An Analysis of ISP Backbone Availability

Katsushi Kobayashi
ikob@ni.aist.go.jp

- All results in this talk are based only with the IS-IS messages provided by Internet2 observatory. Therefore, the results of specific links and nodes in this presentation are not directly reflect the quality of its service, and/or of its equipment.

How much availability in ISP infrastructure.

- Your ISP offers 99.9% SLA for intra-ISP,
  - really premium ?
  - valuable to pay more ?
- Just presenting infrastructure availability, not taking into account :
  - Any convergence delay of routing protocol
  - Packet behavior

Internet infrastructure : viewpoint from Routing

- Breakdown network failures into its causes:
  - Routing and centralized-NMS (Labovitz ’99)
- A lot of BGP activities
  - BGP failures affects world wide Internet system
  - BGP can be seen by other ISP’s
  - BGP continues to be recorded as UO’s RouteViews

ISP infrastructure : viewpoint from IGP

- Fewer IGP activities than BGP
  - IS-IS on Qwest , Alaettinoglu (‘02)
  - OSPF on Michi-Net, Watson (‘03)
- required to install collector ISP network inside.
- IGP dataset will disclose ISP backbone quality.
- or, It is not a news network is working fine :)
- IGP message represents infrastructure events
  - Lost adjacency, ext. route : circuit / switch / interface down
  - Est. adjacency, ext. route : circuit / switch / interface up
  - Lost LSP/LSA : router down
  - Reset LSP/LSA seq. : router up

IS-IS collector in Abilene

- IS-IS collector is part of I2 Abilene observatory activity.

  http://ndb2-blmt.abilene.ucaid.edu/isis/

  Contributed by Shu Zhang [ZK06]

- Deployed all Abilene nodes for multi observation points.
- Synchronized with CDMA timer (GPS based)
- From Aug. ’04 to Apr. ’07 data set is available.

[ZK06] S. Zhang and K. Kobayashi, “Rtanaly: A System to Detect and Measure IGP Routing Changes”

