K-root anycast deployment

Operational issues, status and plans

Agenda

- Deployment status
- Covering prefix announcement
- Current load data and performance
- Future deployment considerations
- Deployment strategy

Deployment status

Current deployment

- 5 global nodes (prepended)
  - LINX
  - AMS-IX (7/2004)
  - Tokyo (5/2005)
  - Miami (7/2005)
  - Delhi (8/2005)

- 12 local nodes (announced with no-export)
  - Frankfurt, Athens, Doha, Milan, Reykjavik, Helsinki, Geneva, Poznan, Budapest, Abu Dhabi, Brisbane, Novosibirsk

Covering prefix announcement

"Oh K can you see"

- Problem pointed out by Randy Bush
- http://www.merit.edu/mail.archives/nanog/2005-10/msg01226.html

- Nasty interaction of no-export with anycast
  - We use no-export to prevent local nodes from leaking
  - If we have a customer AS
    - Whose providers all peer with a local node
      - And honour no-export
    - They might see no route at all!

"Oh K can you see" (2)

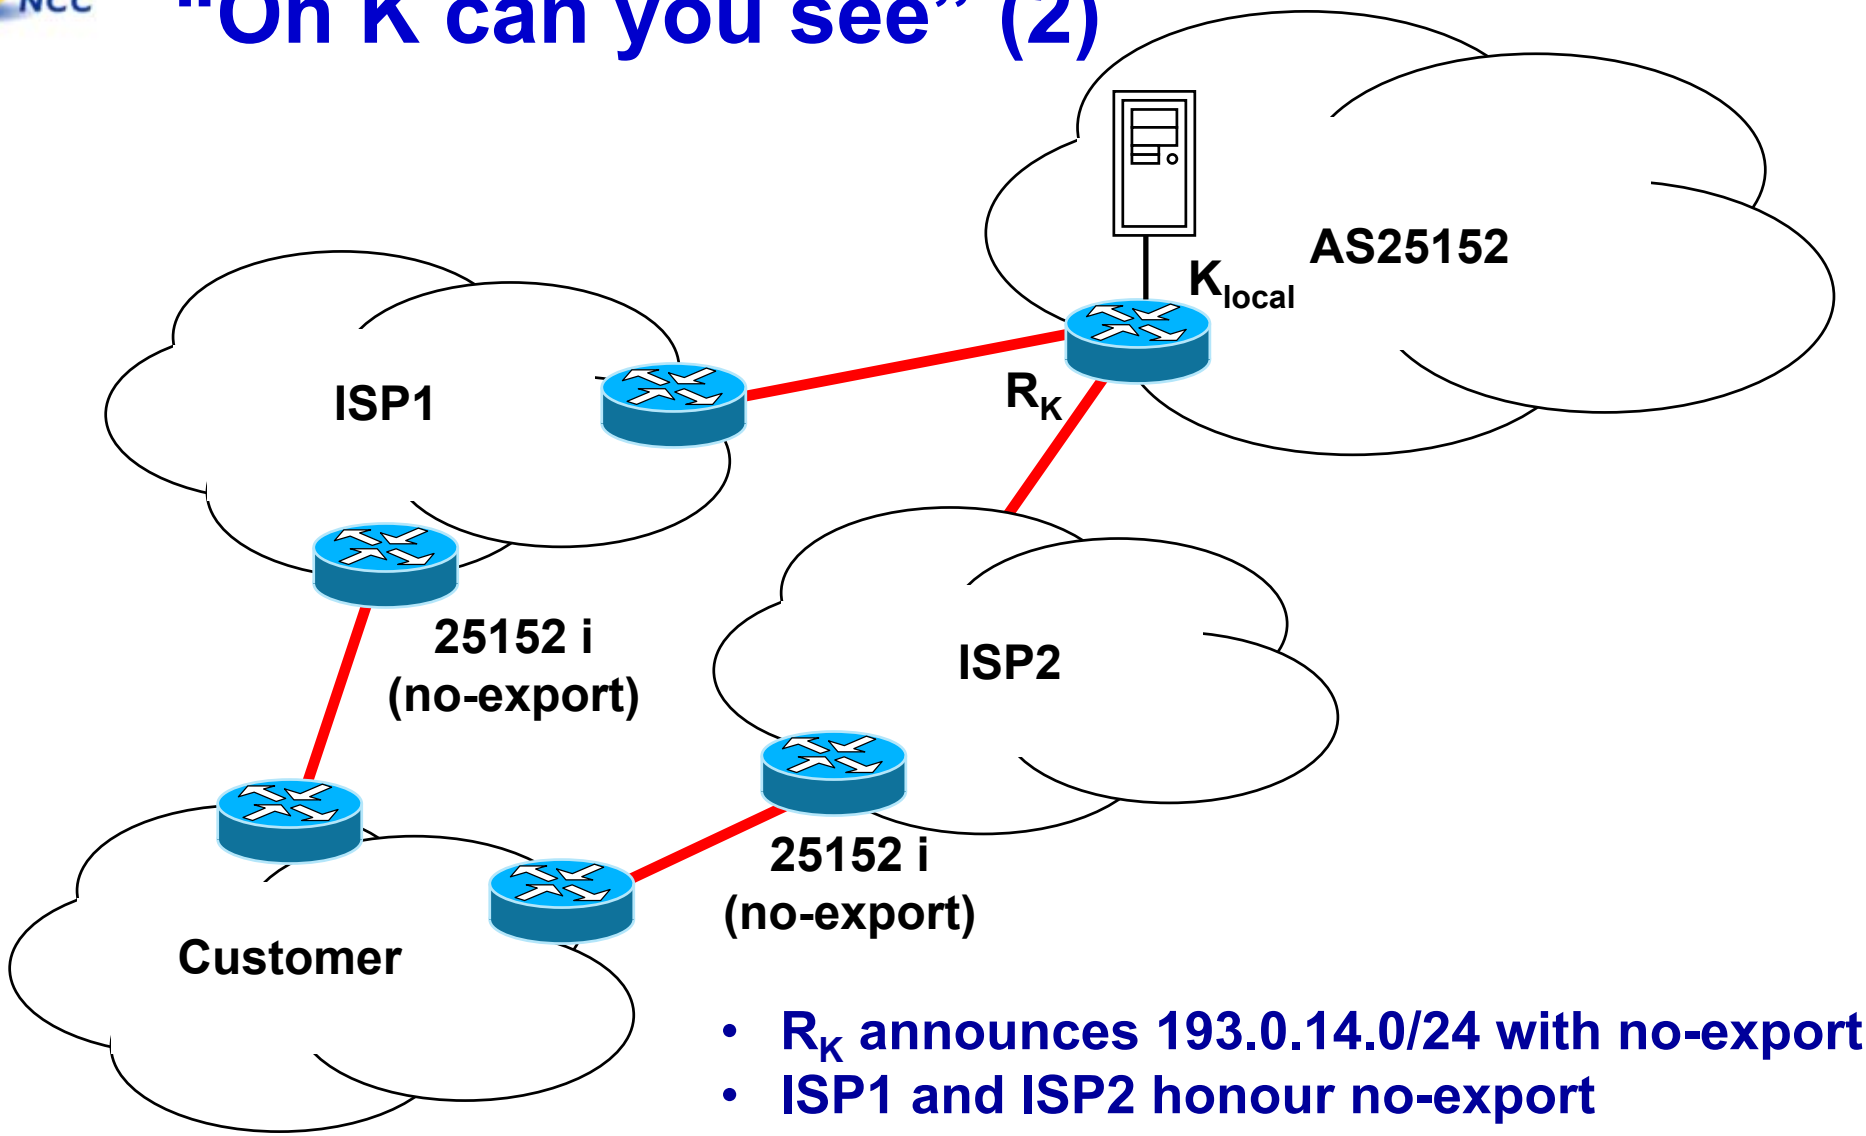

- R_K announces 193.0.14.0/24 with no-export
- ISP1 and ISP2 honour no-export
- Customer has no route to 193.0.14.0/24

Extent of the problem

- Solution: announce 193.0.14.0/23 without no-export @AMS-IX
- Was this a problem?
- See what happened when prefix was announced

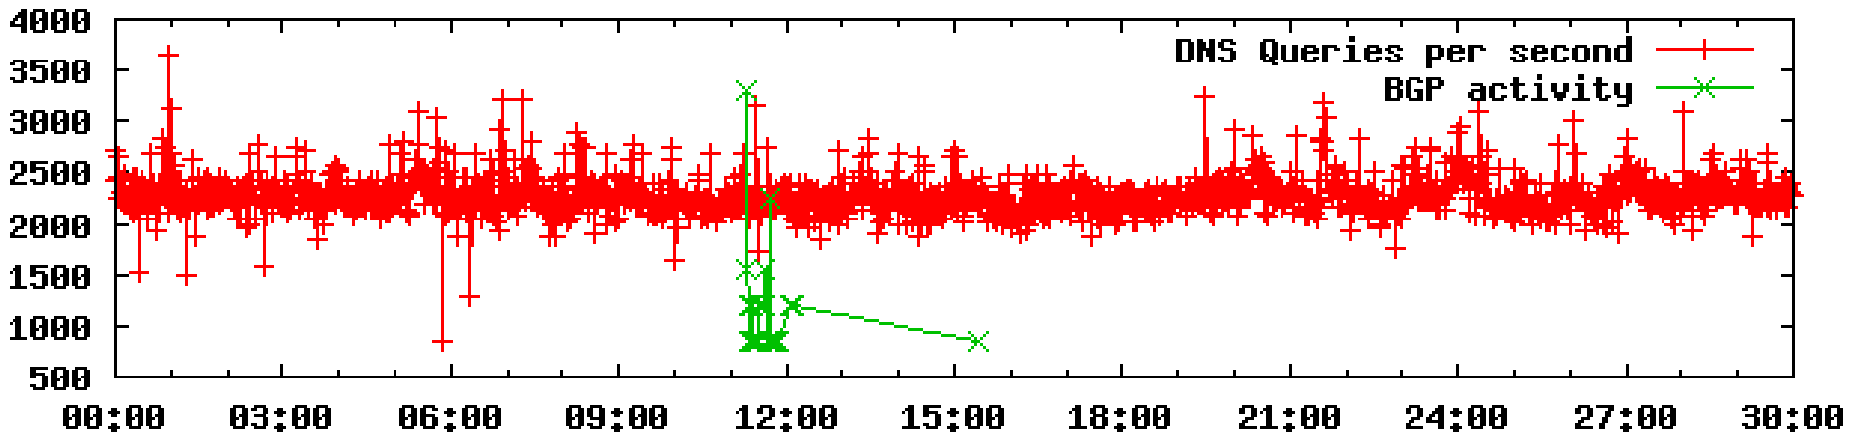

- Red: AMS-IX queries per second
- Green: BGP activity
- “Nothing here”

