


Atom-based Routing

Patrick Verkaik, Andre Broido, Young Hyun, kc claffy

CAIDA / NLnet Labs / RIPE NCC

<http://www.caida.org/projects/routing/atoms/>



Motivation

- Observation: many prefixes share AS path in all RouteViews / RIPE peers
- BGP policy atom: set of prefixes that share AS path
- Routed the same

Motivation

Fewer atoms than prefixes + atoms relatively stable:

- 8 May 2003 RouteViews data:
 - around 33000 atoms
 - covering around 113000 prefixes
 - (15400 ASes)
- Stability over 8 hours:
 - 4.9% of atoms undergo prefix membership change
 - 2-3% of prefixes change atom membership
(Tel Aviv University, 2002)

Apply to routing?

- Summarise prefixes of atom into one routed object
- Incorporate into BGP

Reduce number of routed objects in Default-Free Zone (DFZ):

- Shrink routing tables and forwarding tables
- Perform routing updates per atom, not per prefix
- Potentially improved convergence

Remainder of talk

- Architecture
- Incremental deployment
- Status and unresolved issues


Architecture

- Origination of atoms
- Overview of routing and forwarding
- Routing outside DFZ
- Routing inside DFZ
- Forwarding inside DFZ
- Forwarding outside DFZ

Architecture — Origination of atoms

- To be *declared* by origin ASes
- These ASes partition prefixes into atoms and announce
- Prefixes can be IPv4 or IPv6

Architecture — Example of declared atom



Origin links and declared atoms

- Origin AS with neighbours to which an atom is announced
 - 40% of atoms have just one origin link (single-homed)
 - 37% of atoms have two origin links
 - the total is 24600 origin link sets
 - approximates the number of declared atoms
 - recall the number of full-AS-path atoms is 33000

Architecture — Overview of routing

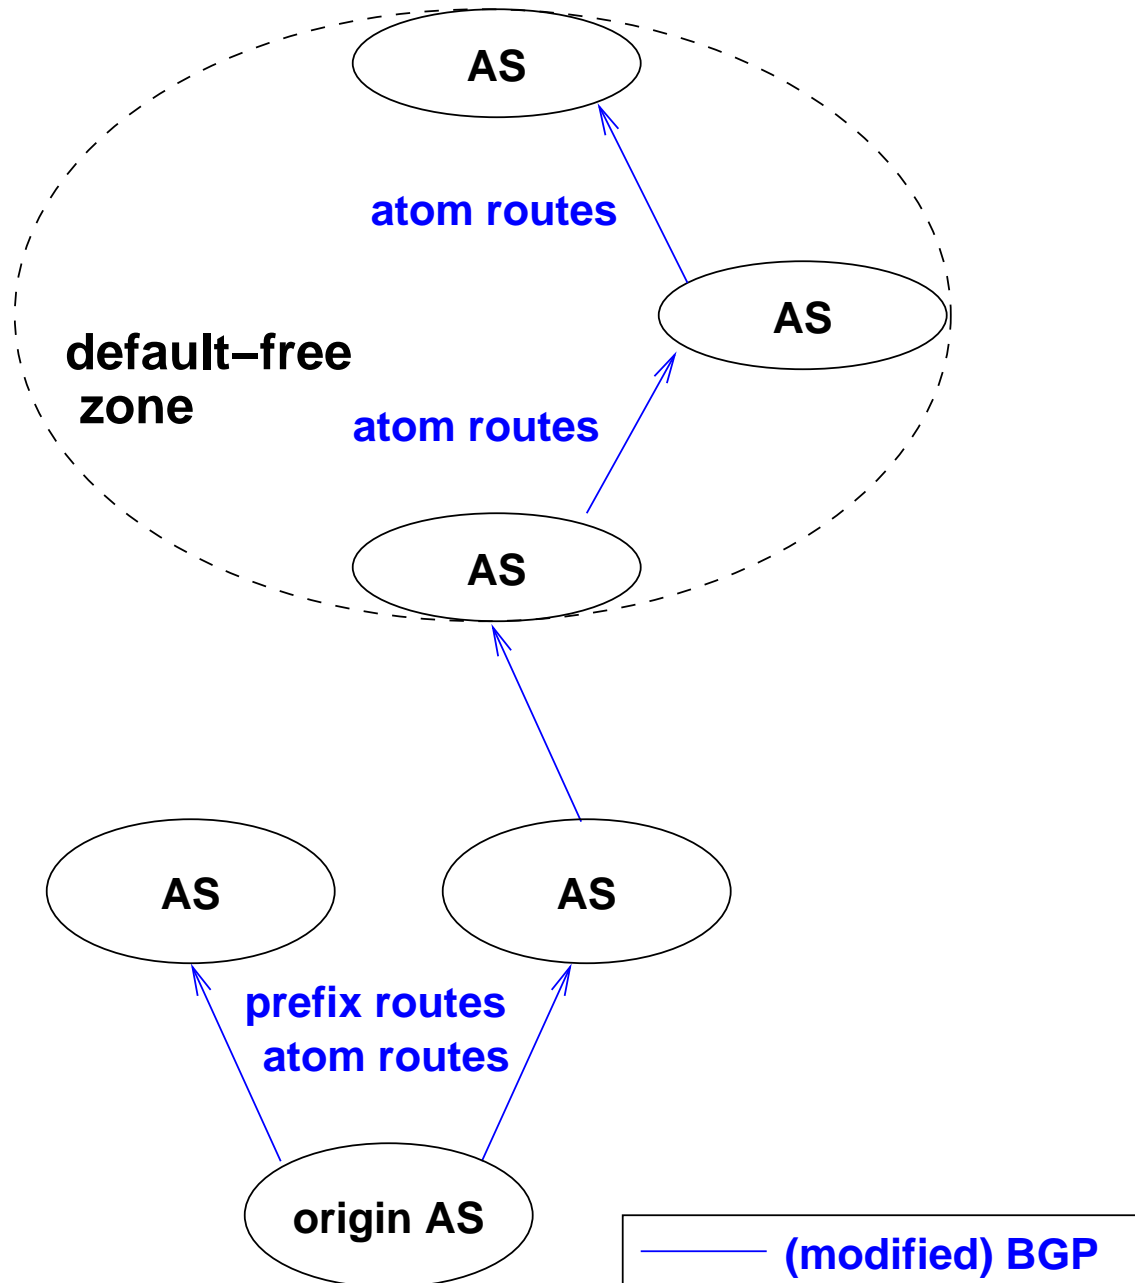
Two functions:

- Atom routing
 - Atom is represented by an atom ID
 - Atom ID syntactically a prefix (unrelated to prefixes in atom)
 - Reason: BGP can route atom IDs
- Atom membership
 - Distributes {atom ID \leftrightarrow prefix} mapping

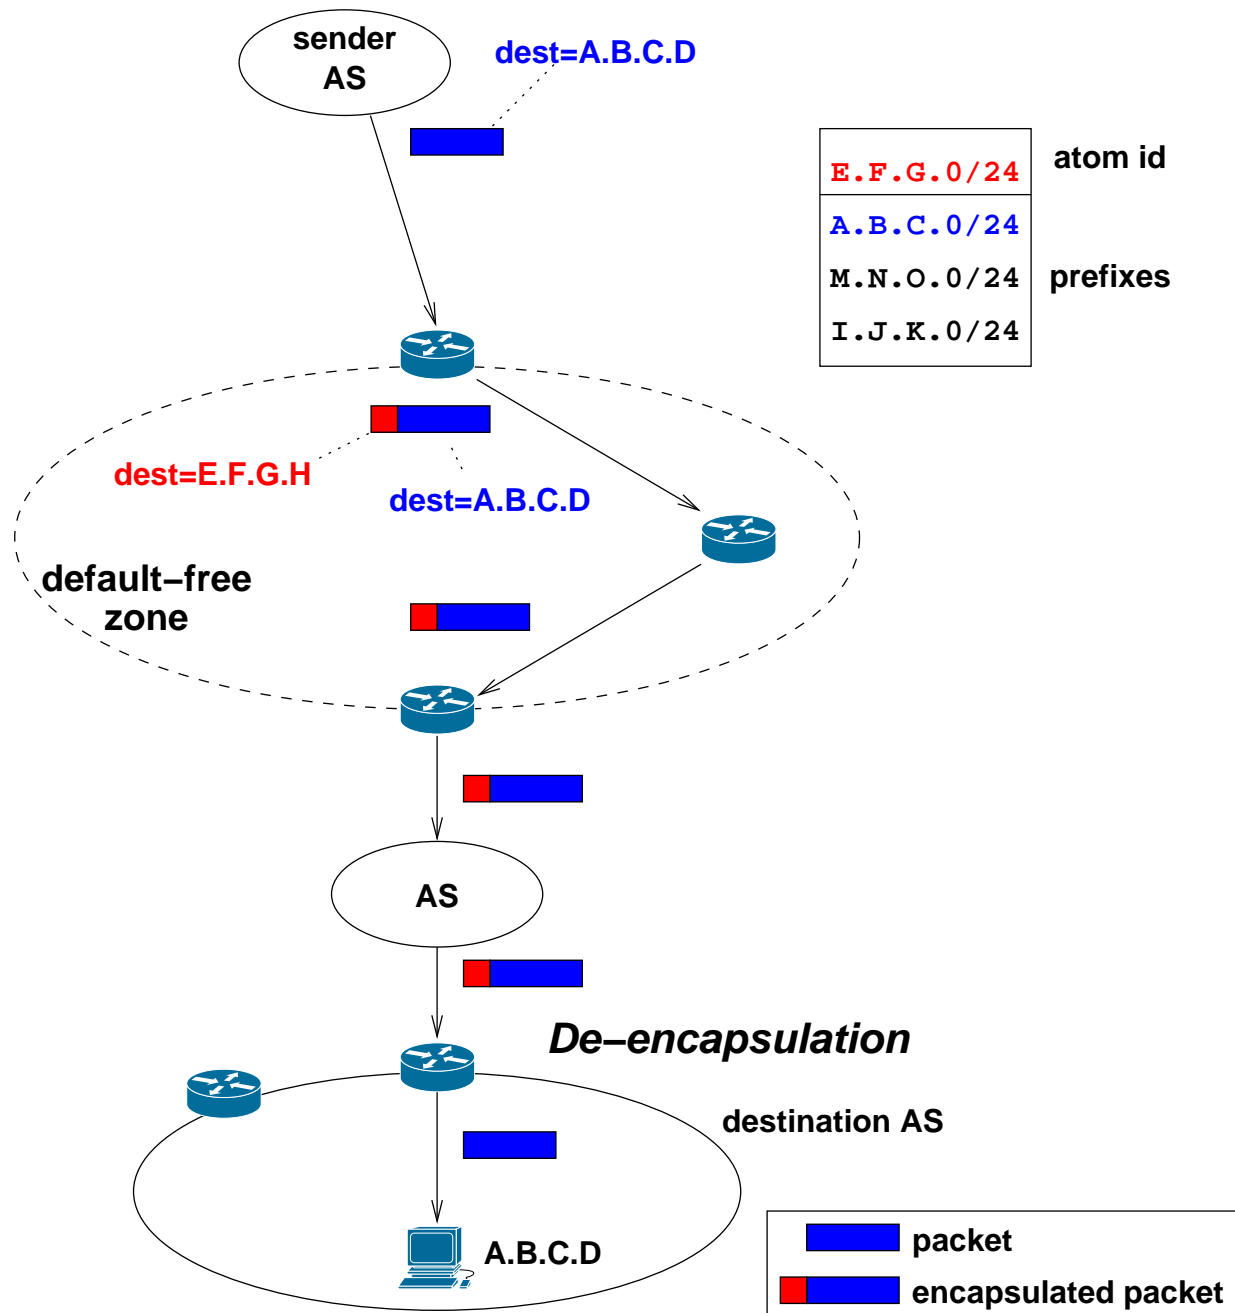
Membership technically not a routing protocol:

- No routing computations
- Location-independent (compare DNS)
- Future optimisation: separate protocol

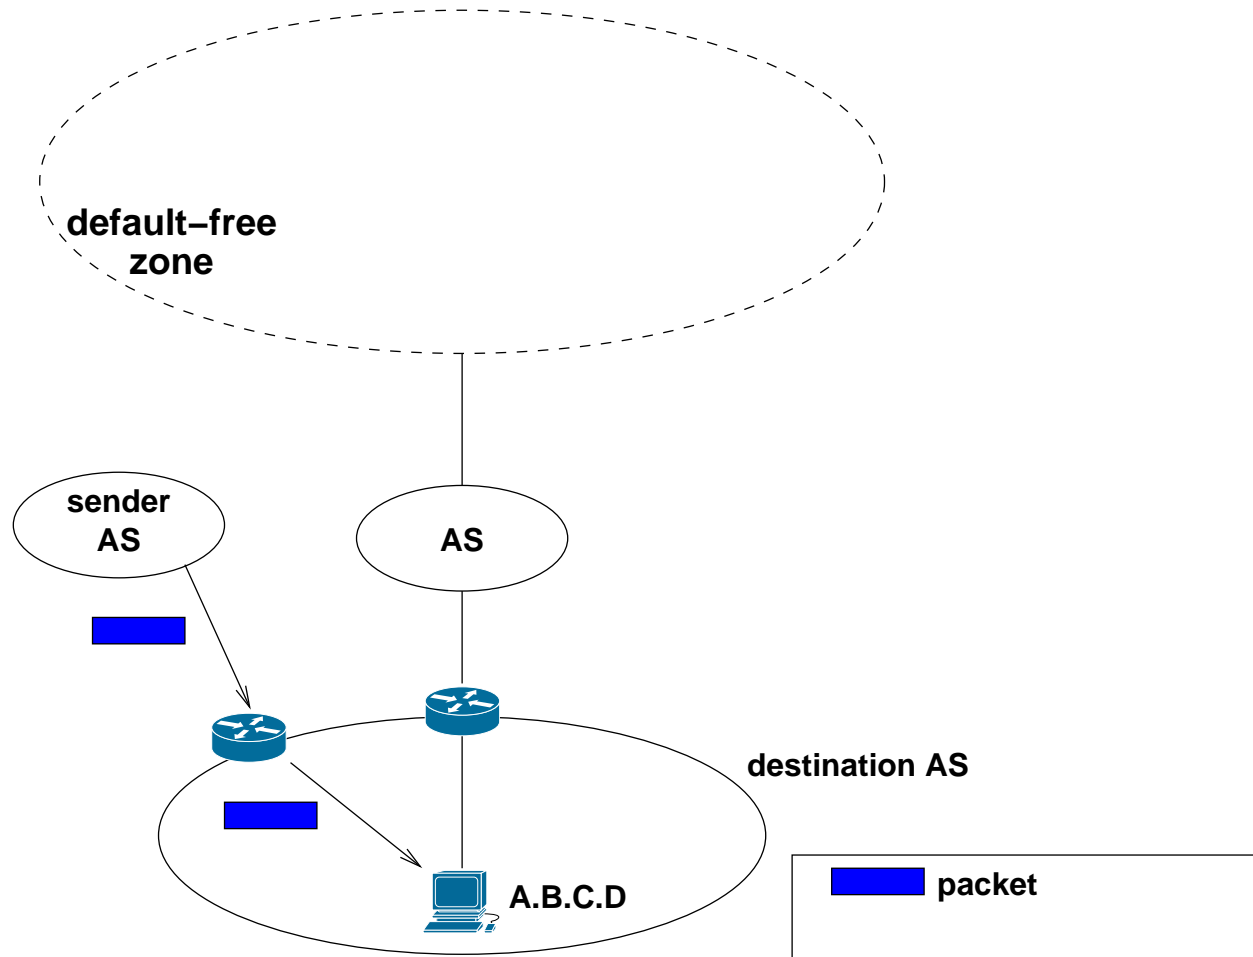
Architecture — Overview of routing



Architecture — Overview of forwarding



Architecture — Overview of forwarding

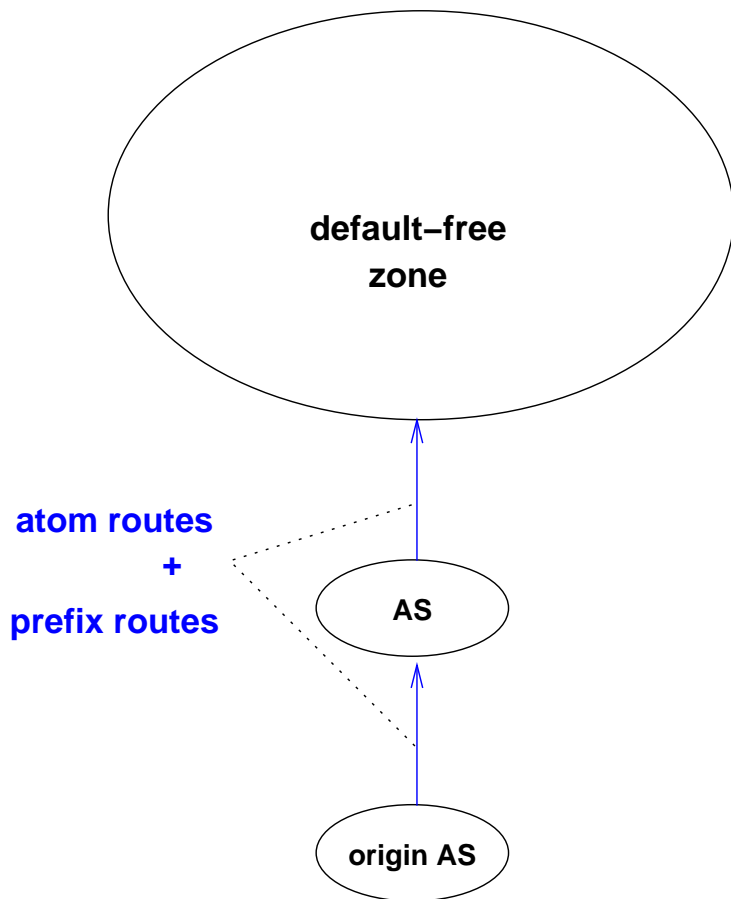


Architecture — Overview of forwarding

Comparing atom-based forwarding and MPLS:

- Atom ID similar to Forwarding Equivalence Class
- Encapsulation rather than label swapping
- MPLS not applied interdomain

Architecture — Routing outside DFZ



Announce atom

Withdraw	
Announce	atom ID
Membership attribute	prefixes

Announce prefixes

Withdraw	
Announce	prefixes
Membership attribute	"atomised"

