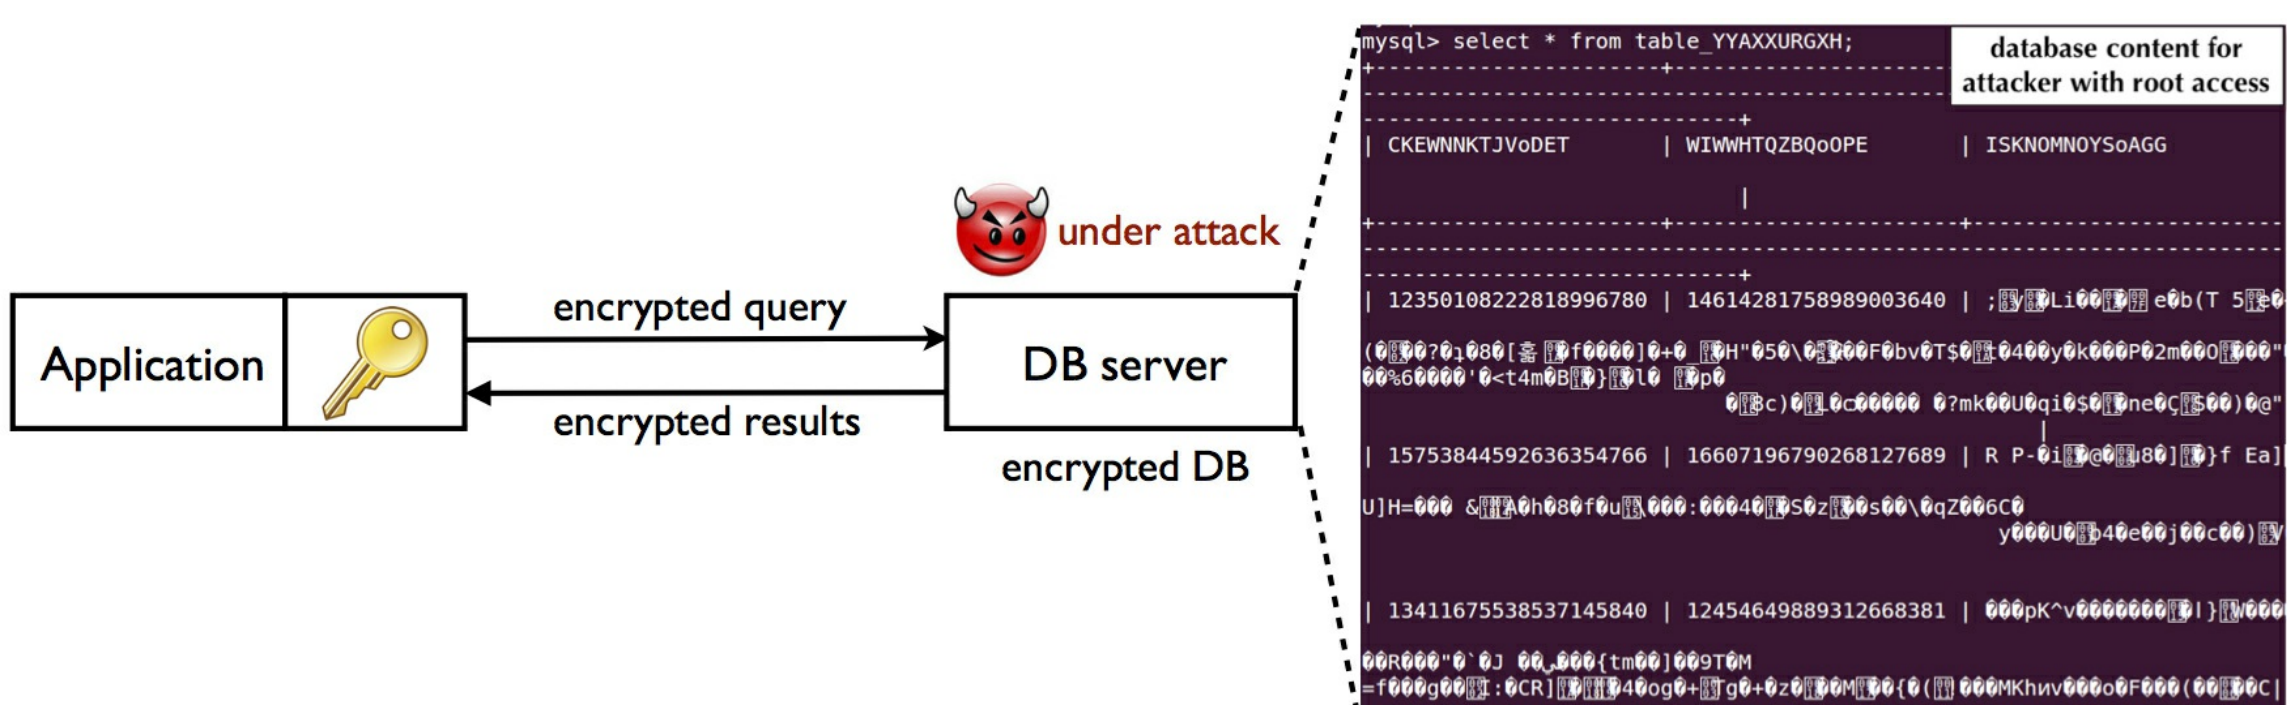


CryptDB

Online applications are vulnerable to theft of sensitive information because adversaries can exploit software bugs to gain access to private data, and because curious or malicious administrators may capture and leak data. CryptDB is a system that provides practical and provable confidentiality in the face of these attacks for applications backed by SQL databases. It works by *executing SQL queries over encrypted data* using a collection of efficient SQL-aware encryption schemes. CryptDB can also *chain encryption keys to user passwords*, so that a data item can be decrypted only by using the password of one of the users with access to that data. As a result, a database administrator never gets access to decrypted data, and even if all servers are compromised, an adversary cannot decrypt the data of any user who is not logged in. An analysis of a trace of 126 million SQL queries from a production MySQL server shows that CryptDB can support operations over encrypted data for 99.5% of the 128,840 columns seen in the trace. Our evaluation shows that CryptDB has low overhead, reducing throughput by 14.5% for phpBB, a web forum application, and by 26% for queries from TPC-C, compared to unmodified MySQL. Chaining encryption keys to user passwords requires 11-13 unique schema annotations to secure more than 20 sensitive fields and 2-7 lines of source code changes for three multi-user web applications.



People	Publications	Software	Impact	Press
------------------------	------------------------------	--------------------------	------------------------	-----------------------

People

- [Raluca Ada Popa](#)
- [Catherine Redfield](#)
- [Stephen Tu](#)
- [Hari Balakrishnan](#)
- [Frans Kaashoek](#)
- [Sam Madden](#)
- [Nickolai Zeldovich](#)
- [Aaron Burrow](#)

Publications

- Raluca Ada Popa, Catherine M. S. Redfield, Nickolai Zeldovich, and Hari Balakrishnan. [CryptDB: Protecting Confidentiality with Encrypted Query Processing](#). In *Proceedings of the 23rd ACM Symposium on Operating Systems Principles (SOSP)*, Cascais, Portugal, October 2011. *(This is the main paper describing CryptDB.)*
