

# Odin: Microsoft's Scalable Fault-Tolerant CDN Measurement System

*Matt Calder*

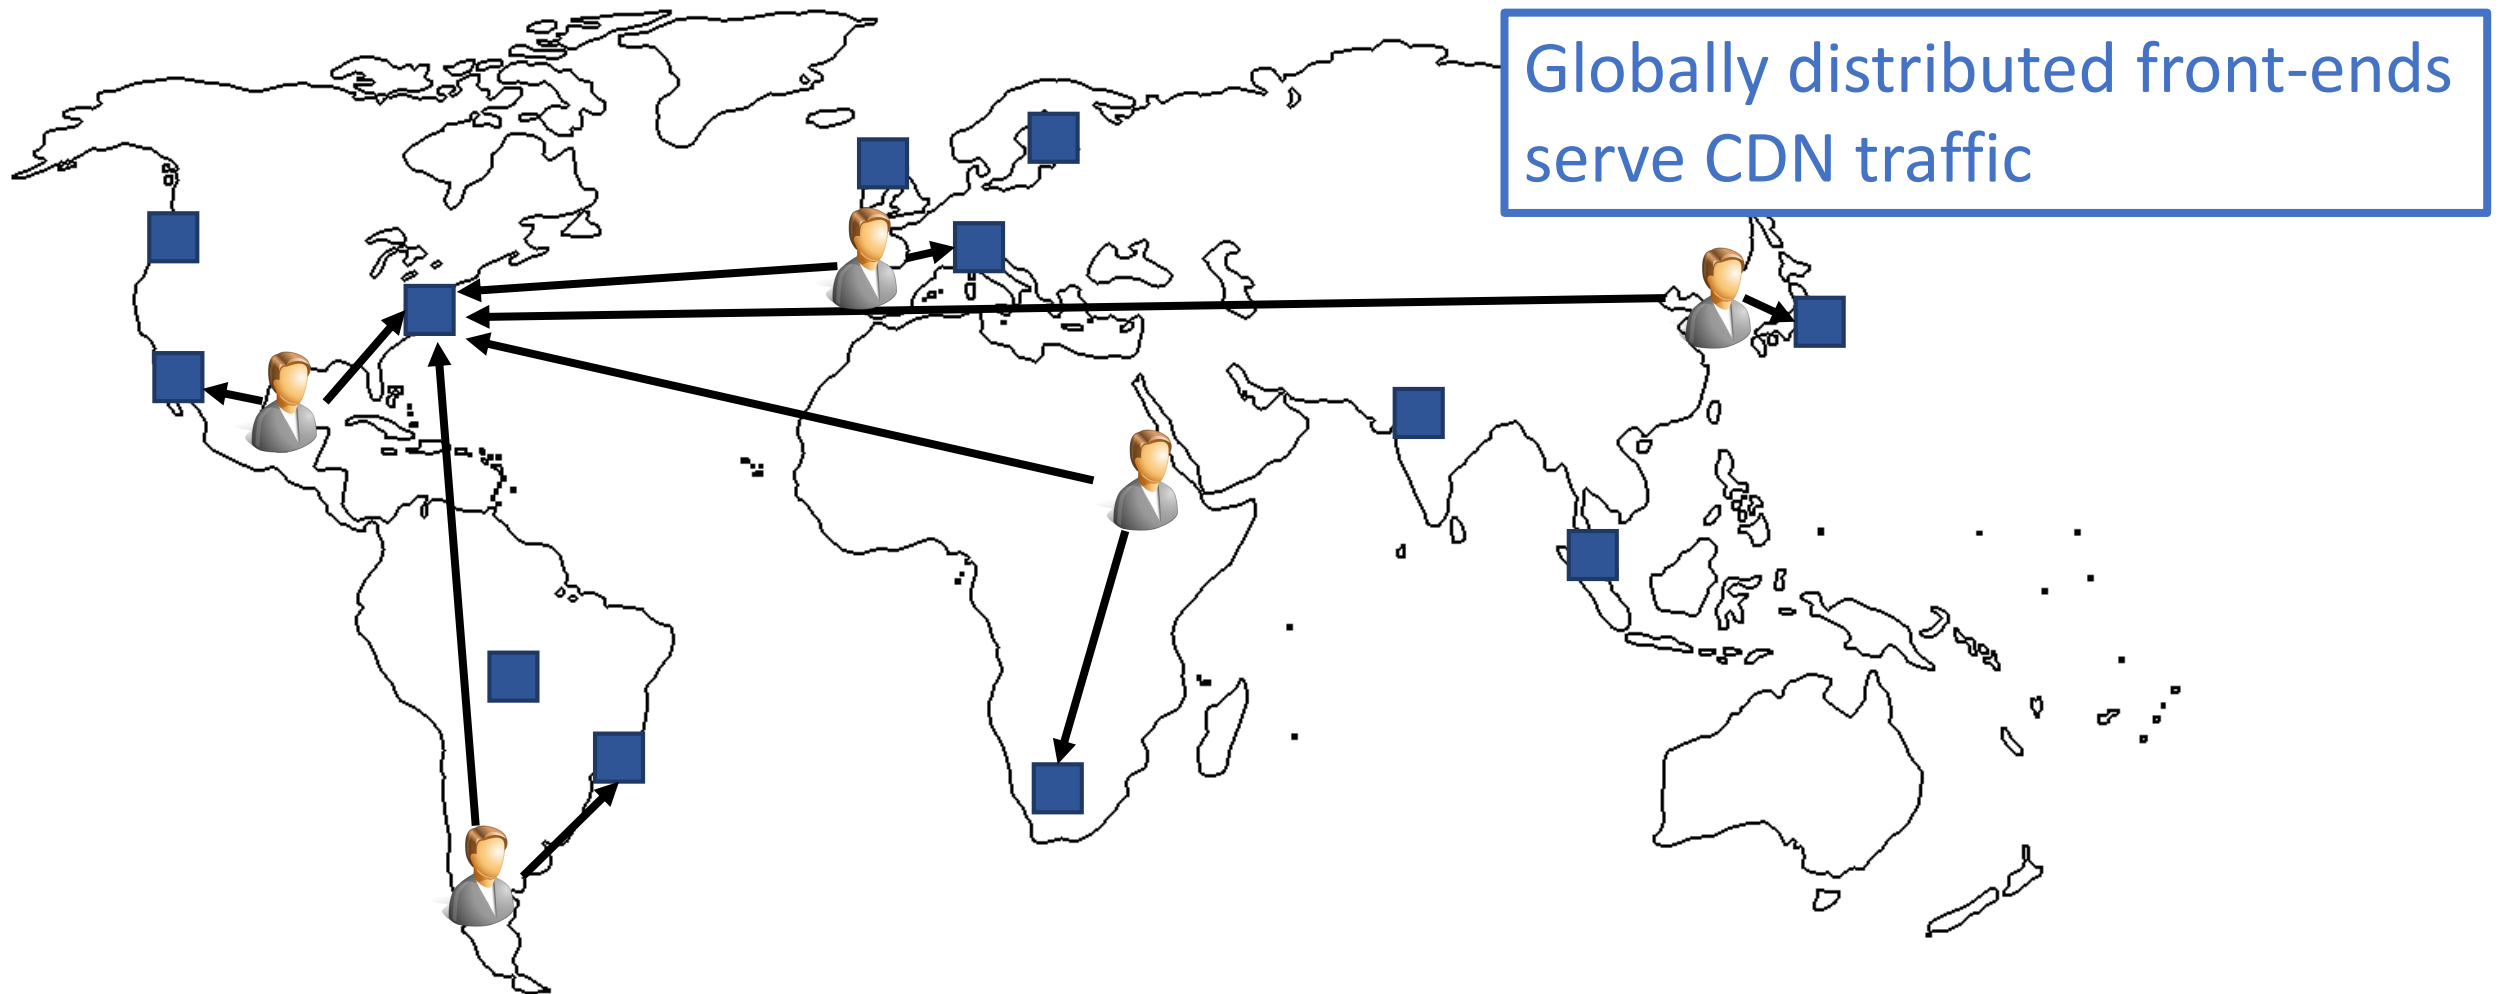
Manuel Schröder, Ryan Gao, Ryan Stewart, Jitendra Padhye, Ratul Mahajan,  
Ganesh Ananthanarayanan, Ethan Katz-Bassett

NSDI, April 2018

# CDN Overview

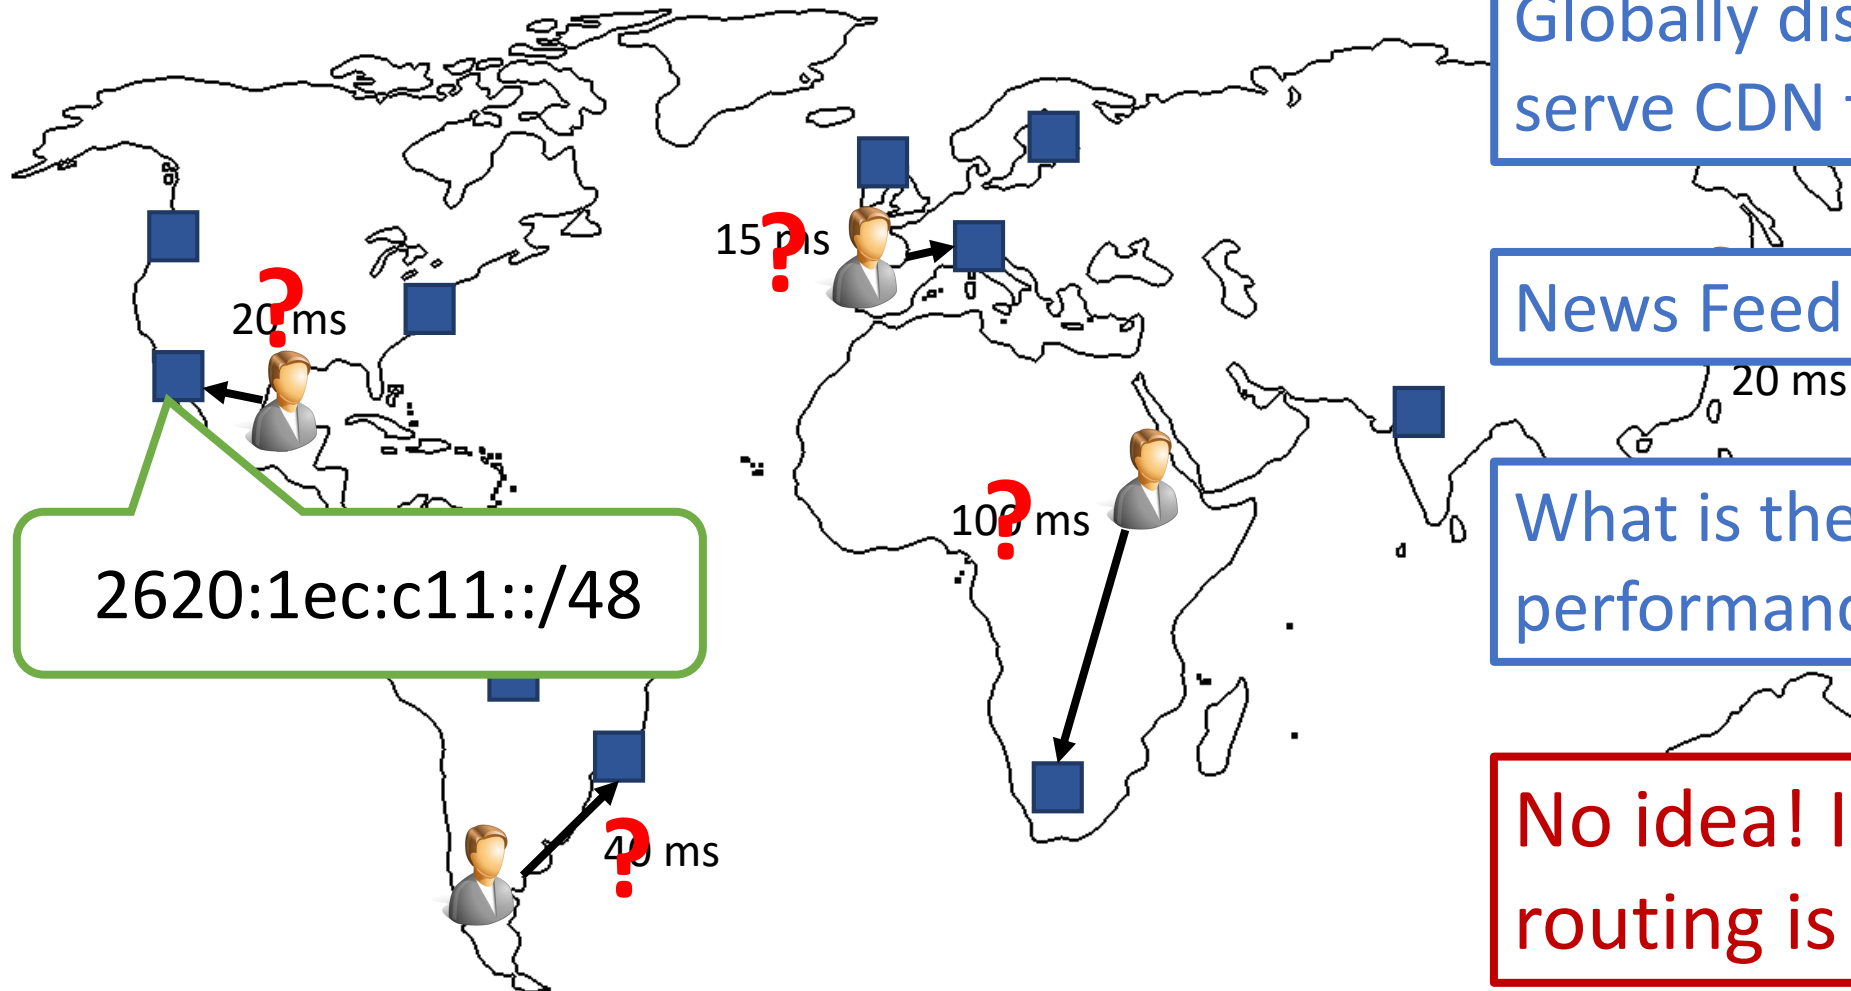
Want low latency service for all clients

Globally distributed front-ends  
serve CDN traffic



# Problem: What is the impact of change in a CDN?

## Microsoft News Feed Application



Globally distributed front-ends serve CDN traffic

News Feed app needs IPv6

What is the global performance impact?

No idea! IPv4 and IPv6 routing is often different

# Problem: What is the impact of change in a CDN?

Microsoft News Feed Application

Need to understand performance  
impact of CDN change

Difficult to evaluate on a global scale

Globally distributed front-ends

ts IPv6

t?

IPv6

routing is often different

2620:1ec:c

20 ms

# Large-scale what-if scenarios

Odin can evaluate these types of changes easily

1. Safely – don't touch production traffic
2. Coverage of Microsoft users
3. Understand impact before deployment

# Outline

1. Introduction
- 2. CDN Operations**
3. Existing Solutions
4. Odin Design
5. CDN Operations with Odin
6. Conclusion

# Internet is diverse and dynamic

- Tens of thousands of ISPs
- Factors outside of direct CDN control
  - Congestion
  - Routing changes
  - Outages

# Internet is diverse and dynamic

- Tens of thousands of ISPs
- Factors affecting performance
  - Congestion
  - Routing
  - Outage

CDN network operations require constant insight into user performance

Internet measurement can provide this insight

# Operational need for measurement

- When services become unavailable or slow, want to alert and reroute quickly
- Need diagnostic capabilities to find the root cause of issues
  - Comcast in Seattle is having trouble reaching my CDN. Are they able to reach other networks?
- Want to measure impact of changes on end-users
  - Want to take a front-end offline for maintenance. What is the performance impact on that front-end's users?

# Outline

1. Introduction
2. CDN Operations
- 3. Existing Solutions**
4. Odin Design
5. CDN Operations with Odin
6. Conclusion

# Layer 3 measurements from CDN infrastructure

- Traceroute, Ping with ICMP
- Launch from Front-ends to Internet destinations
- Issue 1: Poor coverage
  - 74% of end-users unresponsive
- Issue 2: Missing layer 7 behaviors
  - HTTP redirection
  - SSL/TLS

# Layer 3 measurements from CDN infrastructure

Requirements	Layer 3 Measurement from CDN		
Coverage of Microsoft Users	✗		
Application Layer Measurements	✗		

# Server-side Instrumentation

- Client connections are instrumented at servers
  - Collect TCP and application layer metrics
  - Very useful but alone insufficient
- Issue 1: No explicit outage signal
- Issue 2: Alternate path exploration adds risk

# Server-side Instrumentation

Requirements	Layer 3 Measurement from CDN	Server-side Instrumentation	
Coverage of Microsoft Users	✗	✓	
Application Layer Measurements	✗	✓	
Explicit Outage Signal	✓	✗	

# Third-party measurement platforms

- Operate set of vantage points on the Internet
- Run measurements on customer's behalf
- Examples: ThousandEyes, Catchpoint, Cedexis, Dynatrace-Keynote, ...

All have inadequate coverage:  
Fewer than 10 measurements per day from 88% of  
Microsoft customer networks

# Third-party measurement platforms

Requirements	Layer 3 Measurement from CDN	Server-side Instrumentation	Third-party Measurement Platforms
Coverage of Microsoft Users	✗	✓	✗
Application Layer Measurements	✗	✓	✓
Explicit Outage Signal	✓	✗	✓

