

Building Lasting Trust:

The Game Dynamics of the Respect Trust Framework

A Connect.Me White Paper
03 November 2011



Executive Summary

The Respect Trust Framework was listed with Open Identity Exchange in May 2011 to lay the legal foundation for a new Internet-scale peer-to-peer trust network. This network layers over existing social networks and adds a new form of contextual reputation called social vouching. To be successful, such a network must defend against gaming, particularly the Sybil attack (bogus accounts). This white paper explains the game dynamics of establishing four levels of trust through which members progress based on social verification of identity and relationships. It also examines the key role of members serving as trust anchors: the “Wikipedia admins” of the network.

Table of Contents

INTRODUCTION	2
WHY GAME DYNAMICS?	2
WHY PEER-TO-PEER SOCIAL VOUCHING?	3
THE FOUR TRUST LEVELS	4
LEVEL 1: UNVERIFIED	5
LEVEL 2: VERIFIED	5
LEVEL 3: TRUSTED.....	5
LEVEL 4: TRUST ANCHOR.....	5
TRUST ANCHORS	6
FOUNDING TRUST ANCHORS.....	6
HOW TO BECOME A FOUNDING TRUST ANCHOR.....	7
DISTINGUISHED TRUST ANCHORS.....	7
NEGATIVE REPUTATION: COMPLAINTS	8
THE GOAL: A SUSTAINABLE SOCIAL WEB-OF-TRUST	8
HOW TO JOIN THE CONNECT.ME BETA	8

Introduction

The purpose of the [Respect Trust Framework](http://openidentityexchange.org/trust-frameworks/respect-trust-framework),¹ announced at the European Identity Conference in May 2011 (where it won the [Privacy Award](http://blog.connect.me/connectme-wins-2011-eic-privacy-award)),² is to create the international legal foundation for a peer-to-peer reputation network that protects and promotes the ethical use of personal data.

The Respect Trust Framework is publicly listed with the [Open Identity Exchange](http://www.openidentityexchange.org/)³ (OIX), an international non-profit organization for the advancement of open digital trust frameworks. The first public review draft announced in May 2011 has already benefited greatly from the comments received. These resulted in a second draft, [Version 1 Beta](http://openidentityexchange.org/sites/default/files/respect-trust-framework-v1-beta-2011-08-15_0.pdf),⁴ being listed with OIX on August 15, 2011.

The purpose of this white paper is to explain the game dynamics at the heart of the Respect Trust Framework's peer-to-peer reputation system. Most (but not all) features of this system have now been implemented by [Connect.Me](http://connect.me/),⁵ and our goal is for users, bloggers, industry analysts, and press to all have a clear and transparent understanding of this mechanism, its strengths, and its weaknesses so we all can contribute to improving it, just like open source software.

Why Game Dynamics?

The rise of the social web has led to increasing appreciation for game dynamics (also called [game mechanics](http://en.wikipedia.org/wiki/Game_mechanics))⁶ as a tool for driving engagement in online games, particularly [social games](http://en.wikipedia.org/wiki/Social_games).⁷ But why would a trust network have a game dynamic—*trust is not a game*.

On the contrary, trust between real people—online or offline—is in fact a perfect example of [game theory](http://en.wikipedia.org/wiki/Game_theory), which, to quote the Wikipedia page, “reflects calculated circumstances (games) where a person's success is based upon the choices of others”.⁸ Since the trust you earn from others is always their choice, and since the reward can be anything from friends to jobs to political office, trust is *the* game that all of us as social animals are playing throughout our lives.

From this perspective, the goal of the Respect Trust Framework is simply to emulate the rules of the “real life game of trust” in an online network so that we may extend the benefits of the trust we earn in real life to our digital lives and interactions.

¹ <http://openidentityexchange.org/trust-frameworks/respect-trust-framework>

² <http://blog.connect.me/connectme-wins-2011-eic-privacy-award>

³ <http://www.openidentityexchange.org/>

⁴ http://openidentityexchange.org/sites/default/files/respect-trust-framework-v1-beta-2011-08-15_0.pdf. The name “Version 1 Beta” is intentional: the trust framework is expected to evolve and improve just like open source software.