- Raluca Ada Popa. [Building Practical Systems that Compute on Encrypted Data](#). Ph.D. thesis, 2014. *(This thesis elaborates on various aspects of CryptDB.)*
- Raluca Ada Popa, Nickolai Zeldovich, and Hari Balakrishnan. [Guidelines for Using the CryptDB System Securely](#). In *Cryptology ePrint Archive*, Report 2015/979.
- Raluca Ada Popa, Frank H. Li, and Nickolai Zeldovich. [An Ideal-Security Protocol for Order-Preserving Encoding](#). In *Proceedings of the 34th IEEE Symposium on Security and Privacy (IEEE S&P/Oakland)*, San Francisco, CA, May 2013. *(This paper constructs the encryption scheme that computes order queries in CryptDB.)*
- Stephen Tu, M. Frans Kaashoek, Samuel Madden, and Nickolai Zeldovich. [Processing Analytical Queries over Encrypted Data](#). In *Proceedings of the 39th International Conference on Very Large Data Bases (VLDB)*, Riva del Garda, Italy, August 2013. *(This paper extends CryptDB's basic design to complex analytical queries and large data sets.)*
- Raluca Ada Popa and Nickolai Zeldovich. [Cryptographic treatment of CryptDB's Adjustable Join](#). Technical Report MIT-CSAIL-TR-2012-006, Computer Science and Artificial Intelligence Laboratory, Cambridge, MA, March 2012. *(A formal description and analysis of CryptDB's adjustable join cryptographic scheme.)*
- Carlo Curino, Evan P. C. Jones, Raluca Ada Popa, Nirmesh Malviya, Eugene Wu, Sam Madden, Hari Balakrishnan, and Nickolai Zeldovich. [Relational Cloud: A Database-as-a-Service for the Cloud](#). In *Proceedings of the 5th Biennial Conference on Innovative Data Systems Research (CIDR 2011)*, Pacific Grove, CA, January 2011. *(A paper describing how CryptDB can help with hosting databases in the cloud.)*
- Raluca Ada Popa, Nickolai Zeldovich, and Hari Balakrishnan. [CryptDB: A Practical Encrypted Relational DBMS](#). Technical Report MIT-CSAIL-TR-2011-005, Computer Science and Artificial Intelligence Laboratory, Cambridge, MA, January 2011. *(An earlier technical report on CryptDB, which has been superseded by the SOSP paper above.)*

