IBM Research

Towards Automated Identification of Security Zone Classification in Enterprise Networks

IBM: Hari Ramasamy, Birgit Pfitzmann, Nikolai Joukov, Jim Murray

Georgia Tech: Cheng-Lin Tsao

USENIX Hot-ICE 2011, Boston

Background: Firewalls, Security Zones

- Enterprise network infrastructures are divided into *zones* of varying criticality
- Zone: set of devices of same security requirements
  - Guarded by boundary firewalls
  - Security requirements defined in *enterprise policy*, (hopefully) enforced by *network configuration*

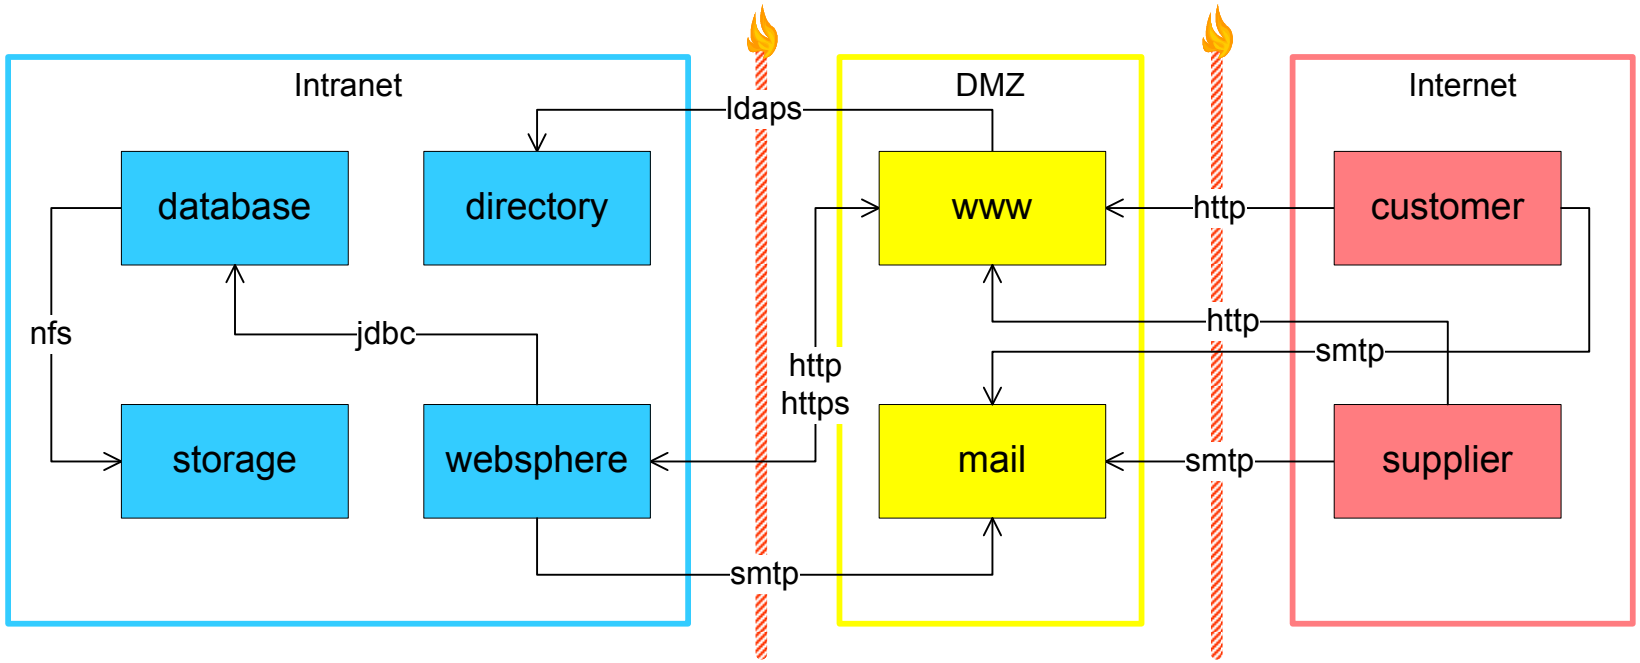

Network Model and Policy Model

$$Host \in Subnet \subseteq Zone \in Classification(Color)$$

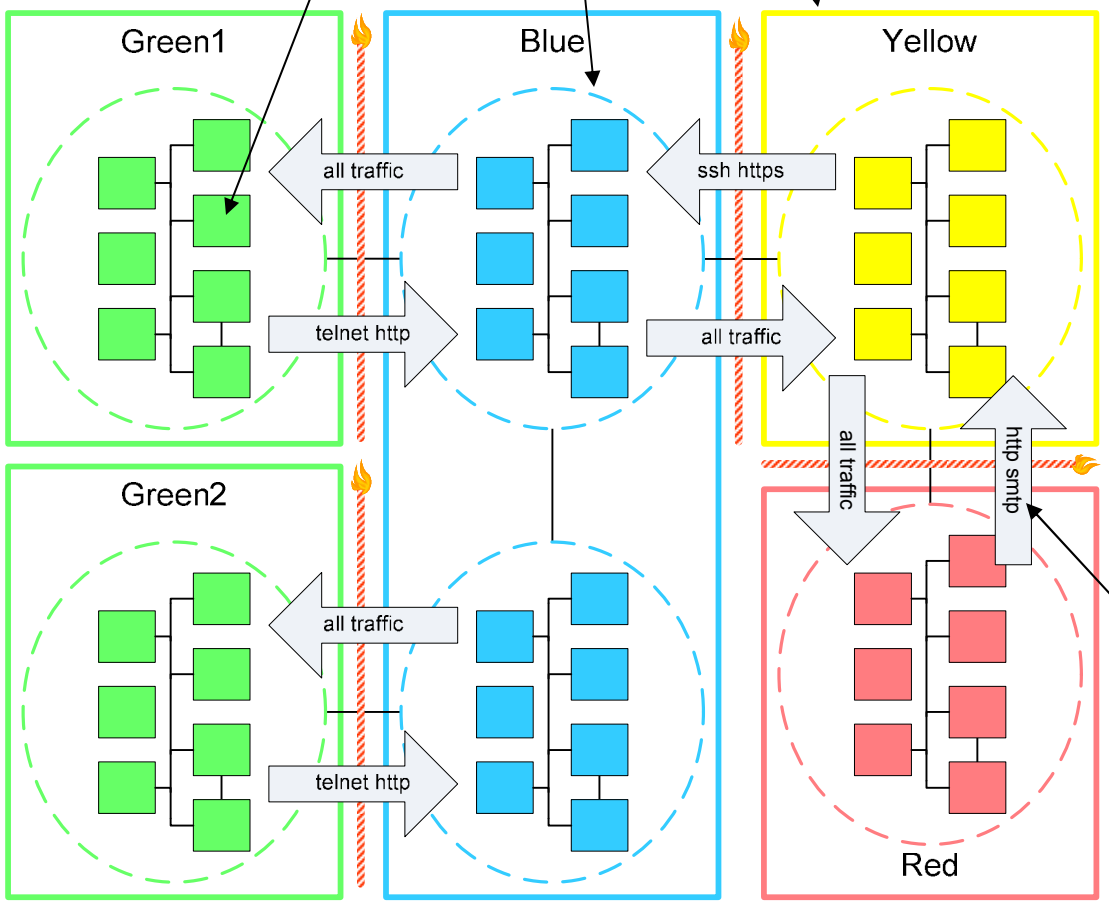

To / From	Blue	Green	Yellow	Red
Blue	All	All	All	All
Green	Auth	All	All	All
Yellow	S.Auth	S.Auth	All	All
