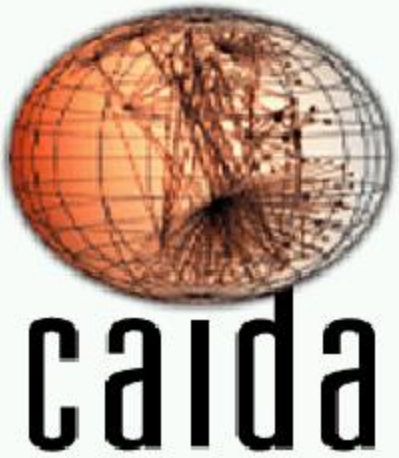

thoughts on measurement and management of the DNS system

kc claffy
Cooperative Association for Internet Data Analysis (CAIDA)
(UC's San Diego Supercomputer Center)

www.caida.org

before i give this talk

- admit not clear to me what problem we are trying to solve
  - DNS as searching and navigation system? (is that serious?)
- assess effect on name assignment, addressing, & searching
  - of growth in users and sites
  - of growth in embedded computing devices
  - of growth in personal and object identifiers
- evaluate technologies that can affect Internet searching
  - addition of generic TLDs,
  - new name assignment, addressing, and indexing schemes
  - new directory structures for locating information/sites
  - improved user interfaces for accessing info on Internet
  - navigate trademark, monopoly hell
  - evolution: competition; stability; portability

instead of giving this talk

we are living on severely borrowed time in using the DNS system for Anything At All

- much less a directory service
- no security
- no hierarchy (except 1 thin layer at top with:)
- 13 points of failure
  - 10 in US
  - 6 in same city (DC), another 4 in California
  - 3 behind military bureaucracies
  - 1 in chapter 11
  - no authority
  - no standard performance evaluation or monitoring
  - no recourse for underperformance
- a technical authority noone trusts
- a policy authority noone trusts

but hey it works (noone's more surprised than we)

“This is a crude version of a more advanced utility that has never been written.”

-- X-windows xwud(1) man-page

Problems that remain persistently insolvable
should always be suspected as
questions asked in the wrong way.

-- Alan Watts

outline of talk i could (will) give

- caida macroscopic DNS measurement activities
  - –skitter for rssac (this talk)
  - –passive measurements of gTLDs/roots from clients (nevil/evi)
  - –root server traffic analysis (evi)
- root name servers: background
- rssac project: background
- target list
- measurements
- high latency destinations
- conclusions

in case i get cut off or you lose conscoiusnesss

upshot relevant to this committee

- if you want to assess performance of the DNS system better
  - or even if you don't
- or if you do put another layer of middleware in

---> don't have 13 points of [root] failure

or if you do

---> make managing those points integral to the architecture
  - (management/modeling/modulation/measurement)

root name servers: background

existing root name servers (listed alphabetically).

Host name	Controlling organization	Location
A.ROOT-SERVERS.NET	VeriSign	Herndon, Virginia, USA
B.ROOT-SERVERS.NET	ISI	Marina del Rey, California, USA
C.ROOT-SERVERS.NET	PSInet	Herndon, Virginia, USA
D.ROOT-SERVERS.NET	University of Maryland	College Park, Maryland, USA
E.ROOT-SERVERS.NET	NASA	Moffett Field, California, USA
F.ROOT-SERVERS.NET	ISC	Palo Alto, California, USA
G.ROOT-SERVERS.NET	DISA	Vienna, Virginia, USA
H.ROOT-SERVERS.NET	ARL	Aberdeen, Maryland, USA
I.ROOT-SERVERS.NET	NORDUnet	Stockholm, Sweden
J.ROOT-SERVERS.NET	IANA	Herndon, Virginia, USA
K.ROOT-SERVERS.NET	RIPE	London, United Kingdom
L.ROOT-SERVERS.NET	IANA	Marina del Ray, California, USA
M.ROOT-SERVERS.NET	WIDE	Tokyo, Japan

highlighted root servers are monitored by CAIDA.

