Extreme Cloud

Rick Schlichting

Executive Director, Software and Distributed Systems Research

AT&T Labs

Florham Park, NJ 07932, US

Cautionary Language Concerning Forward-Looking Statements

Information set forth in this document contains forward-looking statements and evaluation analysis that are subject to risks and uncertainties. AT&T disclaims any obligation to update or revise statements contained in this document based on new information, developments, or otherwise.

Special thanks to Mary Fernandez and KK Ramakrishnan

What should the next-generation Cloud eco-system look like?

Cloud?

Eco-system?

Extreme Cloud

Next-generation?

Should?

Services/applications?

Eco-system/infrastructure?

Scenario: Multimodal Interfaces

Scenario: Multimodal Interfaces

Scenario: Collaborative Services

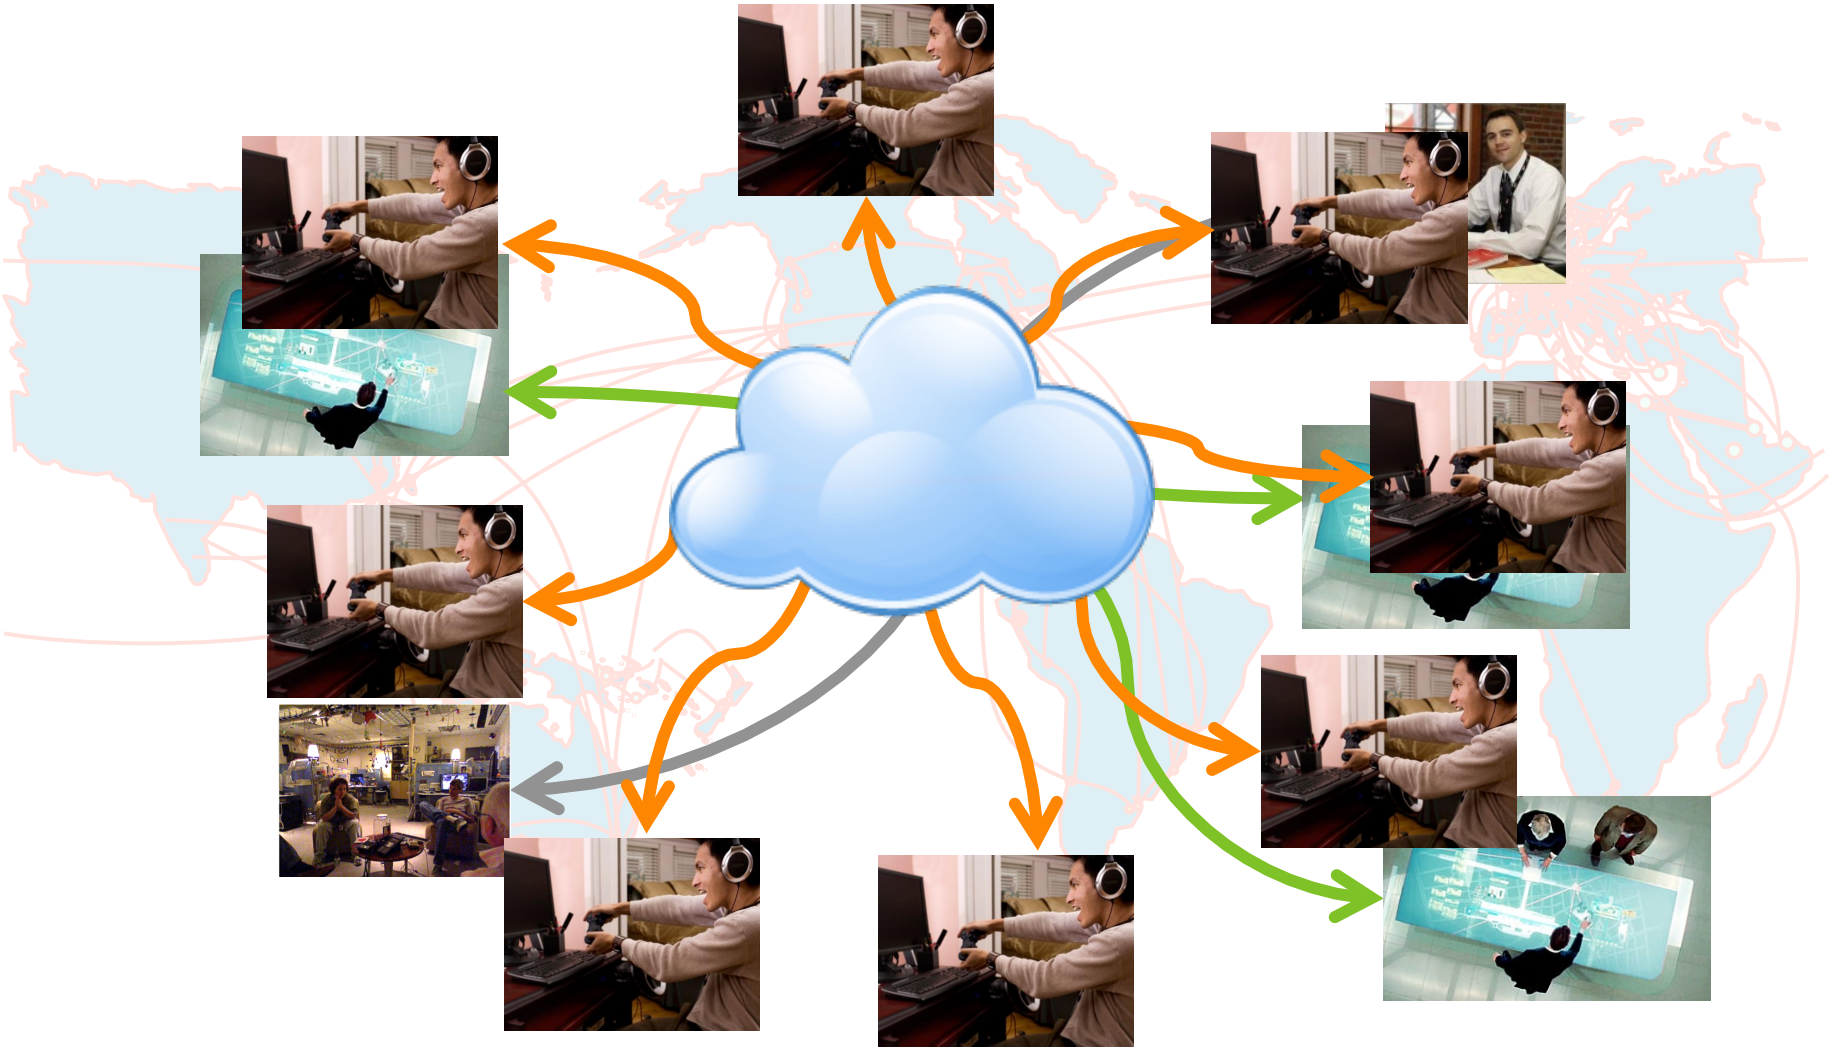

Scenario: Personal Proxies

Virtual desktops (VDI)
Virtual Individual Servers (VIS)

