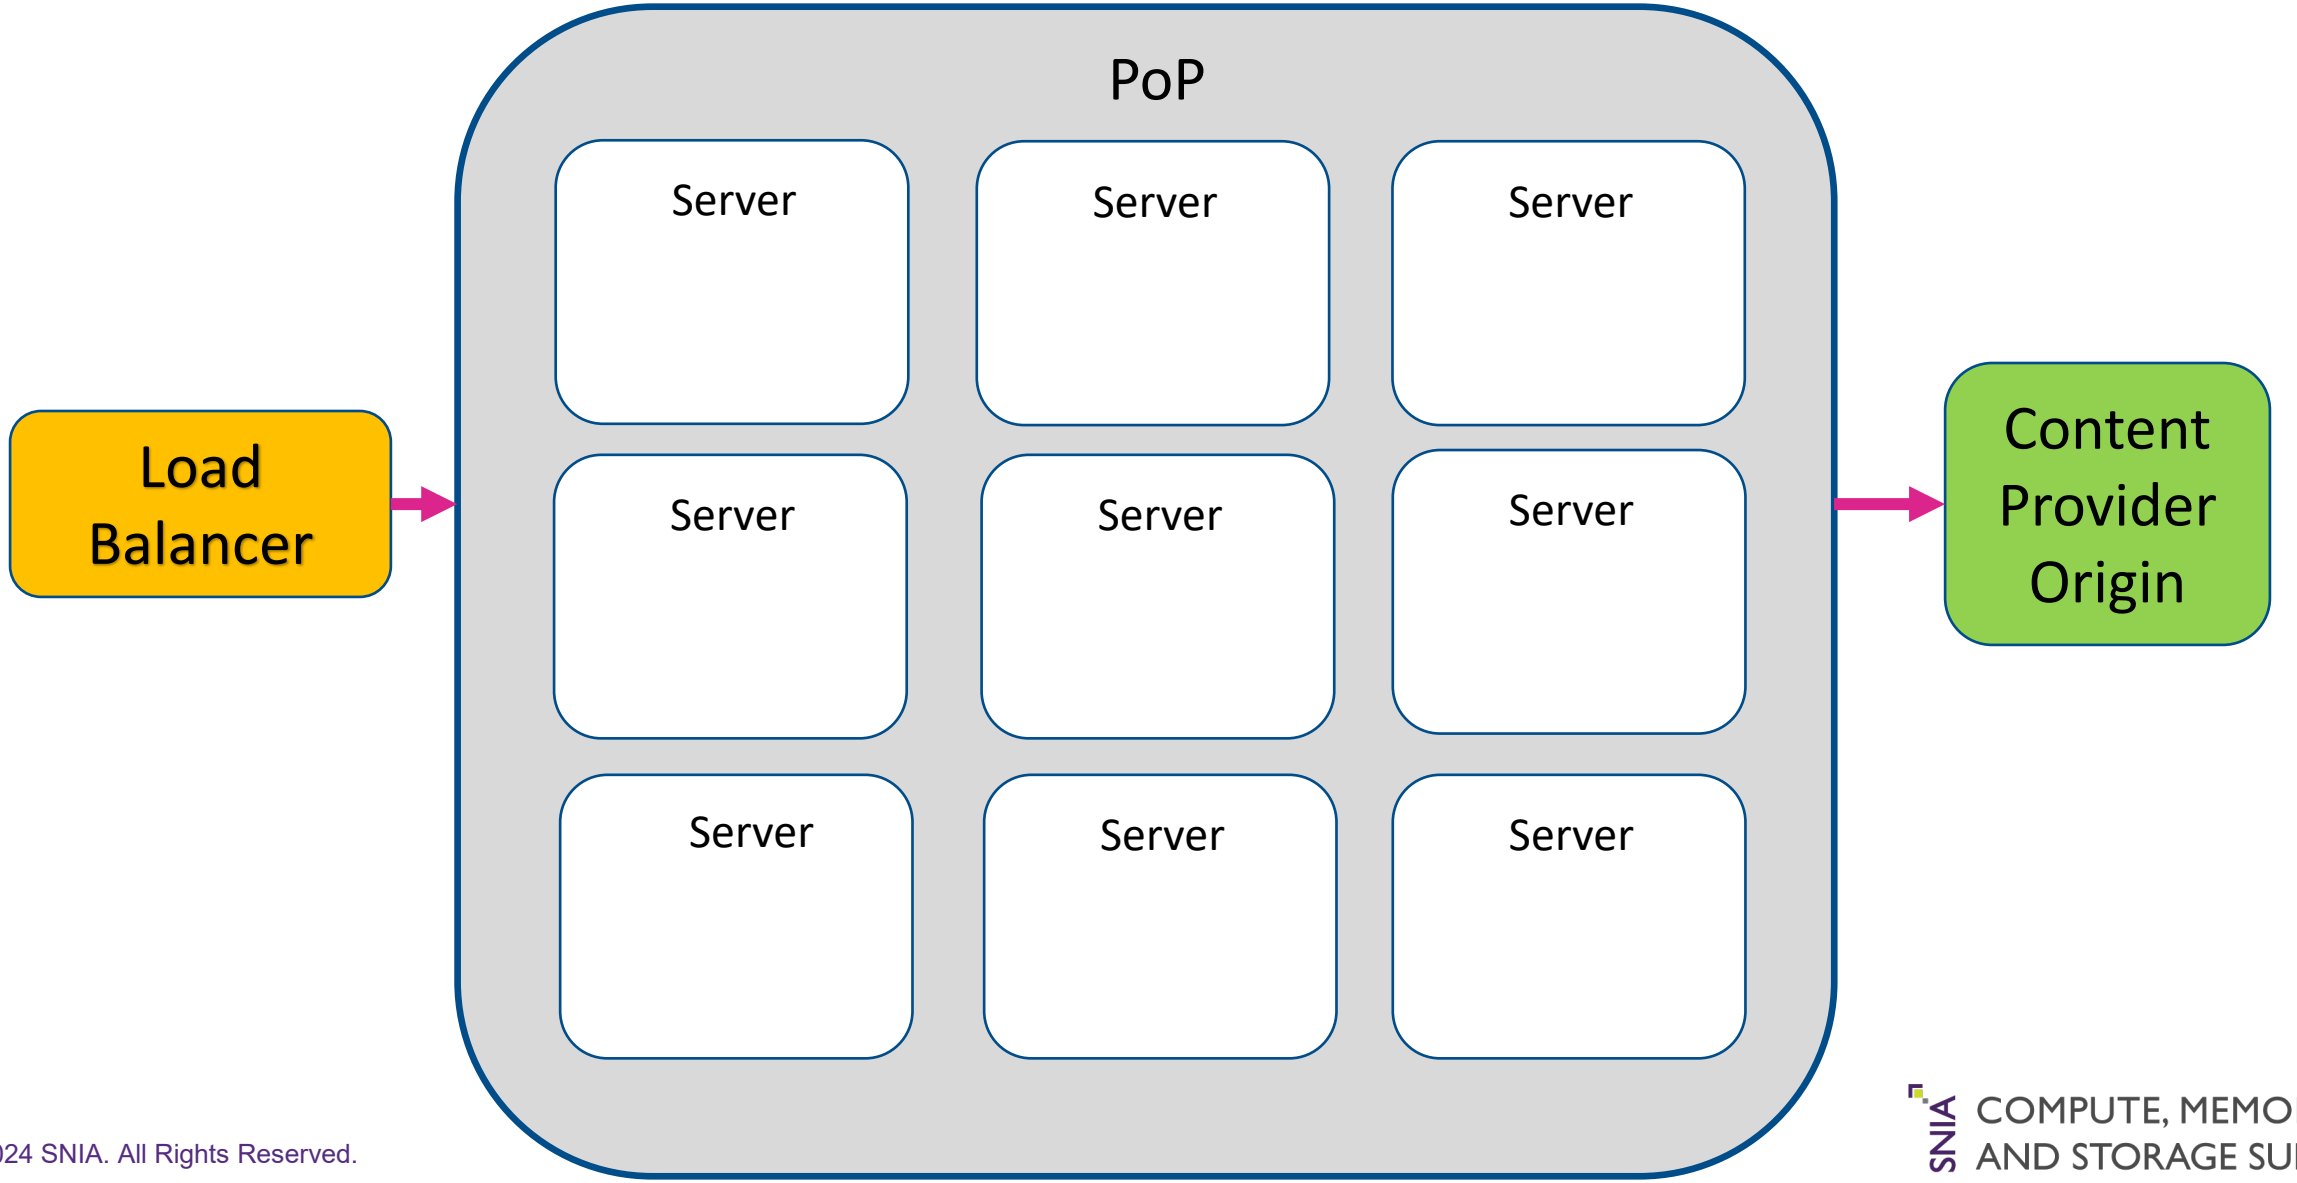

CDN PoP – Compute + Memory

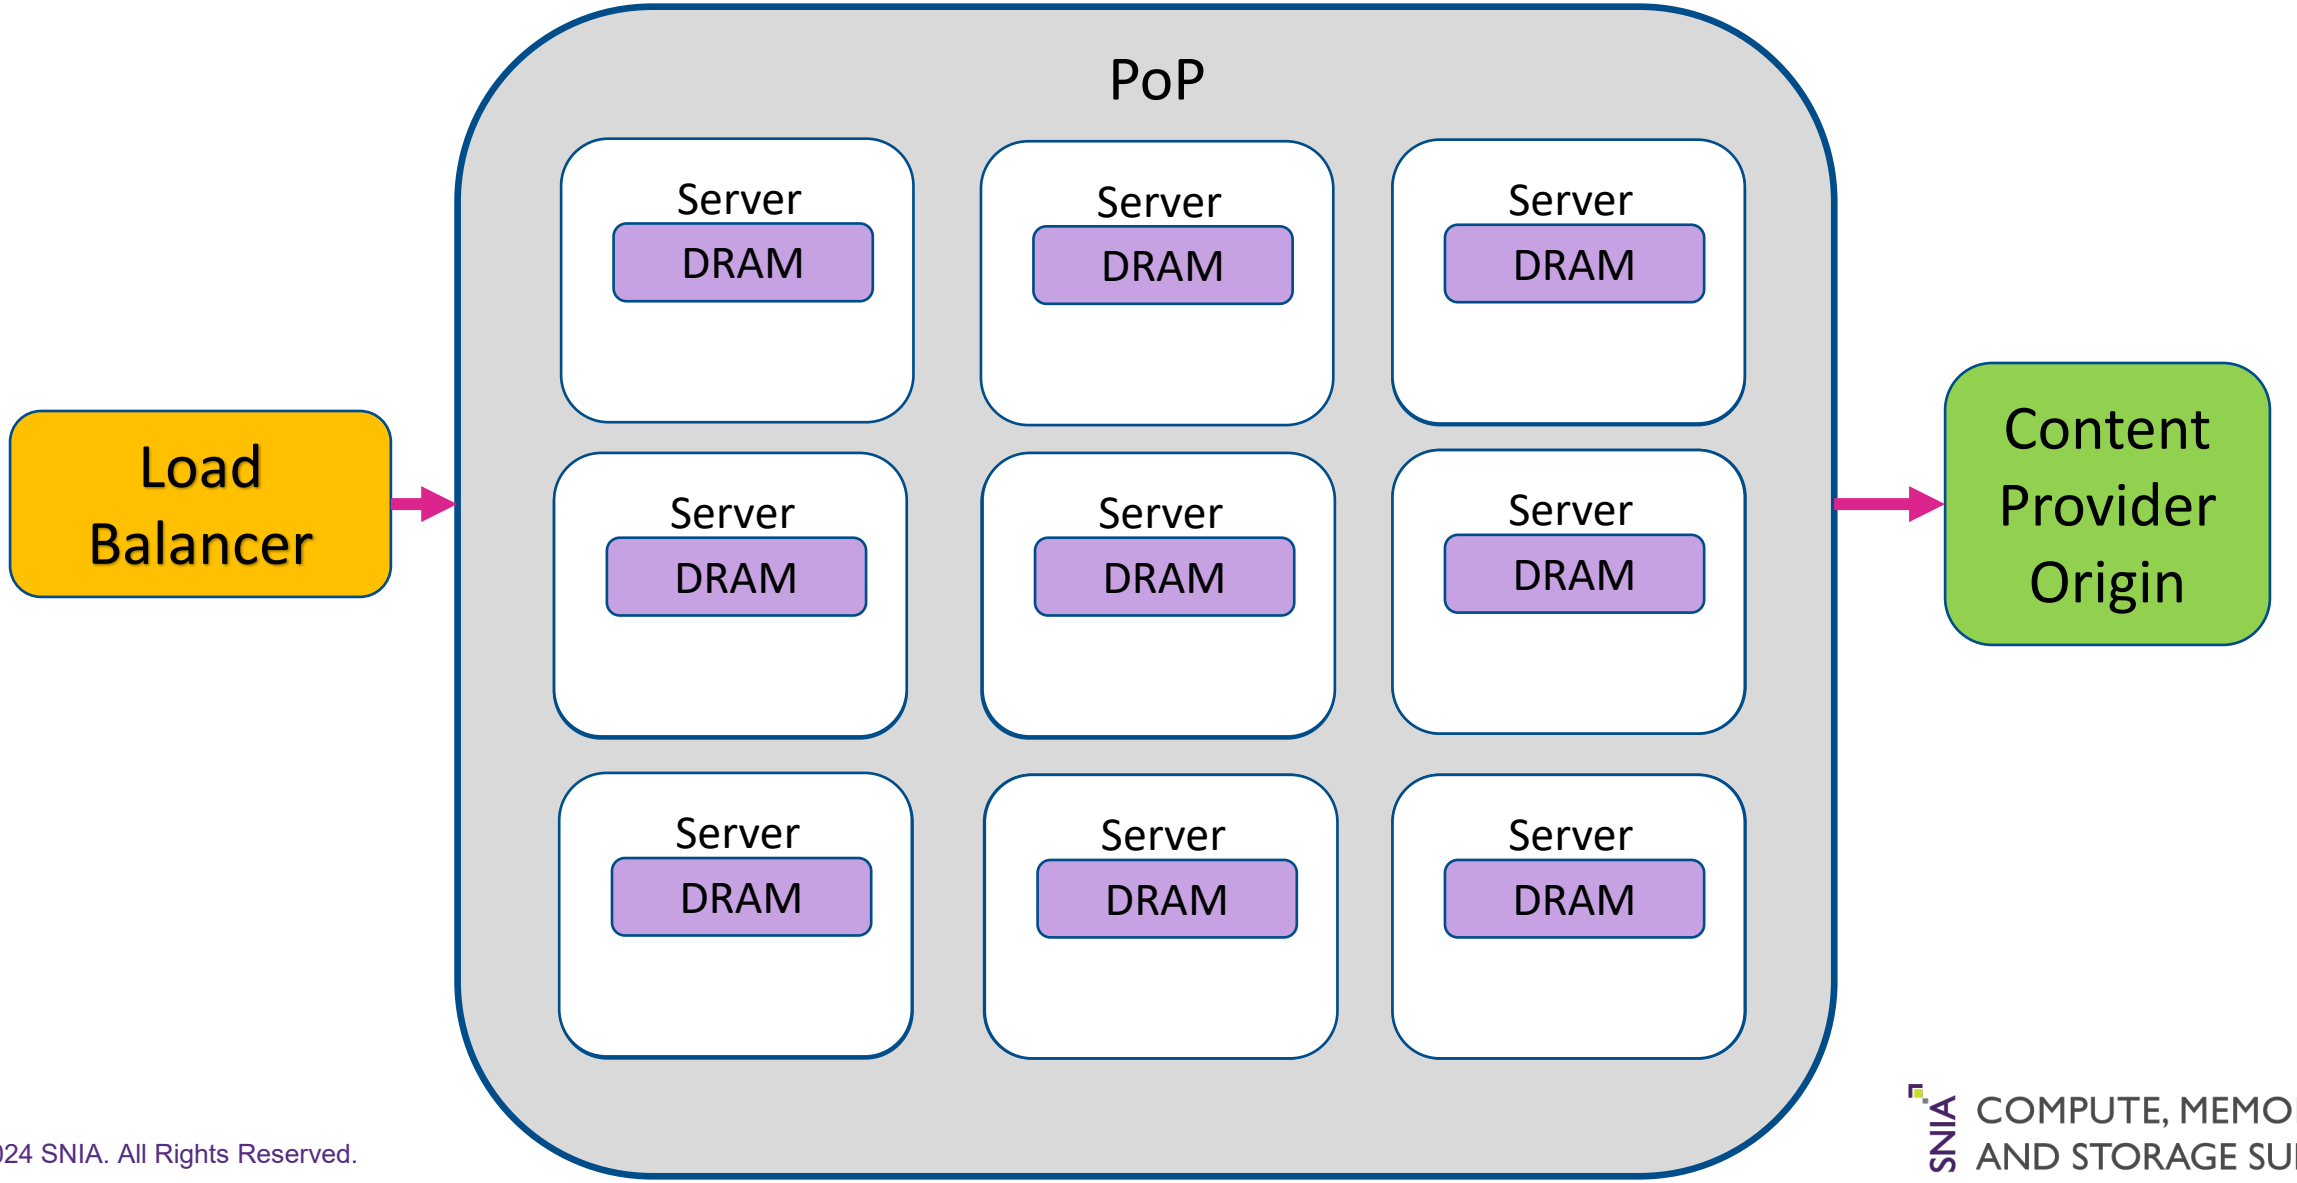

CDN PoP – Compute + Memory + Storage

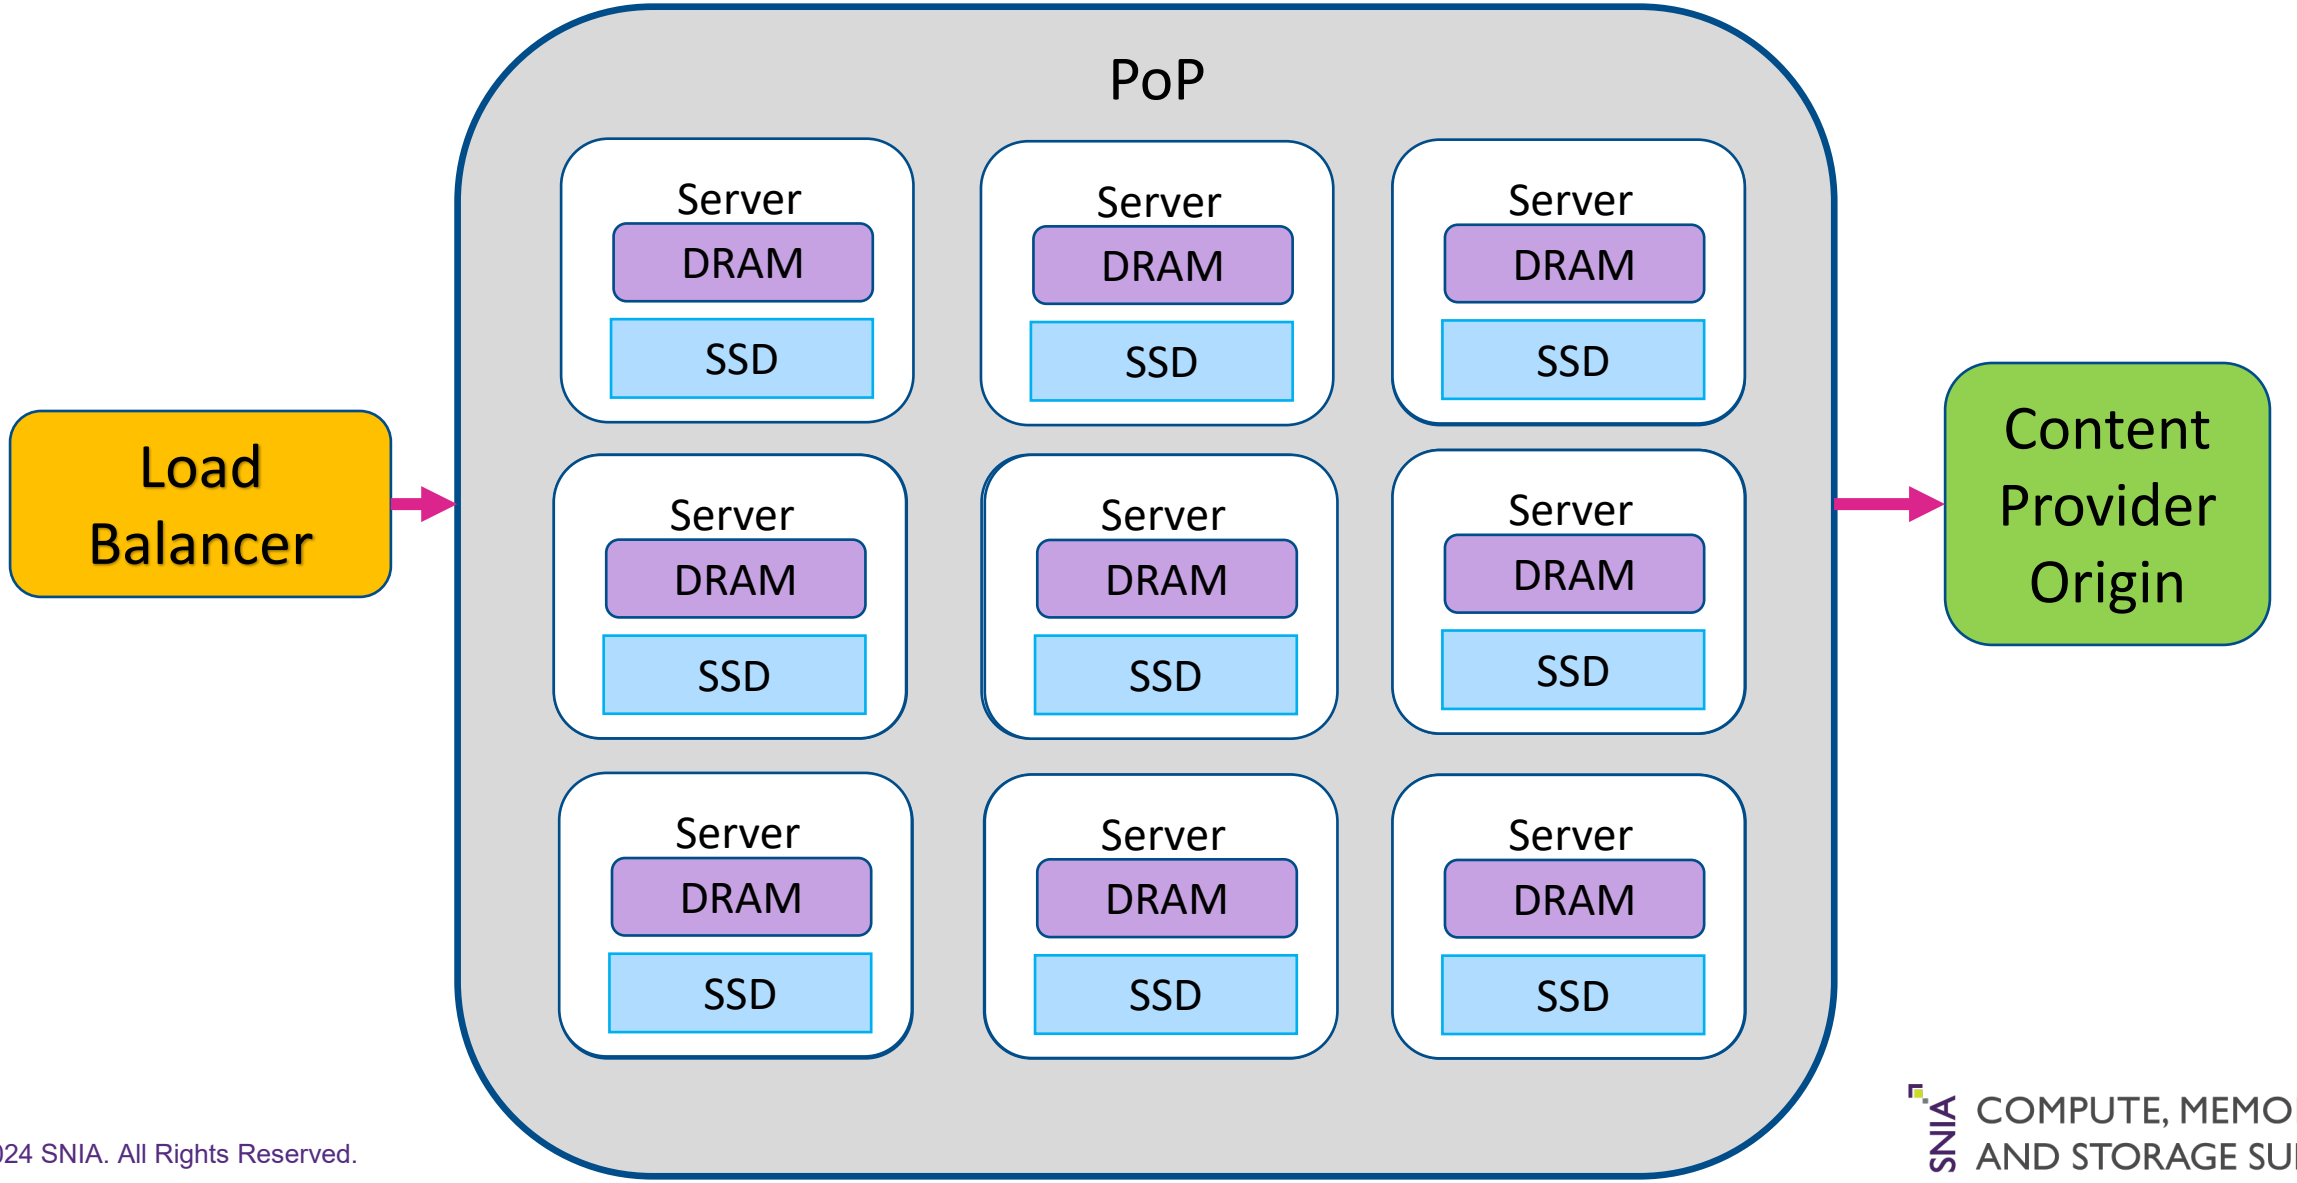

CDN PoP – Compute + Memory + Storage

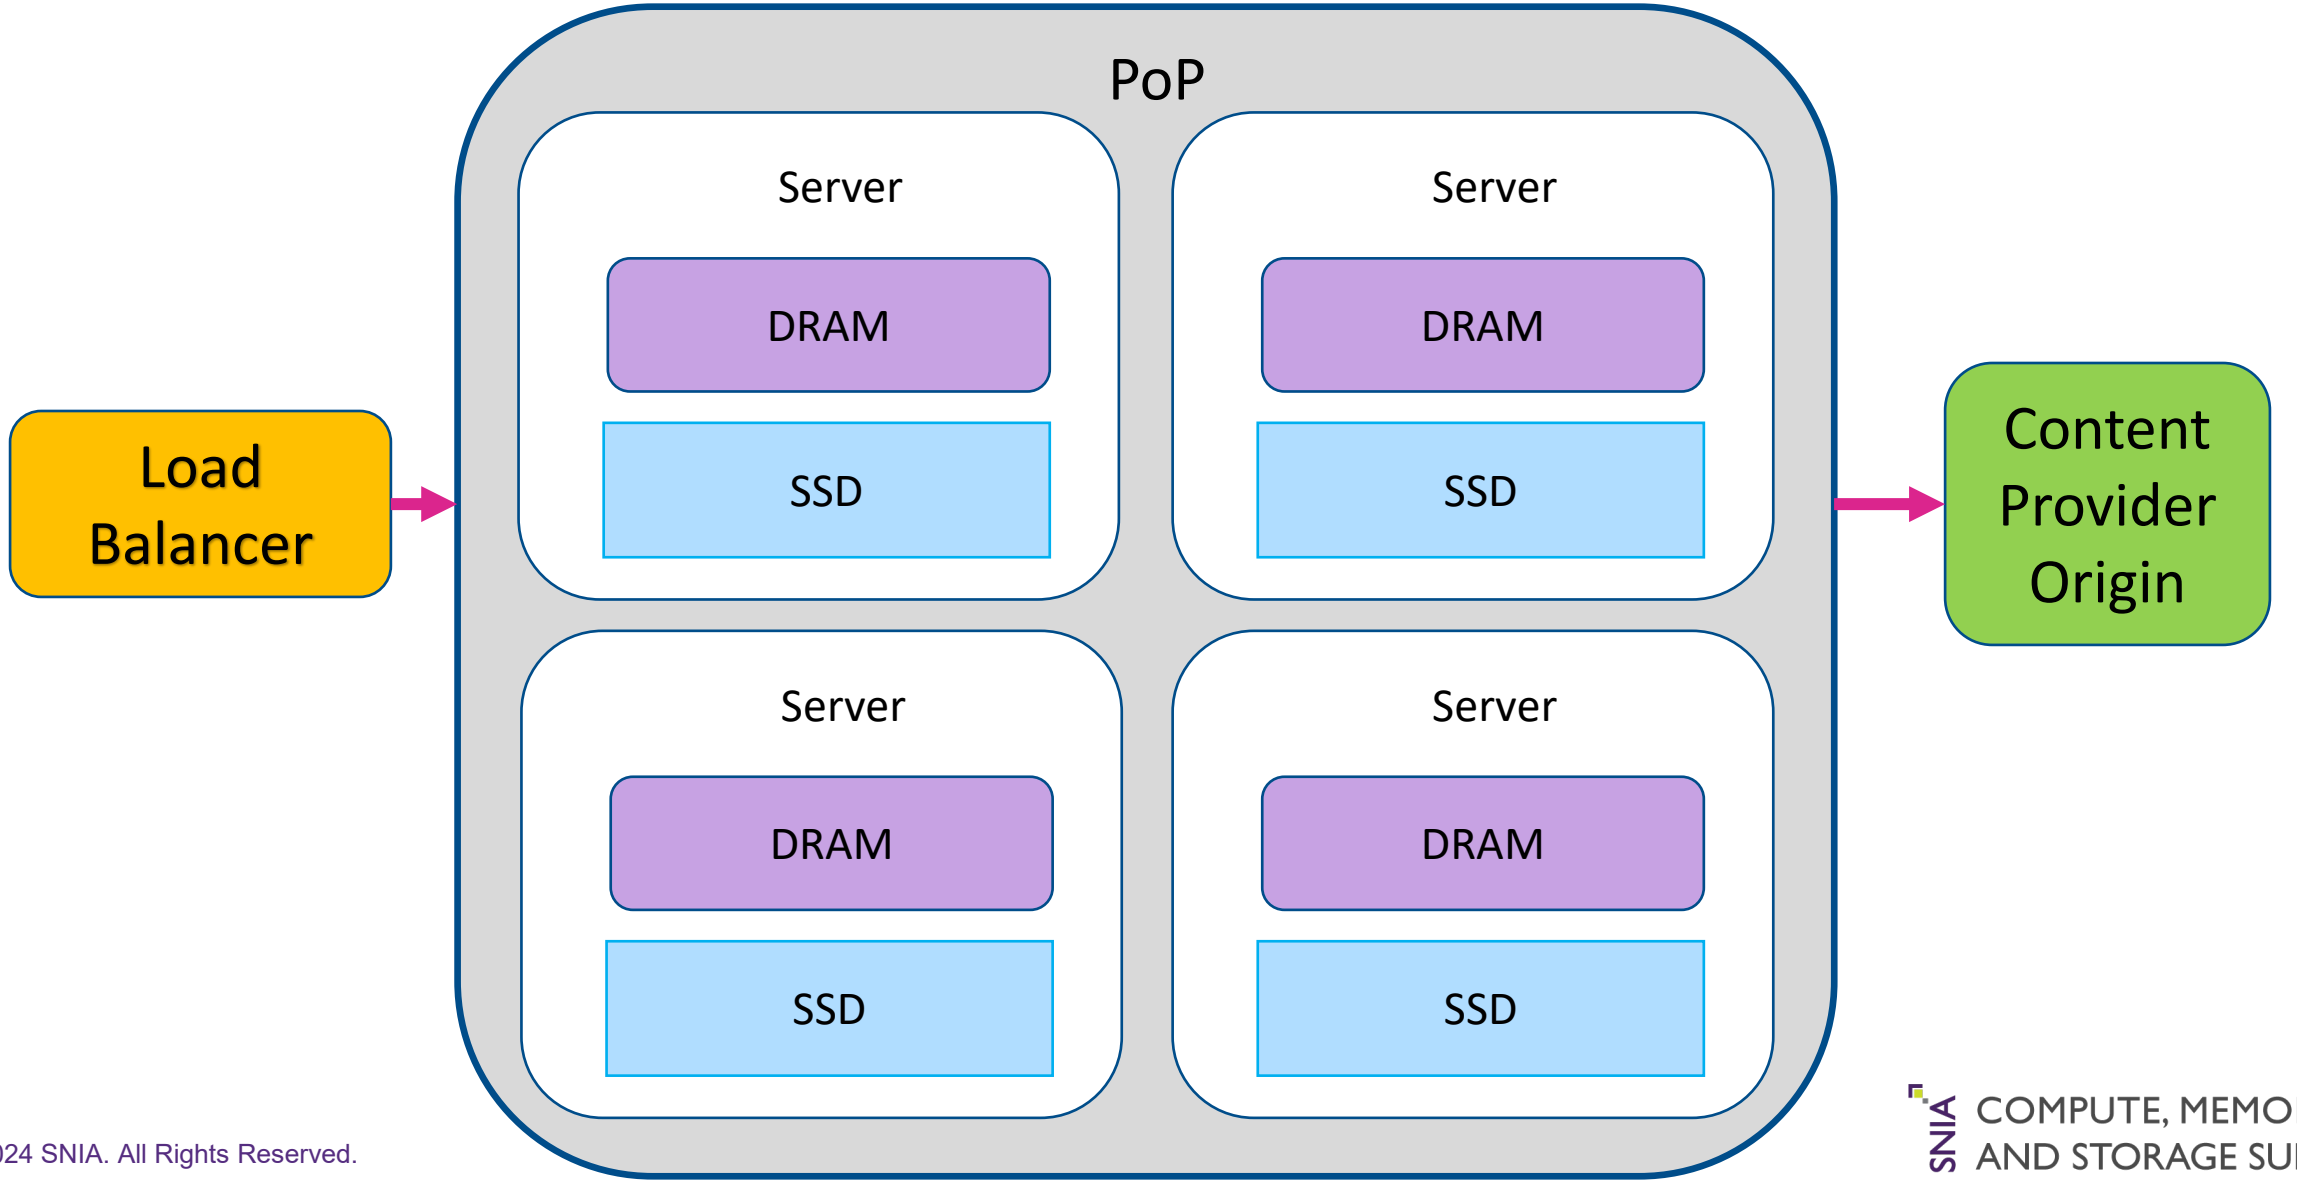

