



Constructing Large Scale Distributed Registries

David Ortiz



Motivation

- ① Increase of EHR systems in clinical / research setting
- ② Opportunity to study rare diseases across enough patients



What is Carranet?

- 60 participating pediatric rheumatology centers collecting data
 - Data is stored across 60 virtual machines
 - All site data is stored i2b2 data store hosted
- Federate queries using SHRINE and SPIN
 - SPIN federates queries
 - SHRINE i2b2 interface to spin
 - Harvard SHRINE installation
 - ◆ 5 Nodes and fully messed



Challenges

- Linking up nodes to share data is actually hard
 - How do we make it scalable?
 - How do we make it secure?
 - What kind of topology?
- Fault Tolerance as number of nodes increases
 - Raid Example: 60 drives, likelihood that one fails is HIGH
- Cost of each node link
 - Mutual Identity Trust (cert exchange)
 - Firewall Rule (exhausting)
 - Service Endpoints (routing rules)




Network Topology Approaches

- P2P “fully meshed”
- Hub Spoke
- Hub Spoke, with fault tolerance
- Multiple Overlay Networks




Adding a node in P2P Model






Adding a node in P2P Model






Adding a node in P2P Model





Adding a node in P2P Model






P2P Approach

- 60 hospitals fully meshed = 1170 links
- Good:
 - single site failure, no problem!
 - Best choice for small networks
- Limits:
 - Limited in ability to scale to large numbers of nodes
 - High technical cost
- Cost
 - Mutual Identity Trust → 1170 cert exchanges
 - Firewall Rule → good luck
 - Service Endpoints → 60 copies of routes, DRY!




Adding a Node in Hub and Spoke





Adding a Node in Hub and Spoke






Hub Spoke Approach:

- 60 hospital hub = 60 links
- Good: Less links, single entry point, access control implemented at the hub
- Limits: If our hub fails, we're in trouble
- Cost
 - Mutual Identity Trust → 60 cert exchanges
 - Firewall Rules → each site opens access to the hub
 - Service Endpoints → one master copy, no need to distribute



Multiple Overlays





Hub Spoke with Fault Tolerance

- Fault Tolerance $T = 5$
- 60 hospital hub = $60 * 5 = 300$ links
- Good: relatively simple with few links
- Limits: Cost of running 5 additional “super nodes”
- Cost
 - Mutual Identity Trust → 300
 - Firewall Rules → 300
 - Service Endpoints → 5 master copies



What's Next? Self Scaling Network

- 60 hospital network = 60+ links
- Improved trust exchange
 - We need to use certificate authorities(CA) instead of self signed certificates to establish trust.
 - ◆ What happens if a private key is lost?
 - Deploying a CA doesn't have an easy button
 - CA requires delegating trust.
- Multiple hubs
 - Institutions can collaborate on “study hubs”
 - Sort of a “roll you own Carranet”
- Gossip protocols?



Thank You

- Justin Quan
- Bill Simons
- Andy McMurry
- Clint Gilbert
- Marc Natter