

# Security Evaluation of Free/Open Source Software Powered by a Peer-to-Peer Ecosystem

Jean-Marc Seigneur

Université de Genève  
Jean-Marc.Seigneur@cui.unige.ch

**Position Paper.** The EU-funded EDOS project [1, 2] aims at providing an environment to improve the development and distribution of Free/Libre and Open Source Software (F/LOSS or FOSS). At a time when FOSS is considered for mission critical applications, it becomes crucial to correctly monitor the security of open source software. This short paper depicts how the EDOS ecosystem can be used to improve the evaluation and monitoring of FOSS security.

A major tool in the EDOS tool box is the Project Management Interface (PMI) [3] because it allows any F/LOSS community member to query and monitor the whole F/LOSS ecosystem. The reference PMI implementation has been being developed as an Eclipse plug-in. Distribution providers can use the PMI to give maintainer role to known members and display information about their project members and packages. Two Linux distribution providers are part of the EDOS project, namely, Mandriva [4] and Caixa Mágica [5]. Distribution contributors are noticed when any tasks that they can contribute to are posted and can prove to other distributions that they are knowledgeable in what they have contributed to before. From a security point of view, roles may be given based on reputation and computational trust in the contributors [6]. Some information may be restricted to specific roles and projects. The end-users of any of the F/LOSS projects (beyond distribution projects) should be able to retrieve information about their FOSS of interest: for example, they should be allowed to see the liveness of a project or its Business Readiness Rating<sup>TM</sup> (BRR). Being better informed, it becomes easier for the end-users to select the most suitable subproject among a set of subprojects with apparently the same functionalities. It is a F/LOSS digital ecosystem where the fittest parties prosper. However, the ecosystem is under threat if the fitness metric, for example based on BRR, is compromised. A question remains about the trustworthiness of the BRR displayed by the PMI if the trustworthiness of the underlying technical infrastructure is not considered. If the other building blocks of the EDOS project are compromised, any BRR could be compromised. The metric of interest may also be the security metric: in this case, the user wants to find the most secure FOSS.

Firstly, if we assume that all the underlying technical infrastructure is perfectly secure, the evaluation of specific open source software may be based on external information, for example, from the Advogato [7] Web site. In a service like Advogato, the users rate themselves with regard to their expertise in software development and then when they rate a FOSS project their expertise is taken into account in the global rating of the FOSS project. The rating given from one user to a second user may be considered as the *recommending trustworthiness* [8] of the second user. The global rating may be more or less resilient to attacks, for example,

the Advogato rating mechanism is argued to be quite attack-resistant. However, this rating remains static over time if no user updates the ratings. Another more dynamic approach would be to use the technical information that is implicitly given when the users report a FOSS bug or when a specific FOSS crashes and a crash report is propagated. For example, the rating would be based on the number of FOSS bugs or the number of times that the FOSS has crashed. Ideally, the dynamic information based on objective evidence count (such as the number of bugs or crashes) and the subjective evidence given by the user's manual rating should be combined to improve the security rating. Once one has the rating of the sub-components that compose the FOSS under consideration, the rating of the high-level FOSS may be inferred from the ratings of the sub-components. For example, the rating may correspond to a *technical trust value* [8] between 0 (untrustworthy) and 1 (fully trustworthy), which is the result of the multiplication of the technical trust values of all its  $N$  sub-components:

$$FOSSTrustValue = \prod_{i=1}^N SubComponent_iTrustValue = CryptoLibraryTrustValue \times \dots$$

Secondly, if we assume that the underlying technical infrastructure is prone to attacks and is not perfectly secure, the rating impact of the information supposed to be given by a specific user or peer must be taken into account. For example, if the peer/machine of the user is known to be outdated with regard to security patches, then the information given by the peer should be discounted in the global rating process.

The notion of peer is important in EDOS because the F/LOSS community members will share their resources in a peer-to-peer fashion to improve FOSS distribution performance. Thus, the EDOS building block that improves the performance and scalability of the distribution of the ecosystem information and packages, namely, peer-to-peer technology, is the root of a significant security risk that needs to be evaluated to be taken into account in the global rating process and ideally mitigated. The risk is due to two parts of the EDOS distribution platform. The first part, that is, KadoP [9], enables the publication and discovery of content in a peer-to-peer way based on a Distributed Hash Table (DHT) for indexing content in the network. As for other peer-to-peer DHTs, a major risk is that an attacker uses multiple faked identities to take control of a specific part of the DHT, which is related to Douceur's work on the Sybil attack [10]. The risk can be mitigated when known Certification Authorities (CAs) are trusted in advance and it is the case in EDOS when there is a well-known F/LOSS publisher, such as Mandriva [4] or Caixa Mágica [5], which can act as the root CA. A more difficult alternative may be to consider a fully decentralised identity management system [11] that would be needed if we open EDOS to the fully decentralised world of the F/LOSS community, beyond the boundaries of a specific publisher, including any other project stakeholders, such as Nuxeo [12] enterprise content management or Nexedi [13] enterprise resource planning (two other industrial partners of the EDOS consortium). The potential problem of the use of ActiveXML [14] and its intentional data feature – parts of the XML document consists of method calls that can be triggered on demand to return new XML information – has been mitigated by constraining the intentional data feature of ActiveXML in EDOS to only specific and well-known harmless method calls. However, the problem that a peer may not answer correctly (either due to maliciousness or reliability issues) to ActiveXML queries remains. The second part of

the distribution platform that contributes to its risk significance is called IDiP [15]. IDiP enables the efficient dissemination of data to users with heterogeneous needs based on a clustering subsystem that groups users having similar requests. The problem is that this approach might lead to intelligent network-engineered topology attacks [8] using the knowledge of the clustering topology.