Abilene Network Map

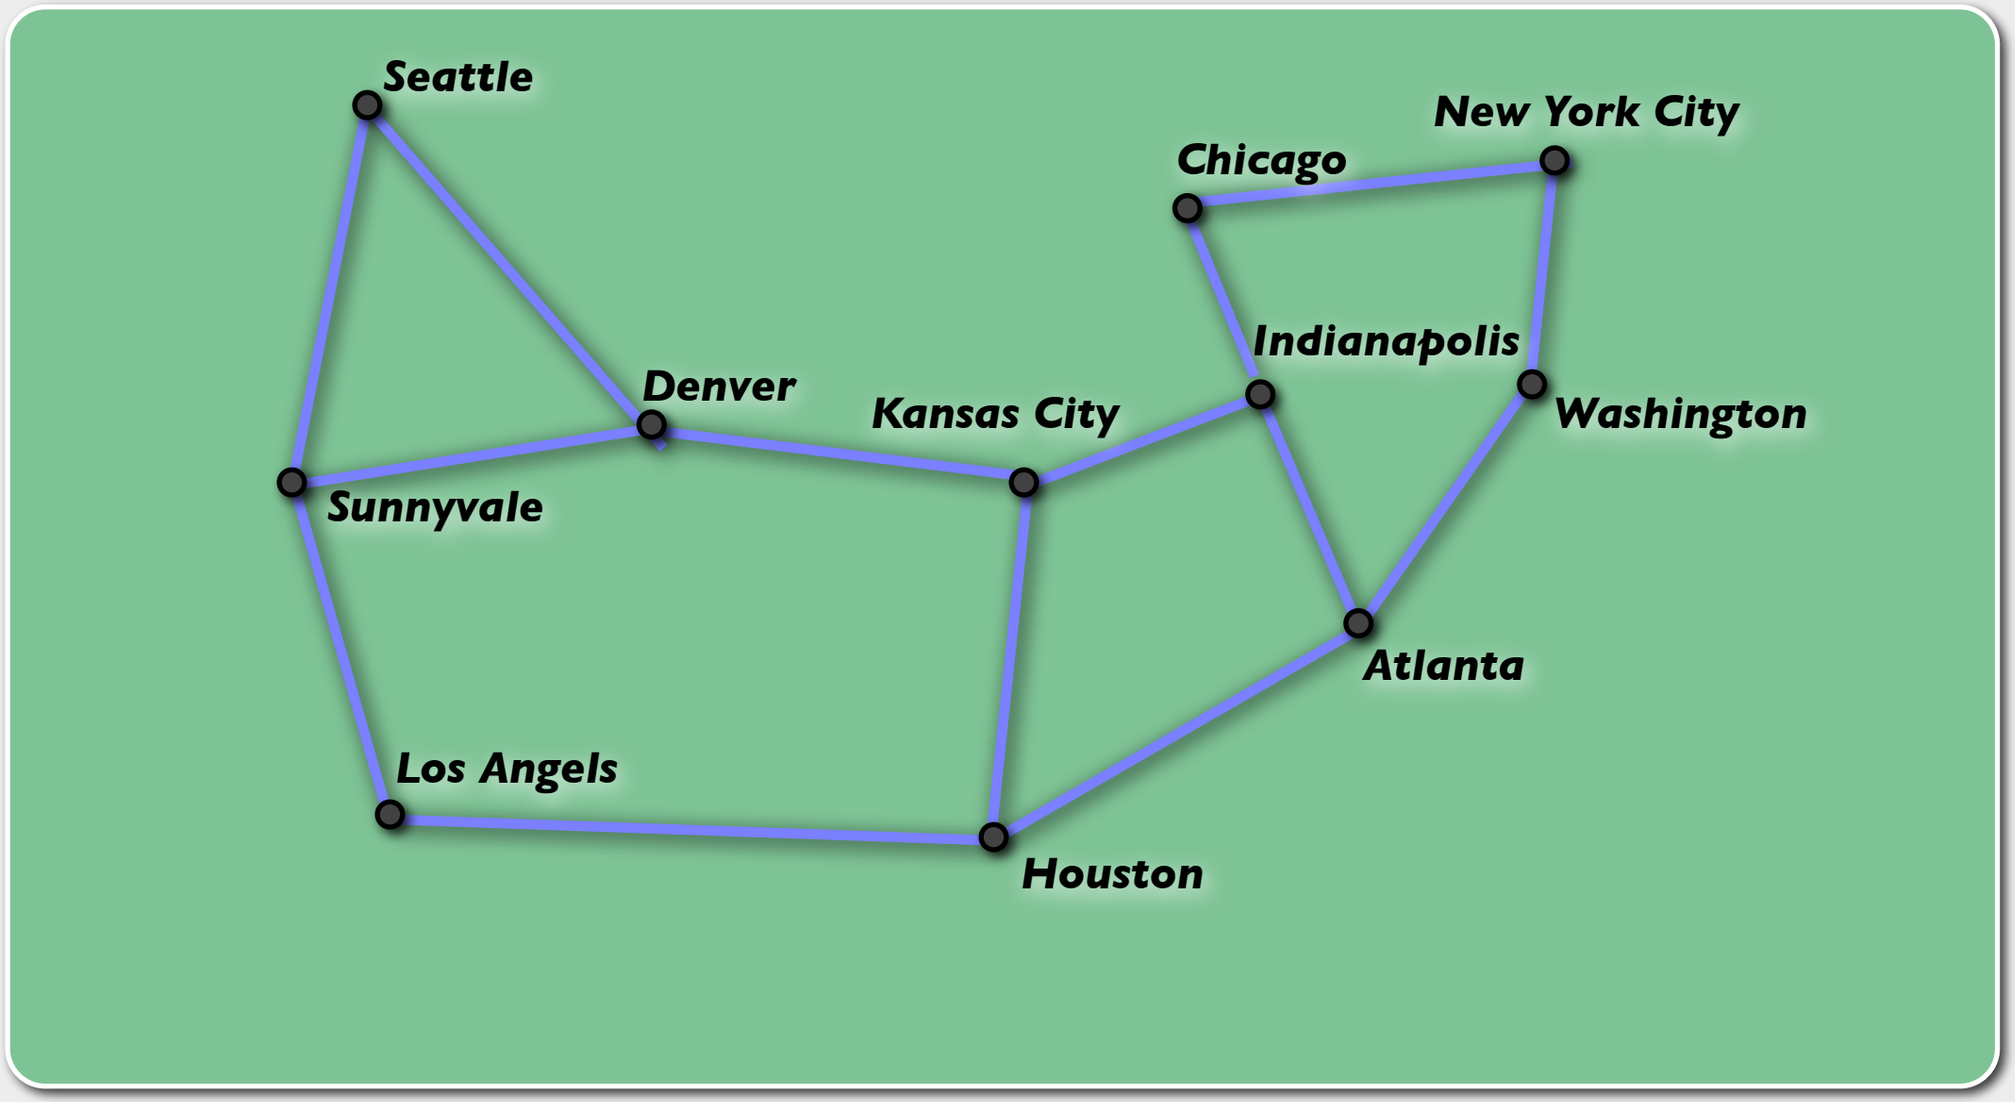

11 nodes with T640 routers, and 14 OC192 circuits.

Abilene IS-IS operation

- 9 sec. Hello interval, lost ISIS adjacency after missing 3 hellos
  - 22.5 sec. failure detection delay is supposed.
    - More faster failure detection is possible, e.g., shorter hello interval, BFD, carrier loss with circuit failure.
- IGP maintains infrastructure information only.
  - Minimize IGP database
  - Not import any BGP route into IS-IS.

- Network availability in hereafter:
  - All network works without any failure.
    - From Network operator’s viewpoint.
