The State of BGP Routing

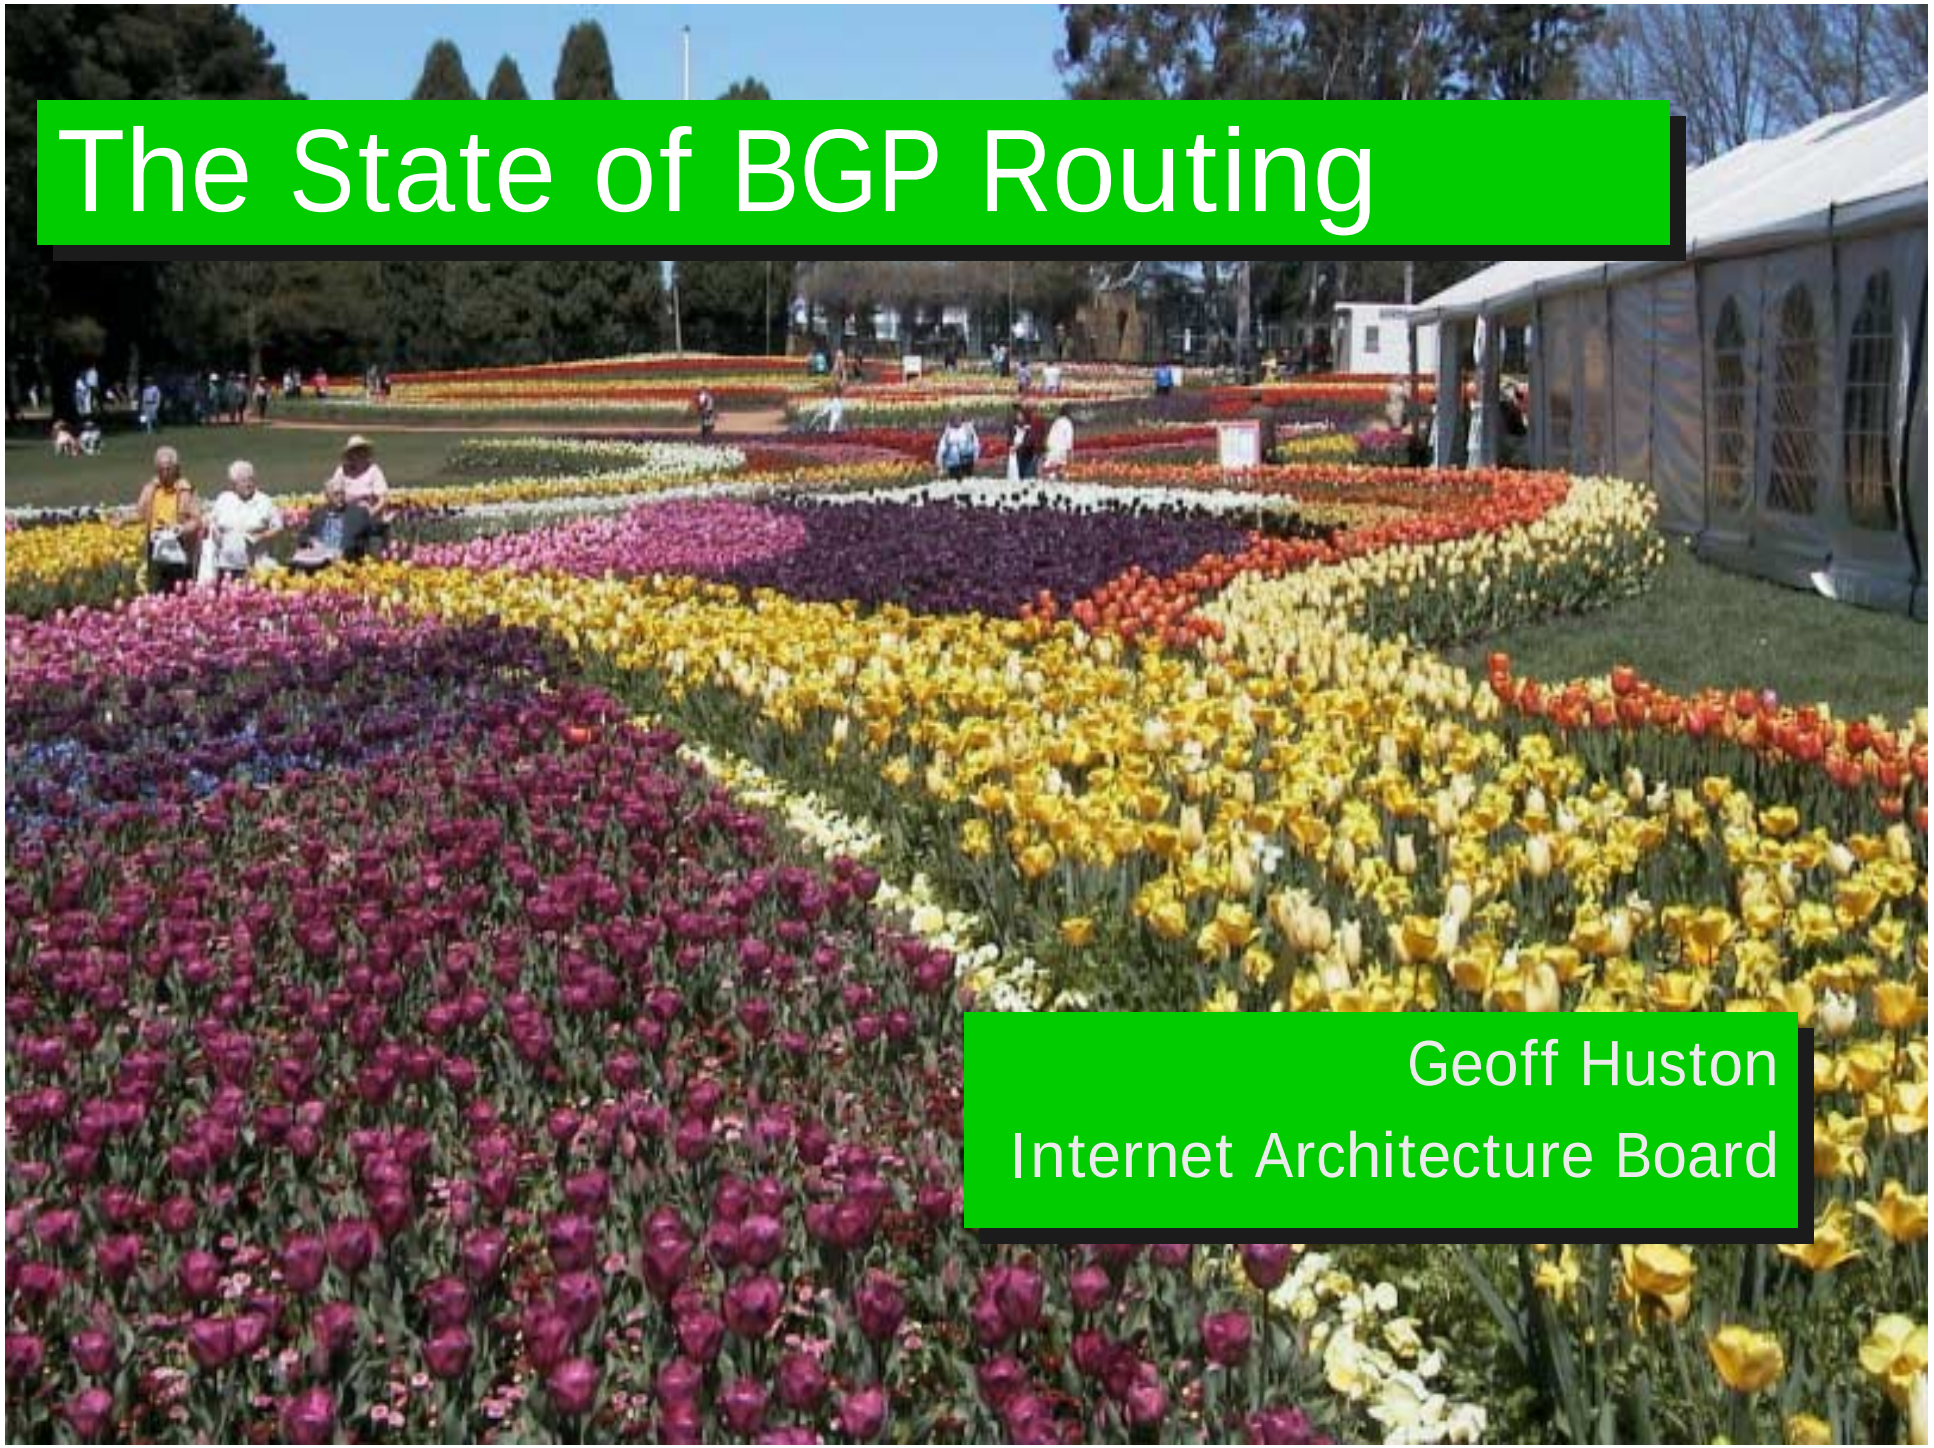

Geoff Huston
Internet Architecture Board

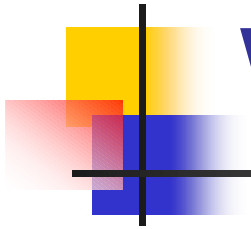

Why BGP?

- BGP describes the structure of the Internet, and an analysis of the BGP routing table can provide information to help answer the following questions:
  - What is changing in the deployment environment?
  - Are these changes sustainable?
  - How do inter-domain routing protocols interact with the changes in the structure of the Internet?

