Shadow Configurations: A Network Management Primitive

Richard Alimi, Ye Wang, Y. Richard Yang

Laboratory of Networked Systems
Yale University

Configuration is Complex

*"**80%** of IT budgets is used to maintain the status quo."*

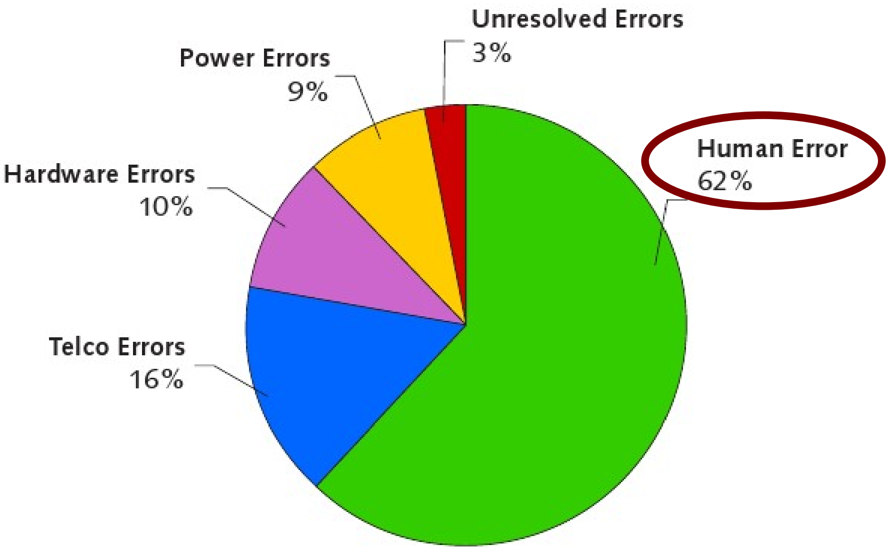

Source: The Yankee Group, 2004

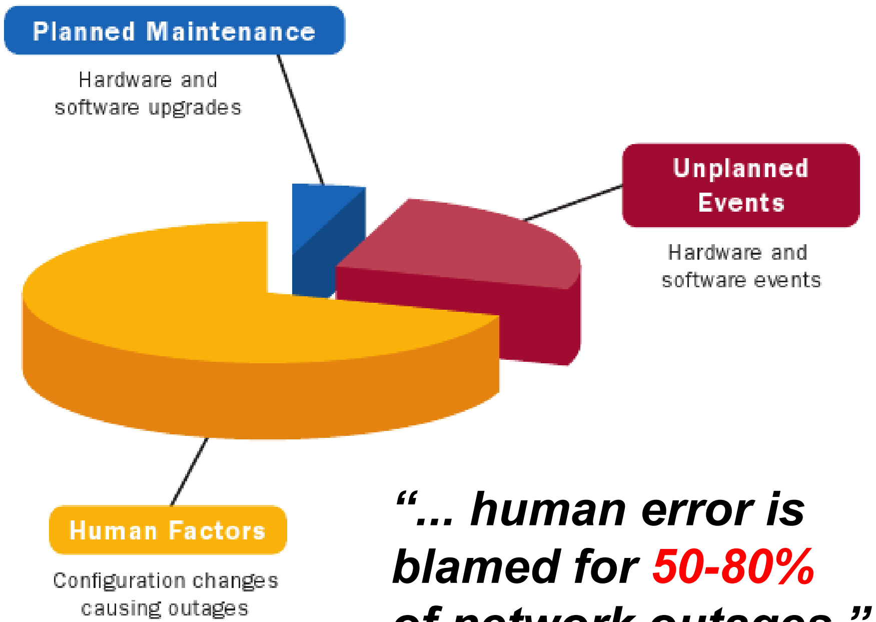

"… human error is blamed for 50-80% of network outages."

Source: Juniper Networks, 2008

Configuration is Complex

"80% of IT budgets is used to maintain the status quo."

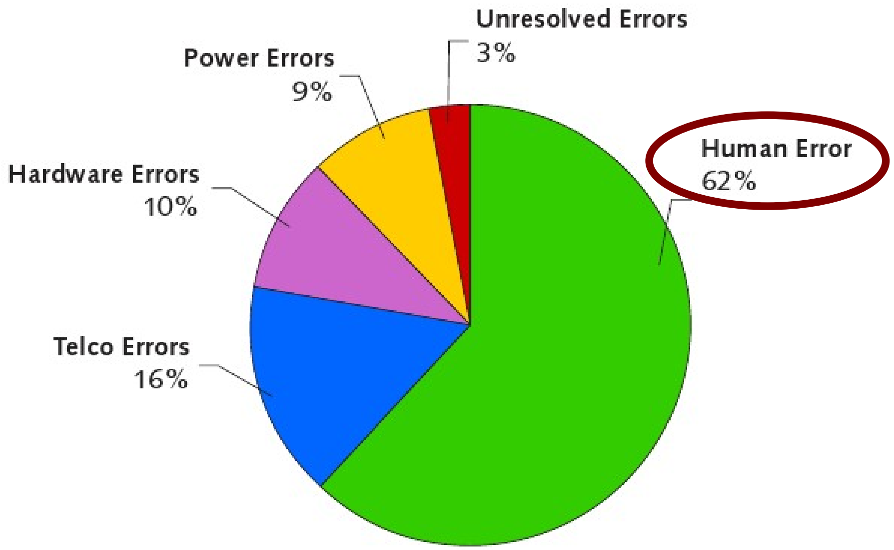

Source: The Yankee Group, 2004

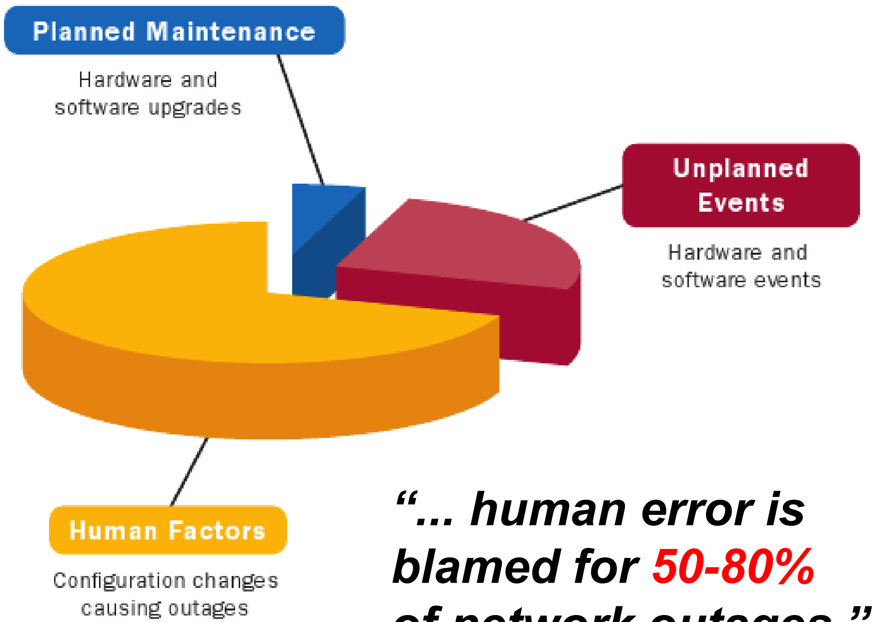

"... human error is blamed for 50-80% of network outages."

Source: Juniper Networks, 2008

Why is configuration hard today?

Configuration Management Today

Simulation & Analysis

- Depend on simplified models
  - Network structure
  - Hardware and software
- Limited scalability
- Hard to access real traffic

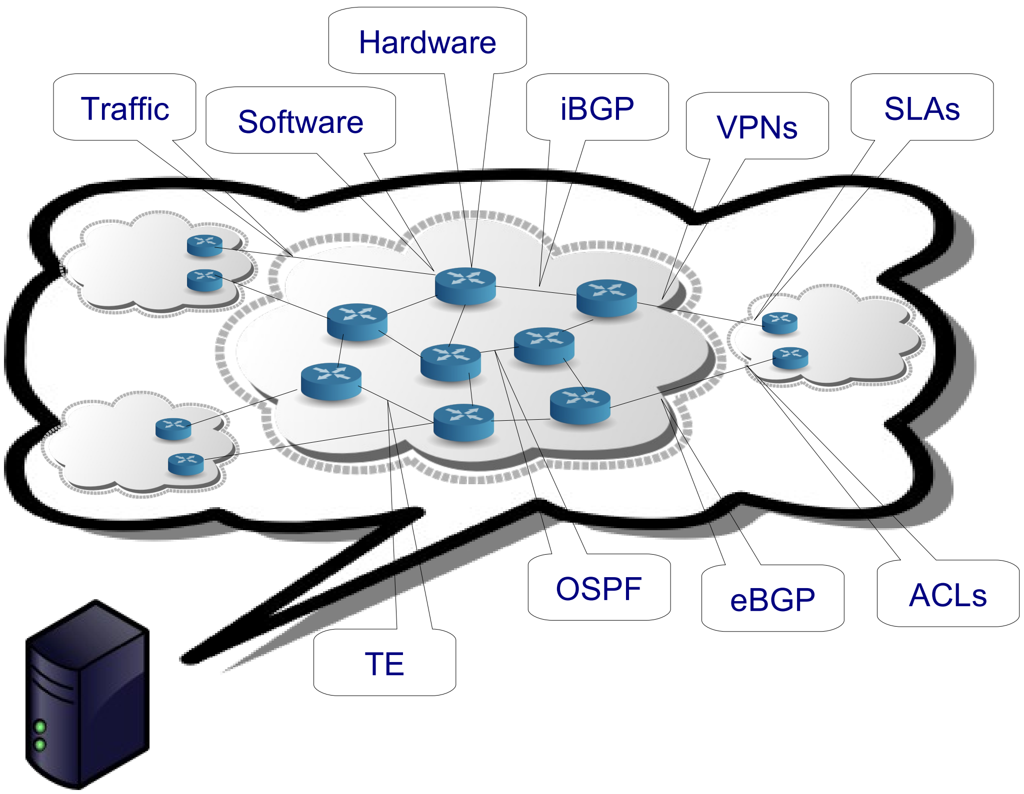

Configuration Management Today

Simulation & Analysis

- Depend on simplified models
  - Network structure
  - Hardware and software
- Limited scalability
- Hard to access real traffic

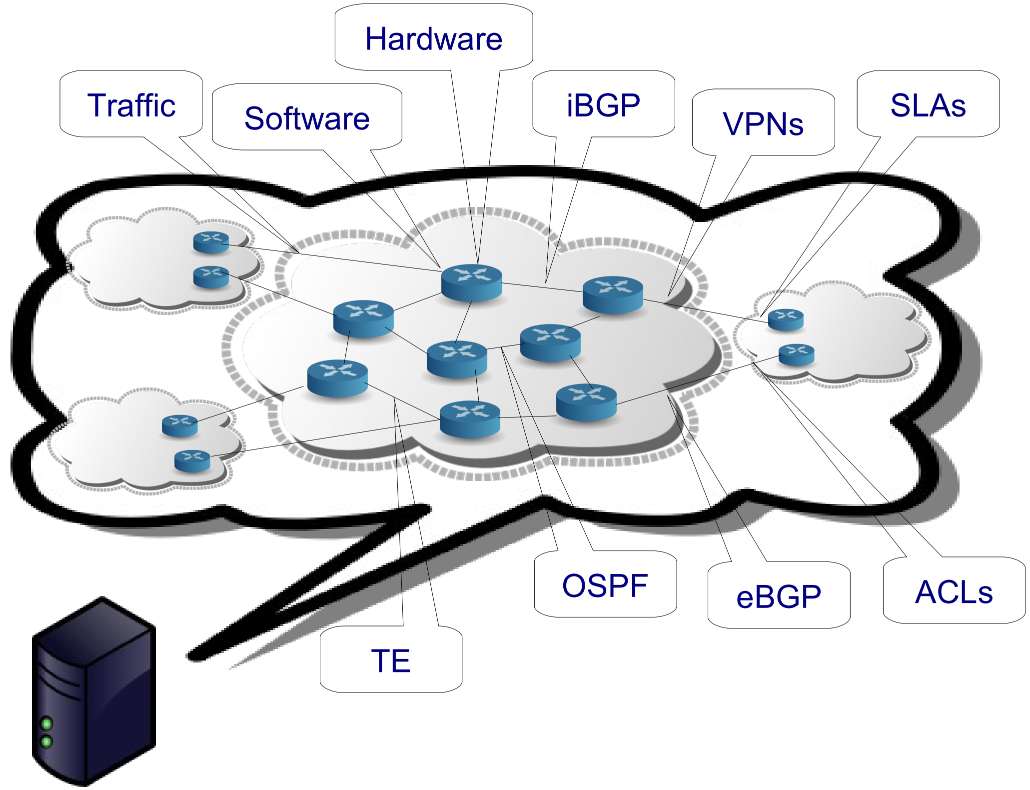

Test networks

- Can be prohibitively expensive

Configuration Management Today

Simulation & Analysis

- Depend on simplified models
  - Network structure
  - Hardware and software
- Limited scalability
- Hard to access real traffic

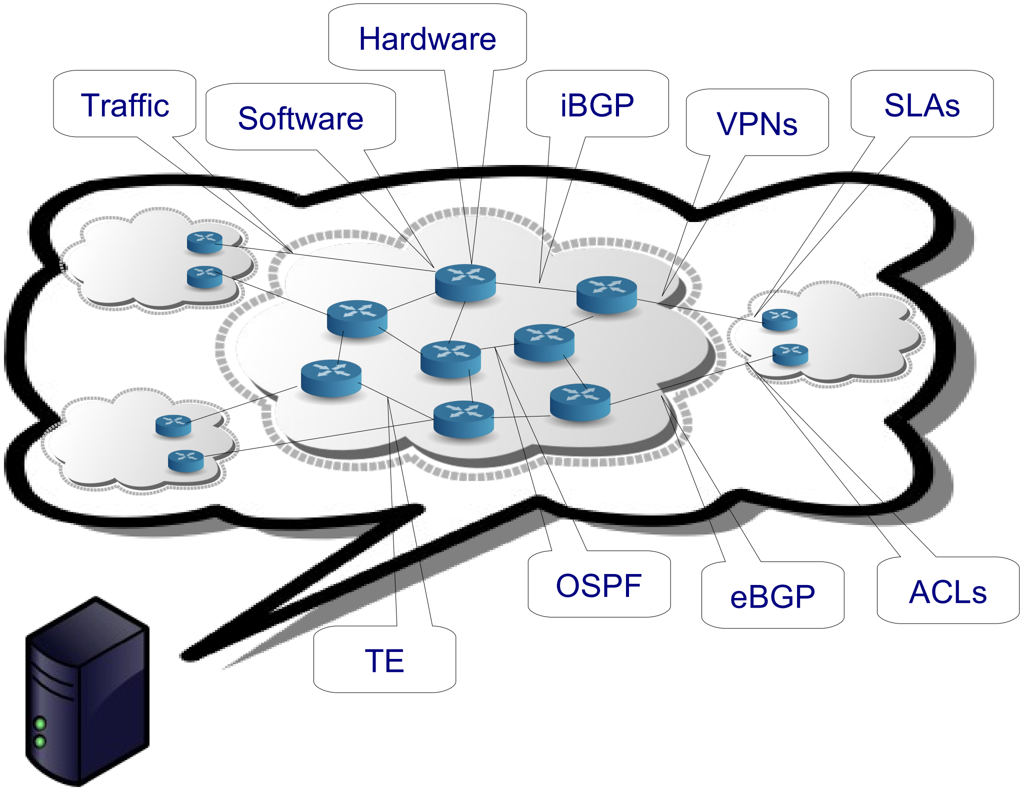

Test networks

- Can be prohibitively expensive

Why are these not enough?