Red	None	None	All	All

Enterprise Policy versus Network Configurations

Problem Statement – Zone Discovery

- Input
  - Devices and policy
    - Color of some devices known a priori
- Output
  - zones, colors, interconnections between zones

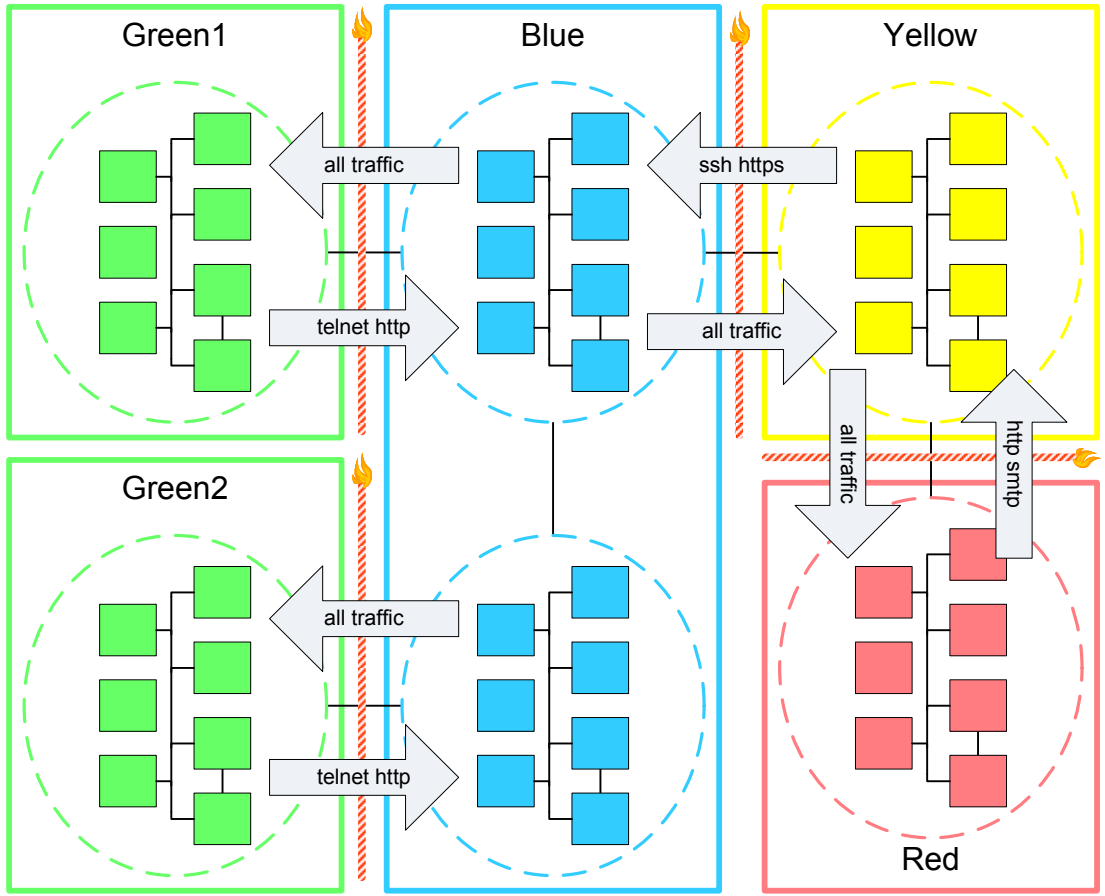

Motivation for Security Zone Discovery

- Even medium sized enterprises may have hundreds of security zones
- Information about zones is *required* in many IT management situations
  - System Migration and Storage Consolidation
  - End-to-end Security Assessment
  - Network Rearrangement or Optimization
- An enterprise-wide inventory of zones is simply absent in many enterprises
- Information about zones is synthesized manually, and often incomplete
- Existing tools can analyze network configs, but don't yield zone information