BGP Table Growth – 12 year history

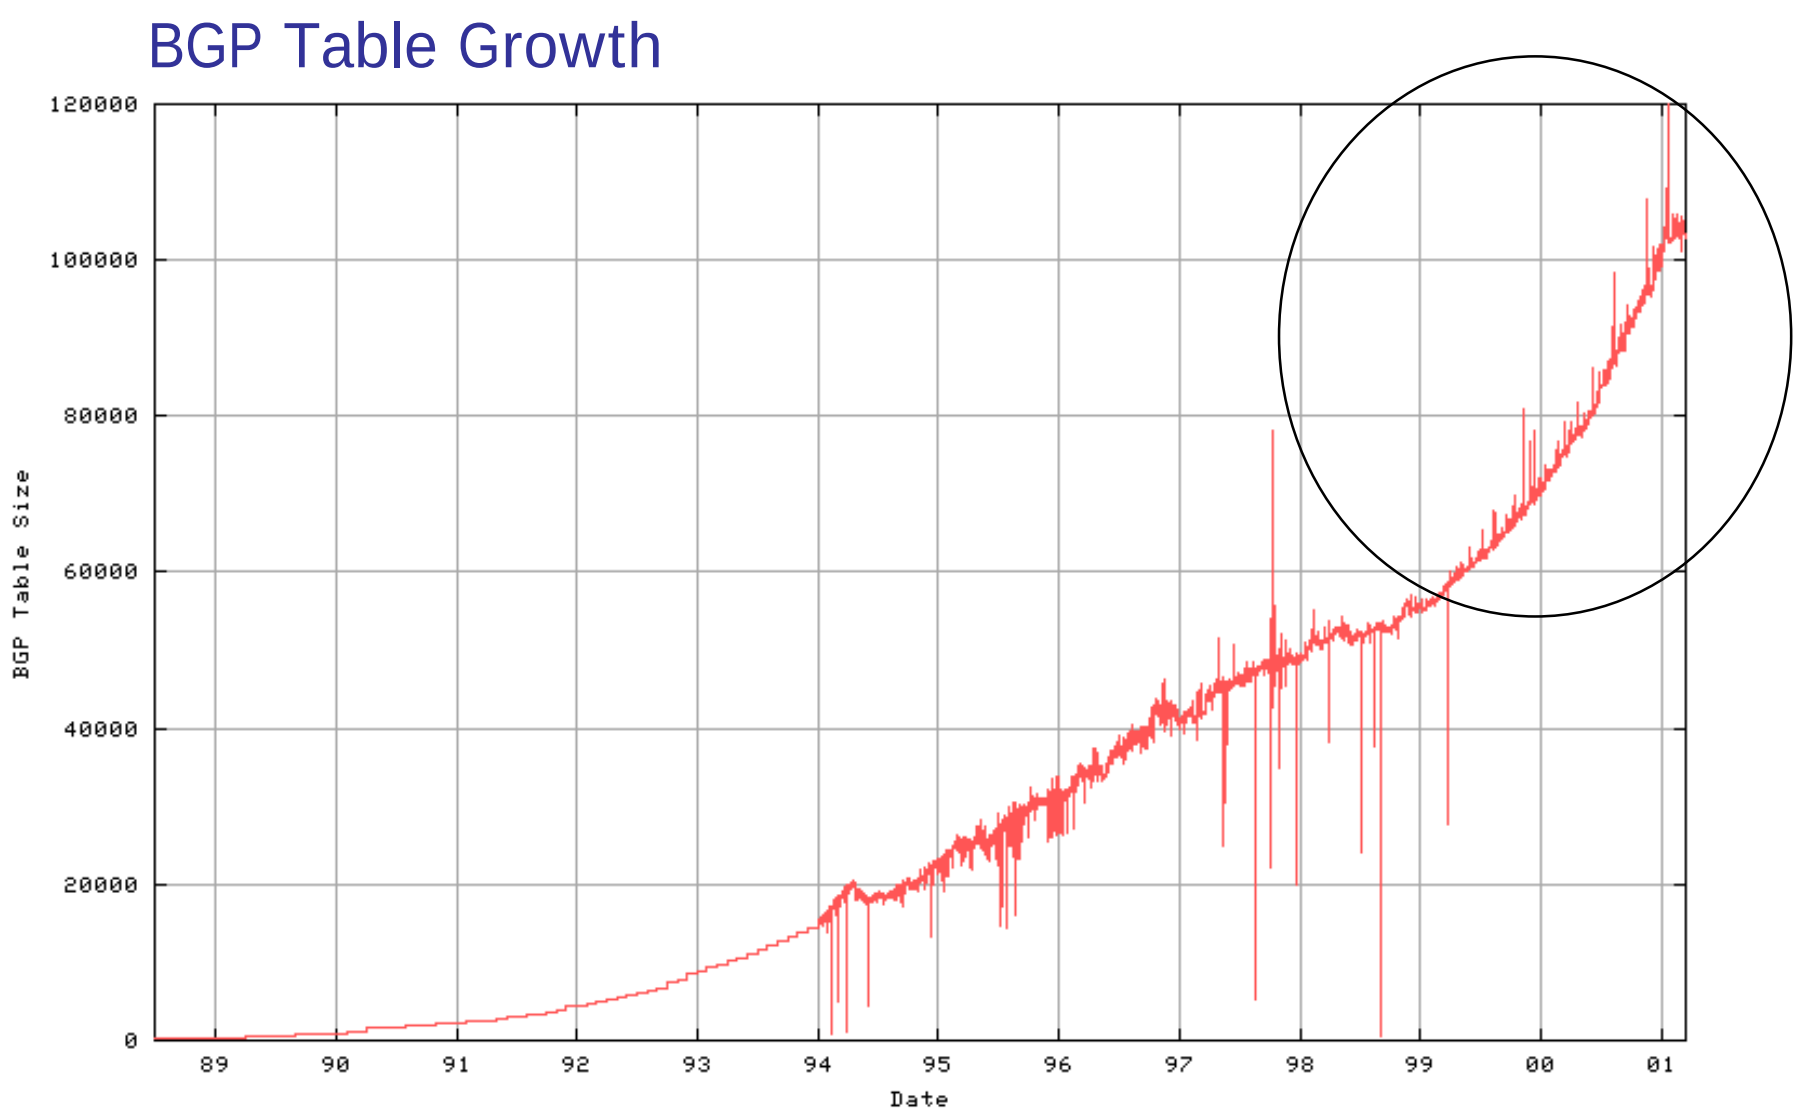

BGP Table Growth – 2 year history

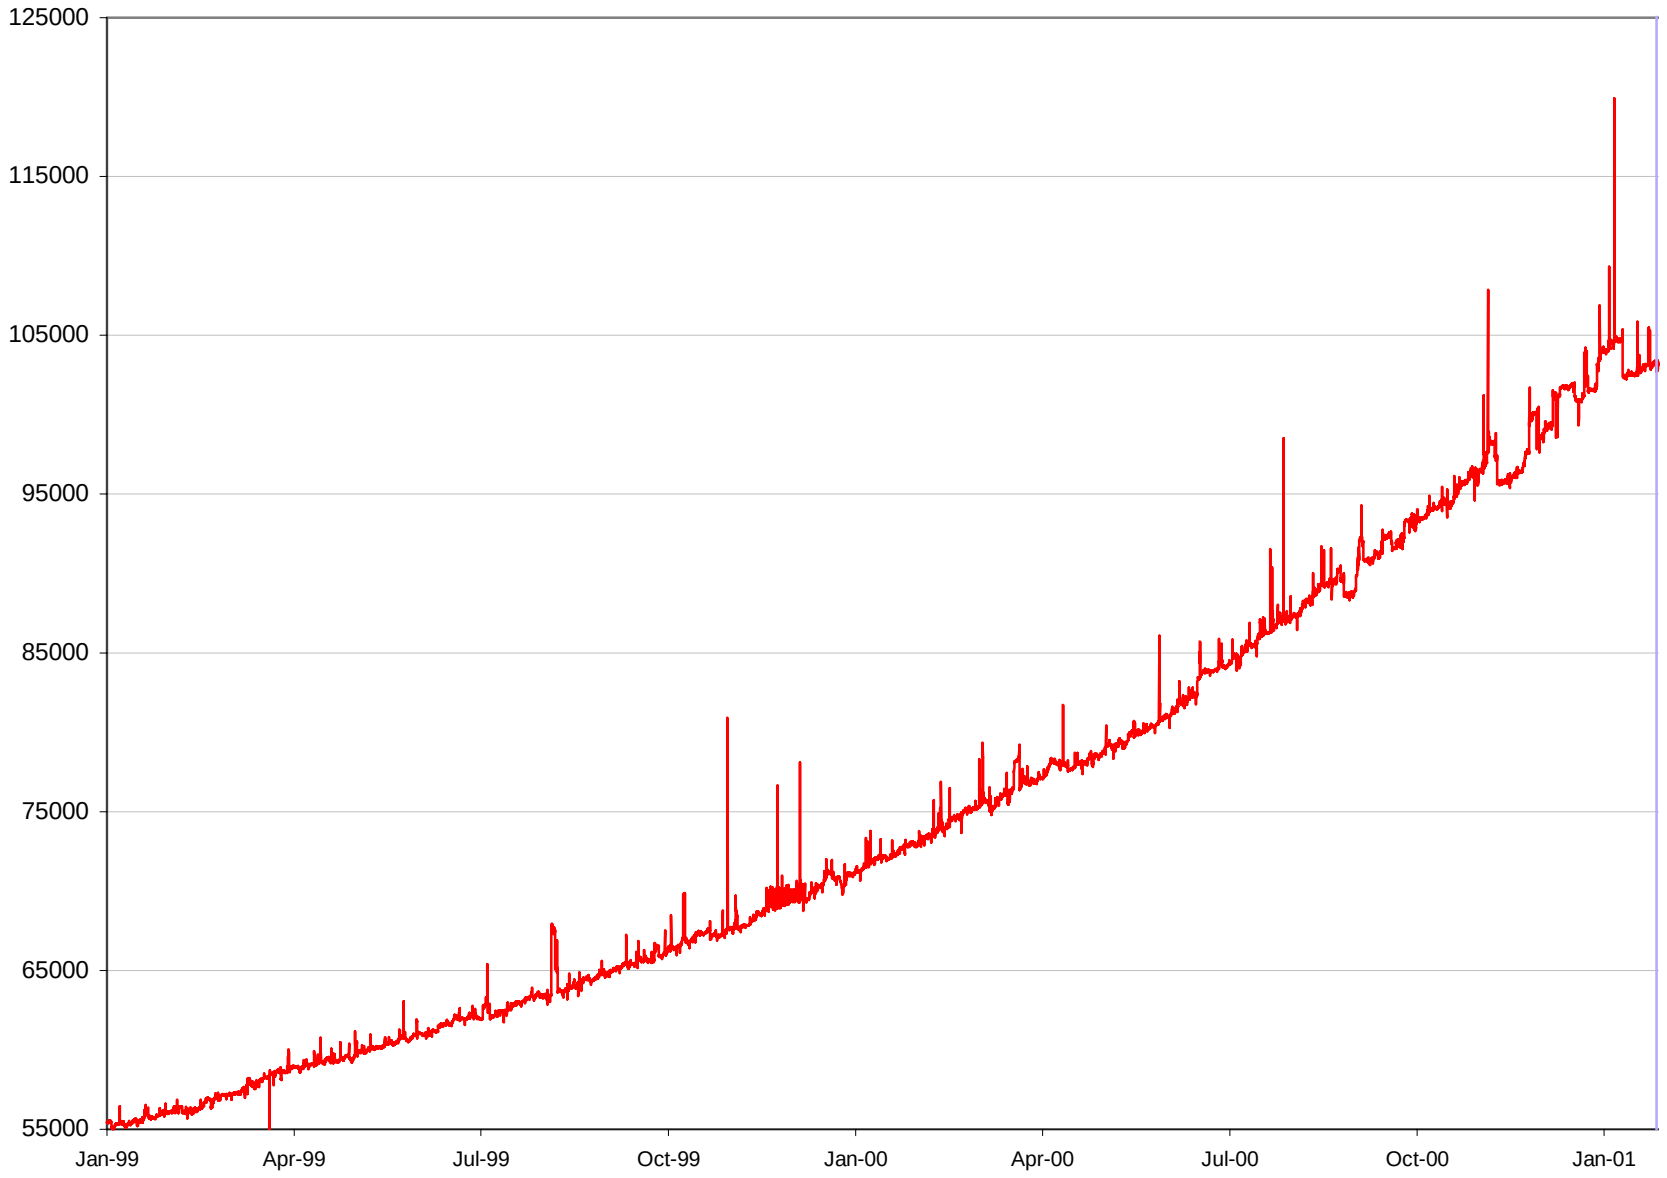

BGP Table Growth – 2 year & 6 month trends