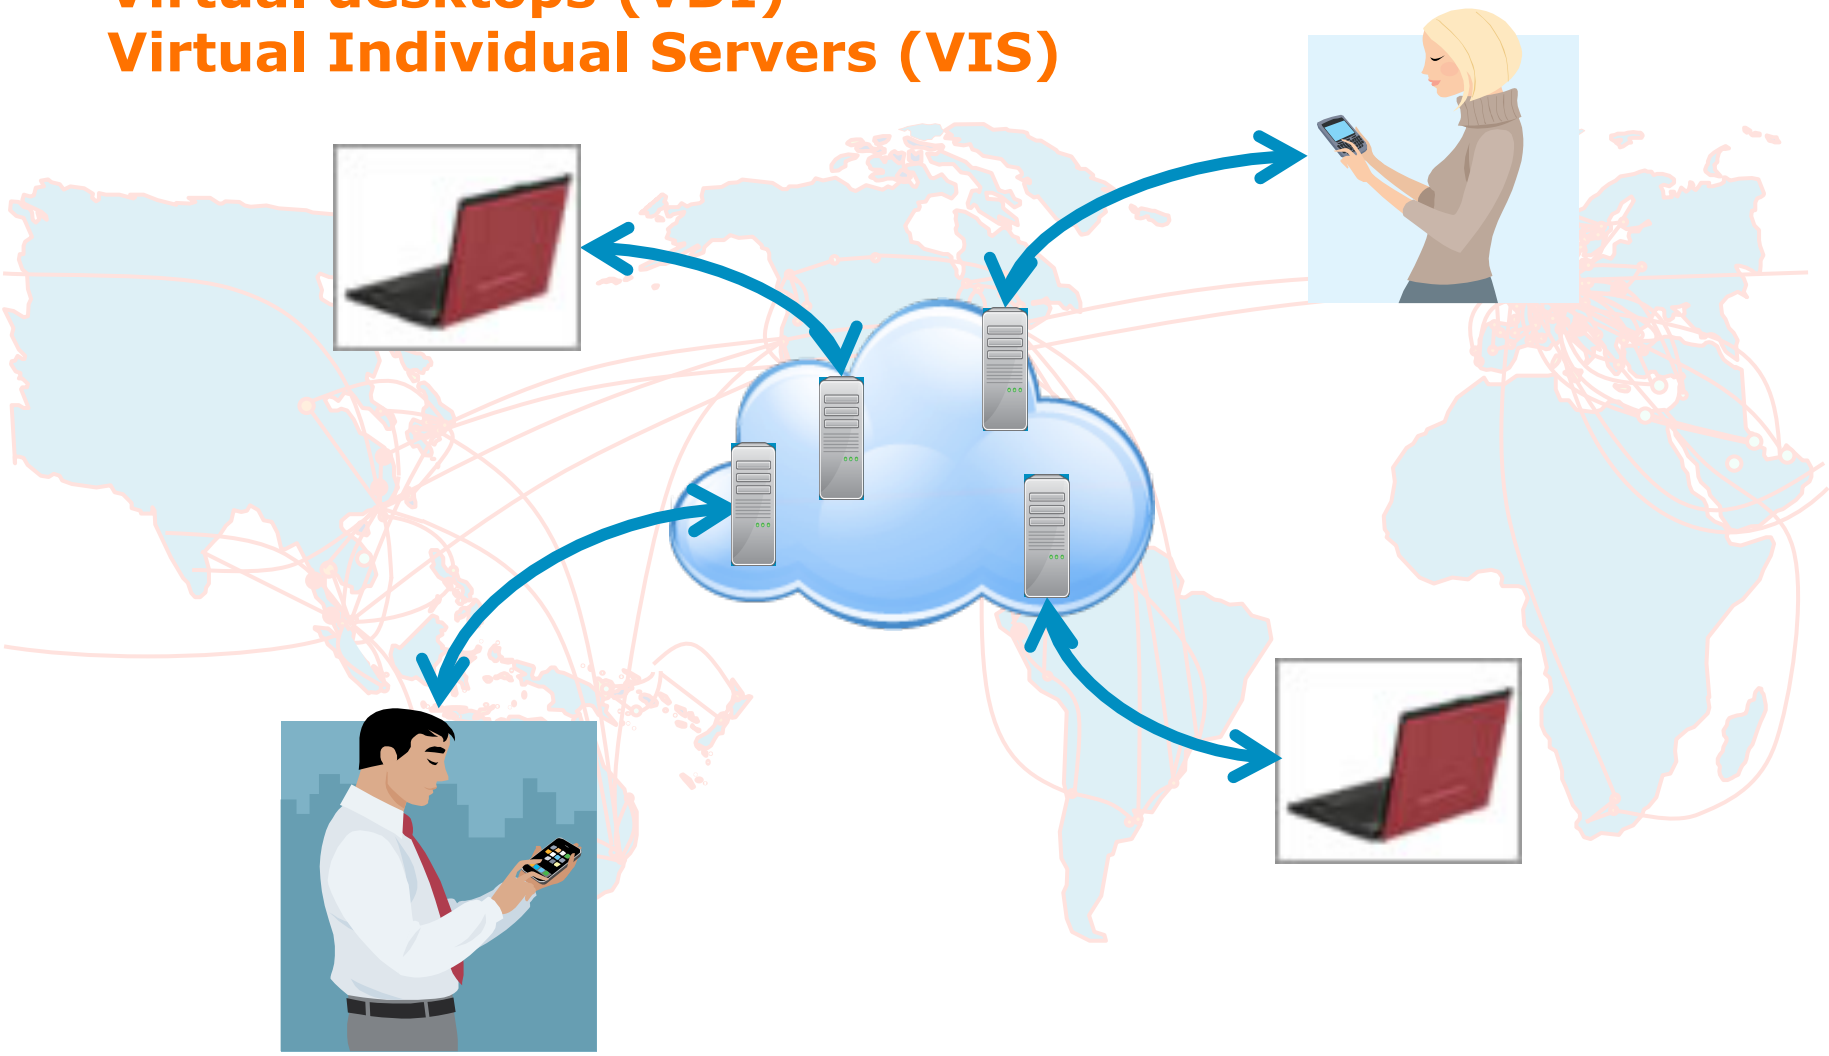

Scenario: Location-Based Services

Air Graffiti™

Two fundamental aspects:

1.Associate text, pictures, music, video with a location

2. Browse what has already been added to the location

Scenario: Optimizing Mobile Content Delivery

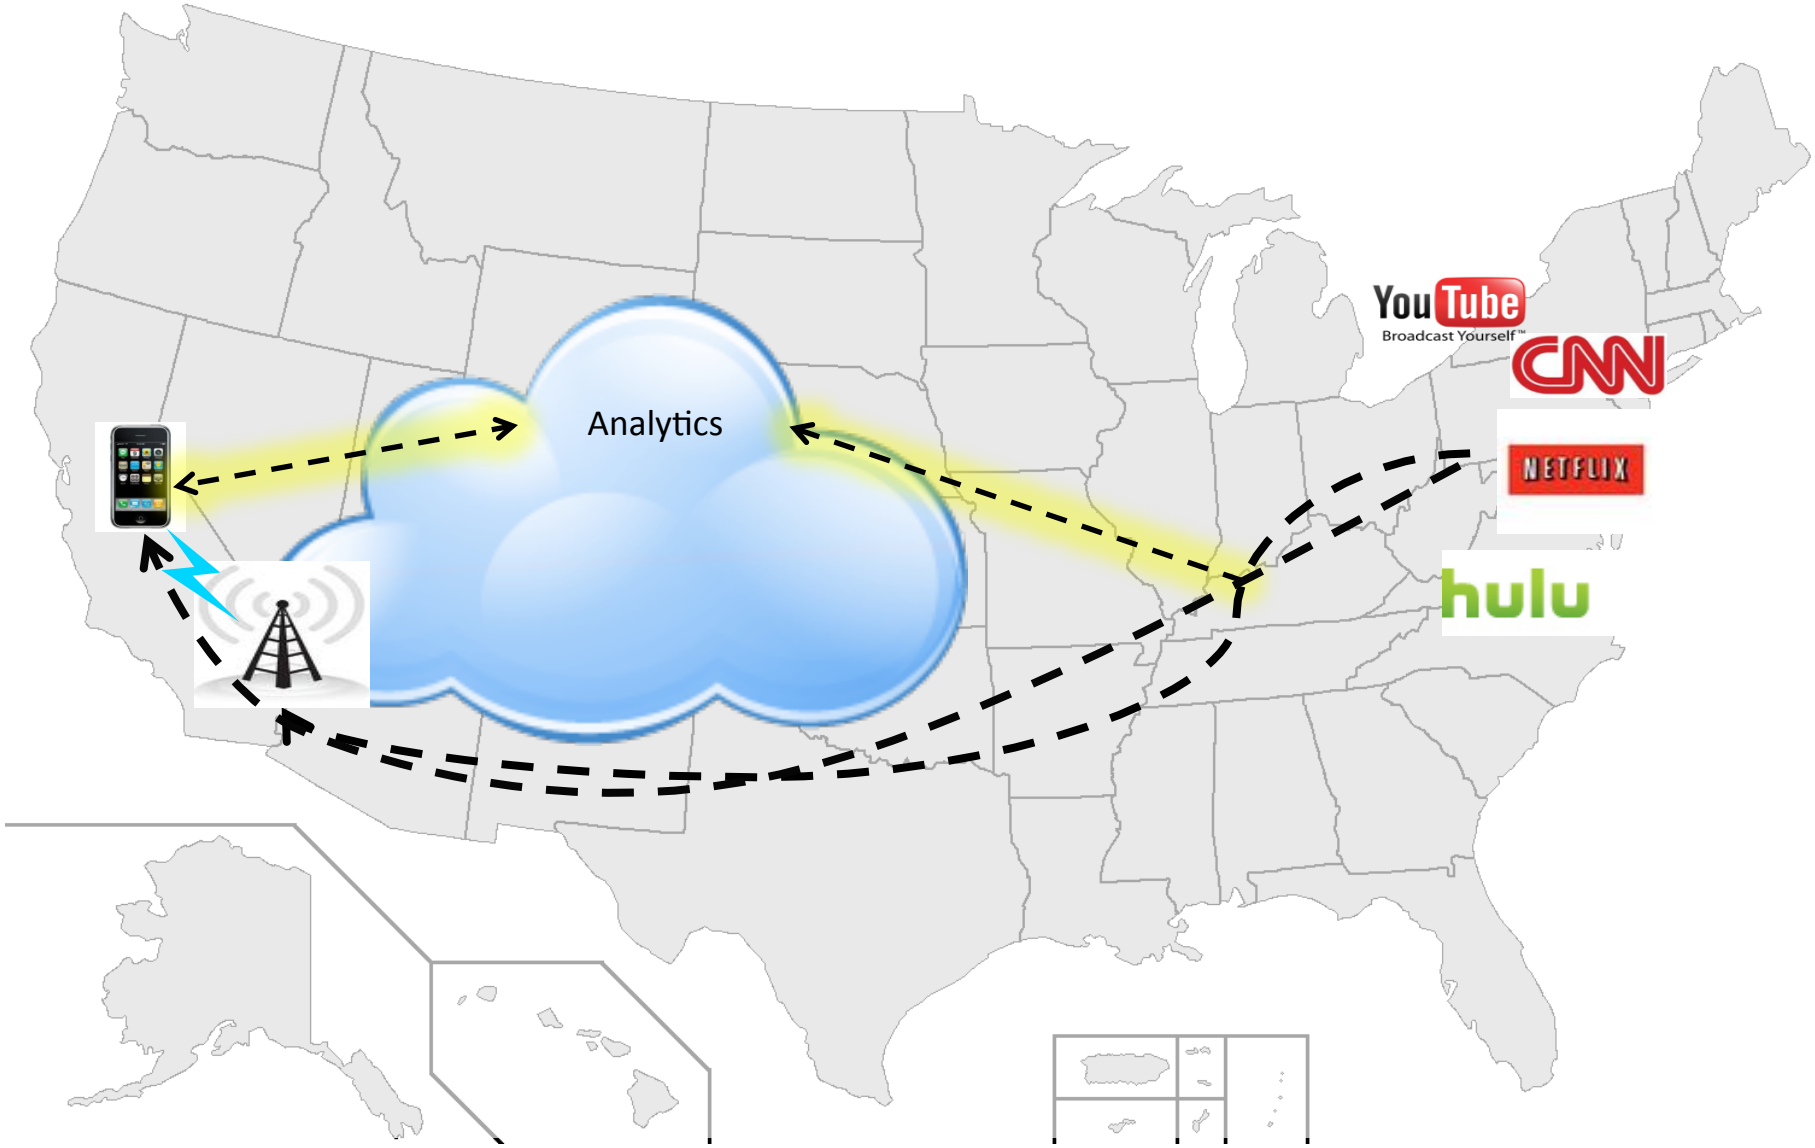

Observations

Mobile endpoints

Offloaded computation

Thin clients

Ubiquitous access

Different types of connectivity

Low latency/responsiveness

Other QoS properties

Local consumption of locally-produced data

Heterogeneous resources

…

Extreme Cloud

Ubiquitous

- Seamless access to storage and computational resources anywhere and at any time from any type of device

Heterogeneous

- Structured from different types of computational, storage and network resources

Enhanced properties

- Provides properties such as responsiveness and location awareness.
- Supports services with QoS requirements in the areas of low latency, high throughput, high availability, …

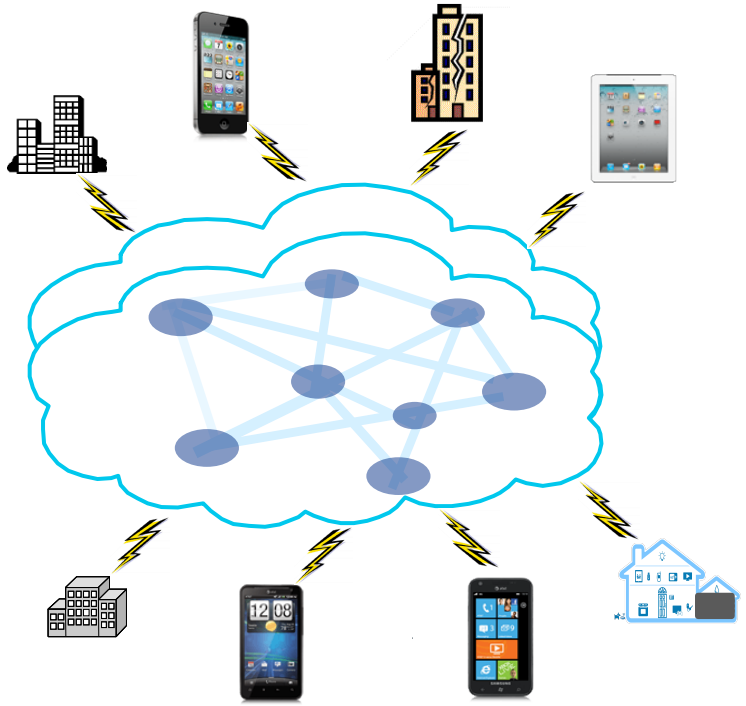

Extreme Cloud puts cloud resources "into" devices, "at" locations

Extremely responsive, Extremely dependable, Extremely secure, ….

