

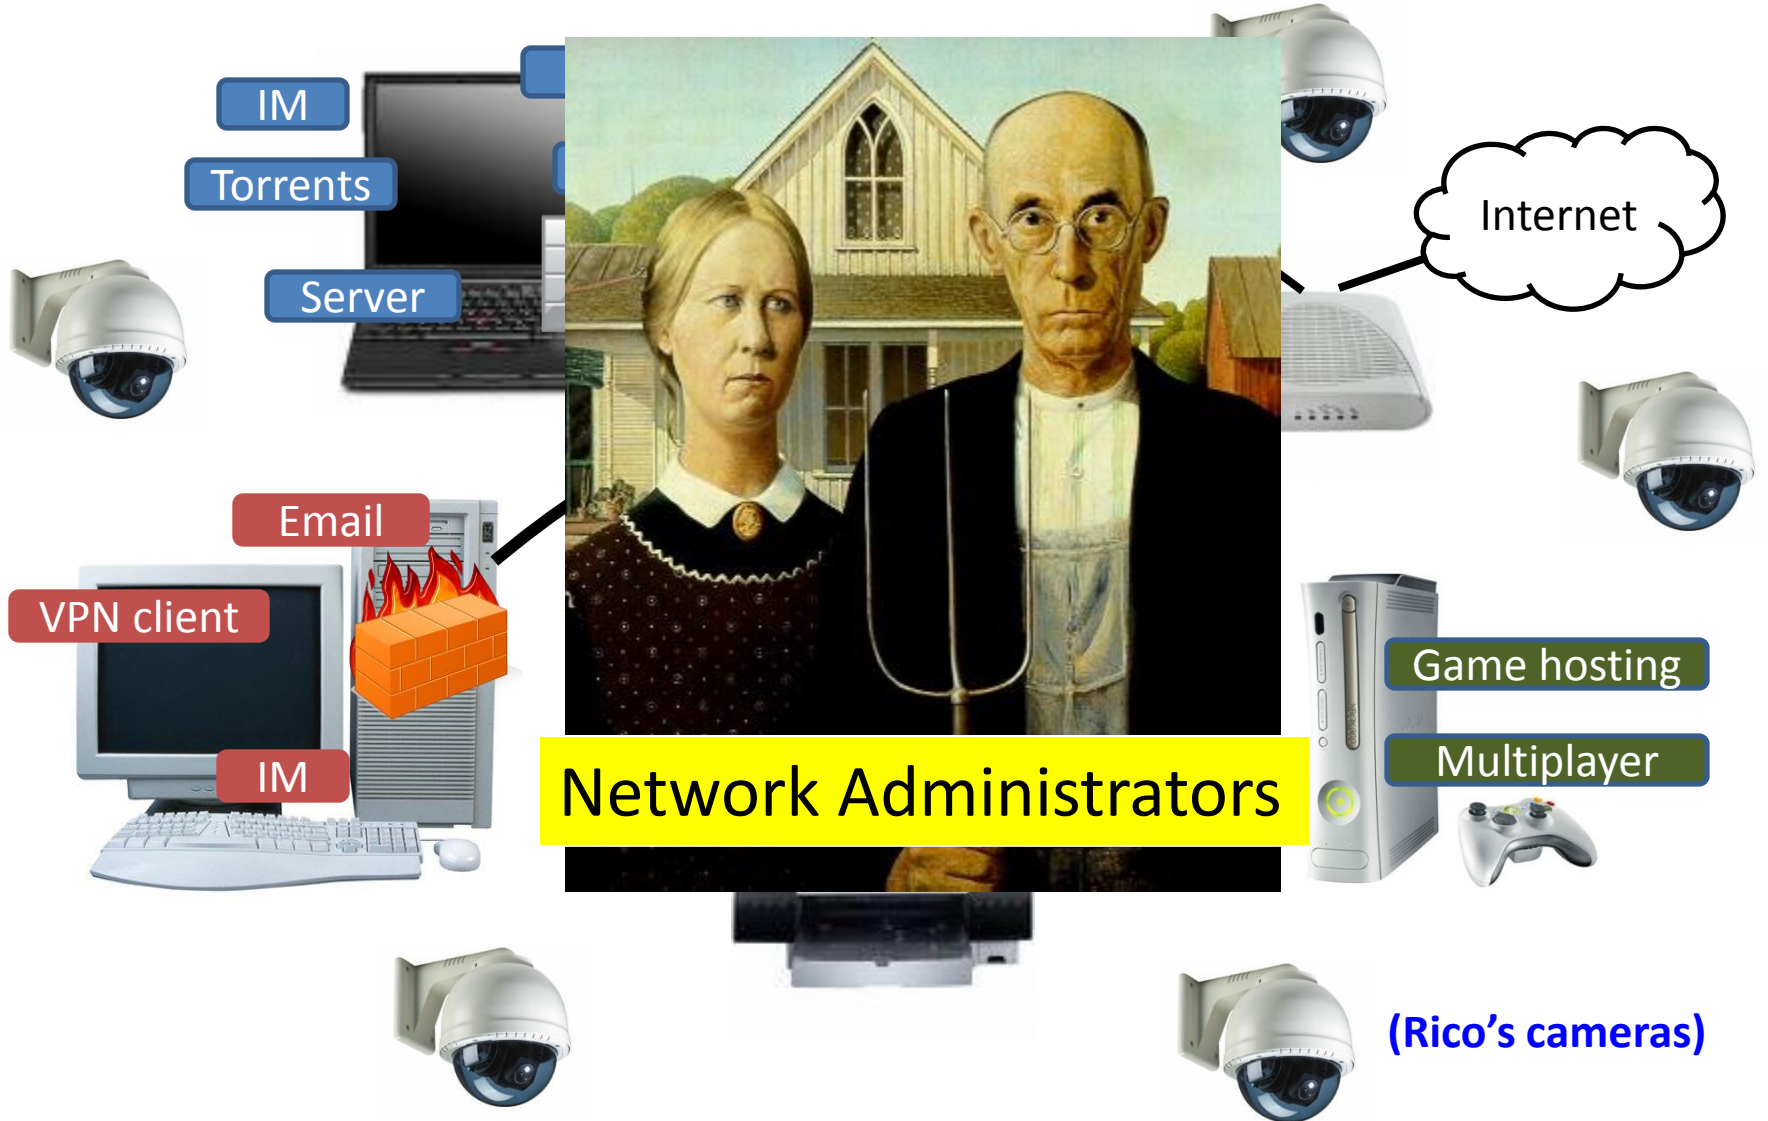
# NetPrints

## Diagnosing Home Network Misconfigurations using Shared Knowledge

*Bhavish Aggarwal, **Ranjita Bhagwan**, Siddharth Eswaran, IIT Delhi*  
*Tathagata Das, Venkat Padmanabhan **Geoffrey M. Voelker**, UCSD*  
**Microsoft Research India**

(Paper appeared in NSDI'09)

# Typical Home Network



# Examples of Problems

Problem	Solution
VPN client does not connect from home	Turn on PPTP passthrough on router, use a subnet that is either 192.168.0.x or 192.168.1.x
XBOX doesn't connect to the Live service	Turn up your MTU above 1300, NAT settings to full-cone, turn on UPnP