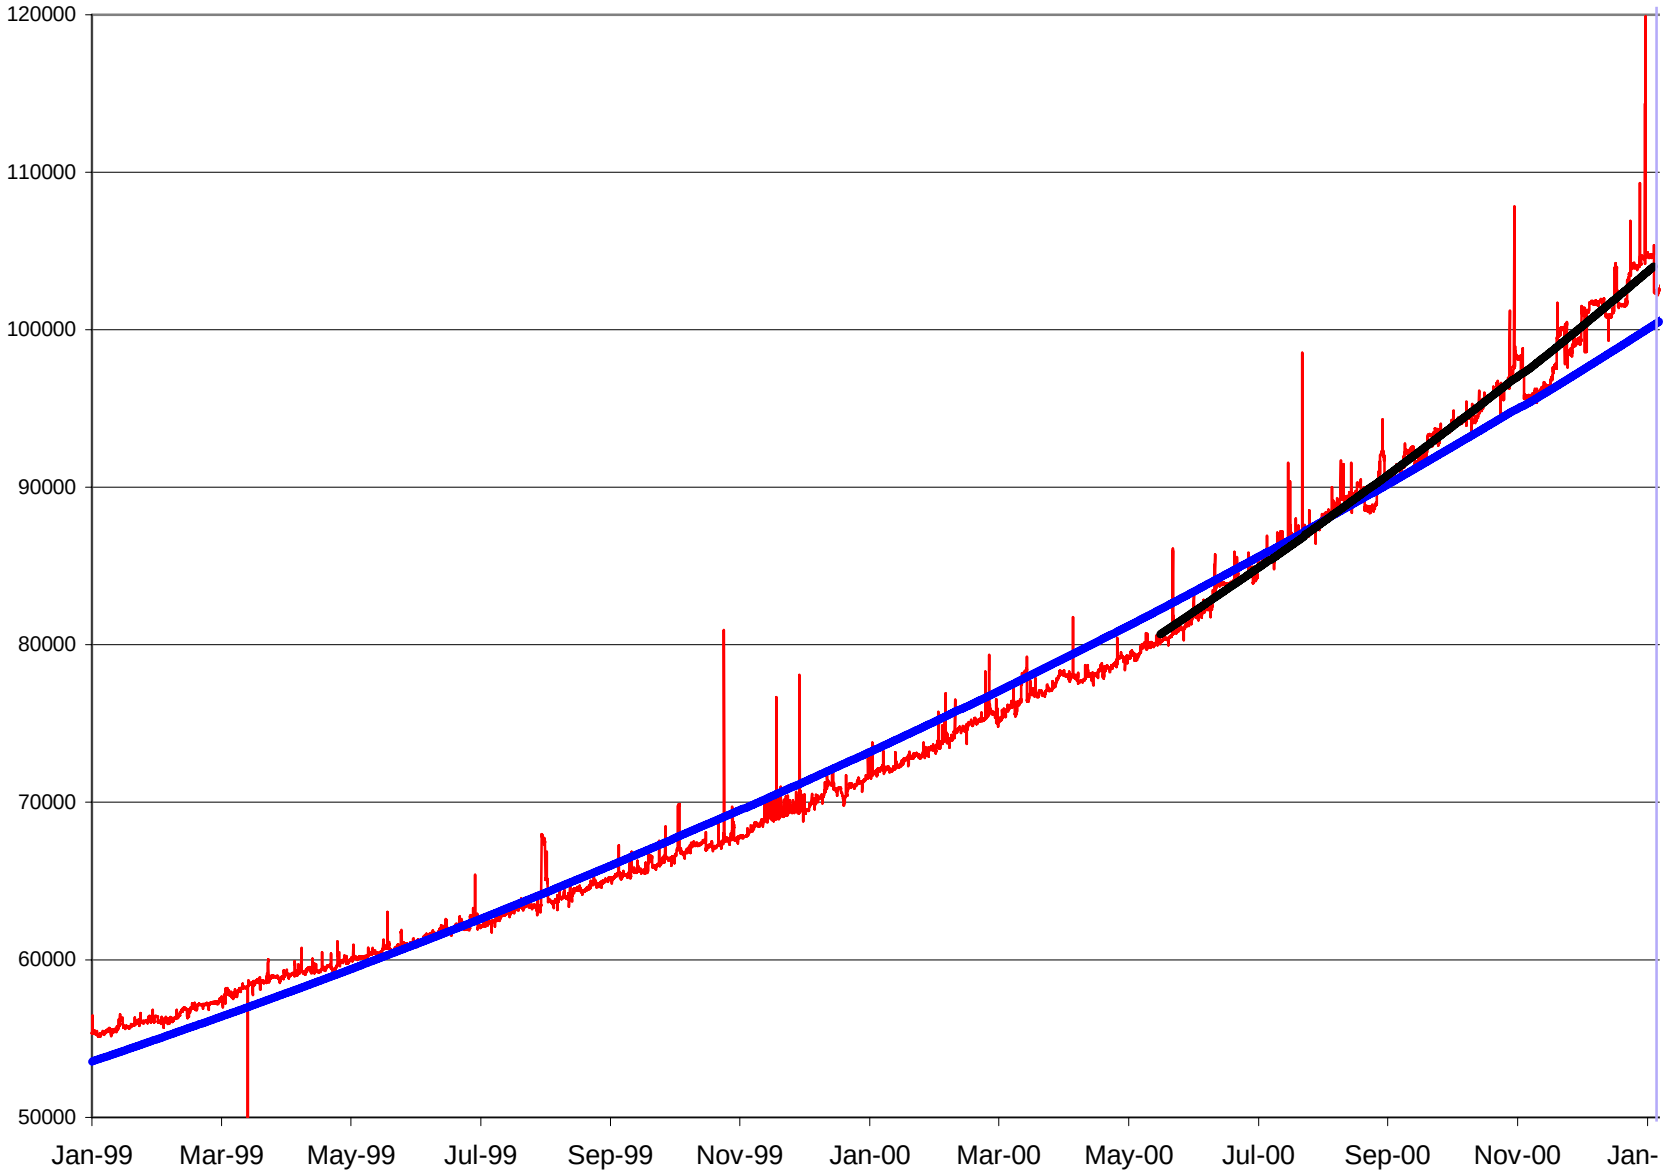

BGP Table Growth – Projections

