IBM Research

Enterprise IT Infrastructure Transformation: **Towards Automated Identification of Security Zone Classifications**

Hari Ramasamy

Services Research

IBM T.J. Watson Research Center, Hawthorne, NY

Joint Work with

Cheng-Lin Tsao, Birgit Pfitzmann, Nikolai Joukov, and Jim W. Murray

59th Meeting of the 10.4 WG on Dependable Computing and Fault Tolerance

Snowmass Village, Colorado, US

Jan 12-16th, 2010

Background: IT Service Management

Intersection of people, process, information, and technology

Objective: Effective and efficient delivery of IT services in support of business goals

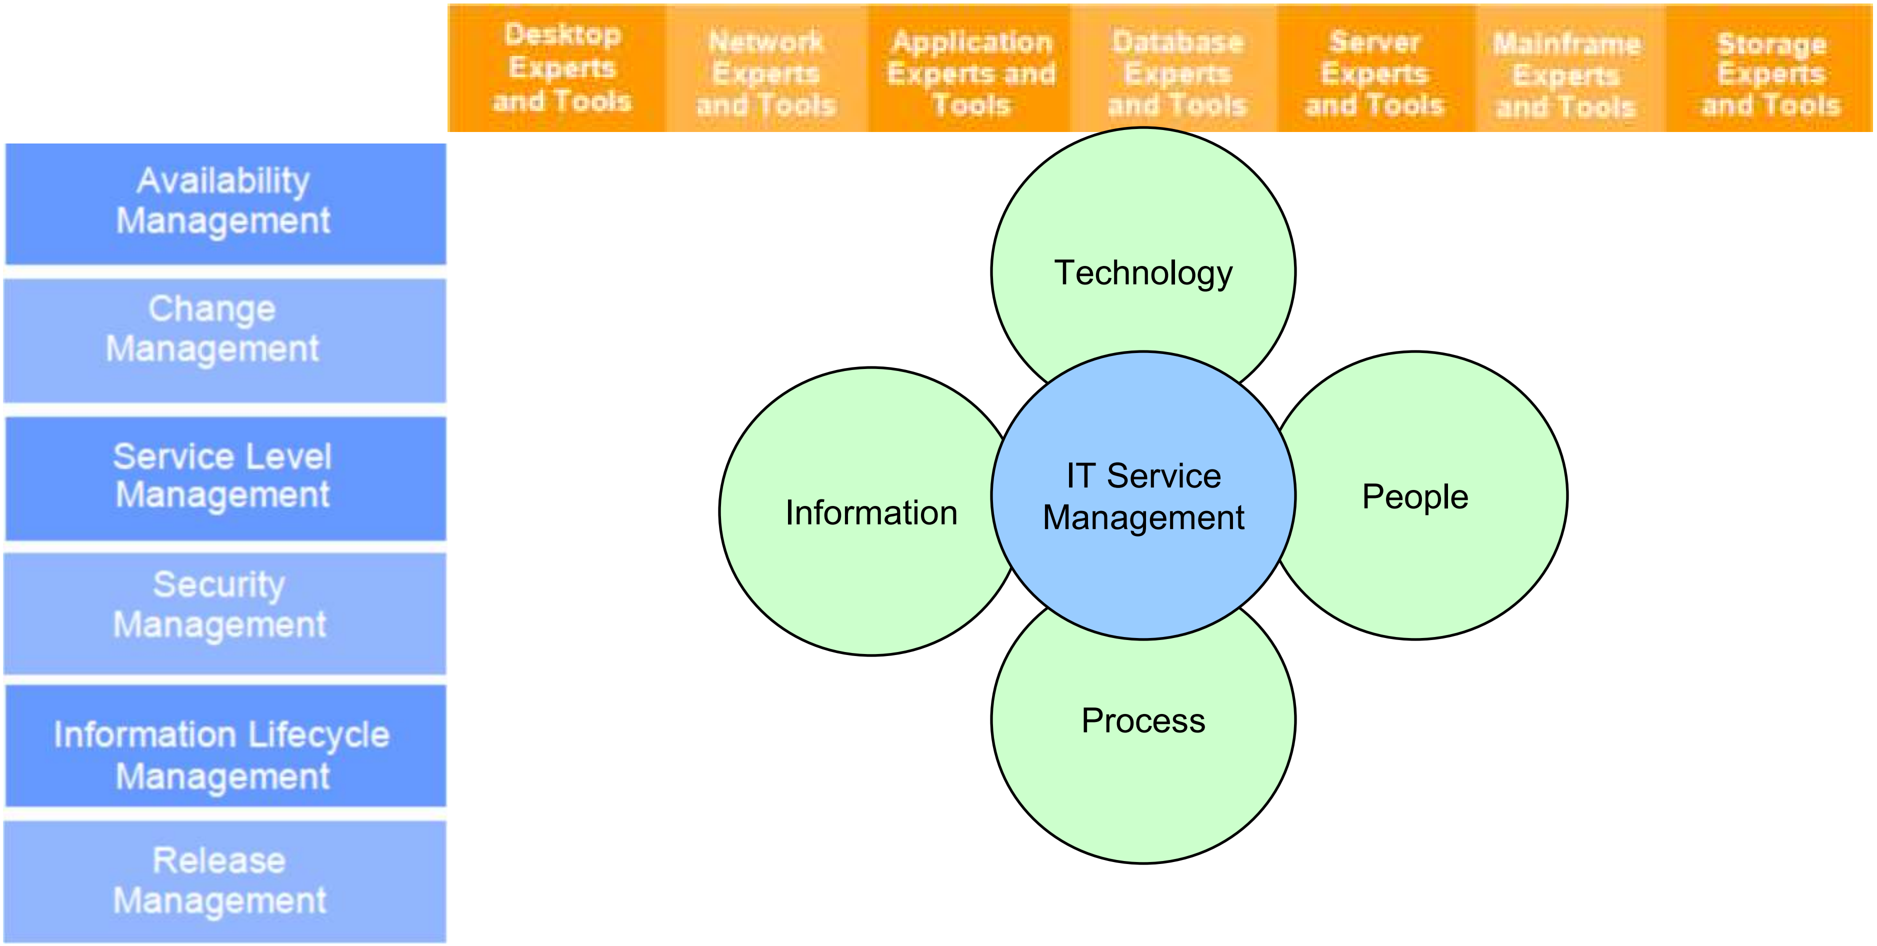

Desktop Experts and Tools

Network Experts and Tools

Application Experts and Tools

Database Experts and Tools

Server Experts and Tools

Mainframe Experts and Tools

Storage Experts and Tools

Availability Management

Change Management

Service Level Management

Security Management

Information Lifecycle Management

Release Management

Technology

Information

IT Service Management

People

Process

Background: Enterprise IT Infrastructure

- **Consists of hardware, software, and services**
- **Connects users and systems to each other and the outside world**
- **Significant expenditure for enterprises**
  - Large financial firms spend 3.5% of their revenue on IT infrastructure
  - Network infrastructure alone accounts for 12-18% of Fortune 500 companies' expenses
- **Large environmental footprint**
  - Servers alone consume 1.2% of the electricity produce in US

Background: Enterprise IT Infrastructure Transformation

- **Drivers of Transformation**
  - New business needs
  - Pressures to reduce cost
- **Fraught with risks and challenges**
  - Can’t disrupt running production workloads
  - System structure not fully known
  - System heterogeneity
  - Many “moving parts”

