Security for Free/Open Source Software Powered by Peer-to-Peer

Jean-Marc Seigneur

Position Paper

The EDOS [1, 2] project aims at providing an environment to improve the development and distribution of Free/Libre and Open Source Software (F/LOSS or FOSS). This short paper depicts the roadmap that has been followed with regard to security evaluation in EDOS.

A risk analysis on the whole EDOS building blocks has brought to light that the building block to improve distribution performance and scalability, namely, peer-to-peer technology, is the root of a significant security risk that needs to be mitigated.

The risk is due to two parts of the EDOS distribution platform. The first part, that is, KadoP [3], 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 [4]. 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 [5] or Caixa Mágica [6], which can act as the root CA. A more difficult alternative may be to consider a fully decentralised identity management system [7] 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 [8] enterprise content management or Nexedi [9] enterprise resource planning. The potential problem of the use of ActiveXML [10] 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. 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 [11]. 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 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 [12], which is not yet implemented with the complex security mechanisms proposed to secure Pastry [13]. CPS [14] 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 a secure assignment 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 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 is 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 start, in turn, to become compromised and this vicious circle may end up with the total collapse of the EDOS ecosystem. This is the motivation to consider a consistent, cross-building blocks, solution to security in EDOS. Ideally, the Project Management Interface (PMI) [15] building block will be used to manage behavioural peer information. 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 will be examined 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 [16], stored by the KadoP/ActiveXML/IDiP/SecurePastry peer-to-peer building block extracted by the reference Eclipse PMI plug-in. To summarise, the main milestones of the roadmap are:

1. P2P-based F/LOSS Ecosystem Security Preliminary Design (June 06)
2. PMI-based Preliminary Implementation and Evaluation (Sept. 06)
3. Integration Report in the EDOS Security Library Section D4.2.2 (Sept. 06)
4. Final Evaluation during the Integration with the EDOS Reference Implementation (March 07)

This work is sponsored by the European Union, which funds the FP6-IST-004312 EDOS project.

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] KadoP, http://gemo.futurs.inria.fr/projects/KadoP/, access date: 07/04/2006.

[4] J. R. Douceur, "The Sybil Attack," vol. Proceedings of the 1st International Workshop on Peer-to-Peer Systems, 2002.

[5] Mandriva, http://www.mandriva.com, access date: 07/04/2006.

[6] Caixa Mágica, http://www.caixamagica.pt, access date: 07/04/2006.

[7] J.-M. Seigneur, "Decentralized Identity for the Digital Business Ecosystem," in *ERCIM News*, 2005.

[8] Nuxeo, http://www.nuxeo.com/, access date: 07/04/2006.

[9] Nexedi, http://www.nexedi.com/, access date: 07/04/2006.

[10] ActiveXML, http://activexml.net/, access date: 07/04/2006.

[11] T. Milo, A. Sagi, and E. Verbin, "Compact Samples for Data Dissemination," Tel Aviv University, 2005.

[12] 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.

[13] M. Castro, P. Druschel, A. Ganesh, A. Rowstron, and D. S. Wallach, "Secure routing for structured peer-to-peer overlay networks," 2002.

[14] CSP, http://www.csp.it/, access date: 07/04/2006.

[15] M. Pawlak, "Project Management Interface (PMI)," EDOS Project Deliverable 5.5.1, 2005.

[16] R. Di Cosmo, B. Durak, X. Leroy, F. Mancinelli, and J. Vouillon, "Maintaining large software distributions: new challenges from the FOSS era," in *EASST Newsletter*, vol. Proceedings of the FRCSS 2006 workshop, 2006.