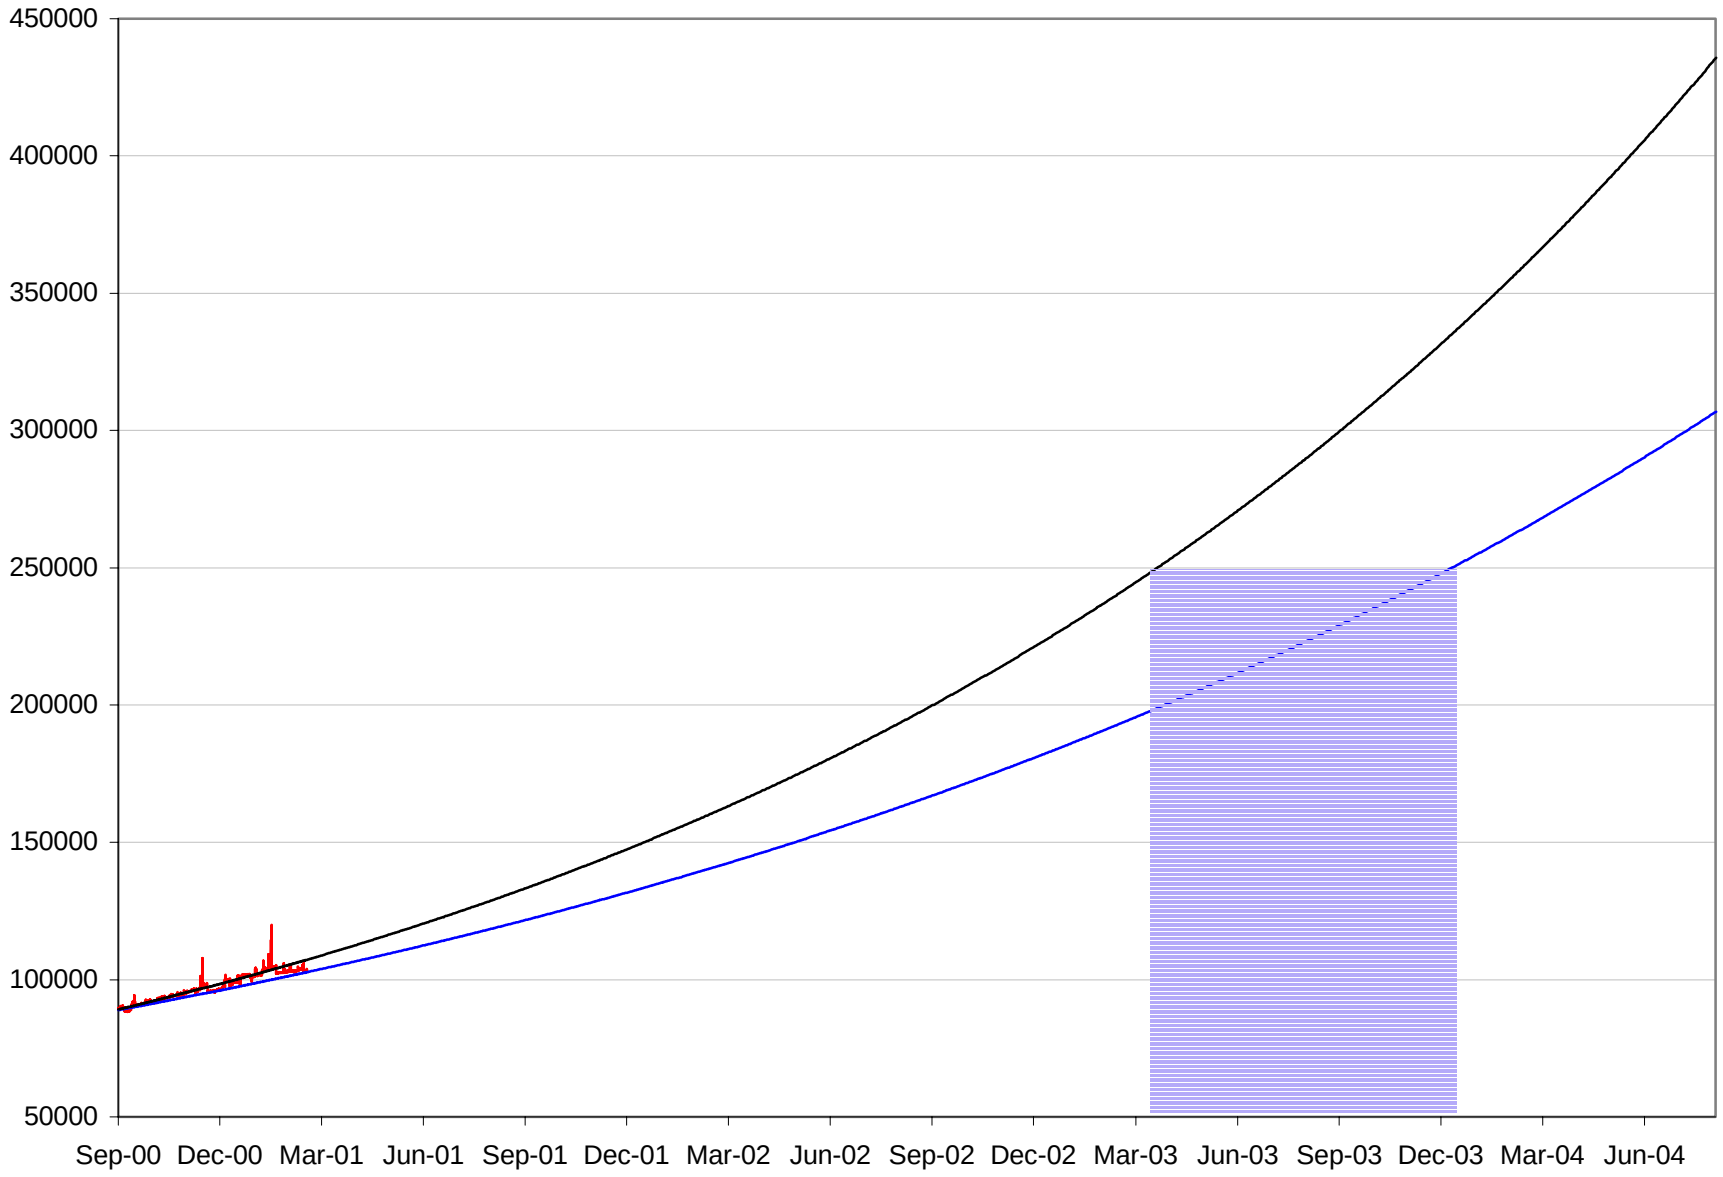

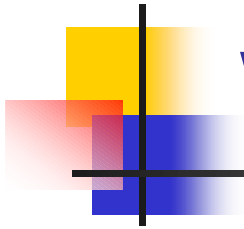

Why the recent growth?

