

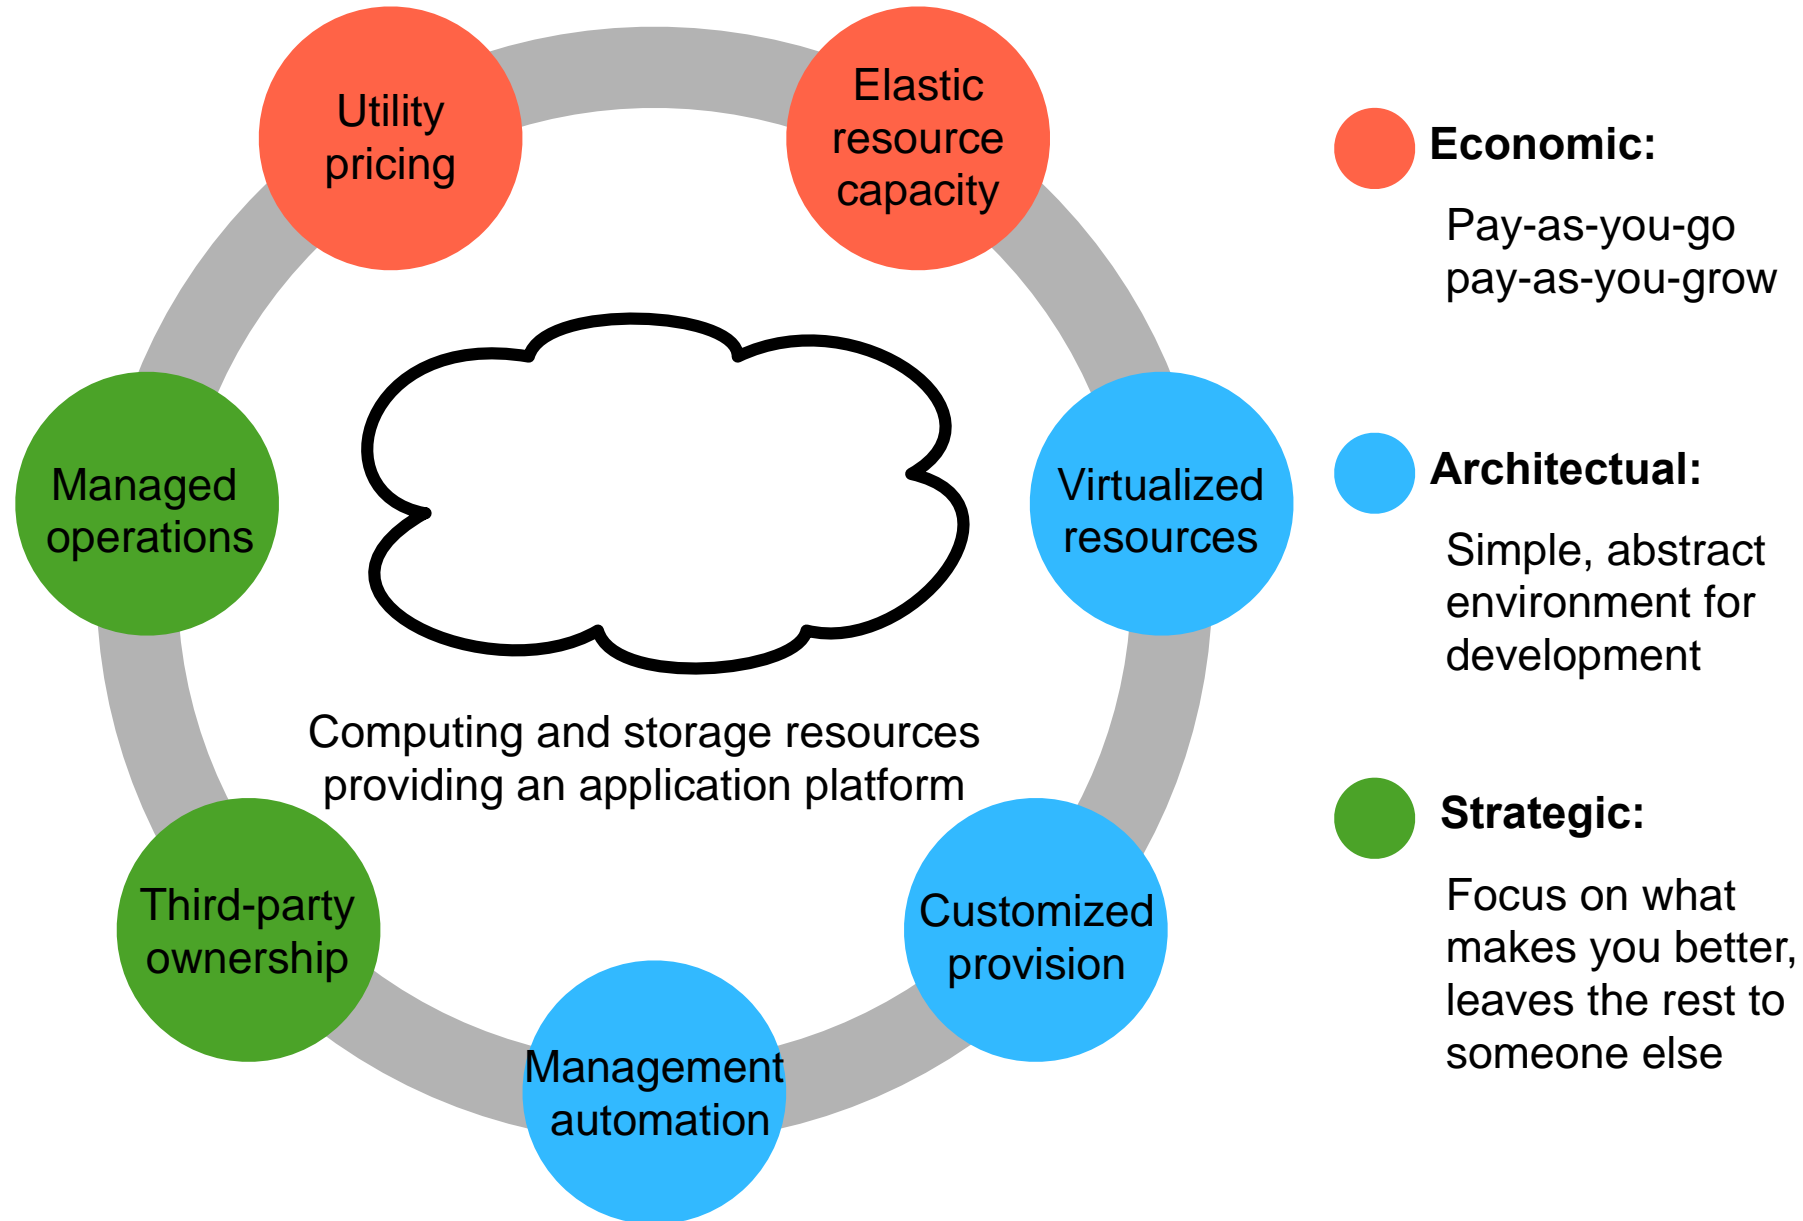
The Internet-based Computing and Datacenter Networks

