



Ethane: Taking Control of the Enterprise

Martín Casado, Michael J. Freedman,
Justin Pettit, Jianying Luo,
and Nick McKeown
Stanford University

Scott Shenker
U.C. Berkeley and ICSI

SIGCOMM'07, August 27–31, 2007,

Outline

- ▶ INTRODUCTION
- ▶ OVERVIEW OF ETHANE DESIGN
- ▶ ETHANE IN MORE DETAIL
- ▶ THE *POL-ETH* POLICY LANGUAGE
- ▶ PROTOTYPE AND DEPLOYMENT
- ▶ PERFORMANCE AND SCALABILITY
- ▶ CONCLUSIONS



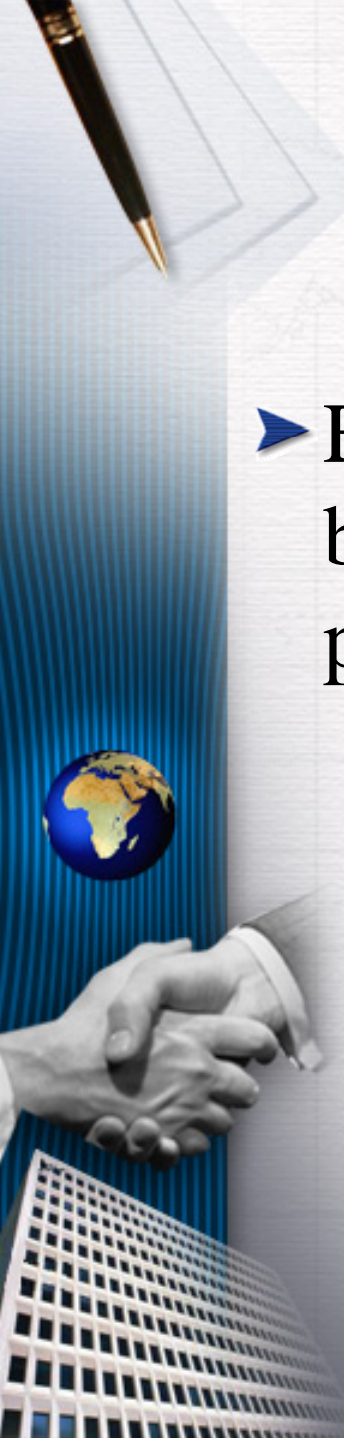
INTRODUCTION(1/2)

- ▶ Enterprise networks
 - ▶ large, variety of applications and protocols,
 - ▶ and typically operate under strict reliability and security constraints;
- ▶ Middle-box
- ▶ Provide tools for diagnosis
 - ▶ By adding a new layer of protocol



INTRODUCTION(2/2)

- *How could we change the enterprise network architecture to make it more manageable?*
- **Ethane** is built around three fundamental principles
 - *The network should be governed by policies declared over high-level names.*
 - *Policy should determine the path that packets follow*
 - *The network should enforce a strong binding between a packet and its origin.*
- *A centralized control architecture.*



OVERVIEW OF ETHANE DESIGN(1/5)

- ▶ Ethane does not allow any communication between end-hosts without explicit permission.
 - ▶ *Central Controller*
 - ▶ *Ethane Switches*

OVERVIEW OF ETHANE DESIGN(2/5)

- ▶ First, Ethane takes over all the binding of addresses.
 - ▶ DHCP-know which switch port the machine is connected.
 - ▶ the packet must come from a machine that is registered on the network.
 - ▶ Users are required to authenticate with the network.
 - ▶ binding users to hosts.

OVERVIEW OF ETHANE DESIGN(3/5)

- ▶ Second, the packet must come from a machine that is registered on the network.
- ▶ Finally, users are required to authenticate themselves with the network.

OVERVIEW OF ETHANE DESIGN(4/5)





Figure 1: Example of communication on an Ethane network. Route setup shown by dotted lines; the path taken by the first packet of a flow shown by dashed lines.



OVERVIEW OF ETHANE DESIGN(5/5)

► Five basic activities

► Registration.

- Hosts-mac, users-name & password, switch-public key.
- at the Controller

► Bootstrapping.

- spanning tree

► Authentication.

- Hosts \leftrightarrow switch \leftrightarrow controller
- Hosts send DHCP request= \rightarrow name to IP, IP to mac, mac to switch port
- Users is bound to hosts

► Flow Setup.