My IM client doesn't work from home	Turn off the DNS proxy on the router
File sharing doesn't seem to work at home	Make sure you and the file server are on the same domain/workgroup
Printing doesn't work from my laptop	Turn on correct firewall rules on server machine
Cannot send large emails	Lower MTU on your router

Router misconfig

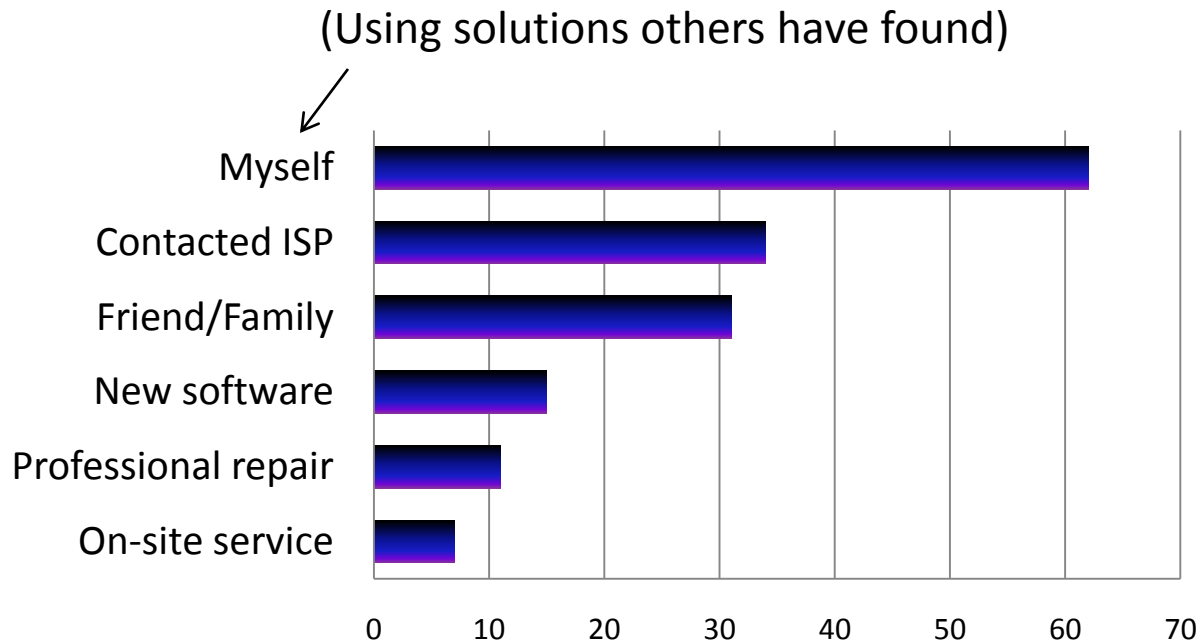
End-host misconfig

Remote problem, local changes

Bottom Line: Amazing anything works at all



# What Do Users Do Today?



Source: *Managing the Digital Home*, a survey of 6,116 U.S. and Canadian home Internet users  
© 2007 Parks Associates

