Syslog Processing for Switch Failure Diagnosis and Prediction in Datacenter Networks

Shenglin Zhang, Weibin Meng, Jiahao Bu, Sen Yang
Dan Pei, Ying Liu, Jun (Jim) Xu, Yu Chen, Hui Dong, Xianping Qu, Lei Song

Network Devices in Data Center Networks

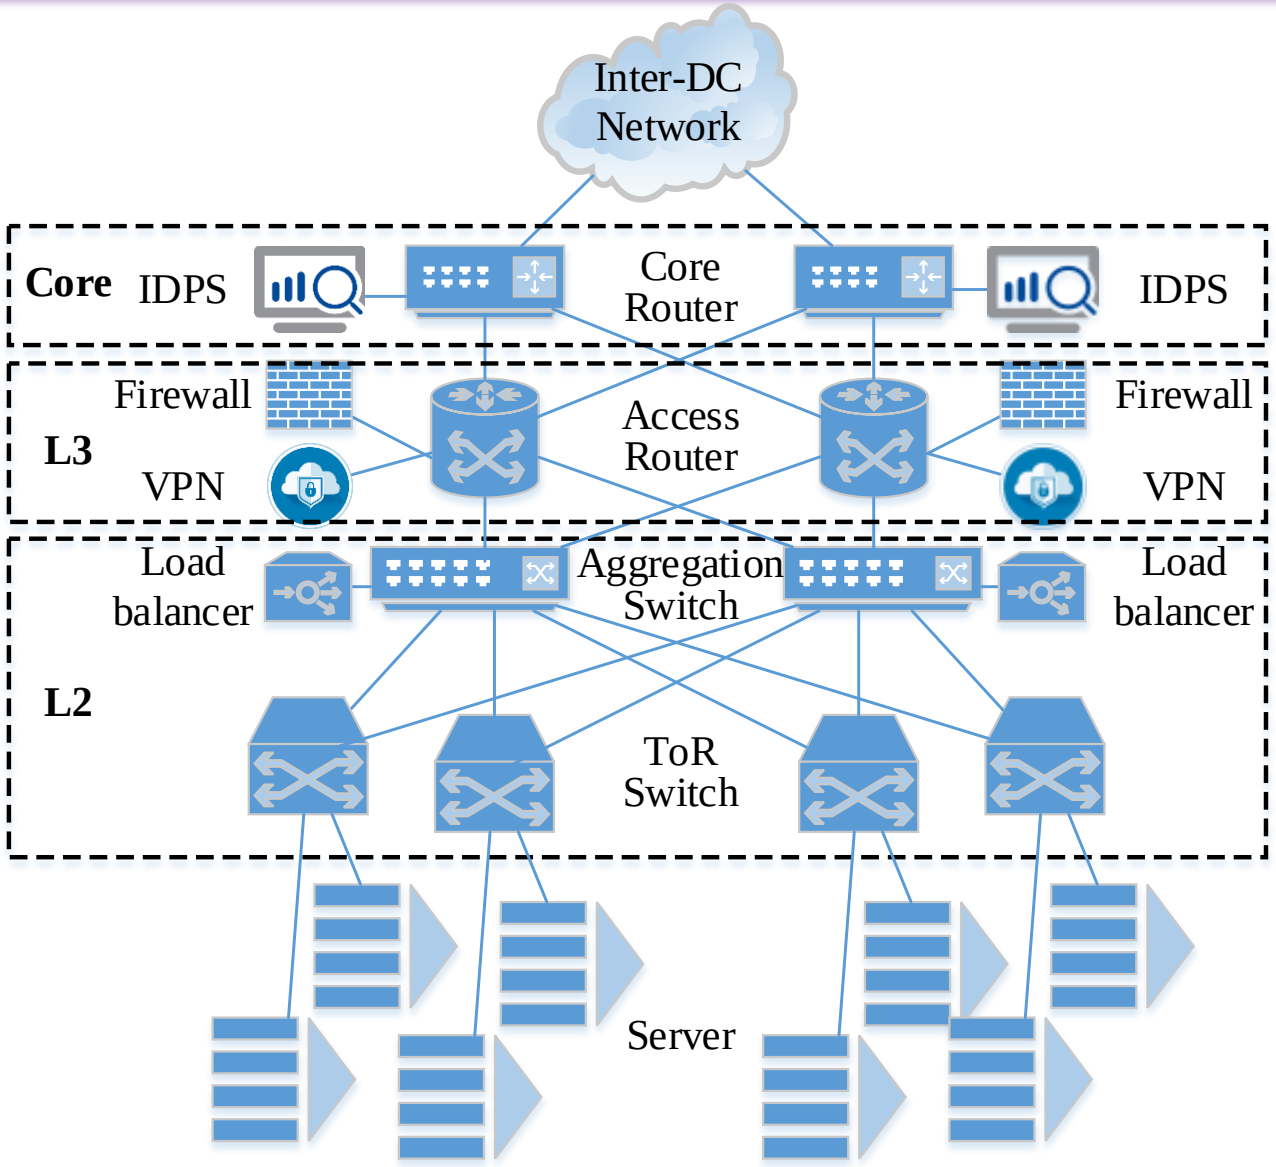

Network Devices in Data Center Networks

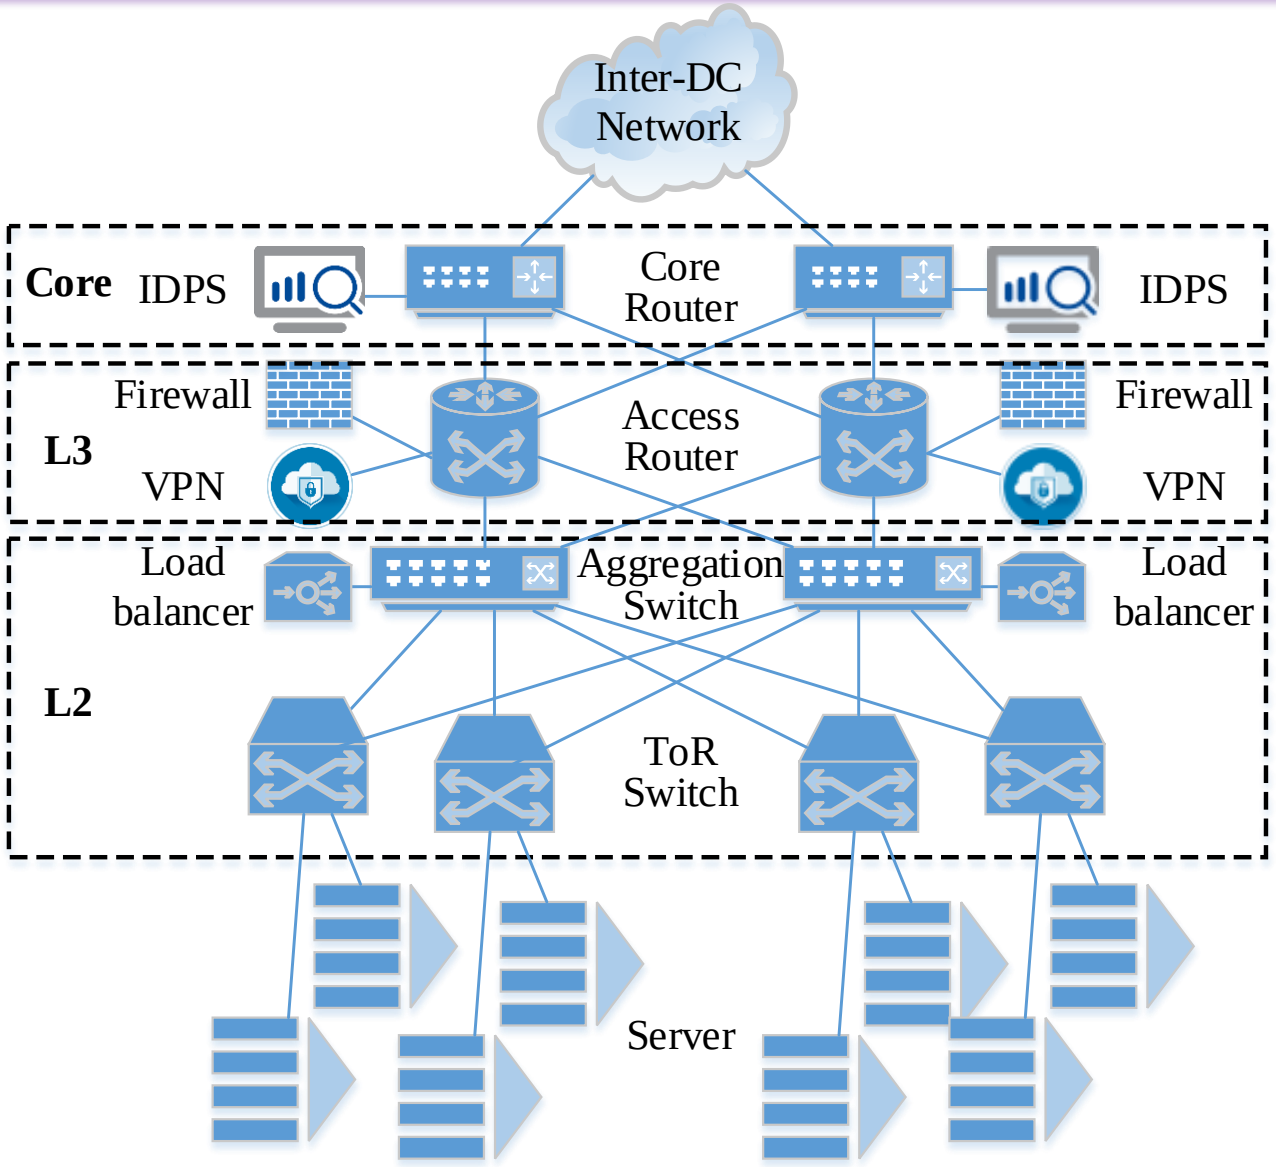

- Switch
  - Top-of-rack switch
  - Aggregation switch
- Router
  - Access router
  - Core router
- Middle box
  - Firewall
  - Intrusion detection and prevention system (IDPS)
  - Load balancer
  - VPN

