




中国科学院计算技术研究所  
INSTITUTE OF COMPUTING TECHNOLOGY, CHINESE ACADEMY OF SCIENCES




中国科学院大学  
University of Chinese Academy of Sciences



# RLive: Robust Delivery System for Scaling Live Streaming Services

Yu Tian, Gerui Lv, Qinghua, Ruili Fang, Yajie Peng, Zhichen Xue, Rui Han,  
Chuanqing, Xiaofei Pang, Ri Lu

# Background: Higher Quality, Higher Pressure

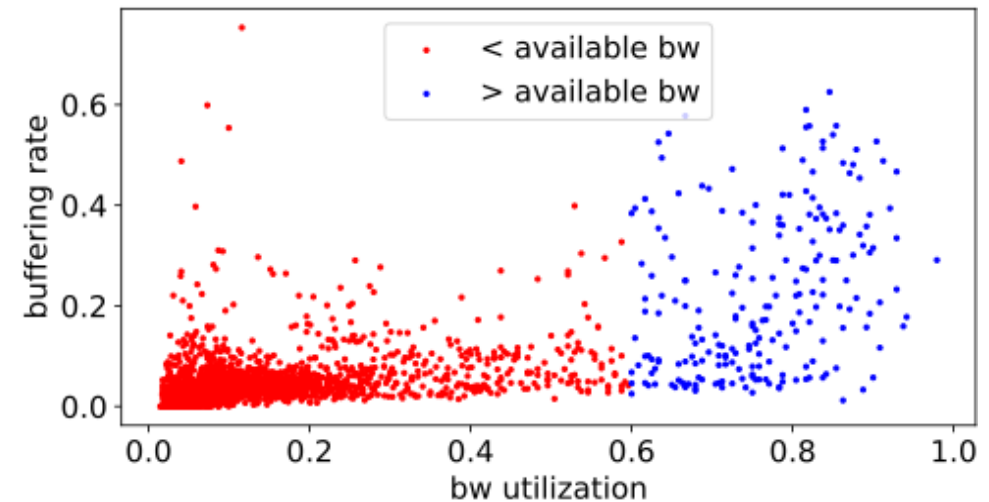


Higher Quality: Per-stream demand grows  $\sim 16\times$  from 480p to 4K.

	2022		For the Year Ended December 31,				
	RMB	%	2023		2024		
			RMB	%	RMB	US\$	%
<b>Cost of revenues:</b>							
Revenue-sharing costs	9,115,351	50.5%	9,507,483	55.6%	10,803,944	1,480,134	59.8%
Content costs	3,496,871	19.4%	3,195,620	18.7%	2,729,520	373,943	15.1%
Server and bandwidth costs	1,752,878	9.7%	1,477,116	8.7%	1,643,678	225,183	9.1%
IP derivatives and others	3,684,772	20.4%	2,905,903	17.0%	2,880,420	394,616	16.0%
<b>Total cost of revenues</b>	<b>18,049,872</b>	<b>100.0%</b>	<b>17,086,122</b>	<b>100.0%</b>	<b>18,057,562</b>	<b>2,473,876</b>	<b>100.0%</b>

Server and bandwidth costs for Bilibili in 2024: **164 million RMB**


Dedicated CDN doesn't scale cheaply.



Rush-Hour Bottlenecks Hurt QoE [1].

[1] Cost-Saving Streaming: Unlocking the Potential of Alternative Edge Node Resources, IMC24

# Background: Higher Quality, Higher Pressure




Higher Quality: Per-stream demand grows  $\sim 16\times$  from 480p to 4K.

	2022		For the Year Ended December 31,				
	RMB	%	2023		2024		
			RMB	%	RMB	US\$	%
<b>Cost of revenues:</b>							
Revenue-sharing costs	9,115,351	50.5%	9,507,483	55.6%	10,803,944	1,480,134	59.8%
Content costs	3,496,871	19.4%	3,195,620	18.7%	2,729,520	373,943	15.1%
Server and bandwidth costs	1,752,878	9.7%	1,477,116	8.7%	1,643,678	225,183	9.1%
IP derivatives and others	3,684,772	20.4%	2,905,903	17.0%	2,880,420	394,616	16.0%
<b>Total cost of revenues</b>	<b>18,049,872</b>	<b>100.0%</b>	<b>17,086,122</b>	<b>100.0%</b>	<b>18,057,562</b>	<b>2,473,876</b>	<b>100.0%</b>

