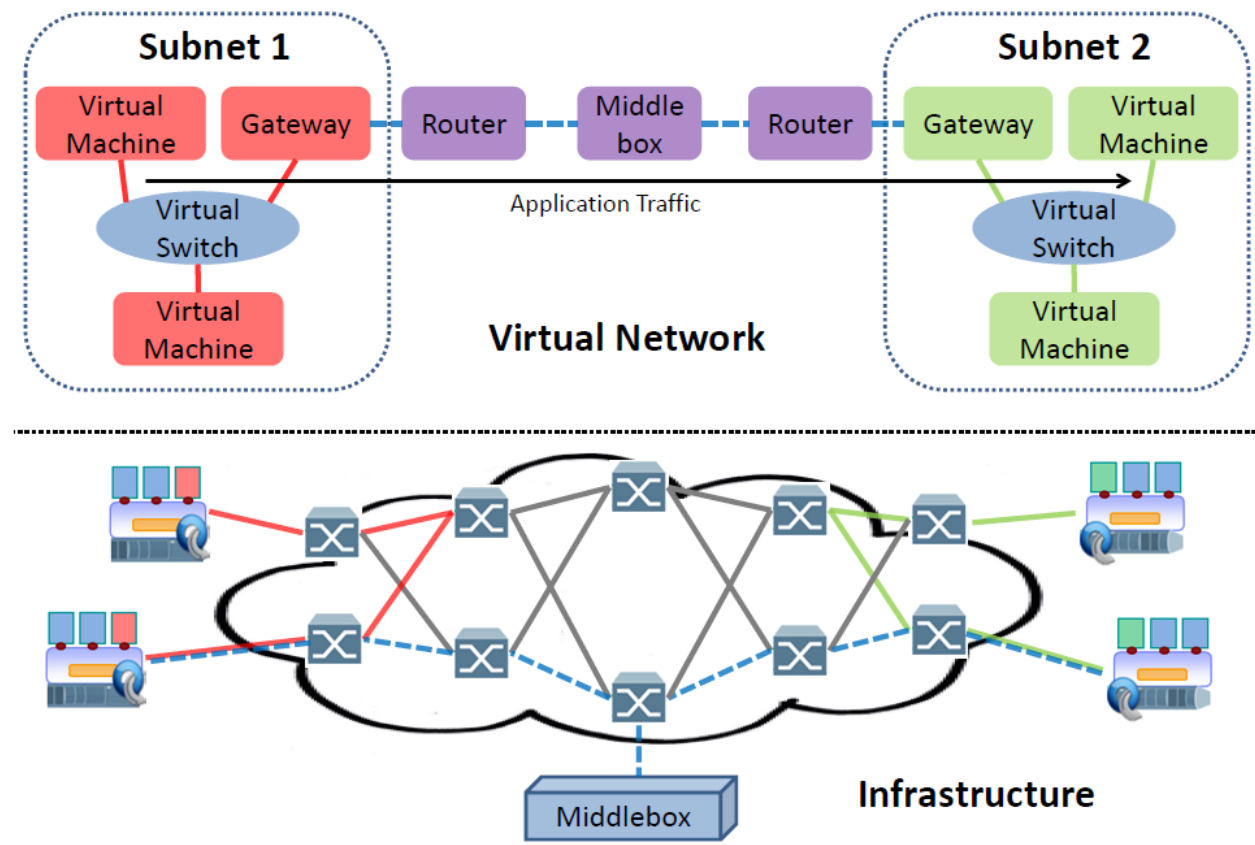


Virtual Networks



Virtual Network Diagnosis

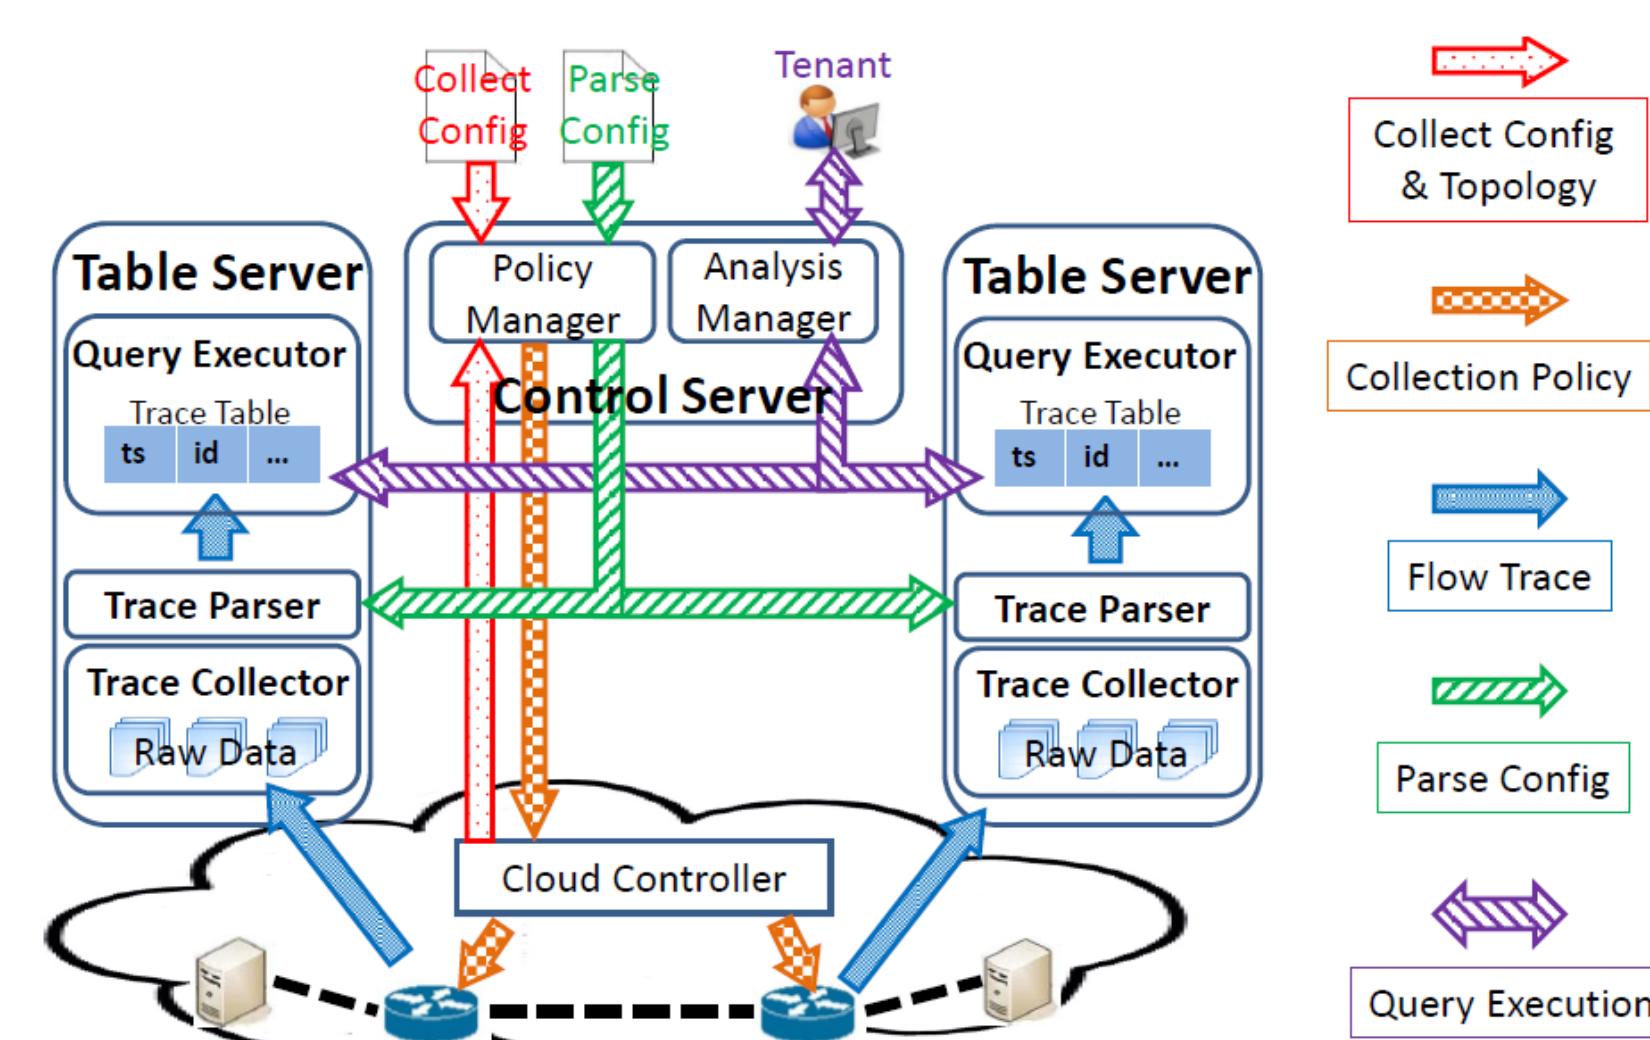
- Multiple layers in the cloud may have various problems
- Isolation and abstraction make it difficult for the tenants to diagnose their virtual networks
 - They cannot touch infrastructure
 - They do not have access to some logic components