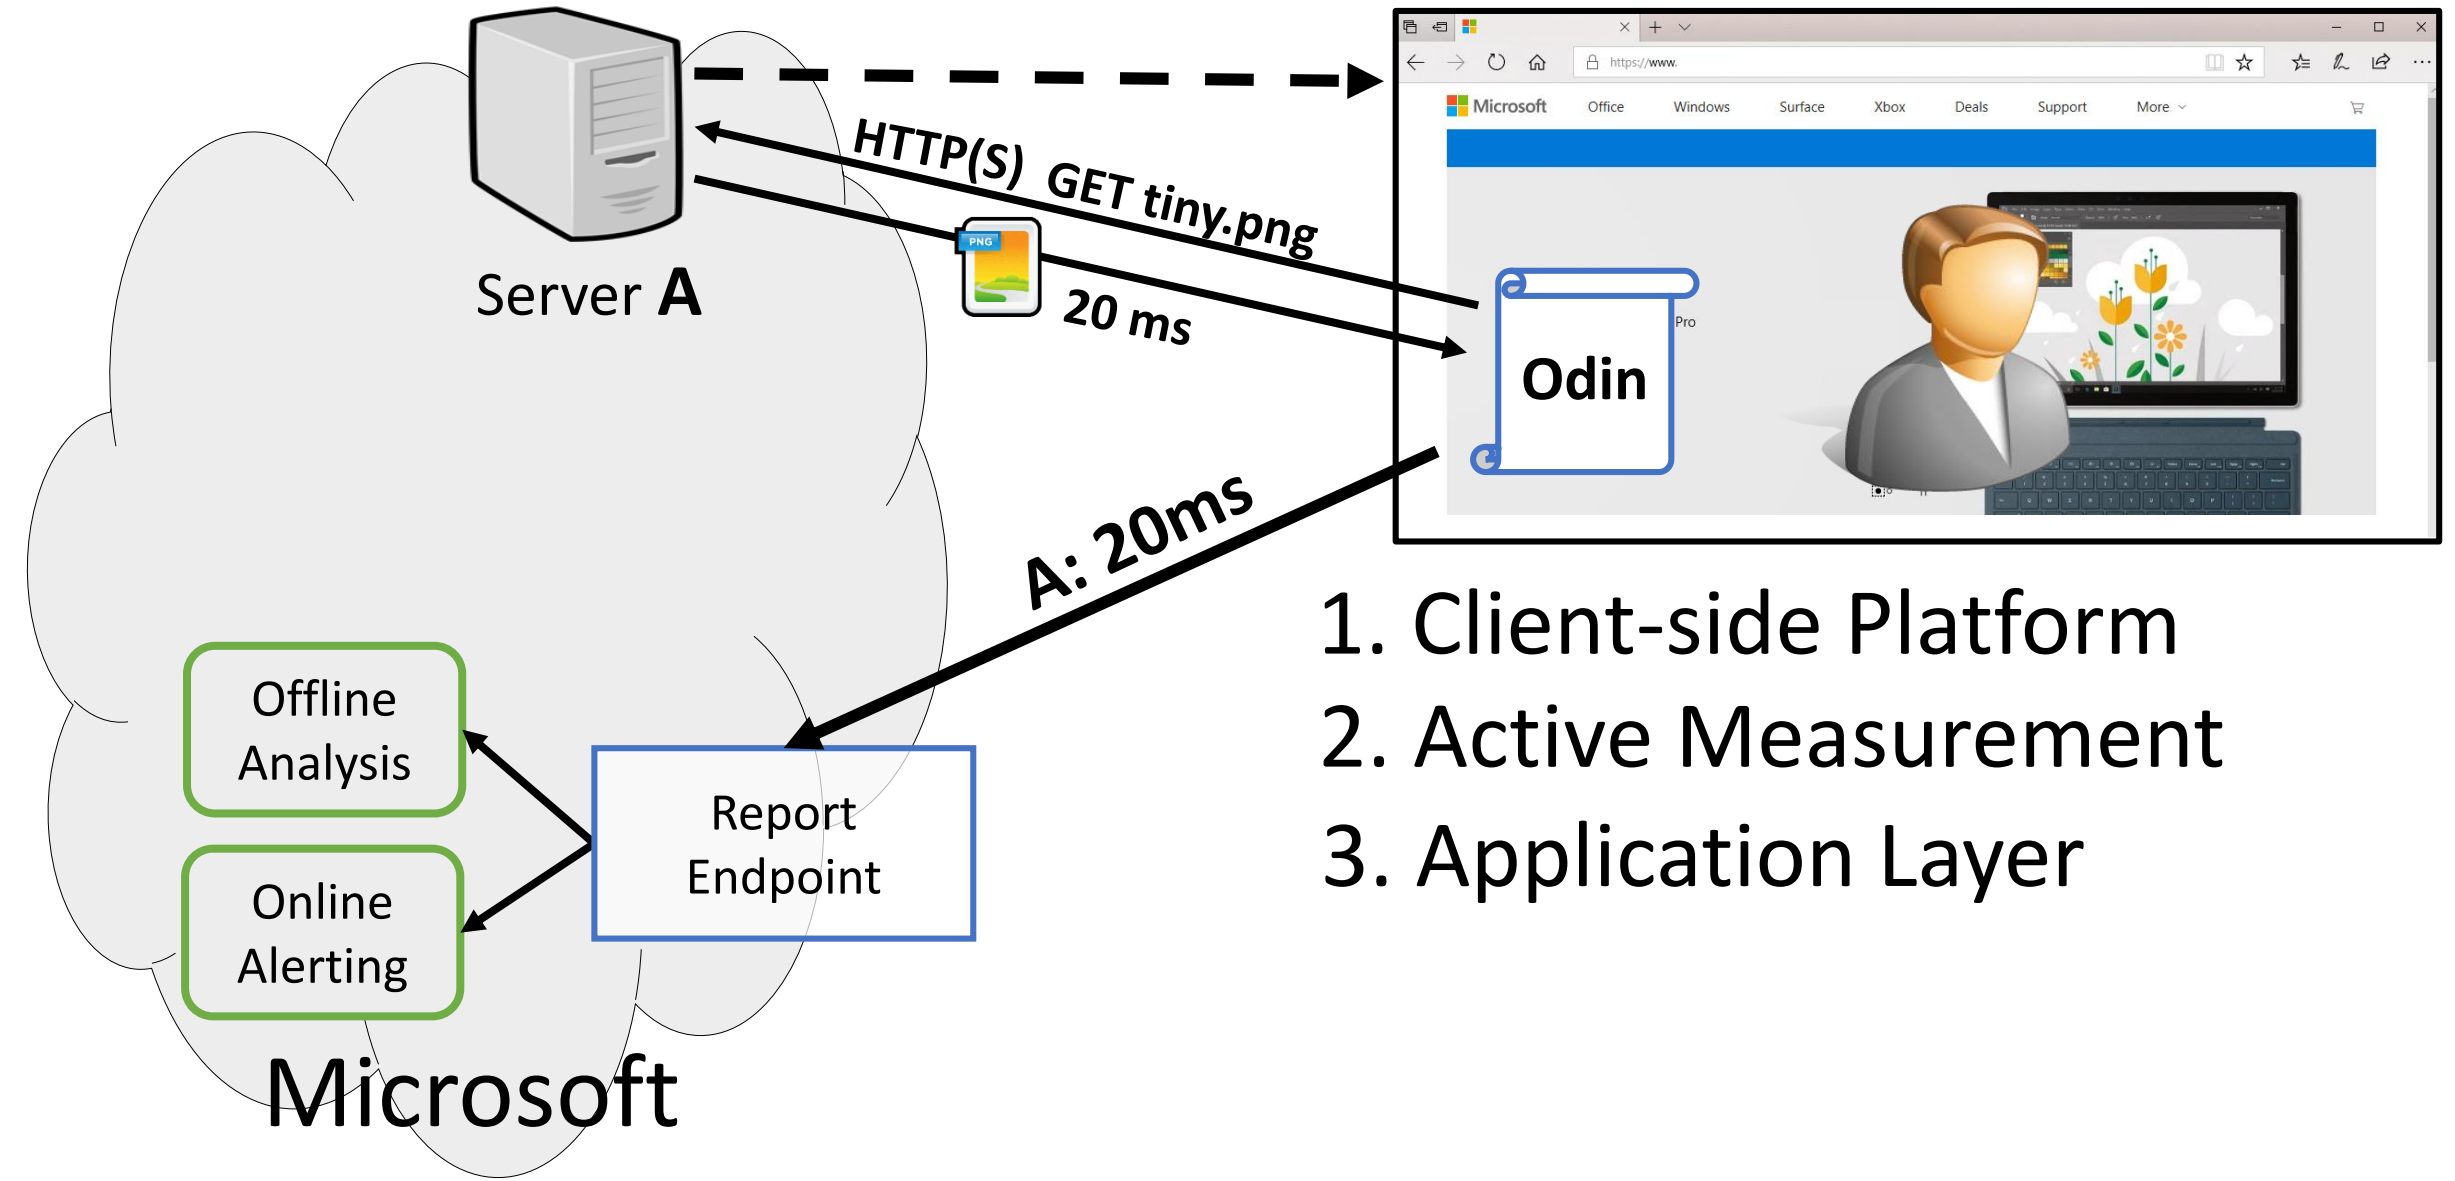
# Existing solutions don't meet our needs

Requirements	Layer 3 Measurement from CDN	Server-side Instrumentation	Third-Party Measurement Platforms	Odin
Coverage of Microsoft Users	✗	✓	✗	✓
Application Layer Measurements	✗	✓	✓	✓
Explicit Outage Signal	✓	✗	✓	✓
... (see paper for others)				✓