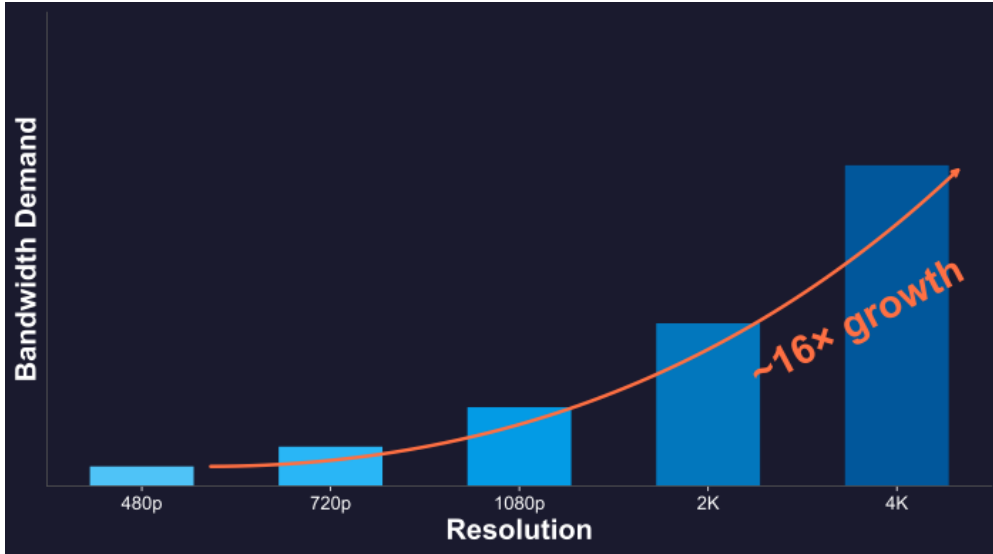
Server and bandwidth costs for Bilibili in 2024: **164 million RMB**

Need lower-cost bandwidth sources.



Rush-Hour Bottlenecks Hurt QoE.

# Background: Higher Quality, Higher Pressure

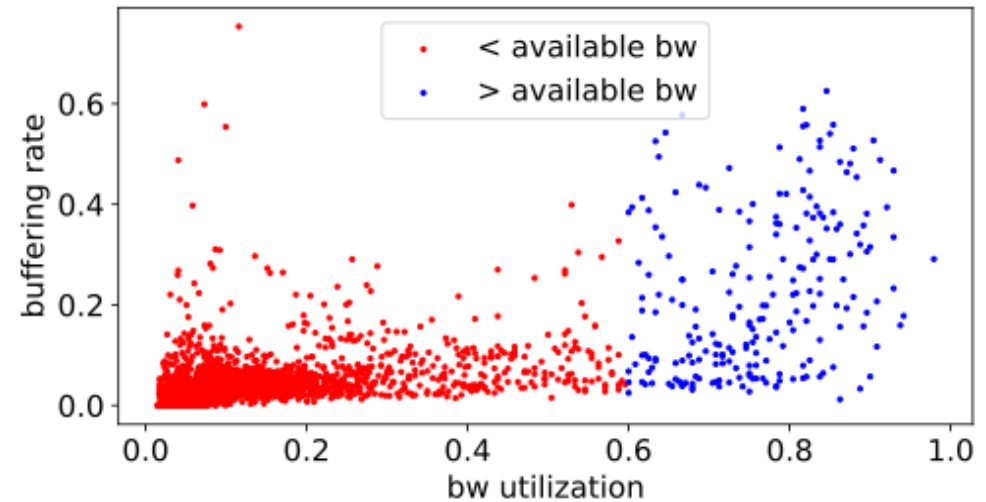


Higher Quality: Per-stream demand grows  $\sim 16\times$  from 480p to 4K.

	2022		For the Year Ended December 31,				
	RMB	%	2023		2024		
			RMB	%	RMB	US\$	%
<b>Cost of revenues:</b>							
Revenue-sharing costs	9,115,351	50.5%	9,507,483	55.6%	10,803,944	1,480,134	59.8%
Content costs	3,496,871	19.4%	3,195,620	18.7%	2,729,520	373,943	15.1%
Server and bandwidth costs	1,752,878	9.7%	1,477,116	8.7%	1,643,678	225,183	9.1%
IP derivatives and others	3,684,772	20.4%	2,905,903	17.0%	2,880,420	394,616	16.0%
<b>Total cost of revenues</b>	<b>18,049,872</b>	<b>100.0%</b>	<b>17,086,122</b>	<b>100.0%</b>	<b>18,057,562</b>	<b>2,473,876</b>	<b>100.0%</b>