  - Don’t care specific source destination path availability.
    - Not from customer’s viewpoint.
- Timeframe:
  - May include more than one event at same time.

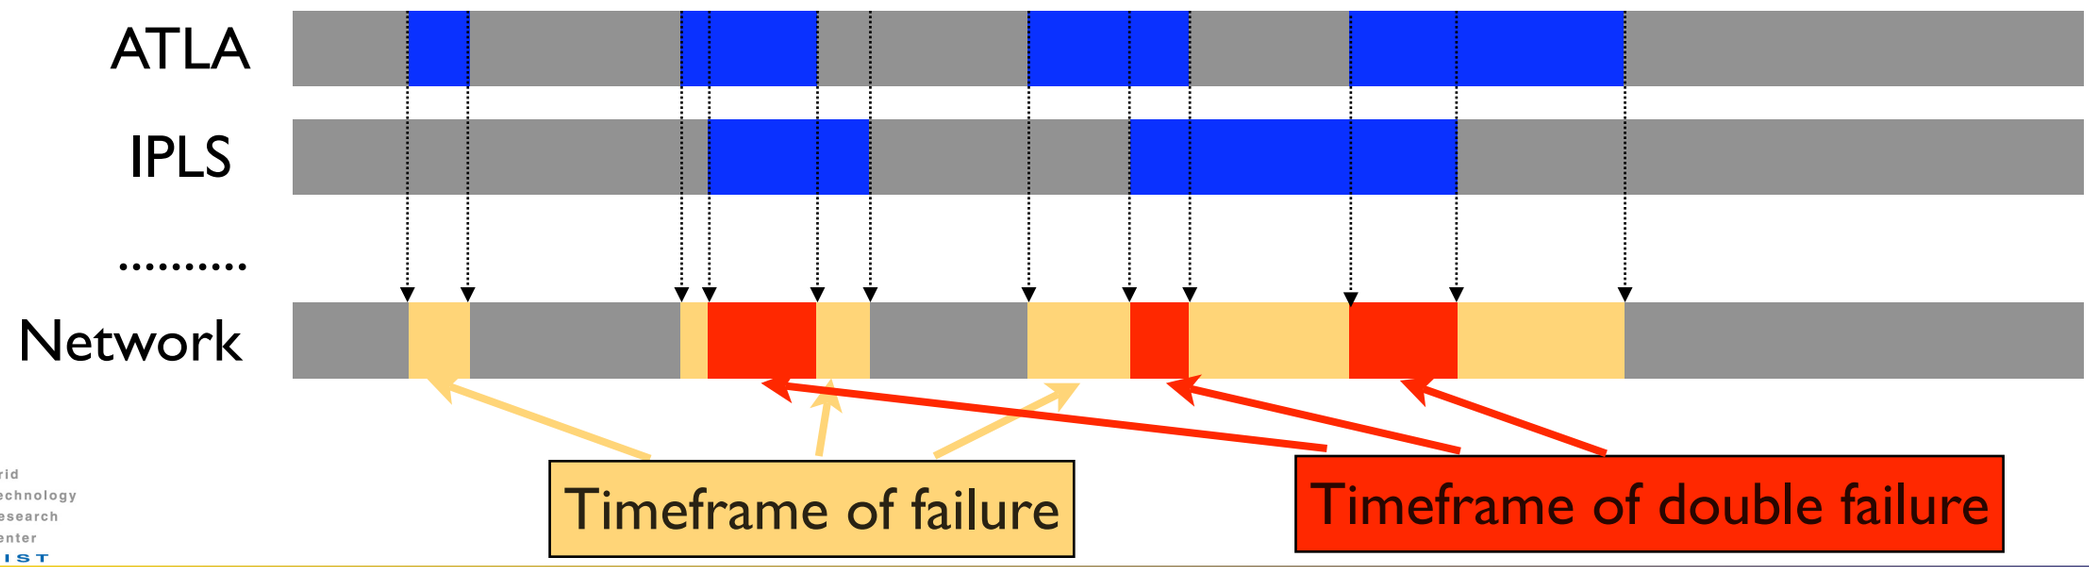

Abilene IS-IS overview ’05-’06

- Node failure: timeout node LSP, or seq. number reset.
  - Only 1 times on ’05 (53 sec. downtime), 2 on ’06 (1,298 sec.)
- Circuit failure: adjacency away from list in LSP
  - Usually found, 635 timeframe on ’05, 513 on ’06.
- Ext. route failure: Route away from LSP
  - Represent edge troubles ?
  - Difficult to identify whether serious or trivial.