topology mapping project: background

- skitter
  - http://www.caida.org/tools/measurement/skitter
  - traceroute–like methodology
    - increments Time–To–Live (TTL)
    - ICMP echo requests
    - small (52–bytes) probe packets
    - slow–paced
- probes measure
  - IP forward path information
  - round trip time (RTT) to destination
  - thousands of destinations
- result
  - a ton of data (millions of paths per day, for years)
  - most comphrehensive macroscopic Internet topology data in world
    - (low bar)

DNS Clients list

- common list to run on all monitor probes:
  - combine individual clients lists from all root name servers
  - stratify routable IPv4 address prefix space
- DNS clents list for this study was created in September 2000
  - 49,374 addresses passively collected from root servers
  - 8,944 addresses from other CAIDA lists

=> cover more than 58,000 prefixes (out of nearly 90,000 in the BGP table)

- augmenting list as new data from root servers available

DNS Clients list: characteristics

Top level domains		Origin ASes		Countries	
com	11345	AS 701, ALTERNET	1660	USA	31172
net	8697	AS 1, BBN Planet	577	Canada	3276
au	1929	AS 7018, AT&T	546	Australia	2645
edu	1763	AS 3561, Cable & Wireless	538	unknown	2373
jp	1376	AS 2914, Verio	472	Germany	1681
ca	1212	AS 1785, Applied Theory Corp.	472	Japan	1285
org	969	AS 1239, Sprint	467	United Kingdom	1061
de	891	AS 1221, AARNET	428	France	981
us	854	AS 2200, INRIA-Rocquencourt	358	Mexico	803
mil	673	AS 2907, SINET	335	South Korea	794

"Top Tens" of the DNS Clients list.

DNS Clients list: characteristics

distribution of destinations in the DNS Clients list by continents

measurements at each monitor

- probes DNS Clients list 7–13 times per day
- reaches from 33,000 to 36,000 destinations per day
  - dips on weekends
  - decreasing by ~2% per month

unique destinations replying per day, march 2001

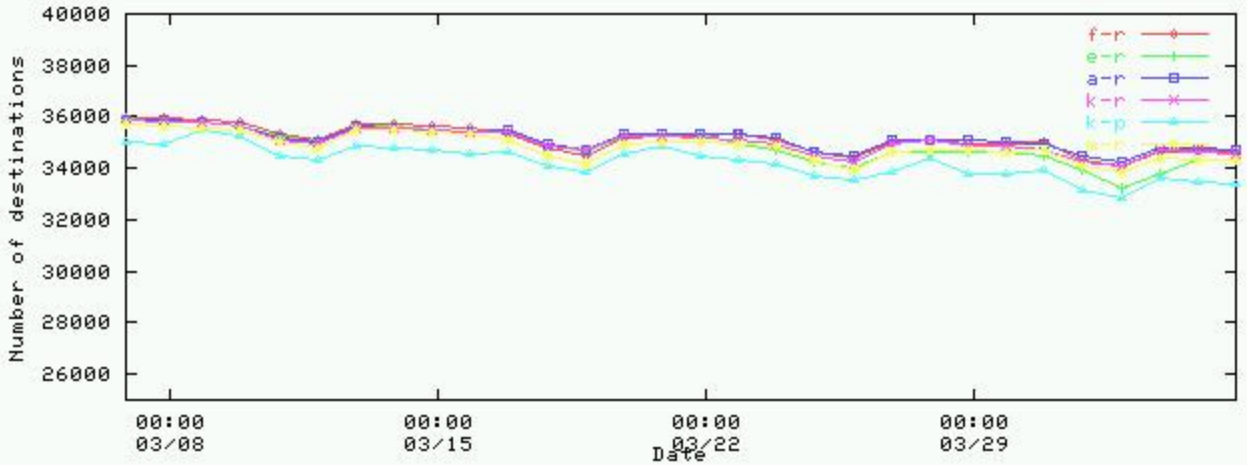

measurements at each monitor (continued)

