

Question 25: Extensible Provisioning Protocol (EPP)

Key Systems GmbH first gained experience working on the client side of the Extensible Provisioning Protocol (EPP) with the .Info registry (launched in 2001). This deep industry knowledge and experience has been transferred to KSregistry GmbH. Launched against early Internet Drafts, a great deal was learned in the first years leading to important protocol revisions. Some of these included changes to prevent the formation of orphaned glue records in an EPP registry and the ability to affect mass internal host updates with a single request. Most importantly, EPP brought the required functionality of populating registration contact data in the registry, allowing the subsequent implementation of a centralized or “thick” Whois service.

EPP was largely motivated by the growth in the number of accredited registrars that occurred beginning in 1999. Technologists working on EPP believed that the emergence of many new TLD registries was imminent and sought to ease the client-to-server implementation work that would flood the registry/registrar community if a standardized protocol was not developed through which to interact with domain registries. The effort was largely successful, although there has been extensive distinct diversion among overlapping EPP extensions between different registries over the years.

These many EPP extension differences first led KSregistry (KSR) to write a EPP proxy server in May 2009. The KSR RFC compliant EPP server implementation has been serving over 300 large and small registrars, passing up to 7000 EPP transactions per minute to over 280 registries. Some of these well known registries include: com, net, org, biz, info.

KSR EPP Server Interface

The EPP server is set up as a cluster to guarantee a high availability solution. The sizing of the servers and databases are calculated to meet the needs of each business case. The system is set up in a scalable way so that an increase of domains or registrars can be handled by adding hardware as needed. The KSR EPP API is offered over TCP on port 700 with mandatory SSL session enforcement for registrars for automated interaction with username and password.

To increase security, a registrar IP address limitation is in place for the EPP servers (both production and OT&E). This API also supports a secure web-based (over https) EPP client for registrars' manual use only. The web-based graphical interface interacts with the EPP server through standard EPP XML queries. The EPP XML responses are in turn displayed in the web interface. This allows the registrars to perform registry transactions through the web-based interface.

The KSR EPP Interface is capable of supporting up to 5000 read transactions per minute and 2000 write transactions (concurrently). It is provisioned in a highly redundant, duplicative environment using stateless, multiple

application instances. Please see Q. 31 for details on expected transaction volumes for the .bio gTLD registry.

A. RFC Relevance to KSR (KSR)

A.1 RFC 5730

This RFC is a base protocol document for EPP. EPP is an XML-text object based client-server protocol, atomic in its transactions, and developed to support multiple transports and lower level security protocols. There are no partial failures; all commands either succeed or fail definitively. Object-to-object associations are standard with limited application of parent-child relationships where delegate relationships are necessary for affected functionality, such as internal host data and its relationship to domain objects. The KSR registry fully implements the service discovery, commands, responses, and the extension framework described.

A.2 RFC 5731

This RFC explains the mapping of the primary EPP registry object, the domain object. It reviews associated attributes and states of the domain object as well as child object relationships (hosts). It also details associations with other contact objects. KSR complies with the full XML examples and descriptions and applies flexibility where permitted. For example, 5731 allows operators to implement the info command with different responses for a “sponsoring registrar” and a “non-sponsoring registrar” in regards to most domain object attributes. KSR implements this as a base protocol document for EPP.

A.3 RFC 5732

KSR implements this as a base protocol document for EPP. KSR notes this RFC describes the mapping of relationships to host objects, which are by definition subordinate to the superordinate domain name object. Host objects that are defined as internal or in the namespace of the registry must be related to a superordinate domain object to be created. Internal hosts, as full child objects, face restrictions associated with the management of their superordinate domain object. External hosts are hosts belonging to another domain namespace and as such are not subordinate in the present namespace. Internal hosts can have a glue or an A record associated with them, external hosts refer to another namespace or zone for the associated A record.

A.4 RFC 5733

Another base RFC implemented in the KSR server, this RFC describes the contact object mappings in EPP. Contact objects are used to contain related data surrounding the standardized contacts types in TLD registries including attributes such as contact type, country, telephone numbers,

email addresses, etc. As a standalone object, a contact object can be created and associated with no domain objects or with any number of domain objects available in the registry. This is used commonly by registrars to update common contact information associated across large numbers of domains in a single transaction. Like the domain object, it can be secured with a passphrase or "authinfo" code. Contact object data represents the definitive data source for authoritative RDDS (WHOIS) in new TLDs.

A.5 RFC 5734

KSR will implement this RFC as the preferred industry transport and in compliance with ICANN's requirements. Early implementations of EPP were considered over BEEP. This RFC describes a standard implementation of TCP incorporating TLS. As mentioned earlier, EPP can be implemented over multiple transports. The transport of choice for the EPP registry community has been TCP. Implementers are encouraged to take precautions against denial of service attacks through the use of standard technologies such as firewall and border router filters. IANA awarded port 700 as the dedicated port for the server side. There is no dedicated port assignment for the client side.