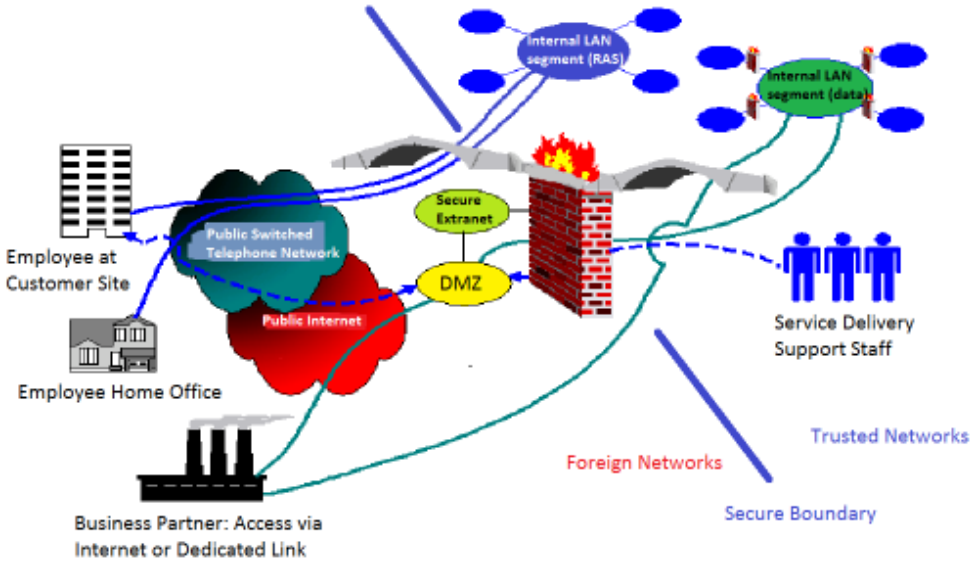

Solution Overview

- Staged approach, where each stage has 2 phases
- Information Collection
  - Collect information about *actually allowed* flows
- Analysis
  - Infer zone colors by comparing actually allowed network flows against policy

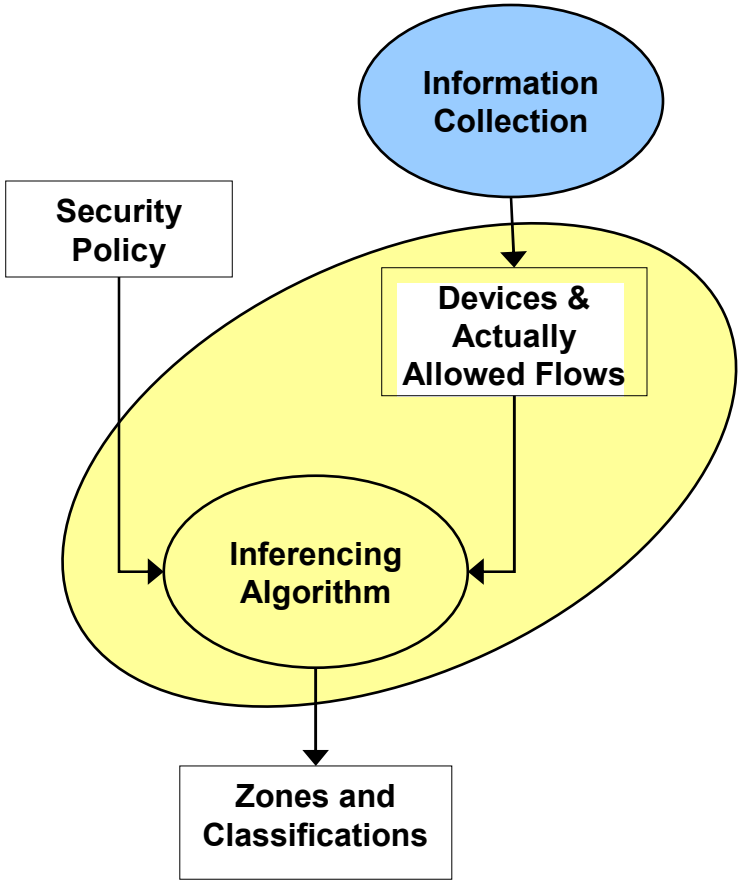

Elimination-Based Inferencing Algorithm

- If color of a zone is Unknown, initially, assign all possible colors (Blue, Green, Red, Yellow)
- Eliminate color if *actually allowed network flows* violates *enterprise policy* for that color
  - Compliance Assumption
- Red zone can send to Unknown
  - Green color is impossible, per policy
  - Blue color is impossible, per policy
- Unknown can send to Blue zone
  - Red color is impossible, per policy
- Only yellow is possible

Enterprise Policy

From \ To	Blue	Green	Yellow	Red
Blue	All	All	All	All
Green	Auth	All	All	All
Yellow	S.Auth	S.Auth	All	All
Red	None	None	All	All