- Inter-Domain routing carries simultaneously:
  - Connectivity information
  - Originating policy information
  - Resiliency demands
  - Traffic engineering preferences
- The growth in the BGP table space is due to a larger number of distinguished originating policies being disseminated more widely across the inter-domain space from more locations
  - More players, fewer natural points of aggregation

Technology trends for communications systems

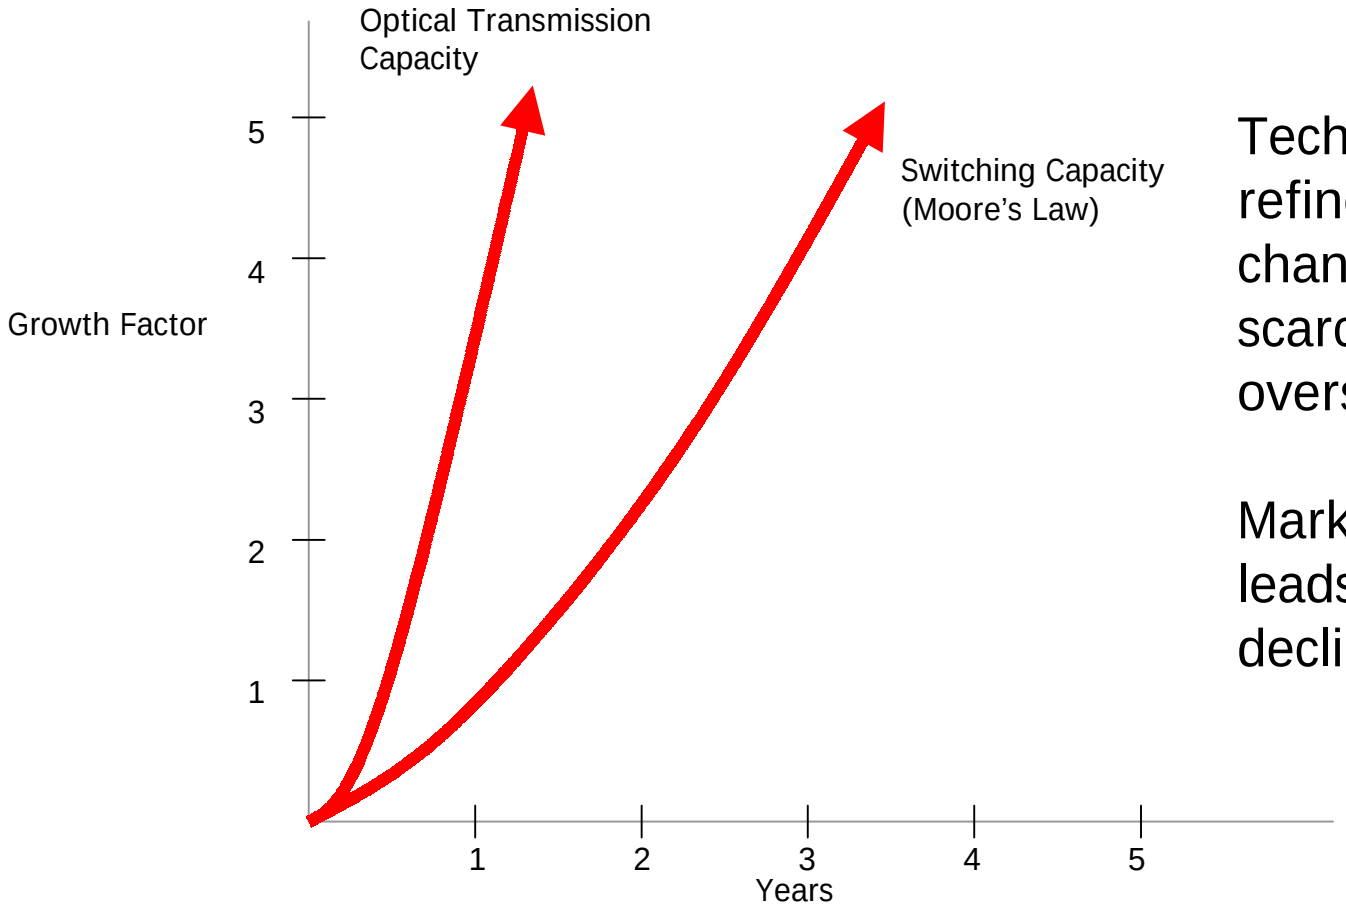

Technology refinements have changed market scarcity to market oversupply

Market oversupply leads to price decline

Cable price movements

Example: Capacity cost of international circuits between Tokyo and the US West Coast

Example Capacity Prices				
Year	Data Rate	Monthly Lease	IRU / Capital Lease	**Unit Price**
1997	E1	$ 54,000	n.a.	$ 27,000
1998	DS3	$ 540,000	n.a.	$ 12,000
1999	DS3	$ 320,000	n.a.	$ 7,111
2000	OC3	$ 200,000	$ 8,000,000	$ 1,290

As communications costs decline...

- More local consumer markets are exposed
- Local market opportunities encourage more local providers
- Local providers can substitute richer connectivity for parts of existing single upstream services
- Customers can multi-home across multiple providers to improve perceived resiliency
- Network hierarchies get replaced by network meshes interconnecting more entities

- How is this richer connectivity and associated richer non-aggregated policy environment expressed today?
  - More finer grained prefixes injected into the BGP routing system

