CoCCA - Recommend Practices | Continuity and Infrastructure

Overview:

In order to mitigate various risks to a TLD registry, the SRS & database should be:

1. Backed up several times a day using full “snapshot” backups. The CoCCA backups use the postgres pg_dump tool, pg_dump backs up the database structure *and* data in a single SQL file (these backups can also be made using the postgres command line).

2. Configured to use streaming replication to replicate all transactions in real-time to one or more servers / virtual machines for zero-loss failover. CoCCA recommends using postgres replication:

   - it does not require users to buy any specific virtualization tools and;
   - It is vendor neutral (it will work for all users regardless of OS or virtualization environment).

3. Where possible, the servers should have both public and private IP addresses, and the replication should be configured to use the private IP ranges.

4. Backup to an offline solution (external drive / USB) for disaster recovery. The CoCCA backups include a dump of the postgres database in binary format, the matching ROOT.war, CA jks and cocca licence key. The backup tar.gz file is all that is required to restore or clone the registry.

   There should also be a daily backup to infrastructure controlled by a trusted third party (or cloud provider) for disaster recovery. CoCCA has tools to easily encrypt and make ICANN and/or CoCCA formatted escrow deposits to third parties (including ICANN approved escrow agents).

Specific database backup advice:

Database should be backed up offline and replicated up to a minimum three instances under control of the registry operator. Two at the primary site (on different physical appliances) and one at a second physical site.

DNS Recommendations:

The master DNS server should be a non-public “hidden master” that is used only to distribute the zone files to the public facing DNS servers. Distribution should be secured using TSIG and IP allow lists. Only trusted IP should be notified / able to pull the zones.

Multiple DNS anycast clouds run by different vendors should be used to complement one or more unicast DNS servers under the direct control of the TLD manager. It is recommended that at least one public facing DNS server in the ROOT be under the direct control of the TLD manager.

To simplify the DNS configuration and DNSSEC signing, all registry zones that are delegated to the TLD servers in the ROOT should be compressed into the top-level zone . This makes it possible to add or remove zones without co-ordination / re-configuration with DNS providers, it also means that for DNSSEC, only one zone needs to be signed.

Dual Signers, the zone should be dual singed for redundancy. TLD manager should sign in-line (CoCCA recommends https://www.knot-dns.cz/) and by a trusted third party, preferably using HSM infrastructure. PCH as well as other provide a signing service to TLD managers.

What TTL to use ? There is some debate re this issue, CoCCA recommends following the advice in the following RIPE document.

https://labs.ripe.net/author/giovane_moura/how-to-choose-dns-ttl-values

CoCCA - Recommend Practices | Continuity and Infrastructure

Public facing WHOIS and RDAP services:

Where possible and practical, WHOIS and RDAP services (which require only read access to the database), should be run against a read-only replicated copy of the database. Using a read-only postgres user, connected to a read only db minimises the risk of unauthorised modification of the registry database from public facing services.

Keep it simple:

Continuous availability is critical on the public facing DNS services, this can be achieved by having multiple TLD DNS servers some operated by the TLD manager and some by third parties.

High availability or any kind of automated failover of the registry systems is not critical and likely to cause more problems than it solves. A temporary outage of the registry does not impact resolution of already registered domains, it simply prevents create and update transactions.

The decision initiates a failover - make one of the replicated registry instances the primary, will require a review of the issue that caused the outage and manual intervention. Once a decision is taken to failover from one instance site to another, the process only takes a few mins.

The process:

- Stop Resin on the failover site.
- Promote a read-only postgres instance that contains a replicated copy of the database to read-write (stop postgres, start postgres)
- Change the IP address A / AAAA records of the registry and rdqp whois portal — or if the issue is appliance failure at the primary site, and you have a second appliance there with replicated data, simple move the IPs to the second appliance.
- Start Resin

CoCCA SRS Security, key points:

- Credentials, tokens and passwords are hashed and salted.
- Only trusted entitles have http or API access.
- 2FA, IP restrictions Browser Fingerprints and Tokens supported.
- Public facing API’s and clients run off read-only db replicas.
- API access (EPP) is secured by IP address, TLD managed client keys, and passwords.
- TLD managers doe not manage DNS for third parties. Resellers and domain registrants make their own DNS arrangements. The registries only responsibility is to ensure the domain delegation information requested by the reseller / registrant is in the zone loaded on the TLD DNS servers (anycast clouds and unicast instances).
- Monitoring using PNGUOT using NAGIOS and cloud based Pingdom for PNGUOT hosts and appliances. Third party Anycast DNS is monitored by the providers, TLD managers do not have the access required to do detailed monitoring of the anycast clouds – or fix any issues that arise.