Analogy with Programming

Programming

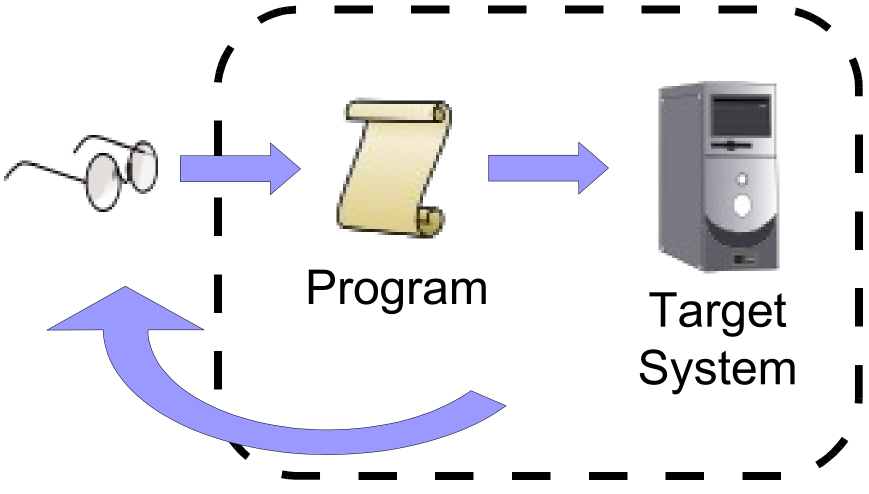

Analogy with Programming

Programming

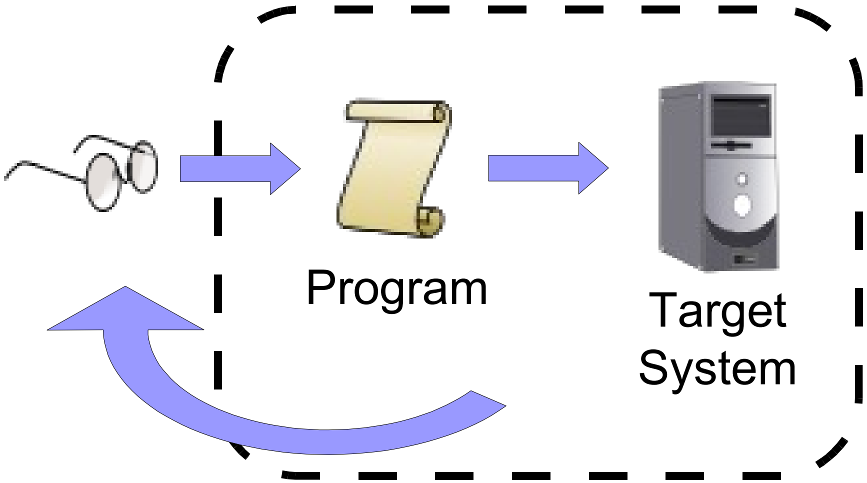

Network Management

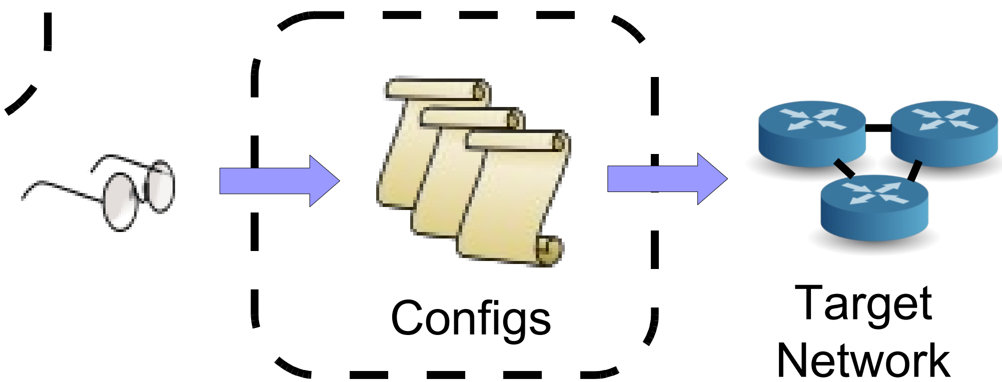

Analogy with Databases

Databases

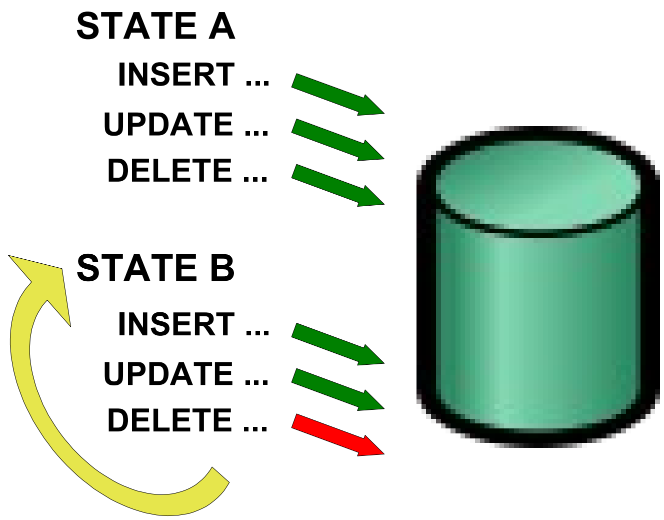

Analogy with Databases

Databases

STATE A

INSERT ...

UPDATE ...

DELETE ...

STATE B

INSERT ...

UPDATE ...

DELETE ...

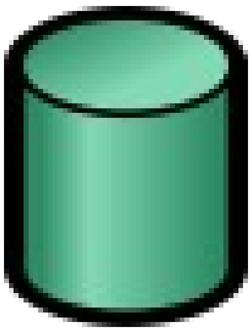

Network Management

STATE A

ip route ...

STATE B

ip addr ...

STATE C

router bgp ...

STATE D

router ospf ...

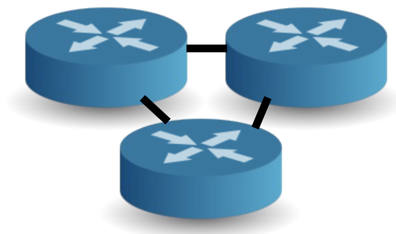

?

Enter, Shadow Configurations

Key ideas

- Allow additional (shadow) config on each router
- In-network, interactive shadow environment
- "Shadow" term from computer graphics

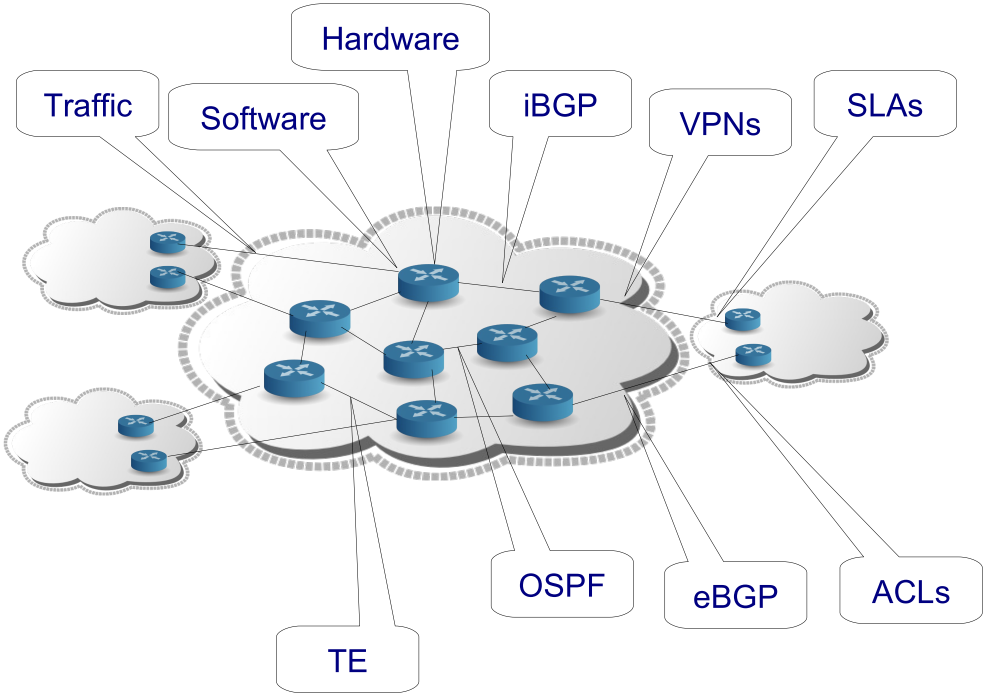

Enter, Shadow Configurations

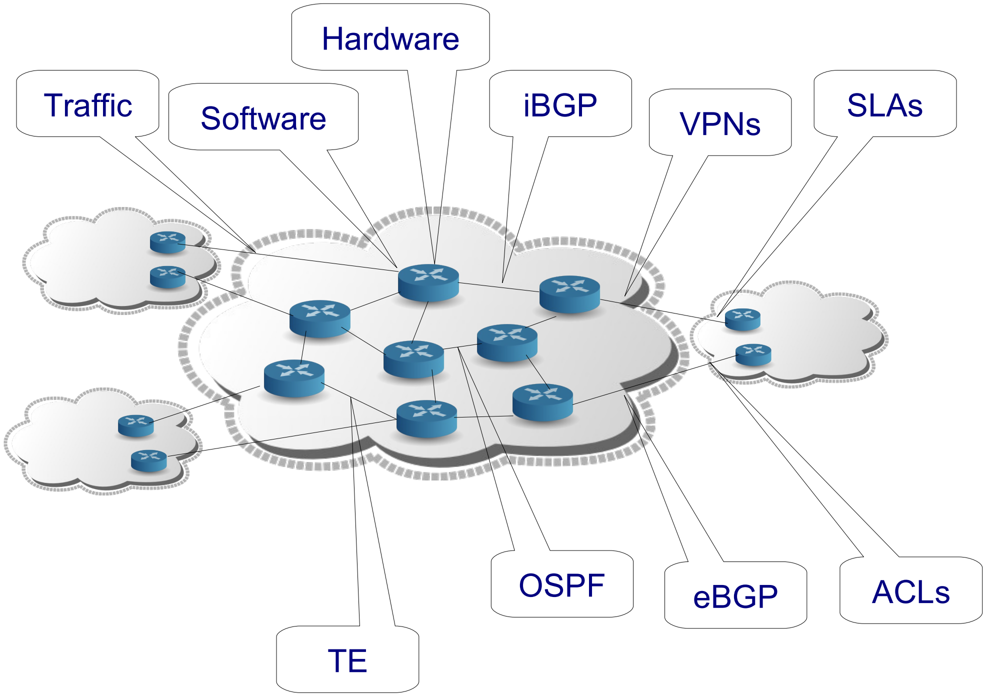

Key ideas

- Allow additional (shadow) config on each router
- In-network, interactive shadow environment
- “Shadow” term from computer graphics

Key Benefits

- Realistic (no model)
- Scalable
- Access to real traffic
- Transactional

Roadmap

Motivation and Overview

System Basics and Usage

System Components

- Design and Architecture
- Performance Testing
- Transaction Support

Implementation and Evaluation

System Basics

What's in the shadow configuration?

- Routing parameters
- ACLs
- Interface parameters
- VPNs
- QoS parameters

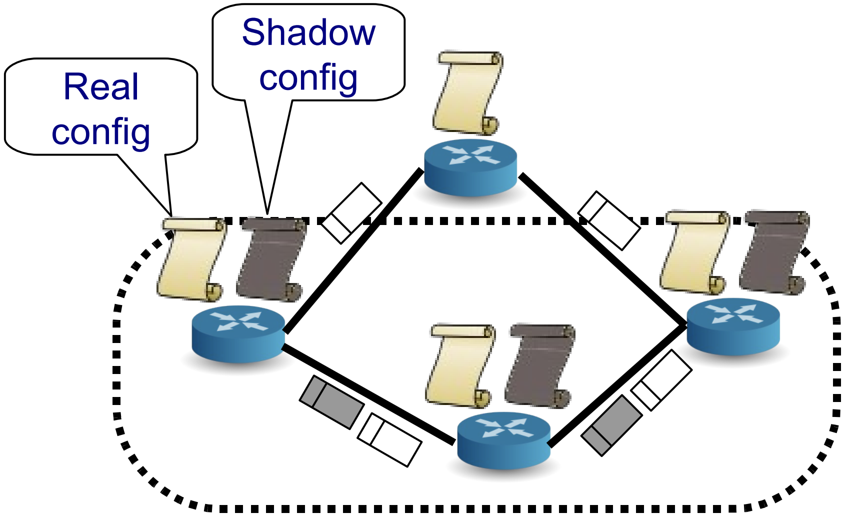

System Basics

What's in the shadow configuration?