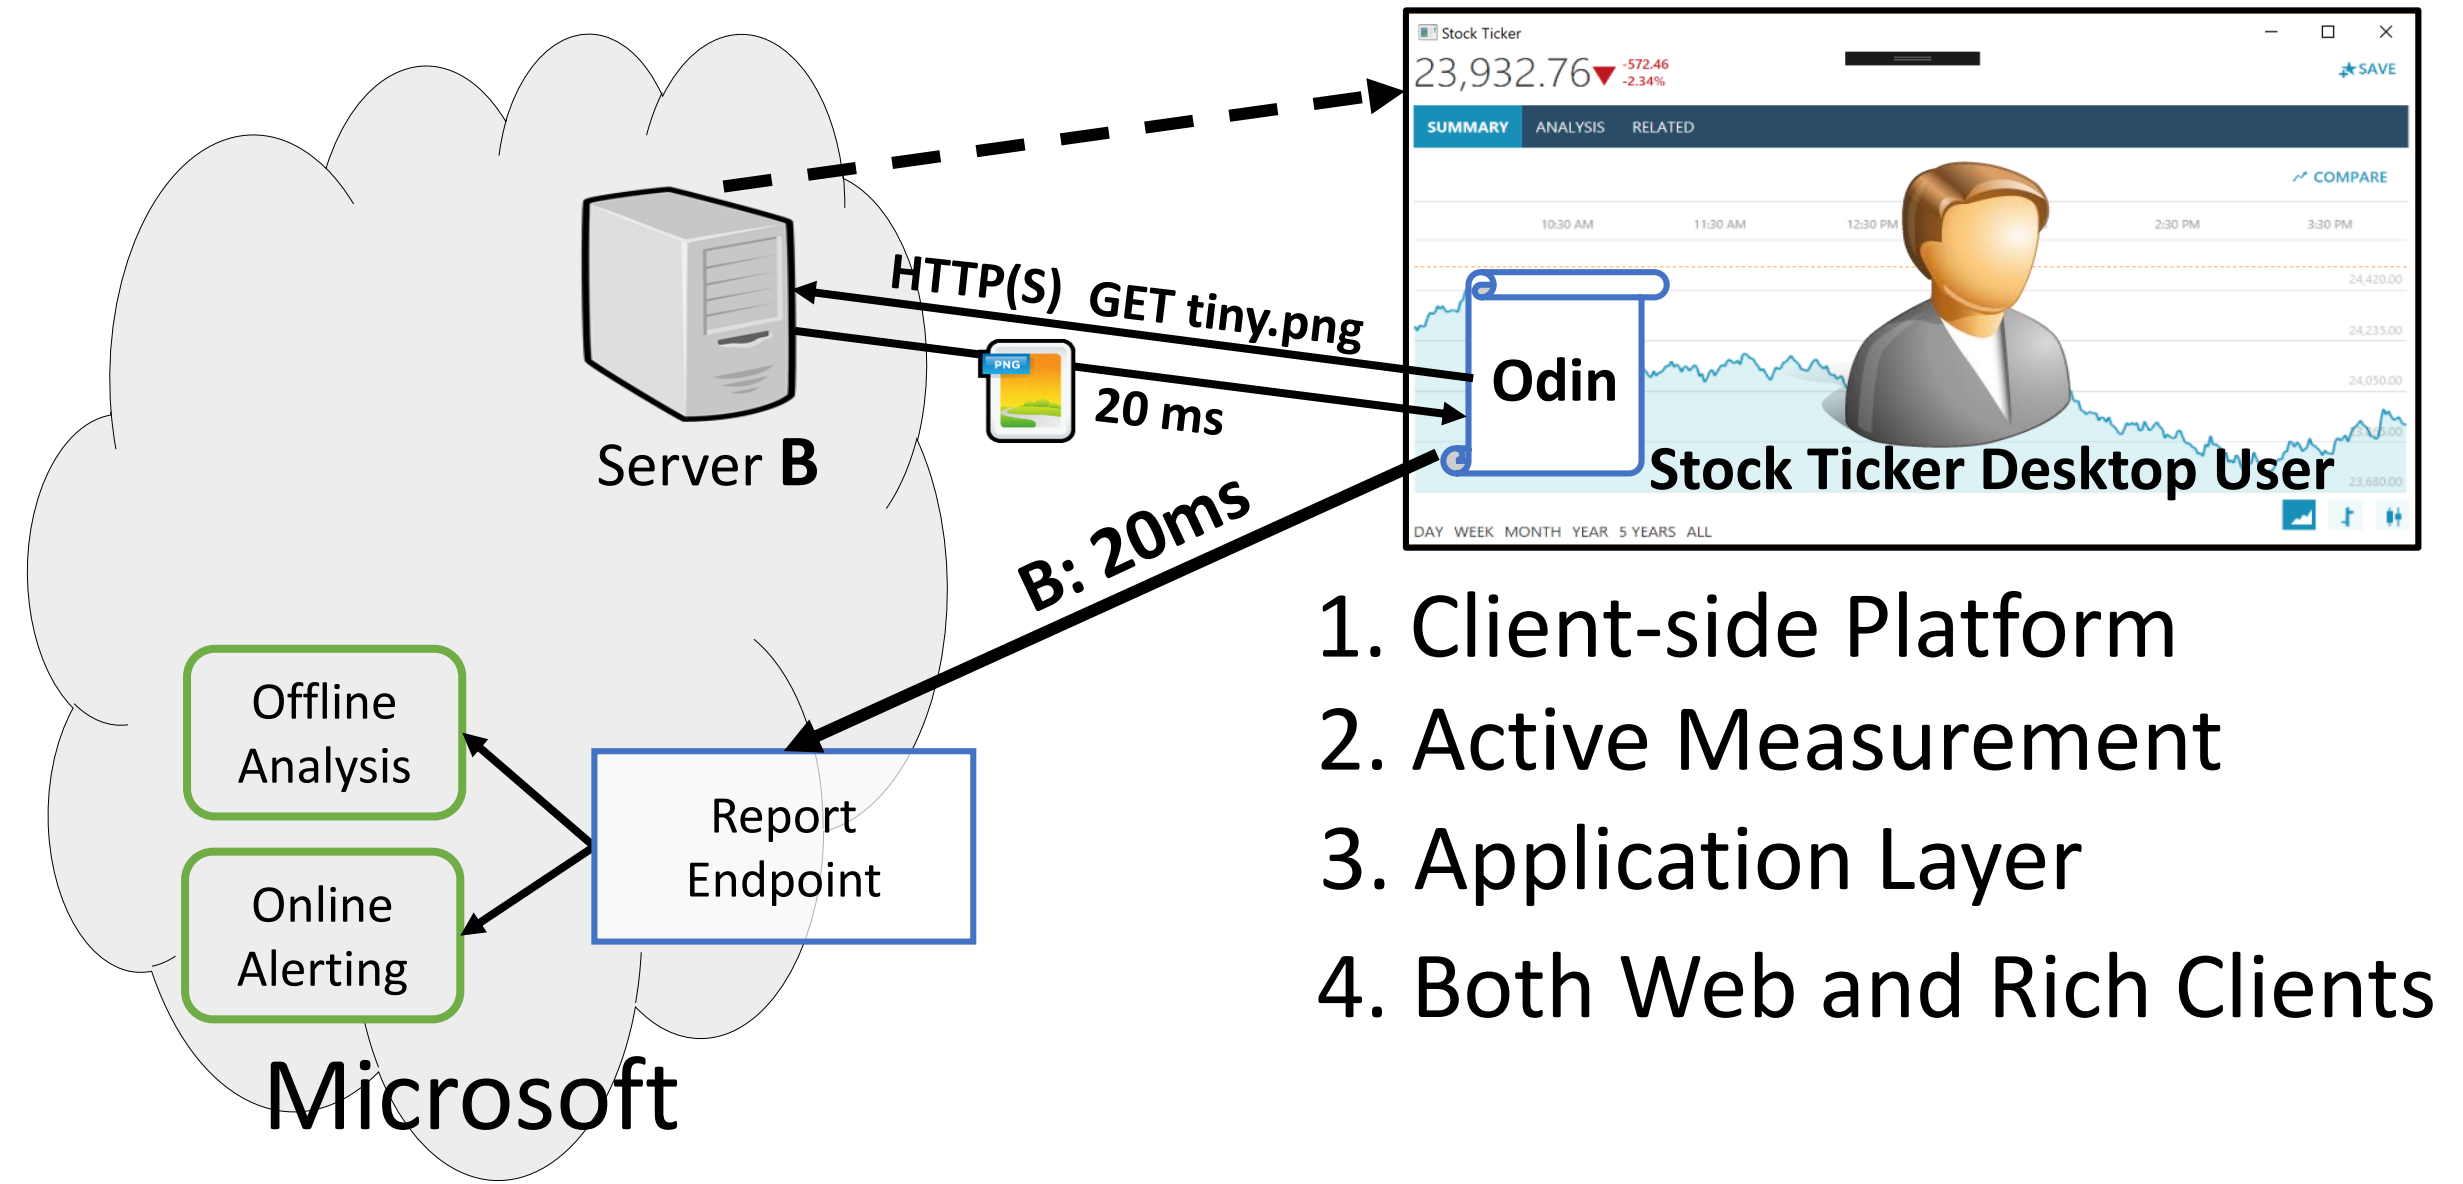
# Outline

1. Introduction
2. CDN Operations
3. Existing Solutions
- 4. Odin Design**
5. CDN Operations with Odin
6. Conclusion

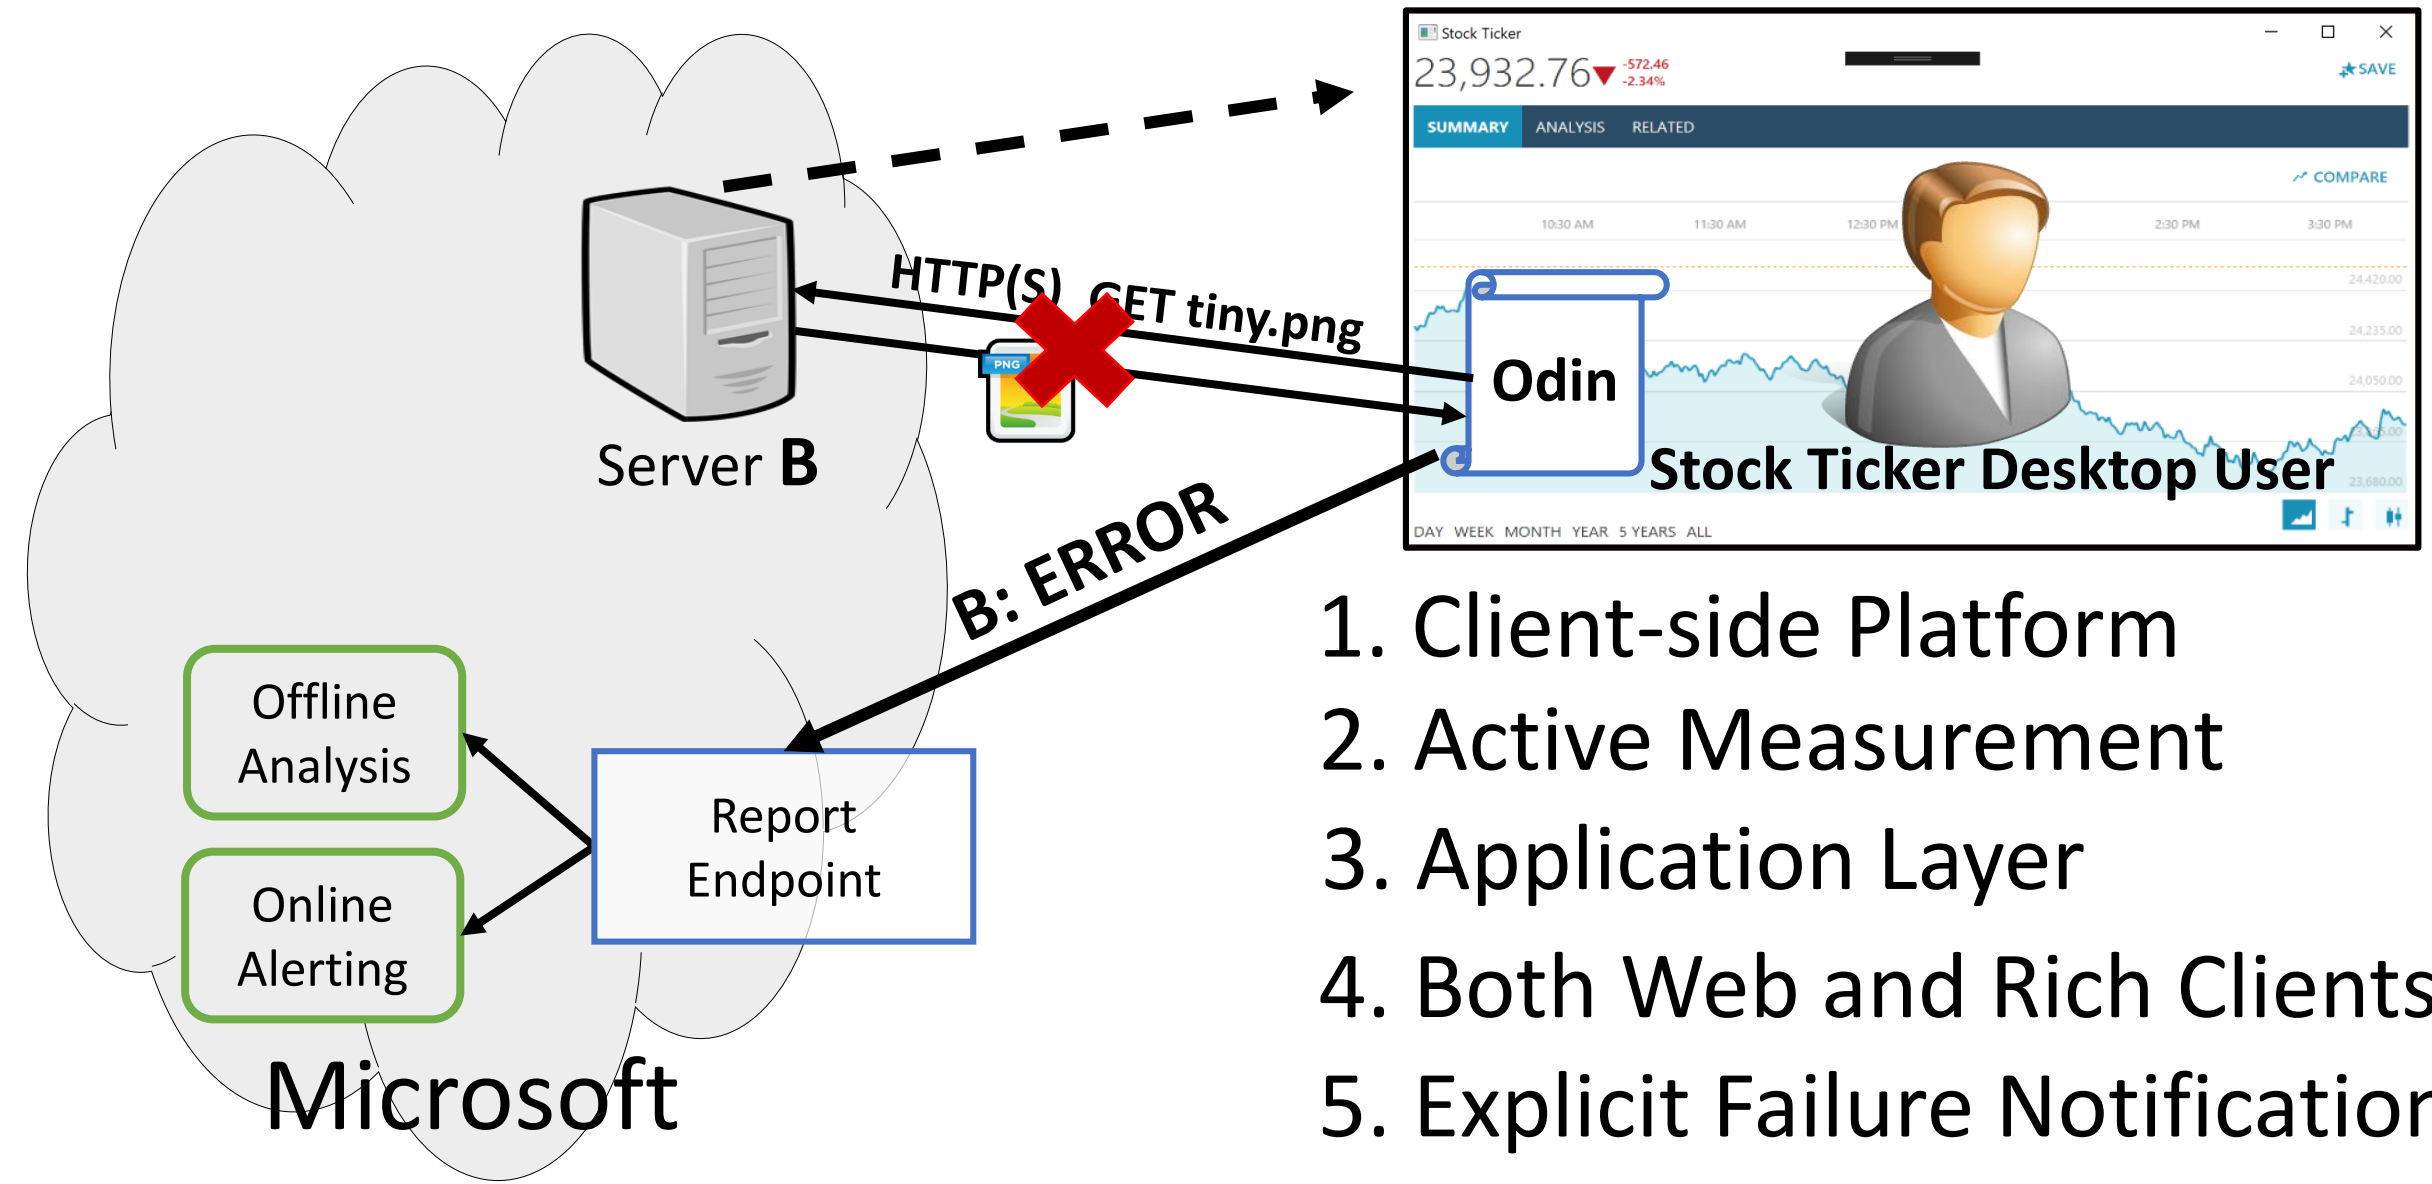
# Odin Design



# Odin Design



# Odin Design

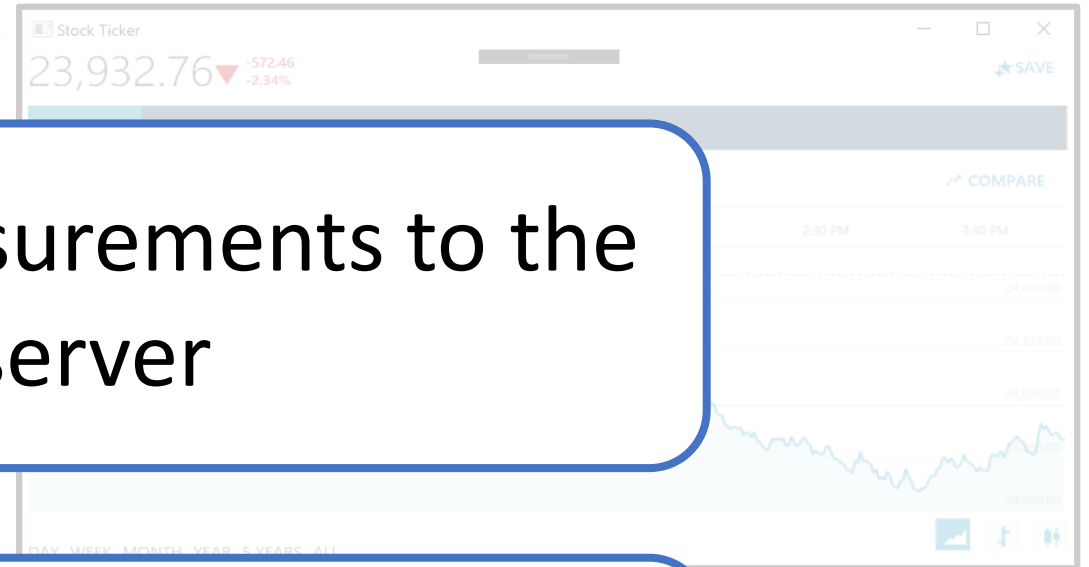


1. Client-side Platform
2. Active Measurement
3. Application Layer
4. Both Web and Rich Clients
5. Explicit Failure Notification