- collects between 250,000 and 450,000 RTT values daily

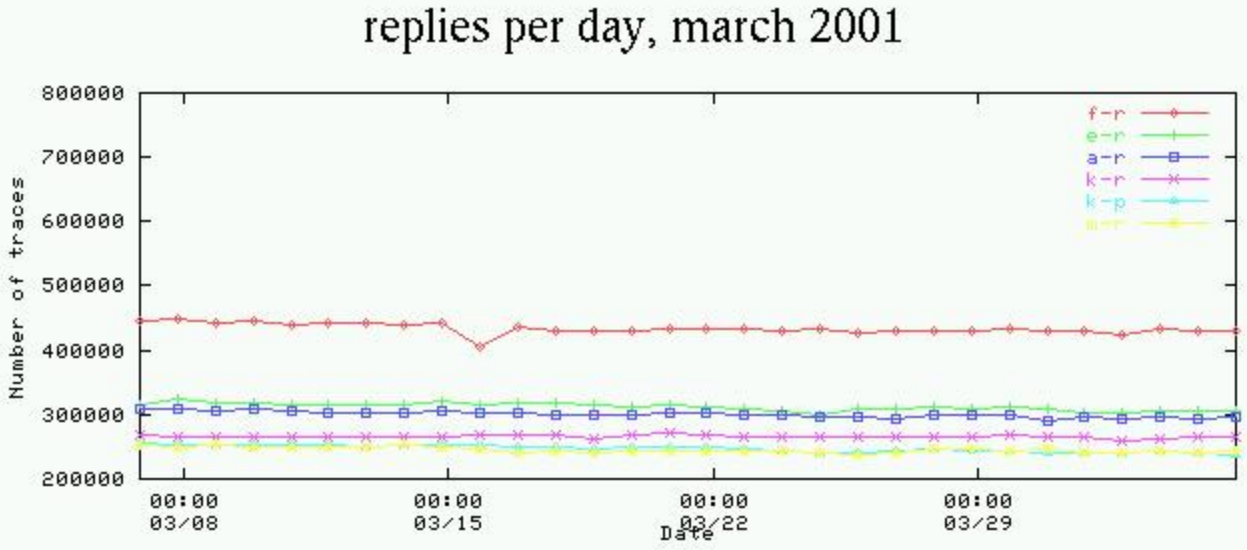

high latency destinations (HLD): definitions

- consider RTT distributions in each cycle of probes
  - large diurnal variations in RTT values
- RTT is high if above 90th percentile in given cycle (.5–1s)
- a destination is high latency on a given day if it had:
  high RTTs in at least half the cycles on all root server monitors
- aggregate two 30–day long sets of data:
  - 1 – 30 December 2000
  - 6 March – 4 April 2001