To focus this failure.

Lost adjacency event

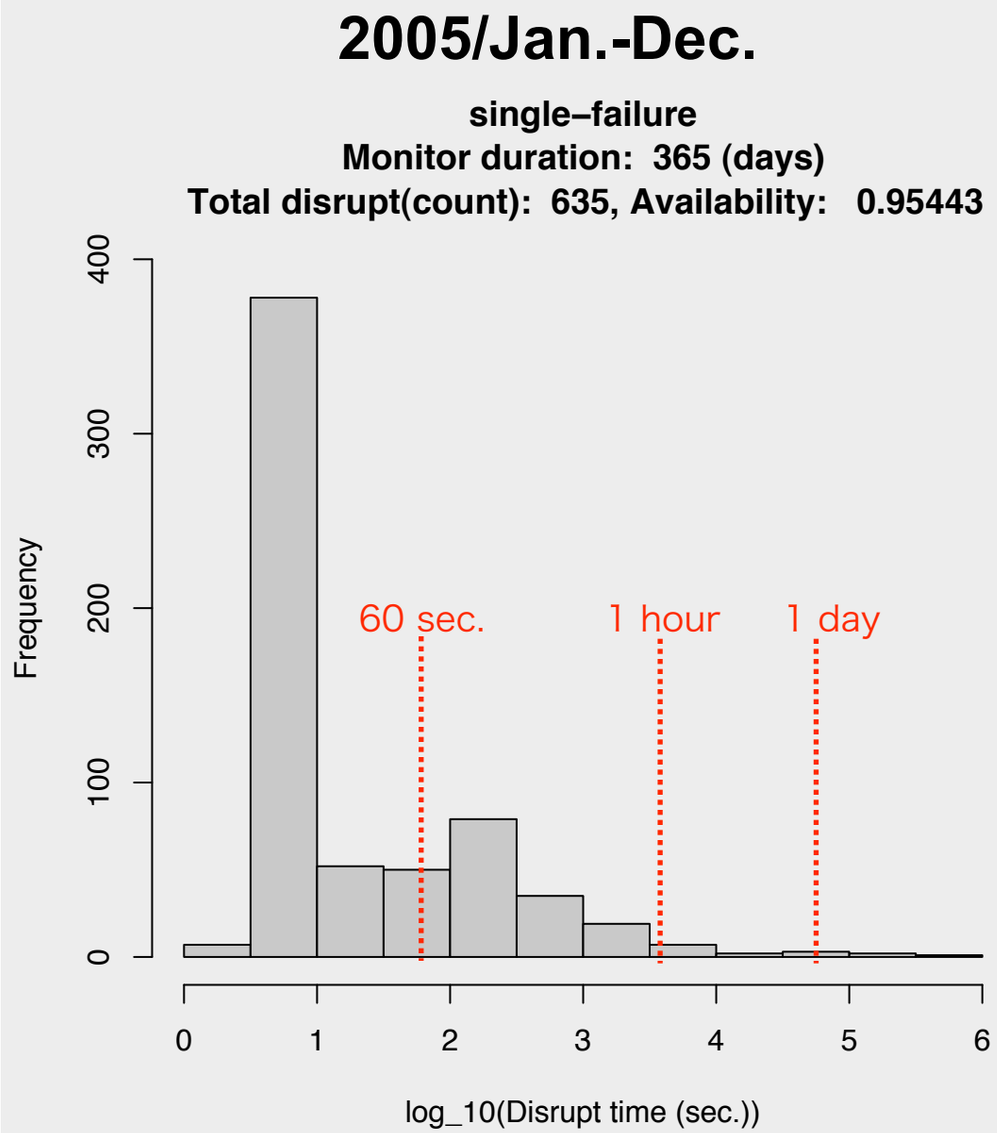

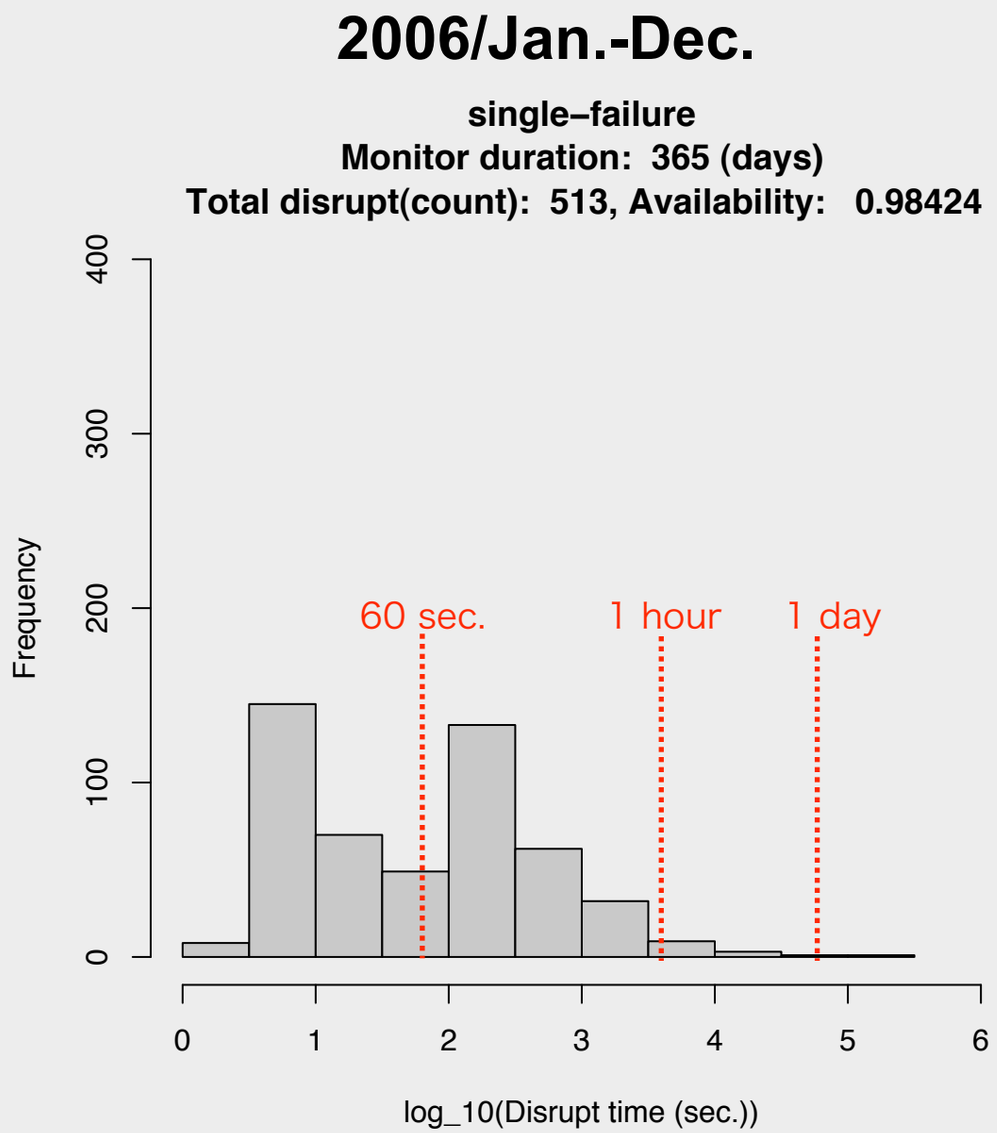

Note that above histograms are drawn with IS-IS captured data at Atlanta.
Few details are different with other IS-IS observatory point.