**Avg time to resolve solutions: 2 hours**

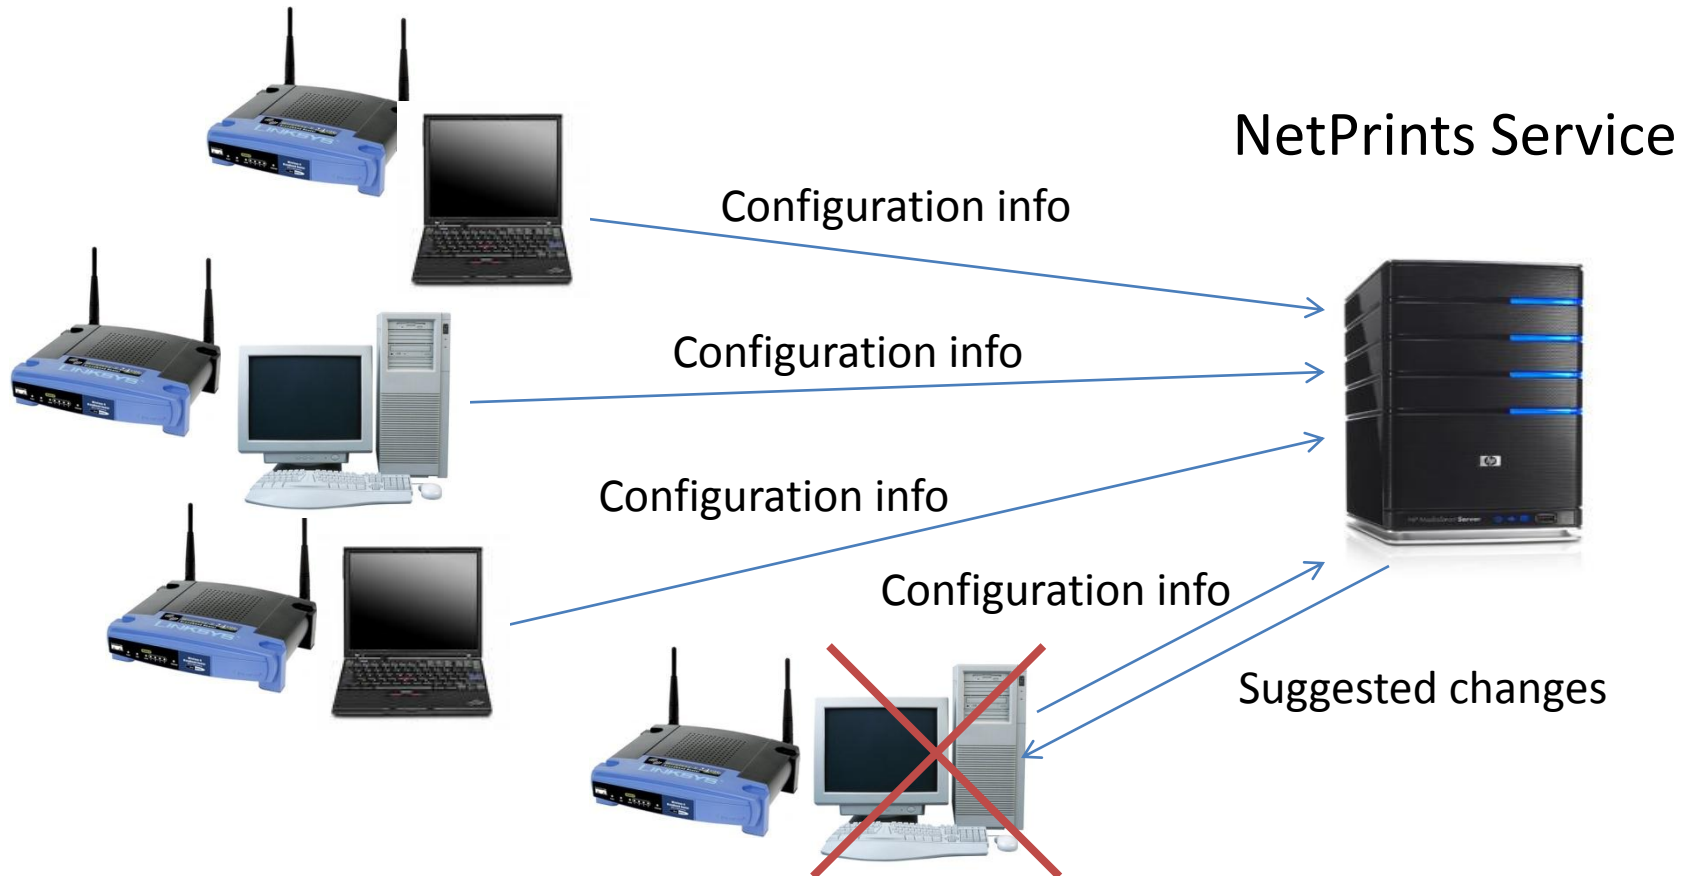
THE TRUTH IS OUT THERE

The image features a dark, atmospheric landscape. In the foreground, a large, dark mountain slope descends from the right side towards the center. The background consists of rolling hills and mountains under a cloudy, overcast sky. The overall color palette is dominated by dark blues, greys, and blacks, creating a somber and mysterious mood. The text "THE TRUTH IS OUT THERE" is centered horizontally across the upper portion of the image.

# NetPrints

NetPrints = Network Problem Fingerprinting

Automate problem diagnosis using **shared knowledge**





# Assumptions

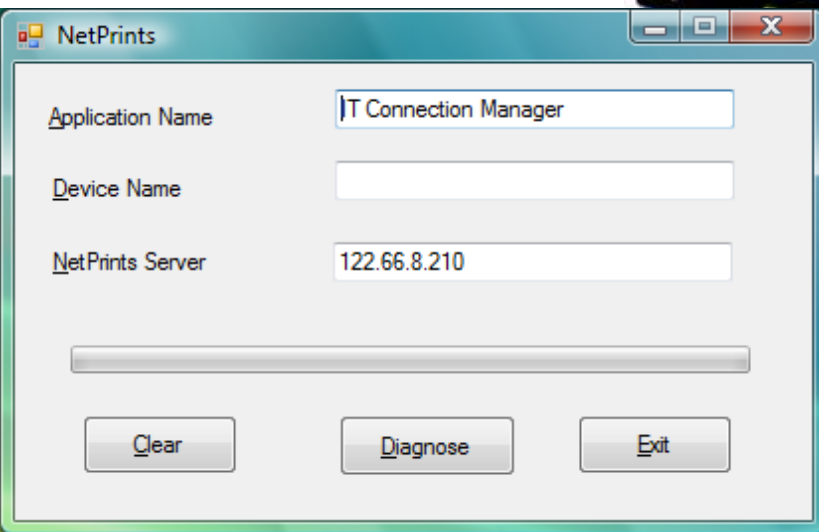
- Current design requires basic connectivity
  - Looking at application-specific problems
  - Not inherent, Knowledgebase can be shipped offline
- Not dealing with performance
  - “good” (working) and “bad” (not working) are the only two states considered