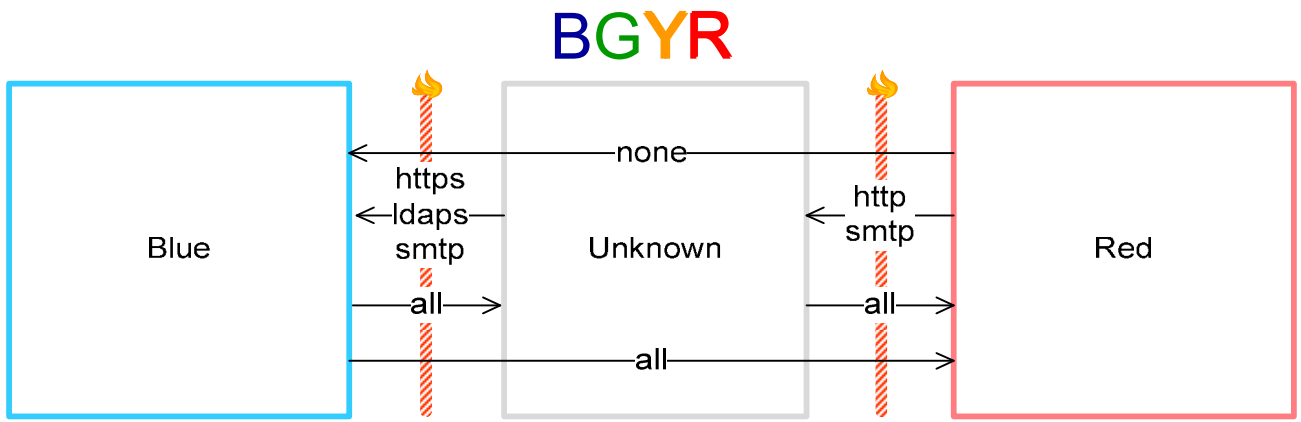

Sample Techniques for Collecting information about Actually Allowed Flows

- Host Config Analysis
  - Routing tables: subnets and groups in the same zone
  - Active connections: app behaviors
- Connectivity Probes
  - Probing with existing app like ping, Telnet, nslookup
- Firewall Config Analysis
  - Parsing firewall configuration files
  - Emulating firewall filtering to find the permitted connections
- Flow Log Analysis
- Network Statistics Analysis
- Packet Analysis

Implemented in BlueGates Tool

Incremental Discovery: Sequence collection methods so that lower interference methods are performed ahead

Case Study: Our approach in action (0 of 5)

- Input
  - Hosts w/ unknown color: X1 ~ X5
  - Hosts w/ known color: B1 (blue) and R1 (red)
  - Policy

To / From	Blue	Green	Yellow	Red
Blue	All	All	All	All
Green	Auth	All	All	All
Yellow	S.Auth	S.Auth	All	All
Red	None	None	All	All

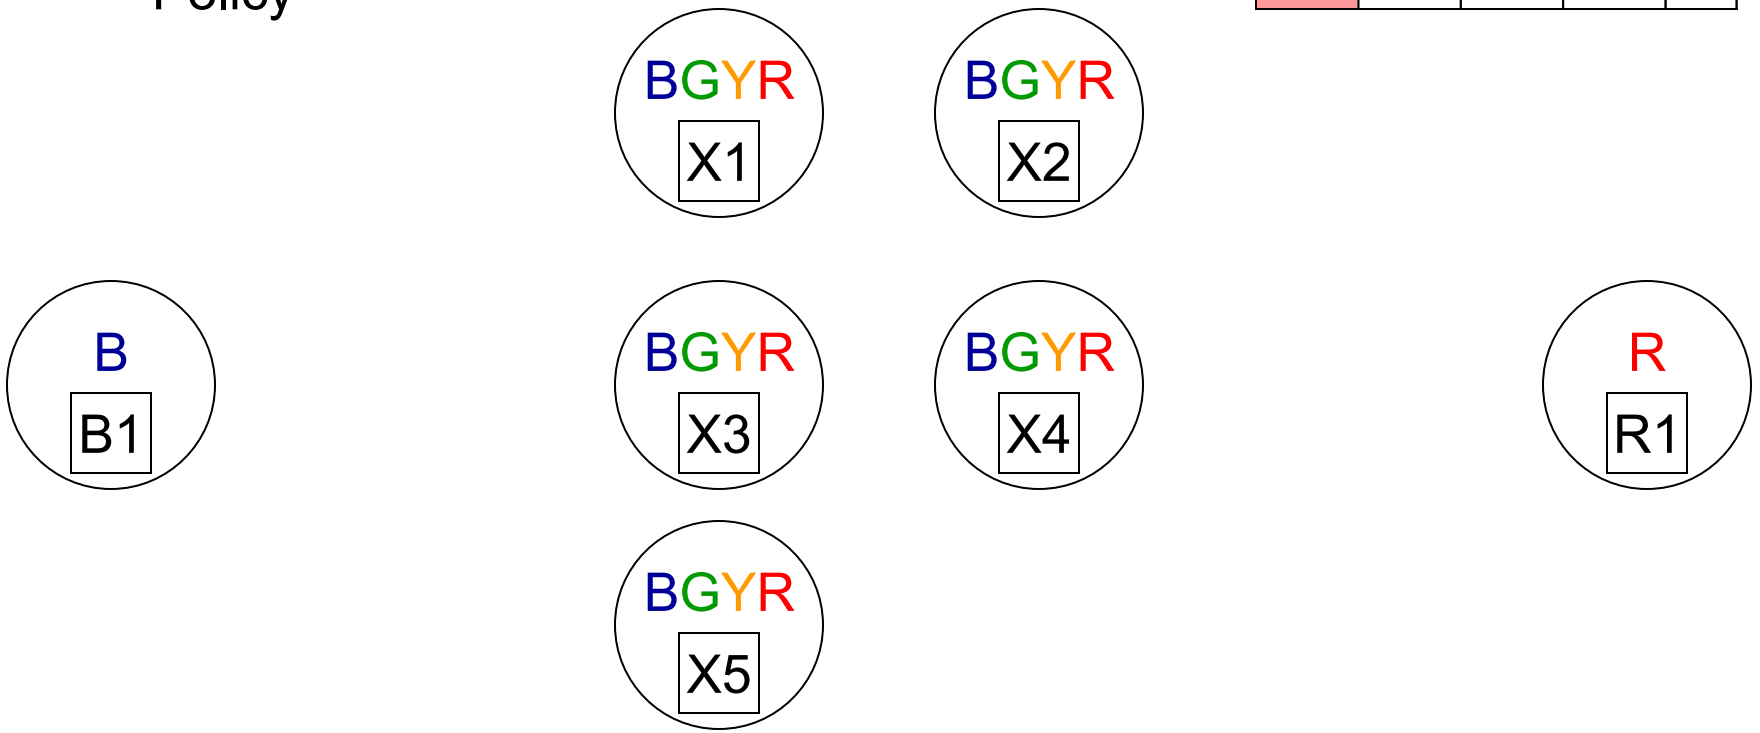

Case Study: Our approach in action (1 of 5)

