Six/One

A Scalable and Traffic-Engineering-Compatible Multi-Homing Solution

Christian Vogt

IST Routing in Next Generation workshop, Madrid
December 13, 2007

Towards Increased Reliability and Performance

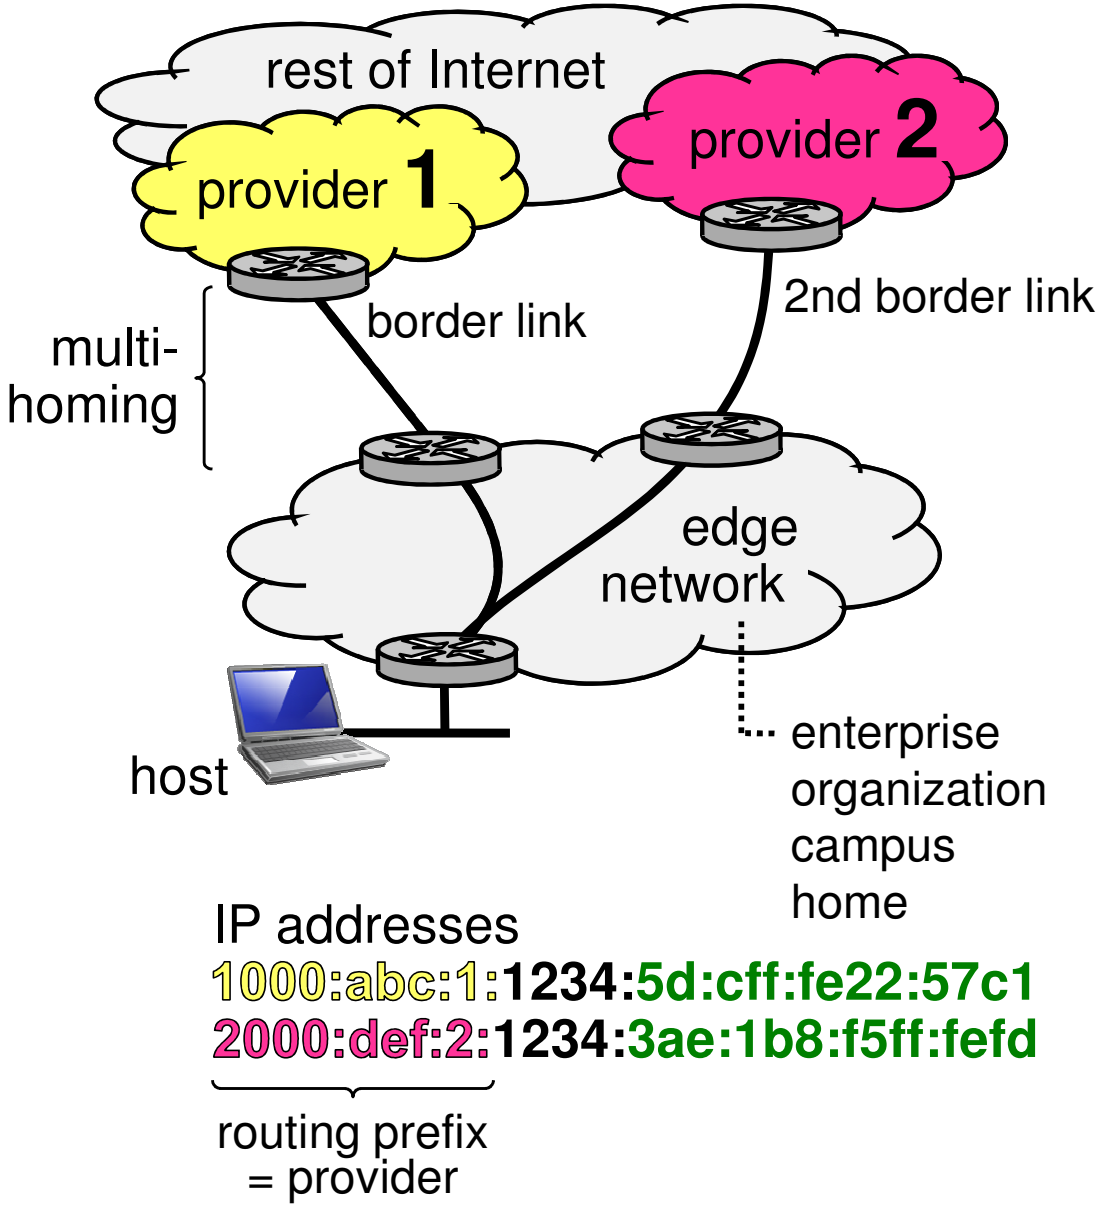

- Dynamic traffic redirection between border links desired

- No support in classic Internet
  - Addresses encode provider
  - Enables aggregation, but redirection ⇒ address changes
  - Applications need stable address
  - Address change ⇒ applications abort

Towards Increased Reliability and Performance

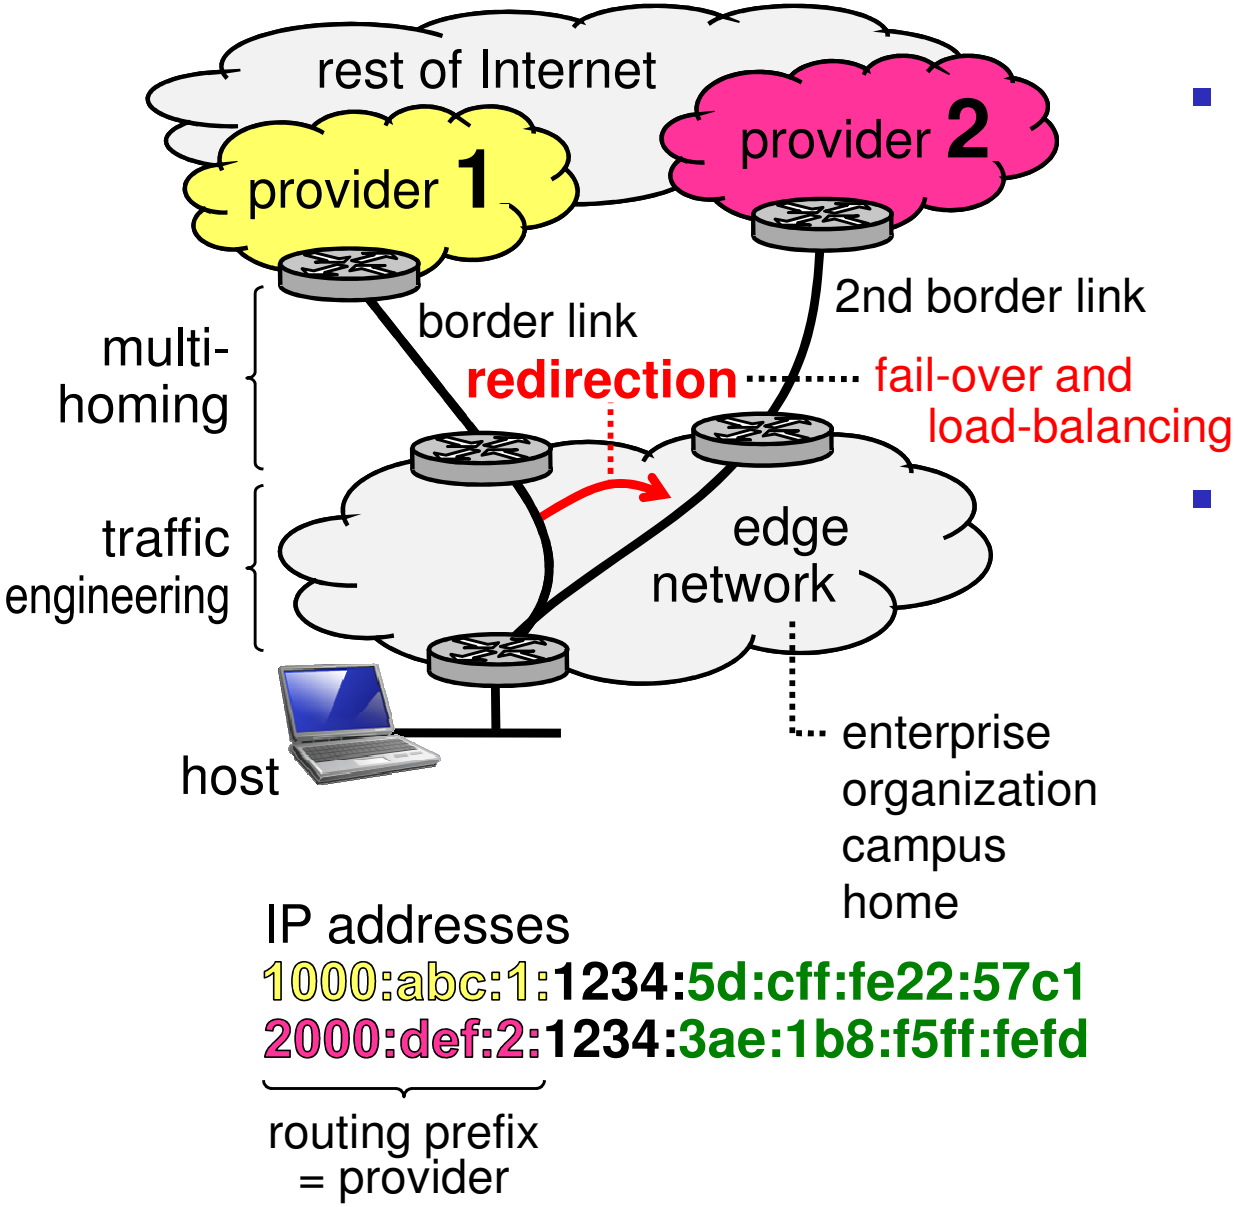

- Dynamic traffic redirection between border links desired
- No support in classic Internet
  - Addresses encode provider
  - Enables aggregation, but redirection ⇒ address changes
  - Applications need stable address
  - Address change ⇒ applications abort

Today‘s Practice: Provider-Independent Addresses

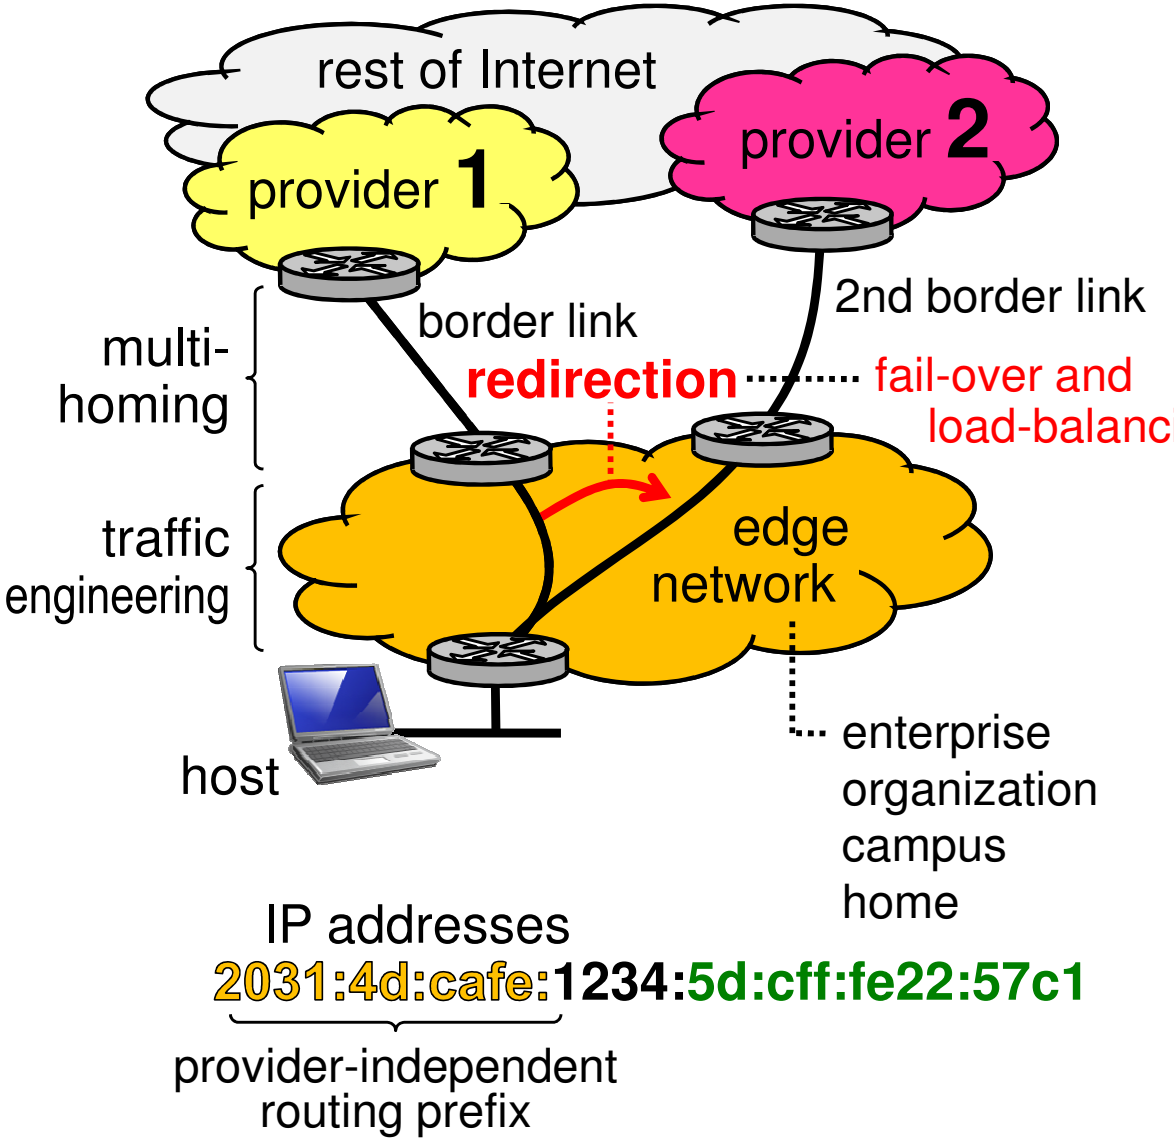

- Simple traffic engineering
  - Egress traffic via any provider
  - Ingress traffic redirection via Border Gateway Protocol
- But increased load on routers
  - Aggregation impossible
  - Separate routing table entries for edge networks
  - Global routing table updates for local traffic redirection
- Challenge: better scalability, same traffic engineering

Related Work

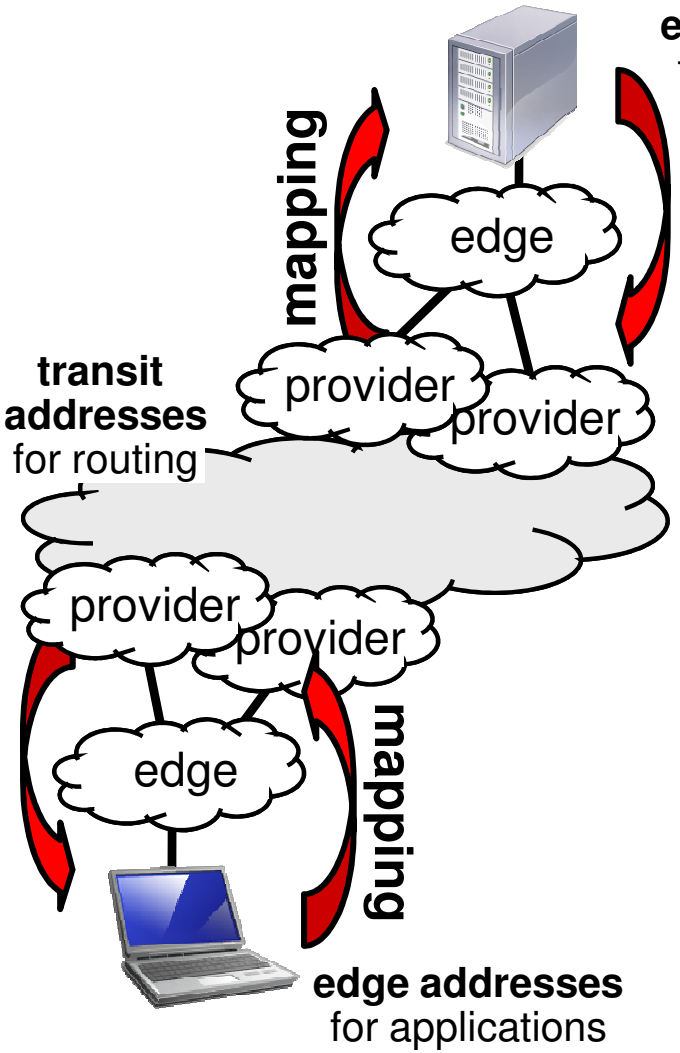

		Shim6 host-controlled address mapping	LISP network-controlled address mapping
scalability	reduced size of routing table	all IP addresses aggregatable	aggregatable after mapping
	fewer updates of routing table	✓	✓
traffic engineering compatibility	reliable fail-over	end-to-end connectivity checks	edge-to-edge connectivity checks
	load-balancing	IP addresses provider-dependent	provider-independent before mapping
incremental deployability	backward-compatibility	no conflicts with legacy Internet	edge addresses unreachable
	incentives for early adoption	if support at both ends	edge addresses unreachable