# NetPrints in Action



```
config.xml
... ptp_pass=1
... ptp_pass=0
...
```



Knowledgebase for VPN client



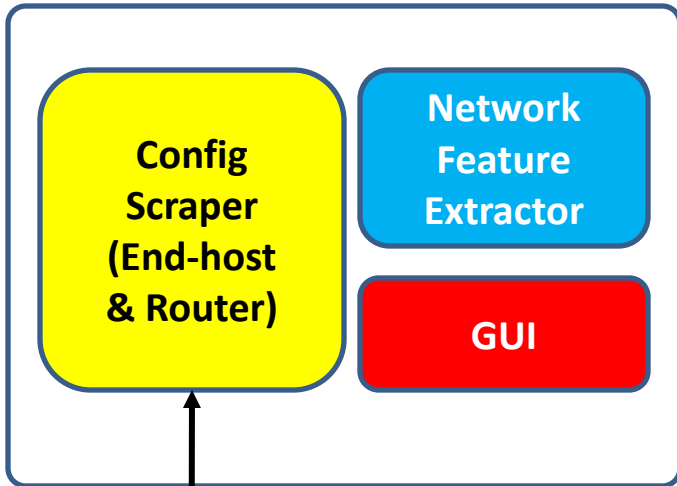
NetPrints server

# Diagnosis Strategies

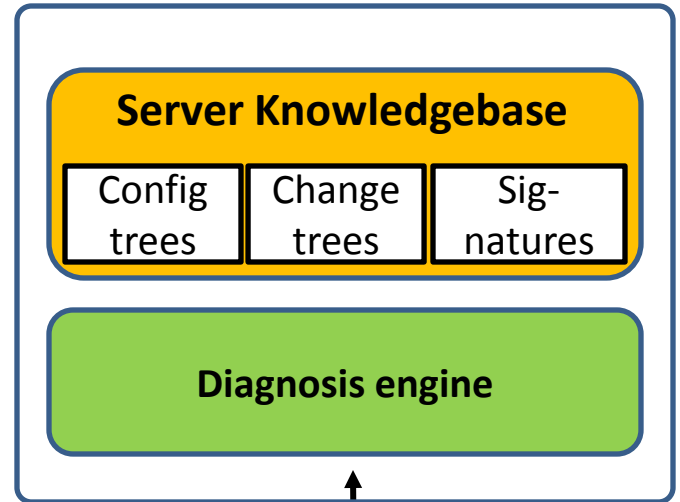
- Snapshot-based
  - Collect config snapshots from different users
- Change-based
  - Collect config changes that a user makes
- Symptom-based
  - Collect problem signatures from network traffic