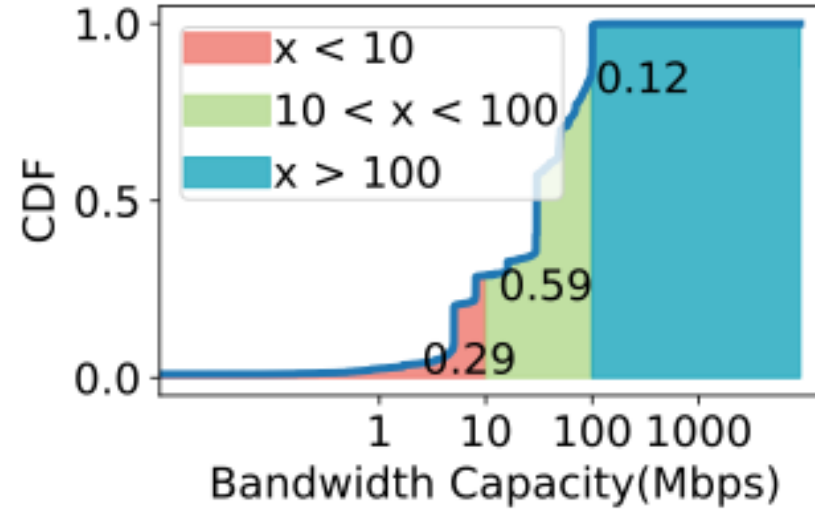
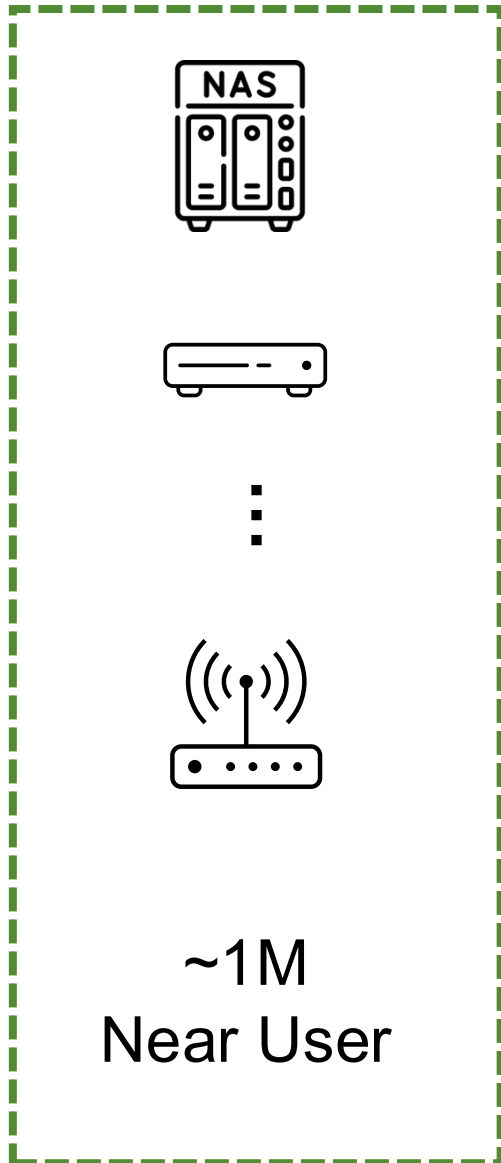
Server and bandwidth costs for Bilibili in 2024: **164 million RMB**

Need lower-cost bandwidth sources.

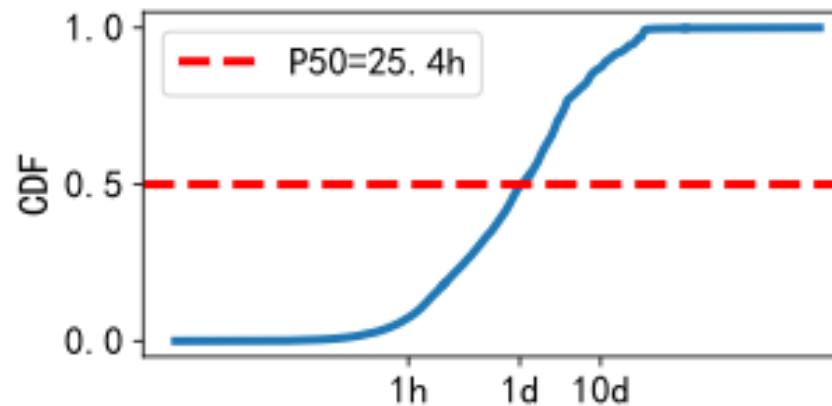


Need massive delivery capacity for traffic surges.