- Most proposed improvements based on address mapping
  - Between stable edge addresses and aggregatable transit addresses
- Not traffic-engineering-compatible, or only at high cost backward-compatible

Contributions of Six/One

- Scalable, traffic-engineering-compatible multi-homing
  - Incrementally deployable, incentives for early adoption
- IP address bunches
  - Address mapping on hosts
  - Routers may rewrite addresses to redirect traffic
    - See also 8+8 and draft-nordmark-shim6-esd
- Extensions
  - Proxies: Deployment without host upgrades
    - See also draft-bagnulo-pshim6
  - Address translation: redirect packets without remote support

Six/One in a Nutshell

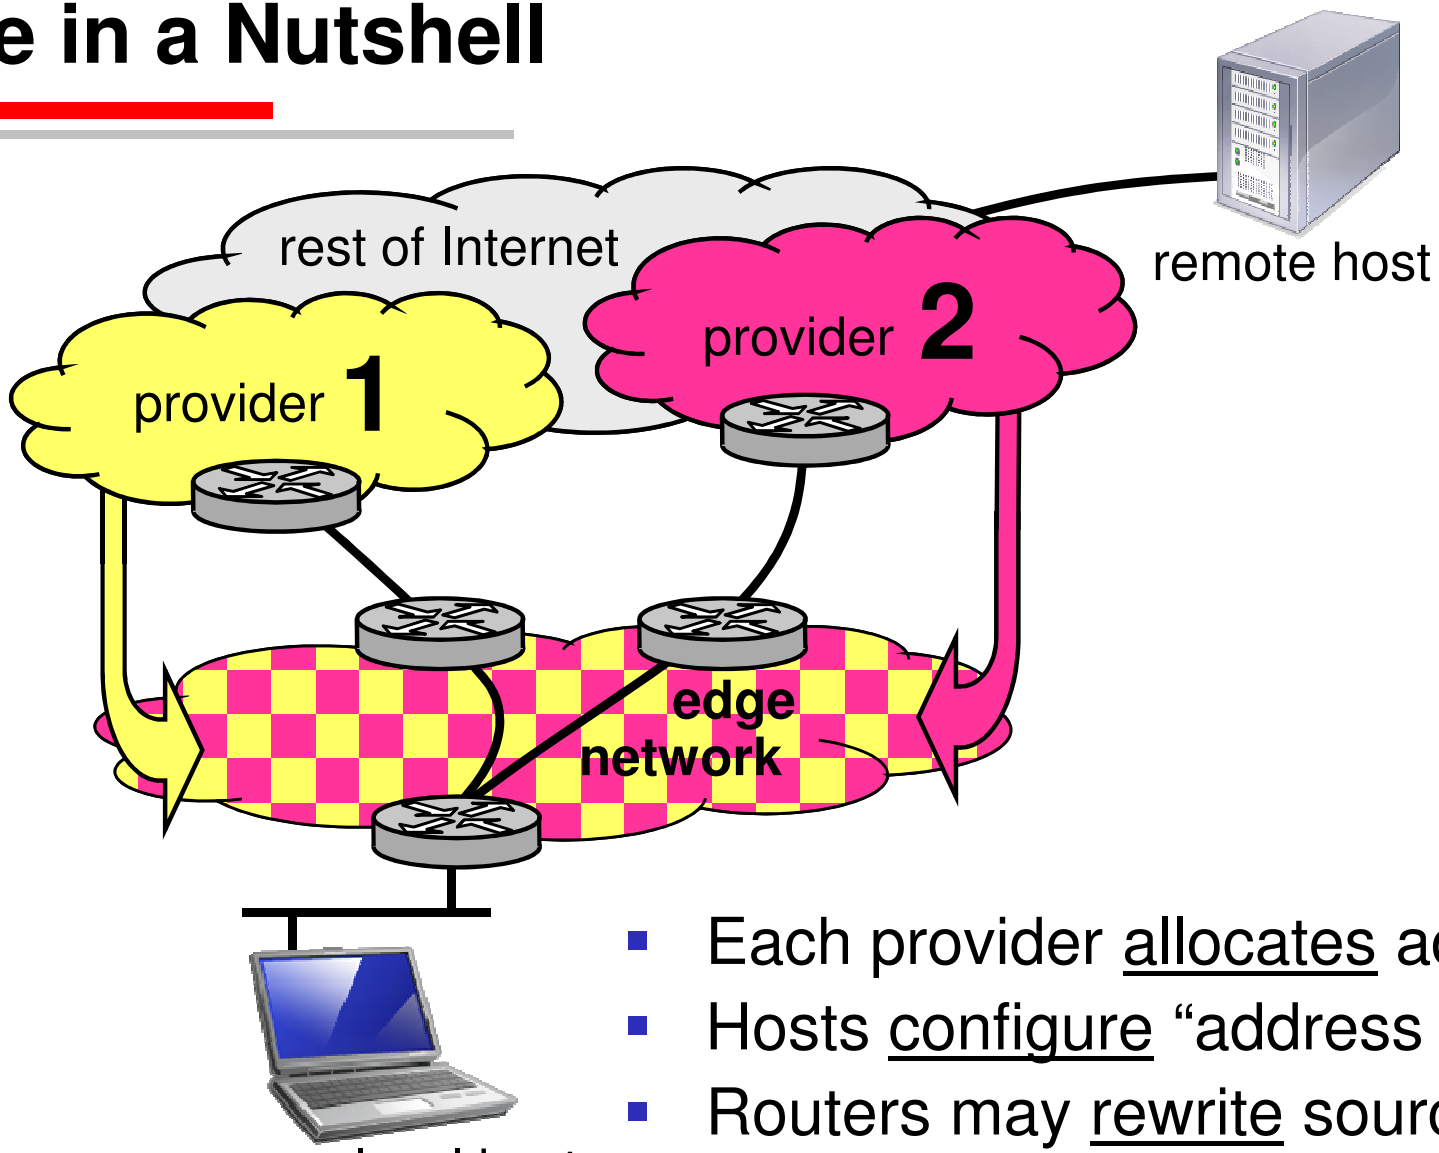

address bunch

1000:abc:1:1234:5d:cff:fe22:57c1
2000:def:2:1234:5d:cff:fe22:57c1

routing prefix | crypto interface ID

- Each provider allocates addresses
- Hosts configure "address bunch"
- Routers may rewrite source address
- Hosts recognize rewrites and adapt
  - Address mapping at hosts' IP layer
  - Address ownership verification via crypto interface ID in address bunch

Six/One in a Nutshell

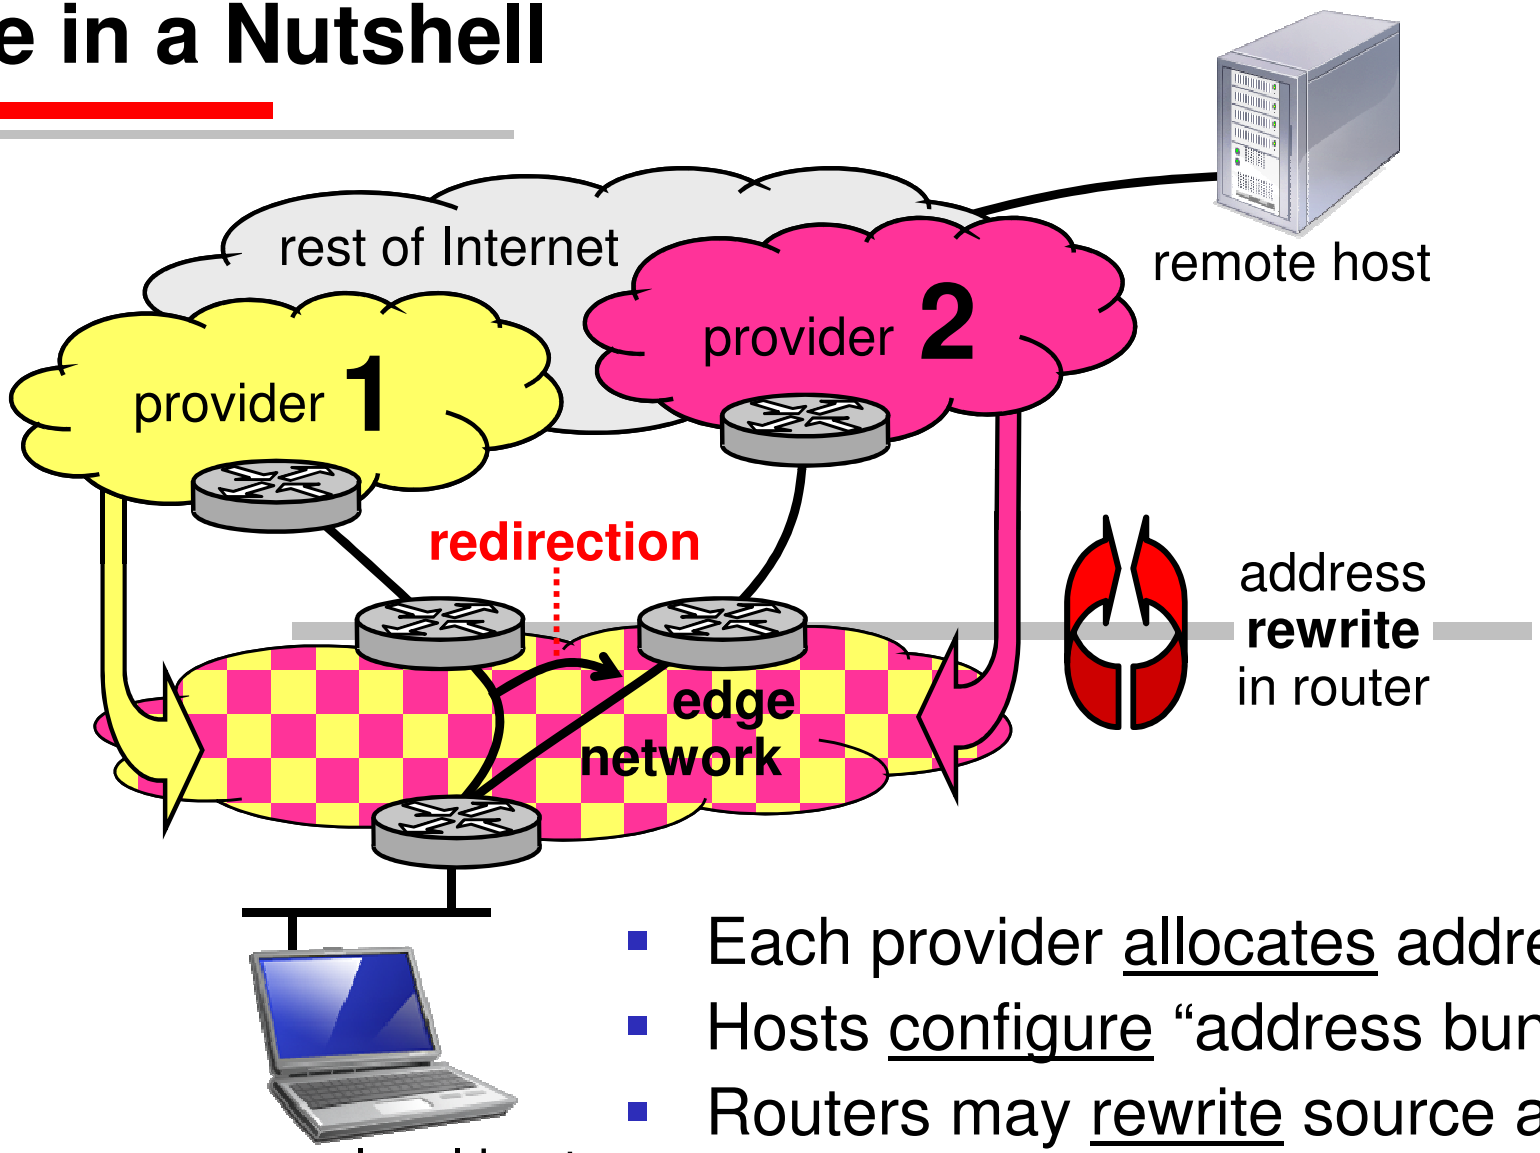

address bunch

1000:abc:1:1234:5d:cff:fe22:57c1
2000:def:2:1234:5d:cff:fe22:57c1

routing prefix — crypto interface ID

- Each provider allocates addresses
- Hosts configure "address bunch"
- Routers may rewrite source address
- Hosts recognize rewrites and adapt
  - Address mapping at hosts' IP layer
  - Address ownership verification via crypto interface ID in address bunch