# System Design


## NetPrints Client



## NetPrints Server




Local-Area Network

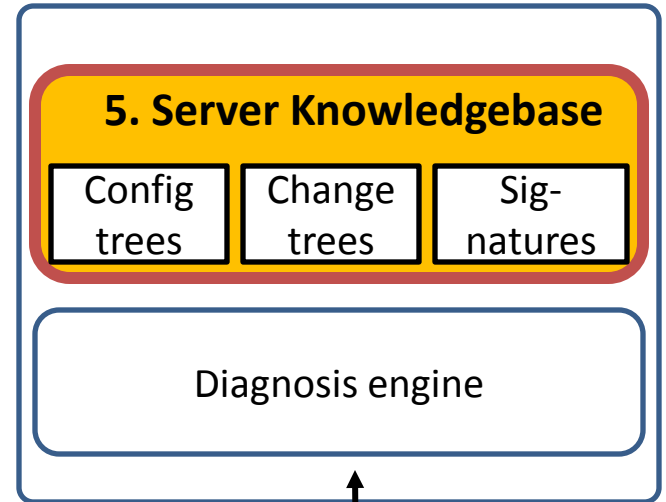


# Normal Mode

## NetPrints Client

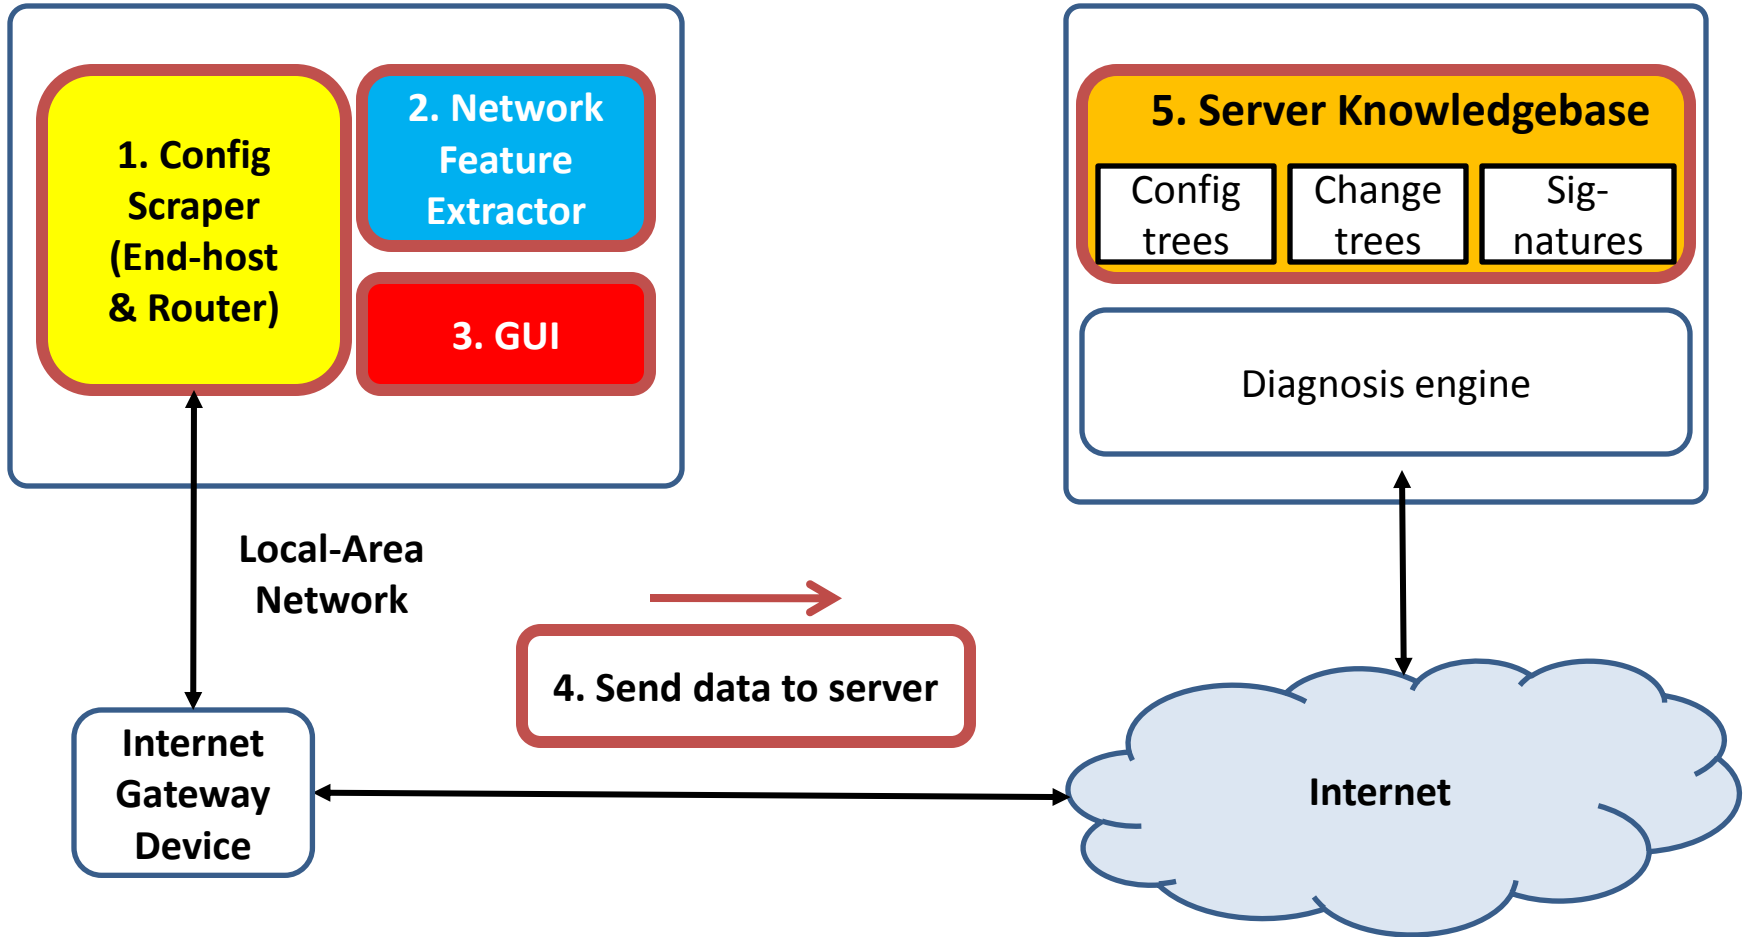



## NetPrints Server



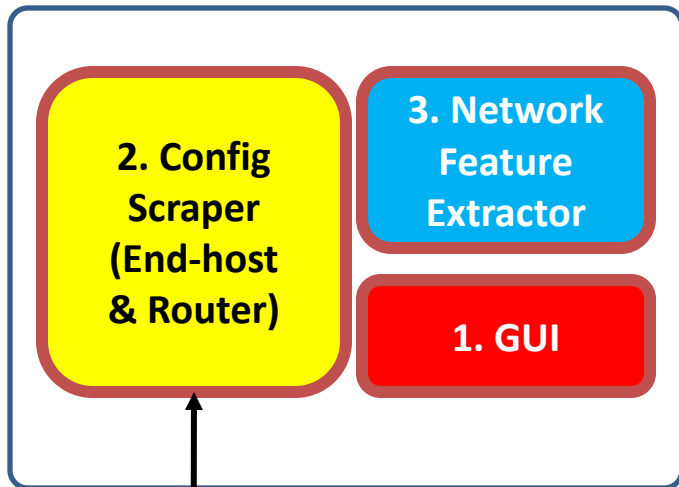
Local-Area Network

4. Send data to server


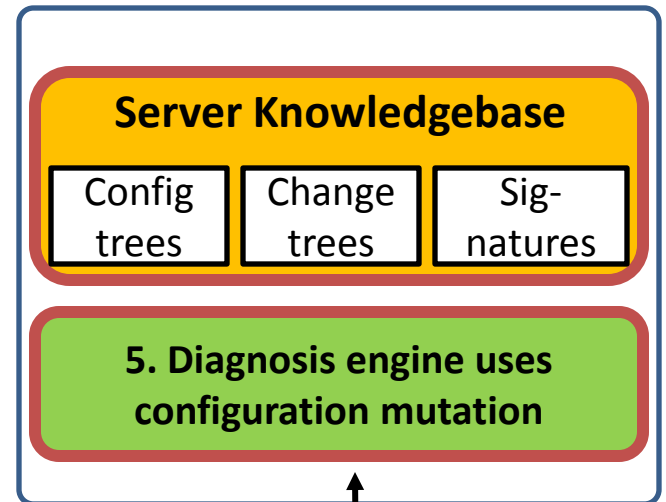


# Diagnose Mode

## NetPrints Client



## NetPrints Server



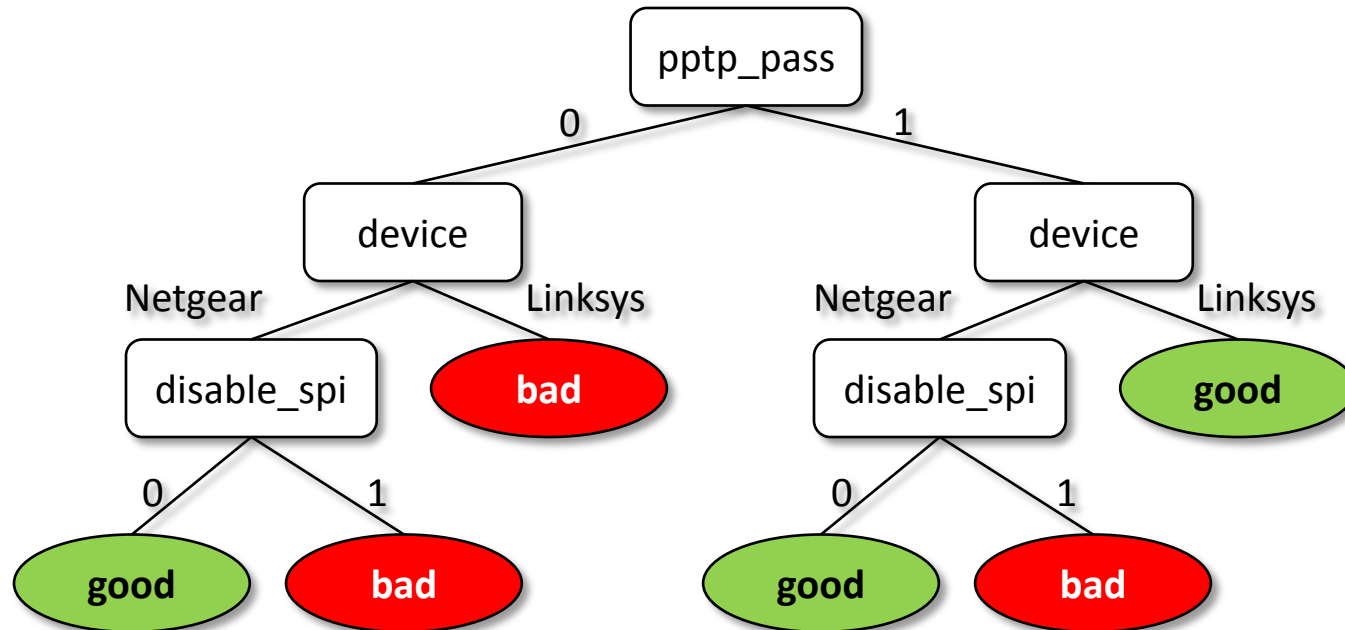
# #1: Configuration Scraper

- Router scraper
  - UPnP
  - Web Interface (HTTP Request Hijacking)
- End-host scraper
  - Interface-specific parameters
