

The SIB Swiss Institute of Bioinformatics Secures Sensitive Data with WekaFS



Swiss Institute of Bioinformatics

“ We were looking for a solution that is easy to manage, delivers high performance, and protects data with encryption in-flight.

This solution ensures secure data access coupled with the necessary data analysis performance required by research groups.

*Dr. Heinz Stockinger,
Head of Core-IT, SIB Swiss
Institute of Bioinformatics*

CHALLENGES

- Prevent unauthorized users from accessing sensitive personal data
- Provide a platform with secure access to data for authorized users
- Find a storage solution that can handle mixed workloads
- Accelerate application performance to enable faster research and discovery

BENEFITS

- In-flight encryption without performance degradation
- Advanced authentication services ensuring only authorized access to data
- Protected, shared accessibility with multi-tenancy interoperability
- Ease of manageability

The Weka File System (WekaFS™), the world’s fastest shared parallel file system from WekaIO™ (Weka), secures data for the SIB Swiss Institute of Bioinformatics (SIB) by encrypting it both in-flight and at rest, with absolutely no application performance degradation. Established in 1998 in Switzerland, SIB is an internationally recognized center of expertise in data science, dedicated to biological data. It provides life scientists and clinicians with essential databases and custom-made tools, computational biology expertise, secure services for sensitive data as well as bioinformatics training. A non-profit foundation, the Institute also represents the Swiss bioinformatics community by federating about 800 data scientists in the country and fostering collaborations among them. As part of its infrastructure activities, SIB develops a nationwide project, called BioMedIT, aiming to implement a secure and cutting-edge IT network to support biomedical research and clinical bioinformatics. One of the three nodes of this network, the Romandie infrastructure at SIB operates a platform which offers biomedical researchers a full service for the processing of sensitive data. In this context, cyber security and trust are of utmost importance.

THE CHALLENGE: ENCRYPT DATA FOR SECURITY AND DELIVER HIGH PERFORMANCE

SIB supports the varied workloads of its healthcare and research community on a private cloud computing service based on OpenStack. The workloads are typically RAM-intensive or I/O-intensive applications with mixed file sizes that have sensitive personal data (e.g., patient data). SIB’s objective to provide a high-performance computing and storage platform that is shareable, interoperable, and accessible by different end-user groups is essential to its mission. However, because of the sensitivity of the personal [patient] data that SIB manages, its IT strategy has to comply with data protection regulations that allow its access to data for authorized users such as research scientists but restrict it for others, including storage administrators.

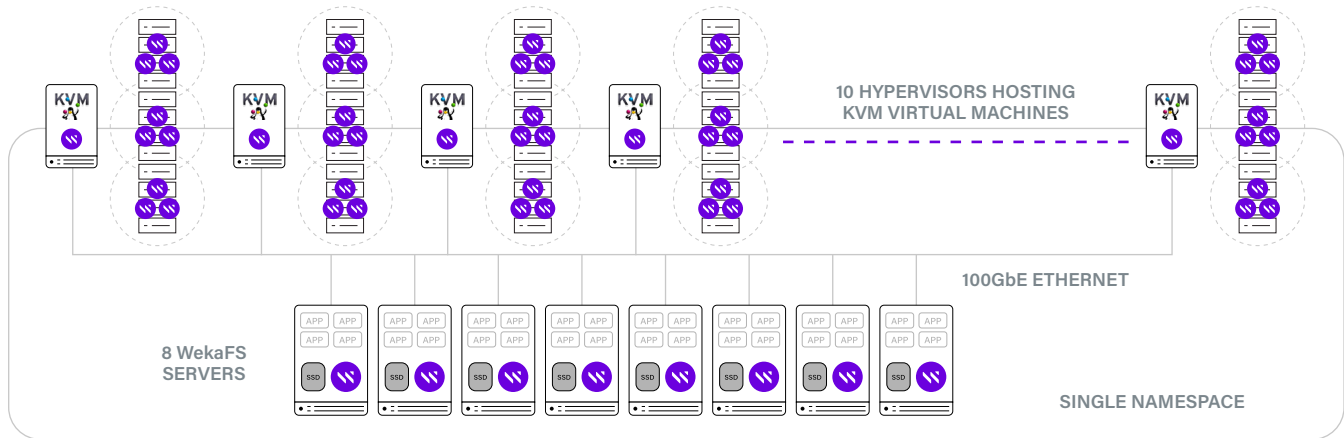
The challenge for SIB was to find a file system solution that provides the combination of data security with encryption, shared data accessibility for authorized users only, and high performance for applications to accelerate research. The IT team at SIB had prior experience with parallel file systems and hierarchical storage managers. Having evaluated other storage alternatives including BeeGFS and ZFS, the team found that none met the data shareability and encryption in-flight requirements without a negative impact on performance. The storage solution that SIB needed had to meet several criteria for technical and biomedical research benefits:

- provide encryption at rest and in-flight for data security
- maintain data interoperability and shared accessibility for authorized users
- deliver high performance and high bandwidth to client nodes
- handle mixed workloads
- provide a stable storage environment with data protection
- offer simplicity and ease of management.

THE SOLUTION: RUN WekaFS ON-PREMISES IN A VIRTUAL ENVIRONMENT

After comparisons with on-premises legacy NFS-based Network Attached Storage (NAS) solutions, SIB selected WekaFS to accelerate research while providing a secure environment for sensitive personal data. WekaFS was installed on an 8-node, on-premises, GPU-enabled, multi-tenant OpenStack cluster. OpenStack pools the virtual resources into a private cloud, and Weka is used to mount it for different tenants, each for a different scientific project. Data is tiered from the NVMe layer to 100 TB of Red Hat Ceph object storage.

Previously, SIB focused on high-throughput life science applications that did not deal with sensitive personal data and therefore could be supported with alternative legacy storage systems. Access restrictions at the file system level were sufficient, and using a shared infrastructure was common practice in the research community. However, new scientific applications related to personal health research require full access control and isolated project spaces for different research groups. Therefore, virtualized environments were introduced to split the physical cluster into smaller components servicing various groups of researchers instead of having one large cluster. Each research group then obtains an isolated WekaFS file system to securely store and access sensitive data.



WekaFS implemented in a virtualized environment

BENEFITS AND RETURN ON INVESTMENT: INCREASE IN THE NUMBER OF RESEARCH PROJECTS SUPPORTED

By choosing WekaFS, SIB was able to realize several benefits and greater return on investment:

- **Workload Increase:** since installing WekaFS, SIB supports a growing number of biomedical research projects thanks to the generated trust
- **Data Security:** Weka's advanced encryption protects confidential data both in flight and at rest, ensuring that only authorized users have access
- **Performance:** even with encryption enabled, Weka still delivers good performance and bandwidth to single client nodes, all while handling mixed workloads
- **Multi-tenancy and restricted access:** Weka's private cloud multi-tenancy capability enables shared data access for authorized user

THE WEKA INNOVATION NETWORK™ (WIN) DELIVERS SOLUTIONS THAT SOLVE BIG PROBLEMS

DALCO, a WIN Accelerator partner headquartered in Volketswil, Switzerland, brought expertise on how to design and build a high-performance cluster to the SIB implementation. DALCO helps customers such as SIB deliver results by architecting systems that solve their difficult compute problems and the challenge of managing ever-growing amounts of data. DALCO delivers unsurpassed customer value by architecting solutions that provide breakthrough density, high performance, efficiency, reliability, and manageability. For more information on DALCO, go to: <http://www.dalco.ch/>

For more information or to locate a partner in the Weka Innovation Network, go to: <https://www.weka.io/partners>.

For more information on the SIB Swiss Institute of Bioinformatics, go to: <https://www.sib.swiss/>



910 E Hamilton Avenue, Suite 430, Campbell, CA 95008 T: 408.335.0085 E: info@weka.io www.weka.io

©2018-2020 All rights reserved. WekaIO, WekaFS, Weka AI, WIN, Weka Innovation Network, Weka Within, Weka AI logo, Weka Within logo, Weka brand mark, Weka logo, and Radically Simple Storage are trademarks of WekaIO, Inc. and its affiliates in the United States and/or other countries. Other trademarks are the property of their respective companies. References in this publication to WekaIO's products, programs, or services do not imply that WekaIO intends to make these available in all countries in which it operates. Product specifications provided are sample specifications and do not constitute a warranty. Information is true as of the date of publication and is subject to change. Actual specifications for unique part numbers may vary.

W07CS202008