Solution & Challenges

- The cloud provider offers a virtual network diagnostic service to tenants
- Tenants use interfaces to analyze their application traffic
- Challenges
 - Preserve abstractions
 - Low overhead
 - Scalability
 - Flow correlation

Architecture

- One Control Server
 - Communication hub
 - Decide data collection policy
- Multiple Table Servers
 - Data Collection
 - Data parse
 - Query Execution

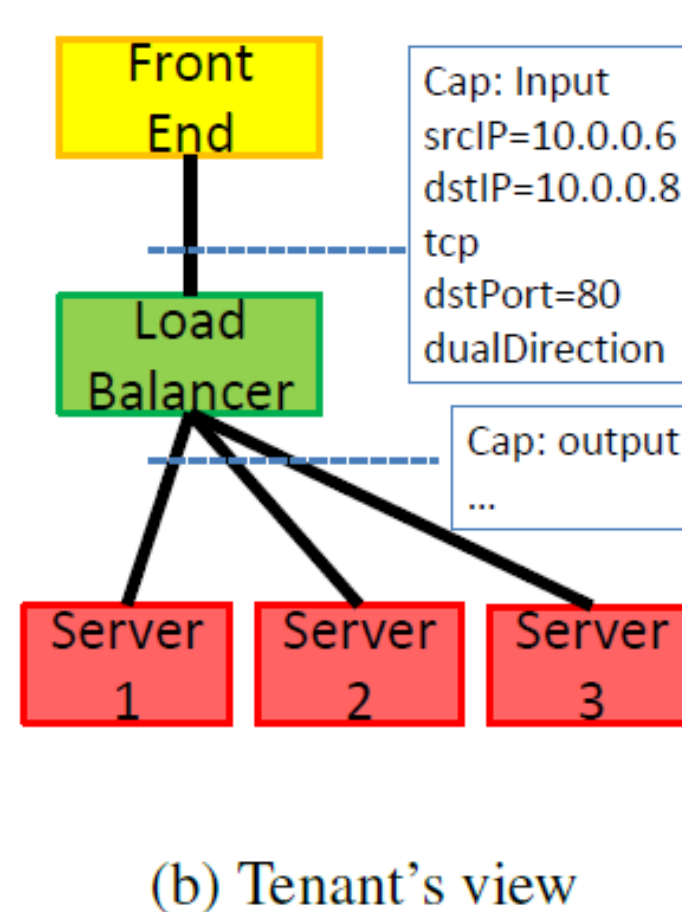


Data Collection

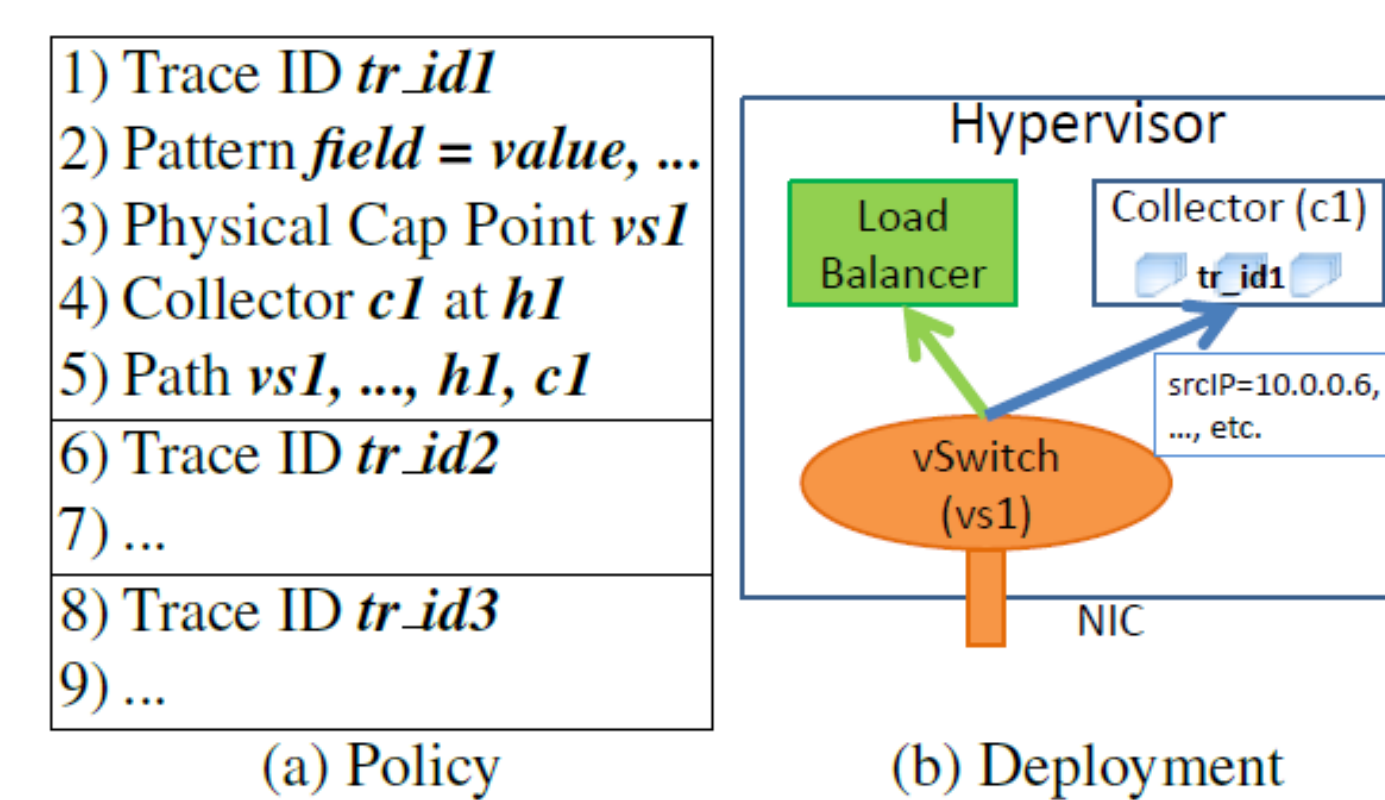
- The tenant submits a diagnosis request

```

1) Appliance Node : lb
2) Cap input
3) srcIP=10.0.0.6/32
4) dstIP=10.0.0.8/32
5) proto=TCP
6) srcPort=*
7) dstPort=80
8) dualDirect=True
9) Cap output
10) ...
11) Appliance ...
12) ...
    
```



- Translate the request into a diagnosis policy
- Set up collectors and configure routing



Data Parse

- Parse packets and extract packet fields
- Dump results into trace tables locally

```

Trace ID tr_id1
Table ID tab_id1
Filter exp
Fields field_list
Table ID tab_id2
...
exp = not exp | exp and exp |
exp or exp | (exp) | prim,
prim = field ∈ value_set,
field_list = field (as name)
(.field (as name))*
    
```