A.6 RFC 5735

KSR will implement this RFC as applicable to any extensions it utilizes as this RFC provides specific and detailed guidance on EPP extensions. An important principle in creating extensions to, as opposed to modifying, the EPP protocol was to fully preserve the integrity of the existing protocol schema. Additionally, a valid extension itself should be extensible. Another important requirement in the RFC is to include announcements of all available extensions in the EPP server greeting element before establishing an interactive client session.

A.7 RFC 3915

KSR will support this extension since the .bio gTLD implements the grace period implementation known as the Redemption Grace Period or "RGP". When RGP is in use, domains are deleted into the RGP where Registrars may request a restoration of the domain. This is a billable event and requires a three-step process: placement of the domain into a pending restore state, submission of a restore report explaining why the domain is being restored, and finally the restoration of the domain. The RFC extends the domain update command, adds related domain statuses, such as "redemptionPeriod" and "pendingRestore," and extends the responses of domain info and other details. The RFC provides a lifecycle description of the RGP and defines the format and content for client to server submission of the associated restore reports.

A.8 RFC 5910

KSR will support DNSSEC from the initiation of the .bio gTLD and therefore

will also support this extension from initiation of the registration process. DNSSEC is a mechanism for cryptographically verifying that each delegate zone in the DNS hierarchy has been referred to or is referring to its genuine parent or child zone respectively. Since TLD zone files are generated from authoritative registry data, this extension specifically provides the ability to add elements to the domain-create and domain-update functions and to the domain-info responses, allowing registrars to submit associated delegated signer information of the child zone indicating it is digitally signed and that the parent zone recognizes the indicated key as a valid zone key for the child zone.

B. Extensions used by KSR and Related Internet Drafts

B.1 Draft-tan-epp-launchphase-01 (Launch Phase Mapping for the EPP)

KSR intends to use this EPP internet draft to facilitate Sunrise phases during the initiation of the .bio gTLD registry. This internet draft proposes an extension mechanism that supports the organization of Sunrise related domain applications. The extension considers the following elements:

<lp:phase>

This element allows a Sunrise application submission to be marked by the EPP client as a particular Sunrise application type, in respect to running different types of Sunrise applications during a concurrent submission period. KSR will use this to identify Sunrise A and Sunrise B application types.

<lp:status>

This element allows the EPP server to assign one of a number of statuses indicating what stage the Sunrise application is in. These statuses can be expressed through the domaininfo command response and, optionally, through the RDDS service if applicable. The statuses listed below can be assigned uniquely or in combinations where appropriate:

<pvrc>

The Pre-Validation Result Code, an opaque string issued by a third-party validation agent

<claimIssuer>

contains the ID of a contact object (as described in RFC 5733 [RFC5733]) identifying the contact information of the authority which issued the right (for example, a trade mark office or company registration bureau)

<claimName>

identifying the text string in which the applicant is claiming a prior right

<claimNumber>

the registration number of the right (ie trademark number or company registration number)

<claimType>

indicates the type of claim being made (eg trademark, symbol, combined mark, company name)

<claimEntitlement>

indicates the applicant's entitlement to the claim (ie, owner or licensee)

<claimRegDate>

the date of registration of the claim

<claimExDate>

the date of expiration of the claim

<claimCountry>

indicates the country in which the claim is valid

<claimRegion>

indicates the name of a city, state, province or other geographic region in which the claim is valid. This may be a two-character code from [WIPO.ST3]

The complete draft is described in attachment Q25_Figure3.pdf.

B.2 Draft-obispo-epp-idn-00 (Internationalized Domain Name Mapping Extension for the EPP)

KSR intends to use this EPP internet draft to facilitate the usage of provisioning Internationalized Domain Names (IDNs). This internet draft extends the EPP domain name mapping to provide additional features that are required to implement registrations of domain names in character sets other than ASCII.

The extension considers the following element in both the domain create request and the domain-info response:

<idn:language>

This element allows for association of a domain name to a language tag, as defined in the code division of the Unicode code chart.

This element allows the registrar submitting the registration to identify a language tag associated with a punycode registration. The language tag refers the server to consider a declared set of table- and/or algorithmic-driven policies regarding a set and/or combination of defined unicode points, including variations of the punycode registration.

Insert Attachment with entire RFC including full xml schema examples

The complete draft is described in attachment Q25_Figure4.pdf.

C. KSR EPP Server

C.1 KSR EPP Command and Elements and Overview

Attachment Q25_Figure1.pdf contains the table with the supported EPP commands and the EPP object relationship.

Note: There are at least 2 name servers required for an active domain. Otherwise, the domain will be in the inactive status.