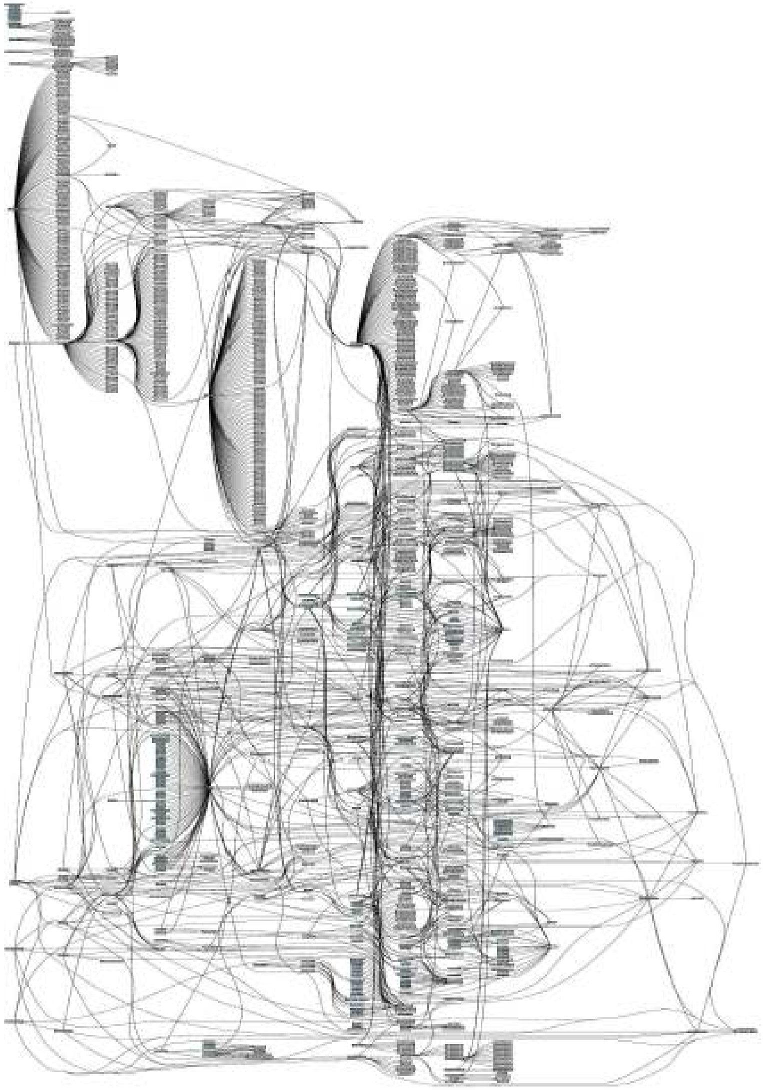

Reasons for Enterprise IT Transformation

- **Technical Reasons**
  - Reduction in Cost/Complexity through Network/Server Consolidation
  - Server and Desktop Virtualization, Migration to Cloud
  - Virtualized Network Services
  - Unified Messaging
  - Enterprise Mobility
- **Business Reasons**
  - Acquisitions and Mergers
  - Infrastructure Outsourcing
  - Regulation and Compliance

Sample Enterprise IT Transformation Scenario

- **Enterprise Firewall Infrastructure Consolidation**
  - Over time, the firewall infrastructure of an enterprise grows organically, eventually resulting in management, maintenance, and cost issues
  - Objective: Reduce the number of firewalls, while enhancing security