Realization

Strategically locate cloud resources in the network, making it possible to deploy new services that are responsive and location-sensitive

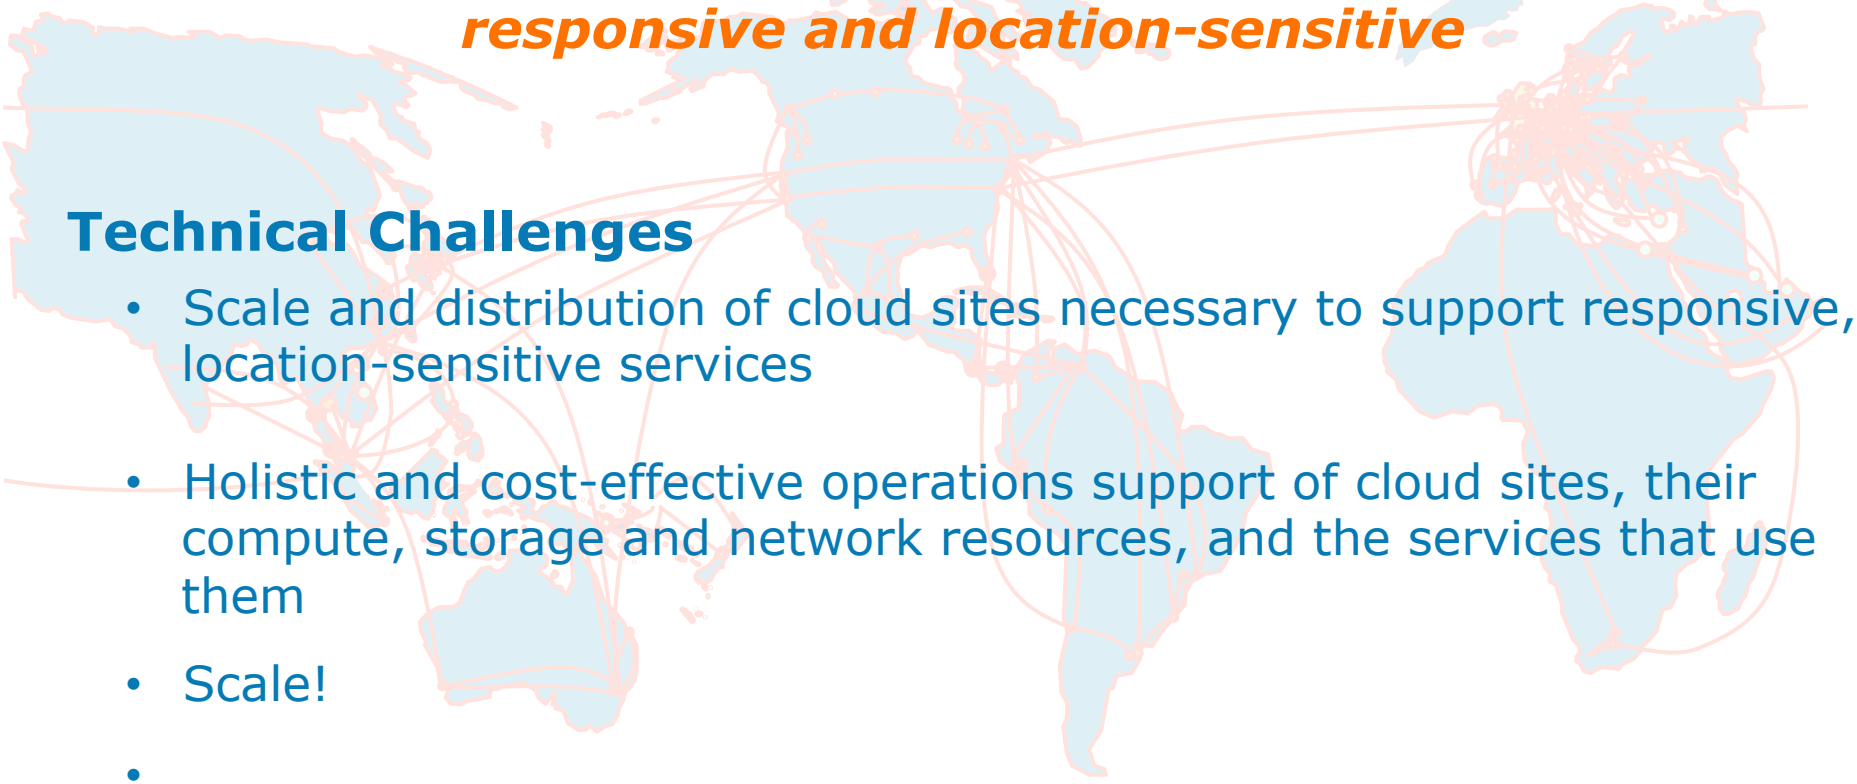

Technical Challenges

- Scale and distribution of cloud sites necessary to support responsive, location-sensitive services
- Holistic and cost-effective operations support of cloud sites, their compute, storage and network resources, and the services that use them
- Scale!
- …

Many of these challenges relate to dependability and QoS in general!

Extreme Cloud Research Program

Building, operating, and using a large-scale, network-centric, highly decentralized cloud infrastructure

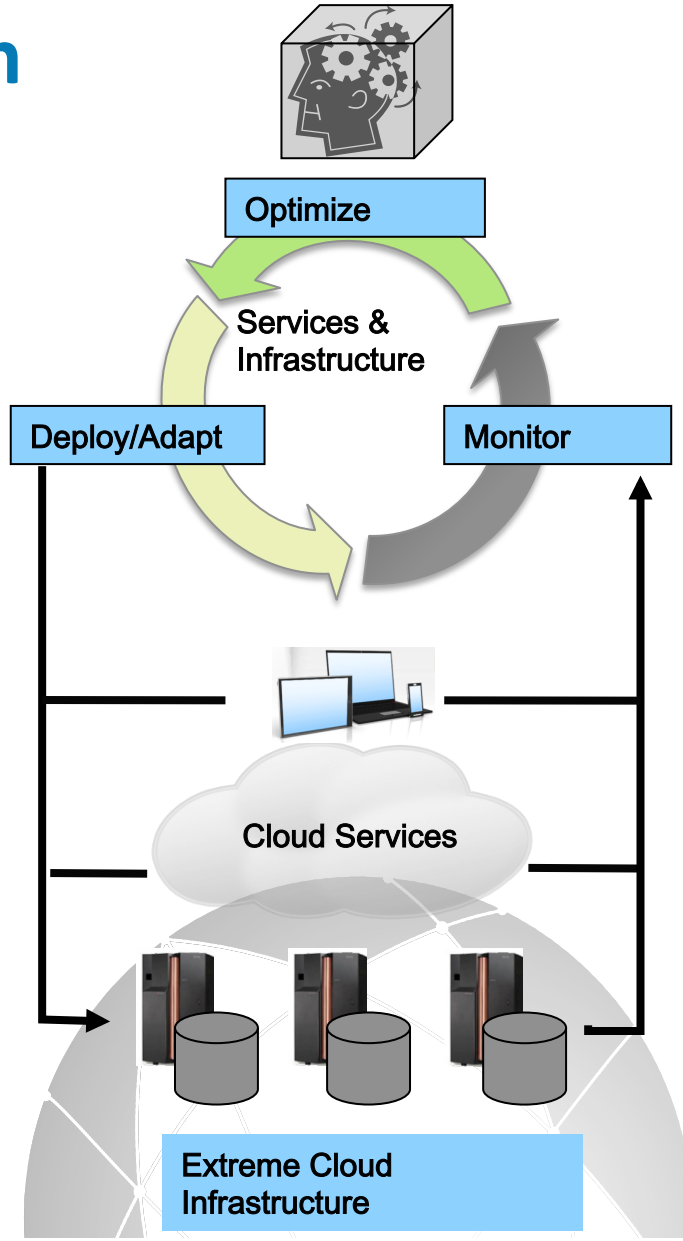

Enhanced infrastructure

Providing enhanced properties

- Distributed storage (security, dependability), PipeCloud (dependability), Self-Service Cloud (security)

Cloud control plane

Holistic control of the Extreme Cloud

- Orchestration: Tropic
- Control loop and optimization: Mistral

Cloud operating systems

For building cloud services

Example applications and services

Enhancing Cloud Technology

Three projects that aim to improve Cloud technologies in different ways.

TROPIC → **Cloud orchestration**
Transactional manipulation of compute, storage, and network resources

Mistral → **Optimized Cloud resource management**
Savings from improved utilization.