# Best-effort Edge Nodes: Cheap but Unreliable





Limited  
Bandwidth  
Capacity




Unstable  
Life Span

# Strawman Fails: Single-Source Doesn't Work

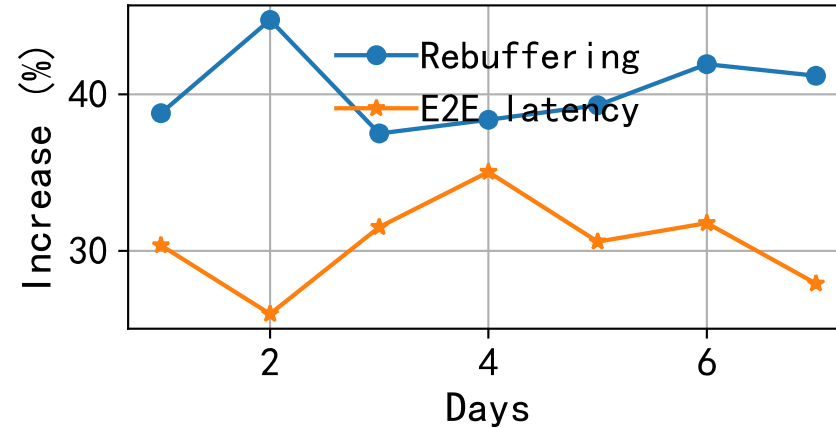
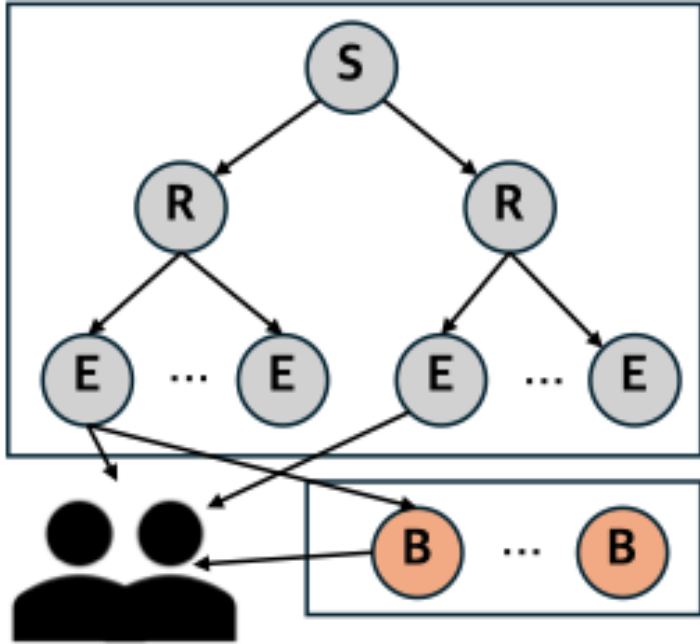


higher latency and more rebuffering.




< 0.1M additional clients.

# Strawman Fails: Single-Source Doesn't Work



higher latency and more rebuffering.



< 0.1M additional clients.


How to fully exploit all best-effort resources, including the weak ones, while delivering CDN-grade QoE?

# Key Insight: Multi-source Multi-Substream Transmission


## Source Failure



All subsequent frames are lost  
Slow recovery, poor experience



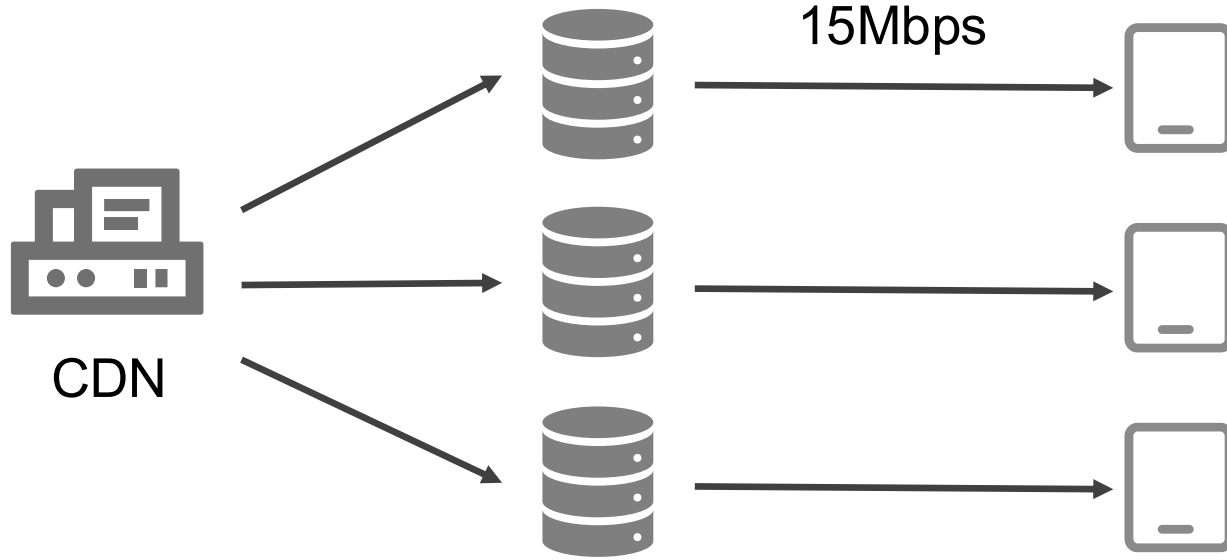
...