Network Devices in Data Center Networks

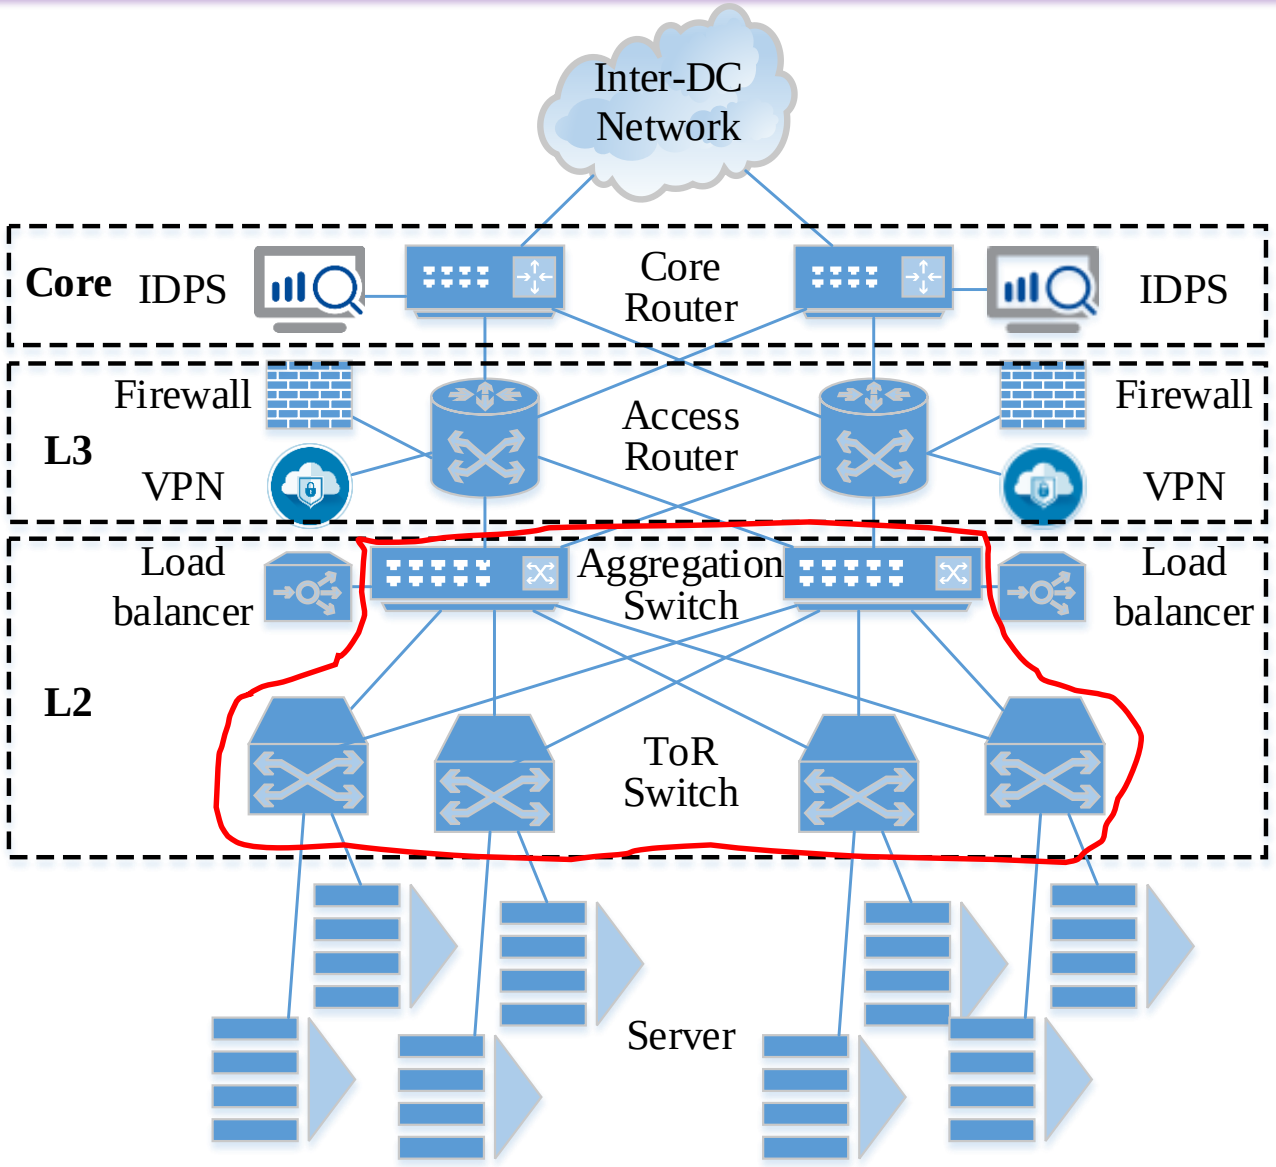

- Switch
  - Top-of-rack switch
  - Aggregation switch
- Router
  - Access router
  - Core router
- Middle box
  - Firewall
  - Intrusion detection and prevention system (IDPS)
  - Load balancer
  - VPN

Scale of Network Devices in Datacenter

Microsoft (C. Guo, et al., SIGCOMM'15)

- Hundreds of thousands to millions of servers
- Hundreds of thousands of switches
- Millions of cables and fibers

Scale of Network Devices in Datacenter

Microsoft (C. Guo, et al., SIGCOMM'15)	• Hundreds of thousands to millions of servers • Hundreds of thousands of switches • Millions of cables and fibers
Baidu	• Hundreds of thousands of servers • Tens of thousands of switches

Scale of Network Devices in Datacenter

Microsoft (C. Guo, et al., SIGCOMM'15)	• Hundreds of thousands to millions of servers • Hundreds of thousands of switches • Millions of cables and fibers
Baidu	• Hundreds of thousands of servers • Tens of thousands of switches
Swich failures are the norm rather than the exception (P. Gill, et al., SIGCOMM'11)	• More than 400 switch failures per year

Switch Failures Lead to Outages

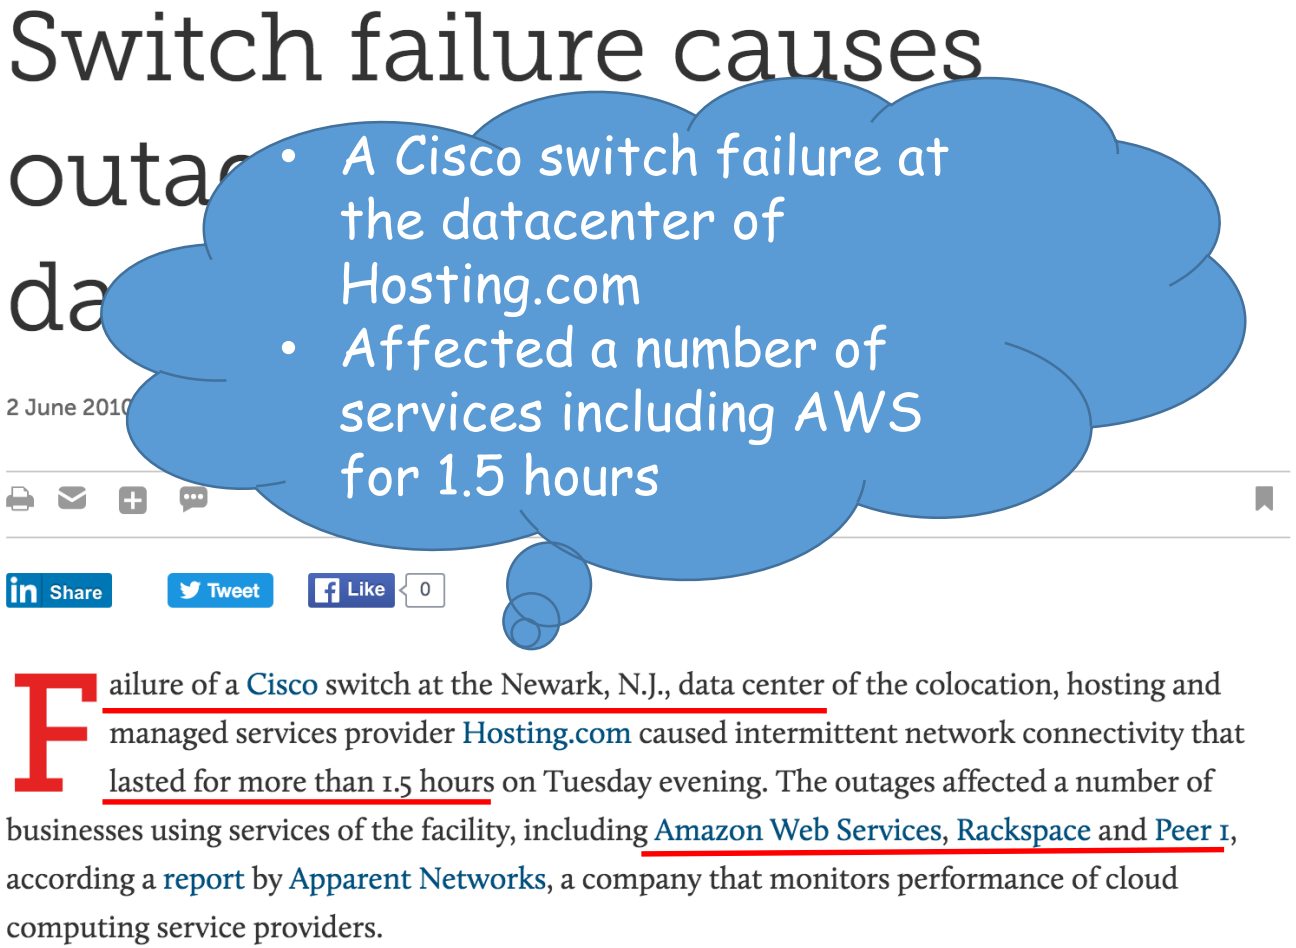

Switch failure causes outa
da

2 June 2010

Share Tweet Like 0

Failure of a Cisco switch at the Newark, N.J., data center of the colocation, hosting and managed services provider Hosting.com caused intermittent network connectivity that lasted for more than 1.5 hours on Tuesday evening. The outages affected a number of businesses using services of the facility, including Amazon Web Services, Rackspace and Peer 1, according a report by Apparent Networks, a company that monitors performance of cloud computing service providers.