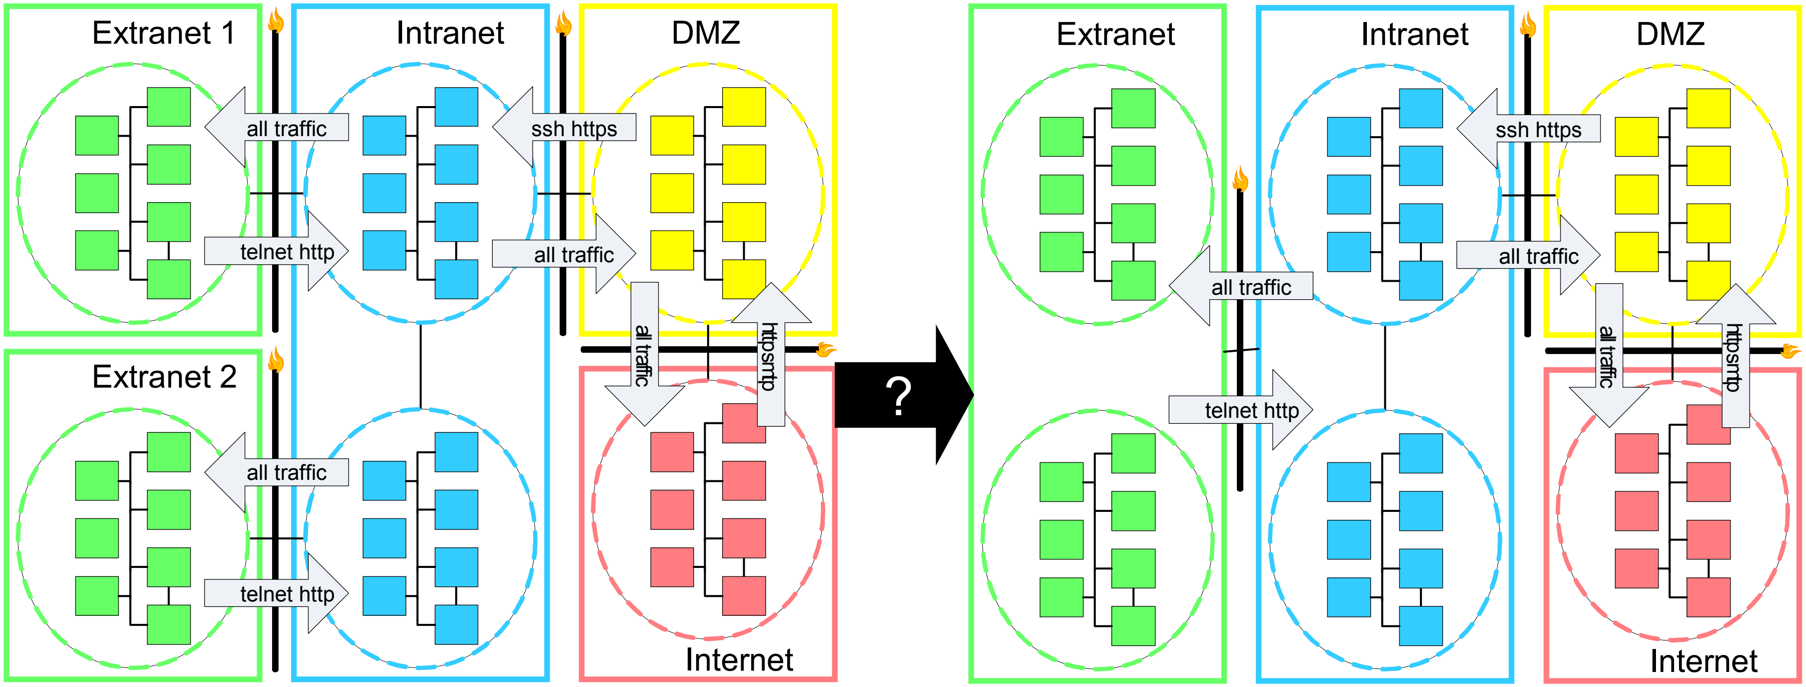

Gain vs. Pain of Transformation

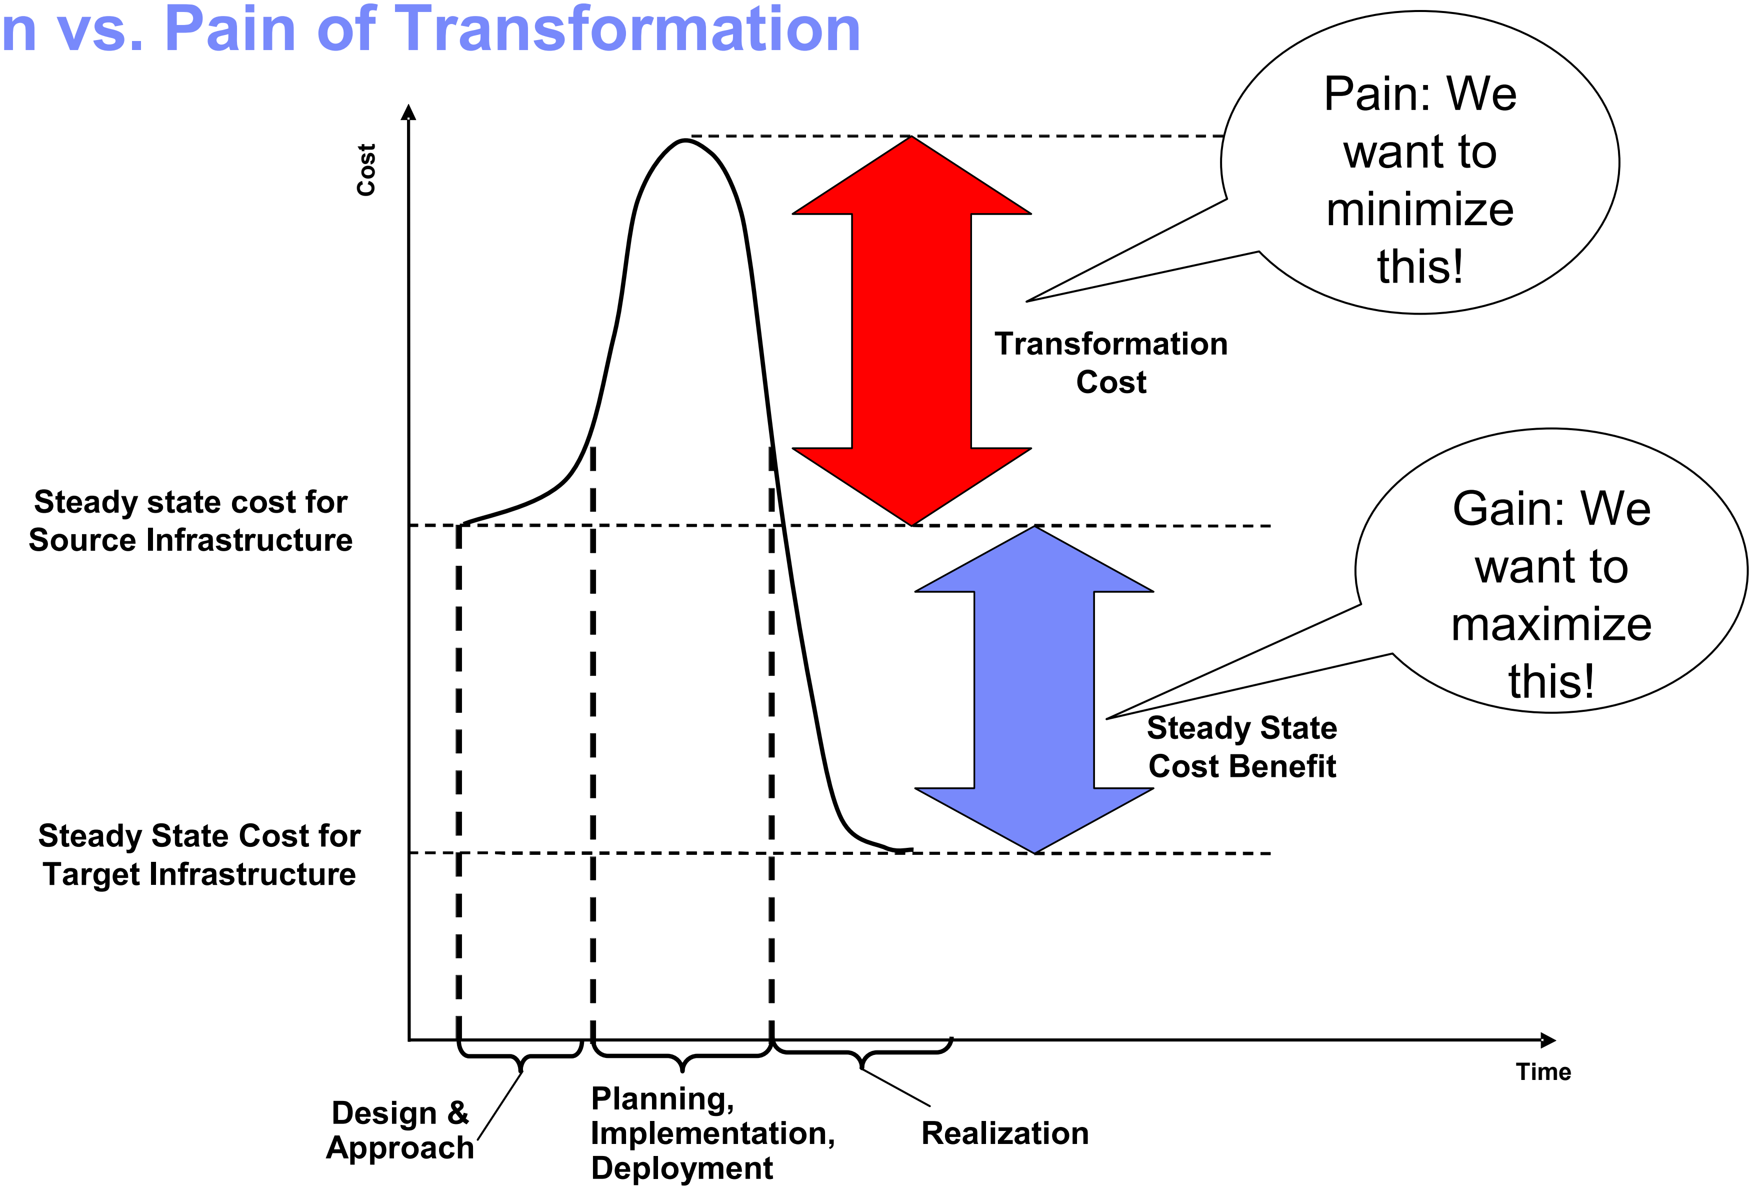
Cost
Pain: We want to minimize this!
Transformation Cost
Steady state cost for Source Infrastructure
Gain: We want to maximize this!
Steady State Cost Benefit
Steady State Cost for Target Infrastructure
Time
Design & Approach
Planning, Implementation, Deployment
Realization

Abstract Architecture for Enterprise IT Transformation

Specific problem encountered:

Security Zone Identification

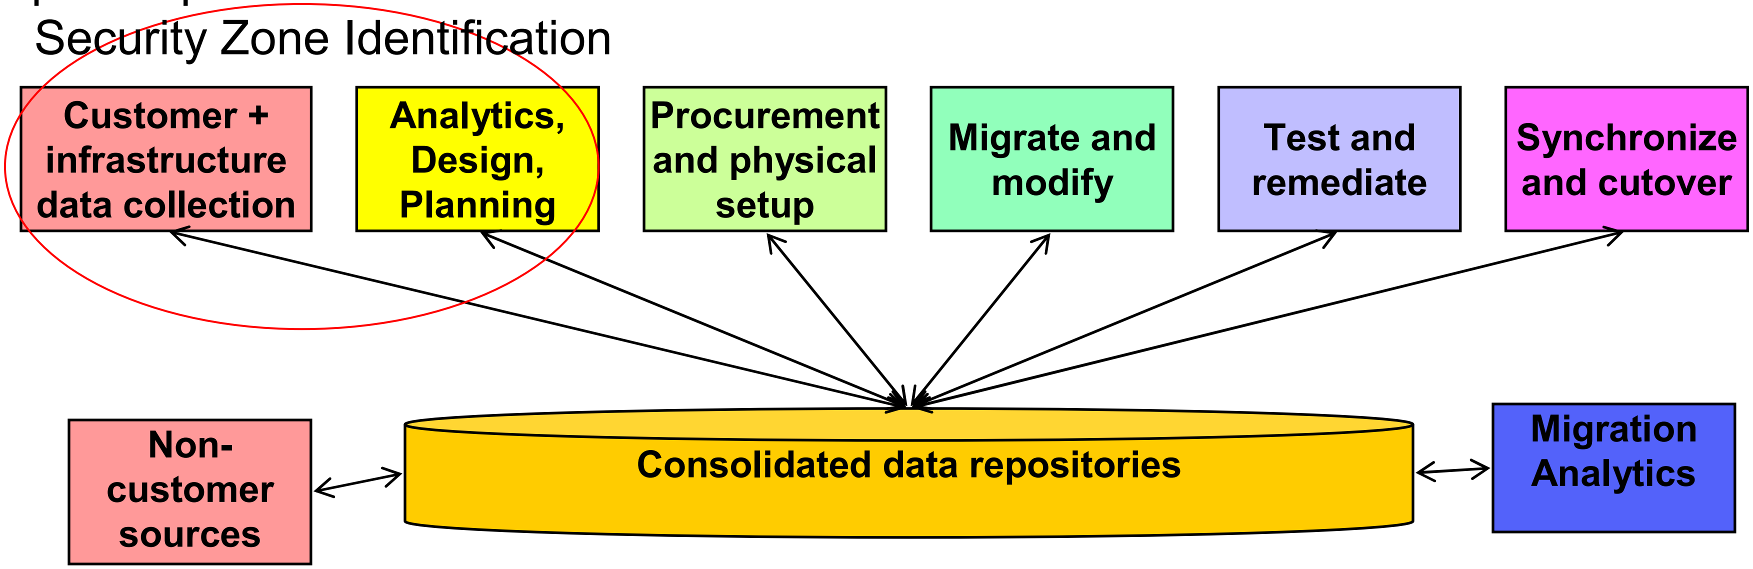

Customer + infrastructure data collection

Analytics, Design, Planning

Procurement and physical setup

Migrate and modify

Test and remediate

Synchronize and cutover

Non-customer sources

Consolidated data repositories

Migration Analytics

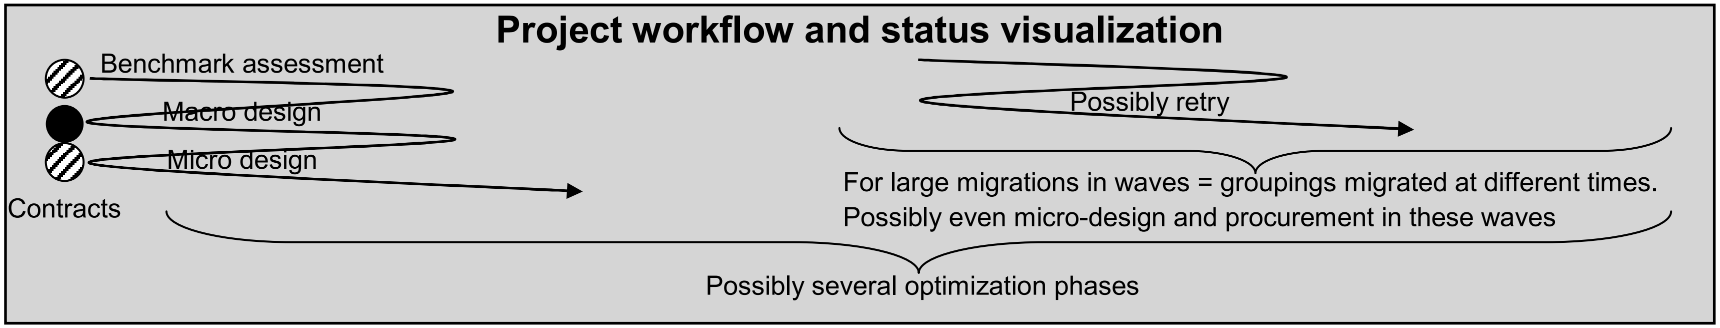

Project workflow and status visualization

Benchmark assessment

Macro design

Micro design

Contracts

Possibly retry

For large migrations in waves = groupings migrated at different times. Possibly even micro-design and procurement in these waves

Possibly several optimization phases

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*

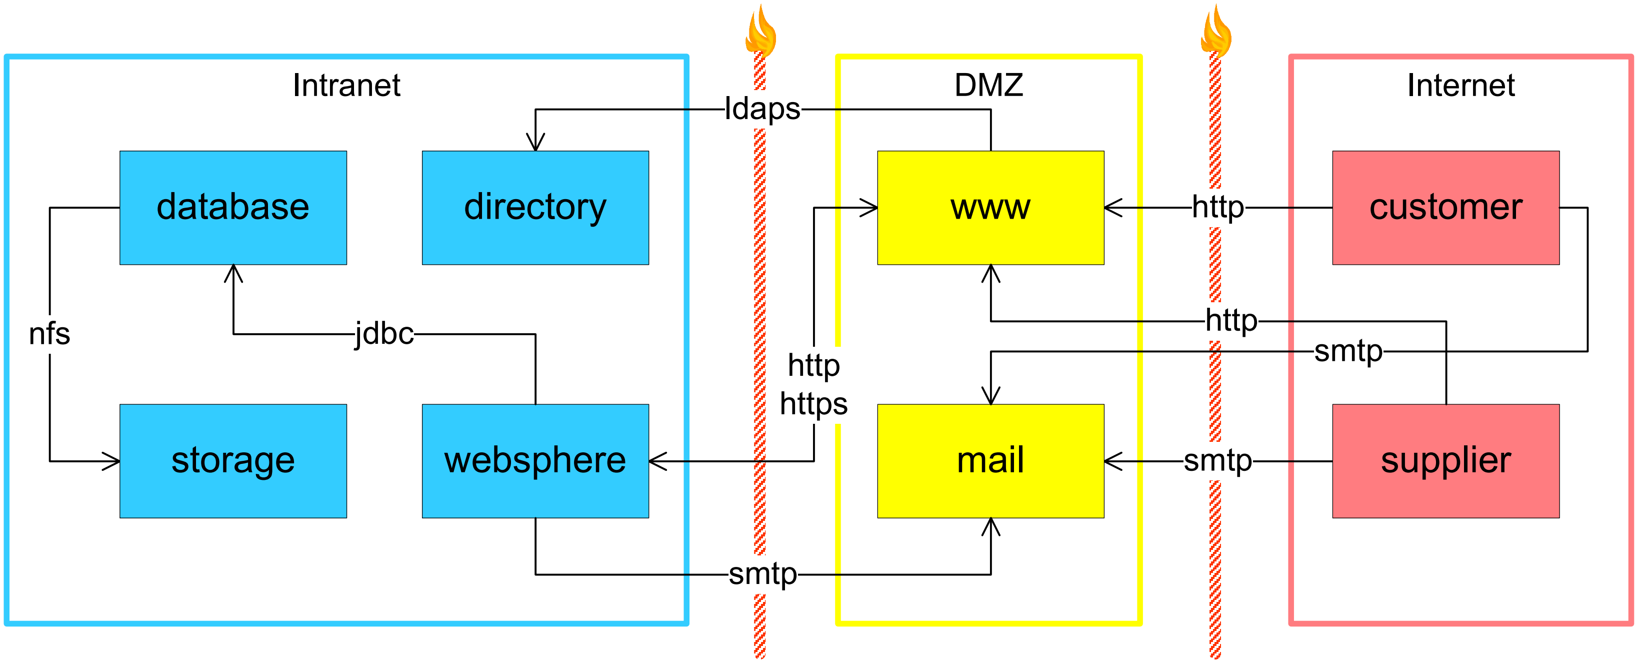

Network Model and Policy Model

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

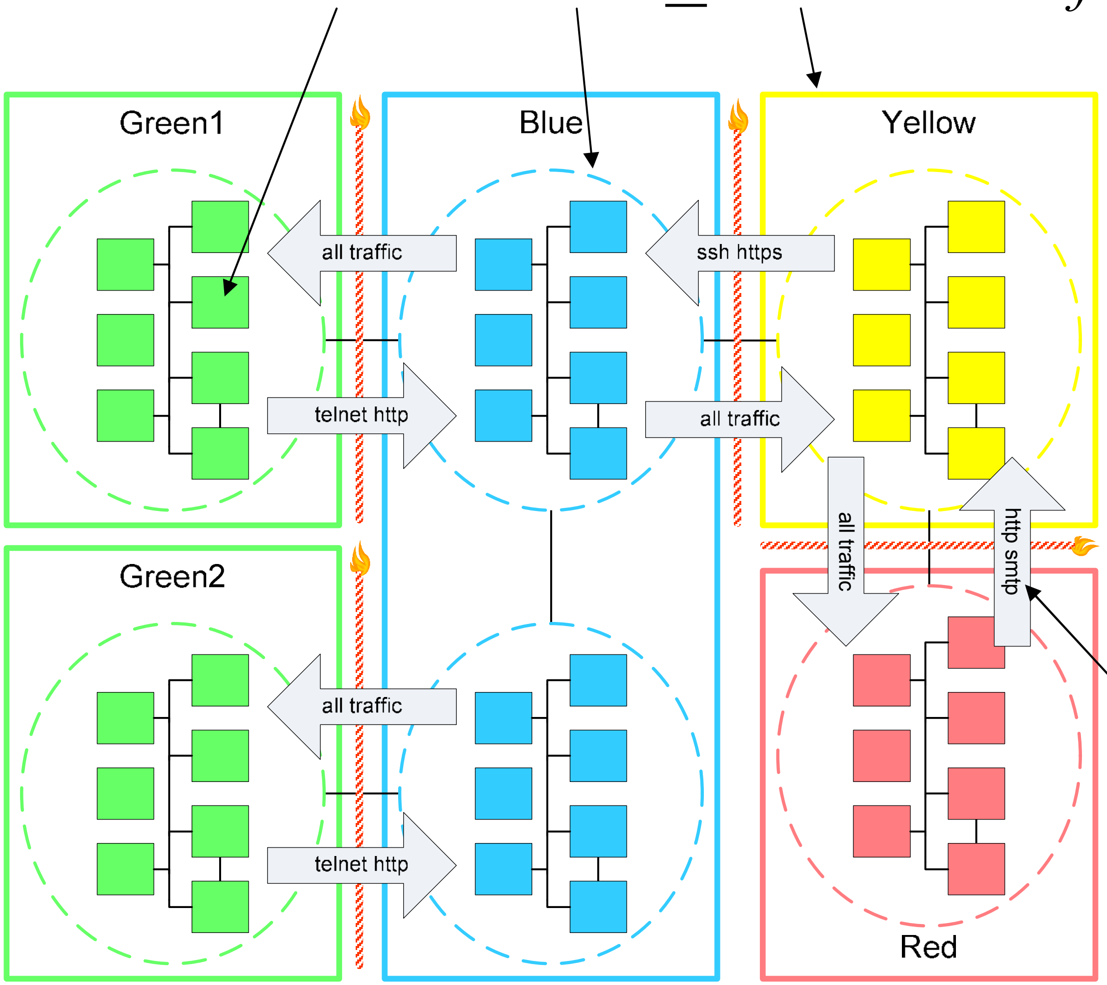

To / From	Blue	Green	Yellow	Red
Blue	All	All	All	All
Green	Auth	All	All	All
Yellow	S.Auth	S.Auth	All	All
Red	None	None	All	All