- Routing parameters
- ACLs
- Interface parameters
- VPNs
- QoS parameters

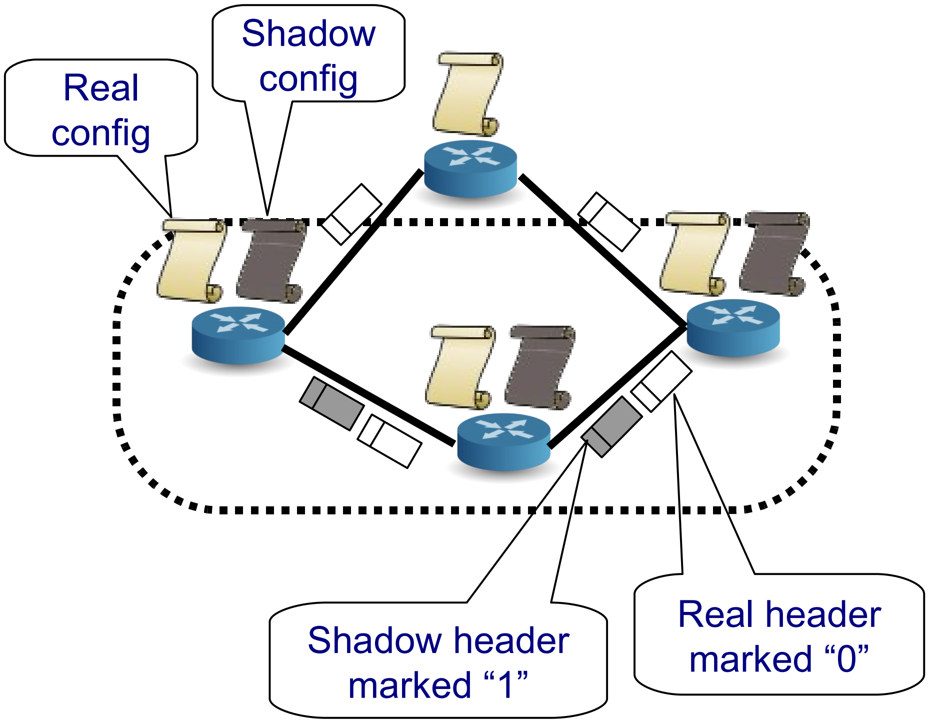

Example Usage Scenario: Backup Path Verification

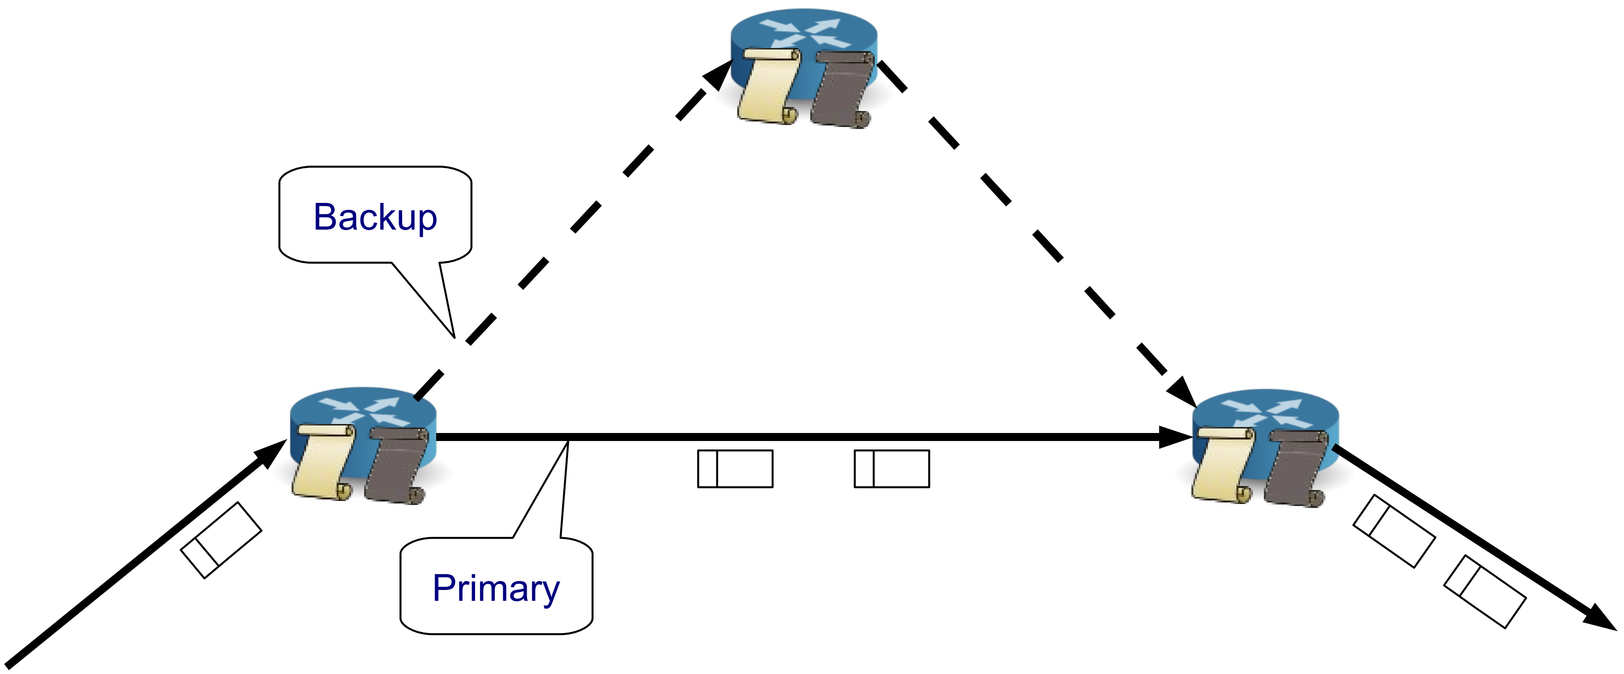

Example Usage Scenario: Backup Path Verification

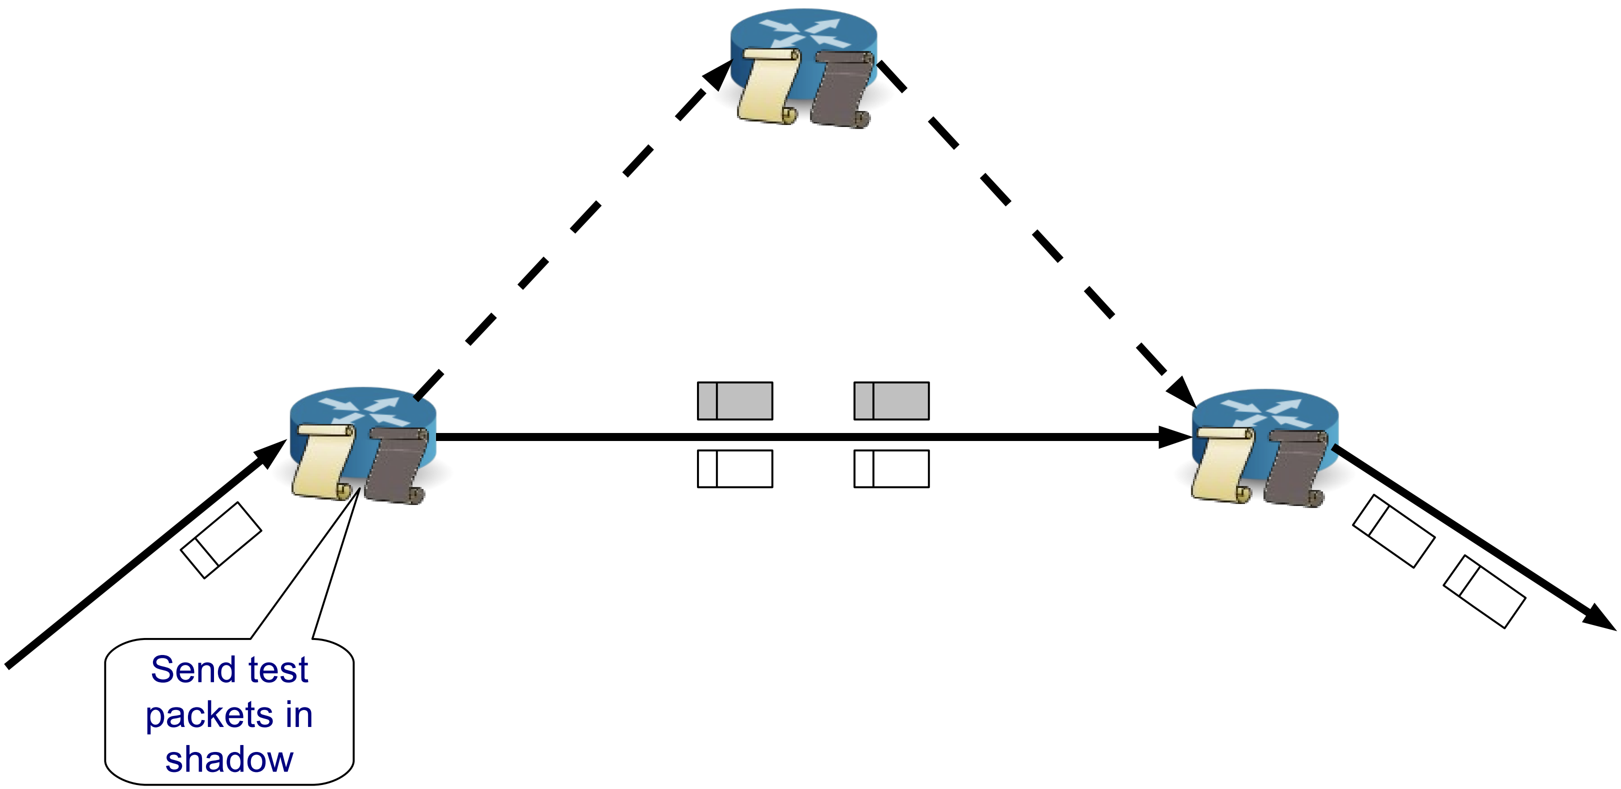

Example Usage Scenario: Backup Path Verification

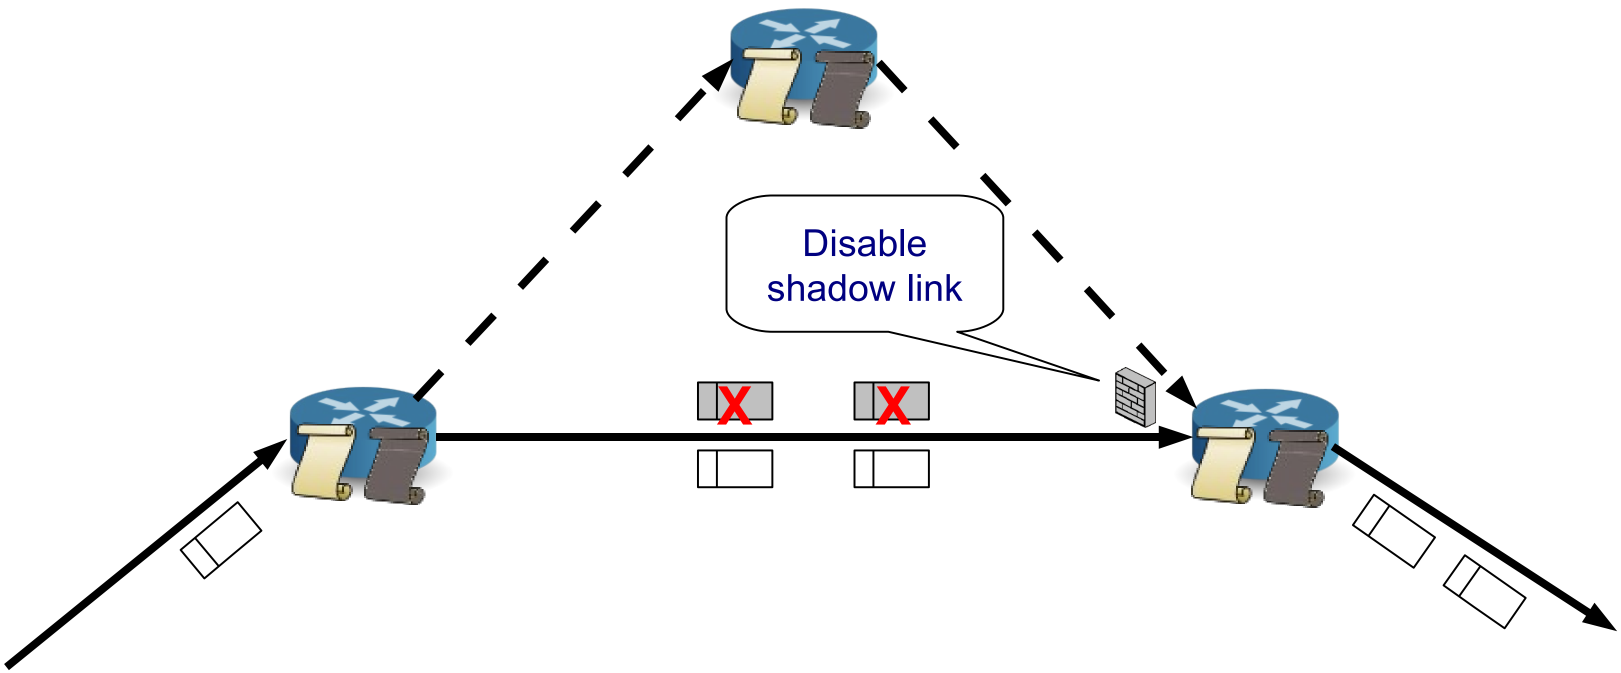

Example Usage Scenario: Backup Path Verification