CoCCA - Recommend Practices | Continuity and Infrastructure

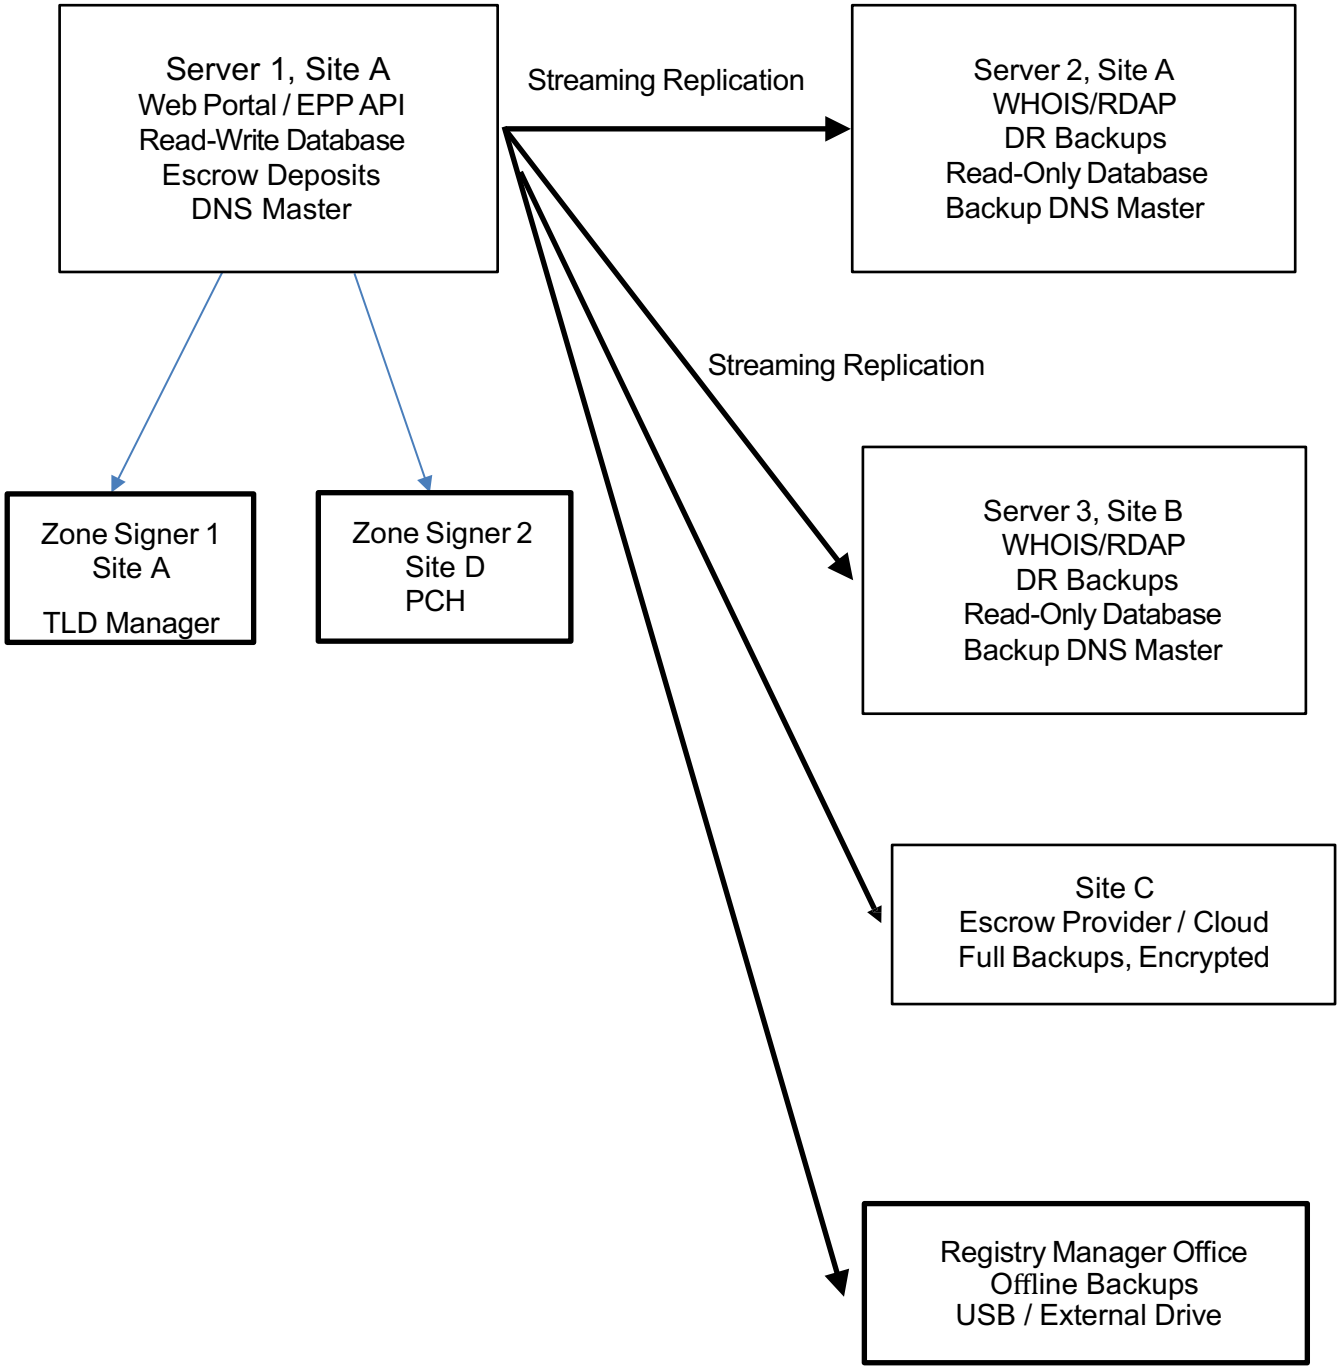

CoCCA - Recommend Practices | Continuity and Infrastructure

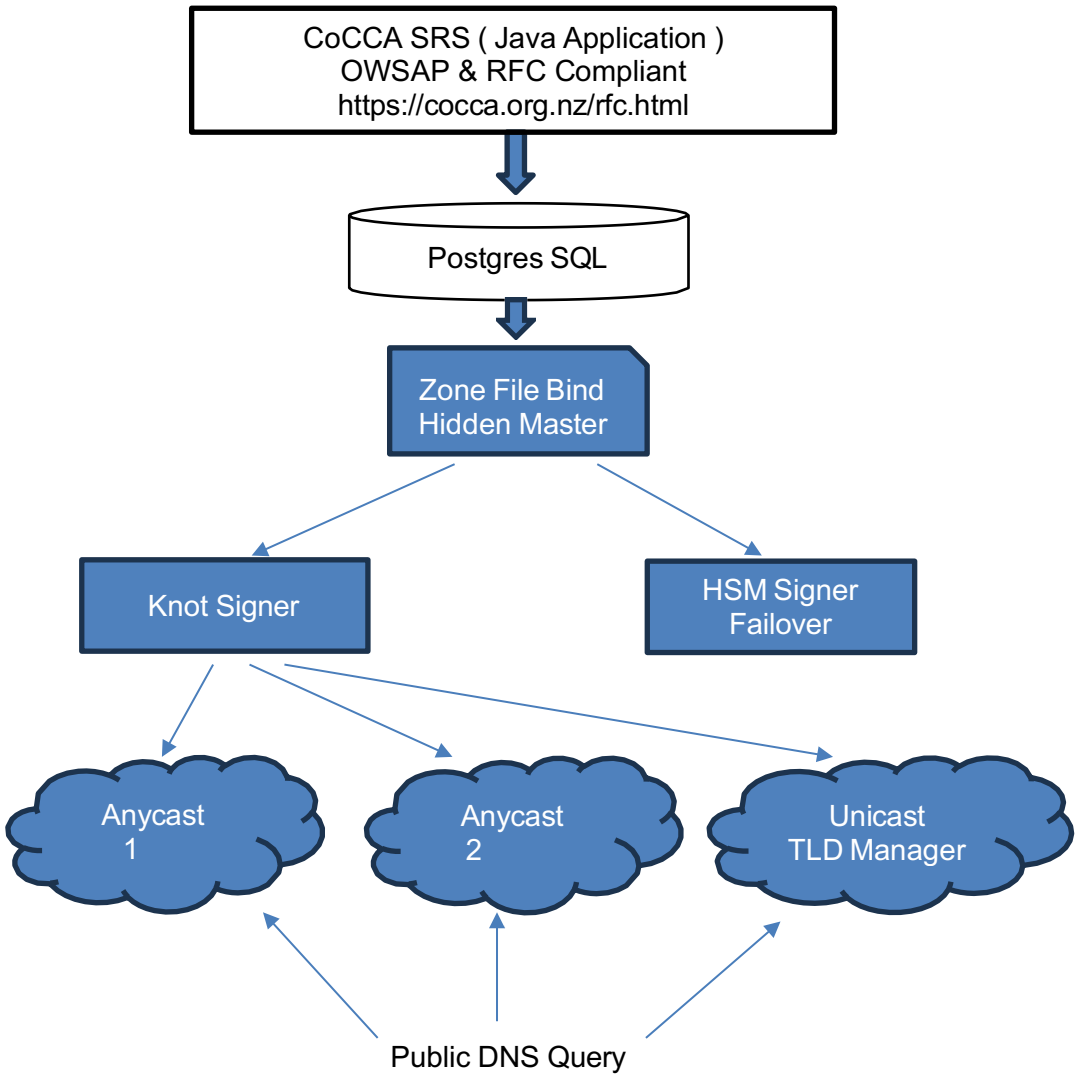

Minimum Recommendations, two anycast one local unicast.

PCH https://www.pch.net/services/anycast - 180 nodes

GRANSY https://anycast.systems/tld-anycast/ - 90 Nodes

TLD Manager (at least 1 signer and 1 public unicast node)

PCH Failover Signer