- A Cisco switch failure at the datacenter of Hosting.com
- Affected a number of services including AWS for 1.5 hours

Switch Failures Lead to Outages

Switch failure causes outa... da...

2 June 2010

- A Cisco switch failure at the datacenter of Hosting.com
- Affected a number of services including AWS for 1.5 hours

Share Tweet Like 0

Failure of a Cisco switch at the Newark, N.J., data center of the colocation, hosting and managed services provider Hosting.com caused intermittent network connectivity that lasted for more than 1.5 hours on Tuesday evening. The outages affected a number of businesses using services of the facility, including Amazon Web Services, Rackspace and Peer 1, according a report by Apparent Networks, a company that monitors performance of cloud computing service providers.

Switch failure shuts down computer network at data center

AP By The Associated Press
May 24, 2016 8:49 am

CHESTER, Va. (AP) — The computer network of a data center in Chester went dark after a switch failure.

The Richmond Times-Dispatch (http://bit.ly/20v8U5T) reports that Saturday's outage at the Commonwealth Enterprise Solutions Center affected access to the network by almost every executive branch agency the center serves, including the Department of Motor Vehicles.

Email, cellphones and agency computer servers in the center went dark, causing outage for inbound and outbound ... the DMV ...

V...

- The datacenter network went dark after a switch failure
- Almost every executive branch agency are affected for a few hours

Switch Failure Diagnosis and Proactive Detection

Frameworks

- SyslogDigest (IMC 2010)
- Spatio-temporal Factorization (INFOCOM 2014)
- Proactive Failure Detection (CNSM 2015)

Based on analyzing syslogs

Syslog Structure

Switch ID	Message timestamp	Message type	Detailed message
Switch 1	Jun 12 19:03:03 2014	SIF	Interface te-1/1/59, changed state to down
Switch 2	Jul 15 11:05:07 2015	OSPF	Neighbour(rid:10.231.0.43, addr:10.231.39.61) on vlan23, changed state from Exchange to Loading
Switch 3	Jan 12 21:03:01 2016	%%SLOT	SFP te-1/1/33 is plugged in, vendor: BROCADE, serial number: AAA210383148232

The detailed message field

Describe events occurring on switches

- Interface up/down
- Plug in/out of slot
- DDoS attack
- Operator log in/out

Important to failure diagnosis and proactive detection

Extracting events from the detailed message field

- Pre-processing for failure diagnosis
- Pre-processing for proactive failure detection

Syslog Messages Under the Type "SIF"

1. Interface ae3, changed state to down
2. Vlan-interface vlan22, changed state to down
3. Interface ae3, changed state to up
4. Vlan-interface vlan22, changed state to up
5. Interface ae1, changed state to down
6. Vlan-interface vlan20, changed state to down
7. Interface ae1, changed state to up
8. Vlan-interface vlan20, changed state to up

Syslog Messages Under the Type "SIF" Before A Failure

1. Interface *, changed state to down
2. Vlan-interface *, changed state to down
3. Interface *, changed state to up
4. Vlan-interface *, changed state to up

Common practice for syslog pre-processing:
Extracting templates from syslog messages
Matching syslog messages to templates

Syslog Messages Under the Type "SIF" Before A Failure

1. Interface *, changed state to down
2. Vlan-interface *, changed state to down
3. Interface *, changed state to up
4. Vlan-interface *, changed state to up

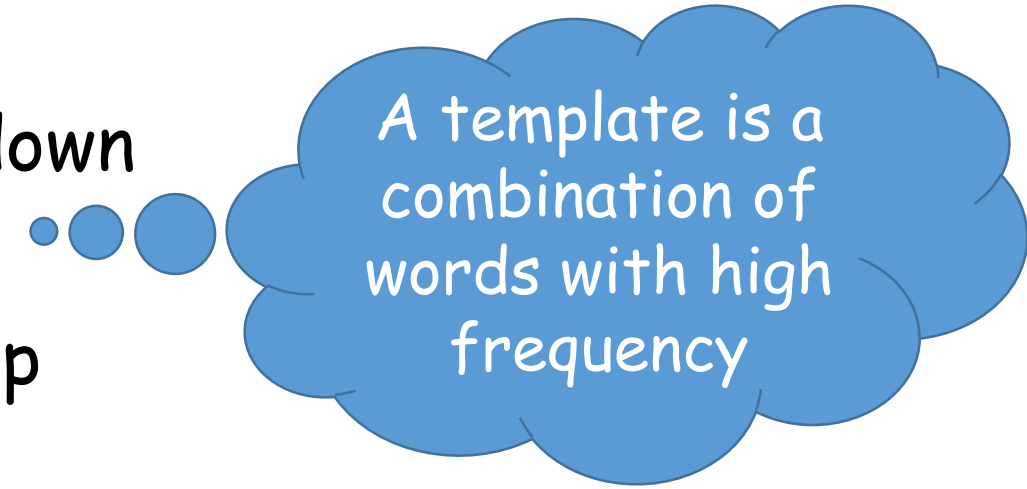

Common practice for syslog pre-processing:
Extracting templates from syslog messages
Matching syslog messages to templates

Outline

- Background and Motivation
- Challenges
- Key Ideas
- Results
- Conclusion

Challenges

Unstructured texts

Huge amount of syslog messages

- Tens of millions everyday
- Long period of historical data for training (two years)

Diverse types of syslog messages

- Operator log in/out
- Interface up/down
- Plug in/out of slot

Challenges

Unstructured texts

Huge amount of syslog messages

- Tens of millions everyday
- Long period of historical data for training (two years)