high latency destinations: persistence

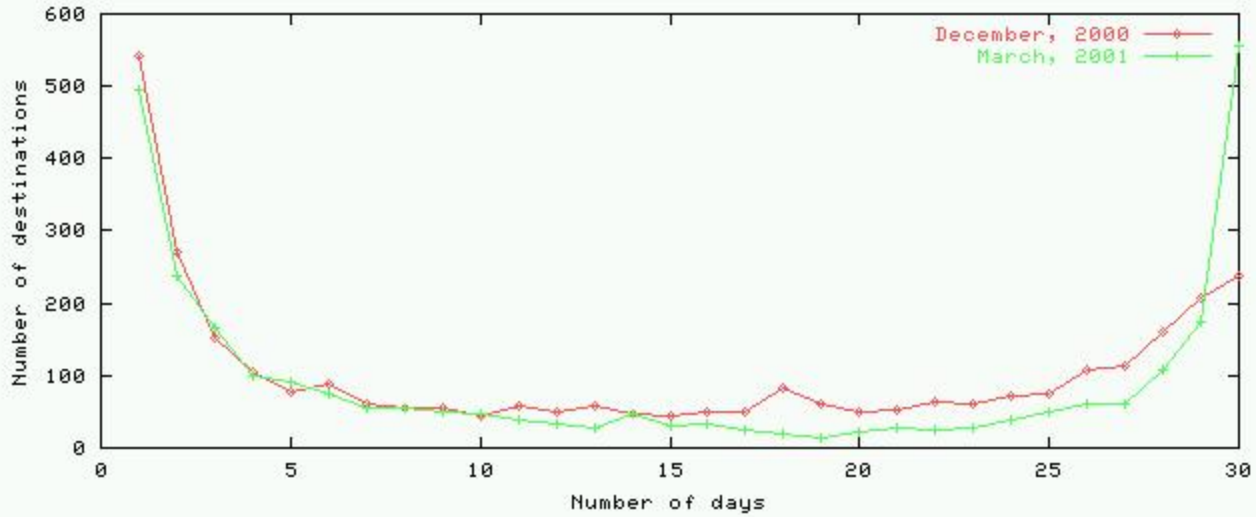

- left peak: random variations in connectivity
- right peak: consistently high latency (RTT) destinations

high latency destinations: by origin ASes

	December 2000			March 2001		
AS number	# of targets	# with high latency	% with high latency	# of targets	# with high latency	% with high latency
AS 3741, Internet Solution	92	55	60	102	49	48
AS 4755, APNIC	204	49	24	174	22	13
AS 7545, APNIC	128	30	23	138	30	22
AS 2905, TICSA-ASN	38	24	63	38	21	55
AS 2277, ECUANET	35	19	54	32	21	66
AS 7633, APNIC	28	19	68	29	19	66
AS 10530, Interpacket Group	101	18	18	72	25	35
AS 11127, NetSat Express	39	18	46	28	15	54
AS 6140, IMPSAT ARGENTINA	59	16	27	55	14	25
AS 6471, ENTEL CHILE	55	15	27	55	16	29
AS 6453, Teleglobe	54	14	26	68	18	26
AS 3132, Red Cientifica Peruana	23	14	61	24	10	42
AS 8143, Publicom	29	13	45	26	15	58
AS 7087, COLOMSAT	25	13	52	24	10	42
AS 9241, APNIC	15	12	80	15	11	73
AS 2018, UNINET-ZA	44	30	68			
AS 4621, APNIC	21	17	81			
AS 2614, RIPE	18	14	78			
AS 1239, SprintLink	398	10	3			
AS 6429, AT&T Chile Internet				137	73	53
AS 3255, RIPE				34	24	71

high latency destinations: by countries

			December 2000		March 2001	
Continent	Country	# of targets	# with high latency	% with high latency	# with high latency	% with high latency
Asia	India	382	94	25	75	20
	Indonesia	165	35	21	36	22
	Pakistan	88	18	20	21	24
	Russia	437	14	3	15	3
	Thailand	315	25	8		
	Jordan	31	11	35		
	Georgia	15	10	67		
	Turkey	175			15	9
	Bangladesh	13			11	85
Europe	Romania	377	86	23	55	15
	Ukraine	185	30	16	71	38
	Bulgaria	203			14	7
North America	USA	31172	71	0	74	0
	Costa Rica	35	12	34		
South America	Ecuador	90	34	38	40	44
	Chile	375	30	8	142	38
	Argentina	592	27	5	19	3
	Colombia	213	25	12	21	11
	Peru	88	19	22	17	19
	Brazil	411	14	3		
Oceania	Australia	2645	29	1	23	1
	Fiji	13	10	77		
Africa	South Africa	268	124	46	79	29

high latency destinations: differences between two data sets

- number of HLDs in India, Romania and South Africa has decreased by 20%, 36% and 36%, correspondingly.
- number of HLDs in Ukraine more than doubled, and in Chile it increased almost 5−fold.
- Thailand, Jordan, Georgia, Costa Rica, Brazil and Fiji contributed each more than 1% of the HLD subset in December 2000.
- Bangladesh, Turkey, Bulgaria and Nigeria contributed each more than 1% of the HLD subset in March 2000.