Finer grained policy generates finer grained advertisements

Prefix distribution in the BGP table

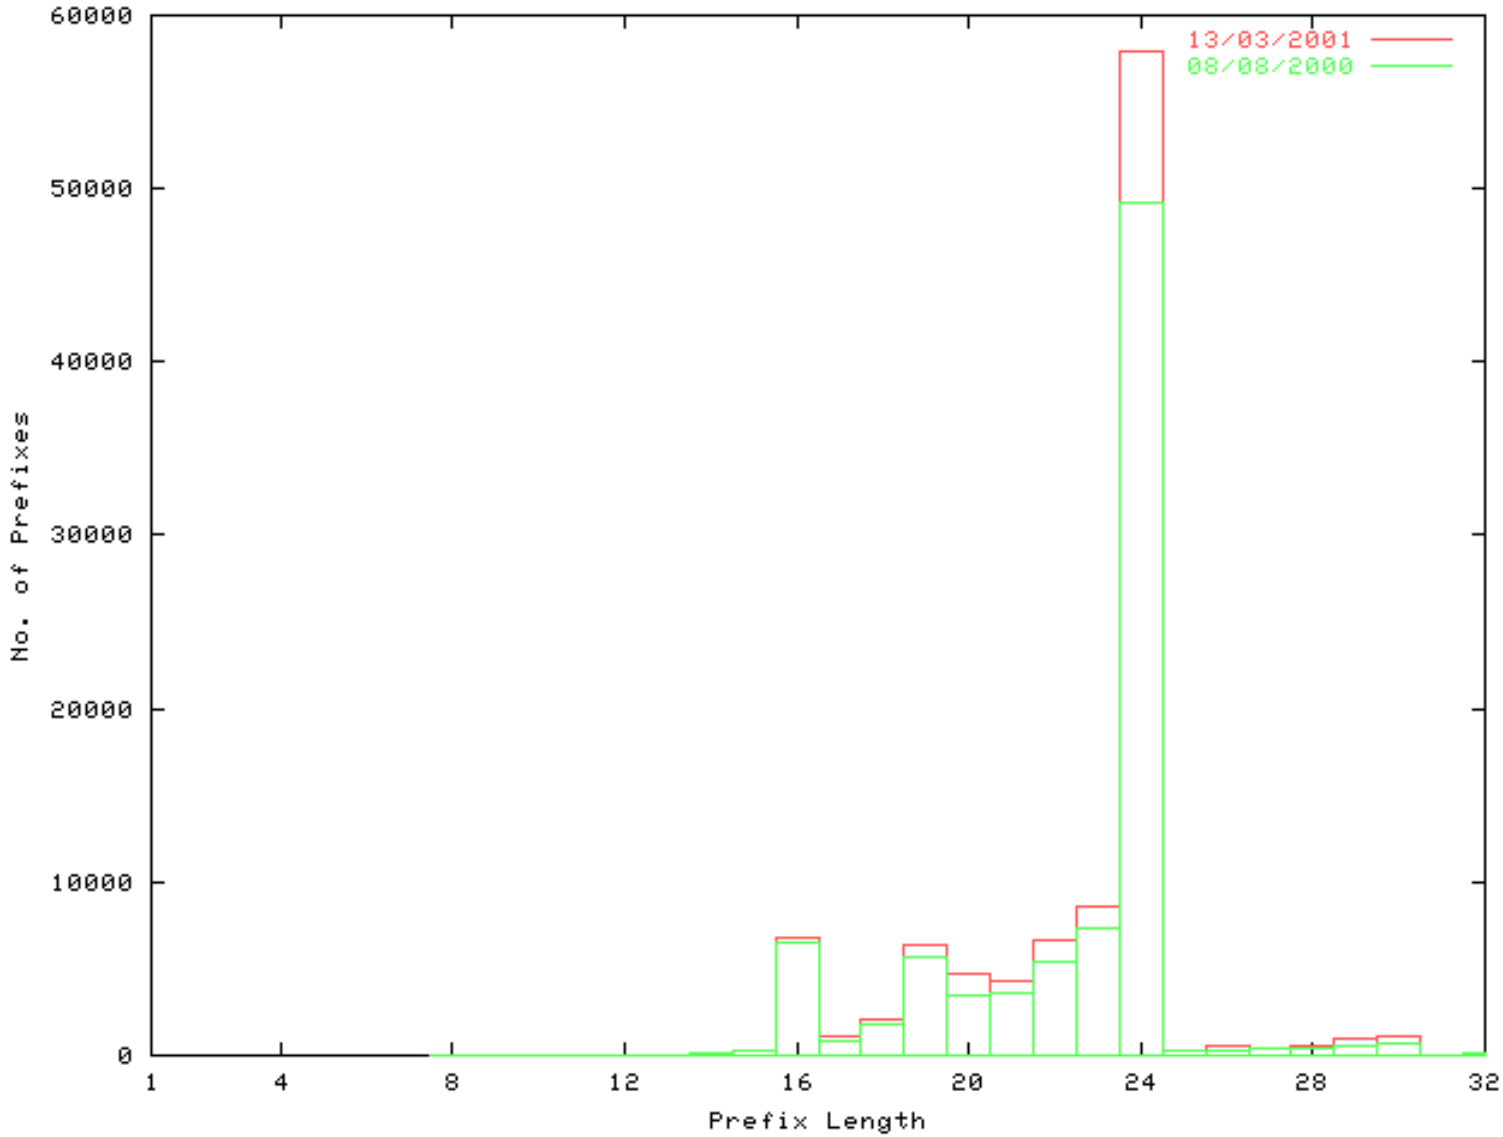

/24 is the fastest growing prefix length

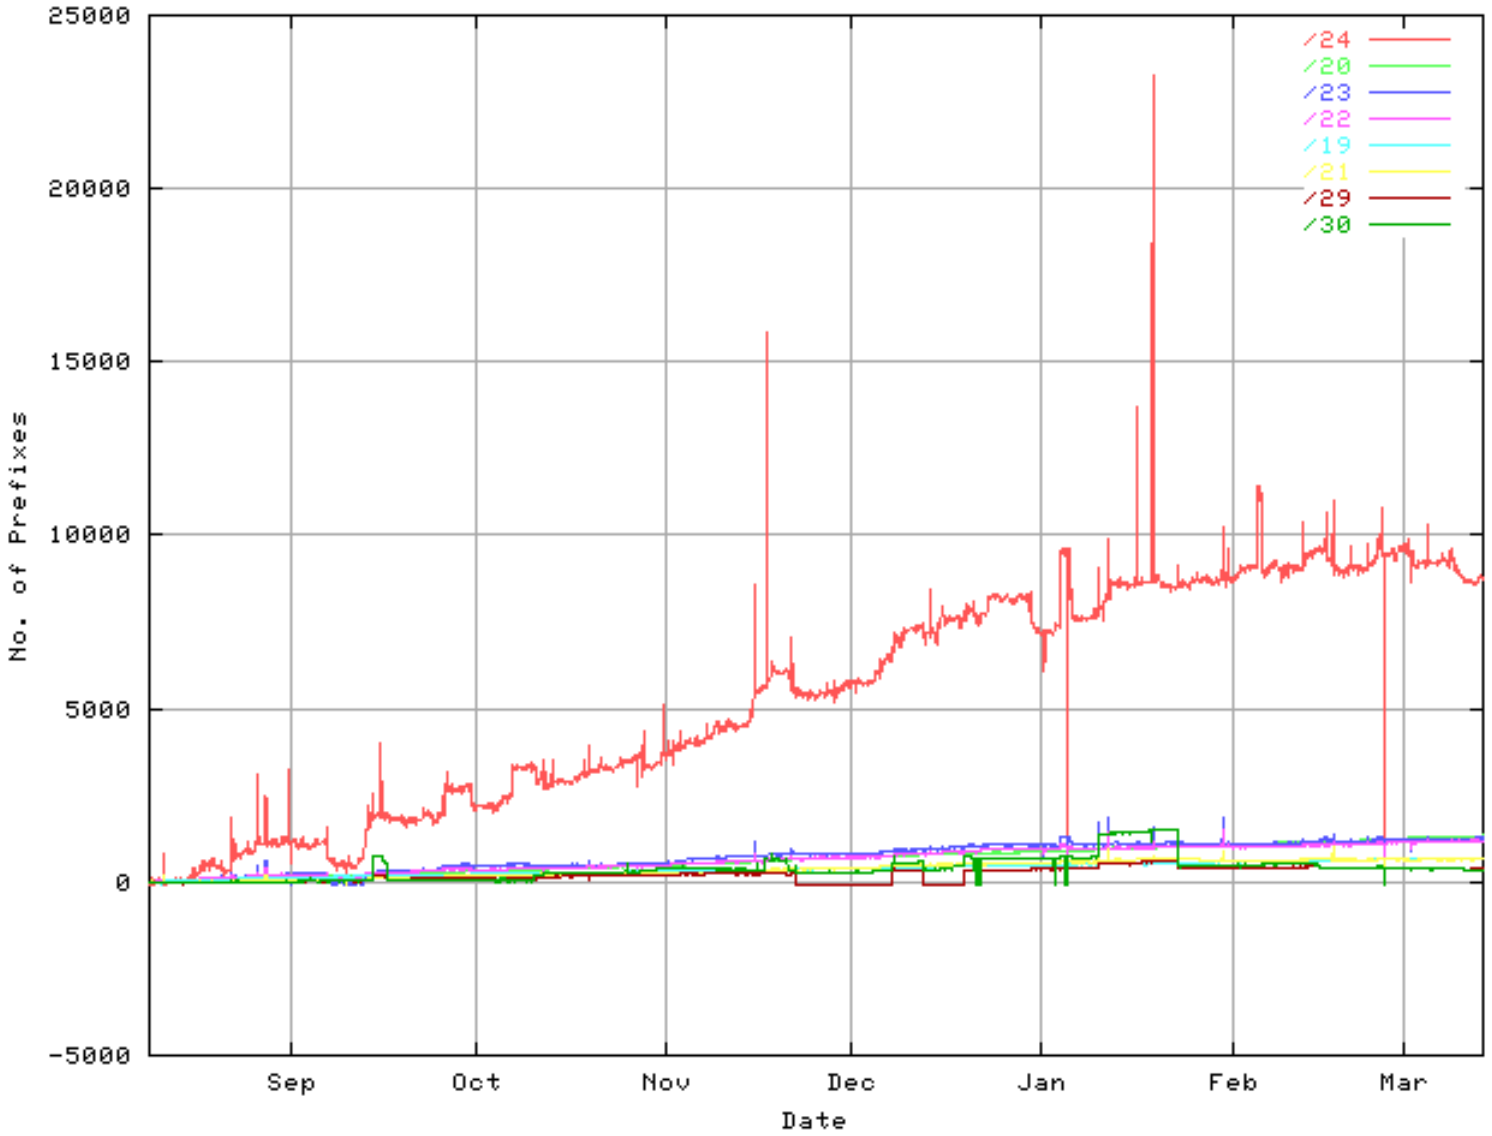

Multi-homing on the rise?