Six/One in a Nutshell

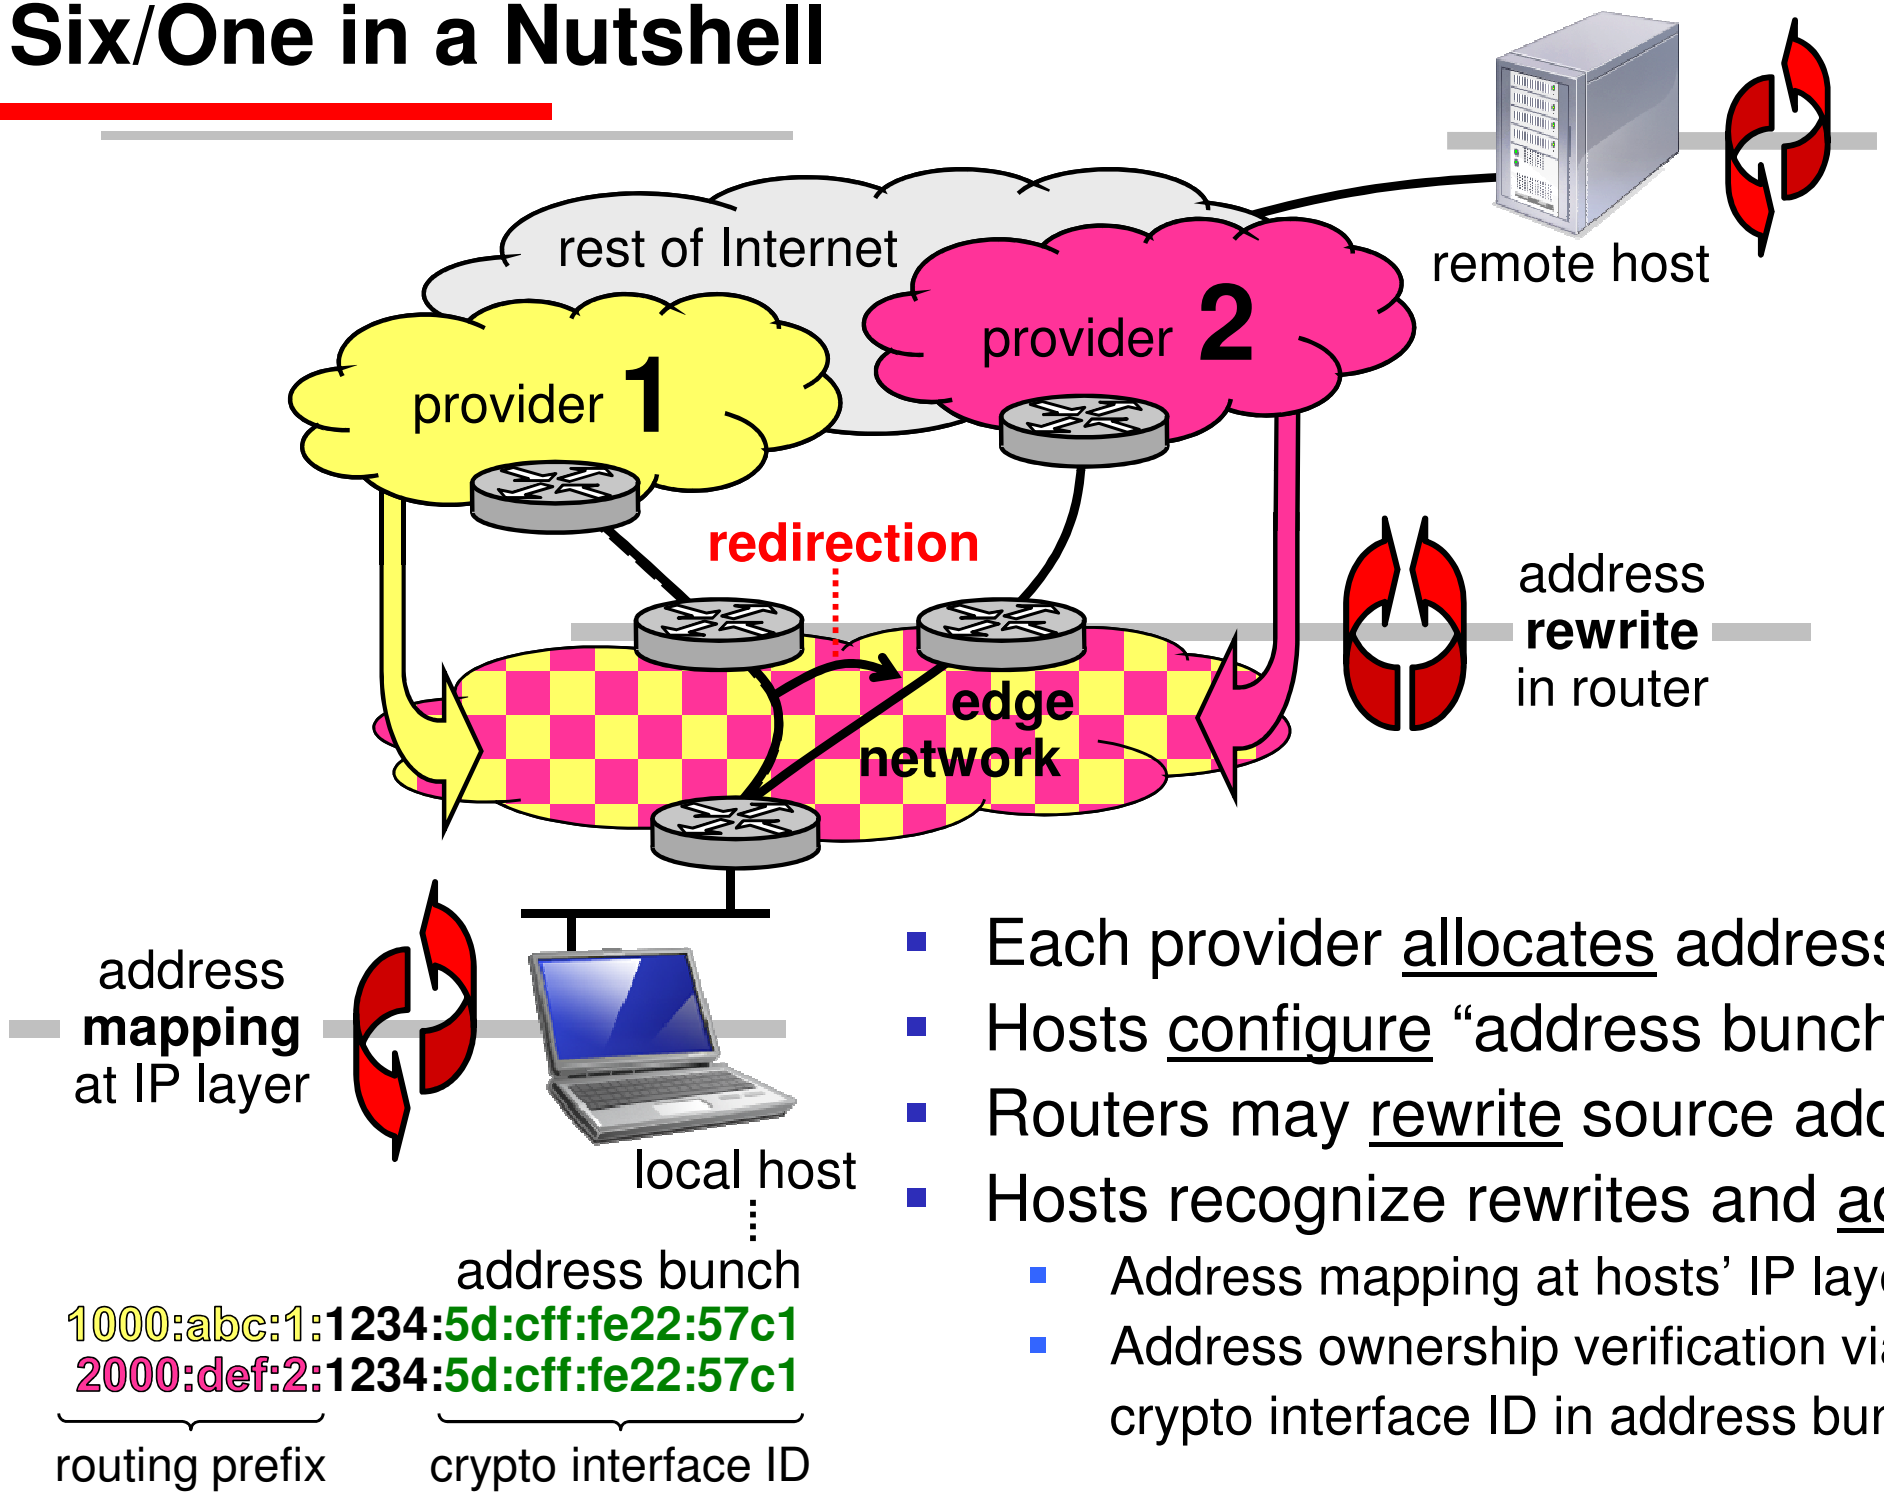

- Each provider allocates addresses
- Hosts configure "address bunch"
- Routers may rewrite source address
- Hosts recognize rewrites and adapt
  - Address mapping at hosts' IP layer
  - Address ownership verification via crypto interface ID in address bunch

Six/One: Mapping and Rewriting in Detail

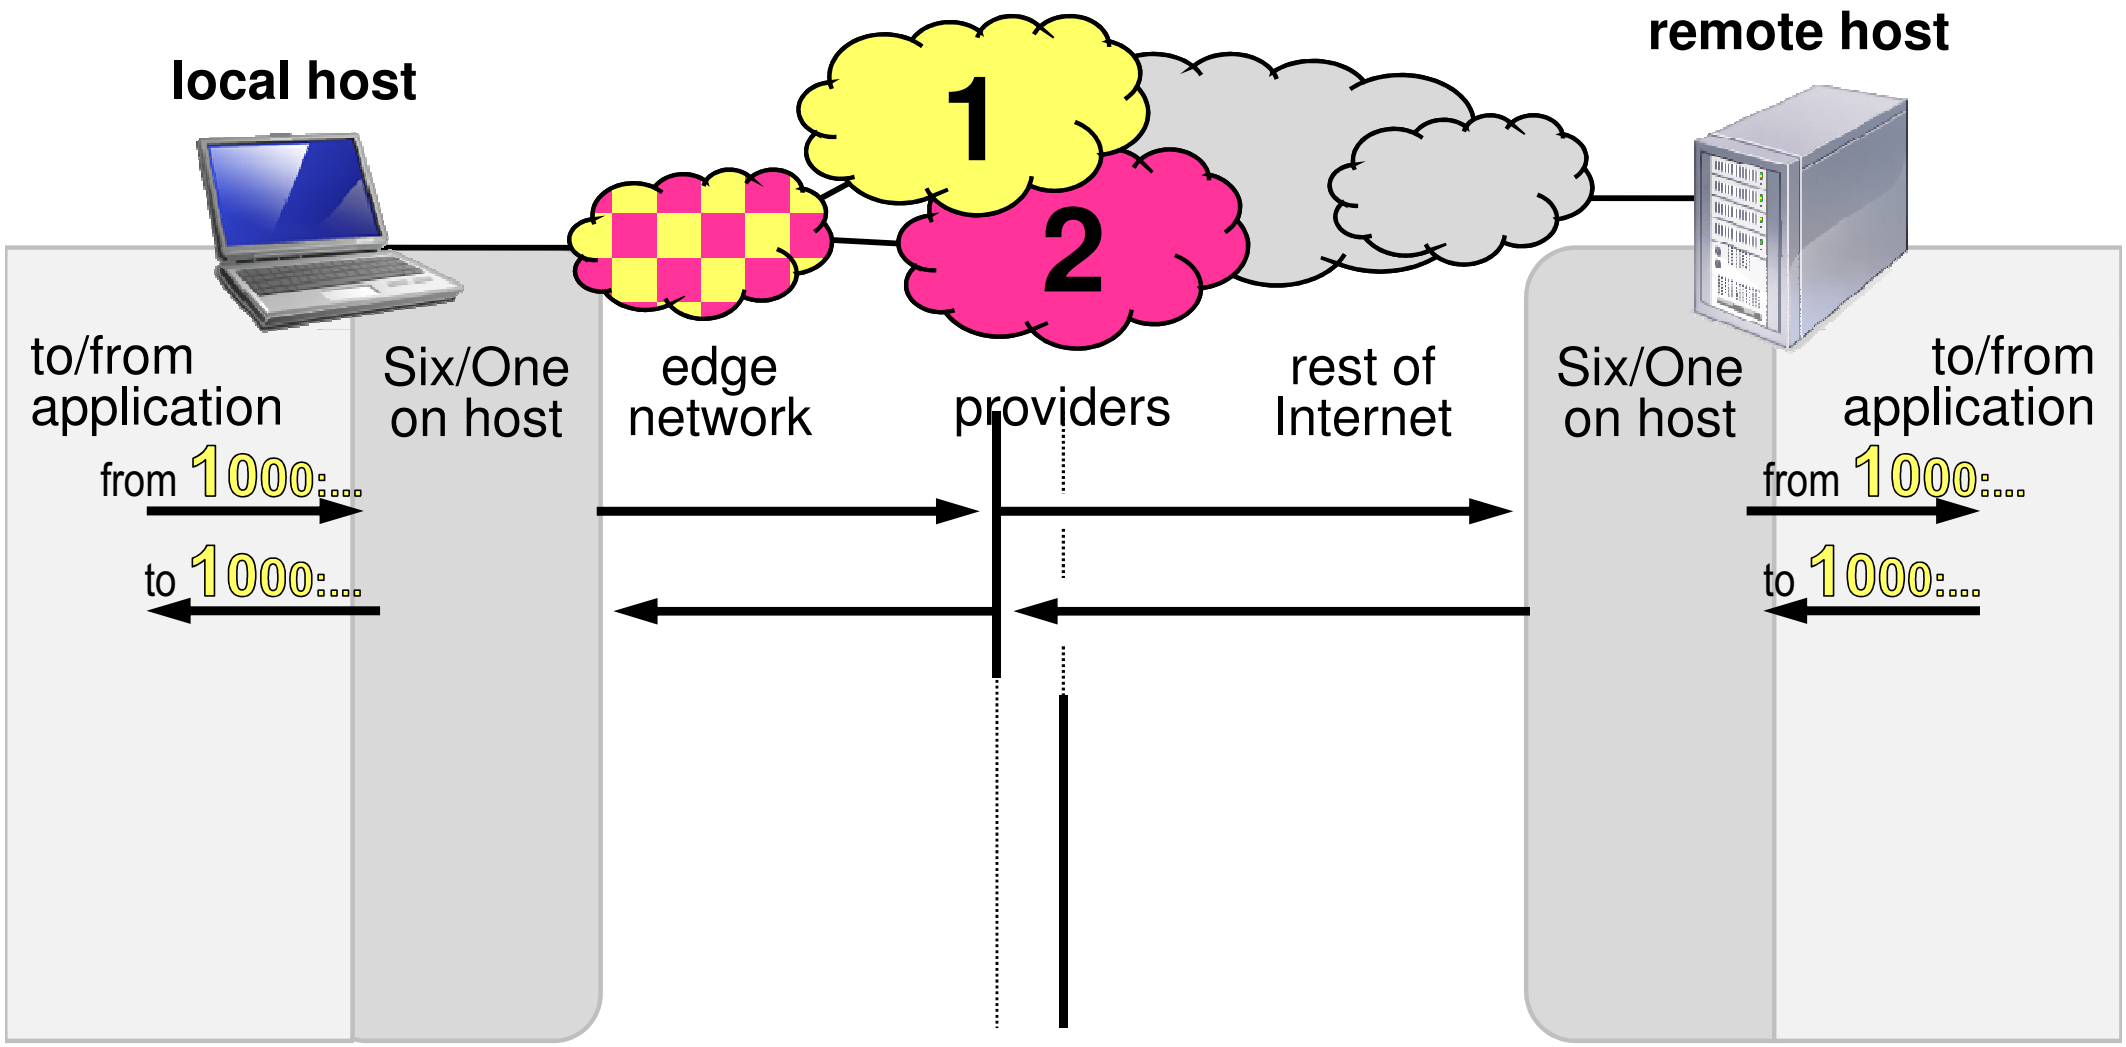

Case 1: no rewriting

- Host selects source address
- It thereby suggests provider
- Routers accept host selection

