

SAP Prototypes String Compression Algorithm

Open FPGA Stack and Intel FPGA Acceleration Boards and Platforms provide the infrastructure to accelerate development.

At a glance:

- SAP SE is a German multinational software corporation based in Walldorf, Baden-Württemberg that develops enterprise software to manage business operations and customer relations.
- By leveraging open source [Open FPGA Stack \(OFS\)](#) and the Intel® FPGA PAC D5005 to prototype the Re-Pair compression workload, SAP was able to use their preferred configuration in their cloud infrastructure using Garden Linux running in Docker containers. [Download the one-page summary](#)

Executive Summary

Developers at SAP wished to create a proof-of-concept (PoC) of cloud-based Compression as a Service (CaaS). They needed to use FPGAs to accelerate the computationally intensive Re-Pair compression algorithm and hoped to employ Docker containers in SAP's HANA cloud employing SAP's own Garden Linux operating system (OS).

The Open FPGA Stack (OFS) eases the development and deployment of custom boards and workloads using Intel or third-party platforms powered by Intel® FPGAs. SAP developers used OFS to expedite deployment of their string compression workload onto an Intel FPGA PAC D5005. Furthermore, they were able to leverage Docker containers by following the OFS deployment flow. This was facilitated by the fact that OFS Device Feature List (DFL) FPGA drivers have been included in all versions of the Garden Linux kernel from release 5.15 onwards.

By using OFS, SAP is able to leverage workload portability across Intel FPGA-based devices, a growing ecosystem of OFS-enabled partner boards and workloads, flexibility in bare-metal/virtualized/containerized deployments, and upstreamed and open-sourced kernel drivers and user space code. OFS source code and technical documentation is open source in the [OFS Repository](#) on GitHub.

Background & Challenge

SAP SE is a German multinational software corporation based in Walldorf, Baden-Württemberg that develops enterprise software to manage business operations and customer relations.

SAP HANA is a relational database management system developed and marketed by SAP SE. It is an enterprise-grade database server that can store and retrieve data as requested by higher-level applications leveraging its in-memory and columnar storage for hybrid transaction/analytical processing.

Columnar data in SAP HANA is encoded with dictionaries, mapping any domain value to a fixed size value. String dictionaries in particular may contain vast amounts of textual data that needs to be compressed to minimize memory requirements. There are many different compression algorithms available (LZ77, LZR, LZSS, LZMA, ZStandard, etc.), but these algorithms are typically employed to compress large quantities of information into a single block. If a dictionary were compressed in this way, the entire file would have to be decompressed to access a single entry, which would be extremely inefficient in terms of time, computation, and power consumed. Alternatively, using these algorithms to compress each dictionary entry individually would present its own inefficiencies because they are not optimized to compress small amounts of data.

"We are now able to deploy our compression algorithm in Docker Containers running on our Garden Linux distribution in a matter of minutes using the OFS framework and the Intel FPGA PAC D5005. Using Intel's platform acceleration technology, SAP can now provide developers with the benefits of FPGA re-programmability in our HANA Cloud."

— SAP SE - Dr. Norman May HANA central (database) architect

Re-Pair is a compression algorithm that is well suited for applications such as string dictionaries that require random accesses to compressed data. Unfortunately, Re-Pair is a computationally intensive and expensive algorithm that has not enjoyed widespread use in the data management community due to its prohibitively high compression and decompression times when implemented on central processing units (CPUs). However, the programmable fabric in field-programmable gate arrays (FPGAs) can be configured to perform algorithmic processing in a massively parallel fashion. This means that algorithms like Re-Pair can be executed quickly while consuming relatively little power.

Garden Linux is a Debian GNU/Linux derivative crafted to provide small, auditable Linux images for use by cloud service providers (CSPs) and bare metal deployments. SAP has its own Garden Linux distribution.

In the context of computing, a container is a fully functional and portable cloud or non-cloud computing environment that includes the application along with any libraries and other dependencies. Using containers facilitates moving applications from one server to another because everything that is required to run that application is already inside the container. Docker is a common type of container used by many CSPs.

The challenge was to prototype the Re-Pair compression workload deployment using Docker containers in SAP's HANA Cloud employing its Garden Linux operating system (OS) in conjunction with a high-performance PCI Express (PCIe)-based FPGA acceleration card.