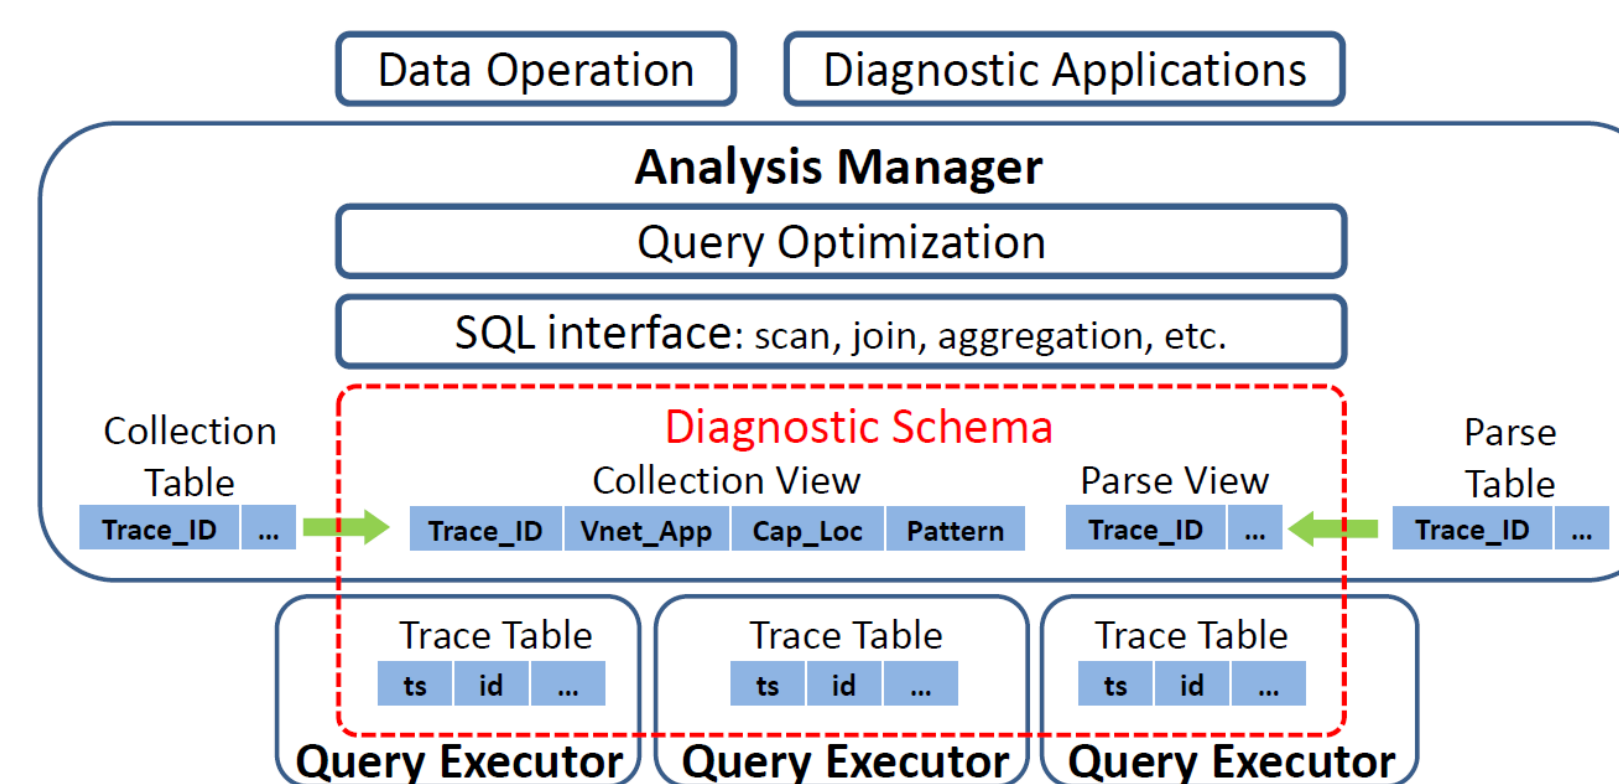
(a) Configuration

```

Trace ID all
Filter: ip.proto = tcp
or ip.proto = udp
Fields: timestamp as ts,
ip.src as src_ip,
ip.dst as dst_ip,
ip.proto as proto,
tcp.src as src_port,
tcp.dst as dst_port,
udp.src as src_port,
udp.dst as dst_port
    
```

(b) Example

- View all tables as a distributed database
- Provide SQL interface to the tenant