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

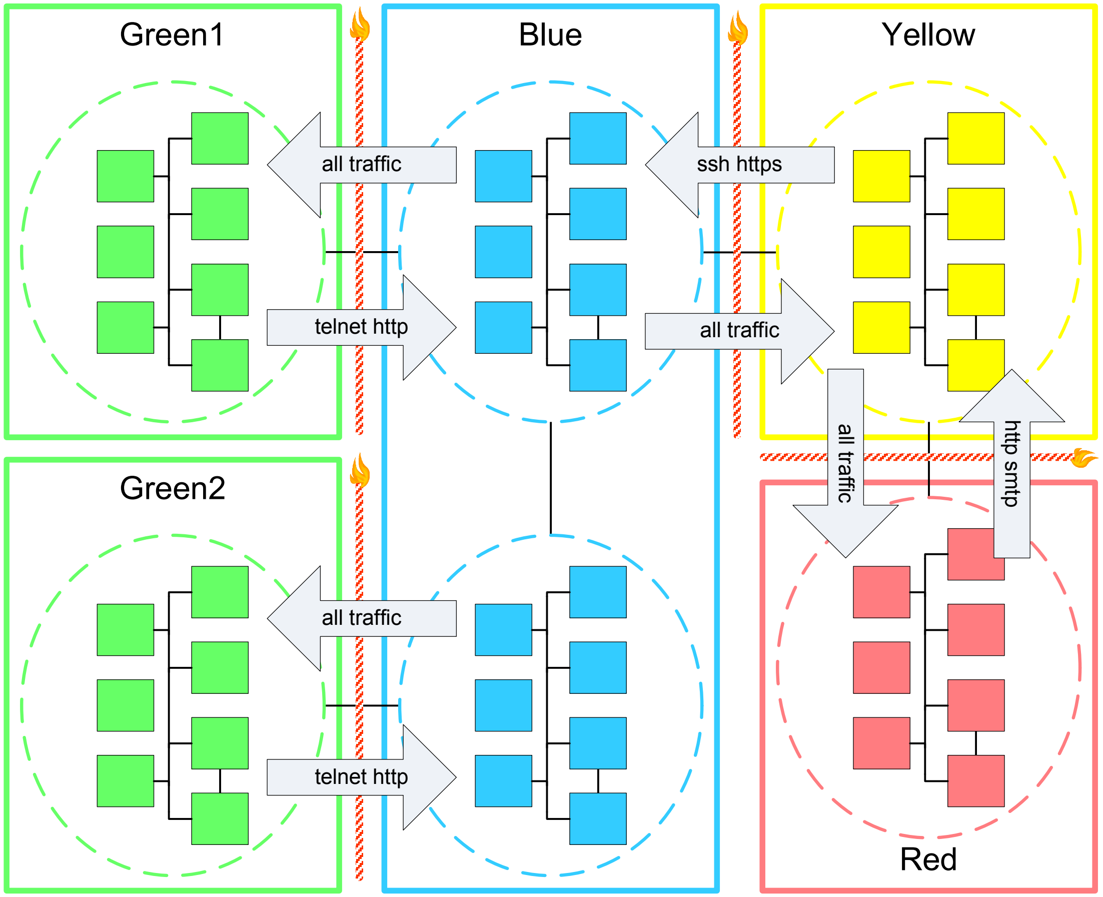

Motivation for Security Zone Discovery

- 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

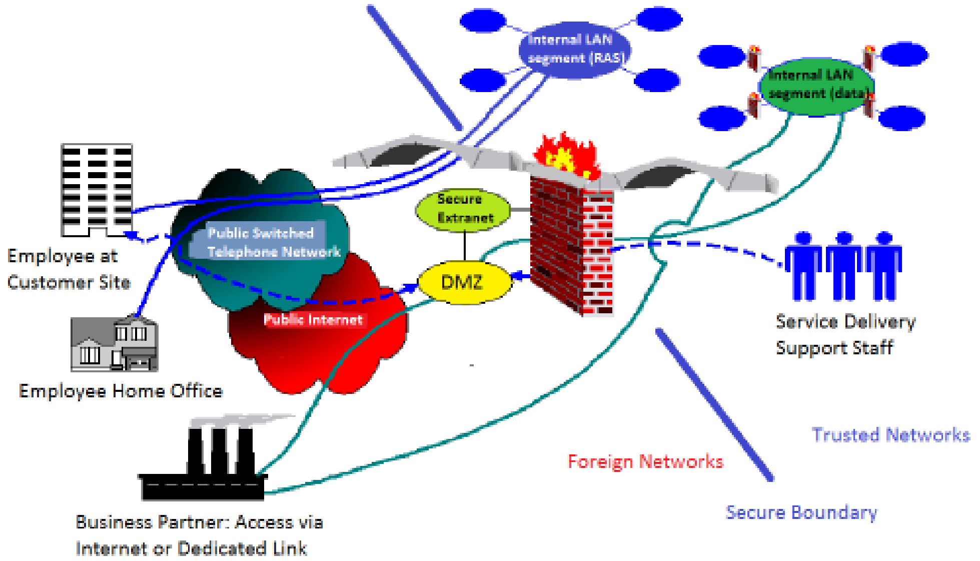
Internal LAN segment (RAS)
Internal LAN segment (data)
Secure Extranet
DMZ
Public Switched Telephone Network
Public Internet
Employee at Customer Site
Employee Home Office
Service Delivery Support Staff
Trusted Networks
Foreign Networks
Secure Boundary
Business Partner: Access via Internet or Dedicated Link

Flow Control Specification (Policy) vs. Implementation (Configuration)

- Flow control policy (or reachability policy) of an enterprise defines which packets can get through to which devices in the enterprise
- Access Control Lists (ACLs) are the actual implementation of enterprise flow control policy
  - ACLs are placed in networking devices (routers, firewalls) and hosts
- An ACL is a sequential collection of rules. Each rule is a *permit condition* or a *deny condition*
- A packet passing through a device interface is matched against each rule successively
  - Testing stops with the first match, so order of rules is important
  - If no match found, an implicit “deny any” rule is assumed, and packet rejected