1/K lost frames  
Fast loss recovery

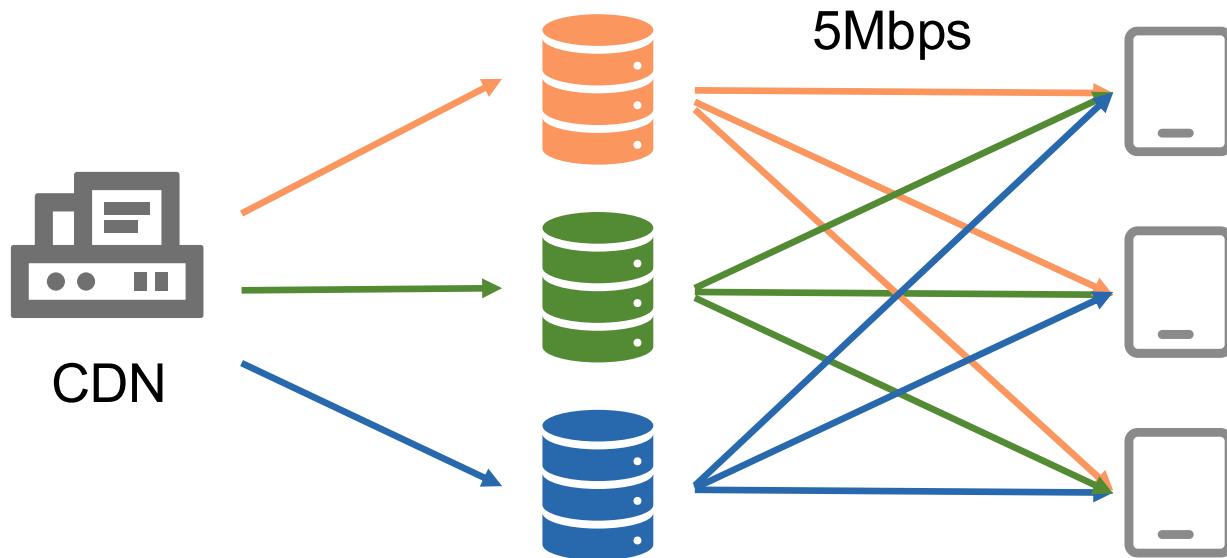
K sources

# Key Insight: Multi-source Multi-Substream Transmission



**Low Traffic Expansion Rate**

$$\gamma = \frac{\text{serving traffic}}{\text{backward traffic}}$$
$$= \frac{45}{45} = 1$$



**High Traffic Expansion Rate**

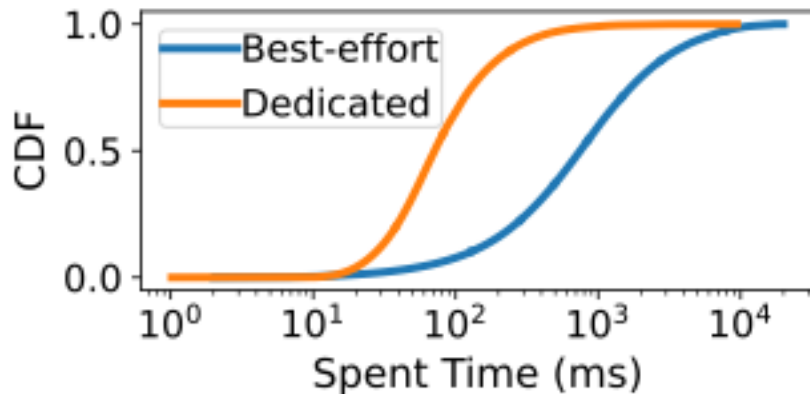
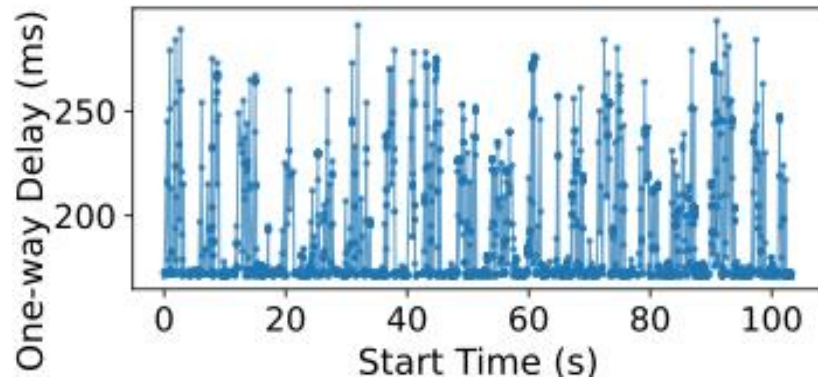
$$\gamma = \frac{45}{15} = 3$$

Save bandwidth, better scalability

# Design Challenges

**Table 1.** Live streaming service overview.

	6 am	12 pm	6 pm	12 am	max
#streams	~0.70M	~1.60M	~1.75M	~1.38M	~2.47M
#nodes	~0.9M	~0.95M	~1.0M	~1.02M	~1.05M