PipeCloud → **Cloud-based disaster recovery**
Fast remote DR with consistency guarantees

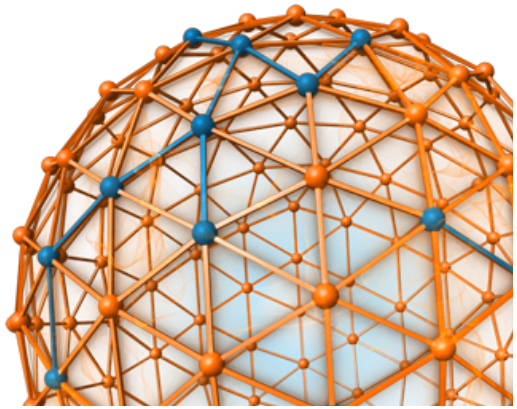

TROPIC: Transactional Cloud Orchestration

(C. Liu, Y. Mao, X. Chen, M. Fernandez, B.T. Loo, J. Van der Merwe, USENIX ATC, 2012, with U. Penn)

***Orchestration* involves provisioning, configuring, and decommissioning virtual resources across a distributed set of physical resources**

Challenges

- Seamless integration of cloud resources and VPNs → ***CloudNet***
- Resource provisioning tools tightly coupled to individual vendors' devices
- No interoperability; Low-level control interfaces

TROPIC is platform for ***defining, provisioning, managing*** cloud services

- Manages distributed virtual & physical resources holistically
- Supports migration, mobility of services across network enabling latency-sensitive, location-aware services, e.g., gaming

TROPIC cloud services are

- ***Robust***: Transactional recovery from mis-configuration, device errors, etc.
- ***Safe***: Enforces service & engineering constraints
- ***Highly available***: Hot fail-over of controller for cloud service continuity

TROPIC: Architecture

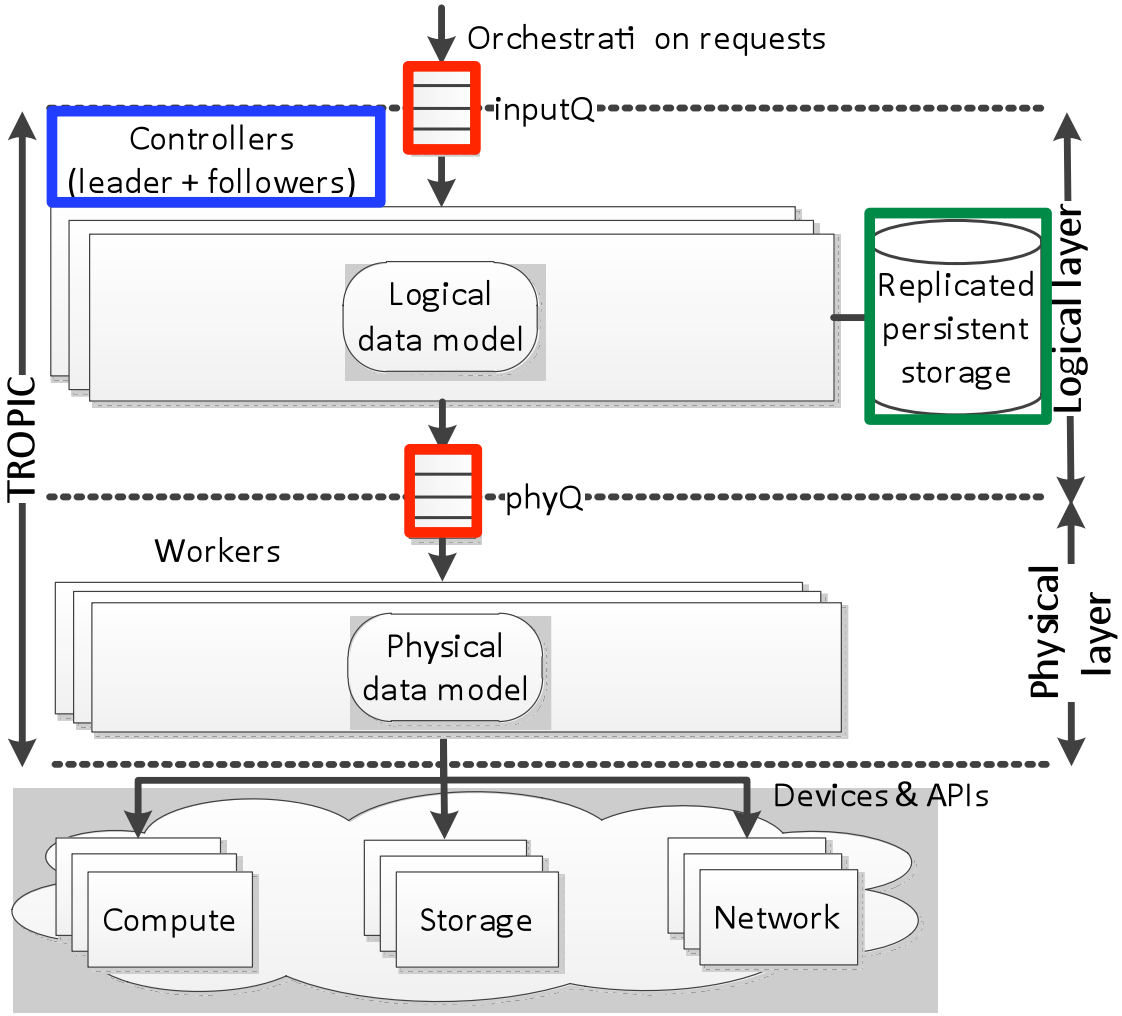

Logical / physical layer

- Replica: weak, eventual consistency
- Transaction manager: scheduling, simulation, concurrency control