  - Patches and software versions
  - Firewall rules
- Remote scraper
  - Composition of local and remote configs

## #2: Server Knowledgebase

- Per-application decision trees constructed using labeled configuration snapshots
  - Decision trees aid interpretability
  - C4.5 decision tree learning algorithm

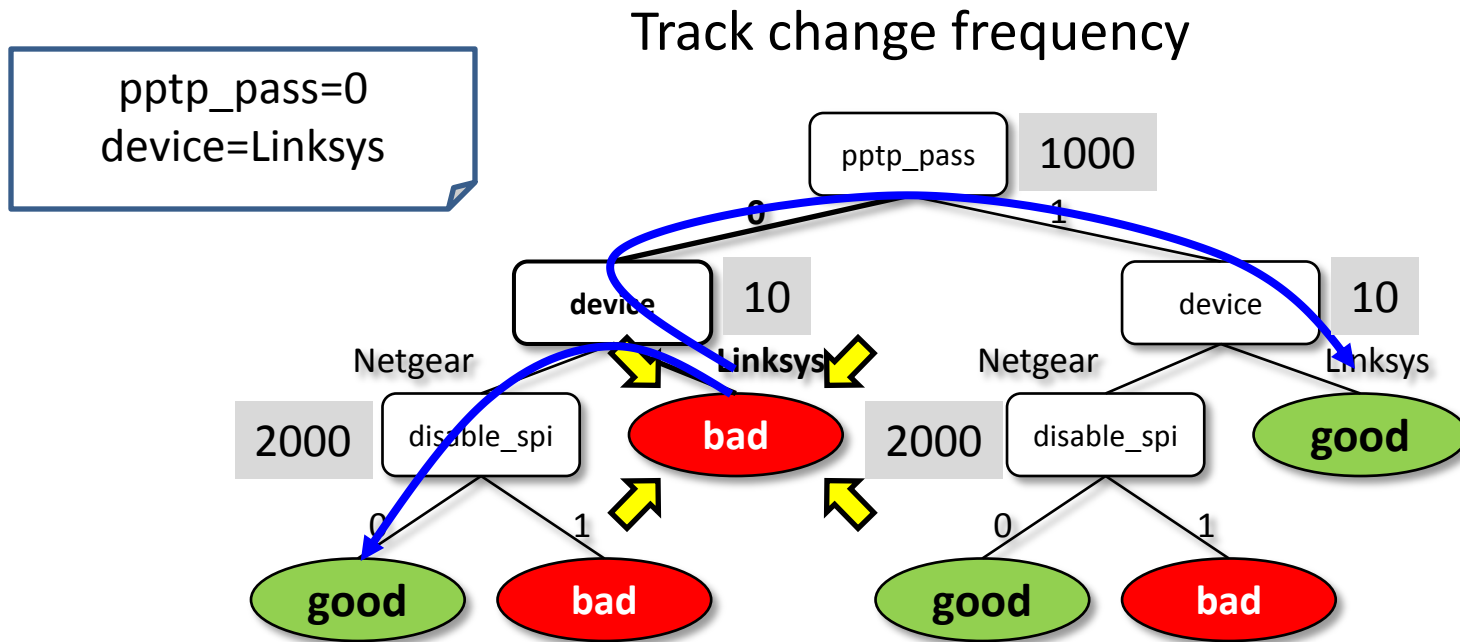
# Example of Configuration Tree



**Simplified Configuration Tree for VPN Client**



# #3: Configuration Mutation



- Preference for mutations involving frequently changing parameters
- Assumption: higher the frequency, less disruptive the change