- Host Config Analysis
  - Routing table analysis: X1 and X2 belongs to the same subnet

To / From	Blue	Green	Yellow	Red
Blue	All	All	All	All
Green	Auth	All	All	All
Yellow	S.Auth	S.Auth	All	All
Red	None	None	All	All

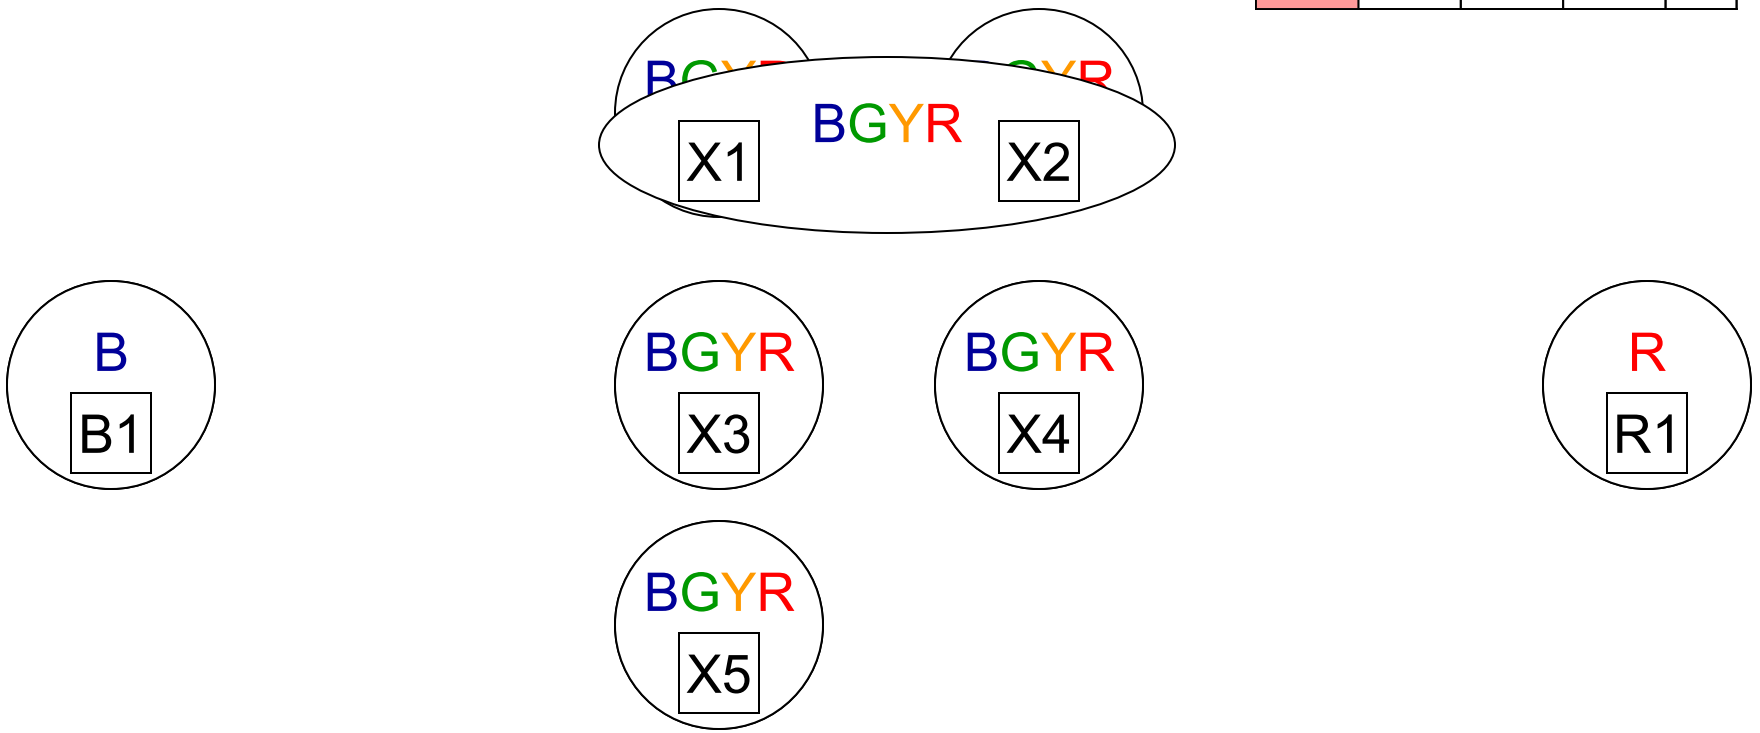

Case Study: Our approach in action (2 of 5)

- Host Config Analysis
  - Routing table analysis: X1 and X2 belongs to the same subnet
  - Active connections analysis: HTTP from R1 to X2 and X4

To / From	Blue	Green	Yellow	Red
Blue	All	All	All	All
Green	Auth	All	All	All
Yellow	S.Auth	S.Auth	All	All
Red	None	None	All	All

