TMCH Solution Feedback

1 Summary

ARI Registry Services (ARI) and Neustar oppose the current proposed TMCH model as it currently stands—being overly complicated, unnecessarily costly, and seemingly ignoring key reporting requirements, which are required to assess effectiveness of the initiative. The following response identifies the current proposed model's critical issues, and provides an alternate draft model. In constructing this response we have collected input from policy, account management and technical staff alike, thus comments address issues with the current proposed model from many viewpoints.

2 Issues with the current proposed solution

2.1 Current Proposed Sunrise Model

2.1.1 Participation in Sunrise

The current proposed sunrise model requires that the registry's half of sunrise codes be obtained prior to commencement of the sunrise period. Rights holders that have not completed validation in the TMCH at the time of requesting these marks will seemingly be unable to participate in the sunrise, irrespective of the mark's actual registration date. It seems there is no consideration for a process on how someone who is either not 'approved' by the clearinghouse yet, or who hasn't even applied yet, can participate in the Sunrise process. It seems that those that miss this window may be required to compete with potential registrants who do not have a mark during a land-rush phase, which could result in an auction, thus contradicting the goal of providing a cost-effective solution for rights holders. Depending on the amount of 'advertising' that ICANN/TMCH performs, most mark holders will probably discover the TLD sunrise long before they do the clearinghouse; it seems highly likely that many people will first attempt to participate in the Sunrise and then be told that they need to enter their marks into the clearinghouse. In previous sunrise scenarios registries accepted sunrise registrations based on trademark registration dates not on some arbitrary time that one's mark was entered into a clearinghouse. Whilst some believe that this may be partially mitigated by giving mark holders several months to validate their mark in the TMCH, as mentioned above, ARI believe that, especially during the early stages of TLD launch, most mark holders first interaction will be with TLD registries (via registrars, or in most cases by smaller 'resellers' that specialize in online brand management), only then would it be directed to the clearinghouse—not the other way around.

The pre-generation of all sunrise codes also introduces a cost that does not consider registration volume. Single-registrant, single-user registries (e.g. for .brands), that may not even register any names during the compulsory sunrise

period) may have to pay fees similar to an open TLD that is expecting large registration volumes. Presumably, this can be mitigated by allowing each TLD to specify specific criteria, and only generating codes for marks that match that criteria. However in the case of an open TLD, the potential is for there to be millions of marks in the clearing house, but at most only 100,000 registrations may be received; it is a lot of work to generate and distribute codes for millions of marks for each registry when less than 10% of them are actually used. This is very inefficient, and inefficiency is unnecessary cost.

2.1.2 Access to Trademark Data

The current proposed sunrise model does not provide registries with access to the trademark data being used by the registration as the reason they are eligible to participate in the sunrise. This is problematic for a number of reasons:

Firstly, registries may restrict eligibility to rights holders that have a mark in a particular jurisdiction or class of goods (or indeed some other criteria applied to the mark). In order to validate such eligibility, under the current proposed sunrise model, a registry will require the rights holder to submit additional information about the mark they are using during the registration process. This additional information will then need to be validated by the registry, the cost of which will be passed on to the potential registrant. However this information is the same information that has already been supplied and validated by the clearing house, thus this increase in cost and effort is unnecessary. It has been suggested that this issue can be mitigated by requesting that the clearinghouse only provide sunrise codes for marks that meet a certain criteria. In this instance registries, without having means to verify would rely on the clearinghouse to have done this correctly. Registries would seek legal indemnifications from the clearinghouse to protect them against any issues arising from incorrect allocation of sunrise codes to non-eligible parties; something that is likely to only increase the TMCH costs and indemnification, thus preferable to avoid..

Secondly, registries may have an allocation method that preferences jurisdiction of mark or class of goods. An example is a geo-TLD giving preference to marks registered within its locality over marks from other jurisdictions when resolving contention between names. This differs from the above in that whilst a particular set of marks is eligible, smaller subsets of those marks are given higher preference than others. In the current proposed sunrise model, as above, additional data will also be required to be collected to meet this requirement and, as above, this will result in increasing costs for all, and the requirement of further indemnifications from the TMCH.