⁵ <http://connect.me/>. See the final section, *How to Join the Connect.Me Beta*.

⁶ http://en.wikipedia.org/wiki/Game_mechanics

⁷ http://en.wikipedia.org/wiki/Social_games

⁸ http://en.wikipedia.org/wiki/Game_theory

Why Peer-to-Peer Social Vouching?

Many organizations and institutions with whom we have trust relationships could be in a position to help us extend that trust online. This is the premise of [federated identity management](#)⁹ as embodied by several initiatives over the past decade, including the Liberty Alliance, Kantara, Shibboleth, and OpenID. More recently, OAuth-based social login services are leveraging the trust relationships we have with social networks like Facebook, Twitter, and LinkedIn.

However in real life we have an order of magnitude more personal relationships than organizational relationships. Since Connect.Me and the Respect Trust Framework begin with the empowerment of individuals, our trust fabric begins with personal trust relationships that are independent of any organization.¹⁰

We refer to the person-to-person contextual reputation statements upon which the Respect Trust Framework is based as *social vouching* because:

- They build on the existing fabric of friend and follower relationships already established on social networks like Facebook, Twitter, and LinkedIn.
- A vouch is a personal assurance of another person's qualities and trustworthiness.

Each vouch from one person to another is a "+1" in a specific context, called a **tag**. Every person can tag him/herself, and may also suggest tags for others by vouching for them with that tag. Each person also controls the visibility of his/her own tags, i.e., if you receive a vouch on a new tag, it is not public unless you decide to display it. (If you do not want that tag, you may simply delete the vouch.)

Figure 1 shows an example Connect.Me card for Flora H. displaying six tags and the number of vouches Flora has received for each tag.




Figure 1: An example Connect.Me card

⁹ http://en.wikipedia.org/wiki/Federated_Identity_Management

¹⁰ This does not mean that organizations or governments will not be welcomed into the network, only that the starting point for trust is individuals.

The Four Trust Levels

The key challenge for any peer-to-peer (p2p) trust network is to prevent gaming attacks. The best known is the [Sybil attack](#),¹¹ where an attacker creates multiple fake accounts (“sock puppets”), then uses them to artificially manipulate reputations. Even if fake accounts can be prevented, a p2p network must still deal with *conspiracy attacks*, where groups of members collaborate to manipulate reputations by making dishonest reputation statements (positive or negative).

An effective trust model must guard against both types of attacks. The Respect Trust Framework achieves this by having members progressing through a four-level game dynamic based on social verification of identity and relationships (Figure 2). Each level increases assurance that:

1. An individual member is a real person complying with the **One-Person-One-Account** rule.¹²
2. The member is vouching honestly and not trying to game the system.




Figure 2: The four trust levels for individual members in the Respect Trust Framework

Table 1 lists the specific requirements for each trust level. Note that because this trust model uses social verification, it does *not* require “real names”, i.e., members are free to use online pseudonyms to protect their real-world identity.¹³

¹¹ http://en.wikipedia.org/wiki/Sybil_attack, named after the famous case of multiple personality disorder.

¹² One of six basic Accountability Rules – see page 10 of the Respect Trust Framework V1 Beta.

¹³ See <http://my.nameis.me> for why social software should support online pseudonyms.