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

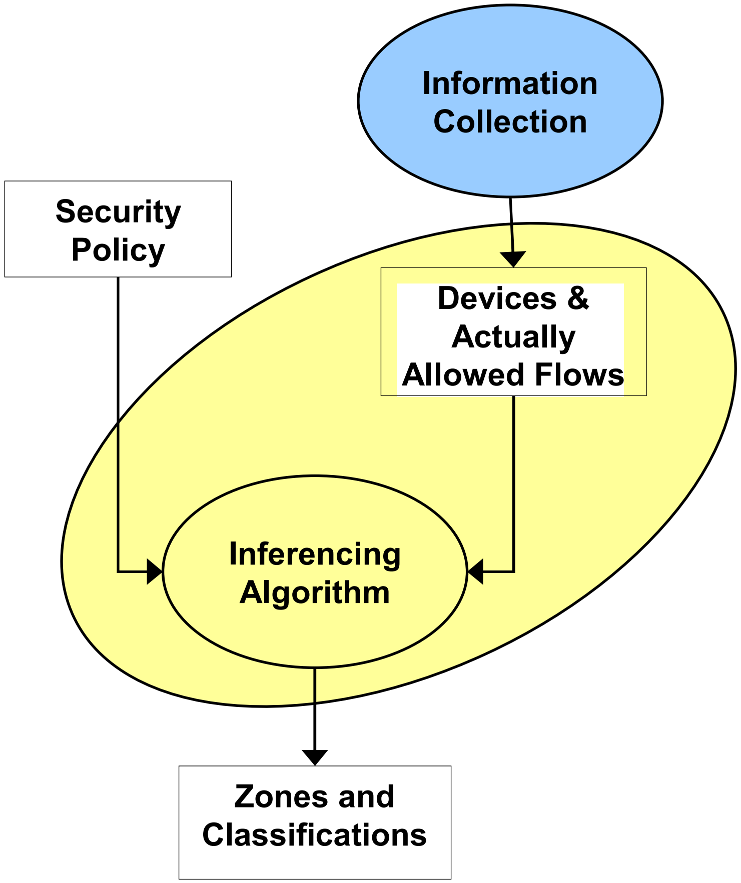

Elimination-Based Inferencing Algorithm

- 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

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

BGYR

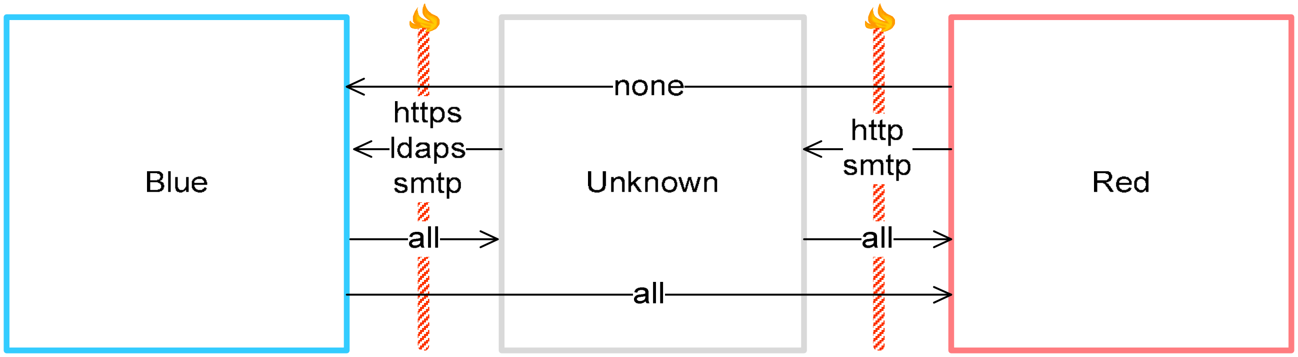

Sample Techniques for Collecting information about Actually Allowed Flows

- 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

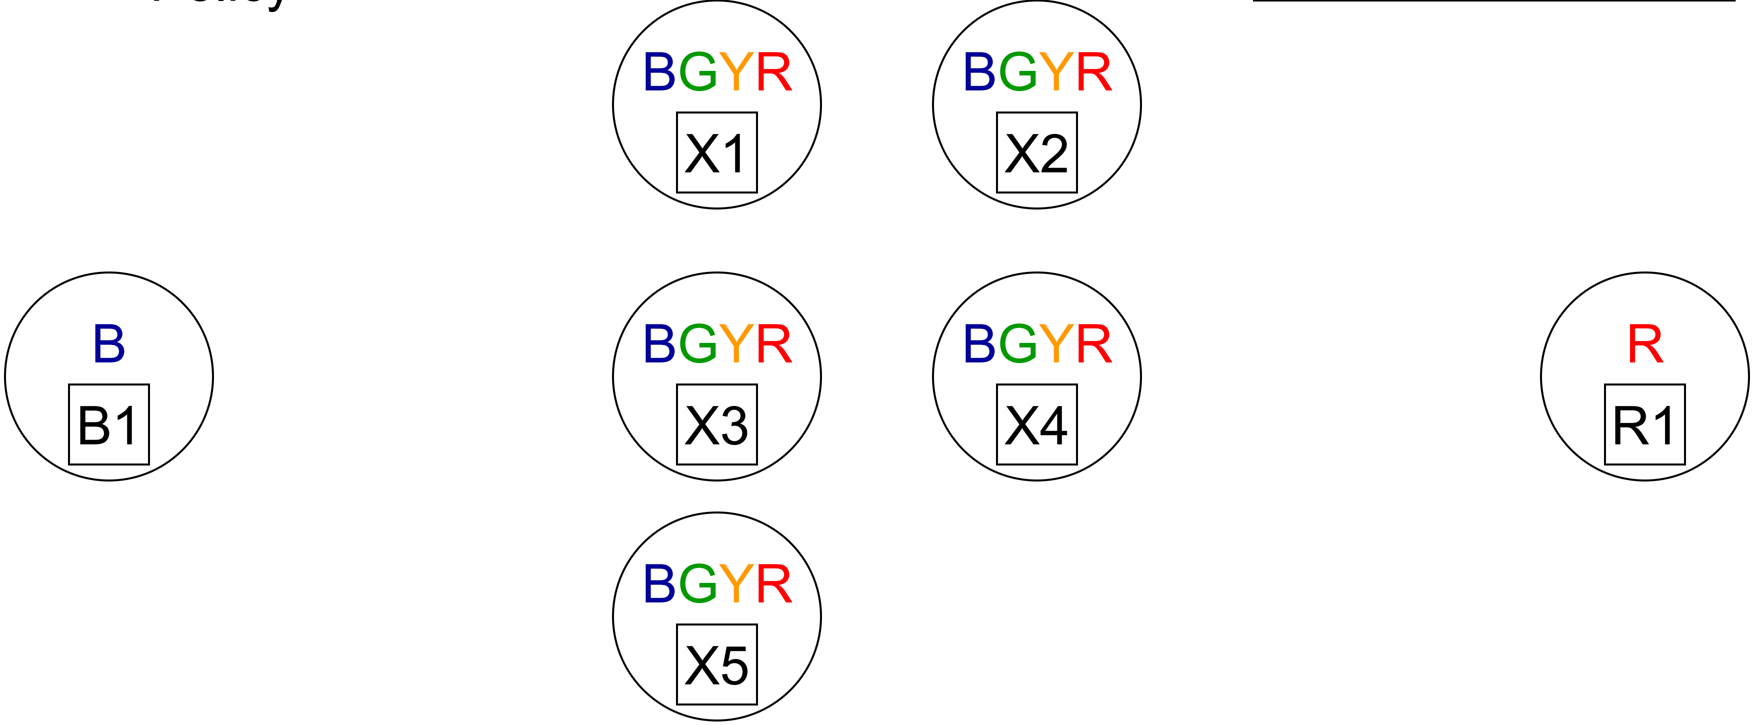

Case Study: Our approach in action (1 of 5)

- 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

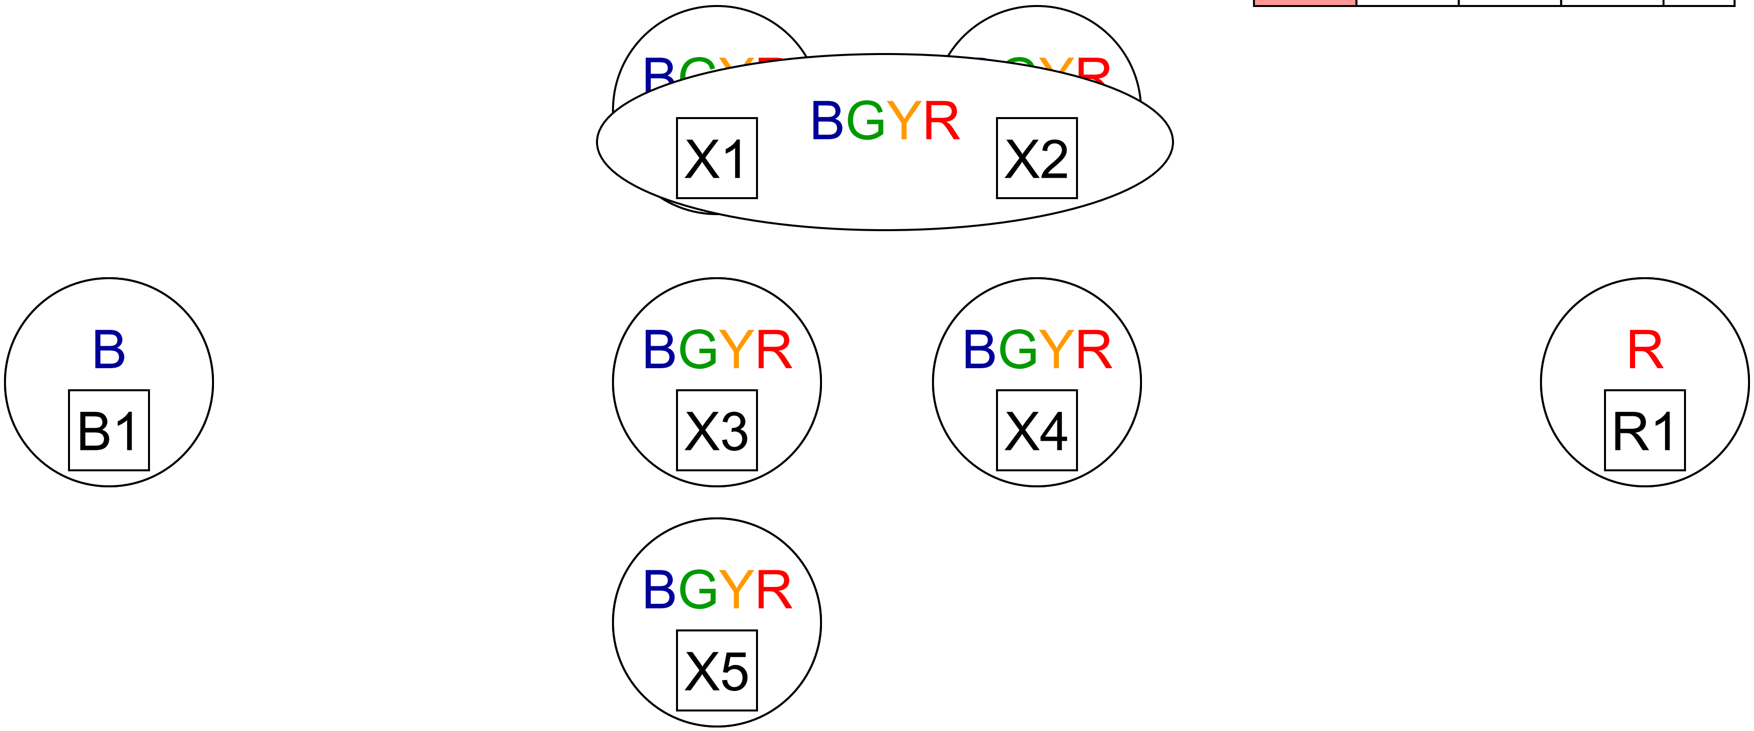

Case Study: Our approach in action (2 of 5)