Finally, lack of trademark data reduces transparency of domains allocated through the sunrise process. Previous sunrises have included mark data in the WhoIs responses of domain names registered during sunrise processes. An example is observable by performing a query for microsoft.info, taking notice of the trademark information in the beginning of the response. This provides visibility for those external to the sunrise process, including other rights holders,

when making a decision to challenge the eligibility of a registrant to the given domain name. By not publishing this information in the registry WhoIs, it will be difficult for entities to decide if they have a case against a domain name registrant, thus most likely increasing the number of pointless cases being raised through UDRP and other circumstances.

2.1.3 Encryption / Cryptographic Obfuscation / Hashing

ARI believes that authenticating a sunrise 'code' of some sort is an appropriate model for sunrise validation; however the current proposed model for managing and validating sunrise codes is unnecessarily complex. Three reasons support this claim:
Firstly, with the code alone, one cannot determine to which registry a particular sunrise code is valid; therefore it is difficult for rights holders to manage their sunrise codes, because the potential number of codes they will have to manage is the number of TLDs multiplied by the number of registered marks. Rights holders may need to develop systems to manage and maintain their sunrise codes. Further it is also impossible, by looking at the code, to tell which code belongs to which mark.
Secondly, the obfuscation/encryption used makes it impossible for a registry operator and/or a registrar to provide any level of support to the rights holders. In the case of a rights holder supplying an incorrect code, neither registry nor registrar will be able to inform the rights holder whether:

- The code is not valid for the TLD
- The code is not for the provided name
- The code was copied incorrectly
- The registry has incorrect data for the provided name
- There was some other error

Neither registry nor registrar can determine the error or look at the code to determine the TLD and mark name. The rights holder has no choice but to receive support through the TMCH and even then they will have difficulty resolving issues as there are multiple parties involved. If the current proposed sunrise model is adopted in its current form, the TMC would be required to provide 24/7 support to registrars, registries and mark holders and stringent SLAs for resolution of issues, which increases costs.
Finally, the obfuscation prevents registries from calculating variants of the mark names. Because the mark name is not known, the registry must either:

- Check only the name provided for registration; or
- Check the name provided for registration and all variant names including blocked variants (which for some names can number in the thousands); or
- Check the name provided for registration and all names that will be automatically activated only (potentially not correct).

The first option is undesirable because some mark holders may (through restrictions in systems and input methods) submit a sunrise application for a variant of the mark name. Because the names would not match, the rights holder would be prevented from participating in the sunrise.
The second option eliminates difficulties that arise in the first option; however it introduces security risks to registries where malicious names (those resulting in the computation of millions of potential variant names) could be used to perform a denial of service attack against the Registry, as well as being impractical in some situations.
The third option strikes a balance between the two aforementioned options; however it still exposes the rights holder to potential difficulties during registration should one of the names not match the name provided to (or used by) the TMC in the calculation of the sunrise code. This third option does not consider scenarios where variant domain names could be activated post registration.

The second approach provides a solution least likely to cause confusion for rights holders and others involved in the process, however due to the risks it creates to the performance, stability and integrity of registry systems, it cannot be adopted. Most registries have already developed solutions that mitigate these issues from the general operation of a registry. These solutions require that both the provided name and the names in the lookup (comparison) data set are known. Technically, this involves computing a canonical form of each name for comparison, where a property of the canonical form is that it is unique for a name and all its variants.

2.2 Current Proposed Claims Model

As a registry operator and backend services provider, ARI has significant concerns about the claims data being located at the registry. These concerns are borne out of the responsibility of maintaining this data and the uncertainty of the data itself.