CDN PoP – Compute + Memory + Storage

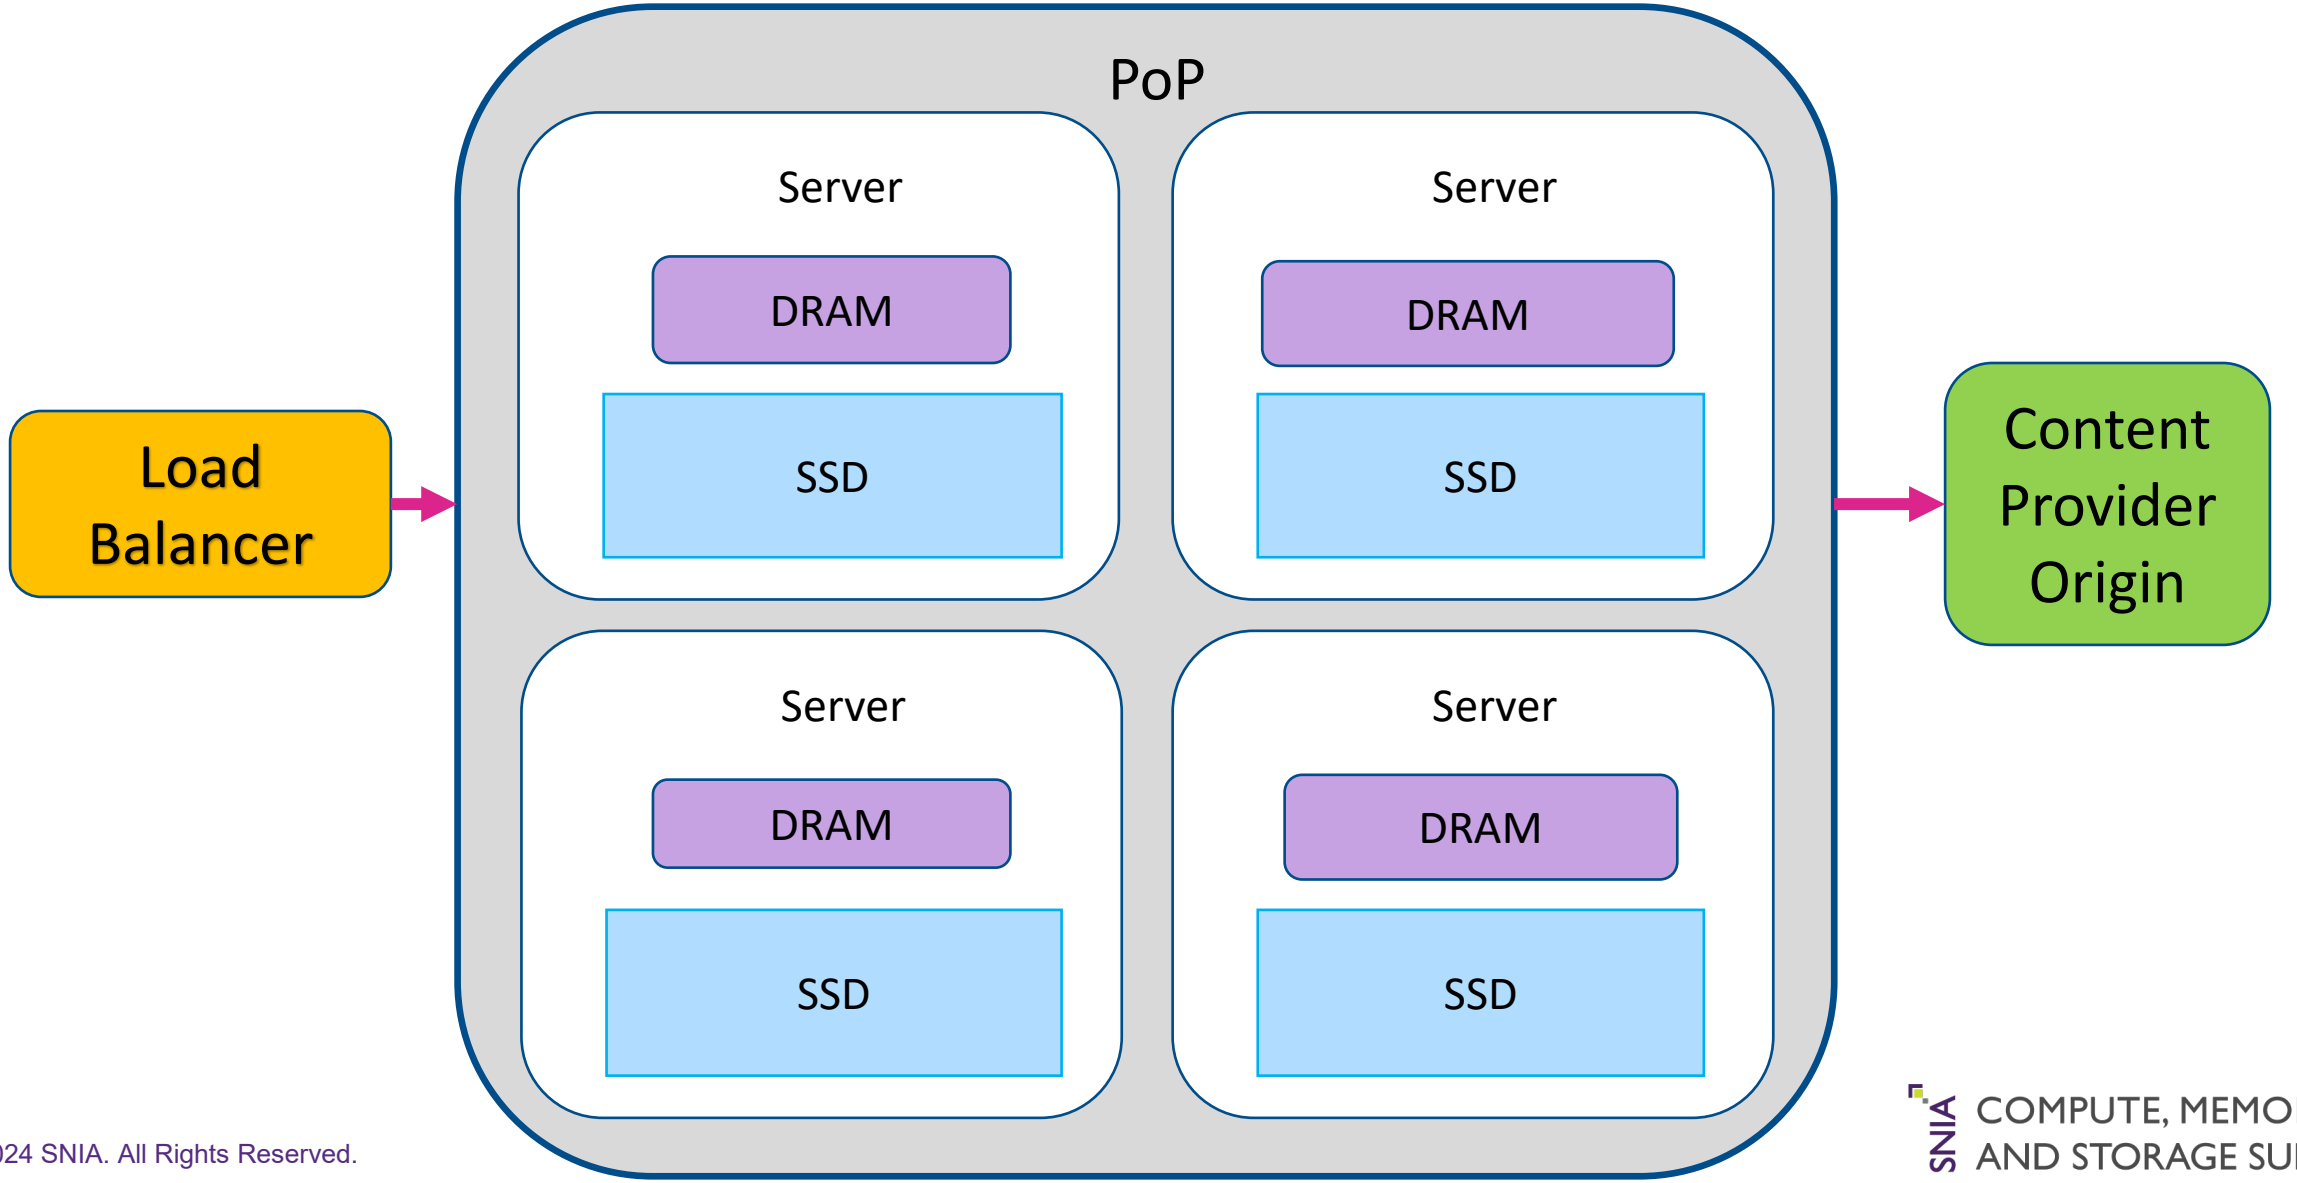

CDN PoP – Compute + Memory + Storage

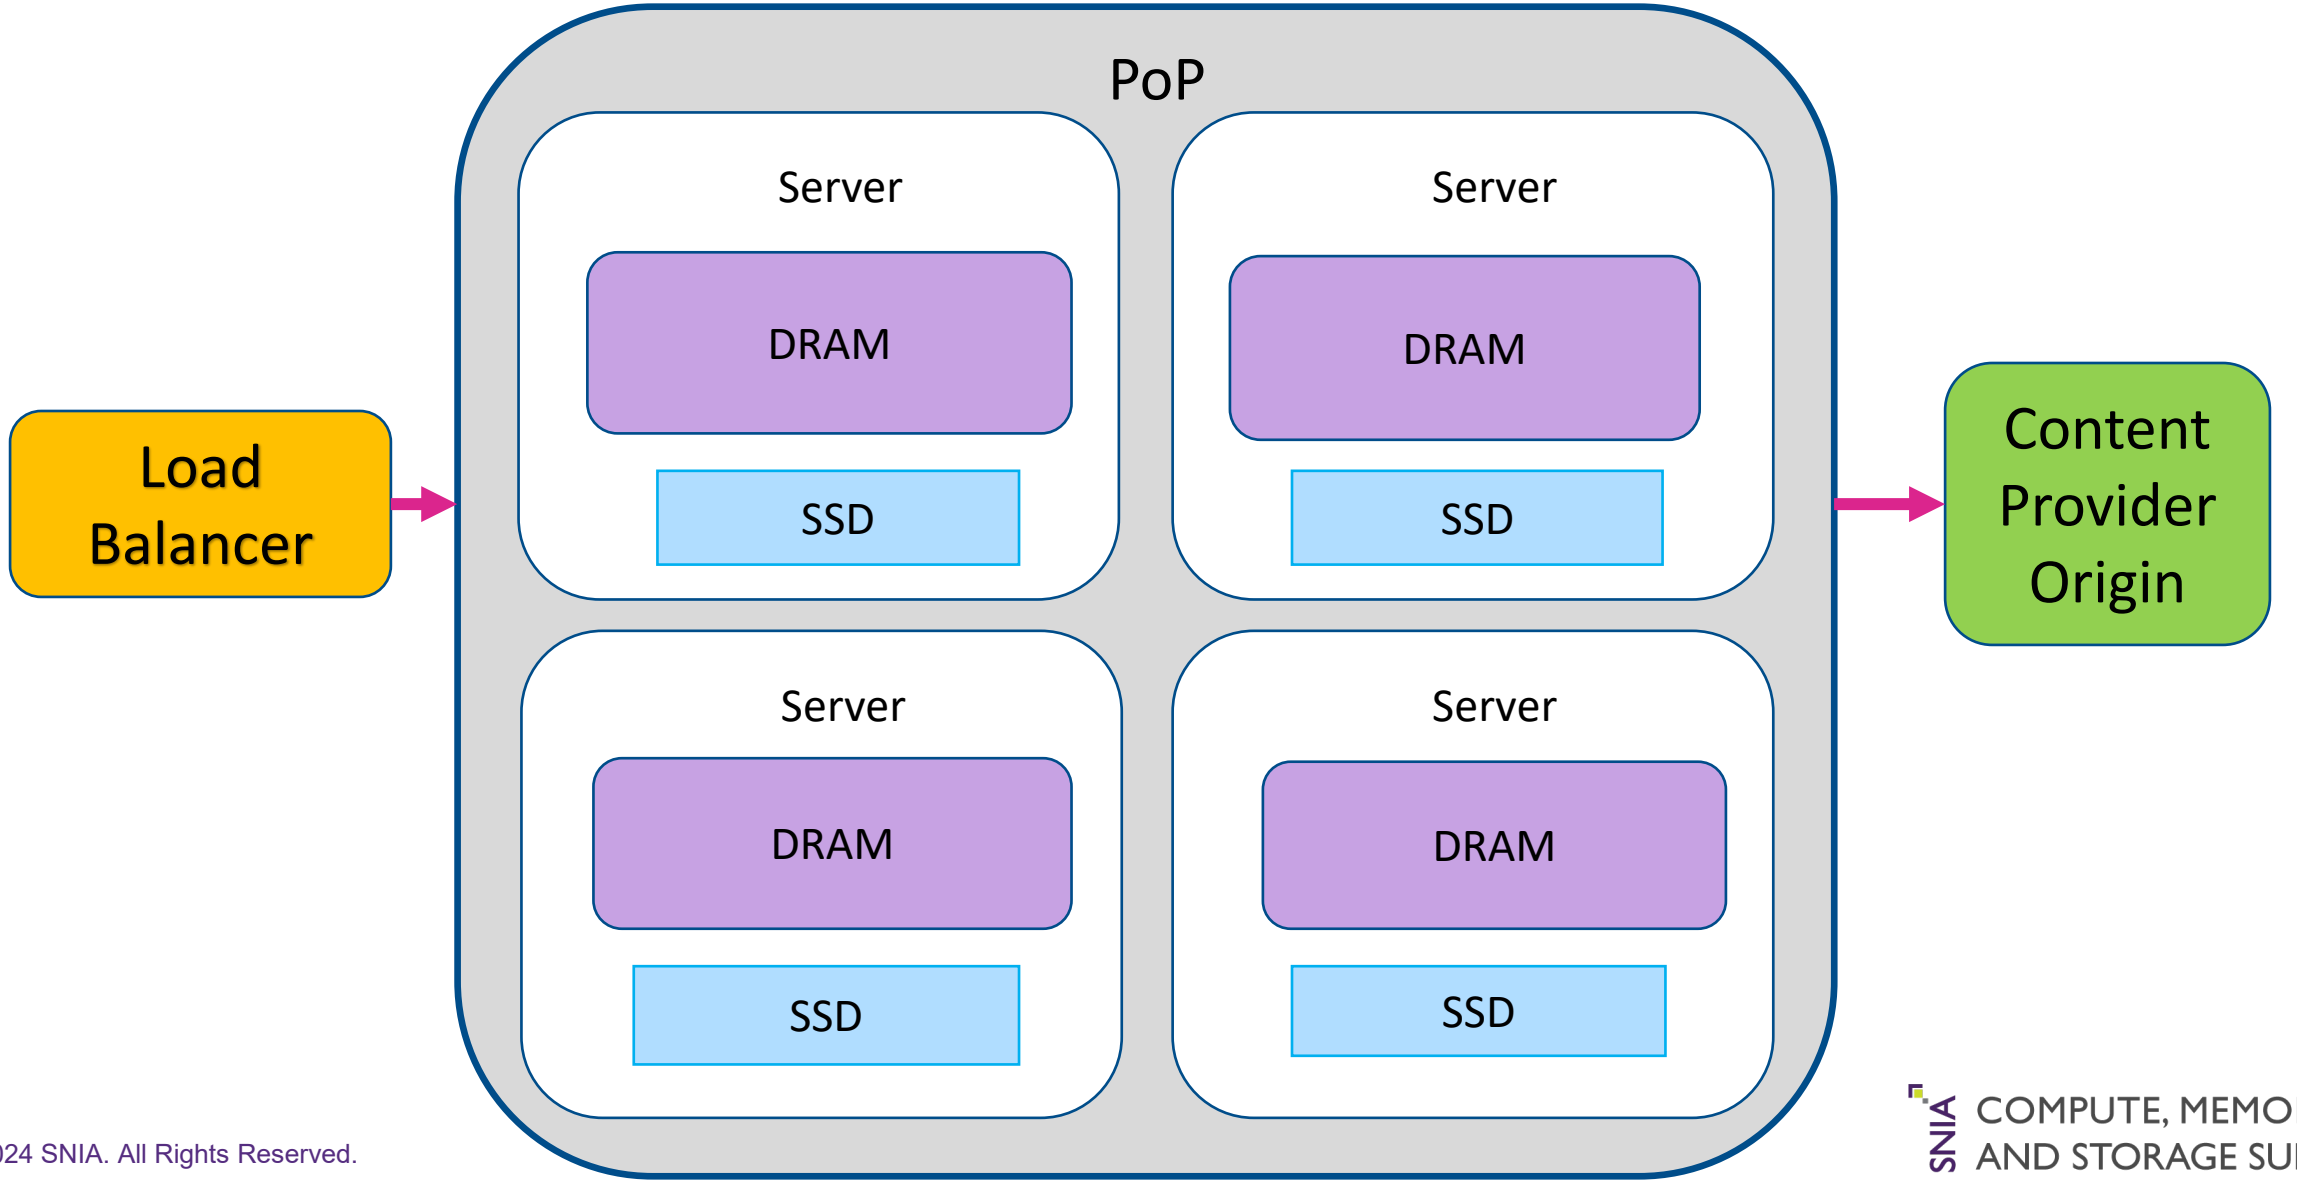

CDN PoP – Compute + Memory + Storage

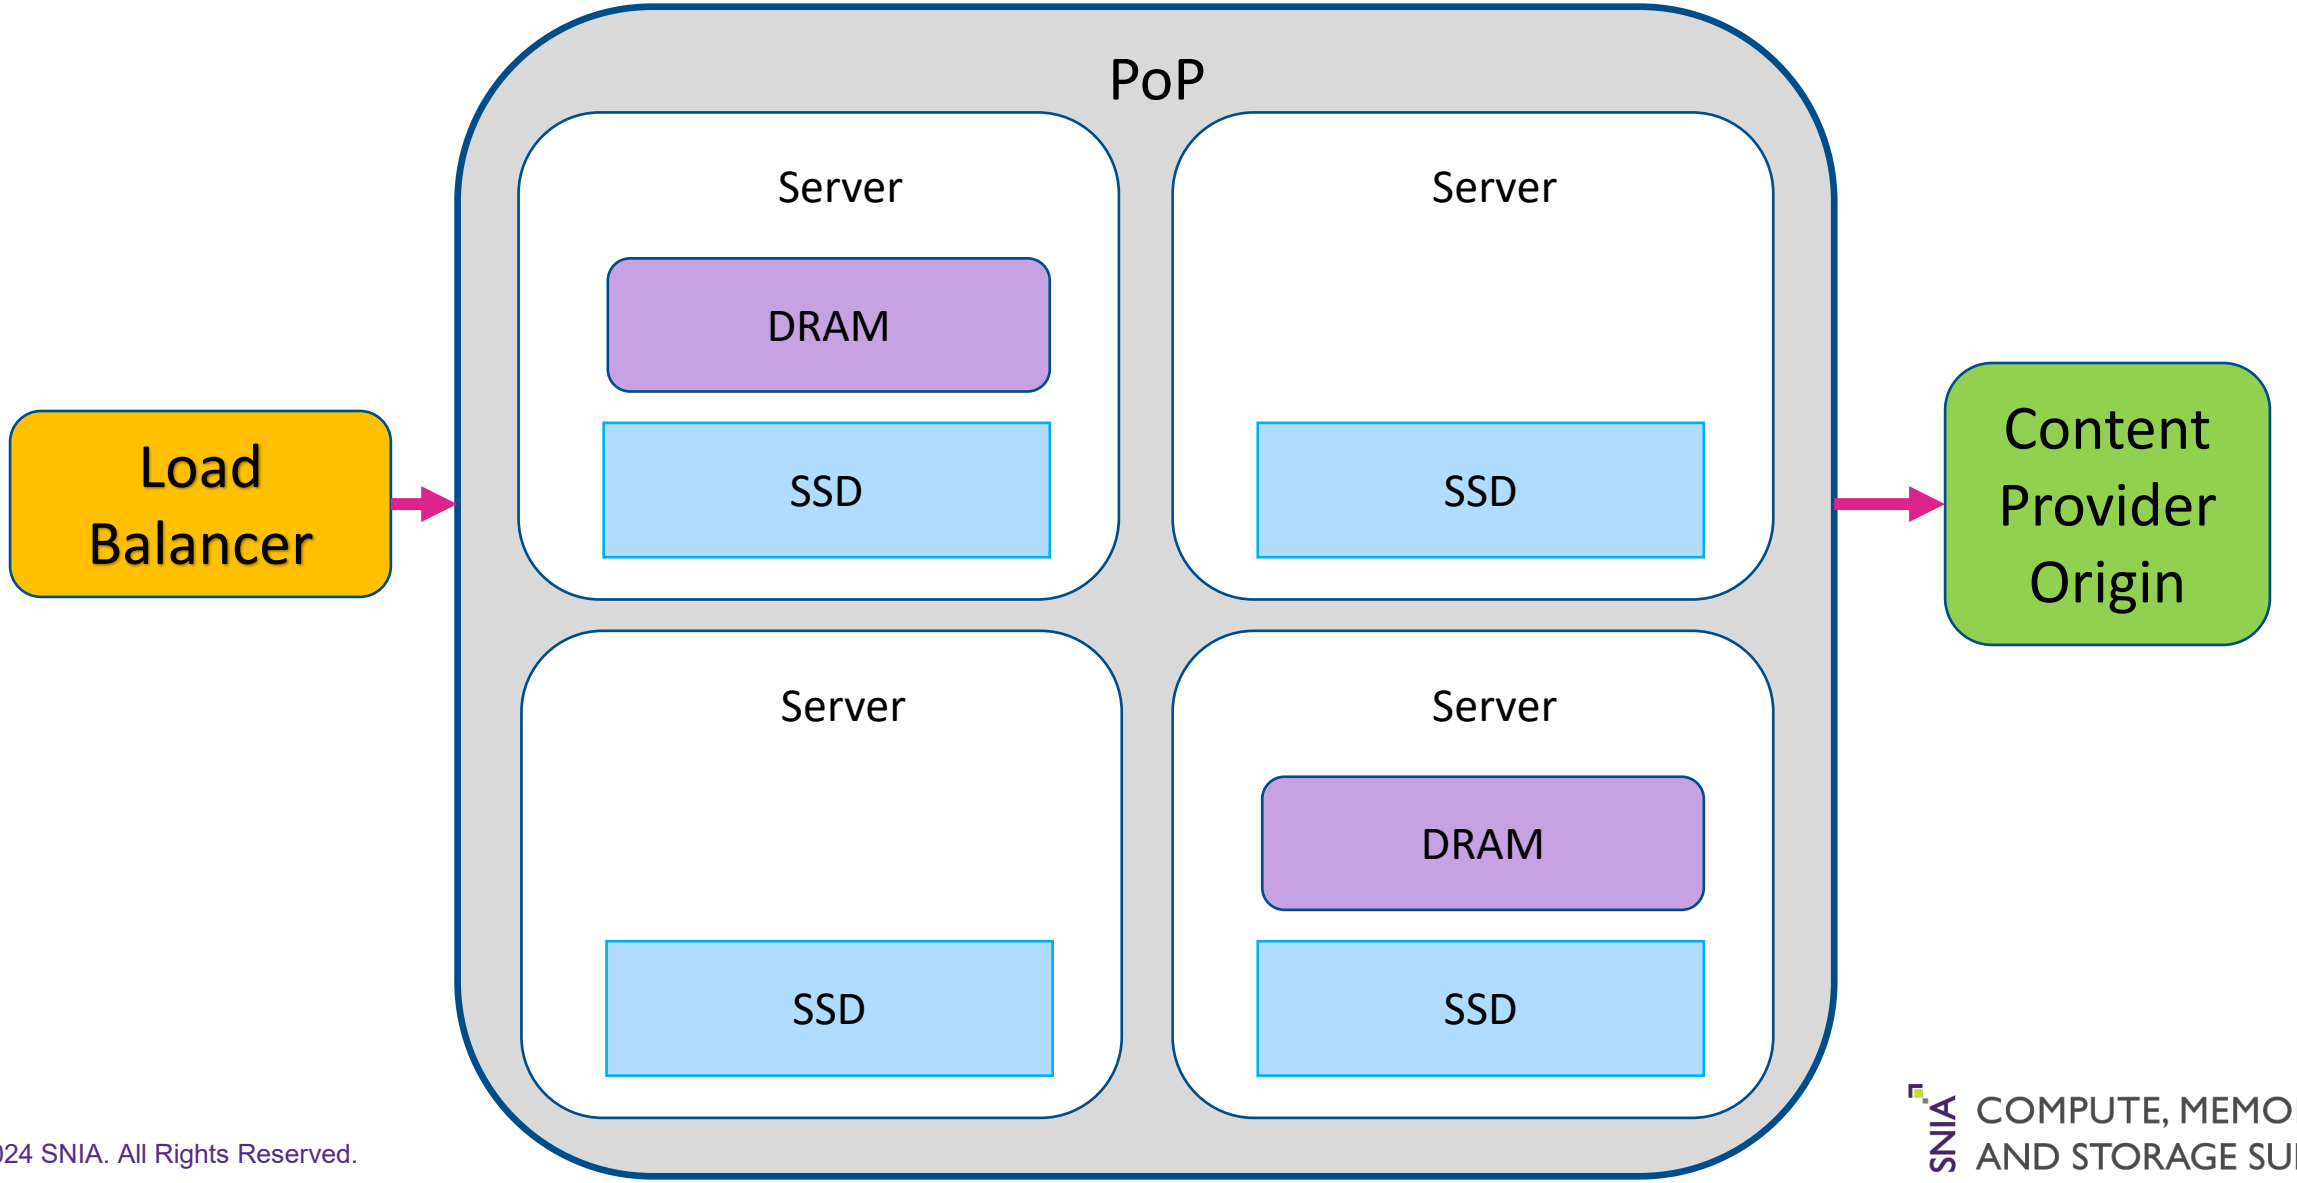

CDN PoP – Compute + Memory + Storage