Replicated components

- Multiple controllers
- Distributed queues
- Persistent storage

18

Cross-Site Service Management

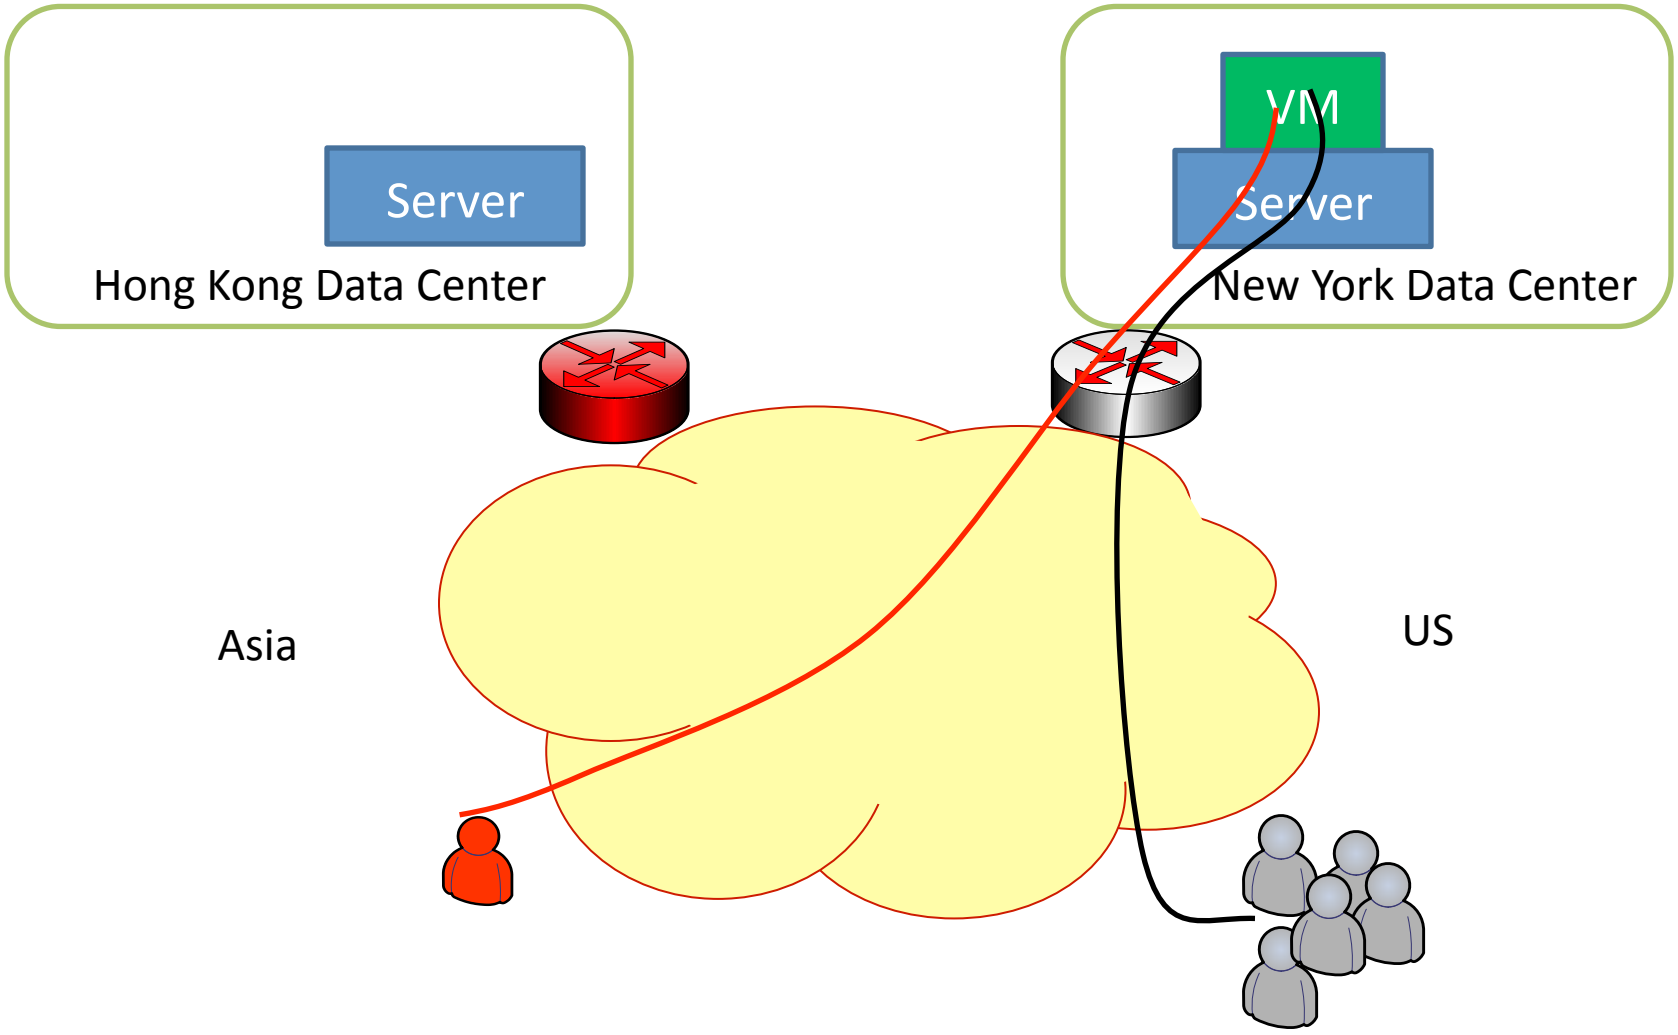

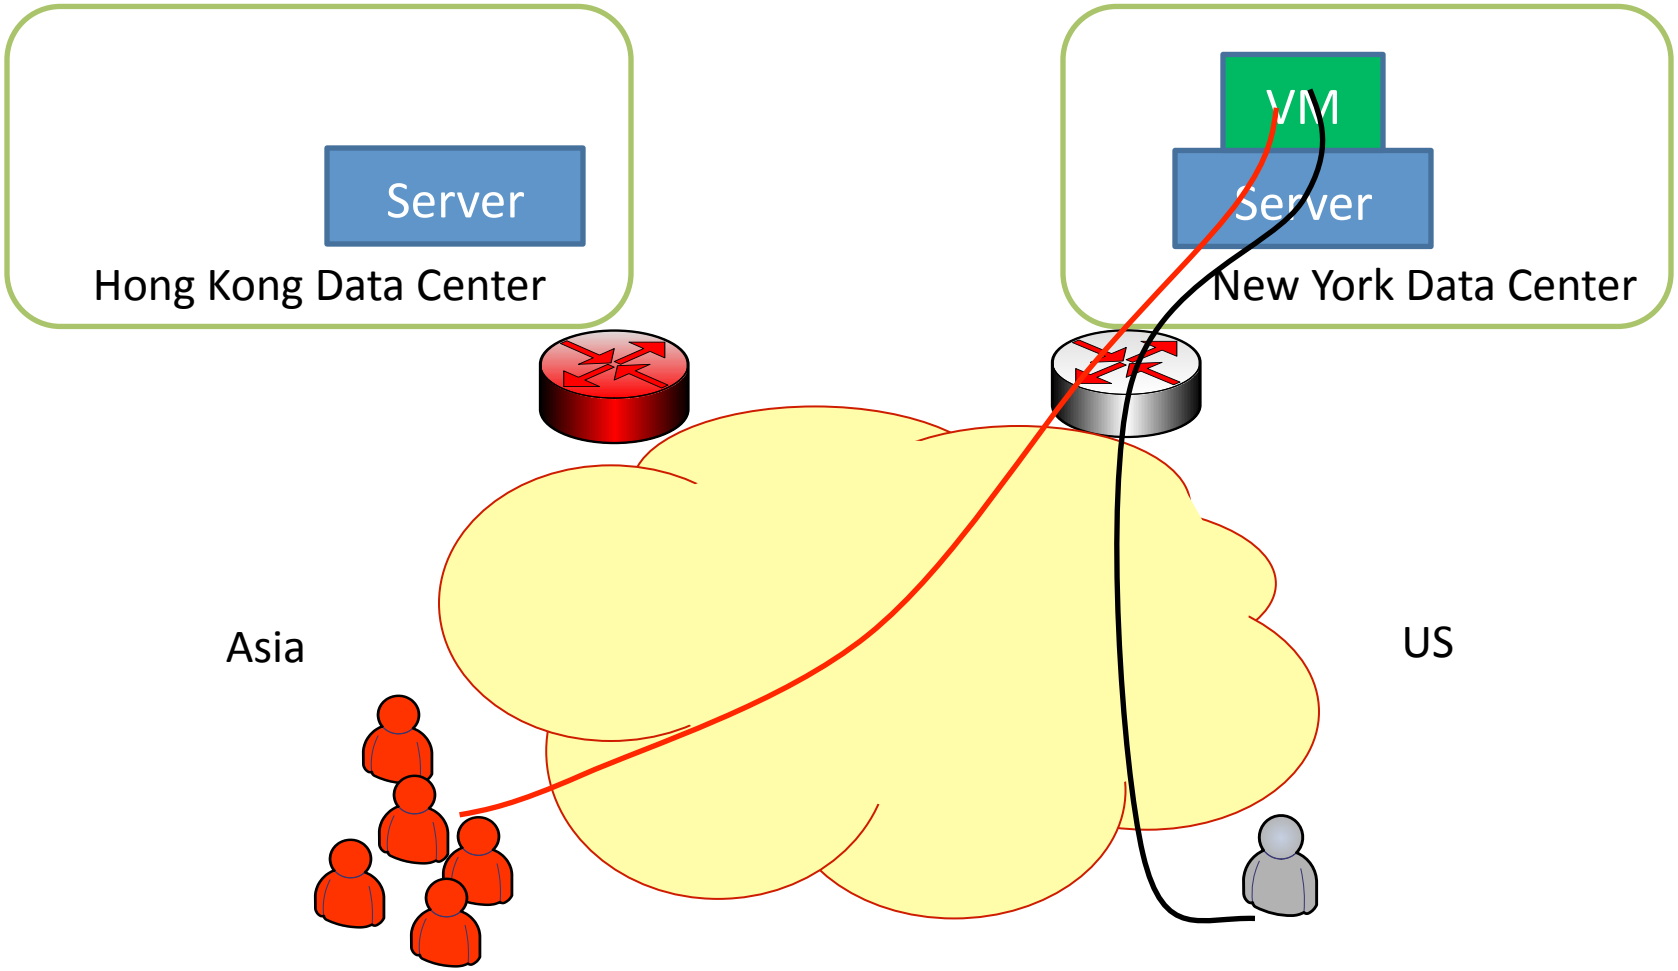

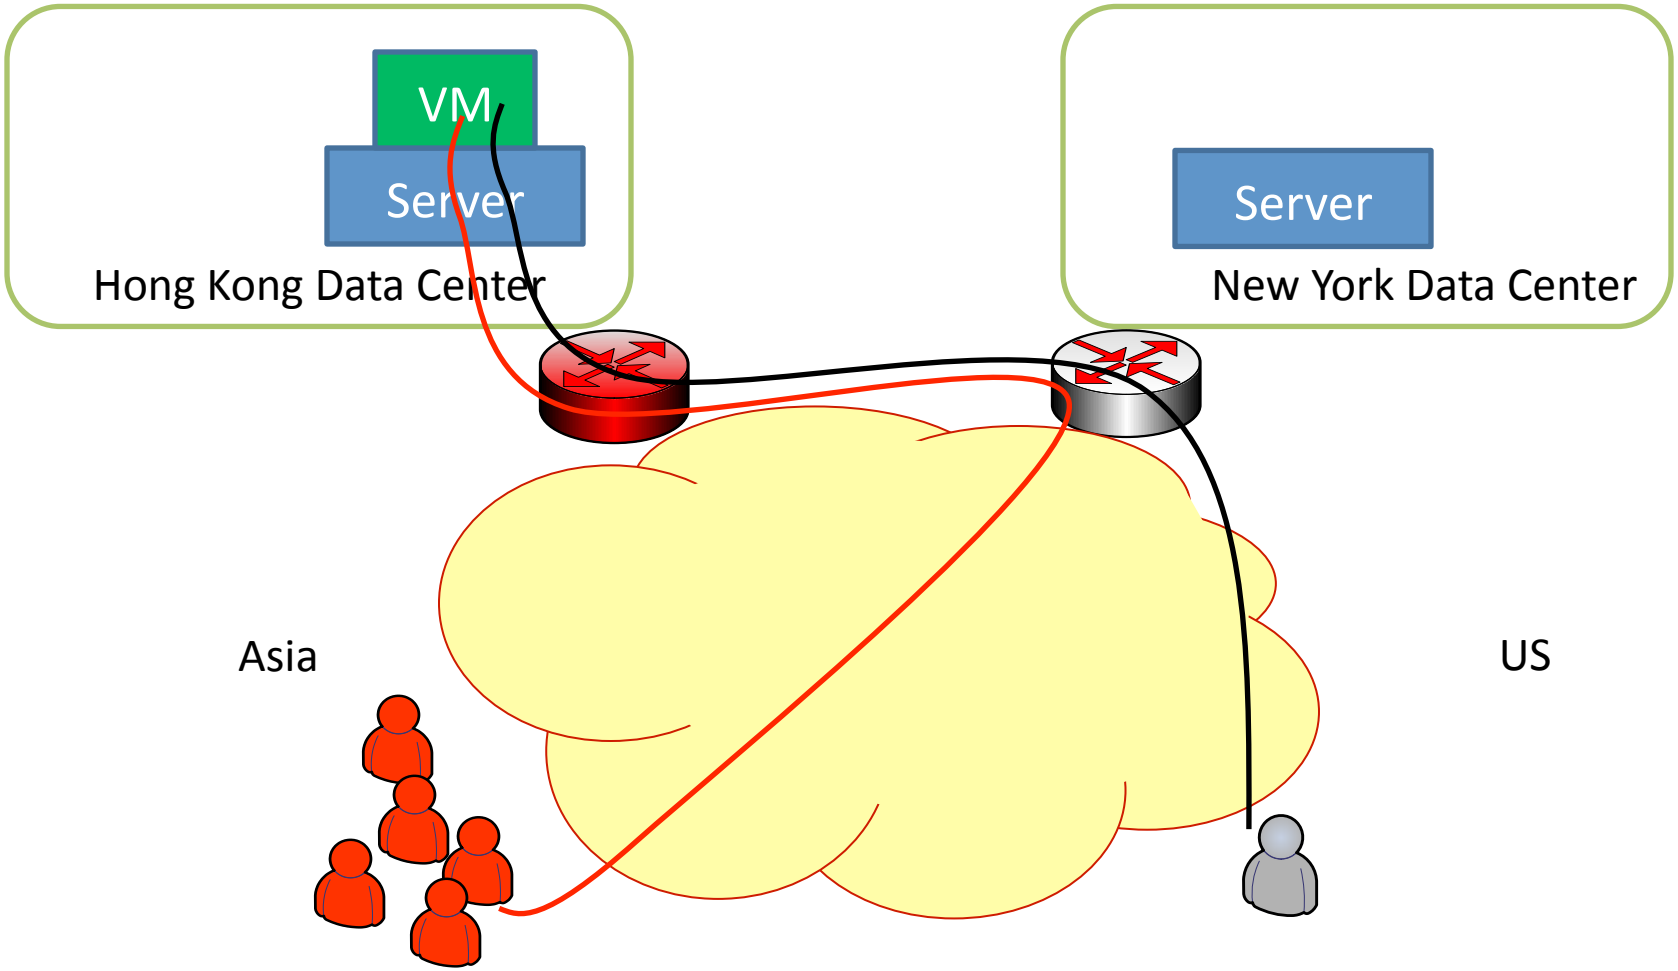

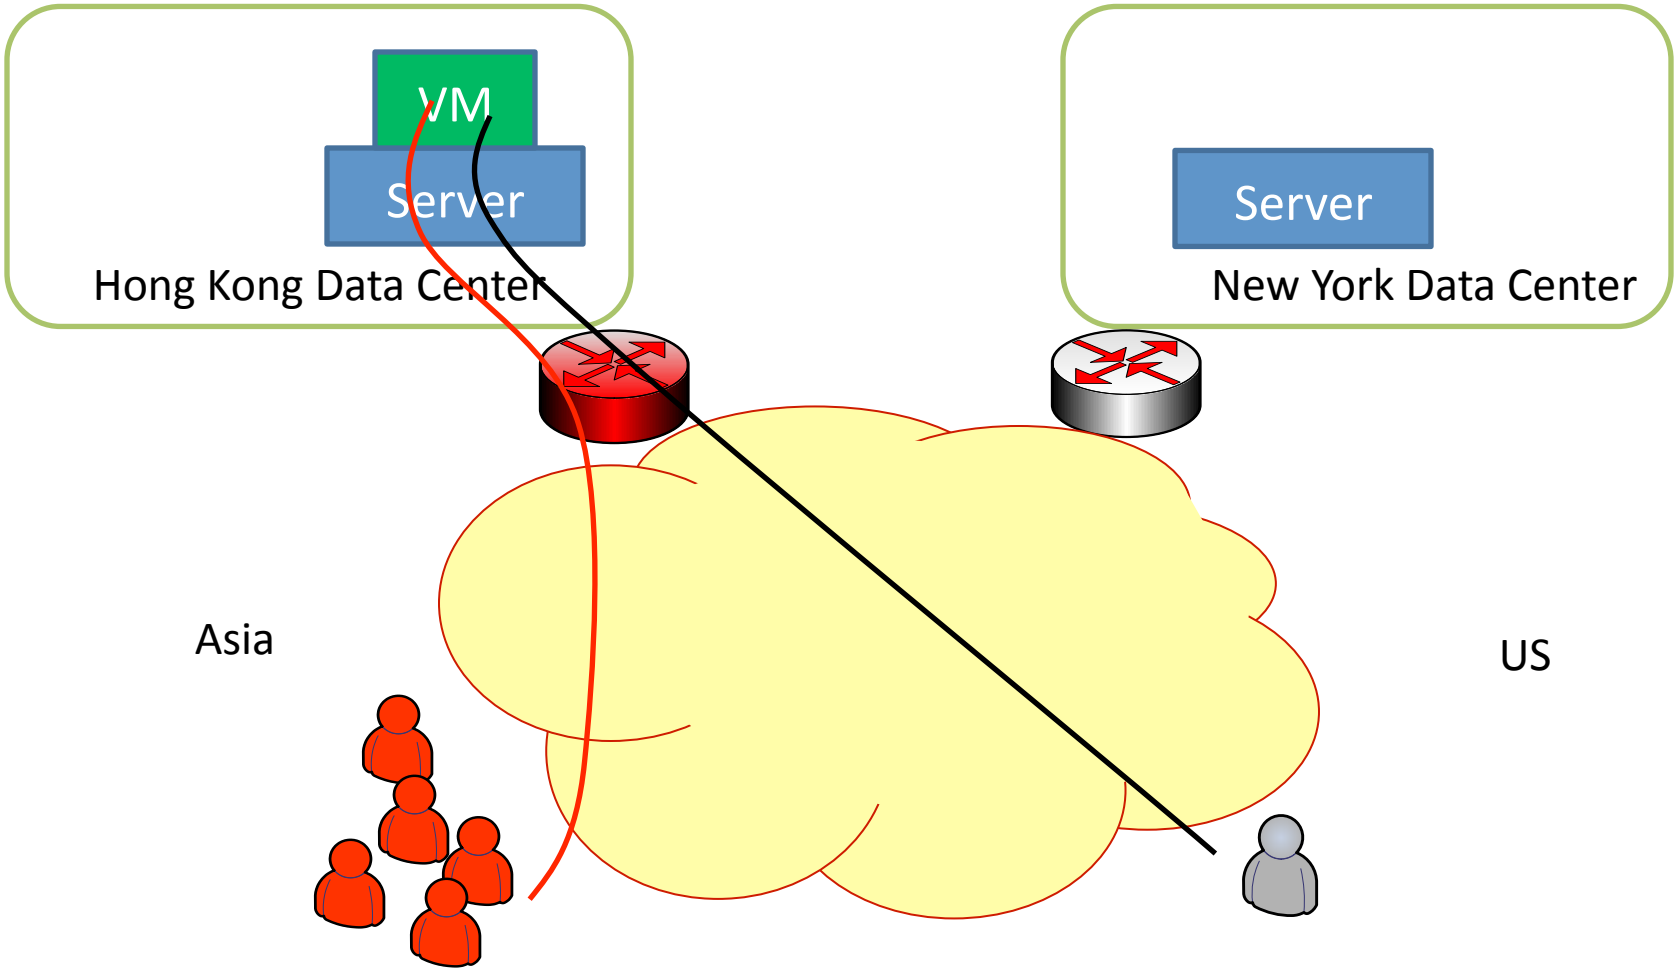
VM
Server
Hong Kong Data Center
Server
New York Data Center
Asia
US

TROPIC: Key Points

Addresses the control plane issues of assigning virtual to physical resources and orchestrating changes

Provides a high level of abstraction:

- A data model to deal with heterogeneity of resources
- A domain-specific language for specifying constraints and actions
- Transactional execution semantics
- Provides useful guarantees for cloud services, while still allowing highly concurrent operation

Architecture makes TROPIC highly available and scalable

➨ Currently working with the OpenStack community to integrate transactional features into the open source platform.

Mistral: An Optimizing Control Plane

(G. Jung, M. Hiltunen, K. Joshi, R. Schlichting, C. Pu. ICDCS 2010, with Georgia Tech)