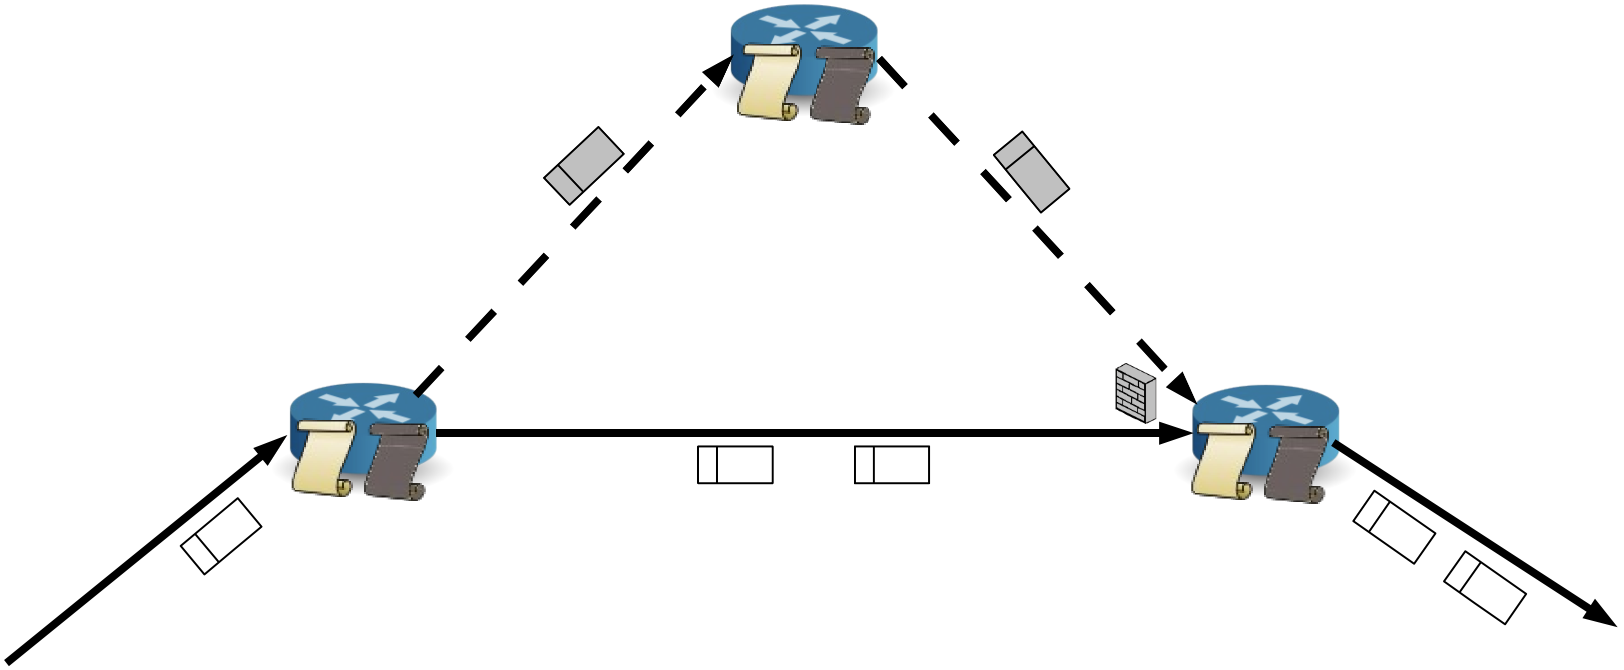

Example Usage Scenario: Configuration Evaluation

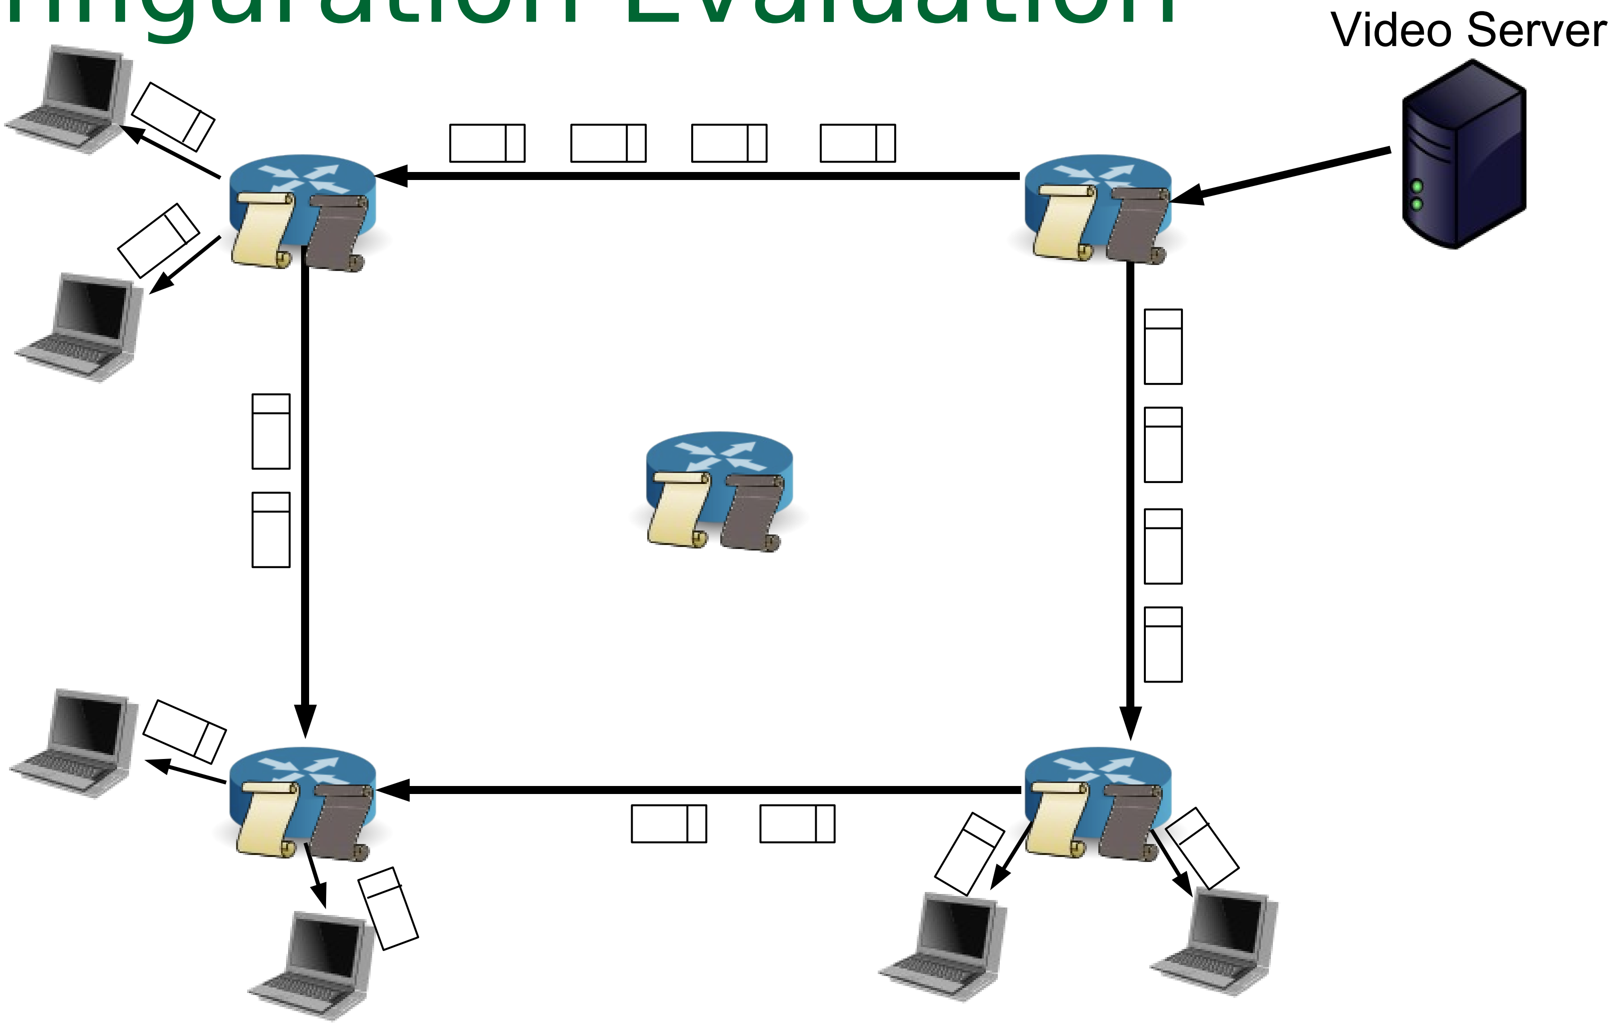

Example Usage Scenario: Configuration Evaluation

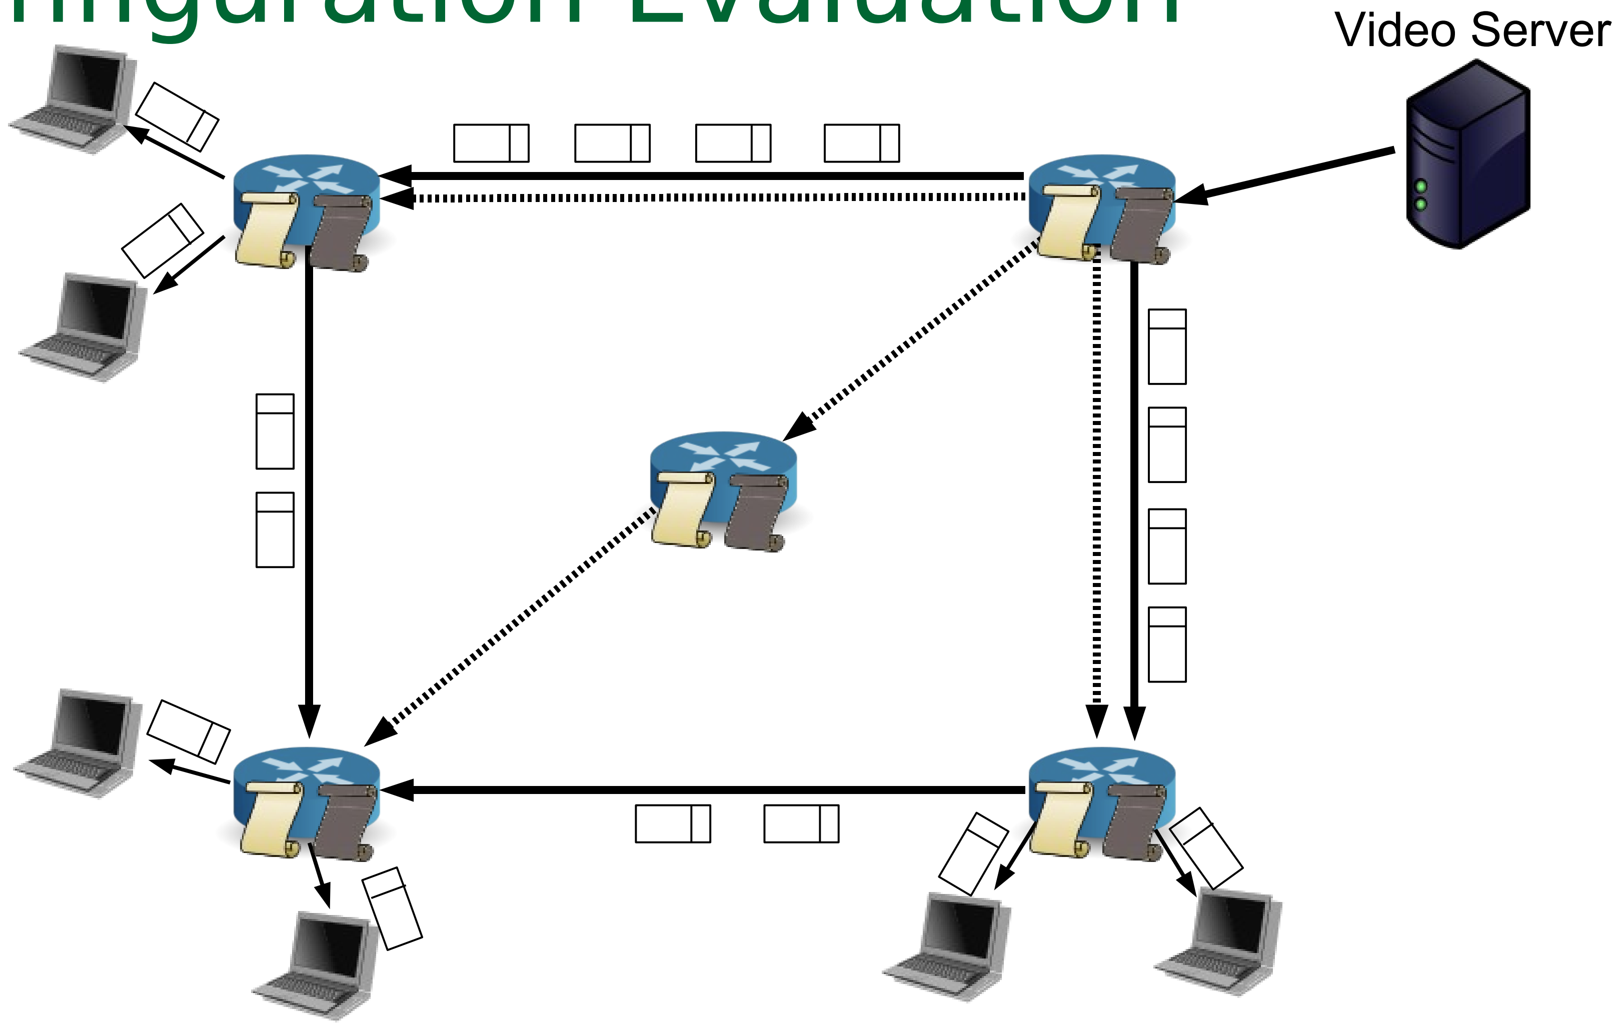

Example Usage Scenario: Configuration Evaluation

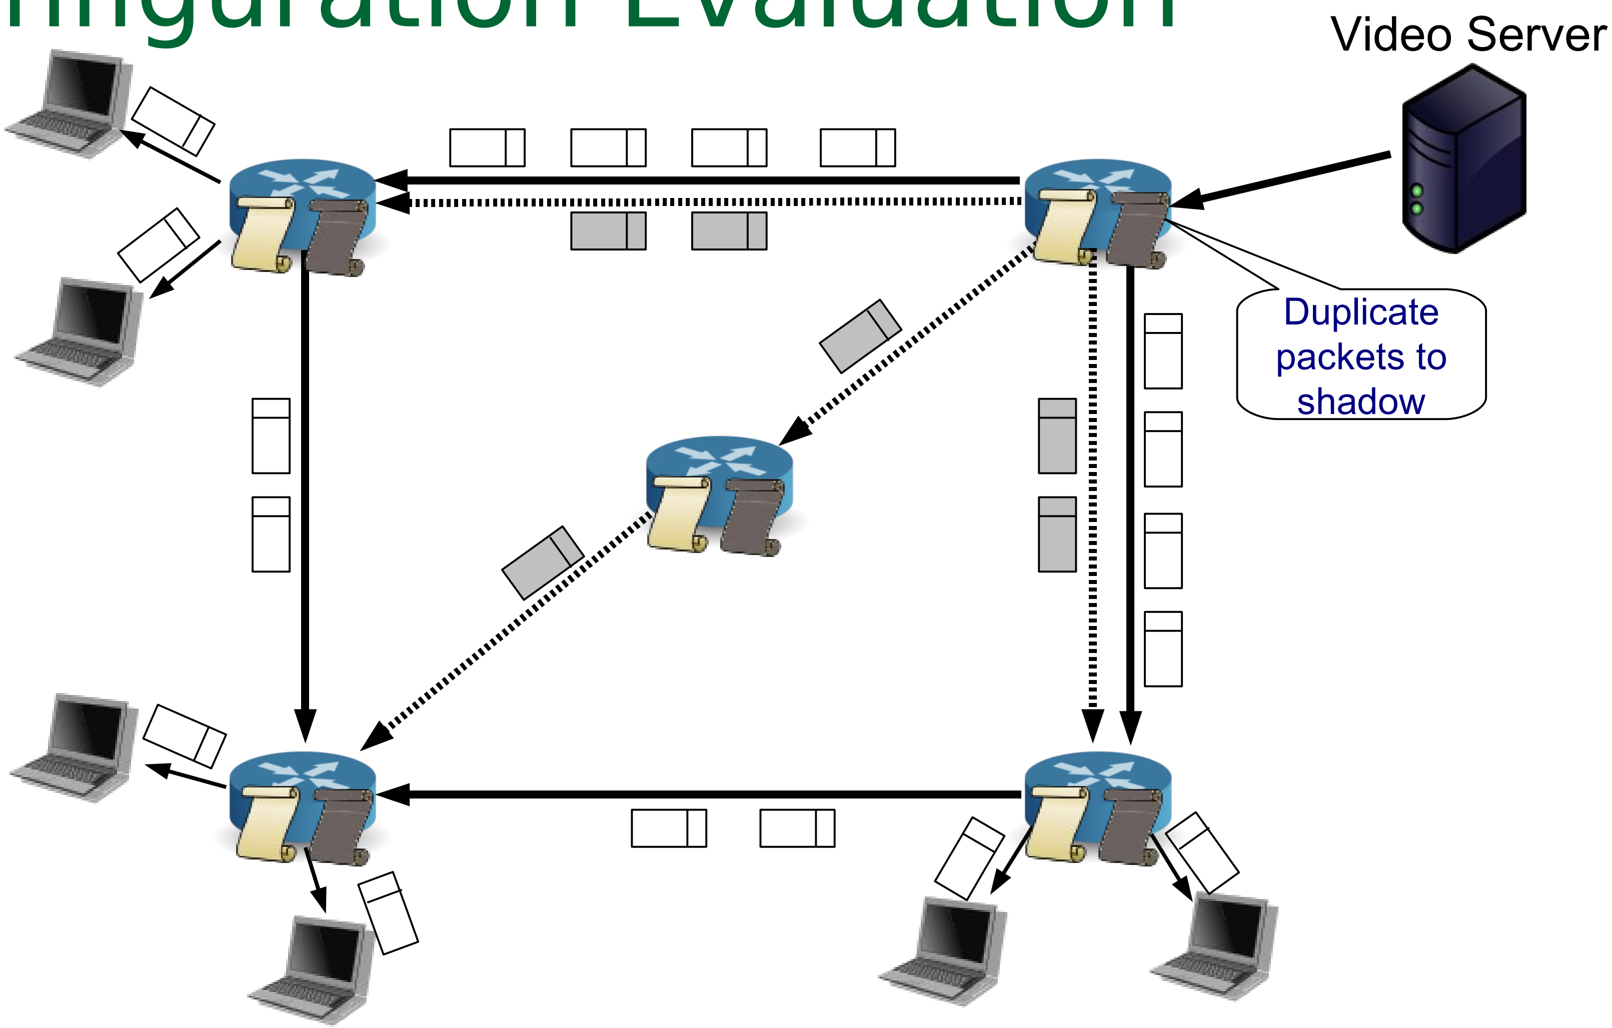

Roadmap

Motivation and Overview

System Basics and Usage

System Components

- ***Design and Architecture***
- ***Performance Testing***
- ***Transaction Support***