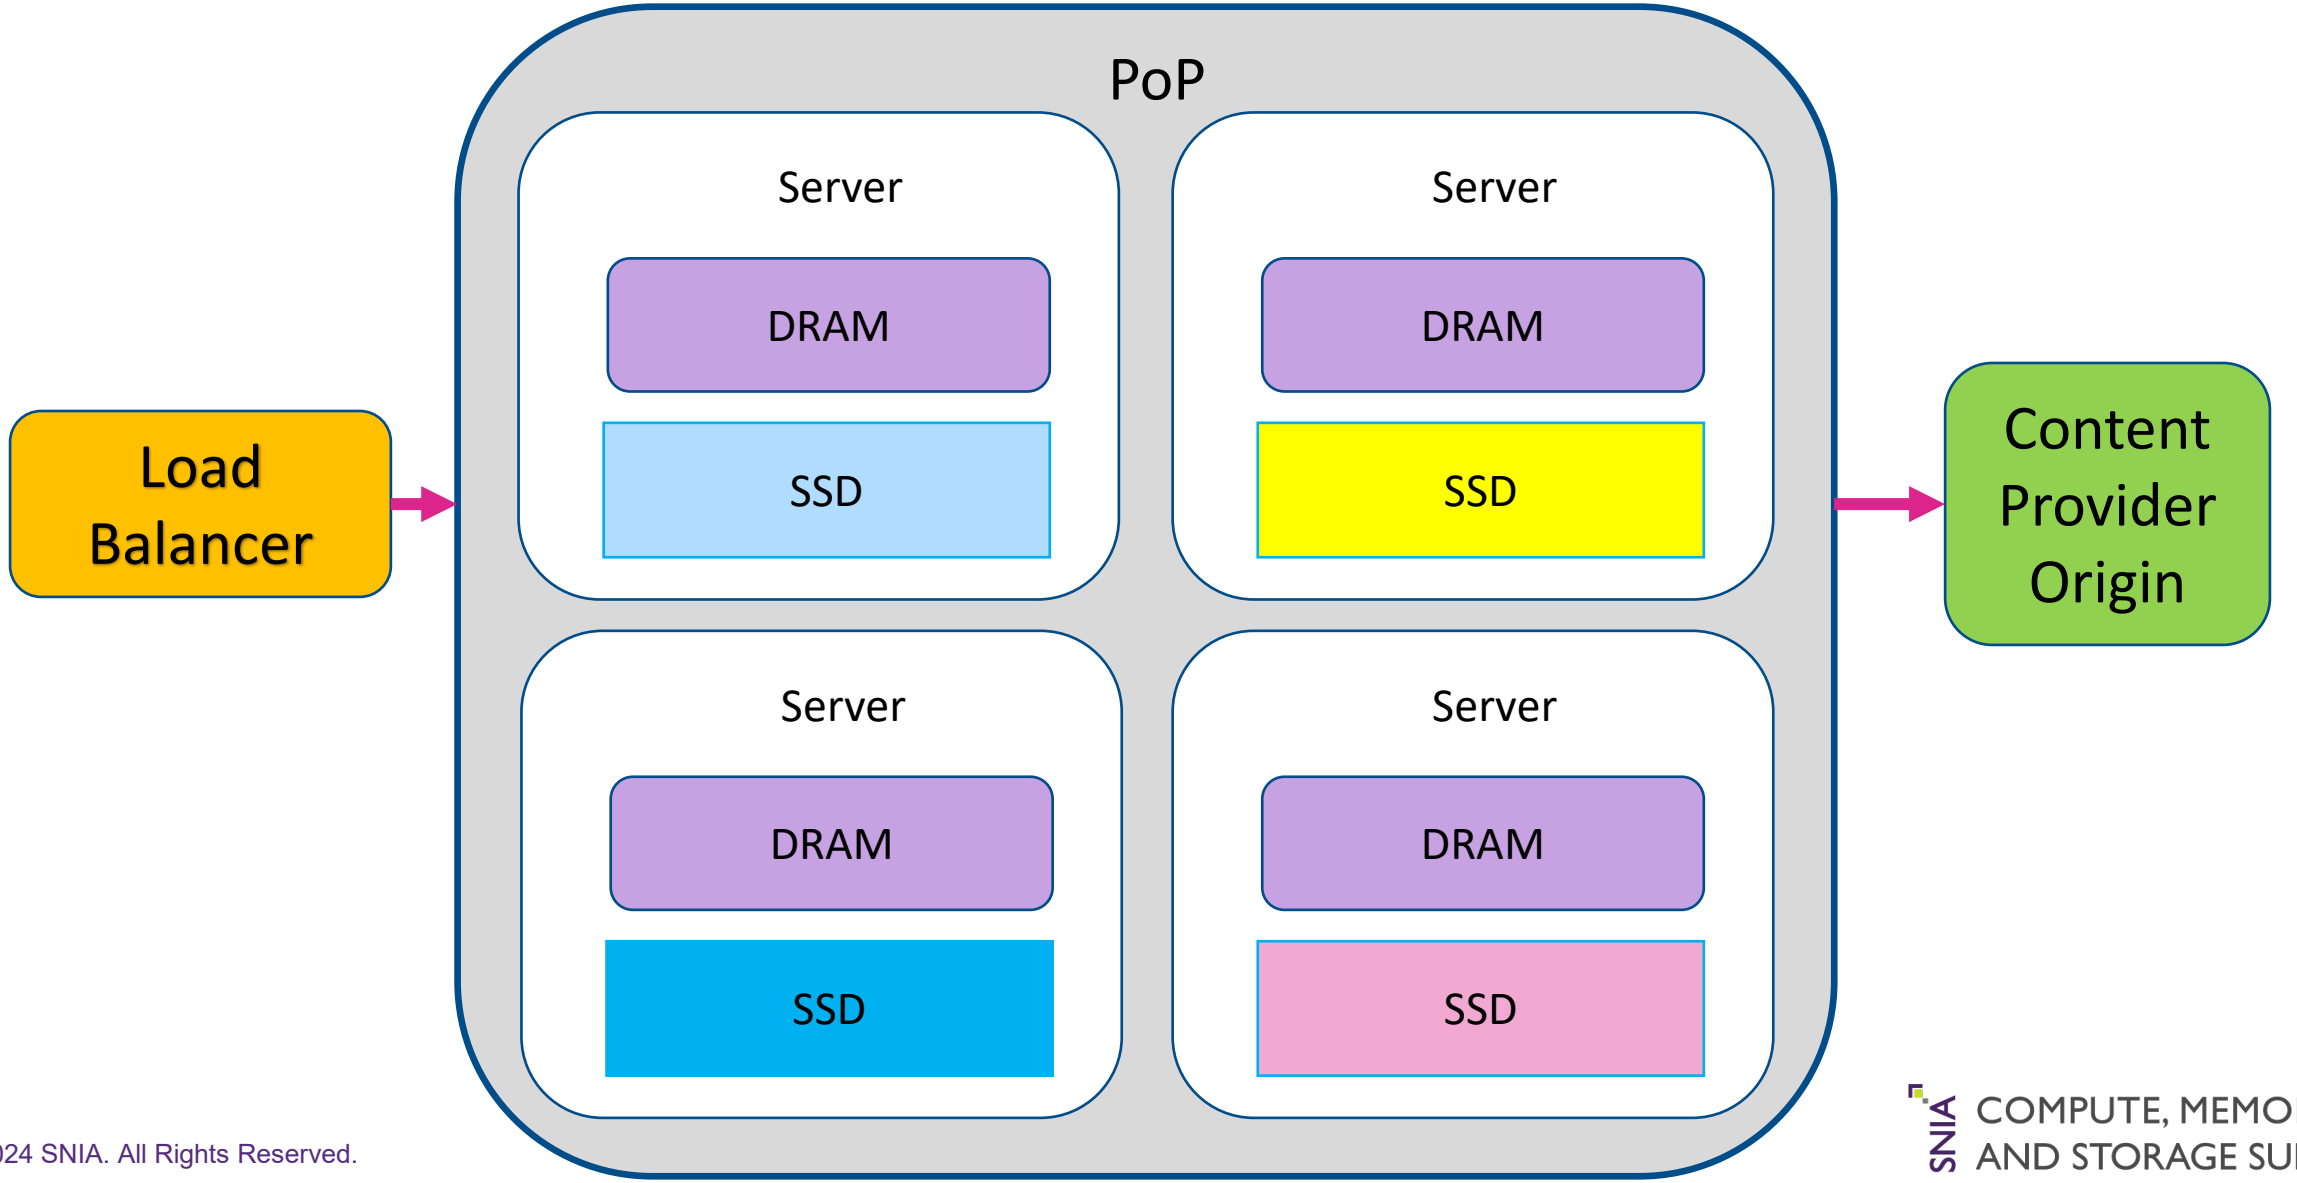

CDN PoP – Compute + Memory + Storage

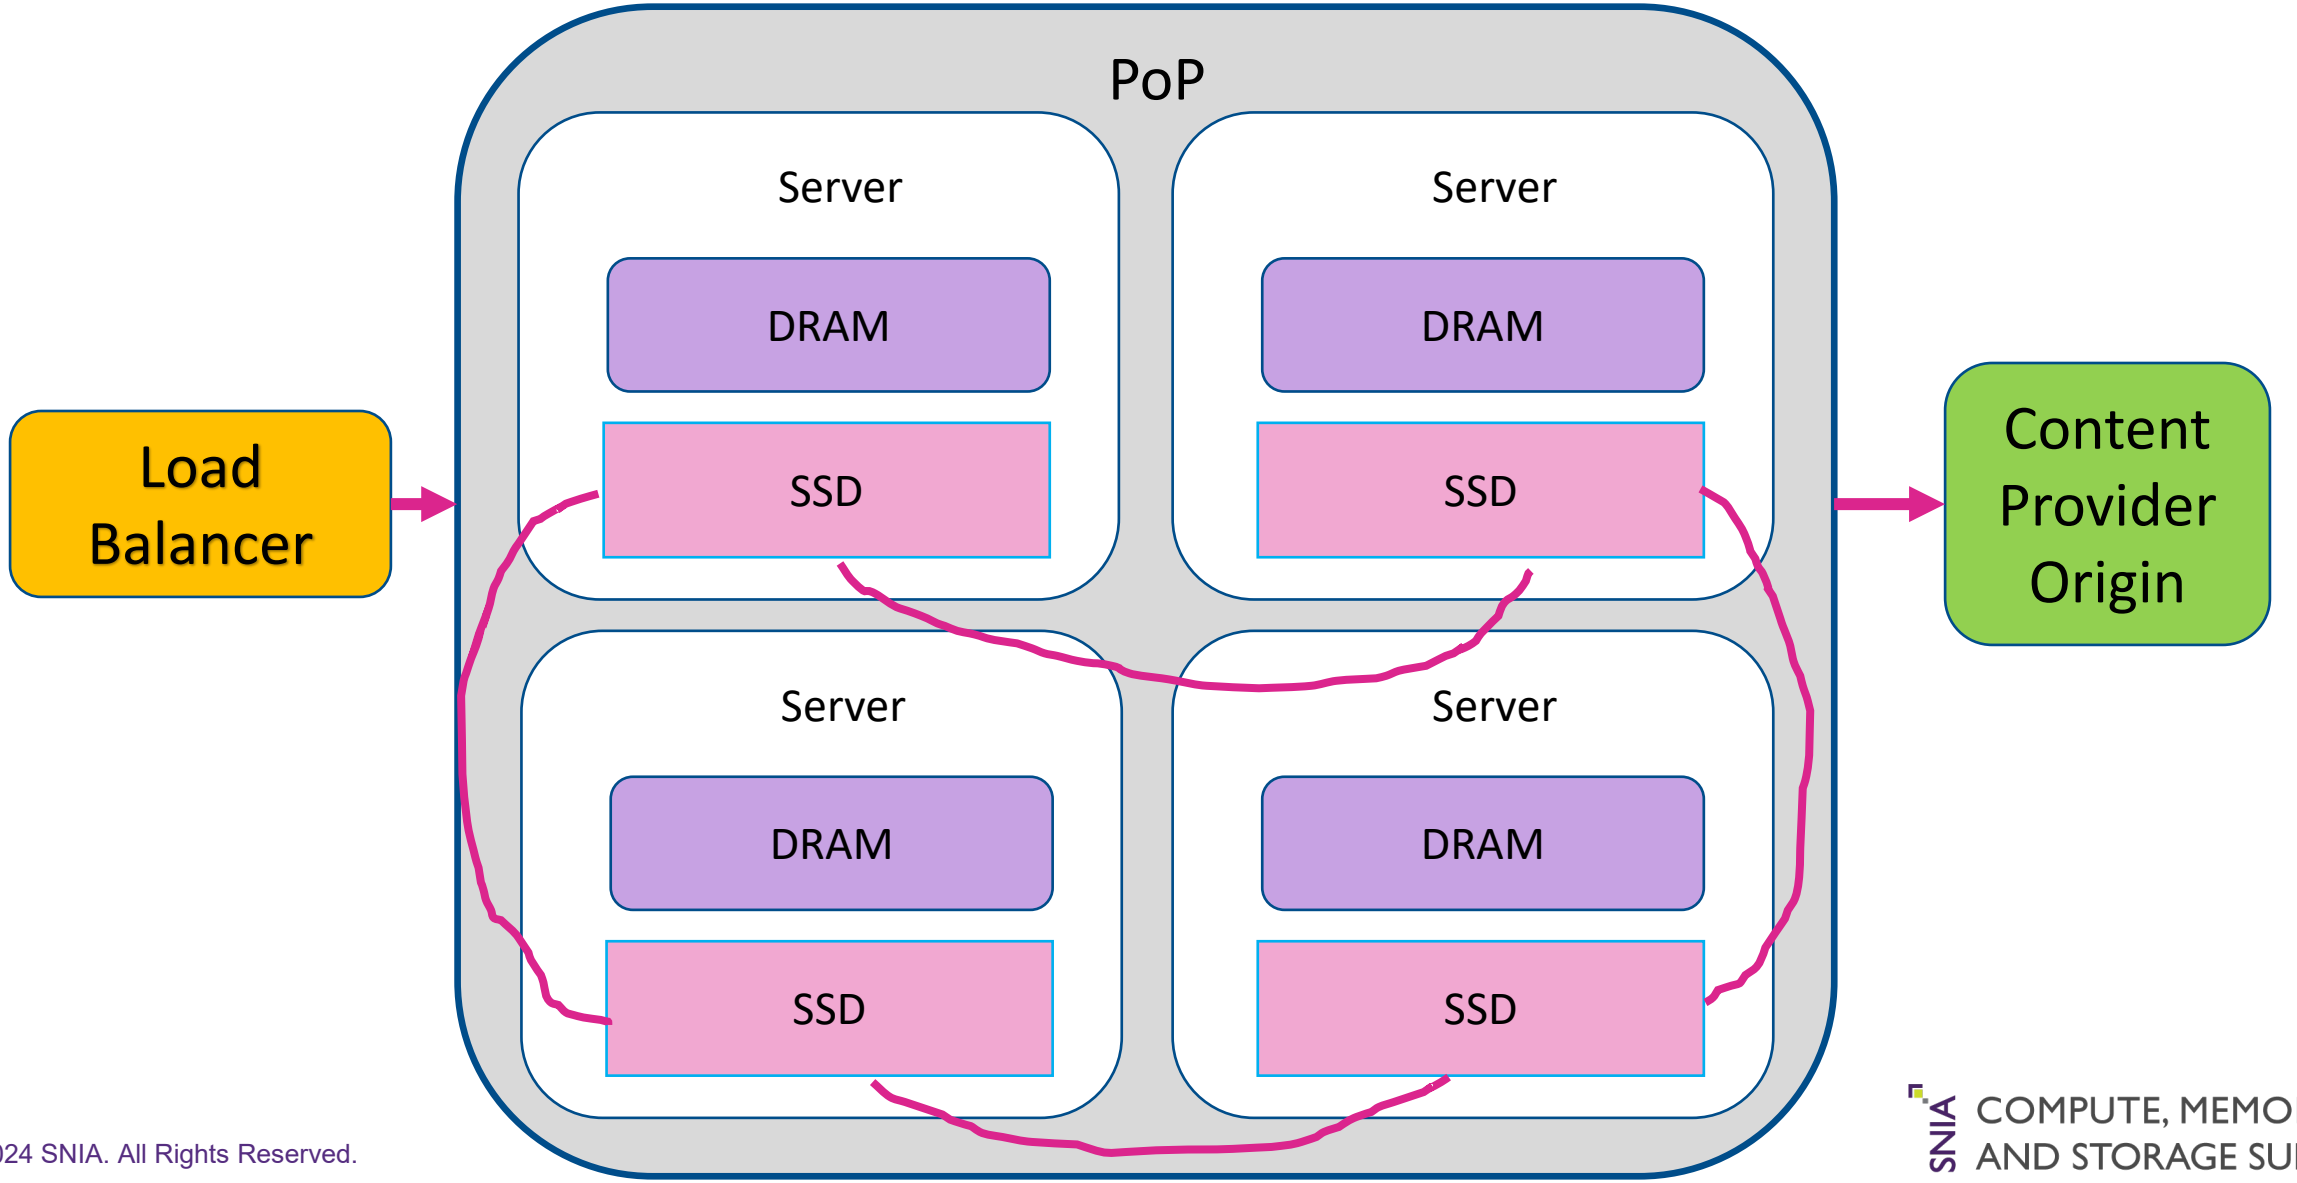

CDN PoP – Compute + Memory + Storage

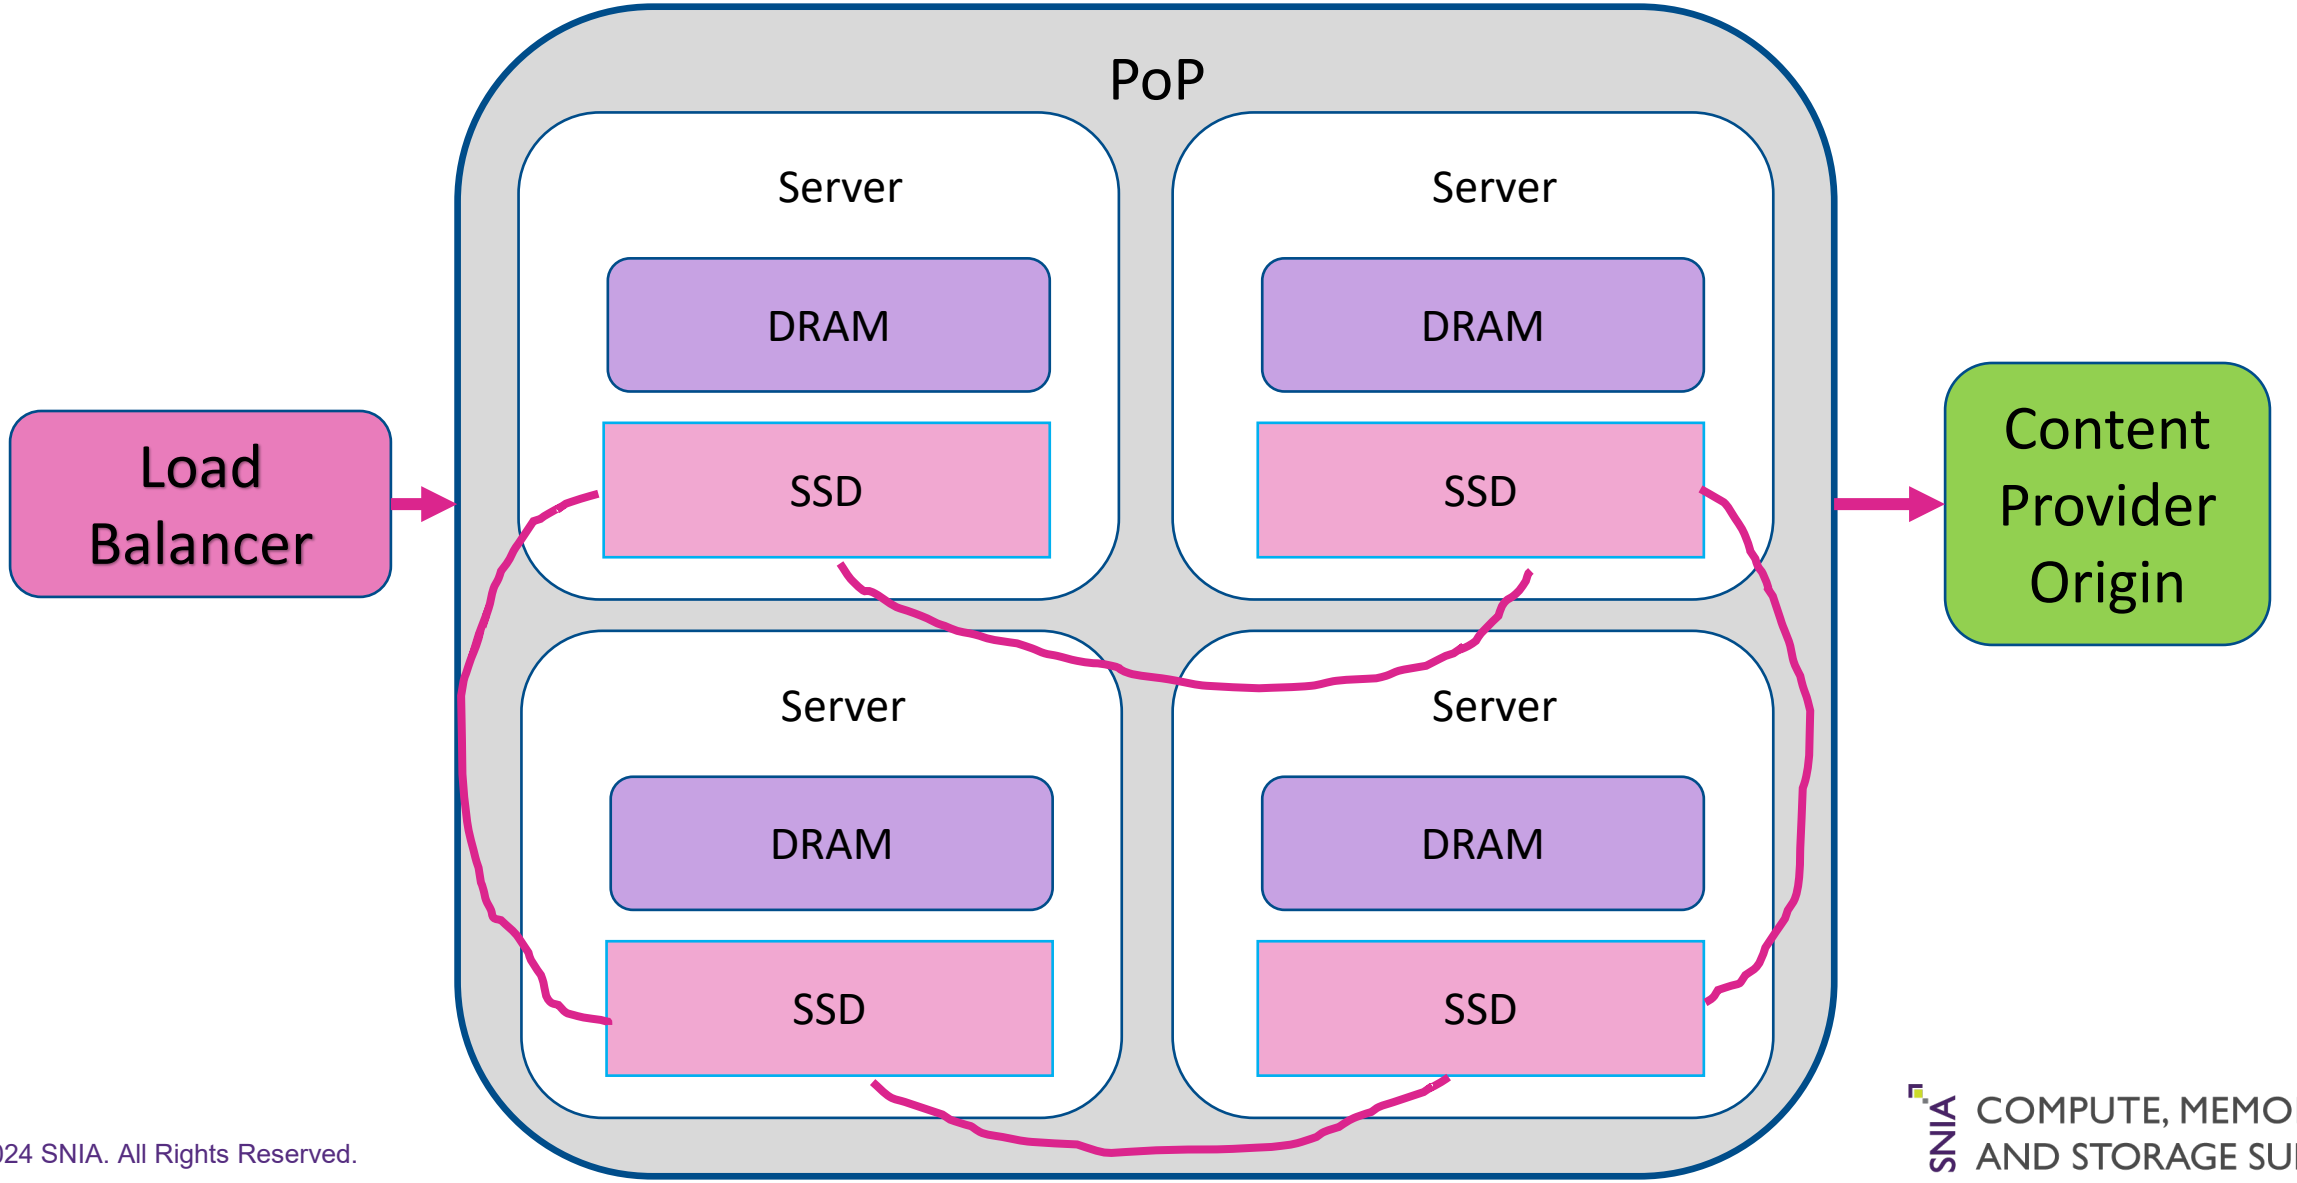

CDN PoP – Compute + Memory + Storage

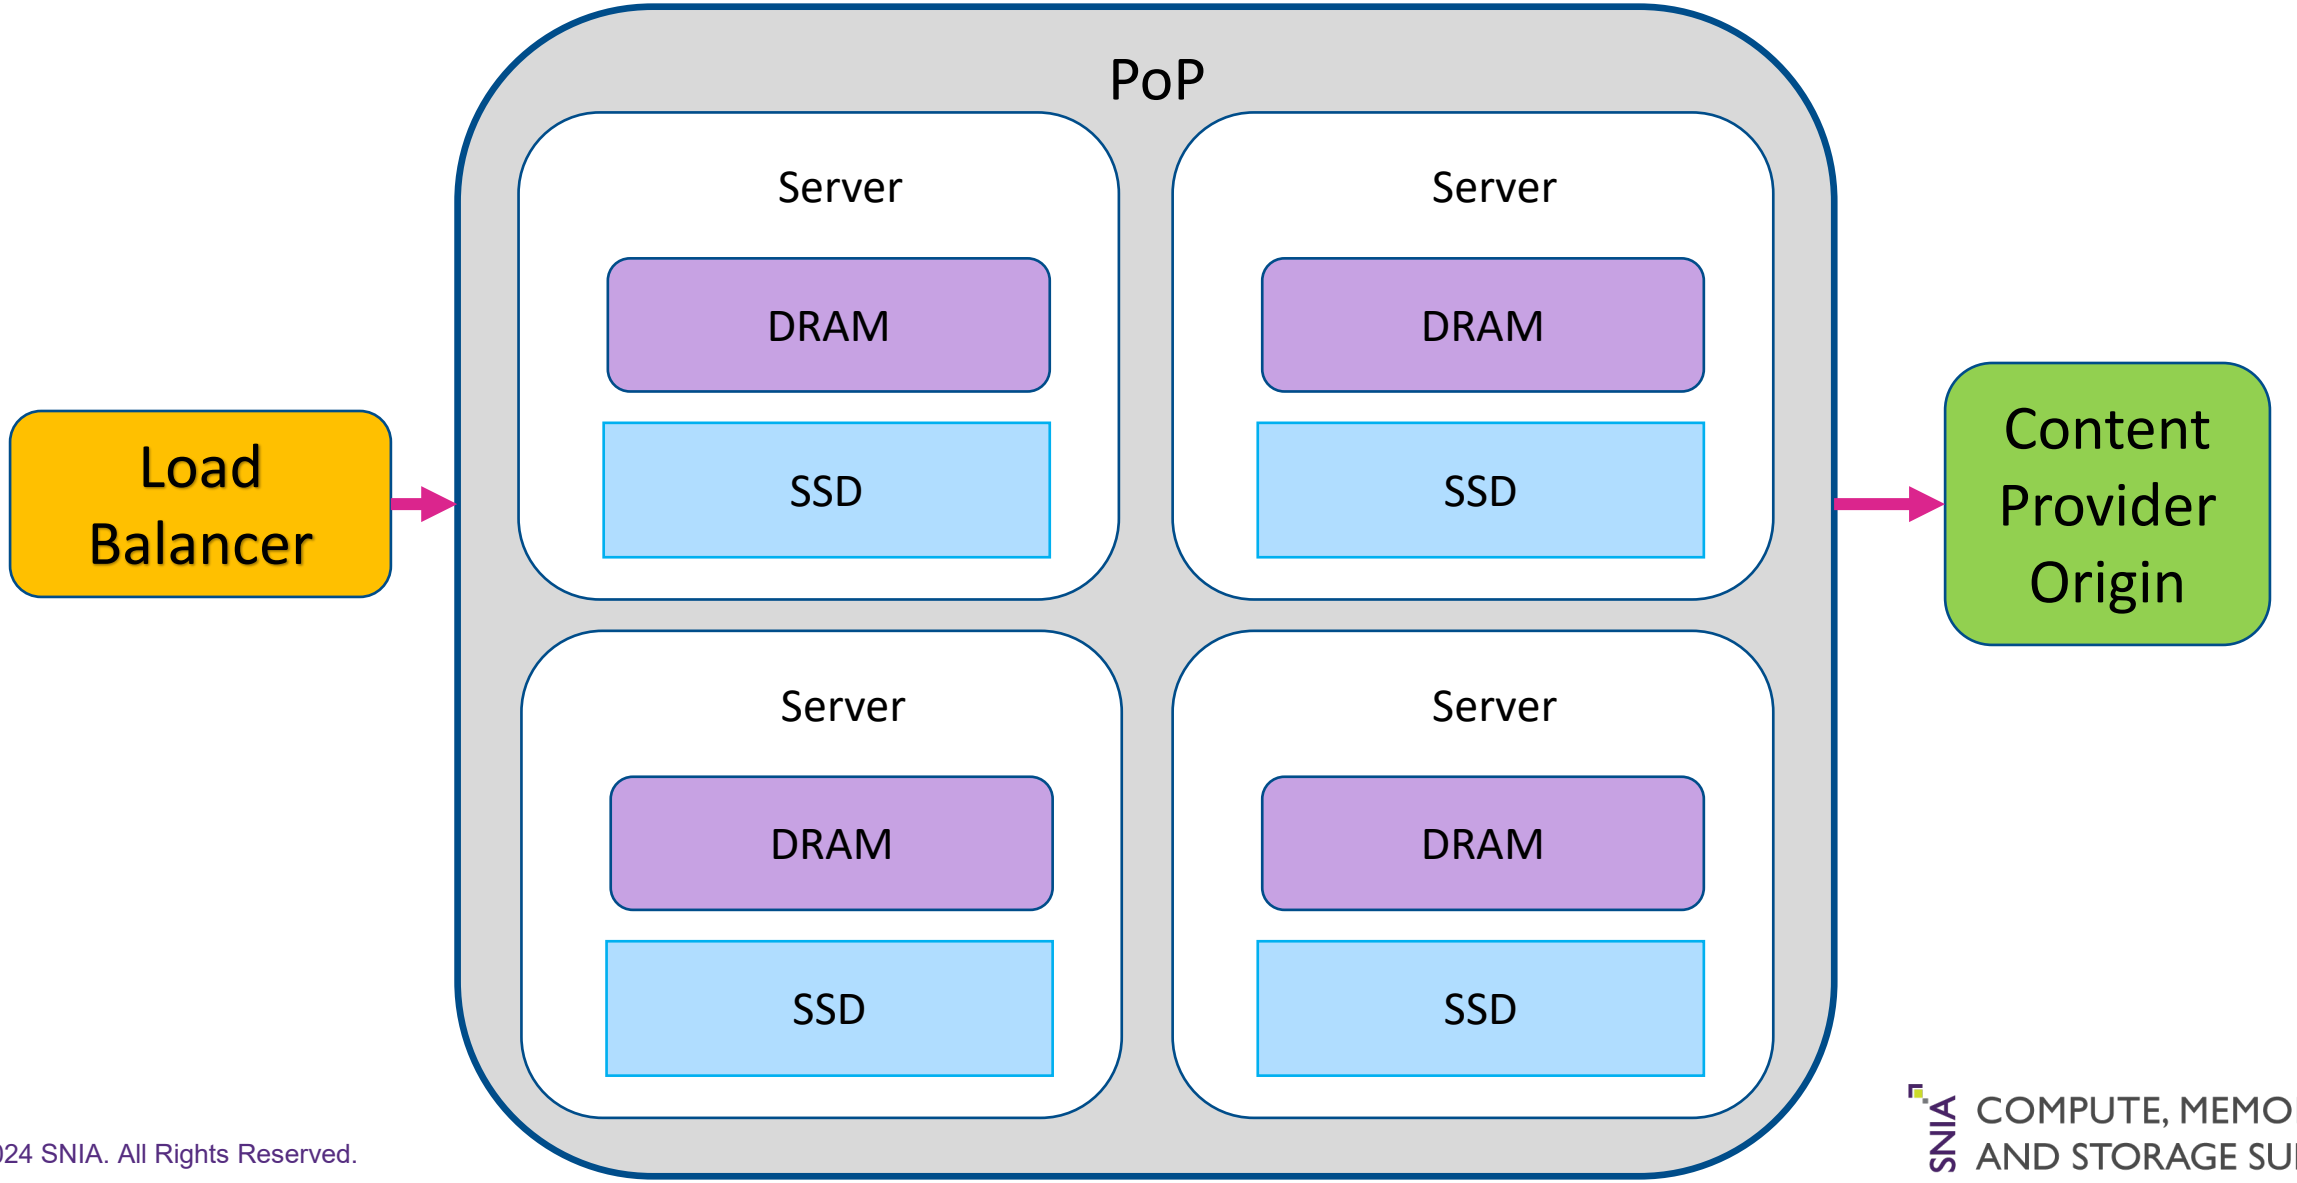