Diverse types of syslog messages

- Operator log in/out
- Interface up/down
- Plug in/out of slot

Templates should be updated periodically

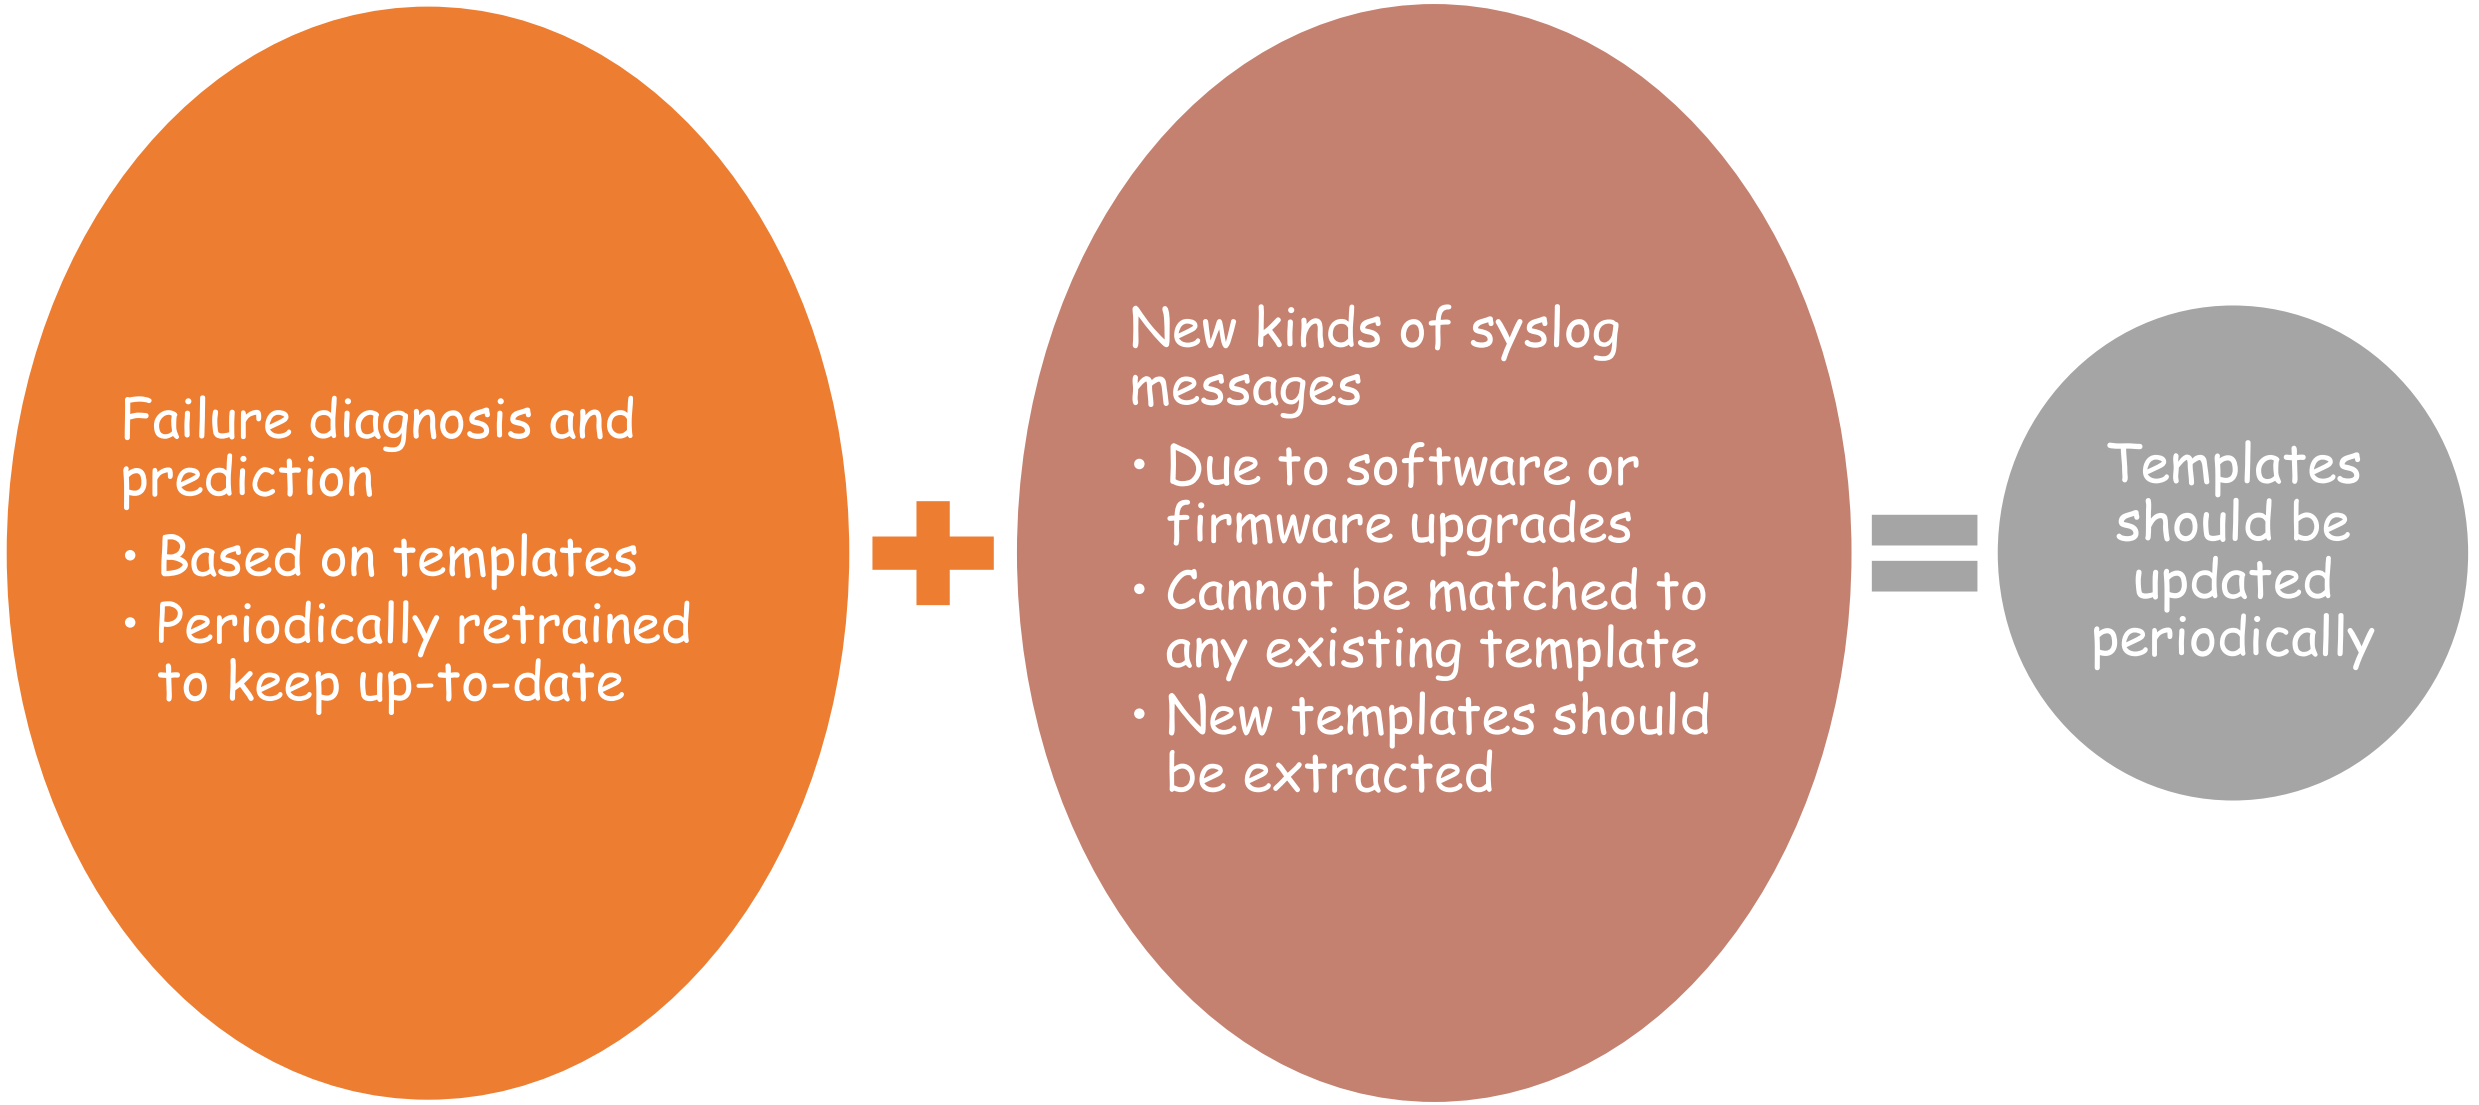

Incrementally re-trainable

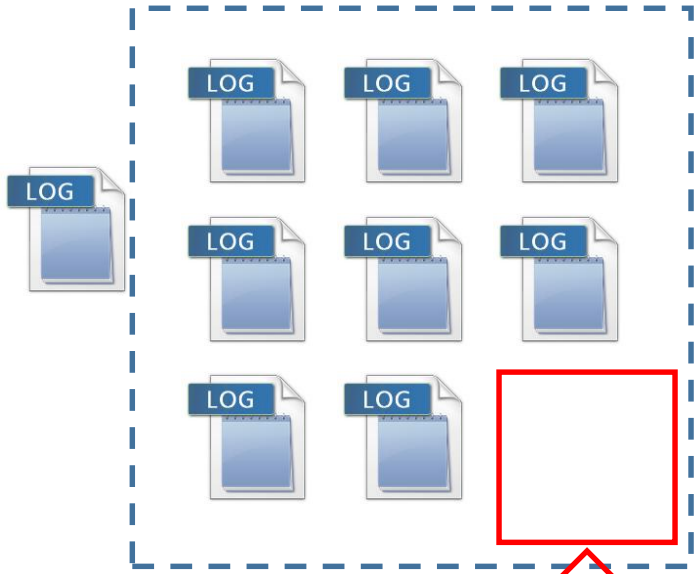

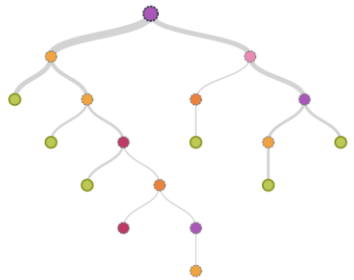

Not incrementally re-trainable

Template extraction method

Computationally efficient

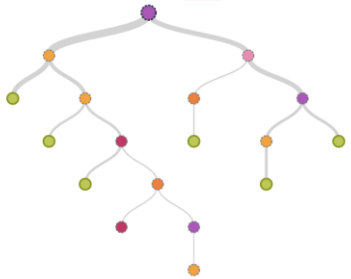

Incrementally re-trainable

Template extraction method

Existing template extraction methods

Method	Conference	Merits	Drawbacks
Signature Tree	IMC 10	Accurate	Not incrementally re-trainable
STE	INFOCOM 14	None	Inaccurate and not incrementally re-trainable
LogSimilarity	CNSM 15	Learn incrementally	Inaccurate

Our goal

Accurate, incrementally re-trainable, efficient template extraction method

Outline

- Background and Motivation
- Challenges
- Key Ideas
- Results
- Conclusion

Construct FT-tree

- Support: if a word W appears in some message, (the support of W) ++

Construct FT-tree

- Support: if a word W appears in some message, (the support of W) ++
- Scan all the messages, order all of the words into a map M in the descending order of support

Construct FT-tree

- Support: if a word W appears in some message, (the support of W) ++
- Scan all the messages, order all of the words into a map M in the descending order of support

M

Words	Support
“changed”, “state”, “to”	8
“Interface”, “Vlan-interface”, “up”, “down”	4
“vlan20”, “vlan22”, “ae1”, “ae3”	2

1. Interface ae3, changed state to down
2. Vlan-interface vlan22, changed state to down
3. Interface ae3, changed state to up
4. Vlan-interface vlan22, changed state to up
5. Interface ae1, changed state to down
6. Vlan-interface vlan20, changed state to down
7. Interface ae1, changed state to up
8. Vlan-interface vlan20, changed state to up

Construct FT-tree

- Order words in each message in the descending order of support
  - Interface ae3, changed state to down
  - ➢V1 = {"changed", "state", "to", "Interface", "down", "ae3"}
  - Vlan-interface vlan22, changed state to down
  - ➢V2 = {"changed", "state", "to", "Vlan-interface", "down", "vlan22"}
  - Interface ae3, changed state to up
  - ➢V3 = {"changed", "state", "to", "Interface", "up", "ae3" }
  - Vlan-interface vlan22, changed state to up
  - ➢V4 = {"changed", "state", "to", "Vlan-interface", "up", "vlan22"}
  - …