Six/One: Mapping and Rewriting in Detail

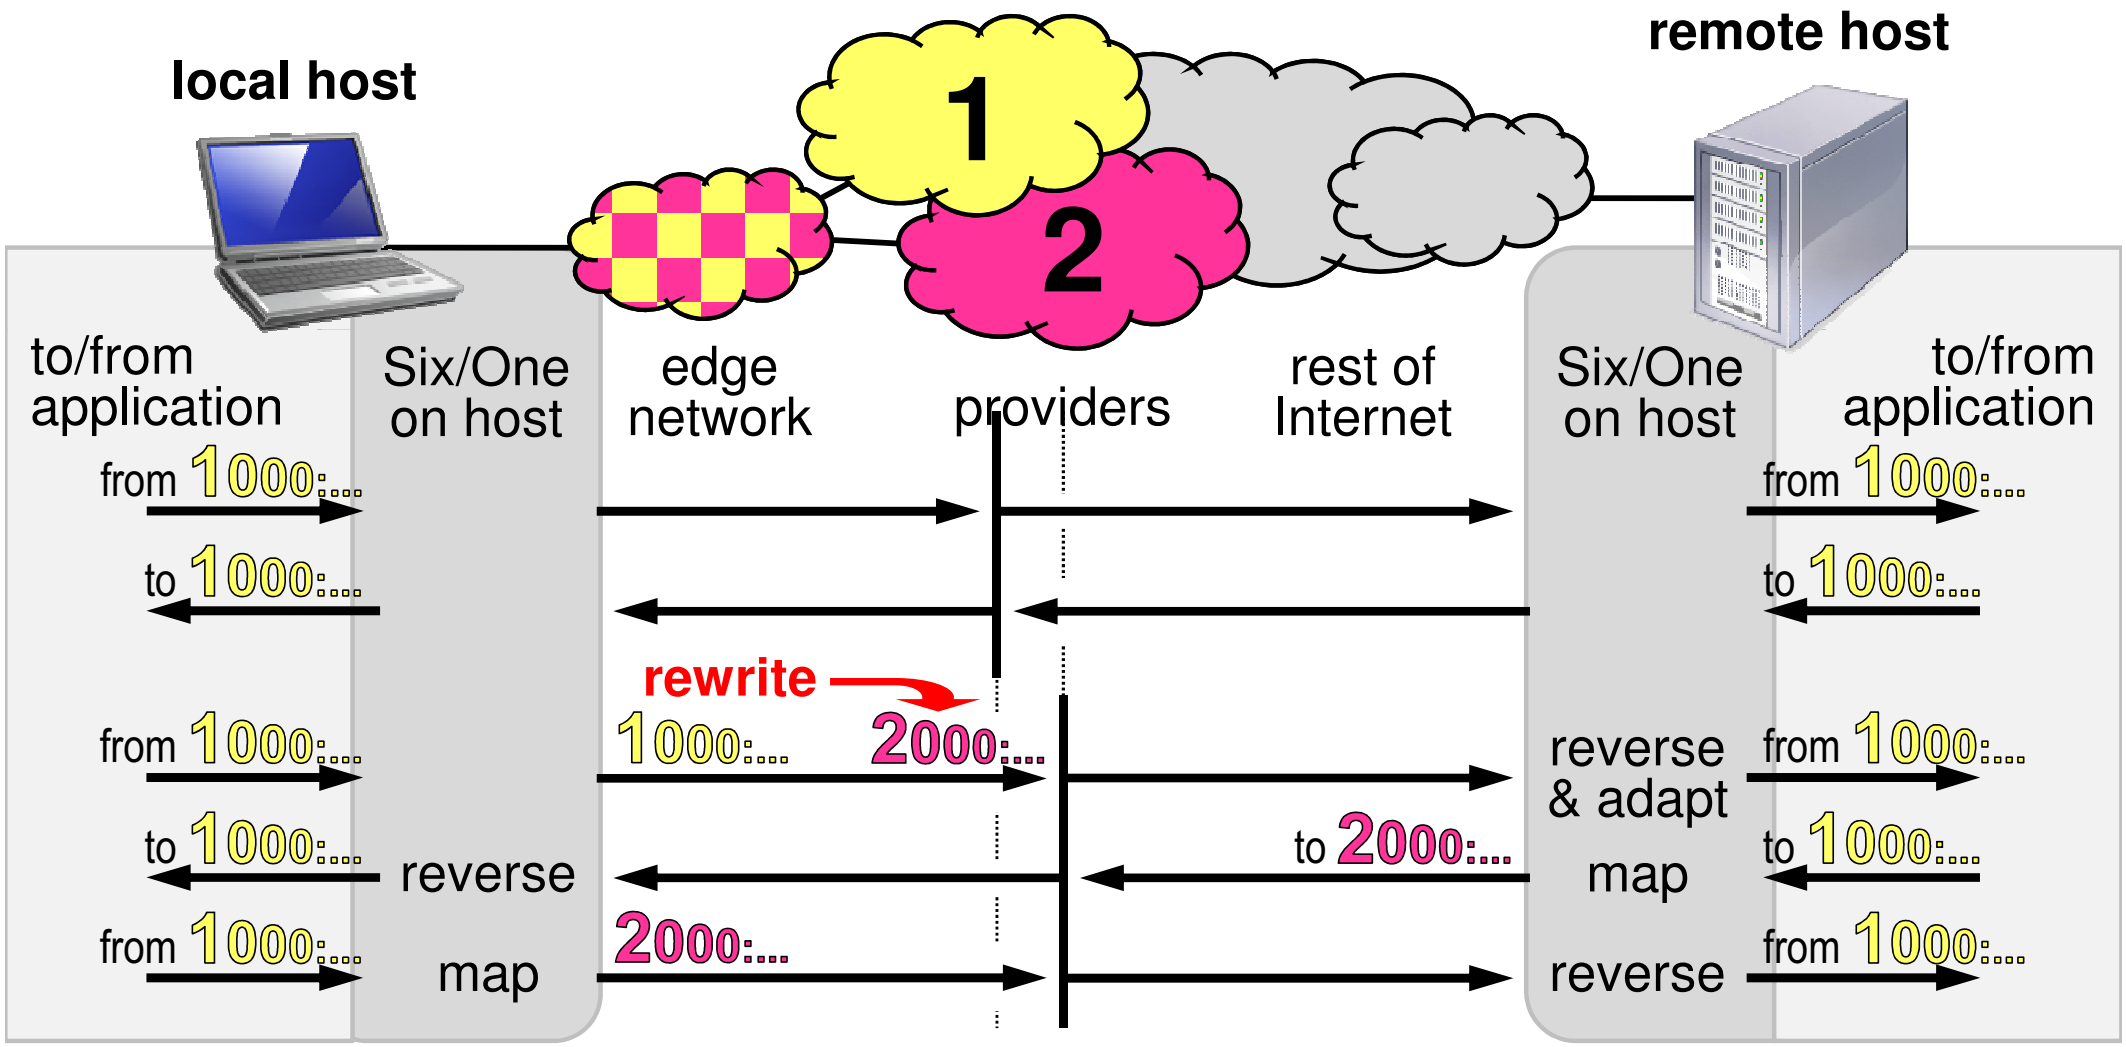

Case 1: no rewriting

- Host selects source address
- It thereby suggests provider
- Routers accept host selection

Case 2: rewriting in edge network

- Router rewrites source address
- Hosts learn new address and adapt
- No address change in application

Six/One: Context Establishment

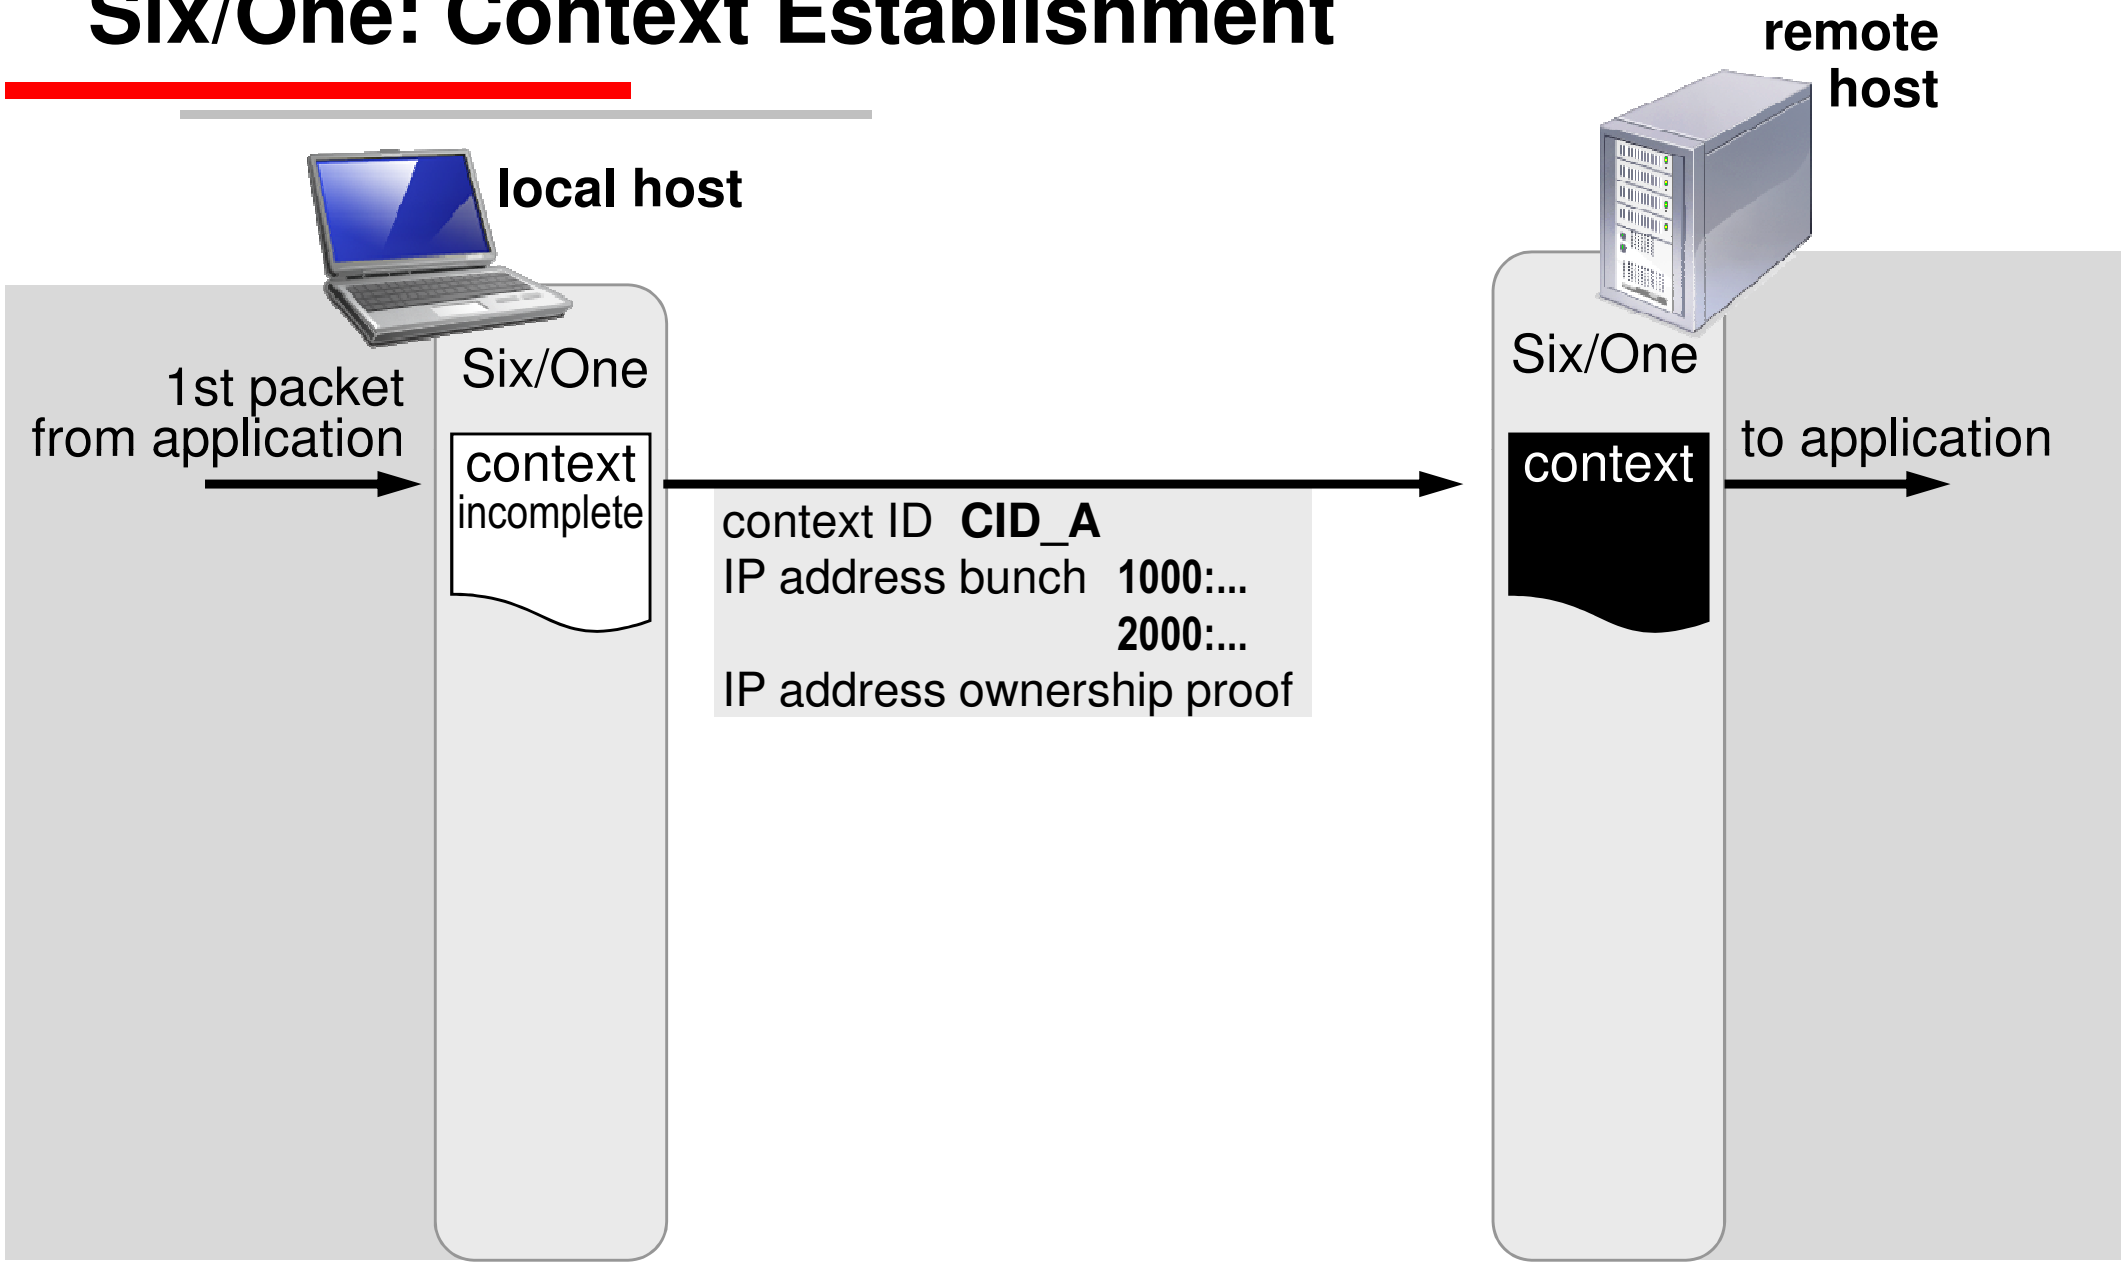

- Context establishment when hosts initiate first communication session
- Context IDs for subsequent look-up