Current load and performance data

Global node load

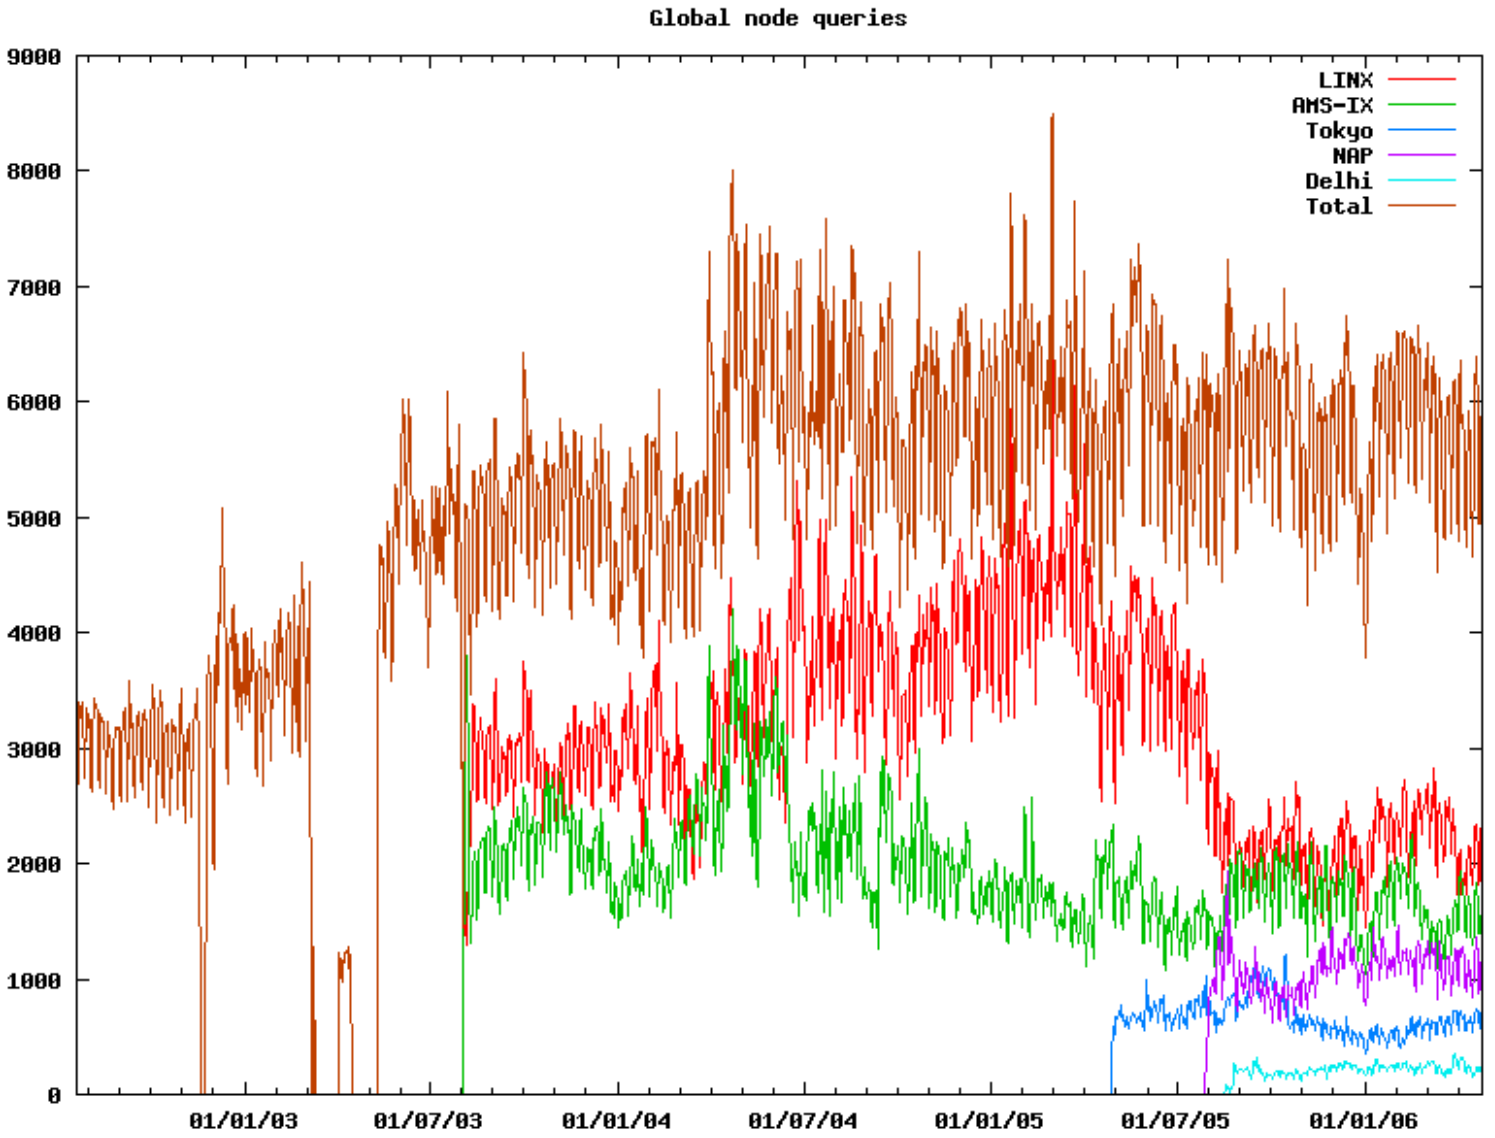

Local node load

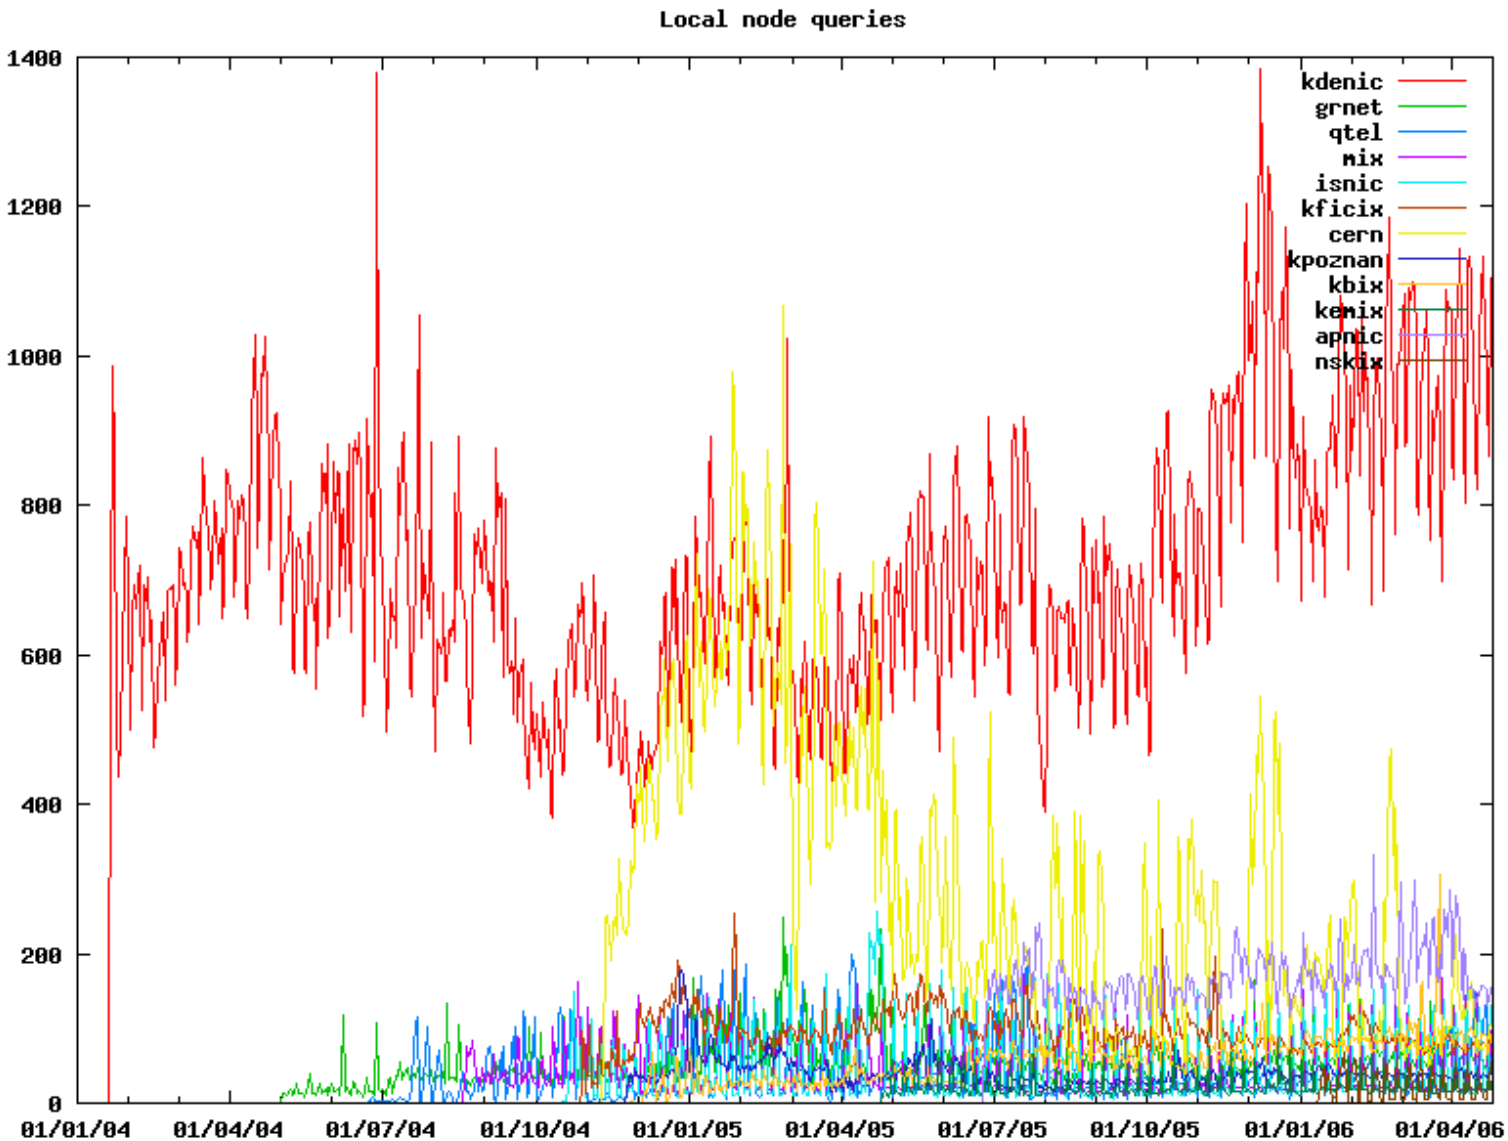

Local node load, cumulative

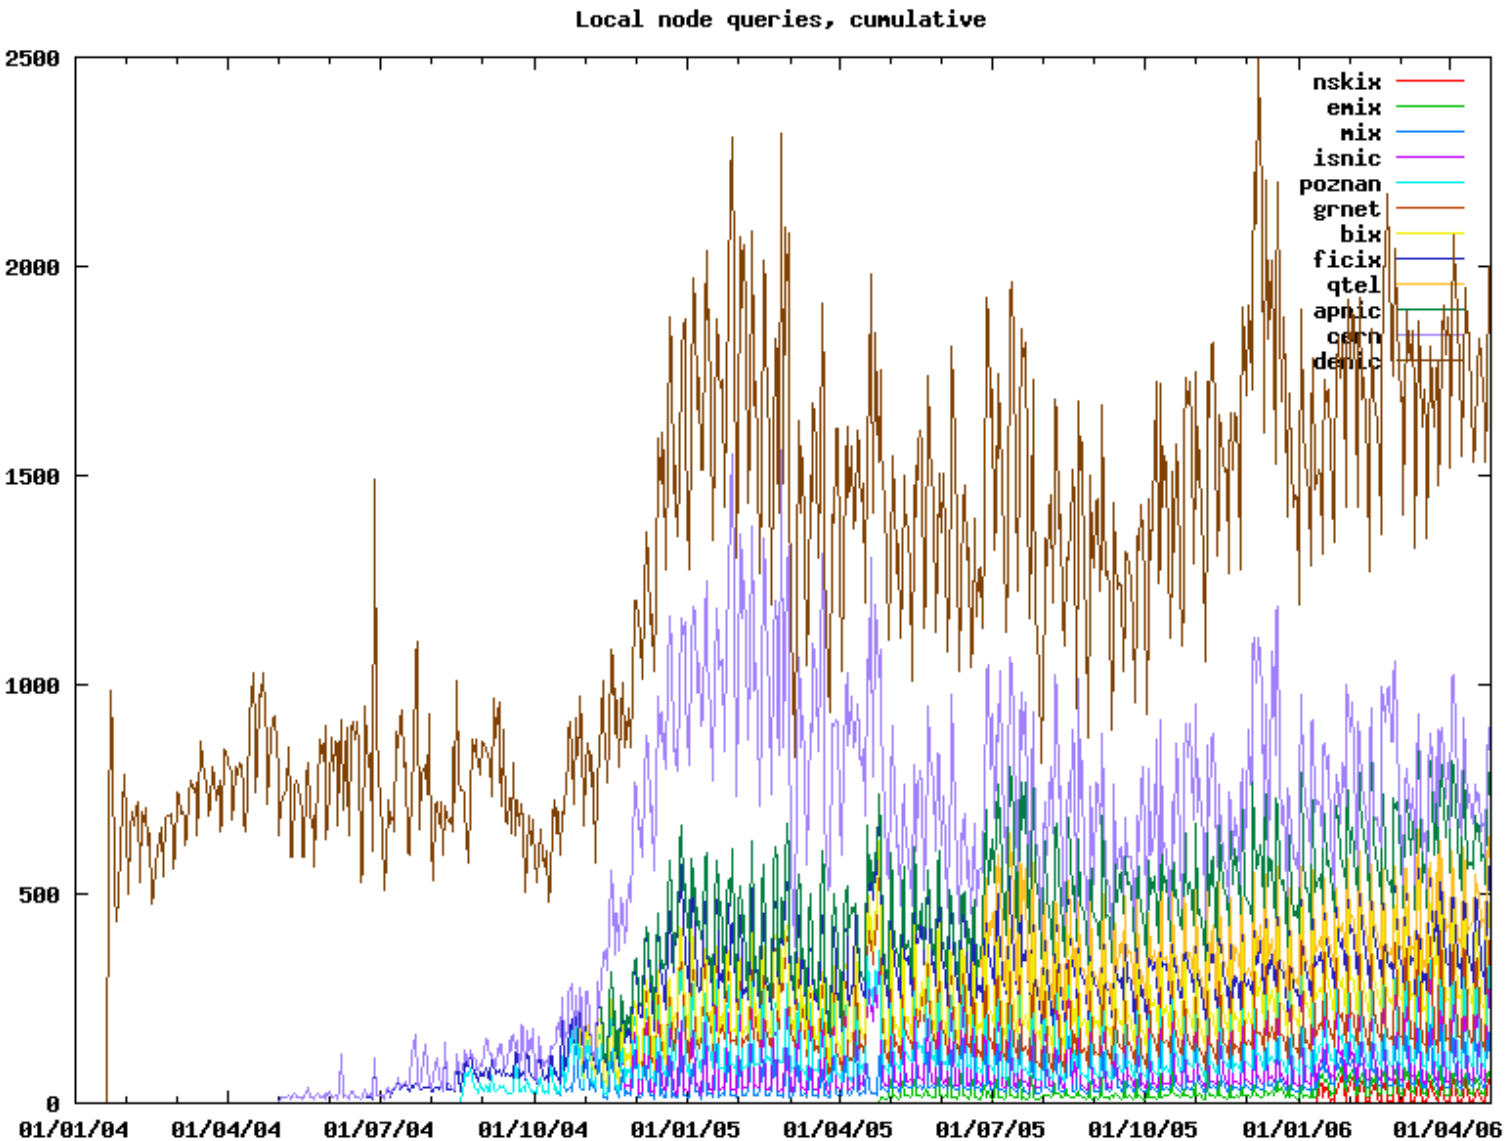

Local vs global load

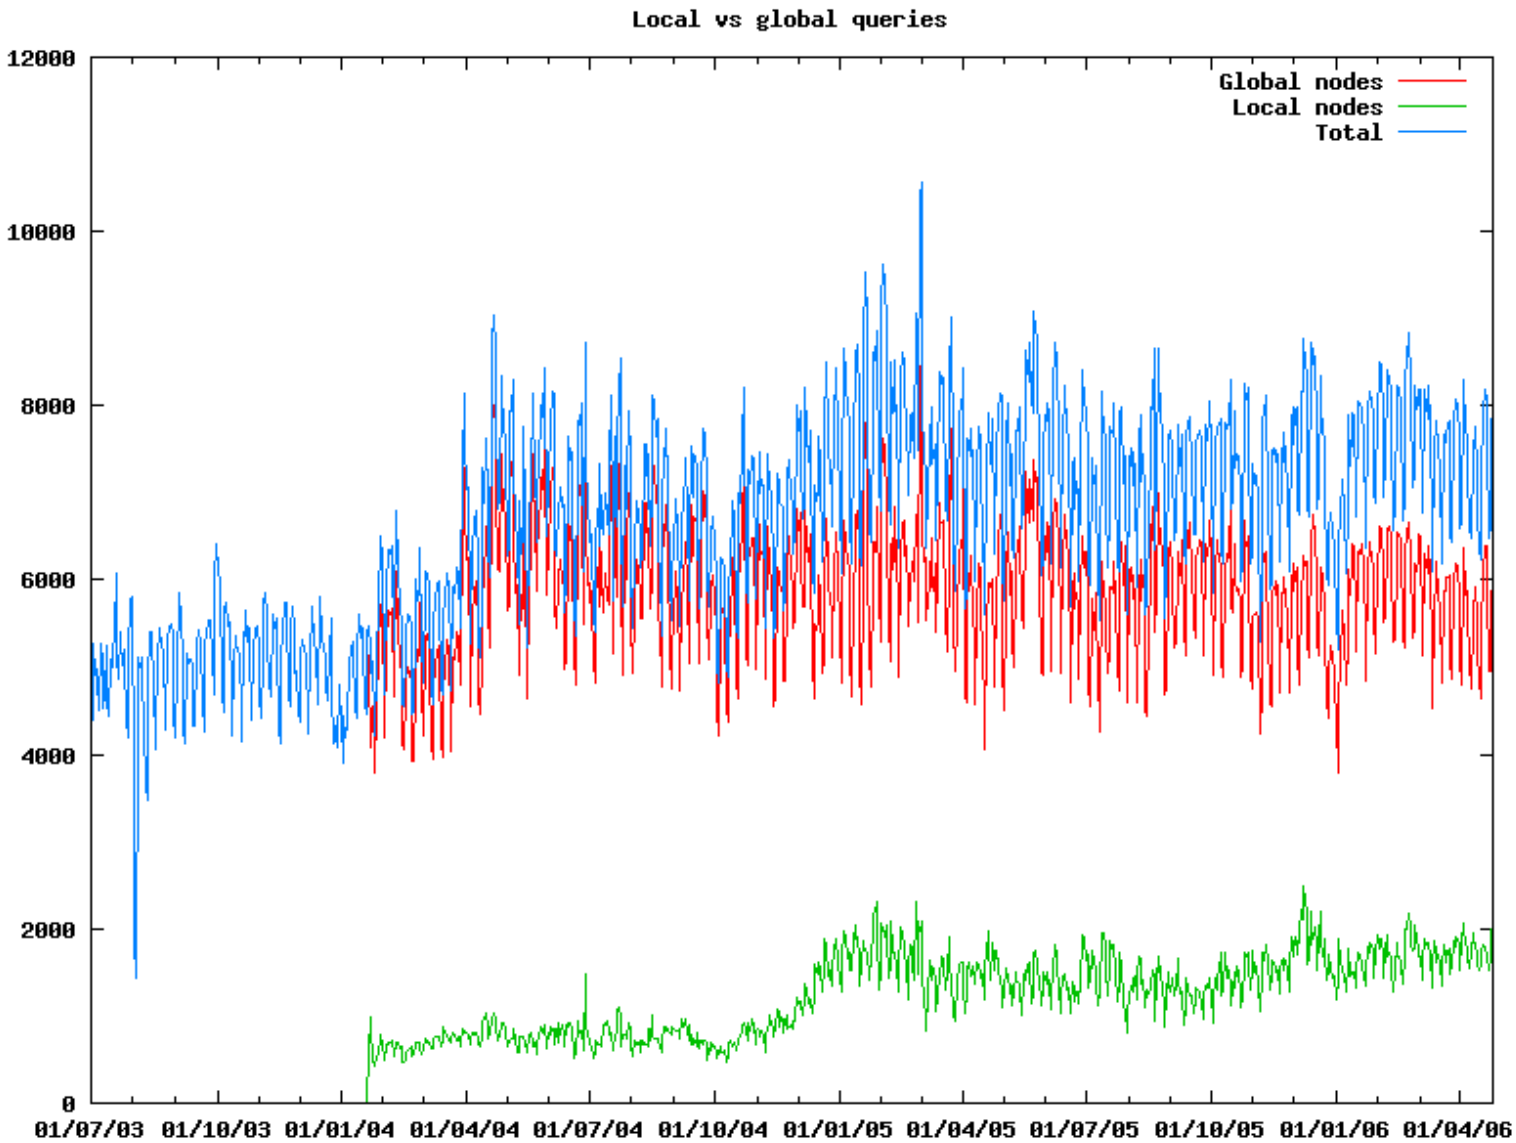

Benefit of global nodes

Europe

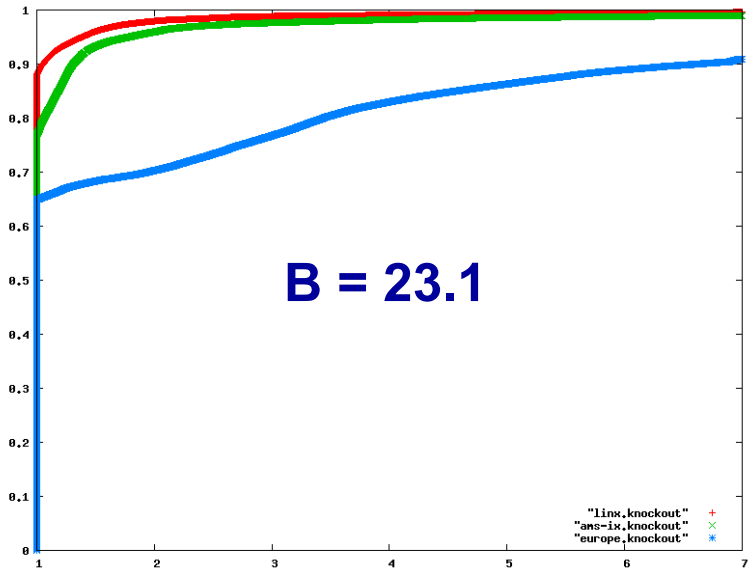

B = 23.1

B = 14.1

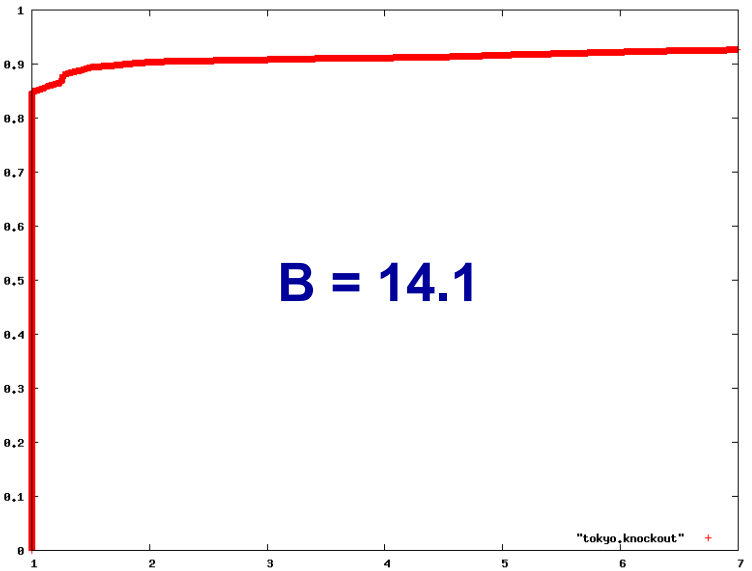

Tokyo

NAP