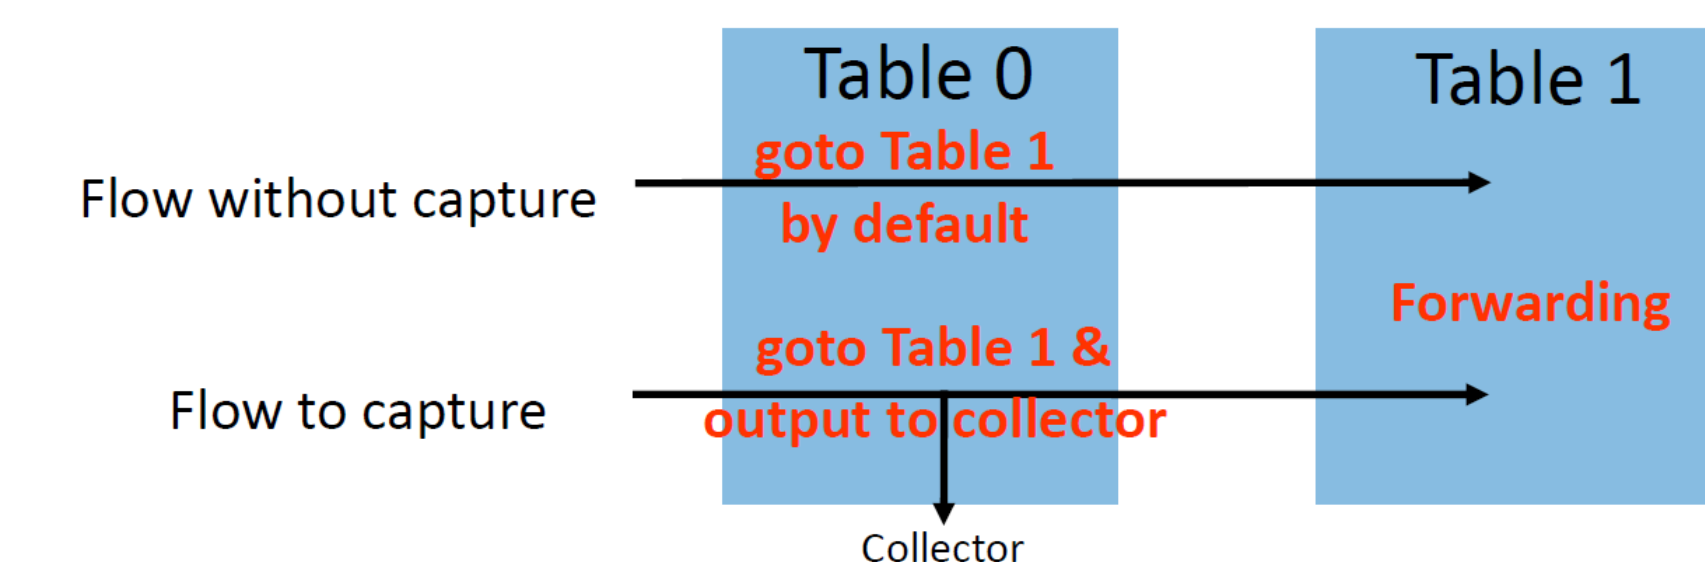


Data Analysis

- Throughput
 - # assume the timestamp unit is second
 - $select\ ceil(ts),\ sum(payload_length)\ from\ table\ group\ by\ ceil(ts)$
- RTT
 1. create view F as $select\ * \ from\ T\ where\ srcIP=IP1\ and\ dstIP=IP2$
 2. create view B as $select\ * \ from\ T\ where\ dstIP=IP1\ and\ srcIP=IP2$
 3. create view RTT as $select\ F.ts\ as\ t1,\ B.ts\ as\ t2\ from\ F,\ B\ where\ F.seq + F.length = B.ack$
 4. $select\ avg(t2-t1)\ from\ RTT\ group\ by\ ceil(t1)$

Optimizations

- Put data collectors locally with capture points
- Separate flow collection rules from routing rules
 - Openflow multi-table feature