Issue: Managing resources in a cloud to meet response time SLAs and maximize utilization

Cloud provider's point of view:

- How to maximize utility while minimizing costs (including power) while hosting numerous applications.

Cloud user's point of view:

- How to minimize cost for my application while meeting users' response time expectations.

Focus on generating **adaptation policies** for multi-tier enterprise applications.

Techniques: modeling, optimization, planning, prediction.

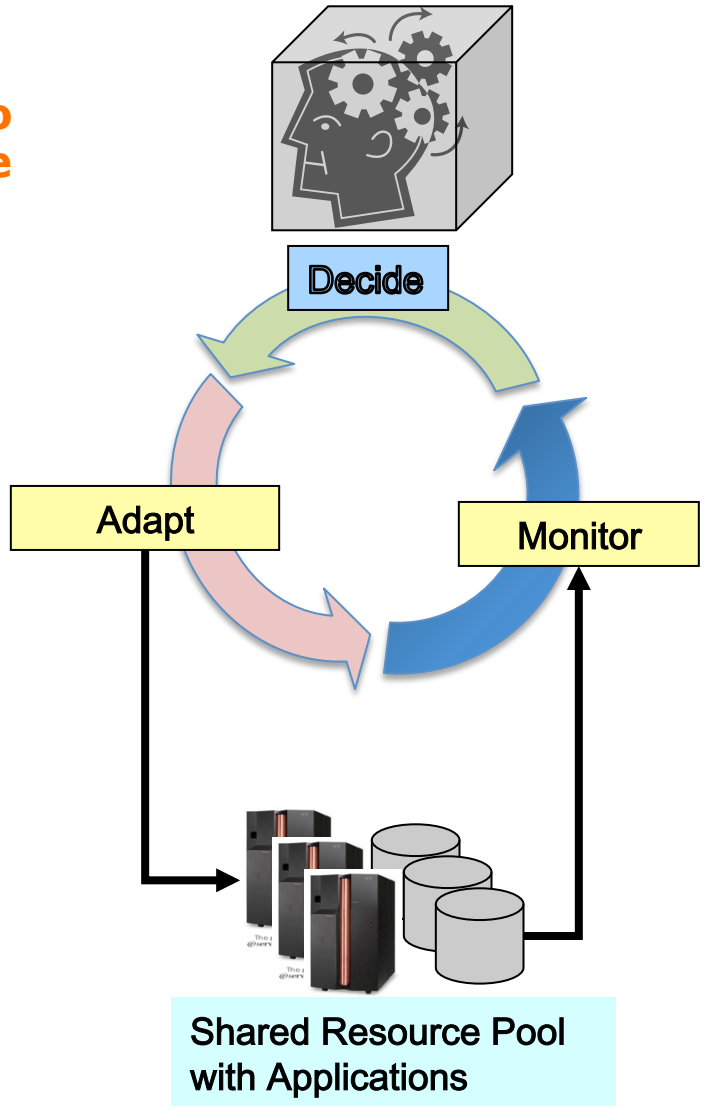

Runtime Resource Management

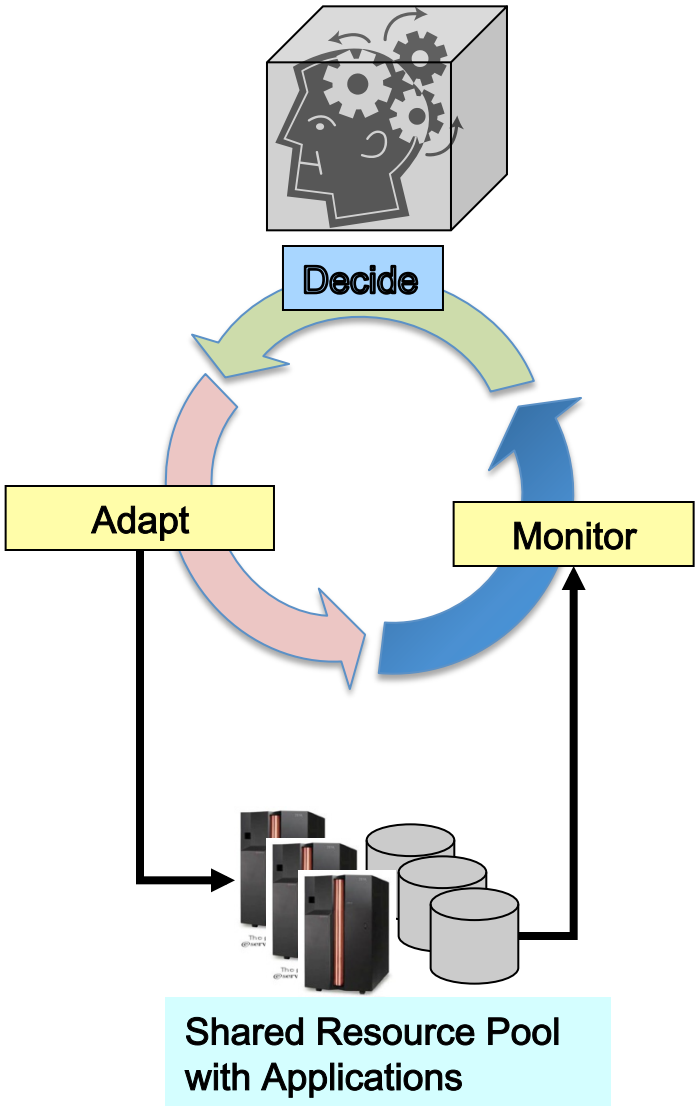

Monitor:

- Request rates of the different applications and their different transaction types (workload), resource utilization, response times

Actions:

- Start/stop Virtual Machines (VMs) (e.g., adjust replication degree of a component).
- Migrate VMs
- Adjust VM CPU fraction

Goal:

- Optimize resource utilization (maximize utility), minimize SLA violations (mean response time)
- Consider impact of adaptation
- Focus on managing multiple, multi-tier enterprise applications

How to develop rules?