# Odin Design

Examples showed measurements to the application server

Want richer measurements

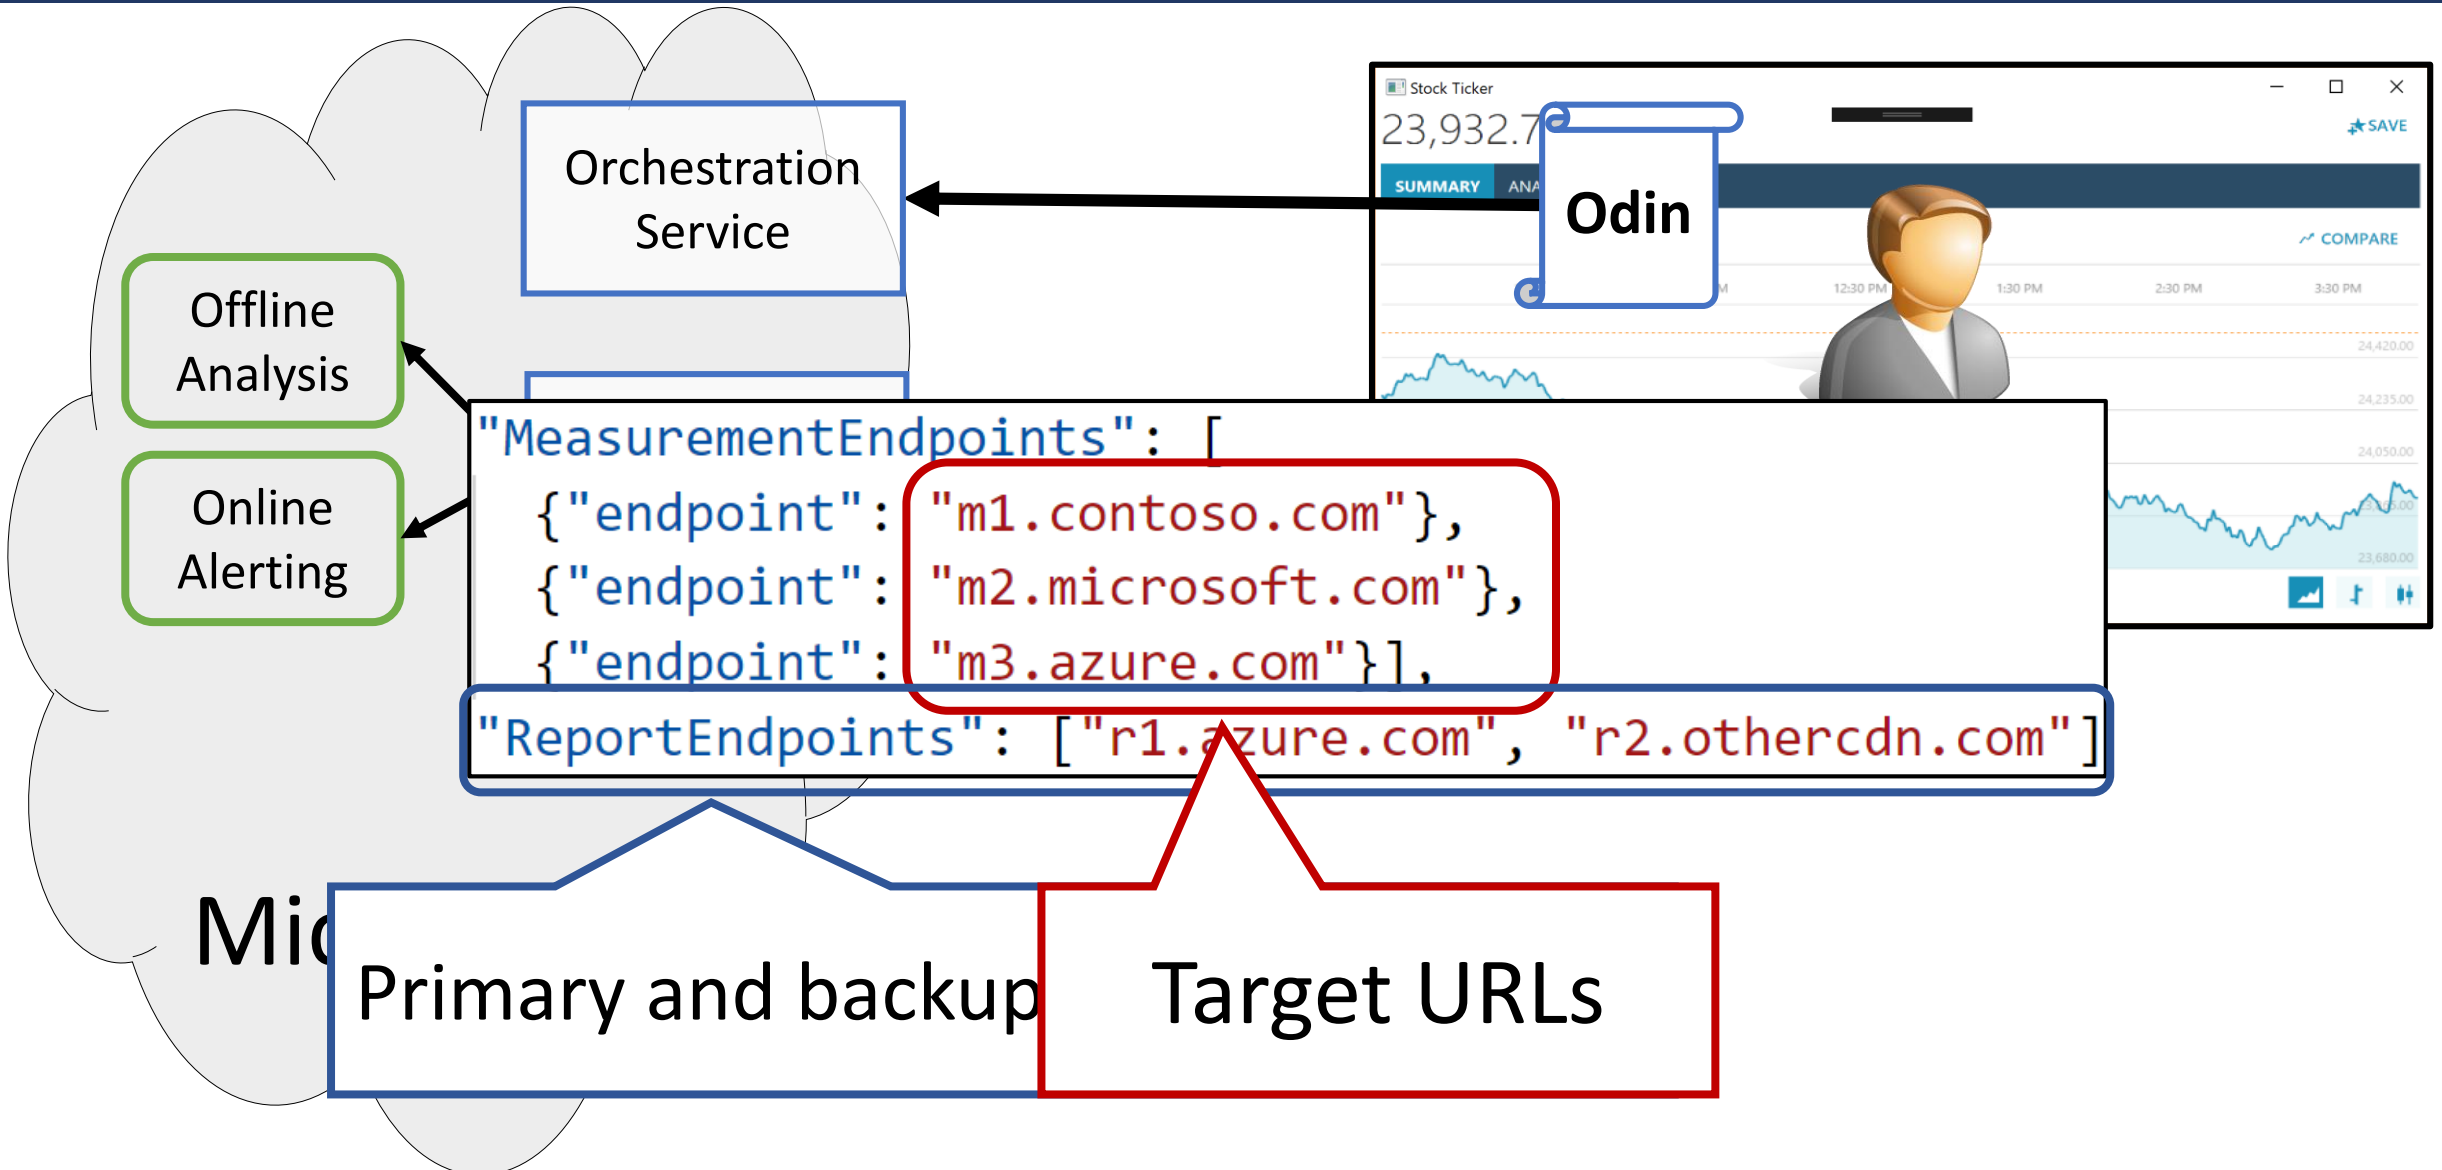


Offline  
Analysis

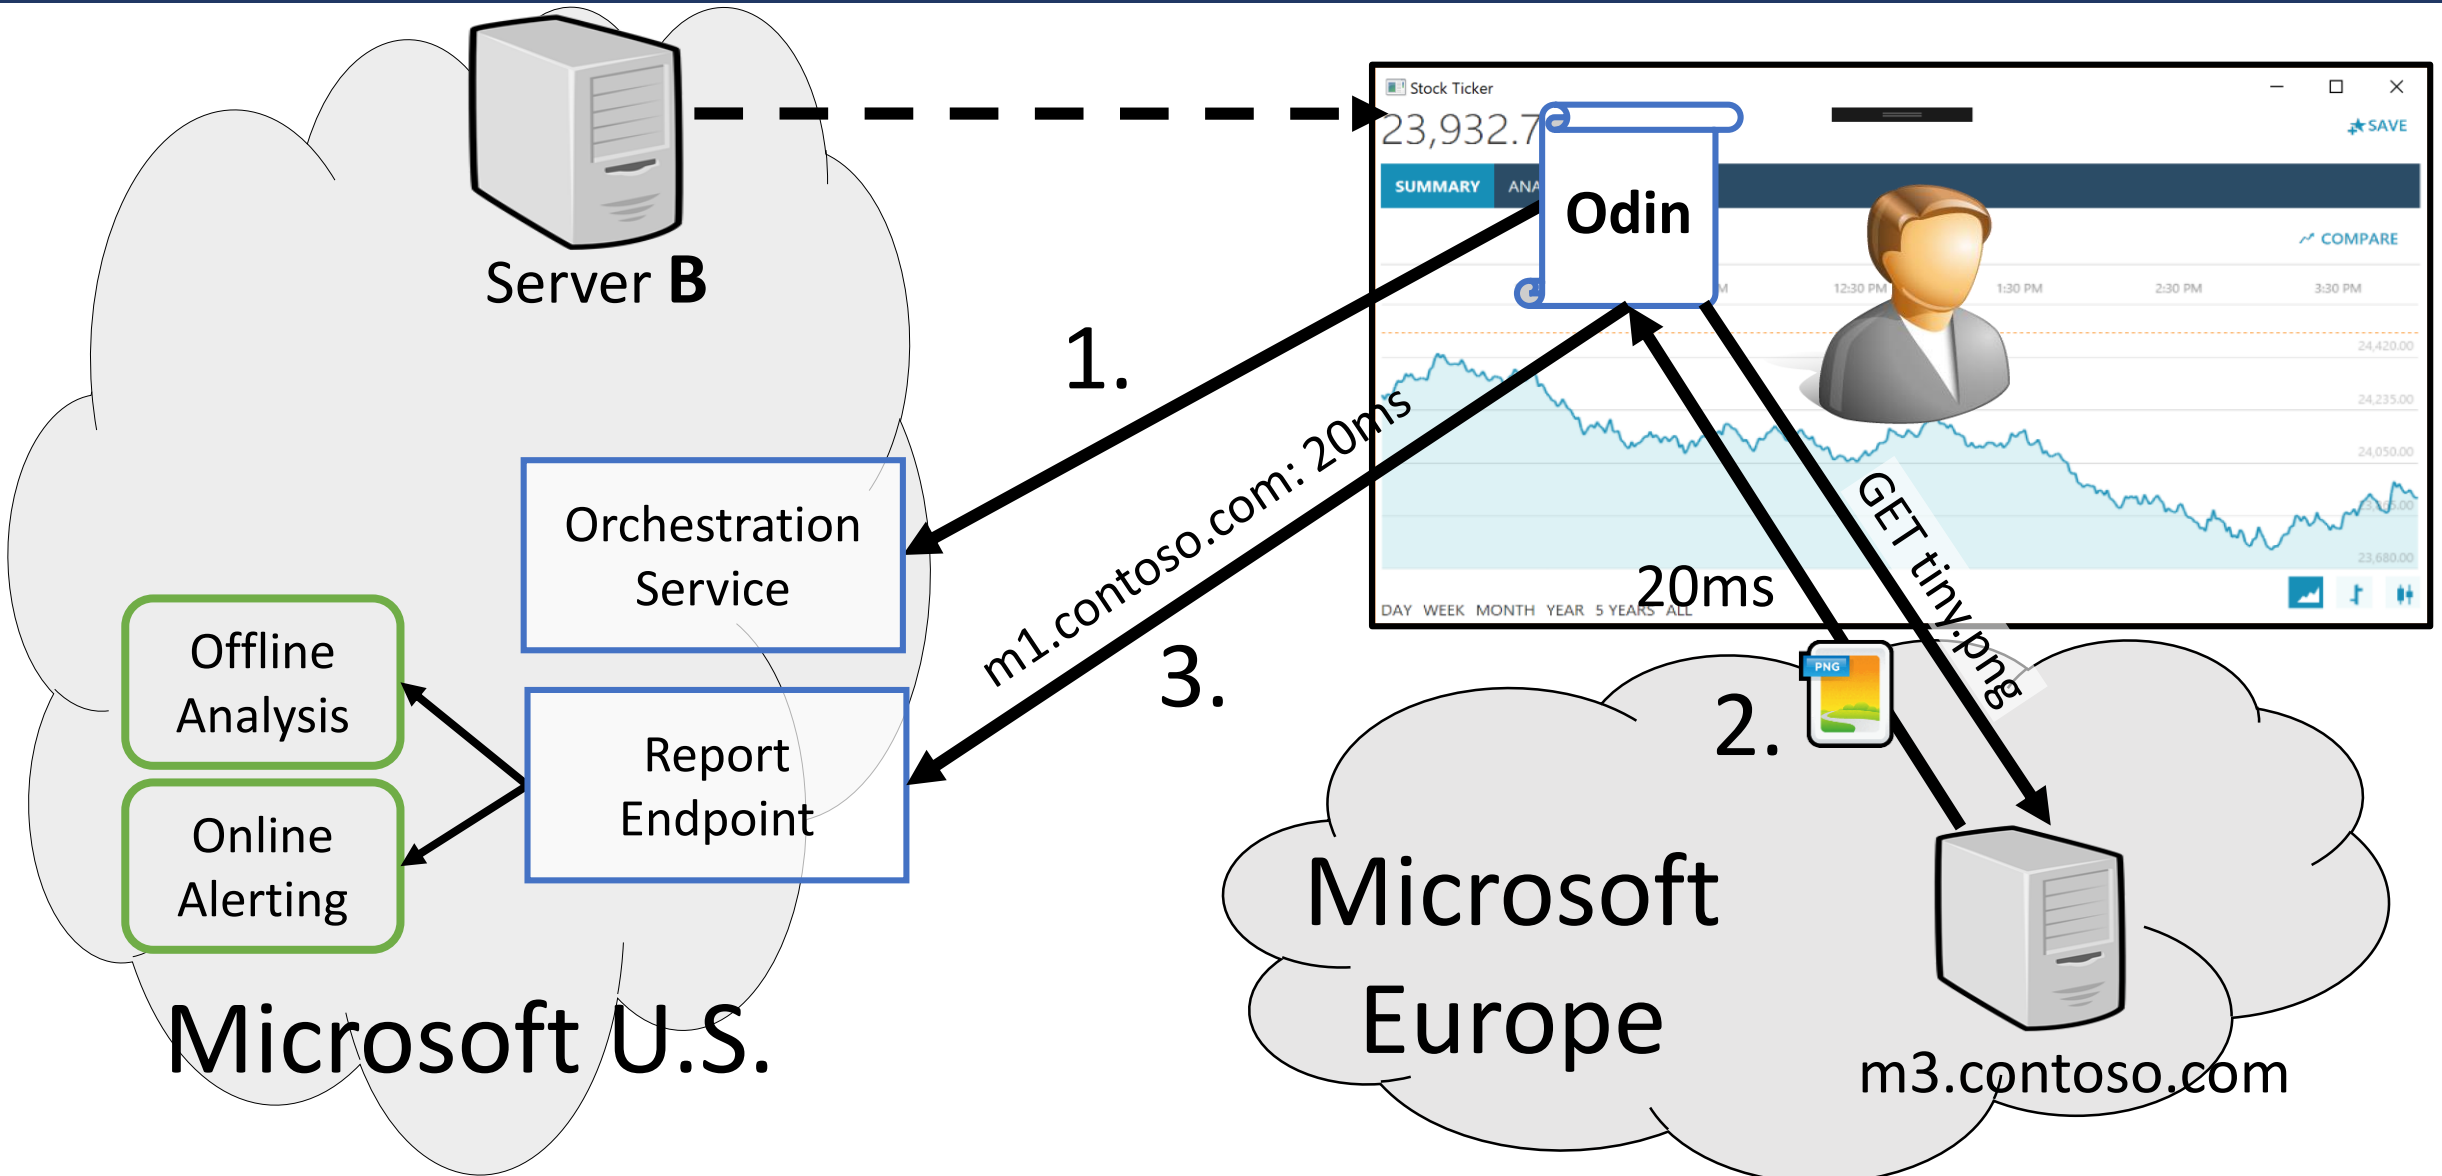
Online  
Alerting

Microsoft

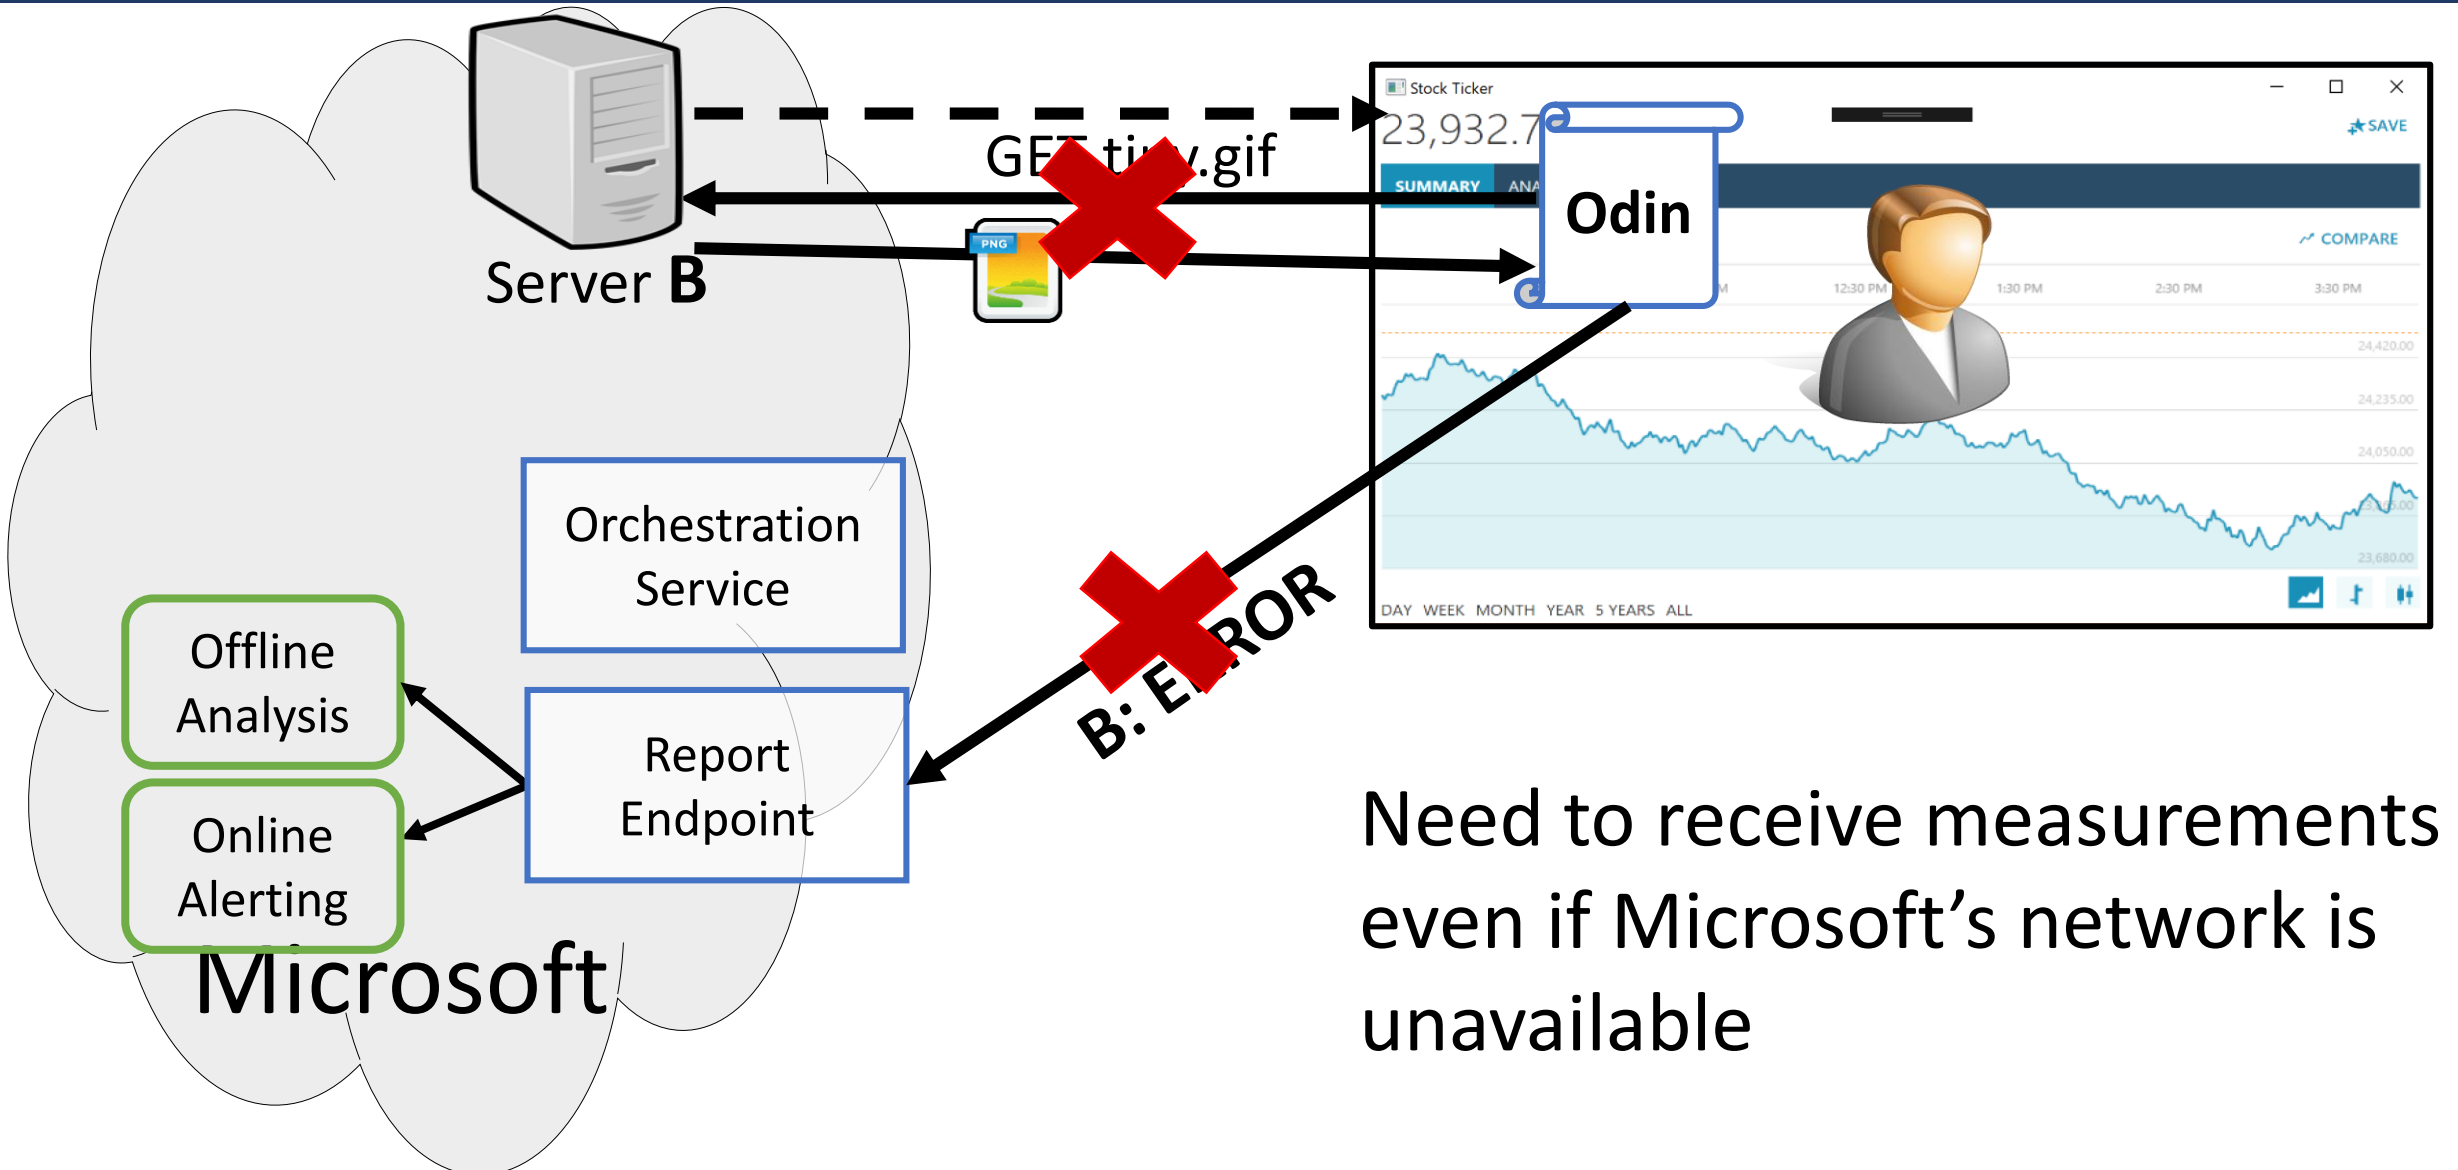
# Odin Design



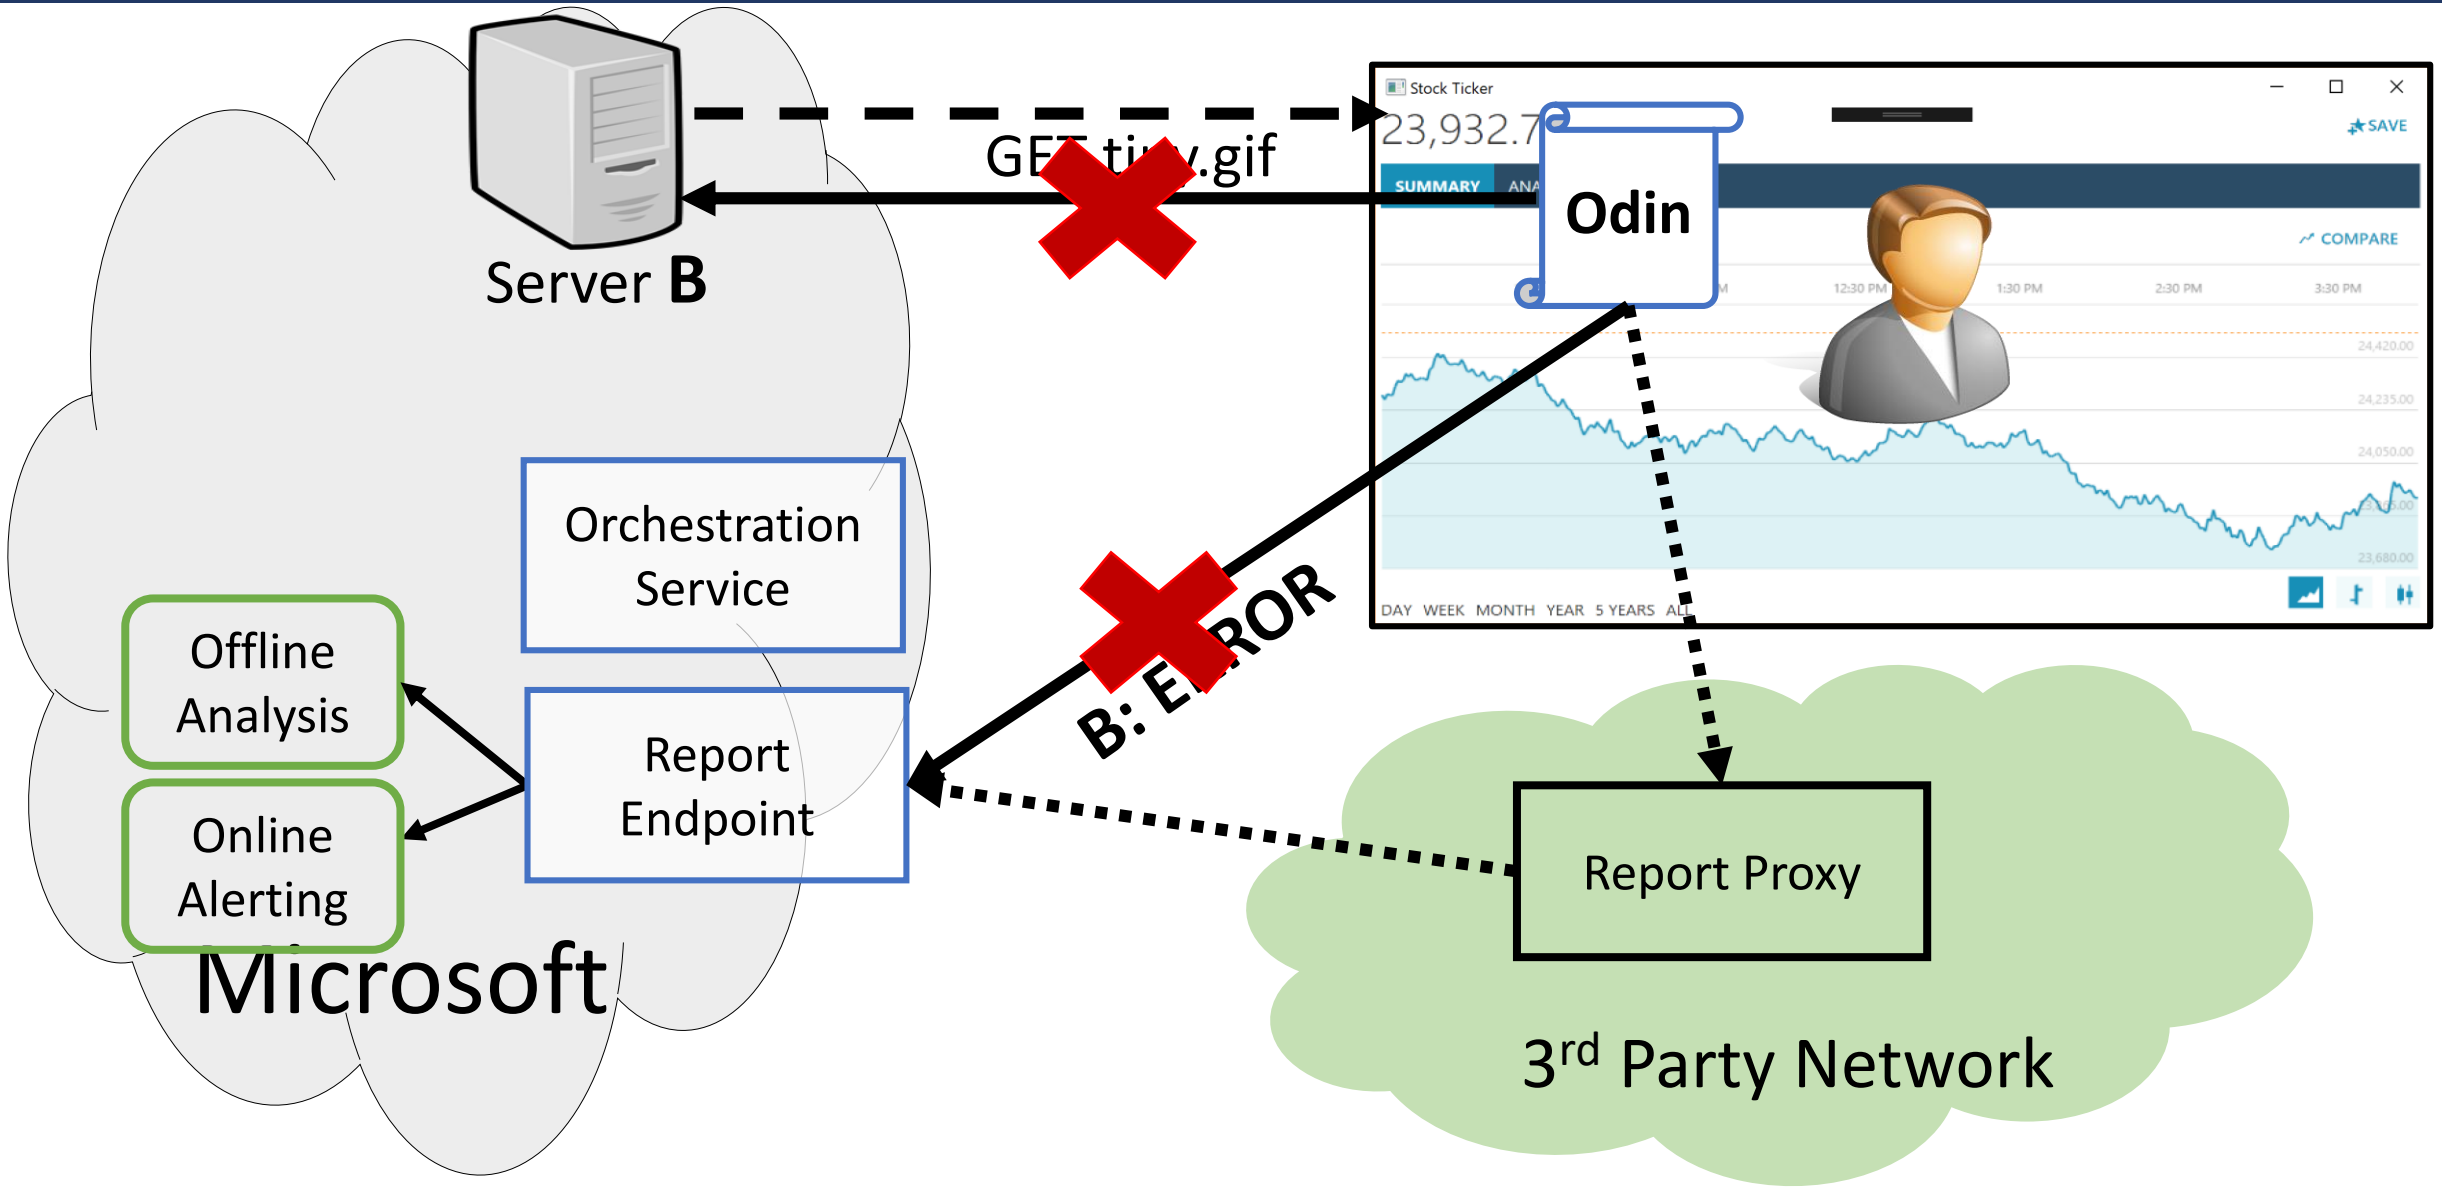
# Odin Design



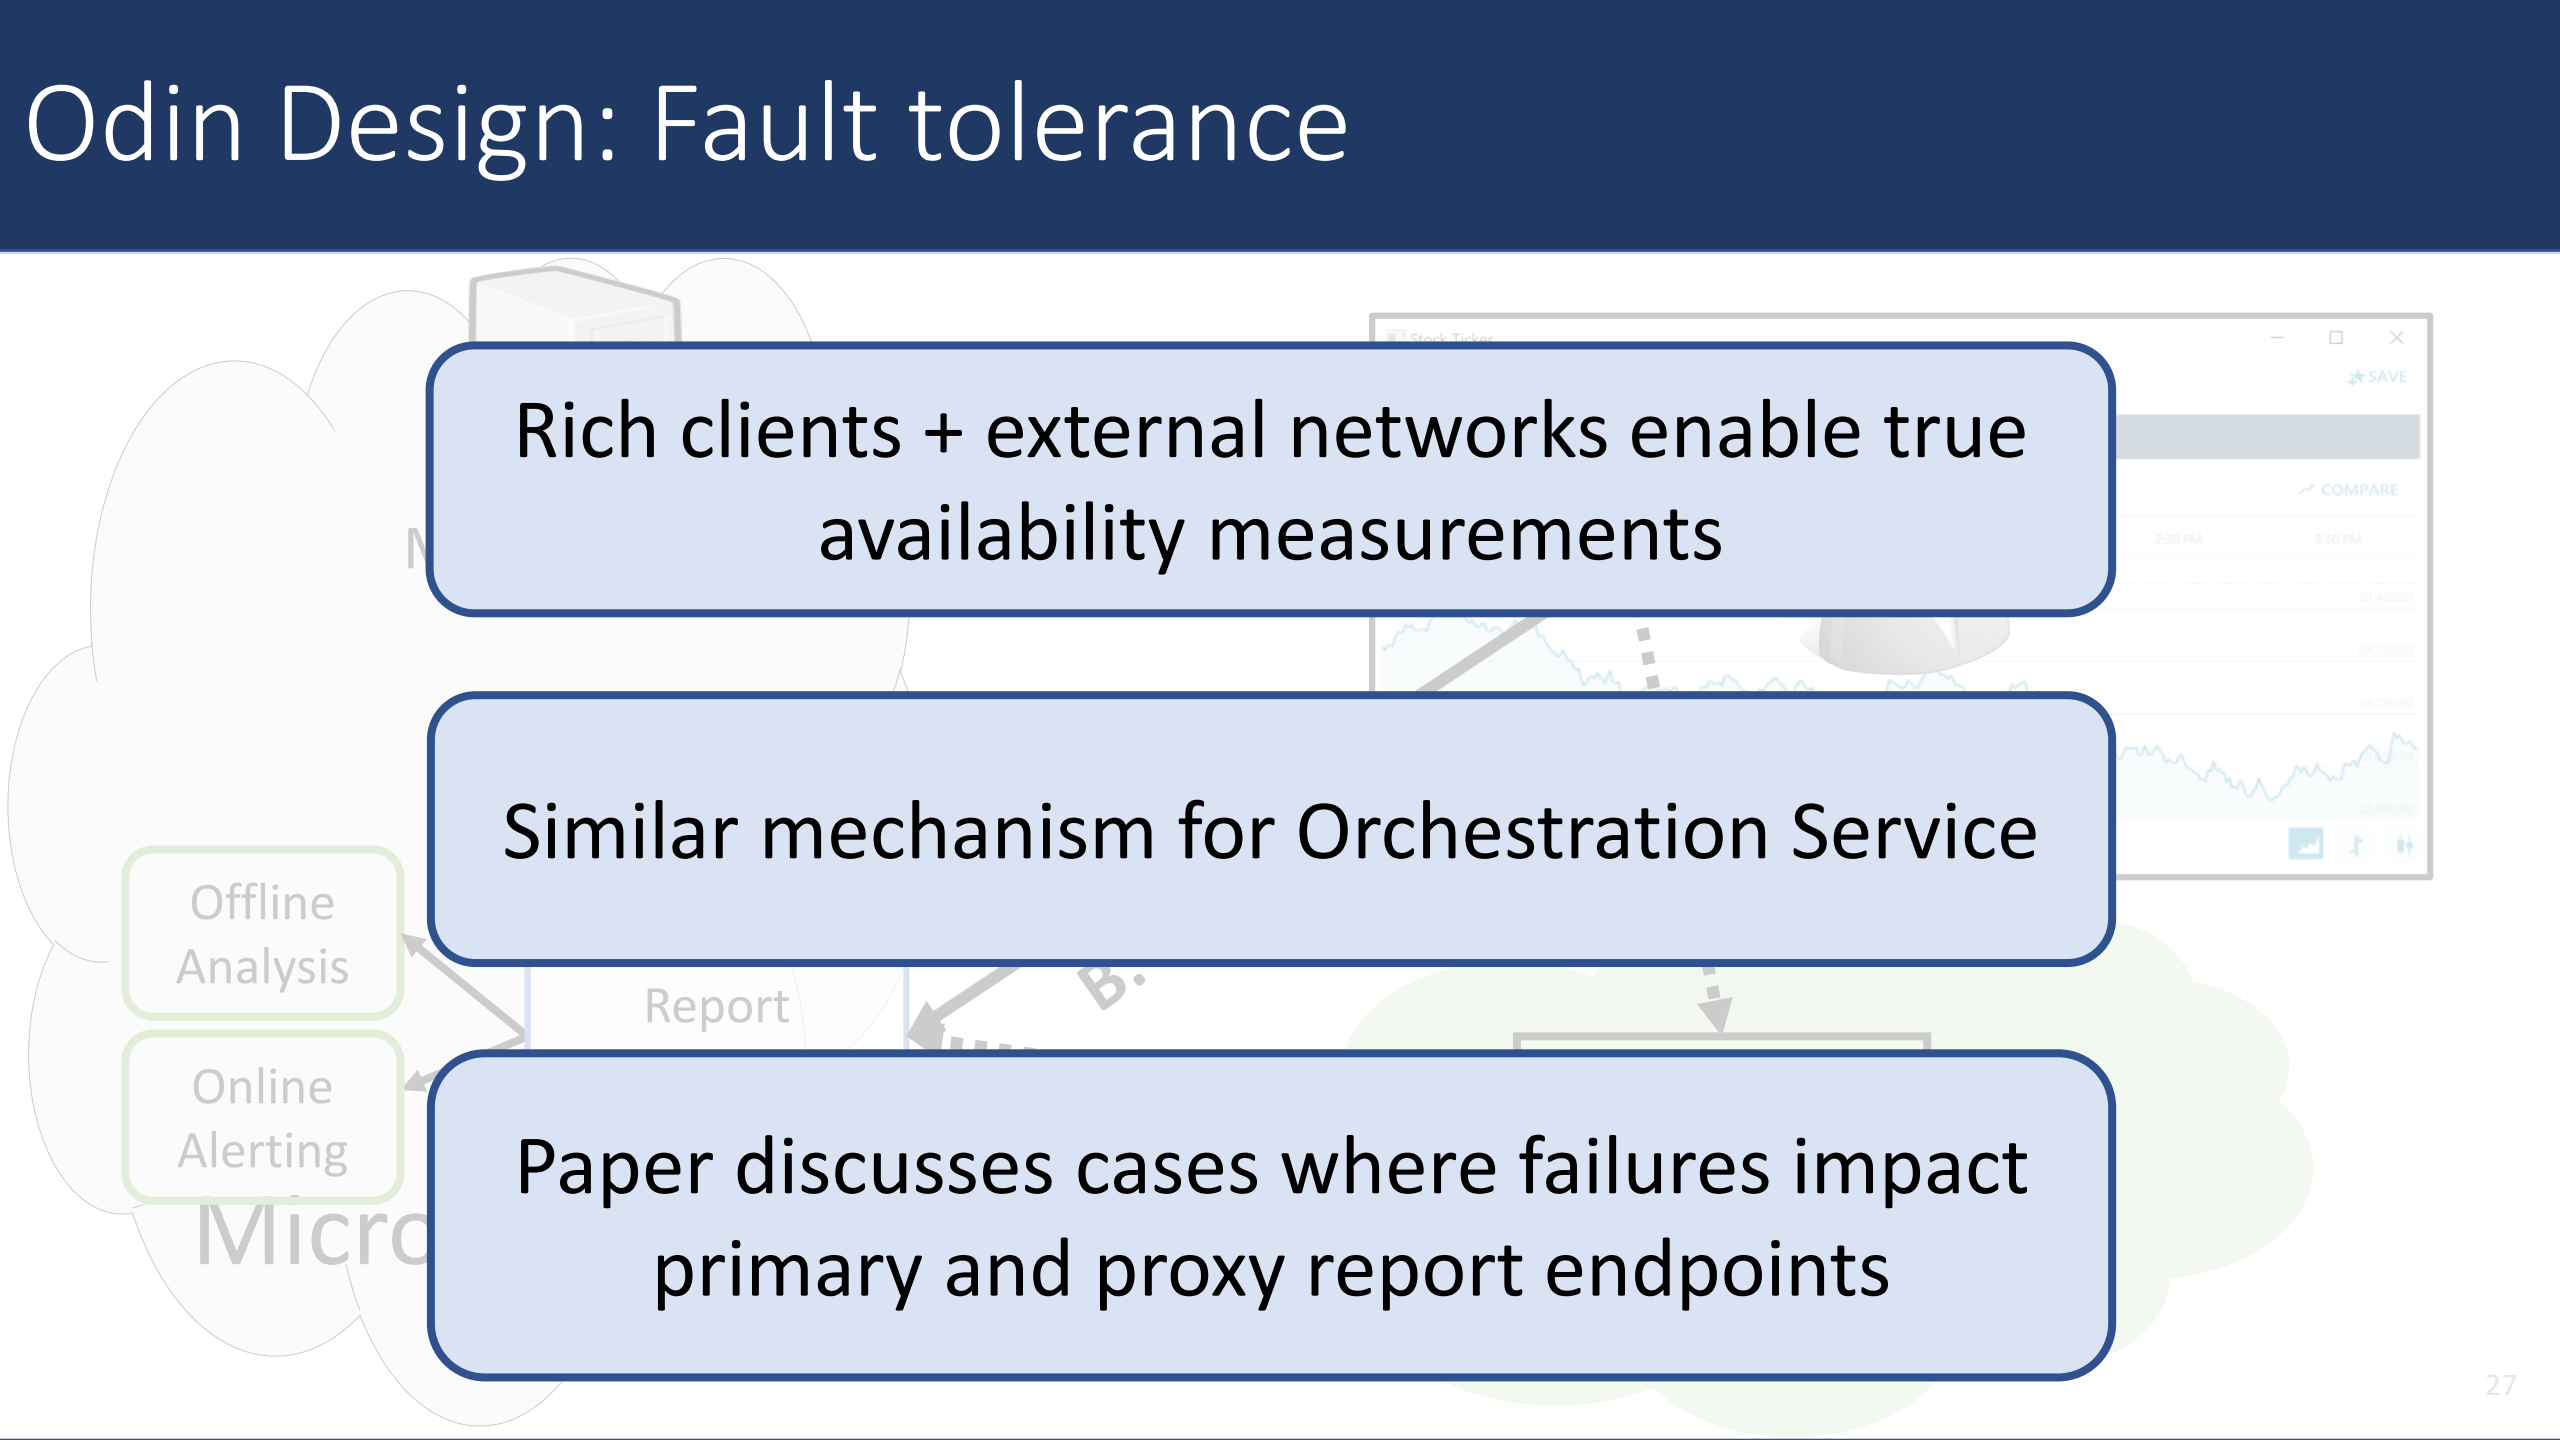