Withdraw atom

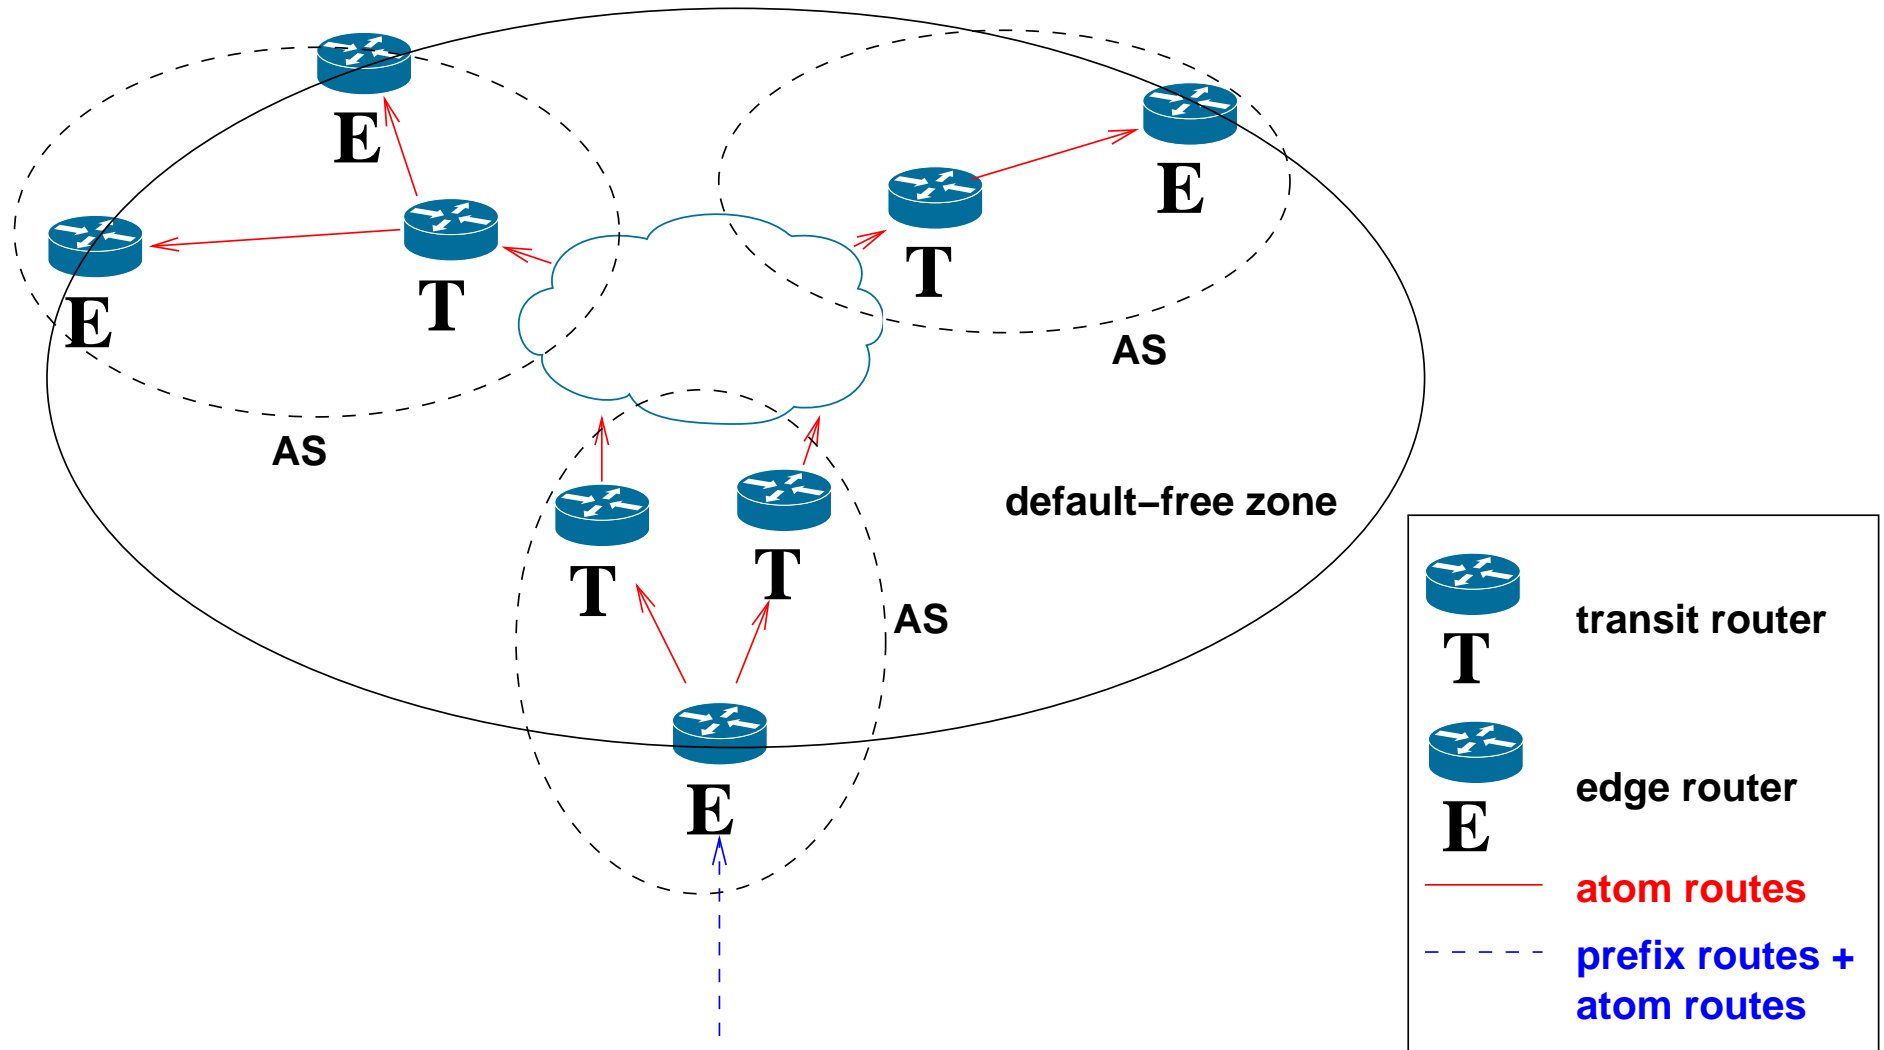
Withdraw	atom ID
Announce	
Membership attribute	