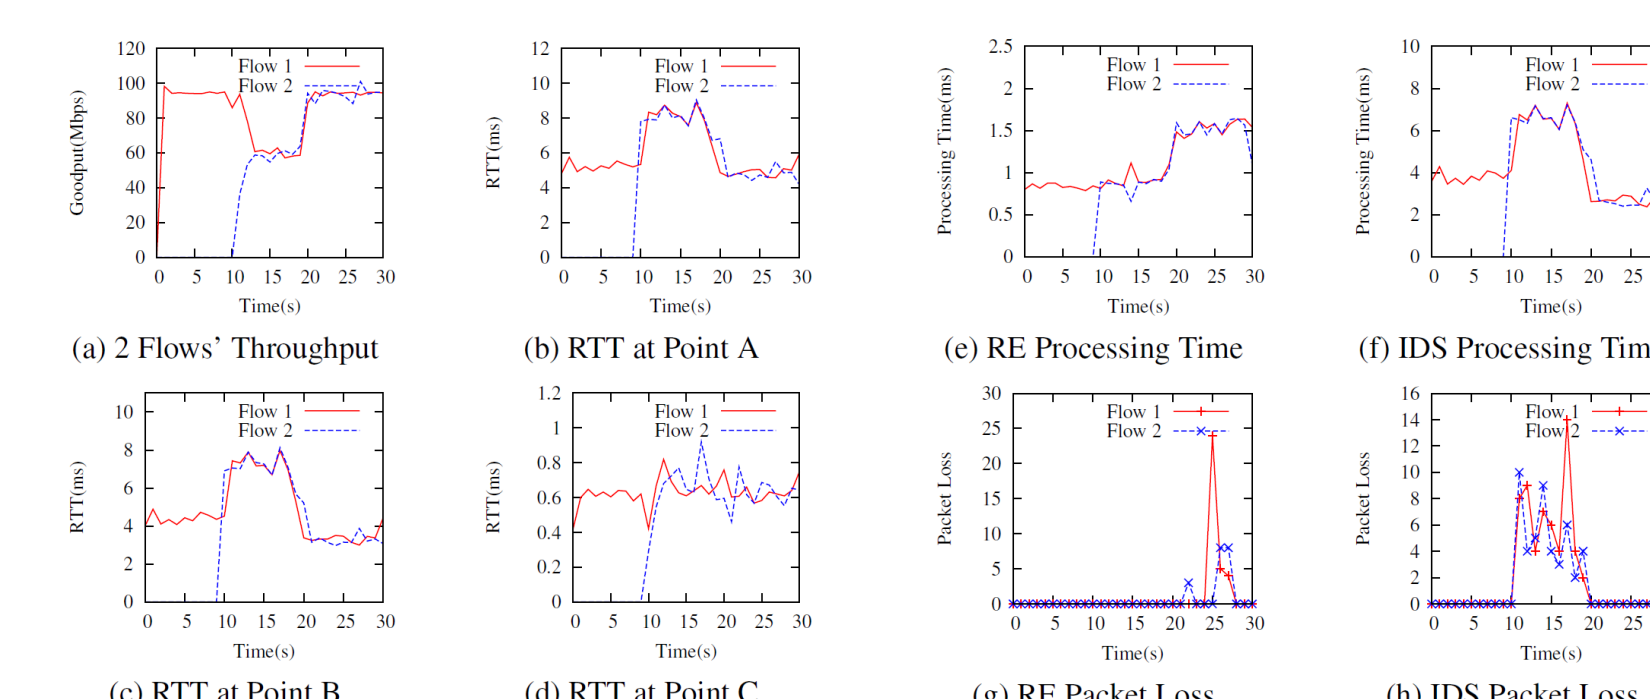
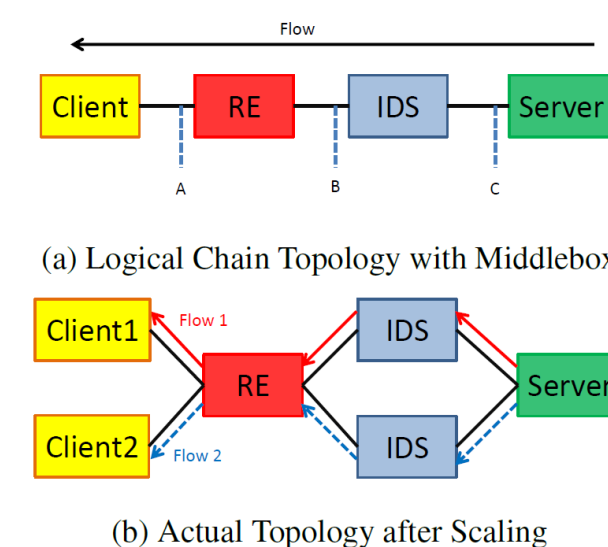


Flow Correlation

- Make use of unique fields in packet headers
 - IP identification, TCP sequence number
- Use the packet payload as the fingerprint
 - hash(payload)
- Other cases
 - layer-4 load balancer: connection built sequence