Trust Level	Cumulative Requirements	Notes
Level 1: Unverified	<ul style="list-style-type: none"> • Successfully register using one social networking account (as of October 2011: Facebook, Twitter, or LinkedIn) • Agree to honor the Respect Trust Framework. 	<p>This is the lowest level of assurance of compliance to the One-Person-One-Account rule because Twitter allows individuals to register multiple accounts. Vouches issued by an Unverified Member will eventually expire unless the member becomes Verified.</p>
Level 2: Verified	<p>Account Verification:</p> <ul style="list-style-type: none"> • Connect a Facebook and/or LinkedIn account with a combined total of 100 or more connections, OR • Verify a unique mobile phone number <p>Vouching (“10/3 rule”):</p> <ul style="list-style-type: none"> • Vouch for 10 other members at any trust level • Receive vouches from 3 other members at any trust level 	<p>The Terms of Service (TOS) for Facebook and LinkedIn require using a “real identity”, and relatively few people have multiple mobile phones, so the combination of account verification and vouching significantly increases assurance of compliance with the One-Person-One-Account rule. It also establishes an initial base of references from other members.</p>
Level 3: Trusted	<p>Vouching (“25/25 rule”):</p> <ul style="list-style-type: none"> • Vouch for 25 other members at any trust level • Receive vouches from 25 other Verified or higher members 	<p>The requirement of 25 vouches from other Verified members broadens and hardens the base of references.</p>
Level 4: Trust Anchor	<ul style="list-style-type: none"> • Receive Trust Anchor vouches from 3 other Trust Anchors (see the next page) 	<p>This highest level of assurance is based on an explicit public chain of trust starting from individuals whose identity has been publicly verified. See the following pages.</p>

Table 1: The requirements for each of the four trust levels

Trust Anchors

After tags and vouching, the most important concept in the Respect Trust Framework is *trust anchors*. This is because a 2005 paper by Alice Cheng and Eric Friedman called [Sybilproof Reputation Mechanisms](#)¹⁴ proves mathematically that it is impossible to prevent Sybil attacks in a p2p reputation system unless it establishes a known set of trusted accounts. These accounts literally provide the “anchor points” from which chains of trust can be calculated.

In the Respect Trust Framework, these trust anchors are *people*, and the chains of trust are formed by having the initial trust anchors vouch for other people with whom they have the strongest trust bonds, “deputizing” them as trust anchors. This *trust anchor vouching* is just like regular vouching except it takes place in a pre-defined context whose meaning is specified in the Respect Trust Framework:¹⁵

Every Trust Anchor agrees that a Vouch in the Trust Anchor Context means that the Voucher personally knows the Recipient and has good reason to believe that the Recipient will honor and abide by the Respect Promise: “I promise to uphold the purpose, principles, and rules of the Respect Trust Framework.”

Once a person at the Trusted level receives three vouches as a Trust Anchor, they are automatically promoted to the Trust Anchor level and may begin vouching for other trust anchors. In this manner, the population of trust anchors grows organically, similar to the way Wikipedia maintains its population of editors.

This person-to-person trust fabric stands in contrast to most hierarchical PKI (Public Key Infrastructure) systems where the trust anchors are corporations or governments because they can afford the attendant liabilities. In a p2p trust network, these liabilities are spread across the reputations of millions of individuals, each of whom has a personal stake in defending it.

Founding Trust Anchors

One key challenge remains in this design: how to bootstrap the initial seed population of trust anchors, called the *Founding Trust Anchors*. In the Respect Trust Framework the solution is very pragmatic: direct personal enrollment, beginning with active members of the Internet identity, privacy, and security communities.

Connect.Me began enrolling Founding Trust Anchors at four Internet identity and privacy conferences held during May 2011: Internet Identity Workshop (Mountain View, CA); European Identity Conference (Munich); Telco 2.0/Personal Data 2.0 (London); and Privacy/Identity/Innovation (Santa Clara, CA).

¹⁴ <http://www.sigcomm.org/sigcomm2005/paper-CheFri.pdf>

¹⁵ Respect Trust Framework Version 1 Beta, page 10, “Accountability for Trust Anchor Vouches”.