- **Key challenge:** Predicting response times and resource utilization
- **Approach:** Models

Architecture

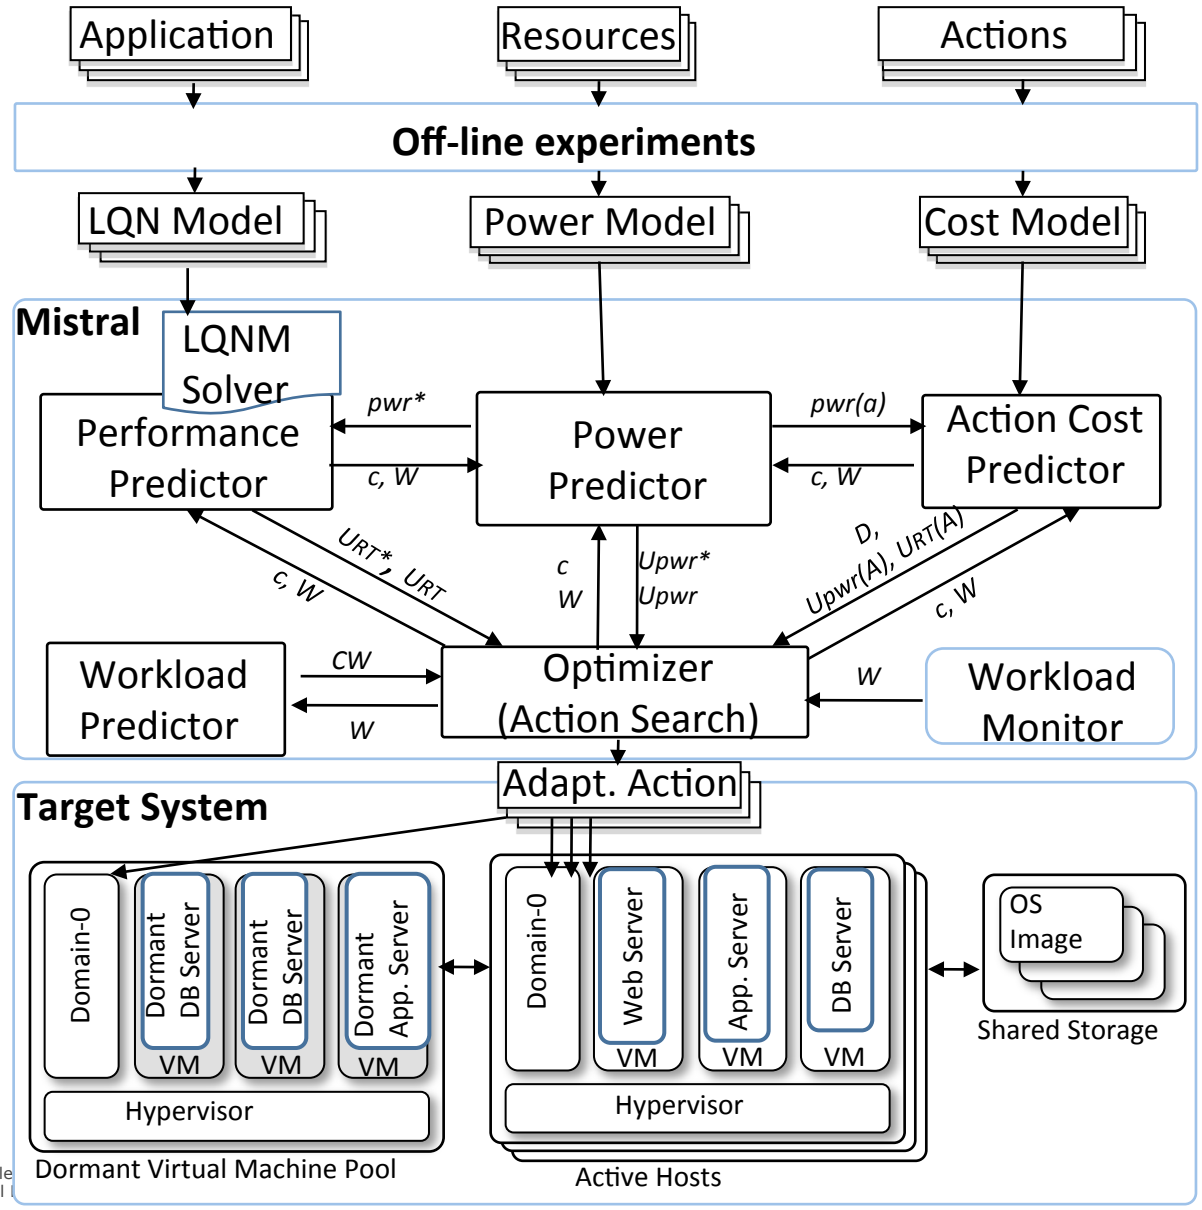

Cumulative Utility ($$$)

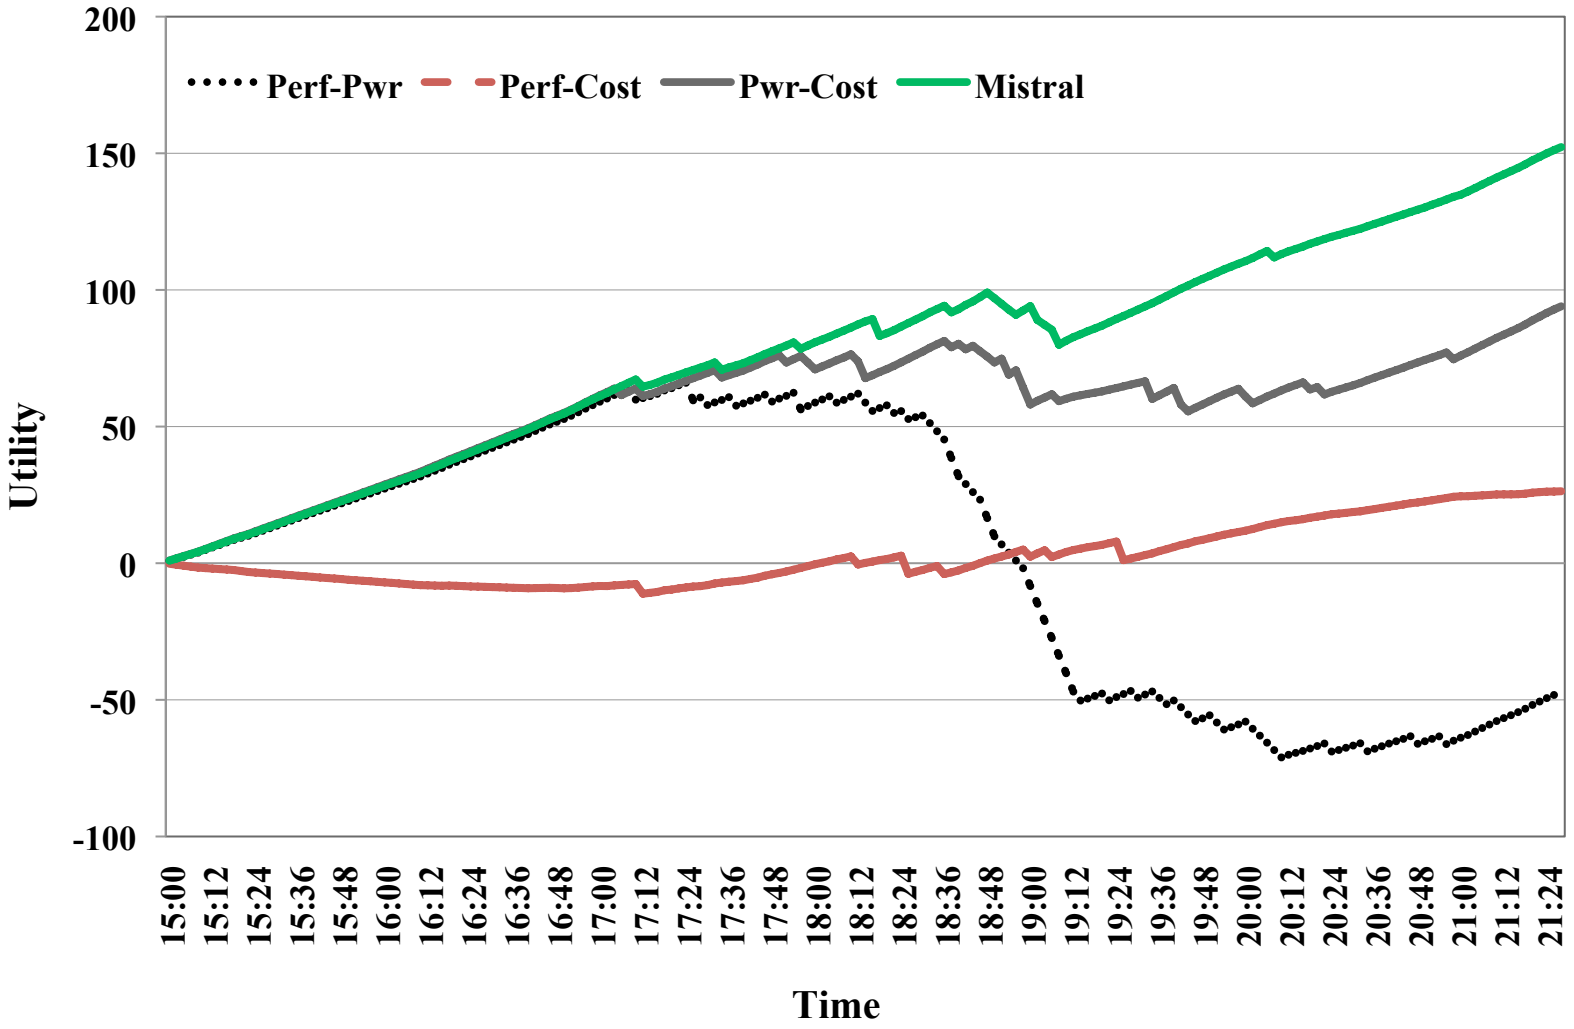

Mistral: Key Points

Addresses the control plane issue of optimizing resource utilization though adaptive actions

- Policy component
- Multi-dimensional optimization problem
- Approach based on offline modeling
- Limitations:
  - Collections of multi-tier enterprise applications
  - Modeling requirement
  - Optimization step

➨ Scalability addressed using multiple cooperating controllers operating at different time scales

PipeCloud: Disaster Recovery as a Cloud Service

(T. Wood, A. Lagar-Cavilla, KK Ramakrishnan, P. Shenoy, J. Van der Merwe, ACM SOCC 2011, with U. Mass)

Key challenge: providing DR services to support Business Continuity (BC), allowing applications to rapidly come back online after a failure occurs

Current DR services - expensive

- Performance: Come either at very high cost (Synch replication) or weak guarantees on amount of data loss (Asynch replication)
- Significant time required to restart operation after a failure (large *Recovery Point Objective* (RPO)).

Cloud computing platforms can be well suited for DR as a service

- pay-as-you-go pricing model can lower costs
- use of automated virtual platforms can minimize recovery time after a failure.