2.2.1 Replication of Data

The current proposed claims model involves the ENTIRE TMCH database being replicated to each registry (TLD). This has a number of issues:

Firstly, each TLD has its own distinct replicated copy of the database, using different cryptography keys, so registry providers, managing registry services for multiple (in some cases hundreds of TLDs), will have to maintain (download, store and update) hundreds of distinct copies of the same data for no clear reason. This is unnecessarily costly. The current proposed claims model requires registries to store claims data that grows linearly to the number of TLDs. This is further exacerbated given that expected data set sizes and update frequencies have not been approximated.
Next, registries are required to keep secure and contractually protect a large data set of which they only really need access to a small portion of. If there are millions of entries in the clearinghouse, however only 100,000 domains are registered during the claims period, the registry unnecessarily had access to, and

had to protect, more than 90% of the TMCH data. In a .brand scenario it is highly likely that 99.99% of the data is useless to the registry and will never be used. The current proposed claims model is fundamentally flawed, goes against basic security principles, and by burdening registries with unnecessarily taking on this responsibility, most registries will seek indemnification against the exposure of this data; whether by mining by a third-party, security breaches or any other method.
Registries are also concerned with the unknown nature of the problem space. Currently the expected volume of data (both initially and updates), frequency of updates, and other metrics have not been provided nor discussed, affecting the ability of a registry to plan for the deployment and operations of the critical launch phases of new TLDs.

2.2.2 The TMCH 'Outage' Problem

It is VERY important to note that the proposed model does not solve the 'TMCH' outage problem, it simply changes it form. Previous proposals have revolved around the TMCH maintaining a 'query' interface of some sort for the TMCH data. This solution has its issues when we consider the situation that arises when that query interface is no longer available due to an issue (a solution to which we propose below). The current proposed solution does not eliminate this situation, consider the following scenario: What if during the rsync process with the TMCH, the TMCH has a malfunction of some sort that results in the incorrect data being rsynced to the registry? This is indeed extremely feasible, especially with the TMCH being asked to managed one, encrypted, obsurfacted, not easily inspected or verifiable model of the data for EACH REGISTRY. Also given the encryption, in some case the registry would have a hard time even knowing that the data that they have is incorrect (or gibberish). Depending on the encryption used, they will either not be able to decrypt the data or, decrypt it and just get garbage as the response.
Should this situation arise, the registry cannot have an unplanned outage and still needs to continue taking registrations? What are they to do? This is exactly the same issue as when an online query interface is not available, there still MUST be an acceptable action plan for registries when the TMCH data is not available and this CANNOT be 'stop taking registrations' (unless the TMCH wants to financially compensate registries for the outage time and ICANN will relax SLA commitments for the period the data is unavailable, however this is still extremely unfavorable).

2.2.3 Use of Encryption

The current proposed encryption model is pointless. The encryption only exists to stop registries reading the entire trademark clearinghouse database; it serves no other purpose. This assumes registries are non-trusted parties who are bad actors. Assuming a bad-actor registry, it is relatively trivial to gain access to most (if not all) of the database. Because the encryption key is made up of something the registry already has (the registry part), and something that is determined from each domain name; all that is needed to decrypt the data is a list of domain names. The zone file access service which provides access to the .com and .net

zone files, gives access to 100,000 million domain names to check against the database. Thus running the following 'pseudo code' would likely result in the decryption of most, if not all data.

for each name in source of domain names (.com & .net zone file)
determine if name is subject to trademark claims
if subject to trademark claims then decrypt claims data
done

Given that most marks will be registered in .com (whether by the mark holder or not) the probability of getting greater than 90% of the database with this method is high.
Additionally, the use of encryption does not align with the goals of a registry operator, whose interest is to provide the lowest round-trip time for queries due to contractual penalties applying if round-trip-times exceed a specific threshold. It should be expected that a pre-computation of the TMC data would be considered by registries, to remove the cost of decryption from the round-trip-time of claims lookups. Having access to the decrypted data also allows a registry operator to verify the integrity of claims data, and to have knowledge of the data that the registry will be transmitting (maximum length of claims notices, average length of claims notice)—required for testing of software and capacity planning. Such upfront analysis removes the possibility of run-time failures (such as failure to decrypt data) affecting domain name registration. In short, registries, especially back end registry companies, may have legitimate reasons for using the process described above to decrypt the information provided by the clearinghouse.
Further to this, the encryption does not prevent data mining by parties other than registries. Instead, distribution of claims data to multiple registries will actually increase the ability of calculated attempts to mine claims data. Malicious users can distribute their queries amongst a combination of registrars and registries, never providing enough information to any one party to determine mining. Should mining be detected, damage would be done before the application of remedies across each of the registries and/or registrars.