Software

Our source code is available. The source code has not been maintained since March 2014 and hence, it has only been tested up to Ubuntu 13.04. We might resume maintaining the source code in 2015. You can access it using [git](#), as follows:

```
git clone -b public git://g.csail.mit.edu/cryptdb
```

To install, read [doc/README](#).

We will announce any significant changes to CryptDB on the cryptdb-announce mailing list.

If you are interested in using CryptDB's source code in any way or to receive announcements about CryptDB, we encourage you to subscribe to the cryptdb-announce mailing list below.

Mailing lists

Please subscribe to the cryptdb-announce mailing list to receive announcements about updates to CryptDB.








To subscribe, fill out the following form. You will be sent email requesting confirmation, to prevent others from gratuitously subscribing you.

Email address:

We no longer maintain the mailing list cryptdb-users. The list was used to ask questions about CryptDB, get help, and offer any suggestions. To see the collection of prior postings to the list, visit the [archive](#).

Impact

A few companies or organizations used or adopted CryptDB, or were inspired by CryptDB. Most of these companies got in touch with us and gave credit to CryptDB.

 SAP AG's system SEED	SAP AG developed a system called SEED, which implements CryptDB's design on top of their HANA database system. SEED uses most of the building blocks of CryptDB as well as the adjustable encryption (onion) strategy. Here are some references: Project SEED , white paper .
 Google's Encrypted BigQuery	Google has developed an experimental extension of the BigQuery client, known as Encrypted BigQuery , which was informed and motivated by the CryptDB paper. It offers client-side encryption for a subset of query types, using encryption building blocks similar to the RND, HOM, and DET used in CryptDB. Their code is available here .
 Lincoln Laboratory	Lincoln Labs added the CryptDB design on top of their D4M Accumulo no-SQL engine (using the RND, DET, OPE and HOM building blocks).
 Microsoft's Always Encrypted SQL Server	Microsoft's Always Encrypted SQL Server enables administrators to encrypt columns with RND and DET. Before this service, the database in the SQL Server was in plaintext during processing. Some applications can support a lot of fields with RND and a set of other fields with DET, thus giving a significant security increase as compared to no encryption for these fields. The service is now distributed as part of the SQL Server. The authors of Microsoft's Cipherbase system led this effort; Cipherbase is a successor of CryptDB which enhances CryptDB with trusted hardware support for queries not supported on encryption.
 Skyhigh Networks	Skyhigh networks seems to be using most of the encryption building blocks in CryptDB. Skyhigh discusses these schemes here .
 sql.mit.edu	sql.mit.edu is a SQL server at MIT hosting many MIT-ran applications. Volunteering users of Wordpress switched to running Wordpress through CryptDB, using our source code.
 Startups based on CryptDB	Privic, a startup in Silicon Valley, and Cryptonor, a startup in Europe, are both based on CryptDB's design. CryptonorDB targets no-SQL databases.

All the companies above except two have been in touch with us already and confirmed the relationship of their system to CryptDB. We have not yet been in touch with Skyhigh networks and Microsoft's Always Encrypted team.


Press

 [An MIT Magic Trick: Computing On Encrypted Databases Without Ever Decrypting Them](#), Andy Greenberg, Forbes Magazine, December 2011.

 [You want to crunch top-secret data securely? CryptDB may be the app for that](#), Barb Darrow, GigaOM, April 2013.

 [MIT Software Allows Queries On Encrypted Databases](#), Slashdot, December 2011.

 [CryptDB: Encrypted MySQL](#), Hackernews, December 2011.

 [Encryption's holy grail is getting closer, one way or another](#), ZDNet, July 2015.