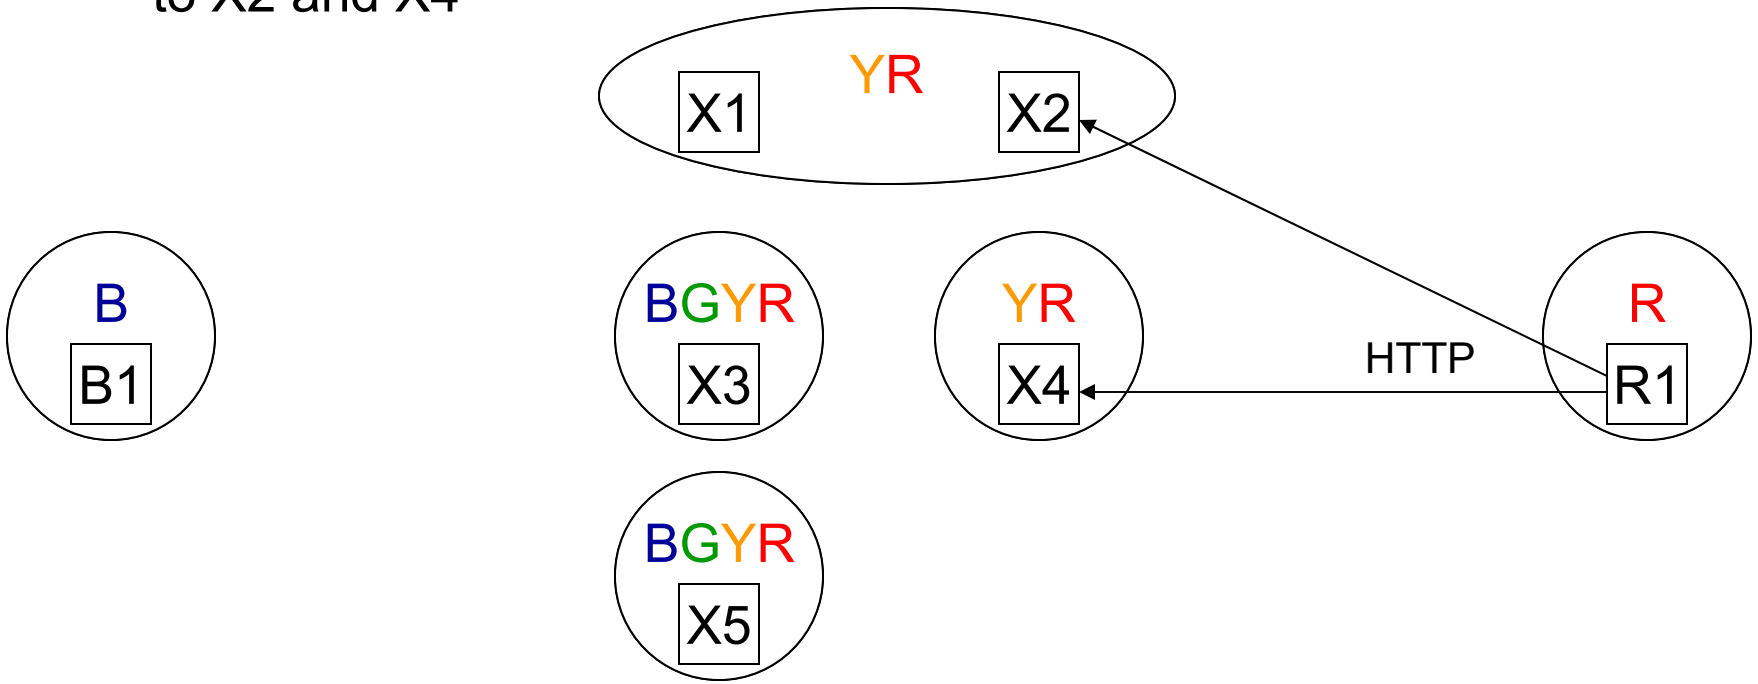

Case Study: Our approach in action (3 of 5)

- Host Config Analysis
  - Routing table analysis: X1 and X2 belongs to the same subnet
  - Active connections analysis: HTTP from R1 to X2 and X4
  - Active connections analysis: HTTPS from X1 to B1

To / From	Blue	Green	Yellow	Red
Blue	All	All	All	All
Green	Auth	All	All	All
Yellow	S.Auth	S.Auth	All	All
Red	None	None	All	All