# Odin Design: Fault tolerance



# Odin Design: Fault tolerance



# Odin Design: Fault tolerance



Rich clients + external networks enable true availability measurements

Similar mechanism for Orchestration Service

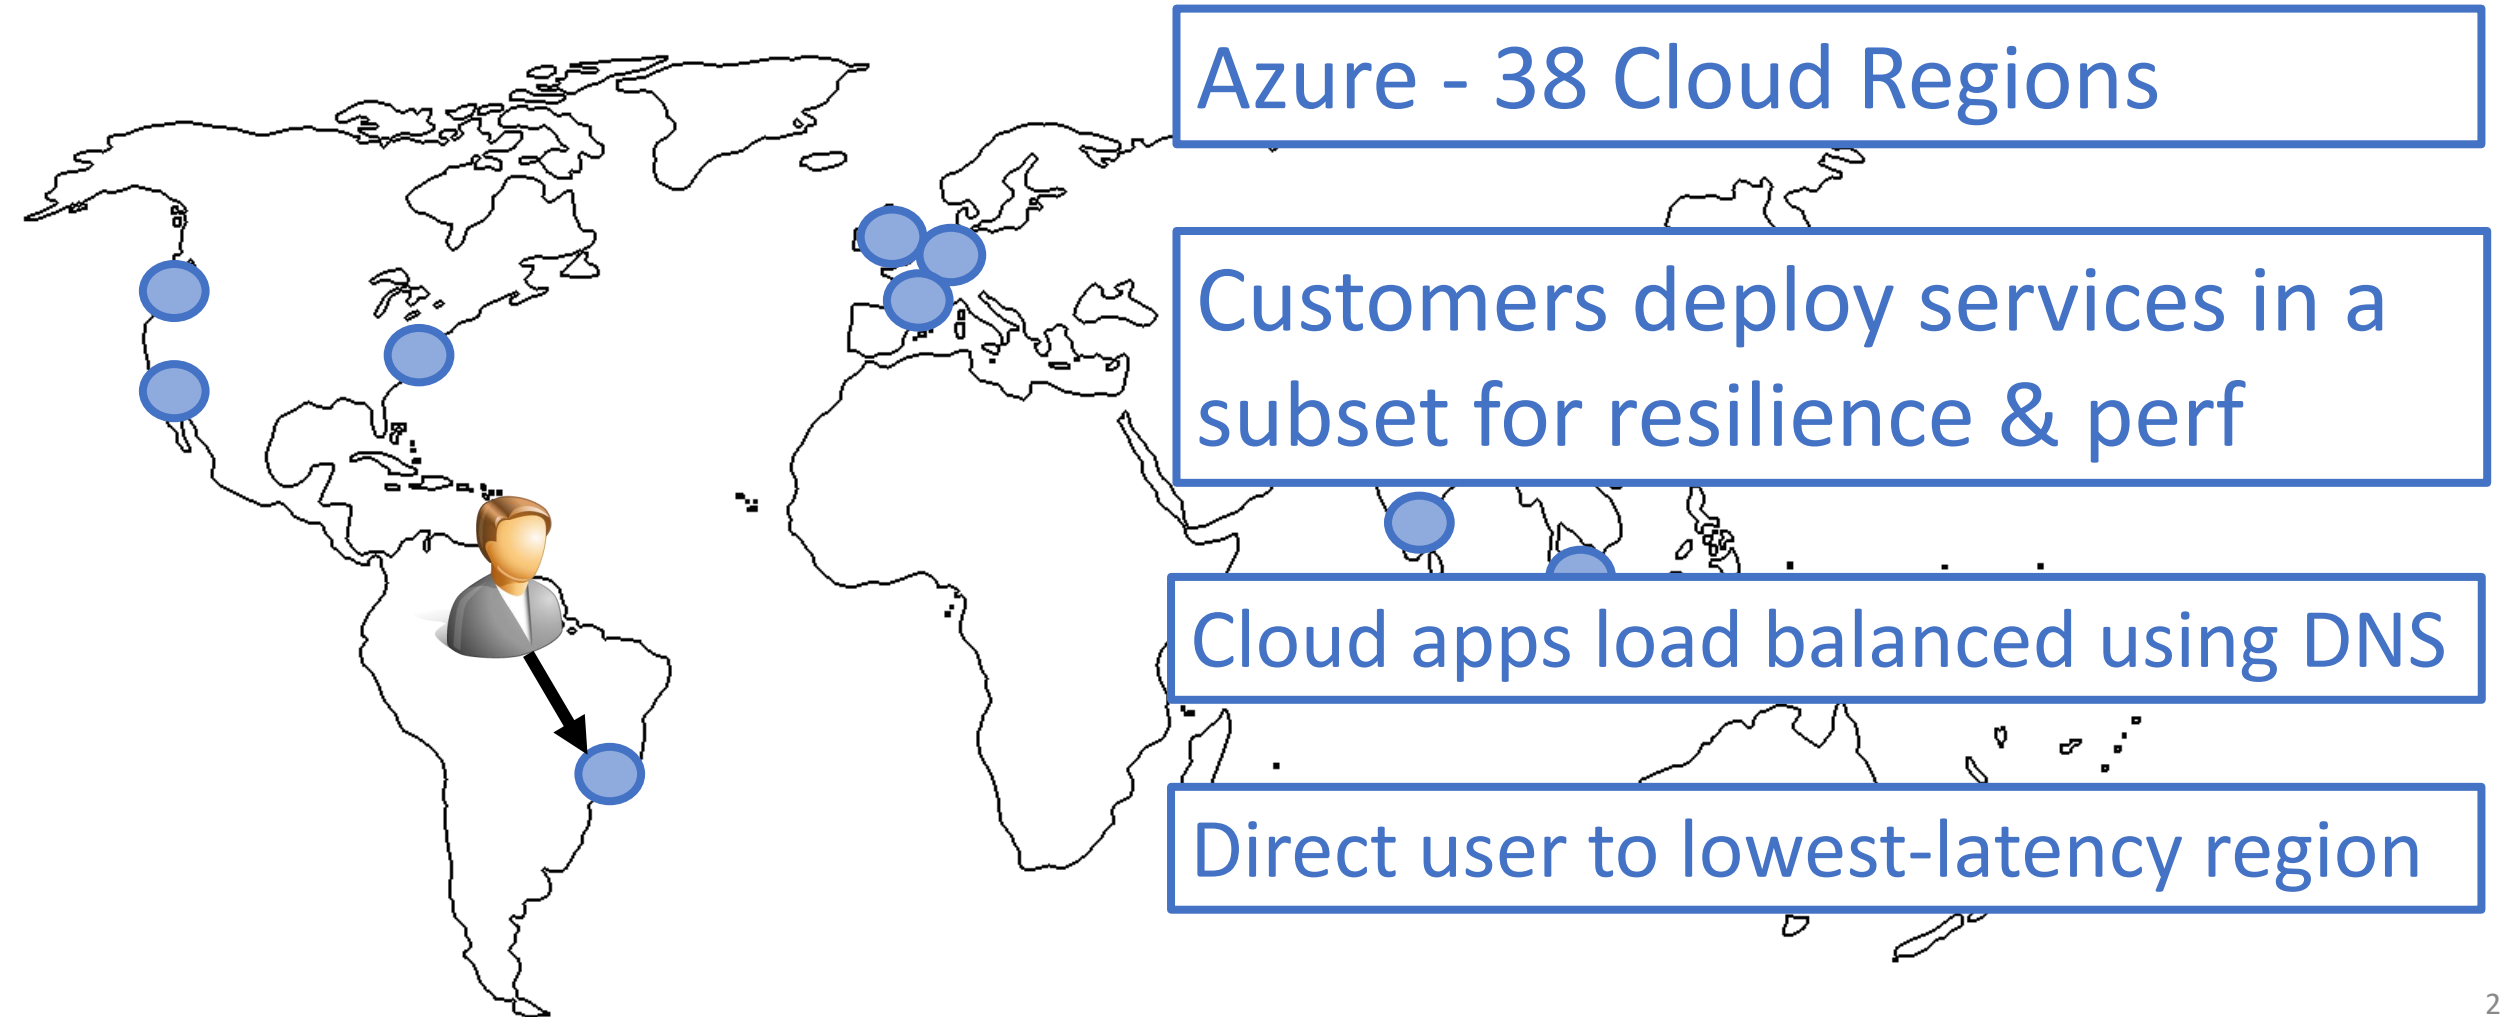
Paper discusses cases where failures impact primary and proxy report endpoints

