IRIS - An Overview for the ICANN Whois Task Force

Andrew Newton, VeriSign Labs
Leslie Daigle, VeriSign Labs
April 29, 2004

Level Set - The term “whois”

- It’s meaning to a technical person:
  - The Nicname/Whois protocol as defined by RFC 954 operating on port 43.
  - This protocol is used by three types of registries.
- It’s meaning to a policy person:
  - The meta-data about a domain registration.
  - Domain registration data is delivered to most end users via a web page.
- Both are right in their respective contexts.
  - The web servers get their data from port 43.

Port 43 - Nicname/Whois

- Nicname/Whois was first described in RFC 812 in 1982.
  - RFC 812 describes Whois over NCP, not IP.
  - It predates the modern Internet.
- By comparison, the first RFC to describe DNS was published in 1983.
  - It was never intended to describe DNS or distributed repositories of information.
  - RFC 954, the most current specification for Whois, spends more text describing who from ARPANET and MILNET should be in the database than describing the protocol itself.

Yesterday's Protocol, Today's Problem

- Because Nicname/Whois is doing a job for which it was not designed, it does not serve all of our needs.
- Lacking in authentication.
  - It can do weak authentication via source IP address, but this has problems with dynamic IP, NATs, network reassignments, mobile users, etc...
- No structure.
  - So no I18N/L10N, rules for query distribution, navigation, entity distinction, etc...

Moves To Deprecate RFC 954

- There have been requests from the technical community to move the Nicname/Whois protocol to “Historical” status.
  - it is not historic because it is still being used
- A 954bis document is currently before the IESG.
  - removes the outdated cruft
  - basically says, stuff goes in, stuff comes out
- Nicname SRV to Informational
  - With standards work on SRV/NAPTR to be focused on IRIS cohabitation.

Enter IRIS

- Text (XML) based protocol designed to allow registries of Internet resources to express query and result types specific to their needs while providing a framework for authentication, structured data, entity references and search continuations
- Encompasses the following
  - a decentralized system using DNS hierarchies where possible for location
  - multiple authentication mechanisms
  - built upon standard Internet building blocks
  - does not impose any informational trees or matrices
  - may be used with multiple application transports
  - rules for query distribution
  - has structure for Internationalization and Localization
  - etc...

Policy Neutral

- IRIS is policy neutral.
  - Access can be anonymous and/or authenticated.
  - Data can be given to some users and/or not others.
  - Trust can be based locally, regionally, globally, or all of the above.
  - Information can be centralized, distributed, or centrally indexed but distributed or all of the above.
- Since policy is not in the protocol, it can be differ between servers or sets of servers.
- Policy makers now have more tools.

Navigation of Servers and Data

- Navigation of DNS to help find an authoritative server.
- Query Distribution with entity references and search continuations.
- Relay bags to enable common index servers, trusted authentication, etc...
- Structured queries and results give clients the knowledge to display relationships.