high latency destinations: differences between two data sets

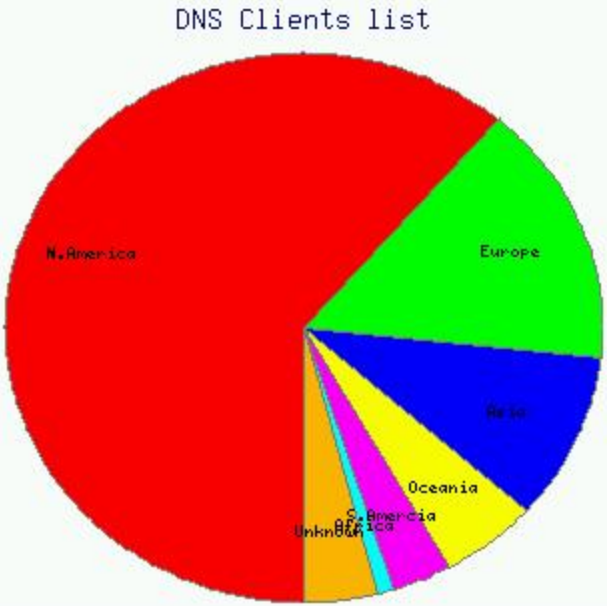

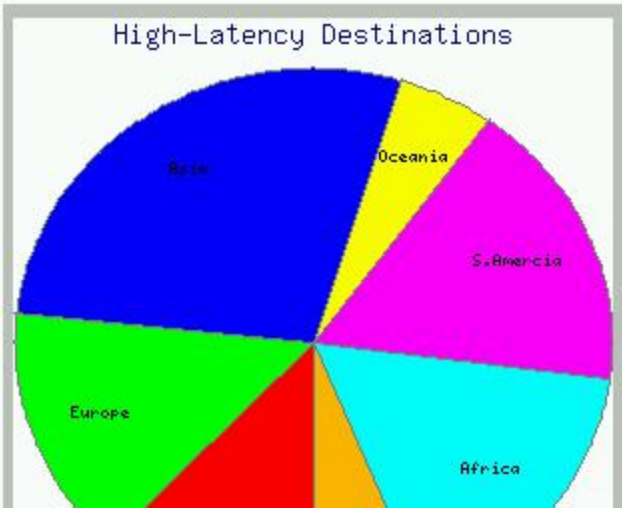

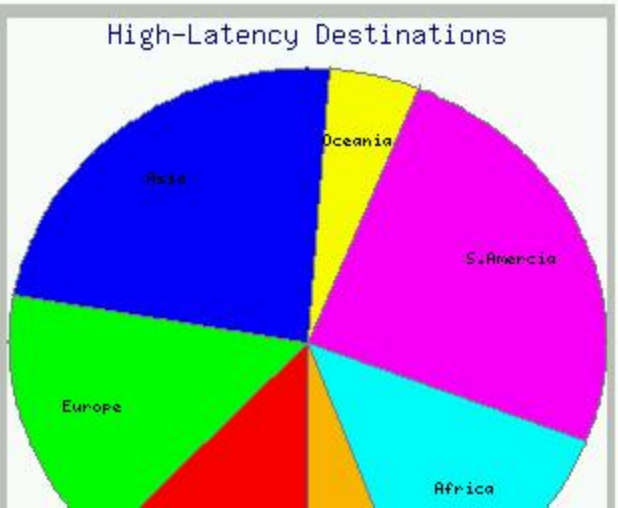

High latency destinations: by continents

- general geographic pattern same in both data sets
- number of HLDs:
  - in Asia decreased slightly
  - in South Africa increased slightly
- largest proportions of HLD (relative to the target list):
  - Africa
  - South America
  - Asia

conclusions

- topology & performance data scant
- need to monitor ALL 13 root servers to minimize bias in identifying high–latency destinations
- high latency: last mile bandwidth or topology deficiency?
  - further examination with other tools to assess cause of the high latency
- need to expand to gTLD servers
- future root/gTLD server candidate sites should run a monitor for at least 6 months

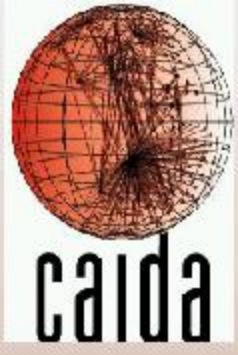

www.caida.org/outreach/presentations/

kc claffy
UCSD/SDSC/CAIDA
kc@caida.org
www.caida.org