# Outline

1. Introduction
2. CDN Operations
3. Existing Solutions
4. Odin Design
- 5. CDN Operations with Odin**
6. Conclusion

# Supporting Operations: Odin's DNS redirection maps

## Azure Global Application Traffic Management



# Supporting Operations: Odin's DNS redirection maps

Azure Global Application Traffic Management

Azure - 38 Cloud Regions



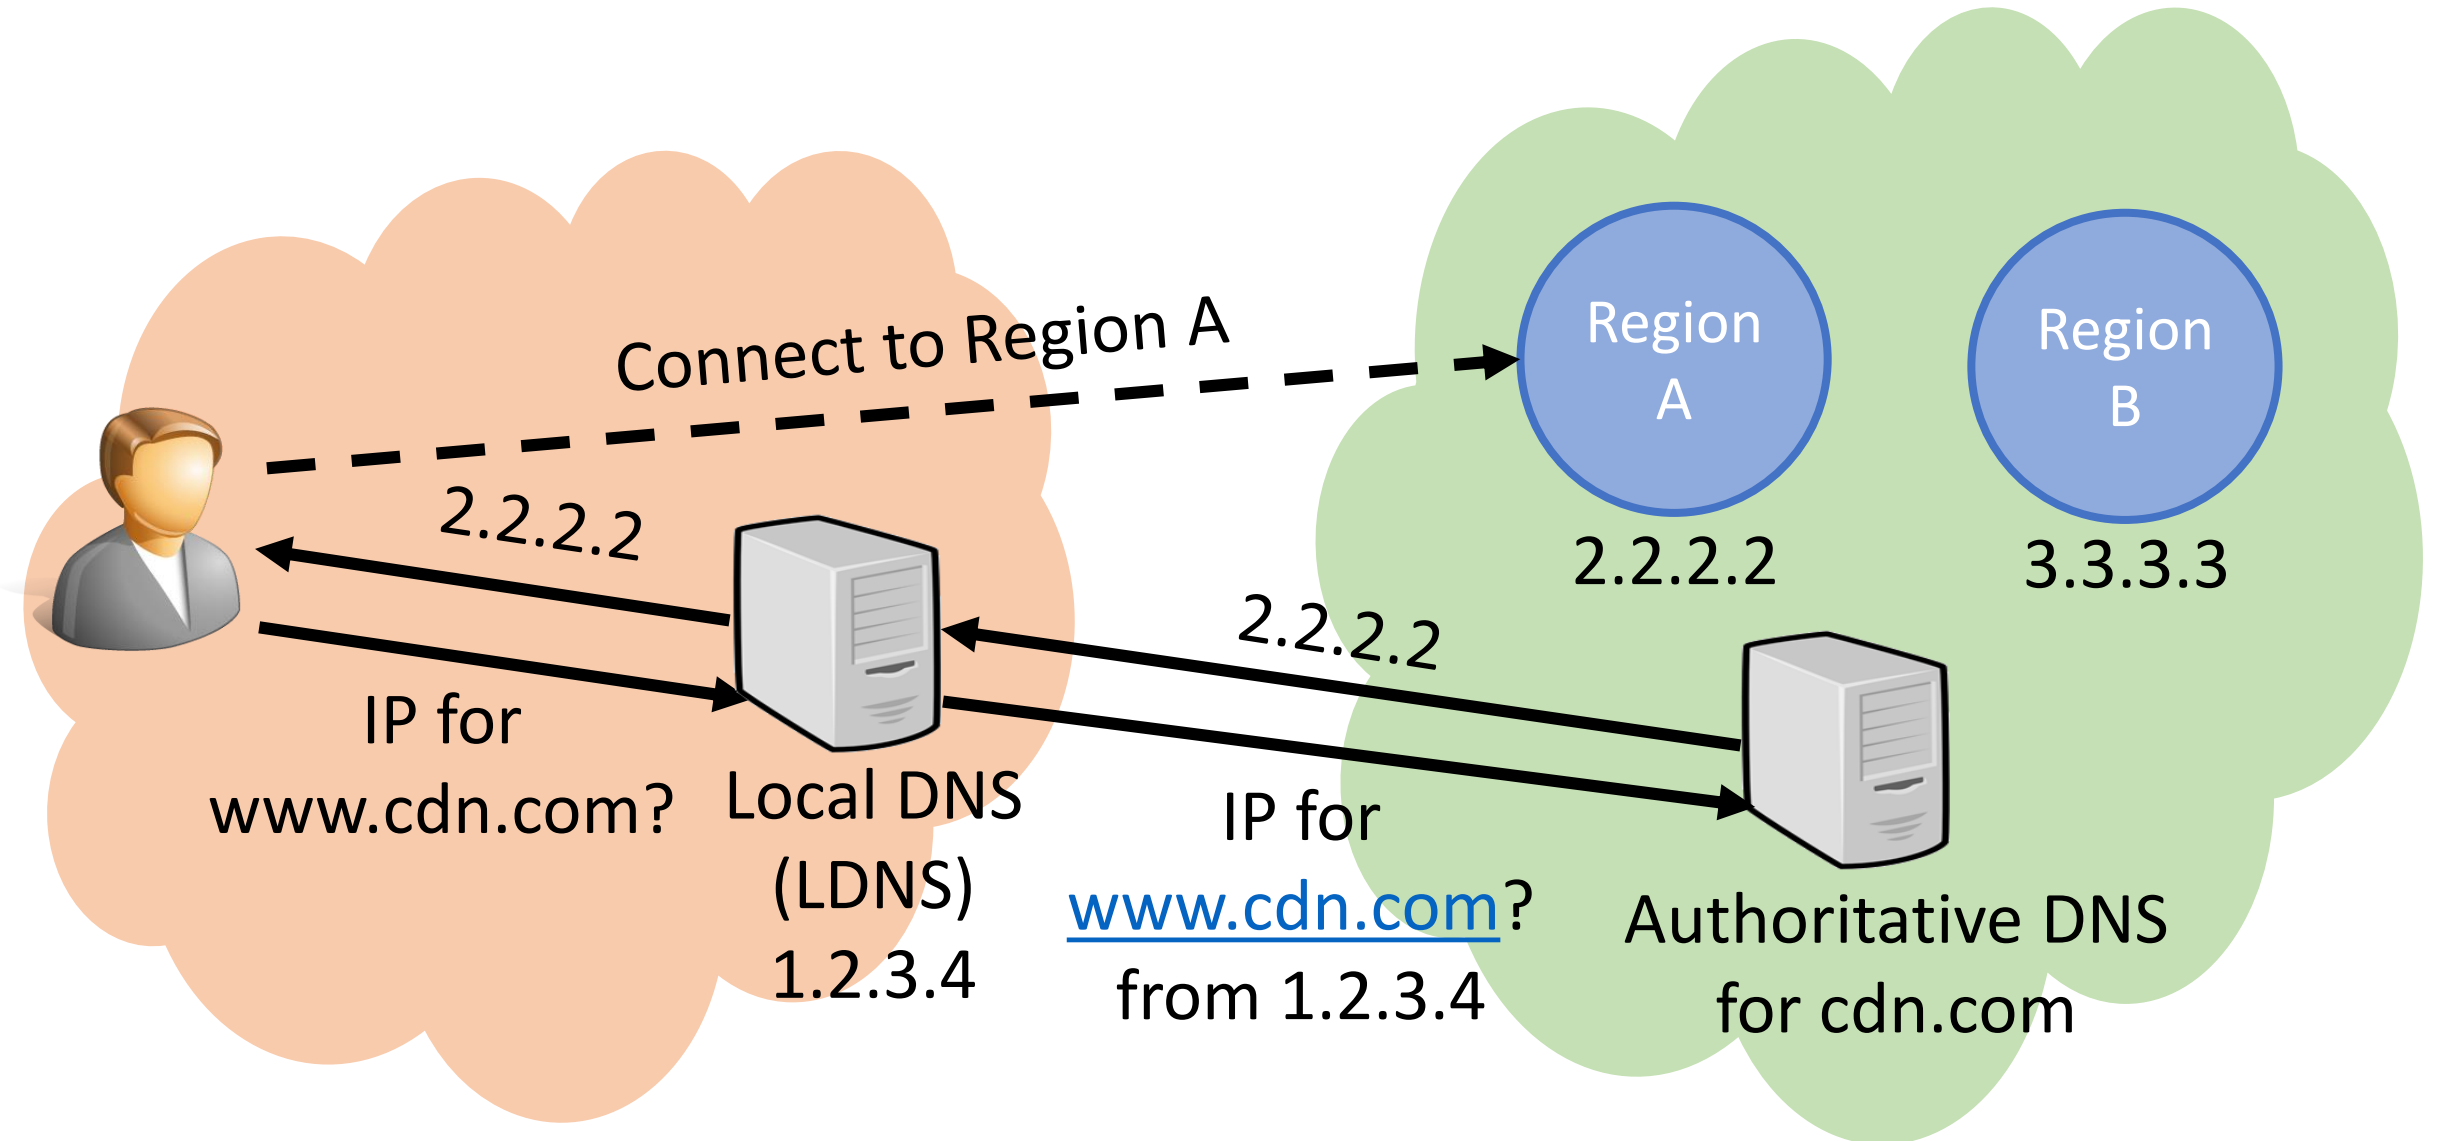
**Odin can find the lowest latency region for users**

Services in a  
& perf

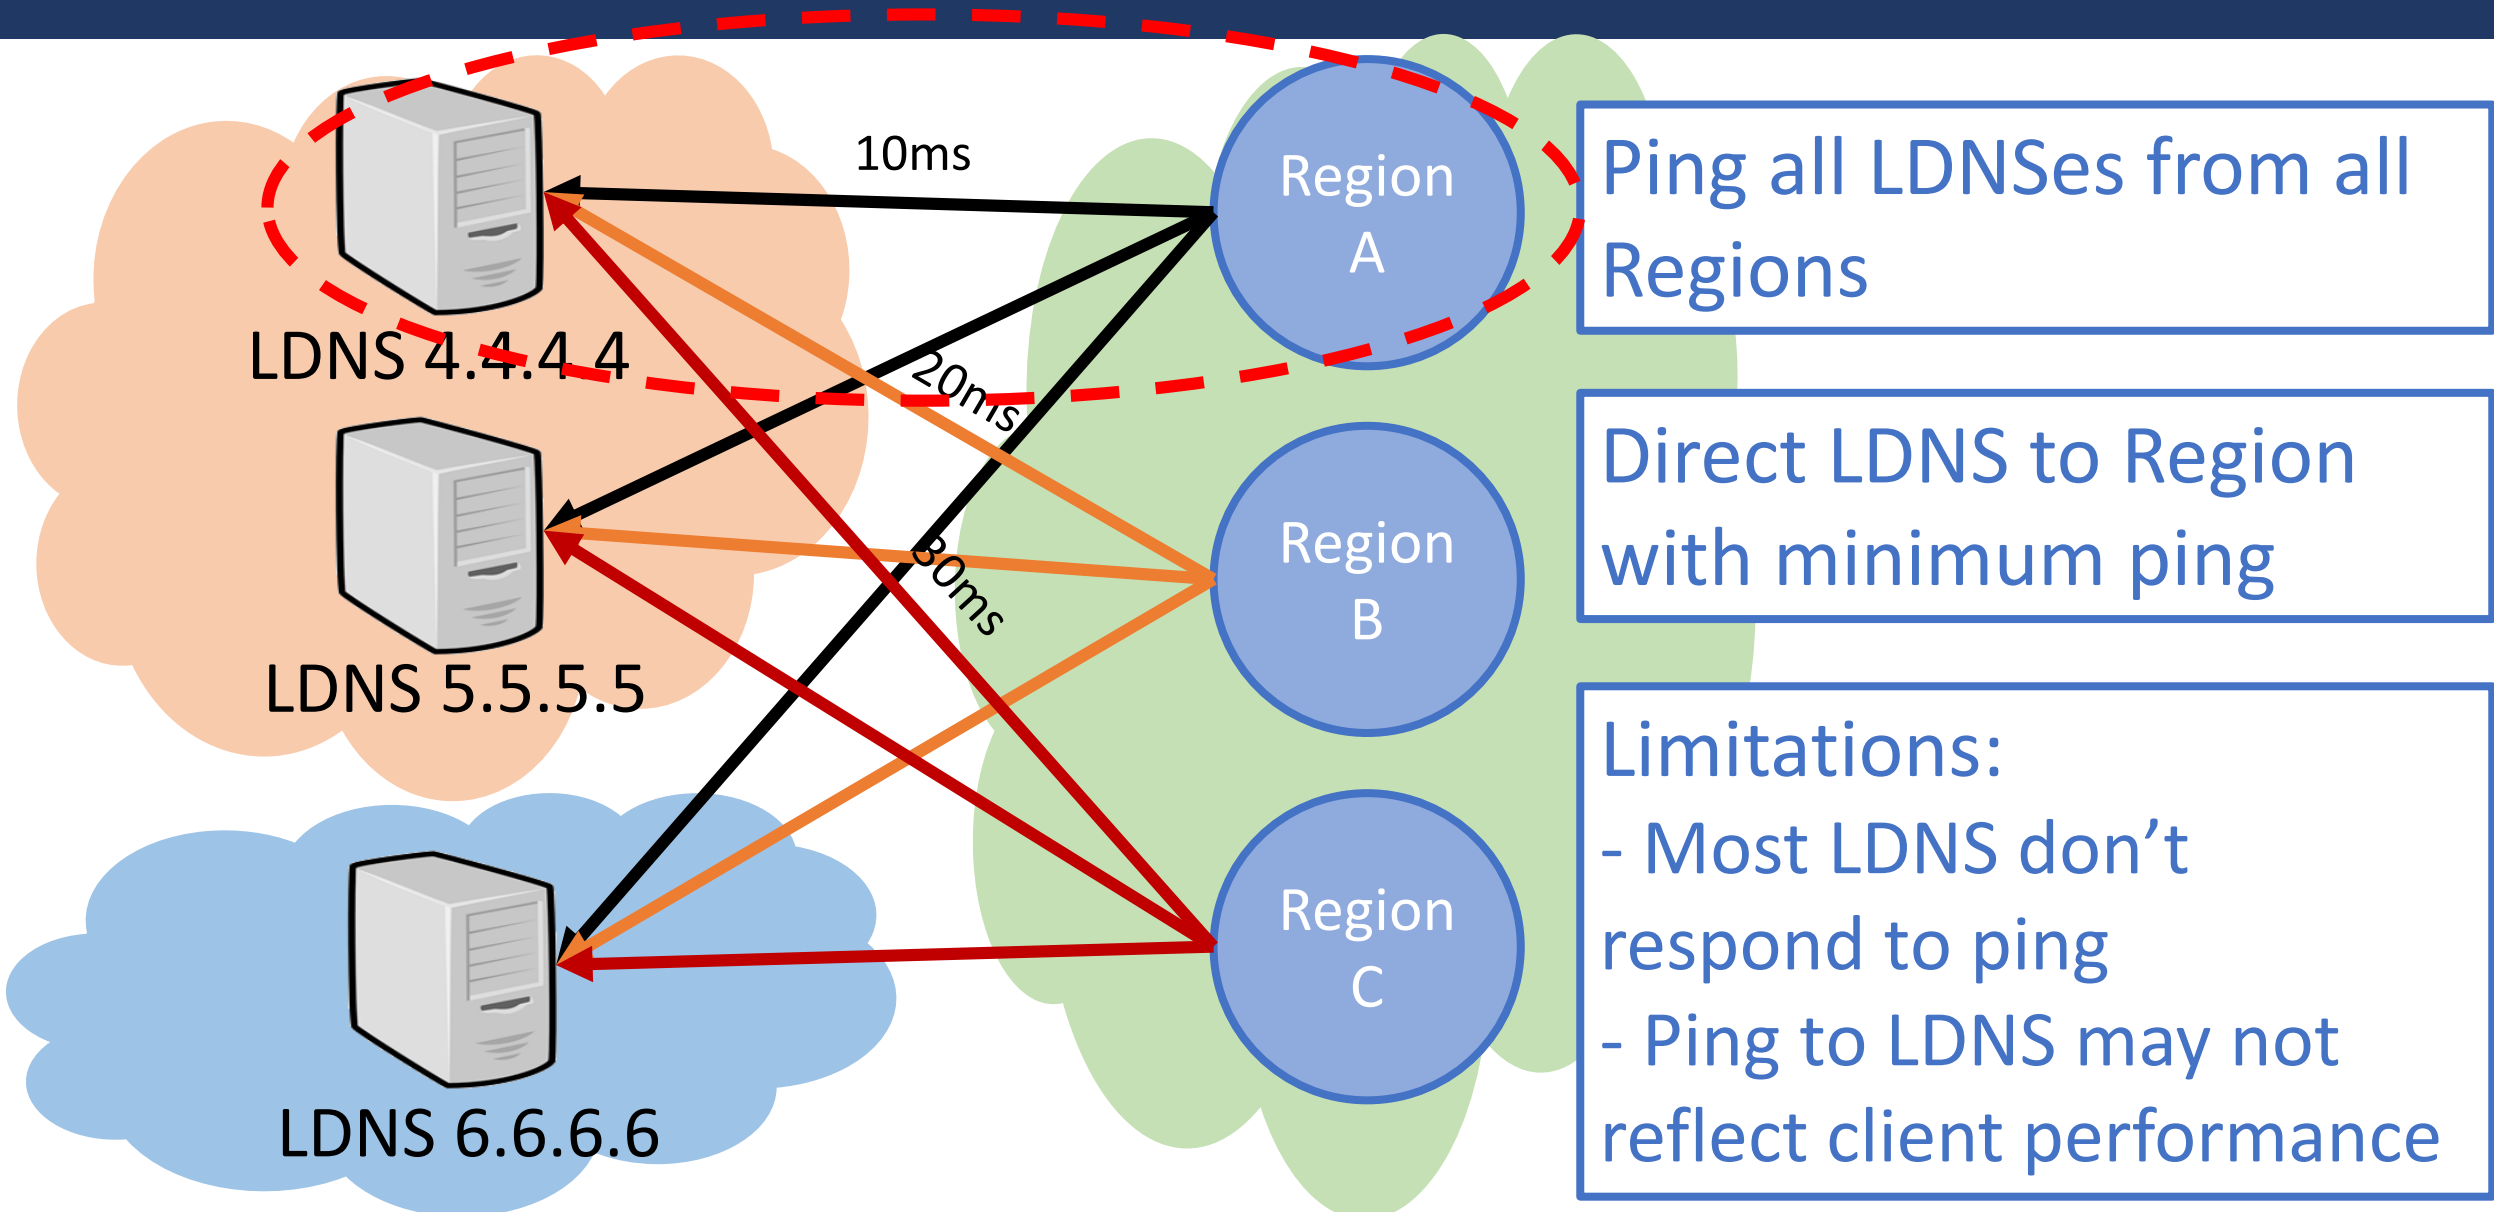
ed using DNS

Direct user to lowest-latency region

# DNS-based redirection for CDNs

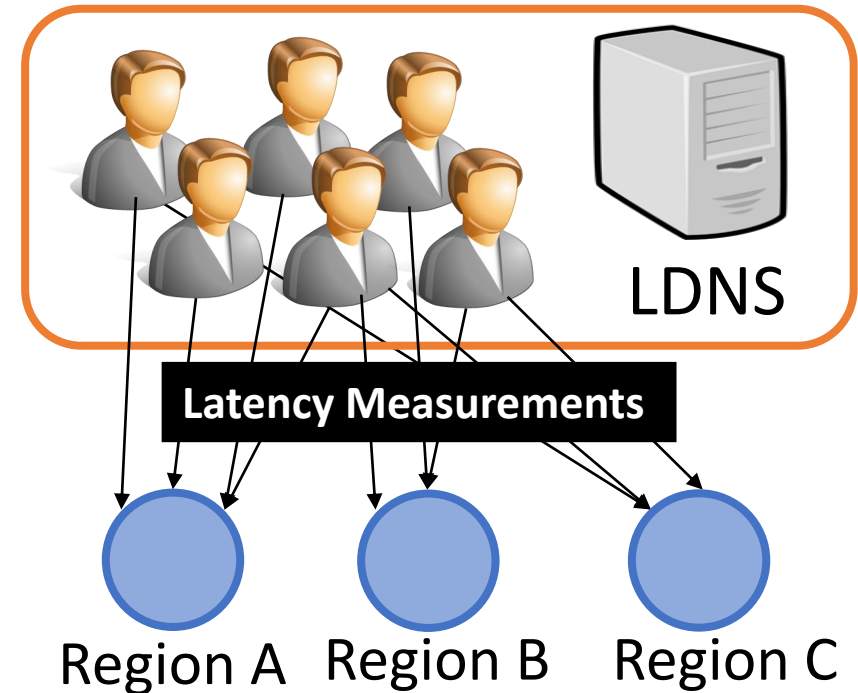



# Previous Approach: Predicting best region for LDNS



# Odin's Approach: Predicting best region for LDNS

- Odin knows client => LDNS
- Know all clients served by an LDNS
- Associate client to region measurements with client's LDNS
- Compare median latency toward each region
- Select lowest latency region as best