Bo Li 李波, Professor

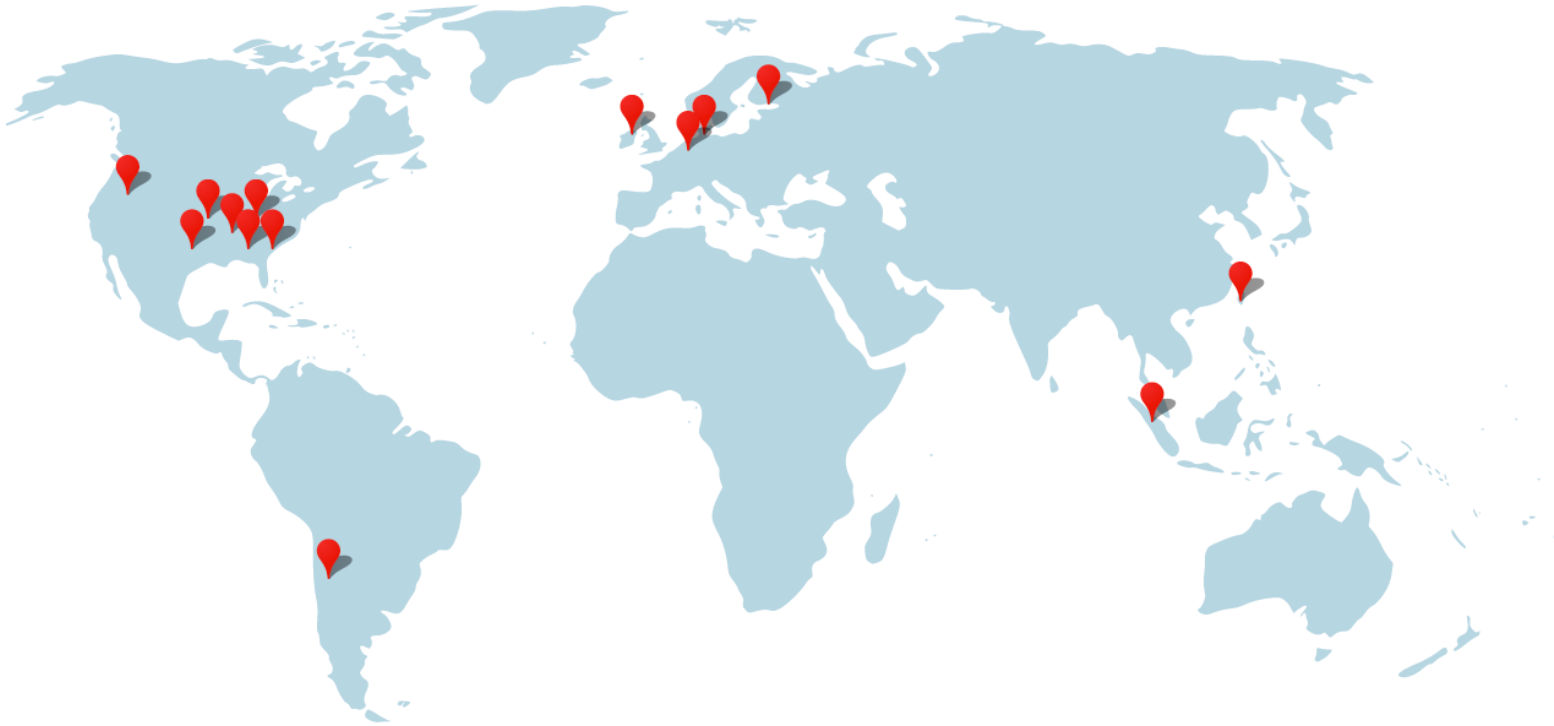
Research Background and Interests

- I have been with HKUST for a “while”
- Research coverages:
 - Networking: from wireline (LAN-WAN) to wireless, from infrastructure to device, from applications to link-level transmission, from sensing to data processing, and etc.
- Major contributions:
 - Internet video broadcast – Test-of-Time Paper Award from IEEE INFOCOM 2015
 - Fellow of IEEE
 - State Natural Science Award (2nd class) 国家自然科学基金二等奖
 - Cheung Kong Scholar 长江学者
 - Five Best Paper Awards from IEEE
- Recent research interests
 - Internet-based computing – cloud computing
 - Datacenter networking

Cloud and Internet-based Computing - Paradigm




Datacenters (Google's View)



- 14 datacenters around the world, >10 Gbps link
- Running about 2,376,640 (estimated) servers in early 2013
- Supporting web search and Youtube video, etc


Inside a Datacenter Network



- ❖ Clos topology – dense connection
 - ❖ Many equal cost paths for any source and destination
- ❖ Research problems:
 - ❖ Routing and traffic engineering:
 - How to distribute network traffic among all the available paths?
 - ❖ Transport layer protocol:
 - Design data center customized congestion control algorithms
 - ❖ Dynamic Network control:
 - Update forwarding rules to ensure policy consistency