Six/One: Context Establishment

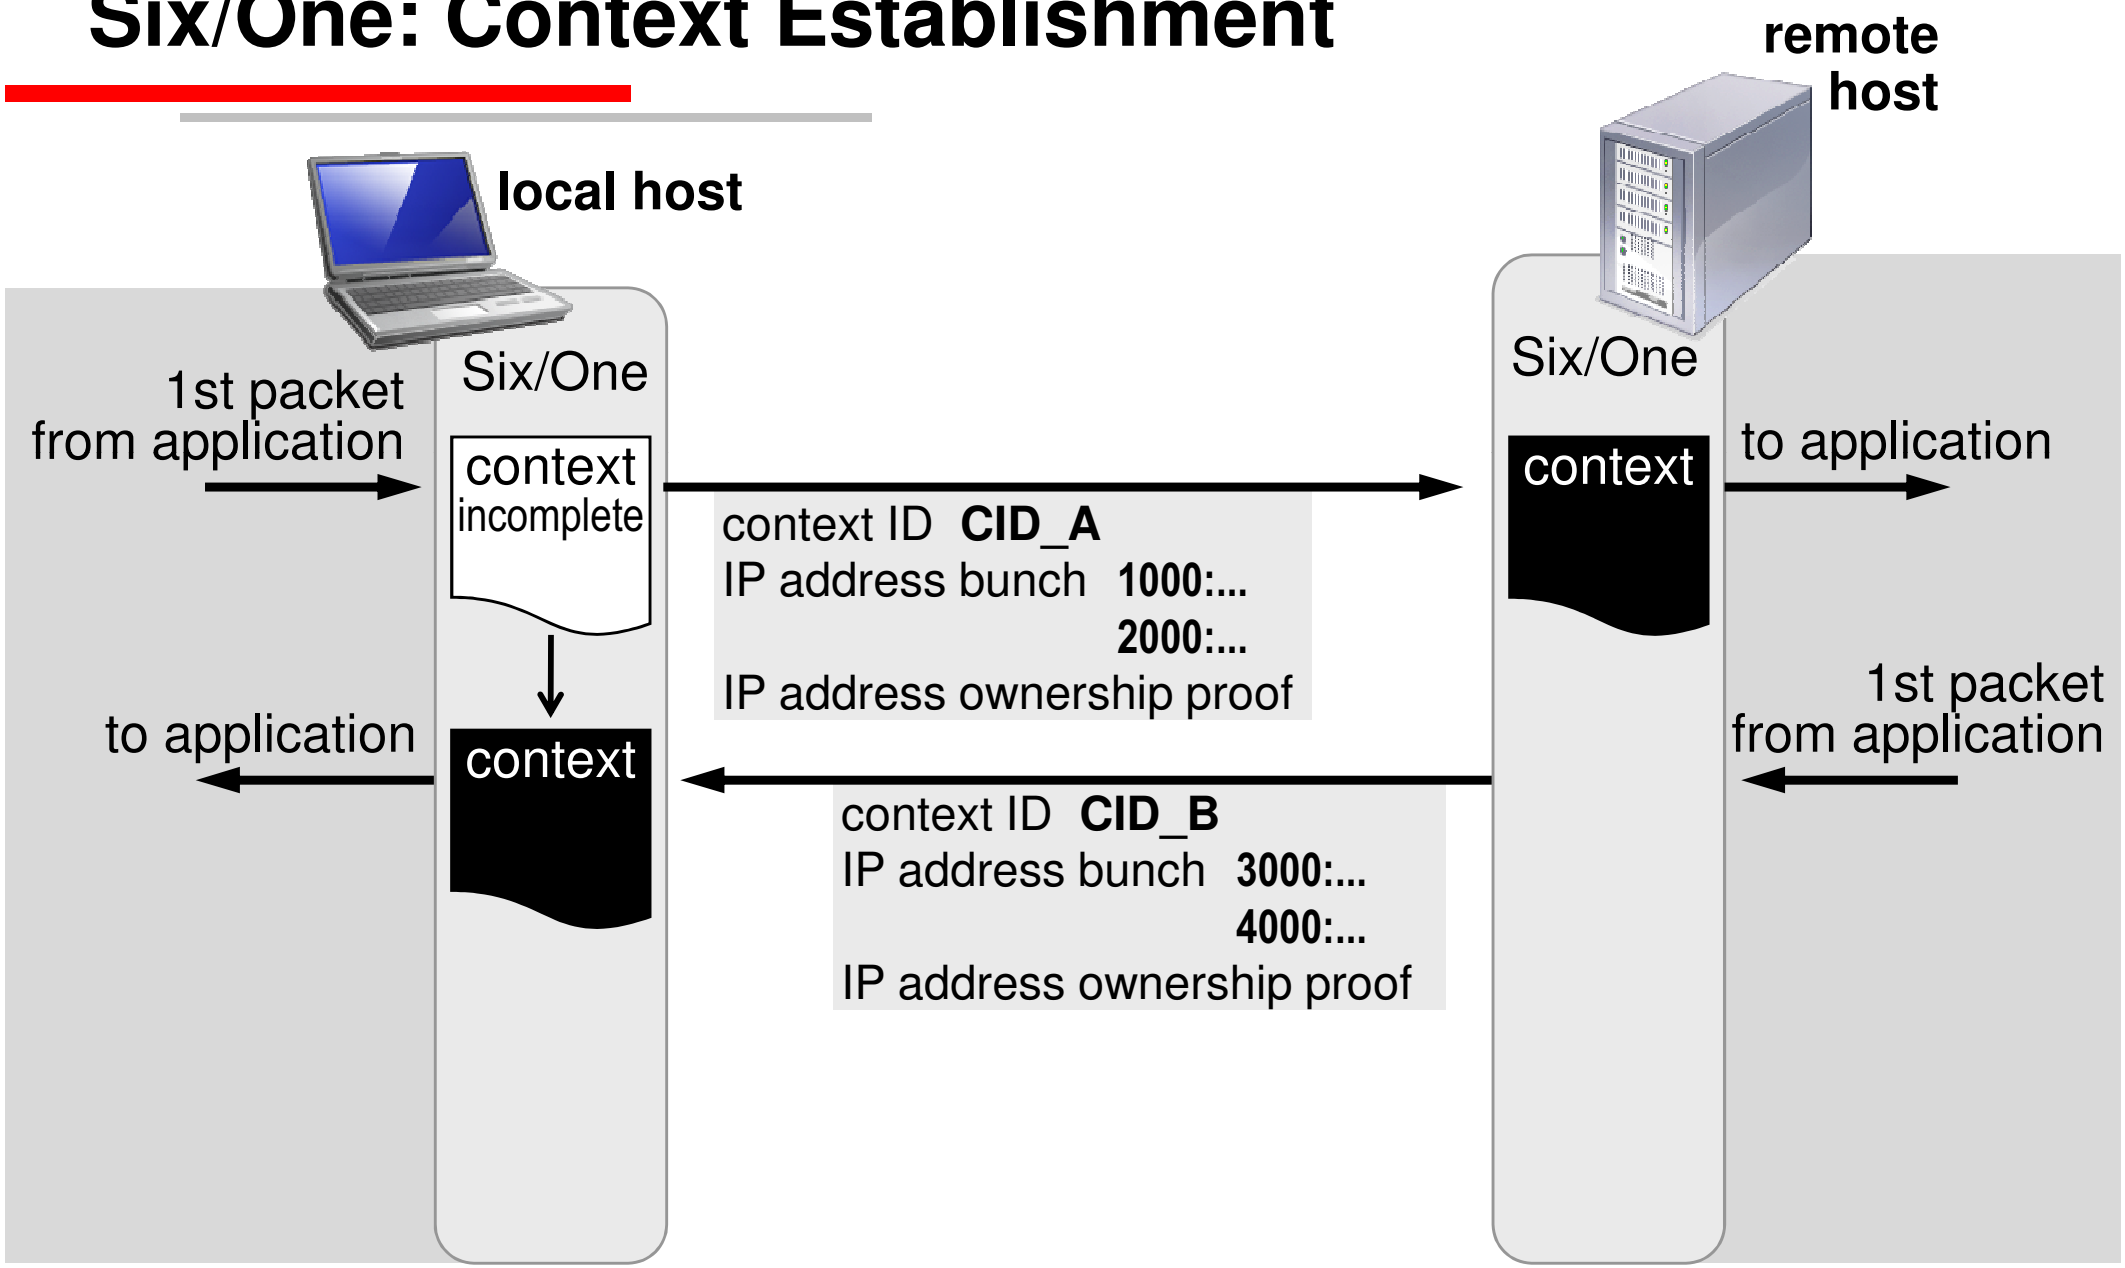

- Context establishment when hosts initiate first communication session
- Context IDs for subsequent look-up

Six/One: Context Establishment

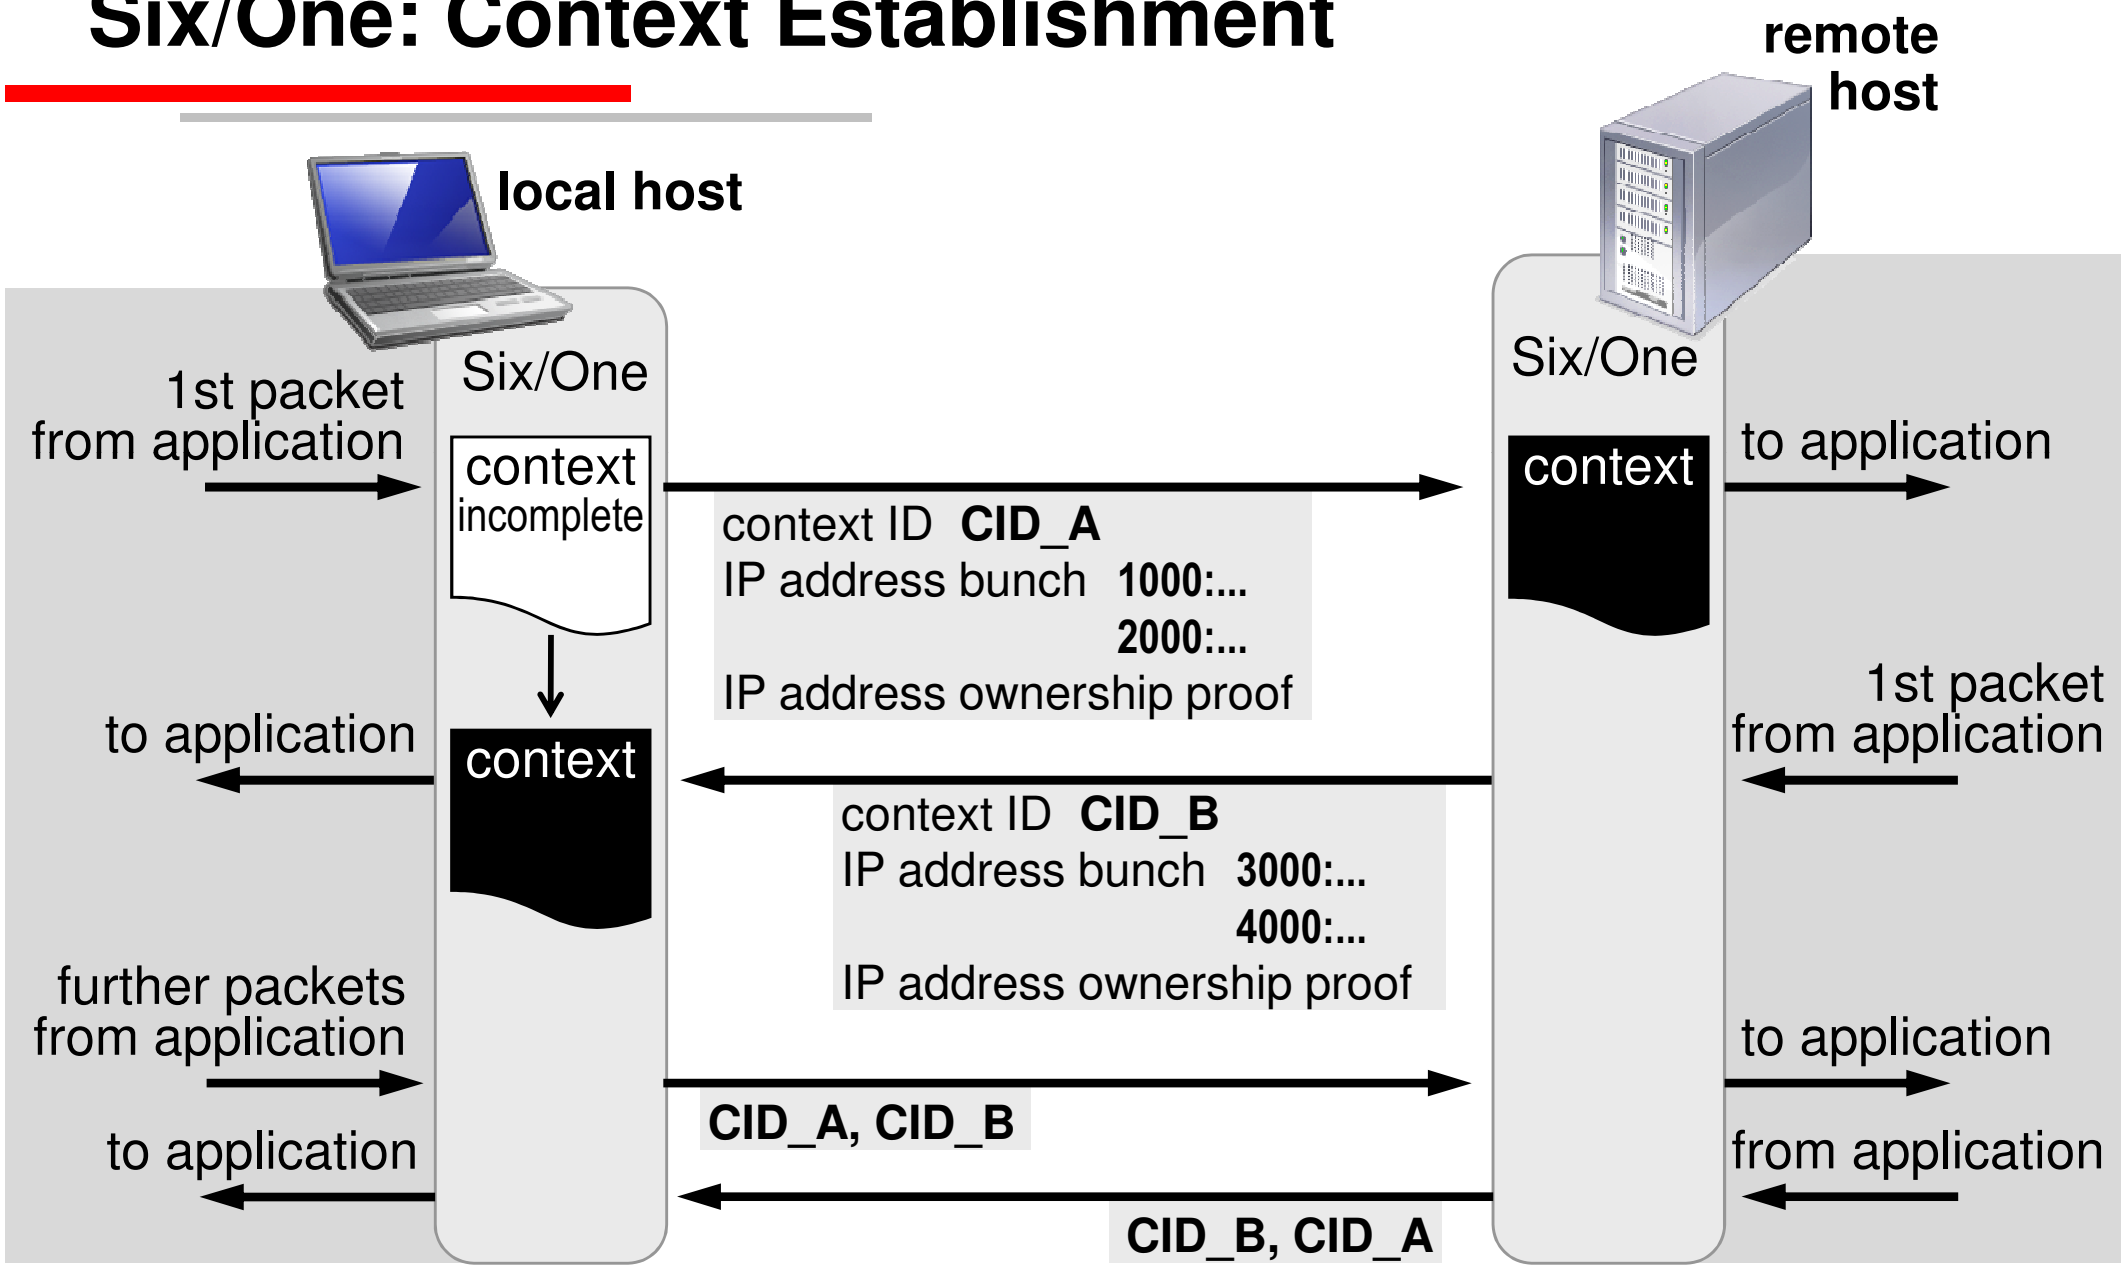

- Context establishment when hosts initiate first communication session
- Context IDs for subsequent look-up
- Routers do not rewrite IP address prefixes before context established

Review

		Shim6 host-controlled address mapping	LISP network-controlled address mapping	Six/One address mapping by host+network
scalability	reduced size of routing table	all IP addresses aggregatable	aggregatable after mapping	all IP addresses aggregatable
	fewer updates of routing table	✓	✓	✓
traffic engineering compatibility	reliable fail-over	end-to-end connectivity checks	edge-to-edge connectivity checks	end-to-end connectivity checks
	load-balancing	IP addresses provider-dependent	provider-independent before mapping	address rewriting by routers
incremental deployability	incremental deployability	no conflicts with legacy Internet	edge addresses unreachable	no address rewrite for legacy hosts
	incentives for early adoption	if support at both ends	edge addresses unreachable	if support at both ends

Extensions

1. **Proxy support**
   - Six/One deployment without host upgrades
   - Increased deployment flexibility

2. **Address translation**
   - Redirection without Six/One support on remote side
   - Increased benefits