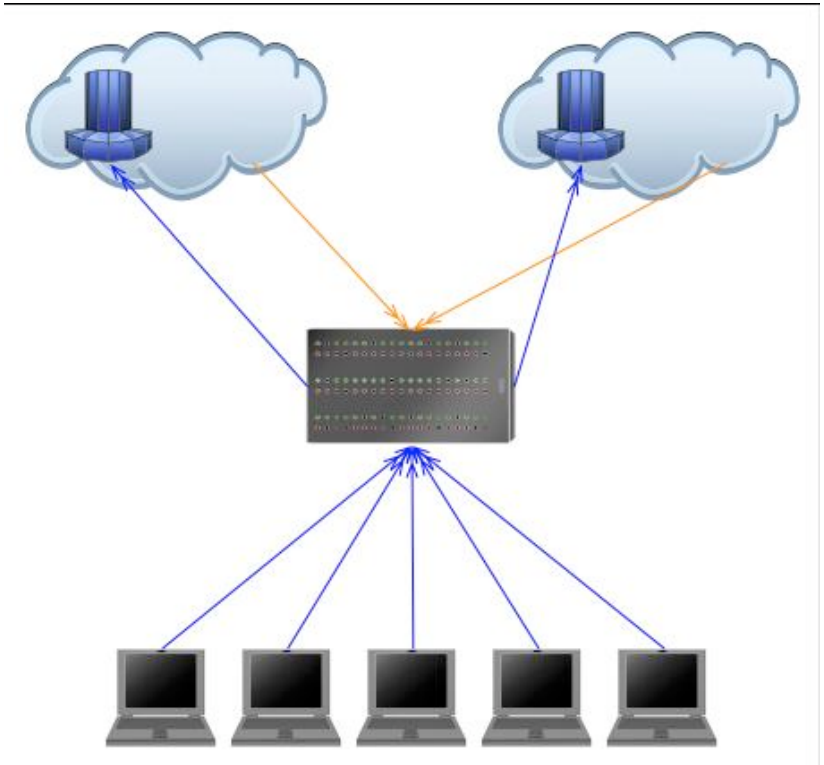

Tiered Access in IRIS

- Designed for distributed data repository architecture, with defined methods for finding the right server
- Ability to control who gets the info
- Critical need for network administration and law enforcement, etc...

```
$iris kosters.net
  Kosters, Mark
  US

$iris -cert fbi.cert kosters.net
  Kosters, Mark
  13121 Fox Shadow Lane
  Clifton, VA 20124  US
  703-948-3362
```

Authentication vs. Authorization

- Authentication – the process used to verify the identity of a user
- Authorization – the access policies applied to a user based on authentication
- Authentication mechanisms facilitate authorization schemes.

Modern Authentication and Authorization

- Authentication mechanisms
  - passwords, one-time passwords, digital certificates, references
- Authorization schemes
  - user-based, sequence-based, chain-based, attribute-based, time-based, referee-based

Distribution of Authentication Lists

- One of the challenges with tiered access is giving the right users access to the right information without overburdening the servers with the constant need to sync user lists.
- Digital certificates can off-load this burden.
  - Chains of trust.
    - A sender doesn’t trust the user, but does trust the entity that issued the certificate to the user.
  - User-based attributes.
    - A sender doesn’t trust the user, but trusts that a user of a certain type based on data in the certificate.

Referrals

- The IRIS protocol allows a server to pass extra information via a client to a referent server.
- This information may contain authentication data.
- The information could even contain the authorization policy.

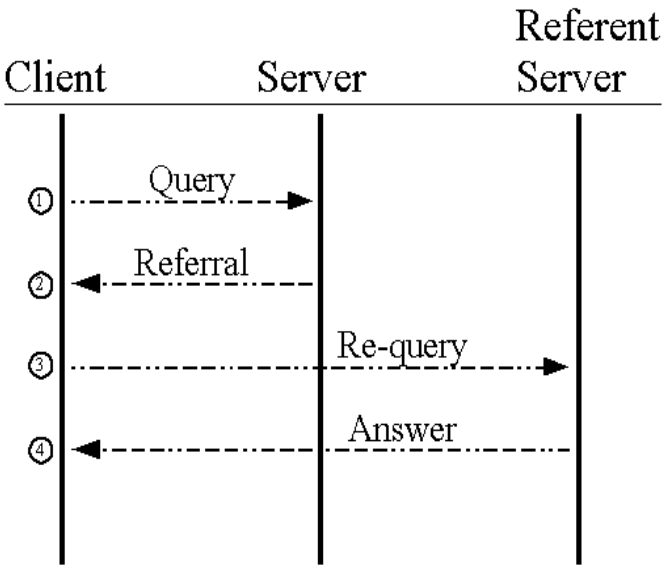

Structure

- XML gives data structure.
- Queries are well understood because they are structured.
- The information has distinguished entities and normalization.
  - TLDs can have differing models without confusing clients.
- For Internationalization:
  - datatypes are given well known tags for localization by the clients
  - data with multiple locales are given language tags

Structure Does More

- Structure enables the power user within the scope of policy.
  - They are no longer reliant on your web pages.