M

Words	Support
"changed", "state", "to"	8
"Interface", "Vlan-interface", "up", "down"	4
"vlan20", "vlan22", "ae1", "ae3"	2

Construct FT-tree

SIF

Construct FT-tree

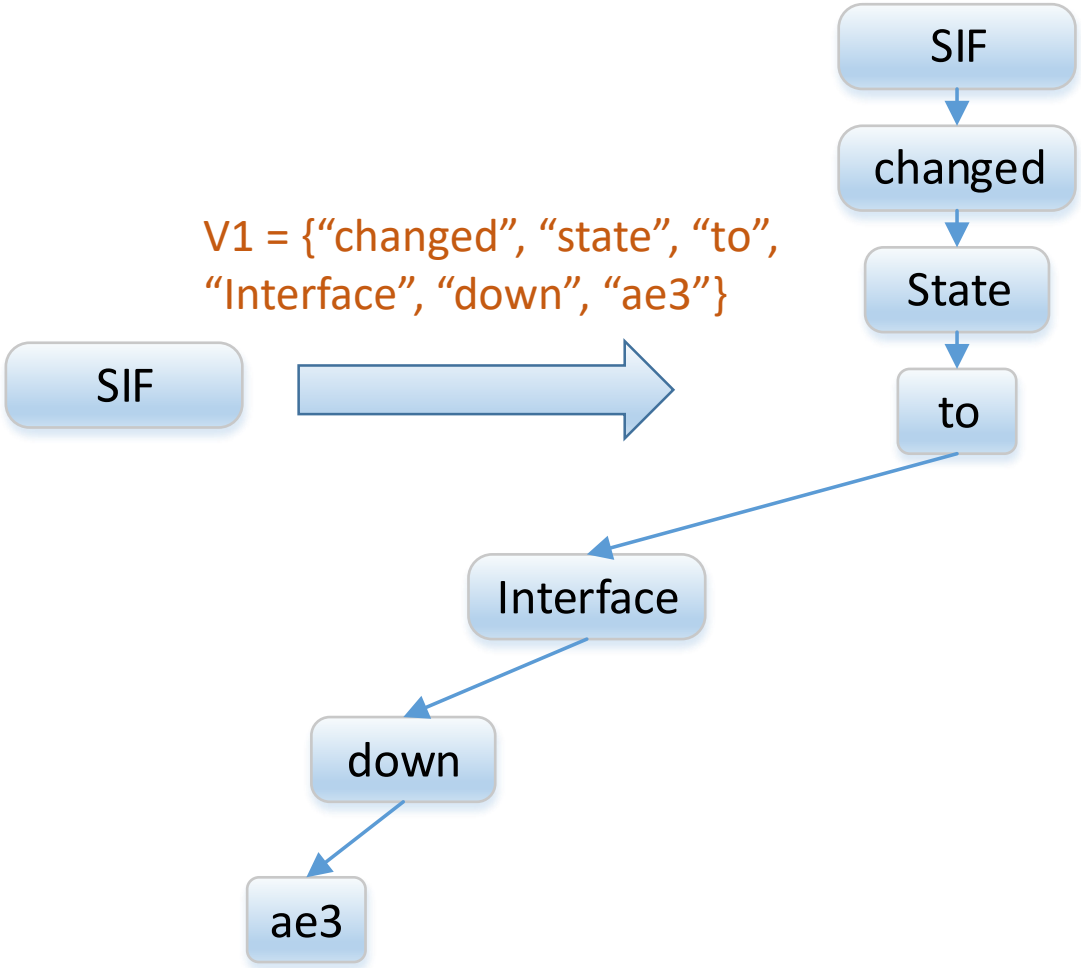

Construct FT-tree

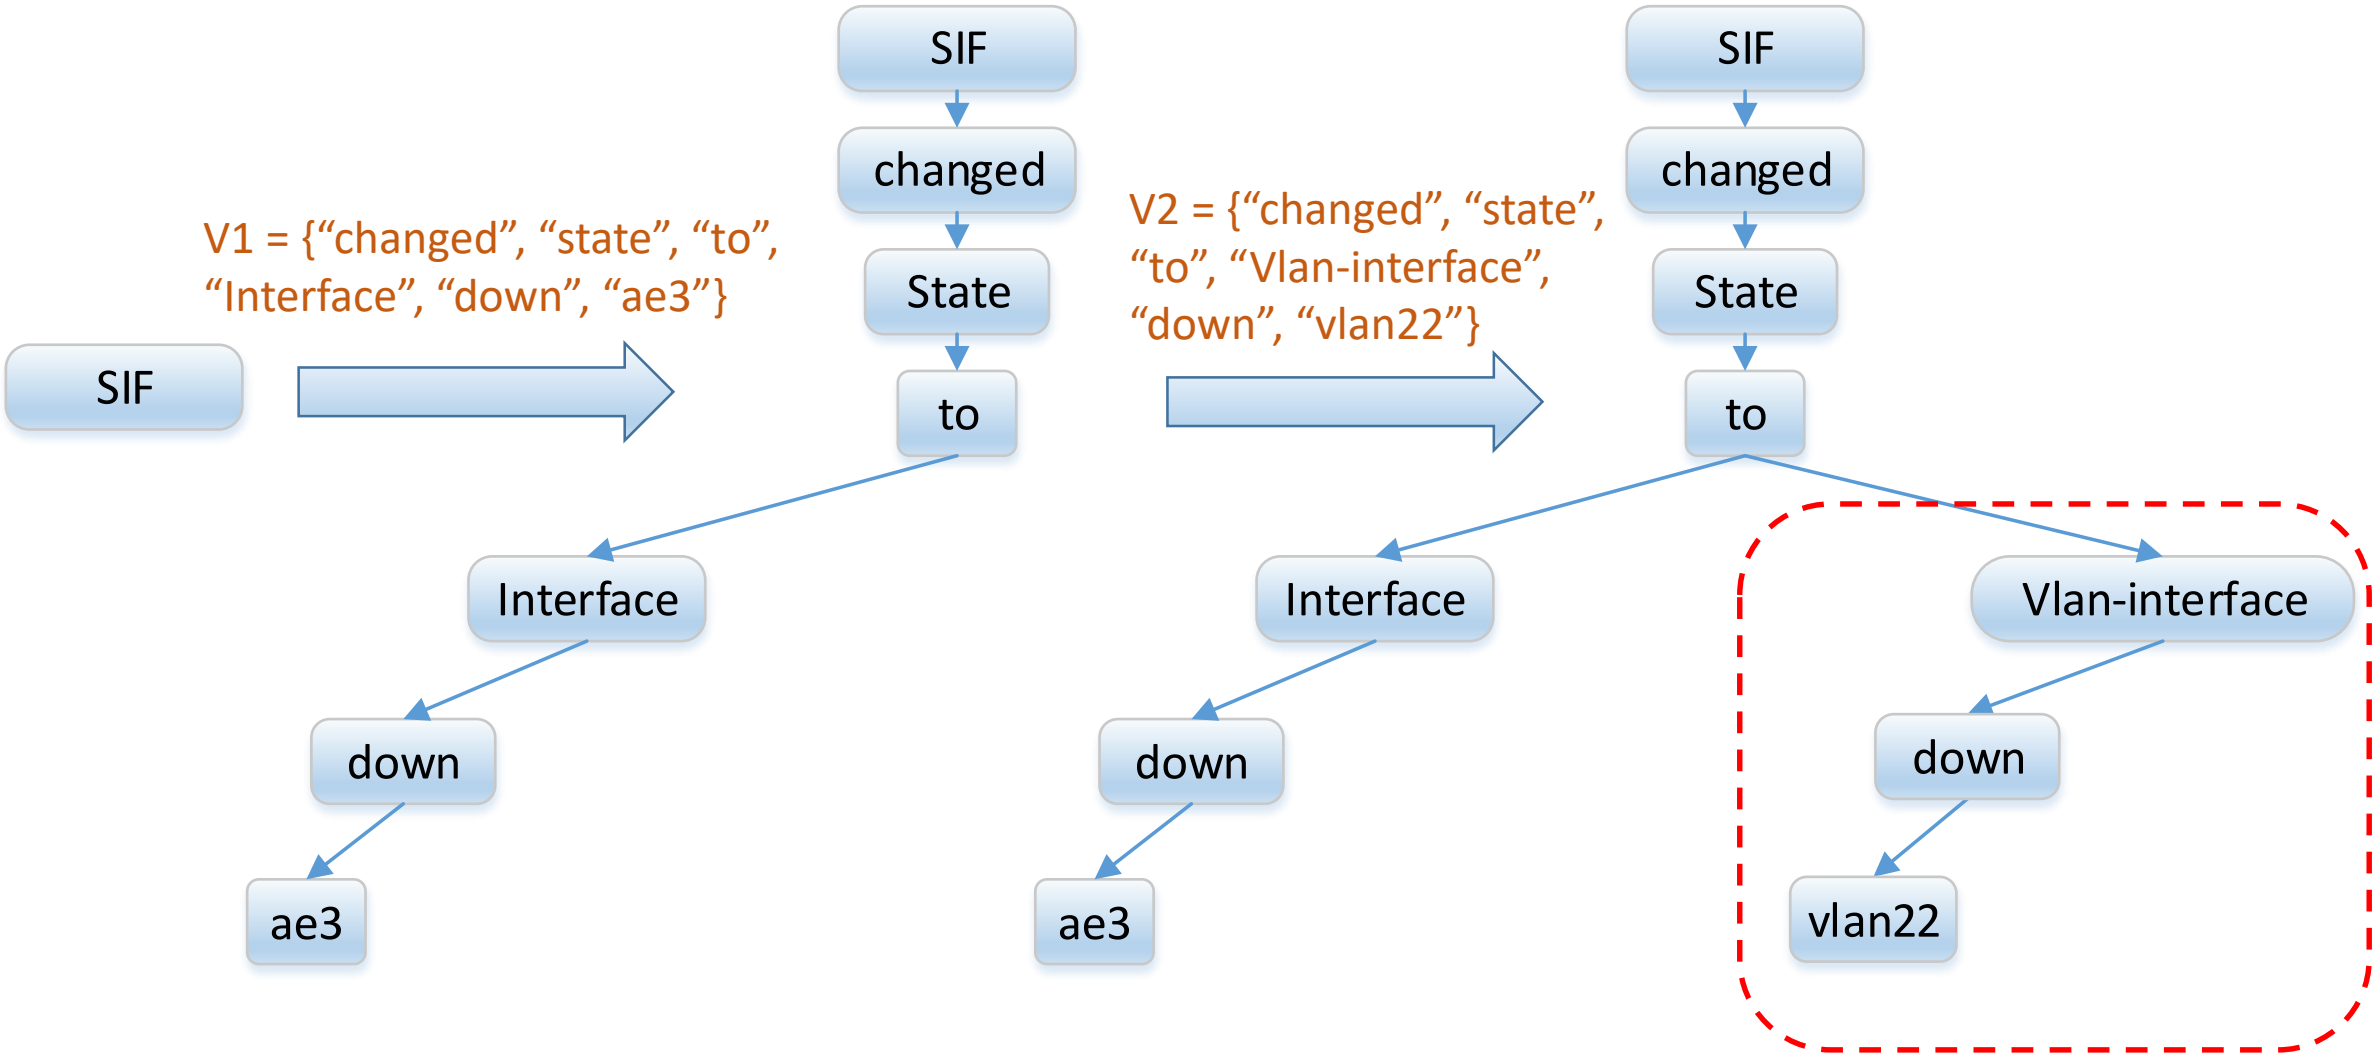

Construct FT-tree

V3 = {“changed”, “state”, “to”, “Interface”, “up”, “ae3” }

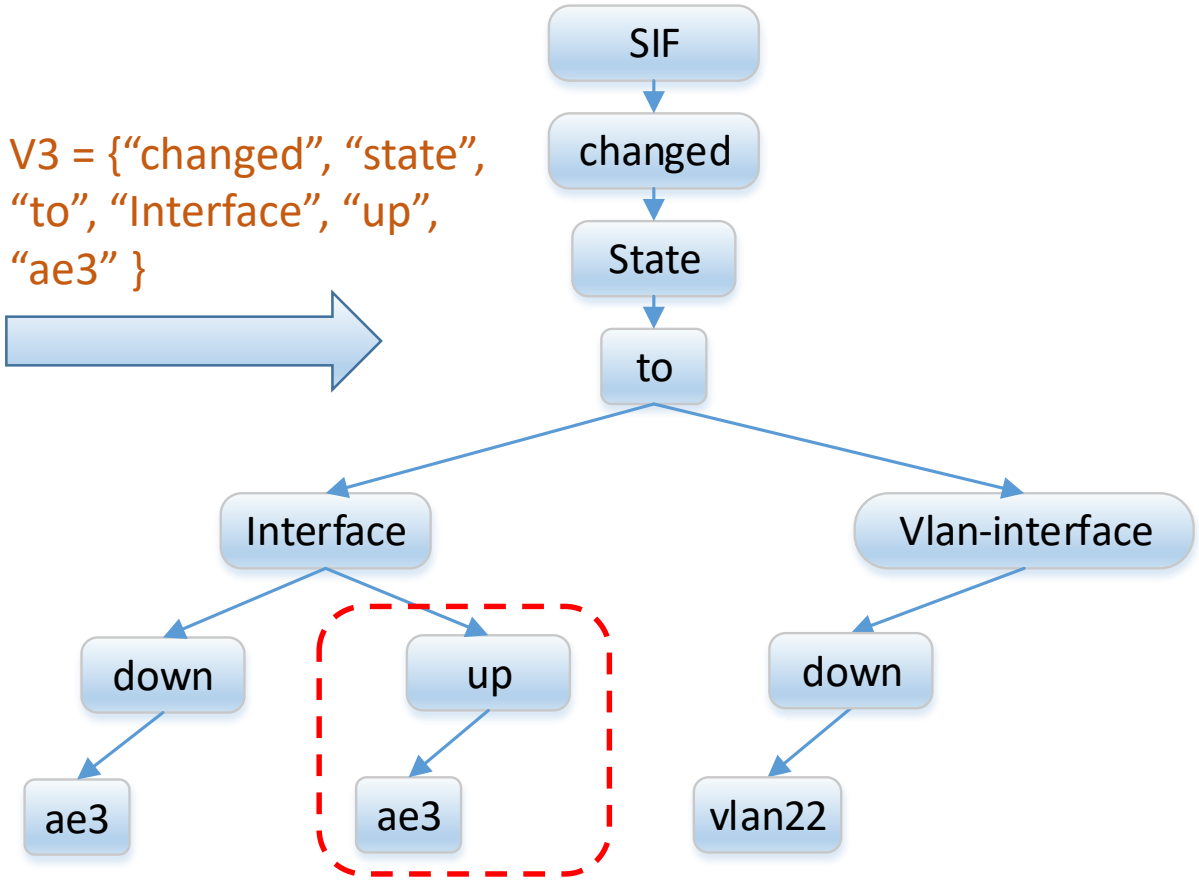

Construct FT-tree