Proxy Support

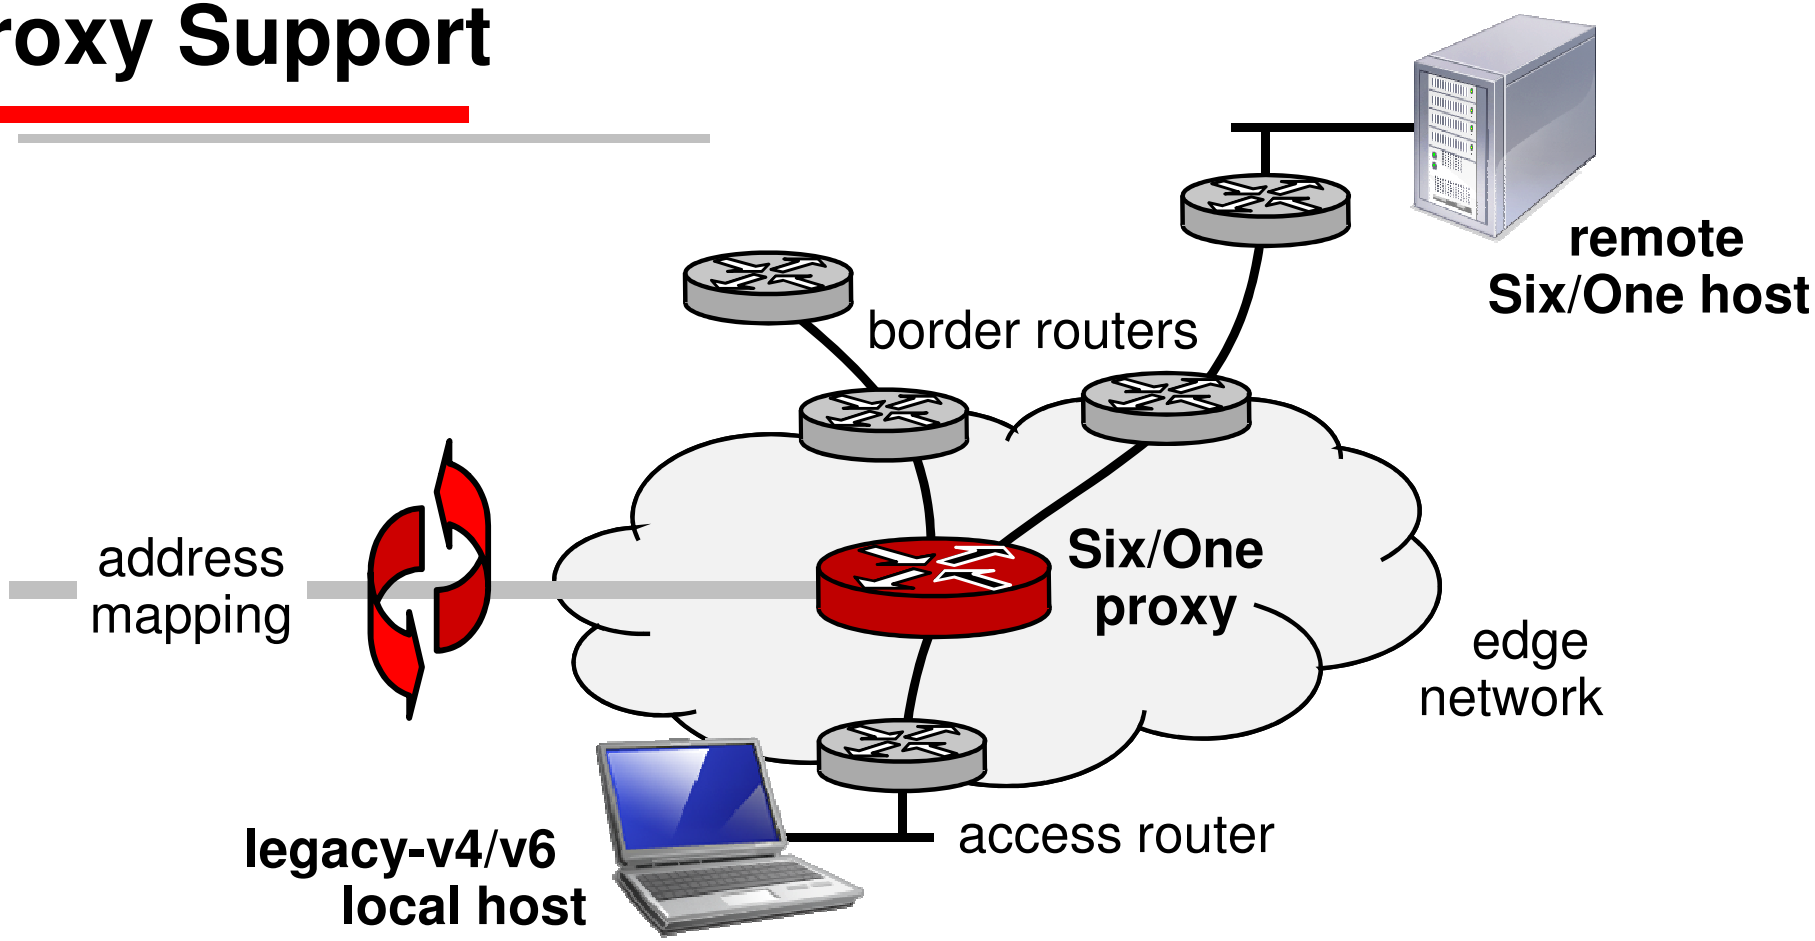

- Six/One deployment without upgrades to local hosts
- Increased deployment flexibility
  - Don't wait for host upgrades
  - For hosts with limited capabilities
- Designed for interoperability
  - Transparent to local legacy hosts
  - Works with local Six/One hosts
  - Interoperable with remote Six/One hosts and proxies
- Free placement of proxies

Proxy Operation for Local Legacy-v6 Host

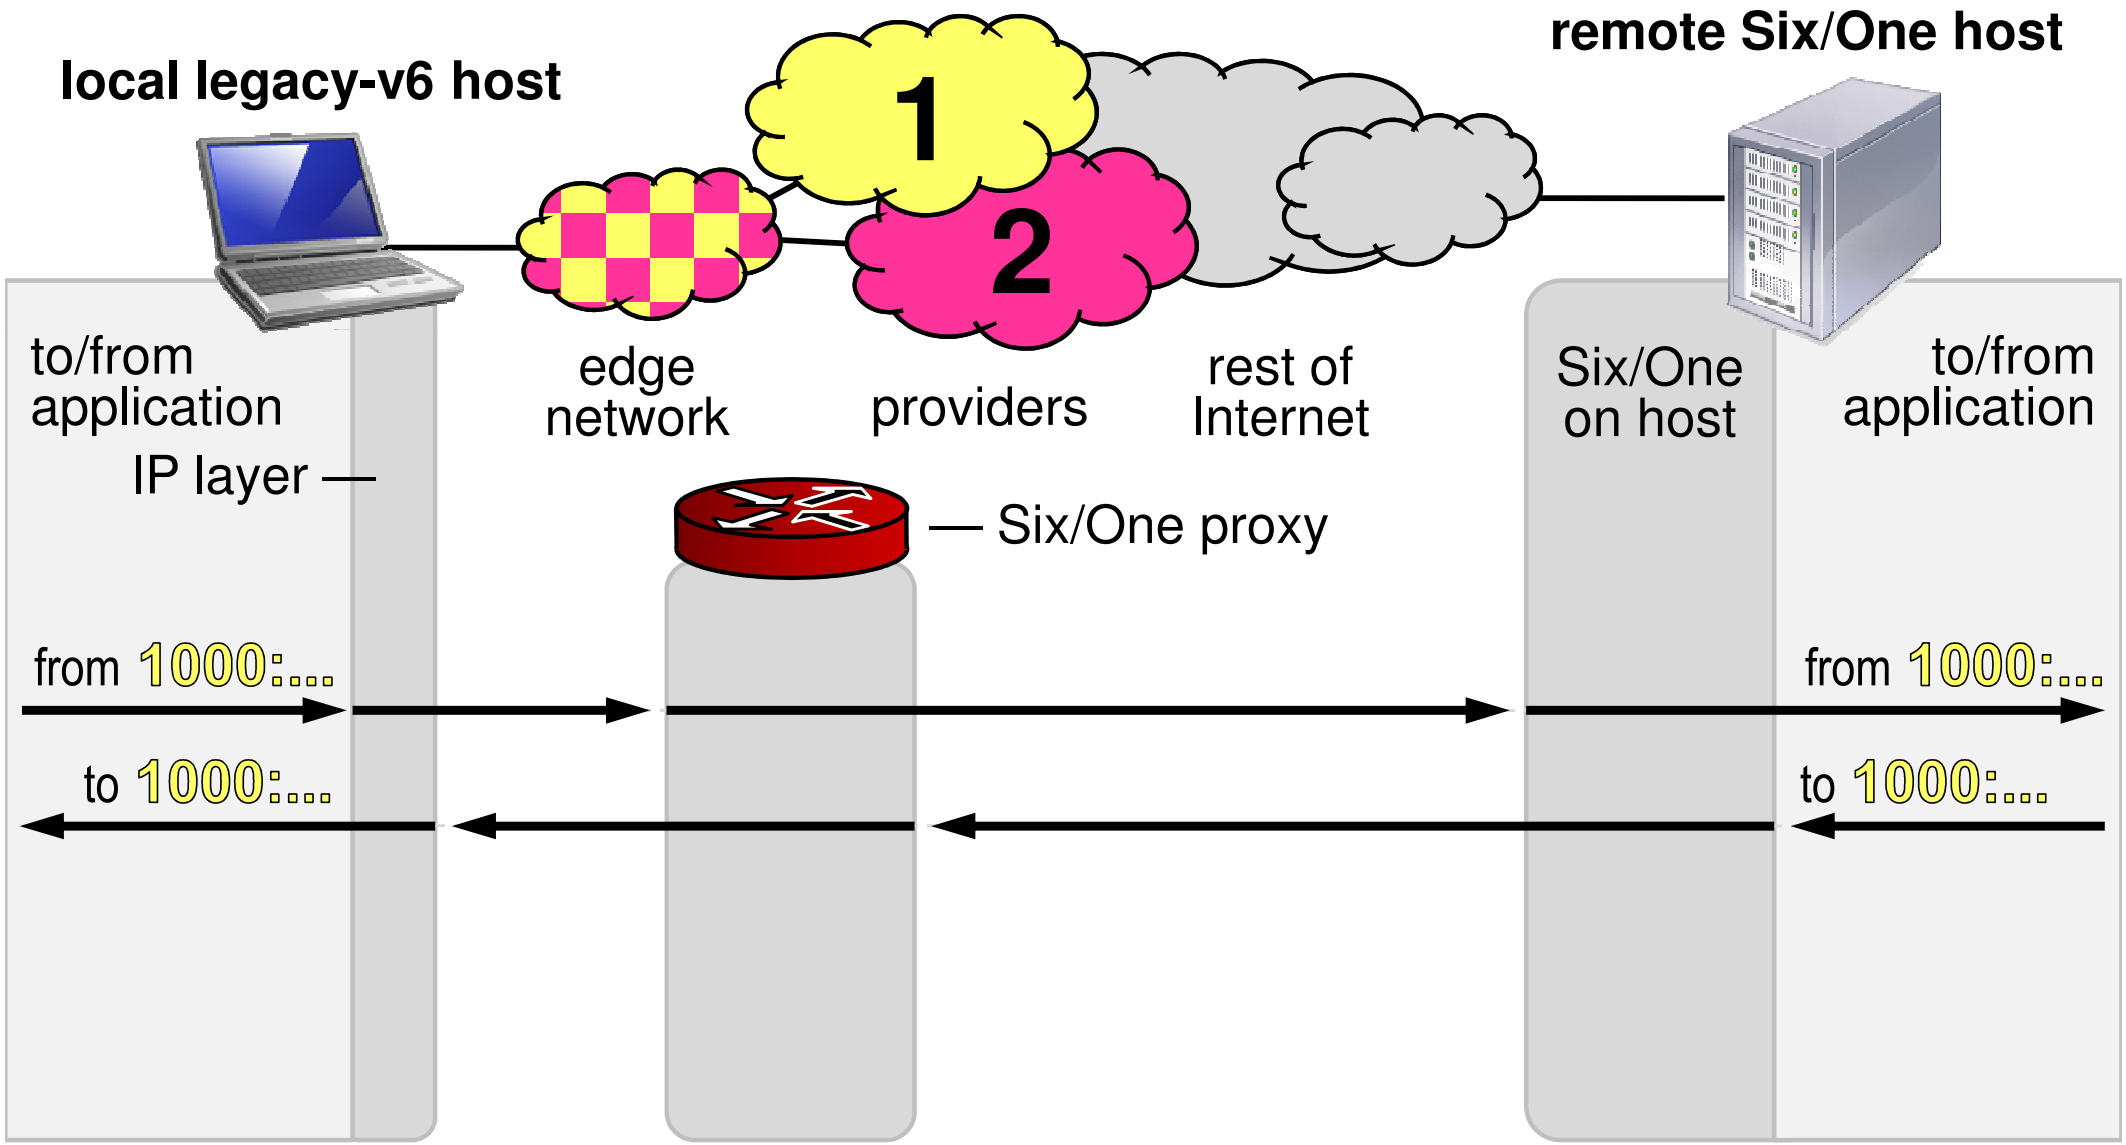

address bunch of local host
1000:abc:1:**1234:5d:cff:fe22:57c1**
2000:def:2:**1234:5d:cff:fe22:57c1**