Scalable user-to-node mapping

**Collaborative** Control

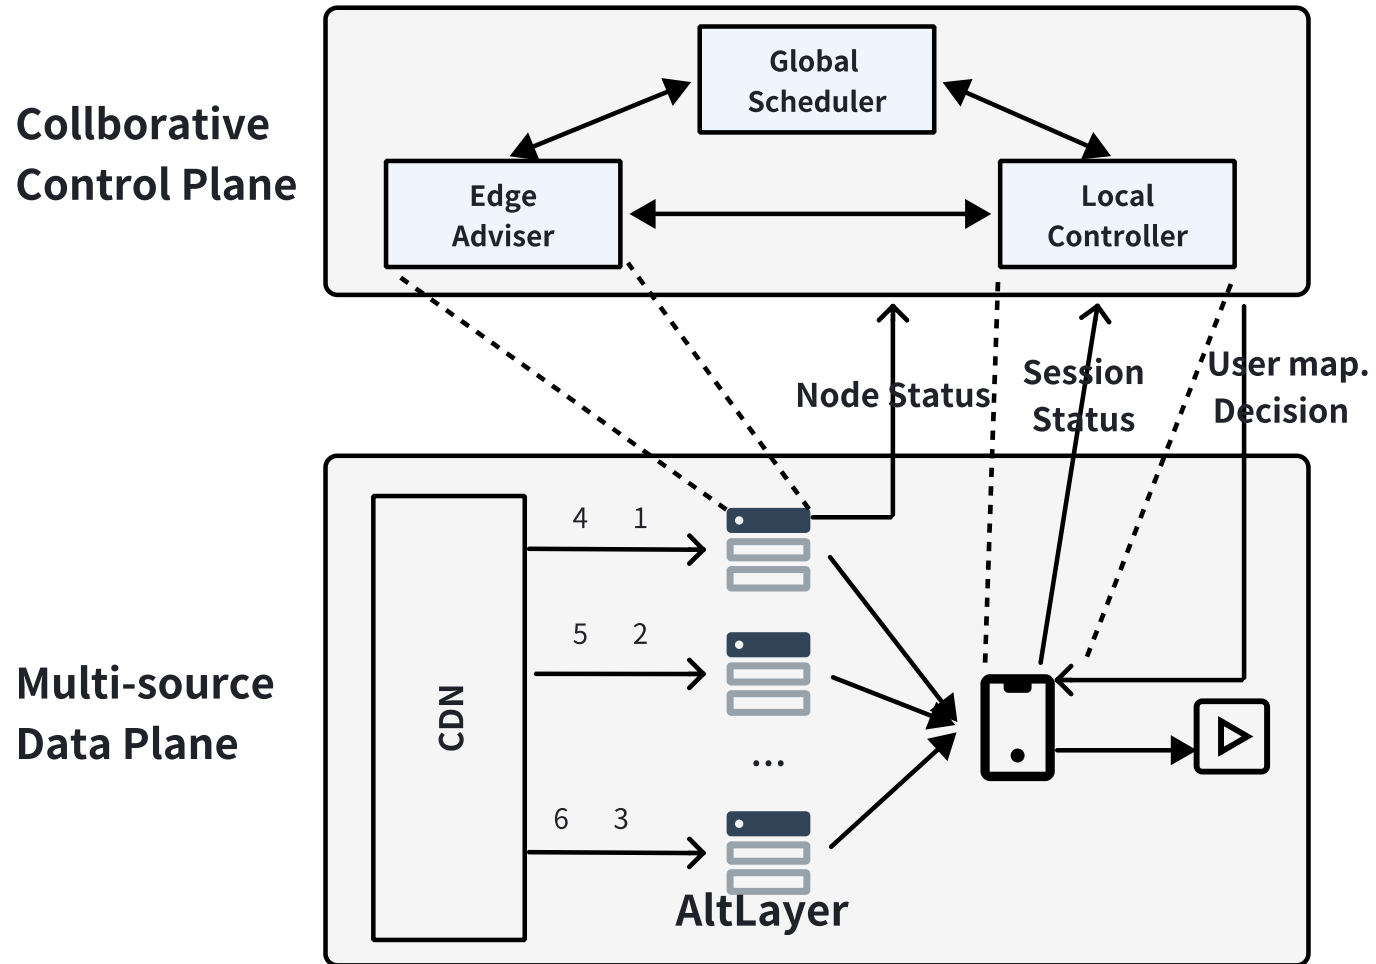
Fast frame re-ordering

**Distributed** Frame Sequencing

Cost-effective loss recovery