V3 = {“changed”, “state”, “to”, “Interface”, “up”, “ae3” }

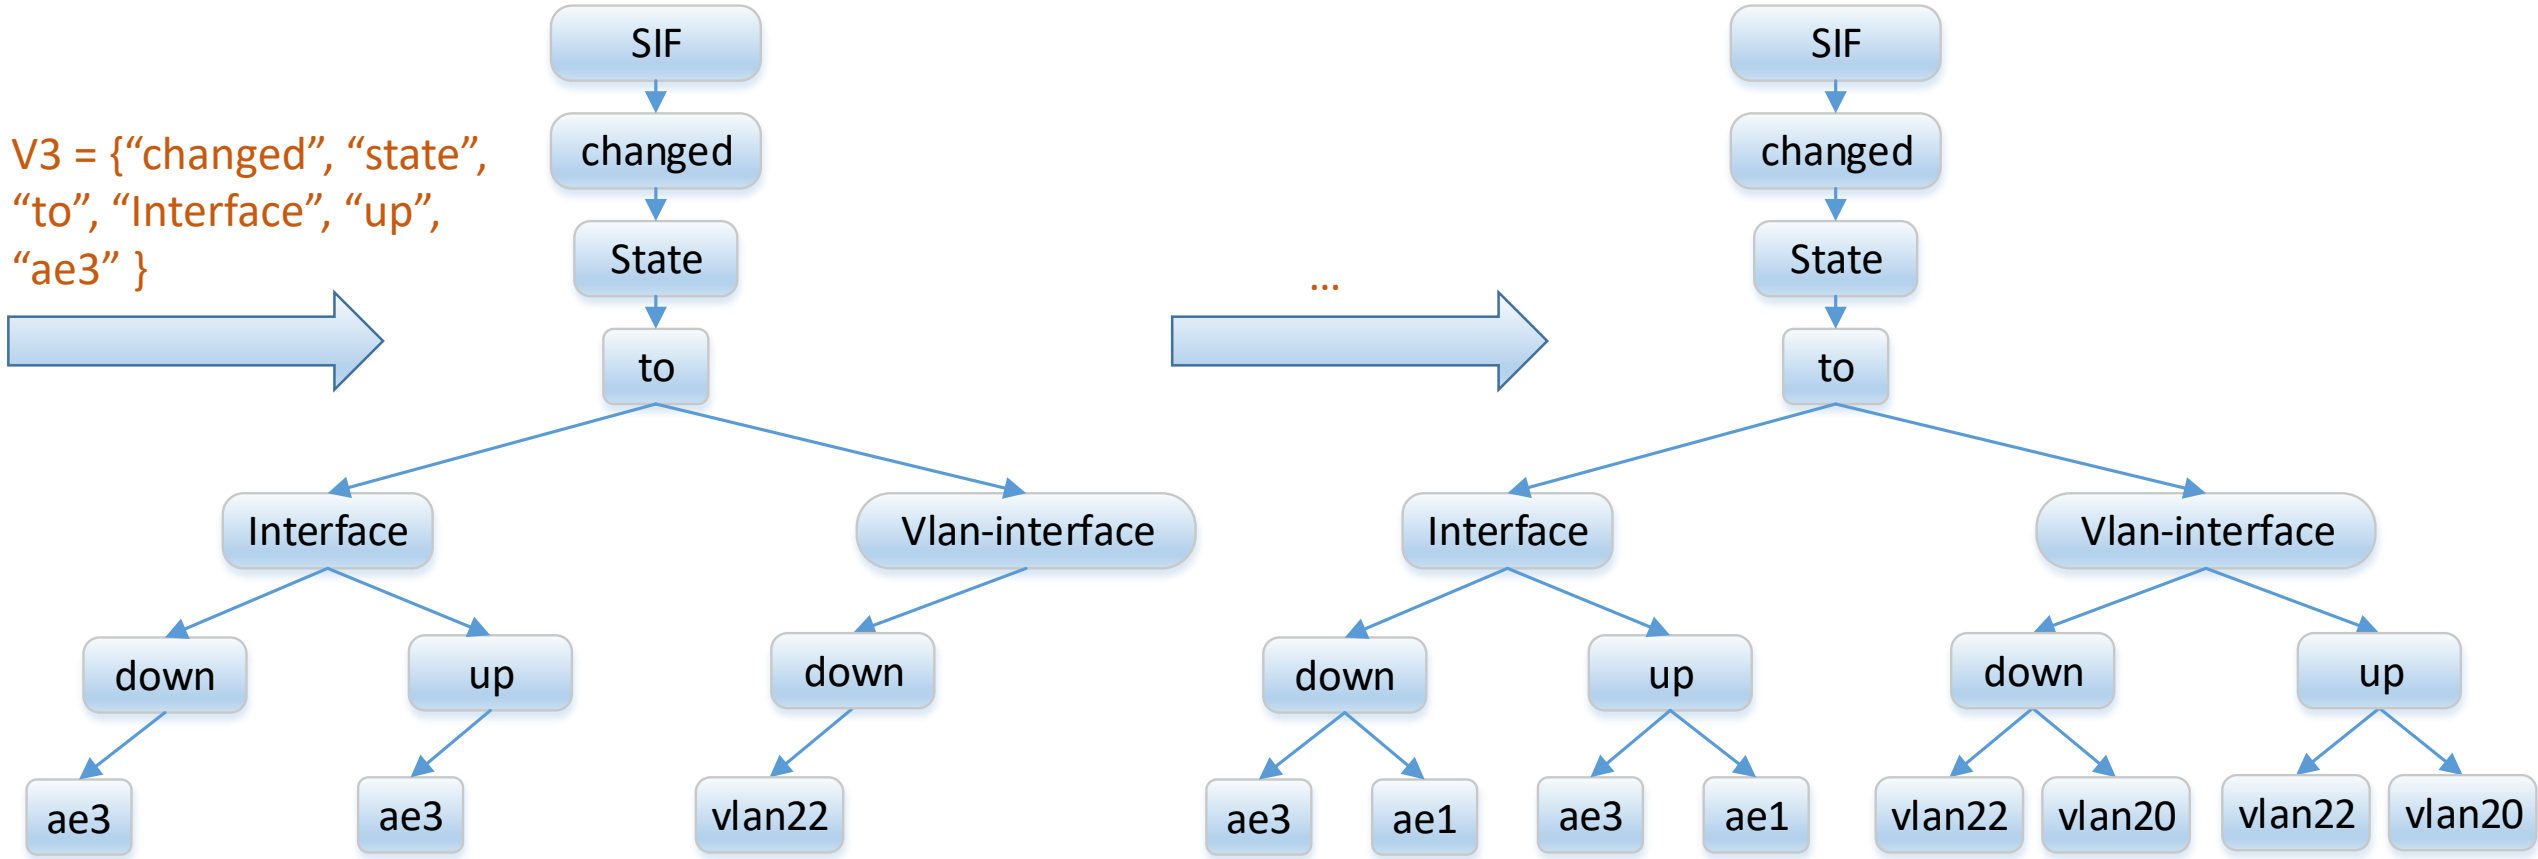

FT-tree Definition

- The item in the root node is syslog message type
- Each node in the tree has one field, word

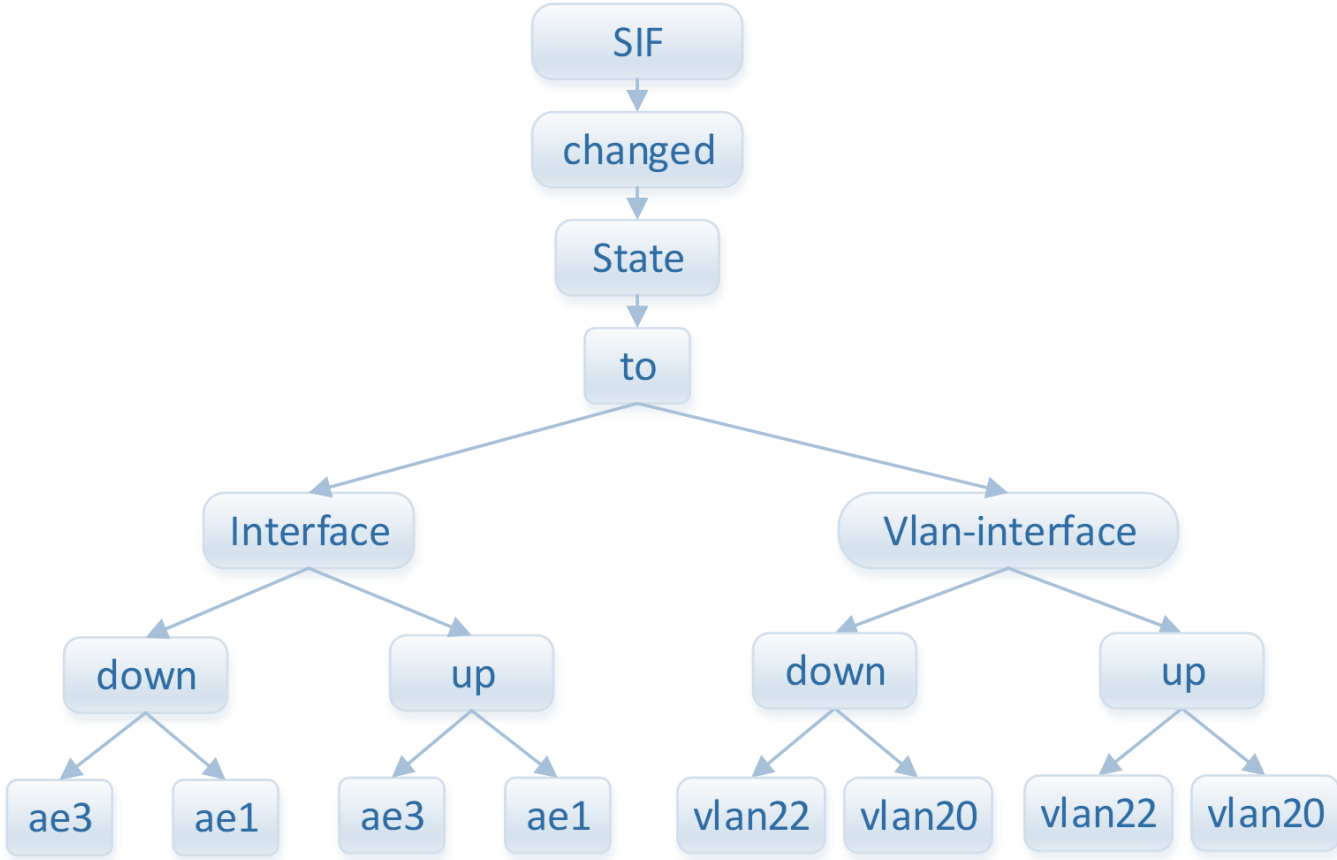

FT-tree Definition

- The item in the root node is syslog message type
- Each node in the tree has one field, word
- Prune the FT-tree
  - A parent node has P+ children
  - The children of this node will be pruned