CDN PoP – Compute + Memory + Storage

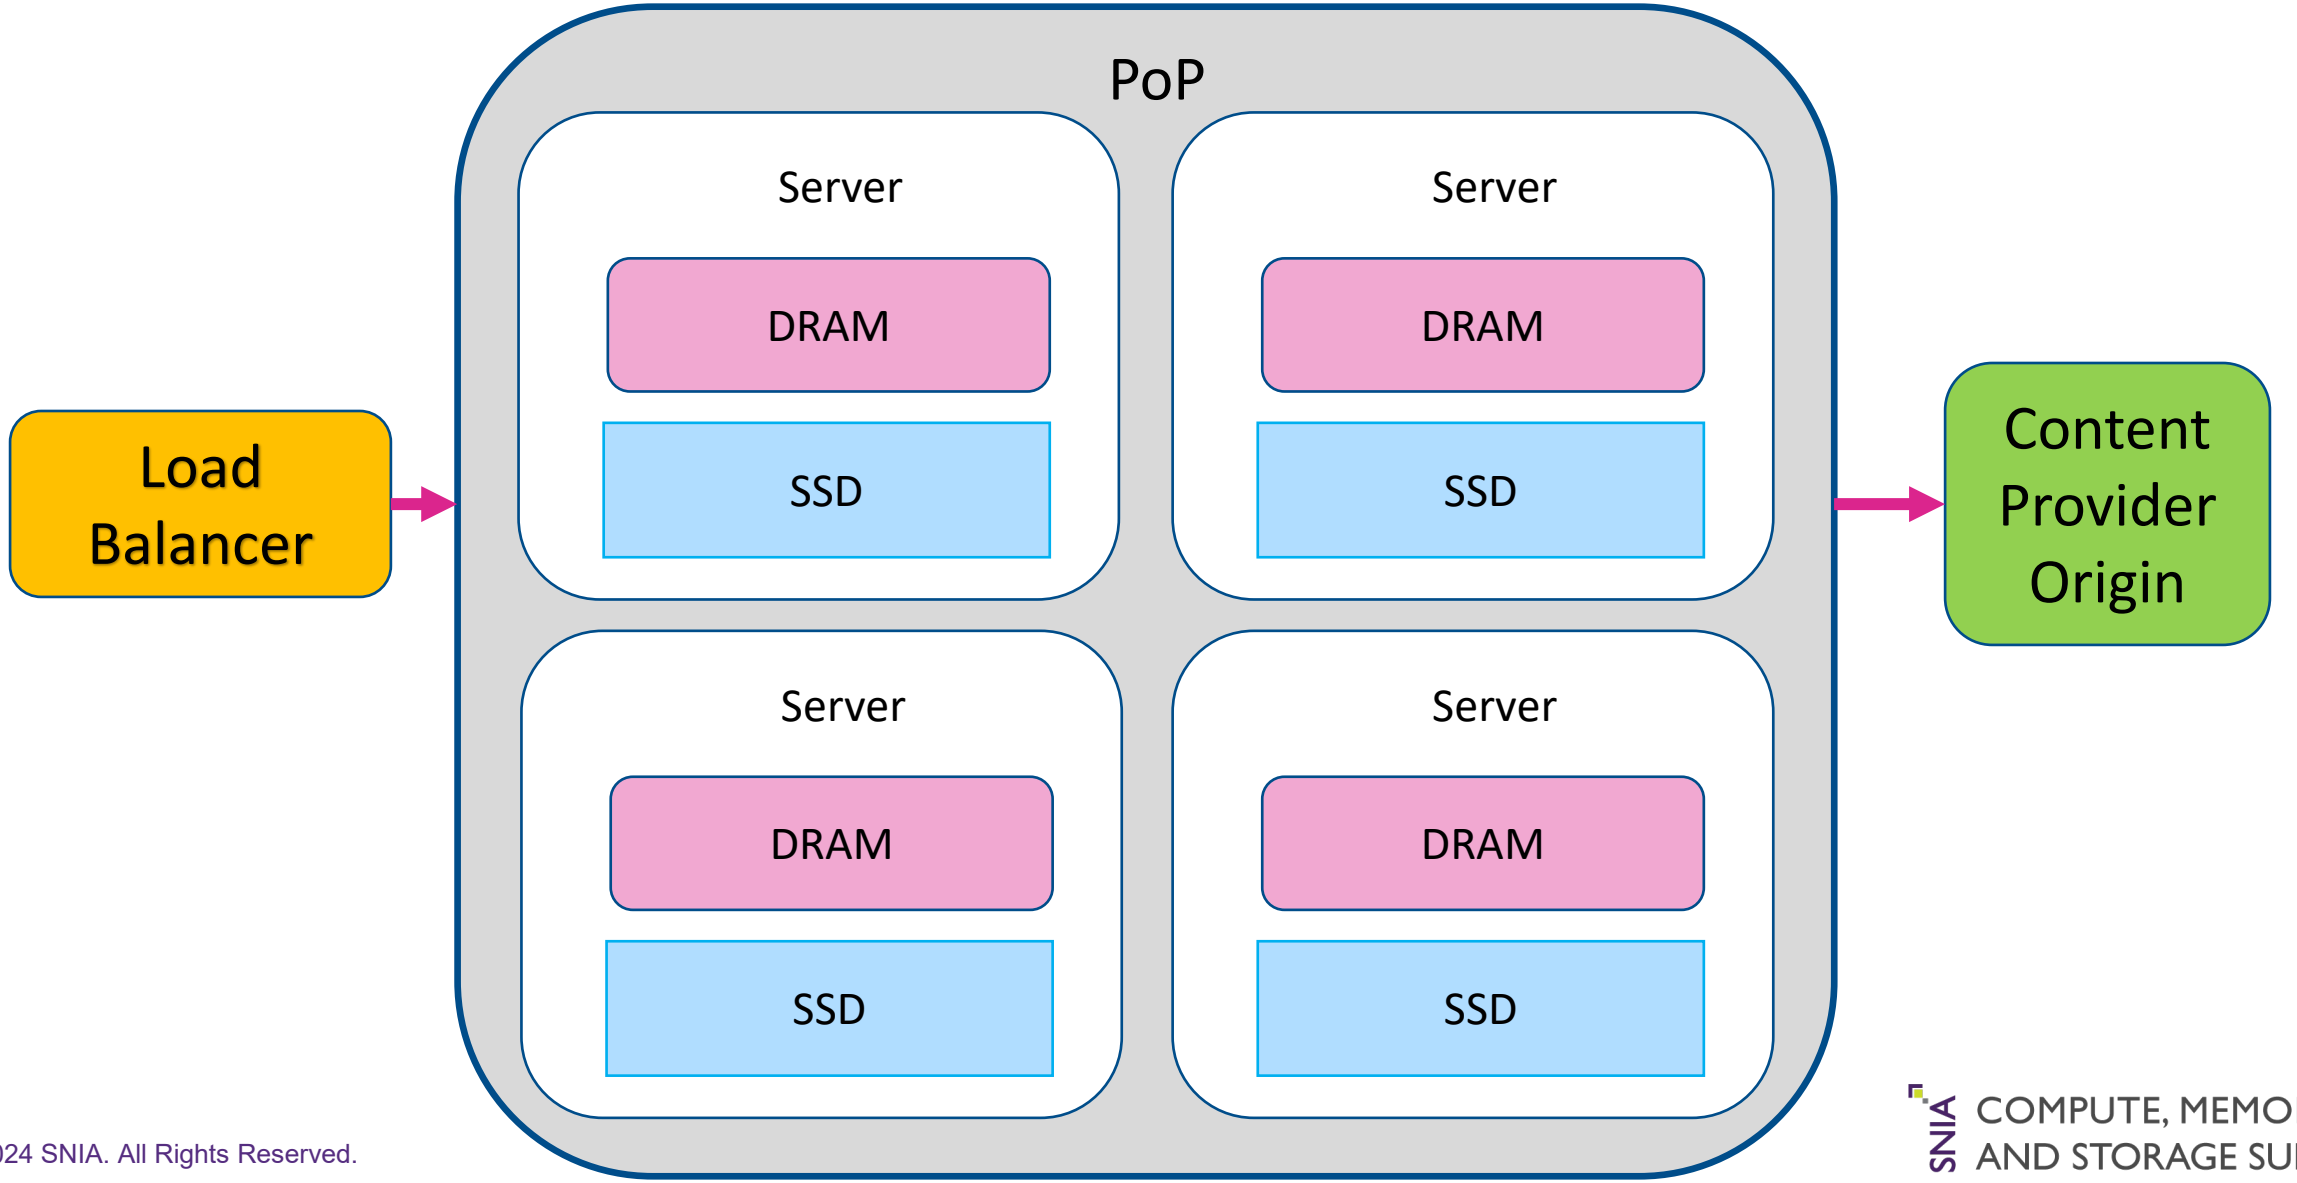

CDN PoP – Compute + Memory + Storage

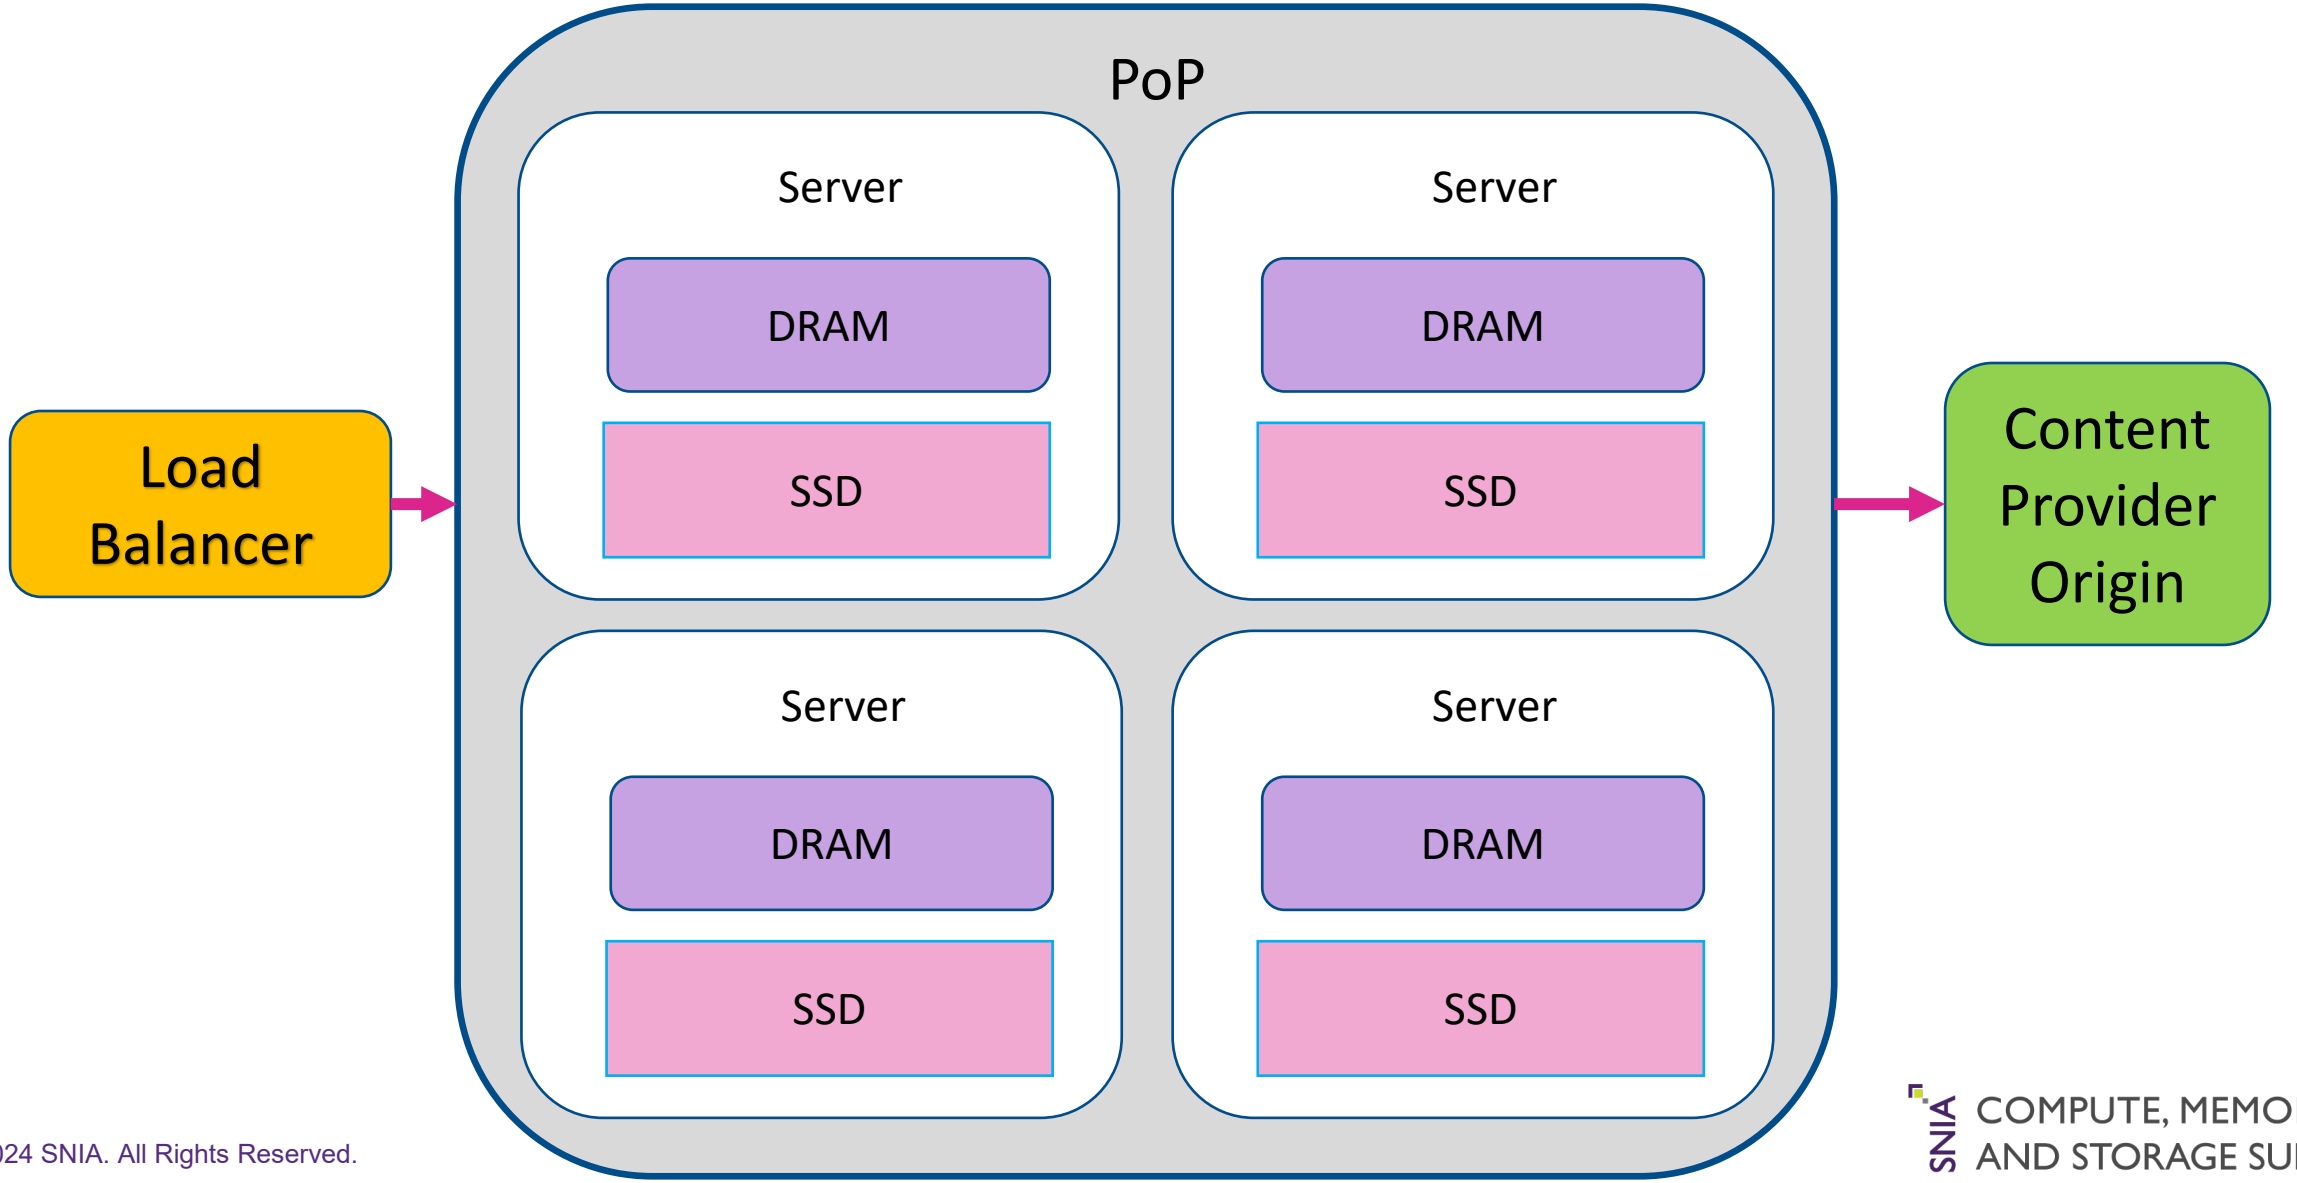

CDN PoP – Compute + Memory + Storage

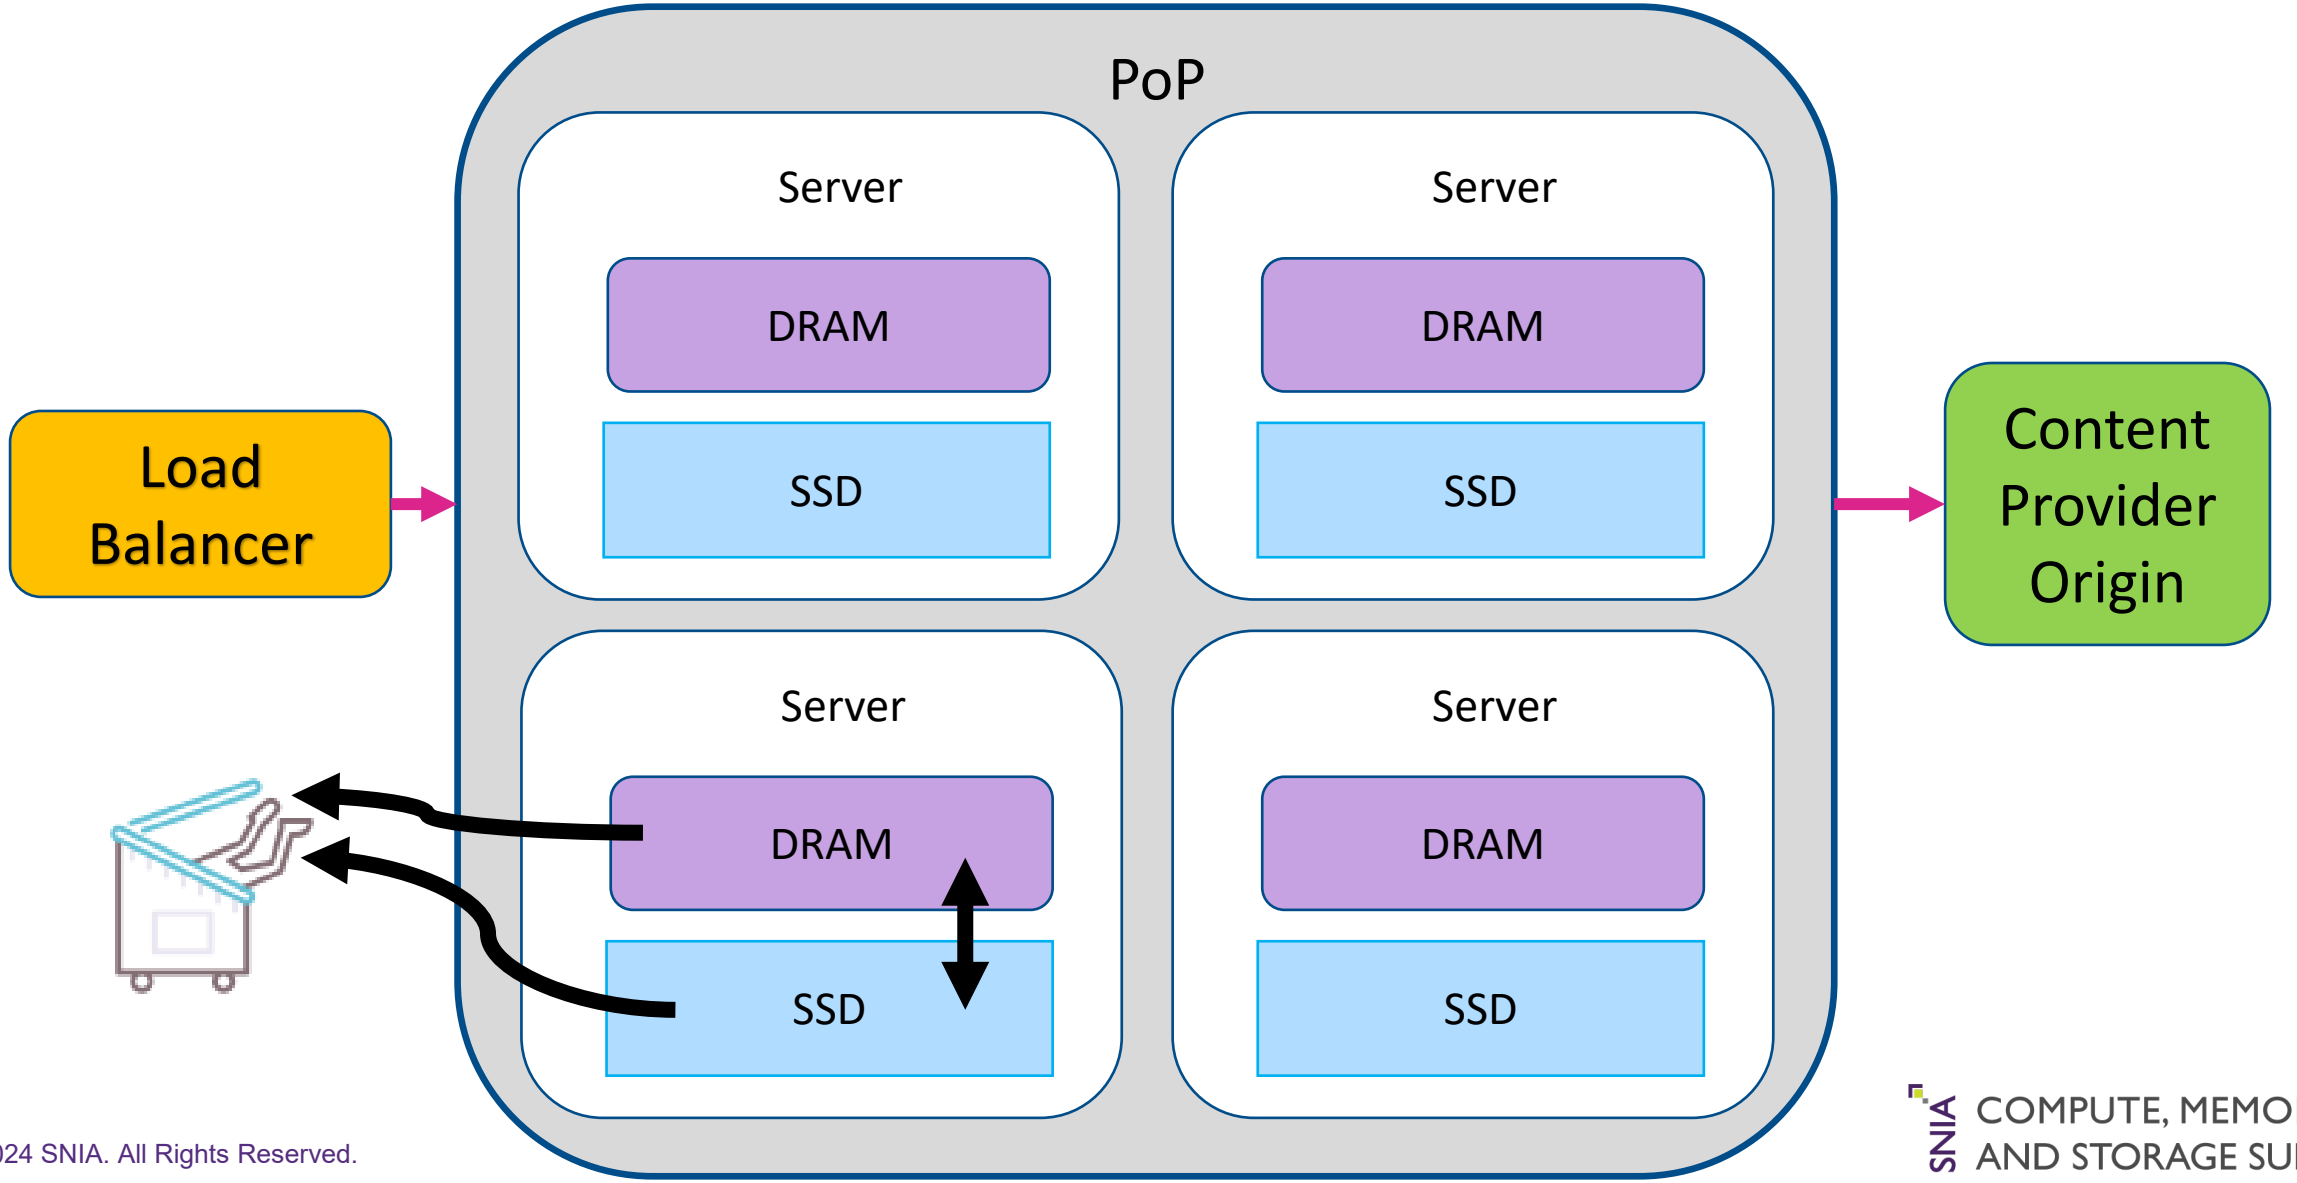

Optimization Nightmares