# Shortcoming of Configuration Trees

- Some config info may not be learned
  - Traversal of config tree may end in a “good” leaf even if config is problematic
- Reasons
  - Insufficient data
    - e.g., a new router enters the market
  - Hidden configurations
    - e.g., application-specific parameters
- Use **Network Signatures + Change Trees**
  - See paper for full details

# Experimental Evaluation


- Testbed
  - 7 routers: Netgear, Linksys, D-Link, Belkin
  - 4 applications: VPN, FTP, file sharing, Xbox Live
  - Various clients hosts: in lab, over DSL, an Xbox
- Workload
  - Generate many different configurations
  - Reproduce real problems (from Web, our own)
  - Parameter sweep for 11 config fields commonly found in problem reports
- Robustness Experiments
  - Mislabeling (accidental or malicious)
  - Parameter skew (some fields changed frequently)

(Looks like a home network, yes?)



# Findings

- Intuitive inferences
  - VPN: **If pptp\_pass==1 then GOOD**
- Surprising inferences
  - VPN: **If stateful==off and pptp\_pass==0 and ipsec\_pass==0 and l2tp\_pass==0 then GOOD**
  - File sharing: **If Guest ∈ ACL and User ∉ Local**  
(guests can access files, but not local user)
- Robust to mislabeling and skew
- Full results in paper...



# Summary

- Home network diagnostics is challenging
  - diversity of apps and configs
  - absence of an admin
- NetPrints leverages community info to perform *automated* diagnosis
  - decision tree based learning
  - configuration trees, network traffic signatures and change trees
- (Visit MSR India, it's a great place!)

<http://research.microsoft.com/netprints>

# Putting NetPrints in Context

Rule-based techniques, Learning-based

Windows Diagnostics Framework  
Network Magic  
Apple's Diagnostics

Strider+PeerPressure  
Autobash  
SVM-based performance debugger

*Resolve basic connectivity issues*  
*Resolves local configuration issues*  
*(Application specific: too many rules)*

## NetPrints

- Distributed configuration information
- Unstructured, heterogeneous environment
- Problems due to interaction of multiple configurations



# Composing Local & Remote Configs

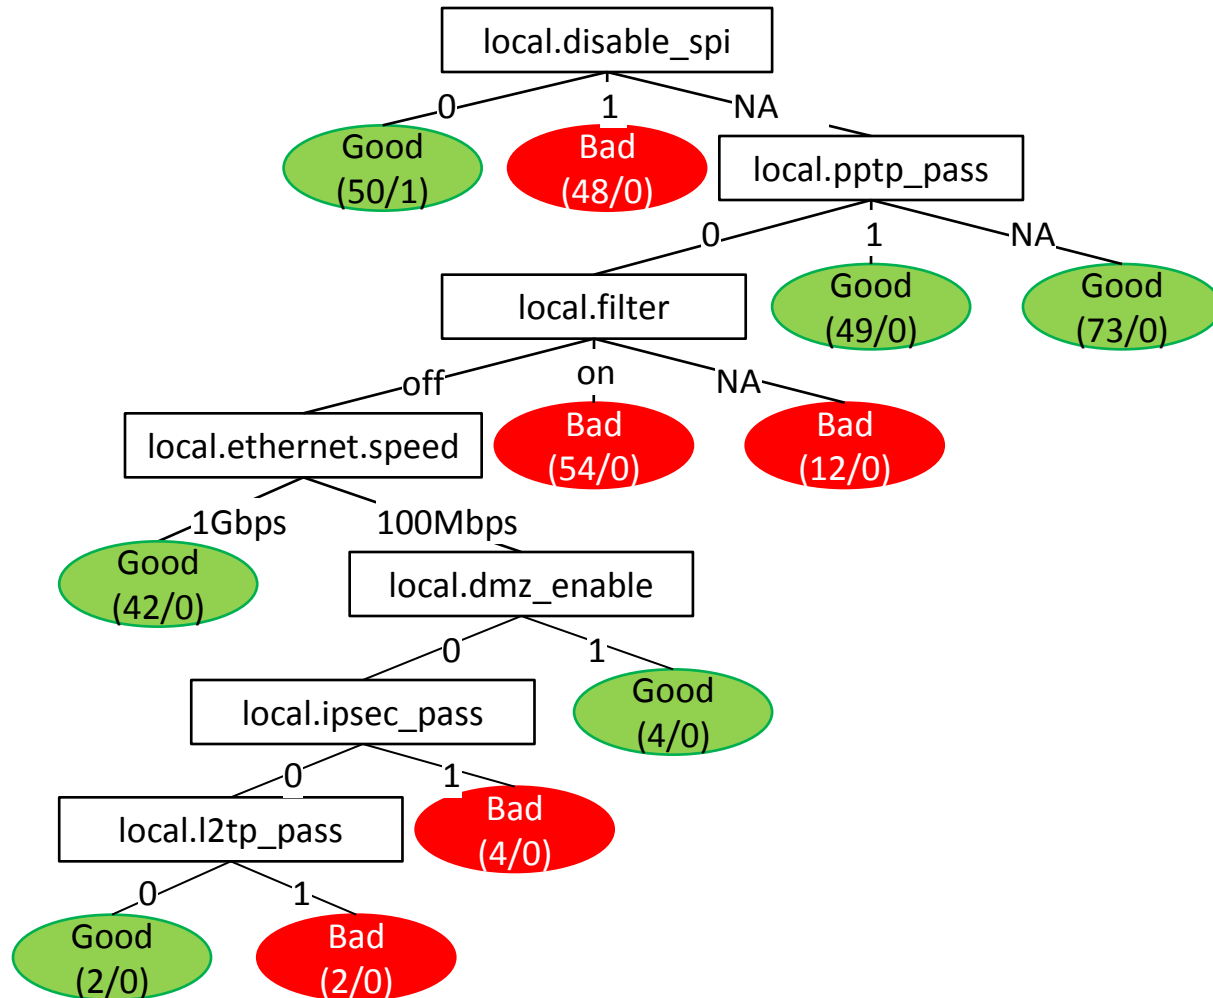
Problem	Solution
VPN client does not connect from home	Turn on PPTP passthrough on router, use a subnet that is either 192.168.0.x or 192.168.1.x
XBOX doesn't connect to the Live service	Turn up your MTU above 1365, change NAT settings to full-cone, turn on UPnP
My IM client doesn't work from home	Turn off the DNS proxy on the router
File sharing doesn't seem to work at home	Make sure client and the server are on the same domain/workgroup.
Printing doesn't work from my laptop	Turn on correct firewall rules on print server machine
Cannot send large emails	Turn down MTU on your router