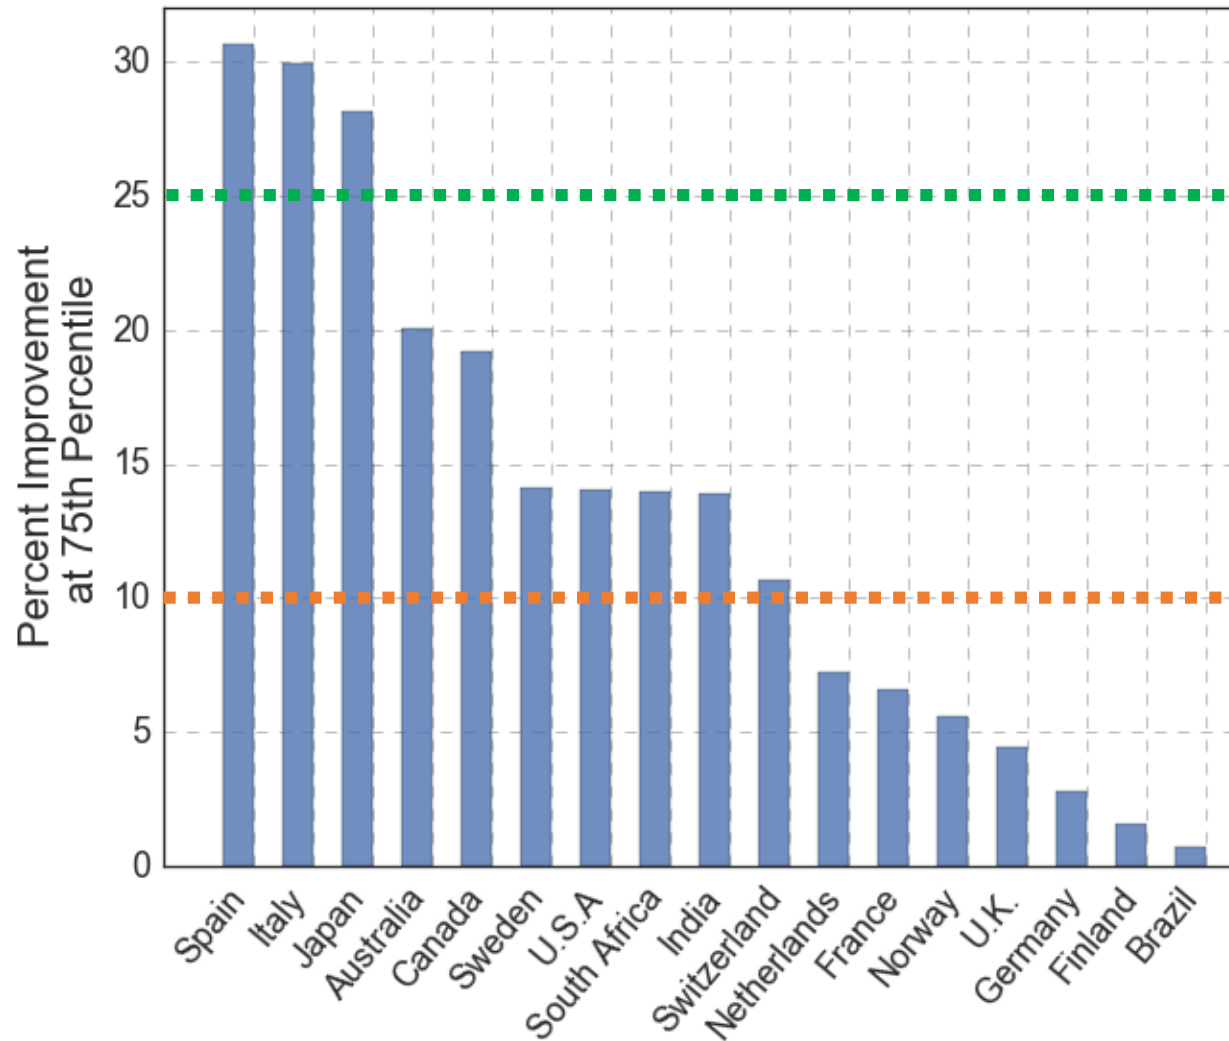




LDNS

Region A	Region B	Region C
25 ms	14 ms	40 ms
...	...	...
30 ms	17 ms	37 ms
35 ms	21 ms	42 ms

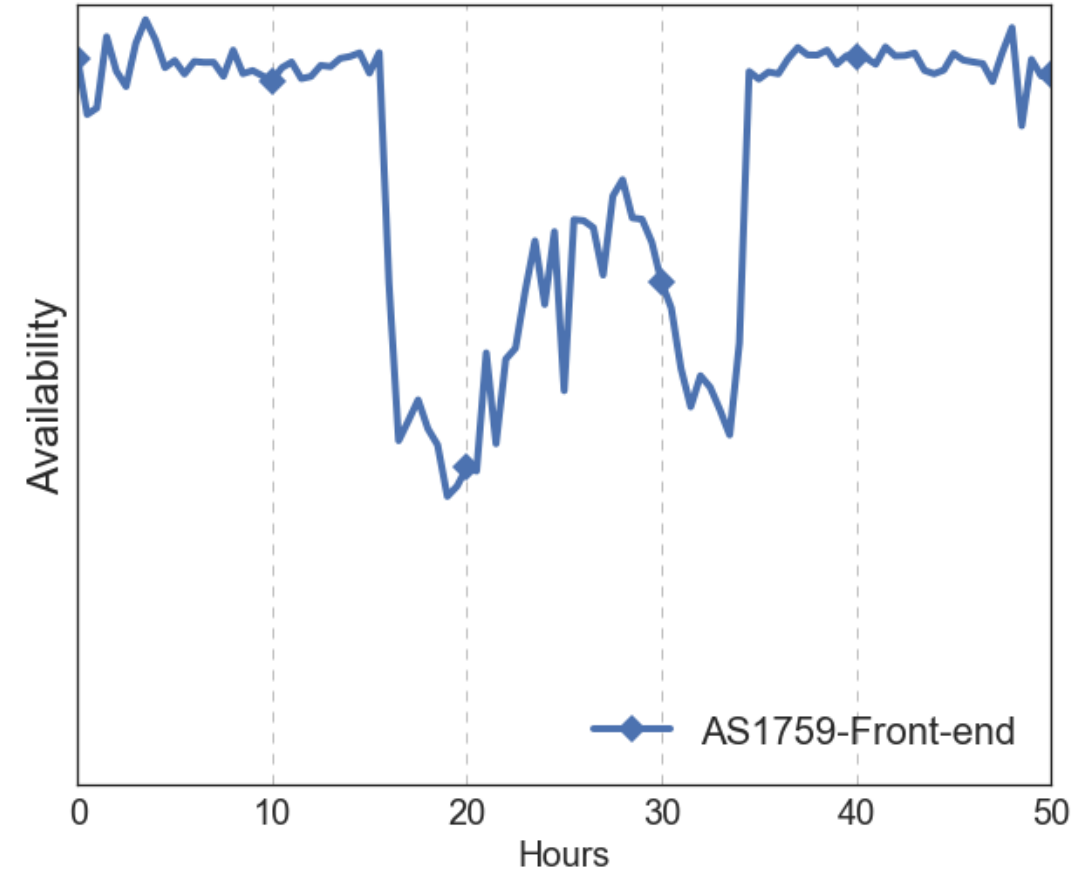
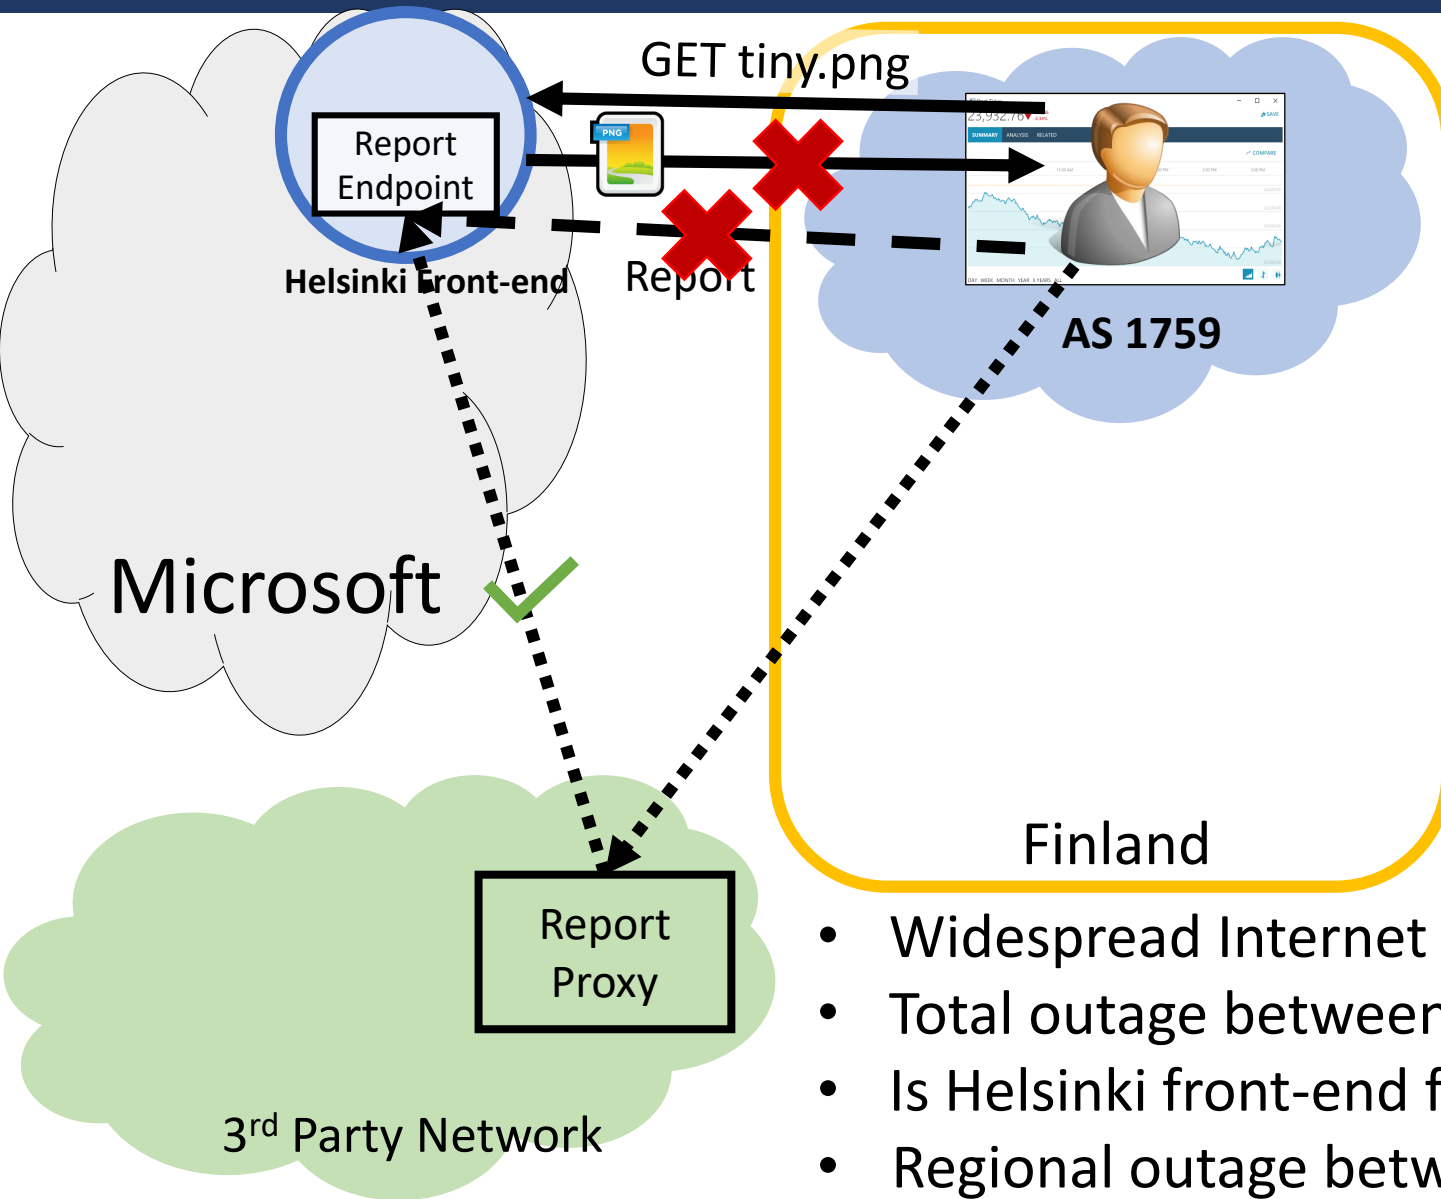
# Improving cloud performance with Odin



Japan, Italy, Spain over 25% latency improvement

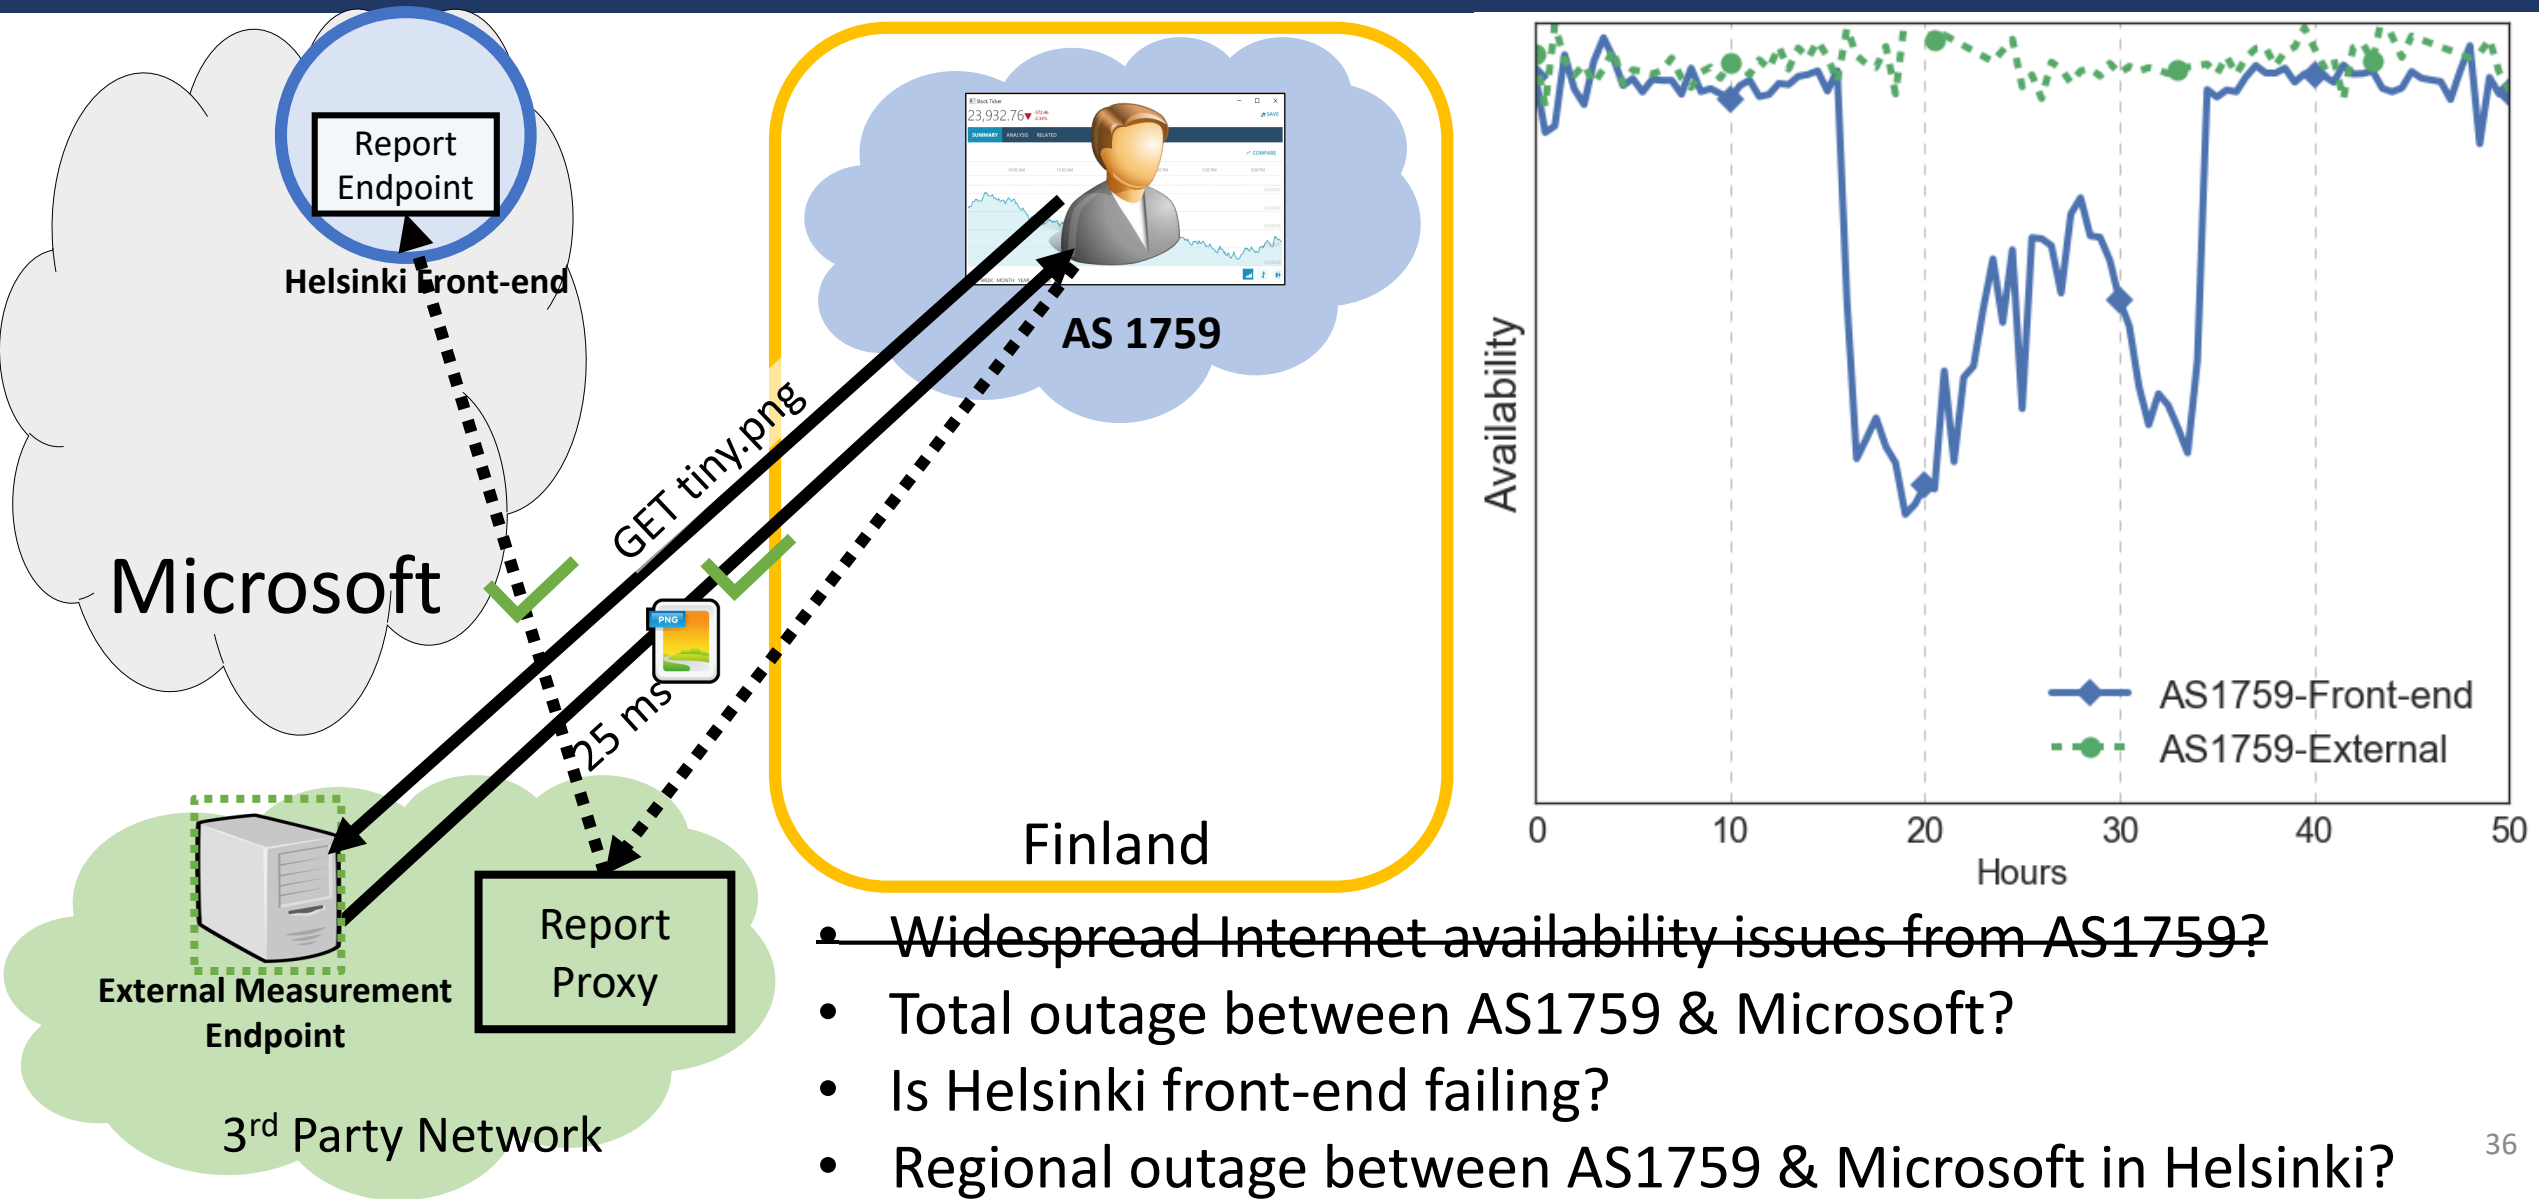
10 countries see > 10% improvement in latency