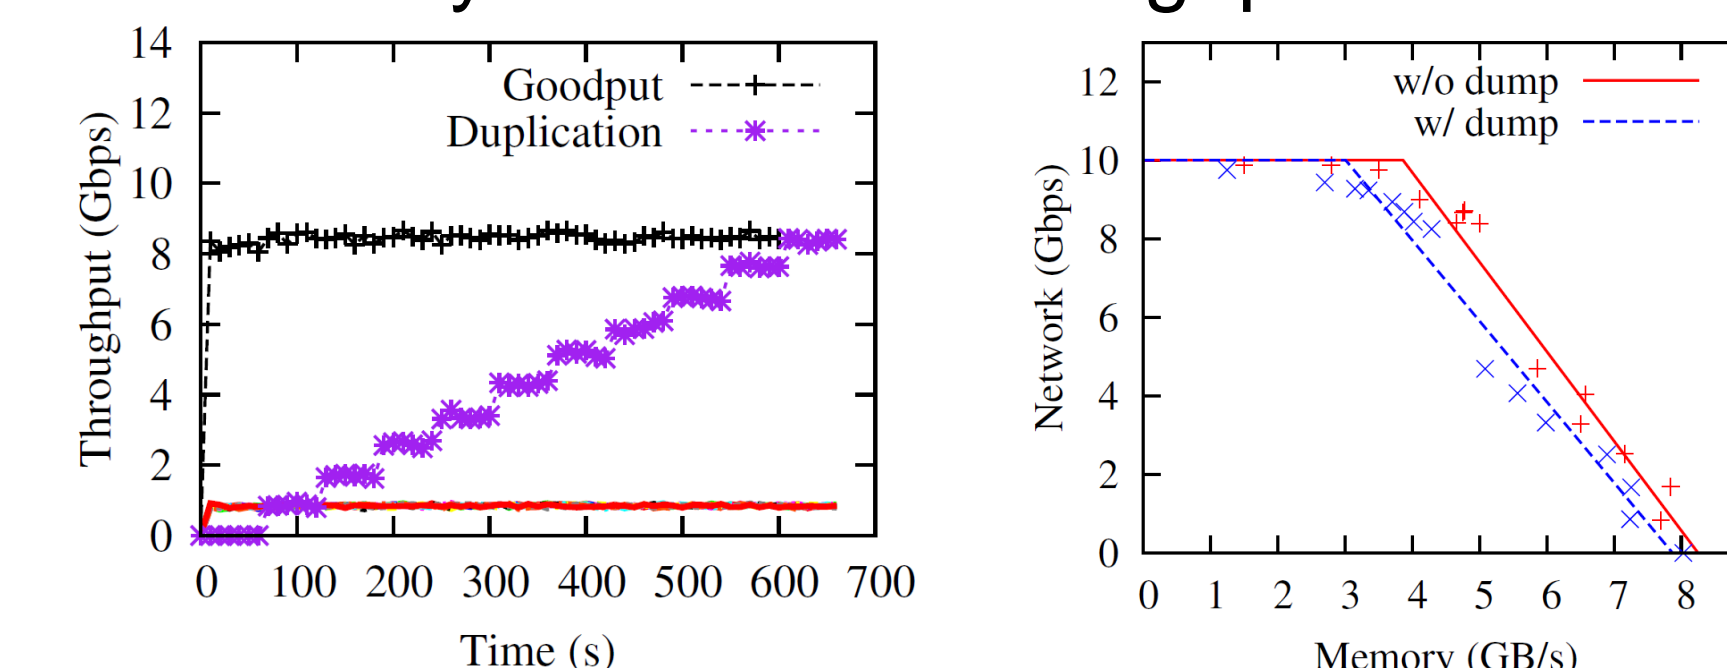
Functional Validation

- Start flow 1 in a chain topology (0s)
- Start flow 2 (10s)
- Find bottleneck
- Scale IDS (20s)



Overhead

- Transfer data between VMs on two hypervisors
 - Flow capture – Throughput
- VMs perform data transfer and memory copy
 - Memory – Network Throughput



- Database storage and query traffic are negligible

Conclusions

- Cloud providers should offer the tenants a virtual network diagnosis service
- We propose VND framework (architecture, interfaces and operations) to provide this service
- VND optimizations make it scalable to many tenants
- Our implementation, experiments and simulation demonstrate the feasibility of VND framework