SNIA COMPUTE, MEMORY, AND STORAGE SUMMIT
Solutions, Architectures, and Community
VIRTUAL EVENT, MAY 21-22, 2024

How to optimize all these factors

- Compose simulations of complex memory and storage
- Break the simulation into components
- Allows the components to be assembled like building blocks
- Provide reasonable but constrained set of variables
- Run simulations with synthetic data or actual IO traces

Deterministic Behavioral Simulations

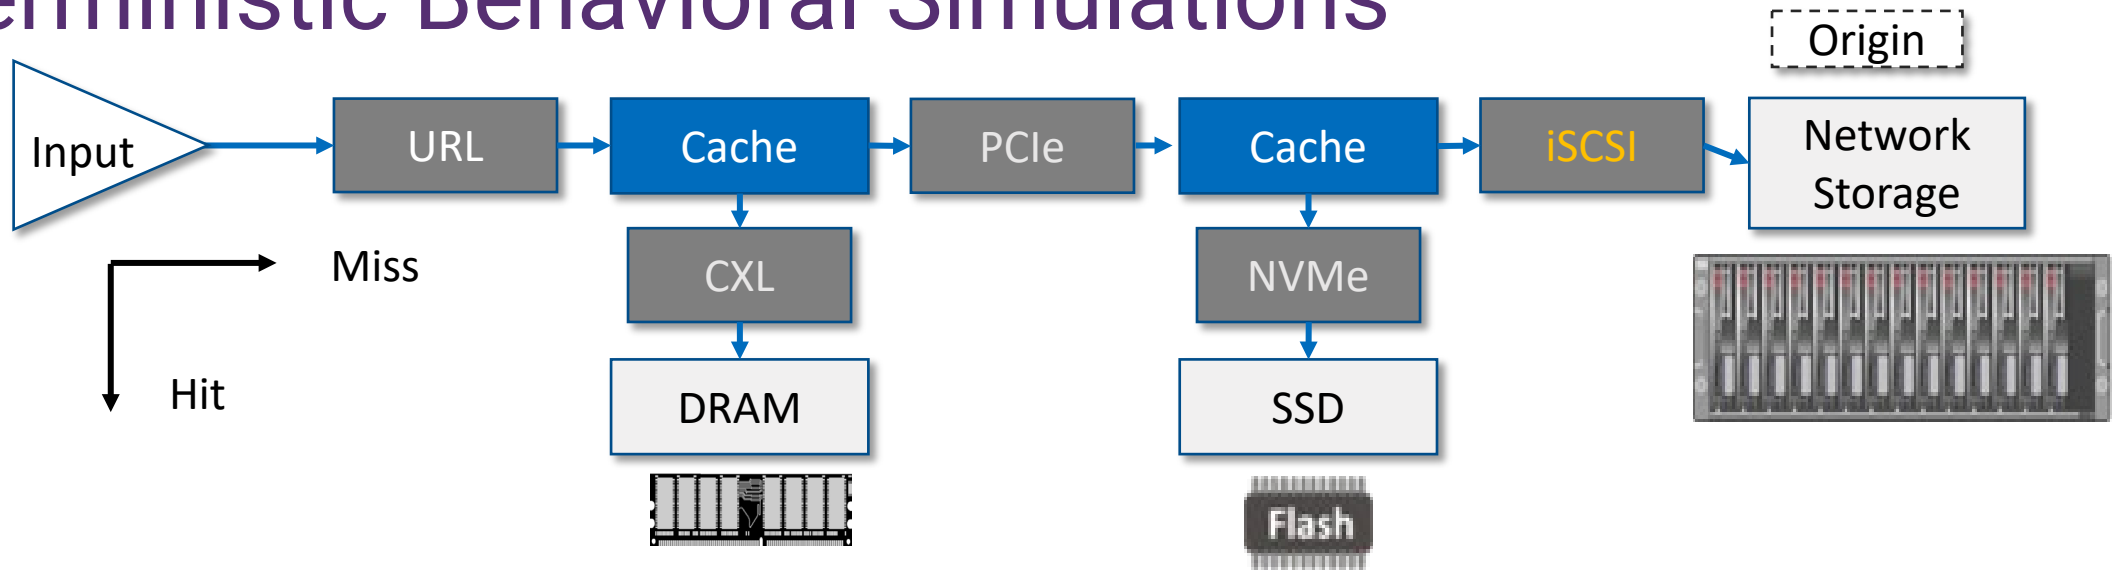

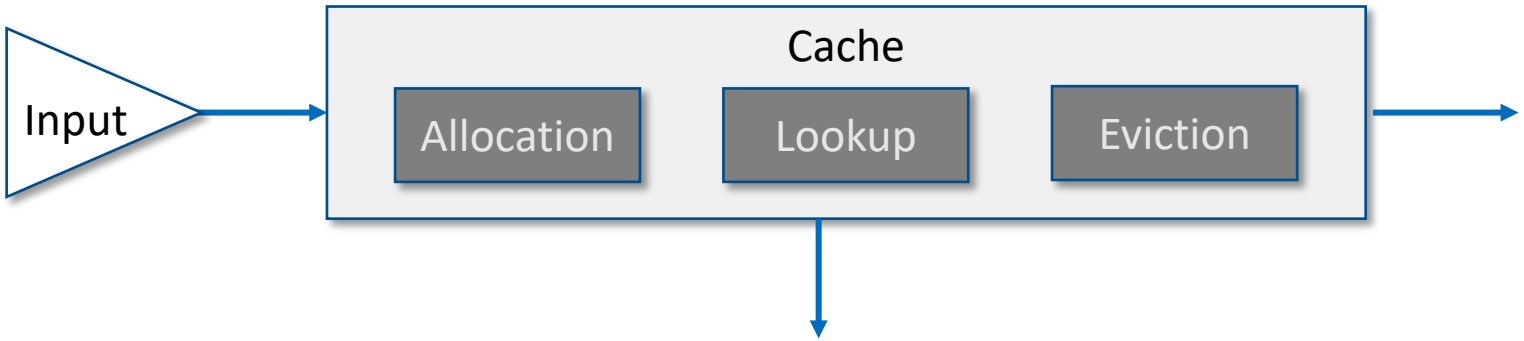

Case Studies

- Cloudflare and Wikimedia
- Genuine workloads tested
- Variants of baseline algorithms
  - Number of L1 and L2 caches
  - Promotion and demotion policies
  - Eviction policies of L1 and L2
  - DRAM/SSD ratios
  - Load balancer algorithm
- ~125k variants, each run against a day's worth of traffic