Sometimes it is the *combination* of local and remote configs that is the problem


# Methodology



- Testbed comprising 7 different routers
  - various makes: Netgear, Linksys, D-Link, Belkin
- Clients running the VPN sent configurations to the NetPrints service
  - Roughly 6000 config parameters per snapshot
- Learned configuration trees w/ C4.5 algorithm

# Configuration Tree for VPN Client




# Summary of Diagnosis Procedure



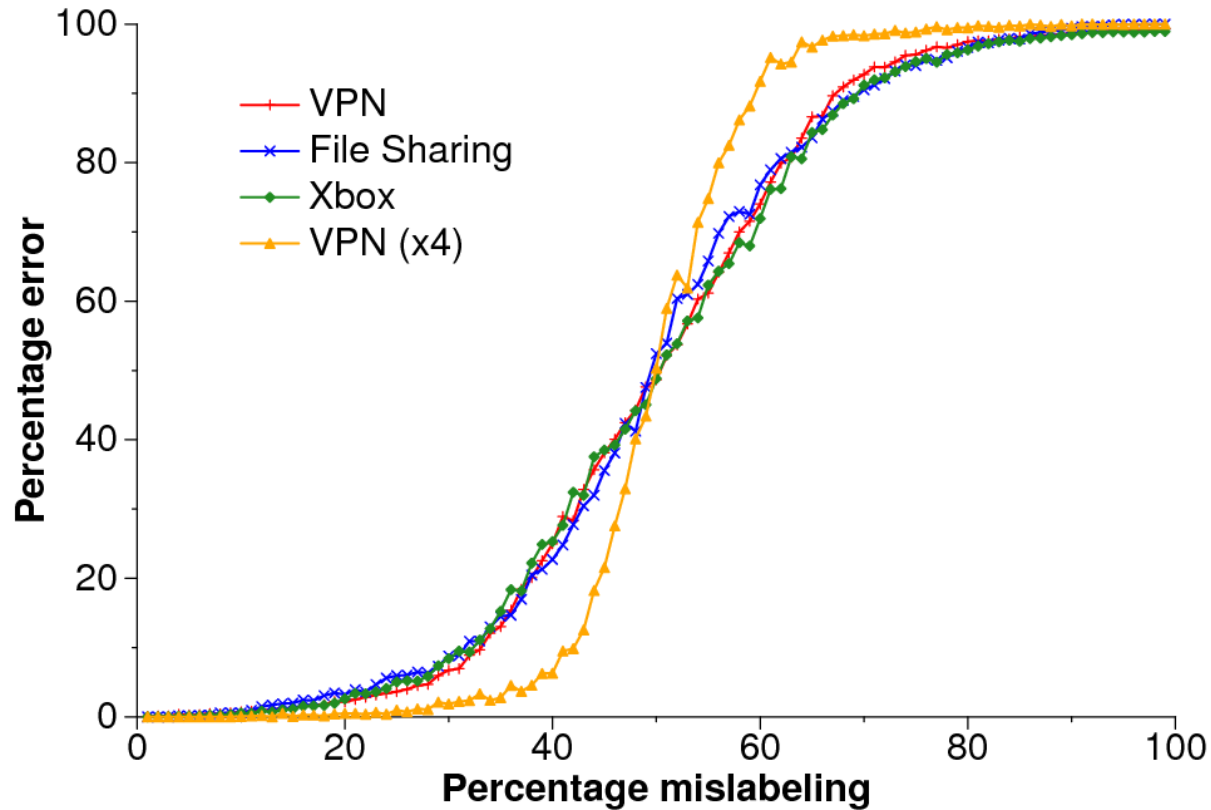
Network traffic signature	Change trees
1XXXXXX	
0XXX X1X	

# Tolerance to Mislabeling



13-17% mislabeling  $\Rightarrow$  1% error in diagnosis

# Tolerance to Mislabeling



13-17% mislabeling  $\Rightarrow$  1% error in diagnosis

# Experimental Evaluation

- Testbed comprising 7 different routers
  - various makes: Netgear, Linksys, D-Link, Belkin