Withdraw prefixes

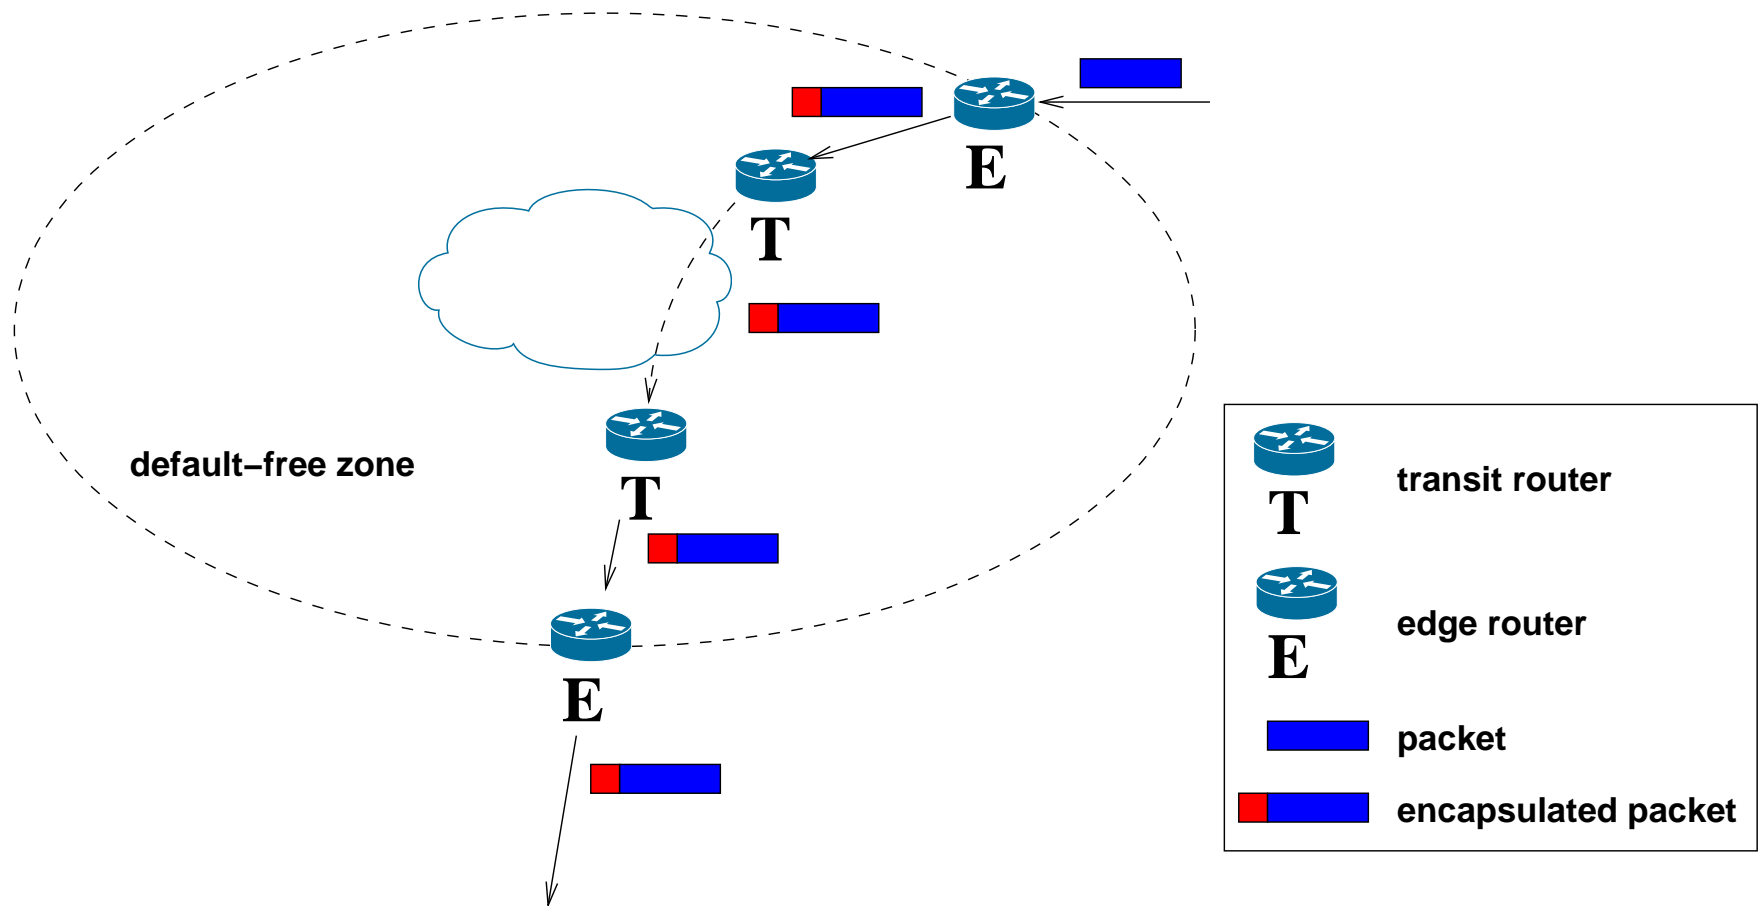
Withdraw	prefixes
Announce	
Membership attribute	

- Atom announcement carries prefix membership in attribute
- Prefixes also routed separately