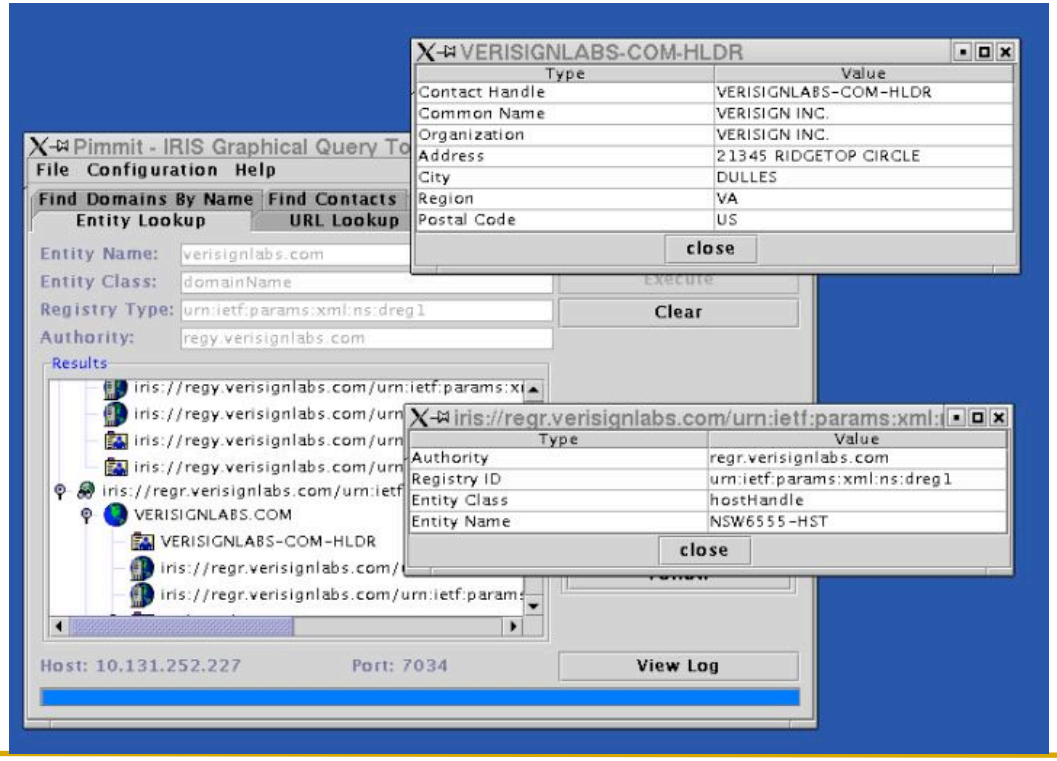

Extensibility & Layers

- IRIS is a layered protocol
  - Clear lines of responsibility in each layer.
  - Makes re-use of components simple.
  - Enables future extensibility.

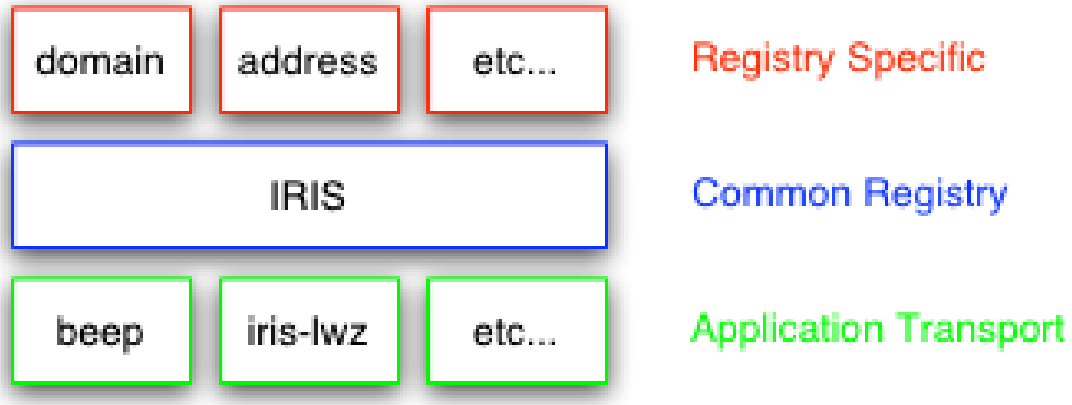

Responsibilities of the Layers

- Registry Specific
  - Defines queries, results, and entity classes of a specific type of registry. Each specific type of registry is identified by a URN.
  - AREG, DREG
- Common Registry
  - Defines base operations and semantics common to all registry types such as search sets, result sets, referrals, etc. It also defines the syntaxes for talking about specific registry types.
  - IRIS
- Application Transport
  - Defines the mechanisms for authentication, message passing, connection and session management, etc. It also defines the URI syntax specific to the application-transport mechanism.
  - BEEP

Future Need

- SASL - for future authentication mechanisms
- XML - for future data models
- S-NAPTR - for future transports

IRIS Status

- Prime focus of CRISP working group of the IETF
- A new specification for use by registries of Internet resources
  - Requirements are done
  - Protocol selection is done
  - Working Group last call is done
  - Approval by the IESG is next
- Opens source tools available
  - http://iris.verisignlabs.com/
    - multiple clients
    - multiple servers

Things To Come

- IRIS over UDP
  - IRIS-LWZ
  - Enables faster transactions and smaller server loads. UDP is one of the reasons DNS has a comparatively low overhead.
- Domain Availability Check
  - IRIS-DCHK
  - A scaled down version of DREG.
    - Implementations can use the same code as DREG.
    - Can be put on separate boxes with different SLAs.
    - When combined with IRIS-LWZ, very fast.