Breakdown in ‘05

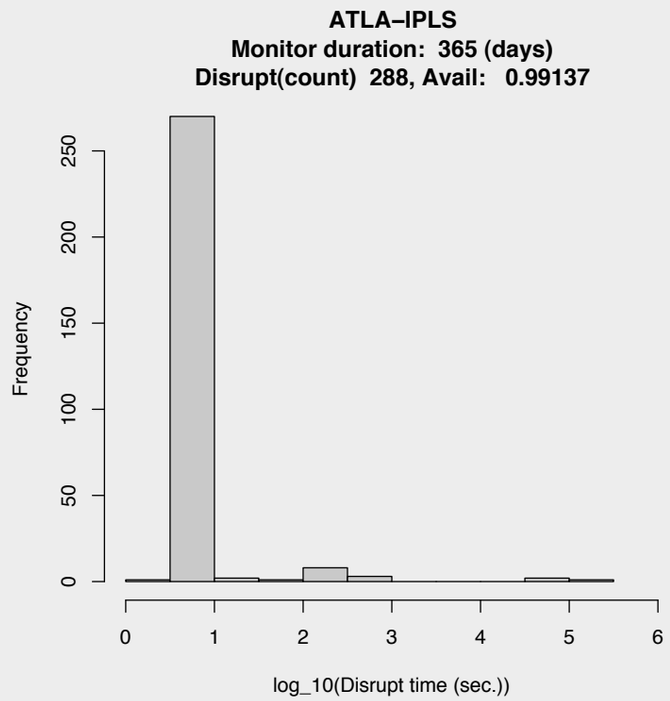

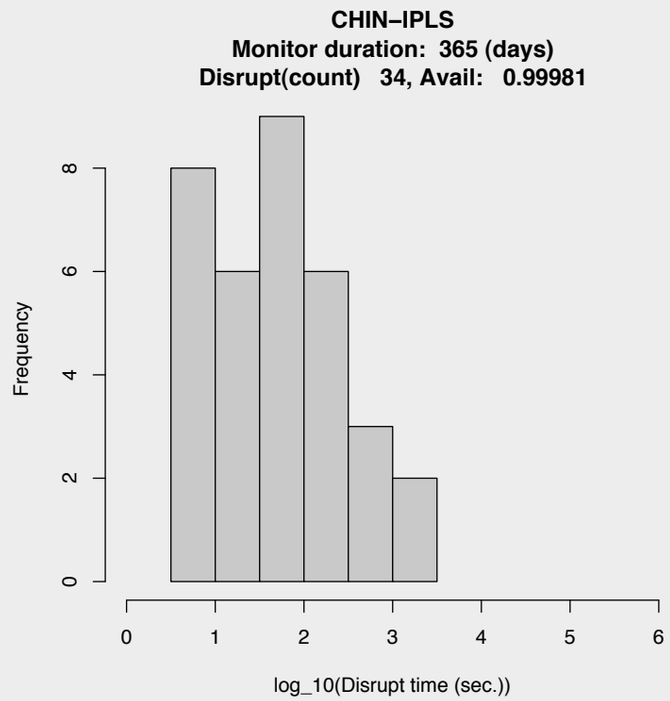

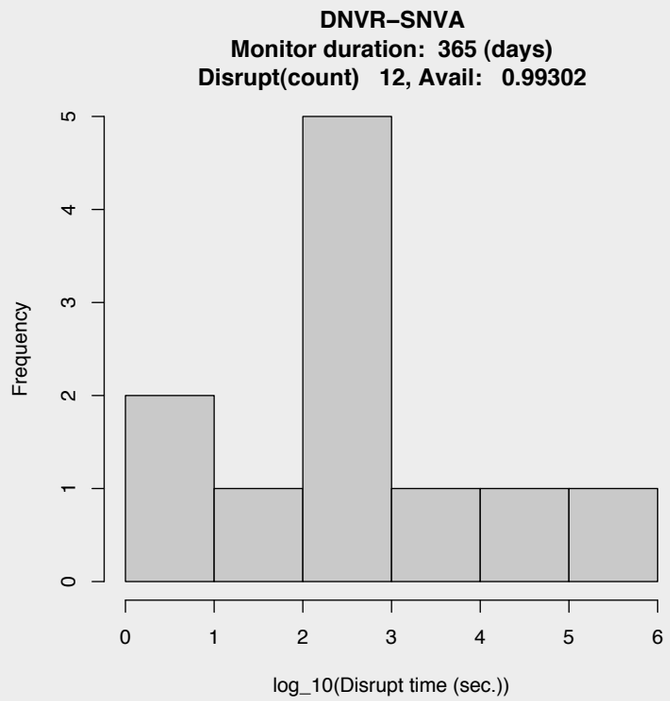

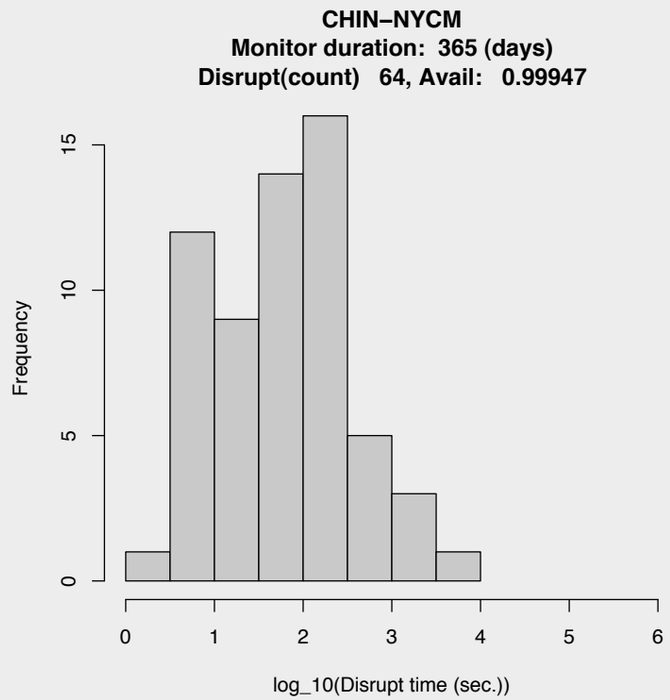

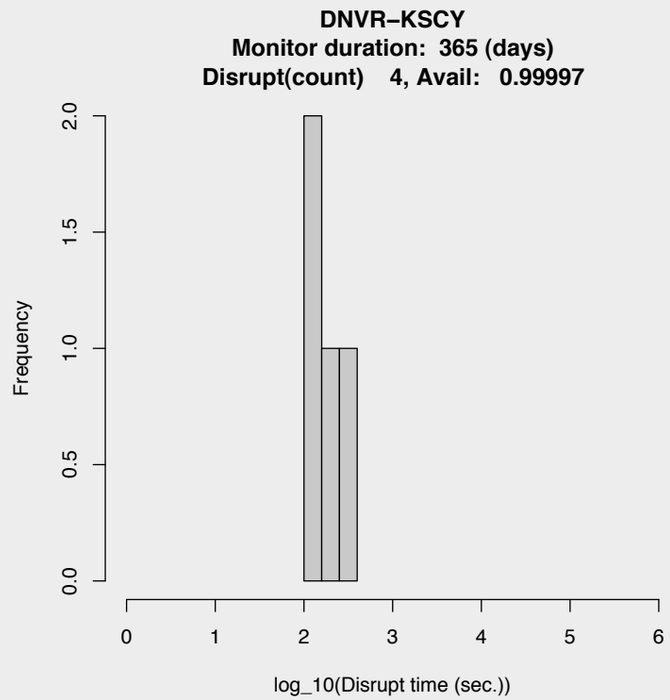

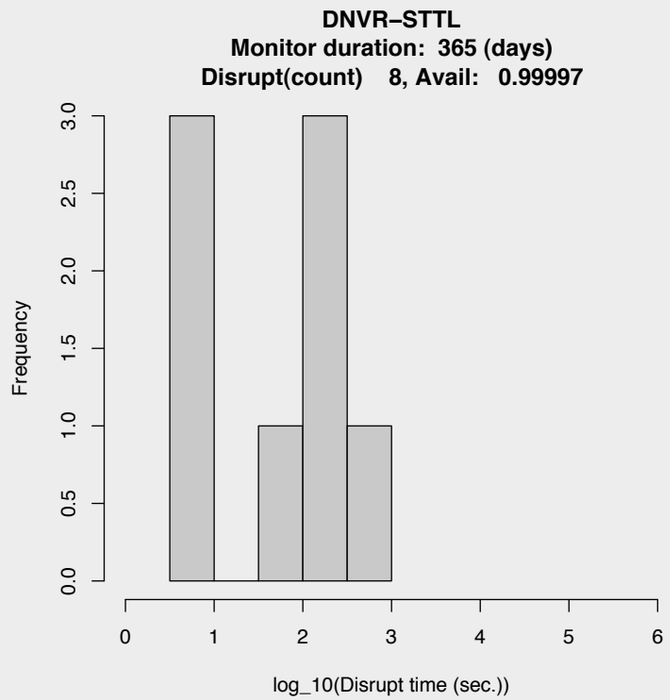

Availability Map (05/01-12)

Availability / Disrupt count / Longest down time (sec.)

Yearly summary ’05 - ’06

	2005/Jan.- Dec.			2006/Jan.- Dec.		
	Avail.	Disrupt cnt.		Avail.	Disrupt cnt.	
ATLA-HSTN	0.9738	24		0.9990	39	
ATLA-IPLS	0.9914	288		0.9975	48	
ATLA-WASH	0.9998	12		0.9994	25	
CHIN-IPLS	0.9998	34		0.9998	14	
CHIN-NYCM	0.9995	64		0.9999	30	
DNVR-KSCY	1.0000	4		0.9999	18	
DNVR-SNVA	0.9930	12		0.9922	51	
DNVR-STTL	1.0000	8		0.9999	5	
HSTN-KSCY	0.9993	18		0.9990	19	
HSTN-LOSA	0.9991	121		0.9996	40	
IPLS-KSCY	0.9998	10		0.9998	17	
LOSA-SNVA	0.9997	54		0.9993	128	
NYCM-WASH	0.9993	17		0.9989	113	
SNVA-STTL	1.0000	11		1.0000	129	
Total(*)	0.9544	677		0.9842	676	

Critical events.

- 2 or more lost adjacency at same timeframe
  - Some combination makes serious impact. But, not all event lead split graph condition.
- 32 timeframes (47 disrupt) in ’05, 58 (61) in ’06
- 26/47 timeframes in ’05, 49/61 in ’06, are attributed as missing a node in LSP database.