Variables

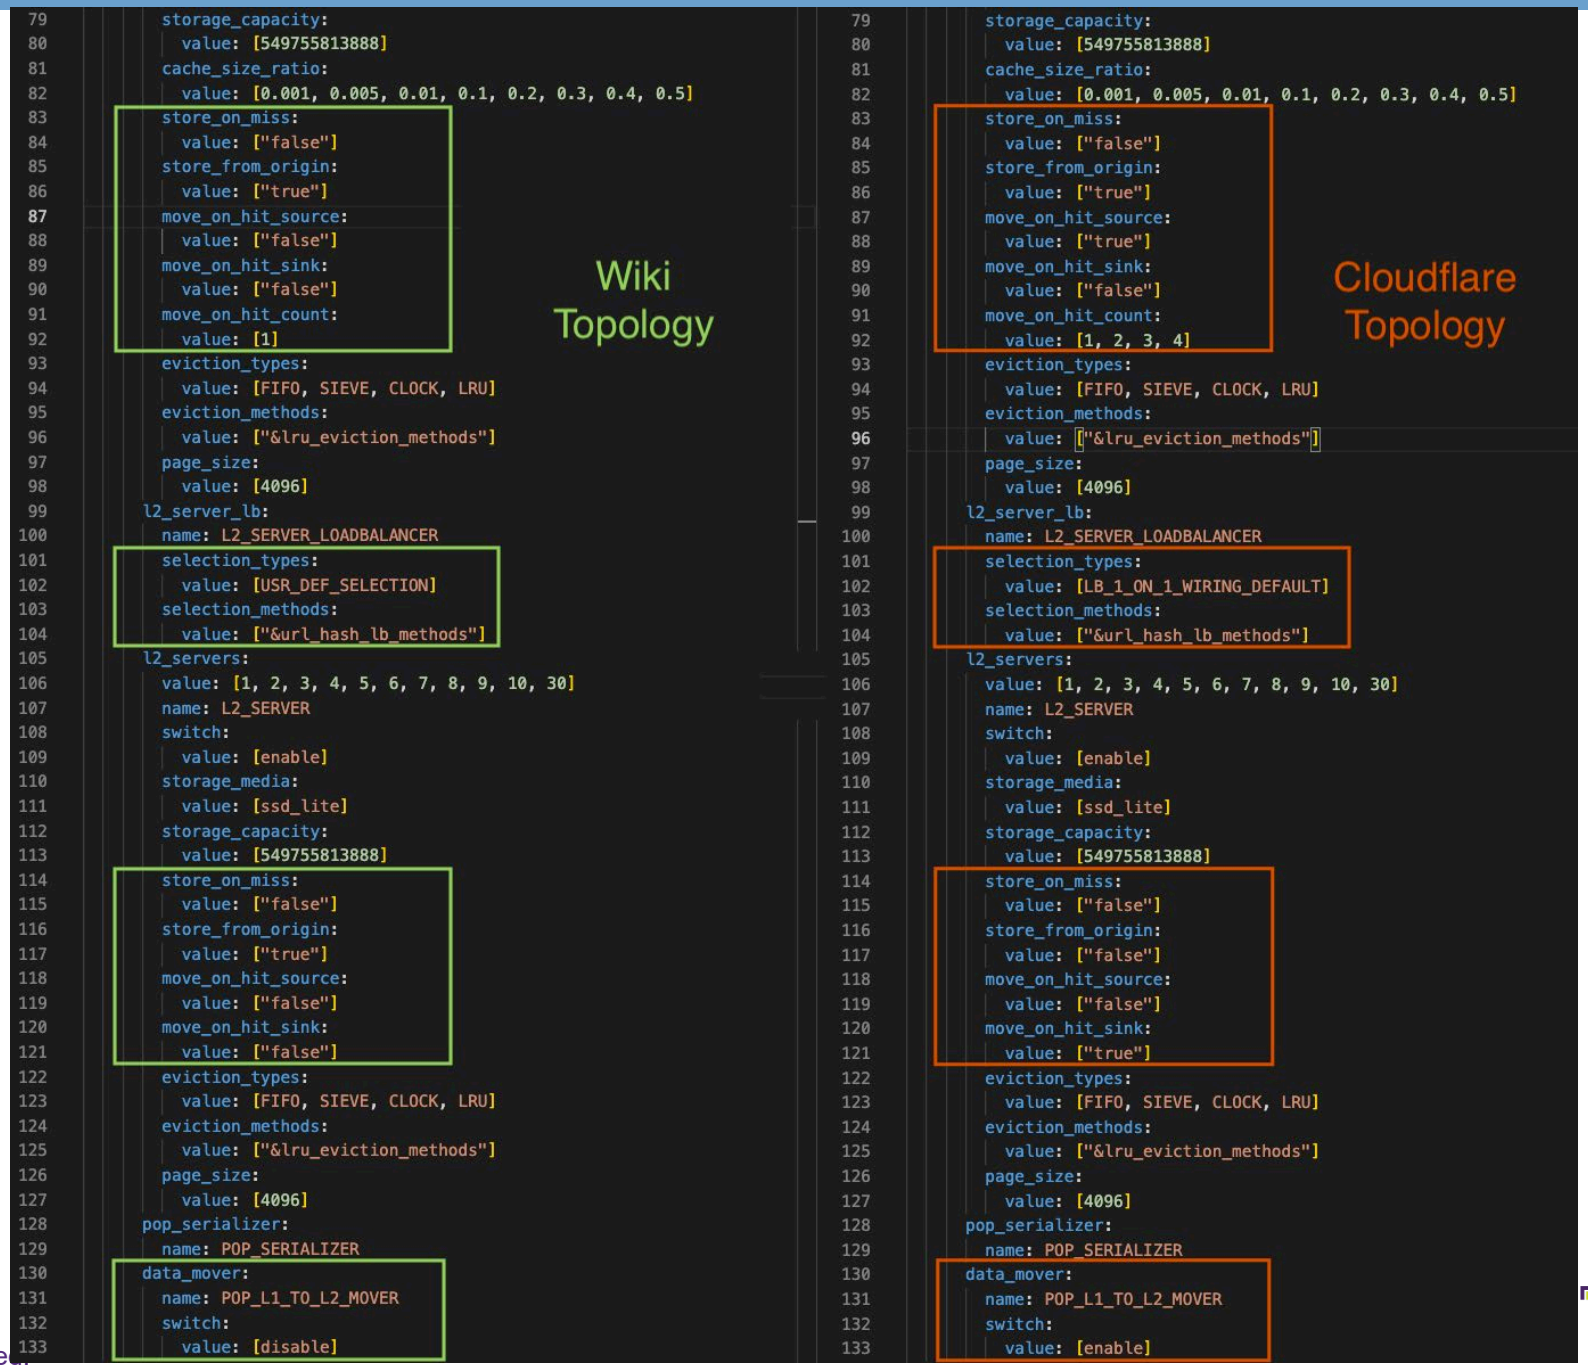

Case Study Results

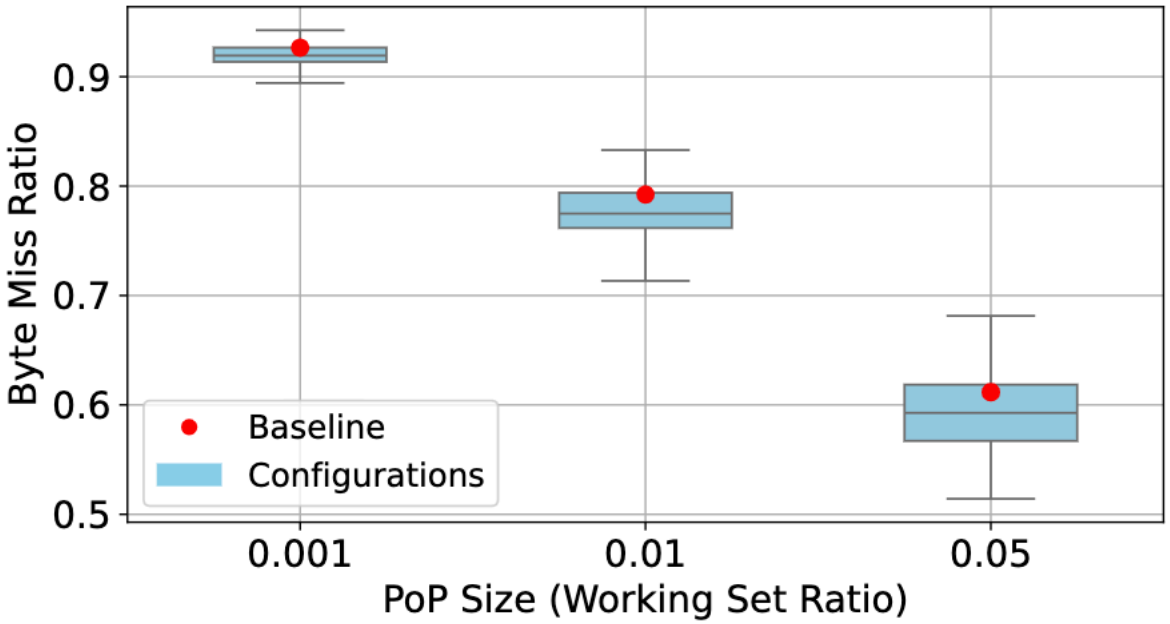

Wikimedia

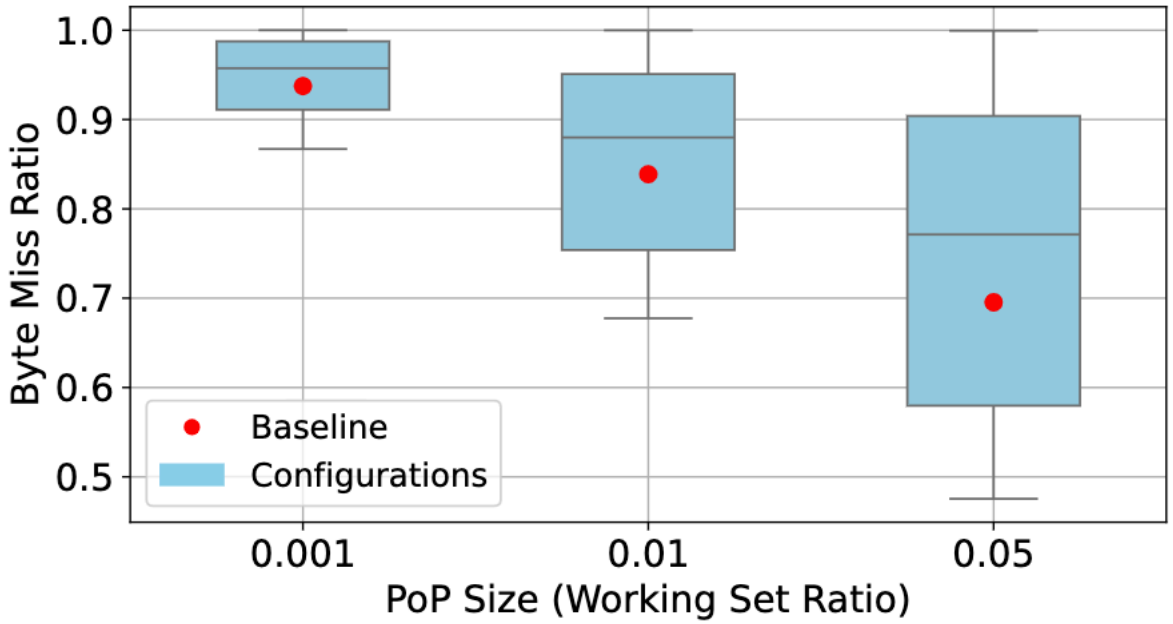

Cloudflare

Case Study Results

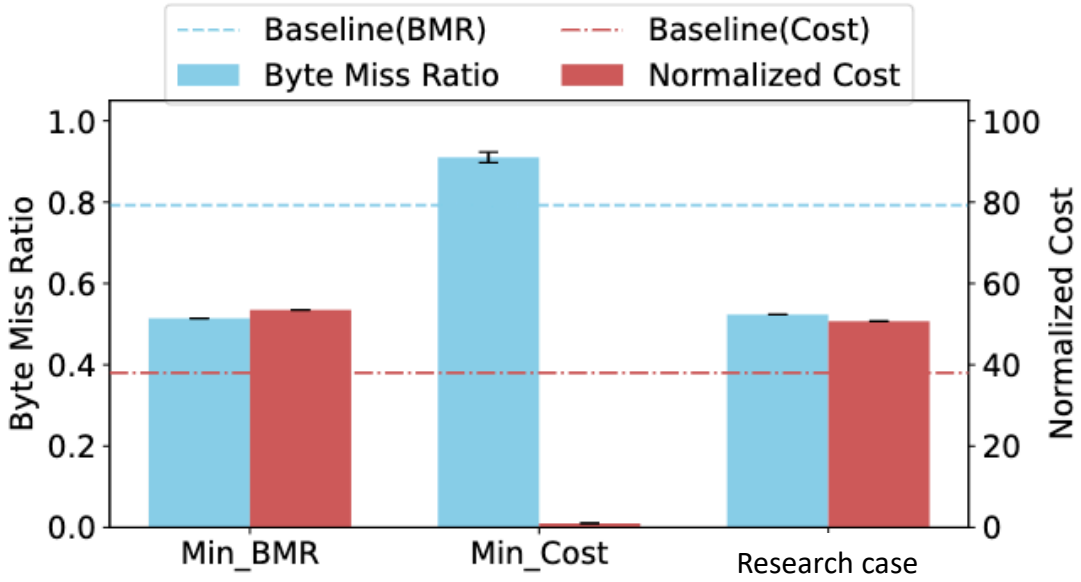

Wikimedia

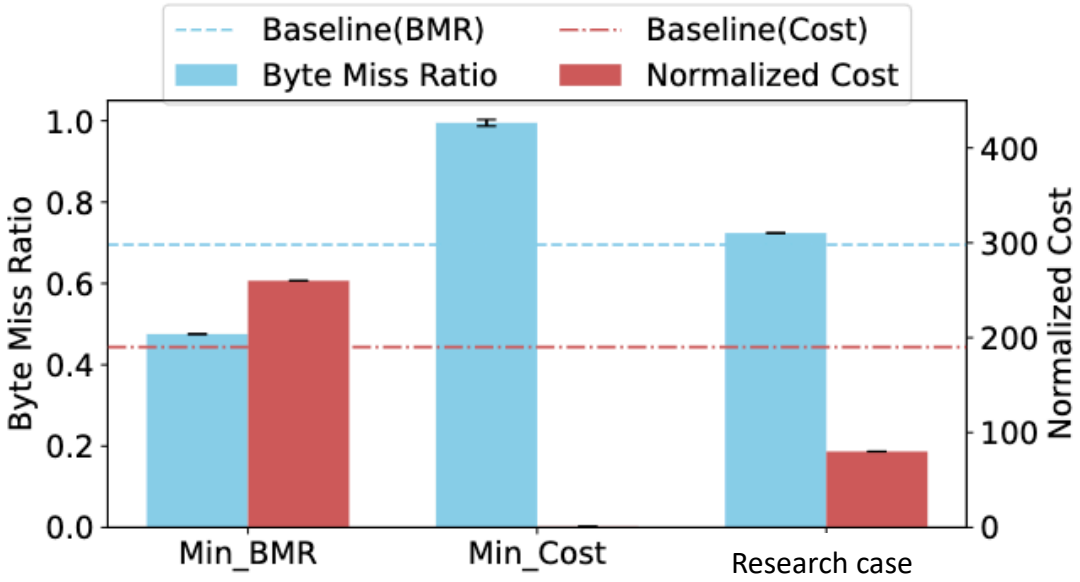

Cloudflare

Case Study Results

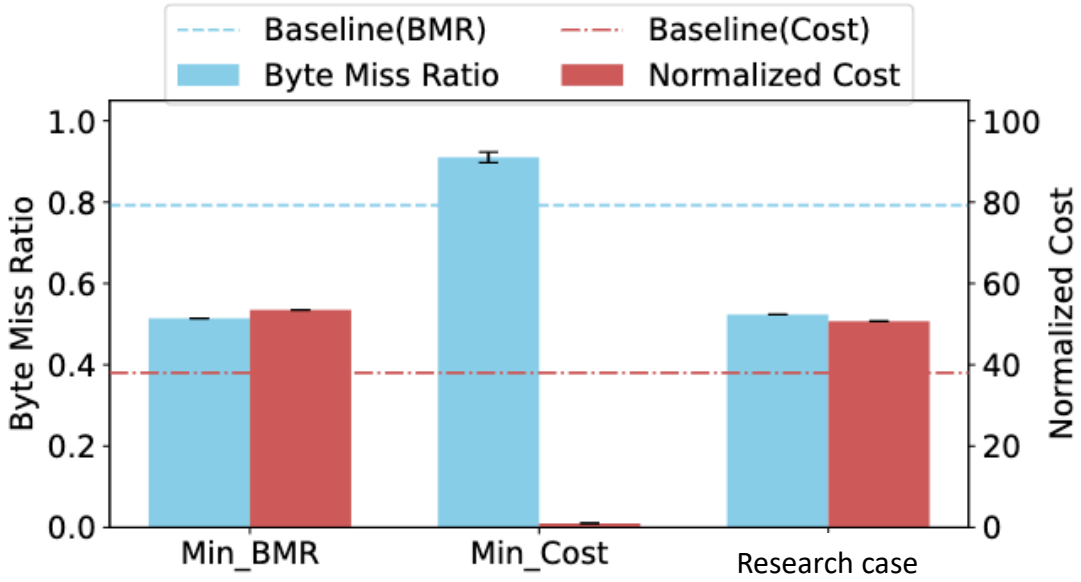

Wikimedia

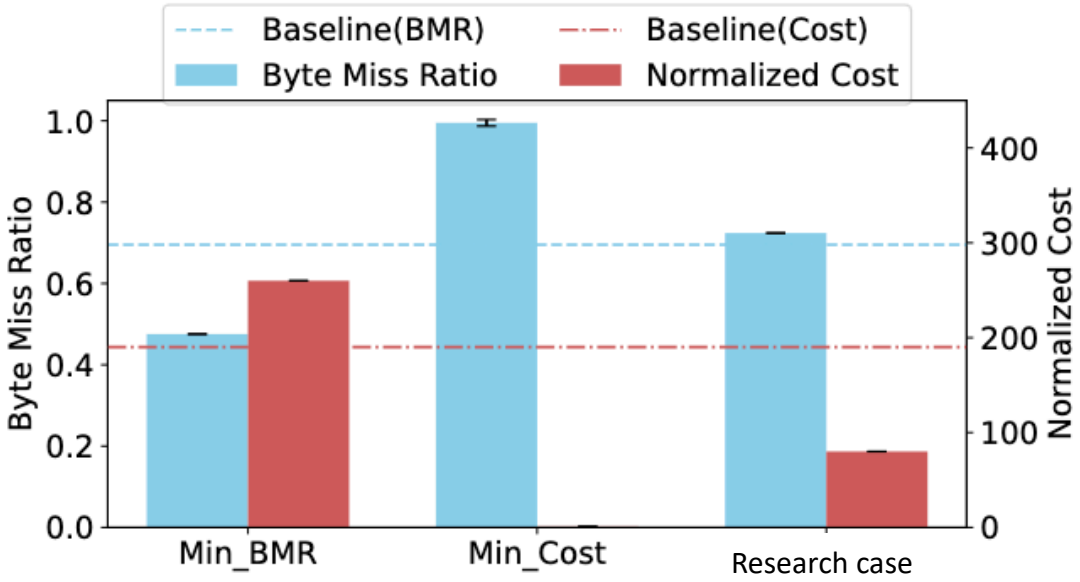

Cloudflare

https://cachemon.github.io/SIEVE-website/

RESULTS WITH MAGNITION

PROVEN IN MARKET TODAY

As an example, a current customer has achieved the following measurable outcomes with Magnition:

Experiments **per day per engineer**:

- Without Magnition**: 2**
- With Magnition: **50,000+**

Parameter variations tested **before prod release**:

- Without Magnition**: 50**
- With Magnition: **1,000,000+**

Workload performance improvement using our products to find **optimal out-of-the-box settings**: **10-50%+**

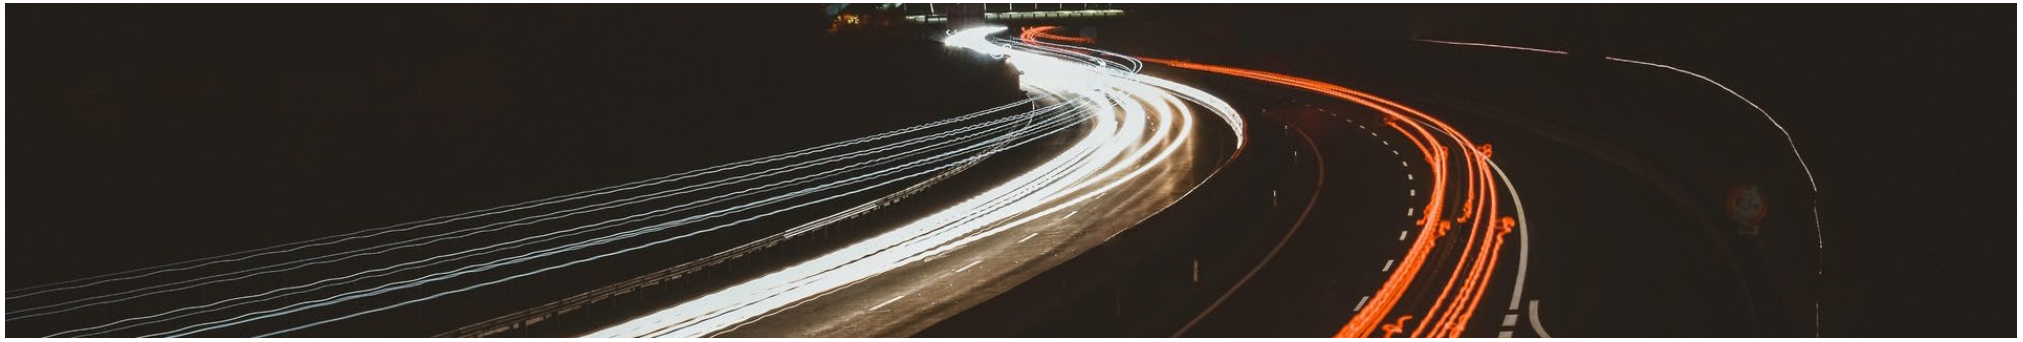

Please take a moment to rate this session.

Your feedback is important to us.

Transition Slide