

## Predicting Failures in Distributed Cloud-Based Systems

Sebastian Moreno, Andrew Newell, Rahul Potharaju, Cristina Nita-Rotaru, and Jennifer Neville  
*{smorenoa, newella, rpothara, crism, neville}@cs.purdue.edu*

### Introduction

- We analyze a cloud networking provider supplying high quality networking services to customers.
- The system consists of a set of geographically distributed routers in an overlay network (a virtual network on top of the Internet), which relay packets to/from clients.
- The system aims to deliver packets within 200ms to/from any clients, 100% of the time—if a routing path goes down for even just a short period of time (e.g., seconds), it will negatively impact their performance.
- Any outage is fixed in a short period time with human intervention, but there are no techniques to accurately predict these outages.


### Goals of the project

1. Analyze logs files to explore behavior and correlation in data.
2. Identify signatures of “outage” events to determine examples to use for learning.
3. Learn classification models to predict upcoming outage events based on temporal behavior.


### Data

- The 2Tb dataset consists of multiple log files with different types of statements that are output at varying frequencies.
- *Local logs* are recorded at the source and destination clients.
- *Remote logs* are recorded in the overlay routers that are connected to the clients.
- *Routing logs* are recorded locally for every overlay router in each overlay network.

### Preliminary analysis



**Overlay network:**  
 Edges represent connections among servers (ovals). Clients (squares) connect to multiple servers. Server node size represents number of active clients (blue squares).





**Server connections:**  
 Variation in number of clients connected to each server over time (clients switch servers in response to network delays).

### Event identification

- Serious “outage” event: Clients try to change between servers then leave the network for connection problems among servers.
- Eight serious events were detected: two network resets, two local problems, and four important failures.

#### Outage event:

The majority of network connections have gone down (red edges). Servers do not have any active clients. Several clients migrate from the overlay network to deliver their content.



#### Event identification:

Number of clients connected to the overlay network drops to zero (green), while number of edges that fail in the routing network increase considerably (blue), for a non-trivial amount of time.

### Event prediction

- Current work: develop automated methods to discriminate among different types of failures and predict possible failures.
- Given the positive events identified above, we can learn (i.e., train) models based on temporal information in the client and server logs.
- Models that can accurately predict possible failures with even a few seconds of lead time will allow for automated adaptation mechanisms to prevent large scale failures.