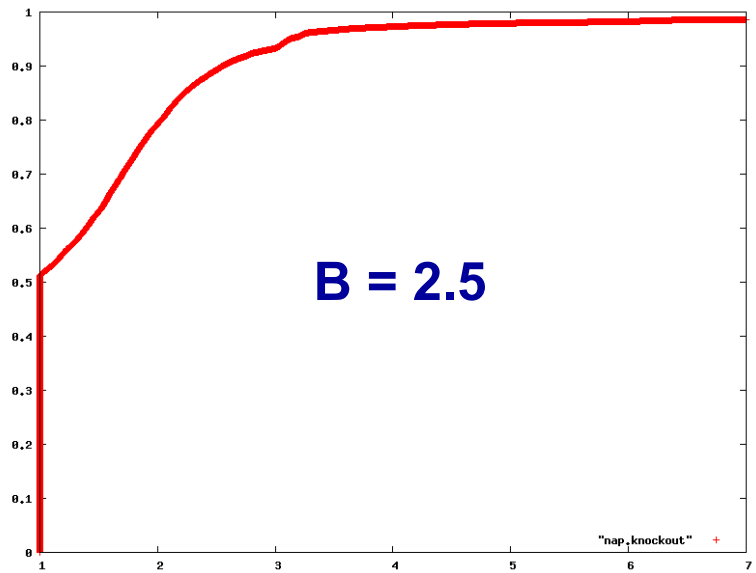

B = 2.5

B = 1.01

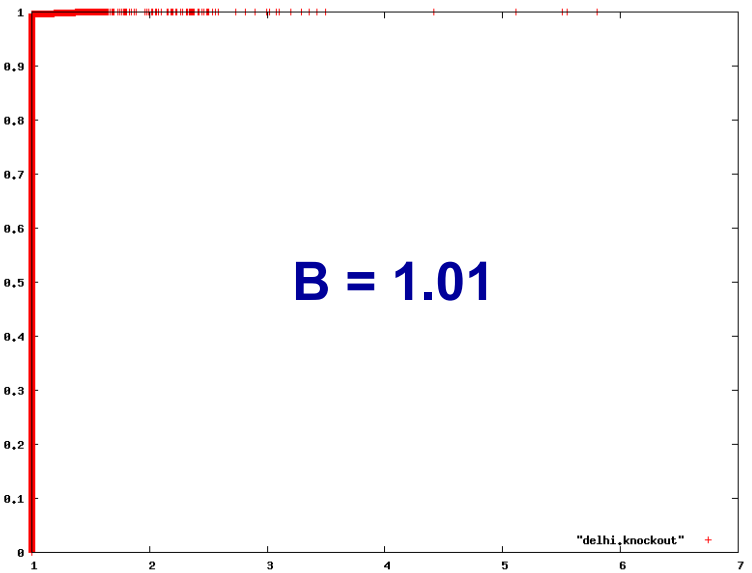

Delhi

Load considerations

- Good global load-balancing (except Delhi)
- New global nodes mostly take traffic from LINX

- Wide variation in load
  - 1 order of magnitude for global nodes
  - 2 orders of magnitude for local nodes
  - Caused by request-driven deployment

- ~25% of load handled by local nodes

Future deployment considerations

Basics