Using the Cloud

Long distance data backups + failover mechanism

- Replicate disk writes
- Switch to backup site after disaster

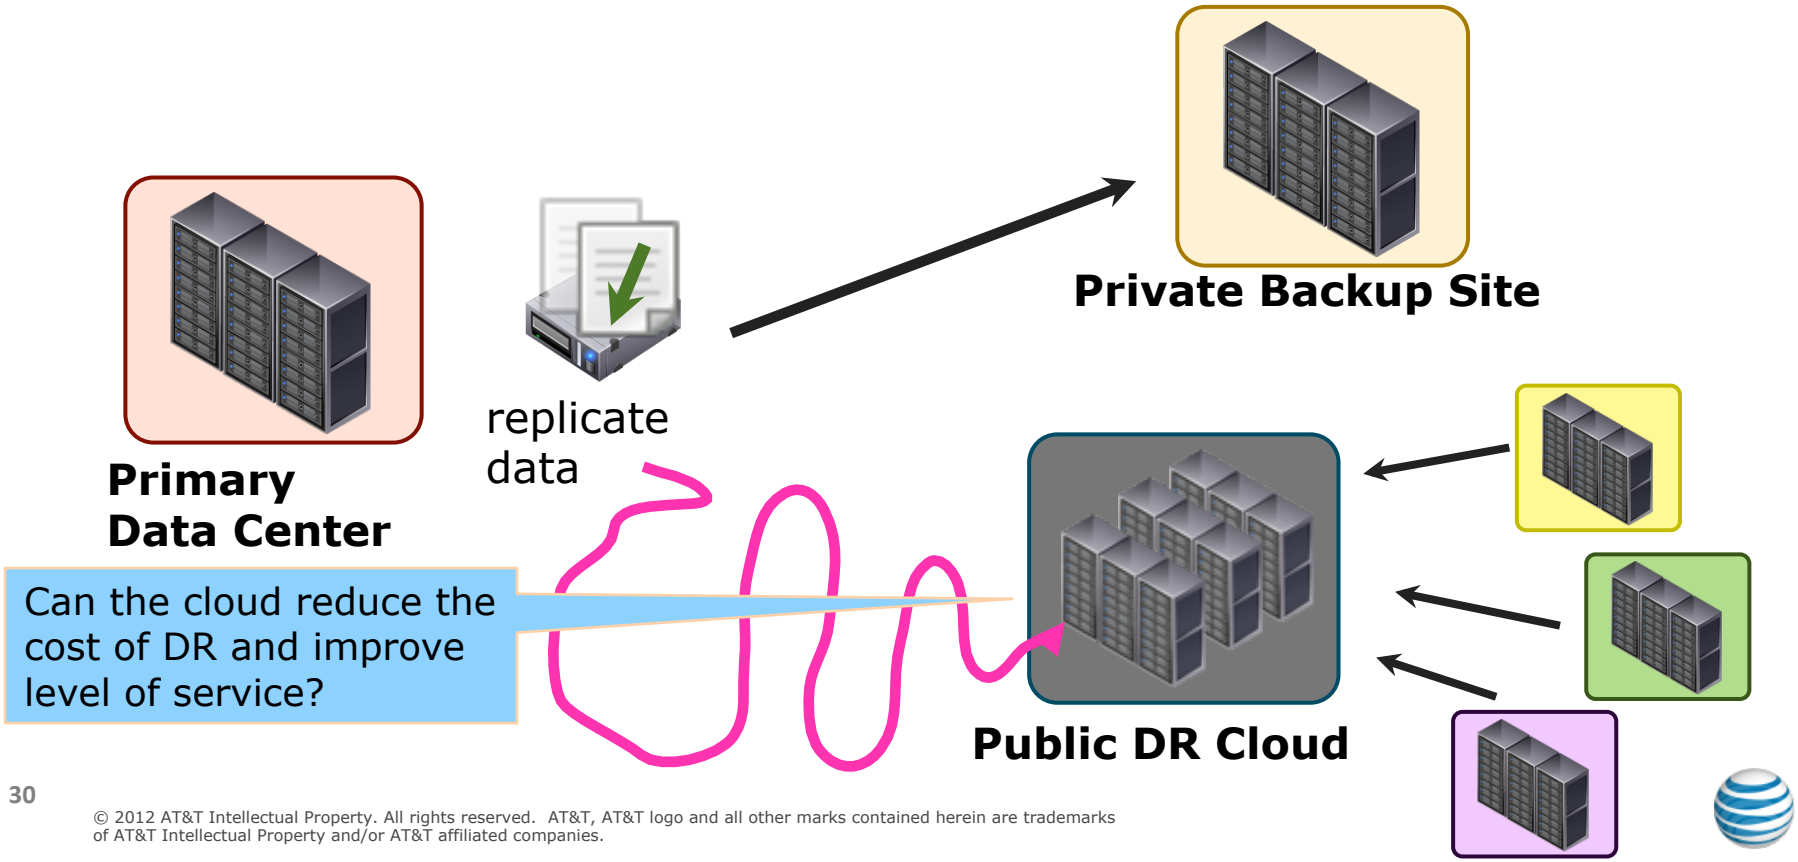

Using the Cloud

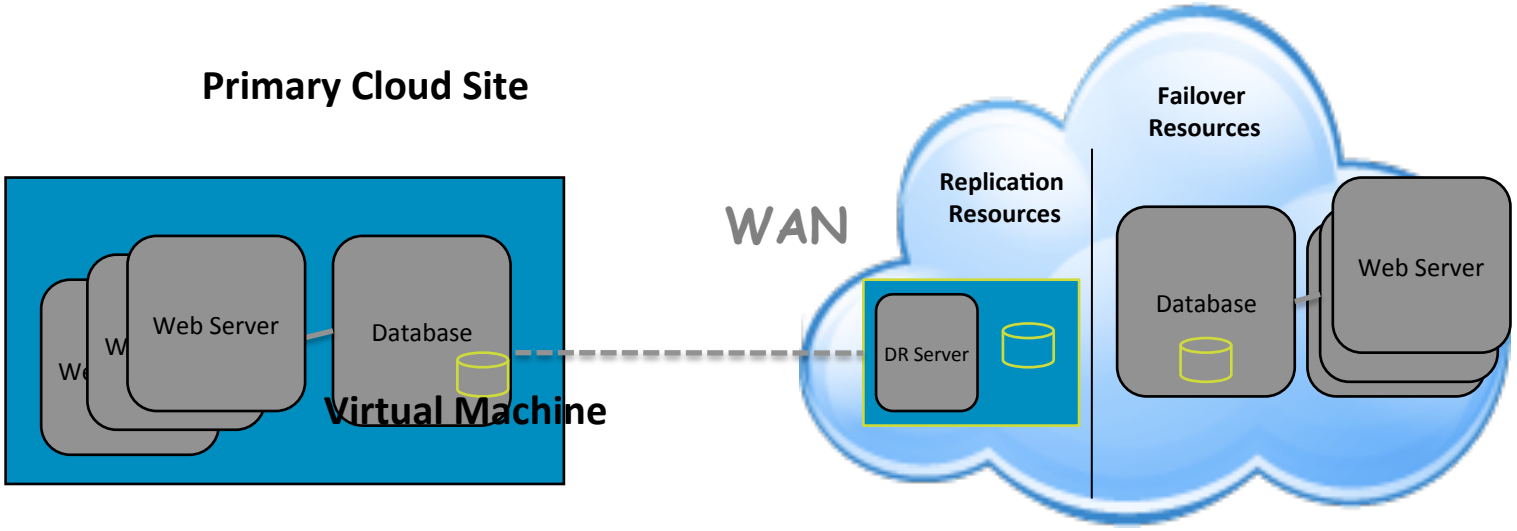

- Resources in the Cloud during non-failure mode to replicate state can be shared and be low-cost (statistical multiplexing)
- Failover Mode resources used only as needed, and available on-demand from Cloud Provider
- **Automation makes BC== DR** with nearly immediate failover
- Cost Analysis: considerable savings for enterprise

Tracking Writes and Replies

Treat VM as black box

- **No modifications to application**
- Track disk writes and network calls

Hypervisor intercepts all I/O

- Replicate disk writes to backup
- Hold network packets until prior disk writes are committed
- Count pending and committed writes
- **Enforces total ordering of events – even with multi-tier applications**

Achieves considerable performance improvement with near-zero RPO

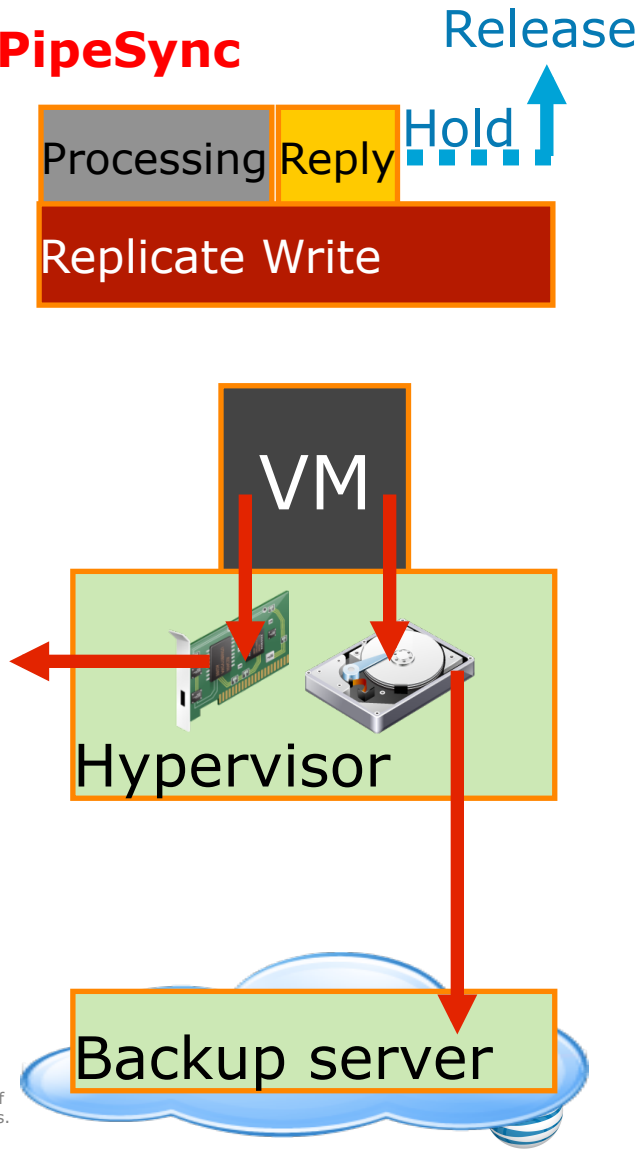

PipeCloud: Key Points

Provides cost-effective cloud-based DR with strong guarantees

- Implements high-level abstraction of continuous service for enterprise applications (high availability)
- Relies on and exploits capabilities of cloud infrastructure
- No modifications to the VM/application

➨ Demonstrates a compelling approach to realizing enhanced properties (ie., QoS attributes) for applications that use the cloud.

Conclusions

An Extreme Cloud:

- Is ubiquitous
- Supports heterogeneous compute, storage and network resources
- Provides enhanced properties and service support
- Realized through a highly decentralized infrastructure

Enables new types of services and applications

➨ **A true "computer utility"?**

Rethink Possible