As of November 2011 approximately 200 Founding Trust Anchors have been appointed. They can be identified by the special trust level ribbon that appears on their Connect.Me card as shown in Figure 3:

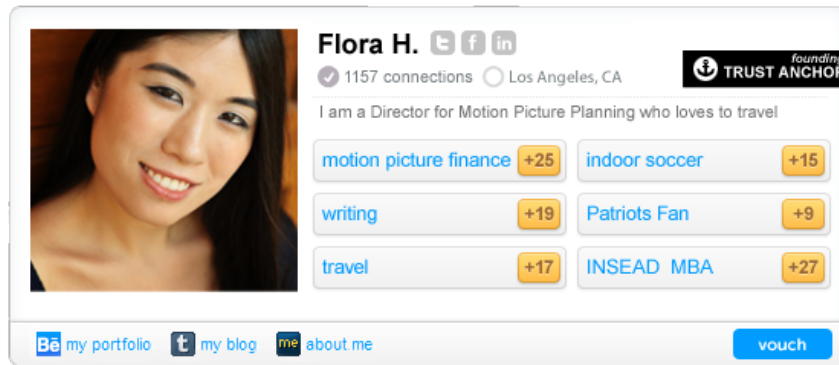


Figure 3: An example Connect.Me card for a Founding Trust Anchor

How to Become a Founding Trust Anchor

Enrollment of Founding Trust Anchors will continue as long as required to provide a large enough seed population. Nominations may be made in three ways:

1. Founding Trust Anchors may nominate other Founding Trust Anchors.
2. Members reaching the Trusted level may request nomination.
3. Individuals interested in serving as a Founding Trust Anchor may contact Connect.Me directly by sending email to **anchors@connect.me**.

Distinguished Trust Anchors

Certain Founding Trust Anchors have made special contributions to the development of user-centric Internet identity, privacy, and trust infrastructure. To honor this work and shine a light on these individuals as examples of what it means to be a trust anchor, Connect.Me has created the Distinguished Trust Anchor program. Recipients of this honor embody the following definition:

A Distinguished Trust Anchor is an individual whose words and deeds exemplify the spirit and principles of the Respect Trust Framework.

The first Distinguished Trust Anchors will be announced when Connect.Me introduces trust anchor vouching (expected before the end of 2011). The complete list will be available at connect.me/trust-anchors.

As with Founding Trust Anchors, additional Distinguished Trust Anchors will be nominated by the initial group and announced on a regular basis beginning in 2012.

Negative Reputation: Complaints

No matter how well a p2p reputation system rewards good behavior, there will still be bad actors. While arguably this could be dealt with by withdrawing positive reputation statements (deleting vouches) from the offending members, a more efficient enforcement mechanism is allowing members to assert negative reputation statements, called **complaints**.

To prevent complaints from being used to game the system (e.g., to try to harm competitors), the cost of submitting a false complaint must be very high (i.e., it should backfire and damage the submitter's reputation). Making this determination is quintessentially a human decision (e.g., judges and juries) so the Respect Trust Framework assigns complaint moderation to the members at the highest trust level: trust anchors. See pages 11 and 12 of the Respect Trust Framework for details.

The Goal: A Sustainable Social Web-of-Trust

Wikipedia has proved that using the wisdom of the crowd to protect against the most sophisticated attacks on a peer-driven trust community can work at Internet scale. As the [Reliability of Wikipedia](#) article states:¹⁶

The Wikipedia model allows anyone to edit, and relies on a large number of well-intentioned editors to overcome issues raised by a smaller number of problematic editors. It is inherent in Wikipedia's editing model that misleading information can be added, but over time quality is anticipated to improve in a form of group learning as editors reach consensus, so that substandard edits will very rapidly be removed.

The goal of the Respect Trust Framework is to apply the Wikipedia model to a p2p reputation network. By having trust anchors vouch for other trust anchors and moderate complaints, they essentially serve as the “Wikipedia admins” of the network. Like Wikipedia, we do not expect it to be perfect. But like Wikipedia, we do believe it can build a self-regulating and self-healing social web-of-trust that can foster many new forms of innovation and value.

How to Join the Connect.Me Beta

To join the beta, you must either be signed up on the Connect.Me beta invite list or be vouched for by someone already in the network. To sign up for the beta invite list, please visit [connect.me](#) and follow the instructions. If you are interested in serving as a Founding Trust Anchor, please send email to anchors@Connect.Me.

¹⁶ http://en.wikipedia.org/wiki/Reliability_of_Wikipedia, 2011-04-20