- 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

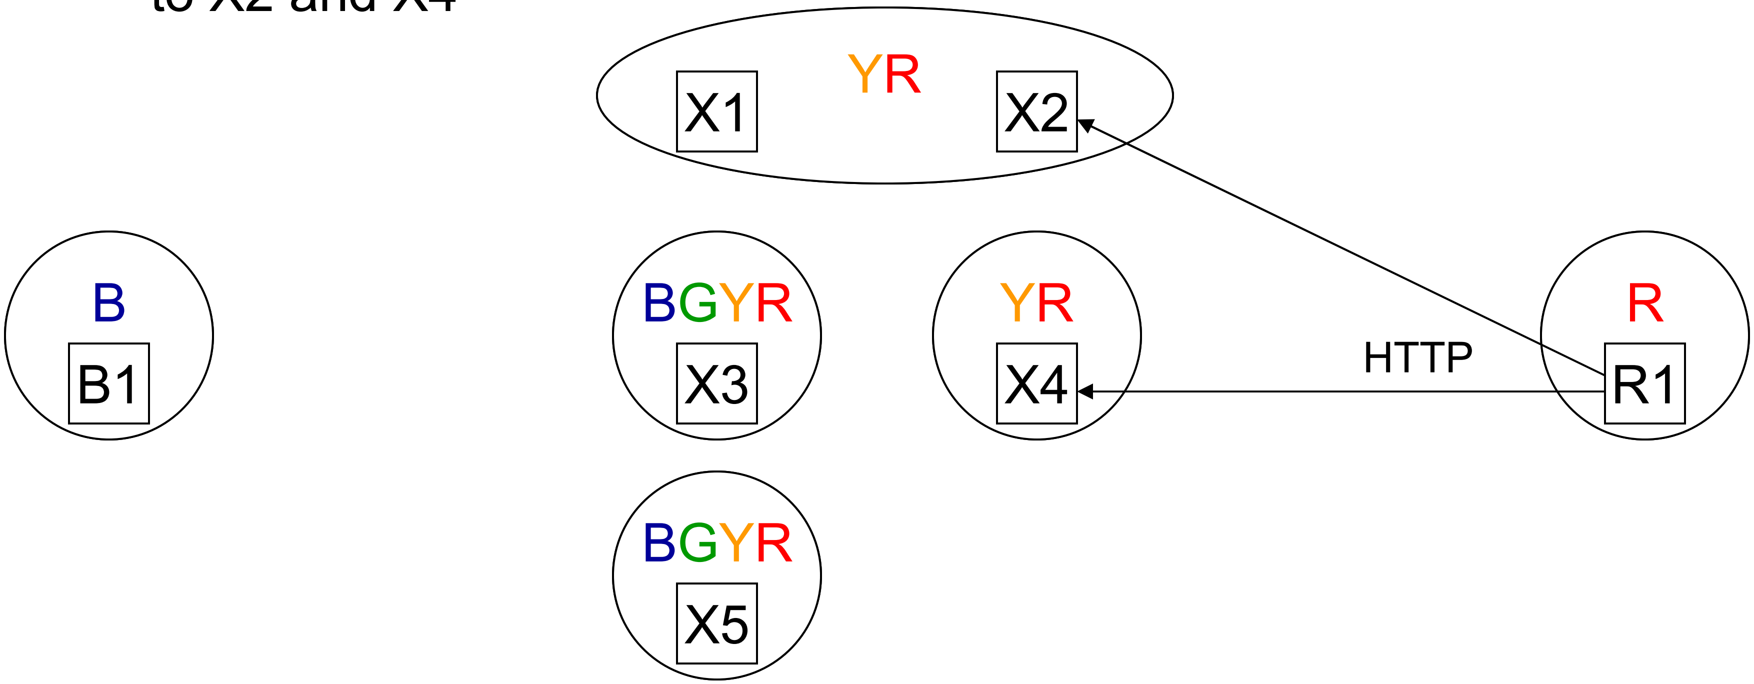

Case Study: Our approach in action (3 of 5)

- 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

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

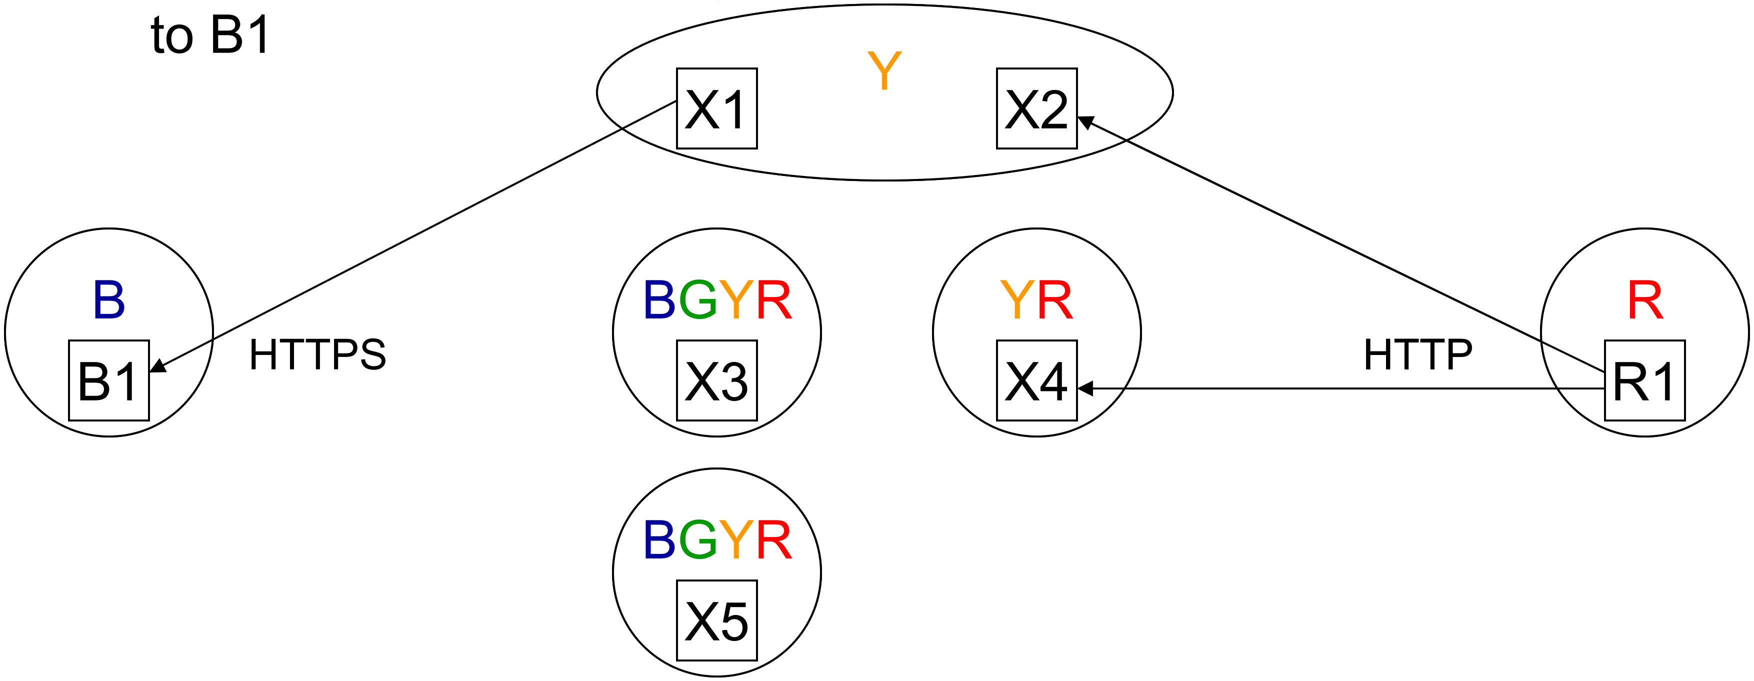

Case Study: Our approach in action (4 of 5)