Implementation and Evaluation

Design and Architecture

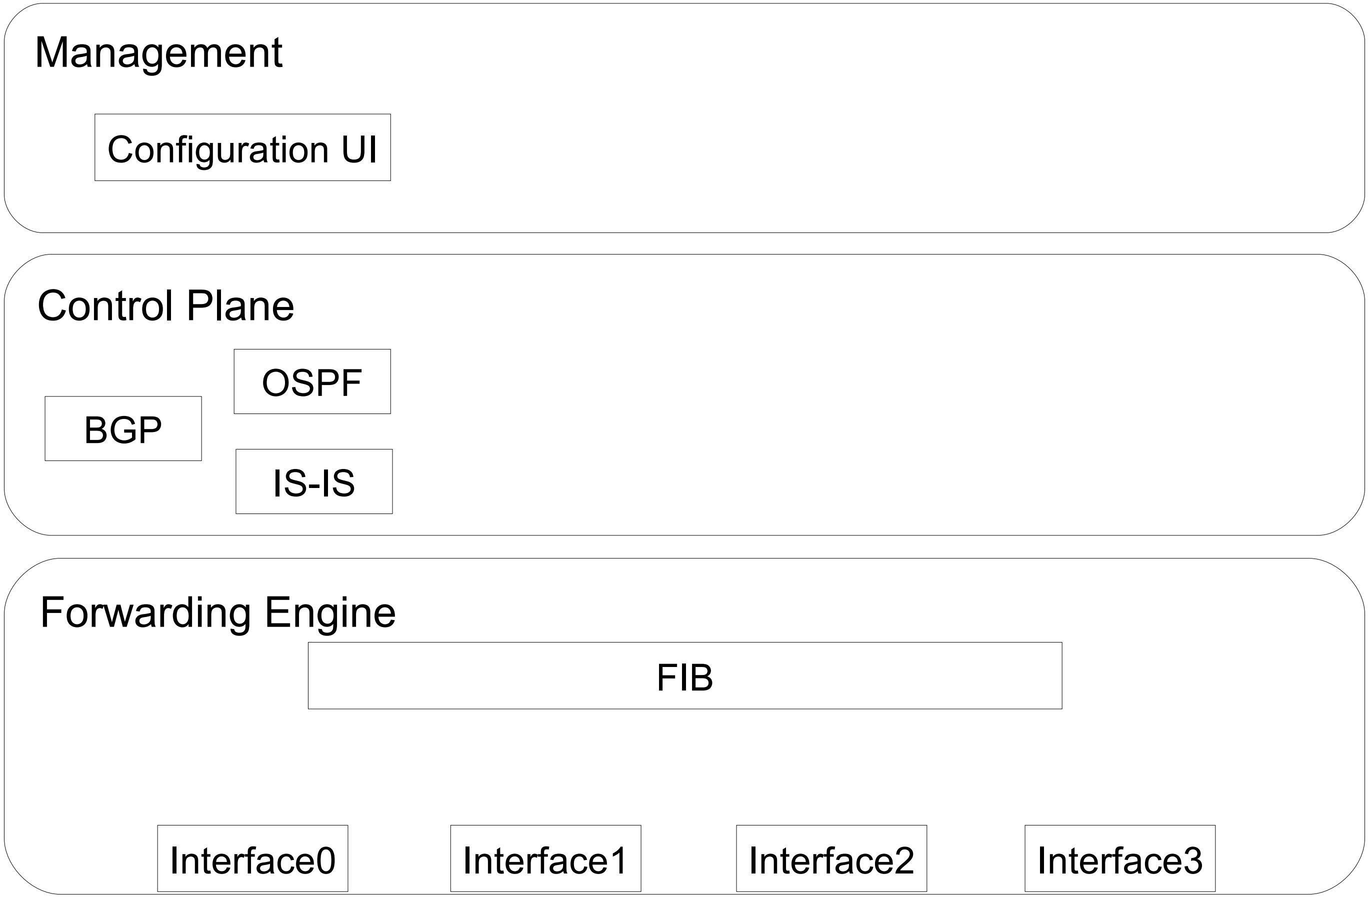

Design and Architecture

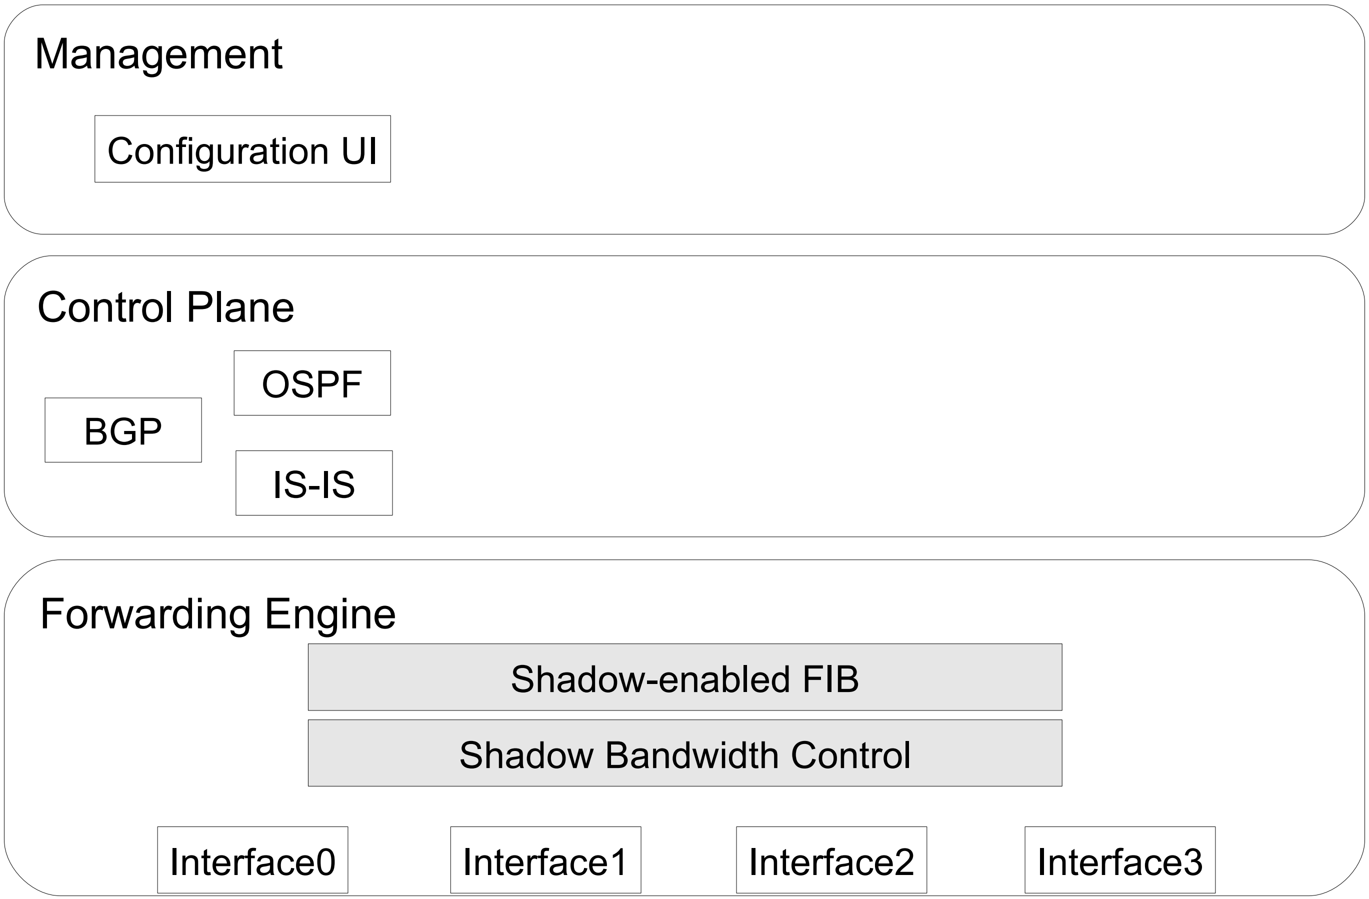

Design and Architecture

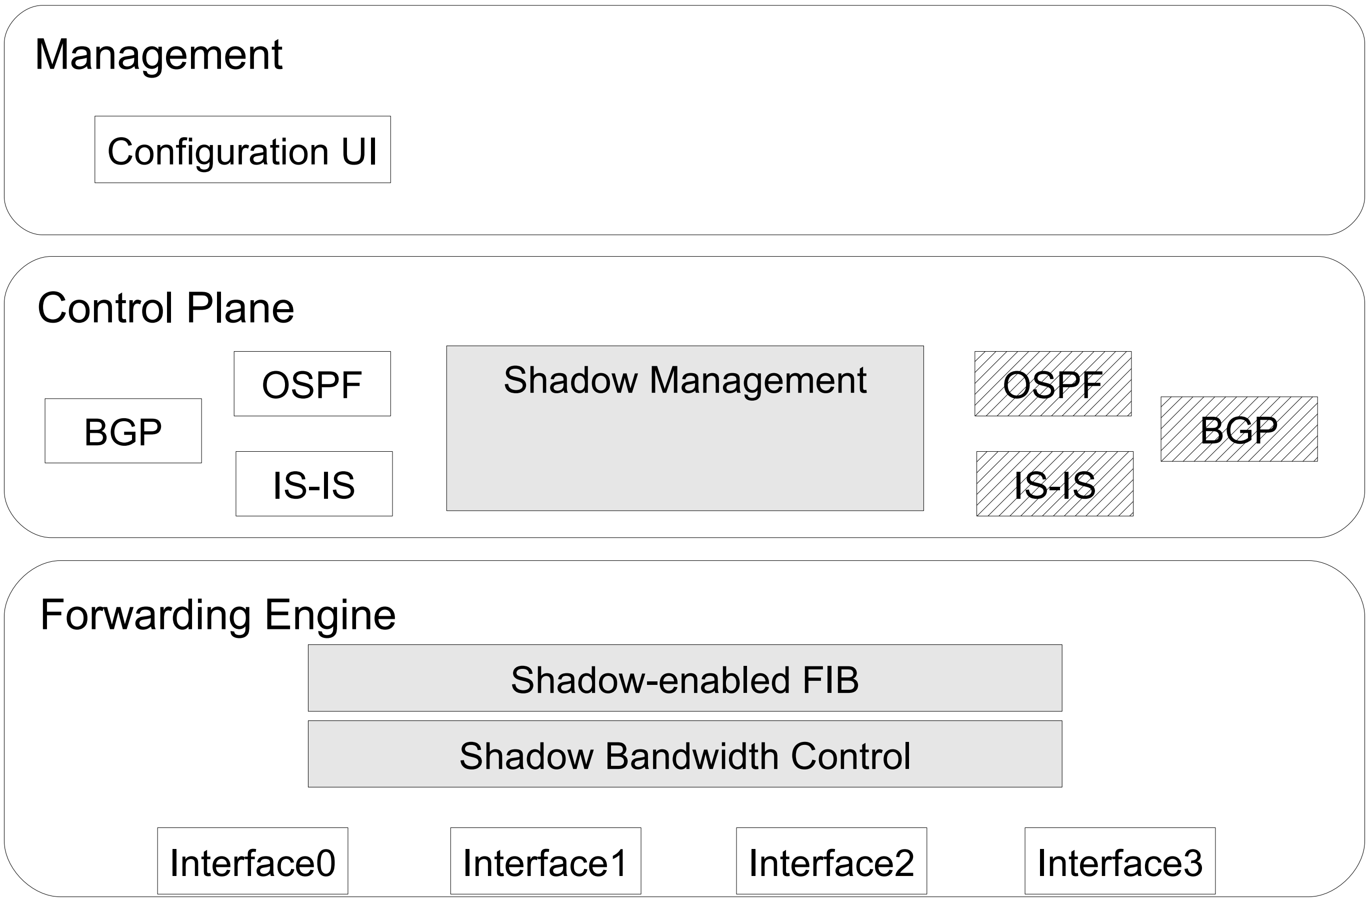

Design and Architecture

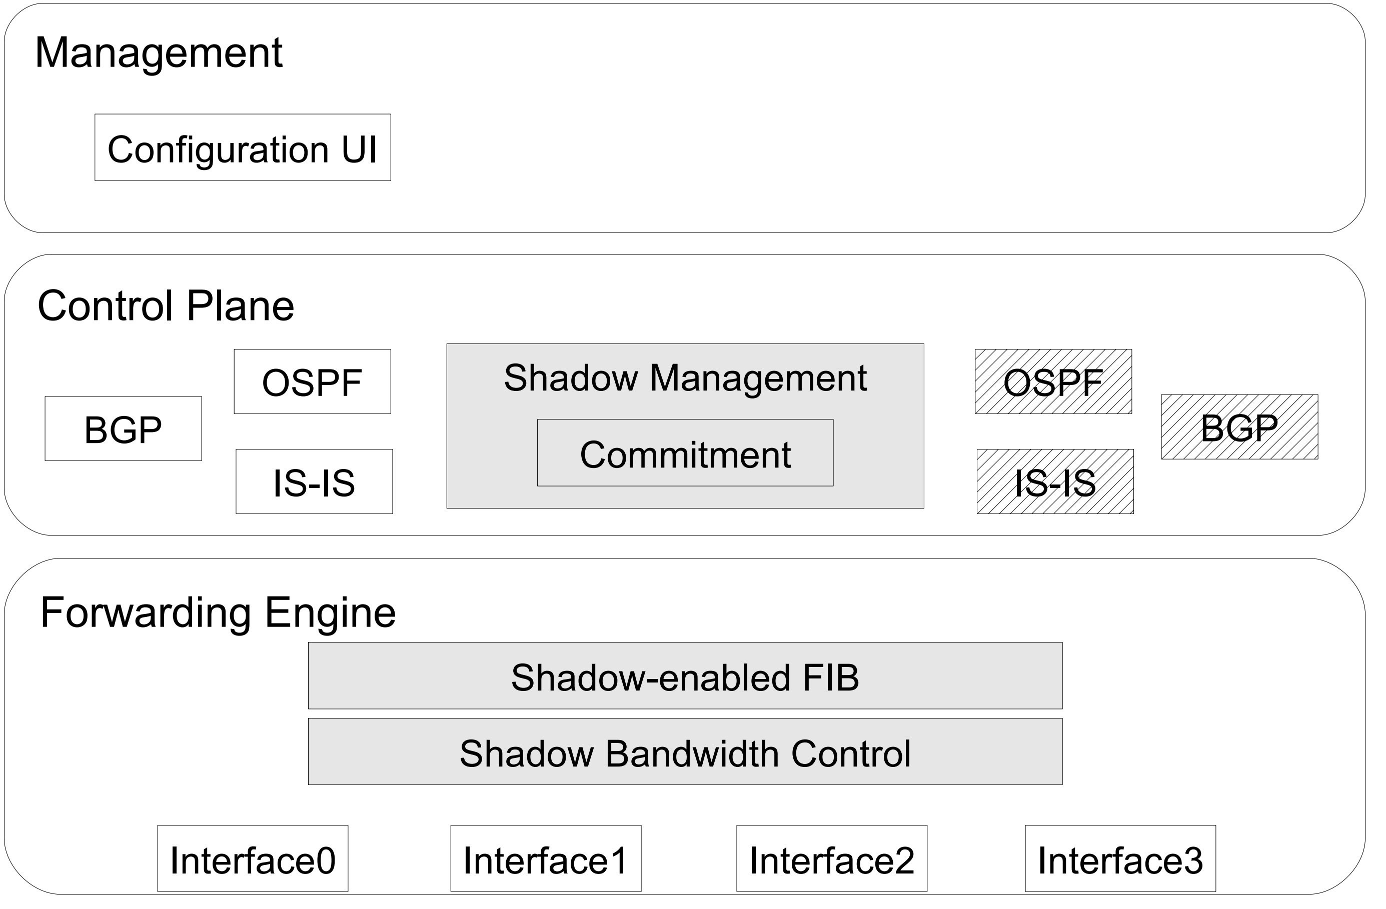

Design and Architecture

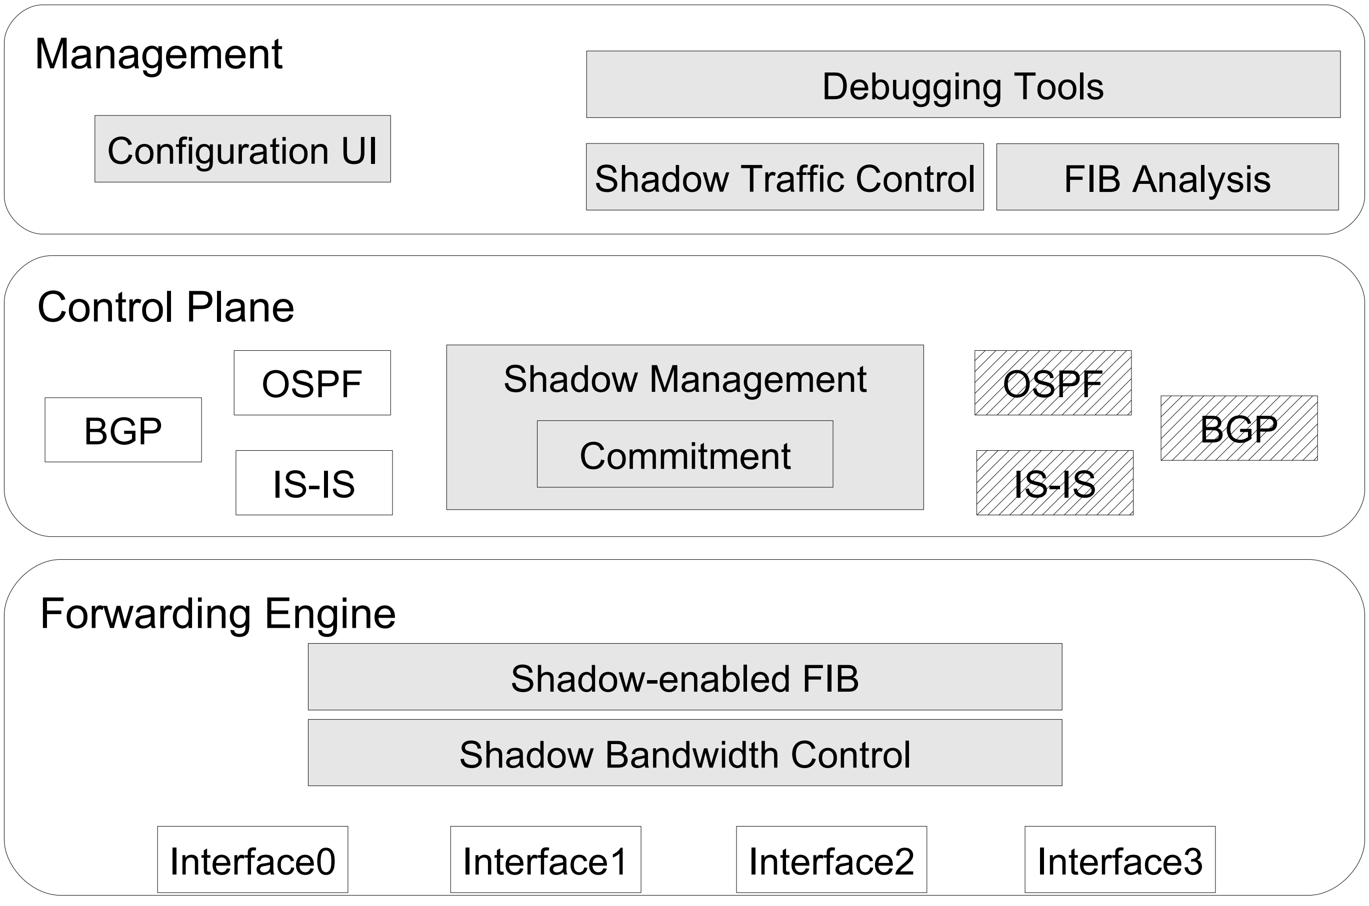

Design and Architecture

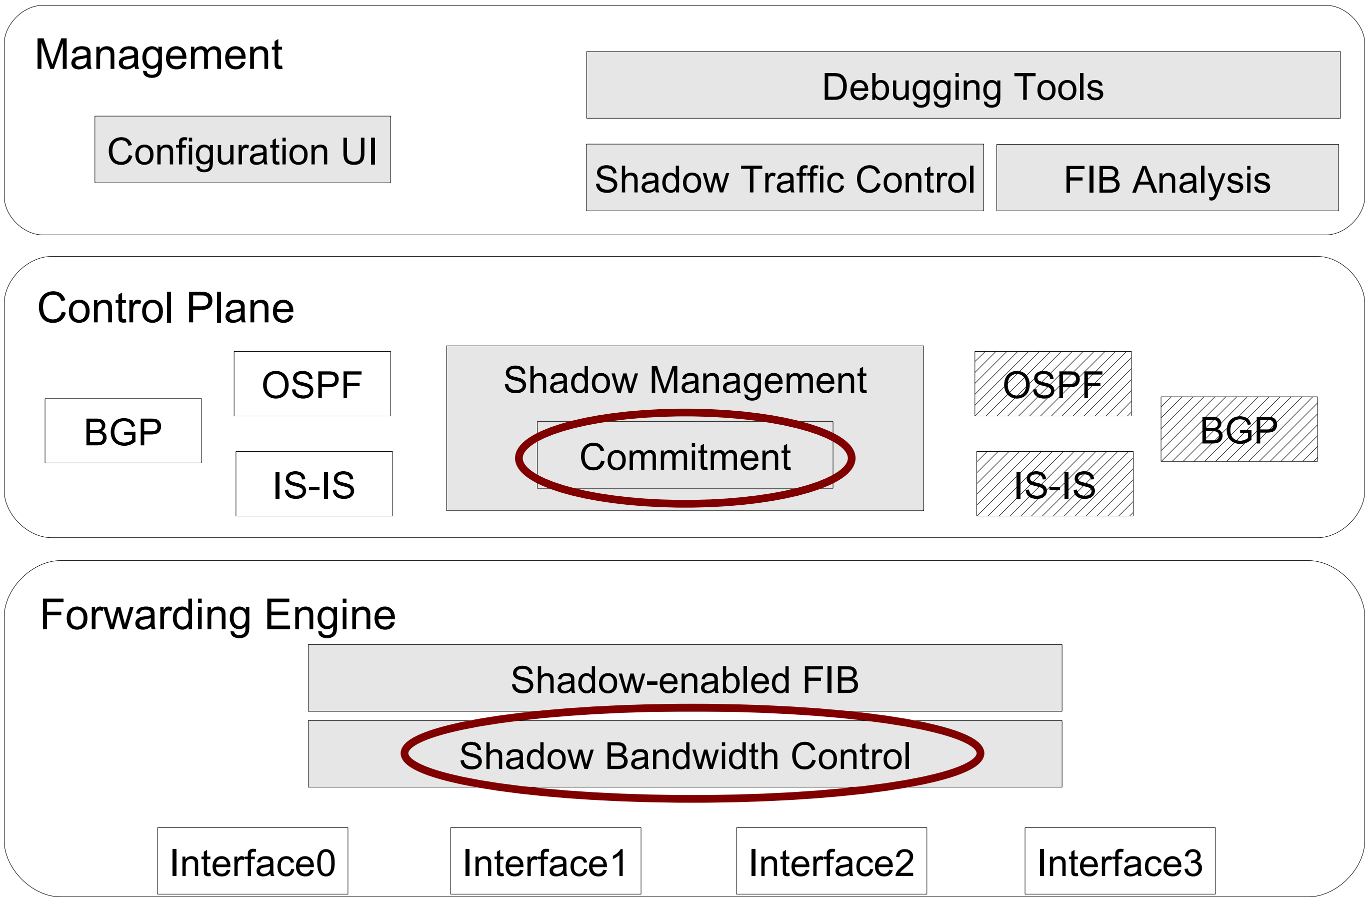

Shadow Bandwidth Control

Requirements

- Minimal impact on real traffic
- Accurate performance measurements of shadow configuration

Shadow Bandwidth Control

Requirements
- Minimal impact on real traffic
- Accurate performance measurements of shadow configuration

Supported Modes
- Priority