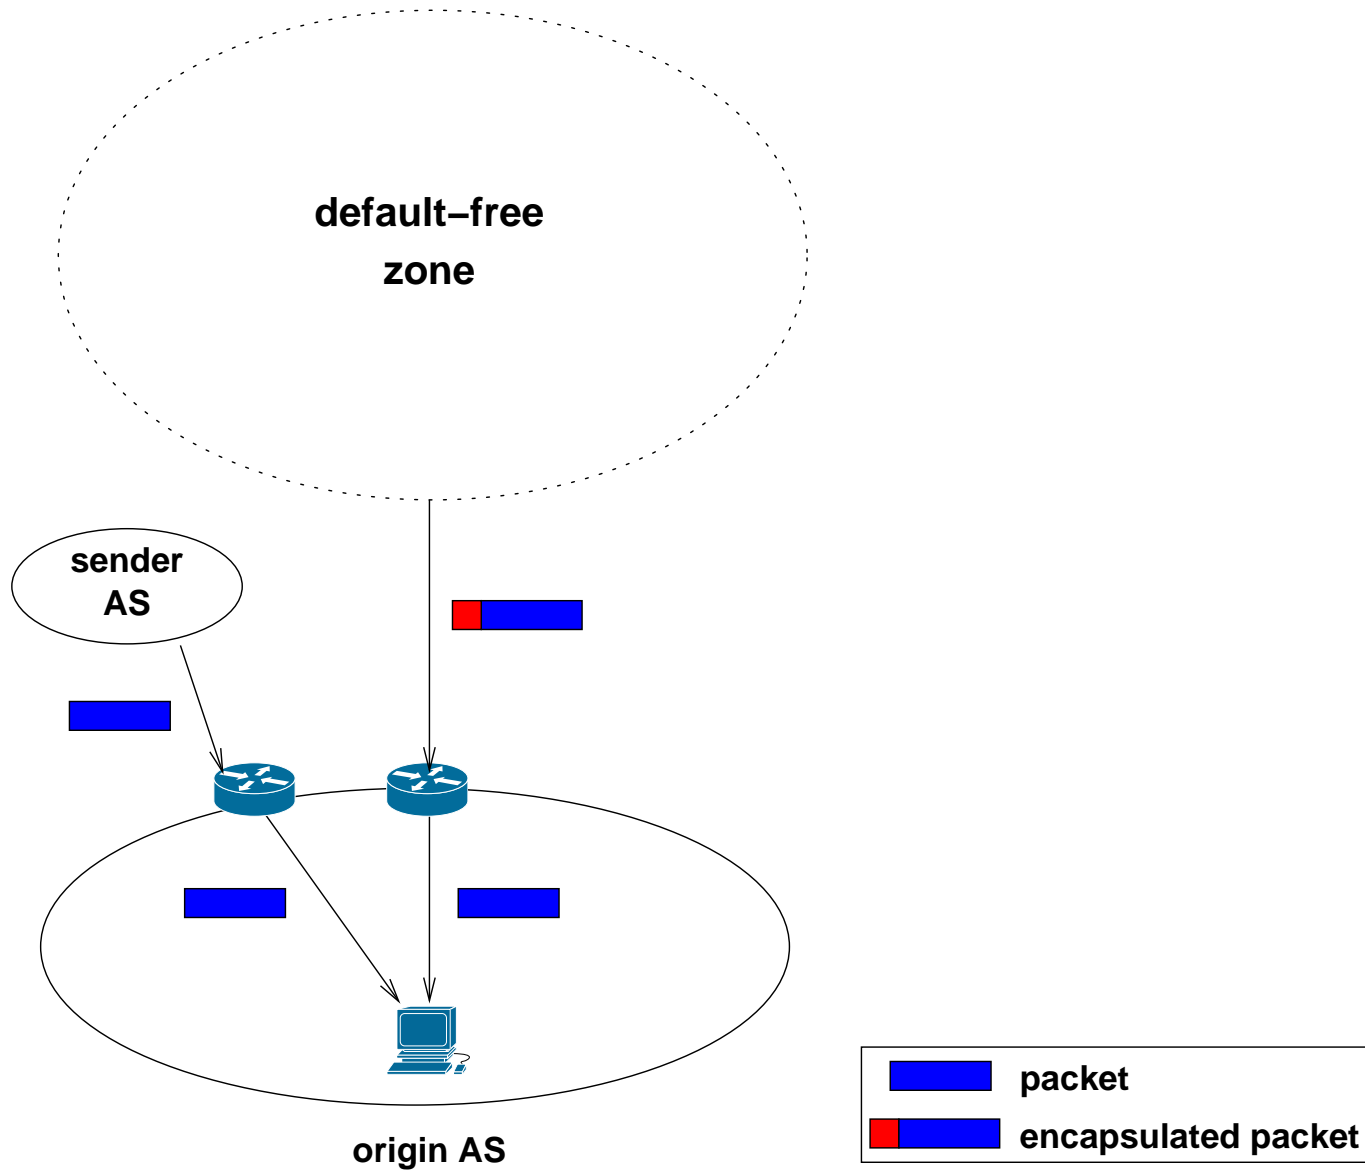
Architecture — Routing inside DFZ



Architecture — Forwarding inside DFZ



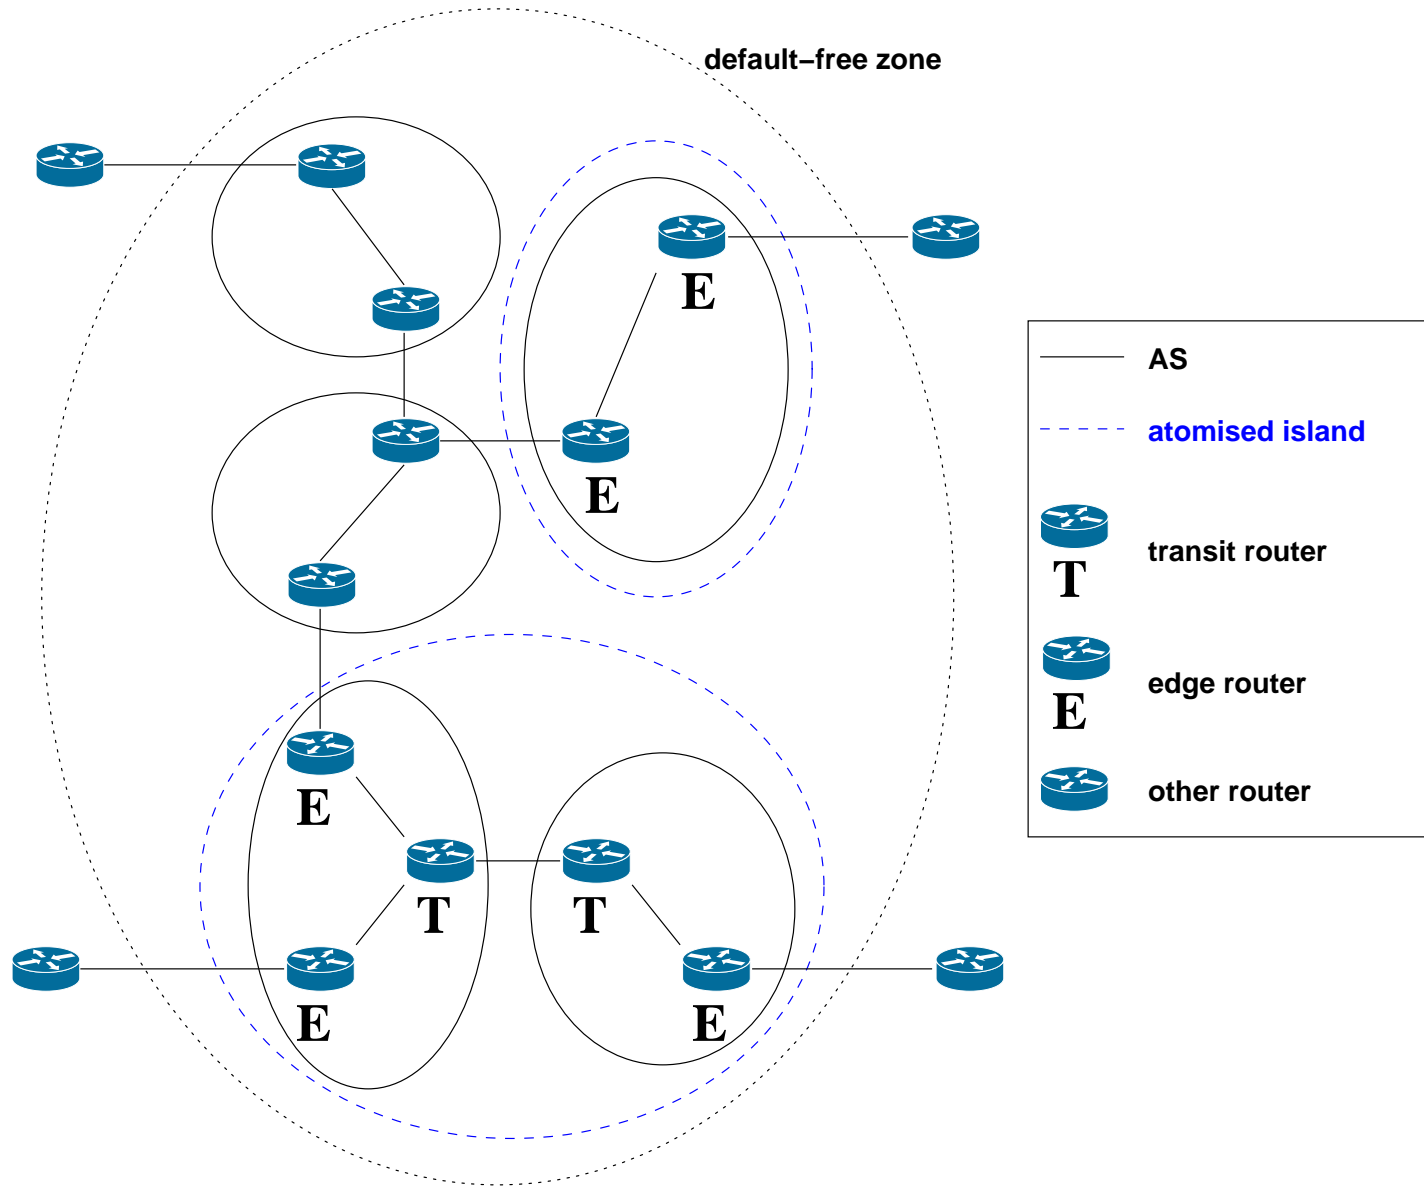
Architecture — Forwarding outside DFZ



Incremental deployment

- Atomless prefixes
- ASes outside DFZ that are not atomised:
 - Nothing special happens between origin AS and DFZ
 - Forwarding is transparent
 - BGP attributes are *optional transitive*
- DFZ ASes that are not atomised.
Solution: *islands* of atomised ASes

Incremental deployment — Islands



Implementation

Status of implementation:

- Implemented previous architecture of atomised routing:
 - No notion of default-free zone
 - All routers prefix-aware
 - Implementing new architecture (fall 2003)
- In Zebra: free routing software (GNU license)
- Atoms declared manually in router configuration language (routemap)

TO DO

- Further analysis into the dynamics of atoms
- A DFZ membership protocol
- Look at non-AS path policies
- Declare atoms upstream from origin AS
- Failures that force atom re-declaration

Questions?

Acknowledgements

Andrew Lange	Jeffrey Haas
Andrew Partan	Maarten van Steen
Bill Woodcock	Nevil Brownlee
Bradley Huffaker	Mike Lloyd
CAIDA folks	Omer Ben-Shalom
Cengiz Alaettinoglu	Pedro Roque Marques
Daniel Karrenberg	Ronald van der Pol
Dave Meyer	Sean Finn
Evi Nemeth	Senthilkumar Ayyasamy
Frances Brazier	Ted Lindgreen
Frank Kastenholz	Teus Hagen
Geoff Huston	Vijay Gill
Henk Uijterwaal	Wytze van der Raay

<http://www.caida.org/projects/routing/atoms/>

Architecture — Overview of forwarding

