DNSSEC Disaster Recovery

Edward Lewis
ed.lewis@neustar.biz
DNSSEC Workshop @ ICANN 46
April 10, 2013

In 5 minutes or less

- I'll assume we have a plan for disaster recovery for mundane items like loss of a data center, machine, network, power, staff, operating funds and so on
- What is at risk that is unique to DNSSEC?
- It's the secrecy of the private key!

How can a secret be exposed?

- It doesn’t matter.

- Maybe an employee walks.
- Maybe someone just guesses it.
- It doesn’t matter, you have to deal with it.

- A poorly managed registry can have it’s secret exposed
- A well run registry can have it’s secret exposed

- Don’t be concerned with “why did this happen”

What happens next?

- Some one can forge data and then poison caches "Kaminsky style"
- The global percentage of caches poisoned will be low
  - Could be damaging, but not "globally"
- Eventually you (the "owner") find out
  - Mean time to discovery…can be lowered with monitoring
- What's the worst thing you can do?

Panic - don’t!

- The only workable solution is to perform an unscheduled but otherwise normal key roll
- Sudden disruption will harm the majority of caches that weren’t already hit
  - You can then force as many poisoned caches to reload, but think of the consequences
  - Not all data in all caches are poisoned
  - You don’t know all of the caches
- All you can do is an unscheduled key roll!

MTTR

- Mean time to recovery is what you can manage
- How quickly can you change a key?
  - Lowering TTLs quickens changes, but raises normal query rates
  - Once you have a problem, it's too late to change TTLs
  - Shortening signature durations limits the time a key can be abused, but raises the amount of signing you must do
  - Once you have a problem, it's too late to change durations
- At "design time" you have to decide these numbers
  - Balance "sunny day" costs with "rainy day" risks