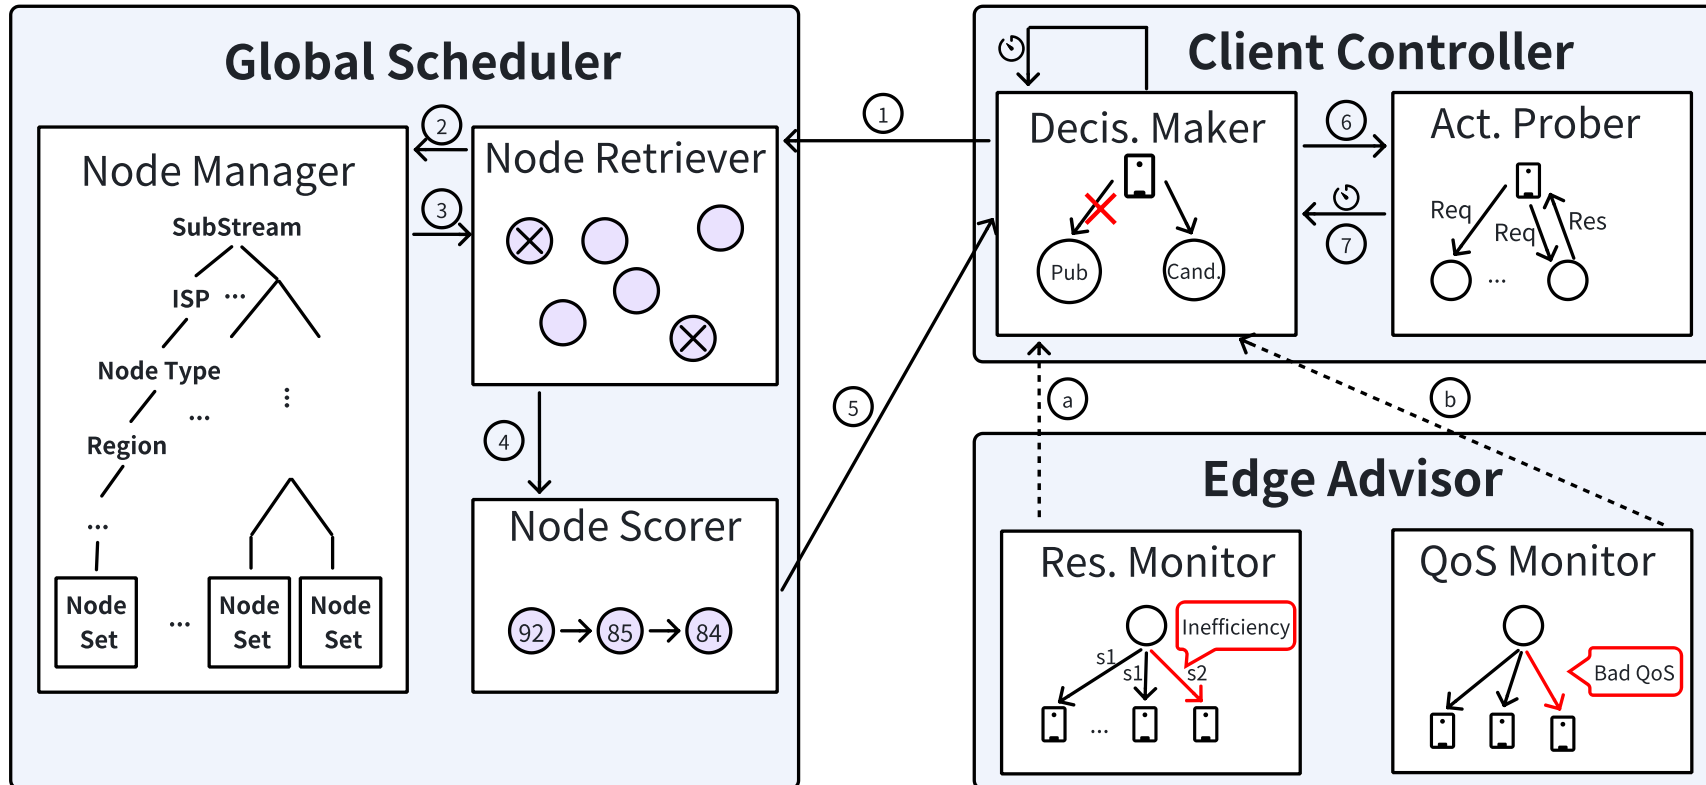
**QoE-driven** Loss Recovery

# Rlive System Overview



# Multi-Granularity Collaborative Control

- Three controllers, three timescales, one coordinated decision.