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

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

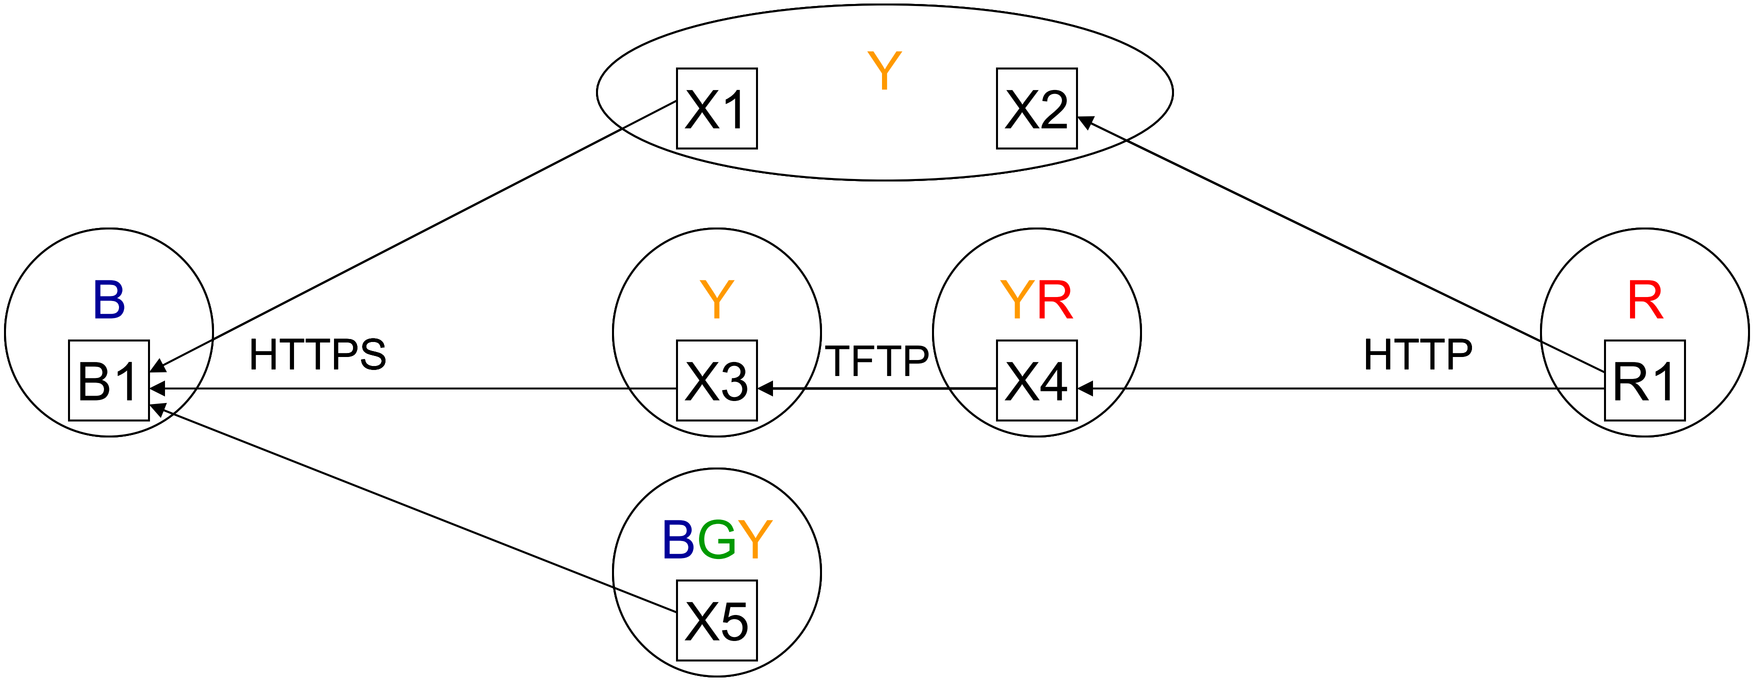

Case Study: Our approach in action (5 of 5)

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

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

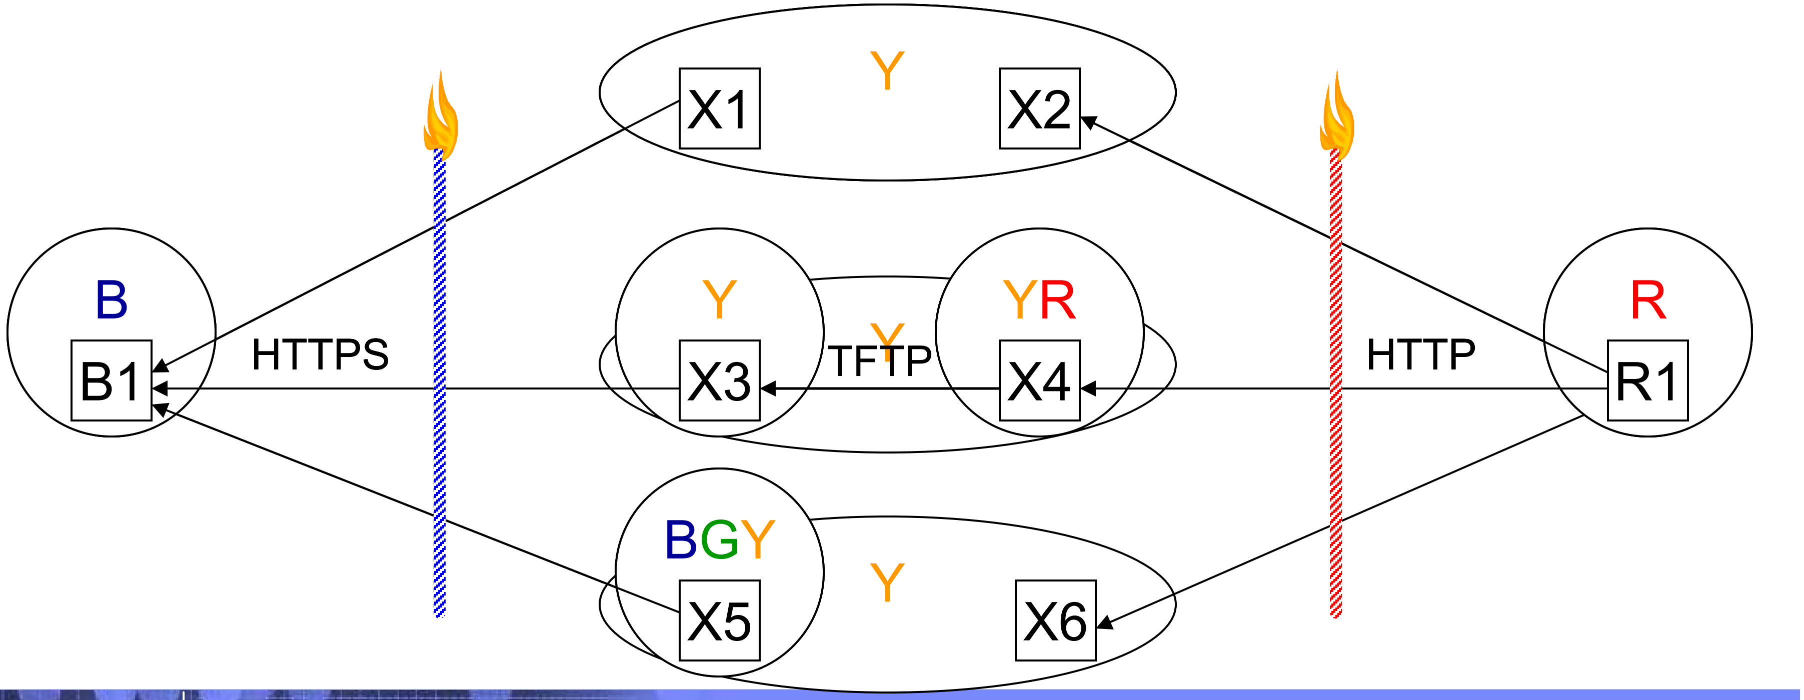

Summary and Conclusion

- Enterprise IT Infrastructure Transformation
  - Challenging endeavour due to ground realities
  - Structured solutions are evolving
  - Several interesting research problems
- Systematic and semi-automated approach for discovering security zone classifications of devices
  - Staged approach to information collection
  - Elimination-based inferencing
  - Future work
    - Loosening the compliance assumption
    - Evaluating the approach in large-scale infrastructures