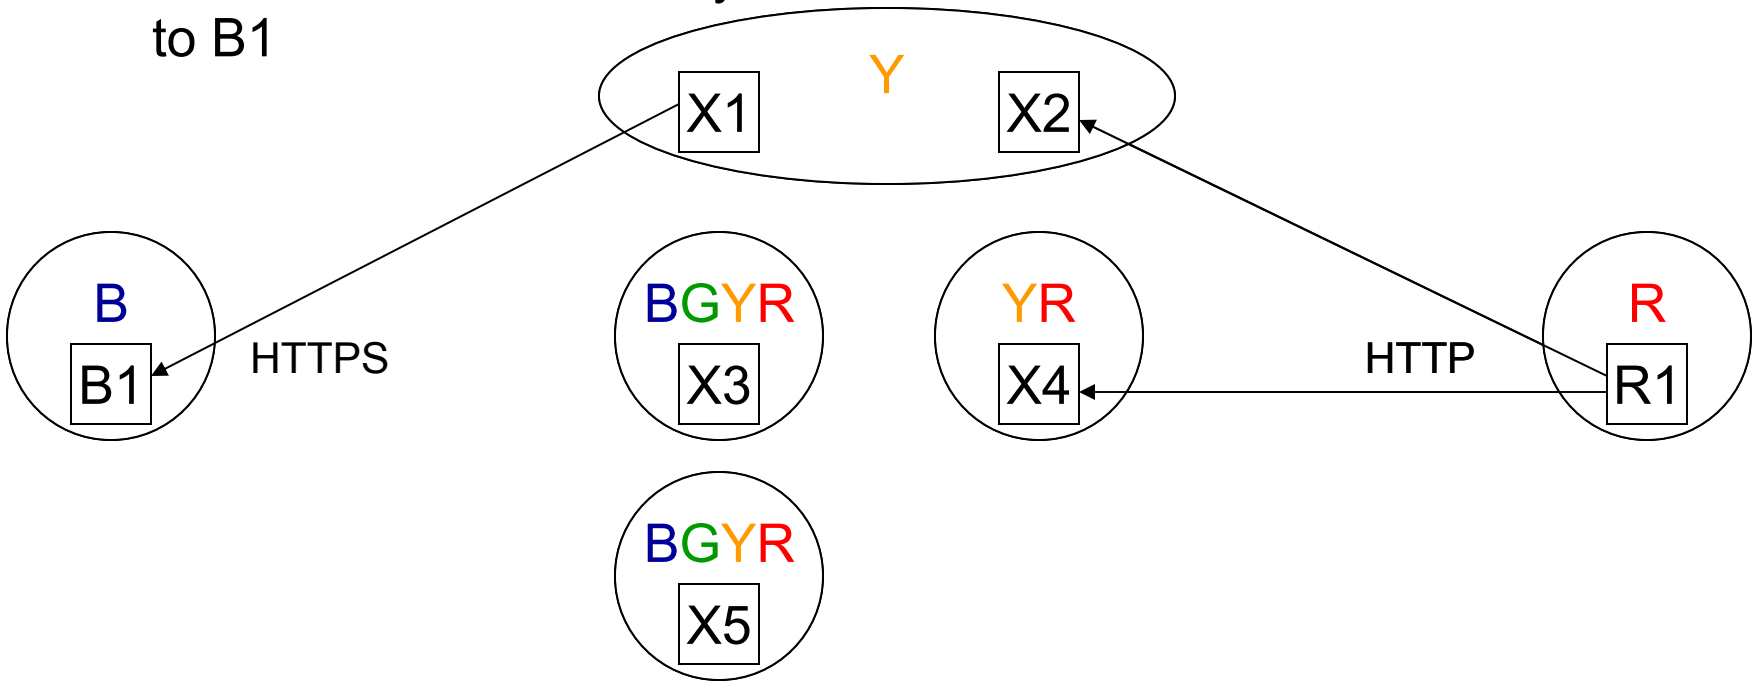

Case Study: Our approach in action (4 of 5)

- Connectivity Probing
  - HTTPS traffic allowed from X3 and X5 to B1
  - TFTP traffic allowed from X4 to X3

From \ To	Blue	Green	Yellow	Red
Blue	All	All	All	All
Green	Auth	All	All	All
Yellow	S.Auth	S.Auth	All	All
Red	None	None	All	All