  - The parent node will become a leaf node

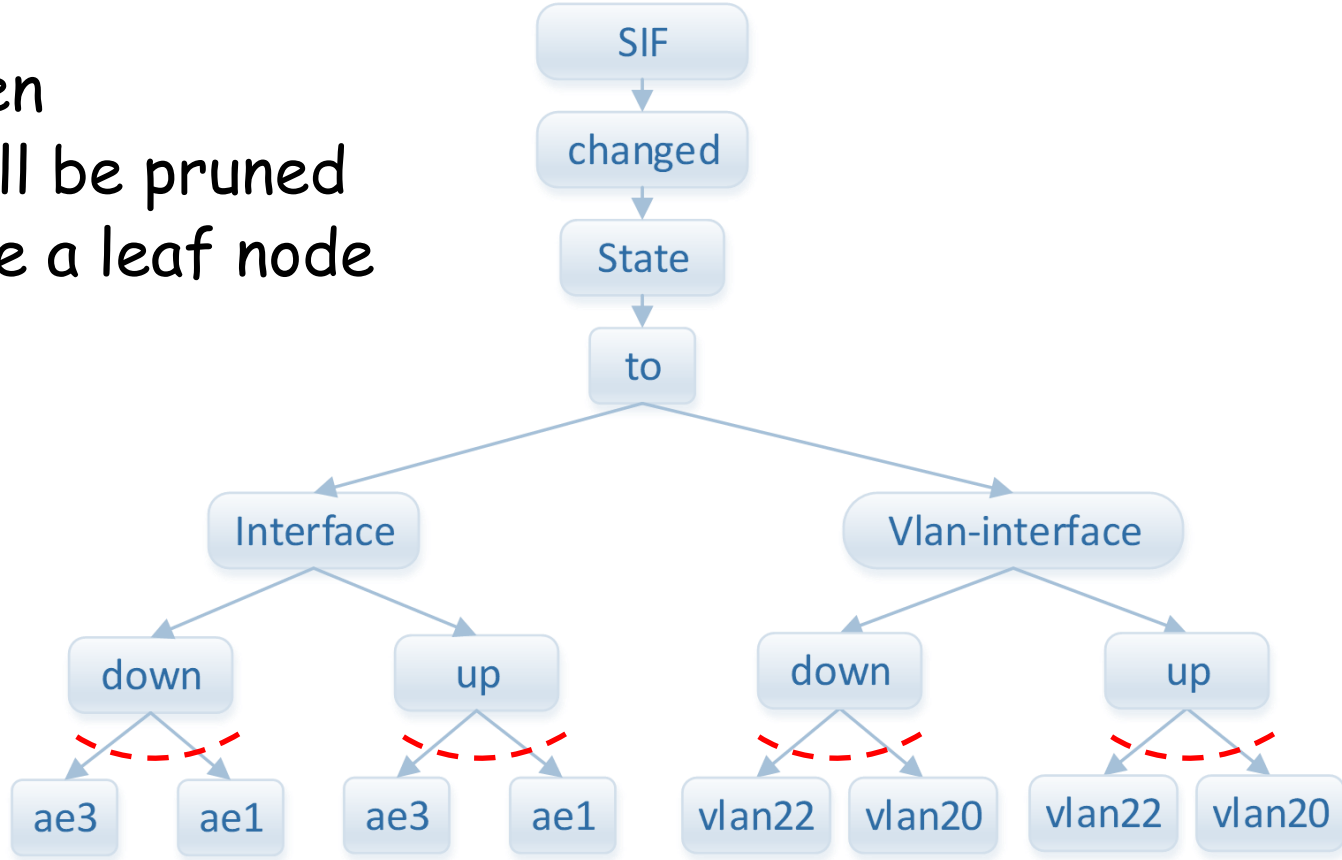

FT-tree Definition

- The item in the root node is syslog message type
- Each node in the tree has one field, word
- Prune the FT-tree
  - A parent node has P+ children
  - The children of this node will be pruned
  - The parent node will become a leaf node
- Words on the path from each leaf to root constitute a template