Encryption also prevents registries from calculating mark name variants. The issues are identical to those described in response to the sunrise process above, thus the recommendation is to remove the encryption of mark names.

ARI recommends that should data still be provisioned with a registry, that data be provided without encryption. Most registries do not have any interest in this data, however encrypting it assumes that registries are bad actors. However any registry that is a bad actor can easily get access to the data anyway, as described above. As was discussed on the May 14 call, encryption is not enough to restrict access from the registry and contractual agreements will need to be put in place. If there is reliance on contractual provisions, then encryption should be removed to allow registries to better meet their requirements.

2.2.4 Allocation

The proposed claims model appears to only consider allocation based on a first-come first-served model. Consider a land-rush process where potential registrants submit applications for domain names and resolve contention via auction. This would create a separation between the submission of the application, and subsequent allocation (if any) of the requested domain name. The concern is that the submission of notices to rights holders before allocation could be considered front-running if the rights holder subsequently filed an application for the said name. In response to this, applications might be submitted in the closing hours of the application window, negating the purpose of reducing the rush of applications in a short period. At a minimum, ARI expect clarification of the process for auction-based allocation methods, which should involve the removal or otherwise definition of the term 'registration'.

2.3 Handling of Variants

ARI is concerned that rights holders have the belief that they can rely on registry operators for the protection of their mark name and its variants. This assumption leads to issues for two reasons; firstly the variant mappings are not guaranteed to be consistent with the rulings of various trademark bodies and courts for equivalent characters; secondly IDN tables and their variant mappings can differ between registries for the same sets of characters.

The developers of IDN tables may have considered visual similarity, audio similarity, input methods of various devices, or any other measure when determining those characters that are to be treated as equivalent. While these tables can be considered to include some definition of similarity, this is not guaranteed to be equivalent to the definition of identical or confusingly similar as used by trademark bodies or courts.

Consider an IDN table allowing the use of an accented *e*, as used in the term café, without defining it as a variant of the non-accented *e*. A mark containing the accented *e* may be validated through the clearing house; however a domain name registration of a similar name, where the accent has been dropped from the *e*, will not trigger any notifications or alerts. This may contradict rules in trademark practice within several jurisdictions where the two names would be considered identical or confusingly similar. Note that this fictional scenario is without knowledge of trademark practice, however should suffice to illustrate the potential issues with the current model.

Registries are not mandated to follow any particular rules regarding the registration and management of variant domain names. In fact ICANN has not come to consensus on the issue of internationalised domain names and variant names and, given the wide variety and interpretation of issues, it is highly likely that there will never be a consistent application of variant rules. Registries can use tables allowing all Unicode code points that have the same 'script' property, or alternatively use tables that correspond to a set of characters commonly used in a particular region. Due to these differences there is no guarantee that variant concepts will be consistent among registries.

Furthermore, there is no requirement that registries manage variants. Registries may place the responsibility of identifying and registering variants onto the

registrar, who is arguably in the best position to identify variants based on sound, meaning or string similarity. The rules around marks of simplified and traditional Chinese trademarks are another example.
ARI believe that the best path forward is to be explicit regarding the mark names that are protected. The TMC, either through internal or external process, should identify whether the mark, for example for café, also covers the mark cafe. This behavior of explicitly identifying variants of mark names ensures consistent registration and notification practices amongst all TLDs.
To summarise, mark holders should submit their mark, and any names they consider equivalent to their mark, for the TMCH to verify and enter. Registries should only be required to perform matches to this list, as domain name variant issues are very different to trademark issues. Any disputes between mark holders about domain names being similar to their marks should be handled by UDRP, courts or similar processes which are better equipped (and more qualified) to deal with these issues.

3 Proposed Alternate Model

The following outlines, at a high level, a more appropriate implementation model, that meets all of the requirements of all parties and is significantly cheaper and easier to implement for all parties.