C.2 EPP Compliance Assurance:

KSR is committed to ensuring and maintaining compliance with the aforementioned EPP RFCs and, to this end, employs numerous mechanisms as listed below:

Quality Assurance Program

KSR runs a robust Quality Assurance (QA) program with multiple dedicated QA engineers. KSR has developed complete unit, regression, and stress based automated test suites for positive and negative use case testing. The test suites are self-developed and optimized for the relevant use cases. KSR reviews its use cases regularly with the entire development and registry operations teams for consideration as additional test cases. Additionally, KSR hosts periodic events where we bring together registrar engineers to discuss these use cases and seek new cornerstone cases that registrars may be able to offer from their experiences and points of view. KSR provides its QA team with a robust production grade testing environment with client load emulation capabilities that far exceed the load (through rate limiting) permitted on the KSR production environment.

OT&E

All new candidate EPP application versions will be released to a pre-candidate Registrar Operational and Testing Environment (OT&E) before promotion. Minor revisions, defined as new optional functionality, will have a minimum 30 day period in OT&E. Major changes, defined as requiring changes on the registrar client side, will have a minimum 90 day period in OT&E.

Inline XML Validator

The KSR EPP application uses the following XML validator in its server implementation. (Perl library XML::LibXML) XML errors or malformed XML will fail EPP transactions with the client atomically and the server will detail the failure state in the returned error message as well as the incorrect XML.

Third Party Validation

KSR is partnered with another Registry Service Provider (RSP), Internetwire, and has a bilateral agreement for each party to independently test and verify each other's EPP RFC compliance. KSR may also opt to engage other third parties for compliance testing.

D. Resources and Roles

D.1 Resources

Key-Systems GmbH has gathered experience in various roles in the domain business for more than ten years and has access to extensive knowledge in the domain industry. This deep industry knowledge and experience has been transferred to KSregistry GmbH, the technical provider of the KSregistry system (KSR), and is evident in many trusted persons serving in different roles throughout the company.

All employees, contractors, and consultants that have access to or control of the KSregistry system are trusted persons.

Each role is staffed with multiple human resources for backup and capacity purposes.

Prior to commencement of employment in a trusted role, KSregistry GmbH performs the following background checks on a prospective candidate:

- Criminal records bureau check
- Verification of previous employment
- Check of professional references

D.2 Roles

Designated Engineering Role

The designated engineering role includes the software developers of the entire SRS and all related interfaces (EPP, RDDS (Whois), escrow, etc.). All engineers are also integrated into 3rd level support of the SRS and related interfaces. The members of the engineering role are located in two geographically separate locations in Germany (St. Ingbert and Munich).

System Administration Role

The system administrators take care of the infrastructure of the SRS system. This includes the entire network, hardware, and system installations, as well as the cluster setup of the databases and all data backups which are made. Further, the installation of the hardware security module is performed by this role. This includes network setup, operating system installation, and HSM activation.

Support Role

The support role covers the first and second levels of support (service desk). All registrars and SRS customers may contact the support role as the first line of contact. Requests can be submitted via email or phone and can be placed 24 x 7. The support role is located in two geographically separate locations: one in Germany and the other in Mexico.

The first level of support receives all incoming requests from registrars and SRS customers and establishes the first contact. All problems that arise due to improper usage of the system or a misunderstanding of procedures will be resolved by the first level support. In addition, the first level points the customers to the online wiki and knowledge bases to prevent such requests in the future.

The second level support takes care of all problems which could not be solved by the first level. If problems are traced back to an erroneous system behavior as the cause, all available data are gathered and a problem report is generated and handed over to the change management team.

Quality Management Role

The Quality Management (QM) role takes care of each software component which is integrated into the SRS system and related interfaces. After a development cycle is finished, the QM performs full integration testing of the entire system. The testing is performed on a separate testing system (OT&E) which mirrors the production system. The QM ensures production readiness of each and every software upgrade, including emergency software patches. No software or system change will be promoted to production without the explicit approval of the QM.

Change Management (CM) / Project Management (PM)

The CM and PM role ensures that all steps in the development and system change processes are assessed, approved, implemented and reviewed in a controlled manner. This role filters requests so that only useful, valid and approved changes are implemented. They are also responsible for managing development efforts and changes to ensure that the changes are applied in accordance with predefined processes. They also chair the Change Advisory Board (CAB) and the Emergency Change Advisory Board (ECAB). These boards are comprised of selected people from other functions within the company. The project and change management role reviews and closes requests for change and reports to management.

The table in (fig. 2 in attachment Q25_Figure2.pdf) shows how the roles described above are planned for the SRS system. The calculations differ between the project phase and the years after the operational start. The project phase requires more resources as there is much planning, management, and development required. All human resources are only engaged in the domain industry and are experts in their area.

However, as the resources are shared and are not dedicated exclusively to one SRS project, the columns contain the number of resources available for this role and the percentage of all people working for the .bio gTLD. This percentage is the guaranteed time the resources for this SRS project assure.

E. List of Attachments

- Q25_Figure1.pdf
- Q25_Figure2.pdf
- Q25_Figure3.pdf
- Q25_Figure4.pdf