# Supporting Operations: Using Odin to identify outages

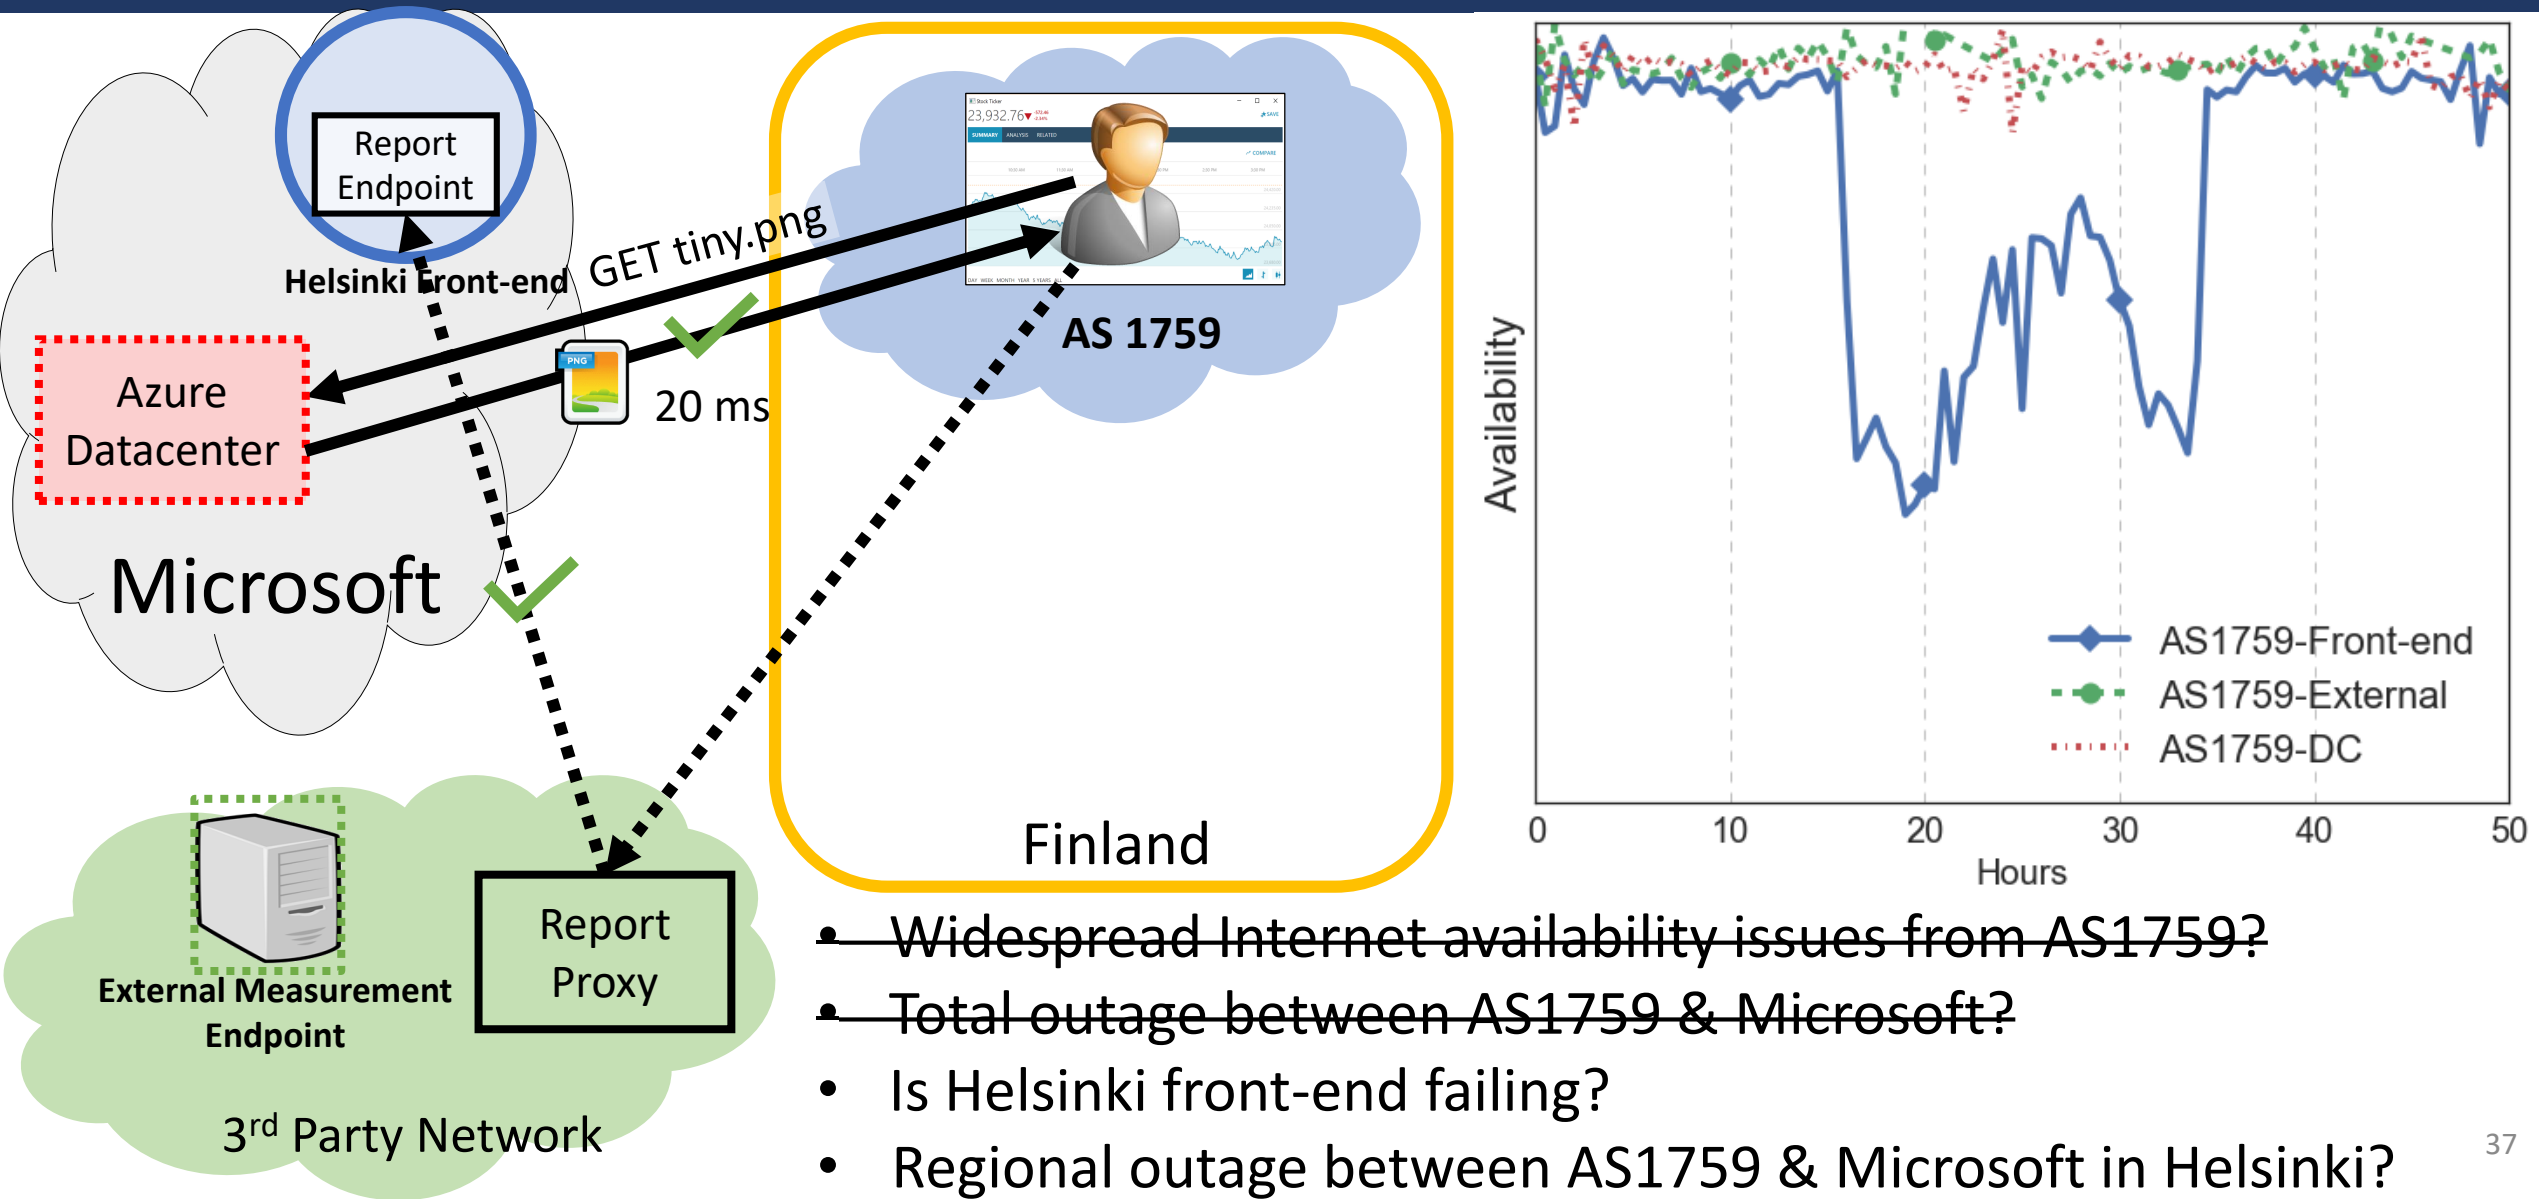


- Widespread Internet availability issues from AS1759?
- Total outage between AS1759 & Microsoft?
- Is Helsinki front-end failing?
- Regional outage between AS1759 & Microsoft in Helsinki?

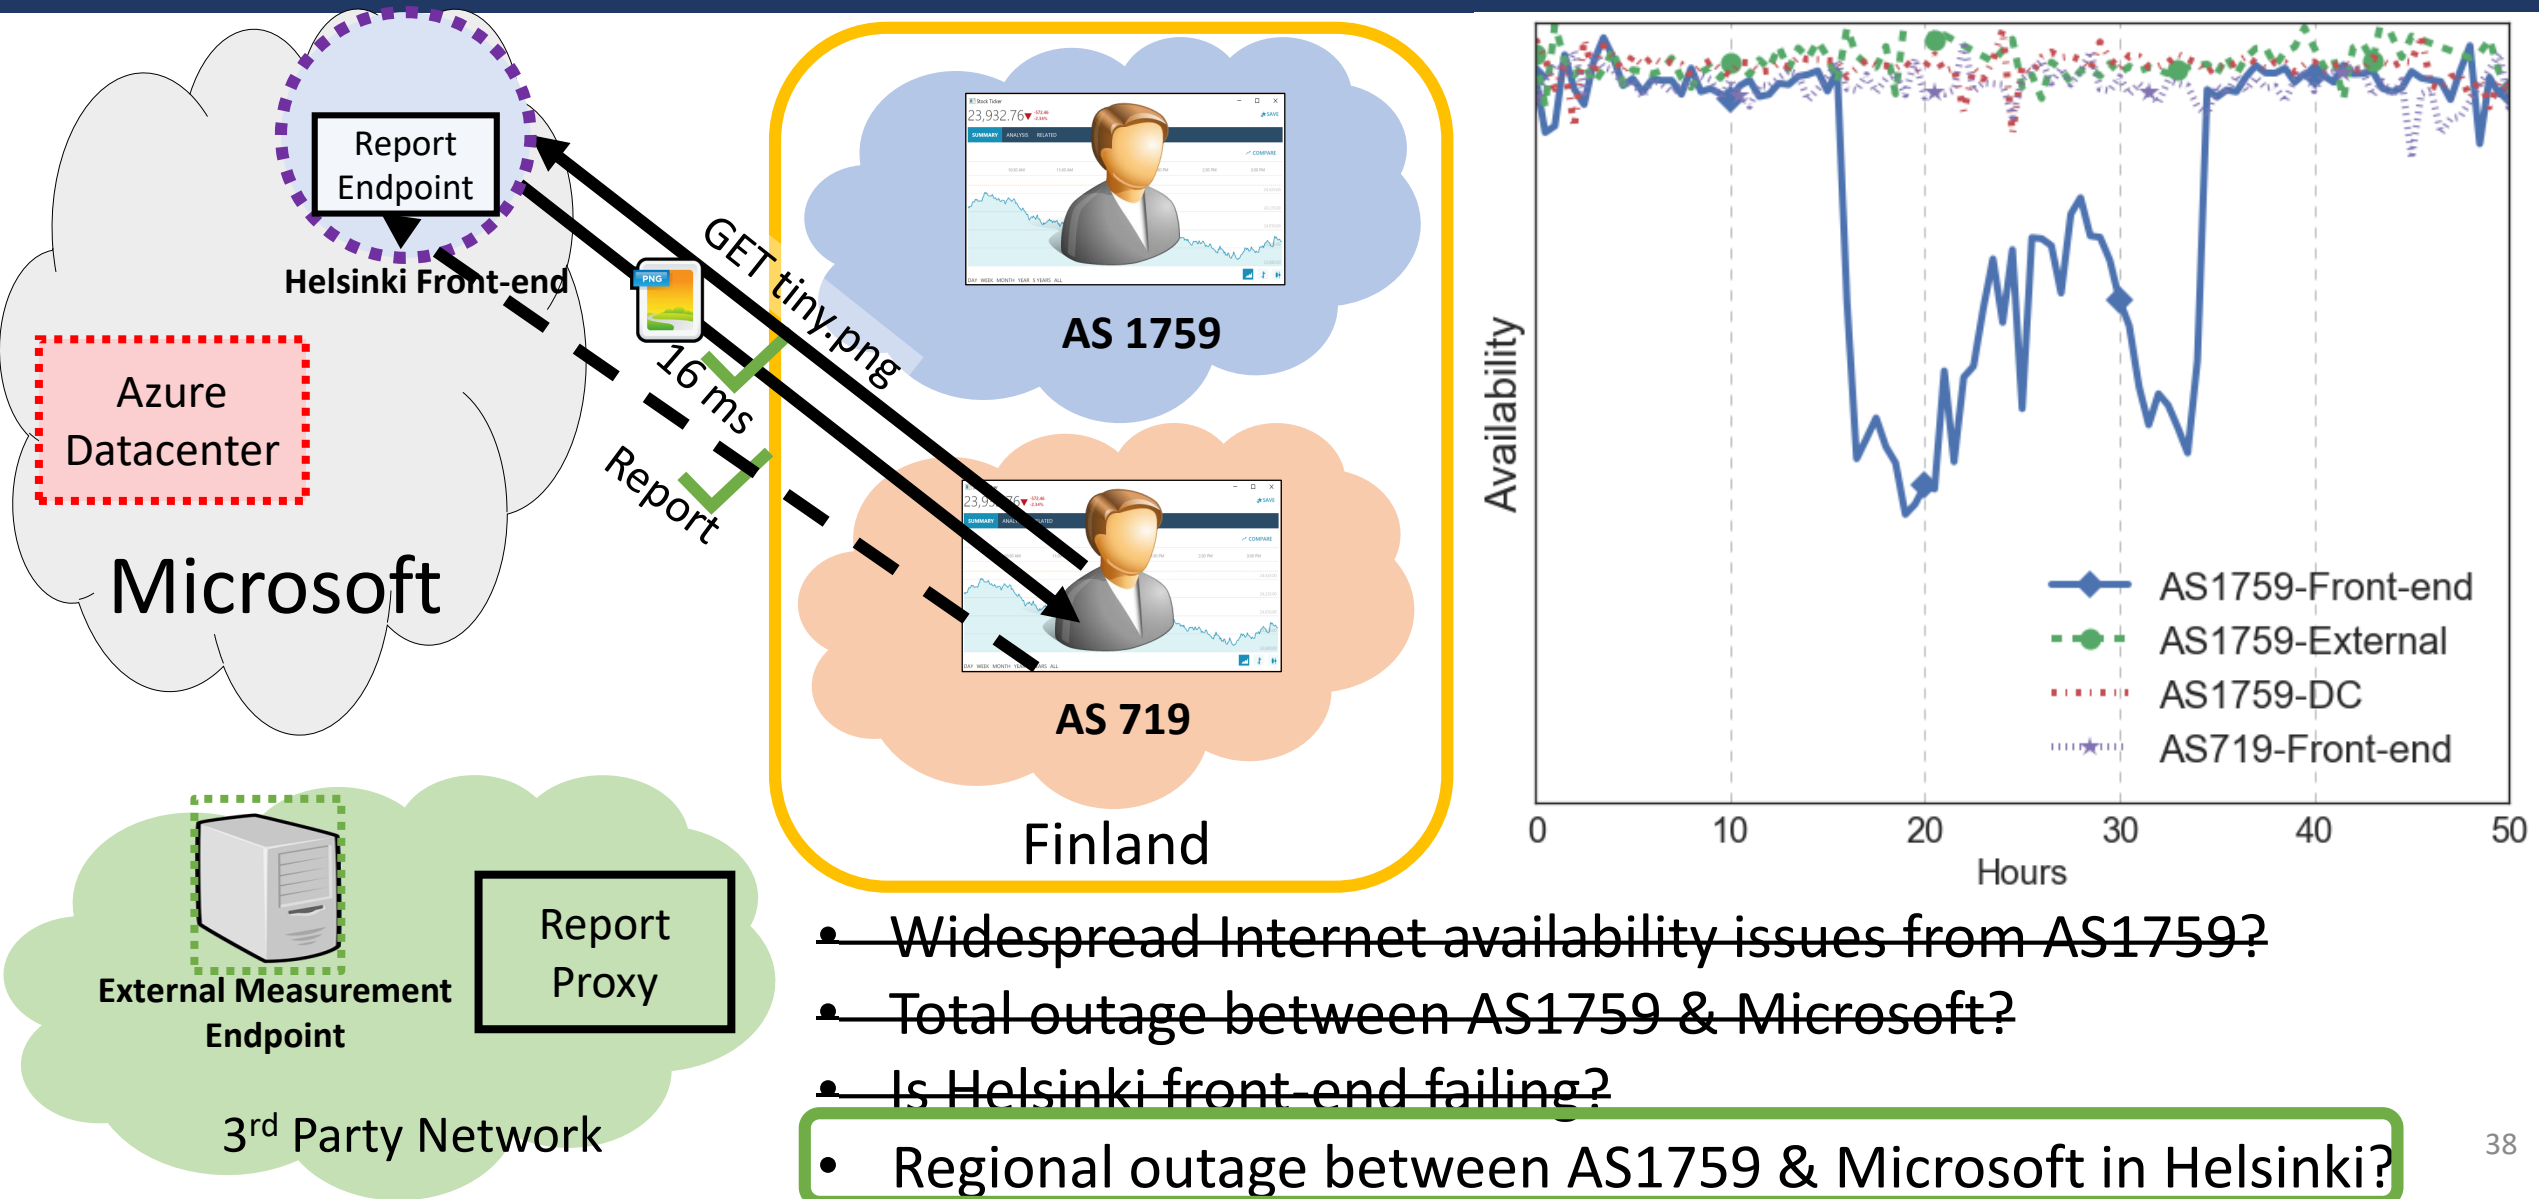
# Supporting Operations: Using Odin to identify outages



# Supporting Operations: Using Odin to identify outages



# Supporting Operations: Using Odin to identify outages



# Conclusion

- **Odin** is a client-side, active measurement platform to support Microsoft CDN operations
  - Running in production for over 2 years
  - Deployed in key Microsoft end-user applications
  - Billions of measurements per day
- Overcomes limitations of existing approaches
- Improves CDN operations
  - Detect and diagnose outages
  - Odin-based DNS-redirection for cloud customers