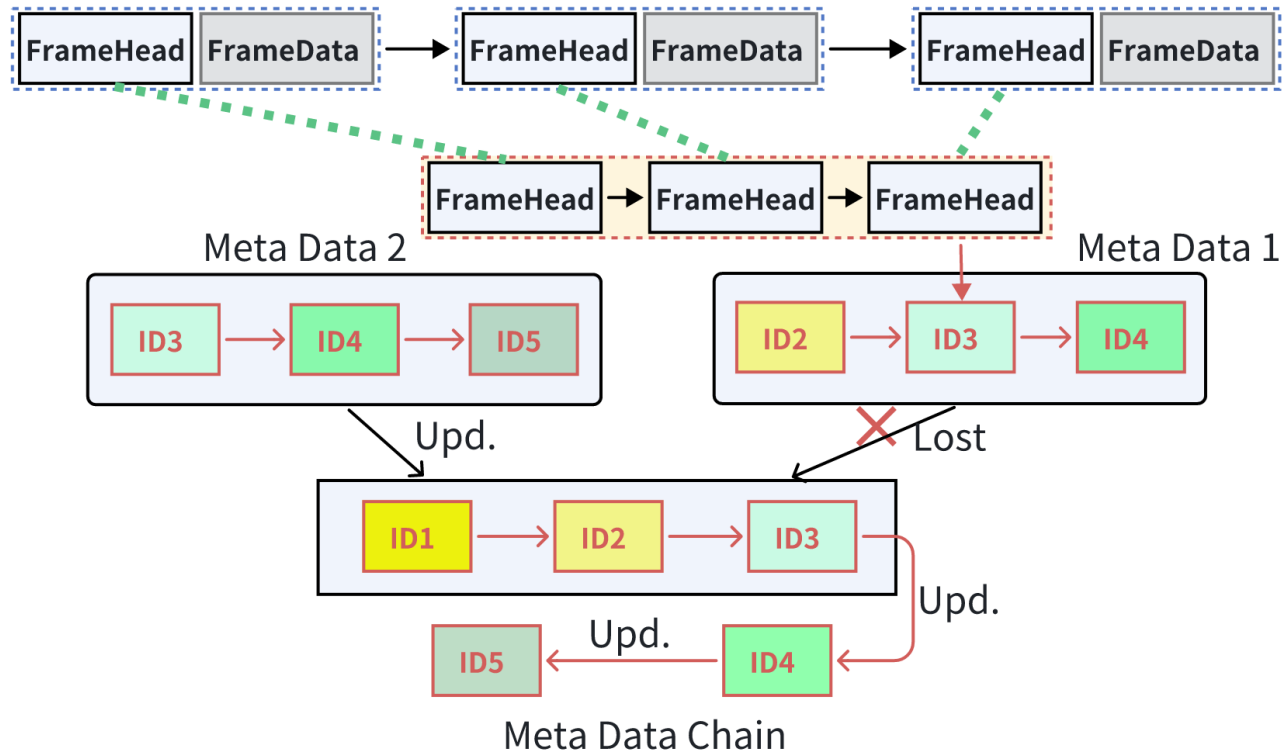


Node recommendation based on global view (second-level updates)

Probing-based real-time node selection (millisecond-level observations)

Proactive QoS and cost-aware feedback

# Distributed Frame Sequencing




Best-effort node independently generates **local frame chains**


embedded directly into **data packets**

combine them into a **global frame chain** for correct playback order

# Evaluation: A/B Test



Rebuffering times per hundred seconds.



Video bitrate.

- Deployed in ByteDance production CDN for **14 days**, covering **billions** of views daily
- **3× capacity expansion** by fully leveraging heterogeneous best-effort edge nodes
- **15% fewer rebuffering events, 11% higher bitrate** through robust multi-source data transmission

# Lessons Learned

## ❖ Multi-Layer Feedback over Centralized Optimization

- ▶ Global explore-exploit balancing + local blacklisting of failing nodes outperforms purely centralized control.

## ❖ Embed Control Metadata in the Data Path

- ▶ Co-locating frame ordering information with data packets eliminates centralized metadata bottlenecks and single points of failure.

## ❖ Favor Resilience over Optimality at Scale

- ▶ In hyperscale heterogeneous environments, simple and robust strategies consistently outperform complex precise optimizations.

# Conclusion

## ❖ Collaborative Control Plane

- ▶ Three-layer coordination (global, client, edge) for scalable and responsive user-to-node mapping

## ❖ Redundancy-Free Multi-Source Data Plane

- ▶ Distributed frame sequencing + QoE-driven loss recovery for robust delivery

## ❖ Deployed in ByteDance CDN for 3+ years

- ▶ **3×** capacity expansion, **15%** fewer rebuffering events, **11%** higher bitrate

Unreliable edge resources, when orchestrated properly, can deliver CDN-grade live streaming at scale.



中国科学院计算技术研究所  
INSTITUTE OF COMPUTING TECHNOLOGY, CHINESE ACADEMY OF SCIENCES



中国科学院大学  
University of Chinese Academy of Sciences



**Thanks for listening / Q&A**  
[tianyu21b@ict.ac.cn](mailto:tianyu21b@ict.ac.cn)