Proxy Operation for Local Legacy-v6 Host

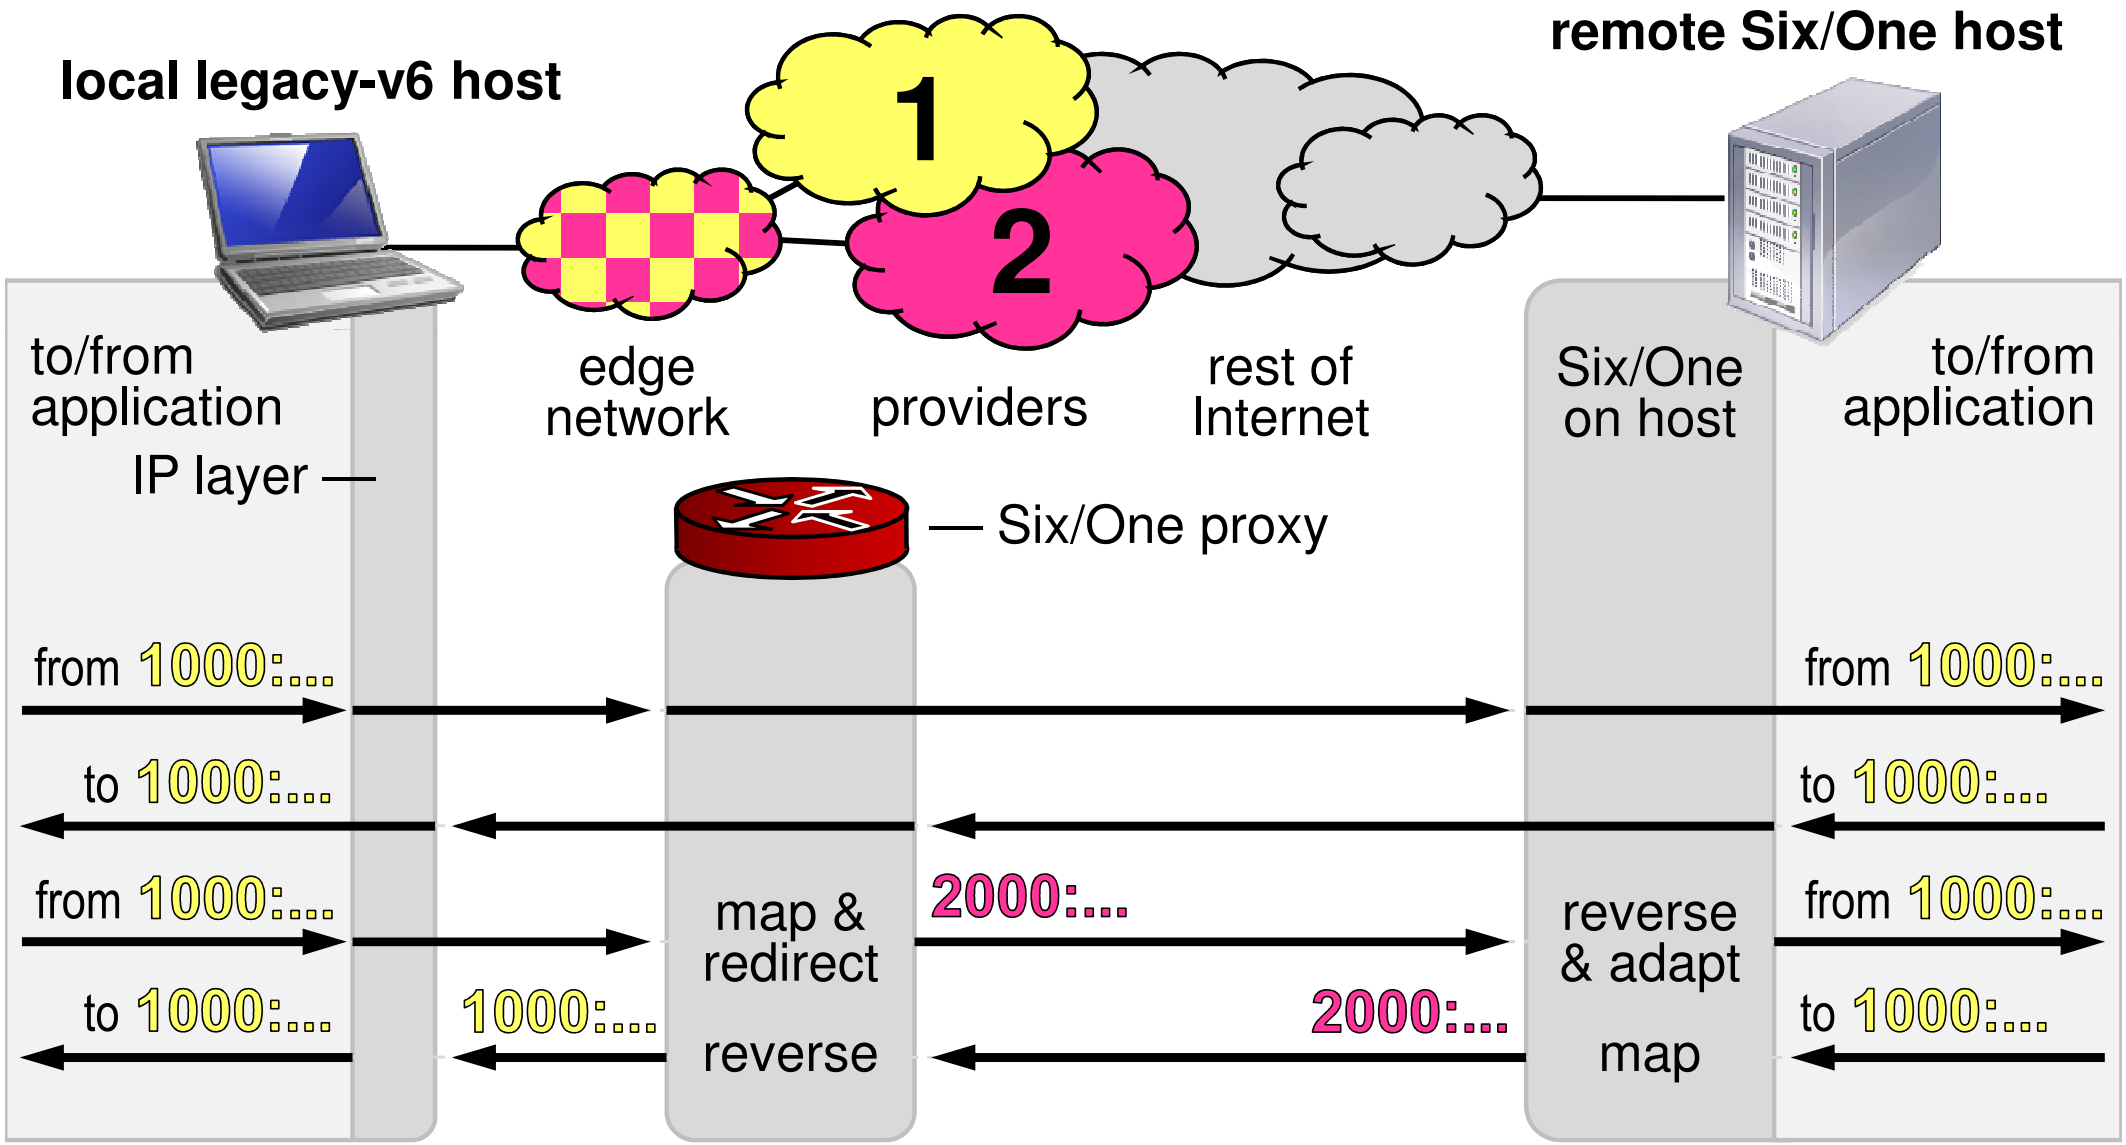

address bunch of local host

1000:abc:1:**1234:5d:cff:fe22:57c1**

2000:def:2:**1234:5d:cff:fe22:57c1**

Address Bunch Configuration with Proxies

- DHCP server generates address bunch
  - Crypto parameters made by DHCP server
  - Proxy retrieves crypto parameters from DHCP server
- Flexible support for Six/One-upgraded hosts
  - Disable by mandating standard DHCP
  - Permit by allowing Six/One hosts to autonomously configure address bunches in Router Advertisements
  - Permit by communicating crypto parameters to hosts, e.g., via new DHCP option

Address Translation

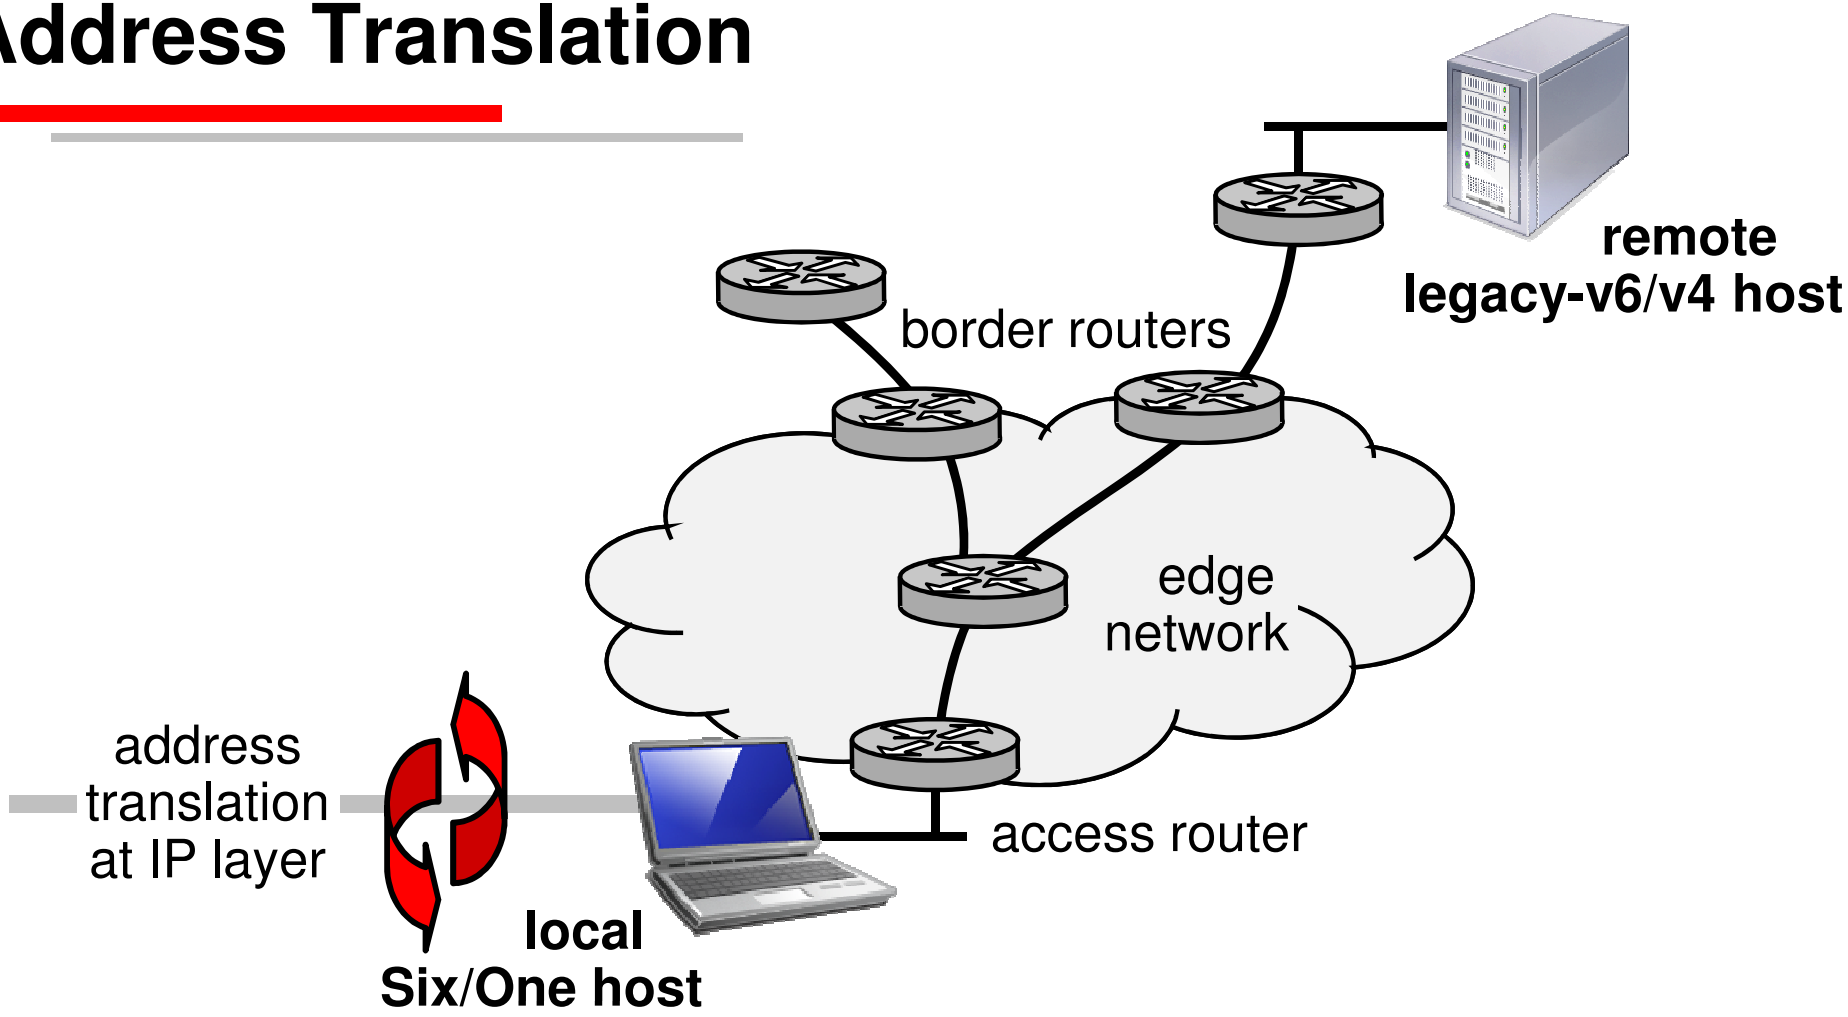

- Redirection without Six/One support on remote side
- Like a NAT: Replace addresses in packets without reversal
  - But addresses are of same host

- Address translation OK with or without reverse mapping
  - Static redirection of all packets if no reverse mapping
  - Dynamic redirection only with reverse mapping