- Given wide variation in node benefit, we want to place nodes only where they are useful
  - In first approximation, this is determined by number of queries
- What determines the numbers?
  - Node placement
  - Peering agreements
- Node placement is obvious. What about peering?

Impact of peering (1)

- IGP metrics lose out to AS-path length
  - If an AS peers with one K node, the AS will use that node only
    - Clients in large Ases prefer to traverse the whole AS rather than take transit through someone else
- “Peer with the same people everywhere”
  - So IGPs and MEDs can make good decisions

Impact of peering (2)

- Prefer-customer routing
  - If a tier-3 ISP in Tokyo provides transit for K,
    - Their tier-2 ISP will prefer their route to a transit or peer route
    - European customers of the tier-2 might go to Tokyo
      - ... instead of taking transit to Europe or Miami
- “Peer with big networks”

Deployment strategy

Deployment strategy

- So:
  - Use global nodes for maximum effectiveness
  - Place the nodes in geographic regions of high client density
  - Peer with big networks
  - Peer with the same networks everywhere
  - … and let IGPs sort it out
- We need to turn our attention to peering
  - Outreach

Geographic distribution

- Currently, geographic distribution is quite good
- Areas we're missing:
  - Africa
    - South Africa?
    - Is there demand?
  - US West coast
    - Good candidate
- Idea: evaluate candidate locations
  - Ping client population from a host there to estimate B

Local node deployment

- Can continue to deploy local nodes
  - But only where it makes sense to do so
    - Evaluate exchange traffic
    - Expression of interest by exchange members?
      - Peering is the key
- Must be careful of leakage

Questions?