► Forwarding.

ETHANE IN MORE DETAIL(1/4)





Figure 2: An example Ethane deployment.

ETHANE IN MORE DETAIL(2/4)

- Switches
- A Ethane Switch is like a simplified Ethernet switch.
 - has several interfaces that send and receive standard Ethernet packets.
 - much simpler.



ETHANE IN MORE DETAIL(3/4)

- Flow Table and Flow Entries.
 - The Switch datapath is a managed flow table.
 - Flow entries contain
 - a Header, an Action, Per-Flow Data
 - Action-two common types of entry
 - Per-flow entries describing application flows that should be *forwarded*,
 - Per-host entries that describe misbehaving hosts whose packets should be *dropped*.
 - Only the Controller can add entries to the flow table.

ETHANE IN MORE DETAIL(4/4)

► Controller

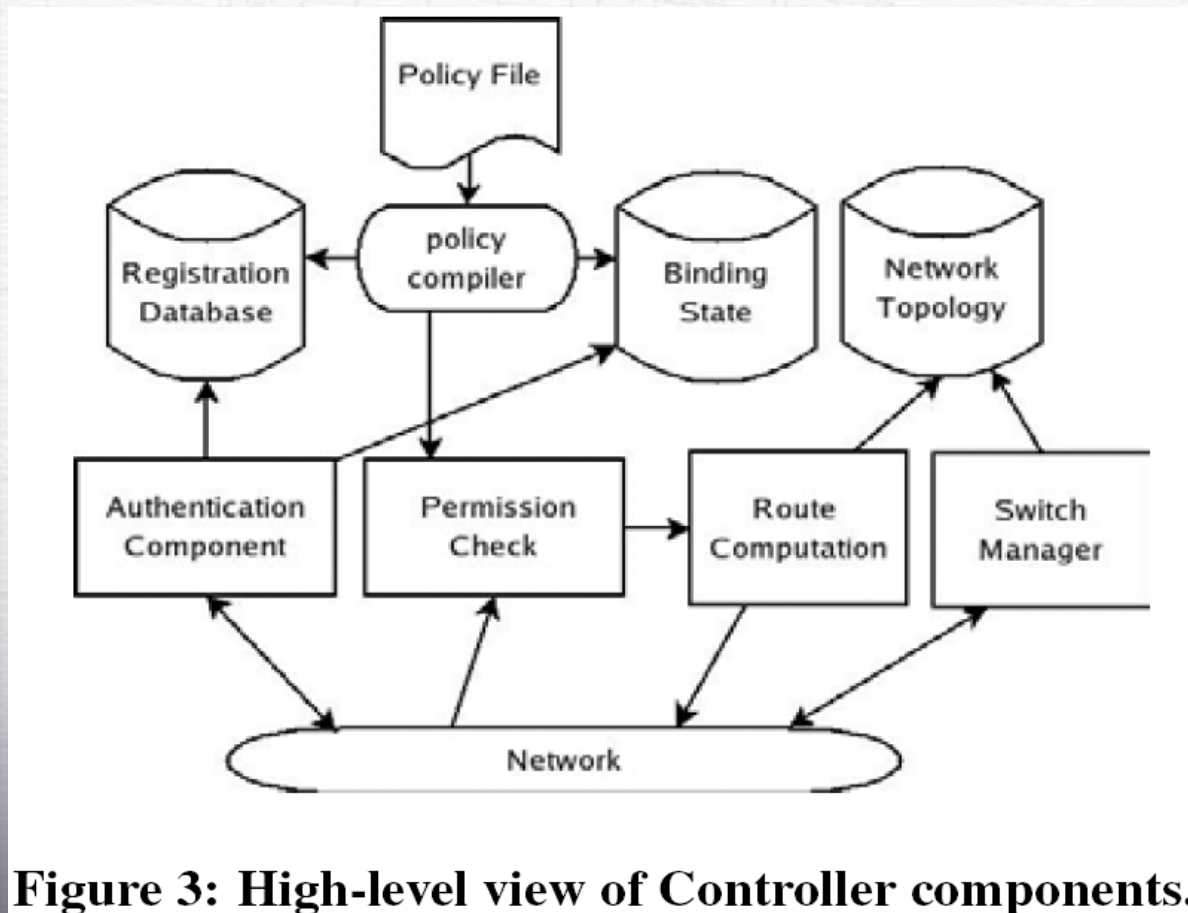




Figure 3: High-level view of Controller components.