3.1 Proposed Alternate Model for Sunrise

The following proposed alternate model is offered for sunrise, improving the current process while meeting all requirements. This proposed alternate model attempts to simplify the process by decreasing the coupling between the TMC and registries:

1. The TMC generates and maintains a public-private key pair (this could be global, or per registry) and transmits the public key to the registry.
2. The TMC signs sunrise (trademark) data with its private key and distributes this digitally signed information to the mark holder.
3. The registrar accepts this digitally signed data (which you could consider the 'authcode') along with the request for registration.
4. The registry verifies the signature with the public key and verifies the mark matches the name being registered.
5. The registry notifies the TMC of the use of the digitally signed data (for reporting purposes).

In this proposed alternate model the TMC only makes available a public-key instead of providing each registry with a shared secret, and a database of sunrise codes (1). The registry still has the responsibility of validating mark registration data during domain name registration; however the decoupling of the two systems simplifies scenarios when the sunrise period is changed/extended or participants are not present in the clearinghouse at commencement of the sunrise. Also registrars can use the public key to pre-verify data before it is submitted to the registry, enabling them to notify the customer and correct the

issue on the spot. Registrars are in the best position to do this, because they ultimately own and manage the relationship with the customer (registrant).

Note: The TMC can decide whether keys are unique per registry or shared across several registries, based on perceived threat of loss or forgery. However having keys per registry puts undue burden on mark holders to keep track of which signed data bundle belongs to which registry, even if methods below allow registrars and registries to detect and assist with these issues.

The signed sunrise data (2) is similar to the authcode mechanism—this will be referred to as a 'sunrise code'. This signed data meets a pre-defined format (such as XML) similar to that described in the current model for claims purposes. The data format should prescribe:

- Serial or version for this code
- The registry(s) that will accept this mark
- The label(s) that the mark covers
- The mark name
- The registration date of the mark
- Other mark information – such as jurisdiction and class
- Owner information – such as organization name and address
- Signature

A simple example could look as follows (illustrative purposes only):

```
<sunrisecode serial='123456'>
        <registry>.example</registry>
        <label>mark-and-mark</label>
        <label>markandmark</label>
        <mark>
                <name>Mark & Mark</name>
                <registrationDate>2001-01-01</registrationDate>
                <class>XYZ</class>
        </mark>
        <owner>
                <contact-name>Mr contact</contact-name>
                <organization>Mark & Mark Company</organization>
                <address>1 some street, someplace</address>
        </owner>
        <signature>SHA256 HASH of  labels or something similar</signature>
</sunrisecode>
```

This information can be easily understood my mark holders, especially those juggling many sunrise codes.
This format can be extended to include any other information such as TMC validation date, and any other information deemed relevant. This may provide benefit in steps 3 and 4 where eligibility, or allocation preference, depends on this additional mark data.

The mark holder would submit this sunrise code to a registrar for registration (3). Registrars may be able to streamline this process such that mark holders upload only the sunrise code and select the names (where there are variant labels) for registration. Registrars would be able to identify many issues early, such as malformed input, registration date following the cutoff date, and incorrect signatures should the corresponding public key be distributed to registrars. This has the potential to greatly simplify the registration process for mark holders for those registrars determined to put in the effort.
The registry is ultimately responsible for verification of the data once the registrar has issued the create request (4). This would involve validating that the signature is valid for the code, and could involve further checks such as that the registration date is not later than the cutoff or so forth. The registrar knows the data has been unmodified and has been independently verified by the trademark clearing house thus does not have to verify it themselves. The information available in the sunrise code may be used to supplement responses to WhoIs queries, providing visibility into the mark that was used to establish eligibility. As the entity that owns the mark, has voluntarily submitted the information to the registry, they have remained in full control of their own private data, they have elected to submit the data to the registry to enable them to register a domain name, and there are no privacy issues with this approach.
The registry could then optionally (to the model, not to the registry) notify the trademark clearinghouse (5) of the use of the sunrise code by serial number and registered labels. This information allows the TMC to identify the relevant mark holder and notify them (and anyone else) as appropriate, and for reporting purposes.
This model supports deferred requests for a sunrise code, meaning that codes only need to be generated for clients that require them. A benefit of this model is that the TMC can choose to pre-generate the sunrise codes, or alternatively generate a code on request. This solution lends itself well to a user pays charging model.
Mark holders will pay to enter their names into the clearinghouse, and need only pay to get 'sunrise codes' for the TLDs they wish to register in. If a single public key is used, then no additional payment should be required as there is no work for anyone to do.
Alternatively registries could pay the clearinghouse each time a sunrise code is used to register a name, however ARI propose that registries will simply build this cost into the name registration and that ultimately the registrant will be pay this bill anyway.