- Bandwidth Partitioning
- Packet Cancellation

Packet Cancellation

Observation

- Content of payload may not important in many network performance testing scenarios
- Only payload size may matter

Packet Cancellation

Observation

- Content of payload may not important in many network performance testing scenarios
- Only payload size may matter

Idea: only need headers for shadow traffic

Piggyback shadow headers on real packets

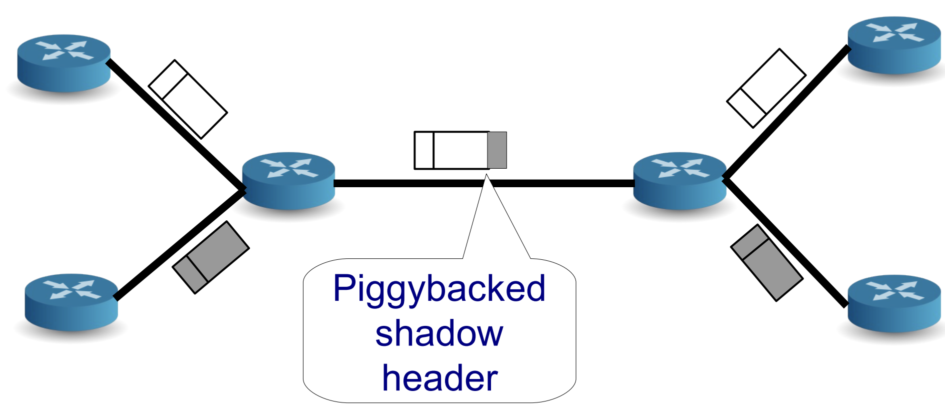

Packet Cancellation Details

Output interface maintains real and shadow queues

Packet cancellation scheduling

- If real queue non-empty
  - Grab real packet
  - Piggyback shadow header(s) if available
- Else if shadow queue non-empty
  - Send full shadow packet