THE *POL-ETH* POLICY LANGUAGE(1/2)

- 
- 
- ▶ *Pol-Eth* is a language for declaring policy in an Ethane network.
 - ▶ $[(\text{usrc}=\text{"bob"}) \wedge (\text{protocol}=\text{"http"}) \wedge (\text{hdst}=\text{"websrv"})]:\text{allow};$
 - ▶ Conditions
 - ▶ domains include $\{\text{usrc}, \text{udst}, \text{hsrc}, \text{hdst}, \text{apsrc}, \text{apdst}, \text{protocol}\}$
 - ▶ Actions
 - ▶ *allow*, *deny*, *waypoints*, and *outbound-only* (for NAT-like security).
 - ▶ *waypoints*("ids", "webproxy").

THE *POL-ETH* POLICY LANGUAGE(2/2)

```
# Groups —
desktops = ["griffin", "roo"];
laptops = ["glaptop", "rlaptop"];
phones = ["gphone", "rphone"];
server = ["http_server", "nfs_server"];
private = ["desktops", "laptops"];
computers = ["private", "server"];
students = ["bob", "bill", "pete"];
profs = ["plum"];
group = ["students", "profs"];
waps = ["wap1", "wap2"];
%%
# Rules —
[(hsrc=in("server")^(hdst=in("private")))] : deny;
# Do not allow phones and private computers to communicate
[(hsrc=in("phones")^(hdst=in("computers")))] : deny;
[(hsrc=in("computers")^(hdst=in("phones")))] : deny;
# NAT-like protection for laptops
[(hsrc=in("laptops"))] : outbound-only;
# No restrictions on desktops communicating with each other
[(hsrc=in("desktops")^(hdst=in("desktops")))] : allow;
# For wireless, non-group members can use http through
# a proxy. Group members have unrestricted access.
[(apsrc=in("waps"))^(user=in("group"))] : allow;
[(apsrc=in("waps"))^(protocol="http")] : waypoints("http-proxy");
[(apsrc=in("waps"))] : deny;
[] : allow; # Default-on: by default allow flows
```

Figure 4: A sample policy file using *Pol-Eth*



PROTOTYPE AND DEPLOYMENT()

► Ethane

- connects over 300 registered hosts and several hundred users.
- 19 Switches of three different types:
 - Ethane wireless access points and
 - Ethane Ethernet switches in two flavors



PERFORMANCE AND SCALABILITY(1/6)

- ▶ Our primary question:
 - ▶ How many Controllers are needed for a network of a given size?
- ▶ Consider the question:
 - ▶ How big does the flow table need to be in the Switch?

PERFORMANCE AND SCALABILITY(2/6)




Figure 5: Frequency of flow-setup requests per second to Controller over a 10-hour period (top) and 4-day period (bottom).

PERFORMANCE AND SCALABILITY(3/6)




Figure 6: Flow-setup times as a function of Controller load. Packet sizes were 64B, 128B and 256B, evenly distributed.

PERFORMANCE AND SCALABILITY(4/6)




Figure 7: Active flows for LBL network [19].




Figure 8: Flow-request rate for Stanford network.

PERFORMANCE AND SCALABILITY(5/6)

► Performance During Failures

Failures	0	1	2	3	4
Completion time	26.17s	27.44s	30.45s	36.00s	43.09s

Table 1: Completion time for HTTP GETs of 275 files during which the primary Controller fails zero or more times. Results are averaged over 5 runs.

PERFORMANCE AND SCALABILITY(6/6)

► Flow Table Sizing

- holding 8K–16K entries, each entry is 64B, such a table requires about 1MB of storage, or as much as 4MB with two-way hashing scheme




Figure 9: Active flows through two of our deployed switches

A vertical decorative strip on the left side of the slide. It contains, from top to bottom: a black pen with a gold tip, a small globe showing Africa and Europe, a black and white image of two hands shaking, and a stylized grid pattern representing a building facade.

CONCLUSIONS

- It much easier to manage the Ethane network than we expected.
- It is natural and fast to add new policy rules in a single location.