One simplified standardised EPP extension needs to be produced to allow this code to be transmitted to the registry; no complex check extensions are required. Alternatively, most registries have an existing key-value pair extension they can use to facilitate transfer of the code.

The registry would communicate the registered names and used sunrise 'tokens' via a simple file upload in a specific format. The registry could generate and upload the file to a secure FTP or some other site (could be over HTTP if the TMCH so desired) on a periodic frequency, create time stamps may also be

required. The frequency should be no more frequent than 4 times a day, however, realistically we see no reason why it couldn't be once daily.

This alternate process does not introduce new actors, nor does it impose additional requirements on the registries and other parties outside of the TMCH.

3.2 Proposed Alternate Claims Model

The current proposed claims model forces an implementation that is not only costly to implement and maintain, but does not take into consideration the business requirements of registries and registrars. The following response presents an alternate solution that reduces the number of parties involved and greatly simplifies the process (thus reducing cost and complexity).

The TMC should provide the content for the claims notice, eliminating the need for any registry to have access to claims data and simplifying reporting and abuse monitoring mechanisms by the centralisation of data access. This proposed alternate claims model is summarised as follows:

1. Registrant approaches registrar to register a domain name.
2. Registrar checks if the name is available using the normal EPP Check mechanism.
3. Registrar checks if the name matches a claimed mark by performing a DNS query against a defined set of DNS servers (see DNS list of marks below). If there is a match this process continues, otherwise registration completes as per normal and this process ends. On match the DNS returns a random unique string that the registry uses in the HTTP query below.
4. If a match is found the registrar then obtains the claims notice content from the clearinghouse using the HTTP protocol (see HTTP claims notice below). As part of getting this information the registrar provides the matching name and the TLD it is being registered in. The claims notice contains all relevant information plus a short code (8 chars maximum) generated from a hash of the matching domain name and a shared secret between the clearinghouse and the registry (identified by the TLD) or alternatively cryptographically generated using the Clearinghouse public/private ley pair. This code (as displayed to the registrant) could be the first 8 chars, it is not required that the codes are unique, only that they are related to the domain name and cryptographically verifiable by the registry.
5. The registrant reads the notice which instructions them to input the code into an input box (or provide it via some other method) to indicate acceptance of the notice.
6. The registrar then sends the create command to the registry including the entered code and the whole signature if public/private key pair was used. Again a very simple EPP extension is required for this, and most registries already have key-value pair extensions they can use for this.
7. The registry verifies that the name is subject to claim and that the entered code/signature is valid for the name and shared secret between the registry and clearinghouse. If valid the registration succeeds, if not it fails.

8. The registry periodically notifies the TMCH of all names registered and codes utilised (for optional verification by the TMCH). The TMCH can then report on and notify mark holders. This should be create time stamped and use the same processes as described in sunrise above.

This proposed alternate model, has the following features:

- Protects mark data against mining, as:
    - The URLs for claims notices are not predictable
    - The TMCH can see all requests for claims notices and can monitor, rate limit or whatever else it deems necessary to protect the data
    - Data is centralized only at the TMCH, so all access is known and controlled by them
- Allows registrars (the entities with the customer relationship) to provide feedback to potential registrants immediately, especially when checking names across multiple namespaces
- Is technically efficient and very simple to implement
- By having the registrar obtain a code from the TMCH, then making the registrant read the notice to 'find' the code, input it into an field on the registrars site, then send this code to the registry, then finally the registry sending it full circle back to the clearinghouse and enforcing each party to do these things contractually this provides a clearly auditable chain of events that if later can be used to prove that notices where generated, and viewed by the registrant, who had to read them to get the code.