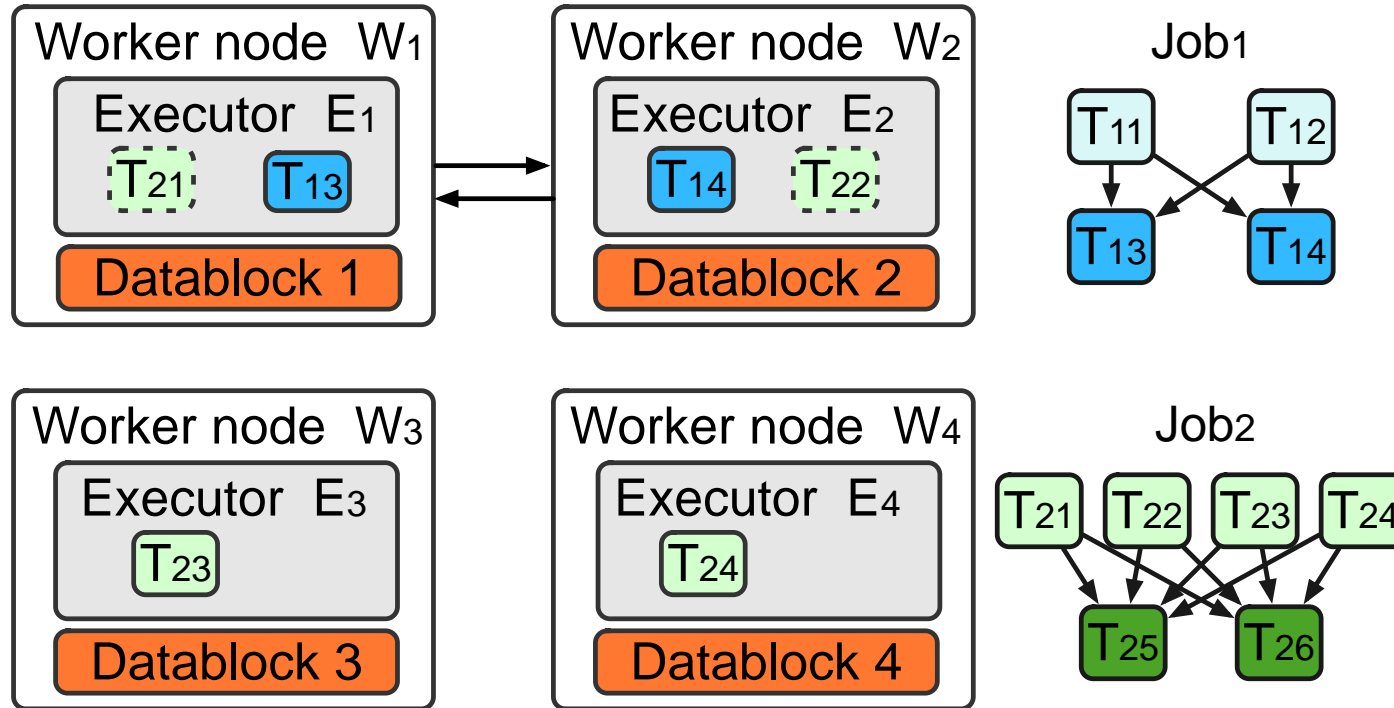
Data-parallel computing

- ❖ Use cases:
 - ❖ Web search queries
 - ❖ Iterative machine learning algorithms (e.g. PageRank)
 - ❖ Recommendation algorithms for social networks and e-Commerce
- ❖ Popular data-parallel computing platforms: Spark, Hadoop
- ❖ Consume a large amount of resources
 - ❖ CPU, memory, storage
- ❖ Research problems:
 - ❖ How to efficiently allocate network resources?
 - ❖ Flow scheduling
 - ❖ How to efficiently allocate computation resources?
 - ❖ Task placement



Scheduling for data-parallel tasks

❖ How to place tasks to servers in and across datacenters



❖ Objective

- ❖ High utilization of computation and network resources
- ❖ Short job completion time
- ❖ Low WAN link capacity

Network Control - Software-Defined networking (SDN)

- ❖ Motivation:
 - ❖ Decouple control plane and data plane
 - ❖ Dynamic and consistent control
- ❖ Challenges:
 - ❖ Reduce control overhead
 - ❖ Physically distributed controllers
 - ❖ Dynamic network control:
 - ❖ Consistently update network policies without temporary congestion
 - ❖ Flexible bandwidth allocation
 - ❖ Ensure high network utilization
 - ❖ Reduce latency
 - ❖ Application-level performance