- Currently 55% of all route advertisements are routing “holes” inside existing route aggregates

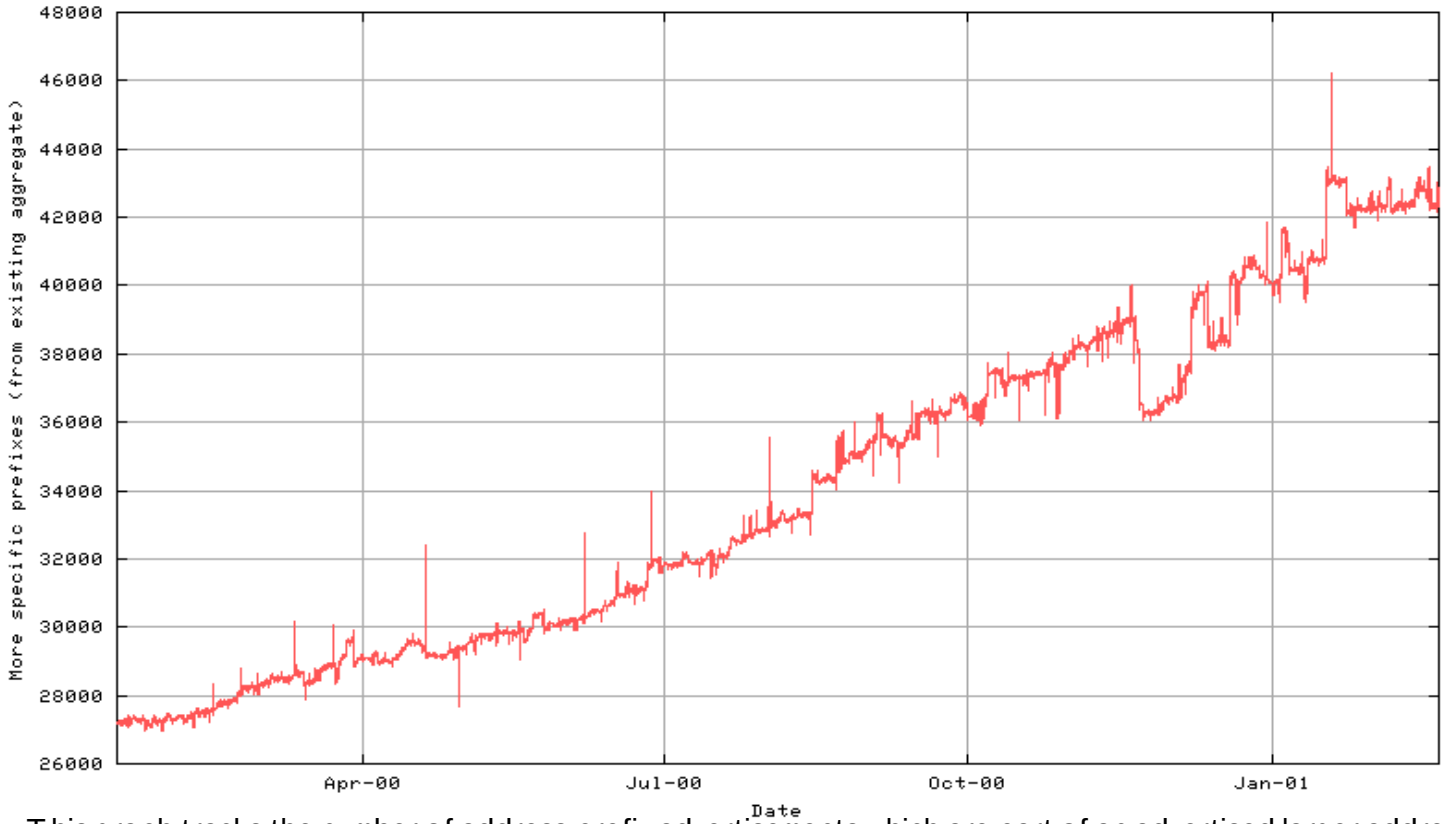

T his graph tracks the number of address prefix advertisements which are part of an advertised larger address prefix

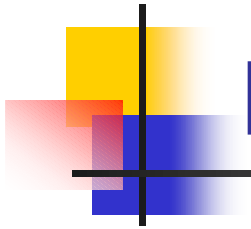

Historical scaling assumptions

- Inter-Domain routing negotiates injected **connectivity** and **policy** information
- The scaling assumption was **aggregation of policy** using provider-based hierarchies
- Aggregation of local policy was a natural outcome of high communications costs, carriage efficiency drivers within IP itself, local regulatory constraints* and knowledge and technology capability barriers

* where applicable

Change the environment and you change the Internet

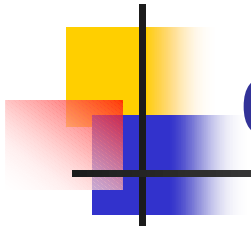

- These aggregation limitations are being dismantled:
  - Widespread deregulation, lower communications costs, greater dissemination of knowledge and higher levels of local demand
- The current environment features widespread multi-homing and traffic engineering across the entire Internet
- These environmental features are not susceptible to our scaling toolset of provider-based hierarchical policy aggregation techniques

The nature of the Inter-Domain routing space today

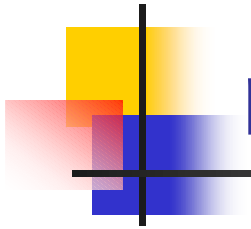

- More AS's, more small routing prefixes and shorter AS paths are all symptoms of the changing environment
  - Addressing these particular symptoms will not alter the underlying industry drivers
- A denser interconnection mesh of more visible unique entities with distinguished reachability policies is the challenge we face in inter-domain routing

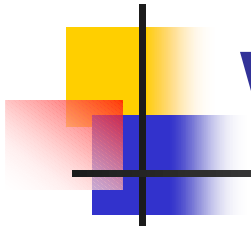

What can we do now?

- Modify – where we can - the current toolset and operational practices to meet immediate demands of continued table growth
- But - the external environment will not change, so we need to adapt our inter-domain routing tools to match the environment's needs
- Investigate the requirements of operating a very large dense mesh of connectivity, overlayed by a rich set of distinguished originating policies, modified by a collection of dynamic traffic engineering constraints

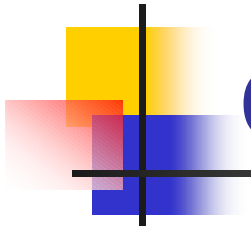

Commentary

- draft-iab-bgparch-00.txt
  - Exponential growth of BGP tables has resumed
  - Multi-Homing and Trafic Engineering are significant drivers
  - Supporting a richer set of originating policies in a denser connectivity mesh
  - What are the inter-domain routing protocol evolutionary requirements?