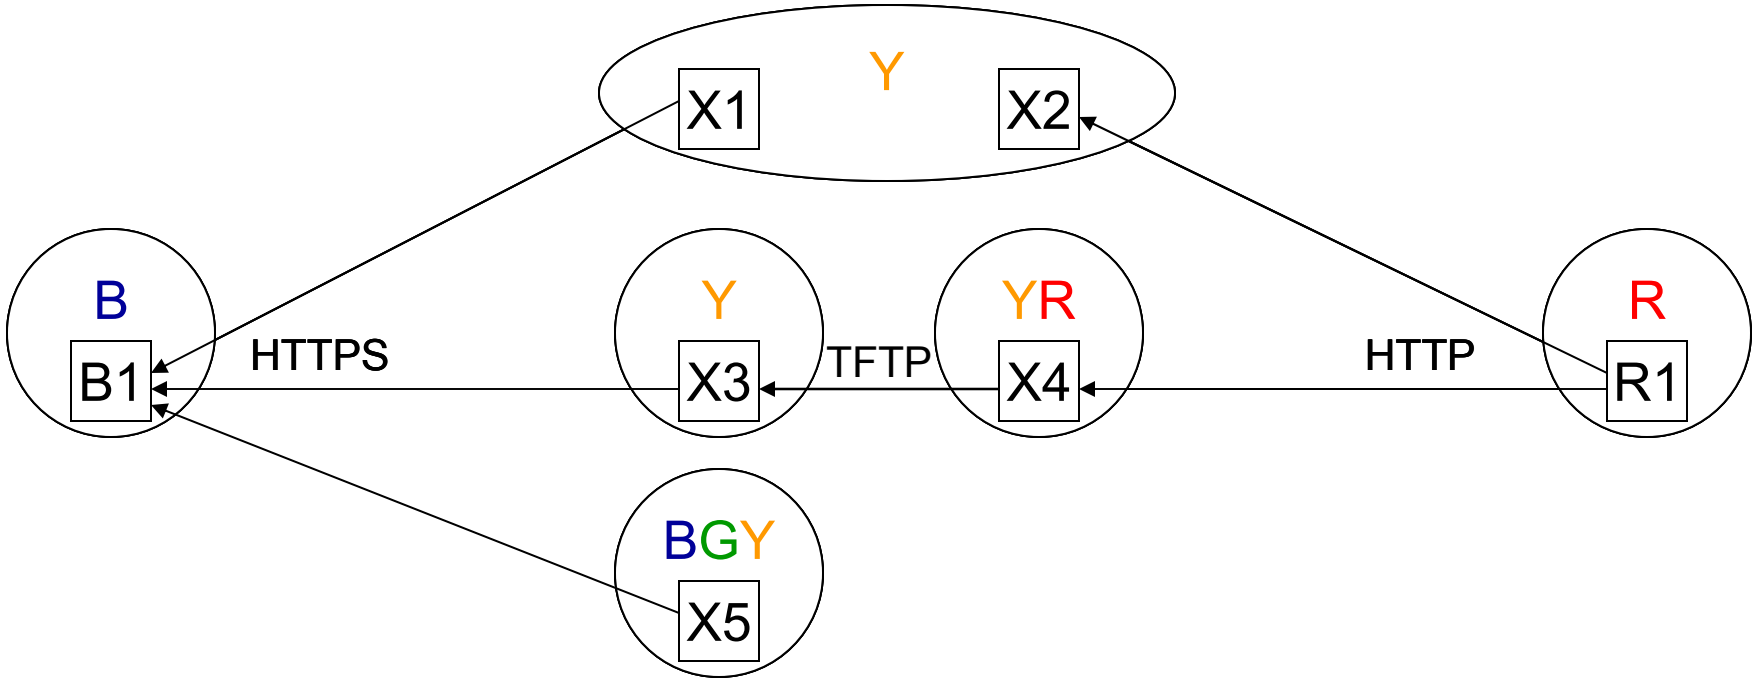

Case Study: Our approach in action (5 of 5)

- Firewall Config Analysis
  - No firewall between X3 and X4
  - HTTP traffic between R1 and new host X6
  - X5 and X6 in same subnet

From \ To	Blue	Green	Yellow	Red
Blue	All	All	All	All
Green	Auth	All	All	All
Yellow	S.Auth	S.Auth	All	All
Red	None	None	All	All

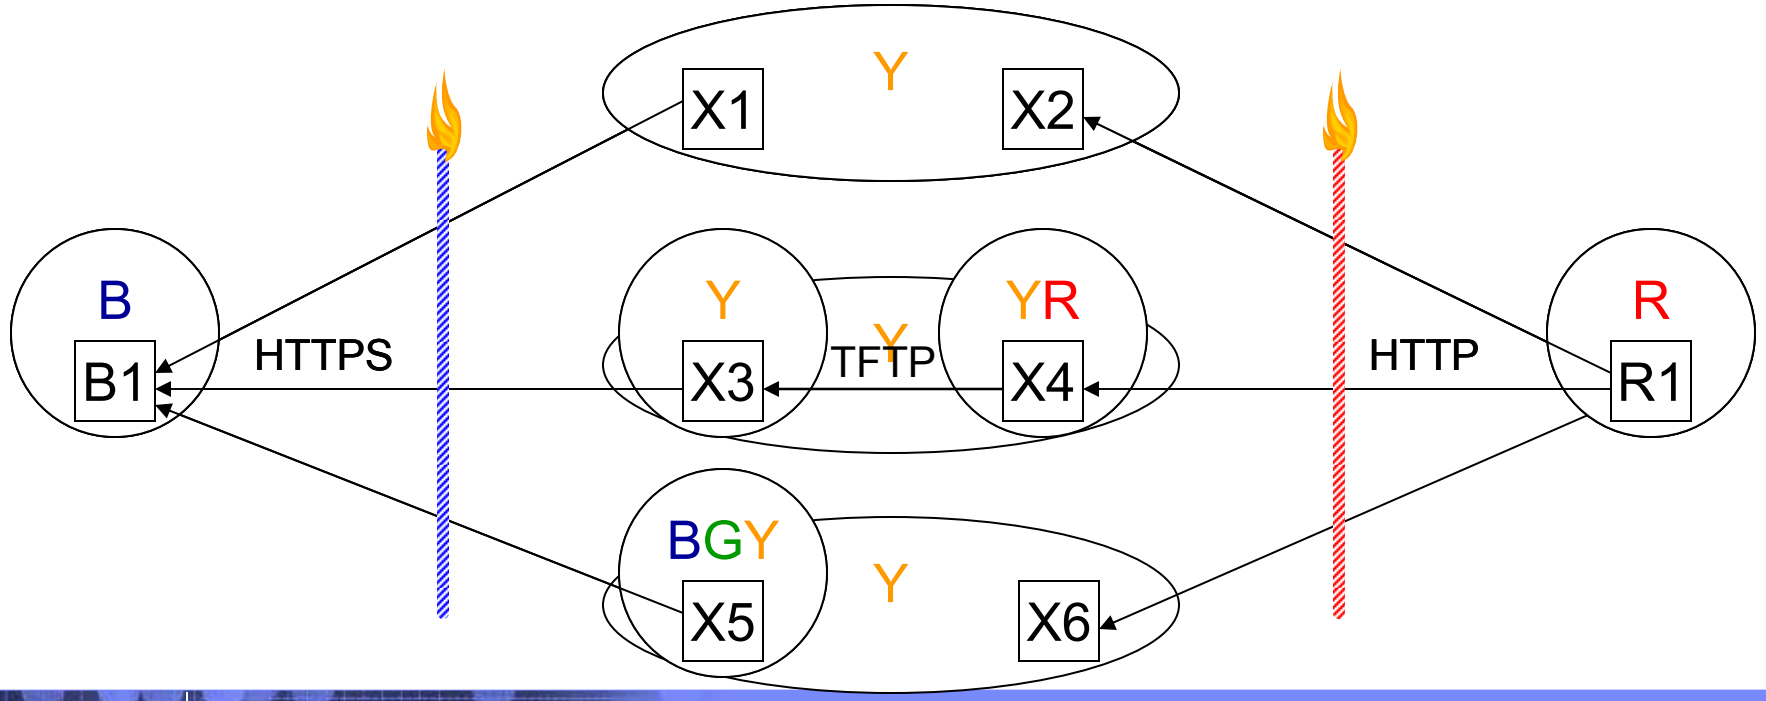

Conclusion

- Systematic and semi-automated approach for discovering security zone classifications of devices
  - Staged approach to information collection
  - Elimination-based inferencing
  - Generalization as a Constraint Satisfaction Problem
- Future (on-going) work
  - Loosening the compliance assumption
  - Evaluating the approach in large-scale infrastructures