To mitigate this risk, although adjunct cryptographic security and a thorough modification of the chosen DHT implementation are required, it seems that it is not sufficient. In fact, the chosen DHT implementation for EDOS is FreePastry [16], which is not yet implemented with the complex security mechanisms proposed to secure Pastry [17]. CPS [18] has been working on delivering a secure implementation of a DHT based on Pastry for EDOS. Anyway, these security mechanisms to secure routing in Pastry require more than secure assignments of node identifiers as mentioned above: secure routing table maintenance and secure message forwarding are also required. Unfortunately, even with these mechanisms, secure routing cannot be maintained with more than 25% of malicious participating nodes. One may argue that few attackers would be powerful enough to spend a few millions to gain control of a large-enough part of the peer-to-peer network. However, a pessimistic view is adopted given the power of a few non-F/LOSS business parties and the dramatic impact of a massive F/LOSS collapse. Even if the collapse is short, reputation is said to drop very quickly and be hard to be rebuilt, especially at a level where F/LOSS is considered for mission critical applications. For example, if the cost to acquire a certificate is 20, it (only) costs 12 millions to control 10% of a network of 6 millions of peers, which roughly corresponds to the current number of all Mandriva-based user machines.

To reach an acceptable level of risk mitigation, the investigation of the use of peer behavioural information extracted from other EDOS building blocks is underway. However, if the peer-to-peer building block becomes compromised, it might compromise the integrity of the other EDOS building blocks because the peer-to-peer building block is one link of the whole EDOS ecosystem and security is said to be as good as the weakest link. Thus, if the peer-to-peer building block starts to become compromised, it might be the case that the building blocks that provide behavioural information for further security rating start, in turn, to become compromised. This vicious circle might end up with the total collapse of the digital ecosystem. To unambiguously avoid this, EDOS is given a consistent, cross-building blocks, solution to security. Ideally, the PMI [3] building block will be used to manage behavioural peer information for security ratings. Again, due to the cyclic relationship with regard to security between the peer-to-peer building block and the other EDOS building blocks – the peer-to-peer building block stores the information used by the PMI building block – the security aspect of the PMI may have to be revised in light of its relationship with the peer-to-peer building block. It is expected to evaluate the overall ecosystem security solution with the real F/LOSS information, for example based on the EDOS F/LOSS dependency and test tools [19], stored by the KadoP/ActiveXML/IDiP/SecurePastry peer-to-peer building block extracted by the reference Eclipse PMI plug-in. As mentioned above, we will also consider the use of BRR from a security point of view and how a BRR could be set up to reflect a specific security aspect.

This work is sponsored by the European Union, which funds the FP6-IST-004312 EDOS project. The views expressed in this paper are solely the views of the author and do not necessarily reflect the official views of the EDOS consortium.

## References

- [1] Environment for the development and Distribution of Open Source software (EDOS), <http://www.edos-project.org/>, access date: 07/04/2006.
- [2] S. Abiteboul, X. Leroy, B. Vrdoljak, R. Di Cosmo, S. Fermigier, S. Laurière, F. Lepied, R. Pop, F. Villard, J.-P. Smets, C. Bryce, K. R. Dittrich, T. Milo, A. Sagi, Y. Shtossel, and E. Panto, "EDOS: Environment for the Development and Distribution of Open Source Software," 2005.
- [3] M. Pawlak, "Project Management Interface (PMI)," EDOS Project Deliverable 5.5.1, 2005.
- [4] Mandriva, <http://www.mandriva.com>, access date: 07/04/2006.
- [5] Caixa Mágica, <http://www.caixamagica.pt>, access date: 07/04/2006.
- [6] Trustcomp, <http://www.trustcomp.org/>, access date: 08/04/2006.
- [7] Advogato, <http://www.advogato.org>, access date: 09/05/2006.
- [8] J.-M. Seigneur, "Trust, Security and Privacy in Global Computing," Trinity College Dublin, PhD Thesis Technical Report TCD-CS-2006-02, 2005, <http://www.cs.tcd.ie/publications/tech-reports/reports.06/TCD-CS-2006-02.pdf>.
- [9] KadoP, <http://gemo.futurs.inria.fr/projects/KadoP/>, access date: 07/04/2006.
- [10] J. R. Douceur, "The Sybil Attack," in Proceedings of the 1st International Workshop on Peer-to-Peer Systems, 2002.
- [11] J.-M. Seigneur, "Decentralized Identity for the Digital Business Ecosystem," in *ERCIM News*, 2005.
- [12] Nuxeo, <http://www.nuxeo.com/>, access date: 07/04/2006.
- [13] Nexedi, <http://www.nexedi.com/>, access date: 07/04/2006.
- [14] ActiveXML, <http://activexml.net/>, access date: 07/04/2006.
- [15] T. Milo, A. Sagi, and E. Verbin, "Compact Samples for Data Dissemination," Tel Aviv University, 2005.
- [16] A. Rowstron and P. Druschel, "Pastry: Scalable, distributed object location and routing for large-scale peer-to-peer systems," presented at International Conference on Distributed Systems Platforms, 2001.
- [17] M. Castro, P. Druschel, A. Ganesh, A. Rowstron, and D. S. Wallach, "Secure routing for structured peer-to-peer overlay networks," 2002.
- [18] CSP, <http://www.csp.it/>, access date: 07/04/2006.
- [19] R. Di Cosmo, B. Durak, X. Leroy, F. Mancinelli, and J. Vouillon, "Maintaining large software distributions: new challenges from the FOSS era," in Proceedings of the FRCSS 2006 workshop, EASST Newsletter, 2006.