If the 'circular' proof of notice described in the bullet point above is not required the process could be simplified even further as follows:

1. Registrant approaches registrar to register a domain name.
2. Registrar checks if the name is available using the normal EPP Check mechanism.
3. Registrar checks if the name matches a claimed mark by performing a DNS query against a defined set of DNS servers (see DNS list of marks below). If there is a match this process continues, otherwise registration completes as per normal and this process ends. On match the DNS returns a random unique string that the registry uses in the HTTP query below.
4. If a match is found the registrar then obtains the claims notice content from the clearinghouse using the HTTP protocol (see HTTP claims notice below). As part of getting this information the registrar provides the matching name and the TLD it is being registered in. The claims notice contains all relevant information including the mark and TLD and also contains a signature of the notice generated with the TMCH private key (see Sunrise process above)
5. The registrant reads the notice and chooses to continue with the registration.
6. The registrar then sends the create command to the registry including the signature from the clearinghouse. This allows the registry to verify that he claims notice was retrieved from the clearinghouse. Again a very simple

EPP extension is required for this, and most registries already have key-value pair extensions they can use for this.

7. The registry verifies that the name is subject to claim and that the supplied signature is valid. If valid the registration succeeds, if not it fails.
8. The registry periodically notifies the TMCH of all names registered and signatures utilised (for optional verification by the TMCH). The TMCH can then report on and notify mark holders. This should be create time stamped and use the same processes as described in sunrise above.

3.2.1 DNS List of Marks

Given that the TMCH has already deduced the list of marks and the labels that represent those marks, publishing that list as a 'zone file' of DNS names in a known zone (say tmch.icann.org) is relatively trivial. The choice of DNS makes sense as all the clearinghouse need do it provide a master service that registries can transfer the zones from and not have to deal with large number of queries. Also registries are already familiar with DNS technologies. ARI are confident that registries will volunteer to secondary the claims 'zone' and should do so at no cost given most registries already have large DNS infrastructures that could easily accommodate the estimated query volumes, and meet the availability requirements (100%).
Registries should only allow registrars to query the DNS infrastructure and get access to the random string codes needed to get claims information, further protecting the clearinghouse from mining. Larger registrars may be allowed to secondary the zone as well, bringing the information closer to where it is required.
One possibility is that the zone file contains text records that will contain a random string needed to complete the predictable URL for retrieving records from the clearing house. Alternatively the clearinghouse could just maintain a list of IP addresses and link them back to known registrars, or they could do both. Either of these methods will protect the data from the general public.

3.2.2 HTTP Claims Notice

If the codes are to be used, when a claim is matched the registrar now has all the information it requires to construct a URL, e.g.:

http://claimsnotice.tmch.icann.org/claims/generatNotice?label=<mark>&tld=<tld>&randomCode=<code-from-dns>

As stated above, alternatively the web server could maintain ACLs for each registrar, however that may become unmanageable (and until now the TMCH has had no need to deal directly with registrars). This URL is for a simple 'rest like' web service. The TMCH would then return the notice in a JSON or XML format, which the registrar can frame or use AJAX like practices to present to the potential registrant.
As mentioned above, contained within the notice will be a code unique to that mark/TLD combination which must be returned to the registry.

3.2.3 Issues with this Model and How to Address These

This proposed alternate model comes with one potential drawback, explored below, however ARI propose that despite this drawback, the alternate model is a better approach than the one currently suggested by ICANN.

This proposed alternate model places technical requirements on the TMCH, to provision claims notices over a HTTP based web service. However the provision of this simple (web) service does not differ greatly from the provision of the mark registration and validation service they will already need to supply. However the ICANN requirements currently impose that registrations cannot proceed unless relevant claims notices have been presented and accepted. This places the ability to register names within the hands of the TMCH.