2 or more links failure (2)
- Missing node -

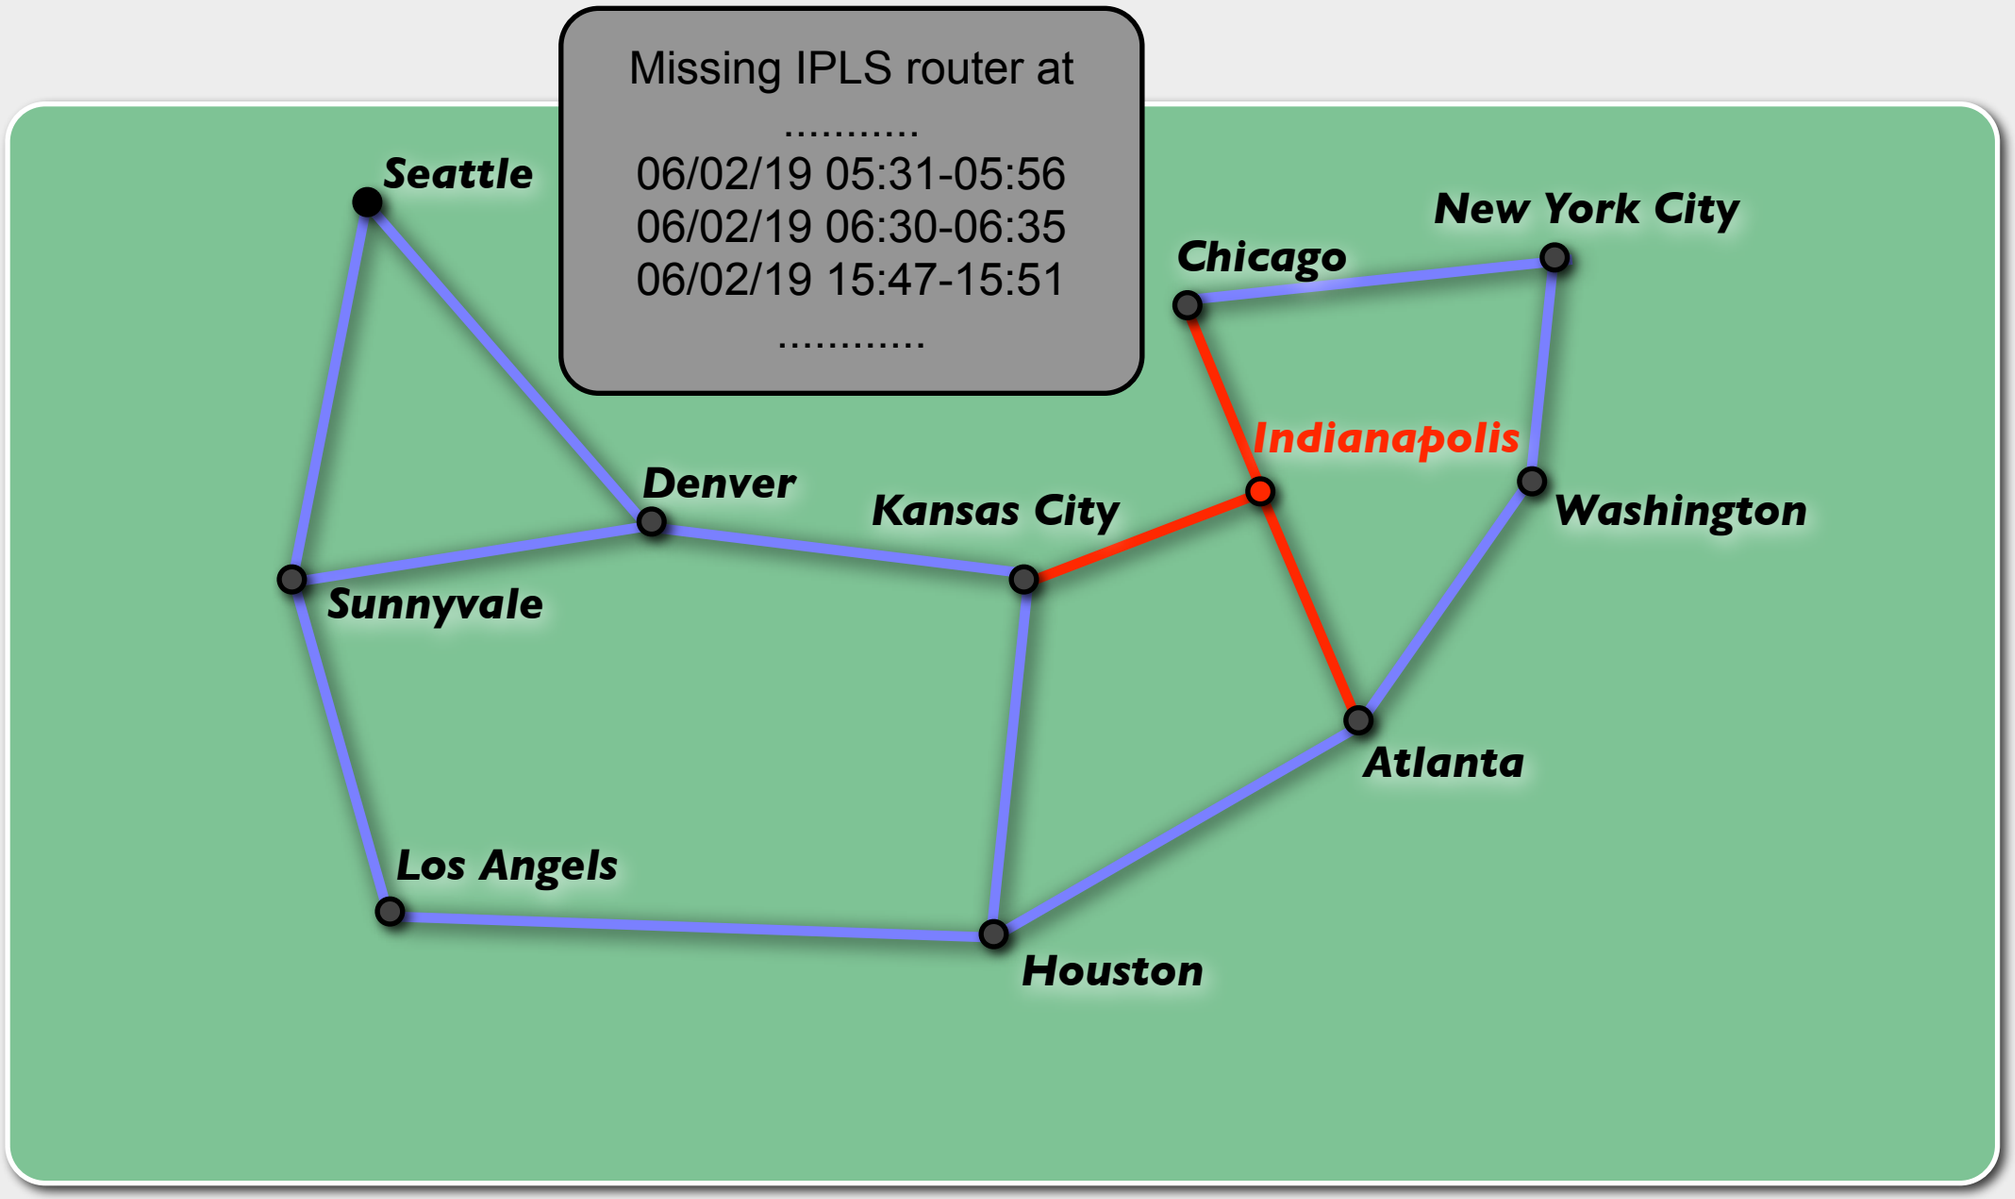

Two or more failure in ‘05

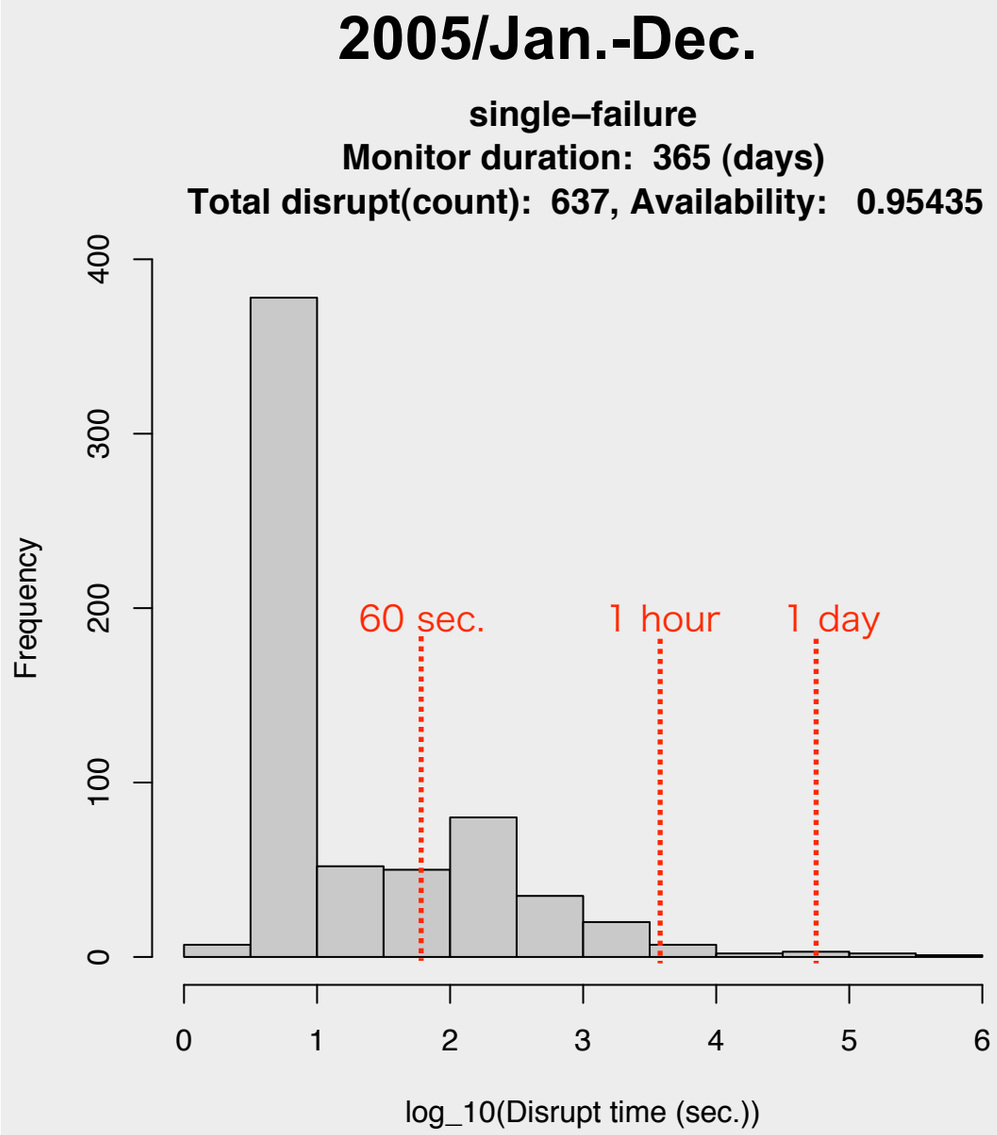

All lost adjacency events

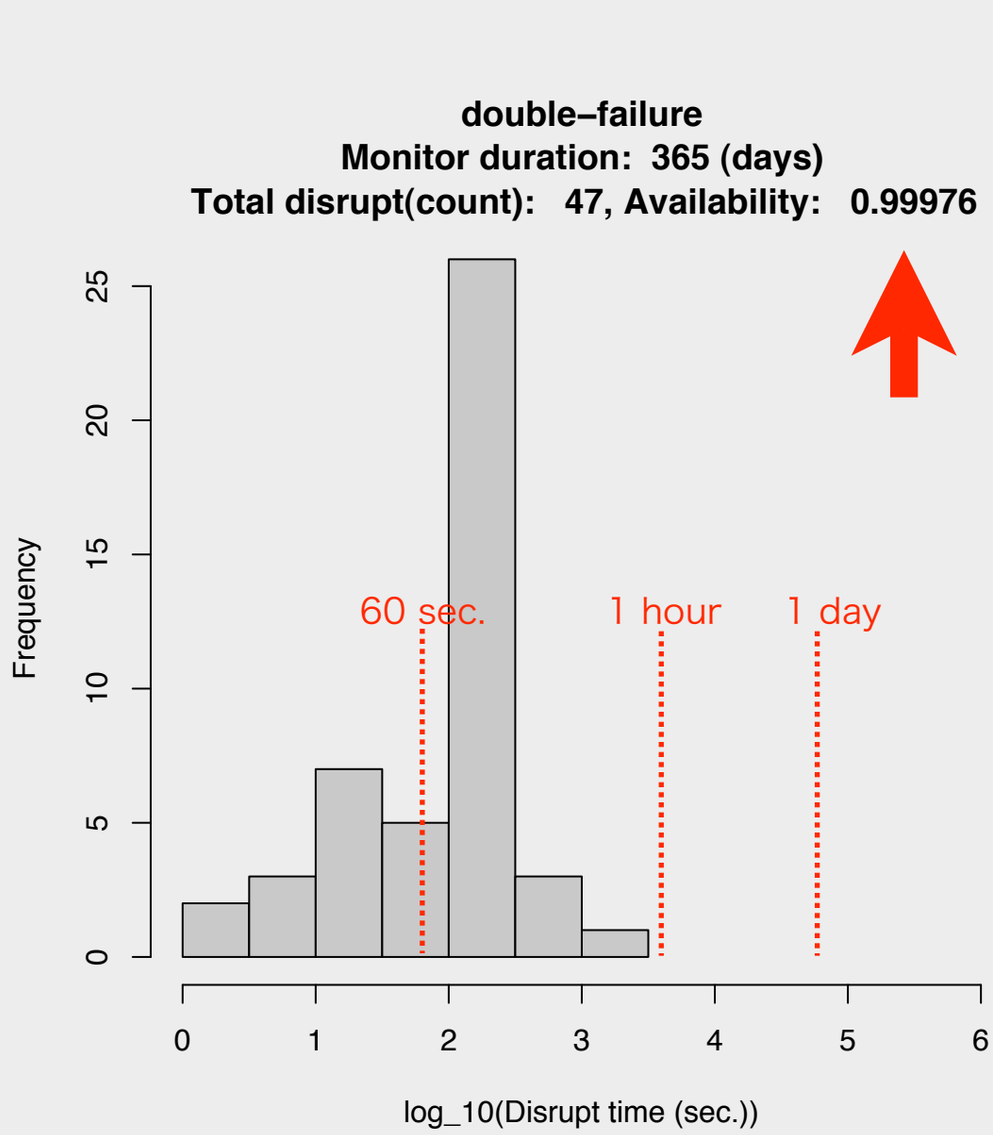

Two or more missing

Two or more failure in ‘06

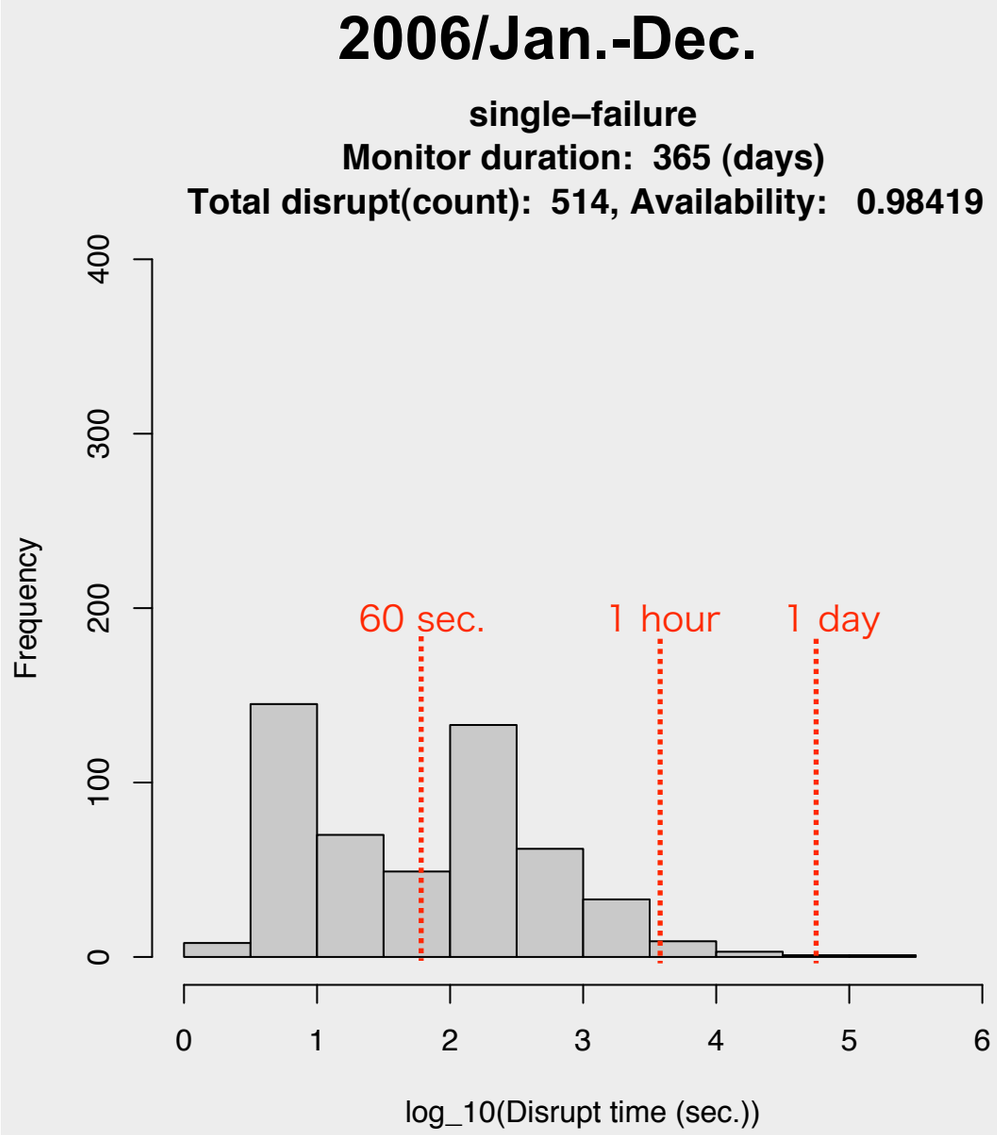

All lost adjacency events

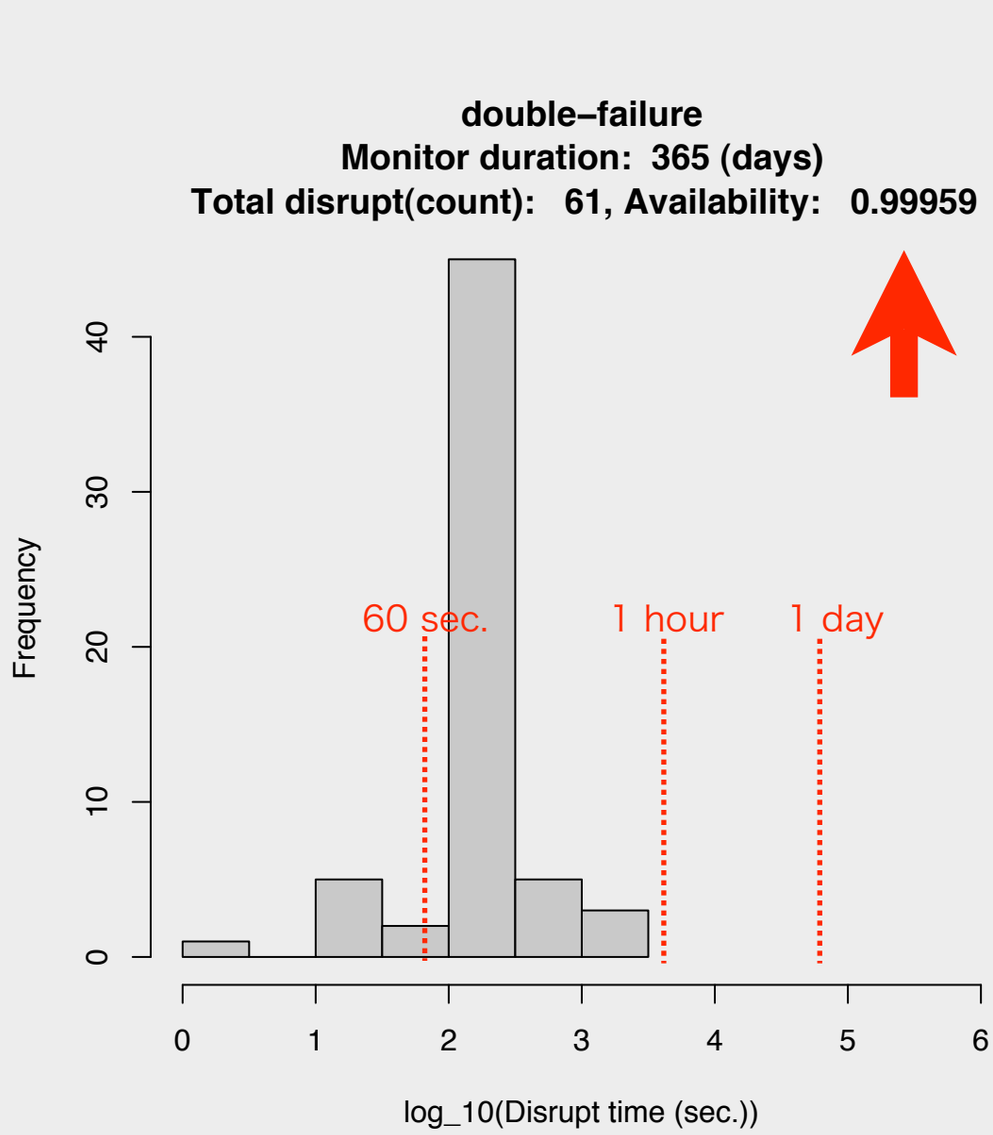

Two or more missing

Single link failure is trivial ? (1)

- Lost two or more adjacency events are rare, more than 99.95% availability, < 5 hours/year downtime.
- More than 500 lost single adjacency are founded.
  - 637 times in ’05, and 514 in ’06
- 3-4 hours/year downtime are estimated:
  - Only suppose 22 sec. downtime for each lost adjacency.
  - Other delays, i.e., routing convergence, degrade it more.

Single link failure is trivial ? (2)

- 22 sec. downtime for each lost adjacency is overestimated ?
  - Router can detect circuit failure more faster triggered with lower layer information, e.g., loss of optical, framer error.
  - IGP timer hack or BFD provide faster failure detection as sub-second or less [AC02].
    - Sub-second is derived from propagation delay limit, impossible to reduce it.
  - IP FRR would help more.

Conclusion

- ’05-’06 Full-year availability evaluation using Abilene ISIS trace data:
  - > 99.95 % backbone network viewpoint from IGP.
    - Better than real one.
      - routing convergence delay / access link
    - Abilene backbone is over-provisioned bandwidth.
    - It is not a news network worked fine :-)

- Thanks for Shu Zhang, Randy Bush, and Xing Li