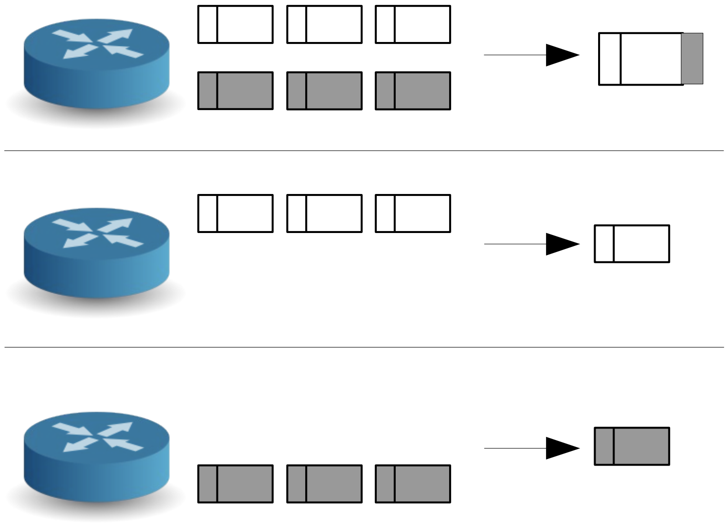

Commitment

Objectives

- Smoothly swap real and shadow across network
  - Eliminate effects of transient states due to config changes
- Easy to swap back

Commitment

Objectives

- Smoothly swap real and shadow across network
  - Eliminate effects of transient states due to config changes
- Easy to swap back

Issue

- Packet marked with *shadow* bit
  - 0 = Real, 1 = Shadow
- Shadow bit determines which FIB to use
- Routers swap FIBs asynchronously
- Inconsistent FIBs applied on the path

Commitment Protocol

Idea: Use tags to achieve consistency

- Temporary identifiers

Basic algorithm has 4 phases

Commitment Protocol

Idea: Use tags to achieve consistency

- Temporary identifiers

Basic algorithm has 4 phases

- Distribute tags for each config
  - **C-old** for current real config
  - **C-new** for current shadow config

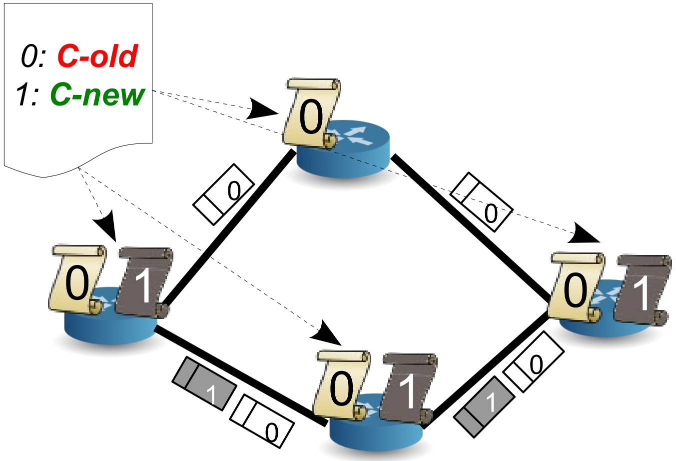

Commitment Protocol

Idea: Use tags to achieve consistency

- Temporary identifiers

Basic algorithm has 4 phases

- Distribute tags for each config
  - **C-old** for current real config
  - **C-new** for current shadow config
- Routers mark packets with tags

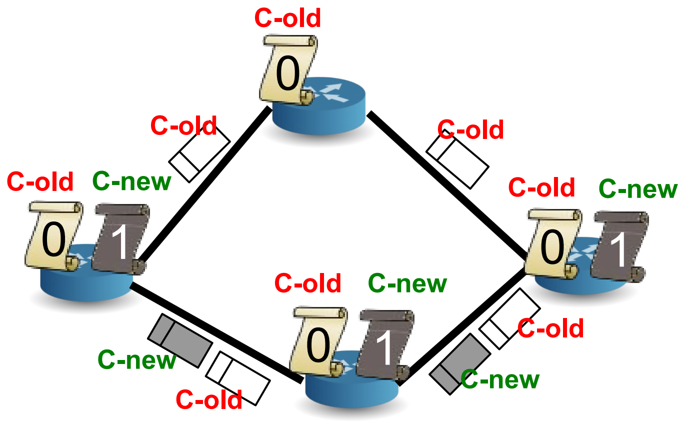

Commitment Protocol

Idea: Use tags to achieve consistency

- Temporary identifiers

Basic algorithm has 4 phases

- Distribute tags for each config
  - **C-old** for current real config
  - **C-new** for current shadow config
- Routers mark packets with tags
- Swap configs (tags still valid)

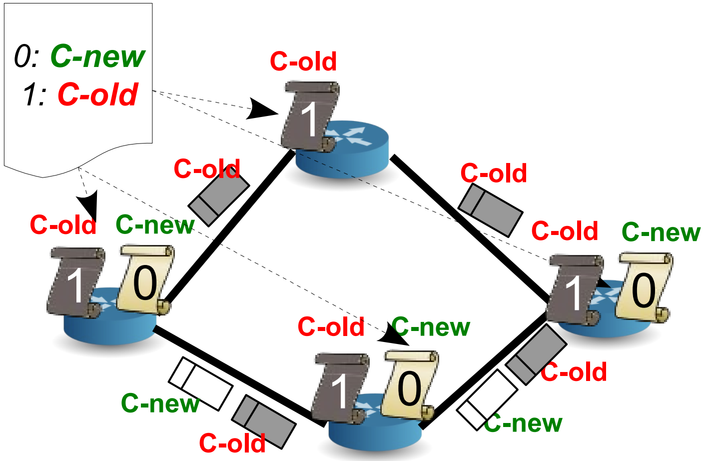

Commitment Protocol

Idea: Use tags to achieve consistency

- Temporary identifiers

Basic algorithm has 4 phases

- Distribute tags for each config
  - **C-old** for current real config
  - **C-new** for current shadow config
- Routers mark packets with tags
- Swap configs (tags still valid)
- Remove tags from packets
  - Resume use of shadow bit

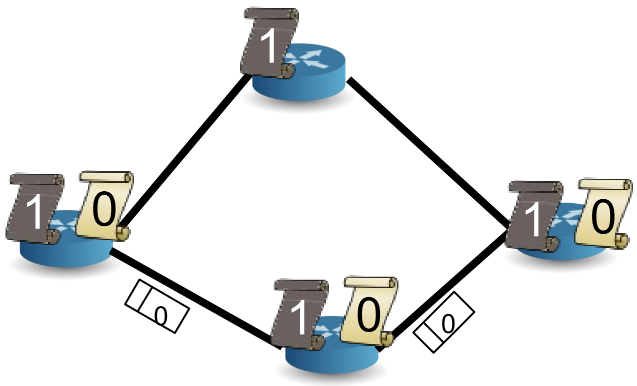

Commitment Protocol

Idea: Use tags to achieve consistency

- Temporary identifiers

Basic algorithm has 4 phases

- Distribute tags for each config
  - **C-old** for current real config
  - **C-new** for current shadow config
- Routers mark packets with tags
- Swap configs (tags still valid)
- Remove tags from packets
  - Resume use of shadow bit
- *For more details, see paper*

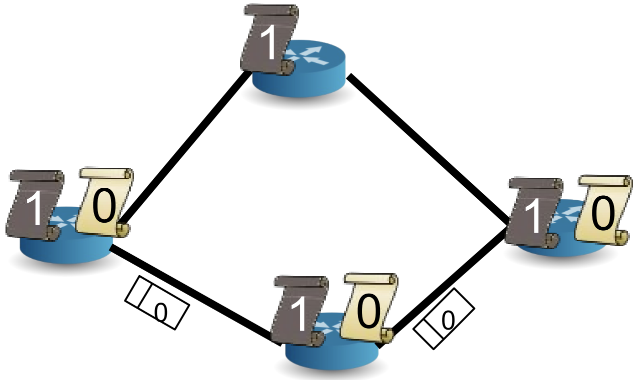

Roadmap

Motivation and Overview

System Basics and Usage

System Components

- Design and Architecture
- Performance Testing
- Transaction Support

Implementation and Evaluation

Implementation

Kernel-level (based on Linux 2.6.22.9)

- TCP/IP stack support
- FIB management
- Commitment hooks
- Packet cancellation

Tools

- Transparent software router support (Quagga + XORP)
- Full commitment protocol
- Configuration UI (command-line based)

Evaluated on Emulab (3Ghz HT CPUs)