However, the requirement preventing a registration if claims information cannot be determined is strongly encouraged to be relaxed. To prevent the abuse of this relaxation, ARI propose that registrations can be submitted if a TMCH outage is declared. It is a **must** to keep in perspective that if the TMCH provider is held to strict SLAs with financial penalty (say 99.99%), this means a registration window of 86 minutes across the entire 60 day period. Given the DNS lookup mechanism registrations in this window will still be displayed, a generic notice indicates that the name matched marks in the clearing house, and the TMCH is still notified that the registration went ahead (so they can still notify the mark holders); the only difference is that in this situation the complete information about the conflicting marks is not displayed.

The generic claims notice displayed would also include information that states that an email will be received by the registrant from the clearinghouse with the complete claims data. The notice would include warnings of how important it is that the registrant contact email address is correct and functioning. This notice could say that once they receive the email with the information if they decide they do not want the name any more they can contact their registrar to return the name for a refund (assuming they do so within the add grace period), however this will require more discussion with Registrars before it becomes a 'feature' of the model. The notice could alternatively say that they proceed at their own risk, and if later when they see the claims information, if they are not happy with it, there is no guarantees about what happens (they may loose the name, may decide to delete it to avoid risk or so forth) and that if they are not happy with taking such risk, to return later when the TMCH service is up. More thought and discussion with Registrars is required on this topic.

The email sent by the clearinghouse would include all the relevant marks that match the name and similar instructions. The clearinghouse would know to send the email based on a notification of registration being sent to them with a missing or hard coded string such as GENERIC-NOTICE in place of the 'claims notice code'.

This does however mean that the Registry will need to also send the email address of the Registrant to the clearinghouse, this shouldn't be an issue though as that information is publically available in the WhoIs interface anyway.

This is considered acceptable based on the assumptions that:

1. Registrants that intend to register names in bad faith will continue to do so, regardless of the presentation of claims data.
2. Mark holders that choose not to participate in the claims process are not disadvantaged in any available options for remedy, thus those that did participate but are not protected for this short period of time can still remedy the situation, URS and UDRP will still be available to them during this limited window, and if they are in the TMCH they will still be notified of the registration.
3. Registration processes will most likely not include the claims process beyond the mandatory minimum period of 60 days.
4. The entire TMCH and URS situation is a considerable improvement for rights holders over the current situation, thus a compromise needs to be made somewhere (it is also a significant change, which brings risk). We believe this process is a good first step and that learnings from this can be applied to a second gTLD release round when that occurs.
5. As described in 2.2.2 above no model is going to be able to guarantee that the TMCH data is available thus we need an appropriate solution to the situation, registries simply cannot just stop taking registrations.

This would effectively eliminate registration 'downtime' that may otherwise occur from unavailability of the service. The financial penalty mentioned above is important as it ensures that the TMCH takes the system seriously and motivates them to invest the correct amount of resources into quality development, systems and testing.

3.3 Fees

ARI proposes the following simple billing model for the TMCH to recover costs:

- Entry into clearing house:
  - Entry and validation of information into clearinghouse by mark holders
- Sunrise:
  - Either request of sunrise code by mark holder, or at notification of registration using sunrise code by the registry (only one not both)
- Trademark Claims:
  - At notification of names being registered that match claims during the claims period. Each month, the TMCH could bill the registry for all claims notices issued for names were the registration actually proceeded (where the code was sent to the TMCH).

These seem to match up with when the clearinghouse needs to perform work, and coincide with how registries receive revenue.

3.4 Benefits:

Overall ARI see the benefits of this model as:

- Fast (using DNS)
- Highly available (sunrise, and claims), and supportive of continued registration operations should the TMCH be unavailable for an extended period of time.
- No back and forth
- No complicated EPP extensions—in fact for most registries/registrars there is nothing to do from an EPP perspective
- Registrars still have to do a lot of work, however they may do less with this alternate model
- The TMCH and a mark holder are the only two entities with information mark holder consents to give to a registry only when they have a reason to (and can do so at there will).
- The TMCH sees all data access and can detect and prevent mining, effectively protect data
- No privacy issues
- Less costly for all parties (ultimately all costs will make it back to registrants)
- Simple