Solution

The Intel Programmable Solutions Group offers a wide variety of industry-leading FPGAs and SoC FPGAs. The Intel Programmable Solutions Group also offers a wide range of high-performance PCIe-based FPGA acceleration cards, such as the Intel® Stratix® 10 FPGA-based Intel FPGA PAC D5005.

Complementing these FPGA acceleration cards is OFS, a scalable, open source hardware and software infrastructure that addresses the challenges associated with designing FPGA-based acceleration platform solutions deployed in [Intel® Xeon®-processor-based servers](#).

OFS enables software, hardware, and application developers to use standard interfaces and application programming interfaces (APIs) to accelerate workload development and enable code reuse. OFS also allows applications to be deployed bare metal, virtualized, or containerized.

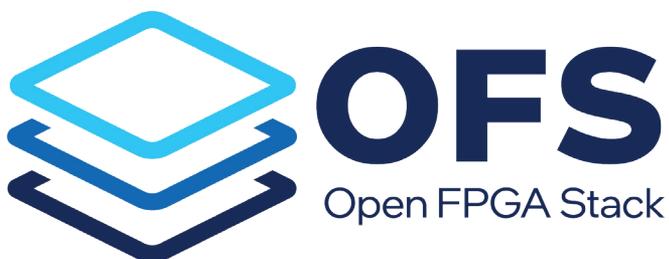
OFS provides the hardware and software infrastructures required to let users focus on their own unique applications. In this case study, the SAP developers leveraged the provided infrastructure to quickly port their Re-Pair compression workload to their Intel FPGA-based accelerator of choice. By following the OFS deployment flow, they were also able to leverage the high-level design (HLD) shim, which is a collection of hardware and software components to enable HLD-based workload support.

OpenCL™ and oneAPI are HLD frameworks used for heterogeneous computing across different compute accelerator architectures, including CPUs, graphics processing units (GPUs), digital signal processors (DSPs), FPGAs, and artificial intelligence (AI) accelerators. The initial proof-of-concept (PoC) of SAP's cloud-based Compression as a Service (CaaS) currently runs using OpenCL, with plans to adopt oneAPI in future iterations.

OFS also provides flexibility for different OS distributions, which—by extension—facilitates support in management and orchestration frameworks. In this case, SAP developers were able to deploy their workload using Docker containers. Additionally, OFS allows you to take full advantage of FPGA re-programmability by offering two configuration options: flat designs or designs that support partial reconfiguration (PR). PR provides a high level of flexibility by allowing portions of the FPGA to be reconfigured while the device is running—keeping the overall infrastructure intact and operating—thereby allowing changes to be made with no interruption to the system as a whole. OFS enabled SAP developers to leverage PR for their PoC inside their Docker containers.

By leveraging OFS and the Intel FPGA PAC D5005 to prototype the Re-Pair compression workload, SAP was able to use their preferred configuration in their cloud infrastructure using Garden Linux running in Docker containers. All of this was facilitated by the fact that OFS Device Feature List (DFL) FPGA drivers have been included in all versions of the Garden Linux kernel since release 5.15 onwards.

This case study provides a great example of how the infrastructure and flexibility provided by OFS enabled SAP to use their desired setup and port their workload in a short amount of time.



Results

This case study demonstrated how SAP benefits from FPGA-based re-programmability inside Docker containers and easy deployment in its own cloud.

This deployment was expedited by using the provided OFS reference infrastructure, source code, documentation, and the Intel FPGA PAC D5005 hardware reference platform. SAP was able to deploy their PoC in their cloud infrastructure, enabling them to plan the deployment of FPGAs in production. OFS also provides the flexibility for SAP to migrate to other Intel or third-party Intel® Stratix® 10 FPGA and [Intel Agilex® FPGA-based platforms](#) using OpenCL or oneAPI.

(It is important to note that the solution discussed in this case study is intended only as a PoC of cloud-based CaaS, and that this solution is not available for use by SAP customers at the time of this writing.)

Resources

- [Read the white paper to learn more about how SAP prototyped a containerized compression workload using OFS](#)
- [Download the PDF](#)



OpenCL and the OpenCL logo are trademarks of Apple Inc. used by permission by Khronos.

Intel technologies may require enabled hardware, software, or service activation.

No product or component can be absolutely secure.

© Intel Corporation. Intel, the Intel logo, and other Intel marks are trademarks of Intel Corporation or its subsidiaries. Other names and brands may be claimed as the property of others.