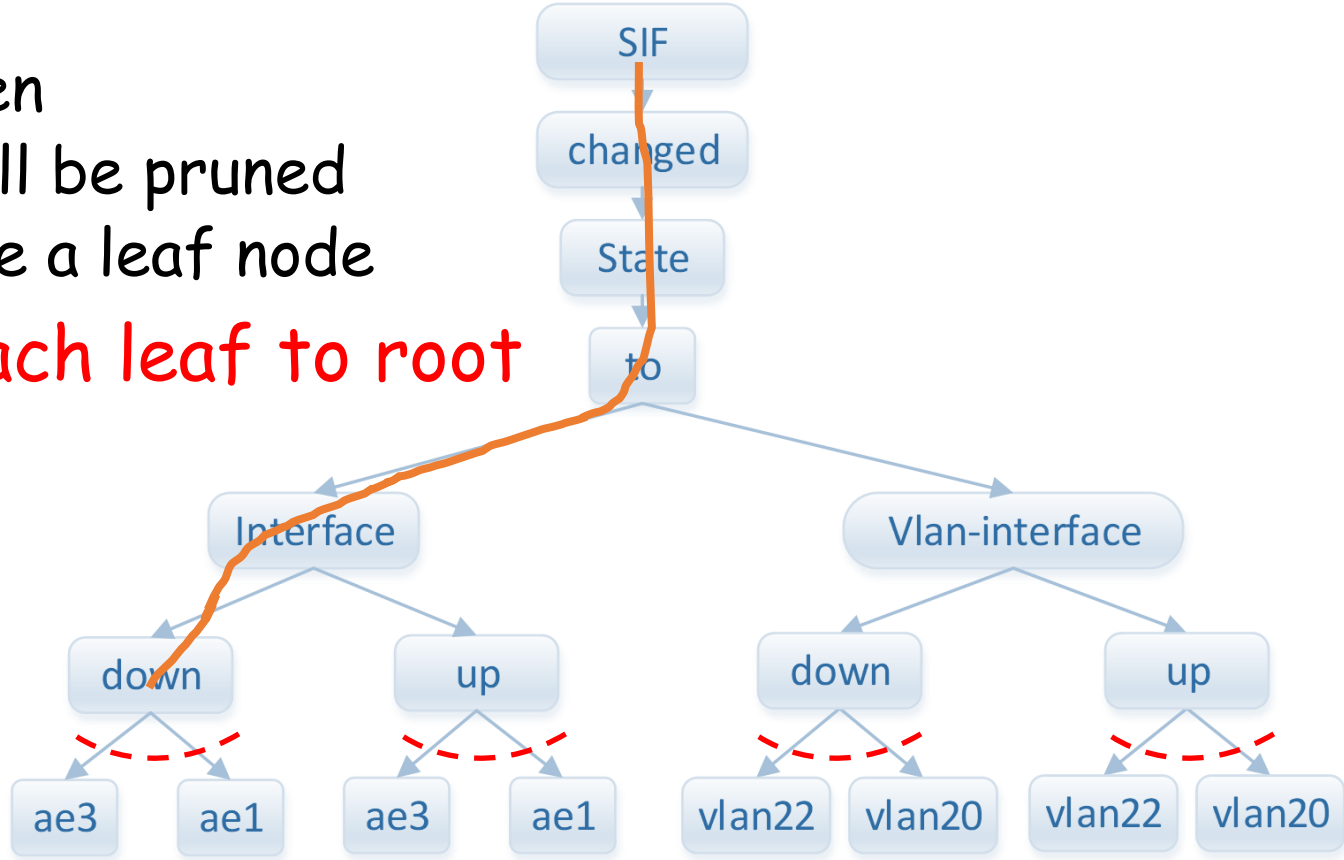

FT-tree: accurate and incrementally re-trainable

Based on word frequency

A template is a combination of words with high frequency

Accurately extract events from syslog messages

Naturally incrementally built

Incrementally re-trainable

Outline

- Background and Motivation
- Challenges
- Key Ideas
- Results
- Conclusion

Evaluation

Dataset

Syslogs & failure tickets

2000+ switches

10+ datacenters

Two-year period

Benchmark methods

- Signature Tree (IMC 10)
- STE (INFOCOM 14)
- LogSimilarity (CNSM 15)

Evaluation

- Compare accuracy
  - Based on manual labels by operators
  - Four types of syslog messages

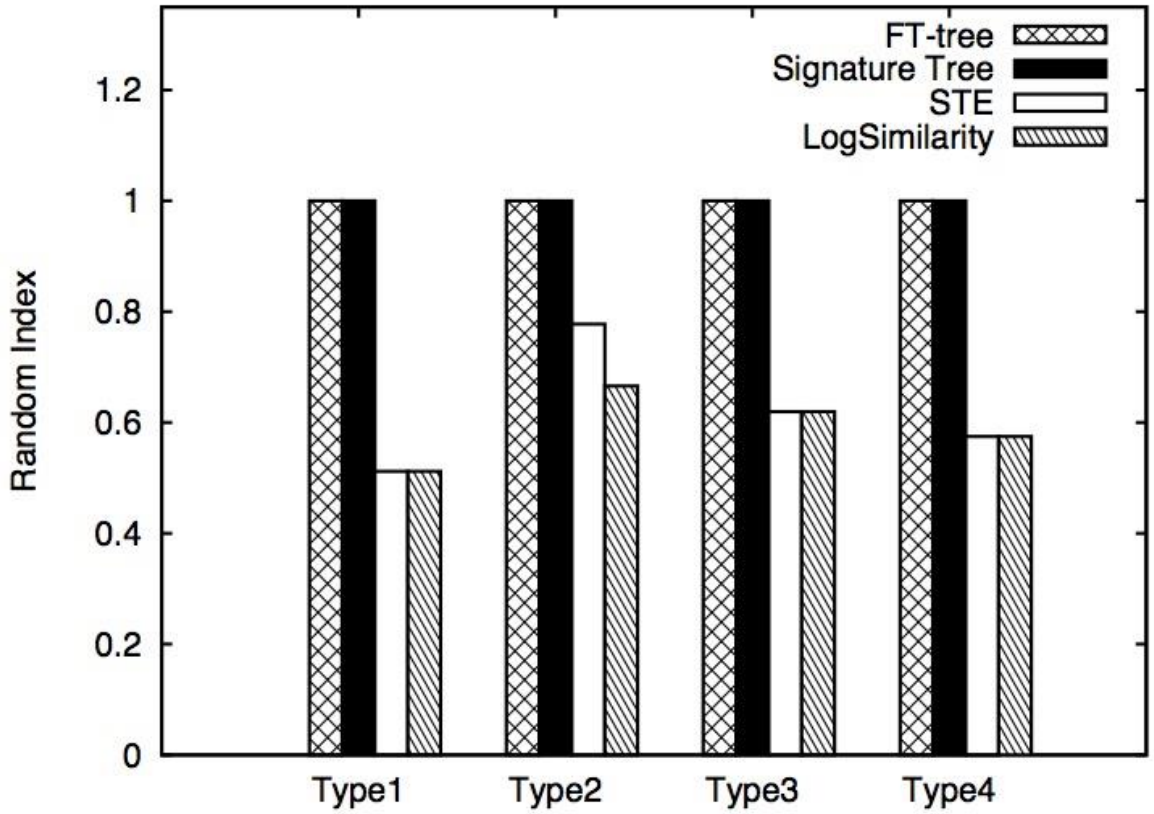

Evaluation

- Compare failure prediction accuracy
  - Hidden Semi-Markov Model (HSMM) as the failure prediction framework
  - 10-fold cross validation

Method	Precision	Recall	F1 measure
FT-tree	32.27%	95.3%	48.21%
Signature Tree	32.27%	95.3%	48.21%
STE	9.14%	99.6%	16.75%
LogSimilarity	10.67%	83.5%	18.93%

Evaluation

- Compare computational efficiency
  - 10 million syslog messages per day
  - Retrained everyday to match new syslog messages
  - The same type of CPU core

Method	FT-tree	Signature Tree	STE	LogSimilarity
Training time	51 mins	628 hours	100 hours	80 mins

Outline

- Background and Motivation
- Challenges
- Key Ideas
- Results
- Conclusion

Conclusion

Challenges of template extraction

- Unstructured texts
- Huge amount of syslogs
- Diverse types of syslogs

FT-tree

- Accurately extract events from syslogs
- Incrementally re-trainable

Evaluation

- Real-world data

Future work

- Switch failure prediction

Thank you!

Q&A

slzhangsd@gmail.com

Q&A

slzhangsd@gmail.com