Translator Operation in Detail

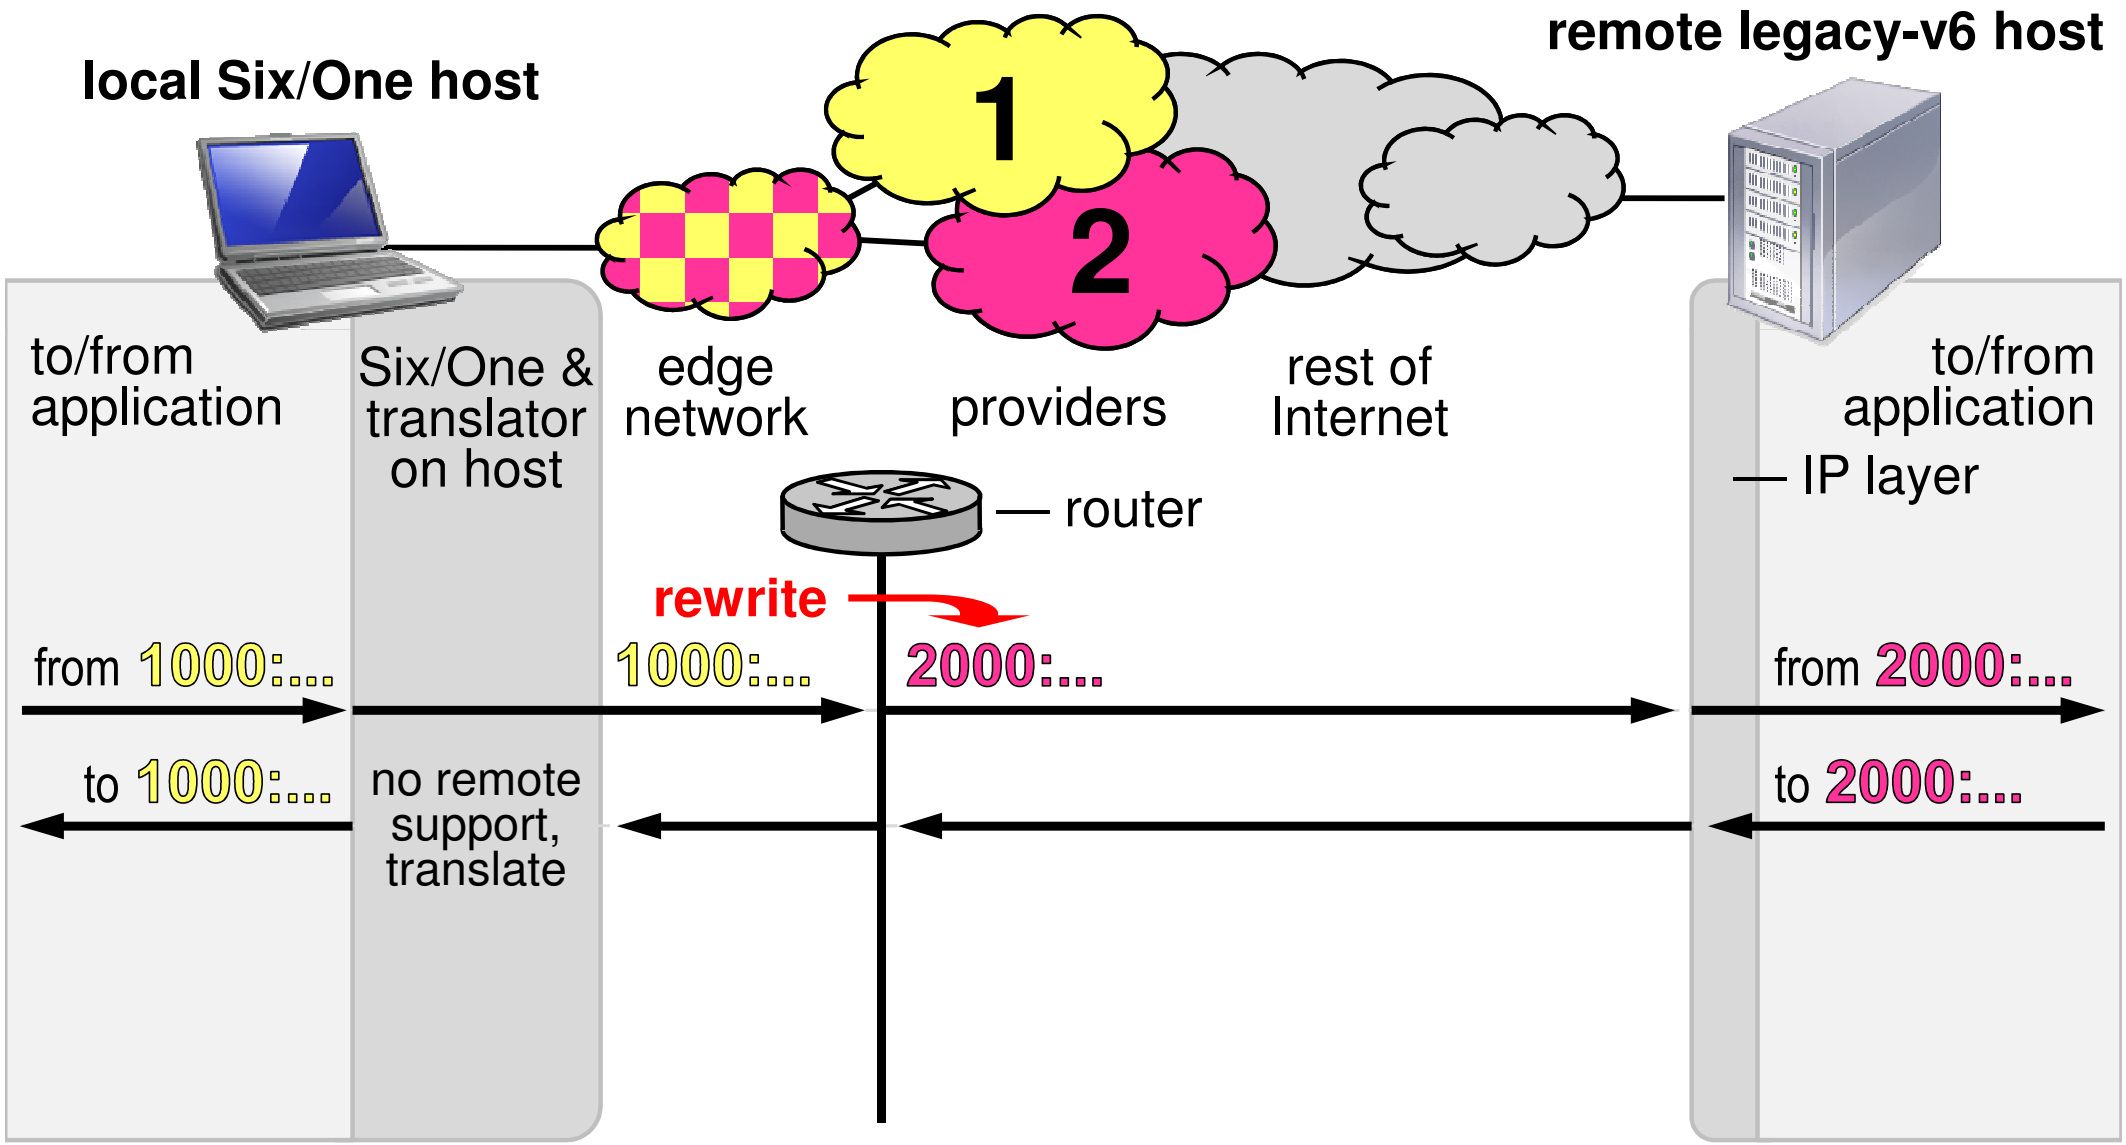

address bunch of local host

1000:abc:1:**1234:**5d:cff:fe22:57c1

2000:def:2:**1234:**5d:cff:fe22:57c1

Translator Operation in Detail

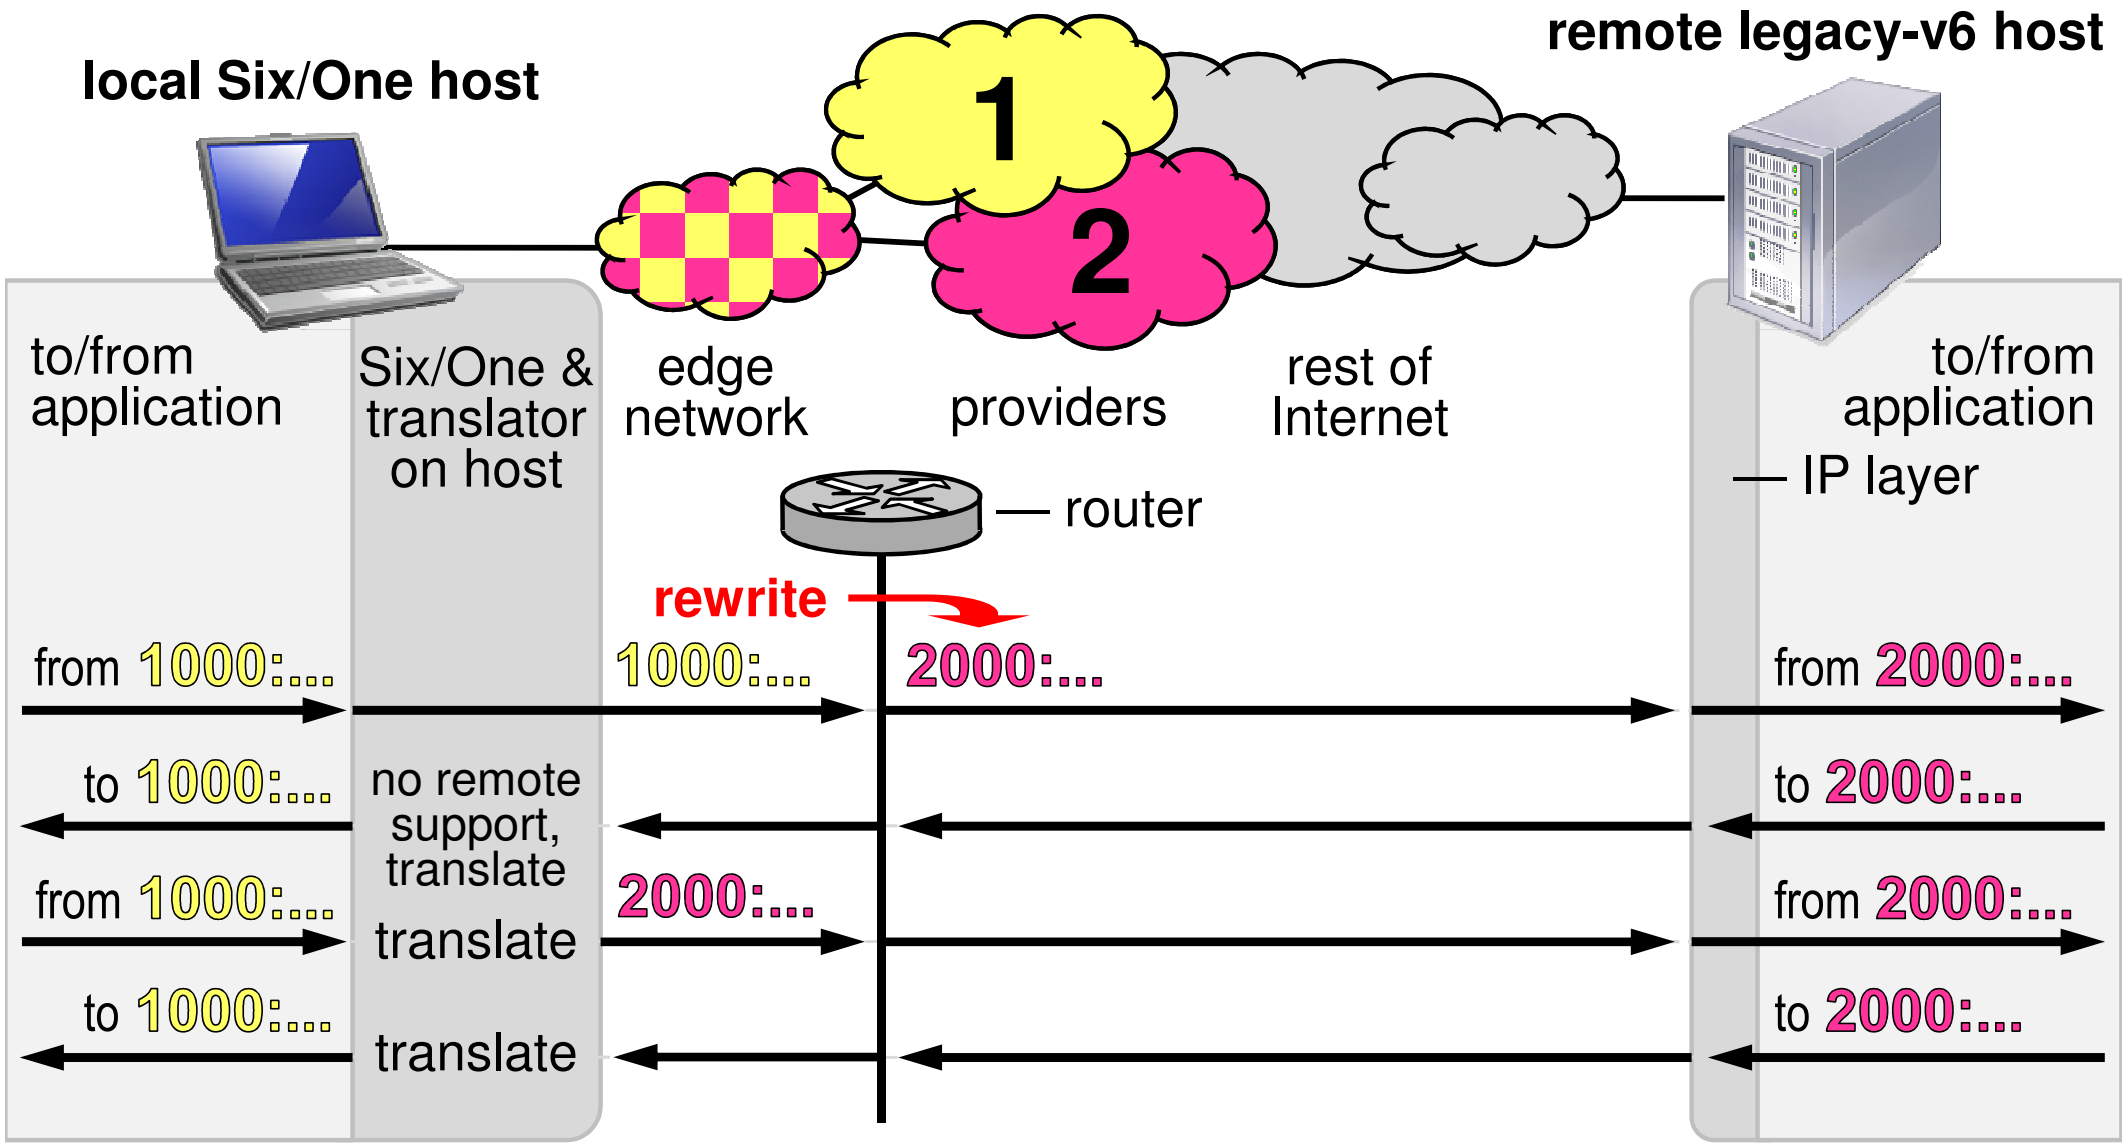

address bunch of local host

1000:abc:1:**1234**:5d:cff:fe22:57c1

2000:def:2:**1234**:5d:cff:fe22:57c1

Combined Proxy and Translator for v4 Support

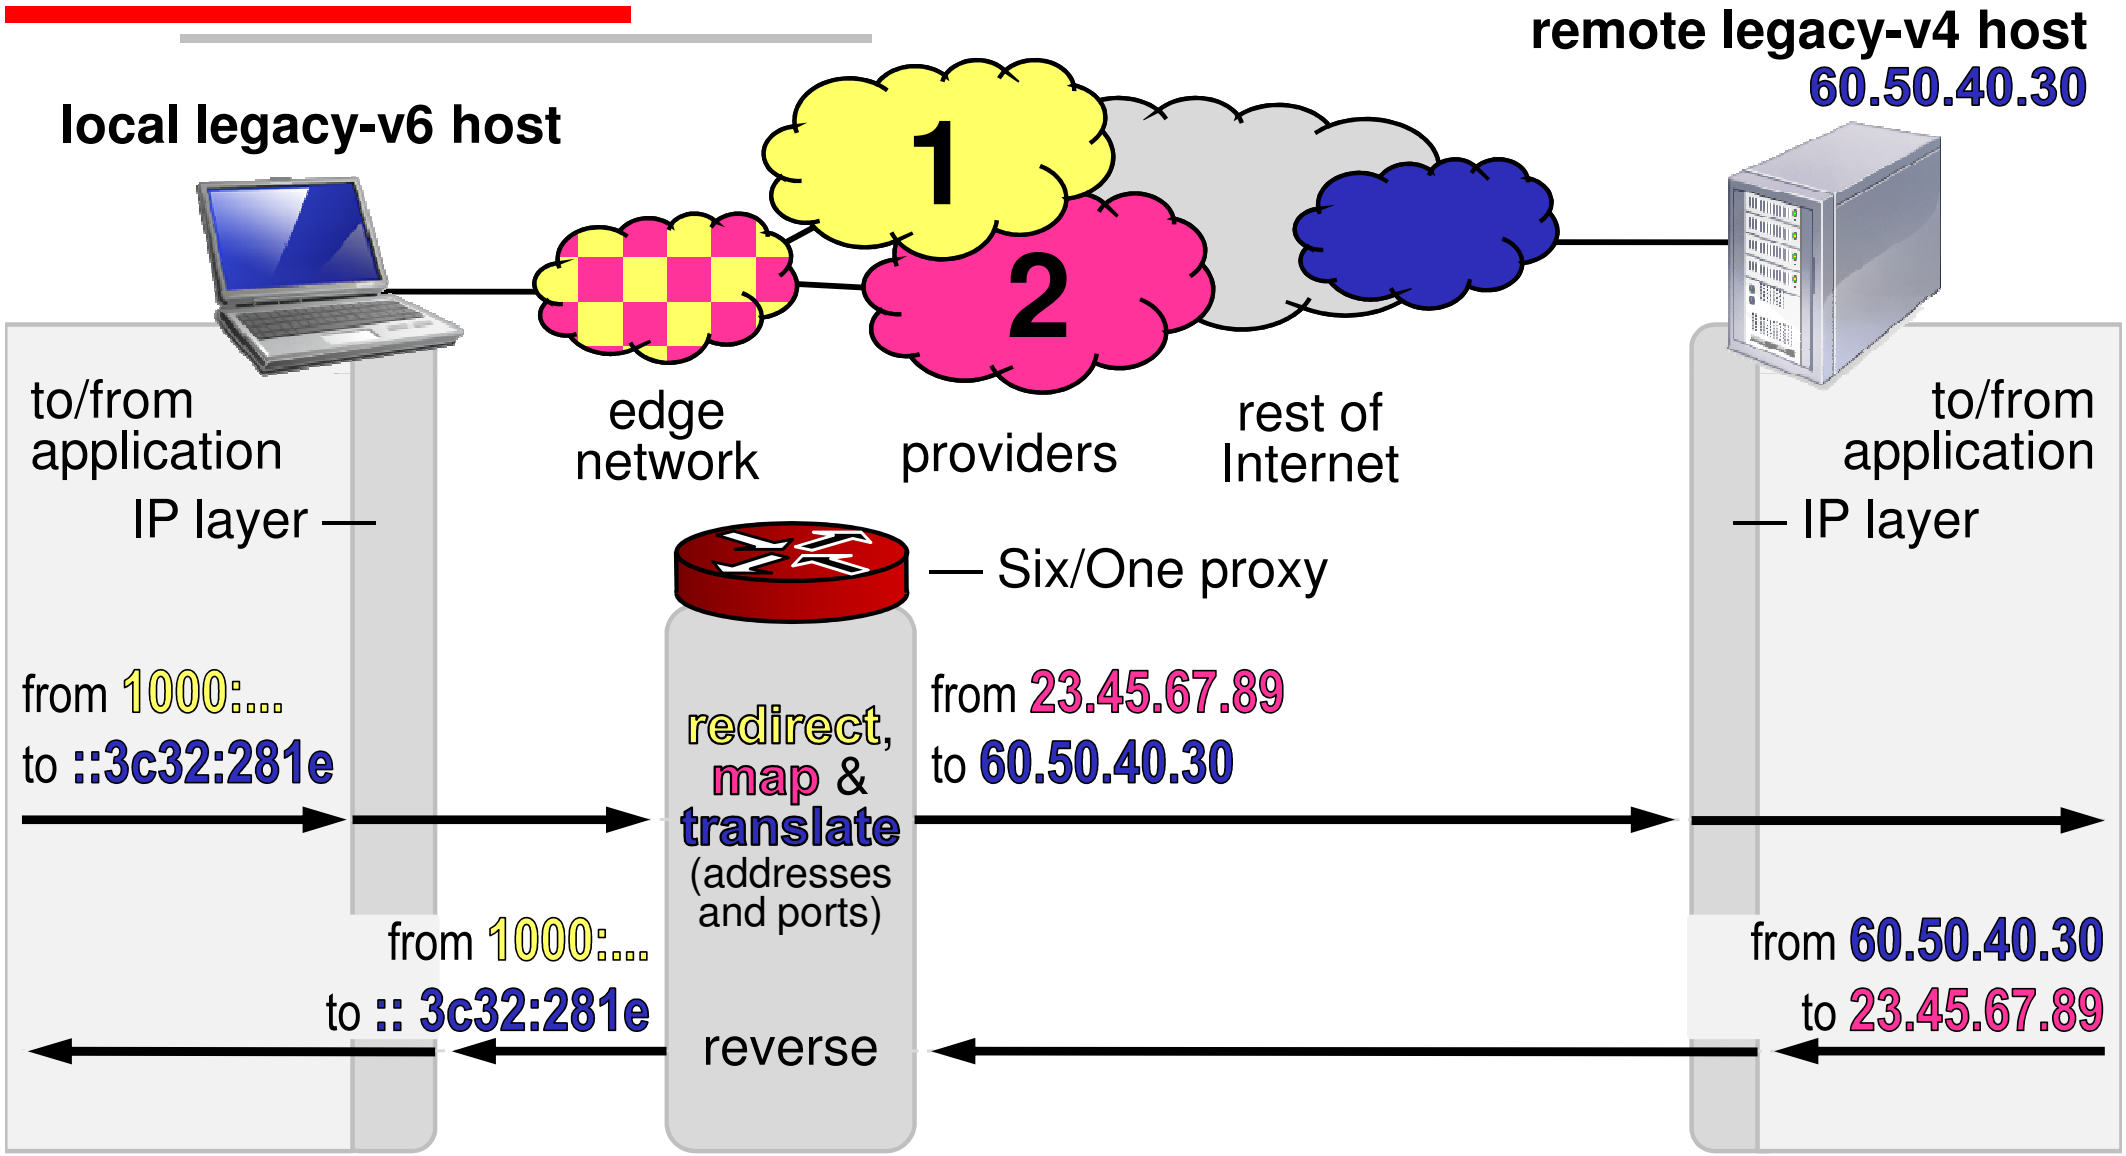

address bunch of local host
1000:abc:1:1234:5d:cff:fe22:57c1
2000:def:2:1234:5d:cff:fe22:57c1

external addresses of Six/One proxy
12.34.56.78
23.45.67.89

Conclusion

- Scalable and traffic-engineering-compatible multi-homing
  - Incrementally deployable, incentives for early adoption
  - Idea: IP address bunches
  - Packet redirection through address rewrite in network
  - Address adaptation and mapping on host
- Extensions
  - Proxies: Deployment without host upgrades
  - Address translation: redirect packets without remote support
- Ongoing work
  - Refined protocol specification
  - Combined HIP-Six/One implementation