Evaluation: CPU Overhead

Static FIB

- 300B pkts
- No route caching

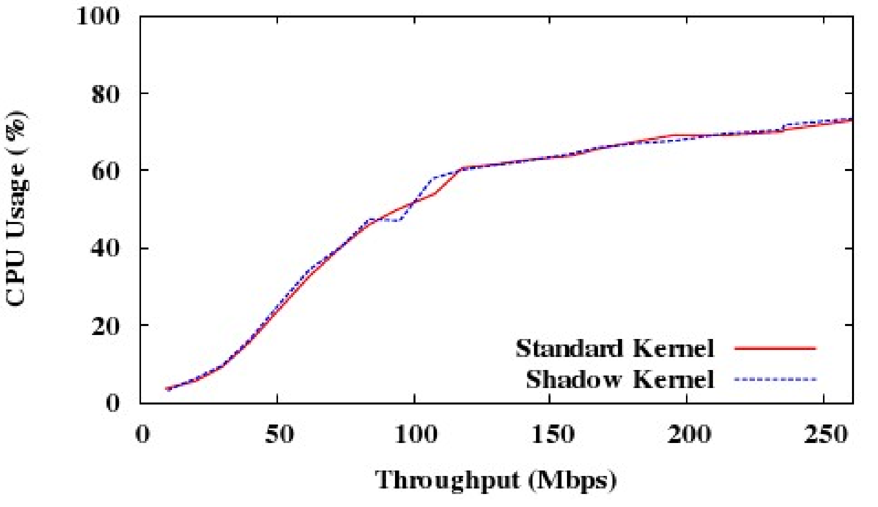

With FIB updates

- 300B pkts @ 100Mbps
- 1-100 updates/sec
- No route caching

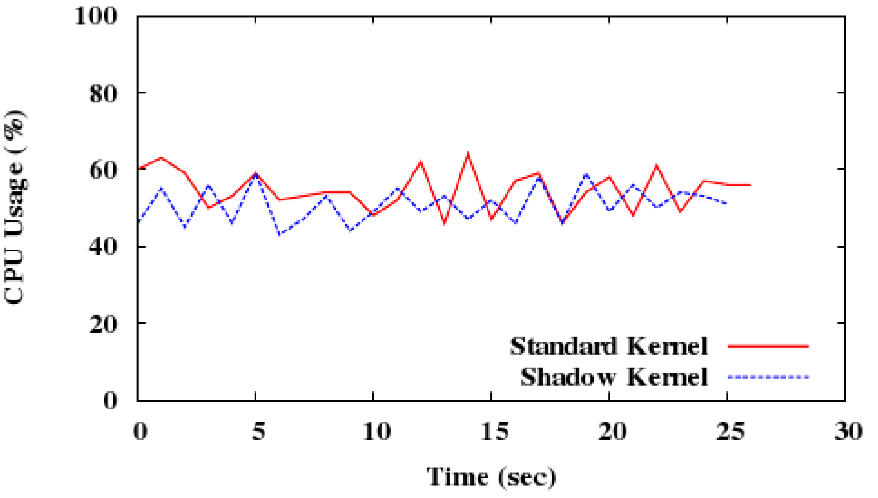

Evaluation: Memory Overhead

FIB storage overhead for US Tier-1 ISP

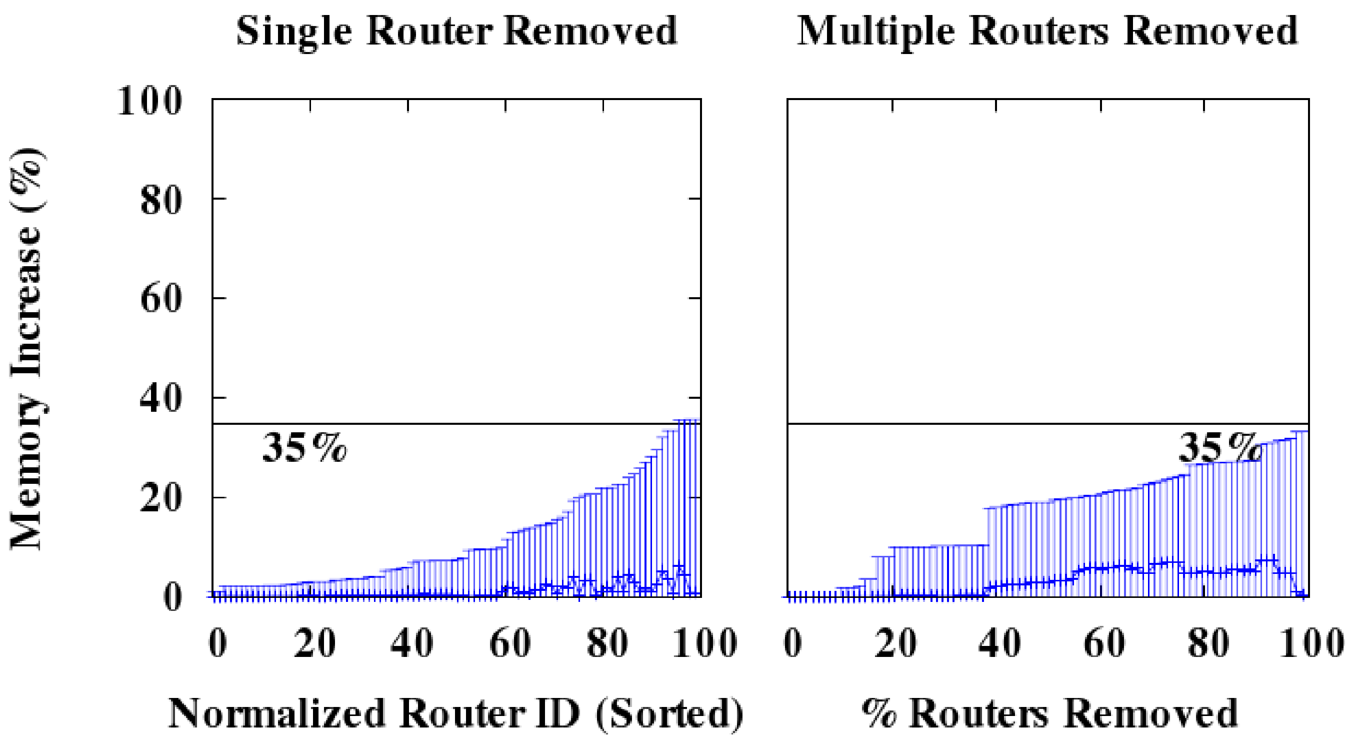

Evaluation: Packet Cancellation

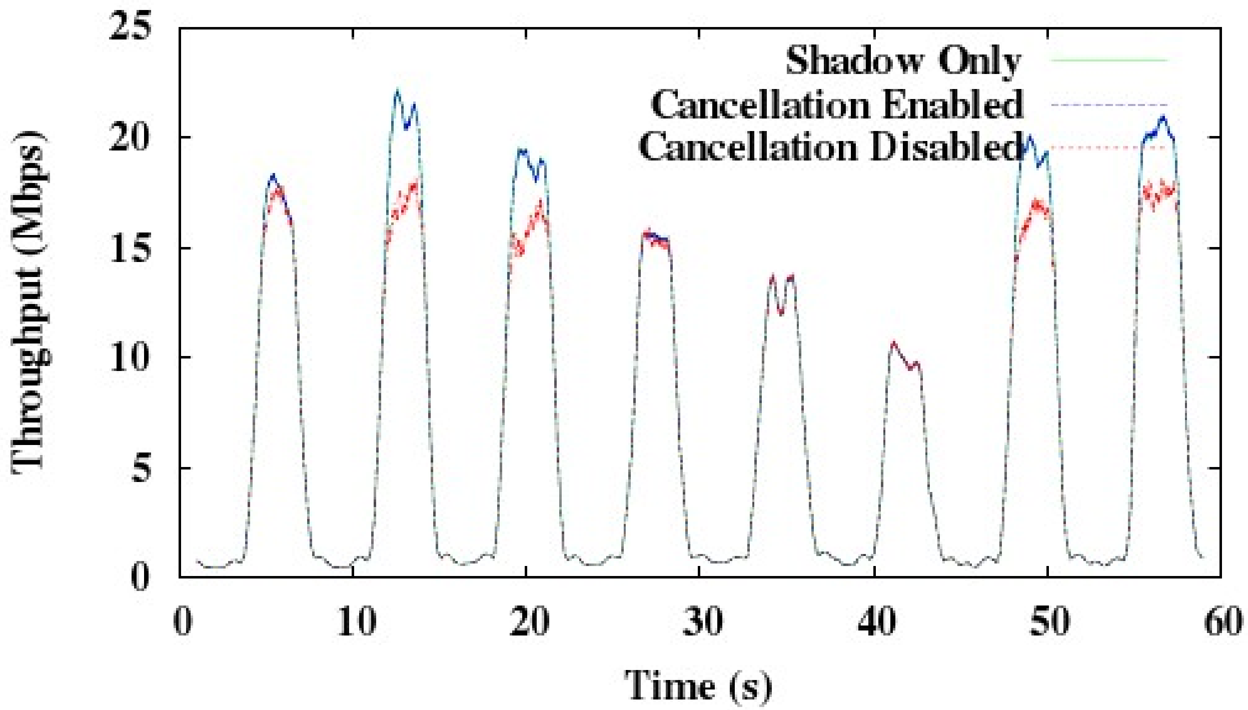

Accurate streaming throughput measurement

- Abilene topology
- Real transit traffic duplicated to shadow
- Video streaming traffic in shadow

Evaluation: Packet Cancellation

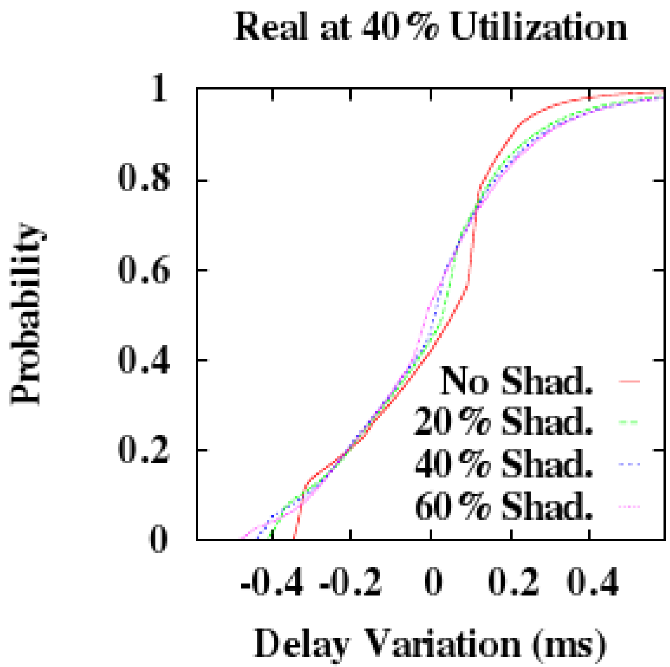

Limited interaction of real and shadow

- Intersecting real and shadow flows
  - CAIDA traces
- Vary flow utilizations

Evaluation: Packet Cancellation

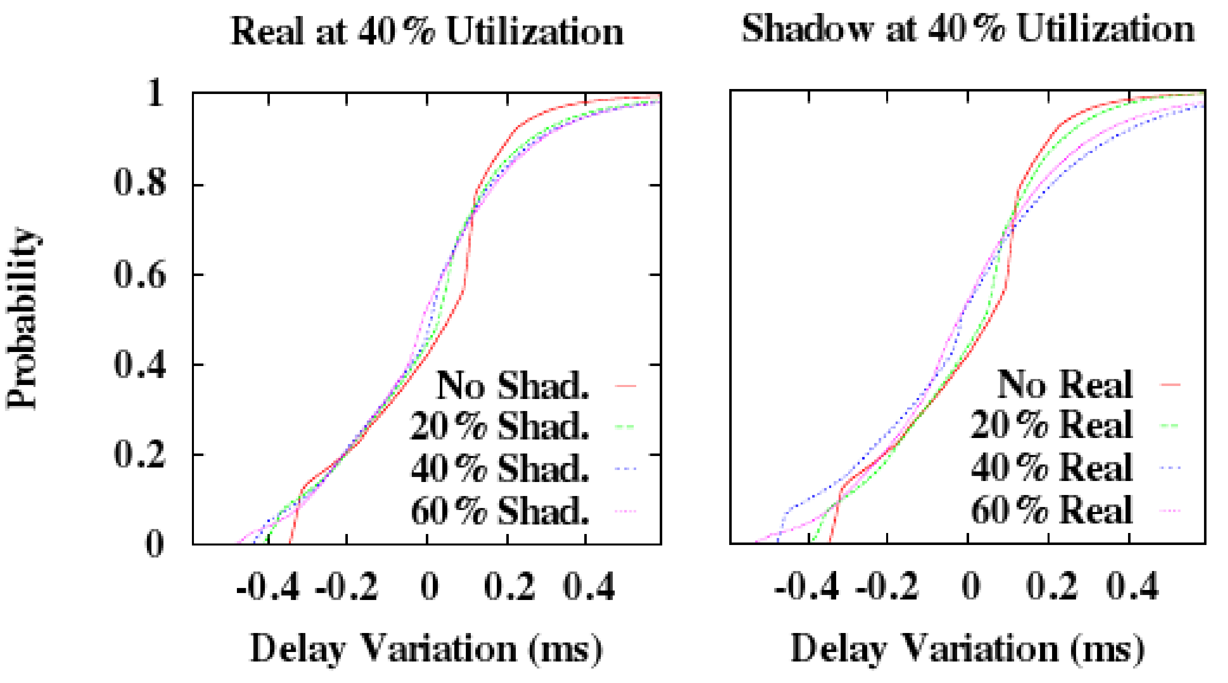

Limited interaction of real and shadow

- Intersecting real and shadow flows
  - CAIDA traces
- Vary flow utilizations

Evaluation: Commitment

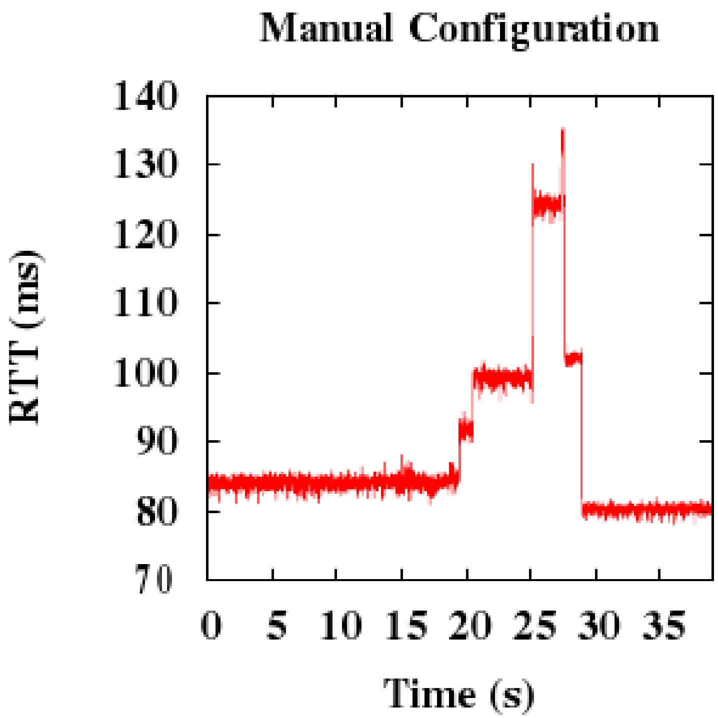

Applying OSPF link-weight changes

- Abilene topology with 3 external peers
  - Configs translated to Quagga syntax
  - Abilene BGP dumps

Evaluation: Commitment

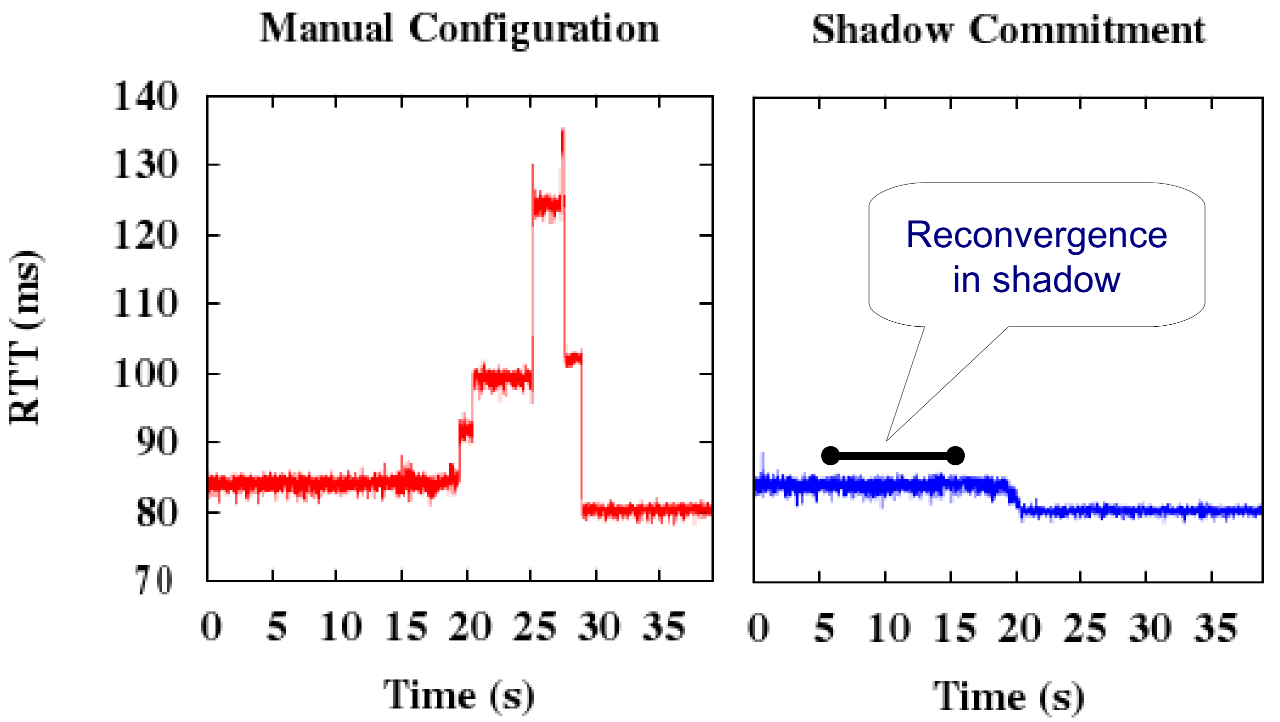

Applying OSPF link-weight changes

- Abilene topology with 3 external peers
  - Configs translated to Quagga syntax
  - Abilene BGP dumps

Conclusion and Future Work

Shadow configurations is new management primitive

- Realistic in-network evaluation
- Network-wide transactional support for configuration

Future work

- Evaluate on carrier-grade installations
- Automated proactive testing
- Automated reactive debugging

Thank you!

Backup Slides

Evaluation: Router Maintenance

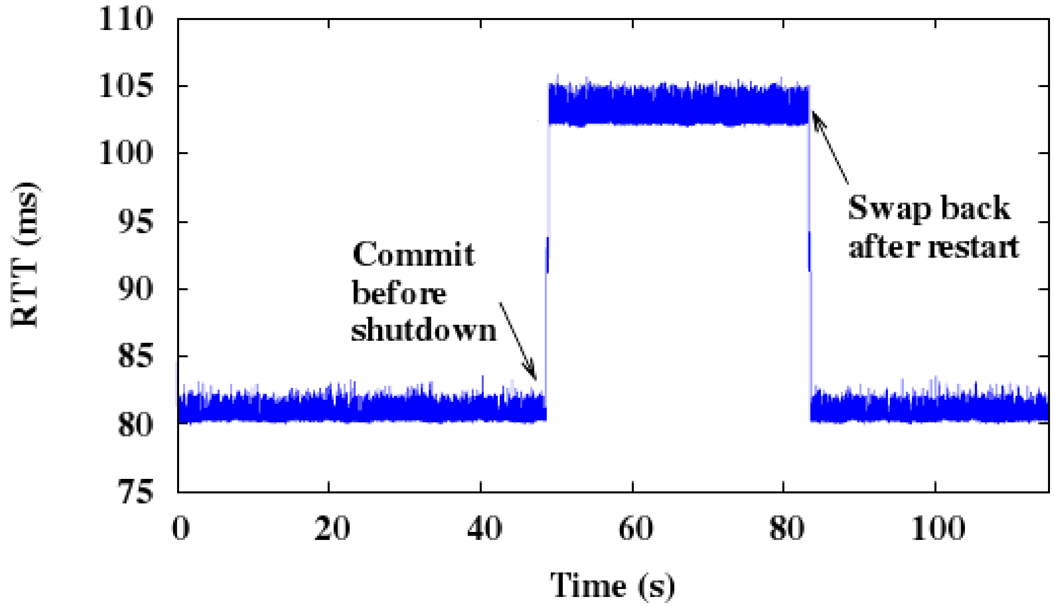

Setup

- Abilene topology with 3 external peers
  - Configs translated to Quagga syntax
  - Abilene BGP dumps